# Windows Server 2003/2003 R2

**Copyright© 2016 Microsoft Corporation**

The content in this document is retired and is no longer updated or supported. Some links might not work. Retired content represents the latest updated version of this content.

# Windows Server 2003 R2 and Windows Server 2003

Updated: February 26, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Welcome to the Windows Server 2003 R2 and Windows Server 2003 Technical Library. To find content, browse using the sections provided below.

| **Important** |
| --- |
| Mainstream support for Windows Server 2003 ended on July 13, 2010, and extended support ends on July 14, 2015. For more information, see the Windows Server 2003 end of support website.<br><br>For the most recent version of Windows Server, see the Windows Server TechNet Library. |

**Were you looking for...?**

| **Windows Server 2003 archived topics** | **Related products and technologies** | **Related topics** |
| --- | --- | --- |

| | | |
|---|---|---|
| • Comparison of Windows Server 2003 Editions<br>• Windows Server 2003: Operations<br>• Windows Server 2003: Security and Protection<br>• Windows Server 2003: Technical Reference<br>• Windows Server 2003 Glossary<br>• Windows Server 2003: System Requirements<br>• Windows Server 2003: Getting Started<br>• Windows Server 2003: Planning and Architecture<br>• Windows Server 2003: Deployment<br>• Windows Server 2003: Product Help<br>• Windows Server 2003 Service Pack 1 Technical Library<br>• Windows Server 2003 Service Pack 2 Technical Library<br>• Install Windows Server 2003 with Service Pack 2<br>• R2: Install Windows Server 2003 R2<br>• R2: Product Evaluation<br>• R2: Getting Started (R2 only)<br>• R2: Planning and Architecture (R2 only)<br>• R2: Deployment (R2 only)<br>• R2: Product Help (R2 only) | • Command-Line Reference<br>• Group Policy in Windows Server<br>• Group Policy Object Editor<br>• IIS 7.0 Health Model<br>• IIS 7.5 Health Model<br>• Metabase Property Reference<br>• Remote Desktop Services (Terminal Services) TechCenter<br>• Windows Home Server (mainstream support ends April 2016)<br>• Windows Server Update Services (WSUS)<br>• Windows Vista (extended support ends April 2017)<br>• Windows XP (extended support ends April 2014) | • Change a dynamic disk back to a basic disk<br>• Configure bitmap caching<br>• Configuring the Management Service in IIS 7<br>• Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042)<br>• How DNS Support for Active Directory Works<br>• Office 2007 Office Customization Tool<br>• Office 2007 Resource Kit<br>• Office 2010 Resource Kit<br>• Troubleshoot Specific Remote Desktop Problems |

**See also**

- Windows Server 2012 R2 and Windows Server 2012
- Windows Server 2008 R2 and Windows Server 2008

# Changes to Functionality in Windows Server 2003 Service Pack 1

Applies To: Windows Server 2003 with SP1

This document specifically focuses on the changes between earlier versions of Windows Server 2003 and Windows Server 2003 with Service Pack 1 (SP1) and reflects the current thinking of Microsoft about Service Pack 1 and its implications for IT professionals and developers. Examples and details are provided for several of the technologies that are experiencing the biggest changes.

The goal for Service Pack 1 is to build on the Trustworthy Computing efforts of Microsoft that were applied in Windows Server 2003. For an overview of the Microsoft Trustworthy Computing initiative, see "Trustworthy Computing Defined," on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=20970.

Windows Server 2003 SP1 incorporates many of the changes in functionality introduced with Windows XP Service Pack 2 (XP SP2), so there are many changes that were documented for Windows XP SP2 that are reiterated in this document.

Information that is not directly related to a technology area is contained in the Appendices section.

A downloadable copy of this document is available on the Microsoft Download Center at http://go.microsoft.com/fwlink/?LinkId=49637.

**Document History**
This is the third release of this document for Windows Server 2003 Service Pack 1 (SP1).

- Added new sections: Distributed File System (DFS) in Windows Server 2003 Service Pack 1, Distributed Transaction Coordinator, Internet Information Services (IIS) in Windows Server 2003 Service Pack 1, .NET Framework for x64-based operating systems, Terminal Services in Windows Server 2003 Service Pack 1

- Revised sections: Access-based Enumeration, Data Execution Prevention, DCOM Security Enhancements, Post-Setup Security Updates, Remote Installation Services (RIS) for 64-bit operating systems, Windows Firewall

**Other Resources and Feedback**
If you have any other questions that are not answered by this paper, "Windows Server 2003 Resources for IT Professionals" on TechNet has links to many other resources regarding Windows Server 2003. This page is updated periodically with the most recent information that is available.

In addition, we appreciate feedback on our documentation and our product. The following resources are available to you for providing feedback:

**Windows Server 2003 Newsgroups**
Newsgroups are a great place to ask questions of other users and find general information about other users' experiences with Windows Server 2003. You can examine some of the newsgroups using a Web browser on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=32745.

**Product Support Services**
If you are having a problem with your computer after installing Windows Server 2003 Service Pack 1, check the product support Web site on Microsoft.com first to see if your issue has been identified in Frequently Asked Questions or by a Knowledge Base article. If not, you can contact Product Support Services to get help with your issue. To start, see "Windows Server 2003 Support Center" on the Microsoft Web site at http://go.micorosoft.com/fwlink/?LinkId=42095.

**Microsoft Wish Program**
If you have a suggestion about how to improve a feature in Windows that you would like considered for the next service pack or major Windows version, you can contact the Microsoft Wish Program and tell them your idea. To find out how to send a comment or suggestion, see "How to Contact the Microsoft Wish Program" on the Microsoft Knowledge Base Web site at http://go.microsoft.com/fwlink/?LinkId=32748.

**Documentation feedback**
If you have any comments or suggestions about Windows Server 2003 Service Pack 1 documents, on the Web version of this document on the Microsoft Web site at the bottom of the page, click **Comments** and tell us what you think. Note that this is only for comments on the documentation, not the product itself.

# About this document

Applies To: Windows Server 2003 with SP1

In Windows Server 2003 Service Pack 1, Microsoft is delivering improvements in several technologies to help enhance security and help customers protect themselves against malicious software and other risks to their computer. These technologies are not intended to replace periodic security updates as they are released, but rather to help strengthen the overall defenses against malicious attacks.

The changes to functionality in Windows Server 2003 Service Pack 1 are in the following technology areas:

- Access-based Enumeration
- Active Directory in Windows Server 2003 Service Pack 1
    - Active Directory Preparation Tool (adprep.exe)
    - Domain Controller Diagnostics Tool (dcdiag.exe))
    - Active Directory Directory Services Maintenance Utility (ntdsutil.exe)
    - Security Account Manager (SAM)
- Add or Remove Programs Filter
- DCOM Security Enhancements
- Data Execution Prevention
- Device\PhysicalMemory Object
- Distributed File System (DFS) in Windows Server 2003 Service Pack 1
- Distributed Transaction Coordinator
- Download, Attachment, and Authenticode Enhancements
- Internet Explorer Feature Control Settings in Group Policy
- Internet Explorer URL Action and Advanced Security Settings in Group Policy
- Internet Explorer Using Feature Control Registry Settings with Security Zone Settings
- Internet Explorer Add-on Management and Crash Detection
- Internet Explorer Binary Behaviors Security Setting
- Internet Explorer BindToObject Mitigation
- Internet Explorer Information Bar
- Internet Explorer Local Machine Zone Lockdown
- Internet Explorer MIME Handling Enforcement
- Internet Explorer Network Protocol Lockdown

- Internet Explorer Object Caching
- Internet Explorer Pop-up Blocker
- Internet Explorer Untrusted Publishers Mitigations
- Internet Explorer Window Restrictions
- Internet Explorer Zone Elevation Blocks
- Internet Information Services (IIS) in Windows Server 2003 Service Pack 1
- Kernel patch protection for x64-based operating systems
- Outlook Express
- Post-Setup Security Updates
- Remote Access Quarantine
- Remote Installation Services (RIS) for 64-bit Operating Systems
- Resultant Set of Policy (RSoP)
- RPC Interface Restriction
- Security Configuration Wizard in Windows Server 2003 Service Pack 1
- Setup
- Terminal Services in Windows Server 2003 Service Pack 1
- Transmission Control Protocol/Internet Protocol (TCP/IP)
- WebDAV Redirector
- Windows Firewall in Windows Server 2003 Service Pack 1
- Windows Media Player
- Wireless Provisioning Services (WPS)
- .NET Framework version 1.1 for 64-bit operating systems

# Access-based Enumeration

Applies To: Windows Server 2003 with SP1

Access-based Enumeration is a new feature included with Windows Server 2003 Service Pack 1. This feature allows users of Windows Server 2003–based file servers to list only the files and folders to which they have access when browsing content on the file server. This eliminates user confusion that can be caused when users connect to a file server and encounter a large number of files and folders that they cannot access.

**What does Access-based Enumeration do?**
Access-based Enumeration filters the list of available files and folders on a server to include only those that the requesting user has access to.

**Who does this feature apply to?**
This feature applies to:

- Domain-joined computers.
- IT professionals who want to control the user's experience.

**Detailed description**
Access-based Enumeration allows users to see only files and folders that they have access to on a file server. This feature is not enabled by default.

To enable this feature, a property must be set on a file share to allow access-based enumeration. To enable this feature on your server, you can download a shell extension that provides both a graphical user interface for enabling access-based enumeration and a command-line interface for managing this feature. When this download is installed, a wizard will run that can automatically enable Access-based Enumeration on the shared folders on your computer. This download includes a whitepaper that provides further details about the shell extension, command line interface, and the `NetShareSetInfo` application programming interface (API). This download is available on the Microsoft Download Center at http://go.microsoft.com/fwlink/?LinkId=46228.

If you want to develop a tool yourself, you can use the `NetShareSetInfo` API. This property is an attribute of the `NetShareSetInfo` (API). For more information about the `NetShareSetInfo` API, see the Platform SDK and the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=46511. To enable Access-based Enumeration, you need to set a flag that points to the `SHARE_INFO_1005` structure. For more information about the `SHARE_INFO_1005` structure, see the Platform SDK and the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=45504.

The new flag to enable Access-based Enumeration is

```
#define SHI1005_FLAGS_ENFORCE_NAMESPACE_ACCESS  0x0800
```

This flag is only applicable to Windows Server 2003 Service Pack 1 and will have no effect on other versions of the Windows operating system.

After the feature is enabled, a listing of the content in that share will present the content that the user has access to.

**Why is this change important?**
This change is important because this allows users to see only those files and directories that they have access to and nothing else. This mitigates the scenario where unauthorized users might otherwise be able to see the contents of a directory even though they don't have access to it.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
The `SHI1005_FLAGS_ENFORCE_NAMESPACE_ACCESS` flag has been added to the NetShareSetInfo API. The flag enables you to turn on the Access-based Enumeration feature.

# Active Directory in Windows Server 2003 Service Pack 1

Applies To: Windows Server 2003 with SP1

**What does Active Directory do?**
Active Directory® is a directory service that stores information about objects on a network and makes this information available to users and network administrators. Active Directory objects typically include shared resources such as servers, volumes, printers, and the network user and computer accounts.

Active Directory is composed of the following:

- Schema. This is a set of rules that defines the classes of objects and attributes contained in the directory, the constraints and limits on instances of these objects, and the format of their names.
- Global catalog. This data store contains information about every object in the directory. This allows users and administrators to find directory information regardless of which domain in the directory actually contains the data.
- Query and index. Using this mechanism objects and their properties can be published and found by network users or applications.
- Replication service. This service distributes directory data across a network. All domain controllers in a domain participate in replication and contain a complete copy of all directory information for their domain. Any change to directory data is replicated to all domain controllers in the domain.
- Active Directory client software. The Active Directory client enables many of the Active Directory features available on Windows 2000 Professional or Windows XP Professional clients for computers running Windows 95, Windows 98, and Windows NT 4.0.

**Who does this feature apply to?**
The changes in Active Directory for Windows Server 2003 Service Pack 1 (SP1) will be of interest to:

- IT professionals who support Active Directory, such as Active Directory administrators, Active Directory schema administrators, Domain Name System (DNS) administrators, and domain controller administrators.
- Help desk professionals.
- Application developers.
- System integrators.

**What functionality is changing in Windows Server 2003 Service Pack 1?**
**Directory service backup reminders**
A new event message, event ID 2089, provides the backup status of each directory partition that a domain controller stores, including application directory partitions and Active Directory Application Mode (ADAM) partitions. If halfway through the backup latency interval (tombstone lifetime) a partition has not been backed up, this event is logged in the Directory Service event log and continues daily until the partition is backed up.

**Added replication security and fewer replication errors**
Replication metadata for domain controllers from which Active Directory has been removed is no longer retained by default, although a waiting period can be configured. This change improves replication security and eliminates replication error messages that are caused by failed attempts to replicate with decommissioned domain controllers. For more information about preserving replication metadata, see "How the Active Directory Replication Model Works" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=46510.

**Install from Media improvement for installing DNS servers**
**Install from Media** improvements make it easier to create a new domain controller that is a DNS server by providing the new option to include application directory partitions in the backup media that is used to install the new domain controller. This option eliminates the requirement for replication of the DomainDNSZones and ForestDNSZones application directory

partitions before the DNS server is operational.

**Enhancements for replication and DNS testing**
The Dcdiag.exe command-line tool, which is available in Windows Support Tools, provides new reporting on the overall health of replication with respect to Active Directory security. This test provides a summary of results along with detailed information for each domain controller that is tested and a diagnosis of any security errors. Dcdiag.exe also has new DNS tests for connectivity, service availability, forwarders and root hints, delegation, dynamic update, locator record registrations, external name resolution, and enterprise infrastructure. These tests can be performed on one domain controller or on all domain controllers in a forest. For more information about the changes to Dcdiag.exe, see the Dcdiag.exe section of this article.

**Support for running domain controllers in virtual machines**
On a single physical server that is running Windows Server 2003 and Microsoft Virtual Server 2005, you can install multiple Windows Server 2003 or Windows 2000 Server domain controllers in separate virtual machines. This platform is well suited for test environments. By using virtual machines, you can effectively host multiple domains, multiple domain controllers for the same domain, or even multiple forests on one physical server that is running a single operating system. Windows Server 2003 SP1 also provides protection against directory corruption that can result from improper backup and restoration of domain controller images. For more information about running domain controllers in virtual machines, see "Running Domain Controllers in Virtual Server 2005" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=38330.

**Operations master health and status reporting**
If an operation that requires a domain controller that holds an operations master role (also known as flexible single-master operations (FSMO)) cannot be performed, events are now logged in the Directory Service event log. Events identify role holders that do not exist, exist but are not available, or are available but have not replicated recently with the contacting domain controller. For more information about operations masters, see "How Operations Masters Work" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=38333.

**Extended storage of deleted objects**
The default period that a copy of a deleted object is retained in Active Directory, called the tombstone lifetime, is extended from 60 days to 180 days. Longer tombstone lifetime decreases the chance that a deleted object remains in the local directory of a disconnected domain controller beyond the time when the object is permanently deleted from online domain controllers. The tombstone lifetime is not changed automatically when you upgrade to Windows Server 2003 with SP1, but you can change the tombstone lifetime manually after the upgrade. New forests that are installed with Windows Server 2003 with SP1 have a default tombstone lifetime of 180 days. For more information about tombstone lifetime, see "How the Data Store Works" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=38339.

**Improved domain controller name resolution**
In response to Domain Name System (DNS) name resolution failures that may be encountered during location of replication partners and global catalog servers, domain controllers running Windows Server 2003 with SP1 request other variations of the server name that might be registered, which results in fewer failures due to DNS delays and misconfiguration. For more information about DNS name resolution, see "How DNS Support for Active Directory Works" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=38335.

**Simplified process for server metadata removal**
The Ntdsutil.exe command-line tool for managing the Active Directory database has new commands that make it easier to remove domain controller metadata. Preliminary steps, such as connecting to a server, domain, and site, are no longer required. You simply specify the server to remove. You can also specify the server on which to make the deletion. For more information about the changes to Ntdsutil.exe, see the Ntdsutil.exe section of this article.

**Improved security to protect confidential attributes**
To prevent Read access to confidential attributes, such as a Social Security number, while allowing Read access to other object attributes, you can designate specific attributes as confidential by setting a search flag on the respective attributeSchema object. By default, only domain administrators have Read access to confidential attributes, but this access can be delegated. For more information about access to attributes, see "How Security Descriptors and Access Control Lists Work" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=45972.

**Retention of SID history on tombstones**
The sIDHistory attribute has been added to the set of attributes that are retained on an object tombstone when the object is deleted. If a tombstoned object is reactivated (undeleted), the sIDHistory attribute is now restored with the object. For more information about tombstones, see "How the Data Store Works" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=45973.

**Adprep.exe improvements for Windows 2000 Server upgrades**
The Adprep.exe tool has been improved to reduce the impact of File Replication service (FRS) synchronization that results from updating SYSVOL files during upgrade. Adprep.exe is used to upgrade the Windows 2000 Server schema to the Windows Server 2003 schema and to update some forest- and domain-specific configuration, including SYSVOL, that is required for a

Windows Server 2003 domain controller to be operational. The tool now allows performing SYSVOL operations in a separate step when preparing the domain for upgrade. A new switch, `/gpprep`, has been added to accommodate the SYSVOL updates, which can be performed at a convenient time following the upgrade. The `adprep /domainprep` command, which formerly performed both directory and SYSVOL updates, now updates only the directory. Adprep.exe also now detects third-party schema extensions that block an upgrade, identifies the blocking extensions, and recommends fixes. Microsoft Exchange Server schema objects are also detected so that the Exchange Server schema can be prepared appropriately to accommodate InetOrgPerson naming. For more information about the changes to Adprep.exe, see the Adprep.exe section of this article.

**Changes in dragging and dropping objects in Active Directory Users and computers**
In Windows Server 2003, Service Pack 1 two changes to the drag and drop behavior in the Active Directory Users and Computers Microsoft Management Console (MMC) snap-in were made in response to customer feedback.

First, by default there is now a confirmation dialog when dragging and dropping objects in Active Directory Users and Computers Microsoft Management Console (MMC) snap-in. In Windows Server 2003 drag and drop support in Active Directory Users and Computers was enabled. However, it did not provide any confirmation dialog when moving objects. This made it easier to move objects, but also made it easier to inadvertently move an object to the wrong location and cause client workstations to lose access to critical resources. By adding a confirmation dialog to the drag and drop behavior, the administrator has a chance to correct an unintentional error before it impacts the organization. When the confirmation dialog is displayed there is a check box for **Don’t show this warning while this snap-in is open**.

If the user selects the **Don’t show this warning while this snap-in is open** checkbox, then the confirmation dialog will no longer be shown throughout the current snap-in session. Subsequent drag and drop attempts in that snap-in session will occur without any confirmation.

If the user doesn't select the **Don’t show this warning while this snap-in is open** checkbox, then the warning message will be shown every time the user tries to drag or drop an object.

Second, an administrator can choose to disable dragging and dropping completely by setting the **flags** attribute on the **Display Specifiers** container. The display specifiers container is in the directory at: **CN=DisplaySpecifiers,CN=Configuration,DC=**<*insert domain name*>. This attribute can be set using ADSIedit.msc, which is available in Windows Support Tools.

The overall behavior is:

- If the **flags** attribute is set to any value, then drag and drop is disabled. This is not the default.
- If the **flags** attribute is not set (default case), then the user will be able to use drag and drop to move objects in the Active Directory Computers and Users MMC snap-in.

# Active Directory Preparation Tool (adprep.exe)

Applies To: Windows Server 2003 with SP1

**What does ADPrep.exe do?**
Adprep.exe is a command-line tool used to prepare a Microsoft Windows 2000 forest or a Windows 2000 domain for the installation of Windows Server 2003 domain controllers.

**Who does this feature apply to?**
The changes in ADPrep.exe for Windows Server 2003 Service Pack 1 will be of interest to:

- IT professionals who support Active Directory, such as Active Directory administrators, Active Directory Schema administrators, Domain Name System (DNS) administrators, and domain controller administrators.
- Help desk professionals.
- Application developers.
- System integrators.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Adprep.exe enhancement to detect conflicting Exchange Server schema objects**
**Detailed description**
When Microsoft Exchange Server is deployed in an organization, Exchange Server uses Active Directory as a data store and it extends the Windows 2000 Active Directory schema to enable it to store objects specific to Exchange Server. The `ldapDisplayName` of the attribute schema `ms-Exch-Assistant-Name`, `ms-Exch-LabeledURI`, and `ms-Exch-House-Identifier` defined by Exchange Server conflicts with the `iNetOrgPerson` schema that Active Directory uses in Windows Server 2003. When Windows Server 2003 Service Pack 1 is installed, Adprep.exe will be able to detect the presence of the schema conflict and block the upgrade of the schema until the issue has been resolved.

**Why is this change important?**
Upgrading the Active Directory schema from Windows 2000 to Windows 2003 when these schema objects are present causes the `ldapDisplayName` to become corrupted and results in issues with Active Directory replication. Fixing the Exchange Server schema objects before the upgrade occurs results in a much smoother upgrade experience.

**What works differently?**
The Windows 2000 Active Directory schema cannot be upgraded to the Windows Server 2003 schema until the required Exchange Server schema objects are fixed.

**How do I resolve these issues?**
If Adprep.exe detects the presence of the conflicting Exchange Server schema objects, you can use the following procedure to fix these objects and enable Adprep.exe to successfully upgrade your Active Directory schema.

### To fix conflicting Exchange Server schema objects

1. Log on to the computer that holds the Schema Operation Master role. By default, the first domain controller that you install in your forest is the Schema Operation Master. You must log on using an account that is a member of the Schema Admins security group.
2. Click **Start**, click **Run**, type **notepad.exe** in the **Open** box, and then click **OK**.
3. Create the InetOrgPersonPrevent.ldf script by copying the following text including the trailing hyphen after "schemaUpdateNow: 1" to Notepad:

```
dn: CN=ms-Exch-Assistant-Name,CN=Schema,CN=Configuration,DC=X

changetype: Modify

replace:LDAPDisplayName

LDAPDisplayName: msExchAssistantName

-
```

```
dn: CN=ms-Exch-LabeledURI,CN=Schema,CN=Configuration,DC=X

changetype: Modify

replace: LDAPDisplayName

LDAPDisplayName: msExchLabeledURI

-

dn: CN=ms-Exch-House-Identifier,CN=Schema,CN=Configuration,DC=X

changetype: Modify

replace: LDAPDisplayName

LDAPDisplayName: msExchHouseIdentifier

-

dn:

changetype: Modify

add: schemaUpdateNow

schemaUpdateNow: 1

-
```

4. On the **File** menu, click **Save**. In the **Save As** dialog box, follow these steps to save the InetOrgPersonPrevent.ldf script:
   - In **File name**, type the following:

     **\%userprofile%\InetOrgPersonPrevent.ldf**

   - In **Save as type**, click **All Files**.
   - In **Encoding**, click **Unicode**.
   - Click **Save**.
   - Close Notepad.

5. Run the InetOrgPersonPrevent.ldf script using the following steps:
   - Click **Start**, click **Run**, type **cmd** in the **Open** box, and then click **OK**.
   - At a command prompt, type the following, and then press ENTER:

     **cd %userprofile%**

   - Type the following command

     **c:\documents and settings\%username%ldifde -i -f inetorgpersonprevent.ldf -v -c DC=X** "*domain_name_path_for_forest_root_domain*"

6. Verify that the `ldapDisplayNames` for the `CN=ms-Exch-Assistant-Name`, `CN=ms-Exch-LabeledURI`, and `CN=ms-Exch-House-Identifier` attributes in the schema naming context now appear as `msExchAssistantName`, `msExchLabeledURI`, and `msExchHouseIdentifier`.

**Note**

In step 5 of the previous procedure note the following details:

**DC=X** is a case-sensitive constant.

The domain name path for the root domain must be enclosed in quotation marks. For example, the command syntax for an Active Directory forest whose forest root domain is Contoso.com would be:

**c:\documents and settings\administrator>ldifde -i -f inetorgpersonprevent.ldf -v -c DC=X "dc=contoso,dc=com"**

**Adprep.exe enhancement to detect other conflicting schema objects**
**Detailed description**
Many applications use Active Directory as a data store and extend the Windows 2000 Active Directory schema to enable it to store objects specific to the application. If an application defined a non-RFC compliant schema object, such that the ldapDisplayName, object identifier (OID) or other schema attributes conflict with the Windows 2003 Active Directory schema, when Adprep.exe is run it will detect the conflict and display a generic error.

**Why is this change important?**
Upgrading the Active Directory schema from the Windows 2000 Server schema to the Windows Server 2003 schema when these schema objects are present causes the object to become corrupted and results in issues with Active Directory replication. Fixing the schema objects before the upgrade occurs results in a much smoother upgrade experience.

**What works differently?**
The Windows 2000 Active Directory schema cannot be upgraded to the Windows Server 2003 schema until the conflicting schema objects are fixed. In Windows Server 2003 Service Pac 1, Adprep.exe will identify that a conflicting schema object exists, prevent the corruption of the schema object by blocking the upgrade, and if possible identify which objects are in conflict so that you can resolve the conflict.

**Adprep.exe enhancement to perform SYSVOL operations in a separate step**
**Detailed description**
In previous versions of Windows Server 2003 running **adprep /domainprep** resulted in the addition of an inheritable access control entry (ACE) to all Group Policy objects (GPOs) in the SYSVOL folder. This ACE gives Enterprise domain controllers read access to the GPOs to support Resultant Set of Policy (RSoP) functionality for site base policy. The addition of this ACE is detected by the file replication service (FRS) and initiates an FRS synchronization of all GPOs in the SYSVOL folder.

In Windows Server 2003 Service Pack 1, the addition of the ACE to the GPOs in the SYSVOL folder is not performed while running **adprep /domainprep**. Instead, a new switch (**/gpprep**) has been added to **adprep** to add the inheritable ACE to the GPO folders in the SYSVOL directory. This allows administrators to update the ACE of the GPO objects at their convenience.

**Why is this change important?**
If an organization has a large number of files contained in the GPOs or slow links to replication servers, the FRS synchronization triggered by the **/domainprep** operation could adversely affect the deployment schedule for Windows Server 2003. By putting this operation at the discretion of the administrator, the impact of this operation can be planned and scheduled as part of the deployment.

**What works differently?**
The deployment of a Windows Server 2003 domain controller can occur after running **adprep /forestprep** and **adprep /domainprep**. Resultant Set of Policy (RSoP) functionality will only be operational after running **adprep /domainprep /gpprep**.

# Domain Controller Diagnostics Tool (dcdiag.exe)

Updated: February 17, 2010

Applies To: Windows Server 2003 with SP1

**What does DCDiag.exe do?**
This command-line tool analyzes the state of one or all domain controllers in a forest and reports any problems to assist in troubleshooting. DCDiag.exe consists of a variety of tests that can be run individually or as part of a suite to verify domain controller health.

**Tool location**
The `DCDiag` command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Tool requirements**

- Except as noted below, all commands in `DCDiag` can be run on Windows XP Professional and Windows Server 2003 family (member servers and domain controllers).
- The new `DCDIAG /TEST:DNS` command can validate DNS health of Windows 2000 Server (SP3 or later) or Windows Server 2003 family domain controllers when run from the console of Windows XP or Windows Server 2003 member computers or Windows Server 2003 domain controllers.

**Who does this feature apply to?**
This feature is of interest to the following audiences:

- DNS administrators
- Domain controller administrators
- DCDiag.exe users

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
There are two significant improvements to DCDiag in Windows Server 2003 Service Pack 1:

- `DCDIAG /TEST:DNS` to validate DNS health.
- `DCDIAG /CheckSecurityError` to detect security configurations that can cause Active Directory replication to fail.

The details of these new enhancements are described below.

**New DNS diagnostic tests**
**Detailed description**
DCDiag.exe has been enhanced for Windows Server 2003 Service Pack 1 to include new DNS functionality for reporting on the overall DNS health of domain controllers. There are seven new DNS-related tests that can be run individually or simultaneously. These tests may be performed on one or all domain controllers in an Active Directory forest. When the tests have completed, DCDiag.exe presents a summary of the results, along with detailed information for each domain controller tested.

| **Note** |
| --- |
| The new DNS tests require Enterprise Admin credentials.<br><br>The new DNS tests can be run only against Windows 2000 Server (SP3 or later) or Windows Server 2003 family domain controllers. |

**Command line syntax**

Windows Server 2003 SP1 `dcdiag` uses the same basic syntax as previous versions of `dcdiag`. The syntax for running the new DNS tests is as follows:

`Dcdiag /test:DNS [/DnsBasic | /DnsForwarders | /DnsDelegation | /DnsDynamicUpdate | /DnsRecordRegistration | /DnsResolveExtName [/DnsInternetName:` *InternetName* `] | /DnsAll] [/f:` *Logfile* `] [/ferr:` *Logerr* `] /S:` *DCName* `[/e] [/v]`

| Parameter | Description |
|---|---|
| `/test:DNS` | Performs all seven subtests except the `/DnsInternetName` test against the scoped set of domain controllers. The most common DCDIAG command line arguments are `DCDIAG /TEST:DNS /V /S:`*DCNAME* to run the six default DNS subtests against a single domain controller (DC) or `DCDIAG /TEST:DNS /V /E` to run the six default DNS subtests against all DCs in the console computer's test forest. `DCDIAG /TEST:DNS` is identical to the `/DnsAll` command when individual subtests are not defined. |
| `/test:DNS [DNS test]` | Performs the specified DNS test. If no test is specified, defaults to `/DnsAll`. |
| `/DnsBasic` | Performs basic DNS tests, including network connectivity, DNS client configuration, service availability, and zone existence. |
| `/DnsForwarders` | Performs the `/DnsBasic` tests, and also checks the configuration of forwarders. |
| `/DnsDelegation` | Performs the `/DnsBasic` tests, and also checks for proper delegations. |
| `/DnsDynamicUpdate` | Performs the `/DnsBasic` tests, and also determines whether dynamic update is enabled in the Active Directory zone. |
| `/DnsRecordRegistration` | Performs the `/DnsBasic` tests, and also checks whether the A, CNAME, and well-known SRV records are registered. In addition, creates an inventory report based on results. |
| `/DnsResolveExtName [/DnsInternetName:` *InternetName* `]` | Performs the `/DnsBasic` tests, and also attempts to resolve a sample intranet or Internet Name. If `/DnsInternetName` is not specified, then the command attempts to resolve the name www.microsoft.com. If `/DnsInternetName` is specified, then the command attempts to resolve the Internet name supplied by the user. |
| `/DnsAll` | Performs all tests, except for the `DnsResolveExtName` test, and generates a report. |
| `/f:` *Logfile* | Redirects output to the log file supplied by the user. |
| `/ferr:` *Logerr* | Redirects fatal error output to a separate log file. |
| `/s:` *DCName* | Specifies the domain controller against which to run the tests. |
| `/e` | All tests specified by `/test:DNS` are run against all domain controllers in the Active Directory forest. |
| `/v` | Verbose. Presents information about successful test results, in addition to information about errors and warnings. (When the `/v` parameter is not used, only error and warning information is presented.) Microsoft recommends using the `/v` switch when errors or warnings are reported in the summary table |

**Enterprise DNS Infrastructure Test (/e)**

- When `/test:DNS` is run in conjunction with the `/e` parameter, all tests specified by `test:/DNS` are run against all domain controllers in the Active Directory forest.

| **Note** |
|---|
| Run times for DNS tests can be significant in large enterprises when the `/e` parameter is used. Domain controllers and DNS servers that are offline will increase run time due to long time out periods for RPC and other protocols. |

**Connectivity test**

- The connectivity test is a mandatory test and runs automatically before any other `dcdiag` test is run.

- The connectivity test determines whether domain controllers are registered in DNS, can be pinged, and have LDAP/RPC connectivity.
- If the connectivity test fails on a given controller, no other tests are run against that domain controller.

| Note |
| --- |
| The connectivity test has not been changed in SP1, but is included in this document for reference. |

**Basic DNS Test (/DnsBasic)**

- The basic DNS test confirms that the following essential services are running and available on domain controllers tested by `dcdiag`:
    - DNS client service
    - Netlogon service
    - KDC service
    - DNS Server service (if DNS is installed on the domain controller)
- The basic DNS test confirms network connectivity for each domain controller by confirming that DNS servers on all adapters are reachable.
- The basic DNS test confirms that the A record of each domain controller is registered on at least one of the DNS servers configured on the client.
- If a domain controller is running the DNS Server service, the basic DNS test confirms that the Active Directory domain zone and SOA record for the Active Directory domain zone are present.
- The basic DNS test checks whether the root (.) zone is present.

**Forwarder test (/DnsForwarders)**

| Note |
| --- |
| This test runs only if the domain controller being tested is running the Microsoft DNS Server service. |

- The forwarder test determines whether recursion is enabled.
- If forwarders or root hints are configured, the forwarder test confirms that all forwarders or root hints on the DNS server are functioning, and also confirms that the _ldap._tcp.<Forest root domain> DC Locator record is resolved. (Resolution of the _ldap_tcp.<Forest root domain> DC Locator record is not attempted for forwarders or root hints configured on the forest root domain controller.)

**Delegation test (/DnsDelegation)**

| Note |
| --- |
| This test runs only if the domain controller being tested is running the Microsoft DNS Server service. |

- The delegation test confirms that the delegated name server is a functioning DNS Server.
- The delegation test checks for broken delegations by ensuring that all NS records in the Active Directory domain zone in

which the target domain controller resides have corresponding glue A records.

**Dynamic Update Test (/DnsDynamicUpdate)**

- The dynamic update test confirms that the Active Directory domain zone is configured for secure dynamic update and performs registration of a test record (_dcdiag_test_record). The test record is subsequently deleted.

**Record Registration Test (/DnsRecordRegistration)**

- The record registration test verifies the registration of all essential DC Locator records on all DNS Servers configured on each adapter of the domain controllers. This test returns the following records.

| Record | Description |
|---|---|
| CNAME GUID | The GUID registered as the canonical name (CNAME) of the DNS server. |
| A | The host address (A) resource record. Maps a DNS domain name to an Internet Protocol (IP) version 4 32-bit address. |
| LDAP SRV | The service locator (SRV) resource record for the LDAP service. |
| GC SRV | The service locator (SRV) resource record for the global catalog (GC) server. |
| PDC SRV | The service locator (SRV) resource record for the primary domain controller (PDC). |

**External Name Resolution Test (/DnsResolveExtName)**

| **Note** |
|---|
| The external name resolution test is run only if specified explicitly (using `/DnsResolveExtName`); it is not run as part of `/DnsAll`. |

- The external name resolution test verifies basic resolution of external DNS from a given client, using a sample Internet name (www.microsoft.com), or user-provided Internet name.

- The external name resolution test cannot resolve external Internet names in an environment where a proxy server is being used.

- You can test name resolution using either intranet or Internet names.

- To resolve a user-provided Internet or intranet name (rather than the default name of www.microsoft.com), the `/DnsInternetName` parameter must be used.

**How to read the output of DNS enhanced dcdiag**

The following steps summarize how to interpret the results provided by DNS-enhanced `dcdiag`:

1. Run `dcdiag test:DNS /e /f:dns.txt`. Microsoft recommends always using the `/v` switch to obtain verbose information.

2. Open the report in Notepad or a compatible editor.

3. Scroll to end of the report and read the summary table.

4. Identify servers that returned "warn" or "fail" status for any subtest in the summary table.

5. Review the section of output for that server to see what problem was detected (hint: use the **Find** command on the **Edit** menu to search on the string "`DC:` *DC_computername*" (without quotes) to locate the detailed section for a given DC.

6. Resolve problems on DNS clients or DNS server(s) as required.

7. Run `dcdiag /test:DNS /v /e` (or `/s:`*DCName*) again to verify the fix. Repeat steps 1 through 6 as required until all failures are understood and reconciled.

**Warnings and Errors**

`Dcdiag` takes a conservative approach by identifying DNS client or DNS server configurations that may be problematic, do not conform to best practice configurations, or that `dcdiag` cannot fully validate. Therefore, the summary and detailed sections of `dcdiag` may report warnings for DNS configurations that are currently functional. Administrators should investigate and validate such configurations when identified by `dcdiag`.

The tables below contain the configurations that can trigger `dcdiag` to report warnings or errors for each of the DNS subtests.

**Basic**

| Warning | Additional information |
|---|---|
| Warning: Adapter <adapter name> has dynamic IP address | Static IP addresses are recommended for all DNS servers. |
| Warning: Adapter <adapter name> has invalid DNS server: <name> <IP address> | DNS server may not be reachable. |
| Warning: No DNS RPC connectivity (error or non Microsoft DNS server is running) | Disregard this warning if the DNS server is a BIND or other non-Microsoft DNS server. |
| Warning: The Active Directory zone on this DC/DNS server was not found | N/A |
| Warning: Root zone on this DC/DNS server was found | N/A |

| Error | Additional information |
|---|---|
| Error: Authentication failed with specified credentials | DCDIAG requires Enterprise Admin credential to run all the tests. |
| Error: No LDAP connectivity | N/A |
| Error: No DS RPC connectivity | N/A |
| Error: No WMI connectivity | DNS test requires WMI connectivity to run on the remote computer. |
| Error: Can't read operating system version through WMI | This might be caused by the lack of a WMI connection on the remote computer. |
| Error: <Operating system name> not supported (this tool is supported on Windows 2000, Windows XP, and Windows Server 2003 only) | N/A |
| Error: Open Service Control Manager failed | Unable to find whether the service is running or not. |
| Error: Kdc/netlogon/DNS/dnscache is not running | Some of the key services are not running. |
| Error: Can't read network adapter information through WMI | N/A |
| Error: All DNS servers are invalid | DNS servers that the client is pointing to are either not reachable, not a DNS server, or have invalid IP addresses. |
| Error: The A record for this DC was not found | Every DC should register an A record. Make sure A records are registered on all the DNS servers the client is pointing to. |
| Error: Enumeration of zones failed to find root and AD zone | N/A |

| | |
|---|---|
| Error: Could not query DNS zones on this DC | Make sure that the zone in which the DC is supposed to register is present. |

**Forwarder**

| Error | Additional information |
|---|---|
| Error: Forwarders list has invalid forwarder: <IP address of the forwarder> | Forwarders configured on the DNS server have an invalid IP address or are not a DNS server, or name resolution is not working (that is, cannot resolve forest root domain SRV record if it is a non-root domain DC). |
| Error: Both root hints and forwarders are not configured. Please configure either forwarders or root hints | Make sure either forwarders or root hints are configured on the DNS server unless it hosts root zone. |
| Error: Root hints list has invalid root hint server: <IP address of Root hint server> | Root hint servers configured on the DNS server have invalid IP address or are not a DNS server, or name resolution not working (that is, cannot resolve forest root domain SRV record if it is a non root domain DC). |
| Error:<Root hint server Name> IP: <Unavailable> Status:<status of the server> | Configured root hint servers don't have corresponding IP address. Status field will tell you the status of the server |
| Error:<Root hint server Name> IP: <Unavailable> Status: A record not found | Configured root hint servers don't have A record. |
| Error: Enumeration of Root hint servers failed on <DNS server name> | Couldn't list the root hint servers on the target DNS server. |

**Delegation**

| Warning | Additional information |
|---|---|
| Warning: DNS server: <DnsServer name> IP: <Ipaddress> Failure: Missing glue A record | The configured delegation is missing glue A record. |

| Error | Additional information |
|---|---|
| DNS server: <Server name> IP:<IP address> Error: Broken delegation -verbose | Delegation is configured but the name server is not responding. |
| DNS server: <Server name> IP:<IP address> Error: Broken delegated domain <Delegated domain name> -non-verbose | N/A |
| Error: Failed to enumerate the records at the zone root on the server | N/A |

**DynamicUpdate**

| Warning | Additional information |
|---|---|
| Warning: Dynamic update is enabled on the zone but not secure <zone name> | Secure dynamic updates are recommended. |
| Warning: Failed to add test record _dcdiag_test_record with error <error code> in zone <zone name> | Test adds a dummy record dynamically |
| Warning: Failed to delete test record _dcdiag_test_record with error <error code> in zone zone <zone name> | Deletes the added record as well. |

| Error | Additional information |
|---|---|
| Error: Dynamic update is not enabled on the zone <zone name> | Dynamic update is not enabled on the Active Directory zone so client cannot register its records. |

**Record registration**

| Warning | Additional Information |
| --- | --- |
| Warning: Missing DC SRV record at DNS server <record name> | Ignore the error if `DNSAvoidRegisterRecord` registry key or its Group Policy has been configured to prevent registration of this record. |
| Warning: Missing GC SRV record at DNS server <record name> | Ignore the error if `DNSAvoidRegisterRecord` registry key or its Group Policy has been configured to prevent registration of this record. |
| Warning: Missing PDC SRV record at DNS server <record name> | Ignore the error if `DNSAvoidRegisterRecord` registry key or its Group Policy has been configured to prevent registration of this record. |
| Warning: Record Registrations not found in some network adapters | N/A |

| Error | Additional information |
| --- | --- |
| Error: Missing A record at DNS server <DNS Server IP address> : <A record name> | DC hasn't registered its A record on the specified DNS server. |
| Error: Missing CNAME record at DNS server <DNS Server IP address> : <CNAME record name> | DC hasn't registered its CNAME record on the specified DNS server. |
| Error: Missing DC SRV record at DNS server <DNS Server IP address> : <SRV record name> | DC hasn't registered its DC SRV record on the specified DNS server. |
| Error: Missing GC SRV record at DNS server <DNS Server IP address> : <SRV record name> | DC hasn't registered its GC SRV record on the specified DNS server. |
| Error: Missing PDC SRV record at DNS server <DNS Server IP address> : <SRV record name> | DC hasn't registered specified PDC SRV record on the specified DNS server. All these records can be registered by stopping and starting the netlogon service. |
| Error: Record registrations cannot be found for all the network adapters | If there are multiple network adaptors the test checks whether all the records are present on all the DNS servers configured on each adaptor. This error occurs if the record registration is missing on the DNS server. |

**External name resolution**

| Error | Additional information |
| --- | --- |
| Error: Internet name <name> cannot be resolved | Specified Internet name cannot be resolved. Make sure the proxy client, servers, root hints, and forwarders are configured properly. |

**Enterprise DNS infrastructure tests**

| Warning | Additional information |
| --- | --- |
| Warning: Neither forwarders nor root hints are configured from subordinate domain to parent domain | Forwarder or root hints need to be configured in the DNS servers of either the parent or subordinate domains that are hosting the authoritative zones for their respective domain to enable name resolution to work. |

| Error | Additional information |
| --- | --- |
| Error: Delegation is not configured on the parent domain | Delegation should be configured from parent to subordinate domain. |
| Error: Delegation is present but the glue record is missing | Delegation is configured but the name servers are missing their glue record. |
| Error: Forwarders are misconfigured from parent domain to subordinate domain | Forwarders must be configured from subordinate domain to parent domain. |

| | |
|---|---|
| Error: Root hints are misconfigured from parent domain to subordinate domain | Root hints must be configured from subordinate domain to parent domain. |
| Error: Forwarders are configured from subordinate to parent domain but some of them failed DNS server tests (See DNS servers section for error details) | Forwarders configured have an invalid IP address or are not a valid DNS server, or name resolution is not working (cannot resolve forest root domain SRV record if it is in the non-root domain). |
| Error: Root hints are configured from subordinate to parent domain but some of them failed DNS server tests (See DNS servers section for error details) | Root hints configured have an invalid IP address or are not a valid DNS server, or name resolution is not working. |

**Examples:**
The following examples illustrate the use of Windows Server 2003 SP1 `dcdiag`. You should replace the parameters in italics with those appropriate for your environment:

- To run all DNS tests on a single domain controller in non-verbose mode:

  `Dcdiag /test:DNS /s:`*TargetDCName* `/f:`*LogFileName*

- To run all DNS tests on a single domain controller in verbose mode:

  `Dcdiag /test:DNS /s:`*TargetDCName* `/v /f:`*LogFileName*

- To run all DNS tests on an entire forest in non-verbose mode:

  `Dcdiag /test:DNS /e /f:`*LogFileName*

- To run all DNS tests on an entire forest in verbose mode:

  `Dcdiag /test:DNS /v /e /f:`*LogFileName*

- To run the DNS basic test on a single domain controller:

  `Dcdiag /test:DNS /DnsBasic /s:`*TargetDCName* `/f:`*LogFileName*

- To run the DNS forwarders test on a single domain controller:

  `Dcdiag /test:DNS /DnsForwarders /s:`*TargetDCName* `/f:`*LogFileName*

- To run the DNS delegation test on a single domain controller:

  `Dcdiag /test:DNS /DnsDelegation /s:`*TargetDCName* `/f:`*LogFileName*

- To run the DNS dynamic update test on a single domain controller:

  `Dcdiag /test:DNS /DnsDynamicUpdate /s:`*TargetDCName* `/f:`*LogFileName*

- To run the DNS record registration test on a single domain controller:

  `Dcdiag /test:DNS /DnsRecordRegistration /s:`*TargetDCName* `/f:`*LogFileName*

- To resolve a sample Internet or intranet name:

  `Dcdiag /test:DNS /DnsResolveExtName /DnsInternetName:`*InternetName* `/f:` *LogFileName*

**Note**

When an individual test is run, the `/DnsBasic` tests are run by default before running the individual test specified.

If no individual test is specified, all DNS tests (except `/DnsResolveName`) are run by default.

**New Active Directory replication security tests**
**Detailed description**
DCDiag.exe has been enhanced for Windows Server 2003 Service Pack 1 to include new functionality to identify security configurations that can cause Active Directory replication to fail.

The new `CheckSecurityError` test may be performed on one or all domain controllers in an Active Directory forest. The test performs the following operations:

- Checks for the availability of a Key Distribution Center (KDC) in both the destination and source domain controller's domains.
- Verifies that the destination DC can transmit and receive sufficiently large UDP-formatted packets (used by Kerberos).
- Verifies that system clock of the destination DC is no more than 5 minutes different from the system time of the KDC in the destination and source domain, and the source DC.
- Confirms that the root of each naming context on the source domain controller is configured with the necessary permission.
- Confirms that the source and destination DC computer accounts are not disabled, are trusted for delegation, and contain all required service principal names.

When the test has completed, DCDiag.exe presents a summary of the results for each domain controller tested and the diagnosis of the security errors encountered

This test can be run from the command-line using the following syntax:

```
Dcdiag /test:CheckSecurityError
```

Optionally, you can add the switch `/ReplSource:`*SourceDC* to the command to identify a specific domain controller as a source in a replication attempt. The domain controller specified in the `/replsource:` parameter does not need to be a current source domain controller that the domain controller being tested currently replicates from (one that the destination domain controller currently has an inbound connection object from).This test will collect information from the domain controller, key distribution center (KDC) source and destination servers, and Active Directory.

**Note**

`Dcdiag /test:CheckSecurityError` can be executed on the console of a member computer (using the `/e` or `/s:`*servername* commands) as well as a domain controller. For best results, run `Dcdiag /test:CheckSecurityError` on the console of each domain controller that is failing inbound Active Directory replication due to a suspected security error.

**Why is this change important?**
If replication is not working and the error is a security error (such as "Access Denied", "The target account name is incorrect", or "The RPC server is unavailable") there are many different factors that could be causing the issue. This test automates the diagnosis by looking at the most common sources of these errors and reporting them so that you can resolve the issue.

# Active Directory Directory Services Maintenance Utility (ntdsutil.exe)

Applies To: Windows Server 2003 with SP1

**What does Ntdsutil.exe do?**
Ntdsutil.exe is a command-line tool that provides management facilities for Active Directory. You can use Ntdsutil.exe to perform database maintenance of Active Directory, manage and control single master operations, create application directory partitions, and remove metadata left behind by domain controllers that were not successfully demoted using the Active Directory Installation wizard (DCPromo.exe).

**Who does this feature apply to?**
This feature applies to the Ntdsutil.exe utility, and is of interest to Active Directory administrators only.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**General Improvements**
**Detailed description**
Ntdsutil.exe no longer requires the administrator to perform the following tasks in the `ntdsutil metadata cleanup` command:

- Connect to specific domain controller using `ntdsutil metadata cleanup connections`command.
- List and select the Active Directory domain, site, and server using the `ntdsutil metadata cleanup Select Operation Target` command.

Two new variations of this command are introduced in Windows Server 2003 Service Pack 1:

- `Ntdsutil "metadata cleanup" "remove selected server"`*ServerObject*

  When using this command, specify the distinguished name (DN) path of the server object (*ServerObject*) of the domain controller whose metadata you want to remove. The server object is the parent of the NTDS settings object in the configuration container. For example, for the domain controller named DC1 located in the default-first-site-name of the contoso.com forest, the DN path of the server object would be cn=DC1,cn=servers,cn=default-first-site-name ,cn=configuration,dc=contoso,dc=com. If the DN path contains any spaces, enclose the entire DN path in quotes.

- `Ntdsutil "metadata cleanup" "remove selected server"`*ServerObject* `on` *TargetDC*

  This command is identical to the one above, except it allows the administrator to specify the domain controller (*TargetDC*) on which the removal is performed. *TargetDC* must be entered as the DNS or NetBIOS name of the domain controller.

**Why is this change important? What threats does it mitigate?**
This change significantly improves the usability of this command for removing metadata.

**What works differently?**
From the "`metadata cleanup`" menu, the user no longer has to go into the "`connections`" menu or the "`select operations target`" menu to set up the appropriate state.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Improved Metadata Cleanup**
**Detailed description**
The metadata cleanup command has been improved in Windows Server 2003 Service Pack 1 to clean up metadata in Active Directory.

**What works differently? Are there any dependencies?**
The existing "`remove selected server`" command in the "`metadata cleanup`" menu of Ntdsutil.exe has been enhanced with new functionality.

Prior to Service Pack 1, this command only performed the following operations:

- Delete the NTDS settings object for the domain controller (DC).
- Delete all manual and automatic inbound connections to the DC being removed.
- Delete the corresponding DC's FRS member object from the sysvol replica set.

With the release of Service Pack 1, the following additional operations are performed as part of this command:

- Delete the computer account for the DC being deleted, including FRS subscriber objects.
- Delete all manual and automatic outbound Active Directory connections from the DC being removed.
- Delete inbound and outbound FRS connections from any non-sysvol FRS replica sets that the DC being deleted is a member of.
- Check whether the DC being removed holds any operations masters roles. If yes, this command will attempt to reassign (seize) the roles to an active DC that meets criteria for the operations master role(s).

# Security Account Manager (SAM)

Applies To: Windows Server 2003 with SP1

**What does the Security Account Manager (SAM) do?**
The Security Account Manager (SAM) is a database present on servers running Windows Server 2003 that stores user accounts and security descriptors for users on the local computer.

**Who does this feature apply to?**
This feature applies to IT professionals who want to troubleshoot or understand the SAM component's behavior in deployment. Troubleshooting may require that additional tools be downloaded. This topic also applies to developers that have licensed the SAMR and LSAR protocols and developers that use the LSA trust application programming interfaces (APIs) from MSDN.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**SAM WPP Logging**
**Detailed description**
Security Accounts Manager (SAM) debug logs can be collected by Windows software trace preprocessor (WPP) during deployment. WPP can be used to gather information about what the SAM component is doing during a period of time when the Windows system is not behaving as expected. This information can be used by Microsoft Product Support Services to help you troubleshoot issues in your deployment.

**Why is this change important? What threats does it mitigate?**
This may reduce the number of live debug sessions if the log information is sufficient to determine what is occurring.

**What works differently?**
Nothing works differently. A new feature is enabled to generate logging. Following are the logman commands that can be used to enable logging:

```
logman create trace samlog -p "{f2969c49-b484-4485-b3b0-b908da73cebb}" 3
logman start samlog
rem repeat action that is interesting and that should be captured in log
logman stop samlog
```

This will generate an extended transaction log (ETL), which your Product Support Engineer will be able to parse using the debug symbol set.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
A new Event Tracing for Windows (ETW) entry exists, **f2969c49-b484-4485-b3b0-b908da73cebb**. This entry reflects whether or not logging has been enabled for the SAM component. The following sample output from `tracelog -enumguid` includes the new entry.

```
    Guid                        Enabled  LoggerId Level Flags
-------------------------------------------------------------
1046d4b1-fce5-48bc-8def-fd33196af19a     FALSE  0     0     0
5007c7b1-1444-4303-bdbe-359c79fc032a     FALSE  0     0     0
7e4b70ee-8296-4f0f-a3ba-f58ef7bb4e96     FALSE  0     0     0
77db410c-561e-4358-8b0e-af866e91bb89     FALSE  0     0     0
dd5ef90a-6398-47a4-ad34-4dcecdef795f     FALSE  0     0     0
196e57d9-49c0-4b3b-ac3a-a8a93ada1938     FALSE  0     0     0
1540ff4c-3fd7-4bba-9938-1d1bf31573a7     FALSE  0     0     0
94a984ef-f525-4bf1-be3c-ef374056a592     FALSE  0     0     0
3121cf5d-c5e6-4f37-be86-57083590c333     FALSE  0     0     0
94335eb3-79ea-44d5-8ea9-306f49b3a04e     FALSE  0     0     0
4a8aaa94-cfc4-46a7-8e4e-17bc45608f0a     FALSE  0     0     0
f33959b4-dbec-11d2-895b-00c04f79ab69     FALSE  0     0     0
8e598056-8993-11d2-819e-0000f875a064     FALSE  0     0     0
f2969c49-b484-4485-b3b0-b908da73cebb     FALSE  0     0     0
cc85922f-db41-11d2-9244-006008269001     FALSE  0     0     0
c92cf544-91b3-4dc0-8e11-c580339a0bf8     FALSE  0     0     0
bba3add2-c229-4cdb-ae2b-57eb6966b0c4     FALSE  0     0     0
8fc7e81a-f733-42e0-9708-cfdae07ed969     FALSE  0     0     0
cddc01e2-fdce-479a-b8ee-3c87053fb55e     FALSE  0     0     0
6acd39eb-4cb0-486b-83fa-307aa23767b1     FALSE  0     0     0
65f67abd-ecd2-4501-9b10-d48db2300e6c     FALSE  0     0     0
28cf047a-2437-4b24-b653-b9446a419a69     FALSE  0     0     0
fc4b0d39-e8be-4a83-a32f-c0c7c4f61ee4     FALSE  0     0     0
```

```
fc570986-5967-4641-a6f9-05291bce66c5     FALSE  0    0    0
39a7b5e0-be85-47fc-b9f5-593a659abac1     FALSE  0    0    0
dab01d4d-2d48-477d-b1c3-daad0ce6f06b     FALSE  0    0    0
58db8e03-0537-45cb-b29b-597f6cbebbfe     FALSE  0    0    0
58db8e03-0537-45cb-b29b-597f6cbebbfd     FALSE  0    0    0
```

**Preventing SAM and LSA handles from hijacking**
**Detailed description**
The server side implementation of the SAMR and LSAR protocols now implements security checks to ensure that the current caller is the same caller that opened the first handle returned from `SamConnect` and `LsaOpenPolicy` respectively.

The Security Account Manager Remote Procedure Call (RPC) protocol (SAMR) is an integral subsystem that is used to perform remote Service Account Manager operations, such as user account management and manipulation. The SAMR interface defines the remote Security Account Manager (SAM) methods that are called by the client. `SamConnect` is the function that is used to connect to the SAM database.

**Why is this change important?**
This change is related to the RPC changes that help prevent elevation of privilege attacks on your system. Implementing this change in the Active Directory interfaces helps make your system more secure by default.

**What works differently?**
If your application uses the SAMR or LSAR protocols, access checks are performed on each call received and verify that the client identity that opens the context handle is the same as the client identity that is making the call. If your application does not use that convention, it will no longer function after installation of Windows Server 2003 Service Pack 1 (SP1).

**How do I fix these issues?**
All calls to the SAMR and LSAR methods must be in the same security context as the call that generated the context handle being used in the call. If they are not you must modify your application to comply with this requirement.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
Most applications will not need to be changed. However, if the application's code switches security contexts while using context handles obtained from the SAMR and LSAR interface it will need to be modified. If the application is using any of the following APIs, verify with the application developer that the calling security context doesn't change between the call `LsaOpenPolicy` and any subsequent Lsa API call that uses the handle returned from `LsaOpenPolicy`.

- `LsaOpenPolicy`
- `LsaQueryInformationPolicy`
- `LsaSetInformationPolicy`
- `LsaQueryDomainInformationPolicy`
- `LsaSetDomainInformationPolicy`
- `LsaEnumerateTrustedDomains`
- `LsaLookupNames`
- `LsaLookupNames2`
- `LsaLookupSids`
- `LsaEnumerateAccountsWithUserRight`
- `LsaEnumerateAccountRights`
- `LsaAddAccountRights`
- `LsaRemoveAccountRights`

- LsaOpenTrustedDomainByName
- LsaQueryTrustedDomainInfo
- LsaSetTrustedDomainInformation
- LsaDeleteTrustedDomain
- LsaQueryTrustedDomainInfoByName
- LsaSetTrustedDomainInfoByName
- LsaEnumerateTrustedDomainsEx
- LsaCreateTrustedDomainEx
- LsaQueryForestTrustInformation
- LsaSetForestTrustInformation
- LsaForestTrustFindMatch
- LsaStorePrivateData
- LsaRetrievePrivateData

# Add or Remove Programs Filter

Applies To: Windows Server 2003 with SP1

**What does the filter for Add or Remove Programs do?**
The filter for **Add or Remove Programs** provides a means for the user to select whether or not updates, such as security updates downloaded from the Microsoft Web site, are displayed under the **Currently installed programs** list.

**Who does this feature apply to?**
**Add or Remove Programs** can be used by any user with Administrator credentials on the local computer. Although some applications can be installed or removed by non-administrators, most do require administrative credentials.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Filtering out updates from the Change or Remove Programs list**
**Detailed description**
The **Change or Remove Programs** list in **Add or Remove Programs** displays installed programs that the user can change or remove. This list also shows updates to Windows or other programs that have been installed.

In Windows Server 2003 Service Pack 1, the user is able to choose whether to show or hide updates for Windows and other programs in this view. A new **Show updates** check box appears above the list, which enables the user to toggle between showing or hiding installed updates.

**Why is this change important?**
Software vendors are creating more software updates and releasing them more frequently than ever before. These frequent updates help to increase the reliability and security of user's systems. However, by showing every update in the Change or Remove Programs list in Add or Remove Programs, the list of installed programs is overwhelmed by the list of installed updates. A new option to filter out the updates from the list and only show installed programs makes this list easier for users to read.

**What works differently?**
By default, **Change or Remove Programs** will not show installed updates to Windows. To see the updates that have been installed, you can select the **Show Updates** check box at the top of the list.

Any program can take advantage of this feature by marking their updates so that they are hidden when appropriate. Windows programs that were written before the release of Windows XP Service Pack 2 will be shown regardless of the filter option selection. Windows XP Service Pack 2 and Windows Server 2003 Service Pack 1 will also always be shown in the list

**How do I resolve these issues?**
To turn off the filter feature on a single computer, use the following procedure.

### To turn off filter for Add or Remove Programs

1. Open the Registry Editor by clicking **Start**, then clicking **Run**, typing `regedit`, and pressing **Enter**.
2. Navigate to the following registry key: `\\HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall`
3. Select the **DontGroupPatches** key.
4. By default, the **DWORD** value is equal to **0**. Change the **DWORD** value to **1** to disable the filter feature.

| ⚠Caution |
|---|
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

In an enterprise environment, you can create a Group Policy object to modify the registry setting that controls the filter feature to make **Add or Remove Programs** work as it did previously.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
**Filter for Add or Removed Programs Settings**

| Setting Name | Location | Previous Default Value | Default Value | Possible Values |
|---|---|---|---|---|

| | | | | |
|---|---|---|---|---|
| DontGroupPatches | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\ CurrentVersion\ Policies\Uninstall | N/A | REG_DWORD:0 | REG_DWORD:1 |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
Programs do not need to change in order to continue to work with the filter for **Add or Remove Programs**. If a program is not changed to use the new feature, it will continue working as it did previously.

A program can take advantage of the new filtering option by marking its updates so that they are not shown by default. Details on how to mark programs as updates are available on MSDN in "Working with Add or Remove Programs in the Control Panel Application" which is located on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=35318.

# DCOM Security Enhancements

Applies To: Windows Server 2003 with SP1

**What does DCOM do?**
The Microsoft Component Object Model (COM) is a platform-independent, distributed, object-oriented system for creating binary software components that can interact. The Distributed Component Object Model (DCOM) allows applications to be distributed across locations that make the most sense to you and to the application. The DCOM wire protocol transparently provides support for reliable, secure, and efficient communication between COM components. For more information, see "Component Object Model" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=20922.

**Who does this feature apply to?**
If you only use COM for in-process COM components, this section does not apply to you.

This feature applies to you if you have a COM server application that meets one of the following criteria:

- The access permission for the application is less stringent than the launch permission that is necessary to run the application.
- The application is usually activated by a remote COM client without using an administrative account.
- The application is only meant to be used locally. This means you can restrict your COM server application so it is not remotely accessible.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Computerwide restrictions**
**Detailed description**
A change has been made in COM to provide computerwide access controls that govern access to all call, activation, or launch requests on the computer. The simplest way to think about these access controls is as an additional AccessCheck call that is done against a computerwide access control list (ACL) on each call, activation, or launch of any COM server on the computer. If the AccessCheck fails, the call, activation, or launch request will be denied. (This is in addition to any AccessCheck that is run against the server-specific ACLs.) In effect, it provides a minimum authorization standard that must be passed to access any COM server on the computer. There will be a computerwide ACL for launch permissions to cover activate and launch rights, and a computerwide ACL for access permissions to cover call rights. These can be configured through the Component Services Microsoft Management Console (MMC).

These computerwide ACLs provide a way to override weak security settings specified by a specific application through CoInitializeSecurity or application-specific security settings. This provides a minimum security standard that must be passed, regardless of the settings of the specific server.

These ACLs are checked when the interfaces exposed by RPCSS are accessed. This provides a method to control who has access to this system service.

These ACLs provide a centralized location where an administrator can set general authorization policy that applies to all COM servers on the computer.

By default, Windows Server 2003 SP1 computer restriction settings are as follows.

**DCOM Computer Restriction Settings**

| Permission | Administrator | Distributed COM Users (Built-in Group) | Everyone | Anonymous |
|---|---|---|---|---|
| Launch | Local Launch<br>Local Activation<br>Remote Launch<br>Remote Activation | Local Launch<br>Local Activation<br>Remote Launch<br>Remote Activation | Local Launch<br>Local Activation | N/A |
| Access | N/A | Local Access<br>Remote Access | Local Access<br>Remote Access | Local Access<br>Remote Access |

**Note**

**Distributed COM Users** is a new built-in group included with Windows Server 2003 Service Pack 1 to expedite the process of adding users to the DCOM computer restriction settings.

**Why is this change important?**
Many COM applications include some security-specific code (for example, calling CoInitializeSecurity), but use weak settings, often allowing unauthenticated access to the process. There is currently no way for an administrator to override these settings to force stronger security in earlier versions of Windows.

COM infrastructure includes the RPCSS, a system service that runs during system startup and always runs after that. It manages activation of COM objects and the running object table and provides helper services to DCOM remoting. It exposes RPC interfaces that can be called remotely. Because some COM servers allow unauthenticated remote access (as explained in the previous section), these interfaces can be called by anyone, including unauthenticated users. As a result, RPCSS can be attacked by malicious users using remote, unauthenticated computers.

In earlier versions of Windows, there was no way for an administrator to understand the exposure level of the COM servers on a computer. An administrator could get an idea of the exposure level by systematically checking the configured security settings for all the registered COM applications on the computer, but, given that there are about 150 COM servers in a default installation of Windows, that task was daunting. There was no way to view the settings for a server that incorporates security in the software, short of reviewing the source code for that software.

DCOM computerwide restrictions mitigate these three problems. They also give an administrator the capability to disable incoming DCOM activation, launch, and access calls.

**What works differently?**
By default, the **Everyone** group is granted local launch, local activation, and local access permissions. This should enable all local scenarios to work without modification to the software or the operating system.

By default, the **Everyone** and **Anonymous** groups are granted remote access permissions. This enables most COM client scenarios, including the common case where a COM client passes a local reference to a remote server, in effect turning the client into a server.

Also by default, only members of the **Administrators** group are granted remote activation and launch permissions. This disables remote activations by non-administrators to installed COM servers.

**How do I resolve these issues?**
If you implement a COM server and expect to support remote activation by a non-administrative COM client or remote unauthenticated calls, then you should consider whether the risk associated with enabling this process is acceptable or if you should modify your implementation to not require remote activation by a non-administrative COM client.

If the risk is acceptable and you want to enable remote activation by a non-administrative COM client, you will need to change the default configuration for this feature.

You can change the configuration settings using either the Component Services Microsoft Management Console (MMC) or the Windows registry.

If you use the Component Services MMC snap-in, these settings can be configured on the **COM Security** tab of the **Properties** dialog box for the computer you are managing. The **COM Security** tab has been modified to enable you to set computerwide limits in addition to the standard default settings for COM servers. Additionally, you can provide separate ACL settings for local and remote access under both limits and defaults.

Alternatively, you can configure these ACL settings using the registry.

These ACLs are stored in the registry at the following locations:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Ole\MachineAccessRestriction= ACL
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Ole\MachineLaunchRestriction= ACL
```

These are named values of the REG_BINARY type that contain data describing the ACL of the principals that can access any COM class or COM object on the computer. The access rights in the ACL are:

```
COM_RIGHTS_EXECUTE 1
COM_RIGHTS_EXECUTE_LOCAL 2
COM_RIGHTS_EXECUTE_REMOTE 4
COM_RIGHTS_ACTIVATE_LOCAL 8
COM_RIGHTS_ACTIVATE_REMOTE 16
```

These ACLs can be created using normal security functions.

| **Note** |
| --- |
| COM_RIGHTS_EXECUTE rights must always be present; absence of this right will generate an invalid security descriptor. Only users with Administrator rights can modify these settings. |

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**RPCSS runs as a network service**
**Detailed description**
RPCSS is a key service for the RPC Endpoint Mapper and DCOM infrastructure that is network facing. This service ran as Local System in previous versions of Windows. To reduce the attack surface of Windows and provide defense in depth, the RPCSS service functionality was split into two services. The RPCSS service with all the original functionality that did not require Local System privileges now runs under the Network Service account. A new DCOMLaunch service that includes functionality that requires Local System privileges runs under the Local System account; however, this service is not network facing.

**Why is this change important?**
This change reduces the attack surface and provides defense in depth for the RPCSS service because an elevation of privilege in the RPCSS service will now be limited to the Network Service privilege.

**What works differently?**
This change should be transparent to users because the combination of the RPCSS and DCOMLaunch services is equivalent to the previous RPCSS service provided in prior versions of Windows.

**More specific COM permissions**
**Detailed description**
COM server applications have two types of permissions: launch permissions and access permissions. Launch permissions control authorization to start a COM server during COM activation if the server is not already running. These permissions are defined as security descriptors that are specified in registry settings. Access permissions control authorization to call a running COM server. These permissions are defined as security descriptors provided to the COM infrastructure through the `CoInitializeSecurity` API, or using registry settings. Both launch and access permissions allow or deny access based on principals, and make no distinction as to whether the caller is local to the server or remote.

The first change distinguishes the COM access rights, based on distance. The two distances that are defined are Local and Remote. A Local COM message arrives by way of the Lightweight Remote Procedure Call (LRPC) protocol, while a Remote COM message arrives by way of a remote procedure call (RPC) host protocol like transmission control protocol (TCP).

COM activation is the act of getting a COM interface proxy on a client by calling `CoCreateInstance` or one of its variants. As a side effect of this activation process, sometimes a COM server must be started to satisfy the client's request. A launch permissions ACL asserts who is allowed to start a COM server. An access permissions ACL asserts who is allowed to activate a COM object or call that object after the COM server is already running.

The second change is that the call and activation rights are separated to reflect two distinct operations and to move the activation right from the access permission ACL to the launch permission ACL. Since activation and launching are both related to acquiring an interface pointer, activation and launch access rights logically belong together in one ACL. Because you always specify launch permissions through configuration (as compared to access permissions, which are often specified programmatically), putting the activation rights in the launch permission ACL provides the administrator with control over activation.

The Launch Permission access control entries (ACEs) are broken into four access rights:

- Local Launch (LL)
- Remote Launch (RL)
- Local Activation (LA)
- Remote Activation (RA)

The Access Permission security descriptor is split into two access rights:

- Local Access calls (LC)

- Remote Access calls (RC)

This allows the administrator to apply very specific security configurations. For example, you can configure a COM server so that it accepts local calls from everyone, while only accepting remote calls from Administrators. These distinctions can be specified through changes to the COM Permissions security descriptors.

**Why is this change important? What threats does it help mitigate?**
Earlier versions of the COM server application have no way to restrict an application so that it can only be used locally without exposing the application on the network by way of DCOM. When a user has access to a COM server application, he or she has access for both local and remote use.

A COM server application might want to expose itself to unauthenticated users to implement a COM callback scenario. In this scenario, the application must also expose its activation to unauthenticated users, which might not be desirable because malicious users could use that scenario to gain unauthorized access to that server.

Precise COM permissions give flexibility to the administrator to control a computer's COM permission policy. These permissions enable security for the described scenarios.

**What works differently? Are there any dependencies?**
To provide backwards compatibility, existing COM security descriptors are interpreted to allow or deny both local and remote access simultaneously. That is, an access control entry will either allow both local and remote, or deny both local and remote.

There are no backwards-compatibility issues for call or launch rights. There is, however, an activation rights compatibility issue. If, in the existing security descriptors for a COM server, the configured launch permissions are more restrictive than the access permissions and are more restrictive than what is minimally needed for client activation scenarios, then the Launch Permissions ACL must be modified to give the authorized clients the appropriate activation permissions.

For COM applications that use the default security settings, there are no compatibility issues. For applications that are dynamically started using COM activation, most will have no compatibility issues, since the launch permissions must already include anyone who is able to activate an object. Otherwise, such applications would have generated activation failures even before applying Service Pack 1, when callers without launch permission try to activate an object and the COM server is not already running.

The applications of most concern for compatibility issues are COM applications that are already started by way of some other mechanism, such as Windows Explorer or Service Control Manager. You can also start these applications by way of a previous COM activation, which overrides the default access and launch permissions and specifies launch permissions that are more restrictive than the call permissions. For more details on addressing this compatibility issue, see "How do I resolve these issues?" in the next section.

If a system that was updated to Windows Server 2003 Service Pack 1 is rolled back to an earlier service pack, any access control entry that was edited to allow local access, remote access, or both, will be interpreted to allow both local and remote access. Any ACE that was edited to deny local access, remote access, or both, will be interpreted to deny both local and remote access. Whenever you uninstall a service pack, you should ensure that no newly set ACEs will cause applications to stop working.

**How do I resolve these issues?**
If you implement a COM server and you override the default security settings, confirm that the application-specific launch permissions ACL grants activation permission to appropriate users. If it does not, you will need to change your application-specific launch permission ACL to give appropriate users activation rights so applications and Windows components that use DCOM do not fail. These application-specific launch permissions are stored in the registry.

The COM ACLs can be created or modified using normal security functions.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**

| ⚠**Caution** |
| --- |
| Improper use of these settings can cause applications and Windows components that use DCOM to fail. |

In the following table, these abbreviations are used:

LL - Local Launch

LA - Local Activation

RL - Remote Launch

RA - Remote Activation

LC - Local Access calls

RC - Remote Access calls

**DCOM Settings**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| MachineLaunch Restriction | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft\Ole\ | Everyone: LL, LA, RL, RA<br><br>Anonymous: LL, LA, RL, RA<br><br>(This is a new registry key. Based on existing behavior, these would be the effective values.) | Administrator: LL, LA, RL, RA<br><br>Everyone: LL, LA<br><br>Distributed COM users: LL, LA, RL, RA | ACL |
| MachineAccess Restriction | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft\Ole\ | Everyone: LC, RC<br><br>Anonymous: LC, RC<br><br>(This is a new registry key. Based on existing behavior, these would be the effective values.) | Everyone: LC, RC<br><br>Anonymous: LC, RC | ACL |
| CallFailure LoggingLevel | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft \Ole\ | Not applicable. | This registry key is not present; however, a missing key or value is interpreted as 2.<br><br>This event is not logged by default. If you change this value to 1 to start logging this information to help you troubleshoot an issue, be sure to monitor the size of your event log, as this is an event that can generate a large number of entries. | 1 - Always log event log failures when COM infrastructure finds an invalid security descriptor.<br><br>2 - Never log event log failures when COM infrastructure finds an invalid security descriptor. |
| InvalidSecurity Descriptor LoggingLevel | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft\Ole\ | Not applicable. | This registry key is not present, however a missing key or value is interpreted as 1.<br><br>This event is logged by default. It should rarely occur. | 1 - Always log event log failures when COM infrastructure finds an invalid security descriptor.<br><br>2 - Never log event log failures when COM infrastructure finds an invalid security descriptor. |

| | | | | |
|---|---|---|---|---|
| DCOM:Machine Launch Restrictions in Security Descriptor Definition Language (SDDL) Syntax | (Group Policy object) Computer Configuration \Windows Settings \Local Policies \Security Options | Not applicable. | Not defined. | Access Control List in SDDL Format. Existence of this policy, overrides, values in MachineLaunch Restriction, above. |
| DCOM:Machine Access Restrictions in Security Descriptor Definition Language (SDDL) Syntax | (Group Policy object) Computer Configuration \Windows Settings \Local Policies \Security Options | Not applicable. | Not defined. | Control List in SDDL Format. Existence of this policy, overrides, values in MachineAccess Restriction, above. |

# Data Execution Prevention

Applies To: Windows Server 2003 with SP1

**What does data execution prevention do?**
Data execution prevention (DEP) is a set of hardware and software technologies that perform additional checks on memory to help protect against malicious code exploits. In Windows Server 2003 with Service Pack 1, DEP is enforced by both hardware and software.

**Hardware-enforced DEP**
Hardware-enforced DEP marks all memory locations in a process as non-executable unless the location explicitly contains executable code. There is a class of attacks that attempt to insert and execute code from non-executable memory locations. DEP helps prevent these attacks by intercepting them and raising an exception.

Hardware-enforced DEP relies on processor hardware to mark memory with an attribute that indicates that code should not be executed from that memory. DEP functions on a per-virtual-memory-page basis, usually changing a bit in the page table entry (PTE) to mark the memory page.

The actual hardware implementation of DEP and marking of the virtual memory page varies by processor architecture. However, processors that support hardware-enforced DEP are capable of raising an exception when code is executed from a page marked with the appropriate attribute set.

Both Advanced Micro Devices (AMD) and Intel Corporation have defined and shipped Windows-compatible architectures that are compatible with DEP.

32-bit versions of Windows Server 2003 with Service Pack 1 utilize the no-execute page-protection (NX) processor feature as defined by AMD or the Execute Disable bit (XD) feature as defined by Intel. In order to use these processor features, the processor must be running in Physical Address Extension (PAE) mode. The 64-bit versions of Windows use the NX or XD processor feature on 64-bit extensions processors and certain values of the access rights page table entry (PTE) field on IPF processors.

It is hoped that all future 32-bit and 64-bit processors will provide support for hardware-enforced DEP. Microsoft continues to work with processor vendors to encourage the adoption and development of DEP technologies.

**Software-enforced DEP**
An additional set of DEP security checks has been added to Windows Server 2003 with Service Pack 1. These checks, known as software-enforced DEP, are designed to mitigate exploits of exception handling mechanisms in Windows. Software-enforced DEP runs on any processor that is capable of running Windows Server 2003 with Service Pack 1. By default, software-enforced DEP protects only limited system binaries, regardless of the hardware-enforced DEP capabilities of the processor.

**Who does this feature apply to?**
Application and driver developers should be aware of DEP and the requirements of software running on a supporting platform. Applications that perform just-in-time (JIT) code generation or execute memory from the default process stack or heap should pay careful attention to DEP requirements.

Driver developers are encouraged to be aware of PAE mode on platforms supporting hardware-enforced DEP. PAE mode behavior on systems running Windows Server 2003, Standard Edition with Service Pack 1, is changed to improve driver compatibility.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Data execution prevention on 32-bit and 64-bit versions of Windows and applications**
**Detailed description**
**Hardware-enforced DEP**
To provide consistency for application and driver developers, the memory protection model (including DEP) is designed to be the same for both 32-bit and 64-bit versions of Windows.

Application developers should be aware of DEP behavior in user mode. A user-mode DEP exception results in a `STATUS_ACCESS_VIOLATION (0xc0000005)` on Windows systems. The first parameter of `ExceptionInformation` that is located inside the `EXCEPTION_RECORD` structure contains the type of access violation that occurred. A value of 8 for `ExceptionInformation[0]` indicates the access violation was an execution violation.

In most processes, the `STATUS_ACCESS_VIOLATION` exception will be an unhandled exception and result in termination of the process.

DEP is also applied to drivers in kernel mode. DEP for memory regions in kernel mode cannot be selectively enabled or

disabled. On 32-bit versions of Windows, DEP is applied to the stack by default. This differs from kernel-mode DEP on 64-bit versions of Windows, where the stack, paged pool, and session pool have DEP applied.

Device drivers are not permitted to execute code from the stack when DEP is enabled. A DEP access violation in kernel mode will result in an error `0xFC: ATTEMPTED_EXECUTE_OF_NOEXECUTE_MEMORY`.

**Software-enforced DEP**

Software-enforced DEP performs additional checks on exception handling mechanisms in Windows. If the program's image files are built with Safe Structured Exception Handling (SafeSEH), software-enforced DEP ensures that before an exception is dispatched, the exception handler is registered in the function table located within the image file.

If the program's image files are not built with SafeSEH, software-enforced DEP ensures that before an exception is dispatched, the exception handler is located within a memory region marked as executable.

**DEP application close behavior**

The majority of applications will not encounter a problem with DEP. However, when an application does encounter a problem with DEP, a Data Execution Prevention message is presented to the user, alerting them to the problem.

The Data Execution Prevention message indicates that a DEP problem occurred with the application and provides the ability for the user to learn more about DEP and optionally disable DEP for the application that was closed.

**Important**

If a DEP problem occurs with an application, Microsoft recommends contacting the application vendor for an update. The security implications of disabling DEP for an application should be thoroughly considered before disabling DEP for an application.

The ability to change DEP settings for a closed application using the **Change Settings** button on the Data Execution Prevention message window depends on the system-wide DEP configuration. Changes to DEP protection for an application can be made only if the system-wide DEP configuration is set to the **OptOut** mode.

The Data Execution Prevention message is presented immediately before a Windows Error Reporting window, which provides the opportunity to submit a report about the DEP problem to Microsoft.

On Windows Server 2003 with Service Pack 1, the Data Execution Prevention message will be presented the next time an Administrator logs onto the system interactively. The behavior is changed on Windows Server 2003 with Service Pack 1 from Windows XP with Service Pack 2 because Windows Error Reporting is configured in queued mode by default. Queued mode causes error reporting messages to be queued until the next time an administrator interactively logs onto the system.

To learn more about Windows Error Reporting, or to configure Windows Error Reporting such that DEP and Windows Error Reporting messages are displayed immediately following an application problem, see the "Using Windows Server 2003 in a Managed Environment: Windows Error Reporting" article on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=38443.

Windows Error Reporting generates an error signature when an application is closed due to DEP. The error signature can be viewed by following the "click here" link on the Windows Error Reporting dialog. The error signature for a DEP problem has the following attributes:

| Parameter | Example Value | Description |
|---|---|---|
| EventType | BEX | Indicates a buffer overflow (/GS) or DEP exception (BEX64 indicates a buffer overflow (/GS) or DEP exception on 64-bit versions of Windows) |
| P1 | DEPDemo.exe | Name of the application executable that encountered the problem |
| P2 | 5.1.2600.2180 | Version of the application executable that encountered the problem |
| P3 | 416725f2 | Faulting application stamp |
| P4 | DEPDemo.exe | Faulting module name |
| P5 | 5.1.2600.2180 | Faulting module version |
| P6 | 416725f2 | Faulting module stamp |

| | | |
|---|---|---|
| P7 | 00002060 | Fault offset (Instruction address if a module is not loaded at the faulting address) |
| P8 | C0000005 | Indicates a STATUS_ACCESS_VIOLATION exception (If this parameter is c0000409, the problem is a /GS-related fault) |
| P9 | 00000008 | Indicates an execution STATUS_ACCESS_VIOLATION (00000002 indicates an execution STATUS_ACCESS_VIOLATION on 64-bit versions of Windows for the Intel Itanium architecture) |

Finally, the Data Execution Prevention message might not be shown for some applications when they encounter a problem with DEP, regardless of the Windows Error Reporting configuration. These applications handle the STATUS_ACCESS_VIOLATION exception raised by DEP, or they install an unhandled exception filter (UEF), which overrides the default Win32 UEF. The default Win32 UEF is responsible for triggering both the Data Execution Prevention and Windows Error Reporting messages. The Data Execution Prevention message may also not be shown if an application has called the `SetErrorMode()` function with the `SEM_NOGPFAULTERRORBOX` flag.

**Why is this change important? What threats does it help mitigate?**
The primary benefit of DEP is that it helps to prevent code execution from data pages such as the default heap, various stacks, and memory pools. In normal operations of the system, code is not typically executed from the default heap and stack. Hardware-enforced DEP detects code that is running from these locations and raises an exception when execution occurs. If the exception is unhandled, the process will be terminated. Execution of code from protected memory in kernel mode results in an error.

Although terminating a process or causing the system to fail with an error do not appear to be ideal experiences, this helps prevent malicious code from executing. Preventing malicious code from executing on the system may prevent damage to the system or propagation of malicious code whose harmful effects could easily exceed those of a terminated process or system error.

DEP can help prevent some exploits in which a virus or other attack has injected a process with executable code and then attempts to execute the injected code. On a system with DEP, execution of the injected code should result in an exception. Software-enforced DEP can help mitigate exploits of exception handling mechanisms within Windows.

A secondary benefit of DEP relates to good engineering and best practices for application and driver developers. DEP forces developers to avoid executing code out of data pages without explicitly marking the pages as executable.

**What works differently?**
**Application Compatibility**
Some application behaviors are expected to be incompatible with DEP. Applications that perform dynamic code generation (such as just-in-time code generation) and that do not explicitly mark generated code with Execute permission might have compatibility problems with DEP. Applications that are not built with SafeSEH must have their exception handlers located in executable memory regions.

Applications that attempt to violate DEP will receive an exception with status code `STATUS_ACCESS_VIOLATION (0xC0000005)`. If an application requires executable memory, it must explicitly set this attribute on the appropriate memory by specifying `PAGE_EXECUTE`, `PAGE_EXECUTE_READ`, `PAGE_EXECUTE_READWRITE` or `PAGE_EXECUTE_WRITECOPY` in the memory protection argument of the `Virtual*` memory allocation functions. Heap allocations using the `malloc()` and `HeapAlloc()` functions are non-executable.

**Driver compatibility**
Driver compatibility issues with DEP mostly center on PAE mode-induced compatibility issues.

| **Note** |
|---|
| PAE is required only on systems running 32-bit versions of Windows with processors that support hardware-enforced DEP. |

On its own, DEP might create compatibility problems with drivers that perform code generation or use other techniques to generate executable code in real time. Although many drivers with such behavior would have been fixed — as DEP is "always on" for drivers loaded on 64-bit versions of Windows — there is no guarantee that all drivers have been updated. However, there are few drivers that employ these techniques, and it is not expected that DEP alone will cause a large quantity of driver compatibility problems.

The primary driver compatibility concern is running Physical Address Extension (PAE) mode on 32-bit systems. PAE mode enables processors to address greater than 4 gigabytes (GB) of memory. The primary difference between PAE memory paging and non-PAE memory paging schemes is the extra level of paging that is required in PAE mode (three levels instead of two).

Some drivers might fail to load if PAE is enabled, because the device might be unable to perform 64-bit addressing or the

drivers might assume that PAE mode requires more than 4 GB of random access memory (RAM). Such drivers expect that they will always receive 64-bit addresses when in PAE mode and that they or their device are incapable of interpreting the address.

Other drivers might load in PAE mode but cause system instability by directly modifying system page table entries (PTEs). These drivers expect 32-bit PTEs, but receive 64-bit PTEs in PAE mode instead.

The largest driver PAE compatibility issue involves direct memory access (DMA) transfers and map register allocation. Many devices that support DMA, usually 32-bit adapters, are not capable of performing 64-bit physical addressing. When run in 32-bit mode, the device can address all physical address space. In PAE mode, it is possible that data would be present at a physical address greater than 4 GB. To allow devices with these constraints to function in this scenario, Windows 2000 Server and later operating systems provide double-buffering for the DMA transaction by providing a 32-bit address that is indicated by a map register. The device can perform the DMA transaction to the 32-bit address and the kernel copies the memory to the 64-bit address that is provided to the driver.

When the system runs with PAE disabled, drivers for 32-bit devices never require their map registers to be backed by real memory. This means that double-buffering is not necessary, since all devices and drivers are contained within the 32-bit address space. Based on testing of drivers for 32-bit devices on 64-bit processor–based computers, it is expected that most client-tested, DMA-capable drivers expect unlimited map registers.

To constrain compatibility issues, Windows Server 2003, Standard Edition with Service Pack 1, includes hardware abstraction layer (HAL) changes that mimic the 32-bit HAL DMA behavior. The altered HAL grants unlimited map registers when the system is running in PAE mode. In addition, the kernel memory manager ignores any physical address above 4 GB. Any system RAM beyond the 4 GB barrier would be made unaddressable by Windows and be unusable in the system. By limiting the address space to 4 GB, devices with 32-bit DMA bus master capability will not see a transaction with an address above the 4 GB barrier. Because these changes remove the need to double-buffer the transactions, they avoid a class of bugs in some drivers related to proper implementation of double buffering support.

Note that the PAE behavior of Windows Server 2003, Enterprise Edition with Service Pack 1, and Windows Server 2003, Datacenter Edition with Service Pack 1, is unchanged from the versions without the service pack.

As a result of these changes to the HAL and memory manager, the impact to device driver compatibility is expected to be minimal on systems running Windows Server 2003 with Service Pack 1 with DEP enabled.

**System compatibility**

A final DEP compatibility concern derives from systems with PAE mode enabled, even though they may not be designed for more than 4 GB of physical RAM. During internal testing Microsoft has noticed that some systems with processors that support hardware-enforced DEP fail to start up or have other stability issues when the processor is running in PAE mode.

PAE mode is a requirement for leveraging the NX processor feature. Therefore, system designers and firmware engineers should be aware that even though the system's chipset and firmware may not have been designed to support more than 4 GB of physical RAM, the system may be running in PAE mode.

Of particular concern is system firmware that interprets page table entries to determine instructions executed by the operating system. Page table entries are extended to 64 bits in length when the processor is running in PAE mode. System designers and firmware developers are encouraged to contact their processor and chipset vendors for more information about how to safely determine instructions executed by the operating system.

System designers working with AMD processors can obtain more information in the "BIOS and Kernel Developer's Guide for AMD Athlon 64 and AMD Opteron Processors." To obtain this paper, go to the AMD Athlon 64 Web site at http://go.microsoft.com/fwlink/?LinkId=28165 and click "BIOS and Kernel Developer's Guide for AMD Athlon 64 and AMD Opteron Processors."

Intel does not make detailed information about System Management Mode (SMM) available publicly. System designers working with Intel processors are encouraged to contact Intel directly for more information.

For more information regarding Windows support for PAE mode, see "Physical Address Extension - PAE Memory and Windows" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=45912.

**How do I resolve these issues?**

Applications that require executable regions of memory must use the `PAGE_EXECUTE`, `PAGE_EXECUTE_READ`, `PAGE_EXECUTE_READWRITE`, or `PAGE_EXECUTE_WRITECOPY` attributes when allocating memory. Additionally, applications cannot execute from the default process heap or the stack. Most applications that perform actions incompatible with DEP will need to be updated to be compatible. Applications must also be built with SafeSEH or ensure their exception handlers are located in memory explicitly marked as executable.

An application can use the `VirtualAlloc()` application programming interface (API) function to allocate executable memory

with the appropriate memory protection options. At a minimum, the `PAGE_EXECUTE` memory protection option should be used. After the executable code has been generated, it is recommended that the application set memory protections to disallow write access to the allocated memory. Applications can disallow write access to allocated memory using the `VirtualProtect()` API function. Disallowing write access ensures maximum protection for executable regions of process address space.

If a malicious process attempts to insert code into an executable region, the access would result in a `STATUS_ACCESS_VIOLATION` write exception. The application should attempt to make the executable regions of its address space as small as possible. This would result in a smaller attack surface through which executable memory could be injected into the process address space and be executed.

Additionally, sophisticated applications can control the layout of their virtual memory and create executable regions. These applications should attempt to locate executable regions in a lower memory space than non-executable regions. The purpose of locating executable regions below non-executable regions is to protect a buffer overflow from overflowing into executable memory.

A small number of executables and libraries may contain executable code in a data section of the image file. In some cases, applications may place small segments of code (commonly referred to as thunks) in the data sections. However, DEP will mark sections of the image file loaded in memory as non-executable unless the section has the executable attribute applied.

Therefore, executable code in data sections should be moved to a code section, or the data section containing the executable code should be explicitly marked as executable. The executable attribute, `IMAGE_SCN_MEM_EXECUTE (0x20000000)`, should be added to the Characteristics field of the corresponding section header for sections that contain executable code.

The Microsoft linker that is distributed with Microsoft Visual Studio products can add the executable attribute to a section using the `/SECTION` linker option. The `/SECTION` linker option has the following format:

**/SECTION:** *Name* **,[E][R][W][S][D][K][L][P][X][,ALIGN=#]**

The `E` value indicates the executable attribute (0x20000000). More information about `/SECTION` and other Microsoft linker options is available on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=28167.

Additionally, the Microsoft COFF Binary File Editor (Editbin.exe) utility can be used to change the section attributes of an existing image. The Editbin utility has a `/SECTION` option with the following format:

**/SECTION:** *Name* **[=** *newname* **][,[[!]{CDEIKOMPRSUW}][A{1248PTSX}]]**

The `C` and `E` values indicate code and executable attributes respectively. For more information about the Editbin utility and the `/SECTION` option, see the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=28168.

Microsoft has provided service packs for Microsoft .NET Framework version 1.0 and version 1.1 to take advantage of DEP in Windows Server 2003 with Service Pack 1. Applications that use the Microsoft .NET Framework will continue to function normally, but will not benefit from DEP if it is enabled unless the appropriate Microsoft .NET Framework Service Pack has been installed.

Microsoft encourages application developers who redistribute the Microsoft .NET Framework to update to Microsoft .NET Framework version 1.0 Service Pack 3 or version 1.1 Service Pack 1, which take advantage of DEP.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
**System-wide configuration of DEP**
The primary difference in DEP behavior on Windows Server 2003 Service Pack 1 as compared to Windows XP Service Pack 2 (SP2) is that on the server operating system the default configuration is to protect all applications and services. In Windows XP SP2, DEP was turned on by default only for essential Windows operating system programs and services.

DEP configuration for the system is controlled through Boot.ini switches. Additionally, changes to System in Control Panel have been made to enable end users to easily configure DEP settings if they are logged onto the system as an administrator.

System DEP configuration settings apply only for 32-bit applications and processes when running on 32-bit or 64-bit versions of Windows. On 64-bit versions of Windows, if hardware-enforced DEP is available it is always applied to 64-bit processes and kernel memory spaces and there are no system configuration settings to disable it.

Windows supports four system-wide configurations for both hardware-enforced and software-enforced DEP.

**DEP Configuration**

| Configuration | Description |
|---|---|

| | |
|---|---|
| OptIn<br><br>(default for Windows XP SP2 and Windows XP 64-bit edition) | On systems with processors capable of hardware-enforced DEP, DEP is enabled by default for limited system binaries and applications that opt in. With this option, only Windows system binaries are covered by DEP by default. |
| OptOut<br><br>(default for Windows Server 2003 Service Pack 1) | DEP is enabled by default for all processes. Users can manually create a list of specific applications that do not have DEP applied using System in Control Panel. IT pros can use the Application Compatibility Toolkit to opt-out one or more applications from DEP protection. System Compatibility Fixes ("shims") for DEP do take effect. |
| AlwaysOn | This provides full DEP coverage for the entire system. All processes always run with DEP applied. The exceptions list for exempting specific applications from DEP protection is not available. System Compatibility Fixes ("shims") for DEP do not take effect. Applications that have been opted-out using the Application Compatibility Toolkit run with DEP applied. |
| AlwaysOff | This does not provide any DEP coverage for any part of the system, regardless of hardware DEP support. However, the processor will run in PAE mode with 32-bit versions of Windows unless the `/noexecute=alwaysoff` option is replaced with the `/execute` option in the boot entry. |

Hardware-enforced and software-enforced DEP are configured in the same manner. If the system-wide DEP policy is set to OptIn, the same Windows core binaries and applications will be protected by both hardware and software-enforced DEP. If the system is not capable of hardware-enforced DEP, the Windows core binaries and applications will be protected only by software-enforced DEP.

Similarly, if the system-wide DEP policy is set to OptOut, applications that have been exempted from DEP protection will be exempted from both hardware and software-enforced DEP.

The four system-wide DEP configurations are controlled through Boot.ini switches. The Boot.ini settings are as follows:

`/noexecute=policy_level`

where *policy_level* is defined as `AlwaysOn`, `AlwaysOff`, `OptIn`, or `OptOut`.

Any existing `/noexecute` setting in the Boot.ini file is not changed when Windows Server 2003 Service Pack 1 is installed or if a Windows operating system image is moved across computers with and without hardware-enforced DEP support.

During installation of Windows Server 2003 Service Pack 1, the `OptOut` policy level is enabled by default unless a different policy level is specified in an unattended installation. If the `/noexecute=`*policy_level* setting is not present in the boot entry for a version of Windows that supports DEP, the behavior is the same as if the `/noexecute=OptIn` option was included.

End users who are logged on as administrators can manually configure DEP between the `OptIn` and `OptOut` policies using the **Data Execution Prevention** tab inside the **Performance Options** dialog box. The following procedure describes how to manually configure DEP on the computer:

**To configure DEP settings**

1. Click **Start**, click **Control Panel**, and then double-click **System**.
2. Click the **Advanced** tab. Then, under **Performance**, click **Settings**.
3. Click the **Data Execution Prevention** tab.
4. Click **Turn on DEP for essential Windows programs and services only** to select the **OptIn** policy.
5. Click **Turn on DEP for all programs and services except those I select** to select the **OptOut** policy.
6. If you selected the **OptOut** policy, click **Add** and add the applications that you do not want to use DEP with.

IT professionals can control system-wide DEP configuration with a variety of methods. The Boot.ini file can be modified directly with scripting mechanisms or with the Bootcfg.exe tool, which is included as part of Windows Server 2003 Service Pack 1.

For unattended installations of Windows Server 2003 with Service Pack 1, you can use the Unattend.txt file to prepopulate a specific DEP configuration. You can use the `OSLoadOptionsVar` entry in the `[Data]` section of the Unattend.txt file to specify a

system-wide DEP configuration.

**Per-application DEP configuration**
For the purposes of application compatibility when DEP is set to the `OptOut` policy level, it is possible to selectively disable DEP for individual 32-bit applications. However, DEP is always enabled for 64-bit applications.

For end users, the **Data Execution Prevention** tab in **System Properties** can be used to selectively disable DEP for an application.

For IT professionals, a new application compatibility fix named DisableNX is included with Windows Server 2003 Service Pack 1. The DisableNX compatibility fix disables DEP for the program it is applied to.

The DisableNX compatibility fix can be applied to an application by using the Application Compatibility Toolkit. For more information about Windows application compatibility, see "Windows Application Compatibility" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=23302.

# Device\PhysicalMemory Object

Applies To: Windows Server 2003 with SP1

**What does \Device\PhysicalMemory Object do?**
The \Device\PhysicalMemory section object is used by some applications to access physical memory. Common examples of this use are management applications that attempt to read System Management BIOS (SMBIOS) or other BIOS data that is stored in the lowest 1 MB of physical memory.

**Who does this feature apply to?**
This feature will be of interest to hardware and hardware management application developers.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Device\PhysicalMemory section object**
**Detailed description**
In Windows Server 2003 SP1, user-mode access to the \Device\PhysicalMemory object is not permitted. All forms of access (read, write) are refused when the \Device\PhysicalMemory object is accessed from a user-mode application. Access to the \Device\PhysicalMemory object is refused regardless of the user context (Administrators, Users, Local System, etc.) the application is running in.

Kernel-mode (driver) access to the \Device\PhysicalMemory object is unchanged for Windows Server 2003 SP1.

There are no means to revert this changed behavior in Windows Server 2003 SP1.

**Why is this change important? What threats does it help mitigate?**
This change was made to help prevent security exploits that might leverage the functionality of the \Device\PhysicalMemory object from user-mode.

In previous versions of Windows, access to the \Device\PhsyicalMemory object was protected by an access control list (ACL). However, the ACL may be inadvertently changed. Refusing all access to the \Device\PhysicalMemory object from user-mode ensures this scenario cannot occur.

**What works differently? Are there any dependencies?**
While the majority of applications do not attempt to access physical memory, some management applications that require information from the BIOS may be affected by this change. Many of these management applications may be attempting to read SMBIOS data.

**How do I fix these issues?**
Management application vendors are encouraged to review their applications for access to the \Device\PhysicalMemory object and implement one of the workarounds proposed in this section.

| **Note** |
| --- |
| The following methods can be used to read low physical memory and firmware table data from an application in Windows Server 2003 SP1. Applications that require write access to low physical memory must do so from a kernel-mode driver. |

**Reading SMBIOS information**
Windows Management Instrumentation (WMI) is the preferred method for reading SMBIOS information in Windows. WMI supports reading individual SMBIOS properties (such as the BIOS version or the system manufacturer) from the Win32 classes contained in the root\cimv2 namespace.

WMI also supports reading the entire contents of SMBIOS data in a single buffer using the MSSMBios_RawSMBiosTables class inside of the root\wmi namespace. The SMBiosData property returns a buffer containing the entire SMBIOS data table, except the SMBIOS Structure Table Entry Point as defined in section 2.1.1 of the SMBIOS specification.

WMI supports both a COM and script API to access WMI information. Additionally, WMI is supported in the Windows XP SP2 and Windows Server 2003 SP1 versions of the Windows Preinstallation Environment (WinPE).

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
If your application uses the \Device\PhysicalMemory object you will need to change your code for your application to work with Windows Server 2003 SP1. There are two new user-mode APIs that can be used by applications to read various system firmware tables. These APIs can be used to read SMBIOS and ACPI data, as well as raw BIOS data from low physical memory. The details of the API implementation are available in the Platform SDK at the MSDN Web site. The following are the two new APIs:

- `EnumSystemFirmwareTables`, described on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=45562.

- `GetSystemFirmwareTable`, described on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=45563.

**For more information**

The following Web resources may also be of interest to you if you are affected by this change in functionality:

- The Windows Management Instrumentation (WMI) section of the Windows Platform SDK available on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=159.

- The System Management BIOS (SMBIOS) standard specification provide by the Distributed Management Task Force, Inc. (DMTF) which is available on the DMTF Web site at http://go.microsoft.com/fwlink/?LinkId=38721.

- The Microsoft® Windows® Preinstallation Environment (Windows PE) tool, which is available for download on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=38722.

# Distributed File System (DFS) in Windows Server 2003 Service Pack 1

Applies To: Windows Server 2003 with SP1

**What does DFS do?**
Distributed File System (DFS) allows administrators to group shared folders located on different servers by transparently connecting them to one or more DFS namespaces. A DFS namespace is a virtual view of shared folders in an organization. Using the DFS tools, an administrator selects which shared folders to present in the namespace, designs the hierarchy in which those folders appear, and determines the names that the shared folders show in the namespace. When a user views the namespace, the folders appear to reside on a single, high-capacity hard disk. Users can navigate the folders in the namespace without needing to know the server names or shared folders hosting the data. DFS also provides many other benefits, including fault tolerance and load-sharing capabilities, making it ideal for all types of organizations

**Who does this feature apply to?**
Storage administrators managing a DFS namespace, or administrators considering deploying DFS, should be aware of the changes incorporated in SP1.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**DFS Consolidation Roots**
**Detailed description**
New functionality in Distributed File System (DFS) maintains the original Universal Naming Convention (UNC) path of files after they are moved to a new server. This reduces the impact of file server consolidation and migration, saves end users time spent searching for files, and ensures that line-of-business applications keep running.

**Why is this change important?**
There are many scenarios where you will want Universal Naming Convention (UNC) paths to remain unchanged when the underlying files are moved to other servers or to other paths. For example, you may want to preserve the UNC paths that users are accustomed to if you migrate or consolidate your existing file servers to new computers running Microsoft Windows Server 2003. The paths may be embedded in links, in line-of-business applications, and in other places where the names are difficult to change

**What works differently?**
The consolidation root functionality is most easily configured using the File Server Migration Toolkit, available on the Microsoft Web site at http://go.microsoft.com/fwlink/LinkId=37029.

This functionality can also be configured manually by modifying keys in the Windows Registry, as described in the article about using the Distributed File System update to support consolidation roots at the Microsoft Web site at http://go.microsoft.com/fwlink/LindId=43006.

**Move and Rename Support for DFS Namespaces**
**Detailed description**
Administrators can modify the DFS namespace to correct mistakes or to adjust the hierarchy as business needs change or as new folders are added to the namespace.

**Why is this change important?**
Prior to Windows Server 2003 Service Pack 1, there was no way for administrators to rename elements of their namespace. Renaming or moving an element near the top of the namespace was particularly difficult as it required manually deleting and recreating every element below it.

**What works differently?**
DFS administrators may now use the command-line utility **dfscmd** to move or rename links. The syntax for this new command is as follows:

**dfscmd /move ** *DFSName* **** *DFSShareName* **** *Path* **** *ServerName* **** *ShareName* **** *Path*

The following examples illustrate the different ways that this command can be used:

- To create a new folder, move the contents of an existing folder into it, and adjust any links under the folder as appropriate:

```
dfscmd /move \\domain\root\old_folder \\domain\root\new_folder
```

- To make an existing folder a subfolder of a new top-level folder and adjust any links under the folder as appropriate:

```
dfscmd /move \\domain\root\old_folder \\domain\root\new_top_folder\old_folder
```

- To move the contents of one folder to another existing folder

```
dfscmd /move \\domain\root\folder \\domain\root\other-folder\
```

**Client Failback**
**Detailed description**
Client failover in DFS namespaces is the process in which clients attempt to access another server in a referral after one of the servers fails or is removed from the namespace. DFS administrators can specify that clients "failback" to a preferred, local server when it is restored. The target failback policy may be set for the entire root, or for individual links.

**Why is this change important? What threats does it mitigate?**
Prior to Windows Server 2003 Service Pack 1, there was no way for administrators to configure clients to fail back to their local servers; this behavior can be problematic in branch office environments when clients continue to access the hub server even after the branch server is restored.

**What works differently?**
The failback settings for roots and links can be set using the **/TargetFailback** option of **dfsutil**. In addition, the DFS Client Failback hotfix for Windows XP and Windows Server 2003 must be installed to enable the client to perform the failback.

Setting the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters\SysvolNetlogonTargetFailback` registry key and restarting the DFS service on the domain controller enables target failback for SYSVOL/NETLOGON referral requests handled by that controller.

**Target Priority**
**Detailed description**
Target priority enables DFS administrators to explicitly assign the priority of a DFS link or root target.

**Why is this change important? What threats does it mitigate?**
Prior to Windows Server 2003 Service Pack 1, administrators had no way of designating preference between servers that host the same data in a namespace. To work around this, administrators created special subnets for such servers and artificial sites in Active Directory with high link cost. This workaround is time consuming, complex, and leads to increased support costs.

DFS management APIs have been added to configure whether a server appears first or last in a referral, which is the list of servers returned to the client for a given folder (link) in the namespace. Assigning target priority is useful in many scenarios, such as “hot-standby” scenarios where one server is considered the server of last resort. In this scenario, the administrator specifies that the standby server always appears last in referrals, and clients will fail over to this server only if all the other servers fail or become unavailable due to network outages.

**What works differently?**
Targets in a DFS referral response are grouped into sets based on the cost of accessing targets from a client. Within each set, targets are shuffled randomly and these sets are ordered by increasing cost from the DFS client. Prior to Windows Server 2003 SP1, the number of sets in the referral response depended on whether or not Windows Server 2003 site costing was enabled. With site costing disabled there are exactly two sets, one set containing targets in the same site as the DFS client and the other set containing the remaining targets. With site costing enabled sets are grouped on the basis of actual site cost of the target from the DFS client, and hence there could be many more of these sets. Note that there is no information in the current DFS V1, V2, or V3 referral responses that allows a DFS client to identify these sets.

With server target priority, these sets will still be based on the cost of accessing targets. Server target priority simply extends the cost sorting criteria for targets, so the sets are now those targets having the same site cost and server target priority. In addition, the new V4 protocol will indicate where the set boundaries are in the target list.

Server target priority is broken out into two values: a priority class and a priority rank. Priority classes are defined at two levels: locally, within sets of targets of equal site cost, and globally. Within each of these, there is a coarse ordering of high, normal and low priority targets. This gives us five priority classes in the following order:

- Global high

- Site-cost high

- Site-cost normal
- Site-cost low
- Global low

Note that there is no separate “global normal”. If global high or global low is not specified then the referral is placed in the "global normal" set. Priority rank is evaluated by a simple integer ranking – 0, 1, 2, and so forth.

In ordering a referral, the complete process now works as follows:

- The sets of global high and global low targets are identified, along with the remaining “global normal” targets.
- These three sets are placed in order – global high, normal and low.
- If set, INSITE is applied to the “normal” set and targets outside the client’s site are removed. It is not applied to global high and global low.
- Within each of these three sets, the targets are ordered by the site cost mechanism (either local site or by full site costing), producing sets of targets of equal site cost
- Within the sets of “global normal” targets of equal site cost, targets are ordered by priority class – site cost high, normal and low.
- Within the sets of targets of equal site cost and priority class, targets are now ordered by priority rank (0 being the highest).
- Within the sets of targets of equal site cost, priority class, and priority rank, targets are randomly shuffled for load balancing.

The **/targetpriority** option of **dfsutil** can be used to set the priority options listed above.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Better Delegation**
**Detailed description**
As of Windows Server 2003 Service Pack 1, administrators may delegate the ability to manage domain DFS roots by setting rights to the DFS object in Active Directory. Administrators may also delegate the ability to manage stand alone DFS roots by setting rights to the DFS object in the registry. This is in addition to the existing delegation model of adding DFS administrators to the local Administrators group on every server that hosts the namespace.

**Why is this change important?**
In the previous versions of Windows Server 2003, administrators who want to delegate the ability to manage a namespace must make the user a member of the local Administrators group on every server that hosts the namespace. This requirement can lead to an inconsistent management experience if the group memberships are not assigned correctly, and it also allows the user full access to the local file system of each namespace server.

**What works differently?**
The following procedures have been updated as a result of this change in functionality:

**To grant a user permission to administer only a domain DFS namespace:**

1. In the **Active Directory Users and Computers** MMC snap-in, on the **View** menu, click **Advanced Features**.
2. In the console tree, double-click the **System** folder to expand it.
3. Click the **DFS-Configuration** folder.

   Any existing root objects appear in the details pane.
4. Right-click the root object that you want to allow the user to administer, and then click **Properties**.

5. On the **Security** tab, click **Add**.
6. Type the name of the user, and then click **OK**.
7. Verify that the user is granted the **Full Control** permission, and then click **OK**.

| **Note** |
| --- |
| The user's rights as defined by the access control list applied to the DFS object in Active Directory are evaluated at the individual root targets. Because of this, rights granted to well known users and groups are ignored due to potential differences between the membership of the well known groups in Active Directory and at each root target. For example, granting Power Users full control over the DFS root object in Active Directory will not enable Power Users on the individual root targets to manage the corresponding DFS root. Additionally, if a well known user or group is granted rights to a root target but is denied access to the DFS root object in Active Directory, that denial will be ignored by the DFS root targets. |

**To grant a user permission to administer a standalone DFS namespace**

1. Using regedit, locate the following key:

   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\DFS\Roots\Standalone**

2. Right-click the root object you want to allow the user to administer, and then click **Permissions...**
3. On the **Security** tab, click **Add**.
4. Type the name of the user, and then click **OK**.
5. Verify that the user is granted the **Full Control** permission, and then click **OK**.

**Referral Time to Live Values Are Not Renewed for DFS Clients**
**Detailed description:**
For DFS clients that are not running Windows XP with SP2 or Windows Server 2003 with SP1, the Time to Live for a referral determines the earliest time that a client will request a new referral, but only if the existing referral expires before it is accessed again. Clients that use a cached referral will renew the Time to Live of the referral and thus use the referral indefinitely until the client's referral cache is cleared or the client is restarted.

This behavior has changed for clients running Windows XP with SP2 and Windows Server 2003 with SP1. Specifically, the Time to Live value is not renewed each time a client accesses a target using a cached referral. Instead, the referral expires after the Time to Live value lapses.

**Why is this change important?**
This change allows DFS clients running Windows XP with SP2 or Windows Server 2003 with SP1 to honor the actual Time to Live value of a referral. One benefit of this change is that DFS clients running these service packs will more quickly react to changes to links and roots. For example, if the link targets of a link named Current are changed daily, DFS clients without these service packs would refresh the Time to Live value each time they access the Current link, causing them to continue to reference stale link targets well beyond the Time to Live value associated with the initial referral request.

**What works differently?**
This change has several effects:

- Clients running Windows XP with SP2 or Windows Server 2003 with SP1 will request referrals more frequently than other DFS clients, which can cause moderately increased load on the domain-based DFS root servers, stand-alone root servers, and domain controllers.
- Because they request new referrals more frequently, clients running Windows XP with SP2 or Windows Server 2003 with SP1 will discover namespace updates more quickly than other DFS clients.
- In terms of DFS failover, this new behavior can cause the DFS client to fail over to a new target if the previously active target is not in the new referral, such as when the namespace is changed to remove the target that the client had accessed.

**Service Disabled By Default**
**Detailed description**
The DFS service is disabled by default in Windows Server 2003 Service Pack 1.

**Why is this change important?**
Unnecessary services are disabled by default in Windows Server 2003 SP1 to reduce the attack surface of the server.

**What works differently?**
Computers that are new installations of Windows Server 2003 with Service Pack 1 will have the DFS service disabled by default.

Upgrading a computer that is running Windows Server 2003 by installing Service Pack 1 will result in the DFS service being disabled unless the server is either a domain controller or is already hosting an existing DFS root.

When a server running Windows Server 2003 Service Pack 1 is promoted to domain controller status, the DFS service is enabled.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| Server Consolidation Retry | HKEY_LOCAL_MACHINE \SYSTEM \CurrentControlSet \Services\Dfs \Parameters\Replicated | N/A | 0 | 0,1 |
| LogServer Consolidation | HKEY_LOCAL_MACHINE \SYSTEM \CurrentControlSet \Services\Dfs \Parameters\Replicated | N/A | 0 | 0,1 |
| Sysvol Netlogon Target Failback | HKEY_LOCAL_MACHINE \SYSTEM \CurrentControlSet \Services\Dfs \Parameters | N/A | 0 | 0,1 |

# Distributed Transaction Coordinator

Updated: January 12, 2011

Applies To: Windows Server 2003 with SP1

**What does Distributed Transaction Coordinator do?**
The Distributed Transaction Coordinator (DTC) service coordinates transactions that update two or more transaction-protected resources, such as databases, message queues, files systems, and so on. These transaction-protected resources may be on a single computer or distributed across many networked computers.

**Who does this feature apply to?**

- Users of any computers that participate in DTC transactions, either directly or through other computers.
- System administrators of networks that use DTC components to perform transactions across networks.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Securing all network communication by default**
**Detailed description**
In Windows Server 2003 Service Pack 1, DTC provides the administrator with greater control over the network communication between computers. By default, all network communication is disabled.

In order to manipulate the communication settings, the DTC security settings properties page has been enhanced. To see the page, use the following procedure:

**To open the DTC security settings properties page**

1. Open the **Component Services** snap-in Microsoft Management Console (MMC).
2. In the console tree, click the **Computers** folder.
3. In the results pane, right click **My Computer** and then click **Properties**.
4. Click the **MSDTC** tab, and then click **Security Configuration**.

The table below defines the new fields in the property page, along with the registry keys affected for the different settings. All the registry keys related to MSDTC are located in the following registry key:

`MyComputer\HKEY_LOCAL_MACHINE\Software\Microsoft\MSDTC`

| **Caution** |
|---|
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. These registry keys might not be supported in future releases. |

The following table tells you where to find the MSDTC key specific values.

| Setting | Description | Corresponding registry value |
|---|---|---|
| Network DTC Access | Determines whether DTC on the local computer is allowed to access the network. This setting must be enabled in combination with one of the other settings to enable network DTC transactions.<br><br>Default setting: Off | Security\NetworkDtcAccess<br><br>0 = Off<br><br>1 = On |

| | | |
|---|---|---|
| Allow inbound | Allows a distributed transaction that originates from a remote computer to run on this computer.<br><br>Default setting: Off | To enable this setting you must set the following registry key values to 1:<br><br>`Security\NetworkDtcAccess`<br><br>`Security\NetworkDtcAccessTransactions`<br><br>`Security\NetworkDtcAccessInbound`<br><br>To disable this setting, you only need to set the following registry key value to 0:<br><br>`Security\NetworkDtcAccessInbound` |
| Allow Outbound | Allows the local computer to initiate a transaction and run it on a remote computer. | To enable this setting, you need to set the following registry key values to 1:<br><br>`Security\NetworkDtcAccess`<br><br>`Security\NetworkDtcAccessTransactions`<br><br>`Security\NetworkDtcAccessOutbound`<br><br>To disable this setting, you only need to set the following registry key value to 0:<br><br>`Security\NetworkDtcAccessOutbound` |
| Mutual Authentication Required | Adds support for mutual authentication in future versions and is the highest secured communication mode. In the current versions of Windows and Windows Server, it is functionally equivalent to the Incoming Caller Authentication Required setting. This is the recommended transaction mode for clients running Windows XP SP2 and servers running a member of the Windows Server 2003 family.<br><br>**Warning**<br>You cannot use the **Mutual Authentication Required** transaction mode with computers that are in a clustered environment, or any computers that are negotiating transactions with such computers. In that context, you can use the **Incoming Caller Authentication Required** transaction mode instead. In a clustered environment, the computer account for the Distributed Transaction Coordinator service specifies the cluster node's host name instead of the transaction node's host name, which prevents the authentication request from succeeding when the **Mutual Authentication Required** transaction mode is enabled. | AllowOnlySecureRpcCalls = 1<br><br>FallbackToUnsecureRPCIfNecessary = 0<br><br>TurnOffRpcSecurity = 0 |
| Incoming Caller Authentication Required | Requires the local DTC to communicate with a remote DTC using only encrypted messages and mutual authentication. This setting is recommended for servers running Windows Server 2003 that are operating in a cluster.<br><br>Only Windows Server 2003 and Windows XP SP2 support this feature, so you should only use this if you know that the DTC on the remote computer runs either the Windows Server 2003 or Windows XP SP2 operating system. | AllowOnlySecureRpcCalls = 0<br><br>FallbackToUnsecureRPCIfNecessary = 1<br><br>TurnOffRpcSecurity = 0 |

| | | |
|---|---|---|
| No Authentication Required | Provides system compatibility between previous versions of the Windows operating system. When enabled, communication on the network between DTCs can fall back to a non-authentication or non-encrypted communication if a secure communication channel cannot be established. This setting should be used if the DTC on the remote computer runs a Windows 2000 operating system or a Windows XP operating system earlier than SP2. This setting is also useful when the DTCs that are involved are located on computers that are in domains that do not have an established trust relationship or if the computers are part of a Windows workgroup. | AllowOnlySecureRpcCalls = 0<br><br>FallbackToUnsecureRPCIfNecessary = 0<br><br>TurnOffRpcSecurity = 1 |

**Why is this change important? What threats does it help mitigate?**
These changes are important in order to secure any communication coming into or going out from the computer. By default, after installing Windows Server 2003 Service Pack 1, the computer will not accept or issue any network traffic and therefore will be less vulnerable to network attacks.

Additionally, the online network protocol has been upgraded to support a more securely encrypted and mutually authenticated communication mode. This helps to ensure that attackers can not intercept or take over communications between DTCs.

**What works differently?**
After installing Windows Server 2003 Service Pack 1, all network communication coming out of or getting into DTC is disabled. For example, if a COM+ object attempts to update a SQL database on a remote computer using a DTC transaction, the transaction fails. Conversely, if your computer is hosting a SQL database that components from remote computers try to access using a DTC transaction, their transactions fail.

**How do I fix these issues?**
If your transactions fail because of network connectivity, you can use MSDTC security properties, as described previously in this document, select the **Network DTC Access** check box, and then select the **Allow Inbound** and **Allow Outbound** check boxes, as appropriate.

If you want to change these setting programmatically as part of your Windows Server 2003 Service Pack 1 deployment, you can directly change the registry values that correspond to your desired setting as described in the table in "Securing all network communication by default," earlier in this document. After you have changed the registry settings, you must restart the MSDTC service.

If you are using Windows Firewall to protect the computers in your organization, you must add MSDTC into the exception list in the Windows Firewall settings. To do so, use the following steps:

1. In **Control Panel**, open **Windows Firewall**.
2. Click the **Exceptions** tab, and then click **Add Program**.
3. Click **Browse**, and then add **c:\windows\system32\msdtc.exe**.
4. In **Programs and Services**, select the **Msdtc.exe** check box, and then click **OK**.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| NetworkDtcAccess | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft \MSDTC \Security | 1 | 0 | 0,1 |
| NetwordDtcAccessTransactions | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft \MSDTC \Security | 1 | 0 | 0,1 |
| NetworkDtcAccessInbound | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft \MSDTC \Security | n/a | 0 | 0,1 |
| NetworkDtcAccessOutbound | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft \MSDTC \Security | n/a | 0 | 0,1 |

| | | | | |
|---|---|---|---|---|
| AllowOnlySecureRpcCalls | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft \MSDTC | n/a | 1 | 0,1 |
| FallbackToUnsecureRPCIfNecessary | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft \MSDTC | n/a | 0 | 0,1 |
| TurnOffRpcSecurity | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft \MSDTC | n/a | 0 | 0,1 |

# Download, Attachment, and Authenticode Enhancements

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Notification Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What do the download, attachment and Authenticode enhancements do?**
In Windows Server 2003 with Service Pack 1, the prompts that are used for file downloads, mail attachments, shell process execution, and program installation have been modified to be more consistent and clearer than they were in previous versions of Windows Server. In addition, the publisher information will be shown before a file type that is signable and can potentially harm the user's machine is opened. (Common examples of signable file types that can potentially harm the user's machine are .exe, .dll, .ocx, .msi, and .cab.)

There is a new application programming interface (API), which allows application developers to make use of this new user interface. For more information regarding the API, see "AES API Integration," in the section of this document on changes to e-mail features in Windows Server 2003 Service Pack 1.

**Who does this feature apply to?**
Application developers will be able to call the new Attachment Manager dialog box from their Windows applications by using the API that is described in the "Attachment Manager API Integration" topic in the "Outlook Express" section of this document.

Application developers should also be aware that, in certain scenarios, such as attempting to open an attachment or downloading a file that is potentially dangerous, file types that can potentially harm a user's computer will have their digital signatures checked before they are opened. The signature information is presented to the user to help inform the user of the file's publisher.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Internet Explorer file download prompt**
**Detailed description**
When a user uses Internet Explorer to download a file, the dialog box that appears has the following changes:

- A file handler icon has been added.
- A new information area has been added to the bottom of the dialog box that provides slightly different information, depending on whether the downloaded file type is of higher or lower risk.

All file types that are signable and that can potentially harm a user's computer are checked for publisher information. This information will be shown to the user before opening the file.

The publisher information is shown before opening a file type that is signable and that can potentially harm the user's computer. The **Authenticode** dialog box presents this information to the user, who can then make a more informed decision about running the file.

**Why is this change important?**
This change helps bring consistency and clarity to the experience of downloading files and code to a user's computer. The publisher check provides crucial information when a signature is found in a file and provides a systematic way to prevent files that are from suspicious publishers from compromising the security of a computer.

**What works differently?**
Files with blocked publishers are not allowed to run.

**How do I resolve these issues?**
You can unblock a publisher of an add-on by using **Manage Add-ons** in Internet Explorer. To unblock a publisher to enable

the download of a specific file, you can remove the publisher from the **Untrusted Publishers** list. To do this, in Internet Explorer, on the **Tools** menu, click **Internet Options**, click the **Content** tab, click the **Publishers** button and then remove the publisher's name from the **Untrusted Publishers** list.

**Outlook Express e-mail attachment prompt**
**Detailed description**
The Outlook Express e-mail attachment prompt uses the same procedures as file downloads and leverages the AES API Integration. As a result e-mail attachments in Outlook Express show the publisher information for files types that can potentially harm a user's computer and any file whose publisher has been blocked will not be allowed to run.

**Why is this change important?**
This change helps bring consistency and clarity to the experience of downloading files and code to a user's computer.

**Add-on install prompt**
**Detailed description**
The Internet Explorer add-on install prompt has been simplified and only displays the file name and publisher information from the digital signature. It provides a warning about the risk associated with installing the add-on in order to help the user make a good decision about installing the add-on. Also, additional functionality was added to the prompt so that users always block a publisher, indicating that Windows should never trust anything from the publisher. This blocks the publisher from running code on the computer.

**Why is this change important?**
This change helps bring consistency and clarity to the experience of downloading files and code to a user's computer. In addition, the user can choose not to trust a publisher when the user is prompted to install the add-on. This gives users more control over their experience.

**What works differently? Are there any dependencies?**
When you install an add-on, the user interface is more clear and concise.

**How do I resolve these issues?**
By default, Internet Explorer will not allow users to run invalid or unsigned ActiveX controls. The Information Bar will provide an alternative way for the user to choose to install a blocked control. For more information see Internet Explorer Information Bar.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
Users now have the ability to block a publisher from running code on their computer.

# Internet Explorer Feature Control Settings in Group Policy

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Notification Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does Internet Explorer Feature Control Settings in Group Policy do?**
Windows XP Service Pack 2 introduced new registry keys and values for Internet Explorer security features called Feature Controls. These security features have been incorporated in Windows Server 2003 Service Pack 1. The specific behavior of the feature control registry settings is discussed with each security feature throughout this section.

A modified Inetres.adm file contains the feature control settings as policies. Administrators can manage the feature control policies by using Group Policy objects (GPOs). When Internet Explorer is installed, the default preferences settings for these feature controls are registered on the computer in `HKEY_LOCAL_MACHINE`. In Group Policy, the Administrator can set them in either `HKEY_LOCAL_MACHINE (Computer Configuration)` or `HKEY_CURRENT_USER (User Configuration)`.

**Who does this feature apply to?**
Group Policy administrators can uniformly configure the Internet Explorer Feature Control settings for the computers and users that they manage.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Group Policy Internet Explorer settings**
**Detailed description**
The new feature control policies are:

- Binary Behavior Security Restriction
- MK Protocol Security Restriction
- Local Machine Zone Lockdown Security
- Consistent MIME Handling
- MIME Sniffing Safety Feature
- Object Caching Protection
- Scripted Window Security Restrictions
- Protection from Zone Elevation
- Information Bar
- Restrict ActiveX Install
- Restrict File Download
- Add-on Management
- Network Protocol Lockdown

In the Group Policy Management Console, the local computer policies for Feature Controls are in **\Computer Configuration\Administrative Templates\Windows Components\Internet Explorer\Security Features**.

The current user policies for Feature Controls are in **\User Configuration\Administrative Templates\Windows Components\Internet Explorer\Security Features**.

The policy for the feature needs to be enabled for the process — for example, IExplore.exe — before the zones' individual security setting policies or preferences will be applied. For more about the behavior of Feature Controls keys and setting them for a process, see the section on each feature and the section Internet Explorer Using Feature Control Registry Settings with Security Zone Settings. For more information about specific security settings by zone, see Internet Explorer URL Action and Advanced Security Settings in Group Policy.

Administrators of Group Policy can manage these new policies in the Administrative Templates extension to the Group Policy Management Console. When configuring these policies, the administrator can enable or disable the security feature for explorer processes (Internet Explorer and Windows Explorer), for executable processes that they defined, or for all processes that host the WebOC.

Users cannot see the feature control policies or preference settings through the Internet Explorer user interface, except for Local Machine Zone Lockdown Security. Feature control policies can only be set using the Group Policy Management Console, and feature control preference settings can only be changed programmatically or by editing the registry.

**Configuring policies and preferences**

Group Policy is the recommended tool for managing Internet Explorer for client computers on a corporate network. Internet Explorer supports Group Policy management for the IE feature controls included in Windows XP Service Pack 2 and Windows Server 2003 Service Pack 1 as well as for Security page settings or URL actions. Administrators of Group Policy can manage these policy settings in the Administrative Templates extension of the Group Policy Management Console.

When implementing policy settings, it is recommended to configure template policy settings in one Group Policy object (GPO) and configure any related individual policy settings in a separate GPO. You can then use Group Policy management features (for example, precedence, inheritance, or enforce) to apply individual settings to specific client computers.

Policies can be read by users but can only be changed by Group Policy management or by an administrator. Preference settings can be changed programmatically, by editing the registry, or in the case of URL actions and Local Machine Zone Lockdown Security, by using Internet Explorer. Note that settings that are associated with policies will take precedence over settings specified using Internet Explorer preferences.

**IEAK/IEM**

For operating systems prior to Windows XP SP2 and Windows Server 2003 SP1 and previous Internet Explorer versions, Internet Explorer Administration Kit (IEAK) 6 Service Pack 1 remains the recommended tool for solution providers and application developers to customize Internet Explorer for their end users. IEAK support and the IEAK/IEM process does not change for Internet Explorer versions prior to Windows XP Service Pack 2. The process also has not changed for using IEAK/IEM to set user setting preferences in Internet Explorer versions prior to and including Windows Server 2003 SP1. This includes the new Internet Explorer 6 in Windows XP Service Pack 2 and Windows Server 2003 SP1 preference settings. However, the true policy settings incorporated by this feature can only be managed within Group Policy. For more information about IEAK, see "Microsoft Internet Explorer 6 Administration Kit Service Pack 1" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=26002.

In summary, the IEAK can still be used as before for all Internet Explorer versions prior to Windows XP Service Pack 2, and is still the tool to use for branding in Windows XP Service Pack 2 and Windows Server 2003 SP1. IEM/IEAK can be still be used to set user preference settings, but true policies should be set using the Group Policy Management Console.

**Frequently asked questions for existing users of the IEAK**

| Question | Answer |
|---|---|
| I currently use the IEAK with a corporate license to configure Internet Explorer on desktop computers and I don't have Active Directory in my organization. How can I configure Internet Explorer if the IEAK doesn't work with Windows XP SP2 or Windows Server 2003 SP1? | You can still use Group Policy to configure settings even if you don't use Active Directory. You can use Group Policy to create a local Group Policy object (GPO) with your settings and then deploy that GPO. After you have configured your GPO for Internet Explorer, you can deploy it using your standard deployment methods. For example, you might use startup or logon scripts, Systems Management Server scripts, or you might send links to users in e-mail. |

| | |
|---|---|
| I currently use the IEAK with a corporate license to configure Internet Explorer on desktops and I don't have Active Directory in my organization. What happens if I keep running my IEAK 6 SP1 packages against Windows XP SP2 or Windows Server 2003 SP1? | If you install an IEAK 6 SP1 package on a computer running either Windows XP SP2 or Windows Server 2003 SP1, the settings from Internet Explorer 6 SP1 will be updated, but the security settings will not be configurable, since IEAK 6 SP1 wasn't designed to deploy those settings. |
| I currently use the IEAK with an Internet service provider (ISP) license to brand Internet Explorer bits and setup connections for my ISP customers. Will my IEAK 6 SP1 packages still be able to apply settings on Windows XP SP2 and Windows Server 2003 SP1? | The branding settings in your ISP license IEAK 6 SP1 package should be applied correctly. |

**Why is this change important?**

By adding the Internet Explorer Feature Controls policies to Group Policy, administrators can manage these true policies to establish standard security settings for all the computers that they configure.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**

Internet Explorer in Windows Server 2003 SP1 adds policies to Group Policy but does not change how policies are managed. Developers need to be aware of how each feature control setting affects security-related behavior for their applications. The effects of the different security-related behaviors on application development are discussed within this document in the specific sections for each feature.

# Internet Explorer Using Feature Control Registry Settings with Security Zone Settings

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Notification Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What do Feature Control Registry Settings and Security Zone Settings do?**
Feature Control registry settings are provided for Internet Explorer so that a specific process can be configured to opt-in to a particular security feature. Each security feature has a corresponding registry key that you can use to opt-in or opt-out of the security feature.

When a process has been configured to use a security feature, the security feature is running. Once the feature is running, there might be corresponding security zone settings that can be applied for more precision. Some security features do not have additional security zone settings.

In the **Security Settings** tab of **Internet Options**, the user can adjust these settings for many of the new feature controls. If you select **Enable**, it lowers the security settings and allows the behavior to run less securely, or in the same manner as it did in previous version of Internet Explorer. The feature control can be applied again by setting the security zone setting to **Disable**, which blocks the less-secure behavior while the feature control is enabled for that process.

Each of the Feature Controls is discussed in more detail in this document. For more information about URL action settings and how they relate to security zones, see "About URL Security Zones Templates" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=26001.

Using security zone settings for a feature provides additional precision in control for security features in Internet Explorer and can help manage application compatibility for organizational intranet applications. A user or administrator can select different behaviors based on risk.

**Who does this feature apply to?**
Web application developers need to be aware that the Internet Explorer security settings are dependent on the zone in which an application is run. Therefore, you should assign security zones carefully; this should be a part of your information security considerations. The security zones that you use should also be considered when assessing application compatibility.

Administrators of Group Policy may want to adjust the default values for each zone to suit the particular environments in their organization.

Unless prevented by policies in Group Policy, users can manage the values for these security zone settings (or URL actions) for each zone through **Internet Options** in Control Panel. Note that the Local Machine zone is not available through Control Panel. To access the security settings for a zone, click **Start**, click **Control Panel**, click **Internet Options**, click the **Security** tab, click a Web security zone, and then click **Custom Level**.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Feature control registry settings**
**Detailed description**
Windows Server 2003 Service Pack 1 introduces new feature control registry settings.

For many of these features, when the registry setting is on, users can configure the security settings (also known as URL action flags) to fine tune the feature control in each individual security zone

If you choose **Enable** as the action to take for an Internet Explorer feature control, the zone is secured as it was for the previous version of Internet Explorer. Relevant security control features will not apply in this zone; the security zone will run without the added layer of security provided by this feature.

If you choose to disable the security zone setting, the actions that may be harmful cannot run; this Internet Explorer security feature will be turned on in this zone, as dictated by the feature control setting for the process.

Security settings are often applied to a zone by a URL security zone template. The default values for the security settings and the settings by zone template are listed in the section Internet Explorer URL Action and Advanced Security Settings in Group Policy.

**Why is this change important? What threats does it help mitigate?**
As originally envisioned, each feature control setting would either be on or off for all security zones. Customer feedback indicated that more precise tuning with the settings was necessary for some features. For example, the internal workflow of some organizations depends on intranet applications. A feature control that protects users in the Internet zone may cause an intranet application to stop working. Because of this, Microsoft has incorporated the ability to control many security settings by zone.

**What works differently?**
Adding security settings by zone provides more flexibility in applying the new security features. This flexibility will provide a more manageable implementation of this new security feature, particularly in intranet scenarios.

**How do I resolve these issues?**
If the feature control setting is suspected of causing problems for an application, changing the feature control setting in the zone where the application is running to **Enable** allows the administrator or user to return to the previous behavior in that zone for that specific feature while maintaining the more secure behavior in other security zones. For some security settings, additional configuration options such as **Prompt** and **Admin-approved** are available, as well as **Enable** and **Disable**.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
If the code uses the default URLmon security manager, the developer must call `CoInternetIsFeatureEnabledForURL` to check the security settings for a particular zone.

# Internet Explorer URL Action and Advanced Security Settings in Group Policy

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Notification Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the Enhanced Security Configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does Internet Explorer Settings in Group Policy do?**
Windows XP Service Pack 2 introduced true policies for the configurable actions in the Internet Explorer **Security** tab settings. In addition to incorporating these policies into Internet Explorer in Windows Server 2003 Service Pack 1, additional policies were created for selected configurable actions in the Internet Explorer **Advanced** tab, as well as for URL action policies in Locked-Down zones used only by the Network Protocol Lockdown security feature. In this release, these security settings are managed using the Group Policy Management Console and, if set, can only be changed by a Group Policy object (GPO) or by an administrator.

An updated Inetres.adm file contains a list of settings as policies, including Advanced settings, which are also found in the Internet Explorer user interface as preferences. Administrators can manage the new feature control policies by using Group Policy objects (GPOs). When Internet Explorer is installed, the default HKEY_CURRENT_USER preferences settings for these settings are registered on the computer as they were in previous versions. The Administrator has to use the Group Policy Management Console (GPMC) to add these settings as policies.

**Who does this feature apply to?**
Group Policy administrators can uniformly configure the new Internet Explorer Advanced setting policies, as well as policies for Locked-Down security zones, for the computers and users that they manage. It is important to inform the end-user which actions are controlled by policy, as these actions will override user preference settings.

| Note |
| --- |
| The Internet Options control panel will display policy settings when opened and users can interact with user interface and appear to change their preferences. However, these preferences will not actually override Group Policy settings, which may cause a confusing user experience. The administrator can also set a policy to disable the **Advanced** page user interface so that it is clearer to the user that these settings are not available to be changed. This is not an issue for the Locked-Down zones' settings as they are not accessible through the user interface. |

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Group Policy Internet Explorer advanced settings**
**Detailed description**
The following definitions apply to Internet Explorer settings for Windows Server 2003 with Service Pack 1:

- Security zones: Locked-Down Intranet Zone, Locked-Down Trusted Sites Zone, Locked-Down Internet Zone, and Locked-Down Restricted Sites zone.

- Templates: Standard settings for all URL actions in these security zones. Templates can be applied in any zone, and settings will provide a range of choices from low security, medium-low, medium, and up to high security for the zone.

- URL actions: Security settings in the registry that identify the action to take for that feature in the security zone where the URL resides. URL action settings include enable, disable, prompt, and others as appropriate.

- URL action policies: URL action policies can be added individually by enabling the desired URL action policy, then selecting the setting for the policy registry key value. They can also be set by zone template.

Internet Explorer will look for a policy in the following order:

- HKEY_LOCAL_MACHINE policy hive
- HKEY_CURRENT_USER policy hive
- HKEY_CURRENT_USER preference hive
- HKEY_LOCAL_MACHINE preference hive

If Internet Explorer finds a policy in the HKEY_LOCAL_MACHINE policy hive, it stops and does not continue; that is the setting it respects. If Internet Explorer does not find a policy in HKEY_LOCAL_MACHINE policy hive, it looks in the HKEY_CURRENT_USER policy hive, and so on. The administrator can set a policy for one or more URL actions in one or more zones, and allow the end user to manage preferences for URL actions that do not require policy-level security management.

**Policy values for URL action**
The new URL action policies have the same numeric values as their related preference keys. The following table provides a reference to these URL actions.

| URL action flag name | Security setting UI | Numeric name |
|---|---|---|
| URLACTION_DOWNLOAD_SIGNED_ACTIVEX | Download signed ActiveX controls | 1001 |
| URLACTION_DOWNLOAD_UNSIGNED_ACTIVEX | Download unsigned ActiveX controls | 1004 |
| URLACTION_ACTIVEX_RUN | Run ActiveX controls and plugins | 1200 |
| URLACTION_ACTIVEX_OVERRIDE_OBJECT_SAFETY | Initialize and script ActiveX controls not marked as safe | 1201 |
| URLACTION_SCRIPT_RUN | Active scripting | 1400 |
| URLACTION_SCRIPT_JAVA_USE | Scripting of Java applets | 1402 |
| URLACTION_SCRIPT_SAFE_ACTIVEX | Script ActiveX controls marked safe for scripting | 1405 |
| URLACTION_CROSS_DOMAIN_DATA | Access data sources across domains | 1406 |
| URLACTION_SCRIPT_PASTE | Allow paste operations via script | 1407 |
| URLACTION_HTML_SUBMIT_FORMS | Submit non-encrypted form data | 1601 |
| URLACTION_HTML_FONT_DOWNLOAD | Font download | 1604 |
| URLACTION_HTML_USERDATA_SAVE | Userdata persistence | 1606 |
| URLACTION_HTML_SUBFRAME_NAVIGATE | Navigate sub-frames across different domains | 1607 |
| URLACTION_HTML_META_REFRESH | Allow META REFRESH | 1608 |
| URLACTION_HTML_MIXED_CONTENT | Display mixed content | 1609 |
| URLACTION_SHELL_INSTALL_DTITEMS | Installation of desktop items | 1800 |
| URLACTION_SHELL_MOVE_OR_COPY | Drag and drop or copy and paste files | 1802 |
| URLACTION_SHELL_FILE_DOWNLOAD | File download | 1803 |
| URLACTION_SHELL_VERB | Launching applications and files in an IFRAME | 1804 |
| URLACTION_SHELL_POPUPMGR | Use Pop-up blocker | 1809 |

| | | |
|---|---|---|
| URLACTION_NETWORK_MIN | Logon | 1A00 |
| URLACTION_CLIENT_CERT_PROMPT | Don't prompt for client certificate selection when no certificates or only one certificate exists | 1A04 |
| URLACTION_JAVA_PERMISSIONS | Java permissions | 1C00 |
| URLACTION_CHANNEL_SOFTDIST_PERMISSIONS | Software channel permissions | 1E05 |
| URLACTION_BEHAVIOR_RUN | Script and Binary Behaviors | 2000 |
| URLACTION_MANAGED_SIGNED | Run .NET Framework-reliant components signed with Authenticode | 2001 |
| URLACTION_MANAGED_UNSIGNED | Run .NET Framework-reliant components not signed with Authenticode | 2004 |
| URLACTION_FEATURE_MIME_SNIFFING | Open files based on content, not file extension | 2100 |
| URLACTION_FEATURE_ZONE_ELEVATION | Web sites in less privileged Web content zones can navigate into this zone | 2101 |
| URLACTION_FEATURE_WINDOW_RESTRICTIONS | Allow script-initiated windows without size or position constraints | 2102 |
| URLACTION_AUTOMATIC_DOWNLOAD_UI | Automatic prompting for file downloads | 2200 |
| URLACTION_AUTOMATIC_ACTIVEX_UI | Automatic prompting for ActiveX controls | 2201 |
| URLACTION_ALLOW_RESTRICTEDPROTOCOLS | Allow active content over restricted protocols to access my computer | 2300 |

For more information about using URL action flags, see "URL Action Flags" on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=32776.

The following table provides a reference to the setting options available for each URL action.

| **Numeric Name** | **URL Action Policy Setting Options** |
|---|---|
| 1001 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1004 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1200 | "Administrator approved"=0x00010000<br>"Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1201 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |

| | |
|---|---|
| 1400 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1402 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1405 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1406 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1407 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1601 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1604 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1606 | "Enable"=0x00000000<br>"Disable"=0x00000003 |
| 1607 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1608 | "Enable"=0x00000000<br>"Disable"=0x00000003 |
| 1609 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1800 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |

| | |
|---|---|
| 1802 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1803 | "Enable"=0x00000000<br>"Disable"=0x00000003 |
| 1804 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 1809 | "Enable"=0x00000000<br>"Disable"=0x00000003 |
| 1A00 | "Anonymous logon"=0x00030000<br>"Automatic logon only in Intranet zone"=0x00020000<br>"Automatic logon with current user name and password"=0x00000000<br>"Prompt for user name and password"=0x00010000 |
| 1A04 | "Enable"=0x00000000<br>"Disable"=0x00000003 |
| 1C00 | "High safety"=0x00010000<br>"Medium safety"=0x00020000<br>"Low safety"=0x00030000<br>"Custom"=0x00800000<br>"Disable Java"=0x00000000 |
| 1E05 | "High Safety"=0x00010000<br>"Medium Safety"=0x00020000<br>"Low Safety"=0x00030000 |
| 2000 | "Enable"=0x00000000<br>"Administrator approved"=0x00010000<br>"Disable"=0x00000003 |
| 2001 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 2004 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |

| | |
|---|---|
| 2100 | "Enable"=0x00000000<br>"Disable"=0x00000003 |
| 2101 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |
| 2102 | "Enable"=0x00000000<br>"Disable"=0x00000003 |
| 2200 | "Enable"=0x00000000<br>"Disable"=0x00000003 |
| 2201 | "Enable"=0x00000000<br>"Disable"=0x00000003 |
| 2300 | "Enable"=0x00000000<br>"Disable"=0x00000003<br>"Prompt"=0x00000001 |

**Key for numeric translation of URL policy settings**

| Value | DWORD | Setting |
|---|---|---|
| 0 | 0x00000000 | Enable |
| 1 | 0x00000001 | Prompt |
| 3 | 0x00000003 | Disable |
| 65536 | 0x00010000 | High Safety |
| 131072 | 0x00020000 | Medium Safety |
| 196608 | 0x00030000 | Low Safety |

For descriptions for each of the URL policy settings, see "URL Action Flags" on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=32777.

**Default settings for each URL action in zones and templates**

Each URL action has a default that is set in each zone and set when a specified template is applied. The default settings for each zone are described in the following table.

**URL action default settings**

| URL action numeric name | Locked-Down Restricted zone | Locked-Down Internet zone | Locked-Down Intranet zone | Locked-Down Trusted zone |
|---|---|---|---|---|
| 1001 | 3 | 1 | 1 | 0 |
| 1004 | 3 | 3 | 3 | 3 |
| 1200 | 3 | 3 | 3 | 3 |
| 1201 | 3 | 3 | 3 | 3 |
| 1400 | 3 | 3 | 3 | 3 |
| 1402 | 3 | 0 | 0 | 0 |
| 1405 | 3 | 0 | 0 | 0 |

| | | | | |
|---|---|---|---|---|
| 1406 | 3 | 3 | 1 | 0 |
| 1407 | 3 | 0 | 0 | 0 |
| 1601 | 1 | 1 | 0 | 0 |
| 1604 | 1 | 0 | 0 | 0 |
| 1606 | 3 | 0 | 0 | 0 |
| 1607 | 3 | 0 | 0 | 0 |
| 1608 | 3 | 0 | 0 | 0 |
| 1609 | 1 | 1 | 1 | 1 |
| 1800 | 3 | 1 | 1 | 0 |
| 1802 | 1 | 0 | 0 | 0 |
| 1803 | 3 | 0 | 0 | 0 |
| 1804 | 3 | 1 | 1 | 0 |
| 1809 | 0 | 0 | 3 | 3 |
| 1A00 | 65536 | 131072 | 131072 | 0 |
| 1A04 | 3 | 3 | 3 | 3 |
| 1C00 | 0 | 0 | 0 | 0 |
| 1E05 | 65536 | 131072 | 131072 | 196608 |
| 2000 | 3 | 65536 | 65536 | 65536 |
| 2001 | 3 | 3 | 3 | 3 |
| 2004 | 3 | 3 | 3 | 3 |
| 2100 | 3 | 3 | 3 | 3 |
| 2101 | 3 | 3 | 3 | 3 |
| 2102 | 3 | 3 | 3 | 3 |
| 2200 | 3 | 3 | 3 | 3 |
| 2201 | 3 | 3 | 3 | 3 |
| 2300 | 3 | 1 | 1 | 1 |

**Group Policy Settings Paths**

These paths locate the available Advanced settings in the Group Policy Management Console:

- HKEY_LOCAL_MACHINE policies for Advanced settings:

```
\Computer Configuration\Administrative Templates\Windows Components\Internet Explorer\Internet Control
Panel\Advanced Page
```

- HKEY_CURRENT_USER policies for Advanced settings:

```
\User Configuration\Administrative Templates\Windows Components\Internet Explorer\Internet Control
Panel\Advanced Page
```

These paths locate the security zone settings in the Group Policy Management Console:

- HKEY_LOCAL_MACHINE policies by security zone for URL actions:

```
\Computer Configuration\Administrative Templates\Windows Components\Internet Explorer\Internet Control
Panel\Security Page
```

- HKEY_CURRENT_USER policies by security zone for URL actions:

```
\User Configuration\Administrative Templates\Windows Components\Internet Explorer\Internet Control
Panel\Security Page
```

These paths locate the Advanced settings in policy and in preference in the Windows registry (in either HKEY_LOCAL_MACHINE or HKEY_CURRENT_USER):

| **Advanced setting UI** | **Preference key name** | **Policy key name** |
|---|---|---|
| Install on Demand (Internet Explorer) | HKCU \Software \Microsoft\Internet Explorer\Main \NoJITSetup | Software\Policies\Microsoft\Internet Explorer\Main\NoJITSetup |
| Install on Demand (Other) | HKCU\Software\Microsoft\Internet Explorer\Main \NoWebJITSetup | Software\Policies\Microsoft\Internet Explorer\Main\NoWebJITSetup |
| Third-party Browser Extensions | HKCU\Software \Microsoft\Internet Explorer\Main\Enable Browser Extensions | Software\Policies \Microsoft\Internet Explorer\Main\Enable Browser Extensions |
| Automatically check for IE Updates | HKCU\Software \Microsoft\Internet Explorer\Main \NoUpdateCheck | Software\Policies \Microsoft\Internet Explorer\Main \NoUpdateCheck |
| Play Animations in Web Pages | HKCU\Software \Microsoft\Internet Explorer\Main \Play_Animations | Software\Policies \Microsoft\Internet Explorer\Main \Play_Animations |
| Play Sounds in Web Pages | HKCU\Software \Microsoft\Internet Explorer\Main \Play_Background_Sounds | Software\Policies \Microsoft\Internet Explorer\Main \Play_Background_Sounds |
| Play Videos in Web Pages | HKCU\Software \Microsoft\Internet Explorer\Main\Display Inline Videos | Software\Policies \Microsoft\Internet Explorer\Main\Display Inline Videos |
| Allow software to run or install even if the signature is invalid | HKCU\Software \Microsoft\Internet Explorer\Download \RunInvalidSignatures | Software\Policies \Microsoft\Internet Explorer\Download \RunInvalidSignatures |
| Allow active content from CDs to run on user machines | HKCU\Software \Microsoft\Internet Explorer\Main \FeatureControl \FEATURE_LOCALMACHINE_LOCKDOWN \Settings \LocalMachine_CD_Unlock | \Software\Policies \Microsoft\Internet Explorer\Main \FeatureControl \FEATURE_LOCALMACHINE_LOCKDOWN \Settings \LocalMachine_CD_Unlock |
| Check for Server Certificate Revocation | HKCU\Software \Microsoft\Internet Explorer\Download \CertificateRevocation | Software\Policies \Microsoft\Windows \CurrentVersion \InternetSettings \CertificateRevocation |
| Check for Signatures on Downloaded Programs | HKCU\Software \Microsoft\Internet Explorer\Main\ CheckExeSignatures | Software\Policies \Microsoft\Internet Explorer\Main\ CheckExeSignatures |
| Do Not Save Encrypted Pages to Disk | HKCU\Software \Microsoft\Windows \CurrentVersion \InternetSettings \DisableCachingOfSSLPages | Software\Policies \Microsoft\Windows \CurrentVersion \InternetSettings \DisableCachingOfSSLPages |
| Empty Temporary Internet Files Folder When Browser is Closed | HKCU\Software \Microsoft\Internet Explorer\Cache \Persistent | Software\Policies \Microsoft\Windows \CurrentVersion \InternetSettings\Cache \Persistent |

These paths locate the security zone settings in policy and in preference in the Windows registry (in either

HKEY_LOCAL_MACHINE or HKEY_CURRENT_USER):

- Location of Locked-Down Intranet zone policy values:

```
Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\Lockdown_Zones\1
```

- Location of Locked-Down Trusted Sites policy:

```
Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\Lockdown_Zones\2
```

- Location of Locked-Down Internet zone policy values:

```
Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\Lockdown_Zones\3
```

- Location of Locked-Down Restricted Sites policy values:

```
Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\Lockdown_Zones\4
```

- Location of Locked-Down Intranet zone template:

```
Software\Policies\Microsoft\Windows\CurrentVersion\Intranet Lockdown Settings
```

- Location of Locked-Down Trusted Sites template:

```
Software\Policies\Microsoft\Windows\CurrentVersion\Trusted Sites Lockdown Settings
```

- Location of Locked-Down Internet zone template:

```
Software\Policies\Microsoft\Windows\CurrentVersion\Internet Lockdown Settings
```

- Location of Locked-Down Restricted Sites template:

```
Software\Policies\Microsoft\Windows\CurrentVersion\Restricted Sites Lockdown Settings
```

**Configuring policies and preferences**
Group Policy is the recommended tool for managing Internet Explorer for client computers on a corporate network. Internet Explorer supports Group Policy management for all new Internet Explorer Feature Controls in Windows Server 2003 Service Pack 1, and for Security page settings or URL actions. Administrators of Group Policy can manage these new policy settings in the Administrative Templates extension of the Group Policy Management Console.

When implementing policy settings, it is recommended that you configure template policy settings in one Group Policy object (GPO) and configure any related individual policy settings in a separate GPO. You can then use Group Policy management features (for example, precedence, inheritance, or enforce) to apply individual settings to specific client computers.

Policies can be read by users but can only be changed by via Group Policy management or by an administrator. Preference settings can be changed programmatically, by editing the registry, or in the case of URL actions, by using Internet Explorer. Settings specified by Group Policy take precedence over settings specified using preferences.

**Why is this change important?**
By adding the new Advanced setting policies and Locked-Down security policies to Group Policy, administrators can manage these true policies to establish standard settings for all the computers that they configure. The administrator can control these settings in such a way that they cannot be changed except through Group Policy or by a user with administrator privileges, thus ensuring that security and certain Advanced settings are not set by end users.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
Windows Server 2003 Service Pack 1 adds new policies to Group Policy but does not change how policies are managed. Developers need to be aware of how each Feature Control and URL action setting or setting combination affects security-related behavior for their applications in each security zone.

For greater security, the administrator should enable policies for all zones, so that there is a known configuration set by policy rather than an unknown setting read from HKEY_LOCAL_MACHINE or HKEY_CURRENT_USER preference settings not set by

policy. If the administrator sets policies for all zones, we recommend that the policy to disable the **Security** page be enabled, which will make the user interface in Internet Explorer unavailable.

**Feature Control Policies**
The administrator should also understand the Feature Control policy settings. Some of the URL action settings will not be valid unless the corresponding Feature Control policy is enabled. Internet Explorer checks to see whether the feature is enabled, and if it is, then looks for the setting for the action based on the security zone of the URL.

**Zone Map Policies**
The method for adding Zone Map keys to policy is as follows:

1. To set computer policy, go to \Computer Configuration\Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel\Security Page within Group Policy. To set user policy, go to \User Configuration\Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel\Security Page within Group Policy.

2. Select the **Site to Zone Assignment List** policy.

3. Select **Enabled** and click **Show...**

4. For each site you would like to map:

   a. Click **Add...**

   b. Enter the name, IP address, or IP range of the site you want to map (for example, http://www.contoso.com, www.contoso.com, 127.0.0.1, 127.0.0.1-10)

   c. Enter the value identifying the zone to which this site should be mapped. The choices are (1) Intranet zone, (2) Trusted Sites zone, (3) Internet zone, (4) Restricted Sites zone.

   d. Click **OK.**

   e. The site name and value should appear in the list.

5. Click **OK** in the **Show Contents** window.

6. Click **OK** again to close the **Site to Zone Assignment List Properties** window.

| **Note** |
| --- |
| Policies created by following these instructions are ignored by computers with the Windows Server 2003 Internet Explorer Enhanced Security Configuration component installed. To set zone map policy on a computer with Windows Server 2003 Internet Explorer Enhanced Security Configuration component installed, use the Internet Explorer Maintenance (IEM) snap-in to Group Policy. When using the IEM to create a Group Policy object to apply to a computer with the Windows Server 2003 Internet Explorer Enhanced Security Configuration component installed, you must be using a computer with the Windows Server 2003 Internet Explorer Enhanced Security Configuration component installed. |

| **Note** |
| --- |
| For more information about using Group Policy, see "Implementing Registry-based Group Policy" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=28188. For more information about using Internet Explorer security zone and privacy settings, see "Description of Internet Explorer Security Zones Registry Entries" on the Microsoft Knowledge Base Web site at http://go.microsoft.com/fwlink/?LinkId=28195. |

# Internet Explorer Add-on Management and Crash Detection

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Information Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does Internet Explorer Add-on Management and Crash Detection do?**
These are two new, closely-related features that are included in Internet Explorer.

Internet Explorer Add-on Management allows users to view and control the list of add-ons that can be loaded by Internet Explorer with more detailed control than before. It also shows the presence of some add-ons that were previously not shown and could be very difficult to detect.

Internet Explorer Add-on Crash Detection attempts to detect crashes in Internet Explorer that are related to an add-on. When the add-on is successfully identified, this information is presented to the user. The user has the option of disabling add-ons to diagnose crashes and improve the overall stability of Internet Explorer.

**Who does this feature apply to?**
Users will be able to view, enable, and disable the add-ons used by Internet Explorer, and identify add-ons that might be related to Internet Explorer crashes. Administrators can enforce a list of add-ons that are allowed or disallowed and restrict the ability of users to manage add-ons.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Internet Explorer Add-on Management**
**Detailed description**
Internet Explorer Add-on Management allows users to view and control the list of add-ons that can be loaded by Internet Explorer with more detailed control than before. It also shows the presence of some add-ons that were previously not shown and could be very difficult to detect. These add-ons might provide undesired functionality or services and, in some cases, might present a security risk.

For example, a user might unintentionally install an add-on that secretly records all Web page activity and reports it to a central server. Previously, specialized software and deep technical knowledge might have been required to identify and remove that add-on. Internet Explorer Add-on Management provides an easier way to detect and disable that add-on.

Add-ons include:

- Browser help objects
- ActiveX controls
- Toolbar extensions
- Browser extensions

Add-ons can be installed from a variety of locations and in several ways, including:

- Download and installation while viewing Web pages.
- Installation by the user by way of an executable program.
- As pre-installed components of the operating system.

- As pre-installed add-ons that come with the operating system.

#### Manage Add-ons

Users can enable and disable each add-on individually and view information about how often the add-ons have been used by Internet Explorer. To do this, use one of the following procedures to open Manage Add-ons.

**Open Manage Add-ons Using Internet Explorer**

1. Click **Start**, and then click **Internet Explorer**.
2. On the **Tools** menu, click **Manage Add-ons**.

**Open Manage Add-ons using the Control Panel**

1. Click **Start**, and then click **Control Panel**.
2. Double-click **Internet Options**.
3. Click the **Programs** tab, and then click **Manage Add-ons**.

**Manage Add-ons** has several options that allow you to change your add-on configuration.

You can use the **Show** drop-down list in **Manage Add-ons** to control the way in which the add-ons list is displayed. It has two options:

- **Add-ons currently loaded in Internet Explorer.** This option lists the add-ons that have been instantiated (or loaded into memory) within the current Internet Explorer process and those that have been blocked from instantiating. This includes ActiveX controls that were used by Web pages that were previously viewed within the current process.

- **Add-ons that have been used by Internet Explorer.** This option lists all add-ons that have been referenced by Internet Explorer and are still installed.

The list of add-ons shows all installed add-ons of the types listed previously in the detailed description section. To enable or disable an installed add-on, click the add-on in the list, then click **Enable** or **Disable**.

If you click an ActiveX control in the list, then click **Update ActiveX**, Windows searches for an update at the location where the original control was found. If a newer version is found at that location, Internet Explorer attempts to install the update.

The list of add-ons also contains signed add-ons that were blocked from installation because their publisher was untrusted. After selecting one of these controls, the user can unblock the control by clicking **Allow**. Caution should be exercised when doing this, because clicking **Allow** removes the publisher from the **Untrusted** list.

#### Blocked Add-on status bar icon

A **Blocked Add-on** icon appears in the status bar when a Web page attempts to instantiate an ActiveX control that is disabled or blocked because its publisher is untrusted. You can double click the icon to open **Manage Add-ons**. The status bar icon is accompanied by a balloon tip the first five times it appears.

#### Add-on notification balloon tip

When a Web page attempts to instantiate a disabled add-on and there is no current Blocked Add-on status bar icon, a message appears to tell the user that the current Web page is requesting an add-on that is disabled. The user can click the message for more details on blocking add-ons.

You can use the Internet Options Control Panel to suppress the message.

#### Why is this change important?

Windows Error Reporting data has shown that add-ons are a major cause of stability issues in Internet Explorer. These add-ons significantly affect the reliability of Internet Explorer. These add-ons can also pose a security risk, because they might contain malicious and unknown code.

Many users are unaware of the add-ons they have installed on their computer. Some add-ons are loaded whenever Internet Explorer is started, but cannot be detected unless the user searches the registry. When users experienced crashes, there was no easy way to diagnose whether the issue was related to an add-on. Even if they suspected that the problem stemmed from recently-installed software, it was difficult to isolate the cause and often impossible to resolve if the software did not provide an uninstall option.

Internet Explorer Add-on Management, together with Add-on Crash Detection, gives users the ability to improve the security and stability of their systems by identifying and disabling problematic add-ons. Administrators are also provided with a powerful administrative tool to control add-on use in their organization.

**What works differently?**
**Behavior when add-ons are disabled**
Disabling an add-on does not remove it from the computer. It only prevents Internet Explorer from instantiating the object and executing its code. There is no guarantee that the disabled add-on will never be loaded, since an add-on that is considered by Internet Explorer to be disabled can still be used by another component in the system. The behavior that is displayed by disabling different object types varies.

- If an ActiveX control is disabled, Web pages that rely on the control might not work as expected. They behave as if the user has uninstalled the control from the computer and declined to install it. Users are not prompted to upgrade controls that have been disabled.

- If a browser helper object is disabled, functionality that depends on the object is not available, and there is no visual indication that a component is disabled.

- If a browser extension is disabled, toolbar buttons and menu entry points are not shown for that extension. Internet Explorer behaves as if the extension was not installed.

- If a toolbar extension is disabled, the toolbar does not appear in Internet Explorer and, on the View menu, the Toolbars item is disabled. Internet Explorer behaves as if the toolbar was not installed.

The concept of a disabled add-on only applies to instances of Internet Explorer (Iexplore.exe) and Windows Explorer (Explorer.exe) by default. Currently, other programs based on Internet Explorer components, such as the `WebBrowser` control, do not respect the disabled state. However, you can use the `featurecontrol` key to extend this functionality to other applications.

Some software programs depend on a combination of multiple add-ons to work correctly, and disabling any one of them might cause problems. Caution should be exercised when deciding to disable one or more add-ons.

**Uninstallation**
If the user disables a non-ActiveX add-on and subsequently uninstalls and then re-installs it, the add-on might remain in a disabled state. This is because Internet Explorer is not notified of application installations and does not detect any application state changes. However, if Internet Explorer is started while the add-on is not installed, it detects a change and automatically clears the disabled state.

If the user disables an ActiveX control and then uninstalls it, the next time a Web page attempts to use the control, Internet Explorer detects that the control is no longer present and clears the disabled state. However, if the ActiveX control is reinstalled using an executable file (as opposed to a Web page download) before there are any attempts to instantiate the control, then it remains disabled. This is because Internet Explorer does not detect a state change.

**How do I resolve these issues?**
In the event that disabling an add-on causes a lack of functionality, it can be restored by enabling the add-on in **Manage Add-ons**. Internet Explorer must be restarted for new settings to take effect, with the exception of ActiveX controls, where reloading the affected page might be sufficient.

**Internet Explorer Add-on Management for Administrators**
**Detailed description**
**Disabling the Crash Detection feature**
To disable the Crash Detection feature of Add-on Management, see "What settings are added or changed in Windows Server 2003 Service Pack 1?" below. When Crash Detection is disabled, a crash in Internet Explorer exhibits previous behavior, which is usually to invoke Windows Error Reporting. All policies for Windows Error Reporting continue to apply.

**Disabling Add-on Management user interface**
To disable the Add-on Management user interface, see "What settings are added or changed in Windows Server 2003 Service Pack 1?" below. When the Add-on Management user interface is disabled, the **Enable** and **Disable** options are unavailable in **Manage Add-ons**.

**Deny all add-ons unless specifically allowed in the Add-on list**
This policy setting allows administrators to ensure that any Internet Explorer add-ons not listed in the Add-on List policy setting will be denied.

To set this policy, an administrator can modify the `RestrictToList` registry key in either of the following locations:

- `HKCU\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Ext\`
- `HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Ext\`

**Key reference**
Name: `RestrictToList`

Type: **DWORD**

Value:

- 1 (Anything not on the Add-on list is considered disabled.)
- 0 (Anything not on the Add-on list works as it would without policy.)

**Add-on List**
Administrators can control the use of specific add-ons through the add-on list policy. Administrators can choose to enable or disable an add-on as well as allow a specific add-on to be managed by the user.

To set this policy, an administrator can create a registry value based on the GUID of the add-on in either of the following keys and then set the desired value:

- `HKCU\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Ext\CLSID`
- `HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Ext\CLSID`

Each add-on is a value in this registry key with the following properties.

**Key reference**
Name: *GUID* of add on

Type: **REG_SZ**

Value:

- 0 - Add-on is disabled and cannot be managed by the end user.
- 1 - Add-on is allowed and cannot be managed by the end user.
- 2 - Add-on is allowed and can be managed by the end user.

The Add-on (CLSID) lists are empty by default.

**Behavior of Management user interface when policies are applied**
When an Add-on Management policy is in effect, and the user selects an add-on from the management list that is disabled by policy, **Enable** and **Disable** are unavailable.

**Why is this change important?**
This feature allows administrators to control the usage of the new features.

**What works differently?**
The new features for allowing and disallowing add-ons work in conjunction with existing policies for managing ActiveX controls. Add-on disabling is applied on top of existing checks and does not replace other security restrictions that might be in place. For example, if an ActiveX control is blocked by its ActiveX compatibility flags, it will always be blocked, regardless of the add-on management settings.

Using the "Deny all add-ons unless specifically allowed in the Add-on List" policy will disable script and other controls necessary for some Web pages to function properly. For a list of CLSIDs that might need to be enabled for certain Web sites to function correctly, see the article on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=45658

**How do I resolve these issues?**
If you are using the "Deny all add-ons unless specifically allowed in the Add-on list" policy some Web applications might break

due to disabled scripting and other disabled controls. For information about enabling scripting and other commonly used Web controls, see the article on the Microsoft Web site at http://go.microsoft.com/fwlink/?linkid=45658

In the event that these controls do not address the issue and adding these policies continues to remove functionality that is required for a Web application that you want to use, remove the policies that were applied and restart Internet Explorer.

**Internet Explorer Add-on Crash Detection**
**Detailed description**
Whenever Internet Explorer stops unexpectedly, Windows starts the Add-on Crash Detection program. Add-on Crash Detection is an error analysis program that examines the state of the Iexplore.exe (Internet Explorer) process. It collects the list of dynamic link libraries (DLLs) that are loaded, and the value of the instruction pointer register (EIP) at the time of the crash. Add-on Crash Detection then attempts to find the DLL whose memory range the EIP lies within. This DLL is often the cause of the crash. If a DLL is found, it is not a system DLL, and the DLL is the COM server for an Internet Explorer add-on, the Internet Explorer Add-on Crash Detection dialog box appears. This dialog box contains information that indicates which add-on caused the crash, the name of the company associated with the add-on, and the description of the DLL file that contains the add-on code. To display Manage Add-ons, which you can then use to disable the identified add-on, click **Advanced**. After you review the information and click **Continue**, the standard Windows Error Reporting window opens.

**Why is this change important? What threats does it help mitigate?**
For this information, see "Internet Explorer Add-on Management for Users," earlier in this subject.

**What works differently?**
Since this feature only runs when Internet Explorer stops operating, there should be no changes to normal operation.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
**Internet Explorer Add-on Management and Crash Detection Settings**

| Setting name | Location | Default value | Possible values |
|---|---|---|---|
| Disable Crash Detection | HKCU {or HKLM} \Software\Policies \Microsoft\Internet Explorer \Restrictions<br><br>Name: NoCrashDetection<br><br>Type: DWORD | 0 | 0 — Off,<br><br>1 — On |
| Deny all add-ons unless specifically allowed in the Add-on List | HKCU {or HKLM} \Microsoft\Windows\CurrentVersion \Policies\Ext\<br><br>Name: RestrictToList<br><br>Type: DWORD | 0 | 0 — Off,<br><br>1 — On |
| Add-on List | HKCU {or HKLM} \SOFTWARE\Microsoft\Windows \CurrentVersion\Policies\Ext\CLSID<br><br>Name: GUID of the control<br><br>Type: REG_SZ | Not available | 0 – Add-on is disabled and cannot be managed by the end user.<br><br>1 – Add-on is allowed and cannot be managed by the end user.<br><br>2 – Add-on is allowed and CAN be managed by the end user. |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
Your code does not need to change to work with Internet Explorer Add-on Crash Detection or Add-on Management.

# Internet Explorer Binary Behaviors Security Setting

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Notification Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does binary behaviors security setting do?**
Internet Explorer contains dynamic binary behaviors: components that encapsulate specific functionality for HTML elements to which they were attached. These binary behaviors are not controlled by any Internet Explorer security setting, allowing them to work on Web pages in the Restricted Sites zone. In Windows Server 2003 Service Pack 1, there is a new Internet Explorer security setting for binary behaviors. This new setting disables binary behaviors in the Restricted Sites zone by default. In combination with the Local Machine Lockdown security feature, it also requires administrative approval for binary behaviors to run in the Local Machine zone by default. This new binary behaviors security setting provides a general mitigation to vulnerabilities in Internet Explorer binary behaviors.

For more information about binary behaviors, such as how they work and how to implement them, see "Cutting Edge: Binary Behaviors in Internet Explorer 5.5" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=21862. Note that binary behaviors, which are defined in C++ and compiled, are different from attached behaviors and element behaviors, which are defined in script.

**Who does this feature apply to?**
Application developers whose applications use Internet Explorer functionality in the restricted sites or local machine zones should review this feature to plan to adopt changes in their applications. For example, e-mail applications that render HTML e-mail in the Restricted Sites zone might need to be modified.

Users can only be affected by applications that do not completely render HTML content with this new setting. These applications will typically alert the user that some active behavior has been blocked from display. For example, when Outlook Express encounters this situation, it informs the user that it has restricted active content in the e-mail.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**New Internet Explorer security setting**
**Detailed description**
A new URL action setting, Binary and Script Behaviors, is in each Internet Explorer security zone. The default value for this setting is **Enable** for all zones except the Restricted Sites zone and the Locked-Down Local Machine zone. In the Restricted Sites zone, the default value is **Disable**. In the Locked-Down Local Machine zone, the default value is **Administrator approved**.

**Why is this change important? What threats does it help mitigate?**
This new setting helps mitigate attacks in which binary behaviors were being used maliciously and allows the user to control the use of binary behaviors on a per-zone basis.

**What works differently?**
Any use of any binary behaviors for HTML rendering from the Restricted Sites zone is blocked.

**How do I resolve these issues?**
To use binary behaviors from the Restricted Sites zone, an application will have to implement a custom security manager. (For more information, see "Creating a Customized URL Security Manager" in "Introduction to URL Security Zones" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=21863.)

When the binary behaviors URL action is exercised from a custom security manager, the URL action will pass in a string representation of the particular binary behaviors that can be enabled by that custom security manager as needed for application compatibility. The following process takes place when this URL action is exercised:

- Internet Explorer calls into a custom security manager (if available), using the `ProcessUrlAction` method with a `dwAction`

of `URLACTION_BEHAVIOR_RUN`.

- The *pContext* parameter points to a `LPCWSTR` that contains the behavior that a policy is being queried for. For example, `#default#time`.

- You set `*pPolicy =URLPOLICY_ALLOW` for your **SmartTag** behavior, from within your custom security manager, as appropriate.

In the absence of the custom security manager, the default action is to disallow running behaviors in the Restricted Sites zone, and to disallow running most behaviors in the Local Machine zone.

If you are a desktop administrator you can decide which binary behaviors to allow in the Locked-down Local Machine zone. To enable a behavior in the Locked-down Local Machine zone, you can add it to the list of administrator-approved behaviors as follows, replacing the namespace and behavior variables as appropriate to your environment:

`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedBehaviors`

`#%` *Namespace* `%#%` *Behavior* `%=dword:00000001`

Behaviors that are defined in this list will also be used for any other zone where the binary behavior restriction setting is configured to "Admin-Allowed" (65536).

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
None. This is only a setting to turn on or off the existing binary behaviors functionality.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
**Internet Explorer Binary Behaviors Settings**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| * | HKEY LOCAL MACHINE [or Current User] \Software\Microsoft \Internet Explorer\Main \Feature Control \FEATURE_BEHAVIORS | None | 1 | 0 - Off<br>1 - On |
| 2000 | HKEY_CURRENT_USER \Software \Microsoft \Windows \CurrentVersion \Internet Settings\Zones [or Lockdown_Zones] *\ | None | 3 - Disabled (for Restricted Sites zone)<br>65536 - Admin-approved (for the Locked-down Local Machine zone)<br>0 - Enabled (for all other zones) | 3 - Disabled<br>65536 - Admin-approved<br>0 - Enabled |

**Note**

* is used in the preceding table to represent that all processes are opted-in for this feature control setting by default.

The binary behaviors setting can also be modified through Group Policy as part of the Internet Explorer Security Zones and Content Ratings setting.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
If your code uses binary behaviors in the Restricted Sites zone, then you will need to change your code by implementing a custom security manager for your application. If your code uses binary behaviors in the Local Machine zone, then you will need to either implement a custom security manager, add your behaviors to the list of approved behaviors, or use Mark of the Web to load your pages in less restrictive zones. For more information, see the "Creating a Customized URL Security Manager" section in "Introduction to URL Security Zones" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=21863.

# Internet Explorer BindToObject Mitigation

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Notification Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does BindToObject Mitigation do?**
In Windows Server 2003 with Service Pack 1, the ActiveX security model is applied in all cases where URL binding is used to instantiate and initialize an object. The ActiveX security model allows controls to be marked as "safe for scripting" and "safe for initialization" and provides users with the ability to block or allow ActiveX controls by security zone, based on those settings. This allows greater flexibility and control of active content in Internet Explorer.

**Who does this feature apply to?**

- Web developers and network administrators need to be aware of these new restrictions to plan changes or workarounds for any possible impact to their Web site.
- Application developers should review this feature to plan to adopt changes in their applications.
- Users could be affected by sites that are not compatible with these stricter rules.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
None. Existing security functionality is being extended.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**ActiveX security model applied to URL object initializations**
**Detailed description**
The most effective way to remove ActiveX safety vulnerabilities is to apply security policies consistently at the source of the URL binding: URLMON. Declaring an ActiveX control in an HTML page using the `<object>` tag and `CODEBASE` attribute is one commonly known example of using `BindToObject`. The same functionality is used by any component that wants to resolve a URL and get back a stream or object. The ActiveX security model is now applied to all object initializations with a URL as a source.

**Why is this change important?**
In the case of ActiveX controls, the ActiveX security model allows controls to be marked as "safe for scripting" or "safe for initialization" and provides users with the ability to block or allow ActiveX controls by zone, based on those settings. In earlier versions of Windows, this security framework was not applied in all cases where URL binding took place. Instead, the calling code was responsible for assuring the integrity and security of the control, which could often result in security vulnerabilities. There are now a number of public exploit variations that expose this exact issue by going through Internet Explorer to compromise vulnerabilities in the calling code.

**What works differently?**
The ActiveX security model is applied to all object initializations with a URL as a source, and the "Safe for initialization" tag is applied to all objects. This mitigation only applies to cases where Internet Explorer resolves a URL, instantiates an object, and initializes the object with data retrieved from that URL.

**How do I resolve these issues?**
Application compatibility problems should be minimal. Applications can opt out if they have their own security manager. For more information about opting out of this security model, see "Security Considerations: URL Security Zones API," on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=21814.

Applications can also opt in or out of this mitigation using the feature control key FEATURE_SAFE_BINDTOOBJECT, as

described in the topic Internet Explorer Using Feature Control Registry Settings with Security Zone Settings.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
**Internet Explorer Object Caching**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| IExplore.exe<br><br>Explorer.exe<br><br>WMPlayer.exe | HKEY_LOCAL_MACHINE (or Current User)\Software \Microsoft \Internet Explorer\Main \FeatureControl \FEATURE_SAFE_BINDTOOBJECT | None | 1 | 0 - Off<br><br>1 - On |

# Internet Explorer Information Bar

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Information Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does the Internet Explorer Information Bar do?**
The Internet Explorer Information Bar in Windows Server 2003 Service Pack 1 replaces many of the common dialog boxes that prompted users for information in previous versions and provides a prominent area for displaying information that users may want to view or act upon. Examples of dialog boxes that have been replaced by Information Bar notifications include blocked ActiveX installs, downloads, and active content. The Information Bar will provide information similar to the notification area in Microsoft Outlook 2003, which informs users of blocked content.

**Who does this feature apply to?**
This feature applies to the following audiences:

- Users who need to understand how the new behavior will affect their Web browsing experience.
- System administrators, who need to know how to turn this functionality on or off for the client computers in their organization.
- Designers of Web sites that rely on add-ons, which will provide a different user experience.
- Developers of Web-based applications who need to understand how their experience changes. For example, this affects the development of ActiveX controls. ActiveX controls that are updates to controls that are currently installed on other computers will only be treated as updates if the GUID of the new control matches the current GUID.
- Developers of applications hosting the Web browser control will need to know how to use the new application programming interface (API) to take advantage of this new functionality.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Information Bar user interface**
**Detailed description**
The Information Bar looks and acts in a similar way to the Outlook 2003 blocked content notification area. It appears below Internet Explorer toolbars and above the Web page in view when a notification is present and disappears on the next navigation. The text in the Information Bar varies, depending on the notification that is provided, and wraps to two lines if the text exceeds the bounds of the notification area. If a user controls the focus (that is, which object in the browser can receive input) by pressing the TAB key, the Information Bar gets focus after the toolbar and before the Web page.

Either clicking or right-clicking the Information Bar brings up a menu that relates to the notification that is presented. This menu always contains a link to Information Bar Help, which provides more detailed information about the notification. Additional menu items related to the notification appear above the **Help** menu item.

Users can configure the Information Bar to play a sound when it appears; the default setting for the sound is **On**. When the Information Bar appears, the Windows trust icon appears in place of the **Error on page** notification on the status bar.

In certain cases, more than one action can be blocked. For example, a pop-up window may be blocked at the same time that an add-on install is blocked. In those cases, the text becomes more generic and the menus merge to show each action blocked at the top level, with each action's menus in a submenu.

There is a custom security zone setting for the Information Bar that enables users to change the settings of the Information Bar

by security zone. Users can choose to be notified with the Information Bar or to go back to the previous behavior and get a less prominent notification for file and code downloads.

**Why is this change important?**
In Windows Server 2003 Service Pack 1, Internet Explorer may block content that is necessary to complete certain online tasks. The Information Bar provides a prominent notification that informs users how to get Web pages they trust working again without the more intrusive prompt that was provided in Windows Server 2003.

**What works differently?**
For this information, see the next section, "What existing functionality is changing in Windows Server 2003 Service Pack 1?"

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Add-on install prompts**
**Detailed description**
In previous versions of Internet Explorer included with Windows Server 2003, when a Web page refers to an ActiveX control that is not currently on the computer, users are asked whether they want ActiveX controls to be downloaded. In the current version of Internet Explorer, this prompt is displayed in the Information Bar. The following table describes the elements in the Information Bar. Text in italics will be replaced by the specific items appropriate to the situation.

**Information Bar details for Add-on Install Prompts**

| Information Bar element | Message text |
|---|---|
| Information Bar text | This site might require the following ActiveX control: *Control_Name* from *Publisher_Name*. Click here to install... |
| Short text | Installation Blocked |
| Menu options | Install ActiveX Control...<br><br>What's the risk? |

Trusted publishers will work as they did in previous versions. The controls that are provided by these publishers install without requiring additional configuration.

Blocked publishers display the status bar icon. The control provided by these publishers will not install on the computer and does not go into the Information Bar.

Add-on upgrades work as they did in Windows XP SP1. Internet Explorer uses the following criteria when determining whether the control is an upgrade:

- The file that is registered as the ActiveX control must be signed with Authenticode technology. (This file is referenced from `HKEY_CLASSES_ROOT\CLSID\`*Control_clsid*`\InProcServer32`, where *Control_clsid* is the *CLSID* specified by the `OBJECT` tag.)

- The publisher name in the digital signature of the new control matches the publisher name in the digital signature of the existing control.

- If the ActiveX control is packaged in a CAB file, the CAB file must be signed. The DLL or OCX to be installed should also be signed in order for subsequent upgrades to bypass the Information Bar.

**Why is this change important?**
Providing add-on install prompts in the Information Bar rather than a dialog box reduces the occurrences of users inadvertently installing code on their computer.

**What works differently?**
Certain Web pages currently rely on users installing code to function correctly. Some sites redirect the user to a separate page that explains how to install the ActiveX control. If a site automatically redirects the page without providing the control on the new page, the Information Bar will appear on the redirect page to offer the opportunity to install the control. If, however, the site closes the window that attempted to install the ActiveX control, the customer might not get a chance to install the control.

**How do I resolve these issues?**
Web authors should ensure that the ActiveX control is also available on the page to which the user is redirected. This will ensure that users have ample opportunity to install the control.

Web page authors should not suggest that users lower their security settings, because it will not help in this situation. For additional guidance for Web authors on the changes introduced with Windows XP SP2 and subsequently included in Windows Server 2003 Service Pack 1, see "Fine-Tune Your Web Site for Windows XP Service Pack 2" on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=32775.

**Pop-up blocked notification**
**Detailed description**
Internet Explorer displays a notification in the Information Bar when a pop-up is blocked. This becomes a more obvious entry point to the Pop-up Blocker functionality, such as replaying the pop-up, adding the site to an "allow" list for pop-ups, or navigating to Pop-up Blocker settings. The Information Bar also provides a top level entry point to turn off the Information Bar for pop-ups if the user decides the notification is too big for this event.

**Information Bar details for pop-up blocked notification**

| Information Bar element | Message text |
|---|---|
| Information Bar text | Pop-up blocked. To see this pop-up or additional options, click here... |
| Short text | Pop-up Blocked |
| Menu options | Temporarily Allow Pop-ups<br><br>Allow Pop-ups for this Site<br><br>Settings |

**Why is this change important?**
Showing the pop-up blocked notification in the Information Bar gives higher priority to that notification. Users have a better understanding of where to go to see a blocked pop-up window or to see their Pop-up Blocker settings.

**What works differently?**
Turning off the Information Bar for the Pop-up Blocker causes the Pop-up Blocker to return to notifying users with the status bar icon. All the same menu items are accessible from this status bar icon if the bar is disabled for pop-ups. For more information, see "Internet Explorer Pop-up Blocker," later in this document.

**Automatic download prompts**
**Detailed description**
File download prompts that are launched automatically now appear in the Information Bar.

The Information Bar includes descriptive text that explains why the action was taken and provides a context sensitive menu that you can use to respond to the notification. The following table identifies the text that will appear in the Information Bar and the actions that you can select from the menu.

**Information Bar details for automatic download prompts**

| Information Bar element | Message text |
|---|---|
| Information Bar text | To help protect your security, Internet Explorer blocked this site from downloading files to your computer. Click here for more options... |
| Short text | File Download Blocked |
| Menu options | Download Software...<br><br>What's the Risk? |

**Why is this change important?**
By moving download prompts to the Information Bar, users can be prevented from installing unwanted code on their computers. Previously, sites could overwhelm users with file download prompts and, as a result, users could accidentally run unwanted software on their computer. With this change, file download prompts that are launched automatically are more likely the result of a user's deliberate click and not an accidental action.

**What works differently?**
Any time a site refers to a file download prompt without a user action, such as clicking on an element of the page, the prompt appears in the Information Bar.

**How do I resolve these issues?**
Web authors should ensure there is a link on the Web page that a user can click to get to the file download. If you use a script to navigate to the resource, the script should run synchronously within the context of the `OnClick` event handler for the link.

For additional guidance for Web authors on the changes introduced with Windows XP SP2 and subsequently included in Windows Server 2003 Service Pack 1, see "Fine-Tune Your Web Site for Windows XP Service Pack 2" on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=32775.

# Internet Explorer Local Machine Zone Lockdown

Updated: January 5, 2009

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Notification Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What Does Local Machine Zone Lockdown do?**

When Internet Explorer opens a Web page, it places restrictions on what the page can do, based on the page's Internet Explorer security zone. There are several possible security zones, each with different sets of restrictions. The security zone for a page is determined by its location. For example, pages that are located on the Internet will normally be in the more restrictive Internet security zone. They might not be allowed to perform some operations, such as accessing the local hard drive. Pages that are located on your corporate network would normally be in the Intranet security zone, and have fewer restrictions. The precise restrictions that are associated with most of these zones can be configured by the user through **Internet Options** on the **Tools** menu.

Prior to Windows XP Service Pack 2, the content on the local file system, aside from that cached by Internet Explorer, was considered to be secure and was assigned to the Local Machine security zone. This security zone normally allows content to run in Internet Explorer with relatively few restrictions. However, attackers often try to take advantage of the Local Machine zone to elevate privilege and compromise a computer.

Many of the exploits that involve the Local Machine zone were mitigated by other changes to Internet Explorer in Windows XP SP2. These changes were incorporated into Internet Explorer in Windows Server 2003 Service Pack 1. However, attackers may still be able to figure out ways to exploit the Local Machine zone. Currently, Internet Explorer further protects the user by locking down the Local Machine zone by default. Local HTML hosted in other applications will run under the less restrictive settings of the Local Machine zone used in previous version of Internet Explorer unless that application makes use of Local Machine Zone Lockdown.

Administrators will be able to use Group Policy to manage Local Machine Zone Lockdown and more easily apply it to groups of computers.

**Who does this feature apply to?**

All application developers should review this feature. Applications that host local HTML files in Internet Explorer are likely to be affected. Developers of standalone applications that host Internet Explorer will want to modify their applications to make use of Local Machine Zone Lockdown.

By default, Local Machine Zone Lockdown is only enabled for Internet Explorer. Developers will need to register their applications to take advantage of the changes. Applications that do not use this mitigation should independently review their applications for Local Machine zone attack vectors.

Software developers with applications that host Internet Explorer should use this feature by adding their process name to the registry as described later in this document. In the future, Microsoft might implement this feature using an "opt out" policy rather than an "opt in" policy. Applications that host Internet Explorer should be tested to ensure that they function properly with Local Machine Zone Lockdown enabled for their process.

Network Administrators might have local scripts that will be affected by these restrictions. Administrators should review the available solutions to enable their local scripts without compromising the security of their users' client computers.

Developers of Web sites that are hosted on the Internet or Local Intranet zones should not be affected by changes to the Local Machine zone, except when loading those files from the local machine during development.

Users could be affected by applications that are not compatible with these more stringent restrictions.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**

**Changes to Local Machine zone security settings**

**Detailed description**
The Local Machine zone is now more restrictive than the Internet zone. Any time that content attempts one of the following actions in this zone, the Information Bar will appear in Internet Explorer with the following text:

**To help protect your security, Internet Explorer has restricted this file from showing active content that could access your computer. Click here for options...**

The user can click the Information Bar to remove the lockdown from the restricted content.

The security settings that control the privileges that are granted to content running in the Local Machine zone are known as URL actions. When Local Machine Zone Lockdown is applied to a given process, it changes the behavior of URL actions from the previous Local Machine zone setting of **Enabled** to **Disabled**. As a result, scripts and ActiveX controls will not run. The default URL actions changed are:

- URLACTION _SCRIPT_RUN
- URLACTION_DOWNLOAD_UNSIGNED_ACTIVEX
- URLACTION_ACTIVEX_RUN
- URLACTION_ACTIVEX_OVERRIDE_OBJECT_SAFETY (to **Prompt**, not **Disabled**)
- URLACTION_CLIENT_CERT_PROMPT
- URLACTION_BEHAVIOR_RUN
- URLACTION_JAVA_PERMISSIONS
- URLACTION_BEHAVIOR_RUN (to **Administrator approved**, not **Disabled**)
- URLACTION_FEATURE_MIME_SNIFFING
- URLACTION_FEATURE_WINDOWS_RESTRICTIONS
- URLACTION_AUTOMATIC_DOWNLOAD_UI
- URLACTION_AUTOMATIC_ACTIVEX_UI

| **Note** |
| --- |
| URLACTION_FEATURE_ZONE_ELEVATION is set to **Disabled** in the Local Machine zone with or without this feature. |

For Local Machine Zone Lockdown, these settings are stored under a separate registry key:

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Internet
Settings\Lockdown_Zones\0
```

The default Local Machine zone URL action settings are found under:

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Internet Settings\Zones\0
```

**Why is this change important?**
This change helps prevent content on a user's computer from elevating privilege. Code with such elevated privilege can then run any code through an ActiveX control or read information with a script.

**What works differently?**
If a Web page uses any of the restricted types of content that were previously listed, Internet Explorer displays the Information Bar, as previously described.

HTML files that are hosted on the res: protocol on the local computer will automatically run under the security settings for the Internet zone. For more information about what these templates allow, see "Introduction to URL Security Zones" on the MSDN

Web site at http://go.microsoft.com/fwlink/?LinkId=26003.

**How do I resolve these issues?**
You can allow ActiveX and scripts to always run in Web pages that are launched from a CD by clicking **Yes** when presented with the following message:

**Active content can harm your computer or disclose personal information. Are you sure that you want to allow CDs to run active content on your computer?**

If your Web page needs to run ActiveX or scripting, you can add a Mark of the Web comment in the HTML code. This Internet Explorer feature allows the HTML files to be forced into a zone other than the Local Machine zone so that they can then run the script or ActiveX code based on the security template that would be applied to the URL identified in the comment. For example, if the URL specified is www.contoso.com and that URL is present in your Local Intranet sites list, the page uses the security template for the Local Intranet Sites zone. However, if www.contoso.com is listed in the Trusted Sites zone, the page is treated as if it were part of the Internet zone. This is by design for security purposes in Windows XP Service Pack 2 and later. This setting works in Internet Explorer 4 and later. To insert a Mark of the Web comment into your HTML file, add one of the following comments:

`<!-- saved from url=` *(0022)http://www.example.com* `-->`

Use this comment when you are inserting a Mark of the Web into a page whose domain is identified, replacing *http://www.example.com* with the URL of the Internet or intranet domain that the page is hosted by. Include the length of the URL in parenthesis used for the Mark of the Web before the URL, for example *(0022)*.

If you want your Web page to always be treated as though it were part of the Internet zone, you can use the following Mark of the Web:

```
<!-- saved from url=(0014)about:internet -->
```

Use this comment when you need to generically insert a Mark of the Web. The `about:internet` portion will place the page in the Internet zone.

Beginning with Windows Server 2003 Service Pack 1 and Windows XP Service Pack 2, this HTML comment can also be used with .mht files, known as multipart HTML or .xml files. Mark of the Web will be ignored for .mht or .xml files in earlier versions of Internet Explorer.

As another option, you can create a separate application that hosts the HTML content in the Internet Explorer Web Object Control (WebOC). The HTML is then no longer bound by the same rules that apply to content run in Internet Explorer. When the HTML content runs in the other process, it can have full rights as defined by the developer or the zone policy for that process.

An easy way to do this is to save your content as an .hta (HTML application) file and try to run the file again in the Local Machine zone. An .hta file is hosted in a different process and therefore is not affected by the mitigation. However, .hta files run with full privileges, so you should not allow code that is not trusted to run in this manner.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
Developers should test their applications and enable the lockdown in order to offer enhanced levels of security. Developers of standalone applications should plan to adopt these changes in their applications that host Internet Explorer.

Developers of ActiveX controls that previously allowed elevated privileges in the Local Machine zone should not change their controls to allow elevated privileges in another zone. Instead, these controls should be converted to run only from an HTML application (.hta file) or a standalone application that runs outside of Local Machine Zone Lockdown.

By default, Local Machine Zone Lockdown is not enabled for non-Internet Explorer processes. Developers must explicitly register their applications to take advantage of the changes. Application developers that do not use this mitigation should independently review their applications for Local Machine zone attack vectors. To enable Local Machine Zone Lockdown for your application, go to the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Internet Explorer\Main\FeatureControl\FEATURE_LOCALMACHINE_LOCKDOWN
```

Add a `REG_DWORD` value to this key named for your application (for example, MyApplication.exe) and set it to 1. Any other setting for this value will disable Local Machine Zone Lockdown for the application.

To control whether Local Machine Zone Lockdown is applied to Web pages launched from a CD, go to the following registry key and value:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Internet
Explorer\Main\FeatureControl\FEATURE_LOCALMACHINE_LOCKDOWN\Settings\LOCALMACHINE_CD_UNLOCK
```

Setting this value to 1 disables this feature for Web pages launched from a CD on the user's computer.

# Internet Explorer MIME Handling Enforcement

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Information Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does MIME Handling Enforcement do?**
Internet Explorer uses Multipurpose Internet Mail Extensions (MIME) type information to decide how to handle files that have been sent by a Web server. For example, when there is a Hypertext Transfer Protocol (HTTP) request for .jpg files, when they are received, they will generally be displayed to the user in an Internet Explorer window. If Internet Explorer receives an executable file, Internet Explorer generally prompts the user for a decision on how to handle the file.

Internet Explorer in Windows Server 2003 with Service Pack 1 will follow stricter rules than Internet Explorer in Windows Server 2003. These rules are designed to protect users from accidentally downloading or executing a dangerous file due to misleading MIME or file name extension information.

**Who does this feature apply to?**
Web developers need to be aware of these new restrictions to plan changes or workarounds for any possible impact to their Web site.

Application developers should review this feature to plan to adopt changes in their applications. The feature is not enabled for non-Internet Explorer processes by default and developers will need to register their applications to take advantage of the changes.

End users will be affected by sites or applications that are not compatible with these stricter rules.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**MIME-handling file type agreement enforcement**
**Detailed description**
When files are served to the client, Internet Explorer uses the following pieces of information to decide how to handle the file:

- File name extension, the corresponding ProgID and CLSID for the registered handler of that file name extension.
- Content-Type from the HTTP header (MIME type), the corresponding ProgID, and CLSID for the registered handler of that Content or MIME type.
- Content-Disposition from the HTTP header.
- Results of the MIME sniff.

Internet Explorer is more restrictive about executing a downloaded file that could be dangerous than it was in Windows Server 2003.

Internet Explorer will enforce consistency between how a file is handled in the browser and how it is handled in the Windows shell. As the file is downloaded into the cache, Internet Explorer will compare the MIME-type of the cache file to the extension of the cache file. If there is a mismatch between the MIME-type and the file name extension, Internet Explorer will attempt to reconcile that mismatch by renaming the file in the cache.

Before a file is loaded in its MIME handler or executed by its extension handler, Internet Explorer will compare the CLSIDs of the MIME handler and the extension handler. If there is still a mismatch between the two handlers, Internet Explorer will impose a mandatory prompt for the user to confirm that the user wants to load the file in the MIME handler. If the MIME handler rejects the mismatched file, Internet Explorer will show a download error dialog and not automatically execute the file in the Windows

shell extension handler but the dialog will allow saving the file.

There is a related but separate change to prevent execution of potentially corrupt files in their shell extension handlers. Internet Explorer will show the download error dialog for any file that is rejected by its MIME handler with the error code `E_Cannot_Load_Data` and will not execute that file in its shell extension handler regardless of MIME-type or file name extension.

These changes do not affect cases where a `"Content-disposition=attachment"` HTTP header is used for the file. In those cases, the file name or extension suggested by the server is considered final and the file can be executed regardless of MIME/extension mismatch if the user chooses to accept the file download prompt.

**Why is this change important?**
If file type information is misreported by the server and that information is saved to the computer, a dangerous file could be incorrectly executed later. For example, Internet Explorer might download a file that appears to be a text file. If the file can't be loaded by its MIME handler and has a .doc file name extension, the .doc file might run in an application such as Microsoft Word without prompting the user. In Microsoft Word, the file may be able to use active content, such as a macro, to run a program (such as a virus) on the user's machine.

**What works differently?**
Internet Explorer will now attempt to rename downloaded files in the Internet Explorer cache to have matching content types and extensions to protect against files that mislead the user about their type.

Internet Explorer will prompt the user to download the file and will no longer execute MIME and extension mismatched files that are rejected by the registered MIME handler.

Internet Explorer will also not execute a file in its shell handler if the `E_Cannot_Load_Data` error code was reported by its MIME handler.

Instead of executing such a file in the shell-handler, Internet Explorer will show an error dialog and give the user an option to save the file.

Web developers can isolate nonworking applications due to this behavior by switching off the functionality, as covered in the Settings section later in this document.

**How do I resolve these issues?**
Web developers must change their Web servers to host files, using consistent content-type headers and file name extensions. If this is not possible, Web developers can use the `"Content-disposition=attachment"` HTTP header to directly send the file to its extension handler rather than to the MIME handler. Note that file downloads with the `"Content-disposition=attachment"` header will prompt the user to open or save the file.

If you have developed a MIME handler and intentionally rely on Internet Explorer to execute files that your MIME handler rejects, you will need to make changes in your MIME handler to accommodate this change. The most secure change would be to natively handle the file directly in the MIME handler rather than rejecting the file.

For some scenarios, it may not be possible to change the behavior of the MIME handler to natively handle downloaded files. In those scenarios, there are a few options.

You may choose to develop a MIME handler and file name extension handler that are both part of the same CLSID; Internet Explorer will accept the CLSID match and therefore not prompt the user to download the file or block the file from execution in the extension handler.

If the MIME handler does not need to be loaded and will cause errors due to MIME and extension mismatch, the developer can mark the MIME handler to be ignored by Internet Explorer in the case of a MIME and extension mismatch. For example, if the MIME handler for a certain media MIME type has a mismatched extension and needs to be executed directly to be played properly, the developer can mark the ProgID of their MIME handler to be ignored on the mismatch when the media file name extension belongs to a different ProgID. To do this, the developer sets the following value in the registry with the MIME handler to ignore:

```
HKEY_CLASSES_ROOT\PROG_ID_OF_MIMEHANDLER_TO_IGNORE\"PreferExecuteOnMismatch"=DWORD:00000001
```

If neither of these solutions is viable, developers should notify their users of the incompatibility and explain to the user how to save the mismatched file to the file system and then launch it manually.

If your scenario is affected by unwanted file-download prompts because of an irreconcilable MIME/extension mismatch, you can register you MIME handler ProgID to bypass all download prompts including the new prompt on mismatch.

Before doing this, you should confirm that your MIME handler can securely deal with any file that is delegated to it. For example, you should confirm that your handler would never allow an attacker to gain more privilege than allowed by the zone

of the originating file. This should be done through threat modeling, code review for secure failure modes, and buffer overruns. If you determine that your MIME handler is capable of safely handling any file that might be delegated to it, you can register your handler to circumvent download prompts by adding a new key to one of the following registry settings:

```
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\InternetSettings\Secure_Mime_Handlers
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\InternetSettings\Secure_Mime_Handlers
```

The key should be named with the ProgID of your MIME handler and have a DWORD=00000001.

**MIME sniffing file type elevation**
**Detailed description**
One of the backup criteria for determining a file type is the result of the MIME sniff. By examining (or sniffing) a file, Internet Explorer can recognize the bit signatures of certain types of files. In Windows Server 2003 Service Pack 1, Internet Explorer MIME sniffing will not promote files of type text\plain to more dangerous file types in the Restricted Sites zone. For example, files that are received as plain text but that include HTML code will not be promoted to the HTML type, which could contain active content.

**Why is this change important?**
This change provides users additional defense in depth against malicious content posted on a friendly Web server where the server serves files with content-type=text\plain for a file but an attacker has managed to load HTML with active content into the file.

**What works differently?**
Web servers that do not include the correct Content-Type header with their files and that use nonstandard file name extensions for HTML pages now may have their pages rendered as plain text rather than HTML.

**How do I resolve these issues?**
You should configure Web servers to use the correct Content-Type headers or you can name the files with the appropriate file name extension for the application that should handle the file.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| IExplore.exe<br>Explorer.exe | HKEY_LOCAL_MACHINE(or Current User)\Software \Microsoft \Internet Explorer\Main \FeatureControl \FEATURE_MIME_HANDLING\ | None | 1 | 0 - Off<br>1 - On |
| IExplore.exe<br>Explorer.exe | HKEY_LOCAL_MACHINE(or Current User)\Software \Microsoft \Internet Explorer\Main \FeatureControl\FEATURE_MIME_SNIFFING\ | None | 1 | 0 - Off<br>1 - On |

**MIME Sniffing behavior per-zone settings**
The new MIME sniffing restriction is controlled by the **Open files based on content, not file extension** security setting which can be enabled or disabled for individual security zones. The following table provides a reference of the default settings per security zone:

| Security zone | "Open files based on content, not file extension" Default security setting |
|---|---|
| Restricted Sites zone | Disable |
| Internet zone | Enable |
| Intranet zone | Enable |
| Trusted Sites zone | Enable |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
You should configure Web servers to use the correct Content-Type headers. You can also name the files with the appropriate file name extension for the application that should handle the file.

# Internet Explorer Network Protocol Lockdown

Applies To: Windows 7, Windows Server 2003 with SP1, Windows Vista, Windows XP

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Notification Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What Does Network Protocol Lockdown Do?**
Internet Explorer can be configured to lock down HTML content from particular network protocols in additional zones besides the Local Machine zone. This feature allows an administrator to extend the same restrictions of the Local Machine Zone Lockdown (which is described previously in this document) to be applied to any content on any arbitrary protocol in any security zone. For example, an administrator can configure Internet Explorer to lock down HTML content hosted on the Shell: protocol if it is in the Internet zone. Since the Shell: protocol's most common use is for local content and not Internet content, this mitigation can reduce the attack surface of the browser against possible vulnerabilities in protocols less commonly used than HTTP.

**Who does this feature apply to?**
By default, Network Protocol Lockdown is not enabled for any application.

All application developers should review this feature. Applications that host HTML files over non-HTTP protocols in Internet Explorer may be affected in organizations where administrators elect to apply additional restrictions. Developers of standalone applications that host Internet Explorer might want to modify their applications to make use of Network Protocol Lockdown.

Developers who chose to opt in to Network Protocol Lockdown should register their applications to take advantage of the changes. Applications that do not use this mitigation should independently review their applications for support for arbitrary protocols.

Software developers with applications that host Internet Explorer can use this feature by adding their process name to the registry as described later in this document. In the future, Microsoft might implement this feature with certain uncommonly used protocols restricted by default and with an "opt-out" policy for applications rather than the current "opt-in" policy for applications. Applications that host Internet Explorer should be tested to ensure that they function properly with Network Protocol Lockdown enabled for their process.

Network Administrators should consider adding unused protocols to the restricted protocol list on managed desktop machines. If the network administrator enables this restriction, there may be HTML files that will be affected.

Developers of Web sites that are hosted on the HTTP protocol should not be affected by restrictions to other protocols.

Users are most likely to be affected by these more stringent restrictions if their Network Administrator chose to restrict certain protocols for their desktop.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Changes to security settings for restricted protocols**
**Detailed description**
With Windows Server 2003 Service Pack 1, HTML content in an application that has "opted in" to use the Network Protocol Lockdown feature that is served on one of the restricted protocols will be restricted to run at a higher security level. Any time the restricted protocol content attempts to use a restricted feature, such as ActiveX controls, the Information Bar will appear in Internet Explorer with the following text (text may be different for other blocked URL actions):

**Internet Explorer has blocked this site from using an ActiveX control in an unsafe manner. As a result this page may not display correctly.**

The user can click the Information Bar to remove the lockdown from the restricted content. The change in setting using the Information Bar is per session only, unless the policies are changed in the registry.

The security settings that are locked down for the content on the restricted protocols are the same as the settings enforced for

the Local Machine zone lockdown, which is described earlier in this document. Please consult that section to review exactly which security settings are enforced for the content on the restricted protocols.

**Restricted protocols feature is off by default for Internet Explorer and all applications**
**Detailed description**
The behavior of the Network Protocol Lockdown is controlled per-process by a new Internet Explorer Feature Control setting. Since this feature is designed to provide an additional layer of defense-in-depth for network administrators, the default Internet Explorer processes, IExplore.exe and Explorer.exe are not opted in by default. To opt in to the Network Protocol Lockdown, network administrators or developers should add a DWORD at either of the following locations where the name is their process name and the value is set to 1 to have the mitigation apply to them. To forcibly opt out, set the value of the key to 0. If the administrator decides to put the setting under the Policies hive, set a REG_SZ instead of a DWORD.

- `HKEY_LOCAL_MACHINE\Software\(Policies)\Microsoft\Internet Explorer\Main\FeatureControl\FEATURE_PROTOCOL_LOCKDOWN`

- `HKEY_CURRENT_USER\Software\(Policies)\Microsoft\Internet Explorer\Main\FeatureControl\FEATURE_PROTOCOL_LOCKDOWN`

Applications may want to proactively opt out to prevent the mitigation from being applied to them when a wildcard is used to force the mitigation into Opt-out mode.

**Behavior per-zone when an application has opted in**
For an opted in process, the behavior of the Network Protocol Lockdown is also controlled per zone by a new Internet Explorer security setting or URL action called URLACTION_ALLOW_RESTRICTEDPROTOCOLS. This URL action will be set to the following values.

| Security Zone | Default behavior for restricted protocols | Example user situation |
|---|---|---|
| Restricted Sites Zone | Disallow | Since ActiveX is never allowed in the Restricted Sites zone by default, the Information Bar is not shown when a restricted protocol is encountered. The Information Bar might be shown in the case where a URL action that was previously allowed in the Restricted Sites zone is now disallowed under network protocol lockdown. In this case, the user will NOT be able to click the Information Bar to allow the action. |
| Internet Zone | Prompt | If the administrator locks down the file:// protocol, HTML that uses script over the file:// protocol is restricted, but users can click the Information Bar to allow it. |
| Intranet Zone | Prompt | If the administrator locks down the local:// protocol, HTML that uses Java over the local:// protocol is restricted, but users can click the Information Bar to allow it. |
| Trusted Sites Zone | Prompt | If the administrator locks down the Shell:// protocol, HTML that uses Binary Behaviors over the Shell:// protocol is restricted, but users can click the Information Bar to allow it. |
| Local Machine Zone | Prompt | If local machine lockdown is enabled, its settings will supersede those established by network protocol lockdown settings. |

**Per-Zone Protocol Lock Down**
The list of protocols that are restricted is defined separately for each zone to allow some protocols to be locked down in some zones but run without restrictions in other zones. Protocols can be restricted for a given zone by writing the protocol name to the restricted list for a particular security zone.

| Security Zone | Registry location of the list of restricted protocols for each zone | Security settings applied to restricted protocol content |
|---|---|---|

| | | |
|---|---|---|
| Restricted Sites Zone | HKEY_LOCAL_MACHINE<br>-or-<br>HKEY_CURRENT_USER<br>\Software\(Policies) \Microsoft\Windows \CurrentVersion\Internet Settings \RestrictedProtocols\4 | HKEY_CURRENT_USER \Software\Microsoft \Windows\CurrentVersion \Internet Settings \Lockdown_Zones\4 |
| Internet Zone | HKEY_LOCAL_MACHINE<br>-or-<br>HKEY_CURRENT_USER<br>\Software\(Policies) \Microsoft\Windows \CurrentVersion\Internet Settings \RestrictedProtocols\3 | HKEY_CURRENT_USER \Software\Microsoft \Windows\CurrentVersion \Internet Settings \Lockdown_Zones\3 |
| Intranet Zone | HKEY_LOCAL_MACHINE<br>-or-<br>HKEY_CURRENT_USER<br>\Software\(Policies) \Microsoft\Windows \CurrentVersion\Internet Settings \RestrictedProtocols\2 | HKEY_CURRENT_USER \Software\Microsoft \Windows\CurrentVersion \Internet Settings \Lockdown_Zones\2 |
| Trusted Sites Zone | HKEY_LOCAL_MACHINE<br>- or -<br>HKEY_CURRENT_USER<br>\Software\(Policies)<br>\Microsoft\Windows \CurrentVersion\Internet Settings<br>\RestrictedProtocols\1 | HKEY_CURRENT_USER \Software\Microsoft \Windows\CurrentVersion \Internet Settings \Lockdown_Zones\1 |
| Local Machine Zone | HKEY_LOCAL_MACHINE<br>- or -<br>HKEY_CURRENT_USER<br>\Software\(Policies)<br>\Microsoft\Windows \CurrentVersion\Internet Settings<br>\RestrictedProtocols\0 | HKEY_CURRENT_USER \Software\Microsoft \Windows\CurrentVersion \Internet Settings \Lockdown_Zones\0 |

**Protocols to consider for lock down**

The default list of restricted protocols is blank. Network administrators should add additional protocols to the lockdown that they know are not needed in their organization for a particular zone. Network administrators should consider restricting some of the following default Windows protocols on managed desktop machines and other protocols that are not needed for rendering HTML with active content in the organization.

- local://
- file://
- shell://

- hcp://
- ftp://

**Why is this change important?**
This change provides general defense-in-depth against vulnerabilities in less frequently used protocols. For example, an ActiveX control running under the local:// protocol might assume that it is loaded in the Local Machine zone and it may grant elevated privilege to the hosting page.

**What works differently?**
If a Web page served on a protocol that is restricted for a given zone uses any restricted content, such as ActiveX, Internet Explorer will display the Information Bar, as previously described.

**How do I resolve these issues?**
If your Web page needs to run ActiveX or scripting on a protocol that should be restricted for your intranet, you might allow the HTML to render correctly by moving the domain for that HTML to the trusted sites zone on the managed desktop machines. As a long term solution, you can look for ways to move the content off of the restricted protocol or if that's not possible, you might remove the active content from the restricted protocol pages entirely by performing needed computations on the server using a server-side script such as an Active Server Page.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
Since this feature is off by default, you will probably not need to change your HTML content unless it runs over a protocol that is restricted by a network administrator for your organization.

# Internet Explorer Object Caching

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Information Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does Object Caching do?**
In previous versions of Windows Server 2003 with Internet Explorer, some Web pages could access objects cached from another Web site. In Windows Server 2003 Service Pack 1, a reference to an object is no longer accessible when the user navigates to a new domain.

**Who does this feature apply to?**
Web developers should review this feature and plan to adopt changes to their Web site.

Application developers should review this feature and plan to adopt changes in their applications.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
None. Existing functionality has been extended.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Security context is invalidated upon navigation to a different domain**
**Detailed description**
For Windows Server 2003 Service Pack 1, there is now a new security context on all scriptable objects so that access to cached objects (except for ActiveX controls) is blocked. In addition to blocking access when navigating across domains, access is also blocked when navigating within the same domain. (In this context, a domain is defined as a fully qualified domain name, or FQDN.) A reference to an object is no longer accessible after the context has changed due to navigation.

**Why is this change important? What threats does it help mitigate?**
Prior to Internet Explorer 5.5, navigations across HTML pages (or to subframes) purged instances of MSHTML, which is the Microsoft HTML parsing and rendering engine. With the Internet Explorer 5.5 Native Frames architecture, an instance of MSHTML lives across navigations. This introduced a new class of vulnerabilities, because objects could be cached across navigations. If an object can be cached and provide access to the contents of a Web page from another domain, there is a cross-domain hole.

Once you can get to properties on the inner document, script outside of a page's domain can access the contents of an inner page. This is a violation of the Internet Explorer cross-domain security model.

For example, you can use this method to create scripts that listen to events or content in another frame, such as credit card numbers or other sensitive data that is typed in the other frame.

**What works differently? Are there any dependencies?**
In those few classes that don't already have them, four more bytes are added for the cached markup. There should be no noticeable impact on speed.

**How do I resolve these issues?**
For most of these classes of vulnerabilities, Internet Explorer 5 would have crashed, so the application compatibility risk of resolving the exploit should be small. Other applications might need to be addressed on a case by case basis.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
**Internet Explorer Object Caching**

| Setting name | Location | Previous default value | Default value | Possible values |
| --- | --- | --- | --- | --- |

| | | | | |
|---|---|---|---|---|
| IExplore.exe<br><br>Explorer.exe | HKEY_LOCAL_MACHINE (or Current User)\Software \Microsoft \Internet Explorer\Main \FeatureControl \FEATURE_OBJECT_CACHING | None | 1 | 0 - Off<br><br>1 - On |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
If your application is attempting to use a cached object, you might encounter Access Denied errors. In this instance you must recache the object before you access it using a script.

In the following example, the security context is invalidated when the `designMode` property is set on a document object.

Broken script example

```
var d = myFrame.document;
d.designMode = "On";
d.open();  <-------------------------causes permission denied error
```

Fixed script example

```
var d = myFrame.document;
d.designMode = "On";
d = myFrame.document;   // re-establish pointer to document object.
d.open();
```

Also, when you compare the values of the two `frame.frames` properties of an object, the results may be incorrect, or the values may not be retained. This is because the frames object is now wrapped with a security wrapper. When the object caching feature is enabled, this security wrapper is applied. Therefore, access to all cached objects is blocked. This can cause the frames object comparison to return false even though the frames are equal. To resolve this problem, you can use the following method to compare the value of the frames:

```
a.name == parent.frames[1].frames.name
```

# Internet Explorer Pop-up Blocker

Applies To: Windows Server 2003 with SP1

| **Note** |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Information Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does Pop-up Blocker do?**
Pop-up Blocker blocks most unwanted pop-up windows from appearing. Pop-up windows that are opened when the end user clicks a link will not be blocked.

End users and IT administrators can let specific domains open programmatic pop-up windows. Developers will be able to use or extend the pop-up functionality in Internet Explorer for applications hosting Internet Explorer.

**Who does this feature apply to?**
For most end users, browsing the Web will be less annoying, because unwanted pop-up windows will not automatically appear.

For Web developers, Pop-up Blocker affects the behavior of windows opened by Web sites, for example, by using the `window.open()` and `showHelp()` methods

For application developers, there is a new user interface called `INewWindowManager`.

Applications that use the rendering engine in Internet Explorer to display HTML can choose to use or extend the Pop-up Blocker functionality.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
The Pop-up Blocker is a new feature for Internet Explorer which can be broken down into three sections:

- User experience changes, defaults, and advanced options.
- Changes in behavior of current application programming interfaces (APIs), such as `window.open` and `showHelp`.
- The new `INewWindowManager` interface, which allows applications to use the pop-up technology in Internet Explorer.

**Pop-up Blocker features**
**Detailed description**
**Defaults**
Pop-up Blocker is turned on by default. There are restrictions on the size and position of pop-up windows, regardless of the Pop-up Blocker setting. Pop-up windows cannot be opened larger than or outside the viewable desktop area. For more information, see "Windows Restrictions" in this document.

When this functionality is enabled, automatic and background pop-up windows are blocked, but windows that are opened by a user click will still open in the usual manner. Note that sites in the Trusted Sites and Local Intranet zones do not have their pop-up windows blocked by default, as they are considered safe. This setting can be configured in the **Security** tab in **Internet Options**.

**Enabling Pop-up Blocker**
Pop-up Blocker is enabled by default. You can change this in the **Tools** menu, with the **Pop-up Blocker** item, or in the Information Bar when a pop-up is blocked.

**When a pop-up window is blocked**
If a site opens a pop-up window that is blocked by Internet Explorer, a notification appears in the Information Bar and status bar and a sound is played. If you click the notification in the Information Bar or status bar, you see a menu with the following

options:

- **Temporarily Allow Pop-ups**. Reloads the page, allowing pop-up windows.
- **Always Allow Pop-ups from This Site**. Adds the current site to the Allow list.
- **Settings**. Shows more Pop-up Blocker settings menu items and gives access to the **Pop-up Blocker Settings** window.

| **Note** |
| --- |
| You can allow pop-up windows to open by pressing CTRL while the pop-up is opening |

**Advanced options**

Internet Explorer provides advanced configuration options through **Pop-up Blocker Settings**. To access these settings, open **Internet Options**, click **Privacy** and then in the **Pop-up Blocker** area click **Settings** to open **Pop-up Blocker Settings**. You can configure the following options:

- **Address of Web site to allow**. Enter a URL and click **Add** to add sites to the **Allowed sites** list. Any site on the list can open pop-up windows.
- **Filter Level**. There are three different filter levels that you can use with Pop-up Blocker:
    - **High: Block All Pop-ups**. The default behavior of Pop-up Blocker allows sites to open a pop-up window when the user clicks a link. This setting changes that behavior by blocking windows that are opened from a link. If this setting is enabled, you can allow pop-up windows to open by pressing the CTRL key at the same time that you click the link to launch the pop-up.
    - **Medium: Block most automatic pop-ups**. This is the default Pop-up Blocker setting. This setting blocks most pop-ups that launch automatically when a Web site is loaded into your browser, but doesn't block pop-ups that are opened when you click a link.
    - **Low: Allow pop-ups from secure sites**. This setting allows secure Web sites (those that use the https:// protocol) to automatically launch pop-up windows without requiring that you add them to the **Allowed sites** list.
- **Configure Sound**. You can toggle whether or not Pop-up Blocker plays a sound when a pop-up is blocked through **Pop-up Blocker Settings**. To do this, open **Internet Options**, click **Privacy** and then in the **Pop-up Blocker** area click **Settings** to open **Pop-up Blocker Settings**. In the **Notification and Filter Level** area, select or clear the check box next to **Play a sound when a pop-up is blocked** as appropriate.

  You can also change the sound that plays. To do this, click **Start**, click **Control Panel**, and then double-click the **Sounds and Audio Devices** icon to open its properties sheet. Then click the **Sounds** tab, click the **Blocked Pop-up Window** program event and choose the sound to play in the **Sounds** drop-down list.

| **Note** |
| --- |
| The Pop-up Blocker is used by default in the Internet and the Restricted Sites security zones. You can expand the scope of the Pop-up Blocker to include the Local Intranet or Trusted Sites security zone by clicking the zone and then either moving the slider in the **Security level for this zone** area to **Medium** or **High** or clicking **Custom level** and then changing the setting for **Use Pop-up Blocker** to **Enabled**. |

**When will you see pop-up windows while Pop-up Blocker is enabled?**

You will still see pop-ups windows opened in the following cases:

- The pop-up is opened by a link that the user clicked.
- The pop-up is opened by software that is running on the computer.

- The pop-up is opened by ActiveX controls that are instantiated from a Web site.
- The pop-up is opened from the Trusted Sites or Local Intranet zones and you have not expanded the scope of the Pop-up Blocker to include those zones.

**Why is this change important?**
Pop-ups have been misused in many ways. By blocking pop-ups, you have more control over your browsing experience.

**INewWindowManager**
**Detailed description**
By default, the Pop-up Blocker functionality does not apply to applications that host the `WebBrowser` control or MSHTML. These applications have the ability to use or extend Pop-up Blocker, use their own Pop-up Blocker, or disable pop-up management for their application through the `INewWindowManager` interface.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Methods: window.open(), window.external.navigateAndFind(), showHelp()**
**Detailed description**
If one of these functions normally returns a window object, then the function will return null when a window is blocked. Web developers can check for null to determine whether the window they attempted to open was blocked.

Windows that are outside the viewable screen when they are opened are positioned onto the viewable area.

Windows that are larger than the viewable screen when they are opened are resized to the viewable area.

For more information, see "Internet Explorer Window Restrictions" later in this document.

**What works differently?**
In the Internet zone, the Pop-up Blocker blocks windows that are automatically opened by these methods without the user clicking a link. Windows that are opened by these methods by clicking a link might also be blocked if the customer has enabled the more restrictive blocking setting.

**How do I resolve these issues?**
Ensure that all windows that are opened with `window.open()` are opened through user interaction and not automatically through your code.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
**Pop-up Blocker Settings**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| URLname | HKEY_CURRENT_USER \Software\Microsoft \Internet Explorer\New Windows\Allow | None | Empty | URL names of trusted sites |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
Web page authors should check for a NULL return value for any windows that you open. This will indicate whether the pop-up window opened successfully and allow you to handle either case.

If your software opens windows automatically, they will be blocked. Look for alternative ways of doing the same thing as described earlier in this document. The best way to open a window is to have the customer click a link or graphical element.

# Internet Explorer Untrusted Publishers Mitigations

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Information Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does Untrusted Publishers Mitigations do?**
This feature allows the user to block all signed content from a given publisher without showing the Authenticode dialog box to the user while doing so. This stops code from the blocked publisher from being installed. This feature also blocks installation of code with invalid signatures.

**Who does this feature apply to?**
This feature applies to all users, since it deals with installation and running of applications that are signed.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Blocked publisher**
**Detailed description**
Through Authenticode, the user can block content for a given publisher from installing or running. To do this, the user selects the **Never install software from** *PublisherName* check box in the Authenticode dialog box. If selected, the user is never prompted when code that is identified with the publisher's digital signature is trying to install itself on the system. It will be automatically blocked without showing the Authenticode dialog box.

**Why is this change important?**
This feature was designed to help users block ActiveX controls and other signed file formats from repeatedly prompting them on the Web. Users had no way of saying, "I don't want content from this publisher. Do not ask me again." Because they didn't have this feature, many users installed applications or content just to keep from encountering repeated prompts.

**What works differently?**
Previously, the Authenticode dialog box only supported selecting the **Always trust content from** *PublisherName* check box, which allowed the automatic installation of code from a specified publisher without prompting the user. Now the user can perform the opposite action and designate a publisher as untrusted. No application compatibility issues should be encountered for trusted code.

**How do I resolve these issues?**
You can unblock a publisher of an add-on by using **Manage Add-ons** in Internet Explorer. To unblock a publisher to enable the download of a specific file, you can remove the publisher from the **Untrusted Publishers** list. To do this, in Internet Explorer, on the **Tools** menu, click **Internet Options**, click the **Content** tab, click the **Publishers** button and then remove the publisher's name from the **Untrusted Publishers** list.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Blocking invalid signatures**
**Detailed description**
By default, Windows blocks the installation of signed code if it has an invalid digital signature.

**Why is this change important? What threats does it help mitigate?**
If code has an invalid signature, it usually means that the code has been changed since it was signed. When this happens, Internet Explorer considers the code to be unsigned, because someone might have tampered with it. By default, Internet Explorer blocks ActiveX applications that are unsigned that come from the Internet zone. This extends that functionality so that it applies to all code with invalid signatures.

**What works differently?**
By default, code with invalid signatures cannot be installed.

**How do I resolve these issues?**
To revert to previous functionality and allow unsigned code to run, see the `RunInvalidSignatures` setting in the "What settings

are added or changed in Windows Server 2003 Service Pack 1?" section below.

**One prompt per control per page**
**Detailed description**
Internet Explorer only prompts once per ActiveX control per page.

**Why is this change important? What threats does it help mitigate?**
This change helps defend against the social engineering trick of prompting the user a number of times for the same control. Even though users repeatedly refuse, they cannot get out of the loop, and they might eventually accept the installation out of frustration.

**What works differently?**
The user only sees one prompt per page per control.

**Ellipsis placed on text for application description and publisher name**
**Detailed description**
When the text that is given for the application description, file name, or publisher name is wider than the dialog box in width, Internet Explorer places an ellipsis on the text. This helps indicate to the user that there is more text that they are not seeing.

**Why is this change important? What threats does it help mitigate?**
This reduces the ability of control authors from placing marketing text and EULAs in the dialog box or using other social engineering tricks to overwhelm the users and get them to install the control.

**What works differently?**
Application description, file names, and publisher names will contain an ellipsis if the text is longer than the width of the dialog box. No applications or Web pages should need to be modified.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| RunInvalidSignatures | HKEY_CURRENT_USER \Software\Microsoft \Internet Explorer \Download<br><br>HKEY_LOCAL_MACHINE \Software\Microsoft \Internet Explorer \Download | None | 0<br><br>(Controls with invalid signatures will be blocked, regardless of zone.) | 1<br><br>(Controls with invalid signatures will be allowed to run, regardless of zone.) |

# Internet Explorer Window Restrictions

Applies To: Windows Server 2003 with SP1

| **Note** |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Information Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does Window Restrictions do?**

Internet Explorer provides the capability for scripts to programmatically open additional windows of various types, and to resize and reposition existing windows. The Window Restrictions security feature, formerly called UI Spoofing Mitigation, restricts two types of script-initiated windows that have been used by malicious persons to deceive users:

- HTML popup windows created by the `window.createPopup()` method; the appearance of these HTML pop-ups is determined completely by the caller.
- New Internet Explorer frame windows (also referred to as "pop-up windows") created by the `window.open()` method. These new frame windows can either show or not show interface elements (such as a title bar, status bar, address bar, and so on) depending on the `sFeatures` parameter of the `window.open()` call.

The Window Restrictions feature also constrains script-initiated movement of the frame window to prevent repositioning or resizing the frame window in such a manner that key elements are outside the visible display area. This affects the following methods:

- `moveTo`
- `moveBy`
- `resizeTo`
- `resizeBy`

and the following properties:

- `Left`
- `Top`
- `Width`
- `Height`

The Window Restrictions feature also forces the status bar to be displayed on all windows created by the `window.open()` method.

**Who does this feature apply to?**

Web developers should be aware of these new restrictions to plan changes or workarounds for any possible impact to their Web site.

Application developers should review this feature to plan to adopt changes in their applications. This feature is only enabled by default for Internet Explorer processes. Developers must register non-Internet Explorer applications to take advantage of the changes

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Script positioning of Internet Explorer windows**
**Detailed description**
Script-initiated placement of new Internet Explorer frame windows and script-initiated positioning of existing frame windows are constrained to ensure that key security-related interface components (the title and status bars, and address bar if displayed) remain visible after the operation completes.

Scripts cannot position windows so that the title bar or address bar are above the visible top of the display.

Scripts cannot position windows such that the status bar is below the visible bottom of the display.

**Why is this change important?**
Without this change, existing-window movement using the `moveTo` and `moveBy` methods and the `Left` and `Top` properties, and new windows that are created by the `window.open()` method can be called by scripts and used to spoof a user interface or desktop or to hide malicious information or activity by one of the three following methods:

- Positioning the window such that the title bar, status bar, or address bar are off-screen.
- Positioning the window to hide important elements of the user interface from the user.
- Positioning the window so that it is entirely off-screen.

The visible security features of Internet Explorer windows provide information to the user to help the user ascertain the source of the Web page and the security of the communication that uses that page. When these elements are hidden from view, users might think they are on a more trusted page or interacting with a system process when they are actually interfacing with a malicious host. Malicious use of window relocation can present false information to the user, obscure important information, or otherwise "spoof" important elements of the user interface in an attempt to motivate the user to take unsafe actions or to divulge sensitive information.

**What works differently?**
This change places constraints on script-initiated positioning of existing Internet Explorer frame windows and of new frame windows created using the `window.open()` method, to ensure that the title bar and status bar in these windows are always visible to the user. Scripts cannot move a window off-screen, although the user can still move a window off-screen. If you maintain a script that creates off-screen windows in Internet Explorer, you need to change your code.

**How do I resolve these issues?**
If your script creates or moves a window off-screen, you should examine this requirement and alternate ways to accomplish your goal.

**Script sizing of Internet Explorer windows**
**Detailed description**
Script-initiated resize operations on Internet Explorer frame windows are constrained to ensure that the title bar and status bar remain visible after the operation completes.

Scripts cannot resize existing frame windows or create new frame windows in such a manner that the title bar, address bar, or status bar cannot be seen.

When creating a window, the definition of the `fullscreen=yes` specification is changed to mean "show the window as maximized," which will keep the title bar, address bar, and status bar visible.

**Why is this change important?**
Without this change, existing Internet Explorer frame windows can be resized or new frame windows can be created using the `window.open()` method and used to spoof a user interface or desktop or to hide malicious information or activity by sizing the window so that the status bar is not visible.

Internet Explorer windows provide visible security information to the user to help them ascertain the source of the Web page and the security of the communication with that page. When these elements are not in view, the user might think they are on a more trusted page or interacting with a system process when they are actually interacting with a malicious host. Malicious uses of window sizing can obscure important security-related information, and otherwise spoof important elements of the user interface in an attempt to motivate the user to take unsafe actions or to divulge sensitive information

**What works differently?**
With this change, there are constraints on script-initiated resizing operations on existing Internet Explorer frame windows and on the size of the new frame windows created using the `window.open()` method, to ensure that the title bar and status bar of these windows is always visible to the user. The result is that a script cannot open a window in kiosk mode, a mode that does not display the title bar, address bar, and status bar, which present important security information to the user.

The user can choose to display a window in kiosk mode. This election is still persistent.

**How do I resolve these issues?**
Script-initiated windows will be displayed fully, with the Internet Explorer title bar and status bar. The user or the site administrator can manually change this state.

**Script management of Internet Explorer status bar**
**Detailed description**
Internet Explorer has been modified to always display the status bar in Internet Explorer frame windows created using the `window.open()method`.

**Why is this change important?**
Without this change, windows that are created using the `window.open()` method can be called by scripts and spoof a user interface or desktop or hide malicious information or activity by hiding important elements of the user interface from the user.

The status bar is a security feature of Internet Explorer windows that provides Internet Explorer security zone information to the user. This zone cannot be spoofed, and lets the user know exactly what security zone the displayed content is in. When the status bar is hidden from view, the user might think they are on a more trusted page when they are actually interacting with a malicious host.

**What works differently?**
On all windows created by the `window.open()` method, the status bar will be displayed so that the security zone is visible to the user. The '`status=no`' or '`status=0`' specifications in the `sFeatures` parameter of the `window.open()` method are ignored. Application impact depends on the operation carried out on the window as follows:

- `window.open()` method calls will not need to be modified, because the optional width and height values passed in the `sFeatures` parameter specify the size of the content area of the windows and do not include the title bar, status bar, and other window attributes.

- `resizeTo()` method calls might need to be modified because the size parameters of the `resizeTo()` method are for the entire Internet Explorer frame window. Application that have been creating new frame windows with no status bar using the `windows.open()` method and subsequently resizing them using the `resizeTo()` method will need to be modified to account for the fact that the windows now have a status bar.

**Internet Explorer HTML pop-up window placement**
**Detailed description**
HTML pop-up windows are now constrained so that they:

- Do not extend above the top or below the bottom of the content window from which they are created. The "content window" is the top-level DOM window object for the page; visually it is the area where the HTML content is displayed, and extends from the bottom of the lowest displayed interface component at the topic of the Internet Explorer frame window (the title bar, menu, tool bar, or address bar) to the top of the status bar).

- Are not taller in height than the content window.

- Overlap the content window horizontally.

- Appear immediately above the content window, so that other windows (such as a dialog box) cannot be hidden.

- Are automatically repositioned to satisfy the constraints above if the content window moves.

**Why is this change important?**
Pop-up windows are created by the `window.createPopup()` method and are also called chromeless windows because they do not have the border "chrome" components, such as the address bar, title bar, status bar, and toolbars. Without the constraints previously described, these windows:

- Can be opened on top of a dialog box and obscure or replace important elements.
- Can be used to overlay the address bar with a different address.
- Can simulate a full-screen Windows desktop with a password dialog box.

Unrestricted chromeless windows can deceive the user in several ways:

- A chromeless pop-up window that is opened on top of a dialog box can obscure or replace important elements of the dialog box, such as warning text and selection or action controls. (These include check boxes, option buttons, and so on.) This might lead the user to a response that might be inappropriate or harmful.
- A chromeless pop-up window can overlay the address bar with an address that is different from the actual address of the page, which gives the user a false sense of security. In the same way, it can overlay the status notification area, so it might indicate that Internet Explorer is displaying a secure Web page (which displays a URL beginning with https://) Because of this, the user might think that security is in effect for the page when no such security exists.
- A chromeless pop-up can use the entire display. With this method, a malicious user can simulate a full-screen Windows desktop with a password dialog box, with a malicious script that captures the user's private authentication information.

**What works differently?**
HTML pop-up windows are constrained horizontally, vertically, and in order of placement on top of other windows.

An HTML pop-up window must appear between the top and bottom of its parent window's chrome, so it does not overlap the Internet Explorer address bar, title bar, status bar, or toolbars.

Horizontally, an HTML pop-up window must always overlap some area of its parent window.

An HTML pop-up window must stay immediately on top of its parent, so it cannot be placed over other windows.

These constraints might affect the appearance of an HTML pop-up window if it has been designed to display in an area that is larger or separate from its parent window. The HTML pop-up windows might be repositioned and might also be truncated, which might obscure some of the information displayed in that window.

**How do I resolve these issues?**
Redesign the HTML pop-up window to fit into the constraints of this mitigation.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
There is only one setting for this feature. This setting either enables the Windows Restrictions or does not enable them. For application compatibility, this feature is not enabled by default for non-Internet Explorer processes.

**Internet Explorer Windows Restrictions Settings**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| IExplore.exe<br>Explorer.exe<br>Msimn.exe<br>WMPlayer.exe | HKEY_LOCAL_MACHINE (or Current User)\Software \Microsoft \Internet Explorer\Main \FeatureControl \FEATURE_WINDOW_RESTRICTIONS\ | Not applicable | 1 | 0 - Off<br>1 - On |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
The script will call the same methods for the creation of an Internet Explorer window with chrome (using the `window.open()` method) or an Internet Explorer chromeless pop-up window (using the `window.createPopup()` method). However, the design might need to be reviewed to ensure that pop-up windows are appropriately visible to the user and that the status bar contains accurate information.

# Internet Explorer Zone Elevation Blocks

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| The Microsoft Windows Server 2003 Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings that disable scripts, ActiveX components, and file downloads for resources in the Internet security zone. As a result, many of the security enhancements included in the latest release of Internet Explorer will not be as noticeable in Windows Server 2003 Service Pack 1. For example, the new Internet Explorer Notification Bar and Pop-up Blocker features will not be used unless the site is in a zone whose security setting allows scripting. If you are not using the enhanced security configuration on your server, these features will function as they do in Windows XP Service Pack 2. |

**What does Zone Elevation Blocks do?**
When a Web page is opened in Internet Explorer, Internet Explorer puts restrictions on what the page can do, based on where that Web page came from: the Internet, a local intranet server, a trusted site, and so on. For example, pages on the Internet have stricter security restrictions than pages on a user's local intranet. Web pages on a user's computer are in the Local Machine security zone, where they have the fewest security restrictions. This makes the Local Machine security zone a prime target for malicious users. Zone Elevation Blocks makes it harder to get code to run in this zone. In addition, Local Machine Zone Lockdown makes the zone less vulnerable to malicious users by changing its security settings.

**Who does this feature apply to?**
Web developers must plan changes or workarounds for any possible impact to their Web site.

Application developers should review this feature to plan to adopt changes in their applications that run in the Local Machine security zone. Because the feature is not enabled for processes other than Internet Explorer by default, developers must register their applications to take advantage of the changes.

End users might be affected by sites that are not compatible with these stricter rules and settings.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Zone Elevation Blocks**
**Detailed description**
Internet Explorer prevents the overall security context for any link on a page from being higher than the security context of the root URL. This means, for example, that a page in the Internet zone cannot navigate to a page in the Local Intranet zone. A script, for example, could not cause this navigation. For the purpose of this mitigation, the security context ranking of the zones, from highest security context to lowest, is: Restricted Sites zone, Internet zone, Local Intranet zone, Trusted Sites zone, and Local Machine zone.

Zone Elevation Blocks also disables JavaScript navigation if there is no security context.

If a user clicks a link that causes the Web site to attempt to navigate to a higher zone, navigation is blocked for navigation to the Local Machine zone, but a dialog box will appear in Internet Explorer when a Web page attempts to open a page in a security zone that has a higher security context and you will be prompted as in the following message. The italicized portion changes, according to the security zone that the Web page is attempting to navigate to.

- The current Web page is trying to open a site in your *Trusted* sites list. Do you want to allow this?

In any case, the default action does not allow the zone elevation. The user must explicitly allow the requested zone elevation.

**Why is this change important?**
Elevation of privilege is one of the most exploited vulnerabilities in Internet Explorer, with the ultimate goal of running malicious code in the Local Machine zone. Zone Elevation Blocks helps mitigate many privilege escalation attacks.

**What works differently?**
Navigation from one zone to a "higher" zone is blocked. This means that Web pages that automatically call more privileged Web pages fail.

**How do I resolve these issues?**
If you have a trusted Web application that is impacted by this change because it navigates between different security zones without user interaction, you should map the domains that the Web application uses into the security zone with the least

privilege necessary to perform the task for which the application was designed.

# Internet Information Services (IIS) in Windows Server 2003 Service Pack 1

Applies To: Windows Server 2003 with SP1

**What does Internet Information Services do?**
Windows Server 2003 Service Pack 1 includes Internet Information Services (IIS), Version 6.0, which makes it possible for you to host your own Web site on the Internet or your intranet.

IIS is an optional component of Windows Server 2003, is not enabled by default, and must be installed separately.

**Who does this feature apply to?**
This feature applies to the following audiences:

- IT professionals that use IIS to host and administer a Web site.
- Web developers that use IIS to develop Web content.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Metabase auditing**
Beginning with Windows Server 2003 Service Pack 1 (SP1), IIS 6.0 includes a metabase auditing feature that allows tracking of each change that is made to the metabase. Metabase auditing is enabled by enabling an audit access control entry (ACE) on a node in the metabase. After the ACE is enabled, whenever a metabase change takes place on that node, an audit event is published in the Security event log. Using metabase auditing, you can keep track of:

- What was changed (metabase node, property, and old and new values).
- When the change was made (date and time).
- Who made the change (domain and user name).
- Success or failure of the change attempt (HRESULT).
- When a change is made remotely (client IP number).

| **Note** |
| --- |
| To avoid disclosing sensitive information, such as passwords, values of secure properties do not appear in audit event log entries. |

**Request tracing**
The Windows operating system includes the Event Tracing for Windows (ETW) infrastructure to help individuals troubleshoot problems in the operating system, including problems in HTTP components. If an HTTP request in IIS fails or becomes unresponsive while ETW is enabled, you can view ETW trace data, called events, to determine which component caused the failure. IIS includes the following tracing features:

- IIS Currently-executing Requests Tracing: This tracing feature provides general statistics and details about all requests executing on the server at the moment tracing was started.
- IIS Request-Based Tracing: This tracing feature tracks HTTP requests as they move through IIS components.

Windows Server 2003 Service Pack 1 or later also includes a provider for tracing the IIS Admin service during startup and shutdown.

**Kernel-mode SSL.**
You can run Secure Sockets Layer (SSL) in kernel mode, instead of the default user mode. Running in kernel mode means that components or processes run in the core address space of the operating system. Moving encryption and decryption operations

to the kernel improves SSL performance by reducing the number of transitions between kernel mode and user mode. Enabling kernel-mode SSL requires setting a new registry key, `EnableKernelSSL`.

**SSL host headers**
IIS 6.0 now supports using SSL to secure Web sites that use host headers — a security feature that many users want to have. SSL host header support requires obtaining a wildcard server certificate and specifying the SSL port number on the SecureBindings metabase property.

**Running 32-bit applications on 64-bit Windows**
Windows Server 2003 Service Pack 1 enables IIS 6.0 to run 32-bit Web applications on 64-bit Windows using the Windows-32-on-Windows-64 (WOW64) compatibility layer. IIS 6.0 using WOW64 is intended to run 32-bit personal productivity applications needed by software developers and administrators, including 32-bit IIS Web applications.

**W3C centralized logging**
World Wide Web Consortium (W3C) centralized logging is a global configuration on the server where all Web sites write data to a single log file. Data is stored in the log file using the W3C Extended log file format. The log file can be viewed in a text editor, unlike IIS Centralized Binary Logging, which writes data in binary format and requires a parsing tool to view the data.

**Secure configuration for Web servers**
Windows Server 2003 Service Pack 1 includes a Security Configuration Wizard (SCW), which is a role-based tool you can use to create a policy that enables the services, inbound ports, and settings required for a selected server to perform a specific role. If you select the Web Server role in the wizard, SCW configures IIS 6.0 to help further reduce the attack surface of your Web server.

# Kernel patch protection for x64-based operating systems

Applies To: Windows Server 2003 with SP1

**What does kernel patch protection do?**
Kernel patch protection prohibits kernel-mode drivers that extend or replace kernel services through undocumented means. This feature describes changes in policy related to patching the kernel for Microsoft Windows Server 2003 Service Pack 1 x64-based operating systems.

**Who does this feature apply to?**
This feature primarily applies to driver developers, but IT professionals may also find this information useful.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Patching policy for x64-based systems**
**Detailed description**
Microsoft Windows Server 2003 SP1 and later versions of Windows for x64-based systems do not allow the kernel to be patched except through authorized Microsoft-originated hotfixes.

**Why is this change important?**
Kernel-mode drivers that extend or replace kernel services through undocumented means (such as hooking the system service tables) can interfere with other software and affect the stability of the operating system. For x86-based systems, Microsoft discourages such practices but does not prevent them programmatically, because doing so would break compatibility for a significant amount of released software. A similar base of released software does not exist for x64-based systems, so it is possible to add this level of protection to the kernel without breaking compatibility.

**What works differently?**
Many system structures are protected on x64-based systems, including the system service dispatch tables, the interrupt descriptor table (IDT), and the global descriptor table (GDT). The operating system also does not allow third-party software to allocate memory "on the side" and use it as a kernel stack. If the operating system detects one of these modifications or any other unauthorized patch, it will generate a bug check with the stop code 0x109 and shut down the system.

**How do I fix these issues?**
For compatibility with Windows for x64-based systems, drivers cannot modify the kernel. The following actions are blocked in Windows Server 2003 SP1 for x64-based systems:

- Modify system services tables, for example, by hooking the KeServiceDescriptor table
- Modify the IDT
- Modify the GDT
- Use kernel stacks that are not allocated by the kernel
- Patch any part of the kernel (detected on AMD64-based systems only)

Kernel patch protection might be extended in future system updates to protect against additional malicious patching techniques as new vulnerabilities are detected. To avoid compatibility issues with these updates, drivers should not attempt to update the kernel using other mechanisms.

Drivers for other platforms should also avoid patching the kernel to help ensure stability and reliability of the operating system and a better experience for users.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
Windows Server 2003 SP1 for x64-based systems is a new platform that requires new drivers. These new drivers must conform to x64 patch policy documented on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=38447.

For additional information see The Microsoft Windows Driver Development Kit (DDK) on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=38448.

# Outlook Express

Applies To: Windows Server 2003 with SP1

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Plain text mode**
**Detailed description**
The plain text mode feature of Outlook Express provides users with the option to render incoming mail messages in plain text instead of Hypertext Markup Language (HTML). When Outlook Express is running in plain text mode, the rich edit control is used instead of the MSHTML control. You avoid some security issues that result from the use of MSHTML by using the rich edit control.

**Why is this change important?**
The use of the rich edit control provides an additional barrier to malicious code that is transmitted using e-mail. Computers running earlier versions of Windows XP had a vulnerability to malicious code because Outlook Express processes HTML header scripts in the HTML content. The MSHTML control automatically executes these scripts. The rich edit control does not execute HTML scripts, so this is mitigated. Because plain text e-mail does not require HTML header processing to be displayed properly, there is usually little visible difference from this processing change in standard message formats. Portions of e-mail messages that do not appear to render correctly are relying on HTML rendering and could present a danger to your system.

**What works differently?**
The following Outlook Express features are not available when running in plain text mode:

- Changing text size to a larger or smaller font.
- Full text searching through the body of a mail message.

You can configure plain text mode in several ways, including:

- Reading a message.

  In Outlook Express, on the **Tools** menu, click **Options**, and then click the **Read** tab. Select the **Read all messages in plain text** check box.

- Composing a message.

  In Outlook Express, on the **Tools** menu, click **Options**, and then click the **Send** tab. Under **Mail Sending Format**, select the **Plain Text** option.

- With a new menu option.

  On the **View** menu, click **Message in HTML**.

  This new menu item switches the current message view to HTML if it is currently in plain text view, both in the preview display as well as in the full message display.

**How do I resolve these issues?**
If you are sure that the source of an e-mail message can be trusted and you want to use the full feature set that is provided with the MSHTML control to support rich HTML e-mail for reading or composing, you can switch to the HTML mode by using the **View** menu option procedure as described above in "With a new menu option."

**Limit external HTML content downloads**
**Detailed description**
This Outlook Express feature helps users to avoid getting repeated spam mailings by preventing the user from unknowingly validating his or her e-mail address to spam originators. Businesses that use spam as a marketing technique typically include references to images that reside on their Web servers inside the e-mail message. Some of these spam e-mail messages contain single pixel images that are not visible to the recipient of the e-mail so that the recipient will not be aware that there is any content that is malicious. When the user opened an e-mail that contains the image, previous versions of Outlook Express automatically contacted the Web server to download and display the images. When the request for the image was made to the Web server, it could ascertain that a spam e-mail message was received by an active e-mail account, which validated the e-mail

address in the spam originator's mailing list. Now, when the **Block images and other external content in HTML e-mail** setting is enabled, the default behavior of Outlook Express changes so that it does not contact the Web server to download external content, which helps prevent the verification of the e-mail address with the spam originator. This download behavior is configurable and is enabled by default when you install Windows Server 2003 Service Pack 1.

This feature also helps to minimize a common problem that is experienced by people whose computers use dial-up network connections. Prior to implementing this feature, if users downloaded mail messages and then disconnected their network connection, when they attempted to view an HTML messages that included pictures or other external Internet content, their modem would automatically start to dial out to download the external content.

**Why is this change important?**
This feature increases the privacy that is provided to users of Outlook Express. Their e-mail address is not automatically validated by the Web server of spam originators without their knowledge when a spam e-mail message is opened. Using this feature may result in the following advantages:

- The user receives less spam.
- The user is less distracted by the receipt of spam.
- Automatic attempts by a user's modem to reconnect to the Internet after receiving HTML e-mail decrease.

**What works differently?**
Implementing this feature in Outlook Express helps prevent the rendering of pictures in HTML e-mail if the pictures must be retrieved from servers that are in either the Internet or Restricted Web content zones. This new default behavior results in the user's name not being validated by the Web site hosting the pictures, which makes the user's e-mail name less useful to spam senders. This may result in the user getting less spam over time.

To communicate that these pictures are missing, there is now an External Message Information Bar that is displayed in both the Outlook Express message window as well as in the preview area. This External Message Information Bar appears whenever the message contains references to external Internet content, such as images or script and the options are set to render the message in HTML.

When Outlook Express blocks content, the actual image is replaced with the standard placeholder for the blocked image in the text of the mail message. Images are the only blocked items that provide a visual cue that something is not being displayed. For sounds, IFrames, and other content, there is no visual indication in the body of the mail message. When users print an HTML e-mail that has blocked content, Outlook Express prints the e-mail exactly as it appears on the screen, with a placeholder for the blocked images. The external content is not downloaded.

An added benefit of this feature is that it minimizes a common dial-up user problem: undesired automatic dial-up network connection attempts. When viewing an HTML e-mail message off-line, previous versions of Outlook Express would automatically dial out to connect to the Internet and retrieve any reference images. However, because almost all external HTML references in e-mail messages point to resources on the Internet that are part of the Internet zone, the content is not displayed by default and a dial-up network connection is not requested.

**How do I resolve these issues?**
To turn off all external content blocking, on the **Tools** menu, point to **Options**, and then click **Security**. Clear the **Block images and other external content in HTML e-mail** check box. From that point, no content is blocked, which returns Outlook Express to the prior behavior of automatically downloading external content.

To explicitly download external content for an e-mail message, click the External Message Information Bar to download the external content that was included with the message.

**Attachment Manager API integration**
**Description**
Outlook Express now integrates a new set of application programming interfaces (APIs), called the Attachment Manager, to check e-mail attachments. This allows applications to eliminate custom code that performs similar safety checks and instead rely on this centrally-managed API set. The use of Attachment Manager provides a consistent user experience across all applications that check the security of an attachment.

**Why is this change important?**
It is important to have a more unified approach for attachment security across all Windows applications. This helps to ensure that users get a consistent experience with regard to the security check performed on attachments.

**What works differently?**

Apart from the consistent user experience, this feature does not provide any visible change to the user.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
There are several differences in functionality that a developer should be aware of.

When API names are provided, they are the Attachment Manager API. If the **Do not allow attachments to be saved or opened that could potentially be a virus** setting is disabled, Outlook Express calls `SetReferer()` and passes http://URL as a parameter. This is done so that the subsequent call to `CheckPolicy()` considers Outlook Express to be in the Internet Web content zone. Attachment Manager discriminates differently, depending on whether the caller is in the context of the Internet or Restricted security zones. The following sections provide overviews of different behaviors that the Attachment Manager API supports.

**Behavior when previewing a message that includes an attachment**
Before the preview area is rendered, `CheckPolicy()` is called to determine the state of the menu options associated with the attachment icon in the preview area header, and the corresponding actions as follows:

- If `CheckPolicy()` returns `E_Fail` (dangerous attachment), `S_OK`, or `S_False` (safe attachment), there is no change to the previous functionality of Outlook Express.
- Opening the attachment saves the file as a temporary file and then calls `Execute()` to execute the file instead of the currently used `ShellExecute()` call.
- If `Execute()` fails, subsequent user actions are handled by Attachment Manager.
- When the **Save Attachments** dialog box is opened, the list of attachments contains items that are enabled in the menu. Blocked attachments do not appear in the **Save Attachments** dialog box. When the user selects the destination folder and clicks **Save**, Outlook Express saves the files to the specified folder and then calls `Save()` on each of the saved files.

In the case of previewing mail with multiple attachments, `CheckPolicy()` is called on each of the attachments. Depending on whether the return value is `E_Fail`, or `S_OK`, or `S_False`, Outlook Express either disables or enables the attachment name in the menu.

In future implementations, `Save()` could fail if `CheckPolicy()` does not return `S_OK`. In this case, Outlook Express will display the error message "The following attachments were not saved because they could not be verified as being safe", followed by a list of failed files.

**Behavior when reading a message that includes an attachment**
Before the Outlook Express message window is rendered, `CheckPolicy()` is called for every attachment to determine which attachments are shown and which are blocked from access to the user.

- If `CheckPolicy()` returns `E_Fail` (dangerous attachment), `S_OK` or `S_False` (safe attachment), Outlook Express behaves just as it did in the past. Double-clicking the attachments that are displayed in the Attach area of the message window follows the exact same steps as described when executing attachments from the preview area.
- When the user clicks **Save As**, selects the destination folder and file name, and then clicks **Save**, Outlook Express saves the attachment in the specified folder and then calls `Save()` to sync.
- Selecting **Print** is similar to running the attachment, except that, instead of calling `Execute()` without any parameters, Outlook Express issues a call to `Execute("print")`. All other tasks, such as saving the file to a temporary file remain the same as when executing the attachment.
- If the **Do not allow attachments to be saved or opened that could potentially be a virus** setting is disabled, Outlook Express calls `SetReferer()` and passes http://URL as a parameter. The subsequent call to `CheckPolicy()` then considers Outlook Express to be in the Internet Web content zone.

**Behavior when moving a message that includes an attachment**
If the user moves an item to a location outside Outlook Express — for example, dragging a message containing an attachment to the desktop — Outlook Express performs these actions:

- Generates a temporary file with `HDROP`.

- Saves a temporary file
- Calls `Save()` on the temporary file

If it is successful, `HDROP` is made available

If it fails, `HDROP` is not made available and the drop target is disabled.

# Post-Setup Security Updates

Applies To: Windows Server 2003 with SP1

**What does Post-Setup Security Updates do?**
Post-Setup Security Updates is designed to help protect a new server installation from risk of infection between the time the server is first connected to the network and the application of the most recent security updates from Windows Update.

Post-Setup Security Updates is a user interface that appears the first time an administrator logs onto the new server and provides links for you to apply updates to your server and to configure automatic updates. Post-Setup Security Updates also informs the administrator that all inbound connections other than those specifically opened during setup or by policy settings, were blocked. If the administrator set exceptions to the firewall through Group Policy or by an unattended setup script, inbound connections assigned to these exceptions remain open.

Post-Setup Security Updates is not available from the **Start** menu and is only available under specific conditions as described later in this document.

| **Note** |
| --- |
| Post-Setup Security Updates does not appear when the server is being upgraded from the following operating systems:<br><br>- Windows NT Server 4.0 to Windows Server 2003 with Service Pack 1<br><br>- Windows 2000 Server to Windows Server 2003 with Service Pack 1<br><br>- Windows Server 2003 to Windows Server 2003 with Service Pack 1 |

**Who does this feature apply to?**
Post-Setup Security Updates applies to Windows server administrators who are performing a full installation of Windows Server 2003 that includes Service Pack 1 or later (such as a slip-stream version of Windows Server 2003 with Service Pack 1). This feature does not apply if either of the following statements is true:

- Windows Firewall is enabled or disabled using an unattended-setup script for operating system installation.
- Windows Firewall is enabled or disabled by application of Group Policy before Post-Setup Security Updates is displayed.

This feature does not apply if the administrator is updating an existing Windows Server 2003 operating system by adding a service pack or if the administrator is upgrading an existing Windows 2000 Server operating system to Windows Server 2003 with Service Pack 1.

**Why is this change important?**
Security updates that mitigate virus threats may have been released by Microsoft since the release of the operating system files being installed. If the new server is connected to the network and a firewall is not enabled, the server may be infected with a virus before the security updates can be downloaded and installed. Post-Setup Security Updates uses the Windows Firewall to mitigate this risk.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
Post-Setup Security Updates is a new feature in Windows Server 2003 Service Pack 1.

**Post-Setup Security Updates**
**Detailed description**
If Windows Server 2003 with Service Pack 1 or later is installed as a new installation and Windows Firewall is not explicitly enabled or disabled using an unattended-setup script during the installation or by application of Group Policy, Windows Firewall will be enabled by default on first startup and logon in order to allow the administrator to securely download and install updates from Windows Update, and the Windows Server Post-Setup Security Updates screen will be shown. The Post-Setup Security Updates screen informs you that all inbound connections other than those specifically opened during setup or by policy settings, were blocked.

Windows Firewall blocks all inbound connections with the following exceptions:

- If Remote Desktop was enabled using an unattended-setup script during installation, port 3389 is not blocked.

- If Group Policy is used to apply policy settings that do not enable or disable Windows Firewall, but define exceptions to the firewall, exceptions defined by the policy settings are not blocked.

Post-Setup Security Updates offers links to Windows Update to allow you to download any security updates released since this operating system version was released and, if you have not already done so, provides the opportunity for you to configure Automatic Updates to help protect this server in the future.

**What happens when Post-Setup Security Updates is closed?**
If Windows Update or any other configuration change causes a restart before you click the **Finish** button on Post-Setup Security Updates, it reopens the next time an administrator logs on to the server.

If you close Post-Setup Security Updates using ALT+F4 or Task Manager, no change is made to the configuration of Windows Firewall. The tests the server uses to determine whether Post-Setup Security Updates should be displayed run again the next time a user logs on.

When you click the **Finish** button on the **Post-Setup Security Updates** dialog box, a dialog box explaining the consequences of closing Post-Setup Security Updates is displayed. In order to provide correct information, the following steps are taken to determine the current status of Windows Firewall:

- If you made no changes to the Windows Firewall configuration since Post-Setup Security Updates appeared, a confirmation dialog box appears explaining that inbound connections will now be opened and giving you the opportunity to confirm that you are done with any post-setup security updates. When the action is confirmed, Post-Setup Security Updates attempts to disable Windows Firewall and stop and disable the Windows Firewall/Internet Connection Sharing service.

- If Windows Firewall is disabled successfully, a registry value is set to suppress Post-Setup Security Updates in the future. It is possible that Windows Firewall is disabled successfully, but the attempt to stop the Windows Firewall/Internet Connection Sharing service fails.

- If Windows Firewall settings cannot be changed, a dialog box appears explaining that no changes will be made to inbound connection settings. Post-Setup Security Updates is not suppressed and the tests to determine whether Post-Setup Security Updates should be displayed will be run again the next time a user logs on.

- If Windows Firewall was explicitly enabled or disabled since Post-Setup Security Updates appeared, a dialog box appears explaining that no changes will be made to inbound connection settings. These changes could have been made by the application of Group Policy settings or by opening the Windows Firewall control panel and clicking **OK** to confirm the firewall settings. A registry value is set to suppress Post-Setup Security Updates in the future.

- If the Windows Firewall/Internet Connection Sharing service was stopped or disabled since Post-Setup Security Updates appeared, a dialog box appears explaining that no changes will be made to inbound connection settings. A registry value is set to suppress Post-Setup Security Updates in the future.

- If Internet Connection Sharing was enabled since Post-Setup Security Updated appeared, a confirmation dialog box appears explaining that inbound connections will now be opened and giving you the opportunity to confirm that you are done with any post-setup security updates. When the action is confirmed, Post-Setup Security Updates attempts to disable Windows Firewall. The service shared between Windows Firewall and Internet Connection Sharing is not turned off.

- If the state of the firewall cannot be determined, a dialog box appears explaining that no changes will be made to inbound connection settings. Post-Setup Security Updates is not suppressed and the tests to determine whether Post-Setup Security Updates should be displayed will be run again the next time a user logs on.

| **Note** |
| --- |
| The text on Post-Setup Security Updates is not refreshed if the firewall status changes after the initial display. If the status of the firewall changes after it appears and before the **Finish** button is clicked, the text may state that all inbound connections are blocked when, in fact, they are not. When you click **Finish**, Post-Setup Security Updates checks the status of the firewall again before displaying a dialog box explaining any changes to be made on closure. |

**When will the Post-Setup Security Updates screen be displayed?**
Because this feature runs automatically and cannot be started on request, you can use the following information to determine whether your server will display the Post-Setup Security Updates feature.

The following tests are run to determine whether or not to display Post-Setup Security Updates.

| Test | Positive Result | Negative result |
|---|---|---|
| Is the logged-on user an administrator? | Continue on to the next test | Skip the remaining tests and do not display Post-Setup Security Updates. These tests run again the next time a user logs on. |
| Is this is a new installation of a version of Windows Server 2003 that includes Service Pack 1 or later (not an upgrade) | Continue on to the next test | Skip the remaining tests and do not display Post-Setup Security Updates. The registry value is set to suppress Post-Setup Security Updates in the future. |
| Has Post-Setup Security Updates been suppressed in the registry? | Skip the remaining tests and do not display Post-Setup Security Updates | Continue on to the next test |
| Is the Windows Firewall/Internet Connection Sharing service running? | Continue on to the next test. | Repeat this test for two minutes. If the service has still not started, do not display Post-Setup Security Updates. These tests are run again the next time a user logs on. |
| Has Windows Firewall been explicitly enabled or disabled for the current Windows Firewall profile?<br><br>(The firewall may have been enabled or disabled using an unattended-setup script at the time of installation or through the application of Group Policy settings or by opening the Windows Firewall control panel and clicking OK to confirm the firewall settings.) | Skip the remaining tests and do not display Post-Setup Security Updates. A registry value is set to suppress Post-Setup Security Updates in the future. | If Windows Firewall is enabled and the user did not enable it, display Post-Setup Security Updates.<br><br>If the status of Windows Firewall cannot be determined do not display Post-Setup Security Updates. These tests are run again the next time a user logs on. |

**What works differently?**
Manage Your Server is not automatically displayed until Post-Setup Security Updates closes.

Post-Setup Security Updates does not cause any applications to work differently.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
Windows Firewall (previously known as Internet Connection Firewall) was not enabled by default at the end of a new installation unless the administrator enabled it using an unattended-setup script. Under the circumstances described earlier in this document, Windows Firewall is now enabled automatically until Post-Setup Security Updates is finished.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
No new policy settings were created relating to Post-Setup Security Updates. The following value in the registry was added. This key does not affect firewall settings.

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| DontLaunchSecurityOOBE<br><br>(DWORD) | HKEY_LOCAL_MACHINE \SOFTWARE \Microsoft \Windows \Current Version \ServerOOBE \SecurityOOBE | N/A | This key does not exist by default. | The key can exist or not exist. If the key exists, Post-Setup Security Updates does not display. The numerical value of this setting is irrelevant. |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**

If you do new installations of a version of Windows Server 2003 that includes a service pack by using an unattended-setup script and you want to suppress Post-Setup Security Updates, it is recommended that you explicitly enable or disable Windows Firewall in either your setup script or by Group Policy. This change automatically suppresses Post-Setup Security Updates.

# Remote Access Quarantine

Applies To: Windows Server 2003 with SP1

**What does Remote Access Quarantine do?**
Remote Access Quarantine control provides network administrators the ability to validate the configuration of remote client computers before they are permitted to access the corporate network. Typical remote access connections only validate the credentials of the remote access user. Therefore, the computer used to connect to a private network can often access network resources even when its configuration does not comply with organization network policy. For example, a remote access user with valid credentials could connect to a network with a computer that does not have the following:

- The correct service pack or the latest security patches installed.
- The correct antivirus software and signature files installed.
- Routing disabled. A remote access client computer with routing enabled might pose a security risk, providing an opportunity for a malicious user to access corporate network resources through the client computer, which has an authenticated connection to the private network.
- Firewall software installed and active on the Internet interface.
- A password-protected screensaver with an adequate wait time.

Despite the efforts made within organizations to ensure that computers used internally comply with network policy, those used from employees' homes for remote access can still present significant risk to the network.

Remote Access Quarantine, a new feature in Windows Server 2003 Service Pack 1, delays normal remote access to a private network until the configuration of the remote access computer has been examined and validated by an administrator-configured script (included in the connection settings). When a remote access computer initiates a connection to a remote access server, the user is authenticated and the remote access computer is assigned an IP address. However, the connection is placed in quarantine mode, in which network access is limited. The administrator-configured script is run on the remote access computer. When the script notifies the remote access server that it has successfully run and the remote access computer complies with current network policies, quarantine mode access restrictions are removed and the remote access computer is granted normal remote access.

The quarantine restrictions placed on individual remote access connections can consist of the following:

- A set of quarantine packet filters that restrict the traffic that can be sent to and from a quarantined remote access client.
- A quarantine session timer that restricts the amount of time the client can remain connected in quarantine mode before being disconnected.

You can use either restriction, or both, as needed. The administrator also has the option to help the client to remedy the configuration, by say updating the signature file for the antivirus software, through the validation script.

The components required for this Remote Access Quarantine solution include the following:

- The Remote Access Quarantine Service (RQS or Listener) to be run on the routing and remote access (RRAS) server that listens for requests from the remote clients for removal of quarantine restrictions.
- A RADIUS server (or the RRAS server itself) where a quarantine policy can be defined for applying IP filters or session timeouts to remote connections.
- A configuration validation script that performs the validation checks to verify that the remote access client computer conforms to the minimum security guidelines required to access the corporate network.
- A Connection Manager profile configured to run the Remote Access Quarantine Client (RQC) as a Post-Connect action on the remote client computer. The Quarantine CM profile will update the validation scripts from an administrator-specified

share path and run the validation scripts. If the minimum requirements are verified by the scripts to have been met, RQC will notify RQS and request removal of the quarantine restrictions.

- Remote access clients configured to run the Remote Access Quarantine Client and the validation script (distributed through the Connection Manager Profile).

| ⚠Caution |
| --- |
| Remote Access Quarantine is not a security solution. It is designed to help prevent computers with unsafe configurations from connecting to a private network; not to protect a private network from malicious users who have obtained a valid set of credentials. |

**Who does this feature apply to?**
This feature applies to:

- Remote access servers, running Windows Server 2003 with Service Pack 1.
- Remote access client computers connecting to the corporate network from remote locations, running Windows 2000 or Windows XP.
- Network Administrators who want to validate configuration of client computers before they are allowed access to the corporate network.

**Why is this change important?**
Remote Access Quarantine provides network administrators with a mechanism to quarantine remote access clients by providing VPN access to limited parts of the private network, and allowing administrators to validate that the computer meets the minimum security requirements. After the computers have been verified to meet the guidelines for accessing the network, quarantine restrictions can be lifted allowing the client computers to have normal access to the network resources.

This mitigates the threat to a private or corporate network from vulnerable client computers that are at remote locations or are not domain-joined and thus outside the administrator's purview.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**

| Setting name | Location | Previous default value | Default value | Possible values |
| --- | --- | --- | --- | --- |
| AllowedSet, REG_MULTI_SZ | HKEY_LOCAL_MACHINE \System \CurrentControlSet \Services\Rqs | N/A | RASQuarantineConfigPassed | N/A |
| Port, REG_DWORD | HKEY_LOCAL_MACHINE \System \CurrentControlSet \Services\Rqs | N/A | Not set<br><br>When this key is not set the server listens for client notification on port 7250 | port number to listen to |
| Authenticator, REG_SZ | HKEY_LOCAL_MACHINE \System \CurrentControlSet \Services\Rqs | N/A | Not set | NULL |
| Verifier, REG_SZ | HKEY_LOCAL_MACHINE \System \CurrentControlSet \Services\Rqs | N/A | Not set | NULL |

# Remote Installation Services (RIS) for 64-bit Operating Systems

Applies To: Windows Server 2003 with SP1

**What does Remote Installation Services (RIS) do?**
The Remote Installation Services feature simplifies the task of installing an operating system on computers throughout an organization. It provides a mechanism for computers to connect to a network server during the initial boot process, while the server controls a local installation of Windows. This can be used either to install the correct configuration of the operating system on a new computer, or to restore a failed computer to a known operating system configuration.

**Who does this feature apply to?**
This feature applies to IT professionals who want to use RIS to deploy x64-based operating system images in an enterprise.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1**
**Support for x64 operating system images**
**Detailed description**
Remote Installation Services in Windows Server 2003 supported the installation of Windows Server 2003 Itanium-based editions using RISETUP. With Windows Server 2003 SP1, Remote Installation Services now supports the deployment of both RISETUP and RIPREP images for the x64 architecture.

There are two methods for adding operating system images to a Remote Installation Services (RIS) computer – RISETUP and RIPREP. RISETUP images are images built from the Windows CD media. When adding the images, a path to the Windows setup files, either on CD or on a network share, must be specified. RISETUP images are generally added by one of three methods:

- Using the **Remote Installation Services Setup** option from the **Start** menu of the RIS server. To start **Remote Installation Services Setup**, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Add a new OS image to this remote installation server**.

- Using the **Add a new installation image** option available from the **Properties** tab of the RIS server in the Active Directory Users and Computers Microsoft Management Console (MMC) snap-in.

- Running the `RISETUP /add` command from a command prompt on the RIS server.

RIPREP-generated images are also file system based images located on a RIS server. However, unlike RISETUP images, RIPREP images more closely resemble the hard disk images created using a third-party disk-imaging tool and the Windows System Preparation tool (Sysprep). RIPREP images are created by running the RIPREP wizard (Riprep.exe) on a master computer that has the operating system configuration, applications and settings, and desktop customizations that are to be captured and deployed to client computers in the organization.

With the changes in Windows Server 2003 SP1, RISETUP and RIPREP images created for x64-based operating systems can be deployed from a RIS server of any architecture – be that a 32-bit operating system, an x64-based operating system, or an Itanium-based operating system.

**Why is this change important?**
This change enables the rapid deployment of x64-based versions of Windows Server 2003 and Windows XP throughout an enterprise by leveraging Remote Installation Services (RIS) as the image deployment mechanism.

**What works differently?**
In order to support x64-based RIPREP and RISETUP images, several pieces of functionality were added to Remote Installation Services (RIS).

**RISETUP**
Beginning with Windows Server 2003 SP1, RISETUP now allows for the addition of x64-based images to the RIS server. Prior to Windows Server 2003 SP1, trying to add a RISETUP x64-based image would result in the following error message:

**A file that is needed to set up the installation image on the server was not found.**

**This may indicate that the image source is corrupt or that the source is not a valid Windows installation source.**

**Verify that the path you entered points to a valid Windows installation source.**

This message occurred because prior to Windows Server 2003 SP1 RISETUP did not include the x64-based image structure when examining the source media. RISETUP would find an i386 directory in the CD image source; however, RISETUP would not

find all of the files defined for the source. Therefore, the program determined that the image source was missing files and displayed the appropriate error condition.

Changes to RISETUP were also necessary to populate the *****RIS_Server***\RemInst\Admin\amd64** folder with the x64-based version of RIPREP. The x64-based version of RIPREP only ships on x64-based editions of Windows (both client and server). However, it is possible to deploy x64-based images from x86-based and Itanium-based Windows Server 2003 SP1 RIS servers. In these cases, the x64-based RIPREP files would not natively be on the x86-based and Itanium-based servers.

Each RIPREP image requires a supporting RISETUP image. So, before a RIPREP image of an x64-based version of Windows Server 2003 SP1 can be made, a RISETUP image of an x64-based version of Windows Server 2003 SP1 must first be loaded to the RIS server. On the initial RISETUP load of the x64-based image, the necessary 64-bit files for RIPREP – Riprep.exe, Riprep.inf, and Setupcl.exe – are copied from the x64-based RISETUP image location (either flat image or CD) and placed in the **\RemoteInstall\Admin\amd64** folder. In addition, a file version checking mechanism has been put in place to ensure that the most up-to-date versions of Riprep.exe, Riprep.inf, and Setupcl.exe are in place on the RIS server. If the file versions in the copied x64-based RISETUP image are newer than the existing x64-based versions on the RIS server, the RIS server versions will be updated to contain the latest available versions of the files. This is the desired behavior since Riprep.exe, Riprep.inf, and Setupcl.exe are backwards compatible. If the existing x64-based versions in the copied x64-based RISETUP image are the same as or older than the existing x64-based version on the RIS server, the files on the RIS server will remain untouched.

**Client Installation Wizard**

Remote Installation Services (RIS) provides a set of Client Installation Wizard (CIW) screens that prompt a user through the installation of an operating system. These screens are maintained as simple text files in an OS Chooser Markup Language (OSCML) format and stored with an .osc extension.

You can customize and tailor these screens to include Help Desk phone numbers, additional contact information, or any other additional information using Hypertext Markup Language (HTML) 2.0. The Client Installation Wizard screens are kept in the **RemoteInstall\OSChooser\Language** directory.

Windows Server 2003 SP1 includes changes to the Client Installation Wizard (CIW) screens to allow for the selection of 32-bit or 64-bit operating systems when both are available as installation options.

A new .osc screen, x8664.osc, was added. This new screen, in conjunction with its associated registry settings, allows the system administrator to restrict or filter operating system selection based on architecture for x64-based RIS clients.

With the x64 architecture, computers can run either 32-bit or 64-bit operating systems. To assist in the selection process, the new .osc screen, x8664.osc, is placed in the **RemoteInstall\OSChooser\Language** directory upon copy of the first Windows Server 2003 SP1 RISETUP image. If the Windows Server 2003 SP1 RISETUP image is the first image to be uploaded to the RIS server, the new x8664.osc screen will be copied by default. If there are existing images on the RIS server at the time the first Windows Server 2003 SP1 image is added, the installer must choose the option to overwrite (or backup and overwrite) the .osc screens during the installation to get the new screen.

| **Note** |
| --- |
| If custom actions or automations exist in the .osc screen flow, choosing to overwrite the .osc screens upon RISETUP of an image may cause those customizations to be lost. |

The x8664.osc screen is used to determine whether the operating system choices displayed to the end-user will be 32-bit x86-based or 64-bit x64-based when the booting client computer is x64-based. The behavior is controlled using manual selection of the appropriate installation choice on x8664.osc. Alternatively, a new registry setting, `DefaultPlatformforX8664`, found at `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\BINLSVC\Parameters`, may be used to force all x64-based clients to see only one architecture of images, x86-based or x64-based. For example, if the registry value is set to `i386`, only i386 operating system images will be displayed in OSChoice.osc. If the registry value is set to `amd64`, only x64-based operating system images will be displayed in OSChoice.osc. If the registry value is not set, then the default behavior will cause the x8664.osc screen to prompt for user input. At that point the user can choose one of two available choices:

- **Install Microsoft Windows 32-bit Edition**
- **Install Microsoft Windows 64-bit Edition**

Choosing the first option limits installation image choices to only 32-bit operating systems. If only one 32-bit operating system image exists, the CIW will proceed directly on to warning.osc. Choosing the second option limits installation image choices to only 64-bit operating systems. If only one 64-bit operating system image exists, the CIW will proceed directly on to warning.osc.

Some other points of clarification regarding the new x8664.osc screen behavior:

- The x8664.osc screen will only be seen by x64-based client computers. Computers running x86-based and Itanium-based operating systems will not see the new screen.

- If no Windows Server 2003 SP1 RISETUP images (or Windows XP Professional x64 Edition images) exist on the RIS server, the new behavior will not be in place. This is because the new x8664.osc screen is not copied to the server until the first upload of a Windows Server 2003 SP1 (or Windows XP Professional x64 Edition) RISETUP image.

- If the new x8664.osc screen was not copied on upload of the first x64 RISETUP image (for example, if the Administrator chose to keep the current .osc screens and not copy the new .osc screens on x64-based RISETUP image creation), the new functionality will not be in place. To obtain the new functionality to restrict x64-based RIS client computers to only view x86-based 32-bit or x64-based 64-bit operating system images, the new .osc screens must be in place and the registry value must be configured appropriately.

- A third option, **Install Default Windows**, can be seen in the x8664.osc screen source. This option allows the x8664.osc screen to be automatically processed with no user interaction. This option will be used in the following situations:

    - If the booting client computer is not x64-based. If the client computer is x86-based or Itanium-based, installation using the Client Installation Wizard will continue automatically and the user will not be prompted for input in x8664.osc.

    - If a Windows Server 2003 SP1 x86-based or Itanium-based RISETUP image is placed on a RIS server that is running Windows Server 2003 without Service Pack 1. In this case, it is possible that the new x8664.osc will be placed on a Windows Server 2003 RIS server if the option to replace .osc screens was chosen upon image upload. Since the RIS server has not been updated it does not have the ability to process the new **Install 32-bit Windows** and **Install 64-bit Windows** options in x8664.osc. Since the default option is the only menu choice populated, it is processed and the screen is bypassed.

**Maintenance and Troubleshooting menu**

64-bit maintenance and troubleshooting tools can be provided to users through the Client Installation Wizard (CIW). To do this, the existing tools.osc screen was repurposed and a new CIW screen, Tlchoice.osc, was implemented.

There are two paths that can be taken when choosing an operating system image to deploy from within the Client Installation Wizard. The first path involves selection of either **Automatic Setup** or **Custom Setup** from the Choice.osc screen. When choosing either of these options, a computer account, if one does not already exist, will be created in Active Directory. In the former case, Remote Installation Services creates the account on behalf of the user based on policy as stored in Active Directory. In the latter case, the user may choose both the computer name and location (for example, Organizational Unit) when creating the computer account. The second path involves selection of the **Maintenance and Troubleshooting** menu option from the Choice.osc screen. This installation method differs from the first in that no computer account will be created. The most common use of starting an operating system from the **Maintenance and Troubleshooting** menu option is starting the Windows Preinstallation Environment (Windows PE) to aid in operating system deployment. Because Windows PE is usually an "interim" operating system – used to start a computer temporarily before applying an operating system image – creating a computer account in Active Directory at this point is not useful.

There are two ways that Windows PE images can be made available from the **Maintenance and Troubleshooting** menu option on the Choice.osc screen:

- Setting `'Imagetype=WinPE'` in the `[OSChooser]` section of Ristndrd.sif answer file.

- Creating a \Tools folder at the same level as \Images (for example, \RemoteInstall\Setup*Language*\Tools) and placing the folder structure containing the Windows PE image into the \Tools folder.

X64-based computers can run either 32-bit or 64-bit maintenance and troubleshooting tools. To assist the user in the selection process, the Tools.osc screen has been updated to present the user a choice of two options:

- **Display 32-bit Tools**

- **Display 64-bit Tools**

Choosing the first option displays only 32-bit tools in the new Tlchoice.osc screen. Selecting the second option displays only 64-bit tools in the Tlchoice.osc screen. Functionality that was previously in the Tools.osc screen has been moved to the Tlchoice.osc screen.

Some points of clarification regarding the new Tools.osc and Tlchoice.osc behavior:

- The Tools.osc screen will only be seen by x64-based clients. X86-based and Itanium-based clients will not see the screen.

- If no Windows Server 2003 SP1 RISETUP images (or Windows XP Professional x64 Edition images) exist on the RIS server, the new behavior will not be in place. This is because the new Tlchoice.osc screen and the updated Tools.osc screen will not be copied to the server until the first upload of Windows Server 2003 SP1 (or Windows XP Professional x64 Edition images) RISETUP image.

- If the new Tlchoice.osc screen and updated Tools.osc screen were not copied on upload of the first x64 RISETUP image (for example, the administrator chose to keep the current .osc screens and not copy the new .osc screens on x64 RISETUP image creation), the new functionality will not be in place. To obtain the new functionality to restrict x64 RIS client machines to only view x86 32-bit or x64 64-bit maintenance and troubleshooting tools, the new .osc screens must be in place and the registry value must be configured appropriately.

- The same registry value that controls the x8664.osc behavior in the **Automatic Setup** and **Custom Setup** choices, `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ BINLSVC\Parameters\DefaultPlatformForX8664 (REG_SZ)`, also controls the **Maintenance and Troubleshooting** behavior.

In some situations the Tools.osc screen will be automatically processed with no user interaction required. This is accomplished through the third option **Display Default Tools** that can be seen in the Tools.osc screen source. This option is used in the following situations:

- The booting client computer is not x64-based. If the client computer is x86-based or Itanium-based, installation via the Client Installation Wizard will continue automatically and the user will not be prompted for input in Tools.osc.

- A Windows Server 2003 SP1 x86 or Itanium RISETUP image is placed on a Windows Server 2003 RTM RIS server. In this case, it is possible that the new Tools.osc is placed on a RIS server running Windows Server 2003 RTM if the option to replace .osc screens was chosen upon image upload of a Windows Server 2003 SP1 image. Because the RIS server program has not been updated to Windows Server 2003 Service Pack1, the RIS server does not have the logic to process the new **Display 32-bit Tools** and **Display 64-bit Tools** options in the Tools.osc screen. Because the default option is the only menu choice pre-populated, it is processed and the screen is bypassed.

| **Note** |
| --- |
| As a consequence of the changes in Windows Server 2003 SP1, the screen menu order for the **Maintenance and Troubleshooting** menu option has changed as follows: |

1. Choice.osc

2. Tools.osc

3. TLchoice.osc

**What settings are added or changed in Windows Server 2003 Service Pack 1?**

| Setting name | Location | Previous default value | Default value | Possible values |
| --- | --- | --- | --- | --- |

| | | | | |
|---|---|---|---|---|
| DefaultPlatformForX8664 | HKLM\System \CurrentControlSet \Services\BINLSVC \Parameters | N/A | Not set | Not set<br><br>(Platform type to be selected by users in x8664.osc)<br><br>i386<br><br>(Only display x86-based images in OSChoice.osc)<br><br>amd64<br><br>(Only display x64-based images in OSChoice.osc) |

# Resultant Set of Policy (RSoP)

Applies To: Windows Server 2003 with SP1

**What does Resultant Set of Policy do?**
Group Policy Resultant Set of Policy (RSoP) reports Group Policy settings that are applied to a user or computer. Group Policy Results in Group Policy Management Console (GPMC) requests RSoP data from a target computer and presents this in a report in HTML format. Group Policy Modeling requests the same type of information, but the data reported is from a service that simulates RSoP for a combination of computer and user. This simulation is performed on a domain controller running Windows Server 2003 and is then returned to the computer running GPMC for presentation. Finally, the RSoP Microsoft Management Console (MMC) provides an alternative way to display this information, although Group Policy Results is generally the preferred method.

**Who does this feature apply to?**
Group Policy administrators in an Active Directory domain environment. In addition, an IT professional who needs to plan or validate the application of Group Policy might be interested in RSoP.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**RSoP Use with Windows Firewall Enabled**
**Detailed description**
In Windows XP Service Pack 2 (SP2), Windows Firewall is enabled by default. Incoming requests against unopened ports—as opposed to responses to requests originated from the computer—are blocked by Windows Firewall. In Windows Server 2003 Service Pack 1 (SP1), Windows Firewall is not enabled by default.

If you elect to use Windows Firewall, you should be aware of the impact on its use on RSoP across the network.

For more information about Windows Firewall, see "Windows Firewall," in this document.

**Why is this change important?**
Enabling a firewall, such as Windows Firewall, provides more protection from many network-based attacks. For example, if Windows Firewall had been enabled the recent MSBlaster attack would have been greatly reduced in impact, regardless of whether users were up-to-date with patches.

**What works differently?**
There are two important changes to RSoP in Windows Server 2003 SP1.

- After Windows Firewall is installed on a computer, remote access to RSoP data no longer works from that target computer.
- If Windows Firewall is enabled, when GPMC is run for the purpose of using Group Policy Results or Group Policy Modeling to retrieve RSoP data it will be unable to retrieve this data.

**How do I resolve these issues?**
The following table summarizes the changes necessary to fully support remote RSoP tasks when running Windows XP SP2 or Windows Server 2003 SP1 with Windows Firewall enabled. Please see the sections below for further details.

| Task | Target Computer | Administrative Computer |
|---|---|---|
| Generate Group Policy results | Enable **Windows Firewall Allow remote administration exception** Group Policy setting.<br><br>This Group Policy setting is located in Computer Configuration \Administrative Templates\Network \Network Connections\Windows Firewall\[Domain \| Standard] Profile\. | GPMC with SP1.<br><br>No action required.<br><br>RSoP snap-in.<br><br>Enable Windows Firewall: Define program exceptions. Configure the program exception list with the full path to Unsecapp.exe so that the WMI messages can be transmitted. In a default installation Unsecapp.exe is located in the C:\Windows\System32\Wbem folder.<br><br>Enable **Windows Firewall: Define port exception policy** to open Port 135. |

| | | |
|---|---|---|
| Delegate access to Group Policy results | Enable **Windows Firewall: Allow remote administration exception** Group Policy setting.<br><br>Configure the following DCOM security settings:<br><br>DCOM: Machine access restrictions...<br><br>DCOM: Machine launch restrictions...<br><br>These policy settings are located in Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options. | No changes necessary |
| Remotely edit a Local Group Policy object | Enable **Windows Firewall: Allow file and printer sharing administration exception** policy setting.<br><br>This policy setting is located in Computer Configuration \Administrative Templates\Network \Network Connections\Windows Firewall\[Domain \| Standard] Profile\. | No changes necessary. |

**Administering Remote RSoP with GPMC SP1**

The initial release of GPMC used a callback mechanism when waiting for the results of a Group Policy Results or Modeling request. The administrative computer must be "listening" for this response. If Windows Firewall is enabled, Windows will block these responses. Although opening the appropriate ports can address this issue, using the updated Group Policy Management Console (GPMC) with Service Pack 1 completely removes the use of the callback mechanism. We recommend that you install GPMC with Windows Server 2003 Service Pack 1, because this allows Group Policy Results and Modeling to continue to work without opening up ports on the administrative computer. To install GPMC with Windows Server 2003 Service Pack 1, see "Group Policy Management Console with Service Pack 1" on the Microsoft Download Center at http://go.microsoft.com/fwlink/?LinkId=23529.

In order to administer RSoP remotely, you must enable the **Windows Firewall: Allow remote administration exception** Group Policy setting on target computers.

**Administering Remote RSoP with the RSoP MMC snap-in**

In order to administer RSoP remotely using the RSoP MMC snap-in, the target computer must listen on the appropriate network ports to ensure that incoming RSoP requests can be serviced. This can be managed through Group Policy using the following policy settings:

- Enable the **Windows Firewall: Define program exceptions** Group Policy setting to permit Unsecapp.exe. Make sure you enter the full path to Unsecapp.exe.

- Enable the **Windows Firewall: Define port exception** Group Policy setting and open Port 135. Click **Show** and enter **135:TCP:*:Enabled:135**.

| **Caution** |
|---|
| Enabling the **Windows Firewall: Define port exception** Group Policy setting may also allow unwanted data to be accepted on this port. Be sure to fully review this Group Policy setting before enabling it in your environment.<br><br>Enabling this policy setting is not necessary if the **Windows Firewall: Allow remote administration exception** Group Policy setting is enabled on the administrative computer. |

**Delegating access to Group Policy Results**

By default, Group Policy Results and the RSoP snap-in can only be used remotely when the person originating the request is a local administrator on the target computer. Beginning in Windows Server 2003, a delegation model is available that allows this right to be delegated to users who are not Administrators on the target computer. This is a common scenario when help desk personnel require access to computers without being made Administrators on those computers.

In Windows XP Service Pack 2 and Windows Server 2003 Service Pack 1, the security model for DCOM authentication (on which RSoP relies) has been strengthened. Even if RSoP delegation has been configured correctly, this strengthening prevents local non-administrators from retrieving RSoP information from a target computer. Note that this issue does not impact Group Policy Modeling, since the request for simulated RSoP data is made against a domain controller running Windows Server 2003, which, by definition, is not running Windows XP.

You can manage the list of users and groups associated with DCOM authentication through Group Policy. To allow continued use of delegated RSoP, users to whom you want to grant this right must also have access through the DCOM authentication model. For more information about the security changes to DCOM in Windows Server 2003 Service Pack 1, see "DCOM Security Enhancements" earlier in this document.

Use the following procedure to delegate access to Group Policy Results:

**To delegate access to Group Policy Results**

1. Enable the **Windows Firewall: Allow remote administration exception** Group Policy setting on target computers.
2. Set the following DCOM security policy settings on target computers. (They are located in Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options.)

   **DCOM: Machine access restrictions in Security Descriptor Definition Language (SDDL) syntax**

   **DCOM: Machine launch restrictions in Security Descriptor Definition Language (SDDL) syntax**
3. Right-click the Group Policy object, and then click **Properties**.
4. Click **Edit Security**. **Access Permission opens**.
5. Click **Add**, and then **Select Users, Computers, or Groups opens**.
6. Enter the desired delegation targets.

**Remotely editing a local Group Policy object**

In order to remotely edit a local Group Policy object on a target computer that has Windows Firewall enabled, you need to enable the following policy setting: **Windows Firewall: Allow file and printer sharing administration**.

The policy setting is located in Computer Configuration\Administrative Templates\Network\Network Connections\Windows Firewall\[Domain|Standard] Profile\.

# RPC Interface Restriction

Applies To: Windows Server 2003 with SP1

**What does RPC Interface Restriction do?**
A number of changes have been made in the Remote Procedure Call (RPC) service for Windows Server 2003 with Service Pack 1 that help make RPC interfaces secure by default and reduce the attack surface of Windows Server 2003. The most significant change is the addition of the `RestrictRemoteClients` registry key. This key enables you to modify the behavior of all RPC interfaces on the system and can be used to eliminate remote anonymous access to RPC interfaces on the system, with some exceptions. Additional changes include the `EnableAuthEpResolution` registry key and three new interface registration flags.

**Who does this feature apply to?**
This feature applies to RPC application developers. System administrators should also be familiar with this change to RPC.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**RestrictRemoteClients Registry Key**
**Detailed description**
When an interface is registered using **RpcServerRegisterIf**, RPC allows the server application to restrict access to the interface, typically through a security callback. The `RestrictRemoteClients` registry key forces RPC to perform additional security checks for all interfaces, even if the interface has no registered security callback.

RPC clients that use the named pipe protocol sequence (**ncacn_np**) are exempt from all restrictions discussed in this section. The named pipe protocol sequence cannot be restricted due to several significant backwards compatibility issues.

The `RestrictRemoteClients` registry key can have one of three DWORD values that can also be controlled programmatically in rpcdce.h. If the key is not present, it is equivalent to setting the DWORD=0 value (RPC_RESTRICT_REMOTE_CLIENT_NONE) for server SKUs and DWORD=1 value (RPC_RESTRICT_REMOTE_CLIENT_AUTH) for client SKUs.

The following table provides a reference to the `RestrictRemoteClients` registry key:

| Type | Description |
|---|---|
| Key name | `RestrictRemoteClients` |
| Type | DWORD |
| Configurable through User Interface | Yes. This key can be configured using the Group Policy Object Editor. |
| Key values | 0 (Default for server SKUs)<br><br>This value is the default value in Windows Server 2003 Service Pack 1 for Server SKUs. It causes the system to bypass the RPC interface restriction. This corresponds to the value RPC_RESTRICT_REMOTE_CLIENT_NONE in rpcdce.h. It is entirely the responsibility of the server application to impose appropriate RPC restrictions. This setting is equivalent to the behavior in previous versions of Windows.<br><br>1<br><br>This value is the default value in Windows XP Service Pack 2 and client SKUs built from the SRSP1 code base. It restricts access to all RPC interfaces. All remote anonymous calls are rejected by the RPC runtime except calls coming in through named pipes (ncacn_np). This corresponds to the value RPC_RESTRICT_REMOTE_CLIENT_DEFAULT in rpcdce.h. If an interface registers a security callback and provides the RPC_IF_ALLOW_CALLBACKS_WITH_NO_AUTH flag, then this restriction does not apply to that interface.<br><br>2<br><br>All remote anonymous calls are rejected by the RPC runtime with no exemptions. This corresponds to the value RPC_RESTRICT_REMOTE_CLIENT_HIGH in rpcdce.h. When this value is set, a system cannot receive remote anonymous calls using RPC. |

**Why is this change important? What threats does it help mitigate?**

It is much more difficult to attack an interface if you require calls to perform authentication, even a relatively low level of authentication. Enabling `RestrictRemoteClients` can be a particularly useful mitigation against worms that rely on exploitable buffer overruns that can be invoked remotely through anonymous connections.

**What works differently?**
If your RPC application expects to receive calls from remote anonymous RPC clients, using this feature might not allow your application to run correctly. As a result, applications that use DCOM might not work correctly if this value is set.

Because secure RPC calls over connectionless protocols such as User Datagram Protocol (UDP) and Internetwork Packet Exchange (IPX) (`ncadq_ip_udp` and `ncadg_ipx`) use a lower level of security than calls over connection-oriented protocols, these calls are always considered nonsecure for the purposes of this policy. As a result, RPC calls over connectionless protocols will fail if this key is enabled in Windows Server 2003 with Service Pack 1.

To allow RPC client calls using connectionless protocols, keep the `RestrictRemoteClients` value set to 0 (RPC_RESTRICT_REMOTE_CLIENT_NONE), the default setting for Windows Server 2003 with Service Pack 1.

**How do I resolve these issues?**
There are the following options that may allow you to use `RestrictRemoteClients` on your server:

- Require your RPC clients to use RPC security when contacting your server application. This is the best method to mitigate security threats.

- Exempt your interface from requiring authentication by setting the RPC_IF_ALLOW_CALLBACKS_WITH_NO_AUTH flag during interface registration. This configures RPC to allow anonymous connections to only your application's interface.

**EnableAuthEpResolution Registry Key**
**Detailed description**
An RPC interface that is remotely and anonymously accessible and is registered by default on Windows Server 2003 presents a significant attack surface. RPC itself must register such an interface to provide endpoint resolution for calls using dynamic endpoints.

If you enable the `RestrictRemoteClients` flag, the RPC Endpoint Mapper interface will not be accessible anonymously. This is a significant security improvement, but it changes the task of resolving an endpoint. Currently, an RPC client that attempts to make a call using a dynamic endpoint will first query the RPC Endpoint Mapper on the server to determine what endpoint it should connect to. This query is performed anonymously, even if the RPC client call itself is performed using RPC security. Anonymous calls to the RPC Endpoint Mapper interface will fail on Windows Server 2003 with Service Pack 1 if the `RestrictRemoteClients` key is set to 1 or higher . This makes it necessary to modify the RPC client runtime to perform an authenticated query to the Endpoint Mapper. If the `EnableAuthEpResolution` key is set, the RPC client runtime will use NTLM to authenticate to the endpoint mapper. This authenticated query will take place only if the actual RPC client call uses RPC authentication.

**Why is this change important?**
This change enables an RPC client to make a call to an RPC server that has registered a dynamic endpoint on a system running Windows Server 2003 with Service Pack 1 with `RestrictRemoteClients` enabled. The client computer must set this registry key so that it will perform an authenticated query to the RPC Endpoint Mapper.

**What works differently?**
This registry key is used to enable the specific scenario described in the previous section. When this key is turned on, all RPC Endpoint Mapper queries that are performed on behalf of authenticated calls are performed using NTLM authentication.

This setting can also be specified using the Group Policy Object Editor to configure the Group Policy object located in Computer Configuration\Administrative Templates\System\Remote Procedure Call\RPC Endpoint Mapper Client Authentication.

**New RPC Interface Registration Flags**
**Detailed description**
Three new interface registration flags have been created to make it easier for an application developer to secure an RPC interface.

- RPC_IF_ALLOW_CALLBACKS_WITH_NO_AUTH

  When this flag is registered, the RPC runtime invokes the registered security callback for all calls, regardless of the call security settings. Without this flag, RPC rejects all unauthenticated calls before they reach the security callback. This flag works only when a security callback is registered.

- RPC_IF_SEC_NO_CACHE

  A security callback is registered for an interface in order to restrict access to that interface. The typical security callback impersonates the client to determine whether the client has sufficient rights to make a call to the interface. If a particular client identity passes a security callback once, it usually passes the same security callback every time.

  The RPC runtime takes advantage of this pattern by remembering when an individual client identity passes a security callback and skips the security callback for subsequent calls by that client to the same interface. This feature is called security callback caching and has existed since the Microsoft Windows 2000 family of operating systems. For Windows Server 2003 with Service Pack 1, you can use the RPC_IF_SEC_NO_CACHE flag to disable security callback caching for a given interface. This is useful if the security check might change, possibly rejecting a client identity that was previously permitted.

- RPC_IF_LOCAL_ONLY

  When an interface is registered with this flag, RPC rejects calls made by remote RPC clients. In addition, local calls over all `ncadg_*` protocol sequences and all `ncacn_*` protocol sequences (except for named pipes, using `ncacn_np`) are also rejected. If a call is made on`ncacn_np`, RPC allows the call only if it does not come from SRV, which filters out all remote calls. `Ncalrpc` calls are always allowed through.

**Why is this change important?**
This change provides RPC application developers with additional security tools to help secure their RPC interface.

**What works differently?**
These flags will not change any existing Windows Server 2003–based application or cause it not to run correctly. The use of these new flags is at the discretion of the application developer.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
**RPC Settings**

| Setting name | Location | Default value | Possible values |
|---|---|---|---|
| RestrictRemoteClients | HKEY_LOCAL_MACHINE\ SOFTWARE\Policies\ Microsoft\Windows NT\RPC<br><br>-or-<br><br>(Group Policy object)<br><br>Computer Configuration\ Administrative Templates\System\Remote Procedure Call\Restrictions for Unauthenticated RPC Clients | 0 | 0 – None<br><br>1 – XP SP2 and SRSP1 client SKUs Default<br><br>2 – High |
| EnableAuthEpResolution | HKEY_LOCAL_MACHINE\ SOFTWARE\Policies \Microsoft\Windows NT\RPC<br><br>-or-<br><br>(Group Policy object)<br><br>Computer Configuration\ Administrative Templates\System\Remote Procedure Call\RPC Endpoint Mapper Client Authentication | 0 | 0 – Disabled<br><br>1 – Enabled |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
You may need to change your code to work with Windows Server 2003 Service Pack 1 if you decide to enable `RestrictRemoteClients`. For more information about application changes that might be required, see the previous sections on `RestrictRemoteClients` and `EnableAuthEpResolution`.

# Security Configuration Wizard in Windows Server 2003 Service Pack 1

Applies To: Windows Server 2003 with SP1

**What does Security Configuration Wizard do?**
Security Configuration Wizard (SCW) is a new feature in Windows Server 2003 with Service Pack 1 that provides guided attack surface reduction for your servers. SCW is highly recommended for creating security policies for servers based on their roles. SCW includes the following features:

- Remote or local role-based lockdown ensures that all necessary services are turned on, minimizing risk of breakage.
- Any services not specifically needed by the role(s) performed by the target system can be disabled.
- Firewall and IPsec configuration ensures that the server presents the minimum possible attack profile based on the needs of the selected role(s).
- Automated selection of key security settings based on a guided wizard to minimize compatibility problems and maximize security.
- Selection of appropriate audit settings.
- Inclusion of standard security templates for detailed customization of security policy.
- Integration with Group Policy for Active Directory-based deployment of security policies.
- Fully functional and scriptable command line interface for automated deployment to multiple servers across the network.
- Rollback of policies for testing and troubleshooting purposes.
- Declarative extensibility of role knowledge base to allow authoring of new roles.

When you run SCW you will be asked a series of questions to determine the functional requirements of your server. Any functionality that is not required based on your selections can be disabled automatically.

**Who does this feature apply to?**
SCW can be used with any system running Windows Server 2003 with Service Pack 1 and in any network configuration. SCW will be of interest to:

- Security professionals responsible for authoring corporate security policies.
- Information technology (IT) professionals responsible for deploying, configuring, and managing servers.
- Developers of server-based applications who want their applications to be managed with SCW.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Security Configuration Wizard**
**Detailed description**
Security Configuration Wizard uses a roles-based metaphor driven by an extensible XML knowledge base that defines the services, ports, and other functional requirements for almost 200 different system roles, including roles for Windows Server System applications such as Microsoft ISA Server and SQL Server.

SCW uses this extensible XML knowledge base to perform role discovery, solicit user input, and author security policies that disable services, block ports, modify registry values, and configure audit settings. Even ports that are left open can be restricted to specific subnets or systems using Internet Protocol security (IPsec). SCW also allows you to roll back previously applied policy settings. It includes a command line tool you use with administrative scripts and other administrative utilities to apply a

security configuration and perform compliance analysis on groups of servers in your organization. SCW also integrates with Active Directory to support deployment of SCW-generated policy settings through Group Policy.

**Summary of SCW security coverage**
Security Configuration Wizard allows users to easily:

- Disable unnecessary services.
- Protect Internet Information Services (IIS).

| **Note** |
| --- |
| IIS policies cannot be deployed with Group Policy since Group Policy does not currently configure IIS. If a policy containing IIS settings is converted to a Group Policy object the IIS settings are lost. |

- Block unused ports, including support for multi-homed scenarios.
- Secure ports that are left open using IPsec.
- Reduce protocol exposure for Lightweight Directory Access Protocol (LDAP), LAN Manager, and server message block (SMB).
- Configure audit settings with a high signal-to-noise ratio.
- Import Windows security templates for coverage of settings that are not configured by the wizard.

**Summary of SCW operational features**
In addition to roles-based guided security policy authoring, SCW also supports:

- **Rollback**. Enables you to return your server to the state it was in before you applied the SCW security policy.
- **Analysis**. Enables you to check that servers are in compliance with expected policies.
- **Remote** configuration and analysis support. All SCW functionality can be performed on both local and remote systems.
- **Command line support**. A command-line tool is provided for scripting use.
- **Active Directory integration**. Supports deploying SCW policies using Group Policy.
- **Editing**. A security policy created using SCW can be modified when necessary, such as when machines are repurposed or when a system configured with a particular policy does not behave as expected.
- Reports. Provides the ability to view the data stored in the knowledge base, policies, and analysis results XML files.

**Why is this change important?**
Attack surface reduction is a fundamental security best practice, yet it is often difficult for server administrators to find the time to properly secure, test, and deploy a Windows server without breaking required functionality, which can lead to vulnerable servers within an organization.

Security configuration guides (such as the Windows Server 2003 Security Guide on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=14845) provide general settings that support a broad range of systems but that do not provide an optimal security versus functionality tradeoff for a specific class of systems. SCW automates the lockdown process of systems providing specific functionality and is fully tested and supported by Microsoft. Reducing the attack surface of Windows servers can minimize the number of servers that need to be immediately patched when a vulnerability is discovered, as a given vulnerability will not necessarily be present or exploitable in all configurations.

**What works differently?**

Today, Windows administrators typically define security policies using the Security Configuration Editor on their own, in conjunction with documented guidance, or with existing security templates designed for specific scenarios. In contrast, Security Configuration Wizard is an authoring tool that allows you to create a custom security policy by answering a series of questions. For settings that are not configured by the wizard, SCW allows the administrator to import existing security templates.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
No, but SCW is extensible so that developers can create their own SCW role definitions for their own applications. A whitepaper about how to extend the SCW knowledge base will be available when Windows Server 2003 Service Pack 1 is released.

**What do I need to change in my environment to deploy Windows Server 2003 Service Pack 1?**
Nothing, however SCW can be used during the deployment process to ensure servers are deployed with the expected security policy.

If unattended setup is used to deploy servers, consider the following:

- Install the SCW optional component automatically during unattended setup by adding the following entry to the `[Components]` section of unattend.txt: `SCW = On`.
- To apply an SCW policy during the unattended installation, also carry out the following steps:
  - Create the policy file on a server that already has SCW installed.
  - Either create a Cmdlines.txt file or modify the existing one so that it has a `[Commands]` section containing the following line: `scwcmd configure /p:SCWPolicy.xml`
  - Copy the Cmdlines.txt file and the previously created policy file (SCWPolicy.xml in this case) to the $OEM$ directory.
- If an imaging solution is used to deploy servers, you can apply an SCW on the reference machine that will be imaged prior to creating the image.

For additional information on Security Configuration Wizard go to the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=45503.

# Setup

Applies To: Windows Server 2003 with SP1

**What does Setup do?**
Setup is the program that installs and configures the operating system on your computer. When you originally install Windows Server 2003, Setup installs and configures the operating system to work properly with the computer.

Often, any program that is used to install and configure software programs is called Setup. However, in Windows Server 2003, there are two different programs that install and configure programs after the operating system is operational. When Windows Server 2003 Service Pack 1 (SP1) or any other software update is installed on Windows Server 2003, the Package Installer for Windows updates and changes the existing installation of Windows Server 2003. If an application, such as Microsoft Office, is installed, configured, or upgraded, then Windows Installer is the program that is responsible for that task.

**Who does this feature apply to?**
All users and administrators who install or deploy Windows Server 2003 Service Pack 1 should become familiar with the changes in this feature.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Package Installer for Windows**
**Detailed description**
The Package Installer for Windows (Update.exe) installs Windows Server 2003 Service Pack 1. This is different installation technology from what is used to install the complete operating system. For a complete discussion of the Package Installer, see "The Package Installer for Windows and Windows Components" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=26004.

Changes in Package Installer in Windows Server 2003 Service Pack 1 include:

- Command line switches

  Microsoft has moved to a standard for Installer switches, which will be available later this year. There are no new switches to replace the **/D** and **/S** switch functions; the old switches are still valid. The following table describes the former, new, and legacy switches for use with the Package Installer.

| Old Switch | New Switch | Description |
|---|---|---|
| None | /Uninstall | Used to automate removal of a service pack. |
| /? | /Help | Displays help text. |
| /D:*FolderName* | No change | For service pack installations, backs up files to the specified folder. |
| None | /ER | Enables extended return codes. (For more information about extended return codes, see Appendix F in the "The Package Installer for Windows and Windows Components" white paper.) |
| /F | /F | Forces other applications to close when the computer restarts after the installation. |
| None | /Forcerestart | Forces restart after installation finishes. |
| /L | /L | Lists installed hotfixes (Windows updates only). |
| /N | /N | Does not back up files for removing the service pack or hotfix. In Add or Remove Programs, there is no Remove button for the hotfix, so it cannot be uninstalled. |
| /O | /O | Overwrites OEM files without prompting. |
| /Q | /Quiet | Uses quiet mode, which is the same as unattended mode, except that the user interface is hidden. No prompts appear during the installation. |

| | | |
|---|---|---|
| /Intergrate:*FolderName* | No change | Combines the operating system image with Windows Server 2003 Service Pack 1 in a shared distribution folder for an integrated installation. When using this switch, type the name of the shared distribution folder after the colon. |
| /U | /Passive | Uses unattended setup mode. Only critical errors and a progress bar appear during setup. |
| /Z | /Norestart | Does not restart the computer after the installation finishes. |

- **/O** command line switch

  The **/O** command line switch prevents Update.exe from performing an exhaustive search for original equipment manufacturer (OEM) supplied files, such as video drivers. Use this switch only if you want to overwrite all files supplied by the OEM on your system. By default, OEM files are not overwritten.

- Do not apply updates during installation

  If you are running in an environment that uses a utility to scan and apply updates or security updates automatically, the utility may conflict with the service pack installation. This conflict can occur when a service pack is in the process of installation or uninstallation and has not fully completed due to the need to restart and then run the post-installation processes. If the update utility attempts to apply updates to the system at shutdown or restart, it may attempt to install updates prior to completing the service pack installation, which may result in the computer being left in an unusable state.

- Improved recovery following installation failure

  In the event that there is a failure to complete the service pack installation, in most cases, automatic recovery will occur. For example, if power failure occurs during the latter parts of installation, when you restart the computer, the service pack uninstall procedure automatically starts and rolls back all the changes that were made during the service pack installation.

- Antivirus software

  Although antivirus software is a critical factor in keeping a computer secure, it can affect the overall time required to install a service pack. Disabling your antivirus software during installation of the service pack could reduce the installation time by up to 20%. If you decide to disable the antivirus software, be aware of the risks involved and be sure to enable it after the service pack is installed.

**What works differently?**

A new prerequisite to service pack installation is a check as to whether the file SpUpdSvc.inf exists on the computer. This file will be present only if Windows Server 2003 Service Pack 1 was previously installed on the computer. If this file is found on your computer by the SP1 installer you will be presented with the message:

**Setup has detected an .INF file from the beta installation that will prevent successful installation of the Service Pack. Please see Deleting SpUpdSvc.inf before install in the Readme document accompanying this release for details on how to unblock the Windows Server 2003 Service Pack 1 installation.**

There is also a prerequisite that blocks installation of the service pack if the Distributed Transaction Coordinator (also know as MSDTC) service status is **Disabled** and **Stopped** at the time of installation. In this case the installation will stop and you will be presented with the message:

**Setup has detected that the Distributed Transaction Coordinator (MSDTC) service is disabled. The MSDTC service is required for the successful installation of Microsoft ® Windows Server™ 2003 Service Pack 1 (SP1). Please enable the MSDTC service, and then restart the installation of Windows Server 2003 SP1. For more information about enabling the MSDTC service, contact your network administrator.**

**How do I resolve these issues?**

**To manually delete the SPUpdSvc.INF file**

1. Locate the SPUpdSvc.INF file in the Windows directory.

2. Rename the file from SPUpdSvc.INF to SPUpdSvc.BAK.
3. Restart your computer.
4. Attempt to install the service pack
5. If the service pack successfully installs, delete the SPUpdSvc.BAK file.

**To enable the MSDTC service**

1. Open a command prompt.
2. At the command prompt, type the following command:

   **sc config MSDTC start=demand**

   This will change the status of the MSDTC service from disabled to enabled, but will not start the service.

# Terminal Services in Windows Server 2003 Service Pack 1

Applies To: Windows Server 2003 with SP1

**What does Terminal Services do?**
On Windows Server 2003 operating systems, the Terminal Server feature gives users at client computers throughout your network access to Windows-based programs installed on terminal servers. With Terminal Server, you can provide a single point of installation that allows multiple users access to Windows Server 2003 operating system desktops, where they can run programs, save files, and use network resources, all from a remote location, as if these resources were installed on their own computers.

Terminal Services is ideal for rapidly deploying Windows-based applications to computing devices across an enterprise—especially applications that are frequently updated, infrequently used, or hard to manage. Terminal Server lets you deliver Windows-based applications, or the Windows desktop itself, to virtually any computing device—including those that cannot run Windows.

Windows Server 2003 Service Pack 1 includes several new features designed to maximize both the speed and efficiency of Terminal Services administration, and the security of communications between Terminal Services clients and servers.

**Who does this feature apply to?**
The features described here will be of interest to Terminal Server client users as well as IT professionals who deploy and configure Terminal Services.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**New fallback printer driver capability**
**Detailed Description**
With the release of Windows Server 2003 with Service Pack 1 (SP1), you can make local printing more accessible for Terminal Server clients by configuring Terminal Services to default to a printer driver compatible with PostScript (PS) or Printer Control Language (PCL). The new fallback printer driver capability is exceptionally useful if a terminal server does not have a printer driver installed that matches the Terminal Server client user's specific printer brand and model.

A new Group Policy setting, **Terminal Server fallback printer driver behavior**, allows you to specify the location and file name of a fallback printer driver, in the event that no printer drivers installed on a terminal server are compatible with the local printer for a Terminal Server client.

By default, the Terminal Server fallback printer driver is disabled. If the terminal server does not have a printer driver that matches the client's printer, no printer will be available for the terminal server session.

If the fallback printer driver is enabled, Terminal Server's default behavior is to locate a suitable printer driver. If one is not found, the client user cannot print Terminal Server session documents to a local printer. The Group Policy setting allows you to select one of four options to modify Terminal Server printing behavior:

- **Do nothing if one is not found**. This is the default setting. In the event of a printer driver mismatch, the server attempts to find a suitable driver. If one is not found, the client's printer is unavailable during the Terminal Server session.

- **Default to PCL if one is not found**. If no suitable printer driver can be found, Terminal Server uses the Hewlett-Packard compatible Printer Control Language (PCL) fallback printer driver.

- **Default to PS if one is not found**. If no suitable printer driver can be found, Terminal Server uses the Adobe PostScript (PS) fallback printer driver.

- **Show both PCL and PS if one is not found**. In the event that no suitable driver can be found, show both PS-based and PCL-based fallback printer drivers.

If this setting is disabled or not configured, Terminal Server does not use a fallback printer driver.

Printing Terminal Server session documents may still be disabled for some client computers, if the fallback printer driver's vendors have deviated from PS or PCL specifications.

**Note**

If the Group Policy setting **Do not allow client printer redirection** is enabled, any configuration for the **Terminal Server fallback printer driver behavior** policy setting is ignored, and the fallback driver is disabled.

**Why is this change important?**
This change simplifies local printing for Terminal Server client users. The new Group Policy setting allows client users to print documents locally, if the printer driver installed on the terminal server to which they're connected is incompatible with their local printers, provided their printers are compatible with either a PCL or a PS printer driver.

**Authentication and encryption for Terminal Services connections**
**Detailed description**
In Windows Server 2003 SP1, you can enhance the security of Terminal Server by configuring Terminal Services connections to use Secure Sockets Layer (SSL)/Transport Layer Security (TLS) 1.0 for server authentication, and to encrypt terminal server communications. The version used by Terminal Services in Windows Server 2003 SP1 is TLS 1.0.

**Server prerequisites**
For SSL (TLS) authentication to work correctly, the terminal server must meet both the following prerequisites:

- The terminal server must be running Windows Server 2003 SP1.
- You must obtain a certificate for the terminal server. To obtain a certificate, use one of the following methods:
  - Visit the Web site for your certification authority. For example, visit http://*servername*/certsrv.
  - Run the Windows Server 2003 Certificate Request Wizard or the Windows 2000 Server Certificate Request Wizard.
  - Obtain a certificate from a third-party certification authority, and then manually install the certificate.

If you plan to obtain a certificate by using the certification authority Web site or the Certificate Request Wizard, a public key infrastructure (PKI) must be configured correctly to issue SSL-compatible X.509 certificates to the terminal server. Each certificate must be configured as follows:

- The certificate is a computer certificate.
- The intended purpose of the certificate is server authentication.
- The certificate has a corresponding private key.
- The certificate is stored in the terminal server's personal store. You can view this store by using the Certificates snap-in.
- The certificate has a cryptographic service provider (CSP) that can be used for the SSL (TLS) protocol (for example Microsoft RSA SChannel Cryptographic Provider).

For more information, see Microsoft Cryptographic Service Providers (http://go.microsoft.com/fwlink/?LinkID=40983).

**Client prerequisites**
In order for SSL (TLS) authentication to function correctly, clients must meet the following prerequisites:

- Clients must run Windows 2000 or Windows XP.
- Clients must be upgraded to use the Remote Desktop Protocol (RDP) 5.2 (Windows Server 2003) client. You can install this client-side Remote Desktop Connection package by using the %systemdrive\system32\clients\tsclient\win32\msrdpcli.msi file. The msrdpcli.msi file is located on Windows Server 2003 terminal servers. Installing this file from the terminal server installs the 5.2 version of Remote Desktop Connection to the %systemdrive\Program files\Remote Desktop folder on the destination computer. For more information, see Remote Desktop Connection for Windows Server 2003 [5.2.3790] (http://go.microsoft.com/fwlink/?LinkID=41068).

- Clients must trust the root of the server's certificate. That is, clients must have the certificate of the certification authority (CA) that issued the server certificate in their Trusted Root Certification Authorities store. You can view the certificate by using the Certificates snap-in.

| **Important** |
| --- |
| Because RDP runs on port 3389, when using SSL (TLS) to secure RDP, SSL (TLS) will run on port 3389. |

**Why is this change important?**
By default, Terminal Server uses native Remote Desktop Protocol (RDP), which provides data encryption, but does not provide authentication to verify the identity of a terminal server.

For more information about Terminal Services and security protocol settings, see the following:

- Configure Authentication and Encryption (http://go.microsoft.com/fwlink/?LinkId=45407)
- How to configure a Windows Server 2003 terminal server to use TLS for server authentication (http://go.microsoft.com/fwlink/?LinkId=64593)

**New Group Policy settings for Terminal Services Licensing**
**Detailed Description**
Windows Server 2003 SP1 includes new Group Policy settings for Terminal Services Licensing described as follows.

**Set the Terminal Server licensing mode**
The new Group Policy setting **Set the Terminal Server licensing mode** determines the type of Terminal Server client access license (CAL) a device or user requires to connect to a terminal server.

When this setting is enabled, you can choose one of the following two licensing modes:

- **Per User**: Each user connecting to the terminal server requires a Per User Terminal Server CAL.
- **Per Device**: Each device connecting to the terminal server requires a Per Device Terminal Server CAL.

If you enable this policy setting, the licensing mode that you specify overrides the licensing mode specified during setup, or in Terminal Services Configuration (TSCC.msc).

If you disable or do not configure this policy setting, Terminal Services uses the licensing mode specified during setup or found in TSCC.msc.

To configure the Terminal Services licensing mode on a specific terminal server using TSCC.msc, see Configure the Terminal Server licensing mode (http://go.microsoft.com/fwlink/?LinkId=45592) in the Terminal Services Help.

**Use the specified Terminal Server license servers policy setting**
The Group Policy setting **Use the specified Terminal Server license servers** determines whether terminal servers must first attempt to locate Terminal Server license servers that are specified in this policy setting before attempting to locate license servers elsewhere on the network.

During the automatic discovery process, terminal servers attempt to contact license servers in the following order:

1. Enterprise license servers or domain license servers that are specified in the LicenseServers registry key.
2. Enterprise license servers that are specified in Active Directory.
3. Domain license servers.

If you enable this policy setting, terminal servers attempt to locate license servers that are specified in this setting, before following the automatic license server discovery process.

If you disable or do not configure this policy setting, terminal servers follow the automatic license server discovery process.

You can configure a specific terminal server to locate a Terminal Server license server using TSCC.msc. For more information,

see Set preferred Terminal Server license servers (http://go.microsoft.com/fwlink/?LinkId=45410) in the Terminal Server Licensing Help.

**Show ToolTips for licensing problems on Terminal Server policy setting**
This Group Policy setting allows you, after successfully logging on to a terminal server as an administrator, to display ToolTips that show any licensing problems with the terminal server, and also display the expiration date of the terminal server's licensing grace period. If this Group Policy setting is not configured, ToolTip display is defined by registry settings.

**Why is this change important?**
Specifying the name of a preferred licensing server in Group Policy saves time and may eliminate roadblocks to successful configuration of your terminal servers. With the name of a specific licensing server added to Group Policy, Terminal Services does not need to search the network for a licensing server.

Using ToolTips to view Terminal Server license statistics at a glance speeds administration tasks. By configuring Group Policy to show ToolTips for Terminal Server licenses, you do not need to open the **Properties** dialog box for specific licenses to view status and expiration information.

Allowing administrators to configure a global Terminal Server licensing mode makes it possible for them to implement unified license policies regardless of the configuration of Terminal Services client computers. With the new Group Policy setting, differences in configuration between terminal servers and clients can be resolved by defining a global policy that overrides other settings.

For more information about Terminal Server Licensing, see the following:

- Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkId=45409)
- Set preferred Terminal Server license servers (http://go.microsoft.com/fwlink/?LinkId=45410)

**Update to Group Policy setting for starting a program on connection to a terminal server**
**Detailed Description**
The Group Policy setting **Start a program on connection** configures Terminal Services to run a specified program automatically when a client connects to a terminal server.

By default, Terminal Services sessions provide access to the full Windows desktop, unless the server administrator has otherwise specified using this policy setting, or unless the user has specified during configuration of the client connection. Enabling this Group Policy setting overrides the **Start program** settings made by the server administrator or user. The **Start** menu and Windows Desktop are not displayed, and when the user exits the program, the Terminal Server session is automatically logged off.

If the **Start a program on connection** policy setting is enabled, Terminal Services sessions automatically run the specified program and use the specified working folder (or the program default folder, if a working folder is not specified) as the working folder for the program.

If this policy setting is disabled or not configured, Terminal Services sessions start with the full desktop, unless the server administrator or client user specifies otherwise.

| **Note** |
| --- |
| This setting appears in both **Computer Configuration** and **User Configuration**. If both settings are configured, the **Computer Configuration** setting overrides local user settings. |

You can configure a specific terminal server to start a program when a client successfully logs on using TSCC.msc. For more information, see Specify a program to start automatically when a user logs on (http://go.microsoft.com/fwlink/?linkid=64608) in the Terminal Services Configuration Help.

**Why is this change important?**
Before the release of Windows Server 2003 with Service Pack 1 (SP1), this policy setting could only be edited in Group Policy if the computer was a domain controller, and it was necessary to access Group Policy by opening **Active Directory Users and Computers**. Now, you can modify the **Start program on connection** policy setting in Group Policy for the local policy object, meaning that you can configure this policy setting for individual terminal servers within a domain.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
The following table lists the Group Policy settings that have changed for Terminal Services in Windows Server 2003 with SP1, and provides their locations in Group Policy.

| Setting name | Location | Default value | Possible values |
|---|---|---|---|
| Terminal Server Fallback Printer Driver Behavior | Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection | Not configured | Enabled, disabled, not configured |
| Set the Terminal Server licensing mode | Administrative Templates\Windows Components\Terminal Services | Not configured | Enabled, disabled, not configured |
| Use the specified Terminal Server license servers | Administrative Templates\Windows Components\Terminal Services | Not configured | Enabled, disabled, not configured |
| Show ToolTips for licensing problems on Terminal Server | Administrative Templates\Windows Components\Terminal Services | Not configured | Enabled, disabled, not configured |
| Start a program on connection | Administrative Templates\Windows Components\Terminal Services | Not configured | Enabled, disabled, not configured |

# Transmission Control Protocol/Internet Protocol (TCP/IP)

Applies To: Windows Server 2003 with SP1

**What does TCP/IP do?**
Transmission Control Protocol/Internet Protocol (TCP/IP) is a suite of standard protocols for connecting computers across networks. TCP/IP enables Windows-based computers to connect and share information with other Microsoft and non-Microsoft systems.

**Who does this feature apply to?**
All users who use TCP/IP to connect and communicate information over a network should be aware of the changes incorporated in Windows Server 2003 Service Pack 1.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**SYN attack protection is enabled by default**
**Detailed Description**
When a SYN attack is detected, TCP/IP in Windows Server 2003 and Windows XP lowers the number of retransmissions of the SYN-ACK segment and does not allocate memory or table entry resources for the connection until the TCP three-way handshake has been completed.

A TCP Synchronize (SYN) attack is a denial-of-service attack that exploits the retransmission and time-out behavior of the Synchronize-Acknowledgement (SYN-ACK) segment during the TCP three-way handshake to create a large number of half-open TCP connections. Depending on the TCP/IP protocol implementation, a large number of half-open TCP connections could do any of the following:

- Use all available memory.
- Use all possible entries in the TCP Transmission Control Block (TCB), an internal table used to track TCP connections. When the half-open connections use all the entries, further connection attempts are responded to with a TCP connection reset.
- Use all available half-open connections. When all the half-open connections are used, further connection attempts are responded to with a TCP connection reset.

**Why is this change important? What threats does it help mitigate?**
To mitigate the impact on a host experiencing a SYN attack, TCP/IP minimizes the amount of resources devoted to incomplete TCP connections and reduces the amount of time before abandoning the connection. When a SYN attack is detected, TCP/IP in Windows Server 2003 and Windows XP lowers the number of retransmissions of the SYN-ACK segment and does not allocate memory or table entry resources for the connection until the TCP three-way handshake has been completed.

You can control SYN attack protection through the `SynAttackProtect` registry setting at `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters`(REG_DWORD type). You set `SynAttackProtect` to 0 to disable SYN attack protection and to 1 to enable it.

For TCP/IP in Windows XP (all versions) and Windows Server 2003 with no service packs installed, `SynAttackProtect` is set to 0 by default. For TCP/IP in Windows Server 2003 SP1, `SynAttackProtect` is set to 1 by default.

**New SYN attack notification IP Helper APIs**
**Detailed Description**
To allow an application to notify network administrators that a SYN attack is taking place, the IP Helper API supports new SYN attack notification APIs named `NotifySecurityHealthChange` and `CancelSecurityHealthChangeNotify`. Information about these new APIs has not yet been published in the Microsoft Developer Network (MSDN). A link to the MSDN topics describing these new APIs will be posted here when available.

**Smart TCP port allocation**
**Detailed Description**
TCP/IP in Windows Server 2003 SP1 has implemented a smart TCP port allocation algorithm. When an application requests any available TCP port, TCP/IP first attempts to find an available port that does not correspond to a connection in the TIME WAIT state. If a port cannot be found, then it picks any available port.

**Why is this change important? What threats does it mitigate?**
When a TCP peer initiates a TCP connection termination and the connection termination completes, the TCP connection enters

the TIME WAIT state. When the TIME WAIT state is reached, TCP must wait twice the maximum segment lifetime (MSL) before a connection with the same set of socket addresses can be created. The set of socket addresses consists of the combination of the source and destination IP addresses and source and destination TCP ports. The MSL is the maximum amount of time a TCP segment can exist in an internetwork, and its recommended value is 120 seconds. This delay prevents a new connection's TCP segments using the same set of socket addresses from being confused with duplicated TCP segments of the old connection.

The TCP port for a connection in the TIME WAIT state is considered an available port and can be assigned for use by an application. This can lead to the following series of events:

1. An application requests any available TCP port.
2. TCP/IP assigns a TCP port to use for the application socket.
3. The application attempts to open a socket with a specific destination IP address.
4. The application establishes a TCP connection and sends data.
5. The application terminates the TCP connection.
6. TCP/IP places the application's TCP connection in the TIME WAIT state until 2*MSL has passed.
7. The same application requests another available TCP port.
8. TCP/IP assigns a TCP port to use for the application socket. Because the port for the connection in the TIME WAIT state is considered open, it can be chosen as the next port to assign to the requesting application. Assume that TCP/IP assigns the same TCP port number.
9. The application attempts to open a socket with the same destination IP address.
10. Because the connection is using the same set of socket addresses as the connection in the TIME WAIT state, TCP/IP indicates an error to the application.

Mitigating this error situation required adjusting the registry settings in one of the following ways:

- Setting the `MaxFreeTWTcbs` registry setting at `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters` (REG_DWORD type) to a lower value. The value of `MaxFreeTWTcbs` controls the number of connections that can be in the TIME WAIT state. When this number is exceeded, the oldest connection is automatically removed from the TIME WAIT state.
- Setting the `TcpTimedWaitDelay` registry setting at `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters` (REG_DWORD type) to a lower value. The value of `TcpTimedWaitDelay` determines the length of time that a connection stays in the TIME WAIT state.

However, lowering the value of these registry settings is contrary to the original design of TCP and the MSL and is not recommended.

Therefore, Windows Server 2003 SP1 includes this port allocation algorithm change to prevent an application from creating a connection with the same set of socket addresses of a connection that is in a TIME WAIT state if another port is available.

This new behavior makes it much more unlikely that an application will be assigned a TCP port that is in the TIME WAIT state when connecting to the same destination, so that it is no longer necessary to modify the values of the `MaxFreeTWTcbs` and `TcpTimedWaitDelay` registry keys to avoid the error situation described above.

**Improved receive Winsock APIs**
**Detailed description**
Winsock APIs now support a mechanism that allows applications to specify that a receive request should not be completed until the receive buffer is full. This helps reduce network latency and better supports high performance applications. For more information about the Winsock implementation, see "About the Winsock SPI" on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=45684.

**Winsock self-healing**
**Detailed description**
Winsock, the consortium-based networking application programming interface (API), is extensible by a mechanism known as the Winsock Service Provider Interface (SPI). The SPI provides a mechanism for layering providers on top of each other. Providers that are layered in this fashion are called Winsock Layered Service Providers (LSPs). Winsock LSPs are available for a wide range of useful purposes, including parental controls and Web content filtering. The specific layering order of all providers is kept in the Winsock Catalog. In previous versions of Windows Server 2003, improperly removed LSPs (possibly due to poorly implemented LSP install/removal code) could result in corruption of the Winsock Catalog, potentially resulting in a loss of all network connectivity. Winsock now has the ability to detect improperly removed LSPs and automatically repair itself.

**Why is this change important? What threats does it help mitigate?**
Customers need to be able to safely remove LSPs from their systems.

**New Winsock Netsh commands**
**Detailed description**
Two new Netsh commands are available in Windows Server 2003 Service Pack 1:

- `netsh winsock reset catalog`

  This command resets the Winsock catalog to the default configuration. This can be useful if a malformed LSP is installed that results in loss of network connectivity. While use of this command can restore network connectivity, it should be used with care because any previously installed LSPs will need to be reinstalled.

- `netsh winsock show catalog`

  This command displays the list of Winsock LSPs that are installed on the computer.

**Why is this change important? What threats does it help mitigate?**
These commands provide additional management capabilities for maintaining and troubleshooting Winsock LSPs and can be used in a script to aid in recovering from a widespread installation of malformed LSPs.

**New Netstat parameter**
The following new parameter can be used with the **netstat** command:

- `netstat -b`

Using this parameter with the **netstat** command displays the executable involved in creating each connection or listening port. In some cases well-known executables host multiple independent components, and in these cases the sequence of components involved in creating the connection or listening port is displayed with components being displayed prior to the executable that called the component. The executable name will be denoted by brackets ([]) on either side of the executable name.

| **Note** |
| --- |
| Using this parameter can take a considerable amount of time and might fail if the account it is running under does not have sufficient permissions to gather the information requested. |

**What new settings are added in Windows Server 2003 Service Pack 1?**
A new registry key has been added as shown in the table below. This registry key sets the maximum number of ICMP host routes that can be added. This key should not be changed unless ICMP is being used to add a large number of host routes.

| Setting name | Location | Default value |
| --- | --- | --- |
| MaxICMPHostRoutes | HKEY_LOCAL_MACHINE \SYSTEM \CurrentControlSet \Services\Tcpip\Parameters | 1000 |

# WebDAV Redirector

Applies To: Windows Server 2003 with SP1

**What does WebDAV Redirector do?**
The WebDAV Redirector (DAVRdr) allows computers running Windows Server 2003 to use WebDAV (Web-based Distributed Authoring and Versioning) servers, such as Windows SharePoint Services and MSN Communities, as if they were standard file servers. It consists of a kernel component that connects to a Windows NT remote file system stack, and a user-level component (Web client service) that translates file system requests into WebDAV requests.

**Who does this feature apply to?**
This feature is used by people who access WebDAV servers through the remote file system. WebDAV Redirector is implemented in the remote file system stack. Client administrators, and users who are concerned with the security of their computer credentials, need to be aware of this change if they access remote files on a WebDAV server using a Universal Naming Convention (UNC) path, for example, a file request such as \\ServerName\ShareName\File.txt will be processed by WebDAV Redirector and will be affected by this change in functionality.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Disabling Basic Authentication over a clear channel**
**Detailed description**
WebDAV is an extension of Hypertext Transfer Protocol (HTTP), and as such includes the use of Basic Authentication (BasicAuth). BasicAuth is one form of user authentication, or means by which a user is identified to the server. With BasicAuth, the client transmits the user's credentials (user name and password) to the server. If the channel is unencrypted, such as with normal HTTP traffic, any computer on the network can see the user's user name and password and therefore steal their identity. The DAVRdr does not support encrypted HTTP (HTTPS or SSL), and will transmit the user's credentials in the clear (without encryption) if the server supports basic authentication. Although a server most likely would not be configured to use basic authentication, it would be possible to set up the server expressly to obtain users' credentials.

Because of this possibility, Windows Server 2003 Service Pack 1 (SP1) includes the ability to enable or disable the use of BasicAuth by the DAVRdr. By default, use of BasicAuth is disabled with SP1. When BasicAuth is disabled, the client will either use a different authentication method (if the server supports one) or fail the request.

**Why is this change important?**
Users can log on to WebDAV servers for remote file access without fear of transmitting their password in the clear.

**What threats does it help mitigate?**
Imagine a corporate user at Contoso Corporation who routinely accesses the file share \\Contoso_Server\Sales outside the corporation on a public network, and uses an application that attempts to access that share as part of normal background activity. Since the user's portable computer is outside the corporate network, the request should fail. However, the DAVRdr will transmit a request to see if there is a DAV server named Contoso_Server, even though the actual server that the portable computer is attempting to access is an SMB server.

An attacker can be operating on that same public network with a computer that spoofs WINS requests, returning a pointer to itself in response to any WINS request. The portable computer will then try to access a DAV share on that rogue server. If the rogue server responds with BasicAuth as the authentication method, a dialog box appears that asks for the user's credentials. The dialog box identifies the server as Contoso_Server, leading the user to believe the request is legitimate. If the user enters their user name and password, the client transmits that information in the clear and the attacker thus gains access to that user's login information. There is no indication to the user that the channel is not secure, that the request is being handled by the DAVRdr, or that the portable computer will transmit the user name and password in the clear. Note that the current default Windows authentication methods never transmit a user's password in the clear.

**What works differently?**
Since the change to default behavior only affects the DAVRdr, the only scenarios that fail to work are those that require basic authentication, and that use the DAVRdr. An example is using Notepad.exe to access a Web site that only allows BasicAuth. This scenario will no longer work. Also, even if the server was configured to only use basic authentication, other applications such as Microsoft Office will continue to work, since they use a different DAV client.

**How do I resolve these issues?**
You can enable BasicAuth by adding the following registry key and setting it to a non-zero value:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services \WebClient\Parameters\UseBasicAuth
(DWORD)
```

If you delete the registry key or set it to 0, the behavior reverts to the default, or disabling the use of BasicAuth.

**WinINet: Disabling Basic Authentication over a clear channel**
**Detailed description**
Because the DAVRdr is part of the remote file system stack, a computer is open to attack whenever an attempt is made to remotely access files. Although the threat to other applications that use the Internet APIs is less severe than it is for the DAVRdr, a similar attack is possible whenever an application (or the user) attempts to access a URL. For this reason, WinINet is exposing the mechanism by which the DAVRdr disables BasicAuth to other users of the Internet APIs.

There are two ways to block the use of Basic Authentication over clear (or unencrypted) channels:

- Create the following registry key and set it to a non-zero value.

  `HKEY_CURRENT_USER\SOFTWARE\Microsoft\Windows \CurrentVersion\ Internet Settings\DisableBasicOverClearChannel` (DWORD)

  This prevents WinINet from attempting to use BasicAuth unless the channel is secured (using HTTPS or SSL).

- The application can disable the use of BasicAuth for its connections by setting the AUTH_FLAG_DISABLE_BASIC_CLEARCHANNEL flag (0x4) in the value supplied in the call to **InternetSetOption** using INTERNET_OPTION_AUTH_FLAGS.

**Why is this change important?**
Users can log on to WebDAV servers for remote file access without fear of transmitting their password in the clear.

**What threats does it help mitigate?**
Imagine a corporate user who routinely accesses the Web site http://www.contoso.com/sales. While outside the corporation on a public network, the user attempts to access that site using Internet Explorer. Since the laptop is outside the corporation, the request should fail with a "Server not found" message. An attacker can run on that same public network with a computer that spoofs WINS requests, returning a pointer to itself in response to any WINS lookup. The laptop will then try to send the HTTP request to load the page from the rogue server. If the rogue server responds with BasicAuth as the authentication method, the laptop responds to the user, asking for his or her credentials. It identifies the site http://www.contoso.com/sales, leading the user to believe the request is legitimate. If the user enters his or her user name and password, the client will transmit that information in the clear, and the attacker thus gains access to that user's login information. In particular, there is no indication to the user that the channel is insecure, or that the laptop will transmit the user name and password in the clear.

**What works differently?**
By default, there is no change in behavior for WinINet applications (except for the DAVRdr as noted above). If this setting is disabled, the user is unable to connect to HTTP servers that only support Basic Authentication.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**
**WebDAV Redirector Settings**

| Setting name | Location | Previous default value (if applicable) | Default values | Possible values |
|---|---|---|---|---|
| UseBasicAuth | HKEY_LOCAL_MACHINE\System \CurrentControlSet \Services \WebClient \Parameters \UseBasicAuth | Not applicable. | Key doesn't exist.<br><br>(BasicAuth disabled for DAVRdr) | 0, non-zero |
| DisableBasicOverClearChannel | HKCU\SOFTWARE \Microsoft \Windows \CurrentVersion \Internet Settings \DisableBasicOverClearChannel | Not applicable. | Key doesn't exist. (BasicAuth enabled for everything else) | 0, non-zero |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
No changes are necessary. Developers who write applications that use the Internet APIs and want to disable BasicAuth, such as the DAVRdr, can add a call to **InternetSetOptions()**.

# Windows Firewall in Windows Server 2003 Service Pack 1

Applies To: Windows Server 2003 with SP1

**What does Windows Firewall do?**
Windows Firewall (previously called Internet Connection Firewall or ICF) is a software-based, stateful filtering firewall for Microsoft Windows XP and Microsoft Windows Server 2003. Windows Firewall provides protection for computers that are connected to a network by preventing unsolicited incoming traffic through TCP/IP version 4 (IPv4) and TCP/IP version 6 (IPv6). Configuration options include:

- Configuring and enabling port-based exceptions
- Configuring and enabling program-based exceptions
- Configuring basic ICMP options
- Logging dropped packets and successful connections

Windows Firewall in Windows Server 2003 Service Pack 1 is not enabled by default when the update is applied to your server. It will only be enabled in the following situations:

- If Internet Connection Sharing was previously enabled.
- If Internet Connection Firewall was previously enabled.
- If the server is a new installation of Windows Server 2003 with Service Pack 1 (also known as a slipstream installation).

The best resources to help you fully understand how Windows Firewall works and how it can be used in your environment are the Windows Firewall Information and Help topics on the Windows Server 2003 Tech Center Web site at http://go.microsoft.com/fwlink/?LinkId=48911 and the Windows Firewall Operations Guide on the Windows Server 2003 TechCenter Web site at http://go.microsoft.com/fwlink/?LinkId=48912.

| **Note** |
| --- |
| If you decide to use Windows Firewall with your server, it is strongly recommended that you restart your servers after turning on and configuring the firewall. Windows Firewall in Windows Server 2003 with Service Pack 1 now supports application exceptions and needs to maintain the state of those applications. As a result, any applications or services that you add to the firewall exceptions list that were running prior to the firewall starting will still fail. After the server is restarted, the firewall will be running before any of the applications on the exceptions list and will be able to successfully maintain the state of the applications and handle them correctly. |

**Who does this feature apply to?**
This feature applies to:

- All computers that are connected to a network, including the Internet.
- All programs (applications and services) that listen on the network.
- All programs that do not work with stateful filtering.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Integration of Internet Connection Firewall and IPv6 Internet Connection Firewall into Windows Firewall**
**Detailed description**
The version of Internet Connection Firewall that was introduced with Windows XP filtered only IPv4 traffic. IPv6 Internet Connection Firewall was introduced with the Advanced Networking Pack for Windows XP. With Windows Server 2003 Service Pack 1, Internet Connection Firewall and IPv6 Internet Connection Firewall are integrated into a single component called

Windows Firewall.

With this change, any configuration change applies to both IPv4 and IPv6 traffic. For example, when a static port is opened, it is opened for both IPv4 and IPv6 traffic.

**Why is this change important?**
This allows for easier configuration management and application compatibility.

**What works differently?**
The Internet Connection Firewall service is removed from the system and replaced with the Windows Firewall service, which filters both IPv4 and IPv6 traffic. All firewall APIs are superseded by new APIs introduced with Windows Server 2003 Service Pack 1.

**How do I resolve these issues?**
For more information, see "Do I need to change my code to work with Windows Server 2003 Service Pack 1?" later in this document.

**On-by-default for new installations of Windows Server 2003 that include a service pack**
**Detailed description**
Windows Firewall is on by default only during new installations of Windows Server 2003 that include a service pack (also known as a slipstream release). Windows Firewall provides network protection while users update their system with the latest patches using the new Post-Setup Security Updates feature. As soon as the updates are finished the firewall is turned off unless it was explicitly enabled.

If a server running Windows Server 2003 is updated or upgraded to Service Pack 1 the firewall is off by default and the Post Setup Security Updates feature is not used.

**Why is this change important? What threats does it help mitigate?**
By enabling Windows Firewall by default on new installations, the computer has more protection from many network-based attacks while it is being set up and configured. For example, if Windows Firewall had been enabled by default, the MSBlaster attack would have been greatly reduced in impact, whether or not users had installed the relevant updates on their computers.

**What works differently?**
After a new installation of a slipstream version of Windows Server 2003 with Service Pack 1, Windows Firewall is enabled by default and incoming traffic is blocked until after Post-Setup Security Updates have been completed. This might create application or service incompatibility if the application or service does not work with stateful filtering by default.

**How do I resolve these issues?**
Complete Post-Setup Security Updates, which will automatically turn off the firewall, before proceeding with any other server configuration tasks.

It is also possible to configure the firewall to work with applications or services you need to use, if you don’t want to complete Post-Setup Security Updates until a later time.

**Configuration by the Security Configuration Wizard**
**Detailed description**
The recommended means of turning on Windows Firewall and performing its initial configuration for Windows Server 2003 with Service Pack 1 is to use the Security Configuration Wizard (SCW). SCW will automatically turn on Windows Firewall and create the appropriate settings based on the needs of your server. For more information about SCW, see "Security Configuration Wizard", in this document.

**Why is this change important?**
Some server components and applications should not be used with Windows Firewall or should be used in very specific configurations. SCW has been designed to help you determine the recommended settings for the Windows Firewall based on your environment.

**Boot-time security**
**Detailed description**
In earlier versions of Windows, there is a period of time between when the network stack comes up and when Internet Connection Firewall provides protection. This results in the ability for a packet to be received and delivered to a service without Internet Connection Firewall providing filtering and potentially exposes the computer to vulnerabilities. This was due to the firewall driver not starting to filter until the firewall user-mode service was loaded and had applied appropriate policy settings. The firewall service has a number of dependencies, which causes the service to wait until those dependencies are cleared before it pushes the policy down to the driver. This time period is based upon the speed of the computer.

In Windows Server 2003 Service Pack 1, the IPv4 and IPv6 firewall drivers have a static rule to perform stateful filtering. This static rule is called a boot-time policy. This allows the computer to perform basic networking functions such as DNS and DHCP

and communicate with a domain controller to obtain policy settings. After the Windows Firewall service is running, it loads and applies the runtime policy settings. The boot-time policy cannot be configured.

There is no boot-time security if the Windows Firewall service (which is listed as Windows Firewall/Internet Connection Sharing (ICS) in the Service Control Manager) is set to either Manual or Disabled.

**Why is this change important? What threats does it help mitigate?**
With this change, the computer is open to fewer attacks during startup and shutdown.

**What works differently?**
If the Windows Firewall service fails to start, boot-time security remains in effect. This means that all incoming connections are blocked. In this case, an administrator will not be able to remotely troubleshoot the issue because all the ports will be closed, including the port used by Remote Desktop.

If a service attempts to start before the firewall service a "race condition" might result. If a necessary service is blocked by this condition you will need to disable Windows Firewall.

**How do I resolve these issues?**
To turn off boot-time security, stop the Windows Firewall/Internet Connection Sharing (ICS) service and set its startup type to either Manual or Disabled.

If the computer is in boot-time security mode because the firewall service has not started, an administrator must log on to the computer, resolve the cause of the failure, and then manually start the firewall service.

**Running in safe mode (Safe mode firewall)**
**Detailed description**
The firewall state is maintained when the server is started in safe mode.

**Why is this change important?**
With this change your computer is less vulnerable to attack when starting in safe mode with network connectivity.

**What works differently?**
In previous versions, Internet Connection Firewall was not available when running in safe mode.

**Global configuration**
**Detailed description**
In earlier versions of Windows, Internet Connection Firewall was configured on a per-interface basis. This meant that each network connection had its own set of firewall settings, for example, one set of settings for wireless, another set of settings for Ethernet. This made it difficult to synchronize firewall settings between connections. Additionally, new connections would not have any of the configuration changes that had been applied to the existing connections. Non-standard network connections, such as those created by proprietary dialers (for instance, ISP-configured dial-up networking connections) could not be protected.

With global configuration in Windows Firewall, whenever a configuration change occurs, it automatically applies to all network connections in the Network Connections folder, as well as any non-Microsoft dialers. When new connections are created, the configuration is applied to them as well. Configuration can still be performed on a per-interface basis. Non-standard network connections will have only global configuration. Configuration changes also apply to both IPv4 and IPv6.

**Why is this change important?**
Having global configuration makes it easier for users to manage their firewall policy across all network connections and enables configuration through Group Policy. It also allows you to enable applications to work on any interface with a single configuration option.

**What works differently?**
In earlier versions of Windows Server, firewall configuration was on a per-interface basis. In Windows Server 2003 Service Pack 1, the configuration is global and applies to both IPv4 and IPv6.

**How do I resolve these issues?**
If your application or service requires static openings to work, you should open the ports globally, as described later in this topic, in "Do I need to change my code to work with Windows Server 2003 Service Pack 1?"

**Audit logging**
**Detailed description**
Audit logging enables you to track changes that are made to Windows Firewall settings and to see which applications and services asked your computer to listen on a port. After audit logging is enabled, audit events will be logged in the security event log. Audit logging can be enabled on client computers running Windows XP Service Pack 2 and servers running Windows Server 2003 Service Pack 1. You can use the following procedure to enable audit logging on your computer.

**To enable audit logging**

1. Log on using an account that is a local administrator.
2. Click **Start**, click **Control Panel**, and then click **Administrative Tools**.
3. In **Administrative Tools**, double-click **Local Security Policy** to open the **Local Security Settings** console.
4. In the console tree of the **Local Security Settings** console, click **Local Policies**, and then click **Audit Policy**.
5. In the details pane of the **Local Security Settings** console, double-click **Audit policy change**. Select **Success** and **Failure**, and then click **OK**.
6. In the details pane of the **Local Security Settings** console, double-click **Audit process tracking**. Select **Success** and **Failure**, and then click **OK**.

You can also enable audit logging for multiple computers in an Active Directory directory service domain using Group Policy by modifying the **Audit policy change** and **Audit process tracking** settings at **Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy** for the Group Policy objects in the appropriate domain system containers.

After audit logging is enabled, you can use the Event Viewer snap-in to view audit events in the security event log.

Windows Firewall uses the following event IDs:

- 848 - Displays the startup configuration of Windows Firewall.
- 849 - Displays an application exception configuration.
- 850 - Displays a port exception configuration.
- 851 - Displays a change made to the application exceptions list.
- 852 - Displays a change made to the port exceptions list.
- 853 - Displays a change made to the Windows Firewall operation mode.
- 854 - Displays a change made to Windows Firewall logging settings.
- 855 - Displays a change made to ICMP settings.
- 856 - Displays a change made to the Prohibit unicast response to multicast or broadcast requests setting.
- 857 - Displays a change made to the Remote Administration setting.
- 858 - Displays the application of Windows Firewall Group Policy settings.
- 859 - Displays the removal of Windows Firewall Group Policy settings.
- 860 - Displays a change made to a different profile.
- 861 - Displays an application attempting to listen for incoming traffic.

**Why is this change important?**
Auditing the activity of Windows Firewall is part of a defense in depth strategy because it can be used to alert you to malicious software that is attempting to modify firewall settings. Auditing also generally helps administrators determine the network needs of their applications and design an appropriate policy for deployment to large numbers of users.

**Traffic scoping for exceptions**
**Detailed description**
ICF allowed excepted traffic to come from any IPv4 address. With Windows Firewall in Windows Server 2003 with Service

Pack 1, you can also configure an exception to allow incoming traffic only from addresses that are directly reachable by selecting the **My network (subnet) only** scope option (based on entries in the IPv4 and IPv6 routing table), or from specific IPv4 address ranges by selecting the **Custom list** scope option.

For computers in a workgroup, some exceptions are restricted to locally reachable addresses by default. These exceptions are those needed for file and printer sharing and the UPnP framework. Additionally, when these exceptions are opened for locally reachable addresses on an Internet Connection Sharing (ICS) host, the exceptions will not be opened on the ICS public interface. If you enable these exceptions for all possible addresses they will be opened on the ICS public interface, which is not recommended. When the File and Printer Sharing built-in exception is enabled with the NetShare application programming interface (API), with the Network Setup Wizard, or through the Windows Firewall user interface, incoming file and printer sharing connection requests can come only from directly reachable addresses by default.

If you are enabling Windows Firewall on a server that is already configured for file and printer sharing, the File and Printer Sharing exception might be enabled automatically. It is recommended that you apply the locally reachable addresses restriction to any exception that is used for communicating on the local network. It can be done programmatically at the command line using Windows Firewall Netsh Helper, or by clicking **Windows Firewall** in **Control Panel**.

**Note**

As a best practice, identify custom scopes with specific addresses or subnets for the exceptions that you specify for Windows Firewall.

When you configure and enable an exception, you are instructing Windows Firewall to allow specific unsolicited incoming traffic sent from the specified scope (from any address, from an address that can be reached directly, or from a custom list). For any scope, enabling an exception makes the computer accessible to attacks based on incoming unsolicited traffic from computers that are assigned the allowed addresses and from malicious computers that spoof traffic. There is no way to prevent spoofed attacks from the Internet on connections assigned public IPv4 addresses except by disabling the exception. Therefore, you should try to configure scope options so that the number of computers that are allowed to send unsolicited traffic through an exception is kept to a minimum. This will reduce, but not eliminate, the likelihood of a spoof attack.

If your organizational security policy requires you to ensure that no one outside your network can access a resource, then you should consider using an approach such as IPsec that supports network-level peer authentication, data origin authentication, data integrity, data confidentiality (encryption), and replay protection.

**Why is this change important? What threats does it help mitigate?**
Some applications need to communicate only with other computers on the local network and not computers on the Internet. Configuring Windows Firewall to allow only traffic from locally reachable addresses or from specific address ranges corresponding to locally attached subnets restricts the set of addresses from which unsolicited incoming traffic can be accepted. This mitigates, but does not eliminate, attacks that can occur for enabled exceptions.

**What works differently?**
When the File and Printer Sharing or the UPnP framework built-in exception is enabled using the Control Panel on a computer that is a member of a workgroup, the locally reachable addresses scope is applied to the ports opened. If an application or service also uses these ports, it will be able to communicate only with other nodes that are assigned locally reachable addresses. However, if the computer is a member of a domain, the global scope is applied.

If these exceptions are enabled using an API call or using Netsh.exe instead of from Control Panel, the default scope setting is locally reachable addresses, regardless of whether the computer is a member of a workgroup or a domain.

**How do I resolve these issues?**
If your application or service does not work with this type of restriction, you should open the port for any computer, as described in "Do I need to change my code to work with Windows Server 2003 Service Pack 1?" later in this document.

**Command-line support**
**Detailed description**
The Windows Firewall Netsh Helper was added to Windows XP in the Advanced Networking Pack. This helper applied only to IPv6 Windows Firewall. With Windows Server 2003 Service Pack 1, the structure and syntax of the helper changed and expanded to include support for configuring IPv4 as well. With the Netsh Helper, you can fully configure Windows Firewall, including:

- Configure the default state of Windows Firewall (Off, On, On with no exceptions).
- Configure the ports that must be open.

- Configure the ports to enable global access or to restrict access to the local subnet.
- Set ports to be open on all interfaces or only on a specific interface.
- Configure the logging options.
- Configure the Internet Control Message Protocol (ICMP) handling options.
- Configure and enable program-based exceptions.

Windows Firewall configuration and status information can be retrieved at the command line by using the Netsh.execontext: `firewall`.

To use this context, type `netsh firewall` at a command prompt, and then use additional Netsh commands as needed.

The following commands are useful for gathering firewall status and configuration information and can be useful for troubleshooting the operation of your firewall:

- `Netsh firewall show state`
- `Netsh firewall show config`

The following commands can be used to modify the configuration of Windows Firewall.

| **Command** | **Description** |
|---|---|
| add allowedprogram | Used to add excepted traffic by specifying the program's file name. |
| set allowedprogram | Used to modify the settings of an existing allowed program exception. |
| delete allowedprogram | Used to delete an existing allowed program exception. |
| set icmpsetting | Used to specify allowed ICMP traffic. |
| set logging | Used to specify logging options. |
| set notifications | Used to specify whether notifications to the user when programs try to open ports are enabled. |
| set opmode | Used to specify the operating mode of Windows Firewall either globally or for a specific connection (interface). |
| add portopening | Used to add excepted traffic by specifying a TCP or UDP port. |
| set portopening | Used to modify the settings of an existing open TCP or UDP port exception. |
| delete portopening | Used to delete an existing open TCP or UDP port exception. |
| set service | Used to enable or drop RPC and DCOM traffic, Remote Desktop, file and printer sharing, and UPnP traffic. |
| reset | Resets firewall configuration to Default. This provides the same functionality as the **Restore Defaults** button in Control Panel/Windows Firewall. |

The following table details the show commands supported for Windows Firewall.

| **Command** | **Description** |
|---|---|
| show allowedprogram | Displays the allowed programs. |

| | |
|---|---|
| show config | Displays the detailed local configuration information. |
| show currentprofile | Displays the current profile. |
| show icmpsetting | Displays the ICMP settings. |
| show logging | Displays the logging settings. |
| show notification settings | Displays the current settings for notifications. |
| show opmode | Displays the operational mode for profiles and interfaces. |
| show portopening | Displays the excepted ports. |
| show service | Displays the services exception settings. |
| show state | Displays the current state information. |

You can compare the output from these commands with the output from the `netstat -ano` command to identify the programs that may have listening ports open and that do not have corresponding exceptions in the firewall configuration.

**Why is this change important?**
Providing a command-line interface provides administrators with a method to configure Windows Firewall without going through the graphical user interface. The command-line interface can be used in logon scripts and remote management.

**What works differently?**
Any script that was created with the Netsh Helper that was made available with the Advanced Networking Pack for Windows XP no longer works and must be updated.

**How do I resolve these issues?**
Update any scripts you might have so that they include the new firewall context and syntax.

**"On with no exceptions" operational mode**
**Detailed description**
Windows Firewall can be configured for exceptions to allow specific unsolicited incoming traffic during normal use. Typically, this is because key scenarios, like file and printer sharing, must be enabled. If a security issue is discovered in one or more of the listening services or applications that are running on the computer, it may be necessary for the computer to switch into a client-only mode, which is called "On with no exceptions." Switching into this client-only mode configures Windows Firewall to prevent all unsolicited incoming traffic without having to reconfigure the firewall.

When in this mode, all exceptions are temporarily disabled and any existing connections are dropped. Any application interface that calls into Windows Firewall to create an exception is allowed and the requested firewall configuration is stored, but it is not enabled until the operational mode switches back to normal operation. All listen requests by applications are also ignored and notification dialogs are not displayed, effectively blocking the application from listening on a port while the computer is in this operational mode.

**Why is this change important?**
When a network system is under attack by viruses, worms, and other attackers, the attacker looks for services to exploit. The “On with no exceptions” operational mode provides a way for you to quickly lock-down your system in the event of an attack so that valid exceptions cannot be used to circumvent the protection provided to your computer by Windows Firewall.

**What works differently?**
When in this operational mode, the computer cannot listen for requests that originate from the network. Any existing incoming connections are terminated. Outgoing connections are the only connections that succeed.

**How do I resolve these issues?**
When in this operational mode, it is expected that some functionality will fail because of the strict network security in place. You can restore functionality by returning the operational mode to On. This action should be performed by the user only after the threat has been identified and mitigated, because the security of the computer is reduced by performing this action.

**Program-based exceptions**
**Detailed description**
Some programs (applications or services) act as both network clients and servers. When they act as servers, they must allow unsolicited incoming traffic, because they do not know in advance who the peer will be.

In earlier versions of Windows, a program needed to call the firewall APIs to enable the necessary listening ports to be open. This proved difficult in peer-to-peer situations when the port was not known in advance. It was up to the program to close the

port again after communication was completed. If the program terminated unexpectedly this could result in unnecessary open ports in the firewall.

An additional issue with the previous method of opening firewall ports was that ports could be opened only if programs were running in the security context of a local administrator. This violated the basic information security principle of least privilege by requiring programs to run in an administrative context, rather than only with the minimum necessary privileges.

In Windows Server 2003 with Service Pack 1, a program that needs to listen to the network can be added to the Windows Firewall exceptions list. If a program is enabled on the Windows Firewall exceptions list, Windows Firewall opens and closes the necessary listening ports automatically, regardless of the program's security context. For more information about adding programs to the Windows Firewall exceptions list, see "How do I resolve these issues?" later in this document.

Programs that work with stateful filtering do not need to be placed on the Windows Firewall exceptions list. Only administrators can add a program to the Windows Firewall exceptions list.

**Why is this change important? What threats does it help mitigate?**
When a program is on the Windows Firewall exceptions list, only the necessary ports are opened, and they are opened only for the duration that the program is listening on those ports.

**What works differently?**
If a program needs to listen on the network, it must be enabled on the Windows Firewall exceptions list. If it is not, then the necessary port in Windows Firewall is not opened and the program will not be able to receive unsolicited inbound traffic.

**How do I resolve these issues?**
A program can be placed on the Windows Firewall exceptions list in five ways:

1. **Programmatically**. It is recommended that independent software vendors (ISVs) place their programs on the Windows Firewall exceptions list during installation. For more information about how to programmatically add a program to the exceptions list, see "Do I need to change my code to work with Windows Server 2003 Service Pack 1?" later in this section.

2. **Command-line interface**. This method can be used by IT administrators who manage Windows XP and Windows Server 2003 systems using scripts or other command-line tools.

3. **Group Policy settings**. This method can be used by IT administrators to add the program to the exceptions list through Group Policy.

4. **Windows Firewall notification message**. A user with Administrator rights can interact with the Windows Firewall notification message and add the application to the exceptions list.

   When an application performs a TCP listen or UDP bind to a non-wildcard port, the network stack passes the application name and port to Windows Firewall. Windows Firewall looks up the application name on the exceptions list. If the application is on the exceptions list and enabled, then the corresponding port is opened in the firewall. If the application is on the exceptions list and disabled, then the corresponding port is not opened. If the application is not on the exceptions list, then users are asked to make a choice. If the users have administrative rights, they can:

   - Unblock the application to allow it to listen on the network. It is added to the exceptions list as Enabled and the port is opened.

   - Block the application from listening on the network. It is added to the exceptions list as Disabled and the port is not opened.

   - Choose to be asked again later. The application is not added to the exceptions list and the port is not opened.

   If the user does not have administrative rights, the user is notified that the application is not allowed to listen on the network and that an Administrator must enable the program exception. If the user selects the **Do not ask me again** check box, the application is listed in the exceptions list as Disabled.

**Note**

| |
|---|
| Notification messages can only be used with applications. They cannot be used with services. |

5. **Manual configuration**. Administrators can decide to enable a program manually in the Windows Firewall control panel by selecting it from a list that is populated from the list of programs in the **Start** menu or by browsing for the program.

**Multiple profiles**
**Detailed description**
Multiple profile support in Windows Firewall allows you to create two sets of firewall policy settings: one for when the computer is connected to a managed network and one for when the computer is not. You can specify settings that are less strict when the computer is connected to the corporate network to enable line-of-business applications to work. You can also have more aggressive security policy settings that will be enforced when the computer leaves the corporate network, which helps to protect mobile users.

| **Note** |
|---|
| Multiple profiles for Windows Firewall apply only to computers that are joined to an Active Directory domain. Computers that are in a workgroup use only one profile. |

**Why is this change important? What threats does it help mitigate?**
For a mobile computer, it is desirable to have more than one firewall configuration. Often, a configuration that is safe on a corporate network is likely to be susceptible to attack on the Internet. Therefore, being able to have ports opened on the corporate network and not on other networks is critical to ensuring that only the necessary ports are exposed at any given time.

**What works differently?**
If an application needs to be listed in the Windows Firewall exceptions list in order to work correctly, it might not work on both networks as the two profiles might not have the same set of policy settings. For an application to work on all networks, it must be listed in both profiles. For more information about the Windows Firewall exceptions list, see the earlier section.

**How do I resolve these issues?**
If the computer is joined to a domain, you must ensure that the application is listed in both firewall configurations. Consider creating exceptions though the command-line interface or Group Policy as you will only have access to the currently running profile through **Windows Firewall** in **Control Panel**.

**RPC support for System Services**
**Detailed description**
In earlier versions of Windows, Internet Connection Firewall blocked remote procedure call (RPC) communication. While Internet Connection Firewall could be configured to allow network traffic to the RPC Endpoint Mapper, the port that an RPC server used was unknown and the application would still fail.

Many enterprise applications and components fail if RPC is not allowed to communicate over the network. Some examples include, but are not limited to, the following:

- Remote administration, such as the Computer Management feature and the **Select User, Computers, and Groups** dialog box, which is used by many applications
- Remote Windows Management Instrumentation (WMI) configuration
- Scripts that manage remote clients and servers

RPC opens several ports and then exposes many different servers on those ports. It then requests that Windows Firewall create associated exceptions for these ports. If Windows Firewall is configured to allow such requests, the required ports will be opened for as long as RPC needs the exception (similar to a program exception).

**Why is this change important? What threats does it help mitigate?**
In order to enable remote administration scenarios, many enterprise-wide deployments require that the system services that use RPC work with Windows Firewall by default.

**What works differently?**
By default, RPC does not function through Windows Firewall. All system services that use RPC are affected. However, Windows Firewall can be configured to allow RPC to work for these services using the remote administration setting. This setting also

enables exceptions for the RPC Endpoint Mapper (TCP 135), SMB over TCP (TCP 445), and ICMP echo requests.

**How do I resolve these issues?**
See "Do I need to change my code to work with Windows Server 2003 Service Pack 1?" later in this document.

**Restore defaults**
**Detailed description**
Previously, there was no way for a user to reset the configuration of Internet Connection Firewall (ICF). Over time, the firewall might be configured to allow unsolicited incoming traffic to ports no longer used by other applications. This might make it difficult for the user to easily and quickly go back to a default configuration.

This option enables the user to restore Windows Firewall settings to their original defaults. In addition, the Windows Firewall defaults can be modified by original equipment manufacturers (OEMs) and businesses to provide custom default configuration options.

**Why is this change important?**
This option allows end-users to restore their Windows Firewall settings to the out-of-the-box defaults.

**What works differently?**
No functional changes in Windows Firewall result from this addition. However, use of this feature disables Internet Connection Sharing and Network Bridge.

**Unattended setup support**
**Detailed description**
In earlier versions of Windows, it was not possible to configure Internet Connection Firewall during installation. This made it difficult for OEMs and businesses to preconfigure Internet Connection Firewall before distributing a computer to their end users. In Windows Server 2003 with Service Pack 1, you can configure the following options of Windows Firewall through unattended setup:

- Operational mode
- Applications on the Windows Firewall exception list
- Static ports on the exception list
- ICMP options
- Logging options

**Why is this change important?**
A method to preconfigure Windows Firewall allows Windows resellers and large enterprises more flexibility and customization options for Windows Firewall.

**What works differently?**
This feature adds configuration flexibility to Windows Firewall. No functional changes in Windows Firewall result from this addition.

The syntax used to enable or disable ICF in an unattend script has been replaced with the new syntax for Windows Firewall.

The sections of the Unattend.txt file for Windows Firewall configuration consist of the following:

- `[WindowsFirewall]`

  A required section that defines which profiles to use and Windows Firewall log file settings.

- `[WindowsFirewall.profile_name]`

  The domain profile section, `[WindowsFirewall.Domain]`, contains settings for when a computer is connected to a network that contains domain controllers for the domain of which the computer is a member. The standard profile, `[WindowsFirewall.Standard]`, contains settings for when a computer is not connected to a network that contains domain controllers for the domain of which the computer is a member. If you do not want Windows Firewall to be used you can specify `Profiles = WindowsFirewall.TurnOffFirewall`

The `[WindowsFirewall.profile_name]` section is a user-defined section that is referenced by the `[WindowsFirewall]` section to make changes to Windows Firewall's default configuration, including programs, services, ports, and ICMP settings.

- `[WindowsFirewall.program_name]`

  A user-defined section that adds a program to the Windows Firewall exceptions list.

- `[WindowsFirewall.service_name]`

  A user-defined section that adds a predefined service to the Windows Firewall exceptions list (such as file and printer sharing, UPnP framework, Remote Desktop service, and Remote Administration).

- `[WindowsFirewall.portopening_name]`

  A user-defined section that adds a port to the Windows Firewall exceptions list.

- `[WindowsFirewall.icmpsetting_name]`

  A user-defined section that adds ICMP message types to the Windows Firewall exceptions list.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
**Enhanced multicast and broadcast support**
**Detailed description**
Multicast and broadcast network traffic differs from unicast traffic because the response comes from an unknown host. As such, stateful filtering prevents the response from being accepted. This stops a number of scenarios from working, ranging from streaming media to discovery.

To enable these scenarios, Windows Firewall will allow a unicast response for three seconds from any directly reachable source address on the same port from which the multicast or broadcast traffic originated.

**Why is this change important? What threats does it help mitigate?**
This allows applications and services that use multicast and broadcast for communicating to work without either the user or application/service needing to alter the firewall policy. This is important for things like NETBIOS over TCP/IP, so that sensitive ports such as port 135 are not exposed.

**What works differently? Are there any dependencies?**
In Windows Server 2003, Internet Connection Firewall statefully filtered multicast and broadcast traffic, which required the user to manually open the port to receive the response. In Windows Server 2003 Service Pack 1, Windows Firewall accepts the response to the multicast or broadcast traffic without additional configuration.

**Updated user interface**
**Detailed description**
The firewall user interface is updated in Windows Server 2003 Service Pack 1 to accommodate the new configuration options and the integration of IPv6 Internet Connection Firewall. The new Windows Firewall interface provides the user with the ability to change the operational states, the global configuration, logging options, and ICMP options.

The primary entry to the user interface has been moved from the **Properties** dialog box of the connection to a Control Panel icon. A link from the old location is still provided. Additionally, Windows Server 2003 Service Pack 1 creates a link from the Network Connections folder.

**Why is this change important?**
The functionality that is added in Windows Server 2003 Service Pack 1 required updates to the user interface.

**What works differently?**
The user interface is moved from the **Advanced** tab of the network connection's **Properties** dialog box to a specific Windows Firewall icon in Control Panel.

**New Group Policy support**
**Detailed description**
In earlier versions of Windows, Internet Connection Firewall had a single Group Policy object (GPO): Prohibit Use of Internet Connection Firewall on your DNS domain network. With Windows Server 2003 Service Pack 1, every global configuration

option can be set through Group Policy. Examples of the new configuration options available include:

- Define program exceptions
- Allow local program exceptions
- Allow ICMP exceptions
- Prohibit notifications
- Allow file and printer sharing exception
- Allow logging

Each of these objects can be set for both the corporate and standard profile. For more information about Group Policy options, see "Deploying Windows Firewall Settings for Microsoft Windows XP with Service Pack 2" in the Microsoft Download Center at http://go.microsoft.com/fwlink/?linkid=23277. This document also covers developments in Windows Server 2003 Service Pack 1.

**Why is this change important?**
It is important for administrators to centrally manage Windows Firewall policy settings to enable applications and scenarios to work in the corporate environment.

**What works differently?**
The IT administrator can now decide the default Windows Firewall policy set. This can either enable or disable applications and scenarios. This allows more control, but the policies do not change the underlying functionality of Windows Firewall.

**What settings are added or changed in Windows Server 2003 Service Pack 1?**

| Setting name | Location | Previous default value | Default value | Possible values |
|---|---|---|---|---|
| Protect all network connections | (Group Policy object) Computer Configuration \Administrative Templates \Network\Network Connections\ \Windows Firewall | Not applicable | Not configured | Enabled<br>Disabled |
| Do not allow exceptions | (Group Policy object) Computer Configuration \Administrative Templates \Network\Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br>Disabled |
| Define program exceptions | (Group Policy object) Computer Configuration \Administrative Templates \Network\Network Connections \Windows Firewall | Not applicable | Not configured | Enabled [Program path] [Scope]<br>Disabled |
| Allow local program exceptions | (Group Policy object) Computer Configuration \Administrative Templates \Network\Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br>Disabled |
| Allow remote administration exception | (Group Policy object) Computer Configuration \Administrative Templates \Network\Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br>Disabled |
| Allow file and printer sharing exception | (Group Policy object) Computer Configuration \Administrative Templates \Network\Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br>Disabled |

| | | | | |
|---|---|---|---|---|
| Allow ICMP exceptions | (Group Policy object) Computer Configuration \Administrative Templates \Network\Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br><br>Once enabled, select which of the following message types to allow:<br><br>[Allow outbound destination unreachable]<br><br>[Allow outbound source quench]<br><br>[Allow redirect]<br><br>[Allow inbound echo request]<br><br>[Allow outbound time exceeded]<br><br>[Allow outbound parameter problem]<br><br>[Allow inbound timestamp request]<br><br>[Allow inbound mask request]<br><br>[Allow outbound packets too bug]<br><br>Disabled |
| Allow remote desktop exception | (Group Policy object) Computer Configuration \Administrative Templates \Network\Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br><br>Disabled |
| Allow UPnP framework exception | (Group Policy object) Computer Configuration \Administrative Templates \Network\Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br><br>Disabled |
| Prohibit notifications | (Group Policy object) Computer Configuration \Administrative Templates\Network \Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br><br>Disabled |
| Allow logging | (Group Policy object) Computer Configuration \Administrative Templates\Network \Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br><br>Disabled |
| Prohibit unicast response to multicast or broadcast requests | (Group Policy object) Computer Configuration \Administrative Templates\Network \Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br><br>Disabled |
| Define port exceptions | (Group Policy object) Computer Configuration \Administrative Templates\Network \Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br><br>Disabled |
| Allow local port exceptions | (Group Policy object) Computer Configuration \Administrative Templates\Network \Network Connections \Windows Firewall | Not applicable | Not configured | Enabled<br><br>Disabled |

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**

We recommend that you use the Security Configuration Wizard to configure Windows Firewall for use with Windows

Server 2003 Service Pack 1. SCW is designed to accommodate the requirements of different server roles and workloads and configure the firewall settings correctly. If you are going to manually configure your firewall settings, review the following information for how your applications might be affected.

**Outbound connections**
**Description**
For typical consumer and office computers, the computer is a client on the network. Software on the computer connects out to a server (an outbound connection) and gets responses back from the server. Windows Firewall allows all outbound connections, but applies rules to the types of communication that are allowed back into the computer.

Some examples of tasks involving Microsoft applications that might work this way include:

- Surfing the Web using Microsoft Internet Explorer.
- Checking e-mail in Outlook Express.
- Chatting in MSN Messenger or Windows Messenger.

**Action Required**
None. Windows Firewall will automatically allow all outbound connections, regardless of the program and the user context.

| **Note** |
| --- |
| When a computer initiates a TCP session request to a target computer, it will accept a response only from that target computer.<br><br>When the computer sends UDP packets, Windows Firewall allows UDP responses to the port from which the UDP packets were sent from any IP address for approximately 90 seconds.<br><br>Unicast responses to multicast and broadcast traffic are allowed through Windows Firewall for three seconds if the responses are to the port from which the traffic was sent and are from IP addresses on the same subnet as the computer. A setting in the firewall controls this behavior, which is enabled by default. |

**Unsolicited inbound connections for applications**
**Description**
This scenario covers an application that completes a listen operation on a TCP socket or successfully binds to a specific UDP socket through Winsock. For this scenario, Windows Firewall can automatically open and close ports as needed by the application.

Some examples of tasks involving Microsoft applications that might work this way include:

- Using audio and video in MSN Messenger or Windows Messenger.
- Transferring files in MSN Messenger or Windows Messenger.
- Hosting a multiplayer game.

**Action required**
If you are developing an application that needs to listen on a port (or ports) Microsoft requests that you update your code to ask the users to indicate whether they want to allow the application to open ports in the firewall:

- If the user consents to this, then the application can use the `INetFwAuthorizedApplication` API to add itself to the AuthorizedApplications collection as Enabled.
- If the user does not consent, then the application can use the `INetFwAuthorizedApplication` API to add itself to the AuthorizedApplications collection as Disabled.

When using the `INetFwAuthorizedApplication` API to add an application to the AuthorizedApplications collection, the following values are required:

- Image File Name. This is the file that calls Winsock to listen for network traffic. This must be a fully-qualified path, but it

might contain environment variables.

- Friendly Name. This is the description for the application that will be shown to users in the Windows Firewall user interface.

For more information about the `INetFwAuthorizedApplication` API, see "INetFwAuthorizedApplication" in the Microsoft Platform Software Development Kit (SDK) on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=32000.

Windows Firewall monitors Winsock to see when applications start and stop listening on ports. As a result, ports are automatically opened and closed for applications after their entries have been enabled in the Windows Firewall exceptions list. This means that no action is required by Winsock applications to actually open and close ports in the firewall.

| **Note** |
| --- |
| An application must be running in the context of a user with Administrator rights to add itself to the Windows Firewall exceptions list.<br><br>Ports are automatically opened and closed in the firewall for allowed Winsock applications, regardless of the user context in which the applications are running.<br><br>Applications should get user consent before adding themselves to the `INetFwAuthorizedApplications` collection.<br><br>Svchost.exe cannot be added to the `INetFwAuthorizedApplications` collection. |

**Inbound connections for services using fixed ports**
**Description**
While developers are advised to use the `INetFWAuthorizedApplication` APIs for all other scenarios, the use of global port APIs in Windows Firewall is recommended for services that listen on fixed ports. Because these ports are always open, there is minimal benefit to dynamically opening the ports. Instead, users gain the ability to customize the firewall settings for these fixed ports when the global port APIs are used.

Some examples of services that require inbound connections are:

- File and printer sharing.
- UPnP architecture.
- Remote Desktop.

**Action Required**
When a service needs to listen on a fixed port, it should ask the user whether it should allow the service to open ports in the firewall. Ideally this should be done when the service is being installed.

If the user consents to this, then the service should use the `INetFwOpenPort` API to add rules to Windows Firewall for the fixed port (or ports) needed by the service. These rules should be enabled.

If the user does not consent, then the service should still use the `INetFwOpenPort` API to add rules to Windows Firewall for the fixed port or ports needed by the service. These rules, however, should not be enabled.

When using the `INetFwOpenPort` API to add a port opening to Windows Firewall, the following values are required:

- Protocol. Specifies the network protocol that is used by the service, either TCP or UDP.
- Port. This is the number of the port to be opened.
- Friendly Name. This is the description for the port opening that will be shown to users in the Windows Firewall user interface.

For more information about the `INetFwOpenPort` API, see "InetFwOpenPort" in the Platform Software Development Kit on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=35316.

When a service is disabled, it should use the `INetFwOpenPort` API to close the static ports that it opened, whenever possible. This can be easily done if it is the only service that uses the ports. If the service potentially shares the ports with other services, however, it should not close the ports unless it can verify that none of the other services are using the ports.

An application must be running in the context of a user with Administrator rights to statically open ports in Windows Firewall.

**Note**

When statically opening ports through the `INetFw` API, a service should limit itself to traffic from the local subnet whenever possible.

Services should get user consent before statically opening ports in Windows Firewall. A service should never just automatically open ports without first warning the user.

**Inbound connections on RPC and DCOM ports for system services**

**Description**

Some system services require the use of RPC ports, either through DCOM or RPC directly, for inbound connections. Because of the significant security implications when opening RPC ports, these ports are handled as a special case, and developers should try to enable RPC for system services through Windows Firewall only when absolutely necessary.

**Action required**

Windows Firewall can be configured to enable the automatic opening and closing of RPC and DCOM ports for system services. By default, however, RPC will be blocked by Windows Firewall. This means that applications that use the RPC ports to transfer data to system services will need to configure Windows Firewall appropriately. When an application needs to enable this feature, it should ask the user whether it should allow the services to open ports in the firewall. Ideally, this should be done at installation time.

If the user consents to allowing the RPC ports to be opened, then the service should use the `INetFwRemoteAdminSettings` API to open the ports that are needed by the service.

If the user does not consent to allowing the RPC ports to be opened, then the application or service should not configure Windows Firewall to allow the RPC ports.

For more information about the `INetFwRemoteAdminSettings` API, see "INetFwAuthorizedApplication" on the MSDN Web site at http://go.microsoft.com/fwlink/?linkid=32000 and, in the table of contents, click "RemoteAddresses Property of InetFwAuthorizedApplication."

**Note**

To enable or disable the automatic opening of RPC ports in Windows Firewall, an application or service must be running in the context of a user with Administrator rights.

An application or service should try to allow the RPC ports through Windows Firewall only when absolutely necessary.

If the RPC ports are already allowed, then the application or service does not need to do anything in order to function correctly. You can determine which ports are already opened using the `IsPortAllowed` API.

The RPC ports setting works only for RPC servers that run in the context of local system, network service, or local service. Ports opened by RPC servers running in other user contexts will not be enabled through this setting. Instead, those RPC servers should use the Windows Firewall exceptions list.

# Windows Media Player

Applies To: Windows Server 2003 with SP1

**What does Windows Media Player do?**
Windows Media Player is an application that provides media content, plays media files, and helps organize your stored media files.

**Who does this feature apply to?**
All users who use Windows Media Player should be aware of the changes incorporated in Windows Server 2003 Service Pack 1.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Windows Media Player 10**
**Detailed description**
Windows Media Player 10 is installed with most versions of Windows Server 2003 Service Pack 1 (it is not included with the Itanium-based versions). This version of Windows Media Player includes security fixes and new functionality.

During Windows Server 2003 Service Pack 1 installation, if you select the option to archive files, you can remove Service Pack 1 later. If you perform a new installation of Windows Server 2003 with Service Pack 1, the operating system is replaced, and Windows Media Player 10 cannot be removed. For more information, see the Windows Media Knowledge Center at on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=32062.

**Why is this change important? What threats does it help mitigate?**
Windows Media Player 10 is the most secure version of Windows Media Player and it has been updated to work with the other security enhancements contained in Windows Server 2003 Service Pack 1.

.

**What works differently?**
If you uninstall Windows Server 2003 Service Pack 1, you might need to reacquire some licenses in order to play the content that you have previously licensed. This applies if you have upgraded your computer from a member of the Windows 2000 Server or Windows Server 2003 family to Windows Server 2003 Service Pack 1, because Windows Media Player 10 handles digital content licenses differently than earlier versions.

**How do I resolve these issues?**
To ensure that your existing licensed content remains available to Windows Media Player after you remove Windows Server 2003 Service Pack 1, do one of the following:

- Before you upgrade a computer to Windows Server 2003 Service Pack 1, back up the licenses for your digital media files. (You can do this through License Management in Windows Media Player.) Then, before removing the service pack, back up any additional licenses you have acquired. After you remove the service pack, restore all of the licenses.

- After you remove Windows Server 2003 Service Pack 1, you can install Windows Media Player 10 from the Microsoft Windows Media Download Center.

# Wireless Provisioning Services (WPS)

Updated: October 13, 2010

Applies To: Windows Server 2003 with SP1

**What does Wireless Provisioning Services do?**
An increasing number of users are accessing the Internet through a growing number of public wireless networks, or wireless fidelity (Wi-Fi) hotspots. Using Wireless Provisioning Services (WPS) provides wireless users with a consistent experience and seamless connectivity to public Wi-Fi hotspots through automatic provisioning of the client and seamless roaming. WPS enables Wireless Internet Service Providers (WISPs) to use a standards-based and integrated platform to provide Wi-Fi hotspots with enhanced security that are easy to use and manage. In addition, WPS enables enterprises to easily provide guest access with enhanced security to private wireless networks.

With WPS, WISPs and enterprises can send provisioning and configuration information to mobile clients as they connect to the Internet or the corporate network. This in turn allows seamless, automatic and secure configuration of mobile clients, enabling a uniform sign-up experience in the enterprise and across different public network providers and hotspot locations.

**Who does this feature apply to?**
Wireless Provisioning Services is designed for three types of organizations:

- Hotspot Service Provider (HSP)

  HSPs deploy wireless access points in public places, such as shopping malls and airports, but HSPs are not Internet Service Providers (ISPs). Instead, the HSP contracts with one or more ISPs, and offers users one or more service plans to choose from when they want to establish an account for Internet access.

- Wireless Internet Service Provider (WISP)

  WISPS are ISPs that either deploy Wi-Fi hotspots in public places or outsource Wi-Fi hotspot services to an HSP.

- Enterprise

  Enterprises can use WPS technology to provide managed guest access on their networks.

**What new functionality is added to this feature in Windows Server 2003 Service Pack 1?**
**Wireless Provisioning Services**
**Detailed description**
Wireless Provisioning Services is an extension to the existing wireless services and user interfaces within Windows XP and Windows Server 2003. It builds on the wireless features already in Windows, such as Wireless Zero Configuration, and the wireless security features, such as Protected Extensible Authentication Protocol (PEAP) and Wi-Fi Protected Access (WPA). WPS also includes modifications to Windows Server 2003. The Windows Server 2003 Internet Authentication Service (IAS) component was modified to include guest authentication of the clients in the provisioning process.

WPS includes a provisioning service component that allows for Wireless Internet Service Providers (WISPs) and enterprises to send provisioning and configuration information to a mobile client that is trying to connect to the Internet or the corporate network. By using Wireless Provisioning Services, WISPs can offer services at multiple network locations and use multiple network names (service set identifiers, or SSIDs). After users have signed up to a WISP in one location or are preprovisioned and have downloaded the provisioning information, they can automatically connect to the Internet on subsequent occasions using the network provided by the WISP in their different hotspot locations. The Wireless Zero Configuration (WZC) service will automatically choose the correct network belonging to the WISP based on the provisioning files supplied. WSP also enables automatic and seamless roaming between different providers.

Further, when WPS is used the client computer automatically keeps the provisioning information stored on the client computer up to date. This allows the provider to change the network settings, add new locations, and so on, without disrupting the service or having users reconfigure their systems.

When a user connects his computer to a WISP and establishes an account for the first time, the following four stages occur:

- The computer discovers the WISP network at a Wi-Fi hotspot.

- The user is authenticated using a guest account and the computer is connected to the Wi-Fi network.
- The mobile client is provisioned and the user establishes an account with the WISP.
- The user is authenticated on the Wi-Fi network using the new user account credentials.

Each of these stages is discussed in detail in the following scenario.

A user arrives at a Wi-Fi hotspot with a portable computer running Windows XP with Service Pack 2 or Windows Server 2003 with Service Pack 1 and Wireless Provisioning Services. When the computer comes within range of the WISP access point beacon the following occurs:

1. The Wireless Zero Configuration (WZC) service on the client computer detects the beacon information from the access point, which is enabled with a broadcast service set identifier (SSID). The SSID is equivalent to the network name.
2. The user is informed by Windows that a wireless network is available. The user views information in Windows, including the network's friendly name. In this example, the user possesses a promotion code to use for account establishment, and proceeds by clicking **Connect**. This causes the WPS client to connect the user's computer to the wireless network using a guest account with limited privileges.

When the guest account is authenticated by the Wi-Fi network, the following occurs:

1. WZC uses 802.1x and Protected Extensible Authentication Protocol (PEAP) to connect and authenticate as guest to the WISP network through the access point, automatically passing a blank user name and password to the WISP Internet Authentication Service (IAS) server (IAS is also known as the Microsoft RADIUS server). The access point is connected to a gateway device that allows traffic from the client to pass to the provisioning services in the network to complete the sign-up process, but blocks the client from accessing the Internet.
2. The IAS server (or RADIUS server) is the PEAP authenticator and Transport Layer Security (TLS) endpoint for users who connect as guest. The TLS tunnel is created between the client and the IAS server. All subsequent messages between client and server pass through this tunnel, which traverses the access point and the gateway device.
3. Server authentication is performed when the IAS server verifies its identity to the client computer using a certificate that contains the Server Authentication purpose in Enhanced Key Usage (EKU) extensions. This certificate is issued by a public trusted root certification authority (CA) that the client computer trusts.
4. The IAS server authenticates and authorizes the user as Guest. In the Access-Accept message that the IAS server sends to the client is a container with a URL to the provisioning information. This URL provides the Wireless Provisioning Services engine running on the client, with the location of the XML master file.

When the client is provisioned and the user creates an account, the following occurs:

1. On the client computer, the Wireless Provisioning Services downloads the XML master file and sub-files from the provisioning server. The master file contains pointers to XML subfiles that control the client's progress through the process. When the XML sign-up schema is downloaded, the sign-up wizard is launched on the client to allow the user to create and pay for an account with the WISP.
2. Using the sign-up wizard on the client computer, the user steps through the process of signing up for an account. The user enters the promotion code as well as personal data such as name, address, and credit card number. The data entered by the user is converted by the Wireless Provisioning Services client into an XML document.
3. The XML document containing the user's sign-up data is sent to the Web application on the WISP provisioning server.
4. The Web application checks the promotion code entered by the user against the promotion code database (for example, a SQL Server database). If the promotion code is valid, the Web application continues processing the user's data.
5. The Web application processes the user's payment information. Once payment is verified and sign-up information is

completed successfully, the Web application reads the domain and security group information from the promotion code database and creates a user account in identity services (such as Active Directory) and adds the account to the security group. The Web application also enters the new user name in the promotion code database.

6. An XML document containing the new account credentials is sent from the WISP provisioning server to the Wireless Provisioning Services client on the client computer. The client computer uses the credentials to configure WZC and 802.1x under the name of the WISP. The connection is re-initiated with the new user account password-based credentials (user name and password).

The re-initiated connection process is as follows:

1. The Wireless Zero Configuration (WZC) service on the client computer restarts the association to the SSID for the WISP.

2. WZC finds the correct 802.11 profile which was downloaded with the other WISP information in the XML master file. WZC reassociates with the access point using the correct profile.

3. WZC uses 802.1x to start the authentication process using a combination of the Protected Extensible Authentication Protocol and the Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP-MSCHAPv2) using the new account credentials passed to 802.1x by the Wireless Provisioning Services client.

4. As the client starts the authentication process with PEAP-MSCHAPv2 authentication, a TLS channel is created between the user's client computer and the WISP IAS server.

5. In the second stage of PEAP-MSCHAPv2 authentication, the WISP IAS server authenticates and authorizes the connection request against the new account in the user accounts database (for example, Active Directory). The IAS server sends an Access-Accept message to the access point. Included in the Access-Accept message are attributes that specify the user can now get access to the Internet.

6. The access point instructs the gateway device to assign the client to the logical segment network with access to the Internet.

**Why is this change important?**
Wireless Provisioning Services makes it easier to use wireless hotspots without compromising security. WPS, with Windows Server 2003 Service Pack 1, and Microsoft IAS (also known as a RADIUS server), allows users' computers to more easily discover, connect and roam between wireless hotspots with enhanced security.

- The current connection model for WISP signup and use is not secured. Most Wi-Fi hotspots are configured for open authentication and without data encryption. Users are generally required to launch a Web browser to initially sign-up to the WISP service and for subsequent logins. WSP mitigates this threat by adding encryption and authentication to the communications between the client and the wireless network.

- Browser redirection-based deployment has many usability issues. Users may not even know they have to launch their browser to get connected. Another example of what can happen is when the browser is set to use proxy settings to access the Internet and the user is connected directly to the corporate network. In this case, browser redirection does not work and the user would have to know to disable the proxy settings to connect to the hotspot. This can cause costly support calls to the WISP or the enterprise helpdesk.

- Browser based deployment is vulnerable to man-in-the-middle attacks, for example, by a malicious front-end server using a rogue access point. Users queried by this access point might unknowingly be giving away personal identification and credit card information. By eliminating the need for a Web login WSP reduces the vulnerability of WISP users to this type of attack.

- Without additional hotspot client software users cannot easily detect hotspots and do not have a unified mechanism to sign-up to them. It is not easy for users to find out information about the WISP or search for the hotspot locations for that WISP. If users sign-up at one hotspot, they are not necessarily configured to automatically use the other hotspots. In addition, there is no standard mechanism to keep their provisioning and configuration information up-to-date.

- Add-on hotspot client software can help the user access that specific WISP's network. However, add-on software can also conflict with the wireless services native to the operating system or client software from other providers, potentially causing interoperability problems, even instability of the system as they all attempt to control the wireless settings of the entire system. Updates to the WISP configuration usually require updates to the client software. For these reasons, many corporate IT departments are reluctant to deploy third party hotspot client software to their users.

- There is no standardized mechanism across WISPs to process user sign-ups and update their configurations. As a result, the user experience is fragmented and automatic and seamless roaming across providers can be difficult.

**Wireless Network Registration Wizard**
**Detailed description**
The Wireless Network Registration Wizard provides the user interface to sign-up for a wireless hotspot and guides the user through the provisioning process. The wizard builds content from provisioning information (XML files) provided by the WISP. The provisioning information can be dynamically downloaded or preinstalled on the client system. Preinstallation can be provided by an OEM for new systems, by the IT department within an organization, or from a WISP Web site. The WISP owns and creates the provisioning information and drives the users' sign-up and provisioning experience. The following example presents a simple Wireless Network Registration Wizard experience where the user has prepaid for an access code. The XML schema and wizard are flexible and can enable more complex sign-up experiences.

First, the user can either right-click the wireless network icon in the notification area and then click **View Available Wireless Networks**, or the user can respond to the notification message in the notification area that indicates availability of a new wireless network in range. When **Choose a wireless network** appears, the user selects a new wireless network and places that network on the preferred networks list.

The user then selects a network name (an SSID) and clicks **Connect** to connect to the wireless network. With a WPS-based Wi-Fi hotspot, the client detects that there is more provisioning information in form of XML files that is available about the network and the provider. It then confirms with the user whether the provisioning information should be downloaded. With a non-WPS network, the experience would be the same as with Windows XP today: either the users are prompted for a security key when connecting to a secure network or the users are warned that the network they are trying to connect to is unsecured, and they are asked if they still want to connect to it.

After the download is complete, the Wireless Network Registration Wizard automatically launches and guides the user through the sign-in process. The first screen displays a customized logo (or banner) and content from the provider.

The subsequent screens may include selecting a subscription plan, entering credit card information, personal information and so on. In this example there is just one plan and the user is asked to enter a prepaid or promotional code to get access to the network. Next, Wi-Fi Hotspot Deployment displays information about the selected plan, such as the terms of the service agreement and privacy statement.

On the last screen, the wizard asks the users for their connectivity preferences for this connection. These default preferences can be set by the provider but can be overridden by the user. For example, if the users select a monthly subscription with unlimited data, they probably want to always connect to the network automatically whenever in range. If the users choose a "pay-as-you-go" plan, they probably want to control when to connect and choose a manual connection option as their preference.

The second option determines whether the client keeps the provisioning information automatically up to date. For example, if the provider adds new network names, adds new locations, or changes the network or security settings, the client can automatically update the information without any user interaction required while connected to the network.

On subsequent visits to hotspots made available by the provider or by their roaming partners in the same or different locations, if automatic connection is selected, all the user has to do is to turn the mobile computer back on or resume operations from standby, and the user is automatically connected. When connected, instead of showing a cryptic network name or SSID in the **Choose wireless network** dialog box (which opens from the **View Available Wireless Networks** notification window), a friendly name of the provider will be shown, along with a logo of the provider.

From this dialog box, users can also search for available hotspot locations or view the help and support information provided by the WISP. Both the help and hotspot location information is downloaded as part of the provisioning information. The location information can be searched and viewed online or offline.

**What existing functionality is changing in Windows Server 2003 Service Pack 1?**
The wireless user interface has changed – a new **View Available Wireless Networks** dialog box will replace the existing dialog box.

**Do I need to change my code to work with Windows Server 2003 Service Pack 1?**
Wireless Provisioning Service does not require any changes to existing applications. There are two new APIs with WPS. One of the new APIs provides for adding to and queries through the XML data on the computer. This API can be used to preprovision the client from the WISP Web site by the user (using a standalone application), by OEMs, or IT departments.

**Additional Resources**
For more information about WPS, see

1. Deploying Wireless Provisioning Services (WPS) Technology, available in Word format on the Microsoft Download Center, at http://go.microsoft.com/fwlink/?LinkId=203315.

2. Using the Wireless Provisioning Services (WPS) Technology Authoring Tool, available in Word format on the Microsoft Download Center, at http://go.microsoft.com/fwlink/?LinkId=203316.

# .NET Framework version 1.1 for 64-bit operating systems

Applies To: Windows Server 2003 with SP1

**What does the .NET Framework do?**
The .NET Framework is a new computing platform designed to simplify application development in the highly distributed environment of the Internet. The .NET Framework has several main components: the common language runtime, the .NET Framework class library, and the runtime hosts. The common language runtime is the foundation of the .NET Framework. It manages code execution and provides core services such as memory management, thread management, and remote execution. The .NET Framework class library is a comprehensive, object-oriented collection of reusable classes that you can use to develop a wide range of applications, including ASP.NET applications and XML Web services. The .NET Framework also provides several runtime hosts, including Windows Forms and ASP.NET, which work directly with the runtime to provide a scalable, server-side environment for managed code.

In the previous release of the 64-bit versions of the Microsoft Windows Server 2003 operating systems and Microsoft Windows XP Professional x64 Edition operating system the .NET Framework was not supported. In the 64-bit versions of Windows Server 2003 Service Pack 1 and Windows XP Professional, the .NET Framework is supported. The Microsoft .NET Framework 1.1 is built into all products in the Windows Server 2003 family except for the 64-bit versions.

The 32-bit version of .NET Framework 1.1 is supported by 64-bit versions of Windows Server operating systems and can be installed for 32-bit applications running on 64-bit versions of Windows operating systems (WOW64).

To install .NET Framework 1.1, visit Windows Update and select the .NET Framework 1.1 package or see "Microsoft .NET Framework Version 1.1 Redistributable Package" at the Microsoft Download Center Web site at http://go.microsoft.com/fwlink/?LinkId=17153. To install the service pack, visit Windows Update or see "Microsoft .NET Framework 1.1 Service Pack 1" at the Microsoft Download Center Web site at http://go.microsoft.com/fwlink/?LinkId=39058.

**Important performance and scalability information for Itanium-based systems**
If a 32-bit version of the .NET Framework is installed on 64-bit Itanium-based systems, applications created for the 32-bit version of the Framework will bind and run against it. Therefore, applications created with the Microsoft .NET Framework 1.1 run only in a 32-bit process within the WOW64 compatibility environment. Due to the design of x86 emulation and the WOW64 subsystem for the Itanium processor family applications are restricted to execution on one processor.

Single processor execution and x86 emulation reduce the performance and scalability of 32-bit .NET Framework applications running on Itanium-based systems. We recommend that applications using the .NET Framework 1.1 be used for interactive, client applications. We do not recommend this environment for .NET Framework 1.1 applications that demand high performance or scalability, such as high-load ASP.NET applications. We suggest that application developers instead evaluate the .NET Framework 2.0, which includes native 64-bit support for Itanium-based systems, for increased performance and scalability.

**Who does this feature apply to?**
This feature applies to users with applications written to use the .NET Framework 1.1 and also to application developers developing .NET Framework 1.1 applications that will be run on both 64-bit and 32-bit operating systems.

# Appendices

Applies To: Windows Server 2003 with SP1

**Appendix A: Document History**

| Release Date | Changes to Document |
|---|---|
| October 6. 2004 | First version of document released |
| December 6. 2004 | Added new sections: Access-based Enumeration, Active Directory, ADPrep.exe, DCDiag.exe, Ntdsutil.exe, Security Account Manager (SAM), Device\PhysicalMemory Object, Kernel patch protection for x64-based operating systems, Remote Access Quarantine, Remote Installation Services (RIS) for 64-bit Operating Systems, Setup<br><br>Revised sections: Data Execution Prevention, DCOM, Internet Explorer features, Windows Firewall |
| April 18, 2005 | Added new sections: Distributed File System (DFS) in Windows Server 2003 Service Pack 1, Distributed Transaction Coordinator, Internet Information Services (IIS) in Windows Server 2003 Service Pack 1, .NET Framework for x64-based operating systems, Terminal Services in Windows Server 2003 Service Pack 1<br><br>Revised sections: Access-based Enumeration, Data Execution Prevention, DCOM Security Enhancements, Post-Setup Security Updates, Remote Installation Services (RIS) for 64-bit operating systems, Windows Firewall |

**Appendix B: Group Policy Administrative Templates**

Many operating system components have added support for Group Policy administrative templates, which support many new settings. These new settings have been created in response to customer feedback.

Microsoft has published a reference spreadsheet that lists Group Policy settings described in Administrative Template (.adm) files that are included with Windows Server 2003 Service Pack 1. To download the reference spreadsheet, see "Group Policy Settings Reference for Windows Server 2003 with Service Pack 1" on the Microsoft Download Center at http://go.microsoft.com/fwlink/?LinkId=22031.

# Windows Server 2003 Service Pack 1 Installation and Deployment Guide

Updated: June 8, 2005

Applies To: Windows Server 2003 R2


Introduction to Service Pack 1 (SP1)

Before You Deploy SP1

Planning Your Deployment

The Standalone Installation

The Integrated Installation

Command-Line Options for Installing SP1

Uninstalling SP1

Symbols for Debugging Service Packs

Additional Resources


**Introduction to Service Pack 1 (SP1)**

This guide explains how you can:

- **Deploy Microsoft® Windows Server™ 2003 Service Pack 1 (SP1) to multiple computers** that are already running the Windows Server 2003 operating systems. If you have only one server to update, you can see the Readme for Windows Server 2003 Service Pack 1 (on the Microsoft TechCenter Web site) instead.

- **Integrate SP1 with Windows Server 2003** for deploying to multiple computers running an earlier version of Windows® or with no operating system installed.

- **Uninstall SP1** from computers.

This guide does not explain:

- **How to combine SP1 with other updates for Windows Server 2003**. For information about how to do this, see Microsoft Windows Server 2003 Service Pack 1 Installing and Deploying Updates (on the Microsoft TechCenter Web site).

- **New SP1 functionality**. For important information about functional changes to SP1 that you should understand before you deploy SP1, review the information that applies to your setup on the Microsoft TechCenter Web site.

| Note |
| --- |
| You cannot install SP1 on any x64-based versions of Windows Server 2003. For more information about these versions, see Windows Server 2003 x64 on the Microsoft Web site. |

- **How to deploy SP1 by using Software Update Services (SUS)**

  For details about using SUS, see Windows Server System on the Microsoft Web site.

**Before You Deploy SP1**

Before you deploy SP1, you should review the information about the service pack on the Microsoft TechCenter Web site so that you can:

- Understand what functionality and security technologies SP1 includes and how you can implement them in your

environment.

- Find out how to plan, prepare, test, and troubleshoot your SP1 deployment.

- Read other essential installation information, and check minimum hard disk requirements in the Readme for Windows Server 2003 Service Pack 1 (on the Microsoft TechCenter Web site).

**List of updates**
For information about updates and fixes included in SP1, see article 824721, "Windows Server 2003 Service Pack 1 List of Updates," in the Microsoft Knowledge Base.

**Release notes**
Before you install SP1, see article 889101, "Release Notes for Windows Server 2003 Service Pack 1," in the Microsoft Knowledge Base.

**Files and programs in SP1**
You can obtain installation files to deploy SP1 to multiple computers from either the SP1 CD or the Microsoft Download Center.

For 32-bit versions of Windows Server 2003 (which run on x86-based computers), the SP1 installation file is named WindowsServer2003-KB889101-SP1-x86-*LLL*.exe, where *LLL* is the language version of SP1. For example, ENU represents the English-language version.

For Itanium-based versions of Windows Server 2003, the SP1 installation file is named WindowsServer2003-KB889101-SP1-ia64-*LLL*.exe.

You can update Itanium-based versions with SP1 for only English, German, Japanese, and French locales.

The instructions and examples throughout this document use the name of the installation file for 32-bit versions (unless a particular difference between 32-bit and Itanium-based installations must be called out).

**Contents of the SP1 CDs**
For each language version, two CDs are available to update Windows Server 2003 operating systems with SP1: one for 32-bit versions of Windows Server 2003 and one for Itanium-based versions. The CD for Itanium-based versions is labeled as such.

Each SP1 CD contains the following files:

- A service pack installation file

- The Deploy.cab file. This is a compressed file that contains a collection of tools and accompanying online Help. (See Deploy.chm and Ref.chm when you expand Deploy.cab.) It is located on the SP1 CD in the \Support\Tools folder. Deploy.cab is not automatically loaded when you install SP1.

- The Support.cab file. This is a compressed file that also contains a collection of tools and online Help. It is located in the \Support\Tools folder. Technical support personnel and network administrators can use the updated support tools in this file to manage networks and troubleshoot problems. Support.cab is not automatically loaded when you install SP1.

**Download Center**
You can update multiple servers by downloading SP1 installation files from the Microsoft Download Center. You can download installation files for 32-bit versions or Itanium-based versions of Windows Server 2003.

Because downloading SP1 installation files takes a substantial amount of time, we recommend that you save the file on your computer. This way, if you must restart the installation or reinstall SP1, you will not need to download the file again. Because the installation file is very large, if your computer has more than one hard disk, you should save the file on a disk other than your system disk.

**Planning Your Deployment**
To ensure a successful deployment, you should review the information presented in Before You Deploy SP1.

After you have reviewed this information, you must complete the following set of planning tasks:

- Choose an installation method.

- Remove all prerelease versions of SP1 from your servers.
- Determine whether you want to disable antivirus software during installations from the CD.
- Identify the deployment tools and files.
- Determine update and upgrade options.
- Check space requirements.
- Test and troubleshoot the deployment in your environment.
- Review the post-installation information near the end of this guide.

For detailed information about deployment-planning tasks and strategies, see the Windows Server 2003 Deployment Guide (on the Microsoft TechCenter Web site). This information can help you get started, including describing how to create a deployment plan and how to run a pilot project.

**Choosing an installation method**
You can either:

- Deploy SP1 on a network to multiple servers that are already running Windows Server 2003. (Standalone Installation)
- Integrate a Windows Server 2003 operating system with SP1, and then install the operating system and SP1 together on multiple servers using one procedure. This type of installation is also known as "slipstreaming." You can either install an integrated version on a server that does not have an operating system, or you can use the integrated version to upgrade computers that are running previous versions of Windows to Windows Server 2003 with SP1 (following the same path you would use to upgrade those computers to Windows Server 2003).
- Follow the instructions in the Windows Server 2003 Service Pack 1 Installing and Deploying Updates (on the Microsoft TechCenter Web site) to combine SP1 with an installation of Windows Server 2003, with other updates, or with both the operating system and other updates, and then install them together using one procedure.

You can perform a standalone or integrated installation from a shared distribution folder on a network. Even if you decide to do this, we recommend storing SP1 source files locally.

If you want to save disk space, you can extract the files from the service pack installation file. With this method, Windows accesses the distribution folder whenever it requires a service pack file, such as when Windows File Protection must restore a corrupted file or when an optional component is being configured. This step saves disk space on local computers because, if source files for the service pack are required, the installation will point to the distribution folder for the folder called ServicePackFiles instead of creating this folder locally.

**Important**

Because the service pack source files are stored remotely on the network, this installation method is suitable only for computers that are permanently connected to the network. For example, this method is not suitable for updating laptop computers.

You must use a shared distribution folder that was created specifically for SP1. Do not reuse a folder that was set up for a previous service pack. Shared distribution folders for service pack source files must be permanent to ensure that all of the files that a computer might need to replace are available.

For more information about moving the shared distribution folder, see article 271484, "Files and Folders Are Added to Your System After Service Pack Is Installed," in the Microsoft Knowledge Base.

**Removing prerelease versions of SP1**
Unless you have a supplemental support agreement with Microsoft, you must uninstall any prerelease versions of SP1 before you proceed. Microsoft does not support installing this version over the release candidate or beta versions of SP1.

**Antivirus software and installing the service pack**

| **Important** |
| --- |
| Antivirus software is a critical part of keeping a computer secure. But the time that it takes for the software to scan files on your system can add significantly to the time required to install the service pack: as much as an hour or more on slower systems. Disabling your antivirus software before you install the service pack eliminates this extra installation time. If you disable your antivirus software before you install the service pack, be sure that you understand the risks involved, and be sure to enable the antivirus software after you install the service pack. |

**Checking space requirements**

Check the system requirements for installing SP1 in the article 892807, "Hard Disk Space Requirements for Windows Server 2003 Service Pack 1" in the Microsoft Knowledge Base.

**Overview of the deployment tools and files**

After you select your installation method, review the corresponding scenario in The Standalone Installation or The Integrated Installation later in this guide. This section will help you identify the tools and files that you need for your deployment.

Depending on your environment, you might need one or more of the following deployment tools and files:

- **Microsoft Windows Installer**

  Windows Installer is a Windows component that standardizes and simplifies how you install and manage software programs (such as the service pack) on multiple computers. You can use this installation service to manage the installation, modification, repair, and removal of programs. Because Windows Installer facilitates consistency in your deployments, it enables you to manage shared resources, customize installation processes, make decisions about application usage, and resolve configuration problems.

- **The Unattend.txt answer file** (for integrated installations only)

  The Unattend.txt answer file identifies how Windows Server 2003 Setup interacts with the distribution folders and files that you create and provides information about your preinstallation requirements. This answer file also supplies Setup with all of the information that you are prompted to provide during a typical Windows Server 2003 installation.

  You can create or modify the Unattend.txt file by using either a text editor or Setup Manager. Make sure that only the system administrator has Write access to this file.

  For more information, see the Unattend.txt topic in the Microsoft Windows Preinstallation Reference Help (Ref.chm) in Deploy.cab in the \Support\Tools folder on your SP1 CD.

- **Microsoft Systems Management Server**

  Microsoft Systems Management Server (SMS) provides a variety of tools to help you deploy the service pack in your organization. With SMS 2003, you can automatically update servers in your organization with the new service pack.

  You can allow administrators to run the service pack installation whenever they like, or you can schedule the service pack installation to run at a specific time.

The following resources provide more detailed information about specific deployment tools and files.

- Windows Server 2003 Deployment Kit (on the Microsoft TechCenter Web site).

- Article 892778, "Updated System Preparation tool for Windows Server 2003 Service Pack 1," in the Microsoft Knowledge Base.

- The Microsoft Windows Corporate Deployment Tools User's Guide (Deploy.chm) and the Microsoft Windows Preinstallation Reference Help (Ref.chm) in Deploy.cab in the \Support\Tools folder on your CD.

- The Windows Support Tools Help (Suptools.chm) in Support.cab in the \Support\Tools folder on your CD. This Help file describes the support tools and gives examples of how to use them.

- The Release Notes (Readme.htm) for Windows Support Tools for Microsoft Windows Server 2003, which is located in the \Support\Tools folder on your CD.

- Article 892777, "Windows Server 2003 Service Pack 1 Support Tools," in the Microsoft Knowledge Base.

**Testing the deployment in your computing environment**
Although Microsoft has a high degree of confidence in SP1, it is not possible to test all hardware configurations and applications that might be present in all environments where you might install SP1.

We therefore recommend that you test the service pack in your computing environment before you deploy it to all of your servers.

Testing the service pack can include the steps described in following procedure, but might need to include more steps, depending on your environment.

**To test the service pack in your environment**

1. Your test environment should include a cross-section of the types of computers on which you want to deploy the service pack. Make sure that the test computers are equipped with the software and the hardware devices that are typically used in your organization.

2. Install the service pack on each test computer just as you would in the environment where you plan to use the service pack. Make sure that you complete each of the following tasks as appropriate for your environment:

    - Update test computers that are running Windows Server 2003 to Windows Server 2003 with SP1.

    - Upgrade test computers that are running Windows NT Server 4.0 and Windows 2000 Server to Windows Server 2003 with SP1.

    - Install Windows Server 2003 with SP1 onto test computers that have no other operating system installed (that is, perform a "clean installation").

3. Verify that the software and hardware continue to work as expected for a variety of scenarios.

If you are installing Windows Server 2003 for the first time in your organization, you might want to set up a pilot group to test your deployment and verify whether it works in your environment as expected.

**The Standalone Installation**
During the standalone installation, the service pack is installed on a computer that is already running Windows Server 2003. When you run the Update.exe program to install SP1, it installs the updated system files and makes the necessary registry changes. After the computer is restarted and the post-boot processes have completed, the installation is complete, and the operating system runs with an updated file set.

This section describes four scenarios for installing the service pack.

- Scenario 1: Installing the service pack so that computers use local source files for the service pack (recommended)

- Scenario 2: Installing the service pack so that computers share remote source files for the service pack

- Scenario 3: Using Systems Management Server (SMS) to install the service pack

- Scenario 4: Using Windows Installer and Group Policy to install the service pack

| **Important** |
| --- |
| If a system file on a computer becomes corrupted or must be replaced for some other reason, you must have the source files for that service pack to replace that system file. |

You can distribute the service pack through several types of media, including CD-ROMs, shared folders on networks, and Web downloads. Because this guide is intended for corporate audiences, it focuses on using shared distribution folders on networks, which is the most common distribution medium for this audience.

**Svcpack.log file creation**
When you run a standalone installation, the Update.exe program creates a log file named Svcpack.log in systemroot. This log file contains the following information:

- The command line that initiated the Update.exe program.
- A list of the files that the Update.exe program copied to your computer.
- Error messages if any errors were encountered.
- The command-line options that were used to install the service pack.

**Scenario 1: Installing the service pack so that computers use local source files for the service pack (recommended)**
By using this method, you can set up a shared distribution folder on a network for the service pack installation file. (The 32-bit version of this file is called WindowsServer2003-KB889101-SP1-x86-*LLL*.exe.) If you set up a shared distribution method, you can run the service pack installation from a central location and then store source files for the service pack locally.

To install SP1 from a shared folder, copy the service pack installation file to the folder from which you want to install SP1. You can download the SP1 installation file from the Microsoft Download Center, or you can copy it from the SP1 CD:

- WindowsServer2003-KB889101-SP1-x86-*LLL*.exe.
- WindowsServer2003-KB889101-SP1-ia64-*LLL*.exe.

When you run WindowsServer2003-KB889101-SP1-x86-*LLL*.exe, the updated system files are extracted to the local machine. By default, Update.exe then automatically installs the updated system files and makes the necessary changes in the registry. After the computer is restarted, the installation is complete, and Windows Server 2003 runs with an updated file set.

Unless you have a supplementary support agreement with Microsoft, you must uninstall any prerelease versions of SP1 before you proceed. Microsoft does not support the installation of this version over the release candidate or beta versions of SP1.

| **Note** |
| --- |
| In the following procedure, *Drive*: represents the drive letter of the network or computer where your distribution folder is located. |

**To update servers by installing SP1 from a shared folder so that computers use local source files for the service pack (recommended)**

1. Connect to the network or computer on which you want to create the distribution folder.
2. In the shared folder on the network, create a distribution folder for the service pack.

   For example:

   To install SP1 for a Windows Server 2003 operating system, create a distribution folder named WS2003SP1 by typing the following:

   **mkdir** *Drive*:**\WS2003SP1**

   (In the example, *Drive:* represents the drive where you want to place the distribution folder.)
3. Copy WindowsServer2003-KB889101-SP1-x86-*LLL*.exe from the SP1 CD or download it from the Microsoft Download Center into the distribution folder.
4. To verify that you have assigned appropriate permissions to users who are installing the service pack, do the following:

a. Navigate to the appropriate folder, right-click it, and then click **Properties**.

b. Click the **Sharing** tab, and then click **Share This Folder**.

c. In **Share Name**, type a name for the folder.

d. Click **Permissions**, and then assign permissions that allow users to install the service pack from this folder.

e. If you are using an NTFS file system partition, click the **Security** tab, confirm that the permissions listed there do not conflict with those on the **Sharing** tab, and then click **OK**.

5. To install the service pack from the distribution folder (in this example, the folder is named WS2003SP1), at the command prompt, type the following:

   *****ServerName***\WS2003SP1\WindowsServer2003-KB889101-SP1-x86-***LLL***.exe**

   You can include command-line options as needed. For example, if you want to install updates immediately after you install SP1, specify the **/Z** option in the installation command. This option prevents your computer from restarting automatically after SP1 is installed.

6. Follow the instructions in the Setup Wizard.

7. After the installation is complete, do one of the following:

   - If you do not want to install additional updates right after the service pack installation, restart the computer immediately.

   - If you want to install additional updates, install them before you restart your computer. If you use this approach, SP1 and the updates will become fully operational at the same time.

8. After you restart the computer, restart the antivirus software if you disabled it.

| **Important** |
| --- |
| SP1 will not be fully operational until you restart your computer. |

**# Scenario 2: Installing the service pack so that computers share remote source files for the service pack**
Use this method to set up a shared distribution folder on a network, and then install the service pack from the shared folder onto computers that are already running Windows Server 2003.

With this method, source files for the service pack are stored in the shared distribution folder rather than on local computers.

| **Important** |
| --- |
| Because the source files for the service pack are stored remotely on the network, this method is suitable only for computers that are permanently connected to the network. For example, this method is not suitable for updating laptop computers. |

You must use a shared distribution folder that was created specifically for SP1. Do not reuse a folder that was set up for a previous service pack. Shared distribution folders for service pack source files must be permanent to ensure that all of the files that a computer might need to replace are available.

If you want the installation to use source files from the distribution folder, you must first extract them from the installation file. This method saves disk space on local computers because, if source files for the service pack are required, the installation will point to the distribution folder for the folder called ServicePackFiles instead of creating this folder locally. The ServicePackFiles folder is needed whenever the operating system requires a service pack file, such as when Windows File Protection must restore a corrupted file, or when an optional component is being configured.

For more information about moving the shared distribution folder, see article 271484, "Files and Folders Are Added to Your System After Service Pack Is Installed," in the Microsoft Knowledge Base.

| **Note** |
| --- |
| In the following procedure, *Drive:* represents the drive letter of the network or computer where your distribution folder is located. |

**To install the service pack so that computers share remote source files for the service pack**

1. Connect to the network or computer on which you want to create the distribution folder.

2. In the shared folder on the network, create a distribution folder for the service pack. For example, to create a distribution folder named WS2003SP1, type the following at the command prompt:

   **mkdir** *Drive*:**\WS2003SP1**

3. To extract files to the shared distribution folder, type the following at the command prompt:

   **WindowsServer2003-KB889101-SP1-x86-***LLL***.exe /X:***Path*

   where *Path* is the location of the shared distribution folder.

   If you use **/U** (for unattended), no confirmation dialog box appears after the files have been extracted.

4. For each computer that you plan to update, back up the files and close programs before you continue (unless you plan to force programs to close during the installation).

5. To install the service pack from the shared distribution folder, run Update.exe instead of the installation file (for example, WindowsServer2003-KB889101-SP1-x86-*LLL*.exe). For example, to install the service pack from a distribution folder named WS2003SP1, type the following at the command prompt:

   *Drive***:\WS2003SP1\i386\Update\Update.exe.**

   You can use the same command-line options for Update.exe as you would for the installation file.

6. The installation continues as described in Scenario 1 earlier in this section.

We recommend that you restart the computer immediately to complete the installation after Update.exe has completed.

If you want to install updates, critical updates, or security updates, see the Guide for Installing and Deploying Updates for Microsoft Windows Server 2003 Service Pack 1 (on the Microsoft TechCenter Web site).

If you disabled your antivirus software, enable it again after you restart the computer.

| **Important** |
| --- |
| SP1 will not be fully installed until you restart your computer. |

**# Scenario 3: Using SMS to install the service pack**

This scenario explains how to use Microsoft Systems Management Server (SMS) to install the service pack from a shared distribution folder on a network. This installation method installs the service pack on SMS servers that are already running Windows Server 2003.

This scenario assumes that you have an understanding of SMS and a working knowledge of software distribution.

| **Important** |
| --- |
| To perform this type of installation, your computer must be running SMS 2.0 with Service Pack 5, SMS 2003, or SMS 2003 with Service Pack 1. |

**To use SMS to install the service pack**

1. Create the SMS package by importing the package definition file for the service pack. In the package, provide the path to the source files for the service pack.

2. Distribute the SMS package to the distribution points.

3. Create the advertisement to notify SMS servers that the service pack is available.

For details about deploying SP1 using SMS, see article 894712, "How to install Windows Server 2003 Service Pack 1 (SP1) by using Systems Management Server," in the Microsoft Knowledge Base.

**# Scenario 4: Using Windows Installer and Group Policy to install the service pack**
This scenario explains how to use Windows Installer to install the service pack from a shared distribution folder on a network. This method installs the service pack on servers that are already running Windows Server 2003.

To use this method, you should have a thorough understanding of Windows Installer and a working knowledge of Group Policy and Active Directory. For more information about Windows Installer, Group Policy, and Active Directory, see the Windows Server 2003 Deployment Guide (on the Microsoft TechCenter Web site).

**# Using Windows Installer**
The Windows Installer package file (Update.msi) contains all of the information that Windows Installer requires to install or remove the service pack and to run Setup.

This package file describes the relationships between service pack components and resources. It also contains an installation database from which you can view a summary (known as a summary information stream) of the properties that the installation program uses. To view this summary, right-click the installation database, and then click **Properties**.

You must use the machine-assigned distribution method with Update.msi. No other methods are available.

| **Important** |
| --- |
| We recommend that you use the installation file WindowsServer2003-KB889101-SP1-x86-*LLL*.exe or WindowsServer2003-KB889101-SP1-ia64-*LLL*.exe when deploying SP1 at the command prompt. (*LLL* represents the language code; for example, ENU represents English). If you use Update.msi to distribute service packs, you must use a server-based (machine-assigned) Group Policy object (GPO) only. User-based Group Policy deployments are not supported with Update.msi. For more information, see article 278503, "Best Practices for Using Update.msi to Deploy Service Packs," in the Microsoft Knowledge Base. |

To upgrade servers to SP1, you must join them to the same domain as the server where the Windows Installer (MSI) file resides. After you assign the package, Windows Installer automatically installs the service pack the next time that the server is started (as long as it is connected to the network).

| **Note** |
| --- |
| Only a network administrator or someone who is logged on to a local server as an administrator can remove the assigned software. |

**# Assigning Service Pack 1 to Servers**
The procedure in this section explains how to assign the service pack to servers that are managed by a GPO. For your installation, you might want to associate the GPO with a new Active Directory container.

**To use Group Policy to assign the service pack**

1. Create a shared network distribution folder.

2. Create a GPO for SP1 deployments.

3. Deploy the SP1 Update.msi from the shared distribution folder as machine-assigned. Do not deploy it as a user deployment.

   When the servers are restarted, they are updated to SP1.

| **Important** |
| --- |
| Check the properties of each server to ensure that the update has completed. You might need to restart each server more than once to complete the update. |

For general information about how to apply your software deployment policy to users who are outside the boundaries of organization units, see article 324750 "How To Assign Software to a Specific Group By Using a Group Policy in the Windows Server 2003 Family," in the Microsoft Knowledge Base.

**The Integrated Installation**

You should perform the integrated installation of Windows Server 2003 with SP1 on computers that do not have an operating system or that are running an earlier version of Windows, such as Windows 2000. You would use this method if you already purchased Windows Server 2003 and want to apply the service pack and the operating system at the same time.

When you perform an integrated installation, you install the operating system and the service pack simultaneously. This means that you do not need to perform separate installations of the operating system and the service pack.

To perform an integrated installation, first use the **/integrate** option to update Windows Server 2003 with SP1 installation files in a shared distribution folder. After you do this, run Windows Server 2003 Setup. During the integrated installation process, Windows Server 2003 Setup installs the operating system with the service pack already applied.

The path that you use to upgrade previous versions of Windows to Windows Server 2003 with SP1 is the same path you would use to upgrade to Windows Server 2003.

To complete the procedure for integrating Windows Server 2003 with SP1, you must use a computer that is already running a Windows Server 2003 or Windows XP operating system.

| **Note** |
| --- |
| In the following procedure, *Drive:* represents the drive letter of the network or computer where your distribution folder is located. |

**Spslpsrm.log file creation**

The Update.exe program creates a log file, named Spslpsrm.log, for integrated installations. This log file is created in *systemroot* and contains the following information:

- The command line that initiated the Update.exe program.
- A list of the files that the Update.exe program copied to your computer.
- Error messages if any errors were encountered.
- The command-line options that were used to install the service pack.

**Scenario 1: To create an integrated installation of the service pack with Windows Server 2003**

1. Connect to the network or computer on which you want to create the distribution folder.
2. In the shared folder on the network, create a distribution folder for the Windows Server 2003 installation files.

   For example:

   - To install SP1 with a 32-bit version of Windows Server 2003, create a distribution folder named WS2003\x86 by typing the following:

     **mkdir** *Drive***:\WS2003\x86**

   - To install SP1 with a 64-bit version of Windows Server 2003, create a distribution folder named WS2003\64bit by

typing the following:

**mkdir** *Drive***:\WS2003\64bit**

| **Important** |
| --- |
| Ensure that only system administrators have full access to this folder. Other users should have only Read and Execute permissions. |

3. To verify that you have assigned appropriate permissions to users who are installing the service pack, do the following:

   a. Navigate to the appropriate folder, right-click it, and then click **Properties**.

   b. Click the **Sharing** tab, and then click **Share This Folder**.

   c. In **Share Name**, type a name for the folder.

   d. Click **Permissions**, and then assign permissions that allow users to install the service pack from this folder.

   e. If you are using an NTFS file system partition, click the **Security** tab, confirm that the permissions listed there do not conflict with those on the **Sharing** tab, and then click **OK**.

4. Insert your Windows Server 2003 product CD or DVD into the CD-ROM drive, and then copy the contents of the CD to the distribution folder that you created in step 2.

   For example, to copy the installation files and folders from the Windows Server 2003 product CD in the CD-ROM drive (D:) to the distribution folder named WS2003\x86 on the E: drive, you would type the following:

   **xcopy /E D:\ E:\WS2003\x86**

5. Remove the product media from the CD-ROM drive, and then insert the SP1 CD. Be sure to use the correct SP1 CD (for either 32-bit or Itanium-based versions of Windows Server 2003).

6. Do one of the following:

   - To apply the source files for the service pack to the installation files for a 32-bit version of Windows Server 2003 in the folder named WS2003\x86, type the following:

     **WindowsServer2003-KB889101-SP1-x86-***LLL***.exe /integrate:***Drive***:\WS2003\x86**

   - To apply the source files for the service pack to the installation files for an Itanium-based version of Windows Server 2003 in the folder named WS2003\64bit, type the following:

     **WindowsServer2003-KB889101-SP1-ia64-***LLL***.exe /integrate:***Drive***:\WS2003\64bit**

   For a list of additional command-line options, see Command-Line Options for Installing SP1.

7. Service pack files are extracted and then applied to the operating system files to update your shared distribution folder. You can either use the files from this network shared folder to make a CD (32-bit versions only) to upgrade your servers, or you can customize Setup for Windows Server 2003 as necessary. For more information about how to customize Setup, you can view Help by entering the following command from the shared distribution folder:

   **\i386\winnt32.exe /?**

8. You can now deploy Windows Server 2003 with SP1 to your computers from the shared distribution folder in either

attended or unattended Setup mode. During the standard installation process, Setup for Windows Server 2003 (Winnt.exe or Winnt32.exe) installs the updated operating system with the service pack already applied.

| **Important** |
| --- |
| When you run the command given in Step 6, the Update.exe program automatically creates a Svcpack.log file in systemroot on the computer that is running the Update.exe program. If you plan to update more than one operating system from this computer, rename the Svcpack.log file after you update each version. This step ensures that you do not overwrite the current log file when you update additional operating systems. |

**Scenario 2: Using RIS to install Windows Server 2003 integrated with the service pack**

This section explains how to use Remote Installation Services (RIS) to install Windows Server 2003 operating systems integrated with the service pack. However, this section does not explain how to set up RIS. For more information about RIS, see Remote Installation Services (on the Microsoft TechCenter Web site).

**Types of RIS images**

RIS supports two types of operating system images: CD-based images (images created by using RISetup) and Remote Installation Preparation (RIPrep) images.

Using the CD-based option is similar to setting up a server operating system directly from the Windows Server 2003 product CD, except that the source files reside on a RIS server. You can create a RIPrep image from a server that is running Windows Server 2003 with SP1 and that is configured with applications and settings that you specify. You can then deploy this image to other similar servers connected to the network.

**Adding a CD-based image for an operating system that is integrated with a service pack to a RIS server**

After you install RIS, perform the following steps to add an image of Windows Server 2003 with SP1 to your RIS server.

**To copy Windows Server 2003 with SP1 to your RIS server**

1. Create an integrated installation of Windows Server 2003 with SP1. For more information, see Scenario 1: Creating an installation of Windows Server 2003 integrated with the service pack earlier in this guide.

2. To create an image, click **Start**, click **Run**, and then type the following:

   **risetup.exe -add**

3. The Remote Installation Services Setup Wizard appears. Follow the instructions in the wizard. When the **Installation Source Files Location page** prompts you for the image source, enter the path to the shared folder that contains the integrated installation that you created in Step 1.

4. A box appears showing the progress of the installation. After the RISetup image is complete, you can install the image on each server as described in Remote Installation Services (on the Microsoft TechCenter Web site).

**Updating existing RIPrep images to SP1**

If you want to add additional programs or settings to your integrated installation image, you must have a CD–based image of Windows Server 2003 with SP1 on your RIS server. (For more information, see the preceding section, Adding a CD-based image for an operating system that is integrated with a service pack to a RIS server.)

**To update existing RIPrep images**

1. Use RIS to install the RISetup image on your computer.

2. Add the additional programs and settings that you want to include in the image.

3. Run RIPrep to create an image on the RIS server.

   For instructions on how to run RIPrep on RIS servers running Windows Server 2003, see Remote Installation Services and Automating and Customizing Installations (on the Microsoft TechCenter Web site).

4. You can now deploy your RIPrep image to other servers. For details about how to perform this task, see Remote

Installation Services (on the Microsoft TechCenter Web site).

| Note |
| --- |
| You cannot use the **/integrate** option to integrate SP1 directly into an existing CD-based image or a RIPrep image. |

**Command-Line Options for Installing SP1**

You can define the way SP1 is installed by typing one or more of the following command-line options when you run the installation file (for example, WindowsServer2003-KB889101-SP1-x86-*LLL*.exe) or Update.exe. You can precede the options with either a forward slash (/) or a hyphen (-). The options are not case-sensitive.

| Command-line option | Description |
| --- | --- |
| **/?**<br>or<br>**/Help** | Displays help text. |
| **/D:** *FolderName* | Back up to a specified folder the files that are required to remove the service pack.<br>If you back up files but do not specify a folder, the default location is:<br>$ntservicepackuninstall$<br>Any location that you specify must be on the local drive. |
| **/ER** | Enables extended return codes that return a detailed error message if the update installation fails. For more information, see the white paper, The Package Installer (Formerly Called Update.exe) for Microsoft Windows Operating Systems and Windows Components, on the Microsoft Web site. |
| **/F** | After the installation is complete, you must close all open applications before you restart the computer. This option forces applications to close without saving files. You can use this option with other command-line options, except **/integrate**, **/L**, and **/Z**. |
| **/Forcerestart** | Restarts the computer after the update installation is complete. This option follows standard shutdown behavior. It does not force applications to close. |
| **/integrate:** *Path* | Integrates this software update into the path that you specify. Be sure to provide the absolute path to the installation files for your operating system. |
| **/L** | Lists installed updates. |
| **/N** | We do not recommend this option.<br>Does not back up files for removing SP1. An entry for SP1 does not appear in Add or Remove Programs if you use this option. |
| **/O** | Overwrites OEM files without prompting. |
| **/Q**<br>or<br>**/Quiet** | Uses quiet mode. Shows no user interface during SP1 installation. This is the same as unattended mode, except that the user interface is hidden. No prompts will appear during the installation process. |
| **/U**<br>or<br>**/Passive** | Uses unattended Setup mode and default SP1 options. Requires no user interaction during the installation of the updates and shows only critical errors and a progress bar. |

| | |
|---|---|
| **/Z**<br>or<br>**/Norestart** | Does not restart the computer after the installation is completed. |

**Command-Line Options for Installing the Service Pack Only**

The following table identifies additional command-line options that only the installation program for the service pack supports.

| Command-line option | Description |
|---|---|
| **/X** | Extracts service pack files without starting Update.exe. You are prompted to provide the path for the folder in which you want the extracted service pack installation files to be placed. |
| **/U /X:** *FolderName* | Extracts service pack files and places them in the *FolderName* folder without prompting and without starting Update.exe. |

| **Important** |
|---|
| SP1 will not be fully functional until you restart your computer. |

**Qfecheck.exe**

For more information about the Qfecheck.exe program, see the following articles in the Microsoft Knowledge Base:

- 282784: "Qfecheck.exe Verifies the Installation of Windows 2000 and Windows XP Updates"
- 304864: "Qfecheck Hotfix Tool Reports False Need to Reinstall Freshly Installed Updates"

Although these articles were written for other Windows operating systems, they also apply to updates for Windows Server 2003.

**Uninstalling SP1**

If you created backup files when you installed SP1, you can use Add or Remove Programs or the command prompt to remove SP1 at any time and restore your computer to its previous state.

You can uninstall SP1 only if it was installed as an update to an existing operating system. You cannot remove SP1 if you installed it as part of an integrated installation. In this case, you must uninstall the entire operating system together with SP1.

**To remove SP1 using a command prompt**

1. Open a command prompt and type:

   **%systemroot%\$NtServicePackUninstall$\spuninst**

2. Type **Spuninst.exe /***Option*. Replace *Option* with any of the options in the following table.

3. Press ENTER to continue, and then follow the instructions that appear on your screen.

| Option | Description |
|---|---|
| **/U** | Removes the service pack in unattended mode. If you use this option, only critical error prompts appear while you uninstall SP1. |
| **/Q** | Removes SP1 in quiet mode, which is the same as unattended mode except that the user interface is hidden. If you use this option, no prompts appear while you uninstall SP1. |
| **/Z** | Does not restart the computer after you uninstall SP1. |
| **/F** | Forces other applications to close when the computer restarts after you uninstall SP1. |

**Important**

If you choose to remove SP1, a dialog box displays a list of the applications that you installed after you installed SP1. If you continue to remove SP1, these applications might not work correctly.

Multiple descriptions of some applications might appear instead of a single generic description. You should ignore these additional descriptions.

**Symbols for Debugging Service Packs**

If you are a system administrator (someone who diagnoses problems on computers that are running Windows Server 2003), we recommend that you download the symbol files for debugging service packs. However, you do not need to install these files for your computer to run properly.

If you want to debug Windows Server 2003 with SP1, you must first download and install the symbols for debugging Windows Server 2003 from the download page on the Internet, and then you must add symbols for debugging SP1 to the same folder. For more information, see Debugging Tools for Windows — Overview at the Microsoft Web site.

You can download symbols for debugging SP1 from one of the following locations:

- From the Download Windows Symbol Packages page on the Microsoft Web site.
- On demand from the Internet Symbol Server. For more information about using the Internet Symbol Server, see Debugging Tools and Symbols: Getting Started on the Microsoft Web site.
- The English version of the Windows Server 2003 Service Pack 1 CD, which includes the symbols package. The package (Symbols.exe) is in the \Support\Symbols folder on the service pack CD. The symbols package is included only on the English version of the service pack CD, but this package is language-independent and can be downloaded for debugging any Windows Server 2003 locale.

**Note**

When you install SP1, the current versions of the symbol files overwrite any previous versions that are in the symbol installation folder.

**Additional Resources**

- For information about Windows Server 2003 and SP1, see the Microsoft TechCenter Web site.
- The Microsoft Knowledge Base provides updated information about the Windows Server 2003 operating systems, SP1, and other Microsoft products.
- For the most up-to-date information about supported hardware, see the hardware compatibility information at the Windows Server Catalog Web site.
- To obtain the most recent product updates, see the Windows Update Web site.

# Windows Server 2003 Service Pack 1 Installing and Deploying Updates

Updated: March 28, 2005

Applies To: Windows Server 2003 R2


About this Guide

What Is an Update?

The Standalone Update Installation

The Combination Installation

Installing SP1 and updates with Windows Server 2003

Removing Updates

Copyright


**About this Guide**

This guide was formerly called "Hotfix Deployment and Installation Guide."

This guide is geared for system administrators who are installing software updates on Microsoft® Windows Server™ 2003 Service Pack 1 (SP1).

Unless specified, this guide uses the term "update" to refer to any software update, including an update rollup, service pack, feature pack, critical update, security update, or hotfix that improves or fixes a software product that is released by Microsoft Corporation. For definitions of each of these types of software updates, see article 824684, "Description of the standard terminology that is used to describe Microsoft software updates," in the Microsoft Knowledge Base.

This guide includes technical information, procedures, and recommendations for installing updates on multiple computers in a small business or corporate environment. This includes installing updates separately or in combination with Windows Server 2003, with or without a service pack.

If you want to deploy either the standalone version of SP1 or the version of SP1 that is integrated with the operating system, see the Windows Server 2003 Service Pack 1 Installation and Deployment Guide (on the Microsoft TechCenter Web site).

Although this guide includes some basic operating system information, it assumes that you have a working knowledge of Windows Server 2003 and update installations. This guide is designed to complement formal training and other sources of information, such as the Windows Server 2003 Help documentation, not to replace them.

Although updates are available for both 32-bit (x86) and 64-bit versions of Windows Server 2003, throughout this guide, all examples will use the format for 32-bit updates.

**What Is an Update?**

An update is a file or collection of files that you can apply to your operating system to correct a specific problem. You can install updates separately from the operating system or in combination with Windows Server 2003, with or without a service pack.

An update is an executable (.exe) file. When you install an update, backup files are automatically created so you can remove the update later if you want to. The installation process also copies the new files to specific folders and updates registry settings.

Windows Server 2003 update files are named according to the following convention:

- For the 32-bit (x86) versions of the Windows Server 2003 operating systems:

  WindowsServer2003-KB######-x86-*LLL*.exe

- For the Itanium-based versions of the Windows Server 2003 operating systems

  WindowsServer2003-KB######-ia64-*LLL*.exe

- For the x64-based versions of the Windows Server 2003 operating systems:

WindowsServer2003-KB######-x64-*LLL*.exe

The symbols ###### represent the Microsoft Knowledge Base article number that explains the details about the update, and *LLL* represents the language version of the update.

Updates are self-installing. There are two main types of update installations:

- The standalone update installation

  Use this to install Windows updates on computers that are already running Windows Server 2003.

- The combination installation

  Use this to perform an unattended installation of Windows updates in combination with an installation of Windows Server 2003, SP1, or both.

| **Important** |
| --- |
| Updates are applied only to software that is already installed. If you remove a component and later reinstall it, you must reinstall any updates that apply to that component. In addition, if you add further components to your computer that require this update, you must reinstall the update for it to apply to the new component.<br><br>However, if you install updates by installing SP1, the updates are applied to all components that you add or reinstall. You would not need to reinstall SP1. |

**The Standalone Update Installation**

During the standalone update installation, Windows updates are applied to a computer that is already running Windows Server 2003. When you run the update program, it automatically installs updated system files and makes the necessary registry changes. After the computer is restarted, the installation is complete, and the operating system runs with an updated file set.

You can install the updates by running one of the update program files listed previously, which extracts the update files and runs the Update.exe installation program. Update.exe then checks the service pack version that you are currently using. If the service pack version was released before the updates, and the language is the same, Update.exe automatically installs the updates. If your service pack version was released after the updates, the installation is not completed, and an error message appears stating that the version of the updates is incorrect.

For more information about Update.exe, see the white paper, "Inside Update.exe - The Package Installer for Windows and Windows Components," on the Microsoft Web site.

| **Note** |
| --- |
| If your service pack version was released after the updates and you were completing an unattended installation using either the **/U** (unattended setup) or **/Q** (quiet setup) option, the installation does not continue and no message appears. For more information about command-line options, see Command-line options for the update package, later in this document. |

If the language of the updates does not match the language that you selected for your Windows Server 2003 operating system, Setup is interrupted. If there are no conflicts with the language version, Setup continues with the installation, and Update.exe registers the updates under the following registry keys:

- HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Updates\Server2003\SP2\KB######

- HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Hotfix\KB######

During the installation, information for removing the updates is stored in a hidden folder named %systemroot%\$NtUninstallKB######$. For information about removing updates, see Removing Updates later in this document.

**Methods for installing updates**
There are several methods that you can choose from to install an update, such as:

- Run Update.exe manually using a combination of installation options.
- Use Systems Management Server (SMS).
- Use Windows Installer.

You can deploy updates to multiple computers on a network by using a shared network distribution folder or by downloading the updates from the Web. Because this guide is intended primarily for corporate users, the procedures in the following section apply to the distribution of updates on a shared network. This is the most common way of distributing updates in a corporate environment.

The instructions in the rest of this section explain how to install updates on computers that are already running Windows Server 2003.

**Using Qchain.exe**
Windows Server 2003 updates have Qchain.exe functionality built in as part of Update.exe. This means that you can install SP1, and then install any number of post–SP1 updates without having to restart the computer between installations.

If multiple updates replace the same file, Update.exe ensures that the correct version of the file is retained. If you install multiple updates, be sure to use the **/Z** option that is described in "Command-line options for the update package," the next section in this guide.

**Command-line options for the update package**

| **Command-line option** | **Description** |
|---|---|
| **/?**<br>or<br>**/Help** | Displays help text. |
| **/ER** | Enables extended return codes that return a detailed error message if the update installation fails. For more information, see the white paper, "Inside Update.exe - The Package Installer for Windows and Windows Components," on the Microsoft Web site. |
| **/F** | Forces other applications to close after you install the update package and before you restart the computer. |
| **/Forcerestart** | Restarts the computer after the update installation is complete. This option follows standard shutdown behavior. It does not force applications to close. |
| **/integrate:** *path* | Integrates this software update into the path you specify. Be sure to provide the absolute path to your operating system source files. |
| **/L** | Lists installed updates. |
| **/N** | Does not back up files for removing updates. An entry for the update does not appear in **Add or Remove Programs** if you use this option. |
| **/O** | Overwrites OEM files without prompting. |
| **/Q**<br>or<br>**/Quiet** | Uses quiet mode. Shows no user interface during the installation of the updates. This is the same as unattended mode, except that the user interface is hidden. No prompts will appear during the installation process. |

| | |
|---|---|
| **/U**<br>or<br>**/Passive** | Uses unattended Setup mode. Requires no user interaction during the installation of the updates and shows only critical errors and a progress bar. |
| **/Z**<br>or<br>**/Norestart** | Does not restart the computer after the installation is completed. |

**Installing updates on computers running Windows Server 2003**

To install an update on a single computer, run the update program, WindowsServer2003-KB######-x86-*LLL*.exe, on that computer.

If you want to install the update on more than one computer, you can create a distribution folder for the update on the network. The following procedure describes how to install one or more updates using this method.

| **Important** |
|---|
| You must install SP1 either in combination with or before you install updates released after SP1. Installing more recently released updates before you install SP1 can cause problems with your system. |

| **Note** |
|---|
| For the following procedure, *Drive*:\ represents the drive of the network or computer on which your distribution folder is located. |

**To install an update by running the Update.exe program**

1. Connect to the network or computer on which you want to create the distribution folder.
2. In the shared network distribution folder, create a distribution folder for the update files.

   For example, to create a distribution folder named Update, type the following:

   **mkdir** *Drive***:\Update**

3. Copy the Windows Server 2003 Update.exe program to the distribution folder that you created in Step 2. For example, to copy the SP1 Update.exe program to the distribution folder named Update, type the following:

   **xcopy C:\WindowsServer2003-KB** ###### **-x86-** *LLL* **.exe** *Drive* **:\Update**

4. To install the update from the shared network distribution folder, run the WindowsServer2003-KB######-x86-*LLL*.exe program.

   For example, to install the update from the distribution folder named Update, type the following:

   *Drive* **:\Update\WindowsServer2003-KB** ###### **-x86-** *LLL* **.exe**

5. To ensure that your updates take effect, restart your computer after you finish installing all of the updates.

| **Note** |
|---|
| The WindowsServer2003-KB######-x86-*LLL*.exe file supports the same command-line options as Update.exe. You can use the command-line options that are described in the table that precede this procedure. |

**Installing several updates together**

You can group multiple updates together in a batch file and install them as a unit. This makes it unnecessary for you to restart your computer after you install each update. The update installer provides options that you can use when performing this type of installation.

The following code sample is a batch file that installs updates and ensures that the correct files are replaced after the computer is restarted.

```
@echo off
setlocal
set PATHTOFIXES=Drive:\update
%PATHTOFIXES%\WindowsServer2003-KB######-x86-LLL.exe /quiet /norestart
%PATHTOFIXES%\WindowsServer2003-KB######-x86-LLL.exe /quiet /norestart
%PATHTOFIXES%\WindowsServer2003-KB######-x86-LLL.exe /quiet /norestart
```

**Important**

To ensure that the updates that you installed take effect, restart the computer if the batch file does not automatically restart it for you. If you want your computer to restart automatically, in the last update that you install, replace the option **/norestart** with the option **/forcerestart**.

**The Combination Installation**

The combination installation uses both the update installation and integrated installation processes to install SP1 with one or more of the following:

- Windows Server 2003
- Windows Server 2003 updates
- Additional Microsoft and non-Microsoft software that might be included on the SP1 CD

**Note**

An integrated installation is one in which the operating system and the service pack are installed together as a single installation.

For a combination installation, you can include the components that you want to install as entries in the Svcpack.inf file, along with the updates. You can also choose to install SP1 at the same time.

This section provides instructions for using unattended Setup mode to simultaneously install Windows Server 2003 and updates.

**Installing SP1 and updates with Windows Server 2003**

This section explains how to perform a combination installation that includes SP1 and post–SP1 updates integrated with Windows Server 2003 in a shared distribution folder on a network. A combination installation makes it unnecessary for you to perform separate installations of Windows Server 2003, SP1, and the Windows Server 2003 updates.

If you want to install Windows Server 2003 with updates, you must use this method. This process installs the updates during Windows Server 2003 Setup.

**Preparing for the combination installation**

Before you deploy the installation of Windows Server 2003 and the Windows Server 2003 updates across a network, you must copy the Windows Server 2003 and Windows Server 2003 update installation files to a shared distribution folder and complete the steps described in the following procedure, "To create and set up the required folders and files for Windows Server 2003."

If you are an original equipment manufacturer (OEM) and you plan to install additional OEM files (such as device driver, application, or component files) on the destination computers, create a \$1 subfolder within the \$OEM$ subfolder. The \$1 subfolder corresponds to *systemdrive*, which is the destination drive for the Windows Server 2003 installation.

**Note**

In the following procedure, *Drive***:** represents the drive name of the network or computer where your distribution folder is located.

**To create and set up the required folders and files for Windows Server 2003**

1. Connect to the network or computer on which you want to create the distribution folder.
2. Create a distribution folder by typing the following (using the appropriate drive letter):

**mkdir** *Drive* **:\WServer2003DIST**

3. Use the command **xcopy** to copy the files and subfolders from the integrated installation product CD or DVD that includes Windows Server 2003 and SP1 to the *Drive*:\WServer2003DIST folder.

   For example, if D is your CD-ROM drive, type the following:

   **xcopy /E /I /V D:** *Drive* **:\WServer2003DIST**

4. Use the **/integrate** option to run the update that you want to install by typing the following:

   **WindowsServer2003-KB** ###### **-x86-** *LLL* **.exe /integrate: Drive:\WServer2003DIST**

   Repeat this step for each update you want to include your deployment.

| **Important** |
| --- |
| If you want to include SP1 *and* updates in your deployment of Server 2003, and you do not have a version of the operating system that includes SP1, first you must integrate Windows Server 2003 with SP1 as described in the Service Pack Deployment Guide. Then you can integrate the updates. If you integrate the updates before SP1, you will not be able to include SP1 afterwards. |

For more information about how to integrate updates, see article 828930, "How to integrate software updates into your Windows installation source files," in the Microsoft Knowledge Base.

**Deploying the combination installation**

You can deploy the installation of the Windows Server 2003 updates and the Windows Server 2003 operating system to multiple computers from a shared distribution folder on a network. During the standard installation process, Windows Server 2003 Setup installs the operating system and applies the updates.

### To deploy the installation

1. Verify that the installation and executable files for the Windows Server 2003 updates and Windows Server 2003 exist in your distribution folder.

2. Customize Windows Server 2003 Setup as necessary. For more information about customizing Windows Server 2003 Setup, see the Winnt32.exe command syntax topic in the Microsoft Windows Server 2003 Professional Resource Kit Documentation on the Windows Deployment and Resource Kits Web site.

3. Run Windows Server 2003 Setup. This deploys the installation of the Windows Server 2003 updates and Windows Server 2003 to multiple computers from the shared distribution folder.

| **Note** |
| --- |
| For more information about the preceding procedure, see the Microsoft Windows Corporate Deployment Tools User's Guide (Deploy.chm) included in Deploy.cab in the \Support\Tools folder on your SP1 CD. |

**Removing Updates**

You can remove an update for Windows Server 2003 at a command prompt or by using **Add or Remove Programs** in **Control Panel**.

If you want to remove updates, you must remove them in the reverse order in which they were installed. So, you must remove the most recently installed update first. For example, assume that you installed **UpdateA** before you installed **UpdateB**. To return your computer to its state before the installation of **UpdateA**, you must remove **UpdateB** before you remove **UpdateA**. You should follow this reverse order when you uninstall updates because some updates share files.

| **Important** |
| --- |

If you try to remove the updates in the wrong order, a warning appears that lists all of the updates and programs that have been installed since you installed the update that you are trying to remove. The warning names the update you are trying to remove and warns you that if you continue, the programs listed might not run correctly. Click **No** if you do not want to remove the update and if you prefer to remove the programs and updates that are listed before you continue. Otherwise, click **Yes**.

You cannot remove updates that were installed by using the combination installation method unless you reinstall the operating system. Reinstalling the operating system overwrites the updates.

**To remove an update at the command line**

1. Specify this location at the command line:

   **%systemroot%\$NtUninstallKB** ###### **$\spuninst**

2. Type **Spuninst.exe /***option*. Replace *option* with any of the following options:

| Option | Description |
|---|---|
| **/U** | Removes the update in unattended mode. If you use this option, only critical error prompts appear while you uninstall the update. |
| **/Q** | Removes the update in quiet mode, which is the same as unattended mode except that the user interface is hidden. If you use this option, no prompts appear while you uninstall the update. |
| **/Z** | Does not restart the computer after the process of uninstalling the update is complete. |
| **/F** | Forces other applications to close when the computer restarts after uninstalling the update. |

**To remove an update for Windows Server 2003 by using Add or Remove Programs**

1. Open **Control Panel**.
2. Click **Add or Remove Programs**.
3. Select the **Show updates** check box.
4. Click the update that you want to remove, and then click **Change or Remove**.
5. Follow the instructions that appear on your screen

**Copyright**
Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2005 Microsoft Corporation. All rights reserved.

Microsoft, MS-DOS, Windows, Windows NT, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

# Readme for Windows Server 2003 Service Pack 1 (Readmesp)

Updated: March 28, 2005

Applies To: Windows Server 2003 R2


Introduction to Service Pack 1 (SP1)

Components of SP1

Installation Options

Deploying SP1 to Multiple Servers

Release Notes

Before You Install SP1

Installing SP1 from the CD

Uninstalling SP1

Additional Resources

Copyright


**Introduction to Service Pack 1 (SP1)**

This document provides important information about Microsoft® Windows Server™ Service Pack 1 (SP1). SP1 contains updates that can help improve the security, reliability, and performance of your operating system. We recommend that you install SP1 because it includes updates that can improve how some programs run with Windows Server 2003 operating systems.

It is important to read this document before you install SP1. This document explains how to install SP1, and it provides specific information that might apply to your setup.

You can install SP1 to update the following versions of Windows Server 2003 operating systems:

- Windows Server 2003, Standard Edition
- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition
- Windows Server 2003, Web Edition
- Itanium-based versions of Windows Server 2003

**Note**

If you are running Windows Server 2003, Datacenter Edition, contact your OEM before you install SP1. OEMs are required to test Windows Server 2003, Datacenter Edition with service packs, such as SP1. If you subscribe to updates from your OEM, they should provide you with SP1 and provide or direct you to updated adapter drivers and utilities. They should also provide a custom HAL if your system requires such additional items. It is not possible to update the x64 versions of Windows Server 2003. For information about support for these products, contact your OEM, and see the Windows Server 2003 x64 Web site.

**Components of SP1**

SP1 contains updates and additional functionality for Windows Server 2003. You can read about these changes on the Microsoft TechCenter Web site. Help and Support Center has also been updated for SP1. For information about new functionality, see the updated version of Help and Support Center posted on the Microsoft TechCenter Web site. You can update Itanium-based versions of Windows Server 2003 with SP1 only for English, German, Japanese, and French locales.

**List of updates**

For information about updates and fixes included in SP1, see the article 824721, "Windows Server 2003 Service Pack 1 List of Updates," in the Microsoft Knowledge Base.

**Installation Options**
You can install SP1 from the CD, the Microsoft Download Center, or from the Windows Update Web site.

For 32-bit versions (x86) of Server 2003, the name of the SP1 installation file is WindowsServer2003-KB889101-SP1-x86-*LLL*.exe

The examples throughout this document use the 32-bit name for the SP1 installation file.

For Itanium-based versions of Windows Server 2003, the name of the SP1 installation file is WindowsServer2003-KB889101-SP1-ia64-*LLL*.exe, where *LLL* is the language version of SP1. For example, ENU represents the English-language version.

You can update Itanium-based versions with SP1 for only English, German, Japanese, and French locales.

**Using the CD to install SP1**
You can use the CD to install SP1 onto a single server, or follow the instructions in the Windows Server 2003 Service Pack 1 Installation and Deployment Guide (on the Microsoft TechCenter Web site) to deploy SP1 to multiple servers on a network.

There are two SP1 CDs available to update Windows Server 2003 operating systems with SP1: one for 32-bit versions of the Windows Server 2003 operating system, and one for Itanium-based versions of the Windows Server 2003 operating system. The CD for Itanium-based operating systems is labeled as such.

**Using the Web to install SP1**
There are two options available for installing SP1 from the Web.

| **Important** |
| --- |
| During the download, check the URL of the Web site to ensure that you are downloading files from Microsoft. |

**Windows Update**
You can use Windows Update to install SP1 on a single server running a 32-bit version of Windows Server 2003. You can update Itanium-based versions with SP1 only by using the CD or by downloading it from the Microsoft Download Center.

**Download Center**
You can update a single server or multiple servers by downloading SP1 from the Microsoft Download Center. You can choose to download SP1 for 32-bit (x86) versions or Itanium-based versions of the Windows 2003 operating system.

Because downloading the SP1 installation takes a substantial amount of time, we recommend that you save the SP1 installation file on your computer. This way, if you must restart the installation or reinstall SP1, you will not need to download the installation file again. Because the SP1 installation file is very large, if your computer has more than one hard disk, you should save WindowsServer2003-KB889101-SP1-x86-*LLL*.exe on a disk other than your system disk.

**Deploying SP1 to Multiple Servers**
You can obtain SP1 installation files from the SP1 CD or from the Download Center. Follow the instructions in the Windows Server 2003 Service Pack 1 Installation and Deployment Guide (on the Microsoft TechCenter Web site) and do one of the following:

- Deploy SP1 to multiple servers on a network that are already running Windows Server 2003 operating systems.
- Integrate a Windows Server 2003 operating system with SP1, and then install the operating system and SP1 together on multiple servers using one procedure. This type of procedure is also known as "slipstreaming."

Follow the instructions in the Windows Server 2003 Service Pack 1 Installing and Deploying Updates (on the Microsoft TechCenter Web site) to combine SP1 with an installation of Windows Server 2003, updates, or both the operating system and updates, and install them together using one procedure.

**Release Notes**
Before you install SP1, see article 889101, "Release Notes for Windows Server 2003 Service Pack 1." The information in this section will also help you prepare for your installation of SP1.

**Microsoft .NET Framework version 1.1**
The Microsoft .NET Framework is included in all of the 32-bit versions of the Windows Server 2003 operating systems. It is not included in 64-bit versions of the Windows Server 2003 operating systems.

The 32-bit version of .NET Framework 1.1 is supported by 64-bit versions of Windows Server operating systems and can be installed for 32-bit applications running on 64-bit versions of Windows operating systems (WOW64).

To install .NET Framework 1.1, see "Microsoft .NET Framework Version 1.1 Redistributable Package" at the Microsoft Download Center Web site. To install .NET Framework version 1.1 Service Pack 1, see "Microsoft .NET Framework 1.1 Service Pack 1" at the Microsoft Download Center Web site.

**Important performance and scalability information**
If a 32-bit version of the Microsoft .NET Framework is installed on the 64-bit Itanium-based systems, applications created for the 32-bit version of the .NET Framework will bind and run against it. Therefore, applications created with the .NET Framework version 1.1 run only in a 32-bit process within the WOW64 compatibility environment. Due to the design of x86 emulation for Itanium processors and the WOW64 subsystem, applications are restricted to execution on one processor.

Single processor execution and x86 emulation reduce the performance and scalability of 32-bit .NET Framework applications running on Itanium-based systems. We recommend that applications using the .NET Framework version 1.1 be used for interactive, client applications. With .NET Framework applications that demand high performance and scalability, such as high-load ASP.NET applications, we do not recommend this environment. We suggest that application developers instead evaluate the .NET Framework version 2.0, which includes 64-bit support for Itanium-based systems, for increased performance and scalability.

**Dell OpenManage support for Windows Server 2003 SP1**
In order to support the new enhancements and features in Windows Server 2003 SP1, Dell plans to release Dell OpenManage version 4.4, including Dell Server Assistant version 8.6, to fully support SP1. For current customers who are running Dell OpenManage version 4.3 or earlier, go to the Dell Web site for Dell OpenManage and Service Pack 1 upgrade and deployment information.

| **Note** |
| --- |
| Web addresses can change, so you might be unable to connect to the Web site mentioned here. |

**Before You Install SP1**
You can install SP1 only on a computer that is running a Windows Server 2003 operating system.

Before you install SP1, review the following information about Windows Server 2003 SP1.

1. Review the planning and deployment information for Windows Server 2003 SP1, available on the Microsoft TechCenter Web site.

   Because SP1 includes enhanced security, increased reliability, and a simplified administration, we recommend that you review these functionality changes and how they can affect your server setup before you deploy SP1. For example, SP1 includes the Security Configuration Wizard (SCW) that you can use to create security templates for servers based on their roles. Before you deploy SP1, you should determine how to incorporate SCW into your deployment process. An overview of this and other information you should know before you deploy SP1 is included in "Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1" on the Microsoft TechCenter Web site.

2. Check the system requirements for installing SP1 in the article, 892807, "Hard Disk Space Requirements for Windows Server 2003 Service Pack 1" in the Microsoft Knowledge Base.

3. Review article 824721, “Windows Server 2003 Service Pack 1 List of Updates” in the Microsoft Knowledge Base to find out about updates and fixes included in SP1.

4. Unless you have a supplemental support agreement with Microsoft, you must uninstall any prerelease versions of SP1 before you proceed. Microsoft does not support the installation of this version over the release candidate or beta versions of SP1.

5. Fully back up your servers. Your backup should include all data and configuration information that is necessary for that computer to function. It is important to perform a backup of configuration information for servers, especially those that provide network infrastructure, such as Dynamic Host Configuration Protocol (DHCP). When you perform the backup, be sure to include the boot and system partitions and the System State. Another way to back up configuration information is to create a backup set for Automated System Recovery.

6. If you are not connected to a network or the Internet, and you want to minimize the installation time for the SP1 update, you can choose to disable your antivirus software. You can usually do this by right-clicking your antivirus-software icon (which typically appears in the bottom right-hand corner of the computer screen), and clicking **Disable**. Be sure to enable the antivirus software again after you install SP1. Do this only if you are installing from the CD.

| **Important** |
| --- |
| Antivirus software is a critical part of keeping a computer secure, but the time that it takes for the software to scan files on your system can add significantly to the time required to install the service pack—as much as an hour or more on slower systems.<br><br>Disabling your antivirus software before you install the service pack eliminates this extra installation time.<br><br>If you disable your antivirus software before you install the service pack, be sure that you are aware of the risks involved, and be sure to enable it after the service pack is installed. |

**Installing SP1 from the CD**

The SP1 CD contains the program, WindowsServer2003-KB889101-SP1-x86-*LLL*.exe, which is a self-extracting executable file that contains the source files for SP1. The CD also includes tools for advanced users, technical support staff, and administrators.

**To install SP1 from the CD**

1. To install SP1, you must be logged on to the computer as an administrator.
2. Insert the SP1 CD into your CD-ROM drive. It might take a few moments for the installation document to display. This document is titled, "Installing Microsoft Windows Server 2003 Service Pack 1."

   If this document does not automatically open when you insert the CD, open it using the following steps:

   a. Click **Start**, and then click **Run**.

   b. Type the following, and then press ENTER:

      **D:Srsp1.cmd**

      D is the drive letter associated with your CD-ROM drive.
3. To complete a default installation (recommended), in "Installing Microsoft Windows Server 2003 Service Pack 1," look for the link to install SP1 and click that link to begin. Otherwise, to customize your installation, use the command line. To view command-line options for a customized installation, type:

   **WindowsServer2003-KB889101-SP1-x86-** *LLL* **.exe /?**
4. If the Setup Wizard prompts you to run WindowsServer2003-KB889101-SP1-x86-*LLL*.exe or to save it to disk, click **Run this program from its current location**, and then click **OK**.
5. Follow the instructions in the Setup Wizard.
6. During the installation, backup files are created by default so you can remove SP1 later if you want to. If you install from the command line, you can choose not to back up files by using the option **/N**. We do not recommend this option. If you install from the user interface, the Setup wizard allows you to choose where you want to create backup files.
7. After the installation is complete, restart your computer, and then turn on the antivirus software if you disabled it for the installation.

| **Important** |
| --- |
| If you choose not to create backup files, you will be unable to remove SP1. |

**Uninstalling SP1**

If you chose to create backup files when you installed SP1, you can use **Add or Remove Programs** or the command prompt

to remove SP1 at any time and restore your computer to its previous state.

It is only possible to uninstall SP1 if it was installed as an update to an existing operating system. It is not possible to remove SP1 if you installed it as part of an integrated installation. In this case, you can only uninstall the entire operating system together with SP1.

**To remove SP1 using a command prompt**

1. Open a command prompt and type:

   **%systemroot%\$NtServicePackUninstall$\spuninst**

2. Type **Spuninst.exe /***option*. Replace *option* with any of the options in the following table.

3. Press ENTER to continue, and then follow the instructions that appear on your screen.

| Option | Description |
|---|---|
| **/U** | Removes the service pack in unattended mode. If you use this option, only critical error prompts appear while you uninstall SP1. |
| **/Q** | Removes SP1 in quiet mode, which is the same as unattended mode except that the user interface is hidden. If you use this option, no prompts appear while you uninstall SP1. |
| **/Z** | Does not restart the computer after the process of uninstalling SP1 is complete. |
| **/F** | Forces other applications to close when the computer restarts after uninstalling SP1. |

**Important**

If you choose to remove SP1, a dialog box displays a list of the applications that you installed after you installed SP1. If you continue to remove SP1, these applications might not work correctly.

Multiple descriptions of some applications might appear instead of a single generic description. You should ignore these additional descriptions.

**Additional Resources**

- For information about Windows Server 2003 and SP1, see the Microsoft TechCenter Web site.
- The Microsoft Knowledge Base provides updated information about the Windows Server 2003 operating systems, SP1, and other Microsoft products.
- For the most up-to-date information about supported hardware, see the hardware compatibility information at the Windows Server Catalog Web site.
- To obtain the latest product updates, see the Windows Update Web site.

**Copyright**
This document supports a preliminary release of a software product that may be changed substantially prior to final commercial release. This document is provided for informational purposes only and Microsoft makes no warranties, either express or implied, in this document. Information in this document, including URL and other Internet Web site references, is subject to change without notice. The entire risk of the use or the results from the use of this document remains with the user. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this

document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2005 Microsoft Corporation. All rights reserved.

Microsoft, MS-DOS, Windows, Windows NT, Windows Server, Windows Media, FrontPage, and Outlook are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

# Windows Server 2003 Service Pack 2 Technical Library

Applies To: Windows Server 2003 with SP2

The technical library includes documentation, tools, and other resources to help you evaluate, plan, deploy, and manage Windows Server 2003 Service Pack 2 on the servers in your organization.

**Evaluating and planning**

- Readme for Windows Server 2003 Service Pack 2
- What's New in Windows Server 2003 Service Pack 2
- Release Notes for Microsoft Windows Server 2003 Service Pack 2

**Testing and deploying**

- Windows Server 2003 Service Pack 2 Installation and Deployment Guide
- Installing and Deploying Updates for Microsoft Windows Server 2003

For more information, see http://go.microsoft.com/fwlink/?LinkID=64780.

# What's New in Windows Server 2003 Service Pack 2

Applies To: Windows Server 2003 with SP2

This document outlines some changes in functionality between earlier versions of the Microsoft® Windows Server® 2003 and Microsoft Windows® XP Professional x64 Edition operating systems, and Windows Server 2003 with Service Pack 2 (SP2) and Windows XP Professional x64 Edition with SP2. Examples and details are provided for several of the technologies that are experiencing the biggest changes.

Unless otherwise specified, these notes apply to the following versions and editions:

- Windows Server 2003 with SP2
- Windows Server 2003 R2 with SP2
- Windows Server 2003 x64 Editions with SP2
- Windows Server 2003 R2 x64 Editions with SP2
- Windows Server 2003 for Itanium-based Systems with SP2
- Windows XP Professional x64 Edition with SP2

**Clustering**

A new event log event has been created to address certain situations in which the Cluster service account becomes excessively restricted by domain policy.

The new event ID is 1239. The event text includes troubleshooting information. You can also refer to article 871236 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=69284) for further information.

**Data access components**

XmlLite is new with Windows Server 2003 SP2. XmlLite is a fast, low-level, native XML parser with a small memory footprint. For more information, including the Programmer's Guide and API reference, see the MSDN Web site (http://go.microsoft.com/fwlink/?LinkID=69285).

**Distributed systems**

New options have been added to the Dcdiag.exe Domain Name System (DNS) tests. These new options are **/x** and **/xsl:xslfile.xsl** or **/xsl:xsltfile.xslt**. They generate XML tags when the tests are run with the /test:dns option. You can use this new output mechanism to more easily parse the verbose log that the DNS tests generate.

To direct the XML output file to XMLLog.xml, use the **/x** option. For example:

**dcdiag /test:dns /v /e /x:XMLLog.xml**

Note: The **/x**: option only works with the **/test:dns** option.

To add the processing instructions that reference the specified style sheet, use the **/xsl:xslfile.xsl** or **/xsl:xsltfile.xslt** option. For example:

**dcdiag /test:dns /v /e /x:XMLLog.xml; /xsl:xslfile.xsl**

**dcdiag /test:dns /v /e /x:XMLLog.xml; /xsl:xsltfile.xslt**

Note: The **/xsl:xslfile.xsl** or **/xsl:xsltfile.xslt** option only works with the **/test:dns /x:XMLLog.xml** option.

**File systems**

Icacls.exe is an upgrade of the Cacls.exe tool in Windows Server 2003 SP2, and can be used to reset the access control lists (ACLs) on files from Recovery Console, and to back up ACLs. Also, unlike Cacls.exe, Icacles.exe correctly propagates the creation of inherited ACLs and changes to them.

**Microsoft Message Queuing**

The default storage limit for message queuing has been changed to 1 gigabyte (GB). If you choose to have a storage limit of more than 1 GB, you can change the storage limit setting in Microsoft Management Console (MMC) on the **General** tab of

Message Queuing Properties.

**Networking and communications**

- This version of Windows Server 2003 SP2 includes an update that enables you to simplify the creation and maintenance of Internet Protocol security (IPsec) policy. This update enables you to use an IPsec "Simple Policy". For most environments, the installation of this update allows you to reduce the number of IPsec filters that are required for a Server Isolation deployment or for a Domain Isolation deployment. You can reduce the number of IPsec filters from many hundreds of filters to only two filters. For more information about this update for Windows Server 2003 and Windows XP, see article 914841 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=69286).

- Group Policy support for non-broadcasting networks and Wi-Fi Protected Access 2 (WPA2) settings has been added to the Windows wireless client in Windows Server 2003 SP2. This update allows the Windows wireless client to accept additional wireless Group Policy configuration options. These new settings include support for WPA2 parameters and non-broadcast networks.

- The Windows wireless client now supports WPA2, which enables you to take advantage of high levels of standards-based connection and encryption security. New security features include:

    - Non-broadcast network profiles are now marked with a flag to improve the security of the Windows wireless client.

    - Windows will not automatically connect to a peer-to-peer network, even if it has been automatically saved in the preferred network list. You must manually connect to a peer-to-peer network profile.

**Windows Deployment Services**

Starting with this version of Windows Server 2003 with SP2, Remote Installation Services is replaced by Windows Deployment Services. You can use Windows Deployment Services to set up new computers through a network-based installation without having to be physically present at each computer and without having to install directly from DVD media. For more information about Windows Deployment Services, see the Windows Deployment Services Update Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkID=69289).

# Windows Server 2003 Service Pack 2 Installation and Deployment Guide

Updated: February 19, 2009

Applies To: Windows Server 2003 with SP2

This guide is for system administrators who are installing Microsoft® Windows Server® 2003 operating systems with Service Pack 2 (SP2). It includes technical information, procedures, and recommendations for installing SP2 on multiple computers in a small business or corporate environment. Although this guide includes some basic operating system information, it assumes that you have a working knowledge of Windows Server 2003.

This guide does *not* explain:

- **How to combine SP2 with other updates for Windows Server 2003**. For information about how to do this, see "Installing and Deploying Updates for Microsoft Windows Server 2003" (http://go.microsoft.com/fwlink/?LinkID=64779).
- **New functionality in SP2**. For important information about functional changes to SP2, see "What's New in Windows Server 2003 Service Pack 2" (http://go.microsoft.com/fwlink/?LinkID=62452).
- **How to deploy SP2 by using Windows Server Update Services (WSUS)**. For more information, see the WSUS Web site (http://go.microsoft.com/fwlink/?LinkID=94242).

**Obtain the installation files**

You can install SP2 from the SP2 CD or from the Web. To order an SP2 CD, see the Microsoft TechNet Web site (http://go.microsoft.com/fwlink/?LinkId=78087). To download the installation files, see the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=78088).

Because downloading SP2 installation files takes a substantial amount of time, we recommend that you save the files on your computer. This way, you can reinstall or repair the installation without downloading the file again. You should ensure that there is enough disk space on your hard disk drive before starting the download.

The installation file is named one of the following. XXX stands for the type of operating system (for example x86) and LLL stands for the language (for example ENU).

- WindowsServer2003-KB914961-SP2-XXX-LLL.exe
- WindowsServer2003.WindowsXP-KB914961-SP2-XXX-LLL.exe

**Plan your deployment**

This section includes information about how to plan your SP2 deployment, including pre-deployment tasks, testing the deployment, and additional resources.

**Pre-deployment tasks**

Before you install SP2, do the following:

- **Review the following documentation.**
  - Article 914962 in the Microsoft Knowledge Base for a list of the updates included in this service pack (http://go.microsoft.com/fwlink/?LinkId=74380).
  - "Release Notes for Windows Server 2003 Service Pack 2" for a list of the known issues with this version of SP2 (http://go.microsoft.com/fwlink/?LinkId=74381).
  - "Installing and Deploying Updates for Microsoft Windows Server 2003" for information about combining this service pack with other updates (http://go.microsoft.com/fwlink/?LinkID=64779).
- **Check the disc space requirements.** See article 926029 in the Microsoft Knowledge Base

(http://go.microsoft.com/fwlink/?LinkId=74454).

- **Back up your servers**. Your backup should include all data and configuration information that is necessary for that server to function. It is important to perform a backup of configuration information for servers, especially those that provide network infrastructure, such as Dynamic Host Configuration Protocol (DHCP). When you perform the backup, be sure to include the boot and system partitions and the system state. Another way to back up configuration information is to create a backup set for Automated System Recovery.

- **Disconnect UPS devices.** If you have an uninterruptible power supply (UPS) connected to your destination computer, disconnect the connecting serial cable before running Setup. Setup automatically attempts to detect devices connected to serial ports, and UPS equipment can cause problems with the detection process.

- **Disable your antivirus software.** If you are not connected to a network or the Internet and you want to minimize the installation time, you can disable your antivirus software. The time that it takes for the antivirus software to scan each file can add significant time to the installation (an hour or more on slower computers). Be sure to enable the antivirus software after you install SP2.

- **Uninstall prerelease versions of SP2.** Unless you have a supplementary support agreement with Microsoft, you must uninstall any prerelease versions of SP2 before you proceed. Microsoft does not support upgrading to this version from any prerelease version.

- **Close all open programs.**

**Testing the deployment**

We recommend that you test SP2 in your computing environment before you deploy it to all of your servers. If you are installing Windows Server 2003 for the first time, you should set up a pilot group to test your deployment and verify that the deployment works as expected. For more detailed information about testing your deployment and creating a pilot program, see the Windows Server 2003 Deployment Guide (http://go.microsoft.com/fwlink/?LinkId=44319). To test the deployment, do the following:

1. Create a test environment. Make sure that the test computers contain the operating systems, software, and the hardware devices that are used in your organization.

2. Install the service pack on each test computer. As needed, test the following scenarios:

    - Update computers that are running Windows Server 2003.

    - Update computers that are running Windows Server 2003 with SP1.

    - Update computers that are running Windows Server 2003 R2.

    - Upgrade test computers that are running Windows NT® Server 4.0 and Windows 2000 Server.

    - Perform a clean installation—that is, install Windows Server 2003 with SP2 onto computers that do not have an operating system.

3. Verify that the software and hardware perform as expected in each scenario.

**Additional resources**

The following resources contain detailed information that can help you plan your deployment.

- Windows Server 2003 Deployment Guide (http://go.microsoft.com/fwlink/?LinkId=44319).

- Article 926028 in the Microsoft Knowledge Base: "Updated System Preparation Tool for Windows Server 2003 Service Pack 2." This article contains instructions on how to use the Sysprep tool to deploy system images (http://go.microsoft.com/fwlink/?LinkId=78090).

- Deploy.cab and Support.cab. These are compressed files that contain a collection of tools and online **Help**. Technical support personnel and network administrators can use these updated support tools to manage networks and troubleshoot problems. These files are located on the SP2 CD in the \Support\Tools folder and on the Microsoft Download Center.

    - To download Deploy.cab, search the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=78088) for "Windows Server 2003 Service Pack 2 Deployment Tools."

    - To download Support.cab, search the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=78088) for "Windows Server 2003 Service Pack 2 Support Tools." Support.cab is included in the WindowsServer2003-KB914961-SupportTools-XXX-LLL.exe file.

- Article 926027 in the Microsoft Knowledge Base: "Windows Server 2003 Service Pack 2 Support Tools" (http://go.microsoft.com/fwlink/?LinkId=78091).

- Qfecheck.exe.Qfecheck.exe is a command-line tool that gives network administrators increased ability to track and verify installed hotfixes. For more information about the Qfecheck.exe program, see article 282784 in the Microsoft Knowledge Base: "Qfecheck.exe Verifies the Installation of Windows 2000 and Windows XP Updates" (http://go.microsoft.com/fwlink/?LinkId=44323). Although this article was written for other operating systems, it also applies to Windows Server 2003.

**Installation options**
You have two installation options:

- Standalone installation. Deploys SP2 to servers that are already running Windows Server 2003.

- Integrated installation. Integrates SP2 into Windows Server 2003, and then installs the operating system and SP2 simultaneously. You can use the integrated installation to upgrade servers and to perform clean installations on servers that do not have an operating system. For more information about creating an integrated version, see "Installing and Deploying Updates with Microsoft Windows Server 2003" (http://go.microsoft.com/fwlink/?LinkID=64779).

To perform these installations, first copy the installation files to a network share. Then extract the SP2 source files either to either each computer locally or to a network share (to save disk space on each computer).

When you extract the files to a network share, it is called a UNC install. For UNC installs, each computer uses the files on the network share to install the software. UNC installs are only suitable for computers that are permanently connected to the network because the service pack source files are stored remotely. For example, you should not use this method for updating laptop computers because if connectivity is lost during the installation, the computer may be left in an unusable state. In addition, you must use a network share that was created specifically for SP2 so that the source files are always available for computers. If you need to move a shared distribution folder, see article 271484 in the Microsoft Knowledge Base: "Files and Folders Are Added to Your System After Service Pack Is Installed" (http://go.microsoft.com/fwlink/?LinkId=9273).

**The standalone installation**
In a standalone installation, you install the service pack on computers that already have an operating system. The service pack installation is not complete until the computer is restarted. When you run a standalone installation, Update.exe creates a log file named Svcpack.log in the %systemroot% directory.

There are four scenarios for installing the standalone service pack:

- Scenario 1: Installing SP2 using local source files

- Scenario 2: Installing SP2 using source files that are located on a network share

- Scenario 3: Installing SP2 using Systems Management Server (SMS)

- Scenario 4: Installing SP2 using Windows Installer (Update.msi) and Group Policy

**Scenario 1: Installing SP2 using local source files**
First set up a shared distribution folder for the service pack installation file. Then to install the service pack, save the installation files locally on each computer.

### To install SP2 using installation files that are saved locally on each computer

1. Connect to the network or computer where you want to create the distribution folder.
2. In the shared folder, create a distribution folder for the service pack. For example, to create a distribution folder named WS2003SP2, type **mkdir Drive:\WS2003SP2** at the command prompt(Drive: represents the drive where you want to store the distribution folder).
3. Copy WindowsServer2003-KB914961-SP2-XXX-LLL.exe into the distribution folder.
4. To verify that you have assigned appropriate permissions to users who are installing the service pack:
   a. Navigate to the appropriate folder, right-click it, and then click **Properties**.
   b. Click the **Sharing** tab, and then click **Share This Folder**.
   c. In **Share Name**, type a name for the folder.
   d. Click **Permissions**, and then assign permissions that allow users to install the service pack from this folder.
   e. If you are using an NTFS file system partition, click the **Security** tab, confirm that the permissions listed there do not conflict with those on the **Sharing** tab, and then click **OK**.
5. To install the service pack at the command prompt, type the following using the desired command-line options:

   **\\ServerName\WS2003SP2\WindowsServer2003-KB914961-SP2-XXX-LLL.exe**
6. Follow the instructions in the Setup Wizard.
7. After the installation is complete, if you want to install additional updates, install them before you restart your computer.
8. Restart the computer and then restart the antivirus software if you disabled it.

**Scenario 2: Installing SP2 using source files that are located on a network share**
To install SP2 from the network, first extract the installation files to a shared distribution folder. Then install SP2 on each computer using these files.

### To install the service pack using a network share

1. Connect to the network or computer where you want to create the distribution folder.
2. In the shared folder, create a distribution folder. For example, to create a distribution folder named WS2003SP2, type **mkdir Drive:\WS2003SP2** at the command prompt (Drive: represents the drive where you want to store the distribution folder).
3. To extract the files to the distribution folder, type the following at the command prompt, where *Path* is the location of the distribution folder:

   **WindowsServer2003-KB914961-SP2-XXX-LLL.exe /X:Path**
4. For each computer that you plan to update, back up the files and close programs before you continue (unless you plan to force programs to close during the installation).
5. To install the service pack from the shared distribution folder, run Update.exe. For example, to install from a distribution folder named WS2003SP2, type the following at the command prompt using the desired command-line options:

   **Drive:\WS2003SP2\i386\Update\Update.exe**
6. Follow the instructions in the Setup Wizard.
7. After Update.exe has completed, restart the computer to complete the installation. SP2 will not be fully installed until you

restart your computer.

8. If you disabled your antivirus software, enable it after you restart the computer.

**Scenario 3: Installing SP2 using Systems Management Server (SMS)**
This scenario explains how to use Microsoft Systems Management Server (SMS) to install the service pack from a shared folder. You can use this method to install the service pack on SMS servers that are already running Windows Server 2003.

This scenario assumes that you have an understanding of SMS and a working knowledge of software distribution. For more information about deploying SP2 using SMS, see article 926030 in the Microsoft Knowledge Base: "How to Install Windows Server 2003 Service Pack 2 (SP2) by Using Systems Management Server" (http://go.microsoft.com/fwlink/?LinkId=78092).

### To use SMS to install the service pack

1. Ensure that your computer is running SMS 2003 with Service Pack 1 or Service Pack 2.

2. Create the SMS package by importing the package definition file for the service pack. In the package, provide the path to the source files for the service pack.

3. Distribute the SMS package to the distribution points.

4. Create the advertisement to notify SMS servers that the service pack is available.

**Scenario 4: Installing SP2 using Windows Installer (Update.msi) and Group Policy**
This scenario explains how to install SP2 from a shared folder using the Windows Installer package file (Update.msi). You can use this method to install the service pack on servers that are already running an operating system.

Update.msi contains all of the required information to install or remove the service pack. This file defines the relationships between service pack binary files. It also contains an installation database where you can view a summary of the properties that the installation program uses. To view this summary, right-click the installation database, and then click **Properties**. You must use the machine-assigned distribution method with Update.msi.

Before you begin, please note:

- To use this method, you should have a thorough understanding of Windows Installer and a working knowledge of Group Policy and Active Directory. For more information about Windows Installer, Group Policy, and Active Directory, see the Windows Server 2003 Deployment Guide (http://go.microsoft.com/fwlink/?LinkId=44319).

- For general information about how to apply your software deployment policy to users who are outside the boundaries of organization units, see article 324750 in the Microsoft Knowledge Base: "How To Assign Software to a Specific Group By Using a Group Policy in the Windows Server 2003 Family" (http://go.microsoft.com/fwlink/?LinkId=43972).

- We recommend that you use the SP2 installation file (for example, WindowsServer2003-KB914961-SP2-XXX-LLL.exe) when deploying SP2 at the command prompt. If you use Update.msi to distribute service packs, you must use a server-based ("machine-assigned") Group Policy object (GPO) only. User-based Group Policy deployments are not supported with Update.msi. For more information, see article 278503 in the Microsoft Knowledge Base: "Best Practices for Using Update.msi to Deploy Service Packs" (http://go.microsoft.com/fwlink/?LinkId=28711).

**#Assigning Service Pack 2 to servers using Group Policy**
This procedure explains how to assign the service pack to servers that are managed by a Group Policy object (GPO). For your installation, you might want to associate the GPO with a new Active Directory container. Only a network administrator or someone who is logged on to a local server as an administrator can remove the assigned software.

### To use Group Policy to assign SP2

1. Ensure that the server you want to update is joined to the same domain as the server where Update.msi resides.

2. Create a shared network distribution folder.

3. Create a GPO for SP2 deployments.

4. Deploy Update.msi from the shared distribution folder as "machine-assigned." Do not deploy it as a user deployment.

5. Restart each server. When you restart each server, it is updated to SP2. You should check the properties of each server to ensure that the update has completed. You might need to restart each server more than once to complete the update.

**The integrated installation**
You should perform the integrated installation if you want to install SP2 and Windows Server 2003 at the same time. To perform this type of installation, first use the **/integrate**option to update Windows Server 2003 with the SP2 installation files. To do this, you need to use a computer that is already running a Windows Server 2003 or Windows XP operating system.

After you have the integrated version, you can run Windows Server 2003 Setup. The Update.exe program creates a log file (named Spslpsrm.log) for integrated installations. This log file is created in the %systemroot% directory. If you plan to update more than one operating system from this computer, rename the Svcpack.log file after you update each version. Otherwise Update.exe will overwrite this file.

This section describes three scenarios for installing the integrated installation.

- Scenario 1: Creating an integrated installation
    - For Windows Server 2003
    - For Windows Server 2003 R2
- Scenario 2: Using RIS to install an integrated installation
- Scenario 3: Using Windows Deployment Services to install an integrated installation

**Scenario 1: Creating an integrated installation**
This section contains instructions for creating integrated installations of SP2 with Windows Server 2003 and Windows Server 2003 R2.

**For Windows Server 2003**
**To create an integrated installation of Windows Server 2003 with SP2**

1. Connect to the network or computer where you want to create the distribution folder.
2. Create the shared folder and ensure that only system administrators have full access to this folder. (Other users should have only Read and Execute permissions.)
3. In the shared folder, create a distribution folder for the Windows Server 2003 installation files. For example, to create a distribution folder named WS2003\x86, type:

   **mkdir Drive:\WS2003\x86**
4. To verify that you have assigned appropriate permissions to users who are installing the service pack:
    a. Navigate to the appropriate folder, right-click it, and then click **Properties**.
    b. Click the **Sharing** tab, and then click **Share This Folder**.
    c. In **Share Name**, type a name for the folder.
    d. Click **Permissions**, and then assign permissions that allow users to install the service pack from this folder.
    e. If you are using an NTFS file system partition, click the **Security** tab, confirm that the permissions listed there do not conflict with those on the **Sharing** tab, and then click **OK**.
5. Insert your Windows Server 2003 product disc into the CD-ROM drive and copy the contents to the distribution folder. For example, to copy the contents in the CD-ROM drive (D:) to the distribution folder named WS2003\x86 on the E: drive, type:

   **xcopy /E D:\ E:\WS2003\x86**
6. Remove the Windows Server 2003 product disc, and insert the SP2 CD.
7. Apply the source files for the service pack to the installation files. For example, for a 32-bit version of Windows Server 2003, in the folder named WS2003\x86, type:

**WindowsServer2003-KB914961-SP2-XXX-LLL.exe /integrate:Drive:\WS2003\x86**

8. Service pack files are extracted and then applied to the operating system files to update your shared distribution folder. You can either use the files from this folder to make a CD (32-bit versions only) to upgrade your servers, or you can customize Setup for Windows Server 2003 as necessary. For more information about how to customize Setup, view **Help** by typing **i386\winnt32.exe /?** at the command prompt in theshared distribution folder.

9. You can now deploy Windows Server 2003 with SP2 to your computers from the shared folder in either attended or unattended Setup mode.

**For Windows Server 2003 R2**

This section applies only if you have the first release of Windows Server 2003 R2 where disc 1 contains Windows Server 2003 with Service Pack 1. If the product discs that you have do not specify which service pack is included, the discs contain Windows Server 2003 with SP1. If you have the second release of Windows Server 2003 R2 where disc 1 contains Windows Server 2003 with SP2, you do not need to integrate disc 1 and SP2 because disc 1 already contains SP2.

**To create an integrated installation of Windows Server 2003 R2 and SP2**

1. Connect to the network or computer where you want to create the distribution folder.

2. Create the shared folder and ensure that only system administrators have full access to this folder. (Other users should have only Read and Execute permissions.)

3. In the shared folder, create a distribution folder for the Windows Server 2003 R2 disc 1 installation files. For example, to create a distribution folder named WS2003\x86\CD1, type:

   **mkdir Drive:\WS2003\x86\CD1**

4. To verify that you have assigned appropriate permissions to users who are installing the service pack:

   a. Navigate to the appropriate folder, right-click it, and then click **Properties**.

   b. Click the **Sharing** tab, and then click **Share This Folder**.

   c. In **Share Name**, type a name for the folder.

   d. Click **Permissions**, and then assign permissions that allow users to install the service pack from this folder.

   e. If you are using an NTFS file system partition, click the **Security** tab, confirm that the permissions listed there do not conflict with those on the **Sharing** tab, and then click **OK**.

5. Insert disc 1 into the CD-ROM drive, and copy the contents to the distribution folder. For example, to copy the contents in the CD-ROM drive (D:) to the distribution folder named WS2003\x86\CD1 on the E: drive, type:

   **xcopy /E D:\ E:\WS2003\x86\CD1**

6. Remove the Windows Server 2003 R2 disc 1, and insert the SP2 CD.

7. Apply the source files for the service pack to the installation files. For example, for a 32-bit version of Windows Server 2003 in the folder named WS2003\x86\CD1, type:

   **WindowsServer2003-KB914961-SP2-XXX-LLL.exe /integrate:Drive:\WS2003\x86\CD1**

8. Service pack files are extracted and then applied to the operating system files to update your shared distribution folder. You can either use the files from this folder to make a CD (32-bit versions only) to upgrade your servers, or you can customize Setup for Windows Server 2003 as necessary. For more information about how to customize Setup, view **Help** by typing **i386\winnt32.exe /?** at the command prompt in theshared distribution folder.

9. Insert your Windows Server 2003 R2 disc 2 into the CD-ROM drive, and copy the contents to the distribution folder. For example, to copy the contents in the CD-ROM drive (D:) to the distribution folder named WS2003\x86\CD2 on the E: drive, type:

   **xcopy /E D:\ E:\WS2003\x86\CD2**

10. Apply the source files for the service pack to the installation files. For example, for a 32-bit version of Windows Server 2003 in the folder named WS2003\x86\CD2, type:

    **WindowsServer2003-KB914961-SP2-XXX-LLL.exe /integrate:Drive:\WS2003\x86\CD2**

11. You can now deploy the integrated version of disc 1 and the installation files for disc 2 to your computers from the shared folder in either attended or unattended Setup mode.

**Scenario 2: Using RIS to install an integrated installation**
This section explains how to use Remote Installation Services (RIS) to install an integrated version of Windows Server 2003 with SP2. However, this section does not explain how to set up RIS. For more information about RIS, see Remote Installation Services (http://go.microsoft.com/fwlink/?LinkId=36940).

RIS supports two types of operating system images:

- CD-based images (images created by using RISetup)
- Remote Installation Preparation (RIPrep) images.

The CD-based option is similar to setting up a server operating system directly from the Windows Server 2003 product CD, except that the source files reside on a RIS server. You can create a RIPrep image from a server that is running Windows Server 2003 with SP2, and you can then deploy this image to other similar servers that are connected to the network.

**Creating a CD-based (RISetup) image to add to a RIS server**
After you install RIS, perform the following steps to create the image and then add the image to your RIS server.

**To create an image of Windows Server 2003 with SP2 and add it to your RIS server**

1. Create an integrated version of the operating system with SP2. For more information, see Scenario 1.
2. To create an image, click **Start**, click **Run**, and then type:

   **risetup.exe /add**

3. The Remote Installation Services Setup Wizard appears. Follow the instructions in the wizard. When the **Installation Source Files Location page** prompts you for the image source, enter the path to the shared folder that contains the integrated version that you created in Step 1.
4. A box appears showing the progress of the installation. After the RISetup image is complete, you can install the image on each server as described in Remote Installation Services on the Microsoft TechCenter Web site (http://go.microsoft.com/fwlink/?LinkId=36940).

**Creating a Remote Installation Preparation (RIPrep) image to deploy to other servers**
If you want to add additional programs or settings to your integrated installation image, you must have a CD–based image of Windows Server 2003 with SP2 on your RIS server. Note that you cannot use the **/integrate** option to integrate SP2 directly into an existing CD-based image or a RIPrep image.

**To update existing RIPrep images**

1. Use RIS to install the RISetup image on your computer.
2. Add the additional programs and settings that you want to include in the image.
3. Run RIPrep to create an image on the RIS server. For instructions on how to run RIPrep on RIS servers running Windows Server 2003, see:
   - Remote Installation Services (http://go.microsoft.com/fwlink/?LinkId=36940)
   - Automating and Customizing Installations (http://go.microsoft.com/fwlink/?LinkId=44324)
4. You can now deploy your RIPrep image to other servers.

**Scenario 3: Using Windows Deployment Services to install an integrated installation**
This section explains how to use Windows Deployment Services to install an integrated version of Windows Server 2003 with

SP2. However, this section does not explain how to set up Windows Deployment Services. For more information, see the documentation that is included with the tool and the Windows Deployment Services Update Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkId=66145).

You can also deploy SP2 by using RIS functionality from a Windows Deployment Services Server. For example, first ensure that the server is in either Legacy or Mixed mode. Then follow the procedures in Scenario 2 to deploy SP2. To check the operating mode that you are in, run **WDSUTIL /get-server /show:config**. The output will contain a line with the mode that you are in. For example, "WDS operational mode: Mixed."

To add an image of Windows Server 2003 with SP2 to your Windows Deployment Services server, you must first create a .wim install image. Install images are built from reference installations of the Windows operating system and are deployed when Windows Deployment Services is used to install Windows. A reference installation may be a standard Windows installation or may be configured for a specific environment or use prior to creating the image.

**To create an install image**

1. Create a reference computer that contains Windows Server 2003, SP2, and any programs and settings that you want to include in the image.
2. Download the updated deployment tools (Deploy.cab). To do this, search the Microsoft Download Center for "Windows Server 2003 Service Pack 2 Deployment Tools" (http://go.microsoft.com/fwlink/?LinkId=78088).
3. From a command prompt on the reference computer, change directories to the directory that contains the updated version of Sysprep.exe (which is included in Deploy.cab). For more information about using this tool, see article 926028 in the Microsoft Knowledge Base: "Updated System Preparation Tool for Windows Server 2003 Service Pack 2" (http://go.microsoft.com/fwlink/?LinkId=78090).
4. Type **sysprep.exe /reseal /reboot**.
5. When the reference computer reboots, press **F12**.
6. In **Windows Boot Manager**, choose to boot to the appropriate capture image.
7. On the **Image Capture Wizard** welcome page, click **Next**.
8. On the **Image Capture Source** page, use the **Volume to Capture** drop-down menu to choose the appropriate volume, and then provide a name and description for the image. Click **Next** to continue.
9. On the **Image Capture Destination** page, click **Browse** and browse to the location where you want to store the captured image.
10. In the **Filename** text box, type a name for the image using the .WIM file name extension, and then click **Save**.
11. Click **Upload image to WDS server**.
12. Type the name of the Windows Deployment Services server and click **Connect**.
13. If prompted for credentials, provide a user name and password for an account with sufficient privilege to connect to the Windows Deployment Services server.
14. Using the drop-down **Image Group** menu, choose an image group in which to store the image.
15. Click **Finish**.

**Command-line options**

You can use the following options to define the way SP2 is installed. You can use the following options through Update.exe or when you run the installation file (for example, **WindowsServer2003-KB914961-SP2-XXX-LLL.exe**). You can precede the following options with a forward slash (/) or a hyphen (-). The options are not case-sensitive.

| **Note** |
| --- |
| You can also use the /Log and /Verbose options. For more information, see "The Package Installer (Formerly Called Update.exe) for Microsoft Windows Operating Systems and Windows Components" (http://go.microsoft.com/fwlink/?LinkId=42759). |

| Command-line option | Description |
| --- | --- |
| **/D:Path** | Backs up the files that are required to remove the service pack to *Path*. Any location that you specify must be on the local drive. If you back up files but do not specify a folder, the default location is:<br><br>$ntservicepackuninstall$ |
| **/ER** | Enables extended codes that return a detailed error message if the update installation fails. For more information, see "The Package Installer (Formerly Called Update.exe) for Microsoft Windows Operating Systems and Windows Components" (http://go.microsoft.com/fwlink/?LinkId=42759). |
| **/F**<br><br>or<br><br>**/Forceappsclose** | Forces other applications to close after you install the update package and before the computer is restarted. You cannot use this option with **/integrate**, **/L**, and **/Z**. |
| **/Forcerestart** | Restarts the computer after the installation is complete, even if a restart is not required to complete the installation. This option does not force applications to close. |
| **/Integrate:Path** | Integrates the SP2 source files into the path you specify. Be sure to provide the absolute path to your operating system source files. This option is the same as the **/S** option that was used by previous installers. |
| **/L** | Lists installed updates. |
| **/N**<br><br>or<br><br>**/Nobackup** | Does *not* create backup files during setup. We do not recommend this option because backup files allow you to remove SP2 in the future. An entry for SP2 does not appear in **Add or Remove Programs** if you use this option. |
| **/O**<br><br>or<br><br>**/Overwriteoem** | Overwrites OEM files without prompting. |
| **/Promptrestart** | If a restart is required, a dialog box notifies the user that a restart is required to complete setup. Intended for use with /Quiet. |
| **/Q**<br><br>or<br><br>**/Quiet** | Uses quiet mode. This mode shows no user interface during the installation of the updates. This is the same as unattended mode, except that the user interface is hidden. No prompts will appear during the installation process. |
| **/U**<br><br>or<br><br>**/Passive** | Installs the service pack in unattended mode. Only critical error prompts and a progress bar will appear during the installation. By default, the program invokes the **/Warnrestart** option. |
| **/Warnrestart: [Seconds]** | If a restart is required, a dialog box notifies the user that the computer will restart in the specified seconds. The user can click **Cancel** or **Restart Now**. The default is 60 seconds. |
| **/X**<br><br>or<br><br>**/Extract** | Extracts service pack files without starting Update.exe. You are prompted to provide the path for the folder in which you want the extracted service pack installation files to be placed. For more information, see "The Package Installer (Formerly Called Update.exe) for Microsoft Windows Operating Systems and Windows Components" (http://go.microsoft.com/fwlink/?LinkId=42759). |
| **/Z**<br><br>or<br><br>**/Norestart** | Does not restart the computer after the installation is complete, even if a restart is required to complete the installation. |

| | |
|---|---|
| **/?**<br>or<br>**/Help** | Displays Help. |

**Uninstalling SP2**

It is possible to uninstall SP2 only if it was installed as an update to an existing operating system. It is not possible to remove only SP2 if you installed it as part of an integrated installation. In addition, you can only uninstall the last service pack that you installed. For example, if you install SP1 and then you install SP2, you can only uninstall SP2.

By default, backup files are created during the installation so you can remove SP2. Unless you specified the /nobackup option when installing SP2, then you can remove SP2 using **Add or Remove Programs** or the command prompt.

If you remove SP2, a dialog box displays a list of the applications that you installed after you installed SP2. If you continue, these applications might not work correctly. Multiple descriptions of some applications might appear instead of a single generic description. You can ignore these additional descriptions.

**To remove SP2 using a command prompt**

1. Open a command prompt, type the following, and then press ENTER (*Option* can be any option listed in the following table):

   %systemroot%\$NtServicePackUninstall$\spuninst\Spuninst.exe /*Option*

2. Follow the instructions that appear on your screen.

| Option | Description |
|---|---|
| **/F**<br>or<br>**/Forceappsclose** | Forces applications to close when the computer restarts. |
| **/Q**<br>or<br>**/Quiet** | Removes SP2 in quiet mode, which is the same as unattended mode except that the user interface is hidden and no prompts appear. |
| **/U**<br>or<br>**/Passive** | Removes the service pack in unattended mode. Only critical error prompts appear while SP2 is being uninstalled. |
| **/Z**<br>or<br>**/Norestart** | Does not restart the computer after the process of uninstalling SP2 is complete. |
| **/?**<br>or<br>**/Help** | Displays Help. |

**Debugging symbols**

If you want to debug the operating system, first download and install the symbols for debugging Windows Server 2003, and then add the symbols for debugging SP2 to the same folder. When you install SP2, the current versions of the symbol files overwrite any previous versions that are in the symbol installation folder. For more information, see Debugging Tools for Windows—Overview (http://go.microsoft.com/fwlink/?LinkId=30565).

You can download symbols for debugging SP2 from one of the following locations:

- The Download Windows Symbol Packages page on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=30539).

- On demand from the Internet Symbol Server. For more information about using the Internet Symbol Server, see "Debugging Tools and Symbols: Getting Started" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=30524).

- The English version of the SP2 CD. This CD contains the symbols package (Symbols.cab) in the \Support\Symbols folder. This package is language-independent and can be downloaded for debugging any Windows Server 2003 locale.

**Additional resources**

- For information about SP2, see the Microsoft TechCenter Web site (http://go.microsoft.com/fwlink/?LinkId=78093).

- The Microsoft Knowledge Base provides updated information about the Windows Server 2003 operating systems, SP2, and other Microsoft products (http://go.microsoft.com/fwlink/?LinkId=4441).

- For the most up-to-date information about supported hardware, see the hardware compatibility information at the Windows Server Catalog Web site (http://go.microsoft.com/fwlink/?LinkId=16906).

- To obtain the most recent product updates, see the Windows Update Web site (http://go.microsoft.com/fwlink/?LinkId=22377).

**Copyright**
Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2006 Microsoft Corporation. All rights reserved.

Microsoft, MS-DOS, Windows, Windows NT, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

# Installing and Deploying Updates for Microsoft Windows Server 2003

Applies To: Windows Server 2003 with SP2

**About this guide**
This guide is for system administrators who are installing software updates on Microsoft® Windows Server® 2003. It includes technical information, procedures, and recommendations for installing updates on multiple computers in a small business or corporate environment. Although this guide includes some basic operating system information, it assumes that you have a working knowledge of Windows Server 2003 and update installations. All of the examples in this guide use the format for 32-bit updates.

If you want to deploy the standalone version of a service pack or the service pack that is integrated with the operating system, see the Windows Server 2003 Service Pack Installation and Deployment Guide that applies to your service pack. (For example, the guide for Service Pack 2 is located at http://go.microsoft.com/fwlink/?LinkID=64780.)

**What is an update?**
This guide uses the term "update" to refer to any software update, including an update rollup, service pack, feature pack, critical update, security update, or hotfix that is released by Microsoft Corporation. For descriptions of the standard terminology that is used to describe Microsoft software updates, see article 824684 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=43258).

Although this guide uses update as a broad term, it is important that you know the following differences between hotfixes and service packs:

- Hotfixes can be un-installed in the order they are installed.
- Only the last installed service pack can be uninstalled.
- Hotpatch command-line options are valid only for hotfixes, not for service packs.

You can install updates separately from the operating system or in combination with Windows Server 2003 (with or without a service pack). When you install an update, backup files are automatically created so you can remove the update later if you want to. The installation process also copies the new files to specific folders and updates registry settings.

| **Important** |
| --- |
| Updates are applied only to software that is already installed. For example, if you remove a component and later reinstall it, you must reinstall updates that apply to that component. In addition, if you further add components to your computer that require this update, you must reinstall the update for it to apply to the new component. |

Windows Server 2003 update files are named according to the following format: WindowsServer2003-KB######-XXX-LLL.exe. The symbols###### represent the Microsoft Knowledge Base article that explains the details about the update, XXX represents the version of the operating system (for example, x86), and LLL represents the language version of the update.

Updates are self-installing. There are two main types of update installations:

- The standalone update installation  Use this installation to update computers that are already running Windows Server 2003.
- The combination installation  Use this unattended installation to install an update in combination with an installation of Windows Server 2003, a service pack, or both.

**The standalone update installation**
During the standalone update installation, updates are applied to a computer that is already running an operating system. When you run the update program, it installs updated system files and makes the necessary registry changes. The installation is complete after the computer is restarted.

You can install the updates by running the WindowsServer2003-KB######-XXX-LLL.exe file, which extracts the update files

and runs Update.exe. Update.exe registers the updates under the following registry keys:

- HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Updates\Server2003\SP2\KB######
- HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Hotfix\KB######

For more information about Update.exe, read this white paper about the package installer for Windows and Windows components (http://go.microsoft.com/fwlink/?LinkId=42759).

During the installation, information for removing the updates is stored in a hidden folder named %systemroot%\$NtUninstallKB######$. For information, see Removing updates.

**Using Qchain.exe**
Windows Server 2003 updates have Qchain.exe functionality built in as part of Update.exe. This means that you can install the service pack, and then install any number of updates without having to restart the computer between installations.

If multiple updates replace the same file, Update.exe ensures that the correct version of the file is retained. If you install multiple updates, be sure to use the **/Z** option so the computer does not restart after each update.

**Command-line options for the update package**
To customize your installation, you can use the following command-line options with WindowsServer2003-KB######-x86-LLL.exe and Update.exe.

| **Note** |
| --- |
| If your service pack version was released after the updates and you were completing an unattended installation using either the **/U** (unattended setup) or **/Q** (quiet setup) option, the installation does not continue and no message appears. |

| **Command-line option** | **Description** |
| --- | --- |
| **/ER** | Enables extended codes that return a detailed error message if the update installation fails. For more information, see "The Package Installer (Formerly Called Update.exe) for Microsoft Windows Operating Systems and Windows Components" (http://go.microsoft.com/fwlink/?LinkId=42759). |
| **/F**<br>or<br>**/Forceappsclose** | Forces other applications to close after you install the update package and before the computer is restarted. You cannot use this option with **/integrate**, **/L**, and **/Z**. |
| **/Forcerestart** | Restarts the computer after the installation is complete, even if a restart is not required to complete the installation. This option does not force applications to close. |
| **/Integrate:Path** | Integrates this software update into the path you specify. Be sure to provide the absolute path to your operating system source files. This option is the same as the **/S** option that was used by previous installers. |
| **/L** | Lists installed updates. |
| **/N**<br>or<br>**/Nobackup** | Does *not* create backup files during setup. We do not recommend this option because backup files allow you to remove SP2 in the future. An entry for the update does not appear in **Add or Remove Programs** if you use this option. |
| **/O**<br>or<br>**/Overwriteoem** | Overwrites OEM files without prompting. |

| | |
|---|---|
| **/Q**<br>or<br>**/Quiet** | Uses quiet mode. This mode shows no user interface during the installation of the updates. This is the same as unattended mode, except that the user interface is hidden. No prompts will appear during the installation process. |
| **/U**<br>or<br>**/Passive** | Installs the update in unattended mode. Only critical error prompts and a progress bar will appear during the installation. By default, the program invokes the **/Warnrestart** option. |
| **/Z**<br>or<br>**/Norestart** | Does not restart the computer after the installation is complete, even if a restart is required to complete the installation. |
| **/?**<br>or<br>**/Help** | Displays Help. |

**Installing updates on multiple computers**

Although there are several methods that you can use to install an update (like using Systems Management Server or Windows Installer), the following procedure outlines how to run Update.exe manually using a distribution folder.

| **Important** |
|---|
| You must install the service pack either in combination with or before you install updates that were released after the service pack. Installing more recently released updates before you install the service pack can leave the system in an unstable state. |

**To install an update by running Update.exe**

1. Connect to the network or computer on which you want to create the distribution folder.
2. In the shared distribution folder, create a folder for the update files. For example, to create a distribution folder named Update, type the following (using the appropriate drive letter):

   mkdir *Drive*:\Update

3. Copy the Windows Server 2003 Update.exe program to the distribution folder that you created in Step 2. For example, to copy the Update.exe program to the distribution folder named Update, type:

   xcopy C:\WindowsServer2003-KB######-x86-LLL.exe *Drive*:\Update

4. To install the update from the distribution folder, run the WindowsServer2003-KB######-x86-*LLL*.exe program. For example, to install the update from the distribution folder named Update, type:

   *Drive*:\Update\WindowsServer2003-KB######-x86-LLL.exe

5. To ensure that your updates take effect, restart your computer after you finish installing all of the updates.

**Installing several updates together**

You can group multiple updates together in a batch file and install them as a unit. In this case, you do not need to restart your computer after you install each update. The update installer provides options that you can use when performing this type of installation. The following code is a sample batch file that installs updates and ensures that the correct files are replaced after the computer is restarted.

```
@echo off
setlocal
set PATHTOFIXES=Drive:\update
%PATHTOFIXES%\WindowsServer2003-KB######-x86-LLL.exe /quiet /norestart
%PATHTOFIXES%\WindowsServer2003-KB######-x86-LLL.exe /quiet /norestart
%PATHTOFIXES%\WindowsServer2003-KB######-x86-LLL.exe /quiet /norestart
```

| **Important** |
| --- |
| To ensure that the updates that you installed take effect, restart the computer if the batch file does not automatically restart it for you. If you want your computer to restart automatically, in the last update that you install, replace the option **/Norestart** with the option **/Forcerestart**. |

**The combination installation**

The combination installation uses the update installation and integrated installation processes to install a service pack with one or more of the following:

- Windows Server 2003
- Other updates (such as feature packs, critical updates, or hotfixes)
- Additional Microsoft and non-Microsoft software that might be included on a service pack CD

For a combination installation, you can include the components that you want to install as entries in the Svcpack.inf file along with the updates. You can also choose to install the service pack at the same time. This section provides instructions for using unattended setup mode to simultaneously install Windows Server 2003 and updates.

**Installing a service pack and updates with Windows Server 2003**

This section explains how to perform a combination installation from a distribution folder. This option installs a service pack and updates that are integrated with Windows Server 2003. A combination installation makes it unnecessary for you to perform separate installations of Windows Server 2003, a service pack, and the updates because it installs all of them together.

**Preparing for the combination installation**

Before you deploy the installation of Windows Server 2003 and the updates across a network, you must copy Windows Server 2003 and the update installation files to a shared folder and complete the steps described in To create and set up the required folders and files for Windows Server 2003.

If you are an original equipment manufacturer (OEM) and you plan to install additional OEM files (such as device driver, application, or component files), create a \$1 subfolder within the \$OEM$ subfolder. The \$1 subfolder corresponds to %systemdrive%, which is the destination drive for the Windows Server 2003 installation.

| **Important** |
| --- |
| If you want to include a service pack *and* updates in your deployment of Windows Server 2003, first you must integrate Windows Server 2003 with the service pack as described in the Windows Server 2003 Service Pack Installation and Deployment Guide that applies to your service pack. (For example, the guide for Service Pack 2 is located at http://go.microsoft.com/fwlink/?LinkID=64780.) Then you can integrate the updates, which will support the service pack. |

For more information about how to integrate updates, see article 828930 in the Microsoft Knowledge Base about how to integrate software updates into your Windows installation source files (http://go.microsoft.com/fwlink/?LinkId=42761).

**To create and set up the required folders and files for Windows Server 2003**

1. Connect to the network or computer on which you want to create the distribution folder.
2. Create a distribution folder by typing the following at the command prompt (using the appropriate drive letter):

   mkdir *Drive*:\WServer2003DIST
3. Use the command **xcopy** to copy the files and subfolders from the integrated installation product disc that includes Windows Server 2003 and the service pack to the *Drive*:\WServer2003DIST folder.

   For example, if D is your CD-ROM drive, type:

   xcopy /E /I /V D:\ *Drive*:\WServer2003DIST
4. Use **/Integrate** to run the update that you want to install by typing the following:

   WindowsServer2003-KB######-x86-LLL.exe /integrate: *Drive*:\WServer2003DIST

5. Repeat this step for each update you want to include your deployment.

**Deploying the combination installation**

You can deploy the integrated installation of Windows Server 2003 and the updates to multiple computers from a distribution folder. During the standard installation process, Setup installs the operating system and applies the updates.

| **Note** |
| --- |
| For more information about the following procedure, see the Microsoft Windows Corporate Deployment Tools User's Guide (Deploy.chm) included in Deploy.cab in the \Support\Tools folder on your service pack CD. |

### To deploy the installation

1. Verify that the installation and executable files for Windows Server 2003 and the updates exist in your distribution folder.
2. Customize Windows Server 2003 Setup as necessary. For more information about customizing Windows Server 2003 Setup, see the **Winnt32.exe** command syntax topic in the Microsoft Windows Server 2003 Professional Resource Kit Documentation on the Windows Deployment and Resource Kits Web site (http://go.microsoft.com/fwlink/?LinkId=40616).
3. Run Windows Server 2003 Setup. This deploys the installation of Windows Server 2003 and the updates to multiple computers from the distribution folder.

**Removing updates**

You can remove an update for Windows Server 2003 at a command prompt or by using **Add or Remove Programs** in **Control Panel**. If you want to remove updates, you must remove them in the reverse order of installation (that is, you must remove the most recently installed update first). For example, if you installed UpdateA before you installed UpdateB, you will need to remove UpdateB first to return your computer to its original state. This is because some updates share files.

| **Note** |
| --- |
| You cannot remove updates that were installed using the combination installation method. |

### To remove an update at the command prompt

1. Specify this location at the command prompt:

   %systemroot%\$NtUninstallKB######$\spuninst\

2. Type "Spuninst.exe */option*". Replace *option* with any of the following options:

| Option | Description |
| --- | --- |
| **/U**<br>or<br>**/Passive** | Removes the update in unattended mode.<br>(If you use this option, only critical error prompts appear while you uninstall the update.) |
| **/Q**<br>or<br>**/Quiet** | Removes the update in quiet mode.<br>(This is the same as unattended mode except that the user interface is hidden. If you use this option, no prompts appear while you uninstall the update.) |
| **/Z**<br>or<br>**/Norestart** | Does not restart the computer after the process of uninstalling the update is complete. |

| | |
|---|---|
| **/F**<br>or<br>**/Forceappsclose** | Forces other applications to close when the computer restarts after uninstalling the update. |
| **/?**<br>or<br>**/Help** | Displays Help. |

**To remove an update for Windows Server 2003 by using Add or Remove Programs**

1. Open **Control Panel**.
2. Click **Add or Remove Programs**.
3. Select the **Show updates** check box.
4. Click the update that you want to remove, and then click **Change or Remove**.
5. Follow the instructions that appear on your screen.

**Copyright**
Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2006 Microsoft Corporation. All rights reserved.

Microsoft, MS-DOS, Windows, Windows NT, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

# Readme for Windows Server 2003 Service Pack 2

Applies To: Windows Server 2003 with SP2

This document contains information that you should know before you install Microsoft® Windows Server® 2003 Service Pack 2 (SP2). There are three versions of SP2: one for 32-bit operating systems, one for Itanium-based operating systems, and one for x64-based operating systems. SP2 contains the latest collection of updates to help improve the security, reliability, and performance of the following operating systems:

- 32-bit versions of Microsoft Windows Server 2003
- Itanium-based versions of Microsoft Windows Server 2003
- x64-based versions Microsoft Windows Server 2003
- Windows Server 2003 R2 Editions
- KN versions of Windows Server 2003 R2
- Windows® XP Professional x64 Edition

| **Note** |
| --- |
| If you are running Windows Server 2003, Datacenter Edition, and subscribe to the Windows High Availability Support Program through a qualified original equipment manufacturer (OEM), do not install this service pack. Instead, you should contact your OEM for the service pack, updated adapter drivers, utilities, a custom hardware abstraction layer (HAL) if needed, and support. OEMs are required to test Windows Server 2003 Datacenter Edition with service packs. If you do not subscribe to this program, you can continue installing this service pack. |

**Before you install**

Before you install SP2:

- **Review the following documentation.**
  - Article 914962 in the Microsoft Knowledge Base, for a list of the updates included in this service pack (http://go.microsoft.com/fwlink/?LinkId=74380).
  - "Release Notes for Windows Server 2003 Service Pack 2" (http://go.microsoft.com/fwlink/?LinkId=74381), for a list of the known issues with this version of SP2.
  - "Installing and Deploying Updates for Microsoft Windows Server 2003" (http://go.microsoft.com/fwlink/?LinkID=64779), for information about combining this service pack with other updates.
  - "Windows Server 2003 Service Pack 2 Installation and Deployment Guide" (http://go.microsoft.com/fwlink/?LinkID=64780), for more deployment options.
- **Check the disc space requirements.** See article 926029 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=74454).
- **Back up your servers**. Your backup should include all data and configuration information that is necessary for that server to function. It is important to perform a backup of configuration information for servers, especially those that provide network infrastructure, such as Dynamic Host Configuration Protocol (DHCP). When you perform the backup, be sure to include the boot and system partitions and the system state. Another way to back up configuration information is to create a backup set for Automated System Recovery.
- **Disconnect UPS devices.** If you have an uninterruptible power supply (UPS) connected to your destination computer, disconnect the connecting serial cable before running Setup. Setup automatically attempts to detect devices connected to

serial ports, and UPS equipment can cause problems with the detection process.

- **Disable your antivirus software.** If you are not connected to a network or the Internet and you want to minimize the installation time, you can disable your antivirus software. The time that it takes for the antivirus software to scan each file can add significant time to the installation: as much as an hour or more on slower computers. Be sure to enable the antivirus software after you install SP2.

- **Uninstall prerelease versions of SP2.** Unless you have a supplementary support agreement with Microsoft, you must uninstall any prerelease versions of SP2 before you proceed. Upgrading to this version from any prerelease version is not supported.

- **Close all open programs.**

**How to install SP2 from the Web**

You can install SP2 from:

- **Windows Update:** You can use Microsoft Update to install SP2 on a single server running Windows Server 2003. You can update Itanium-based versions with SP2 only by using the CD or by downloading it from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=77236).

- **Microsoft Download Center:** You can download SP2 from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=77236). We recommend that you save the SP2 installation file locally on the computer. Then if you have to restart the installation or reinstall SP2, you will not need to download the installation file again. Because the SP2 installation file is very large, if your computer has more than one hard disk, save the installation file on a disk other than your system disk.

**Deploying SP2 to multiple servers**

You can use the SP2 installation files to deploy SP2 to multiple servers that are already running one of the supported operating systems. And you can use the SP2 files to integrate a Windows Server 2003 operating system with SP2 and then install the operating system and SP2 together on multiple servers. This procedure is known as "slipstreaming." For more information, see:

- Installing and Deploying Updates for Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkID=64779)

- Windows Server 2003 Service Pack 2 Installation and Deployment Guide (http://go.microsoft.com/fwlink/?LinkID=64780)

**Command-line options**

You can define the way SP2 is installed by typing one or more of the following command-line options when you run the installation .exe file. You can precede the options with either a forward slash (/) or a hyphen (-). The options are not case-sensitive.

| **Note** |
|---|
| You can also use the **/Log** and **/Verbose** options. For more information, see "The Package Installer (Formerly Called Update.exe) for Microsoft Windows Operating Systems and Windows Components" (http://go.microsoft.com/fwlink/?LinkId=42759). |

| Command-line option | Description |
|---|---|
| **/D:Path** | Backs up the files that are required to remove the service pack to *Path*. Any location that you specify must be on the local drive. If you back up files but do not specify a folder, the default location is:<br><br>$ntservicepackuninstall$ |

| | |
|---|---|
| **/ER** | Enables extended codes that return a detailed error message if the update installation fails. For more information, see "The Package Installer (Formerly Called Update.exe) for Microsoft Windows Operating Systems and Windows Components" (http://go.microsoft.com/fwlink/?LinkId=42759). |
| **/F**<br>or<br>**/Forceappsclose** | Forces other applications to close after you install the update package and before the computer is restarted. You cannot use this option with **/integrate**, **/L**, and **/Z**. |
| **/Forcerestart** | Restarts the computer after the installation is complete, even if a restart is not required to complete the installation. This option does not force applications to close. |
| **/Integrate:Path** | Integrates the SP2 source files into the path you specify. Be sure to provide the absolute path to your operating system source files. This option is the same as the **/S** option that was used by previous installers. |
| **/L** | Lists installed updates. |
| **/N**<br>or<br>**/Nobackup** | Does *not* create backup files during setup. We do not recommend this option because backup files allow you to remove SP2 in the future. An entry for SP2 does not appear in **Add or Remove Programs** if you use this option. |
| **/O**<br>or<br>**/Overwriteoem** | Overwrites OEM files without prompting. |
| **/Promptrestart** | If a restart is required, a dialog box notifies the user that a restart is required to complete setup. Intended for use with /Quiet. |
| **/Q**<br>or<br>**/Quiet** | Uses quiet mode. This mode shows no user interface during the installation of the updates. This is the same as unattended mode, except that the user interface is hidden. No prompts will appear during the installation process. |
| **/U**<br>or<br>**/Passive** | Installs the service pack in unattended mode. Only critical error prompts and a progress bar will appear during the installation. By default, the program invokes the **/Warnrestart** option. |
| **/Warnrestart: [Seconds]** | If a restart is required, a dialog box notifies the user that the computer will restart in the specified seconds. The user can click **Cancel** or **Restart Now**. The default is 60 seconds. |
| **/X**<br>or<br>**/Extract** | Extracts service pack files without starting Update.exe. You are prompted to provide the path for the folder in which you want the extracted service pack installation files to be placed. For more information, see "The Package Installer (Formerly Called Update.exe) for Microsoft Windows Operating Systems and Windows Components" (http://go.microsoft.com/fwlink/?LinkId=42759). |
| **/Z**<br>or<br>**/Norestart** | Does not restart the computer after the installation is complete, even if a restart is required to complete the installation. |
| **/?**<br>or<br>**/Help** | Displays Help. |

**Uninstalling SP2**

It is possible to uninstall SP2 only if it was installed as an update to an existing operating system. It is not possible to remove

only SP2 if you installed it as part of an integrated installation. In addition, you can only uninstall the last service pack that you installed. For example, if you install SP1 and then you install SP2, you can only uninstall SP2.

By default, backup files are created during the installation so you can remove SP2. Unless you specified the /nobackup option when installing SP2, then you can remove SP2 using **Add or Remove Programs** or the command prompt.

| **Important** |
| --- |
| If you remove SP2, a dialog box displays a list of the applications that you installed after you installed SP2. If you continue, these applications might not work correctly. Multiple descriptions of some applications might appear instead of a single generic description. You can ignore these additional descriptions. |

**To remove SP2 using a command prompt**

1. Open a command prompt, type the following, and then press ENTER (*Option* can be any option listed in the following table):

   %systemroot%\$NtServicePackUninstall$\spuninst\Spuninst.exe /*Option*

2. Follow the instructions that appear on your screen.

| Option | Description |
| --- | --- |
| **/F**<br>or<br>**/Forceappsclose** | Forces applications to close when the computer restarts. |
| **/Q**<br>or<br>**/Quiet** | Removes SP2 in quiet mode, which is the same as unattended mode except that the user interface is hidden and no prompts appear. |
| **/U**<br>or<br>**/Passive** | Removes the service pack in unattended mode. Only critical error prompts appear while SP2 is being uninstalled. |
| **/Z**<br>or<br>**/Norestart** | Does not restart the computer after the process of uninstalling SP2 is complete. |
| **/?**<br>or<br>**/Help** | Displays Help. |

**Additional resources**

- For more information about Windows Server 2003 SP2, see the Microsoft TechCenter Web site (http://go.microsoft.com/fwlink/?LinkID=64780)

- For the most up-to-date information about supported hardware, see the hardware compatibility information at the Windows Server Catalog Web site (http://go.microsoft.com/fwlink/?LinkID=62603)

- To obtain the latest product updates, see the Microsoft Update Web site (http://go.microsoft.com/fwlink/?LinkID=54836)

**Copyright**

Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2006 Microsoft Corporation. All rights reserved.

Microsoft, MS-DOS, Windows, Windows NT, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

# Release Notes for Microsoft Windows Server 2003 Service Pack 2

Applies To: Windows Server 2003 with SP2

These release notes address late-breaking issues and information about the Microsoft® Windows Server® 2003 operating systems with Service Pack 2 (SP2) and Microsoft Windows® XP Professional x64 Edition operating system with SP2. Unless otherwise specified, these notes apply to the following operating systems:

- Windows Server 2003 with SP2
- Windows Server 2003 R2 with SP2
- x64-based versions of Windows Server 2003 with SP2
- x64-based versions of Windows Server 2003 R2 with SP2
- Itanium-based versions of Windows Server 2003 with SP2
- Windows XP Professional x64 Edition with SP2

**Additional information**

- For information about issues that have been resolved in Windows Server 2003 with SP2, see article 914962 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=82243).
- For information about Windows Server 2003 SP2 support tools, see article 926027 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=82244).
- For information about how to obtain the latest service pack, see article 889100 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=82245).
- To obtain an updated version of the System Preparation tool for Windows Server 2003 SP2, see article 892778 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=82247).

**Active Directory**

If you are running the initial version of Active Directory Application Mode (ADAM) that was available as a download, you will not be able to directly update the computer to Windows Server 2003 Service Pack 2 (SP2). If you do so, the ADAM instances may stop functioning.

**Preventing the issue**

To prevent this, do one of the following before you update the computer to Windows Server 2003 SP2:

- Download and install Windows Server 2003 Active Directory Application Mode Service Pack 1 (ADAM SP1). Follow the instructions provided in Microsoft Knowledge Base article 902838 (http://go.microsoft.com/fwlink/?LinkId=82253). Then, install Windows Server 2003 SP2.
- If you have the first release of Windows Server 2003 R2, then update ADAM as follows. (Disc 1 contains SP1. If the product discs that you have do not specify which service pack is included, then the discs contain SP1.)

    1. Install both disks.
    2. Update the ADAM instances by installing ADAM using **Add/Remove Windows Components** in **Add or Remove Programs** in Control Panel.

3. Install Windows Server 2003 SP2.

**Correcting the issue**
If you have already installed Windows Server 2003 SP2 (for example, from Windows Update or by using Update.exe) without first updating ADAM, you can correct this issue with the following procedure:

### To repair ADAM

1. Uninstall Windows Server 2003 SP2.
2. Download and install ADAM SP1 by following the instructions provided in Microsoft Knowledge Base article 902838 (http://go.microsoft.com/fwlink/?LinkID=82253).
3. Reinstall Windows Server 2003 SP2.

**Updating ADAM**
If you have already installed the integrated version of Windows Server 2003 with SP2 (for example, if you installed Disc 1 of the second release of Windows Server R2, which contains SP2), and you did not update ADAM first, the ADAM instances will continue to function properly.

However, you can update ADAM by doing one of the following:

- Download and install ADAM SP1 by following the instructions provided in Microsoft Knowledge Base article 902838 (http://go.microsoft.com/fwlink/?LinkID=82253).
- Install Windows Server 2003 R2 Disc 2, and then install ADAM by using the **Add/Remove Windows Components** in Control Panel.

| **Note** |
| --- |
| For more information about the options for upgrading to the second release of Windows Server 2003 R2, see the upgrade diagram in Upgrading When You Have the SP2 Version of Windows Server 2003 R2 (http://go.microsoft.com/fwlink/?LinkId=85407). |

**Installation**
This release note applies to you if you are upgrading to SP2 from a computer running SP1 or installing a different edition of Windows Server 2003 with SP2 on a computer already running SP2.

If you are installing SP2 on a computer running SP1, to ensure that certain applications and components (such as Windows .NET Framework 2.0) function properly, install SP2 by using Microsoft Update (http://go.microsoft.com/fwlink/?LinkId=85856).

If you are installing a different edition of SP2 on a computer already running SP2 (for example, if you are installing Windows Server 2003, Enterprise Edition with SP2 on a computer running Windows Server 2003, Standard Edition with SP2), there is no workaround. If any applications or components stop functioning, you may need to reinstall them.

**Internet Explorer**
You do not need to uninstall Internet Explorer 7 prior to installing Windows Server 2003 SP2.

If you have previously installed Internet Explorer 7 on a computer running Windows Server 2003 with SP1 and are now installing Windows Server 2003 SP2 using Microsoft Update, the downloadable SP2 update from the Microsoft Download Center, or the SP2 update CD, your computer will retain Internet Explorer 7 as expected.

If you have previously installed Internet Explorer 7 on a computer running Windows Server 2003 with SP1 and are now installing the integrated version of Windows Server 2003 SP2 (that is, the full operating system integrated with the service pack), the computer will have Internet Explorer 6 as the default browser in Windows Server 2003 SP2. You will need to reinstall Internet Explorer 7.

For more information about options for installing Windows Server 2003 SP2, see http://go.microsoft.com/fwlink/?LinkId=85899.

**Microsoft Management Console**
If you start certain snap-ins for Microsoft Management Console 3.0 and then subsequently uninstall Windows Server 2003 SP2, you may not be able to start some or all of the snap-ins in Microsoft Management Console 2.0 afterwards. If

this occurs, remove all the files in the %APPDATA%\Microsoft\MMC\ folder.

**Migration**
If you use Active Directory Migration Tool (ADMT) version 2.0 to migrate Windows 2000 Professional domain accounts between computers that are running Windows Server 2003 SP1 or Windows Server 2003 SP2, you may encounter an "RPC server unavailable" error message.

To avoid this, download and use Active Directory Migration Tool version 3.0 from http://go.microsoft.com/fwlink/?LinkID=68791.

**Printing**
To avoid interrupting print services, schedule the installation of Windows Server 2003 SP2 for a time when the print spooler will not be in high demand.

For more details about interruption to printing, see article 832219 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=82276).

**Scheduled tasks**
If you uninstall Windows Server 2003 SP2, scheduled tasks may subsequently fail because of the loss of administrative credentials. To verify that this issue is the cause of tasks failing, check the Task Scheduler log for error 0x8004130f.

**Preventing the issue**
To avoid this issue, ensure that you stop the Task Scheduler service before you uninstall Windows Server 2003 SP2.

A hotfix for this issue is available in article 933550 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=85408). The hotfix only applies to certain configurations and scenarios, so be sure to read the hotfix documentation carefully.

**Correcting the issue**
If your scheduled tasks have already failed as a result of this issue, you can restore them by re-entering the appropriate administrative credentials for each task using the Task Scheduler.

**Windows Small Business Server 2003**
Before you install Windows Server 2003 SP2 on a computer running Windows Small Business Server 2003 (Windows SBS), see article 939421 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=93765).

**Windows Media Player**
If you uninstall Windows Server 2003 SP2, you may not be able to access some Windows Media content that is protected by Digital Rights Management. If this occurs, you can restore access by doing the following:

| Operating system | Action |
|---|---|
| All Windows Server 2003 versions | Reinstall Service Pack 2 |
| Windows XP Professional x64 Edition | Do either of the following:<br>• Reinstall Service Pack 2<br>• Install Windows Media Player 11 |

For more detailed information about this issue, see article 932542 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=83575).

**Copyright**
Information in this document, including URL and other Internet Web site references, is subject to change without notice and is provided for informational purposes only. The entire risk of the use or results from the use of this document remains with the user, and Microsoft Corporation makes no warranties, either express or implied. Unless otherwise noted, the companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted in examples herein are fictitious. No association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject

matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2007 Microsoft Corporation. All rights reserved.

Microsoft, Windows, Windows XP, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

All other trademarks are property of their respective owners.

# Comparison of Windows Server 2003 Editions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following table describes features supported by the Microsoft® Windows Server™ 2003 operating system, and illustrates which editions of the operating system support which features.

● = Feature included

◒ = Feature partially supported

○ = Feature not included

| Category | Feature and Description | Standard Edition | Enterprise Edition | Datacenter Edition | Web Edition |
|---|---|---|---|---|---|
| Hardware Specifications | **64-bit Support for Intel Itanium-based Computers**(1)<br>Support for 64-bit processing delivers far higher scalability than 32-bit file servers by providing a greatly enlarged virtual address space and paged pool area, the ability to handle increased numbers of users and connections, and increased hardware reliability through predictive error checking and notification of failures. | ○ | ● | ● | ○ |
| Hardware Specifications | **Hot Add Memory**(2, 3)<br>Hot Add Memory allows ranges of memory to be added to a computer and made available to the operating system and applications as part of the normal memory pool. This does not require rebooting the computer and involves no downtime. This feature only operates on servers that have hardware support for adding memory while the server is operating. Most existing servers do not have such hardware support and can be damaged if memory is installed while the power is on. It is recommended that you consult your server operator's manual for more information. | ○ | ● | ● | ○ |
| Hardware Specifications | **Non-Uniform Memory Access**(3)<br>As processor clock rates continue to increase and put pressure on processor bus architectures, scaling is addressed by implementing multiple processor buses. This can result in an architecture consisting of processors and memory arranged in smaller subsystems called nodes. Processor access time to memory in other nodes is longer than access time to memory in the same node. This results in Non-Uniform Memory Access (NUMA) across the system. The longer access times to other nodes can degrade software performance. The operating system attempts to limit degradation by scheduling threads from the same process on processors that are in the same node, and allocating all memory requests within the same node as the processor making the request. In addition, an API is included to make NUMA architecture information available to application software. These features ensure that memory accesses are local to a node wherever possible, and limit software degradation caused by the NUMA architecture. | ○ | ● | ● | ○ |

| | | | | | |
|---|---|---|---|---|---|
| Hardware Specifications | **Datacenter Program**<br><br>The Datacenter Program provides customers with an integrated hardware, software, and service offering, delivered by Microsoft and qualified server vendors such as original equipment manufacturers (OEMs). | ○ | ○ | ● | ○ |
| Hardware Specifications | **2-GB RAM Maximum**<br><br>RAM facilitates improved system scalability and performance. The more RAM added to a server beyond minimum requirements, the more memory available for applications to use. Designed for building and hosting Web applications, Web pages, and XML Web services, Windows Server 2003, Web Edition supports new systems with up to 2 GB of RAM. | ○ | ○ | ○ | ● |
| Hardware Specifications | **4-GB RAM Maximum**<br><br>RAM facilitates improved system scalability and performance. The more RAM added to a server beyond minimum requirements, the more memory available for applications to use. Designed for small organizations and departmental use, Windows Server 2003, Standard Edition supports new systems with up to 4 GB of RAM. | ● | ○ | ○ | ○ |
| Hardware Specifications | **32-GB RAM Maximum**<br><br>RAM facilitates improved system scalability and performance. The more RAM added to a server beyond minimum requirements, the more memory available for applications to use. Designed for demanding enterprise applications, Windows Server 2003, Enterprise Edition supports new systems with up to 32 GB of RAM. | ○ | ● | ○ | ○ |
| Hardware Specifications | **64-GB RAM Maximum**(4)<br><br>RAM facilitates improved system scalability and performance. The more RAM added to a server beyond minimum requirements, the more memory available for applications to use. Designed for mission-critical applications, the 32-bit version of Windows Server 2003, Datacenter Edition supports up to 64 GB of RAM on x86-based computers. The 64-bit version of Windows Server 2003, Enterprise Edition supports systems with up to 64 GB of RAM. | ○ | ◒ | ● | ○ |
| Hardware Specifications | **512-GB RAM Maximum**(5)<br><br>RAM facilitates improved system scalability and performance. The more RAM added to a server beyond minimum requirements, the more memory available for applications to use. The 64-bit version of Windows Server 2003, Datacenter Edition supports systems with up to 512 GB of RAM. | ○ | ○ | ◒ | ○ |

| | | | | | |
|---|---|---|---|---|---|
| Hardware Specifications | **2-Way Symmetric Multiprocessing**<br>The Windows Server 2003 family supports single or multiple CPUs that conform to the symmetric multiprocessing (SMP) standard. Using SMP, the operating system can run threads on any available processor, which makes it possible for applications to use multiple processors when additional processing power is required to increase the capability of a system. New features include SMP locking performance, improved registry performance, and increased Terminal Server sessions. Designed for building and hosting Web applications, Web pages, and XML Web services, Windows Server 2003, Web Edition supports new systems with up to two-way SMP. | ● | ● | ○ | ● |
| Hardware Specifications | **4-Way Symmetric Multiprocessing**<br>The Windows Server 2003 family supports single or multiple CPUs that conform to the symmetric multiprocessing (SMP) standard. Using SMP, the operating system can run threads on any available processor, which makes it possible for applications to use multiple processors when additional processing power is required to increase the capability of a system. New features include SMP locking performance, improved registry performance, and increased Terminal Server sessions. Designed for small organizations and departmental use, Windows Server 2003, Standard Edition supports new systems with up to four-way SMP. | ● | ● | ○ | ○ |
| Hardware Specifications | **8-Way Symmetric Multiprocessing**<br>The Windows Server 2003 family supports single or multiple CPUs that conform to the symmetric multiprocessing (SMP) standard. Using SMP, the operating system can run threads on any available processor, which makes it possible for applications to use multiple processors when additional processing power is required to increase the capability of a system. New features include SMP locking performance, improved registry performance, and increased Terminal Server sessions. Designed for demanding enterprise applications, Windows Server 2003, Enterprise Edition supports new systems with up to eight-way SMP. This includes the 64-bit edition of Windows Server 2003, Enterprise Edition. | ○ | ● | ● | ○ |
| Hardware Specifications | **32-Way Symmetric Multiprocessing**<br>The Windows Server 2003 family supports single or multiple CPUs that conform to the symmetric multiprocessing (SMP) standard. Using SMP, the operating system can run threads on any available processor, which makes it possible for applications to use multiple processors when additional processing power is required to increase the capability of a system. New features include SMP locking performance, improved registry performance, and increased Terminal Server sessions. Designed for mission-critical applications, Windows Server 2003, Datacenter Edition supports new systems with 8-way to 32-way SMP. This includes the 64-bit edition of Windows Server 2003, Datacenter Edition. | ○ | ○ | ● | ○ |

| | | | | | |
|---|---|---|---|---|---|
| Hardware Specifications | **64-Way Symmetric Multiprocessing**<br><br>The Windows Server 2003 family supports single or multiple CPUs that conform to the symmetric multiprocessing (SMP) standard. Using SMP, the operating system can run threads on any available processor, which makes it possible for applications to use multiple processors when additional processing power is required to increase the capability of a system. New features include SMP locking performance, improved registry performance, and increased Terminal Server sessions. Designed for mission-critical applications, Windows Server 2003, Datacenter Edition supports new systems with up to 64-way SMP. The support for over 32 processors is only available with the 64-bit edition of Windows Server 2003, Datacenter Edition.<br><br>Microsoft also offers a 128-way SKU for Windows Server 2003, Datacenter Edition so Windows can run on a 128-processor computer. However, the largest partition supported would be 64 processors. | ○ | ○ | ● | ○ |
| Directory Services | **Active Directory® Directory Service**<br><br>Active Directory is the directory service for Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. It stores information about objects on the network and makes this information easy for administrators and users to find and use. Active Directory uses a structured data store as the basis for a logical, hierarchical organization of directory information. | ● | ● | ● | ◒ |
| Directory Services | **Microsoft Identity Integration Server Support**<br><br>Microsoft Identity Integration Server (MIIS) is a centralized service that stores and integrates identity information from multiple directories in an organization. The goal of this metadirectory is to provide an organization with a unified view of all known identity information about users, applications, and network resources. A metadirectory solves important business issues that result from having information stored in multiple, disparate data repositories throughout an organization. MIIS is available via Microsoft Consulting Services (MCS) or via an MIIS partner engagement. | ○ | ● | ● | ○ |
| Security Services | **Internet Connection Firewall**<br><br>Internet Connection Firewall (ICF) provides Internet security in the form of a firewall. Designed for use in the home and small business, ICF provides protection on computers directly connected to the Internet. This feature is available for LAN or dial-up networks, virtual private networks (VPNs), and Point-to-Point Protocol over Ethernet (PPPoE) connections. It also prevents scanning of ports and resources (such as file and printer shares) from external sources. | ● | ● | ○ | ○ |

| | | | | | |
|---|---|---|---|---|---|
| Security Services | **Public Key Infrastructure, Certificate Services, and Smart Cards**<br><br>By using Certificate Services and certificate management tools, you can deploy your own public key infrastructure (PKI). With a PKI, you can implement standards-based technologies, such as smart card logon capabilities, client authentication through Secure Sockets Layer (SSL) and Transport Layer Security (TLS), secure e-mail, digital signatures, and secure connectivity using Internet Protocol security (IPSec). By using Certificate Services, you can set up and manage certification authorities (CAs) that issue and revoke X.509 v3 certificates. This means that you do not have to depend on commercial client authentication services, although you can integrate commercial client authentication into your PKI if you choose. | ◒ | ● | ● | ◒ |
| Terminal Services | **Remote Desktop for Administration**<br><br>With Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode), you can administer a computer from virtually any computer on your network. Based on Terminal Services technology, Remote Desktop for Administration is specifically designed for server management. | ● | ● | ● | ● |
| Terminal Services | **Terminal Server**<br><br>Terminal Server lets you deliver Windows-based applications — or the Windows desktop itself — to virtually any computing device, including those that cannot run Windows. For example, a user can access a virtual Windows XP Professional desktop and x86-based Windows applications from hardware that cannot run the software locally. Terminal Server provides this capability for both Windows-based and non–Windows-based client devices. When a user runs an application on Terminal Server, all of the application execution takes place on the server, and only keyboard, mouse, and display information traverses the network.<br><br>**Note**<br><br>• Terminal Server mode is not included on computers running Microsoft Windows Server 2003, Web Edition; however, Remote Desktop for Administration is available on Windows Server 2003, Web Edition. | ● | ● | ● | ○ |
| Terminal Services | **Terminal Server Session Directory**<br><br>Terminal Server Session Directory allows users to easily reconnect to a disconnected session in a load-balanced Terminal Server farm. Session Directory is compatible with the Windows Server 2003 Network Load Balancing service, and is supported by third-party external load-balancer products from manufacturers such as F5 Networks (formerly F5 Labs) and Radware.<br><br>**Note**<br><br>• The Session Directory Service runs on all editions of Windows Server 2003; however, to participate in a Session Directory, the server must be running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition (including the 64- bit editions of the Windows Server 2003 family). | ○ | ● | ● | ○ |

| | | | | | |
|---|---|---|---|---|---|
| Clustering Technologies | **Network Load Balancing**<br><br>Previously known as Windows NT Load Balancing Service (WLBS), Network Load Balancing distributes incoming TCP/IP traffic among multiple servers. Your clustered applications, especially Web server applications, can handle more traffic, provide higher availability, and provide faster response times. | ● | ● | ● | ● |
| Clustering Technologies | **Cluster Service**<br><br>A cluster is a group of independent computers, called nodes, that work together to run a common set of applications and provide high availability. If one node on the cluster fails, the application can be failed over to the next node. Server clusters of up to eight nodes are available only in Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition. | ○ | ● | ● | ○ |
| Communications and Networking Services | **Virtual Private Network Support**<br><br>You can give users ready access to your organization's network even when they are out of the office — and reduce the cost of this access — by implementing a virtual private network (VPN). The VPN connection creates a secure tunnel across the Internet into the private network. There are two types of VPN technologies in the Windows Server 2003 family: Point-to-Point Tunneling Protocol (PPTP), which employs user-level Point-to-Point Protocol (PPP) authentication methods and Microsoft Point-to-Point Encryption (MPEE) for data encryption; and Layer Two Tunneling Protocol (L2TP) with Internet Protocol security (IPSec). L2TP employs user-level PPP authentication methods and computer-level certificates with IPSec for data encryption. On Windows Server 2003, Web Edition and Windows Server 2003, Standard Edition, you can create up to 100 PPTP ports and up to 100 L2TP ports. However, Windows Server 2003, Web Edition, can accept only one VPN connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections via the ports. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000. Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition support unlimited concurrent users. | ● | ● | ● | ◒ |
| Communications and Networking Services | **Internet Authentication Service**<br><br>Internet Authentication Service (IAS) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy. As a RADIUS server, IAS performs centralized connection authentication, authorization, and accounting for many types of network access including wireless, authenticating switch, and remote access dial-up and virtual private network (VPN) connections. As a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers. You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range. | ● | ● | ● | ○ |

| | | | | | |
|---|---|---|---|---|---|
| Communications and Networking Services | **Network Bridge**<br><br>Network Bridge provides an easy way for you to connect different LAN segments, allowing users to bridge connections between different computers and devices on their network, even when they connect to the network through different methods. Network bridge is supported in Windows Server 2003, Standard Edition and Windows Server 2003, Enterprise Edition. | ● | ● | ○ | ○ |
| Communications and Networking Services | **Internet Connection Sharing**<br><br>By using the Internet Connection Sharing (ICS) feature of Network Connections, you can connect your home network or small office network to the Internet. For example, you might have a home network that connects to the Internet through a dial-up connection. By enabling ICS on the computer that uses the dial-up connection, you can provide network address translation (NAT), addressing, and name resolution services for all the computers on your network. | ● | ● | ○ | ○ |
| Communications and Networking Services | **IPv6**<br><br>Internet Protocol version 6 (IPv6) is a suite of Internet standard protocols that will become the next generation of network layer protocols for the Internet. IPv6 is designed to solve many of the problems of the current version of IP (known as IPv4) with regard to address depletion, security, autoconfiguration, extensibility, and more. | ● | ● | ● | ● |
| File and Print Services | **Distributed File System**<br><br>Distributed File System (DFS) is enhanced for Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition by allowing multiple DFS roots on a single server. You can use this feature to host multiple DFS roots on a single server, reducing administrative and hardware costs of managing multiple namespaces and multiple replicated namespaces. Using Active Directory, DFS shares can be published as volume objects and administration can be delegated. Windows Server 2003, Standard Edition supports only one DFS root. DFS is partially supported in Windows Server 2003, Web Edition, where it is possible to access DFS files and act as a node in DFS tree, but only 10 concurrent incoming server message block (SMB) connections are permitted. | ● | ● | ● | ◒ |
| File and Print Services | **Encrypting File System**<br><br>Encrypting File System (EFS) complements other access controls and provides an additional level of protection for your data. EFS runs as an integrated system service on all disks (including clustered disks), making it easy to manage, difficult to attack, and transparent to the user. | ● | ● | ● | ● |
| File and Print Services | **Shadow Copies for Shared Folders**<br><br>Shadow Copy Restore (Previous Versions) provides point-in-time copies for network folders. Users can easily access previous versions of their files through Windows Explorer by right-clicking a file or folder. | ● | ● | ● | ● |

| | | | | | |
|---|---|---|---|---|---|
| File and Print Services | **Removable Storage**<br><br>Removable Storage makes it easy to track your removable storage media (tapes and optical discs) and to manage the hardware libraries (such as changers and jukeboxes) that contain them. Because removable optical discs and tapes are less expensive per megabyte than hard disks, Removable Storage and Remote Storage can decrease your costs. | ● | ● | ● | ● |
| File and Print Services | **Remote Storage**<br><br>Remote Storage uses criteria that you specify to automatically copy little-used files to removable media. If hard disk space drops below specified levels, Remote Storage removes the cached file content from the disk. If the file is needed later, the content is automatically recalled from storage. Remote Storage now supports migration to magneto-optical media. Remote Storage is not available in Windows Server 2003, Standard Edition or Windows Server 2003, Web Edition. | ○ | ● | ● | ○ |
| File and Print Services | **Fax Service**<br><br>Fax Service lets users send and receive faxes by using a modem or a fax board. You can also send faxes by using your network. You can print to fax from any application, send a cover page, and track and monitor fax activity. New wizards simplify configuration and fax sending. Also, developers can write applications using the new fax APIs to automatically send faxes from their applications. | ● | ● | ● | ○ |
| File and Print Services | **Services for Macintosh**<br><br>Services for Macintosh provides Macintosh users with access to files stored on a computer running Windows Server 2003. The file server is accessible over TCP/IP networks and over AppleTalk networks. Print services allow Macintosh clients to print to Windows NT or Windows 2000–based print shares via the AppleTalk protocol. In addition to the print server, there is a 300 dpi PostScript RIP engine that allows Macintosh-generated PostScript print jobs to be sent to non-PostScript printers, such as inkjet printers. | ● | ● | ● | ○ |
| Management Services | **IntelliMirror® Management Technologies**<br><br>To help reduce costs, administrators need high levels of control over portable and desktop systems. IntelliMirror management technologies provide this control on client systems running Windows 2000 Professional or Windows XP Professional. You can use IntelliMirror to define policies based on business roles, group memberships, and locations. With these policies, Windows 2000 Professional desktops and Windows XP Professional desktops are automatically reconfigured to meet a specific user's requirements each time that user logs on to the network, regardless of where the user logs on. | ● | ● | ● | ◒ |

| | | | | | |
|---|---|---|---|---|---|
| Management Services | **Group Policy Results**<br><br>Group Policy Results allows administrators to see the effect a Group Policy setting has on a targeted user or computer. Included in the Group Policy Management Console, Group Policy Results provides administrators with a powerful and flexible base-level tool to plan, monitor, and troubleshoot Group Policy.<br><br>Group Policy Results uses Windows Management Instrumentation (WMI), making the information available to administrators via the console and directly to other management applications via WMI. | ● | ● | ● | ◒ |
| Management Services | **Windows Management Instrumentation Command Line**<br><br>Windows Management Instrumentation (WMI) provides unified access to the management functions of local and remote systems. By adding command-line access to WMI, administrators can directly access these management functions and create queries based on this data. You can monitor local and remote Windows Server 2003 and Windows XP systems directly, to view results through the command line or to retrieve the management data in XML format and processed into built-in or custom XSL output formats. WMI follows the specifications of the industry-standard Common Information Model (CIM) defined by the Distributed Management Task Force (DMTF). | ● | ● | ● | ● |
| Management Services | **Remote OS Installation**<br><br>Remote OS Installation uses Group Policy, Remote Installation Services, and Pre-Boot eXecution Environment (PXE) server hardware to re-image a server with a clean install of a Windows Server 2003–based environment. This feature can also be used to re-image Windows 2000 and Windows XP desktops.<br><br>Administrators can use Remote OS Installation and IntelliMirror together to simplify the task of exchanging or bringing new computers into the network environment:<br><br>• Remote OS Installation can establish a full initial working set image directly to the computer hardware.<br><br>• IntelliMirror can restore policy-based settings for data, settings, and software use.<br><br>Used together with IntelliMirror or on its own, Remote OS Installation can increase the efficiency of computer management in your organization while simplifying the task of maintaining corporate-standard environments on Windows-based servers and desktops. | ● | ● | ● | ● |
| Management Services | **Remote Installation Services**<br><br>You can use Remote Installation Services (RIS) to create automated installation images of operating systems or of complete computer configurations. You can then make these installation images available to users at client computers. You can also specify which RIS server will provide installations to a given client computer. The client computers must support remote booting with the Pre-Boot eXecution Environment (PXE) ROM, or they must be started with a remote-startup floppy disk. | ● | ● | ● | ○ |

| | | | | | |
|---|---|---|---|---|---|
| Management Services | **Windows System Resource Manager**<br><br>Windows System Resource Manager (WSRM) provides resource management and enables the allocation of resources, including processor and memory resources, among multiple processes based on business priorities. An administrator sets targets for the amount of hardware resources that running applications or users (typically in a Terminal Server environment) are allowed to consume. It also creates utilization accounting records for management, service level agreement (SLA) tracking, and enabling charge-backs. | ○ | ● | ● | ○ |
| .NET Application Services | **.NET Framework**(2)<br><br>The .NET Framework enables your developers to create great Web applications with the help of ASP.NET and other technologies. It also helps them build the same type of applications they design and develop today. The .NET Framework is language-neutral; virtually any programming language can target it. Developers can build .NET-based applications and services in a number of languages, including Visual C++, Visual Basic .NET, JScript, and Visual C#. Integrated into the Windows Server 2003 family, the .NET Framework is the infrastructure for .NET. The .NET Framework incorporates the common language runtime and a unified set of class libraries that include Windows Forms, ADO.NET, ASP.NET, and other capabilities. The .NET Framework provides a fully managed, protected, and feature-rich application execution environment, simplified development and deployment, and seamless integration with a wide variety of programming languages. | ● | ● | ● | ● |
| .NET Application Services | **Internet Information Services 6.0**<br><br>Internet Information Services (IIS) 6.0 is a full-featured Web server that provides the foundation for the Windows Server 2003 family and existing Web-based applications and XML Web services. IIS 6.0 offers dedicated application mode, which runs all application code in an isolated environment. IIS 6.0 also supports Web gardens, in which a set of equivalent processes on a computer each receive a share of the requests normally served by a single process, achieving better multiprocessor scalability. | ● | ● | ● | ● |
| .NET Application Services | **ASP.NET**(2)<br><br>ASP.NET is the engine for Web-based applications and XML Web services. It brings rapid application development to the server. Part of the class library in the .NET Framework, ASP.NET pages use a compiled, event-driven programming model that improves performance and enables the separation of application logic and user interface. | ● | ● | ● | ● |

| | | | | | |
|---|---|---|---|---|---|
| .NET Application Services | **Enterprise UDDI Services**<br><br>Universal Description, Discovery, and Integration (UDDI) is an industry specification for publishing and locating information about Web services. Some products in the Windows Server 2003 family include UDDI Services, a Web service that provides UDDI capabilities for use within an enterprise or across organizations. UDDI Services is not included with Windows Server 2003, Web Edition. In addition, Windows Server 2003, Standard Edition, supports only stand-alone installations of UDDI Services. Distributed installation support is available with Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition. In a stand-alone installation of UDDI Services, both the UDDI Web server component and the UDDI database component are installed on a single server. In a distributed installation, UDDI components are distributed across multiple servers. | ● | ● | ● | ○ |
| Multimedia Services | **Windows Media® Services**<br><br>Windows Media Services provide streaming audio and video over corporate intranets and the Internet. In Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, Windows Media Services delivers advanced streaming functionality such as multicasting, wireless network support, Internet authentication, server plug-ins, and cache/proxy APIs. | ◒ | ● | ● | ○ |

([1]) Applies to 64-bit versions only.

([2]) Not supported in 64-bit versions of Windows Server 2003.

([3]) Might be limited by lack of support by OEM hardware.

([4]) Both the 32-bit version of Windows Server 2003, Datacenter Edition and the 64-bit version of Windows Server 2003, Enterprise Edition support up to 64 GB of RAM.

([5]) The 64-bit version of Windows Server 2003, Datacenter Edition supports up to 512 GB of RAM.

# Printer Connectivity Technical Overview

Updated: March 28, 2003

Applies To: Windows Server 2003 with SP1

Network printing connectivity has a substantial impact on the performance and scalability of Microsoft® Windows® print servers in two main areas. The first of these are the options available for connecting client workstations to a print server. The second area of impact is the printer port type used to de-spool print jobs from the server to the print device. This paper provides a technical examination of network and local printer ports types available to Windows Server 2003, and the value of Standard Port Monitor to a high performance print serving environment.

**In This White Paper**


Client-Server Connection

Printer Port Selection

Network Printer Ports

Local Printer Ports

Internet Printing Protocol (IPP)

Summary of Printer Connectivity Technical Overview

# Client-Server Connection

Updated: March 28, 2003

Applies To: Windows Server 2003 with SP1

Windows clients typically connect to a print server through a remote procedure call (RPC) or SMB service. RPC connections are the preferred connection type, as they support the enhanced Point-and-Print features. SMB connections are associated primarily with Windows 98 and earlier systems. It is also possible to connect from Microsoft Windows NT®–based clients by creating a local printer port that is redirected to the printer share on a Windows print server. HTTP offers another possible connection type from a client to a server. Windows 2000 and Windows Server 2003 implement version 1.0 of the Internet Printing Protocol (IPP) standard. This protocol allows a client to view printers and their status as well as connect to shared printers and submit jobs using the HTTP print provider.

# Printer Port Selection

Updated: March 28, 2003

Applies To: Windows Server 2003 with SP1

Windows Server 2003 provides built-in support for the following network port types:

- Standard TCP/IP port monitor
- LPR print monitor
- AppleTalk print monitor
- HTTP print provider

| **Note** |
| --- |
| The Hewlett-Packard printer port is not available in Windows Server 2003. |

Windows Server 2003 also supports the following local printer ports:

- USB 1.0 and 2.0
- Parallel (IEEE 1284)
- IrDA
- IEEE1394
- Serial port
- SCSI print

Internet printing support was added in Windows 2000 Server as well. Although there are client-side and server-side pieces involved, we'll discuss this topic based on the full scenario rather local vs. network.

Additionally, some printer manufactures choose to provide their own local or network printer port type.

# Network Printer Ports

Updated: March 28, 2003

Applies To: Windows Server 2003 with SP1

**Standard TCP/IP Port Monitor**

For network-connected print devices, the standard TCP/IP port monitor is the best choice. Standard port monitor is the successor to line printer remote (LPR), that has been widely adopted as the de facto standard in network printing for the past several years. Standard port monitor is faster, more scalable, and bidirectional. In contrast, LPR is limited in all of these areas. Although Windows NT 4 and later provided registry modifications to help extend the capabilities of LPR printing, these changes do not compare with the benefits of using standard port monitor.

Standard port monitor uses TCP/IP as the transport protocol, which makes it the preferred port monitor in Windows Server 2003. Standard port monitor also uses the Simple Network Management Protocol (SNMP) to configure and monitor the printer status.

Standard port monitor sends documents to a printer using either the RAW or LPR printing protocols. Together, these protocols support most current TCP/IP printers.

| **Note** |
| --- |
| Do not confuse these print protocols with the transport protocols such as TCP/IP or Data Link Control (DLC). |

The RAW protocol is the default for most print devices. To send a RAW-formatted job, the print server opens a TCP stream to the printer's network interface. For many devices this will be port 9100. When the TCP/IP port is created, Windows uses SNMP to query the device for its object identifier according to RFC 1759 (Printer MIB). If the device returns a value the system file, tcpmon.ini, is parsed for a match. If the printer manufacturer has provided special configuration information for the particular device, it is created with the configuration settings in place. For example, some external print server interfaces support multiple printers (for example, the Hewlett Packard JetDirect EX with three parallel port connections). The manufacture can use different ports to designate which printer a job should be submitted to (for example, 9102 for port 1, 9103 for port 2; and so on). This ability is helpful for print server interfaces that require the use of specific port names, such as PASS on some IBM network printers.

Standard port monitor can be configured to comply more closely with the LPR standard (RFC 1179). Specifically, you can use the **Configure Port** button in the **Server Properties, Ports** tab to enable byte counting. This setting can be helpful if a line printer daemon (LPD) requires an accurate byte count to be sent, which is compliant with the RFC.

By default, standard port monitor deviates from the LPR standard in two ways:

- First, standard port monitor does not conform to the RFC 1179 requirement that the source TCP port lie between port 721 and port 731. Standard port monitor uses ports from the general, unreserved pool of ports (ports 1024 and above).

- Second, the LPR standard states that print jobs must include information about the size of the job the port monitor sends. Sending a print job with job size information requires that the port monitor spool the job twice, once to determine size, and once to send the job to the spooler. Spooling the job only once improves printing performance, so standard port monitor sends the job to the spooler without determining the actual job size, and claims the job is a default size, regardless of the job's actual size.

The table below compares the standard TCP/IP port monitor to the LPR port monitor.

| **LPR (RFC 1179 Compliant)** | **Standard TCP/IP Port Monitor** |
| --- | --- |
| Limited to 721–731 source ports | Source ports from ephemeral ports (9100 default) |
| Four-minute timeout per port | No source port timeout |
| Unidirectional—single error status reporting | Uses SNMP for rich status and error reporting |
| Control file requires double spooling for accurate byte count | Single-file spooling enabled by default. (configurable through Port Properties) |

| | |
|---|---|
| Prints to destination port 515 | Prints to destination port 515 in LPR mode |

**LPR**
Although Windows Server 2003 uses the standard TCP/IP Port monitor by default for a network printer, it still supports LPR/LPD printing for interoperability with legacy systems and RFC-compliant LPR/LPD clients and devices.

An LPR utility lets an application on one computer print to a spooler on a remote computer. The receiving component is called a line printer daemon (LPD). The LPR/LPD combination was developed for UNIX computers but is widely used for many operating systems. Both utilities are included in Print Services for UNIX.

**Default Data Types**
Print Services for UNIX sets the print data type when it sends the document to the spooler. This type is derived from the control command included in a print job from an LPR client. It might be necessary to change the default data type at the client to avoid processing PCL or PostScript print jobs as TEXT format when they are actually in RAW format.

If the control command is *f* or *p*, the data type is TEXT, and the spooler edits the document file to print correctly. If the command is *l*, the data type is RAW, and no editing is done. If the command is *o*, the document is already formatted in PostScript code and is assigned the RAW data type.

Some UNIX systems usually send the f command by default, resulting in the following symptoms:

- Output includes PCL or PostScript code
- Extended characters are printed incorrectly
- Printer's default font is used
- Extra page is printed at the end

You can make a registry entry to force Windows 2000 to process the data as RAW format. For more information, see the following Knowledge Base articles:

- Q150930, “LPD Server Adds and Print Control Codes”
- Q168457, “Configuring Individual Printers to Passthrough LPR Print Jobs”
- Q243729, “Preventing Print Spooler CR or LF to CR/LF Character Translation”

LPR must include an accurate byte count in the control file, but it cannot get it from the local print provider. After LPRMON receives a document from the local print provider, it spools it a second time as a temporary file in the System32 subfolder, finds the size of that file, and sends the size to the LPD print server. The standard TCP/IP port monitor does not adhere to this RFC requirement and sends a very large byte count to the LPD to begin printing. After the job is complete, it simply closes the connection. This step reduces the spool time and disk I/O by eliminating the temporary spool file creation.

The LPR protocol does not return a detailed error status report. If there is a problem, the message is always ERROR and Windows reports ”Printer Error.” As discussed above, through the use of SNMP and the standard printer MIB (RFC 1759), the standard TCP/IP port provides much more detailed error reporting.

Almost all current printer models are compatible with the standard TCP/IP port monitor. In most cases, Print Services for UNIX (LPR/LPD) is required only for interoperability with other computer systems that require strict compliance with the LPR/LPD RFC (1179). Some LPD servers require only the accurate byte count as specified in the RFC. For these systems, simply configuring the standard TCP/IP port with the **LPR Byte Counting Enabled** option eliminates any errors. Microsoft recommends using the standard TCP/IP port whenever possible for performance and status monitoring.

**AppleTalk**
**To install Apple Talk protocol**

1. From Control Panel, select Add or Remove Programs and then select Add or Remove Windows Components
2. Select Other Network File and Print Services, and then click Details

3. Then select Optional Windows Component Print Server for Macintosh. Print Server for Macintosh

This option allows the server to receive jobs from Macintosh clients and print to devices that use the AppleTalk protocol.

When creating an AppleTalk printer port, you are prompted to capture the port. If you capture the AppleTalk print device, it accepts print jobs exclusively from this print server. This option can be helpful in providing tighter device management and control. For a higher level of interoperability or in a mixed environment with multiple servers or clients printing to the device, the port should not be captured, which enables a job-based connection.

In addition, when Printer Services for Macintosh is installed, Macintosh clients can connect to printer shares on a Windows print server. These printers can be locally connected to the server or connected through any support network port, such as standard port monitor, LPR, and so forth.

# Local Printer Ports

Updated: March 28, 2003

Applies To: Windows Server 2003 with SP1

**USB Printing**
Windows Server 2003 supports printing to Universal Serial Bus (USB) printers. USB consists of an external bus architecture for connecting USB-capable peripheral devices to a host computer, as well as a communication protocol that supports serial data transfers between a host system and USB-capable peripherals.

**Parallel Port**
Support for parallel port printing was enhanced in the release of Windows XP and Windows Server 2003. Traditionally, parallel port devices did not benefit from the introduction of Plug and Play several years ago. Windows XP and Windows Server 2003 perform periodic polling of the parallel port to provide an experience similar but not identical to Plug and Play. USB and IEEE 1394 are fully Plug-and-Play aware and should be used whenever printing performance and system performance are important. Parallel printing is slower and more CPU-intensive than these other local printing options.

**IrDA**
Infrared Data Association (IrDA) is a system of exchanging information between computers using infrared transmissions that do not require a cable connection. IrDA can occur between any two devices that support IrDA (such as computers and printers). Windows Server 2003 supports printing using IrDA.

IrDA is a point-to-point protocol based on TCP/IP and WinSock application programming interfaces (APIs). Typically, IrDA printers support Plug and Play, automatically installing the driver if it is available on the workstation. Additionally, the Add Printer Wizard may be used by selecting the IrDA port under **Available Ports**.

Connectivity Capacities:

| Technology | Variant | Speed |
|---|---|---|
| Parallel (1284) | (Centronics) | 250 Kbps |
| | ECP/EPP | 3–5 Mbps |
| IrDA | | 4 Mbps |
| FireWire (1394) | (up to 63 devices) | 400 Mbps |
| USB | 1.0 | 12 Mbps |
| | 2.0 | 480 Mbps |
| Bluetooth | | 2–3 Mbps |

**IEEE 1394**
Windows Server 2003 supports the IEEE 1394 bus, which is designed for high-bandwidth devices, such as digital camcorders, digital cameras, digital VCRs, and storage devices. IEEE 1394 is a serial protocol supporting speeds ranging from 100 to 400 megabits per second (Mbps), depending on the implementation. It provides a high-speed Plug and Play-capable bus that eliminates the need for peripheral devices to have their own power supply, and provides support for asynchronous data transfer.

You can connect up to 63 devices to one IEEE 1394 bus and interconnect up to 1,023 buses to form a very large network with more than 64,000 devices. Each device can have up to 256 terabytes of memory addressable over the bus. A built-in mechanism ensures equal access to the bus for all devices.

Because of the very high data transfer rates that IEEE 1394 can handle, it is ideal for scanning or printing large, high-resolution data.

# Internet Printing Protocol (IPP)

Updated: March 28, 2003

Applies To: Windows Server 2003 with SP1

Windows 2000 Server introduced support for the Internet Printing Protocol 1.0 specification. Windows Server 2003 continues this support for the 1.0 standard. There are two main components to IPP:

- Web-based printer management
- Internet printing

**Web-based Printer Management**

Web-based printer management is very simple to use. The only prerequisite is that IIS be installed and running on the print server. It is on by default for Windows 2000 Server versions and can be installed as an optional component on Windows 2000 Professional. Web-based printer management can be configured on Windows Server 2003 by using the IIS Lockdown Wizard and selecting the **Enable Active Server Pages** and **Enable Internet Printing** check boxes. When a printer is created and shared, it becomes available for management through the Internet browser. To access the Web-based feature, use the following URL, substituting your print server's computer name for *printserver*:

http://*printserver*/printers

Users must have adequate permissions in order to pause printers, purge print jobs, and so on. Most Windows clients in use today have built-in support for IPP or can download IPP support from the Microsoft Web site (http://www.microsoft.com). When a client has IPP support installed, they can view and connect to printers from within their Web browser. Microsoft strongly recommends the use of this technology, particularly for corporations that want to deploy a Web-based solution for finding and connecting to printers.

| **Note** |
| --- |
| Windows XP provides support for Web-based management of locally hosted printers by default. In the future, this option might be disabled by default—as is true for Windows Server 2003 — and enabled by a process similar to the IIS Lockdown Wizard used on the Windows Server 2003 family of products. |

**Internet Printing**

When a user connects to a printer through the Web-based printer management, one of two things can happen: Either Windows creates an IPP printer connection using an HTTP port, or Windows creates an RPC "true-connect" printer connection.

**HTTP Printer Connection Through Web Browser**

If the client's Internet Explorer security settings for the print server are set to medium or higher, Windows creates an IPP printer connection using an HTTP port. Because this technology requires the installation of a local printer queue (unlike a true Point-and-Print connection, which uses RPC), the client must have either **Administrator** or **Power User** status with the added **Load and unload driver** privilege on the local machine. (On Windows 2000 Server, this connection also works if the client has only Power User privilege.)

This solution is recommended only for Internet sites–specifically, untrusted sites viewed by a browser in which the security settings must be set to medium or higher. This type of printing works very well to print data in one location on a page at another location that does not share security or network infrastructure. However, HTTP printer connections are more limited than RPC true-connections (described below) in a Windows printing environment, as they do not support enhanced metafile (EMF) features.

**RPC Printer Connection Through Web Browser – Preferred Method**

If the client's Internet Explorer settings are set to medium-low or lower, when the client chooses the **connect** option in a Web browser, Windows automatically creates an RPC true-connect printer connection. True-connect, or UNC, connections have many benefits over true IPP printer connections and should be used for intranet printing (local to the company) where the security setting in Internet Explorer can be set to medium-low safely. True-connect connections are supported through the Windows remote spooler, which is where the vast majority of future Microsoft innovation will be—in the RPC printer connection rather than the IPP printer connection state.

For additional information about RPC connections, see the "Point and Print Technical Overview" at

http://www.microsoft.com/printserver.

# Summary of Printer Connectivity Technical Overview

Updated: March 28, 2003

Applies To: Windows Server 2003 with SP1

Use Standard TCP/IP Printer Ports and connect using Windows 2000 or later clients by using RPC instead of TCP/IP. (Named Pipes implementation)

Using the portconv.vbs script or Print Migrator 3.0 or later lets you easily convert existing LPR ports to standard TCP/IP ports.

Microsoft is evaluating the standard TCP/IP port support for SNMP through the standard printer MIB (management information base). In future releases, it might be possible to extend the implementation to include even more status and error reporting. Other new technologies such as Universal Plug & Play (UPnP) along with printer driver enhancements will enable a new era of device installation and management for future enterprise printing environments.

**Related Links**
See the following resources for further information:

- For more information about deployment and management of Windows print servers, see *Windows Server 2003 Deployment Resource Kit* (not yet released when this paper was published).

- For more information about the Windows 2000 and Windows Server 2003 architecture, see the "Display and Print Devices" sections in the Windows Driver Development Kit at http://go.microsoft.com/fwlink/?LinkID=123179.

- For more information about Windows Print Servers in general, visit the Print Server Web site at http://www.microsoft.com/printserver.

- For the latest information about Windows Server 2003, see the Windows Server 2003 Web site at http://www.microsoft.com/windowsserver2003.

# Using Clustering for a Highly Available Web Site: An Example

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: September 1, 2003

**In This White Paper**

Introduction (Using Clustering for a Highly Available Web Site: An Example)

Topics Discussed in This White Paper (Using Clustering for a Highly Available Web Site: An Example)

Microsoft.com Objectives

Reasons for Using Network Load Balancing

Microsoft.com Cluster Overview Diagram

Set of Network Load Balancing Clusters Running IIS 6.0

Network Load Balancing Host Diagram

Configuration of the Network Load Balancing Clusters at Microsoft.com

Cluster Monitoring

Network Load Balancing Best Practices (Using Clustering for a Highly Available Web Site: An Example)

Summary (Using Clustering for a Highly Available Web Site: An Example)

Related Links (Using Clustering for a Highly Available Web Site: An Example)

# Introduction (Using Clustering for a Highly Available Web Site: An Example)

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

High availability, high reliability and high performance are important objectives for Web sites. This white paper describes in detail how Microsoft uses Network Load Balancing to achieve these objectives for its own Microsoft.com Web site. You can use information in this document to help you achieve similar goals for your own Web site.

This white paper discusses the Network Load Balancing (NLB) clusters that support the Microsoft.com Web site. It provides detailed information about hardware, software, and Network Load Balancing configurations, and also reasons why the configurations were chosen. This white paper includes a list of best practices for using Network Load Balancing, as well as information about several operational procedures that Microsoft uses to help make its Web site highly reliable and highly available. Web site administrators can use Microsoft.coms implementation as an example for their own Network Load Balancing implementation.

| **Note** |
| --- |
| This is the first of two white papers that describe the use of clustering at Microsoft. The second white paper, which is called Using Clustering with Exchange 2003: An Example, describes how the internal IT department at Microsoft has implemented server clusters to manage their Exchange servers. This white paper is available at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=18652). |

**Focus: Network Load Balancing Clusters**

The primary focus of this white paper is on the Network Load Balancing clusters that the Microsoft.com Web site uses. Other aspects of Microsoft.com, such as the domain structure and the Active Directory configuration, are only described to the extent that they affect the clustering implementations.

**Use This White Paper as an Example Only**

The information in this white paper about the Network Load Balancing clusters that are used by the Microsoft.com Web site is presented as an example. The actual steps and configuration for computers and devices in your own network will be different. This document shows only the configuration that specifically relates to the objectives described later in this document. This document does not cover other procedures and devices that are required in a production environment.

# Topics Discussed in This White Paper (Using Clustering for a Highly Available Web Site: An Example)

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

- For information about Microsoft.coms technical and business goals, see Microsoft.com Objectives.
- For diagrams of Microsoft.com's Network Load Balancing clusters and hosts, see the following:
    - Microsoft.com Cluster Overview Diagram
    - Set of Network Load Balancing Clusters Running IIS 6.0
    - Network Load Balancing Host Diagram
    - For a brief discussion of the specific clustering solution that Microsoft.com has implemented, as well as the reasons this particular solution was chosen, see Reasons for Using Network Load Balancing.
- For detailed configuration information about the load-balanced servers running IIS, see Configuration of the Network Load Balancing Clusters at Microsoft.com.
- For recommendations and product tips on Network Load Balancing, including practices used by Microsoft.com, see Network Load Balancing Best Practices.

# Microsoft.com Objectives

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This document describes how Microsoft uses Network Load Balancing and related operational procedures to achieve the following technical and business objectives for the Microsoft.com Web site:

- **High availability**. Provide Web site content that is always available to users.
- **Efficiency**. Deliver Web content to users quickly, regardless of the number of concurrent sessions.
- **Scalability**. Allow system administrators to add additional servers as necessary in a way that is transparent to users.
- **Ease of maintenance**. Allow system administrators to maintain the Web site easily through a process that is transparent to users.

These objectives apply to the set of Network Load Balancing clusters described in this white paper. These objectives are similar to those of most Web sites. Therefore, the clustering implementations presented here should be valuable to businesses with similar goals.

# Reasons for Using Network Load Balancing

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft uses Network Load Balancing to help achieve its technical and business objectives for the Microsoft.com Web site. This section discusses specific features that Network Load Balancing offers that help Microsoft.com achieve these objectives.

This white paper discusses the set of Network Load Balancing clusters at Microsoft.com that are running IIS 6.0. This set of clusters consists of six Network Load Balancing clusters. Each of these clusters has 10 servers running IIS.

Network Load Balancing provides improved performance, the ability to add additional servers as necessary, and the ability to perform maintenance on any of the servers without any customer impact or hardware configuration changes. By using the combination of Network Load Balancing and multiple, redundant clusters, Microsoft.com can add and remove clusters with minimal customer impact.

The following is a detailed list of Network Load Balancing features:

- **Scalability**
  - Network Load Balancing can load balance requests for individual TCP/IP requests across the cluster.
  - Network Load Balancing can load balance multiple server requests, from the same client or from several clients across multiple hosts in the cluster.
- **High availability**
  - Network Load Balancing can automatically detect and recover from a failed or offline computer.
  - Network Load Balancing can automatically rebalance the network load when hosts are added or removed.
- **Manageability**
  - Network Load Balancing can specify the load-balancing behavior for a single IP port or group of ports using port management rules.
- **Ease of use**
  - Network Load Balancing is installed as a standard Windows networking driver component.
  - Network Load Balancing requires no hardware changes to enable and run.
  - Network Load Balancing allows clients to access the cluster with a single logical Internet name and virtual IP address (also known as the cluster IP address) while retaining individual names for each computer. This feature can be used for administration purposes.
  - Network Load Balancing allows server applications to run in a Network Load Balancing cluster without requiring modification.
  - Network Load Balancing can be configured to automatically add a failed host to the cluster after it is brought back online. The added host will then be able to start handling new server requests from clients.
  - Network Load Balancing allows system administrators to take computers offline for preventive maintenance without disturbing cluster operations on the other hosts.

**Reason for Using Multiple Clusters**

In addition to Network Load Balancing, Microsoft.com uses multiple clusters to make the Web site easier to maintain. For

example, if an entire virtual IP (VIP) begins to experience problems, administrators can stop routing traffic to that entire cluster altogether (a process external to Network Load Balancing). Administrators can then perform maintenance or corrective procedures on the software and equipment associated with the failed VIP while the other VIPs continue to respond to Internet client requests.

For more information about the design and architecture of the Network Load Balancing clusters and individual hosts, see Microsoft.com Cluster Overview Diagram, Set of Network Load Balancing Clusters Running IIS 6.0, and Network Load Balancing Host Diagram later in this document. For details on the hardware and software used in this cluster, as well as the configuration of the hardware and software used to take advantage of these features, see Configuration of the Network Load Balancing Clusters at Microsoft.com later in this document.

# Microsoft.com Cluster Overview Diagram

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Diagram 1 is a general overview of how three sets of clusters work together to support the Microsoft.com Web site. This white paper describes the Network Load Balancing clusters of Web servers. This set of clusters consists of six clusters. Each of these clusters is made up of 10 servers running IIS 6.0.




**Diagram 1**

# Set of Network Load Balancing Clusters Running IIS 6.0

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Diagram 2 depicts the Network Load Balancing clusters of Web servers that support the Microsoft.com Web site. This set of clusters consists of six Network Load Balancing clusters. Each cluster has 10 servers running IIS 6.0.




**Diagram 2**

Each of the six clusters has a unique virtual IP address (VIP). However, all six of these clusters are addressable by the URL www.microsoft.com. Internet traffic is initially routed to one of the six VIPs by round robin. Round robin is a simple mechanism used by DNS servers to share and distribute loads for network resources. Here, round robin is used to rotate the order of VIPs that are used in response to a query from the Internet to www.microsoft.com.

Each host within each of the six clusters has two network adapters: the cluster adapter, which connects to the Internet and is bound to one of the VIPs, and the management adapter, which connects to the Microsoft corporate network by a perimeter network. All 10 of the hosts in each individual cluster have the same VIP that is bound to the hosts cluster adapter (shown by the red line). In contrast, the blue line represents the set of 10 unique IP addresses (one for each host in the cluster) that are bound to its management adapter. For more information about the two network adapters, see Diagram 3.

For more information about the design considerations and configuration details of these Network Load Balancing Clusters, see Reasons for Using Network Load Balancing, earlier in this document, and Configuration of the Network Load Balancing Clusters at Microsoft.com later in this document.

# Network Load Balancing Host Diagram

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Diagram 3 depicts the network-adapter configuration for each Network Load Balancing host. These hosts make up the Network Load Balancing clusters of Web servers for the Microsoft.com Web site.




**Diagram 3**

This set of clusters consists of six individual Network Load Balancing clusters. Each cluster has 10 hosts. Each host in the cluster has two network adapters. The cluster adapter, or front-end adapter (labeled FE by the administrators at Microsoft.com), is attached to the Internet. The management adapter, or back-end adapter (labeled BE by the administrators at Microsoft.com), is attached to the Microsoft corporate network.

The FE adapter is bound to both the virtual IP address (VIP) and a unique dedicated IP address (DIP). All Internet traffic comes through the FE adapter by using the VIP. The DIP on this adapter is typically used only by ClusterSentinel, which is the cluster management and monitoring tool used by Microsoft.com. Other traffic that is intended specifically for a particular host is addressed to the BE adapter.

The BE adapter on each host has a unique static IP address and is used for cluster and host management traffic, which typically originates from the Microsoft corporate network.

For more information about the design consideration and configuration details of Network Load Balancing Cluster 1, see Reasons for Using Network Load Balancing, earlier in this document, and Configuration of the Network Load Balancing Clusters at Microsoft.com later in this document.

# Configuration of the Network Load Balancing Clusters at Microsoft.com

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section describes the details of the hardware, software, and Network Load Balancing configurations of the clusters of Web servers that support the Microsoft.com Web site. Because the purpose of this document is to focus on the details of the clustering implementations, only limited information is provided regarding general network hardware configuration and other configuration settings that are not specifically related to Network Load Balancing. However, each Network Load Balancing setting is discussed in depth.

- For information about hardware configuration, see Hardware Configuration for Network Load Balancing Hosts.
- For information about the software configuration, see Software Configuration for Network Load Balancing Hosts.
- For information about the Network Load Balancing configuration, see Network Load Balancing Configuration.

**Hardware Configuration for Network Load Balancing Hosts**

Each host in the Network Load Balancing clusters at Microsoft.com has the same hardware. The servers are Compaq ProLiant ML570s (7U) with the following standard hardware:

| | |
|---|---|
| **Processors** | 4 x Pentium III XEON 700 Mhz; 1-MB cache processors |
| **Memory** | 2 GB, 100 Mhz SDRAM |
| **Controller** | Smart array controller 5302/64: PCI; 64-bit/66 MHz; 2-channel; U160; 64-MB cache |
| **Hard disk drives** | 10 x 18.2 GB hot-plug; 15K RPM drives<br>2 x 36.4 GB hot-plug; 15K RPM drives |
| **Network adapters** | 2 x NC3134 Fast Ethernet network adapters; 32/64-bit PCI; 33/66 MHz; dual-base 10/100 |
| **Power supply** | Hot-plug redundant power supply |

**Explanation**

This hardware provides the necessary processor capacity for unexpected spikes in Web-site traffic at Microsoft.com. At the time of this writing, these are currently some of the fastest four-processor computers available. With 2 gigabytes (GB) of memory, the computers can handle several IIS 6.0 processes at once, in addition to providing plentiful caching capacity.

**Storage Configuration for Each Host**

Each host has the following storage configuration:

- 120-GB content capacity
- 34-GB logging capacity
- Virtual disk 1: RAID-1 mirrored system volume that contains the operating system and applications
- Virtual disk 2: RAID-1 mirrored logging volume that contains IIS log files
- Virtual disk 3: RAID-5 striped content volume that contains all Web content

**Storage configuration for each host**




**Diagram 4**

Diagram 4 shows the storage configuration for each host. Each host has 12 physical disks, which are configured into three virtual disks. Virtual disks 1 and 2 are RAID-1 configurations and the third is a RAID-5 configuration. Virtual disk 3 (RAID-5) is distributed across a two-channel RAID controller.

**Explanation**

Microsoft.com implemented this storage configuration for each host to achieve high capacity and redundancy. Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, support RAID-0, RAID-1, and RAID-5 configurations. The purpose of RAID-1 and RAID-5 configurations is to guard against the loss of data in the event of a catastrophic hard-disk failure. Microsoft.com used RAID-1 and RAID-5 configurations to provide the required fault tolerance at a minimum cost.

The virtual disks and hard disks were split across two channels because of the limitations of hard disk space. Although each RAID channel can only accommodate six hard disks, each server has space for 12. Therefore, these are split between two channels.

This design is a compromise between minimizing cost and optimizing performance. With only two channels and a single controller to use, Microsoft.com chose to put the Web content on a single channel and therefore, had to put the operating system and the IIS logs together on the other channel. The ideal high-availability solution, which includes storing the mirrored volumes on separate channels, requires additional hardware costs.

**Software Configuration for Network Load Balancing Hosts**

Each host in the Network Load Balancing clusters at Microsoft.com has the following software:

- Windows Server 2003, Enterprise Edition
- IIS 6.0
- Microsoft Operations Manager (MOM)

**Network Adapters**

Each server has two network adapters installed. The cluster adapter is labeled FE (by Microsoft.com administrators) for front-end. The management adapter is labeled BE for back-end.

The network adapters have the following:

| | **"FE" network adapter** | **"BE" network adapter** |
|---|---|---|
| **Protocols** | Internet protocol (TCP/IP)<br><br>Network monitor driver | Internet protocol (TCP/IP)<br><br>Network monitor driver |
| **Clients** | Network Load Balancing | Client for Microsoft networks<br><br>File and printer sharing for Microsoft networks |
| **IP address** | Static virtual IP address (VIP)<br><br>Static dedicated IP address (DIP) | Static IP address |

**Explanation**

The adapters are labeled for easy identification. The FE adapter is the cluster adapter, which is used for Internet traffic, while the

BE adapter is the management adapter, which is used for content replication and administration. As a best practice, you should use two network adapters in all Network Load Balancing implementations. In this case, the two network adapters provide additional security by physically separating the Internet traffic from the perimeter network (also known as the screened subnet) traffic. Also, they address the restriction that a Network Load Balancing cluster that is operating in unicast mode cannot distinguish among single adapters on individual hosts. Therefore, to make communication among cluster hosts possible, each cluster host must have at least two network adapters.

**TCP/IP Configuration**

- Each network adapter is given a static IP address and a static DNS server address. The static IP address on the FE adapter is known as the dedicated IP address (DIP).
- In addition to the DIP, the FE adapter also has the virtual IP address (also known as the cluster IP address, or VIP), which is used by all Internet clients connecting to Microsoft.com.
- NetBIOS is disabled on the FE adapter through the advanced TCP/IP **Properties** dialog box.

**Explanation**
Each adapter that is bound to Network Load Balancing (in this case the FE adapter on each host) has a dedicated IP address (DIP) in addition to the virtual IP address (see Diagram 5). The dedicated IP address must be a static IP address that is unique for each host. The DIP cannot be a DHCP address. The DIP is only used to allow remote connections to a specific host through the Network Load Balancing Adapter. Traffic addressed to the DIP is never load balanced.

Be aware that you should always enter the dedicated IP address first in TCP/IP properties, before the virtual IP address.




**Diagram 5**

Network Load Balancing references the dedicated IP address (DIP) only when a single network adapter is used to handle both client-to-cluster traffic and other network traffic that must go specifically to the dedicated IP address; for example, Telnet access to a specific host within the cluster. At Microsoft.com, the DIP is primarily used by ClusterSentinel. All other administrative communication, such as content replication, occurs through the BE adapter.

Communication over the direct hosting port, as well as other ports that are considered unnecessary, is restricted by router access control lists (ACLs). The NetBIOS is disabled on the FE adapter in order to provide a secondary layer of protection to protect against malicious activity behind the ACLs. However, even with the NetBIOS disabled, direct hosting of Server Message Block (SMB) traffic is still active. For more information about direct hosting, see article 204279, Direct Hosting of SMB Over TCP/IP, in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=4441).

**Network Load Balancing Configuration**
This section describes the parameters that are configured for the Network Load Balancing clusters as a whole, the parameters that are specific to each host in the Network Load Balancing cluster, and the port rules that control how the cluster functions.

**Cluster Parameters**
These parameters appear on the **Cluster Parameters** tab that is available when you configure or view Network Load Balancing properties. These parameters must be set identically on each host of the cluster. The easiest way to achieve this is to use Network Load Balancing Manager when implementing and configuring the hosts.

| Configuration setting | Explanation |
|---|---|

| | |
|---|---|
| Network Load Balancing is enabled on the network adapter labeled FE. | Network Load balancing is only enabled on the FE adapter because the BE adapter is used exclusively for administrative network traffic and communication among the hosts. The BE adapters also handle content replication across the hosts. The BE adapter does not handle any cluster traffic from clients. |
| **IP address** is set to the IP address of the cluster (the VIP). In this document, these addresses are 0.0.0.x.* See Set of Network Load Balancing Clusters Running IIS 6.0 (Diagram 2) and the Network Load Balancing Host Diagram (Diagram 3), earlier in this document, for more information.<br><br>**Subnet mask** is set to the subnet mask of the cluster's IP address.<br><br>* These are not the actual Microsoft.com VIPs. | Note that each cluster has only one virtual IP address. Microsoft.com has not used the Virtual Cluster feature of Windows Server 2003, Enterprise Edition, which allows you to enable multiple virtual IP addresses for each cluster. |
| **Unicast** is selected. | Unicast mode is typically the preferred cluster operation mode. While multicast mode can sometimes be used as an alternative and does offer certain advantages, it is not supported by all routers and often presents additional complexity. Unicast mode with multiple adapters is a model that is suitable for a cluster such as the one described in this document, in which ordinary network communication among cluster hosts is necessary or desirable and a second management adapter is available for that communication. |
| **Remote control** is enabled. | Remote control is enabled to allow the administrators to use ClusterSentinel for controlling the cluster. |

**Note**

Although used at Microsoft.com, the remote control option is generally discouraged because it presents significant security exposures including the possibility of data tampering, denial of service and information disclosure. We strongly recommend that sites use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI) instead of remote control. If you choose to enable remote control, it is vital that you restrict access by specifying a strong remote control password. In addition to the password, it is also imperative that you use a firewall to protect the Network Load Balancing UDP control ports (the ports receiving remote control commands) to shield them from outside intrusion (see Diagram 6). By default, these are ports 1717 and 2504 at the virtual IP address. Microsoft.com has implemented these safeguards.




**Diagram 6:**

**Host Parameters**

These parameters appear on the **Host Parameters** tab that is available when you configure Network Load Balancing. Some of these parameters are unique for each host, but in most cases are identical across all of the hosts in the cluster.

| Configuration setting | Explanation |
|---|---|

| | |
|---|---|
| **Priority (Unique Host ID)** is set to the host ID that the particular server represents in the cluster. | The **Priority (Unique Host ID)** determines which host handles certain cluster traffic. The host with the lowest numerical priority among the current members of the cluster handles all of the cluster's network traffic that is not covered by a port rule. Each host must have a different **Priority** value. If Network Load Balancing Manager is used to configure the hosts, Network Load Balancing Manager will ensure that each host has a unique priority. |
| **Initial host state** is set to stopped and the **Remain suspended** check box is not selected. | **Initial host state** is set to **stopped** to allow the Microsoft.com administrators to perform a variety of tests on the host once it is restarted after a shutdown, but prior to allowing it to rejoin the cluster. This allows the administrators to ensure that the host is operating properly before it is included in the Network Load Balancing rotation. |
| **Dedicated IP address** and **Subnet mask** are set to the IP address and subnet mask of the FE network adapter. | Once again, it is important to distinguish between the dedicated IP address (the DIP), which is configured on this tab, and the virtual IP address (VIP), which is the cluster's IP address and is configured on the **Cluster Parameters** tab. If you use Network Load Balancing Manager to configure your cluster, the dedicated IP address parameter will already contain the correct address. The addresses configured on this tab (and the **Cluster Parameters** tab) are assigned to the FE network adapter. |

**Port Rules**

Port rules are set up to control how each port handles network traffic addressed to the cluster. The number and type of rules must be exactly the same for each host in the cluster. We recommend that you use Network Load Balancing Manager to configure port rules. If you are using Network Load Balancing Manager, when you add additional hosts, they will automatically inherit the cluster port rules from the initial host. If you are not using Network Load Balancing Manager and a host attempts to join the cluster with a different number of rules, or with different rules from the other hosts, it is not accepted as part of the cluster. The rest of the cluster then continues to handle the traffic as before. At the same time, a message is entered into the Windows event log.

Be aware that all network traffic that is addressed to the cluster and is not governed by port rules is handled by the host with the highest host priority among the current members of the cluster. This single host handles all of the cluster network traffic not governed by a port rule, with another host taking over the traffic only in the event that the highest priority host fails or goes offline. This default behavior ensures that Network Load Balancing does not affect cluster network traffic for ports that you do not specifically manage with the Network Load Balancing load-balancing mechanisms. It also provides high availability for handling your cluster network traffic.

**Port Rules Configured at Microsoft.com**

Microsoft.com has only two port rules configured:

- Port range: 80 to 80
    - Virtual IPs: all
    - Protocols: TCP
    - Filtering mode: multiple hosts
    - Affinity: none
    - Load weight: equal
- Port range: 443 to 443
    - Virtual IPs: all
    - Protocols: TCP
    - Filtering mode: multiple hosts
    - Affinity: none

- Load weight: equal

**Explanation**

The **port range** that is specified in each of the rules includes only ports that are handled by the cluster's virtual IP address (VIP). Therefore, they are the only ports handling traffic managed by Network Load Balancing. Be aware, however, that this does not prevent network traffic from being sent to and from each host's dedicated IP address (DIP) or other ports on the VIP. Microsoft.com has used router access control lists (ACLs) to manage this kind of traffic. You can also disable the remaining ports by creating port rules that include all remaining ports, and then using the **Disable** setting in the **Port Rules Configuration** dialog box. While this configuration provides additional protection for the host, do not use port rules as a substitute for a firewall.

Microsoft.com does not make use of virtual clusters in this document, thus the port rules are configured as global port rules by choosing **All** for the virtual IP address.

TCP is the only type of traffic that this cluster handles.

Setting the **filtering mode** to **multiple hosts** allows all of the hosts in the cluster to handle network traffic for the associated port rule. This filtering mode provides scaled performance and fault tolerance by distributing the network load among multiple hosts. You can specify whether the load is equally distributed among the hosts or that each host handle a specified load weight. Microsoft.com has specified **equal load weight** for each host.

**Affinity** is set to **None**, which specifies that multiple connections from the same client IP address can be handled by different cluster hosts (that is, there is no client affinity).

| **Note** |
| --- |
| Although not used by Microsoft.com, Class C or single affinity is typically recommended for Secure Sockets Layer (SSL) transmissions on port 443 to make the most efficient use of server resources. SSL connections are associated with a particular session ID, which (within a specified time constraint) can be reused when communicating with a particular host. If the client computer presents the session ID to a different host, the session will be rejected and a new session will be negotiated. (This process is transparent to the user). The negotiation requires additional host resources and increases the server overhead. Therefore, if you use Class C or single affinity, you can minimize how often the session must be renegotiated. |

This concludes the configuration section of this document. As described earlier, information provided here is meant to cover only those settings directly related to the Network Load Balancing implementation at Microsoft.com. Other software and hardware details that are necessary for a production Web site have not been described.

# Cluster Monitoring

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses the tools that the administrators at the Microsoft.com Web site use to monitor the condition and configuration of their clusters and cluster hosts. Microsoft Operations Manager (MOM) is installed on each server running IIS and is used to monitor event logs and run scripts remotely. In addition to MOM, Microsoft.com uses a variety of other tools for monitoring and reporting including ClusterSentinel, HTTP Monitoring Tool, Data Center Warehouse, and RMS. Microsoft developed these four monitoring tools for use within the company. However, other commercially available products, such as Microsoft Application Center, provide similar functionality.

**ClusterSentinel**
ClusterSentinel is a tool used to monitor and control the Network Load Balancing cluster hosts. Microsoft.com administrators use ClusterSentinel to:

- Determine and report the state of the Network Load Balancing cluster hosts, such as converged or suspended.
- Manually control cluster hosts by using Network Load Balancing commands, such as **Stop**, **Start**, **Drain** and **Drainstop**.
- Monitor and report on application health. If an application is determined to be having problems on a particular host, that host will automatically be removed from the cluster.

ClusterSentinel traffic is handled by the dedicated IP address on the FE network adapter on the cluster hosts. For more information about the network adapters configured in this cluster, see Network Load Balancing Host Diagram earlier in this document.

Network Load Balancing Manager, a new feature in the Windows Server 2003 family of products, can be used for the first two of the tasks performed by ClusterSentinel. Network Load Balancing Manager does not, however, have the ability to monitor application condition.

**HTTP Monitoring Tool**
HTTP Monitoring Tool (HTTPMon) is a Resource Kit utility that is included in Microsoft Windows NT 4.0 and in the *Microsoft Windows 2000 Resource Kit*. HTTPMon is a multithreaded tool that monitors HTTP activity on your servers and can notify you if there are changes in the amount of activity. You can also configure it to automatically remove a host from a Network Load Balancing cluster depending on the availability of specific applications, such as IIS, or the availability of a particular Web site.

HTTPMon runs through a series of remote application level tests and collects information from each server every 30 seconds, looking for failures in the application layer, which are more common than system failures, but much harder to detect. Application-layer problems can range from IIS being overloaded by requests, to an application not responding and requiring a server reboot. HTTPMon looks for errors consistent with RFC 1945. For example, a status code of 200 indicates successful performance, while a status code of 500 indicates a server error. The results of the tests are stored in a SQL database.

If HTTPMon determines that a server is having problems at the application layer (for example, with IIS), HTTPMon directs Network Load Balancing to remove the host from the cluster. One minute later, HTTPMon checks again to see if the server has recovered. If it has, the server is returned to service. If the server still does not respond, a technician is alerted to attend to the computer, determine the nature of the problem so it can be corrected, and bring it back online.

**Data Center Warehouse**
The Data Center Warehouse (DCW) is a tool that compiles all of the information from ClusterSentinel, HTTPMon and a variety of performance monitors, such as ASP requests per second, and presents the data on a single Web page. It displays a global view of each server and cluster. Administrators can then review the information collected in its entirety or they can limit their view to only certain error conditions, such as servers with failing applications or well-functioning servers that are not active in the cluster.

Microsoft.com Data Center Warehouse - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Search Favorites Media

Address http://dcw/

Current View Web Servers | Refresh In : 26 | Microsoft.com Data Center Warehouse - Last Data Update : 1:08:40 PM

| Site/Cluster | ASP Req/sec | ASP Req Exec | ASP Req Queued | ASP Req Rejected | Mem Avail (MB) | % Disk | % Proc | Web Conn Att/sec | Current Users | SMTP Queue | SMTP M/sec | System Uptime | IIS Uptime | CS | HTTPMon | NLB Status |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| NLB Cluster 1 | .00 | .00 | .00 | 0 | 360 | .20 | 27.60 | 159.40 | 2052 | 0 | .00 | 12.60 | 7.80 | | | |
| Host01 | 0 | 0 | 0 | 0 | 170 | N/A | 33 | 43 | 496 | 0 | 0 | 16 | 8 | Bad | 0 | Converged |
| Host02 | 0 | 0 | 0 | 0 | 160 | 0 | 30 | 165 | 537 | 0 | 0 | 13 | 8 | Good | 200 | Default |
| Host03 | 0 | 0 | 0 | 0 | 194 | 0 | 16 | 231 | 67 | 0 | 0 | 13 | 2 | Good | 200 | Converged |
| Host04 | 0 | 0 | 0 | 0 | 640 | 1 | 29 | 170 | 529 | 0 | 0 | 8 | 8 | Bad | 200 | Default |
| Host05 | 0 | 0 | 0 | 0 | 639 | 1 | 30 | 188 | 423 | 0 | 0 | 13 | 13 | Good | 200 | Converged |
| NLB Cluster 2 | 66.57 | .00 | .14 | 0 | 1566 | .00 | 5.85 | 34.14 | 11523 | -7 | -1.00 | 13.14 | 13.14 | | | |
| Host01 | 69 | 0 | 0 | 0 | 352 | 0 | 8 | 36 | 1815 | N/A | N/A | 12 | 12 | Good | 200 | Default |
| Host02 | 74 | 0 | 1 | 0 | 356 | 0 | 10 | 25 | 2097 | N/A | N/A | 12 | 12 | Good | 200 | Converged |
| Host03 | 83 | 0 | 0 | 0 | 3446 | 0 | 7 | 32 | 1923 | N/A | N/A | 12 | 12 | Good | 200 | Converged |
| Host04 | 82 | 0 | 0 | 0 | 132 | 0 | 3 | 49 | 1647 | N/A | N/A | 12 | 12 | Good | 200 | Default |
| Host05 | 0 | 0 | 0 | 0 | 275 | 0 | 1 | 0 | 0 | N/A | N/A | 12 | 12 | Bad | 0 | Suspended |
| Host06 | 74 | 0 | 0 | 0 | 3222 | 0 | 6 | 49 | 2044 | N/A | N/A | 12 | 12 | Good | 200 | Converged |
| Host07 | 84 | 0 | 0 | 0 | 3183 | 0 | 6 | 48 | 1997 | N/A | N/A | 20 | 20 | Good | 200 | Converged |

**Figure 1:**

Figure 1 shows the type of information available in DCW.

**RMS**

RMS is another application used to monitor software configuration. RMS compares the configuration of the cluster and each cluster host to a predefined, standard configuration. The standard configuration includes information such as registry settings, program versions, IIS metadata, operating systems versions, and service packs that are installed. If differences are found between the actual cluster or host and the predefined configuration, Microsoft Operations Manager (MOM) is used to push that information to a SQL database.

By retaining and monitoring the data collected by all of these tools, Microsoft.com can quickly address current problems and also use the historical availability reports to monitor and track trends that might indicate future problems.

# Network Load Balancing Best Practices (Using Clustering for a Highly Available Web Site: An Example)

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Microsoft.com Web site implements many of the following best practices, as illustrated throughout this document. Additional best practices have also been included. Web site administrators should review these practices to ensure that they are using Network Load Balancing most effectively.

**Properly Secure the Network Load Balancing Hosts and the Load-Balanced Applications.**

- Network Load Balancing does not provide additional security for the load-balanced hosts and cannot be used as a firewall. Therefore, it is important to properly secure the load-balanced applications and hosts. To find out about security procedures for your applications, see the documentation for each particular application. For example, if you are using Network Load Balancing to load balance a cluster of servers running Internet Information Services (IIS), follow the procedures and guidelines for securing IIS. To view the IIS product documentation, install IIS, open the IIS User Interface (the IIS snap-in), and then click **Help**.

- To avoid interference from unauthorized heartbeat packets, you must physically protect the Network Load Balancing subnet from intrusion by unauthorized computers and devices.

- If you use the optional host list with Network Load Balancing Manager, ensure that only users in the local Administrators group have access to the host list file.

**Use Two or More Network Adapters in Each Cluster Host Whenever Possible. Two Network Adapters, However, are not Required.**

- If the cluster is operating in unicast mode (the default), Network Load Balancing cannot distinguish among single adapters on individual hosts. Therefore, to make communication among cluster hosts possible, each cluster host must have at least two network adapters.

- Although you are able to configure Network Load Balancing on more than one network adapter, make sure that you install Network Load Balancing on only one adapter (called the cluster adapter) if you use a second network adapter to address this best practice.

**Use Only the TCP/IP Network Protocol on the Cluster Adapter.**

- Do not add any other protocols (for example, IPX) to this adapter.

**Use Network Load Balancing Manager to Configure Options.**

- You can configure many Network Load Balancing options using either Network Load Balancing Manager or the **Network Load Balancing Properties** dialog box that you can access through **Network Connections**. However, Network Load Balancing Manager is the preferred method. We recommend that you do not use both Network Load Balancing Manager and **Network Connections** together to change Network Load Balancing properties. Also, you should not use two instances of Network Load Balancing Manager to configure the same cluster.

**Do not Enable Network Load Balancing Remote Control.**

- The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. We strongly recommend that you do not enable remote control. Instead, use Network Load Balancing Manager or other remote management tools such as Windows Management Instrumentation (WMI).

  If you choose to enable remote control, it is vital that you restrict access by specifying a strong remote control password.

It is also imperative that you use a firewall to protect the Network Load Balancing UDP control ports (the ports that receive remote control commands) to shield them from outside intrusion (see Diagram 7). By default, these are ports 1717 and 2504 at the virtual IP address. Use remote control only from a secure, trusted computer within your firewall.




**Diagram 7**

**Enable Network Load Balancing Manager Logging.**

- You can configure Network Load Balancing Manager to log each Network Load Balancing Manager event. This log can be very useful for troubleshooting problems or errors when using Network Load Balancing Manager. Enable Network Load Balancing Manager logging by clicking **Log Settings** in the Network Load Balancing Manager **Options** menu. Select the **Enable logging** check box and specify a name and location for the log file.

  The Network Load Balancing Manager log file contains potentially sensitive information about the Network Load Balancing cluster and hosts, so it must be properly secured. By default, the log file inherits the security settings of the directory in which it is created, so you might have to change the explicit permissions on the file to restrict read and write access to those individuals who do not need full control of the file. Be aware that the individual using Network Load Balancing Manager does require full control of the log file.

**Verify that the Configuration Options for Cluster Parameters, Port Rules and Host Parameters are Set as Follows:**

- Make sure that all hosts in a cluster have identical cluster parameters and port rules for each unique virtual IP address.

  Each unique virtual IP address must be configured with the same port rules for every host that services that virtual IP address. However, if you have multiple virtual IP addresses configured on a host, each of those virtual IP addresses can have a different set of port rules.

- Verify that port rules are set for all ports that are used by the load-balanced application. For example, FTP uses port 20, port 21, and ports 102465535.

- Verify that the dedicated IP address is unique for each host and that the virtual IP address is added to each cluster host.

- If fragmentation of the network traffic is a possibility and you are using **UDP** or **Both** (UDP and TCP/IP) for your protocol setting, verify that the affinity setting is set to **Single** or **Class C**.

**Verify that any Load-Balanced Application is Started on all Cluster Hosts on which the Application is Installed.**

- Network Load Balancing does not start or stop applications.

**Verify that the Configuration Options for Dedicated IP Address and Virtual IP Address are Set as Follows:**

- Make sure that the virtual IP address (VIP) and, optionally the DIP, are entered during Setup in the **Network Load Balancing Properties** dialog box and also in the **Internet Protocol (TCP/IP) Properties** dialog box. Make sure that the addresses are the same in both places. This is most easily achieved by using Network Load Balancing Manager to create and manage your cluster.

  If you omit this step, the cluster will converge and appear to be working properly, but the cluster host will not accept and handle cluster traffic.

For more information about convergence, you can view Help and Support Center topics on the Web in the Product Documentation for Windows Server 2003 section of the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4299). You can also open Help and Support Center on a computer running a product in the Windows Server 2003 family by clicking **Start**, and then clicking **Help and Support**. Look for the topics on Network Load Balancing, which can be found by navigating to Availability and Scalability\Windows Clustering\Network Load Balancing or searching Help.

- Verify that the DIP (if present) is always listed first (before the virtual IP address) in the **Internet Protocol (TCP/IP) Properties** dialog box. This ensures that responses to connections that originate from a host will return to the same host.

- Verify that both the dedicated IP address and the virtual IP address are static IP addresses. They cannot be DHCP addresses.

**Ensure that All Hosts in a Cluster Belong to the Same Subnet and that the Clusters Clients are Able to Access this Subnet.**

**Perform Moves of a Cluster Host According to the Following Guidelines:**

- If you move a cluster host from one cluster to another on the same subnet by changing the virtual IP address, first remove Network Load Balancing and remove the virtual IP address from TCP/IP. Then, re-enable Network Load Balancing after changing the IP address. This will prevent you from experiencing an IP address conflict.

**Verify that all Cluster Hosts are Operating in Either Unicast or Multicast Mode, but not Both.**

- All network adapters within a cluster must be configured identically with regards to unicast or multicast.

**Do not Enable Network Load Balancing on a Computer that is Part of a Server Cluster.**

- Network Load Balancing can interfere with network adapters used by server clusters. Microsoft does not support the same computer being a host in a Network Load Balancing cluster and also being a node within a server cluster.

**Avoid Uninstalling Network Load Balancing.**

- There is typically no need to uninstall this feature. Network Load Balancing is an integral part of the products in the Windows Server 2003 family and does not need to be installed or uninstalled separately.

# Summary (Using Clustering for a Highly Available Web Site: An Example)

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Microsoft.com Web sites use of Network Load Balancing is a typical implementation of this clustering technology, so it serves as a useful example. Web site administrators can use the details of Microsoft.coms host, cluster and port rule configuration as an example for their own Network Load Balancing implementation. Web site administrators should also review the Best Practices section closely to ensure that they are using Network Load Balancing most effectively.

# Related Links (Using Clustering for a Highly Available Web Site: An Example)

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

See the following resources for more information:

- For updated information about Network Load Balancing, you can search the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=4441). Search using both Network Load Balancing and NLB.

- For more information about how the internal IT department at Microsoft has implemented a related technology, server clusters, see the white paper called Using Clustering with Exchange 2003: An Example, which is available at the Microsoft Web site (http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/clustering/scenep2.mspx).

- For information about Network Load Balancing and server clusters, see Availability and Scalability/Windows Clustering in Help and Support Center on any product in the Windows Server 2003 family.

  To view Help and Support Center in the Windows Server 2003 family, click **Start**, and then click **Help and Support**.

- For more information about the following topics, see Planning Server Deployments, at the Microsoft Windows Server 2003 Deployment Kit Web site (http://go.microsoft.com/fwlink/?LinkId=4298):

  - The Network Load Balancing design process

  - Deploying Network Load Balancing, unattended installations, and upgrading existing Network Load Balancing clusters

  - Planning for high availability (only available on this Web site)

- For the latest information about the Windows Server 2003 family, see the Windows Server 2003 section of the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4299).

# Step-by-Step Guide for Setting Up IPv6 in a Test Lab

Updated: July 27, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides detailed information about how you can use five computers to create a test lab with which to configure and test IPv6 with the Microsoft® Windows® XP Professional with Service Pack 2 (SP2) operating system and the 32-bit versions of the Windows Server™ 2003 with Service Pack 1 (SP1) operating system. These step-by-step instructions take you through the configuration required for setting up IPv6 in the test lab.

| **Note** |
| --- |
| The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network. |

**Setting up the IPv6 infrastructure**

The infrastructure for the IPv6 test lab network consists of five computers performing the following services:

- A computer running Windows Server 2003 with SP1, Standard Edition, that is used as a Domain Name System (DNS) server. This computer is named DNS1.
- A computer running Windows XP Professional with SP2 that is used as a client. This computer is named CLIENT1.
- A computer running Windows Server 2003 with SP1, Standard Edition, that is used as a router. This computer is named ROUTER1.
- A computer running Windows Server 2003 with SP1, Standard Edition, that is used as a router. This computer is named ROUTER2.
- A computer running Windows XP Professional with SP2 that is used as a client. This computer is named CLIENT2.

The following figure shows the configuration of the IPv6 test lab.




In the preceding figure, there are three network segments:

- Subnet 1 uses the private IP subnet prefix of 10.0.1.0/24 and global subnet prefix of 2001:DB8:0:1::/64.
- Subnet 2 uses the private IP subnet prefix of 10.0.2.0/24 and global subnet prefix of 2001:DB8:0:2::/64.
- Subnet 3 uses the private IP subnet prefix of 10.0.3.0/24 and global subnet prefix of 2001:DB8:0:3::/64.

All computers on each subnet are connected to a separate common hub or Layer 2 switch. The two router computers, ROUTER1 and ROUTER2, have two network adapters installed.

For the IPv4 configuration, each computer is manually configured with the appropriate IP address, subnet mask, default gateway, and DNS server IP address. For the IPv6 configuration, link-local addresses are used initially. Dynamic Host Configuration Protocol (DHCP) and Windows Internet Name Service (WINS) servers are not used.

The configuration of the lab routing infrastructure goes through three phases.

1. IPv4 routing. This phase demonstrates IPv4 connectivity and the automatic configuration of IPv6 local-link addresses - similar to Automatic Private IP Addressing (APIPA).

2. IPv6 routing. This phase demonstrates all test lab nodes can be reached by using IPv6 traffic.

3. ISATAP routing. ISATAP is an address assignment and automatic tunneling technology that is used to provide unicast IPv6 connectivity between IPv6/IPv4 hosts over an IPv4 intranet. This phase removes IPv6 connectivity for Subnet 2 and Subnet 3 and restores it using ISATAP.

The following sections describe how to configure each of the computers in the test lab with the basic IPv6 infrastructure. To create this test lab, configure the computers in the order presented.

**Consider using Virtual PC or Virtual Server**
Microsoft Virtual PC or Virtual Server allow you to create the computer lab used in this document using only one physical computer. After the virtual lab is configured, you can switch between the five virtual computers needed for this lab with the click of a button. For more information, see the following resources:

- Virtual PC 2004 Product Information (http://go.microsoft.com/fwlink/?LinkId=69217)

- Virtual Server Product Information (http://go.microsoft.com/fwlink/?LinkId=69220)

- Do More With Less: Exploring Virtual Server 2005 (http://go.microsoft.com/fwlink/?LinkId=69221)

- TechNet Webcast: Virtual Server 2005 - Setting Up a Virtual Test and Development Environment—Level 200 (http://go.microsoft.com/fwlink/?LinkId=69222)

**DNS1**
DNS1 is a computer running Windows Server 2003 with SP1, Standard Edition. It is providing DNS Server services for the testlab.microsoft.com DNS domain. To configure DNS1 for this service, perform the following steps.

**Configure DNS1 to provide DNS Services**

1. Install Windows Server 2003 with SP1, Standard Edition, as a standalone server. Set the Administrator password.

2. After restarting, log on as Administrator.

3. Configure the TCP/IP protocol with the IP address of 10.0.1.2, the subnet mask of 255.255.255.0, and the default gateway of 10.0.1.1.

**Install the DNS Server service**

1. Open Windows Components Wizard. To open the Windows Components Wizard, click **Start**, click Control Panel, double-click **Add or Remove Programs**, and then click **Add/Remove Windows Components**.

2. In **Components**, select the **Networking Services** check box, and then click **Details**.

3. In **Subcomponents of Networking Services**, select the **Domain Name System (DNS)** check box, click **OK**, and then click **Next**.

4. If prompted, in **Copy files from**, type the full path to the distribution files, and then click **OK**.

Define a forward lookup zone named testlab.microsoft.com that allows dynamic updates.

**Define a forward lookup zone**

1. Open DNS. To open DNS, click **Start**, select **Administrative Tools**, and then click **DNS**.

2. In the console tree, click and then right-click the DNS server **DNS1**, and then click **New Zone** to run the New Zone Wizard.

3. On the **Welcome to the New Zone Wizard** page, click **Next**.

4. On the **Zone Type** page, the option to create a primary zone is selected by default. Click **Next**.

5. On the **Forward or Reverse Lookup Zone** page, the option to create a forward lookup zone is selected by default. Click **Next**.

6. On the **Zone Name** page, type **testlab.microsoft.com** in **Zone name**. Click **Next**.

| **Note** |
| --- |
| The domain name testlab.microsoft.com is used here for example purposes only. You can use any domain name in your test lab configuration. |

7. On the **Zone File** page, the new zone file name, **testlab.microsoft.com.dns**, is automatically derived from the zone name entered on the previous page, as shown in the following figure. Click **Next**.




8. On the **Dynamic Update** page, select **Allow both nonsecure and secure dynamic updates**. Click **Next**.

9. Click **Next**. On the **Completing the New Zone Wizard** page, click **Finish**.

### Install IPv6

- On DNS1, at the command prompt, type:

  **netsh interface ipv6 install**

#### CLIENT1

CLIENT1 is a computer that is being used as a client.

### Configure CLIENT1 as a client computer

1. Install Windows XP Professional with SP2 as a workgroup computer. Set the Administrator password.

2. After restarting, log on as Administrator.

3. At the command prompt, install the IPv6 protocol by typing:

   **netsh interface ipv6 install**

4. Configure the TCP/IP protocol with the IP address of 10.0.1.3, the subnet mask of 255.255.255.0, a default gateway of 10.0.1.1, and the DNS server IP address of 10.0.1.2.

5. Append the DNS suffix **testlab.microsoft.com**. Click **Start**, click **Control Panel**, click **Network and Internet**

**Connections**, and then click **Network Connections**. Right-click the network connection that you want to configure, and then click **Properties**. On the **General** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**. Click **Advanced**, and then click the **DNS** tab. Add **testlab.microsoft.com** to the list of DNS suffixes, as show in the following dialog box:




Windows Firewall is automatically turned on in Windows XP Professional with SP2. You need to configure an exception to allow CLIENT2 to ping CLIENT1.

**Configure Windows to allow communication between client computers**

1. Click **Start**, point to **Control Panel**, and then click **Security Center**.
2. Click **Windows Firewall**, and then in the **Windows Firewall** dialog box, click the **Advanced** tab.
3. Click **Settings** for **ICMP**, and then click **Allow incoming echo request**.
4. Click **OK** twice to close **Windows Firewall**.

**ROUTER1**

ROUTER1 is a computer that is being used as a router between Subnet 1 and Subnet 2.

**Configure ROUTER1 as a router**

1. Install Windows Server 2003 with SP1, Standard Edition, as a workgroup computer. Set the Administrator password.
2. After restarting, log on as Administrator.
3. At the command prompt, install the IPv6 protocol by typing:

   **netsh interface ipv6 install**

4. In **Control Panel-Network Connections**, rename the LAN connection connected to Subnet 1 to **Subnet 1 Connection** and rename the LAN connection connected to Subnet 2 to **Subnet 2 Connection**.
5. For **Subnet 1 Connection**, configure the TCP/IP protocol with the IP address of 10.0.1.1, the subnet mask of 255.255.255.0, and the DNS server IP address of 10.0.1.2.
6. For **Subnet 2 Connection**, configure the TCP/IP protocol with the IP address of 10.0.2.1, the subnet mask of 255.255.255.0, and a default gateway of 10.0.2.2.
7. Start the registry editor (Regedit.exe) and set HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\\IPEnableRouter to 1.

This step enables IPv4 routing between Subnet 1 and Subnet 2.

8. Restart the computer.

**ROUTER2**
ROUTER2 is a computer that is being used as a router between Subnet 2 and Subnet 3.

**Configure ROUTER2 as a router**

1. Install Windows Server 2003 with SP1, Standard Edition, as a workgroup computer. Set the Administrator password.
2. After restarting, log on as Administrator.
3. At the command prompt, install the IPv6 protocol by typing:

   **netsh interface ipv6 install**

4. Open Network Connections, and rename the LAN connection connected to Subnet 2 to **Subnet 2 Connection**, and rename the LAN connection connected to Subnet 3 to **Subnet 3 Connection**.
5. For **Subnet 2 Connection**, configure the TCP/IP protocol with the IP address of 10.0.2.2, the subnet mask of 255.255.255.0, and a default gateway of 10.0.2.1.
6. For **Subnet 3 Connection**, configure the TCP/IP protocol with the IP address of 10.0.3.1 and the subnet mask of 255.255.255.0.
7. Start the registry editor (Regedit.exe) and set HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\IPEnableRouter to 1.

   This step enables IPv4 routing between Subnet 2 and Subnet 3.

8. Restart the computer.

**CLIENT2**
CLIENT2 is a computer that is being used as a client.

**Configure CLIENT2 as a client computer**

1. Install Windows XP Professional with SP2 as a workgroup computer. Set the Administrator password.
2. After restarting, log on as Administrator.
3. At the command prompt, install the IPv6 protocol by typing:

   **netsh interface ipv6 install**

4. Configure the TCP/IP protocol with the IP address of 10.0.3.2, the subnet mask of 255.255.255.0, a default gateway of 10.0.3.1, and the DNS server IP address of 10.0.1.2, and append the DNS suffix **testlab.microsoft.com**.
5. Verify the integrity of the IPv4 routing infrastructure by pinging 10.0.1.3 from the CLIENT2 computer. On CLIENT2, type the following command:

   **ping 10.0.1.3**

   You should be able to successfully ping CLIENT1.

**Link-local addresses**
Link-local addresses, identified by the beginning address block FE80, are equivalent to APIPA IPv4 addresses autoconfigured on computers running current Microsoft Windows operating systems using the 169.254.0.0/16 prefix. One of the most useful aspects of IPv6 is its ability to automatically configure itself without the use of a stateful configuration protocol, such as Dynamic Host Configuration Protocol for IPv6 (DHCPv6). By default, an IPv6 node configures a link-local address for each interface. Link-local addresses are used by nodes when communicating with neighboring nodes on the same link (a network segment bounded by routers). The scope of a link-local address is the local link. An IPv6 router never forwards link-local traffic beyond the link.

**Link-local ping**
At this point in the lab configuration, only IPv4 traffic across the subnets is routed through an IPv4 routing infrastructure. However, you should be able to successfully ping neighboring link-local addresses. After the lab is configured to route IPv6

traffic, you will be able to successfully ping using global addresses.

**Determine and ping the link-local address of CLIENT1 from DNS1**

1. On CLIENT1, type the **ipconfig** command to obtain the link-local address and interface index of the interface named Local Area Connection.




In this example, the link-local address is FE80::203:FFFF:FEE1:2A73 and the interface index is 5.

2. On DNS1, type the **ipconfig** command to obtain the interface index of the interface named Local Area Connection.

   In this example, the interface index of Local Area Connection is 4.




3. On DNS1, ping the IPv6 link-local address of CLIENT1. You must use the interface index of the sending interface, the Local Area Connection interface of DNS1. For example, using the example configuration in step 1:

   **ping FE80::203:FFFF:FEE1:2A73%4**

   A successful ping demonstrates that CLIENT1 and DNS1 are neighbors on the same link.

**Creating a static IPv6 routing infrastructure**

Configure a static IPv6 routing infrastructure so that all test lab nodes can be reached by using IPv6 traffic.

**Create a static IPv6 routing infrastructure**

1. On ROUTER1, type the **ipconfig** command to obtain the link-local addresses of the interfaces connected to Subnet 1 Connection and Subnet 2 Connection.

```
Command Prompt

C:\Documents and Settings\Administrator>ipconfig

Windows IP Configuration

Ethernet adapter Subnet 1 Connection:

   Connection-specific DNS Suffix  . :
   IP Address. . . . . . . . . . . . : 10.0.1.1
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   IP Address. . . . . . . . . . . . : fe80::203:ffff:fee1:fa74%4
   Default Gateway . . . . . . . . . :

Ethernet adapter Subnet 2 Connection:

   Connection-specific DNS Suffix  . :
   IP Address. . . . . . . . . . . . : 10.0.2.1
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   IP Address. . . . . . . . . . . . : fe80::203:ffff:fefc:fa75%5
   Default Gateway . . . . . . . . . : 10.0.2.2

Tunnel adapter Teredo Tunneling Pseudo-Interface:

   Connection-specific DNS Suffix  . :
```

For this example, the information is as follows:

Subnet 1 Connection - FE80::203:FFFF:FEE1:FA74

Subnet 2 Connection - FE80::203:FFFF:FEFC:FA75

2. On ROUTER2, type the **ipconfig** command to obtain the link-local addresses of the interfaces connected to Subnet 2 Connection and Subnet 3 Connection.

```
Command Prompt

C:\Documents and Settings\Administrator>ipconfig

Windows IP Configuration

Ethernet adapter Subnet 2 Connection:

   Connection-specific DNS Suffix  . :
   IP Address. . . . . . . . . . . . : 10.0.2.2
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   IP Address. . . . . . . . . . . . : fe80::203:ffff:fee0:fa76%4
   Default Gateway . . . . . . . . . : 10.0.2.1

Ethernet adapter Subnet 3 Connection:

   Connection-specific DNS Suffix  . :
   IP Address. . . . . . . . . . . . : 10.0.3.1
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   IP Address. . . . . . . . . . . . : fe80::203:ffff:fefd:fa77%5
   Default Gateway . . . . . . . . . :

Tunnel adapter Teredo Tunneling Pseudo-Interface:

   Connection-specific DNS Suffix  . :
```

For this example, the information is as follows:

Subnet 2 Connection - FE80::203:FFFF:FEE0:FA76

Subnet 3 Connection - FE80::203:FFFF:FEFD:FA77

3. On ROUTER1, type the following commands:

   a. **netsh interface ipv6 set interface "Subnet 1 Connection" forwarding=enabled advertise=enabled**

   b. **netsh interface ipv6 set interface "Subnet 2 Connection" forwarding=enabled advertise=enabled**

      Steps a. and b. configure the Subnet 1 and 2 Connection interfaces to route (forward IPv6 packets not addressed to them) and advertise subnet prefixes (for address autoconfiguration).

   c. **netsh interface ipv6 add route 2001:db8:0:1::/64 "Subnet 1 Connection" publish=yes**

   d. **netsh interface ipv6 add route 2001:db8:0:2::/64 "Subnet 2 Connection" publish=yes**

      Steps c. and d. add routes to the IPv6 routing table corresponding to the subnet prefix for Subnet 1 and 2 and configures the router to advertise these routes in Router Advertisement messages.

e. **netsh interface ipv6 add route ::/0 "Subnet 2 Connection" nexthop=** *ROUTER2AddressOnSubnet2* **publish=yes**

Step e. adds the default route to the IPv6 routing table and configures the router to advertise itself as a default router in the Router Advertisement messages.

In the preceding command, *ROUTER2AddressOnSubnet2* represents the link-local address assigned to the Subnet 2 Connection interface on ROUTER2. For example, using the connection in this lab the command would be:

**netsh interface ipv6 add route ::/0 "Subnet 2 Connection" nexthop= FE80::203:FFFF:FEE0:FA76 publish=yes**

4. On ROUTER2, type the following commands:

   a. **netsh interface ipv6 set interface "Subnet 2 Connection" forwarding=enabled advertise=enabled**

   b. **netsh interface ipv6 set interface "Subnet 3 Connection" forwarding=enabled advertise=enabled**

   c. **netsh interface ipv6 add route 2001:db8:0:2::/64 "Subnet 2 Connection" publish=yes**

   d. **netsh interface ipv6 add route 2001:db8:0:3::/64 "Subnet 3 Connection" publish=yes**

   e. **netsh interface ipv6 add route ::/0 "Subnet 2 Connection" nexthop=** *ROUTER1AddressOnSubnet2* **publish=yes**

   In the preceding command, *ROUTER1AddressOnSubnet2* represents the link-local address assigned to the Subnet 2 Connection interface on ROUTER1. For example, using the connection in this lab the command would be:

   **netsh interface ipv6 add route ::/0 "Subnet 2 Connection" nexthop= FE80::203:FFFF:FEFC:FA75 publish=yes**

**Global addresses**

At this point, you have now created static IPv6 routes with global subnet prefixes. Global addresses are equivalent to the IPv4 public address space. Unlike link-local addresses, global addresses are not automatically configured and must be assigned either through stateless or stateful address configuration processes. You can determine which address is the global address by the "ff:fe" portion of the public address in the sixth and seventh blocks of the address, indicating an EUI-64-based interface ID. The Extended Unique Identifier (EUI)-64 address is a newer 64-bit MAC address.

**Global ping**

**Verify the IPv6 routing structure**

1. On CLIENT1, type the **ipconfig** command to check for a new global IPv6 address.




   In this example, the CLIENT1 global address is 2001:DB8:0:1:203:FFFF:FEE1:2A73.

2. On CLIENT2, type the following commands:

   **ping** *CLIENT1GlobalAddress*

**tracert -d** *CLIENT1GlobalAddress*

In this example, the CLIENT1 global address is 2001:DB8:0:1:203:FFFF:FEE1:2A73.

A successful ping and tracert demonstrates that IPv6 static routes have been created and are functioning.

3. You can view the entries in the ROUTER1 neighbor cache for CLIENT1 and ROUTER2, by typing the following on ROUTER1:

   **netsh interface ipv6 show neighbors**

**Using name resolution**

To resolve host names to IPv6 addresses, you must first configure DNS. On DNS1, create an AAAA record for CLIENT2 with the DNS name client2.testlab.microsoft.com for its global IPv6 address using the IPv6 Host resource record type.

**Configure DNS to resolve names to IPv6 addresses**

1. On DNS1, click **Start**, select **Administrative Tools**, and then click **DNS**.
2. In the console tree, click and then right-click **testlab.microsoft.com** in the **Forward Lookup Zones** folder, and then click **Other New Records**.
3. Click **IPv6 Host (AAAA)**, and then click **Create Record**.
4. In the **Host** text box type **client1**, and then in the **IP version 6 host address** text box, enter the IPv6 global address for CLIENT1.

   In this example, the CLIENT1 global address is 2001:DB8:0:1:203:FFFF:FEE1:2A73.




5. Click **OK**, and then click **Done** to add the new IPv6 host record to the zone.
6. On CLIENT 2, type the following command:

   **ping client1**

   A successful ping demonstrates that host names can resolve to IPv6 addresses.

**Configuring the test lab to use ISATAP**

ISATAP is an address assignment and automatic tunneling technology that is used to provide unicast IPv6 connectivity between IPv6/IPv4 hosts on an IPv4 intranet, which are also known as ISATAP hosts. ISATAP is described in RFC 4214. More information about ISATAP can be found in the IPv6 Transition Technologies white paper (http://go.microsoft.com/fwlink/?LinkId=67210). By default, the IPv6 protocol for Windows XP Professional with SP2 and Windows Server 2003 with SP1, Standard Edition, configures a link-local ISATAP address on the Automatic Tunneling Pseudo-Interface for each IPv4 address assigned to a computer.

To configure global ISATAP addresses, or to communicate beyond the logical subnet defined by the IPv4 intranet, you need an ISATAP router. An ISATAP router performs the following functions:

- Advertises its presence and address prefixes, enabling global ISATAP addresses to be configured.
- Optionally forwards IPv6 packets between ISATAP hosts on the IPv4 intranet and IPv6 hosts beyond it.

An ISATAP router is typically configured to perform both functions, but can perform either individually. Most often, an ISATAP router acts as the forwarder between ISATAP hosts on an IPv4 intranet and IPv6 hosts on an IPv6-enabled portion of an intranet.

To demonstrate the use of an ISATAP router between IPv6 and IPv4 intranets, the following steps first separate the lab into a portion that has IPv4 and IPv6 connectivity and another that has IPv4 connectivity only. Then, ROUTER1 is configured as an ISATAP router so that hosts on the IPv4-only portion of the intranet can communicate with hosts on the IPv6-enabled portion of the intranet.

To do this, we will disable IPv6 forwarding and advertising on the Subnet 2 Connection interface of ROUTER1 and both interfaces of ROUTER2. This emulates an intranet in which a portion is IPv6-enabled (Subnet 1) and a portion is not (Subnet 2 and Subnet 3).

**Configure the test lab to use ISATAP**

1. Disable forwarding and advertising on the Subnet 2 Connection interface of ROUTER1 and ROUTER2.

   a. On ROUTER1, at the command prompt, type:

      **netsh interface ipv6 set interface "Subnet 2 Connection" forwarding=disabled advertise=disabled**

   b. On ROUTER2, at the command prompt, type:

      **netsh interface ipv6 set interface "Subnet 2 Connection" forwarding=disabled advertise=disabled**

2. Disable forwarding and advertising on the Subnet 3 Connection interface. On ROUTER2, at the command prompt, type:

   **netsh interface ipv6 set interface "Subnet 3 Connection" forwarding=disabled advertise=disabled**

3. Renew CLIENT2 configuration to remove the global address and route. On CLIENT2, type the following command:

   **netsh interface ipv6 renew**

4. Ping CLIENT1 from CLIENT2. On CLIENT2, type the following command:

   **ping** *CLIENT1GlobalAddress*

   or

   **ping client1**

   In this example, the CLIENT1 global address is 2001:DB8:0:1:203:FFFF:FEE1:2A73.

   This Ping command fails because there is no longer IPv6 connectivity between CLIENT2 and CLIENT1.

5. Configure ROUTER1 as an ISATAP router advertising the 2001:DB8:0:10::/64 route to the logical ISATAP subnet consisting of Subnet 2 and Subnet 3.

   a. On ROUTER1, enable forwarding and advertising on the Automatic Tunneling Pseudo-Interface, using the following commands:

      **netsh interface ipv6 isatap set router 10.0.2.1**

      **netsh interface ipv6 set interface "Automatic Tunneling Pseudo-Interface" forwarding=enabled advertise=enabled**

   b. On ROUTER1, add a route for the subnet prefix of the logical subnet of the IPv4 intranet to the Automatic Tunneling

Pseudo-Interface and configure it to be published. Use the following command:

**netsh interface ipv6 add route 2001:db8:0:10::/64 "Automatic Tunneling Pseudo-Interface" publish=yes**

6. Add an ISATAP address (A) resource record to the DNS server for the IPv4 address of 10.0.2.1.

   a. On DNS1, click **Start**, select **Administrative Tools**, and then click **DNS**.

   b. In the console tree, click and then right-click **testlab.microsoft.com** in the **Forward Lookup Zones** folder, and then click **New Host (A)**.

   c. In the **Name** text box, type **ISATAP**, and in the **IP address** text box, type **10.0.2.1**, as shown in the following figure.




   d. Click **Add Host**, click **OK**, and then click **Done** to add the new host record to the zone.

7. Renew the CLIENT2 configuration to recognize the ISATAP router IPv4 address of ROUTER1 and obtain an ISATAP address with the global prefix of 2001:DB8:0:10::/64. On CLIENT2, type the following command:

   **netsh interface ipv6 renew**

8. Ping CLIENT1 from CLIENT2. On CLIENT2, type the following command:

   **ping** *CLIENT1GlobalAddress*

   or

   **ping client1**

   In this example, the CLIENT1 global address is 2001:DB8:0:1:203:FFFF:FEE1:2A73.

   This ping command succeeds because IPv6 connectivity using ISATAP now exists between CLIENT2 and CLIENT1. On the link between CLIENT1 and ROUTER1, the IPv6 traffic is encapsulated as IPv4 packets.

| **Note** |
| --- |
| It may take a few minutes for IPv6 connectivity between CLIENT2 and CLIENT1 to be established. You may need to restart CLIENT2 and run **netsh interface ipv6 renew** on CLIENT1. |

**Conclusion**

This guide described the steps required to configure IPv6 in a test lab with five computers. It also described how to use ISATAP to provide IPv6 connectivity between IPv6/IPv4 hosts on an IPv4-only portion of an intranet and IPv6 hosts on an IPv6-enabled portion of an intranet.

**Additional references**

For more information, see the following resources:

- Microsoft IPv6 Web site (http://go.microsoft.com/fwlink/?LinkId=24350)
- Introduction to IPv6 (http://go.microsoft.com/fwlink/?LinkId=69223)
- Support WebCast: Introduction to Internet Protocol Version 6 (http://go.microsoft.com/fwlink/?LinkId=69224)
- Using Windows Tools to Obtain IPv6 Configuration Information (http://go.microsoft.com/fwlink/?LinkId=69225)
- Windows Server 2003 Web site (http://go.microsoft.com/fwlink/?LinkId=5221)

# Step-by-Step Guide for Setting Up VPN-based Remote Access in a Test Lab

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides detailed information about how you can use five computers to create a test lab with which to configure and test virtual private network (VPN) remote access with the Microsoft® Windows® XP Professional operating system with Service Pack 2 (SP2) and the 32-bit versions of the Microsoft Windows Server™ 2003 operating system with Service Pack 1 (SP1). These instructions are designed to take you step-by-step through the configuration required for a Point-to-Point Tunneling Protocol (PPTP) connection, a Layer Two Tunneling Protocol (L2TP) with Internet Protocol security (L2TP/IPsec) connection, and a VPN connection that uses certificate-based Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication.

| **Note** |
| --- |
| The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to show the desired functionality clearly. This configuration is designed to reflect neither best practices nor a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed to work only on a separate test lab network. |

**Setting Up the Test Lab for PPTP, L2TP/IPsec, and EAP-TLS Remote Access VPN Connections**

The infrastructure for the VPN test lab network consists of five computers performing the following services:

- A computer running Windows Server 2003 with SP1, Enterprise Edition, named DC1 that is acting as a domain controller, a Domain Name System (DNS) server, a Dynamic Host Configuration Protocol (DHCP) server, and a certification authority (CA).
- A computer running Windows Server 2003 with SP1, Standard Edition, named IAS1 that is acting as a Remote Authentication Dial-In User Service (RADIUS) server.
- A computer running Windows Server 2003 with SP1, Standard Edition, named IIS1 that is acting as a Web and file server.
- A computer running Windows Server 2003 with SP1, Standard Edition, named VPN1 that is acting as a VPN server. VPN1 has two network adapters installed.
- A computer running Windows XP Professional with SP2 named CLIENT1 that is acting as a VPN client.

The following diagram shows the configuration of the VPN test lab.




There is a network segment representing a corporate intranet and a network segment representing the Internet. All computers on the corporate intranet are connected to a common hub or Layer 2 switch. All computers on the Internet are connected to a separate common hub or Layer 2 switch. Private addresses are used throughout the test lab configuration. The private network of 172.16.0.0/24 is used for the intranet. The private network of 10.0.0.0/24 is used for the simulated Internet. Windows

Firewall is set up and configured on the RADIUS server (IAS1), the Web and file server (IIS1), and the client computer (CLIENT1). Windows Firewall should not be turned on or configured on either the domain controller (DC1) or the VPN server (VPN1). In addition, the Windows Firewall/Internet Connection Sharing (ICS) service should be disabled on VPN1.

IIS1 obtains its IP address configuration using DHCP. CLIENT1 uses DHCP for its IP address configuration; however, it is also configured with an alternate IP configuration so that it can be placed on either the intranet network segment or the simulated Internet. All other computers have a manual IP address configuration. There are no Windows Internet Name Service (WINS) servers present.

To reconstruct this test lab, configure the computers in the order presented, beginning with the PPTP-based remote access VPN connection. Additional sections of this guide describe L2TP/IPsec-based and EAP-TLS-based remote access VPN connections.

**PPTP-based Remote Access VPN Connections**

The following sections describe how to set up and configure each of the computers in the test lab for a PPTP-based remote access VPN connection. PPTP is typically used when there is no public key infrastructure (PKI) to issue computer certificates that are required for L2TP/IPsec connections.

**DC1**

DC1 is a computer running Windows Server 2003 with SP1, Enterprise Edition, that is providing the following services:

- A domain controller for the example.com Active Directory® domain.
- A DNS server for the example.com DNS domain.
- A DHCP server for the intranet network segment
- The enterprise root certification authority (CA) for the example.com domain.

| **Note** |
| --- |
| Windows Server 2003 with SP1, Enterprise Edition, is used so that autoenrollment of user certificates for EAP-TLS authentication can be configured. This is described in the "EAP-TLS-based Remote Access VPN Connections" section of this guide. |

**Configure DC1**

1. Install Windows Server 2003 with SP1, Enterprise Edition, as a stand-alone server.
2. Configure the TCP/IP protocol with the IP address of 172.16.0.1 and the subnet mask of 255.255.255.0.

**Configure DC1 as a domain controller**

1. To start the Active Directory Installation Wizard, click **Start**, click **Run**, type **dcpromo**, and then click **OK**.
2. In the **Welcome to the Active Directory Installation Wizard** dialog box, click **Next**.
3. In the **Operating System Compatibility** dialog box, click **Next**.
4. Verify that **Domain controller for a new domain** option is selected, and then click **Next**.
5. Verify that **Domain in a new forest** is selected, and then click **Next**.
6. Verify that **No, just install and configure DNS on this computer** is selected, and then click **Next**.
7. On the **New Domain Name** page, type **example.com**, and then click **Next**.
8. On the **NetBIOS Domain Name** page, confirm that the Domain NetBIOS name is **EXAMPLE**, and then click **Next**.
9. Accept the default **Database and Log Folders** directories, as shown in the following figure, and then click **Next**.




10. In the **Shared System Volume** dialog box, shown in the following figure, verify that the default folder location is correct. Click **Next**.




11. On the **Permissions** page, verify that the **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems** check box is selected, as shown in the following figure. Click **Next**.




12. On the **Directory Services Restore Mode Administration Password** page, leave the passwords blank, and then click

**Next**.

13. Review the information that appears on the **Summary** page, and then click **Next**.




14. On the **Completing the Active Directory Installation Wizard** page, click **Finish**.

15. When prompted to restart the computer, click **Restart Now**.

**Raise the domain functional level**

1. Open the **Active Directory Domains and Trusts** snap-in from the **Administrative Tools** folder, and then right-click the domain computer **dc1.example.com**.

2. Click **Raise Domain Functional Level**, select **Windows Server 2003** on the **Raise Domain Functional Level** page, and then click **Raise**, as shown in the following figure.




**Install and configure DHCP**

1. In Control Panel, double-click **Add or Remove Programs**, and then install DHCP as a Networking Services component.

2. Open the **DHCP** snap-in from the **Administrative Tools** folder.

3. Click **Action**, and then click **Authorize** to authorize the DHCP service.

4. In the console tree, right-click **dc1.example.com**, and then click **New Scope**.

5. On the **Welcome** page of the New Scope Wizard, click **Next**.

6. On the **Scope Name** page, type **CorpNet** in **Name**. This is shown in the following figure.




7. Click **Next**. On the **IP Address Range** page, type **172.16.0.10** in **Start IP address**, **172.16.0.100** in **End IP address**, and **24** in **Length**. This is shown in the following figure.




8. Click **Next**. On the **Add Exclusions** page, click **Next**.

9. On the **Lease Duration** page, click **Next**.

10. On the **Configure DHCP Options** page, click **Yes, I want to configure DHCP options now**. This is shown in the following figure.




11. Click **Next**. On the **Router (Default Gateway)** page, click **Next**.

12. On the **Domain Name and DNS Servers** page, type **example.com** in **Parent domain**. Type **172.16.0.1** in **IP address**, and then click **Add**. This is shown in the following figure.




13. Click **Next**. On the **WINS Servers** page, click **Next**.

14. On the **Activate Scope** page, click **Yes, I want to activate this scope now**. This is shown in the following figure.




15. Click **Next**. On the **Completing the New Scope Wizard** page, click **Finish**.

**Install Certificate Services**

1. In Control Panel, double-click **Add or Remove Programs**, and then install the Certificate Services component as an enterprise root CA with the name **Example CA**.

2. Select **Enterprise root CA**, as shown in the following figure, and then click **Next**.




3. Type **Example CA** for the **Common name for this CA**, as shown in the following figure, and then click **Next**.




4. Click **Next** to accept the default **Certificate Database Settings** shown in the following figure.




5. Click **Finish**.

**Add computers, users, and groups to the domain**

1. Open the Active Directory Users and Computers snap-in.
2. In the console tree, open **example.com**.
3. Right-click **Users**, point to **New**, and then click **Computer**.
4. In the **New Object - Computer** dialog box, type **IAS1** in **Computer name**. This is shown in the following figure.




5. Click **Next**. In the **Managed** dialog box, click **Next**. In the **New Object - Computer** dialog box, click **Finish**.
6. Use steps 3 through 5 to create additional computer accounts with the following names: **IIS1**, **VPN1**, and **CLIENT1**.
7. In the console tree, right-click **Users**, point to **New**, and then click **User**.
8. In the **New Object - User** dialog box, type **VPNUser** in **First name**, and type **VPNUser** in **User logon name**. This is shown in the following figure.




9. Click **Next**.
10. In the **New Object - User** dialog box, type a password of your choice in **Password** and **Confirm password**. Clear the **User must change password at next logon** check box and select the **Password never expires** check box. This is shown in the following figure.




11. In the **New Object** - **User** dialog box, click **Finish**.

12. In the console tree, right-click **Users**, point to **New**, and then click **Group**.

13. In the **New Object - Group** dialog box, type **VPNUsers** in **Group name**, and then click **OK**. This is shown in the following figure.




14. In the details pane, double-click **VPNUsers**.

15. Click the **Members** tab, and then click **Add**.

16. In the **Select Users, Contacts, Users, or Groups** dialog box, type **vpnuser** in **Enter the object names to select**. This is shown in the following figure.




17. Click **OK**. In the **Multiple Names Found** dialog box, click **OK**. The VPNUser user account is added to the VPNUsers group. This is shown in the following figure.




18. Click **OK** to save changes to the VPNUsers group.

**IAS1**

IAS1 is a computer running Windows Server 2003 with SP1, Standard Edition, that is providing RADIUS authentication, authorization, and accounting for VPN1.

**Configure IAS1 as a RADIUS server**

1. Install Windows Server 2003 with SP1, Standard Edition, as a member server named IAS1 in the example.com domain.
2. For the intranet local area connection, configure the TCP/IP protocol with the IP address of 172.16.0.2, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.
3. In Control Panel, double-click **Add or Remove Programs**, and then install Internet Authentication Service as a Networking Services component.
4. Open the Internet Authentication Service snap-in from the **Administrative Tools** folder.
5. Right-click **Internet Authentication Service**, and then click **Register Server in Active Directory**. When the **Register Internet Authentication Server in Active Directory** dialog box appears, click **OK**. This is shown in the following figure.

6. In the console tree, right-click **RADIUS Clients**, and then click **New RADIUS Client**.
7. On the **Name and Address** page of the New RADIUS Client wizard, for **Friendly name**, type **VPN1**. In **Client address (IP or DNS)**, type **172.16.0.4**. This is shown in the following figure.




8. Click **Next**. On the **Additional Information** page of the New RADIUS Client wizard, for **Shared secret**, type a shared secret for VPN1, and then type it again in **Confirm shared secret**. This is shown in the following figure.




9. Click **Finish**.

10. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**.

11. On the **Welcome to the New Remote Access Policy Wizard** page, click **Next**.

12. On the **Policy Configuration Method** page, type **VPN remote access to intranet** in **Policy name**. This is shown in the following figure.




13. Click **Next**. On the **Access Method** page, select **VPN**. This is shown in the following figure.




14. Click **Next**. On the **User or Group Access** page, click **Group**. This is shown in the following figure.




15. Click **Add**. In the **Select Groups** dialog box, click **Locations**, select **example.com** as the location, and then click **OK**.

16. Type **vpnusers** in **Enter the object names to select**. This is shown in the following figure.




17. Click **OK**. The VPNUsers group in the example.com domain is added to the list of groups on the **User or Group Access** page. This is shown in the following figure.




18. Click **Next**. On the **Authentication Methods** page, the **Microsoft Encrypted Authentication version 2 (MS-CHAPv2)** authentication protocol is selected by default. This is shown in the following figure.




19. Click **Next**. On the **Policy Encryption Level** page, clear the **Basic encryption** and **Strong encryption** check boxes, leaving only **Strongest encryption** selected. This is shown in the following figure.




20. Click **Next**. On the **Completing the New Remote Access Policy** page, click **Finish**.

**Configure Windows Firewall on IAS1**

1. In Control Panel, double-click **Windows Firewall**.
2. In the **Windows Firewall** dialog box, click the **Exceptions** tab.
3. Click **Add Port**, and in the **Add a Port** dialog box add the following port exceptions:

| **Note** |
|---|
| You must click **Add Port** on the **Exceptions** tab for each port exception. |

| Name | Port Number | Protocol |
|---|---|---|
| Legacy RADIUS | 1645 | UDP |
| Legacy RADIUS | 1646 | UDP |
| RADIUS Accounting | 1812 | UDP |
| RADIUS Authentication | 1813 | UDP |

4. Verify that the four port exceptions that you added are selected on the **Exceptions** tab, as shown in the following example.




5. Click the **Advanced** tab, and then click **Settings** for **Security Logging**.

6. In the **Log Setting** dialog box, select **Log dropped packets** and **Log successful connections**. Note the path and file name in **Name**.

   The log file allows you to see where connection errors occur, as well as which source and destination ports the errors occurred on. This log file should provide you with the information needed in case you need to add more ports to the exception list.

7. Click **OK** twice to close Windows Firewall.

**IIS1**

IIS1 is a computer running Windows Server 2003 with SP1, Standard Edition, and Internet Information Services (IIS). It is providing Web and file server services for intranet clients. To configure IIS1 as a Web and file server and to configure Windows Firewall on IIS1, perform the following steps.

**Configure IIS1 as a Web and file server**

1. Install Windows Server 2003 with SP1, Standard Edition, as a member server named IIS1 in the example.com domain.

2. In Control Panel, double-click **Add or Remove Programs**, and then install Internet Information Services (IIS) as a subcomponent of the Application Server component in the Windows Components Wizard.

3. On IIS1, use Windows Explorer to create a new share for the root folder of drive C using the share name ROOT with the default permissions.

4. To determine whether the Web server is working correctly, run Internet Explorer on IAS1. If the Internet Connection Wizard prompts you, configure Internet connectivity for a LAN connection. In Internet Explorer, in **Address**, type **http://IIS1.example.com/iisstart.htm**. You should see a message saying the Web site is under construction.

5. To determine whether file sharing is working correctly, on IAS1, click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**. You should see the contents of the root folder of drive C on IIS1.

**Configure Windows Firewall on IIS1**

1. In Control Panel, double-click **Windows Firewall**.

2. In the **Windows Firewall** dialog box, click the **Exceptions** tab.

3. Select **File and Print Sharing**, and then click **Add Program**.

4. In the **Add a Program** dialog box, select **Internet Explorer**, and then click **OK**.
5. Click **Add a Port**.
6. In the **Add a Port** dialog box, type **World Wide Web Publishing Service** for the **Name**, type **80** for the **Port number**, select **TCP** as the type of traffic processed by the port, and then click **OK**.
7. Verify that **File and Print Sharing**, **Internet Explorer**, and **World Wide Web Publishing Service** are all selected in the **Exceptions** dialog box, and then click the **Advanced** tab.
8. Click **Settings** for **Security Logging**.
9. In the **Log Setting** dialog box, select **Log dropped packets** and **Log successful connections**. Note the path and file name in **Name**.
10. Click **OK** twice to close Windows Firewall.

**VPN1**
VPN1 is a computer running Windows Server 2003 with SP1, Standard Edition, that is providing VPN server services for Internet-based VPN clients.

**Configure the VPN server**

1. Install Windows Server 2003 with SP1, Standard Edition, as a member server named VPN1 in the example.com domain.
2. Open the **Network Connections** folder.
3. For the intranet local area connection, rename the connection to **CorpNet**. For the Internet local area connection, rename the connection to **Internet**. This is shown in the following figure.




4. Configure the TCP/IP protocol for the CorpNet connection with the IP address of 172.16.0.4, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.
5. Configure the TCP/IP protocol for the Internet connection with the IP address of 10.0.0.2 and the subnet mask of 255.255.255.0.

Windows Firewall and Routing and Remote Access cannot run simultaneously on VPN1. If Windows Firewall is turned on, you will need to turn it off; if the Windows Firewall/Internet Connection Sharing (ICS) service has started or is set to automatic before you configure Routing and Remote Access, you must disable it.

**Disable the Windows Firewall/Internet Connection Sharing (ICS) service**

1. Click **Administrative Tools**, and then click **Services**.
2. In the **Services** details pane, right-click **Windows Firewall/Internet Connection Sharing (ICS)** service, and then click **Properties**.
3. If the service **Startup Type** is either **Automatic** or **Manual**, change it to **Disabled**.
4. Click **OK** to close the **Windows Firewall/Internet Connection Sharing (ICS)** dialog box, and then close the **Services**

page.

**Configure Routing and Remote Access**

1. Run the **Routing and Remote Access** snap-in from the **Administrative Tools** folder.
2. In the console tree, right-click **VPN1**, then and click **Configure and Enable Routing and Remote Access**.
3. On the **Welcome to the Routing and Remote Access Server Setup Wizard** page, click **Next**.
4. On the **Configuration** page, **Remote access (dial-up or VPN)** is selected by default. This is shown in the following figure.




5. Click **Next**. On the **Remote Access** page, select **VPN**. This is shown in the following figure.




6. Click **Next**. On the **VPN Connection** page, click the **Internet** interface in **Network** interfaces. This is shown in the following figure.




7. Click **Next**. On the **IP Address Assignment** page, **Automatically** is selected by default. This is shown in the following figure.




8. Click **Next**. On the **Managing Multiple Remote Access Servers** page, click **Yes, set up this server to work with a RADIUS server**. This is shown in the following figure.




9. Click **Next**. On the **RADIUS Server Selection** page, type **172.16.0.2** in **Primary RADIUS server** and the shared secret in **Shared secret**. This is shown in the following figure.




10. Click **Next**. On the **Completing the Routing and Remote Access Server Setup Wizard** page, click **Finish**.

11. You are prompted with a message describing the need to configure the DHCP Relay Agent. This is shown in the following figure.




12. Click **OK**.

13. In the console tree, open **VPN1 (local)**, then **IP Routing**, and then **DHCP Relay Agent**. Right-click **DHCP Relay Agent**, and then click **Properties**.

14. In the **DHCP Relay Agent Properties** dialog box, type **172.16.0.1** in **Server address**. This is shown in the following figure.




15. Click **Add**, and then click **OK**.

**CLIENT1**

CLIENT1 is a computer running Windows XP Professional with SP2 that is acting as a VPN client and gaining remote access to intranet resources across the simulated Internet.

**Configure Client1 as a VPN client for a PPTP connection**

1. Connect CLIENT1 to the intranet network segment.
2. On CLIENT1, install Windows XP Professional with SP2 as a member computer named CLIENT1 of the example.com domain.

| **Note** |
| --- |
| Installing Windows XP Professional with SP2 also installs and automatically turns on Windows Firewall. Leave Windows Firewall turned on for this scenario. You will not need to configure any port or program exceptions. |

3. Add the VPNUser account in the example.com domain to the local Administrators group.
4. Log off and then log on using the VPNUser account in the example.com domain.
5. In Control Panel, open the **Network Connections** folder, obtain properties on the Local Area Network connection, and then obtain properties on the Internet protocol (TCP/IP).
6. Click the **Alternate Configuration** tab, and then click **User configured**.
7. In **IP address**, type **10.0.0.1**. In **Subnet mask**, type **255.255.255.0**. This is shown in the following figure.




8. Click **OK** to save changes to the TCP/IP properties. Click **OK** to save changes to the Local Area Network connection.
9. Shut down the CLIENT1 computer.
10. Disconnect CLIENT1 from the intranet network segment, and connect it to the simulated Internet network segment.
11. Restart CLIENT1 and log on using the VPNUser account.
12. On CLIENT1, in Control Panel, open the **Network Connections** folder.
13. In **Network Tasks**, click **Create a new connection**.
14. On the **Welcome to the New Connection Wizard** page of the New Connection Wizard, click **Next**.
15. On the **Network Connection Type** page, click **Connect to the network at my workplace**. This is shown in the following figure.




16. Click **Next**. On the **Network Connection** page, click **Virtual Private Network connection**. This is shown in the following figure.




17. Click **Next**. On the **Connection Name** page, type **PPTPtoCorpnet** in **Company Name**. This is shown in the following figure.




18. Click **Next**. On the **VPN Server Selection** page, type **10.0.0.2** in **Host name or IP address**. This is shown in the following figure.



19. Click **Next**. On the **Connection Availability** page, click **Next**.

20. On the **Completing the New Connection Wizard** page, click **Finish**. The **Connect PPTPtoCorpnet** dialog box appears. This is shown in the following figure.

21. Click **Properties**, and then click the **Networking** tab.

22. On the **Networking** tab, in **Type of VPN**, click **PPTP VPN**. This is shown in the following figure.




23. Click **OK** to save changes to the **PPTPtoCorpnet** connection. The **Connect PPTPtoCorpnet** dialog box appears.

24. In **User name**, type **example\VPNUser**. In **Password**, type the password you chose for the VPNUser account.

25. Click **Connect**.

26. When the connection is complete, run Internet Explorer.

27. If prompted by the Internet Connection Wizard, configure it for a LAN connection. In **Address**, type **http://IIS1.example.com/iisstart.htm**. You should see a message saying the Web page is under construction.

28. Click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**. You should see the contents of the local drive (drive C) on IIS1.

29. Right-click the **PPTPtoCorpnet** connection, and then click **Disconnect**.

**L2TP/IPsec-based Remote Access VPN Connections**

L2TP/IPsec-based remote access VPN connections require computer certificates on the VPN client and the VPN server. L2TP/IPsec is typically used when there are stronger requirements for security and a public key infrastructure (PKI) is in place to issue computer certificates to VPN clients and servers.

**DC1**

**Configure DC1 for autoenrollment of computer certificates**

1. Open the Active Directory Users and Computers snap-in.
2. In the console tree, double-click **Active Directory Users and Computers**, right-click the **example.com** domain, and then click **Properties**.
3. On the **Group Policy** tab, click **Default Domain Policy**, and then click **Edit**.
4. In the console tree, open **Computer Configuration**, open **Windows Settings**, open **Security Settings**, open **Public Key Policies**, and then open **Automatic Certificate Request Settings**. This is shown in the following figure.




5. Right-click **Automatic Certificate Request Settings**, point to **New**, and then click **Automatic Certificate Request**.
6. On the **Welcome to the Automatic Certificate Request Setup Wizard** page, click **Next**.
7. On the **Certificate Template** page, click **Computer**. This is shown in the following figure.




8. Click **Next**. On the Completing the **Automatic Certificate Request Setup Wizard** page, click **Finish**. The Computer certificate type now appears in the details pane of the Group Policy Object Editor snap-in. This is shown in the following

figure.




9. Type **gpupdate** at a command prompt to update Group Policy on DC1.

**VPN1**

**Update Group Policy on VPN1**

- To immediately update Group Policy and request a computer certificate, type **gpupdate** at a command prompt.

After updating VPN1 with the new certificates you need to stop and restart the IPsec Policy Agent and Routing and Remote Access services.

**Restart IPsec Policy Agent and Routing and Remote Access**

1. Click **Start**, point to **Administrative Tools**, and then click **Services**.
2. In the details pane, point to **IPSEC Services**, point to **Action**, and then click **Restart**.
3. In the details pane, point to **Routing and Remote Access**, point to **Action**, and then click **Restart**.

**CLIENT1**

To obtain a computer certificate on CLIENT1 and then configure an L2TP/IPsec-based remote access VPN connection, perform the following steps.

**Obtain a computer certificate and configure an L2TP/IPsec-based remote access VPN connection**

1. Shut down the CLIENT1 computer.
2. Disconnect CLIENT1 from the simulated Internet network segment, and connect it to the intranet network segment.
3. Restart CLIENT1 and log on using the VPNUser account. The computer and user Group Policy is automatically updated.
4. Shut down CLIENT1.
5. Disconnect CLIENT1 from the intranet network segment, and connect it to the simulated Internet network segment.
6. Restart CLIENT1 and log on using the VPNUser account.
7. On CLIENT1, in Control Panel, open the **Network Connections** folder.
8. In **Network Tasks**, click **Create a new connection**.
9. On the **Welcome to the New Connection Wizard** page of the **New Connection Wizard**, click **Next**.
10. On the **Network Connection Type** page, click **Connect to the network at my workplace**. This is shown in the following figure.




11. Click **Next**. On the **Network Connection** page, click **Virtual Private Network connection**. This is shown in the following figure.




12. Click **Next**. On the **Connection Name** page, type **L2TPtoCorpnet** in **Company Name**. This is shown in the following figure.



13. Click **Next**. On the **Public Network** page, click **Do not dial the initial connection**. This is shown in the following figure.




14. Click **Next**. On the **VPN Server Selection** page, type **10.0.0.2** in **Host name or IP address**. This is shown in the following figure.




15. Click **Next**. On the **Connection Availability** page, click **Next**.

16. On the **Completing the New Connection Wizard** page, click **Finish**. The **Connect L2TPtoCorpnet** dialog box appears. This is shown in the following figure.

17. Click **Properties**, and then click the **Networking** tab.

18. On the **Networking** tab, in **Type of VPN**, click **L2TP IPSec VPN**. This is shown in the following figure.




19. Click **OK** to save changes to the **L2TPtoCorpnet** connection. The **Connect L2TPtoCorpnet** dialog box appears.

20. In **User name**, type **example\VPNUser**. In **Password**, type the password you chose for the VPNUser account.

21. Click **Connect**.

22. When the connection is established, run the Web browser.

23. In **Address**, type **http://IIS1.example.com/iisstart.htm**. You should see a message saying the Web site is under construction.

24. Click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**. You should see the contents of the local drive (drive C) on IIS1.

25. Right-click the **L2TPtoCorpnet** connection, and then click **Disconnect**.

**EAP-TLS-based Remote Access VPN Connections**

EAP-TLS-based remote access VPN connections require a user certificate on the VPN client and a computer certificate on the IAS server. EAP-TLS is for authenticating your VPN connection with the most secure user-level authentication protocol. Locally installed user certificates, enabled in the following steps, make it easier to set up a test lab. In a production environment, it is

recommended that you use smart cards, rather than locally installed user certificates, for EAP-TLS authentication.

**DC1**

**Configure DC1 for autoenrollment of user certificates**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Snap-in**, double-click **Certificate Templates**, click **Close**, and then click **OK**.
4. In the console tree, click **Certificate Templates**. All of the certificate templates will be displayed in the details pane. This is shown in the following figure.




5. In the details pane, click the **User** template.
6. On the **Action** menu, click **Duplicate Template**.
7. In the **Template display name** box, type **VPNUser**.
8. Verify that the **Publish Certificate in Active Directory** check box is selected. This is shown in the following figure.




9. Click the **Security** tab.
10. In the **Group or user names** list, click **Domain Users**.

11. In the **Permissions for Domain Users** list, select the **Read**, **Enroll,** and **Autoenroll** check boxes so that these permissions are allowed. This is shown in the following figure.




12. Click the **Subject Name** tab.

13. Clear the **Include E-mail name in subject name** and **E-mail name** check boxes. Because you did not configure an e-mail name for the VPNUser user account, you must clear these check boxes to allow a user certificate to be issued. This is shown in the following figure.




14. Click **OK**.

15. Open the **Certification Authority** snap-in from the **Administrative Tools** folder.

16. In the console tree, open **Certification Authority**, open **Example CA**, and then open **Certificate Templates**. This is shown in the following figure.




17. On the **Action** menu, point to **New**, and then click **Certificate Template to Issue**.

18. Click **VPNUser**. This is shown in the following figure.




19. Click **OK**.

20. Open the Active Directory Users and Computers snap-in.

21. In the console tree, double-click **Active Directory Users and Computers**, right-click the example.com domain, and then click **Properties**.

22. On the **Group Policy** tab, click **Default Domain Policy**, and then click **Edit**.

23. In the console tree, open **User Configuration**, open **Windows Settings**, open **Security Settings**, and then open **Public Key Policies**. This is shown in the following figure.




24. In the details pane, double-click **Autoenrollment Settings**.

25. Click **Enroll certificates automatically**. Select the **Renew expired certificates, update pending certificates, and remove revoked certificates** check box. Select the **Update certificates that use certificate templates** check box. This is shown in the following figure.

26. Click **OK**.

**IAS1**

**Configure IAS1 with a computer certificate for EAP-TLS authentication**

1. Restart IAS1 to ensure that IAS1 has autoenrolled a computer certificate.

2. Open the Internet Authentication Service snap-in.

3. In the console tree, click **Remote Access Policies**.

4. In the details pane, double-click **VPN remote access to intranet**. The **VPN remote access to intranet Properties** dialog box appears. This is shown in the following figure.




5. Click **Edit Profile**, and then click the **Authentication** tab. This is shown in the following figure.




6. On the **Authentication** tab, click **EAP Methods**. The **Select EAP Providers** dialog box appears. This is shown in the following figure.




7. Click **Add**. The **Add EAP** dialog box appears. This is shown in the following figure.




8. Click **Smart Card or other certificate**, and then click **OK**.

9. Click **Edit**. The **Smart Card or other Certificate Properties** dialog box appears. This is shown in the following figure.

10. The properties of the computer certificate issued to the IAS1 computer are displayed. This step verifies that IAS1 has an acceptable computer certificate installed to perform EAP-TLS authentication. Click **OK**.

11. Click **OK** to save changes to EAP providers. Click **OK** to save changes to the profile settings.

12. When prompted to view help topics, click **No**. Click **OK** to save changes to the remote access policy.

These configuration changes will allow the **VPN remote access to intranet** remote access policy to authorize VPN connections using the EAP-TLS authentication method.

**CLIENT1**

**Obtain a user certificate on CLIENT1, and then configure an EAP-TLS-based remote access VPN connection**

1. Shut down the CLIENT1 computer.

2. Disconnect CLIENT1 from the simulated Internet network segment, and connect it to the intranet network segment.

3. Restart CLIENT1 and log on using the VPNUser account. The computer and user Group Policy is automatically updated.

4. Shut down CLIENT1.
5. Disconnect CLIENT1 from the intranet network segment, and connect it to the simulated Internet network segment.
6. Restart CLIENT1 and log on using the VPNUser account.
7. On CLIENT1, in Control Panel, open the **Network Connections** folder.
8. In **Network Tasks**, click **Create a new connection**.
9. On the **Welcome to the New Connection Wizard** page of the New Connection Wizard, click **Next**.
10. On the **Network Connection Type** page, click **Connect to the network at my workplace**.
11. Click **Next**. On the **Network Connection** page, click **Virtual Private Network** connection.
12. Click **Next**. On the **Connection Name** page, type **EAPTLStoCorpnet** in **Company Name**.
13. Click **Next**. On the **Public Network** page, click **Do not dial the initial connection**.
14. Click **Next**. On the **VPN Server Selection** page, type **10.0.0.2** in **Host name or IP address**.
15. Click **Next**. On the **Connection Availability** page, click **Next**.
16. On the **Completing the New Connection Wizard** page, click **Finish**. The **Connect EAPTLStoCorpnet** dialog box appears. This is shown in the following figure.

17. Click **Properties**, and then click the **Security** tab.
18. On the **Security** tab, click **Advanced**, and then click **Settings**. The **Advanced Security Settings** dialog box appears.
19. In the **Advanced Security Settings** dialog box, click **Use Extensible Authentication Protocol (EAP)**. This is shown in the following figure.




20. Click **Properties**. In the **Smart Card or other Certificate Properties** dialog box, click **Use a certificate on this computer**. This is shown in the following figure.




21. Click **OK** to save changes to the **Smart Card or Other Certificate** dialog box. Click **OK** to save changes to the **Advanced Security Settings**. Click **OK** to save changes to the **Security** tab. The connection is immediately initiated using the installed user certificate. The first time you try to connect, it may take several attempts to successfully make a connection.

22. When the connection is successful, run the Web browser.

23. In **Address**, type **http://IIS1.example.com/iisstart.htm**. You should see a message saying the Web site is under construction.

24. Click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**. You should see the contents of the local drive (drive C) on IIS1.

25. Right-click the **EAPTLStoCorpnet** connection, and then click **Disconnect**.

**Summary**

This guide described in detail the steps required to configure secure VPN remote access using PPTP, L2TP/IPsec, and EAP-TLS

in a test lab with five computers simulating an organization intranet and the Internet.

**See Also**
**Other Resources**
Microsoft VPN Web site
Windows Server 2003 Web site

# Step-by-Step Guide for Setting Up a PPTP-based Site-to-Site VPN Connection in a Test Lab

Updated: April 15, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides an example with detailed information about how you can use five computers in a test lab environment to configure and test a Point-to-Point Tunneling Protocol (PPTP)-based site-to-site VPN connection using any 32-bit version of the Microsoft® Windows Server™ 2003 operating system with Service Pack 1 (SP1) as well as the Microsoft Windows® XP Professional operating system with Service Pack 2 (SP2).

You can use this example deployment to learn about Windows Server 2003 with SP1 site-to-site VPN functionality before you deploy a site-to-site VPN connection in a production environment. This test lab configuration simulates a deployment of a PPTP-based site-to-site VPN connection between the Seattle and New York offices of an organization.

| **Note** |
| --- |
| The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the functionality. This configuration is designed to reflect neither best practices nor a recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed to work only on a separate test lab network. |

**Setting Up the Test Lab**

The infrastructure for a PPTP-based site-to-site VPN deployment test lab network consists of five stand-alone computers performing the roles shown in Table 1. Each computer is part of a workgroup. None of the computers are joined to a domain.

In this test lab scenario, Windows Firewall is installed and turned-on automatically on the client computers running Windows XP Professional with SP2. You will configure a Windows Firewall exception on CLIENT1, allowing communication between the two client computers. On the three computers with Windows Server 2003 with SP1, Standard Edition, Windows Firewall is automatically installed, but it is not turned-on by default. On these computers, Windows Firewall will remain turned-off. In addition, the Windows Firewall/Internet Connection Sharing (ICS) service should be disabled on each of these computers.

**Table 1 Test Lab Computer Setup**

| Computer | Role |
| --- | --- |
| CLIENT1 running Windows XP Professional with SP2 | • Client computer |
| ROUTER1 running Windows Server 2003 with SP1, Standard Edition | • VPN Server<br>• Answering router |
| INTERNET running Windows Server 2003 with SP1, Standard Edition | • Internet router |
| ROUTER2 running Windows Server 2003 with SP1, Standard Edition | • VPN server<br>• Calling router |
| CLIENT2 running Windows XP Professional with SP2 | • Client computer |

In addition to these five computers, the test lab also contains four hubs (or Layer 2 switches):

- A hub that connects the Seattle office (CLIENT1) to the answering router (ROUTER1).
- A hub that connects the New York office (CLIENT2) to the calling router (ROUTER2).

- A hub that connects the answering router (ROUTER1) to the Internet router (INTERNET).
- A hub that connects the calling router (ROUTER2) to the Internet router (INTERNET).

| **Note** |
| --- |
| Because there are only two computers on each subnet, the hubs can be replaced by Ethernet crossover cables. |

The configuration of this test lab is shown in the following figure:




The IP addresses for the test lab configuration are shown in Tables 2, 3, and 4.

**Table 2 IP Addresses for the Seattle Office Subnet**

| Computer/Interface | IP Addresses |
| --- | --- |
| CLIENT1 | 172.16.4.3 |
| ROUTER1 (to the Seattle intranet) | 172.16.4.1 |

**Table 3 IP Addresses for the Internet Subnets**

| Computer/Interface | IP Addresses |
| --- | --- |
| ROUTER1 (to the Internet) | 10.1.0.2 |
| INTERNET (to ROUTER1, the answering router) | 10.1.0.1 |
| ROUTER2 (to the Internet) | 10.2.0.2 |
| INTERNET (to ROUTER2, the calling router) | 10.2.0.1 |

**Table 4 IP Addresses for the New York Office Subnet**

| Computer/Interface | IP Addresses |
| --- | --- |
| ROUTER2 (To the New York intranet) | 172.16.56.1 |
| CLIENT2 | 172.16.56.3 |

Configure your test lab by performing the following tasks:

1. Configure the client computers in the Seattle and New York offices.
2. Configure the computers performing as the answering and calling routers.
3. Configure the computer performing as the Internet router.

**Configuration for CLIENT1**

CLIENT1 is a client on the Seattle office subnet, running Windows XP Professional with SP2.

### Configure TCP/IP properties

1. Open **Network Connections**, right-click the network connection you want to configure, and then click **Properties**.
2. On the **General** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.

3. Click **Use the following IP address**, type **172.16.4.3** for the **IP address**, type **255.255.255.0** for the **Subnet mask**, and then type **172.16.4.1** for the **Default gateway**.

Installing Windows XP Professional with SP2 turns on Windows Firewall by default. You will need to configure port exceptions on the firewall allowing for communication between CLIENT1 and CLIENT2.

**Configure Windows Firewall on CLIENT1**

1. Click **Start**, point to **Control Panel**, and then click **Security Center**.
2. Click **Windows Firewall**, and then in the **Windows Firewall** dialog box, click the **Advanced** tab.
3. Click **Settings** for **ICMP**, and then click **Allow incoming echo request**.
4. Click **OK** twice to close **Windows Firewall**.

**Configuration for CLIENT2**

CLIENT2 is a client on the New York office subnet, running Windows XP Professional with SP2. Installing Windows XP Professional with SP2 turns on Windows Firewall by default. However, because CLIENT2's only role is as a calling computer, Windows Firewall does not need to be enabled with any exceptions. Leave the default Windows Firewall settings on CLIENT2.

**Configure TCP/IP properties**

1. Open **Network Connections**, right-click the network connection you want to configure, and then click **Properties**.
2. On the **General** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
3. Click **Use the following IP address**, and then type **172.16.56.3** for the **IP address**, type **255.255.255.0** for the **Subnet mask**, and then type **172.16.56.1** for the **Default gateway**.

**Configuration for the Answering and Calling Routers**

This section describes the setup for the routers in the test lab. For information about configuring routing and remote access for the answering router (ROUTER1) and the calling router (ROUTER2), see "Configuring a PPTP-based Site-to-Site VPN Connection" in this guide.

**ROUTER1**

ROUTER1 is a computer on the Seattle office subnet, running Windows Server 2003 with SP1, Standard Edition. ROUTER1 is acting as the answering router.

**Configure TCP/IP properties**

1. Open **Network Connections**, right-click the network connection you want to configure, and then click **Properties**.
2. On the **General** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
3. Configure the IP address and subnet mask with the following values:
   a. On the **To the Internet** interface, type **10.1.0.2** for the **IP address**, type **255.255.0.0** for the **Subnet mask**, and then type **10.1.0.1** for the **Default gateway**.
   b. On the **To the Seattle intranet** interface, type **172.16.4.1** for the **IP address**, type **255.255.255.0** for the **Subnet mask**. Leave the **Default gateway** clear.

Windows Firewall and Routing and Remote Access cannot run simultaneously on VPN1. If Windows Firewall is turned on, you will need to disable it. If the Windows Firewall/Internet Connection Sharing (ICS) service has started or is set to automatic before configuring Routing and Remote Access, you must disable it.

**Disable the Windows Firewall/Internet Connection Sharing (ICS) service**

1. Click **Administrative Tools**, and then click **Services**.
2. In the **Services** details pane, right-click **Windows Firewall/Internet Connection Sharing (ICS)** service, and then click **Properties**.
3. If the service **Startup Type** is either **Automatic** or **Manual**, change it to **Disabled**.
4. Click **OK** to close the **Windows Firewall/Internet Connection Sharing (ICS)** dialog box, and then close the **Services**

page.

#### ROUTER2

ROUTER2 is a computer on the New York office subnet, running Windows Server 2003 with SP1, Standard Edition. ROUTER2 is acting as the calling router.

**Configure TCP/IP properties**

1. Open **Network Connections**, right-click the network connection you want to configure, and then click **Properties**.
2. On the **General** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
3. Configure the IP address and subnet mask with the following values:
    a. On the **To the Internet** interface, type **10.2.0.2** for the **IP address**, type **255.255.0.0** for the **Subnet mask**, and then type **10.2.0.1** for the **Default gateway**.
    b. On the **To the Seattle intranet** interface, type **172.16.56.1** for the **IP address**, type **255.255.255.0** for the **Subnet mask**. Leave the **Default gateway** empty.

As with ROUTER1, you must turn off Windows Firewall on ROUTER2 and disable the Windows Firewall/ICS service.

**Disable the Windows Firewall/Internet Connection Sharing (ICS) service**

1. Click **Administrative Tools**, and then click **Services**.
2. In the **Services** details pane, right-click **Windows Firewall/Internet Connection Sharing (ICS)** service, and then click **Properties**.
3. If the service **Startup Type** is either **Automatic** or **Manual**, change it to **Disabled**.
4. Click **OK** to close the **Windows Firewall/Internet Connection Sharing (ICS)** dialog box, and then close the **Services** page.

#### Configuration for the Internet Router

INTERNET is a computer running Windows Server 2003 with SP1, Standard Edition.

**Configure TCP/IP properties**

1. Open **Network Connections**, right-click the network connection you want to configure, and then click **Properties**.
2. On the **General** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
3. Configure the IP address and subnet mask with the following values:
    a. On the **To Router1** interface, type **10.1.0.1** for the **IP address**, and then type **255.255.0.0** for the **Subnet mask**.
    b. On the **To Router2** interface, type **10.2.0.1** for the **IP address**, and then type **255.255.0.0** for the **Subnet mask**.
4. Click **Administrative Tools**, and then click **Routing and Remote Access** to open the Routing and Remote Access Microsoft Management Console (MMC) snap-in.
5. In the Routing and Remote Access snap-in, right-click **INTERNET (local)** in the console tree, and then click **Configure and Enable Routing and Remote Access**.
6. The Routing and Remote Access Server Setup Wizard opens. Click **Next**.
7. On the **Configuration** page, select **Custom configuration**, as shown in the following figure.




8. Click **Next**. On the **Custom Configuration** page, select **LAN routing**, as shown in the following figure.




9. Click **Next**. On the **Completing the Routing and Remote Access Server Setup** page, click **Finish**, and then click **Yes** to start the service.

Make sure that Windows Firewall is turned off and that the Windows Firewall/Internet Connection Sharing (ICS) service is disabled.

**Disable the Windows Firewall/Internet Connection Sharing (ICS) service**

1. Click **Administrative Tools**, and then click **Services**.
2. In the **Services** details pane, right-click **Windows Firewall/Internet Connection Sharing (ICS)** service, and then click **Properties**.
3. If the service **Startup Type** is either **Automatic** or **Manual**, change it to **Disabled**.
4. Click **OK** to close the **Windows Firewall/Internet Connection Sharing (ICS)** dialog box, and then close the **Services** page.

**Verify the routing infrastructure**

1. On **ROUTER1**, ping the IP address 10.2.0.2. This should be successful.
2. On **CLIENT2**, ping the IP address 172.16.4.3. This should be unsuccessful because CLIENT1 cannot be reached by

CLIENT2 across the simulated Internet until the site-to-site VPN connection is created.

**Configuring a PPTP-based Site-to-Site VPN Connection**
To create a PPTP-based VPN connection, perform the following tasks:

1. Configure VPN on the answering router.
2. Configure the demand-dial interface on the answering router.
3. Configure VPN on the calling router.
4. Configure the demand-dial interface on the calling router.
5. Confirm remote access policy configuration on the answering and calling routers.
6. Initiate the VPN connection.
7. Test the VPN connection.

**Configure VPN on the answering router (ROUTER1)**

1. On **ROUTER1**, click **Administrative Tools**, and then click **Routing and Remote Access**.
2. In Routing and Remote Access, right-click **ROUTER1 (local)** in the console tree, and then click **Configure and Enable Routing and Remote Access**.
3. The **Routing and Remote Access Server Setup Wizard** appears. Click **Next**.
4. On the **Configuration** page, select **Remote access (dial-up or VPN)**, as shown in the following figure.




5. Click **Next**. On the **Remote Access** page, select **VPN**, as shown in the following figure.




6. Click **Next**. On the **VPN Connection** page, select **To the Internet**, and then verify that the **Enable security on the selected interface by setting up static packet filters** check box is selected, as shown in the following figure.




7. Click **Next**. On the **IP Address Assignment** page, select **From a specified range of addresses**, as shown in the following figure.




8. Click **Next**. On the **Address Range Assignment** page, click **New**, as shown in the following figure.




9. In the **New Address Range** dialog box, do the following:

   a. In the **Start IP address** box, type **172.16.100.1**.

   b. In the **End IP address** box, type **172.16.100.2**.

   c. In the **Number of Addresses** box, accept the displayed value of **2**, as shown in the following figure.




10. Click **OK**. On the **Address Range Assignment** page, click **Next**.

11. On the **Managing Multiple Remote Access Servers** page, select **No, use Routing and Remote Access to authenticate connection requests**, as shown in the following figure.




12. Click **Next**. On the **Completing the Routing and Remote Access Server Setup** page, click **Finish**.

13. Click **OK** to close the message box prompting you to configure the DHCP Relay Agent. For this scenario the DHCP Relay Agent will not be configured.

**Configure the demand-dial interface on the answering router (ROUTER1)**

1. In the Routing and Remote Access snap-in, expand **ROUTER1**, and then right-click **Network Interfaces**.

2. Click **New Demand-dial Interface** to open the Demand-Dial Interface Wizard, and then click **Next**.

3. On the **Interface Name** page, type **VPN_NewYork**, as shown in the following figure. The interface name must match the user account name of the calling router.




4. Click **Next**. On the **Connection Type** page, select **Connect using virtual private networking (VPN)**, as shown in the following figure.




5. Click **Next**. On the **VPN Type** page, select **Point-to-Point Tunneling Protocol (PPTP)**, as shown in the following figure.




6. Click **Next**. On the **Destination Address** page, type **10.2.0.2** in the **Host name or IP address** box, as shown in the following figure.



7. Click **Next**. On the **Protocols and Security** page, do the following:

   a. Select **Route IP packets on this interface**.

   b. Select **Add a user account so a remote router can dial in**, as shown in the following figure.




8. Click **Next**. On the **Static Routes for Remote Networks** page, click **Add**, as shown in the following figure.




9. In the **Static Route** dialog box, do the following:

   a. In the **Destination** box, type **172.16.56.0**.

   b. In the **Network Mask** box, type **255.255.255.0**.

   c. In the **Metric** box, accept the displayed value **1**, as shown in the following figure.




10. Click **OK**. On the **Static Routes for Remote Networks** page, click **Next**.

11. On the **Dial In Credentials** page, type a password for the **VPN_NewYork** user account, and then retype the password in the **Confirm password** box. The **User name** box is automatically populated with the value **VPN_NewYork**.




12. Click **Next**. On the **Dial Out Credentials** page, do the following:

    a. In the **User name** box, type **VPN_Seattle**.

    b. In the **Domain** box, type **ROUTER2**.

    c. In the **Password** box, type a password for the **VPN_Seattle** user account.

    d. In the **Confirm password** box, retype the password for the **VPN_Seattle** user account, as shown in the following figure.




13. Click **Next**. On the last **Demand-Dial Interface Wizard** page, click **Finish**.
14. Click **OK** to close the message box prompting you to configure the DHCP Relay Agent. For this scenario the DHCP Relay Agent will not be configured.

**Configure VPN on the calling router (ROUTER2)**

1. On **ROUTER2**, click **Administrative Tools**, and then click **Routing and Remote Access**.
2. In Routing and Remote Access, right-click **ROUTER2 (local)** in the console tree, and then click **Configure and Enable Routing and Remote Access**.
3. The Routing and Remote Access Server Setup Wizard appears. Click **Next**.
4. On the **Configuration** page, select **Remote access (dial-up or VPN)**, and then click **Next**.
5. On the **Remote Access** page, select **VPN**, and then click **Next**.
6. On the **VPN Connection** page, select **To the Internet**, verify that the **Enable security on the selected interface by setting up static packet filters** check box is selected, and then click **Next**.
7. On the **IP Address Assignment** page, select **From a specified range of addresses**, click **Next**, and then on the **Address Range Assignment** page, click **New**.
8. In the **New Address Range** dialog box, do the following:
   a. In the **Start IP address** box: type **172.56.200.1**.
   b. In the **End IP address** box, type **172.56.200.2**.
   c. In the **Number of Addresses** box, accept the displayed value of **2**, and then click **OK**.
9. On the **Address Range Assignment** page, click **Next**.
10. On the **Managing Multiple Remote Access Servers** page, select **No, use Routing and Remote Access to authenticate connection requests**. Click **Next**.
11. On the **Completing the Routing and Remote Access Server Setup** page, click **Finish**.
12. Click **OK** to close the message box prompting you to configure the DHCP Relay Agent.

**Configure the demand-dial interface on the calling router (ROUTER2)**

1. In the Routing and Remote Access snap-in, expand **ROUTER2 (local)**, and then right-click **Network Interfaces**.
2. Click **New Demand-dial Interface** to open the Demand-dial Interface Wizard. To complete the Demand-Dial Interface

Wizard, click **Next**.

3. On the **Interface Name** page, type **VPN_Seattle**. The interface name must match the user account name of the answering router. Click **Next**.
4. On the **Connection Type** page, select **Connect using virtual private networking (VPN)**. Click **Next**.
5. On the **VPN Type** page, select **Point-to-Point Tunneling Protocol (PPTP)**. Click **Next**.
6. On the **Destination Address** page, type **10.1.0.2**, and then click **Next**.
7. On the **Protocols and Security** page, do the following:
   a. Select **Route IP packets on this interface**.
   b. Select **Add a user account so a remote router can dial in**, and then click **Next**.
8. On the **Static Routes for Remote Networks** page, click **Add**.
9. In the **Static Route** dialog box, do the following:
   a. In the **Destination** box, type **172.16.4.0**.
   b. In the **Network Mask** box, type **255.255.255.0**.
   c. In the **Metric** box, accept the displayed value **1**, and then click **OK**.
10. On the **Static Routes for Remote Networks** page, click **Next**.
11. On the **Dial In Credentials** page, type the password for the **VPN_Seattle** user account, and then retype the password in the **Confirm password** box. The **User name** box is pre-populated with the value **VPN_Seattle**. Click **Next**.
12. On the **Dial Out Credentials** page, do the following:
    a. In the **User name** box, type **VPN_NewYork**.
    b. In the **Domain** box, type **ROUTER1**.
    c. In the **Password** box, type the password for the **VPN_NewYork** user account created on **ROUTER1**.
    d. In the **Confirm password** box, retype the password for the **VPN_NewYork** user account, and then click **Next**.
13. On the last **Demand-Dial Interface Wizard** page, click **Finish**.

**Confirm the remote access policy configuration on the answering and calling routers**

1. On **ROUTER2**, in Routing and Remote Access, click **Remote Access Policies**.
2. In the details pane, right-click **Connections to Microsoft Routing and Remote Access server**, and then click **Properties**.
3. On the **Settings** tab, select **Grant remote access permission**, and then click **OK** to save changes.
4. Repeat steps 1 through 3 on **ROUTER1**.

Initiate the VPN connection by performing the following steps on ROUTER2.

**Initiate the VPN connection**

1. On **ROUTER2**, in the console tree of the Routing and Remote Access snap-in, click **Network Interfaces**.
2. In the details pane, right-click **VPN_Seattle**, and then click **Connect**.
3. Confirm that the connection state of **VPN_Seattle** is connected.

**Test the VPN connection**

1. On **CLIENT2**, at the command prompt, type **ping 172.16.4.3**.

   This is the IP address for CLIENT1. Pinging CLIENT1 from CLIENT2 will test whether the Seattle subnet is now reachable.

2. To confirm that the packets crossed the VPN connection, at the command prompt, type **tracert 172.16.4.3**. Note that you must use the IP address of CLIENT1, rather than its computer name, because a DNS server is not configured in this test lab scenario.

   Results that are similar to the following indicate that the connection is working.

```
Tracing route to 172.16.4.3 over a maximum of 30 hops:
     1    <1 ms    <1 ms    <1 ms  [172.16.56.1]
     2     1 ms    <1 ms    <1 ms  [172.56.200.2]
     3     1 ms     1 ms     1 ms  [172.16.4.3]
Trace complete.
```

| **Note** |
| --- |
| 172.16.56.1 is the IP address of the ROUTER2 interface that connects to the New York intranet. 172.56.200.2 is the IP address that ROUTER2 assigned to ROUTER1. The presence of this IP address in the Tracert output indicates that packets are moving across the site-to-site VPN connection. 172.16.4.3 is the IP address of CLIENT1. |

**Summary**

The preceding tasks describe how to configure a PPTP-based site-to-site VPN connection in a test lab with five computers. These five computers simulate two remote sites using VPN to connect over the Internet.

**See Also**

**Other Resources**

Virtual Private Networks for Windows Server 2003
Windows Server 2003 Deployment Guide
Windows Server 2003 Web site

# System Services for the Windows Server 2003 Family and Windows XP Operating Systems

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In This White Paper**


Alerter (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Application Layer Gateway Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Application Management (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

ASP .NET State Service

Automatic Updates (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Background Intelligent Transfer Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Certificate Service

Client Service for NetWare (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Clipbook (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Cluster Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

COM+ Event System

COM+ System Application

Computer Browser (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Cryptographic Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

DHCP Client (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

DHCP Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Distributed File System (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Distributed Link Tracking Client (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Distributed Link Tracking Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Distributed Transaction Coordinator (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

DNS Client (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

DNS Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Error Reporting Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Event Log (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Fast User Switching Compatibility

Fax (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

File Replication (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

File Server for Macintosh (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

FTP Publishing Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Help and Support (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

HTTP SSL

Human Interface Device Access (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

IAS Jet Database Access (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

IIS Admin Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

IMAPI CD-Burning COM Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Indexing Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Infrared Monitor (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Internet Authentication Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Internet Connection Firewall (ICF) / Internet Connection Sharing (ICS) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Intersite Messaging (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

IPv6 Helper Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

IPSEC Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Kerberos Key Distribution Center (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

License Logging (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Logical Disk Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Logical Disk Manager Administrative Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Message Queuing (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Message Queuing Down Level Client Support (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Message Queuing Triggers (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Messenger (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Microsoft POP3 Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Microsoft Software Shadow Copy Provider (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

MSSQL$UDDI

MSSQLServerADHelper

Net Logon (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

NetMeeting Remote Desktop Sharing (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Network Connections (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Network DDE (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Network DDE DSDM (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Network Location Awareness (NLA) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Network News Transfer Protocol (NNTP) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

NT LM Security Support Provider (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Performance Logs and Alerts (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Plug and Play (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Portable Media Serial Number

Print Server for Macintosh (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Print Spooler (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Protected Storage (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Access Auto Connection Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Access Connection Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Administration Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Desktop Help Session Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Installation (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Procedure Call (RPC) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Procedure Call (RPC) Locator (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Registry (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Server Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Remote Storage Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Removable Storage (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Resultant Set of Policy Provider (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Routing and Remote Access (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

SAP Agent (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Secondary Logon (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Security Accounts Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Shell Hardware Detection (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Simple Mail Transfer Protocol (SMTP) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Simple TCP/IP Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Single Instance Storage Groveler (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Smart Card (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

SNMP Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

SNMP Trap Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Special Administration Console Helper (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

SQLAgent$UDDI

System Event Notification (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Task Scheduler (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

TCP/IP NetBIOS Helper (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

TCP/IP Print Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Telephony (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Telnet (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Terminal Server Licensing (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Terminal Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Terminal Services Session Directory (System Services for the Windows Server 2003 Family and Windows XP Operating

Systems)

Themes (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Trivial FTP Daemon (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Uninterruptible Power Supply (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Upload Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Virtual Disk Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Volume Shadow Copy (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

WebClient (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Web Element Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Windows Audio (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Windows Image Acquisition (WIA) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Windows Installer (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Windows Internet Name Service (WINS) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Windows Management Instrumentation (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Windows Management Instrumentation Driver Extensions (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Windows Media Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Windows System Resource Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Windows Time (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

WinHTTP Web Proxy Auto-Discovery Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Wireless Configuration (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

WMI Performance Adapter (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Workstation (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

World Wide Web Publishing Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

# Alerter (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Alerter

**Executable Name:** svchost.exe -k LocalService

**Log On As:** LocalService

**Description:** The Alerter service notifies selected users and computers of administrative alerts.

Use the Alerter service to send alert messages to specified users that are connected on your network.

Alert messages warn users about security, access, and user session problems. Alert messages are sent as messages from a server to a user's computer. The Messenger service must be running on the user's computer for the user to receive alert messages.

If this service is stopped or disabled, applications that use the NetAlertRaise or NetAlertRaiseEx APIs will be unable to notify a user or computer (by a message box from the messenger service) that the administrative alert took place.

*See also* **Messenger**.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Workstation

**The following system components depend on this service:** None

**Port Numbers used:**

TCP: 2869, dynamic

UDP: 1900

# Application Layer Gateway Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** ALG

**Executable Name:** ALG.EXE

**Log On As:** LocalService

**Description:** This subcomponent of the Internet Connection Sharing (ICS) / Internet Connection Firewall (ICF) service provides support for independent software vendors (ISVs) to write protocol plug-ins that allow their proprietary network protocols to pass through the firewall and work behind ICS. Application Layer Gateway plug-ins have the power to open ports and change data (such as ports and IP addresses) embedded in packets. File Transfer Protocol (FTP) is the only network protocol that has a plug-in shipping with Windows Server 2003 Standard Edition and Windows Server 2003 Enterprise Edition.

The service listens for outgoing FTP traffic from an FTP client. It extracts the port that the FTP client is expecting to receive data from and creates an appropriate dynamic port mapping for the FTP data channel

If this service is disabled, the Internet Connection Firewall and Internet Connection Sharing service will not start. The ALG service will start when set to manual if the Internet Connection Firewall and Internet Connection Sharing service is started.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

None

**The following system components depend on this service:**

Internet Connection Firewall (ICF) / Internet Connection Sharing (ICS)

**IP Port Numbers used:**

TCP: 21, dynamically allocated

# Application Management (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** AppMgmt

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Provides software installation services, such as Assign, Publish, and Remove. This service processes requests to enumerate, install, and remove applications deployed via a corporate network. When you click Add in Add/Remove Programs on a computer joined to a domain, the program calls this service to retrieve the list of your deployed applications. The service is also called when you use Add/Remove Programs to install or remove an application, and in cases when a component (such as the shell or Component Object Model (COM)), makes an install request for an application to handle a file extension, COM class, or ProgID that is not present on the computer. The service is started by the first call made to it—it does not terminate once started.

If this service is disabled, deployed application information will not be retrieved nor will this information appear with Add/Remove Programs, Add New Programs. No programs are available on the network will be displayed in the Add programs from your network dialog box. Stopping this service is not possible once started. This service must be disabled to prevent it from starting if not required.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:** None

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# ASP .NET State Service

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Aspnet_state

**Executable Name:** aspnet_state.exe

**Log On As:** Network Service

**Description:** ASP State Service provides support for out-of-process session states for ASP. ASP has a concept of session state – a listing of values associated with the client session is accessible from ASP pages through the Session property. There are three options provided to store session data: In process; SQL database; and out-of-process. The ASP State Service stores session data out-of-process. The service communicates with ASP using sockets.

If this service is stopped or disabled, out of process requests will not be processed.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed Through:** Add/Remove Windows Components, Application Server, ASP.NET

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:** None

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 42424

# Automatic Updates (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Wuauserv

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Enables the download of updates from Microsoft's Windows Update Web site.

This service keeps your computer up-to-date automatically with the latest updates, drivers and enhancements from Microsoft. You no longer have to search for critical updates and information; Windows Update delivers them directly to your computer if configured. Windows recognizes when you are online and uses your Internet connection to search for downloads from the Windows Update page on the Microsoft Web site. A message appears on your desktop to let you know when new updates are available.

The Automatic Update feature can be turned off through the Systems setting in the Control Panel or by right mouse clicking My Computer and selecting Properties.

You can also use the MMC Group Policy Object Editor snap in administrative template to configure an intranet server to host updates from the Microsoft Update Web sites. This setting lets you specify a server on your network to function as an internal update service. The Automatic Updates client will search this service for updates that apply to the computers on your network.

Stopping or disabling this service has no effect on the rest of the operating system. The user will have to update their computer manually from the Windows Update Web site located at http://v4.windowsupdate.microsoft.com/en/default.asp.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Installed By:** Default

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:** None

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 80

# Background Intelligent Transfer Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** BITS

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Use Background Intelligent Transfer Service (BITS) to transfer files asynchronously between a client and an HTTP server. BITS is a background file transfer mechanism and queue manager. Requests to the BITS service are submitted and the files are transferred uses idle network bandwidth so that other network related activities, such as browsing, are not affected.

BITS suspends the transfer if a connection is lost or if the user logs off. BITS persists transfer information while the user is logged off, across network disconnects, and during machine restarts. When the user logs on, BITS resumes the user's transfer job.

BITS uses a queue to manage file transfers. You can prioritize transfer jobs within the queue and specify whether the files are transferred in the foreground or background. Background transfers are optimal in that BITS uses idle network bandwidth to transfer the files and will increase or decrease the rate (throttle) at which files are transferred based on the amount of idle network bandwidth available. If a network application begins to consume more bandwidth, BITS decreases its transfer rate to preserve the user's interactive experience.

BITS provides one foreground and three background priority levels that you can use to prioritize transfer jobs. Higher priority jobs preempt lower priority jobs; whereas, jobs at the same priority level share transfer time (round-robin scheduling prevents a large job from blocking the transfer queue). Lower priority jobs do not receive transfer time until all higher priority jobs are complete or in an error state.

The BITS service can be demand started when set to manual and the first job is submitted. When all outstanding jobs are completed, the BITS service will stop.

When stopped: If the service is stopped, features such as Automatic Update, and MSN Explorer will be unable to automatically download programs and other information until the BITS service has started.

When disabled: If this service is disabled, any services that explicitly depend on this service will fail to transfer files if they do not have a fail safe mechanism to transfer files directly through other methods such as Internet Explorer.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Installed By:** Default

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 80

# Certificate Service

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** CertSvc

**Executable Name:** CERTSVC.EXE

**Log On As:** LocalSystem

**Description:** Part of the core operating system that enables a business to act as if its own certificate authority (CA), and issue and manage digital certificates for applications such as Secure/Multipurpose Internet Mail Extensions (S/MIME), Secure Sockets Layer (SSL), Encrypting File System (EFS), IP Security (IPSEC), and smartcard log on.

If this service is stopped or disabled, certificate requests will not be accepted and the Certificate Revocation Lists (CRLs) and delta CRLs will not be published. If this service is paused or stopped long enough for CRLs to expire, validation of existing certificates will fail.

For more information about S/MIME and SSL, see "Certificate Services" in the Windows Server Distributed Services Guide.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition.

**Installed through:** Add/Remove Windows Components, Certificate Service

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:** None

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 1117

# Client Service for NetWare (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** NWCWorkstation

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalService

**Description:** Provides access to files and directories as well as resources on NetWare networks.

With Client Service for Netware, you can access file and print resources on Netware Servers that are running Novell Directory Services (NDS) or bindery security (Netware versions 3.x or 4.x) from your computer.

The Client Service for Netware does not support the IP protocol and therefore cannot be used to interoperate with NetWare 5.x in an IP-only environment. To do this, you must load the IPX protocol onto the Netware 5.x server, or use a redirector that is compatible with Netware Core Protocol (NCP) and supports native IP.

If this service is stopped or disabled, access to file and print resources on NetWare networks will no longer function unless the Novell Client for NetWare is installed.

**Available on:** Windows XP Home Edition; Windows XP Professional Edition; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add-on via Network Connections, Client

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:** None

**The following system components depend on this service:** None

**IP Port Numbers used:**

UDP: 213 (IPX Over IP)

# Clipbook (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** ClipSrv

**Executable Name:** clipsrv.exe

**Log On As:** LocalSystem

**Description:** Enables the Clipbook Viewer to create and share "pages" of data to be viewed by remote computers.

This service depends on the (NetDDE) Network Dynamic Data Exchange service to create the actual file shares that other computers can connect to, while the Clipbook application and service allow users to create the pages of data to share.

This service is disabled by default. When this service is disabled and Clipbrd.exe is launched, Clipbrd.exe can still be used to view the local Clipboard (where data is stored when a user highlights text and then goes to the **Edit** menu and selects **Copy**, or types **Ctrl+C**).

**Available on:** Windows XP Home Edition; Windows XP Professional Edition; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Installed by:** Default

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Network DDE

Network DDE DSDM

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# Cluster Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** ClusSvc

**Executable Name:** Clussvc.exe

**Log On As:** Domain account specified when cluster service is configured

**Description:** Server clusters provide high availability and scalability for mission-critical applications such as databases, messaging systems, and file and print services. If one of the nodes in a cluster becomes unavailable either due to planned downtime for maintenance or unplanned downtime due to failure of a node, the operating system or an application, another node takes over to provide the service to the end-user—a process known as failover. When failover occurs, users who are accessing the cluster service continue to access the service, and are unaware that it is now being provided from a different server (node).

Supports for up to 8-node server clusters in both Enterprise Server and Datacenter Server, however, a cluster cannot be made up of nodes running both Windows Server 2003 Enterprise Server and Windows Datacenter Server, they must all be running either Enterprise server or Datacenter server. A server cluster can be operated with some nodes running Windows 2000 and others running Windows Server as part of a rolling upgrade scenario.

**Server clusters can be set up as one of three different configurations:**

- Single node server clusters can be configured with, or without, external cluster storage devices. For single node clusters without an external cluster storage device, the local disk is configured as the cluster storage device. Single node configurations can be used to develop cluster aware applications or can be used in production to provide the local health monitoring and restart capabilities to applications.

- Single quorum device server clusters have two or more nodes and are configured so that every node is attached to one or more cluster storage devices. The cluster configuration data is stored on a single cluster storage device known as the quorum disk.

- Majority node set server clusters have two or more nodes where the nodes may or may not be attached to one or more cluster storage devices. The cluster configuration data is stored on multiple disks across the cluster and the Cluster service makes sure that this data is kept consistent across the different disks.

If this service is stopped or disabled, the cluster service itself and any applications or services hosted by the cluster service will be stopped.

**Available on:** Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Not started until the first node of a cluster is configured.

**This service depends on the following system components:**

Network Connections

Remote Procedure Call

Windows Time

Network Cluster Driver

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 3343

UDP: 1204, 1094, 2036, 2008, 3343

# COM+ Event System

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** EventSystem

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Provides automatic distribution of events to subscribing COM (Component Object Model) components. COM+ Events extend the COM+ programming model to support late-bound events or method calls between the publisher or subscriber and the event system. Instead of repeatedly polling the server, the event system notifies interested parties as information becomes available.

COM+ Events handle most of the event semantics for the publisher and subscriber. Publishers offer to publish event types, and subscribers request event types from specific publishers. Subscriptions are maintained outside both the publisher and subscriber and are retrieved when needed. This simplifies the programming model. The subscriber does not need to contain the logic for building subscriptions—building a subscription is as easy as building a COM component. The life cycle of the subscription is separate from that of either the publisher or the subscriber. Subscriptions can be built prior to either the subscriber or publisher being made active.

If the service is disabled, you will be prompted to also stop the SENS (System Event Notification) service as well since SENS is dependent upon this service and logon and log off notifications will not occur and other inbox applications, such as Volume Snapshot service, will not work correctly.

**Available on:** Windows XP, Home Edition; Windows XP, Professional Edition; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:**

System Event Notification

Window Internet Name Service (WINS)

DHCP Server

COM+ System Application

**IP Port Numbers used:** None

# COM+ System Application

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** COMSysApp

**Executable Name:** dllhost.exe /Processid:{02D4B3F1-FD88-11D1-960D-00805FC79235}

**Log On As:** LocalSystem

**Description:** The COM+ system application hosts COM+ services and manages COM+ application configuration and tracking.

COM+ applications will not work if the COM+ system application is disabled, also you will not be able to administer anything to do with COM+ or OLE registrations on the system.

If you disable the MS DTC (Microsoft Distributed Transaction) service, any COM+ system applications will fail because this service needs local transactions to maintain its state.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call

COM+ Event System

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# Computer Browser (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Browser

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Maintains an up-to-date list of computers on your network, and supplies the list to programs that request it. The Computer Browser service is used by Windows-based computers that need to view network domains and resources. Computers designated as browsers maintain browse lists, which contain all shared resources used on the network. Earlier versions of Windows applications, such as My Network Places, the NET VIEW command, and the Microsoft Windows NT Explorer, all require browsing capability.

There are several different roles a computer may perform in a browsing environment. Under some conditions (such as failure or shutdown of a computer designated for a specific browser role) browsers—or potential browsers—may change to a different role of operation. Windows NT assigns the following special roles to computers running the Computer Browser service:

- **Domain Master Browser** - Used only in domain environments. By default, the DC (domain controller) for a domain operates in this role. The domain master browser collects and maintains the master browse list of available servers for its domain, in addition to any names for other domains and workgroups used in the network. It also distributes and synchronizes the master browse list for master browsers on other subnets that have computers belonging to the same domain.

- **Master Browser** - Collects and maintains the list of available network servers in its subnet. The master browser fully replicates its listed information with the domain master browser to obtain a complete browse list for the network, and distributes it to backup browsers located on the same subnet.

- **Backup Browser** - The backup browser receives a copy of the browse list from the master browser for its subnet, and distributes it to other computers upon request.

- **Potential Browser** - Capable of becoming a backup browser when instructed to by the subnet's master browser, the potential browser operates similarly to a non-browser under normal conditions.

- **Nonbrowser** - A nonbrowser is configured so it cannot become a browser, and it does not maintain a browse list. It can operate as a browse client, requesting browse lists from other computers operating as browsers on the same subnet. When the Computer Browser service is turned off there is no mechanism to discover other computers to populate the My Network Places, and so on.

If this service is stopped or disabled, computer browsing features will not be available.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Server

Workstation

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 139

UDP: 137(browsing requests of NETBIOS over TCP/IP), 138 (browsing datagram responses of NetBios over TCP/IP)

# Cryptographic Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** CryptSvc

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Provides key management services for your computer. The Cryptographic Service is comprised of three management services: Catalog Database Service, Protected Root Service, and Key Service.

- Catalog Database Service - This part of the service is responsible for adding, removing, and looking up catalog files. Catalog files are used to bulk sign all the files in the operating system. So, this service is used by Windows File Protection (WFP), Driver Signing, and setup to verify signed files.

- Protected Root Service - This part of the service is responsible for adding and removing Trusted Root Certification Authority certificates for the current user. It displays a service message box with the certificate's name and thumbprint. If the user clicks OK, the certificate is added to or removed from the current user's list of trusted root authorities. Only Local System has write access to this list. If this part of the service is turned off, the current user will not be able to add or remove Trusted Root Certificate Authority certificates.

- Key Service - This part of the service allows administrators to enroll for certificates on behalf of the local machine account. To that end, it provides several pieces of functionality required for enrollment -- enumeration of available Certification Authorities, enumeration of available machine templates, the ability to create and submit a certificate request in the local machine context, and so on. Only administrators may enroll on behalf of the local machine account.

If this service is disabled, administrators will not be able to enroll for machine certificates and auto enrollment will not be able to automatically acquire the default set of machine certificates. Also, WFP and driver signing will not be able to check the signatures of the operating system files.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# DHCP Client (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Dhcp

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** NetworkService

**Description:** Dynamic Host Configuration Protocol (DHCP) Client manages network configuration by registering and updating IP addresses and Domain Name Server (DNS) names for your computer. You do not have to manually change the IP settings when a client, such as a roaming user, wanders throughout the network. The client is automatically given a new IP address regardless of the subnet it reconnects to—as long as a DHCP server is accessible from each of those subnets. There is no need to manually configure settings for DNS or Windows Internet Name Service (WINS). The DHCP server can give these settings to the client, as long as the DHCP server has been configured to issue such information. To enable this option on the client, simply select the Obtain DNS Server Address Automatically option button. There are no conflicts caused by duplicate IP addresses.

If this service is stopped or disabled, you will not be able to obtain an IP address from a DHCP Server and you will have to configure a static IP address on the local machine. If this service is set to manual it will not demand start.

See also **DHCP Server**, **Manage Your Server**, **Configure Your Server**.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

AFD Networking Support Environment

TCP/IP Protocol Driver

IPSEC Driver

**The following system components depend on this service:**

WinHTTP Web Proxy Auto-Discovery Service

**IP Port Numbers used:**

TCP: 68

UDP: 67, 68, 1029

# DHCP Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** DHCPServer

**Executable Name:** tcpsvcs.exe

**Log On As:** LocalSystem

**Description:** The DHCP Server service allocates IP addresses and allows the advanced configuration of network settings such as DNS servers, WINS servers to DHCP clients automatically.

DHCP uses a client-server model. The network administrator establishes one or more DHCP servers that maintain TCP/IP configuration information and provide it to clients. The server database includes the following:

- Valid configuration parameters for all clients on the network.
- Valid IP addresses maintained in a pool for assignment to clients, plus reserved addresses for manual assignment.
- Duration of a lease offered by the server. The lease defines the length of time for which the assigned IP address can be used.

Dynamic Host Configuration Protocol (DHCP) is an IP standard designed to reduce the complexity of administering address configurations by using a server computer to centrally manage IP addresses and other related configuration details used on your network. The Windows Server family provides the DHCP service, which enables the server computer to perform as a DHCP server and configure DHCP-enabled client computers on your network as described in the current DHCP draft standard, RFC 2131.

DHCP includes the Multicast Address Dynamic Client Assignment Protocol (MADCAP) which is used to perform multicast address allocation. When registered clients are dynamically assigned IP addresses through MADCAP, they can participate efficiently in the data stream process, such as for real-time video or audio network transmissions.

With a DHCP server installed and configured on your network, DHCP-enabled clients can obtain their IP address and related configuration parameters dynamically each time they start and join your network. DHCP servers provide this configuration in the form of an address-lease offer to requesting clients. With a DHCP server installed and configured on your network, DHCP-enabled clients can obtain their IP address and related configuration parameters dynamically each time they start and join your network. DHCP servers provide this configuration in the form of an address-lease offer to requesting clients.

If the DHCP Server service is stopped or disabled off, DHCP clients will not automatically receive IP addresses or network settings. See also **DHCP Client**, **Configure Your Server**, **Manage Your Server**.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Configure Your Server, Manage Your Server, Add/Remove Windows Components

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Event Log

Remote Procedure Call

Security Accounts Manager

COM+ Event System

TCP/IP Protocol Driver

IPSEC Driver

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 7, 9, 13, 17, 19, 135 (DHCP Manager) 515, 2535, 2856

UDP: 67, 68

# Distributed File System (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** DFS

**Executable Name:** dfssvc.exe

**Log On As:** LocalSystem

**Description:** The Distributed File System (DFS) service manages logical volumes distributed across a local or wide area network. DFS is a single hierarchical file system, the contents of which are distributed across the network. DFS provides a logical tree structure for file system resources that may be anywhere on the network. Since the DFS tree is a single point of reference, regardless of the actual location of the underlying resources, you can access network resources based on a meaningful representation of the data. You no longer need to know and specify the actual physical location of files distributed across the network.

If the DFS service is stopped or disabled, users will be unable to access network data through the logical namespace and any attempt to configure a DFS root using the MMC DFS snap in will result in an RPC error. In order to for users to access the data, users will need to know the names of all the servers and shares in the namespace, and access each of these targets independently.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Server

Workstation

Remote Procedure Call (RPC)

Security Account Manager

MUP

DFS Driver

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 137 (name lookup), 139

UDP: None

# Distributed Link Tracking Client (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Trkwrks

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Maintains links between the NTFS file system files within a computer or across computers in a network domain. The DLT Client service ensures that shortcuts and OLE (Object Linking and Embedding) links continue to work after the target file is renamed or moved. When a shortcut to a file on an NTFS v5 volume is created, distributed link tracking stamps a unique object identifier (ID) into the target file, known as the link source. Information about the object ID is also stored within the referring file, known as the link client.

Distributed link tracking can use this object ID to locate the link source file in any combination of the following scenarios that occur within a Windows domain:

- The link source file is renamed.
- The link source file is moved to another folder on the same volume or to a different volume on the same computer.
- The link source file is moved from one NTFS volume to another within the same domain. (The NTFS volumes must be on computers running Windows 2000 or later. The NTFS volumes cannot be on removable media.)
- The computer containing the link source file is renamed.
- The shared network folder containing the link source file is renamed.
- The volume containing the link source file is moved to another computer within the same domain.

Distributed link tracking also attempts to maintain links even when they do not occur within a domain, such as cross-domain, within a workgroup, or on a single computer that is not connected to a network. Links can always be maintained in these scenarios when a link source is moved within a computer, or when the network shared folder on the link source computer is changed. Typically, links can be maintained when the link source is moved to another computer, though this form of tracking is less reliable over time.

Distributed link tracking uses different services for client and server:

- The DLT Client service runs on all Windows 2000-based or later computers. In non-networked computers, the DLT Client service performs all activities related to link tracking.
- The DLT Server service runs on Windows 2000 Server or later domain controllers. The service maintains information relating to the movement of link source files. Because of this service and the information it maintains, links within a domain are more reliable than those outside a domain. For computers that run in a domain, the DLT Client service takes advantage of this information by communicating with the DLT Server service.

| **Note** |
| --- |
| The DLT Client service monitors activity on NTFS volumes and stores maintenance information in a file called Tracking.log, which is located in a hidden folder called System Volume Information at the root of each volume. This folder is protected by permissions that allow only the system to have access to it. The folder is also used by other Windows services, such as Indexing Service.<br><br>If the DLT Client service is disabled, you won’t be able to track links. Likewise, users on other computers won’t be able to track links for documents on your computer. |

*See also* **Distributed Link Tracking Server**.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# Distributed Link Tracking Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** TrkSrv

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The distributed link tracking server runs on each domain controller in a domain. The service accepts notifications of file and volume moves from the tracking service on a computer and allows the distributed link tracking client to query a link source's current location.

This server service maintains information in the DC about volumes and files which have been moved. The information on moves cannot grow above a certain size and it is automatically removed if it becomes unnecessary.

The distributed link-tracking services are available only on NTFS, and are only available for link sources on NTFS 5.0 and later volumes. Thus if a link source is moved to a FAT volume, or if a link source is moved to a computer running Windows NT 4.0, the tracking information is lost. Additionally, if a link source is moved even between NTFS 5.0 volumes, but the computer performing the move is running an earlier version of Windows NT or Windows 95/98/Me, the link tracking information is lost. When the link tracking information is lost, no harm is done to the link-source file itself, it is simply not trackable by the distributed link-tracking services.

Links to files on removable media are not maintained. Also, the tracking service does not recognize a new NTFS volume until the system is rebooted. A new volume might become available because of repartitioning, reformatting a FAT volume to NTFS, or connecting a new external drive. The service accepts notifications of file and volume moves from the tracking service on a computer and allows the distributed link tracking client to query a link source's current location.

If the DLT server service is stopped or disabled, links maintained by the DLT Client service will be unreliable. The "NtfsDisableDomainLinkTracking" policy should be enabled in the File system policy group to prevent DLT clients from repeatedly trying to reach the disabled service.

*See also* **Distributed Link Tracking Client**.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# Distributed Transaction Coordinator (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** MSDTC

**Executable Name:** msdtc.exe

**Log On As:** NetworkService

**Description:** Coordinates transactions that are distributed across multiple computer systems and/or resource managers, such as databases, message queues, file systems, or other transaction-protected resource managers. The Distributed Transaction Coordinator is necessary if transactional components are going to be configured through Component Services (COM+). It is also required for transactional queues in Message Queuing (MSMQ) and Microsoft SQL Server operations that span multiple systems.

Stopped or disabling this service prevents the transactions described above from occurring.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

Security Accounts Manager

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# DNS Client (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Dnscache

**Executable Name:** svchost.exe -k NetworkService

**Log On As:** NetworkService

**Description:** The Domain Name System (DNS) client service resolves and caches DNS names. The DNS client service must be running on every computer that will perform DNS name resolution. The ability to resolve DNS names is crucial for locating domain controllers in Active Directory domains. The DNS client service is also critical for locating devices identified using DNS name resolution.

The DNS Client service running on Windows implements the following features:

- System-wide caching

  Resource records (RRs) from query responses are added to the client cache as applications query DNS servers. This information is then cached for a set Time to Live (TTL) and can be used again to answer subsequent queries.

- RFC-compliant negative caching support

  In addition to caching positive query responses from DNS servers (which contain resource record information in the answered reply), the DNS Client service also caches negative query responses. A negative response results when a resource record for the queried name does not exist.

  Negative caching prevents the repeating of additional queries for names that do not exist, which can adversely affect client system performance. Any query information negatively cached is kept for a shorter period of time than is used for positive query responses; by default, no more than 5 minutes. This avoids continued negative caching of stale information if the records later become available.

  Negative caching is a new DNS standard specification defined in RFC 2308. For more information, refer to this RFC. For more information on obtaining RFCs, see TCP/IP RFCs.

- Avoidance of unresponsive DNS servers

  The DNS Client service uses a server search list, ordered by preference. This list includes all preferred and alternate DNS servers configured for each of the active network connections on the system. Windows rearranges these lists based on the following criteria:

  - Preferred DNS servers are given first priority.
  - If no preferred DNS servers are available, then alternate DNS servers are used.
  - Unresponsive servers are removed temporarily from these lists.

If the DNS client service is stopped or disabled on your computer, you may not be able to locate the domain controllers of the Active Directory domains and will no longer be able to locate the devices identified using DNS names.

*See also* **DNS Server**.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

TCP/IP Protocol Driver

IPSEC Driver

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 53

# DNS Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** DNS

**Executable Name:** dns.exe

**Log On As:** LocalSystem

**Description:** Enables DNS name resolution by answering queries and update requests for DNS names. Presence of the DNS servers is crucial for locating devices identified using DNS names and locating domain controllers in Active Directory. If there is no authoritative DNS for a particular portion of the namespace, then locating devices in that portion of the namespace will fail. Not having the authoritative DNS server for the DNS namespace used to resolve Active Directory domains results in an inability to locate the domain controllers for such domain.

If this service is stopped or disabled, DNS name or device resolution will not be reliable.

*See also* **DNS Client**, **Configure Your Server**, **Manage Your Server**.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components, Networking Services; Configure Your Server, Manage Your Server.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

AFD Networking Support Environment

TCP/IP Protocol Driver

IPSEC Driver

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 53, 139 (DNS ADMINISTRATION) 1024, 1043, 1067, 1068

UDP: 53, 1046, 1045

# Error Reporting Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** ERSvc

**Executable Name:** svchost.exe –k netsvcs

**Log On As:** LocalSystem

**Description:** The Error Reporting Service provides an infrastructure for collecting, storing and reporting kernel mode, operating system and application faults to Microsoft. Error reporting helps Microsoft track and address errors. You can configure error reporting to send Microsoft specific error information and to generate reports for operating system errors, Windows component errors or program errors. An operating system error causes the computer to display a blue screen with error codes. A program or component error causes the program or component to stop working.

If you have an Internet connection, you can report these errors directly to Microsoft. You can configure error reporting to respond to program errors in one of two ways: as soon as an error occurs, the error reporting dialog box can prompt any user to send the error to Microsoft, or the next time an administrator logs on, the error reporting dialog box can prompt the administrator to send the error report to Microsoft.

Windows treats operating system errors and unplanned shutdowns differently from the way it does program errors. When operating system errors or unplanned shutdowns occur, Windows writes the error information to a log file. The next time an administrator logs on, the error reporting dialog box prompts them to report the error.

When you send an error report to Microsoft over the Internet, you provide technical information that programming groups at Microsoft use to enhance future versions of the product. This data is used for quality control purposes only and is not used for tracking individual users or installations for any marketing purpose. If information is available to help you solve the problem, Windows displays an Error Reporting dialog box with a link to that information.

Alternatively, if your organization has configured Group Policy, administrators in your information technology (IT) department can use Corporate Error Reporting to collect and report to Microsoft only those errors that they think are important. Administrators configure workstations and servers for Corporate Error Reporting by enabling the Report Errors policy setting and configuring the Corporate upload file path to the local file server where the Corporate Error Reporting tool is installed. When errors occur, they are automatically redirected to this file server. Administrators can then review the error information, identify the important data, and submit it to Microsoft using the Corporate Error Reporting tool. You can download the Corporate Error Reporting tool from the Office XP Resource Kit Web site (http://www.microsoft.com/office/).

If error reporting is stopped or disabled, users will not be given the option to report errors. If Display Error Notification is enabled, users will still get a message indicating that a problem occurred, but will not have the option to report this information to Microsoft or a local network share.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 80

# Event Log (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Eventlog

**Executable Name:** services.exe

**Log On As:** LocalSystem

**Description:** This service logs event messages issued by programs and the Windows operating system. Event Log reports contain information that can be useful in diagnosing problems. Reports are viewed in Event Viewer. The Event Log service writes events sent by applications, services, and the operating system to log files. The events contain diagnostic information in addition to errors specific to the source application, service, or component. The logs can be viewed programmatically through the Event Log APIs or through the Event Viewer in an MMC (Microsoft Management Console) snap-in.

By default, a computer running Windows 2000 Server, Windows XP and Windows Server 2003, records events in three kinds of logs:

**Application log**

- The application log contains events logged by applications or programs. For example, a database program might record a file error in the application log. Program developers decide which events to log.

**Security log**

- The security log records events such as valid and invalid logon attempts, as well as events related to resource use such as creating, opening, or deleting files or other objects. For example, if logon auditing is enabled, attempts to log on to the system are recorded in the security log.

**System log**

- The system log contains events logged by Windows system components. For example, the failure of a driver or other system component to load during startup is recorded in the system log. The event types logged by system components are predetermined by the server.

A computer running a Windows Server 2003 family operating system configured as a domain controller records events in two additional logs:

**Directory service log**

- The directory service log contains events logged by the Windows Active Directory service. For example, connection problems between the server and the global catalog are recorded in the directory service log.

**File Replication service log**

- The File Replication service log contains events logged by the Windows File Replication service. For example, file replication failures and events that occur while domain controllers are being updated with information about system volume changes are recorded in the file replication log.

A computer running Windows configured as a Domain Name System (DNS) server records events in an additional log:

**DNS server log**

- The DNS server log contains events logged by the Windows DNS service.

If the event log is disabled, you will be unable to track events, which will significantly reduce the ability to successfully diagnose system problems. In addition security events will not be audited and you will not be able to view previous event logs using the

MMC event viewer snap in.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:** None

**The following system components depend on this service:**

DHCP Server

File Replication

Network News Transfer Protocol (NNTP)

Simple Mail Transfer Protocol (SMTP)

SNMP Service

SNMP Trap Service

Windows Internet Name Services (WINS)

Windows Management Instrumentation

**IP Port Numbers used:**

TCP: 139

# Fast User Switching Compatibility

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** FastUserSwitchingCompatibility

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Provides management services for applications that require assistance in a multiple user environment. This service is responsible for activating the application compatibility shims for Fast User Switching. Fast User Switching allows multiple users to log on and run applications on a single machine. Some applications (usually non-Windows 2000 logo applications) are not able to run in this environment. We have created shims that allow these applications to run in a multiple user environment. This service only starts if there is more than one user logged on to the machine. It does not activate when there is only one user on the machine. Fast User Switching only works on Personal and Professional when joined to a workgroup. Thus, machines joined to a domain are not affected by this service.

If the service is stopped or disabled, the shims* used to make applications work in a multiple user environment will not occur. You will not have application compatibility for known Fast User Switching issues.

*Shims are application compatibility fixes which are used to address application compatibility issues reported to Microsoft.

**Available on:** Windows XP Home, Windows XP Professional.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call

Terminal Services

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# Fax (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Fax

**Executable Name:** fxssvc.exe

**Log On As:** LocalSystem

**Description:** The Fax service, a TAPI-compliant service, provides fax capabilities from your computer. The Fax service allows users to send and receive faxes from their desktop applications using either a local fax device or a shared network fax device. The service offers the following features:

- Send and receive faxes
- Track and monitor fax activity
- Routing of inbound faxes
- Server and device configuration management
- Archiving of sent faxes

If stopped, the Fax service will automatically start when there are fax jobs pending in the queue and will stop when no longer needed to process fax requests. If the print spooler or telephony service is disabled, the Fax service will not start successfully.

*See also* **Print Spooler**, **Telephony**.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Installed by:** Add/Remove Windows Components

**Startup type:** Automatic

**Service status:** Stopped

**This service depends on the following system components:**

Plug and Play

Print Spooler

Remote Procedure Call

Telephony

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# File Replication (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** NtFrs

**Executable Name:** ntfrs.exe

**Log On As:** LocalSystem

**Description:** Enables files to be automatically copied and maintained simultaneously on multiple servers. File Replication Service (FRS) is the automatic file replication service in Windows 2000 and the Windows Server 2003 family and its function is to replicate the volume (Sysvol) on all domain controllers. In addition, it can be configured to replicate files among alternate targets associated with the fault-tolerant Distributed File System (DFS).

This service that provides multimaster file replication for designated directory trees between designated servers running Windows 2000 or Windows Server 2003. The designated directory trees must be on disk partitions formatted with the version of NTFS used within the Windows Server 2003 family. FRS is used by the Distributed File System (DFS) to automatically synchronize content between assigned replicas and by Active Directory to automatically synchronize content of the system volume information across domain controllers.

If this service is stopped or disabled, file replication will not occur and server data will not be synchronized. Stopping the File Replication service on a domain controller may seriously impair a domain controllers ability to function.

For more information about FRS, see "File Replication Service" in the Windows Server Distributed Services Guide.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Event Log

Remote Procedure Call

COM+ Event System

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: dynamically allocated

UDP: 1024 - 65535

# File Server for Macintosh (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** MacFile

**Executable Name:** sfmsvc.exe

**Log On As:** LocalSystem

**Description:** This service enables Macintosh users to store and access files on a NTFS volume on a Windows Server. Encrypted files on an NTFS volume are not accessible from Macintosh clients.

If this service is stopped or disabled, Macintosh clients will not be able to view NTFS shares.

*See also* **Print Server for Macintosh**.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition.

**Installed through:** Add/Remove Windows Components

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Workstation

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 548

UDP: 548

# FTP Publishing Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** MSFtpsvc

**Executable Name:** inetinfo.exe

**Log On As:** NetworkService

**Description:** Provides File Transfer Protocol (FTP) connectivity and administration through the Internet Information Service (IIS) snap-in. Features include bandwidth throttling, security accounts, and extensible logging. This includes the new FTP User Isolation feature, which allows users to access only their files on an FTP site. In addition, there is improved international support.

If the service is stopped or disabled, the computer cannot function as an FTP server.

**Available on:** Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components, Application Server, Internet Information Services (IIS).

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

IIS Admin Service

Remote Procedure Call

Security Accounts Manager

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 20 (data), 21

# Help and Support (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Helpsvc

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Provides Help and Support application and framework functionality.

The help service supports the Help and Support Center application and enables communication between the client application and the help data it accesses. This includes access to stores and services such as the taxonomy database that contains metadata and information about the help topics, the support automation framework that enables data collection for registered support providers, user history and preference information, and the search engine manager. When a user is interacting with the Help and Support Center features like search, index, table of contents, and so on, the service allows for data transaction that supports all of these features.

If this service is set to manual, launching Help and Support in the user interface will start this service. Disabling this service renders the Help and Support Center essentially unusable and the user will receive a message Windows cannot open Help and Support because a system service is not running. The user can access some topics at a high level that might have been cached on the local system, but most of the help and support application features cannot function without the help service enabled. *.HLP and *.CHM files located in the Windows\Help directory can still be viewed, however.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed by:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 80

# HTTP SSL

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** HTTPFilter

**Executable Name:** lsass.exe

**Log On As:** LocalSystem

**Description:** Enables Internet Information Services (IIS) to perform Secure Sockets Layer (SSL) functions. SSL is a proposed open standard for establishing a secure communications channel to prevent the interception of critical information, such as credit card numbers. Primarily, it enables secure electronic financial transactions on the World Wide Web, although it is designed to work on other Internet services as well. If IIS is not installed, the HTTP SSL service will depend on the HTTP driver to perform SSL functions.

If this service is stopped or disabled, If this service is stopped, the IIS service would stop . Any HTTP requests aimed for IIS would not be handled.. This service is called SSL for HTTP.SYS on Windows XP Professional.

For more information about SSL, see "Secure Web Communications", in the Windows Server 2003 Resource Kit Internet Information Services Resource Guide.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

IIS Admin Service

Remote Procedure Call

Security Accounts Manager

HTTP

**The following system components depend on this service:**

World Wide Web Publishing Service

**IP Port Numbers used:**

TCP: 43, 445

UDP: 443

# Human Interface Device Access (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Hidserv

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** This service provides generic access to specific functions contained within controls collections on HID (Human Interface Devices). It enables the use of predefined hot buttons on keyboards, remote controls and other multimedia devices.

If the service is stopped or disable, the buttons on USB keyboards will not function (i.e., back, forward, volume up, down, previous track, next track), nor will the volume buttons on USB speakers.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# IAS Jet Database Access (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** IASjet

**Executable Name:** svchost.exe –k iasjet

**Log On As:** LocalSystem

**Description:** The IAS Jet Database Access service uses the Remote Authentication Dial-in User Service (RADIUS) protocol to provide authentication, authorization, and accounting services. With IAS, you can centrally manage the authentication, authorization, and accounting of users. You can also use IAS to authenticate users in databases on your domain controller running Windows NT 4.0, Windows 2000, or Windows Server operating systems. IAS works equally well in homogeneous and heterogeneous networks running Windows Server operating systems.

Internet Authentication Service (IAS) can be used as a RADIUS proxy to provide the routing of RADIUS messages between RADIUS clients (access servers) and RADIUS servers that perform user authentication, authorization, and accounting for the connection attempt. When used as a RADIUS proxy, IAS is a central switching or routing point through which RADIUS access and accounting messages flow. IAS records information in an accounting log about the messages that are forwarded.

A RADIUS authentication, authorization, and accounting infrastructure consists of the following components:

**Access clients**

An access client is a device that requires some level of access to a larger network. Examples of access clients are dial-up or virtual private network (VPN) clients, wireless clients, or LAN clients connected to a switch.

**RADIUS clients (access servers)**

An access server is a device that provides some level of access to a larger network. An access server using a RADIUS infrastructure is also a RADIUS client, sending connection requests and accounting messages to a RADIUS server. Examples of access servers are:

- **Network access servers** (NASs) that provide remote access connectivity to an organization network or the Internet. An example is a Windows 2000 computer running the Routing and Remote Access service and providing either traditional dial-up or virtual private network (VPN) remote access services to an organization's intranet.

- **Wireless access points** that provide physical layer access to an organization's network, using wireless-based transmission and reception technologies.

- **Switches** that provide physical layer access to an organization's network, using traditional LAN technologies such as Ethernet.

**RADIUS proxies**

A RADIUS proxy is a device that forwards or routes RADIUS connection requests and accounting messages between RADIUS clients (and RADIUS proxies) and RADIUS servers (or RADIUS proxies). The RADIUS proxy uses information within the RADIUS message, such as the User-Name or Called-Station-ID RADIUS attributes, to route the RADIUS message to the appropriate RADIUS server.

A RADIUS proxy can be used as a forwarding point for RADIUS messages when the authentication, authorization, and accounting must occur at multiple RADIUS servers in different organizations.

**RADIUS servers**

A RADIUS server is a device that receives and processes connection requests or accounting messages sent by RADIUS clients or RADIUS proxies. In the case of connection requests, the RADIUS server processes the list of RADIUS attributes in the connection request. Based on a set of rules and the information in the user account database, the RADIUS server either authenticates and authorizes the connection and sends back an Access-Accept message or sends back an Access-Reject message. The Access-Accept message can contain connection restrictions that are implemented by the access server for the duration of the connection.

**User account database**

The user account database is the list of user accounts and their properties that can be checked by a RADIUS server to verify authentication credentials and user account properties containing authorization and connection parameter information.

The user account databases that IAS can use are the local Security Accounts Manager (SAM), a Microsoft Windows NT 4.0 domain, or the Active Directory service. For Active Directory, IAS can provide authentication and authorization for user or computer accounts in the domain in which the IAS server is a member, two-way trusted domains, and trusted forests with domain controllers running Windows Standard Server, Windows Enterprise Server, and Windows Datacenter Server.

If the user accounts for authentication reside in a different type of database, IAS can be configured as a RADIUS proxy to forward the authentication request to a RADIUS server that does have access to the user account database. Different databases for Active Directory include untrusted forests, untrusted domains, or one-way trusted domains.

If this service is disabled, both the Routing and Remote Access Service (RRAS) and Internet Authentication Service (IAS) will fail to start. You will also not be able to administer RRAS or IAS either locally or remotely. If this service is stopped or disabled, users will not be able to connect to the network in cases where network access requires user authentication (e.g., remote access [dial-up, VPN] or wireless LAN [802.1X], Ethernet 802.1X LAN access).

**Available on:** Windows Server 2003, Enterprise Edition (64-bit only) , Windows Server 2003, Datacenter) Edition (64-bit only).

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# IIS Admin Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** IISADMIN

**Executable Name:** inetinfo.exe

**Log On As:** LocalSystem

**Description:** IIS Admin service manages the IIS metabase and updates the Microsoft Windows operating system registry for the WWW service, FTP service, SMTP service, and NNTP service. The metabase is a data store that holds IIS configuration data. IIS Admin service exposes the metabase to other applications, including the core components of IIS, applications that are built on IIS, and third-party applications that are independent of IIS, such as management or monitoring tools.

IIS Admin service runs as Iisadmin.dll and is hosted by Inetinfo.exe.

This service allows administration of Microsoft Internet Information Server components such as FTP, Applications Pools, Web sites, Web service extensions and both NNTP and SMTP virtual servers.

If this service is stopped or disabled, you will not be able to run Web, FTP, Network News Transfer Protocol (NNTP), or Simple Mail Transfer Protocol (SMTP) sites, nor configure IIS and the MMC IIS snap-in will appear blank.

On Windows 2000, the IIS Admin service and related services were installed by default. On the Windows Server 2003 family, IIS must be installed using Add/Remove Windows Components or Configure Your Server, Manage Your Server.

*See also* **World Wide Web Publishing Service**, **FTP Publishing Service**, **HTTP SSL**, **Simple Mail Transfer Protocol (SMTP)**, **Network News Transfer Protocol (NNTP) services**.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default on Windows Server 2003, Web Edition. On Windows Server 2003, Standard Edition, Enterprise Edition and Datacenter Edition, the IIS Admin service must be added via Add\Remove Windows Components, Application Server, Internet Information Services (IIS), Configure Your Server or Manage Your Server.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

Security Accounts Manager

**The following system components depend on this service:**

FTP Publishing Service

Simple Mail Transfer Protocol (SMTP)

World Wide Web Publishing Service

HTTP SSL

Network News Transfer Protocol (NNTP)

Microsoft POP3 Service

**IP Port Numbers used:**

TCP: 80

UDP: 80

# IMAPI CD-Burning COM Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** ImapiService

**Executable Name:** imapi.exe

**Log On As:** LocalSystem

**Description:** This service manages burning CDs through the IMAPI (Image Mastering Applications Programming Interface) COM interface and performs CD-R writes when requested by the user through Windows explorer, WMP (Windows Media Player) or 3rd party applications which use this API.

The Image Mastering API (IMAPI) allows an application to stage and burn a simple audio or data image to CD-R and CD-RW devices. The API supports the following formats: Redbook audio and data discs with both Joliet and ISO 9660. The architecture allows for future expansion of the supported format set.

Stopping or disabling this service may prevent burning a CD-R. If a 3rd party CD-RW application is used, turning off this service may not impact recording to CD-R. If this service is started after log on, you must log off your computer in order to write data to CD-R media via your CD-R using Windows Explorer.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

T**his service depends on the following system components:** None

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# Indexing Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** cisvc

**Executable Name:** cisvc.exe

**Log On As:** LocalSystem

**Description:** Indexes contents and properties of files on local and remote computers and provides rapid access to files through a flexible querying language. It also enables quick searching of documents on local and remote computers and a search index for content shared on the Web. The Indexing Service builds indexes of all textual information in files and documents. Once the initial index build is complete, the Indexing Service maintains its indexes whenever a file is created, modified, or deleted.

Initial indexing can be resource-intensive. On Windows, the Indexing Service will index only when this service is started and the computer is idle. Using the MMC Indexing Service snap in, allows the configuration of resource allocation which are optimized for query or indexing usage patterns.

If this service is stopped or disabled, all search functionality is provided by traversing the folder hierarchy and scanning each file for the requested string and search responses will be slower. If the MMC Indexing Service snap in is used to create a new catalog, the catalog will remain offline until this service is started.

**Available on:** Windows XP Home, Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed by:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped, will start when using Configure Your Server/Manage Your Server to configure the File Server role.

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 80

# Infrared Monitor (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Irmon

**Executable Name:** svchost.exe –k netsvcs

**Log On As:** LocalSystem

**Description:** Supports infrared devices installed on the computer and detects other devices that are in range.

This service is installed by default only if an infrared device is detected during operating system installation. This service is not available on Windows Web, Enterprise or Datacenter Server.

**Available on:** Windows XP Home, Windows XP Professional, Windows Sever 2003 Standard Edition.

**Installed through:** Default operating system installation (only if an infrared port is detected during installation)

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

Terminal Services

IrDa Protocol

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# Internet Authentication Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** IAS

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Performs centralized authentication, authorization, auditing, and accounting of users who are connecting to a network (LAN or remote) using VPN (Virtual Private Network) equipment, RAS (Remote Access Service) equipment, or 802.1X Wireless and Ethernet/Switch Access Points. Internet Authentication Service implements the IETF standard RADIUS (Remote Authentication Dial-In User Service) protocol, which enables use of heterogeneous network access equipment.

If IAS is stopped or disabled, authentication requests will fail over to a backup IAS Server, if available. If there is no other backup IAS Servers are available, users will not be able to connect.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition.

**Installed through:** Add/Remove Programs, Add/Remove Windows Components, Networking Services

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: None

UDP: 1645, 1646, 1812, 1813

# Internet Connection Firewall (ICF) / Internet Connection Sharing (ICS) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SharedAccess

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Provides NAT (network address translation), addressing and name resolution services for all computers on your home or small-office network through a dial-up or broadband connection. When Internet Connection Sharing is enabled, your computer becomes an "Internet gateway" on the network, enabling other client computers to share one connection to the Internet, share files, and use the same printers.

If this service is stopped or disabled, features such as Internet connection sharing, name resolution, and addressing will not be available to clients on the network. Therefore clients on the home or small office network may not be able to access the Internet and their IP addresses will expire, resulting in some clients using APIPA (Automatic Private IP Addressing) for peer-to-peer networking connectivity.

*This service was called Internet Connection Sharing on Windows 2000.

**Available on:** Windows XP Home, Windows XP Professional, Windows Server 2003, Standard Edition, Windows Server 2003, Enterprise Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Application Layer Gateway Service

Network Connections

Remote Access Connection Manager

Remote Access Auto Connection Manager

Remote Procedure Call

Telephony

Plug and Play

Network Location Awareness

AFD Networking Support Environment

TCP/IP Protocol Driver

IPSEC Driver

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 53, 67

UDP: None

# Intersite Messaging (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** IsmServ

**Executable Name:** ismserv.exe

**Log On As:** LocalSystem

**Description:** Intersite Messaging enables messages to be exchanged between computers running Windows Server sites. This service is used for mail-based replication between sites. Active Directory includes support for replication between sites by using SMTP over IP transport. SMTP support is provided by the SMTP service, which is a component of IIS. The set of transports used for communication between sites must be extensible; therefore, each transport is defined in a separate add-in dynamic link library (DLL). These add-in DLLs are loaded into the ISM service, which runs on all domain controllers that are candidates for performing communication between sites. The ISM service directs send requests and receive requests to the appropriate transport add-in DLLs, which then route the messages to the ISM service on the destination computer.

If this service is stopped or disabled, Intersite messaging replication will not work nor will site routing information be calculated for other services.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Security Accounts Manager

Remote Procedure Call

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 1863

# IPv6 Helper Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** 6to4

**Executable Name:** svchost.exe –k netsvcs

**Log On As:** LocalSystem

**Description:** Offers IPv6 connectivity over an existing IPv4 network. 6to4 allows IPv6-enabled sites and hosts to communicate by using IPv6 over an IPv4 infrastructure (for example, the Internet). IPv6 sites and hosts can use their 6to4 address prefix and the Internet to communicate without obtaining an IPv6 global address prefix from an Internet service provider and connecting to the 6bone (the IPv6-enabled portion of the Internet).

6to4 is a tunneling technique that is described in RFC 3056. 6to4 hosts do not require any manual configuration and create 6to4 addresses by using standard auto configuration. 6to4 uses the global address prefix of 2002:WWXX:YYZZ::/48, where WWXX:YYZZ is the colon-hexadecimal representation of a public IPv4 address (w.x.y.z) that is assigned to a site or host. WWXX:YYZZ is the Next Level Aggregator (NLA) portion of a 6to4 address.

The IPv6 Helper service also supports 6over4, also known as IPv4 multicast tunneling, is a tunneling technique that is described in RFC 2529. 6over4 allows IPv6 and IPv4 nodes to communicate by using IPv6 over an IPv4 infrastructure. 6over4 uses the IPv4 infrastructure as a multicast-capable link. In order for 6over4 to work correctly, the IPv4 infrastructure must be IPv4 multicast-enabled.

If this service is stopped or disabled, the computer will only have IPv6 connectivity if connected to a native IPv6 network.

For more information about IPv6, see "Introduction to IPv6" and "Windows IPv6" in the Windows Server Resource Kit Networking Guide.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Network Connections, Protocol

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

Windows Management Instrumentation

Event Log

TCP/IP Protocol Driver

Microsoft IPv6 Protocol Driver

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# IPSEC Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** PolicyAgent

**Executable Name:** lsass.exe

**Log On As:** LocalSystem

**Description:** Provides end-to-end security between clients and servers on TCP/IP networks. Manages IP security (IPSEC) policy, starts the Internet Key Exchange (IKE) and coordinates IPSEC policy settings with the IP security driver. The service is controlled using the net start or net stop policy agent command.

IPsec operates at the IP layer and is transparent to other operating system services and applications. It provides packet filtering, and can negotiate security between sending and receiving computers on IP networks. IPsec can be administratively configured to provide:

- Packet filtering - with actions permit, block, or negotiate security
- Negotiated trust and secure IP communication - The IKE protocol mutually authenticates the sender and receiver of IP data packets based on policy settings. Authentication can use the Kerberos authentication protocol, digital certificates, or a shared secret key (password). IKE automatically generates cryptographic keys, and IPsec security associations.
- Protect IP packets with IPsec secure formats: provides cryptographic integrity, authenticity, and optionally encryption of IP packets.
- Secure end-to-end connections: using IPsec transport mode.
- Secure IP tunnels: using IPsec tunnel mode.

IPsec also provides security for L2TP VPN connections. The IPsec security settings for L2TP are created automatically and require the IPsec service to be running in order to take effect.

*On Windows 2000 this service was called IPSEC Policy Agent.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

TCP/IP Protocol Driver

IPSEC

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 50 (ESP), 51 (AH)

UDP: 500 (ISAKMP)

# Kerberos Key Distribution Center (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Kdc

**Executable Name:** lsass.exe

**Log On As:** LocalSystem

**Description:** Enables users to log on to the network using the Kerberos v5 authentication protocol.

As in other implementations of the Kerberos protocol, the KDC is a single process that provides two services:

- **Authentication Service**

  This service issues ticket granting tickets (TGTs) for connection to the ticket-granting service in its own domain or in any trusted domain. Before a client can ask for a ticket to another computer, it must request a TGT from the authentication service in the clients account domain. The authentication service returns a TGT for the ticket granting service in the target computers domain. The TGT can be reused until it expires, but first access to any domains ticket-granting service always requires a trip to the authentication service in the clients account domain.

- **Ticket-Granting Service (TGS)**

  This service issues tickets for connection to computers in its own domain. When clients want access to a computer, it must request a TGT, and ask for a ticket to the computer. The ticket can be reused until it expires, but first access to any computer always requires a trip to the ticket-granting service in the target computers account domain.

If this service is stopped, users will be unable to log on to the domain and access services.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call

AFD Networking Support Environment

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 88, 544

UDP: 88, 464

# License Logging (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** LicenseService

**Executable Name:** llssrv.exe

**Log On As:** NetworkService

**Description:** Tracks Client Access License usage for server products, such as Internet Information Server (IIS), Terminal Services, File and Print services, as well as other products such as SQL Server and Microsoft Exchange Server.

The License Logging service for Windows Server 2003 operating systems collects site-licensing information on the site license server for a single site. This enables a geographical or organizational unit in an organization to purchase and manage their software licenses separately. However, because each site allocates licenses to users independently of other sites, a user connecting to a product in more than one site might be allocated more than one Client Access License for the same product.

If stopped or disabled, licensing logging for these programs will not be logged.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:** None

**The following system components depend on this service:** None

**IP Port Numbers used:**

TCP: 135 (replication only)

# Logical Disk Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Dmserver

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** Detects and monitors new hard disk drives and sends disk volume information to Logical Disk Manager Administrative Service for configuration. This service watches Plug and Play events for new drives detected and passes volume and/or disk information to the Logical Disk Manager Administrative Service to be configured. The Logical Disk Manager uses an administrator service and a watchdog service. The service should not be disabled if dynamic disks are in the system.

If this service is disabled or stopped, dynamic disk status and configuration information will become out of date and hard disk drives will not be detected. Also, any attempt to use the MMC Disk Management snap-in to configure disks will display the unable to connect to Logical Disk Manager service error message.

See also Logical Disk Manager Administrative Service.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Plug and Play

Remote Procedure Call

**The following system components depend on this service:**

Logical Disk Manager Administrative Service

**IP Port Numbers used:** None

# Logical Disk Manager Administrative Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Dmadmin

**Executable Name:** dmadmin.exe /com

**Log On As:** LocalSystem

**Description:** Performs administrative service for disk management requests. This service is started only when configuring a drive or partition or when a new drive is detected. This service is configured as manual by default, and is activated whenever dynamic disk configuration changes occur or when the Disk Management MMC snap-in is open. Such changes include converting a basic disk to dynamic, recovery of fault tolerant volumes, volume formatting, or changing a page file. The service starts, completes the configuration operation, and then exits.

This service only runs for disk configuration processes once invoked via the MMC Disk Management snap-in and then stops. If this service is disabled, attempts to use the MMC Disk Management snap-in to configure disks will display the unable to connect to Logical Disk Manager service error message.

*See also* **Logical Disk Manager**.

**Available on:** Windows XP Home, Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Logical Disk Manager

Remote Procedure Call

Plug and Play

**The following system components depend on this service:** None

**IP Port Numbers used:** None

# Message Queuing (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** msmq

**Executable Name:** mqsvc.exe

**Log On As:** LocalSystem

**Description:** The Message Queuing service is a messaging infrastructure and development tool for creating distributed messaging applications for Windows. Such applications can communicate across heterogeneous networks and can send messages between computers that might be temporarily unable to connect to each other. Message Queuing provides guaranteed message delivery, efficient routing, security, support for sending messages within transactions, and priority-based messaging. Message Queuing provides both Win32®and COM APIs for all programmatic functionality including administration and management.

If this service is stopped or disabled, distributed messages are unavailable.Disabling the Message Queuing service affects other services including COM+ Queued Component (QC) functionality, some functionality of Windows Management Instrumentation (WMI), and the Message Queuing Triggers service.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components, Application Server, and Message Queuing.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Message Queuing Access Control

NT LM Security Support Provider

Remote Procedure Call

RMCAST <Pgm> Protocol Driver

Security Accounts Manager

**The following system components depend on this service:** None

# Message Queuing Down Level Client Support (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** mqds

**Executable Name:** mqdssvc.exe

**Log On As:** LocalSystem

**Description:** The Message Queuing Down Level Clients Support service provides Active Directory access for Message Queuing clients (Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000, and Windows 2003) on domain controllers. The Message Queuing service optionally uses information published in Active Directory in order to obtain routing information for security-related objects (for example, destination public keys) and to learn about public queues.

If the Message Queuing service is installed in Workgroup mode, Active Directory is never accessed.

If this service is stopped or disabled on a domain controller, then previous versions of Message Queuing clients cannot obtain Active Directory services on the specified domain controller for public queue discoverability, message routing, and site recognition.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components, Application Server, Message Queuing, and Downlevel Client Support

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

IPSEC Driver

Message Queuing

Message Queuing Access Control

NTLM Security Support Provider

Remote Procedure Call (RPC)

RMCAST (Pgm) Protocol Driver

Security Accounts Manager

TCP/IP Protocol Driver

**The following system components depend on this service:** None

# Message Queuing Triggers (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** MSMQTriggers

**Executable Name:** Mqtgsvc.exe

**Log On As:** LocalSystem

**Description:** The Message Queuing Triggers service provides rule-based monitoring of messages arriving in a Message Queuing queue and, when the conditions of a rule are satisfied, invokes a COM component or a stand-alone executable program to process the message.The Message Queuing Triggers service is installed as an integral part of Message Queuing.

If this service is stopped or disabled, you cannot apply rule-based monitoring nor invoke programs to process messages automatically.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components, Application Server, Message Queuing, and Triggers

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Distributed Transaction Coordinator

IPSEC Driver

Message Queuing

Message Queuing Access Control

NTLM Security Support Provider

Remote Procedure Call (RPC)

RMCAST (Pgm) Protocol Driver

Security Accounts Manager

TCP/IP Protocol Driver

**The following system components depend on this service:** None

# Messenger (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Messenger

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Messenger service transmits the net send command and Alerter service messages between clients and servers. For console messages to be received, the Messenger service must be enabled and started on the target computer.This service is not related to any Microsoft Windows or Microsoft® MSN® Messenger applications used for instant messaging.

If this service is stopped or disabled, Alerter service messages are not transmitted — for example, the command-line commands net send and net name cease to function.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Disabled

**This service depends on the following system components:**

NetBIOS Interface

Remote Procedure Call

Plug and Play

Workstation

**The following system components depend on this service:** None

# Microsoft POP3 Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** POP3SVC

**Executable Name:** Pop3svc.exe

**Log On As:** LocalService

**Description:** Microsoft POP3 Service provides e-mail transfer and retrieval services. Administrators can use Microsoft POP3 Service to store and manage e-mail accounts on the mail server. When you install Microsoft POP3 Service on the mail server, users can connect to the mail server and retrieve e-mail to their local computer by using an e-mail client that supports Post Office Protocol 3 (POP3) protocol (such as Microsoft® Outlook®). Microsoft POP3 Service is combined with the Simple Mail Transfer Protocol (SMTP) service — which allows users to send outgoing e-mail — for full e-mail services.Microsoft POP3 Service is the mechanism that allows users to retrieve their e-mail from the mail server. The sender's client computer is connected to the Internet through an Internet service provider (ISP). Using an e-mail client, the sender sends an e-mail, the e-mail is picked up according to the SMTP protocol and it is transferred to the sender's ISP, which then routes the e-mail out onto the Internet. The e-mail moves throughout the Internet and is relayed through a varying number of intermediate servers as it moves toward the recipient. When the e-mail reaches the recipient's ISP, it is dropped into the recipient's mailbox. When the recipient's computer connects to the recipient's ISP, the e-mail is then transferred from the ISP to the recipient's e-mail client on the local computer according to the POP3 protocol.

If this service is stopped or disabled, e-mail transfer and retrieval services do not function.

**Available on:**Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Configure Your Server Wizard or Manage Your Server, or go to Add/Remove Windows Components, E-mail Services, and POP3 Service.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

IIS Admin Service

Remote Procedure Call

Security Accounts Manager

**The following system components depend on this service:** None

# Microsoft Software Shadow Copy Provider (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:**SwPrv

**Executable Name:** svchost.exe -k swprv

**Log On As:** LocalSystem

**Description:** The Microsoft Software Shadow Copy Provider service manages software-based shadow copies taken by the Volume Shadow Copy service. A shadow copy enables you to create a copy (or apparent copy) of a disk volume that represents a consistent read-only point in time for that volume. This point in time then stays constant and allows an application (such as Backup) to copy data from the shadow copy to tape.

There are two general classes of shadow copies:

- **Hardware.** For hardware shadow copies, a mirror of two or more disks is split into separate volumes. One of the two volumes then remains the working set and the other one can be mounted separately.

- **Software.** A software volume uses a copy-on-write scheme to copy all sectors of a volume that change over time into a differential area. When the shadow copy is mounted, all unchanged sectors are read from the original volume and all sectors that have changed are read from the differential area.

Shadow copies can resolve three classic problems with Backup:

- The need to delete share flags in a filter driver in order to back up files that have been opened for exclusive access. The problem of backing up a file as it is changing. Without a shadow copy or a way to suspend the application, the result of doing this is typically a corrupt backup.

- Allowing a system to remain active while the shadow copy is being performed. Tax season is a good example of when organizations cannot afford to take the time to shut a system down long enough to back it up.

- Leveraging the snapshots approach and using the same communication channels to build information transfer between application and backup tools.

If this service is stopped or disabled, software-based volume shadow copies cannot be managed.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** None

# MSSQL$UDDI

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** MSSQL$UDDI

**Executable Name:** sqlservr.exe –sUDDI

**Log On As:** LocalSystem

**Description:** The MSSQL$UDDI service is installed when you install the Universal Description, Discovery, and Integration (UDDI) feature of the Windows Server 2003 family, which provides UDDI capabilities within an enterprise. The SQL Server database engine is the core component of this feature.The MSSQL$UDDI service utilized by UDDI manages all of the files that comprise the databases owned by an instance of SQL Server. It is the component that processes all Transact-SQL statements sent from SQL Server client applications. The MSSQL$UDDI (SQLServer.exe) service allocates computer resources effectively among multiple concurrent users. It also enforces business rules defined in stored procedures and triggers, ensures the consistency of the data, and prevents logical problems such as allowing two people to try to update the same data at the same time.UDDI is an industry specification for the description and discovery of Web services. The UDDI specification builds on the Simple Object Access Protocol (SOAP), XML, and HTTP/S protocol standards developed by the World Wide Web Consortium (W3C) and Internet Engineering Task Force (IETF). UDDI services are standards-based XML Web Services designed for enterprise developers to efficiently publish, discover, share, and reuse Web Services directly through their development tools. Built on the Microsoft .NET Framework, UDDI services use proven SQL Server technology and tools to provide a scalable storage mechanism. IT managers can leverage UDDI services support for standard categorization schemes and Active Directory authentication, allowing for easy integration within an enterprise environment.

For more information about UDDI, see UDDI Services

If this service is stopped, the SQL Server database is unavailable on the computer and clients can no longer query against or access the data in its databases.

If this service is disabled, any services that explicitly depend on this service will fail to start.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

None

**The following system components depend on this service:** SQLAgent$UDDI

# MSSQLServerADHelper

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** MSSQLServerADHelper

**Executable Name:** Sqladhlp.exe

**Log On As:** LocalSystem

**Description:** The MSSQLServerADHelper service enables SQL Server and Microsoft SQL Server Analysis Services to publish information in Active Directory when the services are not running under the LocalSystem account.

There is only one instance of MSSQLServerADHelper service on a computer. All instances of SQL Server and Microsoft SQL Server Analysis Services use it when needed.MSSQLServerADHelper is not a server service and does not service requests from the client. The service does not use a User Datagram Protocol (UDP) or TCP port.

This service cannot be stopped. This service is dynamically started by an instance of SQL Server or Analysis Manager when needed, and is stopped as soon as it has completed its work. This service should always be run from the LocalSystem account and should not be started manually from a console.

If this service is disabled, it might affect adding, updating, or deleting SQL Server–related Active Directory objects if MSSQLSERVER service is not running under LocalSystem.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components and UDDI Services.

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

None

**The following system components depend on this service:** None

# Net Logon (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Netlogon

**Executable Name:** lsass.exe

**Log On As:** LocalSystem

**Description:** The Net Logon service maintains a secure channel between your computer and the domain controller for authenticating users and services. It passes the user's credentials through a secure channel to a domain controller and returns the domain security identifiers and user rights for the user. This is commonly referred to as pass-through authentication.

| **Note** |
| --- |
| The secure channel is not directly used for Kerberos authentication, which is preferred authentication mechanism in Windows 2000 and Windows Server 2003. |

The Net Logon service is started automatically when the computer is a member of a domain. In Windows 2000 Server and Windows Server 2003 family, the Net Logon service publishes service resource records in the Domain Name System (DNS) and uses DNS to resolve names to the IP addresses of domain controllers. The Net Logon service registers DNS records on domain controllers only, not on member servers.

The Net Logon service also implements the replication protocol based on RPC for synchronizing Windows NT 4.0 backup domain controllers (BDCs) and the primary domain controller (PDC).

If this service is stopped or disabled, the computer might not authenticate users and services, and the domain controller cannot register DNS records.Specifically, it might deny NTLM authentication requests and, if this service is stopped or disabled on a domain controller, the domain controller is not discoverable by client computers.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Started (if joined to a domain)

**This service depends on the following system components:**

Workstation

**The following system components depend on this service:** None

# NetMeeting Remote Desktop Sharing (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** mnmsrvc

**Executable Name:** Mnmsrvc.exe

**Log On As:** LocalSystem

**Description:** The NetMeeting Remote Desktop Sharing service allows authorized users to remotely access your Windows desktop from another computer over a corporate intranet by using Microsoft® NetMeeting® conferencing software. The service must be explicitly enabled by NetMeeting, and can be disabled in NetMeeting or shut down by means of an icon in the notification area.

If this service is stopped or disabled, the NetMeeting display driver is unloaded and remote desktop sharing is unavailable.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

None

**The following system components depend on this service:** None

# Network Connections (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Netman

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Network Connections service manages objects in Network Connections, where you can view both network and remote connections. This service is responsible for network configuration of client computers. It displays network status in the notification area on the desktop (the area on the taskbar to the right of the taskbar buttons). Configuration parameters can also be accessed through this service.

This service starts automatically when the startup type is Manual and Network Connections is running.

If this service is stopped, client-side configuration for LAN and dial-up or VPN connections are unavailable.

If this service is disabled, the following might result:

- Connections do not appear in the Connections folder; thus, dial out access is prevented and LAN settings cannot be configured.
- Other services that use Network Connections to check for network-location-aware Group Policy settings do not function properly.
- Events pertaining to media connections and disconnections are not received.
- ICS does not function correctly.
- The ability to configure incoming connections, wireless settings, or your home network is unavailable.
- New connections are not created.
- Any services that explicitly depend on this service will fail to start.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call

**The following system components depend on this service:** Internet Connection Firewall (ICF)/Internet Connection Sharing (ICS)

# Network DDE (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** NetDDE

**Executable Name:** Netdde.exe

**Log On As:** LocalSystem

**Description:** The Network DDE service provides network transport and security for Dynamic Data Exchange (DDE) for programs running on the same computer or on different computers. You can create Network DDE shared folders programmatically or by using Ddeshare.exe on your computer, and make them visible to other applications and computers. Traditionally, the user creating the shared folder will create and run a server process to handle incoming requests from client processes and/or applications (running on the same computer or remotely); after they are connected, these processes can exchange any kind of data over a secure network transport.

This service is disabled by default, and it is only started when invoked by an application using Network DDE, like Clipbrd.exe or Ddeshare.exe. In addition, applications running on remote computers can send messages to other computers that will invoke the Network DDE service.

If this service is stopped, DDE transport and security are unavailable.If this service is disabled, any services that explicitly depend on the service will fail to start.Any applications depending on Network DDE will time out while trying to start the service. If an application on a remote computer tries to start Network DDE on another computer, the remote computer will not be viewable on the network.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Network DDE DSDM

**The following system components depend on this service:** ClipBook

# Network DDE DSDM (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** NetDDEdsdm

**Executable Name:** netdde.exe

**Log On As:** LocalSystem

**Description:** The Network DDE DSDM service manages Dynamic Data Exchange (DDE) shared folders. This service is used only by Network DDE to manage shared DDE conversations. You can create and trust DDE shares by using Ddeshare.exe to allow remote computers and applications to connect and share data. Network DDE DSDM maintains a database of DDE shared folders, including information on which ones are trusted. For each connection request from or to an application, the service queries the database and validates your security settings to determine if the request should be granted.

If this service is stopped, DDE shared folders are unavailable.If this service is disabled, any services that explicitly depend on the service will fail to start. Any applications depending on Network DDE will time out while trying to start the service. If an application on a remote computer tries to start Network DDE on another computer, the remote computer will not be viewable on the network.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

None

**The following system components depend on this service:**

ClipBook

Network DDE

# Network Location Awareness (NLA) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Nla

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Network Location Awareness (NLA) service collects and stores network configuration information such as changes to the names and locations of IP addresses and domain names, and notifies applications when this information changes. The NLA service component supports the Network Location Service. This component provides the infrastructure that informs applications and the operating system of the network connections on a specific computer. This component implements the NS_NLA Windows Sockets namespace.

This service will automatically start when the startup type is Manual.

If this service is disabled, NLA functionality is not available because the NLA service is a demand-start service and it is used by several critical system components, including Winlogon.exe.If this service is stopped, the change-requests of the services that are using NLA for network change notifications are cancelled. These services immediately try to reestablish the change requests, thereby demand-starting the service again almost instantaneously. NLA correctly reports to the Service Control Manager (SCM) thread, which requested the stop, that the service stopped successfully; however, Winlogon.exe or another service immediately restarts before you are able to verify the status in the Services snap-in.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Started

**This service depends on the following system components:**

AFD Networking Support Environment

IPSEC Driver

TCP/IP Protocol Driver

**The following system components depend on this service:**

Internet Connection Firewall (ICF)/Internet Connection Sharing (ICS)

IPSEC Driver

# Network News Transfer Protocol (NNTP) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** NntpSvc

**Executable Name:** Inetinfo.exe

**Log On As:** LocalSystem

**Description:** The Network News Transfer Protocol (NNTP) service provides NNTP, which is a member of the TCP/IP suite of protocols used to distribute network news messages to NNTP servers and clients (newsreaders) on the Internet. NNTP is designed so that news articles are stored on a server in a central database, thus enabling a user to select specific items to read.You can use a news client such as Microsoft Outlook Express to retrieve newsgroups from the server and read headers or bodies of the articles in each newsgroup. You can then post news back to the NNTP server.

| **Note** |
| --- |
| The version of Outlook Express included in Windows 2000 does not support feeds, where two news servers replicate their contents between each other. However, the version included in Exchange Server does include this functionality. |

If this service is stopped or disabled, client computers cannot connect, read, or retrieve posts. If IIS Admin Service is stopped, the NNTP service will stop as well.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components, Application Server, Internet Information Services (IIS), and NNTP Service.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Event Log

IIS Admin Service

Remote Procedure Call (RPC)

Security Accounts Manager

**The following system components depend on this service:** None

# NT LM Security Support Provider (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** NtLmSsp

**Executable Name:** Lsass.exe

**Log On As:** LocalSystem

**Description:** The NT LM Security Support Provider service is the Local Security Authority (LSA) for the system. The NT LM Security Support Provider name used in the Services snap-in is a historical holdover.The LSA is responsible for two major tasks: authentication and management of local security policy objects.The LSA handles all authentication requests on the local system through the use of multiple security service providers. It provides interactive authentication for users in addition to noninteractive authentication for services. The LSA uses a variety of security service providers, which provide several different authentication protocols. Each of these providers is in a different DLL loaded by the LSA:

- Kerberos (Kerberos.dll)
- SChannel (Schannel.dll)
- Digest (Wdigest.dll)
- NTLM — LM, NTLM, and NTLMv2 (Msv1_0.dll)

The second major function of the LSA is to store and manage security objects. These objects include the access tokens for users, local security policy, and the trusted domains. For example, all access control list (ACL) checks use information from the LSA.

This is a critical system process. If it is stopped or disabled, all user and system authentication stops and all security policy checks fail — in short, the system no longer functions. For this reason, the administration tools do not allow an administrator to stop this process.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:** None

**The following system components depend on this service:**

Telnet

Windows Internet Name Service (WINS)

# Performance Logs and Alerts (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SysmonLog

**Executable Name:** Smlogsvc.exe

**Log On As:** NetworkService

**Description:** The Performance Logs and Alerts service collects performance data from local or remote computers based on preconfigured schedule parameters, and then writes the data to a log or triggers an alert. The Performance Logs and Alerts service starts and stops each named performance data collection based on the information contained in the named log collection setting. This service runs only if at least one collection is scheduled.

If this service is stopped or disabled, performance information is not collected. Also, currently running data collections and future scheduled collections will not occur.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:** None

**The following system components depend on this service:** None

# Plug and Play (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** PlugPlay

**Executable Name:** services.exe

**Log On As:** LocalSystem

**Description:** The Plug and Play service enables a computer to recognize and adapt to hardware changes with little or no user input. Plug and Play enables you to add or remove devices without any intricate knowledge of computer hardware, and without being forced to manually configure hardware or the operating system. For example, you can plug in a USB keyboard, and Plug and Play will detect the new device, attempt to find a driver for it, and install it. Or you can dock a portable computer and use the docking station's Ethernet card to connect to the network without changing the configuration. Later, you can undock the same computer and use a modem to connect to the network without making any manual configuration changes.

This service cannot be stopped or disabled through the Services snap-in because of the impact on operating system stability. If this service is stopped using the System Configuration Utility (Msconfig.exe) troubleshooting tool, the Device Manager interface appears empty and no hardware devices are displayed.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:** None

**The following system components depend on this service:**

Fax

Internet Connection Firewall (ICF)/Internet Connection Sharing (ICS)

Logical Disk Manager

Logical Disk Manager Administrative Service

Messenger

Remote Access Auto Connection Manager

Remote Access Connection Manager

Smart Card

Telephony

Virtual Disk Service

Windows Audio

Windows Image Acquisition

# Portable Media Serial Number

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** WmdmPmSN

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Portable Media Serial Number service retrieves the serial number of any portable music player connected to your computer. The service allows Windows Media Device Manager (WMDM) to acquire the serial number from devices so that media content can be copied securely to the device. Without the serial number, you cannot associate the content with the specific device. Thus, protected content might not be allowed to transfer to the device.

To uniquely identify portable media, many storage media manufacturers have implemented a unique serial number stored on a nonvolatile area of the storage. For example, the CompactFlash specification revision 1.3 from the CompactFlash Association (CFA) requires CompactFlash cards to have a unique serial number. Iomega Zip and Jaz disks also have unique serial numbers on them. Recently produced SmartMedia cards also have built in unique IDs. Emerging new portable media, such as MultiMediaCards (MMC), SD Memory Cards, and Sony MemoryStick, are all built with unique IDs by design.

For a portable media reader or adapter to be Windows Media–compliant, it must allow an application to retrieve media serial numbers.

In Windows 2000, this service was called WMDM PMSP Service.

If this service is stopped or disabled, protected content might not be allowed to be transferred to the device and the serial number might not be retrieved from portable media devices.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition (32-bit only)

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:** None

**The following system components depend on this service:** None

# Print Server for Macintosh (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** MacPrint

**Executable Name:** Sfmprint.exe

**Log On As:** LocalSystem

**Description:** The Print Server for Macintosh service enables Macintosh clients to route printer files to a print spooler located on a computer running Windows 2000 or Windows Server 2003. This service also allows any Windows-based server to communicate with a print device by using the AppleTalk protocol.

If this service is stopped or disabled, printing is unavailable to Macintosh clients.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition.

**Installed through:** Add/Remove Windows Components, Other Network File and Print Services, and Print Server for Macintosh.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Apple Talk Protocol

Print Spooler

Remote Procedure Call (RPC)

**The following system components depend on this service:** None

# Print Spooler (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Spooler

**Executable Name:** spoolsv.exe

**Log On As:** LocalSystem

**Description:** The Print Spooler service manages all local and network print queues and controls all print jobs. The print spooler is the center of the Windows printing subsystem. It manages the print queues on the system and communicates with printer drivers and input/output (I/O) components (for example, the USB port and the TCP/IP protocol suite).

For more information about the Print Spooler service, see the *Server Management Guide* of the *Windows Server 2003 Resource Kit* (or see the *Server Management Guide* on the Web at http://www.microsoft.com/reskit).

If this service is stopped, both printing and faxing on the local computer are unavailable.

This service will not automatically start if it is stopped when it has either the Automatic or Manual startup types.

If this service is disabled, any services that explicitly depend on this service will fail to start.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

Fax

Print Server for Macintosh

TCP/IP Print Server

# Protected Storage (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** ProtectedStorage

**Executable Name:** lsass.exe

**Log On As:** LocalSystem, to allow this service to interact with the desktop

**Description:** The Protected Storage service protects the storage of sensitive information, such as private keys, and prevents access by unauthorized services, processes, or users. Protected Storage is a set of software libraries that allows applications to retrieve security and other information from a personal storage location, hiding the implementation and details of the storage itself.

The storage location provided by this service is secure and protected from modification. Protected Storage uses the Hash-Based Message Authentication Code (HMAC) and the Secure Hash Algorithm (SHA1) cryptographic hash functions to encrypt the user's master key. This component requires no configuration.Protected Storage is an outdated service whose functions can now be performed by the Data Protection API (DPAPI), which is currently the preferred method for protected storage. Unlike DPAPI, the interface to Protected Storage is not publicly exposed. Protected Storage is needed to recover keys migrated from earlier versions of the Windows operating system. It is still needed for protecting keys in Internet Explorer and Outlook Express.

If this service is stopped or disabled, some migrations and earlier versions of applications might fail.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:** None

# Remote Access Auto Connection Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** RasAuto

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Remote Access Auto Connection Manager service creates a connection to a remote network whenever a program references a remote DNS or network BIOS (NetBIOS) name or address. This service (sometimes called the autodial service) detects an attempt to resolve the name of a remote computer or shared folder, or an unsuccessful attempt to send packets to a remote computer or shared folder. The service is activated only when there is no network access. In that case, the service brings up a dialog box that offers to make a dial-up or VPN connection to the remote computer.

To assist users, the service maintains a local database of connections that were used in the past to reach named computers or shared folders. When the service detects an unsuccessful attempt to reach a remote computer or shared folder, this service will offer to dial the connection that was last used to reach this remote device.

| **Note** |
| --- |
| Windows Product Activation will start both this service and the Telephony service when you do the product activation for any Windows XP or Windows Server operating system. |

If this service is stopped or disabled, you need to manually set up connections to remote computers.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Plug and Play

Remote Access Connection Manager

Remote Procedure Call (RPC)

Telephony

**The following system components depend on this service:** None

# Remote Access Connection Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** RasMan

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Remote Access Connection Manager service manages dial-up and VPN connections from your computer to the Internet or other remote networks. When you double-click a connection in Network Connections in Control Panel and then click Connect, the Remote Access Connection Manager service dials the connection (or sends a VPN connection request) and handles subsequent negotiations with the remote access server in order to set up the connection.

This service will stop when there are no further requests pending. Network Connections calls on this service to enumerate the set of connections and to display the status of each one. Therefore, except for the case in which there are no connections in Network Connections, the Remote Access Connection Manager service is always able to start.

If this service is stopped or disabled, the computer cannot make dial-up or VPN connections to a remote network, nor accept incoming connections. In addition, Network Connections will not display VPN or dial-up or connections. Any operating system components that depend on this service will fail to operate properly.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Plug and Play

Remote Procedure Call (RPC)

Telephony

**The following system components depend on this service:**

Internet Connection Firewall (ICF)/Internet Connection Sharing (ICS)

# Remote Administration Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SrvcSurg

**Executable Name:** srvcsurg.exe

**Log On As:** LocalSystem

**Description:** Remote Administration Service is responsible for performing the following remote administration tasks when Windows Server 2003 restarts:

- Increment the server boot count.
- Send an alert when the computer is running a backup of the operating system.
- Generate a self-signed certificate.
- Send an alert if the date and time has not been set on the computer.
- Send an alert if Alert E-mail functionality has not been configured.

Remote Administration Service starts performing the appropriate tasks when requested by the Remote Server Manager service through a COM interface. Remote Administration Service runs in the LocalSystem account and accepts requests through the COM interface only from clients running under the Administrators or LocalSystem accounts.

If the startup type is set to Manual, the service will start when it receives a request from the Remote Server Manager service; it can be subsequently stopped without affecting any server functionality.

If this service is stopped or disabled, some remote administration tools or features might not function properly, such as the Web interface for remote administration.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:** None

# Remote Desktop Help Session Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** RDSessMgr

**Executable Name:** Sessmgr.exe

**Log On As:** LocalSystem

**Description:** The Remote Desktop Help Session Manager service manages and controls the Remote Assistance feature within the Help and Support Center application (Helpctr.exe).

If this service is stopped or disabled, Remote Assistance is unavailable.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

Terminal Services

**The following system components depend on this service:**

None

# Remote Installation Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** BINLSVC

**Executable Name:** Tcpsvcs.exe

**Log On As:** LocalSystem

**Description:** Remote Installation Services (RIS) is a Windows deployment feature that is included in Windows Server 2003. With RIS, you can support on-demand, image-based, or script-based operating system installations over a network connection from a RIS server to a client computer. RIS is designed to simplify the deployment of operating systems and applications, and to improve failure recoverability.

You can use RIS in a variety of ways, including the following:

- Provide an operating system to users on demand.

You can use RIS to create automated installation images of Windows Server 2003, Windows XP, and Windows 2000. When a user starts a client computer, even if that computer contains no operating system, the RIS server can respond by installing an operating system over the network; no CD is required. To support this, client computers must use the Pre-boot eXecution Environment (PXE), a remote-boot technology that allows the client computer to begin a boot sequence from BIOS, using the network adapter. For more information about the process through which RIS installs an operating system, see PXE architecture and Technologies and components used by Remote Installation Services in see the *Server Management Guide* of the *Windows Server 2003 Resource Kit* (or see the *Server Management Guide* on the Web athttp://www.microsoft.com/reskit).

- Provide operating system images which include specific settings and applications.

For example, you can use RIS to create an image that complies with a corporate desktop standard and offer it to a particular group of users.

| **Note** |
| --- |
| In Windows 2000, this service was called the Boot Information Negotiation Layer Service. |

If this service is stopped or disabled, PXE-enabled client computers are unable to install Windows remotely or use other RIS-based tools.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition

**Installed through:** Add/Remove Windows Components and Remote Installation Services.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Server

TCP/IP Protocol Driver

**The following system components depend on this service:**

None

# Remote Procedure Call (RPC) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** RpcSs

**Executable Name:** svchost -k rpcss

**Log On As:** LocalSystem

**Description:** The Remote Procedure Call (RPC) service is a secure interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality that resides in a different process. That different process can be on the same computer, on the LAN, or across the Internet. The Remote Procedure Call service serves as the RPC endpoint mapper and Service Control Manager (SCM). Many services depend on the Remote Procedure Call service to start successfully.

For more information about RPC, see the *Internetworking Guide* of the *Windows Server 2003 Resource Kit* (or see the *Internetworking Guide* on the Web at http://www.microsoft.com/reskit).

This service should not be stopped or disabled. Disabling the Remote Procedure Call (RPC) service will result in the operating system not loading.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:** None

**The following system components depend on this service:**

Background Intelligent Transfer Service

Cluster Service

COM+ Event System

COM+ System Application

Cryptographic Services

DHCP Server

Distributed Link Tracking Client

Distributed Link Tracking Server

Distributed Tracking Coordinator

DNS Server

Error Reporting Service

Fax

File Replication

Help and Support

Human Device Interface Access

IIS Admin Service

Indexing Service

Internet Authentication Service

IPSEC Services

IPv6 Helper Service

Kerberos Key Distribution Center

Logical Disk Manager

Logical Disk Administrator Service

Messenger

MS Software Shadow Copy Provider

Network Connections

Print Spooler

Protected Storage

Remote Desktop Help Session Manager

Remote Registry

Removable Storage

Resultant Set of Policy Provider

Routing and Remote Access

Security Accounts Manager

Shell Hardware Detection

Task Scheduler

Telephony

Telnet

Terminal Services

Terminal Services Session Directory

Terminal Services Licensing

Upload Manager

Volume Shadow Copy

Web Element Manager

Windows Audio

Windows Image Acquisition (WIA)

Windows Installer

Windows Internet Name Service (WINS)

Windows Management Instrumentation

Windows Media Services

Wireless Configuration

WMI Performance Adapter

World Wide Web Publishing Service

# Remote Procedure Call (RPC) Locator (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** RpcLocator

**Executable Name:** Locator.exe

**Log On As:** NetworkService

**Description:** The Remote Procedure Call (RPC) Locator service enables RPC clients using the RpcNs family of APIs to locate RPC servers, and manages the RPC name service database.

The RPC Locator service (also known as RPC Named Service) can locate other services on the network. The instance of the RPC Locator service that runs on a domain controller acts as repository for information about other RPC services that are active on the network. If you run the RPC Locator service in Windows 2000 or Windows Server 2003, backup storage for information about other services on the network is Active Directory. The instances of the RPC Locator service that run on individual computers act as computer-wide caches for information about other services on the network. Those services that need to be located through the RPC Locator service register with the RPC Locator service. The RPC APIs that start with the prefix RpcNs are used to register services with the RPC Locator service and query the RPC Locator service about previously registered services.

For more information about the RpcNs family of APIs, see the Software Development Kit (SDK) information in the MSDN Library link on the Web Resources page (http://go.microsoft.com/fwlink/?linkid=291)

If this service is stopped or disabled, RPC clients that use RpcNs APIs might be unable to locate servers or fail to start. RPC clients that rely on RpcNs APIs from the same computer might not be able to find RPC servers that support a given interface. If this service is stopped or disabled on a domain controller, RPC clients that use the RpcNs APIs and the domain controller might experience interruption of service while trying to locate clients.

RpcNs APIs are not used internally in Windows; you need to start this service only if third-party applications require it.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual (Automatic on a domain controller)

**Service status:** Stopped (Started on a domain controller)

**This service depends on the following system components:**

Workstation

**The following system components depend on this service:**

None

# Remote Registry (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** RemoteRegistry

**Executable Name:** svchost.exe -k regsvc

**Log On As:** LocalService

**Description:** The Remote Registry service enables remote users to modify registry settings on the local computer, provided the remote users have the required administrative credentials. The service is primarily used by remote administrators and for performance counters.

If this service is stopped, the registry on the local computer can only be modified locally.

If this service is disabled, it will not affect registry operations on the local computer, but other computers or devices can no longer connect to the local computer's registry.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Remote Server Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** AppMgr

**Executable Name:** Appmgr.exe

**Log On As:** LocalSystem

**Description:** The Remote Server Manager service provides the following functionality:

- Holds the Remote Administration Service alert information.
- Provides an interface for raising, clearing, and enumerating Remote Administration Service alerts.
- Provides an interface for performing Remote Administration Service tasks.

The Remote Server Manager service acts as a Windows Management Instrumentation (WMI) instance provider for Remote Administration Service alert objects and a WMI method provider for Remote Administration Service tasks. The Remote Server Manager service runs in the LocalSystem account and accepts requests through the COM interface only from clients running under the Administrators or LocalSystem accounts.

If the startup type is set to Manual, the service will start when the next request for Remote Administration Service tasks or Remote Administration Service alerts is received.

When this service is stopped, the service is restarted if the Web interface for remote administration is accessed.

If this service is disabled, information about the current Remote Administration Service alerts on the system is lost.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Remote Storage Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Remote_Storage_Server

**Executable Name:** RsServ.exe

**Log On As:** LocalSystem

**Description:** The Remote Storage Server service stores infrequently used files in a secondary storage media. This service allows the Remote Storage application to notify the user when an offline file has been accessed.

Remote Storage is a hierarchical storage management application that migrates data from upper-level storage to lower-level storage. Upper-level storage is commonly known as local storage: data that is frequently accessed is stored locally on high-performance disks. Lower-level storage is also known as remote storage: data that is infrequently accessed is stored on less expensive media until it is needed again. Hierarchical storage management reduces the costs of storing large quantities of data, while ensuring that the data is still accessible.

File data is cached locally so that it can be accessed quickly as needed. When the amount of available space on a managed volume dips below the necessary level, Remote Storage automatically removes the content from the cached files, providing the disk space you need. When data is removed from a file, the disk space used by the file is reduced to zero. Data from cached files is not removed until more disk space is needed. When you need to open a file from which data has been removed, the data is automatically recalled from remote storage.

Because removable tapes or disks in a library are less expensive per megabyte (MB) than hard disks, Remote Storage can be an economical way to provide both maximum data storage and optimal local disk performance.

For more information about the Remote Storage Server service and the Remote Storage application, see the *Server Management Guide* of the *Windows Server 2003 Resource Kit* (or see the *Server Management Guide* on the Web at http://technet.microsoft.com/en-us/windowsserver/bb633748.aspx)

If this service is stopped or disabled, files cannot be moved to or retrieved from the secondary storage media.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components and Remote Storage

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Event Log

Remote Procedure Call (RPC)

Removable Storage

Task Scheduler

**The following system components depend on this service:**

None

# Removable Storage (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** NtmsSvc

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Removable Storage service manages and catalogs removable media and operates automated removable media devices. This service maintains a catalog of identifying information for removable media used by your computer, including tapes and CDs. If the computer also has automated devices for maintaining removable media (such as a tape autoloader or CD jukebox), Removable Storage also automates mount, dismount, and eject media. Removable Storage is used by applications such as Backup and Remote Storage to handle media cataloging and automation.

Removable Storage labels, catalogs, and tracks media; controls library drives, slots, and doors; and provides drive-cleaning operations.

For more information about the Removable Storage service, see the Server Management Guide of the Windows Server 2003 Resource Kit (or see the Server Management Guide on the Web at http://technet.microsoft.com/en-us/windowsserver/bb633748.aspx)

If this service is stopped or disabled, applications that are dependent on the Removable Storage service, such as Backup and Remote Storage, operate more slowly.

The Removable Storage service stops when there is nothing to process. If there are no automated devices attached to the system, the Removable Storage service runs only while there are applications using it; therefore, it is not necessary to stop the service. When started under these circumstances, the Removable Storage service frequently needs to inventory the complete contents of attached autoloaders and jukeboxes, which includes mounting each media in a drive.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Resultant Set of Policy Provider (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** RSoPProv

**Executable Name:** RSoPProv.exe

**Log On As:** LocalSystem

**Description:** The Resultant Set of Policy Provider service enables you to connect to a Windows Server 2003 domain controller, access the Windows Management Instrumentation (WMI) database for that computer, and simulate Resultant Set of Policy (RSoP) for Group Policy settings. The policy settings are applied to a user or computer located in Active Directory on a Windows 2000 or later domain.

This service is used by the RsoP snap-in, and is also accessed by any client of the RSoP-Planning Mode Provider/WMI Provider. RSoP is an addition to Group Policy that makes policy implementation and troubleshooting easier. RSoP is a query engine that polls existing policy settings and planned policy settings, and then reports the results of those queries. It polls existing policy settings based on site, domain, domain controller, and organizational unit. RSoP gathers this information from the Common Information Management Object Model (CIMOM) database (otherwise known as CIM-compliant object repository) through WMI.

RSoP provides details about all policy settings that are configured by an administrator, including Administrative Templates, Folder Redirection, Internet Explorer Maintenance, Security Settings, Scripts, and Software Installation.

When policy settings are applied on multiple levels (for example, site, domain, domain controller, and organizational unit), the results can conflict. RSoP can help you determine a set of applied policy settings and their precedence (that is, the order in which policy settings are applied).

RSoP consists of two modes: planning mode and logging mode. By using planning mode, you can simulate the effect of policy settings that you want to apply to a computer and user. Logging mode reports the existing policy settings for a computer and the user that is currently logged on.

This service is used by the RsoP snap-in, and is also accessed by any client of the RsoP-Planning Mode Provider/WMI Provider.

For more information about RSoP, see the *Distributed Services Guide* of the *Windows Server 2003 Resource Kit* (or see the *Distributed Services Guide* on the Web at http://technet.microsoft.com/en-us/windowsserver/bb633748.aspx)

If this service is stopped or disabled, RSoP planning mode simulation is unavailable on the domain controller.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Routing and Remote Access (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** NtmsSvc

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Routing and Remote Access service provides multiprotocol LAN-to-LAN, LAN-to-WAN, VPN, and network address translation (NAT) routing services. In addition, Routing and Remote Access also provides dial-up and VPN remote access services.

The Routing and Remote Access service (note capitalization) replaces the Routing and Remote Access Service (RRAS) and Remote Access Service (RAS) features from Windows NT 4.0. The Routing and Remote Access service is a single, integrated service that terminates connections from either dial-up or VPN clients or that provides routing (IP, IPX, and Services for Macintosh), or both. With the Routing and Remote Access service, your server can function as a remote access server, a VPN server, a gateway, or a branch-office router.

From a routing perspective, the Routing and Remote Access service supports the Open Shortest Path First (OSPF) and Routing Information Protocol (RIP) routing protocols, and controls the routing tables for the TCP/IP stack-forwarding engine.

If this service is stopped or disabled, the remote access server cannot accept incoming RAS, VPN, or demand-dial connections, and routing protocols are not received or transmitted.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

NetBIOS Group

NetBIOS Interface

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# SAP Agent (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** NwSapAgent

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The SAP Agent service advertises network services on an IPX network by using the Internetwork Packet Exchange Service Advertising Protocol (IPX SAP) protocol. It also forwards advertisements on a multihomed host. Some features, such as Client Services for NetWare, rely on the SAP Agent.

The SAP Agent service requires the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol to be installed.

If this service is stopped or disabled, IPX service discovery does not work; the server does not advertise itself, and clients are not able to reach it.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** In Network Connections, right-click a local area connection, click Properties, click Install, click Service, click Add, and then click Service Advertising Protocol.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

NWLink IPX/SPX/NetBIOS Compatible Transport Protocol

**The following system components depend on this service:**

None

# Secondary Logon (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** seclogon

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Secondary Logon service enables the starting of processes that run under alternate administrative credentials so that the user can create processes in the context of different security principals. A common use of this service is for administrators to log on as a restricted user and then use Secondary Logon to temporarily run an application as Administrator.

Another feature of this service is the Runas command. Using the Runas command, you can run programs (*.exe files), saved MMC consoles (*.msc files), shortcuts to programs, or Control Panel items as an Administrator while you are logged on to your computer as a member of another group, such as the Users or Power Users group. As long as you provide the appropriate user account and password information, the user account has the ability to log on to the computer, and the program, MMC console, or Control Panel item is available on the system and to the user account.

| **Note** |
| --- |
| In Windows 2000, this services was called the RunAs Service. |

If this service is stopped or disabled, the Run as snap-in generates the error message, “Windows cannot access the specified device, path or file. You may not have the appropriate permissions to access them.” Any calls to the CreateProcessWithLogonW API fail. Specifically, this causes the Run as snap-in and the run as command to function improperly.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

None

**The following system components depend on this service:**

None

# Security Accounts Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SamSs

**Executable Name:** Lsass.exe

**Log On As:** LocalSystem

**Description:** The Security Accounts Manager (SAM) service is a protected subsystem that manages user and group account information. In Windows 2000 and Windows Server 2003, workstation security accounts are stored by SAM in the registry of the local computer, and domain controller security accounts are stored in Active Directory. In Windows NT 4.0, both local and domain security accounts are stored in the registry.

The startup of this service signals other services that SAM is ready to accept requests.

This service cannot be stopped.

If you disable this service, other services in the system are not notified when SAM is ready, which might cause those services to fail to start correctly. This service should not be disabled.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

DHCP Server

Distributed File System

Distributed Transaction Coordinator

FTP Publishing Service

HTTP SSL

IIS Admin Service

Intersite Messenging

Message Queuing

Message Queuing Downlevel Client Support

Message Queuing Triggers

Microsoft POP3 Service

Network News Transfer Protocol (NNTP)

Simple Mail Transfer Protocol (SMTP)

Windows Internet name Services (WINS)

World Wide Web Publishing Service

# Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** lanmanserver

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Server service provides RPC support, file printing, and named pipe sharing over the network. It allows the sharing of local resources (such as disks and printers) so that other users on the network can access them. It also allows named pipe communication between applications running on other computers and your computer, which is used for RPC.

Named pipe communication is memory reserved for the output of one process to be used as input for another process. The input-accepting process does not need to be local to the computer.

If this service is stopped, you cannot share files and printers on this computer with other computers on the network, and the computer cannot satisfy RPC requests. You also cannot serve named pipes to other computers on the network.

If this service is disabled, any services that explicitly depend on this service will fail to start.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

None

**The following system components depend on this service:**

Computer Browser

Distributed File System

Net Logon

# Shell Hardware Detection (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** ShellHWDetection

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Shell Hardware Detection service monitors and provides notification for AutoPlay hardware events. AutoPlay is a feature that detects content such as pictures, music, or video files on removable media and removable devices. AutoPlay then automatically starts applications to play or display that content. This simplifies the use of specialized peripheral devices such as Moving Picture Experts Group (MPEG) Audio Layer-3 (MP3) players and digital photo readers. It also makes it easier for users who are unfamiliar with the software needed to access various content types.

AutoPlay supports a variety of media content types and applications. Both independent hardware vendors (IHVs) and independent software vendors (ISVs) can extend this support to include their hardware devices and applications. A user can configure a different AutoPlay action for any combination of pictures, music files, and video.

Supported media and device types include the following:

- Removable storage media, such as Zip disks by Iomega
- Flash media, such as the CompactFlash specification from the CompactFlash Association, the SmartMedia specification from the SSFDC Forum, and the Memory Stick specification developed by Sony
- PC Cards
- External USB or IEEE 1394 fixed drives

Supported content types include the following:

- Pictures (.jpg, .bmp, .gif, .tif)
- Music files (.mp3, .wma)
- Video (.mpg, .asf)

If this service is stopped, you lose the hardware AutoPlay functionality, My Computer icons and labels are displayed in Classic View, and shell performance is also affected.

If this service is disabled, any services that explicitly depend on this service will fail to start.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

Windows Image Acquisition (WIA)

# Simple Mail Transfer Protocol (SMTP) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SMTPSVC

**Executable Name:** Inetinfo.exe

**Log On As:** LocalSystem

**Description:** The Simple Mail Transfer Protocol (SMTP) service transports e-mail across the network. It is used as an e-mail submission and relay agent. It can accept and queue email for remote destinations and retry at specified intervals. Windows domain controllers use the SMTP service for intersite e-mail–based replication. The Collaboration Data Objects (CDO) for Windows Server 2003 COM component can use the SMTP service to submit and queue outbound e-mail.

Other applications — for example, Microsoft Exchange Server — might use the SMTP service as the basis for their SMTP support.

In Windows 2000, the SMTP service was installed by default. In Windows Server 2003, this service is now an optional component installed through Add/Remove Windows Components.

If this service is stopped or disabled, e-mail–based replication and forwarding do not function.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Event Log

IIS Admin Service

Remote Procedure Call (RPC)

Security Accounts Manager

**The following system components depend on this service:**

None

# Simple TCP/IP Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SimpTcp

**Executable Name:** Tcpsvcs.exe

**Log On As:** NetworkService

**Description:** Simple TCP/IP Services implements support for the Echo, Discard, Character Generator, Daytime, and Quote of the Day protocols.

- **Echo (port 7, RFC 862)** This protocol echoes back data from any messages it receives on this server port. Echo can be useful as a network debugging and monitoring tool.

- **Discard (port 9, RFC 863)** This protocol discards all messages received on this port without response or acknowledgment. It can serve as a null port for receiving and routing TCP/IP test messages during network setup and configuration or, in some cases, it can be used by programs as a message discard function.

- **Character Generator (CHARGEN) (port 19, RFC 864)** This protocol sends data made up of the set of 95 printable ASCII characters. It is useful as a debugging tool for testing or troubleshooting line printers.

- **Daytime (port 13, RFC 867)** This protocol returns messages containing the day of the week, month, day, year, current time (in hh:mm:ss format), and time-zone information. Some programs can use the output from this service for debugging or monitoring variations in system clock time or on a different host.

- **Quote of the Day (port 17, RFC 865)** This protocol returns a quotation as one or more lines of text in a message. Quotations are taken at random from the following file: *systemroot*\System32\Drivers\Etc\Quotes. A sample quote file is installed with Simple TCP/IP Services. If this file is missing, the quote service fails.

To find RFC standards, see the Request For Comments (RFC) link on the Web Resources page at http://technet.microsoft.com/windowsserver/bb633748.aspx.

After this service is installed and started, all five protocols are enabled on all network adapters. There is no provision for selectively enabling specific protocols or enabling this service on a per–network adapter basis.

Stopping or disabling this service has no effect on the rest of the operating system.

You should not install Simple TCP/IP Services unless you specifically need this service to support communication with other systems that use this protocol service.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

AFD Networking Support Environment

**The following system components depend on this service:**

None

# Single Instance Storage Groveler (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Groveler

**Executable Name:** grovel.exe

**Log On As:** LocalSystem

**Description:** The Single Instance Storage Groveler (SIS) service is an integral component of Remote Installation Services (RIS). SIS reduces the overall storage required on the RIS volume. The SIS drivers contain a feature called a "groveler agent" that scans the RIS volume for duplicate files. If the SIS groveler finds duplicate files, it copies the original file into the SIS and leaves a link file in its place. The link file contains information about the original file, such as its current location, size, and attributes. If an image contains duplicate files, those duplicates are copied into the store; as a result, less disk space is required on the RIS server.

SIS has two limitations:

- It cannot act upon any files that are referenced through junction points.
- It cannot be used with any file system except NTFS, which is the only file system supported on RIS servers.

If this service is stopped or disabled, files are no longer automatically linked in this manner, but the existing linked files are still accessible.

New RIS installation images will consume their full image size, and conserve little space.

If SIS is no longer needed on the system, use Add/Remove Windows Components to remove Remote Installation Services.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components and Remote Installation Services. (When you install Remote Installation Services, this service is automatically installed.)

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Single Instance Storage

**The following system components depend on this service:**

None

# Smart Card (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SCardSvr

**Executable Name:** SCardSvr.exe

**Log On As:** LocalService

**Description:** The Smart Card service manages and controls access to a smart card inserted into a smart card reader attached to your computer. The smart card subsystem is based on personal computer/smart card (PC/SC) consortium standards and consists of the following components:

**Resource manager.** This component manages access to readers and smart cards. To manage these resources, it performs the following functions:

- Identifies and tracks resources.
- Allocates readers and resources across multiple applications.
- Supports transaction primitives for accessing services available on a given card.

The resource manager also exposes a subset of the Win32 API to provide applications with access to these functions.

**Card/Reader Selection UI.** This component allows simple applications that work with smart cards to access a card and reader with minimum coding.

If this service is stopped or disabled, your computer cannot read smart cards.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Plug and Play

**The following system components depend on this service:**

None

# SNMP Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SNMP

**Executable Name:** Snmp.exe

**Log On As:** LocalSystem

**Description:** The SNMP service allows incoming Simple Network Management Protocol (SNMP) requests to be serviced by the local computer. SNMP includes agents that monitor activity in network devices and report to the network console workstation. SNMP provides a method of managing network hosts such as workstation or server computers, routers, bridges, and hubs from a centrally located computer running network management software. SNMP performs management services by using a distributed architecture of management systems and agents.

SNMP can be used to:

- Configure remote devices. Configuration information can be sent to each networked host from the management system.
- Monitor network performance. You can track the speed of processing and network throughput, and collect information about the success of data transmissions.
- Detect network faults or inappropriate access. You can configure trigger alarms on network devices when certain events occur. When an alarm is triggered, the device forwards an event message to the management system. Common types of events that need alarms include a device being shut down and restarted, a link failure being detected on a router, and inappropriate access.
- Audit network usage. You can monitor both overall network usage to identify user or group access, and types of usage for network devices and services.

The SNMP service also includes an SNMP agent that allows remote, centralized management of computers running the following:

- Microsoft® Windows® XP Home Edition
- Microsoft® Windows® XP Professional
- Microsoft® Windows® 2000 Professional
- Microsoft® Windows® 2000 Server
- Microsoft Windows Server 2003 family

The SNMP agent also allows management of the following services:

- Windows XP or the Windows Server 2003 family and Windows 2000–based WINS
- Windows XP or the Windows Server 2003 family and Windows 2000–based DHCP
- Windows XP or the Windows Server 2003 family and Windows 2000–based IIS
- LAN Manager

If this service is stopped or disabled, the computer cannot respond to SNMP requests. If the computer is being monitored by network management tools, those tools cannot collect data from the computer nor control its functionality through the SNMP

service.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components, Management and Monitoring Tools, and Simple Network Management Protocol.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Event Log

**The following system components depend on this service:**

None

# SNMP Trap Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SNMPTRAP

**Executable Name:** Snmptrap.exe

**Log On As:** LocalService

**Description:** SNMP Trap Service receives trap messages generated by local or remote Simple Network Management Protocol (SNMP) agents and forwards the messages to SNMP management programs running on your computer.

The SNMP service, when configured for an agent, generates trap messages if any specific events occur. These messages are sent to a trap destination. For example, an agent can be configured to initiate an authentication trap if a request for information is sent by an unrecognized management system.

Trap destinations consist of the computer name or the IP or IPX address of the management system. The trap destination must be a network-enabled host that is running SNMP management software. Trap destinations can be configured by a user, but the events (such as restarting a computer) that generate a trap message are internally defined by the SNMP agent.

If this service is stopped or disabled, SNMP-based programs on your computer cannot receive SNMP trap messages. If this computer is being used to monitor network devices or server applications by using SNMP traps, significant system events are lost.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Add/Remove Windows Components, Management and Monitoring Tools, and Simple Network Management Protocol. (When you install Simple Network Management Protocol, the SNMP Trap Service is automatically installed.)

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Event Log

**The following system components depend on this service:**

None

# Special Administration Console Helper (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Sacsvr

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** You can use the Special Administration Console (!SAC) Helper service to perform remote management tasks if Windows Server 2003 stops functioning because of a Stop error.

The main functions of !SAC are to:

- Redirect Stop error message explanatory text.
- Restart the system.
- Obtain computer identification information.

The Special Administration Console (!SAC) Helper service is an auxiliary Emergency Management Services command-line environment that is hosted by Windows Server 2003. It also accepts input and sends output through the out-of-band port. !SAC is a separate entity from both the Special Administration Console (SAC) and Windows Server 2003 command-line environments. After a specific failure point is reached, Emergency Management Services components determine when the shift should be made from SAC to !SAC. !SAC becomes available automatically if SAC fails to load or is not functioning.

If this service is stopped or disabled, !SAC services are not available.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

None

**The following system components depend on this service:**

None

# SQLAgent$UDDI

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SQLAgent$UDDI

**Executable Name:** SQLagent.exe –iUDDI

**Log On As:** Network Service or Domain Account

**Description:** The SQLAgent$ UDDI service is a job scheduler and monitoring service. It also moves information between computers running Microsoft SQL Server and is used heavily for backups and replication.

If this service is stopped or disabled, replication does not occur. Also, all scheduled jobs and alert/event monitoring, and auto restart of the SQL Server service, will be disrupted.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition (32-bit only)

**Installed through:** Add/Remove Windows Components and UDDI Services.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

MSSQL$UDDI

**The following system components depend on this service:**

None

# System Event Notification (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** SENS

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The System Event Notification service (SENS) tracks system events such as Windows logon network and power events, and notifies the COM+ Event System service subscribers of these events.

If this service is stopped or disabled, COM+ Event System service subscribers do not receive system event notifications. They also encounter the following problems:

- Win32 APIs IsNetworkAlive() and IsDestinationReachable() APIs do not work. These APIs are mostly used by mobile applications on portable computers.
- ISens* interfaces — specifically, SENS Logon/Logoff notifications — do not work.
- SyncMgr (mobsync.exe) does not work properly, because it depends on connectivity information and also on Network Connect/Disconnect and Logon/Logoff notifications from SENS.
- COM+ EventSystem tries to notify SENS of some events, but it does not succeed.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition.

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

COM+ Event System

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Task Scheduler (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Schedule

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Task Scheduler service allows you to perform automated tasks on a computer. By using Task Scheduler, you can schedule any script, program, or document to run at a time that is most convenient for you.

By using Task Scheduler, you can:

- Schedule a task to run daily, weekly, monthly, or at certain times (such as system startup).
- Change the schedule for a task.
- Stop a scheduled task.
- Customize how a task runs at a scheduled time.

Each scheduled task you create is stored as a .job file in the C:\Windows\Tasks folder. The .job file contains the properties and configuration information for the task. You can create a scheduled task on your computer and then drag the .job object to a remote computer. After you drag a task from one computer to another, you must update the account information for the task before it will run. Removing a scheduled task removes only the .job file from the schedule. The program file the task runs is not removed from the hard disk.

When you first install Windows Server 2003, the Task Scheduler service is started. By default, you must be a member of the Administrators, Backup Operators, or Server Operators group on the local computer to view, add, edit, or delete scheduled tasks, or to disable, pause, or restart the Task Scheduler service.

The Task Scheduler service can be started from the Services snap-in or set to automatically start. Task Scheduler can be accessed from the Task Scheduler graphical user interface by using the Task Scheduler API as described in the Microsoft Platform Software Development (SDK), or by using the Schtasks.exe tool.

For more information about Task Scheduler APIs, see the Microsoft Platform SDK link on the Web Resources page (http://go.microsoft.com/fwlink/?linkid=291).

If this service is stopped or disabled, tasks do not run at their scheduled times. For example, jobs that are scheduled to run do not run at their designated time or interval. If there are no jobs listed in the *systemroot*\Tasks\ folder, the effect of disabling this service is minimal; otherwise, jobs that are required to run might be severely impacted.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# TCP/IP NetBIOS Helper (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** LMHosts

**Executable Name:** svchost.exe -k LocalService

**Log On As:** LocalService

**Description:** The TCP/IP NetBIOS Helper service provides support for the NetBIOS over TCP/IP (NetBT) service and NetBIOS name resolution for clients on your network, enabling users to share files, print, and log on to the network. The TCP/IP NetBIOS Helper service provides support for the NetBT service by performing DNS name resolution and pinging a set of IP addresses that return a list of accessible IP addresses.

If this service is stopped or disabled, NetBT, RDR, SRV, Net Logon, and Messenger service clients might not be able to share files or printers, or log on to computers.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

AFD Networking Support Environment

NetBios over Tcpip

**The following system components depend on this service:**

None

# TCP/IP Print Server (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** LPDSVC

**Executable Name:** Tcpsvcs.exe

**Log On As:** LocalSystem

**Description:**The TCP/IP Print Server service enables TCP/IP-based printing by using the Line Printer Daemon protocol. The TCP/IP Print Server service on the print server receives documents from native Line Printer Remote (LPR) utilities running on UNIX computers.

If this service is stopped or disabled, TCP/IP-based printing is unavailable.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition

**Installed through:** Add/Remove Windows Components, Other Network File and Print Services, and Print Services for Unix. (The TCP/IP Print Server and LPRMON services are installed when you install Print Services for Unix.)

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

IPSEC Driver

Print Spooler

Remote Procedure Call (RPC)

TCP/IP Protocol Driver

**The following system components depend on this service:**

None

# Telephony (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** TapiSrv

**Executable Name:** svchost.exe -k tapisrv

**Log On As:** LocalSystem

**Description:** The Telephony service provides Telephony API (TAPI) support. Clients use the Telephony service to run programs that control telephony devices and IP-based voice connections on your computer and throughout the network. The Telephony service enables applications to act as clients to telephony equipment such as PBXs, telephones, and modems. Telephony supports the TAPI by which different wire protocols that communicate with telephony equipment can be supported. These protocols are implemented in Telephony service providers (TSPs).

If this service is stopped or disabled, any services that explicitly depend on this service, such as modem support, fail to start.

The Telephony service cannot be stopped if there is another dependent service, such as RAS, currently active. If there are no other dependent services running and you stop the Telephony service, it will be restarted when any application makes an initialization call to the TAPI interface.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Plug and Play

Remote Procedure Call (RPC)

**The following system components depend on this service:**

Fax

Internet Connection Firewall (ICF)/Internet Connection Sharing (ICS)

Remote Access Auto Connection Manager

Remote Access Connection Manager

# Telnet (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** TlntSvr

**Executable Name:** Tlntsvr.exe

**Log On As:** LocalService

**Description:** The Telnet service supports connections from various TCP/IP Telnet clients, including UNIX-based and Windows-based computers. Telnet Server for Windows provides ASCII terminal sessions to Telnet clients. Telnet Server supports two types of authentication and supports four types of terminals: ANSI, VT100, VT52, and VTNT.

Telnet also allows a remote user to log on to the system and run console programs by using the command line.

If this service is stopped, remote users cannot connect to the computer by using Telnet.

If this service is disabled, remote users are unable to connect to this computer through the Telnet protocol, and users cannot connect to the computer or run console-based applications.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

IPSEC Driver

NTLM Security Support Provider

Remote Procedure Call (RPC)

TCP/IP Protocol Driver

**The following system components depend on this service:**

None

# Terminal Server Licensing (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** TermServLicensing

**Executable Name:** Lserver.exe

**Log On As:** LocalSystem

**Description:** The Terminal Server Licensing service installs a license server and provides registered client licenses when a computer is connecting to a server that has Terminal Server enabled. The Terminal Server Licensing service is a low-impact service that stores the client licenses that have been issued for a terminal server and tracks the licenses that have been issued to client computers or terminals.

If this service is stopped or disabled, the computer running the Terminal Server Licensing service is unable to issue Terminal Server licenses to clients when they request them. If another Terminal Server License Server is installed on a domain controller in the forest, the requesting terminal server attempts to employ it.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Add/Remove Windows Components and Terminal Server Licensing.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Event Log

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Terminal Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** TermService

**Executable Name:** svchost.exe -k termsvcs

**Log On As:** LocalSystem

**Description:** Terminal Services provides a multisession environment that allows users to access a virtual Windows desktop session and Windows-based programs running on a remote computer. Terminal Services allows multiple users to be connected interactively to a terminal server (a server that has Terminal Server enabled), allows users to access desktops and applications on remote computers, and allows users to access another user's desktop for administration purposes.

If this service is stopped or disabled, the computer might become unreliable. To prevent remote use of your computer, clear the **Allow Remote Assistance** and **Allow Remote Desktop** check boxes on the **Remote** tab of the **System Properties** dialog box.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** By default, Remote Desktop for Administration is installed on computers running Windows Server 2003. To install Terminal Server, use the Configure Your Server Wizard or Add/Remove Windows Components.

**Startup type:** Manual

**Service status:** Started

**This service depends on the following system components:**

Infrared Monitor

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Terminal Services Session Directory (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Tssdis

**Executable Name:** Tssdis.exe

**Log On As:** LocalSystem

**Description:** The Terminal Services Session Directory service provides a multisession environment that allows client devices to access a virtual Windows desktop session and Windows-based programs running on a Windows Server 2003–based computer.

The Terminal Services Session Directory allows clusters of load-balanced terminal servers (servers that have Terminal Server enabled) to properly route a user's connection request to the server where the user already has a session running. Users are routed to the first-available terminal server, regardless of whether they have a running session elsewhere in the cluster. Load balancing pools the processing resources of several servers by using the TCP/IP networking protocol. You can use this service with a cluster of terminal servers to scale the performance of a single terminal server by distributing sessions across multiple servers. Session Directory keeps track of disconnected sessions on the cluster, and ensures that users are reconnected to those sessions.

If this service is stopped or disabled, load balancing for Terminal Services does not work.

Connection requests are routed to the first available server. Users are routed to the first-available terminal server, regardless of whether they have a running session elsewhere in the cluster.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Themes (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Themes

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Themes service provides user-experience theme-management services and rendering support for the new Windows Server 2003 graphical user interface (GUI). A desktop theme is a predefined set of icons, fonts, colors, sounds, and other window elements that give the computer desktop a unified and distinctive look. You can select from different themes, create your own theme by modifying a theme and then saving it with a new name, or restore the traditional Windows Classic look as your theme.

If this service is stopped or disabled, the new Windows Server 2003 visual style (windows, buttons, scrollbars, and other controls) reverts to the Windows Classic visual style.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

None

**The following system components depend on this service:**

None

# Trivial FTP Daemon (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** tftpd

**Executable Name:** tftpd.exe

**Log On As:** LocalSystem

**Description:** The Trivial FTP Daemon (TFTPD) service implements support for the Trivial File Transfer Protocol (TFTP) defined by the following RFCs:

- RFC 1350, "The TFTP Protocol (Revision 2)"
- RFC 2347, "TFTP Option Extension"
- RFC 2348, "TFTP Blocksize Option"
- RFC 2349, "TFTP Timeout Interval and Transfer Size Options"

To find RFC standards, see the Request For Comments (RFC) link on the Web Resources page (http://go.microsoft.com/fwlink/?linkid=291). TFTP does not require a user name or password and is an integral part of the Remote Installation Services (RIS).

A RIS server uses the TFTPD service to download the initial files required for the remote installation process to begin. The most common file downloaded to the client by using the TFTPD service is Startrom.com, which is responsible for bootstrapping the client computer. If the user then presses F12 when prompted, the Client Installation Wizard is downloaded to begin the remote installation process.

If this service is stopped, the server fails to reply to the client request.

If this service is disabled by using the Services snap-in, this server fails to send the initial files to the client computer.

The correct way to disable this service is to uninstall RIS.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Add/Remove Windows Components and Remote Installation Services. (When you install RIS, the Trivial FTP Daemon service is automatically installed.)

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

IPSEC Driver

TCP/IP Protocol Driver

**The following system components depend on this service:**

None

# Uninterruptible Power Supply (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** UPS

**Executable Name:** Ups.exe

**Log On As:** LocalService

**Description:** The Uninterruptible Power Supply service manages an uninterruptible power supply (UPS) connected to your computer by a serial port.

| **Note** |
| --- |
| If you have a USB UPS, you should not install this service. |

By default the startup type is Manual and the default status is Stopped, unless you install and configure a serial UPS. After you install and configure a serial UPS, the startup type changes to Automatic and the default status changes to Started.

If this service is stopped or disabled, communications with the UPS are lost. If power is lost on the AC line, the UPS is unable to pilot the computer to shut down as the UPS battery discharges towards a critically low state. This could potentially result in loss of data.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Automatic, if you install and configure a serial UPS on the **UPS** tab of the **Power Options Properties** dialog box.

**Service status:** Started

**This service depends on the following system components:**

None

**The following system components depend on this service:**

None

# Upload Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** uploadmgr

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Upload Manager service manages synchronous and asynchronous file transfers between clients and servers on the network. Driver data is anonymously uploaded from customer computers to Microsoft and then used to help users find the drivers required for their systems. The Microsoft Driver Feedback Server asks the client's permission to upload the computer's hardware profile and then search the Internet for information about how to obtain the appropriate driver or get support from Microsoft or a third party.

Information uploaded from your computer for the purposes of finding driver information includes the hardware identification numbers for the device, the time at which the Windows Hardware Wizard was finished, and an ID for the Windows operating system you are running. The system information uploaded cannot be tracked to a user, computer, corporation, IP address, or any other source information.

The collected data is used to track which devices do not have easily obtainable drivers. If additional device driver information exists, it will be available after the device identification number is uploaded. If additional driver information is not available, Microsoft records the device's identification number to work with hardware vendors to increase the availability of drivers necessary to run devices on Windows, or provide information about driver availability and device support.

If this service is stopped or disabled, synchronous and asynchronous file transfers between clients and servers on the network do not occur.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Manual on Windows Server 2003, Standard Edition and Windows Server 2003, Enterprise Edition

Disabled on Windows Server 2003, Datacenter Edition and on Windows Server 2003, Web Edition

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Virtual Disk Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** VDS

**Executable Name:** Vds.exe

**Log On As:** LocalSystem

**Description:** Virtual Disk Service (VDS) provides a single interface for managing block storage virtualization, whether done in operating system software, redundant array of independent disks (RAID) storage hardware subsystems, or other virtualization engines.

VDS provides a vendor-neutral and technology-neutral interface for the management of logical volumes (in software) and logical units (in hardware). Management operations include binding, topology discovery and tracking, volume status and fault tracking, and performance monitoring.

Virtual disks should not be confused with snapshots. Unlike the Volume Shadow Copy service, VDS does not coordinate with applications or the file system. Thus, the data contained on a volume will not be synchronized prior to a volume or disk configuration operation. VDS might be used to configure a mirror plex, but a snapshot provider is necessary to perform the coordination when removing the plex and surfacing the snapshot. That use is outside the scope of this document with two exceptions:

- VDS coordinates with the file system prior to extending or shrinking volumes.
- Full-copy snapshots appear as plexes to VDS.

If this service is stopped or disabled, VDS is no longer available.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

Plug and Play

**The following system components depend on this service:**

None

# Volume Shadow Copy (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** VSS

**Executable Name:** Vssvc.exe

**Log On As:** LocalSystem

**Description:** The Volume Shadow Copy service manages volume shadow and is used by backup applications. When a backup application attempts to start a backup by using the new volume shadow infrastructure, the backup application calls methods to determine the number of writers that are running on the service. It queries each writer to gather the required metadata. After it has obtained the metadata, the backup application can collect the volumes that need to be shadow copied to ensure a successful backup session.

Volume Shadow Copy allows you to create shadow copy backups of volumes — exact copies of files, including all open files. For example, databases that are held open exclusively and files that are open due to operator or system activity are backed up during a volume shadow copy backup. In this way, files that have changed during the backup process are copied correctly.

Shadow copy backups ensure that:

- Applications can continue to write data to the volume during a backup.
- Files that are open are not omitted during a backup.
- Backups can be performed at any time, without locking out users.

Some applications manage storage consistency differently while files are open, which can affect the consistency of the files in the backup. For critical applications, consult the application vendor.

If this service is stopped or disabled, shadow copies are unavailable for backup and the backup process might not succeed.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Maunal

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# WebClient (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** WebClient

**Executable Name:** svchost.exe -k LocalService

**Log On As:** LocalService

**Description:** The WebClient service allows Win32 applications to access documents on the Internet. It extends the networking capability of Windows by allowing standard Win32 applications to create, read, and write files on Internet file servers by using the Web Document Authoring & Versioning (WebDAV) protocol.

The WebDAV protocol is a file access protocol described in XML that travels over HTTP. By using standard HTTP, WebDAV runs over an existing Internet infrastructure (for example, firewalls and routers).

If this service is stopped or disabled, you cannot use the Web Publishing Wizard to publish data to the Internet for locations that use the WebDAV protocol.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

WebDav Client Redirector

**The following system components depend on this service:**

None

# Web Element Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** ELEMENTMGR

**Executable Name:** Elementmgr.exe

**Log On As:** LocalSystem

**Description:** The Web Element Manager service is responsible for serving Web user interface elements for the Administration Web site at port 8098. These elements are used to determine the following information:

- Tabs to display on the administration Web site
- Remote Administration tasks available to a member of the local Administrators group on the Server Appliance
- Table of Contents
- Help topics
- Remote Administration alerts that can be displayed

A member of the local Administrators group can remotely manage a server by connecting to the server at https://*servername*:8098. When a connection is received on this Web site, the default Active Server Pages (ASP) code queries the Web Element Manager service on the Microsoft® Windows® Powered Web server appliance for each type of information listed above. After all the information has been collected, the appropriate Web page is displayed to the administrator.

The Web Element Manager service loads up all the information at startup time; the client (in this case, the ASP code) requests the Web user interface elements through a COM interface. The service runs in the LocalSystem account and requests through the COM interface are accepted only from clients running under the Administrator or LocalSystem accounts. If the service is stopped or its startup type is Manual, this service will start when the next request for Web user interface elements is received.

If this service is stopped, the service is automatically restarted when the Web interface for remote administration is accessed.

If this service is disabled, the Remote Administration Tools Web user interface for server administration does not function properly.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Windows Audio (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** AudioSrv

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Windows Audio service provides support for sound and related Windows Audio event functions. This service manages Plug and Play events for audio devices such as sound cards and global audio effects (GFX) for Windows audio application program interfaces. Examples of GFXs are equalization (EQ), bass enhancement, and speaker correction. The service handles the loading, unloading, and state save/restore of the GFXs on a per-session basis.

Through the Multimedia item under Components in System Information, the user can do the following:

- Enable or disable a GFX.
- Select among several GFX filters if more than one GFX is available that is designed for the audio hardware in question. A GFX driver's INF file specifies the target hardware for the GFX.

This service cannot be stopped after it has started. If this service is disabled, audio functionality might be unable to provide sound or process GFXs.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Automatic on Windows Server 2003, Standard Edition

Disabled on Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows Server 2003, Web Edition

**Service status:** Started

**This service depends on the following system components:**

Plug and Play

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Windows Image Acquisition (WIA) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** StiSvc

**Executable Name:** svchost.exe -k imgsvc

**Log On As:** LocalService

**Description:** The Windows Image Acquisition (WIA) service provides image acquisition services for scanners and cameras.

Windows supports still-image devices through the WIA service, which uses the Windows Driver Model (WDM) architecture. The WIA service provides robust communication between applications and image-capture devices, allowing you to capture images efficiently and transfer them to your computer for editing and use. The service is needed to capture events generated by imaging devices.

The WIA service supports small computer system interface (SCSI), IEEE 1394, USB, and serial digital still-image devices. Support for infrared, parallel, and serial still-image devices — which are connected to standard COM ports — comes from the existing infrared, parallel, and serial interfaces. Image scanners and digital cameras are examples of still-image devices. The WIA service also supports Microsoft® DirectShow®–based Webcams and digital video (DV) camcorders to capture frames from video.

If this service is stopped or disabled, events from imaging devices are not captured and processed. If a WIA device is installed on the computer, the service will restart automatically when you restart the computer. Also, the service restarts any time a WIA-enabled application is started.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Disabled

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

Shell Hardware Detection

**The following system components depend on this service:**

None

# Windows Installer (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** MSIServer

**Executable Name:** msiexec.exe /V

**Log On As:** LocalSystem

**Description:** The Windows Installer service manages the installation and removal of applications by applying a set of centrally defined setup rules during the installation process. These setup rules define the installation and configuration of the installed application. In addition, this service can modify, repair, or remove an existing application. The Windows Installer technology consists of the Windows Installer service for the Windows operating systems and the package (.msi) file format used to hold information regarding the application setup and installations.

Windows Installer is not only an installation program; it is also an extensible software management system. Windows Installer manages the installation, addition, and deletion of software components, monitors file resiliency, and maintains basic disaster recovery by means of rollbacks. Additionally, Windows Installer supports installing and running software from multiple sources, and can be customized by developers who want to install custom applications.

If startup type is Manual, an application that uses the installer will start this service.

If this service is stopped or disabled, the installation, removal, repair, and modification of applications that use Windows Installer do not succeed. A number of applications make use of this service while they are running, and those applications might not work correctly.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Windows Internet Name Service (WINS) (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** WINS

**Executable Name:** Wins.exe

**Log On As:** LocalSystem

**Description:** Windows Internet Name Service (WINS) enables NetBIOS name resolution. The presence of the WINS server(s) is crucial for locating the network resources identified by using NetBIOS names. WINS servers are required unless all domains have been migrated to Active Directory and all servers on the network are running Windows Server 2003.

| **Note** |
| --- |
| When performing a clean installation of any member of the Windows Server 2003 family, you have to explicitly select WINS in order to have it installed on the computer. By default, WINS is not installed. |

For more information about WINS, see the Networking Guide of the Windows Server 2003 Resource Kit (or see the Networking Guide on the Web at (http://technet.microsoft.com/en-us/windowsserver/bb633748.aspx)).

If this service is stopped or disabled, the following occurs:

- Locating Windows NT 4.0 domains and domain controllers does not succeed.
- Locating Windows 2000 or Windows Server 2003 Active Directory domains and domain controllers from Windows NT 4.0 clients does not succeed.
- NetBIOS name resolution fails, unless the device whose name must be resolved is on the same subnet as the device attempting name resolution. The device must be configured to attempt NetBIOS name resolution by using broadcasts.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Configure Your Server Wizard, or go to Add/Remove Windows Components, Networking Services, and Windows Internet Name Service (WINS)

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

COM+Event System

Event Log

NTLM Security Support Provider

Remote Procedure Call (RPC)

Security Accounts Manager

**The following system components depend on this service:**

None

# Windows Management Instrumentation (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** winmgmt

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Windows Management Instrumentation (WMI) service provides a common interface and object model to access management information about operating system, devices, applications, and services. WMI is an infrastructure for building management applications and instrumentation shipped as part of the current generation of Microsoft operating systems.

The WMI infrastructure is a Windows operating system component that moves and stores information about managed objects. The WMI infrastructure is made of two components: the WMI service and the WMI repository. The WMI service acts as an intermediary between the providers, management applications, and the WMI repository, placing information from a provider into the WMI repository. The WMI service also accesses the WMI repository in response to queries and instructions from management applications. Finally, the WMI service can pass information directly between a provider and a management application. In contrast, the WMI repository acts as a storage area for information passed in by the various providers.

WMI provides accessible information about applications and services, including management events those applications and services might generate.

WMI provides access to the management data through a number of interfaces, including COM API, scripts, and command-line interfaces. WMI is compatible with previous management interfaces and protocols, such as SNMP.

If this service is stopped or disabled, most Windows-based software does not function properly and management data is unavailable.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Event Log

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Windows Management Instrumentation Driver Extensions (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** Wmi

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Windows Management Instrumentation Driver Extensions service monitors all drivers and event trace providers that are configured to publish WMI or event trace information.

| **Note** |
| --- |
| In Windows Server 2003, the Windows Management Instrumentation Driver Extensions service is primarily a stub service that was created to deal with Windows 2000 service compatibility issues. |

If this service is stopped or disabled, any services that explicitly depend on it will fail to start.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

None

**The following system components depend on this service:**

None

# Windows Media Services (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** WMServer

**Executable Name:** Wmserver.exe

**Log On As:** NetworkService

**Description:** Windows Media Services provides streaming media services over IP-based networks. It replaces the four separate services that comprised Windows Media Services versions 4.0 and 4.1:

- Windows Media Monitor Service
- Windows Media Program Service
- Windows Media Station Service
- Windows Media Unicast Service

Windows Media Services is now a single service that runs on the Windows Server 2003 family. Its core components were developed by using the Component Object Model (COM), creating a flexible architecture that is easily customized for specific applications. It supports a greater variety of control protocols, including Real Time Streaming Protocol (RTSP), Microsoft Media Server (MMS) protocol, and HTTP.

The Windows Media Service platform is compliant with the following industry standards:

- WMI for server event notification and messaging
- SNMP for networking components
- XML, Synchronized Multimedia Integration Language (SMIL) 2.0, and the Document Object Model (DOM) for playlist implementation
- MPEG 1 and 2 for audio and video formats

Most streaming scenarios can be accomplished by using the core components installed with Windows Media Services. However, more advanced scenarios might require that you incorporate some custom programming and integration work. For developers and systems integrators, the Windows Media Services Software Development Kit (SDK) provides access to all elements of the server through a combination of plug-ins, a fully documented object model, and a rich set of external event notifications, all designed to be easily customizable.

If this service is stopped or disabled, streaming media services might be unavailable.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Configure Your Server Wizard or Manage Your Server, or go to Add/Remove Windows Components and Windows Media Services.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Windows System Resource Manager (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** WindowsSystemResourceManager

**Executable Name:** Wrm.exe

**Log On As:** LocalSystem

**Description:** The Windows System Resource Manager (WSRM) service is a tool to help deploy applications into consolidation scenarios. It provides policy-based management of CPU and memory consumption of processes running on a single operating system instance. Planned scenarios include multiple heterogeneous server applications, multiple Terminal Services users, multiple instances of SQL Server, multiple IIS 6.0 application pools, or Exchange Server and IIS 6.0 running together on the same computer.

The WSRM service enables a system administrator to do the following:

- Configure CPU and memory allocation policy settings on applications. This includes selecting processes to be managed and setting resource usage targets or limits.
- Manage CPU utilization (percent CPU in use).
- Limit the process working set size (physical resident pages in use).
- Manage committed memory (paging file usage).
- Apply policy settings to users or groups on a Terminal Services application server.
- Apply policy settings on a Date and Time schedule.
- Generate, store, view, and export resource utilization accounting records for management, service level agreement (SLA) tracking, and charge-back purposes.

The WSRM service maintains an updatable exclusion list of processes that should not be managed because of the negative system impact such management could create. The WSRM service also applies limits to process working set size and committed memory consumption. The WSRM service does not manage address windowing extensions (AWE) memory, large page memory, locked memory, or operating system pool memory.

The primary option for CPU management is bandwidth targets (expressed as a percent of system CPU). Targets are maintained by dynamically monitoring and adjusting process priorities. The WSRM service also provides affinity management, which is provided by using per-process APIs for hard affinity.

Memory management options include working set limits and committed memory maximum applied per process. Working set limits are set in policy and applied by the WSRM service by means of a kernel API. Subsequently, the kernel memory manager will apply and maintain the limits on working set size by paging the process as necessary. Committed memory is simply monitored against an upper limit. When the upper limit is exceeded, the process is stopped or an event is logged (at the user's discretion).

Additional features include full calendar features for scheduling desired policies, sophisticated pattern matching for identifying processes at runtime, WSRM service–specific counters, and a basic job accounting system.

The WSRM service must be installed on each managed computer. For administration of the service, an MMC snap-in and command-line programs are provided. These client components can be installed and run on any Windows 2000, Windows XP Professional, or Windows Server 2003 operating system. The service can be installed and run only on Windows Server 2003, Datacenter Edition and Windows Server 2003, Enterprise Edition.

If this service is stopped or disabled:

- You will not receive any management (WSRM provides runaway process protection by default).
- You will not receive any job accounting data.
- You will not be able to administer WSRM by means of the snap-in or CLI.

**Available on:** Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition

**Installed through:** Windows System Resource Management CD

**Startup type:** Automatic

**Service status:** Started or Stopped based on user check box in setup

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Windows Time (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** W32Time

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Windows Time service maintains date and time synchronization on all computers running on a Windows-based network. It uses the Network Time Protocol (NTP) to synchronize computer clocks so that an accurate clock value, or time stamp, can be assigned to network validation and resource access requests. The implementation of NTP and the integration of time providers make the Windows Time service a reliable and scalable time service for enterprise administrators. For computers that are not joined to a domain, the Windows Time service can be configured to synchronize time with an external time source.

If this service is stopped or disabled, date and time synchronization is unavailable in the forest or an external NTP server.

There are two scenarios:

1. Stopping the Windows Time service on a workstation prevents the workstation from synchronizing its time with another source, but has no effect on any other external server.

2. Stopping the Windows Time service on a domain controller has the same effect as in item 1, but it also prevents any domain members from synchronizing time with it. This might adversely impact time synchronization in the enterprise.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

None

**The following system components depend on this service:**

Cluster Service

# WinHTTP Web Proxy Auto-Discovery Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** WinHttpAutoSvc

**Executable Name:** svchost.exe –k LocalService

**Log On As:** LocalService

**Description:** WinHTTP Web Proxy Auto-Discovery Service (WinHttpAutoSvc) implements proxy configuration discovery for WinHTTP clients. It performs the lookup done by WinHttpGetProxyForUrl() as documented in Microsoft® MSDN®. Applications use the service to isolate the running of an autoproxy script in a sandbox from the client process.

The WinHttpGetProxyForUrl() function implements the Web Proxy Auto-Discovery (WPAD) protocol for automatically configuring the proxy settings for an HTTP request. The WPAD protocol downloads a Proxy Auto-Configuration (PAC) file, which is a script that identifies the proxy server to use for a given target URL. PAC files are typically deployed by the IT department within a corporate network environment.

To make configuring the proxy settings for WinHTTP-based applications easier, WinHTTP now implements the Web Proxy Auto-Discovery (WPAD) protocol, often referred to as autoproxy. This is the same protocol that Web browsers implement to automatically discover the proxy configuration without requiring an end user to specify a proxy server manually. This feature is available starting with WinHTTP version 5.1 in Windows 2000 Service Pack 3, Windows XP Service Pack 1, and Windows Server 2003.

Stopping or disabling this service might cause WinHTTP clients to be unable to automatically detect proxy settings, and might cause other WinHTTP clients to automatically detect proxy settings within their own process.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Started

**This service depends on the following system components:**

AFD Networking Support Environment

DHCP Client

IPSEC Driver

TCP/IP Protocol Driver

**The following system components depend on this service:**

None

# Wireless Configuration (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** WZCSVC

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Wireless Configuration service enables automatic configuration for IEEE 802.11 wireless adapters for wireless communications.

Microsoft partnered with 802.11 network adapter vendors to improve the wireless roaming experience on Windows by automating the process of configuring the network adapter to associate with an available network.

The wireless network adapter and its Network Driver Interface Specification (NDIS) driver need to do very little beyond supporting a few new NDIS Object Identifiers (OIDs) used for the querying and setting of device and driver behavior. The network adapter scans for available networks and passes those to a computer running Windows Server 2003. The wireless configuration service takes care of configuring the network adapter with an available network. In the case in which there are two networks covering the same area, the user can configure a preferred network order and the computer will try each network in order until it finds an active one. It is also possible to limit association to only the configured, preferred networks.

If this service is stopped or disabled, automatic wireless configuration is unavailable.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Automatic on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Manual on Web Edition

**Service status:** Started

**This service depends on the following system components:**

NDIS Usermode I/O Protocol

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# WMI Performance Adapter (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** WmiApSrv

**Executable Name:** Wmiapsrv.exe

**Log On As:** LocalSystem

**Description:** The WMI Performance Adapter service provides performance library information from WMI HiPerf providers. Applications and services that need to provide performance counters today can do so in two ways: by writing a WMI High Performance provider, or by writing a performance library. Consumers of high performance data also have two ways to request performance data: through WMI, or through the Performance Data Helper (PDH) APIs. There are mechanisms in place so that the two models interact together, and clients accessing counters through each model can still see the counters provided by the other model. The Reverse Adapter is one of those mechanisms.

The WMI Performance Adapter service transforms performance counters supplied by WMI High Performance providers into counters that can be consumed by PDH through the Reverse Adapter Performance Library. This way, PDH clients (for example, System Monitor) can consume performance counters surfaced by any WMI High Performance providers on the computer.

The startup type is Manual and the service is not running by default. It runs on demand when a performance client (for example, System Monitor) uses PDH to query performance data. After the client disconnects, the service stops.

This service runs on demand when a performance client, (for example, System Monitor) uses PDH to query performance data. After the client disconnects, the service stops.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Manual

**Service status:** Stopped

**This service depends on the following system components:**

Remote Procedure Call (RPC)

**The following system components depend on this service:**

None

# Workstation (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** lanmanworkstation

**Executable Name:** svchost.exe -k netsvcs

**Log On As:** LocalSystem

**Description:** The Workstation service creates and maintains client network connections to remote servers. The Workstation service is a user-mode layer for the Microsoft Networks redirector. It loads and performs configuration functions for the redirector, provides support for making network connections to remote servers, provides support for the WNet APIs, and furnishes redirector statistics.

If this service is stopped, you cannot establish connections to remote servers and access files through named-pipes. This prevents you from accessing files and printers located on other remote computers, but does not affect TCP/HTTP connectivity. Internet browsing and Web Client access will still work.

If this service is disabled, any services that explicitly depend on this service will fail to start.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** Default operating system installation

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

None

**The following system components depend on this service:**

Alerter

Computer Browser

Distributed File System

File Server for Macintosh

Messenger

Net Logon

Remote Procedure Call (RPC) Locator

# World Wide Web Publishing Service (System Services for the Windows Server 2003 Family and Windows XP Operating Systems)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service Name:** W3SVC

**Executable Name:** svchost.exe -k iissvcs

**Log On As:** LocalSystem

**Description:** World Wide Web Publishing Service provides Web connectivity and administration through IIS Manager.

World Wide Web Publishing Service provides HTTP services for applications on the Windows platform. The service contains a process manager and a configuration manager. The process manager controls the processes in which custom applications and simple Web sites reside. The configuration manager reads the stored system configuration and ensures that Windows is configured to route HTTP requests to the appropriate application pools or operating system processes.

This service can monitor the processes in which custom applications reside and provide recycling services for these applications. Recycling is a configuration property of an application pool and can be done on the basis of memory limits, request limits, processing time, or time of day. The service will queue HTTP requests if custom applications stop responding, and will also attempt to restart custom applications.

The service depends on IIS Admin Service and kernel TCP/IP support.

| **Note** |
| --- |
| In Windows 2000 Server, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server, IIS was installed by default. This has changed on Windows Server 2003: IIS is installed by default only on Windows Server 2003, Web Edition. |

If this service is stopped or disabled, the Windows Server 2003 operating system will not be able to serve Web pages or any form of Web request.

**Available on:** Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition and Windows Server 2003, Web Edition

**Installed through:** By default on Windows Server 2003, Web Edition.

In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Server Edition, use the Configure Your Server Wizard, or go to Add/Remove Windows Components, Application Server, Internet Information Services (IIS), and World Wide Web Service.

**Startup type:** Automatic

**Service status:** Started

**This service depends on the following system components:**

HTTP SSL

IIS Admin Service

Remote Procedure Call (RPC)

Security Accounts Manager

**The following system components depend on this service:**

None

# Server and Domain Isolation

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Server and domain isolation enable you to implement network authentication based on policy rather than network topology, so that your Windows environment is dynamically segmented into more secure, isolated logical networks.

By using server and domain isolation, you can better:

- Reduce the scope of network security threats
- Safeguard sensitive data and intellectual property
- Extend the value of existing infrastructure

Server and domain isolation builds on Internet Protocol security (IPsec) and Active Directory® Domain Services (AD°DS), which are available in Windows Vista®, Windows Server® 2008, Windows Server® 2003, Windows XP, and Windows 2000. This removes the need to make potentially disruptive changes to your network topology, existing applications, or requiring costly hardware upgrades.

**Available documentation**
The following resources can help you to effectively design, deploy, and manage server and domain isolation scenarios in your network.

- Introduction to Server and Domain Isolation

  This paper introduces the basic concepts of server and domain isolation, and describes common scenarios in which each can be used to improve network security.

- Server Isolation with Microsoft Windows Explained

  This paper provides a detailed overview of server isolation in the Microsoft Windows operating systems. It explains how server isolation protects servers and describes the benefits of deploying server isolation in your IT environment. It also provides a brief overview of how to deploy server isolation.

- Domain Isolation with Microsoft Windows Explained

  This paper provides a detailed overview of domain isolation in the Microsoft Windows operating systems. It explains how domain isolation protects domain member computers and describes the benefits of deploying domain isolation in your IT environment. It also provides a brief overview of how to deploy domain isolation.

# Introduction to Server and Domain Isolation

Applies To: Windows Server 2008, Windows Vista

The proliferation of networking technologies that allow computers to freely communicate has ushered in a new age of productivity for businesses, governments, and individuals. People have around-the-clock access to e-mail, messages, files, databases, and Web pages. They can access this data through multiple types of connections, including wired (such as Ethernet), wireless, and remote access.

Unfortunately, the ease and ubiquity of connection poses new risks. The same ease of connectivity that allows users to access networked resources at any time from almost anywhere also allows malicious programs (such as viruses and worms) and malicious users to attack computers or their resources at any time and from anywhere. Although your company needs to allow access to resources, this access should be permitted to only authenticated and authorized users and computers. Protecting your network requires a layered defense-in-depth security model. Your network must be isolated, not only from the Internet, but also from unauthorized and unmanaged computers on your intranet.

To isolate the authorized and managed computers from the other computers on your network, you can create an isolated network; a set of network nodes whose grouping is independent of the physical network topology. You can create an isolated network based on the Physical layer of the Open Systems Interconnection (OSI) model, in which you run a separate cabling system for the isolated network. However, this type of isolated network is costly and difficult to maintain for large networks. You can also create an isolated network based on the Data Link layer of the OSI model, in which you use Layer 2 switches and virtual LAN (VLAN) technology to create logical network segments by grouping computers regardless of their physical connection to a set of switches. With VLAN technology, you can also create an isolated network based on the Network layer of the OSI model, in which you create logical subnets and define the routing between the subnets. However, for VLAN-based isolated networks, your switching fabric must support VLANs and the logical segments or subnets must be maintained over time as the topology of the network changes and computers move to different switch ports.

With the Microsoft® Windows® operating systems, you can logically isolate your domain and server resources to limit access to authenticated and authorized computers. For example, you can create a logical network consisting of computers that share a common Windows-based security framework and a set of requirements for secure communication. Each computer on the logically isolated network can provide authentication credentials to the other computers on the isolated network to prove its membership. Requests for communication that originate from computers that are not part of the isolated network are ignored.

Windows-based network isolation occurs at the Network layer of the OSI model. Therefore, the isolated network can span hubs, switches, and routers across the physical and geographical boundaries of your organization network.

By leveraging the authentication infrastructure of Windows-based computers, you can create an isolated network that does not depend on a separate cabling system or on the features of your switching or routing fabric. Additionally, Windows-based isolated networks do not require ongoing maintenance based on changes in network topology or computers moving to different switch ports. The result is an isolated network that leverages your existing Windows infrastructure with no additional costs associated with network reconfiguration or hardware upgrades.

Figure 1 shows an example of an isolated network.




In Figure 1, the entire organization network is isolated from the Internet by a firewall, a proxy server, or other types of security systems. A subset of the network's computers is located on the isolated network. The computers on the isolated network are protected from the other computers on the organization network.

For example, many types of viruses and worms cannot propagate into an isolated network. Malicious users and software from outside the isolated network cannot successfully attack isolated network computers because they lack the authentication credentials with which to establish communication.

The security requirements of the isolated network can optionally include data encryption. By requiring data encryption for the traffic exchanged between isolated network members, you can satisfy business partner and regulatory requirements to encrypt

data when it traverses a network.

**Server and Domain Isolation Solutions**

To create an isolated network, you need to separate the various types of computers on the organization network according to the type of access you want the computers to have. The communication requirements are the following:

- Computers on the isolated network can initiate communications with all of the computers on the organization network, including those that are not located on the isolated network.

- Computers that are not on the isolated network can initiate communications only with other computers that are not on the isolated network. They cannot initiate communications with computers on the isolated network.

Computers on the isolated network will ignore all requests to initiate communication from computers that are not on the isolated network. This helps provide protection against malicious programs and users outside the isolated network and decreases the exposure of computers on the isolated network to various types of network attacks (such as port scanning) and malicious software (such as viruses and worms). In addition, you can specify that traffic sent between the computers on the isolated network be encrypted.

To create an isolated network, you need the following elements:

- Credentials

  The computers on the isolated network use security credentials when initiating communications to prove their identity and authenticate themselves to the other computers on the isolated network.

- Network policy settings

  The computers on the isolated network use network policy settings to require authentication for incoming communication requests, to secure communications and, if needed, to provide encryption.

In an organization network, it is impractical to configure and administer credentials and network policy settings locally on individual computers. To create an isolated network that can be centrally configured and that scales as your organization grows, you need to manage the computers on your organization network. To centrally manage your network, you need an authentication and policy distribution infrastructure.

On a Windows-based network that is managed by a central IT department, computers use the authentication infrastructure and policy distribution mechanism to implement Server and Domain Isolation. A managed environment not only makes Server and Domain Isolation easier, it also provides the infrastructure that allows you to configure other security activities, such as distributing operating system updates and installing antivirus software.

In Windows, a computer can be a member of an Active Directory® directory service domain. During the process of joining an Active Directory domain, the computer is issued a set of credentials. On an ongoing basis after joining the domain, the computer receives centrally configured network policies through Group Policy. Network administrators use Group Policy to distribute computer and user settings to the member computers of an Active Directory domain. With the appropriate Group Policy settings to require authentication before communication, a domain member computer sends its credentials to authenticate a communication attempt, which can then be verified by any domain controller.

By leveraging Active Directory membership and Group Policy settings, everything that you need to create an isolated network is already available on computers running Windows Vista®, Windows Server® 2008 (currently in beta), Microsoft® Windows® XP, Microsoft® Windows Server® 2003, and Microsoft® Windows® 2000 Server operating systems. All that you need to do is to ensure that computers are members of your domain and to configure the appropriate Group Policy settings to require authentication for incoming communication attempts, to secure data traffic, and optionally, to encrypt data traffic. After you have configured and applied the appropriate Group Policy settings, you add a new computer to the isolated network by making it a member of the Active Directory domain. In contrast to cable or VLAN-based isolated networks, no new hardware or maintenance of the cabling system or switching fabric is required.

Figure 2 shows an example of an isolated network using an Active Directory domain.




In Figure 2, the computers on the isolated network are members of an Active Directory domain, which includes computers that are locally connected to the organization network—through Ethernet or Institute of Electrical and Electronic Engineering (IEEE) 802.11 wireless LAN connections, for example—or are remotely connected to the organization network, through a remote access dial-up connection or a virtual private network (VPN) connection across the Internet (not shown). The computers on the organization network that are not part of the isolated network include stand-alone computers running Windows—such as those that are members of workgroups or other untrusted domains—or other computers that do not support Active Directory, such as Apple Macintosh computers or UNIX-based computers.

If business operational requirements exist, options are available to enable non-Windows hosts to participate in the isolated network. For example, workstations and departmental servers running Linux, Mac OS X, and Solaris can be configured to communicate directly with an isolated domain. Alternatively, using ISA Server as an IPsec proxy, communication to isolated domains can also be configured for systems that do not natively support IPsec or for mainframes for which IPsec is not normally implemented.

Server and Domain Isolation provides an extra layer of security and access control that compliments other host- and network-based security technologies, such as antivirus, anti-spyware, firewall, 802.1X and intrusion detection, to enable greater resiliency in the presence of network security threats. This solution also further extends the value of Active Directory and the operational efficiency benefits of Group Policy. Logical isolation policies are created distributed and managed centrally using Active Directory Group Policy and existing Active Directory-based credentials (e.g. Kerberos or X.509 certificates). This results in a zero-touch deployment experience for IT administrators and an unchanged experience for end-users. No additional end-user training is required nor is there a need to install new software or visit each computer during deployment.

After you have deployed an Active Directory domain, you can configure Server and Domain Isolation.

**Domain Isolation**

To isolate a domain, you use an Active Directory domain and domain membership to enforce the following network policy: domain member computers accept only authenticated and secured communications from other domain member computers. This network policy isolates domain member computers from non-domain-member computers. With Domain Isolation, the isolated network consists, in the majority of cases, of the set of computers that belong to an Active Directory domain, as Figure 2 shows.

To configure Domain Isolation, use Group Policy settings to require that all incoming communication requests be authenticated using Active Directory and that domain member computers can initiate unprotected communications with non-domain-member computers. Optionally, you can require that all communication within the isolated domain be encrypted. You can also configure exemptions so that specific trusted computers that are not domain members can initiate unprotected communications with domain member computers.

Domain Isolation provides many benefits by:

- Restricting incoming communications to domain member computers.

  Domain member computers use their domain credentials to authenticate communication attempts and network policy settings to secure traffic with each other. This helps mitigate the risk of rogue or unmanaged devices from exploiting potentially unpatched vulnerabilities, propagating malicious software (malware) threats like viruses, worms or spyware, or disrupting business operations through denial-of-service (DoS) attacks. Non-domain-member computers do not have domain credentials and, therefore, cannot authenticate communication attempts with protected computers.

- Supplementing other security mechanisms designed to prevent unwanted communications.

  Domain Isolation provides end-to-end security that supplements the security mechanisms already deployed on your network, providing defense-in-depth. For example, if you deployed Domain Isolation and your firewall was compromised, malicious Internet users could not directly initiate communications with protected computers.

- Encouraging domain membership.

  By placing critical organization servers, such as database servers, on the isolated network, you prevent users on the organization network from connecting to them from a non-domain-member computer. To receive valid domain credentials for communicating with organization servers, computers must be joined to the domain. After a computer has been joined to the domain, you can manage it in other ways, such as by ensuring that it has the latest operating system and antivirus updates. These proactive measures also help increase system reliability and security, while reducing the risk of network-based attacks and lowering ongoing maintenance costs.

- Securing traffic between domain member computers.

  Traffic sent between domain member computers is secured so that the receiving computer can verify that an authenticated computer sent the packet and that the packet was not modified in transit. Optionally, the traffic between domain member computers can be encrypted, providing protection from malicious users on your organization network who attempt to capture and interpret network traffic.

For more information about Domain Isolation, see "Domain Isolation with Microsoft Windows Explained" at http://go.microsoft.com/fwlink/?LinkId=44642. For examples of how customers have deployed Domain Isolation, see the "Case Studies"section at http://go.microsoft.com/fwlink/?LinkId=79428.

**Server Isolation**

To isolate a specific server or servers, sensitive data, and associated clients, you use an Active Directory domain and domain membership to enforce the following network policy: specific server computers that are domain members accept authenticated and secured communications only from other domain member computers. This network policy isolates specific servers from non-domain-member computers. For example, to protect database traffic, you would configure and deploy Server Isolation Group Policy settings to require secured traffic between domain member client computers and their database servers. With Server Isolation, the isolated network consists of the server computers and domain member client computers, both of which belong to an Active Directory domain.

You can also create the following group-specific server isolation network policy: specific server computers that are domain members will accept authenticated and secured communications only from other domain member computers that are members of specific Active Directory security groups. Group-specific Server Isolation provides an additional layer of authorization and isolates specific servers from both non-domain-member computers and unauthorized domain member computers. Only an authorized domain member computer that has the business need can access the isolated server. With group-specific server isolation, the isolated network consists of the server computers and the group of authorized domain member client computers.

For example, you can configure group-specific Server Isolation settings so that a server that contains sensitive medical information allows secure communications only with computers that meet the following criteria:

- They are domain members.

- They are members of the ConfidentialMedical Active Directory security group.

After you have configured the appropriate Group Policy and server settings, you can allow a new computer to access this server by joining the computer to the domain and then adding the computer account of the new computer to the ConfidentialMedical Active Directory security group.

You can also use group-specific Server Isolation in conjunction with Domain Isolation to define and enforce tiered network access restrictions based on business objectives and policy versus network topology. For example, you can use Domain Isolation to better protect your entire Windows environment from rogue or unmanaged computers. Then, you can add another layer of protection through Server Isolation to further restrict access to specific networked resources to only the user requiring this access for business reasons.

Server Isolation provides many benefits by:

- Supplementing other security mechanisms designed to prevent unwanted communications.

- Encouraging domain membership.

  Non-domain-member computers cannot communicate with critical isolated servers unless they join the Active Directory

domain.

- Protecting traffic sent to and from isolated servers.

  The isolated server can verify that an authenticated computer sent the packet and that it was not modified in transit. Optionally, traffic to and from the isolated servers can be encrypted, providing protection from malicious users on your organization network who attempt to capture and interpret network traffic.

- Protecting applications that cannot protect themselves.

  Applications running on isolated servers that do not have facilities for enforcing access control or security can benefit from server isolation to enforce authentication, authorization, and communication security.

For more information about Server Isolation, see "Server Isolation with Microsoft Windows Explained" at http://go.microsoft.com/fwlink/?LinkId=44641.

**Overview of Server and Domain Isolation Documents**

This paper is the first in a series of papers that describes Server and Domain Isolation and provides guidelines for planning their deployment. The other papers include:

- "Domain Isolation with Microsoft Windows Explained" at http://go.microsoft.com/fwlink/?LinkId=44642

  This paper provides a detailed overview of Domain Isolation. It explains how Domain Isolation protects domain member computers and the benefits of deploying Domain Isolation. It also provides a brief overview of how to deploy Domain Isolation. This paper is intended for IT professionals in organizations that are investigating using the Microsoft implementation of Internet Protocol security (IPsec) in Windows to deploy Domain Isolation. It assumes that you are somewhat familiar with the Microsoft implementation of IPsec and would like more detailed information about using that technology to deploy Domain Isolation.

- "Server Isolation with Microsoft Windows Explained" at http://go.microsoft.com/fwlink/?LinkId=44641

  This paper provides a detailed overview of Server Isolation. It explains how Server Isolation protects isolated servers and the benefits of deploying Server Isolation. It also provides a brief overview of how to deploy Server Isolation. This paper is intended for IT professionals in organizations that are investigating using the Microsoft implementation of IPsec in Windows to deploy Server Isolation. It assumes that you are somewhat familiar with the Microsoft implementation of IPsec and would like more detailed information about using that technology to deploy Server Isolation.

- "Domain Isolation Planning Guide for IT Managers" at http://go.microsoft.com/fwlink/?LinkId=44645

  Designed for enterprise IT managers who are investigating using IPsec in Microsoft Windows to deploy Domain Isolation, this paper will help you and your IT staff to gather the information required to develop a Domain Isolation deployment plan and to design your IPsec policies. It includes an overview of the deployment process, a step-by-step guide to the planning process, and links to resources that you can use to plan and design your deployment. It does not explain how to deploy Domain Isolation.

- "A Guide to Domain Isolation for Security Architects" at http://go.microsoft.com/fwlink/?LinkId=44643

  Designed for security architects of enterprise organizations that are investigating the use of IPsec in Microsoft Windows to deploy Domain Isolation, this paper describes the implications of deploying Domain Isolation in an enterprise environment and explains how to assess the enterprise environment and plan Domain Isolation. Read this guide after you have developed a working knowledge of Domain Isolation.

- "Setting Up IPsec Server and Domain Isolation in a Test Lab" at http://go.microsoft.com/fwlink/?LinkId=44646

  This paper demonstrates how to set up IPsec Server and Domain Isolation in a limited test environment. It provides procedures for setting up a basic deployment, which you can use as the basis for your own deployment. This paper is designed for network architects who are investigating the use of IPsec in Microsoft Windows to deploy Server and Domain Isolation.

The size of your environment and IT staff will determine your planning and deployment requirements, which, in turn, will determine which of the server and domain deployment papers you will need to read.

**For Enterprises and Large Businesses**
For the purposes of these papers, an enterprise is a business or organization that has more than 100 servers and more than 500 desktop computers, comprises multiple organizational units, and has distributed branch offices. A large business is a business or organization that has more than 100 servers and more than 500 desktop computers and that might have multiple organizational units.

Typically, deployments in enterprises and large businesses require greater planning and more precisely defined processes. If you are planning Server and Domain Isolation deployment in an enterprise or large business, read the following papers in the order listed:

- "Domain Isolation with Microsoft Windows Explained" at http://go.microsoft.com/fwlink/?LinkId=44642
- "Server Isolation with Microsoft Windows Explained" at http://go.microsoft.com/fwlink/?LinkId=44641
- "Domain Isolation Planning Guide for IT Managers" http://go.microsoft.com/fwlink/?LinkId=44645
- "A Guide to Domain Isolation for Security Architects" at http://go.microsoft.com/fwlink/?LinkId=44643
- "Setting Up IPsec Server and Domain Isolation in a Test Lab" at http://go.microsoft.com/fwlink/?LinkId=44646

**For Medium and Small Businesses**
For the purposes of these papers, a medium business is a business or organization that has between 4 and 100 servers and 25 to 500 desktop computers, which might have only one organizational unit, and might have one or two distributed branch offices. A small business is a business or organization that has between 1 and 3 servers and 1 and 25 desktop computers, might not have a domain or has only one domain, and has no branch offices.

Typically, deployments in these organizations require less planning and less highly developed processes than larger deployments. If you are planning Server and Domain Isolation deployment in a medium or small business, read the following papers in the order listed:

- "Domain Isolation with Microsoft Windows Explained" at http://go.microsoft.com/fwlink/?LinkId=44642
- "Server Isolation with Microsoft Windows Explained" at http://go.microsoft.com/fwlink/?LinkId=44641
- "Domain Isolation Planning Guide for IT Managers" http://go.microsoft.com/fwlink/?LinkId=44645

Optionally, you might want to read the following papers:

- "A Guide to Domain Isolation for Security Architects" at http://go.microsoft.com/fwlink/?LinkId=44643
- "Setting Up IPsec Server and Domain Isolation in a Test Lab" http://go.microsoft.com/fwlink/?LinkId=44646

**See Also**
**Other Resources**
Server and Domain Isolation
Server and Domain Isolation Datasheet
Server and Domain Isolation Using IPsec and Group Policy
Improving Security with Domain Isolation
Windows Firewall Web site
Windows Server 2003

# Server Isolation with Microsoft Windows Explained

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As described in "Introduction to Server and Domain Isolation with Microsoft Windows," server isolation uses an Active Directory® directory service domain, domain membership, and Group Policy to enforce a network policy requiring specific server computers that are domain members to accept only authenticated and secured communications from other domain member computers. This network policy isolates specific servers from computers that are not domain members.

For example, sensitive data on servers is typically protected by access control security at the Application layer. Before accessing the files on a file server that contains sensitive data, a user must provide security credentials, such as a user name and a password. However, by specifying access control lists (ACLs) at the Application layer, you do not protect the server from attacks against other services running on the server or other types of network-level attacks—such as denial of service, viruses, or worms—launched from computers that are not domain members on your intranet.

To provide another layer of protection for servers that store sensitive data, you can isolate them from computers that are not domain members by implementing additional authentication and security at the Internet layer with Internet Protocol security (IPsec). By using IPsec-based server isolation, computers that are not domain members cannot initiate Internet Protocol (IP)-based communication with isolated servers. Organizations that send sensitive data over their networks and that must provide extra protection for sensitive data assets — such as those in the financial services, health care, or government sectors — require this additional level of protection.

Server isolation provides many benefits by:

- Restricting incoming communications to domain member computers.

  In the Microsoft® Windows® operating systems, you can manage a computer that is a member of an Active Directory domain by centrally configuring settings in Group Policy and applying those settings to all domain member computers. You can also apply other types of security updates, such as operating system updates and antivirus software signatures, to domain member computers. Domain member computers use their domain credentials to authenticate communication attempts with isolated servers. Computers that are not domain members — such as stand-alone, those that do not run Windows, unknown, and guest computers — do not have domain credentials and, therefore, cannot authenticate communication attempts with isolated servers.

- Supplementing security mechanisms designed to prevent unwanted communications.

  Server isolation supplements the security mechanisms already deployed on your network, such as ACLs, Institute of Electrical and Electronic Engineers (IEEE) 802.1X, Secure Sockets Layer (SSL), and firewalls. For example, if you isolated your servers and then your Internet firewall was compromised, malicious users from the Internet could not initiate direct communications with the isolated servers.

- Encouraging domain membership.

  By isolating critical organization servers, such as e-mail servers, you can prevent network users from accessing critical resources from a computer that is not a domain member. To receive valid domain credentials for authenticating with isolated servers, computers that are not domain members must join the domain. After being joined to the domain, these computers can be managed in other ways. For example, you can ensure that they have the latest operating system and antivirus updates.

- Protecting traffic sent to and from isolated servers.

  Traffic sent to and from isolated servers is cryptographically protected, so that the receiving computer can verify that an authenticated computer sent the packet and that the packet was not modified in transit. Optionally, the isolated server traffic can be encrypted, providing protection from malicious network users attempting to capture and interpret network traffic.

- Protecting applications that cannot protect themselves.

  Applications running on servers running Windows that cannot enforce access control or security at the Application layer can use server isolation to enforce authentication, authorization, and communication security at the Internet layer.

To isolate a server, you can configure Group Policy settings to require that all communication with isolated servers must be authenticated and protected by using IPsec. IPsec protects traffic from address spoofing, data injection, session hijacking, replay attacks, and other types of data tampering. Optionally, you can specify that packets are encrypted. You can also configure exceptions to specify that trusted computers, computers that are not domain members, or computers that are known as *exempted computers*, can initiate unprotected communications with isolated servers.

**How Server Isolation Works**

This section describes how Active Directory, Group Policy, and IPsec work together to implement server isolation.

**Components of Server Isolation**

To isolate a server, you need the following components:

- **An Active Directory domain.** A domain includes domain controllers and the appropriate trust relationships to establish trust with other domains or the directory trees of an organization network.
- **Member computers.** These are computers that have joined the Active Directory domain and received domain credentials.
- **Group Policy settings.** These computer and user settings are automatically downloaded to member computers.
- **Active IPsec policy settings.** These Group Policy settings determine the server isolation behavior of domain member computers.

In a simplified server isolation deployment, you configure and activate an IPsec policy with rules that define specific types of traffic and how the traffic should be handled. You then activate the IPsec policy for the appropriate Active Directory containers, such as sites, domains, and organizational units. The member computers in the Active Directory containers to which the Group Policy settings apply automatically download the Group Policy settings.

After the domain member computers have downloaded and applied the Group Policy settings, they have both the correct IPsec policy for server isolation and domain credentials that allow them to communicate securely with isolated servers. Computers that are not domain members, which do not have domain credentials or the correct IPsec settings, cannot initiate communications with isolated servers.

**Communication Processes**

When you implement server isolation by configuring domain member computers with the appropriate IPsec policy settings, communication between computers in your network differs depending on which type of computer (domain member or non-domain-member) initiates communication and which type of computer receives communication. This section describes how communication occurs:

- When a domain member computer initiates communication with an isolated server.
- When a computer that is not a domain member initiates communication with an isolated server.
- When an exempted computer initiates communication with an isolated server.

The following figure shows the types of communication that occur when you deploy server isolation.




**Legend**




**Communication with an isolated server initiated by a domain member computer**

When a domain member computer with both Active Directory credentials and server isolation IPsec policy settings (for example, COMPUTER1 in the figure) initiates communication with an isolated server (for example, SERVER1), the following occurs:

1. The initial communication packet sent by COMPUTER1—for example, a Transmission Control Protocol (TCP) Synchronize (SYN) segment destined for the IP address of an e-mail server that is isolated—matches the rule of the active IPsec policy that specifies that the initiating computer must secure the traffic with IPsec.

2. COMPUTER1 uses IPsec to perform mutual authentication with SERVER1 and to negotiate the use of IPsec protection.

3. Because both COMPUTER1 and SERVER1 have domain credentials, the IPsec authentication process succeeds. Because COMPUTER1 has IPsec policy settings that match SERVER1, negotiation of IPsec protection also succeeds.

4. COMPUTER1 sends the initial communication packet to SERVER1 with IPsec protection.

5. SERVER1 sends the response to the initial communication packet—for example, a TCP SYN-Acknowledgement (SYN-ACK) segment—to COMPUTER1 with IPsec protection.

6. Subsequent packets sent between COMPUTER1 and SERVER1 are protected by IPsec.

Domain member computers that have the appropriate IPsec policy settings authenticate and protect with IPsec communications to all isolated servers.

**Communication with an isolated server initiated by a non-domain-member computer**

When a non-domain-member computer (for example, COMPUTER2 in the figure) initiates communication with an isolated server (for example, SERVER1), the following occurs:

1. Because COMPUTER2 does not have IPsec policy settings, it sends its initial communication packet—for example, a TCP SYN segment—without IPsec protection to SERVER1.

2. On SERVER1, the initial communication packet sent by COMPUTER2 matches the IPsec policy rule for the server that requires IPsec protection for all incoming packets.

3. Because SERVER1 does not accept unprotected packets, it silently discards the TCP SYN segment sent by COMPUTER2.

4. SERVER1 also discards subsequent communication initiation packets sent by COMPUTER2.

5. Eventually, COMPUTER2 ends its attempt to communicate with SERVER1.

The IPsec policy on SERVER1 causes unprotected communications to be dropped. Non-domain-member computers that attempt to communicate with an isolated server never receive a response and, therefore, cannot connect to the server. Even if a user on a non-domain-member computer was able to duplicate the IPsec policy settings of a domain member computer, communications with SERVER1 would fail because the non-domain-member computer does not have valid domain

credentials.

**Communication with an isolated server initiated by an exempted computer**
When an exempted computer (for example, COMPUTER3 in the figure) initiates communication with an isolated server (for example, SERVER1), the following occurs:

1. Because COMPUTER3 does not have IPsec policy settings, it sends its initial communication packet—for example, a TCP SYN segment—without IPsec protection to SERVER1.

2. On SERVER1, the initial communications packet sent by COMPUTER3 matches the IPsec policy rule for the server that permits unsecured communication to its own IP address from the IP addresses of specific exempted computers.

3. SERVER1 sends a response—for example, a TCP SYN-ACK segment—to COMPUTER3 without IPsec protection.

4. Subsequent packets sent between COMPUTER3 and SERVER1 are sent without IPsec protection.

The specified IPsec policy allows exempted computers to communicate with isolated servers without requiring IPsec protection.

**Group-Specific Server Isolation**
The server isolation configuration described thus far allows all domain member and exempted computers to communicate with an isolated server. However, servers differ in the level of sensitivity of their data and whether they permit universal access or access only from specific computers. For example, e-mail servers typically must be available to all domain member computers in order to allow a user on any domain member computer to access e-mail. However, finance or legal department servers should only be available to a specific subset of computers.

To further isolate servers that store sensitive data and prevent unauthorized domain member computers from communicating with them, you can separate these servers and authorized computers by providing them with an IP address on a separate subnet, or by using a different authentication method, such as digital certificates. However, both solutions require additional administrative overhead.

For servers running Windows, you can use another method to enforce authorization. By applying the **Access this computer from the network** user right (a local Group Policy setting), you can specify the computer accounts or Active Directory security groups that you want to allow to access a server over the network. When IPsec processes the credentials of the computer requesting communications, Windows checks the **Access this computer from the network** user right.

Therefore, to further isolate servers that store sensitive data based on Active Directory group membership, you can do the following:

1. Ensure that the servers have been added to the existing IPsec rule for isolated servers.

2. Create an Active Directory security group and add the computer accounts of the authorized computers to the group.

3. Configure the local Group Policy settings of each server that stores sensitive data and change the **Access this computer from the network** user right as follows:

    a. Click **Computer Configuration**, click **Windows Settings**, and then click **Security Settings**.

    b. Click **Local Policies**, and then click **User Rights Assignment** so that it contains only the Active Directory security group created in step 2.

For a file server that contains sensitive financial data, add the financial server to the IPsec rule of the isolated servers, create a ConfidentialFinancial security group, and then add the authorized computer accounts. Finally, change the **Access this computer from the network** user right on the financial file server to contain only the ConfidentialFinancial security group.

**Overview of Server Isolation Deployment**
To deploy server isolation, do the following:

1. Determine the state of your network infrastructure.

    Before you can begin planning for server isolation, you must assess your organization's network. In your assessment, identify and document your network's physical topology (such as client and server computer configurations), logical

topology (such as your Active Directory infrastructure including trust relationships and system container structure), and current use of Group Policy settings. You must also determine which computers to exempt.

2. Design and test server isolation IPsec policy in a lab network.

   Create a scaled-down version of your network in a physically isolated lab that is not connected to your production network. Your test lab network should include domain member client computers, client computers that are not a member of the domain, and exempted computers. Then, configure the IPsec rules required to implement server isolation for your network. Use the test lab to ensure that the policies work as expected. Fine-tune your policy settings, as needed.

3. Perform a pilot using a subset of computers.

   After verifying the IPsec policy settings in the test lab, configure and activate the server isolation IPsec policy on a subset of computers on your production network to test their behavior. For example, you might want to activate the IPsec policy for the computers in a specific Active Directory organizational unit.

4. Roll out the IPsec policy in phases.

   After the pilot program is complete, begin activating the IPsec policy for other parts of your domain infrastructure in a phased roll out.

For more information about deploying server isolation, see "Server and Domain Isolation Using IPsec and Group Policy."

**See Also**
**Other Resources**
Introduction to Server and Domain Isolation with Microsoft Windows
Domain Isolation with Microsoft Windows Explained
Server and Domain Isolation Using IPsec and Group Policy
Improving Security with Domain Isolation
Windows Server 2003 IPsec Web site
Windows Server 2003 Web site

# Domain Isolation with Microsoft Windows Explained

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As described in "Introduction to Server and Domain Isolation with Microsoft Windows," domain isolation uses an Active Directory® directory service domain, domain membership, and Group Policy settings to enforce a network policy that requires domain member computers to accept incoming communication requests only from computers that can authenticate themselves with domain credentials. This network policy isolates domain member computers from non-domain-member computers.

By isolating your domain, you provide an additional layer of protection for your network traffic. Security technologies, such as Institute of Electrical and Electronic Engineers (IEEE) standard 802.1X, require a computer to authenticate itself before sending frames on a network. However, 802.1X does not protect the traffic sent by an 802.1X-authenticated computer after it is on the network. Secure Sockets Layer (SSL) provides computer authentication and data confidentiality (encryption) for SSL-enabled client and server applications. However, SSL works only if the client and server application support SSL. Whereas 802.1X works at the Data Link layer (of the Open Systems Interconnections [OSI] model) and SSL works at the Application layer, domain isolation works at the Network layer, providing additional protection for IP-based traffic. Organizations that routinely send sensitive data on their networks and, therefore, must provide extra protection for data assets—such as those in the financial services or health care industries or government institutions—need this additional level of protection at the Network layer.

Domain isolation provides many benefits by:

- Restricting incoming connections to domain member computers.

  In Microsoft® Windows®, you can manage a computer that is a member of an Active Directory domain by centrally configuring settings in Group Policy and by applying those settings to all domain member computers. You can also apply other types of security updates, such as operating system updates and antivirus software signatures, to domain member computers. Domain member computers use their domain credentials to authenticate communication attempts with each other. Computers that are not domain members — stand-alone, unknown, and guest computers — do not have domain credentials and, therefore, cannot authenticate communication attempts with domain member computers.

- Supplementing other security mechanisms designed to prevent unwanted communications.

  Domain isolation supplements the security provided by security mechanisms already deployed on your network. For example, if you isolated your domain and then your firewall was compromised, malicious users from the Internet could not initiate direct communications with the isolated hosts.

- Encouraging domain membership.

  By placing critical organization servers, such as e-mail servers, on the isolated domain, you prevent network users from connecting to any of those servers from a computer that is not a domain member. To receive valid domain credentials for authenticating with isolated organization servers, computers that are not domain members must be joined to the domain. After they have been joined to the domain, you can manage these computers in other ways, such as ensuring that they have the latest operating system and antivirus updates.

- Protecting traffic between domain member computers.

  Traffic sent between domain member computers is cryptographically protected so that the receiving computer can verify that an authenticated computer sent the packet and that the packet was not modified in transit. Optionally, the traffic between domain member computers can be encrypted, providing protection from malicious network users who attempt to capture and interpret network traffic.

To deploy domain isolation, you configure Group Policy settings to require that all incoming connection requests and subsequent data be authenticated and protected by using Internet Protocol security (IPsec). IPsec protects traffic from address spoofing, data injection, session hijacking, replay attacks, and other types of data tampering. Optionally, you can specify that packets be encrypted. You also can configure exemptions so that specific computers that are not domain members can initiate communications with isolated hosts.

**How Domain Isolation Works**

This section describes how Active Directory, Group Policy, and IPsec work together to implement domain isolation.

**Components of Domain Isolation**

To isolate a domain, you need the following components:

- **An Active Directory domain.** The domain includes domain controllers and the appropriate trust relationships to establish trust with other domains or the directory trees of an organization network.
- **Member computers.** These are computers that have joined the Active Directory domain and received domain credentials.
- **Group Policy settings.** These computer and user settings are automatically downloaded to domain member computers.
- **Active IPsec policy settings.** These Group Policy settings determine the domain isolation behavior of domain member computers.

In a simplified domain isolation deployment, you configure and activate an IPsec policy with rules that define specific types of traffic and how the traffic will be handled. You then activate the IPsec policy for the appropriate Active Directory containers, such as sites, domains, and organizational units. The member computers in the Active Directory containers to which the Group Policy settings apply automatically download the Group Policy settings.

After the domain member computers have downloaded and applied the Group Policy settings, they have both the correct IPsec policy for domain isolation and the domain credentials that will allow them to communicate securely with each other and to communicate without security with non-domain-member computers. Computers that are not domain members, which do not have the correct IPsec policy settings for domain isolation or domain credentials, cannot initiate communication with isolated hosts.

**Communication Processes**

When you implement domain isolation, communication between hosts in your network differs depending on which type of host (isolated or non-isolated) initiates communication and which type of host the initiating host attempts to communicate with. This section describes how communication occurs:

- When an isolated host initiates communication with another isolated host.
- When a non-isolated host initiates communication with an isolated host.
- When an isolated host initiates communication with a non-isolated host.

The following figure shows the types of communication that occur when you deploy domain isolation.




**Legend**




**Communication with an isolated host initiated by another isolated host**

When an isolated host with both Active Directory credentials and domain isolation IPsec policy settings (for example, COMPUTER1 in the figure) initiates communication with another isolated host (for example, COMPUTER2), the following occurs:

1. The initial communication packet sent by COMPUTER1—for example, a Transmission Control Protocol (TCP) Synchronize

(SYN) segment—matches the rule of the active IPsec policy that specifies that the initiating host must attempt to secure the traffic with IPsec.

2. COMPUTER1 uses IPsec to perform mutual authentication with COMPUTER2 and to negotiate the use of IPsec protection.

3. Because COMPUTER2 has domain credentials, the IPsec authentication process succeeds. Because COMPUTER2 has IPsec policy settings that match those on COMPUTER1, negotiation of IPsec protection also succeeds.

4. COMPUTER1 sends the initial communication packet to COMPUTER2 with IPsec protection.

5. COMPUTER2 sends the response to the initial communication packet—for example, a TCP SYN-Acknowledgement (SYN-ACK) segment—to COMPUTER1 with IPsec protection.

6. Subsequent packets sent between COMPUTER1 and COMPUTER2 are also protected by IPsec.

Because they are domain members and have IPsec policy settings, isolated hosts authenticate and protect with IPsec communications initiated with other isolated hosts.

**Communication with a non-isolated host initiated by an isolated host**
When an isolated host with both Active Directory credentials and domain isolation IPsec policy settings (for example, COMPUTER1 in the figure) initiates communication with a non-isolated host (for example, COMPUTER3), the following occurs:

1. The initial communication packet being sent by COMPUTER1—for example, a TCP SYN segment—matches the rule of the IPsec policy that specifies that the initiating host must attempt to secure the traffic with IPsec.

2. COMPUTER1 attempts to use IPsec to authenticate COMPUTER3 and to negotiate the use of IPsec protection.

3. Because COMPUTER3 does not have domain credentials, the IPsec authentication attempt fails.

4. Because the rule matched in Step 1 allows unsecured communication with computers that fail the IPsec authentication, COMPUTER1 sends the initial communication packet without IPsec protection.

5. COMPUTER3 sends the response to the initial communication packet sent by COMPUTER1 without IPsec protection.

6. COMPUTER1 and COMPUTER3 send subsequent packets without IPsec protection.

An isolated host tries to authenticate non-isolated hosts. If it cannot authenticate a host, an isolated host sends packets without IPsec protection, allowing isolated hosts to initiate communications with non-isolated hosts.

**Communication with an isolated host initiated by a non-isolated host**
When a non-isolated host (for example, COMPUTER3 in the figure) initiates communication with an isolated host (for example, COMPUTER2), the following occurs:

1. Because COMPUTER3 does not have IPsec policy settings, it sends its initial communication packet—for example, a TCP SYN segment—without IPsec protection to COMPUTER2.

2. On COMPUTER2, the initial communications packet sent by COMPUTER3 matches the rule of the IPsec policy that requires incoming communication attempts to be protected by IPsec.

3. Because the rule does not allow COMPUTER2 to accept incoming communication attempts that are not protected by IPsec, COMPUTER2 discards the initial communications packet sent by COMPUTER3.

4. COMPUTER2 also discards subsequent incoming communication attempts from COMPUTER3.

5. COMPUTER3 fails in its attempt to communicate with COMPUTER2.

Isolated hosts discard all initial communication packets sent by non-isolated hosts.

**Overview of Domain Isolation Deployment**

To deploy domain isolation, use the following basic procedure:

1. Determine the state of your network infrastructure.

   Before you can begin planning for domain isolation, you must assess your organization's network. In your assessment, identify and document your network's physical topology (such as client and server computer configurations), logical topology (such as your Active Directory infrastructure including trust relationships and system container structure), and current use of Group Policy settings. You must also determine which computers to exempt.

2. Design and test domain isolation IPsec policy in a lab network.

   Create a scaled-down version of your network in a physically isolated lab that is not connected to your production network. Then, configure the IPsec rules required to implement domain isolation for your network. Be sure to include isolated and non-isolated client computers in your test lab network. Use the test lab to ensure that the IPsec policies work as expected. Fine-tune your policy settings, as needed.

3. Perform a pilot with a subset of computers.

   After verifying IPsec policy settings in the test lab, configure and activate the domain isolation IPsec policy on a subset of computers on your production network to test their behavior. For example, you might want to activate the IPsec policy settings for the computers in a specific Active Directory organizational unit.

4. Roll out the IPsec policy in phases.

   After the pilot program is complete, begin activating the IPsec policy for other parts of your domain infrastructure in a phased rollout.

For more information about deploying domain isolation, see "Server and Domain Isolation Using IPsec and Group Policy." For an example of how Microsoft deployed domain isolation, see "Improving Security with Domain Isolation."

**See Also**

**Other Resources**

Introduction to Server and Domain Isolation with Microsoft Windows
Server Isolation with Microsoft Windows Explained
Server and Domain Isolation Using IPsec and Group Policy
Improving Security with Domain Isolation
Windows Server 2003 IPsec Web site
Windows Server 2003 Web site

# What's New in Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Microsoft® Windows Server™ 2003 R2, a software update to the Windows Server 2003 operating system, makes it easier and more cost effective to extend connectivity and control to identities, locations, data, and applications throughout and beyond your organization.

Built on Windows Server 2003, Windows Server 2003 R2 takes advantage of the stability and security of the code base while extending connectivity and control into new areas. With Windows Server 2003 R2, you get the following improvements:

- Identity and access management
- Branch office server management
- Storage setup and management
- Application development inside and outside your organization's traditional boundaries

For more information about Windows Server 2003 R2, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45560).

This topic describes the following new components that you can install with Windows Server 2003 R2:

**Server Manageability**

- Administration Tools Packs
- Hardware Management
- MMC 3.0

**Features for Active Directory**

- Active Directory Application Mode (ADAM)
- Active Directory Federation Services

**Disk and File Management Features**

- Branch Office
    - Distributed File System (DFS)
- Common Log File System (CLFS)
- File Server Management
- Microsoft Services for Network File System
- Storage Management for SANs
- File Server Resource Manager

**Printer and Protocol Support**

- Print Management

**Microsoft .NET Framework**

- Microsoft .NET Framework 2.0

**Internet and E-Mail Services and Features**

- Windows Sharepoint Services

**UNIX Interoperability**

- Identity Management for UNIX
- Subsystem for UNIX-based Applications

**Server Manageability**
To help you manage your server, Windows Server 2003 R2 provides Administration Tools Packs and the following improvements to hardware management and Microsoft Management Console (MMC) 3.0.

**Administration Tools Packs**
To make the remote management of your servers easier, you can install the Administration Tools Packs for Windows Server 2003 R2 located on Disc 2 in the \mgmttls directory. Each of the following components has an Administration Tools Pack:

- Microsoft Management Console (MMC) 3.0
- Print Management
- File Server Resource Manager
- DFS Replication
- Identity Management for UNIX
- File Server Management

In order to administer components, you must first install the MMC 3.0 Administration Tools Pack. The Administration Tools Packs can be installed on computers running Windows® XP Professional with Service Pack 2 (SP2). For specific installation instructions, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=54293).

**Hardware Management**
You can use Windows Remote Management (WinRM) to manage server hardware remotely across firewalls and monitor conditions on servers that are offline.

This new feature in Windows Server 2003 R2 is implemented using WS-Management, an industry-standard, Web services-based remoting protocol. WS-Management uses HTTPS as the secure transport that enables firewall traversal.

WinRM provides a command-line interface for common management tasks and a scripting application programming interface (API) for writing Windows Script Host-based system administration scripts.

Using WinRM, you can write scripts to monitor and control the state of server hardware by communicating with a Baseboard Management Controller (BMC). A BMC is a separate micro-controller with its own network adapter that is connected to the processor board of a server and can monitor conditions even when the server is off or malfunctioning. A new Windows Management Interface (WMI) provider for Intelligent Platform Management Interface (IPMI) exposes six new WMI classes for accessing BMC information with scripts.

WinRM is not installed by default with Windows Server 2003 R2. To install it, open Add or Remove Programs from Control Panel, and click **Add/Remove Windows Components**.

For more information about hardware management, see "Hardware Management in Microsoft Windows Server 2003 R2" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45204).

Back to Top

**MMC 3.0**

Microsoft Management Console (MMC) 3.0 supports richer functionality in snap-ins designed for the MMC 3.0 infrastructure. In addition, the following improvements apply to all MMC 3.0 consoles:

- **Action pane.** The Action pane appears on the right side of the snap-in console. It lists the actions that are currently available to you, based on the currently selected items in the tree or the results pane.

  To show or hide the action pane, click the **Show/Hide Action Pane** button on the toolbar, which is similar to the **Show/Hide Console Tree** button.

- **New Add or Remove Snap-in dialog box.** The new **Add or Remove Snap-in** dialog box makes it easy to add, organize, and remove snap-ins. You can control which extensions are available, and whether to automatically enable snap-ins that may be installed later. You can nest snap-ins and rearrange the snap-ins in the tree.

- **Improved error handling.** MMC 3.0 notifies you of errors in snap-ins that could cause MMC to fail and provides several options for responding to those errors.

MMC 3.0 is automatically installed when you install Windows Server 2003 R2.

Back to Top

**Features for Active Directory**

The following Active Directory components are included in Windows Server 2003 R2.

**Active Directory Application Mode (ADAM)**

Active Directory Application Mode (ADAM) is an independent mode of the Active Directory® directory service, without infrastructure features, that provides directory services for applications. It provides a data store and services for accessing the data store. It uses standard application programming interfaces (APIs) for accessing the application data. ADAM operates either as a standalone data store, or it operates with replication. Its independence enables local control and autonomy of directory services for specific applications. It also facilitates independent, flexible schemas and naming contexts.

Back to Top

**Active Directory Federation Services**

The fundamental purpose of Active Directory Federation Services (ADFS) is to leverage single user sign-on to authenticate the user to multiple related Web applications over the life of a single online session. ADFS does this by securely sharing digital identity and entitlement rights across security and enterprise boundaries.

Back to Top

**Disk and File Management Features**

Windows Server 2003 R2 includes the following disk and file management features.

**Branch Office**

Branch offices are the remote locations of an organization that connect to a headquarters or hub over a wide-area network (WAN). The technologies included in Windows Server 2003 R2 can simplify how you manage a variety of activities and tasks in branch offices. These activities and tasks include:

- Publishing files from centralized hubs to branch offices.

- Replicating files from branch locations for backup, fault-tolerance, or cross-publishing.

- Collaborating on documents that are shared among branches or between hubs and branches.

- Efficiently managing printers, printer drivers, and print queues in branch offices.

These improvements are supported by the enhanced print-management tools in Windows Server 2003 R2. They are also supported by Distributed File System (DFS) in Windows Server 2003 R2.

For more information about these technologies in Windows Server 2003 R2, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51679).

Back to Top

**Distributed File System (DFS)**
The Distributed File System (DFS) solution in Windows Server 2003 R2 provides simplified, fault-tolerant access to files and WAN-friendly replication. Distributed File System consists of two technologies:

- **DFS Namespaces.** Formerly known as Distributed File System, DFS Namespaces allows administrators to group shared folders located on different servers and present them to users as a virtual tree of folders known as a namespace. A namespace provides numerous benefits, including increased availability of data, load sharing, and simplified data migration.

- **DFS Replication.** The successor to File Replication service (FRS), DFS Replication is a new state-based, multimaster replication engine that supports scheduling and bandwidth throttling. DFS Replication uses a new compression algorithm known as Remote Differential Compression (RDC). RDC is a "diff over the wire" protocol that can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, re-arrangements of data in files, enabling DFS Replication to replicate only the deltas (changes) when files are updated.

For more information about Distributed File System in Windows Server 2003 R2, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51679).

Back to Top

**Common Log File System (CLFS)**
Common Log File System (CLFS) is a loadable driver that provides kernel- or user-mode applications with a robust logging subsystem. CLFS is a unique Windows technology that can be used to develop applications and middleware that depend on durably writing and reading sequential data. Examples include replication agents, auditing agents, databases, and other transactional resource managers.

Features include the ability to create log files with a single stream of data or with multiple streams of data for shared use by one or more clients; both circular and linear logging; guaranteed ability to flush buffered data by pre-reserving space in the log; policy-based log size and space management; sharing of a single log by both kernel and user clients; a notification mechanism to allow different users within the same log to coordinate their log use; flexible buffering of log data; archiving APIs do not interfere with normal operations; atomic multi-sector writes; and torn-write detection.

CLFS is optimized for performance. All writes to the log file are buffered until an explicit flush, an opportunity to share a write with another client, or the buffer is filled. Log data is written directly to the hard disk from the log buffers without copying. Multiple streams of data can be written during the same I/O operation, resulting in only one disk seek for what normally takes multiple seeks and writes. Reads are cached to save disk accesses during normal operation or bursts of read activity.

For more information, see the "Platform SDK" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=44489).

Back to Top

**File Server Management**
You can use File Server Management to perform many tasks, such as formatting volumes, creating shares, defragmenting volumes, creating and managing shares, setting quota limits, creating storage utilization reports, replicating data to and from the file server, managing Storage Area Networks (SANs), and sharing files with UNIX and Macintosh systems.

Back to Top

**File Server Resource Manager**
With the increasing demand on storage resources, as organizations rely more heavily on data than ever before, IT administrators face the challenge of overseeing a larger and more complex storage infrastructure, while at the same time, tracking the kind of information available in it.

File Server Resource Manager is a new Microsoft Management Console (MMC) snap-in that provides a suite of tools for administrators to understand, control, and manage the quantity and type of data stored on their servers. By using File Server Resource Manager, administrators can place quotas on folders and volumes, actively screen files, and generate comprehensive storage reports.

**File Server Resource Manager quotas vs. NTFS disk quotas**
The Windows 2000 and Windows Server 2003 operating systems support disk quotas, which are used to track and control disk

usage per user on NTFS volumes. The following table outlines the advantages of using the quota management tools in File Server Resource Manager.

| Quota features | File Server Resource Manager | NTFS disk quotas |
|---|---|---|
| Quota tracking | By folder or by volume | Per user on a volume |
| Disk usage calculation | Actual disk space | Logical file size |
| Notification mechanisms | E-mail, custom reports, command execution, event logs | Event logs only |

Back to Top

**Microsoft Services for Network File System**

Microsoft Services for Network File System (NFS) is a component of Windows Server 2003 R2 that provides a file-sharing solution for enterprises that have a mixed Windows and UNIX environment. Microsoft Services for NFS allows users to transfer files between Windows Server 2003 R2 and UNIX computers using the Network File System (NFS) protocol.

Microsoft Services for NFS is an update to the NFS components that were previously available in Services for UNIX 3.5. Microsoft Services for NFS includes the following new features:

- Support for 64-bit

  Microsoft Services for NFS components can be installed on all editions of Windows Server 2003 R2, including 64-bit versions.

- Updated Microsoft Management Console (MMC) Microsoft Services for NFS Administration snap-in

- Enhanced reliability

- Support for UNIX special devices (mknod)

Microsoft Services for NFS allows you to support a mixed environment of Windows- and UNIX-based operating systems. It also allows you to update your organization's computers while supporting older technology during the transition phase. The following scenarios are examples of how enterprises can benefit from deploying Microsoft Services for NFS.

- Allows UNIX clients to access resources on computers running Windows Server 2003 R2. Your company may have UNIX clients accessing resources, such as files, on UNIX file servers. To take advantage of new Windows Server 2003 features, such as Shadow Copies for Shared Folders, you can move resources from your UNIX servers to computers running Windows Server 2003 R2, and then set up Microsoft Services for NFS to enable access by UNIX clients that are running NFS software. All of your UNIX clients will be able to access the resources using the NFS protocol with no changes required.

- Allows computers running Windows Server 2003 R2 to access resources on UNIX file servers. Your company may have a mixed Windows and UNIX environment with resources, such as files, stored on UNIX file servers. You can use Microsoft Services for NFS to enable computers running Windows Server 2003 R2 to access these resources when the file servers are running NFS software.

Back to Top

**Storage Management for SANs**

Storage Manager for SANs is a new Microsoft Management Console (MMC) snap-in that helps you create and manage logical unit numbers (LUNs) on fibre channel and iSCSI disk drive subsystems in your storage area network (SAN). Storage Manager for SANs can be used on storage subsystems that support Virtual Disk Server (VDS).

Use Storage Manager for SANs to create and assign LUNs, manage connections between LUNs and the servers in your SAN, and set the security properties for iSCSI storage subsystems.

Back to Top

**Printer and Protocol Support**

Windows Server 2003 R2 provides an updated MMC snap-in for Print Management.

**Print Management**
Print Management is an updated Microsoft Management Console (MMC) snap-in that you can use to view and manage printers and print servers in your organization. You can use Print Management from any computer running Windows Server 2003 R2, and you can manage all network printers on print servers running Windows 2000 Server, Windows Server 2003, or Windows Server 2003 R2.

Print Management provides up-to-the-minute details such as the queue status, printer name, driver name, and server name. You can also set custom views by using the Print Management filtering capability. For example, you can create a view that only displays printers in a particular error state. You can also configure Print Management to send e-mail notifications or run scripts when a printer or print server needs attention. The filtering capability also allows you to bulk edit print jobs, such as canceling all print jobs at once. You can also delete multiple printers at the same time.

Administrators can install printers remotely by using the automatic detect feature, which finds and installs printers on the local subnet to the local print server. Administrators can log on to a local server at a branch location by using Remote Desktop and then easily install printers remotely.

When a printer model has a Web page feature, Print Management can display troubleshooting details, such as the location of the paper jam or printer toner level. Some printer models have Web pages that provide options for remote control functions that can help resolve problems at branch locations.

For more information, see the "Print Management Step-by-Step Guide" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50141).

Back to Top

**Microsoft .NET Framework**
New features for .NET Framework 2.0 are available in Windows Server 2003 R2.

**Microsoft .NET Framework 2.0**
The .NET Framework version 2.0 is the latest version of the .NET Framework, a Windows component for building, deploying and running applications. The .NET Framework 2.0 is focused on delivering developer productivity, operational excellence, and the ability to preserve and enhance existing software investments. Core enhancements include support for 64-bit processors for increased performance and scalability, transactional enhancements and optimizations, productivity improvements in ASP.NET, and support for WS-I BP 1.1 and SOAP 1.2 with Web Services.

For a comprehensive list of the new .NET Framework 2.0 features, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50322).

Back to Top

**Internet and E-Mail Services and Features**
Windows SharePoint® Services is included in Windows Server 2003 R2.

**Windows SharePoint Services**
Microsoft Windows SharePoint Services is an integrated portfolio of collaboration and communication services designed to connect people, information, processes, and systems both within and beyond the organizational firewall. Windows SharePoint Services Service Pack 2 is included in this release of Windows Server 2003. Now you can install Windows SharePoint Services directly from the Configure Your Server wizard or Manage Your Server. When you install Windows SharePoint Services, you get the following features:

**SharePoint sites – file storage plus collaboration**
Web sites based on Microsoft Windows SharePoint Services 2.0 provide a place where teams can communicate, share documents, and work together on a project. SharePoint sites include:

- Team collaboration features, such as event calendars, contacts, Web links, discussions, issues lists, announcements, and much more.
- Document libraries, which are places where users can store and retrieve documents while taking advantage of rich features such as check-in and check-out, version history, custom metadata, and flexible, customizable views.
- Web Parts, which can provide data access, Web services, and many other applications and content to SharePoint sites.

Site users can contribute to the site by using nothing more than a Web browser. However, if users have Windows SharePoint Services-compatible client programs, such as Microsoft Office 2003, installed on their computers they can work seamlessly with the site, saving files to libraries, editing documents in the client program, and moving or linking that information to the

site.

**SharePoint Central Administration – Web browser interface for managing your server**
You can manage a single server or an entire server farm running Windows SharePoint Services from a Web browser interface called SharePoint Central Administration. Use SharePoint Central Administration to extend a virtual server, create sites or turn on Self-Service Site Creation so users can create their own sites, manage security settings, manage the list of servers in your server farm, and so on. If you prefer, you can also use the Stsadm.exe command-line utility to manage your servers running Windows SharePoint Services.

**Applications for Windows SharePoint Services – download and install ready-to-use applications to get working quickly**
Boost worker productivity by downloading and installing an application for Windows SharePoint Services. These applications are tailored to address specific processes or tasks. There are many to choose from, including recruiting, project management, help desk issue tracking, timesheet and scheduling, event planning, and more. Use the applications as soon as you install them or customize them further to meet your organization's specific needs and requirements.

For more information, see "Applications for Windows SharePoint Services" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47987).

**Support for advanced extranet configurations**
A URL zone mapping feature (new in Windows SharePoint Services, Service Pack 2) allows for easier interactions with a proxy server or firewall. Windows SharePoint Services uses absolute URLs to generate some hyperlinks in its Web pages and e-mail messages. This prevented Windows SharePoint Services from supporting certain advanced extranet configurations (such as SSL termination, host header modification, and port translation) where a reverse proxy server is deployed in front of the Windows SharePoint Services server. Now you can set up URL zones, such as Internet, intranet, and extranet, and map incoming and outgoing URLs to those zones, so that these configurations can work correctly.

For more information about Windows SharePoint Services in Windows Server 2003 R2, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45560).

Back to Top

**UNIX Interoperability**
For improved interoperability with UNIX, the following features are available in Windows Server 2003 R2.

**Identity Management for UNIX**
Windows Services for UNIX delivers Windows and UNIX integration with updated identity management solutions. The following solutions help provide uninterrupted user access and efficient management of network resources across operating systems:

- **Server for NIS.** Helps integrate Windows and UNIX-based Network Information Service (NIS) servers by enabling an Active Directory domain controller to act as a master NIS server for one or more NIS domains.

- **Password Synchronization.** Helps integrate Windows and UNIX servers by simplifying the process of maintaining secure passwords. With Password Synchronization, users do not need to maintain separate passwords for their Windows and UNIX accounts or remember to change the password in multiple locations.

Back to Top

**Subsystem for UNIX-based Applications**
Subsystem for UNIX-based Applications (SUA) is a source-compatibility subsystem for compiling and running custom UNIX-based applications on a computer running a Windows server-class operating system. You can make your UNIX applications fully interoperable with Windows in SUA with little or no change to your original source code.

Subsystem for UNIX-based Applications provides an operating system for Portable Operating System Interface (POSIX) processes. SUA, along with its package of support utilities (such as shells and a telnet client), provides a complete UNIX environment. The package of support utilities includes a comprehensive set of scripting utilities and a software development kit (SDK) designed to fully support the development capabilities of SUA.

SUA also supports case-sensitive file names, job control, compilation tools, and the use of over 300 UNIX commands, utilities, and shell scripts. Because the subsystem installs separately from the Windows kernel, it offers true UNIX functionality without any emulation.

New features in this release include the following:

- **Database (OCI/ODBC) library connectivity.** SUA supports connectivity to Oracle and SQL Server from database

applications by using the Oracle Call Interface (OCI) and the Open Database Connectivity (ODBC) standard.

- **Microsoft® Visual Studio® Debugger Extension for debugging POSIX applications.** SUA includes support for debugging your POSIX processes using Visual Studio IDE.

- **Utilities based on SVR-5 and BSD UNIX environments.** The SUA download package supports two different UNIX environments: SVR-5 and BSD.

- **Support for 64-bit applications.** Using a process called thunking, SUA provides support not only for 64-bit applications running on a 64-bit operating system, but also default support for 32-bit binaries running on a 64-bit operating system.

To download the support utilities, see "Installing and Using Utilities and SDK for UNIX-based Applications" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=38959).

Back to Top

# Overview of ADFS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Introduction to ADFS**
Active Directory® Federation Services (ADFS) is a component in Windows Server® 2003 R2 that provides Web single-sign-on (SSO) technologies to authenticate a user to multiple Web applications over the life of a single online session. ADFS accomplishes this by securely sharing digital identity and entitlement rights, or "claims," across security and enterprise boundaries.

ADFS is not:

- A database or repository for employee or customer identity data.
- An extension of the Active Directory schema.
- A type of Windows domain or forest trust.

ADFS in Windows Server 2003 R2 supports the WS-Federation Passive Requestor Profile (WS-F PRP).

**Key features of ADFS**
The following are some of the key features of ADFS in Windows Server 2003 R2:

- Federation and Web SSO

  When an organization uses the Active Directory™ directory service, it currently experiences the benefit of SSO functionality through Windows-integrated authentication within the organization's security or enterprise boundaries. ADFS extends this functionality to Internet-facing applications, which enables customers, partners, and suppliers to have a similar, streamlined, Web SSO user experience when they access the organization's Web-based applications. Furthermore, federation servers can be deployed in multiple organizations to facilitate business-to-business (B2B) federated transactions between partner organizations.

- Web Services (WS)-* interoperability

  ADFS provides a federated identity management solution that interoperates with other security products that support the WS-* Web Services Architecture. ADFS does this by employing the federation specification of WS-*, called WS-Federation. The WS-Federation specification makes it possible for environments that do not use the Windows identity model to federate with Windows environments.

- Extensible architecture

  ADFS provides an extensible architecture that supports the Security Assertion Markup Language (SAML) token type and Kerberos authentication (in the Federated Web SSO with Forest Trust scenario). ADFS can also perform claim mapping, for example, modifying claims using custom business logic as a variable in an access request. Organizations can use this extensibility to modify ADFS to coexist with their current security infrastructure and business policies.

**Extending Active Directory to the Internet**
Active Directory serves as a primary identity and authentication service in many organizations. With Windows Server 2003 Active Directory, forest trusts can be created between two or more Windows Server 2003 forests to provide access to resources that are located in different business units or organizations. For more information about forest trusts, see How Domain and Forest Trusts Work on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

However, there are scenarios in which forest trusts are not a viable option. For example, access across organizations may need to be limited to only a small subset of individuals, not every member of a forest.

By employing ADFS, organizations can extend their existing Active Directory infrastructures to provide access to resources that are offered by trusted partners across the Internet. These trusted partners can include external third parties or other

departments or subsidiaries in the same organization.

ADFS is tightly integrated with Active Directory. ADFS retrieves user attributes from Active Directory, and it authenticates users against Active Directory. ADFS also uses Windows Integrated Authentication.

ADFS works with both Active Directory and Active Directory Application Mode (ADAM). Specifically, ADFS works with both enterprise-wide deployments of Active Directory or instances of ADAM. When it works with Active Directory, ADFS can take advantage of the strong authentication technologies in Active Directory, including Kerberos, X.509 digital certificates, and smart cards. When it works with ADAM, ADFS uses Lightweight Directory Access Protocol (LDAP) Bind as a means to authenticate users.

ADFS supports distributed authentication and authorization over the Internet. ADFS can be integrated into an organization's or department's existing access management solution to translate the terms that are used in the organization into claims that are agreed on as part of a federation. ADFS can create, secure, and verify the claims that move between organizations. It can also audit and monitor the activity between organizations and departments to help ensure secure transactions.

**Federation scenarios**

ADFS supports federated identity scenarios that use Web Services Federation (WS-Federation), WS-F PRP, and WS-Federation Passive Requestor Interoperability Profile specifications. The ADFS solution helps administrators deal with federated identity management challenges by making it possible for organizations to securely share a user's identity information over federation trusts. The following three descriptions of deployment scenarios illustrate how you can use ADFS server roles to federate identities, depending on the needs of your organization.

**Federated Web SSO**

The ADFS Federated Web SSO scenario involves secure communication that often spans multiple firewalls, perimeter networks, and name resolution servers, in addition to the entire Internet routing infrastructure. Communication over a federated Web SSO environment can help foster more efficient and secure online transactions between organizations that are joined by federation trust relationships.

As shown in the following illustration, a federation trust relationship can be established between two businesses. In this scenario, federation servers route authentication requests from user accounts in Tailspin Toys to Web-based applications that are located in the network of Online Retailer.




Federation servers authenticate requests from trusted partners based on the credentials of the partners. Representations of the credentials are exchanged in the form of security tokens.

For enhanced security, federation server proxies can be used to relay requests to federation servers that are not directly accessible from the Internet.

**Federated Web SSO with Forest Trust**

The ADFS Federated Web SSO with Forest Trust scenario involves two Active Directory forests in a single organization, as shown in the following illustration. One of the forests is located in the organization's perimeter network (also known as a demilitarized zone (DMZ) or a screened subnet). The other forest is located in the internal network. A one-way, forest trust is established so that the forest in the perimeter network trusts the forest in the internal network. Federation servers are deployed in both networks. A federation trust is established so that accounts in the internal forest can be used to access a Web-based application in the perimeter network, whether the accounts access the site from the intranet forest or from the Internet.




**A. Datum Corp.**

In this scenario, external users, such as customers, can access the Web application by authenticating to the external account federation server, which is located in the perimeter network. External users have user accounts in the perimeter-network Active Directory forest. Internal users, such as employees, can also access the Web application by authenticating to the internal account federation server, which is located in the internal network. Internal users have accounts in the internal Active Directory forest.

If the Web-based application is a Windows NT token-based application, the ADFS Web Agent that is running on the Web application server intercepts requests and creates Windows NT security tokens, which are required by the Web application to make authorization decisions. For external users, this is possible because the Web server that hosts the Windows NT® token-based application is joined to the domain in the external forest. For internal users, this is enabled through the forest trust relationship that exists between the perimeter forest and the internal forest.

If the Web-based application is a claims-aware application, the ADFS Web Agent that is running on the Web application server does not have to create Windows NT security tokens for the user. The ADFS Web agent can expose the claims that come across, which allows the application to make authorization decisions based on the contents of the security token that is provided by the account federation server. As a result when it deploys claims-aware applications, the Web server does not have to be joined to the domain, and the external-forest-to-internal-forest trust is not required.

**Web SSO**

In the ADFS Web SSO scenario, users must authenticate only once to access multiple Web-based applications. In this scenario all users are external, and no federation trust exists. Because the Web servers must be Internet accessible and also joined to the Active Directory domain, they are connected to two networks; that is, they are multihomed. The first network is Internet facing (the perimeter network) to provide the needed connectivity. The second network contains the Active Directory forest (the protected network), which is not directly Internet accessible. The federation server proxy is also multihomed to provide the necessary connectivity to the federation server and the Internet. In this scenario, placing the federation server on a network that is not directly accessible from the Internet greatly reduces the risk to the federation server.




**Partner organizations**
When you plan for cross-organizational (federation-based) collaboration using ADFS, you first determine whether your organization will host a Web resource to be accessed by other organizations across the Internet — or the reverse. This determination affects how you deploy ADFS, and it is fundamental in the planning of your ADFS infrastructure.

For federation scenarios such as Federated Web SSO and Federated Web SSO with Forest Trust (but not the Web SSO scenario), ADFS uses terms such as "account partner" and "resource partner" to help differentiate the organization that hosts the accounts (the account partner) from the organization that hosts the Web-based resources (the resource partner). The term "federation trust" is used in ADFS to characterize the one-way, nontransitive relationship that is established between the account partner and the resource partner.

The following sections explain some of the concepts that are related to account partners and resource partners.

**Account partner**
An account partner represents the organization in the federation trust relationship that physically stores user accounts in either an Active Directory store or an ADAM store. The account partner is responsible for collecting and authenticating a user's credentials, building up claims for that user, and packaging the claims into security tokens. These tokens can then be presented across a federation trust to access Web-based resources that are located in the resource partner organization.

In other words, an account partner represents the organization for whose users the account-side Federation Service issues security tokens. The Federation Service in the account partner organization authenticates local users and creates security tokens that are used by the resource partner in making authorization decisions.

In relation to Active Directory, the account partner in ADFS is conceptually equivalent to a single Active Directory forest whose accounts need access to resources that are physically located in another forest. Accounts in this example forest can access resources in the resource forest only when an external trust or forest trust relationship exists between the two forests and the resources to which the users are trying to gain access have been set with the proper authorization permissions.

| **Note** |
| --- |
| This analogy is meant strictly to emphasize how the relationship between account and partner organizations in ADFS is similar, in concept, to the relationship between an account forest and a resource forest in Active Directory. External trusts and forest trusts are not required for ADFS to function. |

**Resource partner**
A resource partner is the second organizational partner in the federation trust relationship. A resource partner is where the Web servers that host one or more Web-based applications (the resources) reside. The resource partner trusts the account partner to authenticate users. Therefore, to make authorization decisions, the resource partner consumes the claims that are packaged in security tokens coming from users in the account partner.

In other words, a resource partner represents the organization whose Web servers are protected by the resource-side Federation Service. The Federation Service at the resource partner uses the security tokens that are produced by the account partner to make authorization decisions for Web servers that are located in the resource partner.

To function as an ADFS resource, Web servers in the resource partner organization must have the ADFS Web Agent component of ADFS installed. Web servers that function as an ADFS resource can host either claims-aware applications or Windows NT token-based applications.

| Note |
| --- |
| If the application that is hosted on the Web server is a Windows NT token–based application, a resource account may be required for the Active Directory forest in the resource partner organization. |

In relation to Active Directory, the resource partner is conceptually equivalent to a single forest whose resources are made available over an external trust or forest trust relationship to accounts that are physically stored in another forest.

| Note |
| --- |
| This analogy is meant strictly to emphasize how the relationship between account and partner organizations in ADFS is similar, in concept, to the relationship between an account forest and a resource forest in Active Directory. External trusts and forest trusts are not required for ADFS to function. |

#### Enhanced identity privacy

**Enhanced identity privacy** is an optional setting that can be configured on a resource partner in the trust policy. If the **Enhanced identity privacy** option is enabled, this setting hashes the user-name portion of outgoing UPN claims and e-mail claims. It substitutes the common name with a random value.

The purpose of this feature is to prevent:

- The resource partner from correlating identity claims to personally identifiable user information.
- Collusion between partners in correlating identity claims to personally identifiable user information. This setting creates a unique hash per partner so that identity claim values are different across different trusting realm partners but consistent across sessions for a single partner.
- Simple dictionary attacks against the hash by salting the user value with data that is in the trust policy — data that is not known by the resource partners.

#### How enhanced identity privacy works

If the **Enable enhanced identity privacy** setting is enabled, the user's identity is transformed by the account Federation Service before it is sent to the resource realm. The account Federation Service does this by hashing a combination of the privacy key (for salting), the resource partner Uniform Resource Identifier (URI), and the identity claim itself. Therefore, the resulting transformed identity will be the same when the following values are the same:

- Privacy key
- Resource partner URI
- Identity claim (user principal name (UPN) or e-mail)

The hashing algorithm is case sensitive. Therefore, the hashed results are different if the case of the URI or the user-name portion of the UPN claim or e-mail claim is different.

#### ADFS server roles

ADFS can operate only when the servers running Windows Server 2003 R2 are configured appropriately. Depending on the environment in your organization, specific ADFS server roles must be deployed. The following sections describe the server roles that can be used to provide an ADFS federated identity management solution.

#### Federation server

Federation servers host the Federation Service component of ADFS. They are used to route authentication requests that are made from user accounts in other organizations (in Federated Web SSO scenarios) or from clients that can be located anywhere on the Internet (in the Web SSO scenario).

Federation servers also host a security token service that issues tokens that are based on the credentials (for example, user name and password) that are presented to it. After the credentials are verified (by the user logging on), claims for the user are collected through examination of the attributes for the user that are stored in Active Directory or ADAM.

In Federated Web SSO scenarios, claims can then be modified by claim mappings for a specific resource partner. The claims are built into a token that is sent to a federation server in the resource partner. After a federation server in the resource partner

receives the claims as incoming claims, it maps them into their organization claims. The organization claims are then built into a new token that is sent to the ADFS Web Agent.

The role that a federation server plays in either of the Federated Web SSO scenarios (Federated Web SSO or Federated Web SSO with Forest Trust) can depend on whether your organization is designated as the account partner or the resource partner:

- Federation servers in the account partner are used to log on local user accounts in either an Active Directory store or an ADAM store. Federation servers also issue initial security tokens that the local user accounts can use to access Web-based applications that are hosted in the resource partner. In addition, federation servers in the account partner issue cookies to users to maintain login status. These cookies include claims for those users. These cookies enable SSO capabilities so that users do not have to enter credentials each time that they visit different Web-based applications in the resource partner.

- Federation servers at the resource partner validate the security tokens that are issued by the federation servers at the account partner. Federation servers at the resource partner also issue security tokens that are meant for the Web-based applications in the resource partner. In addition, federation servers in the resource partner issue cookies to the user accounts, which come from the account partner. These cookies enable SSO capabilities so that users do not have to log in again at their federation servers in the account partner when users attempt to access different Web-based applications at the resource partner.

**Federation server proxy**

Federation server proxies host the Federation Service Proxy component of ADFS. Federation server proxies can be deployed in an organization's perimeter network (also known as a DMZ or a screened subnet) to forward requests to federation servers that are not accessible from the Internet.

| **Note** |
| --- |
| Although you can deploy separate servers to host the Federation Service Proxy component, it is not necessary to deploy a separate server to act as a federation server proxy in the intranet forest of either the account partner or the resource partner. A federation server performs this role automatically. |

The role that a federation server proxy plays in your organization can depend on whether your organization is the account partner or the resource partner:

- Federation server proxies at the account partner act as proxies for user logons to federation servers that are located in the intranet. Federation server proxies also act as proxies for security tokens that are issued by the account partner federation server for both its own tokens and those tokens that are destined for resource partners.

- Federation service proxies in the resource partner act as proxies for user security tokens, which are issued from federation servers in both the account partner and the resource partner, to Web-based applications in the resource partner.

**Web server**

In ADFS, Web servers in the resource forest host the ADFS Web Agent component to provide secure access to the Web applications that are hosted on those Web servers. The ADFS Web Agent manages security tokens and authentication cookies that are sent to a Web server. The Web server requires a relationship with a Federation Service so that all authentication tokens come from that Federation Service.

The ADFS Web Agent supports two types of applications: claims-aware applications and Windows NT token–based applications.

**Federation trusts**

You can use ADFS to enable efficient and secure online transactions between partner organizations that are joined by federation trust relationships. In other words, a federation trust is the embodiment of a business-level agreement or partnership between two organizations.

As shown in the following illustration, you can establish federation trust relationships between two partner organizations when both of the organizations deploy at least one ADFS federation server and they configure their Federation Service settings appropriately. The one-way arrow signifies the direction of the trust, which — like the direction of Windows trusts — always points to the account side of the forest. This means that authentication flows from the account partner organization to the resource partner organization.




| Note |
| --- |
| No communication occurs over the network between the account Federation Service and the resource Federation Service. |

After you create the federation trust, users who are located in the account partner organization can send authentication requests successfully through the federation trust to the Web server in the resource partner organization. A federation trust is created when the account partner organization and the resource partner organization both install the Federation Service component of ADFS and they both use the Active Directory Federation Services snap-in to configure the account partner and resource partner appropriately.

If one side of a federation trust (either the account partner or the resource partner) is not configured or if it is configured incorrectly by the administrator for either organization, the federation trust will not be created successfully. For detailed guidance about how to create federation trusts, look for ADFS step-by-step or deployment content on the ADFS home page (http://go.microsoft.com/fwlink/?LinkId=79542).

| Note |
| --- |
| Federation trusts are not used in the Web SSO scenario. |

**Federation Service**

The Federation Service is a component of ADFS that can be installed independently from other ADFS components. The Federation Service functions as a security token service. The act of installing the Federation Service component on a computer makes that computer a federation server. It also makes the Active Directory Federation Services snap-in available in the **Administrative Tools** menu on that computer.

The Federation Service is designed to use Active Directory to provide tokens in response to requests for security tokens. This enables Active Directory domains and forests to function as:

- Identity providers that can federate with compliant account partners and resource partners. As an identity provider, the Federation Service can project Active Directory identities across the Internet to interact with applications at compliant service providers.

- Service providers that can federate with compliant account partners and resource partners. As a service provider, the Federation Service can allow identities from other organizations to access a partner's Windows-based and Microsoft ASP.NET-based applications.

- Security token providers for applications that are compliant with the WS-F PRP specification.

As an account partner, the Federation Service allows Active Directory users to access resources at partner organizations. In response to a request from a resource partner, the Federation Service collects and verifies user credentials against Active Directory or ADAM. The Federation Service can then populate a set of organization claims based on the LDAP attributes of the user account. The organization claims are then mapped to appropriate claims for the resource partner and packaged into a security token that is signed by the Federation Service's token-signing certificate. The resultant security token is posted as the response to the resource partner's original request. The resource partner then uses the token to allow access for the user.

As a resource partner, the Federation Service plays the opposite role. When a user attempts to access an ADFS-protected application, the Federation Service determines which account partner should authenticate the user. It then sends an authentication request to that partner. When the user returns with a security token, the Federation Service verifies that the token has been correctly signed by the partner. It then extracts the claims from the token. The claims are mapped to organization claims, and the filtering policy for the specific application is applied. The filtered organization claims are packaged

into a security token that is either signed by the Federation Service's token-signing certificate or protected by a Kerberos session key for the Web application. The resultant security token is posted back to the original application Uniform Resource Locator (URL). The application then uses the token to allow access for the user.

ADFS uses the WS-F PRP protocol to carry claims in security tokens that are issued by the Federation Service to the Web application.

These claims are populated initially from account stores, either Active Directory or ADAM account stores. The Federation Service issues tokens based on the credentials that are presented. After the credentials are verified using the account store, the claims for the user are generated according to the rules of the trust policy. The obtained inbound claims are mapped into outbound claims that are appropriate for a resource partner. The resulting claim mappings are added to a security token that is issued to the resource partner.

After the token is verified, an authentication cookie is issued and written to the client browser. Each time that the client needs to be authenticated, this cookie is used by the Federation Service so that the client does not have to enter credentials again, thereby enabling SSO.

**Federation Service Web pages**
The Federation Service provides a Web page that prompts the user to select an appropriate account partner to which the user can authenticate. The Federation Service also provides a Web page that prompts for the user's credentials, such as a user name and password for forms-based authentication. A Web page is also provided that supports Windows Integrated Authentication.

Behind the Web pages, the Federation Service provides an ASP.NET Web service that processes requests from the client or the federation server proxy. The federation server proxy is located in the perimeter network. It acts as an intermediary between an Internet client and a Federation Service in the intranet. There are two basic types of requests to which the Federation Service responds:

- Requests to issue security tokens
- Requests to retrieve trust policy data

**Account partner discovery**
Account partner discovery is the process by which a client can identify what account partner it prefers for authentication in the event that more than a single account partner is configured. The Federation Server presents this choice to the client browser as a drop box containing the account partner names as they are configured in the trust policy.

One mechanism that you can use to avoid account partner discovery is to include the **`whr`** parameter in the query string for the resource being accessed, for example,

`https://webserver/testapp/testpage.aspx?whr=urn:federation:<` *accountpartner*`>`, where `<`*accountpartner*`>` indicates the account partner realm of the client.

When the **`whr`** parameter is used, the Web server removes the parameter and writes a cookie to the client browser to remember this setting for future requests. Then, the request proceeds in the same way as if it had not been provided.

**Federation Service Proxy**
The Federation Service Proxy is a component of ADFS in Windows Server 2003 R2 that can be installed independently from other ADFS components. The Federation Service Proxy functions as a proxy in a perimeter network (also known as a DMZ or a screened subnet) for the Federation Service. The act of installing the Federation Service Proxy component on a computer makes that computer a federation server proxy. It also makes the Active Directory Federation Services snap-in available on that computer in the **Administrative Tools** menu.

A federation server proxy participates in the WS-F PRP protocol by communicating with a protected Federation Service on the client's behalf. When the federation server proxy is protecting an account partner, it collects user credential information from browser clients. When the federation server proxy is protecting a resource partner, it relays requests by and for Web applications to the Federation Service.

The federation server proxy also stores Hypertext Transfer Protocol (HTTP) cookies on clients when necessary to facilitate SSO. The federation server proxy writes all three types of cookies: authentication cookies, account partner cookies, and sign-out cookies.

**ADFS Web Agent**
The ADFS Web Agent is a component of ADFS. It is used to consume security tokens and either allow or deny a user access to a Web application. To accomplish this, the Web server requires a relationship with a resource Federation Service so that it can direct the user to the Federation Service as needed.

The ADFS Web Agent can be used for two different types of applications:

- Claims-aware applications: An ASP.NET application that is written to published ADFS objects that allow the querying of ADFS security token claims. The application makes authorization decisions based on these claims. Security token failures for this type of application result in the client seeing an "Access Denied" message and events in the Federation Service event log.

- Windows NT token–based applications: An application that uses Windows-based authorization mechanisms. The ADFS Web Agent supports conversion from an ADFS security token to an impersonation-level Windows NT access token.

The Web server also stores HTTP cookies on clients where necessary to facilitate SSO.

**Terminology used in ADFS**

ADFS uses terminology from several different technologies, including certificate services, Internet Information Services (IIS), Active Directory, ADAM, and Web Services (WS-*). The following table describes these terms.

| Term | Description |
|---|---|
| account partner | A federation partner that is trusted by the Federation Service to provide security tokens. The account partner issues these tokens to its users (that is, users in the account partner realm) so that they can access Web-based applications in the resource partner. |
| Active Directory Federation Services (ADFS) | A Windows Server 2003 R2 component that provides Web SSO technologies to authenticate a user to multiple Web applications over the life of a single online session. ADFS accomplishes this by securely sharing digital identity and entitlement rights across security and enterprise boundaries. ADFS in Windows Server 2003 R2 supports the WS-F PRP. |
| claim | A statement that an issuer makes (for example, name, identity, key, group, privilege, or capability) about a client. |
| claim mapping | The act of mapping, removing or filtering, or passing claims between various claim sets. |
| claims-aware application | An ASP.NET application that performs authorization based on the claims that are present in an ADFS security token. |
| client account partner discovery Web page | The Web page that is used to interact with the user to determine which account partner the user belongs to when ADFS cannot automatically determine which of the account partners should authenticate the user. |
| client logoff Web page | When ADFS performs a logoff operation, a Web page that is executed to provide visual feedback to the user that the logoff has occurred. |
| client logon Web page | When ADFS collects client credentials, a Web page that is executed to perform the user interaction. The client logon Web page may use any necessary business logic to determine the type of credentials to collect. |
| federation | A pair of realms or domains that have established a federation trust. |
| Federation Service | A security token service that is built into Windows Server 2003 R2. The Federation Service provides tokens in response to requests for security tokens. |
| Federation Service Proxy | A proxy to the Federation Service in the perimeter network (also known as a DMZ or a screened subnet). The Federation Service Proxy uses WS-F PRP protocols to collect user credential information from browser clients and Web applications and send the information to the Federation Service on their behalf. |
| organization claims | Claims in intermediate or normalized form within an organization's namespace. |

| | |
|---|---|
| passive client | An HTTP browser, capable of broadly supported HTTP, that can make use of cookies. ADFS in Windows Server 2003 R2 supports only passive clients, and it adheres to the WS-F PRP specification. |
| resource partner | A federation partner that trusts the Federation Service to issue claims-based security tokens. The resource partner contains published Web-based applications that users in the account partner can access. |
| security token | A cryptographically signed data unit that expresses one or more claims. |
| security token service (STS) | A Web service that issues security tokens. An STS makes assertions based on evidence that it trusts, to whoever trusts it (or to specific recipients). To communicate trust, a service requires proof, such as a signature, to prove knowledge of a security token or set of security tokens. A service itself can generate tokens or it can rely on a separate STS to issue a security token with its own trust statement. This forms the basis of trust brokering. In ADFS, the Federation Service is an STS. |
| server farm | In ADFS, a collection of load-balanced federation servers, federation server proxies, or Web servers hosting the ADFS Web Agent. |
| single sign-on (SSO) | An optimization of the authentication sequence to remove the burden of repeated logon actions by an end user. |
| token-signing certificate | An X509 certificate whose associated public/private key pair is used to provide integrity for security tokens. |
| Uniform Resource Identifier (URI) | A compact string of characters that identifies an abstract resource or physical resource. URIs are explained in RFC 2396 (http://go.microsoft.com/fwlink/?LinkId=48289). In ADFS, URIs are used to uniquely identify partners and account stores. |
| Web Services (WS-*) | The specifications for a Web Services Architecture that is based on industry standards such as Simple Object Access Protocol (SOAP); XML; Web Service Description Language (WSDL); and Universal Description, Discovery, and Integration (UDDI). WS-* provides a foundation for delivering complete, interoperable business solutions for the extended enterprise, including the ability to manage federated identity and security.<br><br>The Web services model is based on the idea that enterprise systems are written in different languages, with different programming models, which run on and are accessed from many different types of devices. Web services are a means of building distributed systems that can connect and interact with one another easily and efficiently across the Internet, regardless of what language they are written in or what platform they run on. |
| Web Services Security (WS-Security) | A series of specifications that describes how to attach signature and encryption headers to SOAP messages. In addition, WS-Security describes how to attach security tokens, including binary security tokens such as X.509 certificates and Kerberos tickets, to messages. In ADFS, WS-Security is used when Kerberos signs security tokens. |
| Windows NT token–based application | A Windows application that relies on a Windows NT token to perform authorization of users. |
| WS-Federation | A specification that defines a model and set of messages for brokering trust and the federation of identity and authentication information across different trust realms.<br><br>The WS-Federation specification identifies two sources of identity and authentication requests across trust realms: active requestors, such as SOAP-enabled applications, and passive requestors, which are defined as HTTP browsers capable of supporting broadly supported HTTP, for example, HTTP 1.1. |
| WS-Federation Passive Requestor Profile (WS-F PRP) | An implementation of the WS-Federation specification that proposes a standard protocol for how passive clients (such as Web browsers) apply the federation framework. Within this protocol, Web service requestors are expected to understand the new security mechanisms and be capable of interacting with Web service providers. |

**ADFS resources**

You can find detailed documentation about Active Directory Federation Services (ADFS) on the ADFS home page (http://go.microsoft.com/fwlink/?LinkId=79542).

For more information about ADFS or for other ADFS-related information, see the following Web resources:

- Web Services and Other Distributed Technologies Developer Center on the MSDN Web site (http://go.microsoft.com/fwlink/?LinkId=44189)
- Web Services Protocol Workshops on the MSDN Web site (http://go.microsoft.com/fwlink/?LinkId=44190)
- Web Services Specifications on the MSDN Web site (http://go.microsoft.com/fwlink/?LinkId=44191)
- Web Services Interoperability Organization (WS-I) (http://go.microsoft.com/fwlink/?LinkId=34328)

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

| Tip |
| --- |
| Distributed File System is enhanced in Windows Server 2012. Explore the Evaluation Guide and download the Windows Server 2012 Trial. |

One of the goals of information technology (IT) groups in medium and large organizations is to manage file servers and their resources efficiently while keeping them available and secure for users. As organizations expand to include more users and servers—whether they are located in one site or in geographically distributed sites—administrators find it increasingly difficult to keep users connected to the files they need. On one hand, storing files on distributed servers makes files available to more users and decreases latency and bandwidth use when the servers are located near users. On the other hand, as the number of distributed servers increases, users have difficulty locating files they need, and operational costs increase.

Administrators who manage these distributed, remote servers need a solution that helps them limit network traffic over slow WAN connections, ensure the availability of files during WAN outages or server failures, and ensure that branch servers are backed up correctly. The Distributed File System solution in the Microsoft® Windows Server™ 2003 R2 operating system helps administrators address these challenges by providing two technologies, DFS Namespaces and DFS Replication, which, when used together, offer simplified, fault-tolerant access to files, load sharing, and WAN-friendly replication.

- DFS Replication is a new state-based, multimaster replication engine that supports replication scheduling and bandwidth throttling. DFS Replication uses a new compression protocol called Remote Differential Compression (RDC), which can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, and re-arrangements of data in files, thereby enabling DFS Replication to replicate only the changes when files are updated. Additionally, a function of RDC called *cross-file RDC* can help reduce the amount of bandwidth required to replicate new files.

- DFS Namespaces, formerly known as Distributed File System, allows administrators to group shared folders located on different servers and present them to users as a virtual tree of folders known as a *namespace*. A namespace provides numerous benefits, including increased availability of data, load sharing, and simplified data migration.

For the latest information about DFS Namespaces and DFS Replication, see DFS Namespaces and DFS Replication overview.

The following figure illustrates how DFS Namespaces and DFS Replication work together. The processes marked 1 and 2 are described in more detail following the figure.




As the figure shows, when a user attempts to access a folder in the namespace (1), the client computer contacts a namespace server. The namespace server sends the client computer a referral that contains a list of servers that host the shared folders (called folder targets) associated with the folder. The client computer caches the referral and then contacts the first server in the referral (2), typically a server in the client's own site unless no same-site servers exist or the administrator configures target priority.

The highlighted folder in the figure shows that it is hosted by shared folders on two servers, one in New York and one in London, to provide users in those sites with fast, reliable access to files. The shared folders are kept synchronized by DFS Replication. The fact that multiple servers host the folder is transparent to users, who see only a single folder in the namespace. If one of the servers becomes unavailable, the client computer fails over to the remaining server.

**Key Scenarios**
Although DFS Namespaces and DFS Replication are two separate technologies, when used together they provide solutions for a number of key scenarios in medium and large organizations.

**Data collection**
The data collection scenario helps address the need to eliminate the use of tape backup in branch offices and significantly reduce operational costs for branch offices. To accomplish this, data is replicated from a server in a branch office to a server in a hub office or data center. Administrators at the hub office can use backup software to back up the branch server's data from a hub server, eliminating the often error-prone process of having non-IT administrators performing the backups at branch offices that are not staffed by trained IT personnel. Centralizing backups at the hub office also allows organizations to consolidate backup hardware and related management tasks.

Thanks to RDC, DFS Replication replicates only the differences (or changes) between the two servers; as a result, bandwidth used during replication is minimized, an important consideration for branch offices that use low-bandwidth WAN connections to the hub office. In addition, replication schedules and bandwidth throttling can be used to set up replication windows during off-hours and regulate the amount of bandwidth used during replication, providing administrators more control over WAN traffic.

When DFS Replication is used in conjunction with DFS Namespaces, administrators can configure a namespace so that branch clients always connect to the branch server. If the branch server becomes unavailable, branch clients fail over to the hub server. And, using a DFS Namespaces enhancement known as *client failback*, branch clients running a client failback hotfix will fail back to the branch server after it is restored.

**Data distribution**
DFS Namespaces and DFS Replication can be used to publish documents, software, and line-of-business data to users throughout an organization. Although DFS Replication alone is sufficient to distribute data, using DFS Namespaces allows administrators to configure the namespace so that a folder in the namespace is hosted by multiple servers, which increases data availability and distributes the client load across servers. When browsing the namespace, users see a single folder and are not aware the folder is hosted by multiple servers. When a user opens the folder, the client computer is automatically referred to a server in its own site; if no same-site servers are available, administrators can configure the namespace to refer the client to a server that has the lowest connection cost as defined in the Active Directory® directory service. And, using a DFS Namespaces enhancement known as *target priority*, administrators can specify the priority of servers so that a specific server is always placed first or last in the list of servers (known as a referral) the client receives when accessing a folder in the namespace.

**Sharing files across branch offices**
In large organizations that have branch offices located around the world, users in one branch office often need access to shared folders and files stored at another branch or at the hub office or data center. Although in some cases the need is collaborative, in other cases users simply need access to those files. If low latency and data availability during WAN outages is worth the tradeoff of network traffic and disk space used in the branch office, administrators can use DFS Replication to replicate files between branch servers, providing users fast access to files in their respective branch offices. Having files in multiple branch offices also benefits users who travel from one branch office to another. The users can make changes to their files in one branch office, and those changes are replicated back to the user's branch office.

When the shared folders in a branch office are defined as folders in a namespace, users access them by browsing what appears to be a single server, giving the users a consistent place to begin browsing for files. For example, a user who travels from one branch office to another will access what appears to be the same folder in the namespace, but the user is actually referred to the closest server that contains the data. If that data is replicated using DFS Replication, the user's changes are replicated back to the user's branch office.

Collaboration scenarios using DFS Replication are recommended only if users can tolerate some inconsistency of the files as the changes are replicated to other branch servers. In addition, because DFS Replication uses a heuristic of last writer wins for changes made to the same file on multiple servers, this scenario works best when a file is updated from a single server or by a single user. Setting shared folder permissions can help ensure that conflicting changes don't occur. For collaboration scenarios that require file locking, we recommend using Microsoft® Windows® SharePoint® Services.

| **Note** |
| --- |
| Because DFS Replication only replicates a file after it is closed, DFS Replication is not recommended for replicating database files or any files that are held open for long periods of time. |

**DFS Replication Benefits in Windows Server 2003 R2**
DFS Replication is a new replication engine that provides substantial improvements over the File Replication service (FRS). The benefits of using DFS Replication are described below.

**Simplified process for replicating discrete folders to the same set of servers**
The process of setting up replicated folders is simplified in Windows Server 2003 R2 by the introduction of *replication groups* and *replicated folders*, which are illustrated in the following figure.




As this figure shows, a *replication group* is a set of servers, known as *members*, that participates in the replication of one or more replicated folders. A *replicated folder* is a folder that is kept synchronized on each member. In the previous figure, there are two replicated folders, Projects and Proposals. As data changes in each replicated folder, the changes are replicated across *connections* between the members. The connections between all members form the replication *topology*.

Creating multiple replicated folders in a single replication group simplifies the process of deploying replicated folders because the topology, schedule, and bandwidth throttling for the replication group are applied to each replicated folder. Each replicated folder also has its own settings, such as file and subfolder filters, so that administrators can filter out different files and subfolders for each replicated folder. To deploy additional replicated folders, administrators use a brief wizard to define the local path and permissions for each new replicated folder.

The replicated folders stored on each member can be located on different volumes in the member, and the replicated folders do not need to be shared folders or part of a namespace, though the DFS Management snap-in makes it easy to share replicated folders and optionally publish them in an existing namespace.

**Differential replication of changes to files**

DFS Replication uses RDC to replicate only the differences (or changes) between the members. This allows branch offices with slow WAN connections to participate in replication using minimal bandwidth.

RDC is especially efficient when small changes are made to large files. For example, a change that is made to a 2-megabyte (MB) PowerPoint® presentation can result in only 60 kilobytes (KB) being sent across the network, a 97 percent savings in bytes transferred. A test was run on a mix of 780 Office files (.doc, .ppt, and .xls) replicating from a source server to a target server using DFS Replication with RDC. The goal was to determine how effective RDC is for Office files on which typical edits are made. The target server had version *x* of the files and the source server had version *x+*, and the two versions differed with significant edits. The percent savings in bytes transferred was on average 50 percent and significantly better for large files.

Also, by using RDC and compression, DFS Replication provides bandwidth savings for distributing software files. For example, the IT department at Microsoft uses DFS Replication to replicate software files to Microsoft offices in several countries/regions. Although bandwidth used varies over time, in one typical case, a 63.14 percent reduction in bandwidth was measured in network traffic, specifically, 16.22 gigabytes (GB) replicated instead of 44.0 GB.

| **Note** |
| --- |
| RDC is not used on files smaller than 64 KB and might not be beneficial on high-speed LANs where network bandwidth is not contended. RDC can be disabled on a per-connection basis using the DFS Management snap-in. |

**Reduction of bandwidth used for replicating new files**

Another function of RDC, known as cross-file RDC, uses a heuristic to identify files that are similar to the file that needs to be replicated. This is useful when a new file is created on one server and needs to be replicated to another server. Instead of replicating the entire file, DFS Replication can use portions from files that are similar to the replicating file to minimize the amount of data transferred over the WAN.

Cross-file RDC is available only if one of the servers (in a replicating pair of servers) is running Windows Server 2003 R2, Enterprise Edition; Windows Server 2003 R2, Datacenter Edition; or Windows® Storage Server R2, Enterprise Edition. For example, if a branch office server is running Windows Server 2003 R2, Standard Edition, and the hub server is running Windows Server 2003 R2, Enterprise Edition, cross-file RDC will be used.

**Efficient and scalable replication**

When two members of a replication group begin to synchronize with each other, they use an efficient algorithm for determining which files need to be replicated. The amount of metadata exchanged is minimal, and the possibility of sending changes unnecessarily (due to the order the changes occur) is eliminated because the synchronization is state-based instead of event-based, as in File Replication service (FRS).

The introduction of state-based synchronization, along with RDC, allows DFS Replication to support replicating more files to more members than FRS. The tested scalability figures are as follows:

- Each server can be a member of up to 256 replication groups.
- Each replication group can contain up to 256 replicated folders.
- Each server can have up to 256 connections (for example, 128 incoming connections and 128 outgoing connections).
- On each server, the number of replication groups multiplied by the number of replicated folders multiplied by the number of simultaneously replicating connections must be kept to 1024 or fewer. If the replication schedule is staggered, you do not need to count the connections that are not replicating due to a closed schedule.
- A replication group can contain up to 256 members.
- A volume can contain up to 8 million replicated files, and a server can contain up to 1 terabyte of replicated files. These are tested numbers and recommended guidelines for performance and scalability reasons.

**Flexible scheduling and bandwidth throttling**
DFS Replication supports replication scheduling and bandwidth throttling in 15-minute increments during a 7-day period. When specifying a replication interval, administrators choose the start and stop times as well as the bandwidth to use during that interval. The settings for bandwidth usage range from 16 kilobits per second (Kbps) to 256 megabits per second (Mbps) as well as full (unlimited) bandwidth. Administrators can configure a default schedule and bandwidth that applies to all connections between members and, optionally, create a custom schedule and bandwidth for individual connections.

To help administrators configure replication windows for servers in different time zones, administrators can set the schedule so that the server that initiates replication interprets the schedule as being in either Coordinated Universal Time (UTC) or in the server's local time.

**Supported in stand-alone and domain-based namespaces and on individual folders**
DFS Replication can be used in both stand-alone and domain-based namespaces as well as for folders that are not part of any namespace. The folders to be replicated can be shared or unshared.

**Self-healing after USN journal wraps and database corruption**
DFS Replication provides "self-healing" for update sequence number (USN) journal wraps and Jet database corruption. Although replication temporarily stops during this healing process, the service recovers without any administrator intervention. To repair itself, DFS Replication scans the file system and re-creates the DFS Replication database that contains metadata associated with the files in the replicated folder. The database must then be synchronized with a database on another member. During the synchronization process, the amount of metadata sent across the network is dictated by the number of files under the replicated folder's local path (that is, the number of ID records in the database) and the size of the metadata to be sent per file. The metadata size for a file is two times the length of the file name plus approximately 144 bytes. Additional RPC and TCP overhead results in a roughly 5 percent overhead; therefore, in the worst case, for 1 million files in the database with an average file name size of 50 bytes, approximately 194 megabytes (MB) of metadata is sent across the network.

**Ease of member recovery**
DFS Replication stores its global configuration settings, such as the topology and replication schedule, in Active Directory. These settings are also cached in a local .xml file on each member; DFS Replication can rebuild this file (using the settings stored in Active Directory) if the file becomes corrupted or if the member is restored after a failure. This type of self-healing allows for greater server uptime and reliability and makes it easier to rebuild a member of a replication group during disaster recovery.
DFS Replication also uses the .xml file to store member-specific settings, such as debug log settings or RPC port settings that are configured using Windows Management Instrumentation (WMI).

**Simple and flexible prestaging of new servers**
Before adding a new server to a replication group, administrators can prestage the replicated folders on the target servers by copying the data to the servers, restoring a backup, or copying files from a tape, DVD, or removable hard disk. As described earlier, the synchronization process is very efficient in terms of bandwidth usage and metadata exchanged, resulting in minimal WAN traffic during the initial synchronization for files that are the same on the source server (known as the *primary member*) and target servers. If the files on the target server are out-of-date, DFS Replication will use RDC to replicate only the changes that occurred since the data was prestaged. Any prestaged files that existed on the target servers, but not on the source server, are moved to the PreExisting folder under the target server's replicated folder path.

**New management tools**
Administrators can use the DFS Management snap-in to configure both DFS Namespaces and DFS Replication. The snap-in provides integration between the two Distributed File System components so that administrators can:

- Select an existing folder in a namespace and configure DFS Replication on the folder targets (shared folders) associated with

the folder.

- Add a replicated folder to an existing namespace.

Administrators can also perform administrative tasks at the command line using Dfsradmin.exe or Dfsrdiag.exe. Both tools are part of Windows Server 2003 R2. Additional configurations can be made programmatically using WMI.

**Delegation of management tasks**
Administrators who are not part of the Domain Admins group can be delegated the ability to create new replication groups in a domain, manage an existing replication group, or both. Members of the Domain Admins group can use either the DFS Management snap-in or the Dfsradmin.exe command-line tool to perform this delegation.

**Built-in health metrics and diagnostic events**
DFS Replication provides built-in WMI providers for monitoring the health of DFS Replication. For example, the WMI providers can report USN journal wraps, database loss, insufficient disk space, network connectivity issues, sharing violations, excessive replication, and clock skew between members. These events are also reported in the DFS Replication event log, which is used exclusively for storing events related to replication.

There are two ways of monitoring DFS Replication: a built-in diagnostic report and the Windows DFS Replication Management Pack for Microsoft Operations Manager (MOM). The diagnostic report is an .html file generated using either the DFS Management snap-in or the Dfsradmin.exe command-line tool. The report contains a wealth of DFS Replication information, including error and warning events; service state and uptime; replication efficiency based on RDC and stream compression; backlogged sending and receiving transactions; free disk space, and more. An example health report is shown in the following figure:

# DFS Replication Health Report

(Show All)

| | |
|---|---|
| **Data collected on:** | Tuesday, October 04, 2005 at 7:00:04 PM (GMT-8:00) |
| **Replication Group:** | SOFTWARE (na.corp.contoso.com) |
| **Reference member:** | DDS-1 (DDS-1.na.corp.contoso.com) |
| **Server scope:** | Selected 30 of 30 servers |
| **DFS Replication efficiency savings:** | 66.98% reduction (36.53 GB replicated instead of 111 GB) |
| **Server health:** | Servers with no errors or warnings (14)<br>Servers with DFS Replication errors (5)<br>Servers unavailable for reporting (2)<br>Servers with DFS Replication warnings (14) |

**ERRORS (5 servers with errors)** (Show All)

- **B01** (1 error) (View Server Details)
  - DFS Replication service stopped replication on replicated folder Software\Betas\Developers\Visual Studio.
- **C01** (1 error) (View Server Details)
  - DFS Replication unable to replicate files for replicated folder Software\OS\Windows XP due to insufficent disk space.
- **H01** (1 error) (View Server Details)
  - DFS Replication unable to replicate files for replicated folder Software\OS\Windows Server 2003 due to insufficent disk space.
- **O01** (6 errors) (View Server Details)
- **SRV-1** (7 errors) (View Server Details)

**WARNINGS (14 servers with warnings)** (Show All)

**SERVERS UNAVAILABLE FOR REPORTING (2 servers not reporting)** (Show All)

**SERVER DETAILS (30 servers)** (Show All)

**C-SRV-1** (Show All)

| | |
|---|---|
| DNS name: | c-srv-1.europe.corp.contoso.com |
| Domain name: | europe.corp.contoso.com |
| Reference domain controller: | dc-01.europe.corp.contoso.com |
| IP address: | 151.67.11.58 |
| Site: | EU-CORP |
| Time zone: | (GMT-8:00) |

**ERRORS (There are no errors to report)**

**WARNINGS (There are no warnings to report)**

**INFORMATIONAL**

- Service state: Running
- DFS Replication service uptime: 4 days 9 hr.
- Summary of replicated folder status

  The following table provides a high-level overview of replicated folder status on this server.

| Replicated Folder | Status | Backlogged Sending Transactions | Backlogged Receiving Transactions | # of Files Received | DFS Replication Efficiency Savings |
|---|---|---|---|---|---|
| Software\Applications\Server\ISA Server 2004 | Normal | 120 | 289 | 568 | 93.13% |
| Software\Applications\Server\SQL Server | Normal | 84 | 416 | 632 | 65.54% |

The Windows DFS Replication Management Pack is a real-time monitoring tool that monitors the health state and replication progress of DFS Replication on a per-member basis. Administrators can use this management pack to monitor the state of the DFS Replication service, replication groups, replicated folders, and the volumes where replicated folders are stored. Conditions that cause full or partial replication failures will temporarily change the state of objects; these objects will change back to a healthy state when the problem is resolved, either automatically by the service (as in the case of intermittent connection failures) or by the administrator.

The management pack is supported on Microsoft Operations Manager 2005 and will be available for download on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55220).

**DFS Namespaces Enhancements in Windows Server 2003 R2**

As described earlier, the Distributed File System technology in Windows Server™ 2003 has been renamed to DFS Namespaces. Although the underlying service and basic functionality are unchanged, there are a number of DFS Namespaces enhancements exposed in Windows Server 2003 R2. These enhancements, introduced as updated APIs in Windows Server 2003 Service Pack 1 (SP1), provide easier management and more flexibility for namespaces used in branch offices.

### New and updated management tools

The new DFS Management snap-in in Windows Server 2003 R2 provides an improved graphical user interface for managing namespaces and DFS Replication. The snap-in allows administrators to configure DFS Namespaces enhancements, such as target priority, delegation, and client failback, as well as existing features that in Windows Server 2003 were configurable only by using Dfsutil.exe. For example, administrators can use the DFS Management snap-in to configure how servers are ordered in a referral, such as by lowest cost or restricted to the same site as the client. Administrators can also enable root scalability mode, which reduces the load on the primary domain controller (PDC) emulator in large namespaces.

The command-line tool Dfsutil.exe, which is part of the Windows Support Tools in Windows Server 2003 SP1, is updated to include the DFS Namespaces enhancements. The operating system tool Dfscmd.exe is also updated to allow administrators to move or rename folders in the namespace.

### Client failback

Client failover in DFS Namespaces is the process in which clients attempt to access another server in a referral after one of the servers fails or is removed from the namespace. Unless client failback is configured, clients will continue using the server they failed over to unless the client is restarted or the client's referral cache is cleared. When client failback is configured, and clients have the appropriate client failback hotfix installed, clients will fail back to a preferred, local server when it is restored.

### Target priority

When a client accesses a namespace, the client receives a referral that contains a list of targets associated with the namespace root or folder. These targets are listed according to the current ordering method for the namespace or folder. To fine-tune how particular targets are ordered, administrators can specify whether a server appears first or last in a referral. Assigning target priority is useful in many scenarios, such as “hot-standby” scenarios where one server is considered the server of last resort. In this scenario, the administrator can specify that the standby server always appears last in referrals, and clients will fail over to this server only if all the other servers fail or become unavailable due to network outages.

### Better delegation

Administrators can easily delegate the ability to create domain-based namespaces and manage individual stand-alone and domain-based namespaces. The DFS Management snap-in sets the appropriate permissions on either the DFS Namespace configuration objects in Active Directory or in the namespace server’s registry (depending on the namespace type).

### Ability to restructure the namespace

Renaming or moving folders in the namespace is easy using the DFS Management snap-in. Administrators can restructure the namespace to correct mistakes or to adjust the hierarchy as business needs change or as new folders are added to the namespace. Administrators can also move namespace folders by using the updated version of the command-line tool, Dfscmd.exe.

#### Configuration Requirements

The following sections describe how servers must be configured to support DFS Replication and DFS Namespaces in Windows Server 2003 R2.

### DFS Replication requirements

Before DFS Replication can be deployed, administrators must configure servers and storage as follows:

- The Active Directory schema must be updated to include the new DFS Replication objects. These schema changes are provided on disc 2 of the Windows Server 2003 R2 operating system installation discs. This schema can be applied to domain controllers running Microsoft® Windows® 2000 Server, Windows Server 2003, and Windows Server 2003 R2.

- The servers that will participate in DFS Replication must run Windows Server 2003 R2. After you install Windows Server 2003 R2, you must install the DFS Replication Service on each server that will take part in replication, and you must install the DFS Management snap-in to manage replication. For information about installing the DFS Management snap-in on a client computer running Windows XP with Service Pack 2 (SP2), see "Administering Distributed File System from a computer running Windows XP" later in this document.

- Antivirus software must be compatible with DFS Replication; contact your antivirus software vendor to check for compatibility.

- Servers in a replication group must be in the same forest. You cannot enable replication across servers in different forests.

- Replicated folders must be stored on NTFS volumes.

- On server clusters, replicated folders should be located in the local storage of a node because the Distributed File System Replication service is not cluster aware and the service will not fail over to another node.

| **Important** |
| --- |
| DFS Replication is not supported for SYSVOL replication in Windows Server 2003 R2. Do not attempt to configure DFS Replication on SYSVOL by disabling FRS and setting up a replication group for SYSVOL. Continue to use FRS for SYSVOL replication on domain controllers running Windows Server 2003 R2. FRS and DFS Replication can co-exist on the same member server or domain controller. |

**DFS Namespaces requirements**

To enable all features in DFS Namespaces, you must configure lab servers and clients as follows:

- Servers where the DFS Management snap-in will be used must run Windows Server 2003 R2 or Windows XP SP2. For information about installing the DFS Management snap-in on a client computer running Windows XP with SP2, see "Administering Distributed File System from a computer running Windows XP" later in this document.
- To support new namespace features, all servers that host namespaces must run Windows Server 2003 SP1 or Windows Server 2003 R2.
- To support new namespace features, all domain controllers must run Windows Server 2003 with SP1 or Windows Server 2003 R2.
- Namespaces must be created on NTFS volumes.
- Clients that access namespaces can run any of the supported client operating systems, but only clients running the following operating systems, service packs, and the appropriate client failback hotfix can be configured for client failback:
    - Windows XP SP2 and the Client Failback hotfix.
    - Windows Server 2003 SP1 and the Client Failback hotfix.

For more information about the Client Failback hotfix, see article 898900 in the Microsoft Knowledge Base on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=53202.

For more information about supported operating systems for DFS Namespaces clients, see "Evaluating Client and Server Compatibility" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=37455.

It is possible to use DFS Namespaces when domain controllers and namespace servers run a mix of Windows 2000 Server, Windows Server 2003 with SP1, and Windows Server 2003 without SP1, but some functionality is disabled or available inconsistently, depending on the operating systems on the servers. Some examples of mixed-mode behavior are as follows:

- If domain controllers or namespace servers are running Windows Server 2003 without SP1, they cannot provide referrals that support target priority or client failback.
- If domain controllers or namespace servers are running Windows 2000 Server, they cannot provide referrals that support target server priority or client failback, nor can they order targets by lowest cost in referrals. Additional configuration is required to enable these namespace servers and domain controllers to detect the site of each target server in the namespace. For details, see the answer to the question "What are the issues to consider when I use multiple servers running Windows 2000 Server and Windows Server 2003 to host a domain-based DFS root?" in the DFS FAQ on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=39465.
- If the DFS Management snap-in connects to a namespace server that is not running Windows Server 2003 SP1 or Windows Server 2003 R2, none of the new configuration settings (such as client failback and target priority) can be enabled. Renaming or moving folders will not work, and delegation will not be effective.

**Administering Distributed File System from a computer running Windows XP**

You can also manage DFS Namespaces and DFS Replication from a computer running Windows XP with SP2 by installing the Administration Tools Packs for Windows Server 2003 R2. When this pack is installed, the DFS Management snap-in is available as

part of the File Server Management snap-in. For more information about installing this pack, see the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=55225.

**See Also**
**Other Resources**
Microsoft Branch Offices Web Site
Branch Office TechCenter Web site
Distributed File System Web Site

# DFS Replication: Frequently Asked Questions (FAQ)

Published: October 16, 2006

Updated: October 9, 2013

Applies To: Windows Server 2003 R2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

This FAQ answers questions about Distributed File System (DFS) Replication (also known as DFS-R or DFSR) for the Windows Server 2012 R2, Windows Server 2012, Windows Server® 2008 R2, Windows Server 2008, and Windows Server 2003 R2 operating systems.

For information about DFS Namespaces, see DFS Namespaces: Frequently Asked Questions.

For information about what's new in DFS Replication, see the following topics:

- DFS Namespaces and DFS Replication Overview (in Windows Server 2012)
- What's New in Distributed File System topic in Changes in Functionality from Windows Server 2008 to Windows Server 2008 R2
- Distributed File System topic in Changes in Functionality from Windows Server 2003 with SP1 to Windows Server 2008

For a list of recent changes to this topic, see the Change History section of this topic.

**Interoperability**

- Can DFS Replication communicate with FRS?
- Can DFS Replication replace FRS for SYSVOL replication
- Can I upgrade from FRS to DFS Replication without losing configuration settings?
- Can I use DFS Replication in a mixed Windows/UNIX environment?
- Can I use the Volume Shadow Copy Service with DFS Replication?
- Can I use Windows Backup (Ntbackup.exe) to remotely back up a replicated folder?
- Do file system policies impact DFS Replication?
- Does DFS Replication replicate mailboxes hosted on Microsoft Exchange Server?
- Does DFS Replication support file screens created by File Server Resource Manager?
- Is DFS Replication cluster aware?
- Is DFS Replication compatible with Data Deduplication?
- Is DFS Replication compatible with RIS and WDS?
- Is it possible to use DFS Replication with Offline Files?
- What antivirus applications are compatible with DFS Replication?
- What are the benefits of using DFS Replication instead of Windows SharePoint Services?

**Limitations and requirements**

- Can DFS Replication replicate between branch offices without a VPN connection?
- Can DFS Replication replicate files encrypted with the Encrypting File System?
- Can DFS Replication replicate Outlook .pst or Microsoft Office Access database files?
- Can I use DFS Replication in a workgroup?
- Can more than one folder be replicated on a single server?
- Does DFS Replication require DFS Namespaces?
- Does DFS Replication require time synchronization between servers?
- Does DFS Replication support replicating an entire volume?
- Does DFS Replication support RPC over HTTP?
- Does DFS Replication work across wireless networks?
- Does DFS Replication work on ReFS or FAT volumes?
- Does DFS Replication work with sparse files?
- Do I need to log in as administrator to replicate files?
- How can I upgrade or replace a DFS Replication member?
- Is DFS Replication suitable for replicating roaming profiles?
- Is there a file character limit or limit to the folder depth?
- Must members of a replication group reside in the same domain?
- What are the supported limits of DFS Replication?
- When should I not use DFS Replication?
- Why is a schema update required for DFS Replication?

**Monitoring and management tools**

- Can I automate the health report to receive warnings?
- Can I use Microsoft System Center Operations Manager to monitor DFS Replication?
- Does DFS Replication support remote management?
- Do Ultrasound and Sonar work with DFS Replication?
- How can files be recovered from the ConflictAndDeleted or PreExisting folders?
- Is there a way to know the state of replication?

**Performance**

- Does DFS Replication support dial-up connections?
- Does DFS Replication perform bandwidth sensing?
- Does DFS Replication throttle bandwidth per schedule, per server, or per connection?
- Does DFS Replication use Active Directory Domain Services to calculate site links and connection costs?
- How can I improve replication performance?
- How does DFS Replication avoid saturating a connection?
- How does DFS Replication performance compare with FRS?
- How frequently does DFS Replication replicate data?
- How much of my server's system resources will DFS Replication consume?
- What happens if a WAN link fails during replication?

**Remote Differential Compression details**

- Are changes compressed before being replicated?
- Can an administrator turn off RDC or change the threshold?
- Does RDC work on all file types?
- How does RDC work on a compressed file?
- Is cross-file RDC enabled when upgrading to Windows Server Enterprise Edition or Datacenter Edition?
- Is RDC true block-level replication?
- What happens if I rename a file?
- What is cross-file RDC?
- What is RDC?
- When is RDC used for replication?
- Which editions of the Windows operating system support cross-file RDC?

**Replication details**

- Can I change the path for a replicated folder after it is created?
- Can I configure which file attributes are replicated?
- Can I control which member is replicated?
- Can I seed a replication group member with data prior to the initial replication?

- Does DFS Replication overcome common File Replication Service issues?
- Does DFS Replication replicate files in chronological order?
- Does DFS Replication replicate files that are being used by another application?
- Does DFS Replication replicate NTFS file permissions, alternate data streams, hard links, and reparse points?
- Does DFS Replication replicate timestamp changes if no other changes are made to the file?
- Does DFS Replication replicate updated permissions on a file or folder?
- Does DFS Replication support merging text files in the event of a conflict?
- Does DFS Replication use encryption when transmitting data?
- Is it possible to disable the use of encrypted RPC?
- How are simultaneous replications handled?
- How do I force replication or polling?
- Is it possible to configure a quiet time between replications for files that change frequently?
- Is it possible to configure one-way replication with DFS Replication?
- Is there a way to force a complete replication of all files including unchanged files?
- What happens if the primary member suffers a database loss during initial replication?
- What happens if the replication schedule closes while a file is being replicated?
- What happens when two users simultaneously update the same file on different servers?

**Staging**

- Does DFS Replication continue staging files when replication is disabled by a schedule or bandwidth throttling quota, or when a connection is manually disabled?
- Does DFS Replication prevent other applications from accessing a file during staging?
- Is it possible to change the location of the staging folder with the DFS Management Tool?
- When are files staged?
- What happens if a file is changed after it is staged but before it is completely transmitted to the remote site?

**Interoperability**

**Can DFS Replication communicate with FRS?**

No. DFS Replication does not communicate with File Replication Service (FRS). DFS Replication and FRS can run on the same server at the same time, but they must never be configured to replicate the same folders or subfolders because doing so can cause data loss.

**Can DFS Replication replace FRS for SYSVOL replication**

Yes, DFS Replication can replace FRS for SYSVOL replication on servers running Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008. Servers running Windows Server 2003 R2 do not support using

DFS Replication to replicate the SYSVOL folder.

For more information about replicating SYSVOL by using DFS Replication, see the SYSVOL Replication Migration Guide: FRS to DFS Replication.

**Can I upgrade from FRS to DFS Replication without losing configuration settings?**
Yes. To migrate replication from FRS to DFS Replication, see the following documents:

- To migrate replication of folders other than the SYSVOL folder, see DFS Operations Guide: Migrating from FRS to DFS Replication and FRS2DFSR – An FRS to DFSR Migration Utility (http://go.microsoft.com/fwlink/?LinkID=195437).

- To migrate replication of the SYSVOL folder to DFS Replication, see SYSVOL Replication Migration Guide: FRS to DFS Replication.

**Can I use DFS Replication in a mixed Windows/UNIX environment?**
Yes. Although DFS Replication only supports replicating content between servers running Windows Server, UNIX clients can access file shares on the Windows servers. To do so, install Services for Network File Systems (NFS) on the DFS Replication server.

You can also use the SMB/CIFS client functionality included in many UNIX clients to directly access the Windows file shares, although this functionality is often limited or requires modifications to the Windows environment (such as disabling SMB Signing by using Group Policy).

DFS Replication interoperates with NFS on a server running a Windows Server operating system, but you cannot replicate an NFS mount point.

**Can I use the Volume Shadow Copy Service with DFS Replication?**
Yes. DFS Replication is supported on Volume Shadow Copy Service (VSS) volumes and previous snapshots can be restored successfully with the Previous Versions Client.

**Can I use Windows Backup (Ntbackup.exe) to remotely back up a replicated folder?**
No, using Windows Backup (Ntbackup.exe) on a computer running Windows Server 2003 or earlier to back up the contents of a replicated folder on a computer running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 is not supported.

To back up files that are stored in a replicated folder, use Windows Server Backup or Microsoft® System Center Data Protection Manager. For information about Backup and Recovery functionality in Windows Server 2008 R2 and Windows Server 2008, see Backup and Recovery. For more information, see System Center Data Protection Manager (http://go.microsoft.com/fwlink/?LinkId=182261).

**Do file system policies impact DFS Replication?**
Yes. Do not configure file system policies on replicated folders. The file system policy reapplies NTFS permissions at every Group Policy refresh interval. This can result in sharing violations because an open file is not replicated until the file is closed.

**Does DFS Replication replicate mailboxes hosted on Microsoft Exchange Server?**
No. DFS Replication cannot be used to replicate mailboxes hosted on Microsoft Exchange Server.

**Does DFS Replication support file screens created by File Server Resource Manager?**
Yes. However, the File Server Resource Manager (FSRM) file screening settings must match on both ends of the replication. In addition, DFS Replication has its own filter mechanism for files and folders that you can use to exclude certain files and file types from replication.

The following are best practices for implementing file screens or quotas:

- The hidden DfsrPrivate folder must not be subject to quotas or file screens.

- Screened files must not exist in any replicated folder before screening is enabled.

- No folders may exceed the quota before the quota is enabled.

- You must use hard quotas with caution. It is possible for individual members of a replication group to stay within a quota before replication, but exceed it when files are replicated. For example, if a user copies a 10 megabyte (MB) file onto server A (which is then at the hard limit) and another user copies a 5 MB file onto server B, when the next replication occurs, both servers will exceed the quota by 5 megabytes. This can cause DFS Replication to continually retry replicating

the files, causing holes in the version vector and possible performance problems.

**Is DFS Replication cluster aware?**
Yes, DFS Replication in Windows Server 2012 R2, Windows Server 2012 and Windows Server 2008 R2 includes the ability to add a failover cluster as a member of a replication group. For more information, see Add a Failover Cluster to a Replication Group (http://go.microsoft.com/fwlink/?LinkId=155085).The DFS Replication service on versions of Windows prior to Windows Server 2008 R2 is not designed to coordinate with a failover cluster, and the service will not fail over to another node.

| **Note** |
| --- |
| DFS Replication does not support replicating files on Cluster Shared Volumes. |

**Is DFS Replication compatible with Data Deduplication?**
Yes, DFS Replication can replicate folders on volumes that use Data Deduplication in Windows Server.

**Is DFS Replication compatible with RIS and WDS?**
Yes. DFS Replication replicates volumes on which Single Instance Storage (SIS) is enabled. SIS is used by Remote Installation Services (RIS), Windows Deployment Services (WDS), and Windows Storage Server.

**Is it possible to use DFS Replication with Offline Files?**
You can safely use DFS Replication and Offline Files together in scenarios when there is only one user at a time who writes to the files. This is useful for users who travel between two branch offices and want to be able to access their files at either branch or while offline. Offline Files caches the files locally for offline use and DFS Replication replicates the data between each branch office.

Do not use DFS Replication with Offline Files in a multi-user environment because DFS Replication does not provide any distributed locking mechanism or file checkout capability. If two users modify the same file at the same time on different servers, DFS Replication moves the older file to the DfsrPrivate\ConflictandDeleted folder (located under the local path of the replicated folder) during the next replication.

**What antivirus applications are compatible with DFS Replication?**
Antivirus applications can cause excessive replication if their scanning activities alter the files in a replicated folder. For more information, Testing Antivirus Application Interoperability with DFS Replication (http://go.microsoft.com/fwlink/?LinkId=73990).

**What are the benefits of using DFS Replication instead of Windows SharePoint Services?**
Windows® SharePoint® Services provides tight coherency in the form of file check-out functionality that DFS Replication does not. If you are concerned about multiple people editing the same file, we recommend using Windows SharePoint Services. Windows SharePoint Services 2.0 with Service Pack 2 is available as part of Windows Server 2003 R2. Windows SharePoint Services can be downloaded from the Microsoft Web site; it is not included in newer versions of Windows Server. However, if you are replicating data across multiple sites and users will not edit the same files at the same time, DFS Replication provides greater bandwidth and simpler management.

**Limitations and requirements**
**Can DFS Replication replicate between branch offices without a VPN connection?**
Yes—assuming that there is a private Wide Area Network (WAN) link (not the Internet) connecting the branch offices. However, you must open the proper ports in external firewalls. DFS Replication uses the RPC Endpoint Mapper (port 135) and a randomly assigned ephemeral port above 1024. You can use the **Dfsrdiag** command line tool to specify a static port instead of the ephemeral port. For more information about how to specify the RPC Endpoint Mapper, see article 154596 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=73991).

**Can DFS Replication replicate files encrypted with the Encrypting File System?**
No. DFS Replication will not replicate files or folders that are encrypted using the Encrypting File System (EFS). If a user encrypts a file that was previously replicated, DFS Replication deletes the file from all other members of the replication group. This ensures that the only available copy of the file is the encrypted version on the server.

**Can DFS Replication replicate Outlook .pst or Microsoft Office Access database files?**
DFS Replication can safely replicate Microsoft Outlook personal folder files (.pst) and Microsoft Access files only if they are stored for archival purposes and are not accessed across the network by using a client such as Outlook or Access (to open .pst or Access files, first copy the files to a local storage device). The reasons for this are as follows:

- Opening .pst files over network connections could lead to data corruption in the .pst files. For more information about why .pst files cannot be safely accessed from across a network, see article 297019 in the Microsoft Knowledge Base

(http://go.microsoft.com/fwlink/?LinkId=125363).

- .pst and Access files tend to stay open for long periods of time while being accessed by a client such as Outlook or Office Access. This prevents DFS Replication from replicating these files until they are closed.

**Can I use DFS Replication in a workgroup?**
No. DFS Replication relies on Active Directory® Domain Services for configuration. It will only work in a domain.

**Can more than one folder be replicated on a single server?**
Yes. DFS Replication can replicate numerous folders between servers. Ensure that each of the replicated folders has a unique root path and that they do not overlap. For example, D:\Sales and D:\Accounting can be the root paths for two replicated folders, but D:\Sales and D:\Sales\Reports cannot be the root paths for two replicated folders.

**Does DFS Replication require DFS Namespaces?**
No. DFS Replication and DFS Namespaces can be used separately or together. In addition, DFS Replication can be used to replicate standalone DFS namespaces, which was not possible with FRS.

**Does DFS Replication require time synchronization between servers?**
No. DFS Replication does not explicitly require time synchronization between servers. However, DFS Replication does require that the server clocks match closely. The server clocks must be set within five minutes of each other (by default) for Kerberos authentication to function properly. For example, DFS Replication uses time stamps to determine which file takes precedence in the event of a conflict. Accurate times are also important for garbage collection, schedules, and other features.

**Does DFS Replication support replicating an entire volume?**
Yes. However, you must first install Windows Server 2003 Service Pack 2 or the hotfix. For more information, see article 920335 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=76776). Additionally, replicating an entire volume can cause the following problems:

- If the volume contains a Windows paging file, replication fails and logs DFSR event 4312 in the system event log.

- DFS Replication sets the System and Hidden attributes on the replicated folder on the destination server(s). This occurs because Windows applies the System and Hidden attributes to the volume root folder by default. If the local path of the replicated folder on the destination server(s) is also a volume root, no further changes are made to the folder attributes.

- When replicating a volume that contains the Windows system folder, DFS Replication recognizes the %WINDIR% folder and does not replicate it. However, DFS Replication does replicate folders used by non-Microsoft applications, which might cause the applications to fail on the destination server(s) if the applications have interoperability issues with DFS Replication.

**Does DFS Replication support RPC over HTTP?**
No.

**Does DFS Replication work across wireless networks?**
Yes. DFS Replication is independent of the connection type.

**Does DFS Replication work on ReFS or FAT volumes?**
No. DFS Replication supports volumes formatted with the NTFS file system only; the Resilient File System (ReFS) and the FAT file system are not supported. DFS Replication requires NTFS because it uses the NTFS change journal and other features of the NTFS file system.

**Does DFS Replication work with sparse files?**
Yes. You can replicate sparse files. The **Sparse** attribute is preserved on the receiving member.

**Do I need to log in as administrator to replicate files?**
No. DFS Replication is a service that runs under the local system account, so you do not need to log in as administrator to replicate. However, you must be a domain administrator or local administrator of the affected file servers to make changes to the DFS Replication configuration.

For more information, see "DFS Replication security requirements and delegation" in the Delegate the Ability to Manage DFS Replication (http://go.microsoft.com/fwlink/?LinkId= 182294).

**How can I upgrade or replace a DFS Replication member?**
To upgrade or replace a DFS Replication member, see this blog post on the Ask the Directory Services Team blog: Replacing

DFSR Member Hardware or OS.

**Is DFS Replication suitable for replicating roaming profiles?**
Yes. Certain scenarios are supported when replicating roaming user profiles. For information about the supported scenarios, see Microsoft's Support Statement Around Replicated User Profile Data (http://go.microsoft.com/fwlink/?LinkId=201282).

**Is there a file character limit or limit to the folder depth?**
Windows and DFS Replication support folder paths with up to 32 thousand characters. DFS Replication is not limited to folder paths of 260 characters.

**Must members of a replication group reside in the same domain?**
No. Replication groups can span across domains within a single forest but not across different forests.

**What are the supported limits of DFS Replication?**
The following list provides a set of scalability guidelines that have been tested by Microsoft on Windows Server 2012 R2:

- Size of all replicated files on a server: 100 terabytes.
- Number of replicated files on a volume: 70 million.
- Maximum file size: 250 gigabytes.

| **Important** |
| --- |
| When creating replication groups with a large number or size of files we recommend exporting a database clone and using pre-seeding techniques to minimize the duration of initial replication. For more information, see DFS Replication Initial Sync in Windows Server 2012 R2: Attack of the Clones. |

The following list provides a set of scalability guidelines that have been tested by Microsoft on Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008:

- Size of all replicated files on a server: 10 terabytes.
- Number of replicated files on a volume: 11 million.
- Maximum file size: 64 gigabytes.

| **Note** |
| --- |
| There is no longer a limit to the number of replication groups, replicated folders, connections, or replication group members. |

For a list of scalability guidelines that have been tested by Microsoft for Windows Server 2003 R2, see DFS Replication scalability guidelines (http://go.microsoft.com/fwlink/?LinkId=75043).

**When should I not use DFS Replication?**
Do not use DFS Replication in an environment where multiple users update or modify the same files simultaneously on different servers. Doing so can cause DFS Replication to move conflicting copies of the files to the hidden DfsrPrivate\ConflictandDeleted folder.

When multiple users need to modify the same files at the same time on different servers, use the file check-out feature of Windows SharePoint Services to ensure that only one user is working on a file. Windows SharePoint Services 2.0 with Service Pack 2 is available as part of Windows Server 2003 R2. Windows SharePoint Services can be downloaded from the Microsoft Web site; it is not included in newer versions of Windows Server.

**Why is a schema update required for DFS Replication?**
DFS Replication uses new objects in the domain-naming context of Active Directory Domain Services to store configuration information. These objects are created when you update the Active Directory Domain Services schema. For more information, see Review Requirements for DFS Replication (http://go.microsoft.com/fwlink/?LinkId=182264).

**Monitoring and management tools**
**Can I automate the health report to receive warnings?**
Yes. There are three ways to automate health reports:

- Use the DFSR Windows PowerShell module included in Windows Server 2012 R2 or DfsrAdmin.exe in conjunction with Scheduled Tasks to regularly generate health reports. For more information, see Automating DFS Replication Health Reports (http://go.microsoft.com/fwlink/?LinkId=74010).

- Use the DFS Replication Management Pack for System Center Operations Manager to create alerts that are based on specified conditions.

- Use the DFS Replication WMI provider to script alerts.

**Can I use Microsoft System Center Operations Manager to monitor DFS Replication?**
Yes. For more information, see the DFS Replication Management Pack for System Center Operations Manager 2007 in the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=182265).

**Does DFS Replication support remote management?**
Yes. DFS Replication supports remote management using the DFS Management console and the **Add Replication Group** command. For example, on server A, you can connect to a replication group defined in the forest with servers A and B as members.

DFS Management is included with Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003 R2. To manage DFS Replication from other versions of Windows, use Remote Desktop or the Remote Server Administration Tools for Windows 7.

| **Important** |
| --- |
| To view or manage replication groups that contain read-only replicated folders or members that are failover clusters, you must use the version of DFS Management that is included with Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, the Remote Server Administration Tools for Windows 8, or the Remote Server Administration Tools for Windows 7. |

**Do Ultrasound and Sonar work with DFS Replication?**
No. DFS Replication has its own set of monitoring and diagnostics tools. Ultrasound and Sonar are only capable of monitoring FRS.

**How can files be recovered from the ConflictAndDeleted or PreExisting folders?**
To recover lost files, restore the files from the file system folder or shared folder using File History, the **Restore previous versions** command in File Explorer, or by restoring the files from backup. To recover files directly from the ConflictAndDeleted or PreExisting folder, use the `Get-DfsrPreservedFiles` and `Restore-DfsrPreservedFiles` Windows PowerShell cmdlets (included with the DFSR module in Windows Server 2012 R2), or the RestoreDFSR sample script from the MSDN Code Gallery. This script is intended only for disaster recovery and is provided AS-IS, without warranty.

**Is there a way to know the state of replication?**
Yes. There are a number of ways to monitor replication:

- DFS Replication has a management pack for System Center Operations Manager that provides proactive monitoring.

- DFS Management has an in-box diagnostic report for the replication backlog, replication efficiency, and the number of files and folders in a given replication group.

- The DFSR Windows PowerShell module in Windows Server 2012 R2 contains cmdlets for starting propagation tests and writing propagation and health reports. For more information, see Distributed File System Replication Cmdlets in Windows PowerShell.

- Dfsrdiag.exe is a command-line tool that can generate a backlog count or trigger a propagation test. Both show the state of replication. Propagation shows you if files are being replicated to all nodes. Backlog shows you how many files still need to replicate before two computers are in sync. The backlog count is the number of updates that a replication group member has not processed. On computers running Windows Server 2012 R2, Windows Server 2012 or Windows Server 2008 R2, Dfsrdiag.exe can also display the updates that DFS Replication is currently replicating.

- Scripts can use WMI to collect backlog information—manually or through MOM.

**Performance**

**Does DFS Replication support dial-up connections?**

Although DFS Replication will work at dial-up speeds, it can get backlogged if there are large numbers of changes to replicate. If small changes are made to existing files, DFS Replication with Remote Differential Compression (RDC) will provide a much higher performance than copying the file directly.

**Does DFS Replication perform bandwidth sensing?**

No. DFS Replication does not perform bandwidth sensing. You can configure DFS Replication to use a limited amount of bandwidth on a per-connection basis (bandwidth throttling). However, DFS Replication does not further reduce bandwidth utilization if the network interface becomes saturated, and DFS Replication can saturate the link for short periods. Bandwidth throttling with DFS Replication is not completely accurate because DFS Replication throttles bandwidth by throttling RPC calls. As a result, various buffers in lower levels of the network stack (including RPC) may interfere, causing bursts of network traffic.

**Does DFS Replication throttle bandwidth per schedule, per server, or per connection?**

If you configure bandwidth throttling when specifying the schedule, all connections for that replication group will use that setting for bandwidth throttling. Bandwidth throttling can be also set as a connection-level setting using DFS Management.

**Does DFS Replication use Active Directory Domain Services to calculate site links and connection costs?**

No. DFS Replication uses the topology defined by the administrator, which is independent of Active Directory Domain Services site costing.

**How can I improve replication performance?**

To learn about different methods of tuning replication performance, see Tuning Replication Performance in DFSR on the Ask the Directory Services Team blog.

**How does DFS Replication avoid saturating a connection?**

In DFS Replication you set the maximum bandwidth you want to use on a connection, and the service maintains that level of network usage. This is different from the Background Intelligent Transfer Service (BITS), and DFS Replication does not saturate the connection if you set it appropriately.

Nonetheless, the bandwidth throttling is not 100% accurate and DFS Replication can saturate the link for short periods of time. This is because DFS Replication throttles bandwidth by throttling RPC calls. Because this process relies on various buffers in lower levels of the network stack, including RPC, the replication traffic tends to travel in bursts which may at times saturate the network links.

DFS Replication in Windows Server 2008 includes several performance enhancements, as discussed in the Distributed File System topic in Changes in Functionality from Windows Server 2003 with SP1 to Windows Server 2008.

**How does DFS Replication performance compare with FRS?**

DFS Replication is much faster than FRS, particularly when small changes are made to large files and RDC is enabled. For example, with RDC, a small change to a 2 MB PowerPoint® presentation can result in only 60 kilobytes (KB) being sent across the network—a 97 percent savings in bytes transferred.

RDC is not used on files smaller than 64 KB and might not be beneficial on high-speed LANs where network bandwidth is not contended. RDC can be disabled on a per-connection basis using DFS Management.

**How frequently does DFS Replication replicate data?**

Data replicates according to the schedule you set. For example, you can set the schedule to 15-minute intervals, seven days a week. During these intervals, replication is enabled. Replication starts soon after a file change is detected (generally within seconds).

The replication group schedule may be set to Universal Time Coordinate (UTC) while the connection schedule is set to the local time of the receiving member. Take this into account when the replication group spans multiple time zones. Local time means the time of the member hosting the inbound connection. The displayed schedule of the inbound connection and the corresponding outbound connection reflect time zone differences when the schedule is set to local time.

**How much of my server's system resources will DFS Replication consume?**

The disk, memory, and CPU resources used by DFS Replication depend on a number of factors, including the number and size of the files, rate of change, number of replication group members, and number of replicated folders. In addition, some resources are harder to estimate. For example, the Extensible Storage Engine (ESE) technology used for the DFS Replication database can consume a large percentage of available memory, which it releases on demand. Applications other than DFS Replication can be hosted on the same server depending on the server configuration. However, when hosting multiple applications or server roles on a single server, it is important that you test this configuration before implementing it in a production environment.

**What happens if a WAN link fails during replication?**

If the connection goes down, DFS Replication will keep trying to replicate while the schedule is open. There will also be

connectivity errors noted in the DFS Replication event log that can be harvested using MOM (proactively through alerts) and the DFS Replication Health Report (reactively, such as when an administrator runs it).

**Remote Differential Compression details**

**Are changes compressed before being replicated?**

Yes. Changed portions of files are compressed before being sent for all file types except the following (which are already compressed): .wma, .wmv, .zip, .jpg, .mpg, .mpeg, .m1v, .mp2, .mp3, .mpa, .cab, .wav, .snd, .au, .asf, .wm, .avi, .z, .gz, .tgz, and .frx. Compression settings for these file types are not configurable in Windows Server 2003 R2.

**Can an administrator turn off RDC or change the threshold?**

Yes. You can turn off RDC through the property page of a given connection. Disabling RDC can reduce CPU utilization and replication latency on fast local area network (LAN) links that have no bandwidth constraints or for replication groups that consist primarily of files smaller than 64 KB. If you choose to disable RDC on a connection, test the replication efficiency before and after the change to verify that you have improved replication performance.

You can change the RDC size threshold by using the **Dfsradmin Connection Set** command, the DFS Replication WMI Provider, or by manually editing the configuration XML file.

**Does RDC work on all file types?**

Yes. RDC computes differences at the block level irrespective of file data type. However, RDC works more efficiently on certain file types such as Word docs, PST files, and VHD images.

**How does RDC work on a compressed file?**

DFS Replication uses RDC, which computes the blocks in the file that have changed and sends only those blocks over the network. DFS Replication does not need to know anything about the contents of the file—only which blocks have changed.

**Is cross-file RDC enabled when upgrading to Windows Server Enterprise Edition or Datacenter Edition?**

The Standard Editions of Windows Server do not support cross-file RDC. However, it is automatically enabled when you upgrade to an edition that supports cross-file RDC, or if a member of the replication connection is running a supported edition. For a list of editions that support cross-file RDC, see Which editions of the Windows operating system support cross-file RDC?

**Is RDC true block-level replication?**

No. RDC is a general purpose protocol for compressing file transfer. DFS Replication uses RDC on blocks at the file level, not at the disk block level. RDC divides a file into blocks. For each block in a file, it calculates a signature, which is a small number of bytes that can represent the larger block. The set of signatures is transferred from server to client. The client compares the server signatures to its own. The client then requests the server send only the data for signatures that are not already on the client.

**What happens if I rename a file?**

DFS Replication renames the file on all other members of the replication group during the next replication. Files are tracked using a unique ID, so renaming a file and moving the file within the replica has no effect on the ability of DFS Replication to replicate a file.

**What is cross-file RDC?**

Cross-file RDC allows DFS Replication to use RDC even when a file with the same name does not exist at the client end. Cross-file RDC uses a heuristic to determine files that are similar to the file that needs to be replicated, and uses blocks of the similar files that are identical to the replicating file to minimize the amount of data transferred over the WAN. Cross-file RDC can use blocks of up to five similar files in this process.

To use cross-file RDC, one member of the replication connection must be running an edition of Windows that supports cross-file RDC. For a list of editions that support cross-file RDC, see Which editions of the Windows operating system support cross-file RDC?

**What is RDC?**

Remote differential compression (RDC) is a client-server protocol that can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, and rearrangements of data in files, enabling DFS Replication to replicate only the changes when files are updated. RDC is used only for files that are 64 KB or larger by default. RDC can use an older version of a file with the same name in the replicated folder or in the DfsrPrivate\ConflictandDeleted folder (located under the local path of the replicated folder).

**When is RDC used for replication?**

RDC is used when the file exceeds a minimum size threshold. This size threshold is 64 KB by default. After a file exceeding that threshold has been replicated, updated versions of the file always use RDC, unless a large portion of the file is changed or RDC is disabled.

**Which editions of the Windows operating system support cross-file RDC?**

To use cross-file RDC, one member of the replication connection must be running an edition of the Windows operating system

that supports cross-file RDC. The following table shows which editions of the Windows operating system support cross-file RDC.

**Cross-file RDC availability in editions of the Windows operating system**

| Operating System Version | Standard Edition | Enterprise Edition | Datacenter Edition |
|---|---|---|---|
| Windows Server 2012 R2 | Yes* | Not available | Yes* |
| Windows Server 2012 | Yes | Not available | Yes |
| Windows Server 2008 R2 | No | Yes | Yes |
| Windows Server 2008 | No | Yes | No |
| Windows Server 2003 R2 | No | Yes | No |

* You can optionally disable cross-file RDC on Windows Server 2012 R2.

**Replication details**

**Can I change the path for a replicated folder after it is created?**

No. If you need to change the path of a replicated folder, you must delete it in DFS Management and add it back as a new replicated folder. DFS Replication then uses Remote Differential Compression (RDC) to perform a synchronization that determines whether the data is the same on the sending and receiving members. It does not replicate all the data in the folder again.

**Can I configure which file attributes are replicated?**

No, you cannot configure which file attributes that DFS Replication replicates.

For a list of attribute values and their descriptions, see File Attributes on MSDN (http://go.microsoft.com/fwlink/?LinkId=182268).

The following attribute values are set by using the `SetFileAttributes dwFileAttributes` function, and they are replicated by DFS Replication. Changes to these attribute values trigger replication of the attributes. The contents of the file are not replicated unless the contents change as well. For more information, see SetFileAttributes Function in the MSDN library (http://go.microsoft.com/fwlink/?LinkId=182269).

- FILE_ATTRIBUTE_HIDDEN
- FILE_ATTRIBUTE_READONLY
- FILE_ATTRIBUTE_SYSTEM
- FILE_ATTRIBUTE_NOT_CONTENT_INDEXED
- FILE_ATTRIBUTE_OFFLINE

The following attribute values are replicated by DFS Replication, but they do not trigger replication.

- FILE_ATTRIBUTE_ARCHIVE
- FILE_ATTRIBUTE_NORMAL

The following file attribute values also trigger replication, although they cannot be set by using the `SetFileAttributes` function (use the `GetFileAttributes` function to view the attribute values).

- FILE_ATTRIBUTE_REPARSE_POINT

> **Note**
>
> DFS Replication does not replicate reparse point attribute values unless the reparse tag is IO_REPARSE_TAG_SYMLINK. Files with the IO_REPARSE_TAG_DEDUP, IO_REPARSE_TAG_SIS or IO_REPARSE_TAG_HSM reparse tags are replicated as normal files. However, the reparse tag and reparse data buffers are not replicated to other servers because the reparse point only works on the local system.

- FILE_ATTRIBUTE_COMPRESSED
- FILE_ATTRIBUTE_ENCRYPTED

| **Note** |
| --- |
| DFS Replication does not replicate files that are encrypted by using the Encrypting File System (EFS). DFS Replication does replicate files that are encrypted by using non-Microsoft software, but only if it does not set the FILE_ATTRIBUTE_ENCRYPTED attribute value on the file. |

- FILE_ATTRIBUTE_SPARSE_FILE
- FILE_ATTRIBUTE_DIRECTORY

DFS Replication does not replicate the FILE_ATTRIBUTE_TEMPORARY value.

**Can I control which member is replicated?**
Yes. You can choose a topology when you create a replication group. Or you can select **No topology** and manually configure connections after the replication group has been created.

**Can I seed a replication group member with data prior to the initial replication?**
Yes. DFS Replication supports copying files to a replication group member before the initial replication. This "prestaging" can dramatically reduce the amount of data replicated during the initial replication.

The initial replication does not need to replicate contents when files differ only by real attributes or time stamps. A real attribute is an attribute that can be set by the Win32 function `SetFileAttributes`. For more information, see SetFileAttributes Function in the MSDN library (http://go.microsoft.com/fwlink/?LinkId=182269). If two files differ by other attributes, such as compression, then the contents of the file are replicated.

To prestage a replication group member, copy the files to the appropriate folder on the destination server(s), create the replication group, and then choose a primary member. Choose the member that has the most up-to-date files that you want to replicate because the primary member's content is considered "authoritative." This means that during initial replication, the primary member's files will always overwrite other versions of the files on other members of the replication group.

For information about pre-seeding and cloning the DFSR database, see DFS Replication Initial Sync in Windows Server 2012 R2: Attack of the Clones.

For more information about the initial replication, see Create a Replication Group.

**Does DFS Replication overcome common File Replication Service issues?**
Yes. DFS Replication overcomes three common FRS issues:

- Journal wraps: DFS Replication recovers from journal wraps on the fly. Each existing file or folder will be marked as journalWrap and verified against the file system before replication is enabled again. During the recovery, this volume is not available for replication in either direction.
- Excessive replication: To prevent excessive replication, DFS Replication uses a system of credits.
- Morphed folders: To prevent morphed folder names, DFS Replication stores conflicting data in a hidden DfsrPrivate\ConflictandDeleted folder (located under the local path of the replicated folder). For example, creating multiple folders simultaneously with identical names on different servers replicated using FRS causes FRS to rename the older folder(s). DFS Replication instead moves the older folder(s) to the local Conflict and Deleted folder.

**Does DFS Replication replicate files in chronological order?**
No. Files may be replicated out of order.

**Does DFS Replication replicate files that are being used by another application?**
If an application opens a file and creates a file lock on it (preventing it from being used by other applications while it is open), DFS Replication will not replicate the file until it is closed. If the application opens the file with read-share access, the file can still be replicated.

**Does DFS Replication replicate NTFS file permissions, alternate data streams, hard links, and reparse points?**

- DFS Replication replicates NTFS file permissions and alternate data streams.
- Microsoft does not support creating NTFS hard links to or from files in a replicated folder – doing so can cause replication issues with the affected files. Hard link files are ignored by DFS Replication and are not replicated. Junction points also are not replicated, and DFS Replication logs event 4406 for each junction point it encounters.
- The only reparse points replicated by DFS Replication are those that use the IO_REPARSE_TAG_SYMLINK tag; however, DFS Replication does not guarantee that the target of a symlink is also replicated. For more information, see the Ask the Directory Services Team blog.
- Files with the IO_REPARSE_TAG_DEDUP, IO_REPARSE_TAG_SIS, or IO_REPARSE_TAG_HSM reparse tags are replicated as normal files. The reparse tag and reparse data buffers are not replicated to other servers because the reparse point only works on the local system. As such, DFS Replication can replicate folders on volumes that use Data Deduplication in Windows Server 2012, or Single Instance Storage (SIS), however, data deduplication information is maintained separately by each server on which the role service is enabled.

**Does DFS Replication replicate timestamp changes if no other changes are made to the file?**
No, DFS Replication does not replicate files for which the only change is a change to the timestamp. Additionally, the changed timestamp is not replicated to other members of the replication group unless other changes are made to the file.

**Does DFS Replication replicate updated permissions on a file or folder?**
Yes. DFS Replication replicates permission changes for files and folders. Only the part of the file associated with the Access Control List (ACL) is replicated, although DFS Replication must still read the entire file into the staging area.

| **Note** |
| --- |
| Changing ACLs on a large number of files can have an impact on replication performance. However, when using RDC, the amount of data transferred is proportionate to the size of the ACLs, not the size of the entire file. The amount of disk traffic is still proportional to the size of the files because the files must be read to and from the staging folder. |

**Does DFS Replication support merging text files in the event of a conflict?**
DFS Replication does not merge files when there is a conflict. However, it does attempt to preserve the older version of the file in the hidden DfsrPrivate\ConflictandDeleted folder on the computer where the conflict was detected.

**Does DFS Replication use encryption when transmitting data?**
Yes. DFS Replication uses Remote Procedure Call (RPC) connections with encryption.

**Is it possible to disable the use of encrypted RPC?**
No. The DFS Replication service uses remote procedure calls (RPC) over TCP to replicate data. To secure data transfers across the Internet, the DFS Replication service is designed to always use the authentication-level constant, `RPC_C_AUTHN_LEVEL_PKT_PRIVACY`. This ensures that the RPC communication across the Internet is always encrypted. Therefore, it is not possible to disable the use of encrypted RPC by the DFS Replication service.

For more information, see the following Microsoft Web sites:

- RPC Technical Reference
- About Remote Differential Compression
- Authentication-Level Constants

**How are simultaneous replications handled?**
There is one update manager per replicated folder. Update managers work independently of one another.

By default, a maximum of 16 (four in Windows Server 2003 R2) concurrent downloads are shared among all connections and replication groups. Because connections and replication group updates are not serialized, there is no specific order in which updates are received. If two schedules are opened, updates are generally received and installed from both connections at the same time.

**How do I force replication or polling?**
You can force replication immediately by using DFS Management, as described in Edit Replication Schedules. You can also force replication by using the `Sync-DfsReplicationGroup` cmdlet, included in the DFSR PowerShell module introduced with Windows Server 2012 R2, or the **Dfsrdiag SyncNow** command. You can force polling by using the `Update-DfsrConfigurationFromAD` cmdlet, or the **Dfsrdiag PollAD** command.

**Is it possible to configure a quiet time between replications for files that change frequently?**
No. If the schedule is open, DFS Replication will replicate changes as it notices them. There is no way to configure a quiet time for files.

**Is it possible to configure one-way replication with DFS Replication?**
Yes. If you are using Windows Server 2012 or Windows Server 2008 R2, you can create a read-only replicated folder that replicates content through a one-way connection. For more information, see Make a Replicated Folder Read-Only on a Particular Member (http://go.microsoft.com/fwlink/?LinkId=156740).

We do not support creating a one-way replication connection with DFS Replication in Windows Server 2008 or Windows Server 2003 R2. Doing so can cause numerous problems including health-check topology errors, staging issues, and problems with the DFS Replication database.

If you are using Windows Server 2008 or Windows Server 2003 R2, you can simulate a one-way connection by performing the following actions:

- Train administrators to make changes only on the server(s) that you want to designate as primary servers. Then let the changes replicate to the destination servers.

- Configure the share permissions on the destination servers so that end users do not have Write permissions. If no changes are allowed on the branch servers, then there is nothing to replicate back, simulating a one-way connection and keeping WAN utilization low.

**Is there a way to force a complete replication of all files including unchanged files?**
No. If DFS Replication considers the files identical, it will not replicate them. If changed files have not been replicated, DFS Replication will automatically replicate them when configured to do so. To overwrite the configured schedule, use the WMI method **ForceReplicate()**. However, this is only a schedule override, and it does not force replication of unchanged or identical files.

**What happens if the primary member suffers a database loss during initial replication?**
During initial replication, the primary member's files will always take precedence in the conflict resolution that occurs if the receiving members have different versions of files on the primary member. The primary member designation is stored in Active Directory Domain Services, and the designation is cleared after the primary member is ready to replicate, but before all members of the replication group replicate.

If the initial replication fails or the DFS Replication service restarts during the replication, the primary member sees the primary member designation in the local DFS Replication database and retries the initial replication. If the primary member's DFS Replication database is lost after clearing the primary designation in Active Directory Domain Services, but before all members of the replication group complete the initial replication, all members of the replication group fail to replicate the folder because no server is designated as the primary member. If this happens, use the **Dfsradmin membership /set /isprimary:true** command on the primary member server to restore the primary member designation manually.

For more information about initial replication, see Create a Replication Group.

| ⚠**Caution** |
| --- |
| The primary member designation is used only during the initial replication process. If you use the **Dfsradmin** command to specify a primary member for a replicated folder after replication is complete, DFS Replication does not designate the server as a primary member in Active Directory Domain Services. However, if the DFS Replication database on the server subsequently suffers irreversible corruption or data loss, the server attempts to perform an initial replication as the primary member instead of recovering its data from another member of the replication group. Essentially, the server becomes a rogue primary server, which can cause conflicts. For this reason, specify the primary member manually only if you are certain that the initial replication has irretrievably failed. |

**What happens if the replication schedule closes while a file is being replicated?**
If remote differential compression (RDC) is enabled on the connection, inbound replication of a file larger than 64 KB that began replicating immediately prior to the schedule closing (or changing to **No bandwidth**) continues when the schedule

opens (or changes to something other than **No bandwidth**). The replication continues from the state it was in when replication stopped.

If RDC is turned off, DFS Replication completely restarts the file transfer. This can delay when the file is available on the receiving member.

**What happens when two users simultaneously update the same file on different servers?**
When DFS Replication detects a conflict, it uses the version of the file that was saved last. It moves the other file into the DfsrPrivate\ConflictandDeleted folder (under the local path of the replicated folder on the computer that resolved the conflict). It remains there until Conflict and Deleted folder cleanup, which occurs when the Conflict and Deleted folder exceeds the configured size or DFS Replication encounters an Out of disk space error. The Conflict and Deleted folder is not replicated, and this method of conflict resolution avoids the problem of morphed directories that was possible in FRS.

When a conflict occurs, DFS Replication logs an informational event to the DFS Replication event log. This event does not require user action for the following reasons:

- It is not visible to users (it is visible only to server administrators).
- DFS Replication treats the Conflict and Deleted folder as a cache. When a quota threshold is reached, it cleans out some of those files. There is no guarantee that conflicting files will be saved.
- The conflict could reside on a server different from the origin of the conflict.

**Staging**
**Does DFS Replication continue staging files when replication is disabled by a schedule or bandwidth throttling quota, or when a connection is manually disabled?**
No. DFS Replication does not continue to stage files outside of scheduled replication times, if the bandwidth throttling quota has been exceeded, or when connections are disabled.

**Does DFS Replication prevent other applications from accessing a file during staging?**
No. DFS Replication opens files in a way that does not block users or applications from opening files in the replication folder. This method is known as "opportunistic locking."

**Is it possible to change the location of the staging folder with the DFS Management Tool?**
Yes. The staging folder location is configured on the **Advanced** tab of the **Properties** dialog box for each member of a replication group.

**When are files staged?**
Files are staged on the sending member when the receiving member requests the file (unless the file is 64 KB or smaller) as shown in the following table. If Remote Differential Compression (RDC) is disabled on the connection, the file is staged unless it is 256 KB or smaller. Files are also staged on the receiving member as they are transferred if they are less than 64 KB in size, although you can configure this setting between 16 KB and 1 MB. If the schedule is closed, files are not staged.

**The minimum file sizes for staging files**

| | RDC enabled | RDC disabled |
|---|---|---|
| Sending member | 64 KB | 256 KB |
| Receiving member | 64 KB by default | 64 KB by default |

**What happens if a file is changed after it is staged but before it is completely transmitted to the remote site?**
If any part of the file is already being transmitted, DFS Replication continues the transmission. If the file is changed before DFS Replication begins transmitting the file, then the newer version of the file is sent.

**Change History**

| Date | Description | Reason |
|---|---|---|
| October 9th, 2013 | Updated the What are the supported limits of DFS Replication? section with results from tests on Windows Server 2012 R2. | Updates for the latest version of Windows Server |
| January 30th, 2013 | Added the Does DFS Replication continue staging files when replication is disabled by a schedule or bandwidth throttling quota, or when a connection is manually disabled? entry. | Customer questions |

| | | |
|---|---|---|
| October 31st, 2012 | Edited the What are the supported limits of DFS Replication? entry to increase the tested number of replicated files on a volume. | Customer feedback |
| August 15, 2012 | Edited the Does DFS Replication replicate NTFS file permissions, alternate data streams, hard links, and reparse points? entry to further clarify how DFS Replication handles hard links and reparse points. | Feedback from Customer Support Services |
| June 13, 2012 | Edited the Does DFS Replication work on ReFS or FAT volumes? entry to add discussion of ReFS. | Customer feedback |
| April 25, 2012 | Edited the Does DFS Replication replicate NTFS file permissions, alternate data streams, hard links, and reparse points? entry to clarify how DFS Replication handles hard links. | Reduce potential confusion |
| March 30, 2011 | Edited the Can DFS Replication replicate Outlook .pst or Microsoft Office Access database files? entry to correct the potential impact of using DFS Replication with .pst and Access files.<br><br>Added How can I improve replication performance? | Customer questions about the previous entry, which incorrectly indicated that replicating .pst or Access files could corrupt the DFS Replication database. |
| January 26, 2011 | Added How can files be recovered from the ConflictAndDeleted or PreExisting folders? | Customer feedback |
| October 20, 2010 | Added How can I upgrade or replace a DFS Replication member? | Customer feedback |

# Security Configuration Wizard Quick Start Guide

Updated: March 2, 2005

Applies To: Windows Server 2003

This guide is designed to get you up and running quickly with Security Configuration Wizard (SCW), a tool for reducing the attack surface of computers running Windows Server® 2003 with Service Pack 1 (SP1). It provides system requirements, installation instructions, steps for getting started with SCW, and instructions for troubleshooting simple problems.

SCW determines the minimum functionality required for a server's role or roles, and disables functionality that is not required. Specifically, SCW:

- Disables unneeded services.
- Blocks unused ports.
- Allows further address or security restrictions for ports that are left open.
- Prohibits unnecessary Internet Information Services (IIS) Web extensions, if applicable.
- Reduces protocol exposure to server message block (SMB), LanMan, and Lightweight Directory Access Protocol (LDAP).
- Defines a high signal-to-noise audit policy.

SCW guides you through the process of creating, editing, applying, or rolling back a security policy based on the selected roles of the server. The security policies that are created with SCW are XML files that, when applied, configure services, network security, specific registry values, audit policy, and if applicable, Internet Information Services (IIS).

**Requirements for Installing and Running SCW**
SCW is an optional component included with Windows Server 2003 SP1. You can install and run SCW only on computers running Windows Server 2003 with SP1. The computers you target with SCW (for prototyping to create security policy or for application of SCW-created security policy) must also run Windows Server 2003 with SP1.

Several security-related IIS settings can be configured by using SCW. You need a server running IIS if you want to do this.

SCW is not used with Windows XP or other client operating systems or Microsoft Windows Small Business Server 2003.

**Securing Windows Small Business Server 2003**
Instead of SCW, Windows Small Business Server 2003 uses the default settings in Setup and in the Configure E-mail and Internet Connection Wizard to help secure your server.

If you have not already run the Configure E-mail and Internet Connection Wizard, you should run it to help secure your server.

### To start the Configure E-mail and Internet Connection Wizard on the computer running Windows Small Business Server 2003

1. Click **Start**, and then click **Server Management**.
2. In the console tree, click **Internet and E-mail**.
3. In the details pane, click **Connect to the Internet**.

**Getting Help**
SCW Help is installed with Windows Server 2003 SP1, and it contains information beyond what is in this Quick Start Guide, including help for every page of SCW. After you install Windows Server 2003 SP1, you can access SCW Help through Help and Support Center, or at the command line.

**Viewing SCW Help topics**
The SCW Help is available even though SCW itself is not installed by default.

### To access SCW help through Help and Support Center

1. Click **Start**, and then click **Help and Support**.
2. In **Search**, type **SCW** or type **Security Configuration Wizard**, and then press ENTER.
3. Click one of the listed SCW Help topics.

**To access SCW help at the command line**

1. Click **Start**, and then click **Run**.
2. Type **hh scwhelp.chm**, and then press ENTER.

| **Note** |
| --- |
| In some cases, you must be connected to the Internet to use the links in SCW Help. If your computer is not connected to the Internet, you can find the same topic in Help and Support Center by searching for the link text. |

# Getting Started with SCW

Updated: December 16, 2008

Applies To: Windows Server 2003

This topic provides information about getting started with SCW. For procedures to install and use SCW, see SCW Procedures.

**SCW Components**
There are three main components that you need to know about in order to get started using SCW:

- SCW User Interface
- Scwcmd command-line tool
- Security Configuration Database

**SCW User Interface**
SCW guides you through the process of creating a security policy, based on the roles performed by a given server. Once a policy is created, it can be edited or applied to one or more similarly configured servers. Applied policies can be rolled back in order to undo changes that have caused problems. To edit, apply, or roll back a security policy, you must have created the policy with SCW.

You can use the SCW user interface for the following tasks:

- Create a new security policy.
- Edit an existing SCW-generated security policy.
- Apply an existing SCW-generated security policy.
- Roll back the last applied SCW policy.

**Scwcmd Command-Line Tool**
SCW includes the Scwcmd.exe command-line tool. You can use Scwcmd for the following tasks:

- Configure one or many servers with an SCW-generated policy.
- Analyze one or many servers with an SCW-generated policy.
- View analysis results in HTML format.
- Roll back SCW policies.
- Transform an SCW-generated policy into native files that are supported by Group Policy.
- Register a Security Configuration Database extension with SCW.

When you use Scwcmd to configure, analyze, or roll back a policy on a remote server, SCW must be installed on the remote server.

**To get basic help about the Scwcmd tool**

1. Install SCW, as described in SCW Procedures.
2. Open a command prompt, and type:

   **Scwcmd**

**Security Configuration Database**

The Security Configuration Database consists of a set of XML documents that list services and ports that are required for each server role that is supported by SCW. These files are installed in %Systemroot%\Security\Msscw\KBs. After you select a server, on the **Processing Security Configuration Database** page, the server is scanned to determine the following:

- Roles that are installed on the server.
- Roles that are likely being performed by the server.
- Services that are installed but not part of the Security Configuration Database.
- IP addresses and subnets that are configured for the server.

SCW combines this server-specific information into a single XML file named Main.xml. The Security Configuration Wizard displays Main.xml if you click **View Security Configuration Database** on the **Processing Security Configuration Database** page.

The directory %Systemroot%\Security\Msscw\transformfiles contains .xsl transform files. These are applied to the .xml policy file for the rendering process when you view analysis results through the **scwcmd /view** command.

# Best Practices for SCW

Updated: March 2, 2005

Applies To: Windows Server 2003

This topic tells how to get the most out of SCW by using Best Practices.

This topic tells how to get the most out of SCW by using Best Practices.

**Identify and target similar servers**
SCW helps to reduce the attack surface of servers by creating a security policy that is specifically designed for their specific roles. Administrators can simplify policy authoring and distribution by identifying groups of servers that perform the same or similar tasks. Here are ways you can do this:

- **Author one policy for a group of servers**. SCW authors a security policy based on the roles, tasks, and functions performed by a server. Others servers that perform the same, or very similar, functions can be configured with the same security policy. Administrators can use SCW once to author a security policy, save it, and apply it to all servers that perform the job function.

- **Group similar servers in one organizational unit (OU) in Active Directory**. The SCW transform operation can apply a security policy to a domain or OU by using Group Policy. To simplify policy distribution, an administrator could group servers that perform similar job functions, and use the same security policy, into a single OU. A new security policy can be distributed quickly and easily to the server OU by using the SCW transform operation.

- **Create policies for similar platforms**. For services or ports specific to 64-bit computers, create the policies on a 64-bit computer. Then deploy these policies only to other 64-bit computers (not 32-bit computers) to ensure that the services are properly identified and configured.

- **Model your prototype server after the target servers to be configured at the service level**. It is highly recommended that the prototype server from which the security policy will be created matches the target servers to be configured at the service level. The security policy disables any service on the server that is contained in the Security Configuration Database but was not present on the prototype server when the policy was created.

  For example, if the DCOM Server Process Launcher service is listed in the Security Configuration Database, but is not present on the prototype server, the security policy created based on the prototype server will set the DCOM Server Process Launcher state to disabled. When you apply the security policy to other servers, the DCOM Server Process Launcher service will be disabled on those servers. You can configure unnecessary services in SCW (you can disable the service or leave the startup mode of the service unchanged), but only services that are not in the Security Configuration Database, and therefore are not defined in the security policy that you create with SCW.

**Test new security policies offline before deployment**
The settings configured in the new security policies may cause compatibility issues with applications or services. Therefore, thoroughly test new security policies in a test environment before applying the policies to production servers.

**Create one complete security policy**
SCW should be used to author a single security policy that contains all desired security settings for a server or a group of servers. This will simplify configuration, rollback, and analysis. For simple configuration and rollback, a single security policy for a computer, or group of computers, is much easier to understand and update than a series of policies. If a security policy defines all the desired settings for a server, a compliance report can be generated by running one scan, which facilitates analysis when using the **scwcmd /analyze** command. For more information about Scwcmd, see Getting Started with SCW.

# Deploying SCW

Updated: March 2, 2005

Applies To: Windows Server 2003

This topic discusses deploying SCW.

In addition to using Add or Remove Programs, there are two additional principal mechanisms for deploying SCW:

- **Microsoft Systems Management Server**. Systems Management Server can be used to deploy SCW to multiple servers.
- **Unattended installation of SCW**. Unattended installation of SCW is possible. See SCW Help for more information about unattended installation.

For detailed information about SCW deployment options, see SCW Deployment V2 on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=45183).

In addition to using SCW, you can use Group Policy to apply the security policies SCW creates to either a single server or multiple servers. To do this, you first create the policy by using SCW user interface, and then use the command-line tool to create a Group Policy-compatible version.

# SCW Procedures

Updated: March 2, 2005

Applies To: Windows Server 2003

This topic walks you through five basic SCW tasks.

This topic walks you through five basic SCW tasks.

**Basic SCW Tasks**
These five tasks are presented in the order in which they would typically be carried out.

1. Install SCW
2. Create a security policy based on a prototype server
3. Apply a security policy to a server
4. Analyze and view security policy for a server
5. Save an SCW security policy in native Group Policy format

The SCW user interface is used for the second and third tasks, and the Scwcmd command-line tool for the last two tasks, but the third and fourth tasks could be done by using either the command-line tool or the user interface.

**Install SCW**
After you have installed Windows Server 2003 SP1, you are ready to install SCW.

**To install SCW**

1. In Control Panel, double-click **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**, select the check box for **Security Configuration Wizard**, and then click **Next**.

| **Note** |
| --- |
| You can also install SCW on individual computers by using an unattended installation. Consult the SCW Help for information about unattended installation of SCW. You can deploy SCW to multiple computers by using Microsoft Systems Management Server. |

**Create a security policy based on a prototype server**
Prototype, or model, servers are always used as a basis for creating SCW security policies for groups of similarly configured servers. Use the following procedure to create a security policy that you can test on a prototype server before applying it to production servers.

**To create a security policy based on a prototype server**

1. Click **Start**, click **Administrative Tools**, and then click **Security Configuration Wizard**.
2. Read the **Welcome** page, and click **Next**.
3. Select **Create a new security policy**, and then click **Next**.
4. Type the name of the prototype server, and then click **Next**.
5. When processing is complete, click **Next**.
6. For each of the next five wizard pages, just click **Next**:
    - **Role-Based Service Configuration** page.

- **Select Server roles** page.
- **Select Client Features** page.
- **Select Administration and Other Options** page.
- **Select Additional Services** page.

7. On the **Handling Unspecified Services** page, select either **Do not change the startup mode of the service** (default) or **Disable the service**, and then click **Next**.

| **Note** |
| --- |
| The settings on the **Handling Unspecified Services** page control how SCW treats services that it finds on the prototype server, but that are not defined in the Security Configuration Database, and thus are not known to SCW. For information about extending the database, see Extending the Security Configuration Wizard on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=45183). |

8. For each of the next 20 wizard pages, just click **Next**:
   - **Confirm Service Changes** page.
   - **Network Security** page.
   - **Open Ports and Confirm Applications** page.
   - **Confirm Service Changes** page.
   - **Confirm Port Configuration** page.
   - **Registry Settings** page.
   - **Require SMB Security Signatures** page.
   - **Require LDAP Signing** page.
   - **Outbound Authentication Methods** page.
   - **Outbound Authentication Methods using Domain Accounts** page.
   - **Registry Settings Summary** page.
   - **Audit Policy** page.
   - **System Audit Policy** page.
   - **Audit Policy Summary** page.
   - **Internet Information Services** page.
   - **Select Web Service Extensions for Dynamic Content** page.
   - **Select the Virtual Directories to Retain** page.
   - **Prevent Anonymous Users from Accessing Content Files** page.

- **IIS Settings Summary** page.
- **Save Security Policy** page.

9. On the **Security Policy File Name** page, type a name for the prototype policy, and then click **Next**.

| **Caution** |
| --- |
| Do not use the name of the prototype computer because scwcmd.exe uses *computername.xml* to save analysis results, and you do not want the policy to be created to have the same name. |

| **Note** |
| --- |
| The security policy settings that you can configure within SCW are a subset of those that can be set by using security templates (.inf files). On the Security Policy File Name page, you can include a security template if you want to add settings that cannot be configured directly from SCW. If you attach a security template, and it contains settings that conflict with some SCW-configured settings, the SCW-configured settings have precedence. |

10. On the **Completing the Security Configuration Wizard** page, click **Finish**.

#### Apply a security policy to a server

The following procedure can apply a security policy to either a single server or multiple servers.

**To apply a security policy to a server**

1. Click **Start**, click **Administrative Tools**, and then click **Security Configuration Wizard**.
2. Read the **Welcome** page, and then click **Next**.
3. On the **Configuration Action** page, select **Apply an existing security policy**, type the full path and file name of the policy, and then click **Next**.
4. On the **Select Server** page, type the name of the server to which the policy will be applied, and then click **Next**.

| **Note** |
| --- |
| To configure multiple servers with a policy, you can use **scwcmd configure /p:***PolicyFile* **/i:***MachineList* at the command prompt, rather than this SCW UI procedure. Type **scwcmd configure** at the command prompt to learn about the parameters. |

5. On the **Apply Security Policy** page, click **Next**.
6. On the **Applying Security Policy** page, wait for processing to finish, and then click **Next**.
7. On the **Completing the Security Configuration Wizard** page, click **Finish**.

#### Analyze and view security policy for a server

Use the following procedure to analyze and view security policy for a computer from the command line.

**To analyze and view security policy for a server**

1. At the command prompt, type

   **scwcmd analyze /m:** *MachineName* **/p:** *PathAndPolicyFileName* **/o:** *OutputDirectory*

| **Note** |
| --- |
| You should first replace the italic parameters with your specific ones. When **scwcmd analyze** is finished processing, you will find that it has saved *MachineName.xml*. This is the analysis result for that server, saved as XML. |

2. When **scwcmd analyze** processing is complete, type:

   **scwcmd view /x:** *MachineName.xml* **/s:scwanalysis.xsl**

Scwanalysys.xsl is one of the files installed with SCW. It formats the analysis results for display.

**Save an SCW security policy in native Group Policy format**

The following procedure makes the security policy available for used in Group Policy.

| **Note** |
| --- |
| You might not always decide to save SCW security policy in Group Policy format, because Security policy applied through Group Policy cannot be rolled back. |

**To save an SCW security policy in native Group Policy format**

- At the command prompt, type:

  **scwcmd transform /p:** *PathAndPolicyFileName* **/g:** *GPODisplayName*

  where *PathAndPolicyFileName* is the policy you created earlier with SCW, including its .xml file name extension and *GPODisplayName* is the name that the Group Policy object (GPO) will show when you view it in Group Policy Object Editor or in Group Policy Management Console (GPMC).

When the **scwcmd transform** command has completed, the GPO will have been created in Active Directory, but the policy it contains will not be applied until the GPO is linked to a site, domain, or organizational unit. For instructions about linking GPOs, see the GPMC Help.

# Troubleshooting SCW

Updated: March 2, 2005

Applies To: Windows Server 2003

This topic discusses troubleshooting SCW.

This topic lists some problems and error messages and provides suggestions for handling them. For more detailed information about troubleshooting SCW, see SCW Troubleshooting Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=45183).

**Common SCW Problems**
This section lists circumstances under which SCW might not work as you expect.

**SCW does not run on a specific computer**
SCW is supported only on Windows Server 2003 with SP1. For details about SCW requirements, see "Requirements for Installing and Running SCW" in Security Configuration Wizard Quick Start Guide.

**SCW security policies do not apply to some computers**
The security policies that are created with SCW should only be applied to servers and groups of servers running Windows Server 2003 with SP1. The security policies should not be applied to operating systems such as Windows XP or Windows Small Business Server.

**IIS role is not detected by SCW**
If an SCW security policy is applied to a server before Internet Information Services (IIS) is installed, the server will be configured with the Hypertext Transport Protocol (HTTP) Secure Sockets Layer (SSL) service disabled, as HTTP SSL is not required on a computer that is not a Web server. If IIS is then installed on the server and SCW is run, the Administrator must first ensure that IIS is started, in order for IIS to be detected by SCW. IIS cannot run without the HTTP SSL service.

**Error messages**
Here are some errors you might encounter with SCW.

**Cannot Process XML Data**
The Microsoft XML core services are corrupted or not installed.

**Cannot Process Security Configuration Database**
There is a problem with Main.XML. Check the Application log, which might have more information.

**The Security Configuration Wizard cannot continue because the security configuration database processing failed**
There is a problem with Main.XML. Check the Application log, which might have more information.

**You do not have administrator privileges on the selected server. Click Specify User Account to provide an administrator account on the selected server**
You must be an administrator on the local server to run SCW, or you must specify Administrative credentials on a remote server. If you are trying to apply a policy to a server, the policy must be available on the server running SCW. You must also be authenticated as an administrator on the computer that is receiving the policy.

**Cannot rollback last applied security policy**
Check the Application log, which might have more information.

**The selected security policy file has a format that is not valid. Select another security policy file and try again**
You can receive this message if you browse to, or type in, an XML file that is not a security policy file. If this is a security policy file that you created, try running SCW again to recreate it.

**Not Found!**
In the Network Security section of SCW, this message appears if SCW cannot verify that an approved application exists at the specified path. This is not necessarily a problem, as you might be configuring a policy for a remote server. If the application will exist at the specified path at the time the policy is applied, the Not Found! message can be ignored.

**Cannot determine the IP address from the computer name. The lookup service is not available**
In the Network Security section of SCW, the IP address of the computer name you entered must be resolvable, or you cannot continue. If the computer has a static IP address (not assigned by DHCP), then you can enter the IP address instead of the name.

# Quick Start Guide for Server Clusters

Applies To: Windows Server 2003 R2

This guide provides system requirements, installation instructions, and other, step-by-step instructions that you can use to deploy server clusters if you are using Microsoft® Windows Server™ 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, operating systems.

The server cluster technology in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, helps ensure that you have access to important server-based resources. You can use server cluster technology to create several cluster nodes that appear to users as one server. If one of the nodes in the cluster fails, another node begins to provide service. This is a process known as "failover." In this way, server clusters can increase the availability of critical applications and resources.

**Copyright**
This document is provided for informational purposes only and Microsoft makes no warranties, either express or implied, in this document. Information in this document, including URL and other Internet Web site references, is subject to change without notice. The entire risk of the use or the results from the use of this document remains with the user. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

Copyright © 2005 Microsoft Corporation. All rights reserved.

Microsoft, Windows, Windows NT, SQL Server, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

**Requirements and Guidelines for Configuring Server Clusters**
The section lists requirements and guidelines that will help you set up a server cluster effectively.

**Software requirements and guidelines**

- You must have either Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, installed on all computers in the cluster. We strongly recommend that you also install the latest service pack for Windows Server 2003. If you install a service pack, the same service pack must be installed on all computers in the cluster.

- All nodes in the cluster must be of the same architecture. You cannot mix x86-based, Itanium-based, and x64-based computers within the same cluster.

- Your system must be using a name-resolution service, such as Domain Name System (DNS), DNS dynamic update protocol, Windows Internet Name Service (WINS), or Hosts file. Hosts file is supported as a local, static file method of mapping DNS domain names for host computers to their Internet Protocol (IP) addresses. The Hosts file is provided in the systemroot\System32\Drivers\Etc folder.

- All nodes in the cluster must be in the same domain. As a best practice, all nodes should have the same domain role (either member server or domain controller), and the recommended role is member server. Exceptions that can be made to these domain role guidelines are described later in this document.

- When you first create a cluster or add nodes to it, you must be logged on to the domain with an account that has administrator rights and permissions on all nodes in that cluster. The account does not need to be a Domain Admin level account, but can be a Domain User account with Local Admin rights on each node.

**Hardware requirements and guidelines**

- For Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, Microsoft supports only complete server cluster systems chosen from the Windows Catalog. To determine whether your system and hardware components are compatible, including your cluster disks, see the Microsoft Windows Catalog at the Microsoft Web site. For a geographically dispersed cluster, both the hardware and software configuration must be certified and listed in the Windows Catalog. For more information, see article 309395, "The Microsoft support policy for server clusters, the Hardware Compatibility List, and the Windows Server Catalog," in the Microsoft Knowledge Base.

- If you are installing a server cluster on a storage area network (SAN), and you plan to have multiple devices and clusters sharing the SAN with a cluster, your hardware components must be compatible. For more information, see article 304415, "Support for Multiple Clusters Attached to the Same SAN Device," in the Microsoft Knowledge Base.

- You must have two mass-storage device controllers in each node in the cluster: one for the local disk, one for the cluster storage. You can choose between SCSI, iSCSI, or Fibre Channel for cluster storage on server clusters that are running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. You must have two controllers because one controller has the local system disk for the operating system installed, and the other controller has the shared storage installed.

- You must have two Peripheral Component Interconnect (PCI) network adapters in each node in the cluster.

- You must have storage cables to attach the cluster storage device to all computers. Refer to the manufacturer's instructions for configuring storage devices.

- Ensure that all hardware is identical in all cluster nodes. This means that each hardware component must be the same make, model, and firmware version. This makes configuration easier and eliminates compatibility problems.

**Network requirements and guidelines**

- Your network must have a unique NetBIOS name.

- A WINS server must be available on your network.

- You must use static IP addresses for each network adapter on each node.

| **Important** |
| --- |
| Server clusters do not support the use of IP addresses assigned from Dynamic Host Configuration Protocol (DHCP) servers. |

- The nodes in the cluster must be able to access a domain controller. The Cluster service requires that the nodes be able to contact the domain controller to function correctly. The domain controller must be highly available. In addition, it should be on the same local area network (LAN) as the nodes in the cluster. To avoid a single point of failure, the domain must have at least two domain controllers.

- Each node must have at least two network adapters. One adapter will be used exclusively for internal node-to-node communication (the private network). The other adapter will connect the node to the client public network. It should also connect the cluster nodes to provide support in case the private network fails. (A network that carries both public and private communication is called a mixed network.)

- If you are using fault-tolerant network cards or teaming network adapters, you must ensure that you are using the most recent firmware and drivers. Check with your network adapter manufacturer to verify compatibility with the cluster technology in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

| **Note** |
| --- |

Using teaming network adapters on all cluster networks concurrently is not supported. At least one of the cluster private networks must not be teamed. However, you can use teaming network adapters on other cluster networks, such as public networks.

**Storage requirements and guidelines**

- An external disk storage unit must be connected to all nodes in the cluster. This will be used as the cluster storage. You should also use some type of hardware redundant array of independent disks (RAID).

- All cluster storage disks, including the quorum disk, must be physically attached to a shared bus.

**Note**

This requirement does not apply to Majority Node Set (MNS) clusters when they are used with some type of software replication method.

- Cluster disks must not be on the same controller as the one that is used by the system drive, except when you are using boot from SAN technology. For more information about using boot from SAN technology, see "Boot from SAN in Windows Server 2003 and Windows 2000 Server" at the Microsoft Web site.

- You should create multiple logical unit numbers (LUNs) at the hardware level in the RAID configuration instead of using a single logical disk that is then divided into multiple partitions at the operating system level. We recommend a minimum of two logical clustered drives. This enables you to have multiple disk resources and also allows you to perform manual load balancing across the nodes in the cluster.

- You should set aside a dedicated LUN on your cluster storage for holding important cluster configuration information. This information makes up the cluster quorum resource. The recommended minimum size for the volume is 500 MB. You should not store user data on any volume on the quorum LUN.

- If you are using SCSI, ensure that each device on the shared bus (both SCSI controllers and hard disks) has a unique SCSI identifier. If the SCSI controllers all have the same default identifier (the default is typically SCSI ID 7), change one controller to a different SCSI ID, such as SCSI ID 6. If more than one disk will be on the shared SCSI bus, each disk must also have a unique SCSI identifier.

- Software fault tolerance is not natively supported for disks in the cluster storage. For cluster disks, you must use the NTFS file system and configure the disks as basic disks with all partitions formatted as NTFS. They can be either compressed or uncompressed. Cluster disks cannot be configured as dynamic disks. In addition, features of dynamic disks, such as spanned volumes (volume sets), cannot be used without additional non-Microsoft software.

- All disks on the cluster storage device must be partitioned as master boot record (MBR) disks, not as GUID partition table (GPT) disks.




**Deploying SANs with server clusters**

This section lists the requirements for deploying SANs with server clusters.

- Nodes from different clusters must not be able to access the same storage devices. Each cluster used with a SAN must be deployed in a way that isolates it from all other devices. This is because the mechanism the cluster uses to protect access to the disks can have adverse effects if other clusters are in the same zone. Using zoning to separate the cluster traffic from other cluster or non-cluster traffic prevents this type of interference. For more information, see "Zoning vs. LUN masking" later in this guide.

- All host bus adapters in a single cluster must be the same type and have the same firmware version. Host bus adapters are the interface cards that connect a cluster node to a SAN. This is similar to the way that a network adapter connects a server to a typical Ethernet network. Many storage vendors require that all host bus adapters on the same zone—and, in some cases, the same fabric—share these characteristics.

- In a cluster, all device drivers for storage and host bus adapters must have the same software version. We strongly recommend that you use a Storport mini-port driver with clustering. Storport (Storport.sys) is a storage port driver that is provided in Windows Server 2003. It is especially suitable for use with high-performance buses, such as Fibre Channel buses, and RAID adapters.

- Tape devices should never be used in the same zone as cluster disk storage devices. A tape device could misinterpret a bus reset and rewind at inappropriate times, such as when backing up a large amount of data.

- In a highly available storage fabric, you should deploy server clusters with multiple host bus adapters using multipath I/O software. This provides the highest level of redundancy and availability.

| **Note** |
| --- |
| Failover software for host bus adapters can be version sensitive. If you are implementing a multipath solution for your cluster, you should work closely with your hardware vendor to become fully aware of how the adapter interacts with Windows Server 2003. |

**Creating a Cluster**

It is important to plan the details of your hardware and network before you create a cluster.

If you are using a shared storage device, ensure that when you turn on the computer and start the operating system, only one

node has access to the cluster storage. Otherwise, the cluster disks can become corrupted.

In Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, logical disks that are not on the same shared bus as the boot partition are not automatically mounted and assigned a drive letter. This helps prevent a server in a complex SAN environment from mounting drives that might belong to another server. (This is different from how new disks are mounted in Microsoft Windows® 2000 Server operating systems.) Although the drives are not mounted by default, we still recommend that you follow the procedures provided in the table later in this section to ensure that the cluster disks will not become corrupted.

The table in this section can help you determine which nodes and storage devices should be turned on during each installation step. The steps in the table pertain to a two-node cluster. However, if you are installing a cluster with more than two nodes, the Node 2 column lists the required state of all other nodes.

| Step | Node 1 | Node 2 | Storage | Notes |
|---|---|---|---|---|
| Set up networks | On | On | Off | Verify that all storage devices on the shared bus are turned off. Turn on all nodes. |
| Set up cluster disks | On | Off | On | Shut down all nodes. Turn on the cluster storage, and then turn on the first node. |
| Verify disk configuration | Off | On | On | Turn off the first node, turn on second node. Repeat for nodes three and four if necessary. |
| Configure the first node | On | Off | On | Turn off all nodes; and then turn on the first node. |
| Configure the second node | On | On | On | After the first node is successfully configured, turn on the second node. Repeat for nodes three and four as necessary. |
| Post-installation | On | On | On | All nodes should be turned on. |

**Preparing to create a cluster**

Complete the following three steps on each cluster node before you install a cluster on the first node.

- Install Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, on each node of the cluster. We strongly recommend that you also install the latest service pack for Windows Server 2003. If you install a service pack, the same service pack must be installed on all computers in the cluster.
- Set up networks.
- Set up cluster disks.

All nodes must be members of the same domain. When you create a cluster or join nodes to a cluster, you specify the domain user account under which the Cluster service runs. This account is called the Cluster service account (CSA).

**Installing the Windows Server 2003 operating system**

Install Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, on each node of the cluster. For information about how to perform this installation, see the documentation you received with the operating system.

Before configuring the Cluster service, you must be logged on locally with a domain account that is a member of the local administrators group.

| **Important** |
|---|
| If you attempt to join a node to a cluster that has a blank password for the local administrator account, the installation will fail. For security reasons, Windows Server 2003 operating systems prohibit blank administrator passwords. |

**Setting up networks**

Each cluster node requires at least two network adapters and must be connected by two or more independent networks. At least two LAN networks (or virtual LANs) are required to prevent a single point of failure. A server cluster whose nodes are connected by only one network is not a supported configuration. The adapters, cables, hubs, and switches for each network must fail independently. This usually means that the components of any two networks must be physically independent.

Two networks must be configured to handle either **All communications (mixed network)** or **Internal cluster communications only (private network)**. The recommended configuration for two adapters is to use one adapter for the private (node-to-node only) communication and the other adapter for mixed communication (node-to-node plus client-to-cluster communication).

You must have two PCI network adapters in each node. They must be certified in the Microsoft Windows Catalog and supported by Microsoft Product Support Services. Assign one network adapter on each node a static IP address, and assign the other network adapter a static IP address on a separate network on a different subnet for private network communication.

Because communication between cluster nodes is essential for smooth cluster operations, the networks that you use for cluster communication must be configured optimally and follow all hardware compatibility-list requirements. For additional information about recommended configuration settings, see article 258750, "Recommended private heartbeat configuration on a cluster server," in the Microsoft Knowledge Base.

You should keep all private networks physically separate from other networks. Specifically, do not use a router, switch, or bridge to join a private cluster network to any other network. Do not include other network infrastructure or application servers on the private network subnet. To separate a private network from other networks, use a cross-over cable in a two-node cluster configuration or a dedicated hub in a cluster configuration of more than two nodes.

**Additional network considerations**

- All cluster nodes must be on the same logical subnet.
- If you are using a virtual LAN (VLAN), the one-way communication latency between any pair of cluster nodes on the VLAN must be less than 500 milliseconds.
- In Windows Server 2003 operating systems, cluster nodes exchange multicast heartbeats rather than unicast heartbeats. A heartbeat is a message that is sent regularly between cluster network drivers on each node. Heartbeat messages are used to detect communication failure between cluster nodes. Using multicast technology enables better node communication because it allows several unicast messages to be replaced with a single multicast message. Clusters that consist of fewer than three nodes will not send multicast heartbeats. For additional information about using multicast technology, see article 307962, "Multicast Support Enabled for the Cluster Heartbeat," in the Microsoft Knowledge Base.

Determine an appropriate name for each network connection. For example, you might want to name the private network "Private" and the public network "Public." This will help you uniquely identify a network and correctly assign its role.

The following figure shows the elements of a four-node cluster that uses a private network.




**Setting the order of the network adapter binding**

One of the recommended steps for setting up networks is to ensure the network adapter binding is set in the correct order. To do this, use the following procedure.

**To set the order of the network adapter binding**

1. To open **Network Connections**, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
2. On the **Advanced** menu, click **Advanced Settings**.
3. In **Connections**, click the connection that you want to modify.

4. Set the order of the network adapter binding as follows:
    - External public network
    - Internal private network (Heartbeat)
    - [Remote Access Connections]
5. Repeat this procedure for all nodes in the cluster.

**Configuring the private network adapter**
As stated earlier, the recommended configuration for two adapters is to use one adapter for private communication, and the other adapter for mixed communication. To configure the private network adapter, use the following procedure.

**To configure the private network adapter**

1. To open **Network Connections**, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
2. Right-click the connection for the adapter you want to configure, and then click **Properties**. **Local Area Properties** opens and looks similar to the following figure:




3. On the **General** tab, verify that the **Internet Protocol (TCP/IP)** check box is selected, and that all other check boxes in the list are clear.
4. If you have network adapters that can transmit at multiple speeds and that allow you to specify the speed and duplex mode, manually configure the Duplex Mode, Link Speed, and Flow Control settings for the adapters to the same values and settings on all nodes. If the network adapters you are using do not support manual settings, contact your adapter manufacturer for specific information about appropriate speed and duplex settings for your network adapters. The amount of information that is traveling across the heartbeat network is small, but latency is critical for communication. Therefore, if you have the same speed and duplex settings, this helps ensure that you have reliable communication. If the adapters are connected to a switch, ensure that the port settings of the switch match those of the adapters. If you do not know the supported speed of your card and connecting devices, or if you run into compatibility problems, you should set all devices on that path to 10 megabytes per second (Mbps) and Half Duplex.

    Teaming network adapters on all cluster networks concurrently is not supported because of delays that can occur when heartbeat packets are transmitted and received between cluster nodes. For best results, when you want redundancy for the private interconnect, you should disable teaming and use the available ports to form a second private interconnect. This achieves the same end result and provides the nodes with dual, robust communication paths.

    You can use Device Manager to change the network adapter settings. To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.

Right-click the network adapter you want to change, and then click **Properties**. Click **Advanced** to manually change the speed and duplex mode for the adapter. The page that opens looks similar to the following figure:




5. On the **General** tab in **Network Connections**, select **Internet Protocol (TCP/IP)**, and click **Properties**.

   Internet Protocol (TCP/IP) Properties opens and looks similar to the following figure:




6. On the **General** tab, verify you have selected a static IP address that is not on the same subnet or network as any other public network adapter. You should put the private network adapter in one of the following private network ranges:

   - 10.0.0.0 through 10.255.255.255 (Class A)
   - 172.16.0.0 through 172.31.255.255 (Class B)
   - 192.168.0.0 through 192.168.255.255 (Class C)

7. On the **General** tab, verify that no values are defined in **Default Gateway** under **Use the following IP address**, and no values are defined under **Use the Following DNS server addresses**. After you have done so, click **Advanced**.

8. On the **DNS** tab, verify that no values are defined on the page and that the check boxes for **Register this connection's addresses in DNS** and **Use this connection's DNS suffix in DNS registration** are clear.

9. On the **WINS** tab, verify that no values are defined on the page, and then click **Disable NetBIOS over TCP/IP.**

   **Advanced TCP/IP Settings opens and looks similar to the following figure:**




10. After you have verified the information, click **OK**. You might receive the message "This connection has an empty primary WINS address. Do you want to continue?" To continue, click **Yes**.

11. Repeat this procedure for all additional nodes in the cluster. For each private network adapter, use a different static IP address.

**Configuring the public network adapter**

If DHCP is used to obtain IP addresses, it might not be possible to access cluster nodes if the DHCP server is inaccessible. For increased availability, static, valid IP addresses are required for all interfaces on a server cluster. If you plan to put multiple network adapters in each logical subnet, keep in mind that the Cluster service will recognize only one network interface per subnet.

- **Verifying connectivity and name resolution**. To verify that the private and public networks are communicating properly, "ping" all IP addresses from each node. To "ping" an IP address means that you search for and verify it. You should be able to ping all IP addresses, both locally and on the remote nodes. To verify the name resolution, ping each node from a client using the node's computer name instead of its IP address. It should only return the IP address for the public network. You might also want to try using the `PING -a` command to perform a reverse name resolution on the IP addresses.

- **Verifying domain membership**. All nodes in the cluster must be members of the same domain, and they must be able to access a domain controller and a DNS server. They can be configured as member servers or domain controllers. You should have at least one domain controller on the same network segment as the cluster. To avoid having a single point of failure, another domain controller should also be available. In this guide, all nodes are configured as member servers, which is the recommended role.

  In a two-node server cluster, if one node is a domain controller, the other node must also be a domain controller. In a four-node cluster, it is not necessary to configure all four nodes as domain controllers. However, when following a "best practices" model of having at least one backup domain controller, at least one of the remaining three nodes should also be configured as a domain controller. A cluster node must be promoted to a domain controller before the Cluster service is configured.

  The dependence in Windows Server 2003 on DNS requires that every node that is a domain controller must also be a

DNS server if another DNS server that supports dynamic updates is not available.

You should consider the following issues if you are planning to deploy cluster nodes as domain controllers:

- If one cluster node in a two-node cluster is a domain controller, the other node must also be a domain controller.
- There are performance implications associated with the overhead of running a computer as a domain controller. There is increased memory usage and additional network traffic from replication because these domain controllers must replicate with other domain controllers in the domain and across domains.
- If the cluster nodes are the only domain controllers, they each must be DNS servers as well. They should point to themselves for primary DNS resolution and to each other for secondary DNS resolution.
- The first domain controller in the forest or domain will assume all Operations Master Roles. You can redistribute these roles to any node. However, if a node fails, the Operations Master Roles assumed by that node will be unavailable. Because of this, you should not run Operations Master Roles on any cluster node. This includes Schema Master, Domain Naming Master, Relative ID Master, PDC Emulator, and Infrastructure Master. These functions cannot be clustered for high availability with failover.
- Because of resource constraints, it might not be optimal to cluster other applications such as Microsoft SQL Server™ in a scenario where the nodes are also domain controllers. This configuration should be thoroughly tested in a lab environment before deployment.

Because of the complexity and overhead involved when cluster nodes are domain controllers, all nodes should be member servers.

- **Setting up a Cluster service user account**. The Cluster service requires a domain user account that is a member of the Local Administrators group on each node. This is the account under which the Cluster service can run. Because Setup requires a user name and password, you must create this user account before you configure the Cluster service. This user account should be dedicated to running only the Cluster service and should not belong to an individual.

| **Note** |
| --- |
| It is not necessary for the Cluster service account (CSA) to be a member of the Domain Administrators group. For security reasons, domain administrator rights should not be granted to the Cluster service account. |

The Cluster service account requires the following rights to function properly on all nodes in the cluster. The Cluster Configuration Wizard grants the following rights automatically:

- Act as part of the operating system
- Adjust memory quotas for a process
- Back up files and directories
- Restore files and directories
- Increase scheduling priority
- Log on as a service

You should ensure that the Local Administrator Group has access to the following user rights:

- Debug programs
- Impersonate a client after authentication

- Manage auditing and security log

You can use the following procedure to set up a Cluster service user account.

**To set up a Cluster service user account**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the folder to which you want to add a user account.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder*

3. Point to **New**, and then click **User.**
4. **New Object - User** opens and looks similar to the following figure:




5. Type a first name and last name (these should make sense but are usually not important for this account)
6. In **User logon name**, type a name that is easy to remember, such as ClusterService1, click the UPN suffix in the drop-down list, and then click **Next**.
7. **In Password** and **Confirm password**, type a password that follows your organization's guidelines for passwords, and then select **User Cannot Change Password** and **Password Never Expires**. Click **Finish** to create the account.

   If your administrative security policy does not allow the use of passwords that never expire, you must renew the password and update the Cluster service configuration on each node before the passwords expire.

8. In the console tree of the Active Directory Users and Computers snap-in, right-click **Cluster**, and then click **Properties**.
9. Click **Add Members to a Group**.
10. Click **Administrators**, and then click **OK**. This gives the new user account administrative permissions on the computer.

**Setting up disks**

This section includes information and step-by-step procedures you can use to set up disks.

| **Important** |
| --- |
| To avoid possible corruption of cluster disks, ensure that both the Windows Server 2003 operating system and the Cluster service are installed, configured, and running on at least one node before you start the operating system on another node in the cluster. |

**Quorum resource**

The quorum resource maintains the configuration data necessary for recovery of the cluster. The quorum resource is generally accessible to other cluster resources so that any cluster node has access to the most recent database changes. There can only be one quorum disk resource per cluster.

The requirements and guidelines for the quorum disk are as follows:

- The quorum disk should be at least 500 MB in size.
- You should use a separate LUN as the dedicated quorum resource.
- A disk failure could cause the entire cluster to fail. Because of this, we strongly recommend that you implement a hardware RAID solution for your quorum disk to help guard against disk failure. Do not use the quorum disk for anything other than cluster management.

When you configure a cluster disk, it is best to manually assign drive letters to the disks on the shared bus. The drive letters should not start with the next available letter. Instead, leave several free drive letters between the local disks and the shared disks. For example, start with drive Q as the quorum disk and then use drives R and S for the shared disks. Another method is to start with drive Z as the quorum disk and then work backward through the alphabet with drives X and Y as data disks. You might also want to consider labeling the drives in case the drive letters are lost. Using labels makes it easier to determine what the drive letter was. For example, a drive label of "DriveR" makes it easy to determine that this drive was drive letter R. We recommend that you follow these best practices when assigning driver letters because of the following issues:

- Adding disks to the local nodes can cause the drive letters of the cluster disks to be revised up by one letter.
- Adding disks to the local nodes can cause a discontinuous flow in the drive lettering and result in confusion.
- Mapping a network drive can conflict with the drive letters on the cluster disks.

The letter Q is commonly used as a standard for the quorum disk. Q is used in the next procedure.

The first step in setting up disks for a cluster is to configure the cluster disks you plan to use. To do this, use the following procedure.

**To configure cluster disks**

1. Make sure that only one node in the cluster is turned on.
2. Open Computer Management (Local).
3. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
4. When you first start Disk Management after installing a new disk, a wizard appears that provides a list of the new disks detected by the operating system. If a new disk is detected, the Write Signature and Upgrade Wizard starts. Follow the instructions in the wizard.
5. Because the wizard automatically configures the disk as dynamic storage, you must reconfigure the disk to basic storage. To do this, right-click the disk, and then click **Convert To Basic Disk**.
6. Right-click an unallocated region of a basic disk, and then click **New Partition**.
7. In the New Partition Wizard, click **Next**, click **Primary partition**, and then click **Next**.
8. By default, the maximum size for the partition is selected. Using multiple logical drives is better than using multiple partitions on one disk because cluster disks are managed at the LUN level, and logical drives are the smallest unit of failover.
9. Change the default drive letter to one that is deeper into the alphabet. For example, start with drive Q as the quorum disk, and then use drives R and S for the data disks.
10. Format the partition with the NTFS file system.
11. In **Volume Label**, enter a name for the disk; for example, "Drive Q." Assigning a drive label for cluster disks reduces the time it takes to troubleshoot a disk recovery scenario.

After you have finished entering values for the new partition, it should look similar to the following figure:




**Important**

Ensure that all disks are formatted as MBR; GPT disks are not supported as cluster disks.

After you have configured the cluster disks, you should verify that the disks are accessible. To do this, use the following procedure.

**To verify that the cluster disks are accessible**

1. Open Windows Explorer.
2. Right-click one of the cluster disks, such as "Drive Q," click **New**, and then click **Text Document**.
3. Verify that the text document was created and written to the specified disk, and then delete the document from the cluster disk.
4. Repeat steps 1 through 3 for all cluster disks to verify that they are all accessible from the first node.
5. Turn off the first node, and then turn on the second node.
6. Repeat steps 1 through 3 to verify that the disks are all accessible from the second node.
7. Repeat again for any additional nodes in the cluster.
8. When finished, turn off the nodes and then turn on the first node again.

**Creating a new server cluster**

In the first phase of creating a new server cluster, you must provide all initial cluster configuration information. To do this, use the New Server Cluster Wizard.

**Important**

Before configuring the first node of the cluster, make sure that all other nodes are turned off. Also make sure that all cluster storage devices are turned on.

The following procedure explains how to use the New Server Cluster Wizard to configure the first cluster node.

**To configure the first node**

1. Open Cluster Administrator. To do this, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
2. In the **Open Connection to Cluster dialog box**, in **Action**, select **Create new cluster, and then click OK**.

3. The New Server Cluster Wizard appears. Verify that you have the necessary information to continue with the configuration, and then click **Next** to continue.

4. In **Domain**, select the name of the domain in which the cluster will be created. In **Cluster name**, enter a unique NetBIOS name. It is best to follow the DNS namespace rules when entering the cluster name. For more information, see article 254680, "DNS Namespace Planning," in the Microsoft Knowledge Base.




5. On the **Domain Access Denied** page, if you are logged on locally with an account that is not a domain account with local administrative permissions, the wizard will prompt you to specify an account. This is not the account the Cluster service will use to start the cluster.

| **Note** |
| --- |
| If you have the appropriate credentials, the **Domain Access Denied** screen will not appear. |




6. Since it is possible to configure clusters remotely, you must verify or type the name of the computer you are using as the first node. On the **Select Computer** page, verify or type the name of the computer you plan to use.

| **Note** |
| --- |
| The wizard verifies that all nodes can see the cluster disks. In some complicated SANs, the target IDs for the disks might not match on all the cluster nodes. If this occurs, the Setup program might incorrectly determine that the disk configuration is not valid. To address this issue, click **Advanced**, and then click **Advanced (minimum) configuration**. |




7. On the **Analyzing Configuration** page, Setup analyzes the node for possible hardware or software issues that can cause installation problems. Review any warnings or error messages that appear. Click **Details** to obtain more information about each warning or error message.




8. On the **IP Address** page, type the unique, valid, cluster IP address, and then click **Next**. The wizard automatically associates the cluster IP address with one of the public networks by using the subnet mask to select the correct network. The cluster IP address should be used for administrative purposes only, and not for client connections.




9. On the **Cluster Service Account** page, type the user name and password of the Cluster service account that was created during pre-installation. In **Domain**, select the domain name, and then click **Next**. The wizard verifies the user account and password.




10. On the **Proposed Cluster Configuration** page, review the information for accuracy. You can use the summary information to reconfigure the cluster if a system recovery occurs. You should keep a hard copy of this summary information with the change management log at the server. To continue, click **Next**.

| **Note** |
| --- |
| If you want, you can click **Quorum** to change the quorum disk designation from the default disk resource. To make this change, in the **Quorum resource** box, click a different disk resource. If the disk has more than one partition, click the partition where you want the cluster-specific data to be kept, and then click **OK**. |




11. On the **Creating the Cluster** page, review any warnings or error messages that appear while the cluster is being created. Click to expand each warning or error message for more information. To continue, click **Next**.




12. Click **Finish** to complete the cluster configuration.

| **Note** |
| --- |
| To view a detailed summary, click **View Log**, or view the text file stored at the following location: |

%SystemRoot%\System32\LogFiles\Cluster\ClCfgSrv.Log




**Validating the cluster installation**

You should validate the cluster configuration of the first node before configuring the second node. To do this, use the following procedure.

### To validate the cluster configuration

1. Open Cluster Administrator. To do this, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

2. Verify that all cluster resources are successfully up and running. Under **State**, all resources should be "Online."

**Configuring subsequent nodes**

After you install the Cluster service on the first node, it takes less time to install it on subsequent nodes. This is because the Setup program uses the network configuration settings configured on the first node as a basis for configuring the network settings on subsequent nodes. You can also install the Cluster service on multiple nodes at the same time and choose to install it from a remote location.

| Note |
| --- |
| The first node and all cluster disks must be turned on. You can then turn on all other nodes. At this stage, the Cluster service controls access to the cluster disks, which helps prevent disk corruption. You should also verify that all cluster disks have had resources automatically created for them. If they have not, manually create them before adding any more nodes to the cluster. |

After you have configured the first node, you can use the following procedure to configure subsequent nodes.

**To configure the second node**

1. Open Cluster Administrator. To do this, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
2. In the **Open Connection to Cluster** dialog box, in **Action**, select **Add nodes to cluster**. Then, in **Cluster or server name**, type the name of an existing cluster, select a name from the drop-down list box, or click **Browse** to search for an available cluster, and then click **OK** to continue.
3. When the Add Nodes Wizard appears, click **Next** to continue.
4. If you are not logged on with the required credentials, you will be asked to specify a domain account that has administrator rights and permissions on all nodes in the cluster.
5. In the **Domain** list, click the domain where the server cluster is located, make sure that the server cluster name appears in the **Cluster name** box, and then click **Next**.
6. In the **Computer name** box, type the name of the node that you want to add to the cluster. For example, to add Node2, you would type Node2.




7. Click **Add**, and then click **Next**.
8. When the Add Nodes Wizard has analyzed the cluster configuration successfully, click **Next**.
9. On the **Cluster Service Account** page, in **Password**, type the password for the Cluster service account. Ensure that the correct domain for this account appears in the **Domain** list, and then click **Next**.
10. On the **Proposed Cluster Configuration** page, view the configuration details to verify that the server cluster IP address, the networking information, and the managed disk information are correct, and then click **Next**.
11. When the cluster is configured successfully, click **Next**, and then click **Finish**.

**Configuring the server cluster after installation**
**Heartbeat configuration**
After the network and the Cluster service have been configured on each node, you should determine the network's function within the cluster. Using Cluster Administrator, select the **Enable this network for cluster use** check box and select from among the following options.

| Option | Description |
| --- | --- |
| Client access only (public network) | Select this option if you want the Cluster service to use this network adapter only for external communication with other clients. No node-to-node communication will take place on this network adapter. |
| Internal cluster communications only (private network) | Select this option if you want the Cluster service to use this network only for node-to-node communication. |
| All communications (mixed network) | Select this option if you want the Cluster service to use the network adapter for node-to-node communication and for communication with external clients. This option is selected by default for all networks. |

This guide assumes that only two networks are in use. It explains how to configure these networks as one mixed network and one private network. This is the most common configuration.

Use the following procedure to configure the heartbeat.

**To configure the heartbeat**

1. Open Cluster Administrator. To do this, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
2. In the console tree, double-click **Cluster Configuration**, and then click **Networks**.
3. In the details pane, right-click the private network you want to enable, and then click **Properties**. **Private Properties** opens and looks similar to the following figure:




4. Select the **Enable this network for cluster use** check box.
5. Click **Internal cluster communications only (private network)**, and then click **OK**.
6. In the details pane, right-click the public network you want to enable, and then click **Properties**. **Public Properties** opens and looks similar to the following figure:




7. Select the **Enable this network for cluster use** check box.

8. Click **All communications (mixed network)**, and then click **OK**.

**Prioritize the order of the heartbeat adapter**

After you have decided the roles in which the Cluster service will use the network adapters, you must prioritize the order in which the adapters will be used for internal cluster communication. To do this, use the following procedure.

**To configure network priority**

1. Open Cluster Administrator. To do this, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

2. In the console tree, click the cluster you want.

3. On the **File** menu, click **Properties**.




4. Click the **Network Priority** tab.

5. In **Networks used for internal cluster communications**, click a network.

6. To increase the network priority, click **Move Up**; to lower the network priority click **Move Down.**

7. When you are finished, click **OK.**

| **Note** |
| --- |
| If multiple networks are configured as private or mixed, you can specify which one to use for internal node communication. It is usually best for private networks to have higher priority than mixed networks. |

**Quorum disk configuration**

The New Server Cluster Wizard and the Add Nodes Wizard automatically select the drive used for the quorum device. The wizard automatically uses the smallest partition it finds that is larger then 50 MB. If you want to, you can change the automatically selected drive to a dedicated one that you have designated for use as the quorum. The following procedure explains what to do if you want to use a different disk for the quorum resource.

**To use a different disk for the quorum resource**

1. Open Cluster Administrator. To do this, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

2. If one does not already exist, create a physical disk or other storage-class resource for the new disk.

3. In the console tree, click the cluster name.

4. On the **File** menu, click **Properties, and then click the Quorum tab**. The quorum property page opens and looks similar to the following figure:




5. On the **Quorum** tab, click **Quorum resource**, and then select the new disk or storage-class resource that you want to use as the quorum resource for the cluster.

6. In **Partition**, if the disk has more than one partition, click the partition where you want the cluster specific data kept.

7. In **Root path**, type the path to the folder on the partition; for example:

   **\MSCS**

**Testing the Server Cluster**

After Setup, there are several methods you can use to verify a cluster installation.

- **Use Cluster Administrator**. After Setup is run on the first node, open Cluster Administrator, and then try to connect to the cluster. If Setup was run on a second node, start Cluster Administrator on either the first or second node, attempt to connect to the cluster, and then verify that the second node is listed.

- **Services snap-in**. Use the Services snap-in to verify that the Cluster service is listed and started.

- **Event log**. Use Event Viewer to check for **ClusSvc** entries in the system log. You should see entries that confirm the Cluster service successfully formed or joined a cluster.

#### Testing whether group resources can fail over

You might want to ensure that a new group is functioning correctly. To do this, use the following procedure.

**To test whether group resources can fail over**

1. Open Cluster Administrator. To do this, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

2. In the console tree, double-click the **Groups** folder.

3. In the console tree, click a group.

4. On the **File** menu, click **Move Group**. On a multi-node cluster server, when using **Move Group**, select the node to move the group to. Make sure the **Owner** column in the details pane reflects a change of owner for all of the group's dependencies.

5. If the group resources successfully fail over, the group will be brought online on the second node after a short period of time.

#### SCSI Drive Installations

This section of the guide provides a generic set of instructions for parallel SCSI drive installations.

| **Important** |
| --- |
| If the SCSI hard disk vendor’s instructions differ from the instructions provided here, follow the instructions supplied by the vendor. |

The SCSI bus listed in the hardware requirements must be configured before you install the Cluster service. This configuration applies to the following:

- The SCSI devices.

- The SCSI controllers and the hard disks. This is to ensure that they work properly on a shared SCSI bus.

- The termination of the shared bus. If a shared bus must be terminated, it must be done properly. The shared SCSI bus must have a terminator at each end of the bus. It is possible to have multiple shared SCSI buses between the nodes of a cluster.

In addition to the following information, refer to documentation from the manufacturer of your SCSI device.

#### Configuring SCSI devices

Each device on the shared SCSI bus must have a unique SCSI identification number. Because most SCSI controllers default to SCSI ID 7, configuring the shared SCSI bus includes changing the SCSI ID number on one controller to a different number, such as SCSI ID 6. If there is more than one disk that will be on the shared SCSI bus, each disk must have a unique SCSI ID number.

#### Storage Area Network Considerations

Fibre Channel systems are required for all server clusters running 64-bit versions of Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. It is also best to use Fibre Channel systems for clusters of three or more nodes. Two methods of Fibre Channel-based storage are supported in a cluster that is running Windows Server 2003: arbitrated loops and switched fabric.

| **Note** |
| --- |
| To determine which type of Fibre Channel hardware to use, read the Fibre Channel vendor's documentation. |

**Fibre Channel arbitrated loops (FC-AL)**
A Fibre Channel arbitrated loop (FC-AL) is a set of nodes and devices connected into a single loop. FC-AL provides a cost-effective way to connect up to 126 devices into a single network.

Fibre Channel arbitrated loops provide a solution for a small number of devices in a relatively fixed configuration. All devices on the loop share the media, and any packet traveling from one device to another must pass through all intermediate devices. FC-AL is a good choice if a low number of cluster nodes is sufficient to meet your high-availability requirements.

FC-AL offers the following advantages:

- The cost is relatively low.
- Loops can be expanded to add storage; however, nodes cannot be added.
- Loops are easy for Fibre Channel vendors to develop.

The disadvantage of FC-ALs is that they can be difficult to deploy successfully. This is because every device on the loop shares the media, which causes the overall bandwidth of the cluster to be lower. Some organizations might also not want to be restricted by the 126-device limit. Having more than one cluster on the same arbitrated loop is not supported.

**Fiber Channel switched fabric (FC-SW)**
With Fibre Channel switched fabric, switching hardware can link multiple nodes together into a matrix of Fibre Channel nodes. A switched fabric is responsible for device interconnection and switching. When a node is connected to a Fibre Channel switching fabric, it is responsible for managing only the single point-to-point connection between itself and the fabric. The fabric handles physical interconnections to other nodes, transporting messages, flow control, and error detection and correction. Switched fabrics also offer very fast switching latency.

The switching fabric can be configured to allow multiple paths between the same two ports. It provides efficient sharing (at the cost of higher contention) of the available bandwidth. It also makes effective use of the burst nature of communications with high-speed peripheral devices.




Other advantages to using switched fabric include the following:

- It is easy to deploy.
- It can support millions of devices.
- The switches provide fault isolation and rerouting.
- There is no shared media, which allows faster communication in the cluster.

**Zoning vs. LUN masking**
Zoning and LUN masking are important to SAN deployments, especially if you are deploying a SAN with a server cluster that is running Windows Server 2003.

**Zoning**
Many devices and nodes can be attached to a SAN. With data stored in a single storage entity (known as a "cloud") it is important to control which hosts have access to specific devices. Zoning allows administrators to partition devices in logical volumes, thereby reserving the devices in a volume for a server cluster. This means that all interactions between cluster nodes and devices in the logical storage volumes are isolated within the boundaries of the zone; other non-cluster members of the SAN are not affected by cluster activity. The elements used in zoning are shown in the following figure:




You must implement zoning at the hardware level with the controller or switch, not through software. This is because zoning is a security mechanism for a SAN-based cluster. Unauthorized servers cannot access devices inside the zone. Access control is implemented by the switches in the fabric, so a host adapter cannot gain access to a device for which it has not been configured. With software zoning, the cluster would not be secure if the software component failed.

In addition to providing cluster security, zoning also limits the traffic flow within a given SAN environment. Traffic between ports is routed only to segments of the fabric that are in the same zone.

**LUN masking**

A logical unit number (LUN) is a logical disk defined within a SAN. Server clusters see LUNs and act as though the LUNs are physical disks. With LUN masking, which is performed at the controller level, you can define relationships between LUNs and cluster nodes. Storage controllers usually provide the means for creating LUN-level access controls that allow one or more hosts to access a given LUN. With access control at the storage controller, the controller itself can enforce access policies to the devices.

LUN masking provides security at a more detailed level than zoning. This is because LUNs allow for zoning at the port level. For example, many SAN switches allow overlapping zones, which enable a storage controller to reside in multiple zones. Multiple clusters in multiple zones can share the data on those controllers. The elements used in LUN masking are shown in the following figure:




**Other Resources**

For a comprehensive list of hardware and software supported by Windows operating systems, see one of the following:

- Windows Catalog at the Microsoft Web site.
- Hardware and software compatibility information in the Windows Server Catalog at the Microsoft Web site

For the latest information about Windows Server 2003, see the Windows 2003 Server Web site at the Microsoft Web site.

For interactive help in solving a problem with your computer or to research your problem see Product Support Services at the Microsoft Web site.

For additional information about cluster deployment, see "Designing and Deploying Clusters" at the Microsoft Web site.

For information about troubleshooting, see "Troubleshooting cluster node installations" at the Microsoft Web site.

For information about quorum configuration, see "Quorum Drive Configuration Information" at the Microsoft Web site.

For information about private heartbeat configuration, see "Recommended private 'Heartbeat' configuration on a cluster server" at the Microsoft Web site.

For information about network failure, see "Network Failure Detection and Recovery in a Server Cluster" at the Microsoft Web site.

For information about quorum disk designation, see "How to Change Quorum Disk Designation" at the Microsoft Web site.

For additional information about Storage Area Networks, see "Microsoft Windows Clustering: Storage Area Networks" at the Microsoft Web site.

For information about geographically dispersed clusters, see "Geographically Dispersed Clusters in Windows Server 2003" at the Microsoft Web site.

# What's New in Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Microsoft® Windows Server™ 2003 R2, a software update to the Windows Server 2003 operating system, makes it easier and more cost effective to extend connectivity and control to identities, locations, data, and applications throughout and beyond your organization.

Built on Windows Server 2003, Windows Server 2003 R2 takes advantage of the stability and security of the code base while extending connectivity and control into new areas. With Windows Server 2003 R2, you get the following improvements:

- Identity and access management
- Branch office server management
- Storage setup and management
- Application development inside and outside your organization's traditional boundaries

For more information about Windows Server 2003 R2, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45560).

This topic describes the following new components that you can install with Windows Server 2003 R2:

**Server Manageability**

- Administration Tools Packs
- Hardware Management
- MMC 3.0

**Features for Active Directory**

- Active Directory Application Mode (ADAM)
- Active Directory Federation Services

**Disk and File Management Features**

- Branch Office
    - Distributed File System (DFS)
- Common Log File System (CLFS)
- File Server Management
- Microsoft Services for Network File System
- Storage Management for SANs
- File Server Resource Manager

**Printer and Protocol Support**

- Print Management

**Microsoft .NET Framework**

- Microsoft .NET Framework 2.0

**Internet and E-Mail Services and Features**

- Windows Sharepoint Services

**UNIX Interoperability**

- Identity Management for UNIX
- Subsystem for UNIX-based Applications

**Server Manageability**
To help you manage your server, Windows Server 2003 R2 provides Administration Tools Packs and the following improvements to hardware management and Microsoft Management Console (MMC) 3.0.

**Administration Tools Packs**
To make the remote management of your servers easier, you can install the Administration Tools Packs for Windows Server 2003 R2 located on Disc 2 in the \mgmttls directory. Each of the following components has an Administration Tools Pack:

- Microsoft Management Console (MMC) 3.0
- Print Management
- File Server Resource Manager
- DFS Replication
- Identity Management for UNIX
- File Server Management

In order to administer components, you must first install the MMC 3.0 Administration Tools Pack. The Administration Tools Packs can be installed on computers running Windows® XP Professional with Service Pack 2 (SP2). For specific installation instructions, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=54293).

**Hardware Management**
You can use Windows Remote Management (WinRM) to manage server hardware remotely across firewalls and monitor conditions on servers that are offline.

This new feature in Windows Server 2003 R2 is implemented using WS-Management, an industry-standard, Web services-based remoting protocol. WS-Management uses HTTPS as the secure transport that enables firewall traversal.

WinRM provides a command-line interface for common management tasks and a scripting application programming interface (API) for writing Windows Script Host-based system administration scripts.

Using WinRM, you can write scripts to monitor and control the state of server hardware by communicating with a Baseboard Management Controller (BMC). A BMC is a separate micro-controller with its own network adapter that is connected to the processor board of a server and can monitor conditions even when the server is off or malfunctioning. A new Windows Management Interface (WMI) provider for Intelligent Platform Management Interface (IPMI) exposes six new WMI classes for accessing BMC information with scripts.

WinRM is not installed by default with Windows Server 2003 R2. To install it, open Add or Remove Programs from Control Panel, and click **Add/Remove Windows Components**.

For more information about hardware management, see "Hardware Management in Microsoft Windows Server 2003 R2" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45204).

Back to Top

**MMC 3.0**

Microsoft Management Console (MMC) 3.0 supports richer functionality in snap-ins designed for the MMC 3.0 infrastructure. In addition, the following improvements apply to all MMC 3.0 consoles:

- **Action pane.** The Action pane appears on the right side of the snap-in console. It lists the actions that are currently available to you, based on the currently selected items in the tree or the results pane.

  To show or hide the action pane, click the **Show/Hide Action Pane** button on the toolbar, which is similar to the **Show/Hide Console Tree** button.

- **New Add or Remove Snap-in dialog box.** The new **Add or Remove Snap-in** dialog box makes it easy to add, organize, and remove snap-ins. You can control which extensions are available, and whether to automatically enable snap-ins that may be installed later. You can nest snap-ins and rearrange the snap-ins in the tree.

- **Improved error handling.** MMC 3.0 notifies you of errors in snap-ins that could cause MMC to fail and provides several options for responding to those errors.

MMC 3.0 is automatically installed when you install Windows Server 2003 R2.

Back to Top

**Features for Active Directory**

The following Active Directory components are included in Windows Server 2003 R2.

**Active Directory Application Mode (ADAM)**

Active Directory Application Mode (ADAM) is an independent mode of the Active Directory® directory service, without infrastructure features, that provides directory services for applications. It provides a data store and services for accessing the data store. It uses standard application programming interfaces (APIs) for accessing the application data. ADAM operates either as a standalone data store, or it operates with replication. Its independence enables local control and autonomy of directory services for specific applications. It also facilitates independent, flexible schemas and naming contexts.

Back to Top

**Active Directory Federation Services**

The fundamental purpose of Active Directory Federation Services (ADFS) is to leverage single user sign-on to authenticate the user to multiple related Web applications over the life of a single online session. ADFS does this by securely sharing digital identity and entitlement rights across security and enterprise boundaries.

Back to Top

**Disk and File Management Features**

Windows Server 2003 R2 includes the following disk and file management features.

**Branch Office**

Branch offices are the remote locations of an organization that connect to a headquarters or hub over a wide-area network (WAN). The technologies included in Windows Server 2003 R2 can simplify how you manage a variety of activities and tasks in branch offices. These activities and tasks include:

- Publishing files from centralized hubs to branch offices.

- Replicating files from branch locations for backup, fault-tolerance, or cross-publishing.

- Collaborating on documents that are shared among branches or between hubs and branches.

- Efficiently managing printers, printer drivers, and print queues in branch offices.

These improvements are supported by the enhanced print-management tools in Windows Server 2003 R2. They are also supported by Distributed File System (DFS) in Windows Server 2003 R2.

For more information about these technologies in Windows Server 2003 R2, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51679).

Back to Top

**Distributed File System (DFS)**
The Distributed File System (DFS) solution in Windows Server 2003 R2 provides simplified, fault-tolerant access to files and WAN-friendly replication. Distributed File System consists of two technologies:

- **DFS Namespaces.** Formerly known as Distributed File System, DFS Namespaces allows administrators to group shared folders located on different servers and present them to users as a virtual tree of folders known as a namespace. A namespace provides numerous benefits, including increased availability of data, load sharing, and simplified data migration.

- **DFS Replication.** The successor to File Replication service (FRS), DFS Replication is a new state-based, multimaster replication engine that supports scheduling and bandwidth throttling. DFS Replication uses a new compression algorithm known as Remote Differential Compression (RDC). RDC is a "diff over the wire" protocol that can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, re-arrangements of data in files, enabling DFS Replication to replicate only the deltas (changes) when files are updated.

For more information about Distributed File System in Windows Server 2003 R2, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51679).

Back to Top

**Common Log File System (CLFS)**
Common Log File System (CLFS) is a loadable driver that provides kernel- or user-mode applications with a robust logging subsystem. CLFS is a unique Windows technology that can be used to develop applications and middleware that depend on durably writing and reading sequential data. Examples include replication agents, auditing agents, databases, and other transactional resource managers.

Features include the ability to create log files with a single stream of data or with multiple streams of data for shared use by one or more clients; both circular and linear logging; guaranteed ability to flush buffered data by pre-reserving space in the log; policy-based log size and space management; sharing of a single log by both kernel and user clients; a notification mechanism to allow different users within the same log to coordinate their log use; flexible buffering of log data; archiving APIs do not interfere with normal operations; atomic multi-sector writes; and torn-write detection.

CLFS is optimized for performance. All writes to the log file are buffered until an explicit flush, an opportunity to share a write with another client, or the buffer is filled. Log data is written directly to the hard disk from the log buffers without copying. Multiple streams of data can be written during the same I/O operation, resulting in only one disk seek for what normally takes multiple seeks and writes. Reads are cached to save disk accesses during normal operation or bursts of read activity.

For more information, see the "Platform SDK" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=44489).

Back to Top

**File Server Management**
You can use File Server Management to perform many tasks, such as formatting volumes, creating shares, defragmenting volumes, creating and managing shares, setting quota limits, creating storage utilization reports, replicating data to and from the file server, managing Storage Area Networks (SANs), and sharing files with UNIX and Macintosh systems.

Back to Top

**File Server Resource Manager**
With the increasing demand on storage resources, as organizations rely more heavily on data than ever before, IT administrators face the challenge of overseeing a larger and more complex storage infrastructure, while at the same time, tracking the kind of information available in it.

File Server Resource Manager is a new Microsoft Management Console (MMC) snap-in that provides a suite of tools for administrators to understand, control, and manage the quantity and type of data stored on their servers. By using File Server Resource Manager, administrators can place quotas on folders and volumes, actively screen files, and generate comprehensive storage reports.

**File Server Resource Manager quotas vs. NTFS disk quotas**
The Windows 2000 and Windows Server 2003 operating systems support disk quotas, which are used to track and control disk

usage per user on NTFS volumes. The following table outlines the advantages of using the quota management tools in File Server Resource Manager.

| Quota features | File Server Resource Manager | NTFS disk quotas |
|---|---|---|
| Quota tracking | By folder or by volume | Per user on a volume |
| Disk usage calculation | Actual disk space | Logical file size |
| Notification mechanisms | E-mail, custom reports, command execution, event logs | Event logs only |

Back to Top

**Microsoft Services for Network File System**

Microsoft Services for Network File System (NFS) is a component of Windows Server 2003 R2 that provides a file-sharing solution for enterprises that have a mixed Windows and UNIX environment. Microsoft Services for NFS allows users to transfer files between Windows Server 2003 R2 and UNIX computers using the Network File System (NFS) protocol.

Microsoft Services for NFS is an update to the NFS components that were previously available in Services for UNIX 3.5. Microsoft Services for NFS includes the following new features:

- Support for 64-bit

  Microsoft Services for NFS components can be installed on all editions of Windows Server 2003 R2, including 64-bit versions.

- Updated Microsoft Management Console (MMC) Microsoft Services for NFS Administration snap-in

- Enhanced reliability

- Support for UNIX special devices (mknod)

Microsoft Services for NFS allows you to support a mixed environment of Windows- and UNIX-based operating systems. It also allows you to update your organization's computers while supporting older technology during the transition phase. The following scenarios are examples of how enterprises can benefit from deploying Microsoft Services for NFS.

- Allows UNIX clients to access resources on computers running Windows Server 2003 R2. Your company may have UNIX clients accessing resources, such as files, on UNIX file servers. To take advantage of new Windows Server 2003 features, such as Shadow Copies for Shared Folders, you can move resources from your UNIX servers to computers running Windows Server 2003 R2, and then set up Microsoft Services for NFS to enable access by UNIX clients that are running NFS software. All of your UNIX clients will be able to access the resources using the NFS protocol with no changes required.

- Allows computers running Windows Server 2003 R2 to access resources on UNIX file servers. Your company may have a mixed Windows and UNIX environment with resources, such as files, stored on UNIX file servers. You can use Microsoft Services for NFS to enable computers running Windows Server 2003 R2 to access these resources when the file servers are running NFS software.

Back to Top

**Storage Management for SANs**

Storage Manager for SANs is a new Microsoft Management Console (MMC) snap-in that helps you create and manage logical unit numbers (LUNs) on fibre channel and iSCSI disk drive subsystems in your storage area network (SAN). Storage Manager for SANs can be used on storage subsystems that support Virtual Disk Server (VDS).

Use Storage Manager for SANs to create and assign LUNs, manage connections between LUNs and the servers in your SAN, and set the security properties for iSCSI storage subsystems.

Back to Top

**Printer and Protocol Support**

Windows Server 2003 R2 provides an updated MMC snap-in for Print Management.

**Print Management**
Print Management is an updated Microsoft Management Console (MMC) snap-in that you can use to view and manage printers and print servers in your organization. You can use Print Management from any computer running Windows Server 2003 R2, and you can manage all network printers on print servers running Windows 2000 Server, Windows Server 2003, or Windows Server 2003 R2.

Print Management provides up-to-the-minute details such as the queue status, printer name, driver name, and server name. You can also set custom views by using the Print Management filtering capability. For example, you can create a view that only displays printers in a particular error state. You can also configure Print Management to send e-mail notifications or run scripts when a printer or print server needs attention. The filtering capability also allows you to bulk edit print jobs, such as canceling all print jobs at once. You can also delete multiple printers at the same time.

Administrators can install printers remotely by using the automatic detect feature, which finds and installs printers on the local subnet to the local print server. Administrators can log on to a local server at a branch location by using Remote Desktop and then easily install printers remotely.

When a printer model has a Web page feature, Print Management can display troubleshooting details, such as the location of the paper jam or printer toner level. Some printer models have Web pages that provide options for remote control functions that can help resolve problems at branch locations.

For more information, see the "Print Management Step-by-Step Guide" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50141).

Back to Top

**Microsoft .NET Framework**
New features for .NET Framework 2.0 are available in Windows Server 2003 R2.

**Microsoft .NET Framework 2.0**
The .NET Framework version 2.0 is the latest version of the .NET Framework, a Windows component for building, deploying and running applications. The .NET Framework 2.0 is focused on delivering developer productivity, operational excellence, and the ability to preserve and enhance existing software investments. Core enhancements include support for 64-bit processors for increased performance and scalability, transactional enhancements and optimizations, productivity improvements in ASP.NET, and support for WS-I BP 1.1 and SOAP 1.2 with Web Services.

For a comprehensive list of the new .NET Framework 2.0 features, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50322).

Back to Top

**Internet and E-Mail Services and Features**
Windows SharePoint® Services is included in Windows Server 2003 R2.

**Windows SharePoint Services**
Microsoft Windows SharePoint Services is an integrated portfolio of collaboration and communication services designed to connect people, information, processes, and systems both within and beyond the organizational firewall. Windows SharePoint Services Service Pack 2 is included in this release of Windows Server 2003. Now you can install Windows SharePoint Services directly from the Configure Your Server wizard or Manage Your Server. When you install Windows SharePoint Services, you get the following features:

**SharePoint sites – file storage plus collaboration**
Web sites based on Microsoft Windows SharePoint Services 2.0 provide a place where teams can communicate, share documents, and work together on a project. SharePoint sites include:

- Team collaboration features, such as event calendars, contacts, Web links, discussions, issues lists, announcements, and much more.
- Document libraries, which are places where users can store and retrieve documents while taking advantage of rich features such as check-in and check-out, version history, custom metadata, and flexible, customizable views.
- Web Parts, which can provide data access, Web services, and many other applications and content to SharePoint sites.

Site users can contribute to the site by using nothing more than a Web browser. However, if users have Windows SharePoint Services-compatible client programs, such as Microsoft Office 2003, installed on their computers they can work seamlessly with the site, saving files to libraries, editing documents in the client program, and moving or linking that information to the

site.

**SharePoint Central Administration – Web browser interface for managing your server**
You can manage a single server or an entire server farm running Windows SharePoint Services from a Web browser interface called SharePoint Central Administration. Use SharePoint Central Administration to extend a virtual server, create sites or turn on Self-Service Site Creation so users can create their own sites, manage security settings, manage the list of servers in your server farm, and so on. If you prefer, you can also use the Stsadm.exe command-line utility to manage your servers running Windows SharePoint Services.

**Applications for Windows SharePoint Services – download and install ready-to-use applications to get working quickly**
Boost worker productivity by downloading and installing an application for Windows SharePoint Services. These applications are tailored to address specific processes or tasks. There are many to choose from, including recruiting, project management, help desk issue tracking, timesheet and scheduling, event planning, and more. Use the applications as soon as you install them or customize them further to meet your organization's specific needs and requirements.

For more information, see "Applications for Windows SharePoint Services" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47987).

**Support for advanced extranet configurations**
A URL zone mapping feature (new in Windows SharePoint Services, Service Pack 2) allows for easier interactions with a proxy server or firewall. Windows SharePoint Services uses absolute URLs to generate some hyperlinks in its Web pages and e-mail messages. This prevented Windows SharePoint Services from supporting certain advanced extranet configurations (such as SSL termination, host header modification, and port translation) where a reverse proxy server is deployed in front of the Windows SharePoint Services server. Now you can set up URL zones, such as Internet, intranet, and extranet, and map incoming and outgoing URLs to those zones, so that these configurations can work correctly.

For more information about Windows SharePoint Services in Windows Server 2003 R2, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45560).

Back to Top

**UNIX Interoperability**
For improved interoperability with UNIX, the following features are available in Windows Server 2003 R2.

**Identity Management for UNIX**
Windows Services for UNIX delivers Windows and UNIX integration with updated identity management solutions. The following solutions help provide uninterrupted user access and efficient management of network resources across operating systems:

- **Server for NIS.** Helps integrate Windows and UNIX-based Network Information Service (NIS) servers by enabling an Active Directory domain controller to act as a master NIS server for one or more NIS domains.

- **Password Synchronization.** Helps integrate Windows and UNIX servers by simplifying the process of maintaining secure passwords. With Password Synchronization, users do not need to maintain separate passwords for their Windows and UNIX accounts or remember to change the password in multiple locations.

Back to Top

**Subsystem for UNIX-based Applications**
Subsystem for UNIX-based Applications (SUA) is a source-compatibility subsystem for compiling and running custom UNIX-based applications on a computer running a Windows server-class operating system. You can make your UNIX applications fully interoperable with Windows in SUA with little or no change to your original source code.

Subsystem for UNIX-based Applications provides an operating system for Portable Operating System Interface (POSIX) processes. SUA, along with its package of support utilities (such as shells and a telnet client), provides a complete UNIX environment. The package of support utilities includes a comprehensive set of scripting utilities and a software development kit (SDK) designed to fully support the development capabilities of SUA.

SUA also supports case-sensitive file names, job control, compilation tools, and the use of over 300 UNIX commands, utilities, and shell scripts. Because the subsystem installs separately from the Windows kernel, it offers true UNIX functionality without any emulation.

New features in this release include the following:

- **Database (OCI/ODBC) library connectivity.** SUA supports connectivity to Oracle and SQL Server from database

applications by using the Oracle Call Interface (OCI) and the Open Database Connectivity (ODBC) standard.

- **Microsoft® Visual Studio® Debugger Extension for debugging POSIX applications.** SUA includes support for debugging your POSIX processes using Visual Studio IDE.

- **Utilities based on SVR-5 and BSD UNIX environments.** The SUA download package supports two different UNIX environments: SVR-5 and BSD.

- **Support for 64-bit applications.** Using a process called thunking, SUA provides support not only for 64-bit applications running on a 64-bit operating system, but also default support for 32-bit binaries running on a 64-bit operating system.

To download the support utilities, see "Installing and Using Utilities and SDK for UNIX-based Applications" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=38959).

Back to Top

# Step-by-Step Guide for the Distributed File System Solution in Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**About This Guide**
Welcome to the Step-by-Step Guide for the Distributed File System (DFS) solution in Microsoft® Windows Server™ 2003 R2. This guide provides system requirements, installation instructions, and step-by-step walkthroughs for deploying the DFS Namespaces and DFS Replication features in the Distributed File System solution in Windows Server 2003 R2 in a lab environment.

If you are not familiar with DFS Namespaces and DFS Replication in Windows Server 2003 R2, we recommend that you read the document titled "Overview of the Distributed File System Solution in Windows Server 2003 R2" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=46801). The overview describes the benefits of DFS Replication and the improvements it offers over File Replication service (FRS). It also describes the DFS Namespaces enhancements exposed in Windows Server 2003 R2. These enhancements, introduced as updated application programming interfaces (APIs) in Windows Server 2003 with Service Pack 1 (SP1), provide easier management and more flexibility for namespaces used in branch offices.

**Lab Requirements**
Two step-by-step sections are provided in this guide. The requirements for completing each of the step-by-step sections are as follows:

**Lab Requirements for Completing the Step-by-Step Guide to Deploying a Namespace**
To complete all tasks in this section, you need a minimum of two servers configured in the test lab as follows:

- One server must run Windows Server 2003 R2. This is the server on which you will install the DFS Management snap-in to perform the tasks in this guide.
- The second server must run Windows Server 2003 SP1 or Windows Server 2003 R2.
- To create domain-based namespaces in these tasks, you must have Active Directory® deployed in the test lab. You must also be a member of the Domain Admins group or have been delegated the ability to create domain-based namespaces. For more information about delegation, see "Security requirements for creating and managing namespaces" in the DFS Management snap-in Help.
- To deploy DFS Replication in the namespace, you must have extended the schema to include the new DFS Replication objects in Active Directory. Refer to the section "Installing Windows Server 2003 R2 and Distributed File System" later in this guide for specific configuration requirements.

You can complete a subset of tasks if you have a single server or if you do not have Active Directory deployed in the test lab.

**Lab Requirements for Completing the Step-by-Step Guide to Deploying DFS Replication**
To complete all tasks in this section, you need to configure the test lab as follows:

- You need a minimum of three file servers. All three servers must have Windows Server 2003 R2 and the DFS Replication service installed. One of the servers must have the DFS Management snap-in installed. Follow the procedures in "Installing Windows Server 2003 R2 and Distributed File System" later in this guide to install the service and snap-in.
- The test lab must have Active Directory installed, and you must extend the schema using the instructions described in "Installing Windows Server 2003 R2 and Distributed File System" later in this guide.

Before beginning either step-by-step section, it is important to review the following requirements and limitations of DFS Replication.

- Servers in a replication group must be in the same forest. You cannot enable replication across servers in different forests.
- Replicated folders must be stored on NTFS volumes.
- On server clusters, replicated folders should be located in the local storage of a node, because the DFS Replication service is not cluster aware and the service will not fail over to another node.
- Antivirus software must be compatible with DFS Replication. Contact your antivirus software vendor to check for compatibility.
- DFS Replication might not work across firewalls because it uses the RPC dynamic endpoint mapper. Additionally, configuring DFS Replication using the DFS Management snap-in does not work when a firewall is enabled. To allow DFS Replication to work through a firewall, you can define a static port using the Dfsrdiag.exe command-line tool.

**Getting Started**

The following sections provide instructions on installing Windows Server 2003 R2 and the DFS Management snap-in.

**Installing Windows Server 2003 R2 and Distributed File System**

During Setup, follow the on-screen prompts to install Windows Server 2003 R2. Refer to the section "Lab Requirements" earlier in this guide for details about which servers must run Windows Server 2003 R2 and which servers can run Windows Server 2003 SP1.

If you plan to use DFS Replication, you must first update the Active Directory schema to install the Active Directory components of DFS Replication. To upgrade the schema, on the schema operations master, run **adprep.exe /forestprep**. The adprep.exe command-line tool is available in the Cmpnents\R2\Adprep folder on the second Windows Server 2003 R2 installation CD.

After Windows Server 2003 R2 is installed and the schema is updated as necessary, you can install the Distributed File System (DFS) components and open the DFS Management snap-in by using the following procedures.

| **Note** |
| --- |
| Two methods for installing Distributed File System are provided. The method using the **Manage Your Server** file server tool allows you to install Distributed File System as a part of the file server role. This method also installs other file server tools, such as File Server Resource Manager and File Server Management. The method using **Add or Remove Programs** allows you to install either the DFS Management snap-in, the DFS Replication Service, or both. |

Use the following procedure if the file server role has not been added.

**To install Distributed File System using Manage Your Server**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Manage Your Server**.
2. Click **Add or remove a role**.
3. In the **Configure Your Server Wizard**, click **Next**. On the **Server Role** page, click **File Server**, and then click **Next.** Click **Next** again.
4. In the **Add File Server Role Wizard**, click **Next**, and then click **Replicate data to and from this server**. Follow the steps in the wizard to complete the installation. You will be asked to restart the server when the installation is complete.

Use the following procedure if the file server role has already been installed and you see an option called **Update this role** in **Manage Your Server**.

**To update the file server role using Manage Your Server**

1. In **Manage Your Server**, click **Update this role**.
2. In the **File Server Role Wizard**, click **Next**.
3. On the **File Server Environment** page, click **Replicate data to and from this server**.
4. Follow the steps in the wizard to complete the installation. You will be asked to restart the server when the installation is complete.

**To install Distributed File System using Add or Remove Programs**

1. In **Control Panel**, double-click **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**, highlight **Distributed File System**, and then click **Details**.
3. Click all three options, and then click **OK**.

| **Note** |
| --- |
| Installing DFS Management also installs Microsoft .NET Framework 2.0, which is required to run the DFS Management snap-in. |

**To open the DFS Management snap-in**

- Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **DFS Management**.

**Overview of the DFS Management Snap-in**

The DFS Management snap-in is the graphical user interface tool for managing DFS Namespaces and DFS Replication. This snap-in is new and differs from the Distributed File System snap-in in Windows Server 2003. Therefore, before you begin using DFS Namespaces and DFS Replication, you might want to review the components of this snap-in, which are shown in the following figure and described in the sections that follow.




**Console Tree**

The console tree has two nodes, Namespaces and Replication, from which you can manage namespaces and DFS Replication.

**Namespaces node**

The following figure shows the elements under the Namespaces node in the console tree.




As the figure shows, the Namespaces node contains the namespaces you create as well as any existing namespaces you add to the console display. In the previous figure, one namespace is shown, **\\Contoso.com\Public**. Under each namespace is a hierarchical view of folders. Folders that have targets use a special icon to differentiate them from folders that do not have targets.

| **Note** |
| --- |
| If you are not familiar with namespace terminology, see the section "Introduction to Namespaces" later in this guide. |

**Replication node**

The following figure shows the elements under the Replication node in the console tree.




As the figure shows, the Replication node contains the replication groups you create as well as any existing replication groups that you add to the console display. A replication group represents a group of servers that participates in the replication of data. For more information about replication groups, see "Introduction to DFS Replication" later in this guide.

**Details Pane**

The contents of the details pane change according to what you have selected in the console tree. For example, if you select a namespace in the console tree, you see tabs named **Namespace**, **Namespace Servers**, and **Delegation** in the details pane. If you select a replication group, you see tabs named **Memberships**, **Connections**, **Replicated Folders**, and **Delegation**. You can double-click objects in the details pane to view their properties.

**Action Pane**

The Action pane shows two types of tasks: common tasks and tasks that apply to the selected object. If the Action pane is not visible, you can open it using the following steps: click the **View** menu, click **Customize**, and then click the **Action pane** option in the **Customize View** dialog box.

**Step-by-Step Guide to Deploying a Namespace**

The following sections introduce namespaces and walk you step-by-step through the process of deploying a namespace in a test lab. The tasks in these sections are designed for administrators who are new to DFS Namespaces as well as administrators who have experience using DFS in Windows Server 2003 and Windows® 2000 Server. These tasks walk you through deployment steps and point out aspects of the DFS Management snap-in that are new or significantly different from previous DFS management tools.

If you have not used DFS in Windows Server 2003 or Windows 2000 Server, we recommend that you read the introduction section that follows to

learn more about namespaces before you begin the tasks.

**Introduction to Namespaces**

DFS Namespaces allows you to group shared folders located on different servers by transparently connecting them to one or more namespaces. A *namespace* is a virtual view of shared folders in an organization. When you create a namespace, you select which shared folders to add to the namespace, design the hierarchy in which those folders appear, and determine the names that the shared folders show in the namespace. When a user views the namespace, the folders appear to reside on a single, high-capacity hard disk. Users can navigate the namespace without needing to know the server names or shared folders hosting the data.

The path to a namespace is similar to a Universal Naming Convention (UNC) path of a shared folder, such as \\Server1\Public\Software\Tools. If you are familiar with UNC paths, you know that in this example the shared folder, Public, and its subfolders, Software and Tools, are all hosted on Server1. Now, assume you want to give users a single place to locate data, but you want to host data on different servers for availability and performance purposes. To do this, you can deploy a namespace similar to the one shown in the following figure. The elements of this namespace are described after the figure.




- **Namespace server**. A namespace server hosts a namespace. The namespace server can be a member server or a domain controller.

- **Namespace root**. The root is the starting point of the namespace. In the previous figure, the name of the root is Public, and the namespace path is \\Contoso\Public. This type of namespace is known as a domain-based namespace, because it begins with a domain name (for example, Contoso) and its metadata is stored in Active Directory. Although a single namespace server is shown in the previous figure, a domain-based namespace can be hosted on multiple namespace servers.

- **Folder**. Folders help build the namespace hierarchy. Folders can optionally have folder targets. When users browse a folder with targets in the namespace, the client computer receives a referral that directs the client computer to one of the folder targets.

- **Folder targets**. A folder target is a UNC path of a shared folder or another namespace that is associated with a folder in a namespace. In the previous figure, the folder named **Tools** has two folder targets, one in London and one in New York, and the folder named **Training Guides** has a single folder target in New York. A user who browses to \\Contoso\Public\Software\Tools is transparently redirected to the shared folder \\LDN-SVR-01\Tools or \\NYC-SVR-01\Tools, depending on which site the user is in.

To view an online demo that describes the purpose and benefits of Distributed File System in Windows Server 2003, click the link "Distributed File System: Simplifying Complexity" on the Distributed File System and File Replication services page on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=39468).

**Deploying a Namespace**

The tasks in this section walk you through the process of deploying a namespace that looks similar to the namespace shown in the figure that appears in "Introduction to Namespaces" earlier in this guide.

**Task 1: Create a Namespace**

In this task, you create a new namespace using the DFS Management snap-in.

**To create a namespace**

1. In the console tree of the DFS Management snap-in, right-click the **Namespaces** node, and then click **New Namespace**.

2. Follow the steps in the **New Namespace Wizard** and supply the information described in the following table.

| New Namespace Wizard Page | What to Enter |
|---|---|
| **Namespace Server** | Enter the name of the server to host the namespace. The server can be a domain controller or a member server. |
| **Namespace Name and Settings** | In **Name**, type **Public**. |
| **Namespace Type** | If Active Directory is deployed in your test lab and you are a member of the Domain Admins group or have been delegated permission to create domain-based namespaces, choose **Domain-based namespace**. Otherwise, choose **Stand-alone namespace**.<br><br>To learn how a member of the Domain Admins group can delegate permission to create domain-based namespaces, see "Security requirements for creating and managing namespaces" in DFS Management Help. |

| | |
|---|---|
| **Review Settings and Create Namespace** | Click **Create** to create the namespace. |
| **Confirmation** | Click **Close** to close the wizard. |

When the wizard finishes, your new namespace will be added to the console tree. Double-click the **Namespaces** node, if necessary, to view your namespace, which should be similar to the following figure.

To browse the new namespace, type the following command in the Run dialog box, substituting either the server name (if you created a stand-alone namespace) or the domain name (if you created a domain-based namespace) as appropriate:

**** *server_or_domain* **\Public**

**Task 2: Add a Namespace Server**

If you created a domain-based namespace, perform this task to specify an additional server to host the namespace. Doing so increases the availability of the namespace and allows you to place namespace servers in the same sites as users. If you created a stand-alone namespace, you must skip this task because stand-alone namespaces only support a single namespace server.

**To add a namespace server**

1. In the console tree of the DFS Management snap-in, right-click *****domain***\Public**, and then click **Add Namespace Server**.
2. In **Namespace server**, type the name of another server to host the namespace, and then click **OK**.

After you finish this procedure, click the *****domain***\Public** namespace in the console tree and review the contents of the **Namespace Servers** tab in the details pane, which should look similar to the following figure. Notice that two UNC paths are listed. The site of each namespace server is also displayed.




**Task 3: Delegate Management Permissions**

You can delegate management permissions so that users who are not members of the Domain Admins group can create domain-based namespaces, and you can delegate management permissions so that users or groups can manage existing namespaces. In this section, you will delegate permissions to manage the namespace you created in the previous task.

**To delegate permission to manage an existing namespace**

1. In the console tree of the DFS Management snap-in, right-click *****server_or_domain***\Public**, and then click **Delegate Management Permissions**.
2. Type the name of a user or group that you want to manage the namespace, and then click **OK**.

After you finish this procedure, review the contents of the **Delegation** tab in the details pane. It should look similar to the following figure.




Notice that the user or group you added shows "Explicit" in the **How Permission Is Granted** column. "Explicit" means that you can remove the user or group from the delegation list by right-clicking the user or group, and then clicking **Remove**. Any users or groups that show "Inherited" have inherited management permissions from Active Directory, and you cannot remove them from the delegation list using the DFS Management snap-in.

| **Note** |
|---|
| To delegate the ability to create domain-based namespaces, see "Security requirements for creating and managing namespaces" in DFS Management Help. |

**Task 4: Add Folders to the Namespace**

In this task, you add three folders to the namespace. Two of the folders will have folder targets. The hierarchy of the folders you will add is shown in the "Elements of a Namespace" figure earlier in this guide.

**To create a folder named Software in the namespace**

1. In the console tree of the DFS Management snap-in, right-click *****server_or_domain***\Public**, and then click **New Folder**.

2. In **Name**, type **Software**, and then click **OK**.

Note that the previous procedure creates a new folder in the namespace to build depth in the namespace hierarchy. You are not specifying the name of an existing folder, nor will you store data in this folder. This folder will not have folder targets that direct clients to other servers.

After you finish this procedure, the Software folder is added to the console tree as shown in the following figure. (You might need to double-click the *****server_or_domain***\Public** root to display the Software folder.




Next, you add two folders with targets to the namespace. You create one folder named Tools within the Software folder, and you create another folder named Training Guides directly under the root named Public.

**To create a folder named Tools in the namespace**

1. In the console tree of the DFS Management snap-in, right-click the **Software** folder, and then click **New Folder**.

2. In **Name**, type **Tools**.

3. Click **Add** to add a folder target.

4. Click **Browse** to open the **Browse for Shared Folders** dialog box.

5. In **Server**, enter the name of the server that will host the Tools shared folder.

6. Click **New Shared Folder**.

7. In the **Create Share** dialog box, in the **Share name** box, type **Tools**, and then enter the local path where you want the shared folder to be created. If the folder does not exist, you are prompted to create it. Click **OK** to close all dialog boxes.

After you finish this procedure, the Tools folder is added to the console tree as shown in the following figure. (You might need to double-click the Software folder to display the Tools folder.) Notice the icon next to the Tools folder and how it differs from the Software folder's icon. This icon appears next to all folders that have targets to differentiate them from folders that do not have targets.




Now, select the Tools folder and review the contents of the **Folder Targets** tab in the details pane. Notice there is a single path shown. This means that only one server hosts the folder target that corresponds to the Tools folder. If that server becomes unavailable, the shared folder is also unavailable.

To increase the availability of the Tools folder, you can add a second folder target.

**To add a second folder target to the Tools folder**

1. In the console tree of the DFS Management snap-in, right-click the **Tools** folder, and then click **Add Folder Target**.

2. Click **Browse** to open the Browse for Shared Folders dialog box.

3. In **Server**, enter the name of another server that will host the Tools shared folder. Be sure to enter a different server from the one you specified in the previous procedure.

4. Click **New Shared Folder**.

5. In the **Create Share** dialog box, in the **Share name** box, type **Tools**, and then enter the local path where you want the shared folder to be created. If the folder does not exist, you are prompted to create it. Click **OK** to close all dialog boxes.

6. You are prompted to choose whether to create a replication group for these folder targets. For now, click **No**. You will enable DFS Replication on this folder in a later task.

**To create a folder named Training Guides in the namespace**

1. In the console tree of the DFS Management snap-in, right-click *****server_or_domain***\Public**, and then click **New Folder**.

2. In **Name**, type **Training Guides**.

3. Click **Add** to add a folder target.

4. Click **Browse** to open the Browse for Shared Folders dialog box.

5. In **Server**, enter the name of the server that will host the Training Guides shared folder.

6. Click **New Shared Folder**.

7. In the **Create Share** dialog box, in the **Share name** box, type **Training Guides**, and then enter the local path where you want the shared folder to be created. If the folder does not exist, you are prompted to create it. Click **OK** to close all dialog boxes.

When you finish these procedures, your namespace will look like the following figure.




**Task 5: Change How Targets Are Ordered in Referrals**

In this task, you change settings that optimize how targets are ordered in referrals. If you are not familiar with referrals, a *referral* is a list of targets that a client computer receives from a domain controller or namespace server when the user accesses a namespace root or folder with targets in the namespace. The referral tells the client which servers host the associated root target or folder target. So, for example, when a client navigates to *server_or_domain*\Public, the client receives a root referral that contains a list of root targets on the namespace servers. When the client then navigates to the Tools folder, which has folder targets, the client receives a folder referral that contains a list of folder targets that correspond to the Tools folder.

When a client requests a referral, the Distributed File System service takes into account the site of the client and the site of the target and provides a referral with targets that are ordered according to the current referral ordering method. By default, targets in a client's site are listed first in a referral in random order, followed by a list of targets outside of the client's site, sorted by lowest cost.

To fine-tune how targets outside of a client's site are ordered, you can change the ordering method for an entire namespace or for individual folders with targets. Changing the ordering method is an important consideration in namespaces whose targets span sites. For example, there might be situations in which you want to prevent the client from accessing targets outside of its own site. If so, you can configure the namespace root or folder with targets so that clients receive referrals only for targets within their own site.

To further optimize how targets are listed in referrals, you can set target priority, which overrides the ordering method. For example, you can specify that a target is always first or last in a referral, regardless of the client's site, or you can specify that a target is always first or last among the targets that have the same connection cost. One common scenario for using target priority is when you have a "hot standby" server that is considered the server of last resort. In this scenario, you can specify that the standby server always appears last in referrals, and clients will fail over to this server only if all the other servers fail or become unavailable due to network outages.

In the following procedures, you verify the referral ordering method for the namespace and choose target priority of a folder target.

**To verify the referral ordering method for a namespace**

1. In the console tree of the DFS Management snap-in, right-click *****server_or_domain***\Public**, and then click **Properties**.

2. On the **Referrals** tab, in **Ordering method**, verify that **Lowest cost** is selected.

In the lowest cost ordering method, also called least expensive target selection or site costing in previous documentation, targets in a referral are ordered as follows:

1. Targets in the same site as the client are listed in random order at the top of the referral.

2. Targets outside of the client's site are listed in order of lowest cost to highest cost. Referrals with the same cost are grouped together and within each group the targets are listed in random order.

This method ensures that clients do not traverse expensive wide area network (WAN) links to access targets when lower-cost targets are available. This ordering method works in both stand-alone and domain-based namespaces, as long as all namespace servers and all domain controllers are running Windows Server 2003.

| **Note** |
| --- |
| If you do not want clients to access folder targets outside of their site, you can override the ordering method for individual folders. To do this, right-click a folder with targets in the console tree, click **Properties**, click the **Referrals** tab, and then click **Exclude targets outside of the client's site**. Note that if no same-site targets are available, the client fails to access the folder because no folder targets are returned in the referral. |

In the next procedure, you change the priority of one of the folder targets of the Tools folder.

**To change the priority of a folder target**

1. In the console tree of the DFS Management snap-in, click the **Tools** folder.

2. In the details pane, on the **Folder Targets** tab, right-click one of the folder targets, and then click **Properties**.

3. On the **Advanced** tab, click **Override referral ordering**, and then click **Last among all targets**.

**Task 6: Rename and Move a Folder**

You can use the DFS Management snap-in to rename folders or move folders to another location in the namespace. This is useful if you need to change a folder name or restructure the namespace.

In this task, you rename the Training Guides folder to Training Demos and move it to the Software folder. Currently, your namespace should look similar to the following figure.




**To rename the Training Guides folder**

1. In the console tree of the DFS Management snap-in, right-click the **Training Guides** folder, and then click **Rename Folder**.
2. In the **Rename Folder** dialog box, in **New name**, type **Training Demos**.

**To move the Training Demos folder**

- In the console tree of the DFS Management snap-in, click the **Training Demos** folder, and then drag it to the **Software** folder.

After you finish these procedures, your namespace should look like this:

**Task 7: Replicate a Folder in the Namespace Using DFS Replication**

In this task, you enable DFS Replication on the Tools folder. If you recall from "Task 4: Add Folders to the Namespace," you created two folder targets for the Tools folder. Because users can be directed to either one of the folder targets, you need to ensure that the contents of the folders are kept synchronized.

If you are familiar with File Replication service (FRS) in Windows Server 2003, you know that FRS is only supported in domain-based namespaces. In Windows Server 2003 R2, you can use DFS Replication in both stand-alone and domain-based namespaces. Therefore, you can complete this task regardless of the type of namespace you created in "Task 1: Create a Namespace."

| **Important** |
|---|
| To perform this task, you need to have Active Directory deployed in your test lab environment, and you must be a member of the Domain Admins group or have been delegated the ability to create replication groups to perform this task. |

| **Important** |
|---|
| After you complete this task, replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member. |

**To enable DFS Replication on the Tools folder**

1. In the console tree of the DFS Management snap-in, right-click the **Tools** folder, and then click **Replicate Folder**.
2. Follow the steps in the Replicate Folder Wizard and supply the information in the following table.

| Replicate Folder Wizard Page | What to Enter |
|---|---|
| **Replication Group and Replicated Folder Name** | Accept the defaults. |
| **Replication Eligibility** | Accept the defaults. |
| **Primary Member** | If the folder targets are empty, choose either member. If both folder targets contain content, choose the member that has the most up-to-date content. |
| **Topology Selection** | Select **Full mesh.** |
| **Replication Group Schedule and Bandwidth** | Select **Replicate continuously using the specified bandwidth**. |
| **Review Settings and Create Replication Group** | Click **Create** to create the replication group. |
| **Confirmation** | Click **Close** to close the wizard. |
| **Replication Delay** | Click **OK** to close the dialog box that warns you about the delay in initial replication. |

After you finish the previous procedure, navigate to the **Replication** node in the console tree. Notice that a new replication group has been

created, as shown in the following figure.




If you are not familiar with DFS Replication terminology, a *replication group* is a set of servers, known as members, that participates in the replication of one or more replicated folders. A *replicated folder* is a folder that is kept synchronized on each member. When you enable DFS Replication on a folder with targets, the servers that host the folder targets become members of the replication group, and the folder targets are associated with the replicated folder. The name of the replication group matches the namespace path (Contoso.com\Public\Software\Tools), and the name of the replicated folder matches the folder name (Tools).

From the **Replication** node, you can manage aspects of DFS Replication, such as the schedule and bandwidth usage, file and subfolder filters, and the topology (a framework of replication paths between members). On the **Replicated Folders** tab in the details pane, you can also view the namespace path that corresponds to the replicated folder, as shown in the following figure.




If you navigate back to the **Tools** folder in the **Namespaces** node, notice that the **Replication** tab in the details pane shows that the Tools folder is being replicated using DFS Replication.




If one of the folders targets contained data when you enabled DFS Replication, you can verify that replication has completed by clicking the **Folder Targets** tab, right-clicking the folder target that initially held no data, and then clicking **Open in Explorer**. After the initial replication delay, the files in this folder target should match the files in the target that initially held the data.

Another way to view the status of replication is to create a diagnostic report. You will do this in the following task.

**Task 8: Create a Diagnostic Report**

In this task, you create a diagnostic report to check the status of replication. The diagnostic report is an .html file that includes error and warning events, replication statistics, backlogged files, and other information for each member of the replication group.

**To create a diagnostic report**

1. In the console tree of the DFS Management snap-in, under the **Replication** node, right-click the *domain*\Public\Software\Tools replication group, and then click **Create Diagnostic Report**.

2. Follow the steps in the Diagnostic Report Wizard and supply the information in the following table.

| Diagnostic Report Wizard Page | What to Enter |
|---|---|
| **Path and Name** | Accept the defaults. |
| **Members to Include** | Accept the defaults. |
| **Options** | Ensure that **Yes, count backlogged files in this report** is selected, and also click the **Count the replicated files and their sizes on each member** check box. |
| **Review Settings and Create Report** | Click **Create** to create the diagnostic report. |
| **Confirmation** | The wizard closes automatically, and the diagnostic report appears. |

Review the diagnostic report created for the Tools replication group. In particular, take a look at the following sections:

- Note the DFS Replication efficiency savings. This savings will change over time as files are added and changed.

- Review any errors or warnings, if any, for the members. These are typically event log errors that appear in the member's respective DFS Replication event log.

- In the informational section for each member, review the replicated folder status (the status will be "Normal" after initial replication is

complete) and other information. Notice that the primary member will show different statistics from the non-primary member; this is because data originated from the primary member and replicated to the non-primary member during initial replication.

**Task 9: Browse the Namespace**

In this task, you browse the namespace you created and view the referrals in the client's referral cache as you browse portions of the namespace. Viewing referrals cached on the client is useful in troubleshooting scenarios. The following procedures assume you are performing these tasks from a client computer running Windows XP or Windows Server 2003.

**To browse to the namespace root and view the root referral**

1. Click **Start**, click **Run**, type *****server_or_domain***\Public**, and then click **OK**.

   Windows Explorer opens and your view of the namespace looks similar to the following figure:




2. In Windows Explorer, click the **Folders** button to display the Public root in the folder tree.
3. In the folder tree, right-click **Public**, and then click **Properties** to open the **Properties** dialog box.
4. On the **DFS** tab, review the paths listed under **Referral list**. These are the root targets in the root referral that the client received when it accessed *****server_or_domain***\Public**. These should match the root targets you created earlier in this guide. The target marked Active is the target currently connected to your client computer.
5. Click **OK** to close the dialog box.

**To browse to the Tools and Training Demos folders and view their folder referrals**

1. In Windows Explorer, double-click the **Software** folder. You should see two folders, **Tools** and **Training Demos**.
2. Double-click the **Tools** folder to open it.
3. In the folder tree, right-click the **Tools** folder, and then click **Properties**.
4. On the **DFS** tab, review the paths listed under **Referral list**. These are the folder targets in the folder referral that the client received when it accessed *****server_or_domain***\Public\Software\Tools**. These should match the folder targets you created earlier in this guide, and the target you marked as **Last among all targets,** when you set the target priority, should be listed last. The target marked Active is the target currently connected to your client computer.
5. Click **OK** to close the dialog box.
6. Click the **Training Demos** folder in the folder tree to open it.
7. Right-click the **Training Demos** folder in the folder tree, click **Properties**, and then click the **DFS** tab. Notice that only one folder target is listed in the referral list. Your client computer is currently connected to this folder target.

**Task 10: Test Failover**

In this task, disable the network card or turn off the server that hosts one of the root targets for the *****domain***\Public** namespace. Do the same for a server that hosts one of the folder targets for the Tools folder. After the network cards are disabled or the servers are turned off, repeat the procedures in "Task 9: Browse the Namespace." The procedures should work because another server continues to host the *****domain***\Public** namespace and the Tools folder.

**Step-by-Step Guide to Deploying DFS Replication**

This section walks you through the process of deploying DFS Replication in a test lab. Although it is not necessary to have completed the tasks in "Step-by-Step Guide to Deploying a Namespace," it is helpful to complete the previous section first so that you have an existing namespace in which to publish a replicated folder.

**Introduction to DFS Replication**

DFS Replication is the new state-based, multimaster replication engine in Windows Server 2003 R2. Although some DFS Replication concepts and processes are similar to the concepts and processes in File Replication service (FRS), there are several important differences that you should be aware of before you deploy DFS Replication.

**DFS Replication Concepts**

First, let's review the basic concepts of DFS Replication. These concepts—replication groups, connections, members, and replicated folders—are illustrated in the following figure.




As this figure shows, a *replication group* is a set of servers, known as *members*, that participates in the replication of one or more replicated folders. A *replicated folder* is a folder that is kept synchronized on each member. In the previous figure, there are two replicated folders, Projects and Proposals. As data changes in each replicated folder, the changes are replicated across *connections* between the members. The connections between all members form the replication *topology*.

Creating multiple replicated folders in a single replication group simplifies the process of deploying replicated folders, because the topology, schedule, and bandwidth throttling for the replication group are applied to each replicated folder. To deploy additional replicated folders, you can use a short wizard to define the local path and permissions for the new replicated folder. Each replicated folder also has its own settings, such as file and subfolder filters, so that you can filter out different files and subfolders for each replicated folder.

The replicated folders stored on each member can be located on different volumes in the member, and the replicated folders do not need to be shared folders or part of a namespace, though the DFS Management snap-in makes it easy to share replicated folders and optionally publish them in an existing namespace. You will do both in one of the tasks later in this guide.

**What to Expect During Initial Replication**

When you first set up replication, you must choose a primary member. Choose the member that has the most up-to-date files that you want replicated to all other members of the replication group, because the primary member's content is considered "authoritative." This means that during initial replication, the primary member's files will always win the conflict resolution that occurs when the receiving members have files that are older or newer than the same files on the primary member.

The following explanations will help you better understand the initial replication process:

- Initial replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

- Initial replication always occurs between the primary member and the receiving replication partners of the primary member. After a member has received all files from the primary member, then that member will replicate files to its receiving partners as well. In this way, replication for a new replicated folder starts from the primary member and then progresses out to the other members of the replication group.

- When receiving files from the primary member during initial replication, the receiving members that contain files that are not present on the primary member move those files to their respective DfsrPrivate\PreExisting folder. If a file is identical to a file on the primary member, the file is not replicated. If the version of a file on the receiving member is different from the primary member's version, the receiving member's version is moved to the Conflict and Deleted folder and remote differential compression (RDC) can be used to download only the changed blocks.

- To determine whether files are identical on the primary member and receiving member, DFS Replication compares the files using a hash algorithm. If the files are identical, only minimal metadata is transferred.

- After the initialization of the replicated folder, the "primary member" designation is removed. Initialization takes place after all files that exist before DFS Replication picks up the configuration are added to the DFS Replication database. The member that was previously the primary member is then treated like any other member and its files are no longer considered authoritative over those of other members that have completed initial replication. Any member that has completed initial replication is considered authoritative over members that have not completed initial replication.

**Deploying DFS Replication**

The tasks in this section walk you through the process of deploying DFS Replication, adding a member to a replication group, publishing a replicated folder in a namespace, and creating a diagnostic report.

**Task 1: Create a Multipurpose Replication Group and Two Replicated Folders**

To enable DFS Replication, you use the **New Replication Group Wizard** to specify the members, topology, and default schedule and bandwidth for the replication group. In this task, you create a replication group named Data Distribution and two replicated folders named Antivirus Signatures and LOB Data.

| **Important** |
| --- |
| When you create a new replication group, replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and each member's long polling interval (60 minutes). |

Before you enable replication, you will create two folders on one of the servers to be added to the replication group. You will then add files to the folders.

**Create folders named Antivirus Signatures and LOB Data**

1. On one server, create two separate (non-overlapping) folders named Antivirus Signatures and LOB Data. Add some files to each folder but do not exceed the recommended limits described in "DFS Replication limits" in DFS Management Help.

2. Optionally, you can create the same folders on the second server. You can add the same files or different files from the primary member. If you add the same files, the files will be used for prestaging and will not be re-replicated. If you add files that don't exist on the primary member, those files will be moved to the PreExisting folder on the second member. (This folder is in the DfsrPrivate folder under the replicated folder's local path.) If you want to observe how the primary member's content becomes authoritative during initial replication, use updated versions of the files you added to the primary member. These updated files will be moved to the Conflict and Deleted folder on the non-primary members.

Next, create a replication group to replicate files between the two servers.

**To create a replication group and two replicated folders**

1. In the console tree of the DFS Management snap-in, right-click the **Replication** node, and then click **New Replication Group**.

2. Follow the steps in the New Replication Group Wizard and supply the information in the following table.

| New Replication Group Wizard Page | What to Enter |
|---|---|
| **Replication Group Type** | Select **Multipurpose replication group**. |
| **Name and Domain** | In **Name of replication group**, type **Data Distribution**. |
| **Replication Group Members** | Click **Add** to select at least two servers that will participate in replication. The servers must have the DFS Replication Service installed. |
| **Topology Selection** | Select **Full mesh**. |
| **Replication Group Schedule and Bandwidth** | Select **Replicate continuously using the specified bandwidth**. |
| **Primary Member** | Select the member that has the most up-to-date content that you want to replicate to the other member. |
| **Folders to Replicate** | Click **Add** to enter the local path of the LOB Data folder you created earlier on the first server. Use the name **LOB Data** for the replicated folder name. Repeat this procedure and enter the local path of the **Antivirus Signatures** folder. |
| **Local Path of LOB Data on Other Members** | On this page, you specify the location of the LOB Data folder on the other members of the replication group. To specify the path, click **Edit**, and then on the **Edit Local Path** dialog box, click **Enabled**, and then type the local path of the LOB Data folder. |
| **Local Path of Antivirus Signatures on Other Members** | On this page, you specify the location of the Antivirus Signatures folder on the other members of the replication group. To specify the path, click **Edit**, and then on the **Edit Local Path** dialog box, click **Enabled**, and then type the local path of the Antivirus Signatures folder. |
| **Review Settings and Create Replication Group** | Click **Create** to create the replication group. |
| **Confirmation** | Click **Close** to close the wizard. |
| **Replication Delay** | Click **OK** to close the dialog box that warns you about the delay in initial replication. |

After you finish the New Replication Group Wizard, click the new replication group named **Data Distribution** located under the **Replication** node in the console tree as shown in the following figure:




Notice the four tabs in the details pane: **Memberships**, **Connections**, **Replicated Folders**, and **Delegation**. Each of these tabs displays different details about the selected replication group, its members, and its replicated folders. Review the following details about each tab.

- On the **Memberships** tab, notice that entries on the tab are sorted by replicated folder and that there are two replicated folders listed. For example, the rows under the **Replicated Folder: Antivirus Signatures** heading are the members that host the Antivirus Signatures replicated folder. Double-click a member to view per-member, per-replicated folder properties on the **General**, **Replicated Folder**, and **Advanced** tabs. For example, on the **Advanced** tab, you can view the location and size of the staging folder and Conflict and Deleted folder on the selected member.

- On the **Connections** tab, two connections are listed. Each connection is a one-way replication path, so replication between two members requires two connections that replicate data in the opposite direction. Each connection has a schedule and other settings, such as a check box for enabling or disabling remote differential compression (RDC). Double-click a connection to view its settings.

- On the **Replicated Folders** tab, notice that two replicated folders are listed and that they are not published in a namespace. Double-click a replicated folder to view its properties, such as file and subfolder filters.

- On the **Delegation** tab, review the default users and groups granted permissions to manage the replication group. Any users or groups shown as "Inherited" have inherited management permissions from Active Directory, and you cannot remove them from the delegation list using the DFS Management snap-in.

**Task 2: Add a New Member to a Replication Group**

In this task, you add a third server to the Data Distribution replication group and specify where one of the two replicated folders, Antivirus Signatures, will be stored on the new member. You'll use a new feature in DFS Replication to specify that the LOB Data replicated folder is not replicated to the new member. You also will create a custom schedule that applies only to the connections to and from the new member.

| **Important** |
|---|
| Replication does not begin immediately on the new member. The DFS Replication settings for the new member must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency, the short polling interval (5 minutes) on the new member, and the long polling interval (60 minutes) on existing members. |

**To add a new member to the Data Distribution replication group**

1. In the console tree of the DFS Management snap-in, right-click the **Data Distribution** replication group, and then click **New Member**.

2. Follow the steps in the **New Member Wizard** and supply the information in the following table.

| New Member Wizard Page | What to Enter |
|---|---|
| **New Member** | Enter the name of the server to add to the replication group. The server must have the DFS Replication Service installed. |
| **Local Path of Replicated Folders** | Select the **Antivirus Signatures** replicated folder, click **Edit**, click **Enabled**, and then enter the local path of the replicated folder to be created on the new member. When you close the **Edit Local Path** dialog box, notice that the LOB Data replicated folder shows **<Disabled>**, which means that this replicated folder will not be replicated to the new member. Because you only want the Antivirus Signatures folder to be replicated to the new member, you can ignore the warning message that appears. |
| **Connections** | Under **Available members**, click a member, and then click **Add**. Repeat this step to add the second member. The new member will replicate directly with both existing members. |
| **Replication Schedule** | Select **Custom connection schedule**, and then click **Edit Schedule**. In the Edit Schedule dialog box, click **Details** to expand the schedule, and then select the entry that begins **Sunday 12:00 AM** and then click **Edit**. In the Edit Schedule dialog box, under **Bandwidth usage**, click **128 Mbps**. |

| | |
|---|---|
| **Review Settings and Create Member** | Click **Create** to add the new member to the Data Distribution replication group. |
| **Confirmation** | Click **Close** to close the wizard. |
| **Replication Delay** | Click **OK** to close the dialog box that warns you about the delay in initial replication. |

After you finish the wizard, click **Data Distribution** in the console tree, and then review the contents of the **Connections** tab. It should look similar to the following figure:




Notice that in the **Schedule Type** column, connections to and from the new member show **Custom Connection Schedule** instead of **Replication Group Schedule**. These show **Custom Connection Schedule** because you chose a custom schedule when you added the new member. Creating custom schedules for individual connections allows you to fine-tune the replication interval and bandwidth used when replicating to specific members. Although it isn't obvious in the user interface, each connection marked **Custom Connection Schedule** is a separate schedule. You can modify one schedule marked **Custom Connection Schedule**, but the other custom schedules are not affected.

| **Note** |
|---|
| To change how the items are grouped, click a column heading. For example, to group the items by schedule type, click the **Schedule Type** heading. |

Entries marked **Replication Group Schedule** use the default replication schedule; this schedule is applied to all connections in the replication group that do not have a custom schedule. To modify the default replication schedule, right-click the **Data Distribution** replication group in the console tree, click **Properties**, and then click **Edit Schedule**. To change a connection schedule from a custom connection schedule to the replication group schedule or vice versa, on the **Connections** tab in the details pane, double-click the connection, click the **Schedule** tab, and then click **Replication group schedule** or **Custom connection schedule**.

#### Task 3: Share and Publish Replicated Folders in a Namespace

When you created replicated folders in the previous tasks, you specified the local path of a folder on each member of the replication group. Unless the local path on each server was previously shared, users cannot access the replicated folders after they are created. To make replicated folders available to users, you must share them and, optionally, publish them in an existing namespace.

In this task, you publish the LOB Data replicated folder in the *****server_or_domain***\Public** namespace that you created in the DFS Namespaces step-by-step section. If you did not complete the previous step-by-step section or do not have a namespace in your test lab, skip this procedure.

**To share the LOB Data replicated folder and publish it in a namespace**

1. In the console tree of the DFS Management snap-in, under the **Replication** node, click the **Data Distribution** replication group.
2. In the details pane, click the **Replicated Folders** tab, right-click the **LOB Data** replicated folder, and then click **Share and Publish in Namespace**.
3. Follow the steps in the **Share and Publish Replicated Folder Wizard** and supply the information in the following table.

| **Share or Publish Replicated Folder Wizard Page** | **What to Enter** |
|---|---|
| **Publishing Method** | Select **Share and publish the replicated folder in a namespace**. |

| | |
|---|---|
| **Share Replicated Folders** | For each member that shows **[Shared Folder Needed]** in the **Action** column, select the member, and then click **Edit** to create the new shared folder and adjust shared folder permissions if necessary. If the **Action** column shows **Create shared folder: LOB Data** or **Existing Shared Folder**, you can click **Next**. |
| **Namespace Path** | In **Parent folder in namespace**, type *****server_or_domain***\Public\Software**. |
| **Review Settings and Share Replicated Folder** | Click **Share** to share the replicated folders and publish the LOB Data replicated folder in the namespace. |
| **Confirmation** | Click **Close** to close the wizard. |

After you finish the wizard, review the console tree and the **Replicated Folders** tab in the details pane. First, notice in the console tree that an LOB Data folder was added to the namespace and the folder icon indicates that the folder is replicated, as shown in the following figure.




Next, review the namespace path listed in the **Replicated Folders** tab, which should look similar to the following figure.




You can see that the LOB Data replicated folder is published in a namespace but the Antivirus Signatures replicated folder is not. To access the LOB Data folder in the namespace, in the Run dialog box, type *****server_or_domain***\Public\Software\LOB Data**.

| **Note** |
|---|
| If you want to stop publishing the LOB Data replicated folder in the namespace, you can right-click the replicated folder and then click **Remove from Namespace**. |

#### Task 4: Create a Replication Group for Collection Purposes

In this task, assume that you have a hub server in a central hub or data center location and a branch server in a remote office. The branch server contains folders named Projects and Proposals that are very important to the branch office, but there is a concern that the backups performed at the branch office are performed incorrectly. You want to back up the Projects and Proposals folders from the data center to ensure that the backups are successful.

To accomplish this goal, you will set up a replication group for data collection purposes. This type of replication group consists of two members and one replicated folder for each folder that you want to back up from the hub server at the data center. The permissions that are set on the replicated folders on the branch server will be applied to the replicated folders on the hub server. You specify a single folder on the hub server under which subfolders for the replicated folders will be created. This allows you to back up multiple replicated folders from a single location on the hub server.

| **Important** |
|---|
| When you create a new replication group, replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and each member's long polling interval (60 minutes). |

**To create a replication group to replicate the Projects and Proposals folders from a branch server to a hub server**

1. In Windows Explorer or from the command prompt, create a folder named **Projects** and a folder named **Proposals** on a server that will act as the branch server. The folders should be separate folders (that is, not nested in one another).
2. Add some data to the **Projects** and **Proposals** folders on the branch server.
3. In Windows Explorer or from the command prompt, create a folder named **Branch_Backups** on a server that will act as the hub server. Do not put data in this folder.
4. In the console tree of the DFS Management snap-in, right-click the **Replication** node, and then click **New Replication Group**.
5. Follow the steps in the **New Replication Group Wizard** and supply the information in following table.

| New Replication Group Wizard Page | What to Enter |
|---|---|

| | |
|---|---|
| **Replication Group Type** | Select **Replication group for data collection**. |
| **Name and Domain** | In **Name of replication group**, type **Branch Backups**. |
| **Branch Server** | Type the name of a server that will act as the branch server. |
| **Replicated Folders** | Click **Add**. In the **Add Folder to Replicate** dialog box, type the local path of the Projects folder you created in Step 1. Repeat this step for the local path of the Proposals folder. |
| **Hub Server** | Type the name of a server that will act as the hub server. This is the server where you can back up the Projects and Proposals folders using backup software. |
| **Target Folder on Hub Server** | In **Target folder**, type the path of the folder you created in Step 3. |
| **Replication Group Schedule and Bandwidth** | Select **Replicate continuously using the specified bandwidth**. |
| **Review Settings and Create Replication Group.** | Click **Create** to create the replication group. |
| **Confirmation** | Click **Close** to close the wizard. |
| **Replication Delay** | Click **OK** to close the dialog box that warns you about the delay in initial replication. |

After you finish the wizard, click the **Branch Backups** replication group in the console tree and view the **Memberships** tab in the details pane. Notice that two replicated folders were created, **Projects** and **Proposals**, as shown in the following figure:




In the previous figure, notice that the path of the Projects and Proposals replicated folders on server BO-08 (the hub server) are both within the C:\Branch_Backups folder. This allows you to back up both replicated folders from a single location on the hub.

#### Task 5: Create a Diagnostic Report

In this task, you create a diagnostic report to check the status of replication. The diagnostic report is an .html file that includes error and warning events, replication statistics, backlogged files, and so forth for each member of the replication group.

**To create a diagnostic report**

1. In the console tree of the DFS Management snap-in, under the **Replication** node, right-click the **Branch Backups** replication group, and then click **Create Diagnostic Report**.
2. Follow the steps in the Diagnostic Report Wizard and supply the information in the following table.

| Diagnostic Report Wizard Page | What to Enter |
|---|---|
| **Path and Name** | Accept the defaults. |
| **Members to Include** | Accept the defaults. |
| **Options** | Ensure that **Yes, count backlogged files in this report** is selected, and also click the **Count the replicated files and their sizes on each member** check box. |
| **Review Settings and Create Report** | Click **Create** to create the diagnostic report. |
| **Confirmation** | The wizard closes automatically, and the diagnostic report appears. |

Review the diagnostic report created for the Tools replication group. In particular, take a look at the following sections:

- Note the DFS Replication efficiency savings. This savings will change over time as files are added and changed.
- Review any errors or warnings, if any, for the members. These are typically event log errors that appear in the member's respective DFS Replication event log.
- In the informational section for each member, review the replicated folder status (the status will be "Normal" after initial replication is

complete) and other information. Notice that the primary member will show different statistics from the non-primary member; this is because data originated from the primary member and replicated to the non-primary member during initial replication.

**See Also**
**Other Resources**
Microsoft Branch Offices Web Site
Branch Office Tech Center
Distributed File System Web Site

# Extending Your Active Directory Schema in Windows Server 2003 R2

Applies To: Windows Server 2003 R2

The schema defines objects and attributes and the types of data they can store for Active Directory. You can extend the schema to support new operating system features and applications, or to add RFC-defined attributes that are commonly found in other Lightweight Directory Access Protocol (LDAP) directories. You might also extend the schema so that the directory can store information that is interesting to your organization.

Extending the schema for a new operating system version is a two-phase operation that consists of a one-time schema extension on the domain controller that hosts schema operations master role (an operations master is also known as flexible single master operations or FSMO), and a one-time update on the domain controller that holds infrastructure master role in each domain in the forest. Both operations are performed with the ADPREP tool located on Windows Server 2003 R2 Disc 2.

These topics focus on using the adprep /forestprep command to add the Windows Sever 2003 R2 schema extensions that are required to support the following Windows Server 2003 R2 features or server roles:

- Distributed File System (DFS) Replication Service
- Domain controllers running Windows Server 2003 R2
- Identity Management for UNIX
- The Deployed Printers functionality of the Print Management console

These schema additions are automatically installed when you use Windows Server 2003 R2 installation media to create the first domain controller in a new Active Directory forest. However, if you are planning to add a domain controller that runs Windows Server 2003 R2 to an existing forest, you must extend the schema first. The following topics explain the Windows Server 2003 R2 schema update, interoperability considerations for x86-based versions and x64-based versions of Windows Server 2003 R2, and the steps for extending the schema:

- Active Directory Schema Update
- Interoperability of Adprep for x86-based Versions and x64-based Versions of Windows Server 2003 R2
- Steps for Extending the Schema

# Active Directory Schema Update

Applies To: Windows Server 2003 R2

Updating the Active Directory schema is an irreversible operation that must be performed on the schema master. The adprep /forestprep command adds Windows Server 2003 R2 object classes and attributes to Active Directory so that they can be used by the operating system, server roles, and applications.

You must update the schema before you upgrade an existing domain controller to Windows Server 2003 R2 (using Winnt32.exe) or before you add the Active Directory domain controller role on a computer running Windows Server 2003 R2 by using the Active Directory Installation Wizard (Dcpromo.exe). You must extend the schema before you perform either of these actions for the first time in an existing forest.

Typically, schema updates are approached cautiously because they can only be reversed by a forest recovery. However, Windows Server 2003 R2 schema updates have been designed and tested for compatibility with Windows client, server, and domain controller roles. Adding the Windows Server 2003 R2 schema updates does not affect the compatibility of client or server operating systems that are deployed inside or outside the forest.

You should test the Windows Server 2003 R2 schema updates in a lab environment to ensure they will not conflict with any applications that run in your environment. There should not be any conflicts if your applications use RFC-compliant object and attribute definitions.

Adprep skips redundant updates. Conflicting updates, such as the introduction of duplicate object identifiers, cause Adprep to stop until the conflicts are reconciled by an administrator.

In Windows Server 2003 R2, three new features, and the domain controller server role if you are adding a Windows Server 2003 R2 domain controller to an existing forest, require that you extend the existing schema before you attempt to deploy them. The following table explains why each of these features requires a schema extension.

| **Note** |
| --- |
| Microsoft recommends that you add the Windows Server 2003 R2 schema updates to your existing Active Directory installation before you deploy computers running Windows Server 2003 R2, regardless of whether you plan to use these new Windows Server 2003 R2 features or not. |

| | |
| --- | --- |
| DFS Replication Service | Requires that new object classes be added to the domain directory partition to contain replication groups and content sets. Also requires that new object classes be added to the **computer** object class in the domain directory partition to contain computer-specific settings to support Distributed File System (DFS) replication. |
| Domain controller server role | Requires new object classes be added to support objects that are defined for new features in Windows Server 2003 R2. |
| Identity Management for UNIX | Requires new object classes to represent items in a Network Information Service (NIS) map. |
| The Deployed Printers functionality of the Print Management console | Requires a new object class to contain information about printers so that their connections can be deployed through Group Policy. |

When you run adprep /forestprep to add the Windows Server 2003 R2 schema updates, you do not have to upgrade your existing domain controllers to Windows Server 2003 R2; they can continue to run Windows NT® Server 4.0, Windows® 2000 Server, or Windows Server 2003. Note, however, that Identity Management for UNIX must be installed on a Windows Server 2003 R2 domain controller. DFS Replication Service and Print Management Console can run on any computer that runs Windows Server 2003 R2.

The new object classes and attributes for the Windows Server 2003 R2 schema update are defined in a file named Sch31.ldf, which is located in the \cmpnents\R2\adprep folder on the Windows Server 2003 R2 Disc 2 (not in the \I386 folder of the Windows Server 2003 R2 Disc 1). You can view Windows Server 2003 R2 .ldf files in a text editor, but under no circumstances should their contents be modified. Modification of the .ldf files can cause irreversible changes and forest-wide failure. Use the

A**dprep** command to make all schema updates. Do not use import tools such as LDIFDE to import the Windows Server 2003 R2 .ldf files.

# Interoperability of Adprep for x86-based Versions and x64-based Versions of Windows Server 2003 R2

Applies To: Windows Server 2003 R2

You must run the version of Adprep that is compatible with the version of the Windows server operating system that is running on the schema operations master. Specifically, you must run the x86-based version of Adprep on an operations master that runs an x86-based version of Windows. Similarly, you must run the x64-based version of Adprep on an operations master that runs an x64-based version of Windows.

| **Windows server version running on the schema master** | **Version of Windows Server 2003 R2 Adprep that must be run to extend the schema** |
|---|---|
| x86-based version | x86-based version of Windows Server 2003 R2 Adprep.exe |
| x64-based version | x64-based version of Windows Server 2003 R2 Adprep.exe |
| IA64-based version | IA64-based version of Windows Server 2003 R2 Adprep.exe |

Windows Server 2003 R2 Disc 2 includes a single version of Adprep that is compatible with the operating system architecture for that version of Windows (either x86-based or x64-based).

If you have another domain controller that runs a compatible version of Windows but that domain controller does not hold the schema operations master role, transfer the role to that domain controller and then run the adprep /forestprep command that is on Windows Server 2003 R2 Disc 2.

If the version of Adprep on the Windows Server 2003 R2 Disc 2 is not compatible with your current schema master or your other domain controllers that might host the role, see article 919151 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=82345) to obtain a version of Adprep that is compatible with the version of Windows that runs on your schema master.

**Example**

An administrator attempts to deploy the first Windows Server 2003 R2 x64-based domain controller, or a Windows Server 2003 R2 feature, in an existing forest that contains domain controllers that run only x86-based versions of Windows 2000 and Windows Server 2003.

When the administrator runs x64-based version of Windows Server 2003 R2 adprep /forestprep from Disc 2 on the x86-based schema master, the adprep /forestprep command fails with the error "Adprep.exe is not a valid 32 bit application."

Because the x64-based version of Windows Server 2003 R2 Disc 2 only contains the x64-based version of Adprep and there is no x-64-based domain controller that can host the schema master role, the administrator must run an x86-based version of Adprep. In this case, the administrator must obtain hotfix that contains a compatible version of Adprep, as described in article 919151 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=82345). There are no charges associated with obtaining the hotfix.

Files with the .ldf file name extension do not depend on the type of processor (x86-based or x64-based) on the schema master. These files run on all versions of Windows Server 2003 R2.

# Steps for Extending the Schema

Updated: September 28, 2009

Applies To: Windows Server 2003 R2

Before you install one of the new features that is described in Active Directory Schema Update or before you add a domain controller running Windows Server 2003 R2 to a forest for the first time (unless it is the first domain controller in a new forest), you must first extend the schema with the Adprep tool. Perform the following steps to extend the schema:

- Verify Active Directory functionality before you apply the schema extension
- Apply the schema extension
- Verify the schema extension

**Verify Active Directory functionality before you apply the schema extension**

Verify Active Directory functionality before you update the schema to help ensure that the schema extension proceeds without error. At a minimum, ensure that all domain controllers for the forest are online and performing inbound replication.

**To verify Active Directory functionality before you apply the schema extension**

1. Log on to an administrative workstation that has the Windows Support Tool Repadmin.exe installed.

   | Note |
   | --- |
   | The Support Tools are located on the operating system installation media in the Support\Tools folder. |

2. Open a command prompt, and then change directories to the folder in which the Windows Support Tools are installed.
3. At a command prompt, type the following, and then press ENTER:

   **repadmin /replsum /bysrc /bydest /sort:delta**

   All domain controllers should show **0** in the **Fails** column, and the largest deltas (which indicate the time that has elapsed since the last successful replication) should be less than or roughly equal to the replication frequency of the site link that is used by the domain controller for replication. The default replication frequency is 180 minutes.

For more information about additional steps that you can take to verify Active Directory functionality before you apply the schema extension, see article 325379 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=71057).

**Apply the schema extension**

Use the following procedure to apply the Windows Server 2003 R2 schema extension to the Active Directory schema. In this procedure, you run the version of Adprep.exe that is included in the Windows Server 2003 R2 installation media on an existing domain controller (the schema operations master) in your forest.

**To apply the Windows Server 2003 R2 schema extension to the Active Directory schema**

1. Log on to the computer that holds the schema master operations role (also known as flexible single master operations or FSMO) as a member of the Schema Admins group and the Enterprise Admins group. If you are not sure which computer holds the schema master operations role, type the following at a command prompt, and then press ENTER:

   **netdom query FSMO**

   | Note |
   | --- |
   | The built-in Administrator account in the forest root domain is a member of the Schema Admins group by default. |

2. Verify that the schema master has performed inbound replication of the schema directory partition since the last time that the server restarted. Type the following at a command prompt, and then press ENTER:

   **repadmin /showrepl**

3. Locate the version of Adprep, either in the \cmpnents\R2 folder of the Windows Server 2003 R2 Disc 2 or from Microsoft hotfix 919151, that is compatible with the version of Windows that runs on your schema master.

   Each version of Windows Server 2003 R2 (x86-based or x64-based) ships with a single version of Adprep on Disc 2 that is compatible only with operation masters that run that version of Windows Server 2003 R2 (x86-based or x64-based).

   If your schema master is running run an x86-based version of Windows, run the x86-based version of Adprep.

   If your schema master is running run an x64-based version of Windows, run the x64-based version of Adprep.

   If your schema master does not run a version of Windows that is compatible with the version of Adprep that you plan to run, but your forest contains a domain controller that does run a compatible version of Windows, transfer the schema master role to that domain controller. Continue to step 4, and transfer the role back to the original role holder after the schema update is complete.

   If you do not have a compatible domain controller, obtain the hotfix described in article 919151 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=82345).

   To determine the version of the Windows operating system that is running on the schema master, type the following at a command prompt, and then press ENTER:

   **winver**

| **Important** |
| --- |
| Be sure to use the version of Adprep that is on Windows Server 2003 R2 Disc 2 or hotfix 919151, not the version of Adprep that is on Windows Server 2003 R2 Disc 1. |

4. Run **adprep /forestprep**. Change directories to the location that contains the appropriate Adprep version. Type the following command at the command prompt, and then press ENTER:

   **cd cmpnents\R2\ADPREP**

   **adprep /forestprep**

| **Note** |
| --- |
| When you change the schema on the schema master, the changes are automatically propagated to all other domain controllers in the forest. Therefore, it is not necessary to perform this operation on other domain controllers. Also, there is no need to run **adprep /domainprep** in any child domain where you have already installed a domain controller running Windows Server 2003 with Service Pack 1 (SP1). The necessary domain partition updates were performed when the domain controller running Windows Server 2003 SP1 was installed. |

**Verify the schema extension**

After you run Adprep, you can use the Windows Support tool ADSI Edit to verify the schema extension.

**To verify the schema extension**

1. Log on to an administrative workstation that has ADSI Edit installed.
2. Click **Start**, click **Run**, type **adsiedit.msc**, and then click **OK**.
3. Double-click **Configuration Container**, and then double-click **CN=Configuration,DC=***forest_root_domain*

   where *forest_root_domain* is the fully qualified domain name (FQDN) of your forest root domain.
4. Double-click **CN=ForestUpdates**.
5. Right-click **CN=Windows2003Update**, and then click **Properties**.
6. Verify that the **Revision** attribute value is **9**, and then close the **Properties** dialog box.
7. Double-click **Schema**.
8. Right-click **CN=Schema,CN=Configuration,DC=***forest_root_domain*

where *forest_root_domain* is the FQDN of your forest root domain.

9. Click **Properties**.

10. On the **Attribute Editor** tab, for **Select a property to view**, select **objectVersion**, and verify that the attribute **Value(s)** equals **31**.

# File Server Management Overview

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**File Server Management Overview**
File Server Management provides a centralized tool for managing your file server. A file server provides a central location on your network where you can store and share files with users across your network.

Using File Server Management, you can perform many file server management tasks, such as formatting volumes, creating shares, defragmenting volumes, creating and managing shares, setting quota limits, creating storage utilization reports, replicating data to and from the file server, managing Storage Area Networks (SANs), and sharing files with UNIX and Macintosh systems.

The tasks you can perform depend on the features selected when you installed the file server role. For information about how to install the file server role or to add optional features, see File server role: Configuring a file server in the Configure Your Server online Help (CYS.chm).

The following features are typically loaded with File Server Management:

**Share Folder Management**
Share Folder Management includes:

- **Shared Folders**

  Shared Folders enables you to manage shared resources over a network. With Shared Folders, you can:

  - Control user access permissions, session activity, and shared resource properties.
  - Create, view, and set permissions for shared resources.
  - View a summary of connections and resource use for local and remote computers.
  - View a list of all users who are connected over a network to the computer and disconnect one or all of them.
  - View a list of files that are opened by remote users and close one or all of the open files.

**Disk and Volume Management**
Disk and Volume Management includes:

- **Disk Defragmenter**

  Disk Defragmenter analyzes local volumes and consolidates, or defragments, fragmented files and folders so that each occupies a single, contiguous space on the volume. As a result, your system can access files and folders and save new ones more efficiently. By consolidating your files and folders, Disk Defragmenter also consolidates a volume's free space, making it less likely that new files will be fragmented. The process of consolidating fragmented files and folders is called defragmentation. Disk Defragmenter can defragment volumes that are formatted with the file allocation table (FAT) file system, the FAT32 file system, and the NTFS file system.

- **Disk Management**

  Disk Management is a system utility for managing hard disks and the volumes or partitions that they contain. With Disk Management, you can initialize disks, create volumes, format volumes with the FAT, FAT32, or NTFS file systems, and create fault-tolerant disk systems. Disk Management enables you to perform most disk-related tasks without restarting the system or interrupting users; most configuration changes take effect immediately.

**File Server Resource Manager**
File Server Resource Manager includes:

- **Quota Management**

  Quota management allows you to create quotas to limit the space allowed for a volume or folder and generate notifications when the quota limits are approached or exceeded. You can also define quota templates that can be easily applied to new volumes or folders and that can be used across an organization and auto apply quota templates to all existing folders in a volume or folder, as well as to any new subfolders created in the future.

- **File Screening Management**

  With File Screening Management you can create file screens to control the types of files that users can save and to send notifications when users attempt to save blocked files, define file screening templates that can be easily applied to new volumes or folders and that can be used across an organization, and create file screening exceptions that extend the flexibility of the file screening rules.

- **Storage Reports Management**

  Storage Reports Management enables you to schedule periodic reports that help you to identify trends in disk usage. You can also monitor attempts to save unauthorized files for all users or a selected group of users.

**DFS Management**

DFS Management includes:

- **DFS Management snap-in**

  You can use the DFS Management snap-in to configure and manage the two technologies in the Distributed File System solution: DFS Namespaces and DFS Replication. Together, you can use these technologies to collect and distribute data and for loose collaboration.

- **DFS Namespaces**

  You can create a namespace to group shared folders located on different servers and present them to users as a virtual tree of folders. A namespace provides numerous benefits, including increased availability of data, load sharing, and simplified data migration. When browsing the namespace, users see a single folder and are not aware the folder is hosted by multiple servers. When a user opens the folder, the client computer is automatically referred to a server in its own site; if no same-site servers are available, you can configure the namespace to refer the client to a server that has the lowest connection cost as defined in the Active Directory® directory service.

- **DFS Replication**

  DFS Replication is a new multimaster replication engine that is used to keep folders synchronized on multiple servers. Replicating data to multiple servers increases data availability and gives users in remote sites fast, reliable access to files. DFS Replication uses a new compression algorithm known as remote differential compression (RDC). RDC is a “diff over the wire” protocol that can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, re-arrangements of data in files, enabling DFS Replication to replicate only the deltas (changes) when files are updated.

# IIS 6.0 Content Revision Summary

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

This topic summarizes the changes made to the content in the IIS 6.0 Deployment Guide, the IIS 6.0 Operations Guide, and the IIS 6.0 Technical Reference since the original release of the documentation.

The table below gives a description of the area that each change affects.

| Area | Description |
|---|---|
| Application Configuration | Provides information about changes made to IIS 6.0 content about configuring applications. |
| Architecture | Provides information about changes made to IIS 6.0 content about IIS architecture. |
| Error Messages | Provides information about changes made to IIS 6.0 content about error messages. |
| Installation | Provides information about changes made to IIS 6.0 content about installation of IIS. |
| Introduction/Overview | Provides information about changes made to introductory content about IIS 6.0. |
| Logging | Provides information about changes made to IIS 6.0 content about logging. |
| Performance | Provides information about changes made to IIS 6.0 content about performance and scalability. |
| Reference | Provides information about changes made to IIS 6.0 reference content, such as metabase properties, registry entries, etc. |
| Security | Provides information about changes made to IIS 6.0 content about security. |
| Server Administration | Provides information about changes made to IIS 6.0 content about administering an IIS server. |
| Site Administration | Provides information about changes made to IIS 6.0 content about Web sites, Web directories, File Transfer Protocol (FTP) sites, Network News Transfer Protocol (NNTP) sites, and Simple Mail Transfer Protocol (SMTP) sites. |
| Troubleshooting | Provides information about changes made to IIS 6.0 troubleshooting content. |
| Web Applications | Provides information about changes made to IIS 6.0 content about Web applications. |

**Application Configuration**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Deployment Guide | Appendix D: Running 32-bit Applications on 64-bit Windows | Added new appendix with information about using the WOW64 emulator to run 32-bit Web applications on 64-bit Windows. |
| IIS 6.0 Operations Guide | Configuring IIS to Run 32-bit Applications on 64-bit Windows and subordinate topics | Added new topic about configuring IIS to run 32-bit Web applications on 64-bit Windows. |
| IIS 6.0 Operations Guide | Configuring the Maximum Number of Concurrent CGI Requests | Added topic with information about configuring the maximum number of Common Gateway Interface (CGI) requests allowed to run at one time by setting the MaxConcurrentCgisExecuting registry key. |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Configuring CGI Application Timeouts | Added topic with information about configuring CGI application time-out values. |
| IIS 6.0 Operations Guide | Configuring ASP.NET Applications in Worker Process Isolation Mode | Updated which **processModel** settings still apply in worker process isolation mode. |
| IIS 6.0 Operations Guide | Configuring Applications to Use COM+ Services | Clarified that some configurations of launch permissions may prohibit COM+ objects from being created, which may affect the running of the app/extension/etc. in which the object resides. |
| IIS 6.0 Operations Guide | Configuring Applications to Use COM+ Services | Corrected list of Windows Components that support Side-by-Side assemblies. |
| IIS 6.0 Operations Guide | Configuring Applications to Use COM+ Services | Added note that batch Logon is required for COM+ server-activated application identities. |
| IIS 6.0 Operations Guide | Installing ISAPI Filters | Added notes for the following issues:<br>• In IIS 6.0, an Internet Server API (ISAPI) filter added at the site level will not load until the Web server has served at least one request to that site, so the filter's status won't display anything until that time.<br>• At the Web site level, you will not see any global filters inherited from the Web server's master properties. |
| IIS 6.0 Operations Guide | Recycling Worker Processes with IIS 6.0 and subordinate topics | Added information about worker process recycling including how it works, when to use it, how to configure it, and logging the recycling events. |

**Architecture**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Technical Reference | Evaluating Application Isolation Modes | Added information about how SF_NOTIFY_READ_RAW_DATA notifications work depending on the application mode in which the server is running. |
| IIS 6.0 Technical Reference | HTTP Protocol Stack | Added a note about improvements made to HTTP.sys for Windows Server 2003 SP1 and a link to the HTTP API in Windows Server 2003 SP1 white paper. |

**Error Messages**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Operations Guide | Configuring Custom Error Messages | Added note that the following errors are not customizable: 400, 403.9, 411, 414, 500, 500.11, 500.14, 500.15, 501, 503, 505. |
| IIS 6.0 Operations Guide | Error Messages in IIS 6.0 | Added note that HTTP.sys does not serve custom errors, so standard HTTP 1.1 error messages will be returned to the browser if a request does not reach user mode. |

**Installation**

| Content Type | Topic | Description |
|---|---|---|

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Installing IIS | Added link to information about answer files and unattended setup parameters in the IIS 6.0 Technical Reference. |
| IIS 6.0 Technical Reference | Creating an Answer File | Documented the following unattended setup parameters:<br>• *InetPrint*<br>• *SvcManualStart*<br>• *ExtensionFile* |

**Introduction/Overview**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Deployment Guide | IIS 6.0 Benefits and Features | Added information about HTTP compression. |
| IIS 6.0 Deployment Guide | Appendix C: Choosing a Cluster Technology for IIS 6.0 and subordinate topics | Added topics with information about the different clustering technologies available for IIS 6.0, and included links for additional information. |
| IIS 6.0 Operations Guide | What's Changed | Clarified the meaning of "extension" in the context of the Known Extensions feature of IIS 6.0, and added information about how to lock down DLLs, executables, and file name extensions. |
| IIS 6.0 Operations Guide | What's Changed<br>Does ASP debugging work in IIS 6.0 worker process isolation mode?<br>Important Changes in ASP | Added information about configuring Launch and Access permissions to enable ASP debugging for Script Debugger and Visual InterDev when using IIS 6.0 Worker Process Isolation mode. |

**Logging**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Operations Guide | Enabling W3C Centralized Logging | Added topic describing how to enable W3C centralized logging. |
| IIS 6.0 Operations Guide | Auditing the Metabase | Added new topic with information about auditing in the IIS metabase. |
| IIS 6.0 Operations Guide | Auditing in IIS 6.0 | Added link to information about disabling auditing of successful logons in the security event log. |
| IIS 6.0 Technical Reference | W3C Centralized Logging | Added topic about the new W3C centralized logging feature. |
| IIS 6.0 Technical Reference | Metabase Auditing | Added new topic with information about auditing in the IIS metabase. |

**Performance and Scalability**

| Content Type | Topic | Description |
|---|---|---|

| | | |
|---|---|---|
| IIS 6.0 Deployment Guide | Enabling HTTP Compression IIS 6.0<br><br>Configure Compression | Added topics with information about HTTP compression. |
| IIS 6.0 Operations Guide | Using HTTP Compression<br><br>Enabling HTTP Compression<br><br>Customizing the File Types IIS Compresses | Added topics with information about HTTP compression. |
| IIS 6.0 Operations Guide | Throttling Bandwidth | Clarified that the value in the **Maximum bandwidth (in kilobytes per second)** box must be at least 1024 Kbps. |
| IIS 6.0 Technical Reference | Events and Conditions That Disable HTTP.sys Response Caching | Added information about uncached responses when accessing a Web site by requesting the root folder with no specific file, and about Web sites with IP Restriction lists. |
| IIS 6.0 Technical Reference | Using HTTP Compression for Faster Downloads | Added information about configuring HTTP compression for individual sites and site elements. |

**Reference**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Operations Guide | NTAuthenticationProviders Metabase Property | Added information about the state of the Integrated Windows Authentication setting following installation of Windows Server 2003 Service Pack 1 (SP1) as part of a slipstream installation of a Windows Server 2003 operating system. |
| IIS 6.0 Operations Guide | MD_ETAG_CHANGENUMBER Metabase Property | Added information about setting this property in an environment with clustered Web servers. |
| IIS 6.0 Operations Guide | AppPoolQueueLength Metabase Property | Added note that the default value of this property changed with Windows Server 2003 Service Pack 1 (SP1) from 4000 to 1000, with a recommendation to lower the value if IIS is already installed on the server prior to applying the service pack. |
| IIS 6.0 Operations Guide | FTP Service Events | Added new FTP Service event message 14401: "The server failed to create a data channel connection to a client on local interface *ip_address*. The FTP data port *port_number* may be in use by another service or application." |
| IIS 6.0 Operations Guide | ConnectionTimeout Metabase Property | Added information about the reasons a connection can be considered inactive and disconnected by IIS. |
| IIS 6.0 Operations Guide | LogEventOnRecycle Metabase Property | Clarified that the flags for this property only determine whether or not IIS will log the type of recycle event in the flag. |
| IIS 6.0 Operations Guide | AspSxsName Metabase Property<br><br>AspAppServiceFlags Metabase Property | Corrected list of Windows Components that support Side-by-Side assemblies. |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | ASP Registry Entries | Created the topic ASP Registry Entries and documented the following ASP registry keys:<br>• AspMaxPropertyStringLength<br>• AspMaxResponseHeaderLength<br>• DisableCachedResponseOnUNCAccessFailure<br>• DisableComPlusCpuMetric<br>• DisableOOMRecycle<br>• EnableChangeNotificationForUNC<br>• F5AttackDetectionEnabled<br>• FileMonitoringEnabled<br>• FileMonitoringTimeoutSeconds<br>• HangDetConsecIllStatesThreshold<br>• HangDetEnabled<br>• HangDetRequestThreshold<br>• HangDetThreadHungThreshold<br>• MaxCPU<br>• MaxCSR<br>• ThreadMax |
| IIS 6.0 Operations Guide | MaxProcesses Metabase Property | Removed erroneous information about setting the value of this property to zero in order to create an unmanaged application pool. |
| IIS 6.0 Operations Guide | FTP Service Registry Entries | Documented the following FTP service registry keys:<br>• AnonymousHomeDirInfoGrace<br>• AnonymousHomeDirInfoInterval<br>• CacheAnonymousHomeDirInfo |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Global Registry Entries | Documented the following global registry keys:<br>• DataSetCacheSize<br>• DoDirMonitoringForUNC<br>• ErrorLoggingDir<br>• FileAttributeCheckThreshold<br>• ThreadPoolMaxCPU<br>• ThreadPoolStartDelay<br>• ThreadPoolStartupThreadCount<br>• ThreadPoolMaxContextSwitch<br>• ThreadPoolRefTraceCounter<br>• ThreadPoolExactThreadCount |
| IIS 6.0 Operations Guide | WWW Service Registry Entries | Documented the following World Wide Web Publishing Service (WWW Service) registry keys:<br>• EnableTraceMethod<br>• LogErrorRequests<br>• MaxConcurrentCgisExecuting<br>• UseSharedWPDesktop |
| IIS 6.0 Operations Guide | DisableSocketPooling Metabase Property | Clarified that setting this metabase property has no effect in IIS 6.0, and included link to information about how to use the new utility, Httpcfg.exe, to allow applications to use the same port on different IP addresses. |
| IIS 6.0 Operations Guide | ServerListenTimeout Metabase Property<br><br>ConnectionTimeout Metabase Property<br><br>HeaderWaitTimeout Metabase Property<br><br>MinFileBytesPerSec Metabase Property | Clarified that **ServerListenTimeout** metabase property is not used by the WWW service, but is used by the FTP, SMTP, and NNTP services.<br><br>Added information about using the **ConnectionTimeout**, **HeaderWaitTimeout**, and **MinFileBytesPerSec** metabase properties instead of **ServerListenTimeout** for Web service nodes. |
| IIS 6.0 Operations Guide | SMPProcessorAffinityMask Metabase Property | Added note to warn that setting the **SMPProcessorAffinityMask** property to zero can create an error condition. |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Metabase Property Reference | Documented the following metabase properties:<br>• **ServerID**<br>• **LogInUTF8**<br>• **IsapiRestrictionList**<br>• **CgiRestrictionList**<br>• **RestrictionListCustomDesc**<br>• **SuppressDefaultFTPBanner**<br>• **WebDAVMaxAttributesPerElement** |
| IIS 6.0 Operations Guide | Metabase Property Reference | Corrected documentation for the following metabase properties:<br>• **DoDynamicCompression**<br>• **DoStaticCompression**<br>• **HcDoDynamicCompression**<br>• **HcDoStaticCompression**<br>• **IdleTimeout**<br>• **NotDeletable**<br>• **PeriodicRestartTime**<br>• **RevocationURLRetrievalTimeout**<br>• **RevocationFreshnessTime** |
| IIS 6.0 Operations Guide | NotDeletable Metabase Property | Clarified that the **NotDeletable** metabase property is invalid in IIS 6.0. |
| IIS 6.0 Technical Reference | Global Registry Entries | Corrected the following registry key references:<br>• DigestContextCacheTTL<br>• DigestPartialContextCacheTTL<br>• MaxConcurrency<br>• MaxPoolThreads<br>• PoolThreadLimit<br>• ThreadTimeout |

**Security**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Deployment Guide | Reviewing Security Policies, Processes, and Procedures | Corrected Table 3.16, "Firewall and Router Security," to refer to port 445 instead of 455 for NetBIOS/SMB. |
| IIS 6.0 Operations Guide | Using Isolation to Secure Applications | Added topic with information about giving each application pool its own account identity to isolate applications on the Web server. |
| IIS 6.0 Operations Guide | Configuring Subauthentication and subordinate topics | Corrected references to the **AnonymousPasswordSync** metabase property. |
| IIS 6.0 Operations Guide | Encrypting Data When Running WMI–Based Remote Administration Scripts | Added important note that Windows Management Instrumentation (WMI) based scripts must be edited to encrypt authentication credentials and data after upgrading to SP1. |
| IIS 6.0 Operations Guide | Configuring Constrained Delegation for Kerberos | Added information about Kerberos authentication failing with a 401.3 error if not configured properly. |
| IIS 6.0 Operations Guide | Integrated Windows Authentication in IIS 6.0<br><br>Configuring Integrated Windows Authentication in IIS 6.0<br><br>Windows Authentication Provider in ASP.NET<br><br>Authentication in ASP.NET | Added information about the state of the **Integrated Windows Authentication** setting following installation of Windows Server 2003 Service Pack 1 (SP1) as part of a slipstream installation of a Windows Server 2003 operating system. |
| IIS 6.0 Operations Guide | Securing Sites with IP Address Restrictions | Added note about the disabling of the HTTP.sys cache when IP address restriction is used on a Web site. |
| IIS 6.0 Operations Guide | Enabling Kernel-Mode SSL | Added new topic about kernel-mode Secure Sockets Layer (SSL). |
| IIS 6.0 Operations Guide | Basic Authentication in IIS 6.0 | Added information about administering users and adding them to groups with specific permissions when using Basic authentication. |
| IIS 6.0 Operations Guide | Requesting a New Server Certificate from an Online CA<br><br>Configuring SSL on a Web Server or Web Site | Clarified the following issues:<br>• The Web Server Certificate Wizard is used to request new server certificates.<br>• The request file generated by the Web Server Certificate Wizard is saved by default as C:\Certreq.txt.<br>• SSL can be configured for a Web server or a Web site. |
| IIS 6.0 Operations Guide | Installing Server Certificates | Added note that a password must be specified before the certificate can be assigned to a Web server. |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Configuring SSL on a Web Server or Web Site | Added note explaining that users will receive an error message when browsing to a site that does not have SSL completely setup and the **Require Secure Channel (SSL)** check box is selected for the site. |
| IIS 6.0 Operations Guide | IIS 6.0 Security Best Practices | Added recommendation to use an encrypted remote connection with WMI-based administration scripts. |
| IIS 6.0 Operations Guide | Mapping Client Certificates with Directory Service Mapping | Added procedure for enabling directory sService mapping. |
| IIS 6.0 Operations Guide | Encrypting Data When Running WMI–Based Remote Administration Scripts | Added topic about establishing an encrypted connection when running WMI-based administration scripts on an IIS server over a remote connection. |
| IIS 6.0 Operations Guide | Configuring Subauthentication and subordinate topics | Corrected references to **AnonymousPasswordSync** metabase property and added IIsWebService node as recommended configurable location. |
| IIS 6.0 Operations Guide | Anonymous Authentication in IIS 6.0 | Clarified that the IUSR_*Computername* account, which is a member of the **Guests** group on computers running IIS, has the same access as members of the **Users** group by default. |
| IIS 6.0 Operations Guide | Configuring SSL on a Web Server or Web Site | Added note that users will receive an error when browsing to a Web site that requires SSL if the process of setting up SSL is not complete. |
| IIS 6.0 Operations Guide | Enabling and Disabling Dynamic Content in IIS 6.0 | Added reminder to disable Web service extensions when features are no longer required. |
| IIS 6.0 Operations Guide | Enabling Web Service Extensions<br><br>Enabling and Disabling Dynamic Content in IIS 6.0 | Added important note that allowing unknown ISAPI extensions or CGI extensions to run on a Web server can leave the server susceptible to attack. |
| IIS 6.0 Operations Guide | Configuring Subauthentication on a New Installation of IIS 6.0<br><br>Configuring Subauthentication in IIS 5.0 Isolation Mode | Added node at which the **AnonymousPasswordSync** metabase property should be set. |
| IIS 6.0 Operations Guide | Creating an Authorization Policy Store in Authorization Manager<br><br>Setting the Scope for the Application | Clarified items in the procedure that must be typed exactly as shown in order for the authorization policy store for IIS 6.0 URL Authorization to be configured successfully. |
| IIS 6.0 Operations Guide | Setting Metabase Properties for URL Authorization | Removed semi-colons from sample script. |
| IIS 6.0 Operations Guide | Adding Urlauth.dll as a New Web Service Extension | Added procedure for adding Urlauth.dll as a new Web service extension. |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Adding the IIS Worker Process to the Readers Role | Added procedure for adding the IIS worker process to the **Reader** role. |
| IIS 6.0 Operations Guide | Exporting a Client Certificate for One-to-One Mapping | Added note that client certificate mapping programs for IIS 6.0 and Windows Active directory services accept certificates formatted as either binary or Base-64 encoded. |
| IIS 6.0 Operations Guide | Setting NTFS Permissions for Directories or Files and subordinate topics | Clarified that IIS Manager is used to set permissions on Web sites, directories, and virtual directories, and that to set permissions for individual files you must use a programmatic solution, such as Adsutil.vbs. |
| IIS 6.0 Operations Guide | Advanced Digest Authentication in IIS 6.0<br><br>Digest Authentication in IIS 6.0<br><br>UseDigestSSP Metabase Property | Added note that the WWW service must be restarted for a change in the **UseDigestSSP** metabase property to take effect. |
| IIS 6.0 Operations Guide | Configuring SSL Host Headers<br><br>Obtaining and Installing a Wildcard Server Certificate<br><br>Configuring Server Bindings for SSL Host Headers<br><br>Advanced Digest Authentication | Added new topics that explain how to obtain and install a wildcard server certificate that can be used for Web sites that have the same IP address:port binding, as well as how to use Adsutil.vbs to configure SSL host headers and how to ensure that secure content is not served for an HTTP request. |
| IIS 6.0 Technical Reference | Advanced Digest Authentication | Added information about configuring the registry keys DigestPartialContextCacheTTL and DigestContextCacheTTL. |
| IIS 6.0 Technical Reference | .NET Passport Authentication | Added information about configuring Web sites to conform to .NET Passport privacy guidelines. |
| IIS 6.0 Technical Reference | Digest Authentication | Added information about the issue where logging on with a user principal name (UPN) in the form *username@domain.top-leveldomain* fails. |
| IIS 6.0 Technical Reference | Integrated Windows Authentication | Added information about the state of the **Integrated Windows Authentication** setting following installation of Windows Server 2003 Service Pack 1 (SP1) as part of a slipstream installation of a Windows Server 2003 operating system. |
| IIS 6.0 Technical Reference | Authentication Methods Supported in IIS 6.0 | Added information about the importance of using authentication with IIS when resources on the Web server need to be restricted to approved users, and clarified that IIS goes in order from the most secure to the least secure authentication method when multiple authentication methods are configured. |
| IIS 6.0 Technical Reference | Integrated Windows Authentication | Added information that a user logged on to the local computer as a domain user will not be prompted to authenticate if they access a network computer in that same domain. |
| IIS 6.0 Technical Reference | Controlling Access to a Database on a Web Server | Added topic with information about controlling access to databases from Web applications. |
| IIS 6.0 Technical Reference | Access Control Lists | Added three tables outlining the default Access Control Lists (ACLs) and privileges for the **IIS_WPG** group. |

| | | |
|---|---|---|
| IIS 6.0 Technical Reference | Kernel-Mode SSL | Added topic about running SSL in kernel mode instead of user mode (the default) in Windows 2003 Server SP1. |
| IIS 6.0 Technical Reference | Configuring Application Pool Identity in IIS 6.0 | Added section about the **IIS_WPG** group. |
| IIS 6.0 Technical Reference | SSL and SF_NOTIFY_READ_RAW_DATA Filters | Added information about how SF_NOTIFY_READ_RAW_DATA works with SSL when running worker process isolation mode. |
| IIS 6.0 Technical Reference | Using URL Authorization | Removed semi-colons from sample script. |
| IIS 6.0 Technical Reference | Digest Authentication<br><br>Advanced Digest Authentication | Added note that the WWW service must be restarted for a change in the **UseDigestSSP** metabase property to take effect. |

**Server Administration**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Operations Guide | Configuring Launch Permissions for COM+ Server Objects | Added important note about the IUSR_*computername* and IWAM_*computername* accounts no longer being granted launch permissions to COM+ objects by default after installations of IIS on Windows Server 2003 SP1. |
| IIS 6.0 Operations Guide | Implementing Incremental Web Site Numbering | Added new topic with information about incremental Web site numbering. |
| IIS 6.0 Operations Guide | Renaming a Web Service Extension | Added topic with information about renaming Web Service extensions. |
| IIS 6.0 Operations Guide | Administering Servers Remotely in IIS 6.0 | Mde references to the Remote Administration (HTTP) Tool consistent. |
| IIS 6.0 Operations Guide | Uninstalling IIS | Added topic with procedure on uninstalling IIS. |
| IIS 6.0 Operations Guide | Configuring Application Pool Identity with IIS 6.0 | Added JScript code example to show how to programmatically change the account under which an application pool runs. |
| IIS 6.0 Operations Guide | Setting Active Directory User Isolation Using lisftp.vbs | Changed references to the following parameters:<br>From *FTPDir\|FTPRoot* to *msIIS-FTPDir\|msIIS-FTPRoot*. |
| IIS 6.0 Operations Guide | Reliability Features<br><br>Configuring Web Gardens with IIS 6.0 | Clarified information about how web gardens benefit reliability and performance. |
| IIS 6.0 Operations Guide | Listing IIS Backup Configurations Using lisback.vbs | Added note that **iisback /list** uses Coordinated Universal Time (UTC) for date and time stamps. |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Enabling ASP.NET | Added note to clarify the difference between using the Configure Your Server wizard or Add or Remove Programs to install IIS 6.0 and enable ASP.NET. |
| IIS 6.0 Operations Guide | Configuring Application Pool Identity with IIS 6.0 | Added the following procedures:<br>• To change the account under which an application pool runs using the Adsutil.vbs administration script<br>• To programmatically change the account under which an application pool runs |
| IIS 6.0 Operations Guide | Enabling and Disabling Dynamic Content in IIS 6.0 | Added information about needing to enable FrontPage Server Extensions and ASP in order for Visual Studio .NET 2003 applications to work on IIS 6.0. |
| IIS 6.0 Operations Guide | Setting Up Virtual Servers for Clustering | Corrected the procedure step explaining how to copy IIS configuration from one node to other nodes in a server cluster. |
| IIS 6.0 Technical Reference | Configuring Rapid-Fail Protection | Updated information about Rapid-Fail Protection. |
| IIS 6.0 Technical Reference | Creating Application Pools | Removed erroneous reference to an unmanaged application pool. |

**Site Administration**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Deployment Guide | Creating Web Sites and Home Directories<br>Create a Web Site | Added information about storing Web site content in a location where "/bin" will not be part of the requested URL or a 404 error will be returned. |
| IIS 6.0 Operations Guide | Creating Global MIME Types | Added information about creating Multipurpose Internet Mail Exchange (MIME) types for files without extensions, and corrected steps in the procedure. |
| IIS 6.0 Operations Guide | Creating MIME Types for a Web site or directory | Added information about creating MIME types for files without extensions. |
| IIS 6.0 Operations Guide | Editing Global MIME Types<br>Editing MIME Types for a Web site or directory<br>Removing Global MIME Types | Created topics with procedures to edit and remove MIME types. |
| IIS 6.0 Operations Guide | Removing Global MIME Types<br>Removing MIME Types from a Web site or directory | Added information about removing MIME types from the registry at HKEY_CLASSES_ROOT\extension after removing the MIME type mapping. |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Renaming Virtual Directories | Added new topic about renaming virtual directories. |
| IIS 6.0 Operations Guide | Setting FTP Messages and Directory Output Style | Added procedure for suppressing the default FTP banner using the **SuppressDefaultFTPBanner** metabase property. |
| IIS 6.0 Operations Guide | Redirecting FTP Requests to Directories or Network Shares | Added information and procedure about configuring an FTP site so that users browsing the site can see a directory entry for a virtual directory. |
| IIS 6.0 Operations Guide | Creating FTP Sites Using lisftp.vbs | Added an example and changed references to the following parameters:<br>• From **/domain** *DomainName* to **/ADdomain** *DomainName*.<br>• From **/Admin** [*Domain*\]*User* to **/ADadmin** [*Domain*\]*User*.<br>• From **/AdminPwd** *Password* to **/ADpass** *Password*. |
| IIS 6.0 Operations Guide | Using FrontPage Server Extensions to Extend Web Sites | Clarified that the _vti_bin virtual directory of new FrontPage Server Extensions virtual servers must be assigned to the MSSharePointPool application pool in order to remotely administer a Windows® SharePoint® Services (STS) Web site. |
| IIS 6.0 Operations Guide | Working with MIME Types and subordinate topics | Added important note about not using wildcard MIME types on production servers, and revised wording in procedures about providing valid MIME types. |
| IIS 6.0 Operations Guide | Setting FTP Messages and Directory Output Style | Added note that although the FTP exit message box can contain up to 30,000 characters, the message size is limited to 1024 characters (including the reply code, blank, carriage return, and null). |
| IIS 6.0 Operations Guide | Hosting Multiple Web Sites<br>Naming Web Sites | Added note that host headers for Web sites cannot contain Unicode characters because the Domain Name System (DNS) standard doesn't recognize Unicode characters. |
| IIS 6.0 Technical Reference | Metabase Structure in IIS 6.0 | Added important note about enabling sequential numbering of Web sites in IIS 6.0. |

**Troubleshooting**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Operations Guide | FTP server terminates user connections and returns MSFTPSVC event message 8 | Added topic with information about hosting a high-load FTP site with the FTP anonymous home directory on a UNC share, and receiving FTP service message 8. |
| IIS 6.0 Operations Guide | Cannot launch COM objects with IUSR_computername or IWAM_computername accounts | Added topic with information about the IUSR_*computername* and IWAM_*computername* accounts no longer being granted launch permissions by default on installations of IIS on Windows Server 2003 SP1. |
| IIS 6.0 Operations Guide | UPN logons not accepted by IIS 6.0 server in Windows 2000 domain | Added topic with information about the correct format for logins when the IIS 6.0 server is a member of a Windows 2000 domain, and is configured for both Digest authentication and subauthentication. |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Install certificate after deleting the pending certificate request | Added topic with information about how to install a certificate if the pending certificate request was deleted from the server. |
| IIS 6.0 Operations Guide | CGI Processes Will Not Start | Added information about the default anonymous user account, IUSR_*computername*, being denied access to executable programs in the *windir*\System32 folder and how to resolve the problem if failures occur when attempting to call external programs from CGI code. |
| IIS 6.0 Operations Guide | Cannot make immediate request for a Web server certificate | Added topic with information about not being able to request a Web server certificate immediately after installing Certificate Services on a server and configuring it to be either an Enterprise Root certification authority (CA) or a stand-alone certification authority. |
| IIS 6.0 Operations Guide | Client cannot renegotiate request and returns an HTTP 413 error | Clarified that clients may receive an HTTP 413 error when uploading large files if client certificates are enabled on a Web server, Web site, or on individual directories or files on the site. |
| IIS 6.0 Operations Guide | Configuring Custom Error Messages | Added information about storing a custom error URL in the application pool that directs the request to the custom error URL so that the worker process can serve the custom error page successfully. |
| IIS 6.0 Operations Guide | Clients browsing to a Web site receive HTTP 403.4 - Forbidden: SSL is required to view this resource error, but the Web site is not configured to use SSL | Added information about troubleshooting an issue where clients receive an error message about requiring SSL when the Web site is not required to use SSL. |
| IIS 6.0 Operations Guide | Client requests receive a "503: Service Unavailable" error | Updated information about the "503: Service Unavailable" error message, which occurs when the server is unable to process a request, and linked to a topic with more information about increasing the application pool queue length limit when the Httperr.log file includes the 503 error code along with the message "QueueFull". |
| IIS 6.0 Operations Guide | When adding a new Web site, the Web site is not added, and the error message "No more memory is available for security information updates" is displayed (also, error 1131 might appear in the system Event Log) | Added information about issue and steps to correct the problem. |
| IIS 6.0 Operations Guide | All requests with /bin in the URL are rejected and return a 404 error | Added information about issue and steps to correct the problem. |
| IIS 6.0 Operations Guide | Client requests time out or error out | Clarified that **ServerListenTimeout** metabase property is not used by the WWW service, but is used by the FTP, SMTP, and NNTP services. |

| | | |
|---|---|---|
| IIS 6.0 Operations Guide | Requests for static files return 404 error | Added content outlining the following options for serving files with unknown file name extensions:<br>• Add the file name extension to the MIME map on specific Web sites or directories.<br>• Turn off the Known Extensions feature on the specific Web sites or directories.<br>• Turn off the Known Extensions feature on the Web server.<br>• Added important note that turning off Known Extensions on the Web server lessens the security of the server. |
| IIS 6.0 Operations Guide | When attempting to download a file from an FTP site that is hosted on IIS 6.0, a "File not found" or similar error is returned, and the file name in the error message contains a caret (^) character | Added information about issue and steps to correct the problem. |
| IIS 6.0 Operations Guide | Applications are denied access to resources | Added note that when IIS 6.0 worker processes run under the default low-privileged Network Service user account, application events, including CGI events, are written to the Application event log not the System event log. |
| IIS 6.0 Operations Guide | Client receives 403.16 error when IIS cannot process a complete certificate chain | Added information about adding a certificate to the Intermediate Certification Authorities folder on the local computer when the client receives a 403.16 error. |
| IIS 6.0 Operations Guide | Computer runs slowly when scripts using ADSI scripts to query IIS 6.0 | Added information about memory usage when running a program that uses multiple threads and Active Directory Service Interfaces (ADSI) scripts to query Internet Information Services (IIS). |
| IIS 6.0 Operations Guide | Advanced Troubleshooting with IIS Tracing and subordinate topics | Added multiple topics that describe how to configure and execute new IIS tracing features. |
| IIS 6.0 Operations Guide | Tracing Events Reference | Added new topic that describes tracing events available with IIS Admin tracing, IIS Request-based Tracing, and IIS Currently-executing request tracing. |

| | | |
|---|---|---|
| IIS 6.0 Technical Reference | Advanced IIS Troubleshooting with Tracing and Debugging and subordinate topics | Added multiple topics that address the new tracing features in IIS 6.0 SP1. These features include the following:<br>• Currently-executing requests tracing: allows tracing of all requests executing on the IIS server. This tracing feature is helpful when requests have become slow or unresponsive and it is unclear which Web site or application is processing the requests.<br>• Request-based tracing: allows tracing of requests for a specific site or for a specific component of IIS, such as ASP pages. This tracing feature helps narrow troubleshooting efforts to a specific site or IIS component.<br>• IIS Admin tracing: allows tracing of the IIS Admin service during startup and shutdown events. If the IIS Admin service hangs during startup and shutdown, this tracing feature can help determine why. |
| IIS 6.0 Technical Reference | Enabling Debugging for IIS | Replaced the script for OrphanActionExe. |

**Web Applications**

| Content Type | Topic | Description |
|---|---|---|
| IIS 6.0 Operations Guide | ASP.NET Settings Schema Reference | Linked topic to MSDN for information about the individual elements in the schema and removed the subordinate topics. |
| IIS 6.0 Operations Guide | Process Model Modes | Updated which **processModel** settings still apply in worker process isolation mode. |
| IIS 6.0 Operations Guide | What's New in ASP | Corrected list of Windows Components that support Side-by-Side assemblies. |

# IIS 6.0 Typographical Conventions

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The following graphics and text conventions are used throughout the IIS 6.0 Deployment Guide and the IIS 6.0 Technical Reference.

**Flowchart Symbols**

Use the following table of symbols as a resource for understanding the flowcharts included in the IIS 6.0 Deployment Guide.

| Symbol | Meaning | Symbol | Meaning |
|---|---|---|---|
| | Step or component process | | Data stored to a database |
| | Predefined process or subroutine | | Flowchart beginning or end |
| | Decision point | 1 | Intra-chart connector: Flow continues to next page |
| | Output to a document or input from a document | 1 | Intra-chart connector: Flow continues from previous page |
| | Data transfer to a file on disk | A | Inter-chart connector: Indicates an exit point to another flowchart |
| | Data transfer to a data store | B | Inter-chart connector: Indicates an entry point from another flowchart |

**Art Symbols**

Use the following table of the art symbols as a resource for understanding the graphics included in this guide.

| Symbol | Meaning | Symbol | Meaning |
|---|---|---|---|
| | Workstation | | Macintosh client |
| | Portable computer | | Tablet computer |
| | Terminal | | Cellular phone |
| | Portable digital assistant (PDA) | | Document |
| | File folder | | E-mail |
| | Chart | | Wireless network adapter |
| | Modem | | Video camera |
| | Network adapter | | Digital camera |
| | Facsimile | | Printer |
| | Telephone | | Scanner |
| | Hard disk | | Tape drive |
| | Database | | Tape |
| | Compact disc | | Security key |
| | Digital certificate | | Padlock |
| | Padlock | | Uninterruptible power supply |

| | | | |
|---|---|---|---|
| | Access token | | Hub |
| | Modem bank | | Automated library |
| | Windows NT–based server | | Generic server |
| | Mainframe computer | | Host |
| | Data jack | | Input/output (I/O) filter |
| | Shadowed router | | Windows 2000–based router |
| | Router | | Switch |
| | Server farm | | Clustered servers |
| | Interface | | Packets |
| | Transceiver | | Script |
| Internet | Internet | Intranet | An intranet |
| | Firewall | | Tunnel |
| | Process or communication failure | “ · ” | DNS root |
| | Directory tree root | | Root |
| | Organization | | Organizational unit |
| | Active Directory domain | | User group |
| | Common name | | Generic node |
| Windows 2000 Domain | Windows 2000 domain | | Site or Windows NT 4.0 domain |
| | Active Directory™ | | |

# Windows Server 2003 Terminal Server Capacity and Scaling Abstract

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft Windows Terminal Server lets users run Windows-based applications on a remote computer running one of the Windows Server 2003 operating systems. This white paper contains testing methodologies, results, analysis, and sizing guidelines for Windows Server 2003 Terminal Server. See Windows Server 2003 Terminal Server Capacity and Scaling.

# Microsoft IT SP1 Deployment Case Study Abstract

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This case study provides an inside look at how Microsoft Information Technology (Microsoft IT) deployed Microsoft® Windows Server™ 2003 Service Pack 1 (SP1) into the Microsoft global environment. This document describes Microsoft IT's high-level processes for deploying SP1 and the best practices they learned as a result. Microsoft IT minimized the cost of the deployment by planning carefully, creating a ramped rollout schedule, and using an installation sequencer to deploy a standardized software package. IT architects can use this case study to help plan their SP1 deployments.

To view this document, see the Microsoft IT SP1 Deployment Case Study.

# Mixed Workload Consolidation Guide Abstract

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This document provides detailed information you can use to identify appropriate applications for mixed workload consolidation, validate compatibility of applications, and plan for a mixed workload consolidation deployment. This document discusses mixed workload consolidation processes primarily as they pertain to Microsoft Windows Server 2003 Enterprise Edition and Microsoft Windows Server 2003 Data Center Edition, which are the preferred operating systems for mixed workload consolidation.

For information about this topic, see Mixed Workload Consolidation Guide.

# Domain Isolation Planning Guide for IT Managers

Updated: June 8, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

Plan your domain isolation using the step-by-step planning information in this guide. This guide is designed to help you work with your IT staff to gather the necessary information, design your IPsec polices, and to create a deployment plan.

This document is intended for IT professionals who are investigating using IPsec in Microsoft® Windows® to deploy domain isolation in their environments. This guide is designed to help you work with your IT staff to gather the necessary information, design your IPsec polices, and to create a deployment plan.

This guide includes an overview of the deployment process, a step-by-step guide to the planning process, and links to other resources that you can use in the planning and design process. The material in this guide covers the planning of the deployment but not the actual deployment process itself.

This guide provides material relevant only to Windows-based computers and is not intended as a guide for deploying domain isolation on operating systems other than Windows. This guide does not provide background information about IPsec and related technologies.

#### Terminology Used in This Guide

These terms are defined so that you can more clearly understand how they are used in this guide:

**Domain isolation** - The use of Internet Protocol security (IPsec) to require authentication, encryption, or both, among members of a Windows domain as well as between members of the domain and unknown or unauthorized computers.

**Isolated host** - A computer that is a member of an isolated domain.

**Non-isolated computer** - A computer that is not a member of an isolated domain.

**Isolated domain** - A domain or network that requires authentication, encryption, or both, by using IPsec before allowing any domain member to communicate with any other computer.

**Boundary computer** - A computer in the boundary group of an isolated domain that accepts secure communication from members of the isolated domain and unsecured communication from computers that are not members of the isolated domain.

**Proxy server** - A firewall component that manages Internet traffic to and from a local area network (LAN) and that can provide other features, such as document caching and access control. A proxy server can improve performance by supplying frequently requested data, such as a popular Web page, and it can filter and discard requests that the owner does not consider appropriate, such as requests for unauthorized access to proprietary files.

#### IT Roles Used in This Guide

The IT roles used in this guide are generalizations derived from the IT industry and should closely approximate those in your IT structure. Your roles might differ from these, but they are offered to help you assign tasks to your staff members.

| Role | Primary Responsibilities |
|---|---|
| IT Manager | Coordinates the deployment effort and communicates status to upper management. Manages the IT staff responsible for the infrastructure, desktop and server deployment, and server administration and operations. Evaluates the impact of the technology solution on core business and IT resources. Uses this guide to assign tasks to staff members, collect and collate information, guide the overall process of planning, and make the major decisions. |
| Systems Architect | Provides information and assists in IPsec policy design. Responsible for designing the overall server infrastructure for all systems. Develops server deployment strategies and policies. Contributes to networking connectivity design. Ensures that deployment policies are followed. Provides overall architectural guidance and assists in designing the isolated domains and policies for establishing these segments. |

| | |
|---|---|
| Security Manager | Provides information, assists in IPsec policy design, and assists IT Manager in planning. Responsible for determining and implementing server security policy. Evaluates new technology and its effects upon security and availability. Responsible for security issues and solutions. Is the primary source of information about what security levels are needed for segments, how segments might be designed around security needs, what current security policies and tools can be made obsolete by domain isolation, and what security concerns might arise in the design and implementation of domain isolation. |
| Windows Systems Administrator | Provides information and implements polices on Windows-based computers. Responsible for determining and implementing server security policy. Evaluates new technology and its effects upon security and availability. Responsible for security issues and solutions. Is the primary source of information about what security levels are needed for segments, how segments might be designed around security needs, what current security policies and tools can be made obsolete by domain isolation, and what security concerns might arise in the design and implementation of domain isolation. |
| UNIX Systems Administrator | Provides information and assists Helpdesk and Security Managers in planning. Responsible for configuration and administration of UNIX servers, including upgrades, backups, capacity monitoring, and planning and interoperability issues with Windows systems. Is the primary source of information about IPsec interoperability between UNIX and Windows. |
| Network Administrator | Provides information and assists in IPsec policy design and implementation. Responsible for overall connectivity for the entire network, including hardware. Manages connectivity between heterogeneous systems (Windows and UNIX). Troubleshoots all performance issues across the network. Provides significant input regarding how the network will influence the design of, or be affected by, IPsec policies. |
| Database Administrator | Provides information and assists in IPsec policy design. Also assists in the configuration and management of database solutions. Evaluates build images, deploys new databases and changes, and conducts server integration testing. Is the primary source of information about how domain isolation might affect database access and performance. |
| Desktop Configuration Manager | Provides information and assists in IPsec policy design and implementation. Responsible for provisioning desktop PCs and deploying service packs and updates to these PCs. Involved in setting the strategic direction for the desktop operating system and applications. Is the primary source of information about how IPsec polices might affect desktop configuration. |
| Helpdesk Manager | Provides information and creates user education and notification materials. Responsible for all Helpdesk operations. Is the primary contact for information about how users might be affected, or respond to, access issues during or after deployment, and how user education can help mitigate any issues that might arise. |

**Introduction to Deployment Planning Phases**

This section provides a brief overview of the different phases involved in the domain isolation planning process. This process is suggested as a way of making domain isolation deployment as effective and efficient as possible and to suggest how you can work with your IT team to gather the needed information, discuss domain isolation issues, create a deployment plan, design IPsec policies, and test/refine these polices to reduce any user and operations issues that might arise from domain isolation.

**Collect Information About Your IT Environment**

You and your team will gather information about network topology, security policy and implementation, server operating systems and applications, service level agreements (SLAs), user types, any interoperability issues or concerns, and regulations or other external constraints. This information will be used along with other information, such as IT polices and guidelines and any business needs, to determine what domain isolation needs you have and then to design the IPsec policies that will be used to fulfill these domain isolation needs.

**Determine Your Domain Isolation Needs**

You and your team will use the collected information and determine what kind of isolation needs you have, based upon business needs, regulatory influences, security requirements, Service Level Agreements, the IPsec technology, user needs, and other factors.

**Design Your IPsec Policies**

This is probably the most crucial phase and requires close attention to the details for designing IPsec filter lists, filter actions, rules, and policies for each segment. Carefully-designed policies will make the deployment process smooth and efficient, the isolation effective, the protection of your assets solid, and it will keep user problems to a minimum while still meeting all SLAs, regulatory requirements, and other criteria.

**Deploy the Policies in a Test Environment**

Your team can test the domain isolation deployment and discover any refinements that should be made to the IPsec policies and the deployment process before deploying to a large or business-critical segment. The test environment can be designed specifically for deployment testing, or it can be a small, non-business-critical domain environment.

**Refine Policies**
The test phase might highlight some connectivity, security, or administration issues that can be addressed by refinements to your IPsec policies, by adjustments in operations and administration, or by improvements in user education. This information is important for a smooth and effective deployment.

**Create a Deployment Schedule**
Once you have completed your plans and IPsec policies, your team can discuss and solidify how, when, and where you will implement domain isolation. This phase is where you can discuss any potential problems with the schedule and agree on the best plan for actual deployment.

**Prepare for User and Infrastructure Support**
Before deployment you can develop plans, documents, and tools to assist your Helpdesk staff to deal with pre-deployment notification and education, user issues during deployment, and any post-deployment issues that might arise as a result of domain isolation deployment.

**Inform Team Members About IPsec**
The planning and design of domain isolation will be more efficient and effective if your team has a good understanding of what IPsec is and how it can be used for domain isolation. The following table lists the type of IPsec information that might be useful before you begin the domain isolation planning phases.

| Action | Owner |
|---|---|
| Review IPsec concepts. | As needed |
| Review examples of domain isolation deployment. | IT Manager, Systems Architect, Security Manager |
| Review IPsec limitations. | IT Manager, Systems Architect, Security Manager, others as needed |
| Review IPsec interoperability information. | IT Manager, Systems Architect, Security Manager, others as needed |

**Phase 1: Collect Information About Your IT Environment**
**Collect Computer Information**
Because not all computers in your environment will implement IPsec the same way, you must document the operating systems and service pack versions on your computers, both servers and desktop computers.

| Action | Owner |
|---|---|
| Identify IPsec-compatible Windows-based computers (Microsoft Windows 2000, Microsoft Windows XP, and Microsoft Windows Server™ 2003). | Windows Systems Administrator |
| Identify IPsec-incompatible Windows-based computers (Microsoft Windows 98, Microsoft Windows Millennium Edition, and Microsoft Windows NT® 4.0). | Windows Systems Administrator, Desktop Configuration Manager |
| Identify IPsec-incompatible computers using operating systems other than Windows (UNIX, BSD UNIX, Macintosh OS X, Linux, IBM OS390, Sun Microsystems Solaris 8, etc.). | UNIX Systems Administrator, Desktop Configuration Manager |

**Collect Network Device Information**
Network devices, such as routers or firewalls, are also involved in the domain isolation, and their IPsec capabilities must be documented to ensure that deployment planning takes them into account.

| Action | Owner |
|---|---|
| Identify IPsec-compatible network devices (Cisco IOS, Lucent VPN Firewall Bricks, Lucent Access Point IP services routers, 3Com SuperStack 3 Firewall, Nortel VPN Gateways, etc.). | Network Administrator |
| Identify IPsec-incompatible network devices - these devices cannot allow IPsec-protected packets to traverse the network. | Network Administrator, Systems Architect, UNIX Systems Administrator |

**Collect Active Directory/Domain Information**
Your IPsec policy design will be strongly influenced by your domain structure. Document your domain structure to identify

organizational and business units that might require different levels of security and other information.

| Action | Owner |
|---|---|
| Document your domain structure. | Systems Architect |
| Document your organizational units (OUs). | Systems Architect, IT Manager |
| Document your global security groups. | Security Manager |

**Collect Regulation and Other External Constraint Information**

Your organization might be affected by regulations that require enhanced security, access to information, or other actions. By being aware of these impacts during the IPsec policy design, you can ensure compliance or plan for compliance issues.

| Action | Owner |
|---|---|
| Identify any government regulations that might affect planning and design. | IT Manager, Security Manager |
| Identify any business partner regulatory requirements that might affect planning and design. | IT Manager |
| Identify any international regulatory requirements that might affect planning and design. | IT Manager |
| Identify any company policies that might affect planning and design. | IT Manager |

**Collect Security Information**

Your IPsec policy design will be strongly influenced by your current security policies and practices, such as firewall configurations. IPsec policies will also affect and expand your security policies.

| Action | Owner |
|---|---|
| Document your firewall configurations. | Security Manager |
| Document logon scripts or policies that might affect IPsec policies. | Security Manager |
| Document your software update policies and processes. | Security Manager |
| Document your application deployment policies. | Security Manager |
| Document your security response plan. | Security Manager |

**Collect Service Level Agreement Information**

Your domain isolation planning process must be designed and deployed in a way that complies with your SLAs. By being aware of these SLAs during the planning and design process, you can ensure compliance or make plans for mitigating and dealing with any compliance problems.

| Action | Owner |
|---|---|
| Identify any internal/direct SLA requirements that might affect planning and design. | IT Manager, UNIX Systems Administrator, Database Administrator, Desktop Configuration Manager, Network Administrator |
| Identify any business partner SLA requirements that might affect planning and design. | IT Manager |
| Identify any international SLA requirements that might affect planning and design. | IT Manager |

**Collect User and Partner Connectivity Information**

You need to consider how domain isolation will affect user connectivity during and after the deployment. This information can guide you in designing and deploying IPsec policies or in training staff and informing users before and after deployment.

| Action | Owner |
|---|---|
| Identify IT administrative staff that must have uninterrupted connectivity. | IT Manager |

| | |
|---|---|
| Identify knowledge workers. | Desktop Configuration Manager |
| Identify Helpdesk staff members that require detailed knowledge of the deployment. | Helpdesk Manager |
| Identify internal or external partners that require uninterrupted connectivity. | IT Manager |
| Document any applications that require uninterrupted connectivity. | Desktop Configuration Manager, Windows Systems Administrator, Database Administrator |
| Identify security personnel that must monitor security during the deployment. | Security Manager |

**Collect Interoperability Information**

Understanding how your IPsec policy design and deployment might affect IPsec-incompatible computers and computers with other implementations of IPsec will help you to plan the security for these computers and to determine how they will connect to IPsec-incompatible computers.

| Action | Owner |
|---|---|
| Identify any servers that are incompatible with the Microsoft implementation of IPsec but must access resources on Windows servers. | UNIX Systems Administrator, Desktop Configuration Manager |
| Identify any Windows servers that are incompatible with servers configured for a non-Microsoft implementation of IPsec but must access resources on these servers. | Windows Systems Administrator |
| Identify any applications that might require IPsec policy exemptions for business reasons. | Desktop Configuration Manager, Windows Systems Administrator |

**Phase 2: Determine Your Domain Isolation Needs**

**Business Needs**

Your IPsec policies must also take into account your business needs and the financial impact that the deployment might have.

| Action | Owner |
|---|---|
| Document how knowledge workers, internal customers, and partners will be affected by domain isolation deployment and the business impact of this. Document how these effects will be mitigated. | Helpdesk Manager, Desktop Configuration Manager, Security Manager |
| Document how the operations infrastructure will be affected by domain isolation deployment and the business impact of this. | Security Manager, Systems Architect, IT Manager |

**Regulation Needs**

Government and international regulations might influence your IPsec policy designs. You should enumerate the regulations and restrictions you must comply with and take these into account during the design phase.

| Action | Owner |
|---|---|
| Document the regulations you must comply with. | IT Manager, Security Manager |
| Determine the time-frame you have to comply. | IT Manager |
| Determine how you can use domain isolation to comply with these regulations. | IT Manager, Security Manager, Systems Architect |

**Security Management Needs**

You might have sensitive data or servers that require additional security in your environment. You should enumerate these and take into account what levels and types of security you want domain isolation to provide.

| Action | Owner |
|---|---|

| | |
|---|---|
| Document any effects of deploying domain isolation upon current security polices. | Security Manager |
| Identify whether any additional security administration will be required. | Security Manager, Desktop Configuration Manager, IT Manager |
| Determine whether any of your current security technologies are incompatible with domain isolation. | Security Manager, Network Administrator, Systems Architect |
| Identify where you will need boundary computers. | Security Manager, Network Administrator, UNIX Systems Administrator |
| Determine and document what additional security methods will be used to protect boundary hosts. | Security Manager, Network Administrator, Systems Architect |

**Service Level Agreement Needs**

During the planning process, determine whether the deployment will negatively affect your SLAs.

| Action | Owner |
|---|---|
| Determine how domain isolation will affect deployment and administration of SLAs. | IT Manager |
| Document how any effects on deployment and administration of SLAs will be measured. | IT Manager |
| Document how any effects on deployment and administration of SLAs will be mitigated or corrected. | IT Manager, Security Manager |
| Determine how any effects upon SLAs will be communicated to the appropriate parties. | IT Manager, Helpdesk Manager |

**IPsec Technology Needs**

Some of the policy design decisions are based on the IPsec technology itself and how it secures traffic and data.

| Action | Owner |
|---|---|
| Document which of the four IPsec-negotiated security modes will be used, where they will be used, and why.<br><br>The four modes are:<br><br>• **Request Mode**. A host responds to both IPsec and unauthenticated (non-IPsec) requests. It initiates communications with IPsec and, if that fails, allows unauthenticated communications.<br><br>• **Secure Request Mode**. A host responds to requests secured by IPsec and ignores unauthenticated requests. It initiates communications with IPsec and, if that fails, returns to unauthenticated communication.<br><br>• **Secure Require Mode**. A host requires IPsec-secured communications for both incoming and outgoing requests.<br><br>• **Default Response**. A host responds to IPsec requests, but never initiates IPsec. | Security Manager |
| Identify where IPsec tunnel mode will be needed. | Security Manager, Network Administrator |
| Identify where data integrity using Authenticated Header (AH) will be needed. | Security Manager, Systems Architect |
| Identify where data integrity and encryption (using ESP) will be needed. | Security Manager, Network Administrator |

| | |
|---|---|
| Document which forms of encryption will be used and where. | Security Manager |
| Document which forms of authentication will be used and where. | Security Manager |
| Identify which infrastructure applications or servers, such as DNS and DCHP servers, will need to be added to the default exemptions so that all clients can access them. | Security Manager, Systems Architect |
| Identify which ports/protocols will need to be opened in firewalls for IPsec. | Security Manager, Network Administrator |
| Identify where no IPsec protection will be needed. | Security Manager, Network Administrator, UNIX Systems Administrator, Desktop Configuration Manager |

**User and Partner Needs**

Your domain isolation design and deployment plans should also take into account how the process might affect the ability of users and partners to access information stored on your network.

| Action | Owner |
|---|---|
| Determine how users might be affected by the deployment. | Helpdesk Manager |
| Determine what user education steps can be taken to prepare users for the deployment and any possible issues resulting from it. | Helpdesk Manager |
| Determine how partners might be affected by the deployment. | IT Manager |
| Determine what steps can be taken to prepare partners for the deployment and any possible issues resulting from it. | IT Manager, Helpdesk Manager |

**Interoperability Needs**

If your environment includes computers that either cannot implement IPsec or whose implementation of IPsec is not the same as the Microsoft implementation, then you need to determine how, or even if, you will allow these computers to communicate with IPsec-incompatible computers.

| Action | Owner |
|---|---|
| Determine how IPsec-incompatible servers and Macintosh clients will communicate. | UNIX Systems Administrator, Security Manager |
| Determine how IPsec-compatible computers will communicate with IPsec-incompatible Windows clients. | Desktop Configuration Manager, Security Manager |
| Document how any effects of denied communications will be mitigated or corrected. | UNIX Systems Administrator, Desktop Configuration Manager, Security Manager |
| Determine which Windows services cannot be used with higher levels of IPsec protection. | Security Manager, Windows Systems Administrator |
| Determine whether there are any current IPsec policies (local or global) that might conflict with ones being designed. | Security Manager |

**Phase 3: Design Your IPsec Policies**

**Review IPsec Policy Design Documentation**

| Action | Owner |
|---|---|
| Designing IPSec Policies | Security Manager, Systems Architect |
| Understanding Default IPSec Policies | Security Manager, Systems Architect |
| Determining Your IPSec Needs | Security Manager, Systems Architect |

| | |
|---|---|
| Special IPSec considerations | Security Manager, Systems Architect |
| Weighing IPSec Tradeoffs | Security Manager, Systems Architect |
| Establishing an IPSec security plan | Security Manager |

**Create a Naming Convention**

You can create a naming convention for policies, filter lists, and filter actions. A naming convention can make backing up, restoring, and managing changes to policies, much easier.

Policy names should include the isolated domain and the date issued, for example, "Accounting_3.28.2005." Filter list names should describe the type of network traffic they match, for example, "All ICMP Traffic." Filter action names should describe the level of security they provide and the type of negotiation they use, for example, "Request Security."

| Action | Owner |
|---|---|
| Determine a naming convention for policies. | Security Manager, Systems Architect |
| Determine a naming convention for filter lists. | Security Manager, Systems Architect |
| Determine a naming convention for filter actions. | Security Manager, Systems Architect |

**Create an IPsec Policy Management Process**

A policy management process can reduce confusion and make backing up, restoring, and managing changes to policies much easier.

If you plan to create policies for a Windows Server 2003-based computer, you should keep in mind that Windows Server 2003 incorporates some new features that are not available in Windows XP or Windows 2000. For more information, see "IPsec Policy Compatibility Considerations" in Using Microsoft Windows IPsec to Help Secure an Internal Corporate Network Server.

| Action | Owner |
|---|---|
| Document how policies will be created. | Systems Architect |
| Document where policies will be backed up. | Systems Architect |
| Document how policies will be backed up. | Systems Architect |
| Document how policies will be changed, adjusted, and deployed. | Systems Architect |
| Determine how changes and backups will be documented. | Systems Architect |
| Determine how policies will be secured. | Security Manager |

**Review Example IPsec Policies**

Microsoft provides default IPsec policies with Windows, and also in the Windows Server 2003 Security Guide, that you might be able to use, with modifications, in your environment. These can also be used as examples in your policy design process.

| Action | Owner |
|---|---|
| Review the default policies. | Security Manager |
| Determine whether any of these policies can be used with modifications for your environment. | Security Manager |
| Review the policies provided in the Windows Server 2003 Security Guide. | Security Manager |
| Determine whether any of these policies can be used with modifications for your environment. | Security Manager |

**Design IPsec Filter Actions**

If the intranet traffic is to be secured, filter actions specify an ordered set of security methods (such as which integrity and encryption methods are used) and other settings. A combination of a filter action and a filter list make up a rule in an IPsec policy.

Try to create the fewest filter actions that meet your needs. For example, Microsoft IT was able to use only three for its environment. Your environment might be more diverse and might require more.

| Action | Owner |
| --- | --- |
| Determine where encryption of packet contents is required. | Security Manager |
| If encryption is used, determine which encryption algorithm is appropriate. | Security Manager |
| If encryption is not being used, determine whether packet integrity (signing) is necessary or desired. | Security Manager |
| Determine whether the correct action is to block, permit, or negotiate security. | Security Manager, Systems Architect |
| Determine whether the negotiation allows unsecured connections with IPsec-incompatible computers or IPsec-compatible computers with which IPsec negotiations fail. | Security Manager, Systems Architect |

**Design IPsec Filters and Filter Lists**

IP filters define matching criteria for a computer or a group of computers by specifying source and destination IP addresses, IP protocols, and source/destination TCP or UDP ports. Filter lists are a collection of one or more IP filters that logically belong together as a unit and that should have only one filter action associated with them. A rule combines a filter list with a filter action.

As a best practice use the **Any IP address** setting rather than the **My IP address** setting to mitigate problems with DHCP changing IP addresses. Additional best practices are provided in Improving Security with Domain Isolation.

| Action | Owner |
| --- | --- |
| Determine which IP addresses or subnets should be included in the filter lists. | Systems Architect, Security Manager |
| Determine which protocol/port combinations belong in the filter lists. | Systems Architect, Security Manager |
| Determine where the filters within the filter lists should and should not be mirrored. | Systems Architect, Security Manager |

**Design IPsec Policy Rules**

An IPsec policy rule combines a filter list with a filter action. If the filter action requires security, then the rule also specifies authentication methods, tunnel mode settings, and the types of interfaces to which this rule applies.

| Action | Owner |
| --- | --- |
| Pair the filter lists with the appropriate filter actions to define the set of rules for the IPsec policy. | Systems Architect, Security Manager |
| Determine whether the default response rule needs to be enabled or disabled. | Systems Architect, Security Manager |
| For rules requiring security, determine which authentication methods the rule uses to establish trust. | Systems Architect, Security Manager |

**Design IPsec Policies**

IPsec policies are a collection of one or more rules. Policies should group together all the rules that are appropriate for distribution to an organizational unit (OU), domain, or security group of the Active Directory® directory service.

Each segment will typically have an associated IPsec policy. The same policy might apply to many segments. Each computer can have only one policy assigned (active) at a time.

| Action | Owner |
| --- | --- |
| Determine where the policy will be deployed. | Security Manager, Systems Architect |
| Determine which rules need to added to the policy. | Security Manager, Systems Architect |
| Determine which computers or subnets will need to be added to an exemption list. | Security Manager, Systems Architect |

**Phase 4: Deploy the Policies in a Test Environment**

**Determine the Appropriate Test Environment**

The test environment will help you find and resolve any issues that could arise from your domain isolation deployment. The

more closely the test environment represents your actual IT environment, the more effective this testing will be.

| Action | Owner |
|---|---|
| Determine whether the current test environment is appropriate for testing domain isolation. | Security Manager, Systems Architect, Network Administrator |
| If the test environment is not appropriate, determine what changes need to be made to the test environment to properly test policy design and deployment. | Security Manager, Systems Architect, Network Administrator |
| Determine the cost/benefit of making changes to the test environment. | IT Manager, Security Manager, Systems Architect, Network Administrator |
| Determine whether testing can be accomplished in a smaller, non-critical domain or subdomain. | IT Manager, Security Manager, Systems Architect, Network Administrator |

**Deploy the Policies to the Test Environment**

You should deploy the policies using the plan you have created. If you are also deploying non-Microsoft IPsec solutions, such as those for UNIX or Apple Macintosh computers, you should deploy the one with the largest operating base first and refine and stabilize it before deploying the other solutions. If you are using these solutions, be sure to test them all together before deploying to your IT environment.

| Action | Owner |
|---|---|
| Deploy the least restrictive policy. | Security Manager |
| Monitor the communications in this segment for failures, etc. | Network Administrator |
| Correct any policy design issues (see "Refine Policies" below). | Security Manager, Systems Architect, Network Administrator |
| Continue to deploy, monitor, and correct until the deployment is successful. | Security Manager, Systems Architect, Network Administrator |
| Deploy other policies in order of increasing restriction. | Security Manager, Systems Architect, Network Administrator |

**Phase 5: Refine Policies**

During your test deployment, you might have issues that require a change to existing policies or the addition of new policies. The testing and refining process might require more than one cycle to find all the issues and redesign your policies to fit your environment.

For more information about troubleshooting IPsec issues, see "Testing and Monitoring Successful IPsec Operation" in Using Microsoft Windows IPsec to Help Secure an Internal Corporate Network Server.

| Action | Owner |
|---|---|
| Determine whether any computers that should be able to connect are blocked. | Network Administrator, Windows Systems Administrator, Desktop Configuration Manager |
| Determine whether any computers that should be blocked can connect. | Network Administrator, Security Manager |
| Determine whether any computer's performance is significantly affected by the policies. | Network Administrator, Database Administrator, Windows Systems Administrator |
| Determine whether any computers already have a conflicting IPsec policy implemented. | Network Administrator, Security Manager, Windows Systems Administrator, Desktop Configuration Manager |

| | |
|---|---|
| Determine whether any computers need to be updated so they can implement IPsec properly (for example, computers running Windows 2000, Windows XP with no service packs installed, and Windows XP with SP1 need to be updated for IPsec NAT-T support). | Network Administrator, Security Manager, Windows Systems Administrator, Desktop Configuration Manager |
| Determine whether any computers or network devices that were thought to be IPsec-compatible are not. | Network Administrator, Security Manager |
| Determine whether any computers that can use IPsec must have policies changed to work correctly (for example, any VPN servers that are not domain members, and therefore cannot use Kerberos v5 authentication, must use a certificate or preshared key). | Network Administrator, Security Manager |
| Determine where VPN or other remote connections do not work with IPsec. | Network Administrator, Security Manager |
| Determine whether there are any features or configurations that will not work with IPsec. | Network Administrator, Security Manager |

**Phase 6: Create a Deployment Schedule**

The policy testing and refinement processes will provide you with valuable information about how the deployment is best implemented for a given segment. You can also use the answers to the planning questions earlier to help you determine the best sequence of, and time frame for, IPsec deployment to your environment. Information based on the actions listed below might help you determine your deployment schedule.

As a best practice, Microsoft IT found that it worked for them to deploy to smaller, non-critical domains first, then to larger domains, and finally to mission-critical domains. They also deployed "Request Mode" first and then "Secure Mode." Additional best practices are provided in Improving Security with Domain Isolation.

| Action | Owner |
|---|---|
| Document the order in which the segments will be deployed. | IT Manager, Systems Architect, Security Manager |
| Document the best date and time for the deployment. | IT Manager, Systems Architect |
| Document how you will monitor the segment to make sure it is working properly. | Systems Architect, Security Manager |
| Document a contingency plan if connectivity is blocked. | Systems Architect, Security Manager |
| Document how you will back-out a change if something goes wrong. | Systems Architect, Security Manager |
| Determine when all parties should be informed of a pending change. | Helpdesk Manager |
| Determine how you will know that the deployment is sound enough to be implemented on the next segment. | Systems Architect, Security Manager, Helpdesk Manager |
| Determine how you will know that the entire deployment is sound and the goals have been achieved. | IT Manager, Systems Architect, Security Manager, Helpdesk Manager |

**Phase 7: Prepare for User and Infrastructure Support**

After you have finalized your deployment schedule, you can finalize when and how you will inform the operations staff, department heads, server owners, application owners, users, and partners of the pending changes. The availability of information about the process is important to a smooth deployment and the information should be available online (and in written form) well ahead of the actual deployment. The actions listed below might help you determine your Helpdesk needs.

| Action | Owner |
|---|---|
| Inform the Helpdesk staff of the changes, what they need to do to prepare for them, what they can expect to experience, and who to contact in case of problems. | Helpdesk Manager |

| | |
|---|---|
| Train Helpdesk staff using simulation drills with problems that are likely to arise. | Helpdesk Manager |
| Inform the Helpdesk staff about IPsec and the deployment process, possible problems that might arise, resources for helping users with these problems, and escalation of problems. | Helpdesk Manager |
| Inform the following groups of the changes, what they need to do to prepare for them, what they can expect to experience, and who to contact in case of problems:<br>• Department heads<br>• IT Operations staff<br>• IT Security staff<br>• Server owners<br>• Application owners<br>• Users<br>• Partners | Helpdesk Manager, Windows Systems Administrator |
| Document confirmation of compliance from all parties, signifying that they understand and have made all changes necessary for compliance. | Helpdesk Manager, IT Manager |

**Other Resources**
**IPsec concepts and overview**

- How IPSec Works
- Windows Server 2003 IPSec Concepts

**Examples of IPsec deployments**

- Improving Security with Domain Isolation
- Using Microsoft Windows IPsec to Help Secure an Internal Corporate Network (Foundstone paper)
- Server and Domain Isolation Using IPsec and Group Policy

**IPsec interoperability**

- Description of the Microsoft L2TP/IPsec Virtual Private Networking Client for Earlier Clients - Microsoft Support Article 324915
- "Soft Associations" Between IPsec-Enabled and Non-IPsec-Enabled Computers - Microsoft Support Article 234580

**IPsec limitations**

- IPsec Default Exemptions Can Be Used to Bypass IPsec Protection in Some Scenarios (Windows 2000) - Microsoft Support Article 811832
- Traffic That Can--and Cannot--Be Secured by IPsec (Windows 2000) - Microsoft Support Article 253169

**Designing IPsec policies**

- Windows Server 2003 Security Guide (The downloaded guide is in PDF format. Sample scripts for building IPsec policies from the command line are available in the "Tools and Templates\Security Guide\Sample Scripts" folder of the downloaded guide.)

- Designing IPSec Policies

**IPsec implementation details**

- How To Use IPsec

- Deploying IPSec Across Your Network

**IPsec testing**

- Testing Your Policies in a Test Lab

- Testing Your Policies in a Pilot Project

**IPsec troubleshooting**

- How to Disable IPSEC for Clients That Are Running an Earlier Version of Windows - Microsoft Support Article 323311

- IPSec Troubleshooting

# A Guide to Domain Isolation for Security Architects

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Vista

Learn how to assess the enterprise environment, plan domain isolation, and assess the implications of deploying domain isolation in an enterprise environment. Designed for security architects of enterprise organizations that are investigating using IPsec in Microsoft Windows to deploy domain isolation.

*Domain isolation* is the use of Internet Protocol security (IPsec) to require authentication, encryption, or both, among members of a Windows domain as well as between members of the domain and unknown or unauthorized computers. The goal of domain isolation is to isolate secured computers from unsecured computers by providing an additional level of protection for the traffic that is sent on a network. By using IPsec you add an additional defense-in-depth layer to the mechanisms that currently secure your network. IPsec-based communication changes the format of the TCP/IP packets transmitted over your network, which requires careful study and planning to mitigate risks and ensure success. This paper presents an overview of the three steps a security architect should follow to prepare for IPsec deployment:

1. Determine the current state of your IT infrastructure
2. Plan isolation groups
3. Consider operational issues for IPsec-based traffic

**Knowledge Prerequisites**

Prior to reading this paper you should read Domain Isolation with Microsoft Windows Explained at http://go.microsoft.com/fwlink/?LinkId=44642.

In addition, make sure that you have a thorough understanding of the following topics before you begin IPsec domain isolation:

- TCP/IP and core networking principles.
- Windows authentication, including use of the Kerberos V5 protocol and public key infrastructure (PKI).
- Active Directory concepts, including Active Directory structure and tools; users, groups, and other Active Directory objects; and Group Policy.
- Windows system security, including security concepts such as users, groups, auditing, and access control lists (ACLs).
- A thorough understanding of Windows Server 2003 IPsec policies. For more information, see the Windows Server 2003 Technical Reference at http://go.microsoft.com/fwlink/?LinkId=21711.

**Determining the Current State of Your IT Infrastructure**

Before starting the planning process for a domain isolation project, you need to collect and analyze up-to-date information about the current state of your IT infrastructure. A thorough and accurate assessment of your infrastructure provides the basis for planning what to isolate and how isolation should be implemented. Perform your data collection in three phases:

- Identify network configuration
- Map Active Directory infrastructure
- Identify host configuration

**Identifying Network Configuration**

Determining the current state of your network involves gathering information about network segmentation and infrastructure devices, and then performing an analysis of network communication.

**Network Segmentation**
Determine what networks are in place and what hosts exist on those networks. Typically, this information is illustrated by using schematic diagrams.

Figure 1 shows a very simple network diagram.




**Network Infrastructure Devices**
Next in your assessment, review the devices that make up the network infrastructure — routers, switches, load balancers, and firewalls. These devices are required to process IPsec-based traffic after the isolation solution is implemented. For this reason, it is important to determine whether these devices can handle the technical and physical requirements of the isolation design. Some of the characteristics to examine include:

- **Response to datagram change**. IPsec without encryption changes the offset for the destination ports and protocols within the TCP/IP packet. These changes might limit the functionality of traffic analysis applications, such as Cisco NetFlow.

- **Normal and peak utilization**. Peak use information can help determine whether the device is capable of using IPsec or if a memory or firmware upgrade is needed. If problems arise after IPsec is implemented, normal utilization data can assist in determining if the cause is related to high utilization of the device.

- **Whether the device is performing network address translation (NAT)**. To allow IPsec-based traffic to pass through a NAT device, you must ensure that IPsec NAT-T is supported on your IPsec peer computers.

- **Internal vs. external subnets and firewall configuration**. When a firewall or filtering router exists between IPsec peers, it must be configured to forward IPsec traffic on UDP source and destination port 500 for Internet Security Association and Key Management Protocol (ISAKMP), UDP source and destination port 4500 (IPsec NAT-T), IP protocol 50 for Encapsulating Security Payload (ESP), or IP protocol 51 for Authentication Header (AH).

For specific considerations for network infrastructure devices, see “Considering Operational Issues for IPsec-based Traffic" in this paper.

**Network Communication**
A thorough analysis of the communication flow on your network will provide three vital pieces of information:

- What communication is currently taking place between hosts

- What communication should be restricted

- What communication is at risk of failing when secured with IPsec

When you examine traffic flow, look closely at how all network devices interact. Keep in mind the following questions:

- What protocols and ports do servers and clients use to communicate with each other?

- Are there any security devices or projects that are implemented or planned that could impact the isolation project?

- Are certificates used to provide authentication?

- Do you have subnets that are dedicated to specific types of traffic (for example, Macintosh workstations and UNIX

workstations, or mainframe-to-terminal connections)?

**Mapping Active Directory Infrastructure**

Active Directory is the next source from which to gather information. Make sure that you understand the forest structure, including domain layout, organizational unit (OU) architecture, Group Policy, and site topology. This information provides the logical placement of computers, their configuration, and the impact to Active Directory as a result of implementing IPsec. For Active Directory, determine the following:

- **Number of forests**. The forest is the security boundary in Active Directory. Therefore, make sure that you understand the current Active Directory architecture because IPsec communication across forests is possible only if external trusts are configured.

- **Names and number of domains**. Active Directory domains are the building blocks of domain isolation. Therefore, computers that are domain members are assumed to support IPsec using Kerberos V5 authentication.

- **Number and types of Active Directory trusts**. Trusts affect the logical boundaries of isolation and can define how IKE authentication occurs. Active Directory trusts can affect Kerberos-based IKE authentication.

- **Names and number of sites**. Site architecture should be aligned with network architecture.

- **OU structure and Group Policy**. Organizational units can be used to apply IPsec policy. Review the OU structure to determine how Group Policy is used and how OUs are organized.

- **Existing use of IPsec policy**. Multiple IPsec policies that are assigned in different Group Policy objects (GPOs) are not merged. The IPsec policy applied is the one assigned to the GPO that is closest to the computer in the Active Directory domain structure. More IPsec policies in your environment means more complexity in your domain isolation design.

**Host Configuration**

The final piece of your current network environment assessment is a review of client and server host computers. IPsec-based communication depends on host systems meeting specific operating system and application requirements. Useful host information to gather includes:

- **Operating system, service pack, and hotfix versions**. The operating system version is a key factor in determining the ability of a host to communicate with IPsec. It is also important to track the current state of service packs and hotfixes that can be installed because these are often used to confirm that the minimum security standards have been met.

- **Domain membership**. This information is used to determine whether a computer can obtain IPsec policy from Active Directory or whether it must use a local IPsec policy. Hosts that are not domain members can use a local IPsec policy applied via scripts.

- **Physical location**. This information is simply the location of the device in your organization. It can be used to determine whether a device should participate in a particular isolation group based on its location or the location of the hosts that it communicates with regularly.

- **Hardware type and role**. This information can help determine the eligibility of a particular computer to participate in isolation and the isolation group in which to place the computer.

**For More Information**

Various methods can be used to gather data from the devices on your network. These methods range from high-end, fully automated systems to completely manual data collection. Generally, the use of automated methods to gather data, such as Microsoft® Systems Management Server, is preferred over manual methods for reasons of speed and accuracy. For more information about how SMS 2003 can help perform asset management, see the demonstration and the asset management datasheet on the SMS 2003 Asset Management page at http://go.microsoft.com/fwlink/?LinkId=45078.

For more information about determining your current network configuration, see "Chapter 3: Determining the Current State of Your IT Infrastructure" available at http://go.microsoft.com/fwlink/?LinkId=45130.

**Planning Domain Isolation**

With a good understanding of your company's network and host configuration, you are ready to specify your isolation domain. You can do this in four phases:

- Define isolation
- Determine isolation groups
- Specify communication paths and IPsec policies
- Complete implementation steps

**Defining Isolation**

*Domain isolation* is the use of Internet Protocol security (IPsec) to require authentication, encryption, or both, among members of a Windows domain as well as between members of the domain and unknown or unauthorized computers. In order to be a member of an isolated domain, a host must meet these minimum requirements:

- **Domain membership**. The host must be a member of an Active Directory domain.
- **Ipsec-compatible operating system.** The host must run an operating system that is IPsec-compatible such as Windows 2003, Windows XP, or Windows 2000 SP4.

Your organization might decide to add further requirements for security such as:

- A properly installed and configured antivirus program with up-to-date signature files
- Current security updates
- Current service packs

Hosts failing to meet minimum requirements — and whatever additional requirements your organization adds — cannot be included in the isolated domain. These hosts will be restricted from communicating with domain isolation members.

**Determining Isolation Groups**

The primary goal of domain isolation is to restrict the communication between isolated and non-isolated hosts. For planning purposes, it is useful to combine hosts into logical groups called isolation groups. An isolation group is a device for grouping host computers that share the same inbound and outbound network traffic requirements.

In Windows, isolation groups are implemented by placing hosts in organizational units (OUs) or security groups, and then by applying policy or by assigning access permissions. An isolation group's inbound and outbound access controls can be implemented by IPsec policy using permit/block actions, or by using IPsec policy in combination with Group Policy network access rights.

Figure 2 illustrates a simple isolation group design.




In this example, four groups are created for domain isolation:

- **Isolation domain**. This is the default group for all isolated computers and will likely be the group with the largest number of hosts. If the domain includes bidirectional trusted domains, members of those domains would be part of the isolation domain. All host-to-host communication within this group is secured by using IPsec.

- **Non-isolated hosts**. This group contains all non-isolated computers not capable of IPsec-based communication, such as Macintosh and UNIX computers.

- **Boundary isolation**. This group contains isolated hosts that will be allowed to communicate with non-isolated and isolated hosts. These computers will be exposed to a higher level of risk because they can receive incoming communications directly from unsecured computers. For example, you might want your Remote Installation Services (RIS) servers, which support downloading operating system images to computers that are not yet IPsec-compatible, to be boundary hosts.

- **Exemption lists**. This group contains computers that are by default in your isolated domain, but that must be available to all systems on the network. For example, you might want your domain controllers, and DNS, DHCP, and WINS servers to be exempted.

All computers on your network that share the same communications security policy are assigned to one of the isolation groups. After the groups are designed and host membership decided, you're ready to plan the network traffic between the isolation groups.

**Specify Communication Paths and IPsec Policies**

At this point in the planning process, the communications traffic requirements that will be allowed to pass between the groups need to be documented, along with the form the communications will take. Figure 3 illustrates communication paths for the isolation groups.




After IPsec policies are in place for the isolation groups, communication between the groups will proceed as shown in the following table.

| Path | Groups | Communication |
|---|---|---|
| 1 | From non-isolated hosts to hosts in the boundary group | Unsecured, plaintext TCP/IP connections. |
| 2 | From boundary group hosts to non-isolated hosts | Hosts attempt to communicate with IPsec-based traffic but can fall back to non-IPsec-based communication (the **Fall back to clear** option). |
| 3 | Between hosts in the isolation domain and boundary groups | Hosts attempt to use IPsec to protect the traffic. This traffic might require encryption, depending on the security requirements. If the IKE negotiation fails, the communications will fail because there is no **Fall back to clear** option when IKE negotiation fails. |
| 4 | From isolated hosts to non-isolated hosts | Hosts attempt to communicate with IPsec-based traffic but will fall back to non-IPsec-based communication. |
| 5 | From non-isolated hosts to isolated hosts | Blocked from initiating communication. |

Though not included in the diagram or table, all hosts are allowed to communicate with hosts in the exemption lists group using unsecured, plaintext connections.

#### Implementation Steps

With a complete understanding of your network and a plan for controlling traffic with isolation groups, you are now ready for the next steps in your utilization of IPsec to secure your network. These steps include:

- **Create IPsec policies**. Creating IPsec policies involves creating filter lists, filter actions, and IPsec policies containing rules that match filter lists with filter actions. For more information about these tasks, see "Chapter 5: Creating IPsec Policies for Isolation Groups" at http://go.microsoft.com/fwlink/?LinkId=45132.

- **Create a test lab**. To ensure that IPsec policy functions as expected and provides the appropriate level of security, test specific IPsec policy configurations on clients and servers in a lab environment, and then conduct pilot or beta tests in a limited operational environment before conducting a full-scale deployment. For more information about creating a test lab, see Setting Up IPsec Server and Domain Isolation in a Test Lab at http://go.microsoft.com/fwlink/?LinkId=44646.

- **Deploy IPsec Policies Across Your Network**. For more information about IPsec deployment, see Deploying IPsec in the Windows Server 2003 Deployment Kit at http://go.microsoft.com/fwlink/?LinkId=31789.

#### Considering Operational Issues for IPsec-based Traffic

IPsec provides greater security at the expense of network and processing performance and compatibility with other services, tools, and features. The following section describes these and other important operational considerations for managing IPsec policies.

#### The Impact of IPsec on Performance

IPsec has some costs in network and processor performance. The exact cost cannot be predicted because it depends on many factors, such as the level of application traffic and the number and configuration of IPsec filters deployed. However, the following factors should be understood and reviewed for your deployment.

#### Network Performance

IPsec impacts network performance in two ways.

**IKE SA establishment slows connection establishment**.

When a security association (SA) is initially established between two computers, a delay of up to three seconds occurs while security is negotiated. This time might vary depending on policy design, the authentication method that is selected, network roundtrip time, and the load placed on the servers required to establish the connection. Accordingly, end users might notice brief initial delays when IPsec-based connections are established. Typically, after SAs are established, no further delays occur.

**Large packet headers reduce throughput and increases network utilization**.

For each IPsec-based packet, the AH and ESP headers add an additional 24 to 36 bytes of data, depending on which method of encapsulation is used and whether IPsec tunneling is used. This decreases the effective IP payload size for upper-layer protocols such as TCP and UDP. As a result, traffic that is secured by IPsec uses more packets than traffic that is not secured by IPsec. The Microsoft IT department observed a mean bandwidth usage increase of 3 percent when they deployed IPsec (ESP-Null) across their network.

Different enterprise environments might experience a different amount of increased usage, depending on the mix of applications present. The increase in bandwidth usage is likely to have a more significant effect in a WAN scenario where the traffic is already near capacity, than in a LAN over Fast Ethernet connections.

#### Processor Performance

IPsec impacts processor performance in two ways.

**SA maintenance affects scalability and increases CPU and memory overhead**.

Performance degradation and lower throughput might be experienced by servers that must establish IPsec-based communications with many clients. In particular, full ESP-encrypted traffic places the highest burden on server CPU and can gain the most benefit from the use of offload network adapters. Typically, for each active IPsec client, 5 to 20 kilobytes (KB) of memory is used on a server. The IKE authentication method has the greatest impact on memory usage for each active IPsec client. For example, Kerberos V5 authentication typically requires less memory per SA than certificate authentication.

**Per-packet processing overhead increases CPU and memory overhead**.

Each IPsec-based packet creates additional performance overhead for encryption, authentication, and other packet processing functions. Each packet that is not secured by IPsec must still pass through all IPsec filters. Therefore, when you configure an IPsec policy, make sure that you create the minimum number of filters required for your environment. The impact of IPsec on throughput varies in proportion to the rate at which traffic is sent or received. Typically, active IPsec policies with 500 or fewer filters result in only a minor impact on throughput. In IPsec for Windows XP and Windows Server 2003, the per-packet processing performance is significantly improved, compared to the per-packet processing performance in IPsec for Windows 2000.

The effects of IPsec on CPU performance are typically negligible on clients because clients often have ample, spare CPU processing power and memory to accommodate IPsec. However, servers that have highly loaded CPUs might incur a significant increase in load, depending on the amount of traffic that must be secured by IPsec, the number of SAs that must be processed, and the cryptographic algorithms that are selected.

If your servers already have high loads on their CPUs, consider using IPsec offload network adapters before deploying IPsec. IPsec offload network adapters accelerate cryptographic operations used to secure IPsec packets. However, these adapters do not accelerate filter processing time. IPsec-enabled network adapters can be used with Windows 2000, Windows XP, and Windows Server 2003. Such network adapters typically complete IPsec operations at such a high rate that there is minimal change in transmission speeds, resulting in minimal performance degradation. However, such adapters generally have a maximum number of SAs that they can support, after which additional SAs are processed by Windows IPsec.

**Crossing Security Boundaries with IPsec Traffic**

When you design an IPsec policy, it is important to consider security boundaries, such as NATs, firewalls, and forests, that IPsec-based traffic must cross. Within a TCP/IP packet, IPsec changes the offsets for the destination ports and protocols. These changes might adversely affect the transmission of data over network devices and should be researched prior to IPsec deployment.

**NAT Traversal**

If traffic between the client and a server must pass through a network address translator (NAT), then IPsec cannot secure the traffic (the IKE negotiation will fail when translated by a network address translator). IPsec NAT traversal (NAT-T) allows IPsec peers that are behind NATs to detect the presence of NATs, negotiate IPsec SAs, and send ESP-protected data, despite the fact that the addresses in the IPsec-protected IPv4 packets change. IPsec NAT-T does not support the use of AH across network address translation. For more information about how IPsec NAT-T works, see IPsec NAT Traversal Overview, the August 2002 Cable Guy article at http://go.microsoft.com/fwlink/?LinkId=45080.

IPsec NAT-T is supported by Windows Server 2003, Windows XP with Service Pack 2 (SP2). IPsec NAT-T is also supported by Windows XP with Service Pack 1 (SP1) and Windows 2000 if used in combination with a free Web download. However, due to the behavior of IPsec and NATs, Windows XP SP2 by default no longer supports establishing IPsec NAT-T SAs to servers that are located behind NATs to avoid a perceived security risk.

If you must send IPsec-based communication across a NAT device, it is recommended that public IP addresses be used for all servers that can be connected to directly from the Internet. If this configuration is not possible, then the default behavior of Windows XP with SP2 can be changed to enable IPsec NAT-T security associations to servers that are located behind a network address translator.

In addition, servers running Windows Server 2003 also require a hotfix so that they can be reached by hosts on the other side of a NAT device. The hotfix will be included in the Windows Server 2003 SP1 release. (This hotfix can currently be obtained from Microsoft Product Support Services.)

To ensure that a server is reachable behind a NAT for IPsec traffic, you must configure the NAT with static translation entries that map UDP port 500 (IKE) and UDP port 4500 (IPsec NAT-T) traffic to the server.

For additional information, see the following articles:

- Problems with Using Network Address Translators – The Cable Guy, October 2004 at http://go.microsoft.com/fwlink/?LinkId=45081.

- IPsec NAT-T is not recommended for Windows Server 2003 computers that are behind network address translators on the Microsoft Support Site at http://go.microsoft.com/fwlink/?LinkId=45083.

- L2TP/IPsec NAT-T update for Windows XP and Windows 2000 on the Microsoft Support Site at http://go.microsoft.com/fwlink/?LinkId=45084.

**Creating Firewall Filters to Permit ISAKMP, AH, and ESP Traffic**
In some cases, IPsec-based traffic might need to pass through a router, firewall, or other filtering device. In the case of a router, unless the router filters TCP and UDP traffic or other upper-level protocol headers, no special configuration is required to permit the IPsec-based traffic to be forwarded. In the case of a filtering router or a firewall, you must configure these devices to allow IPsec-based traffic to be forwarded.

In order for IPsec-based communications to pass through a firewall or other filtering device, you must configure the firewall to permit IPsec traffic on UDP source and destination port 500 (ISAKMP), UDP source and destination port 4500 (IPsec NAT-T) and IP Protocol 50 (ESP). You might also need to configure the firewall to permit IPsec traffic on IP protocol 51 (AH) to permit troubleshooting by IPsec administrators and to allow the traffic to be inspected while it is still IPsec-protected.

For more information, see How to Enable IPsec Traffic Through a Firewall on the Microsoft Support Site at http://go.microsoft.com/fwlink/?LinkId=45085.

**Different Trust Environments**
When you design an IPsec policy, consider any logical boundaries that might affect IPsec-based communications. For example, your domain and trust environment is critical in determining an appropriate IKE authentication method.

Kerberos V5 authentication is recommended for use in a two-way (mutual) domain and forest trust environment. You can use Kerberos V5 for IKE authentication across trusted domains whether the domains are in the same forest or in different forests. If the two domains are in different forests, you must configure two external trusts, one for each direction, between the domains. The external trusts must use the fully qualified domain name (FQDN) of the domains, and you must configure the IPsec policy to allow an IKE initiator to communicate to any domain controller in the other forest, so the initiator can obtain a Kerberos ticket from a domain controller in the responder's domain. If the domains are on opposite sides of a firewall, then you must configure the firewall to allow Kerberos traffic over UDP destination port 88, TCP destination port 88, as well as UDP destination port 389.

For more information, see Active Directory in Networks Segmented by Firewalls, available at http://go.microsoft.com/fwlink/?LinkId=45087.

If the use of Kerberos is not possible because two-way trusts across forests cannot be established as in some large enterprise environments, you can use a PKI and digital certificates to establish IPsec-based communication. For an example of how Microsoft deployed their PKI, see Deploying PKI Inside Microsoft at http://go.microsoft.com/fwlink/?LinkId=45088.

**Operational Considerations**
In addition to the impact of IPsec on security boundaries, other operational areas of your network will likely be affected by an IPsec deployment. These areas include: non-IPsec-compatible support, NLB clusters, and network inspection technologies.

**Support for IPsec-Incompatible Hosts**
In most organizations, some workstations or servers cannot communicate by using IPsec. Such hosts include:

- Non-Windows-based computers and devices, such as Macintosh computers, UNIX computers, and personal digital assistants (PDAs).
- Hosts running older versions of the Windows operating system. Computers running Windows NT, Windows 95, Windows 98, and Windows Millennium Edition do not support Group Policy or IPsec policies.
- Windows-based computers that are not domain members. Stand-alone computers cannot perform IKE authentication using the Kerberos V5 protocol. A computer that is joined to a domain that is not trusted by the forest being used as the trust boundary for IKE authentication cannot participate in a domain isolation solution.
- Other non-Windows-based remote access or VPN clients. If a non-Windows-based IPsec virtual private network (VPN) client is being used for an organization's remote access solution, it typically is not compatible with Windows IPsec. If Windows IPsec cannot be used, the host cannot participate in an IPsec domain isolation solution.

In an IPsec-based network, these hosts would likely be considered non-isolated and communication with them would be blocked. However, business requirements often exist for these computers to communicate with isolated hosts in the isolation domain.

To deal with this situation, you can create a boundary group that contains hosts that will be allowed to communicate with non-isolated hosts. These boundary hosts will be exposed to a higher level of risk because they can receive incoming unsecured communications directly from non-isolated computers. Computers in the boundary group are isolated hosts that can communicate both with other isolated hosts and with non-isolated hosts.

Boundary hosts will attempt to communicate using IPsec by initiating an IKE negotiation to the originating computer. If no IKE response is received within three seconds, the host will use plaintext without IPsec. Because plaintext communication is allowed, these hosts must be secured in other ways or monitored for acceptable risk. For more information about planning and deploying boundary groups, see Chapter 4: Designing and Planning Isolation Groups available at http://go.microsoft.com/fwlink/?LinkId=45131.

**NLB Clusters**

There are challenges in implementing IPsec on Network Load Balancing (NLB) clusters and server clusters. NLB allows multiple servers to be clustered together to provide high availability for a service by providing automatic failover to other nodes in the cluster if an individual server fails. IPsec directly opposes NLB in that it prevents different computers from handling the same client connection. Because IPsec is negotiated between individual hosts before any data is transmitted, the same node in the cluster must respond to established IPsec connections or the traffic will be dropped.

If the particular node in the cluster fails, existing IPsec connections to that server cannot detect the failover and attempt to rebuild the security association until the preset time-out period expires. In Windows 2000 and Windows XP, IPsec must be idle for five minutes before the client attempts to re-authenticate. This results in a two-to-six minute interruption in service for that client.

Windows XP SP1 (with the NAT-T update), Windows XP SP2, and Windows Server 2003 reduce the preset time-out period to one minute. There will always be a delay if a client happens to be connected to a cluster node that fails.

For more information about IPsec and configuring NLB clusters, see the following articles on the Microsoft Support Site:

- How to Configure Network Load Balancing to Work with IPsec at http://go.microsoft.com/fwlink/?LinkId=45089
- IPsec is not designed for failover at http://go.microsoft.com/fwlink/?LinkId=45091
- How to Configure IPsec on an Exchange Server 2003 Back-End Server That Is Running on a Windows Server 2003 Server Cluster at http://go.microsoft.com/fwlink/?LinkId=45092.

**Network Inspection Technologies**

Within a TCP/IP packet, IPsec changes the offsets for the destination ports and protocols. These changes adversely affect many applications running on network devices like routers that monitor and manage traffic on the network.

Some network applications are updated to support IPsec, but many are not yet compatible:

- Router access control lists (ACLs) cannot examine protocol and port fields in ESP-encrypted packets, and therefore the packets are dropped. ACLs based only on IP address are forwarded as usual.
- Cisco NetFlow on routers cannot analyze packets between IPsec members based on protocol or port. Analysis based only on IP address is unaffected.
- Router-based Quality of Service (QoS) cannot use ports to classify packets. Host-based classification of packets supported by Windows is unaffected, as is router-based QoS using only IP addresses.
- Weighted Fair Queuing (WFQ) and other flow-based router traffic prioritization methods might not queue packets as intended resulting in suboptimal network utilization.
- Network monitoring tools might be unable to parse ESP packets that are not encrypted (ESP-Null).

In general, IPsec defeats network-based prioritization and port- or protocol-based traffic management. For encrypted packets, there is no workaround—the host itself must handle any traffic management functions. For unencrypted, authenticated-only packets, the devices and applications must be aware of how IPsec changes packets to be able to do anything with them other than route them to the correct host. If it is not possible to upgrade monitoring or management devices to support IPsec, it is vital that this information be recorded and figured into your isolation design.

Network Monitor added an ESP parser in version 2.1 to aid troubleshooting of unencrypted IPsec packets. Network Monitor is available in a light version in Windows Server 2003 and a full version for SMS. The SMS version adds formatting of data, advanced diagnostic capabilities, and other features for capture management.

The version of Network Monitor that is provided with Windows Server 2003 includes parsers for the ISAKMP (IKE), AH, and

ESP protocols. Network Monitor parsers for ESP can parse inside the ESP packet only if ESP null-encryption is being used. Network Monitor cannot parse the encrypted portions of IPsec-based ESP traffic when encryption is performed in software. However, if encryption is performed by an IPsec hardware offload network adapter, the ESP packets are decrypted when Network Monitor captures them on either the source or the destination and, as a result, they can be parsed and interpreted into the upper-layer protocols. If you need to diagnose ESP software-encrypted communication, you must disable ESP encryption and use ESP-null encryption by changing the IPsec policy on both computers.

**Summary**

The use of IPsec to isolate computers that belong to an Active Directory domain requires careful planning. Three steps presented in this paper will prepare the security architect for the necessary research and cross-team communication. The three steps presented in this paper are: Determine the current state of your IT infrastructure, plan isolation groups, and consider operational issues for IPsec-based traffic.

**See Also**

**Other Resources**

Server and Domain Isolation Using IPsec and Group Policy
Windows Server IPsec page
Deploying IPsec
http://go.microsoft.com/fwlink/?LinkId=31789
Interoperability Considerations for IPsec Server and Domain Isolation
Improving Security with Domain Isolation: Microsoft IT implements IP Security (IPsec)
Using Microsoft Windows IPsec to Help Secure an Internal Corporate Network Server
Problems with Using Network Address Translators
Domain Isolation with Microsoft Windows Explained

# Setting Up IPsec Domain and Server Isolation in a Test Lab

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This document is intended for network architects who are investigating the use of Internet Protocol security (IPsec) in Microsoft® Windows® operating systems to deploy domain and server isolation in their environments. The procedures in this document demonstrate how to set up IPsec domain and server isolation in a limited test environment, which you can use as a basis for your own deployment.

The objective of setting up domain isolation in a test lab is to configure a test lab network in which isolated domain member computers accept only authenticated communications initiated from other domain member computers, while ignoring communications initiated from non-isolated computers outside the isolated domain.

The objective of setting up group-specific server isolation in a test lab is to configure a test lab network in which a server computer accepts only authenticated communications initiated from domain member computers that are members of a specific security group, while ignoring communications initiated from non-isolated computers outside the isolated domain and isolated computers that are not members of a specific security group.

**Overview of the Domain and Server Isolation Test Lab**

The following steps are required for the successful setup and demonstration of the domain and server isolation lab:

1. Setting up four computers that make up the test lab network: two server computers and two client computers.
2. Demonstrating that the clients have full, unrestricted access to the servers.
3. Creating an IPsec policy to isolate the computers in the isolated domain from computers outside of the isolated domain.
4. Demonstrating, after the domain isolation policy is in place, that the isolated computers in the domain can access other isolated computers in the domain, but that non-isolated computers outside of the isolated domain do not have access to domain members.
5. Demonstrating, after the server isolation settings are in place, that specific, authorized isolated computers can access the isolated server, and that unauthorized isolated computers are denied access.

This lab is designed to demonstrate the basic capabilities of IPsec domain and group-specific server isolation in a limited test environment. It does not explore every possible configuration, but rather outlines a specific set of procedures to get you started. If you plan to deploy IPsec server and domain isolation in your production environment, see the "Other Resources" section at the end of this paper for links to IPsec planning and deployment documentation that may address the specifics of your network.

| **Note** |
| --- |
| The procedures in this document use the Microsoft Windows Server™ 2003, Standard Edition operating system. Your server might function differently based on the version and edition of the Windows operating system that is installed, your account permissions, and your menu settings. For example, the steps to open an item in Control Panel vary slightly based on how a server is configured. |

**Domain and Server Isolation Lab Setup**

The following illustration shows the configuration of the IPsec test lab.



Figure 1. Test lab configuration.

The following four computers make up the test lab network:

- **DC1**: a computer running Windows Server 2003 that is used as a domain controller and a Domain Name System (DNS) server.
- **FS1**: a computer running Windows Server 2003 that is used as a file sharing server.
- **CLIENT1**: a computer running Windows XP that is used as a client.
- **CLIENT2**: another computer running Windows XP that is used as a client.

The lab requires one network segment that represents a private network. All computers on the private network are connected to a common hub or Layer 2 switch. Private IP addresses are used throughout the test lab configuration. The private network of 172.16.0.0/24 is used for the test lab network.

In addition, each computer is manually configured with an IP address, subnet mask, and DNS server IP address. Dynamic Host Configuration Protocol (DHCP) and Windows Internet Name Service (WINS) servers are not used.

The following sections describe the configuration for each of the computers in the test lab.

**Setting Up DC1**

DC1 is a computer running Windows Server 2003 that is functioning as:

- The domain controller for the contoso.com Active Directory® domain.
- The DNS server for the contoso.com DNS domain.

| **Note** |
| --- |
| The domain name, contoso.com, is used here for example purposes only. You can use any domain name in your test lab configuration. |

To configure DC1 for these services, perform the following steps.

1. Install Windows Server 2003 as a stand-alone server.
2. Using Network Connections, configure the TCP/IP protocol for the Local Area Network connection with the IP address of **172.16.0.1** and the subnet mask of **255.255.255.0**.
3. Run the **Active Directory Installation Wizard** (**dcpromo.exe**) for a new domain named **contoso.com** in a new forest. Install the **DNS** service when prompted.

**Setting Up FS1**

FS1 is a computer running Windows Server 2003 that is providing file and print sharing services. To configure FS1 as a file server, perform the following steps.

1. Install Windows Server 2003 as a stand-alone server.

2. Using **Network Connections**, configure the TCP/IP protocol for the Local Area Network connection with the IP address of **172.16.0.2**, the subnet mask of **255.255.255.0**, and the DNS server IP address of **172.16.0.1**.

3. Using **System**, join FS1 to the contoso.com domain.

4. After restarting the computer, in **Windows Explorer**, share the root directory of **Local Disk (C:)** using the share name **ROOT**.

**Setting Up CLIENT1**

CLIENT1 is a computer running Windows XP Service Pack 1 (SP1) or SP2 that is acting as a client and member of the contoso.com domain. To configure CLIENT1 as a client, perform the following steps.

1. Install Windows XP SP1 or SP2 as a workgroup computer.

2. Turn off Windows Firewall.

3. Using Network Connections, configure the TCP/IP protocol for the Local Area Network connection with the IP address of **172.16.0.3**, the subnet mask of **255.255.255.0**, and the DNS server IP address of **172.16.0.1**.

4. Using System, join CLIENT1 to the contoso.com domain.

5. After restarting the computer, in Windows Explorer, share the root directory of Local Disk (C:) using the share name ROOT.

**Setting Up CLIENT2**

CLIENT2 is a computer running Windows XP SP1 or SP2 that is acting as a client but is not a member of the contoso.com domain and does not have a local IPsec policy. To configure CLIENT2 as a client, perform the following steps.

1. Install Windows XP SP1 or SP2 as a workgroup computer.

2. Turn off Windows Firewall.

3. Using Network Connections, configure the TCP/IP protocol with the IP address of **172.16.0.4**, the subnet mask of **255.255.255.**0, and the DNS server IP address of **172.16.0.1**.

4. In Windows Explorer, share the root directory of Local Disk (C:) using the share name ROOT.

**Configuring the Domain**

Now that the domain is set up, move the computers to which the domain isolation IPsec policy should be applied from the Computers folder in Active Directory Users and Computers into a new organizational unit (OU). On DC1, perform the following steps using Active Directory Users and Computers.

1. In **Active Directory Users and Computers**, create the **Isolate** OU within the contoso.com domain.

2. In **Active Directory Users and Computers**, move the **FS1** and **CLIENT1** computer accounts to the **Isolate** OU.

**Testing CLIENT1 and CLIENT2 Access to FS1**

Before configuring and assigning IPsec policy, confirm that hosts can successfully communicate.

1. On CLIENT2, from a command prompt, type **net view \\FS1.**

   You should be able to successfully connect to the shared folder on FS1.

2. On CLIENT1, from a command prompt, type **net view \\FS1.**

You should be able to successfully connect to the shared folder on FS1.

**Configuring IPsec Domain Isolation Policy**

Now that your domain computers are configured, you can configure an IPsec domain isolation policy to apply to the contoso.com domain. In the following steps, you will create a domain isolation policy that will allow only isolated computers to be accessed by other isolated computers within the isolated domain. Non-isolated computers outside the isolated domain, making communications requests to isolated computers inside the isolated domain, will be ignored. Perform the following steps using Group Policy.

1. Configure filter lists.
2. Configure policy filter actions.
3. Configure IPsec policy.
4. Assign IPsec policy.
5. Propagate IPsec policy to domain members.

**Configuring Filter Lists**

You will create the following IP filter lists:

- **Domain Controller**
- **Secure Subnet**

**To create the Domain Controller filter list**

1. On DC1, click **Start**, point to **Administrative Tools**, and then click **Domain Controller Security Policy**.
2. Under **Security Settings**, right-click **IP Security Policies on Active Directory**, and then click **Manage IP filter lists and filter actions**.
3. On the **Manage IP Filter Lists** tab, click **Add**.
4. In **IP Filter List**, clear the **Use Add Wizard** check box.
5. In **Name**, type **Domain Controller**, in **Description**, type **Permits all traffic between DC1 and any other computer**, and then click **Add** to add filter properties.
6. In the **Source address** list on the **Addresses** tab, confirm that **Mirrored** is enabled.
7. In **Source address**, select **My IP Address**. In **Destination address**, select **A specific IP Address**.
8. In **IP address**, type the address of DC1, **172.16.0.1**, click **OK**, and then click **OK** again.

**To create the Secure Subnet filter list**

1. On DC1, on the **Manage IP Filter Lists** tab, click **Add**.
2. In **IP Filter List**, confirm that the **Use Add Wizard** check box is cleared.
3. In **Name**, type **Secure Subnet**, in **Description**, type **Matches all traffic between computers on the 172.16.0.0 subnet**, and then click **Add** to add filter properties.
4. On the **Addresses** tab, confirm that **Mirrored** is enabled.

5. In **Source address**, select **My IP Address**. In **Destination address**, select **A specific IP Subnet**.

6. In **IP address**, type the private network ID, **172.16.0.0.** In **Subnet mask**, type **255.255.255.0**, click **OK**, and then click **OK** again.

#### Configuring Policy Filter Actions

IPsec Policy Filter Actions specify how the resulting IPsec policy negotiates security between computers. They indicate whether security will be negotiated, by what means it will be negotiated, and how to handle requests from non-IPsec-compatible computers.

You will create the following IP filter action:

- **Isolate Domain**

**To create the Isolate Domain filter action**

1. On DC1, click the **Manage Filter Actions** tab, clear the **Use Add Wizard** check box, and then click **Add**.

2. On the **Security Methods** tab, in New **Filter Action Properties**, click **Add**.

3. In **New Security Method**, click **Integrity only**, and then click **OK**.

4. In **New Filter Action Properties**, click the **General** tab. In **Name**, type **Isolate Domain**, and in **Description**, type **Prohibits communication between isolated and non-isolated hosts**.

5. Click **OK**, and then click **Close**.

#### Configuring IPsec Policy

Create an IPsec policy with the name "Domain Isolation." Use the following settings.

- IP Filter Lists: Secure Subnet, Domain Controller, All ICMP Traffic

- Filter Actions: Isolate Domain, Permit

- Authentication method: Kerberos

- Tunnel setting: This rule does not specify an IPsec tunnel

- Connection type: All network connections

1. On DC1, under **Security Settings**, right-click **IP Security Policies on Active Directory**, and then click **Create IP Security Policy**. Click **Next**.

2. In **Name**, type **Domain Isolation**, and in **Description**, type **Provides domain isolation**. Click **Next**.

3. Clear the **Activate the default response rule** check box, click **Next**, and then click **Finish**.

4. Clear the **Use Add Wizard** check box, and then click **Add**.

5. On the **IP Filter List** tab, select **Secure Subnet**. On the **Filter Action** tab, select **Isolate Domain**, and then click **OK**.

6. Click **Add**. On the **IP Filter List** tab, select **Domain Controller**. On the **Filter Action** tab, select **Permit**, and then click **OK**.

7. Click **Add**. On the **IP Filter List** tab, select **All ICMP Traffic**. On the **Filter Action** tab, select **Permit**, click **OK**, and then

click **OK** again.

#### Assigning IPsec Policy

Assign the Domain Isolation IPsec policy to the Isolate OU using Group Policy.

1. On DC1, click **Start**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.

2. Expand contoso.com, right-click **Isolate**, and then click **Properties**.

3. Click the **Group Policy** tab, and then click **New**.

4. Type **Domain Isolation** for the new policy name, and then click **Edit**.

5. Under **Computer Configuration**, expand **Windows Settings**, and then expand **Security Settings**.

6. Click **IP Security Policies on Active Directory (contoso.com)**, right-click **Domain Isolation**, and then click **Assign**.

   There should now be a small green diamond on the lower-right corner of the container for the Domain Isolation policy. This diamond indicates that the policy is currently assigned.

7. Close the Group Policy Object Editor, and then click **Close** to close **Isolate Properties**.

#### Propagating IPsec Policy to Domain Members

IPsec policy is updated through the standard Group Policy refresh intervals. To immediately update changes in IPsec policy settings, from a command prompt, run the following commands on FS1 and CLIENT1:

- **gpupdate /force**

- **net stop policyagent**

- **net start policyagent**

#### Verify IPsec Policy Is Applied

To verify that IPsec policy is applied to the isolated domain hosts, FS1 and CLIENT1, do the following:

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.

2. In **MMC**, click **File**, click **Add/Remove Snap-in**, and then click **Add**.

3. In the **Add Standalone Snap-in** dialog box, click **IP Security Monitor**, and then click **Add**.

4. Click **Close**, and then click **OK**.

5. Expand **IP Security Monitor**, expand computer name, and then expand **Quick Mode**.

6. Click **Specific Filters** and verify that three filters are applied to this host. If the filters do not appear, restart FS1 or CLIENT1.

For more information about IPsec Security Monitor, see "Monitor IPSec Activity" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=52598).



Figure 2. IP Security Monitor

**Testing Isolated and Non-Isolated Host Communication**

After configuring and assigning IPsec policy, FS1 and CLIENT1 should be isolated from communication with CLIENT2.

1. On CLIENT2, from a command prompt, type **net view \\FS1.**

   You should not be able to successfully connect to the shared folder on FS1. You will see the following message:

   "System error 53 has occurred. The network path was not found."

2. On CLIENT2, from a command prompt, type **net view \\CLIENT1.**

   You should not be able to successfully connect to the shared folder on CLIENT1.

3. On CLIENT1, from a command prompt, type **net view \\FS1**.

   You should be able to successfully connect to the shared folder on FS1.



Figure 3. Result of domain isolation.

**Configuring Group-Specific Server Isolation**

Now that the entire domain is secured with an IPsec policy, you may want to further isolate specific servers so that only certain isolated computers can communicate with them. Group-specific server isolation requires that a computer that is requesting IPsec authentication belongs to an Active Directory security group. If the requesting computer is a member of the group, it is allowed access to the server. If it is not a member of the group, the IPsec authentication fails.

The following steps build on the domain isolation lab steps, showing how to further isolate server FS2 so that only authorized client computers can communicate with it.

1. Join CLIENT2 to the domain.
2. Create an Active Directory security group named Authorized Computers.
3. Add CLIENT2 to the Authorized Computers group.
4. Configure the "Access this computer from the network" local Group Policy user right on FS1 for the Active Directory group.
5. Determine that CLIENT2 can access FS1 but CLIENT1 cannot.

**Joining CLIENT2 to the Domain**

Using System, join CLIENT2 to the contoso.com domain.

1. On CLIENT2, click **Start**, click **Control Panel**, click **Performance and Maintenance**, and then click **System**.
2. On the **Computer Name** tab, click **Change**.
3. Under **Member of**, click **Domain**, type **contoso.com**, and then click **OK**.
4. You will be prompted to provide a user name and user password to join the computer to the domain.
5. Click **OK** to close the **System Properties** dialog box. You will be prompted to restart your computer to apply your changes.
6. Reboot and log on to the **contoso** domain.
7. On DC1, in **Active Directory Users and Computers**, move the **CLIENT2** computer account to the **Isolate** OU.
8. On CLIENT2, reboot and log on to the **contoso** domain.
9. From a command prompt, type **net view \\FS1.**

You should be able to successfully connect to the shared folder on FS1.

**Creating an Active Directory Security Group Named "Authorized Computers"**

1. On DC1, click **Start**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
2. In the console tree, right-click **contoso.com**.
3. Click **New**, click **Group**, type **Authorized Computers** in **Group Name**, and then click **OK**.

**Adding CLIENT2 to the Active Directory Security Group Named "Authorized Computers"**

1. On DC1, double-click the **Authorized Computers** security group, click the **Members** tab, and then click **Add**.
2. In **Select Users, Contacts, Computers, or Groups**, type **client2; administrator**.
3. Click **Object Types**, click **Computers**, and then click **OK**.
4. Click **Check Names** to verify the object names, click **OK**, and then click **OK** again.

**Configuring the "Access this computer from the network" Local Group Policy Right**

1. On DC1, right-click the **Isolate** OU, click **Properties**, and then click the **Group Policy** tab.

2. Click **New**, type **Restricted FS1 Access**, and then click **Edit**.

3. In the console tree, under **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **User Rights Assignment**.

4. In **Policy**, double-click **Access this computer from the network**.

5. Click **Define these policy settings** to enable this setting.

6. Click **Add User or Group**, and then click **Browse**.

7. Type **Authorized Computers**, click **Check Names** to verify the object name, and then click **OK**.

8. Click **OK** to close the **Add User or Group** properties.

9. Click **OK** to close the **Access this computer from the network** properties.

10. Click **Yes** to accept the warning, and then click **Close** to close **Isolate Properties**.

11. On FS1, from a command prompt, run the following command: **gpupdate /force**.

12. Reboot FS1, CLIENT1, and CLIENT2.

**Testing CLIENT1 and CLIENT2 Access to FS1**
After configuring group specific isolation, FS1 should now be isolated from communication with CLIENT1, even though CLIENT1 is a host in the isolated domain.

1. On CLIENT1, from a command prompt, type **net view \\fs1**.

   You should **not** be able to successfully connect to the shared folder on FS1.

2. On CLIENT2, from a command prompt, type **net view \\fs1**.

   You should be able to successfully connect to the shared folder on FS1.



Figure 5. Result of server isolation.

**Conclusion**
This step-by-step guide has shown how to set up domain isolation and server isolation in a test lab by performing the following steps:

1. Setting up the four computers that make up the domain.

2. Demonstrating that the clients have full access to the servers.

3. Creating an IPsec policy to isolate computers in the isolated domain from computers outside the isolated domain.

4. Demonstrating that computers in the isolated domain can access other computers in the isolated domain, but that computers outside the isolated domain do not have access to members of the isolated domain.

5. Creating an IPsec policy that builds on the domain isolation policy by further isolating a server using group-specific server isolation.

6. Demonstrating that specific, authorized computers can access the isolated server and that unauthorized computers are denied access.

**Other Resources**

For more information, see the following resources:

- Server and Domain Isolation Using IPsec and Group Policy on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=33947).

- Improving Security with Domain Isolation on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=42414).

- Windows Server 2003 IPsec Web site (http://go.microsoft.com/fwlink/?LinkId=42985).

For the latest information about Windows Server 2003, see the Windows Server 2003 Web site.

# Network Load Balancing: Security Best Practices for Windows 2000 and Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

**In This White Paper**


- General Assumptions
- Deployment and Operational Management
- Summary of Security Attributes


**General Assumptions**

There are a number of general assumptions that need to hold around the infrastructure and operational best practices to ensure a secure environment in which to run NLB clusters:

1. Servers and storage are in physically secure locations.
2. Practical security implementations such firewalls, network probes and management tools to detect irregular traffic are in place.
3. Best practices and common sense in terms of security are adhered to in areas like administration, storage of logs, backup and restore, etc.
4. Platform level security best practices are adhered to in terms of assigning administrative permissions, Access Control Limiting resources, and other housekeeping roles.
5. The network infrastructure services such as Active Directory, DNS, DHCP, WINS, etc. must be secure. Any compromise of these infrastructure services can lead to a compromise of the NLB itself.
6. Administration tools or other applications that administer NLB clusters can be run from remote workstations. The administrator must ensure that the applications are run from trusted computers. Any compromise on the computers on which the applications are executing (that the cluster administrator runs), can compromise the cluster. For example, if there are untrusted users with elevated privileges on the workstation where the administration tools are run, untrusted or malicious code can be run against the cluster by the cluster administrator without the cluster administrator realizing.
7. NLB creates and maintains a set of objects in the operating system such as files, devices, registry keys, etc. These objects have a default security setting that ensures non-privileged users cannot impact the cluster configuration or the applications running on the cluster. Changing these security settings to less restrictive security settings can lead to the cluster being compromised and application data being corrupted.

# Deployment and Operational Management

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

**Cluster administrators**
Administrators can specify groups or individuals that are allowed to manage the cluster. Because NLB does not implement a cluster-wide administrative account, the cluster administrator must ensure that an appropriate administrative account with Administrator privileges has been enabled on all cluster hosts. This account is used by NLB management tools, such as NLB Manager and the NLB control program, wlbs.exe, to administer individual NLB hosts. To simplify this process and ensure uniform access on all hosts, it is preferable to use a single domain or global group account for this purpose.

**Using the cluster administrative tools through a firewall**
The NLB control program, wlbs.exe, should only be used on NLB hosts or, when remote control operations have been enabled, only on a trusted, internal computer within the firewall. By default, remote control operations are disabled, and they must only be enabled if the NLB cluster has been protected by a firewall that blocks UDP control ports 1717 and 2504. Otherwise, unauthorized remote control packets could be delivered to NLB hosts and thereby cluster operations.

The NLB Manager administrative tool can safely be used on trusted computers outside the firewall. This tool uses a secure method of communication to access NLB hosts. NLB Manager uses WMI interface to manage NLB hosts. Please refer to WMI/DCOM documentation to see what ports to unblock to allow WMI/DCOM access from outside the firewall. Nevertheless, make sure that this tool is only executed on a trusted, remote computer.

**Using Kerberos Authentication in an NLB Cluster**
NLB does not use its own cluster-wide authentication mechanism for cluster administration. Instead, NLB management tools use an appropriate administrative account, which must be enabled with administrative privileges on each NLB host. If a domain account is used for this purpose, it will use Kerberos authentication to control access to NLB hosts.

**Using NLB to support VPN connections**
NLB can be used to load balance virtual private network connections using the PPTP protocol in Windows 2000 and using both the PPTP or IPSec protocols in Windows Server 2003. When using IPSec, some client sessions may be disrupted when a new NLB host joins the cluster. This issue has been corrected in Windows XP SP1 and in Windows Server 2003.

**Network Security**
To coordinate their actions, all NLB hosts communicate with each other by periodically exchanging heartbeat messages on a common subnet, which is also used to receive incoming client requests. This subnet must be physically protected from intrusion. Otherwise, unauthorized heartbeat messages could be delivered to NLB hosts and disrupt cluster operations. Note that NLB heartbeat messages use a uniquely assigned Ethertype (0x886F) which is not normally routed across subnets. However, the cluster administrator must ensure that unauthorized computers and devices which could emit invalid NLB heartbeat packets are not placed directly on the NLB subnet. The effects of these disruptions are described below in the subsection "Rogue Servers".

**Network Flooding**
NLB cluster hosts may be affected by denial of service attacks, which flood the cluster with invalid network packets. These attacks can create additional CPU and network load on the cluster hosts and thereby delay the handling of valid client requests. In addition, they can cause NLB hosts to increase their memory usage up to parameterized limits. These limits are adjustable by setting registry variables on each NLB host, as described in the Windows 2000 Resource Kit. In extreme cases, a high volume of invalid network packets can disrupt extant client connections and require affected clients to reconnect to the cluster. Note that the Windows network protocol stack and server applications, such as Web servers, also may be affected by denial of service attacks.

**Rogue Servers**
Rogue NLB servers on an NLB subnet could emit heartbeat packets that disrupt cluster operations. Disruptions can include impeding NLBs convergence process such that cluster hosts cannot be added and recovery for a failed host cannot be completed. In addition, disruptions can block some or all of NLBs service to clients. NLB subnets must be physically protected from intrusion by unauthorized computers and devices.

**Best practices**
To summarize, best practices for NLB cluster administrators:

- Cluster administrators should establish a common administrative account on all NLB hosts that have Administrator privileges. To simplify cluster management, it is preferable that a single domain or global group account be used for this

purpose.

- The NLB subnet must be physically protected from intrusion by unauthorized computers and devices to avoid interference from unauthorized heartbeat packets.

- NLBs remote control mechanism is disabled by default. It should remain disabled unless the administrator can ensure that the NLB subnet has been physically protected from intrusion and that UDP control ports 1717 and 2504 have been firewalled to prevent unauthorized remote control.

- If remote control has been enabled, remote control operations using the NLB control program wlbs.exe must only be performed from a trusted computer within the firewall.

- NLB Manager provides a secure method for remote cluster management and should be used instead of wlbs.exe whenever possible. NLB Managers optional host list file should only be saved to a file that is accessible by user accounts (such as the NLB cluster administrators account) that have local Administrator privileges.

# Summary of Security Attributes

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section outlines the various objects and their associated security attributes needed to ensure that NLB runs successfully. This section lists the minimum security requirements. If the security attributes are made more restrictive, NLB may not run. If the security attributes are less restrictive, it may result in malicious attacks being able to compromise the cluster or corrupt end-user data. We recommend that the security attributes on the objects not be changed from the default settings.

**Directory and file protection**

| Object | Description | Minimum recommended security attributes |
|---|---|---|
| **<filespec>** | **NLB Manager host list file** | **BUILTIN\Administrators: Full control** |

**System resources**

The cluster service uses other system resources that are protected by default.

| Object | Description | Minimum recommended security attributes |
|---|---|---|
| **HKLM*** | **HKLM registry hive** | **SYSTEM:Full control**<br><br>**BUILTIN\Administrators: Full control** |

# Kerberos Protocol Transition and Constrained Delegation

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This document refers to features that are included with Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. You can run most of the features that are discussed in this document on the Windows Server 2003, Web Edition operating system, except for Active Directory and Kerberos Key Distribution Center (KDC).

**In This White Paper**


Introduction (Kerberos Protocol Transition and Constrained Delegation)

Authenticating Web Application Users

Windows Server 2003 Kerberos Extensions

Sample Scenario Source Files

Summary (Kerberos Protocol Transition and Constrained Delegation)

Conclusion (Kerberos Protocol Transition and Constrained Delegation)

# Introduction (Kerberos Protocol Transition and Constrained Delegation)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The number of security threats to Web applications that are deployed on the Internet increase every day. An enterprise that deploys a Web application on the Internet has to deal with security issues, such as denial of service attacks, identity spoofing, unauthorized access to program functions, and so on. You can mitigate certain types of security risks, such as unauthorized access to data, through the process of user authentication and authorization. This document examines the requirements that you can use to authenticate the users of Web applications and discusses how the new extensions to the Kerberos authentication protocol in Microsoft Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition can satisfy these requirements. Note that this document is not a specification of the protocol extensions and, therefore, detailed explanations of the protocol message flows and data structures are outside of the scope of this document. If you are familiar with the features of the extensions and want a summary of the implementation of the extensions, you can skip to the "Summary" section at the end of the document.

This article assumes that you have a basic understanding of the Kerberos authentication protocol and Windows security concepts, such as user privileges and tokens. For a Kerberos protocol primer, see "Exploring Kerberos, the Protocol for Distributed Security in Windows 2000" at the Microsoft Web site (http://www.microsoft.com/msj/0899/kerberos/kerberos.aspx).

The sample scenario and sample code in this document assumes some basic knowledge of Microsoft ASP.NET programming.

# Authenticating Web Application Users

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section discusses certain characteristics of Web applications emphasize some of the key requirements for Web application authentication protocols.

Web applications can use a variety of protocols to authenticate users. For example, Web applications can use the authentication protocols that are supported by HTTP or Web applications may depend on custom authentication mechanisms that are used with the existing Web server infrastructure (such as form and cookie-based authentication). However, not all of the authentication mechanisms have the same security attributes; some of the authentication mechanisms have stronger security properties than others. For example, Kerberos supports mutual authentication, while cookie-based authentication schemes do not.

Designers of scalable and reliable Web applications often follow design principles that advocate a tiered-application approach. For example, a Web-based line of business (LOB) application may consist of presentation, business logic, and data tiers. Frequently, the tiered approach introduces the delegation requirement that should be supported by the underlying network and operating system security infrastructure. Delegation allows a service to act on another security principles behalf to gain access to resources throughout the network. N-tiered applications may use delegation to satisfy various application requirements (for example, auditing or making authorization decisions in the original users context).

A more specific example is a Web-based bank account withdrawal application with a business component that resides in the business tier. The application verifies whether the Web user has permissions to withdraw the requested amount of $1,000 US from a given bank account. That business component needs to know the identity of the application user to allow the transaction. In the ideal scenario, the authentication protocol would propagate the original user context throughout the end-to-end application flow, even though the flow may cross physical computer boundaries. Application components that are in the respective tiers can then retrieve the original user context and perform authorization accordingly.

At the time that this document was published, Kerberos was the only widely-adopted authentication protocol that was capable of performing delegation. When n-tier applications are not using an authentication protocol that enables delegation, they typically design in user identifiers that are passed through application message bodies. Without further security measures to protect the application message while it is in transit, this approach is susceptible to message modification which could result in the identifier being modified and, subsequently, result in a security breach (such as identity spoofing). In addition to delegation, Kerberos also offers the benefits of mutually-authenticating communicating parties, as well as protecting the application message when it is in transit by using data encryption.

Kerberos is not the default authentication protocol that is used by Web applications because Kerberos uses port 88, a port that is typically blocked by corporate firewalls. Because of this port restriction, it is difficult to successfully deploy Kerberos in an Internet environment.

Note that:

- Web applications require firewall-friendly authentication at the authentication tiers
- Delegation is a frequently desired security feature for n-tier Web applications

Ideally, Web applications would use an authentication protocol with all of the desired security attributes of Kerberos, while offering ease of deployment over the Internet. To work around some of these limitations, consider the following engineering approaches:

- Fix the existing authentication protocols (like Basic, Digest, and SSL) so that they support delegation and mutual authentication. This is not a viable option because of the breadth of design and engineering work that is required for every authentication protocol that is supported.
- Fix the Kerberos protocol so that it can traverse corporate firewalls. This approach (or at least its functional equivalent) is not available.
- Do enough engineering work so that the Kerberos protocol becomes the default authentication protocol that is used behind corporate firewalls, independent of the protocol that is used at the Web server authentication tier to authenticate the Internet user. This approach has been adopted in the Windows Server 2003 family to help address delegation issues

that Web applications encounter. The new Kerberos protocol extensions provide significant security benefits to Web applications without modifying the functions of Web authentication protocols that are already deployed.

# Windows Server 2003 Kerberos Extensions

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

There are two new extensions in the implementation of the Kerberos protocol in Windows Server 2003:

- **Protocol transition**: The protocol transition extension allows a service that uses Kerberos to obtain a Kerberos service ticket on behalf of a Kerberos principal to the service without requiring the principal to initially authenticate to the Kerberos Key Distribution Center (KDC) with a credential.

- **Constrained delegation**: The constrained delegation extension allows a service to obtain service tickets (under the delegated users identity) to a subset of other services after it has been presented with a service ticket that is obtained either through the TGS_REQ protocol, as defined in IETF RFC 1510, or in the protocol transition extension.

  For more information about RFC 1510, see the IETF Web site (http://www.ietf.org).

| **Note** |
| --- |
| Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here. |

The extensions that are included in Windows Server 2003 are described in the following sections of this document.

**The Protocol Transition Extension**

As mentioned in the previous section, the protocol transition extension allows a service to obtain a Kerberos service ticket to the service on behalf of a Kerberos security principal. No user credential is required for the transition. Applications may transition into Kerberos through two mechanisms that are described in this section.

The first method that you can use to transition to the Kerberos protocol is to call the **LsaLogonUser** Win32 function or directly instantiate a **WindowsIdentity** class object. The new **KerbS4ULogon** logon message type and **KERB_S4U_LOGON** logon information structure have been defined to support protocol transitioning through the **LsaLogonUser** function. A successful call to the **LsaLogonUser** function returns a Windows user token.

There is a new **WindowsIdentity** class constructor in the Microsoft .NET Framework assembly that s the **LsaLogonUser** function to help managed code applications use the Kerberos protocol extension features. The sample in this article uses this new **WindowsIdentity** class.

Two variations of user tokens may be returned, depending on whether the caller process has the **Act as part of the operating system** privilege. When the call is made from a process without the **Act as part of the operating system** privilege, an identification token is returned. When the caller process is assigned the **Act as part of the operating system** privilege, an impersonation token is returned. The Windows Server 2003 family includes a new **SeImpersonatePrivilege** privilege. Only processes that have this privilege can obtain impersonation-level tokens and use the tokens to impersonate the users to perform certain operating system functions. To use the full feature of the impersonation-level token that is obtained through protocol transition, the caller process that uses protocol transition must have both the **SeImpersonatePrivilege** and **Act as part of the operating system** privileges. If the caller process has only the **Act as part of the operating system** privilege, the impersonation-level token is downgraded to an identification-level token when it tries to impersonate by using the token. Basically, an identification token allows the process to determine the identity and privileges that are associated with a user, but an impersonation token also allows the process to take on the token users identity and operate in that users identity. The user token level does not directly determine whether the caller process can subsequently use constrained delegation to call other services. The caller process may use the service ticket that was obtained through protocol transition to obtain service tickets to other services if you correctly configure the service account to use the Kerberos constrained delegation extension. (The constrained delegation extension is discussed in more detail in the next section.)

The programmatic mechanism that is described above allows applications to initiate protocol transition. In addition, you should also configure the applications to run as services that use Kerberos. Programming and configuration tasks are discussed in more details later in this document.

The second method that you can use for applications to transition into Kerberos is through other existing Windows-integrated authentication protocols, such as HTTP basic, digest, or Secure Sockets Layer/Transport Layer Security (SSL/TLS). Some applications are designed to use non-Kerberos, Windows-integrated authentication protocols when the Kerberos protocol may

not be the best option at the user-authentication tier. Subsequent application tiers may switch to Kerberos when it is needed and possible to do so. This concept is illustrated in Figure 1:

- In step 1, an application user is authenticated through a non-Kerberos authentication protocol that is integrated into the Windows operating system.
- In step 2, the security package creates a Windows user token for the authenticated user when the authentication is completed. Note that the user token is indirectly linked to a Kerberos service ticket for the user if the user is authenticated by the Kerberos protocol. If the user is not authenticated by Kerberos, no service ticket is associated with the user token.
- In step 3, the application impersonates the user and connects to a database server that requires Kerberos for authentication.
- In step 4, the Kerberos security package detects that there is no Kerberos ticket in the user token.
- In step 5, the Kerberos Security Support Provider (SSP) automatically initiates protocol transition by requesting a service ticket for the impersonated user.




**Figure 1: Protocol Transition from Windows-Integrated Authentication Protocols**

| **Note** |
| --- |
| Two sample Web scenarios that you can use for Kerberos protocol transition are discussed later in this document. |

The following network and operating system requirements must be met if you are using protocol transition in an Active Directory forest:

- You must run the operating system process that initiates protocol transition on a computer that is running Windows Server 2003
- All domain controllers that are used for protocol transition in the service and user account trust path domains must be running Windows Server 2003. There are two ways to accomplish this, and you can use either one:
    - You can upgrade all domain controllers in the domain to computers that are running Windows Server 2003
    - You can upgrade all domain controllers in Active Directory sites where domain controllers are used for protocol transition to computers that are running Windows Server 2003

To use protocol transition across Active Directory forests, the following additional requirements must also be satisfied:

- Both Active Directory forests must be running at the Windows Server 2003 forest functional level
- The forests must have established two way trusts across forests

**The Constrained Delegation Extension**

The reason why the constrained delegation extension is introduced in Windows Server 2003 can be best explained by describing the limitations in the Windows 2000 implementation of Kerberos delegation. In the Windows 2000 Kerberos delegation model, the Kerberos Key Distribution Center (KDC) does not limit the scope of services to which a Kerberos principal's identity can be delegated. In other words, after a service account is trusted for delegation, it can request service tickets on behalf of an authenticated user to any other service accounts. (See the Windows 2000 Kerberos Authentication white paper on the MSDN Web site for more technical information about the mechanics of delegation at http://www.microsoft.com/technet/prodtechnol/windows2000serv/deploy/confeat/kerberos.mspx.) This delegation method does not provide precise mechanisms for an application to specify a subset of service accounts that it determines to be trustworthy for delegation. Essentially, applications are exposed to broader impersonation risks that may span across resource domains that have different levels of security policy requirements; some of the security policies may not be as strict as the applications security requirements. From the domain administrators point of view, it is too risky to enable unconstrained Kerberos delegation in the enterprise because there is no way to exclude untrusted servers from participating in delegation.

With constrained delegation, domain administrators can configure service accounts so that they delegate only to specific sets of service accounts. On Active Directory that is running Windows Server 2003, a new **Allowed-to-Delegate-to** (**A2D2**) Active Directory attribute is added to service accounts to help enforce constrained delegation. The **A2D2** attribute lists the service principal names (SPNs) of other service accounts to which a given service account is allowed to delegate. When a Windows Server 2003 KDC processes a service ticket request by using the constrained delegation extension, the KDC verifies that the target service account is listed in the **A2D2** attribute. With the addition of constrained delegation, computers that are running Windows Server 2003 support two modes of delegation: Kerberos unconstrained delegation and Kerberos constrained delegation. Kerberos unconstrained delegation is supported if the user initially provides a user credential to obtain a ticket-granting ticket (TGT) that can be forwarded to a service that is trusted for delegation. This behavior is the same as the Windows 2000 behavior for Kerberos implementations. Constrained delegation can be used by a service if the service can obtain a Kerberos service ticket to itself for the user whose security context is to be delegated. It does not matter whether the user obtained the service ticket directly by authenticating through Kerberos, or whether the service obtained a service ticket on behalf of the user through the protocol transition extension. Note that constrained delegation works only in the boundary of a domain.

For a service account to use protocol transition in conjunction with constrained delegation, you must set a new control flag, **Trusted-to-Authenticate-for-Delegation** (**T2A4D**), on the service account in Active Directory. Only domain administrators can change the value of the control flag. If you set the flag for a service account, the account can use the protocol transition extension to obtain a service ticket that it can then use to obtain service tickets to a preconfigured subset of services that use Kerberos. In addition, you must not mark the impersonated user account as a sensitive account that cannot be delegated.

**Enabling User Identification, Impersonation, and Delegation Through Extension Configurations**

You can perform several security-related operations by using the extensions that have been discussed in this document. You can perform these security-related operations, even when a users Kerberos authentication credential is not available. For example, you can:

- Obtain both privilege and authorization information for a user.
- Impersonate a user by running an application request in a users identity.
- Make remote calls to a set of preconfigured Kerberos service instances in a given users identity.

Whether an application can perform one or all of the above operations is based on how you configure the service account. As discussed earlier in this paper, an application can initiate protocol transition by using two mechanisms: by programmatically invoking the **LsaLogonUser** function or automatically through the Kerberos SSP. The necessary conditions for protocol transition to take place are different for the two mechanisms.

Scenario A, Scenario B, Scenario C, and Scenario D in Figure 2a, Figure 2b, Figure 2c, and Figure 2d illustrate the successful results that can occur when you call the **LsaLogonUser** function with different service account configurations.

In Scenario A in Figure 2a, protocol transition allows a service without the **Act as part of the operating system** privilege to obtain an identification-level token that works only for making application authorization decisions. In this scenario, the **Act as part of the operating system** service account does not have the **Trusted-to-Authenticate-for-Delegation** (**T2A4D**)

permission and the service ticket that is issued does not have the Forwardable flag set. A ticket that does not have the Forwardable flag set cannot be used for a constrained delegation operation.




1. The program calls the **LsaLogonUser** function or instantiates the **WindowsIdentity** class object.
2. The Kerberos **Protocol Transition** service ticket.
3. KDC checks the delegation policy for the T2A4D flag on the service account.
4. The Kerberos Protocol Transition service ticket reply.
5. An identity or impersonation-level user token is created when the **LsaLogonUser** function or the **WindowsIdentity** class constructor call returns to the program.

**Figure 2: a Scenario A: Service Account Does Not Have TCB and the T2A4D Flag Is Not Set**

In Scenario B in Figure 2b, the service account has the **Act as part of the operating system** privilege and protocol transition allows the service to obtain an impersonation token. As noted previously, the service process must also have the **SeImpersonatePrivilege** privilege to use the impersonation token so that it can run in the users identity. You do not want any arbitrary service to be able to impersonate and run as other users. Because of this, the service must have both the **Act as part of the operating system** and **SeImpersonatePrivilege** privileges that can be assigned only by members of the local administrators group. Only local administrators can give those services the ability to impersonate users on the computers where the services are running. Like Scenario A in Figure 2a, the Kerberos service ticket that is obtained by the service account in Scenario B does not have the Forwardable flag set and therefore the ticket cannot be used for constrained delegation.




1. The program calls the **LsaLogonUser** function or instantiates the **WindowsIdentity** class object.
2. The Kerberos **Protocol Transition** service ticket.
3. KDC checks the delegation policy for the T2A4D flag on the service account.
4. The Kerberos Protocol Transition service ticket reply.
5. An identity or impersonation-level user token is created when the **LsaLogonUser** function or the **WindowsIdentity** class constructor call returns to the program.

*Impersonation is enabled if the process has the **SeImpersonatePrivilege** privilege.

**Figure 2: b Scenario B: Service Account Has TCB and the T2A4D Flag Is Not Set**

The **T2A4D** control flag is set for the service account on Active Directory by domain administrators and is independent of **Act as part of the operating system** privileges on the local computer. In Scenario C in Figure 2c, the KDC that issues the service ticket does not check to see if the calling process has either the **Act as part of the operating system** or **SeImpersonatePrivilege** privileges when processing a service ticket request by using the protocol transition extension. Because of this, the KDC issues a service ticket with the Forwardable flag set if the **T2A4D** control flag is set. This service ticket is accepted by a KDC for constrained delegation requests if both the requesting and target services are in the same domain and the target service (SPN) is listed in the requesting services **A2D2** attribute in Active Directory.




1. The program calls the **LsaLogonUser** function or instantiates the **WindowsIdentity** class object.
2. The Kerberos **Protocol Transition** service ticket.
3. KDC checks the delegation policy for the T2A4D flag on the service account.
4. The Kerberos Protocol Transition service ticket reply.
5. An identity or impersonation-level user token is created when the **LsaLogonUser** function or the **WindowsIdentity** class constructor call returns to the program.

*Impersonation is enabled if the process has the **SeImpersonatePrivilege** privilege.

**Figure 2: c Scenario C: Service Account Has TCB and the T2A4D Flag Is Set**

In Scenario D in Figure 2d, the service ticket is issued with the Forwardable flag set and the created user token is an identification-level token because the service does not have the **Act as part of the operating system** privilege. Although the service ticket allows constrained delegation, the service cannot use the identification-level token to impersonate the user. Most practical applications require an impersonation-level token to perform network-related operations before authenticating as the delegated user. Because of this, the service cannot impersonate the user and use the service ticket for constrained delegation for most applications.




1. The program calls the **LsaLogonUser** function or instantiates the **WindowsIdentity** class object.
2. The Kerberos **Protocol Transition** service ticket.
3. KDC checks the delegation policy for the T2A4D flag on the service account.
4. The Kerberos Protocol Transition service ticket reply.
5. An identity or impersonation-level user token is created when the **LsaLogonUser** function or the **WindowsIdentity** class constructor call returns to the program.

**Figure 2: d Scenario D: Service Account Does Not Have TCB and the T2A4D Flag Is Set**

In situations where protocol transition is initiated by the Kerberos Security Support Provider (SSP) as in the scenario shown in Figure 1(for example, not by directly calling the **LsaLogonUser** function or the **WindowsIdentity** class object), an impersonation-level user token must already exist for that user before the Kerberos SSP tries to perform protocol transition. The impersonation-level user token could be created previously by any of the Windows-integrated SSPs. The service must also have the **SeImpersonatePrivilege** privilege to impersonate the user and initiate the protocol transition extension through the Kerberos SSP. Subsequently, the services ability to perform constrained delegation is determined by the Active Directory policies and the state of the Forwardable flag in the service ticket that is issued through protocol transition. The Active Directory policy checks that are used by the KDC to issue service tickets that are enabled for constrained delegation are the same as the flags that were previously described (for example, the **T2A4D** and **A2D2** flags).

**Sample Web Service Scenarios**

Figure 3 displays two simple Web service scenarios that illustrate how the Kerberos features work.

The first sample scenario illustrates where protocol transition is initiated by the Kerberos SSP. The scenario consists of a browser client that invokes an ASP.NET Web service that, in turn, calls a stored procedure on a computer that is running Microsoft SQL Server.

The second sample scenario illustrates how protocol transition works by programmatically instantiating a **WindowsIdentity** class object. In this scenario, the browser client is used to navigate through an ASP.NET Web application that invokes the same SQL Server stored procedure that is invoked in the first sample scenario.

Both scenarios are configured to use an Active Directory forest with a single-domain Kerberos Key Distribution Center (KDC). Note that you can run both the ASP.NET Web service and application from different virtual directories of the same computer that is running Microsoft Internet Information Services (IIS).

For the sample scenarios to work, Active Directory must run on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. Also, IIS must run on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows Server 2003, Web Edition. The computer that is running SQL Server 2000 must run Windows 2000 Professional, Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or a later version of Windows.

**Sample Scenario One**
Sample scenario one illustrates the following authentication flow:

1. The user is initially authenticated to the Web service through a non-Kerberos Windows authentication protocol. This sample scenario uses Hypertext Transfer Protocol (HTTP) digest authentication.

2. The Web service impersonates the authenticated Windows user before connecting to the SQL server.

3. The Web service tries to open a connection to the SQL server, and is challenged to prove its identity.

4. The Kerberos authentication package that is running on the same computer as the Web service cannot find a Kerberos ticket for the impersonated user and, therefore, uses protocol transition to obtain a service ticket to the Web service in the impersonated user context.

5. The Web service uses constrained delegation to obtain a service ticket to the SQL server. When this occurs, the authentication is completed in the impersonated user identity.

For demonstration purposes, the Web service uses HTTP digest authentication for initial user authentication. Although the sample uses digest authentication, you can easily modify the sample to support protocol transitioning from other Windows-integrated authentication protocols by changing the IIS directory security configuration.

The sample Web service supports the public **TestKerbSfu** method that the browser client invokes. This Web service method then invokes the **GetUserInfo** SQL server stored procedure that returns the identity of the user who invoked the stored procedure. The Web methods final return value is a string that contains the caller identities of the Web service and stored procedure. If you set up the sample that is described in this document, the two user identities that are in the return string are identical and they reflect the domain user account under which the Web service and stored procedure code operates.

In this scenario:

1. The browser gains accesses to the Web service with a URL such as http://www.fabrikam.com/KerbExtWebSvc/authxws.asmx?op=TestKerbSfu.

2. The browser displays an authentication dialog box that prompts for the user name and password to authenticate to the Web server.

3. The user name and authentication credential is sent to the server on which IIS is running to host the Web service.

4. IIS performs the HTTP digest authentication before forwarding the Web service request to the .NET Framework .asmx handler.

5. Before establishing the connection to the SQL server, the Web service impersonates the authenticated user and transitions into the Kerberos protocol by obtaining a Kerberos service ticket to the Web service.

The Web service uses Kerberos constrained delegation to obtain a service ticket to the SQL server in the authenticated browser users identity.




**Figure 3: Sample Scenarios**

**Sample Scenario Two**

Sample scenario two illustrates a slightly different method that you can use to use the Kerberos protocol transition extension. In this scenario, the Web service is replaced by an ASP.NET Web application that uses ASP.NET form-based authentication to authenticate users. The Web application goes through a series of interactive steps to show you how an application can use an authentication protocol that is not integrated in the Windows operating system to authenticate the Web users and then transition into Kerberos to authenticate to backend services, such as the SQL server that was used in sample scenario one.

The Web application illustrates the following authentication flow:

1. The browser client navigates to the virtual directory of the sample and is redirected to the login Web form.

2. The user types a user name and password in the login Web form, and then clicks **Login**.

3. The login Web page authenticates the user by using a custom authentication mechanism (for example, by checking a hashed password that is stored in a SQL server database), and then issues a login cookie to the browser when the authentication is successful. The user is then redirected to the default Web page for the Web application.

4. The user types a unique principal name (UPN) to be used for Kerberos protocol transition, and then the user clicks **Do protocol transition**.

5. The Web application transitions into Kerberos by using the specified UPN, and then saves the Kerberos security context in the Web application session state.

6. The user has the option to connect to the SQL server if they click **Invoke SQL Server**. When the user clicks this option, the Web application retrieves the user context from the application session state, and then the Web application impersonates the user before connecting to the SQL server.

7. When the Web application is challenged by SQL server to authenticate, the Web application uses Kerberos constrained delegation to obtain a service ticket to the SQL server in the impersonated user identity.

8. The Web application returns a response string that shows the impersonated user identity and the user identity that is authenticated on the SQL server.

# Sample Scenario Source Files

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The sample in this section uses Microsoft Visual Studio .NET and refers to the Windows Server 2003 implementation of the .NET Framework runtime. The version of the .NET Framework runtime in Windows Server 2003 is different than the runtime that is included with Visual Studio .NET.

You can obtain the source code and setup programs in the sample scenario from the Microsoft Download Center Web site at http://download.microsoft.com/download/8/9/7/89734348-63AB-4F55-B11B-8C0AC50612B1/KerberosProtocolExtensions.msi. The file contains the following items:

| Files and folders | Descriptions |
|---|---|
| Authxws* | Contains the implementation of the ASP.NET Web service. You need to create an IIS /Authxws virtual directory that points to the Authxws folder so that the Visual Studio .NET project can successfully start. |
| WebApp* | Contains the implementation of the ASP.NET Web application. You need to create a WebApp IIS virtual directory that points to the WebApp folder so that the Visual Studio .NET project can successfully start. |
| KerbExtWebSvc* | Contains the Visual Studio .NET Deployment project that builds the setup program that you can use to create the ASP.NET Web service. |
| KerbExtWebApp* | Contains the Visual Studio .NET deployment project that builds the setup program that is used to create the ASP.NET Web application. |
| Sql* | Contains a SQL script that sets up a sample database and a stored procedure. |
| SfuIdentity* | Contains managed C++ code that illustrates Kerberos protocol transition through the **LsaLogonUser** function. Note that the rest of the sample does not use code in this folder. This code is sample code for applications that need to use native Win32 functions instead of the **WindowsIdentity** class. |

**Sample Scenarios Configuration Tasks**

| **Note** |
|---|
| To complete the following configuration steps, you must have domain administrator privileges to Active Directory, local administrator privileges to the ASP.NET Web service, and administrative privileges on the SQL server. |

To set up the sample:

1. Configure the domain to run at the Windows Server 2003 functional level:

    a. Open the Active Directory Users and Computers Microsoft Management Console (MMC) Snap-in.

    b. In the console tree, right click **Active Directory User and Computers** *server_name* , point to **All Tasks**, and then click **Raise Domain Functional Level**.

    c. In the **Select an available domain functional level** box, click **Windows Server 2003**.

2. Create two domain user accounts for the IIS process pools that host the ASP.NET Web service and application:

| **Note** |
|---|

| |
|---|
| Although you can run the sample Web service and application in the default Network Service account that all newly-created Web applications use on the server on which IIS is installed, the Network Service account is generally used to run other services and Web applications. This means that other processes that run in this account can also exploit the extra Kerberos extensions-related privileges if you give the Network Service account the ability to use both the protocol transition and constrained delegation extensions. This is not the desired behavior, so you should isolate the ability to use Kerberos extensions to a specific domain user account on computers that are running Windows Server 2003. |

a. Open the Active Directory Users and Computers MMC Snap-in.

b. In the console tree, double-click **Active Directory Users and Computers** *server_name* .

c. In the details pane, right-click **Users**, point to **New**, and then click **User**.

d. In the **New Object User** box, type the appropriate information , and then complete the procedure.

e. Repeat step c and step d for the second user account.

3. Create a domain user account for the SQL server service. To do this, repeat the instructions in step 2.

4. Register Kerberos Service Principal Names (SPNs) for the user accounts that you created in the previous two steps.

   To use the domain account that you created as a Kerberos service account, you must use the Setspn (Setspn.exe) command-line tool to create an SPN for the account. For example, to add an SPN named ws/kerbsfuidentity for the kerbsfuidentity domain account, type the following command at a command prompt:

   **setspn a ws/kerbsfuidentity kerbsfuidentity**

   To create an SPN for the SQL Server service that is running on sqlcomputer.devdomain.fabrikam.com by using the sqlservice service account, type the following command at a command prompt:

   **setspn a MSSQLSvc/sqlcomputer.devdomain.fabrikam.com:1433 sqlservice**

| **Note** |
|---|
| For more information about how to configure SQL Server service for Kerberos authentication, see Microsoft SQL Server 2000 Security on the MSDN Web site. |

5. Create a domain user account that will be used as the user account for Web browser user authentication in scenario one and for protocol transition in scenario two. To create a domain user account, repeat the procedure that is described in step 2.

6. Configure the service accounts that will be trusted for constrained delegation to the SQL Server.

   In the sample scenario, you want to constrain the ability of the service accounts to present delegation credentials to only the specific SQL Server service account. Use the following procedure to add the specified service to the **A2D2** attribute on the service accounts that you created in step 2:

   a. Right-click the user account, and then click **Properties**.

   b. In the **Properties** dialog box, click the **Delegation** tab,, and then click **Trust this user for delegation to specified services only**.

   c. Click **Use any authentication protocol**, and then click **Add**.

   d. Click the SQL Server service SPN that you registered in step 4.

| Note |
| --- |
| The **Delegation Property** tab for a user account is available only if you have created an SPN for the user account. Step 4 describes how to create an SPN. |

7. Give service accounts Read access to the **tokenGroupsGlobalAndUniversal** user object attribute in Active Directory.

   The service accounts that is performing the protocol transition procedure must have Read access to the **tokenGroupsGlobalAndUniversal** user object attribute in Active Directory. To give the service accounts Read access:

   a. Open the Active Directory Users and Computers MMC Snap-in.

   b. In the console tree, click **Builtin**.

   c. In the results pane, right-click **Windows Authorization Access Group**, and then click **Properties**.

   d. Click the **Members** tab, and then click **Add**.

   e. In the **Enter the object names to select (examples)** box, type the name of the first service account that you created in step 2, and then click **OK**.

   f. Repeat step 7d and step 7e for the second service account that you created in step 2.

8. Add the service accounts to the IIS worker process group on the computers that are running the ASP.NET Web service and the application:

   a. Click **Start**, point to **Administrative Tools**, and then click **Computer Management**.

   b. In the console tree, click **Local Users and Groups**, and then click **Groups**.

   c. In the results pane, right-click **IIS_WPG**, and then click **Add to Group**.

   d. Click **Add**, in the **Enter the object names to select (examples)** box, type the name of the service account, and then click **OK**.

| Note |
| --- |
| At the time that this document was published, you also need to manually configure the service account to have access to the %SystemRoot%/Temp folder so that you can run ASP.NET. Confirm that the service account has both **List Folder\Read Data** and **Delete** permissions to the %SystemRoot%/Temp folder. |

9. Deploy the ASP.NET Web service:

   a. Create a new IIS process pool:

      Click **Start**, point to **Administrative Tools**, and then click **Computer Management**.

      In the console tree, double-click **Services and Applications**, and then double-click **Internet Information Services (IIS) Manager**.

      In the console tree, right-click **Application Pools**, point to **New**, and then click **Application Pool**.

      In the **Application Pool ID** box, type a name for the new application pool, and then click **OK**.

   b. Set the process pool identity to one of the user accounts that you created in step 2 for the process pool that hosts

the ASP.NET Web service:

In the **Application Pools** pane, right-click the application pool that you created in the previous step, and then click **Properties**.

Click the **Identity** tab, click **Configurable**, and then type the appropriate user name and password.

10. Create a new virtual directory for the ASP.NET Web service. To do this, run the Setup.exe program in the KerbExtWebSvc folder to start the wizard that will help you set up a virtual directory for the Web Service.

11. Configure the Web service to use digest authentication:

    a. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.

    b. Right-click the Web service virtual directory, and then click **Properties**.

    c. Click the **Directory Security** tab, and then click **Edit** in the **Authentication and access control** box.

    d. In the **Authentication Methods** box, confirm that the **Digest authentication for Windows domain servers** check box is the only check box that is selected, click **OK**, and then click **Apply**.

12. Configure the Web service to use the process pool that you created in step 9a:

    a. In IIS Manager, right-click the Web service virtual directory, and then click **Properties**.

    b. Click the **Virtual Directory** tab, click **Application Pool**, and then click the application pool that you created in step 9a.

13. Configure the name of the computer that is running SQL Server in the Web services Web.config file.

    The Web service uses the application configuration section in Web.config file to find the name of the computer that is running SQL Server. Because of this, you need to modify the value of the **sqlServer** key (in the <appSettings> section) in the Web.config file to point to the computer that is running SQL Server. For example, if the name of the computer that is running SQL Server is *computer_name*, open a text editor, such as Notepad, and then open the Web.config file that is located in the folder where the Web service is installed. Search for <appSettings>, and then make the following changes to the file:

    <**appSettings**>

    <**add key="sqlServer" value="***computer_name***" /**>

    <**/appSettings**>

14. Deploy the ASP.NET Web application:

    a. Assign the **Act as part of the operating system** privilege to one of the user accounts that you created in step 2. This account should be different from the account that you used in step 9b:

    On the computer that is hosting the Web application, click **Start**, point to **Administrative Tools**, and then click **Local Security Policy**.

    In the console tree, double-click **Local Policies**, and then double-click **User Rights Assignment**.

    Right-click the **Act as part of the operating system** policy, point to **Properties**, and then click **Add user or Group**.

Type the user account that you created in step 2 for the IIS process pool that is hosting the ASP.NET Web application, and then click **OK**.

b. Create a new IIS process pool for the Web application. To do this, repeat the procedure in step 9a.

c. Set the process pool identity to the user account that was assigned the **Act as part of the operating system** privilege in step 14a. To set the identity of the process pool that you created in the previous step, repeat the procedure in step 9b.

d. Create a new virtual directory for the ASP.NET Web application. To do this, run Setup.exe in the KerbExtWebApp folder to start the wizard.

e. Enable anonymous access to the Web application:

   Using IIS Manager, right-click the Web application virtual directory, and then click **Properties**.

   Click the **Directory Security** tab, and then, in the **Authentication and access control** box, click **Edit**.

   Confirm that the **Enable anonymous access** check box is the only check box that is selected in the **Authentication Methods** dialog box, click **OK**, and then click **Apply**.

f. Configure the Web application to use the process pool that you created in step 14b. To do this, repeat the procedure in step 12, and then specify the process pool that you created in step 14b when you are prompted to do so.

15. In the Web applications Web.config file, configure the name of the computer that is running SQL Server.

16. Configure the computer that is running SQL Server:

    a. Configure SQL Server to use Integrated Windows Authentication:

       Click **Start**, point to **Programs**, point to **Microsoft SQL Server**, and then click **Enterprise Manager**.

       Right-click the instance of the SQL Server that you want to manage, and then click **Properties**.

       Click the **Security** tab, and then click **Windows only**.

       Click **This account**, type the domain user name and password to configure the startup service account to use the user account that you created in step 3, and then click **OK**.

    b. Set up the sample database and stored procedure by running the Kerbs4u.sql script file:

       To run the Kerbs4u.sql script file, click **Start**, point to **Programs** point to **Microsoft SQL Server**, and then click **Query Analyzer**.

       To connect to the target SQL Server computer, in the **Connect to SQL Server** box, click the SQL Server to which you want to connect, and then click **OK**.

       On the **File** menu, click **Open**. Navigate to the sample folder that contains the Kerbs4u.sql file, click the **Kerbs4u.sql** file, and then click **Open**.

       Press F5 to run the SQL Server script.

       When you run the script, both a KerbS4uDB database and a **GetUserInfo** stored procedure are created. SQL Query Analyzer returns output that is similar to the output that is in Figure 4 when you successfully run the script.

       Grant rights to the user account that you created in step 5.




**Figure 4: Setting up the Sample Database and Stored Procedure**

In the following configuration steps, you will grant the user account login rights to KerbS4uDB, and then you will add permissions so that the user has permissions to run the **GetUserInfo** stored procedure:

1. Open the SQL Server Enterprise Manager MMC Snap-in.

2. Right-click **Logins**, and then click **New Login**.

3. In the **SQL Server Login Properties New Login** dialog box, click **Windows Authentication**, in the **Domain** box, click the user login domain, in the **Name** box, type the user account that you created in step 5, click **Grant access**, and then, in the **Database** box, click **KerbS4uDB** in the **Database** box.

4. Click the **Database Access** tab, select the **KerbS4uDB** check box, in the **Database roles for KerbS4uDB** box, click the **public** row , and then click **Properties**.

5. In the **Database Role Properties public** dialog box, click **Permissions**.

6. In the **Exec** column, select the **GetUserInfo** check box, and then click **OK**.

**Running the Samples**

The following sections describe how to test the sample scenarios.

**Running Sample Scenario One**

To run sample scenario one:

1. Start Internet Explorer, and then type the URL for the Web service method in the **Address** box. For example, type **http://www.fabrikam.com/kerbextwebsvc/authxws.asmx?op=testkerbsfu**.

2. In the **Enter Network Password** dialog box for digest authentication, type the domain user name and password, and then click **OK**.

3. Click **Invoke** to send a test call to the **TestKerbSfu** Web method. When the call successfully runs, you receive a return string that is similar to the following return string:

```
<?xml version="1.0" encoding="utf-8" ?>
<string xmlns=http://fabrikam.com/>web method user context:
FABRIKAM\jeffpike; AuthenticationType: Digest; Sql server user
context: FABRIKAM\jeffpike</string>
```

**Note**

For security reasons, you can invoke the Web method only if the browser client is running on the same computer that is running the Web service in Windows Server 2003.

The return string indicates that the fabrikam\jeffpike domain user has been authenticated by the **Digest** authentication protocol. The user account that is used to authenticate the SQL Server connection is returned by the **GetUserInfo** stored procedure. To confirm that you successfully used constrained delegation to connect to the SQL Server, view the Security event log on the computer that is running SQL Server to verify that the user account successfully logged on to the network by using the Kerberos authentication package. To view the Security event log, right-click **My Computer**, click **Manage**, click to expand **Event Viewer**, and then click **Security**. Figure 5 shows an example of this logon event. At the bottom of the **Event Properties** dialog box, Transitted Services information is displayed. The Transitted Services information displays the identity of the service that uses constrained delegation. Note that the computer must be configured to audit successful logons to capture this event. You can enable this configuration through the **Audit Policy** option in the Local Security Policy MMC Snap-in.



**Figure 5: Kerberos User Logon Event Log**

**Sample Scenario One Code Walkthrough**

In sample scenario one, the user is authenticated through HTTP Digest authentication. The ASP.NET Framework runtime automatically instantiates a **WindowsIdentity** object for the authenticated user. Subsequently, the Web service method impersonates the user before authenticating to the SQL server. The following code sample provides an example for how you can perform protocol transition with programming code for sample scenario one:

```
[WebMethod]
public string TestKerbSfu()
{

    //
    // impersonate the user before connecting to the SQL database
    //

    WindowsIdentity userIdentity = Context.User.Identity as WindowsIdentity;
    WindowsImpersonationContext impContext = userIdentity.Impersonate();

    string userName = ConnectToSql( sqlServer );

    if( impContext != null )
    {
        impContext.Undo();
    }

}

private string ConnectToSql( string sqlServer )
{
```

```
    ...
    try
    {
        string connString = "integrated security=SSPI;initial
catalog=KerbS4uDB;data source=" + sqlServer + ";Connect Timeout=30";
        connection = new SqlConnection( connString );
        connection.Open();
        ...
     }
}
```

Multiple processing steps occur when the **SqlConnection.Open** function is called:

1. The Kerberos SSP is invoked when the Web service connects to the SQL server. The Kerberos SSP finds an impersonation-level user token for the impersonated user; however, the Kerberos SSP does not find a Kerberos ticket that is available for the user.

2. The Kerberos SSP initiates protocol transition, and then obtains a service ticket to the Web service on behalf of the impersonated user.

3. The Kerberos SSP obtains a service ticket to SQL Server through constrained delegation.

**Running Sample Scenario Two**

To run sample scenario two, start Internet Explorer, type the URL for the Web application in the **Address** box (for example, type **http://www.fabrikam.com/KerbExtWebApp**), and then press ENTER.

The Web application redirects the browser to the user logon page. Note that the sample code does not validate the user and password information that is typed in at the logon page. Typically, your Web application that uses ASP.NET form-based authentication verifies that the user name and password are valid before allowing the user to use the Web application. After the user logs on, the browser is redirected to the Web applications default page. The default page is where you type the user account that you want the Web application to use for protocol transition. When protocol transition is successful, the next Web page allows you to connect to the SQL server through constrained delegation. When the SQL stored procedure returns, the browser displays both the impersonated and delegated user identities. Similar to sample scenario one, you can examine the Security event log on the computer that is running SQL server to look at the Kerberos-related events that were logged during the SQL connection authentication.

**Sample Scenario Two Code Walkthrough**

There are two code samples that you should examine for scenario two. The first code sample is used to initiate protocol transition into Kerberos and the second code sample demonstrates how to use constrained delegation to connect to the SQL server.

The following code sample shows how you can initiate protocol transition by creating a new instance of the **WindowsIdentity** class object:

```
private void DoProtocolTransition_Click(
object sender, System.EventArgs e)
{
    ...
    WindowsIdentity newIdentity = new WindowsIdentity( UserAccount.Value );

    if( null == newIdentity )
    {
        Msg.Text = "Cannot do protocol transition, please try again";
        return;
    }

    //
    // Store this user as part of the HTTP session state
    //

    HttpContext.Current.Session.Add( "KerberosIdentity", ( object )newIdentity );
    ...
}
```

Note that the Web browser user can specify the Kerberos security principal for protocol transitioning in this Web application. Typically, a Web application may use a user mapping table to determine the Kerberos security principal that will be used as the user identity for Kerberos protocol transition. The previous code shows that a new **WindowsIdentity** class object is created

for the user that you specify in the **UserAccount.Value** variable. After the class object is created, the **WindowsIdentity** object is stored in the ASP.NET HTTP session state. The Web application can retrieve this **WindowsIdentity** object from the HTTP session state when it needs to impersonate the user and use constrained delegation to authenticate the connection to the SQL server. The following code sample shows how the Web application performs protocol transition and constrained delegation:

```
...
//
// Get the user's Kerberos identity from the HTTP session state
//
WindowsIdentity user = ( WindowsIdentity )HttpContext.Current.Session[
                       "KerberosIdentity" ];
...
//
// impersonate the user before connecting to the SQL database
//

WindowsImpersonationContext impContext = user.Impersonate();
string userName = ConnectToSql( sqlServer );

Identity.Text = "Impersonated user identity = " + user.Name + "; SQL
Server user identity = " + userName;
if( impContext != null )
{
    impContext.Undo();
}
...
```

# Summary (Kerberos Protocol Transition and Constrained Delegation)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section summarizes and defines some of the key points for the new protocol extensions.

**Protocol Transition**
The following items summarize the protocol transition implementation details that are described in this document:

- The protocol transition extension can be initiated in two ways:

    - Programmatically through the **LsaLogonUser** function or the **WindowsIdentity** class object. This approach is recommended when the initial user-authentication mechanism is not one of Windows-integrated authentication protocols. The process must have the **Act as part of the operating system** privilege for the function to create an impersonation-level token. If the process does not have the **Act as part of the operating system** privilege, an identification-level token is created.

    - Using existing Windows-integrated authentication protocols for the initial user logon procedure with credentials, and using Kerberos protocol transition for subsequent user authentication to other services. The Kerberos SSP initiates the protocol transition extension if the user token that you created for initial logon is an impersonation-level token.

- You must run the operating system process that initiates protocol transition on a computer that is running Windows Server 2003.

- You can use protocol transition both in and across Active Directory forests.

- When you use protocol transition in an Active Directory forest, all domain controllers that are located in the user and service accounts trust path and that are used for the protocol transition process must be running Windows Server 2003. To meet this requirement, use one of the following methods:

    - Upgrade all of the domain controllers that are in the domain that is participating in protocol transition to computers that are running Windows Server 2003. This approach requires that all of the domain controllers in the domain are running Windows Server 2003.

      **-or-**

    - Upgrade all of the domain controllers in the Active Directory sites that are participating in protocol transition to computers that are running Windows Server 2003. This approach does not require that all domain controllers that are in the domain are running Windows Server 2003. This approach is recommended if you cannot upgrade all of the domain controllers in the domain to Windows Server 2003, and if you can isolate the applications that require the protocol transition extension to a certain number of Active Directory sites for which you can upgrade all of the domain controllers.

- When you use protocol transition across Active Directory forests, both forests must be operating at the Windows Server 2003 forest functional level and two-way forest trust must be established between the forests.

**Constrained Delegation**
The following items summarize the constrained delegation implementation details that are described in this document:

- The constrained delegation extension is affected by the following domain policies that are set in Active Directory:

    - The requesting service must be trusted for constrained delegation (the T2A4D flag must be set for the requesting service) if Kerberos is not used as the initial user authentication protocol.

    - The requesting service must be authorized to delegate to the target service (the target service must be on the requesting services A2D2 list).

- You cannot use constrained delegation between services whose accounts are in different domains. All domain controllers in the domain must be running Windows Server 2003, and the domain must be operating at the Windows Server 2003 functional level. The accounts of users accessing the services do not have to be in the same domain as the services.

# Conclusion (Kerberos Protocol Transition and Constrained Delegation)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Delegation is a security pattern that occurs frequently in n-tier applications. At the time when this document was published, Kerberos was the only widely-adopted authentication protocol that possessed the delegation property. In this document, the Kerberos extensions on computers that are running Windows Server 2003 are discussed. The document also describes how the extensions allow many applications to use the authentication protocol. The sample code illustrates how n-tier Web services and applications can use the new feature.

Protocol transition provides application designers with increased flexibility and security by enabling applications to support different authentication mechanisms at the user authentication tier, and by switching to the Kerberos protocol for security features, such as mutual authentication and constrained delegation, in the subsequent application tiers.

Constrained delegation gives administrators the ability to specify and enforce application trust boundaries by limiting the scope where application services can act on a users behalf. This flexibility to constrain a services authorization rights helps improve application security design by reducing the opportunities for compromise by untrusted services.

# Role-Based Access Control for Multi-tier Applications Using Authorization Manager

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

**Author:** Dave McPherson, Microsoft Corporation

**Acknowledgements:** Thanks to Jason Rush, Praerit Garg, Don Schmidt, Paul Leach, and Doug Bayer for their support.

In the Windows Server 2003 family, Authorization Manager introduces a new model for application authorization on the Windows platform. Authorization Manager provides an easy-to-use authorization framework for line-of-business applications where you must grant access to users according to the users role within an organization or application. To address this need, applications often create custom implementations of authorization mechanisms, which increases development costs and administrative complexity because of divergent authorization models. Authorization Manager gives applications a role-based access control framework that provides manageable administration and natural development for Web-based or line-of-business applications.

**In This White Paper**

Authorization Policy Storage

Application Administration

Auditing

Application Design Process

Developing Authorization Manager Applications

Performance

Environment-specific Design Issues

Internet Information Services 6.0 URL Authorization

Scenario: Web Expense Application

Scenario: Using Authorization Manager with Custom Principals

**ACLs and line-of-business applications**

Windows operating systems have supported the use of ACLs in access control for applications since the introduction of the Private Object Security API in Microsoft Windows NT Server 4.0. In the ACL model, you can attach discretionary access control lists (DACLs) to a securable object and you make access decisions by calling the **AccessCheck** application programming interface (API), which looks at a users group memberships in a token and compares them to the contents of the ACL to determine if the user has the requested access. The ACL model is ideal for many types of applications. Resource managers that have well-defined persistent objects, such as the NT system registry, use the ACL model appropriately to provide object-level access control. In these types of applications, access control decisions can always be made in terms of a user requesting access to an existing object. In the future, applications that use the ACL model should use AuthzAPI, which was introduced in Windows XP and the Windows Server 2003 family. It provides performance and flexibility enhancements, compared to the Windows NT Private Object Security APIs. For more information about AuthzAPI, see the Microsoft Platform Software Development Kit.

A distinct authorization problem exists for some line-of-business (LOB) applications, such as a Web expense reporting application or a shopping application. For these applications, the authorization decisions do not determine access to well-defined persistent objects. Instead, they verify a workflow or perform multiple distinct operations, such as querying a database and sending mail. In LOB applications, access decisions are often made on business logic, such as the amount submitted in an expense application or verification of workflow completion, and not just based on token group membership. Applications that don't have well-defined persistent objects have no place to put an ACL, so the ACL model can be harder to apply in these applications.

**Role-based access control**

Traditional access control administration models are based on objects. In these models, access control is specified at the object or object container (for example, in an ACL), and the administrator must go to the object to query and specify access to the object. These models require administrators to translate the organizational authorization policy into permissions on objects:

Each object has a list of access permissions that is granted to various users and groups within an organization. Role-based access control (RBAC) simplifies access control administration and provides better manageability in enterprise environments by allowing permissions to be managed in terms of user job roles.

Some of the goals of role-based access control can be accomplished using groups. A group corresponds to an employee role, and application administrators can specify the permissions that are needed by the role by granting the group permission in a DACL for the object. As object collections grow, the number of places where administrators need to manage permissions grows. Diligent use of resource groups and user groups can help minimize this effort, but this requires consistent practice and coordination among administrators and precise definitions of resource groups. These processes slow down the administrative process, so administrators often avoid them.

Additionally, querying the granted access to a particular group or role across an application becomes more difficult as the number of objects grows. To accurately determine what permissions are granted to a user or group, administrators have to examine the permissions on every object. While inheritance appears to simplify this, the ability of each object to avoid inheriting permissions still makes it necessary to look at each object to fully understand the authorization policy. Since there are too many objects to query, the state of access control with respect to a particular group or user is sometimes rarely verified.

Role-based access control attempts to allow administrators to specify access control in terms of the organizational structure of a company. RBAC does this by creating a new object called a role. You assign a user to a role to perform a job function. Unlike groups, however, a role defines the authorization permissions on some set of resources. In the RBAC model, the administrator uses the role to manage permissions and assignments. For example, a company may create a role called Sales Manager that has the permissions that a sales manager needs for their job. When sales managers are hired, they are assigned the Sales Manager role and instantly have all required permissions for that job. When they leave the position of sales manager, they are removed from the Sales Manager role and no longer have Sales Manager access. Since the role allows access to be granted in terms of a company's organizational model, it is more intuitive and natural for administrators to specify access control. Figure 1 identifies the relationship between the role, the users, and the permissions. In this model, the role is the object to which permissions are granted, and the users are assigned to the role.

**Figure 1: Role-based access**




RBAC maps user job roles to application permissions so that access control administration can be accomplished in terms of a users job role. The RBAC system translates a users role membership to application permissions. Since the permissions are granted at the role, permissions can be queried and changed at the role without examining the specific resources. In most environments, once the role permissions are established, changes to a roles permissions will be rare, compared to changes in assignments to the role. This means that administrators will have to set up roles, such as Employee, Manager and Administrator, but once roles are created, administrators will manage membership in the roles; not permissions on objects.

**Authorization Manager conceptual model**

Authorization Manager is a set of COM-based runtime interfaces that allows applications to easily manage and verify a clients requests to perform application operations. Additionally, Authorization Manager provides a Microsoft Management Console (MMC) snap-in that application administrators can use to manage user roles and permissions.

Authorization Manager:

- Simplifies application access control administration in line-of-business applications.
- Provides a simplified and natural development model.
- Enables flexible and dynamic authorization decisions.

Typically, you define access control models with the terms Subject, Resources, and Permissions. A subject is often a user operating upon resources. Access is controlled by a set of permissions that allow the subject to perform specific operations on the resource. Administration of access control rarely deals directly with these constructs but more often involves management of collections of subjects, resources, and permissions. You usually manage subjects or users in groups, but you manage resources in application-specific collections, such as directories or organizational units. Permissions may be granted directly, but are often grouped into higher level permissions such as Full Control. These collections make authorization management possible in environments where there are many subjects, resources, or permissions. Application administrators must use these collections to make an authorization policy for the application so that each subject has the appropriate permissions to use the application resources.

Traditional resource manager applications store authorization policy with the resources that the policy applies to. For example, an ACL on a file is stored by the resource manager and is logically associated with the file. In the Authorization Manager model, authorization policy data is stored separately from the object in an authorization policy store. An application identifies the operations or tasks that the application can perform and declares them in the authorization store when the application is installed. Administrators manage the authorization policy by defining roles in terms of the tasks and operations that are needed to perform a job in an organization. The definitions of roles, tasks and operations, and the assignment of users and groups to the roles are all stored in the authorization store. Applications query the authorization policy at run time to confirm that a client is authorized to perform a requested operation on a resource. Authorization Manager provides APIs for managing authorization policy and validating access control and a user interface for administrators so that they can manage the authorization policy store. The following figure illustrates the composition of the authorization store that Authorization Manager provides for querying applications and managing the authorization store. It also illustrates the relationships between each Authorization Manager object.

**Figure 2: Authorization Manager object relationships**




**Definitions**

This section introduces the objects that Authorization Manager uses to manage authorization policy.

**Operation.** A low-level permission that a resource manager uses to identify security procedures. Several operations may be required to perform a meaningful task. Operations are often not exposed or meaningful to administrators. An example of an operation could be WriteAttributes or ReadAttributes.

**Task.** A collection of low-level operations. The purpose of the task is to determine which low-level operations are required to do some unit of work that is meaningful to administrators. An example of a task might be Change Password. Tasks can also contain other tasks. For example, a task called Manage User Accounts may include the tasks Change Password, Reset Password, Disable Account, and so on.

**Role Definitions.** A collection of permissions that are needed for a particular role, where permissions can be tasks or operations. Notice that this definition is similar to that of a task.

**Role.** The set of permissions that users must have to be able to do their job. The role is applied to a set of objects and users are assigned to the role.

**BizRules.** Also called authorization scripts. Scripts that are attached to a task object that is run at the time of the access request. It can use information that is only available at runtime, such as "time of day" or "dollar amount requested," to make an authentication decision.

**Scope.** A collection of objects or resources with a distinct authorization policy. A scope can represent a physical collection, such as a folder, or a more flexible collection of resources, such as *.doc. Applications can use the scope as necessary to group resources, and must be able to map the requested resource to its scope at the time the application checks the access for a user.

**Application groups.** Groups that are applicable only to an authorization store, an application within an authorization store, or a scope within an application. Authorization Manager implements two types of application groups.

**Application basic groups.** A subset of application groups. A list of members (Active Directory Users or groups or other application groups) that is maintained for a specific set of applications, a single application, or a scope within an application. Application groups also can identify non-members, which allows for exceptions where a large group can be used with a smaller group or particular user excluded.

**LDAP-query groups.** A subset of application groups. Groups that are defined by an Lightweight Directory Access Protocol (LDAP) query on a given Active Directory users account attributes. These provide flexible group membership to be defined in terms of user account attributes and that remains up to date with the users account object.

# Authorization Policy Storage

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

When an application that uses Authorization Manager is initialized, it loads the authorization policy information from the store. Authorization Manager provides support for storage of authorization policy in Active Directory or in .xml files.

| **Important** |
| --- |
| Because the authorization policy information is stored by Active Directory or in an NTFS file, it is secured with the security capabilities of either Active Directory or NTFS, and the auditing support for changes that are made directly against the storage (without using the Authorization Manager runtime interfaces) is handled by the system on which the store is located. Additionally, administrators on the system that contains the authorization policy store have a high degree of access to the store. For this reason, the authorization policy store must be located on a trusted system. |

**Active Directory Authorization Manager storage**

To use Active Directory to store the Authorization Manager policy, the domain must have a functional level of Windows Server 2003. Active Directory in the Windows Server 2003 family contains schema updates for the Authorization Manager objects. When using the Active Directory store, Authorization Manager creates Active Directory objects for the store itself and child objects for each application group, application, operation, task, role, and scope. The scope object can contain tasks, roles, and groups that were created in that scope.

When using the Active Directory authorization policy store, applications control when application objects are loaded. When an application connects to the policy store, the stores global attributes (including store-level groups) are cached locally, along with header information for each application. When an application initializes its authorization policy using the **OpenApplication** method, the applications attributes, roles, tasks, operations, application groups, and header information for each scope are loaded into memory. Once an application is loaded, it remains cached in memory until the application object is freed. Scopes are loaded on demand; either the first time the scope is specified in a call to **AccessCheck** or when the scope is explicitly loaded using the **OpenScope** method. Once a scope is loaded, it remains cached in memory until the application object is freed.

Active Directory does not support enforcement of transactional writes on multiple objects or attributes, so it is possible that the store could be edited simultaneously by two administrative applications. In this situation, a store could become corrupted. Because of this, it is recommended that applications regularly back up their authorization store. For more information about Active Directory backup and restore tools, see "Active Directory" in the Windows Server 2003 Help.

**Deployment considerations for Active Directory administrators**

Application administrators can benefit from using Active Directory to store the authorization policy. The Active Directory store allows applications to take advantage of the existing domain infrastructure to provide the application with storage, availability, redundancy, and distribution of the authorization policy. Active Directory administrators who deploy Authorization Manager should consider the following information.

**Domain mode requirements**

The Authorization Manager Active Directory Store requires that the domain be in Windows Server 2003 functional level.

**Authorization Manager Active Directory impact on infrastructure**

Authorization Manager is a runtime that runs on the application server. It does not run any applications or services on domain controllers. Active Directory can be used to store the authorization policy data.

Applications that use the policy download application policy information when it initializes and downloads application scope policy as needed. When the application reads the policy information, the impact of the read on a domain controller will vary. The size of the authorization policy for an application or scope will vary with the number of roles, tasks, operations, and groups in the scope. Application groups vary in size, but it is recommended that you do not use application groups for large groups so that application performance is not affected. The size of the groups and roles can vary, but may be similar to Active Directory groups. The read of a group or role costs as much in processing time as the read for the members of an Active Directory group with the same amount of members. Application groups can have the same maximum size as Active Directory groups but are usually smaller than Active Directory groups because they are only used by an application. When possible, we recommend that people use Active Directory groups instead of application groups when groups with more than 5000 members are needed.

**Proximity requirements**

Applications need high-quality access to Authorization Manager policy. Even though individual applications may perform

acceptably in low-bandwidth scenarios, unless they have small authorization stores, domain controllers should generally be within the local area network (LAN) of the application server.

**Replication**
The authorization policy store is replicated to all domain controllers in the domain in which the authorization policy is stored and Authorization Manager objects are referenced in each global catalog server in the forest.

A separate domain or forest could be used to store an Authorization Manager policy. You may choose to do this to minimize replication traffic within a domain, to minimize load on domain controllers in an account domain, or to have separate administrators manage the Authorization Manager policy store. Using a separate domain or forest requires that appropriate trust be established between the domain where user accounts are stored (the account domain) and the domain the Authorization Manager policy is stored in (the Authorization Manager storage domain.) The trust configuration that is required depends on what domain the accounts of the administrators of the Authorization Manager policy are in and which domain the application server and any back-end resource servers are in.

**Deploying an Authorization Manager store in Active Directory**
The following gives step-by-step instructions on deploying an Authorization Manager store in Active Directory.

1. Choose the Authorization Manager Active Directory store location.

   To deploy Authorization Manager in an Active Directory store, you must choose a location for the store. It is recommended that the administrator create a new organizational unit (OU) in the domains Program Data container and create a store object in the new OU. This recommendation is not enforced. If the application has data in a container other than Program Data, you can put the Authorization Manager store there.

   Authorization Manager cannot be placed in a non-domain naming context, which is also called an application partition.

2. Create an Authorization Manager store.

   Applications can create the Authorization Manager policy store in Active Directory themselves through the Authorization Manager API, or administrators can create a store using the Authorization Manager MMC. In order to do this, the user or application attempting to create a store must have Create Child Object permissions in the container that was chosen in the previous step and they must know the distinguished name of that container.

   For more information about creating a store using the Authorization Manager MMC snap-in, see Work with Authorization Stores in Authorization Manager Help.

3. Roles and authorization store access

   If the application administrator creates the store using a program, then they should already have permission to manage the store. If the store is created manually by an Active Directory administrator, permission to manage the store must be given to the application administrator who is using the store. Use the Authorization Manager MMC to grant the application administrator access to manage the authorization policy store. The authorization policy within the store is managed by the application administrator. The Authorization Manager Active Directory store allows the application administrator to delegate management of applications and scopes within the store.

   Since application service accounts that use the authorization policy must be able to read the policy store, the application service account must be included in the Readers role in for the appropriate store or application or scope within a store.

4. Enable the Windows Authorization Access group.

   | **Note** |
   | --- |
   | This step is only needed if the application designers use either the **InitializeClientContextFromName** and **InitializeClientContextFromStringSID** methods to create client context objects within the application. Applications that use **InitializeClientContextFromToken** do not require this step. If you are not sure which APIs the application server will use, the application administrator or designer should know. |

   These methods attempt to read the users token-groups-global-and-universal attribute in Active Directory to get the user's Active Directory group membership information. To read this information, the application service account must

have Read access on this attribute on each user object.

If the domain is configured for Pre-Windows 2000 Compatible Access mode (that is, the Everyone group is in the in the Pre-Windows 2000 Compatible Access group), then the Authenticated Users group will have access to call the **InitializeClientContextFromName** or **InitializeClientContextFromStringSID** methods by default. If the domain is set for Windows 2000 and Windows Server 2003 Compatible Access then, by default, only the Authenticated Users group will be in the Pre-Windows 2000 Compatible Access group. In this case, application service accounts will have the required access to call the **InitializeClientContextFromName** or **InitializeClientContextFromStringSID** methods. If the domain has been locked down (that is, Active Directory administrators have made the default access permissions more restrictive) it is possible that the Pre-Windows 2000 Compatibility Access group will be empty, in which case application servers may fail when calling the **InitializeClientContextFromNameOfSID** interfaces.

In this situation, to grant access to the service account, it is recommended that, first, you add the service account to the Windows Authorization Access group, and then grant the Windows Authorization Access group access to each users token-groups-global-and-universal attribute if it does not already have that access. For more information about the Windows Authorization Access group (which is new for the Windows Server 2003 family), search for the term in the Microsoft Knowledge Base at http://support.microsoft.com.

**Authorization Manager storage with XML files**

Authorization Manager supports storage of the authorization policy in an .xml file that is stored on an NTFS volume. The XML store can be kept on the same computer as the server application using Authorization Manager APIs or it can be stored remotely.

In order to support renaming of objects, the XML format contains globally unique identifiers (GUIDs). Because of this, you should not edit the .xml file directly. Furthermore, the XML schema is not published currently. You can edit the store through the Authorization Manager MMC or through the Authorization Manager interfaces which are available to scripting languages, such as Microsoft Visual Basic Scripting Edition (VBScript) and JScript. The XML store does not support delegation because access to the XML file is controlled by the discretionary access control list (DACL) on the file which grants or restricts access to the entire contents of the file.

It is recommended that application administrators back up their authorization store regularly. The NTFS file system does not support applications issuing a sequence of separate write operations as a single logical write to a file when multiple applications write to the same file. (These operations are sometimes known as transactional operations.) Since an Authorization Manager policy store could be edited simultaneously by two administrative applications, a store could become corrupted.

# Application Administration

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

After installing an application, the job of the administrator is to add any new roles required, verify the permissions on the initial roles (if present), and assign users to roles. If the application creates a complete set of tasks with meaningful names, the administrator's job of setting up initial roles is much easier.

The point of role-based access control is make administration more intuitive and manageable. To do so, administrators must devise a strategy for deploying roles within the application that resembles the user job descriptions. Once the roles are established, the administrative focus is primarily managing memberships within the roles, so it is very important that you verify each role's permissions. Generally, it is safer to grant each role only those permissions that are needed immediately and then add permissions as needed, rather than granting too many permissions and then removing unneeded permissions later.

**Managing scopes**

Often, the applications manage the creation of scopes and the administrator then creates and manages the roles within a scope. This way, the administrator does not have to know the applications scope syntax and the application logic that is used to associate user access request to a scope.

When applications allow administrators to create scopes, the application must provide documentation so that the administrator understands how to create scopes that are named correctly, since scopes are defined by each application and the application will map each access request to the appropriate scope.

Scopes allow administrators to group application resources into physical or logical collections and specify different authorization policy in each scope. Since tasks and role definitions can be created so that they are global to the application (in other words, visible to all scopes), roles that are created for different scopes can use the same global role definition and have different memberships. This prevents the need to redefine the permissions in each scope. Such a technique is useful when a role or task, such as Administrator or Full Control will have the same permissions in each scope but different people assigned to it.

| **Important** |
| --- |
| Authorization Manager applications have a default scope. Permission that is granted to roles in the default scope apply in all scopes in the application. For this reason, the applications default scope should only be used for role assignments that apply to all scopes in the application. |

**Role inheritance**

Since the process of establishing roles involves assigning user permissions, it requires special attention. To minimize the validation and analysis that is required, a hierarchal role model can be designed that leverages the work that was done on establishing existing roles when creating new roles.

**Figure 3: Role hierarchy**




Role definition involves tasks and operations that define the permissions for a role. Using role definitions allows roles to be defined hierarchically, using existing roles to create new ones.

**Delegation of administration**

When using the Active Directory Authorization Manager policy store, Authorization Manager supports delegation of administration at the authorization store, application, and scope levels. This allows higher-level administrators to give limited

access to others to manage some subset of the authorization policy. The store, application, and scope objects can each have a set of Administrators and Readers. Administrators of a store, application, or scope can perform all operations that are needed to manage the authorization policy within the delegated scope, and Readers can only read the authorization policy. Additionally, users who have been delegated authority of a scope or application must be added to the Delegated User administrative role of the store. This allows the delegated user to view groups that are created at both the application and the store level.

When a scope has been delegated (because one or more principals are in the scopes Administrator role), BizRules cannot be added to tasks or role definitions that were created in the scope. Delegated scopes can use BizRules that are attached to tasks and role definitions at the application level. Store-level administrators can also globally prevent all BizRules from running by setting the ScriptEngineTimeout store level value to 0. After doing this, **AccessCheck** assumes that existing BizRules will return as false and access will not be granted through task or role definitions containing a BizRule.

**Using Authorization Manager application groups**

Authorization Manager introduces two types of groups which can be used by Authorization Manager applications to facilitate applications that need to maintain groups of security principals: application basic groups and LDAP query groups. (These are discussed in more detail later in this paper.) Authorization Manager application groups can be global to an Authorization Manager policy store and can be used by all applications within the store, specific to an application and used by all scopes in the application, or specific to a single scope. Some applications need to manage groups of users without having to depend on updates to domain groups that may be managed by different people. Domain administrators may have requirements to keep the number of groups low without managing application-level groups. Latency may exist between the application administrators request and the actual changes to the Active Directory groups. Additionally, some groups are only needed for certain applications and, if the application can manage the group memberships, the domain administrators does not have to.

**Application basic groups**

An application basic group is similar to a traditional Windows domain group in that it is a collection of member security principals, but it is unlike traditional groups in that it also contains a non-member or subtractive list of security principals. This allows for exceptions that traditionally use Deny access control entries (ACEs) on objects. Security principals that are added to the non-member list will never be a member of the group, even if they are directly or indirectly granted membership in the member list. Application basic groups store security identifiers (SIDs) for each member and nonmember, so growing application groups will result in growth of the Authorization Manager store. Since applications load portions of the authorization store into memory, large group and role memberships can increase the time required to initialize an Authorization Manager store, application, or scope object.

As application groups become very large:

- Performance for initialization of stores, initialization of applications, and calls to **AccessCheck** are affected.
- The application requires more memory.
- There is more load on domain controllers when applications initialize and refresh the loaded policy.
- The membership of the application groups tend to change, which requires more frequent refreshing of the cache.

Because of this, it is strongly recommended that, as application groups grow to around 2000 members, that they be converted to Active Directory groups. (For more information about performance, see Performance, later in this paper.) A users Active Directory group memberships are determined when the user logs on or when the user context is created and do not affect the time required to initialize an Authorization Manager store, application, or scope.

**LDAP query groups**

Often, groups are maintained in Active Directory for collections of users that have similar attributes on their Active Directory accounts. For example, users who work in a given cost center may be in a group for the cost center, even though the cost center information is maintained as an attribute on each users account in Active Directory. When the user moves from one cost center to another, the administrator must update both the cost center group and the cost center attribute on the user object. To leverage the users properties and relieve administrators of maintaining the same information in multiple places, Authorization Manager provides the LDAP query group.

Membership in an LDAP query group is determined by using an LDAP query on a given user object. This allows applications to use these user attributes as groups in roles and frees administrators from managing memberships of these groups. Application administrators control membership by creating an LDAP query group in terms of the properties on user accounts. The following are some examples of coding for LDAP query groups.

All members of the Managers group:

```
(memberOf= CN=Managers,OU=Distribution Lists,DC=nwtraders,DC=com)
```

All members over 18 years old.

```
(age>=18)
```

All users in Germany

```
(country=Germany)
```

All users over 18 years old in Germany

```
(&(age>=18)(country=Germany))
```

LDAP query groups require that the connected users account is maintained in Active Directory, but they are also supported when using the XML-based Authorization Manager policy store. LDAP query groups are evaluated when **AccessCheck** is called, therefore, the **AccessCheck** call will take more time to complete when the LDAP query is used to populate roles, because evaluating the query will require a LDAP search in Active Directory. While membership in an LDAP query group is cached for the life of the client context, initial **AccessCheck** calls are slower if many LDAP query groups must be evaluated.

**Enterprise-wide roles**

Since Authorization Manager roles are defined in terms of the permissions that a role will require within an application, a role cannot span across applications. To have roles in distinct applications have the same memberships, use common groups to populate application roles. This can be done using Windows groups or application groups if the applications are managed in the same Authorization Manager policy store. These groups are visible to each application and can contain the members for roles that are defined in multiple applications. Then, in each application, these groups can be assigned as members of the roles. By using Active Directory global or universal groups to contain role membership, application roles can have similar memberships across applications and Authorization Manager policy stores in an organization.

# Auditing

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Authorization Manager provides two types of auditing.

- Authorization Manager runtime auditing
- Authorization Manager policy store change auditing.

**Authorization Manager runtime auditing**
Authorization Manager runtime auditing audits:

- Application initialization
- Client context initialization or deletion
- All calls to **AccessCheck** with pass and fail audits

Auditing can be configured at the store or application layers. Authorization Manager runtime auditing is enabled through the **GenerateAudits** property of the **AzAuthorizationStore** interface, which is stored in the Authorization Manager policy store.

In order to use Authorization Manager runtime auditing, the system that is running Authorization Manager must have enabled auditing, and the context that calls the Authorization Manager interfaces must have the SE_AUDIT_NAME or the Generate Audits privilege. If an application requires that auditing is enabled, it can require the audit privilege by passing the AZ_AZSTORE_FLAG_AUDIT_IS_CRITICAL flag to IAzAuthorizationStore::Initialize. This requires the calling application to have SE_AUDIT_PRIVILEGE or else IAzAuthorizationStore::Initialize fails.

**Authorization Manager policy store change auditing**
Authorization Manager policy store change auditing triggers audits when there is a change to the Authorization Manager auditing policy. The Active Directory supports auditing of Active Directory objects so that the Active Directory store changes can be audited using the Active Directory system auditing where auditing settings can be configured at the store, application, and scope levels. Changes to the XML store can be audited as a whole by using the system access control list (SACL) that is on the objects.

# Application Design Process

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Part of application design is making good choices for roles, tasks and operations. This section gives you information that will help you in your design.

**Identifying operations**
An operation is a low-level permission that represents privileged actions or capabilities of an application. Operations should be created for each routine, query, method, and so on that comprises a discrete application function which requires access control. An operation alone may not be enough to perform a high-level task and may be required for more than one high-level task, but the operation itself is always executed as a unit and must be secured.

For example, an operation such as **ReadOrderInfo** alone may not be enough for the high-level task Process Order, but the same routines may be needed for another high-level task, such as Query Order Status. Creating an operation for the set of routines that are required to use the order controls allows finer permission specification. The more precisely the operations are defined, the more flexibility you have when managing permissions. However, operations that are defined too precisely make administration more complex.

For many resource managers, operations correspond to securable routines, procedures, or queries that operate on data or resources. The operations may be low-level and meaningful only to the application developer. To create permissions that are meaningful to administrators, the application developer can group low-level operations into Authorization Manager task objects.

**Figure 4: Operations and tasks**




**Defining tasks**
As stated above, the application should define operations for each routine that will run as a single unit and must be secured. While this allows the application to grant access more precisely, multiple operations may be required to perform a high-level task that the application exposes to the user. In this case, you should determine which operations are required to perform the high-level task and create an Authorization Manager task object for it. This requires some research on the part of the administrator. Task objects are designed to simplify role management by providing meaningful, high-level permission sets that correspond to the application tasks that an administrator would grant a user permission to perform. Therefore, it is important for applications to install with a complete set of tasks required to control access for all uses of the application. Figure 4 illustrates how you can use task objects to combine permissions into high-level tasks.

**BizRules**
The Authorization Manager task object has the additional ability of dynamically qualifying the permissions granted in a task, based on the result of an attached VBScript or JScript associated with the task. This allows for the access control decision to take into account runtime data, such as the dollar amount of an expense that has been requested or the inventory of a requested item. BizRules are run during the **AccessCheck** call, and will run in the context of the Web application in the thread calling **AccessCheck**. If the BizRule returns success, the user receives the requested operations which were associated with the task. If the BizRule fails, the client is not allowed to perform the operations that are associated with the task. The user may be able to perform those operations through a separate task. Parameters are sent to the BizRule through the **AccessCheck** call, which takes corresponding name and value arrays as its parameters. For example, the first element in the *varParameterNames* array is the name corresponding to the first element in the *varParameterValues* array. The BizRule parameter's names array

(*varParameterValues*) must be ordered accordingly so that names and values have the same index.

An applications authorization policy can have many tasks. Each task may or may not have associated BizRules, and new tasks and BizRules may be added by application administrators. At design time, application developers do not know if a call to **AccessCheck** will need to specify BizRule parameters (since BizRules may not be used for any given call to **AccessCheck**), so the application must send in the parameters potentially needed by BizRules. To facilitate BizRule developers, applications should define and publish the set of BizRule parameters that are sent into each **AccessCheck** call.

The type of data that **AccessCheck** should send as BizRule parameters includes information such as user name and restriction data. Restriction data may include items such as amount limits, account or ID numbers, or a users manager. Any data that may be useful in determining whether or not to grant access at runtime can be used in each **AccessCheck** call and must be included in the applications list of published BizRules parameters.

BizRules require some scripting knowledge, and therefore it is not appropriate for most administrators to create and modify them. Usually, BizRules should be developed and shipped with the application or provided by the application vendor or other developer later. Since each BizRule that is run invokes the Windows Script Engine, BizRules should be computationally small tasks, such as comparing given parameters or querying a database.

**Sample BizRules**

Here is a BizRule written in JScript that insures that the time of day is between 9:00 and 17:00, using a 24-hour clock:

```
AzBizRuleContext.BusinessRuleResult = false;
dt = new Date();
hour = dt.getHours();
if (hour > 9 && hour < 17)
{
   AzBizRuleContext.BusinessRuleResult = true;
}
```

To improve performance of BizRules that are called frequently, the results of the BizRule are cached for the life of a client context. Because of this, BizRules are time-sensitive, as in the above example. Applications cache client contexts for appropriately short amounts of time.

Here is a Microsoft Visual Basic Scripting Edition (VBScript) BizRule to make sure an amount that is passed into the BizRule as a parameter named ExpAmount is less than 500:

```
Dim Amount
AzBizRuleContext.BusinessRuleResult = FALSE
Amount = AzBizRuleContext.GetParameter("ExpAmount")
if Amount < 500 then AzBizRuleContext.BusinessRuleResult = TRUE
```

When a BizRule refers to a parameter, **AccessCheck** caches the parameters, values, and result for a particular BizRule for the life of a client context. This allows improved performance in subsequent **AccessCheck** calls on the same client context that references the same BizRule with identical values. This means that, if BizRule results can change in subsequent **AccessCheck** calls without changing the BizRule parameters (as in the time of day BizRule described above), then the cached value may not be the correct value. For this reason, applications using BizRules should free the client context periodically.

**Initial roles**

Some applications have a small number of hard coded roles, such as Administrator, Reader, and Writer. Such applications would call the Authorization Manager API to create these initial roles at startup time.

**Scoping**

Authorization Manager scope objects are collections of resources, and can be either simple collections such as C:\My Documents or more collections, such as *.doc. The application must take into consideration what scoping logic it will use. A scope provides great flexibility to the application in how to apply different access control to different resources. Depending on what the application is controlling access to, the collections may vary considerably in form. For example, if an application performs operations on other computers, then a scope may represent collections of computer objects and be named something like Computers in Sales Dept or "Computers Running Windows XP. The only requirement is that the application must be able to map a user request for access to a resource to the scope containing the resource when the access validation is performed.

In addition to logical collections of resources, scopes give applications the flexibility to address complex authorization problems, such as incorporating a "work on behalf" feature into an application. A "work on behalf" feature allows a user to grant other users permission to do work on their behalf. For example, suppose Bob wants to allow James to do work on their behalf. In an application, this could be implemented by creating a scope for Bob in which a role would be defined that specifies the subset of Bobs permissions that James is allowed to use. (James would be assigned to this role.) When James connects and

requests permission to perform an operation on behalf of Bob, an **AccessCheck** operation is performed against the Bob Work on Behalf scope to verify that James has permission to perform the operations on behalf of Bob.

Since scopes are defined by each application and the application will map each access request to the appropriate scope, the application must provide documentation to allow the administrator understand how to create scopes that are named correctly.

**Installation**

When an application is installed, it can install authorization policy information into a new or existing authorization store by calling the Authorization Manager API. If the application installs its authorization policy into an existing store, the installation process must give the installing user a way to specify the location of the existing store. After initializing the authorization policy store, the install application uses the **CreateApplication** method to create an application object in the policy store. Then, for each operation, the installation application calls the **CreateOperation** method. Initial tasks, scopes and roles are installed using the **CreateTask**, **CreateScope**, and **CreateRole** methods, respectively. The following VBScript creates a store and installs the authorization policy for a simple expense application.

```
'--- Initialize the admin manager object
Dim pAzManStore
Set pAzManStore = CreateObject("AzRoles.AzAuthorizationStore")
'--- Create a new store for expense app
' 0 = Open store for Access Checking.
' AZ_AZSTORE_FLAG_CREATE   = 0x1,
' AZ_AZSTORE_FLAG_MANAGE_STORE_ONLY = 0x2, // cannot do access
checking.
' AZ_AZSTORE_FLAG_BATCH_UPDATE      = 0x4,
pAzManStore.Initialize 1+2+4, "msxml://C:\AzPolicy.xml"
pAzManStore.Submit
Dim App1
Set App1 = pAzManStore.CreateApplication("Expense")
App1.Submit
'--- Create operations -----------------------
Dim Op1
Set Op1=App1.CreateOperation("RetrieveForm")
Op1.OperationID = CLng(61)
Op1.Submit
Set Op1=App1.CreateOperation("EnqueRequest")
Op1.OperationID = CLng(62)
Op1.Submit
Set Op1=App1.CreateOperation("DequeRequest")
Op1.OperationID = CLng(63)
Op1.Submit
Set Op1=App1.CreateOperation("UseFormCotnrol")
Op1.OperationID = CLng(64)
Op1.Submit
Set Op1=App1.CreateOperation("MarkFormApproved")
Op1.OperationID = CLng(65)
Op1.Submit
Set Op1=App1.CreateOperation("SendApprovalNotify")
Op1.OperationID = CLng(66)
Op1.Submit
'--- Create tasks ------------------------------
Dim Task1
Set Task1=App1.CreateTask("Submit Expense")
Task1.AddOperation CStr("RetrieveForm")
Task1.AddOperation CStr("EnqueRequest")
Task1.AddOperation CStr("UseFormCotnrol")
Task1.Submit
Set Task2 = App1.CreateTask("Approve Expense")
Task2.AddOperation CStr("MarkFormApproved")
Task2.AddOperation CStr("SendApprovalNotify")
Task2.AddOperation CStr("DequeRequest")
Task2.BizRuleLanguage = CStr("VBScript")
Task2.BizRule = "Dim Amount" & vbnewline  & _
           "AzBizRuleContext.BusinessRuleResult= FALSE" & vbnewline &_
           "Amount = AzBizRuleContext.GetParameter( " & Chr(34) &_
                             "Amount" & Chr(34) & ")"  & vbNewLine &_
           "if Amount < 500 then AzBizRuleContext.BusinessRuleResult=
TRUE"
Task2.Submit
'--- Create role definitions ------------------------------
Set Task3 = App1.CreateTask("Expense Admin")
Task3.AddTask CStr("Approve Expense")
```

```
Task3.AddTask CStr("Submit Expense")
Task3.IsRoleDefinition = TRUE
Task3.Submit
Set Task4 = App1.CreateTask("Expense User")
Task4.AddTask CStr("Submit Expense")
Task4.IsRoleDefinition = TRUE
Task4.Submit
'--- Create initial scopes and roles ------------------------------
'--- only one scope in this app (we may instead choose to use no scope)
Dim Scope1
Set Scope1 = App1.CreateScope("AllRoutines")
Scope1.Submit
Set RoleA=Scope1.CreateRole("Expense Administrator")
RoleA.AddTask("Expense Admin")
RoleA.Submit
Set RoleB=Scope1.CreateRole("Expense User")
RoleB.AddTask("Expense User")
RoleB.Submit
'--- Demo - add everyone to ExpenseUser Role --------------------------
RoleB.AddMemberName("Everyone")
RoleB.Submit
```

**Application model**

Applications that use Authorization Manager use the trusted subsystem application model described below. Note that, while the two models imply distinct application designs, based on the tradeoff, a particular application may choose to use a hybrid model that contains a trusted subsystem middle tier design and use impersonation to maintain back-end auditing or to interface with legacy back-end resource managers.

**Impersonation model**

Windows 2000 and earlier versions of Windows NT Server supported the Windows NT impersonation model. In this model, server applications obtain a token for the connected client and impersonate the client token before attempting a secure operation, such as opening a file. The ACL on the file is compared to the users token groups to determine whether or not a user has permission to open the file. While the Windows Server 2003 family enhances support for the impersonation model through the Kerberos extensions of protocol transition and constrained delegation.

**Trusted subsystem application model**

Authorization Manager adds support for a new trusted subsystem application model to the Windows security infrastructure.

**Figure 5: Impersonation model and trusted subsystem model**




In the trusted subsystem model, the application server account only has enough access so that it can perform all the operations it exposes to clients. When a client requests an operation, the middle-tier application server authorizes the client request based on authorization policy that is specific to the application. If the client has permission to perform the requested operation, the server performs the operation on behalf of the client. In this model, clients do not have direct access to resources. While the impersonation model has its strengths and is appropriate in many cases, the trusted subsystem model has some advantages over the impersonation model.

**Management of permissions**

In the impersonation model, each client users account is used to access the back-end resource, so the DACLs of all resources

must be maintained by granting each user the appropriate level of access. When the number of resources that are secured separately grows, especially when stored on separate back-end computers, management of the DACLs becomes more tedious. For example, you might organize resources into units that are secured in the same way, such as protecting all files in a subtree the same way or organizing users into groups with similar permissions.

In the trusted subsystem model, only the middle-tier service account is used to access the back-end resources; therefore, the DACLs can be simplified by granting only the service account sufficient access to perform all required operations on the objects. A service account should not be given Full Control rights or be a local administrator, but should have just enough permission to perform the applications complete set of operations. Application designers can document the level of permissions the application need on back-end resources. This reduces the management of DACLs, because only the service account requires access to them. The service is given permissions to the resources, so you must trust that the service will not accidentally let one user access resources of another.

**Rights ion**

With the impersonation model, assigning DACLs to resources to grant users the ability to perform high-level tasks can be difficult. High-level tasks such as Submit Expense or Query Inventory may require several operations on back-end resources. *Rights ion* is the process of determining precisely which low-level resource permissions are required for high level application operations. It can be time consuming, particularly when many secured resource types are used.

In practice, administrators sometimes grant excessive access to resources to save time. A trusted subsystem application model allows the administrator to grant access only to the service account. Since the access to the resources is broadly given to the service account, determining the correct access is easier: usually the server gives just enough access to perform all operations that the server application may need (as defined by the application designer). The administrator does not have to define and maintain different levels of permissions on DACLs.

**Connection pooling**

You often can achieve better scalability by connection pooling. Since the impersonation model maintains the user context through the request to a resource server, a separate connection to the resource server is usually required for each user request. If you create a connection for every client, you prevent reuse of a single connection to a resource server, such as a SQL database. A trusted subsystem server application authorizes the client's request and accesses the remote resource for the client in the context of the applications service account. Since a new connection is not required, the application server can use the same connection with resource servers to service requests from different clients.

**Controlled access points**

Some applications enforce rules and workflow in the application server. In the impersonation model, each client is given appropriate access to each resource. Since the client has access to directly to the resource, they could potentially access the resource though a method other than through the intended applications which could cause undesired access and manipulation of resources.

In the trusted subsystem model, only the application service account has access to resources. This means that users cannot directly gain access to resources through a tool or API. In this model, the application server has more control over how the user sees and manipulates the resources.

**Auditing**

Auditing allows administrators to examine which users have attempted to gain access to particular resources. Auditing is most authoritative if the audits are generated by the same application that manages the resources. The impersonation model does this by maintaining the users context when it requests access on each system. This allows the remote system to authoritatively log the user and the requested access.

When you use a trusted subsystem application model, the back-end resource managers generate audits that log the servers service account as the account requesting access, not the client user on whose behalf the operation was performed. Authorization Manager provides runtime auditing when the access validation is performed (in other words, when the **AccessCheck** API is called), but these audits are generated at the application server. Because of this, mapping an access audit on a back-end resource to the user who made the request requires that the audit logs on the application server be correlated with back-end audits. The application server audits can often be more informative than audits made on the back-end resource. From the application server, you can know more about the high-level task that the client is requesting, so a greater degree of the users intent can be recorded in the audits. On the back-end resource managers, only low-level operation requests can be audited, so it may be difficult to determine the users higher-level activities.

**Other differences**

In the impersonation model, the service account in which the application server runs may have little or no access to resources. To access resources on behalf of the client, the application server must impersonate the clients security context, which the application server can obtain when the client connects to it. When you limit the permissions of the server context, the server is limited to the permissions of the connected client when a successful attack takes place which compromises the server. If an attacker wants to gain broad access to back-end data, they must compromise the server and wait for users to connect. As each

user connects to a compromised server, the access of the attacker grows in proportion to the connected users permissions. If a highly-privileged account logs on to the server, such as the applications administrator, the attacker quickly gains permissions to the application data. If client users who have access to resources other than the applications log on to the server, the attacker also gains access to those resources because of the ability to impersonate the client to those resources.

A trusted subsystem server starts out with a high degree of access to application resources so, if the server is compromised, the attacker has broad access to these resources. But since the client is not impersonated, the attacker is limited to the access of the service context (the application's resources). While the attacker gains broad access to application resources, the attacker cannot access non-application resources.

# Developing Authorization Manager Applications

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Server applications use Authorization Manager runtime interfaces to manage connections to the authorization policy store and validate client requests

**Initialization**
When the application runs, it initializes the Authorization Manager policy store using the **IAzAuthorizationStore::Initialize** interface. The application should have Read access to the authorization store, but does not need Write access. After connecting to the store, the application calls **IAzApplication::Initialize** to initialize an interface to its specific application within the policy store.

**Client connection**
When a client connects to the application, an Authorization Manager context must be created for the user. To do this, use the **IAzApplication::IntiializeClientContextFromToken**, **IAzApplication::AzInitializeClientContextFromStringSID**, or the **IAzApplication::InitializeClientContextFromName** method. Generally, if a token is available, you should use a token when creating a context. A token is usually the result of a Windows logon and will contain logon information, such as whether it was an interactive or a network logon. Context creation from a token is usually faster because it does not have to query a domain controller for group information. In rare cases, it is possible that role assignment may have been given depending on these logon properties, so when possible, use the FromToken form of InitializeClientContext APIs. For example, when a user logs on to the network interactively, the NT AUTHORITY\INTERACTIVE group is added to the resulting token. Using the **InitializeClientContextFromName** (or **SID**) method, an **IAzClientContext** object is created by looking up the users group assignments in Active Directory. Since the user has not logged on, the context will not contain groups such as the NT AUTHORITY\INTERACTIVE group that identifies the logon type. In this case, if membership in a group or role is given to the interactive group, a user context that is created using **InitializeClientContextFromName** does not receive the membership. It is recommended that you do not use these logon property groups to secure resources or specify memberships in groups or roles, in which case **InitializeClientContextFromName** and **InitializeClientContextFromSID** can be used.

The **AzInitializeClientContextFromStringSID** method creates an Authorization Manager context from a given SID in textual form. This behaves in a similar manner as the **InitializeClientContextFromName** method. When the AZ_CLIENT_CONTEXT_SKIP_GROUP flag is used, the **AzInitializeClientContextFromStringSID** method does not attempt to determine the group memberships of the given SID. The resulting client context only contains the specified SID. If the **IAzAccessCheck** method is called from this client context instance, role membership is only granted if the specified SID is used as a member of a role or group assigned to a role.

**Roles query and authorization**
After a client context is created, applications have the option of querying role information about the user in order to present a user interface based on the users role memberships. For example a Web-based application may render a Manager user interface for people in the Manager role, and an employee role for everyone else. The **IAzClientContext::GetRoles** method enumerates the users role memberships. The set of roles that is returned from **GetRoles** is based only on explicit membership assignment; **GetRoles** does not evaluate BizRules.

When the user makes a request, the application maps the request to the required operations and determines the scope in which the user is making the request. Then the application calls **AccessCheck** sending the scope, requested operations, BizRule parameters, and object name. The scope determines the users role membership, which determines whether or not access is granted to the user.

# Performance

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Application initialization and access validation performance can vary, depending on how applications use features such as XML storage, BizRules, and LDAP query groups. Application designers have several options to consider in meeting performance requirements.

**Application initialization performance**
Load times for large stores improve when you use the Active Directory store rather than the XML Authorization Manager storage provider. When an application initializes a connection to an Authorization Manager policy store, authorization policy is loaded into the address space of the application. The amount of policy that is loaded depends on the store type.

- The Active Directory Authorization Manager policy store allows Authorization Manager application objects to be loaded and unloaded on demand and allows the Authorization Manager scope objects to be loaded only when needed. (For more information, see Active Directory Authorization Manager Storage, earlier in this paper.)

- The XML Authorization Manager storage provider loads the entire authorization policy into memory.

Additionally, the time it takes to initialize increases as the number of objects within the Authorization Manager policy store increases or as the number of members in Authorization Manager application basic groups increases. To minimize the size of application groups, use Active Directory groups when possible and then use these Active Directory groups in Authorization Manager application groups or roles.

**AccessCheck performance**
Using the Authorization Manager **AccessCheck** method provides a flexible way to optimize performance using BizRules and LDAP query groups.

The **AccessCheck** method:

1. Identifies scopes that apply to the **AccessCheck** routine, both within the specified scope and the global scope.

2. Enumerates the set of roles in the scopes from step 1.

3. Discards roles that do not refer to requested operations. If **RoleForAccessCheck** is specified, all roles other than those specified are discarded.

4. Verifies the client security contexts membership in each role.

5. If the client is a member of the role in step 4, **AccessCheck** records all requested operations that are assigned to the role as granted.

6. Repeats steps 3-5 for each role in applying scopes.

The above logic is repeated a maximum of three times.

1. On the first pass, **AccessCheck** ignores any tasks containing BizRules and application groups containing LDAP queries. Both BizRules and LDAP queries can take a long time to compute and can require network operations, so the first pass takes place in case they are not needed.

2. In the second pass, **AccessCheck** processes BizRules, but still ignores LDAP query groups. BizRules are supposed to be simple and may not involve network operations, so they are likely to be faster to process than LDAP queries.

3. The third pass processes all data, including LDAP queries, using the cached results from previous iterations. **AccessCheck** terminates as soon as access is granted to all of the requested operations.

Since **AccessCheck** is optimized to evaluate static memberships ahead of dynamic memberships, and tasks without BizRules ahead of tasks with BizRules, you can optimize performance by using static assignments where possible.

Authorization Manager application groups are not evaluated until **AccessCheck** is called and membership in a role is queried. (This process of group evaluation is also called late-bound.) For LDAP query application groups, once membership in an application group is established for a user, the users membership in that LDAP query group is cached for the life of the user context object. Because of this, future **AccessCheck**s that involve group evaluations speed up as group memberships are resolved and cached.

As a best practice, BizRules should be kept computationally small. BizRules offer the greatest degree of flexibility in an **AccessCheck** call, and therefore offer the greatest potential to decrease performance. Network bound BizRules decrease the performance of **AccessCheck**. Often, this cannot be avoided, for example, if a database needs to be queried to check a credit history, but such queries will delay the result to the user.

# Environment-specific Design Issues

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Authorization Manager is designed for applications using the trusted subsystem application model, where access checks are done in the context of the server, and ACLs on back-end server resources only need to grant access to the application server service account. (For more information, see Trusted Subsystem Application Model, earlier in this paper.) To use a trusted subsystem model, usually the application server will run in the context of a dedicated service account. Since the application needs to read Authorization policy information from the Authorization Manager policy store, the security context that the application runs in needs Read permission to the store. You can do this by designating the dedicated service account as a Reader of the Authorization Manager policy store. To do this, you can use the Authorization Manager MMC (Microsoft Management Console).

If impersonation is used in the application, you should not impersonate a client while calling Authorization Manager API. The only exception to this is when you wish to initialize a client context from the impersonated thread context. In that case, call the **IAzApplication::InitializeClientContextFromToken** method and specify 0 for the *ullTokenHandle* parameter. The other Authorization Manager interfaces read the Authorization Manager authorization policy store (which contains user role information) and the **IAzApplicatin::AccessCheck** function runs BizRules and resolves LDAP query groups. Since a user requires a high degree of access to be able to read authorization policy and perform BizRule script actions, you should not call these functions while impersonating the security context of the user. Each Windows programming environment gives you different options to use to manage impersonation and call Authorization Manager from a service account for the application.

**ASP.NET**

Application servers written in ASP.NET can use the Authorization Manager API through the .NET Interop assembly that ships in the \AuthMan directory in the .NET Framework. To create a trusted subsystem model in ASP.NET, you need to run the application server in a service account. To configure your ASP.NET application server to run in a service account, you create a separate IIS 6.0 worker process which runs in the security context of the service account created for the application server. You can then configure your ASP.NET application to use this dedicated IIS worker process. For more information about IIS 6.0 worker processes, see the Microsoft Internet Information Services page on the Microsoft Web site at http://www.microsoft.com/technet/prodtechnol/windows2000serv/technologies/iis/default.mspx. Alternatively, you can use ASP.NET to configure the context in which an application runs using the Web.config file for your ASP.NET application. In that case, you should configure the ASP.NET application to run as the dedicated service account.

When your ASP.NET application is initialized, it can call **AzAuthorizationStore.Initialize** to connect to the Authorization Manager store. It then calls **AzApplication.Initialize** to load an initial percentage of the applications authorization policy.

At run time, when a client connects, you have the option of initializing a user context from the clients account name or from a token representing the client. The token can be retrieved in ASP.NET by way of the ASP.NET **HttpWorkerRequest** object. Authorization Manager can also initialize a context from a clients SID, though that is not usually available in ASP or ASP.NET. It is recommended that you initialize the Authorization Manager context from a token when possible because, when initializing from a token, you do not have to potentially query the users account object. That may require network-bound queries to Active Directory.

The following code calls the Authorization Manager API through ASP.NET, using a .NET Interop assembly to approve an expense approval. This code is a simple Web expense application that runs in a dedicated service account. The Authorization Manager policy store for this application can be created from the VBScript install script that appeared earlier in this paper.

```
<%@ Page Language="C#" Debug="True"%>
<%@ Assembly Name ="Microsoft.Interop.Security.AzRoles"%>
<%@ Import Namespace="Microsoft.Interop.Security.AzRoles"%>
<%@ Import Namespace="System.Runtime.InteropServices" %>
<html>
   <head>
   </head>
   <body>
      <%
      Microsoft.Interop.Security.AzRoles.AzAuthorizationStoreClass
AzManStore =
            new
Microsoft.Interop.Security.AzRoles.AzAuthorizationStoreClass();
      // Keep the authorization store in a safe place out of the Web
space.
      AzManStore.Initialize(0,
```

```
@"msxml://D:\SecureDir\MyWebAppsAzStore.xml",
                           null);
      IAzApplication azApp =
AzManStore.OpenApplication("expense",null);
      HandleRef token = new HandleRef(this, ((HttpWorkerRequest)
((IServiceProvider)Context).GetService(typeof(HttpWorkerRequest))).GetU
serToken());
//-------------- Create Client Context --------------
      IAzClientContext context =
azApp.InitializeClientContextFromToken(
                       (UInt64)token.Handle,
                       0);
      object[] scope = new Object[1];
      scope[0] = "";
      object[] operations = new Object[1];
      operations[0] = 55;
//--------------- Do access check --------------------
//Set up BizRule params. To optimize performance, name/value pairs must
//be placed in array alphabetically
      Object[] BRNames = new Object[1];
      BRNames[0] = "ExpAmount";
      Object[] BRValues = new Object[1];
      BRValues[0] = Convert.ToInt32(Request.QueryString["Amount"]);
      object[] results = (object[]) context.AccessCheck("Approve",
                                                        scope,
                                                        operations,
                                                        BRNames,
                                                        BRValues,
.null,null,null);
      bool bAuthorized = true;
      foreach (int iResCode in results)
      {
         if ( iResCode != 0 ) // zero = no error
         {
           bAuthorized = false;
           break;
         }
      }
      Response.Write((string)context.UserSamCompat + "<br>");
      Response.Write(
             "Approval of " + Convert.ToString(BRValues[0]) + " was");
      if (bAuthorized) {
         %><font color="#008000"><b>APPROVED</b></font><%
      } else {
         %><font color="#FF0000"><b>DENIED</b></font>.<%
      }
      %>
   </body>
</html>
```

**ASP**

Applications that use Authorization Manager usually use the trusted subsystem model. By design, IIS Application Service Providers (ASP) applications impersonate the context of each connected client. To use a trusted subsystem model in ASP, you have two options:

**Option 1: Use Internet Information Services 6.0 URL authorization and impersonate the Web expense service account.**

Create a separate IIS 6.0 worker process which runs in a security context that is created for the ASP application. This runs an instance of IIS and ASP in a dedicated service account for the ASP application. Normally, ASP applications run impersonating the authenticated user. IIS 6.0 URL authorization allows you to configure the ASP application to run in the context of the IIS worker process. Since you set up the IIS worker process to run in the application service account, we can use IIS 6.0 URL authorization to run the ASP application in that security context. IIS integrated authentication still authenticates the client, but instead of impersonating the client, IIS 6.0 URL authorization causes the ASP application to run in the service account. When a client connects and is authenticated by IIS, you can use the ASP server variable AUTH_USER to retrieve the authenticated name of the client and call **AzApplication.InitializeClientContextFromName** to create an Authorization Manager context for the client. For more information about IIS 6.0 URL authorization, see IIS 6.0 URL Authorization later in this paper.

The following code is an ASP application, using Authorization Manager to validate an expense approval. This code assumes that the ASP application is impersonating a dedicated service account, not the client security context. The application is configured to impersonate the IIS worker process using IIS 6.0 URL authorization, as described above.

```
<%@ LANGUAGE="VBSCRIPT" %>
<HTML>
<HEAD>
</HEAD>
<BODY>
<%
Dim CCHandle
Dim Results
Dim Names(5)
Dim Values(5)
Dim Scopes(5)
Dim Operations(10)
Dim AzManStore
Dim App
Dim UserName
Set AzManStore = CreateObject("Azroles.AzAuthorizationStore")
'Set myUser = Session.GetObject("LDAP")
UserName = Request.ServerVariables("LOGON_USER")
AzManStore.Initialize 0,
"msxml://D:\Inetpub\wwwroot\MyWebAppsAzStore.xml"
set App = AzManStore.OpenApplication ("Expense")
'--------------- Create Client Context --------------
Set CCHandle = App.InitializeClientContextFromName(UserName)
'--------------- Do access check --------------------
'Set up BizRule params. To optimize performance, name/value pairs must
be
'placed in array alphabetically
Names(0) = "ExpAmount"
Values(0) = 0
Names(1) = "Param2 for BizRule - Name"
Values(1) = "Param2 for BizRule - value"
Scopes(0) = Empty
Operations(0) = 55
If Len(Request.QueryString("amount")) > 9 Then
    Response.Write "Amount field is too long, please try again."
Else
   Set reg = New RegExp
   reg.Pattern = "\D+" ' Look for something other than 0-9's.
   if reg.Test(Request.QueryString("amount")) = False then
      Values(0) = Clng(Request.QueryString("amount"))
      '-------- AccessCheck ----------------------------------------
      Results = CCHandle.AccessCheck("Approve",
                                     Scopes, Operations, Names, Values)
      If Results(0) = 0 Then ' Zero = NO_ERROR
         Response.Write "Approval of " & Values(0) & " was accepted."
      Else
         Response.Write "Approval of " & Values(0) & " was denied."
      End If
      '-------------------------------------------------------------
   Else
      Response.Write "Must enter number, please try again"
   End if
End If
Set CCHandle = Nothing
set App = Nothing
Set AzManStore = Nothing
%>
</BODY>
</HTML>
```

**Option 2: Create custom COM objects and interfaces to control impersonation.**

Create a custom COM object that allows you to control impersonation from your ASP application. After that, set up the IIS worker process to run in a service account as described above and use the custom COM object to call the RevertToSelf API to return to the IIS worker process context (the ASP applications service account). Then, call the ImpersonateSelf API to impersonate the context of the IIS worker process. (ASP expects to run applications under impersonation.) It is important to save the client token in this step because it will be used later.

By impersonating the IIS worker process from the custom COM object you can call Authorization Manager APIs to initialize the Authorization Manager policy store and the Authorization Manager application in the context of the ASP applications service account (the security context that you configured the worker process to use). This process of initializing the Authorization Manager policy store and application objects might only be done once when the ASP application initializes.

When a client request is made, you must create an Authorization Manager client context to use for checking the clients access. To do this, you can use the ASP server variable AUTH_USER to get the clients user name and call **AzApplication.InitializeClientContextFromName**. But you have the option of using your custom COM object to create a client context using the logon token that was created when IIS authenticated the client user. To do this, expose a function in your custom COM object to retrieve the token that you stored above and call the ImpersonateLoggedOnUser API. This will impersonate the client security context. Then, create an Authorization Manager client context using the **AzApplication.InitializeClientContextFromToken** method with (0,0) as parameters to create an Authorization Manager context from the current thread token. Since you are impersonating the client, the current thread token represents the client and the Authorization Manager client context created from this token will also represent the client. After doing this, use your custom COM object to re-impersonate the ASP application service account process again, as explained above.

This works because, when **AzApplication.InitializeClientContextFromToken** is called with the value of **0** for the token handle, the API creates an Authorization Manager context that is derived from the token of the calling thread. Since, in this case, we are impersonating the client user, the Authorization Manager context that is created represents the client.

Implementing a custom COM interface to manage impersonation requires additional work, but using a custom COM object to control impersonation has the benefit of using the clients token that is maintained in IIS to create a context as described. Methods that use the **AzApplication.InitializeClientContextFromName** interfaces result in querying the user account object in Active Directory. This can result in slower performance than when you use a token, because the token already has this information.

**C, C++, and Visual Basic**

You can control impersonation in native C, C++, or Visual Basic by calling various context management APIs, such as ImpersonateLoggedOnUser, OpenThreadToken, or RevertToSelf. In these languages, the application server can be installed as a Windows System service and run in a dedicated account for the service application. When the application initializes, it first calls **AuthorizationStore.Initialize** to connect and load information from the Authorization Manager policy store, and then **AzApplication.Initialize** to create a COM interface to the application policy within the policy store

When a user logs on to the application, it authenticates the client using whatever means it chooses. If a native Windows authentication mechanism is used, such as Security Service Provider Interface (SSPI), Distributed Component Object Model (DCOM), Remote Procedure Call (RPC), or Named Pipes, then the application server can use the resulting token to create an Authorization Manager context with one of two methods:

- Impersonating the client by using the appropriate impersonation API, such as ImpersonateNamedPipeClient or ImpersonateSecurityContext in SSPI, and calling **IAzApplication::InitializeClientContextFromToken**, specifying **0** for the token handle value.

- Querying the clients token handle, such as **QueryContextAttributes** in SSPI, and using that handle in the call to **IAzApplication::InitializeClientContextFromToken**.

If non-native Windows authentication is used, then applications can map the authenticated client to a Windows account and then use the IAzApplication::InitializeClientContextFromName API, specifying the Windows account name that the client was mapped to after authentication.

# Internet Information Services 6.0 URL Authorization

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0, which comes with the Windows Server 2003 family, integrates with Authorization Manager to implement IIS 6.0 URL authorization. This provides Web application administrators with the ability to control access to URLs based on custom user roles, LDAP queries, and BizRules.

Authorizing user access to Web pages in IIS can require management of many DACLs on resources used by Web applications. Resources for Web applications may include Web page files, database records, registry keys, and so on. Maintaining DACLs requires administrators to know precisely what back-end permissions are needed on each object to perform meaningful tasks in the Web application. IIS 6.0 URL authorization allows administrators to simplify access management by authorizing user access to the URLs that comprise a Web application. When a client requests a URL, IIS 6.0 URL authorization validates the users access based on the user roles. This allows administrators to simplify access control management by controlling all user access to URLs instead of each DACL on each object. The Web application can further restrict access to resources and operations by using the Authorization Manager role-based framework.

IIS 6.0 URL authorization is implemented as an ISAPI interceptor .dll, which is configured as a wildcard extension for an application, virtual directory, or URL. When an application, virtual directory, or URL is configured to use IIS 6.0 URL authorization, each request to a URL is routed to the URL authorization ISAPI interceptor. The IIS 6.0 URL authorization ISAPI interceptor will use the Authorization Manager runtime to authorize access to the requested URL. To do this, the URL (application, virtual directory, or single URL) must be associated with an Authorization Manager policy store that contains the authorization policy for the URL. Once the client has been authorized to access the URL, the IIS 6.0 URL Authorization ISAPI passes the request to the appropriate handler for the URL, such as ASP, ASP.NET ISAPIs, or the static file handler.

**Figure 6: IIS URL authorization**




Using IIS 6.0 URL authorization, an administrator can control access to URLs based on information that is only available at runtime. For example, if you had a Web page that is only viewable to employees in a given cost center or to employees of a certain age, you can assign roles to the correct users based on LDAP queries that check the cost center or age attributes on a users object. If employees can only access certain pages on certain days of the week or during a certain time of day, a BizRule can be created which grants access to the URL based on these values or any value that can be asserted at runtime, including IIS server variables.

**Setting up IIS 6.0 URL authorization**

To use IIS 6.0 URL authorization, set the following metabase attributes on the application, virtual directory, or URL:

- **AzEnable.** Enables URL authorization for the virtual directory, application, or URL that corresponds to the entry in the metabase.

- **AzStoreName.** Associates an Authorization Manager store with the virtual directory, application, or URL.

- **AzScopeName.** Associates the virtual directory, application, or URL with a scope. This scope will be the name of a scope in the IIS 6.0 URL authorization application in the Authorization Manager policy store that the **AzStoreName** attribute refers to. If it specifies either an empty string or no scope at all, the default scope of the IIS 6.0 URL authorization is used.

- **AzImpersonationLevel.** Determines the impersonation behavior for the application. This allows you to configure the Web application to impersonate the client user, the IIS worker process, or the IUSER_* account for the worker process. Each setting significantly changes the environment and implied design of the Web application. Set the value to 0, 1 or 2 corresponding to the desired impersonation style.

  **0** Authenticated client is impersonated. Some applications require client impersonation, such as ASP 2.0 applications. Use this if you want to use URL authorization for applications that expect to run while impersonating the client, such as ASP 2.0 applications.

  **1** The IIS worker process identity is impersonated. This is used when the application runs as a service account, authorizes each client request itself, and performs requested operations on behalf of the client in the trusted service context. For more information, see Trusted subsystem application model, earlier in this paper.

  **2** The anonymous user account is impersonated. This allows you to use URL authorization to authorize the clients access to the URL while still running the Web application as the anonymous account to limit all access of the client through the Web application.

The following script uses ADSI to set these properties on the IIS metabase entry for an application, virtual directory, or URL.

```
'  SetUrlAuth.vbs
'  Configuring UrlAuth with a Vdir.
'  This script uses ADSI to configure a vdir in the IIS 6.0 metabase.
to use
'  IIS 6.0 URL authorization.
'
'  This script assumes that the IIS vdir corresponding to the
application has
'  already been created.
'
'  After running this script to verify the settings, use the Adsutil
tool like so
'
'  inetpub\adminscripts\adsutil enum w3svc/1/root/MyApp
'  (where MyApp is the app name)
'  Use the Authorization Manager MMC snap-in to do the following:
'     Create an application called "IIS 6.0 URL Authorization"
'     Create an operation called "AccessURL" with an operation ID of 1
'     Create a scope corresponding to the AzScopeName specified to this
script.
'       This scope name will be used to manage access to URLs in this
Vdir
'     Make sure that the IIS worker process has read access to the
'       AzMan policy store. You can do this using the AzMan UI (right
click
'       store and select options and then the security tab
Set objArgs = WScript.Arguments
If objArgs.count < 4 then
  wscript.echo "Usage: SetUrlAuth VDirPath AzScopeName AzStoreName AzEnable
[ImpersonationLevel]"
  wscript.echo ""
  wscript.echo "Example:"
  wscript.echo " SetUrlAuth w3svc/1/root/MyApp MyApp
msxml://d:\inetpub\wwwroot\AzStore.xml True 1"
  wscript.echo ""
  wscript.echo "Run with 'cscript' command in cmd.exe to avoid msg
boxes"
Else
wscript.echo objargs(0)
'
' Get interface to the Vdir object
'
```

```
DIM iis
set iis = GetObject("IIS://localhost/" & objArgs(0))
' Set scope name - this scope will preside over URLs referring to this vdir
' This scope must exist in the corresponding Authorization Manager store.
'
iis.AzScopeName = objArgs(1)
'
' Specify location of Authorization Manager policy store
'
iis.AzStoreName = objArgs(2)
'
' Enable urlauth for this Vdir
'
iis.AzEnable = objArgs(3)
'
' Set Impersonation Level
'
If objArgs.count > 4 then
   iis.AzImpersonationLevel = objArgs(4)
End if
'
' Write settings to Metabase
'
iis.SetInfo
End if
```

When Web pages are requested, URL authorization will authorize the client based on the Authorization Manager authorization policy, if the client is allowed to access the URL. The security context that is configured by the **AzImpersonationLevel** attribute to retrieve the URL data (such as an .htm file) requires the correct IIS directory security and NTFS permissions. IIS URL authorization allows the IIS directory security and NTFS permissions to be more easily maintained using the trusted subsystem model. For more information, see "Trusted Subsystem Application Model," earlier in this paper.

When IIS 6.0 URL authorization is configured, you can identify an Authorization Manager policy store by the **AzStoreName** attribute in the IIS metabase entry for the application, virtual directory, or URL. To manage the authorization policy, us the following procedure.

1. Click the **Start** button, click **Run**, type **Azman.msc**, and then press ENTER.

2. In the console tree, right-click **Authorization Manager**, and then click **Open Authorization Store**.

3. In **Store Name**, type the path to the authorization store location, which is identified by the **AzStoreName** attribute, and then click **OK**.

For more information about opening an authorization store, see Authorization stores and applications in the Authorization Manager MMC Help.

IIS 6.0 URL authorization for this store is managed with the IIS 6.0 URL Authorization application. The **AzScopeName** attribute in the metabase entry will be an Authorization Manager scope in the IIS 6.0 URL Authorization application. Use this scope to manage access to the corresponding URL. When configuring an application, virtual directory, or URL for URL authorization, a scope must be created in the authorization policy store with the same name as that specified in the corresponding metabase entry's **AzScopeName** attribute.

The following figure shows the Authorization Manager MMC with an authorization policy store that is used to manage access to a Web Expense application and Web Application 2. This store also contains authorization policy for these applications.

**Figure 7: Authorization Manager managing two Web applications**




Authorization policy for IIS 6.0 URL authorization is stored in an Authorization Manager policy store in either Active Directory or in an .xml file. IIS 6.0 URL authorization allows the authorization policy for URLs to be maintained in the same authorization policy store as that of Web applications that use Authorization Manager to control access to application tasks and operations. Maintaining the authorization policy in the same store allows administrators to manage access to URLs and Web application resources from the same Authorization Manager MMC or custom user interface. This provides a common point of access control administration for URLs and applications and allows for the same Authorization Manager groups and LDAP queries to be used to populate roles for both URL authorization and application authorization.

**Configuring IIS 6.0 URL Authorization Step by Step**

The following steps set up IIS 6.0 URL Authorization for a Web application called WebApp, which is installed and configured in IIS as an application in the *Systemroot*\InetPub\WWWroot\WebApp folder:

| **Note** |
| --- |
| Integrated Windows Authentication is required for URLs using IIS URL Authorization. By default, IIS Web applications and virtual directories enable Integrated Windows Authentication. For information about configuring Integrated Windows Authentication, see Integrated Windows Authorization in Internet Information Services online Help. |

**Set up IIS: Create your application folder under Default Web Site, and then disable anonymous access to the site.**

1. Click **Start**, click **Run**, type **%systemroot%\System32\InetSrv\IIS.msc**, and then press ENTER.

   You can also click **Start**, click **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.

2. In the console tree, double-click **Internet Information Services**, double-click the computer name, double-click **Web Sites**, double-click **Default Web Site**, right-click the folder corresponding to your Web application, click **Properties**, and then click **Create**.

3. Click the **Directory Security** tab and then, in **Authentication and Access Control**, click **Edit**.

4. Confirm that the **Enable Anonymous Access** check box is cleared, confirm that the **Integrated Windows Authentication** check box is selected, click **OK**, and then click **OK** again.

**Set wildcard configuration properties to point to URLAuth.dll**

1. In **Internet Information Services (IIS) Manager**, right-click the folder for the Web application which you created, click **Properties**, and then click **Configuration**.

2. In **Wildcard application maps (order of implementation)**, click **Insert**.

3. In **Add/Edit Application Extension Mapping**, click **Browse** and, in **Files of type**, select **All files (*.*)**.

4. Browse to the *Systemroot*\System32\InetSrv folder.

5. Click **URLAuth.dll**, click **Open**, click **OK**, click **OK**, and then click **OK** again.

   If you get an error saying This executable path is already used then URLAuth.dll is already configured as a wildcard application extension mapping.

**Add URLAuth.dll as a new Web service extension**

1. In the console tree, double-click **Internet Information Services**, double-click the computer name, and then click **Web Service Extensions**.

2. In the details pane, click **Add a new Web service extension**.

3. Click **Add**, click **Browse**, browse to *Systemroot*\System32\InetSrv\, click **URLAuth.dll**, click **Open**, and then click **OK**.

4. In **Extension Name**, type **URL Authorization**.

5. Select the **Set extension status to Allowed** check box, and then click **OK**.

**Set up an authorization policy store in Authorization Manager**

1. Click the **Start** button, click **Run**, type **Azman.msc**, and then press ENTER.

2. In the console tree, right click **Authorization Manager**, and then click **Options**.

3. Confirm that **Developer mode** is selected, and then click **OK**.

   Now we will create a new store. For this example, we will put the store in C:\ and name the store "MyStore.xml."

4. In the console tree, right-click **Authorization Manager** and click **New Authorization Store**.

5. Click **XML file** and, in **Store name**, type **C:\MyStore.xml**.

6. (Optional) In **Description**, type information about the new store.

7. Click **OK**.

8. In the console tree, right-click **MyStore.xml** and click **New Application**.

9. In **New Application**, in **Name**, type **IIS 6.0 URL Authorization**.

10. (Optional) In **Description**, type information about this usage of IIS URL Authorization.

11. Click **OK**.

12. In the console tree, double-click **Authorization Manager**, double-click the store name, double-click the application name, double-click **Definitions**, right-click **Operation Definitions**, and then click **New Operation Definition**.

13. In **New Operation Definition**, in **Name**, type **AccessURL**, and then, in **Operation number**, type **1**.

**Set up the scope for the application**

1. In the console tree, right click **IIS 6.0 URL Authorization**, and click **New Scope**.

2. In **Name**, type **WebApp**, and then click **OK**.

3. In the console tree, double-click **IIS 6.0 URL Authorization**, double-click **Definitions**, right-click **Role Definitions**, and then click **New Role Definition**.

   This role definition could also be created in the WebApp scope.

4. In **Name**, type **Viewer**, and then click **OK**.

5. In the console tree, double-click **Role Definitions**.

6. In the details pane, right-click **Viewer**, and then click **Properties**.

7. Click the **Definition** tab, click **Add**, and then click the **Operations** tab.

8. Select the **AccessURL** check box, click **OK**, and then click **OK** again.

9. In the console tree, under **IIS 6.0 URL Authorization**, under **WebApp**, right click **Role Assignments**, and then click **Assign Roles**.

10. Select the **Viewer** check box, and then click **OK.**

11. In the details pane, right click the role assignment you created above, and then click **Assign Windows Users and Groups**.

12. In **Enter the object names to select (examples)**, type your user name, and then click **OK**.

#### Add the IIS worker process to the store's Readers role

By default, IIS runs in the Network Service account. You can configure an IIS worker process to run in a different account. For this sample, we will add Network Service to the Readers role.

| **Note** |
| --- |
| If you use a remote authorization store (such as Active Directory or a remote XML file based store) and run IIS in the default Network Service context, then the Active Directory account of the Web server running IIS must be added to the store's Readers role. |

#### Add the IIS worker process to the store's Readers role

1. If the Authorization Manager MMC is not already open, click **Start**, click **Run**, type **Azman.msc**, and then press ENTER.

2. In the console tree, right-click **Authorization Manager**, click **Open Authorization Store**, click **Browse**, click **C:\MyStore.xml** (which, for this example, is the store you want to work with), click **OK**, and then click **OK** again.

3. In the console tree, right-click the name of the store, and then click **Properties**.

4. Click the **Security** tab and, in the **Authorization manager user role** list, click **Reader**, and then click **Add**.

5. In the **Enter the object names to select (examples)** box, type **Network Service**, click **OK**, and then click **OK** again.

#### Configure the IIS Metabase to use URL Authorization for the Web Application

1. Click **Start**, click **Run**, type **Cmd**, and press ENTER.

2. Use the above script to create the SetUrlAuth.vbs file and save the file to \InetPub\AdminScripts directory. (By default the

InetPub directory is in the root of the drive which IIS is installed on.)

3. Type the below command and press ENTER.

4. CScript SetUrlAuth.vbs *VDirPathAuthMgrScopeNameAuthMgrStoreNameEnabledImpersonationLevel*

   For example, you might type the following:

   **Cscript SetUrlAuth.vbs W3svc\1\Root\WebApp WebApp msxml://C:\MyStore.xml true 1**

   Now, IIS URL Authorization should be configured and running for the specified Web application. Users who are in the Viewer role should be able to browse to the pages within the application.

| **Important** |
| --- |
| You must append **msxml://** if you are using an .xml file or **msldap://** if you using the Active Directory Authorization Manager store. |

# Scenario: Web Expense Application

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

In this scenario, a Web-based expense application uses the following procedure.

1. A user goes to the URL of the expense application and chooses to enter a new expense report.
2. The user enters the expense report and submits it for approval.
3. The manager receives e-mail notification that the expense report is awaiting approval.
4. The manager goes to the approval URL, which allows the manager to inspect and approve or reject the expense report.
5. If the manager approves the expense report, it is added to a verification queue, where a verifier can inspect the report, collect receipts, mark the expense report Approval-verified, and request a reimbursement.

| **Note** |
| --- |
| For a sample application using Authorization Manager, see *MicrosoftSDKPath*\Samples\Security\Authorization\AzMan in the Microsoft Platform SDK. For information about how to install the Microsoft Platform SDK, see http://www.microsoft.com/downloads/details.aspx?FamilyId=A55B6B43-E24F-4EA3-A93E-40C0EC4F68E5&displaylang=en on the Microsoft Web site. |

**Operations**

In the above expense application scenario, the operations can be defined as:

- **ExecuteExpenseControls.** Use the controls that allow user to fill out the expense form.
- **ExecuteApprovalControls.** Use the controls that allow user to approve the expense.
- **RetrieveExpenseForm.** Retrieve requested expense form from database
- **SaveExpenseForm.** Save the expense form to local machine.
- **EnqueApproval.** Put a copy of the completed form in the approval queue.
- **DequeApproval.** Remove the expense from approval queue.
- **SendRequestNotification.** Send e-mail to the manager that requests approval.
- **SendApprovalNotification.** Send e-mail notification of approval.
- **ApproveDenyExpense.** Mark the expense as approved or denied.
- **VerifyApproval.** Mark the expense as Approval-Verified.
- **SetExpenseLimit.** Set the limit amount that various roles can approve.
- **FwdRembursment.** Send a message to Accounts Payable.
- **ReadApprovals.** Read the expense forms in the approval queue.
- **ReadArchive.** Read the expense forms in approval archive.

**Tasks**
Each of the above operations may be needed in several expense application tasks that a user may need to do. Applications that use tasks should test each task as a stand-alone requirement that does not require any other task or operations permission.

In the expense application scenario, the tasks could be:

- **Create Report.** Routines to fill out an expense report. Operations (ExecuteExpenseControls, RetrieveExpenseForm, SaveExpenseForm.)
- **Submit Report.** Routines that upload the expense report to pending state, Sends e-mail to manager. (EnqueApproval, SendRequestNotification)
- **View Report** Routines that display a pending report. Operations ([Create Report] + AddApprovalQueue + SendNotification )
- **Approve Report** Routines that cause report to move from Pending state to Approved state, Sends e-mail to submitter and notifies payment system.
- **Verify Approval** Operations required to mark expense as Approval-Verified and forward reimbursement request to Accounts Payable. Operations (ExecuteApprovalControls, ApproveExpense SendApprovalNotification, FwdRembursment)
- **Audit Report** Allows reading of reports in approval queue and approval archive (ReadApprovals)
- **Config Limits** Allows the setting of approval limits for employees and managers (SetExpenseLimit)

**BizRules**
In the Web expense application, an Approve Report task would have the following BizRule attached.

```
Dim Amount
BizRuleContext.BusinessRuleResult = FALSE
Amount = BizRuleContext.GetParameter("Amount")
Limit = BizRuleContext.GetParameter(Limit)
If Amount < Limit then BizRuleContext.BusinessRuleResult = true
```

Attaching the above BizRule to the Submit Report task requires that the given expense is less than the limit amount that has been passed. The *Amount* parameter is retrieved by the application from the expense report, while the *Limit* amount is configured by the administrator of the Web expense application.

| **Important** |
|---|
| Every BizRule should immediately set the value of **BusinessRuleResult** to FALSE. If the BizRule prematurely ends due to code error or an environment problem, it is possible that random data could assign the BizRule result value to TRUE, which would grant access. |

**Roles and role definitions**
For the Web expense application, the following roles are defined during installation.

| Roles | Tasks Submit Report | View Report | Approve Report | Verify Approval | Delete Report | Config Limits |
|---|---|---|---|---|---|---|
| User | Submit reports for themselves | Those personally submitted | No | No | No | No |
| Manager | Submit for direct reports | Submitted by direct report | Submitted by direct report | No | No | No |
| Verifier | No | Yes | No | Yes | No | No |

| | | | | | | |
|---|---|---|---|---|---|---|
| Expense Admin | No | All reports | No | No | Yes | Yes |

**Role assignments**
Role assignments are configured by the Web expense application administrator after installation. An administrator might use:

- **User:** A Windows Active Directory group for all employees.
- **Manager:** An LDAP query group that checks title and number of direct reports.

  ```
  (&(Title = Manager)(numReports > 0))
  ```

  The Active Directory schema may be extended to support these properties on the user object. These properties would be maintained by a separate human resources application.
- **Verifier:** An application basic group that lists each user who verifies each expense.
- **Expense Admin:** An Active Directory group that contains the users in the expense accounting department.

After using one of the above techniques to create an Authorization Manager client context, the Web expense application can then call the **AzClientContext.GetRoles** method on the client context interface to determine which roles the user is assigned. This allows you to use the clients role membership to determine what options to display for the user.

If the user is in the Employee role, then the expense controls are displayed, which allows the user to submit an expense. If the user is in the Manager role, the displayed page has an additional area that lists the pending expense forms that have been submitted.

Closely relating role names to application logic makes modifications difficult for administrators, so a convention must be used by the application to customize presentation, based on role memberships. In this example, Employee and Manager roles are not hard-coded into the application. Instead, the role name corresponds to a similarly-named HTML file. Additionally, the **ApplicationData** property of the AzRole object may be used to provide the application with presentation information.

When the user clicks a form button to create a new expense report, the Web expense application calls **AccessCheck**, sending in the operation values that correspond to the **ExecuteExpenseControls**, **RetrieveExpenseForm** and **SaveExpenseForm** operations. By doing this, you prevent restricted users from loading these controls and degrading performance. If the user has access to the controls, the controls are displayed, an expense form is retrieved from a SQL store and loaded into the control, and the user is allowed to enter the expense data. When completed, the user clicks **Submit** and the Web expense application calls **AccessCheck**, specifying the operations that are needed to submit an expense, for example, the operation number corresponding to EnqueApproval or SendRequestNotification. Additionally, the amount of the expense is sent into **AccessCheck** as a BizRule parameter called Amount. If access is allowed, the expense report is uploaded to a SQL queue of pending requests and an e-mail is sent that notifies the users manager of the pending request.

To be able to perform these operations, permissions are needed on several back-end resources. These include the NTFS directory from which the expense form is retrieved, the SQL database for submitted expense reports, and a Web expense Microsoft Exchange account which can be used to send e-mail notifications. Instead of impersonating and performing operations as the client, the Web expense service account should be granted enough access on back-end objects to perform the operations that the application needs. Therefore, back-end DACLs only grant access to the Web expense service account. (For more information, see "Trusted subsystem application model" earlier in this paper.)

When a manager sees the pending expense as described above), **AccessCheck** verifies the ability to use the approval controls. The manager then browses the expense details and clicks **Approve** or **Deny** and **AccessCheck** is called on the operations: **DequeApproval**, **ApproveDenyExpense** and **SendRequestNotification**. Additionally, the **Amount**, **ApprovingMgr**, and **Submitter** parameters are given the values of the amount of the submitted expense, the current user attempting to approve or deny the expense, and the name of the person who submitted the expense, respectively. During **AccessCheck**, a BizRule validates the calling user as a manager of the submitting user by checking the relationship in Exchange or SQL and verifies that the amount is within a specified limit for the current approving manager.

# Scenario: Using Authorization Manager with Custom Principals

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Applications that use custom or non-Windows principals (for example, SQL, Passport, Microsoft Commerce Server, or ADAM principals) can leverage most of Authorization Manager's runtime features.

Whether in Active Directory or in XML, Authorization Manager will store SIDs as members of groups and roles, except for LDAP query groups, which don't have static members. When a SID is added to a role or an application group, it is not verified as a SID corresponding to an Active Directory or Windows NT account. You can add a completely meaningless SID as a member of a role or a group. When **IAzClientContext::AccessCheck** is called, the SIDs in the client context are compared to SIDs in Authorization Manager roles and groups and there is no validation of the SID as an Active Directory account. If a client context has a custom SID and a role has the same custom SID assigned to iteither directly or by way of an application group that has the same custom SIDthe **AccessCheck** function will check if the SID is a member of a role or group with the requested operation and grant the permissions associated with the roles to which the SID is assigned. You can use this behavior to add custom principal types to roles or application basic groups in Authorization Manager by creating a custom SID for each custom principal and assigning that SID to appropriate roles and application basic groups

Textual SIDs have the structure **S**-*Version-AuthorityID-X-Y-Z*, where *Version* is the SID version number, *AuthorityID* is the security authority ID of the issuer of the SID, and *X*, *Y* and *Z* are sub-authority values. (Currently, all SIDs are version 1.) SIDs are variable in length because then number of subauthorites in SIDs can vary. SIDs can have up to **SID_MAX_SUB_AUTHORITIES** sub-authorities. (This is defined as 15 in winnt.h.) Custom SIDs use 1 for the version number, the SECURITY_RESOURCE_MANAGER_AUTHORITY for the AuthorityID value, and custom values for sub-authorities.

When creating custom SIDs you must establish a SID design for your application. For example, you might have **S-1-9-***AppInstanceGUID-UserGuid*, where 9 is the resource manager subauthority and *AppInstanceGUID* and *UserGUID* are each broken into four sub authorities. You can also use **S-1-9-***AppInstanceGuid-UserRID*, where *UserRID* is a unique number for the user in the scope of the application instance. For sample code that generates a GUID, see the Microsoft Windows Software Development Kit (SDK).

Once you have created a custom SID, you can add the SID to Authorization Manager roles or groups using the **AddMember** method of the IAzRole or IAzApplicationGroup interfaces.

To initialize a client context with a custom SID, you can call **IAzApplication::InitializeClientContextFromSID** and specify a custom SID in string form. This creates a valid client context that is suitable for performing access validation for the custom principal. In practice, after a custom principal authenticates, an application maps the authenticated principal to a custom SID that is previously created for the principal and calls **InitializeClientContextFromSID** to attain a Authorization Manager client context. When **AccessCheck** is called on this client context, this SID will be evaluated for membership in Authorization Manager roles and groups.

#### Drawbacks to custom principals

There are some drawbacks to using custom principals.

- You cannot use LDAP query groups. Since the user is not in Active Directory, it does not make sense to have LDAP query groups.
- You cannot use Active Directory groups or local machine groups.
- You cannot use the Authorization Manager MMC user interface. Since the user interface validates user SIDs when displaying role or group memberships, administrators see textual string SIDs in the user interface and not user names. More importantly, the user interface only obtains SIDs to enter into roles and application groups by way of Object Picker, which only allows you to pick valid accounts or groups.

#### Using custom principals

The following figure illustrates the process of authorizing custom principals.

**Figure 8: Authorizing custom principals**




1. Authenticate the custom principal user. and then map the custom principal to an Authorization Manager custom SID.

2. Use **InitializeClientContextFromSID** and pass the AZ_CLIENT_CONTEXT_SKIP_GROUP flag as the *loptions* parameter to build a client context with the custom SID.

After this is done, you can use **AccessCheck** to verify that the custom principal is authorized to perform the specified operations in the specified scope.

# Using Administrative Template Files with Registry-Based Group Policy

Updated: October 11, 2004

Applies To: Windows Server 2003

Published: October 11, 2004

This white paper explains the concepts, architecture, and implementation details for registry-based Group Policy in Microsoft® Windows® operating systems. It shows how to create custom Administrative Template (.adm) files and includes a complete reference for the .adm language. In addition, it includes information about changes in .adm files for Windows XP with Service Pack 2 (SP2).

**In This White Paper**

Introduction (Using Administrative Template Files with Registry-Based Group Policy)

Overview of Registry-Based Policy and Administrative Template Files

Design Considerations for Creating Policy Settings

How to Write a Simple .Adm File for Registry-based Group Policy

Testing Administrative Template Files

Maintaining and Managing .Adm Files

What's New for Administrative Template Files in Windows XP SP2

Language Reference for Administrative Template Files

Related Links (Using Administrative Template Files with Registry-Based Group Policy)

# Introduction (Using Administrative Template Files with Registry-Based Group Policy)

Applies To: Windows Server 2003

Administrators can use Group Policy to deliver and apply one or more desired configurations or policy settings to a set of targeted users and computers within an Active Directory® directory service environment. The majority of available policy settings are provided through Administrative Template files (.adm files) and are designed to modify specific keys in the registry. This is known as registry-based policy. For many applications, the use of registry-based policy delivered by .adm files is the simplest and best way to support centralized management of policy settings.

Intended for IT administrators and developers, this document describes how to implement registry-based Group Policy by using .adm files.

For detailed instructions about enabling applications for Group Policy, see "Group Policy" in the Microsoft Platform Software Development Kit at http://go.microsoft.com/fwlink/?LinkId=26258.

# Overview of Registry-Based Policy and Administrative Template Files

Applies To: Windows Server 2003

By using registry-based policy, operating system components and applications can respond to registry key settings that administrators can manage centrally with Group Policy. These policy settings determine the behavior of the application for targeted computers or users. As long as a component or application has been policy-enabled (that is, its behavior changes based on registry values indicated in the .adm file), you can manage its features and settings through registry-based policy.

.Adm files are UNICODE text files that Group Policy uses to describe where registry-based policy settings are stored in the registry. All registry-based policy settings appear and are configured in the Group Policy Object Editor under the Administrative Templates node. .Adm files do not apply policy settings; they simply enable administrators to view the policy settings in the Group Policy Object Editor. Administrators can then create Group Policy objects (GPOs) containing the policy settings that they want to use. For example, you might have one GPO that contains various policy settings for managing the Active Desktop feature.

With the release of Microsoft® Windows XP Service Pack 2 (SP2), IT administrators have more than 1,300 Administrative Template policy settings available for their use. In addition, administrators and developers can add their own custom settings.

Registry-based Group Policy uses .adm files in the following manner:

- The Group Policy Object Editor reads the .adm files. By default, when an administrator opens a GPO, a comparison is made between the timestamps of the .adm files stored in the GPO being edited and those on the local computer. If the local .adm files have a more recent timestamp then they are uploaded to the domain controller and replicated throughout the domain.

- The Group Policy Object Editor console (gpedit.msc) displays the settings, and, depending on the .adm file, the policy settings can be displayed in a localized language.

- The Group Policy Object Editor uses .adm files to configure user interface settings such as dialog boxes, radio buttons, and drop-down lists, thereby enabling administrators to manage these settings centrally.

- The Group Policy Management Console (GPMC) uses .adm files to display policy settings when using Group Policy Results or Group Policy Modeling, also known as Resultant Set of Policy (RSoP).

| **Note** |
| --- |
| Windows XP SP2 includes modifications to the LISTBOX ADDITIVE behavior in .adm files (implemented by a new version of gptext.dll, the .dll used by the Group Policy Object Editor.) For details about these changes, see the, "What's New for .Adm Files in Windows XP SP2" section later in this document |

For more information about the architecture of Administrative Templates, see Administrative Templates Extension Technical Reference at http://go.microsoft.com/fwlink/?LinkId=35291.

**Default .adm Files**

The Group Policy Object Editor displays the policy settings within the .adm files that are included with the operating system by default. These .adm files are:

- **System.adm.** Provides policy settings to configure the operating system. System.adm is installed by default in Windows Server™ 2003, Windows XP, and Windows 2000 Server operating systems.

- **Inetres.adm.** Provides policy settings to configure Internet Explorer. Inetres.adm is installed by default in Windows Server 2003, Windows XP, and Windows 2000 Server operating systems.

- **Wuau.adm.** Provides policy settings to configure Windows Update. Wuau.adm is installed by default in Windows Server 2003, Windows XP Service Pack 1 (SP1), and Windows 2000 Server Service Pack 3 (SP3) operating systems.

- **Wmplayer.adm.** Provides policy settings to configure Windows Media Player. Wmplayer.adm is installed by default in Windows Server 2003 and Windows XP operating systems. Wmplayer.adm is not available on 64-bit versions of the Windows Server 2003 operating system and Windows XP 64-Bit Edition.

- **Conf.adm.** Provides policy settings to configure NetMeeting. Conf.adm is installed by default in Windows Server 2003, Windows XP, and Windows 2000 Server operating systems. Conf.adm is not available on 64-bit versions of the Windows Server 2003 operating system and Windows XP 64-Bit Edition.

Most Group Policy settings are contained in the System.adm file. The .adm files that ship with Windows Server 2003, Windows XP Professional, and Windows 2000 Server operating systems are located in the %windir%\inf\ folder. (for example, C:\Windows\inf). For more information about .adm file maintenance, see the "Maintaining and Managing .Adm Files" in this document.

**When to Use Registry-Based Group Policy**

In general, if a policy setting can be configured using a simple user interface, and any configuration input can be stored in the registry as plain text, you should consider using registry-based policy. Specifically, registry-based Group Policy is an appropriate solution for the following scenarios:

- **Creating available and unavailable (on/off, or yes/no) functionality.** You can use registry-based policy as if it were a switch, to turn functionality on or off. For example, you can create a policy setting to allow administrators to control whether a certain item is displayed on the desktop.

- **Defining a set of static modes.** For example, you can set the language used on a computer. You can set up a static list of the possible language selections, and when the policy setting is enabled, the administrator can select a language from the pre-created list. This action is typically shown in the user interface as a drop-down list.

- **Creating a policy setting that requires simple input that can be stored in the registry as plain text.** For example, you can create a policy setting to define the screensaver or bitmap that is displayed on the user's desktop. With this policy setting enabled, Group Policy administrators are provided with a text dialog box into which they can enter the name and path of the bitmap file to be used. This information is then stored in the registry as plain text.

**True Policies vs. Preferences**

Group Policy settings that administrators can fully manage are known as "true policies." In contrast, settings that users configure or that reflect the default state of the operating system at installation time are known as "preferences." Both true policies and preferences contain information that modifies the registry on users' computers. True policy settings take precedence over preference settings.

Registry values for true policies are stored under the approved registry keys as listed in Table 1. Users cannot change or disable these settings.

**Table 1 Approved Registry Key Locations for Group Policy Settings**

| For Computer Policy Settings: | For User Policy Settings: |
|---|---|
| HKLM\Software\Policies (The preferred location) | HKCU\Software\Policies (The preferred location) |
| HKLM\Software\Microsoft\Windows<br>\CurrentVersion\Policies | HKCU\Software\Microsoft\Windows<br>\CurrentVersion\Policies |

Preferences are set by the user or by the operating system at installation time. The registry values that store preferences are located outside the approved Group Policy keys listed in Table 1. They are located in other areas of the registry. Users can typically change their preferences at any time. For example, users can decide to change the location of their local dictionary to a different location, or set their wallpaper to a different bitmap. Most users are familiar with setting preferences that are available to them through the operating system or application user interface.

It is possible for an administrator to write an .adm file that sets registry values outside of the approved Group Policy registry trees included in Table 1. In this case, the administrator is only ensuring that a given registry key or value is set in a particular way. With this approach, the administrator configures preference settings, instead of true policy settings, and marks the registry with these settings (that is, the settings persist in the registry even if the preference setting is disabled or deleted).

If you configure preference settings by using a GPO in this manner, the GPOs that you create do not have Access Control List (ACL) restrictions. As a result, users might be able to change these values in the registry. When the GPO goes out of scope (that is, it is unlinked, disabled, or deleted), these values are not removed from the registry. In contrast to this, true registry policy settings do have ACL restrictions to prevent users from changing them, and the policy values are removed when the GPO that set them goes out of scope. For this reason, true policies are considered to be policy settings that can be fully managed. By default, the Group Policy Object Editor only shows policy settings that can be fully managed. To view preferences in the Group Policy Object Editor, you need to click the **Administrative Templates** node, click **View**, then click **Filtering**, and then clear **Only show policy settings that can be fully managed**.

Although Group Policy settings take priority over preferences, they do not overwrite or modify the registry keys used by the preferences. If a policy setting is deployed that conflicts with a preference, the policy setting takes precedence over the preference setting. If a conflicting policy setting is removed, the original user preference setting is restored.

**How to Use Policy Settings and Preferences**

Applications commonly include a user preference and a policy setting that perform similar or related functions. For example, you might want to offer users the ability to configure part of a component through a user preference setting, and also centrally control this setting by using a registry-based policy setting.

An example of where both a policy and preference can co-exist is the configuration of the wallpaper on a Windows desktop. Users can set their desktop wallpaper to be displayed (or not displayed) by using the **Display** icon in Control Panel. You can also use a policy setting to configure desktop wallpaper. To specify the desktop wallpaper that displays on users' desktops, administrators can use the **Active Desktop Wallpaper** policy setting (found in the Group Policy Object Editor, under the **User Configuration\Administrative Templates\Desktop\Active Desktop**. As a result, the user can choose to display or not display the wallpaper, but the administrator can choose which wallpaper is displayed when the display setting is ON.

Table 2 lists the resultant behavior for Group Policy settings and preferences.

**Table 2 Results of Group Policy Settings and Preferences**

| Scenario | Policy Present | Preference Present | Resultant Behavior |
|---|---|---|---|
| No policy or preference | No | No | Default behavior. |
| Preference Only | No | Yes | Preference configures behavior. |
| Policy only | Yes | No | Policy configures behavior. |
| Both policy and preference | Yes | Yes | Policy configures behavior. Preference is ignored. In all cases, policy overrides preference. |

# Design Considerations for Creating Policy Settings

Applies To: Windows Server 2003

This section addresses essential issues for creating and configuring custom policy settings in .adm files.

Use the following questions as a guide to help you design Group Policy settings.

- What is the default behavior (that is, when the policy is set to **Not Configured**)?
- What is the behavior when the policy is **Enabled**, **Disabled**, or **Not Configured**? The **Enabled** behavior should always be the opposite of the default behavior (that is, **Not Configured**).
- Do administrators need to explicitly disable a feature?
- Do the proposed policy settings affect users or computers or both?
- What are some potential future ramifications of the new policy settings? When new products are released, you must continue to maintain the previous .adm settings to manage computers running legacy software. New products and settings must be able to co-exist with earlier versions.

**When to Create Policy Settings**

An administrator should consider creating a policy setting for the following purposes:

- To help administrators manage and increase security of their desktop computers.
- To hide or disable a user interface that can lead users into a situation in which they must call the helpdesk for support.
- To hide or disable new behavior that might confuse users. A policy setting created for this purpose allows administrators to manage the introduction of new features until after user training has taken place.
- To hide settings and options that might take up too much of users' time.

**Controlling Feature Releases to Users**

An administrator can use .adm files to provide policy settings to manage new features of a new or updated application. By creating a single GPO for all new features, or by creating a GPO for each logical grouping of new features, an administrator can reduce the need for support and potential user frustration. A GPO should also be considered for specific features that administrators need to control after the new features have been enabled.

By enabling policy settings in this area, the administrator can control how and when users get new product features. By grouping related features, the administrator can prevent users from using a new feature set until they have been trained.

**Create Policy Settings to Reduce Need for Support**

To reduce the need for support, administrators can start by determining the top issues that users have and considering ways in which they can use policy settings to prevent support calls. For example, you could use policy settings to control software settings in the following scenarios:

- When proper configuration settings require advanced knowledge of the application.
- When there are complicated or advanced configuration settings that the typical user does not need to use.

In these scenarios, it would be appropriate to use Group Policy to give an administrator the ability to control access to the configuration settings.

**Control Data**

You can create policy settings to populate data for your application. Such data usually exists in small sets in the form of numbers, text strings, and so on. For example, a phone dialer could use policy settings to enable administrators to mandate that certain items exist in the phone directory.

**When Not to Create a Policy**
Although registry-based Group Policy provides an effective way of managing components and applications, there are some circumstances where its use is not recommended. For example:

- Do not create a policy for *all* of your application settings because large applications typically contain hundreds or even thousands of settings, and only a subset of these needs to be managed through Group Policy. Be selective about the features you want to enable or disable. Because Group Policy provides centralized management of the setting, you should evaluate whether administrators would want this kind of management before adding the policy setting.

- Do not create a policy if you do not intend to provide support for the policy setting. Treat each policy as a feature that needs to be tested, validated, and supported.

**User Interface Design**
Effective policy settings should be clearly written and displayed. You must also ensure that the user interface is clear and easy to understand. For example, review these user interface design options for disabling **My Network Places**:

- **Create an error message.** For example, the user clicks My Network Places and the following error message appears: "This option has been disabled by your administrator." In response, the user calls the administrator or support desk to ask why this feature has been disabled.

- **Disable the user interface.** For example, a user interface feature in My Network Places is disabled (grayed-out). This implies to the user that there is a way to enable the user interface. In response, the user might spend a lot of time trying to get this feature to work. In the end, the user might either give up in frustration or call the support desk.

- **Hide the user interface feature.** For example, a user interface feature in My Network Places is hidden. In response, the user does not recognize that anything is missing or unavailable. This is the preferred choice in this scenario.

- **Do nothing.** For example, when a user clicks My Network Places, and the screen does not change (that is, nothing happens). In response, the user assumes that something is wrong and calls the support desk.

**Policy Names**
You set the name of the policy setting at the same time that you create it. The name of the policy setting is displayed in the Group Policy Object Editor. Use the following guidelines for creating policy names:

- Use a verb that reflects the effect of the policy setting. Examples of verbs commonly used in policy setting names include: allow, permit, turn on, prohibit, hide, and prevent.

- Do not use the terms "enabled" or "disabled" in your policy setting names. Instead, consider using the terms "turn on" and "turn off," or "allow" and "prevent."

- Avoid overly technical jargon that might not be understood by administrators who are not experts in a particular component. Include technical details of the policy setting in the Explain text.

- Use short, descriptive names that accurately reflect the function of the policy setting, for example, Turn off Internet File Association service.

| **Note** |
| --- |
| Policy names are limited to 256 characters. However, depending on the font used on the user's workstation, a smaller number of characters are displayed, and all others truncated. On most systems, you can assume that you have the ability to display policy names up to 65 characters. (Using a Resolution of 1024 / 768 with small fonts installed).<br><br>When Group Policy names are translated into additional languages, they typically require additional characters. Therefore, if you are using English to name a Group Policy, limit the name to 49 characters. This limitation allows your title to grow by 33 percent during translation, without causing any truncation issues. |

**Explain Text**
Explain text provides information about the behavior of the policy setting, its interaction with other policy settings, and can

include any other information that you would like administrators to be aware of. When an administrator selects a policy setting, and then clicks the **Explain** tab, the Explain text appears in the policy setting's **Properties** dialog box. Explain text is limited to a maximum of 4096 characters, and Category Explain text is limited to a maximum of 256 characters.

| **Note** |
| --- |
| Do not try to supplement a Group Policy setting title in the user interface with tip text that can be displayed in the user interface. Include any tips for using the policy setting in the Explain text. |

Use the following guidelines when you create the **Explain** text:

- Draft the **Explain** text soon after creating the specification document for the policy setting. This serves as a high-level roadmap for developers, and assists testers in creating a test plan for the policy.
- Make sure your documentation team is available to help create the Explain text.
- Target the Explain text to the Group Policy administrator, rather than the end user.

Use the following template for the Explain text, and make sure that you include the following items:

- Write a one- or two-line description of the policy setting.
- Write a one- or two-line description of the feature that the policy setting affects.
- Use separate paragraphs within your Explain text for each policy setting state, as follows:
- If you enable this policy setting...<describe enabled behavior>.
- If you disable this policy setting...<describe disabled behavior>.
- If you do not configure this policy setting...<describe default behavior>.
- Include tips for using the policy setting.
- Include notes or interactions that this policy has with other settings.

As a best practice, you should also provide information about the following:

- Items that are not covered by this policy setting.
- Any other policy settings that are required for your policy setting to function.
- Any other policy settings that are related to the same component that your policy setting affects and which may have a higher or lower priority. For example, if you have a policy setting to restrict access to the LAN settings for a computer. This policy setting takes priority over more specific policy settings regarding the actual items you can configure within a LAN connection.
- Any related policies that the administrator needs to be aware of.

The following Explain text is from the **Active Desktop Wallpaper** policy setting in the **User Configuration\Administrative Templates\Desktop\Active Desktop**. This policy setting controls the desktop background (wallpaper) that is displayed on users' desktops.

Specifies the desktop background ("wallpaper") displayed on all users' desktops.

This setting lets you specify the wallpaper on users' desktops and prevents users from changing the image or its presentation. The wallpaper you specify can be stored in a bitmap (*.bmp), JPEG (*.jpg), or HTML (*.htm, *.html) file.

To use this setting, type the fully qualified path and name of the file that stores the wallpaper image. You can type a local path, such as C:\Windows\web\wallpaper\home.jpg or a UNC path, such as \\Server\Share\Corp.jpg. If the specified file is not available when the user logs on, no wallpaper is displayed. Users cannot specify alternative wallpaper. You can also use this setting to specify that the wallpaper image be centered, tiled, or stretched. Users cannot change this specification.

If you disable this setting or do not configure it, no wallpaper is displayed. However, users can select the wallpaper of their choice.

Also, see the "Allow only bitmapped wallpaper" in the same location, and the "Prevent changing wallpaper" setting in User Configuration\Administrative Templates\Control Panel.

| **Note** |
| --- |
| You need to enable the Active Desktop to use this setting. |

| **Note** |
| --- |
| This setting does not apply to Terminal Server sessions. |

**Best Practices for Developing Registry-Based Policy Settings**

It is strongly recommended that you do not try to create a single policy setting to control all aspects of your component. Group Policy is easier to implement and use if you have several smaller policy settings. This approach gives you more flexibility in designing and modifying your policy settings.

All policy settings should have an associated behavior for each of the three possible states shown in Table 3:

**Table 3 Policy States and Associated Behaviors**

| Enabled | Disabled | Not Configured |
| --- | --- | --- |
| Turns on the behavior indicated by the policy name | Turns off the behavior indicated by the policy name | Has no effect – default behavior |

Consider the following guidelines when you design your policy settings:

- Policy settings are never removed from the .adm files supplied by Windows operating systems. Even if subsequent versions of Windows no longer support the policy setting, Microsoft will continue to include that policy setting in .adm files for all future Windows operating systems that support Group Policy.

- Computer policy settings should override user policy settings.

- The **Enabled** behavior should be the *opposite* of the default behavior. For example, if a feature is on by default, the policy setting should be named using something like "Turn off <feature>", for example, **Turn off reminder balloons**. By default, reminder balloons are displayed when the Offline Files feature is enabled; they are used to notify users when they have lost the connection to a networked file and are working on a local copy of the file. If **Turn off reminder balloons** is set to **Enabled**, the system hides the reminder balloons, and prevents users from displaying them.

- Consider whether administrators need to explicitly disable a feature. You must understand the differences between the **Not Configured** state (which implies that the administrator does not care) and the **Disabled** state (which means that the administrator cares and wants to implement a very specific behavior).

- Make sure that your documentation team is involved with creating the Explain text. Well-written Explain text can help reduce support calls.

- Each component should expose a user interface that always reflects the policy setting that is applied. For example, if a Group Policy setting removes the ability for the user to set a preference for a component, the user interface should clearly indicate this and access to that particular item in the component should be removed (for example, the item is disabled and grayed-out, or it is not visible to users).

**Creating Custom .Adm Files**

You can create custom .adm files to extend the use of registry-based policy settings to new applications and components. Creating and implementing custom .adm files involves the following tasks:

- The application or component must be enabled to use Group Policy-enabled. Many off-the-shelf applications for Windows are already Group Policy-enabled. For example, Microsoft Office is developed to work with Group Policy settings, and applicable .adm files are included in the Office Resource Kit. If a developer's application is not already Group Policy-enabled, the developer must write code that changes the application so that an administrator can apply the application-specific Group Policy settings. For more information, see Microsoft Platform Software Development Kit at http://go.microsoft.com/fwlink/?LinkId=20540.

- After the application or component is Group Policy-enabled, a developer or administrator must write an .adm file that includes the appropriate settings for that application. Custom .adm files are created as text files that describe policy settings. A framework language is provided for .adm files, as described in "Language Reference for Administrative Template Files" later in this document. When you create a new policy setting in a custom .adm file, it is more efficient if you copy and modify existing policy settings that are similar to the ones that you want to create, rather than write new policy settings from scratch. To use an existing .adm file as a template for a custom .adm file, first copy the original .adm file into a new file with a different file name, and then edit the new file. Similar policy settings can provide:

    - A model to use as a basis for your Group Policy settings.

    - Consistency with available policy settings.

    - Information about possible interactions between policies.

- Administrators use the Group Policy Management Console to view and manage GPOs and use the Group Policy Object Editor to view the **Administrative Templates** node under **Computer Configuration** or **User Configuration**.

| ⚠**Caution** |
| --- |
| Treat the .adm files that ship in the operating system as read-only files. These files are often updated when you install service packs or future releases of the product. Policy settings should never be removed from .adm files that are included in the operating system by default. |

Keep in mind that by itself, the .adm file does not actually apply Group Policy to the client computer. You must have a corresponding component or application that responds to the registry key that is affected by the policy setting.

The following code example illustrates a simple .adm file.

```
CLASS USER

CATEGORY !!DesktopLockDown

    POLICY !!DisableTaskMgr
       EXPLAIN !!DisableTaskMgr_Explain
       VALUENAME "DisableTaskMgr"
       VALUEON NUMERIC 1
       VALUEOFF NUMERIC 0
       KEYNAME "Software\Policies\System"
    END POLICY
END CATEGORY
[strings]
DisableTaskMgr="Disable Task Manager"
DisableTaskMgr_Explain="Prevents users from starting Task Manager"
DesktopLockDown="Desktop Settings"
```

This sample code exists in the System.adm file. It allows you to configure a policy called **Disable Task Manager**, which appears in the Group Policy Object Editor namespace in **User Configuration\Administrative Templates\Desktop Settings**.

This policy setting defines the following behavior:

- If an administrator enables this policy setting, it creates a registry key called **DisableTaskMgr** and set its value to 1. The **VALUEON** tag implements this behavior. In this case, users cannot start Task Manager.

- If an administrator disables this policy setting, it creates a registry key called **DisableTaskMgr** and set its value to 0. The **VALUEOFF** tag implements this behavior. In this case, users can start Task Manager.

- In both cases, the **DisableTaskMgr** registry key is created below **HKCU\Software\Policies\System** in the registry. Note that the key is created under **CLASS USER** and not under **CLASS MACHINE** because this is a user policy setting.

- If an administrator sets this policy setting to **Not Configured**, it deletes the registry key called **DisableTaskMgr.** In this case, users can start **Task Manager**.

Table 4 lists where registry-based policy settings are stored in any of the four approved registry locations for Group Policy.

**Table 4 Approved Registry Key Locations for Group Policy Settings**

| **Computer Policy Settings** | **User Policy Setting:** |
|---|---|
| HKLM\Software\Policies (The preferred location) | HKCU\Software\Policies (The preferred location) |
| HKLM\Software\Microsoft\Windows<br><br>\CurrentVersion\Policies | HKCU\Software\Microsoft\Windows<br><br>\CurrentVersion\Policies |

The registry keys are created when a GPO containing an enabled or disabled registry-based policy setting is applied. If the GPO is later removed (for example, unlinked from an OU in which a user resides), the registry keys associated with it are also removed. Security prohibits a standard user from changing the registry keys. A local administrator can overwrite these registry keys, and thus change or disable the behavior of the policy setting. For this reason, as a Group Policy administrator, you might want to enable the **Registry policy processing** policy setting (located in **Computer Configuration\Administrative Templates\System\Group Policy**) and select the following option: **Process even if the Group Policy objects have not changed**.

Selecting this option updates and reapplies the policy settings even if you -- the Group Policy administrator -- have not changed the policy settings. This provides an additional safeguard in the event that local administrators try to change policy settings through the registry. Policy settings are reapplied during the regular background refresh of Group Policy, which occurs by default every 90 minutes with a randomized delay of up to 30 minutes.

# How to Write a Simple .Adm File for Registry-based Group Policy

Applies To: Windows Server 2003

This section shows you how to create a simple .adm file for delivering registry-based Group Policy settings. The code samples provided are parts of a single .adm file. Please refer to "Language Reference for Administrative Template Files," later in this document for full details about the .adm language.

The sections of the sample .adm file illustrate how to set a registry value to:

- Turn a feature on or off.
- Allow the selection of one or more values from a list.
- Display a list with add and remove buttons.
- Enter EDITTEXT and static text.
- Display a numeric list to the administrator.
- Display a numeric list to the administrator, using a spin control.
- Display an actionable list to an administrator.

**To Set a Registry Value to Turn a Feature ON or OFF**

This sample code displays a check box. The value is set in the registry with the REG_DWORD type. By default, the check box is clear. If the check box is selected, it writes the value 1 to the registry. If the check box is clear, it writes the value 0 to the registry.

```
CLASS USER

CATEGORY!!SampleCategory
    KEYNAME "SOFTWARE\Policies\Microsoft\ADM_Samples"

    POLICY!!Sample_ADM_FeatureOnOff
        #if version >= 4
            SUPPORTED!!SUPPORTED_WindowsXPSP1
        #endif
        EXPLAIN!!Sample_ADM_FeatureOnOff_Help
        VALUENAME "ADM_Sample_FeatureOnOff"
            VALUEON 1
            VALUEOFF 0
    END POLICY
END CATEGORY
```

**To Set a Registry Value to Allow the Selection of One or More Values from a List**

You can use **DROPDOWNLIST** to display a combo box with a drop-down style list. You can pre-populate the list of items that are displayed in the list and the corresponding registry value to be written for each item in the list.

```
POLICY!!Sample_ADM_DropDownList
    #if version >= 4
        SUPPORTED!!SUPPORTED_WindowsXPSP1
    #endif
    EXPLAIN!!Sample_ADM_DropDownList_Help
    PART!!Sample_ADM_DropDownList DROPDOWNLIST REQUIRED
        VALUENAME "Sample_ADM_DropDownList"
            ITEMLIST
             NAME !!Sample_ADM_DropDownList_Always    VALUE NUMERIC 1 DEFAULT
             NAME !!Sample_ADM_DropDownList_WorkStationOnly    VALUE NUMERIC 2
             NAME !!Sample_ADM_DropDownList_ServerOnly    VALUE NUMERIC 3
         END ITEMLIST
    END PART
```

```
END POLICY
```

**To Set a Registry Value to Display a List with Add and Remove Buttons**
You can use registry-based Group Policy to display a list box that contains Add and Remove buttons.

By default, only one column appears in the list box, and for each entry, a value is created where the name and value are the same. For example, a name entry in the list box creates a value called *name* that contains data labeled *name*.

```
POLICY!!Sample_ADM_ListBox
    #if version >= 4
        SUPPORTED !!SUPPORTED_WindowsXPSP1
    #endif
    EXPLAIN!!Sample_ADM_ListBox_Help
    PART!!Sample_ADM_DropDownList LISTBOX
        KEYNAME "Sample_ADM_ListBox"
    END PART
END POLICY
```

**To Set a Registry Value for EDITTEXT and Static Text**
This setting allows the user to type alphanumeric text in an edit field. The text is set in the registry with the REG_SZ type.

The following code is an example of how you can use the **EDITTEXT PART** type.

```
POLICY!!Sample_ADM_EditText
    #if version >= 4
        SUPPORTED !!SUPPORTED_WindowsXPSP1
    #endif
    EXPLAIN!!Sample_ADM_EditText_Help
    PART!!Sample_ADM_EditText EDITTEXT
        VALUENAME "ADM_Sample_EditText"
    END PART
END POLICY
```

**To Set a Registry Value for Displaying a Numeric List to the Administrator**
To display a list of numbers from which administrators can select one of the predefined numeric values, you can use a NUMERIC PART type.

```
POLICY!!Sample_ADM_Numeric

    #if version >= 4

        SUPPORTED !!SUPPORTED_WindowsXPSP1

    #endif

    EXPLAIN!!Sample_ADM_Numeric_Help

    PART!!Sample_ADM_Numeric NUMERIC

        VALUENAME "ADM_Sample_Numeric"

    END PART

END POLICY
```

**To Set a Registry Value for Displaying a Numeric List to the Administrator, Using SPIN Control**
To display a list of numbers from which administrators can select one of the predefined numeric values, you can use a SPIN control.

```
POLICY!!Sample_ADM_Spinner
    #if version >= 4
        SUPPORTED !!SUPPORTED_WindowsXPSP1
    #endif
    EXPLAIN!!Sample_ADM_Spinner_Help
    PART!!Sample_ADM_Spinner NUMERIC
        VALUENAME "ADM_Sample_Spinner"
        MIN 5 MAX 23 DEFAULT 14 SPIN 3
    END PART
END POLICY
```

**To Set a Registry Value to Display an ActionList to an Administrator**

To specify a set of arbitrary registry changes to make in response to a control being set to a particular state, you can use the **ACTIONLIST** option.

```
POLICY!!Sample_ADM_ActionList
    #if version >= 4
        SUPPORTED !!SUPPORTED_WindowsXPSP1
    #endif
    EXPLAIN!!Sample_ADM_ActionList_Help
    ACTIONLISTON
        KEYNAME "SOFTWARE\Policies\Microsoft\ADM_Sample\ActionOnList"
            VALUENAME "Action1"
            VALUE NUMERIC 1
        KEYNAME "SOFTWARE\Policies\Microsoft\ADM_Sample\ActionOnList"
            VALUENAME "Action2"
            VALUE NUMERIC 7
        KEYNAME "SOFTWARE\Policies\Microsoft\ADM_Sample\ActionOnList"
            VALUENAME "Action1"
            VALUE NUMERIC 100
    END ACTIONLISTON
    ACTIONLISTOFF
        KEYNAME "SOFTWARE\Policies\Microsoft\ADM_Sample\ActionOnList"
            VALUENAME "Action1"
            VALUE NUMERIC 0
        KEYNAME "SOFTWARE\Policies\Microsoft\ADM_Sample\ActionOnList"
            VALUENAME "Action2"
            VALUE NUMERIC 0
        KEYNAME "SOFTWARE\Policies\Microsoft\ADM_Sample\ActionOnList"
            VALUENAME "Action1"
            VALUE NUMERIC 0
    END ACTIONLISTOFF
END POLICY
```

# Testing Administrative Template Files

Applies To: Windows Server 2003

This section provides recommended procedures for testing .adm files. You should always test policy settings before deploying them in GPOs. Your test plan should examine and document the following for each policy setting:

- The default behavior
- The behavior when the policy setting is enabled, disabled, or not configured
- Changes in the user interface when a policy setting is enabled, disabled, or not configured
- The possible settings that the policy setting can be configured as
- The associated behavior. For example, does configuring the policy setting affect other policy settings, or is one policy setting dependent on another?
- Any associated preferences
- The behavior of the associated preferences
- The behavior in the event of invalid input by an administrator

First, test your new policy settings individually. Next, test how each policy setting interacts with other policy settings that are similar, or policy settings that are also designed to manage the affected component.

Suppose for example that you create a policy setting to configure the wallpaper that is displayed on your clients' desktops. The list of policy settings that ship with Windows Server 2003 includes other wallpaper policy settings. In this scenario, you should test how these policy settings interact. Any issues that arise should be addressed or documented in the Explain text.

Testing should ensure that the user interface is policy-aware. If the user interface is not aware of the policy setting, the user experience might be confusing. For example, if your policy setting restricts access to a certain item in a component, then no access to this component and its configuration should be available in the user interface. Some ways to achieve this are:

- Removing the item completely visually
- Gray-out the item and disable it

**Creating an .Adm Test File**

To simplify testing, create a sample .adm file for your policy setting. By doing this step, you can isolate testing of the new policy settings until they are ready to be merged into a larger .adm file, if appropriate.

Your test plan should include the following validations:

- Your policy settings display and can be configured in Group Policy Object Editor, and can be set to all pertinent combinations of values.
- The settings are configured properly on the client when configured in a GPO, and the component responds appropriately to the policy setting.
- You can generate settings and Resultant Set of Policy (RSoP) reports for your policy settings by using Group Policy Management Console.
- The Explain text for the policy setting is accurate and thoroughly explains how your policy setting works. It should describe the behavior for all states of the setting: **Enabled**, **Disabled**, and **Not Configured**.

**Loading an .Adm File into the Group Policy Snap-in**

After you create an .adm file, you can load it into the Administrative Templates section of Group Policy Object Editor by performing the following procedure.

**To load your .adm file into Group Policy Object Editor**

1. Open Group Policy Object Editor.

2. Under either **Computer Configuration** or **User Configuration**, right-click **Administrative Templates**, and then click **Add/Remove Templates**.

3. In the **Add/Remove Templates** dialog box, click **Add**.

4. Navigate to the folder containing the .adm file that you would like to add. Select the file, and then click **Open**.

5. Do one of the following:

   - If your .adm file was successfully loaded, in the **Add/Remove Templates** dialog box, click **Close**. Your policy template has been added successfully.

   - If your .adm file was not successfully loaded, a dialog box is displayed, showing the error and line number of the error. Make a note of the errors that were found, and click **OK**. Although your .adm file was not successfully loaded, it still appears in the list of .adm files loaded. Select your .adm file, click **Remove,** and then click **Close**. Edit the .adm file and correct any problems.

# Maintaining and Managing .Adm Files

Applies To: Windows Server 2003

This section provides guidance to ensure that your .adm files are properly configured for your administrative workstations.

Each domain-based GPO maintains a single folder in the Sysvol folder of each domain controller. This folder, known as the Group Policy template, contains all of the .adm files that were used in the Group Policy Object Editor when the GPOs were last created or edited.

Each Windows Server operating system includes a standard set of .adm files that are loaded by default into the Group Policy Object Editor. By default, the following .adm files are pre-loaded with a Windows 2000 or later operating system:

- System.adm
- Conf.adm
- Inetres.adm

Windows Server 2003 and Windows XP include the following additional .adm files:

- Wmplayer.adm
- Wuau.adm

**Importance of Timestamps for .Adm Files**

Each administrative workstation that is used to run the Group Policy Object Editor stores its .adm files in the %windir%\Inf folder. When GPOs are created and first edited, the .adm files from this folder are copied to the Sysvol and replicated to all the domain controllers. This includes the standard .adm files and any custom .adm files that are created by the administrator.

| **Note** |
| --- |
| If you create a GPO and do not edit it, you will create a Group Policy template without any associated .adm files. |

By default, when you edit GPOs, the Group Policy Object Editor compares the timestamps of the .adm files in the workstation's %windir%\Inf folder (local .adm files) with those that are stored in the Sysvol (these are the .adm files used by the GPO that is being edited). If the local .adm files are newer, the Group Policy Object Editor copies these files to the Sysvol, overwriting any existing files of the same name. This comparison occurs whenever the Administrative Templates node (under Computer or User Configuration) is selected in the Group Policy Object Editor, regardless of whether the administrator actually edits the GPO.

| **Caution** |
| --- |
| Because of the importance of timestamps on .adm file management, it is recommended that you do not edit system-supplied .adm files. If a new policy setting is required, it is highly recommended that you create a custom .adm file. This prevents the override of system-supplied .adm files whenever service packs are released. |

**Duplicate CATEGORY Sections**

If the .adm file that you add includes a duplicate **CATEGORY** to one that is used in an existing .adm file, the policy settings are merged.

An error can occur when the existing and the added .adm file contain the same **CATEGORY**, and both of them have a default **KEYNAME** specified (regardless of whether it is the same name). If these conditions are met, the following error message appears.

Key name specified more than once. Likely causes are: 1) the KEYNAME tag is defined multiple times in this category, 2) this KEYNAME is already defined in another category with the same name, 3) this KEYNAME is already defined in another category with the same name in a different .adm file.

**.Adm Files Used by the Group Policy Object Editor**

By default, the first time the Group Policy Object Editor is started for a specified GPO, it copies the System.adm file from the current computer's %windir%\inf\ directory to the GPO.

Subsequently, only those .adm files specified in the list are displayed. Each time you open Group Policy Object Editor, the Group Policy Object Editor also checks the listed .adm files and copies any newer versions from the local computer's %windir%\inf\ directory to the GPO.

#### .Adm Files Used by GPMC

By default, GPMC always uses local .adm files, regardless of their time stamp, and it never copies .adm files to Sysvol. If an .adm file is not found locally, GPMC looks for the .adm file in the Sysvol. In GPMC, you can specify an alternative location for .adm files. If an alternative location is specified, this alternative location takes precedence.

#### Controlling .Adm File Replication

Because .adm files are stored in the Group Policy template by default, they increase the Sysvol folder size. The File Replication Service (FRS) replicates all of the .adm files for GPOs throughout the domain. If you edit GPOs frequently, you might experience a significant amount of replication traffic. You can use a combination of the **Turn off automatic updates of adm files** and **Always use local adm files for Group Policy Object Editor** policy settings to reduce the size of the Sysvol folder and policy-related replication traffic.

#### Turn Off Automatic Updates of .Adm Files

The **Turn off automatic updates of adm files** policy setting is available under **User Configuration\Administrative Templates\System\Group Policy** in Windows Server 2003, Windows XP, and Windows 2000 operating systems. When this policy setting is enabled, the updating of an existing GPO remains local to the administrative workstation, and .adm files are not sent from the local computer to the Sysvol.

| **Note** |
| --- |
| In Windows 2000 operating systems, when you edit a GPO for the first time the local .adm files are uploaded to the Sysvol, without regard to how this policy setting is set. If this policy setting is enabled in Windows XP, .adm files are not uploaded when a GPO is edited for the first time. The first time that the GPO is edited might or might not be when the GPO is created. |

#### Always Use Local .Adm Files for Group Policy Object Editor

The **Always use local ADM files for Group Policy Object Editor** policy setting is available under **Computer Configuration\Administrative Templates\System\Group Policy** in Windows Server 2003 and Windows XP Professional. When a GPO is created, this policy setting has no immediate effect, and the .adm files on the local computer are still uploaded to the Sysvol. However, when you edit an existing GPO, any .adm files that are stored in the Sysvol are ignored, and Group Policy Object Editor uses the .adm files from the local computer only. If a policy setting has been set in the GPO, but the corresponding .adm file that describes the policy setting is not available on the local computer, Group Policy Object Editor does not display that policy setting.

| **Note** |
| --- |
| If this policy setting is enabled, the Turn off automatic updates of adm files policy setting is also implied. |

### To clear the Sysvol folder of .adm files

1. Enable the **Turn off automatic update of adm files** policy setting for all Group Policy administrators who will be editing GPOs and verify that this policy has been applied.

2. Copy any custom .adm templates to the %windir%\Inf folder.

3. Edit existing GPOs, and right-click **Administrative Templates**, and then click **Add/Remove Template** to remove all .adm files from the Sysvol.

4. Enable the **Always use local adm files for Group Policy Object Editor** policy setting for administrative workstations.

#### Maintaining .Adm files on Administrative Workstations

When you use the **Always use local adm files for Group Policy Object Editor** policy setting, make sure that each workstation has the latest version of the default and custom .adm files. If all .adm files are not available locally, some policy settings that are contained in a GPO will not be visible to the administrator. Avoid this by implementing a standard operating system and service pack version for all administrators. If you cannot use a standard operating system and service pack, implement a process to distribute the latest .adm files to all administrative workstations.

**Note**

Because the workstation adm files are stored in the %windir%\Inf folder, any process that is used to distribute these files must run in the context of an account that has administrative credentials on the workstation.

**Multilanguage Administration Issues**

In some environments, policy settings must be presented to the user interface in several languages. Because the GPT folder can store only one set of .adm files, you cannot use the GPT folder to store .adm files for multiple languages.

For Windows 2000 operating systems, the use of local .adm files for the Group Policy Object Editor is not supported. If you are using Windows 2000, use the **Turn off automatic updates of adm files** policy setting. Because this policy setting has no effect on the creation of new GPOs, the local .adm files will be uploaded to the GPT folder in Windows 2000. Creating a GPO in Windows 2000 effectively defines the language of the GPO. If the **Turn off automatic updates of adm files** policy setting is in effect on all computers running Windows 2000, the language of the adm files in the GPT folder will be defined by the language of the computer that is used to create the GPO.

If you are using Windows 2000 workstations, use the **Turn off automatic updates of adm files** policy setting for administrators, and consider the adm files in the GPT folder to be the effective language for all Windows 2000 workstations.

For administrators who are using Windows Server 2003 and Windows XP operating systems, the **Always use local adm files for Group Policy Object Editor** policy setting can be used. For example, this policy setting allows a French administrator to view policy settings by using the French .adm files that are installed locally, regardless of the adm file that is stored in the GPT folder. Note that when you use this policy setting, it is implied that the **Turn off automatic updates of adm files** policy setting is also enabled, to avoid unnecessary updates of the adm files to the GPT folder.

Consider standardizing on Windows Server 2003 for administrative workstations in a multi-language administrative environment. Configure both the **Always use local adm files for Group Policy Object Editor** and **Turn off automatic updates of adm files** policy settings.

**Operating System and Service Pack Release Issues**

Each operating system or service pack release includes a superset of the .adm files provided by earlier releases, including policy settings that are specific to operating systems that are different to those of the new release. For example, the .adm files that are provided with Windows Server 2003 include all policy settings for all operating systems, including those that are only relevant to Windows 2000 or Windows XP Professional. This means that only viewing a GPO from a computer with the new release of an operating system or service pack effectively upgrades the .adm files. Because later releases are typically a superset of previous .adm files this will not typically create problems, assuming that the .adm files that are being used have not been edited.

In some situations, an operating system or service pack release includes a subset of the .adm files that were provided with earlier releases. This has the potential to present an earlier subset of the .adm files, resulting in policy settings no longer being visible to administrators when they use Group Policy Object Editor. However, the policy settings will remain active in the GPO--only the visibility of the policy settings in Group Policy Object Editor is affected. Any active (either **Enabled** or **Disabled**) policy settings are not visible in Group Policy Object Editor, but remain active. Because the settings are not visible, it is not possible for the administrator to view or edit these policy settings. To work around this issue, administrators must become familiar with the .adm files that are included with each operating system or service pack release before using the Group Policy Object Editor on that operating system, keeping in mind that the act of viewing a GPO is enough to update the .adm files in the GPT, when the timestamp comparison determines an update is appropriate.

To plan for this in your environment, it is recommended that you either:

- Define a standard operating system/service pack from which all viewing and editing of GPOs occurs, making sure that the .adm files being used include the policy settings for all platforms.

- Use the **Turn off automatic updates of adm files** policy setting for all Group Policy administrators to make sure that .adm files are not overwritten in the Sysvol by any Group Policy Object Editor session, and make sure that you are using the latest .adm files that are available from Microsoft.

**Note**

The Always use local adm files for Group Policy Object Editor policy is typically used with the Turn off automatic updates of adm files policy setting, (that is, when it is supported by the operating system from which you are running Group Policy Object Editor).

**Keyboard Shortcuts for Administrative Templates Node**
You can use the following keyboard shortcuts to navigate the Administrative Templates namespace:

- SHIFT+asterisk (*) (on the keypad only) automatically expands the current node and all of its child nodes.
- The plus sign (+) on the keypad expands one level.
- Minus (-) on the keypad collapses one level.
- Double-click a policy in the results pane to bring up the **Properties** page.

When the **Properties** page appears, move it out from in front of the Group Policy Object Editor window. Then, click back one of the **Policies** in the results pane. You can now use the cursor keys to navigate up and down the list. Notice that the information in the **Properties** page changes. This method also works for the text on the **Explain** tab. You can also use the Tab key to move back and forth between the tree pane and the results pane, while leaving the **Properties** page open.

# What's New for Administrative Template Files in Windows XP SP2

Applies To: Windows Server 2003

Windows XP SP2 includes modifications to tools used to administer Group Policy. This includes changes to the **LISTBOX ADDITIVE** behavior, which is implemented through an update to gptext.dll. This section includes details of these changes.

For information about managing new features released in Windows XP SP2, see Managing Windows XP Service Pack 2 Features Using Group Policy at http://go.microsoft.com/fwlink/?LinkId=31974.

**Changes to LISTBOX ADDITIVE**

Policy settings defined by the **LISTBOX** keyword allow you to manage multiple values under one registry key. When **LISTBOX** is used, the specified settings are written as a **REG_SZ** registry value type, which allows for multiple values to be stored within one registry key.

An example of the use of **LISTBOX** is the **Windows Firewall: Define Port Exceptions** policy setting, located in **Computer Configuration\Administrative Templates\Network\Network Connections\Windows Firewall**, under both the **Domain Profile** and **Standard Profile** nodes.

By default, when LISTBOX policy settings are used in multiple GPOs, the list of the values written to the registry is defined entirely by the last GPO applied ("last writer wins"). The additive keyword changes the behavior so that the aggregate values from multiple GPOs are applied. For example, suppose that you have two GPOs with the following values for the **Windows Firewall: Define Port Exceptions** policy setting:

- GPO x. Sets values A, B, and C.
- GPO y. Sets values D, E, and F.

Without the additive keyword, only the values D, E, and F in GPO y are applied to the registry of targeted computers. This is because GPO y is the last GPO applied. With the additive keyword, all of the values from both GPOs are applied: A, B, C, D, E, and F.

The LISTBOX ADDITIVE behavior in Windows XP with SP2 gives you more flexibility in determining the values that will be applied to targets in a given Active Directory container. For example, the ability to disable policy settings from lower precedent GPOs is a key tenet of Group Policy management. Prior to Windows XP with SP2, the disabled behavior had no effect for policy settings that used LISTBOX ADDITIVE. Furthermore, prior to Windows XP with SP2, LISTBOX ADDITIVE was only found in several, rarely used policy settings affecting offline files. In Windows XP with SP2, LISTBOX ADDITIVE is used more prominently including policy settings that affect Windows Firewall and Internet Explorer.

Therefore, in Windows XP with SP2, the disabled behavior now functions as expected: removing any entries inherited from other lower precedent GPOs.

For example, suppose GPOs with values A, B, and C is normally applied to all computers in a domain. But, as an OU administrator, you wish to prevent these values from taking effect on computers in your OU. You can create a new GPO to disable values A, B, and C. This is illustrated in the following example:

- GPO x. Enables values A, B, and C for all computers in the domain.
- GPO z. Disables values A, B, and C. for all computers in your OU.
- GPO y. Enables values D, E, and F for all computers in the domain.

In this scenario, the values D, E, and F are applied to the computers in your OU.

Because, the disabled functionality only works when using the latest version of the gptext.dll, all policy settings that use LISTBOX ADDITIVE are enclosed by the #if version > 5 ...#end if construct. This eliminates the possibility of having multiple administrators experiencing different disabled behavior if they are using earlier versions of gptext.dll. Consequently, the policy settings that use LISTBOX ADDITIVE are *not* visible when editing a GPO from a computer running Windows Server 2003,

Windows XP with SP1, or Windows 2000 operating systems.

If you have administrative workstations that are not using Windows XP with SP2, you will need to install a hotfix in order to manage these policy settings. For more information, see the "'String too long...' Hotfix for Earlier Operating Systems or Service Packs" section later in this document.

**Managing Policy Settings on Earlier Operating Systems or Service Packs**
The .adm files that use the LISTBOX ADDITIVE syntax do not fully load on earlier versions of the Group Policy Object Editor (gpedit), which is present by default in Windows Server 2003, Windows XP with SP1, and Windows 2000. If attempted, multiple error messages will appear when the system.adm and inetres.adm files are loaded in earlier versions of gpedit.

The error message is "The following entry in the [strings] section is too long and has been truncated." This occurs because earlier versions of gpedit cannot correctly handle the "#if version >= 5 / #endif" construct in the inetres.adm and system.adm files. Although clicking **OK** on all the pop-up error messages does result in the .adm files loading correctly, the new Windows XP SP2 policy settings that use the LISTBOX syntax will not be displayed. (This problem does not occur on computers running Windows XP with SP2 or computers that have been updated with the latest version of gpedit.)

This issue is of particular significance because of the way .adm files are distributed through a domain. By default, when a GPO is opened, a comparison is made between the timestamps of the .adm files stored in the GPO being edited and those on the local computer. If the local .adm files have a more recent timestamp then they are uploaded to the domain controller and replicated throughout the domain. From that point, all earlier versions of gpedit use the new .adm files. This scenario is illustrated in the following steps.

1. Install Windows XP SP2 on the administrative computer.

2. Open existing GPO in Group Policy Object Editor. The .adm file stored on the administrative computer is uploaded to domain controller. GPO is "upgraded" to Windows XP SP2.

3. This .adm file is replicated to all domain controllers in the domain.

4. If the GPO is opened by other administrative workstations in the domain that are not running Windows XP with SP2 or the latest version of gpedit, error messages appear.

**"String Too Long..." Hotfix for Earlier Operating Systems and Service Packs**
If you or other administrators in your organization are going to manage policy settings on computers running earlier operating systems or service packs (for example, Windows Server 2003 or Windows XP with SP1), you need to install a hotfix in order for policy settings to appear correctly in the Group Policy Object Editor.

These hotfixes are available for the following:

- Windows Server 2003

- Windows XP with SP1

- Windows 2000

To obtain these hotfixes, see article 842933, ""The following entry in the [strings] section is too long and has been truncated" error message when you try to modify or to view GPOs in Windows Server 2003, Windows XP Professional, or Windows 2000," in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkId=4441.

If you are going to manage policy settings from workstation computers running Windows XP with SP2 only, you will be able to manage policy settings without applying any hotfixes. For example, you will be able to run the Group Policy Object Editor and view all the new policy settings delivered with Windows XP SP2.

**Important**

Opening a GPO on a computer running Windows XP with SP2 causes all other administrative workstations to use the new .adm files (note that no changes need be made to the GPO for this to occur). This will generate error messages when earlier versions of gpedit are loaded. For more information about this issue, see article 842933, ""The following entry in the [strings] section is too long and has been truncated" error message when you try to modify or to view GPOs in Windows Server 2003, Windows XP Professional, or Windows 2000," in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkId=4441.

By installing the hotfix for Windows Server 2003, Windows XP with Service Pack 1, and Windows 2000, you ensure that the Windows XP SP2 .adm files load correctly on these platforms.

# Language Reference for Administrative Template Files

Applies To: Windows Server 2003

This section includes a complete reference guide for using the .adm language to create policy settings.

Each .adm file can contain zero or more policy settings, and each policy setting in turn can contain zero or more parts. The .adm language includes the following components:

- Comments
- Strings
- CLASS
- CATEGORY
- POLICY
- PART
- ITEMLIST
- ACTIONLIST

**.Adm File Language Versions**

You can specify that any part of your .adm file be evaluated only in specific versions of the Group Policy editing tools. Table 5 lists the versions of the Group Policy editing tools.

**Table 5 Versions of Group Policy Editing Tools**

| Operating System(s) | Version | Type |
|---|---|---|
| Windows XP SP2 | 5.0 | Group Policy |
| Windows Server 2003 and Windows XP | 4.0 | Group Policy |
| Windows Server 2000 | 3.0 | Group Policy |
| Windows NT® 3.x and 4.x | 2.0 | System Policy |
| Windows 95 | 1.0 | System Policy |

**Comments**

You can use two methods to add comments to an .adm file. You can precede the comment either with a semicolon (;) or with two forward slashes (//). You can place comments at the end of any valid line.

**Strings**

To add strings to an .adm file, precede the text with two exclamation points (!!). At the end of the .adm file, all strings must be defined in the [strings] section. The strings must be enclosed in quotation marks ("). Optionally, you can enclose a variable name or hard-coded string in quotation marks.

**Example**

```
POLICY 34]!!LimitSize
   EXPLAIN!!LimitSize_Explain   ; This string is stored in the strings section
   TIP1 "Limit Profile Size to"   ; This string is hard coded

[strings]
LimitSize="Limit profile size"
LimitSize_Explain="Limits the size of user profiles"
```

**Best Practice**
Place all strings in the [strings] section of the .adm file. This facilitates conversion of the .adm file to other languages (that is, for localization), as you only need to modify the [strings] section of an .adm file to port it to different languages.

**CLASS**
This component defines where your policy setting is displayed in the Group Policy Object Editor.

The first entry in the .adm file is the keyword **CLASS**. This specifies whether the subsequent entries should be displayed under the **Computer Configuration** or **User Configuration** node of Group Policy Object Editor.

**Syntax**
The **CLASS** syntax is as follows:

```
CLASS name
```

**Name**
This defines the name of the **CLASS**, which must be MACHINE or USER.

If the .adm file contains a **CLASS** other than the valid classes (MACHINE or USER), the errors are ignored when loaded in Group Policy Object Editor.

**Example**
The following examples illustrate the use of the CLASS component.

```
CLASS MACHINE
CLASS USER
```

| **Note** |
| --- |
| You can define multiple CLASS USER or CLASS MACHINE sections in an .adm file. When the file is processed, all the CLASS USER sections are merged, and all CLASS MACHINE sections are merged. However, for ease of ongoing .adm file management, it is recommended that you define CLASS USER or CLASS MACHINE once. |

**CATEGORY**
After you define the **CLASS** component, you can use the **CATEGORY** component to display a node name under which your policy setting is displayed in the Group Policy Object Editor.

| **Note** |
| --- |
| You can create child nodes by nesting a CATEGORY within another CATEGORY. |

**Syntax**
To specify a **CATEGORY**, use the following syntax.

```
CATEGORY!!name
    KEYNAME key name
    [policy definition statements]
END CATEGORY
```

**name**
The **CATEGORY** name as it should appear in the Group Policy Object Editor list box. Optionally, you can enclose the variable name in quotation marks ("). Names with spaces must be enclosed in quotation marks.

**key name**
The key name is an optional path to the registry key to use for the **CATEGORY**.

Do not use **HKEY_LOCAL_MACHINE** or **HKEY_CURRENT_USER** in the registry path as the preceding **CLASS** statement specifies the keys to use. If you specify a key name, all child categories, policies, and parts will use this key name, unless they specifically provide a key name of their own. Names with spaces must be enclosed in double quotation marks.

If a key name is not specified and if no higher level category specifies a key name, each policy in this category must specify its own key name: otherwise, the key name for the next category that does specify a key name will be used.

**policy definition statements**
A **CATEGORY** can include zero or more **POLICY** statements. A **POLICY** definition statement cannot appear more than once within a single category, as shown in the following sample code.

**Example**

```
CLASS USER

; The following categories will be displayed
; under User Configuration

CATEGORY !!Desktop
   KEYNAME "Software\Policies\System"

   ; <INSERT POLICIES HERE>

   CATEGORY !!InternalApps
   KEYNAME "Software\Policies\InternalApps"

    ; <INSERT POLICIES HERE>

   END CATEGORY
END CATEGORY

[strings]
Desktop="Desktop Settings"
InternalApps="Line of Business Apps settings"
```

**Supported Tag**

The Group Policy Object Editor uses the **Supported** tag to populate the **Requirements** field. This tag informs the Group Policy administrator about the platforms or applications for which the policy setting is supported. For example, many of the policy settings included in the System.adm file use a **Supported** tag that specifies a specific service pack release. Often, the string used for the Supported tag will make reference to multiple operating system or service packs.

While operating system components generally use an operating system or service pack reference in this field, applications – which can be updated outside the release of a service pack – can refer to a specific version of an application. The **Supported** tag is an essential element in the data presented to Group Policy administrators to ensure they are equipped with the right information to make informed decisions about the use of the policy setting.

Because your .adm file may be localized, it is highly recommended that the **Supported** tag use the **!!Stringname** construct, which allows the referenced string to be localized easily. In addition, since the **Supported** tag is only supported in Windows XP and later operating systems, it should be enclosed within a **Version** construct, as follows (this ensures that the Windows 2000 version of Group Policy Object Editor does not attempt to interpret the **Supported** tag):

```
#if version >= 4
        SUPPORTED!!SUPPORTED_MyApplication
    #endif
```

**CATEGORY Keywords**

The valid keywords for **CATEGORY** are:

- **KEYNAME**
- **CATEGORY**
- **POLICY**
- **END**
- **SUPPORTED**

| **Note** |
| --- |
| If you have a CATEGORY defined with a default KEYNAME in it, and the same category is found again later in the .adm file, that same default KEYNAME is still in effect. This means that you can get an error message about KEYNAME being defined twice, when it was actually just defined in the same category earlier. To remove the error condition, remove the duplicate category entry. |

**POLICY**
To identify a policy setting that the user can modify, use the keyword **POLICY**. The policy and its associated controls are displayed in a dialog box that administrators use to set the state of the policy. You can use multiple **POLICY** key names under one **KEYNAME**.

The following examples illustrate the syntax of POLICY.

**Syntax**

```
POLICY name
   [KEYNAME key name]
   [EXPLAIN help string]
   [VALUENAME value name]
   [CLIENTEXT guid]
   [part definition statements]
END POLICY
```

**name**
The name of the policy as it should be displayed in the Group Policy Object Editor namespace.

**key name**
This is an optional path to the registry key to use for the category. Do not include **HKEY_LOCAL_MACHINE** or **HKEY_CURRENT_USER** in the registry path as the preceding CLASS statement determines which of these keys is used.

If you specify a key name, all **PART** definition statements will use this key name unless they specifically provide a key name of their own.

If a key name is not specified and if no higher level category specifies a key name, each policy in this category must specify its own key name: otherwise, the key name for the next category that does specify a key name will be used.

**help string**
The Help string is the text displayed in the **Explain** tab of the dialog box for the policy setting.

**value name**
Value name is the registry value to modify. Selecting this option sets the value as a REG_DWORD of 1. Clearing the option removes the registry value. To specify values other than the default values, use the **VALUEON** and **VALUEOFF** statements directly following the corresponding **VALUENAME** statement. These statements are specified as follows:

```
VALUEON on value
VALUEOFF off value
```

When you use these statements, the behavior is modified such that if the administrator selects the option, the value is set to on value. If the administrator clears the option, the value is set to off value.

**guid**
This is an optional value that specifies the globally unique identifier (GUID) of the snap-in extension.

**part definition statements**
A policy can contain zero or more **PART** statements to specify various options, including drop-down list boxes, text boxes, and text in the lower pane of the Group Policy Object Editor.

**POLICY Example**

```
CLASS MACHINE

CATEGORY!!DiskQuota

   KEYNAME "Software\Policies\MS\DiskQuota"

   POLICY!!DQ_Enable
      EXPLAIN !!DQ_Enable_Help
      VALUENAME "Enable"
         VALUEON  NUMERIC 1
         VALUEOFF NUMERIC 0
      CLIENTEXT {3610eda5-77ef-11d2-8dc}

      PART!!DQ_EnableTip1    TEXT
      END PART
   END POLICY

END CATEGORY
```

```
[strings]
DiskQuota="Disk Quotas"
DQ_Enable="Enable disk quotas"
DQ_Enable_Help="Enables and disables disk quota management"
DQ_EnableTip1="Enable disk quotas for all NTFS volumes"
```

**POLICY Keywords**
The valid keywords for POLICY are:

- **KEYNAME**
- **PART**
- **VALUENAME**
- **VALUEON**
- **VALUEOFF**
- **ACTIONLISTON**
- **ACTIONLISTOFF**
- **END**
- **HELP**
- **CLIENTEXT**
- **POLICY**

**PART**
Use **PART** to specify various options, such as drop-down list boxes, text boxes, and text in the lower pane of Group Policy Object Editor.

For a simple policy where you only need to set a registry key to either 1 or 0, you do not need to use **PART**. PART allows a richer system administrator experience, and collects more information from the administrator through simple controls.

**Syntax**

```
PART name part-type type-dependent data
  [KEYNAME key name]
  [VALUENAME value name]
END PART
```

**name**
Specifies the **PART** name as it should appear in Group Policy Object Editor. You can enclose it in quotation marks ("). Names with spaces must be enclosed in quotation marks (").

**part-type**
A policy **PART type**. Table 6 lists the valid types for **POLICY**.

**Table 6 Policy PART Types**

| Type | Description |
|---|---|
| CHECKBOX | Displays a check box. The value is set in the registry with the **REG_DWORD** type. The value is other than zero if the check box is selected and zero if it is not selected. |
| COMBOBOX | Displays a combo box. |

| | |
|---|---|
| DROPDOWNLIST | Displays a combo box with a drop-down list style. The user may choose only one of the entries supplied. |
| LISTBOX | Displays a list box with Add and Remove buttons. This is the only PART type that can be used to manage multiple values under one key. |
| EDITTEXT | Displays a text box that accepts alphanumeric text. The text is set in the registry with either the REG_SZ or the REG_EXPAND_SZ type. |
| TEXT | Displays a line of static text. There is no associated registry value with this PART type. |
| NUMERIC | Displays a text box with an optional spin control that accepts a numeric value. The value is set in the registry with the REG_DWORD type. |

**type-dependent data**
This is information about the **PART**.

**key name**
This is an optional path to the registry key to use. Do not include **HKEY_LOCAL_MACHINE** or **HKEY_CURRENT_USER** in the registry path as the preceding **CLASS** statement determines which of these keys is used.

If no key name is specified, the previous key name in the hierarchy is used.

**value name**
The value name indicates the registry value to modify. Selecting this option sets the value to a REG_DWORD of 1, and clearing the option removes the registry value. If you want to specify values other than the default values, use the **VALUEON** and **VALUEOFF** statements directly following the corresponding **VALUENAME** statement. You specify these statements as follows:

```
VALUEON on value
VALUEOFF off value
```

**Keywords**
The valid keywords for **PART** are:

- **CHECKBOX**
- **TEXT**
- **EDITTEXT**
- **NUMERIC**
- **COMBOBOX**
- **DROPDOWNLIST**
- **LISTBOX**
- **END**
- **CLIENTEXT**
- **PART**

**Using PART Types to Add Controls to the User Interface**
Using the valid keywords along with the **PART** component allows you to add text and various user interface controls to the properties page of the policy.

Because much of the syntax is related, the next section presents a task-based approach to writing the syntax for these **PART** types used to create the user interface elements above.

Using the different **PART** types, you can add text and controls to enhance a policy setting. These types need to be used with the **PART** component as previously defined.

**CHECKBOX PART Type**

This **PART** type displays a check box on the **Property** page of a policy setting. The value is set in the registry with the REG_DWORD type. The default behavior is as follows:

- By default the check box is not selected.
- A check box writes the value 1 to the registry if it is selected and 0 if it is not selected.

**Syntax**

```
PART text CHECKBOX
   VALUENAME value name
END PART
```

**text**

This represents the text to be displayed on the right of the check box that you are creating. You can hard code it and enclose it in quotation marks (") or you can make the string a variable by putting **!!** in front of the variable name.

**value name**

Indicates the registry value to which the selected value will be written. Selecting the option sets the value as a REG_DWORD of 1. Clearing the option removes the registry value. To specify values other than the default values, use the **VALUEON** and **VALUEOFF** statements directly following the corresponding **VALUENAME** statement. These statements are specified as follows:

```
VALUEON on value
VALUEOFF off value
```

When you use these statements, the behavior is modified such that if the administrator selects the option, the value is set to on value. If the administrator clears the option, the value is set to off value.

To override the default behavior:

To have the check box selected by default use **DEFCHECKED**. In the preceding sample, the syntax would be:

```
PART !!SampleChkBox_NotChked CHECKBOX
         DEFCHECKED
           VALUENAME "test1"
    END PART
```

You can use **VALUEON** and **VALUEOFF**. This example accomplishes the following::

- Writes the string "Enabled" to the registry when the check box is selected.
- Writes a numeric value of 12 when the check box is not selected.

```
    PART !!SampleChkBox_NotChked CHECKBOX
         VALUENAME "test1"
         VALUEON "Enabled"
         VALUEOFF NUMERIC 12
       END PART
```

To modify more than one registry key, use an ACTIONLIST.

The valid keywords for **CHECKBOX** are:

- KEYNAME
- VALUENAME
- VALUEON
- VALUEOFF
- ACTIONLISTON

- ACTIONLISTOFF
- DEFCHECKED
- CLIENTEXT
- END

**TEXT PART Type**
The **PART** type **TEXT** can be used to display text on the **Property** page of a policy setting. Text uses the following syntax.

```
PART text TEXT
END PART
```

**text**
Text that is to be displayed is entered here. You can hard code it and enclose it in quotation marks ("), or you can make the string a variable by putting **!!** before the variable name.

The following example illustrates the use of **TEXT**. The **Disable Active Desktop** policy deactivates Active Desktop and prevents users from enabling or disabling Active Desktop, or from modifying the configuration.

**TEXT Example**

```
POLICY !!NoActiveDesktop

   KEYNAME "Software\Microsoft\Windows\CurrentVersion\Policies\Explorer"
   EXPLAIN!!NoActiveDesktop_Help
     VALUENAME "NoActiveDesktop"

     PART !!NoActiveDesktop_Tip    TEXT
     END PART

END POLICY
```

The valid keyword for **TEXT** is **END**.

**EDITTEXT PART Type**
The **EDITTEXT** option allows the user to input alphanumeric text into an edit field. The text is set in the registry with the REG_SZ type.

**Syntax**

```
PART !!text EDITTEXT
   VALUENAME value name
END PART
```

**text**
Text to be displayed is entered here. You can hard code it and enclose it in quotation marks (") or you can make the string a variable by putting two explanation points (**!!**) before the variable name. This text is displayed on the left side of the edit box.

**value name**
The value name indicates the registry value to which the users input entered in the Edit Text box will be written.

Table 7 lists the options for **EDITTEXT.**

**Table 7 Options for EDITTEXT**

| Option | Description |
|---|---|
| DEFAULT value | Specifies the initial string to place in the edit field. If this option is not specified, the field is initially empty. |
| MAXLEN value | Specifies the maximum length of a string. The string in the edit field is limited to this length. |
| REQUIRED | Specifies that the Group Policy Object Editor does not allow a policy containing this PART to be enabled, unless a value has been entered for this PART. |

| | |
|---|---|
| OEMCONVERT | Sets the ES_OEMCONVERT style in the edit field so that typed text is mapped from ASCII to OEM and back. ES_OEMCONVERT converts text entered in the edit control. The text is converted from the Windows character set (ASCII) to the OEM character set and then back to the Windows set. This ensures proper character conversion when the application calls the CharToOem <JavaScript:hhobj_1.Click()> function to convert an ASCII string in the edit control to OEM characters. This style is most useful for edit controls that contain file names. |
| EXPANDABLETEXT | Specifies that the text is set in the registry with the REG_EXPAND_SZ type. By default, the text is set in the registry with the REG_SZ type |

The valid keywords for **EDITTEXT** are:

- **KEYNAME**
- **VALUENAME**
- **DEFAULT**
- **REQUIRED**
- **MAXLENGTH**
- **OEMCONVERT**
- **END**
- **EXPANDABLETEXT**
- **CLIENTEXT**

**EDITTEXT Example**

An example of use of the **PART** component with **EDITTEXT** and **TEXT** follows:

```
CLASS USER
CATEGORY !!DesktopLockDown

  KEYNAME "Software\Policies\System"
   POLICY !!Wallpaper
    EXPLAIN !!Wallpaper_Explain

     PART !!Wallpaper_Tip1       TEXT
     END PART

     PART !!Wallpaper_Filename   EDITTEXT
        VALUENAME Wallpaper
        MAXLEN 60
     END PART

   END POLICY

END CATEGORY

[strings]
DesktopLockDown="Desktop Settings"
Wallpaper="Desktop Wallpaper"
Wallpaper_Explain="Used to set the desktop wallpaper"
Wallpaper_FileName="Filename"
Wallpaper_Tip1="Specify UNC Path for selected wallpaper"
```

In the preceding example, the text entered into the edit field is written to the registry key **HKEY_CURRENT_USER\Software\Policies\System\Wallpaper**. The text can be a maximum of 60 characters.

When this policy setting is **Not Configured** or **Disabled**, this key is not written.

**EXPANDABLETEXT Example**
The following example writes a value to registry with data type REG_EXPAND_SZ.

For example:

```
PART!!MyVariable    EDITTEXT EXPANDABLETEXT
VALUENAME ValueToBeChanged
END PART
```

**REQUIRED Example**
The following example generates an error if the user does not enter a value when required.

```
PART!!MyVariable    EDITTEXT REQUIRED
  VALUENAME ValueToBeChanged
END PART
```

**MAXLEN Example**
The following example specifies the maximum length of text.

```
PART!!MyVariable    EDITTEXT
  VALUENAME ValueToBeChanged
  MAXLEN 4
END PART
```

**DEFAULT Example**
The following example specifies a default value. This can be used for text or numeric data.

```
PART!!MyVariable    EDITTEXT
  DEFAULT !!MySampleText
  VALUENAME ValueToBeChanged
END PART
```

**NUMERIC PART Type**
Displays an edit field with an optional spinner control (an up-down control) that accepts a numeric value.

**NUMERIC Syntax**

```
PART text NUMERIC
   VALUENAME value name
   MIN value
   MAX value
   DEFAULT value
   SPIN value
END PART
```

**text**
This represents the text to be displayed on the right of the spin control that you are creating. You can hard code it and enclose it in quotation marks (") or you can make the string a variable by putting !! before the variable name.

**value name**
Indicates the registry value to which the selected value will be written.

**NUMERIC Default Behavior**
The default behavior for the **NUMERIC** PART type is as follows:

- The value is set in the registry as a REG_DWORD type.
- You can optionally have the value written as a REG_SZ type by using the **TXTCONVERT** keyword.

Table 8 shows the options for the NUMERIC type.

**Table 8 Options for NUMERIC**

| Option | Description |
|---|---|
| DEFAULT value | Specifies the initial numeric value for the edit field. If this option is not specified, the field is initially empty. |

| | |
|---|---|
| MAX value | Specifies the maximum value for the number. The default value is 9999. |
| MIN value | Specifies the minimum value for the number. The default value is 0. |
| REQUIRED | Specifies that the Group Policy Object Editor does not allow a policy containing this PART to be enabled unless a value has been entered for this PART. |
| SPIN value | Specifies increments to use for the spinner control. The default is SPIN 1. SPIN 0 removes the spinner control. |
| TXTCONVERT | Writes values as REG_SZ strings ("1", "2", or "128") rather than as binary values. |

The valid keywords for **NUMERIC** are:

- KEYNAME
- VALUENAME
- MIN
- MAX
- SPIN
- DEFAULT
- REQUIRED
- TXTCONVERT
- END
- CLIENTEXT

**Examples of NUMERIC Use**

The following example illustrates use of the **NUMERIC PART** type using the **DEFAULT** option.

```
PART!!MyVariable    NUMERIC
  DEFAULT 5
  VALUENAME ValueToBeChanged
END PART
```

The following example illustrates use of the minimum and maximum valid values for a variable.

```
PART!!MyVariable    NUMERIC
  MIN 100
  MAX 999
  DEFAULT 55
  VALUENAME ValueToBeChanged
END PART
```

The following example illustrates use of the **NUMERIC PART** type using **SPIN**. In this case, increments of 100 are used for the spin control.

```
PART !!ProfileSize    NUMERIC REQUIRED SPIN 100

        VALUENAME "MaxProfileSize"
        DEFAULT 30000
        MAX     30000
        MIN     300
END PART
```

The following example illustrates use of the **NUMERIC PART** type using the **TXTCONVERT** option, which writes values as **REG_SZ** strings (such as "60") instead of binary values.

```
PART !!ScreenSaverTimeOutFreqSpin    NUMERIC DEFAULT 900

        MIN 0 MAX 599940 SPIN 60
        TXTCONVERT
        VALUENAME "ScreenSaveTimeOut"
END PART
```

**COMBOBOX PART Type**

This **PART** type displays a combo box. It accepts the same options as **EDITTEXT**, as well as the **SUGGESTIONS** option, which begins a list of suggestions to be placed in the drop-down list. **SUGGESTIONS** are separated with spaces and must be enclosed in quotation marks (") when a value includes spaces. If a suggestion name includes white space, it must be enclosed in quotation marks. The list ends with **END SUGGESTIONS**.

**Example**

The following example illustrates the use of the **SUGGESTIONS** option.

```
SUGGESTIONS
  Alaska Alabama Mississippi "New York"
END SUGGESTIONS
```

**Keywords**

The valid keywords for **COMBOBOX** are:

- **KEYNAME**
- **VALUENAME**
- **DEFAULT**
- **SUGGESTIONS**
- **REQUIRED**
- **MAXLENGTH**
- **OEMCONVERT**
- **END**
- **NOSORT**
- **EXPANDABLETEXT**
- **CLIENTEXT**
- **END**

**DROPDOWNLIST PART Type**

Displays a combo box with a drop-down list style. The user may choose only one of the entries supplied.

| **Note** |
| --- |
| GPMC requires that you define the key name and value name before you specify DROPDOWNLIST. |

**DROPDOWNLIST Syntax**

**DROPDOWNLIST** uses the following syntax.

```
PART !!text DROPDOWNLIST
   ITEMLIST
     NAME name VALUE value
     ..
     NAME name VALUE value
```

```
   END ITEMLIST
END PART
```

**text**
This represents the text to be displayed on the right of the spin control that you are creating. You can hard code it and enclose it in quotation marks (") or you can make the string a variable by putting **!!** in front of the variable name.

**name**
This is text that will be displayed in the drop-down list for a particular item.

**value**
The value to be written to the specified registry key if this item is selected. Values are assumed to be strings, unless they are preceded by **NUMERIC**. The following example shows both string and numeric values:

```
VALUE "Some value"
VALUE NUMERIC 1
```

The valid keywords for **DROPDOWNLIST** are:

- KEYNAME
- VALUENAME
- REQUIRED
- ITEMLIST
- END
- NOSORT
- CLIENTEXT

**LISTBOX PART Type**
The LISTBOX **PART** component specifies various options such as drop-down list boxes, text boxes, and text in the lower pane of the Group Policy Object Editor. **LISTBOX** accepts the options shown in Table 9.

**Table 9 LISTBOX Options**

| LISTBOX Option | Description |
|---|---|
| ADDITIVE | By default, the content of list boxes overrides any values set in the target registry. This means that a control value is inserted in the policy file that causes existing values to be deleted before the values set in the policy file are merged. If this option is specified, existing values are not deleted, and the values set in the list box is in addition to whatever values exist in the target registry. |
| EXPLICITVALUE | This option makes the user specify the value data and the value name. The list box shows two columns, one for the name and one for the data. This option cannot be used with the VALUEPREFIX option. |
| VALUEPREFIX prefix | The prefix you specify is used in determining value names. If a prefix is specified, the prefix and an incremented integer are used, instead of the default value naming scheme described previously. For example, a prefix of "SampleName" generates the value names "SampleName1", "SampleName2", and so on. The prefix can be empty (""), which causes the value names to be "1", "2", and so on. |

By default, only one column appears in the list box, and for each entry a value is created whose name and value are the same. For instance, a "name" entry in the list box creates a value called "name" that contains data called "name". When using a **LISTBOX**, use the **ADDITIVE** keyword unless you have a specific reason not to do so.

The valid keywords for **LISTBOX** are:

- **KEYNAME**

- **VALUEPREFIX**
- **ADDITIVE**
- **NOSORT**
- **EXPLICITVALUE**
- **EXPANDABLETEXT**
- **END**
- **CLIENTEXT**

| **Note** |
| --- |
| Windows XP SP2 fixed issues relating to the LISTBOX ADDITIVE functionality. For more information, see the "Changes to LISTBOX ADDITIVE" section in this document. |

**ACTIONLIST**
You can use an action list to specify a set of arbitrary registry changes to make in response to a control being set to a particular state.

**Syntax**
The **ACTIONLIST** syntax is as follows:

```
ACTIONLIST
[KEYNAME key name]
VALUENAME value name
VALUE value
END ACTIONLIST
```

**key name**
This is an optional path to the registry key. Do not include HKEY_LOCAL_MACHINE or HKEY_CURRENT_USER in the registry path as the preceding CLASS statement determines which of these keys is used. If no key name is specified, the previous key name in the hierarchy is used.

**value name**
Indicates the registry value to modify. Selecting this option sets the value to a REG_DWORD of 1, and clearing the option removes the registry value. If you want to specify values other than the default values, use the VALUEON and VALUEOFF statements directly following the corresponding VALUENAME statement. You specify these statements as follows:

```
VALUEON on value
VALUEOFF off value
```

**value**
Values are treated as strings unless they are preceded by NUMERIC, as in the following examples:

```
VALUE "Some value"
VALUE NUMERIC 1
```

If **VALUE** is followed by **DELETE** (for example, **VALUE DELETE**), the registry entry is deleted.

Table 10 lists the two variants for **ACTIONLIST** that can be used with **POLICY** and **CHECKBOX**.

**Table 10 Variants for ACTIONLIST**

| Variant | Description |
| --- | --- |
| ACTIONLISTON | Specifies an optional action list to be used if the check box is selected. |
| ACTIONLISTOFF | Specifies an optional action list to be used if the check box is not selected. |

**ACTIONLIST Example**
The following example illustrates the use of **ACTIONLISTON** and **ACTIONLISTOFF**.

```
POLICY  "Deny connections requests"
    EXPLAIN "If enabled, TS will stop accepting connections"
    ACTIONLISTON
       VALUENAME "fDenyTSConnections"    VALUE NUMERIC 1
    END ACTIONLISTON
    ACTIONLISTOFF
       VALUENAME "fDenyTSConnections"    VALUE NUMERIC 0
    END ACTIONLISTOFF
END POLICY
```

**Additional Elements**
The .adm language supports the following elements:

**KEYNAME**
The **KEYNAME** keyword is used within a **CATEGORY** to define which key within the registry is modified as a result of an action here. **KEYNAME** should be followed by the registry path to the key that contains the value that you want to change. Do not include **HKEY_LOCAL_MACHINE** or **HKEY_CURRENT_USER** in the registry path as the preceding CLASS statement determines which of these keys is used.

If the **KEYNAME** contains a space, you must enclose the string in quotation marks (").

**VALUENAME**
Defines the options available within a **POLICY**. First identify the registry value that is to be modified as a result of using the keyword **VALUENAME**. For example, **VALUENAME** MyFirstValue.

The following example illustrates the use of **VALUENAME**. The Disable Boot / Shutdown / Logon / Logoff status messages policy prevents the display of system status messages.

```
POLICY!!DisableStatusMessages
    KEYNAME "Software\Microsoft\Windows\CurrentVersion\Policies\System"
       EXPLAIN!!DisableStatusMessages_Help
       VALUENAME "DisableStatusMessages"
END POLICY
```

Unless you specify otherwise, the value is written in the following format when the user checks or clears the option:

- Checked. Uses a REG_DWORD type with a value of 1.
- Cleared. Removes the value.

You can specify options other than these defaults by using **VALUEOFF** and **VALUEON**. If the option is to be selected within the lower pane of the Group Policy Object Editor, the **VALUENAME** needs to be within a **PART** scope.

**CLIENTEXT**
The **CLIENTEXT** keyword is used to specify which client-side extension to the Group Policy Object Editor needs to process the particular settings on the client computer. By default, the registry extension processes all settings configured under the Administrative Templates node. The **CLIENTEXT** keyword changes the default behavior and causes the specified extension to process these settings after the registry extension has placed them in the registry.

**CLIENTEXT** must be used within either the **POLICY** scope or the **PART** scope and should follow the **VALUENAME** statement.

The following example illustrates use of **CLIENTEXT**.

```
POLICY !!DQ_Enforce

            #if version >= 4
                 SUPPORTED !!SUPPORTED_Win2k
            #endif

            EXPLAIN !!DQ_Enforce_Help
                 VALUENAME "Enforce"
                 VALUEON  NUMERIC 1
            VALUEOFF NUMERIC 0
            CLIENTEXT {3610eda5-77ef-11d2-8dc5-00c04fa31a66}
```

```
        END POLICY
```

The GUID that follows the **CLIENTEXT** keyword is the GUID of the client-side extension. The client-side extensions are listed in the registry under **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\WindowsNT \CurrentVersion\Winlogon\GPExtensions**.

**VALUEON and VALUEOFF**

You can use **VALUEON** and **VALUEOFF** to write specific values based on the state of the option. To enable this functionality, you can write the .adm file as described in the following examples:

```
KEYNAME key name
   POLICY!!MyPolicy
     VALUENAME ValueToBeChanged
     VALUEON "Turned On" VALUEOFF "Turned Off"
   END POLICY
```

```
KEYNAME key name
   POLICY!!MyPolicy
     VALUENAME ValueToBeChanged
     VALUEON 5 VALUEOFF 10
   END POLICY
```

**Using Simple Policies and Policies with the VALUEOFF and VALUEON Statements**

This section presents two examples that illustrate the difference between using the default policy states and specifying **VALUEON** and **VALUEOFF** statements. There is a significant difference between the two example policies.

**Example 1**

In this example, no explicit **VALUEON** or **VALUEOFF** statements are used. This means that the Administrative Templates use the default behavior when the user changes the state of this policy.

```
POLICY!!EnableSlowLinkDetect
    EXPLAIN !!EnableSlowLinkDetect_Help
    KEYNAME "Software\Policies\Microsoft\Windows\System"
    VALUENAME "SlowLinkDetectEnabled"
END POLICY
```

Table 11 lists the default behavior.

**Table 11 Example 1 Policy Defaults**

| State | Behavior |
|---|---|
| Policy setting enabled | A DWORD with the value 1 is written to the registry. |
| Policy setting disabled | The registry value is deleted. |
| Policy setting not configured | Nothing is changed in the registry. |

Note the policy-disabled state. The value is not written to the registry with the value of 0—instead it is explicitly deleted. This means that a component reading the policy will not find the value in the registry, and will fall back to using the default in the code.

**Example 2**

In this example, the state values are explicitly defined, so when the user changes the policy, the Administrative Templates use these values.

```
POLICY!!EnableSlowLinkDetect
    EXPLAIN!!EnableSlowLinkDetect_Help
    KEYNAME "Software\Policies\Microsoft\Windows\System"
    VALUENAME "SlowLinkDetectEnabled"
        VALUEON NUMERIC 1
        VALUEOFF NUMERIC 0
END POLICY
```

Table 12 lists the behaviors in Example 2.

**Table 12 Example 2 Policy Defaults**

| State | Behavior |
|---|---|
| Policy setting enabled | A DWORD with the value 1 is written to the registry. |
| Policy setting disabled | A DWORD with the value 0 is written to the registry. |
| Policy setting not configured | Nothing is changed in the registry. |

**EXPLAIN**

The **EXPLAIN** keyword is used to provide online Help text for a specific Group Policy. In Windows 2000, the **Properties** page for each policy setting includes an **Explain** tab, which provides details about the policy settings.

Each Group Policy that you create should include one **EXPLAIN** keyword, followed by at least one space, and then the **EXPLAIN** string in quotation marks (") or a reference to the Help string. For example:

```
POLICY!!Pol_NoConfigCache
#if VERSION >= 3
EXPLAIN!!Pol_NoConfigCache_Help
#endif
    VALUENAME "NoConfigCache"
    PART!!Lbl_NoConfigCacheHelp1       TEXT
END PART
END POLICY
       .....

[Strings]
Pol_NoConfigCache_Help="Prevents users from changing the automatic
synchronization behavior at logoff."
```

In the preceding example, Help is offered for one of the Offline Files options. The **EXPLAIN** keyword wrapped in the **#if VERSION** allows this .adm file to be used with the Windows 2000 Group Policy Object Editor (version 3).

**Line Breaks**

To start text on a new line or to create a line break, use this syntax:

```
\n = Starts a new line
\n\n = Creates a line break
```

**#If Version for Version Comparison**

The **IF VERSION** conditional statement is used to control the display of certain policy settings and features in the Administrative Templates node, based on the version of the Group Policy Object Editor that you are using. **IF VERSION** allows for part of the .adm files to be conditionally parsed and ignored by earlier versions of the Group Policy Object Editor tool. For example, the SUPPORTED tag is not supported on versions of the Group Policy Object Editor earlier than version 4. For this reason any statement using the SUPPORTED tag should be enclosed by #If Version...#endif.

You can specify that any part of your .adm file be evaluated only in specific versions of the Group Policy editing tools, as shown in Table 5, in the "*.Adm File Language Versions*" section of this document.

To compare versions, use the following syntax:

```
#if Version (operator) x
#endif
```

The valid operators are listed in Table 13.

**Table 13 Valid Operators for the Version Statement Number**

| Operator | Signifies |
|---|---|
| > (GT) | Greater than. For example, a > b means a is greater than b. |
| < (LT) | Less than. For example, a < b means a is less than b. |
| == (EQ) | Equal. For example, a == b means a is equal to b. |
| != (NE) | Not equal. |
| >= (GTE) | Greater than or equal to. For example, a >= b means a is greater than or equal to b. |

| | |
|---|---|
| <= (LTE) | Less than or equal to. For example, a <= b means a is less than or equal to b. |

**.Adm File String/Tag Limits**

Various restrictions apply to .adm files and settings. Table 14 provides a complete list of these restrictions.

**Table 14**

| File String | Tag Limits |
|---|---|
| Maximum string length for Explain text | 4096 |
| Maximum string length for Category Explain text | 255 |
| Maximum string length for EDITTEXT string | 1023 |

# Related Links (Using Administrative Template Files with Registry-Based Group Policy)

Applies To: Windows Server 2003

To learn more about Group Policy, see the following resources:

- Microsoft Group Policy Home Page at http://go.microsoft.com/fwlink/?LinkId=15088.
- Group Policy ADM Files at http://go.microsoft.com/fwlink/?linkid=31057.
- Managing Windows XP Service Pack 2 Features Using Group Policy on the TechNet Web site at http://go.microsoft.com/fwlink/?LinkId=31974.
- Group Policy in the Microsoft Platform SDK at http://go.microsoft.com/fwlink/?LinkId=26258.
- Group Policy Management Console on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=17912.
- Group Policy in Windows Server 2003 on the Microsoft Web site at http://go.microsoft.com/fwlink/?linkid=29269.
- Article 816662, "Recommendations for Managing Group Policy Administrative Template (.adm) Files," in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkId=4441.
- The Core Group Policy Technical Reference collection of the "Windows Server 2003 Technical Reference" at http://go.microsoft.com/fwlink/?LinkId=31788.
- The "Designing a Managed Environment" book of the *Microsoft Windows Server 2003 Deployment Kit* on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=30650.

# Windows Server 2003 Group Policy Infrastructure

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Intended for system administrators, architects, and others who need to create and manage Group Policy settings, this paper explains Group Policy infrastructure and shows how the new Group Policy Management Console (GPMC) fits into this infrastructure. The paper includes detailed information about Group Policy processing as well as many best practices useful to the Group Policy administrator.

Introduced in Windows® 2000 Server, Group Policy provides directory-based desktop configuration management. With Group Policy, you can specify settings for registry-based policies, security, software installation, scripts, folder redirection, Remote Installation Services, and Internet Explorer maintenance. The Windows Server 2003 family of operating systems, extends Group Policy in a number of ways— through GPMC, which includes scripting interfaces, Group Policy Results, Group Policy Modeling, and more.

The Group Policy settings that you create are contained in a Group Policy object (GPO). By linking a GPO with selected Active Directory® service system containers—sites, domains, and organizational units—you can apply these settings to the users and computers in those Active Directory containers. To create GPOs, you use GPMC in conjunction with the Group Policy Object Editor, an MMC snap-in, also known previously as the Group Policy snap-in, Group Policy Object Editor, or GPedit.

**Administrative Requirements for Using Group Policy**

In order to use of all of its features, Group Policy requires Active Directory and client computers running Windows 2000 or later. To set Group Policy for a selected Active Directory container, you must have a Windows 2000 or Windows Server 2003 domain controller installed, and you must have read and write permission to access the system volume of domain controllers (Sysvol folder) and modify rights to the currently selected directory container. The system volume folder is automatically created when you install a domain controller (or promote a server to domain controller).

| **Note** |
| --- |
| Group Policy depends on Active Directory; therefore, it is crucial to understand Active Directory and its structure. It is highly recommended that you familiarize yourself with Active Directory concepts before implementing Group Policy. To learn about Active Directory, see the Active Directory white papers at http://www.microsoft.com/ad. Information about planning and implementing Active Directory is available from the Windows Deployment and Resource Kits page at http://www.microsoft.com/reskit. |

**GPMC System Requirements**

GPMC can manage both Windows 2000 and Windows Server 2003 domains with Active Directory. In either case, the computer on which the tool itself runs must be running Windows Server 2003 or Windows XP Professional (with Windows XP Service Pack 1 and the Microsoft .NET Framework). Note: When installing GPMC on Windows XP Professional with SP1, a post SP1 hotfix is required. This hotfix (Q326469) is included with GPMC. GPMC Setup prompts you to install Windows XP hotfix Q326469 if it is not already present.

**Feedback on this Paper**

If you have any comments about this paper, contact mailto:gpdocs@microsoft.com.

### In This White Paper


What's New About Group Policy in Windows Server 2003

Overview of Group Policy Infrastructure and Mechanics

Migrating GPOs Across Domains

Scripting Group Policy Tasks

Delegating Group Policy

Group Policy Extension Snap-ins

Group Policy Modeling and Results

Group Policy Processing

Group Policy Replication and Domain Controller Selection (Group Policy Infrastructure)

Local Group Policy (Group Policy Infrastructure)

Group Policy Loopback Support

Design Considerations for Organizational Unit Structure and Use of Group Policy Objects

Deployment Considerations

IntelliMirror Features without Active Directory

Migrating Policy-Enabled Clients from Windows NT 4.0 to Windows 2000 or Windows Server 2003

Appendix A: Security Settings and User Rights

Appendix B: Group Policy Storage

Appendix C: WMI Filtering

Glossary (Group Policy Infrastructure)

Related Links (Group Policy Infrastructure)

# What's New About Group Policy in Windows Server 2003

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section summarizes new features in Windows Server 2003 Group Policy. The biggest change for Group Policy in Windows Server 2003 is the introduction of GPMC, the new solution for Group Policy management that helps you manage an enterprise more cost-effectively. It consists of a new Microsoft Management Console (MMC) snap-in and a set of scriptable interfaces for managing Group Policy. GPMC is available for download from the Microsoft GPMC Web site at http://www.microsoft.com/windowsserver2003/gpmc. This paper assumes you are using GPMC.

**Group Policy Management Console**
GPMC simplifies the management of Group Policy by providing a single place for managing core aspects of Group Policy. It addresses the top Group Policy deployment requirements by providing the following functionality:

- A user interface that makes Group Policy much easier to use.
- Backup/restore of GPOs.
- Import/export and copy/paste of GPOs and Windows Management Instrumentation (WMI) filters.
- Simplified management of Group Policy–related security.
- HTML reporting for GPO settings and Resultant Set of Policy (RSoP) data.
- Scripting of policy-related tasks that are exposed within this tool (not scripting of settings within a GPO).

GPMC is used to create, view, and manage GPOs while the Group Policy Object Editor is used to edit GPOs.

More information about GPMC is contained throughout this paper. For additional information about GPMC including step-by-step instructions for completing tasks, see the following resources:

- Group Policy Administration using the Group Policy Management Console. This white paper provides additional technical details of functionality in GPMC.
- GPMC Help. Available when you install GPMC, this provides step-by-instructions for GPMC tasks and addresses key concepts in GPMC.
- Microsoft GPMC Web site at http://www.microsoft.com/windowsserver2003/gpmc/. This site provides links to the latest GPMC resources including downloading information.

**New policy settings**
Windows Server 2003 includes more than 200 new policy settings. The new Windows Server 2003 policy settings allow administrators to control the behavior of:

- System restore, error reporting, PC Health.
- Terminal server.
- Networking such as SNMP, Quality of Service (QoS), personal firewall, and dialup connections.
- DNS and net logon.
- Roaming user profiles and Group Policy.
- Control panel.

- Windows Media® Player.
- Wireless configuration.
- Software restriction policy.

**Administrative Templates Changes**
For Administrative Templates policy settings, Group Policy Object Editor provides explain text directly in the Web view of the console. You also can find this explain text by double-clicking the policy setting and then clicking the **Explain text** tab. In either case, this text shows operating system requirements, defines the policy setting, and includes any specific details about the effect of enabling or disabling the policy setting.

Because new Administrative Template policy settings have been added that only work on specific versions of the operating system such as Windows XP Professional or Windows Server 2003, you can view only the Administrative Template policy settings that might be applied in your users' work environment, based on the "supported" keyword in each Administrative Template (.adm) file. For example, you may want to edit only policy settings that could be applied on client computers running Windows 2000 Service Pack 3. In Group Policy Object Editor, you can specify these options in the **Filtering** dialog box, available by clicking a node in the Administrative Templates section, clicking the **View** menu, and then clicking **Filtering**. For more information, see the section "Using Administrative Templates" in this paper.

**Command Line Refresh of Policy**
Administrators can now refresh policy settings from the command line using Gpupdate, which replaces **secedit /refreshpolicy** in Windows 2000. Gpupdate gives administrators better control and flexibility in refreshing policy. For more information, see the section "Refreshing Policy from the Command Line" in this paper.

**WMI Filtering**
WMI makes a large amount of data, such as hardware and software inventory, settings, and configuration information, available for a target computer. WMI retrieves data from the registry, drivers, file system, Active Directory, SNMP, Windows Installer, SQL, networking, and Exchange. WMI Filtering in Windows Server 2003 allows you to create queries based on this data. These queries (also called WMI filters) determine which users and computers receive all of the policy configured in the GPO where you create the filter. This functionality lets you target Group Policy based on a significant number of different properties of the target. In most organizations only senior administrators would actually create WMI filters; other administrators would simply access the WMI filters that have been created for their domain. For more information, see the section "WMI Filtering" in this paper.

**Tools for Best Practice Organizational Unit Design**
Redirusr.exe (for user accounts) and Redircomp.exe (for computer accounts) are two new tools included with Windows Server 2003 that enable you to change the default location where new user and computer accounts are created so you can more easily scope GPOs directly to newly created user and computer objects. By running Redirusr.exe and Redircomp.exe once for each domain, the domain administrator can specify the organizational units into which all new user and computer accounts are placed at the time of creation. For more information, see the section "Redirecting the Users and Computers Containers in Windows Server 2003 Domains in this paper.

**Forest Trust**
The Windows Server 2003 family introduces a new feature called Forest Trust that enables you to authenticate and authorize access to resources from separate, networked forests. With trusts established between forests, you can manage Group Policy throughout your enterprise, which provides greater flexibility especially in large organizations. For more information, see the section "Using Group Policy Features Across Forests" in this paper.

**Domain Rename**
The ability to rename a domain provides you with the flexibility to make important changes to your forest structure and namespace as the needs of your organization change. Renaming domains can accommodate acquisitions, mergers, name changes, or reorganizations. Domain rename allows you to:

1. Change the DNS and NetBIOS names of any domain in the forest (including the forest root domain).
2. Restructure the position of any domain in the forest (except the forest root domain).

You can only rename domains in a forest where all of the domain controllers are running Windows Server 2003 and the forest functional level has been raised to Windows Server 2003. For more information, see Windows Server 2003 Domain Rename Tools at http://www.microsoft.com/windowsserver2003/downloads/domainrename.mspx.

**Restore GPOs tool**
This is a new command-line tool intended for failure recovery. The tool, dcgpofix.exe, restores the default GPOs to their original state (that is, the default state after initial installation). For more information, see "Troubleshooting Windows Server 2003 Group Policy" available from the Microsoft GPMC Web site at http://www.microsoft.com/windowsserver2003/gpmc/.

**Wireless Support**
A new Wireless Network (IEEE 802.11) Policies Group Policy extension allows you to configure wireless network settings that are part of Group Policy for Computer Configuration. Wireless network settings include the list of preferred networks, Wired Equivalent Privacy (WEP) settings, and IEEE 802.1X settings. These settings are downloaded to targeted domain members, making it much easier to deploy a specific configuration for secure wireless connections to wireless client computers.

**Software Restriction Policy Settings**
Software restriction policy settings address the need to regulate unknown or untrusted software. With the rise in the use of networks, the Internet, and e-mail for business computing, users find themselves exposed to new software in a variety of ways. Users must constantly make decisions about running unknown software. Viruses and Trojan horses often intentionally misrepresent themselves to trick users into running them. It is difficult for users to make safe choices about which software they should run.

With software restriction policy settings, you can protect your computing environment from untrusted software by identifying and specifying which software is allowed to run. You can define a default security level of unrestricted or disallowed for a GPO so that software is either allowed or not allowed to run by default. You can make exceptions to this default security level by creating rules for specific software. For example, if your default security level is set to disallowed, you can create rules that allow specific software to run. For more information, see the section "Software Restriction Policy Settings" in this paper.

**Internet Explorer Enhanced Security Configuration**
Internet Explorer Enhanced Security Configuration, also known as Internet Explorer hardening, is enabled by default on computers running Windows Server 2003. It can be managed using Group Policy in an enterprise environment to ensure consistent trusted sites and security settings on targeted server computers or to disable the feature on specific servers. For example, you may wish to ensure that Internet Explorer Enhanced Security Configuration is reapplied on a specific computer if the local administrator on that computer turns it off using the Optional Component Manager in the **Windows Components Wizard** (available from **Add or Remove Programs**.) In addition, it's likely that you will want to manage computers or groups of computers in your organization by defining a set of trusted sites and/or a specific security level for sites in the **Internet** or **Trusted sites** zones. For more information, see "Using Group Policy and Internet Explorer Enhanced Security Configuration" later in this document and Managing Internet Explorer Enhanced Security Configuration, available from the Microsoft Group Policy Web site at http://www.microsoft.com/grouppolicy.

# Overview of Group Policy Infrastructure and Mechanics

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy uses a document-centric approach to creating, storing, and associating policy settings. Similar to the way in which Microsoft Word stores information in .doc files, Group Policy settings are contained in GPOs. GPOs are linked to the following Active Directory containers: sites, domains, or organizational units. The settings within the GPOs are then evaluated by the affected clients, using the hierarchical nature of Active Directory.

**Linking Group Policy Objects to Active Directory Containers**
GPOs cannot be linked directly to users, computers, or security groups. They can only be linked to sites, domains and organizational units.

A given GPO can be linked to more than one site, domain, or organizational unit. Conversely, a given site, domain, or organizational unit can have multiple GPOs linked to it. In the case where multiple GPOs are linked to a particular site, domain, or organizational unit, you can prioritize the order of precedence in which these GPOs are applied.

By linking GPOs to sites, domains, and organizational units, you can implement Group Policy settings for as broad or as narrow a portion of the organization as you want:

- A GPO linked to a site applies to all users and computers in the site.
- A GPO applied to a domain applies directly to all users and computers in the domain and by inheritance to all users and computers in child organizational units. Note that policy is not inherited across domains.
- A GPO applied to an organizational unit applies directly to all users and computers in the organizational unit and by inheritance to all users and computers in child organizational units.

To link a GPO to a site, domain, or organizational unit, use GPMC. In the console tree, locate the site, domain or organizational unit and then choose **Link an Existing GPO** or **Create and Link a GPO Here**.

Note that GPOs are stored in domains not in organizational units. For example, if you create and link a new GPO for an organizational unit, GPMC is actually completing two steps at once: creating a GPO in the domain and then linking that GPO to that organizational unit. The link is not a component of the GPO; it is a component of the container to which it is linked. Therefore, if you want to delegate the ability to manage links for a given container, it must be delegated on that container, not the GPO. In the GPMC tree view, GPO links on a given container are shown as child nodes of that container.

Although you can link a site, domain, or organizational unit to a GPO in another trusted domain, this is not generally recommended for performance reasons because of the potential delay of processing GPOs at logon.

**Group Policy Hierarchy**
By default, Group Policy is inherited and cumulative, and it affects all computers and users in an Active Directory container. GPOs are processed according to the following order:

1. **Local GPO**. Each computer has exactly one GPO that is stored locally, shared by all users of that computer. This processes for both computer and user Group Policy processing.
2. **Site**. Any GPOs that have been linked to the site that the computer belongs to are processed next. Processing is in the order that is specified by the administrator, on the **Linked Group Policy Objects** tab for the site in GPMC. The GPO with the lowest **link order** is processed last, and therefore has the highest precedence.
3. **Domain**. Processing of multiple domain-linked GPOs is in the order specified by the administrator, on the **Linked Group Policy Objects** tab for the domain in GPMC. The GPO with the lowest **link order** is processed last, and therefore has the highest precedence.
4. **Organizational units**. GPOs that are linked to the organizational unit that is highest in the Active Directory hierarchy are processed first, then GPOs that are linked to its child organizational unit, and so on. Finally, the GPOs that are linked to the organizational unit that contains the user or computer are processed.

At the level of each organizational unit in the Active Directory hierarchy, one, many, or no GPOs can be linked. If several GPOs are linked to an organizational unit, their processing is in the order in which GPOs are linked to the organizational unit. For example, if you link three GPOs to an organizational unit the first GPO you added has the highest precedence and overwrites the settings of all other GPOs. Alternatively, you can specify the order on the **Linked Group Policy Objects** tab for the organizational unit in GPMC.

This order means that the local GPO is processed first, and GPOs that are linked to the organizational unit of which the computer or user is a direct member are processed last, which overwrites settings in the earlier GPOs if there are conflicts. (If there are no conflicts, then the earlier and later settings are merely aggregated.)

In GPMC, you can view the precedence order of inherited GPOs for a given site, domain or organizational unit by navigating to the Group Policy Inheritance tab for any site, domain, or organizational unit.

**Managing Inheritance of Group Policy**

You can further control precedence and how GPO links are applied to specific domains, sites, or organizational units by:

- **Changing the link order.** Within each domain, site, and organizational unit, the link order controls when links are applied. To change the precedence of a link, you can change the link order, moving each link up or down in the list to the appropriate location. The link with the higher order (with 1 being the highest order) has the higher precedence for a given site, domain, or organizational unit. For example, if you add six GPO links and later decide that you want the last one that you added to have highest precedence, you can move the GPO link to the top of the list. However, the link order of an inherited GPO cannot be altered.

- **Blocking Group Policy inheritance.** You can block policy inheritance for a domain or organizational unit. Using block inheritance prevents GPOs linked to higher sites, domains, or organizational units from being automatically inherited by the child-level. By default, children inherit all GPOs from the parent, but it is sometimes useful to block inheritance. For example, if you want to apply a single set of policies to an entire domain except for one organizational unit, you can link the required GPOs at the domain level (from which all organizational units inherit policies by default), and then block inheritance only on the organizational unit to which the policies should not be applied. Blocking does not affect Local GPOs.

- **Enforcing a GPO link.** You can specify that the settings in a GPO link should take precedence over the settings of any child object by setting that link to **Enforced** (formerly known as "no override"). GPO-links that are enforced cannot be blocked from the parent container. Without enforcement from above, the settings of the GPO links at the higher level (parent) are overwritten by settings in GPOs linked to child organizational units, if the GPOs contain conflicting settings. With enforcement, the parent GPO link always has precedence. Note that **Enforce** policy options always take precedence over Block Inheritance.

- **Disabling a GPO link**. By default, processing is enabled for all GPO links. You can completely block the application of a GPO for a given site, domain, or organizational unit by disabling the GPO link for that domain, site, or organizational unit. Note that this does not disable the GPO itself, and if the GPO is linked to other sites, domains or organizational units, they will continue to process the GPO, if their links are enabled.

- **Disabling user and/or computer settings**. A GPO may have its user settings disabled, its computer settings disabled, or all settings disabled. By default, neither user settings nor computer settings are disabled on a GPO.

| **Note** |
| --- |
| A GPO link may be **enforced**, or **disabled**, or both. By default, a GPO link is neither enforced nor disabled. If the link is enforced and disabled, the disabled link has precedence. |

Figure 1 below shows a sample domain structure to illustrate how GPOs can be applied to containers in Active Directory.



**Figure 1. Group Policy and Active Directory**

**Filtering the Scope of the Group Policy Object**
You can further refine which groups of computers and users a particular GPO influences by using security groups or WMI filtering.

**Security Filtering**
Security filtering is a way of refining which users and computers will receive and apply the settings in a GPO. Using security filtering, you can narrow the scope of a GPO so that it applies only to a single group, user, or computer by specifying that only certain security principals within a container where the GPO is linked apply the GPO. Security filtering determines whether the GPO as a whole applies to groups, users, or computers; it cannot be used selectively on different settings within a GPO.

In order for the GPO to apply to a given user or computer, that user or computer must have both **Read** and **Apply Group Policy (AGP)** permissions on the GPO, either explicitly, or effectively though group membership.

By default, all GPOs have **Read** and **AGP** both **Allowed** for the Authenticated Users group. The Authenticated Users group includes both users and computers. This is how all authenticated users receive the settings of a new GPO when it is applied to an organizational unit, domain or site. Therefore, the default behavior is for every GPO to apply to every Authenticated User. By default, Domain Admins, Enterprise Admins, and the local system have full control permissions, without the Apply Group Policy ACE. However, administrators are members of Authenticated Users, which means that they will receive the settings in the GPO by default.

You can change these permissions to limit the scope to a specific set of users, groups, or computers within the organizational unit, domain, or site. Group Policy Management manages these permissions as a single unit, and displays the security filtering for the GPO on the GPO Scope tab. Using GPMC, you can add and remove groups, users, and computers to be used as security filters for each GPO. In addition, security principals used for security filtering also appear on the **Delegation** tab for a GPO as having **Read (from Security Filtering)**, since they have read access to the GPO.

To modify security filtering, you add or remove groups in the Security Filtering section on the Scope tab of a GPO. In practice, you don't have to set the two access control entries (ACEs), because GPMC sets both for you when you set security filtering. In addition, The **Read** and **AGP** permissions are visible separately, and able to be set independently of one another, through the access control list (ACL) editor. In GPMC, the **Security Filtering** section of the **Scope** tab of a GPO shows only whether the GPO will apply. If you want to see the permissions separately, you can open the ACL editor by clicking the **Advanced** button on the **Delegation** tab for the GPO.

To prevent a GPO from applying to a specified group requires removal of the AGP ACE from that group. In the ACL editor, if you remove the AGP ACE (clear the **Allow check box**) for Authenticated Users, you can then explicitly grant this permission to individual security groups that should receive the policy settings. Alternatively, you could set AGP to **Deny** for certain classes of users, such as administrators, that will never need that policy.

| **Note** |
| --- |
| Use the **Deny** ACE with caution. A **Deny** ACE setting for any group has precedence over any **Allow** ACE given to a user or computer because of membership in another group. |

**Best Practice:** If you disallow Apply Group Policy for a GPO for some users, consider also disallowing Read access to those users. When the Read ACE is allowed and the Apply Group Policy is not, the GPO is still processed by the user even though it is not applied to the user. Therefore, to improve performance, you should remove the Read Access Control Entry to prevent the

user from processing the GPO. In addition, removing Read access increases security. With Read access allowed, it is possible for an inquisitive user with considerable knowledge of Active Directory to read the contents of that GPO, even if it's not applied to them. This may not be desirable in some cases, such as a GPO for a human resources group. It might be advisable to limit Read access on GPOs that affect the HR users to only those users.

Security groups and DACLs are also used to delegate control of GPOs, as explained in the section "Delegating Group Policy."

**Note**

Granting Read and AGP is not sufficient to ensure that the GPO is processed for a user or computer. The GPO also has to be linked to a site, domain or organizational unit containing the user or computer, directly or through inheritance.

A GPO with security filtering set to Read and AGP doesn't necessarily apply to all security principals that have security filtering. It only applies to them if those user or computer objects are in the container or child container that is linked to the GPO.

The location of a security group in Active Directory is irrelevant to security filtering and, more generally, irrelevant to Group Policy processing.

**WMI Filtering**

WMI filters allow you to dynamically determine the scope of GPOs based on attributes of the target computer.

When a GPO that is linked to a WMI filter is applied on the target computer, the filter is evaluated on the target computer. If the WMI filter evaluates to false, the GPO is not applied (except if the client computer is running Windows 2000, in which case the filter is ignored and the GPO is always applied). If the WMI filter evaluates to true, the GPO is applied.

WMI makes data about a target computer available for administrative use. Such data can include hardware and software inventory, settings, and configuration information. For example, WMI exposes hardware configuration data such as CPU, memory, disk space, and manufacturer, as well as software configuration data from the registry, drivers, file system, Active Directory, the Windows Installer service, networking configuration, and application data.

A WMI filter consists of one or more queries based on this data. If all queries are true, the GPO linked to the filter will be applied. The queries are written using the WMI Query Language (WQL), a SQL-like language. Queries can be combined with AND and OR logical operators to achieve whatever effect the administrator wants. Each query is executed against a particular WMI namespace. When you create a query, you must specify the namespace. The default is root\CIMv2, which is appropriate for most WMI queries. For more information, see Windows Management Instrumentation in the Microsoft Platform SDK at http://msdn.microsoft.com/library/default.asp?url=/library/en-us/wmisdk/wmi/wmi_start_page.asp.

The WMI filter is a separate object from the GPO in the directory. To apply a WMI filter to a GPO, you link the filter to the GPO. This is shown in the WMI filtering section on the Scope tab of a GPO. Each GPO can have only one WMI filter, however the same WMI filter can be linked to multiple GPOs.

WMI filters, like GPOs, are stored on a per-domain basis. A WMI filter and the GPO it is linked to must be in the same domain.

**Note**

Client support for WMI filters exists only on Windows XP, Windows Server 2003, and later operating systems.
Windows 2000 clients will ignore any WMI filter and the GPO is always applied, regardless of the WMI filter.

WMI filters are only available in domains that have at least one Windows Server 2003 domain controller. In an environment consisting only of Windows 2000 domains, the WMI filter node in GPMC is not shown.

Here are some sample uses of WMI Filters.

- Services – computers where DHCP is turned on.
- Registry – computers that have this registry key populated.
- Hardware inventory – computers with a Pentium III processor.
- Software inventory – computers with Visual Studio® .NET installed.
- Hardware configuration – computers with network adapters on interrupt level 3.

- Software configuration – computers with multi-casting turned on.
- Associations – computers that have any service dependent on SNA service.
- Ping – computers that can ping Server1 in less than 100 milliseconds.

WMI filtering allows administrators to filter the application of a GPO by attaching one set of Windows Query Language query to a GPO. The queries can be written to query WMI for multiple items. If the query returns true for all queried items, then the GPO will be applied to the target user or computer. The WMI filter applies to every setting in the GPO, so you will want to create separate GPOs specifically for WMI filtering.

There are two distinct parts to WMI filtering in Windows Server 2003:

- Server administration of WMI filters. This includes GPMC, changes to the Group Policy Object Editor, and filter specifications.
- Client side processing. (WMI supports batch processing of WMI filters.)
- For more information about WMI filtering including sample filters for specific scenarios, see "Appendix C" in this paper.

**MMC Snap-in Extension Model**

The nodes of the Group Policy Object Editor are also MMC snap-in extensions. These extensions include Administrative Templates, Scripts, Security Settings, Software Installation, Folder Redirection, Remote Installation Services, and Internet Explorer Maintenance. Extension snap-ins may in turn be extended. For example, the Security Settings snap-in includes several extension snap-ins. Developers can also create their own MMC extensions to the Group Policy Object Editor to provide additional policy settings.

For more information about creating MMC extensions, see the Microsoft Management Console section of the Microsoft Platform SDK documentation at: http://www.microsoft.com/msdownload/platformsdk/sdkupdate/.

By default, all the available Group Policy Object Editor extensions are loaded when you start the Group Policy Object Editor. You can modify this default behavior by creating a custom MMC console, or by using policy settings to control the behavior of MMC itself. MMC options are accessed under the **User Configuration\Administrative Templates\Windows Components\Microsoft Management Console** node. For more information, see:

"Specifying Group Policy to Control the Behavior of MMC and Snap-ins", later in this document.

**Group Policy Object Editor Namespace**

The root node of the Group Policy Object Editor is displayed as the name of the GPO and the domain to which it belongs, in the following format:

*GPO Name [DomainName.com] Policy*

For example:

Default Domain Policy [HQ-RES-DC-01.Contoso.com] Policy

**Computer Configuration and User Configuration**

Below the root node, the namespace is divided into two parent nodes: Computer Configuration and User Configuration. These are the parent nodes that you use to configure Group Policy settings. Computer-related Group Policy is applied when the operating system boots and during the periodic refresh cycle, explained later in this document. User-related Group Policy is applied when users log on to the computer and during the periodic refresh cycle.

**Extensions to the Group Policy Object Editor**

Three nodes exist under the Computer Configuration and User Configuration parent nodes: Software Settings, Windows Settings, and Administrative Templates. The Software Settings and Windows Settings nodes contain extension snap-ins that extend either or both of the Computer Configuration or User Configuration nodes. Most of the extension snap-ins extend both of these nodes, but frequently with different options. The Administrative Templates node namespace contains all policy settings pertaining to the registry; it can be extended by using .adm files.

The Group Policy extension snap-ins include:

- **Administrative Templates**. This extension contains all registry-based policy settings, including those for the

Windows 2000 and Windows Server 2003 operating systems and their components as well as any registry-based policy settings provided by applications. You use these policy settings to mandate registry settings that control the behavior and appearance of the desktop, the operating system components, and applications that provide registry-based policy. This node uses .adm files to specify the registry settings that can be modified through the Group Policy Object Editor user interface. For more information about .adm files, see "Administrative Templates" later in this paper.

- **Security Settings**. The Security Settings extension is used to set security options for computers and users within the scope of a GPO. You can define local computer, domain, IP security settings, wireless configuration, and software restriction policy settings. For more information about security settings, see "Security Settings" and "Appendix A: Security Settings and User Rights" later in this paper.

- **Software Installation**. You can use the Software Installation snap-in to centrally manage software in your organization. You can assign and publish software to users and assign software to computers. For more information about software installation, see "Software Installation", later in this document.

- **Scripts**. Scripts are used to automate tasks at computer startup and shutdown, and at user logon and logoff. You can use any language supported by Windows Script Host. These include the Microsoft Visual Basic® development system, Scripting Edition (VBScript), JavaScript, PERL, and MSDOS®-style batch files (.bat and .cmd). See "Scripts", later in this document and Microsoft Windows Script Web site at http://www.microsoft.com/scripting for more information.

- **Remote Installation Services**. Remote Installation Services (RIS) is used to control the behavior of the Remote Operating System Installation feature as displayed to client computers. See "Remote Installation Services", later in this document.

- **Internet Explorer Maintenance**. Internet Explorer Maintenance is used to manage and customize Internet Explorer on computers running Windows 2000 or later. You can also export settings for Windows 95, Windows 98, and Windows NT® 4.0-based client computers (the settings are exported into an .ins and .cab file format for those platforms). Administrators can set options for Browser UI, connections, URLs, proxy settings, security zones, Favorites, and other options. See "Internet Explorer Maintenance", later in this document.

- **Folder Redirection**. You can use folder redirection to redirect special directories on Windows 2000 or Windows Server 2003 from their default user profile location to an alternate location on the network. These special folders include My Documents, Application Data, Desktop, and the **Start** menu. See "Folder Redirection", later in this document.

For more information about extending the functionality of the Group Policy Object Editor see "Implementing Registry-Based Group Policy" at http://www.microsoft.com/windows2000/techinfo/howitworks/management/rbppaper.asp.

**Client-side Extensions to Group Policy**

Some of the Group Policy Object Editor extensions also include *client-side extensions*. These extensions are DLLs that are responsible for implementing Group Policy at the client computers. For more information about the client-side extensions, see the "Client-side Processing of Group Policy" section later in this paper.

**Group Policy Storage**

A GPO is a virtual object. The policy setting information of a GPO is actually stored in two locations: the Group Policy container (GPC) and the Group Policy template (GPT). The Group Policy container is an Active Directory container that stores GPO properties, including information about version, GPO status, and a list of components that have settings in the GPO. The Group Policy template is a directory structure within the file system that stores Administrative Template-based policy settings, security settings, script files, and information regarding applications that are available for Software Installation. The Group Policy template is located in Sysvol in the \Policies sub-directory for its domain. GPOs are identified by their globally unique identifiers (GUIDs) and stored at the domain level. Replication of a GPO to other domain controllers happens through two different mechanisms. The Group Policy container is replicated by using Active Directory replication, whereas the Group Policy template is replicated using File Replication service (FRS). The settings from a GPO are only applied when the Group Policy container and Group Policy template are synchronized. For additional information about storage of Group Policy information, see"Appendix B: Group Policy Storage", later in this paper.

# Migrating GPOs Across Domains

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Taking a GPO in a given domain and creating a new GPO that contains the same set of policies in a different domain is a central function of GPMC.

Although the collection of settings in a GPO is logically a single entity, the data for a single GPO is stored in multiple locations and in a variety of formats; some data is contained in Active Directory and other data (of various types) is stored on the Sysvol file share on the domain controllers. This means that copying GPOs is not as simple as taking a directory and copying it from one computer to another.

In addition to the complex way in which GPO data is stored, certain policy data may be valid in one domain but be invalid in the domain that the GPO is being copied to. For example, Security Identifiers (SIDs) stored in security policy settings are often domain specific. In addition, Universal Naming Convention (UNC) paths for folder redirection or software installation policies may not work properly if the data in the GPO is copied without modification to a different domain.

**Migration Tables**

The solution is to modify these references in the GPO that are domain-specific, during the import or copy operation, so that the settings in the destination GPO are written with the appropriate information for the destination domain. GPMC supports this capability using migration tables.

A migration table is a file that maps references to users, groups, computers, and UNC paths in the source GPO to new values in the destination GPO. A migration table consists of one or more mapping entries. Each mapping entry consists of a type, source reference, and destination reference. If you specify a migration table when performing an import or copy, each reference to the source entry will be replaced with the destination entry when writing the settings into the destination GPO.

The migration table will apply to any references in the settings within a GPO, whether you are performing an import or copy operation. In addition, during a copy operation, if you choose the option to preserve the discretionary access control list (DACL) on the GPO, the migration table will also apply to both the DACL on the GPO and the DACLs on any software installation settings in the GPO.

Migration tables store the mapping information as XML, and have their own file name extension, .migtable. You can create migration tables using the Migration Table Editor (MTE). The MTE is a convenient tool for viewing and editing migration tables without having to work in, or be familiar with, XML. The MTE is associated with the .migtable extension so that when you double click a migration table, it opens in the MTE. The MTE is installed with GPMC. You can also create and edit migration tables using any XML editor or using the GPMC scripting interfaces.

**GPMC as the Solution for Migrating GPOs**

There are four operations that GPMC provides to allow for archival and recovery of GPOs, and for migrating GPOs from one environment to another:

- **Copy**. A copy operation allows you to transfer settings from an existing GPO in Active Directory directly into a new GPO. The new GPO created during the copy operation is given a new GUID and is unlinked. You can use a copy operation to transfer settings to a new GPO in the same domain, another domain in the same forest, or a domain in another forest. Because a copy operation uses an existing GPO in Active Directory as its source, trust is required between the source and destination domains. Copy operations are suited for moving Group Policy between production environments, and for migrating Group Policy that has been tested in a test domain or forest to a production environment, as long as there is trust between the source and destination domains.

- **Backup**. Backing up a GPO copies the data in the GPO to the file system. The backup function also serves as the export capability for GPOs. A GPO backup can be used to restore the GPO to the backed-up state, or to import the settings in the backup to another GPO.

- **Import**. The Import operation transfers settings into an existing GPO in Active Directory using a backed-up GPO in the file system location as its source. Import operations can be used to transfer settings from one GPO to another GPO within the same domain, to a GPO in another domain in the same forest, or to a GPO in a domain in a different forest. The import operation always places the backed-up settings into an existing GPO. It erases any pre-existing settings in the destination GPO. Import does not require trust between the source domain and destination domain. Therefore it is useful for transferring settings across forests and domains that don't have trust. Importing settings into a GPO does not affect its DACL, links on sites domains or organizational units to that GPO, or a link to a WMI filter.

- **Restore**. Restoring a GPO re-creates the GPO from the data in the backup. A restore operation can be used in both of the following cases: the GPO was backed up but has since been deleted, or the GPO is live and you want to roll back to a known previous state.

Each of these operations can be performed through the GPMC user interface, or through the GPMC scripting model.

For more information, see the following resources:

- Group Policy Administration using the Group Policy Management Console White paper. Provides technical details of functionality in GPMC.

- Migrating GPOs Across Domains with GPMC. Explains how to migrate GPOs from one domain to another using GPMC.

- **GPMC Help**. Available when you install GPMC, this provides step-by-instructions for GPMC tasks and addresses key concepts in GPMC.

# Scripting Group Policy Tasks

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

GPMC provides a comprehensive set of COM interfaces for scripting many Group Policy-related operations. The interfaces are documented in the Group Policy Management Console SDK, which is located at **%programfiles%\gpmc\scripts\gpmc.chm** on any computer where you installed GPMC. (The Group Policy Management Console SDK is only available in English.)

When you install GPMC, a set of sample scripts illustrating the use of these interfaces are installed to the **%programfiles%\gpmc\scripts** directory.

The sample scripts address real-world administrative problems and scenarios. You can perform various tasks such as finding all GPOs in a domain that have duplicate names or generating a list of all GPOs in a domain whose settings are disabled or partially disabled.

| **Note** |
| --- |
| Scripted control of individual settings inside a GPO is not provided. |

# Delegating Group Policy

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

One of the features of Active Directory is its ability to delegate control of portions of the directory service. This section explains how Group Policy fits in with the delegation of sites, domains, and organizational units.

With GPMC, the following tasks can be delegated:

- Create GPOs in a domain.
- Set permissions on a GPO.
- Set policy-related permissions on site, domain or organizational unit.
    - Link GPOs to a given site, domain or organizational unit.
    - Perform Group Policy Modeling analyses on a given domain or organizational unit (but not on a site).
    - Read Group Policy Results data for objects in a given domain or organizational unit (but not on a site).
- Create WMI filters in a domain.
- Set permissions on a WMI filter.

GPMC simplifies delegation by managing the various ACEs required for a task as a single bundle of permissions for the task. If you want to see the ACL in detail, you can click the **Advanced** button on the **Delegation** tab.

The underlying mechanism for achieving delegation is the application of the appropriate DACLs to GPOs and other objects in Active Directory. This mechanism is identical to using security groups to filter the application of GPOs to various users, as described earlier in this paper.

You can also specify Group Policy to control the behavior of MMC and MMC snap-ins. For example, you can use Group Policy to manage the rights to create, configure, and use MMC consoles, and to control access to individual snap-ins.

**Using Security Groups to Delegate Group Policy**

The following table lists the default security-permission settings for a GPO:

| Groups or Users | Security permission |
|---|---|
| Authenticated User | Read with Apply Group Policy ACE |
| Domain Admins<br>Enterprise Admins<br>Creator Owner<br>Local System | Full control without Apply Group Policy ACE. |

| Note |
|---|
| By default, administrators are also authenticated users, which means that they have the **Apply Group Policy** attribute set. If this is not desired, administrators have two choices: |

- Remove Authenticated Users from the list on the security tab of the GPO, and add a new security group with the Apply Group Policy and Read attributes set to **Allow**. This new group should contain all the users that this Group Policy is intended to affect.

- Set the Apply Group Policy attribute to **Deny** for the Domain and Enterprise Admins, and possibly the Creator Owner groups. This will prevent the GPO from being applied to members of those groups. Remember that an ACE set to **Deny** always takes precedence over **Allow**. Therefore, if a given user is a member of another group that is set to explicitly **Allow** the Apply Group Policy attribute for this GPO, it will still be denied.

**Managing Group Policy Links for a Site, Domain, or Organizational Unit**

The **Group Policy** tab in the **Properties** page for a site, domain, or organizational unit allows the administrator to specify which GPOs are linked to this site, domain, or organizational unit. This property page stores the user's choices in two Active Directory properties called **gPLink** and **gPOptions**. The **gPLink** property contains the prioritized list of GPOs and the **gPOptions** property contains the **Block Policy Inheritance** setting.

To manage GPO links to a site, domain, or organizational unit, you must have read and write access to the **gPLink** and **gPOptions** properties. By default, Domain Admins have this permission for domains and organizational unit, and only Enterprise Admins and Domain Admins of the forest root domain can manage links to sites.

Active Directory supports security settings on a per-property basis. This means that a non-administrator can be given read and write access to specific properties. In this case, if non-administrators have read and write access to the **gPLink** and **gPOptions** properties, they can manage the list of GPOs linked to that site, domain, or organizational unit.

**Creating GPOs**

By default, only Domain Admins, Enterprise Admins, and Group Policy Creator Owners can create new GPOs. Creating GPOs is a user right of the Group Policy Creator Owners (GPCO) group by default but can be delegated to any group or user. There are two methods to grant a group or user this right:

- Add the user or group to membership of the **Group Policy Creator Owners** group. This was the only method available prior to GPMC.

- Explicitly grant the group or user permission to create GPOs. This method is newly available with GPMC.

You can manage this permission using the **Delegation** tab on the **Group Policy Objects** container for a given domain in GPMC. This tab shows the groups that have permission to create GPOs in the domain, including the GPCO group. From this tab, you can modify the membership of existing groups with this permission, or add new groups.

The ability to grant users permissions to create GPOs without using GPCO was added to facilitate the delegation of GPO creation to users outside the domain. Because the Group Policy Creator Owners group is a domain global group, it cannot contain members from outside the domain. Thus, prior to GPMC, this task could not be delegated to members outside the domain.

It is recommended that for users and groups within the domain, you continue to use the GPCO group to grant them GPO creation rights. If you require that users outside the domain have the ability to create GPOs, then create a new domain local group in the domain ("GPCO - External"), grant that group GPO creation rights in the domain, and then add external domain users to that group.

Adding a user to the membership of GPCO, or granting the user GPO creation permissions directly using the new method available in GPMC, is identical in terms of permissions. Users have the ability to create GPOs in the domain, but do not have permissions on GPOs created by other users. For example, granting a user the ability to create GPOs in the domain does not give the user the ability to edit or delete existing GPOs, or the ability to link the GPO to a site, domain or organizational unit.

Note that when an administrator creates a GPO, the Domain Admins group becomes the Creator Owner of the GPO. The ability to link GPOs to a site, domain or organizational unit is a permission that is specific to that site, domain or organizational unit. When delegating to non-administrators, you should also consider delegating the ability to manage the links for a specific organizational unit. The reason is that by default, non-administrators cannot manage links.

In GPMC, this permission can be managed using the Delegation tab on the site, domain or organizational unit when you click the **Link GPOs** option in the permission drop-down list box. At the individual permission level in Active Directory, this allows Read and Write access to the gPLink and gPOptions attributes on the site, domain, or organizational unit. By default, only Domain Admins and Enterprise Admins have this permission.

**Editing Group Policy Objects**

To edit a GPO, the user must have both read and write access to the GPO. (However, read-only support for opening a GPO is provided in GPMC). To edit a GPO, the user must be one of the following:

- An administrator.
- A Creator Owner.
- A user with delegated access to the GPO. That is, an administrator, or the Creator Owner, must have provided to this user both read and write access to the GPO.

By default, Domain Admins, Enterprise Admins, the operating system, and the GPO Creator Owner can edit GPOs because they have full control of GPOs without the **Apply Group Policy** attribute.

**Delegating an individual GPO**

There are five permission options on GPOs in GPMC user interface. Each corresponds to a set of individual NT permissions. The correspondence is summarized in the following table.

| Option in GPMC user interface | Corresponding NT permission in ACL Editor |
|---|---|
| Read | Allow **Read** Access on the GPO |
| Edit settings | Allow **Read**, **Write**, **Create Child Objects**, and **Delete Child Objects**. |
| Edit, delete, and modify security | Allow **Read**, **Write**, **Create Child Objects**, **Delete Child Objects**, **Delete**, **Modify Permissions**, and **Modify Owner**. This essentially grants full control on the GPO, except that the **Apply Group Policy** permission is not set. |
| Read (from Security Filtering) | This setting cannot be set directly, but appears on the delegation tab if the user has **Read** and **Apply Group Policy** permissions to the GPO. |
| Custom | Any other combination of permissions, including the use of **Deny** will show up as **Custom** in the display. GPMC can only set custom permission sets by clicking the Advanced button and opening the ACL editor. |

Permissions on a GPO are managed from the Delegation tab of that GPO.

**Specifying Group Policy to Control the Behavior of MMC extensions**

Windows Server 2003 Group Policy includes several policy settings designed to control the behavior of MMC snap-ins. For example, you can use Group Policy to manage the rights to use MMC snap-ins.

**Restricting Access to a List of Permitted Snap-ins**

Administrators can specify which MMC snap-ins may be run by the affected user and which may not. This may be specified to be inclusive, which only allows a set of snap-ins to run, or it may be set as exclusive, which does not allow a set of snap-ins to run.

To create a list of permitted snap-ins for users, enable the **Restrict users to the explicitly permitted list of snap-ins** policy. When this policy is enabled, only permitted snap-ins can be run. If this policy is disabled or not configured, all snap-ins are permitted, except those you explicitly prohibit.

This policy is available in the Group Policy console under the **User Configuration\Administrative Templates\Windows Components\Microsoft Management Console** node.

**Controlling Access to a Snap-in**

To restrict or explicitly permit access to a particular snap-in, navigate to **User Configuration\Administrative Templates\Windows Components\Microsoft Management Console\Restricted\Permitted snap-ins\Group Policy** in the console tree. In the details pane, double-click the snap-in that you want to permit or restrict, and then select an option. For more information about these policy settings, select a policy setting, and view the description in the Web view or click the explain tab in the Properties dialog box for the policy setting.

Administrators can enable the **Restrict the user from entering author mode** policy in order to prevent users from using MMC in author mode. This policy is available in the Group Policy console under the **User Configuration\Administrative Templates\Windows Components\Microsoft Management Console** node.

**Creating Custom Group Policy Object Editor Consoles**

You can create custom Group Policy MMC consoles (.msc files), which include only a subset of the Group Policy Object Editor

extensions. You can combine this with the use of the policy settings above to provide a customized tool. For example, you could create a custom Group Policy console that includes only the Security Settings extension. This allows you to define Group Policy settings in a modular fashion.

To set access permissions, use the **Security** tab on the **Properties** page of the selected GPO. These permissions allow or deny specified groups access to the GPO.

# Group Policy Extension Snap-ins

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Group Policy extension snap-ins constitutes the main nodes in the Group Policy Object Editor namespace; they are all loaded by default when the Group Policy Object Editor is started. You can modify which extensions are loaded by creating custom consoles for Group Policy, and by specifying policy settings for MMC. For more information, see "Creating Custom Group Policy Snap-in Consoles" and "Specifying Group Policy to Control the Behavior of MMC and Snap-ins" in this document.

This section presents additional information about the following topics:

- Administrative Templates
- Security Settings
- Software Installation
- Scripts (Startup/Shutdown and Logon/Logoff)
- Folder Redirection
- Internet Explorer Maintenance
- Remote Installation Services

**Administrative Templates**

Administrative templates, (or .adm files), enable administrators to control registry settings using Group Policy. Windows comes with a predefined set of Administrative template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a GPO. These .adm files are stored in two locations by default: inside GPOs in the Sysvol folder and in the %windir%\inf directory on the local computer.

As new versions of Windows are released, new policy settings are added. In addition to supporting these new settings, each successive version of Windows supports all registry policy settings that were available in earlier versions of Windows. For example, the Windows Server 2003 family supports all registry policy settings available in Windows 2000 and Windows XP.

Note that .adm files are Unicode files which consist of a hierarchy of categories and subcategories that define how the options are displayed through the Group Policy Object Editor and GPMC. They also indicate the registry locations where changes should be made if a particular selection is made, specify any options or restrictions (in values) that are associated with the selection, and in some cases, indicate a default value to use if a selection is activated.

It is important to understand that .adm files are not the actual settings that are deployed to client operating systems. The adm file is simply a template file that provides the friendly name for the setting and an explanation. This template file is used to populate the user interface. The settings that are deployed to clients are contained in the registry.pol file inside the GPO. On Windows XP and Windows Server 2003, each registry setting contains a "Supported on" tag that indicates which operating system versions support that policy setting. If a setting is specified and deployed to a client operating system that does not support that setting, the settings are ignored.

Because all successive iterations of .adm files include settings from earlier versions, and because there is no harm if a new setting is applied inadvertently to a computer running an earlier operating system that does not support that setting, it is recommended to always create and edit GPOs from a computer that has the latest .adm files available. Note that the behavior for handling .adm files in GPMC and the Group Policy Object Editor differs.

Windows Server 2003 includes the following .adm files: System.adm, Inetres.adm, Conf.adm, Wmplayer.adm, and Wuau.adm, which contain all the settings initially displayed in the Administrative Templates node.

| .Adm file | Contains | For Use on | Description |
|---|---|---|---|

| | | | |
|---|---|---|---|
| System.adm | Settings to configure the Operating System | Windows 2000 or Windows Server 2003 | Loaded by default. |
| Inetres.adm | Settings to configure Internet Explorer | Windows 2000 or Windows Server 2003 | Loaded by default. |
| Conf.adm | Settings to configure NetMeeting v3 | Windows 2000 or Windows Server 2003.<br>**Note**<br>This tool is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. | Loaded by default. |
| Wmplayer.adm | Settings to configure Windows Media Player | Windows XP, Windows Server 2003.<br>**Note**<br>This tool is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. | Loaded by default. |
| Wuau.adm | Settings to configure Windows Update | Windows 2000 SP3, Windows XP SP1, Windows Server 2003 | Loaded by default. |

**Handling .adm files in Group Policy Object Editor**

- Windows Server 2003, Group Policy Object Editor uses .adm files to display available registry-based policy settings in the Administrative Templates section of a GPO. This includes Group Policy for the Windows Server 2003 operating system and its components and for applications.

- By default it attempts to read .adm files from the GPO (from the Sysvol on the domain controller). Alternatively, the .adm file can be read from the local workstation computer. This behavior can be controlled by a policy setting.

- By default, if the version of the .adm file found on the local computer is newer (based on the time stamp of the file) than the version on the Sysvol, the local version is copied to the Sysvol and is then used to display the settings. This behavior can be controlled by a policy setting.

- If the GPO contains registry settings for which there is no corresponding .adm file, these settings cannot be seen in the Group Policy Object Editor. However, the policy settings are still active and will be applied to users or computers targeted by the GPO.

- Policy settings pertaining to a user who logs on to a given workstation or server are written to the **User** portion of the registry database under **HKEY_CURRENT_USER**. Computer-specific settings are written to the **Local Machine** portion of the registry under **HKEY_LOCAL_MACHINE**.

**Handling .adm files in GPMC**

- GPMC uses .adm files to display the friendly names of policy settings when generating HTML reports for GPOs, Group Policy Modeling, and Group Policy Results.

- By default, GPMC uses the local .adm file, regardless of time stamp. If the file is not found, then GPMC will look in the GPO's directory on Sysvol.

- The user can specify an alternate path for where to find .adm files. If specified, this takes precedence over the previous locations.

- GPMC never copies the .adm file to the Sysvol.

For more information, see Recommendations for Managing Group Policy Administrative Template (.adm) Files at

http://support.microsoft.com/default.aspx?scid=kb;en-us;816662.

**Using Administrative Templates**
For Administrative Templates policy settings, the Group Policy Object Editor provides explain text directly in the Web view of the console. You also can find this explain text by double-clicking the policy setting and then clicking the Explain text tab. In either case, this text shows operating system requirements, defines the policy setting, and includes any specific details about the effect of enabling or disabling the policy setting.

**New policy settings**
Windows Server 2003 includes more than 200 new policy settings. The new Windows Server 2003 policy settings allow administrators to control the behavior of:

- System restore, error reporting, PC Health.
- Terminal Server.
- Networking such as SNMP, QoS, personal firewall, and dialup connections.
- DNS and net logon.
- Roaming user profiles and Group Policy.
- Control panel.

To filter settings based on the "supported on" information: Open the Group Policy Object Editor, click **View**, and then click **Filtering**. Select the versions you want to show and click **OK**.

| **Note** |
| --- |
| Showing only policy settings that can be fully managed is the default setting. You can also show policy settings by supported-on information and show only configured policy settings. In Windows Server 2003, the command **Only show configured policy settings** is now contained in the Filtering dialog box, shown above. |

**True Policy Settings Compared with Group Policy Preferences**
In Windows 2000 and later, all shipping policy settings set registry keys and values in one of the following locations:

- **HKLM\Software\Policies** (preferred location).
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Policies**.
- **HKCU\Software\Policies** (preferred location).
- **HKCU\Software\Microsoft\Windows\CurrentVersion\Policies**.

Policy settings that are stored in these specific locations of the registry are known as *true policies*. Storing settings here has the following advantages:

- These trees are secure and cannot be modified by a non-administrator.
- When Group Policy changes, for any reason, these trees are cleaned, and the new policy settings are then rewritten.

This prevents the behavior that was often present in Windows NT 4.0, whereby System Policies resulted in persistent settings in the user and computer registry. The policy remained in effect until the value was reversed, either by a counteracting policy or by editing the registry. These settings are stored outside the approved registry locations above and are known as *preferences*.

All the policy settings in the System.adm, Inetres.adm, Conf.adm, Wmplayer,adm, and Wuau.adm files use registry settings in the Policies trees of the registry. This means that they will *not* cause persistent settings in the registry when the GPO that applies them is no longer in effect.

By default, only true policy settings are displayed in the Group Policy Object Editor. The following .adm files are loaded:

- System.adm: contains operating system settings
- Inetres.adm: contains Internet Explorer restrictions
- Conf.adm: contains NetMeeting settings
- Wmplayer,adm: contains Windows Media Player settings
- Wuau.adm: contains Windows Update settings

| **Note** |
| --- |
| Because of the persistent nature of non-policy settings, they should be avoided.<br><br>It is still possible for administrators to add an additional .adm file that sets registry values outside of the Group Policy trees mentioned previously. These settings might be more appropriately referred to as preferences because the user, application, or other parts of the system can also change them. In this case, the administrator is ensuring that this registry key or value is set in a particular way. Although it is possible to add any .adm file to the namespace, if you use an .adm file from a previous version of Windows, the registry keys are unlikely to have an effect; or they actually set preference settings and mark the registry with these settings; that is, the registry settings persist. |

**Creating Custom .adm Files**

It is possible to create new .adm files. For example when an application adds Group Policy support, a new .adm file may be necessary to describe the location of the appropriate registry keys and the UI exposed by the Group Policy Object Editor. Through the Group Policy Object Editor, the administrator can optionally add in additional .adm files to the GPO which, by default, will then be copied to the domain controller into the GPO directory.

To view custom .adm files in the Group Policy Object Editor:

- Right click any administrative template node, select **View** and then click **Filtering**. In the filtering dialog box, clear the check box for **Only show policy settings that can be fully managed** and click **OK**.

**Viewing Group Policy Preferences**

By default only the settings that are contained in the genuine Group Policy trees (the trees that correspond to the reserved Group Policy registry areas) are visible in the console.

To eliminate use of non-policies, you can enable the policy setting, **Enforce Show Policies Only**, available in **User Configuration\Administrative Templates**, under the **System\Group Policy** nodes.

If you enable this setting, the **Show Policies Only** command is turned on, and administrators cannot turn it off. As a result, Group Policy displays only true settings; preferences do not appear. If you disable this setting or do not configure it, the **Show Policies Only** command is turned on by default, but administrators can view preferences by turning off the **Show Policies Only** command.

In Group Policy, preferences are indicated by a red icon to distinguish them from true policy settings, which are indicated by a blue icon.

**Impact of GPO Replication**

By default, when you add a new domain to the console, GPMC uses the PDC emulator in that domain to help ensure that all administrators are using the same domain controller. For managing sites, GPMC uses the PDC emulator in the user's domain by default. You can change the default choice of domain controller using the Change Domain Controller dialog box in GPMC. If you are located at a remote site with a slow connection to the default domain controller, you may want to do this.

It is important for administrators to consider the choice of domain controller in order to avoid replication conflicts particularly because both Active Directory and FRS use multi-master replication. This is especially important to consider because GPO data resides in both Active Directory and on Sysvol, and two independent replication mechanisms must be used to replicate GPO data to the various domain controllers in the domain. If two administrators are simultaneously editing the same GPO on different domain controllers, it is possible for the changes written by one administrator to be overwritten by another administrator, depending on replication latency.

| **Important** |
| --- |
| If multiple administrators manage a common GPO, it is recommended that all administrators use the same domain controller when editing a particular GPO, to avoid collisions in FRS.<br><br>Because the Group Policy template is replicated to all domain controllers, the size of the .adm files can have an impact on network bandwidth, particularly where domain controllers are separated by slow links. |

**Security Settings**

You can define a security configuration within a GPO. A security configuration consists of settings applied to one or more security areas supported on Windows 2000 Professional, Windows XP Professional or Windows Server 2003. The specified security configuration is then applied to computers as part of the Group Policy application.

The Security Settings extension of the Group Policy Object Editor complements existing system security tools such as the **Security** tab on the **Properties** page (of an object, file, directory, and so on), and **Local Users** and **Groups** in **Computer Management**. You can continue to use existing tools to change specific settings, whenever necessary.

The security areas that can be configured for computers include the following:

- **Account Policies**. These are computer security settings for password policy, lockout policy, and Kerberos policy in domains on Windows 2000 and Windows Server 2003.

- **Local Policies**. These include security settings for audit policy, user rights assignment, and security options. Local policy allows you to configure who has local or network access to the computer and whether or how local events are audited.

- **Event Log**. This controls security settings for the Application, Security, and System event logs. You can access these logs using the Event Viewer.

- **Restricted Groups**. This allows you to control who should and should not belong to a restricted group, as well as which groups a restricted group should belong to. This allows administrators to enforce security policy settings regarding sensitive groups, such as Administrators or Payroll. For example, it may be decided that only Joe and Mary should be members of the Administrators group. Restricted groups can be used to enforce that policy. If a third user is added to the group (for example, to accomplish some task in an emergency situation), the next time policy is enforced, that third user is automatically removed from the Administrators group.

- **System Services**. These control startup mode and security options (security descriptors) for system services such as network services, file and print services, telephone and fax services, Internet and intranet services, and so on.

- **Registry**. This is used to configure security settings for registry keys including access control, audit, and ownership. When you apply security on registry keys, the Security Settings extension follows the same inheritance model as that used for all tree-structured hierarchies in Windows 2000 and Windows Server 2003 (such as Active Directory and NTFS). Microsoft recommends that you use the inheritance capabilities to specify security only at top-level objects, and redefine security only for those child objects that require it. This approach greatly simplifies your security structure and reduces the administrative overhead that results from a needlessly complex access-control structure.

- **File System**. This is used to configure security settings for file-system objects, including access control, audit, and ownership.

- **Public Key Policies**. You use these settings to:

    - Specify that computers automatically submit a certificate request to an enterprise certification authority and install the issued certificate.

    - Create and distribute a certificate trust list.

    - Establish common trusted root certification authorities.

    - Add encrypted data recovery agents and change the encrypted data recovery policy settings.

- **IP Security Policies on Active Directory**. IP Security (IPSec) policy can be applied to the GPO of an Active Directory object. This propagates that IPSec policy to any computer accounts affected by that GPO.

- **Wireless Networking**. This lets you configure wireless network settings that are part of Group Policy for Computer Configuration. Wireless network settings include the list of preferred networks, WEP settings, and IEEE 802.1X settings. These settings are downloaded to targeted domain members, making it much easier to deploy a specific configuration for secure wireless connections to wireless client computers.

- **Software Restriction Policies**. This lets you protect your computer environment from untrusted code by identifying and specifying which applications are allowed to run. With software restriction policies, you can:

    - Control the ability of programs to run on your system. For example, if you are concerned about users receiving viruses through e-mail, you can apply a policy setting that does not allow certain file types to run in the e-mail attachment directory of your e-mail program.

    - Permit users to run only specific files on multi-user computers. For example, if you have multiple users on your computers, you can set up software restriction policy settings in such a way that users do not have access to any software but those specific files that are necessary for their work.

    - Decide who can add trusted publishers to your computer.

    - Control whether software restriction policy settings affect all users or just certain users on a computer.

    - Prevent any files from running on your local computer, organizational unit, site, or domain. For example, if your system has a known virus, you can use software restriction policy settings to stop a computer from opening the file that contains the virus.

    | **Note** |
    | --- |
    | Software restriction policy settings should not be used as a replacement for antivirus software. |

- You can configure security settings policies in **Computer Configuration\Windows Settings\Security Settings** in the Group Policy Object Editor.

**Default Security Templates**

Windows 2000 and Windows Server 2003 include three default security templates called *Basic*. These new default security settings are applied to Windows 2000 or Windows Server 2003 systems that have been installed onto an NTFS partition. When Windows 2000 or Windows Server 2003 is installed onto a FAT file system, security cannot be applied.

The following Basic security templates are used:

- Basicwk.inf for workstations.

- Basicsv.inf for member servers.

- Basicdc.inf for domain controllers.

The Basic security templates specify default Windows 2000 or Windows Server 2003 security settings for all security areas, with the exception of User Rights and Groups. These templates can be applied to Windows 2000 or Windows Server 2003 systems using the Security Configuration and Analysis MMC snap-in or by using the **Secedit.exe** command-line tool.

**Incremental Security Templates**

Windows 2000 and Windows Server 2003 include several incremental security templates. By default, these templates are stored in %systemroot%\Security\Templates. These predefined templates can be customized using the Security Templates MMC snap-in and can be imported into the Security Settings extension of the Group Policy Object Editor.

These security templates were constructed based on the assumption that they would be applied to computers running Windows 2000 or later and that are configured with Windows 2000 or later default security settings. In other words, these templates *incrementally* modify the default security settings. They do not include the default security settings plus the modifications.

The following table lists the incremental security templates included in Windows 2000 and Windows Server 2003.

| **Security Configuration** | **Computer** | **Templates** | **Description** |
|---|---|---|---|
| Compatible | Workstation, and server | Compatws.inf | For customers who do not want their users to run as Power Users (by default all users are Power Users on Windows 2000 Professional and Windows XP Professional), the Compatible configuration opens up the default permissions for the Users group so that legacy applications are more likely to run. For example, Office 97 should run successfully when users are logged on as a User to a computer running Windows 2000, Windows XP, or Windows Server 2003 that has had the Compatible security template applied over the default settings. Note that this is not considered a secure environment. |
| Secure | Workstation, server, and domain controller | Securews.inf and Securedc.inf | The Secure configuration provides increased security for areas of the operating system that are not covered by permissions. This includes increased security settings for Account Policy, Auditing, and some well-known security-relevant registry keys. Access control lists are not modified by the secure configurations because the secure configurations assume that default Windows 2000, Windows XP, or Windows Server 2003 security settings are in effect. |
| Highly Secure | Workstation, server, and domain controller | Hisecws.inf and Hisecdc.inf | The Highly Secure configuration is provided for Windows 2000, Windows XP, or Windows Server 2003-based computers that operate in native (or pure) Windows 2000 or Windows Server 2003 environments only. In this configuration, it is required that all network communications be digitally signed and encrypted at a level that can only be provided by Windows 2000 or later. Thus, a Windows 2000 highly secure computer cannot communicate with a client running Windows 95, Windows 98, or Windows NT. |

For information about the default security settings contained in the Default Domain Policy GPO and Default Domain Controller Policy GPO, see "Appendix A: Security Settings and User Rights" later in this paper.

**Using Software Installation and Maintenance**

You can use the software installation and maintenance feature to install software applications at computer startup, user logon, or on demand. You can also use this feature to upgrade deployed applications, remove earlier applications that are no longer required, and deploy service packs and operating system upgrades. They can ensure that a person cannot install any software from local media, such as a CD-ROM or disk.

This feature also provides for the following situations:

- If users inadvertently delete files from an application it will repair itself.
- If users move from one computer to another their software will always be available to them.
- If users do not have an application installed on their computer and they try to open a document associated with that application, the application will automatically be installed and the document will open.

Deploying software through Group Policy requires applications to use the Windows Installer service, which provides much more than just the capability to install applications. It also protects the integrity of the application against inadvertent mishaps with local files. For example, if a user attempted to use a copy of Microsoft Word that was missing some essential files, the Windows Installer service would reinstall the files from the install point, the next time that the application is launched.

In addition, Windows Installer-based applications that are deployed using Group Policy can install with elevated privileges, meaning users don't have to be administrators on their local machines to install software that you, as a network administrator, want them to have.

Application repair follows the same logic as on-demand installation. Whenever an application authored by Windows Installer is invoked, the Windows Installer service checks to ensure that the appropriate files are available; if required, files or settings are

repaired automatically.

You use Group Policy to define software installation options that specify which applications are to be deployed, upgraded, or removed from a computer. You can apply software installation policies to groups of users or to groups of computers, depending on your organization's needs. There are two methods by which you can install applications on users' computers: assigning and publishing.

**Assigning Applications**
You can assign applications to either a user or a computer using Group Policy. When you assign applications to a computer, the application is automatically installed the next time the computer is started. When you assign applications to a user with Group Policy, the administrator can choose to either have the application installed on-demand when the user selects the application or in-full when the user next logs on:

- **On Demand.** If the application is installed on demand, the user's computer is set up with a Start menu shortcut, and the appropriate file associations are created in the registry. To the user, it looks and feels as if the application is already present. However, the application is not fully installed until the user needs the application. When the user attempts to open the application or a file associated with that application, Windows Installer checks to make sure that all the files and parameters of the application are present for the application to properly execute. If they are not present, Windows Installer retrieves and installs them from a predetermined distribution point. Once in place, the application opens.

- **Full Install.** The full-install option is useful for specific groups of users such as frequent travelers who might require all available applications to be fully installed before they travel. With full install, a user's applications are installed at logon.

Assigning applications makes them resilient — they are available no matter what the user does; for example, if the user removes an application, it will automatically be reinstalled on demand.

**Publishing Applications**
When you publish an application, it appears in Add or Remove Programsin Control Panel. Users can choose to install published applications. Installation can also be configured to occur automatically when a user attempts to open a file that requires a specific published application. You publish applications when the software is not absolutely necessary for users to perform their jobs.

In order to obtain the full benefits of publishing technology, all published applications should be authored to install using the Windows Installer service. Although you can still publish non-Windows Installer service applications using .ZAP files, you won't get the benefits of elevated privileges as explained earlier, and of course, you won't get the benefits of using Windows Installer either.

A .zap file is a text file that provides a pointer to the setup package, which enables the application to be listed in **Add or Remove Programs**.

For more information, see the "Software Installation and Maintenance" white paper at http://go.microsoft.com/fwlink/?linkid=13293 and the Step-by-Step Guide to Software Installation and Maintenance at http://www.microsoft.com/windows2000/techinfo/planning/management/swinstall.asp.

**Scripts**
With the Scripts extensions, you can assign scripts to run when the computer starts or shuts down or when users log on or off their computers. For this purpose, you can use Windows Script Host to include both Visual Basic® Scripting Edition (VBScript) and Jscript® development software script types.

Windows 2000 and later include Windows Script Host, a language-independent scripting host for 32-bit Windows platforms. For more information about Windows Script Host, see the Microsoft Windows Script Web site at http://www.microsoft.com/scripting.

The names of scripts and their command lines (in the form of registry keys and values) are stored in the Registry.pol file, described in "Registry.pol Files" later in this document.

**Types of Scripts**
The five script types are as follows:

- Group Policy logon scripts.

- Group Policy logoff scripts.

- Group Policy startup scripts.

- Group Policy shutdown scripts.

- Legacy logon scripts (those specified on the User object). This includes support for Windows Script Host scripts. Windows Script Host supports scripts written in VBScript or JavaScript. This means that you can now enter a command line like *sample.vbs* in the logon script path of the user object.

| **Note** |
| --- |
| Windows XP Professional, Windows 2000, and the Windows 98-based clients will properly run .vbs and .js scripts. To run .vbs and .js scripts on Windows NT 4.0 and Windows 95 clients, you must embed the scripts in batch (.bat) files. The scripts continue to run in a normal window. There is a policy that allows for scripts to be run as hidden or minimized. |

**Specifying Policy Settings for Script Behavior**

The following table lists the Group Policy options that are available to control the behavior of scripts.

| **Policy in Computer Configuration\Administrative Templates\System\Logon** | **Description** |
| --- | --- |
| Run logon scripts synchronously | When this option is enabled, the system waits until the script finishes running before it starts Windows Explorer. Note that an equivalent option for this is available under the User Configuration node. The policy setting you specify in the Computer Configuration node has precedence over that set in the User Configuration node. |
| Run startup scripts asynchronously | By default, startup scripts run synchronously and hidden, which means the user cannot logon until the scripts complete. In some corporations, the administrator might want the scripts to run asynchronously since they could take a long time to complete. This policy allows the administrator to change the default behavior. |
| Run startup scripts visible | If this option is enabled, startup scripts run in a command window. |
| Run shutdown scripts visible | If this option is enabled, shutdown scripts run in a command window. |
| Maximum wait time for Group Policy scripts | This policy setting lets you change the default script time out period. (By default, scripts will time-out after 600 seconds). The range is 0 to 32000 seconds. |

| **Policy in User Configuration\Administrative Templates\System\Logon\Logoff** | **Description** |
| --- | --- |
| Run logon scripts synchronously | When you enable this option, Windows waits for the scripts to finish running before it starts Windows Explorer.<br><br>Note that an equivalent option for this is available under the Computer Configuration node. The policy setting you specify in the Computer Configuration node has precedence over that set in the User Configuration node. |
| Run legacy logon scripts hidden | If this option is enabled, legacy logon scripts will run in hidden mode. |
| Run logon scripts visible | If this option is enabled, logon scripts run in a command window. |
| Run logoff scripts visible | If this option is enabled, logoff scripts run in a command window. |

| **Note** |
| --- |
| Scripts that run hidden (and to a lesser degree minimized) can cause an errant script or one that prompts for user input to wait for 600 seconds. This is the default wait-time value and may be changed using a Group Policy. During this time, the system appears to be hung up. In the case of a script running in a minimized window, if the user selects the window, its processing can be stopped. |

**Best Practice**: For easier manageability, it is a good idea to use Group Policy scripts and to avoid using legacy logon scripts, if at all possible. Rather than using a single monolithic script with lots of internal logic branching, Group Policy-based logon scripts allow for use of tiered and modular scripts targeted to the desired set of users.

**Folder Redirection**
The Folder Redirection extension is used to redirect any of the following special folders in a user profile to an alternate location (such as a network file share):

- Application Data
- Desktop
- My Documents
    - My Pictures
- Start Menu

For example, you could redirect a user's My Documents directory to \\Server\Share\%username%. By redirecting the My Documents directory, you can provide the following advantages:

- Ensure that users' documents are available when they roam from one computer to another.
- Reduce the time it takes to log on to and log off from the network. The My Documents folder is part of the Roaming User Profile (RUP), which means that the My Documents folder and its contents are copied back and forth between the client computer and the server when users log on and log off. Relocating the My Documents folder outside of the user profile can significantly decrease that time.
- Store user data on the network (rather than on the local computer). The data can then be managed and protected by the Information Technology department.
- Make users' network-based My Documents folder available to users when they are disconnected from the corporate network by using Offline Folder technologies.

**Folder Redirection Improvements for Windows XP and Windows Server 2003**
This section provides information about the differences between Windows 2000 and Windows Server 2003.

**User Interface changes**
The Folder Redirection user interface has been simplified for Windows Server 2003. The main goals of these changes were to simplify the use of Folder Redirection by removing the requirement that administrators be familiar with environment variables such as %USERNAME%.

In addition to the simplified UI, several new redirection options have been added:

- **Create a directory for each user under the root path**. Rather than having to enter a UNC path such as \\server\share\%username%\MyDocuments, the administrator can simply type in the path to the file share such as \\server\share, and Folder Redirection will automatically append the user name and the directory name when the policy is applied. This removes the need for administrators to be familiar with environment variables, and minimizes the chances of errors and spelling mistakes.
- **Redirect to home directory (My Documents Only)**. Windows Server 2003 and Windows XP allow you to redirect a user's My Documents folder to their home directory. This option is intended only for organizations that have a legacy deployment of home directories and want to transition users to the My Documents metaphor while maintaining compatibility with their existing home directory environment. **You should only select this option if you have already deployed home directories in your organization.**
- Folder redirection treats redirection to the home directory as a special case and certain checks are skipped:
    - Redirection to the home directory is only supported on Windows XP and Windows Server 2003 computers.

Redirection to the home directory policy will fail to apply on Windows 2000 computers.

- No security check is performed and no setting of access permissions is done. The administrator must set restrictive permissions on the directory to ensure that proper access is granted. The **Grant the user exclusive rights to My Documents** check box on the settings page of the property sheet is disabled.

- No ownership checks are made. Normally folder redirection will fail if a user is not the owner of the directory they are being redirected to. Because redirection to the home directory is intended for use in a legacy environment, this ownership check is skipped.

- Users must have the home folder property correctly set on their user object. The folder redirection client side extension retrieves the actual path for the user's home directory from the user object at logon time. Users affected by Folder Redirection Policy must have this path correctly set or folder redirection will not apply.

- **Redirect to a specific path**. This option is intended to allow an administrator to redirect folders to an alternate local drive/partition, or to enter unusual configurations not anticipated by the new user interface. Functionally it works in exactly the same way as the Windows 2000 folder redirection user interface.

- **Redirect to the local user profile**. This option is intended to allow an administrator to redirect the selected folder to the default location in the local user profile, for example: %userprofile%\<Folder Name>. This setting can be used to remove redirection for a particular folder, without removing the GPO. Note: Setting the redirection option to "Not Configured" (or "No Administrative policy specified" in the Windows 2000 UI) **does not** redirect the folder to the local profile, this option means that Folder Redirection is not configured – *if a folder was previously redirected it will continue to be redirected to the previous location*. If an administrator wanted to return the folder to the local user profile they can use this redirection setting.

**My Pictures no longer shown in the Folder Redirection Node**
To simplify the user interface and to help support the best practice that the My Pictures folder should always follow the My Documents folder, the My Pictures folder is not shown in the Folder Redirection node for new GPO's. If you have previously redirected the My Pictures folder separately, the My Pictures node will still appear.

**Redirected Folders automatically made available offline**
By default in Windows XP and Windows Server 2003, any redirected shell folders such as My Documents, Desktop, Start Menu, and Application Data are automatically made available offline. This is in contrast to Windows 2000, which required administrators to configure the "Administratively assigned offline files" policy setting to ensure all files in the redirected folders were always available offline. This setting was difficult to use with advanced folder redirection, and involved extra administrative overhead.

The default behavior can be overridden by enabling the **Do not automatically make redirected folders available offline** policy. This setting can be found in the Group Policy Object Editor in the **User Configuration\Administrative Templates\Network\Offline Files** section.

Note that on Windows Server 2003 Offline files are disabled by default.

More information about Folder Redirection will be available in the white paper, "User Data and Settings Management" at http://www.microsoft.com/grouppolicy.

**Internet Explorer Maintenance**
The Internet Explorer Maintenance extension snap-in includes policy settings to manage the following:

- **Browser User Interface**. You use these options to customize the browser's appearance. For example, you can specify settings for the browser title bar, toolbar button options, and so on.

- **Connection Settings**. You can preset and manage the connection settings, such as local area network (LAN) and dial-up options.

- **Custom URLs**. You can specify which URLs are displayed by the browser, for example, for the Home page, those on the Favorites list, and for the Search page.

- **Security**. You can preset security settings such as security zones, content ratings, and Authenticode. (A browser can be

configured to allow only signed code to be downloaded. Authenticode is the Microsoft version of object signing; it provides a basis for verifying the origin and integrity of an object, as well as links to policies of a certification authority).

- **Program Associations**. You can specify which Internet programs to use by default for Internet-related tasks such as reading e-mail or viewing newsgroups.

**Exporting Internet Explorer Settings for Earlier Clients**

Administrators can export Internet Explorer policy settings into an auto-configuration package (an .ins file and its associated .cab files) to be used to apply these settings to Windows 95, Windows 98, and Windows NT 4.0 clients. The exported packages are auto-configuration packages. Before the original Group Policy Object Editor was created in Windows 2000, Internet Explorer settings were applied to Internet Explorer clients using auto-configuration packages after Internet Explorer installation. Using GPOs is the preferred method of applying Internet Explorer policy settings on clients running Windows 2000 or later, although Windows 2000 and Windows Server 2003 support auto-configuration packages.

**Managing Internet Explorer Maintenance Advanced Settings**

You can manage advanced settings for Internet Explorer such as setting a size limit for users' Temporary Internet files. In order to do this, you need to first enable **Preference Mode** for Internet Explorer Maintenance. By default, the **Preference Mode** option is hidden. You access this option by right-clicking **Internet Explorer Maintenance** node and selecting **Preference Mode** on the shortcut menu.

This adds an **Advanced** node to the results pane. This node contains settings for managing Temporary Internet files and other UI features. Note that switching to **Preference Mode** disables some of the Internet Explorer Maintenance nodes. If a setting name has **Preference Mode** appended to it, it can be used in that mode; otherwise, it means that setting is disabled. For example, the **Connection Settings (Preference Mode)** option under the **Connection** node can be used in **Preference Mode** as indicated by its labeling in the UI, whereas the **User Agent String** option (note the exclusion of **Preference Mode**) cannot be used in **Preference Mode** and this is reflected in its labeling.

For more information, see this Microsoft Knowledge Base article, How to Set Advanced Settings In Internet Explorer by Using Group Policy Objects.

**Using Internet Explorer Customization Wizard and Internet Explorer Profile Manager**

Besides the Internet Explorer Maintenance Group Policy options mentioned above, it is also possible to customize Internet Explorer before deployment and to manage Internet Explorer on other operating systems by using the Internet Explorer Administration Kit (IEAK) at http://www.microsoft.com/windows/ieak/default.asp.These tools provide options for System Policies and restrictions that administrators can use to specify desktop, shell, and security settings, for example.

**Remote Installation Services**

Remote Installation Services is an optional component that is included in the Windows Server operating system and works with other Windows Server 2003 technologies to implement the Remote Operating System Installation feature. Administrators use Remote Operating System Installation to remotely install a copy of the Windows XP Professional operating system on supported computers. (Computers that are PC98-compliant ship with a PXE Remote Boot ROM.) Administrators use the Remote Installation Services extension of Group Policy to specify which options are presented to users by the Client Installation Wizard, for example, Automatic Setup, Custom Setup, and Restart Setup.

Client computers that are enabled with Pre-boot Execution Environment (PXE) remote-boot technology access the RIS server to install the operating system, and then the Remote Installation Services server checks for Group Policy that affects remote installation options defined for the user. The Boot Information Negotiation Layer (BINL) service running on the RIS server performs this work. It impersonates the user who logs on to the RIS client-side pre-boot user interface, and evaluates the GPOs to determine the resulting policy. Based on the resulting policy, it determines which screens to send to the pre-boot RIS client code for display to the user.

# Group Policy Modeling and Results

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy Modeling and Group Policy Results is a feature of Group Policy that makes implementation, troubleshooting, and planning of Group Policy easier. When multiple GPOs apply to a given user or computer, they can contain conflicting policy settings. For most policy settings, the final value of the policy setting is set only by the highest precedence GPO that contains that setting. Group Policy Modeling and Group Policy Results uses the Resultant Set of Policy (RSoP) infrastructure, available on Windows XP and Windows Server 2003, to present the final set of policy that is applied as well as settings that did not apply as a result of policy inheritance.

Specifically, RSoP helps you determine the following:

- The final value of the setting that is applied as a result of all the GPOs.
- The final GPO that set the value of this setting (also known as the winning GPO).

Precedence details that show any other GPOs that attempted to set this setting and the value that each GPO attempted to set for that policy setting.

**Group Policy Results**
This represents the actual policy data that is applied to a given computer and user. It is obtained by querying the target computer and retrieving the RSoP data that was applied to that computer. The Group Policy Results capability is provided by the client operating system and requires Windows XP, Windows Server 2003 or later. Outside of GPMC, Group Policy Results is referred to as RSoP - logging mode.

**Group Policy Modeling**
This is a simulation of what would happen under circumstances specified by an administrator. Group Policy Modeling requires that you have at least one domain controller running Windows Server 2003 because this simulation is performed by a service running on a domain controller that is running Windows Server 2003. With Group Policy Modeling, you can either simulate the RSoP data that would be applied for an existing configuration, or you can perform "what-if" analyses by simulating hypothetical changes to your directory environment and then calculating the RSoP for that hypothetical configuration. For example, you can simulate changes to security group membership, or changes to the location of the user or computer object in Active Directory. Outside of GPMC, Group Policy Modeling is referred to as RSoP - planning mode. Note that although Windows 2000 does not provide the RSoP infrastructure, Group Policy Modeling can be used as an effective way to simulate the affect of Group Policy on Windows 2000 computers.

**Using GPMC Reports**
In GPMC, resultant set of policy data is obtained using Group Policy Modeling or Group Policy Results wizards. GPMC provides an HTML report of the RSoP data. This report shows the final value of the winning settings and the winning GPO that set that value. When you create a Group Policy Modeling or Group Policy Results report, the report is shown in GPMC under the appropriate node. Right-clicking this report and choosing **Advanced View** opens the RSoP snap-in, which provides additional information, enabling you to verify precedence for a policy setting. In the RSoP snap-in, the dialog box for a policy setting contains a Precedence tab, which shows all GPOs that attempted to set a particular setting and the value for each GPO.

**RSoP Architecture**
Figure 2 below shows the high-level architecture of RSoP for Group Policy Results and Group Policy Modeling.




**Figure 2. RSoP high-level architecture**

Take for example, a standard logon procedure in Windows 2000: a client computer logs onto the network and Winlogon runs. The domain controller passes a list of pointers to the GPOs that are to apply. This list is passed to each of the client-side extensions (CSEs) such as Software Installation, Scripts, Security, Administrative Templates, and so on. Each CSE processes this list of GPOs.

Windows Server 2003 uses the same process but improves on Windows 2000 by collecting all the Group Policy processing information and storing it in a Common Information Model Object Management (CIMOM) database on the local computer. This information, such as the list, content and logging of processing details for each GPO, can then be accessed by tools using WMI.

In Group Policy Results, RSoP queries the CIMOM database on the target computer, receives information about the policies and displays it in GPMC. In Group Policy Modeling, RSoP simulates the application of policy using the Group Policy Directory Access Service (GPDAS) on a Domain Controller. GPDAS simulates the application of GPOs and passes them to virtual client-side extensions on the Domain Controller. The results of this simulation are stored to a local CIMOM database on the domain controller before the information is passed back and displayed in GPMC.

**Security and RSoP**
By default, access to Group Policy Results is restricted to enterprise, domain, and local administrators although users can still perform logging on their own computer. In Windows XP, non-administrators can run Group Policy Results for their computer and user account; in Windows Server 2003 non-administrators can only run Group Policy Results for their own user account. Group Policy Modeling is restricted to enterprise and Domain Admins. However, organizations can delegate access to Group Policy Results and Modeling using GPMC. For step-by-step instructions, see Group Policy Help.

**Group Policy Results and Modeling Examples**
Group Policy Results and Modeling allow administrators to solve problems for specific scenarios. Some examples are included below.

**Group Policy Results**
**What is the current state of Folder Redirection for the current user?**
Example: User Paul has four computers, and contacts help desk because they cannot find files on computer D that are on computer A, even though Paul is set up to use folder redirection. The administrator runs Group Policy Results and sees that on computer D, redirection is different than the others because a different GPO applies.

**What is the current state of Folder Redirection for a sampling of users?**
Example: An administrator wants to profile different sets of users. Using Group Policy Results, the administrator picks a sample user from each user group and uses the RSoP information to model redirection within the organization.

**Why did this happen?**
Example: An administrator is confused as to why Paul's documents are being redirected to the SuperUsers server. The administrator uses Group Policy Results to look at the current redirection path, and the GPO and security group that caused the redirection.

The administrator notices that of the three GPOs specifying folder redirection policy, the winning GPO has the advanced option set to redirect users in different security groups to different locations. The administrator notices that Paul is a member of both the VanillaUsers group and the SuperUsers group and realizes that this caused Paul's folders to be redirected to the SuperUsers server.

**Group Policy Modeling**
**Precedence Details**
In processing Group Policy, administrators determine which GPOs were in conflict to configure folder redirection for this user.

**Change of Site**
Example: An administrator can model site and domain changes for individual users or using a sample target to assess what would happen to an entire group of users under different combinations of sites, domains, and so forth. By comparing what should be seen, by what actually exists under the new GPO structure, the administrator can avoid problems before the move actually takes place.

Example: User Jane is going to move from one department to another. The administrator uses Group Policy Modeling to model the move under the different site condition and finds out that a GPO conflict exists that redirects Jane's folder to an alternate location.

**Change of Folder Redirection Mode**
An Administrator wants to configure folder redirection to use the advanced options to redirect users to alternate locations based on their security group membership. The administrator uses Group Policy Modeling to configure Group Policy for the desired folder redirection behavior.

By comparing the current results of folder redirection for the users, with the results of the desired changes, the administrator can avoid problems before the move actually takes place.

**RSoP Schema**

For information about the RSoP schema, see the RSoP SDK, available as part of the Microsoft Windows Platform SDK at http://www.microsoft.com/downloads/details.aspx?FamilyId=A55B6B43-E24F-4EA3-A93E-40C0EC4F68E5&displaylang=en.

# Group Policy Processing

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As described earlier in this paper, Group Policy is processed in the following order: Local Group Policy object (Local GPO), then GPOs linked to containers in this order: site, domain, and organizational units, including any nested organizational units (starting with the organizational unit further from the user or computer object). This means that the local Group Policy object is processed first, and the organizational unit to which the computer or user belongs (the one that it is a direct member of) is processed last. All of this is subject to the following conditions:

- WMI or security filtering that has been applied to GPOs.
- Any domain-based GPO (not local GPO) may be enforced by using the **Enforce** option so that its policies cannot be overwritten. When more than one GPO has been marked as enforced, the GPO that is highest in Active Directory hierarchy takes precedence.
- At any domain or organizational unit, Group Policy inheritance may be selectively designated as **Block Inheritance**. However, blocking inheritance does not prevent policy from **enforced** GPOs from applying; this is because enforced GPOs are always applied, and cannot be blocked.

| **Note** |
| --- |
| Every computer has a single local GPO that is always processed regardless of whether the computer is part of a domain or is a stand-alone computer. The Local GPO can't be blocked by domain-based GPOs. However, settings in domain GPOs always take precedence since they are processed after the Local GPO. |

**Initial Processing of Group Policy**

Group Policy for computers is applied at computer startup. For users, Group Policy is applied when they log on. By default, the processing of Group Policy is synchronous, which means that computer Group Policy is completed before the CTRL+ALT+DEL dialog box is presented, and user Group Policy is completed before the shell is active and available for the user to interact with it. (As explained below, Windows XP with Fast Logon enabled lets users logon while Group Policy is processed in the background.)

**Synchronous and Asynchronous Processing**

Synchronous processes can be described as a series of processes where one process must finish running before the next one begins. Asynchronous processes, on the other hand, can run on different threads simultaneously because their outcome is independent of other processes.

You can change the default processing behavior by using a policy setting for each GPO so that processing is asynchronous instead of synchronous. However, this is not recommended because it can cause unpredictable or undesirable side effects. For example, if the policy has been set to remove the **Run** command from the **Start** menu, it is possible under asynchronous processing that a user could logon prior to this policy taking effect, so the user would initially have access to this functionality. To provide the most reliable operation, it is recommended that you leave the processing as synchronous.

**Fast Logon in Windows XP Professional**

By default in Windows XP Professional, the Fast Logon Optimization feature is set for both domain and workgroup members. This results in the asynchronous application of policies when the computer starts and when the user logs on. This application of policies is similar to a background refresh process and can reduce the length of time it takes for the Logon dialog box to display and the length of time it takes for the shell to be available to the user. An administrator can change the default by using the Group Policy Object Editor.

Fast Logon Optimization is always off during logon under the following conditions:

- When a user first logs on to a computer.
- When a user has a roaming user profile or a home directory for logon purposes.
- When a user has synchronous logon scripts.

Note that under the preceding conditions, computer startup can still be asynchronous. However, because logon is synchronous under these conditions, logon does not exhibit optimization.

The following table summarizes the default processing of policy on Windows XP.

| Client | Application at startup/log on | Application at refresh |
|---|---|---|
| Windows 2000 | Synchronous | Asynchronous |
| Windows XP Professional | Asynchronous | Asynchronous |

Windows XP clients support Fast Logon Optimization in any domain environment. To turn off Fast Logon Optimization, you can use the following policy setting:

Computer Configuration\Administrative Templates\System\Logon\ Always wait for the network at computer startup and logon

| Note |
|---|
| Fast Logon Optimization is not a feature of Windows Server 2003. |

**Folder Redirection and Software Installation Policies**

Note that when logon optimization is on, a user may need to log on to a computer twice before folder redirection policies and software installation policies are applied. This is because application of these types of policies require the synchronous policy application. During a policy refresh (which is asynchronous), the system sets a flag that indicates that the application of folder redirection or a software installation policy is required. The flag forces synchronous application of the policy at the user's next logon.

**Time Limit for Processing of Group Policy**

Under synchronous processing, there is a time limit of 60 minutes for all of Group Policy to finish processing on the client. Any client-side extensions that are not finished after 60 minutes are signaled to stop, in which case the associated policy settings may not be fully applied. An errant extension may not be able to respond; in either case the Group Policy engine goes into asynchronous processing mode. This means that the Group Policy engine is no longer blocked while waiting for a running (likely errant) extension and continues to process; it leaves the extension(s) running and does not terminate it (them). There is no setting to control this time-out period or behavior.

**Background Refresh of Group Policy**

In addition to the initial processing of Group Policy at startup and logon, Group Policy is applied subsequently in the background on a periodic basis, and can also be triggered on demand from the command line.

During a background refresh, a client side extension will by default only reapply the settings if it detects that a change was made on the server in any of its GPOs or its list of GPOs. This is done for performance reasons.

Not all Group Policy extensions are processed during a background refresh. Software Installation and Folder Redirection processing occurs only during computer startup or when the user logs on. This is because processing periodically could cause undesirable results. For example, for Software Installation, if an application is no longer assigned, it is removed. If a user is using the application while Group Policy tries to uninstall it or if an assigned application upgrade takes place while someone is using it, errors would occur.

| Note |
|---|
| The script's extension is processed during background refresh, however the scripts themselves are only ran at startup, shutdown, logon, and logoff, as appropriate. |

**Periodic Refresh Processing**

Group Policy is processed periodically. By default, this is done every 90 minutes with a randomized offset of up to 30 minutes. You can change these default values by using a Group Policy setting in Administrative Templates. Setting the value to zero minutes causes the refresh rate to be set to seven seconds.

| Note |
|---|

Setting a short refresh interval in a production environment is not recommended. This is because a policy refresh causes the Windows shell to be refreshed, which in turn causes all open shortcut menus to close, a brief flicker of the screen, and so on. In addition, it causes computers to contact domain controllers more frequently, increasing the load on the domain controllers. However, setting a shorter interval may be useful in test or demonstration scenarios.

To change the policy refresh interval setting, edit the **Default Domain Controllers** Group Policy object, which is linked to the **Domain Controllers** organizational unit. The **Group Policy Refresh Interval for Computers** setting is located under **Computer Configuration\Administrative Templates\System\Group Policy** node.

For domain controllers, the default period is every five minutes. **Group Policy Refresh Interval for Domain Controllers** setting is available under **Computer Configuration\Administrative Templates\System\Group Policy** node.

**On-Demand Processing**
You can also trigger a background refresh of Group Policy on demand from the client. However, the application of Group Policy cannot be pushed to clients on demand from the server.

**Messages and Events**
When Group Policy is applied, a **WM_SETTINGCHANGE** message is sent, and an event is signaled. Applications that can receive window messages can use it to respond to a Group Policy change. Those applications that do not have a window to receive the message (as with most services) can wait for the event.

**Refreshing Policy from the Command Line**
Gpupdate refreshes local Group Policy settings and Group Policy settings that are stored in Active Directory, including security settings. This command supersedes the now obsolete **/refreshpolicy** option for the **secedit** command.

**Syntax**
Gpupdate [/target:{computer | user}] [/force] [/wait:*Value*] [/logoff] [/boot]

**Parameters**
/target:{computer | user}

Processes only the *Computer* settings or the current *User* settings. By default, both the computer settings and the user settings are processed.

/force

Ignores all processing optimizations and reapplies all settings.

**/wait:** *Value*

Number of seconds that policy processing waits to finish. The default is 600 seconds. 0 equals no wait, and -1 equals wait indefinitely.

/logoff

Logs off after the refresh has completed. This is required for those Group Policy client-side extensions that do not process on a background refresh cycle but that do process when the user logs on, such as user Group Policy Software Installation and Folder Redirection. This option has no effect if there are no extensions called that require the user to log off.

/boot

Restarts the computer after the refresh has completed. This is required for those Group Policy client-side extensions that do not process on a background refresh cycle but that do process when the computer starts up, such as computer Group Policy Software Installation. This option has no effect if there are no extensions called that require the computer to be restarted.

/synch

Causes the next foreground policy application to be done synchronously. Foreground policy applications occur at computer boot and user logon. You can specify this for the user, computer, or both using the /Target parameter. The /Force and /Wait parameters will be ignored if specified.

/?

Displays help at the command prompt.

**Slow Links and Remote Access Issues**
Special considerations apply when processing Group Policy over slow links or remote access.

| Note |
| --- |
| Note that while these issues are related, they are distinct, and the processing of Group Policy is different for each. In particular, *remote access* does not necessarily imply a slow link, nor does a LAN necessarily imply a fast link. A slow link is by default based on the algorithm described in the section below. Windows Server *remote access* is part of the integrated Routing and Remote Access Service; it connects remote or mobile users to corporate networks, allowing users to work as if their computers are physically connected to the network. Users run remote access software to connect to a remote access server, which is a computer running Windows Server and the Routing and Remote Access Service. The remote access server authenticates the user and services sessions until terminated by the user or network administrator. The remote access connection enables all services typically available to a LAN-connected client, such as file and print sharing, messaging, and Web server access. |

**Group Policy and Slow Links**

When Group Policy detects a slow link, it sets a flag to indicate to client-side extensions that a policy setting is being applied across a slow link. Individual client-side extensions can determine whether or not to apply a policy setting over the slow link.

The default settings are as follows:

- Security Settings—ON (and cannot be turned off).
- Administrative Templates—ON (and cannot be turned off).
- Software Installation—OFF.
- Scripts—OFF.
- Folder Redirection—OFF.

For all but the Administrative Templates snap-in and security settings snap-in, a policy is provided for switching the slow link processing settings.

**Setting Policy for Slow-Link Definition**

You can use Group Policy to set the definition of a slow link for computers and users, and for user profiles.

For Group Policy, Windows 2000 and Windows Server 2003 use an IP ping algorithm to ping the server, rather than measuring the file system performance method that was used in Windows NT 4.0. Note: Slow link detection requires the Internet Control Message Protocol (ICMP). If ICMP cannot be used to communicate with the domain controllers, policy processing will not work, in which case you should disable slow link detection.

A slow link is, by default, based on the following algorithm (where ms = milliseconds):

1. Ping the server with 0 bytes of data and time the number of milliseconds. This value is time#1. If it is less than 10 ms, exit (assume a fast link).
2. Ping the server with 2 KB of uncompressible data, and time the number of milliseconds. This value is time#2. The algorithm uses a compressed .jpg file for this.
3. DELTA = time#2 - time#1. This removes the overhead of session setup, with the result being equal to the time to move 2 KB of data.
4. Calculate Delta three times, adding to TOTAL each DELTA value.
5. TOTAL/3 = Average of DELTA, in milliseconds.
6. 2 * (2 KB) * (1000 millisec/sec) / DELTA Average millisec = X
7. X = (4000 KB/sec) / DELTA Average
8. Z Kilobits per second (Kbps) = (4000 KB/sec) / DELTA Average) *(8 bits/byte)

9. Z Kbps = 32000 kbps/Delta Avg.

Two KB of data have moved in each direction (this is represented by the leading factor two on the left side in step six above) through each modem, Ethernet card, or other device in the loop once.

The resulting Z value is evaluated against the policy setting. A default of less than 500 Kbps is considered a slow link; otherwise it is a fast link. This value may be set through Group Policy in the Administrative Templates node.

To specify policy settings for Group Policy slow link detection for computers, you use the **Computer Configuration\Administrative Templates\System\Group Policy** node. To set this policy for users, you use the **User Configuration\Administrative Templates\System\Group Policy** node. The connection speed is set for kilobits per second (Kbps).

For User Profiles, the **Slow network connection time-out for user profiles** policy is located in the **Computer Configuration\Administrative Templates\System\Logon** node. This policy has support for both pinging the server and checking the performance of the file system. This is because user profiles can be stored anywhere, and that server may or may not have IP support. Therefore, the user profile code first tries to ping the server. If the server does not have IP support, it falls back to measuring the file system's performance. You must specify connection speeds in both kilobytes per second (Kbps) and milliseconds (ms) when setting this policy.

**Application of Group Policy During a Remote Access Connection**

Group Policy is applied during a remote access connection as follows:

When using the **Logon using dial-up connection** check box on the logon prompt, both User and Computer Group Policy is applied, provided the computer is a member of the domain that the remote access server belongs to or trusts. However, computer-based software installation settings are not processed. This is because normally computer policy would have been processed before the logon screen, but since no network connection is available until logon, the application of computer policy is done as background refresh at the time of logon.

When the logon is done with cached credentials, and then a remote access connection is established, Group Policy is not applied during logon. For example, if users connecting through a VPN connection are logging in via cached credentials, folder redirection settings will not be processed, because folder redirection policy can only be processed at user logon, not in the background refresh.

Group Policy is not applied to computers that are members of a foreign domain or a workgroup. Although the connection may still be made, access to domain resources may be affected (because of mismatched IPSec security).

**Client-side Processing of Group Policy**

The client-side extensions are loaded on an as-needed basis when a client computer is processing policy. The client computer first gets a list of Group Policy objects. Next, it loops through all the client-side extensions and determines whether each client-side extension has any data in any of the GPOs. If a client-side extension has data in a GPO, the client-side extension is called with the list of Group Policy objects that it should process. If the client-side extension does not have any settings in any of the GPOs, it is not called.

**Computer Policy for Client-Side Extensions**

A computer policy exists for each of the Group Policy client-side extensions (located in Computer Configuration\Administrative Templates\System\Group Policy). Each policy includes a maximum of three options (check boxes). Some of the client-side extensions include only two computer policy options; in those cases, this is because the third option is not appropriate for that extension.

The computer policy options are:

- **Allow processing across a slow network connection**. When a client-side extension registers itself with the operating system, it sets preferences in the registry, specifying whether it should be called when policy is being applied across a slow link. Some extensions move large amounts of data, so processing across a slow link can affect performance (for example, consider the time involved in installing a large application file across a 56 Kbps modem line). An administrator can set this policy to mandate that the client-side extension should run across a slow link, regardless of the amount of data.

- **Do not apply during periodic background processing**. Computer policy is applied at boot time, and then again in the background, approximately every 90 minutes thereafter. User policy is applied at user logon, and then approximately every 90 minutes after that. The **Do not apply during periodic background processing** option gives the administrator the ability to override this logic and force the extension to either run or not run in the background.

| Note |
| --- |
| The Software Installation and Folder Redirection extensions process policy only during the initial run because it is risky to process policy in the background. For example, with Software Installation application upgrades, applications are installed during the initial run and not in the background. If it were done in the background, a user could be running an application, and then have it uninstalled and a new version installed. The application could also have a shared component that is in use by another application. This would prevent the installation from completing successfully. |

- **Process even if the Group Policy Objects have not changed**. By default, if the GPOs on the server have not changed, it is not necessary to continually reapply them to the client, since the client should already have all the settings. However, local administrators may be able modify the parts of the registry where Group Policy settings are stored. In this case, it may make sense to reapply these settings during logon or during the periodic refresh cycle to get the computer back to the desired state.

  For example, assume that you have used Group Policy to define a specific set of security options for a file. Then the user (with administrative credentials) logs on and changes it. The Group Policy administrator may want to set the policy to process Group Policy even if the GPOs have not changed so that the security is reapplied at every boot. This also applies to applications. Group Policy installs an application, but the end user can remove the application or delete the icon. The process gives the administrator the ability to restore the application at the next user logon, even if the Group Policy objects have not changed option.

Note that, by default, security settings are applied every 16 hours (960 minutes) even if a GPO has not changed. It is possible to change this default period by using the following registry key:

```
HKLM\Software\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtentions\{82...}
\MaxNoGPOListChangesInterval, REG_DWORD, in number of minutes.
```

The following table lists the client-side extensions that include only two computer policy options, as well as the reason for this.

| Client-side extension | Missing policy check box | Reason |
| --- | --- | --- |
| Registry | Slow link (Allow processing across a slow network connection) | Registry policy is always applied because it controls the other client-side extensions. |
| Security Settings | Slow link (Allow processing across a slow network connection) | To ensure that security settings are in effect, they must always be applied, even across a slow link. |
| Folder Redirection | Background processing (Do not apply during periodic background processing) | Users' files could be in use while they are logged on. |
| Software Installation | Background processing (Do not apply during periodic background processing) | Users' software could be in use while they are logged on. |

**Policy Settings for Group Policy**

You can use administrative templates to configure how you use Group Policy. Policy settings are located in the following areas of the Group Policy Object Editor:

- **Computer Configuration\Administrative Templates\System\Group Policy**

- **User Configuration\Administrative Templates\System\Group Policy**

- For details on these policy settings, double-click the policy in the details pane, and then in the policy **Properties** dialog box, click the **Explain** tab.

# Group Policy Replication and Domain Controller Selection (Group Policy Infrastructure)

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a domain that contains more than one domain controller, Group Policy information takes time to propagate, or replicate, from one domain controller to another. Low bandwidth network connections between domain controllers slow replication. The Group Policy infrastructure has mechanisms to manage these issues.

Each GPO is stored partly in the Sysvol on the domain controller and partly in Active Directory. GPMC and Group Policy Object Editor present and manage the GPO as a single unit. For example, when you set permissions on a GPO in GPMC, GPMC is actually setting permissions on objects in both Active Directory and the Sysvol. It is not recommended that you manipulate these separate objects independently outside of GPMC and the Group Policy Object Editor. It is important to understand that these two separate components of a GPO rely on different replication mechanisms. The file system portion is replicated through File Replication Service (FRS), independently of the replication handled by Active Directory.

Lack of synchronization between the Group Policy template (data stored on Sysvol) and Group Policy container (data stored in Active Directory) portions of the Group Policy object can occur temporarily because of the differences in the replication schemes used by Active Directory and FRS.

For those Group Policy extensions that store data in only one data store (either Active Directory or Sysvol), this is not an issue, and Group Policy is applied as it can be read. Such extensions include Administrative Templates, Scripts, Folder Redirection, and most of the Security Settings.

For any Group Policy extension that stores data in both storage places (Active Directory and Sysvol), the extension must properly handle the possibility that the data is unsynchronized. This is also true for extensions that need multiple objects in a single store to be atomic in nature, since neither storage location handles transactions.

An example of an extension that stores data in Active Directory and Sysvol is Software Installation. The script files are stored on Sysvol and the Windows Installer package definition is in Active Directory. If the script exists, but the corresponding Active Directory components are not present, then nothing is done. If the script file is missing, but the package is known in Active Directory, application installation fails gracefully and will be retried on the next processing of Group Policy.

The tools used to manage Active Directory and Group Policy, such as GPMC, the Group Policy Object Editor, and Active Directory Users and Computers all communicate with domain controllers. If there are several domain controllers available, changes made to objects like users, computers, organizational units, and GPOs may take time to appear on other domain controllers. The administrator may see different data depending on the last domain controller on which changes were made and which domain controller they are currently viewing the data from.

For example, if you create a GPO on one domain controller and immediately attempt to link it on another domain controller, the operation could fail. In each domain, GPMC uses the same domain controller for all operations in that domain, in order to avoid any synchronization issues. This includes all operations on GPOs, organizational units, and security groups in that domain. In addition, when the Group Policy Object Editor is opened from GPMC, it will also use the same domain controller in use by GPMC. Finally, GPMC uses the same domain controller for all operations on sites within a given forest. This domain controller for sites is used to read and write information about the links to GPOs that exist on any given site; information regarding the GPO itself is obtained from the domain controller of the domain hosting the GPO. This domain controller is used to read and write information about the links to GPOs that exist on any given site; information regarding the GPO itself is obtained from the domain controller of the domain hosting the GPO.

By default, when you add a new domain to the console, GPMC uses the PDC emulator in that domain to help ensure that all administrators are using the same domain controller. For managing sites, GPMC uses the PDC emulator in the user's domain by default. You can change the default choice of domain controller using the Change Domain Controller dialog box in GPMC. If you are located at a remote site with a slow connection to the default domain controller, you may want to do this.

It is important for administrators to consider the choice of domain controller in order to avoid replication conflicts particularly because both Active Directory and FRS use multi-master replication. This is especially important to consider because GPO data resides in both Active Directory and on Sysvol, and two independent replication mechanisms must be used to replicate GPO data to the various domain controllers in the domain. If two administrators are simultaneously editing the same GPO on different domain controllers, it is possible for the changes written by one administrator to be overwritten by another administrator, depending on replication latency.

**Important**

If multiple administrators manage a common GPO, it is recommended that all administrators use the same domain controller when editing a particular GPO, to avoid collisions in FRS.

#### Options governing selection of a domain controller for GPMC

In GPMC, when you right-click a domain or the sites container and click **Change Domain Controller**, you see a **Change Domain Controller** dialog box. The domain controller options for GPMC are:

- **The one with the Operations Master token for the PDC emulator**. This is the default and preferred option.
- **Use any available domain controller.** This is the least safe option.
- **Use any available domain controller that is running Windows Server 2003 or later.** This option is useful if you are restoring deleted GPOs that contain software installation settings. If possible, it is recommended to perform restoration of GPOs containing software installation settings on domain controllers running Windows Server 2003.
- **This domain controller.** This option allows you to choose a specific domain controller from a list of domain controllers in the domain.

If you are changing the domain controller for a site, you can also choose any available trusted domain from the **Look in this domain** drop-down list box in the **Change Domain Controller** dialog box.

When you open the Group Policy Object Editor from GPMC it always uses the same domain controller that is targeted in GPMC for the domain where that GPO is located.

- All of these options may be overridden by a using policy setting, as described next. These settings are available in the **User Configuration\Administrative Templates\System\Group Policy** node of the Group Policy Object Editor.

#### Specifying a Domain Controller by Using Group Policy

Domain Admins can use a policy to specify how Group Policy chooses a domain controller—that is, they can specify which domain controller option should be used. In such cases, the option to choose a domain controller is unavailable since a policy is in place that overrides any setting that the user chooses. This policy allows Domain Admins to mandate that all administrators must use the PDC emulator, for example.

The **Group Policy domain controller selection** policy setting is available in the Administrative Templates node for User Configuration, in the **System\Group Policy** sub-container.

# Local Group Policy (Group Policy Infrastructure)

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can set local Group Policy for any computer, whether or not it participates in a domain. To set local Group Policy, you use the Group Policy Object Editor focused on the local computer. You can access the Group Policy Object Editor tool by typing **mmc** at the command prompt, adding the Group Policy Object Editor to MMC console, and focusing the Group Policy Object Editor on the local computer. Group Policy is processed in this order: local GPO first, followed by Active Directory linked GPOs (site, domain, organizational unit, and any nested organizational units).

**Local Group Policy Object**

On all computers, an Local GPO exists—this is just the Group Policy template portion. The location of the Local GPO is \%SystemRoot%System32\GroupPolicy. Each Group Policy extension snap-in queries the Group Policy engine to get the GPO type, and then decides if it should be displayed.

The following table indicates whether or not the Group Policy Object Editor extensions open when the Group Policy Object Editor is focused on the Local GPO.

| Group Policy Object Editor extension | Loaded when Group Policy Object Editor focused on Local GPO |
|---|---|
| Security Settings | Yes |
| Administrative Templates | Yes |
| Software Installation | No |
| Scripts | Yes |
| Folder Redirection | No |
| Internet Explorer Maintenance | Yes |

**Local Group Policy Object and DACLs**

There is no **Apply Group Policy** ACE for the local GPO; therefore, if you have Read access to the Local GPO, the local GPO applies to you. The implication is that it's difficult to have to choose whom the Local GPO should apply to (for example, the Local GPO also applies to the administrator). Everyone with Read access to the Local GPO who logs on gets the Local GPO. If this is not what you want, a work-around exists. You can set the Read ACE to Deny for a specific user, and then the Local GPO doesn't apply to that user. This is useful for administrators who don't want to be subject to the Local GPO settings. However, without Read access, administrators cannot see the contents of the Local GPO.

**Viewing Policy settings When the Group Policy Object Editor is Focused on the Local Computer**

When administrators run the Group Policy Object Editor focused on a local computer, this shows the information in the local GPO, not the cumulative effect of what has been applied to the computer or user. For Windows Server 2003, it shows the settings that a local administrator has set for that computer and all users of that computer. In the evaluation process, when the computer is joined to a domain, all the policy settings are subject to being overwritten by domain-based policy (any policy set in the site, domain, or organizational unit).

**Local Group Policy Object Processing**

When a computer is joined to a domain with Active Directory and Group Policy implemented, a local Group Policy object is processed. Note that Local GPO policy is processed even when the **Block Policy Inheritance** option has been specified.

Local Group Policy objects are always processed first, and then domain policy is processed. If a computer is participating in a domain and a conflict occurs between domain and local computer policy, domain policy prevails. However, if a computer is no longer participating in a domain, Local GPO policy is applied.

**Modifying the Local GPO on a Domain-based Computer**

If you modify the Local Group Policy object for a computer that is participating in a domain while the computer is disconnected from the network, the change is applied only after the computer is reconnected to the network. This is caused by two facts: the entire domain hierarchy must first be evaluated to find the resultant set of policy settings that apply to the computer and domain user, and domain-based Group Policy settings always take precedence over local Group Policy settings. Therefore, the computer can use only the existing policy settings (no new policy changes can be evaluated) until the computer is reconnected to the network.

# Group Policy Loopback Support

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy is applied to the user or computer, based upon where the user or computer object is located in Active Directory. However, in some cases, users may need policy applied to them, based upon the location of the computer object, not the location of the user object. The Group Policy *loopback* feature gives the administrator the ability to apply user Group Policy, based upon the computer that the user is logging onto.

To describe the loopback feature, we'll use an example. In this scenario, you have full control over the computers and users in this domain because you have been granted Domain Admin privileges.

The following illustration shows the Reskit domain, which is used to work through this example.




**Figure 3. The Reskit domain**

Normal user Group Policy processing specifies that computers located in the Servers organizational unit have the GPOs A3, A1, A2, A4, and A6 applied (in that order) during computer startup. Users of the Marketing organizational unit have GPOs A3, A1, A2, and A5 applied (in that order), regardless of which computer they log on to.

In some cases this processing order may not be what you want to do, for example, when you do not want applications that have been assigned or published to the users of the Marketing organizational unit to be installed while they are logged on to the computers in the Servers organizational unit. With the Group Policy loopback feature, you can specify two other ways to retrieve the list of GPOs for any user of the computers in the Servers organizational unit:

- **Merge mode**. In this mode, the computer's GPOs have higher precedence than the user's GPOs. In this example, the list of GPOs for the computer is A3, A1, A2, A4, and A6, which is added to the user's list of A3, A1, A2, A5, resulting in A3, A1, A2, A5, A3, A1, A2, A4, and A6 (listed in lowest to highest priority).

- **Replace mode**. In this mode, the user's list of GPOs is not gathered. Only the list of GPOs based upon the computer object is used. In this example, the list is A3, A1, A2, A4, and A6.

You can set the loopback feature by using the **User Group Policy loopback processing mode** policy under Computer Settings\Administrative settings\System\Group Policy.

The processing of the loopback feature is implemented in the Group Policy engine, which is the part of Group Policy that runs in the Winlogon process. When the Group Policy engine is about to apply user policy, it looks in the registry for a computer policy, which specifies which mode user policy should be applied in.

**Using Loopback for Terminal Services**

You can apply GPOs to Terminal Servers exclusively with the use of a GPO Loopback policy. This policy directs the system to apply the set of GPOs for the computer to any user who logs on to the computer affected by this policy. This policy is intended for special-use computers, such as those in public places, laboratories, and classrooms, where you must modify the user policy

based on the computer that is being used. Without Loopback, the user's GPOs determine which user policies apply. If this policy is enabled, the location of a user's computer object is the main factor in determining which set of GPOs are to be applied.

**Loopback Processing and Security Filtering**

In security filtering, if you have used the Deny ACL to explicitly prevent a policy setting from applying to a computer, the setting could still apply in loopback replace mode because the user's security principal remains unaffected by the Deny ACL computer settings.

# Design Considerations for Organizational Unit Structure and Use of Group Policy Objects

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses issues you need to consider when planning and implementing your organizational unit structure, and highlights recommendations for the use of GPOs.

**Organizational Unit Structure**
The Group Policy architecture is flexible and allows for many types of design. The guiding principle as you design your organizational unit structure should be to create a structure that is easy to manage and troubleshoot. There are two key reasons to create an organizational unit:

- To enable delegation of administration.
- To scope the application of GPOs.

In general, do not try to model your organizational unit structure based on your business organization. Rather, design your organizational unit structure based on how you administer your business. Information about planning for Active Directory is available in Best Practice Active Directory Design for Managing Windows Networks at http://www.microsoft.com/technet/prodtechnol/ad/windows2000/plan/bpaddsgn.asp.

In most organizations, organizational unit structure is likely to fall into one of the following categories:

- Flat organizational unit structure: 1 or 2 levels
- Narrow organizational unit structure: 3 to 5 levels
- Deep organizational unit structure: more than 5 levels

For organizations with simple administration requirements, it is recommended that administrators use a simple model in which a flat organizational unit structure is used and GPOs are linked at the domain or organizational unit level. Limited use of security groups or WMI filtering to filter GPOs is recommended. If you need additional flexibility, it is suggested that you reconsider your organizational unit structure.

For organizations with moderate administration requirements, it is recommended that administrators use a narrow organizational unit structure and GPOs are linked at the site, domain, or organizational unit level as necessary. Limited use of the Block Policy Inheritance options, the Enforce Policy options, security groups or WMI filtering to filter GPOs is recommended.

For organizations with complex administration requirements, the Active Directory namespace may use flat, narrow, or deep organizational unit structures. In such cases, administrators should consider the following issues:

- Flat organizational unit model: use security groups and DACLs or WMI filtering to filter effects of GPOs as a primary method, and Block Policy Inheritance and Enforce Policy options as secondary methods.
- Narrow organizational unit model: link to GPOs at site, domain, and organizational unit. As a secondary method, use Block Policy Inheritance and Enforce Policy options, and security groups and DACLs, or WMI filtering for filtering effects of GPOs.
- Deep organizational unit model: link to GPOs at site, domain, and organizational unit with security groups filtering and DACLs or WMI filtering. As a secondary method, use Block Policy Inheritance and Enforce Policy options.

**Design Principles**
This section presents general guidelines for using GPOs and policy features, and includes examples of GPO design.

**Administration of Group Policy Objects**
Delegation of authority, separation of administrative duties, central versus distributed administration, and design flexibility are important factors you'll need to consider when designing Group Policy and selecting which scenarios to use for your organization.

How you design your organizational unit structure and GPOs will depend on the administrative requirements and roles in your corporation. For example, if administrators are organized according to their duties (such as security administrators, logon administrators, and so on), you may find it useful to define these policy settings in separate Group Policy objects.

Delegation of authority will depend largely on whether you use centralized or distributed administration in your corporation. Based on their particular corporate requirements, network administrators can use security groups and Discretionary Access Control List permissions to determine which administrator groups can modify policy settings in GPOs. Network administrators can define groups of administrators (for example, Software Installation administrators), and then provide them read and write access to selected GPOs, allowing the network administrator to delegate control of the GPO settings. Administrators who have read and write access to a Group Policy object can by default control all of the contents of that Group Policy object; however, you can restrict access by setting policy to control which MMC snap-ins can be loaded by that user, as described earlier in the "Delegating Group Policy" section.

**Separate Users and Computers into Different organizational units**
It's recommended that you separate users and computers into separate organizational units. This is useful for these reasons:

- This simplifies GPO design because you need to focus on only configuration of either user or computers.
- Typically users and computers are administered differently, perhaps by different groups within your organization, which facilitates administration.
- You can reduce Group Policy processing time because you can disable the unused half of the GPO. It is possible to disable only the User or Computer portion of the GPO. To do this, right-click the GPO, click **Properties**, click either **Disable Computer Configuration settings** or **Disable User Configuration settings**, and then click **OK**. These options are available on the GPO **Properties** page, on the **General** tab.
- This type of design is required to enable loopback processing. See the "Group Policy Loopback Support" section for more information.
- For increased security and ease of administration, you should specify different organizational units for all new user and computer accounts when they are created, as explained below.

**Redirecting the Users and Computers Containers in Windows Server 2003 Domains**
New user and computer accounts are created in the CN=Users and CN=Computers containers by default. It is not possible to apply Group Policy directly to these containers, although they inherit GPOs linked to the domain.

Redirusr.exe (for user accounts) and Redircomp.exe (for computer accounts) are two new tools included with Windows Server 2003 that enable you to change the default location where new user and computer accounts are created so you can more easily scope GPOs directly to newly created user and computer objects. These tools are located in %windir%\system32. By running Redirusr.exe and Redircomp.exe once for each domain, the domain administrator can specify the organizational units into which all new user and computer accounts are placed at the time of creation. This allows administrators to manage these unassigned accounts by using Group Policy before the administrators assign them to the organizational unit in which they are finally placed. You might want to consider restricting the organizational units used for new user and computer accounts by using Group Policy to increase security around these accounts.

For more information about redirecting users and computers, see article 324949, "Redirecting the Users and Computers Containers in Windows Server 2003 Domains," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about the redirusr.exe and redircomp.exe tools, see the Redirecting Users and Computers link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Best Practice Organizational Unit Structure**
Because you cannot apply Group Policy directly to the CN=Users and CN=Computers containers, if you wanted to define policy settings for users and computers to be stored in their default container you needed to do so on the root of the domain. To prevent policy settings that are defined on a superior container (the root of the domain) from applying to users, computers, and groups in subordinate CN and organizational unit containers, you needed to define complex ACLs on the policy setting in

the root of the domain.

The solution for Windows 2000 and Windows Server 2003 domains is to deploy the best-practice organizational unit structure where Users, Computers, Groups, Service Accounts and Admin accounts are each in their own organizational unit.

The following list describes the benefits of using the best-practice organizational unit structure:

- It permits administrators to link GPOs directly to the containers that are hosting users and computers.
- It permits administrators to match GPOs to objects of a common object class. For example, User or Computer policy settings can be linked directly to organizational units that are hosting user or computer accounts.
- It permits non-administrators to apply policy on containers that are not hosting security-sensitive users and groups such as Domain Admins, Schema Admins, or Enterprise Admins.
- It can minimize the effect if an organizational unit is accidentally deleted (this assumes that the parent container is correctly protected).
- It permits you to restore users and groups independently of each other in recovery scenarios. User accounts must exist before the restoration of the group. Having users and groups reside in different containers permits you to restore them and mark them as authoritative independently of each other.

Note that the CN=USERS and CN=COMPUTERS containers are computer-protected objects. You cannot (and must not) remove them for backward compatibility purposes although they can be renamed.

The best-practice organizational unit structure works well for storing existing users, computers, and groups in Active Directory because those objects can be moved into the appropriate organizational unit container on Windows 2000 and Windows Server 2003 domains regardless of its domain or forest functional level. New user accounts, computer accounts, and security groups that are created with earlier-version APIs used by GUI and command-line management tools do not allow administrators to specify a target organizational unit. As a result, these objects will initially be created in the CN=Users and CN=Computers containers until they are moved by the administrator or an administrator-defined script.

For more information about best-practice organizational unit structure see the "Creating an Organizational Unit Design" section of the Best Practice Active Directory Design for Managing Windows Networks at http://www.microsoft.com/technet/prodtechnol/ad/windows2000/plan/bpaddsgn.asp

**Functional Compared with Geographical Organizational Unit Structure**
When organizing organizational units, there are two basic models to start with: functional and then geographical, or geographical and then functional. The key is never to implement a structure that forces an artificial layering, which means that the organizational unit structure for computers may be very different than that for users—it all depends on how they are administered.

**Minimize the Number of Group Policy Objects Associated with Users or Computers**
You should note that the number of GPOs that are applied to a user affects the logon processing time. (Similarly, the number of GPOs applied to a computer affects boot time). The greater the number of associated GPOs, the longer logon will take to process them. During logon time, each GPO from the user's site, domain, and organizational unit hierarchy is applied, provided the user has both the Read ACE and the Apply Group Policy ACE. Note that if the Apply Group Policy ACE is not set, but the Read ACE is, the GPO will still be processed (although not applied), thus impacting logon time. Therefore, if you implement filtering based on security groups, you should also clear Read Access for those users that you clear Apply Group Policy for.

**Minimize the Use of the Block Policy Inheritance Feature**
As mentioned previously, you can prevent Group Policy settings of parent Active Directory containers from affecting users and computers in lower-level parent Active Directory containers. This is a useful and powerful feature that you should use judiciously only when a particular situation requires it. Blocking the inheritance of policy from parent Active Directory containers can complicate troubleshooting policy.

**Minimize the Use of the Enforce Feature**
You can also ensure that the policy settings you specify in a given GPO at a higher-level parent Active Directory container are enforced on lower-level parent Active Directory containers by using the Enforce option. Only use this powerful feature when circumstances require it. Overuse of this feature with other related features, such as Block Policy Inheritance, can complicate troubleshooting policy.

**Use Loopback Processing Only When Necessary**

You can set User Configuration per computer and thus override user-specific policy settings with computer-specific policy settings. This is useful when you want to provide a specific desktop configuration regardless of which users log on to the computer, such as a kiosk or other public terminal. To set User Configuration per computer, you would use the Administrative Templates node under **Computer Configuration** in the Group Policy Object Editor. For more information about this feature, see Group Policy Loopback Support.

**Avoid Using Cross-Domain GPO Assignments**
Although you can assign GPOs from different domains to a single Active Directory container if a particular situation requires it, you should note that in such cases Group Policy processing would be slower. This is because domain boundaries are crossed.

**Avoid Editing the Default Domain GPO**
Instead of editing the default domain GPO, create a new GPO, link it to the domain GPO, and set the new GPO to have precedence over the default domain GPO.

**Design Examples**
This section presents several models of GPO design. These examples are not intended as guidelines, but they do illustrate various ways to approach GPO design. In most corporate environments, administrators may use a combination of these or similar models, tailored to their business requirements.

The key overriding approaches are either functional or geographic models. The rest are usually variants of those.

**Layered GPO Design Model**
The objective of this design model is to create GPOs based on a layered approach. This approach optimizes maintenance of GPOs and facilitates delegation.

The following graphic illustrates an example of this model.



Layered GPO model

**Monolithic GPO Design Model**
The objective of this design is to create GPOs based on a monolithic design—an approach that reduces the number of GPOs that apply to a user and/or computer but may not be optimal for delegation.

The following graphic illustrates an example of the monolithic GPO model.



Monolithic GPO model

**Single Policy Type GPO Design Model**
The objective of this design is to create GPOs that deliver a single type of Group Policy, for example, policy for security settings. Such a design optimizes separation of duties for administrators; however, it may increase the number of GPOs that are applied to a given user or computer.

Each GPO delivers only one type of policy (security GPOs are different from script Group Policy objects, for example). Large corporations often create separate administrator groups based on administrative duties; this scenario would be useful in such corporate environments.

The following graphic illustrates an example of the single policy type GPO model.



Single Policy GPO Model

**Multiple Policy Types GPO Design Model**
The objective of this design is to create GPOs that deliver multiple types of policy. This is a hybrid of the single policy and monolithic models. Each GPO delivers several types of policy settings.

For example, you can create a GPO that includes Group Policy settings for software settings and application deployment and create another GPO that includes security and scripts settings, and so on. A GPO design that supports multiple policy types is useful in delegating administration environments and can reduce the number of GPOs that apply to a user and/or computer.

The following graphic illustrates an example of the multiple policy types GPO model.



Multiple Policy Types GPO

**Teams or Matrix Organizations GPO Model**
This model applies to organizations that leverage the virtual team concept. Individuals within the organization form teams to perform a task or project and each individual is a member of multiple teams. Each team has specific Group Policy requirements. The organizational unit architecture does not reflect the team structure. This model works by using security filtering.

The following graphic illustrates an example of the team GPO design model.




**Public Computing Environment GPO Model**

This scenario applies to environments were you want the computer Group Policy settings to always have precedence over the user Group Policy settings. This scenario is useful for training classes and kiosk-type environments in which you want to provide the same desktop environment regardless of which user logs on to the computer.

The following graphic illustrates an example of the GPO design for a public computing environment. The loopback policy feature with **Replace mode** is used in this example. See "Group Policy Loopback Support" in this document for more information.

Normal Group Policy processing specifies that users in the Sales organizational unit get these GPOs: Domain Policy GPO, Accounting GPO, and Sales GPO. With the loopback policy enabled in **Replace mode**, when users from the Sales organizational unit log on to a computer in the Kiosks organizational unit, the user will process *only* these GPOs: Domain Policy GPO, Accounting GPO, Resources GPO, and Kiosks Loopback Policy GPO—the users' list of GPOs is not gathered in this case. More specifically, the user settings specified in the Kiosks organizational unit (and those inherited) are the *only* GPOs processed for the user logging onto computers in that organizational unit. Those in the Users organizational unit tree are not processed.

**Delegation with Central Control**

This model applies to organizations that choose to delegate administration of GPOs, but would like to enforce certain Group Policy settings throughout the domain (for example, specific security policy settings).




The following graphic illustrates an example of GPO delegation with centralized control, and use of the **Enforce** option.




**Delegation with Distributed Control**

This scenario applies to organizations that want to allow administrators of organizational units to prevent Group Policy settings from being applied to their organizational unit. Administrators of an organizational unit can block Group Policy settings that have been assigned at higher levels in the hierarchy from applying to his or her organizational unit. However, administrators cannot block Group Policy settings that are marked as *Enforce*.

This feature allows organizations to minimize the number of domains without sacrificing autonomy.

# Deployment Considerations

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering a Mix of Windows 2000 and Windows Server 2003 Domains**
GPMC exposes features that are available in the underlying operating system. Because new features have been added to Group Policy since Windows 2000, certain features will only be available in GPMC depending on the operating system that has been deployed on the domain controllers. This section describes these dependencies. In general, there are three key issues that determine whether a feature is available in GPMC:

- Whether the forest supports the Windows Server 2003 schema for Active Directory. Certain features are only available once the schema is upgraded. This is the first step that must be taken before any Windows Server 2003 domain controller can be deployed in an existing Windows 2000 forest. The schema is a forest-wide configuration and is upgraded by running **ADPrep /ForestPrep**. ADPrep is a utility included on the Windows Server 2003 CD. Note that it is possible to have the Windows Server 2003 schema in a forest with all Windows 2000 domain controllers.

- Whether there is at least one domain controller in the forest that is running Windows Server 2003. Group Policy Modeling must be performed on a domain controller running Windows Server 2003.

- Whether a domain contains the Windows Server 2003 domain configuration. This is implemented once **ADPrep /DomainPrep** is run in that domain. This is the first step that must be taken before any Windows Server 2003 domain controller can be deployed in an existing Windows 2000 domain.

Note that there is no dependency from the Group Policy perspective on whether a domain is in native mode or mixed mode.

**Delegation of Group Policy Results and Group Policy Modeling**
In order to delegate either Group Policy Modeling or Group Policy Results, the Active Directory schema in the forest must be the Windows Server 2003 schema. Note that you can use Group Policy Results even without this schema, but only users with local administrative credentials on the target computer can remotely access Group Policy Results data. Thus, if the forest does not have the Windows Server 2003 schema, the delegation pages in GPMC for organizational units and domains will not show these permissions.

**Group Policy Modeling**
Group Policy Modeling is a simulation that is performed by a service that can only run on a domain controller running Windows Server 2003 or later. As long as there is at least one domain controller running Windows Server 2003 in the forest, you can use Group Policy Modeling. GPMC will only show the Group Policy Modeling node in the user interface if the Windows Server 2003 schema is present.

**WMI Filtering**
WMI filters are only available in domains that have the Windows Server 2003 configuration. Although none of the domain controllers need to be running Windows Server 2003, you must have run ADPrep /DomainPrep in this domain. Also note that WMI filters are only evaluated by clients running Windows XP, Windows Server 2003, or later. WMI filters associated with a GPO will be ignored by Windows 2000 clients and the GPO will always be applied on Windows 2000.

If ADPrep /DomainPrep has not been run in a given domain, the WMI Filters node will not be present, and the GPO scope tab will not have a WMI filters section.

**Upgrading Windows 2000 Domains to Windows Server 2003 Domains and Interaction with Group Policy Modeling**
Group Policy Modeling is a new feature of Windows Server 2003 that simulates the resultant set of policy for a given configuration. The simulation is performed by a service that runs on Windows Server 2003 domain controllers. In order to perform the simulation in cross-domain scenarios, the service must have read access to all GPOs in the forest.

In a Windows Server 2003 domain (whether it is upgraded from Windows 2000 or installed as new), the Enterprise Domain Controllers group is automatically given read access to all newly created GPOs. This ensures that the service can read all GPOs in the forest.

However, if the domain was upgraded from Windows 2000, any existing GPOs that were created before the upgrade do not have read access for the Enterprise Domain Controllers group. When you click a GPO, GPMC detects this situation and notifies the user that Enterprise Domain Controllers do not have read access to all GPOs in this domain. To solve this problem, you can use one of the sample scripts provided with GPMC, GrantPermissionOnAllGPOs.wsf. This script can update the permissions for

all GPOs in the domain. To use this script:

Ensure that the person running this script is either a Domain Admin or has permissions to modify security on all GPOs in the domain.

Open a command prompt and navigate to the **%programfiles%\gpmc\scripts** folder by typing:

**CD /D %programfiles%\gpmc\scripts**

Type the following:

**Cscript GrantPermissionOnAllGPOs.wsf "Enterprise Domain Controllers" /Permission:Read /Domain:value**

The value of domain parameter is the DNS name of the domain.

**Using Group Policy Features Across Forests**

The Windows Server 2003 family introduces a new feature called Forest Trust that enables you to authenticate and authorize access to resources from separate, networked forests. With trusts established between forests, you can manage Group Policy throughout your enterprise, which provides greater flexibility especially in large organizations. This section describes Group Policy behavior in an environment with forest trust enabled:

- It is not possible to link a GPO to a domain in another forest.
- With Forest trust, it is possible that a user in Forest B could log onto a computer in Forest A. In this case, when the computer starts up, it will process policy for the computer configuration from Forest A, as usual. When a user from Forest B logs on, where they receive their policy settings from depends on the value of the **Allow Cross-Forest User Policy and Roaming Profiles** policy setting.
    - When this setting is **Not Configured**, no user-based policy settings are applied from the user's forest. Instead, loopback Group Policy processing will be applied, using the GPOs scoped to the computer. Users will receive a local profile instead of their roaming profile.
    - When this setting is **Enabled**, the behavior is exactly the same as Windows 2000 Server: User policy is applied from the user's forest and a roaming user profile is allowed from the trusted forest.
    - When this setting is **Disabled**, the behavior is the same as **Not Configured**.

This setting is available on Windows Server 2003 located at: **Computer Configuration\Administrative Templates\System\Group Policy\Allow Cross-Forest User Policy and Roaming Profiles.**

- It is possible to deploy Group Policy settings to users and computers in the same forest, but have those settings reference servers in other trusted forests. For example, the file shares that host software distribution points, redirected folders, logon scripts, and roaming user profiles could be in another trusted forest.
- Group Policy Modeling requires that both the user and the computer be in the same forest. If you want to simulate a user from Forest A logging on to a computer in Forest B, you must perform two separate Group Policy Modeling simulations: one for the user configuration and the other for the computer configuration.
- Delegation across forests is supported for managing Group Policy. For example, you can delegate to someone in Forest B the ability to perform Group Policy Modeling simulations on objects in Forest A.

**Group Policy and Active Directory Sites**

GPOs that are linked to site containers affect all computers in a forest of domains. Site information is replicated and available between all the domain controllers within a domain and all the domains in a forest. Therefore, any GPO that is linked to a site container is applied to all computers in that site, regardless of the domain (in the forest) to which they belong. This has the following implications:

- It allows multiple domains (within a forest) to get the same GPO (and included policy settings), although the GPO only lives on a single domain and must be read from that domain when the affected clients read their site policy.
- If child domains are set up across wide area network (WAN) boundaries, the site setup should reflect this. If it does not,

the computers in a child domain could be accessing a site GPO across a WAN link.

- To manage site GPOs, you need to be either an Enterprise Admin or Domain Admin of the forest root domain.

- You may want to consider using site-wide GPOs for specifying policy for proxy settings and network-related settings.

In general, it is recommended that you link GPOs to domains and organizational units rather than sites.

**Using Group Policy and Internet Explorer Enhanced Security Configuration**
Windows Server 2003 includes a new default security configuration for Internet Explorer, called Internet Explorer Enhanced Security Configuration, also known as Internet Explorer hardening.

You can manage Internet Explorer Enhanced Security Configuration by:

- **Enabling or disabling Internet Explorer Enhanced Security Configuration.** This is commonly used in situations where you want to ensure that Internet Explorer Enhanced Security Configuration is always enabled. For example, Internet Explorer Enhanced Security Configuration might need to be reapplied on a specific computer if the local administrator on that computer disables it using the Optional Component Manager in the Windows Components Wizard (available from Add or Remove Programs.)

- **Restricting who can manage trusted sites and other Internet Explorer security settings on a server**. This is commonly used when you want to ensure that all servers have the same Internet Explorer Enhanced Security Configuration settings. For example, you might want to configure Internet Explorer Enhanced Security Configuration so that machined-based security settings are applied to each server rather than user-based security settings.

- **Adding trusted Web sites and UNC paths to one of the trusted security zones.** This is commonly used when you want to allow users access to specific Web sites and corporate resources, but still reduce the risk of users downloading or running malicious content.

Enhanced Security Configuration impacts the Security Zones and Privacy settings within the Internet Explorer Maintenance settings of a GPO. The Security Zones and Privacy settings can either be enabled with Enhanced Security Configuration or not.

When you edit settings for Security Zones and Privacy settings in a GPO from a computer where Enhanced Security Configuration is enabled, that GPO will contain Enhanced Security Configuration-enabled settings. When you look at the HTML report for that GPO, the Security Zones and Privacy heading will be appended with the text (Enhanced Security Configuration enabled).

When you edit settings for Security Zones and Privacy settings in a GPO from a computer where Enhanced Security Configuration is not enabled , that GPO will contain Enhanced Security Configuration-disabled settings. ESC is not enabled on any computer running Windows 2000 or Windows XP, nor on computers running Windows Server 2003 where ESC has been explicitly disabled.

Enhanced Security Configuration settings deployed through Group Policy will only be processed on and applied by computers where Enhanced Security Configuration is enabled. Enhanced Security Configuration settings will be ignored on computers where Enhanced Security Configuration is not enabled (all computers running Windows 2000 and Windows XP, and Windows Server 2003 computers where Enhanced Security Configuration has been explicitly disabled). The converse is also true: A GPO that contains non- Enhanced Security Configuration settings will only be processed on and applied by computers where Enhanced Security Configuration is not enabled.

For more information, see Managing Internet Explorer Enhanced Security Configuration, available from the Microsoft Group Policy Web site at http://www.microsoft.com/grouppolicy.

# IntelliMirror Features without Active Directory

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The full functionality of IntelliMirror® management technologies requires Active Directory and Group Policy. However, in an environment without Active Directory and Group Policy, some of the capabilities are available. You can still implement the following IntelliMirror features to manage clients running Windows 2000 or later:

- Roaming User Profiles and Logon Scripts
- Folder Redirection
- Internet Explorer Maintenance
- Administrative Templates (registry-based policy)

**Roaming User Profiles and Logon Scripts**
When using either a Windows NT 4.0 domain or Active Directory, both roaming user profiles and logon scripts are configured on the user object.

**Folder Redirection**
You can redirect special folders to alternate locations, either to a local or network location. You do this by modifying the values under the following registry key:

**HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\Shell Folders**

Each value is of type **REG_SZ**, and the data is the redirected path (either local or UNC). The table below lists the folders that may be redirected and their associated value name.

| Folder | Name |
|---|---|
| My Documents | Personal |
| My Pictures | My Pictures |
| Application Data | AppData |
| Desktop | Desktop |
| Start Menu | Start Menu |

**Internet Explorer Maintenance**
Instead of using Group Policy to control Internet Explorer settings, administrators can use the Internet Explorer Administration Kit (IEAK) to apply settings to Internet Explorer clients using auto-configuration packages.

**Applying Administrative Templates (Registry-Based Policy)**
Domain-based Group Policy processing requires that the User and/or Computer objects be located in Active Directory. If the User or Computer objects are located in a Windows NT 4.0 domain, then Windows NT 4.0 System Policy will be processed for whichever of these objects is located in that domain—this could be the Computer or User object, or both. System Policy is defined as the policy mechanism used natively in Windows NT 4.0; it is a set of registry settings that together define the computer resources available to a group of users or an individual. (Also be aware that the local GPO is always processed prior to any System Policy.)

**Setting Registry-based Policy in a Windows NT 4.0 Domain**
A client running Windows 2000 or Windows XP Professional will process System Policy if either the user or computer account are in a Windows NT 4.0 domain. The client looks for the Ntconfig.pol file used by Windows NT 4.0-style System Policy. By default, it looks for this file in the NETLOGON file share of the authenticating Windows NT 4.0 domain controller.

# Migrating Policy-Enabled Clients from Windows NT 4.0 to Windows 2000 or Windows Server 2003

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses behavior of Group Policy and System Policy in relation to migration to Windows 2000 or Windows Server 2003.

**Windows NT 4.0 and Windows 2000 Policy Setting Comparison**
Group Policy differs greatly from System Policy in Windows NT 4.0. Although Group Policy does include the functionality from Windows NT 4.0 System Policy, it also provides policy settings for scripts, software installation, security settings, Internet Explorer maintenance, folder redirection, and Remote Installation Services.

In Windows NT 4.0 (and Windows 95 and Windows 98), System Policies:

- Are applied to domains.
- May be further controlled by user membership in security groups.
- Are not secure.
- Persist in users' profiles (this is sometimes referred to as *tattooing* the registry), as explained earlier in this paper This means that after a registry setting is set using Windows NT 4.0 System Policies, the setting persists until the specified policy is reversed or the user edits the registry.
- Are limited to desktop lockdown.

In Windows 2000 and Windows Server 2003, Group Policy:

- Represents the primary method for enabling centralized Change and Configuration Management. You can use Group Policy to manage registry-based policy, software installation options, security settings, scripts (for computer startup and shutdown, and for user logon and logoff), Internet Explorer maintenance, folder redirection, and Remote Installation Services.
- Can be linked to sites, domains, and organizational units.
- Affects all users and computers in the specified Active Directory container (site, domain, or organizational unit) by default.
- May be further controlled by user or computer membership in security groups.
- May be further controlled by use of WMI filtering.
- Settings are secure.
- Default policy settings do not persist in the registry.
- Can be used for tightly managed desktop configurations and to enhance the user's computing environment.

The Windows NT 4.0 effect of persistent registry settings can be problematic when a user's group membership is changed. An advantage of Windows 2000Group Policy is that this does not occur. When a GPO no longer applies, registry settings written to the following secure registry locations are removed:

- HKLM\Software\Policies

- HKLM\Software\MS\Windows\CurrentVersion\Policies
- HKCU\Software\Policies
- HKCU\Software\MS\Windows\CurrentVersion\Policies

**Migrating to Windows 2000 or Windows Server 2003**
Migrating Windows NT 4.0-based clients and servers to Windows 2000 or Windows Server 2003 in various combinations causes different behavior for Group Policy. In a pure Windows 2000 or later environment where both the user and computer accounts are in a Windows 2000 or later domain, Windows 2000 or later clients process only Group Policy. System Policy is not processed. However, Windows 2000 or Windows XP clients can process System Policy in cases where either the user account and/or the computer account is not located in a Windows 2000 or Windows Server 2003 domain.

In many organizations it may be impractical to upgrade all Windows NT 4.0-based servers and client computers simultaneously to Windows Server 2003 and Windows XP. In this case, it is important that you know how Group Policy and Windows NT 4.0 System Policy are affected during and after the migration process. This section presents information about the effects of migration on Group Policy.

**Client Computers**
Group Policy applies only to computers running Windows 2000 or later. There is no mechanism to process Group Policy on clients running Windows NT 4.0, Windows 95, Windows 98, and Windows Millennium Edition.

**Upgrading Computer or User Accounts from Windows NT 4.0 to Windows Server 2003**
When migrating from Windows NT 4.0, it's recommended to perform a clean installation of Windows Server 2003. To facilitate a clean installation, you can use the User State Migration Tool to migrate the users' data and settings to the new installation.

For more information about migrating from Windows NT 4 System Policy, see Windows 2000 Group Policy white paper at http://go.microsoft.com/fwlink/?linkid=203.

**Using Group Policy in a Mixed Environment of Windows 2000 and Windows XP Clients**
**Active Directory with Windows 2000 and Windows XP Clients**
This section explains issues to consider when using Group Policy in a Windows 2000 Server or Windows Server 2003 environment where some or all of the clients are running Windows XP Professional or Windows 2000.

**Comparing IntelliMirror Features on Windows 2000 and Windows XP**
The following tables show how IntelliMirror features compare on computers running Windows 2000 Professional and Windows XP Professional.

### Comparing Clients under Windows Server 2003 Active Directory

| Feature | Supported in Windows 2000 Client | Supported in Windows XP Client |
|---|---|---|
| Group Policy | Yes | Yes |
| GPMC | No. But Windows 2000 clients can be managed with GPMC running on Windows Server 2003 | Yes |
| Local Group Policy | Yes | Yes |
| System policy | Yes | Yes |
| Roaming profiles | Yes | Yes |
| Folder redirection | Yes (No home directory redirect) | Yes |
| Software installation | Yes | Yes |
| Internet Explorer Maintenance | Yes | Yes |

| | | |
|---|---|---|
| Security Settings | Yes | Yes |
| Software restriction policies | No | Yes |

**Comparing clients under Windows 2000 Active Directory**

| Feature | Supported in Windows 2000 | Supported in Windows XP |
|---|---|---|
| Group Policy | Yes | Yes |
| GPMC | No. But Windows 2000 clients in a Windows 2000 domain can be managed with GPMC installed on a computer running Windows XP or a member server running Windows Server 2003. | Yes |
| Local Group Policy | Yes | Yes |
| System policy | Yes | Yes |
| Roaming profiles | Yes | Yes |
| Folder redirection | Yes (No home directory redirect) | Yes |
| Software installation | Yes | Yes |
| Internet Explorer Maintenance | Yes | Yes |
| Security Settings | Yes | Yes |
| Software restriction policies | No | Yes (via Local Group Policy Object) |

**Comparing Clients Under Windows NT Server 4.0**

| Feature | Supported in Windows 2000 | Supported in Windows XP |
|---|---|---|
| System policy | Yes | Yes |
| Group Policy | No | No |
| Local Group Policy | Yes | Yes |
| Roaming profiles | Yes | Yes |
| Folder redirection | No | No |
| Software installation | No | No |
| Internet Explorer Maintenance | Yes, with Internet Explorer Administration Kit (IEAK) | Yes, with IEAK. |

| | | |
|---|---|---|
| Security Settings | No | No |
| Software restriction policies | No | No |

**Folder Redirection and Software Installation**
Because background refresh is the default behavior in Windows XP, Folder Redirection and Software Installation may require as many as three logons to apply changes.

This behavior exists because Folder Redirection and Software Installation cannot apply during an asynchronous or background application of policy. Folder Redirection can only apply when processed synchronously.

Here is a sample scenario showing how polices are applied:

1. An administrator deploys a software package to User A.

2. User A logs on fast and receives a background (asynchronous) application of policy.

3. Because the policy application was asynchronous, the software that was set to be installed cannot be installed at this time. Instead the machine is tagged, indicating that software needs to be installed.

4. The next time the user logs on, the machine instead logs on the user synchronously to allow the software package to be installed. (This is the same behavior as Windows 2000). This results in one extra logon for the software to be installed.

In the case of Advanced folder redirection, because policy is evaluated based on security group membership three logons will be required: the first logon to update the cached user object (and security group membership), the second logon for policy to detect the change in security group membership and require a foreground policy application, and the third logon to actually apply folder redirection policy in the foreground.

| **Note** |
|---|
| When a client running Windows XP logs onto a Windows 2000 or Windows Server 2003 Active Directory, all Software Installation policy settings for Windows 2000 clients will be applied and work successfully on the Windows XP client. |

**Internet Explorer Maintenance**
There are no changes in Internet Explorer Maintenance across Windows XP and Windows 2000.

**Roaming Profiles**
Users with roaming profiles can roam between Windows 2000 and Windows XP-based workstations without any changes in behavior. The new profile registry policy settings only work on Windows XP. If you apply these settings to a client running Windows 2000, they will have no effect.

**Security Settings**
Software Restrictions Policies were introduced in Windows XP. If you apply software restriction policy to a client running Windows 2000 it will have no effect. The software restriction policy registry settings will be written to the registry, but the Windows 2000 client will not know how to interpret them.

**64 bit Integration Issues**
If you apply a 64-bit package to Windows 2000 or a 32-bit version of Windows XP, it will not be advertised by default; however, you can override this behavior using the 64-bit deployment options in the Application Deployment Editor (ADE). If you apply a 64-bit package to a 64-bit version of Windows XP, it will be successfully advertised.

# Appendix A: Security Settings and User Rights

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix lists the Security Settings that are defined by default in the Default Domain Policy GPO. This GPO is created when the first domain controller in the domain is installed by the Active Directory Installation Wizard. If this first domain controller is upgraded from a Windows NT 4.0 domain controller, then the values defined for the Windows NT 4.0 domain are used instead.

These domain-wide account policy settings (Password Policy, Account Lockout Policy and Kerberos Policy) are enforced by the domain controller computers in the domain; therefore, all domain controllers always retrieve the values of these account policy settings from the Default Domain Policy GPO.

| **Policy** | **Default Value** | **Comment** |
|---|---|---|
| **Password Policy** | | |
| Enforce password history | 1 password remembered | |
| Maximum password age | 42 days | |
| Minimum password age | 0 days | |
| Minimum password length | 0 characters | |
| Passwords must meet complexity requirements | Disabled | |
| Store password using reversible encryption for all users in the domain | Disabled | |
| **Account Lockout Policy** | | |
| Account Lockout Threshold | 0 | |
| **Kerberos Policy**<br>Since Kerberos support was not available in previous versions of Windows NT, the following Kerberos policy settings are always defined for the first domain controller of a Windows 2000 or Windows Server 2003 domain, regardless of whether it was upgraded or not. | | |
| Enforce user logon restrictions. | Enabled | |
| Maximum lifetime that a user ticket can be renewed | 7 days | |
| Maximum user ticket lifetime | 10 hours | |
| Maximum service ticket lifetime | 600 minutes | |
| Maximum tolerance for synchronization of computer clocks | 5 minutes | |
| **Security Options** | | |

| | | |
|---|---|---|
| Automatically logoff users when logon time expires | Disabled | This is a domain-wide setting even though it appears under the Security Options area. |

**Security Settings in the Default Domain Controllers Policy**

This section lists the Security Settings that are defined by default in the Default Domain Controller Policy GPO. This GPO is created when the first domain controller in the domain is installed via the Active Directory Installation Wizard. If this first domain controller is upgraded from a Windows NT 4.0 domain controller, then the values defined for the Windows NT 4.0 domain are used instead.

By default, these settings apply to all domain controllers in the domain.

| Policy | Default Value | Comment |
|---|---|---|
| **Security Options** | | |
| Digitally sign server-side communication when possible | Enabled | |
| **Audit Policy** | | |
| Audit Account Logon events | No Auditing | |
| Audit Account Management | No Auditing | |
| Audit Directory Service Access | No Auditing | |
| Audit Logon Events | No Auditing | |
| Audit Object Access | No Auditing | |
| Audit Policy Change | No Auditing | |
| Audit Privilege Use | No Auditing | |
| Audit Process Tracking | No Auditing | |
| Audit System Events | No Auditing | |
| **User Rights Policy** | | |
| Access this computer from the network | Administrators, Authenticated Users, Everyone | If the following groups were given this right prior to running the Active Directory Installation Wizard, then they are removed: Backup Operators, Guests, Guest, and Users.<br><br>If a Windows NT 4.0 domain controller is upgraded as the first Windows Server 2003 domain controller, then the Authenticated Users group is automatically given this right. |
| Act as part of the operating system | | |
| Add workstations to the domain | Authenticated Users | This User Right is for the support of legacy APIs. You can also allow users to create computer accounts by using this User Right. Authenticated Users can only create 10 computer accounts using this User Right. |

| | | |
|---|---|---|
| Back up files and directories | Administrators, Backup Operators, Server Operators | |
| Bypass traverse checking | Administrators, Authenticated Users, Everyone | If the following groups were given this right prior to running the Active Directory Installation Wizard, then they are removed: Backup Operators, Users. |
| Change the system time | Administrators, Server Operators | |
| Create a pagefile | Administrators | |
| Create a token object | | |
| Create permanent shared objects | | |
| Debug programs | Administrators | |
| Force shutdown from a remote system | Administrators, Server Operators | |
| Generate security audits | | |
| Increase quotas | Administrators | |
| Increase scheduling priority | Administrators | |
| Load and unload device drivers | Administrators | |
| Lock pages in memory | | |
| Log on as a batch job | | |
| Log on as a service | | |
| Log on locally | Account Operators, Administrators, Backup Operators, Server Operators, Print Operators | If the following groups were given this right prior to running the Active Directory Installation Wizard, then they are removed: Authenticated Users, Guests, Guest, Users, and Everyone. |
| Manage auditing and security log | Administrators | |
| Modify firmware environment variables | Administrators | |
| Profile single process | Administrators | |
| Profile system performance | Administrators | |
| Replace a process-level token | | |
| Restore files and directories | Administrators, Backup Operators, Server Operators | |

| | | |
|---|---|---|
| Shut down the system | Account Operators, Administrators, Backup Operators, Server Operators, Print Operators | If the following groups were given this right prior to running the Active Directory Installation Wizard, then they are removed: Authenticated Users, Guests, Guest, Users, and Everyone. |
| Take ownership of files or other objects | Administrators | |
| Deny Logon Locally | | |
| Deny logon as a batch job | | |
| Deny logon as a service | | |
| Deny Access to this computer from network | | |
| Remove Computer from Docking Station | Administrators | If the following groups were given this right prior to running the Active Directory Installation Wizard, then they are removed: Users. |
| Synchronize directory service data | | |
| Enable computer and user accounts to be trusted for delegation | Administrators | If the following groups were given this right prior to running the Active Directory Installation Wizard, then they are removed: Users. |

**Help for Windows NT 4.0 Administrators**

This section provides information to help administrators who have been using User Manager to configure security policy settings in the past move to the new model of Group Policy for editing and configuring security policy settings.

**Changing Password Policy for the Domain**

**To change password policy for the domain, open the Default Domain GPO from the Administrative Tools menu:**

1. Click **Start**, point to **Programs**, click **Administrative Tools**, and then click **Domain Security Policy**.

2. In the **Domain Security Policy** console, expand **Security Settings**, expand **Account Policies**, expand **Password Policy**, and then select the policy you want to modify in the results pane. You can then make changes.

**Changing Auditing Policy or User Rights for Domain Controllers**

To change the Audit policies or User Rights defined for domain controllers, open the **Default Domain Controllers GPO** from the Administrative Tools menu:

1. Click **Start**, point to **Programs**, click **Administrative Tools**, and then click **Domain Controller Security Policy**.

2. In the **Domain Controller Security Policy** console, expand **Security Settings**, expand **Local Policies**, click either **Audit Policy** or **User Rights Assignment**, and then select the policy you want to modify in the results pane.

**Changing local Password Policy on member Workstations or Servers (Non-Domain Controllers)**

Because the **Default Domain Policy GPO** applies to all computers in the domain and because domain-level policy settings override local policy settings, member workstations and servers apply the Default Domain password policy settings to their local account databases by default. If this does not meet your requirements, then the permissions on the **Default Domain GPO** have to be reconfigured so that member computers that you do not want to receive this policy do not have the **Apply Group Policy** permission on the **Default Domain GPO**. After the permissions are configured so that the member computer does not have access to the default domain policy, local policy settings will no longer be overridden by the password policy settings defined in the **Default Domain GPO**.

**To modify Local Password Policy security settings using the Local Security Policy UI:**

1. Click **Start**, point to **Programs**, click **Administrative Tools**, and then click **Local Security Policy**.

2. In the **Local Security Settings** console, expand **Security Settings**, expand **Account Policies**, click **Password Policy**, and then select in the results pane the policy you want to edit.

**Frequently Asked Questions about Security Settings**

**Is it possible to define different account policies (Password, Lockout, or Kerberos Policies) for different organizational units?**

No. All domain controllers for a domain enforce the account policy settings that are defined in the Default Domain Policy. Domain controllers ignore password, lockout, or Kerberos policy settings defined at an organizational unit or Local GPO level.

**After modifying a local security setting, the change does not take effect. What is happening?**

The Group Policy model specifies that any policy settings configured locally may be overridden by like policy settings specified in the domain. The **Local Security Settings** UI lists the local security setting and the effective security setting for each policy item. (You can access the **Local Security Settings** UI by clicking **Start**, pointing to **Programs**, clicking **Administrative Tools**, and selecting **Local Security Policy**). If the effective security setting is different from the local security setting, it implies that there is a policy from the domain that is overriding your setting.

**After modifying a domain-level-policy security setting, the change does not take effect. What is happening?**

The Group Policy model applies domain-level policy changes periodically; therefore, it is likely that the policy changes made in the directory have not been made to your computer yet. To trigger a policy propagation on a local computer, type the following at the command line:

```
secedit /refreshpolicy MACHINE_POLICY
```

This will cause any changes made to domain-level policy settings to be applied to the local computer. To force a reapplication of policy to domain-level policy settings, regardless of whether there has been a change or not, type the following at the command line:

```
secedit /refreshpolicy MACHINE_POLICY /enforce
```

You can determine whether or not security was applied successfully by viewing the Application Event Log. If an error occurred during the process of applying security policy, you can get detailed information by setting the following **REG_DWORD** to **0x02**:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtensions\
{827D319E-6EAC-11D2-A4EA-00C04F79F83A}\ExtensionDebugLevel
```

When this value is set, the Security Templates will log policy-processing information in the Winlogon.log file at %windir%\Security\Logs\Winlogon.log.

**What is the Add Workstation to Domain Logon right, and how does it relate to delegating similar permissions on the directory?**

The **Add Workstation to Domain** user right is supported for applications that use earlier SAM (Security Accounts Manager) NET APIs to create computer accounts. Users that have this right are allowed to create 10 computer accounts in the Active Directory **Computers** container using these earlier APIs. When a user creates a computer account using this user right, the Domain Admins group becomes the owner of the computer object. Note that this right is *not* recognized when LDAP is used to create computer accounts.

In Windows 2000 and later, the recommended way to allow a user or group to create computer accounts is by granting that user or group the permission to **Create Computer Objects** on the desired container. This can be accomplished in GPMC. When a computer account is created using access control permissions, the actual creator of the object becomes the owner of that object.

| **Note** |
| --- |
| The create-computer-object permission should not be granted indiscriminately. Allowing users to create computers in the domain is similar to allowing users to create user accounts in the domain. Unlike Windows NT 4.0, Windows Server 2003 computer objects can be used to do network authentication and, hence, to access resources over the network. Users that have access permissions to create computer objects are also not subject to any quota restrictions. That is, they can create any number of computer accounts. |

The best security practice would be to grant only trusted users (by using a group) the permission to create computer objects. At the time the computer object is created, the creator can define which users are allowed to use that computer object to join their physical computer to the domain.

# Appendix B: Group Policy Storage

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy objects store information in two locations: a Group Policy container and a Group Policy template.

**Group Policy Container**
The Group Policy container is an Active Directory container that stores GPO properties; it includes sub-containers for computer and user Group Policy information. The Group Policy container has the following properties:

- **Version information.** This is used to ensure that the information is synchronized with the Group Policy template information. Indicates the number of changes made to the GPO.
- **Status information.** This indicates whether the GPO is enabled or disabled.
- **List of components (extensions) that have settings in the GPO.**
- **File System path.** The UNC path to the Sysvol folder.
- **Functionality version.** This is the version of the tool that created the GPO. Currently, this is version 1.

For example, the Group Policy container stores information used by the Software Installation snap-in to describe the state of the software available for installation. This data repository contains data for all applications, interfaces, and APIs that provide for application publishing and assigning.

**Group Policy Template**
Group Policy objects also store Group Policy information in a folder structure called the Group Policy template that is located in the System Volume folder of domain controllers (Sysvol) in the \Policies sub-folder. The Group Policy template is the container where Security Settings, Administrative Template-based policy settings, applications available for Software Installation, and script files are stored.

When you modify a GPO, the directory name given to the Group Policy template is the GUID of the GPO that you modified. For example, assume that you modified a GPO associated with a domain called Seattle. The resulting Group Policy template folder would be named as follows (the GUID is an example):

```
%systemroot%\sysvol\<SYSVOL>\Seattle.yourcompanyname.com\Policies\{47636445-af79-11d0-
91fe-080036644603}
```

where the second *sysvol* is shared as Sysvol. (The default location of the Sysvol folder is %systemroot%).

**Gpt.ini File**
At the root of each Group Policy template folder is a file called Gpt.ini. For local Group Policy objects, the Gpt.ini file stores information indicating the following:

- Which client-side extensions of the Group Policy Object Editor contain User or Computer data in the GPO.
- Whether the User or Computer portion is disabled.
- Version number of the Group Policy Object Editor extension that created the Group Policy object.

For the local GPO, the Gpt.ini file contains the following information:

```
[General]
gPCUserExtensionNames //Includes a list of GUIDs that tells the client
side engine which Client
                      Side Extensions have User data in the GPO.
                      The format is: [{GUID of Client Side Extension}{GUID of MMC
                      extension}{GUID of second MMC extension if appropriate}]
[repeat first
                      section as appropriate].
```

```
GPCMachineExtensionNames //Includes a list of GUIDs that tells the
client side engine which
                         Client Side Extensions have Machine data in the GPO.

Options..//Refers to GPO options such as User portion disabled or
Machine portion disabled.

GPCFunctionalityVersion //The Version number of the Group Policy
extension tool that created
                       the Group Policy object.
```

**Gpt.ini for Active Directory GPOs**
The Gpt.ini file for Active Directory GPOs contains the following entries, which are stored in Active Directory:

```
Version=0  //Version number of the Group Policy Object
DisplayName //Display name of the GPO
```

**Local Group Policy Objects**
A local Group Policy object exists on every computer, and by default it contains only security policy (that is, other types of policy settings are not configured by default). The local GPO is stored in %systemroot%\System32\GroupPolicy, and it has the following ACL permissions:

- Administrators: full control
- Operating system: full control
- User: read

**Group Policy Template Subfolders**
The Group Policy template folder contains the following subfolders:

- **User**. Includes a Registry.pol file that contains the registry settings to be applied to users. When a user logs on to a computer, this Registry.pol file is downloaded and applied to the **HKEY_CURRENT_USER** portion of the registry.

  The **User** folder may contain the following subfolders (depending on the GPO contents):

  - **Applications**. Contains the advertisement files (.aas files) used by the Windows installer. These are applied to users.
  - **Documents and Settings**. Contains the Fdeploy.ini file, which includes status information about the Folder Redirection options for the current user's special folders.
  - **Microsoft\RemoteInstall**. Contains the OSCfilter.ini file, which holds user options for operating system installation through Remote Installation Services.
  - **Microsoft\IEAK**. Contains settings for the Internet Explorer Maintenance Snap-in.
  - **Scripts\Logon**. Contains all the user logon scripts and related files for this GPO.
  - **Scripts\Logoff**. Contains all the user logoff scripts and related files for this GPO.

- **Machine**. Includes a Registry.pol file that contains the registry settings to be applied to computers. When a computer initializes, this Registry.pol file is downloaded and applied to the **HKEY_LOCAL_MACHINE** portion of the registry.

  The **Machine** folder may contain the following subfolders (depending on the GPO):

  - **Scripts\Startup**. Contains the scripts that are to run when the computer starts up.
  - **Scripts\Shutdown**. Contains the scripts that are to run when the computer shuts down.

- **Applications**. Contains the advertisement files (.aas files) used by the Windows installer. These are applied to computers.

- **Microsoft\Windows NT\Secedit**. Contains the Gpttmpl.inf file, which includes the default security configuration settings for a Windows 2000 domain controller.

- **Adm**. Contains all of the .adm files for this GPO.

The User and Machine folders are created at install time, and the other folders are created as needed when policy is set.

**Registry.pol Files**
The Administrative Templates snap-in extension of Group Policy saves information in the Group Policy template in Unicode files referred to as Registry.pol files; they are stored in the Group Policy template. These files contain the customized registry settings that you specify (by using the Group Policy Object Editor) to be applied to the Computer (**HKEY_LOCAL_MACHINE**) or User (**HKEY_CURRENT_USER**) portion of the registry.

Two Registry.pol files are created and stored in the Group Policy template, one for **Computer Configuration**, which is stored in the **\Machine** subdirectory, and one for **User Configuration**, which is stored in the **\User** subdirectory.

When you use the Administrative Templates extension of the Group Policy Object Editor to define customized registry settings, two Registry.pol files are created and stored in the Group Policy template. One Registry.pol file is for Computer Configuration-related registry settings and is stored in the **\Machine** sub-directory, and the other is for User Configuration settings and is stored in the **\User** sub-directory.

The Registry.pol file consists of a header and registry values.

The header contains version information and signature data, both **DWORD** values:

```
REGFILE_SIGNATURE 0x67655250
REGISTRY_FILE_VERSION 00000001 (increments each time the file format changes)
```

The registry values begin with an opening bracket ([) and end with a closing bracket (]):

```
[key;value;type;size;data]
```

where:

*Key* is the path to the registry key to use for the category. Do not include **HKEY_LOCAL_MACHINE** or **HKEY_CURRENT_USER** in the registry path. The location of the file determines which of these keys is used.

The following value has special meaning for this field:

- ****DeleteKeys**—a semi-colon-delimited list of values to delete.

  For example: ****DeleteKeys NoRun;NoFind**.

*Value* is the name of the registry value. The following values have special meaning for this field:

- ****DeleteValues**—a semi-colon-delimited list of values to delete. Use as a value of the associated key.

- ****Del.valuename**—deletes a single value. Use as a value of the associated key.

- ****DelVals**—deletes all values in a key. Use as a value of the associated key.

*Type* is a data type. The field can be any of the standard registry value types, for example:

- REG_DWORD

- REG_EXPAND_SZ

- REG_SZ

Note that although the file format supports all the registry data types (such as **REG_MULTI_SZ**), the Administrative Templates node does not support these registry types: **REG_BINARY, REG_MULTI_SZ.**

*Size* is the size of the data field in bytes. For example, 4.

*Data* is the raw information. For example, 4 bytes of data 0x00000001.

It is possible that the *valuename*, *type*, *data*, and *size* could be missing or 0. In this case, only the key should be created.

This pattern of [] entries continues until the end of the file.

The following special values are used for deleting keys and values:

- ****DeleteKeys** // Semi-colon-delimited list of keys to delete.

  For example: ****DeleteKeys REG_SZ NoRun;NoFind.**

- ****DeleteValues** // Semi-colon-delimited list of values to delete.

  Used as a value of the designated key.

- ****Del.valuename** // Deletes a single value name.

  Used as a value of the designated key.

- ****DelVals** // Deletes all values in a key.

  Used as a value of the designated key.

The Registry.pol file contains data to be written to the registry based on the settings specified with the Group Policy Object Editor, and the names of any scripts and their command lines (in the form of registry keys and values).

**How Registry.pol Files Are Created**
The following section outlines how to form Registry.pol files:

- When you start the Group Policy Object Editor, a temporary registry tree is created that consists of two nodes: **USER** and **MACHINE**.

- As you navigate the Administrative Templates node of the Group Policy Object Editor, .adm file nodes are displayed. The .adm files within the Group Policy Object Editor nodes are loaded dynamically when a particular node is selected, and the .adm file is then cached.

- When a policy is selected in the details pane (the right side of MMC console window), the temporary registry is queried to determine whether the selected policy already has registry values assigned to it; if it does, those values are displayed in the **Policy** dialog box.

  If the selected policy does not have a registry value assigned to it, the default value from the .adm file or from the associated MMC snap-in extension is used.

- After you modify a policy, the registry values that you specify are written to the appropriate portion of the temporary registry (either **MACHINE** or **USER**).

- When you close the Group Policy Object Editor, the temporary registry hives are exported to the Registry.pol files in the appropriate folders of the Group Policy template.

- The next time you start the Group Policy Object Editor for the same Group Policy Object for which you have previously set Group Policy settings, the registry information from the corresponding Registry.pol files is imported into the temporary registry tree. Therefore, when you view the policy settings, they reflect the current state.

# Appendix C: WMI Filtering

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How WMI Works**
The WMI filters are evaluated on the client computer after the list of potential GPOs have been determined and filtered based on Security Group membership.

When a filter is applied to a GPO, it will be evaluated on the client computer. The GPO will only be applied if the entire query results in TRUE. Note that a GPO will be processed if no WMI filter has been selected.

During RSoP, in planning mode evaluation, a simulation of the client side processing occurs, producing a result based on the specified WMI Filters evaluation to TRUE.

For example:

- If no WMI filter is selected, all WMI filters are assumed to evaluate to true; therefore, all GPOs will apply.
- If a filter specified in the RSoP wizard matches a given GPO, the GPO will be evaluated.
- In the RSoP wizard, the wizard will show all the filters based on the list of GPOs that apply to a computer or user. If a filter is removed from the filter list, it is assumed to evaluate to false and the GPO that is associated with the filter will not apply.

**Active Directory Schema additions**
A new property called "gPCWQLFilter" has been added to the properties of a GPO. It includes the namespace path and GUID that represents the WMI filter.

**Using WMI in Mixed Environments**
Clients running Windows 2000 or earlier versions do not have support for WMI filters. They may however be affected by a GPO that has specified a WMI filter. In this case, the client will process the GPO as if the filter evaluated to TRUE because it has no way of knowing if it were FALSE.

However, a mixed environment supports the schema changes required to enable WMI filters. Windows Server 2003-based servers with WMI filter-enabled GPOs will continue to handle Windows 2000 clients as they did previously.

WMI filters are only available in domains that have at least one Windows Server 2003 domain controller. In an environment consisting only of Windows 2000 domains, the WMI filter node in GPMC is not shown.

**Examples of WMI Filters**
This section illustrates some scenarios in which administrators use WMI filtering to achieve a specific goal.

**Software inventory-based targeting (Ored set)**
A company purchases a site license for a new bounds-checker tool that helps software developers write more reliable code. Because the bounds-checker only works with Visual Basic®, Visual C, and Visual C++,® the administrator wants to assign the package only on computers running any of these programs.

The administrator chooses the following filter:

Root\cimv2;Select * from Win32_Product where name = "MSIPackage1" OR name = "MSIPackage2" OR name = "MSIPackage32"

| **Note** |
| --- |
| it may be more reliable to use IdentifyingNumber |

**Software inventory-based targeting( Anded set)**
A software company discovers that the interaction of three software products causes instabilities on the system. The company develops a hot fix but only wants to install it on computers where this interaction is possible.

The administrator chooses the following filter:

root\cimv2;Select * from Win32_Product where name = "MSIpackage1"

root\cimv2;Select * from Win32_Product where name = "MSIPackage2"

root\cimv2;Select * from Win32_Product where name = "MSIPackage3"

| **Note** |
| --- |
| it may be more reliable to use IdentifyingNumber. |

**Operating system-based targeting**
An administrator wants to deploy an enterprise monitoring policy but needs to limit the target set to computers running either Windows 2000 Server or Advanced Server.

The administrator chooses the following filter:

Root\CimV2; Select * from Win32_OperatingSystem where Caption = "Microsoft Windows 2000 Advanced Server" OR Caption = "Microsoft Windows 2000 Server"

**Hardware inventory-based targeting**
An administrator wants to deploy a new connection-manager but needs to avoid wasting space on desktop computers without modems where the connection manager would be useless. An administrator can deploy the package across the enterprise with the following WMI-filter:

Root\CimV2;Select * from Win32_POTSModem

**Resource-based targeting**
To encourage field engineers and consultants to use documentation, a company wants to make Help systems available directly on users' hard disks. But because users complain that the Help files consume too much space, a manager decides to only deploy the documentation on computers that have at least 600 megabytes (MB) available.

An administrator can accomplish this with the following WMI filter:

Root\CimV2; Select * from Win32_LogicalDisk where FreeSpace > 629145600 AND Description <> "Network Connection"

**Computer-based targeting**
An administrator wants to set up a policy to encrypt all "My Documents" folders on notebook computers. The administrator determines that all the company's notebook computers are Toshiba Tecra models 800 and 810.

To set up the policy, an administrator uses the following WMI filter:

Root\CimV2; Select * from Win32_ComputerSystem where manufacturer = "Toshiba" and Model = "Tecra 800" OR Model = "Tecra 810"

**Asset tag-based targeting**
An administrator wants to set hardware inventory monitoring policy for all computers with an "asset tag" of 300,000-355555.

To set up the policy, an administrator uses the following WMI filter:

Root\Cimv2 ; Select * from Win32_SystemEnclosure where SMBIOSAssetTag > '300000' AND SMBIOSAssetTag < '355555'

**Hardware configuration-based targeting**
An administrator wants to target a policy for all computers that have a network adapter on interrupt number 11. The administrator chooses the following filter:

Root\cimv2; Associators of {win32_IRQResource.IRQNumber=11} where resultclass = Win32_NetworkAdapter

**Configuration-based targeting**
An administrator wants to avoid turning on "netmon" on computers enabled with multicasting turned on. The administrator chooses the following filter:

Select * from Win32_NetworkProtocol where SupportsMulticasting = true

**File attribute-based targeting**
An administrator only wants to disable sharing of folders on systems where at least one of "My Documents" directories are not

encrypted. The administrator chooses the following filter:

Root\cimv2 ; Select * from Win32_Directory where filename ='my documents' AND encrypted = false

**Time zone-based targeting**
An administrator needs to target a policy on all servers located on the East Coast of the United States. The administrator chooses the following filter:

Root\cimv2 ; Select * from win32_timezone where bias =-300

**Hot fix-based targeting**
An administrator only wants to apply a policy on computers that have a specific hot fix. The administrator chooses the following filter:

Root\cimv2 ; select * from Win32_QuickFixEngineering where HotFixID = 'q147222'

**Further Information**
The WMI SDK has a tool called cimstudio. This allows users to find a class, searching by name, description, property name, and so forth. Users can then experiment with queries on the class and optimize it before creating a new filter. For more information, see Windows Management Instrumentation in the Microsoft Platform SDK at http://msdn.microsoft.com/library/default.asp?url=/library/en-us/wmisdk/wmi/wmi_start_page.asp.

# Glossary (Group Policy Infrastructure)

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section presents terminology used in this document.

**Active Directory**

The Windows 2000 Server and Windows Server 2003 directory service that stores information about all objects on the computer network and makes this information easy for administrators and users to find and apply. With Active Directory, users can gain access to resources anywhere on the network with a single logon. Similarly, administrators have a single point of administration for all objects on the network, which can be viewed in a hierarchical structure.

**administrative templates (.adm files)**

Template files that provide settings pertaining to Windows 2000, Windows NT 4.0, and Windows 95, Windows 98, and Windows Millennium Edition operating system and registry structure. The .adm file specifies the registry settings that can be modified through the Group Policy Object Editor user interface. The .adm file consists of a hierarchy of categories and subcategories that together define how the options are displayed through the Group Policy Object Editor user interface. It also indicates the registry locations where changes should be made if a particular selection is made, specifies any options or restrictions (in values) that are associated with the selection, and in some cases, specifies a default value to use if a selection is activated.

**Administrative Templates snap-in extension**

A Group Policy Object Editor extension that includes all registry-based Group Policy, which you use to define settings that control the behavior and appearance of the desktop, including the operating system and applications. The Administrative Templates snap-in extension includes functionality for managing disk quotas.

**application assignment**

You can assign applications to either a user or a computer using Group Policy. When you assign applications to a computer, the application is automatically installed the next time the computer is started. When you assign applications to a user with Group Policy, the administrator can choose to either have the application installed on-demand when the user selects the application or in-full when the user next logs on:

- **On Demand.** If the application is installed on demand, the user's computer is set up with a Start menu shortcut, and the appropriate file associations are created in the registry. To the user, it looks and feels as if the application is already present. However, the application is not fully installed until the user needs the application. When the user attempts to open the application or a file associated with that application, Windows Installer checks to make sure that all the files and parameters of the application are present for the application to properly execute. If they are not present, Windows Installer retrieves and installs them from a predetermined distribution point. Once in place, the application opens.

- **Full Install.** The full-install option is useful for specific groups of users such as frequent travelers who might require all available applications to be fully installed before they travel. With full install, a user's applications are installed at logon.

**application publishing**

In Windows 2000 and Windows Server 2003, you can use the Software Installation snap-in extension of the Group Policy Object Editor to *publish* applications to users. Published applications are those that the administrator makes available for on-demand use.

Published applications have no presence on the users' computers. That is, no shortcuts or Start menu references to the application are present on the desktop. A published application is advertised to Active Directory. The advertised attributes are used to locate the application and all the information required for installing it. After the application is advertised in Active Directory, users can activate it by document association, just as an assigned application. Users can also set up the program using the Add or Remove Programs Control Panel tool on their desktop.

**.cab file**

A .cab file contains one or more files, all of which are downloaded together in a single compressed cabinet file. Included in the cabinet is an .inf file that provides further installation information. The .inf file may refer to files in the .cab and to files at other URLs.

**discretionary access control list (DACL)**

A part of the security descriptor that specifies the groups or users that can access an object, as well as the types of access (permissions) granted to those groups or users. *See also* security descriptor.

**disk quotas**

Within the Administrative Templates node of the Group Policy Object Editor are policy options for managing disk quotas, which administrators can use to monitor and limit disk space use for NTFS volumes formatted as NTFS version 5.0. After you enable disk quotas, you can set options for disk quota limits and warnings.

**domain**

A grouping of servers and other network objects under a single name. Domains provide the following benefits:

- You can group objects into domains to help reflect your company's organization in your computer network.
- Each domain stres only the information about the objects located in that domain. By partitioning the directory information this way, Active Directory scales up to as many objects as you need to store information about on your network.
- The administraor of a domain has absolute rights to set policy settings within that domain only.

**domain trees**

You can combine multiple domains into structures called domain trees. The first domain in a tree is called the root of the tree, and additional domains in the same tree are called child domains. A domain immediately above another domain in the same tree is referred to as the parent of the child domain.

All domains within a single domain tree share a hierarchical naming structure. Domains that share a common root share a contiguous namespace. Domains in a tree are joined together through two-way, transitive trust relationships. These trust relationships are two-way and transitive, therefore, a domain joining a tree immediately has trust relationships established with every domain in the tree.

**Folder Redirection snap-in extension**

A Group Policy Object Editor extension that you use to place the Windows 2000 or Windows Server 2003 special folders in network locations other than their default location (%systemroot%\Documents and Settings\%userprofile%) on the local computer.

**globally unique identifier (GUID)**

A 128-bit integer that identifies a particular object class and interface. GUIDs are virtually guaranteed to be unique. A GUID can be generated using either the uuidgen.exe utility from the Platform Software Development Kit, or the GUIDgen tool included in the Microsoft Visual C++® development system. For more information about GUIDs, see the OLE Programmer's Reference, Volume One; the Platform Software Development Kit documentation; and Inside OLE, 2d ed. by Kraig Brockschmidt, Redmond, Wash.: Microsoft Press, 1995.

**Group Policy**

A component used in Windows 2000 and Windows Server 2003 to define options for managed desktop configurations for groups of users and computers. To specify Group Policy options, you use GPMC in conjunction with the Group Policy Object Editor.

**Group Policy engine**

The part of Group Policy that runs in the Winlogon process.

**Group Policy Management Console (GPMC)**

An MMC console to view and edit Group Policy properties, generate reports, copy, import, backup, restore, and to select GPOs for editing. GPMC lets administrators manage Group Policy for multiple domains and sites within a given forest, all in a simplified user interface with drag-and-drop support. Operations are fully scriptable, which lets administrators customize and automate management.

**Group Policy object**

The Group Policy settings that you create by using the Group Policy Object Editor are contained in a GPO, which is in turn associated with selected Active Directory containers: sites, domains, and organizational units (organizational units).

**Group Policy Object Editor**

To edit a specific desktop configuration for a particular group of users and computers, you use the Group Policy Object Editor, also known previously as the Group Policy snap-in, Group Policy Object Editor, or GPedit.

You can specify Group Policy settings for the following:

- Registry-based policy settings—Includes Group Policy for the Windows 2000 and Windows Server 2003 operating systems and their components and for applications. To manage these settings, use the Administrative Templates node of the Group Policy Object Editor.
- Security settings—Includes options for local computer, domain, and network security settings.
- Software Installation and Maintenance options—Used to centrally manage application installation, updates, and removal.
- Script options—Includes scripts for computer startup and shutdown and user logon and logoff.
- Folder Redirection options—Allows administrators to redirect users' special folders to the network.
- Internet Explorer Maintenance—Used to manage and customize Internet Explorer on Windows 2000- and Windows Server 2003-based computers.
- Remote Installation Services—Used to control the behavior of the Remote Operating System Installation feature as displayed to client computers

**Group Policy Modeling**

This is a simulation of what would happen under circumstances specified by an administrator. Group Policy Modeling requires that you have at least one domain controller running Windows Server 2003 because this simulation is performed by a service running on a domain controller that is running Windows Server 2003. With Group Policy Modeling, you can either simulate the RSoP data that would be applied for an existing configuration, or you can perform "what-if" analyses by simulating hypothetical changes to your directory environment and then calculating the RSoP for that hypothetical configuration. For example, you can simulate changes to security group membership, or changes to the location of the user or computer object in Active Directory. Outside of GPMC, Group Policy Modeling is referred to as RSoP - planning mode.

**Group Policy Results**

This represents the actual policy data that is applied to a given computer and user. It is obtained by querying the target computer and retrieving the RSoP data that was applied to that computer. The Group Policy Results capability is provided by the client operating system and requires Windows XP, Windows Server 2003 or later. Outside of GPMC, Group Policy Results is referred to as RSoP - logging mode.

**IntelliMirror**

IntelliMirror refers to the ability to provide users with consistent access to their applications, application settings, roaming user profiles, and user data, from any managed computer—even when they are disconnected from the network. IntelliMirror is delivered via a set of Windows features that enable IT administrators to implement standard computing environments for groups of users and computers.

IntelliMirror can significantly boost user productivity and satisfaction by doing the following:

- Allowing users to continue working efficiently in intermittently connected or disconnected scenarios by enabling uninterrupted access to user and configuration data under these conditions.
- Delivering a consistent computing environment to users from any computer when their desktop or laptop computer is unavailable or in scenarios where users are not assigned a specific computer.
- Minimizing data loss by enabling centralized backup of user data and configuration files by the IT organization.
- Minimizing user downtime by enabling automated installation and repair of applications.
- Implementing IntelliMirror also boosts administrator efficiency and reduces IT costs by doing the following:

- Eliminating the need to manually configure user settings, install applications, or transfer user files to provide users access to their computing environments on any computer.
- Enabling scenarios where users don't have an assigned computer but log in to any available computer in a pool of computers. This helps reduce hardware and administration costs.
- Easing the IT task of implementing centralized backup of user files while satisfying need for these files to be available on the user's computer.
- Reducing support costs by using Windows Installer to automatically repair broken application installations.
- IntelliMirror is implemented by means of a set of Windows features, including Active Directory, Group Policy, Software Installation, Windows Installer, Folder Redirection, Offline Folders, and Roaming User Profiles.

**Internet Explorer Maintenance**

A Group Policy extension snap-in that includes policy settings to manage the following: Browser User Interface, Connection Settings, Custom URLs, Security, and Program Associations.

**Microsoft Management Console (MMC)**

A common console framework for system-management applications. The primary goal of the Microsoft Management Console is to support simplified administration and lower cost of ownership through tool integration, task orientation, support for task delegation, and overall interface simplification. MMC console hosts the administrative tools (these are called MMC snap-ins); the console itself provides no management functionality.

**Migration table**

A migration table is a file that maps references to users, groups, computers, and UNC paths in the source GPO to new values in the destination GPO. A migration table consists of one or more mapping entries. Each mapping entry consists of a type, source reference, and destination reference. If you specify a migration table when performing an import or copy, each reference to the source entry will be replaced with the destination entry when writing the settings into the destination GPO. Migration tables store the mapping information as XML, and have their own file name extension, .migtable. You can create migration tables using the Migration Table Editor (MTE). The MTE is a convenient tool for viewing and editing migration tables without having to work in, or be familiar with, XML. The MTE is associated with the .migtable extension so that when you double click a migration table, it opens in the MTE. The MTE is installed with GPMC.

**MMC snap-in**

Tools that extend MMC console and provide administrative functionality. A snap-in functions independently from other snap-ins.

**MMC extension snap-in**

A tool that enhances the functionality of a parent snap-in. An extension depends on a parent snap-in for contextual data.

**organizational unit (organizational unit)**

A type of directory object contained within domains. organizational units are logical containers into which you can place users, groups, computers, and even other organizational units.

**registry**

A database in which Windows NT internal configuration information and computer- and user-specific settings are stored.

**registry hive**

A section of the registry that is saved as a file. The registry subtree is divided into hives (named for their resemblance to the cellular structure of a beehive). A hive is a discrete body of keys, subkeys, and values.

**Remote Installation Services**

A component that administrators can use to remotely install a local copy of the Windows 2000 Professional or Windows XP Professional on supported computers throughout their organization. Administrators can deploy a new version of an operating system upgrade to large numbers of clients at one time from a centralized location.

Administrators can use Group Policy to specify the client installation options that groups of users can access. These options are determined by the specific Remote operating system Installation Group Policy settings that administrators define for the site, domain, or organizational unit to which the users belong, in conjunction with the specific security group or user account.

**Resultant Set of Policy (RSoP)**

RSoP allows administrators to see the effect of Group Policy on a targeted user or computer. RSoP is an infrastructure leveraged by GPMC to enable Group Policy Results and Group Policy Modeling. In Group Policy Results, administrators assess what has applied to a particular target. In Group Policy Modeling, administrators can see how policy settings would be applied to a target and then examine the results before deploying a change to Group Policy.

**Roaming user profile**

A copy of the local user profile stored on a server share. This profile is downloaded every time that a user logs on to any computer on the network, and any changes made to a roaming user profile are synchronized with the server copy upon logoff. See also user profile.

**schema**

The formal definition of all object classes, and the attributes that make up those object classes, that can be stored in the directory. Active Directory includes a default schema, which defines many object classes, such as users, groups, computers, domains, organizational units, and security policy settings. The Active Directory schema is dynamically extensible; this means that you can modify the schema by defining new object types and their attributes and by defining new attributes for existing objects. You can do this either programmatically with the Schema Manager snap-in tool included with Windows NT Server.

**scripts**

Batch files (.bat) or executable (.exe) files that run when a computer starts up or shuts down or when a user logs on or off at any type of workstation on the network. Windows 2000 and Windows Server 2003 support Windows Script Host Visual Basic Scripting Edition (VBScript) and Jscript, while continuing to support MS-DOS command scripts and executable files.

**security descriptor**

A set of access-control information attached to every container and object on the network. A security descriptor controls the type of access allowed to users and groups. Administrators assign security descriptors to objects stored in Active Directory in order to control access to resources or objects on the network.

A security descriptor lists the users and groups that are granted access to an object (a file, printer, or service, for example), and the specific permissions assigned to those users and groups. *See also* discretionary access control list *and* system access control list.

**Security Settings extension snap-in**

A Group Policy extension snap-in that you use to define security configuration for computers within a GPO. A security configuration consists of settings applied to each security area supported for Windows 2000 or Windows XP Professional and Windows 2000 Server and Windows Server 2003. This configuration is included within a GPO.

**site**

In Windows 2000 and Windows Server 2003 you register your network's physical topology by defining sites. A site is defined as one or more IP subnets. Windows 2000 and Windows Server 2003 uses site information to direct requests from one computer to be fulfilled by another computer at the same site. For example, when a workstation logs on, Active Directory uses the TCP/IP address of the workstation, along with the site information you have entered, to locate a domain controller on the local site. This local controller is used to service the workstation's requests.

**Scripts extension snap-in**

A Group Policy extension snap-in that you use to assign scripts to run at computer startup or shutdown or upon user logon or logoff.

**Software Installation extension snap-in**

A Group Policy extension snap-in that you use to centrally manage software distribution in your organization.

**system access control list (SACL)**

Part of a security descriptor that specifies which user accounts or groups to audit when accessing an object, the access events to be audited for each group or user, and a Success or Failure attribute for each access event, based on the permissions granted in the object's DACL.

**tattooing**

This refers to a registry setting that is set using Windows NT 4.0 System Policies, the setting persists until the specified policy is reversed or the user edits the registry.

**total cost of ownership (TCO)**

Refers to the administrative costs associated with computer hardware and software purchases, deployment and configuration, hardware and software updates, training, maintenance, and technical support.

**user profile**

A user profile describes the desktop computing configuration for a specific user, including the user's environment and preference settings. A profile is created the first time that a user logs on to a computer running Windows Server 2003, Windows XP, Windows 2000, or Windows NT Workstation. A user profile is a group of settings and files that defines the environment that the system loads when a user logs on. It includes all the user-specific configuration settings, such as program items, screen colors, network connections, printer connections, mouse settings, and window size and position. Profiles are not user policies and the user has a profile even if you don't use Group Policy.

**Windows Installer packages (.msi files)**

Packages that contain all the information necessary to describe to the Windows Installer how to set up an application in every conceivable situation: various platforms, different sets of previously installed products, earlier versions of a product, and numerous default installation locations. The Software Installation extension snap-in to the Group Policy Object Editor uses .msi packages.

**Windows Management Instrumentation (WMI)**

A management infrastructure that supports monitoring and controlling system resources through a common set of interfaces and provides a logically organized, consistent model of Windows operation, configuration, and status. WMI Filtering in Windows Server 2003 allows you to create queries based on this data. These queries (also called WMI filters) determine which users and computers receive all of the policy configured in the GPO where you create the filter.

# Related Links (Group Policy Infrastructure)

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

- Microsoft.com Group Policy Home Page at http://www.microsoft.com/grouppolicy. Provides an entry point for Group Policy documentation on the Web. Includes links to documentation, knowledge base articles, support information, and newsgroups.

- Windows Server 2003 Deployment Kit, Designing a Managed Environment. at http://www.microsoft.com/grouppolicy. Available from the Microsoft Group Policy home page, this book describes the technologies in Windows Server 2003 associated with deployment of a managed environment. Has significant coverage of Group Policy and related IntelliMirror technologies. Includes planning, designing, and implementation guidance.

- Troubleshooting Windows Server 2003 Group Policy. at http://www.microsoft.com/grouppolicy. Available from the Microsoft Group Policy home page, this white paper provides a structured guide to troubleshooting Group Policy operations. Covers Group Policy processing itself, dependent technologies and relevant troubleshooting tools.

- Group Policy Administration using the Group Policy Management Console. at http://go.microsoft.com/fwlink/?LinkID=14320. This white paper provides technical details of functionality in GPMC.

- Migrating GPOs Across Domains with GPMC. at http://go.microsoft.com/fwlink/?LinkID=14321. This white paper explains how to migrate GPOs from one domain to another using GPMC.

- **Group Policy Management Console Software Development Kit (SDK)**. Provides information about how to use the COM interfaces of Group Policy Management, which support scripting many of the operations supported by Group Policy Management Console. This is available when you install GPMC. See "gpmc.chm" in the **%programfiles%\gpmc\scripts** directory. The SDK is only available in English.

**Feedback on this Paper**
If you have any comments about this paper, contact mailto:gpdocs@microsoft.com.

**Newsgroups About Group Policy**
If you have a question about Group Policy, you can post to the newsgroup "microsoft.public.windows.group_policy."

# Multiple Forest Considerations in Windows 2000 and Windows Server 2003

Updated: April 8, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The goal of this paper is to present an overview of the effort and infrastructure that might be required to enable varying levels of collaboration among multiple forests that have domain controllers that are running Microsoft® Windows® 2000 Server or Microsoft® Windows Server™ 2003 operating systems. This paper enumerates the scenarios in which a multiforest environment might be necessary or desirable and analyzes the consequences of such an environment on the total cost of ownership of the enterprise.

**In This White Paper**

Who Should Read This Paper

Introduction (Multiple Forest Considerations in Windows 2000 and Windows Server 2003)

Multiforest Deployments

Additional Costs of Deploying Multiple Forests

Additional Configuration for Functionality Across Forests (Multiple Forest Considerations in Windows 2000 and Windows Server 2003)

Multiforest Feature Implementation and Cost Summary (Multiple Forest Considerations in Windows 2000 and Windows Server 2003)

# Who Should Read This Paper

Updated: April 8, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This paper is intended for architects and project managers who are designing an Active Directory deployment and have identified a potential need for multiple Active Directory forests, or who have an existing single-forest environment that must change to accommodate multiple forests. They can use this paper to help them understand the impact of deploying multiple Active Directory forests on total cost of ownership, and draw up their plans accordingly.

This paper is not intended to provide recommendations for when to deploy multiple forests or prescriptive guidelines for how to deploy multiple forests. The goal is to present a high-level view of the configuration changes that are required to enable collaboration between forests. Microsoft intends to issue a deployment guide that provides detailed instructions for implementing these configuration changes.

For more information about the conditions that warrant deploying multiple forests, see "Design Considerations for Delegation of Administration in Active Directory." You can view this document on the Microsoft® TechNet Web site at http://www.microsoft.com/technet/prodtechnol/windows2000serv/technologies/activedirectory/plan/addeladm.mspx.

For general Active Directory deployment planning information, see "Best Practice Active Directory Design for Managing Windows Networks." You can view this document on the Microsoft TechNet Web site at http://www.microsoft.com/technet/prodtechnol/windows2000serv/technologies/activedirectory/plan/bpaddsgn.mspx.

For information about deploying Microsoft® Exchange 2000 Server with Active Directory in a Windows 2000 Server environment, see "Best Practice Active Directory Design for Exchange 2000." You can download this document from the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=10517.

# Introduction (Multiple Forest Considerations in Windows 2000 and Windows Server 2003)

Updated: April 8, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory infrastructure necessarily provides a certain amount of forestwide oversight that enables administrators in the forest root domain to manage schema data, configuration data, and domain controllers. The control granted to these administrators carries with it the potential for misuse by malicious administrators. In addition, where autonomy or isolation is required, the level of security afforded by a single domain might not be adequate. In particular, large organizations might have organizational, legal, or operational requirements that warrant dividing control between separate forests.

The need to deploy multiple forests might be recognized and accommodated during initial deployment planning, or might occur as an unplanned requirement of an organizational change after deployment. Both cases have important ramifications, as follows:

- Planned multiforest deployment: During planning for multiple forests, be sure that you carefully weigh the advantages against the added costs of implementation.
- Unplanned multiforest deployment: If business decisions result in a merger, acquisition, or divestiture that creates a need to add one or more forests to the deployed environment, be prepared for a gradual process of implementing multiple forests and enabling collaboration across forests.

The division of a Windows deployment into multiple Active Directory forests must be undertaken with the understanding that some of the functionality that is available by default in a single Active Directory forest requires additional configuration in a multiforest deployment.

| **Note** |
| --- |
| Some single-forest functionality is not available in multiforest deployments for applications using Active Directory. For example, certain mailbox-related features are not available in multiforest deployments that implement Microsoft Exchange 2000 Server, and require a workaround solution. For information about these features and solutions, see "Features that are Unavailable Across Forests" later in this paper. |

To coordinate the data and services of multiple forests to serve a single organization, the following functionality requires additional configuration:

- Domain Name System (DNS) name resolution across forest boundaries.
- Microsoft Exchange 2000 Server and later functionality, including synchronization of address list information across multiple forests and synchronization of free and busy calendar information.
- Access to resources across the forests.
- Infrastructure data synchronization across the forest boundaries.
- Application data synchronization across multiple forests.

| **Note** |
| --- |
| This paper does not provide step-by-step configuration instructions. Microsoft intends to issue a deployment guide to address deployment information. |

In this paper, the rationale for each additional configuration is explained, as well as an overview of the planned and unplanned conditions for, and effects of, deploying multiple forests.

**What This Paper Does Not Cover**

It is not a goal of this paper to provide recommendations for when to deploy multiple Active Directory forests rather than a single Active Directory forest. If you are planning to deploy multiple forests, it is assumed that you have analyzed operational, organizational, and legal requirements for Active Directory and have determined that one or more requirements are consistent with creating multiple Active Directory forests.

Preliminary reading to examine these conditions includes the following:

- For a detailed explanation of the conditions that warrant partitioning a Windows network between multiple forests to achieve service autonomy, service isolation, or data isolation, see "Design Considerations for Delegation of Administration in Active Directory." You can view this document on the Microsoft TechNet Web site at http://www.microsoft.com/technet/prodtechnol/windows2000serv/technologies/activedirectory/plan/addeladm.mspx.

- For general Active Directory deployment planning information, see "Best Practice Active Directory Design for Managing Windows Networks." You can view this document on the Microsoft TechNet Web site at http://www.microsoft.com/technet/prodtechnol/windows2000serv/technologies/activedirectory/plan/bpaddsgn.mspx.

- In addition, this paper does not provide:

- Prescriptive step-by-step instructions for how to configure features in multiple forests. Microsoft intends to issue a deployment guide with instructions for how to implement configuration changes.

- Recommendations for whether to merge or "bridge" multiple forests when companies merge, or instructions for how to do so.

- Instructions for how to spin off a separate forest in preparation for a divestiture.

# Multiforest Deployments

Updated: April 8, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The decision to deploy multiple forests is often associated with the need to manage directory services independently (autonomy) or to manage data or services without threat of interference (isolation). In Microsoft® Windows NT® 4.0 operating system deployments, the domain is the ultimate component and is discrete in its management and its boundaries. With the introduction of Active Directory® in Microsoft Windows® 2000, a new category of *forest* was introduced, and the forest has service and data implications for all domains.

Availability of service and data isolation or autonomy in an Active Directory domain depends not only on the administrator of that domain and forest administrators, but also on other domain administrators within a forest.

When designing your forest, keep in mind that administrative independence has a cost and that you must carefully consider the tradeoff between autonomy or isolation versus interoperation and collaboration.

**Management Autonomy and Isolation**

Management of Active Directory can be divided into two general categories:

- Service management: Configuration and delivery of the directory service.
- Data management: Creation and management of directory contents (account management, organizational unit administration).

Active Directory provides tools for the delegation of control over these areas of management, and delegation can occur at several levels, including forest, domain, and organizational units. At each of these levels and for each of their respective service or data categories, the need for independence can vary widely, depending on the nature of the organization.

In most cases, an appropriate and adequate level of control can be delegated within a single forest by using domains and organizational units to manage data and by the ability to trust service managers at the forest root level. However, the following requirements for independence can exceed the limitations of single-forest delegation of control:

- Service autonomy: Provides the ability to independently manage and manipulate the schema and configuration containers. The need for this level of control is usually driven by organizational structure or operational requirements. For example, one division of a company wants to install directory-enabled applications that extend the schema without depending on approval by the central Information Technology (IT) department.
- Service isolation: Provides the assurance that no administrator outside the organization can interfere with the operation of the directory service. The need for this level of control is usually driven by legal or operational requirements. For example, suppose a hosting company needs to place domain controllers on a customer's premises. If a breach of domain controller security can affect service delivery in the rest of the forest, that customer's operations can be separated into its own forest to better protect other clients of the directory service.
- Data isolation: Provides the assurance that no service administrators outside of a limited scope of administrators can control or view any subset of data on a domain controller or on member computers joined to the forest. The need for this level of control is usually based on a legal requirement. For example, a financial institution might be required by law to limit access to data that belongs to clients in a particular jurisdiction to the users, computers, and administrators in that jurisdiction.

**Unplanned Deployments**

The need for multiforest deployment is not always planned in advance. A change from a single-forest deployment to a multiforest deployment might be required to accommodate the following business decisions:

- Mergers and acquisitions: Assuming the merged or acquired organizations operate one or more existing forests, you might need to enable collaboration between users and resources from the different parts of the new and original organizations.
- Divestitures: When a decision is made to divest your company of certain portions of the organization, you can create a

forest to accommodate the users and resources in the departments that are being transferred in advance of the actual organizational spin-off and enable collaboration between the original forest or forests and the spun-off forest.

- Grass-roots deployments: In the event that one or more divisions within your organization deploy their own forest without consulting the central IT department, consider the following courses of action:

    - If no central forest has yet been deployed by the IT department, use the existing forest as the main forest for the entire organization, provided the existing forest satisfies your Active Directory design requirements.

    - If an IT forest already exists, merge the IT forest and the grass-roots forest into a single, central IT forest.

    - If an IT forest already exists, implement a multiforest deployment.

**Levels of Collaboration**
The level of collaboration determines which functionality needs to be enabled across the forests and, as a result, determines part of the additional cost associated with deploying multiple forests. As the first step in evaluating the cost of enabling functionality across forests, assess the level of collaboration that your business requires. Use the following example scenarios to help determine the functionality you want to enable on the basis of the needed levels of collaboration.

**Scenario 1: Organizations Share Only Address List Information**
For organizations that require operational and organizational independence, collaboration needs are minimal. For example, a holding company for banks manages businesses that remain competitors. Because the businesses compete, they do not merge to a single forest, modify the messaging system, nor share any resources. Employees of these banks, however, need the ability to look up the phone numbers, office locations, and perhaps managerial information of personnel in the other divisions of the company, and to communicate by e-mail.

These organizations have the following operating conditions and requirements:

- No resource sharing.

- No authentication of users across forests.

- Independent Exchange Server organizations, but commonly viewable global address list between the forests.

In this scenario, the only collaboration requirement is the need to view address information of the users in different forests. The only configuration that is required for this scenario is to enable Exchange address list synchronization between the forests. The cost of this added functionality is that of deploying and operating a directory data synchronization product such as Microsoft Identity Integration Server 2003. An integrated address list synchronization solution works in conjunction with Microsoft Identity Integration Server 2003 to implement address list synchronization.

For more information about implementing address list synchronization, see “Synchronizing Exchange Data Across Forests” later in this paper.

**Scenario 2: Organizations Share Schedule Information**
In this scenario, in addition to the ability to view address list information, users in the separate forests need to schedule meetings with each other and to use other workflow applications that use Exchange public folders. An example of an organization that requires this type of collaboration is a holding company that includes multiple companies (for example, pharmaceutical companies) that compete with each other. For this reason, the companies require separate forests, messaging system, and prevention of resource access from outside of the forests. At the same time, however, they must conduct joint research projects. These joint research projects involve experts from different organizations working together in virtual teams. To enable these experts to schedule meetings and conference calls with each other, the organization decides to synchronize public folders, which also include free and busy calendar information, to make it viewable outside of the forest boundaries.

For businesses that have separate Exchange organizations, this level of collaboration has the following requirements:

- No resource sharing.

- No authentication of users across forests.

- Address list synchronization as described in Scenario 1.

- Free and busy information synchronization (public folder synchronization).

Synchronizing free and busy information involves synchronizing Exchange public folders across forests. The information that is required for scheduling meetings is stored only in public folders on Exchange servers, and not in Active Directory. The solution for enabling this functionality is to use a public folder synchronization tool to synchronize the contents of the Exchange system public folders between forests.

For more information about synchronizing public folders, see "Synchronizing Exchange Data Across Forests" later in this paper.

**Scenario 3: Organizations Share Resources**
In this scenario, the organizations need to be able to assign access to resources in one forest to groups of users in another forest. In addition, users need to be able to look up other users throughout the organization by using the global address list, and to view free and busy calendar information to enable them to schedule meetings with people in different forests.

This scenario is commonly used in the following real-life cases:

1. Two organizations merge and their management decides to eventually merge directory or messaging infrastructures, or both. Because merging directory and messaging infrastructures is a time-consuming process, the IT department is tasked with enabling single sign-on capability for resource access outside of a forest and to simulate the experience of single-forest messaging, in which address lists are viewable and free and busy information is available for all users in both forests.

2. Two organizations merge, but unfortunately domain controllers of one organization are located in less secure location(s) than domain controllers of the other organization. The IT department of the less secure organization does not accept the level of physical security of domain controllers in the more secure organization. For this reason, these administrators decide to not merge their directory and messaging infrastructures. At the same time, businesses of these two organizations need close collaboration, which requires single sign-on capability for resource access outside of a forest and the same messaging experience that is available in a single forest.

3. An enterprise has deployed a test selfhost forest where new configurations, operations, and applications are deployed before they are deployed in the main production forest. All employees are equally trusted and need access to various resources in both forests, and therefore users must have access to the same resources, regardless of which forest contains a user's account. Although deploying a single messaging system is feasible in an environment where employees are equally trusted, this organization chooses to deploy Exchange Server in each forest to completely simulate the main production forest environment. In this way, they can identify any disruptions to messaging before any new configurations, operations, or applications are deployed in the main production forest.

This level of collaboration includes:

- Bridged DNS infrastructures.

- Separate Exchange organizations, but a commonly viewable global address list.

- Free and busy information synchronization between Exchange organizations in separate forests.

- Resource sharing among users across both forests:

    - Controlled sharing or full sharing.

    - Security groups can have members from both forests.

    - Access control lists (ACLs) can list security principals from both forests.

- Single sign-on between forests: User logon experience is the same for access to resources in any forest.

Solutions for this scenario involve configuring DNS to locate servers in the different forests. Enabling resource access requires trust relationships between domains in different forests as well as appropriate groups and permissions for resource sharing.

For the forest functional level of Windows Server 2003, forest trust relationships are available that enable bidirectional, transitive trust relationships between all domains in the respective forests. For Windows 2000 forests or Windows Server 2003 forests that have a functional level other than Windows Server 2003, you must manually create explicit trust relationships between specific domains to enable resource sharing.

| ⚠**Caution** |
| --- |
| External trust relationships between Windows 2000 domains in different forests are created with security identifier (SID) filter quarantining turned off by default. If you plan to implement external trust relationships, see article Q289243, "Forged SID Could Result in Elevated Privileges in Windows 2000," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://go.microsoft.com/fwlink/?LinkId=291. |

For more information about bridging DNS, see “Enabling DNS Name Resolution Across Different Forests” later in this paper. For more information about enabling resource sharing and single sign-on between forests, see “Enabling Resource Access Across Forests” later in this paper. For more information about address list synchronization, see “Synchronizing Exchange Data Across Forests” later in this paper.

**Scenario 4: Organizations are Completely Integrated and Share a Single Exchange Organization**

To reduce the cost of providing a messaging solution, administrators might agree to deploy a single Exchange 2000 organization to serve an entire enterprise, in which case their messaging infrastructure does not need to be synchronized. This solution offers tighter integration because a single Exchange Server organization can span multiple forests without additional configuration.

An example of a tightly integrated multiforest organization is a factory that has assembly lines that must never be stopped because the consequences of stopping the assembly line could be devastating for the organization. Considering the mission-critical role of the assembly line, and because the functionality of the assembly line depends on an Active Directory-integrated application, the IT team deploys a separate assembly-line forest that has no external dependencies. For example, this forest might have no dependency on network connectivity, or perhaps it cannot accept the level of service that is provided for the corporate forest in terms of the time required for the IT department to respond to failures.

Some users have accounts in the assembly-line forest. These users roam throughout the company on a regular basis and they use laptops that are joined to the assembly-line forest. They often need access to file shares on the Distributed File System (DFS) servers in the corporate forest and to use printers that are located in the corporate forest. In addition to the DNS integration and resource sharing that is described in Scenario 3, the following functionalities can also be configured to accommodate the needs of these roaming users:

- Site and subnet synchronization.
- Printer information.

For more information about site and subnet information, see “Synchronizing Sites and Subnets” later in this paper. For more information about making printer information available across forests, see “Locating Printers Across Forests” later in this paper.

# Additional Costs of Deploying Multiple Forests

Updated: April 8, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The most significant disadvantage to deploying multiple forests is the added costs involved with providing additional teams and training for multiple deployments, as well as additional configuration. When you deploy multiple forests as opposed to a single forest, additional costs include:

- Design: Different teams need to design their individual forest, which requires the cost of educating additional IT professionals and the cost of their time invested in the design process itself.
- Implementation: Different teams need to deploy, monitor, and operate a separate forest, which requires the cost of training more IT professionals and the cost of their time invested in these activities.
- Collaboration across forests: Extra cost might be associated with collaborating on the configuration and monitoring of the additional features that are not required for a single-forest deployment, including:
    - Collaboration between forest administrators and possibly domain administrators.
    - Additional software and hardware.
- Training for additional configuration requirements.
- The actual configuration of the specific features that are required to implement single-forest functionality across forests.

If deploying multiple forests is an option rather than a requirement, then prior to deployment you must be sure that the advantages of service autonomy, service isolation, or data isolation outweigh these additional costs.

# Additional Configuration for Functionality Across Forests (Multiple Forest Considerations in Windows 2000 and Windows Server 2003)

Updated: April 8, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although most of the operations performed by users and administrators in a multiforest deployment rely on functionality within their own forests, some operations require additional configuration to extend the underlying functionality across the forest boundaries.

In addition to administrative collaboration that is involved in cooperatively operating multiple forests, the following features require additional configuration to enable collaboration across forests:

- DNS name resolution: Enable DNS name resolution between forests to provide domain controller and resource location functionality.
- Access to resources:
    - Establish trust relationships between user domains and resource domains.
    - Configure the access control lists (ACLs) of resources to allow access to appropriate groups from different forests, and create new groups to accommodate forest roles across forests.
- Synchronization of sites and subnets: Enable each forest to recognize the entire scope of the network. To allow clients in one forest to locate the nearest domain controller in another forest and enable access to DFS file shares with minimum network traffic, configure the same sites and subnets in all forests.
- Storage and retrieval of application data: Provide computers of all forests with the ability to access application data in any forest.
- Printer location: Synchronize printer information.
- Exchange Server integration: Enable address list synchronization and free and busy calendar information.

The topics in this section contain discussion of the essentials of each configuration at a high level but do not provide specific configuration instructions. Microsoft intends to issue a deployment guide with instructions for how to enable the functionality described in this section.

**Enabling DNS Name Resolution Across Different Forests**
Successful Active Directory deployment requires DNS for name resolution. The ability to resolve an IP address from the DNS name of a computer is necessary for the location of domain controllers and resources by users and computers.

**Issue**
With multiple forests, name resolution is required to function between forests so that users and computers in one forest can locate resources and domain controllers in the other forest.

Best practices for designing the DNS namespace and providing DNS support for deployment of a single Active Directory forest are provided in the "Best Practice Active Directory Design for Managing Windows Networks." You can view this document on the Microsoft TechNet Web site at http://go.microsoft.com/fwlink/?LinkId=10478. In the process of designing your multiforest Active Directory deployment, follow these recommendations for each individual forest. After you have completed the DNS design for each individual Active Directory forest, you must analyze the requirements to enable name resolution across the forest boundaries.

The ability to resolve the names from the namespace of a different forest is required by the following activities:

- User or computer needs to locate a resource (for example, an Exchange server, file share, or service) that is hosted on a

computer that is joined to a different forest.

- User or computer needs to locate a domain controller in a different forest for the purpose of authenticating to a resource in that forest.

In most scenarios, you will need to ensure name resolution across different forests.

**Solution**
The need for additional DNS configuration to enable name resolution across forests depends on your current DNS infrastructure and on any DNS design that might be planned for new Active Directory forests. The most common DNS scenarios are described below and relevant configuration requirements are enumerated for each.

If a user in a fictitious organization A (Org A) needs to locate resources in organization B (Org B), or vice versa, different configuration changes apply for each of the following scenarios:

- In Figure 1, Org A has a root DNS server that hosts the root (".") zone. In  In Figure this scenario, a delegation is set up on the root DNS server that hosts the root zone in Org A. This delegation consists of a host (A) resource record and a name server (NS) resource record for the DNS servers that host the highest-level DNS domain in Org B. The effect of these resource records is that queries originating in Org A for names within the namespace of Org B are referred to the DNS servers in Org B. Figure 1 shows the delegation from Org A to Org B.




**Figure 1 Delegation from DNS server hosting root zone in Org A to enable resolution of names in Org B**

- In Figure 2, Org A has no root zone, but DNS servers running a member of the Windows Server 2003 operating systems are configured to forward queries up the namespace, eventually to the Internet. In this scenario, conditional forwarding must be configured on a DNS server in Org A to forward queries ending with a particular suffix (in this case, the suffix of Org B) to DNS servers in Org B. DNS servers in Org A can forward to any level in the namespace of Org B. Figure 2 shows conditional forwarding from Org A to Org B.




**Figure 2 Conditional forwarding from Org A to Org B**

| Note |
| --- |
| Conditional forwarding is available only in Windows Server 2003. |

- In Figure 3, Org B has no root zone and DNS servers in Org B are running a member of the Windows 2000 Server family. To enable DNS servers in Org B to respond to queries for names within the namespace of Org A, DNS servers hosting high-level zones in the namespace hierarchy can be configured to host secondary zones that contain the highest-level zone data of Org A servers. As shown in Figure 3, the Org B servers that host these secondary zones must perform zone transfers from Org A DNS servers to ensure that the secondary zones for names in Org A are updated.




**Figure 3 Secondary zone for Org A in Org B**

**Enabling Resource Access Across Forests**

Access to resources in different forests requires appropriate trust relationships, user and group management, and access control in the trusting domains. Resource sharing between domains is discussed in "Designing the Active Directory Structure" in the *Deployment Planning Guide of the Microsoft® Windows® 2000 Server Resource Kit*. The model of resource access and management for multiple Windows Server 2003 forests is very similar, with a few basic differences.

**Issue**

When users in Forest A need access to resources in Forest B, the resource administrators in Forest B can add the accounts of users in Forest A to the access control list (ACL) on the resource. Assuming that there are a large number of users who need the same access to a resource, it is not feasible to add the individual account of every user to the ACL on every resource that is required by each user. Rather, a more scalable management model is required to allow users in Forest A to access resources in Forest B.

**Solution**

The model that allows users to access resources across forests is very similar to the model for allowing users in different domains in the same forest to have access to resources across domains. The most efficient management model is as follows:

- Create trust relationships between domains in different forests, as follows:

  - To enable resource sharing between all domains in two Windows Server 2003 forests, create a forest trust between the forest root domains. This trust relationship, which is available in Windows Server 2003 forests, enables two-way transitive trust between all domains in two forests. Forest trust relationships are not transitive beyond the two forests that share the trust.

    | **Note** |
    | --- |
    | A transitive trust relationship flows throughout a set of domains, such as a domain tree, and forms a relationship between a domain and all domains that trust that domain. For example, if domain A has a transitive trust with domain B, and domain B trusts domain C, then domain A trusts domain C. Transitive trusts can be one-way or two-way, and they are required for Kerberos-based authentication and Active Directory replication. |

- To enable resource sharing between specific domains in different forests, create external one-way, non-transitive domain trust relationships (explicit trusts) between the specific domains.

- To represent the sets of users who need access to the same type of resources, create role-based global groups in every domain and forest that contains these users.

  | **Note** |
  | --- |
  | You might already have such groups created in some forests. |

- Create universal groups that correspond to the global groups.

- Add multiple global groups to a universal group so that you can assign permissions to related resources in multiple forests.

  | **Note** |
  | --- |
  | Universal groups are not available as security groups in Windows 2000 mixed-mode domains or in Windows Server 2003 domains that have a domain functional level of Windows 2000 mixed. They are available for distribution groups. For information about Active Directory functional levels, see Help and Support Center for Windows Server 2003. |

- Create resource-based domain local groups in every domain for the resources in that domain that need to be accessed from outside the forest.

- Add universal groups (or global groups in mixed-mode domains) from all such forests to the domain local groups across the organization, and use the domain local groups to assign permissions on the resources within the respective domains.

When a new user account needs access to a resource in a different forest, add the account to the respective global group. When a new resource needs to be shared across forests, add the appropriate domain local group to the ACL for that resource. In this way, access is enabled for resources on the basis of group membership.

Figures 4 and 5 compare and contrast this model in single versus multiple forests and in Windows 2000 versus Windows Server 2003. In Figure 4, Forest A has Windows 2000 domains that are in mixed mode and therefore universal groups are not available. Global groups from Domain A and Domain C are added to the domain local group in Domain B, and that group is added to the ACL on the Domain B resource and permissions are assigned to the group that provide the appropriate level of access. The multiforest deployment follows the same model for three forests as is used for three domains in the single-forest deployment.

The scenarios in Figure 4 apply to domains and forests that do not have universal groups available to them. Domains that do not have the ability to use universal groups can be:

- Windows 2000 domains that have the domain mode of mixed.
- Windows Server 2003 domains that have the domain functional level of Windows 2000 mixed.

In Figure 4, group abbreviations are as follows:

- Global group (GG).
- Domain local group (DLG).




**Figure 4 Resource sharing when universal groups are not available**

In Figure 5, multiple global groups are added to a single universal group, which is added to the domain local group. In a single forest, one universal group per role is all that's required. In multiple forests, a universal group can be created for each forest that has users that need access to a resource in a different forest. This approach is useful when multiple global groups from the same forest need access to the same local resource.

The scenarios in Figure 5 apply to domains in which universal groups are available. Domains that have the ability to use universal groups are:

- Windows 2000 domains that have the domain mode of native.
- Windows Server 2003 domains that have the domain functional level of Windows 2000 native or Windows Server 2003.

In Figure 5, group abbreviations are as follows:

- Global group (GG).
- Domain local group (DLG).
- Universal group (UG)




**Figure 5 Resource sharing when universal groups are available**

**Adding Groups to Domain Local Groups vs. Computer Local Groups**

As a general rule, the practice of adding domain local groups to global groups and adding global groups to universal groups is effective when a large number of users must be granted access to several resources. When the number of users is few, you can add them directly to the universal group or to the resource.

In addition, when you have a choice of adding universal groups or global groups to a domain local group or to a computer local group, use the computer local group instead of the domain local group. The reason to do so is that when an access token is constructed, domain local groups, global groups, and universal groups of which the security principal is a member are included in the token. Thus, entries in the token include nested group memberships, which means that if user A is a member of group G and group G is member of group GR, the access token for user A will contain both group G and group GR. A maximum of 1,500 entries can be added to an access token, and if a user belongs to too many groups, this limit can be reached. However, computer local groups are not nested.

Therefore, if resource access is being enabled on a per-computer basis and the number of computers is not very large, using computer local groups avoids potentially nested memberships of a domain local group and is the better practice.

For example, if you want to allow a user to administer an Exchange server, adding their account or group to the computer local group on the Exchange server is more effective than creating a domain local group, adding that domain local group to the Exchange Admins group or to the ACL on the Exchange server object and setting permissions, and then adding the user's account or group to the domain local group.

**Security Considerations for Resource Access Across Forests**

Even when you create multiple forests for data isolation, service isolation, or both, trust relationships with other forests (or individual domains from other forests) can be created and users can be added to groups in other forests so that they can be granted access to resources. However, do not include users from other forests in any of the following groups:

- Groups that are responsible for service management or that can manage the membership of service administrator groups.
- Groups with administrative control over computers that store protected data.
- Groups that have access to protected data, or groups responsible for the management of user or group objects that have access to protected data.

Including users from another forest in any of these groups constitutes a potential security breach of the other forest that can lead to disclosure of protected data or service interruptions.

For information about these administrative groups, see "Design Considerations for Delegation of Administration in Active Directory." You can view this document on the Microsoft TechNet Web site at http://go.microsoft.com/fwlink/?LinkId=10516.

**Synchronizing Sites and Subnets**

Site and subnet information represents the physical topology of the network and is not related to the Active Directory

structure. Active Directory can determine the location of computers on the network according to the relationship of their respective IP addresses to a range of addresses defined by a subnet object in Active Directory. Subnets are associated with sites, and sites identify areas of good network connectivity. Sites are usually connected by Wide Area Network (WAN) links. For this reason, an ideal connection is one that does not cross site boundaries (traffic does not traverse a WAN link). Site and subnet information enables computers that are running Windows 2000, Microsoft® Windows® XP, or a member of the Windows Server 2003 operating systems to locate the following computers and services in the same site or a nearby site:

- Domain controller, global catalog server, or Key Distribution Center (KDC).
- Distributed File System (DFS) server that hosts a requested resource.

**Issue**

When site and subnet information in both forests is identical, or the Netlogon optimization is leveraged, computer and service location is optimized according to site (see article 916474 (http://go.microsoft.com/fwlink/?LinkId=166344) in the Microsoft Knowledge Base). The ability to locate computers and services in the same site or a nearby site depends in part on whether the appropriate site and subnet information is available in Active Directory.

It is important to note that establishing synchronization of site and subnet information is not an immediate requirement of implementing multiple forests. You can establish site and subnet synchronization at any point throughout the lifetime of a multiforest deployment, or you can decide to not establish it. If you decide to establish synchronization of site and subnet information, the question of how soon to establish it depends on the impact of not having an identical site topology between forests or not having Netlogon optimized.

For example, suppose an organization has Forest A and Forest B, and site and subnet information is not synchronized between them. The following scenarios describe the effects on the performance of three common operations that are affected by site awareness:

- Logon process
- File downloads
- Authentication

Use these scenarios to help determine whether the reduction in network traffic or request execution time, or both, outweigh the additional effort that is required to synchronize site and subnet information across the forests.

**Logon Process Across Forests**

When a user from Forest A logs on to a computer in Forest B, the logon process requires location of a domain controller from the user's domain in Forest A. If the site of the computer from Forest B is not specified in Active Directory in Forest A, the computer might locate any (rather than the closest) domain controller from the user's domain in Forest A. If the connection to the domain controller is made over a WAN, then this logon process adds traffic to the WAN (the amount of traffic depends on the number of Group Policies and logon scripts, as well as the size of the roaming profile). This connection generates traffic during logon as well as logoff and usually ranges from 100 kilobytes (KB) to a few hundred KB. Depending on the WAN bandwidth that is available for logon traffic, logon duration over the WAN might be increased. If these drawbacks are acceptable, especially if you anticipate that most users will be logging on to computers in their own forest, then site and subnet information might not be important enough to warrant synchronizing the data between forests.

**File Download Across Forests**

When a user who is logged on to a computer that is joined to Forest B requests a file that is hosted by multiple DFS servers, the nearest one being a server that is joined to Forest A, the DFS server that is contacted for the download depends on whether site and subnet information for Forest A is available in Forest B. If the site of this DFS server is not specified in Forest B, then the file might be downloaded from an arbitrary (potentially remote) DFS server. If the site of the DFS server is available in Forest B, the server in Forest A can be located.

| **Note** |
| --- |
| DFS enhancements in Windows Server 2003 ensure that files are downloaded from the next closest DFS server that hosts the desired file and is joined to Forest B. |

Downloading a file from a remote DFS server increases network traffic over a WAN (the amount of traffic is determined by the size of the file to be downloaded) and potentially increases the download time (the delay depends on bandwidth available to

download the file). If these drawbacks are acceptable, especially if you anticipate that users download files only from the DFS servers in their own forest, then synchronizing site and subnet information might not be important.

**Authentication Across Forests**
When a user needs to authenticate to a resource in a different forest, the user's computer or domain controller (depending on the authentication mechanism that is used) must contact a domain controller in the domain of the resource. Domain controller location is optimized by the closest site across forests only when identical site and subnet information is configured in Active Directory in both forests. However, if the traffic that is generated by authentication of the user does not cause significant delay, then it is not critical that a local domain controller be contacted.

**Solution**
The initial solution for mirroring site and subnet information is to create the same site and subnet objects in all forests. After these objects have been created, two methods can be used to ensure that site and subnet information is maintained identically in both forests:

- Use a directory data synchronization product (for example, Microsoft Identity Integration Server 2003) to synchronize the data when site and subnet information changes in one forest. This approach is characterized by a high level of automation and requires practically no administrator involvement after Microsoft Identity Integration Server 2003 configuration is in place. However, this approach is not acceptable in scenarios where service isolation is required (the service administrators in the destination forest do not trust the service administrators from the source forest).

- Establish a business process by which network administrators inform service administrators in each forest when site and subnet information changes, and the service administrators then update the information in their respective forests. This approach is most appropriate for the customers who have service isolation requirements for separate forests. However, this approach requires a high level of administrator overhead to create mirrored information in multiple forests.

Microsoft intends to issue a deployment guide that will describe the site and subnet objects and attributes that need to be synchronized to mirror network topology information.

**Retrieving Application Data Across Forests**
A rapidly growing family of applications is using Active Directory to store and retrieve application-specific data. Application architecture assumes that application-specific data is stored in a well-known (for service providers and consumers) Active Directory location that is populated by the service providers or administrators.

**Issue**
When clients of the service or application data need to obtain the data or locate the service, they perform a search against Active Directory and retrieve the data. Most clients assume that they are joined to the same forest as the service provider. In multiple forest deployments, this condition is not always true.

For example, service providers for service App1 are running on computers that are joined to Forest A and they register the information that enables clients to locate these service providers only in Forest A. In this case, if a client in Forest B queries Active Directory to discover a service provider for service App1, the query fails and the client cannot locate a service provider.

Clients do not always use the data stored in Active Directory to locate service providers. Sometimes a client or other service provider needs to retrieve data from Active Directory but does not need to locate any other service provider. The following solutions apply to both scenarios.

**Solution**
Two approaches can be used to give clients the ability to find information in Active Directory in a different forest:

- Deploy service providers in each forest.

- Synchronize the data across the forests.

Making a choice between the two approaches depends on the following requirements of the application:

- Load sharing.

- Discovery of the nearest service provider.

- Autonomy or isolation, or both (for security purposes).

**Deploy Additional Service Providers**
The following scenario describes a condition that requires deploying additional service providers on the basis of different application requirements.

**Load requirements**
If service providers that are deployed in one forest cannot sustain the load of requests from clients from multiple forests, then deploying additional service providers might be necessary. Deploying them in all individual forests eliminates a need to synchronize this application-specific data across the forests.

**Synchronize Data Across Forests**
The following scenario describes a condition that requires synchronizing data across multiple forests on the basis of different application requirements.

**Nearest service provider**
For some network-traffic-intensive or response-time-sensitive applications, it is important to ensure that a client locates the nearest service provider. In this case, deploying one or a few service providers to serve an entire network that spans multiple sites might not be appropriate. You might need to deploy at least one service provider per forest per site, but if multiple forests span the overlapping set of sites, then duplicating service providers for each forest might unnecessarily increase the cost of hardware and support of the service providers. In this case, synchronizing the application-specific data across different forests allows clients from one forest to locate the nearest service providers, even though the clients belong to a different forest. For an example of the service providers that can be used by all clients in the same site regardless of their forest affiliation, see "Locating Printers Across Forests" later in this paper.

**No Sharing of Application Data Across Forests**
In the following cases, providing application data across forests is not feasible.

**Autonomy or isolation**
For some applications, in addition to availability of the service and proximity of the service providers, some administrators might have additional requirements of autonomy. If, for example, an application or service owner in one division (that also has its own forest) requires special configuration of the service provider (not matching the configuration in other parts of organization), then users in this forest might not use the service providers in other forests and must rely on only the service providers in the same forest. In this case, there is no reason to synchronize the application-specific data between the forests.

In addition to autonomy or isolation requirements, the owner of a service might also have security reasons for requiring that users and computers use only the service providers from their own forest. Remember that service administrators in a different forest have full control over the services and data stored on the computers joined to their forest. If you cannot allow delivery of your service to depend on another forest (for example, if the other forest is compromised), then you must not synchronize the service-specific data to your forest.

**Considerations for Using Microsoft Identity Integration Server 2003 to Synchronize Application-Specific Data Between Forests**
If you decide to synchronize some application-specific data between forests, the recommended synchronization solution is Microsoft Identity Integration Server 2003. The solution involves the following administrative cooperation:

- The Microsoft Identity Integration Server 2003 administrator configures Microsoft Identity Integration Server 2003 to perform synchronization of the application-specific data.
- The application owner provides the Microsoft Identity Integration Server 2003 administrator with the description of the objects and attributes that need to be synchronized between the forests.
- The forest administrators indicate to Microsoft Identity Integration Server 2003 administrators the directory locations that the information should be read from and written to.

In collaboration with application administrators, administrators of the individual forests must decide relative to each application if they want to participate in data synchronization. This decision should be based on a need to share application-specific data with other forests, to use application-specific data from other forests, and on security requirements.

**Credentials for Writing Application Data**
When a forest administrator decides to allow Microsoft Identity Integration Server 2003 to write application-specific data for a specific application, the administrator must decide which credentials Microsoft Identity Integration Server 2003 should use to write this data. The recommended method is to create a user account that uses its credentials exclusively to allow Microsoft Identity Integration Server 2003 to write this data. The administrator does not have to create separate credentials for each application and might use the same set of credentials for data that is used by different applications.

**Security Considerations for Writing Application Data**

A forest administrator must also add appropriate write permissions to the container into which Microsoft Identity Integration Server 2003 will write the data. Depending on security requirements, the destination forest administrator might need to ensure that these credentials do not enable one to overwrite data from a different forest.

Even if you trust the Microsoft Identity Integration Server 2003 administrator to appropriately configure the synchronization server and you completely eliminate the possibility that the Microsoft Identity Integration Server 2003 administrator could be rogue, the possibility of the administrator being coerced still exists. To minimize the risk that a Microsoft Identity Integration Server 2003 administrator could jeopardize the functionality of an entire individual forest, the forest administrator, in addition to creating a dedicated organizational unit in which synchronized objects are created, should ensure that write permission granted to Microsoft Identity Integration Server 2003 to create, update, and delete data is limited to only this organizational unit, and no other portion of the forest.

Restricting access to a single organizational unit eliminates the possibility that Microsoft Identity Integration Server 2003 can overwrite any data outside of the organizational unit. If this approach is taken, the forest administrator must inform the Microsoft Identity Integration Server 2003 administrator of the location of this organizational unit in the forest hierarchy and ensure that any requirements of the application to read and write to various objects on the directory are fulfilled. For example, some applications expect the data to be in specific locations of the Active Directory hierarchy, and only in those locations.

| **Note** |
| --- |
| If you are using Microsoft Metadirectory Services version 2.2, Microsoft recommends that you hire a Microsoft Consulting Services (MCS) consultant who specializes in Microsoft Metadirectory Services 2.2 to configure Microsoft Metadirectory Services 2.2 to enable data synchronization. |

**Locating Printers Across Forests**

The presence of printer objects in Active Directory enables a computer that has an account in that forest to browse for printers that have specific characteristics (for example, printers that support color printing or stapling) and that are located near the computer.

**Issue**

In the multiforest scenario, a user might be logged on to a computer that is in the same physical location as a printer, but the computer account and printer information exist in different forests. In this case, the user can locate a printer in a different forest only by providing the print server and printer name. You might want to allow users to browse for printers on the basis of their proximity and without regard to forest membership.

**Solution**

It is possible for a user to configure a computer to use a particular printer by specifying the printer name, which requires no additional configuration on the part of the forest administrator. However, if you want to enable users in one forest to browse for printers in a different forest, then consider using Microsoft Identity Integration Server 2003 to synchronize printer data between the forests, as described in “Retrieving Application Data Across Forests” earlier in this document.

Microsoft intends to issue a deployment guide with instructions for how to enable this functionality.

**Synchronizing Exchange Data Across Forests**

Exchange 2000 Server uses Active Directory to store configuration and user information and to provide user authentication and permissions management. In multiforest deployments where more than one Exchange organization is deployed, additional configuration is required to synchronize mail-related data of multiple Exchange organizations among the multiple forests.

**Exchange Organizations and Active Directory**

An *Exchange organization* consists of one or more Exchange servers, and each Exchange organization is specific to one Active Directory forest. Exchange servers rely on access to the global catalog for address information. Because each forest has a separate global catalog, an Exchange organization is associated with only one forest.

However, multiple forests can use the same Exchange organization for mail service. A single Exchange organization that serves multiple forests does not require additional configuration across forests because the Exchange organization uses only one forest for its Active Directory storage and services.

Although user accounts in different forests can use the same Exchange organization to send and receive mail among users in all forests, the servers of a single Exchange organization cannot access information from more than one Active Directory. To enable multiple forests that each deploy separate Exchange organizations to function as a single business organization, additional configuration is required to synchronize the Exchange mail recipients in the respective Active Directories.

The ability to access address information from different forests requires that shadow (or proxy) versions of all mail recipients be created and maintained by synchronization in each relevant directory. With these shadow objects in place, when an

Exchange server queries a single directory, it can view information about mail recipients that originate in its forest as well as those that have been created as shadow objects.

**Single Exchange Organization**

Whether a single Exchange organization serves one forest or more than one forest, the Exchange organization is still associated with only one of the forests, called the *Exchange forest*. Users that have accounts in one forest might have mailboxes in the same forest or in a different forest; however, mailboxes are always in the same forest as the Exchange servers because mailbox data is stored on the Exchange servers.

For example, in a scenario where Forest A and Forest B share a single Exchange organization, Exchange servers are installed only in Forest B. User A has a user account in Forest A and User A's mailbox is stored in forest B. To associate the user in Forest A with the user's mailbox in Forest B, a disabled user account is created for User A in Forest B and the mailbox is associated with it. User A's mail is stored on Exchange servers in Forest B and an attribute on the disabled account references User A's account in Forest A. This type of Exchange environment is known as the *resource forest model*.

Figure 6 shows an Exchange organization that has mailboxes on Exchange servers in one forest and user accounts in a different forest. In this scenario, the user objects in Forest A have corresponding disabled user accounts that represent their respective Exchange mailboxes in Forest B, which hosts the Exchange servers.




Single Exchange Organization




**Figure 6 Single Exchange organization with users and mailboxes in different forests**

Single Exchange organizations can have users in more than one forest, as shown in Figure 7.




**Figure 7 Single Exchange organization that serves multiple forests**

**Multiple Exchange Organizations**

When the same enterprise deploys Exchange servers in multiple forests, each Exchange organization has a separate Exchange forest and therefore a separate global catalog and global address list. Additional configuration is required to enable users in each forest to be able to look up address list information in the other forests. To achieve this goal, mail recipient information from multiple Exchange organizations must be synchronized.

Figure 8 shows a multiforest enterprise that deploys Exchange servers in multiple forests.




**Figure 8 Multiple Exchange organizations and multiple forests**

**Note**

It is not a goal of this document to prescribe best practices to design Active Directory for Exchange 2000 Server. For more information about designing Active Directory for Exchange 2000 Server in a Windows 2000 Server environment, see "Best Practice Active Directory Design for Exchange 2000." You can download this document from the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=10517.

For deployments in which multiple Exchange organizations serve separate forests in the same multiforest enterprise, additional configuration might be needed in order to enable:

- Viewing the global address list in all forests.
- Viewing free and busy calendar information in all forests.

The configuration information in this paper encompasses all scenarios in which Exchange Server is deployed in more than one forest, but not necessarily in every forest in the enterprise.

**Address List Synchronization**

In a multiforest environment with a single Exchange organization, all user accounts reference global catalog servers in the Exchange forest to gain access to the global address list. Although the users are in different forests, Exchange uses a single global address list for these users and therefore no synchronization of address list information is needed.

In a multiforest environment with multiple Exchange organizations, each Exchange forest has its own address list and therefore synchronization is required to present the same view of address list data across the forests.

The global address list is a list of entities in a forest that can receive mail. Mail recipients can be user accounts (both enabled and disabled accounts), contacts, distribution lists, security groups, and folders. Exchange administrators can also create different views of the global list of mail recipient information to generate logical groupings of recipients, which are viewable as address lists.

Every Outlook client is configured with the name of an Exchange server. When an Outlook client user opens the Address Book, or when a user composes a message and types a name or an address in the To: field, the Outlook client uses the global catalog server that is specified by its Exchange server to search the contents of the global address list or other address lists.

**Issue**

Outlook clients must use global catalog servers and Exchange servers in only one forest to look up addresses. For this reason, the clients in one forest can view address lists of another forest in the same Exchange organization, but not of a forest in a different Exchange organization. Therefore, the following issues must be addressed in a multiforest deployment that has multiple Exchange organizations:

- Users in one forest need to be able to look up Address Book data, including mail addresses of objects in a different forest.
- The Exchange client must be able to resolve all addresses in its own forest.

Figure 9 shows two Active Directory forests with Exchange servers in each forest and an example of a mail-enabled user in each forest. The global address list (GAL) of each respective forest does not contain users from the other forest. For this reason, a query by an Exchange client does not find address lists for the other forest.




**Figure 9 Two forests with separate Exchange global address lists (GALs)**

In the configuration that is shown in Figure 9, separate Exchange organizations are deployed in two forests. In this case, before any additional configuration occurs, Forest A does not contain an account for Fred, who resides in Forest B. Likewise, Forest B does not contain an account for John, who resides in Forest A. Therefore, when the user of the Outlook client that is referencing the Exchange server that is configured to use the global address list in Forest A opens the Address Book, Fred is not found in the Address Book. John also cannot send mail to Fred by using an alias, but must use the full Simple Mail Transfer Protocol (SMTP) address for Fred.

The same restriction applies to the distribution lists and resources (for example, conference rooms) that reside in different forests. For example, if a distribution list exists in Forest A, then without this additional configuration, a user whose Outlook client is retrieving global address list information from the global catalog in Forest B does not see this distribution list in the Forest B Address Book.

**Solution**

To allow users to see mail recipients from another forest in the Address Book, the mail recipients must be synchronized between forests. In this scenario, the mail recipients are synchronized in both directions.

**Synchronizing users and groups between forests**

The solution that enables users in different forests to see the identities of each other in the Address Book is to create a proxy, or "shadow," version of mail-enabled objects in the forests that have Exchange 2000 Server deployed. The shadow version is represented by a contact object. In Figure 10, the mail account of John in Forest A is represented by the contact object John in Forest B, while the mail account of Fred in Forest B is represented by the contact object Fred in Forest A. Synchronization is performed by a directory data synchronization product such as Microsoft Identity Integration Server 2003.




**Figure 10 Mailbox synchronization between two forests**

In this example, if administrators of both forests want to allow users from Forest B to view the users from Forest A in the Address Book, then administrators can configure a directory synchronization server, such as Microsoft Identity Integration Server 2003, to perform the following operations:

- Search Forest A for all mail-enabled objects.
- Create corresponding contact objects in Forest B.

These operations can be scheduled to run periodically to dynamically update information in all forests when a mail-enabled object is added, deleted, or modified in its source forest. Upon completion of these operations, Exchange services (running every minute by default) in all forests that have an Exchange server installed will appropriately update existing address lists, including the global address list. Thereafter, when the user Fred in Forest B attempts to find the user John from Forest A by searching the Address Book, Fred successfully finds John because John is represented by a contact object in Forest B.

In addition, because routing information is included in the contact object, a user in Forest B can successfully send messages to John by specifying John's alias. The Exchange server locates the full SMTP address for the alias for the user John in the John contact object and properly forwards the message to the Exchange Server service that contains the mailbox of the user John.

Depending on the decision of the administrators of the source forest (the forest from which data is read) and the destination forest (the forest in which data is updated to match the appropriate data in the source forest), and assuming proper configuration, Microsoft Identity Integration Server 2003 can perform synchronization in one or both directions.

The process is very similar for contact objects that are synchronized out of this forest as well as mail-enabled security and distribution groups. As shown in Table 1, all of these object types are synchronized as contacts into other forests.

**Table 1**

| Source Directory | Target Directory |
|---|---|
| User | Contact |
| Contact | Contact |
| Group | Contact |

**Additional scenarios for address synchronization in forests without Exchange**
The scenarios described earlier apply to multiple forests that deploy separate Exchange organizations. Under the following conditions, address list synchronization might be deployed in forests that do not deploy Exchange servers:

- The attributes on the mailbox-enabled user account need to be propagated back to the forest with the logon user account. For example, if a human resources database typically updates phone numbers on the user account in the Exchange forest and other applications look up this information in the user account forest, then synchronization is required.

- Applications other than Outlook use the Address Book.

**Cost of global address list synchronization**
The cost of synchronizing the global address list and other address lists is that of implementing a directory data synchronization product such as Microsoft Identity Integration Server 2003. A single Microsoft Identity Integration Server 2003 server is required for every synchronization infrastructure that you want to implement. For example, if you decide to use a centralized Microsoft Identity Integration Server 2003 administrator who manages global address list synchronization for all of the forests, then the cost is that of running Microsoft Identity Integration Server 2003 on a single server computer plus a Microsoft Identity Integration Server 2003 administrator who will manage the process. However, if the forest administrators do not trust each other and want to run their own solutions, then each forest must have a Microsoft Identity Integration Server 2003 server with a Microsoft Identity Integration Server 2003 administrator who manages getting data into the forest.

**Which synchronization solution to use**
Any solution that enables synchronization of the objects described above can be used to synchronize address lists across multiple Exchange organizations. However, Microsoft highly recommends using Microsoft Identity Integration Server 2003. Microsoft Identity Integration Server 2003 is capable of synchronizing directory data between multiple Active Directory forests according to the business rules that are specified in the Microsoft Identity Integration Server 2003 configuration.

If you are using Microsoft Metadirectory Services (MMS) version 2.2, it is recommended that you hire a Microsoft consultant to perform the necessary configuration of MMS 2.2 to enable global address list synchronization.

Concurrent with shipping Microsoft Identity Integration Server 2003, Microsoft intends to provide a Global Address List (GAL) Synchronization Solution that automates the configuration of Microsoft Identity Integration Server 2003. The GAL Synchronization Solution significantly simplifies the configuration of Microsoft Identity Integration Server 2003 to perform global address list synchronization, including the creation of contact objects and distribution groups. The GAL Synchronization Solution performs configuration of that is expected to satisfy business rules of the majority of customers. At the same time, customers have the flexibility to modify Microsoft Identity Integration Server 2003 configuration on the basis of their individual business rules regarding global address list synchronization data. Detailed information about the GAL Synchronization Solution is to be included in the Multiforest Deployment Guide, which Microsoft intends to issue in the future.

**Administrative considerations for global address list synchronization**
To enable synchronization of global address list data by Microsoft Identity Integration Server 2003, administrators of the source and destination forests must provide information to the Microsoft Identity Integration Server 2003 administrators, as follows:

- The source forest administrator creates a user account that has read permission on the object that is to be synchronized. Credentials for this account will be used by Microsoft Identity Integration Server 2003 to read the objects from the source forest. The source forest administrator also specifies whether to allow group membership expansion.

- The destination forest administrator creates or identifies an organizational unit in which Microsoft Identity Integration Server 2003 will create the new contact objects. The administrator also creates a user account and grants permissions on the organizational unit to allow the account to create, delete, and update (write) objects in the organizational unit. To eliminate the possibility of tampering with destination and forest data outside of the organizational unit, allow write permissions only in the organizational unit that the administrator created for the purpose of creating synchronized

objects, and to specific attributes of accounts outside the organizational unit that are needed to preserve mail functionality.

| Note |
| --- |
| If the forest administrator decides to synchronize data both into and out of the forest, then the administrator should use the same user account credentials to perform both read and write operations. |

**Risk models for global address list synchronization**

There are two main risk models that apply to enabling global address list synchronization. The two models can coexist within the same organization. These risks apply in a more general context to all synchronization solutions.

- Individual forest administrators trust centralized Microsoft Identity Integration Server 2003 administrators to synchronize data between the forests. They are willing to risk the possibility that a centralized Microsoft Identity Integration Server 2003 administrator might maliciously or accidentally misconfigure Microsoft Identity Integration Server 2003. As a rule, individual forest administrators can and should protect their forest from data disclosure or tampering by ensuring that the user whose credentials are used by Microsoft Identity Integration Server 2003 to perform synchronization does not have permissions that would allow malicious activities. This approach is recommended.

- Individual forest administrators are not willing to risk the possibility that a centralized Microsoft Identity Integration Server 2003 administrator might maliciously or accidentally misconfigure Microsoft Identity Integration Server 2003. In this scenario, each individual forest administrator might deploy a Microsoft Identity Integration Server 2003 server that performs data synchronization from or into their forests, or both. The only cooperation required in this scenario is that the partner forest administrator must provide sufficient credentials to allow reading and writing the appropriate data. This approach is slightly more complex and might have a higher overhead, depending on the number of forests involved.

**Free and Busy Calendar Information Synchronization**

Calendar information enables users to schedule meetings by viewing the free and busy times in the schedules of meeting invitees. If frequent meetings are to be scheduled between members of different forests, the availability of schedule information might warrant synchronizing free and busy calendar data between Exchange organizations in different forests.

**Issue**

John has a mailbox in Forest A and wants to use Outlook to schedule a meeting with Fred, whose mailbox is in Forest B. John also wants to invite a group of people from Forest C. As shown in Figure 11, if John does not see the availability of users and groups in other forests, scheduling a meeting requires contacting them personally.




**Figure 11 Scheduling when calendar data is unavailable**

Free and busy calendar information is stored in the system public folders on Exchange 2000 servers, but not in Active Directory. For this reason, synchronizing global address list information between different forests is not sufficient to enable users that have mailboxes in one forest to view free and busy calendar information of users and resources that have mailboxes in Exchange organizations in other forests.

**Solution**

To enable users that have mailboxes in Forest A to view schedule information of users and resources that have mailboxes in the Exchange organizations in Forest B and Forest C, the system public folders data stored on the Exchange servers in Forest B

and Forest C must be included in the system public folders of the Exchange servers in Forest A. To enable such synchronization, Microsoft recommends using the Microsoft Exchange Interorg Replication Utility (also called the Public Folder [PF] Sync Tool). For information about using this tool, see article Q238573, "Installing and Configuring InterOrg Replication Utility," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://go.microsoft.com/fwlink/?LinkId=291.

In most multiforest environments, a single synchronization server can accommodate multiple forests. Figure 12 shows a multiforest deployment that uses a single server that runs the PF Sync Tool to synchronize free and busy information among all three forests. In this scenario, each forest synchronizes data to and from two other forests.




**Figure 12 Free and busy information synchronized across three forests by one server**

**Migration Considerations**

In the environment that contains multiple Active Directory forests and potentially multiple Exchange organizations in different Active Directory forests, reorganizations or individual employee relocations will almost inevitably require you to migrate one or more of these: user accounts, computer accounts, groups, or mailboxes.

**User Migration**

User migration includes user accounts, groups, and computer accounts. To perform user account migration between forests, use Active Directory Migration Tool (ADMT). For information about using ADMT and to download ADMT, see "Active Directory Migration Tool Overview" in the Active Directory link on the Web Resources page at http://go.microsoft.com/fwlink/?LinkId=291.

**Mailbox Migration**

To perform migration of mailboxes between different Exchange organizations in different forests, use the Exchange Server Migration Wizard. For information about how to perform user mailbox account migration, see the Exchange online documentation.

If you need to migrate both user mailboxes and user accounts, perform the migration of the user account prior to migrating the user's mailbox.

**Considerations for Other Technologies**

Group Policy and Message Queuing features are not directly transferable across multiple Windows forests. Some workarounds are available, as described in "Group Policy" and "Message Queuing" later in this paper.

In addition, if you use trust relationships between domains in different forests or between two Windows Server 2003 forests, consideration for whether you want the ability to use user principal names (UPNs) to log on across forests can affect your decisions about the UPN suffixes you choose and the type of UPNs you create, as described in "Using the User Principal Name to Log On Across Forests" later in this paper.

**Group Policy**

Group Policy provides administrative control over users and computers. Group Policy can be applied to a domain, organizational unit, or site within the same forest. It is not possible to apply the same Group Policy object to users and computers in multiple forests.

If necessary, administrators of the forests can apply identical Group Policy objects to the domains, organizational units, or sites in different forests. For example, if you are managing merged organizations and have applied identical Group Policy objects to affect the users and computers in each domain in the original forest, you might want to apply identical Group Policy objects to all users and computers in each domain in the other (merged) forest.

Another scenario in which Group Policy management across forests is important is in the test-to-production scenario. For

example, for staging purposes, many customers maintain a separate test forest that closely resembles the production forest. After the policy deployment has been tested in the test environment, they want to reproduce this deployment in the production forest.

The Group Policy Management Console (GPMC) enables you to manage Group Policy across multiple forests by importing and copying Group Policy objects across forests. The GPMC will be available on the Microsoft Web site shortly after the release of Windows Server 2003. Although Group Policy Management can run only on the computers running Windows XP Professional and Windows Server 2003 operating systems, it can remotely manage Group Policy objects on domain controllers that are running a member of the Windows 2000 Server family with Service Pack 2 or greater.

**Message Queuing**
If you plan to deploy an application that uses Message Queuing (also known as MSMQ) and runs across the boundaries of multiple Active Directory forests, the application developer should the follow these recommendations:

- Use the direct format name for the target queues so that:
    - Message Queuing does not use Active Directory to resolve the target name.
    - The user cannot use Message Queuing built-in authentication and encryption features.
- If the application will run on a Windows XP or Windows Server 2003 platform, then use Message Queuing over HTTPS with client authentication.

For more information about Message Queuing, see Windows 2000 Server Help or Help and Support Center for Windows Server 2003.

**Using the User Principal Name to Log On Across Forests**
A user principal name (UPN) is a variation of a user account name that looks like an e-mail name but can be used to log on to a domain. The syntax is <*user name*>@<*string*>. UPNs allow you to use the same logon name across different domains in the same forest or in different forests.

UPNs are of two types:

- Implicit: Always of the form *userID@DNSDomainName*. For example, johns@corp.contoso.com is the UPN for the account of John Smith, whose user ID is johns and whose account is a member of the corp.contoso.com forest. The implicit UPN is always associated with the user's account, regardless of whether an explicit UPN is defined.
- Explicit: Always of the form *string@Anystring*, where both "string" and "Anystring" are explicitly defined by the administrator. For example, John Smith might have the UPN ITJS@coneast. Explicit UPNs are useful for situations in which the organization does not want to publicize the name of domains or the forest structure.

A UPN *suffix* is a string that indicates the namespace from which a UPN is derived. A UPN suffix includes everything to the right of the rightmost "at" (@) symbol in the full UPN. Everything to the left of the rightmost @ symbol is considered part of the user name, even if it contains another @ symbol. UPN suffixes can be either flat or hierarchical in structure, although they are typically used to identify flat namespaces.

UPNs are used primarily for interactive logons across domains. With multiple forests, UPN logons are available across forests under some conditions and not others. Under either of the following conditions, UPNs cannot be used to log on to a domain in a different forest:

- Shared UPN suffix: If two or more forests use the same UPN suffix, users cannot use UPNs to log on to a domain in the trusting forest. They can use the UPN only for logons within their own forest.
- Explicit UPN and external trust: When external trust relationships are created between domains in different forests, users in the respective domains cannot use explicitly created UPNs to log on to the trusting domains in the other forest. They can use only implicit UPNs.

Table 2 summarizes a user's ability to use a UPN to log on interactively across forests according to the different kinds of trust relationships that exist between domains in different forests and whether UPN suffixes are shared by multiple forests.

**Table 2 Ability to Use UPNs to Log On Across Forests**

| Trust Type | Explicit UPNs | Implicit UPNs | Shared UPN Suffixes |
|---|---|---|---|
| No trust relationships between domains in different forests | No | No | N/A |
| External trust relationship between domains in different forests | No | Yes | No |
| Forest trust relationship between Windows Server 2003 forests | Yes | Yes | No |

**Note**

Forest trust relationships are available between forests that have an Active Directory forest functional level of Windows Server 2003. For information about Active Directory functional levels, see Help and Support Center for Windows Server 2003.

In summary, if you want to share UPN suffixes among multiple forests but do not need the ability to log on interactively to a domain in another forest, you can use either external trust relationships or forest trust relationships.

To use UPNs to log on interactively across forests, use unique UPN suffixes for each forest. In this case, interactive logons are available as follows:

- If you use forest trust, you can use both explicit and implicit UPNs.
- If you use external domain trusts, you can use only implicit UPNs.

**Features that are Unavailable Across Forests**

Although experience with multiforest deployments shows that the features that cannot be enabled across different forests are not used by most customers in any case, the description of these features is included in this document to ensure a complete picture of the impact of deploying multiple forests.

**Delegating Mail Folders Across Forests**

Within a forest, you can delegate authority for a user to have access to the mail folders in another user's mailbox.

**Issue**

If mailboxes of two users are hosted by Exchange Server organizations in different forests, then a user in one forest cannot access the mail folders of a user in the other forest.

A typical example of this issue is executives who want to delegate the ability to manage their calendars to administrative assistants. If the mailboxes of such an executive and administrative assistant exist in Exchange Server organizations that are in different Active Directory forests, then such administrative delegation of this ability is not possible.

**Workaround**

Ensure that mailboxes of people who need access to the mail folders of others are in the same Exchange Server organization. Also, if a user needs access to mailboxes that are hosted by different Exchange Server organizations, you can create multiple mailboxes in the appropriate forests for the same user who needs access to other mailboxes. For example, in the case of the executives and administrative assistants described earlier, the administrative assistant requires an additional mailbox in the forest of the executive.

# Multiforest Feature Implementation and Cost Summary (Multiple Forest Considerations in Windows 2000 and Windows Server 2003)

Updated: April 8, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After reading this document, use the information in Table 3 to estimate the relative costs of implementing functionality across forests for your organization.

**Table 3 Features and Configuration Costs**

| Enabled Feature | Configuration Costs |
|---|---|
| View address list information | • Deploy Microsoft Identity Integration Server 2003 synchronization server<br>• Train and support synchronization administrator<br>• Create synchronization data container per target forest<br>• Create synchronization account per forest<br>• Enable permissions for synchronization<br>• Run GAL Sync tool |
| View free and busy calendar information | • All of “View address list information” configuration costs, plus:<br>• Deploy PF Sync server<br>• Train and support PF Sync administrator<br>• Confirm MAPI-capable Local Area Network (LAN) connection between Exchange organizations |
| Bridge DNS infrastructures | DNS configuration of one or more of the following:<br>• Delegations<br>• Conditional forwarding (Windows Server 2003 only)<br>• Secondary zones |

| | |
|---|---|
| Resource access | • Configure trust relationships:<br>  • Forest trust (Windows Server 2003 only)<br>  -Or-<br>  • External domain trust (Windows Server 2003 or Windows 2000)<br>• Create groups:<br>  • Role-based global groups<br>  • Resource-based domain local groups<br>  • Universal groups, if available<br>• Add global and/or universal groups to local groups<br>• Apply permissions to resources |
| Application data retrieval | • Configure additional service providers<br>-Or-<br>• Configure Microsoft Identity Integration Server 2003 to synchronize application data (see costs under "View address list information" in this table) |
| Printer location | • Configure clients with print locations<br>-Or-<br>• Configure Microsoft Identity Integration Server 2003 to synchronize printer application data (see costs under "View address list information" in this table) |

# Network Load Balancing: Frequently Asked Questions for Windows 2000 and Windows Server 2003

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

Published: June 1, 2003 | Updated: February 10, 2004

**In This White Paper**

NLB Basics

NLB and Other Clustering Technologies

NLB for Scalability

NLB for High Availability

NLB Fundamentals - FAQ

Deployment Issues

Set up and Configuration

Management and Operations

Support for Sessions

Application and Protocol Support

# NLB Basics

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** What is Network Load Balancing?

**A.** Network Load Balancing is a clustering technology offered by Microsoft as part of all Windows 2000 Server and Windows Server 2003 family operating systems. NLB uses a distributed algorithm to load balance network traffic across a number of hosts, helping to enhance the scalability and availability of mission critical, IP-based services, such as Web, Virtual Private Networking, Streaming Media, Terminal Services, Proxy, etc. It also provides high availability by detecting host failures and automatically redistributing traffic to operational hosts.

**Q.** What is a Cluster?

**A.** A cluster is a group of independent computers that work together to run a common set of applications and provide the image of a single system to the client and application. The computers are physically connected by cables and programmatically connected by cluster software. These connections allow computers to use problem-solving features such as failover in Server clusters and load balancing in Network Load Balancing (NLB) clusters.

# NLB and Other Clustering Technologies

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** What is the Difference Between NLB and Server Clusters?

**A.** A server cluster (MSCS) is a collection of servers that together provide a single, highly available platform for hosting applications. Applications can be failed over to ensure high availability in the event of planned downtime due to maintenance or unplanned downtime due to hardware, Operating System or application failures. Server clusters provide a highly available platform for applications such as SQL Server, Exchange Server data stores, file and print servers, etc. Server clusters are used for stateful applications that rely on some state context from one request to the next.

Network Load Balancing (NLB) clusters dynamically distribute the flow of incoming TCP and UDP traffic among the cluster nodes according to a set of traffic-handling rules. NLB clusters provide a highly available and scalable platform for applications such as IIS, ISA server, etc. NLB is used for stateless applications; i.e. those that do not build any state as a result of a request.

NLB and server clusters compliment each other in complex architectures: NLB is used for load balancing requests between front-end web servers while server clusters provide high availability for backend database access.

NLB and server clusters cannot be used on the same set of servers (see Question Can I use NLB and server clusters on the same set of servers?).

**Q.** Can I Use NLB and Server Clusters on the Same Set of Servers?

**A.** No, Microsoft Server clusters (MSCS) and Network Load Balancing (NLB) are not supported on the same set of nodes. Both Server clusters and Network Load Balancing clusters control and configure network adapters. Since they are not aware of each other, configuring one can interfere with the other.

# NLB for Scalability

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** How Large Should My Cluster Be?

**A.** The size of the cluster is determined by the application that is being hosted and load balanced and on the system resources on the host machines available to service client requests sent to the application. If you notice that client requests are slowing down as more and more clients make connection requests to the application, it may be time to add more hosts to the cluster. A single NLB cluster can have up to 32 hosts in it.

For more information and examples of how to determine the size of your cluster refer to Chapter 10 Designing Network Load Balancing of the Windows Server 2003 Deployment Guide, (http://go.microsoft.com/fwlink/?LinkId=4298).

**Q.** Are There Any Performance Concerns as My Cluster Grows?

**A.** Tests have shown the NLB performance begins to deviate from the linear as the cluster grows beyond 20 to 25 nodes. But, this really depends on the amount of traffic that the application sees in the form of client requests. The following graphs give a better idea of NLB performance as the cluster size and number of users increase.







**Web Load: GET requests for 10KB static web pages**




**Q.** We Expect Our Front-End Cluster to Grow Beyond 32 Nodes. Should I Use NLB Even Though it Has a Cluster Size Limit of 32 nodes?

**A.** NLB can be used to scale beyond 32 machines by using Round Robin DNS between multiple NLB Clusters. This is shown in the following figure.



# NLB for High Availability

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** How Does NLB Detect a Server Failure?

**A.** While servicing client requests, each NLB Cluster host emits heartbeats to the other hosts in the cluster. If a host fails and stops emitting heartbeats, then after a default time period of five seconds, the remaining hosts in the cluster undergo a process called convergence to remove the failed host from the cluster and have new client connection requests mapped to remaining hosts in the cluster. So, new connections made by the client get handled by one of the remaining hosts in the cluster.

**Q.** How Long Does it Take For a Failed Server to be Removed From the Cluster?

**A.** By default, five seconds are required to detect a failed host. Once a failure is detected, the convergence process takes an additional two and a half to three seconds to evict the failed host and redistribute its load to the surviving hosts.

# NLB Fundamentals - FAQ

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** Do the Heartbeat Packets Consume a Lot of Bandwidth?

**A.** No, the heartbeat packets, which are emitted every second by each host, consume less than 1,500 bytes.

**Q.** Do the Heartbeats Need to Go on a Back-End Network?

**A.** No, the heartbeat packets are always sent out on the same network interface on which data packets are received and sent. There is no need for an additional back-end network for the control (heartbeat) packets. In fact, in order to be able to detect connectivity failures on the load-balanced network adapter, NLB actually **requires** that heartbeats traverse the load-balanced network adapter.

**Q.** Is NLB a Kernel Component?

**A.** Yes, NLB has a kernel component called wlbs.sys. This is an intermediate NDIS driver. NLB also has user-mode components for management purposes.

**Q.** What are the Benefits of NLB Over Simple Round Robin Domain Name Service (RRDNS)?

**A.**

- Automatic recovery within 5 seconds
- More even load balancing

**Q.** How Does the NLB Load Balancing Algorithm Work?

**A.** NLB employs a fully distributed filtering algorithm to map incoming clients to the cluster hosts. This algorithm enables cluster hosts to independently and quickly make a load balancing decision for each incoming packet. It is optimized to deliver statistically even load balance for a large client population making numerous, relatively small requests, such as those typically made to Web servers. When the client population is small and/or the client connections produce widely varying loads on the server, the load-balancing algorithm is less effective. However, the simplicity and speed of NLBs algorithm allows it to deliver very high performance, including both high throughput and low response time, in a wide range of useful client/server applications. If *No Affinity* is set, NLB load balances incoming client requests so as to direct a selected percentage of new requests to each cluster host; the load percentage for each host is set in the NLB Properties dialog for each port range to be load balanced. The algorithm does not dynamically respond to changes in the load on each cluster host (such as the CPU load or memory usage). However, the load distribution is modified when the cluster membership changes, and load percentages are renormalized accordingly.

When inspecting an arriving packet, all hosts simultaneously perform a mapping to quickly determine which host should handle the packet. The mapping uses a randomization function that calculates a host priority based on their IP address, port, and other information. The corresponding host forwards the packet up the network stack to TCP/IP, and the other cluster hosts discard it. The mapping remains unchanged unless the membership of cluster hosts changes, ensuring that a given clients IP address and port will always map to the same cluster host. However, the particular cluster host to which the clients IP address and port map cannot be predetermined since the randomization function takes into account the current and past clusters membership to minimize remappings.

In general, the quality of load balance is statistically determined by the number of clients making requests. This behavior is analogous to dice throws where the number of cluster hosts determines the number of sides of a die, and the number of client requests corresponds to the number of throws. The load distribution improves with the number of client requests just as the fraction of throws of an N-sided die resulting in a given face approaches 1/N with an increasing number of throws. As a rule of thumb, with client affinity set, there must be at least five times more clients than cluster hosts to begin to observe even load balance.

The Network Load Balancing client affinity settings are implemented by modifying the statistical mapping algorithms input data. When client affinity is selected in the NLB Properties dialog, the clients port information is not used as part of the mapping. Hence, all requests from the same client always map to the same host within the cluster. Note that this constraint has no timeout value and persists until there is a change in cluster membership. When single affinity is selected, the mapping algorithm uses the clients full IP address. However, when Class C affinity is selected, the algorithm uses only the Class C portion (upper 24 bits) of the clients IP address. This ensures that all clients within the same Class C address space map to the

same cluster host.

**Q.** How Does NLB Cluster Convergence Work?

**A.** Convergence involves computing a new cluster membership list and recalculating the statistical mapping of client requests to the cluster hosts. There are two instances in which cluster traffic has to be remapped due to a change in cluster membership: when a host leaves the cluster and when a host joins the cluster. A convergence can also be initiated when several other events take place on a cluster, such as changing the load balancing weight on a host or implementing port rule changes.

**Removing a Member**. Two situations cause a host to leave the cluster or go offline. First, the host can fail, an event that is detected by the NLB heartbeat. Second, a system administrator can explicitly remove a host out of the load-balancing cluster or stop NLB on that host.

**The NLB heartbeat**. NLB uses a heartbeat mechanism to determine the state of the hosts that are load balanced. This message is an Ethernet-level broadcast that goes to every load-balanced cluster host.

NLB assumes that a host is functioning normally within a cluster as long as it participates in the normal exchange of heartbeat messages between it and the other hosts. If the other hosts do not receive a message from a host for several periods of heartbeat exchange, they initiate convergence. The number of missed messages required to initiate convergence is set to five by default (but can be changed).

During convergence, NLB reduces the heartbeat period by one-half to expedite completion of the convergence process.

**Server Failure**. When a cluster host fails, the client sessions associated with the host are dropped.

After convergence occurs, client connections to the failed host are remapped among the remaining cluster hosts, who are unaffected by the failure and continue to satisfy existing client requests during convergence. Convergence ends when all the hosts report a consistent view of the cluster membership and distribution map for several heartbeat periods.

**Q.** Can NLB Balance Load Based on CPU/Memory Usage?

**A.** No, NLB does not respond to changes in the server load (such as CPU usage or memory utilization) or the health of an application.

# Deployment Issues

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** What is the AOL Proxy or Megaproxy Problem? Does NLB Fix it?

**A.** All connections that originate from large enterprises or megaservices such as AOL or MSN pass through client-side proxies. However, there is no guarantee that two consecutive TCP connections originating from the same client would be sent out through the same proxy. Now, lets assume that an AOL client creates an SSL session with an NLB cluster on the other side of the Internet. To ensure that the session does not break, all TCP connections originating from this client that are part of the session must be sent to the same host in the cluster. To ensure this, the NLB cluster must be configured in Single Affinity Mode.

However, it is possible that multiple TCP connections originating from the same client on the AOL network may be sent to the NLB cluster through different client-side proxies. The NLB Cluster will view these two TCP connections as connections from different clients (because they would now have different source IP addresses) even though they have actually originated from the same host. Now, since NLB load balances connection requests based on the source IP address of the connection, these connections may get handled by different hosts in the cluster even though they are part of the same session, thus breaking the session. NLB does not fix this problem, though the problem may be mitigated by choosing Class C affinity, provided the AOL Proxy happens to use source IP addresses from a single Class C address range. Note that a well designed web application will should not fail if an SSL session is broken; instead it should negotiate a new SSL session, minimizing the impact on clients.

**Q.** Will I Get Even Load Balancing if Most Clients Connect to the NLB Cluster Through a Client-Side Proxy?

**A.** When connections coming from different clients pass through the same client-side proxy, they reach the NLB Cluster bearing the same source IP Address. This makes the NLB cluster believe that these connections are coming from a single machine. So, if the cluster is configured in Single Affinity mode, NLB will use only the Source IP Address of the incoming connections to achieve load balancing and since all these connections seem to come from the same IP address, they will all end up with the same host in the cluster. However, if the cluster is configured in No Affinity mode, NLB will use both the Source IP Address and the Source Port to achieve the load balancing, and so load will be distributed amongst all of the hosts.

**Q.** What is the Difference Between Multicast and Unicast Modes of Operation? Which One Should I Use in My NLB Installation?

**A.** NLB relies on the fact that incoming packets are directed to all cluster hosts and passed up to the NLB driver for filtering. In its default *unicast* mode of operation, this is achieved by NLB reassigning the station (MAC) address of the network adapter for which it is enabled and all cluster hosts are assigned the same MAC (media access control) address. In *multicast* mode, NLB assigns a layer-2 multicast address to the cluster adapter instead of changing the adapters station address.

Both modes of operation have their pros and cons. The advantages of *unicast* mode are that it works seamlessly with all routers and layer-2 switches (and layer-3 switches which are configured to operate in layer-2 mode). The disadvantages are:

- *Unicast* mode induces switch flooding, where all switch ports are flooded with NLB traffic, even ports to which non-NLB servers are attached;

- Since all hosts in the cluster have the same IP Address and the same MAC Address, there is no inter-host communication possible between the hosts configured in *unicast* mode.

*Multicast* allows inter-host communication because it adds a layer two multicast address to the cluster instead of changing it and this makes inter-host communication possible as the hosts retain their original unique MAC addresses and already have unique Dedicated IP Addresses. However, in *multicast* mode, the ARP reply sent out by a host in the cluster, in response to an ARP request, maps the clusters unicast IP Address to its multicast MAC Address. Such a mapping in an ARP reply is rejected by some routers and so administrator must add a static ARP entry in the router mapping the Cluster IP Address to its MAC Address

**Q.** How do I Reduce Switch Flooding Caused by Network Load Balancing?

**A.** For Windows 2000, the only choice is to isolate the NLB hosts. This can be done in one of two ways. First the hosts can be homed to their own LAN or Virtual LAN. This will work for either unicast or multicast modes. Second, for the multicast mode, the switch ports to which the hosts are attached can mapped to the cluster MAC address via static entries in the switchs Content-Addressable Memory (CAM) table.

Windows Server 2003 introduces a new feature, called *IGMP multicast*, as an additional means of limiting switch flooding. This mode can be used if the switch supports IGMP snooping. Note however, that this mode requires the cluster to be configured in

*multicast* mode.

**Q.** Does NLB Require Two Network Cards Per Host?

**A.** NLB does not require more than one network card per host. However, there are several scenarios in which a user may prefer to add another network card:

- **Inter-host communication in unicast mode**

  In *unicast* mode, each host in the cluster has the same IP Address and the same MAC Address making them look identical from a networking perspective. So, *unicast* mode has the side effect of disabling communication among the hosts of the cluster.

- **Separating the front-end traffic from the back-end traffic**

  The network adapter that has NLB bound to it can be used to handle incoming connections and connections to a bank-end database, for example, can be made from a separate back-end network adapter.

**Q.** Can NLB be used with network interface card teaming adapters?

**A.** Network interface card teaming solutions are offered by most major network adapter vendors and provide adapter fault tolerance (AFT). Network interface card teaming permits grouping network adapter ports for a connection to a single physical segment. If connectivity through one port is not working, another port is activated automatically. This operation is transparent to the operating system and other devices on the network.

NLB can be used with certain teamed adapters, however, this requires careful coordination between NLB and the teamed adapter configuration software. Refer to Knowledge Base article INFO: Using NIC Teaming Adapters with Network Load Balancing May Cause Network Problems (http://go.microsoft.com/fwlink/?LinkId=18366), for details.

**Q.** Can NLB be used instead of network interface card teaming software to provide adapter fault tolerance or increase throughput?

**A.** No, NLB cannot be used to team adapters (grouping adapter ports together). NLB is used strictly to load balance traffic among *multiple* hosts in a load-balanced cluster

**Q.** How Do I Configure NLB with Layer 2 Switches?

**A.** If you are connecting NLB hosts to a switch rather than a hub, you need to make sure that the switch does not associate the cluster MAC address with a particular switch port. Knowledge Base article Configuration Options for WLBS Hosts Connected to a Layer 2 Switches (http://go.microsoft.com/fwlink/?LinkId=18367) explains how to configure NLB with Layer 2 switches.

**Q.** How Do I Configure NLB with Layer 3 Switches?

**A.** Layer 3 switches need to be specially configured to work with NLB. A VLAN must be established for the hosts in the cluster, and this VLAN must be configured to operate in Layer 2 mode. All Layer 3 switches may not support this capability, and when they do, the mechanism to setup the Layer-2 VLAN is specific to the particular make and model. Consult the documentation for the switch before attempting to configure such a system.

**Q.** How Do I Remove the Switch as a Single Point-of-Failure?

**A.** Create a subnet that spans two switches and home half of the NLB cluster to each switch. The objective is typically to create a LAN with no single points of failure. To this end you will also need two router uplinks and have one switch uplink to one router. You also need to run a redundant routing protocol such as HSRP for failover of the router links. Finally you will need two cross-over links between the switches (otherwise the cross-over cable is a single point of failure). Each switch has two paths off the network: one via the router it is linked to (which is used by default) and one through the other switch. In this configuration, the loss of a router (or link to a router) causes the affected switch to use the cross link to ship its traffic off the LAN. The loss of a switch cuts the cluster capacity in half.

**Q.** I Have Two Network Adapters on Each Server in My NLB Cluster. How Do I Ensure That All Outbound Traffic Goes Through Non-Load-Balanced Network Adapters?

**A.** Sometimes it is desirable for performance or other reasons to direct all outgoing traffic through a different network adapter that the one that is being load balanced with NLB. This implies that there is more than one network adapter on each host in a cluster: NLB is bound to one network adapter called cluster network interface card, and the other network adapter does not have NLB bound to it. To make sure that the outbound traffic leaves each host through the non-cluster network adapter, do the following:

Set the metric on the cluster network adapter to a higher value than the non-cluster network adapter. For example, if you have two network adapters on each host, set the non-cluster network adapter metric to 1 and cluster network adapter metric to 2. The network adapter with a higher metric means it is more expensive to use than the other one with a lower metric. That will ensure that the outbound traffic will be routed out of the non-cluster network adapter.

If you want to use default gateways on both cluster and non-cluster network adapters, make sure the metric of the default gateway on the cluster network adapter has a higher value than the one on the non-cluster network adapter. If you do not want to route any outgoing traffic out of the cluster network adapter, you should not specify the default gateway for it at all.

**Q.** Can I Have Part of the Cluster Operate in Multicast Mode and the Other in Unicast Mode?

**A.** No, the entire cluster has to be in one mode of operation.

**Q.** Does NLB Support Multiple Virtual IP Addresses?

**A.** Yes, NLB supports multiple, virtual IP addresses. For more details please refer to KB article How to Configure WLBS with Multiple Virtual IP Addresses, (http://go.microsoft.com/fwlink/?LinkId=18368).

**Q.** Is it Possible to Specify Different Port Rules for Different Virtual IP Addresses (VIPs) On the Same Set of Hosts?

**A. Windows 2000:** No, port rules cover all VIPs that are configured. In addition, all of the hosts must have the same VIPs on them, see the KB article How to Configure WLBS with Multiple Virtual IP Addresses, (http://go.microsoft.com/fwlink/?LinkId=18368).

**Windows Server 2003:** *Virtual Clusters* provide per-IP Port Rules capability. Now, it is possible to have different ports rules for different VIPs on the same set of machines.

**Q.** Will NLB Work if Multiple VIPs on the Same Set of Hosts are Added on Different Subnet?

**A.** Yes.

**Q.** Is it Possible to Mix Windows NT 4.0 WLBS, Windows 2000 WLBS and Windows Server 2003 in the Same Cluster?

**A.** Yes, mixing Windows NT 4.0 WLBS, Windows 2000 NLB and Windows Server 2003 is supported. There is no additional work needed and the heartbeat packets from NLB in Windows Server 2003 are backward compatible with WLBS on Windows NT 4.0 and Windows 2000 NLB.

| **Note** |
| --- |
| In mixed mode, you cannot use new features of Windows Server 2003 NLB. For a list of new features of Windows Server 2003 NLB refer to the following document: What's New in Clustering Technologies, (http://go.microsoft.com/fwlink/?LinkId=18369). |

You cannot use Network Load Balancing manager to manage mixed-mode clusters.

**Q.** Is it Possible to Bind NLB to Multiple Interfaces?

**A.** NLB may be bound to multiple interfaces only in Windows Server 2003. The interfaces may not be joined to the same cluster.

**Q.** Can I Have Two NLB Clusters on the Same Subnet?

**A.** Yes. In a switched environment, the clusters will compete for the inbound bandwidth because inbound traffic is flooded to all hosts on the LAN. It is important to keep an eye on the aggregate amount of consumed inbound bandwidth to ensure that the network does not become a bottleneck. Other than that, there is no reason why multiple clusters can't reside on the same subnet.

**Q.** We Need to Span a Cluster Across Buildings. Can We Use NLB to Load-Balance Them?

**A.** Yes, if hosts are part of the same subnet.

**Q.** Does NLB Support Token Ring Networks?

**A.** No. Token Ring networks do not allow multiple hosts to share a common MAC address, which is a fundamental requirement of the NLB algorithm.

# Set up and Configuration

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** Can NLB Clusters Be Configured and Managed From a Single Point?

**A.** A new tool in Windows Server 2003 called the *Network Load Balancing Manager (NLB Manager)* provides a single point of configuration and management of NLB clusters. Some key features of the NLB Manager:

- Creating new NLB clusters and automatically propagating Cluster Parameters and Port Rules to all hosts in the cluster and propagating Host Parameters to specific hosts in the cluster.
- Adding and removing hosts to and from NLB clusters.
- Configuring NLB to load balance multiple web sites or applications on the same NLB Cluster, including adding all Cluster IP Addresses to TCP/IP and controlling traffic sent to specific applications on specific hosts in the cluster.
- Diagnosing misconfigured clusters.

| **Note** |
| --- |
| You cannot use Network Load Balancing Manager to manage Windows 2000 clusters or mixed clusters, including clusters containing only hosts running Windows 2000 Server operating systems or clusters containing both Windows 2000 and Windows Server 2003 Server operating systems. |

**Q.** Can I Use WMI to Remotely Bind and Configure NLB?

**A.** No.

**Q.** Can I Use WMI to Remotely Add an IP Address to TCP/IP?

**A.** Yes, call the following method:

- Namespace: root\cimv2
- Class name: Win32_NetworkAdapterConfiguration
- Method name: EnableStatic

**Q.** What is the MaskSourceMAC Registry Key Used For?

**A.** To work correctly, every packet directed to the virtual IP address must reach all NLB hosts. Each NLB host then filters those packets. When the cluster network adapter of each NLB host connects to a port on a Layer 2 switch, the switch tries to determine the MAC address of the hosts connected to each of its ports. As a result, the switch can associate a MAC address with a given port. Ethernet switches send frames to a MAC address through the port associated with the destination MAC address. If a switch associates the clusters MAC addresses with one of its ports, NLB cannot correctly load-balance the traffic because the traffic will reach only a single NLB host.

In unicast mode NLB utilizes a technique called *MAC address masking*. NLB substitutes MAC addresses when it sends packets through the switch. The switch maps these substitute MAC address with individual ports, but does not associate the cluster MAC address with any port and therefore sends the traffic to the cluster MAC address to all the ports in the switch. Cluster MAC address masking prevents the switch from learning the association of cluster MAC address with a particular port. If a switch does not have a MAC address associated to a port, it sends the frames to all ports. This is known as *port flooding*. MAC address masking is possible only when the MaskSourceMAC registry key is set to its default value of one. When it is set to 0, NLB does not perform automatic MAC address masking. Therefore, other techniques must be used to prevent a Layer 2 switch from learning the cluster MAC address.

**Q.** Where Can I Learn More About NLB Setup and Configuration Options?

**A.** There are many resources to help you out with setting up and configuring NLB clusters.

First, you should read the Microsoft Windows Server 2003 Deployment Kit (http://go.microsoft.com/fwlink/?LinkId=18370) chapters that explain how to design and deploy NLB clusters: Chapter 10 Designing Network Load Balancing and Chapter 11 Deploying Network Load Balancing of the Planning Server Deployments section.

Once you have familiarized yourself with designing and deployment of NLB clusters, you can use a step-by-step checklist that allows you to set up and configure NLB clusters. The checklist is located in the Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=18371) (all versions) online help on your CD, or on the Internet at: http://go.microsoft.com/fwlink/?LinkId=18371.

There are many white papers and helpful documentation at Clustering Services in Windows Server 2003, (http://www.microsoft.com/windowsserver2003/technologies/clustering/default.mspx). There you can find white papers that provide a general overview of Load Balancing, the NLB architecture, configuration information for specific scenarios, troubleshooting and much more.

You should also refer to online Help as often the information you need is located right there.

If you have specific questions that are not answered in the resources listed here, please visit the Clustering Communities page and post your question on one of the NLB or clustering newsgroups. For more information about Microsoft communities, see the Microsoft communities Web site, (http://go.microsoft.com/fwlink/?LinkId=18374).

# Management and Operations

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** How Can I Keep a Record of NLB Manager Activities?

**A.** You can configure Network Load Balancing Manager to log each Network Load Balancing Manager event. This log can be very useful in troubleshooting problems or errors when using Network Load Balancing Manager. Enable Network Load Balancing Manager logging by clicking **Log Settings** in the Network Load Balancing Manager **Options** menu. Check the **Enable logging** box and specify a name and location for the log file.

The Network Load Balancing Manager log file contains potentially sensitive information about the Network Load Balancing cluster and hosts, so it must be properly secured. By default, the log file inherits the security settings of the directory in which it is created, so you might have to change the explicit permissions on the file to restrict read and write access to those individuals who do not need full control of the file. Be aware that the individual using Network Load Balancing Manager does require full control of the log file.

**Q.** Can I Manage an NLB Cluster Remotely Using WLBS.EXE?

**A.** Yes, but this is not recommended for security reasons. WLBS.EXE can be used to control participation of nodes in the cluster, but it cannot be used to alter the configuration of the cluster. In Windows Server 2003, use NLB Manager to configure and manage NLB clusters.

**Q.** I am Seeing NLBS Event Number ### and am Not Sure What To Do About it. Where Can I Find Detailed Information About NLB Events in the Event Log?

**A.** The document containing detailed descriptions of all NLB events, their causes and possible resolutions can be found at http://www.microsoft.com/technet/support/ee/ee_advanced.aspx.

**Q.** How Do I Deal with Denial of Service (DOS) Attacks on my NLB Cluster?

**A.** NLB utilizes TCP/IP Denial of Service attack protection to protect itself from denial of service attacks. Specifically, NLB is prone to SYN attacks because of the state that it maintains for TCP connection affinity. To protect an NLB cluster from SYN attacks, make sure that TCP/IP SYN attack protection is enabled. See the following KB articles that explain how to enable TCP/IP DOS attack protection: HOW TO: Harden the TCP/IP Stack Against Denial of Service Attacks in Windows 2000, (http://go.microsoft.com/fwlink/?LinkId=18377) and HOW TO: Harden the TCP/IP Stack Against Denial of Service Attacks in Windows Server 2003, (http://go.microsoft.com/fwlink/?LinkId=18376).

NLB can also fail because of heartbeat loss due to insufficient network resources. This may cause partitioning of the cluster, which can cause connectivity problems for clients, as clients may be serviced by multiple NLB hosts. To soften the heartbeat loss criteria in NLB, use the NLB WMI interface to change the AliveMsgTolerance (from 5 to 100; default is 5) and/or the AliveMsgPeriod (from 100 to 10000 in milliseconds; default is 1000). These changes can cause NLB to broadcast heartbeats less often (back-off) and/or allow for more heartbeat loss before logging an event.

| **Note** |
| --- |
| Changing the AliveMsgPeriod requires a re-load of the NLB driver (bind/unbind or reboot); changing AliveMsgTolerance requires only an NLB reload operation. |

**Q.** How Do I Secure My NLB Cluster?

**A.** NLB assumes that the LAN to which it is homed is trusted. Ensuring that only trusted machines are homed to this LAN is the first step to protecting an NLB cluster. As an additional measure, ensure that NLB remote control is disabled on all hosts. If it must be enabled, then it is vital that the firewall or router block traffic to UDP port 1717 and UDP port 2504 except to those clients that need access.

**Q.** How do I configure my cluster to handle load non-uniformly?

**A.** NLB distributes load among hosts within the cluster according to the load's associated port rule. By default, a port rule distributes an equal share of the load to each host. To configure a host to handle more or less than an equal share of the load, edit the port rule to clear the "Equal" load weight check box and enter a load weight number between 1 and 100.

The load weight is a relative weight among converged nodes in the cluster that indicates how much load a host should handle. To determine the percentage load a host will handle, divide that host's load weight by the sum of the load weights from the

other hosts in the cluster for the same port rule. (A port rule with the "Equal" checkbox checked has a load weight of 50.)

For example, consider the case of a two-host cluster. If Host 1's port rule has a load weight of 30 and Host 2's port rule has a load weight of 10, Host 1 will handle 30/(30+10) or 75% of the load. Host 2 will handle 25% of the load.

**Caution for Windows 2000 only:** Modifying the load weight of a port rule on a host causes all current connections handled by that host to reset. This includes connections associated with other port rules, but does not include traffic addressed to the dedicated IP address, if any. This caution does not apply to Windows Server 2003 and later. (Changing any other property of the cluster resets all current connections on that host for all versions of Windows.)

# Support for Sessions

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** Does NLB Support Sticky Sessions?

**A.** NLB supports limited forms of session stickiness:

- **Single Affinity.** When Single affinity is enabled, the clients port number isn't used and the mapping algorithm uses the clients full IP address to determine load distribution. As a result, all requests from the same client always map to the same host within the cluster. Because there is no time-out value (typical of dispatcher-based implementations), this condition persists until cluster membership changes. However, session stickiness may not be preserved if the cluster membership changes, in particular when a new host is added to the cluster, or a host is re-added after being temporarily out of service.

- **Class C Affinity.** As in the case of Single affinity, client port numbers are not used to calculate load distribution. When Class C affinity is enabled, the mapping algorithm bases load distribution on the Class C portion (the upper 24 bits) of the clients IP address. As with Single affinity, sessions may not be preserved if the cluster membership changes.

- **IPSec and VPN Sessions.** On Windows Server 2003, NLB has specialized support for sticky IPSec and VPN sessions. NLB preserves these sessions even if the cluster membership changes.

NLB cannot track the boundaries of UDP streams, since the logical session boundaries are defined by particular applications. Instead, Network Load Balancing affinity settings are used to assist in preserving client sessions. When a cluster host fails or leaves the cluster, its client connections are always dropped. After a new cluster membership is determined by *convergence* (see How Does NLB cluster convergence work?), clients that previously mapped to the failed host are remapped across the surviving hosts. All other client sessions are unaffected by the failure and continue to receive uninterrupted service from the cluster. In this manner, Network Load Balancing load balancing algorithm minimizes disruption to clients when a failure occurs.

When a new host joins the cluster, it induces convergence, and a new cluster membership is computed. When convergence completes, a minimal portion of the clients will be remapped to the new host. NLB tracks TCP connections on each host, and, after their current TCP connection completes, the next connection from the affected clients will be handled by the new cluster host; UDP streams are immediately handled by the new cluster host. This can potentially break some client sessions that span multiple connections or comprise UDP streams, or application sessions that span TCP sessions. Hence, hosts should be added to the cluster at times that minimize disruption of sessions. To completely avoid this problem, session state must be managed by the server application so that it can be reconstructed or retrieved from any cluster host. For example, session state can be pushed to a backend database server or stored in client cookies.

**Q.** How Does Single Affinity Mode Differ From No Affinity Mode? Which One Should I Use to Load Balance My Application?

**A.** In *Single Affinity* mode, NLB load balances traffic based only on the Source IP Address of the incoming connection. So, *Single Affinity* mode ensures that all TCP connections originating from the same client (IP Address) are sent to the same host in the cluster. This will continue indefinitely until a host is either added or removed from the cluster, at which time the connections originating from that IP address may get mapped to a different host in the cluster.

In *No Affinity* mode, NLB load balances traffic based on Source IP Address and Source Port of the incoming connection request. Therefore, in *No Affinity* mode, multiple connections from the same client may be handled by different hosts in the cluster as long as these connections have different source ports.

Which mode to use really depends on the application being load balanced. If the application makes use of sessions which persist over multiple TCP connections, NLB should be configured in *Single Affinity* mode because you want to make sure that all TCP connections which are part of a single session are mapped to the same host in the cluster.

On the other hand, *No Affinity* allows a better load distribution because it does not require client connections be handled by specific servers. Thus for applications that do not use sessions and for which it is acceptable for multiple incoming connections originating from the same client to be handled by different hosts in the cluster, *No Affinity* would be a better mode of operation.

**Q.** If My Clients Use SSL to Connect to My Front-End Web Servers, Can I Still Use NLB to Cluster These Web Servers?

**A.** Yes. For efficiency reasons, configure the port rule in *Single Affinity* mode. However, note the following side effects:

- Load balancing will not be as precise as in *No Affinity* mode.

- Sessions can break when hosts are either added to or removed from a cluster. Note that an SSL session will be transparently renegotiated by the protocol.

- If all of the sessions come through a client-side proxy, you will not see any load balancing since all of the TCP connections will have the same IP Address and they will all be handled by the same computer in the cluster.

# Application and Protocol Support

Updated: February 10, 2004

Applies To: Windows Server 2003 with SP1

**Q.** Can NLB be Used For Load Balancing Terminal Server Clusters?

**A.** Yes. See Load balancing and terminal servers (http://go.microsoft.com/fwlink/?LinkId=18378), in the Windows Server 2003 online Help.

For detailed information about using Terminal Server Session Directory with load balancing refer to the following white paper: Session Directory and Load Balancing Using Terminal Server, (http://go.microsoft.com/fwlink/?LinkId=18379).

**Q.** Does NLB Support WINS Resolution?

**A.** No, WINS names should not be automatically registered for the IP addresses configured on the NLB interface. WINS registration isn't NLB-aware and as a result, both the dedicated and virtual IP addresses will be registered with the servers name. Note however, that the IP addresses can be mapped statically in WINS.

**Q.** Does NLB Support DNS Resolution?

**A.** NLB does not automatically register the cluster name with DNS. The administrator must directly configure the DNS server. Note however, that the IP addresses can be mapped statically in WINS.

Also note that DHCP should not be used when NLB is bound to any network adapter on a host.

**Q.** Can I Use L2TP/IPSec on a NLB Cluster?

**A.** Yes, in Windows Server 2003 NLB supports both PPTP and L2TP VPN sessions. NLB support for both PPTP and L2TP protocol requires that NLB be configured in single affinity mode. For detailed information about configuring NLB with IPSec see the following Knowledge Base article: How to Configure Network Load Balancing Server Cluster to Work with IPSec, (http://go.microsoft.com/fwlink/?LinkId=18380).

**Q.** How Do I Configure NLB with IPSec?

**A.** Refer to the following Knowledge Base article: How to Configure Network Load Balancing Server Cluster to Work with IPSec, (http://go.microsoft.com/fwlink/?LinkId=18380).

**Q.** Can I Use Kerberos with Applications Load-Balanced by NLB? How?

**A.** Yes.

**Q.** Can I Use NLB with Host Header Names?

**A.** Yes, ensure that all host header names are hosted by every node in the NLB cluster.

**Q.** Can I Use NLB with .NET Remoting?

**A.** .NET Remoting pipelines method invocations from client to server over a single TCP connection, meaning that once a connection is established, it is reused for subsequent method invocations and is closed only after the connection remains idle for a pre-configured amount of time. NLB can load-balance these connections, but load-balancing will likely be coarse because the unit of load-balancing is the TCP connection, not a method invocation Since one client gets pinned to a specific server, load will appear well distributed only if you have many clients connected to the cluster at the same time. Each client will get load-balanced but the connection will stay open for a very long time.

**Q.** When Using .NET Remoting and NLB, Method Invocations From One Client are Against the Same Host and Are Not Load-Balanced. Why?

**A.** See question Can I use NLB with .NET Remoting?

**Q.** I Use NLB to Load-Balance COM+ but Load is Not Being Distributed. Why?

**A.** COM+ uses connection pooling on the client, which results in method invocations that are pipelined over existing TCP connections once established. NLB load-balances TCP connections so COM+ requests are not distributed among hosts. Instead of NLB, use COM+ load-balancing which was introduced to address this issue.

**Q.** Can I Use NLB as a Firewall?

**A.** NLB is not designed to be used as a firewall. Use a certified firewall solution in front of the cluster just as you would if the hosts were not part of a cluster.

**Q.** Can I Load-Balance NetBIOS Traffic?

**A.** Yes, it is possible, though not recommended for File and Print Services. Ensure that both of the following are true:

1. NetBIOS requires that the first IP address in TCP/IP must be used to access shares, so this must be the cluster IP address (this means you can't have a dedicated IP address). See the KB article System Error 53 Has Occurred Using Windows NT Load Balancing Service (http://go.microsoft.com/fwlink/?LinkId=18381) for further discussion.

2. Set the DisableStrictNameChecking registry value using the following instructions:

   - Start Registry Editor (Regedit.exe).
   - Locate and click the following key in the registry:
     HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\LanmanServer\Parameters
   - On the **Edit** menu, click **Add Value**, and then add the following registry value:

     Value name: DisableStrictNameChecking

     Data type: REG_DWORD

     Radix: Decimal

     Value: 1
   - Quit Registry Editor.
   - Open a command window.
   - Run net stop server.
   - Run net start server.

This solution is documented in the KB article Connecting to SMB Share on a Windows 2000-Based Server May Not Work with an Alias Name (http://go.microsoft.com/fwlink/?LinkId=18382).

**For More Information**

This is a preliminary document and may be changed substantially prior to final commercial release of the software described herein. The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

# Server Clusters: Frequently Asked Questions for Windows 2000 and Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: January 1, 2003

**In This White Paper**

How is this Document Organized?

What are Most Frequently Asked Questions?

Introduction to Server Clusters

Availability

Manageability and Deployment

Scalability (Server Clusters: Frequently Asked Questions for Windows 2000 and Windows Server 2003)

Server cluster Concepts (Server Clusters: Frequently Asked Questions for Windows 2000 and Windows Server 2003)

Server Requirements

Interconnects

Storage (Server Clusters: Frequently Asked Questions for Windows 2000 and Windows Server 2003)

Storage Area Networks (SAN) questions

Network Attached Storage (NAS) questions

Highly Available File Servers

Highly Available Print Servers

Removing all Single Points of Failure

Active Directory, DNS and Domain Controllers

Security Considerations (Server Clusters: Frequently Asked Questions for Windows 2000 and Windows Server 2003)

In-the-box HA Services

Geographically Dispersed Clusters

Quorum

Cluster-Aware Applications

Miscellaneous Topics

# How is this Document Organized?

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** How is this Document Organized?

**A.** The various sections cover different aspects of server clusters from concepts to deployment to operational procedures:

- Introduction to Server clusters
- Availability
- Manageability and deployment
- Scalability
- Server cluster concepts
- Server requirements
- Interconnects
- Storage
- Highly available file servers
- Highly available print servers
- Removing single points of failure
- Active directory and domain controllers
- Security Considerations
- In-the-box HA services
- Geographically dispersed clusters
- Quorum
- Cluster-aware applications
- Miscellaneous topics

# What are Most Frequently Asked Questions?

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What are Most Frequently Asked Questions?

**A.**

- How do I know if my Server cluster is a supported configuration?
- How many servers can there be in a single cluster?
- How do I benefit from more than two nodes in a cluster?
- Are dynamic disks supported in a cluster?
- Can a cluster disk be extended without rebooting?
- How can I replace a disk that has gone bad in a cluster?
- Why is zoning needed to isolate a Server cluster in a SAN?
- Can a cluster server boot from a SAN?
- Is Kerberos authentication possible for services hosted on a cluster?
- Is IIS cluster-aware?
- Can Server clusters span multiple sites?
- Will Microsoft provide a cluster file system?

# Introduction to Server Clusters

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What is a Server Cluster?

**A.** A Server cluster is a collection of independent servers that together provide a single, highly available platform for hosting applications

**Q.** What are the benefits of Server clusters?

**A.** There are three primary benefits to server clustering: availability, manageability, and scalability:

| | |
|---|---|
| **Availability:** | Server clusters provide a highly available platform for deploying applications. A Server cluster ensures that applications continue to run in the event of planned downtime due to maintenance or unplanned downtime due to failures. Server clusters protect against hardware failures, failures of the Windows operating system, device drivers or application software. Server clusters allow operating system and application software to be upgraded across the cluster without having to take down the application. |
| **Manageability:** | Server clusters allow administrators to quickly inspect the status of all cluster resources and move workloads around onto different servers within the cluster. This is useful for manual load balancing, and to perform "rolling updates" on the servers without taking important data and applications offline. |
| **Scalability:** | Applications that can be partitioned can be spread across the servers of a cluster allowing additional CPU and memory to be applied to a problem. As the problem size increases, additional servers can be added to the cluster. A partitioned application is one where the data (or function) can be split up into independent units. For example, a customer database could be split into two units, one covering customers with names beginning A thru L and the other for customers with names beginning M thru Z. |

**Q.** What are Server clusters used for?

**A**. Server clusters provide a highly available platform for mission critical, line of business applications. These typically include database servers (such as Microsoft SQL Server), mail and collaboration servers (such as Exchange Server) and infrastructure servers such as file servers and print servers.

**Q.** Is Server cluster the same as Microsoft Cluster Server (MSCS) and/or Wolfpack?

**A.** The code name for the original Server cluster project was Wolfpack. When the product was released as part of Microsoft Windows NT 4.0 Enterprise Edition it was named Microsoft Cluster Server or MSCS. In Microsoft Windows 2000, the official name of the product became Server cluster.

**Q.** How do Server clusters fit into the larger Windows high availability story?

**A.** Server clusters are part of a much larger story for highly available applications running on Microsoft products. Windows has two cluster technologies, server clusters and Network Load Balancing (NLB). Server clusters are used to ensure that *stateful* applications such as a database (e.g., SQL Server) can be kept highly available by *failing over* the application from one server to another in the event of a failure. NLB is used to distribute client requests across a set of identical servers. NLB is particularly useful for ensuring that *stateless* applications such as a web server (e.g., IIS) are highly available and can be *scaled-out* by adding additional servers as the load increases.

In addition, Microsoft Application Center 2000 includes Component Load Balancing (CLB) CLB is a technology used to ensure COM+ applications are highly available and can scale-out.

The Windows platform also uses other technologies to ensure high availability and scalability. For example, Active Directory is made highly available through replication of the directory contents across a set of domain controllers.

**Q.** What is Scale-up verses Scale-out?

**A.** Scale-up is a term used to describe the scalability provided by large symmetric multi-processing (SMP) systems. As the load increases, more processors and more memory can be added to a single server, increasing the processing power and memory available for applications. Windows 2000 Datacenter Server and Windows Server 2003 Datacenter Edition provide the support for scaling up to large computers with many CPUs and large physical memory. In a scale-out environment, Server clusters are

deployed to ensure that the failure of the server does not cause the applications or services to fail. Scale-up is typically done for either hosting large applications that cannot be partitioned, or to consolidate a set of applications into a single server (or small number of servers in the case of a server cluster).

Scale-out is a term used to describe the idea that an application can be scaled by partitioning the workload and spreading it across a set of cooperating servers. If the load increases, additional servers can be added to the set to provide additional processing power and memory.

Server clusters are particularly useful for enhancing the availability of applications that can be scaled out across a set of nodes by allowing the individual pieces or partitions of a scale-out solution to be made highly available via failover.

**Q.** Is Server cluster a standard feature of all Windows operating system products?

**A.** Server clusters are not available in all Windows operating system products. Server clusters are available in:

- Windows NT Server, Enterprise Edition
- Windows 2000 Advanced Server and Datacenter server
- Windows Server 2003 Enterprise Edition and Datacenter Edition
- Windows Server 2003, 64 bit Enterprise Edition and Datacenter Edition

Server clusters are also available as part of the Server Appliance Kit, available to OEMs for building embedded solutions based on the Windows operating system.

**Q.** How do I know if my Server cluster is a supported configuration?

**A.** All server clusters must be qualified to be supported by Microsoft. A qualified configuration has undergone extensive testing using a hardware compatibility test provided by Microsoft. All qualified solutions appear on the Microsoft Hardware Compatibility List (HCL) (http://go.microsoft.com/fwlink/?LinkID=67738). Only cluster solutions listed on the HCL are supported by Microsoft.

The complete cluster solution must be listed on the *Cluster* HCL list. The complete solution includes the servers, the storage adapters, the interconnect type, the storage controllers firmware and driver versions. All of the components must match exactly, including any software, driver or firmware versions for the solution to be qualified.

The HCL contains a set of qualified cluster components. A solution built from qualified components does NOT imply that the solution is qualified.

**The cluster component lists have been a source of confusion in the past and we will be removing the cluster component lists (such as Cluster/RAID) from the HCL for Windows Server 2003.**

**Q.** Where are supported configurations listed?

**A.** All qualified solutions appear on the Microsoft Hardware Compatibility list (HCL) (http://go.microsoft.com/fwlink/?LinkID=67738). Only cluster solutions listed on the HCL are supported by Microsoft.

**Q.** Where can I find more documentation on Server clusters?

**A.** Server clusters are extensively covered in the online help in Windows 2000 and Windows Server 2003. Additional information can be found on the Microsoft web site at:

- http://go.microsoft.com/fwlink/?LinkID=67742
- http://go.microsoft.com/fwlink/?LinkID=67740
- http://go.microsoft.com/fwlink/?LinkID=67741

# Availability

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** Can Server clusters provide zero downtime for applications?

**A.**

- No. Server clusters can dramatically reduce planned and unplanned downtime. However, even with Server clusters, a server could still experience downtime from the following events:

- *Failover time:* If a Server cluster recovers from a server or application failure, or if it is used to move applications from one server to another, the application(s) will be unavailable for a non-zero period of time (typically under a minute).

- *Failures from which Server clusters can not recover:* There are types of failure that Server clusters do not protect against, such as loss of a disk not protected by RAID, loss of power when a UPS is not used, or loss of a site when there is no fast-recovery disaster recovery plan. Most of these can be survived with minimal downtime if precautions are taken in advance.

- *Server maintenance that requires downtime:* Server clusters can keep applications and data online through many types of server maintenance, but not all (for example: installing a new version of an application which has a new on-disk data format that requires reformatting preexisting data).

Microsoft recommends that clusters be used as one element in customers' overall programs to provide high integrity and high availability for their mission-critical server-based data and applications.

**Q.** Which types of applications and services benefit from Server clustering?

**A.** There are three types of server applications that will benefit from Server clusters:

1. "*In the box*" *services provided by the Windows platform:* For example: File shares, print queues, Microsoft Message Queue Server (MSMQ) services, and Component Services (formerly known as Transaction Server) services.

2. *Generic applications:* Server clusters include a point-and-click wizard for setting up any well-behaved server application for basic error detection, automatic recovery, and operator-initiated management (e.g., move from one server to the other). A "well behaved" server application is one which keeps a recoverable state on cluster disk(s), and whose client can gracefully handle a pause in service as the application is automatically re-started.

3. *Cluster-aware applications:* Software vendors test and support their application products on Server clusters.

# Manageability and Deployment

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** Do Server clusters provide a single image view for administration?

**A.** Yes, Server clusters provide administrators with a graphical interface (Cluster Administrator) and a command line tool (Cluster.exe available on Windows 2000 and Windows Server 2003) from which they can monitor and manage all of the resources in a cluster as if it was a single system. The administrator can:

- monitor the status of all servers and applications in the cluster
- setup new applications, file shares, print queues, etc. for high availability
- administer the recovery policies for applications and resources
- take applications offline, bring them back online, and move them from one server to another

The ability to graphically move workload from one server to another with only a momentary pause in service (typically less than a minute), means administrators can easily unload servers for planned maintenance without taking important data and applications offline for long periods of time.

**Q.** Can Server clusters be managed remotely?

**A.** Yes, an authorized cluster administrator can run Cluster Administrator from any Windows NT or Windows 2000 Workstation or Server, any Windows XP machine or any Windows Server 2003 on the network. The cluster.exe command line tool can be run remotely from any Windows 2000 Workstation or Server, any Windows XP machine or any Windows Server 2003 on the network.

**Q.** Can Server clusters be setup and configured remotely?

**A.** Yes, with Windows Server 2003, Cluster administrator and the cluster.exe command line tool can also be used to remotely setup and configure a Server cluster (e.g., create a new cluster, add a new server to an existing cluster or remove servers from a cluster).

**Q**. Is the server cluster software installed by default?

**A.** On Windows NT and Windows 2000, the Server cluster software is not installed by default. On Windows NT you must install the Enterprise Edition CD. With Windows 2000, you must install the Server cluster software using the optional component management tool.

With Windows Server 2003, the Server cluster software is installed by default on Enterprise Edition and Datacenter Edition. The cluster software is NOT configured by default and will not start until the server is either used to create a new cluster or the server is added to an existing cluster. In Windows Server 2003, you must use Cluster Administrator or the Cluster.exe command line tool to configure a cluster.

**Q.** Can a Server cluster be installed and configured using unattended installs?

**A.** Yes, to perform an unattended installation for a Windows 2000 Cluster using an Uttend.txt file, please consult the Windows 2000 Deployment Guide located on the Windows 2000 CD-ROM in the following location: %CDROM%\SUPPORT\TOOLS\DEPLOY.CAB\Unattend.doc

Additional information can be found in article 176507 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=67745).

**Q.** Do applications have to be installed separately on all the cluster servers?

**A.** It depends on the application. Each application that can be failed over between servers in a cluster must be available on all servers in the cluster. Historically, application setup has not been cluster-aware and therefore the application must be installed separately on each server.

With more recent versions of some applications, for example SQL Server 2000, the application setup is aware of Server clusters. In the case of Microsoft SQL Server 2000 setup, the appropriate files are copied to all servers and registry settings and other SQL Server configuration is done just once.

**Q.** How can I take advantage of extensibility features of ISA Server?

**A.** A number of third-party vendors offer solutions such as virus detection, content filtering, site categorization, reporting, and administration. Customers and developers also have the ability to create their own extensions to ISA Server. ISA Server includes a comprehensive software development kit for developing tools that build on ISA Server firewall, caching, and management features.

**Q.** Are cluster management operations scriptable?

**A.** Yes, there are a number of options for scripting management operations:

- The Cluster.exe command line tool allows command files to be built that can change the state of cluster resources, move resources around etc.. For more information about the command line, see the online help or type cluster at the command line on a server that has Server clusters installed or has the Windows server admin pack installed.

- In Windows 2000 and Windows Server 2003, there is a scriptable COM interface to the Server cluster APIs. This allows VBScript or any other scripting language supported by Windows to call the cluster APIs. These APIs provide a way to change the state of the cluster as well as return data about the resources or applications in a cluster. For more information about the cluster server APIs, see the Platform SDK, it has comprehensive documentation for the cluster APIs and the COM (automation server) APIs.

- In Windows Server 2003, there is a WMI provider for Server clusters that allows WMI scripts to manage the cluster. For more information about the cluster server WMI schema, see the Platform SDK.

**Q.** Can command line tools be used to manage a cluster?

**A.** Yes, Server clusters have a command line tool cluster.exe that can be used to manage the cluster. With Windows Server 2003, this command line tool can also be used to create new clusters, add servers to existing clusters and remove servers from clusters. Almost all of the functionality provided by Cluster Administrator is available through cluster.exe in Windows Server 2003.

**Q.** Is WMI supported for managing a cluster?

**A.** Yes, in Windows Server 2003, there is a WMI provider that allows a Server cluster to be managed. In addition, all the Server cluster events (such as server up and server down, resource online, offline, failed etc.) are available through WMI events.

**Q.** Does Server clusters support WMI eventing?

**A.** Yes, in Windows Server 2003, all of the Server cluster events (such as server up and server down, resource online, offline, failed etc.) are available through WMI events.

**Q.** Can Group Policy be used on cluster servers?

**A.** You can apply Group Policy to cluster servers. There are a few things to remember. Applications failover from one server to another in a cluster and will typically expect the same policies to be in effect no matter where they are hosted. You should ensure that all cluster servers have the same set of group policies applied to them.

With Windows Server 2003 and Windows 2000 SP3 onwards, virtual servers can be published in active directory. You should NOT apply group policies to a virtual server computer object. The policy will only be applied to the server that is currently hosting the virtual server and all policy settings may be lost if the application fails over.

**Q.** Is Systems Management Server (SMS) supported for deploying applications in a Server cluster?

**A.** No, at this time, SMS is not supported for deploying applications across a cluster.

**Q.** Is Application Center 2000 supported on a Server cluster?

**A.** No, Application Center 2000 provides a set of features that are intended for managing servers hosting web server front-end and middle-tier business logic. Server clusters are typically deployed to ensure high availability of the back-end database engine.

**Q.** Are MOM management packs available for monitoring Server clusters?

**A.** No, at this time, there are no server-aware MOM management packs.

**Q.** How do server clusters help administrators perform rolling upgrades?

**A.** With Server clusters, server administrators no longer have to do all of their maintenance within those rare windows of opportunity when no users are online. Instead, they can simply wait until a convenient off-peak time when one of the servers in

the cluster can be removed for maintenance and the workload distributed across the other cluster servers. By pointing-and-clicking in Cluster Administrator (or using scripts) to move all of the workload onto other servers, one of the servers in the cluster can be upgraded. Once the upgrade is complete and tested, that server can be rebooted where it will automatically re-join the cluster, ready for work. When convenient, the administrator repeats the process to perform maintenance on the other servers in the cluster. This ability to keep applications and data online while performing server maintenance is often referred to as doing "rolling upgrades" to your servers.

Server clusters allows rolling upgrades for a new release and the previous version.

**Q.** What combination of Windows versions supports rolling upgrades?

**A.** Server clusters allow rolling upgrades for a new release and the previous version. The following rolling upgrades are supported:

- Windows NT Enterprise Edition to Windows 2000
- Windows 2000 to Windows Server 2003

**Q.** Can I manage different versions of Server cluster from a single Cluster Administrator tool?

**A.** Yes, the Server cluster tools allow a mixed version cluster to be managed from a single point. During a rolling upgrade, a cluster may contain different versions of the operating system. As an example Windows 2000 and Windows Server 2003 can co-exist in the same cluster. The administration tools provided with Windows Server 2003 allow such a mixed-version cluster to be managed.

**Q.** Are there other manageability benefits with Server clusters?

**A.** Server clusters allow administrators to quickly inspect the status of all cluster resources, and move workload around onto different servers within the cluster. This is useful for manual load balancing, and to perform "rolling updates" on the servers without taking important data and applications offline.

# Scalability (Server Clusters: Frequently Asked Questions for Windows 2000 and Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** Do Server clusters enhance server scalability?

**A.** The primary goal of Server clusters is to provide a highly available platform for applications.

There are some types of applications which can be scaled-out to take advantage of a set of machines. More machines can be added to the set to increase the CPU or memory available to the application. Applications of this type are typically scaled by partitioning the data set. For example, a SQL Server database may be scaled out by partitioning the database into pieces and then using database views to give the client applications the illusion of a single database.

Server clusters do not provide facilities to partition or scale-out an application; however, Server clusters allow such applications to be scaled-out in a highly available environment. Each of the partitions can be failed over independently within the set of machines (or to additional spare nodes) so that in the event of failure, a partition of the application remains available.

**Q.** Do Server clusters provide load balancing of applications?

**A.** Server clusters provide a manual load balancing mechanism. Applications can be moved around the servers of a cluster to distribute the load.

There are no automated tools provided with the operating system to automatically load balance applications and no enhanced failover policies that use load to determine the best placement for an application.

# Server cluster Concepts (Server Clusters: Frequently Asked Questions for Windows 2000 and Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What hardware do you need to build a Server cluster?

**A.** The most important criteria for Server cluster hardware is that it be included in a validated Cluster configuration on the Microsoft Hardware Compatibility List (HCL), indicating it has passed the Microsoft Cluster Hardware Compatibility Test. All qualified solutions appear on the Microsoft HCL (http://go.microsoft.com/fwlink/?linkid=67738). Only cluster solutions listed on the HCL are supported by Microsoft.

In general, the criteria for building a server cluster include the following:

- **Servers:** Two or more PCI-based machines running one of the operating system releases that support Server clusters (see below). Server clusters can run on all hardware architectures supported by the base Windows operating system, however, you cannot mix 32-bit and 64-bit architectures in the same cluster.

- **Storage:** Each server needs to be attached to a shared, external storage bus(es) that is/are separate from the bus containing the system disk, the startup disk or the pagefile disk. Applications and data are stored on one or more disks attached to this bus. There must be enough storage capacity on the shared cluster bus(es) for all of the applications running in the cluster environment. This shared storage configuration allows applications to failover between servers in the cluster.

  Microsoft recommends hardware Redundant Array of Inexpensive Disks (RAID) for all cluster disks to eliminate disk drives as a potential single point of failure. This means using a RAID storage unit, a host-based RAID adapter that implements RAID across disks, etc.

  SCSI is supported for 2-node cluster configurations only. Fibre channel arbitrated loop is supported for 2-node clusters only. Microsoft recommends using fibre channel switched fabrics for clusters of more than two nodes.

- **Network:** Each server needs at least two network cards. Typically, one is the public network and the other is a private network between the two nodes. A static IP address is needed for each group of applications that move as a unit between nodes. Server clusters can project the identity of multiple servers from a single cluster by using multiple IP addresses and computer names: this is known as a *virtual server*.

**Q.** What is a cluster resource?

**A.** A cluster resource is the lowest level unit of management in a Server cluster. A resource represents a physical object or an instance of running code. For example, a physical disk, an IP address, an MSMQ queue, a COM object all of these things are considered to be resources. From a management perspective, resources can be independently started and stopped and each is monitored to ensure that it is healthy.

Server cluster can monitor any arbitrary resource type. This is possible because Server clusters define a resource plug-in model. Each resource type has an associated resource plug-in or *resource dll* that is used to start, stop and provide health information that is specific to the resource type. For example, starting and stopping SQL Server is different from starting and stopping a physical disk. The resource dll takes care of the differences. Application developers and system administrators can build new resource dlls for their applications that can be registered with the cluster service.

Server clusters provides some generic plug-ins that can be used to make existing applications cluster-aware very quickly, known as *Generic Service* and *Generic Application*. With Windows Server 2003, a *Generic Script* resource plug-in was added that allows the resource dll to be written in any scripting language supported by the Windows operating system.

**Q.** What is a resource dependency?

**A.** A complete application actually consists of multiple pieces or multiple resources, some pieces are code and others are physical resources required by the application. The resources are related in different ways; for example, an application that writes to a disk cannot come online until the disk is online. If the disk fails, then, by definition, the application cannot continue to run since it writes to the disk. Resource dependencies can be defined by the application developer or system administrator

to capture these relationships. Resource dependencies define the order that resources are brought online and control how failures are propagated to the various pieces of the application.

**Q.** What is a resource group?

**A.** A resource group is a collection of one or more resources that are managed and monitored as a single unit. A resource group can be started or stopped. If a resource group is started, each resource in the group is started (taking into account any start order defined by the dependencies between resources in the group). If a resource group is stopped, all of the resources in the group are stopped. Dependencies between resources cannot span a group. In other words, the set of resources within a group is an autonomous unit that can be started and stopped independently from any other group. A group is a single, indivisible unit that is hosted on one server in a Server cluster at any point in time and it is the unit of failover.

**Q.** Can I have dependencies between resources in different groups?

**A.** No, resource dependencies are confined to a single group.

**Q.** What is a virtual server?

**A.** A virtual server is a resource group that contains an IP address resource and a network name resource. When an application is hosted in a virtual server, the application can be accessed by clients using the IP address or network name in that resource group. As the resource group fails over across the cluster, the IP address and network name remain the same, therefore the client becomes unaware of the physical location of the application and will continue to work in the event of a failure of one of the servers in the cluster.

**Q.** How can I take advantage of extensibility features of ISA Server?

**A.** A number of third-party vendors offer solutions such as virus detection, content filtering, site categorization, reporting, and administration. Customers and developers also have the ability to create their own extensions to ISA Server. ISA Server includes a comprehensive software development kit for developing tools that build on ISA Server firewall, caching, and management features.

**Q.** What is failover?

**A.** Server clusters monitor the health of the nodes in the cluster and the resources in the cluster. In the event of a server failure, the cluster software re-starts the failed server's workload on one or more of the remaining servers. If an individual resource or application fails (but the server does not), Server clusters will typically try to re-start the application on the same server; if that fails, it moves the application's resources and re-starts it on the other server. The process of detecting failures and restarting the application on another server in the cluster is known as *failover*.

The cluster administrator can set various recovery policies such as whether or not to re-start an application on the same server, and whether or not to automatically "failback" (re-balance) workloads when a failed server comes back online.

**Q.** Is failover transparent to users?

**A.** Server clusters do not require any special software on client computers, so the user experience during failover depends on the nature of the client side of their client-server application. Client reconnection can be made transparent, because the Server clusters software has restarted the applications, file shares, etc. at exactly the same IP address.

If a client is using "state-less" connections such as a standard browser connection, then the client would be unaware of a failover if it occurred between server requests. If a failure occurs while a client is connected to the failed resource, then the client will receive whatever standard notification is provided by the client side of their application when the server side becomes unavailable. This might be, for example, the standard "Abort, Retry, or Cancel?" prompt you get when using the Windows Explorer to download a file at the time a server or network goes down. In this case, client reconnection is not automatic (the user must choose "Retry"), but the user is fully informed of what is happening and has a simple, well-understood method of re-establishing contact with the server. Of course, in the meantime, the cluster service is busily re-starting the service or application so that, when the user chooses "Retry", it re-appears as if it never went away.

**Q.** What is failback?

**A.** In the event of the failure of a server in a cluster, the applications and resources are failed over to another node in the cluster. When the failed node rejoins the cluster (after reboot for example), that node now is free to be used by applications. A cluster administrator can set policies on resources and resource groups that allow an application to automatically move back to a node if it becomes available, thus automatically taking advantage of a node rejoining the cluster. These policies are known as *failback* policies. You should take care when defining automatic failback policies since depending on the application, automatically moving the application (which was working just fine) may have undesirable consequences on the clients using the applications.

**Q.** When an application restarts after failover, does it restore the application state at the time of failure?

**A.** No, Server clusters provide a fast crash restart mechanism. When an application is failed over and restarted, the application

is restarted from scratch. Any persistent data written out to a database or to files is available to the application, but any in-memory state that the application had before the failover is lost.

**Q.** At what level does failover exist?

**A.** At the resource group level.

**Q.** What is a Quorum Resource and how does it help Server clusters provide high availability?

**A.** Server clusters require a quorum resource to function. The quorum resource, like any other resource, is a resource which can only be owned by one server at a time, and for which servers can negotiate for ownership. Negotiating for the quorum resource allows Server clusters to avoid "split-brain" situations where the servers are active and think the other servers are down. This can happen when, for example, the cluster interconnect is lost and network response time is problematic. The quorum resource is used to store the definitive copy of the cluster configuration so that regardless of any sequence of failures, the cluster configuration will always remain consistent.

**Q.** What is active/active verses active/passive?

**A.** Active/Active and Active/Passive are terms used to describe how applications are deployed in a cluster. Unfortunately, they mean different things to different people and so the terms tend to cause confusion.

From the perspective of a single application or database:

- Active/Active means that the same application or pieces of the same service can be run concurrently on different nodes in the cluster. For example SQL Server 2000 can be configured such that the database is partitioned and each node can be running a single instance of the database. SQL Server provides the notion of views to provide a single image of the entire database.

- Active/Passive means that only one node in the cluster can be hosting the given application. For example, a single file share is active/passive. Any given file share can only be hosted on one node at a time.

From the perspective of a set of instances of an application or service:

- Active/Active means that different instances of the same application can be running concurrently on different cluster nodes. For example, each node in a cluster can be running SQL Server against a different database. A single cluster can support many file shares that are hosted on the nodes in a cluster concurrently.

- Active/Passive means that only one instance of a service can be running anywhere in the cluster. For example, there must only be a single instance of the DHCP service running in the cluster at any point in time.

From the perspective of the cluster:

- Active/Active means that all nodes in the cluster are running applications. These may be multiple instances of the same application or different applications (for example, in a 2-node cluster, WINS may be running on one node and DHCP may be running on the other node).

- Active/Passive means that one of the cluster nodes is spare and not being used to host applications.

Server clusters support all of these different combinations; the terms are really about how specific applications or sets of applications are deployed.

With the advent of more than two servers in a cluster, starting with Windows 2000 Datacenter, the term active/active is confusing because there may be four servers. When there are multiple servers, the set of options available for deployment becomes more flexible, allowing different configurations such as N+I.

**Q.** How do I benefit from more than two nodes in a cluster?

**A.** Failover is the mechanism that instance applications and the individual partitions of a partitioned application typically employ for high availability (the term Pack has been coined to describe a highly available, single instance application or partition).

In a 2-node cluster, defining failover policies is trivial. If one node fails, the only option is to failover to the remaining node. As the size of a cluster increases, different failover policies are possible and each one has different characteristics.

**Failover Pairs**

In a large cluster, failover policies can be defined such that each application is set to failover between two nodes. The simple example below shows two applications App1 and App2 in a 4-node cluster.




**Figure 1: Failover pairs**

Configuration has pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting heavy-weight[1] applications such as databases. This configuration ensures that in the event of failure, two applications will not be hosted on the same node. |
| Pro | Very easy to plan capacity. Each node is sized based on the application that it will need to host (just like a 2-node cluster hosting one application). |
| Pro | Effect of a node failure on availability and performance of the system is very easy to determine. |
| Pro | Get the flexibility of a larger cluster. In the event that a node is taken out for maintenance, the buddy for a given application can be changed dynamically (may end up with standby policy below). |
| Con | In simple configurations such as the one above, only 50% of the capacity of the cluster is in use. |
| Con | Administrator intervention may be required in the event of multiple failures. |

[1] A heavy-weight application is one that consumes a significant number of system resources such as CPU, memory or IO bandwidth.

Failover pairs are supported by server clusters on all versions of Windows by limiting the *possible owner list* for each resource to a given pair of nodes.

**Hot-Standby Server**

To reduce the overhead of failover pairs, the spare node for each pair may be consolidated into a single node, providing a hot standby server that is capable of picking up the work in the event of a failure.




**Figure 2: Standby Server**

Configuration has pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting heavy-weight applications such as databases. This configuration ensures that in the event of a single failure, two applications will not be hosted on the same node. |
| Pro | Very easy to plan capacity. Each node is sized based on the application that it will need to host, the spare is sized to be the maximum of the other nodes. |
| Pro | Effect of a node failure on availability and performance of the system is very easy to determine. |
| Con | Configuration is targeted towards a single point of failure. |
| Con | Does not really handle multiple failures well. This may be an issue during scheduled maintenance where the spare may be in use. |

Server clusters support standby servers today using a combination of the possible owners list and the preferred owners list. The preferred node should be set to the node that the application will run on by default and the possible owners for a given resource should be set to the preferred node and the spare node.

**N+I**

Standby server works well for 4-node clusters in some configurations, however, its ability to handle multiple failures is limited. N+I configurations are an extension of the standby server concept where there are N nodes hosting applications and I nodes spare.




**Figure 3: N+I Spare node configuration**

Configuration has pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting heavy-weight applications such as databases or Exchange. This configuration ensures that in the event of a failure, an application instance will failover to a spare node, not one that is already in use. |
| Pro | Very easy to plan capacity. Each node is sized based on the application that it will need to host. |
| Pro | Effect of a node failure on availability and performance of the system is very easy to determine. |
| Pro | Configuration works well for multiple failures. |
| Con | Does not really handle multiple applications running in the same cluster well. This policy is best suited to applications running on a dedicated cluster. |

Server cluster supports N+I scenarios in the Windows Server 2003 release using a cluster group public property **AntiAffinityClassNames**. This property can contain an arbitrary string of characters. In the event of a failover, if a group being failed over has a non-empty string in the **AntiAffinityClassNames** property, the failover manager will check all other nodes. If there are any nodes in the possible owners list for the resource that are NOT hosting a group with the same value in **AntiAffinityClassNames**, then those nodes are considered a good target for failover. If all nodes in the cluster are hosting groups that contain the same value in the **AntiAffinityClassNames** property, then the preferred node list is used to select a failover target.

**Failover Ring**

Failover rings allow each node in the cluster to run an application instance. In the event of a failure, the application on the failed node is moved to the next node in sequence.




**Figure 4: Failover Ring**

Configuration has pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting several small application instances where the capacity of any node is large enough to support several at the same time. |
| Pro | Effect on performance of a node failure is easy to predict. |
| Pro | Easy to plan capacity for a single failure. |

| | |
|---|---|
| Con | Configuration does not work well for all cases of multiple failures. If one Node 1 fails, Node 2 will host two application instances and Nodes 3 and 4 will host one application instance. If Node 2 then fails, Node 3 will be hosting three application instances and Node 4 will be hosting one instance |
| Con | Not well suited to heavy-weight applications since multiple instances may end up being hosted on the same node even if there are lightly-loaded nodes. |

Failover rings are supported by server clusters on the Windows Server 2003 release. This is done by defining the order of failover for a given group using the preferred owner list. A node order should be chosen and then the preferred node list should be set up with each group starting at a different node.

**Random**

In large clusters or even 4-node clusters that are running several applications, defining specific failover targets or policies for each application instance can be extremely cumbersome and error prone. The best policy in some cases is to allow the target to be chosen at random, with a statistical probability that this will spread the load around the cluster in the event of a failure.

Configuration has pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting several small application instances where the capacity of any node is large enough to support several at the same time. |
| Pro | Does not require an administrator to decide where any given application should failover to. |
| Pro | Provided that there are sufficient applications or the applications are partitioned finely enough, this provides a good mechanism to statistically load balance the applications across the cluster in the event of a failure. |
| Pro | Configuration works well for multiple failures. |
| Pro | Very well tuned to handling multiple applications or many instances of the same application running in the same cluster well. |
| Con | Can be difficult to plan capacity. There is no real guarantee that the load will be balanced across the cluster. |
| Con | Effect on performance of a node failure is not easy to predict. |
| Con | Not well suited to heavy-weight applications since multiple instances may end up being hosted on the same node even if there are lightly-loaded nodes. |

The Windows Server 2003 release of server clusters randomizes the failover target in the event of node failure. Each resource group that has an empty preferred owners list will be failed over to a random node in the cluster in the event that the node currently hosting it fails.

**Customized control**

There are some cases where specific nodes may be preferred for a given application instance.

Configuration has pros and cons:

| | |
|---|---|
| Pro | Administrator has full control over what happens when a failure occurs. |
| Pro | Capacity planning is easy, since failure scenarios are predictable. |
| Con | With many applications running in a cluster, defining a good policy for failures can be extremely complex. |
| Con | Very hard to plan for multiple cascaded failures. |

Server clusters provide full control over the order of failover using the preferred node list feature. The full semantics of the preferred node list can be defined as:

| **Preferred Node List** | **Move group to best possible initiated via administrator** | **Failover due to node or group failure** |
|---|---|---|
| Contains all nodes in cluster | Group is moved to highest node in preferred node list that is up and running in the cluster. | **Group is moved to the next node on the preferred node list.** |
| Contains a subset of the nodes in the cluster | Group is moved to highest node in preferred node list that is up and running in the cluster.<br>If no nodes in the preferred node list are up and running, the group is moved to a random node. | **Group is moved to the next node on the preferred node list.**<br>**If the node that was hosting the group is the last on the list or was not in the preferred node list, the group is moved to a random node.** |
| **Empty** | **Group is moved to a random node.** | **Group is moved to a random node.** |

**Q.** How many resources can be hosted in a cluster?

**A.** The theoretical limit for the number of resources in a cluster is 1,674; however, you should be aware that the cluster service periodically polls the resources to ensure they are alive. As the number of resources increases, the overhead of this polling also increases.

# Server Requirements

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** How many servers can be clustered?

**A.** The number of servers in a Server cluster is dependent on the Windows product and the Windows release. The following table lists the server sizes:

| Windows Operating System | Maximum number of servers |
|---|---|
| Windows NT Enterprise Edition | 2 |
| Windows 2000 Advanced Server | 2 |
| Windows 2000 Datacenter Server | 4 |
| Windows Server 2003 Enterprise Edition | 8 |
| Windows Server 2003 Datacenter Edition | 8 |

**Q.** Is it necessary that the servers in a Server cluster be identical?

**A.** The cluster hardware compatibility test does not require that all the servers in a qualified cluster be identical. As cluster sizes increase and therefore the investment in hardware increases, it is likely that different types of servers will appear in a single cluster.

**The qualification process and the listing process are being improved to allow heterogeneous solutions to be more easily defined and qualified. This is particularly important to OEMs where server families change relatively often and therefore the additional qualification passes required increases dramatically with the current process. The server itself has never been an issue during qualification, it is typically the HBA or other piece of the storage subsystem, and therefore there is no real reason to mandate exact servers in a given qualified solution.**

**Q.** Can 32-bit servers and 64-bit servers be mixed in the same Server cluster?

**A.** No, a single Server cluster must contain all 32-bit servers or all 64-bit servers.

**Q.** Can I use any of my servers to make a Server cluster?

**A.** All qualified solutions appear on the Microsoft Hardware Compatibility list (HCL) (http://go.microsoft.com/fwlink/?linkid=67738). Only cluster solutions listed on the HCL are supported by Microsoft. You can use any servers that are listed as part of a configuration to build a complete solution.

# Interconnects

Updated: May 5, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What types of networks are required for cluster interconnects?

**A**. The network interface controllers along with any other components used in certified cluster configurations must have the Windows logo and appear on the Microsoft Hardware Compatibility List. The interconnect itself must support TCP/IP and UDP traffic and must appear as a single, non-routed LAN segment or subnet.

**Q.** Do Server clusters support high-bandwidth, low latency interconnects through WinSock direct?

**A.** The Server cluster code itself does not use WinSock direct path for intra-cluster communication. All communication between cluster nodes is over TCP/IP or UDP. Server clusters can use, but may not take advantage of any high-bandwidth, low latency connection between nodes that looks like a standard NDIS network adapter.

**Q.** How many network adapters should a cluster server have?

**A.** The cluster nodes must be connected by two or more independent networks in order to avoid a single point of failure. The use of two local area networks (LANs) is required. Cluster configurations that have only one network are not supported. At least 2 networks must be configured for cluster communications. Typically, one network is a *private network*, configured for cluster communications only and the other is a *public network*, configured for both client access and cluster communications.

**Q.** Can a cluster support multiple adapters configured as public?

**A.** Yes, however, there are a couple of caveats:

- Each network adapter must be on a different subnet.
- You can create different IP address resources that can be associated with different adapters; however, there is no way to make a resource or application dependent on the two networks so that if one fails, the application will continue to be available. In Server clusters, the dependencies must all be online for the dependent application to be online.

For future versions of the cluster technology, we are working on a feature to allow more flexible dependencies such as: a resource can depend on A OR B for it to be online.

**Q.** Can a cluster support multiple private networks?

**A.** Yes.

**Q.** Why do I see heartbeat packets on the networks marked as client-only?

**A.** In Windows 2000 and beyond, the Server cluster software heartbeats through the public network as well as the private network regardless of the configuration of the public network. This is to ensure that the cluster service can detect failures of the public network adapters and can failover an application if the node currently hosting it cannot communicate with the outside world.

**Q.** How should the different networks be configured?

**A.** The cluster nodes must be connected by two or more independent networks in order to avoid a single point of failure. The use of two local area networks (LANs) is typical. The configuration of a cluster with nodes connected by only one network is not supported.

You should configure the private network as Internal Cluster Communication Only and the public network as All Communications.

**Q.** What type of information is carried over the cluster network?

**A.** The cluster service uses the cluster network for the following intra-cluster traffic:

- Server cluster query/management and control information
- Heartbeats for failure detection
- Intra-cluster communication to ensure tight consistency of cluster configuration

- Cluster join requests when a node reboots or recovers from a failure

**Q.** What types of protocol are supported for applications in a Server cluster?

**A.** Cluster-aware applications must use protocols based on top of TCP/IP. The cluster software only supports the TCP/IP protocols for failover out-of -the-box.

**Q.** Will failover occur if the public network on one server fails?

**A.** Yes, in Windows 2000 and above, additional heartbeats are sent across the public networks to detect public network and/or network adapter failures and to cause applications to failover if the node they are currently hosted on cannot communicate with the other nodes.

**Q.** Is network adapter teaming supported in a Server cluster?

**A.** Yes, however there are caveats. The use of network adapter teaming on all cluster networks concurrently is not supported. At least one of the cluster networks that are enabled for internal communication between cluster nodes must not be teamed. Typically, the un-teamed network is a private interconnect dedicated to this type of communication. The use of network adapter teaming on other cluster networks is acceptable; however, if communication problems occur on a teamed network, Microsoft Product Support Services may require that teaming be disabled. If this action resolves the problem or issue, then you must seek further assistance from the manufacturer of the teaming solution.

**Q.** Is DHCP supported for Server cluster virtual servers?

**A.** No, virtual servers must have static IP addresses.

**Q.** Is DHCP supported for Server cluster nodes?

**A.** Yes, addresses may be assigned to the physical nodes dynamically by DHCP, but manual configuration with static addresses is recommended.

**Q.** Is NetBIOS required to run Server clusters?

**A.** In Windows NT and Windows 2000, NetBIOS is required for Server clusters to function correctly.

In Windows Server 2003, the cluster service does not require NetBIOS. A basic principle of server security is to disable unneeded services. To determine whether to disable NetBIOS, consider the following:

- Some services and applications other than the cluster service use NetBIOS. Review all the ways a clustered server functions before disabling NetBIOS on it.

- With NetBIOS disabled, you will not be able to use the Browse function in Cluster Administrator when opening a connection to a cluster. Cluster Administrator uses NetBIOS to enumerate all clusters in a domain.

- With NetBIOS disabled, Cluster Administrator does not work if a cluster name is specified.

You can control the NetBIOS setting for the cluster through the Cluster IP Address resource property sheet (parameters).

**Q.** Is IPSec supported in a Server cluster?

**A.** Although it is possible to use Internet Protocol Security (IPSec) for applications that can failover in a server cluster, IPSec was not designed for failover situations and we recommend that you do NOT use IPSec for applications in a server cluster.

**Q.** How do Server clusters update router tables when doing IP failover?

**A.** As part of its automatic recovery procedures, the cluster service will issue IETF standard ARP "flush" commands to routers to flush the machine addresses related to IP addresses that are being moved to a different server.

**Q.** How does the Address Resolution Protocol (ARP) cause systems on a LAN to update their tables that translate IP addresses to physical machine (MAC) addresses?

**A.** The ARP specification states that all systems receiving an ARP request must update their physical address mapping for the source of the request. The source IP address and physical network address are contained in the request. As part of the IP address registration process, the Windows TCP/IP driver broadcasts an ARP request on the appropriate LAN several times. This request asks the owner of the specified IP address to respond with its physical network address. By issuing a request for the IP address being registered, Windows can detect IP address conflicts; if a response is received, the address cannot be safely used.

When it issues this request, though, Windows specifies the IP address being registered as the source of the request. Thus, all systems on the network will update their ARP cache entries for the specified address, and the registering system becomes the new owner of the address.

Note that if an address conflict does occur, the responding system can send out another ARP request for the same address, forcing the other systems on the subnet to update their caches again. Windows does this when it detects a conflict with an address that it has successfully registered.

**Q.** Server clusters use ARP broadcasts to re-set MAC addresses, but ARP broadcasts don't pass routers. So what about clients behind the routers?

**A.** If the clients were behind routers, they would be using the router(s) to access the subnet where the cluster servers were located. Accordingly, the clients would use their router (gateway) to pass the packets to the routers through whatever route (OSPF, RIP, etc) is designated. The end result is that their packet is forwarded to a router on the same subnet as the cluster. This router's ARP cache is consistent with the MAC address(es) that have been modified during a failover. Packets thereby get to the correct Virtual server, without the remote clients ever having seen the original ARP broadcast.

# Storage (Server Clusters: Frequently Asked Questions for Windows 2000 and Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are many storage questions, these questions are categorized as general questions, questions about deploying Server clusters on a storage area network (SAN) and network attached storage (NAS) questions.

**General Storage Questions**

**Q.** What storage interconnects are supported by Server clusters?

**A.** Cluster Server does not limit the type of storage interconnects supported; however, there are some requirements of the storage subsystem that limit the types from a practical perspective. For example, all of the cluster nodes should be able to access the storage device. This typically impacts the interconnect since only interconnects that allow multiple initiators (i.e. nodes) can be used. The set of interconnects that are currently part of qualified configurations on the HCL include: SCSI (of various flavors), Fibre channel arbitrated loop and fibre channel switched fabrics.

Remember that only clusters where the full configuration appears on the Cluster HCL are supported by Microsoft.

**Q.** How do I configure the nodes and the storage on a SCSI cluster?

**A.** You must make sure that all of the devices on the SCSI bus have different SCSI Ids. By default, SCSI adapters tend to have Id 7. You should make sure that the adapters in each node have different Ids. Likewise, the disks should be given unique SCSI Ids on the bus.

For a SCSI bus to work correctly it must be terminated. There are many ways to terminate the bus, both internally (at the host adapter) and externally (using Y cables). To ensure that the cluster can survive different types of failures (specifically being able to power down one of the nodes), the SCSI bus should be terminated using passive components such as a Y cable. Internal termination, which requires the adapter to be powered up, is not recommended.

| **Note** |
| --- |
| Microsoft only allows 2-node clusters to be built using SCSI storage interconnects. |

**Q.** Does Server clustering support fibre channel arbitrated loop (FC-AL)?

**A.** Yes, Microsoft only allows 2-node clusters to be built using FC-AL storage interconnects. Multiple clusters on a single fibre channel loop are NOT supported.

**Q.** Can multiple clusters be connected to the same storage controller?

**A.** Yes, there is a special device qualification test for storage controllers that ensures they respond correctly if multiple clusters are attached to the same controller. For multiple clusters to be connected to the same controller the storage controller MUST appear on the Multi-cluster device Hardware Compatibility List (HCL) AND each of the individual end-to-end cluster solutions must appear on the Cluster Hardware Compatibility List. For example: EMC Symmetrix 5.0 is on the multi-cluster device HCL list. Multiple clusters (say a Dell PowerEdge cluster and a Compaq Proliant cluster) can be connected to the EMC Symmetrix storage controller as long as Dell PowerEdge + EMC Symmetrix AND Compaq Proliant + EMC Symmetrix as both on the cluster HCL.

**Q.** Will failover occur if the storage cable is pulled from the host bus adapter (HBA)?

**A.** If the storage cable is pulled from the host bus adapter (HBA), there may be a pause before the adapter reacts to losing the connection, however, once the HBA has detected the communication failure, the disk resources within the cluster using the specific HBA will fail. This will trigger a failover to occur and the resource will be brought back on line on another node in the cluster.

If the storage cable is reconnected, the Windows operating system may not rescan for new hardware automatically (this depends on the driver used for the adapter). You may need to manually initiate a rescan for new devices. Once the rescan is done, the node can host any of the physical disk resources. If failback policies are set, any resources that failed over when the cable was removed will failback to the node when the cable is replaced.

| **Note** |
| --- |

An HBA) is the storage interface that is deployed in the server. Typically this is a PCI card that connects the server to the storage fabric.

**Q.** Will Server clusters protect my disks from hardware failures?

**A.** No, Cluster server protects against server failure, operating system or application failure and planned downtime due to maintenance. Microsoft strongly suggests that application and user data is protected against disk failure using redundancy techniques such as mirroring, RAID or replication either in hardware or in software.

**Q.** Do Server clusters support RAID or mirrored disks?

**A.** Yes, Microsoft strongly suggests that application and user data is protected against disk failure using redundancy techniques such as mirroring, RAID or replication either in hardware or in software.

**Q.** Are dynamic disks supported in a cluster?

**A.** Windows server products shipped from Microsoft do not provide support for dynamic disks in a server cluster environment. The Volume Manager for Windows 2000 add-on product from Veritas can be used to add the dynamic disk features to a server cluster. When the Veritas Volume Manager product is installed on a cluster, Veritas should be the first point of support for cluster issues.

**Q.** Can a cluster disk be extended without rebooting?

**A.** Yes, cluster disks can be extended without rebooting if the storage controller supports dynamic expansion of the underlying physical disk. Many new storage controllers virtualize the *Logical Units* (LUNs) presented to the operating system and these controllers allow LUNs to be grown online from the storage controller management console. Microsoft provides a tool called DiskPart that allows volumes or partitions to be grown online to take advantage of the newly created space on the disk without disruption to applications or users using the disk. There are separate versions of DiskPart for Windows 2000 and Windows Server 2003. The Windows 2000 version is available as a free download on the web and the Windows Server 2003 version is shipped on the distribution media.

| **Note** |
| --- |
| An LUN equates to a disk device visible in Disk Administrator. |

**Q.** Can additional disks be added to a cluster without rebooting?

**A.** Yes, you can insert a new disk or create a new LUN and make that visible to the cluster nodes. You should only make the disk visible to one node in the cluster and then create a cluster resource to protect that disk. Once the disk is protected, you can make the LUN visible to other nodes in the cluster. In some cases you may need to do a rescan in device manager to find the new disk. In other cases (especially with fibre channel) the disk may be automatically detected.

**Q.** Can disks be removed from the cluster without rebooting?

**A.** Yes

**Q.** What types of disks can be used for cluster disks?

**A.** Microsoft recommends that all partitions on the cluster disks be formatted with the NTFS file system. This is for two reasons. Firstly, NTFS provides access controls that can be used to secure the data on the disk. Secondly, NTFS can recover from volumes that have been forcibly dismounted; other file systems may become corrupt if they are forcibly dismounted.

Server clusters only support Master Boot Record (MBR) format disks. Cluster disks cannot be GPT format.

**Q.** Can tape devices or other non-disk devices be attached to the same storage bus as the Cluster Disks?

**A.** It depends on the storage interconnect. Server clusters use SCSI reserve and reset to arbitrate for cluster disks. In Windows NT and Windows 2000 cluster server performs an untargeted bus reset. In Windows Server 2003 it is possible to target the reset, however, it may fallback to an untargeted reset. If a tape device receives a reset, this will typically trigger the tape to rewind.

Server cluster does not provide any arbitration mechanism for tape devices that are visible from multiple servers, therefore tape devices are not protected against concurrent access from multiple servers.

Microsoft does NOT support attaching a tape device to the SCSI bus containing the cluster disks in the case of a SCSI cluster or

the loop in and fibre channel arbitrated loop.

Tape devices can be attached to a switched fabric as long as they are not visible through the same adapters as the cluster disks. This can be achieved by either putting the tape drive in a different zone to the cluster disks or by LUN masking techniques.

**Q.** Are software fault tolerant disks (software RAID or mirroring) supported in a Server cluster?

**A.** Windows server products shipped from Microsoft do not provide support for software RAID or mirroring, however, there are 3rd party products that provide this functionality in a clustered environment.

**Q.** Is the Virtual Snapshot Service (VSS) supported in a Server cluster?

**A.** Yes, the Virtual Snapshot Service is new with Windows Server 2003 and provides basic snapshot capabilities that are used by backup applications to create consistent, single point in time backups. The cluster service has a VSS provider that allows the cluster service configuration to be snapshoted and stored as part of the system state by these backup applications.

**Q.** Do Timewarp snapshots work in a Server cluster?

**A.** No, Timewarp is a new feature in Windows Server 2003 that allows persistent snapshot to be created and exposed to clients. TImewarp makes use of features that are not cluster-aware and it is not supported in a cluster at this time.

**Q.** Are hardware snapshots or business recovery volumes supported in a Server cluster?

**A.** Yes, you can use facilities in the latest storage controllers to create snapshots of existing volumes. Note, however, when you create a snapshot of a disk you should NOT expose the snapshot back to the same cluster as the original disk. The cluster service uses the disk signature to uniquely identify a disk. With a snapshot, the disk and the snapshot have the same disk signature.

If you create a hardware snapshot or a business recovery volume of a cluster disk you should expose the snapshot to another server or cluster (typically a dedicated backup server).

**Q.** What other considerations are there when creating clustered disks?

**A.** Modern storage controllers provide a virtual view of the storage itself. A physical RAID set can be carved into multiple logical units that are exposed to the operating system as individual disks or LUNs. If you intend to carve up physical disks in this way and expose them as independent LUNs to the hosts, you should think carefully about the IO characteristics and the failure characteristics remember underneath, there is only a finite bandwidth to each spindle.

Microsoft recommends that you do not create a LUN for use as the quorum disk from the same underlying physical disks that you will be using for applications. The availability of the cluster is directly related to the availability of the quorum disk. If I/Os to the quorum disk take too long, the cluster server will assume that the quorum disk has failed and initiate a failover of the quorum device. At that point, all other cluster related activity is suspended until the quorum device is brought back online.

**Q.** How can I replace a disk that has gone bad in a cluster?

**A.** The answer depends on the Windows release:

- Windows NT Enterprise Edition
    - FTEdit tool along with some manipulate of the registry. This is covered in article 243195 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=67761).
- Windows 2000
    - DumpCfg.
    - ClusterRecovery Reskit tool provided on Windows Server 2003 Reskit
- Windows Server 2003
    - Automated system recovery.
    - ConfDisk. This is covered in article 280425 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=67762).

- ClusterRecovery Reskit tool provided on Windows Server 2003 Reskit

# Storage Area Networks (SAN) questions

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What is a Storage Area Network (SAN)?

**A.** A storage area network (SAN) is defined as a set of interconnected *devices* (e.g. disks and tapes) and *servers* that are connected to a common communication and data transfer infrastructure such as a fibre channel. The common communication and data transfer mechanism for a given deployment is commonly known as the *storage fabric*. The purpose of the SAN is to allow multiple servers access to a pool of storage in which any server can potentially access any storage unit. Clearly in this environment, management plays a large role providing security guarantees (who is authorized to access which devices) and sequencing or serialization guarantees (who can access which devices at what point in time).

**Q.** Why use a SAN?

**A.** Storage area networks provide a broad range of advantages over locally connected devices. They allow compute units to be detached from storage units, thereby allowing flexible deployment and re-purposing of both servers and storage dynamically to suit the current business needs without having to be concerned about buying the right devices for a given server or without re-cabling a datacenter to attach storage to a given server.

**Q.** Can a cluster be connected to a SAN?

**A.** Yes, Microsoft fully supports storage area networks both as part of the base Windows platform and as part of a complete Windows Clustering high availability solution.

**Q.** Can a cluster co-exist on a SAN with other servers?

**A.** One or more Server clusters can be deployed in a single SAN environment along with standalone Windows servers and/or with other non-Windows platforms.

**Q.** What additional SAN configuration is required to put a cluster on a shared SAN?

**A.**

- The cluster disks for each cluster on a SAN MUST be deployed in their own zone. All host bus adapters (HBAs) in a single cluster must be the same type and at the same firmware revision level. Many storage and switch vendors require that ALL HBAs on the same zone, and in some cases the same fabric, are the same type and have the same firmware revision number.

- All storage device drivers and HBA device drivers in a cluster must be at the same software version.

- When adding a new server to a SAN, ensure that the HBA is appropriate for the topology. In some configurations, adding an arbitrated loop HBA to a switched fibre fabric can result in widespread failures of the storage fabric. There have been real-world examples of this causing serious downtime.

| **Note** |
| --- |
| An HBA is the storage interface that is deployed in the server. Typically this is a PCI card that connects the server to the storage fabric. |

**Q.** What is LUN masking or selective presentation?

**A.** LUN masking (also implemented as selective presentations) allows users to express at the controller level a specific relationship between a LUN and a host. Only the hosts that are configured to access the LUN should be able to see it.

**Q.** What is hardware zoning verses software zoning?

**A.** Zoning can be implemented in hardware/firmware on controllers or on software on the hosts. Microsoft recommends that controller based (or hardware) zoning be used since this allows for uniform implementation of access policy that cannot be interrupted or compromised by node disruption or failure of the software component.

**Q.** Why is zoning needed to isolate a Server cluster in a SAN?

**A.** The cluster uses mechanisms to protect access to the disks that can have an adverse effect on other clusters that are in the same zone. By using zoning to separate the cluster traffic from other cluster or non-cluster traffic, there is no chance of interference.




The diagram shows two clusters sharing a single storage controller. Each cluster is in its own zone. The LUNs presented by the storage controller must be allocated to individual clusters using fine-grained security provided by the storage controller itself. LUNs must be setup so that every LUN for a specific cluster is visible and accessible from all nodes of the cluster. A LUN should only be visible to one cluster at a time. The cluster software itself takes care of ensuring that although LUNs are visible to all cluster nodes, only one node in the cluster accesses and mounts the disk at any point in time.

The multi-cluster device test used to qualify storage configurations for the multi-cluster HCL list tests the isolation guarantees when multiple clusters are connected to a single storage controller in this way.

**Q.** Can a cluster server boot from a SAN?

**A.** Yes, however, there is a set of configuration restrictions around how Windows boots from a storage area network. For more information, see article 305547 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=67837).

Server clusters require that the startup disk, page file disk and system disk be on a different storage bus to the cluster server disks. To boot from a SAN, you must have a separate HBA for the boot, system and pagefile disks than the cluster disks. You MUST ensure that the cluster disks are isolated from the boot, system and pagefile disks by zoning the cluster disks into their own zone.

**Q.** Can I use multiple paths to SAN storage for high availability?

**A.** Microsoft does not provide a generic driver that allows multiple paths to the storage infrastructure for high availability; however, several vendors have built their own proprietary drivers that allow multiple HBAs and SAN fabrics to be used as a highly available storage infrastructure. For a Server cluster that has multi-path drivers to be considered supported, the multipath driver MUST appear as part of the complete cluster solution on the Cluster Hardware Compatibility List (HCL). NOTE: The driver version is VERY important and it MUST match the qualified version on the HCL.

**Q.** Can the startup disk, pagefile disks and the cluster disks be on the same SAN fabric?

**A.** No, in Windows Server 2003, there is a registry key that allows the startup disk, pagefile disks, and cluster disks to be on the same bus. This feature is enabled by a registry key, which helps ensure that it is not accidentally enabled by customers who do not understand the implications of this configuration. It is intended for OEMs to ship qualified and tested configurations and not for a typical end-user or administrator to set up in an ad hoc manner.

In the original release of Windows Server 2003, the registry key is:

**HKLM\SYSTEM\CurrentControlSet\Services\ClusSvc\Parameters\ManageDisksOnSystemBuses 0x01**

In Windows Server 2003 with Service Pack 1 (SP1), the registry key is:

**HKLM\SYSTEM\CurrentControlSet\Services\ClusDisk\Parameters\ManageDisksOnSystemBuses 0x01**

In Windows Server 2003 SP1, the key path was changed to use “Clusdisk” as a subkey instead of “ClusSvc.” This change was made to avoid issues during setup. However, the change is backward compatible, and systems that use the old key locations do not need to be modified.

**Q.** Can serverless backup be performed against cluster disks?

**A.** No, the cluster disk arbitration mechanism uses SCSI reserve and release operations. Once a server arbitrates for a cluster disk, that disk cannot be accessed by any other server on the storage network.

# Network Attached Storage (NAS) questions

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What is Network Attached Storage (NAS)?

**A.** Network attached storage (NAS) is an alternative way to connect storage to servers that is built using standard network components such as Ethernet or other LAN technologies. The application servers access storage using file system functions such as open file, read file, write file, close file, etc.. These higher-level functions are encapsulated in protocols such as CIFS, NFS or AppleShare and run across standard IP-based connections.

**Q.** Can Server clusters use NAS for the shared storage?

**A.** Yes, providing the applications can store data on file shares and the file shares are accessible to the applications as they failover across the cluster, there is no reason why NAS cannot be used as the storage solution in a cluster.

There is currently no support in Windows to use NAS as the quorum resource.

In Windows Server 2003, we are providing a new quorum resource Majority Node Set that can be used to remove the need for a shared disk for the quorum resource. If you combine NAS storage with Majority Node Set quorum, you can build a failover cluster that does not require shared disks in the traditional sense of SCSI or SAN.

# Highly Available File Servers

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** Can I have an active/active file server?

**A.** The Windows 2000 resource kit contains a tool ClusTool that can be used to migrate file share settings from a single node to a cluster environment.

**Q.** Can FRS be used to replicate clustered file shares?

**A.** No, the file replication service (FRS) provided by Windows cannot be used to replicate a clustered file share. This means that clustered file shares cannot be the source or target for redundant links in a DFS tree. See the online documentation for DFS for more details.

**Q.** What file system types are supported in a cluster?

**A.** All partitions on clustered disks should be formatted with NTFS.

**Q.** Does client-side caching (offline folders) work with Server clusters?

**A.** Yes, in Windows Server 2003, you can select client-side caching (also know as offline folders) for clustered file shares.

**Q.** Is the Encrypting File System (EFS) supported on cluster disks?

**A.** With Windows Server 2003, the encrypting file system (EFS) is supported on clustered file shares. To enable EFS on a clustered file share, you must perform a number of tasks to configure the environment correctly:

- EFS can only be enabled on file shares when the virtual server has Kerberos enabled. By default, Kerberos is not enabled on a virtual server. To enable Kerberos you must check the Enable Kerberos Authentication check box on the network name resource that will be used to connect to the clustered file share. NOTE: Enabling Kerberos on a network name has a number of implications that you should ensure you fully understand before checking the box.

- All cluster node computer accounts, as well as the virtual server computer account, must be trusted for delegation. See online help for how to do this.

- To ensure that the users private keys are available to all nodes in the cluster, you must enable roaming profiles for users who want to store data using EFS. See online help for how to enable roaming profiles.

Once the cluster file shares have been created and the configuration steps above carried out, users data can be stored in encrypted files for added security.

**Q.** How many file shares can be hosted on a cluster?

**A.** The number of file shares in a cluster depends on the number of nodes in the cluster and the failure scenarios that you are trying to protect against. A single server has a limit for the number of file shares it can support so you need to take that into account when planning your cluster.

In a 2-node cluster, if one node fails, the remaining node must pick up all of the file shares. Therefore, to ensure the highest availability, the cluster should host the maximum number of shares that can be hosted by a single node.

| **Note** |
| --- |
| 2-node Server clusters are focused on high availability, not scale-out, therefore you should not expect to hold more shares on a 2-node cluster than a single node. |

In a 4-node cluster, you have other options that may be more appropriate, depending on the failure scenarios that you wish to protect against. For example, if you wish to survive one node failing at any point in time, you can configure the shares so that if one node fails, its work is spread across the remaining three nodes. This means that each node could be loaded to 66% of the maximum number of shares and still be within the maximum limit of a single node in the event of a single failure. In this case, the cluster can host three times the number of shares that a single server can host. If you wish to survive two nodes failing, then a 4-node cluster can hold twice as many shares (since if two nodes fail, the remaining two nodes need to pick up the load from the two failed servers) and so on.

In general, as the number of nodes in a cluster increases, the more options you have and the more you can use server clusters to scale-out a highly available infrastructure.

**Q**. What is the maximum capacity of a cluster disk?

**A.** Server cluster does not impose any restrictions on the size of a volume supported.

**Q.** How many disks can Server cluster support?

**A.** In Windows 2000, each clustered disk had to have a drive letter assigned to it, therefore the maximum number of clustered disks in a single cluster was limited to 23 volumes (26 letters of the alphabet minus A and B [Floppy drives] and C [system/boot drive]).

In Windows Server 2003, there is no longer a requirement for a clustered disk to have a driver letter assigned, therefore the number of disks is limited by the number that can be physically attached and the number supported by the underlying operating system.

| **Note** |
| --- |
| Applications can access disks with no drive letters in one of two ways a) directly using the object name associated with the disk or more likely b) by using mount points to link multiple disks together that can be accessed using a single drive letter. |

**Q.** How can a cluster support more disks than there are drive letters?

**A.** Using file system mount points. For more information about using mount points with clustered disks see the online help for Windows Server 2003.

**Q.** Why can I browse shares owned by different virtual servers?

**A.** File shares are not scoped by the virtual server name that is hosting them. If you use a browsing tool (e.g. the NET VIEW command) you will see all the shares that are currently hosted on the physical node.

# Highly Available Print Servers

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** How do I cluster a printer?

**A.** Printers can be clustered using the Print Spooler cluster resource. The Windows 2000 and the Windows Server 2003 online help both give specific examples and the steps necessary to create a highly available print server.

**Q.** Can I have active/active print servers?

**A.** Yes, it is possible to host multiple print spoolers on a single Server cluster. The spoolers can be failed over independently and can run concurrently on multiple nodes in the cluster.

**Q.** How do I migrate printer settings from a single server to a cluster?

A. Microsoft provides a tool (Print Migrator) as part of the ResKit that can be used to migrate printer settings from one node to another or from one node to a Server cluster.

**Q.** How many printer resources can be hosted on a cluster?

**A.** The number of printers is limited by the number of resources the cluster can support, however, as the number of printers increases, so will the time to failover.

# Removing all Single Points of Failure

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What other services does the server cluster rely on?

**A.** The cluster service itself relies on being able to authenticate and sign communications traffic between the cluster nodes. It uses the domain infrastructure to authenticate using the cluster service account. In an environment with Server clusters installed, you must ensure that the domain infrastructure is highly available; any disruption to the infrastructure can result in the clusters becoming unavailable.

**Q.** What other services do I need to think about?

**A.** In order for applications to remain highly available in a clustered environment, any services that the application requires external to the cluster must also be highly available. Many of these services have mechanisms such as replication or being made cluster-aware themselves to protect against failures. Examples of services that you should think about include WINS, DNS, DHCP, the domain infrastructure, firewalls, etc.

**Q.** What other single points of failure should I protect against?

**A.** Server clusters are a mechanism that protects applications against hardware, operating system and application failures. There are some types of hardware failure that you should think about:

- Disk failures you should use RAID or mirroring to protect against disk failures
- Hardware failures multiple hot swap fans in the server, redundant power supplies etc.
- Network failures redundant networks that do not have any shared components
- Site failures disaster recovery plans

# Active Directory, DNS and Domain Controllers

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** Is Kerberos authentication possible for services hosted on a cluster?

**A.** Yes, in Windows 2000 SP3 and above and Windows Server 2003, the cluster service publishes a computer object in Active Directory. This provides the infrastructure with sufficient state to allow Kerberos authentication against applications and services hosted in a virtual server.

For more information about Kerberos and how it works, see the TechNet web site (http://go.microsoft.com/fwlink/?linkid=67842).

**Q.** Can cluster servers also be domain controllers?

**A.** Yes, however, there are several caveats that you should fully understand before taking this approach. We recommend that Server cluster nodes are not domain controllers and that you co-locate a domain controller on the same subnet as the Server cluster public.

If you must make the cluster nodes into domain controllers, consider the following important points:

- If one cluster node in a 2-node cluster is a domain controller, all nodes must be domain controllers. It is recommended that you configure at least two of the nodes in a 4-node Datacenter cluster as domain controllers.

- There is overhead that is associated with the running of a domain controller. A domain controller that is idle can use anywhere between 130 to 140 megabytes (MB) of RAM, which includes the running of Windows Clustering. There is also replication traffic if these domain controllers have to replicate with other domain controllers within the domain and across domains. Most corporate deployments of clusters include nodes with gigabytes (GB) of memory so this is not generally an issue.

- If the cluster nodes are the only domain controllers, they each have to be DNS servers as well, and they should point to each other for primary DNS resolution, and to themselves for secondary DNS resolution. You have to address the problem of the ability to not register the private interface in DNS, especially if it is connected by way of a crossover cable (2-node only). For information about how to configure the heartbeat interface refer to article 258750 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=46549). However, *before* you can accomplish step 12 in KB article 258750, you must first modify other configuration settings, which are outlined in article 275554 (http://go.microsoft.com/fwlink/?LinkID=67844).

  If the cluster nodes are the only domain controllers, they must each be Global Catalog servers, or you must implement domainlets.

- The first domain controller in the forest takes on all flexible single master operation roles (refer to article 197132 at http://go.microsoft.com/fwlink/?LinkID=67847). You can redistribute these roles to each node. However, if a node fails over, the flexible single master operation roles that the node has taken on are no longer available. You can use Ntdsutil to forcibly take away the roles and assign them to the node that is still running (refer to article 223787 at http://go.microsoft.com/fwlink/?LinkID=67851). Review article 223346 at http://go.microsoft.com/fwlink/?LinkID=19807) for information about placement of flexible, single master operation roles throughout the domain.

- If a domain controller is so busy that the Cluster service is unable to gain access to the Quorum drive as needed, the Cluster service may interpret this as a resource failure and cause the cluster group to fail over to the other node. If the Quorum drive is in another group (although it should not be), and it is configured to affect the group, a failure may move all group resources to the other node, which may not be desirable. For more information regarding Quorum configuration, please refer to the article 280345 listed in the "Reference" section (http://go.microsoft.com/fwlink/?LinkID=67855).

- Clustering other programs such as SQL Server or Exchange Server in a scenario where the nodes are also domain controllers, may not result in optimal performance due to resource constraints. You should thoroughly test this configuration in a lab environment prior to deployment.

- You may want to consider making cluster nodes domain controllers (refer to KB article 171390 at http://go.microsoft.com/fwlink/?LinkID=67857 for more information), but if a domain controller is already local, or there is a reliable high-speed connectivity to a domain controller available, Microsoft does not recommend implementing them on cluster nodes.

| **Note** |
| --- |
| You must promote a cluster node to a domain controller by using the Dcpromo tool *prior* to installing Windows Clustering. |

- You must be extremely careful when demoting a domain controller that is also a cluster node. When a node is demoted from a domain controller, the security settings and the user accounts are radically changed (user accounts are demoted to local accounts for example).

**Q.** Are virtual servers published in active directory?

**A.** Yes, in Windows 2000 SP3 and above and in Windows Server 2003, each virtual server has the option of being published in active directory.

Although the network name server cluster resource publishes a computer object in active directory, that computer object should NOT be used for administration tasks such as applying Group Policy. The ONLY role for the virtual server computer object in Windows 2000 and Windows Server 2003 is to allow Kerberos authentication and delegation and for cluster-aware, active directory-aware services (such as MSMQ) to publish service provider information.

**Q.** Is the cluster configuration stored in active directory?

**A.** No, at this time there is no cluster information other than the computer objects for virtual servers published in Active directory.

**Q.** Do Server clusters make domain controllers highly available?

**A.** No, domain controllers use replication across a set of servers to achieve high availability.

**Q.** How should my DNS server be configured to work with Server clusters?

**A.** The cluster service account needs to be able to publish records. In a secure, DNS backed zone, the DNS administrator can chose to restrict the access rights for users. The cluster service account must be granted permission to create records or alternatively, the records can be pre-created. If the records are pre-created, you should not set the zone to dynamic update.

# Security Considerations (Server Clusters: Frequently Asked Questions for Windows 2000 and Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** How do I update the cluster service account password?

**A.** The cluster service account on ALL nodes in the cluster must match to ensure that the intra-cluster communication can be successfully authenticated. The cluster service itself sends messages between cluster nodes under a variety of conditions and if any of those communications fail, the cluster node will be removed from the cluster (i.e. the cluster service will be stopped). It is not possible to determine when the cluster service will establish communication and therefore there is no clear window that allows the cluster service account to be changed in a reliable way while ensuring that the cluster remains running.

**Windows 2000**

On Windows 2000, the cluster account password can only be reliably changed using the following steps:

1. Stop the cluster service on ALL nodes in the cluster
2. Change the password of the cluster service account at the domain controller
3. Update the service control manager password on ALL cluster nodes
4. Re-start the cluster service on all the cluster nodes

**Windows Server 2003**

The cluster.exe command on Windows Server 2003 has the ability to change the cluster account password dynamically without shutting down the cluster service on any of the nodes. The cluster.exe command changes the domain account password and updates the service control manager account information about all nodes in the cluster.

**Cluster /cluster**:*cluster_name1[,cluster_name2,]***/changepassword***[:new_password[,old_password]]* [**/skipdc**] [**/force**] [/options]

For more information refer to the online help for Windows Server 2003.

**Q.** What other security considerations and best practices do I need to worry about for Server clusters?

**A.** For security best practices, see the online help for Windows Server 2003.

# In-the-box HA Services

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What highly available operating system services are provided on the Windows release?

**A.** Server clusters following highly available services by default in the Windows operating system:

- IP Address and Network Name: highly available network configuration to allow clients to location independent and failover unaware
- DHCP: Highly available DHCP server
- MSDTC: Highly available distributed transaction coordinator
- IIS: Highly available web server and FTP server*
- File Share: Highly available file share service and DFS
- Message Queue: Highly available MSMQ service
- MSMQ triggers: Highly available MSMQ trigger service (new for Windows Server 2003)
- Print Spooler: Highly available printer service
- WINS: Highly available WINS service

*In Windows Server 2003, IIS is made cluster-aware using the generic script resource and the scripts provided. There is no specific IIS resource type.

**Q.** Are MSMQ triggers supported in a Server cluster?

**A.** Yes, in Windows Server 2003, the MSMQ trigger service can be made highly available using Server clusters.

**Q.** Is IIS cluster-aware?

**A.** Yes, in Windows 2000, IIS web sites and FTP services can be made highly available using the IIS Server Instance resource type. In Windows Server 2003, the IIS Server Instance resource type was replaced with a set of generic scripts provided in the Windows Server 2003 release (see the online help for more information about converting IIS web servers and FTP servers from Windows 2000 to Windows Server 2003).

Although IIS web servers can be made highly available by failover using Server clustering, Microsoft recommends that you use a load balancing cluster such as provided by Network Load Balancing (NLB), another cluster mechanism provided by the Windows operating system to make IIS highly available and to scale-out a web service or web farm.

Depending on the access characteristics, you may choose either Server clusters or Network Load Balancing clusters to provide highly available FTP servers. Server clustering is good for FTP sites with high update rates or where you want to have a single copy of the FTP content. Network Load Balancing is good for mainly read-only FTP Servers.

**Q.** How is the IIS metabase kept consistent across the cluster?

**A.** The Windows operating system comes with a tool (IISSync) that allows the IIS metabase to be kept in sync across the nodes in the cluster. For more details see the online help.

# Geographically Dispersed Clusters

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** Can Server clusters span multiple sites?

**A.** Yes, Server clusters support a single cluster spanning multiple sites. This is known as a geographically dispersed cluster. All qualified geographically dispersed cluster solutions appear on the Microsoft Hardware Compatibility list (HCL) (http://go.microsoft.com/fwlink/?LinkID=67738). Only cluster solutions listed on the HCL are supported by Microsoft.

**Q.** How is a geographically dispersed cluster defined?

**A.** A geographically dispersed cluster is a Server cluster that has the following attributes:

1. Has multiple storage arrays, at least one deployed at each site. This ensures that in the event of failure of any one site, the other site(s) will have local copies of the data that they can use to continue to provide the services and applications.

2. Nodes are connected to storage in such a way that in the event of a failure of either a site or the communication links between sites, the nodes on a given site can access the storage on that site. In other words, in a two-site configuration, the nodes in site A are connected to the storage in site A directly and the nodes in site B are connected to the storage in site B directly. The nodes in site A con continue without accessing the storage on site B and vice-versa.

3. The storage fabric or host-based software provides a way to mirror or replicate data between the sites so that each site has a copy of the data. Different levels of consistency are available.

The following diagram shows a simple two-site cluster configuration.




**Q.** Will geographically dispersed clusters give me disaster tolerance, or disaster recovery?

**A.** The goal of multi-site Server cluster configurations is to ensure that loss of one site in the solution does not cause a loss of the complete application for business continuance and disaster recovery purposes. Sites are typically up to a few hundred miles apart so that they have completely different power, different communications infrastructure providers, and are placed so that natural disasters (e.g., earthquakes) are extremely unlikely to take out more than one site.

Geographically dispersed clusters do not provide disaster tolerance, since, in some cases, manual intervention is required to restart the applications.

**Q.** Can a geographically dispersed cluster use asynchronous replication between sites?

**A.** Yes, however, there are a couple of caveats:

- The quorum data must be synchronously replicated between the sites. To ensure that the Server cluster guarantees of consistency are met, the cluster database must be kept consistent across all nodes. If the quorum disk is replicated across the sites, it MUST be replicated synchronously.

  In Windows Server 2003, a new quorum resource *Majority Node Set* can be used in these configurations as an alternative to replicating the quorum disk.

- If data is replicated asynchronously, in the event of a site failure, the secondary site will be consistent, but out of date. You should check that the applications can handle going back in time and that the client experience makes sense for the business. Applications such as SQL Server can be hosted on asynchronous replication; however, there are a bunch of restrictions and warning that you should be aware of (see the SQL Server high availability documentation for rules around multi-site replication of SQL Server data).

- If data is replicated at the block level, the replication must preserve the order of writes to the secondary site. Failure to ensure this will lead to data corruption.

**Q.** Does Microsoft provide a complete end-to-end geographically dispersed cluster solution?

**A.** No, Microsoft does not provide a software mechanism for replicating application data from one site to another in a geographically dispersed cluster. Microsoft works with hardware and software vendors to provide a complete solution. All qualified geographically dispersed cluster solutions appear on the Microsoft Hardware Compatibility list (HCL) (http://go.microsoft.com/fwlink/?LinkID=67738). Only cluster solutions listed on the HCL are supported by Microsoft.

**Q.** What additional requirements are there on a Server cluster to support multiple sites?

**A.** The Microsoft server clustering software itself is unaware of the extended nature of geographically dispersed clusters. There are no special features in Server cluster in Windows 2000 or Windows Server 2003 that are specific to these kinds of configuration. The network and storage architectures used to build geographically dispersed clusters must preserve the semantics that the server cluster technology expects. Fundamentally, the network and storage architecture of geographically dispersed server clusters must meet the following requirements:

1. The private and public network connections between cluster nodes must appear as a single, non-routed LAN (e.g., using technologies such as VLANs to ensure that all nodes in the cluster appear on the same IP subnets).

2. The network connections must be able to provide a *maximum guaranteed* round trip latency between nodes of no more than 500 milliseconds. The cluster uses heartbeat to detect whether a node is alive or not responding. These heartbeats are sent out on a periodic basis. If a node takes too long to respond to heartbeat packets, the cluster service starts a heavy-weight protocol to figure out which nodes are really still alive and which ones are dead; this is known as a cluster re-group. The heartbeat interval is not a configurable parameter for the cluster service (there are many reasons for this, but the bottom line is that changing this parameter can have a significant impact on the stability of the cluster and the failover time). 500ms round-trip is significantly below any threshold to ensure that artificial re-group operations are not triggered.

3. Windows 2000 requires that a cluster have a single shared disk known as the quorum disk. The storage infrastructure can provide mirroring across the sites to make a set of disks appear to the cluster service like a single disk, however, it must preserve the fundamental semantics that are required by the physical disk resource:

   - Cluster service uses SCSI reserve commands and bus reset to arbitrate for and protect the shared disks. The semantics of these commands must be preserved across the sites even in the face of complete communication failures between sites. If a node on site A reserves a disk, nodes on site B should not be able to access the contents of the disk. These semantics are essential to avoid data corruption of cluster data and application data.

   - The quorum disk must be replicated in real-time, synchronous mode across all sites. The different members of a mirrored quorum disk MUST contain the same data.

In Windows Server 2003, you can use either a mirrored/replicated quorum disk or a new resource *Majority Node Set* for a multi-site cluster.

**Q.** What additional requirements are there on applications?

**A.** As with the Server cluster itself, applications are unaware of the extended nature of geographically dispersed clusters. There is no topology or configuration information provided to applications to make them aware of the different sites.

Typically, no changes are required to ensure that an application runs, as expected, on a geographically dispersed cluster. However, you should check with the application vendors. In some cases different failure timeout periods may be required since disk accesses and failover times may be longer due to the extended distance between clusters and the need to provide mirroring or replication of data between sites.

# Quorum

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What is the quorum used for?

**A.** Server clusters require a quorum resource to function. The quorum resource, like any other resource, is a resource which can only be owned by one server at a time, and for which servers can negotiate for ownership. Negotiating for the quorum resource allows Server clusters to avoid "split-brain" situations where the servers are active and think the other servers are down. This can happen when, for example, the cluster interconnect is lost and network response time is problematic. The quorum resource is used to store the definitive copy of the cluster configuration so that regardless of any sequence of failures, the cluster configuration will always remain consistent.

**Q.** What options do I have for my quorum resource?

**A.** In Windows 2000, there are two quorum-capable resources:

- **Physical Disk Resource**

  This allows a disk on the shared cluster storage bus to be used as the quorum resource. The cluster code uses SCSI commands to arbitrate for the quorum disk ensuring only one node owns the quorum disk at any point in time. This is the standard quorum resource that Microsoft recommends for all production Windows 2000 clusters.

  | **Note** |
  | --- |
  | Some storage vendors provide their own quorum capable resources for specific hardware solutions (e.g., IBM Shark Storage) or software solutions (such as Veritas Volume Manager). You should use these if they are required for your environment. |

- **Local Quorum**

  This quorum-capable resource allows a single node cluster to be setup without having a second storage bus. This type of cluster is useful for developing cluster-aware software without having to have a multi-node cluster.
  Local Quorum can be used in a production environment if you want to take advantage of the resource health monitoring and local restart facilities provided by Server cluster on a single node.

In Windows Server 2003 we have introduced another quorum capable resource type

**Majority Node Set**

A majority node set is a single quorum resource from a Server cluster perspective; however, all of the quorum data is actually stored on multiple disks across the cluster. The majority node set resource takes care to ensure that the cluster configuration data stored on the majority node set is kept consistent across the different disks.

See *What is Majority Node Set?* for more details.

**Q.** Can other applications share the quorum disk?

**A.** Microsoft recommends that you do NOT use the quorum disk for other applications in the cluster and that the quorum disk is restricted to use by the cluster service itself. If you use the quorum disk for other applications you should be aware of the following:

- The quorum disk health determines the health of the entire cluster. If the quorum disk fails, the cluster service will become unavailable on all cluster nodes. The cluster service checks the health of the quorum disk and arbitrates for exclusive access to the physical drive using standard I/O operations. These operations are queued to the device along with any other I/Os to that device. If the cluster service I/O operations are delayed by extremely heavy traffic, the cluster service will declare the quorum disk as failed and force a regroup to bring the quorum back online somewhere else in the cluster. To protect against malicious applications flooding the quorum disk with I/Os, the quorum disk should be protected. Access to the quorum disk should be restricted to the local administrator group and the cluster service account.

- If the quorum disk fills up, the cluster service may be unable to log required data. In this case, the cluster service will fail, potentially on all cluster nodes. To protect against malicious applications filling up the quorum disk, access should be restricted to the local administrator group and the cluster service account.

- The cluster service itself will always try to bring the quorum disk back online. In doing so, it may violate the failover and failback policies assigned to applications in the same group.

**Q.** Can a NAS device be used as the shared quorum disk?

**A.** No, out-of-the-box, the cluster service supports physical disks on the shared cluster bus or in Windows Server 2003 *Majority Node Set* quorum resources.

**Q.** What is Majority Node Set?

**A.** A majority node set is a single quorum resource from a Server cluster perspective; however, all of the quorum data is actually stored on multiple disks across the cluster. The majority node set resource takes care to ensure that the cluster configuration data stored on the majority node set is kept consistent across the different disks. This allows cluster topologies as follows:




The disks that make up the majority node set could, in principle be local disks physically attached to the nodes themselves or disks on a shared storage fabric. In the majority node set implementation that is provided as part of Server clusters in Windows Server 2003, every node in the cluster uses a directory on its own local system disk to store the quorum data. If the configuration of the cluster changes, that change is reflected across the different disks. The change is only considered to have been committed (i.e. made persistent), if that change is made to:

(<Number of nodes configured in the cluster>/2) + 1

This ensures that a majority of the nodes have an up-to-date copy of the data. The cluster service itself will only start up, and therefore bring resources on line, if a majority of the nodes configured as part of the cluster are up and running the cluster service. If there are fewer nodes, the cluster is said not to have quorum and therefore the cluster service waits (trying to restart) until more nodes try to join. Only when a majority or quorum of nodes, are available, will the cluster service start up the resources be brought online. This way, since the up-to-date configuration is written to a majority of the nodes, regardless of node failures, the cluster will always guarantee that it starts up with the latest and most up-to-date configuration.

# Cluster-Aware Applications

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** What is a Cluster Aware Application?

**A.** A cluster-aware application is an application that calls the cluster APIs to determine the context under which it is running (such as the virtual server name etc.) and can failover between nodes for high availability.

**Can applications that were not written for a cluster be made highly available?**

Yes, Server clusters provide a plug-in environment that allows resource dlls to provide the necessary control and health monitoring functions to make existing applications highly available.

Server clusters provide a set of generic resource types that can be used to make existing applications failover in a cluster. In Windows 2000 there are two generic resource types:

| | |
|---|---|
| **Generic application** | allows any application to be started, stopped and monitored by the cluster service |
| **Generic Service** | allows an existing Windows Service to be started, stopped and monitored. |

These generic services provide very rudimentary health monitoring (for example, is the process that was started still a valid process on the system). It does not check that the application is servicing requests since this requires specific knowledge of the application. The generic resources can be used to make applications failover relatively quickly; however, to provide a more appropriate health check, Microsoft recommends that you build an application-specific resource dll.

In Windows Server 2003, we have provided an additional resource type (**Generic Script**) that allows the start/stop and health monitoring functions to be implemented as scripts rather than using C or C++. This makes the job of building application-specific resource plug-ins much more manageable and easier.

**Q.** How do I build a cluster-aware application?

**A.** Server clusters provides a rich API set that allows applications to recognize and utilize the cluster environment. These APIs are fully documented in the Platform SDK.

**Q.** Should I use the Generic Service or Generic Application resource to make my application highly available?

**A.** The generic services provide very rudimentary health monitoring (for example, is the process that was started still a valid process on the system). It does not check that the application is servicing requests since this requires specific knowledge of the application. The generic resources can be used to make applications failover relatively quickly; however, to provide a more appropriate health check, Microsoft recommends that you build an application-specific resource dll.

**Q.** Does Microsoft validate or logo software products that work with Server clusters?

**A.** Yes, Server clustering is an optional component of the Windows Advanced Server logo program. Applications can be logoed as working on a Server cluster.

**Q.** What Microsoft Applications are cluster-aware?

**A.** The following services shipped as part of the Windows operating system are cluster-aware:

- DHCP Highly available DHCP server
- MSDTC Highly available distributed transaction coordinator
- IIS Highly available web server and FTP server*
- File Share Highly available file share service and DFS
- Message Queue Highly available MSMQ service
- MSMQ triggers Highly available MSMQ trigger service (new for Windows Server 2003)

- Print Spooler Highly available printer service
- WINS Highly available WINS service

The following additional Microsoft Products are cluster-aware:

- SQL Server 6.5, 7.0, 2000 and upwards
- Exchange Server 5.5 and upwards
- Services for Unix 3.0 and upwards

**Q.** Does Exchange 2000 support active/active clusters?

**A.** Yes, there are some caveats to supporting Exchange 2000 in an active/active configuration.

**Q.** Does SQL Server support active/active clusters?

**A.** Yes, SQL Server allows the following:

- Multiple nodes in a cluster hosting different databases. Each database can be failed over independently.
- Multiple nodes in a cluster hosting partitions of a single database using database views to tie the different instances into a single logical database from a client perspective.

**Q.** Where do I find information about writing cluster-aware applications?

**A.** The cluster concepts and APIs are fully documented in the Platform SDK. In addition, there are several examples in the Platform SDK that can be used to demonstrate Server cluster integration.

**Q.** Is Services for Macintosh (SFM) supported in a Server cluster?

**A.** No, Services for Macintosh is not supported in a Server cluster.

**Q.** Is Services for Unix (SFU) supported in a Server cluster?

**A.** Yes, Services for Unix supports highly available NFS shares in SFU 3.0.

# Miscellaneous Topics

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Q.** Can Server clusters and Network Load Balancing be used on the same set of servers?

**A.** No, Microsoft Server clusters (MSCS) and Network Load Balancing (NLB) are not supported on the same set of nodes. Both Server clusters and Network Load Balancing clusters control and configure network adapters. Since they are not aware of each other, configuring one can interfere with the other.

**Q.** Can I use antivirus software with a Server cluster?

**A.** Yes, you should make sure that the vendor has tested their solution in a Server cluster environment. Antivirus software typically layers into the storage stack as a disk driver. This can have an impact on the clusters ability to failover a disk if the driver does not support the required features.

**Q.** Does Microsoft provide a Distributed Lock Manager (DLM)?

**A.** No, at this time, Microsoft has no plans to release a distributed lock manager, however, that may change if there is sufficient customer demand for a DLM service.

| **Note** |
| --- |
| Do not confuse a distributed lock manager with a *cluster file system*. A cluster file system can be built in a number of ways, using a lock manager is one of them. However, just providing a lock manager does not solve the cluster file system problem. |

**Q.** Will Microsoft provide a shared disk cluster file system?

**A.** Microsoft is continually looking at ways to improve the services it provides on the Windows operating system. A shared disk cluster file system is a way to provide a number of attributes in a cluster:

- A single file system namespace visible to all applications on all nodes in the cluster
- High speed access to disks from any node in the cluster

**Q.** How do Server clusters and Fault tolerant servers relate?

**A.** Server clusters address a number of availability issues:

- Hardware failures
- Operating system failures
- Application failures
- Site failures
- Operating system and application upgrades

Fault tolerant servers provide an extremely reliable server platform that addresses hardware failures by providing redundancy at the hardware level. In some cases, fault tolerant servers can be used to address site failures.

In and of themselves, fault tolerant servers do not address issues relating to operating system and application monitoring and failure recovery, nor do they address upgrading the operating system or applications without taking the service down.

In summary, Server clusters are about providing high availability, fault tolerant servers provide high reliability. By combining highly reliable, fault tolerant servers with Server clusters, you get the best of both worlds; a set of reliable servers that can provide high availability in the face of operating system and application failures and upgrades.

**Q.** How do I find Server cluster KB Articles?

**A.** All cluster KB articles can be found at http://go.microsoft.com/fwlink/?LinkID=538

- All Windows NT 4.0 related articles have keyword "MSCS"
- All Windows 2000 articles have keyword "W2000MSCS"

This should allow easy selection of Server cluster related articles.

# Server Clusters : Storage Area Networks - For Windows 2000 and Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: January 1, 2003

**In This White Paper**

Introduction (Server Clusters : Storage Area Networks - For Windows 2000 and Windows Server 2003)

Storage Area Network Components

Management

Deploying Server Clusters in a SAN Environment

Cluster Service Features in Windows Server 2003

Related Issues

Summary (Server Clusters : Storage Area Networks - For Windows 2000 and Windows Server 2003)

Related Links (Server Clusters : Storage Area Networks - For Windows 2000 and Windows Server 2003)

# Introduction (Server Clusters : Storage Area Networks - For Windows 2000 and Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A storage area network (SAN) is defined as a set of interconnected *devices* (e.g. disks and tapes) and *servers* that are connected to a common communication and data transfer infrastructure such as fibre channel. The common communication and data transfer mechanism for a given deployment is commonly known as the *storage fabric*. The purpose of the SAN is to allow multiple servers access to a pool of storage in which any server can potentially access any storage unit. Clearly in this environment, management plays a large role in providing security guarantees (who is authorized to access which devices) and sequencing or serialization guarantees (who can access which devices at what point in time).

SANs evolved to address the increasingly difficult job of managing storage at a time when the storage usage is growing explosively. With devices locally attached to a given server or in the server enclosure itself, performing day-to-day management tasks becomes extremely complex; backing up the data in the datacenter requires complex procedures as the data is distributed amongst the nodes and is accessible only through the server it is attached to. As a given server outgrows its current storage pool, storage specific to that server has to be acquired and attached, even if there are other servers with plenty of storage space available. Other benefits can be gained such as multiple servers can share data (sequentially or in some cases in parallel), backing up devices can be done by transferring data directly from device to device without first transferring it to a backup server.

So why use yet another set of interconnect technologies? A storage area network is a network like any other (for example a LAN infrastructure). A SAN is used to connect many different devices and hosts to provide access to any device from anywhere. Existing storage technologies such as SCSI are tuned to the specific requirements of connecting mass storage devices to host computers. In particular, they are low latency, high bandwidth connections with extremely high data integrity semantics. Network technology, on the other hand, is tuned more to providing application-to-application connectivity in increasingly complex and large-scale environments. Typical network infrastructures have high connectivity, can route data across many independent network segments, potentially over very large distances (consider the internet) and there are many network management and troubleshooting tools.

Storage area networks try to capitalize on the best of the storage technologies and network technologies to provide a low latency, high bandwidth interconnect which can span large distances, has high connectivity and good management infrastructure from the start.

In summary, a SAN environment provides the following benefits:

Centralization of storage into a single pool. This allows storage resources and server resources to grow independently, and allows storage to be dynamically assigned from the pool as and when it is required. Storage on a given server can be increased or decreased as needed without complex reconfiguring or re-cabling of devices.

Common infrastructure for attaching storage allows a single common management model for configuration and deployment.

Storage devices are inherently shared by multiple systems. Ensuring data integrity guarantees and enforcing security policies for access rights to a given device is a core part of the infrastructure.

Data can be transferred directly from device to device without server intervention. For example, data can be moved from a disk to a tape without first being read into the memory of a backup server. This frees-up compute cycles for business logic rather than management related tasks.

Because multiple servers have direct access to storage devices, SAN technology is particularly interesting as a way to build clusters where shared access to a data set is required. Consider a clustered SQL Server environment. At any point in time a SQL Server instance may be hosted on one machine in the cluster and it must have exclusive access to its associated database on a disk from the node on which it is hosted. In the event of a failure or an explicit management operation, the SQL Server instance may *failover* to another node in the cluster. Once failed over, the SQL Server instance must be able to have exclusive access to the database on disk from its new host node.

By deploying multiple clusters onto a single storage area network, all of the benefits of SAN technology described above can be brought to the cluster environment. The rest of this paper describes how clusters can be attached to storage area networks, what the requirements are and what is supported today in Windows 2000.

# Storage Area Network Components

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As previously discussed, the primary technology used in storage area networks today is Fibre Channel. This section provides a basic overview of the components in a fibre channel storage fabric as well as different topologies and configurations open to Windows deployments.

**Fibre Channel Topologies**
Fundamentally, fibre channel defines three configurations:

- Point-to-point
- Fibre Channel Arbitrated Loop (FC-AL)
- Switched Fibre Channel Fabrics (FC-SW).

Although the term fibre channel implies some form of fibre optic technology, the fibre channel specification allows for both fibre optic interconnects as well as copper coaxial cables.

**Point-to-point**
Point-to-point fibre channel is a simple way to connect two (and only two) devices directly together, as shown in Figure 1 below. It is the fibre channel equivalent of direct attached storage (DAS).




**Figure 1: Point to point connection**

From a cluster and storage infrastructure perspective, point-to-point is not a scalable enterprise configuration and we will not consider it again in this document.

**Arbitrated Loops**
A fibre channel arbitrated loop is exactly what it says; it is a set of hosts and devices that are connected into a single loop, as shown in Figure 2 below. It is a cost-effective way to connect up to 126 devices and hosts into a single network.




**Figure 2: Fibre Channel Arbitrated Loop**

Devices on the loop share the media; each device is connected in series to the next device in the loop and so on around the loop. Any packet traveling from one device to another must pass through all intermediate devices. In the example shown, for host A to communicate with device D, all traffic between the devices must flow through the adapters on host B and device C. The devices in the loop do not need to look at the packet; they will simply pass it through. This is all done at the physical layer by the fibre channel interface card itself; it does not require processing on the host or the device. This is very analogous to the way a token-ring topology operates.

When a host or device wishes to communicate with another host or device, it must first arbitrate for the loop. The initiating device does this by sending an arbitration packet around the loop that contains its own loop address (more on addressing later). The arbitration packet travels around the loop and when the initiating device receives its own arbitration packet back, the initiating device is considered to be the loop owner. The initiating device next sends an open request to the destination device which sets up a logical point-to-point connection between the initiating device and target. The initiating device can then send

as much data as required before closing down the connection. All intermediate devices simply pass the data through. There is no limit on the length of time for any given connection and therefore other devices wishing to communicate must wait until the data transfer is completed and the connection is closed before they can arbitrate.

If multiple devices or hosts wish to communicate at the same time, each one sends out an arbitration packet that travels around the loop. If an arbitrating device receives an arbitration packet from a different device before it receives its own packet back, it knows there has been a collision. In this case, the device with the lowest loop address is declared the winner and is considered the loop owner. There is a fairness algorithm built into the standard that prohibits a device from re-arbitrating until all other devices have been given an opportunity, however, this is an optional part of the standard.

**Note**

Not all devices and host bus adapters support loop configurations since it is an optional part of the fibre channel standard. However, for a loop to operate correctly, all devices on the loop MUST have arbitrated loop support1. Figure 3 below shows a schematic of the wiring for a simple arbitrated loop configuration.



**Figure 3: FC-AL wiring schematic**

With larger configurations, wiring a loop directly can be very cumbersome. *Hubs* allow for simpler, centralized wiring of the loop (see section Hubs, Switches, Routers and Bridges). Communication in an arbitrated loop can occur in both directions on the loop depending on the technology used to build the loop, and in some cases communication can occur both ways simultaneously.

Loops can support up to 126 devices, however, as the number of devices on the arbitrated loop increases, so the length of the path and therefore the latency of individual operations increases.

Many loop devices, such as JBODs, have dip switches to set the device address on the loop (known as *hard addressing*). Most, if not all devices, implement hard addresses so it is possible to assign a loop ID to a device, however, just as in a SCSI configuration, different devices must have unique hard IDs. In cases where a device on the loop already has a conflicting address when a new device is added, the new device either picks a different ID or it does not get an ID at all (non-participating).

**Note**

Most of the current FC-AL devices are configured automatically to avoid any address conflicts. However, if a conflict does happen then it can lead to I/O disruptions or failures.

Unlike many bus technologies, the devices on an arbitrated loop do not have to be given fixed addresses either by software configuration or via hardware switches. When the loop initializes, each device on the loop must obtain an Arbitrated Loop Physical Address which is dynamically assigned. This process is initiated when a host or device sends out a LIP; a master is dynamically selected for the loop and the master controls a well defined process where each device is assigned an address.

A LIP is generated by a device or host when the adapter is powered up or when a loop failure is detected (such as loss of carrier). Unfortunately, this means that when new devices are added to a loop or when devices on the loop are power-cycled, all the devices and hosts on the loop can (and probably will) change their physical addresses. This can lead to unstable configurations if the operating system is not fully aware of the changes.

For these reasons, arbitrated loops provide a solution for small numbers of hosts and devices in relatively static configurations.

**Fibre Channel Switched Fabric**

In a switched fibre channel fabric, devices are connected in a many-to-many topology using *fibre channel switches*, as shown in Figure 4 below. When a host or device communicates with another host or device, the source and target setup a point-to-point connection (just like a virtual circuit) between them and communicate directly with each other. The fabric itself routes data from the source to the target. In a fibre channel switched fabric, the media is not shared. Any device can communicate with any other device (assuming it is not busy) and communication occurs at full bus speed (1Gbit/Sec or 2Gbit/sec today depending on

technology) irrespective of other devices and hosts communicating.




**Figure 4: Switched Fibre Channel Fabric**

When a host or device is powered on, it must first login to the fabric. This enables the device to determine the type of fabric (there is a set of characteristics about what the fabric will support) and it causes a host or device to be given a fabric address. A given host or device continues to use the same fabric address while it is logged into the fabric and the fabric address is guaranteed to be unique for that fabric. When a host or device wishes to communicate with another device, it must establish a connection to that device before transmitting data in a way similar to the arbitrated loop. However, unlike the arbitrated loop, the connection open packets and the data packets are sent directly from the source to the target (the switches take care of routing the packets in the fabric).

Fibre channel fabrics can be extended in many different ways such as by federating switches or cascading switches, and therefore, fibre channel switched fabrics provide a much more scalable infrastructure for large configurations. Because device addresses do not change dynamically once a device has logged in to the fabric, switched fabrics provide a much more stable storage area network environment than is possible using an arbitrated loop configuration.

Fibre channel arbitrated loop configurations can be deployed in larger switched SANs. Many of the newer switches from vendors like Brocade incorporate functionality to allow arbitrated loop or point-to-point devices to be connected to any given port. The ports can typically sense whether the device is a loop device or not and adapt the protocols and port semantics accordingly. This allows platforms or specific host adapters or devices which only support arbitrated loop configurations today, to be attached to switched SAN fabrics.

Note that not all switches are created equal. Brocade switches are easy to deploy; Vixel and Gadzoox switches behave more like hubs with respect to addressing.

**Loops verses Fabrics**

Both fibre channel arbitrated loops and switched fabrics have pros and cons. Before deploying either, you need to understand the restrictions and issues as well as the benefits of each technology. The vendor's documentation provides specific features and restrictions; however, the following helps to position the different technologies.

FC-AL

**Pros**

Low cost

Loops are easily expanded and combined with up to 126 hosts and devices

Easy for vendors to develop

**Cons**

Difficult to deploy

Maximum 126 devices

Devices share media thus lower overall bandwidth

Switched Fabric

**Pros**

Easy to deploy

Supports 16 million hosts and devices

Communicate at full wire-speed, no shared media

Switches provide fault isolation and re-routing

**Cons**

Difficult for vendors to develop

Interoperability issues between components from different vendors

Switches can be expensive

**Host Bus Adapters**
A host bus adapter (HBA) is an interface card that resides inside a server or a host computer. It is the functional equivalent of the network adapter in a traditional Ethernet network. All traffic to the storage fabric or loop is done via the HBA.

HBAs, with the exception of older Compaq cards and early Tachyon based cards, support both FC-AL and Fabric (since 1999). However, configuration is not as simple or as automatic as could be supposed. It is difficult to figure out if an HBA configures itself to the appropriate setting. On a Brocade fabric, it is possible to get everything connected, however, some of it might be operating as loop and still appear to work. It is important to verify from the switch side that the hosts are operating in the appropriate mode.

| **Note** |
| --- |
| Be sure to select the correct HBA for the topology that you are using. Although some switches can auto-detect the type of HBA in use, using the wrong HBA in a topology can lead to data loss and can cause many issues to the storage fabric. |

**Hubs, Switches, Routers and Bridges**
Thus far, we have discussed the storage fabric as a generic infrastructure that allows hosts and devices to communication with each other. As you have seen, there are fundamentally different fibre channel topologies and these different topologies use different components to provide the infrastructure.

**Hubs**
Hubs are the simplest form of fibre channel devices and are used to connect devices and hosts into arbitrated loop configurations. Hubs typically have 4, 8, 12 or 16 ports allowing up to 16 devices and hosts to be attached, however, the bandwidth on a hub is shared by all devices on the hub. In addition, hubs are typically half-duplex (newer full duplex hubs are becoming available). In other words, communication between devices or hosts on a hub can only occur in one direction at a time. Because of these performance constraints, hubs are typically used in small and/or low bandwidth configurations.

Figure 5 below shows two hosts and two storage devices connected to the hub with the dark arrows showing the physical loop provided by the hub.



**Figure 5: FC-AL hub configuration**

A typical hub detects empty ports on the hub and does not configure them into the loop. Some hubs provide higher levels of control over how the ports are configured and when devices are inserted into the loop.

**Switches**
A switch is a more complex storage fabric device that provides the full fibre channel bandwidth to each port independently, as shown in Figure 6 below. Typical switches allow ports to be configured in either an arbitrated loop or a switched mode fabric.

When a switch is used in an arbitrated loop configuration, the ports are typically full bandwidth, bi-directional allowing devices and hosts to communicate at full fibre channel speed in both directions. In this mode, ports are configured into a loop, providing performance, arbitrated loop configuration.

Switches are the basic infrastructure used for large, point-to-point, switched fabrics. In this mode, a switch allows any device to communicate directly with any other device at full fibre channel speed (1Gbit/Sec or 2Gbit/sec today).




**Figure 6: Switched fibre configuration**

Switches typically support 16, 32, 64 or even 128 ports today. This allows for complex fabric configurations. In addition, switches can be connected together in a variety of ways to provide larger configurations that consist of multiple switches. Several manufacturers such as Brocade and McData provide a range of switches for different deployment configurations, from very high performance switches that can be connected together to provide a core fabric to edge switches that connect servers and devices with less intensive requirements.

Figure 7 below shows how switches can be interconnected to provide a scalable storage fabric supporting many hundreds of devices and hosts (these configurations are almost certainly deployed in highly available topologies, section Highly Available Solutions deals with high availability).




**Figure 7: Core and edge switches in a SAN fabric**

The core backbone of the SAN fabric is provided by high performance (and typically high port density) switches. The inter-switch bandwidth in the core is typically 8Gbit/sec and above. Large data center class machines and large storage pools can be connected directly to the backbone for maximum performance. Severs and storage with less performance requirements (such as departmental servers) may be connected via large arrays of edge switches (each of which may have 16 to 64 ports).

**Bridges and Routers**

In an ideal world, all devices and hosts would be SAN-aware and all would interoperate in a single, ubiquitous environment. Unfortunately, many hosts and storage components are already deployed using different interconnect technologies. To allow these types of devices to play in a storage fabric environment, a wide variety of *bridge* or *router* devices allow technologies to interoperate. For example, SCSI-to-fibre bridges or routers allow parallel SCSI (typically SCSI-2 and SCSI-3 devices) to be connected to a fibre network, as shown in Figure 8 below. In the future, bridges will allow iSCSI (iSCSI is a device interconnect using IP as the communications mechanism and layering the SCSI protocol on top of IP) devices to connect into a switch SAN fabric.




**Figure 8: SCSI to FibreChannelBridge**

**Storage Components**

Thus far, we have discussed devices being attached to the storage bus as though individual disks are attached. While in some very small, arbitrated loop configurations, this is possible, it is highly unlikely that this configuration will persist. More likely, storage devices such as disk and tape are attached to the storage fabric using a *storage controller* such as an EMC Symmetrix or a Compaq StorageWorks RAID controller. IBM would refer to these types of components as Fibre RAID controllers.

In its most basic form, a storage controller is a box that houses a set of disks and provides a single (potentially redundant and highly available) connection to a SAN fabric. Typically, disks in this type of controller appear as individual devices that map directly to the individual spindles housed in the controller. This is known as a JBOD (just a bunch of disks) configuration. The controller provides no value-add, it is just a concentrator to easily connect multiple devices to a single (or small number for high availability) fabric switch port.

Modern controllers almost always provide some level of redundancy for data. For example, many controllers offer a wide variety of RAID levels such as RAID 1, RAID 5, RAID 0+1 and many other algorithms to ensure data availability in the event of the failure of an individual disk drive. In this case, the hosts do not see devices that correspond directly to the individual spindles; rather the controller presents a virtual view of highly available storage devices to the hosts called *logical devices*.




**Figure 9: Logical devices**

In the example in Figure 9, although there are five physical disk drives in the storage cabinet, only two logical devices are

visible to the hosts and can be addressed through the storage fabric. The controller does not expose the physical disks themselves.

Many controllers today are capable of connecting directly to a switched fabric; however, the disk drives themselves are typically either SCSI, or more common now, are disks that have a built-in FC-AL interface. As you can see in Figure 10 below, the storage infrastructure that the disks connect to is totally independent from the infrastructure presented to the storage fabric.

A controller typically has a small number of ports for connection to the fibre channel fabric (at least two are required for highly available storage controllers). The logical devices themselves are exposed through the controller ports as logical units (LUNs).




**Figure 10: Internal components of a storage controller**

**Highly Available Solutions**

One of the benefits of storage area networks is that the storage can be managed as a centralized pool of resources that can be allocated and re-allocated as required. This powerful paradigm is changing the way data centers and enterprises are built, however, one of the biggest issues to overcome is that of guaranteed availability of data. With all of the data detached from the servers, the infrastructure must be architected to provide highly available access so that the loss of one or more components in the storage fabric does not lead to the servers being unable to access the application data. All areas must be considered including:

No single point of failure of cables or components such as switches, HBAs or storage controllers. Typical highly available storage controller solutions from storage vendors have redundant components and can tolerate many different kinds of failures.

Transparent and dynamic path detection and failover at the host. This typically involves *multi-path* drivers running on the host to present a single storage view to the application across multiple, independent HBAs.

Built-in hot-swap and hot-plug for all components from HBAs to switches and controllers. Many high-end switches and most if not all enterprise class storage controllers allow interface cards, memory, CPU and disk drives to be hot-swapped.

There are many different storage area network designs that have different performance and availability characteristics. Different switch vendors provide different levels of support and different topologies, however, most of the topologies are derived from standard network topology design (after all a SAN is a network, just the interconnect technology is tuned to a given application). Topologies include:

- Multiple independent fabrics
- Federated fabrics
- Core Backbone

**Multiple Independent Fabrics**
In a multiple fabric configuration, each device or host is connected to multiple fabrics, as shown in Figure 11 below. In the event of the failure of one fabric, hosts and devices can communicate using the remaining fabric.




**Figure 11: Multiple independent fabrics**

**Pros**

Resilient to management or user errors. For example, if security is changed or zones are deleted, the configuration on the alternate fabric is untouched and can be re-applied to the broken fabric.

**Cons**

Managing multiple independent fabrics can be costly and error prone. Each fabric should have the same zoning and security information to ensure a consistent view of the fabric regardless of the communication port chosen.

Hosts and devices must have multiple adapters. In the case of a host, multiple adapters are typically treated as different storage buses. Additional *multi-pathing* software such as Compaq SecurePath or EMC PowerPath are required to ensure that the host gets a single view of the devices across the two HBAs.

**Federated Fabrics**
In a federated fabric, multiple switches are connected together, as shown in Figure 12 below. Individual hosts and devices are connected to at least two switches.




**Figure 12: Federated switches for single fabric view**

**Pros**

Management is simplified, the configuration is a highly available, single fabric, and therefore there is only one set of zoning information and one set of security information to manage.

The fabric itself can route around failures such as link failures and switch failures.

**Cons**

Hosts with multiple adapters must run additional *multi-pathing* software such as Compaq SecurePath or EMC PowerPath to ensure that the host gets a single view of the devices where there are multiple paths from the HBAs to the devices.

Management errors are propagated to the entire fabric.

**Core Backbone**
A core backbone configuration is really a way to scale-out a federated fabric environment. Figure 7 shows a backbone configuration. The core of the fabric is built using highly scalable, high performance switches where the inter-switch

connections provide high performance communication (e.g. 8-10GBit/Sec using today's technology). Redundant edge switches can be cascaded from the core infrastructure to provide high numbers of ports for storage and hosts devices.

**Pros**

Highly scalable and available storage area network configuration.

Management is simplified, the configuration is a highly available, single fabric, and therefore there is only one set of zoning information and one set of security information to manage.

The fabric itself can route around failures such as link failures and switch failures.

**Cons**

Hosts with multiple adapters must run additional *multi-pathing* software such as Compaq SecurePath or EMC PowerPath to ensure that the host gets a single view of the devices where there are multiple paths from the HBAs to the devices.

Management errors are propagated to the entire fabric.

[1] Most devices today, except for some McData switches, support FC-AL.

# Management

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you can see from the previous section, storage area networks are increasingly complex and large configurations are becoming more and more common. While storage area networks certainly provide many benefits over direct attach storage, the big issue is how to manage this complexity.

**Zoning**

A storage fabric can have many devices and hosts attached to it. With all of the data stored in a single, ubiquitous cloud of storage, controlling which hosts have access to what data is extremely important. It is also important that the security mechanism be an end-to-end solution so that badly behaved devices or hosts cannot circumvent security and access unauthorized data.

Zoning is a mechanism, implemented at the switch level, which provides an isolation boundary. A port (either host adapters or storage controller ports) can be configured as part of a zone. Only ports in a given zone can communicate with other ports in that zone. The zoning is configured and access control is implemented by the switches in the fabric, so a host adapter cannot spoof the zones that it is in and gain access to data for which it has not been configured.




**Figure 13: Zoning**

In Figure 13 above, hosts A and B can access data from storage controller S1, however host C cannot as it is not in Zone A. Host C can access data from storage S2.

Many switches today allow overlapping zones. This enables a storage controller to reside in more than one zone, thus enabling the devices in that controller to be shared amongst different servers in different zones, as shown in Figure 14 below. Finer precision access controls are required to protect individual disks against access from unauthorized servers in this environment.

Zoning can be implemented in either hardware or software. Hardware zoning is done by the ASIC in the switch ports themselves. Every packet is checked at line speed to ensure that it is authorized. Software zoning is done by the name server or other fabric access software. When a host tries to open a connection to a device, access controls can be checked at that time.




**Figure 14: Storage controller in multiple zones**

Zoning is an extremely important concept. Not only is it a security feature, but it also limits the traffic flow within a given SAN environment. Traffic (I/O requests and other storage requests) between ports is only routed to those pieces of the fabric that are in the same zone. Typically with modern switches, as new switches are added to an existing fabric, the new switches are automatically updated with the current zoning information.

I/Os (either read/write or such things as device reset or LIP) from hosts or devices in a fabric cannot leak out and affect other zones in the fabric causing noise or cross-talk between zones. As we shall see, this is fundamental to deploying Server clusters on a SAN.

**Fine-grain Security and Access Control**

While zoning provides a high-level security infrastructure in the storage fabric, it does not provide the fine-grain level of access control needed for large storage devices. In a typical environment, a storage controller may have many Gigabytes or Terabytes of storage to be shared amongst a set of servers.

Storage controllers typically provide LUN-level access controls that enable an administrator to restrict access to a given LUN to one or more hosts. By providing this access control at the storage controller, the controller itself can enforce access policies to the data.

LUN masking is a host-based mechanism that hides specific LUNs from applications. Although the host bus adapter and the lower layers of the operating system have access to and could communicate with a set of devices, LUN masking prevents the higher layers from knowing that the device exists and therefore applications cannot use those devices. LUN masking is a policy-driven software security and access control mechanism enforced at the host. For this policy to be successful, the administrator has to trust the drivers and the operating systems to adhere to the policies.

**SAN Management**

SAN management is a huge topic on its own and is outside the scope of this document. Different vendors (both vendors that provide SAN fabric components as well as software vendors that provide storage management tools) provide a wide range of tools for setting up, configuring, monitoring and managing the SAN fabric, as well as the state of devices and hosts on the fabric.

**Virtualized View of Storage**

The previous section touched on virtualization of storage when describing various RAID levels. The logical devices presented by the controller to the storage fabric are some composite of the real physical devices in the storage cabinet. Moving forward, the panacea for storage management is that the devices presented to the storage infrastructure are not tied to any physical storage. In other words, the set of spindles in the cabinet is treated as a pool of storage blocks. Logical devices can be materialized from that storage pool with specific attributes such as must survive a single failure, have xyz performance characteristics etc. The storage controller is then free to store the data associated with the logical devices anywhere (and indeed change the placement at will) as long as the desired characteristics are maintained.

At this point, there are no real physical characteristics associated with a logical disk, any physical notions, such as a disk serial number or identity, are purely software-generated virtualized views. See Figure 15 below.

By taking this route, storage vendors can drive many value-added storage management features down into the storage infrastructure itself without having to have host involvement. We are seeing the first few steps down this path today with the

notion of *snapshots* provided by some storage controllers today.




**Figure 15: Storage virtualization by the controller**

# Deploying Server Clusters in a SAN Environment

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section covers best practices for MSCS in a SAN environment. Clusters are supported by Microsoft in a SAN environment; however, there are some specific requirements and restrictions placed on the configurations.

| **Note** |
| --- |
| In a SAN environment, the storage fabric provides access to data for a wide range of applications. If the storage fabric stability is compromised, the availability of the entire data center could be at risk, no amount of clustering can protect against an unstable or unavailable storage fabric. |

**Qualified Configurations**

As with all existing cluster configurations, only *complete cluster solutions* that appear on the Microsoft Hardware Compatibility List (HCL) will be supported by Microsoft. Clusters cannot be arbitrarily built up from device-level components (even those components such as RAID controllers, multi-cluster devices etc. that are qualified as cluster components) and put together into a supported configuration.

A single cluster can be qualified and placed on the HCL using fibre channel storage interconnects and switch technology and there are many examples of complete configurations on the HCL today. This, however, does not really constitute a storage area network (SAN) configuration.

Microsoft fully supports multiple clusters and/or servers deployed on a single fibre channel switched fabric and sharing the same storage controllers as long as the configuration adheres to the following rules:

- The storage controller must be on the Cluster/Multi-Cluster Device HCL list if it is shared between clusters.
- The complete configuration for any individual cluster must be on the Cluster HCL list.

Take, for example the following HCL lists:

Cluster/Multi-cluster device HCL list:

Storage Controller St1

Storage Controller St2

Cluster HCL list

2-node advanced server cluster AS1

Server 1: Server Box S1, 256Mb, 700Mhz PIII, HBA H1

Server 2: Server Box S2, 256Mb, 700Mhz PIII, HBA H1

Storage: Storage Controller St1

4-node advanced server cluster AS2

Server 1: Server Box S5, 512Mb, 1.2Ghz PIV, HBA H1

Server 2: Server Box S6, 512Mb, 1.2Ghz PIV, HBA H1

Server 3: Server Box S7, 512Mb, 1.2Ghz PIV, HBA H1

Server 4: Server Box S8, 512Mb, 1.2Ghz PIV, HBA H1

Storage: Storage Controller St1

2-node advanced server cluster AS3

Server 1: Server Box S10, 256Mb, 700Mhz PIII, HBA H2

Server 2: Server Box S11, 256Mb, 700Mhz PIII, HBA H2

Storage: Storage Controller St2

In this case, the 2-node AS1 and the 4-node AS2 configurations can both be placed on the same storage area network and can in fact share the same storage controller St1. It is also possible to have AS3 on the same storage area network as long as it uses storage controller St2 and not St1.

With Windows 2000, the storage area network fabric itself is not on the HCL and is not qualified directly by Microsoft. When building these configurations, you must ensure that the switches and other fabric components are compatible with the HBAs and the storage controllers.

**Arbitrated Loops and Switched Fabrics**
Fibre channel arbitrated loops can be configured to support multiple hosts and multiple storage devices, however, arbitrated loop configurations typically have restrictions due to the nature of the technology. For example, in some cases, a complete storage controller must be assigned to a given server or cluster. Individual devices in the controller cannot be assigned to different servers or clusters. While manufacturers and vendors allow multiple clusters to be hosted on a single arbitrated loop, due to the configuration restrictions and the mechanisms that the cluster service uses to protect disks in a cluster, Microsoft recommends that only one cluster is attached to any single arbitrated loop configuration and that arbitrated loop configurations are limited to small, relatively static cluster configurations.

Fabrics are fully supported by server clusters for both a single cluster and for multiple clusters and independent servers on the same storage fabric. Fabrics provide a much more stable environment where multiple server clusters are deployed using the same storage infrastructure. Nodes (and indeed storage devices) can leave or enter the SAN independently without affecting other parts of the fabric. Highly available fabrics can be built up, and in conjunction with multi-path drivers, can provide a highly available and scalable storage infrastructure.

**Hints, Tips and Don't Dos**
This section describes the dos and don'ts of deploying one or more clusters in a SAN.

**MUST Do**
Each cluster on a SAN MUST be deployed in its own zone. The cluster uses mechanisms to protect access to the disks that can have an adverse effect on other clusters that are in the same zone. By using zoning to separate the cluster traffic from other cluster or non-cluster traffic, there is no chance of interference. Figure 16 shows two clusters sharing a single storage controller. Each cluster is in its own zone. The LUNs presented by the storage controller must be allocated to individual clusters using fine-grained security provided by the storage controller itself. LUNs must be setup as visible to all nodes in the cluster and a given LUN should only be visible to a single cluster.

The multi-cluster device test used to qualify storage configurations for the multi-cluster HCL list tests the isolation guarantees when multiple clusters are connected to a single storage controller in this way.




**Figure 16: Clusters assigned to individual zones**

All HBAs in a single cluster must be the same type and at the same firmware revision level. Many storage and switch vendors require that ALL HBAs on the same zone, and in some cases the same fabric, are the same type and have the same firmware revision number.

All storage device drivers and HBA device drivers in a cluster must be at the same software version.

SCSI bus resets are not used on a fibre channel arbitrated loop; they are interpreted by the HBA and driver software and cause a LIP to be sent. As previously described, this resets all devices on the loop.

When adding a new server to a SAN, ensure that the HBA is appropriate for the topology. In some configurations, adding an arbitrated loop HBA to a switched fibre fabric can result in widespread failures of the storage fabric. There have been real-world examples of this causing serious downtime.

The base Windows 2000 platform will mount any device that it can see when the system boots. The cluster software ensures that access to devices that can be accessed by multiple hosts in the same cluster is controlled and only one host actually mounts the disk at any one time. When first creating a cluster, make sure that only one node can access the disks that are to be managed by the cluster. This can be done either by leaving the other (to be) cluster members powered off, or by using access controls or zoning to stop the other hosts from accessing the disks. Once a single node cluster has been created, the disks marked as cluster-managed will be protected and other hosts can be either booted or the disks made visible to other hosts to be added to the cluster.

**This is no different to any cluster configuration that has disks that are accessible from multiple hosts.**

| **Note** |
| --- |
| In Windows Server 2003 by using the new command mountvol/n you can disable dynamic scanning. It is recommended that dynamic scanning be disabled before the servers are connected to the SAN in a San environment. New Cluster Setup or adding and removing of nodes in a server cluster should be done while dynamic scanning is turned off. It is recommended that dynamic scanning remains turned off as long as the servers are connected to the storage infrastructure. |

**MUST NOT Do**

NEVER allow multiple hosts access to the same storage devices unless they are in the SAME cluster. If multiple hosts that are not in the same cluster can access a given disk, this will lead to data corruption.

NEVER put any non-disk device into the same zone as cluster disk storage devices.

**Other Hints**

Highly available systems such as clustered servers should typically be deployed with multiple HBAs with a highly available storage fabric. In these cases be sure to ALWAYS load the multi-path driver software. If the I/O subsystem in the Windows 2000 platform sees two HBAs, it will assume they are different buses and enumerate all the devices assuming that they are different devices on each bus; where in fact, the host is seeing multiple paths to the same disks. Failure to load the multi-path driver will lead to data corruption. A simple manifestation of this is that the disk signature is re-written. If the Windows platform sees what it thinks are two independent disks with the same signature, it will re-write one of them to ensure that all disks have unique signatures. This is covered in KB article 293778 *Multiple-Path Software May cause Disk Signature to change*.

| **Note** |
| --- |
| Windows Server 2003 will detect the fact that the same volume is being exposed twice. If such a situation arise Windows Server 2003 will not mount the volumes exposed by controller 2 that have were already exposed by the controller 1. |

Many controllers today provide snapshots at the controller level that can be exposed to the cluster as a completely separate LUN. The cluster does not react well to multiple devices having the same signature. If the snapshot is exposed back to the host with the original disk online, the base I/O subsystem will re-write the signature as in the previous example, however, if the snapshot is exposed to another node in the cluster, the cluster software will not recognize it as a different disk. DO NOT expose a hardware snapshot of a clustered disk back to a node in the same cluster. While this is not specifically a SAN issue, the controllers that provide this functionality are typically deployed in a SAN environment.

**Adding and Removing Disks from a Cluster**

In Windows 2000 (SP3 onwards) and Windows Server 2003, adding a disk to the cluster is simple. Simply add the storage (in a SAN this probably means adding the physical drives to a storage controller and then creating a logical unit that is available in the correct zone and with the correct security attributes).

Once the disk is visible to the operating system, you can make the disk a cluster-managed disk by adding a physical disk resource in cluster administrator. The new disk will appear as being capable of being clustered.

| **Note** |
| --- |
| Some controllers use a different cluster resource than physical disk, for those environments; create a resource of the appropriate type. |

Only Basic, MBR format disks that contain at least one NTFS partition can be managed by the cluster. Before adding a disk, it must be formatted.

Remember that the same rules apply when adding disks as in creating a cluster. If multiple nodes can see the disk BEFORE any node in the cluster is managing it, this will lead to data corruption. When adding a new disk, first make the disk visible to only one cluster node and then once it is added as a cluster resource, make the disk visible to the other cluster nodes.

To remove a disk from a cluster, first remove the cluster resource corresponding to that disk. Once it has been removed from the cluster, the disk can be removed (either the drive can be physically removed or the LUN can be deleted or re-purposed)

There are several KB articles on replacing a cluster-managed disk. While disks in a cluster should typically be RAID sets or mirror sets, there are sometimes issues that cause catastrophic failures leading to a disk having to be rebuilt from the ground up. There are also other cases where cluster disks are not redundant and failure of those disks also leads to a disk having to be replaced. The steps outlined in those articles should be used if you need to rebuild a LUN due to failures.

243195 - *Replacing a cluster managed disk in Windows NT 4.0*

280425 - *Recovering from an Event ID 1034 on a Server Cluster*

**Expanding disks**
Now you can expand volumes dynamically without requiring a reboot. Microsoft provided Diskpart tool can be used to expand volumes dynamically. Diskpart tool is available for each Windows 2000 and Windows Server 2003. You can download Windows 2000 version of Diskpart from the www.microsoft.com web site.

**SAN Backup**
Storage area networks provide many opportunities to offload work from the application hosts. Many of the devices in the SAN (either hosts or storage controllers) have CPUs and memory and are capable of executing complex code paths. In addition, any device can communicate with any other device, the SAN provides a peer-to-peer communication mechanism. This leads to such things as SAN-based backups. A storage controller can easily initiate the backup of a disk device to a tape device on the SAN without host intervention. In some cases, hybrid backup solutions are implemented where file system related information is provide by the host, but bulk copying of the data blocks is done directly from storage controller to tape device.

The cluster software uses disk reservations to protect devices that could be accessed by multiple computers simultaneously. The host that currently owns a disk protects it so that no other host can write to it. This is necessary to avoid writes that are in the pipeline when failover occurs from corrupting the disk. When failover occurs, the new owner protects the disk. This means that a cluster disk is always reserved and therefore can only be accessed by the owning host. No other host or device (including the controller that is hosting the disk) can access the disk simultaneously. This means that SAN-based backup solutions where data transfers from disk to tape are initiated by a 3rd party (i.e. initiated by a device other than the owning host) cannot be supported in a cluster environment.

**Booting from a SAN**
Microsoft supports booting from a SAN in limited environments. There are a set of configuration restrictions around how Windows boots from a storage area network, see KB article 305547.

Windows 2000 Server clusters require that the startup disk, page file disk and system disk be on a different storage bus to the cluster server disks. To boot from a SAN, you must have a separate HBA for the boot, system and pagefile disks than the cluster disks. You MUST ensure that the cluster disks are isolated from the boot, system and pagefile disks by zoning the cluster disks into their own zone.

| **Note** |
| --- |
| Windows Server 2003 will allow for startup disk and the cluster server disks hosted on the same bus. However, you would need to use Storport minioprt HBA drivers for this functionality to work. This is NOT supported configuration with in any other combination (for example, SCSI port miniport or Full port drivers). |

# Cluster Service Features in Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Server 2003 release has a number of enhancements. The following enhancements are specific to supporting Server clusters in a SAN environment:

**Targeted Reset**
Historically, server clusters use the SCSI reservation mechanism to protect disks against access, guaranteeing that only the host that has the disk online can actually access it. To ensure that devices can be failed over in the event of failures, server clusters implements a challenge/response mechanism that can detect dead or unresponsive server nodes even though they may not have crashed and therefore the storage fabric is unaware that the server is not responding. To do this, the reservations are periodically broken by other nodes in the cluster using SCSI bus reset commands. In a SAN fabric, SCSI reset commands can be very detrimental to the fabric since they are typically not implemented the same way by different vendors and they typically result in a LIP command that takes the fabric sometime to re-settle.

In Windows Server 2003, the server cluster code uses a new mechanism to first try targeted device reset, then LUN reset and if all else fails it will fall-back to a full bus reset. This feature requires the storage mini-port drivers to interpret the new control codes. At this time, several HBA vendors are modifying their mini-port drivers to provide this feature, thus enabling much more stable cluster configurations in a switched fabric environment.

This feature requires no administration to enable it. If the device driver supports the targeted reset functions they will be automatically used.

**Single Storage Bus Configurations**
As described previously, in Windows 2000, only storage devices on a different bus to the system disk will be considered eligible as cluster-managed devices. In a SAN environment, the goal is to centralize all storage into a single fabric accessible through a single port (actually in most cases the host will have multiple HBAs and multi-path drivers to provide a single port view to the cluster software).

In Windows Server 2003 Cluster server has a switch that when enabled, allows any disk on the system, regardless of the bus it is on, to be eligible as a cluster-managed disk. Using this, the system disk, startup disk, pagefile disks and any cluster managed disks can be attached to the same HBA. This feature is enabled by setting the following registry key:

```
HKLM\SYSTEM\CurrentControlSet\Services\ClusSvc\Parameters\ManageDisks
OnSystemBuses   0x01
```

This feature is enabled by a registry key to ensure that it is not accidentally enabled by customers that do not understand the implications of this configuration. It is intended for OEMs to ship qualified and tested configurations and not for a typical end-user or administrator to setup in an ad hoc manner.

A single storage bus configuration MUST have device drivers that support the targeted reset functionality previously defined.

# Related Issues

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shared disk verses Shared-nothing**

You may see various documents that use terms like shared disk clusters and non-shared disk or shared-nothing clusters. These terms are very misleading and can cause confusion since they depend on the context of the discussion.

When talking about the physical connectivity of devices, shared disk clusters means that multiple computers have direct physical access to any given storage unit (for example, multiple hosts are directly connected to a disk drive on a SCSI bus that the computers are both connected to). Non-shared disk or shared-nothing clusters in this context means that any given disk is only physically connected to one computer. See Figure 17.




**Figure 17: Physical view of cluster topologies**

In the context of file systems or data access from applications, shared disk means that applications running on multiple computers in a cluster can access the same disk directly at the same time. To support this application, the file system must coordinate concurrent access to a single disk from multiple hosts (e.g. a cluster file system). Clearly, shared physical access is required for this configuration. When talking about application or data access, non-shared disk means that only applications running on one computer can access data on any given disk directly. In this case, the physical disk may or may not be connected to multiple computers, but if it is, then only the connection from one computer is in use at any one time. See Figure 18 below.




**Figure 18: Application view of cluster topologies**

**SAN Versus NAS**

There are two industry-wide terms that refer to externally attached storage:

- Storage Area Networks (SAN)
- Network Attached Storage (NAS)

Having two, similar sounding terms leads to some confusion and therefore it is worth discussing the differences between the two different technologies before delving into storage area network details.

Storage area networks (SANs), see Figure 19 below, are typically built-up using storage-specific network technologies. Fibre channel is the current technology leader in this space. Servers connect to storage and access data at the block level. In other words, to the server, a disk drive out on the storage area network is accessed using the same read and write disk block primitives as though it were a locally attached disk. Typically, data and requests are transmitted using a storage-specific protocol (usually based on the SCSI family of protocols). These protocols are tuned for low latency, high bandwidth data transfers required by storage infrastructure.




**Figure 19: Storage Area Network**

While fibre channel is by far the leading technology today, other SAN technologies have been proposed, for example SCSI over Infiniband, iSCSI (which is SCSI protocol running over a standard IP network). All these technologies allow a pool of devices to be accessed from a set of servers, decoupling the compute needs from the storage needs.

In contrast, network attached storage (NAS), see Figure 20 below, is built using standard network components such as Ethernet or other LAN technologies. The application servers access storage using file system functions such as open file, read file, write file, close file, etc.. These higher-level functions are encapsulated in protocols such as CIFS, NFS or AppleShare and run across standard IP-based connections.




**Figure 20: Network Attached Storage**

In a NAS solution, the file servers hide the details of how data is stored on disks and present a high level file system view to application servers. In a NAS environment, the file servers provide file system management functions such as the ability to back up a file server.

As SAN technology prices decrease and the need for highly scalable and highly available storage solutions increases, vendors are turning to hybrid solutions that combine the centralized file server simplicity of NAS with the scalability and availability offered by SAN as shown in Figure 21 below.




**Figure 21: Hybrid NAS and SAN solution**

The following table contrasts the SAN and NAS technologies

**SAN Versus NAS Technologies**

| | **Storage Area Network** | **Network Attached Storage** |
|---|---|---|
| Application Server Access methods | Block-level access | File-level access |
| Communication protocol | SCSI over Fibre Channel iSCSI (SCSI over IP) | CIFS, NFS; AppleShare |
| Network physical technology | Typically storage-specific<br>(e.g. Fibre-channel) but may be high-speed Ethernet | General purpose LAN<br>e.g. Gigabit Ethernet |
| Example Storage Vendors | Compaq StorageWorks SAN family; EMC Symmetrix | Network Appliance NetApp Filers; Maxtor NASxxxx; Compaq TaskSmart N-series |

There are many camps that believe that in the future, various different technologies will win in the storage space and there are those that believe that in the end there will be a single network interconnect that will cover SAN and NAS needs, as well as basic inter-computer networking needs. Over time, the Windows platform and Windows Clustering technologies will support different interconnect technologies as they become important to end-customer deployments.

# Summary (Server Clusters : Storage Area Networks - For Windows 2000 and Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Storage area networks provide a broad range of advantages over locally connected devices. They allow computer units to be detached from storage units, thereby providing flexible deployment and re-purposing of servers and storage to suit current business needs. You do not have to be concerned about buying the right devices for a given server, or with re-cabling a datacenter to attach storage to a specific server.

Microsoft fully supports storage area networks both as part of the base Windows platform, and as part of a complete Windows Clustering, high availability solution. One or more server clusters can be deployed in a single SAN environment, along with standalone Windows servers and/or non-Windows-based platforms.

# Related Links (Server Clusters : Storage Area Networks - For Windows 2000 and Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

See the following resources for further information:

- For hardware and software compatibility information, see the Windows Server Catalog at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4287).
- For a library of technical articles about Windows Server 2003, see the Windows Server 2003 TechCenter at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4281).
- For the latest information about Windows Server 2003, see the Windows 2003 Server Web site at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48237).
- For interactive help in solving a problem with your computer or to research your problem see Product Support Services at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=281).
- For additional information about cluster deployment, see "Designing and Deploying Clusters" at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48238).
- For information about troubleshooting, see "Troubleshooting cluster node installations" at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48239).

# Server Clusters: Cluster Configuration Best Practices for Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: January 1, 2003

**In This White Paper**

Cluster and Application Deployments

Server Load

Load Balancing

# Cluster and Application Deployments

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Historically, we have described clusters and applications as being active/active or active/passive, which is a natural result of the fact that the earliest products supported only 2-node clusters. In some cases, however, this has led to confusion because the terms can be interpreted in different ways depending on whether the context is 1) how a given application instance is running, 2) how different instances of an application are running or 3) whether the nodes in the cluster are performing useful work. To fully understand the specific application and cluster deployments, we need to understand the following:

- How many instances of the application are actually up and running?
- How many instances of the same data are in the cluster?
- Do instances of the data move around?
- Do different instances of the same application run on different cluster nodes?
- Can different applications run on different cluster nodes?
- What kind of load does the application put on a server, and how will this load be re-distributed after a failover?

Before describing how an application is deployed, we must first define what an application is. An application, for the purposes of this document, is defined as the running code and data required to provide a *single* service to an end-user or client. Consider a couple of different examples to illustrate the point: A single instance of Microsoft Word running on a workstation is a single application instance. If there are multiple instances of Word running, each instance is considered a different application instance. There are, however, more complex applications. For example, with Microsoft SQL Server, a single database is considered to be a single application instance. Independent databases are considered to be different application instances. A single database, however, can be partitioned into multiple SQL Server instances and tied-together using the SQL Server query engine. In this case, the set of SQL Server instances that are tied together to provide a single database image is considered a single application instance.

Fundamentally, there are five different attributes that together provide a complete view of the deployment and can be used to reason about or characterize the deployment:

- Server load How much of a servers resources are consumed by the applications it supports, and how will this resource utilization be affected by failover?
- Application style Is the application a single monolithic application that runs on a single node or is the complete application split into smaller pieces that can be deployed around the cluster?
- Application deployment How are the pieces of the application spread around the cluster in a given deployment?
- Failover policies How is the application configured to behave after a failover?
- Application implementation How is the application itself implemented?

# Server Load

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When deploying an application, it is always important to consider what demands it will make on a servers resources. With clustering, there is a related issue that also needs to be taken into account: how is the load re-distributed after a failover?

**Server Load The Basics**

Consider one of the simplest cases, an active/active, 2-node file server cluster, with node A and node B each serving a single share. If node A fails, its resources will move to node B, placing an additional load on node B. In fact, if node A and B were each running at only 50% capacity before the failure, node B will be completely saturated (100% of capacity) after the failover is completed, and performance may suffer.

While this situation may not be optimal, it is important to remember that having all of the applications still running, even in a reduced performance scenario, is a 100% improvement over what you would have without the high availability protection that clusters provide. But this does bring up the notion of risk, and what amount of it you are willing to accept in order to protect the performance, and ultimately the availability, of your applications.

We have intentionally chosen the worst case (an active/active, 2-node cluster with each node running a single application that consumes half of the servers resources) for the purpose of clarity. With an additional node, the equation changes: there are more servers to support the workload, but if all three nodes are running at 50% capacity and there are *two* failures, the single remaining server will simply not be able to handle the accumulated load of the applications from both of the failed servers. Of course, the likelihood of two failures is considerably less than that of a single failure, so the risk is mitigated somewhat.

Nevertheless, the load/risk tradeoff must be considered when deploying applications on a cluster. The more nodes in a cluster, the more options you have for distributing the workload. If your requirements dictate that all of your clustered applications must run with no performance degradation, then you may need to consider some form of active/passive configuration. But even in this scenario, you must consider the risks of the various configurations. If you cannot accept even the slightest risk of any reduced performance under any conditions whatsoever, you will need a dedicated passive node for each active node.

If, on the other hand, you are convinced that the risk of multiple failures is small, you have other choices. If you have a 4-node, or 8-node cluster, you may want to consider an N+I configuration. N+I, which is discussed in more detail in section 1.4.3, is a variant of active/passive, where N nodes are active, and I nodes are passive, or reserve nodes. Typically, the value for I is less than the value for N, and an N+I cluster topology can handle I failures before any performance degradation is likely. The risk is that with more than I failures, performance will likely decline, but once again, the likelihood of multiple failures is increasingly remote.

For this reason, N+I clusters are a useful configuration that balances the hardware cost of having 100% passive server capacity against the low level risk of multiple cluster node failures.

**Server Load Some More Realistic Configurations**

The scenarios above were intentionally simplistic, assuming that one application imposed a monolithic load on each server, and thus its resource utilization could not be spread among more than one other server in the event of a failover. That is often not the case in the real world, especially for file and print servers, so we will take a look at some additional scenarios with a 4-node cluster, named ABCD, and having nodes A, B, C, and D.

Typically a single server will support the load of more than one application. If under normal conditions, each server was loaded at 25%, then the ABCD cluster could survive the loss of three members before a likely loss of application availability, which would be nearly a worst-case scenario.

The following series of figures illustrates what would happen with the application load in a 4-node cluster for successive node failures. The shaded, or patterned, areas indicate the capacity demands of the running applications. Further, the example below assumes that the application load on any given server is divisible, and can be redistributed among any of the surviving nodes.



**Figure 1.1: Cluster under normal operating conditions (each node loaded at 25%)**



**Figure 1.2: Cluster after a single node failure. Note redistribution of application load.**



**Figure 1.3: Cluster after two node failures. Each surviving node is now approximately 50% loaded.**



**Figure 1.4: After three node failures, single surviving node is at full capacity.**

If each node were running at 75% capacity, then without sensible failover policies, even a single node failure could result in loss of application availability. Depending on the application(s), however, you can specify that, in the event of a server failure, some percentage of the applications should fail over to node A, node B, node C, and node D. If the applications are spread evenly among the surviving nodes, then this cluster could now survive the loss of a single machine, because one third of the failed servers load (one third of 75% is 25%) is allocated to each of the surviving three machines. The result is three fully loaded servers (each of the nodes running at 75% capacity now have an additional 25%), but all applications are still available.

As a variation on the example that was illustrated previously, the following series of figures will depict what happens to a 4-node cluster, each of which is running at approximately 33% capacity. Further, in this case, the application load is indivisible, and can not be spread among multiple other servers in the event of a failover (perhaps it is a single application or multiple applications that depend on the same resource).



**Figure 2.1: Cluster under normal operating conditions (each node loaded at approximately 33%)**



**Figure 2.2: Cluster after a single node failure. Note redistributed application load.**



**Figure 2.3: Cluster after second node failure. Each surviving node is now running at approximately 66% capacity.**

**Note that in the event of another node failure, this cluster will no longer be capable of supporting all four of these applications.**



**Figure 2.4: After third failure, the single surviving server can only support three of the four applications.**

Another style of failover policy is Failover Pairs, also known as Buddy Pairs. Assuming each of the four servers is loaded at 50% or less, a failover buddy can be associated with each machine. This allows a cluster in this configuration to survive two failures. More details on Failover Pairs can be found in section 1.4.1.

Taking the previous Failover Pair example, we can convert it to an active/passive configuration by loading two servers at 100% capacity, and having two passive backup servers. This active/passive configuration can also survive two failures, but note that under ordinary circumstances, two servers remain unused. Furthermore, performance of these servers at 100% load is not likely to be as good as with the Failover Pair configuration, where each machine is only running at 50% utilization. Between these two examples, note that you have the same number of servers, the same number of applications, and the same survivability (in terms of how many nodes can fail without jeopardizing application availability). However, the Failover Pair configuration clearly comes out ahead of active/passive in terms of both performance and economy.

Still using our ABCD 4-node cluster, consider what happens if we configure it as an N+I cluster (explained in more detail in section 1.4.3) where three nodes are running at 100% capacity, and there is a single standby node. This cluster can only survive a single failure. As before, however, comparing it to the example where each server is running at 75% capacity, you again have the same number of servers, applications, and same survivability, but the performance and economy can suffer when you have passive servers backing up active servers running at 100% load.

**Application Style**

Now it is time to consider some aspects of the applications design, and how that affects the way it is deployed. Applications running in a cluster can be characterized as one of:

- **Single Instance**

  In a single instance application, only one application instance is running in the cluster at any point in time. An example of this is the DHCP service in a server cluster. At any point in time, there is only one instance of the DHCP service running in a cluster. The service is made highly available using the failover support provided with server clusters. Single instance applications typically have state that cannot be partitioned across the nodes. In the case of DHCP, the set of IP addresses that have been leased is relatively small, but potentially highly volatile in a large environment. To avoid the complexities of synchronizing the state around the cluster, the service runs as a single instance application.




**Figure 3: Single instance application**

  In the example above, the cluster has four nodes. The single instance application, by definition, can be running on only one node at a time.

- **Multiple Instance**

  A multiple instance application is one where multiple instances of the same code or different pieces of code that cooperate to provide a single service can be executing around the cluster. Together these instances provide the illusion of

a single service to an end-user or client computer. There are two different types of multiple instance applications that depend on the type of application and data:

- **Cloned Application**. A cloned application is a single logical application that consists of two or more instances of the **same code** running against the **same data** set. Each instance of the application is self-contained, thus enabling a client to make a request to any instance of the application with the guarantee that regardless of the instance, the client will receive the same result. A cloned application is scalable since additional application instances can be deployed to service client requests and the instances can be deployed across different nodes, thus enabling an application to grow beyond the capacity of a single node. By deploying application instances on different nodes, the application is made highly available. In the event that a node hosting application instances fails, there are instances running on other nodes that are still available to service client requests.

  Typically, client requests are load-balanced across the nodes in the cluster to spread the client requests amongst the application instances. Applications that run in this environment do not have long-running in-memory state that spans client requests, since each client request is treated as an independent operation and each request can be load-balanced independently. In addition, the data set must be kept consistent across the entire application. As a result, cloning is a technique typically used for applications that have read-only data or data that is changed infrequently. Web server front-ends, scalable edge-servers such as arrays of firewalls and middle-tier business logic fall into this category.

  While these applications are typically called *stateless applications*[1] they can save client-specific session-oriented state in a persistent store that is available to all instances of the cloned application, however, the client must be given a token or a key that it can present with subsequent requests so that whichever application instance services each request can associate the request with the appropriate client state.

  The Microsoft Windows platform provides Network Load Balancing as the basic infrastructure for building scale-out clusters with the ability to spray client requests across the nodes. Application Center provides a single image view of management for these environments.




**Figure 4: Cloned Application**

  In the example above, the same application App is running on each node. Each instance of the application is accessing the same data set, in this case the dataset A-Z. This example show each instance having access to its own data set (created and kept consistent using some form of staging or replication technique), some applications can share the same data set (where the data is available to the cluster using a file share for example).




**Figure 5: Cloned Application using a file share**

- **Partitioned Applications**. Applications that have long-running in-memory state or have large, frequently updated data sets cannot be easily cloned. These applications are typically called *stateful applications*. The cost of keeping the data consistent across many instances of the application would be prohibitive.

  Fortunately, however, many of these applications have data sets or functionality that can be readily partitioned. For example, a large file server can be partitioned by dividing the files along the directory structure hierarchy or a large customer database can be partitioned along customer number or customer name boundaries (customers from A to L in one database, customers from M to Z in another database, for example). In other cases, the functionality can be partitioned or componentized. Once an application is partitioned, the different partitions can be deployed across a set of nodes in a cluster, thus enabling the complete application to scale beyond the bounds of a single node. In order to present a single application image to the clients, a partitioned application requires an application-dependent decomposition, routing and aggregation mechanism that allows a single client request to be distributed across the set of partitions and the results from the partitions to be combined into a single response back to the client. For example, take a partitioned customer database, a single SQL query to return all of the accounts that have overdue payments requires that the query be sent to every partition of the database. Each partition will contain a subset of customer records that must be combined into a single data set to be retuned to the client. Partitioning an application allows it to scale but does not provide high availability. If a node fails that is hosting a partition of the application, that piece of the data set or that piece of functionality is no longer accessible.

  Partitioning is typically application-specific, since the aggregation mechanism is dependent of the type of application and the type of data returned. SQL Server, Exchange data stores and DFS are all examples of applications and services that can be partitioned for scalability.




**Figure 6: Partitioned Application Data Partitioning**

In the example above, each node is running multiple instances of the same application against different pieces of the complete data set. A single request from a client application can span multiple instances, for example a query to return all the records in the database. This *splitting* of the client across the different application instances and the *aggregation* of the single result passed back to the client can be done either by the applications on the server cooperating amongst themselves (for example the SQLQuery engine) or it can be done by the client (as is the

case with Exchange 2000 data stores).

Applications may also be partitioned along functional lines as well as data sets.




**Figure 7: Partitioned Application - Functional Partitions**

In the above example, each node in the cluster is performing a different function; however, the cluster together provides a single, uniform service. One node is providing a catalog service; one is providing a billing service, etc. The cluster, though, provides a single book buying service to clients.

Computational clusters are built to support explicitly massively parallel applications. In other words, the applications are specifically written to decompose a problem into many (potentially thousands) of sub-operations and execute them in parallel across a set of machines. This type of application is also a partitioned application. Computational clusters typically provide a set of libraries to provide cluster-wide communication and synchronization primitives tailored to the environment (MPI is one such set of libraries). These clusters have been termed *Beowulf* clusters. Microsoft has created the High Performance Computing initiative to provide support for this type of cluster.

**Application Deployments**

Applications can be deployed in different ways across a cluster depending on what style of application is being deployed and the number of applications that are being deployed on the same cluster. They can be characterized as follows:

1. **One, single instance application**

   In this type of deployment, one node of the cluster is executing the application so that it can service user requests, the other nodes are in stand-by, waiting to host the application in the event of failure. This type of deployment is typically suited to 2-node failover clusters and is not typically used for N-node clusters since only 1/N of the total capacity is being used. This is what some people term an active/passive cluster.

2. **Several, single instance applications**

   Several single instance applications can be deployed on the same cluster. Each application is independent from the others. In a failover cluster, each application can failover independently for the others, so in the case of a node hosting multiple applications, if that node fails, the different applications may failover to different nodes in the cluster. This type of deployment is typically used for consolidation scenarios where multiple applications are hosted on a set of nodes and the cluster provides a highly available environment. In this environment, careful capacity planning is required to ensue that in the event of failures, the nodes have sufficient capacity (CPU, memory IO bandwidth etc.) to support the increased load.

3. **A single, multiple instance application**

   In this environment, while the cluster is only supporting one application, various pieces of the application are running on the different nodes. In the case of a cloned application, each node is running the same code against the same data. This is the typical web-front end scenario (all of the nodes in the cluster are identical). In the event of a failure, the capacity is reduced. In a partitioned application, the individual partitions are typically deployed across the nodes in the cluster. If a failure occurs, multiple partitions may be hosted on a single node. Careful planning is required to ensure that, in the event of a failure, the application SLA is achieved. Take a 4-node cluster as an example. If an application is partitioned into four pieces, then in normal running, each node would host a partition. In the event that one node failed, two nodes would

continue to host one partition and the remaining node would end up hosting two partitions, giving the 3rd node twice the load, potentially. If the application were split into 12 partitions, each node would host three partitions in the normal case. In the event of a failure, each partition could be configured to failover to a different node, thereby; the remaining nodes could each host four partitions, spreading the load in the event of a failure. The cost however, is that support for 12 partitions may be more overhead than for four partitions.

4. **Several, multiple instance applications**

   Of course, several multiple instance applications may be deployed on the same cluster (indeed single instance and multiple instances may be deployed on the same cluster). In the case of multiple cloned applications, each node is simply running one instance of each application. In the case of a partitioned application, capacity planning as well as defining failover targets becomes increasingly complex.

**Failover Policies**
Failover is the mechanism that single instance applications and the individual partitions of a partitioned application typically employ for high availability (the term Pack has been coined to describe a highly available, single instance application or partition).

In a 2-node cluster, defining failover policies is trivial. If one node fails, the only option is to failover to the remaining node. As the size of a cluster increases, different failover policies are possible and each one has different characteristics.

**Failover Pairs**
In a large cluster, failover policies can be defined such that each application is set to failover between two nodes. The simple example below shows two applications App1 and App2 in a 4-node cluster.




**Figure 8: Failover pairs**

This configuration has pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting heavy-weight[2] applications, such as databases. This configuration ensures that in the event of failure, two applications will not be hosted on the same node. |
| Pro | Very easy to plan capacity. Each node is sized based on the application that it will need to host (just like a 2-node cluster hosting one application). |
| Pro | Effect of a node failure on availability and performance of the system is very easy to determine. |
| Pro | Get the flexibility of a larger cluster. In the event that a node is taken out for maintenance, the buddy for a given application can be changed dynamically (may end up with standby policy below). |
| Con | In simple configurations, such as the one above, only 50% of the capacity of the cluster is in use.z |
| Con | Administrator intervention may be required in the event of multiple failures. |

Failover pairs are supported by server clusters on all versions of Windows by limiting the *possible owner list* for each resource to a given pair of nodes.

**Hot-Standby Server**

To reduce the overhead of failover pairs, the spare node for each pair may be consolidated into a single node, providing a hot standby server that is capable of picking up the work in the event of a failure.




**Figure 9: Standby Server**

The standby server configuration has pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting heavy-weight applications such as databases. This configuration ensures that in the event of a single failure, two applications will not be hosted on the same node. |
| Pro | Very easy to plan capacity. Each node is sized based on the application that it will need to host, the spare is sized to be the maximum of the other nodes. |
| Pro | Effect of a node failure on availability and performance of the system is very easy to determine. |
| Con | Configuration is targeted towards a single point of failure. |
| Con | Does not really handle multiple failures well. This may be an issue during scheduled maintenance where the spare may be in use. |

Server clusters support standby servers today using a combination of the possible owners list and the preferred owners list. The preferred node should be set to the node that the application will run on by default and the possible owners for a given resource should be set to the preferred node and the spare node.

**N+I**

Standby server works well for 4-node clusters in some configurations, however, its ability to handle multiple failures is limited. N+I configurations are an extension of the standby server concept where there are N nodes hosting applications and I nodes which are spares.




**Figure 10: N+I Spare node configuration**

N+I configurations have the following pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting heavy-weight applications such as databases or Exchange. This configuration ensures that in the event of a failure, an application instance will failover to a spare node, not one that is already in use. |

| | |
|---|---|
| Pro | Very easy to plan capacity. Each node is sized based on the application that it will need to host. |
| Pro | Effect of a node failure on availability and performance of the system is very easy to determine. |
| Pro | Configuration works well for multiple failures. |
| Con | Does not really handle multiple applications running in the same cluster well. This policy is best suited to applications running on a dedicated cluster. |

Server clusters supports N+I scenarios in the Windows Server 2003 release using a cluster group public property **AntiAffinityClassNames**. This property can contain an arbitrary string of characters. In the event of a failover, if a group being failed over has a non-empty string in the **AntiAffinityClassNames** property, the failover manager will check all other nodes. If there are any nodes in the possible owners list for the resource that are NOT hosting a group with the same value in **AntiAffinityClassNames**, then those nodes are considered a good target for failover. If all nodes in the cluster are hosting groups that contain the same value in the **AntiAffinityClassNames** property, then the preferred node list is used to select a failover target.

**Failover Ring**

Failover rings allow each node in the cluster to run an application instance. In the event of a failure, the application on the failed node is moved to the next node in sequence.




**Figure 11: Failover Ring**

This configuration has pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting several small application instances where the capacity of any node is large enough to support several at the same time. |
| Pro | Effect on performance of a node failure is easy to predict. |
| Pro | Easy to plan capacity for a single failure. |
| Con | Configuration does not work well for all cases of multiple failures. If Node 1 fails, Node 2 will host two application instances and Nodes 3 and 4 will host one application instance. If Node 2 then fails, Node 3 will be hosting three application instances and Node 4 will be hosting one instance. |
| Con | Not well suited to heavy-weight applications since multiple instances may end up being hosted on the same node even if there are lightly-loaded nodes. |

Failover rings are supported by server clusters on the Windows Server 2003 release. This is done by defining the order of failover for a given group using the preferred owner list. A node order should be chosen and then the preferred node list should be set up with each group starting at a different node.

**Random**

In large clusters or even 4-node clusters that are running several applications, defining specific failover targets or policies for each application instance can be extremely cumbersome and error prone. The best policy in some cases is to allow the target to be chosen at random, with a statistical probability that this will spread the load around the cluster in the event of a failure.

Random failover policies have pros and cons:

| | |
|---|---|
| Pro | Good for clusters that are supporting several small application instances where the capacity of any node is large enough to support several at the same time. |
| Pro | Does not require an administrator to decide where any given application should failover to. |
| Pro | Provided that there are sufficient applications or the applications are partitioned finely enough, this provides a good mechanism to statistically load-balance the applications across the cluster in the event of a failure. |
| Pro | Configuration works well for multiple failures. |
| Pro | Very well tuned to handling multiple applications or many instances of the same application running in the same cluster well. |
| Con | Can be difficult to plan capacity. There is no real guarantee that the load will be balanced across the cluster. |
| Con | Effect on performance of a node failure is not easy to predict. |
| Con | Not well suited to heavy-weight applications since multiple instances may end up being hosted on the same node even if there are lightly-loaded nodes. |

The Windows Server 2003 release of server clusters randomizes the failover target in the event of node failure. Each resource group that has an empty preferred owners list will be failed over to a random node in the cluster in the event that the node currently hosting it fails.

**Customized control**

There are some cases where specific nodes may be preferred for a given application instance.

A configuration that ties applications to nodes has pros and cons:

| | |
|---|---|
| Pro | Administrator has full control over what happens when a failure occurs. |
| Pro | Capacity planning is easy, since failure scenarios are predictable. |
| Con | With many applications running in a cluster, defining a good policy for failures can be extremely complex. |
| Con | Very hard to plan for multiple, cascaded failures. |

Server clusters provide full control over the order of failover using the preferred node list feature. The full semantics of the preferred node list can be defined as:

| **Preferred Node List** | **Move group to best possible initiated via administrator** | **Failover due to node or group failure** |
|---|---|---|
| Contains all nodes in cluster | Group is moved to highest node in preferred node list that is up and running in the cluster. | **Group is moved to the next node on the preferred node list.** |
| Contains a subset of the nodes in the cluster | Group is moved to highest node in preferred node list that is up and running in the cluster.<br><br>If no nodes in the preferred node list are up and running, the group is moved to a random node. | **Group is moved to the next node on the preferred node list.**<br><br>**If the node that was hosting the group is the last on the list or was not in the preferred node list, the group is moved to a random node.** |
| **Empty** | **Group is moved to a random node.** | **Group is moved to a random node.** |

[1] This is really a misnomer, the applications have state, however, the state does not span individual client requests.

[2] A heavy-weight application is one that consumes a significant number of system resources such as CPU, memory or IO bandwidth.

# Load Balancing

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Load balancing is a term that means many different things. Within a cluster environment, load balancing is typically used in two different ways:

- Spreading client requests across a set of clones In Windows, Network Load Balancing (NLB) provides the capability to spread client TCP/IP requests across a set of nodes (typically Web servers or ISA arrays or terminal server machines). This functionality can also be provided by external network infrastructure such as CISCO local director.

- Optimizing the compute resources across a failover cluster A cluster that supports failover allows application instances to move around the cluster. Historically this has been used to ensure that on failure, the application can be restarted elsewhere. However, using this same mechanism could be used with performance counters and application load metrics to provide a way to balance the load of applications across the nodes in a cluster. This can be done in one of two ways:

    - **Manual load balancing**: Each group in a cluster can be moved around independently of the other cluster groups. By breaking down applications into finely defined groups where it is possible, manual load balancing can be done using the cluster administration tools. If the load on a given node is high relative to the other nodes in the cluster, one or more groups can be manually moved to other cluster nodes.

    - **Dynamic load balancing**: Today, Server clusters do not provide any load balancing mechanism (beyond a randomized failover policy in Windows Server 2003). Dynamic load balancing is an extremely complex issue, while moving an application may at first appear to make sense, the cost of that move can be hard to determine. For example, moving SQL Server may not only incur the cost of having to re-mount the database, but it also loses the in-memory cache and state which may have a significant impact on performance from the client perspective.

# Using Clustering with Exchange 2003: An Example

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: September 1, 2003

**In This White Paper**

Introduction (Using Clustering with Exchange 2003: An Example)

Topics Discussed in This White Paper

Reasons for Using Server Clusters

OTG Objectives

Server Cluster Overview

How OTG Configures the Exchange Server Cluster

How OTG Backs Up the Server Cluster

How OTG Monitors the Server Cluster

Best Practices for Server Clusters

Summary (Using Clustering with Exchange 2003: An Example)

Related Topics (Using Clustering with Exchange 2003: An Example)

Appendix A (Using Clustering with Exchange 2003: An Example)

# Introduction (Using Clustering with Exchange 2003: An Example)

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft uses its own Exchange 2003 to provide communication within the company. At Microsoft, the Operations and Technology Group (OTG) manages the Exchange servers to keep them highly available and highly reliable.

OTG manages the global messaging environment that supports Microsoft facilities in more than 400 IT locations in 65 countries. Message traffic averages over 6.3 million messages per day, with around 2.5 million messages per day traveling to and from the Internet.

OTG uses server clustering to ensure predictable and trustworthy messaging. A server cluster is a collection of servers that together provide a single, highly available system for hosting applications. To manage the high message traffic, OTG has implemented a clustering solution that makes Exchange servers at Microsoft available (in terms of global overall availability) more than 99.9% of the time (including both planned and unplanned outages).

In the server clusters that support messaging at Microsoft, each node runs Exchange 2003 and Windows Server 2003, Enterprise Edition. Windows Server 2003, Enterprise Edition supports up to eight nodes per cluster. Although OTG uses many server clusters in their operations, only one of the clusters used for internal messaging is described in this document. This unique cluster design has four active nodes, one primary passive node, and two alternate passive nodes. This design uses servers efficiently, and provides a reliable messaging environment that supports approximately 16,000 mailboxes.

| **Note** |
| --- |
| This is the second of two white papers that describe how clustering is used at Microsoft. The first white paper, Using Clustering for a Highly Available Web site: An Example, describes how Microsoft uses Network Load Balancing for its Microsoft.com Web site. You can download this paper at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=19497). |

**Focus: Server Clusters**

The primary focus of this white paper is how OTG has implemented clustering to manage the Exchange servers that provide messaging for the company. Other details about how the Exchange servers are implemented are only described to the extent that they affect the server cluster.

Although OTG uses many server clusters in their operations, only one of the clusters used for internal messaging is described in this document. Throughout this document, this cluster will be referred to simply as the server cluster or the Exchange server cluster. Exchange server administrators can use this clustering example to assist them in designing their own Exchange servers.

| **Note** |
| --- |
| The information about the Exchange server cluster in this white paper is presented as an example. The actual steps and configuration for computers and devices in your own network will be different. This document shows only the configuration that specifically relate to the objectives that are described later in this document. |

# Topics Discussed in This White Paper

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This paper provides detailed information about the hardware, software, and configuration of the Exchange server cluster, and explains why OTG chose the configuration. It also includes best practices for using server clusters, as well as information about several operational procedures that OTG uses to help make messaging highly reliable and available. The following topics are discussed.

- Reasons for Using Server Clusters
- OTG Objectives
- Server Cluster Overview
- How OTG Configures the Exchange Server Cluster
- How OTG Backs Up the Server Cluster
- How OTG Monitors the Server Cluster
- Best Practices for Server Clusters

# Reasons for Using Server Clusters

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses specific clustering features that help the Operations and Technology Group (OTG) achieve their objectives.

A server cluster is a collection of servers known as nodes that together provide a single, highly available system for hosting applications. Multiple nodes in a cluster remain in constant communication. A clusters resources and applications are organized into functional units called resource groups. Each resource group is assigned to a node and can only be owned by a single node at any point in time.

There are three primary benefits to using server clusters. Server clusters are highly available, scalable, and easy to manage.

- **Highly available:** Server clusters provide a highly available system for deploying applications. You can use server clusters to protect against failures of hardware, operating systems, device drivers, or applications. If one of the nodes in a cluster is unavailable as a result of failure or maintenance, another node immediately begins providing service (a process known as failover). Server clusters also allow you to upgrade the operating system and application software with minimal downtime.

- **Scalable:** Server clusters provide scalability for important resources and applications. When the overall load of a cluster exceeds its capabilities, you can incrementally add additional nodes to the cluster. In Windows Server 2003, server clusters can include up to eight nodes.

- **Easy to manage:** In a server cluster, you can quickly inspect the status of all cluster resources and move workloads around onto different nodes. Because you can move processing to alternate nodes, you can perform rolling upgrades on the servers. In a rolling upgrade, a server cluster continues to provide service while software is being upgraded on each node until all nodes have been upgraded.

Microsoft has used these cluster features to help achieve its technical and business objectives. These objectives are described in the following section.

# OTG Objectives

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Operations and Technology Group (OTG) at Microsoft has used clustering to achieve the following technical and business objectives for messaging within the company:

- **High availability.** The server cluster discussed in this document provides a messaging solution with a high level of availability to users. OTG strives to make the Exchange servers available 99.99% of the time.

- **Scalability**. In a server cluster, administrators can add additional servers as necessary in a way that is virtually transparent to the users. Because they have implemented clustering to manage the Exchange servers (and for other reasons), OTG has been able to increase the scale limitation of the mail store from 3000 mailboxes (100 MB per mailbox) to 5000 mailboxes (200 MB per mailbox).

- **Manageability**. By using a server cluster, administrators can manage and maintain the Exchange servers through a process that is virtually transparent to users.

- **Consolidation.** Server clusters can help administrators consolidate their resources. Server clusters are helping OTG consolidate their worldwide messaging resources from 75 locations to seven and from 115 servers to 30.

These objectives are similar to those of most businesses who use Exchange 2003. Therefore, the server cluster information presented in this document should be valuable to businesses with similar goals.

# Server Cluster Overview

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section gives a general overview of server clusters and the Exchange server cluster at Microsoft that supports messaging within the company.

The following terminology is used in this section and throughout the document.

- Active node. The cluster server that currently owns cluster group resources and responds to network requests made to those services.
- Alternative passive node. The cluster server that does not currently own cluster group resources but is available if the active fails over and the primary passive node is unavailable.
- Primary passive node. The cluster server that does not currently own cluster group resources but is available if the active node fails over.
- Virtual server. A collection of services that appear to clients as a physical Windows-based server but are not associated with a specific server. All virtual servers must include a Network Name resource and an IP Address resource.

Server clusters can take two forms: active/passive clusters and active/active clusters. In active/passive clustering, the cluster includes active nodes and passive nodes. The passive nodes are only used if an active node fails. In active/active clusters, all nodes are active. In the event of a failover, the remaining active node takes on the additional processing operations, which causes a reduction in the overall performance of the cluster. Active/passive cluster configurations are generally recommended over active/active configurations because they often increase performance, availability, and scalability.

The Exchange server cluster is based on an active/passive model, with a few alterations that maximize the performance of the hardware. The Operations and Technology Group (OTG) has configured a multinode cluster that consists of four active nodes, one primary passive node, and two alternate passive nodes. The four active nodes provide most of the processing operations for the messaging environment. Like most active/passive configurations, the primary passive node is used only when one of the active nodes fails or to support the process of rolling upgrades. The two alternate passive nodes are only used for messaging if failover occurs and the primary passive node is not available. To take full advantage of the available hardware, the alternate passive nodes are also used to facilitate the backup and upgrade processes. This unique configuration supports a total of 16,000 mailboxes.

The following topics are described in this section:

- Diagram 1. Exchange Server Cluster
- Reasons for Using Alternative Passive Nodes
- Diagram 2. Failover Process of the Server Cluster
- How OTG Configures the Exchange Server Cluster




**Diagram 1 Exchange Server Cluster**

Diagram 1 shows the four active nodes, the primary passive node, and the two alternate passive nodes that make up the Exchange server cluster. All of the nodes are connected to two storage area networks (SANs). The two SANs provide communication between the nodes and the shared storage components, which are two redundant Enterprise Virtual Arrays. For more information about how storage is configured in this cluster, see Storage Components later in this document.

Each node in the cluster has two adapters that support a public network and a private network. This is a best practice that is described later in this document. The private network is reserved exclusively for internal communication between cluster nodes. The public network supports both node-to-node and client-to-cluster communication.

**Reasons for Using Alternative Passive Nodes**

The four active nodes and the primary passive node each have four hyperthreading 1.9-gigahertz (GHz) processors and 4 gigabytes (GB) of RAM. The two alternate passive nodes are less expensive servers with dual 2.4-GHz processors and 2 GB of RAM.

Although it is rare, the alternate passive nodes can be used to support Exchange processing if an active node fails and the primary passive node is unavailable. Instead, the alternate passive nodes serve two main functions. First, the alternate passive nodes are primarily used to facilitate stage 2 of the backup process. Second, these nodes help during scheduled rolling upgrades and/or patch application to the operating system and/or the host.




**Diagram 2 Failover Process of the Server Cluster**

Diagram 2 shows how the primary passive nodes and alternate passive nodes are used during failover. The four active nodes provide the primary processing operations for the messaging environment. The primary passive node is used when one of the active nodes fails. The two alternate passive nodes are used for messaging only if failover occurs and the primary passive node is not available.

# How OTG Configures the Exchange Server Cluster

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section describes the hardware, software, and server cluster configuration that the Operations and Technology Group (OTG) uses to manage messaging within Microsoft. Because the purpose of this document is to focus on the details of the clustering implementations, only limited information is provided regarding general network hardware configuration, Exchange 2003 configurations, and other configuration settings that are not specifically related to server clusters. However, each server cluster setting is discussed in depth.

- For information about hardware configuration, see Hardware in the Server Cluster.
- For information about the software configuration, see Software in the Server Cluster.
- For information about the server cluster configuration, see Server Cluster Settings.

**Hardware in the Server Cluster**

This section describes the hardware configuration for each of the following components of the server cluster:

- Active and Primary Passive Nodes
- Alternate Passive Nodes
- Network Adapters
- Storage Components

**Active and Primary Passive Nodes**

The four active nodes and the primary passive node in the server cluster all use the same hardware configuration, which is based on the HP ProLiant DL580-G2 processor with the following specification. Each active node supports 4000 mailboxes. Because there are 4 active nodes, this server cluster supports approximately 16,000 mailboxes.

| | |
|---|---|
| **Processors** | HP DL580-G2 Quad - 1.90 GHz, Hyperthreading-enabled |
| **Memory** | 4 GB of Error Correction Coding (ECC) RAM |
| **Controller** | Array controller - Integrated 5i |
| **Hard disk drives** | 2 x 36 GB RAID-1:<br>C$ Windows Server 2003, Enterprise Edition<br>D$ Exchange Application |

**Explanation**

OTG chose this hardware to provide optimized performance during peak usage periods, which are typically experienced on Monday mornings. OTG enabled hyperthreading on these nodes, which has increased performance of the node by 20-25%. (Hyperthreading is a technology that Intel developed that enables multithreaded software applications to execute threads in parallel on a single multicore processor instead of processing threads in a linear fashion). In addition, by increasing the performance, each node can manage more mailboxes.

The hardware for the primary passive node was chosen for the same reasons as the active node hardware. In the event of a failover, the passive node needs to perform exactly like an active node in order to maintain dependable messaging.

| **Note** |
|---|

| |
|---|
| Users who are implementing Exchange 2003 or Windows Server 2003 should refer to the Windows Server Catalog (http://go.microsoft.com/fwlink/?LinkId=4303) before buying hardware to ensure compatibility. The Windows Server Catalog is the online source for customers and partners to find hardware and software that is compatible with the Windows Server 2003 family. |

**Alternate Passive Nodes**

The two alternate nodes in the server cluster have a hardware configuration that is based on the HP ProLiant DL380-G2 dual processor with the following specification.

| | |
|---|---|
| **Processors** | DL380-G2 dual - 2.4 GHz |
| **Memory** | 2 GB RAM |
| **Controller** | Array controller - Integrated 5i |
| **Hard disk drives** | 2 x 36 GB RAID1:<br><br>C$ Windows Server 2003, Enterprise Edition<br><br>D$ Exchange Application |

**Explanation**

As discussed earlier in this paper, OTG chose hardware for the alternate nodes to increase resource efficiency. With this hardware, the alternate passive nodes can be used to facilitate the backup process updates to the system.

For more information about why OTG uses alternative passive nodes, see Reasons for Using Alternative Passive Nodes.

| **Note** |
|---|
| Users who are implementing Exchange 2003 or Windows Server 2003 should refer to the Windows Server Catalog (http://go.microsoft.com/fwlink/?LinkId=4303) before buying hardware. The Windows Server Catalog is the online source for customers and partners to find hardware and software that is compatible with the Windows Server 2003 family. |

**Network Adapters**

Each node in the server cluster has two network adapters installed. OTG has used NC7770 Gigabit Server Adapters that are connected to 100-MB switches.

The following settings appear on the **General** tab that is available when you configure or view the properties of the adapters in Cluster Administrator. You can use Cluster Administrator to configure a cluster and its nodes, resource groups, and resources.

**Adapter 1**

This network adapter enables node-to-node communication. It is connected to the private network that is used exclusively in the cluster. Adapter 2 carries all other communication.

| Configuration setting | Explanation |
|---|---|
| **Name:**<br><br>Internal | OTG has changed the default name for this network connection to clearly identify that it is used for internal communication within the network. |
| **Enabled for cluster use** is checked. | This adapter is available for the cluster to use. |

**Adapter 2**

This network adapter enables both node-to-node (internal) and client-to-cluster (public) communication. Because this network adapter carries both kinds of communication, it can serve as a backup for Adapter 1.

| Configuration setting | Explanation |
|---|---|
| **Name:**<br><br>Internal + Public | OTG has changed the default name for this network connection to clearly identify that it is used for all communication within the network. |

| | |
|---|---|
| **Enabled for cluster use** is checked. | This adapter is available for the cluster to use. |

**Note**

It is important to configure the private network to be the highest priority for internal cluster communication in the Cluster Administrator. It is also important to make sure the public network is first in the TCP/IP binding order.

**Explanation**

OTG uses two networks to ensure network availability for internal cluster communications. Both adapters are configured with static IP addresses. The use of dynamic IP addresses is not recommended in server clusters.

**Note**

The nodes of a cluster must be on the same subnet, but you can use virtual LAN (VLAN) switches on the interconnects between two nodes. (An interconnect is a private network that connects nodes in a cluster). If you use a VLAN, the point-to-point, round-trip latency must be less than 1/2 second and the link between two nodes must appear as a single point-to-point connection from the perspective of the Windows operating system running on the nodes. To avoid single points of failure, use independent VLAN hardware for the different paths between the nodes.

**Storage Components**

This section will discuss the storage design for the Exchange server cluster and the reasons for this configuration.

The following terminology is used in this section and throughout the document.

- Disk group. A group of physical disks. Virtual disks are created across the disk groups.
- Storage group. An Exchange 2003 component that supports five databases (mail stores) using either a single or multiple data device with a common log device.
- Virtual server. A collection of services that appear to clients as a physical Windows-based server, but are not associated with a specific server. All virtual servers must include a Network Name resource and an IP Address resource.
- Storage area network (SAN). A high-speed network that provides a direct connection between servers and storage, including shared storage, clusters, and disaster-recovery devices.

All of the nodes in the cluster are connected to two SANs, which provide communication between the nodes and the two shared storage components.

The shared storage components are two redundant Enterprise Virtual Arrays (HP StorageWorks EVA 5000s). Diagram 1 illustrates the relationship of the nodes and the storage components. Diagram 3 describes the storage configuration for each Exchange virtual server. OTG configured each EVA to support four Exchange virtual servers and a total of 16,000 200-MB mailboxes.

The following table summarizes the storage configuration for each EVA. Each EVA has five disk groups that are characterized by different numbers of hard disks, data types, and Vraid types.

| EVA disk group | Number of 72-GB disks | Data type | Vraid type |
|---|---|---|---|
| DG1 | 48 | Data/SMTP | 1 |
| DG2 | 48 | Data/SMTP | 1 |
| DG3 | 14 | Log/Quorum | 1 |
| DG4 | 28 | Backup | 5 |
| DG5 | 28 | Backup | 5 |

**Explanation**

OTG designed this cluster for optimum performance to support production requirements for peak loads. They also strived to ensure maximum availability by eliminating as many single points of failure within the environment. OTG achieved optimum

performance by scaling the servers to support a realistic number of mailboxes and optimizing the SAN design for peak disk transfers (4K random read/write activity). OTG analyzed I/O requirements for each 200 MB mailbox and determined that each mailbox would generate between 1.0 and 1.2 disk transfers for peak loads (this is specific to the user environment at Microsoft).

With a plan to support 4000 mailboxes for each Exchange virtual server, OTG determined that each host would need to sustain 4800 disk transfers per second and also have the ability to support any unexpected increase in demand.

OTG needed to optimize the SAN design in order to anticipate and account for peak disk transfers while maintaining optimum read and write latencies. They also needed to ensure that Exchange would not encounter any bottlenecks.

After evaluating the performance of the HP StorageWorks EVA 5000s, OTG found that the EVA supported 12,000 transfers in an optimized configuration with acceptable latencies, even under an aggressive 4 KB random I/O load. As a result, OTG implemented two EVA 5000s per cluster. These EVAs support four Exchange virtual servers and a total of 16,000 200-MB mailboxes.

OTG uses HP StorageWorks Secure Path, which is a multipath application, to manage and maintain continuous data access to the storage system. Specifically, they use Secure Path to assist with optimizing disk I/O activity by ensuring disk distribution across fiber channel adapters (FCAs) and SAN controllers on each node. (Fiber channel adapters are expansion cards used to connect the nodes to SANs). Disk distribution has proven to be beneficial in maintaining an optimized environment. It has greatly reduced peak read/write disk latency and has improved online backup to disk.

OTG also uses Secure Path to assist with eliminating single points of failure between nodes and the storage. Secure Path helps maintain I/O even if a node fails. In the event of a node failure, Secure Path detects the failure and moves disks from the failed path to an available path. This process is called failover. Failover keeps the resources available while requiring no downtime. Once the failed nodes are replaced, the disks can be failed back using HPs Secure Path Manager to restore optimized I/O.




**Diagram 3 Storage Configuration for Each Exchange Virtual Server**

Diagram 3 depicts the storage configuration for one of the Exchange virtual servers. OTG configured each EVA to support four Exchange virtual servers and a total of 16,000 200-MB mailboxes. Each Exchange virtual server hosts four storage groups with five databases each. Each database is configured with a 200-MB limit per mailbox. This design specifies a maximum of 200 mailboxes per database, so each Exchange virtual server contains 4,000 mailboxes.

For additional information about the storage configuration for each virtual server, see Appendix A. For more information about Enterprise Virtual Arrays, see Microsoft Exchange 2003 Storage Design for the EVA5000 and other related white papers in the Storage Array Systems White Papers section of the Hewlett-Packard Web site (http://go.microsoft.com/fwlink/?LinkId=19716).

**Software in the Server Cluster**

The nodes in the server cluster have the following software installed.

| Node Type | Software |
|---|---|
| Active node and primary passive node | • Windows Server 2003, Enterprise Edition<br>• Exchange 2003<br>• Microsoft Operations Manager (MOM) Agent<br>• HP StorageWorks Secure Path |
| Alternative passive node | • Windows Server, 2003 Enterprise Edition<br>• Exchange 2003<br>• Microsoft Operations Manager (MOM) Agent<br>• HP StorageWorks Secure Path<br>• Veritas Backup Exec |

OTG has made the following operating system modifications on each node:

- /3-GB switch set in the Boot.ini file.
- /USERVA=3030 parameter set in the Boot.ini file.
- In the registry, SystemPages is set to 0.
- Cluster event log replication is set to 0 in Event Viewer. Because OTG uses Microsoft Operations Manager to manage their clusters, this setting is set to 0 to eliminate event replication across all nodes in the cluster.
- Mount points are used to support the log, SMTP, and backup drives in order to reduce the number of drive letters used.

**Server Cluster Settings**

This section describes how OTG has configured the settings for the Exchange server cluster and the reasons why they chose the settings. It also explains why some settings are specific to OTG and are not recommended for typical customer deployments. This section describes the settings that are configured for the server cluster as a whole, the settings that are specific to each resource group in the cluster, and each of the resource configurations.

The server cluster contains four Exchange virtual servers, four backup virtual servers, and one cluster virtual server. A virtual server is a collection of services that appear to clients as a physical Windows-based server, but are not associated with a specific server. All virtual servers must include a Network Name resource and an IP Address resource.

Each of the virtual servers contains the following specified resources. For server clusters, a resource is a physical or logical entity that is capable of being managed by a cluster, brought online and taken offline, and moved between nodes. A resource can be owned only by a single node at any point in time.

| Virtual Server | Resources |
|---|---|

| | |
|---|---|
| Exchange virtual servers | • Nine Physical Disk resources<br>• One Network Name resource<br>• One IP Address resource<br>• The following Exchange 2003 resources:<br>  • Exchange Message Transfer Agent instance (this resource is only in the first Exchange instance in the cluster)<br>  • Exchange Routing Service instance<br>  • SMTP virtual server instance<br>  • Exchange HTTP virtual server instance<br>  • Exchange IMAP4 virtual server instance<br>  • Exchange POP3 virtual server instance<br>  • Search Service instance<br>  • Exchange System Attendant<br>  • Exchange Information Store |
| Backup virtual servers | • Four Physical Disk resources<br>• Four File Share resources<br>• One IP Address resource<br>• One Network Name resource |
| Cluster virtual server | • One Physical Disk resource<br>• One IP Address resource<br>• One Network Name resource<br>• One Distribution Transaction Coordinator resource<br>• One Majority Node Set resource |

The following sections describe how OTG has configured the settings in Cluster Administrator. You can use Cluster Administrator to configure a cluster and its nodes, resource groups, and resources.

- Cluster Properties
- Exchange Virtual Server Properties
- Backup Virtual Server Properties

- Cluster Virtual Server Properties
- Resource Settings

**Cluster Properties**

This section describes how OTG has configured the server cluster as a whole.

The following settings appear on the **Quorum** tab that is available when you configure or view the cluster properties.

| Configuration setting | Explanation |
|---|---|
| **Quorum resource:**<br><br>Majority Node Set | In a majority node set server cluster, each node maintains its own copy of the cluster configuration data. The Majority Node Set resource, acting as the Quorum Resource, ensures that the cluster configuration data is kept consistent across the different nodes. The majority node set cluster is only meant for use in targeted scenarios like a geographically dispersed cluster, which is a cluster that spans multiple sites.<br><br>OTG has used a majority node set server cluster primarily for validation and testing purposes. Majority node set clusters also reduce the likelihood of a single point of failure in certain hardware devices (such a host bus adapters or fibre switches), which can cause the entire cluster to fail.<br><br>**Important**<br>Do not configure your cluster as a majority node set cluster unless you have a specific need to do so and it is part of a cluster solution offered by your Original Equipment Manufacturer (OEM), Independent Software Vendor (ISV), or Independent Hardware Vendor (IHV). The single quorum device server cluster will still cover the vast majority of your cluster deployment needs. |
| **Reset quorum log at:**<br><br>4096 KB | The maximum size of the quorum log is 4096 KB. This is enough space to hold the cluster configuration information, such as which servers are part of the cluster, what resources are installed in the cluster, and what state those resources are in (for example, online or offline). |

The following settings appear on the **Network Priority** tab that is available when you configure or view the clusters properties.

| Configuration setting | Explanation |
|---|---|
| **Networks used for internal cluster communications:**<br><br>• Internal<br><br>• Internal + Public | The Internal network has the highest priority. The Internal + Public network will only be used if the Internal network becomes unavailable. |

The following settings appear on the **Security** tab that is available when you configure or view the clusters properties.

| Configuration setting | Explanation |
|---|---|
| **Group or user names:**<br><br>• Administrators have full Control<br><br>• Network Service has full Control<br><br>• System has full Control | These groups have permission to alter the configuration of this cluster. |

**Exchange Virtual Server Properties**

This section describes how OTG has configured the Exchange virtual servers in the server cluster.

OTG has configured four Exchange virtual servers, which each host four storage groups. These storage groups contain a total of 20 databases. Each database is configured with a 200-MB limit per mailbox. This design specifies a maximum of 200 mailboxes per database, so each Exchange virtual server can contain a maximum of 4,000 mailboxes.

The following settings appear on the **General** tab that is available when you configure or view properties for the Exchange virtual servers.

| **Configuration setting** | **Explanation** |
|---|---|
| **Preferred owners:**<br>• Active node 1<br>• Primary passive node<br>• Alternate passive node 1<br>• Alternate passive node 2 | The Exchange virtual server will fail over to the indicated nodes in the specified order. If the active node fails, the Exchange virtual server will fail over to the primary passive node. If the primary passive node is unavailable, the Exchange virtual server will fail over to the alternate passive nodes. If none of the nodes are available, this resource group will fail over to a randomly selected node (this also depends on the **possible owner** settings for its resources). |

The following settings appear on the **Failover** tab that is available when you configure or view properties for the Exchange virtual servers.

| **Configuration setting** | **Explanation** |
|---|---|
| **Threshold** set to 10<br>**Period** set to 6 | The Exchange virtual server is allowed to fail 10 times within 6 hours before it is left in a failed state. The 11th time a resource in the resource group fails, the Cluster service fails all other resources in the group and leaves the entire group offline instead of failing over the group. |

The following setting appears on the **Failback** tab that is available when you configure or view properties for the Exchange virtual servers.

| **Configuration setting** | **Explanation** |
|---|---|
| **Prevent Failback** is checked. | The Exchange virtual server will not fail back when a failed node returns to operation. Instead, the Exchange virtual server will continue to run on the alternate node after the failed node comes back online. This requires administrator intervention to move the resource group back to the original, preferred node. OTG prefers this method because the administrator has complete control over when the failback occurs and can coordinate this to give users the least amount of downtime. |

**Backup Virtual Server Properties**

This section describes how OTG has configured the four backup virtual servers in the server cluster. For more information about how the backup process works, see How OTG Backs Up the Server Cluster later in this document.

The following setting appears on the **General** tab that is available when you configure or view the properties for the backup virtual servers.

| **Configuration setting** | **Explanation** |
|---|---|

| | |
|---|---|
| There are no **preferred owners** for this resource group. | OTG has created various scheduled tasks that control the movement of the backup virtual servers between nodes. The scheduled tasks force the backup virtual servers to move between the passive and active nodes, depending on the stage of the backup process that is being completed. The movement of the backup virtual servers between nodes is also determined by the **possible owner** settings of their resources. |

The following settings appear on the **Failover** tab that is available when you configure or view the properties for the backup virtual servers.

| **Configuration setting** | **Explanation** |
|---|---|
| **Threshold** set to 10<br><br>**Period** set to 6 | This resource group is allowed to fail 10 times within 6 hours before it is left in a failed state. The 11th time a resource in the group fails, the Cluster service fails all other resources in the group and leaves the entire group offline instead of failing over the group. |

The following setting appears on the **Failback** tab that is available when you configure or view the properties for the backup virtual servers.

| **Configuration setting** | **Explanation** |
|---|---|
| **Prevent Failback** is checked. | This resource group will not fail back when a failed node returns to operation. Instead, it will continue to run on the alternate node after the failed node comes back online. This requires administrator intervention to move the group back to the original, preferred node. OTG prefers this method because the administrator has complete control over when the failback occurs and can coordinate this to make sure that users experience the least amount of downtime. |

**Cluster Virtual Server Properties**

This section describes how OTG has configured the cluster virtual server in the server cluster. The cluster virtual server is the owner of the cluster quorum resource. The quorum resource maintains the configuration data that is necessary for recovery of the cluster. This data contains details of all of the changes that have been applied to the cluster database.

The following setting appears on the **General** tab that is available when you configure or view properties for the cluster virtual server.

| **Configuration setting** | **Explanation** |
|---|---|
| **Preferred owners:**<br>• Primary passive node | This resource group will first attempt to fail over to the primary passive node because it is best equipped to host the group. If this node is not available, this group will fail over to a randomly selected node (this also depends on the **possible owner** settings of its resources). |

The following settings appear on the **Failover** tab that is available when you configure or view the properties for the cluster virtual server.

| **Configuration setting** | **Explanation** |
|---|---|
| **Threshold** set to 10<br><br>**Period** set to 6 | This resource group is allowed to fail 10 times within 6 hours before it is left in a failed state. The 11th time a resource in the group fails, the Cluster service fails all other resources in the group and leaves the entire group offline instead of failing over the group. |

The following settings appear on the **Failback** tab that is available when you configure or view the properties for the cluster virtual server.

| Configuration setting | Explanation |
| --- | --- |
| **Prevent Failback** is checked. | This resource group will not fail back when a failed node returns to operation. Instead, it will continue to run on the alternate node after the failed node comes back online. This requires administrator intervention to move the group back to the original, preferred node. OTG prefers this method because the administrator has complete control over when the failback occurs and can coordinate this to make sure that users experience the least amount of downtime. |

**Resource Settings**

This section describes how each of the resources used in the server cluster is configured.

All resource settings in the server cluster are the default values with the following exceptions. These exceptions appear on the **Advanced** tab that is available when you configure or view the properties for the resource:

- **Restart** is not enabled for the Exchange Information Store and the Exchange System Attendant resources.
- **Restart** is enabled for all other resources (default) but the **Affect the Group** option is disabled.

For more information regarding why these changes were made, see each specific setting below

For more information about cluster resources, see the information about standard resource types at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=19502).

For more information about Exchange resources, see Product Documentation for Exchange at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=116208).

| **Note** |
| --- |
| Each resource of a specific type is configured identically unless otherwise noted. |

**IP Address Resource**

The following settings appear on the **General** tab that is available when you configure or view the IP Address resource properties.

| Configuration setting | Explanation |
| --- | --- |
| All nodes are **possible owners**. | This resource can run on any of the available nodes, which allows failover to occur. During failover, this resource will fail over to a node based on the **preferred owner** settings of the resource group it belongs to. |
| **Run as a separate Resource Monitor** is not checked. | This resource does not run on a separate memory space. |

The following setting appears on the **Dependencies** tab that is available when you configure or view the IP Address resource properties.

| Configuration setting | Explanation |
| --- | --- |
| This resource has no **Resource Dependencies**. | The Cluster service uses the dependencies list when bringing resources online and offline. This resource does not require another resource to operate. |

The following settings appear on the **Advanced** tab that is available when you configure or view the IP Address resource properties.

| Configuration setting | Explanation |
| --- | --- |

| | |
|---|---|
| **Restart** is not enabled.<br><br>The **Affect the group** check box is not selected (with a threshold of 3 and a period of 900 seconds). | This resource restarts after a resource failure.<br><br>The entire resource group will not be failed over when this resource fails. When this resource is in a pending state, the Cluster service will attempt to restart the resource 3 times within 900 seconds before the Cluster service assigns the resource to Offline or Failed status.<br><br>**Important**<br>All resources that restart after a resource failure are not configured to **Affect the group**. This is not a default setting. This allows OTG to track and fix any problems within the cluster. If a resource failure occurs, this setting allows OTG administrators to notify the Microsoft developers of possible bugs in the code.<br><br>This setting also inhibits a cluster resource group from alternating between nodes when a resource fails, which would result in significant instability. This problem can result when a failed disk has default properties that cause a series of bus resets. The resource group then attempts to acquire reservation on a disk that does not exist.<br><br>This configuration is only suitable for users with strong monitoring systems who want to control the location of their groups. |
| **Looks Alive poll interval:**<br>• resource type value | The Cluster service will perform a cursory check the default number of times to determine if a resource is in the Online state. |
| **Is Alive poll interval:**<br>• resource type value | The Cluster service will perform an in-depth check the default number of times to determine if the resource is in the Online state. |
| **Pending timeout** is set at 180 seconds. | This resource has 180 seconds when in a pending state (Online Pending or Offline Pending) to resolve its status, or at least report to the Resource Monitor that it is making progress, before the Cluster service will put the resource in Failed status. |

The following settings appear on the **Parameters** tab that is available when you configure or view the IP Address resource properties.

| Configuration setting | Explanation |
|---|---|
| **Address:**<br><br>XXX.XX.XX.XXX | This parameter specifies the unique IP address. This address is a static IP address that is not already in use on the network. This was obtained by the network administrator. |
| **Subnet mask:**<br><br>XXX.XXX.XXX.XXX | This parameter denotes the subnet mask for the IP address specified. This subnet mask is the same subnet mask as the one for the associated cluster network. This was obtained by the network administrator. |
| **Network**:<br><br>Internal + Public | The mixed communication adapter is being used. For more information, see Network Adapters. |
| **Enable NetBIOS for this address** is checked. | The NetBIOS is enabled for the specified IP address. This allows a dependent network name to be published using NetBIOS. |

**Network Name Resource**

The following settings appear on the **General** tab that is available when you configure or view the Network Name resource properties.

| Configuration setting | Explanation |
|---|---|

| | |
|---|---|
| All nodes are **possible owners**. | This resource can run on any of the available nodes, which allows failover to occur. During failover, this resource will fail over to a node based on the **preferred owner** settings of the resource group it belongs to. |
| **Run as a separate Resource Monitor** is not checked. | This resource does not run on a separate memory space. |

The following setting appears on the **Dependencies** tab that is available when you configure or view the Network Name resource properties.

| Configuration setting | Explanation |
|---|---|
| **Resource Dependencies:**<br>• IP Address Resource | The Cluster service uses the dependencies list when bringing resources online and offline. A dependent resource is one that requires another resource to operate. The network name resource is dependent on the IP Address that you want to correspond to the name. |

The following settings appear on the **Advanced** tab that is available when you configure or view the Network Name resource properties.

| Configuration setting | Explanation |
|---|---|
| **Restart** is not enabled.<br><br>The **Affect the group** check box is not selected (with a threshold of 3 and a period of 900 seconds). | This resource restarts after a resource failure.<br><br>The entire resource group will not be failed over when this resource fails. When this resource is in a pending state, the Cluster service will attempt to restart the resource 3 times within 900 seconds before the Cluster service assigns the resource to Offline or Failed status.<br><br>**Important**<br>All resources that restart after a resource failure are not configured to **Affect the group**. This is not a default setting. This allows OTG to track and fix any problems within the cluster. If a resource failure occurs, this setting allows OTG administrators to notify the Microsoft developers of possible bugs in the code.<br><br>This setting also inhibits a cluster resource group from alternating between nodes when a resource fails, which would result in significant instability. This problem can result when a failed disk has default properties that cause a series of bus resets. The resource group then attempts to acquire reservation on a disk that does not exist.<br><br>This configuration is only suitable for users with strong monitoring systems who want to control the location of their groups. |
| **Looks Alive poll interval:**<br>• resource type value | The Cluster service will perform a cursory check the default number of times to determine if a resource is in the Online state. |
| **Is Alive poll interval:**<br>• resource type value | The Cluster service will perform an in-depth check the default number of times to determine if the resource is in the Online state. |
| **Pending timeout** is set at 180 seconds. | This resource has 180 seconds when in a pending state (Online Pending or Offline Pending) to resolve its status, or at least report to the Resource Monitor that it is making progress, before the Cluster service will put the resource in Failed status. |

The following settings appear on the **Parameters** tab that is available when you configure or view the Network Name resource properties.

| Configuration setting | Explanation | |
|---|---|---|
| **DNS Registration must succeed.** | Exchange virtual server:<br><br>Is not checked. | Domain Name System registration failures are tolerated for this Network Name resource. |
| | Backup virtual server:<br><br>Is not checked. | Domain Name System registration failures are tolerated for this Network Name resource. |
| | Cluster virtual server:<br><br>Is not checked. | Domain Name System registration failures are tolerated for this Network Name resource. |
| **Enable Kerberos Authentication** | Exchange virtual server:<br><br>Is checked. | Clients connecting to the network name represented by this resource can use Kerberos Authentication if they chose to. The System Attendant will not come online unless the Kerberos is enabled in the Network Name for the Exchange virtual server. |
| | Backup virtual server:<br><br>Is not checked. | Clients connecting to the network name represented by this resource can only use NTLM authentication. This is the default configuration. |
| | Cluster virtual server:<br><br>Is not checked. | Clients connecting to the network name represented by this resource can only use NTLM authentication. This is the default configuration. |

**Physical Disk Resource**

The following settings appear on the **General** tab that is available when you configure or view the Physical Disk resource properties.

| Configuration setting | Explanation |
|---|---|
| All nodes are **possible owners**. | This resource can run on any of the available nodes, which allows failover to occur. During failover, this resource will fail over to a node based on the **preferred owner** settings of the resource group it belongs to. |
| **Run as a separate Resource Monitor** is not checked. | This resource does not run on a separate memory space. |

The following settings appear on the **Dependencies** tab that is available when you configure or view the Physical Disk resource properties.

| Configuration setting | Explanation | |
|---|---|---|

| | | |
|---|---|---|
| **Resource Dependencies** | Exchange virtual server:<br><br>Dependencies vary. | The Cluster service uses the dependencies list when bringing resources online and offline. This resource can have varying dependencies based on the following:<br><br>• Mount point disks are dependent on their parent drive.<br><br>• Log drives are dependent on the data drives.<br><br>• Data drives have no dependencies. |
| | Backup virtual server:<br><br>Dependencies vary. | The Cluster service uses the dependencies list when bringing resources online and offline. This resource can have varying dependencies, which enable the transfer of data during backup. |
| | Cluster virtual server:<br><br>Has no dependencies. | The Cluster service uses the dependencies list when bringing resources online and offline. This resource does not require another resource to operate. |

The following settings appear on the **Advanced** tab that is available when you configure or view the Physical Disk resource properties.

| Configuration setting | Explanation |
|---|---|
| **Restart** is not enabled.<br><br>The **Affect the group** check box is not selected (with a threshold of 3 and a period of 900 seconds). | This resource restarts after a resource failure.<br><br>The entire resource group will not be failed over when this resource fails. When this resource is in a pending state, the Cluster service will attempt to restart the resource 3 times within 900 seconds before the Cluster service assigns the resource to Offline or Failed status.<br><br>**Important**<br>All resources that restart after a resource failure are not configured to **Affect the group**. This is not a default setting. This allows OTG to track and fix any problems within the cluster. If a resource failure occurs, this setting allows OTG administrators to notify the Microsoft developers of possible bugs in the code.<br><br>This setting also inhibits a cluster resource group from alternating between nodes when a resource fails, which would result in significant instability. This problem can result when a failed disk has default properties that cause a series of bus resets. The resource group then attempts to acquire reservation on a disk that does not exist.<br><br>This configuration is only suitable for users with strong monitoring systems who want to control the location of their groups. |
| **Looks Alive poll interval:**<br><br>• resource type value | The Cluster service will perform a cursory check the default number of times to determine if a resource is in the Online state. |
| **Is Alive poll interval:**<br><br>• resource type value | The Cluster service will perform an in-depth check the default number of times to determine if the resource is in the Online state. |
| **Pending timeout** is set at 180 seconds. | This resource has 180 seconds when in a pending state (Online Pending or Offline Pending) to resolve its status, or at least report to the Resource Monitor that it is making progress, before the Cluster service will put the resource in Failed status. |

**File Share Resource**

The following settings appear on the **General** tab that is available when you configure or view the File Share resource

properties.

| Configuration setting | Explanation |
|---|---|
| All nodes are **possible owners**. | This resource can run on any of the available nodes, which allows failover to occur. During failover, this resource will fail over to a node based on the **preferred owner** settings of the resource group it belongs to. |
| **Run as a separate Resource Monitor** is not checked. | This resource does not run on a separate memory space. |

The following setting appears on the **Dependencies** tab that is available when you configure or view the File Share resource properties.

| Configuration setting | Explanation |
|---|---|
| **Resource Dependencies** vary. | The Cluster service uses the dependencies list when bringing resources online and offline. This resource can have varying dependencies which enable the transfer of data during backup. |

The following settings appear on the **Advanced** tab that is available when you configure or view the File Share resource properties.

| Configuration setting | Explanation |
|---|---|
| **Restart** is not enabled.<br><br>The **Affect the group** check box is not selected (with a threshold of 3 and a period of 900 seconds). | This resource restarts after a resource failure.<br><br>The entire resource group will not be failed over when this resource fails. When this resource is in a pending state, the Cluster service will attempt to restart the resource 3 times within 900 seconds before the Cluster service assigns the resource to Offline or Failed status.<br><br>**Important**<br>All resources that restart after a resource failure are not configured to **Affect the group**. This is not a default setting. This allows OTG to track and fix any problems within the cluster. If a resource failure occurs, this setting allows OTG administrators to notify the Microsoft developers of possible bugs in the code.<br><br>This setting also inhibits a cluster resource group from alternating between nodes when a resource fails, which would result in significant instability. This problem can result when a failed disk has default properties that cause a series of bus resets. The resource group then attempts to acquire reservation on a disk that does not exist.<br><br>This configuration is only suitable for users with strong monitoring systems who want to control the location of their resource groups. |
| **Looks Alive poll interval:**<br>• resource type value | The Cluster service will perform a cursory check the default number of times to determine if a resource is in the Online state. |
| **Is Alive poll interval:**<br>• resource type value | The Cluster service will perform an in-depth check the default number of times to determine if the resource is in the Online state. |
| **Pending timeout** is set at 180 seconds. | This resource has 180 seconds when in a pending state (Online Pending or Offline Pending) to resolve its status, or at least report to the Resource Monitor that it is making progress, before the Cluster service will put the resource in Failed status. |

The following setting appears on the **Parameters** tab that is available when you configure or view the File Share resource properties.

| Configuration setting | Explanation |
| --- | --- |
| The maximum **User Limit** is allowed. | OTG has not limited the number of simultaneous users. |

**Majority Node Set Resource**

The following settings appear on the **General** tab that is available when you configure or view the Majority Node Set resource properties.

| Configuration setting | Explanation |
| --- | --- |
| All nodes are **possible owners**. | This resource can run on any of the available nodes, which allows failover to occur. During failover, this resource will fail over to a node based on the **preferred owner** settings of the resource group it belongs to. |
| **Run as a separate Resource Monitor** is not checked. | This resource does not run on a separate memory space. |

The following setting appears on the **Dependencies** tab that is available when you configure or view the Majority Node Set resource properties.

| Configuration setting | Explanation |
| --- | --- |
| This resource has no **Resource Dependencies**. | The Cluster service uses the dependencies list when bringing resources online and offline. This resource does not require another resource to operate. |

The following settings appear on the **Advanced** tab that is available when you configure or view the Majority Node Set resource properties.

| Configuration setting | Explanation |
| --- | --- |
| **Restart** is not enabled.<br><br>The **Affect the group** check box is not selected (with a threshold of 3 and a period of 900 seconds). | This resource restarts after a resource failure.<br><br>The entire resource group will not be failed over when this resource fails. When this resource is in a pending state, the Cluster service will attempt to restart the resource 3 times within 900 seconds before the Cluster service assigns the resource to Offline or Failed status.<br><br>**Important**<br>All resources that restart after a resource failure are not configured to **Affect the group**. This is not a default setting. This allows OTG to track and fix any problems within the cluster. If a resource failure occurs, this setting allows OTG administrators to notify the Microsoft developers of possible bugs in the code.<br><br>This setting also inhibits a cluster resource group from alternating between nodes when a resource fails, which would result in significant instability. This problem can result when a failed disk has default properties that cause a series of bus resets. The resource group then attempts to acquire reservation on a disk that does not exist.<br><br>This configuration is only suitable for users with strong monitoring systems who want to control the location of their groups. |
| **Looks Alive poll interval:**<br>• resource type value | The Cluster service will perform a cursory check the default number of times to determine if a resource is in the Online state. |
| **Is Alive poll interval:**<br>• resource type value | The Cluster service will perform an in-depth check the default number of times to determine if the resource is in the Online state. |

| | |
|---|---|
| **Pending timeout** is set at 180 seconds. | This resource has 180 seconds when in a pending state (Online Pending or Offline Pending) to resolve its status, or at least report to the Resource Monitor that it is making progress, before the Cluster service will put the resource in Offline status. |

**Distribution Transaction Coordinator Resource**

The following settings appear on the **General** tab that is available when you configure or view the Distribution Transaction Coordinator resource properties.

| **Note** |
|---|
| OTG has implemented a Distribution Transaction Coordinator resource because Exchange 2003 requires this resource in order to operate. |

| Configuration setting | Explanation |
|---|---|
| All nodes are **possible owners**. | This resource can run on any of the available nodes, which allows failover to occur. During failover, this resource will fail over to a node based on the **preferred owner** settings of the resource group it belongs to. |
| **Run as a separate Resource Monitor** is not checked. | This resource does not run on a separate memory space. |

The following setting appears on the **Dependencies** tab that is available when you configure or view the Distribution Transaction Coordinator resource properties.

| Configuration setting | Explanation |
|---|---|
| **Resource Dependencies:**<br>• Network Name<br>• Quorum Physical Disk | The Cluster service uses the dependencies list when bringing resources online and offline. This resource requires the Network Name resource and the Quorum physical disk resource in order to operate. |

The following settings appear on the **Advanced** tab that is available when you configure or view the Distribution Transaction Coordinator resource properties.

| Configuration setting | Explanation |
|---|---|

| | |
|---|---|
| **Restart** is not enabled.<br><br>The **Affect the group** check box is not selected (with a threshold of 3 and a period of 900 seconds). | This resource restarts after a resource failure.<br><br>The entire resource group will not be failed over when this resource fails. When this resource is in a pending state, the Cluster service will attempt to restart the resource 3 times within 900 seconds before the Cluster service assigns the resource to Offline or Failed status.<br><br>**Important**<br>All resources that restart after a resource failure are not configured to **Affect the group**. This is not a default setting. This allows OTG to track and fix any problems within the cluster. If a resource failure occurs, this setting allows OTG administrators to notify the Microsoft developers of possible bugs in the code.<br><br>This setting also inhibits a cluster resource group from alternating between nodes when a resource fails, which would result in significant instability. This problem can result when a failed disk has default properties that cause a series of bus resets. The resource group then attempts to acquire reservation on a disk that does not exist.<br><br>This configuration is only suitable for users with strong monitoring systems who want to control the location of their groups. |
| **Looks Alive poll interval:**<br><br>resource type value | The Cluster service will perform a cursory check the default number of times to determine if a resource is in the Online state. |
| **Is Alive poll interval:**<br><br>resource type value | The Cluster service will perform an in-depth check the default number of times to determine if the resource is in the Online state. |
| **Pending timeout** is set at 180 seconds. | This resource has 180 seconds when in a pending state (Online Pending or Offline Pending) to resolve its status, or at least report to the Resource Monitor that it is making progress, before the Cluster service will put the resource in Failed status. |

**Exchange Resources**

This section describes how OTG has configured the following Exchange resource properties. The Exchange System Attendant resource and Exchange Information Store resource have different configurations, which are described at the end of this section.

- Exchange Message Transfer Agent instance
- Exchange Routing Service instance
- Exchange SMTP virtual server instance
- Exchange HTTP virtual server instance
- Exchange IMAP4 virtual server instance
- Exchange POP3 virtual server instance
- Search Service instance

**Note**

After the Exchange System Attendant resource is created, the cluster automatically generates all other Exchange resources.

The following settings appear on the **General** tab that is available when you configure or view the various Exchange resource properties.

| Configuration setting | Explanation |
|---|---|

| | |
|---|---|
| All nodes are **possible owners**. | This resource can run on any of the available nodes, which allows failover to occur. During failover, this resource will fail over to a node based on the **preferred owner** settings of the resource group it belongs to. |
| **Run as a separate Resource Monitor** is not checked. | These resources do not run on a separate memory space. |

The following setting appears on the **Dependencies** tab that is available when you configure or view the various Exchange resource properties.

| Configuration setting | Explanation |
|---|---|
| **Resource Dependencies:**<br>• System Attendant | The Cluster service uses the dependencies list when bringing resources online and offline. These resources require the System Attendant resource in order to operate. |

The following settings appear on the **Advanced** tab that is available when you configure or view any of the various Exchange resource properties.

| Configuration setting | Explanation |
|---|---|
| **Restart** is not enabled.<br><br>The **Affect the group** check box is not selected (with a threshold of 3 and a period of 900 seconds). | This resource restarts after a resource failure.<br><br>The entire resource group will not be failed over when this resource fails. When this resource is in a pending state, the Cluster service will attempt to restart the resource 3 times within 900 seconds before the Cluster service assigns the resource to Offline or Failed status.<br><br>**Important**<br>All resources that restart after a resource failure are not configured to **Affect the group**. This is not a default setting. This allows OTG to track and fix any problems within the cluster. If a resource failure occurs, this setting allows OTG administrators to notify the Microsoft developers of possible bugs in the code.<br><br>This setting also inhibits a cluster resource group from alternating between nodes when a resource fails, which would result in significant instability. This problem can result when a failed disk has default properties that cause a series of bus resets. The resource group then attempts to acquire reservation on a disk that does not exist.<br><br>This configuration is only suitable for users with strong monitoring systems who want to control the location of their groups. |
| **Looks Alive poll interval:**<br>• resource type value | The Cluster service will perform a cursory check the default number of times to determine if a resource is in the Online state. |
| **Is Alive poll interval:**<br>• resource type value | The Cluster service will perform an in-depth check the default number of times to determine if the resource is in the Online state. |
| **Pending timeout** is set at 180 seconds. | This resource has 180 seconds when in a pending state (Online Pending or Offline Pending) to resolve its status, or at least report to the Resource Monitor that it is making progress, before the Cluster service will put the resource in Failed status. |

The following settings appear on the **General** tab that is available when you configure or view the Exchange System Attendant resource properties.

| Configuration setting | Explanation |
| --- | --- |
| **Possible Owners:**<br>• Active node 1<br>• Primary passive node<br>• Alternate passive node 1<br>• Alternate passive node 2 | This resource drives the failover path for the Exchange virtual servers because this resource can only run on these specified nodes. If the active node fails, the resource group will fail over to the primary passive node. If the primary passive node is unavailable, the group will fail over to the alternate passive nodes. This ensures that this resource group can only become active on its corresponding active nodes and all passive nodes. |
| **Run as a separate Resource Monitor** is not checked. | This resource does not run on a separate memory space. |

The following setting appears on the **Dependencies** tab that is available when you configure or view the Exchange System Attendant resource properties.

| Configuration setting | Explanation |
| --- | --- |
| **Resource Dependencies:**<br>• 5 physical disks<br>• Network Name | The Cluster service uses the dependencies list when bringing resources online and offline. This System Attendant resource is dependent on five mount drive devices (which are dependent on their parent drives) and the Network name resource. The System Attendant resource requires these resources in order to operate. |

The following settings appear on the **Advanced** tab that is available when you configure or view the Exchange System Attendant resource properties.

| Configuration setting | Explanation |
| --- | --- |
| **Restart** is not enabled. | This resource does not restart after a resource failure. This allows OTG to track and fix any problems within the cluster. If a resource failure occurs, this setting allows OTG administrators to notify the Microsoft developers of possible bugs in the code.<br><br>This setting is specific to OTG as it allows for more precise fault isolation that can then be reported to the server cluster developers. It is not recommended for typical customer deployments. |

| | |
|---|---|
| **Looks Alive poll interval:**<br>• resource type value | The Cluster service will perform a cursory check the default number of times to determine if this resource is in the Online state. |
| **Is Alive poll interval:**<br>• resource type value | The Cluster service will perform an in-depth check the default number of times to determine if this resource is in the Online state. |
| **Pending timeout** is set at 180 seconds. | This resource has 180 seconds when in a pending state (Online Pending or Offline Pending) to resolve its status, or at least report to the Resource Monitor that it is making progress, before the Cluster service will put the resource in Failed status. |

The following settings appear on the **General** tab that is available when you configure or view the Exchange Information Store resource properties.

| Configuration setting | Explanation |
|---|---|
| All nodes are **possible owners**. | This resource can run on any of the available nodes, which allows failover to occur. During failover, this resource will fail over to a node based on the preferred owner settings of the resource group it belongs to. |
| **Run as a separate Resource Monitor** is not checked. | This resource does not run on a separate memory space. |

The following setting appears on the **Dependencies** tab that is available when you configure or view the Exchange Information Store resource properties.

| Configuration setting | Explanation |
|---|---|
| **Resource Dependencies:**<br>System Attendant | The Cluster service uses the dependencies list when bringing resources online and offline. This resource requires the System Attendant resource in order to operate. |

The following settings appear on the **Advanced** tab that is available when you configure or view the Exchange Information Store Resource Properties.

| Configuration setting | Explanation |
|---|---|
| **Restart** is not enabled. | This setting is specific to OTG as it allows for more precise fault isolation that can then be reported to the server cluster developers. It is not recommended for typical customer deployments. |
| **Looks Alive poll interval:**<br>• resource type value | The Cluster service will perform a cursory check the default number of times to determine if this resource is in the Online state. |

| | |
|---|---|
| **Is Alive poll interval:**<br>• resource type value | The Cluster service will perform an in-depth check the default number of times to determine if this resource is in the Online state. |
| **Pending timeout** is set at 180 seconds. | This resource has 180 seconds when in a pending state (Online Pending or Offline Pending) to resolve its status, or at least report to the Resource Monitor that it is making progress, before the Cluster service will put the resource in Failed status. |

# How OTG Backs Up the Server Cluster

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Diagram 4 illustrates the process that the Operations and Technology Group (OTG) uses to back up the mailbox data in the Exchange databases. This backup strategy starts with a defined Service Level Agreement (SLA) that requires OTG to maintain a one-hour backup restore window per database.

Backup is divided into two stages. Stage 1 is a disk-to-disk backup, and stage 2 is a disk-to-tape backup. Stage 1 is a traditional full backup using the Exchange Backup API and is performed daily. The active cluster backs up data to locally mounted backup disks. Each storage group on an active node backs up its individual databases in sequential order, starting with database A and finishing with E. The logs get truncated daily after OTG completes a successful backup.

In Stage 2, the backup disks are failed over to the alternate passive node(s) and then backed up to tape through fibre attached tape devices. This backup can run at any period of the day without impacting server performance.

For more information about how the Exchange data storage is structured, see Storage Components. For additional information about the storage configuration for each virtual server, see Appendix A.




**Diagram 4 Backup Process for the Server Cluster**

# How OTG Monitors the Server Cluster

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section briefly discusses the tools that the administrators in the Operations and Technology Group (OTG) use to monitor the condition and configuration of the Exchange server cluster. Proactive monitoring has several key components that include managing performance, configuration, and fault tolerance. The primary monitoring tool OTG uses is Microsoft Operations Manager (MOM).

**Microsoft Operations Manager**
OTG relies on Microsoft Operations Manager (MOM) as a proactive tool to monitor the health of its Exchange environment. MOM consolidates event performance and alert data in a centralized database. Cluster administrators can use this data to anticipate and react to server conditions.

MOM includes predefined downloadable solutions called Management Pack modules. Management Pack modules provide immediate monitoring of specific application and environments such as Exchange 2003. The Exchange module includes built-in knowledge of common problems. In addition, documentation is included to begin the steps to resolve issues when MOM identifies problems.

With the Exchange module, OTG can proactively manage and report on Exchange 2003 installations, which helps to avoid outages. For example, the Management Pack Exchange module performs the following tasks to alert OTG of possible critical conditions:

- Monitors vital performance data, which can indicate that the server is running low on resources.
- Collects important warning and error events from Exchange 2003 servers and alerts operators to problems.
- Monitors disk capacity and alerts operators when disk capacity is running low.
- Provides knowledge about which Exchange files are on the affected drives.
- Monitors the Exchange services that are expected to be running on a specific server.
- Monitors whether an Exchange database can actually be reached by a MAPI client logon, which verifies both the Exchange database and the Active Directory functionality.
- Monitors high queue lengths that are caused when e-mail cannot be sent to a destination server.
- Monitors and alerts operators to a high number of simultaneous connections, which can indicate a denial-of-service attack.
- Monitors configuration errors or resource shortages affecting service levels.

In addition, OTG uses MOM to monitor backup and restore of all Exchange servers from a central console.

**Other Monitoring Tools**
OTG uses a variety of other tools besides MOM to monitor and report the overall performance of servers and Exchange 2003.

- **HP Storage Management Appliance (SMA)**. The SMA is a centralized, appliance-based monitoring and management server for the storage area network (SAN) fabric. The SMA performs management functions without involving the host servers and allows data transfers to proceed independently between servers and storage devices whether the appliance is operating or not. OTG primarily uses the SMA to remotely manage regional and local SANs.
- **HP Open SAN Manager**. This tool provides a Web-based aggregation and entry point (or portal) for all SMA software applications. With this tool, administrators can organize, visualize, configure, and monitor storage from a single navigation point on the SAN. It also provides a centralized launch site for a variety of external management applications (that is, fabric management).

- **Command View EVA (HSV Element Manager)**. This is a graphical storage monitoring and configuration tool for the HP StorageWorks Enterprise Virtual Array (EVA). Administrators can use it to remotely manage HP SANs.

- **HP StorageWorks Secure Path**. Secure Path is multipath software that manages and maintains continuous data access to HP storage systems to ensure that no single point of failure occurs from server to storage. Secure Path is host-resident software that monitors the data paths between server and storage to increase availability of information.

- **Secure Path Manager 4.0x**. This tool provides simplified management of secure paths through a Web-based GUI that can integrate with the Storage Management Appliance. This tool can host and store views of preferred and available paths.

# Best Practices for Server Clusters

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section describes best practices for deploying a server cluster.

**Planning Your Cluster**
Before designing your cluster, analyze how your messaging system is used and design your cluster in a way that will accommodate peak usage periods. For example, the entire configuration of the server cluster that is discussed in this document (hardware, software and cluster configuration) is designed to meet the requirements for peak usage. At Microsoft, Monday mornings are the busiest time for e-mail, as employees attempt to catch up from the weekend. Every server cluster configuration should accommodate the peak usage periods.

**Planning the Hardware for Your Cluster**

- Use identical network adapters in all cluster nodes. Each adapter should be the same make, model, and firmware version.
- Reserve one network exclusively for internal communication between cluster nodes. This is the private network. Use other networks for communication with clients. These are public networks. Having two networks ensures that the network will be available for internal cluster communications.
- Do not use teaming network adapters on the private network. Teaming network adapters, which are network adapters that are grouped on multiple ports to a single physical network segment, are used to provide fault tolerance.
- Create a standard naming system for your resources, nodes, disks, IP Addresses, and backup sets. This will increase clarity for anyone working with the cluster.

**Configuring the Network Adapters**

- Use static IP addresses for all nodes on the private network.
- Use static IP addresses for all nodes on all public networks. We do not recommend using dynamic configuration with Dynamic Host Configuration Protocol (DHCP). Failure to renew a lease of an IP Address could disrupt cluster operation.
- Do not configure DNS servers, WINS servers, or a default gateway on the private network adapters.
- Configure WINS or DNS servers on the public network adapters. If network name resources will be deployed on the public network, then the DNS servers should support dynamic updates.
- In Cluster Administrator, set the network connection order for each node to be:
    - Private network highest priority
    - Public network(s)
    - Remote access connections lowest priority
- Change the default name for each network connection to clearly identify the use of each network. For example, you might change the name of the private network connection from Local Area Connection (x) to Private Cluster Network.
- Isolate the internal network from public network communication. Only nodes that are part of the cluster should be connected to the private subnet. Where there are several clusters, using the same subnet for the private network for all of the clusters is reasonable.

**Configuring the Cluster Service**

- Set the private network role to Internal cluster communications only. Verify that the role for each public network is set to All communications, which is the default value. These settings are available in Cluster Administrator.

- In Cluster Administrator, configure the private network to be the highest priority for internal cluster communication. In addition, make sure the Public network is first in the TCP/IP Binding Order.

**More Information**

- For more server cluster best practices, see Best Practices for Configuring and Operating Server Clusters at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=19500).

- For more information about best practices for Enterprise Virtual Arrays, see EVA Best Practices - Cost, Performance, and Availability and other related white papers in the Storage Array Systems White Papers section of the Hewlett-Packard Web site (http://go.microsoft.com/fwlink/?LinkId=19716).

# Summary (Using Clustering with Exchange 2003: An Example)

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Operation and Technology Group (OTG) at Microsoft uses Windows Server 2003 server cluster technology in a typical way, with a few changes that maximize the performance of the hardware. Because OTG selected these configurations (both server and storage) for deployment at Microsoft, they might not meet your specific deployment needs. IT professionals should use the Microsoft implementation as an example when configuring their own Exchange server cluster. As always, careful analysis, planning, testing, and piloting should precede any Exchange deployment and cluster design decisions. Ensure that you apply all applicable best practices and recommendations when making your own design decisions. In addition, ensure that you have properly characterized your environment before server configuration choices are made and hardware is purchased. When properly used and implemented, server cluster technology can create more reliable and available messaging

# Related Topics (Using Clustering with Exchange 2003: An Example)

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

See the following resources for more information:

- For more information about server clusters, search the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=4441) using the key words server clusters and clustering.
- You can also see Availability and Scalability/Windows Clustering in Help and Support Center on any product in the Windows Server 2003 family.
- To view Help and Support Center in the Windows Server 2003 family, click **Start**, and then click **Help and Support**.
- For more information about deploying server clusters, see Planning for High Availability and Scalability at the Windows Server 2003 Deployment Kit Web site (http://go.microsoft.com/fwlink/?LinkId=4742).
- For information about the Windows Server 2003 family, see Product Documentation for Windows Server 2003 at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=12908).

# Appendix A (Using Clustering with Exchange 2003: An Example)

Updated: September 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following table shows the disk drives, data types, disk groups, Enterprise Virtual Arrays, sizes, Vraid types, and Logical Unit Numbers (LUNs) that correspond to each Exchange virtual server in the Exchange server cluster that is described in this document.

The server cluster contains four Exchange virtual servers (EVS), four backup virtual servers, and one cluster virtual server (CVS). The Exchange virtual servers include all of the resources that are required to form an Exchange server.

| Host name | Disk drive letters | Data type | Disk group | EVA | Size/ GB | RAID type | LUN |
|---|---|---|---|---|---|---|---|
| EVS1 | E:\Exchsrvr (Mount point) | SMTP | DG1 | EVA-1 | 50 | 1 | 1 |
| | E$ | SG1-Data | DG1 | EVA-1 | 350 | 1 | 2 |
| | F$ | SG2-Data | DG1 | EVA-1 | 350 | 1 | 3 |
| | G$ | SG3-Data | DG1 | EVA-1 | 350 | 1 | 4 |
| | H$ | SG4-Data | DG1 | EVA-1 | 350 | 1 | 5 |
| | U$ | SG1-Backup | DG4 | EVA-1 | 350 | 5 | 37 |
| | U:\Dump2 (Mount point) | SG2-Backup | DG4 | EVA-1 | 350 | 5 | 38 |
| | U:\Dump3 (Mount point) | SG3-Backup | DG4 | EVA-1 | 350 | 5 | 39 |
| | U:\Dump4 (Mount point) | SG4-Backup | DG4 | EVA-1 | 350 | 5 | 40 |
| | E:\Log (Mount point) | SG1-Log | DG3 | EVA-1 | 40 | 1 | 6 |
| | F:\Log (Mount point) | SG2-Log | DG3 | EVA-1 | 40 | 1 | 7 |
| | G:\Log (Mount point) | SG3-Log | DG3 | EVA-1 | 40 | 1 | 8 |
| | H:\Log (Mount point) | SG4-Log | DG3 | EVA-1 | 40 | 1 | 9 |
| EVS2 | I:\Exchsrvr (Mount point) | SMTP | DG2 | EVA-1 | 50 | 1 | 10 |
| | I$ | SG1-Data | DG2 | EVA-1 | 350 | 1 | 11 |
| | J$ | SG2-Data | DG2 | EVA-1 | 350 | 1 | 12 |
| | K$ | SG3-Data | DG2 | EVA-1 | 350 | 1 | 13 |
| | L$ | SG4-Data | DG2 | EVA-1 | 350 | 1 | 14 |
| | V$ | SG1-Backup | DG5 | EVA-1 | 350 | 5 | 41 |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| | V:\Dump2 (Mount point) | SG2-Backup | DG5 | EVA-1 | 350 | 5 | 42 |
| | V:\Dump3 (Mount point) | SG3-Backup | DG5 | EVA-1 | 350 | 5 | 43 |
| | V:\Dump4 (Mount point) | SG4-Backup | DG5 | EVA-1 | 350 | 5 | 44 |
| | I:\Log (Mount point) | SG1-Log | DG3 | EVA-1 | 40 | 1 | 15 |
| | J:\Log (Mount point) | SG2-Log | DG3 | EVA-1 | 40 | 1 | 16 |
| | K:\Log (Mount point) | SG3-Log | DG3 | EVA-1 | 40 | 1 | 17 |
| | L:\Log (Mount point) | SG4-Log | DG3 | EVA-1 | 40 | 1 | 18 |
| EVS3 | M:\Exchsrvr (Mount point) | SMTP | DG1 | EVA-2 | 50 | 1 | 19 |
| | M$ | SG1-Data | DG1 | EVA-2 | 350 | 1 | 20 |
| | N$ | SG2-Data | DG1 | EVA-2 | 350 | 1 | 21 |
| | O$ | SG3-Data | DG1 | EVA-2 | 350 | 1 | 22 |
| | P$ | SG4-Data | DG1 | EVA-2 | 350 | 1 | 23 |
| | W$ | SG1-Backup | DG4 | EVA-2 | 350 | 5 | 45 |
| | W:\Dump2 (Mount point) | SG2-Backup | DG4 | EVA-2 | 350 | 5 | 46 |
| | W:\Dump3 (Mount point) | SG3-Backup | DG4 | EVA-2 | 350 | 5 | 47 |
| | W:\Dump4 (Mount point) | SG4-Backup | DG4 | EVA-2 | 350 | 5 | 48 |
| | M:\Log (Mount point) | SG1-Log | DG3 | EVA-2 | 40 | 1 | 24 |
| | N:\Log (Mount point) | SG2-Log | DG3 | EVA-2 | 40 | 1 | 25 |
| | O:\Log (Mount point) | SG3-Log | DG3 | EVA-2 | 40 | 1 | 26 |
| | P:\Log (Mount point) | SG4-Log | DG3 | EVA-2 | 40 | 1 | 28 |
| EVS4 | Q:\Exchsrvr (Mount point) | SMTP | DG2 | EVA-2 | 50 | 1 | 19 |
| | Q$ | SG1-Data | DG2 | EVA-2 | 350 | 1 | 29 |
| | R$ | SG2-Data | DG2 | EVA-2 | 350 | 1 | 30 |
| | S$ | SG3-Data | DG2 | EVA-2 | 350 | 1 | 31 |
| | T$ | SG4-Data | DG2 | EVA-2 | 350 | 5 | 49 |
| | X$ | SG1-Backup | DG5 | EVA-2 | 350 | 5 | 50 |
| | X$:\Dump2 (Mount point) | SG2-Backup | DG5 | EVA-2 | 350 | 5 | 51 |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| | X$:\Dump3 (Mount point) | SG3-Backup | DG5 | EVA-2 | 350 | 5 | 52 |
| | X$:\Dump4 (Mount point) | SG4-Backup | DG5 | EVA-2 | 350 | 1 | 32 |
| | Q:\Log (Mount point) | SG1-Log | DG3 | EVA-2 | 40 | 1 | 33 |
| | R:\Log (Mount point) | SG2-Log | DG3 | EVA-2 | 40 | 1 | 34 |
| | S:\Log (Mount point) | SG3-Log | DG3 | EVA-2 | 40 | 1 | 35 |
| | T:\Log (Mount point) | SG4-Log | DG3 | EVA-2 | 40 | 1 | 36 |
| CVS | Y$ Quorum | Quorum | DG3 | EVA-1 | 10 | 1 | 53 |

# Microsoft Windows Server 2003 TCP/IP Implementation Details

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: June 1, 2003 | Updated: June 13, 2006

Microsoft Corporation

This article describes the implementation of the TCP/IP protocol stack in the Microsoft Windows Server 2003 family and is a supplement to the Windows Server 2003 Help and Support Center and Technical Reference documentation. This article contains an overview of TCP/IP in Windows Server 2003 features and capabilities, a discussion of protocol architecture, and detailed discussions of the core components, network application interfaces, and critical client components and services. The intended audience for this article is network engineers and support professionals who are already familiar with TCP/IP. Except where noted, the TCP/IP implementation for Windows XP is the same as that for Windows Server 2003.

**In This White Paper**

Introduction (Microsoft Windows Server 2003 TCP/IP Implementation Details)

Capabilities and Functionality

Architectural Model

The NDIS Interface and Below

Core Protocol Stack Components and the TDI Interface

Network Application Interfaces

Critical Client Services and Stack Components

Appendix A: TCP/IP Configuration Parameters

Appendix B: NetBIOS over TCP/IP Configuration Parameters

Appendix C: Windows Sockets and DNS Registry Parameters

Appendix D: Tuning TCP/IP Response to Attack

Appendix E: Format of the Daytime Service Response String

Summary - Microsoft Windows Server 2003 TCP/IP Implementation Details

# Introduction (Microsoft Windows Server 2003 TCP/IP Implementation Details)

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft has adopted TCP/IP as the strategic enterprise network transport for its platforms. In the early 1990s, Microsoft started an ambitious project to create a TCP/IP stack and services that would greatly improve the scalability of Microsoft networking. With the release of the Microsoft Windows NT 3.5 operating system, Microsoft introduced a completely rewritten TCP/IP stack. This new stack was designed to incorporate many of the advances in performance and ease of administration that were developed over the past decade. The stack was a high-performance implementation of the industry-standard TCP/IP protocol. It has evolved with each version of Windows based on the Windows NT code base to include new features and services that enhance performance, security, and reliability.

The goals in designing the TCP/IP stack were to make it:

- Standards-compliant and interoperable
- Portable
- Scalable and fast
- Versatile
- Self-tuning and easy to administer

In this article, the Windows Server 2003 TCP/IP protocol suite is examined from the bottom up. Throughout the article, network traces are used to illustrate key concepts. These traces were gathered and formatted using Microsoft Network Monitor 2.0, a software-based protocol tracing and analysis tool included in the Microsoft Systems Management Server product. Windows 2000 Server and Windows Server 2003 include a limited functionality version of Network Monitor. The primary difference between this version and the Systems Management Server version is that the limited version can only capture frames that would normally be seen by the computer that it is installed on, rather than all frames that pass over the network (which requires the network interface card to be in promiscuous mode). It also does not support connecting to remote Network Monitor Agents.

# Capabilities and Functionality

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Overview**
Windows Server 2003 TCP/IP was designed to make it easy to integrate Microsoft systems into large-scale corporate, government, and public networks, and to provide the ability to operate over those networks in a secure manner. The Windows Server 2003 TCP/IP protocol is installed by default and, unlike previous versions of Windows, cannot be uninstalled. However, you can reset the TCP/IP configuration to a default state with the **netsh interface ip reset** command.

**Support for Standard Features**
Windows Server 2003 TCP/IP supports the following standard features:

- Ability to bind to multiple network adapters with different media types
- Logical and physical multihoming
- Internal IP routing capability
- Internet Group Management Protocol (IGMP) version 3 (IP multicasting)
- Duplicate IP address detection
- Multiple default gateways
- Dead gateway detection
- Automatic Path Maximum Transmission Unit (PMTU) discovery
- Internet Protocol security (IPsec)
- Quality of Service (QoS)
- ATM Services
- Virtual Private Networks (VPNs) with the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol with IPsec (L2TP/IPsec)

**Performance Enhancements**
In addition, Windows Server 2003 TCP/IP has the following performance enhancements:

- Protocol stack tuning, including increased default window sizes and new algorithms for high-delay and high-loss links, which increases throughput
- TCP-scalable window sizes (described in RFC 1323)
- Selective acknowledgments (SACK) (described in RFC 2018)
- TCP fast retransmit and fast recovery (described in RFC 2581)
- Round Trip Time (RTT) and Retransmission Timeout (RTO) calculation improvements
- Improved performance for management of large numbers of connections
- Hardware task offload mechanisms including checksum offload and large send offload (LSO)

**Services Available**

The Windows Server 2003 operating system provides the following TCP/IP-related services:

- Dynamic Host Configuration Protocol (DHCP) client and server and DHCP Relay Agent (with the Routing and Remote Access service)
- In the absence of a DHCP server, Automatic Private IP Addressing (APIPA) is used
- Windows Internet Name Service (WINS), a NetBIOS name client and server
- Domain Name System (DNS) client and server, including support for DNS dynamic updates
- Dial-up support using the Point-to-Point Protocol (client and server) and Serial Line Internet Protocol (client only)
- PPTP and L2TP/IPsec, used for remote access and site-to-site VPN connections
- TCP/IP network printing (client only with the Lpr.exe and Lpq.exe tools)
- SNMP agent
- NetBIOS interface
- Network Location Service
- Windows Sockets version 2 (Winsock2) interface
- Remote Procedure Call (RPC) support
- Network Dynamic Data Exchange (NetDDE)
- Computer browsing (My Network Places) across IP routers
- Reliable multicast with the Pragmatic General Multicast (PGM) protocol
- Basic TCP/IP connectivity utilities, including: finger, ftp, rcp, rexec, rsh, telnet, and tftp
- Server and client software for simple network protocols, including: Character Generator, Daytime, Discard, Echo, and Quote of the Day
- Routing Information Protocol (RIP) listener (for Windows XP Professional) and RIP and Open Shortest Path First (OSPF) (with the Routing and Remote Access service)
- Network Address Translator (NAT) capabilities using either the Internet Connection Sharing (ICS) or the NAT/Basic Firewall routing protocol component of the Routing and Remote Access service
- Stateful firewalling capabilities using either the Internet Connection Firewall (for Windows Server 2003 with no service packs installed), Windows Firewall (for Windows Server 2003 Service Pack 1), or the NAT/Basic Firewall routing protocol component of the Routing and Remote Access service
- Multicast forwarding and IGMP router and proxy capabilities with the Routing and Remote Access service
- TCP/IP management and diagnostic tools, including: arp, ipconfig, nbtstat, netsh, netstat, ping, pathping, route, nslookup, and tracert

**New Features for Windows Server 2003 TCP/IP**

The features and improvements of TCP/IP that are new for Windows Server 2003 include the following:

- Windows Server 2003, Windows XP with Service Pack 1, and Windows XP with Service Pack 2 now include a production-quality IPv6 protocol stack. For more information about IPv6, see Windows Server 2003 Help and Support Center or the Microsoft Windows IPv6 Web site (http://go.microsoft.com/fwlink/?LinkID=17074).
- Auto-negotiation of RFC 1323 options (window scaling and TCP timestamps).
- Default support of network interface cards providing large send offload (LSO) and checksum offload.
- IGMP version 3.
- Reliable multicast with PGM.
- Alternate configuration.
- Automatic determination of the interface-related and default route metrics.

The table below lists features and the operating system versions that they are present in as a reference. Features are described in more detail throughout this article.

**Table 1. Feature Comparison Table for Windows TCP/IP Versions**

| Product | Windows 98 | Windows 98 SE | Windows NT 4.0 SP5 | Windows 2000 | Windows Server 2003 |
|---|---|---|---|---|---|
| **Dead gateway detection** | Y | Y | Y | Y | Y |
| **Fast retransmit/recovery** | Y | Y | Y | Y | Y |
| **APIPA** | Y | Y | N | Y | Y |
| **Selective ACK (SACK)** | Y | Y | N | Y | Y |
| **Jumbo frame support** | Y | Y | Y | Y | Y |
| **Large windows** | D | D | N | D | D |
| **DNS dynamic update** | N | N | N | Y | Y |
| **Media sense** | N | N | N | Y | Y |
| **Wake on LAN** | N | N | N | Y | Y |
| **IP forwarding** | N | D | D | D | D |
| **NAT** | N | D | N | D | D |
| **Kerberos v5** | N | N | N | Y | Y |
| **IPsec** | N | N | N | Y | Y |
| **PPTP** | Y | Y | Y | Y | Y |
| **L2TP/IPsec** | N | N | N | Y | Y |
| **IP Helper API** | Y | Y | Y | Y | Y |
| **Winsock2 API** | Y | Y | Y | Y | Y |
| **GQoS API** | Y | Y | N | Y | Y |
| **IP Filtering API** | N | N | N | Y | Y |

| | | | | | |
|---|---|---|---|---|---|
| **Firewall hook** | N | N | N | Y | Y |
| **Packet scheduler** | N | N | N | D | D |
| **Network location** | N | N | N | N | Y |
| **ISSLOW** | Y | Y | N | Y | Y |
| **Personal firewall** | N | N | N | N | D |
| **Block source routing** | N | Y | Y | Y | Y |
| **ICMP Router Discovery** | Y | Y | D | D | D |
| **IPsec offload** | N | N | N | Y | Y |
| **IGMP v3** | N | N | N | N | Y |
| **Reliable multicast (PGM)** | N | N | N | N | Y |
| **Alternate configuration** | N | N | N | N | Y |
| **Auto-determination of routing metrics** | N | N | N | N | Y |
| **Checksum offload** | N | N | N | N | Y |
| **Large send offload** | N | N | N | N | Y |

N=No, Y=Yes, and D=Disabled by Default

**Internet RFCs Supported by Windows Server 2003 TCP/IP**

Requests for Comments (RFCs) are a constantly evolving series of reports, proposals for protocols, and protocol standards used by the Internet community. You can obtain RFCs from the Engineering Task Force Web site (http://go.microsoft.com/fwlink/?LinkID=29138).

**Table 2. RFCs supported by Windows Server 2003 TCP/IP**

| RFC | Title |
|---|---|
| **768** | User Datagram Protocol (UDP) |
| **783** | Trivial File Transfer Protocol (TFTP) |
| **791** | Internet Protocol (IP) |
| **792** | Internet Control Message Protocol (ICMP) |
| **793** | Transmission Control Protocol (TCP) |
| **816** | Fault Isolation and Recovery |
| **826** | Address Resolution Protocol (ARP) |
| **854** | Telnet Protocol (TELNET) |
| **862** | Echo Protocol (ECHO) |
| **863** | Discard Protocol (DISCARD) |
| **864** | Character Generator Protocol (CHARGEN) |
| **865** | Quote of the Day Protocol (QUOTE) |
| **867** | Daytime Protocol (DAYTIME) |
| **894** | IP over Ethernet |

| | |
|---|---|
| **919, 922** | IP Broadcast Datagrams (broadcasting with subnets) |
| **950** | Internet Standard Subnetting Procedure |
| **959** | File Transfer Protocol (FTP) |
| **1001, 1002** | NetBIOS Service Protocols |
| **1065, 1035, 1123, 1886** | Domain Name System (DNS) |
| **1042** | A Standard for the Transmission of IP Datagrams over IEEE 802 Networks |
| **1055** | Transmission of IP over Serial Lines (IP-SLIP) |
| **1112** | Internet Group Management Protocol (IGMP) |
| **1122, 1123** | Host Requirements (communications and applications) |
| **1144** | Compressing TCP/IP Headers for Low-Speed Serial Links |
| **1157** | Simple Network Management Protocol (SNMP) |
| **1179** | Line Printer Daemon Protocol |
| **1188** | IP over FDDI |
| **1191** | Path MTU Discovery |
| **1201** | IP over ARCNET |
| **1256** | ICMP Router Discovery Messages |
| **1323** | TCP Extensions for High Performance |
| **1332** | PPP Internet Protocol Control Protocol (IPCP) |
| **1518** | Architecture for IP Address Allocation with CIDR |
| **1519** | Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy |
| **1534** | Interoperation Between DHCP and BOOTP |
| **1542** | Clarifications and Extensions for the Bootstrap Protocol |
| **1552** | PPP Internetwork Packet Exchange Control Protocol (IPXCP) |
| **1661** | The Point-to-Point Protocol (PPP) |
| **1662** | PPP in HDLC-like Framing |
| **1748** | IEEE 802.5 MIB using SMIv2 |
| **1749** | IEEE 802.5 Station Source Routing MIB using SMIv2 |
| **1812** | Requirements for IP Version 4 Routers |
| **1828** | IP Authentication using Keyed MD5 |
| **1829** | ESP DES-CBC Transform |
| **1851** | ESP Triple DES-CBC Transform |
| **1852** | IP Authentication using Keyed SHA |
| **1994** | PPP Challenge Handshake Authentication Protocol (CHAP) |

| | |
|---|---|
| **1995** | Incremental Zone Transfer in DNS |
| **1996** | A Mechanism for Prompt DNS Notification of Zone Changes |
| **2018** | TCP Selective Acknowledgment Options |
| **2085** | HMAC-MD5 IP Authentication with Replay Prevention |
| **2104** | HMAC: Keyed Hashing for Message Authentication |
| **2131** | Dynamic Host Configuration Protocol |
| **2136** | Dynamic Updates in the Domain Name System (DNS UPDATE) |
| **2181** | Clarifications to the DNS Specification |
| **2236** | Internet Group Management Protocol, Version 2 |
| **2308** | Negative Caching of DNS Queries (DNS NCACHE) |
| **2401** | Security Architecture for the Internet Protocol |
| **2402** | IP Authentication Header |
| **2406** | IP Encapsulating Security Payload (ESP) |
| **2581** | TCP Congestion Control |
| **3208** | PGM Reliable Transport Protocol Specification |
| **3376** | Internet Group Management Protocol, Version 3 |

**New Features for TCP/IP in Windows Server 2003 Service Pack 1**
The features and improvements of TCP/IP that are new for Windows Server 2003 Service Pack 1 include the following:

- Windows Firewall
- The Netstat –b option
- Netsh commands for Windows Sockets
- SYN attack protection is enabled by default
- SYN attack notification IP Helper APIs
- Registry parameter for ICMP host routes
- Smart TCP port allocation
- Registry value for multiple ARP replies

**Windows Firewall**
Windows Firewall replaces the Internet Connection Firewall provided with Windows Server 2003 with no service packs installed. Windows Firewall is a stateful firewall that drops unsolicited incoming traffic that does not correspond to either traffic sent in response to a request of the computer (solicited traffic) or unsolicited traffic that has been specified as allowed (excepted traffic). Windows Firewall provides a level of protection from malicious users and programs that rely on unsolicited incoming traffic to attack computers.

For more information about Windows Firewall in Windows Server 2003 Service Pack 1, see the Microsoft Windows Server 2003 Windows Firewall TechCenter (http://go.microsoft.com/fwlink/?linkid=67902).

**The Netstat –b option**

The Netstat tool displays a variety of information about active TCP connections, ports on which the computer is listening, Ethernet statistics, the IP route table, and IPv4 and IPv6 statistics. In Windows Server 2003 Service Pack 1, the Netstat tool supports a new **–b** option that displays the set of components that are listening on each open TCP and UDP port.

With Windows Server 2003 with no service packs installed, you can use the **–o** option to display the set of ports being listened on and the corresponding process ID (PID). You can then lookup the PID in the display of the **tasklist /svc** command to discover the name of the process that owns the port. However, in some cases, there are multiple services within a single process and it is not possible to determine which service within the process owned the port.

With the **–b** option, Netstat displays the TCP or UDP port, the file names corresponding to the components of the service that owns the port, and the PID. From the file names and PID, you can determine which of the services in the display of the **tasklist /svc** command owns the port.

**Netsh Commands for Windows Sockets**

There are now Windows Sockets (Winsock) Netsh commands to view the set of installed Winsock Layered Service Providers (LSPs) (the **netsh winsock show catalog** command) and to reset the Winsock LSP catalog to a default configuration (the **netsh winsock reset catalog** command). The **netsh winsock reset catalog** command is useful for restoring the Winsock LSP catalog when it has been corrupted by programs or services that install LSPs. However, you must reinstall the programs or services that use LSPs.

**SYN Attack Protection is Enabled by Default**

A TCP Synchronize (SYN) attack is a denial-of-service attack that exploits the retransmission and time-out behavior of the Synchronize-Acknowledgement (SYN-ACK) segment during the TCP three-way handshake to create a large number of half-open TCP connections. Depending on the TCP/IP protocol implementation, a large number of half-open TCP connections could do any of the following:

- Use all available memory.
- Use all possible entries in the TCP Transmission Control Block (TCB), an internal table used to track TCP connections. Once the half-open connections use all the entries, further connection attempts are responded to with a TCP connection reset.
- Use all available half-open connections. Once all the half-open connections are used, further connection attempts are responded to with a TCP connection reset.

To create a large number of TCP half-open connections, attackers send a large number of SYN segments, each from a spoofed IP address and TCP port number. Each spoofed IP address and TCP port number are for a process that does not respond to the SYN-ACKs being sent by the attacked host. SYN attacks are typically used to render Internet servers inoperative.

To mitigate the impact on a host experiencing a SYN attack, TCP/IP minimizes the amount of resources devoted to incomplete TCP connections and reduces the amount of time before abandoning the half-open connection. When a SYN attack is detected, TCP/IP in Windows Server 2003 and Windows XP lowers the number of retransmissions of the SYN-ACK segment and does not allocate memory or table entry resources for the connection until the TCP three-way handshake has been completed.

You can control SYN attack protection through the SynAttackProtect registry entry at HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters (type REG_DWORD). You set SynAttackProtect to 0 to disable SYN attack protection and to 1 to enable it.

For TCP/IP in Windows XP (all versions) and Windows Server 2003 with no service packs installed, SynAttackProtect is set to 0 by default. For TCP/IP in Windows Server 2003 Service Pack 1, SynAttackProtect is set to 1 by default.

**SYN Attack Notification IP Helper APIs**

To allow an application to notify network administrators that a SYN attack is taking place, the IP Helper API supports new SYN attack notification APIs named NotifySecurityHealthChange and CancelSecurityHealthChangeNotify. For information about these new APIs, see the Microsoft Developer Network (MSDN) (http://go.microsoft.com/fwlink/?linkid=67904).

**Registry Parameter for ICMP Host Routes**

TCP/IP for Windows Server 2003 SP1 supports a new registry parameter that restricts the number of host routes that can be added to the local IP route table by receiving ICMP Redirect messages. The new registry parameter is HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\ MaxICMPHostRoutes (REG_DWORD type). MaxICMPHostRoutes has a default value of 1000. You should not change this value unless the computer needs to be able to add a large number of host routes by receiving ICMP Redirect messages. The update to Windows Server 2003 SP1 available from article 898060 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=69231) changes the default

value of MaxICMPHostRoutes to 10000.

**Smart TCP Port Allocation**
When a TCP peer initiates a TCP connection termination and the connection termination completes, the TCP connection enters the TIME WAIT state. Once the TIME WAIT state is reached, TCP must wait twice the maximum segment lifetime (MSL) before a connection with the same set of socket addresses can be created. The set of socket addresses consist of the combination of the source and destination IP addresses and source and destination TCP ports. The MSL is the maximum amount of time a TCP segment can exist in an internetwork, and its recommended value is 120 seconds. This delay prevents a new connection's TCP segments that are using the same set of socket addresses from being confused with duplicated TCP segments of the old connection.

The TCP port for a connection in the TIME WAIT state is considered an available port and can be assigned for use by an application. This can lead to the following situation:

- An application requests any available TCP port.
- TCP/IP assigns a TCP port to use for the application socket.
- The application attempts to open a socket with a specific destination IP address.
- The application establishes a TCP connection and sends data.
- The application terminates the TCP connection.
- TCP/IP places the application's TCP connection in the TIME WAIT state until twice the MSL has passed.
- The same application requests another available TCP port.
- TCP/IP assigns a TCP port to use for the application socket. Because the port for the connection in the TIME WAIT state is considered open, it can be chosen as the next port to assign to the requesting application.
- Assuming that TCP/IP assigns the same TCP port number, the application attempts to open a socket with the same destination IP address.
- Because the connection is using the same set of socket addresses as the connection in the TIME WAIT state, TCP/IP indicates an error to the application.

You can mitigate this situation by setting the TcpTimedWaitDelay registry entry at HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters (REG_DWORD type) to a lower value. The value of TcpTimedWaitDelay determines the length of time that a connection stays in the TIME-WAIT state. However, lowering the value of TcpTimedWaitDelay is contrary to the original design of TCP and the MSL.

To prevent an application from creating a connection with the same set of socket addresses of a connection that is in a TIME WAIT state, TCP/IP in Windows Server 2003 Service Pack 1 has implemented a smart TCP port allocation algorithm. When an application requests any available TCP port, TCP/IP first attempts to find an available port that does not correspond to a connection in the TIME WAIT state. If a port cannot be found, then it picks any available port.

This new behavior makes it much more unlikely that an application will be assigned a TCP port that is in the TIME-WAIT state when connecting to the same destination. You no longer need to modify the TcpTimedWaitDelay registry entry.

**Registry Value for Multiple ARP Replies**
TCP/IP for Windows Server 2003 SP1 supports a new registry value that determines which MAC address is stored in the ARP cache when multiple ARP Reply messages are received. If there are multiple computers that are using the same IP address on a subnet, when a node sends an ARP Request frame for the IP address, it will receive multiple ARP replies. The new registry value HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\StrictARPUpdate (REG_DWORD type) allows you to specify whether TCP/IP in Windows Server 2003 SP1 will store the MAC address of the last ARP reply received (StrictARPUpdate=0 [default]) or the MAC address of the first ARP reply received (StrictARPUpdate=1).

# Architectural Model

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Overview**

The Windows TCP/IP suite contains *core protocol elements*, *services*, and the *interfaces* between them. The Transport Driver Interface (TDI) and the Network Device Interface Specification (NDIS) are public, and their specifications are available from Microsoft.[1] In addition, there are a number of higher-level interfaces available to user-mode applications. The most commonly used are Windows Sockets, remote procedure call (RPC), and NetBIOS.




**Figure 1. The Windows Server 2003 TCP/IP architectural model**

| **Note** |
| --- |
| Figure 1 does not show the IPsec components. For more information, see How IPSec Works (http://go.microsoft.com/fwlink/?linkid=67907). |

**Sending and Receiving IP Packets**

When the source sends an IP packet, the following components analyze or change the packet in the following order:

1. Routing and Remote Access service IP packet filters
2. IPsec

When a router running Windows Server 2003 forwards an IP packet, the following components analyze or change the packet in the following order:

1. Internet Connection Sharing of the Network Connections folder or the NAT/Basic Firewall component of the Routing and Remote Access service
2. Routing and Remote Access service IP packet filters

3. IPsec

When an IP packet is received for forwarding, the following components analyze or change the packet in the following order:

1. IPsec
2. Routing and Remote Access service IP packet filters
3. Internet Connection Sharing of the Network Connections folder or the NAT/Basic Firewall component of the Routing and Remote Access service

When an IP packet is received that is destined for the local host, the following components analyze or change the packet in the following order:

1. IPsec
2. Routing and Remote Access service IP packet filters
3. Internet Connection Firewall (for Windows Server 2003 with no service packs installed), Windows Firewall (for Windows Server 2003 with Service Pack 1), or the NAT/Basic Firewall component of the Routing and Remote Access service
4. TCP/IP filters

This discussion only includes components that are provided with Windows Server 2003. This does not include any Windows Sockets layered service providers or NDIS intermediate miniport drivers.

**Plug and Play**
Windows Server 2003 includes support for Plug and Play. Plug and Play has the following capabilities and features:

- Automatic and dynamic recognition of installed hardware. This includes initial system installation, recognition of static hardware changes that may occur between boots, and response to run-time hardware events, such as dock or undock, and insertion or removal of cards.
- Streamlined hardware configuration in response to automatic and dynamic recognition of hardware, including dynamic hardware activation, resource arbitration, device driver loading, drive mounting, and so on.
- Support for particular buses and other hardware standards that facilitate automatic and dynamic recognition of hardware and streamlined hardware configuration, including Plug and Play ISA, PCI, PCMCIA, PC Card/CardBus, USB, and IEEE 1394. This includes promulgation of standards and advice about how hardware should behave.
- An orderly Plug and Play framework in which driver writers can operate. This includes infrastructure, such as device information (INF) interfaces, APIs, kernel-mode notifications, executive interfaces, and so on.
- Mechanisms that allow user-mode code and applications to learn of changes in the hardware environment so that they can take appropriate actions.

Plug and Play operation does not require Plug and Play hardware. To the degree possible, the first two bullets above apply to legacy hardware, as well as Plug and Play hardware. In some cases, orderly enumeration of legacy devices is not possible because the detection methods are destructive or inordinately time-consuming.

The primary impact that Plug and Play support has on protocol stacks is that network interfaces can come and go at any time. Windows Server 2003 TCP/IP and related components have been adapted to support Plug and Play.

# The NDIS Interface and Below

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft networking protocols use the Network Device Interface Specification (NDIS) to communicate with network card drivers. Much of the OSI model link layer functionality is implemented in the protocol stack. This makes development of network card drivers much simpler.

**Network Driver Interface Specification (3.1 through 5.1)**
NDIS 3.1 supports basic services that allow a protocol module to send raw packets over a network device and allow that same module to be notified of incoming packets received by a network device.

NDIS 4.0 added the following new features to NDIS 3.1:

- Out-of-band data support.
- WirelessWAN Media Extension.
- High-speed packet send and receive (a significant performance win).
- Fast IrDA Media Extension.
- Media Sense (required for the Designed for Windows logo in PC 97 and later Hardware Design Guide). The Windows Server 2003 TCP/IP stack utilizes media sense information, which is described in the "Automatic Client Configuration" section of this article.
- All local packet filter (prevents Network Monitor from monopolizing the CPU).
- Numerous new NDIS system functions (required for miniport binary compatibility across Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003).

NDIS 5.0 includes all functionality defined in NDIS 4.0, plus the following extensions:

- NDIS power options (required for Network Power Options and Network Wake-up).
- Plug and Play.
- Support for Windows Management Instrumentation (WMI), which provides Web-based Enterprise Management (WBEM)–compatible instrumentation of NDIS miniports and their associated adapters.
- Support for a single INF format across Windows operating systems. The INF format is based on the Windows 98 INF format.
- Deserialized miniport for improved performance.
- Task offload mechanisms, such as TCP and UDP checksum and Fast Packet Forwarding.
- TCP segmentation offload, also known as large-send offload (LSO).
- Broadcast Media Extension (needed for Broadcast Services for Windows).
- Connection-oriented NDIS (required to support Asynchronous Transfer Mode [ATM], Asymmetric Digital Subscriber Line [ADSL], and Windows Driver Model–Connection Streaming Architecture [WDM-CSA]).
- Support for Quality of Service (QoS).

- Intermediate Driver Support (required for Broadcast PC, Virtual LANs, Packet Scheduling for QoS, and NDIS support of IEEE 1394 network devices).

| **Note** |
| --- |
| You can disable a network interface card's use of task offload capabilities on the **Advanced** tab of the properties of the network interface card driver in Control Panel-Network or Device Manager. |

NDIS 5.1 enhancements for Windows Server 2003 include:

- **Plug and Play and Power Event Notification** – Enables network interface card miniport drivers to be notified of power or Plug and Play events. This results in cleaner system operation during these events.

- **Support for Send Cancellation** – Allows network protocols to avoid having to wait lengthy amounts of time for network packet send requests to complete.

- **Increased Statistics Capacity (64-bit statistic counters)** – This enhancement enables accurate network statistic displays, even on today's high-speed network media.

- **Performance Enhancements** – Several enhancements were made to speed up critical network data paths and avoid unnecessary packet copies.

- **Wake on LAN Change** – A change was made to Wake on LAN to allow you to limit wake up packets to just magic packets (instead of protocol registered packet patterns). This is now configurable from the properties of a network interface card driver.

- **Miscellaneous Changes** – Several additional changes have been made to support common needs or requests from driver developers or to improve driver integrity.

Remote NDIS is also included as part of the Windows Server 2003 family. Remote NDIS enables the support of USB-attached network devices without the installation of third party drivers. Microsoft supplies the drivers required to communicate with the network devices. This results in easier installation and a lessened chance of system failure because of a poorly built or tested driver.

For more information about NDIS 5.1 and Remote NDIS, see the Windows DDK (http://go.microsoft.com/fwlink/?linkid=38448) and the following Web pages:

- http://go.microsoft.com/fwlink/?linkid=67945

- http://go.microsoft.com/fwlink/?linkid=67946

NDIS can power down network adapters when the system requests a power level change. Either the user or the system can initiate this request. For example, the user may want to put the computer in sleep mode, or the system may request a power level change based on keyboard or mouse inactivity. In addition, disconnecting the network cable can initiate a power-down request if the network interface card supports this functionality. In this case, the system waits a configurable time period before powering down the network interface card because the disconnect could be the result of temporary wiring changes on the network, rather than the disconnection of a cable from the network device itself.

NDIS power management policy is *no network activity*–based. This means that all overlying network components must agree to the request before the network interface card can be powered down. If there are any active sessions or open files over the network, the power-down request can be refused by any or all of the components involved.

The computer can also be awakened from a lower power state, based on network events. A wakeup signal can be caused by:

- Detection of a change in the network link state (for example, cable reconnect)

- Receipt of a network wakeup frame

- Receipt of a Magic Packet.

At driver initialization, NDIS queries the capabilities of the miniport to determine if it supports such things as Magic Packet, pattern match, or link change wakeups, and to determine the lowest required power state for each wakeup method. The network protocols then query the miniport capabilities. At run time, the protocol sets the wakeup policy, using object identifiers (OIDs), such as **Enable Wakeup**, **Set Packet Pattern**, and **Remove Packet Pattern**.

Currently, TCP/IP is the only Microsoft protocol stack that supports network power management. It registers the following packet patterns at miniport initialization:

- Directed IP packet
- ARP broadcast for the station’s IP address
- NetBIOS over TCP/IP broadcast for the station's assigned computer name

NDIS-compliant drivers are available for a wide variety of network adapters from many vendors. The NDIS interface allows multiple protocol drivers of different types to bind to a single network adapter driver and allows a single protocol to bind to multiple network adapter drivers. The NDIS specification describes the multiplexing mechanism used to accomplish this. Bindings can be viewed or changed as advanced settings from the Network Connections folder.

Windows Server 2003 TCP/IP provides support for the following media types:

- Ethernet (using Ethernet II or IEEE 802.3 Sub-Network Access Protocol [SNAP] encapsulation)
- Fiber Distributed Data Interchange (FDDI)
- Token Ring (IEEE 802.5)
- IEEE 802.11 Wireless LAN
- ATM (using LAN emulation [LANE] and Classical IP [CLIP] over ATM)
- Attached Resource Computing Network (ARCnet)
- Dedicated wide area network (WAN) links such as Dataphone Digital Service (DDS) and T-carrier (Fractional T1, T1, T3, E1, and E3)
- Dial-up or permanent circuit-switched WAN services such as analog phone, Integrated Services Digital Network (ISDN), and Digital Subscriber Line (xDSL)
- Packet-switched WAN services such as X.25, Frame Relay, and ATM

**Link Layer Functionality**

Link layer functionality is divided between the network interface card/driver combination and the low-level protocol stack driver. The network card/driver combination filters are based on the destination media access control (MAC) address of each frame.

Normally, the hardware filters out all incoming frames except those containing one of the following destination addresses:

- The address of the adapter
- The all ones broadcast address (0xFF-FF-FF-FF-FF-FF)
- Multicast addresses that a protocol driver on this host has registered interest in, using the **NDISRequest()** function

Because this first filtering decision is made by the hardware, the network interface card discards any frames that do not meet the filter criteria without incurring any CPU processing. All frames (including broadcasts) that pass the hardware filter and Frame Check Sequence (FCS) validation are then passed up to the network interface card driver through a hardware interrupt.[2] The network interface card driver is software that runs on the computer, so any frames that make it this far require some CPU time to process. The network interface card driver brings the frame into system memory from the interface card. Then the frame is indicated (passed up) to the appropriate bound transport driver(s). The NDIS 5.1 specification provides more detail on this process.

Frames are passed up to all bound transport drivers in the order that they are bound.

As a packet traverses a network or series of networks, the source MAC address is always that of the network interface card that placed it on the media, and the destination MAC address is that of the network interface card that is intended to pull it off the media. This means that, in a routed network, the source and destination MAC address changes with each hop through a network-layer device (router or Layer 3 switch).

**Maximum Transmission Unit (MTU)**

Each media type has a maximum frame size that cannot be exceeded. The link layer is responsible for discovering this MTU and reporting it to the protocols above. NDIS drivers may be queried for the local MTU by the protocol stack. Knowledge of the MTU for an interface is used by upper layer protocols, such as TCP, that optimize packet sizes for each media automatically. For details, see the discussion of TCP path Maximum Transmission Unit (PMTU) discovery in the "Transmission Control Protocol (TCP)" section of this article.

If a network interface card driver—such as an ATM driver—uses LAN emulation mode, it may report that it has an MTU that is higher than what is expected for that media type. For example, it may emulate Ethernet but report an MTU of 9180 bytes. Windows Server 2003 TCP/IP accepts and uses the MTU size reported by the adapter, even when it exceeds the normal MTU for a given media type.

Sometimes the MTU reported to the protocol stack may be less than what would be expected for a given media type. For instance, use of the 802.1p standard for QoS over Ethernet often reduces the MTU reported by 4 bytes due to larger link-layer headers (this is hardware dependent).

# Core Protocol Stack Components and the TDI Interface

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The core protocol stack components are those shown between the NDIS and TDI interfaces in Figure 1. They are implemented in the Windows Server 2003 Tcpip.sys driver. The TCP/IP stack is accessible through the TDI interface and the NDIS interface. The Winsock2 interface also provides some support for direct access to the protocol stack.

**Address Resolution Protocol (ARP)**
ARP performs IP address-to-MAC address resolution for outgoing packets. As each outgoing IP datagram is encapsulated in a frame, source and destination MAC addresses must be added. Determining the destination MAC address for each frame is the responsibility of ARP.

ARP compares the next-hop IP address on every outbound IP datagram to the ARP cache for the network interface card over which the frame will be sent. If there is a matching entry, the MAC address is retrieved from the cache. If not, ARP broadcasts an ARP Request frame on the local subnet, requesting that the owner of the IP address in question reply with its MAC address. If the packet is going through a router, the next-hop address is that of a neighboring router and ARP resolves the MAC address for that next-hop router, rather than the final destination host. When an ARP reply is received, the ARP cache is updated with the new information, and it is used to address the packet at the link layer.

**ARP Cache**
You can use the ARP utility to view, add, or delete entries in the ARP cache. Examples are shown below. Entries added manually are static and are not automatically removed from the cache, whereas dynamic entries are removed from the cache (see the "ARP Cache Aging" section for more information).

The **arp** command can be used to view the ARP cache, as shown here:

```
C:\>arp -a

Interface: 157.60.137.88 --- 0x10003
  Internet Address      Physical Address      Type
  157.60.136.1          00-0a-42-b0-54-0a     dynamic
  157.60.137.0          00-b0-d0-e9-41-43     dynamic

Interface: 10.0.0.3 --- 0x10004
  Internet Address      Physical Address      Type
  10.0.0.1              08-00-2b-c4-25-b6     dynamic
```

The computer in this example is *multihomed*—has more than one network interface card—so there is a separate ARP cache for each interface.

In the following example, the command **arp –s** is used to add a static entry to the ARP cache used by the second interface for the host whose IP address is 10.0.0.32 and whose network interface card address is 00-60-8C-0E-6C-6A:

```
C:\>arp -s 10.0.0.32 00-60-8c-0e-6c-6a 10.0.0.3

C:\>arp -a

Interface: 157.60.137.88 --- 0x10003
  Internet Address      Physical Address      Type
  157.60.136.1          00-0a-42-b0-54-0a     dynamic
  157.60.137.0          00-b0-d0-e9-41-43     dynamic

Interface: 10.0.0.3 --- 0x10004
  Internet Address      Physical Address      Type
  10.0.0.1              08-00-2b-c4-25-b6     dynamic
  10.0.0.32             00-60-8c-0e-6c-6a     static
```

**ARP Cache Aging**
Windows Server 2003 adjusts the size of the ARP cache automatically to meet the needs of the system. If an entry is not used by any outgoing datagram for two minutes, the entry is removed from the ARP cache. Entries that are being referenced are removed from the ARP cache after ten minutes. Entries added manually are not removed from the cache automatically. The registry parameter *ArpCacheLife*, described in Appendix A, allows more administrative control over aging.

Use the command **arp –d** to delete entries from the cache, as shown below:

```
C:\>arp -d 10.0.0.32

C:\>arp -a

Interface: 157.60.137.88 --- 0x10003
  Internet Address      Physical Address      Type
  157.60.136.1          00-0a-42-b0-54-0a     dynamic
  157.60.137.0          00-b0-d0-e9-41-43     dynamic

Interface: 10.0.0.3 --- 0x10004
  Internet Address      Physical Address      Type
  10.0.0.1              08-00-2b-c4-25-b6     dynamic
```

ARP queues only one outbound IP datagram for a specified destination address while that IP address is being resolved to a MAC address. If a UDP-based application sends multiple IP datagrams to a single destination address without any pauses between them, some of the datagrams may be dropped if there is no ARP cache entry already present. An application can compensate for this by calling the Iphlpapi.dll routine **SendArp()** to establish an ARP cache entry, before sending the stream of packets. See article 193059 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=67951) or MSDN (http://go.microsoft.com/fwlink/?linkid=67904) for IP Helper API details.

**Internet Protocol (IP)**

IP is the mailroom of the TCP/IP stack, where packet sorting and delivery take place. At this layer, each incoming or outgoing packet is referred to as a *datagram*. Each IP datagram bears the source IP address of the sender and the destination IP address of the intended recipient. Unlike MAC addresses, the IP addresses in a datagram remain the same throughout a packet's journey across an internetwork, unless you are using source routing. IP layer functions are described below.

**Routing**

Routing is a primary function of IP. Datagrams are handed to IP from UDP and TCP above, and from the network interface card(s) below. Each datagram is labeled with a source and destination IP address. IP examines the destination address on each datagram, compares it to a locally maintained route table, and decides what action to take. There are three possibilities for each datagram:

- It can be passed up to a protocol layer above IP on the local host.
- It can be forwarded using one of the locally attached network interface cards.
- It can be discarded.

The route table maintains four different types of routes:

1. Host route (a route to a single, specific destination IP address)
2. Subnet route (a route to a subnet)
3. Network route (a route to an entire network)
4. Default route (used when there is no other match)

To determine the single route to use to forward an IP datagram, IP uses the following process:

1. For each route in the route table, IP performs a bit-wise logical AND between the destination IP address and the netmask. IP compares the result with the network destination for a match. If they match, IP marks the route as one that matches the destination IP address.
2. From the list of matching routes, IP determines the route that has the most bits in the netmask. This is the route that matches the most bits to the destination IP address and is therefore the most specific route for the IP datagram. This is known as finding the longest or closest matching route.
3. If multiple closest matching routes are found, IP uses the route with the lowest metric. If multiple closest matching routes with the lowest metric are found, IP can choose to use any of those routes. For Windows Server 2003, IP uses the route corresponding to the adapter that is the highest in the binding order. You can view and modify the binding order from

the **Adapters and Bindings** tab in the **Advanced Settings** dialog box for the Network Connections folder.

You can use the **route print** command to view the route table from the command prompt, as shown here:

```
C:\>route print
IPv4 Route Table
===========================================================================
Interface List
0x1 ........................... MS TCP Loopback interface
0x10002 ...00 53 45 00 00 00 ...... WAN (PPP/SLIP) Interface
0x10003 ...00 04 5a 56 10 06 ...... Linksys LNE100TX Fast Ethernet
                                    Adapter(LNE100TX v4)
===========================================================================
===========================================================================
Active Routes:
Network Destination        Netmask          Gateway       Interface  Metric
0.0.0.0          0.0.0.0     157.60.136.1    157.60.137.88     20
157.60.136.0    255.255.252.0    157.60.137.88    157.60.137.88     20
157.60.137.88  255.255.255.255        127.0.0.1        127.0.0.1     20
157.60.255.255  255.255.255.255    157.60.137.88    157.60.137.88     20
127.0.0.0        255.0.0.0        127.0.0.1        127.0.0.1      1
224.0.0.0        240.0.0.0    157.60.137.88    157.60.137.88     20
255.255.255.255  255.255.255.255    157.60.137.88    157.60.137.88      1
Default Gateway:      157.60.136.1
===========================================================================
Persistent Routes:
None
```

If the IPv6 protocol is installed, the display of the **route print** command also lists IPv6 routes.

The route table above is for a computer with the IP address of 157.60.137.88, the subnet mask of 255.255.252.0, and the default gateway of 157.60.136.1. It contains the following entries:

- The first entry, to destination 0.0.0.0, is the default route.
- The second entry is for the subnet 157.60.136.0, on which this computer resides.
- The third entry, to destination 157.60.137.88, is a host route for the local host. It specifies the loopback address, which makes sense because a datagram bound for the local host should be looped back internally.
- The fourth entry is for the all-subnets-directed broadcast address corresponding to the original Class B network ID 157.60.0.0.
- The fifth entry is for the loopback network, 127.0.0.0.
- The sixth entry is for IP multicasting, which is discussed later in this article.
- The final entry is for the limited broadcast (all ones) address.

The *Default Gateway* is the currently active default gateway. This is useful to know when multiple default gateways are configured.

On this host, if a packet is sent to 157.60.138.49, the closest matching route is the local subnet route (157.60.136.0 with the mask of 255.255.252.0). The packet is sent via the local interface that is assigned the IP address 157.60.137.88. If a packet is sent to 10.200.1.1, the closest matching route is the default route. In this case, the packet is forwarded to the default gateway at 157.60.136.1.

The route table is maintained automatically in most cases. When a host initializes, entries for the local network(s), loopback, multicast, and configured default gateway are added. More routes may appear in the table as the IP layer learns of them. For instance, the default gateway for a host may advise it of a better route to a specific address using ICMP redirect, which is explained later in this article. Routes also may be added manually using the **route** command, or by a routing protocol. The **-p** (persistent) switch can be used with the route command to specify permanent routes. Persistent routes are stored in the registry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters \PersistentRoutes.

Windows Server 2003 TCP/IP introduces a new **Automatic metric** configuration option for interface-based and default gateway routing metrics. If selected for the interface, automatic metric configuration determines the metric for the routes associated with the interface configuration, such as subnet routes and host routes, based on the speed (bit rate) of the interface. The higher the speed, the lower the metric. For example, routes associated with 10 Mbps Ethernet interfaces have a metric of 30 and routes associated with 100 Mbps Ethernet interfaces have a metric of 20. If selected for the default gateway, automatic metric configuration determines the metric for the default route assigned to the interface, which is also based on the speed of the interface. Automatic metric configuration for both interface metrics and default routes are enabled by default and can be modified from the advanced configuration properties of the TCP/IP protocol for a connection in Network Connections. For more information, see article 299540 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=67955).

DHCP servers can also provide a base metric and a list of default gateways. If a DHCP server provides a base of 100, and a list of three default gateways, the gateways will be configured with metrics of 100, 101, and 102 respectively. A DHCP-provided base metric does not apply to statically configured default gateways.

By default, Windows-based systems do not behave as routers and do not forward IP datagrams between interfaces. However, the Routing and Remote Access service is included in Windows Server 2003 and can be enabled and configured to provide dynamic IP routing services using RIP and OSPF. Windows XP Professional includes support for silent RIP.

When running multiple logical subnets on the same physical network, the following command can be used to have IP treat all subnets as local (all destinations are on the local link):

**route add 0.0.0.0 MASK 0.0.0.0** <*my local ip address*>

Thus, packets destined for non-local subnets are transmitted directly onto the local media instead of being sent to a router. In essence, the local interface card can be designated as the default gateway. This can be useful where several class C networks are used on one physical network with no router to the outside world, or in a proxy-ARP environment.

**Duplicate IP Address Detection**

Duplicate address detection is an important feature. When the stack is first initialized or when a new IP address is added, gratuitous ARP requests are broadcast for the IP addresses of the local host. The number of ARP requests to send is controlled by the *ArpRetryCount* registry parameter described in Appendix A, which defaults to 3. If another host replies to any of these ARP requests, the IP address is already in use. When this happens for an interface with a single manually-configured address, the Windows-based computer still boots; however, the interface containing the offending address is disabled, a system log entry is generated, and an error message is displayed. If the host that is defending the address is also a Windows-based computer, a system log entry is generated, and an error message is displayed on that computer. In order to update the ARP caches on other computers, the offending computer re-broadcasts another ARP request, spoofing the MAC address of the defending computer, to restore the proper values in the ARP caches of the other computers.

A computer using a duplicate IP address can be started when it is not attached to the network, in which case no conflict would be detected. However, if it is then plugged into the network, the first time that it sends an ARP request for another IP address, any computer running a version of Windows with the Windows NT codebase with a conflicting address detects the conflict. The computer detecting the conflict displays an error message and logs a detailed event in the system log. A sample event log entry is shown below:

The system detected an address conflict for IP address 10.199.40.123 with the system having network hardware address 00:DD:01:0F:7A:B5. Network operations on this system may be disrupted as a result.

DHCP-enabled clients inform the DHCP server with a DHCPDecline message when an IP address conflict is detected and, instead of disabling the TCP/IP protocol, they request a new address from the DHCP server and request that the server mark the conflicting address as bad.

**Multihoming**

When a computer is configured with more than one IP address, it is referred to as a multihomed system. Multihoming is supported in three different ways:

- **Multiple IP addresses per network interface card**
    - To add addresses for an interface, obtain properties of the **Internet Protocol (TCP/IP)** protocol in Network Connection, and then click **Advanced**. In the **Advanced Settings** dialog box, click **Add** on the **IP Settings** tab to add IP addresses.
    - NetBIOS over TCP/IP (NetBT) binds to only one IP address per interface card. When a NetBIOS name registration is sent out, only one IP address is registered per interface. This registration occurs over the IP address that is

listed first on the **IP Settings** tab.

- **Multiple network interface cards per physical network.** There are no restrictions, other than hardware limitations on the number of network interface cards.

- **Multiple networks and media types.** There are no restrictions, other than hardware and media support. See the "The NDIS Interface and Below" section for supported media types.

When an IP datagram is sent from a multihomed host, it is passed to the interface with the best apparent route to the destination. Accordingly, the datagram may contain the source IP address of one interface in the multihomed host, yet be placed on the media by a different interface. The source MAC address on the frame is that of the interface that actually transmitted the frame to the media, and the source IP address is the one that the sending application sourced it from, not necessarily one of the IP addresses assigned to the sending interface.

When a computer is multihomed with network interface cards attached to disjoint networks (networks that are separate from and unaware of each other, such as a private network using private addressing and the Internet), routing problems may arise. It is often necessary to set up static routes to the private networks in this situation.

When configuring a computer to be multihomed on two disjoint networks, the best practice is to configure the default gateway on the interface connected to the largest and least-known network, in which the default route summarizes the most destinations. Then, either add static routes or use a routing protocol to provide connectivity to the hosts on the smaller or better-known network. Avoid configuring a different default gateway on each side; this can result in unpredictable behavior and loss of connectivity. For more information, see Default Gateway Behavior for Windows TCP/IP (http://go.microsoft.com/fwlink/?linkid=47705).

| **Note** |
| --- |
| There can only be one active default gateway for a computer at any moment in time. |

More details on name registration, resolution, and choice of network interface card on outbound datagrams with multihomed computers are provided in the "Transmission Control Protocol (TCP)," "NetBIOS over TCP/IP," and "Windows Sockets" sections of this article.

**Classless Inter-Domain Routing (CIDR)**

CIDR, described in RFCs 1518 and 1519, removes the concept of address classes from the IP address assignment and management process. In place of predefined, well-known boundaries, CIDR allocates addresses defined by a network prefix, which makes more efficient use of available space. The network prefix defines the portion of the address that is fixed. For example, an assignment from an ISP to a corporate client might be expressed as 157.60.1.128/25. In this prefix, the first 25 bits are fixed and the last 7 bits can be used for address assignment. This would result in a 128-address block for local use, with the upper 25 bits being the network identifier part of the address. A legacy, class-full prefix would be expressed as *w*.0.0.0/8, *w.x*.0.0/16, or *w.x.y*.0/24. As these are reclaimed, they will be reallocated using classless CIDR techniques.

Given the installed base of classful systems, the initial implementation of CIDR was to summarize portions of the Class C address space. This process was called *supernetting*. Supernetting can be used to consolidate several network IDs into one prefix. For example, the network IDs 131.107.4.0, 131.107.5.0, 131.107.6.0, and 131.107.7.0 can be summarized with the network ID 131.107.4.0 with a subnet mask of 255.255.252.0 (131.107.4.0/22). For example:

```
NET     131.107.4.0   (1100 0111.1100 0111.0000 0100.0000 0000)
NET     131.107.5.0   (1100 0111.1100 0111.0000 0101.0000 0000)
NET     131.107.6.0   (1100 0111.1100 0111.0000 0110.0000 0000)
NET     131.107.7.0   (1100 0111.1100 0111.0000 0111.0000 0000)
MASK  255.255.252.0   (1111 1111.1111 1111.1111 1100.0000 0000)
```

All four of the network IDs share the same high-order 22 bits (bolded).

When routing decisions are made, only the bits covered by the subnet mask are used, thus making all these addresses appear to be part of the same network for routing purposes. Any routers in use must also support CIDR and may require special configuration.

Windows Server 2003 TCP/IP includes support for the all-0's and all-1's subnets as described in RFC 1878.

**IP Multicasting**

IP multicasting is used to provide efficient multicast services to clients that may not be located on the same network segment. Windows Sockets applications can join a multicast group to participate in a wide-area conference, for instance.

Windows Server 2003 TCP/IP is level 2-compliant with RFC 1112 (send and receive). IGMP is the protocol used to track multicast membership on a subnet, which is described later in this article.

**IP over ATM**
Windows Server 2003 includes support for IP over ATM. RFC 1577 (and successors) define the basic operation of an IP over ATM network known as Classical IP over ATM (CLIP), which defines a Logical IP Subnet (LIS). A LIS is a set of IP hosts that can communicate directly with each other. Two hosts belonging to different LISs can communicate only through an IP router that is a member of both subnets. Windows Server 2003 also includes support for ATM LAN Emulation (LANE), which supports broadcasting.

**ATM Address Resolution**
Because an ATM network is non-broadcast, ARP broadcasts (as used by Ethernet or Token Ring) are not a suitable solution. Instead, a dedicated ARP server is used to provide IP-to-ATM address resolution.

One of the stations in a LIS is designated as an ARP server, and the ARP server software is loaded on it. Stations that use the services of the ARP server are referred to as *ARP clients*. All IP stations within a LIS are ARP clients. Each ARP client is configured with the ATM address of the ARP server. When an ARP client starts up, it makes an ATM connection to the ARP server, and sends a packet to the server that contains the client's IP and ATM addresses. The ARP server builds a table of IP-address-to-ATM-address mappings. When a client has an IP packet to be sent to another client (whose IP address is known but whose ATM address is unknown), it first queries the ARP server for the ATM address of the desired client. When it receives a reply that contains the desired ATM address, the client establishes a direct ATM connection to the target client and sends IP packets for that client on this connection.

The clients close any ATM connection, including the connection to the server, if the connections are inactive. All clients refresh their IP and ATM address information with the server periodically (the default is 15 minutes). The server purges an entry that is not refreshed after 20 minutes (by default). The ATM ARP client and ARP server both support a number of adjustable registry parameters, which are listed in Appendix A.

**Internet Control Message Protocol (ICMP)**
ICMP is a maintenance protocol specified in RFC 792 and is normally considered part of the Internet layer. ICMP messages are encapsulated within IP datagrams, so that they can be routed throughout an internetwork. Windows Server 2003 TCP/IP uses ICMP to:

- Report delivery problems encountered by routers or destination hosts.
- Build and maintain route tables.
- Perform router discovery.
- Assist in Path Maximum Transmission Unit (PMTU) discovery.
- Diagnose reachability problems (the ping, tracert, and pathping tools).
- Adjust flow control to prevent link or router saturation.

**ICMP Router Discovery**
Windows Server 2003 TCP/IP can perform router discovery as specified in RFC 1256. Router discovery provides an improved method of configuring and detecting default gateways. Instead of using manually- or DHCP-configured default gateways, hosts can dynamically discover routers on their subnet. If the primary router fails or the network administrators change router preferences, hosts can automatically switch to a backup router.

When a host that supports router discovery initializes, it joins the all-systems IP multicast group (224.0.0.1), and then listens for the router advertisements that routers send to that group. Hosts can also send router-solicitation messages to the all-routers IP multicast address (224.0.0.2) when an interface initializes to avoid any delay in being configured. Windows Server 2003 TCP/IP sends a maximum of three solicitations at intervals of approximately 600 milliseconds.

The use of router discovery is controlled by the *PerformRouterDiscovery* and *SolicitationAddressBCast* registry parameters, and it defaults to DHCP controlled in Windows Server 2003. Setting *SolicitationAddressBCast* to 1 causes router solicitations to be broadcast, instead of multicast, as described in RFC 1256.

**Maintaining Route Tables**
When a Windows-based computer is initialized, the route table normally contains only a few entries. One of those entries specifies a default gateway. Datagrams that have a destination IP address with no better match in the route table are sent to

the default gateway. However, because routers share information about network topology, the default gateway may know a better route to a given address. When this is the case, then upon receiving a datagram that could take the better path, the router forwards the datagram normally. It then advises the sender of the better route, using an *ICMP Redirect* message. These messages typically specify redirection for a specific destination address. When a Windows-based computer receives an ICMP redirect, a validity check is performed to be sure that it came from the first-hop gateway in the current route, and that the gateway is on a directly connected network. If so, a host route with a 10-minute lifetime is added to the route table for that destination IP address. If the ICMP redirect did not come from the first-hop gateway in the current route, or if that gateway is not on a directly connected network, the ICMP redirect is ignored.

In Windows Server 2003 Service Pack 1, the new MaxICMPHostRoutes registry value defines the maximum number of host routes that can be added through the receipt of ICMP Redirect messages. For more information, see Appendix A.

**Path Maximum Transmission Unit (PMTU) Discovery**
TCP employs Path Maximum Transmission Unit (PMTU) discovery, as described in the "Transmission Control Protocol (TCP)" section of this article. The mechanism relies on ICMP Destination Unreachable messages.

**Use of ICMP to Diagnose Problems**
The ping command-line utility is used to send ICMP echo requests to an IP address and wait for ICMP echo responses. Ping reports on the number of responses received and the time interval between sending the request and receiving the response. There are many different options that can be used with the ping utility.

Tracert is a route-tracing utility that can be very useful. Tracert works by sending ICMP echo requests to an IP address, while incrementing the Time to Live (TTL) field in the IP header, starting at 1, and analyzing the ICMP errors that are returned. Each succeeding echo request should get one hop further into the network before the TTL field reaches 0 and the router attempting to forward it returns an ICMP Time Exceeded-TTL Exceeded in Transit error message. Tracert prints out an ordered list of the routers in the path that returned these error messages, including the name and the IP address the nearside interface of each router. If the **-d** (do not do a DNS reverse query on each IP address) switch is used, only the IP address is reported. The example below illustrates using tracert to find the route from a computer dialed in over Point-to-Point Protocol (PPP) to an Internet service provider in Seattle to www.whitehouse.gov.

```
C:\>tracert www.whitehouse.gov
Tracing route to www.whitehouse.gov [128.102.252.1]
over a maximum of 30 hops:
1  300 ms  281 ms  280 ms roto.seanet.com [199.181.164.100]
2  300 ms  301 ms  310 ms sl-stk-1-S12-T1.sprintlink.net [144.228.192.65]
3  300 ms  311 ms  320 ms sl-stk-5-F0/0.sprintlink.net [144.228.40.5]
4  380 ms  311 ms  340 ms icm-fix-w-H2/0-T3.icp.net [144.228.10.22]
5  310 ms  301 ms  320 ms arc-nas-gw.arc.nasa.gov [192.203.230.3]
6  300 ms  321 ms  320 ms n254-ed-cisco7010.arc.nasa.gov [128.102.64.254]
7  360 ms  361 ms  371 ms www.whitehouse.gov [128.102.252.1]
```

Pathping is a command-line utility that combines the functionality of ping and tracert as well as introducing some new features. Along with the tracing functionality of tracert, pathping will ping each hop along the route multiple times and display delay and packet loss information per hop, which can help you determine if there is a high-loss link or router in the path.

**Internet Protocol Security (IPsec)**
Windows Server 2003 includes support for Internet Protocol security (IPsec). IPsec features and implementation details are very complex and are described in detail in a series of RFCs, Internet drafts, and in other Microsoft technical papers. For more information, see the Windows Server 2003 IPsec Web page (http://go.microsoft.com/fwlink/?linkid=63067).

IPsec uses cryptography-based security to provide data integrity, data origin authentication, replay protection, data confidentiality (encryption), and limited traffic-flow confidentiality. Because IPsec is provided at the IP layer, its services are available to the upper-layer protocols in the stack and, transparently, to existing applications.

IPsec enables a system to select security protocols, decide which algorithm(s) to use for the service(s), and establish and maintain cryptographic keys for each security relationship. IPsec can protect paths between hosts, between security gateways, or between hosts and security gateways. The services available and required for traffic are configured using IPsec policy. IPsec policy may be configured locally on a computer or can be assigned through Group Policy mechanisms using the Active Directory directory service. When using Active Directory, hosts detect policy assignment at startup, retrieve the policy, and then periodically check for policy updates. The IPsec policy specifies how computers trust each other. IPsec can use certificates, Kerberos, or preshared keys as authentication methods. The easiest trust to use is the Windows Server 2003 domain trust based on Kerberos.

Each IP datagram processed at the IP layer is compared to a set of filters that are provided by the security policy, which is maintained by an administrator for a computer that belongs to a domain. IP can do one of three things with any datagram:

- Apply IPsec protections to it.
- Allow it to pass unmodified.
- Discard it.

An IPsec policy contains one or more rules, each of which contain a filter, a filter action, authentication methods, a tunnel setting, and a connection type. For example, two stand-alone computers can be configured to use IPsec between them and activate the secure server policy. If the two computers are not members of the same or a trusted domain, trust must be configured using a certificate or preshared key in a secure server mode by:

- Setting up a filter that specifies all traffic between the two hosts
- Choosing an authentication method
- Selecting a negotiation policy (secure server in this case, indicating that all traffic matching the filter(s) must use IPsec)
- Specifying a connection type (LAN, dial-up, or all)

Once the policy has been put in place, traffic that matches the filters uses the services provided by IPsec. When a host directs IP traffic to another host (including something as simple as ping traffic), a Security Association (SA) is established through an Internet Key Exchange (IKE) negotiation using UDP port 500, and then the traffic begins to flow. The following network trace illustrates setting up a TCP connection between two such IPsec-enabled hosts. The only parts of the IP datagram that are unencrypted and visible to Network Monitor after the SA is established are the MAC and IP headers:

```
Source IP      Dest IP       Prot  Description
192.168.10.47  10.197.14.91  UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 216 (0xD8)
10.197.14.91   192.168.10.47 UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 216 (0xD8)
192.168.10.47  10.197.14.91  UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 128 (0x80)
10.197.14.91   192.168.10.47 UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 128 (0x80)
192.168.10.47  10.197.14.91  UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 76 (0x4C)
10.197.14.91   192.168.10.47 UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 76 (0x4C)
192.168.10.47  10.197.14.91  UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 212 (0xD4)
10.197.14.91   192.168.10.47 UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 172 (0xAC)
192.168.10.47  10.197.14.91  UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 84 (0x54)
10.197.14.91   192.168.10.47 UDP   Src Port: ISAKMP, (500); Dst Port: ISAKMP (500);
Length = 92 (0x5C)
192.168.10.47  10.197.14.91  IP    ID = 0xC906; Proto = 0x32; Len: 96
10.197.14.91   192.168.10.47 IP    ID = 0xA202; Proto = 0x32; Len: 96
192.168.10.47  10.197.14.91  IP    ID = 0xCA06; Proto = 0x32; Len: 88
```

Opening one of the IP datagrams sent after the SA is established reveals very little of what is actually in the datagram (a TCP SYN, or connection request). The only clear parts of the packet are the Ethernet and IP headers. Even the TCP header is encrypted and cannot be parsed by Network Monitor if ESP is used.

```
Src IP         Dest IP        Protoc  Description
==================================================
192.168.10.47  10.197.14.91   IP      ID = 0xC906; Proto = 0x32;
                                      Len: 96
+ FRAME: Base frame properties
+ ETHERNET: ETYPE = 0x0800 : Protocol = IP:  DOD Internet Protocol
IP: ID = 0xC906; Proto = 0x32; Len: 96
IP: Version = 4 (0x4)
IP: Header Length = 20 (0x14)
IP: Precedence = Routine
IP: Type of Service = Normal Service
```

```
IP: Total Length = 96 (0x60)
IP: Identification = 51462 (0xC906)
+ IP: Flags Summary = 2 (0x2)
IP: Fragment Offset = 0 (0x0) bytes
IP: Time to Live = 128 (0x80)
IP: Protocol = 0x32
IP: Checksum = 0xD55A
IP: Source Address = 192.168.10.47
IP: Destination Address = 10.197.14.91
IP: Data: Number of data bytes remaining = 76 (0x004C)
```

Using a secure server policy also restricts all other types of traffic from reaching destinations that do not understand IPsec or are not part of the same trusted group. A secure initiator policy provides settings that apply best to servers; traffic security is attempted, but if the client does not understand IPsec, the negotiation falls back to sending clear text packets.

When IPsec is used to encrypt data, network performance generally drops, due to the processing overhead of encryption. One possible method for reducing the impact of this overhead is to offload the processing to a hardware device. Because NDIS 5.1 supports task offloading, it is feasible to include encryption hardware on network adapters. Network adapters supporting IPsec hardware offload are available from several vendors.

IPsec promises to be popular for protecting both public network traffic and internal corporate or government traffic that requires confidentiality. One common deployment is to apply *secure server* IPsec policies only to specific servers that are used to store and/or serve confidential information.

**Internet Group Management Protocol (IGMP)**

Windows Server 2003 TCP/IP provides level 2 (full) support for IP multicasting and IGMP versions 1 through 3, as described in RFCs 1112, 2236, and 3376. The introduction to RFC 1112 provides a good overall summary of IP multicasting. The text reads:

"IP multicasting is the transmission of an IP datagram to a *host group*—a set of zero or more hosts identified by a single IP destination address. A multicast datagram is delivered to all members of its destination host group with the same 'best-effort' reliability as regular unicast IP datagrams; that is, the datagram is not guaranteed to arrive intact to all members of the destination group or in the same order relative to other datagrams.

"The membership of a host group is dynamic; that is, hosts may join and leave groups at any time. There is no restriction on the location or number of members in a host group. A host may be a member of more than one group at a time. A host need not be a member of a group to send datagrams to it.

"A host group may be permanent or transient. A permanent group has a well-known, administratively assigned IP address. It is the address—not the membership of the group—that is permanent; at any time a permanent group may have any number of members, even zero. Those IP multicast addresses that are not reserved for permanent groups are available for dynamic assignment to transient groups that exist only as long as they have members.

"Internetwork forwarding of IP multicast datagrams is handled by multicast routers that may be co-resident with, or separate from, Internet gateways. A host transmits an IP multicast datagram as a local network multicast that reaches all immediately-neighboring members of the destination host group. If the datagram has an IP time-to-live greater than 1, the multicast router(s) attached to the local network take responsibility for forwarding it towards all other networks that have members of the destination group. On those other member networks that are reachable within the IP time-to-live, an attached multicast router completes delivery by transmitting the datagram as a local multicast."

**IP/ARP Extensions for IP Multicasting**

To support IP multicasting, an additional route is defined on the host. The route (added by default) specifies that if a datagram is being sent to a multicast host group, it should be sent to the IP address of the host group through the local network interface card, and not forwarded to the default gateway. The following route (which you can discover using the **route print** command) illustrates this:

```
Network Address  Netmask    Gateway Address  Interface     Metric
224.0.0.0        240.0.0.0  157.60.137.88    157.60.137.88     20
```

IP multicast addresses are easily identified, as they are from the class D range, 224.0.0.0 to 239.255.255.255 (224.0.0.0/4). These IP addresses all have 1110 as their high-order bits.

To send a packet to a host group, using the local interface, the IP address must be resolved to a media access control address. As stated in RFC 1112:

"An IP host group address is mapped to an Ethernet multicast address by placing the low-order 23 bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01-00-5E-00-00-00 (hexadecimal). Because there are 28 significant bits in an IP host group address, more than one host group address may map to the same Ethernet multicast address."

For example, a datagram addressed to the multicast address 225.0.0.5 would be sent to the (Ethernet) MAC address 01-00-5E-00-00-05. This MAC address is formed by the junction of 01-00-5E and the 23 low-order bits of 225.0.0.5 (00-00-05).

Because more than one host group address can map to the same Ethernet multicast address, the interface may indicate hand-up multicasts for a host group for which no local applications have a registered interest. These extra multicasts are discarded by the TCP/IP protocol.

**Multicast Extensions to Windows Sockets**
Internet Protocol multicasting is currently supported only on AF_INET sockets of type SOCK_DGRAM and SOCK_RAW. By default, IP multicast datagrams are sent with a Time to Live (TTL) of 1. Applications can use the setsockopt() function to specify a TTL. By convention, multicast routers use TTL thresholds to determine how far to forward datagrams. These TTL thresholds are defined as follows:

- Multicast datagrams with initial TTL 0 are restricted to the same host.
- Multicast datagrams with initial TTL 1 are restricted to the same subnet.
- Multicast datagrams with initial TTL 32 are restricted to the same site.
- Multicast datagrams with initial TTL 64 are restricted to the same region.
- Multicast datagrams with initial TTL 128 are restricted to the same continent.
- Multicast datagrams with initial TTL 255 are unrestricted in scope.

**Use of Multicast and IGMP by Windows Components**
Some Windows Server 2003 components use IP multicast traffic. For example, router discovery uses multicasts, by default. WINS servers use multicasting when attempting to locate replication partners. The Windows Server 2003 Routing and Remote Access service supports multicast forwarding and the configuration of interfaces to operate in IGMP router or IGMP proxy mode.

**Transmission Control Protocol (TCP)**
TCP provides a connection-based, reliable byte-stream service to applications. Microsoft networking relies upon TCP for logon, file and print sharing, replication of information between domain controllers, transfer of browse lists, and other common functions. It can only be used for one-to-one communications.

TCP uses a checksum on both the TCP header and payload of each segment to reduce the chance that network corruption will go undetected. NDIS 5.1 provides support for task offloading, and Windows Server 2003 TCP/IP takes advantage of this by allowing the network interface card to perform the TCP checksum calculations if the network interface card driver offers support for this function. Offloading the checksum calculations to hardware can result in performance improvements in very high-throughput environments.

Windows Server 2003 TCP/IP has also been strengthened against a variety of attacks that were published over the past couple of years and has been subject to an internal security review intended to reduce susceptibility to future attacks. For instance, the initial sequence number (ISN) algorithm has been modified so that ISNs increase in random increments, using an RC4-based random number generator initialized with a 2048-bit random key upon system startup.

**TCP Receive Window Size Calculation and Window Scaling (RFC 1323)**
The TCP receive window size is the amount of receive data (in bytes) that can be buffered at one time on a connection. The sending host can send only that amount of data before waiting for an acknowledgment and window update from the receiving host. The Windows Server 2003 TCP/IP stack was designed to tune itself in most environments and uses larger default window sizes than earlier versions. Instead of using a hard-coded default receive window size, TCP adjusts to even increments of the maximum segment size (MSS) negotiated during connection setup. Matching the receive window to even increments of the MSS increases the percentage of full-sized TCP segments used during bulk data transmission.

The default advertised TCP receive-window size in Windows Server 2003 depends on the following, in order of precedence:

1. The SO_RCVBUF WinSock option for the connection.
2. The per-interface *TcpWindowSize* registry parameter.
3. The global *TcpWindowSize* registry parameter.

4. Autodetermination based on the NDIS-reported bit rate of the media The stack also tunes itself based on the media speed:

    - Below 1 Mbps: 8 KB
    - 1 Mbps – 100 Mbps: 17 KB
    - Greater than 100 Mbps: 64 KB

For 10 and 100 Mbps Ethernet, the window is normally set to 17,520 bytes (17 KB rounded up to twelve 1460-byte segments.) There are two methods for setting the receive window size to specific values:

- The *TcpWindowSize* registry parameter (see Appendix A)
- The setsockopt() Windows Sockets function (on a per-socket basis)

To improve performance on high-bandwidth, high-delay networks, Windows Server 2003 TCP/IP supports scalable windows support as described in RFC 1323. This RFC details a method for supporting scalable windows by allowing TCP to negotiate a scaling factor for the window size at connection establishment. This allows for an actual receive window of up to 1 gigabyte (GB). Section 2.2 of RFC 1323 provides a good description:

"The three-byte Window Scale option may be sent in a SYN segment by a TCP. It has two purposes: 1. indicate that the TCP is prepared to do both send and receive window scaling, and 2. communicate a scale factor to be applied to its receive window. Thus, a TCP that is prepared to scale windows should send the option, even if its own scale factor is 1. The scale factor is limited to a power of two and encoded logarithmically, so it may be implemented by binary shift operations.

"This option is an offer, not a promise; both sides must send Window Scale options in their SYN segments to enable window scaling in either direction. If window scaling is enabled, then the TCP that sent this option will right-shift its true receive-window values by 'shift.cnt' bits for transmission in SEG.WND. The value *shift.cnt* may be zero (offering to scale, while applying a scale factor of 1 to the receive window).

"This option may be sent in an initial <SYN> segment (in other words, a segment with the SYN bit on and the ACK bit off). It may also be sent in a <SYN,ACK> segment, but only if a Window Scale option was received in the initial <SYN> segment. A Window Scale option in a segment without a SYN bit should be ignored.

"The Window field in a SYN (in other words, a <SYN> or <SYN,ACK>) segment itself is never scaled."

When you read network traces of a connection that was established by two computers that support scalable windows, keep in mind that the window sizes advertised in the TCP headers must be scaled by the negotiated scale factor. The scale factor can be observed in the connection establishment (three-way handshake) packets, as illustrated in the following Network Monitor capture:

```
****************************************************************
Src Addr   Dst Addr  Protocol  Description
10.0.0.1   10.0.0.9  TCP       ....S., len:0, seq:725163-725163,
                               ack:0, win:65535, src:1217 dst:139
+ FRAME: Base frame properties
+ ETHERNET: ETYPE = 0x0800 : Protocol = IP:  DOD Internet
                                             Protocol
+ IP: ID = 0xB908; Proto = TCP; Len: 64
TCP: ....S., len:0, seq:725163-725163, ack:0, win:65535,
              src:1217  dst:139 (NBT Session)
TCP: Source Port = 0x04C1
TCP: Destination Port = NETBIOS Session Service
TCP: Sequence Number = 725163 (0xB10AB)
TCP: Acknowledgement Number = 0 (0x0)
TCP: Data Offset = 44 (0x2C)
TCP: Reserved = 0 (0x0000)
+ TCP: Flags = 0x02 : ....S.
TCP: Window = 65535 (0xFFFF)
TCP: Checksum = 0x8565
TCP: Urgent Pointer = 0 (0x0)
```

```
TCP: Options
+ TCP: Maximum Segment Size Option
TCP: Option Nop = 1 (0x1)
TCP: Window Scale Option
TCP: Option Type = Window Scale
TCP: Option Length = 3 (0x3)
TCP: Window Scale = 5 (0x5)
TCP: Option Nop = 1 (0x1)
TCP: Option Nop = 1 (0x1)
+ TCP: Timestamps Option
TCP: Option Nop = 1 (0x1)
TCP: Option Nop = 1 (0x1)
+ TCP: SACK Permitted Option
****************************************************************
```

The computer sending the packet above is offering the Window Scale option, with a scaling factor of 5. If the target computer responds, accepting the Window Scale option in the SYN-ACK, then it is understood that any TCP window advertised by this computer needs to be left-shifted 5 bits from this point onward (the SYN itself is not scaled). For example, if the computer advertised a 32 KB window in its first send of data, this value would need to be left-shifted (shifting in 0's from the right) 5 bits as shown below:

```
32 KB = 0x7fff =    111 1111 1111 1111
Left-shift 5 bits = 1111 1111 1111 1110 0000 = 0xfffe0 (1,048,544 bytes)
```

As a check, left-shifting a number 5 bits is equivalent to multiplying it by $2^5$, or 32. Performing the calculation in decimal, we get 32767 * 32 = 1,048,544.

The scale factor is not necessarily symmetrical, so it may be different for each direction of data flow.

TCP window scaling is negotiated on-demand in Windows Server 2003, based on the value set for the SO_RCVBUF WinSock option when a connection is initiated. Additionally, the Window Scale option is used by default on a connection if the received SYN segment for a connection initiated by a TCP peer contains the Window Scale option. You can modify this default behavior with the *Tcp1323Opts* registry parameter (see Appendix A).

**Delayed Acknowledgments**

As specified in RFC 1122, TCP uses delayed acknowledgments (ACKs) to reduce the number of packets sent on the media. The Windows Server 2003 TCP/IP stack takes a common approach to implementing delayed ACKs. As data is received by TCP on a connection, it only sends an acknowledgment back if one of the following conditions is met:

- No ACK was sent for the previous segment received.
- A segment is received, but no other segment arrives within 200 milliseconds for that connection.

In summary, normally an ACK is sent for every other TCP segment received on a connection, unless the delayed ACK timer (200 milliseconds) expires. The delayed ACK timer can be adjusted through the *TcpDelAckTicks* registry parameter.

**TCP Selective Acknowledgment (RFC 2018)**

Windows Server 2003 TCP/IP includes support for an important performance feature known as Selective Acknowledgement (SACK). SACK is especially important for connections using large TCP window sizes. Prior to SACK, a receiver could only acknowledge the latest sequence number of contiguous data that had been received, or the left edge of the receive window. When SACK is enabled, the receiver continues to use the ACK number to acknowledge the left edge of the receive window, but it can also acknowledge other non-contiguous blocks of received data individually. SACK uses two different TCP header options:

- The Sack-Permitted option is used to indicate that the sender can receive and interpret the SACK option. This option is sent in TCP segments in which the SYN flag is set.
- The SACK option is used to convey extended acknowledgment information from the receiver to the sender over an established TCP connection. Within the SACK option are up to four pairs of TCP sequence numbers acknowledging up to four non-continuous blocks. Each TCP sequence pair contains the left edge (the sequence number of the first byte in the block) and the right edge (the sequence number of the byte immediately following the last byte in the block) of the block being acknowledged.

When SACK is enabled (the default), a packet or series of packets can be dropped, and the receiver can inform the sender of

exactly which data has been received, and where the holes in the data are. The sender can then selectively retransmit the missing data without needing to retransmit blocks of data that have already been received successfully. SACK is controlled by the *SackOpts* registry parameter and is enabled by default. The Network Monitor capture below illustrates a host acknowledging all data up to sequence number 54857340, plus the data from sequence number 54858789-54861684.

```
+ FRAME: Base frame properties
+ ETHERNET: ETYPE = 0x0800 : Protocol = IP:  DOD Internet
                                            Protocol
+ IP: ID = 0x1A0D; Proto = TCP; Len: 64
TCP: .A...., len:0, seq:925104-925104, ack:54857341, win:32722,
               src:1242  dst:139
TCP: Source Port = 0x04DA
TCP: Destination Port = NETBIOS Session Service
TCP: Sequence Number = 925104 (0xE1DB0)
TCP: Acknowledgement Number = 54857341 (0x3450E7D)
TCP: Data Offset = 44 (0x2C)
TCP: Reserved = 0 (0x0000)
+ TCP: Flags = 0x10 : .A....
TCP: Window = 32722 (0x7FD2)
TCP: Checksum = 0x4A72
TCP: Urgent Pointer = 0 (0x0)
TCP: Options
TCP: Option Nop = 1 (0x1)
TCP: Option Nop = 1 (0x1)
+ TCP: Timestamps Option
TCP: Option Nop = 1 (0x1)
TCP: Option Nop = 1 (0x1)
TCP: SACK Option
TCP: Option Type = 0x05
TCP: Option Length = 10 (0xA)
TCP: Left Edge of Block  = 54858789 (0x3451425)
TCP: Right Edge of Block = 54861685 (0x3451F75)
```

**TCP Timestamps (RFC 1323)**

Windows Server 2003 TCP/IP supports TCP time stamps, as described in RFC 1323. Like SACK, time stamps are important for connections using large window sizes. Time stamps were conceived to assist TCP in accurately measuring round-trip time (RTT) to adjust retransmission time-outs.

The Timestamps option carries two four-byte time stamp fields. The time-stamp value field (TSval) contains the current value of the time-stamp clock of the TCP sending the option. The Timestamp Echo Reply field (TSecr) is only valid if the ACK bit is set in the TCP header; if it is valid, it echoes a timestamp value that was sent by the remote TCP peer in the TSval field of a Timestamps option. When TSecr is not valid, its value must be zero. The TSecr value will generally be from the most recent Timestamp option that was received; however, there are exceptions that are explained below.

A TCP peer may send the Timestamps option (TSopt) in an initial SYN segment (i.e., segment containing a SYN bit and no ACK bit), and may send a TSopt in other segments only if it received a TSopt in the initial SYN segment for the connection.

The Timestamps option field can be viewed in a Network Monitor capture by expanding the TCP options field, as shown below:

```
TCP: Timestamps Option
TCP: Option Type = Timestamps
TCP: Option Length = 10 (0xA)
TCP: Timestamp = 2525186 (0x268802)
TCP: Reply Timestamp = 1823192 (0x1BD1D8)
```

By default, the Timestamp option is only used on a connection if the received SYN segment for a connection initiated by a TCP peer contains the Timestamp option. This default behavior can be modified with the *Tcp1323Opts* registry parameter (see Appendix A).

**Path Maximum Transmission Unit (PMTU) Discovery**

PMTU discovery is described in RFC 1191. When a connection is established, the two hosts involved exchange their TCP maximum segment size (MSS) values. The smaller of the two MSS values is used for the connection. Historically, the MSS for a host has been the MTU at the link layer minus 40 bytes for the IP and TCP headers. However, support for additional TCP options, such as time stamps, has increased the typical TCP+IP header to 52 or more bytes. Figure 2 shows the difference between MTU and MSS.




**Figure 2. MTU versus MSS**

When TCP segments are destined to a non-local network, the Don't Fragment (DF) flag is set in the IP header. Any link along the path can have an MTU that differs from that of the two hosts. If a link has an MTU that is too small for the IP datagram being routed, the router attempts to fragment the datagram. However, the Don't Fragment bit is set in the IP header. At this point, the router should inform the sending host that the datagram couldn't be forwarded further without fragmentation. This is done with an ICMP Destination Unreachable-Fragmentation Needed and DF Set message. Most routers also specify the MTU for the next hop by putting the value for it in the low-order 16 bits of the ICMP header field that is unused in RFC 792. See RFC 1191, section 4, for the format of this message. Upon receiving this ICMP error message, TCP adjusts its MSS for the connection to the specified MTU minus the TCP and IP header size so that any further packets sent on the connection are no larger than the maximum size that can traverse the path without fragmentation.

| **Note** |
| --- |
| The minimum MTU permitted is 88 bytes, and Windows Server 2003 TCP/IP enforces this limit. |

Some noncompliant routers may silently drop IP datagrams that cannot be fragmented or may not correctly report their next-hop MTU. If this occurs, it may be necessary to make a configuration change to the PMTU detection algorithm. There are two registry changes that can be made to Windows Server 2003 TCP/IP stack to work around these problematic devices:

- *EnablePMTUBHDetect*—Adjusts the PMTU discovery algorithm to attempt to detect PMTU black hole routers. PMTU black hole detection is disabled by default.

- *EnablePMTUDiscovery*—Completely enables or disables the PMTU discovery mechanism. When PMTU discovery is disabled, an MSS of 536 bytes is used for all non-local destination addresses. PMTU discovery is enabled by default (this is the recommended setting).

These registry entries are described in more detail in Appendix A.

The PMTU between two hosts can be discovered manually using the **ping** tool with the **-f** (don't fragment) switch, as follows:

**ping -f -n** <*number of pings*> **-l** <*size*> <*destination ip address*>

As shown in the example below, the *size* parameter can be varied until the MTU is found. The *size* parameter used by ping is the size of the data buffer to send, not including headers. The ICMP header consumes 8 bytes, and the IP header is normally 20 bytes. In the case below (Ethernet), the link layer MTU is the maximum-sized ping buffer plus 28, or 1500 bytes:

```
C:\>ping -f -n 1 -l 1472 10.99.99.10
Pinging 10.99.99.10 with 1472 bytes of data:
Reply from 10.99.99.10: bytes=1472 time<10ms TTL=128
Ping statistics for 10.99.99.10:
Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum =  0ms, Average =  0ms

C:\>ping -f -n 1 -l 1473 10.99.99.10
Pinging 10.99.99.10 with 1473 bytes of data:
Packet needs to be fragmented but DF set.
Ping statistics for 10.99.99.10:
Packets: Sent = 1, Received = 0, Lost = 1 (100% loss),
Approximate round trip times in milliseconds:
Minimum = 0ms, Maximum =  0ms, Average =  0ms
```

In the example shown above, the IP layer returned an ICMP error message that ping interpreted. If the router had been a PMTU black hole router, ping would simply not be answered once its size exceeded the MTU that the router could handle. Ping can be used in this manner to detect such a router.

A sample ICMP Destination Unreachable-Fragmentation Needed and DF Set message is shown here:

```
*****************************************************************
Src Addr        Dst Addr      Protocol  Description
10.99.99.10     10.99.99.9    ICMP      Destination Unreachable:
```

```
                                        10.99.99.10   See frame 3
+ FRAME: Base frame properties
+ ETHERNET: ETYPE = 0x0800 : Protocol = IP:  DOD Internet Protocol
+ IP: ID = 0x4401; Proto = ICMP; Len: 56
ICMP: Destination Unreachable: 10.99.99.10   See frame 3
ICMP: Packet Type = Destination Unreachable
ICMP: Unreachable Code = Fragmentation Needed, DF Flag Set
ICMP: Checksum = 0xA05B
ICMP: Next Hop MTU = 576 (0x240)
ICMP: Data: Number of data bytes remaining = 28 (0x001C)
ICMP: Description of original IP frame
```

This error was generated by using **ping -f –n 1 -l 1000** on an Ethernet-based host to send a large datagram across a router interface that only supports an MTU of 576 bytes. When the router tried to place the large frame onto the network with the smaller MTU, it found that fragmentation was not allowed. Therefore, it returned the error message indicating that the largest datagram that could be forwarded is 576 bytes.

**Dead Gateway Detection**

Dead gateway detection is used to allow TCP to detect failure of the default gateway and to adjust the IP route table to use another default gateway. Windows Server 2003 TCP/IP uses the triggered reselection method described in RFC 816, with slight modifications based upon customer experience and feedback.

When a TCP connection routed through the default gateway attempts to send a TCP packet to the destination a number of times (equal to one-half of the registry value *TcpMaxDataRetransmissions*) without receiving a response, the algorithm changes the Route Cache Entry (RCE) for that remote IP address to use the next default gateway in the list as its next-hop address. When 25 percent of the TCP connections have moved to the next default gateway, the algorithm advises IP to change the computer's default gateway to the one that the connections are now using.

For example, assume that there are currently TCP connections to 11 different IP addresses that are being routed through the default gateway. Now assume that the default gateway fails, that there is a second default gateway configured, and that the value for *TcpMaxDataRetransmissions* is at the default of 5.

When the first TCP connection tries to send data, it does not receive any acknowledgments. After the third retransmission, the RCE for that remote IP address is switched to the next default gateway in the list. At this point, any TCP connections to that one remote IP address have switched over, but the remaining connections still try to use the original default gateway.

When the second TCP connection tries to send data, the same thing happens. Now, two of the 11 RCEs point to the new gateway.

When the third TCP connection tries to send data, after the third retransmission, three of 11 RCEs have been switched to the second default gateway. Because, at this point, over 25 percent of the RCEs have been moved, the default gateway for the whole computer is moved to the new one.

That default gateway remains the primary one for the computer until it experiences problems (causing the dead gateway algorithm to try the next one in the list again) or until the computer is restarted.

When the search reaches the last default gateway, it returns to the beginning of the list.

**TCP Retransmission Behavior**

TCP starts a retransmission timer when each outbound segment is handed down to IP. If no acknowledgment has been received for the data in a given segment before the timer expires, the segment is retransmitted. For new connection requests, the retransmission timer is initialized to 3 seconds (controllable using the *TcpInitialRtt* per-adapter registry parameter), and the request (SYN) is resent up to the value specified in *TcpMaxConnectRetransmissions* (the default for Windows Server 2003 is 2 times). On existing connections, the number of retransmissions is controlled by the *TcpMaxDataRetransmissions* registry parameter (5 by default). The retransmission time-out is adjusted on the fly to match the characteristics of the connection, using Smoothed Round Trip Time (SRTT) calculations as described in Van Jacobson's article named "Congestion Avoidance and Control." The timer for a given segment is doubled after each retransmission of that segment. Using this algorithm, TCP tunes itself to the normal delay of a connection. TCP connections over high-delay links take much longer to time-out than those over low-delay links.[3]

The following Network Monitor capture shows the retransmission algorithm for two hosts that are connected over Ethernet on the same subnet. An FTP file transfer was in progress when the receiving host was disconnected from the network. Because the SRTT for this connection was very small, the first retransmission was sent after about one-half second. The timer was then doubled for each of the retransmissions that followed. After the fifth retransmission, the timer was once again doubled. If no acknowledgment was received before it expired, the connection was aborted.

```
Delta  Source IP    Dest IP      Pro Flags Description
```

```
-----------------------------------------------------------
0.000  10.57.10.32  10.57.9.138  TCP .A.., len: 1460, seq: 8043781, ack: 8153124, win:
8760
0.521  10.57.10.32  10.57.9.138  TCP .A.., len: 1460, seq: 8043781, ack: 8153124, win:
8760
1.001  10.57.10.32  10.57.9.138  TCP .A.., len: 1460, seq: 8043781, ack: 8153124, win:
8760
2.003  10.57.10.32  10.57.9.138  TCP .A.., len: 1460, seq: 8043781, ack: 8153124, win:
8760
4.007  10.57.10.32  10.57.9.138  TCP .A.., len: 1460, seq: 8043781, ack: 8153124, win:
8760
8.130  10.57.10.32  10.57.9.138  TCP .A.., len: 1460, seq: 8043781, ack: 8153124, win:
8760
```

There are some circumstances under which TCP retransmits data prior to the time that the retransmission timer expires. The most common of these occurs due to a feature known as *fast retransmit*. When a receiver that supports fast retransmit receives data with a sequence number beyond the current expected one, it assumes that some data was dropped. To help make the sender aware of this event, the receiver immediately sends an ACK, with the ACK number set to the sequence number that it was expecting. It continues to do this for each additional TCP segment that arrives containing data subsequent to the missing data in the incoming stream. When the sender starts to receive a stream of ACKs that are acknowledging the same sequence number and that sequence number is earlier than the current sequence number being sent, it can infer that a segment (or more) must have been dropped. Senders that support the fast retransmit algorithm immediately resend the segment that the receiver is expecting to fill in the gap in the data, without waiting for the retransmission timer to expire for that segment. This optimization greatly improves performance in a busy network environment.

With fast retransmit, the sender retransmits the TCP segments to fill in the gap in the data before their retransmission timers expire. Because the retransmission timers did not expire for the missing TCP segments, the sender can more quickly send additional segments to the receiver. This is known as *fast recovery*. Fast retransmit and fast recovery are described in RFC 2581.

By default, Windows Server 2003 resends a segment if it receives three ACKs for the same sequence number and that sequence number lags the current one. This is controllable with the *TcpMaxDupAcks* registry parameter. See also the "TCP Selective Acknowledgment (RFC 2018)" section in this article.

**TCP Keep-Alive Messages**

A TCP keep-alive packet is simply an ACK with the sequence number set to one less than the current sequence number for the connection. A host receiving one of these ACKs responds with an ACK for the current sequence number. Keep-alives can be used to verify that the computer at the remote end of a connection is still available. TCP keep-alives can be sent once every *KeepAliveTime* (defaults to 7,200,000 milliseconds or two hours) if no other data or higher-level keep-alives have been carried over the TCP connection. If there is no response to a keep-alive, it is repeated once every *KeepAliveInterval* seconds. *KeepAliveInterval* defaults to 1 second. NetBT connections, such as those used by other Microsoft networking components, send NetBIOS keep-alives more frequently, so normally no TCP keep-alives are sent on a NetBIOS connection. TCP keep-alives are disabled by default, but Windows Sockets applications can use the **setsockopt()** function to enable them.

**Slow Start Algorithm and Congestion Avoidance**

When a connection is established, TCP starts slowly at first to assess the bandwidth of the connection, and to avoid overflowing the receiving host or any other devices or links in the path. The send window is set to two TCP segments, and if that is acknowledged, it is incremented to three segments.[4] If those are acknowledged, it is incremented again, and so on until the amount of data being sent per burst reaches the size of the receive window on the remote host. At that point, the slow start algorithm is no longer in use, and flow control is governed by the receive window. However, congestion could still occur on a connection at any time during transmission. If this happens (evidenced by the need to retransmit), a congestion-avoidance algorithm is used to reduce the send window size temporarily and to grow it back towards the receive window size. Slow start and congestion avoidance are described in RFC 2581.

**Silly Window Syndrome (SWS)**

Silly Window Syndrome is described in RFC 1122 as follows:

"In brief, SWS is caused by the receiver advancing the right window edge whenever it has any new buffer space available to receive data and by the sender using any incremental window, no matter how small, to send more data [TCP:5]. The result can be a stable pattern of sending tiny data segments, even though both sender and receiver have a large total buffer space for the connection."

Windows Server 2003 TCP/IP implements SWS avoidance, as specified in RFC 1122, by not sending more data until there is a sufficient window size advertised by the receiving end to send a full TCP segment. It also implements SWS avoidance on the receive end of a connection by not opening the receive window in increments of less than a TCP segment.

**Nagle Algorithm**
Windows Server 2003 TCP/IP implements the Nagle algorithm described in RFC 896. The purpose of this algorithm is to reduce the number of very small segments sent, especially on high-delay (remote) links. The Nagle algorithm allows only one small segment to be outstanding at a time without acknowledgment. If more small segments are generated while awaiting the ACK for the first one, these segments are coalesced into one larger segment. Any full-sized segment is always transmitted immediately, on the assumption that there is a sufficient receive window available. The Nagle algorithm is effective in reducing the number of packets sent by interactive applications, such as Telnet, especially over slow links.

The Nagle algorithm can be observed in the following Network Monitor capture. The trace was captured by using PPP to dial up an Internet service provider. A Telnet (character-mode) session was established, and then the **Y** key was held down on the computer. At all times, one segment was sent, and further **Y** characters were held by the stack until an acknowledgment was received for the previous segment. In this example, three to four **Y** characters were buffered each time and sent together in one segment. The Nagle algorithm resulted in a huge savings in the number of packets sent—the number of packets was reduced by a factor of about three.

```
Delta Source IP     Dest IP       Prot   Description
-----------------------------------------------------------
0.644 204.182.66.83 199.181.164.4 TELNET To Server Port = 1901
0.144 199.181.164.4 204.182.66.83 TELNET To Client Port = 1901
0.000 204.182.66.83 199.181.164.4 TELNET To Server Port = 1901
0.145 199.181.164.4 204.182.66.83 TELNET To Client Port = 1901
0.000 204.182.66.83 199.181.164.4 TELNET To Server Port = 1901
0.144 199.181.164.4 204.182.66.83 TELNET To Client Port = 1901
. . .
```

Each segment contained several of the "y" characters. The first segment is shown more fully parsed below, and the data portion is pointed out in the hexadecimal display at the bottom.

```
****************************************************************
Time  Source IP      Dest IP        Prot    Description
0.644 204.182.66.83  199.181.164.4  TELNET  To Server Port = 1901
+ FRAME: Base frame properties
+ ETHERNET: ETYPE = 0x0800 : Protocol = IP: DOD Internet Protocol
+ IP: ID = 0xEA83; Proto = TCP; Len: 43
+ TCP: .AP..., len: 3, seq:1032660278, ack: 353339017, win: 7766, src: 1901 dst: 23
(TELNET)
TELNET: To Server From Port = 1901
TELNET: Telnet Data
D2 41 53 48 00 00 52 41 53 48 00 00 08 00 45 00  .ASH..RASH....E.
00 2B EA 83 40 00 20 06 F5 85 CC B6 42 53 C7 B5  .+..@. .....BS..
A4 04 07 6D 00 17 3D 8D 25 36 15 0F 86 89 50 18  ...m..=.%6....P.
1E 56 1E 56 00 00 79 79 79                       .V.V..yyy
^^^
data
```

Windows Sockets applications can disable the Nagle algorithm for their connections by setting the TCP_NODELAY socket option. However, this practice should be avoided unless it is absolutely necessary because it increases network utilization. Some network applications may not perform well if their design does not take into account the effects of transmitting large numbers of small packets and the Nagle algorithm. The Nagle algorithm is not applied to loopback TCP connections for performance reasons. Windows Server 2003 NetBT disables Nagling for NetBIOS over TCP connections as well as direct-hosted redirector/server connections, which can improve performance for applications issuing numerous small file manipulation commands. An example is an application that uses file locking/unlocking frequently.

**TCP TIME-WAIT Delay**
When a TCP connection is closed, the socket-pair is placed into a state known as TIME-WAIT. This is done so that a new connection does not use the same protocol, source IP address, destination IP address, source port, and destination port until enough time has passed to ensure that any segments that may have been misrouted or delayed are not delivered unexpectedly. RFC 793 specifies the length of time that the socket-pair should not be reused as two maximum segment lifetimes (2 MSL), or four minutes. This is the default setting for Windows Server 2003 TCP/IP. However, with this default setting, some network applications that perform many outbound connections in a short time may use up all available ports before the ports can be recycled.

Windows Server 2003 TCP/IP offers two methods of controlling this behavior. First, the *TcpTimedWaitDelay* registry parameter can be used to alter this value. Windows Server 2003 TCP/IP allows it to be set as low as 30 seconds, which should not cause problems in most environments. Second, the number of user-accessible ephemeral ports that can be used to source outbound connections is configurable using the *MaxUserPort* registry parameter. By default, when an application requests any socket from the system to use for an outbound call, a port between the values of 1024 and 5000 is supplied. The *MaxUserPort*

parameter can be used to set the value of the uppermost port that the administrator chooses to allow for outbound connections. For instance, setting this value to 10,000 (decimal) would make approximately 9000 user ports available for outbound connections. For more details on this concept, see RFC 793. See also the *MaxFreeTcbs* and *MaxHashTableSize* registry parameters in Appendix A.

TCP/IP for Windows Server 2003 Service Pack 1 has implemented a smart TCP port allocation algorithm. When an application requests any available TCP port, TCP/IP first attempts to find an available port that does not correspond to a connection in the TIME WAIT state. If a port cannot be found, then it picks any available port. For more information, see "Smart TCP Port Allocation" in this article.

**TCP Connections to and from Multihomed Computers**

When TCP connections are made to a multihomed host, both the WINS client and the Domain Name Resolver (DNR) attempt to determine whether any of the destination IP addresses provided by the name server are on the same subnet as any of the interfaces in the local computer. If so, these addresses are sorted to the top of the list so that the application can try them prior to trying addresses that are not on the same subnet. If none of the addresses is on a common subnet with the local computer, behavior is different depending upon the name space. The *PrioritizeRecordData* TCP/IP registry parameter can be used to prevent the DNR component from sorting local subnet addresses to the top of the list.

In the WINS name space, the client is responsible for randomizing or load balancing between the provided addresses. The WINS server always returns the list of addresses in the same order, and the WINS client randomly picks one of them for each connection.

In the DNS name space, the DNS server is usually configured to provide the addresses in a round robin fashion. The DNR does not attempt to further randomize the addresses. In some situations, it is desirable to connect to a specific interface on a multihomed computer. The best way to accomplish this is to provide the interface with its own DNS entry. For example, a computer named *raincity* could have one DNS entry listing both IP addresses (actually two separate records in the DNS with the same name), and also records in the DNS for raincity1 and raincity2, each associated with just one of the IP addresses assigned to the computer.

When TCP connections are made from a multihomed host, things get a bit more complicated. If the connection is a Winsock connection using the DNS name space, once the target IP address for the connection is known, TCP attempts to connect from the best source IP address available. Again, the route table is used to make this determination. If there is an interface in the local computer that is on the same subnet as the target IP address, its IP address is used as the source in the connection request. If there is no best source IP address to use, the system chooses one randomly.

If the connection is a NetBIOS-based connection using the redirector, little routing information is available at the application level. The NetBIOS interface supports connections over various protocols and has no knowledge of IP. Instead, the redirector places calls on all of the transports that are bound to it. If there are two interfaces in the computer and one protocol installed, there are two transports available to the redirector. Calls are placed on both, and NetBT submits connection requests to the stack, using an IP address from each interface. It is possible that both calls succeed. If so, the redirector cancels one of them. The choice of which one to cancel depends upon the redirector *ObeyBindingOrder* registry value[5]. If this is set to 0 (the default value), the primary transport (determined by binding order) is the preferred one, and the redirector waits for the primary transport to time out before accepting the connection on the secondary transport. If this value is set to 1, the binding order is ignored, and the redirector accepts the first connection that succeeds and cancels the other(s).

**Throughput Considerations**

TCP was designed to provide optimum performance over varying link conditions, and Windows Server 2003 TCP/IP contains improvements such as those supporting RFC 1323. Actual throughput for a link depends on a number of variables, but the most important factors are:

- Link speed (bits-per-second that can be transmitted)
- Propagation delay
- Window size (amount of unacknowledged data that may be outstanding on a TCP connection)
- Link reliability
- Network and intermediate device congestion
- Path MTU

TCP throughput calculation is discussed in detail in Chapters 20–24 of *TCP/IP Illustrated*, by W. Richard Stevens[6]. Some key considerations are listed below:

- The capacity of a pipe is bandwidth multiplied by round-trip time. This is known as the *bandwidth-delay product*. If the link is reliable, for best performance the window size should be greater than or equal to the capacity of the pipe so that the sending stack can fill it. The largest window size that can be specified, due to its 16-bit field in the TCP header, is 65535, but larger windows can be negotiated by using window scaling as described earlier in this article. See *TcpWindowSize* in Appendix A.

- Throughput can never exceed window size divided by round-trip time.

- If the link is unreliable or badly congested and packets are being dropped, using a larger window size may not improve throughput. Along with scaling windows support, Windows Server 2003 TCP/IP supports Selective Acknowledgments (SACK; described in RFC 2018) to improve performance in environments that are experiencing packet loss. It also includes support for timestamps (described in RFC 1323) for improved RTT estimation.

- Propagation delay is dependent upon the speed of light, latencies in transmission equipment, and so on.

- Transmission delay depends on the speed of the media.

- For a specified path, propagation delay is fixed, but transmission delay depends upon the packet size.

- At low speeds, transmission delay is the limiting factor. At high speeds, propagation delay may become the limiting factor.

To summarize, Windows Server 2003 TCP/IP can adapt to most network conditions and can dynamically provide the best throughput and reliability possible on a per-connection basis. Attempts at manual tuning are often counter-productive unless a qualified network engineer first performs a careful study of data flow.

**User Datagram Protocol (UDP)**
UDP provides a connectionless, unreliable transport service. It is often used for communications that use broadcast or multicast IP datagrams. Since delivery of UDP datagrams is not guaranteed, applications using UDP must supply their own mechanisms for reliability, if needed. Microsoft networking components use UDP for logon, browsing, and name resolution. UDP can also be used to carry IP multicast streams.

**UDP and Name Resolution**
UDP is used for NetBIOS name resolution by unicast to a NetBIOS name server or subnet broadcasts, and for DNS host name to IP address resolution. NetBIOS name resolution is accomplished over UDP port 137. DNS queries use UDP port 53. Because UDP itself does not guarantee delivery of datagrams, both of these services use their own retransmission schemes if they receive no answer to queries. Broadcast UDP datagrams are not usually forwarded over IP routers, so NetBIOS name resolution in a routed environment requires a name server of some type, or the use of static database files.

**Mailslots over UDP**
Many NetBIOS applications use *mailslot* messaging. A second-class mailslot is a simple mechanism for sending a message from one NetBIOS name to another over UDP. Mailslot messages can be broadcast on a subnet or directed to the remote host. To direct a mailslot message to another host, there must be some method of NetBIOS name resolution available. Microsoft provides WINS for this purpose.

**NetBIOS over TCP/IP (NetBT)**
The Windows Server 2003 implementation of NetBIOS over TCP/IP is referred to as *NetBT*. NetBT uses the following TCP and UDP ports:

- UDP port 137 (name services)

- UDP port 138 (datagram services)

- TCP port 139 (session services)

NetBIOS over TCP/IP is specified by RFC 1001 and RFC 1002. The Netbt.sys driver is a kernel-mode component that supports

the Transport Driver Interface (TDI) interface. Services such as Workstation (redirector) and Server (file server) use the TDI interface directly, but traditional NetBIOS applications have their calls mapped to TDI calls by the Netbios.sys driver. Using TDI to make calls to NetBT is a more difficult programming task, but can provide higher performance and freedom from historical NetBIOS limitations. NetBIOS concepts are discussed further in the "Network Application Interfaces" section of this article.

**Transport Driver Interface (TDI)**
Microsoft developed the Transport Driver Interface (TDI) to provide greater flexibility and functionality than is provided by existing interfaces, such as NetBIOS and Windows Sockets. All Windows transport providers expose TDI. The TDI specification describes the set of primitive functions by which transport drivers and TDI clients communicate and the call mechanisms used for accessing them. Currently, TDI is kernel-mode only.

The Windows Server 2003 redirector and server both use TDI directly, rather than going through the NetBIOS mapping layer. By doing so, they are not subject to many of the restrictions imposed by NetBIOS, such as the legacy 254-session limit.

**TDI Features**
TDI may be the most difficult to use of all Windows network APIs. It is a simple conduit, so the programmer must determine the format and meaning of messages.

TDI includes the following features:

- Most Windows Server 2003 transports support TDI
- An open naming and addressing scheme
- Message and stream-mode data transfer
- Asynchronous operation
- Support for unsolicited indication of events
- Extensibility—clients can submit private requests to a transport driver that understands them.
- Support for limited use of standard kernel-mode I/O functions to send and receive data
- 32-bit addressing and values
- Support for Access Control Lists (ACLs, used for security) on TDI address objects

More information about TDI is available from the Windows DDK.

**Security Considerations**
Network security is a serious consideration for administrators with computers exposed to public networks. The Microsoft TCP/IP stack has been strengthened against many attacks and in its default state handles most of the common attacks. Some additional protection against popular Denial of Service attacks can be added by setting the value of the *SynAttackProtect* parameter in the registry. This key allows the administrator to choose several levels of protection against SYN attacks.

Here are general guidelines that can lower your exposure to attack:

- Disable unnecessary or optional services (for instance, the Client for Microsoft Networks and the File and Printer Sharing for Microsoft Networks components on the network connections of an IIS server).
- Enable TCP/IP filtering and restrict access to only the ports that are necessary for the server to function. See article 150543 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=67962) for a list of ports that Windows services use.
- Disable NetBIOS over TCP/IP on network connections where it is not needed.
- Configure static IP addresses and parameters for network adapters connected to the Internet.
- Configure registry parameters for maximum protection (see Appendix D).

Consult the Microsoft Security Web site (http://go.microsoft.com/fwlink/?linkID=7420) regularly for security bulletins.

# Network Application Interfaces

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are a number of ways that network applications can communicate using the TCP/IP protocol stack. Some of them, such as named pipes, go through the network redirector, which is part of the Workstation service. Many older applications were written to the NetBIOS interface, which is supported by NetBIOS over TCP/IP.

The Windows Sockets interface is the most popular API for writing Windows network applications today. A quick overview of the Windows Sockets Interface and the NetBIOS Interface is presented here.

**Windows Sockets**
Windows Sockets specifies a programming interface that was originally based on the familiar socket interface from the University of California at Berkeley. It includes a set of extensions designed to take advantage of the message-driven nature of Microsoft Windows. Version 1.1 of the specification was released in January 1993, and version 2.2.0 was published in May of 1996. Windows Server 2003 supports version 2.2, commonly referred to as Winsock2.

**Applications**
There are many Windows Sockets applications available. A number of the utilities that ship with Windows Server 2003 are based on Windows Sockets, including the FTP and DHCP clients and servers, Telnet client, and so on. There are also higher-level programming interfaces that rely on Winsock, such as the Windows Internet API (WinInet) used by Internet Explorer.

**Name and Address Resolution**
Traditionally, Windows Sockets applications use the **gethostbyname()** function to resolve a host name to an IP address and the **gethostbyaddr()** function to resolve an IP address to a host name. The **gethostbyname()** and **gethostbyaddr()** functions are specific to IPv4. To accommodate the use of either IPv4 or IPv6, these functions have been deprecated in RFC 2553 and replaced by the new functions **getaddrinfo()** and **getnameinfo()**. Both **getaddrinfo()** and **getnameinfo()** are independent of whether IPv4, IPv6, or both are installed on the computer.

To allow applications to run over both IPv4 and IPv6, application developers must modify their applications to use **getaddrinfo()** and **getnameinfo()** and make other modifications. For more information about modifying applications to run over both IPv4 and IPv6, see the IPv6 Guide for Windows Sockets Applications (http://go.microsoft.com/fwlink/?linkID=68008).

The **gethostbyname()** or **getaddrinfo()** functions use the following (default) name look-up sequence:

1. Check the local host name for a matching name.
2. Check the Hosts file for a matching name entry.
3. If a DNS server is configured, query it.
4. If no match is found, attempt NetBIOS name-resolution.

The **gethostbyaddr()** or **getnameinfo()** functions use the following (default) address look-up sequence:

1. Check the Hosts file for a matching address entry.
2. If a DNS server is configured, query it.
3. Send a NetBIOS Adapter Status Request to the IP address being queried. If it responds with a list of NetBIOS names registered for the adapter, parse it for the computer name.

**Support for IP Multicasting**
Winsock2 provides support for IP multicasting. Multicasting is described in the Windows Sockets 2.0 specification and in the IGMP section of this article. IP multicasting is currently supported only on AF_INET and AF_INET6 sockets of the types SOCK_DGRAM and SOCK_RAW.

Winsock2 and Windows Server 2003 also support IP multicasting on SOCK_RDM sockets, which provide reliable multicast delivery using the Pragmatic General Multicast (PGM) protocol. Installing the Reliable Multicast Protocol through Network

Connections enables PGM and reliable multicast support. The Reliable Multicast Protocol installs a driver, Rmcast.sys, which communicates directly with the TCP/IP protocol to provide its own transport layer. Windows XP does not include the Reliable Multicast Protocol. PGM is described in RFC 3208.

**Backlog Parameter**
Windows Sockets server applications generally create a socket, and then use the **listen()** function on it to listen for connection requests. One of the parameters passed when calling **listen()** is the backlog of connection requests that the application would like Windows Sockets to queue for it. This value controls the number of unaccepted connections that can be queued. Once an application *accepts* a connection, it is moved out of the connection request backlog and no longer counts. The Windows Sockets 1.1 specification indicates that the maximum allowable value for a backlog is 5; however, Windows NT 4.0, Windows 2000 Server, and Windows Server 2003 accept a backlog of 200. Windows XP Professional accepts a backlog of 5.

**Push Bit Interpretation**
By default, Windows Server 2003 TCP/IP completes a **recv()** call when one of the following conditions is met:

- Data arrives with the PUSH bit set
- The user **recv** buffer is full
- 0.5 seconds have elapsed since any data has arrived

If a client application is run on a computer with a TCP/IP implementation that does not set the push bit on send operations, response delays may result. It is best to correct this on the client; however, a configuration parameter (*IgnorePushBitOnReceives*) was added to Afd.sys to force it to treat all arriving packets as though the push bit were set.

**ConnectEx/TransmitPackets and TCP/IP**
The following two functions are Microsoft-specific extensions to the Windows Sockets 2 specification:

- The Windows Sockets **ConnectEx()** function establishes a connection to another socket application and optionally sends the block of data after the connection is established.
- The Windows Sockets **TransmitPackets()** function transmits in memory and/or file data over a connected socket (either datagram or stream). The operating system's cache manager is used to retrieve file data and locks memory for the minimum necessary time required to transmit it. This provides high performance and efficiency for file and memory data transfer over sockets.

**Windows Sockets Direct Path for System Area Networks**
The Windows Server 2003 family contains substantial performance improvements to Windows Sockets Direct (WSDP) for System Area Networks (SANs). WSD enables Windows Sockets applications written for SOCK_STREAM to obtain the performance benefits of SANs without having to make application modifications. The fundamental component of this technology is the WinSock switch that emulates TCP/IP semantics over native SAN service providers. For the Windows 2000 Server family, WSD support was only available for Windows 2000 Advanced Server and Windows 2000 Datacenter Server. WSD support is included for all members of the Windows Server 2003 family.

**NetBIOS over TCP/IP**
NetBIOS defines a software interface and a naming convention, not a protocol. Early versions of Microsoft networking products provided only the NetBIOS Extended User Interface (NetBEUI) protocol with a NetBIOS application-programming interface. NetBEUI is a small, fast protocol with no networking layer; thus, it is not routable and is not suitable for internetworks. NetBEUI relies on multicasts for name resolution and location of services. NetBIOS over TCP/IP provides the NetBIOS programming interface over the TCP/IP protocol, extending the reach of NetBIOS client and server programs to the WAN, and providing interoperability with various other operating systems.

The Workstation, Server, Browser, Messenger, and NetLogon services are all (direct) NetBT clients. They use TDI (described earlier in this article) to communicate with NetBT. Windows Server 2003 also includes a NetBIOS emulator. The emulator takes standard NetBIOS requests from NetBIOS applications and translates them to equivalent TDI primitives.

Windows Server 2003 still uses NetBIOS over TCP/IP to communicate with prior versions of Windows NT and other clients, such as Windows 98. However, the Windows Server 2003 redirector and server components also support *direct hosting* to communicate with other computers capable of direct hosting. Direct hosting uses DNS for name resolution. No NetBIOS name resolution (WINS or broadcast) is used, and the protocol is simpler. Direct hosting uses TCP port 445, instead of the NetBIOS TCP port 139.

By default, both NetBIOS and direct hosting are enabled, and both are tried in parallel when a new connection is established. The first to succeed in connecting is used for any attempt. NetBIOS support can be disabled to force all traffic to use direct hosting.

**To disable NetBIOS support**

1. On the **Start** menu, point to **Control Panel**, and then double-click **Network Connections**. Right-click **Local Area Connection** (or the name of your LAN connection) and then click **Properties**.
2. Select **Internet Protocol (TCP/IP)** in the list of components, and click **Properties**.
3. Click **Advanced**.
4. On the **WINS** tab, select **Disable NetBIOS over TCP/IP**.

Applications and services that depend on NetBIOS no longer function after this is done, so it is important that you verify that any clients and applications no longer need NetBIOS support before you disable it. For example, computers that cannot direct host will be unable to browse, locate, or create file and print share connections to a Windows Server 2003 computer with NetBIOS disabled.

**NetBIOS Names**

The NetBIOS namespace is flat, meaning that all names within the name space must be unique. NetBIOS names are 16 characters in length. Resources are identified by NetBIOS names, which are registered dynamically when computers boot, services or applications start, or users log on. Names can be registered as unique (one owner) or as group (multiple owner) names. A NetBIOS Name Query is used to locate a resource by resolving the name to an IP address.

Microsoft networking components, such as Workstation and Server services, allow the first 15 characters of a NetBIOS name to be specified by the user or administrator, but reserve the sixteenth character of the NetBIOS name to indicate a resource type (0x00-0xFF). Many popular third-party software packages also use this character to identify and register their specific services. Table 3 lists some example NetBIOS names used by Microsoft components.

**Table 3. Examples of NetBIOS names used by Windows components**

| Unique name | Service |
|---|---|
| *computer_name* **[00h]** | Workstation service |
| *computer_name* **[03h]** | Messenger service |
| *computer_name* **[1Fh]** | NetDDE service |
| *computer_name* **[20h]** | Server service |
| *computer_name* **[BEh]** | Network Monitor Agent |
| *computer_name* **[BFh]** | Network Monitor Application |
| *user_name* **[03]** | Messenger service |
| *domain_name* **[1Dh]** | Master Browser |
| *domain_name* **[1Bh]** | Domain Master Browser |
| *domain_name* **[00h]** | Domain name |
| *domain_name* **[1Ch]** | Domain controllers |
| *domain_name* **[1Eh]** | Browser service elections |
| \\--__MSBROWSE__[01h] | Master browser |

To see which names a computer has registered over NetBT, type the following from a command prompt:

nbtstat -n

Windows Server 2003 allows you to re-register names with the WINS server after a computer has already been started. To do this, type the following from a command prompt:

nbtstat –RR

**NetBIOS Name Registration and Resolution**
Windows TCP/IP systems use several methods to locate NetBIOS resources:

- NetBIOS name cache
- NetBIOS name server
- Broadcast NetBIOS Name Query Request messages
- Static Lmhosts file
- Local host name (optional, depends on *EnableDns* registry parameter)
- Static Hosts file (optional, depends on *EnableDns* registry parameter)
- DNS servers (optional, depends on *EnableDns* registry parameter)

NetBIOS name resolution order depends upon the node type and system configuration. The following node types are supported:

- *B-node* uses broadcasts for name registration and resolution.
- *P-node* uses a NetBIOS name server (such as WINS) for name registration and resolution.
- *M-node* uses broadcasts for name registration. For name resolution, it tries broadcasts first, but switches to p-node if it receives no answer.
- *H-node* uses a NetBIOS name server for both registration and resolution. However, if no name server can be located, it switches to b-node. It continues to poll for a name server and switches back to p-node when one becomes available.
- *Microsoft-enhanced* uses the local Lmhosts file or WINS proxies plus Windows Sockets gethostbyname() calls (using standard DNS and/or local Hosts files) in addition to standard node types.

Microsoft ships a NetBIOS name server known as the Windows Internet Name Service (WINS). Most WINS clients are set up as H-nodes; that is, they first attempt to register and resolve names using WINS, and if that fails, they try local subnet broadcasts. Using a WINS server to locate resources is generally preferable to broadcasting for two reasons:

- Routers do not usually forward broadcasts.
- All computers on a subnet process broadcasts.

**NetBIOS Name Registration and Resolution for Multihomed Computers**
As previously described, NetBT binds to only one IP address per physical network interface. From the NetBT viewpoint, a computer is multihomed only if it has more than one network interface card installed. When a name registration packet is sent from a multihomed machine, it is flagged as a multihomed name registration so that it does not conflict with the same name being registered by another interface in the same computer.

If a multihomed machine receives a broadcast name query, all NetBT-interface bindings receiving the query respond with their addresses, and by default the client chooses the first response and connects to the address supplied by it. This behavior can be controlled by the *RandomAdapter* registry parameter described in Appendix B.

When a directed name query is sent to a WINS server, the WINS server responds with a list of all IP addresses that were registered with WINS by the multihomed computer.

Choosing the best IP address to connect to on a multihomed computer is a client function. Currently, the following algorithm is employed, in the order listed:

1. If one of the IP addresses in the name query response list is on the same logical subnet as the calling binding of NetBT on the local computer, that address is selected. If more than one of the addresses meets the criteria, one is picked at random

from those that match.

2. If one of the IP addresses in the list is on the same (classless) network as the calling binding of NetBT on the local computer, that address is selected. If more than one of the addresses meets the criteria, one is picked at random from those that match.

3. If one of the IP addresses in the list is on the same logical subnet as any binding of NetBT on the local computer, that address is selected. If more than one of the addresses meets the criteria, one is picked at random from those.

4. If none of the IP addresses in the list is on the same subnet as any binding of NetBT on the local computer, an address is selected at random from the list.

This algorithm provides a reasonably good way of balancing connections to a server across multiple network adapters, and still favoring direct (same subnet) connections when they are available. When a list of IP addresses is returned, they are sorted into the best order, and NetBT attempts to ping each of the addresses in the list until one responds. NetBT then attempts a connection to that address. If no addresses respond, a connection attempt is made to the first address in the list anyway. This is tried in case there is a firewall or other device filtering ICMP traffic. Windows Server 2003 supports per interface NetBT name caching, and the **nbtstat -c** command displays the name cache on a per-interface basis.

**NetBT Internet/DNS Enhancements and the SMB Device**

It has always been possible to connect from one Windows-based computer to another using NetBT over the Internet. To do so, some means of name resolution had to be provided. Two common methods were to use the Lmhosts file or a WINS server. Several enhancements were introduced in Windows NT 4.0 and carried forward in Windows Server 2003 to eliminate these special configuration needs.

It is possible to connect to a NetBIOS over TCP/IP resource in the following ways:

- Use the command **net use ***ip address******share_name*. This eliminates the need for NetBIOS name resolution configuration.

- Use the command **net use ***FQDN******share_name*. This allows the use of a DNS to connect to a computer using its fully qualified domain name (FQDN).

Examples of using new functionality to map a drive to ftp.microsoft.com are shown here. The IP address listed here is subject to change.

- **net use f: \\ftp.microsoft.com\data**

- **net use \\131.107.232.1\data**

- **net view \\131.107.232.1**

- **dir \\ftp.microsoft.com\bussys\winnt**

In addition, various applications, such as the Event Viewer **Select Computer** option on the **Log** menu, allow you to enter an FQDN or IP address directly. In Windows Server 2003, it is also possible to use direct hosting to establish redirector or server connections between computers running Windows Server 2003 without the use of the NetBIOS namespace or mapping layer at all. By default, Windows attempts to make connections using both methods so that it can support connections to lower-level computers. However, in Windows XP and Windows Server 2003–only environments, you can disable NetBIOS completely for each network connection in Network Connections.

The interface in Windows Server 2003 that makes NetBIOS-less operation possible is termed the *Server Message Block (SMB) device*. It appears to the redirector and server as another interface, much as an individual network adapter/protocol stack combination does. At the TCP/IP stack however, the SMB device is bound to ADDR_ANY, and it uses the DNS namespace natively, like a Windows Sockets application. Calls placed on the SMB device will result in a standard DNS lookup to resolve the (DNS) name to an IP address, followed by a single outbound connection request (even on a multihomed computer) using the best source IP address and interface as determined by the route table. Additionally, there is no NetBIOS session setup on top of the TCP connection, as there is with traditional NetBIOS over TCP/IP. By default, the redirector places calls on both the NetBIOS device(s) and the SMB device, and the file server receives calls on both. The file server SMB device listens on TCP port 445

instead of the traditional NetBIOS over TCP port 139.

**NetBIOS over TCP Sessions**
NetBIOS sessions are established between two names. For example, when a Windows XP Professional-based workstation makes a file-sharing connection to a server using NetBIOS over TCP/IP, the following sequence of events takes place:

1. The NetBIOS name for the server is resolved to an IP address.

2. The next-hop address for the IP address of the server is resolved to a MAC address.

3. A TCP connection is established from the workstation to the server, using port 139.

4. The workstation sends a NetBIOS Session Request to the server name over the TCP connection. If the server is listening on that name, it responds affirmatively, and a session is established.

When the NetBIOS session has been established, the workstation and server negotiate which level of the SMB protocol to use. Microsoft networking uses only one NetBIOS session between two names at any time. Any additional file or print sharing connections are multiplexed over the same NetBIOS session using identifiers within the SMB header.

NetBIOS keep-alives are used on each connection to verify that both the server and workstation are still able to maintain their session. Therefore, if a workstation is shut down ungracefully, the server eventually cleans up the connection and associated resources, and vice versa. NetBIOS keep-alives are controlled by the *SessionKeepAlive* registry parameter and default to once per hour.

If Lmhosts files are used and an entry is misspelled, it is possible to attempt to connect to a server using the correct IP address but an incorrect name. In this case, a TCP connection is still established to the server. However, the server rejects the NetBIOS session request (using the wrong name), because there is no listen posted on that name. An Error 51, "Remote computer not listening," is returned.

**NetBIOS Datagram Services**
Datagrams are sent from one NetBIOS name to another over UDP port 138. The datagram service provides the ability to send a message to a unique name or to a group name. Group names may resolve to a list of IP addresses or a broadcast. For example, the command **net send /d:***mydomain* **test** sends a datagram containing the text "test" to the group name *mydomain*[03]. The *mydomain*[03] name resolves to an IP subnet broadcast, so the datagram is sent with the following characteristics:

- Destination MAC address: broadcast (0xFF-FF-FF-FF-FF-FF).

- Source MAC address: The MAC address of the sending interface.

- Destination IP address: The local subnet broadcast address.

- Source IP address: The IP address of the local computer.

- Destination name: *mydomain*[03] (the messenger service on the remote computers).

- Source name: *username*[03] (the messenger service on the local computer).

All hosts on the subnet pick up the datagram and process it, at least to the UDP protocol. On hosts that are running a NetBIOS datagram service, UDP hands the datagram to NetBT on port 138. NetBT checks the destination name to see if any application has posted a datagram receive on it and if so, passes the datagram up. If no receive is posted, the datagram is discarded.

If support for NetBIOS is disabled in Windows Server 2003 (as described earlier in this section), NetBIOS datagram services are not available.

# Critical Client Services and Stack Components

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The focus of this article is on core TCP/IP stack components, not on the many available services that use it. However, the stack itself relies upon a few services for configuration information and name and address resolution. A few of these critical client services are discussed here.

**Automatic Client Configuration and Media Sense**

One of the most important client services is the Dynamic Host Configuration Protocol (DHCP) client. The DHCP client has an expanded role in Windows Server 2003. Its primary new feature is the ability to automatically configure an IP address and subnet mask when the client is started on a small private network without a DHCP server available to assign addresses (such as a home network). Another new feature is support for *Media Sense*, which can improve the roaming experience for portable device users.

If the TCP/IP is configured to dynamically obtain TCP/IP protocol configuration information from a DHCP server (instead of being manually configured with an IP address and other parameters), the DHCP client service is engaged each time the computer is restarted. The DHCP client service now uses a two-step process to configure the client with an IP address and other configuration information.

When the client is installed, it attempts to locate a DHCP server and obtain a configuration from it. Many TCP/IP networks use DHCP servers that are administratively configured to allocate TCP/IP configuration information to clients on the network. If this attempt to locate a DHCP server fails, the Windows Server 2003 DHCP client uses one of the following:

1. If Automatic Private IP Addressing (APIPA) is selected, TCP/IP automatically selects an IP address from the IANA-reserved class B network 169.254.0.0 with the subnet mask 255.255.0.0[8]. The DHCP client tests (using a gratuitous ARP) to make sure that the IP address that it has chosen is not already in use. If it is in use, it selects another IP address (it does this for up to 10 addresses). Once the DHCP client has selected an address that is verifiably not in use, it configures the interface with this address. APIPA allows single subnet home office or small office networks to use TCP/IP without static configuration or the administration of a DHCP server. Note that APIPA does not configure a default gateway. Therefore, only local subnet traffic is possible.

2. If an alternate configuration is selected, TCP/IP is configured with the alternate configuration settings. Alternate configuration is useful when a computer is used on more than one network, at least one of the networks does not have a DHCP server, and an APIPA configuration is not wanted. For example, if you have a laptop computer that you use both at the office and at home, it is useful to configure TCP/IP for an alternate configuration. At the office, the laptop uses a DHCP-allocated TCP/IP configuration. At home, where there is no DHCP server present, the laptop automatically uses the alternate configuration, which provides easy access to home network devices and the Internet and allows seamless operation on both networks, without the manual reconfiguration of TCP/IP settings.

You can select whether to use APIPA or an alternate configuration from the **Alternate Configuration** tab from the properties of the TCP/IP protocol in the Network Connections folder.

Once configured for either APIPA or an alternate configuration, TCP/IP continues to check for a DHCP server in the background every 5 minutes. If a DHCP server is found, the APIPA or alternate configuration information is abandoned, and the configuration offered by the DHCP server is used instead.

If the DHCP client has previously obtained a lease from a DHCP server, the following modified sequence of events occurs:

1. If the client's lease is still valid (not expired) at boot time, the client tries to renew its lease with the DHCP server. If the client fails to locate a DHCP server during the renewal attempt, it tries to ping the default gateway that is listed in the lease. If pinging the default gateway succeeds, the DHCP client assumes that it is still located on the same network where it obtained its current lease and continues to use the lease. By default, the client attempts to renew its lease in the background when half of its assigned lease time has expired.

2. If the attempt to ping the default gateway fails, the client assumes that it has been moved to a network that has no DHCP services currently available (such as a home network), and autoconfigures itself as described above. Once autoconfigured, it continues to try to locate a DHCP server every 5 minutes, in the background.

*Media Sense* support was added in NDIS 5.0. It provides a mechanism for the network interface card to notify the protocol stack of media connect and media disconnect events. Windows Server 2003 TCP/IP utilizes these notifications to assist in automatic configuration. For instance, in Windows NT 4.0, when a portable computer was located and DHCP was configured on an Ethernet subnet, and then moved to another subnet without rebooting, the protocol stack received no indication of the move. This meant that the configuration parameters became stale, and not relevant to the new network. Additionally, if the computer was shut off, carried home and rebooted, the protocol stack was not aware that the network interface card was no longer connected to a network, and again stale configuration parameters remained. This could be problematic, as subnet routes, default gateways, and so on, could conflict with dial-up parameters.

Media Sense support allows the protocol stack to react to events and invalidate stale parameters. For instance, if a computer running Windows Server 2003 is unplugged from the network (assuming the network interface card supports Media Sense), after a damping period implemented in the stack (currently 3 seconds), TCP/IP will invalidate the parameters associated with the network which has been disconnected. The IP address(es) will no longer allow sends, and any routes associated with the interface are invalidated. You can make the network connection status when connected visible on the taskbar by selecting a connection, right-clicking it, clicking **Properties**, and then selecting the **Show icon in taskbar when connected** check box. The network connection icon will also appear automatically with a red "X" when the adapter is having a connectivity problem.

If an application is bound to a socket that is using an invalidated address, it should handle the event and recover in a graceful way, such as attempting to use another IP address on the system or notifying the user of the disconnect.

**DNS Dynamic Update Client**

Windows Server 2003 includes support for DNS dynamic updates as described in RFC 2136. Every time there is an address event (new address or renewal), the DHCP client sends option 81 and its fully qualified name to the DHCP server, and requests the DHCP server to register a DNS pointer (PTR) resource record (RR) on its behalf. The DHCP client handles the dynamic update of the address (A) RR on its own. This is done because only the client knows which IP addresses on the host map to that name. The DHCP server may not be able to properly do the A RR registration because it has incomplete knowledge. However, the DHCP server can be configured to instruct the client to allow the server to register both records with the DNS. Registry parameters associated with the DNS dynamic update client are documented in Appendix C.

The Windows Server 2003 DHCP server handles option 81 requests as specified in the Internet draft titled "The DHCP Client FQDN Option" (draft-ietf-dhc-fqdn-option-0x.txt). If a Windows Server 2003 DHCP client obtains an IP address configuration from a down-level DHCP server that does not handle option 81, it registers a PTR RR on its own. The Windows Server 2003 DNS server is capable of handling dynamic updates.

Statically configured (non-DHCP) clients register both the A and PTR RRs with the DNS server themselves.

**DNS Resolver Cache Service**

Windows Server 2003 includes a caching DNS resolver service, which is enabled by default. For troubleshooting purposes, this service can be viewed, stopped, and started like any other Windows service. The caching resolver reduces DNS network traffic and speeds name resolution by providing a local cache for DNS queries. Name query responses are cached for the TTL specified in the response (not to exceed the value specified in the *MaxCacheEntryTtlLimit* parameter), and future queries are answered from the cache, when possible. One interesting feature of the DNS Resolver Cache Service is that it supports negative caching. For example, if a query is made to a DNS server for a given host name and the response is negative, succeeding queries for the same name are answered (negatively) from the cache for *NegativeCacheTime* seconds (the default is 300). Another example of negative caching is that if all DNS servers are queried and none are available, for *NetFailureCacheTime* seconds (the default is 30) all succeeding name queries fail instantly, instead of timing out. This feature can save time for services that query the DNS during the boot process, especially when the client is booted from the network.

The DNS Resolver Cache Service has a number of other adjustable registry parameters, which are documented in Appendix C.

# Appendix A: TCP/IP Configuration Parameters

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The TCP/IP protocol suite implementation for Windows Server 2003 obtains all of its configuration data from the registry. This information is written to the registry by the Setup program. Some of this information is also supplied by the Dynamic Host Configuration Protocol (DHCP) client service, if it is enabled. This appendix defines all of the registry parameters used to configure the protocol driver, Tcpip.sys, which implements the standard TCP/IP network protocols.

The implementation of the protocol suite should perform properly and efficiently in most environments using only the configuration information gathered by Setup and DHCP. Optimal default values for all other configurable aspects of the protocols for most cases have been encoded into the drivers. Some customer installations may require changes to certain default values. To handle these cases, optional registry parameters can be created to modify the default behavior of some parts of the protocol drivers.

| **Note** |
| --- |
| The Windows TCP/IP implementation is largely self-tuning. Adjusting registry parameters may adversely affect system performance. |

All of the TCP/IP parameters are registry values located under the registry key

HKEY_LOCAL_MACHINE

\SYSTEM

\CurrentControlSet

\Services:

\Tcpip

\Parameters

Adapter-specific values are listed under subkeys for each adapter identified by the adapter's globally unique identifier (GUID).

To determine the GUID value for an adapter corresponding to a LAN connection in the Network Connections folder, do the following:

1. Open the Network Connections folder and note the name of the LAN connection, such as "Local Area Connection."

2. Click **Start**, click **Run**, type **regedit.exe**, and then click **OK**.

3. Use the tree view (the left pane) of the Registry Editor tool to open the following key:
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Network\{4D36E972-E325-11CE-BFC1-08002BE10318}

4. Under this key are one or more keys for the globally unique identifiers (GUIDs) corresponding to the installed LAN connections. Each of these GUID keys has a Connection subkey. Open each of the *GUID*\Connection keys and look for the Name setting in the contents pane whose value matches the name of your LAN connection from step 1.

5. When you have found the *GUID*\Connection key that contains the Name setting that matches the name of your LAN connection, write down or otherwise note the GUID value.

Depending on whether the system or adapter is DHCP-configured or static override values are specified, parameters may have both DHCP and statically configured values. If any of these parameters are changed using the registry editor, a restart of the system is generally required for the change to take effect. A restart is usually not required if values are changed using the Network Connections folder.

**Parameters Configurable Using the Registry Editor**

The following parameters receive default values during the installation of the TCP/IP components. To modify any of these values, use the Registry Editor (**Regedit.exe**). A few of the parameters are visible in the registry by default, but most must be

created to modify the default behavior of the TCP/IP protocol driver. Parameters configurable from the user interface are listed separately.

**AllowUserRawAccess**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**Valid Range:** 0, 1 (False, True)

**Default:** 0 (False)

**Description:** This parameter controls access to raw sockets. If true, non-administrative users have access to raw sockets. By default, only administrators have access to raw sockets. For more information about raw sockets, see Windows Sockets 2 (http://go.microsoft.com/fwlink/?LinkID=31786).

**ArpAlwaysSourceRoute**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1, or not present (false, true, or not present)

**Default:** not present

**Description:** By default, the stack transmits ARP queries without source routing first and retries with source routing enabled if no reply is received. Setting this parameter to 0 causes all IP broadcasts to be sent without source routing. Setting this parameter to 1 forces TCP/IP to transmit all ARP queries with source routing enabled on Token Ring networks.

**ArpCacheLife**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Number of seconds

**Valid Range:** 0–0xFFFFFFFF

**Default:** In absence of an *ArpCacheLife* parameter, the defaults for ARP cache time-outs are a two-minute time-out on unused entries and a ten-minute time-out on used entries.

**Description:** See *ArpCacheMinReferencedLife*

**ArpCacheMinReferencedLife**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Number of seconds

**Valid Range:** 0–0xFFFFFFFF

**Default:** 600 seconds (10 minutes)

**Description:** *ArpCacheMinReferencedLife* controls the minimum time until a referenced ARP cache entry expires. This parameter can be used in combination with the *ArpCacheLife* parameter, as follows:

- If *ArpCacheLife* is greater than or equal to *ArpCacheMinReferencedLife*, referenced and unreferenced ARP cache entries expire in *ArpCacheLife* seconds.

- If *ArpCacheLife* is less than *ArpCacheMinReferencedLife*, unreferenced entries expire in *ArpCacheLife* seconds, and referenced entries expire in *ArpCacheMinReferencedLife* seconds.

Entries in the ARP cache are referenced each time that an outbound packet is sent to the IP address in the entry.

**ArpRetryCount**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Number

**Valid Range:** 0–3

**Default:** 3

**Description:** This parameter controls the number of times that the computer sends a gratuitous ARP for its own IP address(es) while initializing. Gratuitous ARPs are sent to ensure that the IP address is not already in use on the locally attached subnet. The value controls the actual number of ARPs sent, not the number of retries.

**ArpTRSingleRoute**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** Setting this parameter to 1 causes ARP broadcasts that are source-routed (Token Ring) to be sent as single-route broadcasts, instead of all-routes broadcasts.

**ArpUseEtherSNAP**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** Setting this parameter to 1 forces TCP/IP to transmit Ethernet packets using IEEE 802.3 SNAP encoding. By default, the stack transmits packets in Ethernet II format, also known as Ethernet DIX format. It always receives both formats.

**DatabasePath**
**Key:** Tcpip\Parameters

**Value Type:** REG_EXPAND_SZ—Character string

**ValidRange:** A valid file path

**Default:** %*SystemRoot*%\system32\drivers\etc

**Description:** This parameter specifies the path to the standard Internet database files (Hosts, Lmhosts, Network, Protocols, Services). It is used by the Windows Sockets interface.

**DefaultTTL**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Number of seconds/hops

**ValidRange:** 0–0xff (0–255 decimal)

**Default:** 128

**Description:** Specifies the default time-to-live (TTL) value set in the header of outgoing IP packets. The TTL determines the maximum amount of time that an IP packet may live in the network without reaching its destination. It is effectively a limit on the number of links on which an IP packet is allowed to travel before being discarded.

**DisableDHCPMediaSense**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter can be used to control DHCP Media Sense behavior. If set to 1, the DHCP client ignores Media Sense events from the interface. By default, Media Sense events trigger the DHCP client to take an action, such as attempting to obtain a lease (when a connect event occurs), or invalidating the interface and routes (when a disconnect event occurs).

**DisableIPSourceRouting**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**Valid Range:** 0, 1, 2

0 - forward all packets

1 - do not forward Source Routed packets

2 - drop all incoming Source Routed packets

**Default:** 1 (for Windows Server 2003 with no service packs installed), 2 (for Windows Server 2003 with Service Pack 1)

**Description:** IP source routing is a mechanism allowing the sender to determine the IP route that a datagram should take through the network. The Ping and Tracert tools have options to specify source routing.

**DisableMediaSenseEventLog**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter can be used to disable logging of DHCP Media Sense events. By default, Media Sense events (connection/disconnection from the network) are logged in the event log for troubleshooting purposes.

**DisableTaskOffload**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter instructs the TCP/IP stack to disable offloading of tasks to the network interface card for troubleshooting and test purposes.

**DisableUserTOSSetting**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** This parameter can be used to allow programs to manipulate the Type Of Service (TOS) bits in the header of outgoing IP packets. In Windows Server 2003, this defaults to **True**. In general, individual applications should not be allowed to manipulate TOS bits.

**DontAddDefaultGateway**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0

**Description:** When you make a remote access connection that uses TCP/IP, a new default route is added to the route table with a metric lower than all other existing default routes. You can disable the automatic adding of this new default route by setting this registry parameter to 1. You can also disable this behavior from the properties of the TCP/IP protocol for a remote access connection in Network Connections. After doing so, you may need to configure static routes for destinations that are reachable across the remote access connection.

**EnableAddrMaskReply**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter controls whether the computer responds to an ICMP address mask request.

**EnableBcastArpReply**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** This parameter controls whether the computer responds to an ARP request when the source Ethernet address in the ARP is not unicast. Network Load Balancing Service (NLBS) will not work properly if this value is set to 0.

**EnableDeadGWDetect**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** When this parameter is set to 1, TCP is allowed to perform dead gateway detection. With this feature enabled, TCP informs IP to change to a backup default gateway if a number of connections are experiencing difficulty. Backup default gateways are configured as advanced TCP/IP settings from the Network Connections folder. See the "Dead Gateway Detection" section in this article for details.

**EnableICMPRedirects**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD--BOOLEAN

**Valid Range:** 0, 1 (False, True)

**Default:** 1 (True)

**Recommendation:** 0 (False)

**Description:** This parameter controls whether Windows Server 2003 TCP/IP will update its route table in response to ICMP Redirect messages that are sent to it by network devices such as a routers.

**EnableFastRouteLookup**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** Fast route look-up is enabled if this flag is set. This can make route lookups faster at the expense of non-paged pool memory. This flag is used only if the computer runs Windows Server 2003 and falls into the medium or large class (in other words, contains at least 64 MB of memory). This parameter is created by the Routing and Remote Access service.

**EnableMulticastForwarding**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** The routing service uses this parameter to control whether or not IP multicasts are forwarded. This parameter is

created by the Routing and Remote Access service.

**EnablePMTUBHDetect**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** Setting this parameter to 1 (true) causes TCP to try to detect PMTU black hole routers while doing Path MTU Discovery. A PMTU black hole router does not return ICMP Destination Unreachable messages when it needs to fragment an IP datagram with the Don't Fragment bit set. TCP depends on receiving these messages to perform Path MTU Discovery. With this feature enabled, TCP tries to send segments without the Don't Fragment bit set if several retransmissions of a segment go unacknowledged[91]. If the segment is acknowledged as a result, the MSS is decreased and the Don't Fragment bit is set in future packets on the connection. Enabling PMTU black hole detection increases the maximum number of retransmissions that are performed for a given segment.

**EnablePMTUDiscovery**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** When this parameter is set to 1 (true) TCP attempts to discover the Maximum Transmission Unit (MTU), or largest packet size, over the path to a remote host. By discovering the Path MTU (PMTU) and limiting TCP segments to this size, TCP can eliminate fragmentation at routers along the path that connect networks with different MTUs. Fragmentation adversely affects TCP throughput and network congestion. Setting this parameter to 0 (not recommended) causes an MTU of 576 bytes to be used for all connections that are not to destinations on a locally attached subnet.

**FFPControlFlags**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** If this parameter is set to 1, Fast Forwarding Path (FFP) is enabled. If it is set to 0, TCP/IP instructs all FFP-capable adapters not to do any fast forwarding on this computer. FFP-capable network adapters can receive routing information from the stack and forward subsequent packets in hardware without passing them up to the stack. FFP parameters are located in the TCP/IP registry key, but are actually placed there by the Routing and Remote Access service.

**FFPFastForwardingCacheSize**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Number of bytes

**Valid Range:** 0–0xFFFFFFFF

**Default:** 100,000 bytes

**Description:** This is the maximum amount of memory that a driver that supports fast forwarding path (FFP) can allocate for its fast-forwarding cache if it uses system memory for its cache. If the device has its own memory for fast-forwarding cache, this value is ignored.

**GlobalMaxTcpWindowSize**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Number of bytes

**Valid Range:** 0–0x3FFFFFFF (1073741823 decimal; however, values greater than 64 KB can only be achieved when connecting to other systems that support RFC 1323 window scaling, which is discussed in the TCP section of this article.)

**Default:** This parameter does not exist by default.

**Description:** The *TcpWindowSize* parameter can be used to set the receive window on a per-interface basis. This parameter can be used to set a global limit for the TCP window size on a system-wide basis.

**IPAutoconfigurationAddress**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—String

**ValidRange:** A valid IP address

**Default:** None

**Description:** The DHCP client stores the IP address chosen using APIPA autoconfiguration here. This value should not be altered.

**IPAutoconfigurationEnabled**
**Key:** Tcpip\Parameters, Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** This parameter enables or disables IP autoconfiguration using APIPA. See the "Automatic Client Configuration and Media Sense" section of this article for details. This parameter can be set globally or per interface. If a per-interface value is present, it overrides the global value for that interface.

**IPAutoconfigurationMask**
**Key:** Tcpip\Parameters, Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—String

**ValidRange:** A valid IP subnet mask

**Default:** 255.255.0.0

**Description:** This parameter controls the subnet mask assigned to the client using APIPA autoconfiguration. See the "Automatic Client Configuration and Media Sense" section of this article for details. This parameter can be set globally or per interface. If a per-interface value is present, it overrides the global value for that interface.

**IPAutoconfigurationSeed**
**Key:** Tcpip\Parameters, Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—Number

**Valid Range:** 0–0xFFFF

**Default:** 0

**Description:** This parameter is used internally by the DHCP client and should not be modified.

**IPAutoconfigurationSubnet**
**Key:** Tcpip\Parameters, Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—String

**ValidRange:** A valid IP subnet

**Default:** 169.254.0.0

**Description:** This parameter controls the initial network ID used by APIPA autoconfiguration to pick an IP address for the client. See the "Automatic Client Configuration and Media Sense" section of this article for details. This parameter can be set globally or per interface. If a per-interface value is present, it overrides the global value for that interface.

**IGMPLevel**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Number

**Valid Range:** 0,1,2

**Default:** 2

**Description:** This parameter determines to what extent the system supports IP multicasting and participates in the Internet Group Management Protocol. At level 0, the system provides no multicast support. At level 1, the system can send IP multicast packets but cannot receive them. At level 2, the system can send IP multicast packets and fully participate in IGMP to receive multicast packets.

**IPEnableRouter**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** Setting this parameter to 1 (true) causes the system to forward IP packets that have a destination address that is not assigned to it.

**IPEnableRouterBackup**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** Setup writes the previous value of *IPEnableRouter* to this key. It should not be adjusted manually.

**KeepAliveInterval**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—time in milliseconds

**Valid Range:** 1–0xFFFFFFFF

**Default:** 1000 (one second)

**Description:** This parameter determines the interval between TCP keep-alive retransmissions until a response is received. Once a response is received, the delay until the next keep-alive transmission is again controlled by the value of *KeepAliveTime*. The connection is aborted after the number of retransmissions specified by *TcpMaxDataRetransmissions* have gone unanswered.

**KeepAliveTime**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—time in milliseconds

**Valid Range:** 1–0xFFFFFFFF

**Default:** 7,200,000 (two hours)

**Description:** The parameter controls how often TCP attempts to verify that an idle connection is still intact by sending a keep-alive packet. If the remote system is still reachable and functioning, it acknowledges the keep-alive transmission. Keep-alive packets are not sent by default. This feature may be enabled on a connection by an application.

**MaxForwardBufferMemory**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number of bytes

**ValidRange:** *network MTU*–0xFFFFFFFF

**Default:** 2097152 decimal (2 MB)

**Description:** This parameter limits the total amount of memory that IP can allocate to store packet data in the router packet queue. This value must be greater than or equal to the value of the *ForwardBufferMemory* parameter. See the description of *ForwardBufferMemory* for more details.

**MaxForwardPending**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—number of packets

**Valid Range:** 1–0xFFFFFFFF

**Default:** 0x1388 (5000 decimal)

**Description:** This parameter limits the number of packets that the IP forwarding engine can submit for transmission to a specific network interface at any time. Additional packets are queued in IP until outstanding transmissions on the interface complete. Most network adapters transmit packets very quickly, so the default value is sufficient. A single remote access interface, however, may multiplex many slow serial lines. Configuring a larger value for this type of interface may improve its performance. The appropriate value depends on the number of outgoing lines and their load characteristics.

**MaxFreeTcbs**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 0–0xFFFFFFFF

**Default:** The following default values are used (note that *small* is defined as a computer with less than 19 MB of RAM, *medium* is 19–63 MB of RAM, and *large* is 64 MB or more of RAM. Although this code still exists, nearly all computers are *large* now).

For Windows Server 2003:

- Small system—500
- Medium system—1000
- Large system—2000

For Windows XP:

- Small system—250
- Medium system—500
- Large system—1000

**Description:** This parameter controls the number of cached (pre-allocated) Transport Control Blocks (TCBs) that are available. A TCB is a data structure that is maintained for each TCP connection.

**MaxHashTableSize**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number (must be a power of 2)

**ValidRange:** 0x40–0x10000 (64-65536 decimal)

**Default:** 512

**Description:** This value should be set to a power of 2 (for example, 512, 1024, 2048, and so on.) If this value is not a power of 2, the system configures the hash table to the next power of 2 value (for example, a setting of 513 is rounded up to 1024.) This value controls how fast the system can find a TCB and should be increased if *MaxFreeTcbs* is increased from the default.

**MaxICMPHostRoutes**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**ValidRange:** 0–0x7FFFFFFE

**Default:** 0x3E8 (1000 in decimal) or 0x2710 (10000 in decimal) with the update to Windows Server 2003 SP1, which is available in article 898060 from the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=67906).

**Description:** This value restricts the number of host routes that can be added to the local IP route table by receiving ICMP Redirect messages. You should not change this value unless the computer needs to be able to add a large number of host routes by receiving ICMP Redirect messages.

**MaxNormLookupMemory**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**ValidRange:** Any **DWORD** (0xFFFFFFFF means no limit on memory.)

**Default:** The following default values are used (Small is defined as a computer with less than 19 MB of RAM, Medium is 19–63 MB of RAM, and Large is 64 MB or more of RAM. Although this code still exists, nearly all computers are Large now).

For Windows Server 2003:

- Small system—150,000 bytes, which accommodates 1000 routes
- Medium system—1,500,000 bytes, which accommodates 10,000 routes
- Large system—5,000,000 bytes, which accommodates 40,000 routes

For Windows XP:

- 150,000 bytes, which accommodates 1000 routes

**Description:** This parameter controls the maximum amount of memory that the system allows for the route table data and the routes themselves. It is designed to prevent memory exhaustion on the computer caused by adding large numbers of routes.

**MaxNumForwardPackets**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 1–0xFFFFFFFF

**Default:** 0xFFFFFFFF

**Description:** This parameter limits the total number of IP packet headers that can be allocated for the router packet queue. This value must be greater than or equal to the value of the *NumForwardPackets* parameter. See the description of *NumForwardPackets* for more details.

**MaxUserPort**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—maximum port number

**ValidRange:** 5000–65534 (decimal)

**Default:** 0x1388 (5000 decimal)

**Description:** This parameter controls the maximum port number used when an application requests any available user port from the system. Normally, short-lived ports are allocated in the range from 1024 through 5000. Setting this parameter to a value outside of the valid range causes the nearest valid value to be used (5000 or 65534).

**MTU**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—number

**ValidRange:** 88–*the MTU of the underlying network*

**Default:** 0xFFFFFFFF

**Description:** This parameter overrides the default Maximum Transmission Unit (MTU) for a network interface. The MTU is the maximum IP packet size, in bytes, that can be transmitted over the underlying network. For values larger than the default for the underlying network, the network default MTU is used. For values smaller than 88, the MTU of 88 is used.

| **Note** |
| --- |
| Windows Server 2003 TCP/IP uses PMTU detection by default and queries the network interface card driver to find out what local MTU is supported. Altering the MTU parameter is generally not necessary and may result in reduced performance. See the "Path Maximum Transmission Unit (PMTU) Discovery" section of this article for more details. |

**NumForwardPackets**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**ValidRange:** 1—*some reasonable value smaller than 0xFFFFFFFF*

**Default:** 0x32 (50 decimal)

**Description:** This parameter determines the number of IP packet headers that are allocated for the router packet queue. When all headers are in use, the system attempts to allocate more, up to the value configured for *MaxNumForwardPackets*. This value should be at least as large as the *ForwardBufferMemory* value divided by the maximum IP data size of the networks that are connected to the router. It should be no larger than the *ForwardBufferMemory* value divided by 256 because at least 256 bytes of forward buffer memory is used for each packet. The optimal number of forward packets for a given *ForwardBufferMemory* size depends on the type of traffic that is carried on the network and is somewhere between these two values. This parameter is ignored and no headers are allocated if routing is not enabled.

**NumTcbTablePartitions**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number of TCB table partitions

**Valid Range:** 1–0xFFFF

**Default:** 4

**Description:** This parameter controls the number of TCB table partitions. The TCB table can be portioned to improve scalability on multi-processor systems by reducing contention on the TCB table. This value should not be modified without a careful performance study. A suggested maximum value is (number of CPUs)(2.

**PerformRouterDiscovery**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD

**Valid Range:** 0, 1, 2

0 (disabled)

1 (enabled)

2 (enable only if DHCP sends the Perform Router Discovery option)

**Default:** 2, DHCP-controlled but off by default.

**Description:** This parameter controls whether Windows Server 2003 TCP/IP attempts to perform router discovery per RFC 1256 on a per-interface basis. See also *SolicitationAddressBcast*.

**PerformRouterDiscoveryBackup**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** none

**Description:** This parameter is used internally to keep a back-up copy of the *PerformRouterDiscovery* value. It should not be modified.

**PPTPTcpMaxDataRetransmissions**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number of times to retransmit a PPTP packet

**Valid Range:** 0–0xFF

**Default:** 5

**Description:** This parameter controls the number of times that a PPTP packet is retransmitted if it is not acknowledged. This parameter was added to allow retransmission of PPTP traffic to be configured separately from regular TCP traffic.

**ReservedPorts**
**Key:** Tcpip\Parameters

**Value Type:** REG_MULTI_SZ

**ValidRange:** *xxxx-yyyy* The string uses the format *xxxx-yyyy*. (port range)

**Default:** NA

**Description:** Allows ports to be reserved so that they are not used as part of the 1024 or greater range. This is useful for applications that want a specific port range.

**SackOpts**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** This parameter controls whether or not Selective Acknowledgment (SACK) support, as specified in RFC 2018, is enabled. SACK is described in more detail in the "Transmission Control Protocol (TCP)" section of this article.

**SolicitationAddressBcast**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter can be used to configure Windows to send router discovery messages as broadcasts instead of multicasts, as described in RFC 1256. By default, if router discovery is enabled, router discovery solicitations are sent to the all-routers multicast group (224.0.0.2). See also *PerformRouterDiscovery*.

**StrictARPUpdate**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD

**Valid Range:** 0-1

**Default:** 0

**Description:** Specifies whether TCP/IP in Windows Server 2003 SP1 will store in the ARP cache the MAC address of the last ARP reply received (StrictARPUpdate=0) or the MAC address of the first ARP reply received (StrictARPUpdate=1). With StrictARPUpdate set to 1, TCP/IP will not update the MAC address of an existing ARP cache entry if it receives additional unsolicited ARP replies.

**SynAttackProtect**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD

**Valid Range:** 0, 1

0 (no SYN attack protection)

1 (reduced retransmission retries and delayed RCE [route cache entry] creation if the *TcpMaxHalfOpen* and *TcpMaxHalfOpenRetried* settings are satisfied and a delayed indication to Winsock is made.)

| **Note** |
| --- |
| When the system finds itself under attack the following options on any socket can no longer be enabled: scalable windows (RFC 1323) and per adapter configured TCP parameters (Initial RTT, window size). This is because when protection is functioning the route cache entry is not queried before the SYN-ACK is sent and the Winsock options are not available at this stage of the connection. |

**Default:** 1 - enabled for Windows Server 2003 Service Pack 1, 0 -disabled for Windows Server 2003 with no service packs installed

**Recommendation:** 1

**Description:** SYN attack protection involves reducing the amount of retransmissions for the SYN-ACKS, which will reduce the time for which resources have to remain allocated. The allocation of route cache entry resources is delayed until a connection is made and the connection indication to AFD is delayed until the three-way handshake is completed. Note that the actions taken by the protection mechanism only occur if *TcpMaxHalfOpen* and *TcpMaxHalfOpenRetried* settings are exceeded.

**Tcp1323Opts**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number (flags)

**Valid Range:** 0, 1, 2, 3

0 (disable RFC 1323 options)

1 (window scaling enabled only)

2 (timestamps enabled only)

3 (both options enabled)

**Default:** No value. The default behavior is as follows: do not use the Timestamp and Window Scale options when initiating TCP connections but use them if the TCP peer that is initiating communication includes them in the SYN segment.

**Description:** This parameter controls the use of RFC 1323 Timestamp and Window Scale TCP options. Explicit settings for timestamps and window scaling are manipulated with flag bits. Bit 0 controls window scaling, and bit 1 controls timestamps.

**TcpAckFrequency**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—number

**Valid Range:** 0–255

**Default:** 2

**Description:** Specifies the number of ACKs that will be outstanding before the delayed ACK timer is ignored. Microsoft does not recommend changing this value from the default without careful study of the environment.

**TcpDelAckTicks**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—number

**Valid Range:** 2–6

**Default:** 2

**Description:** Specifies the number of 100-millisecond intervals to use for the delayed-ACK timer on a per-interface basis. By default, the delayed-ACK timer is 200 milliseconds. If you set this value to 0 or 1, the delayed-ACK time is 200 milliseconds. Microsoft does not recommend changing this value from the default without careful study of the environment.

**TcpInitialRTT**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—number

**Valid Range:** 0–0xFFFF

**Default:** 3

**Description:** This parameter controls the initial time-out in seconds used for a TCP connection request and initial data retransmission on a per-interface basis. Use caution when tuning with this parameter because exponential backoff is used. Setting this value to larger than 3 results in much longer time-outs to nonexistent addresses.

**TcpMaxConnectResponseRetransmissions**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 0–255

**Default:** 2

**Description:** This parameter controls the number of times that a SYN-ACK is retransmitted in response to a connection request if the SYN is not acknowledged. If this value is greater than or equal to 2, the stack employs SYN attack protection internally. If this value is less than 2, the stack does not read the registry values at all for SYN attack protection. See also *SynAttackProtect*, *TCPMaxPortsExhausted*, *TCPMaxHalfOpen*, and *TCPMaxHalfOpenRetried*.

**TcpMaxConnectRetransmissions**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**ValidRange:** 0–255 (decimal)

**Default:** 2

**Description:** This parameter determines the number of times that TCP retransmits a connect request (a SYN segment) before aborting the attempt. The retransmission time-out is doubled with each successive retransmission in a given connect attempt. The initial time-out is controlled by the *TcpInitialRtt* registry value.

**TcpMaxDataRetransmissions**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 0–0xFFFFFFFF

**Default:** 5

**Description:** This parameter controls the number of times that TCP retransmits an individual data segment (not connection request segments) before aborting the connection. The retransmission time-out is doubled with each successive retransmission on a connection. It is reset when responses resume. The Retransmission Timeout (RTO) value is dynamically adjusted, using the historical measured round-trip time (Smoothed Round Trip Time) on each connection. The starting RTO on a new connection is controlled by the *TcpInitialRtt* registry value.

**TcpMaxDupAcks**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 1–3

**Default:** 2

**Description:** This parameter determines the number of duplicate ACKs that must be received for the same sequence number of sent data before fast retransmit is triggered to resend the segment that has been dropped in transit. This mechanism is described in more detail in the “Transmission Control Protocol (TCP)” section of this article.

**TcpMaxSendFree**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 0–0xFFFF

**Default:** 5000

**Description:** This parameter controls the size limit of the TCP header table. On machines with large amounts of RAM increasing this setting can improve responsiveness during a SYN attack.

**TcpNumConnections**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 0–0xFFFFFE

**Default:** 0xFFFFFE

**Description:** This parameter limits the maximum number of connections that TCP can have open simultaneously.

**TcpTimedWaitDelay**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—time in seconds

**ValidRange:** 30-300 (decimal)

**Default:** 0xF0 (120 decimal)

**Description:** This parameter determines the length of time that a connection stays in the TIME_WAIT state when being closed. While a connection is in the TIME_WAIT state, the socket pair cannot be reused. This is also known as the 2MSL state because the value should be twice the maximum segment lifetime on the network. See RFC 793 for further details.

**TcpUseRFC1122UrgentPointer**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter determines whether TCP uses the RFC 1122 or RFC 793 specification for urgent data (used by BSD-derived systems). There are two ways to interpret the value of the Urgent Pointer field in the TCP header: RFC 793 defines the value as indicating the first byte of normal data, RFC 1122 defines the value as indicating the last byte of urgent data. These two interpretations are not interoperable. Windows Server 2003 TCP/IP defaults to the RFC 793 interpretation (BSD mode).

**TcpWindowSize**
**Key:** Tcpip\Parameters, Tcpip\Parameters\Interface*interfaceGUID*

**Value Type:** REG_DWORD—number of bytes

**ValidRange:** 0–0x3FFFFFFF (1073741823 decimal). In practice the TCP/IP stack will round the number set to the nearest multiple of maximum segment size (MSS). Values greater than 64 KB can be achieved only when connecting to other systems that support RFC 1323 Window Scaling, which is discussed in the "Transmission Control Protocol (TCP)" section of this article.

**Default:** The smaller of the following values:

- 0xFFFF
- *GlobalMaxTcpWindowSize* (another registry parameter)
- The larger of four times the MSS
- 16384 rounded up to an even multiple of the MSS

The stack also tunes itself based on the media speed:

- Below 1 Mbps: 8 KB
- 1 Mbps – 100 Mbps: 17 KB
- Greater than 100 Mbps: 64 KB

The default can start at 17520 for Ethernet, but may shrink slightly when the connection is established to another computer that supports extended TCP header options, such as Selective Acknowledgements (SACK) and TCP Timestamps, because these options increase the size of the TCP header beyond the usual 20 bytes, leaving slightly less room for data.

**Description:** This parameter determines the maximum TCP receive window size offered. The receive window specifies the number of bytes that a sender can transmit without receiving an acknowledgment. In general, larger receive windows improve performance over high-delay, high-bandwidth networks. For greatest efficiency, the receive window should be an even multiple of the TCP Maximum Segment Size (MSS). This parameter is both a per-interface parameter and a global parameter, depending upon where the registry key is located. If there is a value for a specific interface, that value overrides the system-wide value. See also *GobalMaxTcpWindowSize*.

**TrFunctionalMcastAddress**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** This parameter determines whether IP multicasts are sent using the Token Ring functional address of 0xC0-00-00-04-00-00 (=1) or the MAC-level broadcast address of 0xFF-FF-FF-FF-FF-FF (=0). For more information, see RFC 1469.

**TypeOfInterface**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD

**Valid Range:** 0, 1, 2, 3

**Default:** 0 (allow multicast and unicast)

**Description:** This parameter determines whether the interface gets routes plumbed for unicast, multicast, or both traffic types, and whether those traffic types can be forwarded. If it is set to 0, both unicast and multicast traffic are allowed. If it is set to 1, unicast traffic is disabled. If it is set to 2, multicast traffic is disabled. If it set to 3, both unicast and multicast traffic are disabled. Since this parameter affects forwarding and routes, it may still be possible for a local application to send multicasts out over an interface, if there are no other interfaces in the computer that are enabled for multicast, and a default route exists.

**UseZeroBroadcast**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** If this parameter is set to 1 (true), IP will use the all 0s address for the limited broadcast address (0.0.0.0) instead of the all 1s address (255.255.255.255). Most systems use the all 1s broadcasts, but some systems derived from BSD implementations use the all 0s broadcasts. Systems that use different broadcasts do not interoperate well on the same network.

**Parameters Configurable from the User Interface**
The following parameters are created and modified automatically by configuring TCP/IP properties from Network Connections. There should be no need to configure them directly in the registry.

**DefaultGateway**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—list of dotted decimal IP addresses

**ValidRange:** Any set of valid IP addresses

**Default:** None

**Description:** This parameter specifies the list of gateways to be used to route packets that are not destined for a subnet that the computer is directly connected to, and for which a more specific route does not exist. This parameter, if it has a valid value, overrides the *DhcpDefaultGateway* parameter. There is only one active default gateway for the computer at any time, so adding multiple addresses is only done for redundancy. See the "Dead Gateway Detection" section in this article for details.

**Domain**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—character string

**ValidRange:** Any valid DNS domain name

**Default:** None

**Description:** This parameter specifies the DNS domain name of the interface. In Windows Server 2003, this and NameServer are per-interface parameters, rather than system-wide parameters. This parameter overrides the *DhcpDomain* parameter (filled in by the DHCP client), if it exists.

**EnableDhcp**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** If this parameter is set to 1 (true), the DHCP client service attempts to use DHCP to configure the first IP interface on this adapter.

**EnableSecurityFilters**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** If this parameter is set to 1 (true), IP security filters are enabled. See *TcpAllowedPorts*, *UdpAllowedPorts*, and *RawIPAllowedPorts*. To configure these values, on the **Start** menu, point to **Settings**, then click **Network Connections**, right-click **Local Area Connection**, and then click **Properties**. Select **Internet Protocol (TCP/IP)**, and click **Properties**, then click **Advanced**. Click the **Options** tab, select **TCP/IP filtering**, and click **Properties**.

**Hostname**
**Key:** Tcpip\Parameters

**Value Type:** REG_SZ—character string

**ValidRange:** Any valid DNS hostname

**Default:** The computer name of the system

**Description:** This parameter specifies the DNS host name of the system, which is returned by the **hostname** command.

**IPAddress**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—list of dotted-decimal IP addresses

**ValidRange:** Any set of valid IP addresses

**Default:** None

**Description:** This parameter specifies the IP addresses of the IP interfaces to be bound to the adapter. If the first address in the list is 0.0.0.0, the primary interface on the adapter is configured from DHCP. A system with more than one IP interface for an

adapter is *logically multihomed*. There must be a valid subnet mask value in the *SubnetMask* parameter for each IP address that is specified in this parameter. To add parameters with **Regedit.exe**, select this key and type the list of IP addresses, pressing **Enter** after each one. Then modify the *SubnetMask* value, and type a corresponding list of subnet masks.

**NameServer**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—a space delimited list of dotted decimal IP addresses

**ValidRange:** Any set of valid IP address

**Default:** None (blank)

**Description:** This parameter specifies the DNS name servers that Windows Sockets queries to resolve names. In Windows Server 2003, this and the *DomainName* are per-interface settings.

**RawIpAllowedProtocols**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—list of IP protocol numbers

**ValidRange:** Any set of valid IP protocol numbers

**Default:** None

**Description:** This parameter specifies the list of IP protocol numbers for which incoming datagrams are accepted on an IP interface when security filtering is enabled (*EnableSecurityFilters* = 1). The parameter controls the acceptance of IP datagrams by the raw IP transport, which is used to provide raw sockets. It does not control IP datagrams that are passed to other transports (for example, TCP). An empty list indicates that no values are acceptable. A single value of 0 indicates that all values are acceptable. The behavior of a list containing the value 0 mixed with other, nonzero values is undefined. If this parameter is missing from an interface, all values are acceptable. This parameter applies to all IP interfaces that are configured on a specific adapter.

**SearchList**
**Key:** Tcpip\Parameters

**Value Type:** REG_SZ—space delimited list of DNS domain name suffixes

**Valid Range:** 1-50

**Default:** None

**Description:** This parameter specifies a list of domain name suffixes to append to a name to be resolved through DNS if resolution of the unadorned name fails. By default, only the value of the *Domain* parameter is appended. This parameter is used by the Windows Sockets interface. See also the *AllowUnqualifiedQuery* parameter.

**SubnetMask**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—list of dotted decimal subnet masks

**ValidRange:** Any set of valid subnet masks.

**Default:** None

**Description:** This parameter specifies the subnet masks to be used with the IP interfaces bound to the adapter. If the first mask in the list is 0.0.0.0, the primary interface on the adapter is configured using DHCP. There must be a valid subnet mask value in this parameter for each IP address specified in the *IPAddress* parameter.

**TcpAllowedPorts**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—list of TCP port numbers

**ValidRange:** Any set of valid TCP port numbers

**Default:** None

**Description:** This parameter specifies the list of TCP port numbers for which incoming SYNs are accepted on an IP interface

when security filtering is enabled (*EnableSecurityFilters* = 1). An empty list indicates that no values are acceptable. A single value of 0 indicates that all values are acceptable. The behavior of a list containing the value 0 mixed with other, nonzero values is undefined. If this parameter is missing from an interface, all values are acceptable. This parameter applies to all IP interfaces configured on a specified adapter.

**UdpAllowedPorts**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—list of UDP port numbers

**ValidRange:** Any set of valid UDP port numbers

**Default:** None

**Description:** This parameter specifies the list of UDP port numbers for which incoming datagrams are accepted on an IP interface when security filtering is enabled (*EnableSecurityFilters* = 1). An empty list indicates that no values are acceptable. A single value of 0 indicates that all values are acceptable. The behavior of a list containing the value 0 mixed with other, nonzero values is undefined. If this parameter is missing from an interface, all values are acceptable. This parameter applies to all IP interfaces configured on a specified adapter.

**Parameters Configurable Using the Route Command**
The **route** command can store persistent IP routes as values under the Tcpip\Parameters\PersistentRoutes registry key. Each route is stored in the value name string as a comma-delimited list of the form:

destination,subnet mask,gateway,metric

For example, the command:

route add 10.99.100.0 MASK 255.255.255.0 10.99.99.1 METRIC 1 /p

produces the registry value:

10.99.100.0,255.255.255.0,10.99.99.1,1

The value type is a REG_SZ. There is no value data (empty string). Addition and deletion of these values can be accomplished using the route command. There should be no need to configure them directly.

**Non-Configurable Parameters**
The following parameters are created and used internally by the TCP/IP components. They should never be modified using the Registry Editor. They are listed here for reference only.

**DhcpDefaultGateway**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—list of dotted decimal IP addresses

**ValidRange:** Any set of valid IP addresses

**Default:** None

**Description:** This parameter specifies the list of default gateways to be used to route packets that are not destined for a subnet to which the computer is directly connected and for which a more specific route does not exist. This parameter is written by the DHCP client service, if enabled. This parameter is overridden by a valid *DefaultGateway* parameter value. Although this parameter is set on a per-interface basis, there is always only one default gateway active for the computer. Additional entries are treated as alternatives if the first one is down.

**DhcpIPAddress**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—dotted decimal IP address

**ValidRange:** Any valid IP address

**Default:** None

**Description:** This parameter specifies the DHCP-configured IP address for the interface. If the *IPAddress* parameter contains a first value other than 0.0.0.0, that value overrides this parameter.

**DhcpDomain**

**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—Character string

**ValidRange:** Any valid DNS domain name

**Default:** None (provided by DHCP server)

**Description:** This parameter specifies the DNS domain name of the interface. In Windows Server 2003, this and NameServer are now per-interface parameters, rather than system-wide parameters. If the *Domain* key exists, it overrides the *DhcpDomain* value.

**DhcpNameServer**
**Key:** Tcpip\Parameters

**Value Type:** REG_SZ—A space delimited list of dotted decimal IP addresses

**ValidRange:** Any set of valid IP address

**Default:** None

**Description:** This parameter specifies the DNS name servers to be queried by Windows Sockets to resolve names. It is written by the DHCP client service, if enabled. If the *NameServer* parameter has a valid value, it overrides this parameter.

**DhcpServer**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—dotted decimal IP address

**ValidRange:** Any valid IP address

**Default:** None

**Description:** This parameter specifies the IP address of the DHCP server that granted the lease on the IP address in the *DhcpIPAddress* parameter.

**DhcpSubnetMask**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—dotted decimal IP subnet mask

**ValidRange:** Any subnet mask that is valid for the configured IP address

**Default:** None

**Description:** This parameter specifies the DHCP-configured subnet mask for the address specified in the *DhcpIPAddress* parameter.

**DhcpSubnetMaskOpt**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_SZ—dotted decimal IP subnet mask

**ValidRange:** Any subnet mask that is valid for the configured IP address

**Default:** None

**Description:** This parameter is filled in by the DHCP client service and is used to build the *DhcpSubnetMask* parameter, which the stack actually uses. Validity checks are performed before the value is inserted into the *DhcpSubnetMask* parameter.

**Lease**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—time in seconds

**Valid Range:** 1–0xFFFFFFFF

**Default:** None

**Description:** The DHCP client service uses this parameter to store the time, in seconds, for which the lease on the IP address for this adapter is valid.

**LeaseObtainedTime**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—absolute time, in seconds, since midnight of 1/1/70

**Valid Range:** 1–0xFFFFFFFF

**Default:** None

**Description:** The DHCP client service uses this parameter to store the time at which the lease on the IP address for this adapter was obtained.

**LeaseTerminatesTime**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—absolute time, in seconds, since midnight of 1/1/70

**Valid Range:** 1–0xFFFFFFFF

**Default:** None

**Description:** The DHCP client service uses this parameter to store the time at which the lease on the IP address for this adapter expires.

**LLInterface**
**Key:** Tcpip\Parameters\Adapters*interfaceGUID*

**Value Type:** REG_SZ—Windows Server 2003 device name

**ValidRange:** A legal Windows Server 2003 device name

**Default:** Empty string (blank)

**Description:** This parameter is used to direct IP to bind to a different link-layer protocol than the built-in ARP module. The value of the parameter is the name of the Windows Server 2003 device to which IP should bind. This parameter is used in conjunction with the RAS component, for example. It is only present when ARP modules other than LAN bind to IP.

**NTEContextList**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—number

**Valid Range:** 0–0xFFFF

**Default:** none

**Description:** This parameter identifies the context of the IP address associated with an interface. Each IP address associated with an interface has its own context number. The values are used internally to identify an IP address and should not be altered.

**T1**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—absolute time, in seconds, since midnight of 1/1/70

**Valid Range:** 1–0xFFFFFFFF

**Default:** None

**Description:** The DHCP client service uses this parameter to store the time at which the service first tries to renew the lease on the IP address for the adapter by contacting the server that granted the lease.

**T2**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—absolute time, in seconds, since midnight of 1/1/70

**Valid Range:** 1–0xFFFFFFFF

**Default:** None

**Description:** The DHCP client service uses this parameter to store the time at which the service tries to renew the lease on the

IP address for the adapter by broadcasting a renewal request. Time T2 should only be reached if the service is unable to renew the lease with the original server for some reason.

# Appendix B: NetBIOS over TCP/IP Configuration Parameters

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

All of the NetBIOS over TCP/IP (NetBT) parameters are registry values located under one of two different subkeys of **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services**:

- NetBT\Parameters
- NetBT\Adapters\Interfaces*interface*, in which *interface* refers to the subkey for a network interface to which NetBT is bound

Values under the latter key(s) are specific to each interface. If the system is configured using DHCP, a change in parameters takes effect if you issue the command **ipconfig /renew** from a command prompt. Otherwise, you must reboot the system for a change in any of these parameters to take effect.

**Parameters Configurable Using the Registry Editor**

The following parameters are installed with default values by Control Panel-Network Connections during the installation of the TCP/IP components. They may be modified using the Registry Editor (Regedit.exe). A few of the parameters are visible in the registry by default, but most must be created in order to modify the default behavior of the NetBT driver.

**BacklogIncrement**

**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**ValidRange:** 3–0x14 (1–20 decimal)

**Default:** 3

**Description:** This parameter was added in response to Internet SYN-ATTACK issues. When a connection attempt is made to the NetBIOS TCP port (139), if the number of free connection blocks is below 2, a *BackLogIncrement* number of new connection blocks are created by the system. Each connection block consumes 78 bytes of memory. A limit on the total number of connection blocks allowed can be set using the *MaxConnBackLog* parameter. One connection block is required for each NetBT connection.

**BcastNameQueryCount**

**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 1–0xFFFF

**Default:** 3

**Description:** This value determines the number of times NetBT broadcasts a query for a specific name without receiving a response.

**BcastQueryTimeout**

**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—time in milliseconds

**Valid Range:** 100–0xFFFFFFFF

**Default:** 0x2ee (750 decimal)

**Description:** This value determines the time interval between successive broadcast name queries for the same name.

**BroadcastAddress**

**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—4-byte, little-endian encoded IP address

**Valid Range:** 0–0xFFFFFFFF

**Default:** The 1s-broadcast address for each network

**Description:** This parameter can be used to force NetBT to use a specific address for all broadcast name-related packets. By default, NetBT uses the 1s-broadcast address appropriate for each net (that is, for a network of 10.101.0.0 with a subnet mask of 255.255.0.0, the subnet broadcast address would be 10.101.255.255). This parameter would be set, for example, if the network uses the 0s-broadcast address (set using the *UseZeroBroadcast* TCP/IP parameter). The appropriate subnet broadcast address would then be 10.101.0.0 in the example above. This parameter would then be set to 0x0b650000. This parameter is global and is used on all subnets to which NetBT is bound.

**CachePerAdapterEnabled**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** This value determines whether NetBIOS remote name caching is done on a per-adapter basis. Nbtstat -c has been enhanced to show the per-adapter name cache.

**CacheTimeout**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—time, in milliseconds

**Valid Range:** 0xEA60–0xFFFFFFFF

**Default:** 0x927c0 (600000 milliseconds = 10 minutes)

**Description:** This value determines the time interval that names are cached in the remote name table. The **nbtstat –c** command can be used to view the remaining time for each name in the cache.

**ConnectOnRequestedInterfaceOnly**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This value can be used to allow NetBT connections on the requested adapter only. When the redirector on a multihomed computer calls another *computername*, it places calls on all NetBT transports (protocol/adapter combinations) to which it is bound. Each transport independently attempts to reach the target name. Setting this parameter limits each transport to connecting to other computers that are reachable via its own adapter, preventing crossover traffic. For more details, see the "NetBIOS Name Registration and Resolution for Multihomed Computers" section of this article.

It no longer works no wonder it doesn't make sense.

**EnableDns**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** If this value is set to 1 (true), NetBT queries the DNS server for names that cannot be resolved by WINS, broadcast, or the Lmhosts file.

**EnableProxyRegCheck**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** If this parameter is set to 1 (true), the proxy name server sends a negative response to a broadcast name registration if the name is already registered with WINS or is in the proxy's local name cache with a different IP address. This feature prevents a system from changing its IP address as long as WINS has a mapping for the name. For this reason, it is disabled by default.

**InitialRefreshT.O.**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—time, in milliseconds

**Valid Range:** 960000–0xFFFFFFF

**Default:** 960000 (16 minutes)

**Description:** This parameter specifies the initial refresh time-out used by NetBT during name registration. NetBT tries to contact the WINS servers at one-eighth of this time interval when it is first registering names. When it receives a successful registration response, that response contains the new refresh interval to use.

**LmhostsTimeout**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—time in milliseconds

**Valid Range:** 1000–0xFFFFFFFF

**Default:** 6000 (6 seconds)

**Description:** This parameter specifies the time-out value for Lmhosts and DNS name queries submitted by NetBT. The actual time-out could be as much as twice the value.

**MaxConnBackLog**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**ValidRange:** 2–0x9c40 (1-40,000 decimal)

**Default:** 1000

**Description:** This value determines the maximum number of connection blocks that NetBT allocates. See the *BackLogIncrement* parameter for more details.

**MaxPreloadEntries**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**ValidRange:** 0x3E8–0x7D0 (1000–2000 decimal)

**Default:** 1000 decimal

**Description:** This value determines the maximum number of entries that are preloaded from the Lmhosts file. Entries to preload into the cache are flagged in the Lmhosts file with the #PRE tag.

**MaxDgramBuffering**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number of bytes

**Valid Range:** 0x20000–0xFFFFFFFF

**Default:** 0x20000 (128K)

**Description:** This parameter specifies the maximum amount of memory that NetBT dynamically allocates for all outstanding datagram sends. Once this limit is reached, further sends fail due to insufficient resources.

**MinimumRefreshSleepTime**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 21600000-4294967295

**Default:** 21600000 ms (6 hours)

**Description:** This parameter is used to reset the TTL on the WakeupTimer if ½ of the TTL is less than 6 hours when the machine is put into sleep or hiberate mode.

**MinimumFreeLowerConnections**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 20-500

**Default:** 50

**Description:** This parameter is used allocate the number of free handles that the system has upon boot to accept incoming connections. These handles are allocated in addition to the number of active connections that are being serviced. Once the machine is in a steady state the number of free handles increases to ½ the value of the used handles. The number of free handles is never less than 50 unless specified in the registry.

**NameServerPort**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—UDP port number

**Valid Range:** 0–0xFFFF

**Default:** 0x89

**Description:** This parameter determines the destination port number to which NetBT sends name service-related packets, such as name queries and name registrations, to WINS. The Microsoft WINS Server listens on port 0x89 (138 decimal). NetBIOS name servers from other vendors may listen on different ports.

**NameSrvQueryCount**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 0–0xFFFF

**Default:** 3

**Description:** This value determines the number of times that NetBT sends a query to a WINS server for a specified name without receiving a response.

**NameSrvQueryTimeout**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—time in milliseconds

**Valid Range:** 100–0xFFFFFFFF

**Default:** 1500 (1.5 seconds)

**Description:** This value determines the time interval between successive name queries to WINS for a specified name.

**NodeType**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**ValidRange:** 1, 2, 4, 8 (b-node, p-node, m-node, h-node)

**Default:** 1 or 8 based on the WINS server configuration

**Description:** This parameter determines what methods NetBT uses to register and resolve names. A b-node system uses broadcasts. A p-node system uses only point-to-point name queries to a name server (WINS). An m-node system broadcasts

first, and then queries the name server. An h-node system queries the name server first, and then broadcasts. Resolution through Lmhosts and DNS, if enabled, follows these methods. If this key is present, it overrides the DhcpNodeType key. If neither key is present, the system defaults to b-node if there are no WINS servers configured for the client. The system defaults to h-node if there is at least one WINS server configured.

**NoNameReleaseOnDemand**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter determines whether the computer releases its NetBIOS name when it receives a name-release request from the network. It was added to allow the administrator to protect the machine against malicious name-release attacks.

**RandomAdapter**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter applies to a multihomed host only. If it is set to 1 (true), NetBT randomly chooses the IP address to put in a name-query response from all of its bound interfaces. Usually, the response contains the address of the interface to which the query arrived. This feature would be used for load balancing by a server with two interfaces on the same network.

**RefreshOpCode**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 8, 9

**Default:** 8

**Description:** This parameter forces NetBT to use a specific opcode field in name-refresh packets. The specification for the NetBT protocol is somewhat ambiguous in this area. Although the default of 8 that is used by Microsoft implementations appears to be the intended value, some other implementations, such as those by Ungermann-Bass, use the value 9. Two implementations must use the same opcode field to interoperate.

**ScopeId**
**Key:** Netbt\Parameters

**Value Type:** REG_SZ—character string

**ValidRange:** Any valid DNS domain name consisting of two dot-separated parts or an asterisk (*).

**Default:** None

**Description:** This parameter specifies the NetBIOS name scope for the node. This value must not begin with a period. If this parameter contains a valid value, it overrides the DHCP parameter of the same name. A blank value (empty string) is ignored. Setting this parameter to the value "*" indicates a null scope and overrides the DHCP parameter.

**SessionKeepAlive**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—time in milliseconds

**Valid Range:** 60,000–0xFFFFFFFF

**Default:** 3,600,000 (1 hour)

**Description:** This value determines the time interval between keep-alive transmissions on a session. Setting the value to 0xFFFFFFF disables keep-alives.

**SingleResponse**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter applies to a multihomed host only. If this parameter is set to 1 (true), NetBT supplies only the IP address from one of its bound interfaces in name-query responses. By default, the addresses of all bound interfaces are included.

**Size/Small/Medium/Large**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD

**ValidRange:** 1, 2, 3 (small, medium, large)

**Default:** 1 (small)

**Description:** This value determines the size of the name tables that are used to store local and remote names. In general, a setting of 1 (small) is adequate. If the system is acting as a proxy name server, the value is automatically set to 3 (large) to increase the size of the name cache hash table. Hash table buckets are sized as follows:

- Small: 16
- Medium: 128
- Large: 256

**SMBDeviceEnabled**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** Windows Server 2003 supports a new network transport known as the SMB Device, which is enabled by default. This parameter can be used to disable the SMB device for troubleshooting purposes. See the "NetBT Internet/DNS Enhancements and the SMB Device" section of this article for more details.

**TryAllNameServers**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter controls whether the client continues to query additional name servers from the list of configured servers when a NetBIOS session setup request to one of the IP addresses fails. If this parameter is enabled, attempts are made to query all the WINS servers in the list and connect to all the IP addresses supplied before failing the request to the user.

**TryAllIPAddrs**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** When a WINS server returns a list of IP addresses in response to a name query, they are sorted into a preference

order based on whether any of them are on the same subnet as an interface belonging to the client. This parameter controls whether the client pings the IP addresses in the list and attempts to connect to the first one that responds, or whether it tries to connect to the first IP address in the (sorted) list and fails if that connection attempt fails. By default, the client pings each address in the list and attempts to connect to the first one that answers the ping.

**UseDnsOnlyForNameResolutions**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter is used to disable all NetBIOS name queries. NetBIOS name registrations and refreshes are still used, and NetBIOS sessions are still allowed. To completely disable NetBIOS on an interface, see the *NetbiosOptions* parameter.

**WinsDownTimeout**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—time in milliseconds

**Valid Range:** 1000–0xFFFFFFFF

**Default:** 15,000 (15 seconds)

**Description:** This parameter determines the amount of time that NetBT waits before trying to use WINS again after it fails to contact any WINS server. This feature primarily allows computers that are temporarily disconnected from the network, such as laptops, to proceed through boot processing without waiting to time out each WINS name registration or query individually.

**Parameters Configurable from the Connections UI**
The following parameters can be set using the Control Panel-Network Connections. There should be no need to configure them directly.

**EnableLmhosts**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** If this value is set to 1 (true), NetBT searches the Lmhosts file, if it exists, for names that cannot be resolved by WINS or broadcast. By default, there is no Lmhosts file database directory (specified by Tcpip\Parameters\DatabasePath), so no action is taken. This value is written by the Advanced TCP/IP Configuration dialog box.

**EnableProxy**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** If this value is set to 1 (true), the system acts as a proxy name server for the networks to which NetBT is bound. A proxy name server answers broadcast queries for names that it has resolved through WINS. A proxy name server allows a network of b-node implementations to connect to servers on other subnets that are registered with WINS.

**NameServerList**
**Key:** Netbt\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—space separated, dotted decimal IP address (that is, 10.101.1.200)

**ValidRange:** any list of valid WINS server IP addresses.

**Default:** blank (no address)

**Description:** This parameter specifies the IP addresses of the list of WINS servers configured for the computer. If this

parameter contains a valid value, it overrides the DHCP parameter of the same name. This parameter replaces the Windows NT 4.0 parameters *NameServer* and *NameServerBackup*, which are no longer used.

**NetbiosOptions**
**Key:** Netbt\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—number

**Valid Range:** 1, 2

**Default:** 1

**Description:** This parameter controls whether NetBIOS is enabled on a per-interface basis. On the **Start** menu, point to **Settings**, and click **Network Connections**. Right-click **Local Area Connection**, and click **Properties**. Select **Internet Protocol (TCP\IP)**, and click **Properties**, then click **Advanced**. Click the **WINS** tab. The NetBIOS options are **Enable NetBIOS over TCP\IP**, **Disable NetBIOS over TCP\IP**, or **Use NetBIOS setting from the DHCP server**, the default. When enabled, the value is 1. When disabled, the value is set to 2. If this key does not exist, the *DHCPNetbiosOptions* key is read. If this key does exist, *DHCPNetbiosOptions* is ignored.

**Non-Configurable Parameters**
The following parameters are created and used internally by the NetBT components. They should never be modified using the Registry Edit or it can cause the component to become unstable. They are listed here for reference only.

**DHCPNameServerList**
**Key:** Netbt\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_MULTI_SZ—space separated, dotted decimal IP address (that is, 10.101.1.200)

**ValidRange:** any list of valid WINS server IP addresses.

**Default:** blank (no address)

**Description:** This parameter specifies the IP addresses of the list of WINS servers, as provided by the DHCP service. This parameter replaces the Windows NT 4.0 parameters *DHCPNameServer* and *DHCPNameServerBackup*, which are no longer used. See also *NameServerList*, which overrides this parameter if it is present.

**DHCPNetbiosOptions**
**Key:** Netbt\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD—number

**Valid Range:** 1, 2

**Default:** 1

**Description:** This parameter is written by the DHCP client service. See the *NetbiosOptions* parameter for a description.

**DhcpNodeType**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 1–8

**Default:** 1

**Description:** This parameter specifies the NetBT node type. It is written by the DHCP client service, if enabled. A valid NodeType value overrides this parameter. See the entry for *NodeType* for a complete description.

**DhcpScopeId**
**Key:** Netbt\Parameters

**Value Type:** REG_SZ—character string

**ValidRange:** a dot-separated name string such as microsoft.com

**Default:** none

**Description:** This parameter specifies the NetBIOS name scope for the node. It is written by the DHCP client service, if enabled. This value must not begin with a period. See the entry for *ScopeId* for more information.

**NbProvider**
**Key:** Netbt\Parameters

**Value Type:** REG_SZ—character string

**ValidRange:** _tcp

**Default:** _tcp

**Description:** This parameter is used internally by the RPC component. The default value should not be changed.

**TransportBindName**
**Key:** Netbt\Parameters

**Value Type:** REG_SZ—character string

**Valid Range:** N/A

**Default:** \Device\

**Description:** This parameter is used internally during product development. The default value should not be changed.

# Appendix C: Windows Sockets and DNS Registry Parameters

Updated: July 21, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AFD Registry Parameters**
Afd.sys is the kernel-mode driver that is used to support Windows Sockets applications. When there are three default values, the default is calculated based on the amount of memory detected in the system:

- The first value is the default for smaller computers (less than 19 MB).
- The second value is the default for medium computers (<32 MB on Windows XP Professional, <64 MB on Windows Server 2003).
- The third value is the default for large computers (>32 MB on Windows XP Professional, >64 MB on Windows Server 2003).

For example, if the default is given as 0/2/10, a system containing 12.5 to 20 MB of RAM would default to 2.

The following values can be set under:

HKEY_LOCAL_MACHINE

\SYSTEM

\CurrentControlSet

\Services

\Afd

\Parameters

**DefaultReceiveWindow**
**Value Type:** REG_DWORD

**Default:** 4096/8192/8192

**Description:** The number of receive bytes that AFD buffers on a connection before imposing flow control. For some applications, a larger value here gives slightly better performance at the expense of increased resource utilization. Applications can modify this value on a per-socket basis with the SO_RCVBUF socket option.

**DefaultSendWindow**
**Value Type:** REG_DWORD

**Default:** 4096/8192/8192

**Description:** This is similar to *DefaultReceiveWindow*, but for the send side of connections.

**DisableAddressSharing**
**Value Type:** REG_DWORD

**Default:** 0

**Range:** 0, 1

**Description:** This parameter is used to prevent address sharing (SO_REUSEADDR) between processes so that if a process opens a socket, no other process can steal data from it. A similar effect can be achieved if an application uses the new socket option SO_EXCLUSIVEADDRUSE. This setting allows administrators to secure older applications that are not aware of this option.

**FastCopyReceiveThreshold**
**Value Type:** REG_DWORD

**Default:** 1024

**Description:** When an application posts a receive with a buffer that is smaller than the current packet being buffered by Winsock, AFD can either make an additional copy of the packet and then copy data to the application buffers directly (which is a two-stage copy because application buffers cannot be accessed directly under the lock), or it can lock and map application buffers and copy data once. This value represents a compromise between extra code execution for data copying, and extra code execution in the I/O subsystem and memory manager. The default value was found, by testing, to be the best overall value for performance. Changing this value is not generally recommended.

**FastSendDatagramThreshold**
**Value Type:** REG_DWORD

**Default:** 1024

**Description:** Datagrams smaller than the value of this parameter go through the fast I/O path or are buffered on send. Larger ones are held until the datagram is actually sent. The default value was found by testing to be the best overall value for performance. Fast I/O means copying data and bypassing the I/O subsystem, instead of mapping memory and going through the I/O subsystem. This is advantageous for small amounts of data. Changing this value is not generally recommended.

**IgnorePushBitOnReceives**
**Value Type:** REG_DWORD—Boolean

ValidRange: 0, 1 (false, true)

**Default:** 0 (false)

**Description:** Normally, Windows Server 2003 completes a Windows Sockets **Receive** when one of the following occurs:

- Data arrives with the push bit set.
- The user **recv** buffer is full.
- 0.5 seconds have elapsed since any data arrived.

Setting this parameter to a 1 causes Afd.sys to treat all incoming packets as though the push bit was set. This should only be done when necessary to work around client TCP/IP implementations that are not properly pushing data.

**LargeBufferSize**
**Value Type:** REG_DWORD

**Default:** PAGE_SIZE (4096 bytes on i386, 8192 bytes on Alpha)

**Description:** The size, in bytes, of large buffers used by AFD. Smaller values use less memory and larger values can improve performance.

**LargeBufferListDepth**
**Value Type:** REG_DWORD

**Default:** 0/2/10

**Description:** Depth of large buffer look-aside list.

**MaxFastTransmit**
**Value Type:** REG_DWORD

**Valid Range:** 0–0xffffffff

**Default:** 64 KB

**Description:** This parameter controls the maximum amount of data that is transferred in a **TransmitFile** request on the fast path. Fast I/O is essentially copying data and bypassing the I/O subsystem, instead of mapping memory and going through the I/O subsystem. This is advantageous for small amounts of data. Changing this value is not generally recommended.

**MaxFastCopyTransmit**
**Value Type:** REG_DWORD

**Valid Range:** 0–0xFFFFFFFF

**Default:** 128

**Description:** This parameter controls the maximum size of data that uses copy instead of cached memory on the fast-path. Fast I/O is essentially copying data and bypassing the I/O subsystem, instead of mapping memory and going through the I/O subsystem. This is advantageous for small amounts of data. Changing this value is not generally recommended.

**MediumBufferSize**
**Value Type:** REG_DWORD

**Default:** 1504

**Description:** The size, in bytes, of medium buffers used by AFD.

**MediumBufferListDepth**
**Value Type:** REG_DWORD

**Default:** 4/8/24

**Description:** Depth of medium buffer look-aside list.

**OverheadChargeGranularity**
**Value Type:** REG_DWORD

**Default:** 1 page

**ValidRange:** A power of 2

**Description:** This parameter determines in what increments overhead is actually charged. The default is one page, and the intention is to properly charge and contain attacker type applications that try to run the system out of memory.

**PriorityBoost**
**Value Type:** REG_DWORD

**Default:** 2

**Valid Range:** 0–16

**Description:** The priority boost that AFD gives to a thread when it completes I/O for that thread. If a multithreaded application experiences starvation of some threads, the problem may be remedied by reducing this value.

**SmallBufferListDepth**
**Value Type:** REG_DWORD

**Default:** 8/16/32

**Description:** Depth of the small buffer look-aside list.

**SmallBufferSize**
**Value Type:** REG_DWORD

**Default:** 128

**Description:** The size in bytes of small buffers used by AFD.

**StandardAddressLength**
**Value Type:** REG_DWORD

**Default:** 22

**Description:** The length of TDI addresses that are typically used for the computer. When using an alternate transport protocol, such as TP4, which uses very long addresses, increasing this value results in a slight performance improvement.

**TransmitIoLength**
**Value Type:** REG_DWORD

**Default:** PAGE_SIZE/PAGE_SIZE*2/65536

**Description:** The default size for I/O (reads and sends) performed by TransmitFile(). For Windows XP Professional, the default I/O size is exactly one page.

**TransmitWorker**

**Value Type:** REG_DWORD

**Default:** 0x10

**Valid Range:** 0x10, 0x20

**Description:** This parameter controls how Afd.sys uses system threads. Setting it to 0x10 causes AFD to use system threads to perform IO that results from a long (more than 2 *SendPacketLength* worth of data) TransmitFile request. Setting it to 0x20 causes AFD to use kernel-mode APC for IO and to execute everything in the context of the same thread. This is new in Windows Server 2003 and can improve performance by reducing the number of context switches in long TransmitFile requests.

#### Dynamic Update Registration Parameters

These parameters control behavior of the dynamic update DNS registration client. If a parameter is not present, the default value listed is used.

##### DNSQueryTimeouts

**Key:** Tcpip\Parameters

**Value Type:** REG_MULTI_SZ—list of timeouts terminated by a zero

**ValidRange:** valid list of numbers

**Default:** 1 2 2 4 8 0 (format note: after entering each number hit return and terminate the list with zero)

**Description:** This parameter can be used to change the DNS query timeouts that the DNS client uses. In a controlled non-Internet or low-delay environment this could be used to decrease the time to failure of the query.

##### DefaultRegistrationTTL

**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—seconds

**Default:** 0x4B0 (1200 decimal, or 20 minutes)

**Valid Range:** 0–0xFFFFFFFF

**Description:** This parameter can be used to control the TTL value sent dynamically with DNS registrations.

##### EnableAdapterDomainNameRegistration

**Key:** Tcpip\Parameters\Interfaces*interface*

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter can be used to enable DNS dynamic update registration of a specific adapter's domain name information. This setting is useful when registrations of the adapter address(es) under the adapter's domain name are needed. When this key is set to true and *DisableDynamicUpdate* is false, the given adapter's address(es) is registered under the specific adapter's domain name and under the system's primary domain name.

##### DisableDynamicUpdate

**Key:** Tcpip\Parameters, Tcpip\Parameters\Interfaces*interface*

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false; dynamic update-enabled)

**Description:** This parameter can be used to completely disable DNS dynamic update registration. This parameter is both a per-interface parameter and a global parameter, depending upon where the registry key is located. If the value at the Tcpip\Parameters level is set to 1, dynamic update is disabled for the entire system. If the value at the Tcpip\Parameters level is set to 0, dynamic updates can be disabled on a per-adapter basis.

##### DisableReplaceAddressesInConflicts

**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter is used to turn off the address registration conflict rule that the last writer wins. By default, a computer does not replace any current records on the DNS server that do not appear to have been owned by it at one time.

**DisableReverseAddressRegistrations**
**Key:** Tcpip\Parameters\Interfaces*interface*

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false; registration of PTR records enabled)

**Description:** This parameter can be used to turn off DNS dynamic update reverse address (PTR) record registration. If the DHCP server that configures this computer is running Windows Server 2003, then it is capable of registering the PTR record with the DNS dynamic update protocol. However, if the DHCP server is not capable of performing DNS dynamic update PTR registrations and you do not want to register PTR records with the DNS dynamic update protocol, set this parameter to 1.

**UpdateSecurityLevel**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—flags

**Default:** 0

**Valid Range:** 0,0x00000010, 0x00000020, 0x00000100

**Description:** This parameter can be used to control the security that is used for DNS dynamic updates. It defaults to 0, to try nonsecure update, and if refused, to send Windows Server 2003 secure dynamic updates. Valid values are listed below:

- **0x00000000**—default, nonsecure updates
- **0x00000010**—security OFF
- **0x00000100**—secure ONLY ON

**DNS Caching Resolver Service Registry Parameters**
Windows Server 2003 includes a DNS caching resolver service. This service performs the function of caching DNR answers so that the DNS server does not need to be repeatedly queried for the same information. The service can be stopped using the Service Control Manager MMC snap-in. Registry parameters for this service are located under the **\System\CurrentControlSet\Services\Dnscache\Parameters** key.

**AdapterTimeoutCacheTime**
**Value Type:** REG_DWORD—seconds

**Valid Range:** 0–0xFFFFFFFF

**Default:** 300 (5 minutes)

**Description:** The amount of time that a particular adapter on a multihomed machine is disabled when a DNS query attempt fails (times out) for all of the given adapter's DNS servers. For instance, if you have two adapters and the DNS servers on one of the networks are unreachable, mark the adapter as unusable for this time period. (A Plug and Play event or cache time-out forces the resolver to retry this interface and mark it as disabled, if needed.)

**DefaultRegistrationRefreshInterval**
**Value Type:** REG_DWORD—time in seconds

**Default:** 0x15180 (86400 decimal, or 24 hours)

Range: 0–0xFFFFFFFF

**Description:** This parameter can be used to control the dynamic update DNS registration refresh interval.

**MaxCacheEntryTtlLimit**
**Value Type:** REG_DWORD—time in seconds

**Default:** 0x15180 (86400 decimal)

**ValidRange:** 0–0xFFFFFFFF (suggested value less than one day, to prevent very stale records)

**Description:** This parameter can be used to control the maximum cache entry time-to-live (TTL) value. It overrides any value that may have been set on a specific record that is larger.

**MaxSOACacheEntryTtlLimit**
**Value Type:** REG_DWORD—time, in seconds

**Valid Range:** 0–0xFFFFFFFF

**Default:** 120 (2 minutes)

**Description:** The maximum number of seconds that the resolver cache caches any SOA records. This value overrides any TTL value greater than itself for a specific SOA record that is returned from a DNS query. SOA records are essential for dynamic updates; therefore, they are not cached for long, to ensure that the most up-to-date record data is available for the DNS start of authority.

**NegativeCacheTime**
**Value Type:** REG_DWORD—time, in seconds

**Default:** 0x12c (300 decimal, or 5 minutes)

**ValidRange:** 0–0xFFFFFFFF (the suggested value is less one day, to prevent very stale records)

**Description:** This parameter can be used to control the cache time for negative records.

**NegativeSOACacheTime**
**Value Type:** REG_DWORD—time, in seconds

**Default:** 0x78 (120 decimal, or 2 minutes)

**ValidRange:** 0–0xFFFFFFFF (the suggested value is less than five minutes)

**Description:** This parameter can be used to control the cache time for negative Start of Authority (SOA) records. DNS registrations that fail are retried at five and ten minutes, so if this value is set to five minutes or more, retries are answered negatively from cache, instead of from the server, which could be available.

**NetFailureErrorPopupLimit**
**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Description:** This parameter enables the UI popup to indicate that the DNS resolver was unable to query (reach) the configured DNS servers for a repeated number of query attempts.

**NetFailureCacheTime**
**Value Type:** REG_DWORD—time, in seconds

**Default:** 0x1e (30 decimal)

**ValidRange:** 0–0xFFFFFFFF (suggested value is less than five minutes)

**Description:** This parameter is used to control the general network failure cache time. It prevents the resolver from querying for a period of time when it has been detected that a time-out error is occurring for queries against all known DNS servers. This avoids slowness (caused by time-outs) when the network does not respond.

**Name Resolution Parameters**
The following list of parameters is used by the Domain Name Resolver service.

**AllowUnqualifiedQuery**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0

**Description:** This parameter controls whether or not the Domain Name Resolver queries the Domain Name Server(s) with the host name, followed by a dot (.) only (an unqualified query). For example, if your computer is in mydomain.com and you ping *target*, by default the DNS is queried for *target.mydomain.com* only. When this parameter is set to 1, *target* is also queried.

**DisjointNameSpace**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1

**Description:** This parameter instructs the DNR to treat each interface as a disjoint name space. On a multihomed computer, a query to the DNS server(s) that is/are configured for one interface may result in a name error. This parameter is used to instruct the resolver to try the query against the possible DNS servers that are configured for other interfaces before returning results.

**PrioritizeRecordData**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1

**Description:** This parameter controls whether or not the Domain Name Resolver sorts the addresses that are returned in response to a query for a multihomed host. By default, the DNR sorts addresses that are on the same subnet as one of the interfaces in the querying computer to the top of the list. This is done to give preference to a common-subnet (non-routed) IP address, when possible.

**QueryIpMatching**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1

**Description:** This parameter controls whether or not the IP address of the DNS server queried is matched to the IP address of the server that sent the DNS response. This can be used as a primitive security feature to ensure that the resolver is not being fooled by a random query response from some computer other than the intended DNS server.

**UseDomainNameDevolution**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—binary

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Description:** This parameter can be used to disable domain name *devolution* for unqualified DNS queries. Devolution describes the process of attempting to locate a host in the DNS by first appending the domain suffix of the client to the host name, and then querying for the full string. If that query fails, one label is removed at a time, and the query is resubmitted. For example, if a user or application on the computer *mycomputer.support.microsoft.com* attempts to reach a host named *target*, the DNR by default tries *target.support.microsoft.com*, and *target.microsoft.com*, and possibly *target*, depending on the value of the *AllowUnqualifiedQuery* parameter.

# Appendix D: Tuning TCP/IP Response to Attack

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP Security Settings**
In addition to the settings that are listed above, the following keys can be altered to assist the system to deal more effectively with an attack. It is important to note that these recommendations by no means makes the system impervious to attack and only focuses on tuning the TCP/IP stack's response to an attack. The setting of these keys does not address any of the many other components on the system, which could be used to attack the system. As with any change to the registry, the administrator needs to fully understand how these changes affect the default function of the system and whether they are appropriate in their environment.

**SynAttackProtect**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD

**Valid Range:** 0, 1

0 (no SYN attack protection)

1 (reduced retransmission retries and delayed RCE [route cache entry] creation if the *TcpMaxHalfOpen* and *TcpMaxHalfOpenRetried* settings are satisfied and a delayed indication to Winsock is made.)

| **Note** |
| --- |
| When the system finds itself under attack the following options on any socket can no longer be enabled: scalable windows (RFC 1323) and per adapter configured TCP parameters (Initial RTT, window size). This is because when protection is functioning the route cache entry is not queried before the SYN-ACK is sent and the Winsock options are not available at this stage of the connection. |

**Default:** 1 - enabled for Windows Server 2003 Service Pack 1, 0 - disabled for Windows Server 2003 with no service packs installed

**Recommendation:** 1

**Description:** SYN attack protection involves reducing the amount of retransmissions for the SYN-ACKS, which will reduce the time for which resources have to remain allocated. The allocation of route cache entry resources is delayed until a connection is made and the connection indication to AFD is delayed until the three-way handshake is completed. Note that the actions taken by the protection mechanism only occur if *TcpMaxHalfOpen* and *TcpMaxHalfOpenRetried* settings are exceeded.

**TcpMaxHalfOpen**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 100–0xFFFF

**Default:** 100 (Windows Server 2003, Standard Edition), 500 (Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition)

**Description:** This parameter controls the number of connections in the SYN-RCVD state allowed before SYN-ATTACK protection begins to operate. If SynAttackProtect is set to 1, ensure that this value is lower than the AFD listen backlog on the port that you want to protect (see backlog parameters in Appendix C for more information). See the *SynAttackProtect* parameter for more details.

**TcpMaxHalfOpenRetried**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—number

**Valid Range:** 80–0xFFFF

**Default:** 80 (Windows Server 2003, Standard Edition), 400 (Windows Server 2003, Enterprise Edition; Windows Server 2003,

Datacenter Edition)

**Description:** This parameter controls the number of connections in the SYN-RCVD state for which there has been at least one retransmission of the SYN sent, before SYN-ATTACK attack protection begins to operate. See the *SynAttackProtect* parameter for more details.

**EnablePMTUDiscovery**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Recommendation:** 1

**Description:** When this parameter is set to 1 (true) TCP attempts to discover the Maximum Transmission Unit (MTU or largest packet size) over the path to a remote host. By discovering the Path MTU and limiting TCP segments to this size, TCP can eliminate fragmentation at routers along the path that connect networks with different MTUs. Fragmentation adversely affects TCP throughput and network congestion. Setting this parameter to 0 causes an MTU of 576 bytes to be used for all connections that are not to hosts on the local subnet.

**NoNameReleaseOnDemand**
**Key:** Netbt\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 0 (false)

**Recommendation:** 1

**Description:** This parameter determines whether the computer releases its NetBIOS name when it receives a name-release request from the network. It was added to allow the administrator to protect the machine against malicious name-release attacks.

**EnableDeadGWDetect**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—Boolean

**ValidRange:** 0, 1 (false, true)

**Default:** 1 (true)

**Recommendation:** 0

**Description:** When this parameter is set to 1, TCP is allowed to perform dead-gateway detection. With this feature enabled, TCP may ask IP to change to a backup gateway if a number of connections are experiencing difficulty. Backup gateways may be defined in the advanced properties of the TCP/IP protocol. See the "Dead Gateway Detection" section in this article for details.

**KeepAliveTime**
**Key:** Tcpip\Parameters

**Value Type:** REG_DWORD—time in milliseconds

**Valid Range:** 1–0xFFFFFFFF

**Default:** 7,200,000 (two hours)

**Recommendation:** 300,000

**Description:** The parameter controls how often TCP attempts to verify that an idle connection is still intact by sending a keep-alive packet. If the remote system is still reachable and functioning, it acknowledges the keep-alive transmission. Keep-alive packets are not sent by default. This feature may be enabled on a connection by an application.

**PerformRouterDiscovery**
**Key:** Tcpip\Parameters\Interfaces*interfaceGUID*

**Value Type:** REG_DWORD

**Valid Range:** 0, 1, 2

0 (disabled)

1 (enabled)

2 (enable only if DHCP sends the router discover option)

**Default:** 2, DHCP-controlled but off by default.

**Recommendation:** 0

**Description:** This parameter controls whether Windows Server 2003 attempts to perform router discovery per RFC 1256 on a per-interface basis. See also *SolicitationAddressBcast*.

# Appendix E: Format of the Daytime Service Response String

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

RFC 867 does not specify the format of the response sent by the daytime service. For Windows 2000 and Windows Server 2003, the daytime service sends an ASCII text string containing the current system time and date formatted in a single line as follows:

*TimeString DateString*

with a single ASCII space character separating the time (*TimeString*) and date (*DateString*) components. Each of these components is formatted according to the default format for the selected locale of the computer. For example, the default format for “English (United States)” is:

*TimeString* = h:mm:ss tt

*DateString* = m/dd/yyyy

in which tt is either "AM" or "PM".

Here is example output for the English (United States) locale:

10:23:16 AM 4/11/2003

The default formats for the date and time for any locale may be viewed or configured from Control Panel-Regional Options (for Windows 2000) or Control Panel-Regional and Language Options (for Windows Server 2003).

There is no ability to change the format of the daytime service ASCII text string.

# Summary - Microsoft Windows Server 2003 TCP/IP Implementation Details

Updated: July 6, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The TCP/IP stack in Windows Server 2003 has extensive support for standard features, performance enhancements, and services. Windows Server 2003 TCP/IP consists of the following core protocols: Address Resolution Protocol (ARP), Internet Protocol (IP), Internet Protocol security (IPsec), Internet Group Management Protocol (IGMP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and NetBIOS over TCP/IP (NetBT). Windows TCP/IP applications typically use the Windows Sockets or NetBIOS APIs. Windows Server 2003 TCP/IP includes support for automatic client configuration, media sense, and DNS dynamic update and client-side caching.

**For More Information**
For the latest information about Windows Server 2003, see the Windows Server 2003 Web site at (http://go.microsoft.com/fwlink/?LinkID=17075).

For more information about IPv6, see the Microsoft Windows IPv6 Web site (http://go.microsoft.com/fwlink/?LinkID=68045).

1. Specifications and programming information are included in the Microsoft Windows Driver Development Kits (DDKs) (http://go.microsoft.com/fwlink/?LinkID=38448).

2. Most network interface cards have the ability to be placed into a mode in which the network interface card does not perform any address filtering on frames that appear on the media. Instead, it passes every frame upwards that passes the cyclic redundancy check (CRC). This feature is used by some protocol analysis software, such as Microsoft Network Monitor.

3. Adding [1] to the registry parameter TcpMaxDataRetransmissions or TcpMaxConnectRetransmissions approximately doubles the total retransmission time-out period. If it is necessary to configure longer time-outs, these parameters should be increased very gradually.

4. Instead of sending one TCP segment when starting out, Windows NT/Windows Server 2003 TCP starts with two. This avoids the need to wait for the delayed ACK timer to expire on the first send to the target computer, which improves performance for some applications.

5. See the Microsoft Knowledge Base for redirector registry parameters.

6. Stevens, Richard. TCP/IP Illustrated, Volume 1: The Protocols. Reading, MA: Addison-Wesley Publishing Co., 1993.

7. IP autoconfiguration can be disabled using the IPAutoconfigurationEnabled registry key. The subnet and subnet mask used can be controlled using the IPAutoconfigurationSubnet and IPAutoconfigurationMask registry keys. These keys are described in Appendix A.

# An Overview of Windows Clustering Technologies: Server Clusters and Network Load Balancing

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: January 1, 2003

**In This White Paper**

Introduction (An Overview of Windows Clustering Technologies: Server Clusters and Network Load Balancing)

Cluster Architecture Essentials

Server cluster Architecture

Network Load Balancing Architecture

Component Load Balancing Architecture

Summary (An Overview of Windows Clustering Technologies: Server Clusters and Network Load Balancing)

Related Links (An Overview of Windows Clustering Technologies: Server Clusters and Network Load Balancing)

# Introduction (An Overview of Windows Clustering Technologies: Server Clusters and Network Load Balancing)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft listened to customers and worked steadily to improve the underlying technology architecture of the Windows operating system.

**Microsoft Cluster Technologies**
While Windows 2000 represents a dramatic improvement over its predecessors in terms of total uptime (availability), reduced system failure (reliability) and ability to add resources and computers to improve performance (scalability), Windows Server 2003 takes the availability, reliability and scalability of the Windows operating system to the next level by enhancing existing features and providing new options.

**Three Part Clustering Strategy**
Microsoft clustering technologies are the key to improving availability, reliability and scalability. With Windows 2000 and Windows Server 2003, Microsoft uses a three-part clustering strategy that includes:

- **Server cluster** provides failover support for applications and services that require high availability, scalability and reliability. With clustering, organizations can make applications and data available on multiple servers linked together in a cluster configuration. Back-end applications and services, such as those provided by database servers, are ideal candidates for Server cluster.

- **Network Load Balancing** provides failover support for IP-based applications and services that require high scalability and availability. With Network Load Balancing (NLB), organizations can build groups of clustered computers to support load balancing of Transmission Control Protocol (TCP), User Datagram Protocol (UDP) and Generic Routing Encapsulation (GRE) traffic requests. Web-tier and front-end services are ideal candidates for NLB.

- **Component Load Balancing**, which is a feature of Microsoft Application Center 2000, provides dynamic load balancing of middle-tier application components that use COM+. With Component Load Balancing (CLB), COM+ components can be load balanced over multiple nodes to dramatically enhance the availability and scalability of software applications.

**Protection Against Failure**
Microsoft cluster technologies guard against three specific types of failure:

- **Application/service failure** affects application software and essential services.

- **System/hardware failure** affects hardware components (for example, CPUs, drives, memory, network adapters, power supplies, etc.).

- **Site failure** could be caused by natural disaster, power outages or connectivity outages.

**Clustering TechnologiesPurposes and Requirements**
Each technology has a specific purpose and is designed to meet different requirements.

- **Server cluster** is designed to maintain data integrity and provide failover support.

- **Network Load Balancing** is designed to address bottlenecks caused by front-end Web services.

- **Component Load Balancing** is designed to address the unique scalability and availability needs of middle-tier applications.

Organizations can use Microsoft cluster technologies to increase overall availability, while minimizing single points of failure and reducing costs by using industry-standard hardware and software.

**E-Commerce Scenario**
The clustering technologies outlined above can be (and typically are) combined to architect a comprehensive service offering. The most common scenario where all three solutions are combined is an e-commerce site where front-end Web servers use NLB, middle-tier application servers use CLB, and back-end database servers use Server cluster.

These technologies alone are not enough to ensure the highest levels of availability. To ensure the highest availability for line-of-business applications and mission-critical services, organizations must take an end-to-end service approach to operations and availability. This means treating the entire service offering as a whole and designing, implementing, and operating the service solution following industry-standard best practices, such as those used by Microsoft Enterprise Services.

**Topics Covered**
The main topics covered in the rest of this article include:

- Cluster Architecture Essentials
- Server cluster Architecture
- Network Load Balancing Architecture
- Component Load Balancing Architecture

# Cluster Architecture Essentials

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section introduces the concept of clustering and its benefits and limitations. It then goes on to discuss cluster organization, infrastructure scaling, cluster operating modes and how clustering is used on multiple, geographically dispersed sites.

**The Concept of a Cluster**
The concept of a cluster involves taking two or more computers and organizing them to work together to provide higher availability, reliability and scalability than can be obtained by using a single system. When failure occurs in a cluster, resources can be redirected and the workload can be redistributed. Typically, the end user experiences a limited failure, and may only have to refresh the browser or reconnect to an application to begin working again.

**Cluster Benefits and Limitations**
A Server cluster provides high availability by making application software and data available on several servers linked together in a cluster configuration. If one server stops functioning, a process called failover automatically shifts the workload of the failed server to another server in the cluster. The failover process is designed to ensure continuous availability of critical applications and data.

While clusters can be designed to handle failure, they are not fault tolerant with regard to user data. The cluster by itself does not guard against loss of a user's work. Typically, the recovery of lost work is handled by the application software; the application software must be designed to recover the user's work, or it must be designed in such a way that the user session state can be maintained in the event of failure.

**Solving Three Typical Problems**
Clusters can be used to solve three typical problems in a data center environment:

- **Need for High Availability**. High availability refers to the ability to provide end user access to a service for a high percentage of scheduled time while attempting to reduce unscheduled outages. A solution is highly available if it meets the organization's scheduled uptime goals. Availability goals are achieved by reducing unplanned downtime and then working to improve total hours of service operation.

- **Need for High Reliability**. High reliability refers to the ability to reduce the frequency of system failure, while attempting to provide fault tolerance in case of failure. A solution is highly reliable if it minimizes the number of single points of failure and reduces the risk that failure of a single component/system will result in the outage of the entire service offering. Reliability goals are achieved using redundant, fault tolerant hardware components, application software and systems.

- **Need for High Scalability**. High scalability refers to the ability to add resources and computers while attempting to improve performance. A solution is highly scalable if it can be scaled up and out. Individual systems in a service offering can be scaled up by adding more resources (for example, CPUs, memory, disks, etc.). The service can be scaled out by adding additional computers.

A well-designed service solution uses redundant systems and components so that the failure of an individual server does not affect the availability of the entire service.

**Limitations**
While a well-designed solution can guard against application failure, system failure and site failure, cluster technologies do have limitations. Cluster technologies depend on compatible applications and services to operate properly. The software must respond appropriately when failure occurs. Cluster technology cannot protect against failures caused by viruses, software corruption or human error. To protect against these types of problems, organizations need solid data protection and recovery plans.

**Cluster Organization**
Clusters are organized in loosely coupled groups that are often referred to as farms or packs. In most cases, as shown below on Figure 1, front-end and middle-tier services are organized as farms using clones, while back-end and critical support services such as component routing are organized as packs.

**IT Staff Considerations**

As IT staff architect clustered solutions, they need to look carefully at the cluster organization they plan to use. The goal should be to organize servers according to the way the servers will be used and the applications they will be running. Typically, Web servers, application servers and database servers are all organized differently.




**Figure 1: Clusters are organized and farms or packs**

**Cluster Farm**

A farm is a group of servers that run similar services, but do not typically share data. They are called a farm because they handle whatever requests are passed out to them using identical copies of data that is stored locally. Because they use identical copies of data (rather than sharing data), members of a farm operate autonomously and are also referred to as *clones*.

Front-end Web servers running Internet Information Services (IIS) and using NLB are an example of a farm. With a Web farm, identical data is replicated to all servers in the farm, and each server can handle any request that comes to it using local copies of the data. Because the servers are identical and the data is replicated to all the servers in the Web farm, the servers are also referred to as *clones*.

ExampleA Load Balanced Web Farm

In a load balanced Web farm with ten servers, you could have:

- Clone 1Web server using local data
- Clone 2Web server using local data
- Clone 3Web server using local data
- Clone 4Web server using local data
- Clone 5Web server using local data
- Clone 6Web server using local data
- Clone 7Web server using local data
- Clone 8Web server using local data
- Clone 9Web server using local data

- Clone 10Web server using local data

**Cluster Pack**
A pack is a group of servers that operate together and share partitioned data. They are called a pack because they work together to manage and maintain services. Because members of a pack share access to partitioned data, they have unique operations modes and usually access the shared data on disk drives to which all members of the pack are connected.

Example - A 4-node SQL Server Cluster Pack

An example of a pack is a database Server cluster running SQL Server 2000 and a server cluster with partitioned database views. Members of the pack share access to the data and have a unique chunk of data or logic that they handle, rather than handling all data requests.

In a 4-node SQL Server cluster:

- Database Server 1 may handle accounts that begin with A-F.
- Database Server 2 may handle accounts that begin with G-M.
- Database Server 3 may handle accounts that begin with N-S.
- Database Server 4 may handle accounts that begin with T-Z.

Combining Techniques A Large-Scale E-Commerce Site

Servers in a tier can be organized using a combination of the above techniques as well. An example of this combination is a large-scale e-commerce site that has middle-tier application servers running Application Center 2000 and CLB.

To configure CLB, two clusters are recommended:

- The Component Routing Cluster handles the message routing between the front-end Web servers and the application servers.
- The Application Server cluster activates and runs the components installed on the application servers.

While the Component Routing Cluster could be configured on the Web tier without needing additional servers, a large e-commerce site may want the high availability benefits of a separate cluster. In this case, the routing would take place on separate servers that are clustered using Server cluster. The application servers would then be clustered using CLB.

**Infrastructure Scaling**
With proper architecture, the servers in a particular tier can be scaled out or up as necessary to meet growing performance and throughput needs. Figure 2 below, provides an overview of the scalability of Windows clustering technologies.

**IT Staff Considerations**
As IT staff look at scalability requirements, they must always address the real business needs of the organization. The goal should be to select the right edition of the Windows operating system to meet the current and future needs of the project.

The number of servers needed depends on the anticipated server load, and the size and types of requests the servers will handle. Processors and memory should be sized appropriately for the applications and services the servers will be running, as well as the number of simultaneous user connections.




**Figure 2: Windows cluster technologies can be scaled to meet business requirements**

**Scaling by Adding Servers**

When looking to scale out by adding servers to the cluster, the clustering technology and the server operating system used are both important. As Table 1 below shows, the key difference in the outward scaling capabilities of Advanced Server and Datacenter Server is the number of nodes that can be used with Server cluster.

- Under Windows 2000, the maximum number of Server cluster nodes is four.
- Under Windows Server 2003, the maximum number of Server cluster nodes is eight.

**Table 1 Cluster Nodes Supported by Operating System and Technology.**

| Operating System | Edition | Network Load Balancing | Component Load Balancing | Server cluster |
|---|---|---|---|---|
| **Windows 2000** | | | | |
| | Advanced Server | 32 | 8 | 2 |
| | Datacenter Server | 32 | 8 | 4 |
| **Windows Server 2003** | | | | |
| | Enterprise Edition | 32 | 8 | 8 |
| | Datacenter Edition | 32 | 8 | 8 |

**Scaling by Adding CPUs and RAM**

When looking to scale up by adding CPUs and RAM, the edition of the server operating system used is extremely important.

In terms of both processor and memory capacity, Datacenter Server is much more expandable.

- Windows 2000 Advanced Server supports up to eight processors and eight gigabytes (GB) of RAM.
- Windows 2000 Datacenter Server supports up to 32 processors and 64 GB of RAM.
- Windows Server 2003 Enterprise Edition supports up to eight processors and 32 GB of RAM.
- Windows Server 2003 Datacenter Edition supports up to 32 processors and 64 GB of RAM.

Thus, organizations typically scale up from Advanced Server or Enterprise Edition to Datacenter Server or Datacenter Edition as their needs change over time.

**Cluster Operating Modes**
With NLB and CLB, cluster nodes usually are identical clones of each other. Because of this, all members of the cluster can actively handle requests, and they can do so independent of each other. When members of a cluster share access to data, however, they have unique operating requirements, as is the case with Server cluster.

**IT Staff Considerations**
As IT staff considers the impact of operating modes in their cluster architecture, they need to look carefully at the business requirements and the expected server loads.

With NLB and CLB, all servers are active and the architecture is scaled out by adding additional servers, which typically are configured identically to the existing NLB or CLB nodes.

With Server cluster, nodes can be either active or passive, and the configuration of nodes depends on the operating mode (active or passive), as well as how failover is configured. A server that is designated to handle failover must be sized to handle the workload of the failed and the current workload (if any). Additionally, both average and peak workloads must be considered. Severs need additional capacity to handle peak loads.

**Server cluster Nodes**
Server cluster nodes can be either active or passive.

- **Active Node**. When a node is active, it is actively handling requests.

- **Passive Node**. When a node is passive, it is idle, on standby waiting for another node to fail.

Multi-node clusters can be configured using different combinations of active and passive nodes.

**Architecting Multi-node Clusters**
When architecting multi-node clusters, the decision as to whether nodes are configured as active or passive is extremely important. To understand why, consider the following:

- **If an active node fails and there is a passive node available**, applications and services running on the failed node can be transferred to the passive node. Since the passive node has no current workload, the server should be able to assume the workload of the failed server without any problems (providing all servers have the same hardware configuration).

- **If all severs in a cluster are active and a node fails**, the applications and services running on the failed node can be transferred to another active node. Since the server is already active, the server will have to handle the processing load of both systems. The server must be sized to handle multiple workloads or it may fail as well.

**In a multi-node configuration where there is one passive node for each active node**, the servers could be configured so that under average workload they use about 50% of CPU and memory resources.

In the 4-node configuration depicted on Figure 3 below, where failover goes from one active node to a specific passive node, this could mean two active nodes (A1 and A2) and two passive nodes (P1 and P2)each with four processors and 4GB of RAM. In this example, node A1 fails over to node P1 and node A2 fails over to node P2 with the extra capacity used to handle peak workloads.




**Figure 3: Examples of Active/Passive and Active/Active configurations**

In a multi-node configuration where there are more active nodes than passive nodes, the servers can be configured so that under average workload they use a proportional percentage of CPU and memory resources.

In the 4-node configuration illustrated on Figure 3 above, where nodes A, B, C, and D are configured as active and failover could go to between nodes A and B or nodes C and D, this could mean configuring servers so that they use about 25% of CPU and memory resources under average workload. In this example, node A could fail over to B (and vice versa) or node C could fail over to D (and vice versa).

Because the servers in this example would need to handle two workloads in case of a node failure, the CPU and memory configuration would at least be doubled, so instead of using four processors and four GB of RAM, the servers may use eight processors and eight GB of RAM.

**Shared-Nothing Database Configuration**

When Server cluster has multiple active nodes, data must be shared between applications running on the clustered servers. In most cases, this is handled with a shared-nothing database configuration.

In a shared-nothing database configuration, the application is partitioned to access private database sections. This means that a particular node is configured with a specific view into the database that allows it to handle specific types of requests, such as account names that started with the letters A-F, and that it is the only node that can update the related section of the database. This eliminates the possibility of corruption from simultaneous writes by multiple nodes.

| **Note** |
| --- |
| Both Microsoft Exchange 2000 and Microsoft SQL Server 2000 support multiple active nodes and shared-nothing database configurations. |

**Multiple Sites and Geographically Dispersed Clusters**

Most organizations build disaster recovery and increased availability into their infrastructure using multiple physical sites. Multi-site architecture can be designed in many ways. In most cases, the architecture has a primary site and one or more remote sites. Figure 4 below shows an example of a primary site and a remote site for an e-commerce operation.

The architecture at the remote site mirrors that of the primary site. The level of integration for multiple sites and the level at which components are mirrored between sites, depends on the service-level agreement and the business requirements.

**Full Implementation Design**

With a full implementation, the complete infrastructure of the primary site could be recreated at remote sites. This allows for a remote site to operate independently, or to handle the full load of the primary site if necessary. In this case, the design should incorporate real-time replication and synchronization for databases and applications.

Real-time replication ensures a consistent state for data and application services between sites. If real-time updates are not possible, databases and applications should be replicated and synchronized as rapidly as possible.




**Figure 4: Multiple site architecture**

**Partial Implementation Design**

With a partial implementation only essential components are installed at remote sites to:

- Handle overflow in peak periods.
- Maintain uptime on a limited basis in case the primary site fails.
- Provide limited services as needed.

**Replicating static content on Web sites and read-only data from databases**. This partial implementation technique would allow remote sites to handle requests for static content and other types of data that is infrequently changed. Users could browse sites, access account information, product catalogs, and other services. If they needed to access dynamic content or modify information (add, change, delete, etc.), the sites geographical load balancers could redirect users to the primary site.

**Implement all layers of the infrastructure, but with fewer redundancies in the architecture**, or **implement only core components, relying on the primary site to provide the full array of features**. With either of these partial implementation techniques, the design may need to incorporate near real-time replication and synchronization for databases and applications. This ensures a consistent state for data and application services.

**Geographically Dispersed Clusters**

A full or partial design could also use geographically dispersed clusters running Server cluster. Geographically dispersed clusters use virtual LANs to connect storage area networks (SANs) over long distances.

- A VLAN connection with latency of 500 milliseconds or less ensures that cluster consistency can be maintained.
- If the VLAN latency is over 500 milliseconds, the cluster consistency cannot be easily maintained.

Geographically dispersed clusters are also referred to as stretched clusters and are available in Windows 2000 and Windows Server 2003.

**Majority Node Clustering**
Windows Server 2003 offers many improvements in the area of geographically dispersed clusters, including a new type of quorum resource called a majority node set. Majority node clustering changes the way the cluster quorum resource is used. This allows cluster servers to be geographically separated while maintaining consistency in the event of node failure.

With a standard cluster configuration, as illustrated on Figure 5 below, the quorum resource writes information about all cluster database changes to the recovery logs; this ensures that the cluster configuration and state data can be recovered. The quorum resource resides on the shared disk drives and can be used to verify whether other nodes in the cluster are functioning.




**Figure 5: Comparing local and geographically dispersed clusters**

With a majority node cluster configuration in Windows Server 2003, the quorum resource is configured as a majority node set resource. This new type of quorum resource allows the quorum data, which includes cluster configuration changes and state information, to be stored on the system disk of each node in the cluster. Because the data is localized, even though the cluster itself is geographically dispersed, the cluster can be maintained in a consistent state.

As the name implies, the majority of nodes must be available for this cluster configuration to operate normally. Should the cluster state become inconsistent, IT staff can force the quorum to get a consistent state. An algorithm also runs on the cluster nodes to help ensure the cluster state.

# Server cluster Architecture

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses Server cluster and how to configure it for failover support for applications and services. Resource groups, cluster storage devices, network configuration and storage area networks are also discussed.

**Server cluster**
Server cluster is used to provide failover support for applications and services. A Server cluster can consist of up to eight nodes. Each node is attached to one or more cluster storage devices. Cluster storage devices allow different servers to share the same data, and by reading this data provide failover for resources.

**Connecting Storage Devices**
The preferred technique for connecting storage devices is fibre channel.

- When using three or more nodes, fibre channel is the only technique that should be used.
- When using 2-node clustering with Advanced Server, SCSI or fibre channel can be used to connect to the storage devices.

**Configuring Server clusters**
Server clusters can be setup using many different configurations. Servers can be either active or passive, and different servers can be configured to take over the failed resources of another server. Failover can take several minutes, depending on the configuration and the application being used, but is designed to be transparent to the end-user.

**Server cluster and Failover**
When a node is active, it makes its resources available. Clients access these resources through dedicated virtual servers.

Server cluster uses the concept of virtual servers to specify groups of resources that failover together. When a server fails, the group of resources configured on that server for clustering fails over to another server. The server that handles the failover should be configured for the extra capacity needed to handle the additional workload. When the failed server comes back online, Server cluster can be configured to allow failback to the original server, or to allow the current server to continue to process requests.




**Figure 6: Multi-node clusters with all nodes active**

Figure 6 above shows a configuration where all nodes in a database cluster are active and each node has a separate resource group. With a partitioned view of the database, each resource group could handle different types of requests. The types of requests handled could be based on one or more factors, such as the name of an account or geographic location. In the event of a failure, each node is configured to fail over to the next node in turn.

**Resource Groups**
Resources that are related or dependent on each other are associated through resource groups. Only applications that need high availability should be part of a resource group. Other applications can run on a server cluster, but don't need to be a part

of a resource group. Before adding an application to a resource group, IT staff must determine if the application can work within the cluster environment.

**Cluster-Aware Applications**. Applications that can work within the cluster environment and support cluster events are called cluster-aware. Cluster-aware applications can register with the Server cluster to receive status and notification information.

**Cluster-Unaware Applications**. Applications that do not support cluster events are called cluster-unaware. Some cluster-unaware applications can be assigned to resource groups and can be failed over.

Applications that meet the following criteria can be assigned to resource groups.

- IP-based protocols are used for cluster communications. The application must use an IP-based protocol for their network communications. Applications cannot use NetBEUI, IPX, AppleTalk or other protocols to communicate.
- Nodes in the cluster access application data through shared storage devices. If the application is not able to store its data in a configurable location, the application data will not be available on failover.
- Client applications experience a temporary loss of network connectivity when failover occurs. If client applications cannot retry and recover from this, they will cease to function normally.

#### New Features for Resources and Resource Types

Windows Server 2003 adds new features for resources and resource types. A new resource type allows applications to be made cluster-aware using VBScript and JScript. Additionally, Windows Management Instrumentation (WMI) can be used for cluster management and event notification.

#### Architecting Resource Groups

When architecting resource groups, IT staff should list all server-based applications and services that will run in the cluster environment, regardless of whether they will need high availability. Afterward, divide the list into three sections:

- Those that need to be highly available
- Those that are not part of the cluster and on which clustered resources do not depend
- Those that are running on the cluster servers that do not support failover and on which the cluster may depend.

**Applications and services that need to be highly available should be placed into resource groups**. Other applications should be tracked, and their interactions with clustered applications and services should be clearly understood. Failure of an application or service that is not part of a resource group should not impact the core functions of the solution being offered. If it does, the application or service may need to be clustered.

| **Note** |
| --- |
| In the case of dependent services that do not support clustering, IT staff may want to provide backup planning in case these services fail, or may want to attempt to make the services cluster-aware using VBScript and JScript. Remember that only Windows Server 2003 supports this feature. |

**Focus on selecting the right hardware to meet the needs of the service offering**. A cluster model should be chosen to adequately support resource failover and the availability requirements. Based on the model chosen, excess capacity should be added to ensure that storage, processor and memory are available in the event a resource fails, and failover to a server substantially increases the workload.

**With a clustered SQL Server configuration, IT staff should consider using high-end CPUs, fast hard drives and additional memory**. SQL Server 2000 and standard services together use over 100 megabytes (MB) of memory as a baseline. User connections consume about 24 kilobytes (KB) each. While the minimum memory for query execution is one MB of RAM, the average query may require two to four MB of RAM. Other SQL Server processes use memory as well.

#### Optimizing Cluster Storage Devices

Cluster storage devices should be optimized based on performance and availability needs. While the Windows Datacenter Hardware Compatibility List provides a detailed list of acceptable Redundant Array of Independent Disks (RAID) configurations for clusters, Table 2 below provides an overview of common RAID configurations. The table entries are organized from the highest RAID level to the lowest.

**Table 2 RAID Configurations**

| RAID Level | RAID Type | RAID Description | Advantages & Disadvantages |
|---|---|---|---|
| 5+1 | Disk striping with parity + mirroring | Six or more volumes, each on a separate drive, are configured identically as a mirrored stripe set with parity error checking. | Provides very high level of fault tolerance but has a lot of overhead. |
| 5 | Disk striping with parity | Three or more volumes, each on a separate drive, are configured as a stripe set with parity error checking. In the case of failure, data can be recovered. | Fault tolerance with less overhead than mirroring. Better read performance than disk mirroring. |
| 1 | Disk mirroring | Two volumes on two drives are configured identically. Data is written to both drives. If one drive fails, there is no data loss because the other drive contains the data. (Does not include disk striping.) | Redundancy. Better write performance than disk striping with parity. |
| 0+1 | Disk striping with mirroring | Two or more volumes, each on a separate drive, are striped and mirrored. Data is written sequentially to drives that are identically configured. | Redundancy with good read/write performance. |
| 0 | Disk striping | Two or more volumes, each on a separate drive, are configured as a stripe set. Data is broken into blocks, called stripes, and then written sequentially to all drives in the stripe set. | Speed/Performance without data protection. |

**Optimizing Network Configuration**

The network configuration of the cluster can also be optimized. All nodes in a cluster must be a part of the same domain and can be configured as domain controllers or member servers. Ideally, multi-node clusters will have at least two nodes that act as domain controllers and provide failover for critical domain services. If this is not the case, the availability of cluster resources may be tied to the availability of the controllers in the domain.

**Private and Public Network Addresses**

Typically nodes in a cluster are configured with both private and public network addresses.

- Private network addresses are used for node-to-node communications.
- Public network addresses are used for client-to-cluster communications.

Some clusters may not need public network addresses and instead may be configured to use two private networks. In this case, the first private network is for node-to-node communications and the second private network is for communicating with other servers that are a part of the service offering.

**Storage Area Networks**

Increasingly, clustered servers and storage devices are connected over SANs. SANs use high-performance interconnections between secure servers and storage devices to deliver higher bandwidth and lower latency than comparable traditional networks. Windows 2000 Datacenter Server and Windows Server 2003 Datacenter Edition implement a feature called Winsock Direct that allows direct communication over a SAN using SAN providers.

**SAN providers have user-mode access to hardware transports**. When communicating directly at the hardware level, the individual transport endpoints can be mapped directly into the address space of application processes running in user mode. This allows applications to pass messaging requests directly to the SAN hardware interface, which eliminates unnecessary system calls and data copying.

**SANs typically use two transfer modes**. One mode is for small transfers, which primarily consist of transfer control information. For large transfers, SANs can use a bulk mode whereby data is transferred directly between the local system and the remote system by the SAN hardware interface without CPU involvement on the local or remote system. All bulk transfers are pre-arranged through an exchange of transfer control messages.

**Other SAN Benefits**

In addition to improved communication modes, SANs have other benefits.

- They allow IT staff to consolidate storage needs, using several highly reliable storage devices instead of many.

- They also allow IT staff to share storage with non-Windows operating systems, allowing for heterogeneous operating environments.

# Network Load Balancing Architecture

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses Network Load Balancing and its role in providing failover support for IP-based applications and services that require high scalability and availability.

**Network Load Balancing**
Network Load Balancing provides failover support for IP-based applications and services that require high scalability and availability. NLB allows IT staff to incrementally scale-out up to 32 servers as demand increases. NLB is ideally suited to improving the availability of Web servers, media servers, terminal servers and e-commerce sites. Load balancing these services ensures there is no single point of failure and that there is no performance bottleneck.

Many of the concepts that apply to Server cluster also apply to NLB. NLB nodes work together to provide availability for critical IP-based resources, which can include: TCP, UDP and GRE traffic requests.

| **Note** |
| --- |
| Support for GRE traffic is supported in Windows Server 2003 and not in Windows 2000. |

**Failover and Failback Using a Virtual IP Address**
As shown on Figure 7 below, NLB uses a virtual IP address. Client requests are directed to this virtual IP address which allows for transparent failover and failback. When a load balanced resource fails on one server, the remaining servers in the group take over the workload of the failed server. When the failed server comes back online, the server can automatically rejoin the cluster group and NLB starts to automatically distribute the load to the server. Failover takes less than 10 seconds in most cases.




**Figure 7: Network load balancing overview**

**No Clustered Storage Devices**
NLB does not use clustered storage devices. Each server runs a copy of the IP-based application or service that is being load balanced, and the data necessary for the application or service to run is stored on local drives.

**Directing Traffic to a Specific Server**
Although NLB is normally used to distribute the workload for an application or service, NLB can also be used to direct a specific type of traffic to a particular server. For example, IT staff may want to load balance HTTP and FTP traffic to a group of servers, but may also want a single server to handle media services traffic. In this latter case, NLB allows traffic to flow to a designated server and only reroutes traffic to another server in case of failure.

**No Hardware Changes Required**
NLB runs as a network driver and requires no hardware changes to install and run. Its operations are transparent to the IP networking stack. Because NLB is IP-based, IP networking must be installed on all load-balanced computers.

**NLB Network Adapters**
To provide high performance throughput and responsiveness, NLB normally uses two network adapters.

- Cluster adapterhandles network traffic for the cluster.
- Dedicated adapterhandles client-to-cluster network traffic, and other traffic originating outside the cluster network.

NLB uses unicast or multicast broadcasts to direct incoming traffic to all servers in the cluster. The NLB driver on each host acts as a filter between the cluster adapter and the TCP/IP stack, allowing only traffic bound for the designated host to be received. NLB only controls the flow of TCP, UDP and GRE traffic on specified ports. It does not control the flow of TCP, UDP and GRE traffic on non-specified ports, and it does not control the flow of other incoming IP traffic. All traffic that is not controlled is passed through without modification to the IP stack.

**Using a Single NLB Network Adapter**
NLB can work with a single network adapter. When it does so, there are limitations.

**Unicast mode**. With a single adapter in unicast mode, node-to-node communications are not possible, meaning nodes within the cluster cannot communicate with each other. Servers can, however, communicate with servers outside the cluster subnet.

**Multicast mode**. With a single adapter in multicast mode, node-to-node communications are possible, as are communications with servers outside the cluster subnet. However, the configuration is not optimal for handling moderate-to-heavy traffic from outside the cluster subnet to specific cluster hosts.

For handling node-to-node communications and moderate to heavy traffic, two adapters should be used.

**Optimizing NLB Servers**
As with Server cluster, servers that use NLB can benefit from optimization. Servers should be optimized for their role, the types of applications they will run and the anticipated local storage they will use.

While IT staff may want to build redundancy into the local hard drives on NLB servers, this adds to the expense of the server without significant availability gains in most instances. Because of this, NLB servers often have drives that do not use RAID and do not provide fault tolerance; the idea being that if a drive causes a server failure, other servers in the NLB cluster can quickly take over the workload of the failed server.

**Synchronizing Data**
If it seems odd to not use RAID, keep in mind that servers using NLB are organized as clones with identical copies of data on each server. Because many different servers have the same data, maintaining the data with RAID sets is not as important as it is with server clusters. A key consideration IT staff does have to make when using NLB, however, is data synchronization. The state of the data on each server must be maintained so that the clones are updated whenever changes are made. This need to synchronize data periodically is an overhead that needs to be considered when designing the server architecture.

# Component Load Balancing Architecture

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses CLB and its key structures. Also discussed are routing servers, designing and optimizing CLB clusters, and storage and memory requirements.

**Component Load Balancing**
Unlike Server cluster and NLB, which are built into the Advanced Server and Datacenter Server editions of the Windows operating system, CLB is a feature of Microsoft Application Center 2000. It is designed to provide high availability and scalability for transactional components. CLB is scalable up to eight servers and is ideally suited to building distributed solutions.

CLB makes use of the COM+ Services supplied as part of the Windows 2000 and Windows Server 2003 operating systems. COM+ Services provide:

- Enterprise functionality for transactions
- Object management
- Security
- Events
- Queuing

COM+ components use the Component Object Model (COM) and COM+ Services to specify their configuration and attributes. Groups of COM+ components that work together to handle common functions are referred to as COM+ applications.

**CLBKey Structures**
Figure 8 below provides an overview of CLB. CLB uses several key structures:

- **CLB Software** handles the load balancing and is responsible for determining the order in which cluster members activate components.
- The **router** handles message-routing between the front-end Web servers and the application servers. It can be implemented through component routing lists stored on front-end Web servers, or a component routing cluster configured on separate servers.
- **Application server clusters** activate and run COM+ components. The application Server cluster is managed by Application Center 2000.




**Figure 8: Component load balancing**

**Routing List**
The routing list, made available to the router, is used to track the response time of each application server from the Web servers. If the routing list is stored on individual Web servers, each server has its own routing list and uses this list to periodically check the response times of the application servers. If the routing list is stored on a separate routing cluster, the

routing cluster servers handle this task.

The goal of tracking the response time is to determine which application server has the fastest response time from a given Web server. The response times are tracked as an in-memory table, and are used in round robin fashion to determine which application server should be passed an incoming request. The application server with the fastest response time (and theoretically, the least busy and most able to handle a request) is given the next request. The next request goes to the application server with the next fastest time, and so on.

**Designing CLB Clusters**
The architecture of CLB clusters should be designed to meet the availability requirements of the service offering. With small-to-moderate sized implementations, the front-end Web servers can host the routing list for the application Server cluster. With larger implementations, dedicated routing clusters are desirable to ensure that high availability requirements can be met.

**Optimizing CLB Servers**
As with NLB, servers in CLB clusters should be optimized for their role, the types of applications they will run and the anticipated local storage they will use.

**High Speed Connections**
Routing servers maintain routing lists in memory and need high-speed connections to the network.

**Storage and Memory Requirements**
Whether configured separately or as part of the front-end, CLB does not require a lot of storage, but a limited amount of additional RAM may be required. Application servers on the other hand, typically need a lot of RAM, fast CPUs and limited redundancy in the drive array configuration. If redundant drive arrays are used, a basic configuration, such as RAID 1 or RAID 5, may be all that is needed to maintain the level of availability required.

# Summary (An Overview of Windows Clustering Technologies: Server Clusters and Network Load Balancing)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Cluster technologies are becoming increasingly important to ensure service offerings meet the requirements of the enterprise. Windows 2000 and Windows Server 2003 support three cluster technologies to provide high availability, reliability and scalability. These technologies are: NLB, CLB and Server cluster. These technologies have a specific purpose and are designed to meet different requirements.

- **Server cluster** provides failover support for applications and services that require high availability, scalability and reliability, and is ideally suited for back-end applications and services, such as database servers. Server cluster can use various combinations of active and passive nodes to provide failover support for mission critical applications and services.

- **NLB** provides failover support for IP-based applications and services that require high scalability and availability, and is ideally suited for Web tier and front-end services. NLB clusters can use multiple adapters and different broadcast methods to assist in the load balancing of TCP, UDP and GRE traffic requests.

- **Component Load Balancing** provides dynamic load balancing of middle-tier application components that use COM+ and is ideally suited for application servers. CLB clusters use two clusters. The routing cluster can be configured as a routing list on the front-end Web servers or as separate servers that run Server cluster.

Cluster technologies by themselves are not enough to ensure that high availability goals can be met. Multiple physical locations may be necessary to guard against natural disasters and other events that may cause complete service outage. Effective processes and procedures, in addition to good architecture, are the keys to high availability.

# Related Links (An Overview of Windows Clustering Technologies: Server Clusters and Network Load Balancing)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

See the following resources for further information:

- Windows 2000 Server at www.microsoft.com/windows2000/server/
- Windows Server 2003 at www.microsoft.com/net/
- Application Center 2000 at www.microsoft.com/applicationcenter/
- Server clusters & Load Balancing at www.microsoft.com/windows2000/technologies/clustering/
- Increasing System Reliability and Availability with Windows 2000 at http://www.microsoft.com/windows2000/docs/relavail.doc
- Hardware Compatibility List at www.microsoft.com/whdc/hcl/default.mspx

# Implementing and Administering Certificate Templates in Windows Server 2003

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

**In This White Paper**


Acknowledgements (Implementing and Administering Certificate Templates in Windows Server 2003)

Introduction (Implementing and Administering Certificate Templates in Windows Server 2003)

Certificate Template Overview

Administering Version 1 Certificate Templates

Administering Version 2 Certificate Templates

Deploying Certificate Templates (Implementing and Administering Certificate Templates in Windows Server 2003)

Planning Considerations

Walkthroughs (Implementing and Administering Certificate Templates in Windows Server 2003)

Appendix A: Wireless Certificates

Appendix B: References

Appendix C: Certificate Template Schema Additions


For information about upcoming certificate template-related features, see Implementing and Administering Certificate Templates in Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkID=92522).

# Acknowledgements (Implementing and Administering Certificate Templates in Windows Server 2003)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

David Cross, Program Manager, Microsoft Corporation.

Mike Danseglio, Technical Writer, Microsoft Corporation

Carsten B. Kinder, Senior Consultant, Microsoft Germany

# Introduction (Implementing and Administering Certificate Templates in Windows Server 2003)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

**Scope**

The scope of this white paper is to discuss the best practices in designing, administering, and implementing version 2 Certificate Templates using Windows Server 2003, Enterprise Edition and Enterprise Certification Authorities (CAs).

**Terms Used in this White Paper**

**Application Constraints** A constraint that limits what purposes a certificate can be used for in a qualified subordination configuration. A presented certificate must contain the required application constraint to be accepted by the partner organization.

**Authority Information Access (AIA)** A certificate extension that contains URL locations where the issuing CAs certificate can be retrieved. The AIA extension can contain HTTP, FTP, LDAP, or FILE URLs.

**Certificate Revocation List (CRL)** A digitally signed list issued by a CA that contains a list of certificates issued by the CA that have been revoked. The list includes the serial number of the certificate, the date that the certificate was revoked, and the revocation reason. Applications can perform CRL checking to determine a presented certificates revocation status; also referred to as a base CRL.

**CRL Distribution Point (CDP)** A certificate extension that indicates where the certificate revocation list for a CA can be retrieved. This extension can contain multiple HTTP, FTP, FILE, or LDAP URLs for the retrieval of the CRL.

**Delta Certificate Revocation List (delta CRL)** A type of CRL that contains the list of certificates revoked since the last base CRL was published. Delta CRLs are often used in environments where numerous certificates are revoked to optimize bandwidth utilization.

**Issuance Policy Constraints** A constraint that defines what issuance practices must be followed for certificates to be trusted by your organization. Issuance policy object identifiers in your organization are mapped to the matching object identifiers in a partner organization, so that object identifiers in presented certificates are recognized by your public key infrastructure (PKI).

**Name Constraints** A constraint that limits what names are permitted or excluded in certificate requests submitted to a CA.

**Online Certificate Status Protocol (OCSP)** A protocol that allows real-time validation of a certificates statusthe CryptoAPI makes a call to an OCSP responder and the OCSP responder provides an immediate validation of the revocation status for the presented certificate. Typically, the OCSP responder uses CRL checking for maintaining its status information.

**Public Key Infrastructure (PKI)** A PKI provides an organization with the ability to securely exchange data over a public network using public key cryptography, thus ensuring privacy by preventing the interception of communications. A PKI consists of Certification Authorities (CAs) that issue digital certificates, directories that store the certificates (including Active Directory in Windows 2000 and Windows Server 2003), and X.509 certificates that are issued to security entities on the network. The PKI provides validation of certificate-based credentials and ensures that the credentials are not revoked, corrupted, or modified.

**Security Principal** A user, security group, or computer account that can be assigned permissions in a Windows Server 2003 discretionary access control list (DACL).

**Subject Key Identifier (SKI)** A certificate extension included in CA certificates that contains a hash of the CA certificates public key. This hash is placed in the Authority Key Identifier (AKI) extension of all issued certificates to facilitate chain building.

# Certificate Template Overview

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Windows 2000 introduced the concept of using certificate templates to define the format and content of a certificate. Certificate templates are used by Windows 2000 Enterprise CAs to define what certificates can be issued by the Windows 2000 Enterprise CAs. Associated with the certificate template is a discretionary access control list (DACL) that defines which security principals have permissions to read, enroll, and configure the certificate template. Enterprise CAs are integrated into Active Directory. The certificate templates and the DACLs of the certificate template objects are defined in Active Directory with a forest-wide validity. If more than one Enterprise CA is running in the Windows forest, permission changes would have an impact on all Enterprise CAs.

The certificate templates used by Windows 2000 Enterprise CAs are known as version 1 certificate templates. Windows 2000 shipped with a number of predefined version 1 certificate templates, but modification of these default certificate templates is not allowed. The only modification that is enabled is the changing of permissions to allow enrollment of the certificate template. The version 1 certificate templates are created by default when an Enterprise CA is installed.

Windows Server 2003 extends certificate templates by introducing version 2 templates. Version 2 templates allow customization of most settings in the template. Several preconfigured version 2 templates are supplied in the default configuration and more can be added as necessary. This allows complete configuration flexibility for administrators. Alternatively, a version 1 certificate template can be duplicated, resulting in a version 2 certificate template that can be modified and secured separately.

| **Note** |
| --- |
| Similar to Windows 2000, Windows Server 2003 supports only version 1 templates. Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition support both version 1 and version 2 templates. Certificates based on version 2 templates can only be issued by an Enterprise CA running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. |

When a certificate template is defined, the definition of the certificate template must be available to all CAs in the forest. This is accomplished by storing the certificate template information in the Configuration naming context (CN=Configuration,DC=*ForestRootName*). The replication of this information depends on the Active Directory replication schedule, and the certificate template may not be available at all CAs until replication is completed. This storage and replication is accomplished automatically by Windows Server 2003 family computers.

**Requirements**

To set up a Windows Server 2003 CA, the Active Directory schema must be upgraded to the Windows Server 2003 schema. You cannot install a Windows Server 2003 CA into a Windows 2000based schema.

The schema is updated to the Windows Server 2003 schema by running ADPREP /Forestprep at a Windows 2000 domain controller with the Windows Server 2003 CD-ROM in the CD-ROM drive.

**Upgrading from Version 1 to Version 2 Certificate Templates**

When you install Windows Server 2003 CA into a Windows Server 2003based Active Directory, the current certificate templates are updated during the upgrade process. The update modifies default settings for the Windows 2000 version 1 certificate templates that implement better security defaults. If a Windows Server 2003, Enterprise Edition CA is installed in addition several, version 2 certificate templates are created.

The upgrade process of an Enterprise CA must be performed by an account that is a member of the forest root Domain Admins group and the Enterprise Admins universal group. This is because the upgrade makes modifications to the Configuration naming context in Active Directory. Specifically, the account performing the upgrade must have the following permissions through group memberships (these are the default permissions):

- Full control permissions over the CN=Certificate Templates, CN=Public Key Services,CN=Services,CN=Configuration, DC=*ForestRootDomain*" container

- Full control permissions over the CN=OID,CN=Public Key Services,CN=Services,CN=Configuration, DC=*ForestRootDomain* container

- Full Control permissions for each certificate template object in the CN=Certificate Templates, CN=Public Key Services,CN=Services,CN=Configuration, DC=*ForestRootDomain*" container

| **Note** |
| --- |
| Delegation over the Certificate Templates container will have no effect on individual certificate templates. In other words, the ACL on certificate templates is not inherited from the ACL on the container. |

To upgrade the certificate templates, perform the following procedure after the upgrade for a Certification Authority to Windows Server 2003 or the installation of a new Windows Server 2003 CA on the network:

1. Upgrade to the Windows Server 2003 schema.

2. Log on as a user account that is a member of the forest root Domain Admins group and the Enterprise Admins group.

3. At a Windows Server 2003, Enterprise Edition CA (the CA can be running on Windows Server 2003, Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition configured as a member-server or Domain Controller), run the Certificate Templates MMC console (certtmpl.msc).

   | **Note** |
   | --- |
   | Alternatively, the Certificate Templates MMC console can be run from a Windows XP Professional computer with the Windows Server 2003 Administration Pack (Adminpak.msi) installed. The same permissions apply as noted previously. |

4. When prompted to write new certificate templates, click OK.

To verify that the upgrade is successful, open the Certificate Templates MMC console and confirm that there are 29 certificate templates. The Version # of templates should all exist and be in the format of xxx.xxx, for example, 100.2. Version 1 certificate templates use a single digit for the primary version number, for example, the Administrator certificate template version number is 3.1. Version 2 certificate template primary version numbers are three digits in length. For example, the Key Recovery Certificate Template version number is 105.0.

| **Note** |
| --- |
| An upgrade of the certificate templates is performed run if a new Windows Server 2003 CA is installed in the forest. If a Windows 2000 CA is upgraded to Windows Server 2003, the template upgrade is not performed automatically and will only be performed when the certificate templates MMC snap-in is opened for the first time. You can still verify that the update has taken place, but the process is performed automatically. |

**Default Templates**

Once the upgrade to Windows Server 2003 certificate templates is completed, the following preconfigured certificate templates are listed in the Certificate Templates MMC console.

| **Name** | **Description** | **Key Usage** | **Subject Type** | **Published to AD** |
| --- | --- | --- | --- | --- |
| Administrator | Allows trust list signing and user authentication | Signature and encryption | User | Yes |
| Authenticated Session | Subject can authenticate to a Web server | Signature | User | No |
| Basic EFS | Used by Encrypting File System (EFS) to encrypt data | Encryption | User | Yes |
| CA Exchange | Used to store keys that are configured for private key archival | Encryption | Computer | No |

| | | | | |
|---|---|---|---|---|
| CEP Encryption | Allows the holder to act as a registration authority (RA) for simple certificate enrollment protocol (SCEP) requests | Encryption | Computer | No |
| Code Signing | Used to digitally sign software | Signature | User | No |
| Computer | Allows a computer to authenticate itself on the network | Signature and encryption | Computer | No |
| Cross-Certification Authority | Used in cross-certification and qualified subordination | Signature | CrossCA | Yes |
| Directory E-mail Replication | Used to replicate e-mail within Active Directory | Signature and encryption | DirEmailRep | Yes |
| Domain Controller | All-purpose certificates held by domain controllers | Signature and encryption | DirEmailRep | Yes |
| Domain Controller Authentication | Used to authenticate Active Directory computers and users | Signature and encryption | Computer | No |
| EFS Recovery Agent | Allows the subject to decrypt files previously encrypted with EFS | Encryption | User | No |
| Enrollment Agent | Used to request certificates on behalf of another subject | Signature | User | No |
| Enrollment Agent (Computer) | Used to request certificates on behalf of another computer subject | Signature | Computer | No |
| Exchange Enrollment Agent (Offline request) | Used to request certificates on behalf of another subject and supply the subject name in the request | Signature | User | No |
| Exchange Signature Only | Used by Microsoft Exchange Key Management Service to issue certificates to Exchange users for digitally signing e-mail | Signature | User | No |
| Exchange User | Used by Microsoft Exchange Key Management Service to issue certificates to Exchange users for encrypting e-mail | Encryption | User | Yes |
| IPSEC | Used by IP Security (IPSec) to digitally sign, encrypt, and decrypt network communication | Signature and encryption | Computer | No |
| IPSEC (Offline request) | Used by IP Security (IPSec) to digitally sign, encrypt, and decrypt network communication when the subject name is supplied in the request | Signature and encryption | Computer | No |
| Key Recovery Agent | This certificate can recover private keys archived on the certification authority. | Encryption | KRA | Yes |
| Root Certification Authority | Used to prove the identity of the root certification authority | Signature | CA | Yes |

| **Name** | **Description** | **Key Usage** | **Subject Type** | **Published to AD** | **Template Version** |
|---|---|---|---|---|---|
| Router (Offline request) | Used by a router when requested through SCEP from a CA that holds a CEP Encryption certificate | Signature and encryption | Computer | No | 3.1 |

| | | | | | |
|---|---|---|---|---|---|
| Smartcard Logon | Allows the holder to authenticate using a smart card | Signature and encryption | User | No | 5.1 |
| Smartcard User | Allows the holder to authenticate and protect e-mail using a smart card | Signature and encryption | User | Yes | 9.1 |
| Subordinate Certification Authority | Used to prove the identity of the root certification authority, issued by the parent or root certification authority | Signature | CA | Yes | 4.1 |
| Trust List Signing | The holder can digitally sign a trust list. | Signature | User | No | 2.1 |
| User | Certificate to be used by users for e-mail, EFS, and client authentication | Signature and encryption | User | Yes | 2.1 |
| User Signature Only | Allows users to digitally sign data | Signature | User | No | 3.1 |
| Web Server | Proves the identity of a Web server | Signature and encryption | Computer | No | 3.1 |

# Administering Version 1 Certificate Templates

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

In Windows 2000, certificate management was very limited because only the templates security permissions could be set. This was done through advanced view of the Active Directory Sites and Services MMC snap-in.

With Windows Server 2003, certificate template management is done through the Certificate Templates MMC snap-in rather than through the Active Directory Sites and Services MMC snap-in.

# Administering Version 2 Certificate Templates

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Version 2 certificate templates allow you to define specific attributes for certificates that meet your organization's business needs. For example, certificate templates allow you to

- Define whether the private key associated with a certificate can be exported.
- Define whether the certificate request must be approved by a certificate manager, and define how many managers must approve a request before the certificate is issued.
- Define what cryptographic service providers (CSPs) are supported by a certificate template.
- Define issuance and application policy for issued certificates.

The following sections detail the specific information that can be configured for a version 2 certificate template using the Certificate Templates MMC console (certtmpl.msc). To create a new certificate template, you start by duplicating an existing certificate template that is similar in functionality to the required certificate template. Selecting the correct initial template is important so that the current certificate template settings will serve as a starting point for your certificate template configuration.

| **Note** |
| --- |
| Since certificate templates exist only once within a forest, all changes made apply on a forest-wide level. If several CAs exist within one forest, templates can be assigned to CAs via the publishing mechanism. Nevertheless, all CAs within a forest use one set of templates. |

**The General Tab**

The General tab (Figure 1) is the first tab that appears when you duplicate an existing certificate template.




**Figure 1: The General Tab**

On the General tab, you can configure the following attributes of the certificate template:

- **Template display name** The display name shown in the Certificate Templates console, Certificates console, and in the

Certification Authority console.

- **Template name** The name of the certificate template object created in the CN=Certificate Templates,CN=Public Key Services,CN=Services,DC=*ForestRootDomain* container.

| **Note** |
| --- |
| The template display name and template name attributes cannot be changed once the certificate template is created. |

- **Validity period** Defines the validity period for an issued certificate. The validity period must not be defined to be greater than the validity period of the CA's certificate. The minimum renewal period is 80% of the certificate lifetime or 6 weeks, whichever is greater.

- **Renewal period** The time before the validity period expires when the certificate will be renewed if re-enrollment is supported for the certificate template.

- **Publish options** You can choose whether to publish the certificate to Active Directory based on the certificate template. The certificates are published as an attribute of the security principal contained in the subject of the issued certificate.

- **Reenrollment option** If the certificate template is published to Active Directory, you can prevent re-enrollment if a valid certificate of the same certificate template exists for the security principal indicated in the subject.

| **Note** |
| --- |
| Re-enrollment settings mainly affect auto-enrollment design in a Windows Server 2003 network. See the Windows 2000 Certificate Autoenrollment in Windows XP white paper at http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/certenrl.mspx |

**The Request Handling Tab**

The Request Handling tab (Figure 2) defines the purpose of the certificate template, the supported cryptographic service providers (CSPs), minimum key length, exportability, auto-enrollment settings, and whether strong private key protection should be required.




**Figure 2: The Request Handling Tab**

**Certificate Purposes**

The certificate purpose defines the intended primary use of the certificate. The certificate purpose can be one of four settings:

- **Encryption** A certificate with this purpose will contain cryptographic keys for encryption and decryption.
- **Signature** A certificate with this purpose will contain cryptographic keys for signing data only.
- **Signature and encryption** A certificate with this purpose covers all primary uses of a certificate's cryptographic key, including encryption of data, decryption of data, initial logon, or digitally signing data.
- **Signature and smartcard logon** A certificate with this purpose allows for initial logon with a smart card, and to digitally sign data; it cannot be used for data encryption.

| **Note** |
| --- |
| The certificate purpose setting will determine whether key archival can be enabled for a certificate template. Key archival is only possible if the certificate purpose is set to Encryption or Signature and encryption. The recovery of a private key for digitally signing information may result in identity theft and is not supported. Key archival is not supported by most smart card CSPs. |

**Archive Settings**

When subjects lose their private keys due to user profile corruption or stolen computers, any information that was persistently encrypted with the corresponding public key is inaccessible. To help avoid this, Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition have the ability to archive a subject's keys in its database when certificates are issued. These keys are encrypted and stored by the CA. If subjects lose their keys, the information can be retrieved from the database and securely provided to the subjects. This allows the encrypted information to be recovered instead of lost.

The following Key Archival settings are defined in the Request Handling tab:

- **Archive subject's encryption private key** If the issuing Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition certification authority is configured for key archival, the subject's private key will be archived.
- **Allow private key to be exported** When this option is specified, the subjects private key can be exported for backup or transportation.
- **Deleting revoked or expired certificates (do not archive)** If a certificate is renewed due to expiration or revocation, the previously issued certificate is removed from the subject's certificate store. By default, the certificate is archived.
- **Include symmetric algorithms allowed by the subject** When the subject requests the certificate, they can supply a list of supported symmetric algorithms. This option allows the issuing certification authority to include those algorithms in the certificate, even if they are not recognized or supported by that server. The algorithms are used by applications like secure e-mail.

To enable key archival and recovery, the following configuration settings in the certificate template and at the Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition CA must be made:

- The specific certificate template must be configured to allow key archival.
- One or more key recovery agents must be identified on the certification authority and key recovery agent certificates must be issued to those agents.
- Key archival must be configured on the certification authority.

**User Input Settings**

The Request Handling tab also allows several user input settings to be defined for a certificate template. These settings include:

- **Enroll subject without requiring any user input** This option allows auto-enrollment without any user interaction and is the default setting for both machine and user certificates. This option must not be enabled for machine certificates.

- **Prompt the user during enrollment** Although useful when testing initial auto-enrollment deployments, typically, this option is disabled. By disabling this option, users do not have to provide any input for the installation of a certificate based on the certificate template.

- **Prompt the user during enrollment and require user input when the private key is used** This option enables the user to set a strong private key protection password on the users private key when the key is generated and requires the user to use it whenever the certificate and private key are used. This option must not be enabled for machine certificates or smart card user certificates.

| **Note** |
| --- |
| To enable smart card auto-enrollment, the **Prompt the user during enrollment** option must be enabled so that the user is prompted to insert the smart card in the smart card reader when required. |

| **Important** |
| --- |
| Strong private key protection with auto-enrollment is not supported or enabled for machine certificates and is only available on Windows XP Service Pack 1 client computers. |

In addition, with Windows 2000 Service Pack 4 and Windows XP Service Pack 1, it is possible to force strong private key protection for all CSPs through the registry. Three new keys can be added to "HKLM\Software\Policies\Microsoft\Cryptography" in the registry:

**ForceKeyProtection** This key will force DPAPI to grey-out the option that would allow the user the choice of using a password when UI was selected. When set, the user MUST enter a password.

<**"0"**> = do not force UI on key protection
<**"1"**> = default to UI, but let user change selection
<**"2"**> = force UI on key protection, grey-out option for user

**CachePrivateKeys** Contains a 1, if (and only if) the following registry key is used:

**PrivateKeyLifetimeSeconds** Contains the number of seconds that a key will remain cached without being used. The key cache timer will be reset on every successful use in the CSP.

| **Caution** |
| --- |
| Requiring the use of strong private key protection and user prompting on all new and imported keys will disable some applications, such as EFS and wireless (802.1x authentication) which cannot display UI. For more information, refer to the following Knowledge Base Article: http://support.microsoft.com/default.aspx?scid=kb;en-us;320828 |

#### Other Request Handling Settings

In addition to key archival settings, some general options that affect all certificates, including those that do not enable key archival, can be defined. These include

- **Minimum key size** This specifies the minimum size, in bits, of the key that will be generated for this certificate.

- **Cryptographic service providers** This is a list of cryptographic service providers (CSPs) that will be used to enroll certificates for the given template. Selecting one or more CSPs configures the certificate to only work with those CSPs. If you do not select a specific CSP, the certificate enrollment will work with any installed CSP, but will use the first CSP from the enumeration list. The CSP must be installed on the client workstation for the CSP to be used during enrollment. If a specific CSP is chosen and not available on a client computer, enrollment will fail.




| **Note** |
| --- |
| You can add third-party CSPs by installing the manufacturers CSP-related files. The newly installed CSP will appear in the list of available CSPs when the certificate templates snap-in is opened after CSP installation. |

**The Subject Name Tab**

When establishing a certificate template, the subject name must be defined. This is included in the issued certificate template and must uniquely identify the subject. The subject name can either be built automatically during the certificate request using the authentication name of the subject, or explicitly defined by the subject and included in the certificate request (Figure 3).




**Figure 3: The Subject Name Tab**

A number of options can be included with the subject name, as well as specific configuration settings for the subject name when the subject name is built from Active Directory information during the certificate request process. The format of the subject name can be defined as

- **None** Does not enforce any name format for this field.
- **Common name** The certification authority creates the subject name from the common name (CN) of the requestor obtained from Active Directory. These should be unique within a domain, but may not be unique within an enterprise.
- **Fully distinguished name** The certification authority creates the subject name from the fully distinguished name

obtained from Active Directory. This guarantees that the name is unique within an enterprise.

In addition, you can choose to include the e-mail name in the subject name. This information is populated from the E-mail attribute of an account, and is included with either the common name or fully distinguished name as part of the subject name.

In addition to the subject name, you can also choose what name formats are included in the alternate subject name attributes of issued certificates. The alternate subject name formats that are available include

- **E-mail name** If the e-mail name field is populated in the Active Directory user object, that e-mail name will be used for user accounts.

| **Note** |
| --- |
| The e-mail name is required for user certificates. If the e-mail name is not populated for a user in Active Directory, the certificate request by that user will fail. |

- **DNS name** The fully qualified domain name (FQDN) of the subject that requested the certificate is used for computer accounts.

- **User principal name (UPN)** The user principal name is part of the Active Directory user object and will be used for user accounts.

- **Service principal name (SPN)** The service principal name is part of the Active Directory computer object and will be used for computer accounts.

If the Subject Name option is set to **Supply in the request**, you should set one or more Issuance Requirements for the template to prevent subjects from requesting certificates using any subject name. This would allow subjects to impersonate other subjects easily. Issuance requirements ensure that the certificate request is inspected and validated before the certificate is issued.

The only case in which a subject can request a certificate with a different subject name is when the request holds a certificate based on the Enrollment Agent template. The Enrollment Agent template allows a subject to request a certificate on behalf of another subject.

| **Note** |
| --- |
| No user interface exists on the client to supply a subject name in the certificate request. If you require the ability to provide the subject name, you must perform the certificate request programmatically using the XEnroll.dll ActiveX control. |

**The Issuance Requirements Tab**

The Issuance Requirements tab (Figure 4) allows higher assurance level certificates to be placed in a pending state until the certificate is reviewed before issuance. This allows for multiple signers of a CMC request to exist. For more information about CMC, see RFC 2797 Certificate Management Messages over CMS.




**Figure 4: The Issuance Requirements Tab**

When CA certificate manager approval is enabled, the certificate is placed into a pending state, rather than being issued immediately. In its pending state, the certificate request can be reviewed by certificate managers, ensuring a higher level of assurance for the issued certificate.

The following settings configure the authentication and signature requirements for issuance certificates that are based on a template:

- **CA certificate manager approval** All certificates are placed into the pending container for a certificate manager to issue or deny. Any defined certificate manager can issue or deny a certificate request of this type.

- **This number of authorized signatures** This setting requires the CMC certificate request to be digitally signed by one or more authorized parties before it can be issued.

| **Note** |
| --- |
| Defining the number of authorized signatures enables several other configuration parameters. |

- **Policy type required in signature** This option is only enabled when the number of authorized signatures value is set. This option defines which specific application policy, issuance policy, or both are required in the signing certificate. This is how the certification authority determines whether the signature is appropriate for authorizing the issuance of the subject's certificate.

- **Application policy** This option specifies the application policy that must be included in the signing certificate used to sign the certificate request. It is enabled when Policy type required in signature is set to either Application policy or both application and issuance policy.

- **Issuance policy** This option specifies the issuance policy that must be included in the signing certificate used to sign the certificate request. This option is enabled when Policy type required in signature is set to either Issuance policy or both application and issuance policy.

An example of where issuance requirements are defined is for the issuance of Cross-Certification Authority certificates. This certificate template requires that the signing certificate includes the Qualified Subordination Application Policy object identifier as shown in Figure 5.




**Figure 5: Issuance Requirements for a Cross-Certification Authority Certificate**

As you can see, the Cross-Certification Authority certificate request must be signed by a certificate containing the Qualified Subordination object identifier.

In addition, you can also determine whether the same issuance requirements are upheld for certificate renewal, or if the existence of a valid existing certificate of the same template in the certificate store meets the minimum requirements for certificate issuance.

| **Note** |
| --- |
| For information about including application and issuance policy object identifiers in an issued certificate, see The Extensions Tab section. |

**The Superseded Templates Tab**

The Superseded Templates tab (Figure 6) allows you to supersede existing certificate templates with a modified version 2 certificate template.




**Figure 6: The Superseded Templates Tab**

The Superseded Templates tab allows you to redeploy version 2 certificates and ensure that the updated certificates replace previous versions of the certificate or other certificate templates that were used for the same purposes. You can supersede

- A version 1 certificate template with a version 2 certificate template
- A version 2 certificate template with a version 2 certificate template
- Multiple existing certificate templates with a common version 2 certificate template

| **Note** |
| --- |
| You can force the application of the updated certificate template by forcing all certificate holders to re-enroll the updated certificate template. For more information, see the Re-Enroll Certificate Holders section. |

**The Extensions Tab**

The Extensions tab (Figure 7) allows you to define specific application policies, issuance policies, certificate subject types, and key usage attributes for a certificate template. The following sections detail the specifics on each form of extension defined in a certificate template.




**Figure 7: The Extensions Tab**

**Application Policies**

Application policies enable you to decide which certificates can be used for certain purposes. You can issue certificates widely without being concerned that they are misused for an unintended purpose.

Application policies are settings that inform a target that the subject holds a certificate that can be used to perform a specific task. They are represented in a certificate by an object identifier that is defined for a given application. This object identifier is included in the issued certificate. When a subject presents its certificate, it can be examined by the relying party to verify the application policy and determine if it can perform the requested action.

| **Note** |
| --- |
| Application policies are similar to the *extended key usage* attribute in version 1 certificate templates. Because some implementations of PKI applications may not understand application policies, both application policies and enhanced key usage sections appear in certificates issued by a Microsoft CA. |

The following table shows some of the more commonly used application policies and their related object identifiers.

| Application Policy | Object Identifier |
| --- | --- |
| Client Authentication | 1.3.6.1.5.5.7.3.2 |
| CA Encryption Certificate | 1.3.6.1.4.1.311.21.5 |
| Smart Card Logon | 1.3.6.1.4.1.311.20.2.2 |
| Document Signing | 1.3.6.1.4.1.311.10.3.12 |
| File Recovery | 1.3.6.1.4.1.311.10.3.4.1 |
| Key Recovery | 1.3.6.1.4.1.311.10.3.11 |
| Microsoft Trust List Signing | 1.3.6.1.4.1.311.10.3.1 |
| Qualified Subordination | 1.3.6.1.4.1.311.10.3.10 |
| Root List Signer | 1.3.6.1.4.1.311.10.3.9 |

**Certificate Policies**

Certificate policies, also referred to as issuance policies, define the level of assurance for an issued certificate. A CA processes

each certificate request by a defined set of rules. The certification authority may issue some certificates with no proof of identification and require subjects of another type to submit some proof. This provides different levels of assurance for different certificates. These levels of assurance are represented in certificates as issuance policies.

| **Note** |
| --- |
| Certificate policies refer to the certificate policies extension identifier described in RFC 2459. |

An issuance policy is a group of administrative rules that are implemented when issuing certificates. They are represented in a certificate by an object identifier that is defined at the certification authority. This object identifier is included in the issued certificate. When a subject presents its certificate, it can be examined by the target to verify the issuance policy and determine if that level of issuance policy is sufficient to perform the requested action.

An issuance policy describes the conditions under which a certificate is issued. This provides a level of assurance that the subject's certificate request was verified in a specific way.

One or more issuance policies may be selected for a certificate template. Because these issuance policies are simply text labels with an associated object identifier, no options are associated with them. The only special issuance policy is **All issuance policies**, which indicates that this policy includes all others. This is normally reserved for certificates held by certification authorities.

Issuance policies are often used when configuring Qualified Subordination (cross-certification) between PKI hierarchies to ensure that certificates recognized by another organization's PKI meet issuance requirements required by your organization.

**Certificate Template Information**
The certificate template information, also referred to as the Certificate subject type, defines the purpose of a certificate template. Six subject types are recognized by Windows Server 2003 CAs:

- Key recovery agent
- Directory e-mail replication
- Cross-certified certification authority
- Certification authority (CA)
- Computer
- User

The certificate template information extension cannot be edited. If you require a specific subject type to be applied to a certificate, you should clone from a certificate template that includes the required subject type. Some of the internal flags that are defined for specific subject types limit the display of the certificate template to computers or users. Choosing to clone an incorrect certificate template will prevent the certificate template from being displayed to the desired enrollment audience. For example, a computer certificate template cannot be enrolled for use by a user account.

**Key Usage**
A certificate enables the subject to perform a specific task. To help control the usage of a certificate outside its intended purpose, restrictions are automatically placed on certificates. Key usage is a restriction method and determines what a certificate can be used for. This allows the administrator to issue certificates that can only be used for specific tasks or certificates that can be used for a broad range of functions. If no key usage is specified, the certificate can be used for any purpose.

For signatures, key usage can be limited to one or more of the following purposes:

- Digital signature
- Signature is a proof of origin (nonrepudiation)
- Certificate signing

- CRL signing

For encryption key usage, the following options are available:

- Key exchange without key encryption

- Key exchange only with key encryption

**The Security Tab**
The Security tab (Figure 8) allows you to define the DACL for a specific certificate template. The permissions that you assign to the certificate template define which security principals can read, modify, enroll, or auto-enroll for a specific certificate template.



**Figure 8: The Security Tab**

The permissions that you can assign to a certificate template include

- **Full Control** This permission allows a security principal to modify all attributes of a certificate template, including the permissions for the certificate template.

- **Read** This permission allows a security principal to see the certificate template when enrolling for certificates. It is required for a security principal to enroll or auto-enroll a certificate; it is required by the certificate server to find the certificate templates in Active Directory.

- **Write** This permission allows a security principal to modify the attributes of a certificate template, including the permissions assigned to the certificate template.

- **Enroll** This permission allows a security principal to enroll for a certificate based on the certificate template. To enroll for a certificate, the security principal must also have Read permissions for the certificate template.

- **Autoenroll** This permission allows a security principal to receive a certificate through the auto-enrollment process. Auto-enrollment permissions require that the user has both Read and Enroll permissions in addition to the Auto-enroll permission.

**Note**

It is recommended that certificate template permissions be assigned only to global groups or universal groups. Because the certificate template objects are stored in the Configuration naming context, you cannot assign permissions using domain local groups. To simplify administration, it is never recommended to assign certificate template permissions to individual user or computer accounts.

It should be considered a best practice to keep Read permission assignment for the Authenticated Users group. This permission assignment allows all users and computers to view the certificate templates in Active Directory. This includes the CA running under the SYSTEM context of a machine account to view the certificate templates when issuing certificates to users and computers.

**Note**

A version 2 certificate template can be distributed to Windows Millennium Edition and Windows 2000 with Service Pack 2 clients through the Web enrollment pages.

# Deploying Certificate Templates (Implementing and Administering Certificate Templates in Windows Server 2003)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

After creating a version 2 certificate template, the next step is to deploy the certificate template. Deployment includes publishing the certificate template at one or more Certification Authorities, defining which security principals have Enroll permissions for the certificate template, and deciding whether to configure auto-enrollment for the certificate template.

**Publishing Certificate Templates**

Once a certificate template is created in the Certificate Templates MMC and the certificate template has replicated to all domain controllers in the forest, it can now be published for deployment. The main decision in publishing certificate templates is deciding which Certification Authority or Authorities will issue the certificates based on this certificate template.

| **Note** |
| --- |
| Remember that version 2 certificate templates can only be issued by Enterprise CAs running on Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition. |

An Enterprise CA can only issue certificates based on certificate templates that exist in the Certificate Templates container in the Certification Authority MMC console (Figure 9).




**Figure 9: Defining the Certificate Templates Published at the EntCA Certification Authority**

**Permission Design**

Certificate templates are published to the configuration naming context, which is stored on every domain controller in the forest in the path CN=Certificate Templates, CN=Public Key Services,CN=Services,CN=Configuration,DC=*ForestRootDomain*. Each certificate template exists as an object in the configuration naming context and has an associated DACL, which defines what specific operations a security principal can do with the certificate.

Use the following recommendations for permissions assignments:

- Assign permissions only to global groups or to universal groups. Certificate templates are defined in the configuration naming context. Permissions should only be assigned to global groups or to universal groups. It is not recommended to assign permissions to domain local groups. They are only recognized in the domain where the domain local group exists and can result in inconsistent application of permissions. It is recommended to never assign permissions directly to an individual user or computer account.

- To enable auto-enrollment, a user or computer must belong to domain groups that are granted Read, Enroll, and Autoenroll permissions.

- To enable enrollment via the Certificates MMC console, through Web-based enrollment, or through auto-renewal, set either domain or universal groups with **Read** and **Enroll** permissions.

- For certificate renewal, a user or computer must belong to a domain security group with Read and Enroll permissions. This is true whether the certificate is manually renewed or the renewal is implemented using auto-enrollment.

- Restrict Write and Full Control permissions to CA managers to ensure that the templates are not improperly configured.

**Auto-Enrollment Considerations**

Auto-enrollment allows a CA administrator to configure users and computers to automatically enroll for certificates, retrieve issued certificates, and renew expiring certificates without requiring subject interaction. The auto-enrollment process can be configured to function as a background task that does not require any user input.

To properly configure subject auto-enrollment, the administrator must plan the configuration of the certificate template that will use auto-enrollment. Several settings in the certificate template directly affect the behavior of subject auto-enrollment, as follows:

- **User input requirements** On the Request Handling tab (Figure 2) of the desired certificate template, the Require user input for auto-enrollment check box changes the behavior of auto-enrollment. When this check box is selected, subjects are prompted for any necessary information for obtaining or renewing a certificate. When this check box is cleared, auto-enrollment operates silently, without any notice to the subject.

| **Note** |
| --- |
| Because smart cards must prompt the user for their PIN, this option must be selected when a smart card CSP is selected. |

- **Limit the number of CSPs for an auto-enrollment certificate template** If more than one smart card CSP is selected on the Request Handling tab (Figure 2), users may receive more than one dialog box when a Windows XP client retrieves the auto-enrolled certificate and begins to install it on the smart card. It is recommended that only one CSP be selected from this list for each template.

- **Do not enable subject creation based on request information** If the Supply in the request option is enabled on the Subject Name tab (Figure 3), auto-enrollment is disabled. This is because enabling the option prompts the subject to interactively create the subject name in the request, which will not work with auto-enrollment.

- **Do not require more than one authorized signature for issuance** Configuring more than one authorized signature in the Issuance Requirements tab (Figure 4) of the desired certificate template disables subject auto-enrollment based on this template. If the authorized signatures value is set to one, the requestor must sign the request with a private key from a valid certificate in their certificate store. This certificate must contain the application and/or issuance policies specified in the Application policy and Issuance policies lists on the same tab. If an appropriate certificate exists in the requestor's certificate store, auto-enrollment signs the request with this certificate's private key and will obtain and install the requested certificate automatically. For example, to increase the security for EFS certificate distribution, you can create a custom version 2 certificate template that requires the existence of the Smart Card Logon OID (1.3.6.1.4.1.311.20.2.2) in the signing certificate's Application policy.

- On the Issuance Requirements tab (Figure 4) of the desired certificate template, the Valid existing certificate option may affect subject auto-enrollment. This option tells the CA that the subject does not need to meet issuance requirements when renewing a valid certificate. Subjects who may have been unable to auto-enroll for the initial certificate may be able to use auto-enrollment to renew that certificate. For example, a user whose distinguished name has changed may still auto-enroll a certificate based on an existing valid certificate, rather than on their new credentials.

| **Note** |
| --- |
| For more information about configuring auto-enrollment, see the Certificate Autoenrollment in Windows XP white paper at http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/certenrl.mspx |

**Delegating Template Management**

You can delegate the ability to manage individual certificate templates or to create any certificate templates by defining appropriate permissions to global groups or universal groups that a user belongs to.

To delegate modification of a specific certificate template, assign a global or universal group the Read and Write permissions on the Security tab (Figure 8) of the certificate template. If you also wish to delegate administration permissions to the security groups, assign Full Control permissions so that the security group can modify the DACL for the certificate template.

To delegate the ability to create certificate templates to users who are not members of the Domain Admins group in the forest root domain, or members of the Enterprise Admins group, it is necessary to define the appropriate permissions in the Configuration naming context of the Active Directory.

To delegate the administration of all certificate templates, including the ability to duplicate and create new certificate temples, you must make the following permission assignments to a global or universal group that the user is a member of:

- CN=Certificate Templates,CN=Public Key Services,CN=Services,CN=Configuration,DC=*ForestRoot* container = Full Control permissions

- Full Control permissions to every certificate template in the CN=Certificate Templates,CN=Public Key Services,CN=Services,CN=Configuration,DC=*ForestRoot* container. The permissions assigned to the Certificate Templates container are not inherited by the individual Certificate Templates.

- Full control permissions to the CN=OID,CN=Public Key Services,CN=Services,CN=Configuration,DC= *ForestRoot* container.

| **Note** |
| --- |
| To create or duplicate existing certificate templates, users only need the Create Child permission for the CN= Certificate Templates,CN=Public Key Services,CN=Services,CN=Configuration,DC=*ForestRoot* and CN=OID, CN=Public Key Services,CN=Services,CN=Configuration,DC= *ForestRoot* containers. |

# Planning Considerations

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

This section looks at the topics that must be considered when planning all aspects of version 2 certificate templates. The section provides design guidelines for the creation of certificate templates and best practices to implement when working with certificate templates.

**Design Guidelines**

When creating a version 2 certificate template, consider the following design guidelines:

- **Defining the subject name** The holder of the private key associated with a certificate is known as the *subject*. This can be a user, a computer, a program, or virtually any object or service. Because the subject can vary greatly depending on who or what it is, you need some flexibility when providing the subject name in the certificate request. A Windows Server 2003 CA can either build the subject name automatically or request it from the subject. If it automatically provides the subject name, it obtains the information from Active Directory. You can configure this process to include or exclude information that is useful in the environment. If it is configured to manually provide the subject name, the subject supplies that information in the certificate request using the Web-based enrollment pages.

- **Defining the certificate lifetime** Certificate-based cryptography uses public-key cryptography to protect and digitally sign data. Over time, it is theoretically possible to collect data protected with the public key and attempt to derive the private key from it. Given enough time and resources, this private key could be compromised, effectively rendering all protected data unprotected. Because certificates can be compromised over time, a finite certificate lifetime should be established.

- **Determining certificate usage** It is possible to issue many specific certificates that can only be used for a single purpose or to issue fewer certificates that have broad usage. This decision depends on the environment, the level of administration desired, and the possible effects on the subjects, as well as the effects of multiple certificates on applications that will use them.

  One strategy of certificate administration creates a number of specific templates one for each job function, such as file encryption or code signing. Subjects can then enroll for each certificate as needed for the appropriate function. This allows subjects to start with few certificates and only obtain new certificates that they need over time. The drawback to this strategy is that the subject may end up with a large number of certificates and private keys that become harder to manage over time.

  Alternatively, you could create a few broad certificate templates that encompass job functions for most common groups of subjects. For example, if most employees use their certificates for mail signing and encryption as well as file encryption, create one template that allows all those functions in the same certificate. This allows most subjects to obtain a single, all-purpose certificate. The drawback to this strategy is that there is no detailed control of the usage of the certificates. The administrator cannot decide that subgroups cannot encrypt mail without modifying the template or changing the strategy.

- **Determining which CSP to implement** A version 2 certificate template allows you to define one or more cryptographic service providers (CSPs) as usable by a template. This allows the administrator to control what types of cryptography subjects can use within an enterprise. This is very useful when security is of paramount importance. Because subjects use the CSP for both portions of any cryptographic serviceeither encryption and decryption or signing and confirming signatureit is necessary to ensure that all subjects can use the same CSP. The easiest way to do this is to configure each certificate template to identify exactly one CSP. Which CSP should be identified is up to the administrator and depends on the level of security required, the intended purposes of the certificate, and the presence of security hardware, such as smart cards.

- **Determining key length** Each CSP provides one or more cryptographic algorithms for encryption or digital signature. You can define a minimum key size allowed for a certificate template. In general, larger keys provide more protection over shorter keys for the same algorithm, but larger keys take longer to generate and use. You should select a minimum key size that ensures the necessary amount of protection without affecting performance.

- **Determining smart card usage** Each type of smart card has at least one associated CSP that must be implemented by the certificate template to allow the smart card to be used. If the correct smart card CSP is not associated with the template, the smart card will not be recognized and the template will fail. Ensure that you enable all smart card CSPs for the smart cards deployed in your environment within the certificate template.

- **Planning deployment methods** Certificates are typically deployed manually or automatically. Manual enrollment can take place using either the Web enrollment pages, the Certificates MMC console, or through CryptoAPI or CAPICOM programming interfaces. Auto-enrollment requires the configuration described in the Auto-Enrollment Considerations section.

- **Planning Key Archival** Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition CAs can provide key archival of private keys. When planning key archival settings for a certificate template, consider the following settings:

    - **Enabling archival of the subject's private key** This option is only available when the issuing CA is running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, and the CA is configured for key archival.

    - **Define whether the private key can be exported** If this option is enabled, the subject can export the private key for backup purposes or move the private key and certificate to another computer. If key archival is centralized, you may not want to enable this option because it allows the key to be recovered in a decentralized manner.

**Best Practices**

When planning certificate template deployment, use the following best practices:

- Do not delete the Certificate Publishers security group.

  The Certificate Publishers security group contains each certification authority's computer account and is used when publishing certificate templates to Active Directory. If this group is removed, the certification authority may not publish certificates to Active Directory correctly. To avoid this, the group should not be deleted and its membership should not be modified.

- Add the CA machine accounts to every Cert Publishers group.

  The Cert Publishers group is a domain local group that exists in every domain in the forest. In each domain, all CA computer accounts should be added to the Cert Publishers group.

- Do not exceed the certificate lifetime of the issuing certification authority.

  Certificate lifetimes work as a subset of the certification authority's (CA) certificate lifetime. All certificates, including the CA certificate, have an expiration date after which they are no longer valid. As a result, a certificate cannot be issued with a lifetime that exceeds the lifetime of the issuing CA. Issuing such a certificate would allow it to be valid for longer than the issuing CA certificate, which violates certificate chaining rules. A CA will therefore continue to issue certificates until the CA's certificate expires or until the requested template's renewal period is greater than the CA's certificate remaining lifetime. If a certificate template requires a lifetime greater than the lifetime of the CA certificate, the validity period of the issued certificate is truncated to the amount of time left in the lifetime of the CA certificate.

- Plan certificate templates before deployment.

  Certificates can be issued to subjects in many ways, including manual enrollment, auto-enrollment, and Web enrollment. In addition, there are many certificate strategies including issuing one all-inclusive certificate to all subjects and issuing several application-specific certificates to subjects as needed. Because there are so many options, planning should be done well in advance of certificate deployment.

- Upgrade the certificate templates in Active Directory before upgrading from Windows 2000.

  Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition certification authorities use

newer versions of certificate templates than those used by Windows 2000 CAs. These templates must be upgraded before the certification authorities are upgraded to ensure proper operation.

- Duplicate new templates from existing templates closest in function to the intended template.

  New certificate templates are duplicated from existing templates. Many settings are copied from the original template. Because of this, duplicating one template to another of a totally different type may carry over some unintended settings. When duplicating a template, examine the subject type of the original template and ensure that you duplicate one that has a similar function to that of the intended template. Although most settings for certificate templates can be edited once the template is duplicated, the subject type cannot be changed.

- Determine publication points for certificate templates.

  Determine which CAs will issue specific certificate templates based on both the administration model implemented by the organization and the usage of the certificate template. For example, if the administration model is a project-based management model, it would be appropriate to only assign the certificate template to CAs associated with the project. On the other hand, if the certificate template is widely used throughout the organization, it may be appropriate to have all CAs issue the certificate template to provide fault tolerance if a CA is unavailable.

- Minimize the number of issued certificates.

  Consider using multi-purpose certificates that can be used for more than one job task rather than issuing separate certificates for each job task that must be performed. This reduces the number of issued certificates and reduces the complexity when a user must select which certificate to present to an application.

# Walkthroughs (Implementing and Administering Certificate Templates in Windows Server 2003)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

**Creating a Version 2 Certificate Template**

To create a new version 2 certificate template

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.
2. Open the Certificate Templates MMC console (Certtmpl.msc).
3. In the details pane, right-click an existing certificate that will serve as the starting point for the new certificate, and then click Duplicate Template.
4. On the General tab, enter the Template display name and the template name, and then click OK.
5. Define any additional attributes for the newly created version 2 certificate template.

**Defining Application and Issuance Policies**

When you create a version 2 certificate template, you can define which application and issuance policies are included in the issued certificates. Defining application and issuance policies requires the completion of three tasks:

- Acquiring Object Identifiers for the application and issuance policies
- Defining the application and issuance policies
- Mapping issuance policies between PKI hierarchies

**Acquiring Object Identifiers**

If you define a custom application policy or issuance policy, you must obtain an object identifier for the policy.

To acquire an object identifier

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task as described in the Delegating Template Management section.
2. Open the Certificate Templates MMC console (Certtmpl.msc).
3. In the details pane, right-click the certificate template you wish to modify, and then click Properties.
4. On the Extensions tab, click Application Policies, and then click Edit.
5. In the Edit Application Policies Extension dialog box, click Add.
6. In Add Application Policy, ensure that the application you are creating does not exist, and then click New.
7. In the New Application Policy dialog box, provide the name for the new application policy, note the generated object identifier, and then click OK.

**Note**

You could also add new object identifiers by editing Certificate Policies rather than Application Policies.

**Establishing Application Policies**

Once you have defined any custom application policies, you can then associate the application policy with the certificate template using the following procedure:

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.
2. Open the Certificate Templates MMC console (Certtmpl.msc).
3. In the details pane, right-click the certificate template you want to change, and then click Properties.
4. On the Extensions tab, click Application Policies, and then click Edit.
5. In Edit Application Policies Extension, click Add.
6. In Add Application Policy, click the desired application policy, and then click OK.

**Establishing issuance Policies**

Once you have defined any custom issuance policies, you can then associate the issuance policy with the certificate template using the following procedure:

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.
2. Open the Certificate Templates MMC console (Certtmpl.msc).
3. In the details pane, right-click the certificate template you want to change, and then click Properties.
4. On the Extensions tab, click Certificate Policies, and then click Edit.
5. In Edit Issuance Policies Extension, click Add.
6. In Add Issuance Policy, click New.
7. Provide the requested information.

**Mapping Issuance Policies between PKI Hierarchies**

When performing qualified subordination, it may be necessary to associate issuance policies in your organization with issuance policies defined in another organization. The policy mappings are defined in the policy.inf file used to generate the Cross-Certification Authority certificate.

In the policy.inf file, you must include a [PolicyMappingsExtension] that maps the policies listed in the policy.inf file with policies defined in the other PKI hierarchy. The following code example shows a section of a policy.inf file that maps issuance policies for high, medium, and low assurance between two organizations.

```
[PolicyStatementExtension]
Policies = HighAssurancePolicy, MediumAssurancePolicy, LowAssurancePolicy
CRITICAL = FALSE
[HighAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.256.257.258.259.1.402
[MediumAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.256.257.258.259.1.401
[LowAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.256.257.258.259.1.400
[PolicyMappingsExtension]
1.3.6.1.4.1.311.21.8.256.257.258.259.1.400 =
```

```
1.3.6.1.4.1.311.21.8.354.232.582.111.1.400
1.3.6.1.4.1.311.21.8.256.257.258.259.1.401 =
1.3.6.1.4.1.311.21.8.354.232.582.111.1.401
1.3.6.1.4.1.311.21.8.256.257.258.259.1.402 =
1.3.6.1.4.1.311.21.8.354.232.582.111.1.402
critical = yEs
```

This example maps the object identifiers for the high assurance, medium assurance, and low assurance policies to object identifiers that exist in the other organization's PKI. The other organization must define a policy.inf file that maps the object identifiers in the opposite direction so that the object identifiers are recognized by both organizations.

**Configuring Permissions for a Certificate Template**

This section explains how to define permissions for specific certificate templates and also for delegating permission for the management of certificate templates.

**Allowing for Enrollment**

To define permissions to allow a specific security principal to enroll for certificates based on a certificate template

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.

2. Open the Certificate Templates MMC console (Certtmpl.msc).

3. In the details pane, right-click the certificate template you want to change, and then click Properties.

4. On the Security tab, ensure that Authenticated users is assigned Read permissions.

   This ensures that all authenticated users on the network can see the certificate templates.

5. On the Security tab, click Add. Add a global group or universal group that contains all security principals requiring Enroll permissions for the certificate template, and then click OK.

6. On the Security tab, select the newly added security group, and then assign Allow permissions for the Read and Enroll permissions.

7. Click OK.

**Allowing for Auto-Enrollment**

To define permissions to allow a specific security principal to auto-enroll for certificates based on a certificate template

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.

2. Open the Certificate Templates MMC console (Certtmpl.msc).

3. In the details pane, right-click the certificate template you want to change, and then click Properties.

4. On the Security tab, click Add. Add a global group or universal group that contains all security principals requiring Enroll permissions for the certificate template, and then click OK.

5. On the Security tab, select the newly added security group, and then assign Allow permissions for the Read, Enroll, and Auto-enroll permissions.

6. Click Apply.

| **Note** |
| --- |
| For more information about configuring certificate auto-enrollment, see the Certificate Autoenrollment in Windows XP white paper at http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/certenrl.mspx |

**Allowing Creation and Modification of any Certificate Template**

To delegate administration of all templates (which includes the ability to duplicate and create new templates)

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.

2. Open the ADSIEdit console (Adsiedit.msc).

3. In the console tree, right-click ADSI Edit, and then click Connect to.

4. In the Connection dialog box, in the Connection Point section, click Naming Context, select Configuration Container from the list below Naming Context, and then click OK.

5. In the console tree, expand ADSI Edit.

6. In the console tree, expand Configuration Container.

7. In the console tree, expand CN=Configuration,DC=*ForestRootDomain* (where *ForestRootDomain* is the LDAP distinguished name of your forest root domain).

8. In the console tree, expand CN=Services.

9. In the console tree, expand CN=Public Key Services.

10. In the console tree, right-click CN=Certificate Templates, and then click Properties.

11. In the CN=Certificate Templates Properties dialog box, on the Security tab, click Add. Add a global or universal group that contains the users you wish to delegate certificate creation and management permissions to, and then click OK.

12. On the Security tab, select the newly added security group, ensure that the security group is assigned Allow permissions for the Full Control permission, and then click OK.

13. In the console tree, right-click CN=OID, and then click Properties.

14. In the CN=OID Properties dialog box, on the Security tab, click Add. Add a global or universal group that contains the users you wish to delegate certificate creation and management permissions to, and then click OK.

15. On the Security tab, select the newly added security group, ensure that the security group is assigned Allow permissions for the Full Control permission, and then click OK.

16. Close ADSI Edit.

17. Ensure that the security group assigned full control permissions to the CN=Certificate Templates and CN=OID containers is also assigned full control permissions for all certificate templates listed in the Certificate Templates MMC console (Certtmpl.msc).

**Publishing a Certificate Template**

To define which certificate templates are issued by a Certification Authority

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.

2. From Administrative Tools, open the Certification Authority MMC console.

3. In the console tree, expand *CAName* (where *CAName* is the name of your Enterprise CA).

4. In the console tree, select the Certificate Templates container.

| Note |
| --- |
| If working with a Windows 2000 CA, the container is named Policy Settings. |

5. Right-click Certificate Templates, and then click New, Certificate Template to Issue.

6. In the Enable Certificate Templates dialog box, select the certificate template(s) you wish the CA to issue, and then click OK.

| Note |
| --- |
| If a certificate template is not listed in the Enable Certificate Templates dialog box, the CA is either already configured to issue the certificate template, or replication of the certificate template is not completed to all domain controllers in the forest. |

The newly selected certificate template(s) will appear in the details pane.

**Removing a Certificate Template from a CA**
Removing a certificate template only unlinks a certificate from a CA instead of deleting it physically from the certificate template store.

To remove a certificate template from the certificate templates currently issued by a Certification Authority

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.

2. From Administrative Tools, open the Certification Authority MMC console.

3. In the console tree, expand *CAName* (where *CAName* is the name of your Enterprise CA).

4. In the console tree, select the Certificate Templates container.

| Note |
| --- |
| If working with a Windows 2000 CA, the container is named Policy Settings. |

5. In the details pane, right-click the certificate template you wish to remove from the CA, and then click Delete.

6. In the Disable Certificate Templates dialog box, click Yes.

The certificate template no longer appears in the details pane.

**Replace an Existing Certificate Template with a New Certificate Template**
This process, also referred to as superseding an existing template, defines which existing templates a version 2 certificate is replacing.

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.

2. Open the Certificate Templates MMC console (Certtmpl.msc).

3. In the details pane, right-click the certificate template you want to change, and then click Properties.

4. Click the Superseded Templates tab.

5. Click Add.

6. Click one or more templates to supersede, and then click OK.

**Re-Enroll Certificate Holders**

If you make modifications to a certificate template that you wish implemented immediately for all existing certificate holders, you can force re-enrollment by using the following procedure:

1. Log on as a member of the Enterprise Admins or the forest root domain's Domain Admins group, or as a user who has been granted permission to perform this task.

2. Open the Certificate Templates MMC console (Certtmpl.msc).

3. In the details pane, right-click the certificate template that you wish to re-enroll for all certificate holders, and then click Reenroll all Certificate Holders.

# Appendix A: Wireless Certificates

Updated: January 17, 2011

Applies To: Windows Server 2003 with SP1

Windows XP introduced native support for 802.1x and wireless networks. To enable strong security, both users and computers need authentication certificates to authenticate to a RADIUS (IAS Server) authorization point. Windows 2000 certification authorities support 802.1x certificate requirements for computers with the version 1 Machine certificate template and user certificates with any of the certificate templates that contain the Client Authentication EKU. If version 2 templates are used for computer auto-enrollment, it is important to configure the certificate template properly. When the Machine template is cloned to a version 2 template, the administrator MUST ensure that the DNS name is included in the subject name (CN) of the certificate. The Windows XP wireless client requires the DNS name of the computer to be contained in the subject for proper usage and authentication to the IAS server (RADIUS).

| **Important** |
| --- |
| If the DNS fully qualified domain name is longer than 64 characters, the name will be truncated during certificate enrollment and the name will be invalid for wireless authentication. |

For more information, see the white paper at Securing Wireless LANs with Certificate Services (http://go.microsoft.com/fwlink/?LinkId=14843)

# Appendix B: References

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1


- Troubleshooting Certificate Status and Revocation
  http://www.microsoft.com/technet/security/topics/cryptographyetc/tshtcrl.mspx

- Certificate Autoenrollment in Windows XP
  http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/certenrl.mspx

- PKI Enhancements in Windows XP Professional and Windows Server 2003
  http://www.microsoft.com/technet/prodtechnol/winxppro/plan/pkienh.mspx

- Data Protection and Recovery in Windows XP
  http://www.microsoft.com/technet/prodtechnol/winxppro/deploy/cryptfs.mspx

- Windows Server 2003, Enterprise Edition Help documentation

# Appendix C: Certificate Template Schema Additions

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The Certificate Templates container contains the certificate templates that are defined within an Active Directory forest. Each certificate template is of the class pKICertificate. Each Certificate Template is managed by using the Certificate Templates MMC snap-in. Windows 2000 includes 24 default certificate templates; Windows Server 2003 includes 29 default templates. Each template is stored in the following location in the Configuration naming context:

CN=<name of template>,CN=Certificate Templates,CN=Public Key Services,CN=Services,CN=Configuration,DC=*ForestRootDomain*

**Version 1 Certificate Template Attributes**

The following version 1 certificate templates attributes are defined in the Active Directory schema.

| Attribute | Description |
|---|---|
| **Cn** | Common name of the certificate type. |
| **distinguishedName** | Distinguished name of the certificate type. |
| **displayName** | Display name of a cert type. |
| **pKIExtendedKeyUsage** | Array of enhanced key usage object identifiers. |
| **pKIDefaultCSPs** | Default CSP list. DWORD, CSP name. |
| **pKICriticalExtensions** | List of critical extensions. |
| **revision** | Major version of the templates. |
| **templateDescription** | Obsolete attribute. |
| **flags** | General enrollment flags. |
| **pKIDefaultKeySpec** | Specifications of the Default Key length and construct. |
| **NTSecurityDescriptor** | Security Descriptor name. |
| **pKIKeyUsage** | Key Usage extension. |
| **pKIMaxIssuingDepth** | Basic Constraints. DWORD value. |
| **pKIExpirationPeriod** | Validity period. Negative FILETIME value. |
| **pKIOverlapPeriod** | Renewal period. Negative FILETIME value. |

**Version 2 Certificate Template Attributes**

The following version 2 certificate templates attributes are defined in the Active Directory schema.

| Attribute | Description |
|---|---|
| **msPKI-Template-Schema-Version** | Schema version of the templates. |
| **msPKI-Template-Minor-Revision** | Minor version of the templates. |
| **msPKI-RA-Signature** | Number of RA signatures required on a request referencing this template. |
| **msPKI-Minimal-Key-Size** | Minimal key size required. |

| | |
|---|---|
| **msPKI-Template-Cert-Template-OID** | Object identifier of this template. |
| **msPKI-Supersede-Templates** | Name of the template that this template supersedes. |
| **msPKI-RA-Policies** | RA issuer policy object identifiers required. |
| **msPKI-RA-Application-Policies** | RA application policy object identifiers required. |
| **msPKI-Certificate-Policy** | The certificate issuer policy object identifiers are placed in the OID_CERT_POLICIES extension by the policy module. |
| **msPKI-Certificate-Application-Policy** | Certificate application policy object identifiers. |
| **msPKI-Enrollment-Flag** | Enrollment flags. |
| **msPKI-Private-Key-Flag** | Private key flags. |
| **msPKI-Certificate-Name-Flag** | Subject name flags. |

**Flags**

The following enrollment flags are defined in the Active Directory schema.

| Flag | Description |
|---|---|
| **CT_FLAG_INCLUDE_SYMMETRIC_ALGORITHMS**<br>• **0x00000001** | Include the S/MIME symmetric algorithms in the requests. |
| **CT_FLAG_PEND_ALL_REQUESTS**<br>• **0x00000002** | All certificate requests are pended. |
| **CT_FLAG_PUBLISH_TO_KRA_CONTAINER**<br>• **0x00000004** | Publish the certificate to the KRA (key recovery agent container) in Active Directory. |
| **CT_FLAG_PUBLISH_TO_DS**<br>• **0x00000008** | Publish the resultant certificate to the userCertificate property on the user object in Active Directory. |
| **CT_FLAG_AUTO_ENROLLMENT_CHECK_USER_DS_CERTIFICATE**<br>• **0x00000010** | The Auto-enrollment client will not enroll for a new certificate if the user has a certificate previously published to the userCertificate attribute in Active Directory with the same template name. |
| **CT_FLAG_AUTO_ENROLLMENT**<br>• **0x00000020** | This cert is appropriate for auto-enrollment. |
| **CT_FLAG_PREVIOUS_APPROVAL_VALIDATE_REENROLLMENT**<br>• **0x00000040** | A previously issued certificate will valid subsequent enrollment requests. |

| | |
|---|---|
| **CT_FLAG_DOMAIN_AUTHENTICATION_NOT_REQUIRED**<br>• **0x00000080** | Obsolete. |
| **CT_FLAG_USER_INTERACTION_REQUIRED**<br>• **0x00000100** | User interaction is required to enroll using auto-enrollment. |
| **CT_FLAG_ADD_TEMPLATE_NAME**<br>• **0x00000200** | Obsolete. |
| **CT_FLAG_REMOVE_INVALID_CERTIFICATE_FROM_PERSONAL_STORE**<br>• **0x00000400** | Remove invalid (expired or revoked) certificate from personal store on the local client computer during auto-enrollment. |

The following subject name flags are defined in the Active Directory schema.

| Flag | Description |
|---|---|
| **CT_FLAG_ENROLLEE_SUPPLIES_SUBJECT**<br>• **0x00000001** | The enrolling application must supply the subject name. |
| **CT_FLAG_ENROLLEE_SUPPLIES_SUBJECT_ALT_NAME**<br>• **0x00010000** | The enrolling application must supply the subjectAltName in request. |
| **CT_FLAG_SUBJECT_REQUIRE_DIRECTORY_PATH**<br>• **0x80000000** | Subject name should be full distinguished name based on the Active Directory path. |
| **CT_FLAG_SUBJECT_REQUIRE_COMMON_NAME**<br>• **0x40000000** | Subject name should be the common name. |
| **CT_FLAG_SUBJECT_REQUIRE_EMAIL**<br>• **0x20000000** | Subject name includes the e-mail name. |
| **CT_FLAG_SUBJECT_REQUIRE_DNS_AS_CN**<br>• **0x10000000** | Subject name includes the DNS name as the common name. |
| **CT_FLAG_SUBJECT_ALT_REQUIRE_DNS**<br>• **0x08000000** | Subject alt name includes the DNS name. |
| **CT_FLAG_SUBJECT_ALT_REQUIRE_EMAIL**<br>• **0x04000000** | Subject alt name includes the e-mail name. |

| | |
|---|---|
| **CT_FLAG_SUBJECT_ALT_REQUIRE_UPN**<br>• **0x02000000** | Subject alt name requires UPN. |
| **CT_FLAG_SUBJECT_ALT_REQUIRE_DIRECTORY_GUID**<br>• **0x01000000** | Subject alt name requires directory GUID (used by domain controllers). |
| **CT_FLAG_SUBJECT_ALT_REQUIRE_SPN**<br>• **0x00800000** | Subject alt name requires SPN (service principal name). |

The following template private key flags are defined in the Active Directory schema.

| **Flag** | **Description** |
|---|---|
| - Private Key Flags | |
| CT_FLAG_ALLOW_PRIVATE_KEY_ARCHIVAL<br>• 0x00000001 | Archival of the private key is allowed/required. |
| CT_FLAG_EXPORTABLE_KEY<br>• 0x00000010 | Mark the private key as exportable. |

The following template general flags are defined in the Active Directory schema:

| **Flag** | **Description** |
|---|---|
| CT_FLAG_MACHINE_TYPE<br>• 0x00000040 | Machine cert type. |
| CT_FLAG_IS_CA<br>• 0x00000080 | CA certificate type. |
| CT_FLAG_IS_CROSS_CA<br>• 0x00000800 | Cross-CA certificate type. |
| CT_FLAG_IS_DEFAULT<br>• 0x00010000 | Default cert type that is set on all V1 templates that cannot be changed. |
| CT_FLAG_IS_MODIFIED<br>• 0x00020000 | The type has been modified (read only). |
| CT_MASK_SETTABLE_FLAGS<br>• 0x0000ffff | Obsolete. |

# Server Clusters: Quorum Options - Windows Server 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft Corporation

Published: March 2003

**Abstract**
The goal of this document is to explain the two types of quorums available in Windows Server 2003 clusters – standard and majority node set (MNS) clusters. Currently, the standard quorum that is used by the Microsoft Cluster Service (MSCS) in Microsoft® Windows NT® Server 4.0 and Windows® 2000 Server family of products is sufficient for the majority of situations users will encounter. With the introduction of Windows Server 2003, Microsoft has introduced the MNS quorum, which is primarily for advanced users and other software vendors to create more sophisticated end-to-end solutions.

**In This White Paper**

Background (Server Clusters: Quorum Options - Windows Server 2003)

Types of Quorums

Quorum Uses

Frequently Asked Questions (Server Clusters: Quorum Options - Windows Server 2003)

# Background (Server Clusters: Quorum Options - Windows Server 2003)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before giving an explanation of what a cluster quorum is, it is important to understand the background to Windows Clustering Technologies.

Starting with Microsoft® Windows NT® 4.0 Enterprise Edition, Microsoft introduced the idea of a cluster, which is simply a group of servers that are presented as one virtual server. For example, you can configure two servers, server A and server B, in a cluster, and present them to the outside world as server C (a virtual server). If, for example, server A dies, server B is used to ensure that the virtual server (server C) and the services it offered, are still available to clients, thereby providing transparent access to the user.

Microsoft offers two types of clustering solutions in its Windows server products:

Server clusters – uses a shared-nothing architecture, which means that a resource can be active on only one server in the cluster at any one time. Because of this, it is well suited to applications that maintain some sort of state (for example, a database).

Network Load Balancing (NLB) – uses a load balancing architecture, which means that a resource can be active on ALL servers in the cluster at any one time. Because of this, it is well suited to applications that *do not* maintain state (for example, a Web server).

Quorums are only used in the MSCS now known as server cluster. Hence, any references to clusters in the remainder of this document are referring to the server clusters. The following sections explain what a cluster quorum is, and why it is important.

**What Is It?**
Each cluster has a special resource known as the *quorum resource*. A quorum resource can be any resource that does the following:

- Provides a means for arbitration leading to membership and cluster state decisions.
- Provides physical storage to store configuration information.

A quorum log is simply a configuration database for the server cluster. It holds cluster configuration information such as which servers are part of the cluster, what resources are installed in the cluster, and what state those resources are in (for example, online or offline). The quorum log is located by default in \MSCS\quolog.log.

**Why Is It Important?**
There are two main reasons why the quorum is important in a cluster. These are outlined below.

Consistency

Since the basic idea of a cluster is multiple physical servers acting as a single virtual server, it is critical that each of the physical servers have a consistent view of how the cluster is configured. The quorum acts as the definitive repository for all configuration information relating to the cluster. In the event that the Cluster service is unable to read the quorum log, it will not start, as it is not able to guarantee that the cluster will be in a consistent state, which is one of the primary requirements for a cluster.

Tie-breaker

The quorum is used as the tie-breaker to avoid “split-brain” scenarios. A split-brain scenario happens when all of the network communication links between two or more cluster nodes fail. In these cases, the cluster may be split into two or more partitions that cannot communicate with each other. The quorum is used to guarantee that any cluster resource is only brought online on only one node. It does this by allowing the partition that “owns” the quorum to continue, while the other partitions are evicted from the cluster.

# Types of Quorums

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Standard Quorum**
As mentioned above, a quorum is simply a configuration database for MSCS, and is stored in the quorum log file. A standard quorum uses a quorum log file that is located on a disk hosted on a shared storage interconnect that is accessible by all members of the cluster.

| **Note** |
| --- |
| It is possible to configure Server clusters to use the local hard disk on a server to store the quorum, but this is only supported for testing and development purposes, and should not be used in a production environment. Each member connects to the shared storage through some type of interconnect (for example, SCSI or Fibre Channel), with the storage consisting of either external hard disks (usually configured as RAID disks), or a storage area network (SAN), where logical slices of the SAN are presented as physical disks. |

| **Note** |
| --- |
| It is important that the quorum uses a physical disk resource, as opposed to a disk partition, as the entire physical disk resource is moved during failover. |

Standard quorums are available in Windows NT 4.0 Enterprise Edition, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, Windows Server 2003 Enterprise Edition and Windows Server 2003 Datacenter Edition.

The following is a sample diagram of a standard quorum in a four-node cluster.




*Figure 1: Standard quorum in a four-node cluster*

**Majority Node Set Quorums**
A majority node set (MNS) quorum is a single quorum resource from a server cluster perspective. However, the data is actually stored by default on the system disk of each member of the cluster. The MNS resource takes care to ensure that the cluster configuration data stored on the MNS is kept consistent across the different disks.

Majority node set quorums are available in Windows Server 2003 Enterprise Edition, and Windows Server 2003 Datacenter Edition.

The following is a sample diagram of an MNS quorum in a four-node cluster.




*Figure 2: MNS quorum in a four-node cluster*

While the disks that make up the MNS could in theory be disks on a shared storage fabric, the MNS implementation that is provided as part of Windows Server 2003 uses a directory on each node's local system disk to store the quorum data. If the configuration of the cluster changes, that change is reflected across the different disks. The change is only considered to have been committed, that is, made persistent, if that change is made to:

(<Number of nodes configured in the cluster>/2) + 1

This ensures that a majority of the nodes have an up-to-date copy of the data. The cluster service itself will only start up and therefore bring resources online if a majority of the nodes configured as part of the cluster are up and running the cluster service. If there are fewer nodes, the cluster is said not to have quorum and therefore the cluster service waits (trying to restart) until more nodes try to join. Only when a majority or quorum of nodes are available, will the cluster service start up, and bring the resources online. In this way, because the up-to-date configuration is written to a majority of the nodes regardless of node failures, the cluster will always guarantee that it starts up with the latest and most up-to-date configuration.

In the case of a failure or split-brain, all partitions that do not contain a majority of nodes are terminated. This ensures that if there is a partition running that contains a majority of the nodes, it can safely start up any resources that are not running on that partition, safe in the knowledge that it can be the only partition in the cluster that is running resources (because all other partitions are terminated).

Given the differences in the way the shared disk quorum clusters behave compared to MNS quorum clusters, care must be taken when deciding which model to choose. For example, if you only have two nodes in your cluster, the MNS model is not recommended, as failure of one node will lead to failure of the entire cluster, since a majority of nodes is impossible.

# Quorum Uses

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Standard Quorums**
Standard quorums should be sufficient for the majority of situations that users will encounter. Typical situations include:

- *Highly available data in a single location* – most customers that require their data to be highly available only need this on a per site basis. If they have multiple sites, each site has its own cluster. Typical applications that use this type of cluster include Microsoft SQL Server™, Exchange Server, file shares, printer queues, and network services (e.g. DHCP & WINS).

- *Stateful applications* – applications or NT services that require only a single instance at any time and require some sort of state to be stored, typically use standard quorums, as they already have some sort of shared storage for maintaining the state.

**Majority Node Set Quorums**
While majority node set quorums are an exciting new feature of Windows Server 2003 clusters, they do have very strict requirements to ensure they work correctly, and therefore should only be considered by people who fully understand the issues involved in using MNS based clusters. The following are the key situations that the product team thought of when creating MNS quorums:

- *Geographically Dispersed Clusters* – this involves a single MSCS cluster that has members in multiple geographic sites. While geographic clusters are possible using a standard quorum (for example, there is a separate geographic cluster Hardware Compatibility List (HCL)), a number of issues arise in terms of presenting the quorum as a single, logical shared drive between all sites. Majority Node Set quorums solve these issues by allowing the quorum to be stored on the local hard disk.

- *Clusters with No Shared Disks* – there are some specialized configurations that need tightly consistent cluster features without having shared disks. For example:

    - Clusters that host applications that can failover, but where there is some other, application-specific way to keep data consistent between nodes (e.g. database log shipping for keeping database state up-to-date; file replication for relatively static data, etc.).

    - Clusters that host applications that have no persistent data, but need to cooperate in a tightly coupled way to provide consistent volatile state.

    - Independent Software Vendors – by abstracting storage from the Cluster service, it provides independent software vendors with much greater flexibility in how they design sophisticated cluster scenarios.

# Frequently Asked Questions (Server Clusters: Quorum Options - Windows Server 2003)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What is a Majority Node Set cluster?**

It is a cluster that has a quorum based on local disks in each node, rather than a shared disk.

**How is the quorum information located on the system disk of each node kept in synch?**

The server cluster infrastructure ensures that all changes are replicated and updated on all members in a cluster.

**Can this method be used to replicate application data as well?**

No, that is not possible in this version of clustering. Only Quorum information is replicated and maintained in a synchronized state by the clustering infrastructure.

**Can I convert a standard cluster to an MNS cluster?**

Yes. You can use Cluster Administrator to create a new Majority Node Set resource and then, on the cluster properties sheet **Quorum** tab, change the quorum to that Majority Node Set resource.

**What is the difference between a geographically dispersed cluster and an MNS cluster?**

A geographic cluster refers to a cluster that has nodes in multiple locations, while an MNS-based cluster refers to the type of quorum resources in use. A geographic cluster can use either a shared disk or MNS quorum resource, while an MNS-based cluster can be located in a single site, or span multiple sites.

**What is the maximum number of nodes in an MNS cluster?**

Windows Server 2003 supports 8-node clusters for both Enterprise Edition and Datacenter Edition.

**Do I need special hardware to use an MNS cluster?**

There is nothing inherent in the MNS architecture that requires any special hardware, other than what is required for a standard cluster (for example, there must be on the Microsoft Cluster HCL). However, some situations that use an MNS cluster may have unique requirements (such as geographic clusters), where data must be replicated in real time between sites.

**Does a cluster aware application need to be rewritten to support MNS?**

No, using an MNS quorum requires no change to the application. However, some cluster aware applications expect a shared disk (for example SQL Server 2000), so while you do not need shared disks for the quorum, you do need shared disks for the application.

**Does MNS get rid of the need for shared disks?**

It depends on the application. For example, clustered SQL Server 2000 requires shared disk for data. Remember, MNS only removes the need for a shared disk quorum.

**If I have a 2-node MNS cluster, and lose one server, will the cluster continue to function?**

No, in order to keep an MNS cluster running, it must have a majority of nodes online, i.e. (<Number of nodes configured in the cluster>/2) + 1. The following table shows the various combinations of configured nodes and node failures:

| Number of nodes configured in the cluster | Number of node failures |
|---|---|
| 1 | 0 |
| 2 | 0 |
| 3 | 1 |
| 4 | 1 |
| 5 | 2 |

| | |
|---|---|
| 6 | 2 |
| 7 | 3 |
| 8 | 3 |

**How do I back up and restore an MNS cluster?**

There is a separate white paper on backing up and restoring clustered configurations. It will be published at the same location as this paper.

**Is there a way I can allow for a "minority set" to take control of the cluster?**

Yes, there is a way for a "minority" set to be allowed to start up and initialize the cluster. But this is a manual process and has to be done in a specified manner. Details on the same are available in "help".

**What happens if majority count is lost?**

The cluster service will gracefully shutdown the cluster and all resources. Clients will lose connectivity and access to any of the clustered services. This is done to prevent the "split-brain" scenario.

**Are there any limitations with regard to network latency and topography in an MNS cluster?**

Requirements are the same as that for vanilla clusters. With regard to topography, we envisage the MNS feature to be predominantly used in geographically dispersed clustering scenarios. This normally means specific semantic requirements in terms of storage usable and management. OEMs provide prescriptive guidance on the same and customers should carefully read the same.

For the latest information about Windows Server 2003, see the Windows Server 2003 Web site at http://www.microsoft.com/windowsserver2003.

# Server Clusters: Storage Best Practices for Windows 2000 and Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In This White Paper**


Storage Topologies

Storage Features in Windows Server 2003 and Clustering

Hardware Based Storage Features

Storage Tools and Services

Backup and Restore (Server Clusters: Storage Best Practices for Windows 2000 and Windows Server 2003)

Storage Management

Storage KB Articles

# Storage Topologies

Updated: January 1, 2003

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are two types of Storage I/O technologies supported in Server clusters: Parallel SCSI and Fibre Channel. With the release of Microsoft Windows Server 2003, support will be provided for SCSI interconnects and fibre channel arbitrated loops for two nodes only. For larger configurations (greater than two), you will need to use a switched fibre channel (fabric) environment.

**Issues**

Parallel SCSI

Only supported in Windows 2000 Advanced Server up to two-nodes

SCSI Adaptors and storage solutions need to be certified

SCSI cards that are hosting the interconnect should have different SCSI IDs normally (6 and 7). Ensure device access requirements are in-line with SCSI IDs and priorities

SCSI adaptor BIOS should be disabled

If devices are daisy-chained, ensure that both ends of the shared bus are terminated

Use physical terminating devices and do not use controller-based or device-based termination

SCSI hubs are not supported

Avoid the use of connector converters (e.g. 68-pin to 50-pin)

Avoid combining multiple device types (single ended and differential, etc.)

Fibre Channel

Fibre Channel Arbitrated Loops (FC-AL) supported up to two nodes

Fibre Channel Fabric (FC-SW) supported for all higher combinations

Components and configuration need to be in the Microsoft Hardware Compatibility List (HCL)

Multi-cluster environment

Fault tolerant drivers and components also need to be certified

Virtualization engines need to be certified.

The switch is the only component that is not currently certified by Microsoft and it is recommended that the end user get the appropriate interoperability guarantees from the switch vendor before implementing switch fabric topologies. In complicated topologies, where multiple switches are used and connected through ISLs, it is recommended that the customer work closely with Microsoft and the switch and storage vendors during the implementation phase to ensure that all of the components work well together.

**Supported and Qualified Configurations**

All server clusters must be qualified to be supported by Microsoft. A qualified configuration has undergone extensive testing using a hardware compatibility test, provided by Microsoft. All qualified solutions appear on the HCL, available at http://www.microsoft.com/whdc/hcl/default.mspx. Only cluster solutions listed on the HCL are supported by Microsoft. The complete cluster solution must be listed on the Cluster HCL list. The complete solution includes the servers, the storage adapters, the interconnect type, the storage controllers firmware and driver versions. All of the components must match exactly, including any software, driver or firmware versions, for the solution to be qualified. The HCL contains a set of qualified cluster components. A solution built from qualified components does NOT imply that the solution is qualified.

The cluster components lists have been a source of confusion in the past and Microsoft will be removing the cluster component lists (such as Cluster/RAID) from the HCL for Windows Server 2003.

**Storage Interconnects**

Ensure that all storage interconnects used in Server Clusters are in the HCL. This also applies for any additional software that is

used to provide fault tolerant or load balancing features for adaptors and interconnects.

**Multiple paths to the storage for high availability**: This is a very common feature implemented by almost all storage vendors. This feature allows for users to implement multiple fabrics (normally two), and use them in fault tolerant or load balancing configurations. In the past, each vendor had their own implementations that varied greatly from one another and also required specific configuration/driver combinations. With the release of Windows Server 2003, Microsoft has developed and supplied vendors with a multi-path driver that will be used by them, in place of custom-made drivers. The driver will ship as part of the individual vendor product, but it is expected that all vendor products (that Microsoft will support) will incorporate this driver. If they insist on using their own driver, it will need to be certified and listed in the HCL.

#### Server Cluster and SANs

Storage Area Networks (SANs) are increasingly being used to host storage that is managed by server clustering. There are some specific requirements driven from the implementations of clusters and the fact that all storage on the SANs may not be owned by nodes in a cluster. Some issues that can be translated into best practices are:

Ensure that the SAN configurations are in the Microsoft HCL (multi-cluster section).

When configuring your storage the following **must** be implemented:

**Zoning**- Zoning allows users to sandbox the logical volumes that will be used by a cluster. Any interactions between nodes and storage volumes will be isolated to the zone and other members of the SAN will not be affected by the same. This feature can be implemented at the controller or switch level and it is important that users have this implemented before installing clustering

**LUN masking**- This feature allows users to express at the controller level, a specific relationship between a LUN and a host. In theory, no other host should be able to see that LUN or manipulate it in any way. However, various implementations differ in functionality and as such; one cannot assume that LUN masking will always work. Therefore, it cannot be used *instead* of zoning. One can combine zoning and masking however, to meet some specific configuration requirements.

**Firmware and driver versions**- Some vendors implement specific functionality in drivers and firmware and it is recommended that users pay close attention to what firmware/driver combinations are compatible with the installation they are running. This is valid not only when building a SAN and attaching host to it but also to the entire life span of the system (hosts and SAN components). Close attention should be paid to issues coming out of applying service packs or vendor specific patches and/or upgrades.

**Hardware verses software zoning**- Zoning can be implemented in hardware/firmware on controllers or on software on the hosts. It is recommended that controller-based zoning be used since this allows for uniform implementation of access policy that cannot be interrupted or compromised by node disruption or failure of the software component.

**Hardware versus software LUN masking**- Some vendors also offer software-based masking facilities. Please make sure that any such software that is closely attached to storage and is involved with the presentation of the storage to the operating system needs to be certified. For the same reasons mentioned above, this is also not considered to be a very good thing to do if one cannot guarantee the stability of the software component.

**Boot from SAN**- This is an increasingly demanded feature and will be supported by Microsoft in Windows Server 2003. Some factors to consider are :

Configurations require support from the Host Bus Adaptor (HBA) and storage vendors. The HBA driver needs to be a Storport driver. Storport drivers improve performance over SCSIPORT, both in terms of throughput and in terms of the system resources that are utilized. It also adds a manageability infrastructure for configuration and management of host-based RAID adapters and is a new implementation in Windows Server 2003. These features are required if the operating system is to successfully boot from a SAN, and more importantly, the vendors that have an implementation need to get them certified by Microsoft.

Such solutions have limited scaling capacities and any additional complexities such as storage replication, recovery mechanisms etc. need to be addressed by the hardware vendors.

Server clusters have enabled a feature that will allow the startup disk, pagefile disks and the cluster disks to be hosted on the same channel. There are other performance and operational implications that need to be considered before implementation.

Please also read KB article 305547 *Support for Booting from a Storage Area Network (SAN)* which discusses this feature.

#### Storage Configuration and Setup

If this is a fresh cluster installation, you need to ensure that you do not have applications running. When creating a cluster or adding nodes to a cluster, the Wizard enumerates all of the storage on the node and clusters all storage that is not somehow determined to be non-clusterable. Non-clusterable storage includes all drives on any storage bus that has any system file on it

(boot, system, page file, crash dump, or hibernate file), LDM drives, and those drives that respond positively to the IOCTL IOCTL_SCSI_MINIPORT_NOT_QUORUM_CAPABLE.

If this is an existing cluster that is being upgraded, then all configuration settings are honored and rolling upgrade from Windows 2000 is a fully supported option.

**Server Clusters and Fault-Tolerant Disks (RAID)**
Creation of fault-tolerant volumes such as striped, mirrored, and RAID-5 volumes in Windows Server 2003 requires the volumes to be dynamic disks. Dynamic disks are not supported out-of-the-box in Server cluster, however Veritas has a product (Veritas Volume Manager) that allows for similar functionality and is supported. NTFS is supported as a format type. If disks are not configured as basic and formatted in NTFS, clustering will not recognize the disks and will not be able to manage them. Server clustering also does not support GPT (GUID Partition table) disks, which was introduced as part of the EFI initiative.

Clustering, however, supports hardware based fault-tolerant disks. This means that the configuration of physical disks into fault-tolerant sets (JBOD, RAID 5, RAID 0+1 etc.) is done at the controller level and the set can be made visible either as a whole or carved into smaller pieces (volumes) and made available to hosts. The hosts are completely unaware of the physical implementation and treat the volume as a disk. It is for this reason that we recommend that usage characteristics be taken into consideration before implementing FT sets. For example, if you think the logical volume will host data that is being read and updated constantly, you should perhaps implement the fault-tolerant set as a RAID 0+1 instead of a RAID 5. Another thumb rule to follow would be associating the logical volume size as closely as possible to the physical set size. So carving out fewer logical volumes generally helps contribute to performance.

Volume expansion is also a regular requirement since data growth is not often considered up-front. Cluster disks can be extended without rebooting if the controller supports dynamic LUN expansion. This feature allows for the physical expansion to be implemented without disruption and users can use tools (diskpart) provided by Microsoft to allow for the change to be seamlessly applied at the logical level as well. There are separate versions of diskpart for Windows 2000 and Windows Server 2003. The Windows 2000 version is available as a free download on the web and the Windows Server 2003 version is shipped on the distribution CD.

**Handling Storage Cable Disconnects in a Cluster**
This document has been written up to answer the various queries we get about Microsoft not supporting manual disconnection of storage cable as a valid test for testing failover mechanisms. Microsoft software fully supports the validity of the test but is not responsible for guaranteeing the results. The main reason for this is that the reactions of the concerned devices (HBAs, switches, Storage solutions etc.) are controlled by lower level device drivers that are written by third party vendors. Different vendors implement storage drivers in different ways and it is upto them to implement features that will allow for such tests to work. At the time of writing this document not all vendors have device drivers that support this test. You should check with your vendor (storage and adapter) to ensure that pulling the storage adapter cable is correctly handled by their device drivers. If the drivers are written to specification made available by Microsoft this issue will be handled cleanly and the tests should work.

This document is specific to Fibre Channel based storage and the issue of storage cables being disconnected or broken and has been created to clearly state the position of the cluster product team with regard to the same. One would see similar issues with SCSI bus implementations but since a majority of implementations these days are fibre channel based we will be scoping our discussion that type of storage network.

Essentially there are usually 2 sets of connections between servers and storage. Servers and storage solutions are connected to switches/hubs. As such there can be 2 places disconnects can occur, one between the server (HBA) and the switch, and the other between the storage solution and the switch. And since there are many cluster and non cluster components involved when such an event happens we will try and map the events to the components thus establishing clear areas of responsibilities.

To begin with Server clusters product team fully understand the fact that storage cable disconnects leading to disruption in host connectivity is a valid scenario and can happen at customer sites. The product (Server clusters), fully supports handling such events gracefully and failing over resources (disk and related). However there are set of events that need to happen before Server clusters can detect the event and failover resources. There is also an additional set of events that need to happen for the resources to fail back to the original node. These set of events depend on the behavioral characteristics of certain components (lower level device drivers) that are not part of the server cluster product. Unfortunately there is no standardization of the required characteristics of these drivers and as such the behavioral semantics might differ from configuration to configuration. But the basic behavior linked to a cable disconnect (that of a failure being detected causing the disks to fail over) should always work.

We will state some of the requirements and issues in this document so as to give the reader an understanding on how all this works together and what questions to ask of whom to ensure that their systems have this functionality.

**Disconnecting storage cables (HBA to switch)**

We will first address the issue of a storage cable between the server (HBA and the switch) being disconnected.

Here is a rough diagram of what the storage stack will look like on a server running Windows and Server clusters. The stack could include scsiport drivers supplied by Microsoft and miniport drivers supplied by the HBA vendor or a full port driver that complete functionality and replaces the scsiport drive:




**Case 1**




**Case 2**

In the first case when the cable is disconnected, if the HBA miniport drivers report the correct status (**BusChangeDetected** - This notification indicates that a target device might have been added or removed from a dynamic bus.) the scsiport drivers can do the right thing and the non availability (or availability) of disks is efficiently detected. If a device was reported as being not available PnP will do a scan and tear down all device objects that are not available) and the server cluster disk driver will get the required notifications. This will lead to the cluster service failing over the disk resources to another valid member of the cluster who has access to the devices in question.

If the HBA miniport driver instead reports a generic status (**ResetDetected** - all this indicates is that the HBA has detected a reset on the SCSI bus and after this notification, the HBA miniport driver is still responsible for completing any active requests). In such a case the cluster disk driver will eventually detect the fact that the disk is no longer available in the course of the monitoring tests (LooksAlive, IsAlive1) it perfumes on a periodic basis. This will cause to detect that the disks are no longer accessible leading it to fail over the disks to another member in the cluster that has valid access to the disks in question. The minimum time required to detect this is approximately three seconds when a reserve issued by clusdisk (the cluster disk driver) fails. The maximum time required is not quantifiable since that will depend on the behavioral and operational characteristics of the lower level device drivers. But the non availability of the disk will always be eventually detected causing the disks to fail over to a valid cluster node.

The frequency for running LooksAlive and IsAlive is tunable, meaning one can increase the frequency of the checks so has to have a smaller detection window but this has other (resource utilization) implications that needs to be carefully considered, especially in enterprise environments where server clusters might need to monitor a large number of resources.

The same is also valid if the out-of-the-box scsiport drivers are replaced by a vendor specific scsiport driver or of the SCSIport /Miniport combination is replaced by a monolithic Fullport driver as shown in Case 2. Such drivers are written by vendors to cater to specific requirements their storage solutions might have. Microsoft cannot guarantee the functionality of full port drivers Microsoft support will not be able to debug related issues if full port drivers are used.

**Disconnecting storage cables (Switch to storage):**
Here again the behavior depends on the implementations. Upon cable disconnect or device removal the HBA gets an RSCN from the switch and should notify the operating system of any changes it detects. Detection is a very complex operation and has a lot of external dependencies and may not always work correctly. In the worst case scenario the HBA may fail the IO or return BUSY. In such cases as well the worst case scenario applies which means that the LooksAlive/IsAlive checks will fail and the cluster service will wither failover the disks or fail them completely.

**Reconnecting Storage Cables (HBA to switch):**
Reconnecting the storage cables and expecting the resources to fail back is also a complicated issue. If the HBA miniport driver or the full port driver sends the right status then PnP will rescan the bus create device objects and allocate resources (such as I/O ports, memory addresses, and interrupts) for that device. These resources are supplied back to the miniport, which must then use these resources. In such cases reconnecting the storage cable is detected and is all that needs to be done to allow for the resources to fail back to the original node.

If the miniport driver or the full port driver does not implement the above then again fail back should just work because the physical device objects exist. If there are issues the user will need to initiate a manual scan (by going into Disk Management MMC and initiating a rescan of disks). Once this is done a failback should work

**Reconnecting Storage Cables (Switch to storage controller):**
Getting information (on cable reconnects) back from the switch is dependant on various conditions (have devices changed, has the name server detected this etc.) and does not always work as expected. A manual rescan should also work in such cases.

In either case the full functionality is provided by the miniport or the full port drivers (depending on what is being used) both of which are written by the storage or HBA vendor.

The above is also valid for fault tolerant solutions that implement features like multiple paths from servers to storage solutions. The behavior should be the same in event of **both paths** on a host failing. Single failure instances should be masked by the multi path driver and should not be visible to higher level subsystems.

To sum it up disconnecting storage cables will cause disks managed by server clusters to fail over to another member of a cluster no matter what driver types are used. Failing the disks back to the original node however is a feature that can only be guaranteed by HBA vendors for the simple fact that they own the drivers that incorporates the functionality that is primarily responsible for correctly handling this situation.

**Adding new disks to a cluster**- Adding new disks to a cluster does not require a rebooting nodes or restarting the cluster service post Windows 2000 SP3 and in Windows Server 2003. Any node can scan and mount the new disk to begin with and once it has been formatted to NTFS all that is required is the creation of a new physical disk resource using Cluster

The above is valid post Windows 2000 SP3 (2 hotfixes) and Windows Server 2003. The hotfixes are supplied by Microsoft Support.

# Storage Features in Windows Server 2003 and Clustering

Updated: January 1, 2003

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Volume Mount Points**
Server clustering in Windows Server 2003 supports volume mount points, which are directories on a volume that an application can use to "mount" a different volume, that is, to set it up for use at the location a user specifies. This feature is not supported in Windows NT 4.0 or Windows 2000.

**EFS**
With Windows Server 2003, the encrypting file system (EFS) is supported on clustered file shares. To enable EFS on a clustered file share, you must perform a number of tasks to configure the environment correctly:

EFS can only be enabled on file shares when the virtual server has Kerberos enabled. By default, Kerberos is not enabled on a virtual server. To enable Kerberos you must check the Enable Kerberos Authentication check box on the network name resource that will be used to connect to the clustered file share.

| **Note** |
| --- |
| Enabling Kerberos on a network name has a number of implications that you should ensure you fully understand before checking the box. |

All cluster node computer accounts, as well as the virtual server computer account, must be trusted for delegation. See online help for how to do this.

To ensure that the users private keys are available to all nodes in the cluster, you must enable roaming profiles for users who want to store data using EFS. See online help for how to enable roaming profiles.

Once the cluster file shares have been created and the configuration steps above have been carried out, users data can be stored in encrypted files for added security.

**Disk Quotas**
Configuring disk quotas on shared disks is supported in Windows 2000 and Windows Server 2003.

**Autochk/Chkdsk/Chkntfs** - Every time Windows restarts, Autochk.exe is called by the Kernel to scan all volumes to check if the volume dirty bit is set. If the dirty bit is set, autochk performs an immediate chkdsk /f on that volume.

Chkdsk is a native Windows tool that can determine the extent of file and file system corruption. If Chkdsk runs in write-mode, it will automatically attempt to remedy disk corruption.

The Chkntfs.exe utility is available in Microsoft Windows NT 4.0 Service Pack 2 and later versions of Windows. It is designed to disable the automatic running of chkdsk on specific volumes, when Windows restarts from an improper shutdown. Chkntfs can also be used to unschedule a chkdsk if chkdsk /f was used to schedule a chkdsk on an active volume on the next system restart.

All of the above are supported to run in specific configurations in Server Clusters. The relevant KB articles that explain the procedures in detail are:

174617 - *Chkdsk runs while running Microsoft Cluster Server Setup*

176970 - *How to Run the CHKDSK /F Command on a Shared Cluster Disk*

160963 - *CHKNTFS.EXE: What You Can Use It For*

**Virtual Disk Service** - The Virtual Disk Service will ship in Server 2003. It provides an interface for software volume management, hardware RAID cabinet management, multi-path and device allocation in a SAN. VDS will be supported in Server Clusters. Additional providers (supplied by ISVs like Veritas) are required to enable cluster functionality.

**Virtual Snap Shot and TimeWarp** - VSS is a Windows service that allows users to snapshot data. It does this with the help of hardware or software providers that actually implement the functionality. VSS coordinates requestors, providers, and writers. The Volume Shadow Copy service is abbreviated as VSS. Software providers are implemented as a user-mode component plus a kernel-mode driver. I/Os are intercepted by the driver that instantiates snapshot volumes. The driver may be a storage filter driver or a component of a volume manager. A writer is any application that stores persistent information on one or more

volumes, and participates in snapshot synchronization. Typically this might be a database (e.g. SQL Server) or system service (e.g. Certificate Services) or a backup program (Backup). VSS implementation in NT backup is fully supported by server clusters. Timewarp is another specific use of snapshots for simple file un-delete and is a feature available in Windows Server 2003. This is also fully supported in clustering and the whole feature set can be enabled on clustered and non clustered disks in the same fashion. End user experience remains unchanged.

# Hardware Based Storage Features

Updated: January 1, 2003

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware snapshots or business recovery volumes**- Hardware snapshots are totally transparent to clustering and are supported if implemented correctly. The only requirement we have is that the snapshot copy and the original should not be exposed to the cluster at the same time.

**NAS as shared storage**- We fully support building clustered solutions using the Microsoft NAS offering called the Server Appliance Kit. The kit is supplied to OEMs who will ship predefined solutions, some of which may be clustered (for example, clustered file servers).

**Geographical clusters (storage implications)**- With geographic clusters, their storage configurations can tend to get very complicated depending upon the level of service users require. All solutions need to be in the Geographic Cluster HCT in order to be supported by Microsoft. Features like data replication, storage management, recovery procedures etc. are entirely the responsibility of the hardware vendor implementing the solution. We also have an HCT test available for free download on the web that can be used to validate production systems.

**Replication (synchronous, asynchronous)**- FRS, the OOB replication mechanism, is not really used at the application data replication level. Server cluster is replication agnostic and does not take the feature into consideration when it is running. If ISVs (hardware or software) choose to implement this feature, they should make sure the semantics required by Server clusters at the storage level are honored. It is their responsibility that all application requirements are met with regard to data replication and consistency.

# Storage Tools and Services

Updated: January 1, 2003

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Storage**- Remote storage services are not supported by Server Cluster.

**Removable Storage**- Removable storage services are not supported by Server Cluster.

**Backup**- NTBackup and any backup solution that use the Cluster Backup APIs are supported for Server Cluster Backup scenarios. Using the APIs allows for consistent backup and restore of cluster configuration data. Shared disks can only be backed up on the node that owns the disk at the time of backup.

**Indexing**- Indexing is not supported by Server Cluster. It also has serious implications with regard to resource utilization and careful consideration must be given before turning it on.

**Distributed File System**- DFS is fully supported by Server Cluster. In Windows Server 2003, multiple DFS roots can be hosted in a cluster.

**Offline Files**- Client-side caching can be enabled on clustered file share in Windows Server 2003. This feature is not available in Windows 2000 or NT 4.0. This is meant for user shares only and not for application data.

# Backup and Restore (Server Clusters: Storage Best Practices for Windows 2000 and Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As stated earlier, any backup solution that uses the cluster backup APIs are supported.

**Tape devices or other non-disk devices** the SAN fabric as the cluster: currently we do not support managing any non-disk devices (tapes, optical media etc.) by Server clusters. Tapes or optical media may very well be used as any other SAN component but they cannot be managed by Server clusters as a resource.

**Serverless backup**- Serverless backup, or SAN-based backup as it is sometimes called, is not supported in Server clusters. This is because nodes in a cluster own physical disk volumes and maintain their ownership through reservations, and this blocks serverless backup mechanisms. In addition backup devices translate incorrectly interpret bus resets leading them to rewind tapes loaded in the drives. This will have serious implications in enterprise backup environments where large data sets are being backed up.

**Using backup agents**- Some ISVs implement backup agents which are really filter agents that are used to overcome file locking issues. They are also sometime called open file agents. Stability and reliability of the agents and the operating system is the responsibility of the vendor if such solutions are implemented.

**Hardware Snapshots**- Hardware snapshots are not supported in clustered environments. Clustering software identify disks thru their disk signatures. Disk snapshots or clones are identical copies of the parent disks which means that 2 disks with similar signatures would get exposed to the operating system and clustering software. This could cause unexpected disruptive behavior.

# Storage Management

Updated: January 1, 2003

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Storage management is a very complex area and is implemented by different vendors differently. This can be done on a piecemeal basis (each component managed differently) or as a whole (one management solution for the entire storage infrastructure). The later is still not commercially available and as such, most implementations are of the former kind. Some issues that one should take note of are:

**SAN Management**- Most vendors supply their own tools/applications to manage the SAN. This may or may not include all the components (like switches) that are required to keep the SAN up and running as well. Users need to make sure that no matter which tool they use, they get a consistent view of the SAN from a configuration and operational perspective. Management stations (the nodes that the management software runs on) could be a node connected to the SAN or could be talking to SAN components through a network (using TCP/IP). Each of the above methods have their advantages and disadvantages and users must make sure that management is tightly bound by operational procedures.

**Storage Controller Management**- Some controllers allow for in-band and/or out-of-band management. Users need to makes sure that management procedures at the controller level are in synch with SAN level and node level (application) management procedures. This is especially relevant in clustered environment where disconnected events could trigger failovers or application shut downs.

**Switch Management**- A switch is a central component in a SAN and controls all access from nodes to storage. In addition to what has been mentioned earlier, users also need to make sure that switch state changes are propagated to both the nodes and the storage controllers. These are highly specialized configurations and are normally owned by a single vendor. Users need to make sure the complete operational system is implemented by the vendor before signing off on the installation. Any configuration changes, like firmware or driver updates, need to be validated by the vendor and Microsoft.

**Node Management**- Nodes can be used to manage the SAN or just have a few management components installed to react to specific state changes sent from a SAN. Needless to say, the functionality of the agents needs to be validated in order to make sure that they do not interfere with the operating system or clustering functionality.

# Storage KB Articles

Updated: January 1, 2003

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows NT and Windows 2000**

| | |
|---|---|
| 174617 | Chkdsk runs while running Microsoft Cluster Server Setup |
| 176970 | Run the CHKDSK /F Command on a Shared Cluster Disk |
| 250355 | Antivirus Software May Cause Problems with Cluster Services |
| 280353 | How to Change Quorum Disk Designation |
| 237853 | Dynamic Disk Configuration Unavailable for Server Cluster Disks |
| 259237 | Troubleshooting Event ID 9, 11, and 15 on Cluster Servers |
| 223023 | Enhanced Disk Resource Private Properties Using Cluster Server |
| 263590 | How to Extend the Space of Existing Shared Disks with Windows Clustering |
| 224508 | How to Migrate Dfs Root Configuration to a Windows 2000 Cluster |
| 224969 | How to: Configure Event Log Replication in Windows 2000 Cluster Servers |
| 226108 | Cannot Place Swap File on a Windows NT Cluster Shared Disk |
| 238137 | Considerations when you create clustered file shares on a Windows 2000 or Windows Server 2003 server cluster |
| 249194 | MSCS/Cluster Does Not Form with Error Messages 170 and 5086 |
| 248998 | How to Properly Restore Cluster Information |
| 254219 | Security Considerations When Implementing Clustered File Shares |
| 256926 | Implementing Home Folders on a Server Cluster |
| 257897 | MSCS printer shares may not come online if quorum log is too small |
| 259243 | How to set the startup value for a resource on a clustered server |
| 262797 | Reparse Point Support in Windows 2000-Based Clusters |
| 269137 | Some Cluster Server File Share Resources Are Not Available |
| 280125 | Cluster Administrator Switches for Connecting to a Cluster |
| 278365 | How to Configure Disk Quotas for a Shared Disk in a Cluster |

| | |
|---|---|
| 280743 | Windows Clustering and Geographically Separate Sites |
| 278710 | No Global Groups Are Available Creating File-Share Resource Permissions in Cluster Administrator |
| 278951 | The Cluster Internet Protocol Address May Not Come Online After the Quorum Disk Is Corrupted |
| 284838 | How to Create a Server Cluster File Share with Cluster.exe |
| 293778 | Multiple-Path Software May Cause Disk Signature to Change |
| 245762 | Recovering from a Lost or Corrupted Quorum Log |
| 257389 | Microsoft Cluster Server May Not Share Folders Automatically |
| 223023 | Enhanced Disk Resource Private Properties Using Cluster Server |
| 259267 | Microsoft Cluster Service Installation Resources |

**Windows Server 2003**

| | |
|---|---|
| 283715 | The Features of the Local Quorum Resource on Windows Server 2003 Cluster |
| 284134 | Server Clusters Do Not Support GPT Shared Disks |
| 245762 | Recovering from a Lost or Corrupted Quorum Log |
| 175761 | Dynamic vs. Basic Storage in Windows 2000 |
| 198513 | Clustering Cannot Determine If a Shared Disk Is Working Properly |
| 251186 | Best Practices for NTFS Compression |
| 290987 | Error 3 When Creating a File Share on a Cluster with the Application Wizard |
| 302539 | How to Troubleshoot Printing Issues on a Windows Server 2003 Cluster |
| 301647 | Cluster Service improvements for Storage Area Networks (SANs) |
| 304415 | Support for Multiple Clusters Attached to the Same SAN Device |
| 306640 | A Removed Disk Can Have Its Corresponding Cluster Resource Started |
| 312090 | Cannot Use 32-Bit Resources on a 64-Bit Server Cluster |
| 318534 | Best Practices for Drive-Letter Assignments on a Server Cluster |
| 280297 | How to Configure Volume Mount Points on a Clustered Server |
| 296799 | How to Configure Windows Clustering Groups for Hot Spare Support |
| 301588 | How to Use DFS on Server Cluster to Maintain a Single Namespace |
| 302829 | Cluster Service Does Not Remove File Share Entries from LANMANSRVR |

| | |
|---|---|
| 304736 | How to Extend the Partition of a Cluster Shared Disk |
| 310072 | Adding support for more than eight LUNs in Windows Server 2003 and Windows 2000 |
| 305547 | Support for Booting from a Storage Area Network (SAN) |

[1] **LooksAlive** is a cursory status check that runs every 5 seconds (default). This routine simply checks that the disk status is not marked failed. **LooksAlive** is a cursory status check that runs every 5 seconds (default). This routine simply checks that the disk status is not marked failed. **IsAlive** is a more thorough check and occurs every 60 seconds (default). This routine checks that the disk status is not marked failed. If status is not failed, then **FindFirstFile** runs on the root of the disk to make sure the file system is still mounted and the disk is accessible

# Server Clusters: Security Best Practices, Windows 2000 and Windows Server 2003, General Assumptions

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Published: January 1, 2003

**Also In This White Paper**

Server Clusters: Security Best Practices, Windows 2000 and Windows Server 2003, Deployment and Operations

**General Assumptions**

There are a number of general assumptions and operational best practices that should be in place for the infrastructure to ensure a secure environment in which to run server clusters.

1. Servers and storage are in physically secure locations.

2. Practical security implementations to detect irregular traffic, such as firewalls, network probes and management tools, are in place.

3. Best practices/common sense in terms of security are adhered to in areas like administration, storage of logs, backup and restore etc.

4. Platform level security best practices are adhered to in terms of assigning administrative permissions, ACLing resources, and other housekeeping roles.

5. The network infrastructure services such as Active Directory, DNS, DHCP, WINS etc. must be secure. Any compromise of these infrastructure services can lead to a compromise of the cluster service itself.

6. The cluster administrator must ensure that applications that call the cluster APIs (ClusAPI) are run from trusted computers. Any compromise on the computers on which the applications are executing (that the cluster administrator runs) can compromise the cluster. For example, if there are untrusted users with elevated privileges on the workstation from where the administration tools are run, untrusted or malicious code can be run against the cluster by the cluster administrator without the cluster administrator realizing.

7. Access to the set of objects created and maintained by the cluster service must not be compromised by adjusting the default settings placed on these objects to a less restrictive setting. The cluster service utilizes a set of objects in the operating system such as files, devices, registry keys etc. These objects have a default security setting that ensures non-privileged users cannot impact the cluster configuration or the applications running on the cluster. Changing these security settings to less restrictive security settings can lead to the cluster being compromised and the application data being corrupted.

# Server Clusters: Security Best Practices, Windows 2000 and Windows Server 2003, Deployment and Operations

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Published: January 1, 2003

**Also In This White Paper**

Server Clusters: Security Best Practices, Windows 2000 and Windows Server 2003, General Assumptions

**In This Section**

This section of the white paper contains the following subsections:

Cluster administrators

Remotely managing and configuring clusters

Cluster service account

Changing the Password

Using Kerberos Authentication in a Server cluster

Creating Computer Objects

Disabling Computer Objects

Renaming Computer Objects

Password Rotation

Caveats, Issues and Considerations

Network Security

NTLM V1 and NTLM V2

NetBIOS

IPSec

Cluster Disks

Managing file shares in a cluster

Cluster Server nodes as Domain Controllers

Summary of security attributes

Majority Node Set considerations

Upgrading to Windows Server 2003 considerations

Developing Cluster-aware Applications

Calling the Server cluster APIs

Backup/Restore APIs

**Cluster administrators**

Administrators can specify groups or individuals that are allowed to manage the cluster. In current versions of Server cluster, there are no fine details of control; either a user has rights to administer the cluster or the user does not. To grant a user or group rights to administer the cluster, the user or group must be added to the cluster security descriptor. This can be done through cluster administrator or the cluster.exe command line tool.

**Note**

| |
|---|
| Apart from the local Administrators group on a node, all other members of the cluster security configuration MUST be either domain user accounts or global groups. This is to ensure that the account is the same, well defined and authorized account on all nodes in the cluster. |

By default, the local Administrators group is added to the cluster service security descriptor.

Adding a user or group to the cluster security descriptor means that the user can manage all aspects of the cluster configuration including (but not limited to):

- Taking resources offline and bringing resources online
- Shutting down the cluster service on nodes
- Adding and removing nodes from the cluster
- Adding and removing resources from the cluster

Because of the scope of impact on the applications and services running in the cluster, careful consideration must be taken when adding a user to the cluster security descriptor.

The cluster service runs the code associated with a resource under the cluster service domain user account (this should not be confused with the account used to administer the cluster). Since a cluster administrator can add new resources to a cluster and since those resources run as the cluster service account, a cluster administrator can install code that runs with local administrator rights on the machine.

**Best Practices**

- Cluster administrators should use a different account than the cluster service account to administer the cluster. This allows different policies (such as password expiration etc.) to be applied separately to the cluster service account and the domain account(s) used to administer the cluster.
- You should only add users with local Administrator rights to the cluster service security descriptor.

  | **Note** |
  |---|
  | By adding a domain user or global group to the local Administrator group, that group or account will automatically be a cluster administrator. |

- Do not remove the local Administrator group from the cluster service security configuration.

**Remotely managing and configuring clusters**
Administration tools or other applications that call the Server cluster APIs (ClusAPI) can be run from remote workstations. The general assumption is that the cluster administrator must ensure that the applications are run from trusted computers. Any compromise on the computers on which the applications are executing (that the cluster administrator runs) can compromise the cluster.

When a cluster is created or the configuration is changed (such as adding a new cluster node), the Cluster Configuration Wizard will create a log file on the machine on which the wizard is run so that in the event of failures, the administrator can use the log for debugging and troubleshooting. This log file can contain cluster configuration data such as cluster IP addresses, network names etc. This data could be used to extend the attack surface if it is read by unauthorized users.

**Cluster service account**
The cluster service account is the account under which the cluster service is started. The credentials for this account are stored in the service control manager (SCM) which is the Windows component that is responsible for starting the cluster service when the cluster nodes are booted.

The cluster service account must be the same on all nodes in a cluster and it must be a domain-level account that also has local administrative rights to each node in the cluster. The domain account must exist before the cluster is created and the Cluster Configuration Wizard will prompt you for an existing account to be used. If the account is not already a member of the local

Administrators group, the Cluster Configuration Wizard will automatically add it to the local Administrator group when the cluster is created. Likewise, when nodes are added to the cluster, the cluster service account will be added to the local Administrators group. If a node is evicted from a cluster or the last node is removed, the cluster service account is not removed from the local Administrators group.

You need to be aware of these semantics to avoid unintentionally granting a domain account local administrator rights to a given set of nodes.

| **Note** |
| --- |
| Evicting a node from a cluster does NOT remove the cluster service account from the local Administrators group. When you have removed a node from the cluster, you should manually remove the cluster service account from the local Administrators group to avoid having stale accounts with local administrator rights to a machine. |

The nodes in a Server cluster use authenticated communication mechanisms to ensure that only valid members of the cluster can participate in the intra-cluster protocols. It is essential that each node in the cluster has the same cluster service account in order to provide authentication consistency. It is also a requirement of the cluster service account password utility introduced in Microsoft Windows Server 2003.

**Required Privileges**
As well as being a member of the local Administrator group, the cluster service account requires a set of additional, locally granted privileges:

- Act as part of the operating system (required for Windows 2000 and beyond).
- Back up files and directories.
- Increase quotas.
- Increase scheduling priority.
- Load and unload device drivers.
- Lock pages in memory.
- Log on as a service.
- Restore files and directories.

You should NOT remove any of these privileges from the cluster service account. If you remove any of these required privileges, the cluster service may not start up or operate correctly.

During the cluster server setup process, these privileges are granted locally to the account. If you ever need to manually re-create the cluster service account, you must also grant these additional privileges. *KB article 269229: How to Manually Re-Create the Cluster Service Account* describes the steps necessary to re-create the cluster service account.

**Password Policies**
The cluster service account is just like any other domain account, it has a password that can have password expiration policies associated with it. If the password has expiration policies assigned to it, then the cluster account password must be changed before it expires. Failure to do so will cause the cluster to stop functioning when the password expires (since the intra-cluster communication can no longer be successfully authenticated).

**Changing the Password**
In most production deployments, domain accounts will have password expiration policies that force the password to be changed relatively frequently (for example every 30 days). Changing the cluster service account password requires careful planning.

**Windows 2000**
The cluster service account on ALL nodes in the cluster must match to ensure that the intra-cluster communication can be successfully authenticated. The cluster service itself sends messages between cluster nodes under a variety of conditions and if any of those communications fail, the cluster node will be removed from the cluster (i.e. the cluster service will be stopped). It is

not possible to determine when the cluster service will establish communication and therefore there is no clear window that allows the cluster service account to be changed in a reliable way while ensuring that the cluster remains running.

On Windows 2000, the cluster account password can only be reliably changed using the following steps:

1. Stop the cluster service on ALL nodes in the cluster
2. Change the password of the cluster service account at the domain controller
3. Update the service control manager password on ALL cluster nodes
4. Re-start the cluster service on all of the cluster nodes

**Windows Server 2003**

The cluster.exe command on Windows Server 2003 has the ability to change the cluster account password dynamically without shutting down the cluster service on any of the nodes. The cluster.exe command changes the domain account password and updates the service control manager account information on all nodes in the cluster.

| **Note** |
| --- |
| This command will only work for Windows Server 2003 nodes. If there are any Windows 2000 nodes in the cluster, their cluster service account password will not be changed. |

**Cluster /cluster**:*cluster_name1*[,*cluster_name2*,] **/changepassword**[:*new_password*[,*old_password*]] [**/skipdc**] [**/force**] [**/**options]

The command arguments have the following meanings:

| **Argument** | **Description** |
| --- | --- |
| **cluster**:*cluster_name1* [,*cluster_name2*,] | Identifies the cluster(s) for the account password change. If multiple clusters are specified, they must use the same cluster service account. If some nodes are unavailable, the password is not change on any of the nodes or at the domain controller. |
| **/changepassword** [:*new_password* [,*old_password*]] | Changes the cluster service account password on the domain controller and all cluster nodes from old_password to new_password. If the passwords are not supplied on the command line, you will be prompted to provide them. |
| **/skipdc** | Changes the cluster service account password only on the cluster nodes. |
| **/force** | Forces the execution of the password change on the available nodes even if some nodes are not available.<br>**Note**<br>This command will only work for Windows Server 2003 nodes. If there are any Windows 2000 nodes in the cluster, their cluster service account password will not be changed. |

The details of the command are documented in the on-line help available from the Cluster Administrator tool.

**Best Practices**

- The cluster service account should NOT be a domain administrator account.
    - All accounts should be given the minimal possible rights and privileges to avoid potential security issues if a given account is compromised.
    - Use delegation to give administration rights to specific accounts on ALL nodes of a cluster.
- If there are multiple clusters in a single domain, using the same cluster service account on all nodes makes administration easier.

    - You need to balance the ease of management with the potential security risks associated with using a single account for many clusters. If the account is compromised, the scope of the impact depends on how many clusters are involved.

    - With Windows Server 2003, the cluster service account password can be changed on multiple clusters at the same time.

- If you have set password expiration policies on your cluster service accounts you should NOT use this account for other services.

    - Do not use the cluster service account for SQL Server 2000 or Exchange 2000, for example, if you have set password expiration policies. If you have multiple services using the same account, coordinating the password change across the cluster service and other services is complex and will lead to the complete cluster and/or service being unavailable during a password rotation. It is better to have a dedicated account for each service, each of which can be maintained independently.

    - With Windows Server 2003, the cluster service account password can be changed online without taking the cluster down ONLY if there are no services using the same service account.

- Changing the cluster service account password on Windows 2000 requires the cluster to be shut down completely before the account password is changed. Shutting down and re-starting the cluster may mean that the cluster cannot meet the availability requirements. For example 99.999% uptime requires less than five minutes downtime *per year* for the applications and services. This level of availability cannot be achieved if the cluster is shutdown for password rotation. In these high availability environments, with Windows 2000, the cluster service account should have a password expiration policy of never expiring.

- With Windows 2000 SP3 and Windows Server 2003, the cluster service can publish virtual servers as computer objects in Active Directory. To ensure correct operation, the cluster service account should have appropriate access rights or privileges to be able to create and manipulate these objects in the Active Directory Computers container.

  See section Using Kerberos Authentication in a Server cluster.

- If you intend to deploy multiple clusters using different cluster service accounts, you should create a Global or Universal group which implements all of the policies and has all of the privileges described above. Each cluster service account should then be placed into the group.

    - This eases management of the cluster service accounts by providing a single container for all cluster service accounts and a single point of management for changing account policies.

**Using Kerberos Authentication in a Server cluster**

Kerberos authentication was released as part of the Windows platform in Windows 2000. It is the primary (and default) security mechanism for the Windows platform moving forward and it has a number of benefits over the previous authentication mechanisms (such as NTLM):

- **Provides mutual authentication of client and server**: The server ensures that the client has access to the service or applications provided at the server and the client can be sure that the server it is communicating with really is the server that it was expecting it to be.

- **Allows delegation of authentication across multiple machines**: Allows end-to-end impersonation based on the credentials associated with the original request in a multi-tiered application deployment. For example, a web site may contain an IIS web server front-end, business objects at the middle-tier and a database back-end. Kerberos allows the original client authentication to be carried across all of the tiers from the web server to the database for end-to-end authentication and authorization.

- An industry standard allowing a common authentication mechanism across multiple platforms.

For more information about Kerberos and how it works, see the TechNet web site:
http://www.microsoft.com/windows2000/techinfo/howitworks/security/kerberos.asp

In a Server cluster, applications are deployed in *virtual servers*. A virtual server is a Server cluster resource group that contains an IP Address resource and a Network Name resource as well as any resource required for a given service or application. Clients connect to the clustered service using the IP Address or Network Name associated with the service. When an application fails over from one node to another, so do the IP Address and Network Name resources, thus the client continues to use the same target location for a service or application regardless of which machine in the cluster is currently hosting it.

For Server clusters in Windows 2000 (up to SP2), the only mechanism available for authenticating a client against a virtual server name is NTLM, thus the benefits provided by Kerberos are not available to deployments that contain Server clusters.

In Windows 2000 SP3 and Windows Server 2003, Kerberos authentication will be available to cluster server applications, allowing end-to-end impersonation in a highly available deployment.

**Virtual Server Computer Objects**
Kerberos authentication requires a backing computer object (CO) to be published in Active Directory. In Windows 2000 SP3 and beyond, the network name cluster resource has the ability to publish a computer object in Active Directory for the virtual server. This is controlled through the new resource property **RequireKerberos** with the following values:

- RequireKerberos = 0

  The Network Name resource does not create an object in Active Directory. This provides the same semantics as Windows 2000 (up to SP2) and is the default value for backward compatibility as well as ensuring correct behavior during rolling upgrade.

  (In the case of a network name that is required for MSMQ, RequireKerberos is enabled by default to ensure that MSMQ continues to work as expected after the upgrade is complete.)

- RequireKerberos = 1

  The Network Name resource does create an object in Active Directory thus enabling Kerberos authentication against the virtual server computer name.

The full semantics for this property are described later in this section.

| **Note** |
| --- |
| Although the Network Name resource publishes a computer object in Active Directory, that computer object should NOT be used for administrations tasks such as applying Group Policy. The ONLY role for the virtual server computer object in Windows 2000 and Windows Server 2003 is to allow Kerberos authentication and delegation and for cluster and Active Directory-aware services (such as MSMQ) to publish service provider information. |

**The Lifetime of a Computer Object**
Network Name resource controls the creation and deletion of computer objects as dictated by a number of factors including resource property settings and the resources own lifetime.

**Creating Computer Objects**
The Network Name resource creates a computer object when the RequireKerberos resource property is set to one (the default value is zero). Setting this value to one requires that the cluster service account has access to Active Directory either by having the *Add Workstations to the Domain* privilege or by granting specific access to Active Directory. If the account does not have access, the Network Name resource will fail to come on line.

Domain administrators can pre-create computer objects for virtual servers (in a separate OU if required) to avoid giving the *Add Workstations to the Domain* privilege to the cluster account. In this case, the computer object must have an Access Control Limit (ACL) that allows the cluster service account to reset the password and write to the DnsHostName and ServicePrincipalName attributes (which have individual access rights).

If an orphan computer object or a computer object created by the domain administrator exists in Active Directory corresponding to the virtual computer name, and the cluster service account has the appropriate access rights to the object and the object is not disabled, the network name resource will hijack the existing computer object. This is the same behavior that is visible if a new computer is added to a domain that contains an old and unused computer object with the same name.

The Network Name resource sets the DnsHostName attribute of the computer object to the fully qualified DNS name of the Network Name resources Name property and the primary DNS suffix of the node. (All nodes in a cluster must be in the same NT and DNS domain.)

**Disabling Computer Objects**
Computer objects are disabled when either Kerberos authentication is no longer required (the RequireKerberos property is changed from one to zero) or when the resource itself is deleted. If the computer object still exists, but the Network Name resource is brought online without Kerberos authentication enabled, client authentication will fail. When removing Kerberos authentication from a Network Name resource, the system administrator must make an explicit decision to delete the corresponding computer object. Note that Active Directory-aware applications hosted in the virtual server (such as Exchange server and MSMQ) may have attached their own information to the computer object. If the computer object is deleted, those properties will also be deleted and the applications may no longer function correctly. Extreme care should be taken when removing Kerberos authentication from a network name.

When the Network Name resource is deleted or Kerberos authentication is removed, the Network Name resource makes a one-time best effort to disable the computer object. All disable operations, along with their final status, are logged in the cluster and system event logs. If disabling the computer object fails for any reason, no other attempt is made to disable it.

**Renaming Computer Objects**
The computer object is intimately tied to the virtual network name exposed by the Network Name resource. If the *Name* property of the network name resource is changed, that change must be reflected in the computer object. These changes are tightly coupled and synchronized; for that reason, changing the Name property can occur only when the Network Name resource is offline. If renaming the computer object fails for any reason, the change to the Name property is also failed.

There is one caveat to this: the Network Name resource allows property values to be changed while the Network Name resource is online; in this case, the rename operation is deferred until the name is taken offline. If that rename attempt fails, the operation is retried as part of the online process. The name will not be brought online until the rename operation is successful.

| **Note** |
| --- |
| Many applications and services such as MSMQ and SQL Server do NOT support changing the *Name* property on the Network Name resource. |

**Password Rotation**
The computer object maintains a password used for authentication when issuing Kerberos tickets. In the current implementation for Windows 2000 SP3 and Windows Server 2003, the password on the computer object is not changed once the computer object is created.

**Caveats, Issues and Considerations**
**Orphan Computer Objects**
Many operating system services and applications use the *Negotiate* security package which is a component provided by the Windows platform that allows a client to authenticate with a service. This package hides the particulars of any given authentication scheme and allows a client to try Kerberos authentication, and if that is unavailable, it falls back to NTLM authentication. If there is a computer object in Active Directory that corresponds to the target name that the client is using, the negotiate package will assume that Kerberos authentication is *required*. (In an all Windows Server 2003 domain this will work as expected.) If a virtual computer object exists in Active Directory (disabled or not) future client connections to that virtual server name will NOT fail-back to NTLM.

If authentication fails against a virtual server and that virtual server should be using NTLM, then you should check whether there is an orphaned or old computer object in active directory corresponding to the virtual server name.

**Cluster-Aware and Active Directory-Aware Services**
There are some services that are both cluster-aware and Active Directory-aware, in particular, SQL Server 2000, Exchange 2000 and MSMQ. Some of these services attach service information to the computer objects. MSMQ in particular attaches service information to the virtual server computer object.

When removing Kerberos authentication from a Network Name resource, the computer object will be disabled, leaving the system administrator with an explicit decision to delete the computer object. If the computer object is deleted, properties attached by Active Directory-aware applications will also be deleted and the applications may no longer function correctly. Extreme care should be taken when removing Kerberos authentication from a Network Name resource.

**Active Directory Replication Delays**
Thus far, we have talked about Active Directory as though it were a single, highly consistent infrastructure. In a production environment, multiple domain controllers are deployed to ensure that the domain infrastructure is highly available. Changes to

Active Directory are replicated across nodes and in some cases that replication delay can be large. This can cause some seemingly inconsistent effects that you will need to be aware of:

- **Clients cannot see the computer object for a network name resource**

  Different nodes may access different domain controllers. If a computer object is created on one domain controller, it may not be visible to the clients until it is replicated to other domain controllers.

**Best Practices**

- Plan your use of Kerberos authentication in a cluster carefully:

  - Enable Kerberos authentication on those virtual servers that make sense for your deployment.

  - You should not disable Kerberos authentication on a virtual server unless you fully understand the services using that virtual server and the implications that will have on the service. The computer object and any service information attached to it WILL be deleted (assuming correct privileges) which may lead to a service being unavailable or unrecoverable.

- The cluster service account should have the *Add workstations to the domain* privilege to allow it to create computer objects in Active Directory. If you do not wish to grant that privilege to the cluster service account, you MUST create the computer object by hand in Active Directory before enabling Kerberos authentication.

| **Note** |
| --- |
| The *Add workstations to the domain* privilege allows a default of 10 computer objects to be created from that account. In addition, the *Add workstations to the domain* by itself does NOT allow the account to delete or rename the computer object. |

- The cluster service account should have *Write all properties* access rights to allow it to rename the computer object.

- Many services, including MSMQ and SQL Server 2000 do not support changing the Network *Name* resource Name property. You should only ever change this property if you fully understand the implications. In some cases changing the Network Name resources *Name* property can lead to loss of data or failure of a service.

- You should ensure that the domain controllers have had a chance to replicate newly created computer objects associated with Network Name resources before allowing clients to access them or before failing over the resource to another node in the cluster. Failure to do this can lead to unpredictable results (see section Active Directory Replication Delays).

**Network Security**
**Network Flooding**
The cluster service uses UDP port 3343 for intra-cluster communication. This communication includes heartbeat traffic to detect node failure and cluster control operations. Some of these operations (such as heartbeats) are time-sensitive. If there is significant load on the ports due to network flood attacks, it can result in false node failure detection and therefore cause unnecessary failover operations which will lead to application downtime.

**Port Squatting**
It is possible for a rogue application to hijack or squat on port 3343. This will stop the cluster service from starting up. In this case, the only option is to kill processes using the port. Port 3343 is registered exclusively for the cluster service.

**Rogue Servers**
The cluster service provides remote management capabilities through the cluster APIs so that a cluster can be managed from a management station. The cluster APIs use NTLM to authenticate with the server. This allows the server to authenticate that the client has sufficient rights and privileges to manage the cluster, however, it does not provide mutual authentication. In other words, the client has no cast-iron guarantee that the node or cluster that the client connected to is the real cluster. If a rogue machine were to appear on the network with the same IP address or network name (if the DNS information were compromised), then the rogue machine could masquerade as the cluster. If that rogue machine also appeared to implement the cluster service APIs, any management commands sent to the cluster would be intercepted by the rogue machine. In many cases, this does not represent a threat since the administrator would simply receive a false positive or negative that an action

was performed. However, there are sensitive operations (such as changing the cluster configuration) where an unauthorized recipient could potentially collect cluster configuration data which may help to extend the attack surface.

This type of attack requires a number of factors to occur:

1. The rogue computer must be visible on the same subnet as the target cluster and respond to traffic to the cluster IP address (this may be the IP address of the physical computer or may be the IP address of virtual servers hosted by the computer).

2. In a typical environment, the rogue computer can only take over the IP address if the cluster node or virtual server is not running.

3. The rogue computer must appear to implement the cluster APIs. For a typical management operation, the client application makes several calls to the server and in some cases, returns handles for subsequent calls.

4. The administrator/operational procedure must change the configuration that includes the sensitive data (e.g. the administrator must change the cluster account password).

**Best practices**
Attacks that compromise the cluster typically involve being on the same subnet as the cluster. To protect against these attacks, the subnets should be protected:

- **Client-access networks**

  The subnet used by the cluster nodes should not extend beyond the set of nodes that can be physically secured or trusted. You must ensure that rogue or potentially insecure machines cannot be attached to the subnet containing the clusters. Since this network may contain domain controllers, DNS servers, WINS servers, DHCP servers and other network infrastructure, you must take steps to ensure that these infrastructure servers are also secure.

  Typical network security procedures should be put in place such as firewalls that only allow specific application requests etc. Securing networks and infrastructure servers is beyond the scope of this document.

- **Private networks**

  Only nodes in the cluster should be visible on a private network (multiple clusters can use the same private network). No other network infrastructure servers or other application servers should be on the private subnet. This can be achieved by either:

  - Physically constraining the network itself (e.g. the private network is a LAN that is only connected to the cluster nodes).

  - Isolating the cluster private networks using VLAN-capable switches.

**NTLM V1 and NTLM V2**
**Windows 2000**
Windows 2000 clusters (up to and including SP2) must use NTLM V1 authentication between cluster nodes. There are several lockdown tools and policies (such as HiSecDC) that apply different policies to the nodes, forcing NTLM V2 as the default security profile. If these policies or tools are applied to a Windows 2000 cluster, the cluster service will fail. The default security profile must be reset to Send LM and NTLM responses. This is detailed in KB articles 295091 and 272129.

**Windows Server 2003 and Windows 2000 SP3 (and above)**
In Windows 2000 SP3 and above, as well as Windows Server 2003, the cluster service is capable of using NTLM V2. Security policies that enable NTLM V2 will not compromise the cluster.

**NetBIOS**
Security of a system is dependent on the number of ways into the system. In Windows Server 2003, the cluster service does not require NetBIOS, however a number of services are affected if NetBIOS is disabled. You should be aware of the following:

- By default, when a cluster is configured, NetBIOS is enabled on the cluster IP Address resource. Once the cluster is created you should disable NetBIOS by clearing the check box on the parameters page of the Cluster IP Address resource

property sheet.

- When you create additional IP Address resources you should clear the NetBIOS checkbox.

- With NetBIOS disabled, you will not be able to use the Browse function in Cluster Administrator when opening a connection to a cluster. Cluster Administrator uses NetBIOS to enumerate all clusters in a domain.

- Print and File services are disabled no virtual names are added as redirector endpoints.

- Cluster Administrator does not work if a cluster name is specified. Cluster Administrator calls GetNodeClusterState which uses the remote registry APIs which, in turn, use named pipes based on the virtual name.

**IPSec**

Although it is possible to use Internet Protocol security (IPSec) for applications that can failover in a Server cluster, IPSec was not designed for failover situations and we recommend that you do NOT use IPSec for applications in a Server cluster.

The primary issue is that Internet Key Exchange (IKE) Security Associations (SAs) are not transferred from one server to the other if a failover occurs because they are stored in a local database on each node.

In a connection that is protected by IPSec, an IKE SA is created in phase-I negotiations. Two IPSec SAs are created in phase II. A time-out value is associated with the IKE and IPSec SAs. If Master Perfect Forward Secrecy is not used, the IPSec SAs are created by using key material from the IKE SAs. If this is the case, the client must wait for the default time-out or lifetime period for the inbound IPSec SA to expire and then wait for the timeout or lifetime period that is associated with the IKE SA.

The default time-out for the Security Association Idle Timer is five minutes. In the event of a failover, clients will not be able to reestablish connections until at least five minutes after all resources are online, using IPSec.

Although IPSec is not optimally designed for a clustered environment, it may be used if your business need for secure connectivity outweighs client downtime in the event of a failover.

**Cluster Disks**

In general, cluster disks (i.e. disks that have a corresponding resource in the cluster configuration) are just like any other disks hosted by Windows; however, there are some additional considerations that you need to understand:

**General Best Practices**

- Cluster server only supports the NTFS file system on cluster disks. This ensures that file protection can be used to protect data on the cluster disks. Since the cluster disks can failover between nodes, you must only use domain user accounts (or Local System, Network Service or Local Service) to protect files. Local user accounts on one machine have no meaning on other machines in the cluster.

- Cluster disks are periodically checked to make sure that they are healthy. The cluster service account MUST have write access to the top level directory of all cluster disks. If the cluster account does not have write access, the disk may be declared as failed.

**Quorum Disk**

- The quorum disk health determines the health of the entire cluster. If the quorum disk fails, the cluster service will become unavailable on all cluster nodes. The cluster service checks the health of the quorum disk and arbitrates for exclusive access to the physical drive using standard I/O operations. These operations are queued to the device along with any other I/Os to that device. If the cluster service I/O operations are delayed by extremely heavy traffic, the cluster service will declare the quorum disk as failed and force a regroup to bring the quorum back online somewhere else in the cluster. To protect against malicious applications flooding the quorum disk with I/Os, the quorum disk should be protected. Access to the quorum disk should be restricted to the local Administrator group and the cluster service account.

- If the quorum disk fills up, the cluster service may be unable to log required data. In this case, the cluster service will fail, potentially on all cluster nodes. To protect against malicious applications filling up the quorum disk, access should be restricted to the local Administrator group and the cluster service account.

- For both reasons above, the quorum disk should NOT be used to store other application data.

**Cluster Data Disks**

- As with the quorum disk, other cluster disks are periodically checked using the same technique. If malicious applications flood the cluster application disks with I/Os, the cluster service health check may fail, thus causing the disk (and any applications dependent on the disk) to be failed over to another cluster node. To avoid denial of service attacks like this, access to the cluster disks should be restricted to those applications that store data on the specific disks.

**EFS and Server clusters**
With Windows Server 2003, the encrypting file system (EFS) is supported on clustered file shares. To enable EFS on a clustered file share, you must perform a number of tasks to configure the environment correctly:

1. EFS can only be enabled on file shares when the virtual server has Kerberos enabled. By default, Kerberos is not enabled on a virtual server. To enable Kerberos you must check the Enable Kerberos Authentication check box on the network name resource that will be used to connect to the clustered file share.

   | **Note** |
   | --- |
   | Enabling Kerberos on a network name has a number of implications that you should ensure you fully understand before checking the box. |

2. All cluster node computer accounts, as well as the virtual server computer account, must be trusted for delegation. See online help for how to do this.

3. To ensure that the users private keys are available to all nodes in the cluster, you must enable roaming profiles for users who want to store data using EFS. See online help for how to enable roaming profiles.

Once the cluster file shares have been created and the configuration steps above carried out, users data can be stored in encrypted files for added security.

**Managing file shares in a cluster**
**Normal file shares**
Normal file shares are the most flexible and easily understood in terms of security. The only real difference is that you administer share level security using the cluster user interface instead of Windows Explorer. You administer NTFS security settings on files on the cluster disks using standard Windows tools such as Windows Explorer. For more information about administering cluster file shares, see the online documentation for Server clusters.

File shares created through Cluster Administrator are created with the same default protection as a single node unless the ACL is set using the cluster administration tools. See file services documentation for default ACL for a file share.

| **Note** |
| --- |
| Always use the cluster administration tools to change security of a clustered file share. If you use the file share management tool, any configuration will be lost when the file share is failed over to another node in the cluster. |

**Shared Subdirectories**
Subdirectory shares are available in versions of Windows later than Windows NT 4.0 Service Pack 4. Shared subdirectories allow administrators to rapidly create directories to host large numbers of shares such as home directories with a single cluster resource. A root share is specified, and all of the subdirectories one level below the specified root are created as regular file shares.

There is no way to specify different security attributes for each shared subdirectory and therefore each share inherits the same share level permissions as the root share. Because each share is typically used by a different user, the share-level permissions should be left to everyone, and file-based security at the file system level should be used to control who has access to which files.

**DFS roots**
DFS roots are available in Windows 2000 onwards. Server clusters only support standalone DFS roots as cluster resources. (Note that does NOT mean you cannot create domain roots on cluster nodes since a cluster node is just like any other server,

however, they are not cluster resources.) You can use share level permissions for the root through the Cluster Administrator user interface and you can administer each link through file share permissions on the appropriate server. However, this method of controlling access can be difficult for DFS trees spanning a large number of servers and links. We recommend you administer DFS trees by leaving file share level permissions open and use NTFS file system permissions to restrict access.

**Cluster Server nodes as Domain Controllers**
To have Server clusters function properly (where the cluster service starts on each node), the node that is forming the cluster must be able to validate the cluster service domain account, which is the account that you configure when the cluster is configured. To accomplish this, each node must be able to establish a secure channel with a domain controller to validate this account. If the account cannot be validated, the cluster service does not start. This is also true for other clustered programs that must have account validation in order for services to start, such as Microsoft SQL Server and Microsoft Exchange.

If a cluster deployment is such that there is no link with either a Windows NT 4.0 domain, a Windows 2000 domain or a Windows Server 2003 domain, you have to configure the cluster nodes as domain controllers so that the cluster service account can always be validated to allow for proper cluster functionality.

If the connectivity between cluster nodes and domain controllers is such that the link is either slow or unreliable, consider having a domain controller co-located with the cluster, or configuring the cluster nodes as domain controllers. Microsoft does not recommend using cluster nodes as domain controllers.

If you must configure the cluster nodes as domain controllers, consider the following important notes:

- If one cluster node in a 2-node cluster is a domain controller, all nodes must be domain controllers. It is recommended that you configure at least two of the nodes in a 4-node Datacenter cluster as domain controllers.

- There is overhead that is associated with the running of a domain controller. A domain controller that is idle can use anywhere between 130 to 140 megabytes (MB) of RAM, which includes the running of Server clustering. There is also replication traffic if these domain controllers have to replicate with other domain controllers within the domain and across domains. Most corporate deployments of clusters include nodes with gigabytes (GB) of memory so this is not generally an issue.

- If the Windows 2000 cluster nodes are the only domain controllers, they each have to be DNS servers as well, and they should point to each other for primary DNS resolution, and to themselves for secondary DNS resolution. You have to address the problem of the ability to not register the private interface in DNS, especially if it is connected by way of a crossover cable (2-node only). For information about how to configure the heartbeat interface refer to KB article 258750. However, *before* you can perform step 12 in KB article 258750, you must first modify other configuration settings, which are outlined in KB article 275554.

- If the cluster nodes are the only domain controllers, make them all Global Catalog servers. For background information about the placement of Global Catalog servers, see:

  http://go.microsoft.com/fwlink/?LinkId=91432

- The first domain controller in the forest takes on all flexible, single master operation roles (refer to KB article 197132). You can redistribute these roles to each node. However, if a node fails over, the flexible, single master operation roles that the node has taken on are no longer available. You can use Ntdsutil to forcibly take away the roles and assign them to the node that is still running (refer to KB article 223787). Review KB article 223346 for information about placement of flexible, single master operation roles throughout the domain.

- If a domain controller is so busy that the cluster service is unable to gain access to the quorum drive as needed, the cluster service may interpret this as a resource failure and cause the cluster group to fail over to the other node. If the quorum drive is in another group (although it should *not* be), and it is configured to affect the group, a failure may move all group resources to the other node, which may not be desirable. For more information regarding Quorum configurations, please refer to the KB article 280345 listed in the "Reference" section.

- Clustering other programs such as SQL or Exchange in a scenario where the nodes are also domain controllers, may not result in optimal performance due to resource constraints. You should thoroughly test this configuration in a lab environment prior to deployment.

- You must promote a cluster node to a domain controller by using the Dcpromo tool *prior* to creating a Server cluster or

adding a node to the cluster.

- You must be extremely careful when demoting a domain controller that is also a cluster node. When a node is demoted from a domain controller, the security settings and the user accounts are radically changed (user accounts are demoted to local accounts for example).

**Generic resources**

The generic resources types, Generic Application, Generic Service and Generic Script, allow existing non-cluster aware applications to be monitored and failed over with little or no effort. (Generic script is available on Windows Server 2003 only.) A file name path is used to identify the application or script. The cluster service runs the resource DLL under the cluster service account and therefore they are run with elevated (local administrator) privileges.

When using a generic resource you should ensure that the file containing the code/script and any registry keys used by the script, application or service are protected against attacks.

- Generic Script
    - Make sure script files are secure using NTFS file system protection.
    - Specify the fully qualified path to the script to avoid spoofing issues associated with the PATH variable.
- Generic Application
    - Make sure the application is trusted and that the files, registry keys that need to be checkpointed and any other resources needed for the application to run are secure.
    - Specify the fully qualified path to the application to avoid spoofing issues associated with the PATH variable.
- Generic Service
    - Make sure the service is trusted and that any registry keys or other resources needed for the service to run are secure.

| **Note** |
| --- |
| Typically, parameters and configurations for Windows services are stored under HKEY_LOCAL_MACHINE in the registry. This document assumes that this registry hive is protected against attacks. The default security attributes of the HKLM hive ensure that the hive is secure. |

**Summary of security attributes**

This section outlines the various objects and their associated security attributes needed to ensure that the cluster service runs successfully. The section lists the minimum security requirements. If the security attributes are made more restrictive, the cluster service may not run. If the security attributes are less restrictive, it may result in malicious attacks being able to compromise the cluster or corrupt end-user data. We recommend that the security attributes on the objects not be changed from the default settings.

**Directory and file protection**

| Object | Description | Minimum recommended security attributes |
| --- | --- | --- |
| %windir%\help | Cluster Help files | BUILTIN\Users:R<br>BUILTIN\Administrators:F |

| | | |
|---|---|---|
| %windir%\cluster | Cluster directory (and subdirectories) | BUILTIN\Administrators:F<br>BUILTIN\Administrators:(OI)(CI)(IO)F<br>NT AUTHORITY\SYSTEM:F<br>NT AUTHORITY\SYSTEM:(OI)(CI)(IO)F<br>CREATOR OWNER:(OI)(CI)(IO)F |
| %windir%\cluster* | Files in cluster directory | BUILTIN\Administrators:F<br>NT AUTHORITY\SYSTEM:F |
| <quorum_drive> | Volume holding quorum data | BUILTIN\Administrators:F<br>BUILTIN\Administrators:(OI)(CI)(IO)F<br>NT AUTHORITY\SYSTEM:F<br>NT AUTHORITY\SYSTEM:(OI)(CI)(IO)F |
| <quorum_drive>:\MSCS | Quorum directory (and subdirectories) | BUILTIN\Administrators:F<br>BUILTIN\Administrators:(OI)(CI)(IO)F<br>CREATOR OWNER:(OI)(CI)(IO)F |
| <quorum_drive>:\MSCS* | Files in quorum directory | BUILTIN\Administrators:F |
| | Volumes on cluster disks | BUILTIN\Administrators:F<br>BUILTIN\Administrators:(OI)(CI)(IO)F<br>NT AUTHORITY\SYSTEM:F<br>NT AUTHORITY\SYSTEM:(OI)(CI)(IO)F |

**System resources**

The cluster service uses other system resources that are protected by default.

| Object | Description | Minimum recommended security attributes |
|---|---|---|
| HKLM* | HKLM registry hive | SYSTEM:Full control<br>BUILTIN\Administrators: Full control |

**Cluster Service Account Policies**

- Do not let the password expire. On Windows Server 2003 you can use the password change mechanism to ensure that the password is periodically cycled.
- Use strong passwords for the cluster service account.

**Active Directory**

- To ensure Kerberos authentication functions as expected, the cluster service account must have access to Active Directory. This is covered in the section Using Kerberos Authentication in a Server cluster.

**DNS**

- The cluster service account needs to be able to publish records. In a secure, DNS backed zone, the DNS administrator can chose to restrict the access rights for users. The cluster service account must be granted permission to create records or alternatively, the records can be pre-created. If the records are pre-created, you should not set the zone to dynamic updated.

**Majority Node Set considerations**
Windows Server 2003 provides a new quorum-capable resource known as Majority Node Set. The primary goal of this resource is to keep multiple copies of data spread across a cluster in sync at all times. It can be used to provide a quorum resource in a cluster without using a shared disk for the quorum data and is primarily targeted at the following scenarios:

- Geographically dispersed clusters to provide a single, common quorum mechanism for multi-site cluster configurations to allow vendors to build solutions without having to be concerned about quorum requirements.

- Clusters of appliances set of off-the-shelf machines that have no shared disks being bound into a single cluster using some other data replication technique for application and/or user data.

In a majority node set cluster there is a single majority node set resource. It is responsible for committing changes to the quorum data stored on each cluster node. Each cluster node has its own copy of the data. To access the data on the other nodes in a cluster (i.e. the ones other than the node that the resource is hosted on), the Majority Node Set resource uses the file share infrastructure. When a Majority Node Set resource is created, a file share is created on each cluster node (and on nodes that are added to the cluster).

This file share is a hidden file share that has a name constructed as follows:

```
\\<node_name>\<resource_GUID>$
```

The default security setting for this file share is:

```
BUILTIN\Administrators:F (this folder, subfolders and files)
CREATOR OWNER:F (Subfolders and files only)
```

To ensure the correct operation of a majority node set cluster:

- Do NOT delete the file share.

- Do NOT remove the cluster service account from the set of accounts that have access to the share.

- Do NOT change the default access controls for the file share

- The file share target is a directory in the %windir%\Cluster\MSCS directory. You should not change the default access permissions on the files or the directories.

- You should never put other files into the MNS target directory %windir%\Cluster\MSCS\MNS.<guid-representation>. If you put other files in this directory, they will be deleted on cleanup if the MNS resource is deleted.

**Upgrading to Windows Server 2003 considerations**
Windows 2000 and Windows NT4 have default security attributes on the %windir%\Cluster directory and the quorum directory that allow any authenticated user to read the contents. The security of the directories has been tightened in Windows Server 2003 to stop non-administrator access altogether to ensure that unauthorized users cannot gain information about the cluster configuration. On an upgrade to Windows Server 2003, however, these security attributes are not modified and therefore on an upgraded Windows Server 2003 machine, all authenticated users may have read access to these directories. You may want to consider manually setting the permissions to conform to the minimum access requirements.

**Developing Cluster-aware Applications**
Resource DLLs run within the context of the cluster service account. This account has a number of elevated privileges as well as being a member of the local Administrators group on every node in the cluster. When developing a cluster-aware application, you should consider how best to split the application between the resource DLL and an independent service or executable to ensure that the application runs with the minimal privileges and rights necessary.

Running applications and services at unnecessary privilege levels introduces a potential security risk if that application is compromised.

**Calling the Server cluster APIs**
The Server cluster APIs are protected so that arbitrary, untrusted users cannot affect the state of the cluster or the availability of applications. The cluster service maintains a security descriptor in order to control access. Only accounts that are part of the security descriptor are able to call cluster APIs (actually, the security descriptor controls which accounts can open a handle to the cluster, since other cluster server APIs require a handle to the cluster, this effectively limits access).

There is one caveat with this mechanism. The cluster configuration mechanism provided in Windows Server 2003 requires that a user who adds a node to a cluster has local administrator rights on every node in the cluster. If the user does not have local administrator rights on the existing cluster nodes, that user will not be able to successfully add a new node to the cluster. Thus, to give an account administration rights to a cluster you should:

1. Add the account to the cluster service security descriptor.

2. Make sure that the account is a member of the local Administrators group on ALL machines in the cluster.

These operations can be done programmatically using the cluster APIs and the security APIs. See the platform SDK for more details.

Security checks are done by the OpenCluster API. If the call returns successfully, the caller has a valid handle and can make any changes to the cluster configuration or enumerate the cluster configuration. There is no notion of open for read or open for write.

**Backup/Restore APIs**
The backup path specified to the backup API should be protected. The data returned by the API contains cluster configuration information that can be used to extend the attack surface (e.g. lists all cluster IP addresses).

During a restore operation, the backup API takes the cluster configuration to inject. This data must be protected as it will be used to restore the cluster by overwriting any configurations that currently exists.

**Installing resources**
Cluster aware applications may install cluster resources. Care should be taken when installing resources since they will be executed at elevated privilege and they must, therefore be protected against malicious attack by default when installed:

- Ensure that resource DLLs are protected at the file system level to allow local Administrator and the cluster service account access only.

- Be extremely careful about any paths that are specified. There are several well-known issues with paths; especially paths that contain spaces (follow recommendations on page 419 of Writing Secure Code). When in doubt, specify a fully qualified path to the resource dll when creating the resource type.

# General Assumptions

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are a number of general assumptions and operational best practices that should be in place for the infrastructure to ensure a secure environment in which to run Server clusters.

1. Servers and storage are in physically secure locations.

2. Practical security implementations to detect irregular traffic, such as firewalls, network probes and management tools, are in place.

3. Best practices/common sense in terms of security are adhered to in areas like administration, storage of logs, backup and restore etc.

4. Platform level security best practices are adhered to in terms of assigning administrative permissions, ACLing resources, and other housekeeping roles.

5. The network infrastructure services such as Active Directory, DNS, DHCP, WINS etc. must be secure. Any compromise of these infrastructure services can lead to a compromise of the cluster service itself.

6. The cluster administrator must ensure that applications that call the cluster APIs (ClusAPI) are run from trusted computers. Any compromise on the computers on which the applications are executing (that the cluster administrator runs) can compromise the cluster. For example, if there are untrusted users with elevated privileges on the workstation from where the administration tools are run, untrusted or malicious code can be run against the cluster by the cluster administrator without the cluster administrator realizing.

7. Access to the set of objects created and maintained by the cluster service must not be compromised by adjusting the default settings placed on these objects to a less restrictive setting. The cluster service utilizes a set of objects in the operating system such as files, devices, registry keys etc. These objects have a default security setting that ensures non-privileged users cannot impact the cluster configuration or the applications running on the cluster. Changing these security settings to less restrictive security settings can lead to the cluster being compromised and the application data being corrupted.

# Web and Application Server Infrastructure - Performance and Scalability

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Published: April 1, 2003

By Bill Karagounis, Microsoft Corporation

**In This White Paper**

Introduction (Web and Application Server Infrastructure - Performance and Scalability)

Programming Considerations (Web and Application Server Infrastructure - Performance and Scalability)

General Performance Considerations (Web and Application Server Infrastructure - Performance and Scalability)

Dedicated Application Scenario (Web and Application Server Infrastructure - Performance and Scalability)

Dense Hosting Scenario (Web and Application Server Infrastructure - Performance and Scalability)

Server Consolidation Scenario (Web and Application Server Infrastructure - Performance and Scalability)

Summary (Web and Application Server Infrastructure - Performance and Scalability)

# Introduction (Web and Application Server Infrastructure - Performance and Scalability)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Microsoft Windows Server 2003 incorporates a significantly different Web and application server infrastructure than previous releases of Windows Server. In fact, the key theme of the release is to bring forward the fundamental capabilities of the Web application platform, with a view to laying a solid foundation for existing and future server-based Windows .NET Framework applications and Web services.

A key part of bringing forward the fundamentals of the platform is to also provide significant business benefit for existing applications and Web sites currently running on Microsoft Windows NT 4.0 and Windows 2000, when migrated to Windows Server 2003. An increase in capacity and/or scalability for these applications can be valuable from a Total Cost of Ownership (TCO) perspective (by decreasing the number of servers used or by improving capacity on existing servers). In addition to ensuring that Windows Server 2003 was optimized for Windows .NET Framework applications, specific engineering effort was put into ensuring that existing ISAPI and ASP/COM+ applications would also see a performance and scalability gain on Windows Server 2003. This paper covers performance tuning considerations for ISAPI and ASP/COM+ applications, in addition to providing tuning ideas for the Windows .NET Framework.

**Audiences for This Document**
This document is intended for operations staff, IT administrators, developers, and architects. It is broken out into four main sections:

1. Architectural Discussion on Web performance features and capabilities of Windows Server 2003
2. Programming Considerations
3. General Performance Considerations
4. Application Scenarios

For all audiences, the architectural discussion section is a useful primer. It lays the foundation and introduces concepts elaborated on in the rest of the paper.

Developers and architects should then focus on the Programming Considerations section, followed by the General Performance Considerations section.

IT administrators and operational staff should approach this paper from either the Application Scenarios section (selecting whichever scenario applies), or the General Performance Considerations section.

**Performance and Scalability a New Architecture**
From the outset, the attributes of performance and scalability were some of the key drivers pushing the new application model and design tradeoffs in Windows Server 2003 development.

Much of the application server functionality in Windows is layered on top of Internet Information Services (IIS)[1]. This discussion will begin with a comparison of IIS architecture past and present.



**Figure 1: IIS 5.0 (Windows 2000) processing architecture**

IIS 5.0 was designed to have one process, Inetinfo.exe, be the main Web server process which could farm requests out to one or more out-of-process applications, DLLHost.exe (Figure 1). In comparison, IIS 6.0 has been redesigned into three new components, which allows IIS to segment Web server code from application handling code. These three new components are: a kernel-mode HTTP listener, called HTTP.sys; a user-mode configuration and process manager, called the WWW Service Administration and Monitoring component; and the application handler, which is loaded into each worker process[2]. These worker processes, in turn, service requests for application pools[3] in HTTP.sys.

**HTTP.sys and Kernel-Mode Queuing**

| **Important** |
| --- |
| There is a very subtle point here that has significant importance. All TCP/IP connections established to IIS 6.0 are held by the kernel- mode HTTP driver (HTTP.sys). |

In this new architecture, regardless of what happens to the user mode custom application, the IIS 6.0 kernel component preserves established TCP/IP connections. Therefore, if a rich client application connects to a Web service running on Windows Server 2003, and the Web service (for whatever reason) crashes or throws a fatal exception, IIS 6.0 will maintain the TCP/IP connection and restart the application. It is unlikely that these events will be noticed by the end user (or network load balancers/switches), as the client connections will be preserved.

Windows Server 2003 introduces a new kernel-mode HTTP driver called HTTP.sys, which is specifically tuned to increase Web server performance and scalability. In IIS 6.0, HTTP.sys listens for requests and routes them to the appropriate requests queue. Each request queue corresponds to one application pool. Because no application code runs in HTTP.sys, it cannot be affected by failures in user-mode code that normally affect the status of the Web service. If an application fails, HTTP.sys continues to accept and queue new requests on the appropriate queue until one of the following: the process has been restarted and begins to accept requests, there are no queues available, there is no space left on the queues, or the Web service itself has been shut down by the administrator. Once the WWW service notices the failed application, it starts a new worker process if there are outstanding requests still waiting to be serviced for the worker processs application pool. Thus, while there may be a temporary disruption in user-mode request processing, an end user does not experience the failure, because requests continue to be accepted and queued.

**WWW Service Administration and Monitoring Component**

The WWW Service Administration and Monitoring component makes up a core portion of the WWW service. Like HTTP.sys, no application code runs in the WWW Service Administration and Monitoring component. This component has two primary responsibilities: system configuration and worker process management.

At initialization time, the configuration manager portion of WWW service uses the in-memory configuration metabase[4] to initialize the HTTP.sys namespace routing table. Each entry in the routing table contains information that routes incoming URLs to the application pool that contains the application associated with the URL. These pre-registration steps inform HTTP.sys that there is an application pool that responds to requests in a particular part of the namespace, and that HTTP.sys can request that a worker process be started for the application pool when a request arrives. All pre-registrations are done before HTTP.sys begins to route requests to individual processes. As application pools and new applications are added, the Web service configures HTTP.sys to accept requests for the new URLs, sets up the new request queues for the new application pools, and indicates where the new URLs should be routed. Routing information can change dynamically without requiring a service restart.

In the worker process management role, the WWW Service Administration and Monitoring component is responsible for controlling the lifetime of the worker processes that process the requests. This includes determining when to start, recycle, or restart a worker process, if it is unable to process any more requests (becomes blocked). It is also responsible for monitoring the worker processes, and can detect when a worker process has terminated unexpectedly.

**Worker Process Isolation Mode**




**Figure 2: IIS 6.0 (Windows Server 2003) processing architecture**

IIS 6.0 introduces a new application isolation mode for managing the processing of Web sites and applications: worker process isolation mode (Figure 2). Worker process isolation mode runs all application code in an isolated environment (W3wp.exe) without incurring a performance penalty for that isolation. Requests are pulled directly from the kernel instead of having a user-mode process pull them from the kernel for the application, and then route accordingly to another user-mode process as was the case in IIS 5.0. First, HTTP.sys routes Web site and application requests to the correct application pool queue. Then, the worker process(es) (W3wp.exe) serving the application pool pull requests directly from the application queue in HTTP.sys. This model eliminates the unnecessary process hops encountered when sending a request to an out-of-process DLLHost.exe and back again (as was the case in IIS 4.0 and 5.0), and increases performance.

In worker process isolation mode, administrators can isolate different Web applications or Web sites into separate application pools. Each application pool is separated from other application pools by process boundaries. For example, a departmental server might have Web-HR in one application pool and Web-Financing in another; an Internet Service Provider (ISP) might have http://www.adventure-works.com in one application pool and http://www.blueyonderairlines.com in another. An application that is assigned to one application pool is not affected by other application pools, and cannot be assigned to another application pool while being serviced by the current application pool. Note that applications can be easily assigned to another application pool while the server is running.

It is important to note that, in IIS 6.0, there is no longer any notion of in-process applications. All necessary HTTP application run-time services, such as ISAPI extension support, are equally available in any application pool. This design prevents a malfunctioning Web site or application from disrupting the operation of other Web applications or the server itself. With IIS 6.0 it is now possible to unload in-process components without having to take down the entire Web service. The host worker process can be taken down temporarily without affecting other worker processes serving content. There is also a benefit from being able to leverage other operating system services available at the process level (for example CPU throttling), per application pool. Additionally, Windows Server 2003 has been re-architected to support many more concurrent processes than ever before.

Worker process isolation mode prevents one application or site from stopping another. In addition, separating applications or sites into separate application pools simplifies a number of management tasks, for example: taking a site/application online or offline (independent of all other site/applications running on the system), changing a component the application uses, debugging an application, monitoring counters for an application, and throttling resources used by an application.

The IIS 6.0 request processing architecture delivers very high levels of reliability without sacrificing performance.

- **Increased reliability.** IIS 6.0 worker process isolation mode prevents Web sites and applications from affecting each other or the server as a whole.

- **Fewer server restarts.** The user will likely never need to restart the server or shut down the entire WWW service, due to a failed application or common administration operations, such as upgrading content or components, or debugging Web applications.

- **Increased application availability.** IIS 6.0 supports auto restart of failed applications and periodic restart of leaky/malfunctioning applications, or applications with faulty code.

- **Increased scalability.** IIS 6.0 supports scaling to ISP scenarios, where there may be hundreds to thousands of sites on a server. IIS 6.0 also supports Web gardens, where a set of equivalent worker processes on a server each receive a share of the requests that are normally served by a single worker process.

- **Strong application platform support.** IIS 6.0 supports the application as the unit of administration. This includes making the application the unit of robustness by enabling application isolation, and also enabling resource throttling and scaling based on the application.

The end result is a Web platform that is more reliable and always available even if applications cause their worker processes to terminate. Worker process isolation mode contains the following features which improve its robustness without a performance impact.

**Clean Separation Between User Code and the Server**
All user code is handled by worker processes, which are completely isolated from the core Web server. This improves upon IIS 5.0, because an ISAPI can be, and often is, hosted in-process to the core Web server. If an ISAPI loaded in a worker process fails or causes an access violation, the only thing taken down is the worker process that hosts the ISAPI. Meanwhile, the WWW service creates a new worker process to replace the failed worker process. The other worker processes are unaffected.

**Multiple Application Pools**
With IIS 5.0, applications can be pooled together out-of-process, but only in one application pool, which is hosted by DLLHost.exe. When IIS 6.0 operates in worker process isolation mode, administrators can create up to 2,000 application pools, where each application pool can be configured separately.

**WebGardens**
Multiple worker processes can be configured to service requests for a given application pool. By default, each application pool has only one worker process. However, an application pool can be configured to have a set of N equivalent worker processes that share the workload. This configuration is known as a *Web garden* because it is similar in nature to a Web farm, the difference being that a Web garden exists within a single server. Requests are distributed by HTTP.sys among the set of worker processes assigned to the application pool. The distribution of requests is based on a round-robin scheme, where new connections with requests for the application pool are assigned to specific worker processes in that application pool. A benefit to Web gardens is that if one worker process slows down, such as when the script engine becomes unresponsive, there are other worker processes available to accept and process requests.

**Processor Affinity**
Worker processes can have an affinity to specific CPUs to take advantage of more frequent CPU cache (L1 or L2) hits. Processor affinity, when implemented, forces IIS 6.0 worker processes to run on specific microprocessors or CPUs and applies to all worker processes serving the Web sites and applications of an application pool. Processor affinity can also be used with Web gardens that run on multiprocessor computers where clusters of CPUs have been dedicated to specific application pools. Organizing workloads on a CPU means a more efficient usage of the CPU resources.

**Demand Start**
The new IIS architecture offers the benefit of demand start. Demand start means that the server is listening on an applications request queue, but will not physically initialize the application until a client has specifically connected to the application and requested service. At that point, a worker process will be initialized and the request will be fed into the processing engine for that application.

The WWW Service Administration and Monitoring component does on-demand process starting, and generally controls and monitors the lifecycle of worker processes.

**Idle Time-out**
The converse side of demand start is *idle time-out*. If an application has not had a request arrive for a configured amount of time, it can help to optimize a system if the resources that the application was using are temporarily freed up, until a client decides to use the application at some later point in time.

Idle time-out does not mean that the application is offline (i.e., the applications network endpoint is still active) just the physical process and run-time environment for the application are given back to the operating system.

**Kernel-Mode Caching**
Kernel-mode caching is available when using IIS 6.0, both in worker process isolation mode and in IIS 5.0 isolation mode (see below). As a single point of contact for all incoming (server-side) HTTP requests, HTTP.sys provides high-performance connectivity for HTTP server applications and provides overall connection management, bandwidth throttling, and Web server logging. IIS 6.0 has been built on top of HTTP.sys and has been specifically tuned to increase Web server throughput. In addition, under specific circumstances, HTTP.sys directly processes requests in the kernel. Both static and dynamic content from Web sites and applications can be cached in the HTTP.sys cache for high-performance responses.

IIS 6.0 has advanced heuristics built in to determine the cacheable hot-set of an application or set of sites. Just because an item is cacheable, does not necessarily mean it makes sense to add that item to an in-memory cache, since there is a cost to managing the item and the memory it consumes. Therefore, IIS 6.0 uses a new heuristic to determine which items should be

cached on the basis of the distribution of requests that a particular application receives. This means that the Web server makes better use of the resources available, while sustaining the performance on frequent requests, leading to scalability improvements.

IIS 6.0 also has heuristics built in to monitor the overall state of the server, and makes decisions to increase or reduce concurrency on that basis. The central idea here is to be efficient in using concurrency. For example, when executing processor-bound requests, starting concurrent work is not always the best approach.

# Programming Considerations (Web and Application Server Infrastructure - Performance and Scalability)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section is not designed to be the comprehensive performance programming section for Microsoft Windows Server programming environments. Rather, it simply identifies changes and highlights important considerations for building Web applications or Web services on Windows Server 2003.

**General Application Design Points**
This section contains some general points that are useful to be aware of when designing and developing and application.

**Design for Multiple Instances of an Application**
This is an important attribute for scaling applications. If an application can have multiple instances of itself running on a single server or in a farm scenario, it stands a better chance of being scaled up should the need arise, for example, to get additional throughput out of the application due to changing business needs.

Windows Server 2003 also offers the capability to easily run multiple instances of Web applications and Web services through a basic configuration parameter. The considerations when building and designing for multiple instances are:

- State management
- Centralized content or a strategy for replicating and synchronizing content

**Understand the Paths a Request Will Take Throughout the System**
It is important to understand, from a physical perspective, which processes and subsystems a specific request or message will propagate through on Windows. If you notice that a request or message is bouncing from process to process (or across a server) while its being processed, the application will need further scrutiny from a performance and scalability perspective. The best case scenario is to keep a unit of work within the same process that it started. This is often an impractical goal, but any reconciliation of process hops and minimization of network hops is good for overall performance and response time latency.

**New ISAPI APIs Related to Performance**
With the IIS 6.0 re-architecture, some of the IIS programming interfaces behave a little differently, and the core (base) native code API for IIS is augmented with new performance APIs and features.

**ExecuteURL**
Normally, an HTTP redirect is accomplished when a server responds to a clients request for a resource, with an HTTP 302 response containing the new location of the resource. The client application then does a new HTTP request to access the content (sometimes on the same server).

The ExecuteURL (**HSE_REQ_EXEC_URL**) **ServerSupportFunction** now allows an ISAPI extension to easily redirect a request to another URL. It answers growing demand by ISAPI extension developers to chain together requests of different types. From a performance perspective, it is very useful, because an application can achieve redirection of a request to a new piece of content without the additional network hop, and associated latency, of delegating a redirection to the client browser or rich client application. ExecuteURL also provides functionality to replace almost all ISAPI read raw data filters.

| **Important** |
| --- |
| If your applications require read raw data filters, your server must be running in IIS 5.0 isolation mode. |

Note that the most common scenario for developing read raw data filters is the need to examine or modify the request entity body before the target URL processes it. Currently, the only way to see the entity body of a request (if you are not the target URL) is through read raw data notifications. Unfortunately, writing an ISAPI filter to accomplish this goal can be exceedingly difficult, or even impossible in some configurations. ISAPI extensions, on the other hand, provide functionality for easy retrieval and manipulation of the entity body. ExecuteURL allows an ISAPI extension to process the request entity body and pass it to a child request, meeting the needs of nearly all read raw data filter developers.

Generally speaking, a great deal of care needs to taken when writing ISAPI filters, because the filters are totally synchronous and the code in an ISAPI filter executes on a core IIS thread. For operations that are going to involve blocking or calling off into

other computers (network latency), ISAPI filters can cause problems. Therefore, when ExecuteURL is used in combination with the wildcard ScriptMap API/feature, a developer has enough programming environment infrastructure to replace the majority of ISAPI filters with ISAPI extensions.

**Wildcard Application Map**

*What is an application map?*

When IIS receives a request from a client, the Web server looks at the extension of the file that is named in the request to determine which ISAPI or Common Gateway Interface (CGI) application handles that file. By using wildcard application maps, you can intercept every request before the requested page is sent to its mapped application. The effect is like having an application mapping that handles every file name extension. This is why the term wildcard is used to name this feature. Applications using wildcard application mappping can only be ISAPI applications, which is an advantage because they can have their own thread pool, work asynchronously, and have full access to entity bodies (important when processing HTTP POST as part of a Web Service).

With IIS 6.0, it is possible to specify one (or more) wildcard (*) application maps. As implied, these wildcard application maps get executed for every extension or request coming in to a site or virtual directory. Therefore, these wildcard application maps start to take on the behavior of an ISAPI filter; they get to see every request. Combined with the ExecuteURL function, they can perform some logic to determine what should be done with the request, and then optionally call in to other types of content.

The IIS 6.0 configuration file (**%SYSTEMROOT%\system32\inetsrv\metabase.xml**) contains a section entitled ScriptMaps. This section explains the application mapping of extensions (.aspx, .asp, .shtm, etc.) to physical components for request types. The ScriptMaps section under the /LM/W3SVC node of the metabase is the parent configuration, that is, all sites virtual directories inherit from it if they do not override a value at their level.

To configure a wildcard application map, please see Installing Wildcard Application Mappings in the IIS 6.0 Help.

**VectorSend**

ISAPI developers have only two options if a response is contained in multiple buffers. They can either call the ISAPI *WriteClient()* API multiple times, or they can assemble the response in one big buffer and send that response completely. The first approach is a performance bottleneck, because there is one kernel-mode transition per buffer and potential network latency. The second approach affects scalability, in that it puts more contention on the process heap and is wasteful in terms of memory and CPU cycles. For example, if you had to assemble a 64K response from multiple response parts, to allocate the target memory block, the heap would have to search for a 64K contiguous block for your program (possibly involving an extension of the program heap in the process) and you would have to expend the CPU cycles to copy 64K worth of data into the newly-allocated memory chunk. The VectorSend (**HSE_REQ_VECTOR_SEND**) **ServerSupportFunction** in IIS 6.0 provides a solution to this problem.

VectorSend allows developers to put together a list of buffers and file handles to send, in order, and then hand off to IIS 6.0 to compile the final response. HTTP.sys sends the buffers and/or file handles as one response by assembling and sending the response on the fly from the original data buffers. This frees the ISAPI from having to do any of this buffer construction or multiple WriteClient calls. This gives the additional benefit of higher performance and better scalability. If possible, existing applications should change from using WriteClient to using VectorSend.

**FinalSend**

The VectorSend **ServerSupportFunction** also includes an optional flag that aids the performance of a high throughput site, FinalSend (**HSE_IO_FINAL_SEND**). The **FinalSend** flag is one of the flags that can be passed into VectorSend to alter the behavior.

Specifying the FinalSend flag, tells IIS that it can clean up its request context, create the log file entry, etc. when it is completing the send in kernel mode. Without an ISAPI application specifying FinalSend, IIS deduces that the application is finished sending for the current request by retuning control to the HttpExtensionProc and waiting for it to exit. The downside here is that if IIS waits for the HttpExtensionProc procedure to return, it has to do an extra system call (transition from user mode to kernel mode and back) to facilitate the request cleanup and logging for the request. The use of FinalSend does all of this at the same time the actual response is sent to the client, saving one kernel transition per request. Although not a massive amount of CPU resources, this can be a big waste if you’re trying to sustain tens of thousands of requests per second.

**Caching Responses from Dynamic Content in the Kernel**

One of the key performance features of IIS 6.0 is that it will allow an ISAPI component to cache the dynamically-generated response to a request in the kernel, called the kernel-mode response cache. This feature has the potential to dramatically improve the performance of a server if the content being served lends itself to this form of caching.

*Why would I cache the output of a dynamic request?*

There are many scenarios where it is useful to dynamically generate a response, and allow the response to be a little stale. An online store scenario is a classic example. The product catalog is a relatively expensive page to create, because it needs to call in

to a database, perform a query that is likely to involve some form of database join, and return and render the results for the specific caller. In terms of CPU cycles, creating that page would have cost many millions of CPU cycles across multiple servers. Now, imagine you're a large e-commerce application on the Internet, getting hundreds or thousands of requests per second; what is the value in regenerating this same page a few hundred times per second, when the actual content is only going to change maybe once every ten minutes? The answer is there is little value in doing this for every request. A better strategy is to generate the page once, have the server efficiently cache it, and then give the response a staleness limit (how long the page can be served from the cache without being regenerated). Once its staleness limit expires, have the server call your code again, in the standard way, to regenerate the page and store the new response for the staleness limit.

*What performance gain do I get when putting responses in the kernel?*

To provide an idea of the difference, a simple ISAPI was written that looped 10,000 times to simulate some processing, it then responded to the client with approximately 1K of data. The ISAPIs performance was baselined on an 8P server[5]. Then, the code to cache the response in the kernel was added to the VectorSend() call. The version of the ISAPI that used the kernel cache returned a throughput of over eight times the baseline*846 percent* of the throughput in the non-cached case. Additionally, the latency of responses was dramatically better. When the ISAPI was executed each time, the average TTFB/TTLB[6] was 35.57 / 35.62 milliseconds; when the responses were served from the kernel, the TTFB/TTLB was 2.01 / 2.12 milliseconds.

Of course, results will vary according to the amount of processing done and the size of the response. The point is that leveraging the kernel cache can make a dramatic difference to the performance of an application.

*What types of content can be cached in the kernel?*

Static file requests, ISAPI components, and Microsoft ASP.NET pages can all cache responses in the kernel. A caveat to this is that the only responses that will be cached in the kernel are responses to HTTP GET requests.

**How to Cache ISAPI Responses in the Kernel**
To cache ISAPI responses in the kernel, you must use the VectorSend ServerSupportFunction to send your response. In preparation for sending your response, you must add the following response headers to go out with it:

- **Last-Modified** [7] a necessary field, it must be provided in GMT time format.

- **Expires** controls the amount of time an item is held in the cache. It is recommended that this value be calculated using an algorithm of current time on the server, plus the amount of time you want the cached copy to be served from the cache. This header should also be provided to IIS in GMT time format.

To cache the response, you call the VectorSend **ServerSupportFunction**, specifying at least the following flags:

**HSE_IO_FINAL_SEND | HSE_IO_CACHE_RESPONSE**

By implication, the send that is cached has to be the one and only send to a response.

The ASP.NET v1.1 release uses this facility in IIS 6.0 to host its output cache in the kernel. So, when using the output cache directive in ASP.NET, your responses are also being served from the kernel cache.

**Synchronous or Asynchronous ISAPI Sends**
This is a key area of architectural change for Windows Server 2003 compared to previous releases of Microsoft Server operating systems. In Windows Server 2003, the IIS 6.0 HTTP.sys interfaces with TCPIP.sys in the Windows kernel. Previous versions of IIS used the Winsock user mode library (a more general purpose TCP/IP sockets library).

Before Windows Server 2003, when an ISAPI application called a synchronous API to send a response to an end browser or application, the Winsock library would buffer the complete response in the kernel and free the ISAPI application to do more work. Meanwhile, if the caller getting the response back was on the end of a very slow link, the kernel memory allocated would be held for a relatively long time and the Windows kernel would grow its memory usage.

This functionality no longer exists in Windows Server 2003. If an ISAPI calls the WriteClient or VectorSend API synchronously (blocking call), IIS and HTTP.sys will only buffer the last 2K of a response on the callers behalf, the thread will wait until all but the last 2K bytes have been sent to that Web client (possibly halfway around the world, running over a 28.8 modem link). What eventually occurs is that the IIS thread pool reaches its maximum limit and stops growing, and the performance of the particular ISAPI is likely to be affected. To rectify this situation, it is recommended that existing ISAPI applications be changed to use the asynchronous versions of the WriteClient API. VectorSend should always be called asynchronously.

**COM+ Services in ASP**
In IIS 4.0 and 5.0, Active Server Page (ASP) applications are able to use COM+ services by way of configuring the applications

Web Application Manager (WAM) object in the COM+ configuration store to use transactions and other COM+ features. In IIS 6.0, the IIS and COM+ teams have separated the COM+ services from physical components. This allows ASP applications to use COM+ services directly inline. For example, a component can enter a transaction context inline in its code, and then exit the transaction context while still processing a method. In addition to those services available in COM+ on Windows 2000, a few new services have been added and are supported in ASP.

**Apartment Model Selection**
ASP, through COM+, allows developers to determine which threading model to use when executing the pages in an application. By default, ASP uses the single-threaded apartment (STA) model. However, if the application uses poolable objects, it can be run in the multi-threaded apartment (MTA) model.

This is an important scalability gain for many existing ASP applications, because running the ASP objects in the multi-threaded apartment means fewer threads in the system, fewer context switches, and better resource usage. If your ASP application is only using COM+ objects marked as Both or Free threaded, you can safely turn on the multi-threaded apartment in IIS 6.0 for these applications.

To turn on the multi-threaded apartment for a particular site or application in IIS 6.0, navigate to the applications node in the MetaBase.xml file, and add the value: AspExecuteInMTA=1.

**High Latency ASP Pages**
The ASP programming model is fundamentally synchronous; a request comes in to an ASP script and the ASP script keeps processing until that request is completed and the response is sent. The simple nature of the programming model is such that a developer cannot start an operation that is going to take a long time, and expect the ASP page to exit, and when the operation completes, to restart from a given point.

Consequently, if an ASP page includes a COM+ object that has to perform an operation that is going to take a long time, what will likely occur is that the thread executing the page will block and idle until the I/O operation is complete. If you have a set of ASP scripts that all have high latency operations, you will likely need a large number of threads in the system to ensure good response times.

With any large multiprocessor system, it is better to minimize any surplus threads you have; more threads = more context switches = less scalability. In Windows Server 2003, COM+ has changed to take a gradual approach to starting up new threads. It looks at system parameters before making the decision to start more concurrent work. If there is an ASP-based site or application (running in the COM+ STA) that just needs to have a high number of threads started when it initializes, it is possible to make COM+ revert to the previous behavior with the following registry setting:

**HKEY_LOCAL_MACHINE\Software\Microsoft\COM3\STAThreadPool\CPUMetricEnabled = 0 (REG_DWORD)**

As per the COM+ parameter above, you may want to also increase the number of ASP thread pool threads when you have an ASP page that experiences high latency, but is running in the COM+ MTA, or uses ASP intrinsic objects that will take a long time to execute. The IIS metabase parameter is *AspProcessorThreadMax* and is set to 25 by default. This 25 number can be a little misleading; it means a maximum of 25 ASP threads per processor, so, if you're running on a 2P system, you have a maximum of 50 threads by default. *AspProcessorThreadMax* can be changed by running the following command:

for a specific site: **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/***NumericSiteID***/AspProcessorThreadMax 5**

or, for all sites: **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AspProcessorThreadMax 5**

**COM+ General**
With the new IIS 6.0 architecture, it is important to question some of the existing guidelines where COM+ is concerned. A major consideration is that, before Windows Server 2003, COM+ application components were configured (by default) to run out-of-processes from the caller. The default for COM+ applications is for them to run as Server Applications, executing in a DLLHost.exe process called into from the object instantiate, or over DCOM.

The performance downside of doing this for every method call, is that there are extra threads running on the system, and every call to a method must be marshaled across process boundaries. This is not noticeable on a small implementation with low request/transaction rates, but on a high volume, large multiprocessor, this kind of overhead can greatly decrease the overall scalability of the system.

Therefore, on Windows Server 2003, it is best to change the default configuration for a COM+ Server Application to Library Application[8] to aid scalability of the calling per use of that application.

**ASP.NET Environment**
There are already many existing documents on http://msdn.microsoft.com and http://www.asp.net/Default.aspx?tabindex=0&tabid=1 which discuss coding tips and techniques for building incredibly fast applications using ASP.NET. Most of

this information applies to running the ASP.NET v1.1 release on Windows Server 2003 as well. This document will emphasize some of the points that have particular relevance to Windows Server 2003.

**Kernel-Mode Response Cache**
When the Windows .NET Framework v1.1 is running on Windows Server 2003, ASP.NET will detect the enhanced capabilities of IIS 6.0 and will leverage some of the performance features available in the Windows Server 2003 platform. A key feature being leveraged is the kernel-mode response cache. ASP.NET uses the kernel-mode response cache to store its output cache responses. The table below shows an example of an ASP.NET page that generates a 4K response. One set of figures is when the page is being retrieved from the user-mode output cache in ASP.NET, and the other is when the page is being retrieved from the kernel-mode cache.

**Table 1 The difference between user-mode and kernel-mode output caches for ASP.NET pages.** [9]

| | ASP.NET User-Mode Output Cache | Kernel-Mode Output Cache |
|---|---|---|
| Requests / Sec | 1,394 | 15,416 |
| TTFB / TTLB (msec) | 70.82 / 70.97 | 3.39 / 4.02 |
| User Mode CPU % | 76.57% | 0.78% |
| Kernel Mode CPU % | 22.69% | 99.22% |
| System Calls / Sec | 20,110 | 2,101 |
| Network Util (KB / Sec) | 6,153 | 68,326 |
| Context Switches / Sec | 2,621 | 6,261 |

As you can see, the difference is substantial when the content is being served directly from the kernel-mode response cache. When the requests are going to the user-mode static-file cache, every request transitions the processor from kernel to user mode and back, and changes process page pointers and other operating system context data. If the workload stays in the kernel, a lot of system overhead is avoided.

Additionally, with the ASP.NET kernel-mode response cache, there are no changes required from a developer's perspective. If the existing .aspx page already has the following directive:

**<@ OutputCache Duration="time in seconds VaryByParam=none %>** [10]

the page output is automatically served from the kernel-mode response cache.

On a high volume Web site, setting the output cache duration to as low as one second can also make an important difference to the overall throughput of a server.

**Partial Page Caching**
Another major consideration for a Web application is which sections of dynamically generated pages can be statically cached and which parts need to be regenerated. It can be a big performance boost if expensive-to-create portions of pages can be cached, and only the parts of a page that really need to be regenerated are dynamically generated per page access.

**Exception Handling**
Exception handling should be just that, code that is executed in the exceptional case, not code that is executed all of the time. Many programmers have developed a style where they use an exception as a convenient way to execute cleanup logic on the basis of a standard condition.

A few statistics might help to drive the point of how expensive exceptions can be: on a processor level, the cost of throwing an exception in the CLR[11] can be up to 1.8 million CPU cycles[12]. Again, 1.8 million CPU cycles gone, when one comparison operation would have sufficed. Worse, throwing an exception invariably acquires a system lock for page frame manipulation. The section entitled Software Locks (Resource Contention) explains why this is a bad thing to do. The result is that an application that throws and catches a lot of exceptions will not scale on large multiprocessor servers, in addition to wasting CPU cycles.

To summarize, throwing exceptions should be done very sparingly and certainly not on your common code paths, if you want to have high performing, scalable code.

**COM+ Interop**
Many companies have an existing inventory of applications or components that have been built using COM+ over the years.

When writing new ASP.NET applications, chances are that many of the existing COM+ components will be integrated into these ASP.NET applications (why fix something that isn't broken?). Therefore, this section contains a set of guidelines regarding the use of COM+ interop in ASP.NET.

**Infrequent, Functionally-Rich Calls**

The key to efficiency when using COM+ interop with ASP.NET applications is to minimize the number of times you transition from one environment to the other. An application needs to minimize the number of housekeeping CPU cycles expended on transitioning between the world of managed code and unmanaged code. Therefore, it is better to try and have an interface that has substantial, but less frequent calls, rather than frequent, chatty calls.

**Transactional Attribute on COM+**

Two-phase commit is one of the key features of COM+. Many projects mark all of the objects involved in an application with the **Required** or **Requires New** transaction attribute. This can introduce a substantial processing and network latency cost to the application.

The reason for this is that when a new transaction is started, or when a component enlists itself in a transaction, there is a communication with the Microsoft Distributed Transaction Coordinator Service (MSDTC). If the MSDTC is running on the local server, this will involve extra context switches; if MSDTC is running on a remote server, it will also include additional network latency. Use the transaction attributes *only* when necessary. If performance is a critical concern, consider re-factoring your components appropriately.

**Single-Threaded Apartment**

If your ASP.NET solution uses COM+ components that were built using Microsoft Visual Basic 6.0 or COM+ components marked as apartment threaded, it is recommended that you mark your ASP.NET page with the **AspCompat** directive. At the top of the .aspx page, add the following to the directive line:

```
<%@ Page AspCompat="true" Language="vb" %>
```

The reason for this is that if your .aspx page detects a COM+ object, in addition to interop, it also has to run your component in its own apartment, which means additional threads and data structure marshalling,essentially more operating system housekeeping. You can avoid some of this overhead with the **AspCompat** directive.

# General Performance Considerations (Web and Application Server Infrastructure - Performance and Scalability)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section contains general performance considerations related to processing HTTP.

**Networking**
Networking is a key consideration for a high volume server and is also becoming more of a factor with recent hardware. Larger numbers of processors per server and the increase in speed of each CPU, means that a network I/O subsystem can easily become a bottleneck for a large, powerful Web or application server.

Consequently, the network I/O subsystem of a large server requires design, planning, and monitoring. If a network card becomes a bottleneck, faster processors, more processors, or more RAM will not aid your server in achieving higher throughput.

*How do you know that a network card is a bottleneck on a system?*

The most obvious sign is by looking in the Windows Task Manager under the Networking tab. Select the network interface you would like to view and check its Network Utilization statistic. If it is close to 100 percent, the card is likely a bottleneck in the performance of your application.

Next, check the CPU utilization and context switch per second parameters. If the CPU utilization is low and context/switches per second are low[13], the server is likely to be experiencing a networking bottleneck.

*What kind of maximum throughput can I get out of my network card?*

| Network Card Speed | Megabytes / Sec |
|---|---|
| 100 megabit (100Mbps) | ~11 |
| 1 gigabit (1000Mbps) | ~60 |

These general, rule-of-thumb numbers are derived through IIS team experimentation across a wide variety of hardware and differing network card combinations. Two factors that can also affect the overall throughput of a network card is the speed of the I/O bus and the number of cards put into that I/O bus. For example, on a standard symmetric multiprocessing (SMP) 8-processor server with multiple cards, gigabit cards were plugged into a 33-MHz root bus. Results showed that the gigabit cards maxed out at around 30 MB/sec, where, on a 66-MHz root bus, the gigabit network cards could achieve 50-60 MB/sec.

The bottom line is that if network card throughput is vital, bus speeds, network cards, and testing of the overall system are also vital to understand the possible throughput characteristics of a system.

**Windows, Networking, and Large Multiprocessor Servers**
As discussed in the Networking section, it is important to plan for the amount of data throughput you expect from your server. When running on a large multiprocessor server, you also need to understand a little about the Windows system internals, in order to maximize your throughput.

From a design standpoint, the Windows networking stack performs a large part of its processing of a network cards I/O on a single processor. Therefore, if you have an intensive networking load (many Web scenarios fit into this category, with thousands of requests per second, and very little CPU time spent per request), it is possible to see uneven processor utilization. Worse, it is possible to make one CPU the bottleneck for the whole server, because it is at 100 percent utilization, while the other processors are idly waiting for it to hand off the network I/O so they can process the next request.

It can be the case that the CPU handling the network interrupts is at 100 percent, but the card is not physically saturated. The variables here are the network card manufacturers driver efficiency and quality, the speed of the processor in the server, offloading (or lack of offloading) support in the network card, and the physical capacity of the network card itself. You can rectify this situation by adding an additional network card and splitting the load coming into the server across the two network cards.

There are additional considerations for networking on large multiprocessor servers covered in the Multiprocessor section.

**HTTP Compression**

HTTP Compression is not often considered as a server throughput enhancing tool, but if your sites use large amounts of bandwidth, or if you would like to more effectively use bandwidth, consider enabling HTTP compression. HTTP compression provides faster transmission time between compression-enabled browsers and IIS. You can compress either static files only, or both static files and dynamic application responses. If your network bandwidth is restricted, HTTP compression can be beneficial, at least for static files, unless your processor utilization is already extremely high.

HTTP compression can have a dramatic improvement on the latency of responses, while improving the throughput capacity of the server.

As shown in the Networking section above, it is very easy to make a network card a bottleneck on a server[14]. Compressing the responses will mean fewer bytes to send for the same response. Network cards do not discriminate on the actual bytes, therefore, to serve the same load, the network cards will be able to push through many more compressed requests per second before they become a bottleneck.

For more information about enabling HTTP Compression. see Utilizing HTTP Compression in the IIS 6.0 Help.

**Compressing Static Content**

IIS 6.0 will store the compressed representation of a static file securely on persistent storage and will serve the compressed version for clients that request the static file and ask for the content in compressed form[15]. This is highly efficient in that, for the lifetime of the content, the compressed version of the content will be served to clients that ask for it without the CPU having to compress at each request.

A further optimization for static content is that if the heuristic for caching the item in the kernel is met, the compressed version of the item is stored in the kernel. This has the added advantage that less memory is used to cache the item and it is served in the most efficient way possible.

As an example of the gains that compression can give you, the following test was devised. The hardware used was a 1 processor, 3-GHz HT[16] server with one 100-megabit network card. The content was a 10K static .xml file.

| | **Non-Compressed** | **Compressed** |
|---|---|---|
| Requests / Sec | 1,170 | 2,493 |
| CPU % | 12.2% | 19.7% |
| Network Card % | 99% | 99% |

In this case, the throughput increased by 113 percent. The additional CPU used in the compressed case was not spent on compression (the compression for the file was only done once). Rather, the additional CPU was expended in processing the additional 1,323 requests per second over the non-compressed case.

**Controlling Static File Types to Be Compressed**

When compression is enabled on IIS 6.0, by default, only files with extension .htm, .html, and .txt are compressed. This is quite a limited list, and there are other static file types that can benefit dramatically from being compressed, such as .xml, .css, and .js. You should also ensure that these files are also in the compression list for the IIS compression filter. To do this, execute the following commands at a command prompt:

**cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/Filters/Compression/deflate/HcFileExtensions htm html txt css js xml**

and

**cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/Filters/Compression/gzip/HcFileExtensions htm html txt css js xml**

If your implementation has its own static file types, feel free to add their extensions to this list as well.

| **Note** |
|---|
| Compressed format images (.gif) do not tend to gain much out of being further compressed by IIS, thus the exclusion from this list. If the extensions are excluded from this list, however, they will not be cached in the kernel. This is an area where the extension inclusion should be considered, dependant on the Web sites content. |




**Figure 3: IIS 6.0 Compression Dialog**

**Compressing Dynamic Content**

You can also enable compression for the output of dynamic content. This is enabled through the IIS configuration dialog (see above) by enabling the Compress Application Files option. The downside in using compression for dynamic content is that IIS 6.0 must run its compression routines on every response. This can have a cost in terms of CPU consumption.

From a business justification perspective, you would enable compression for dynamic content if latency for end clients was an issue, and the extra CPU utilization was acceptable to operational guidelines for a server. Implementations with large numbers of slow client connections or mobile device clients where the wireless networks typically have high propagation delays, are some scenarios where dynamic compression could be used.

The below test data provides an indication of the impact of switching on dynamic content compression. In the test, an ISAPI extension was written whose response consisted of a 10K XML-formatted response. The test with compression on was purposely gated at a max of 50 percent CPU utilization to simulate an acceptable level of CPU use on a server. The test without compression was done to see the difference in CPU utilization at the same request rate. Testing was done on a 2-GHz Intel HT computer, 512 MB RAM, one 100-megabit network card.

| | **Non-Compressed** | **Compressed** |
|---|---|---|
| Requests / Sec | 1,051 | 1,074 |
| CPU % | 23% | 50% |
| Network Card | 99% | 31% |

At roughly the same request rate, dynamic content compression (of an XML-formatted response) consumes approximately twice the CPU without compression on. In addition to less data being sent between client and server, the network card is utilized more efficiently.

As you can see, compression for dynamic content is costly and should be carefully evaluated before being enabled on a production server.

**Compression Encoder Selection**

Compression encoder strength refers to which technique Windows uses to compress the data stream. You have the option of using an incredibly strong encoder, which will typically result in a compressed stream that is smaller than what you get when the standard encoder is used. But using a stronger encoder generally means that the amount of CPU time used is also greater. The administrator should consider the overall goals of the Web or application server and the type of content being compressed.

The strength range for the encoder is between 0 and 10, with the following breakdown:

- 0-3 fast encoder

- 3-9 standard encoder
- 10 optimal encoder

By default, static content is compressed at level 10 (most optimal) and dynamic content is compressed at level 0 (1 pass), to be most efficient at CPU utilization, on the basis of how frequently these routines are called. These parameters can be modified in the IIS metabase. To change the static compression encoder strength, open the MetaBase.xml file with a standard text editor and change all instances of the *HcOnDemandCompLevel* parameter to the desired value. Similarly, to change the compression encoder strength on all dynamic content, search for all instances of the *HcDynamicCompressionLevel* parameter and change it to the desired value.

**Compression at a Site Level Rather Than Global**
IIS 6.0 also allows the enabling of compression at a site level rather than at a global level for the whole server. The properties to manipulate at a per site level are:

DoStaticCompression; and

DoDynamicCompression;

From more information about enabling compression on a per site basis, see DoStaticCompression and DoDynamicCompression in the IIS 6.0 Help.

**TCP/IP Configuration and Connections**
Windows Server 2003 can receive Web traffic via different physical TCP/IP addresses, different TCP/IP ports, or over different host headers (DNS names). A host header can be used to differentiate one site (or application) from another.

If working in a server consolidation scenario or a dense hosting scenario, the favored approach from a system resource usage perspective is to use the host header approach, rather than having a separate TCP/IP address for each site or application. Although possible to go higher, you should limit the number of TCP/IP addresses per system to around 2,000, where it is possible to support tens of thousands of host headers per system.

**Caching**
Caching is an important aspect of a Web platform. Caching can greatly improve the throughput on a server and take the pressure off back-end systems by avoiding the work of totally recreating the response on a per request basis. Figure 3 represents the caches available in Windows Server 2003 for Web sites and applications.




**Figure 4: Windows application server caches (kernel and user-mode)**

**HTTP.sys Kernel-Mode Response Cache**
HTTP.sys has the ability to cache responses to HTTP GET requests[17] in the kernel-mode response cache. The kernel-mode response cache is PAE-aware,[18] which means that it can contain a maximum 64 GB of information on an x86 system. For 64-bit systems, HTTP.sys has access to a full 64-bit address space and, at time of writing, Windows Server 2003 64-bit edition will support systems with up to 512 GB of RAM.

If the kernel-mode response cache is enabled (it is enabled by default), HTTP.sys will always consult it for a received request to see if there is a cached response available. If there is a cached response available, it will be served directly, without transitioning the processor to user mode.

HTTP.sys employs a scavenger model for the kernel-mode response cache. The scavenger runs periodically and clears out cache items that have not been recently accessed. If an item is cleared from the kernel-mode response cache, HTTP.sys sends a

subsequent request for the particular item to the user mode program that originated the initial cached response.

Most configuration of HTTP.sys is done through IIS Manager and the IIS metabase. However, the following list of items can be added to the registry to override the behavior of the kernel-mode response cache. The registry location for these values is:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Http\Parameters** [19]

| Parameter | Range | Comment |
|---|---|---|
| *UriEnableCache* | (DWORD) 1=on; 0=off | If non-zero, the kernel-mode response cache is enabled. |
| *UriMaxCacheUriCount* | (DWORD) 0 - 0xFFFFFFFF | Determines how many responses may be cached. If zero, there is no limit. |
| *UriMaxCacheMegabyteCount* | (DWORD) 0 - 0xFFFFFFFF | Determines how much memory may be used by the kernel-mode response cache. If zero, a heuristic will be used. |
| *UriScavengerPeriod* | (DWORD) 10 - 0xFFFFFFFF seconds | Determines the frequency of the cache scavenger. The scavenger runs once every UriScavengerPeriod seconds. Any response or fragment that has not been accessed in at least UriScavengerPeriod seconds will be flushed. |
| *UriMaxUriBytes* | (DWORD) 4096(4k) 16777216(16Mb) bytes | This parameter is used when HTTP.sys tries to cache a response in the Uri cache, hence the name UriBytes. If an item size exceeds this value, it will not be cached. |

**User-Mode Static-File Cache**

In addition to the kernel-mode response cache, IIS 6.0 employs a user-mode static-file cache. The user-mode static-file cache contains all of the policy decisions about what to cache and when to cache. It has a superset of functionality and will act as a backup for the HTTP.sys cache in situations where HTTP.sys is unable to cache an item. For example, when an administrator enables authentication, content in the virtual directory cannot be cached in the kernel, therefore, the user-mode static-file cache is used.

From a performance perspective, the user-mode static-file cache cannot achieve the performance of the kernel-mode response cache, because the code path length in the server is greater and the processor needs to transition modes per request. The below shows data from a simple test performed on a 2-processor, 1-Ghz server with 1 gigabit network card. The test is for the same 2K static file content item. One test is where the item is served from the user-mode static-file cache; in the second case, the file is served from the kernel-mode response cache.

| | User-Mode Static-File Cache 20 | Kernel-Mode Response Cache 21 | Difference |
|---|---|---|---|
| Request / Sec | 8321 | 17689 | + 113% |
| CPU % | 100.00 | 92.83 | - 7.17 |
| KB / Sec | 120906 | 257159 | + 113% |

The user-mode static-file cache is capable of caching everything that is cached in the kernel-mode response cache, but it has additional caching capabilities that are useful to other scenarios. For example, static files that are cacheable and are requested by clients that have been authenticated, are served from the user-mode static-file cache. It would be inaccurate for IIS to serve a static file to a user without doing an authentication check on the content to ensure the user has rights to the content. All access checks are done in user mode.

**Caching Policy**

In direct opposition to what seems logical, it turns out that caching everything that can be cached is not as helpful as caching everything according to need. Therefore, you should only cache what is actually getting lots of requests rather than caching everything. The reason for this is that populating a cache can be quite an expensive operation, and the management of a cache with thousands or tens of thousands of items in it can also be a CPU-intensive or wasteful operation.

IIS 6.0 tracks the hit distribution of content, and if the content proves itself to be popular and worthy of being cached, IIS 6.0 will put the content in a cache. When IIS 6.0 makes this decision to cache content, it first populates the user-mode static-file cache with the content, and then optionally, depending on the circumstance, puts the item in the kernel-mode response cache.

When an item is placed in the kernel-mode response cache, all the hits are served directly by the kernel. What will likely

happen then, is the user-mode static-file cache content will time out (after 30 seconds) and be dropped from the user-mode static-file cache , as the item in the kernel takes the requests.

**Hot Content Determination and Tuning**
IIS 6.0 determines content is hot is by counting the number of times it is requested. If a cacheable static file is requested twice within a time period (referred to as the ActivityPeriod, by default 10 seconds), IIS 6.0 caches the content. As mentioned above, if the item in the user-mode static-file cache does not receive another hit within 30 seconds (default), the item will be dropped from the user-mode static-file cache.

**User-Mode Static-File Cache Parameters**
The modification of these parameters should be done with care, because an incorrect setting can cause a drastic effect on the performance of the server.

The registry location for these values is:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Inetinfo\Parameters**

| Parameter | Range | Default | Comment |
|---|---|---|---|
| *MaxCachedFileSize* | (DWORD) n/a | 256 KB | If you are running large, dedicated Web servers, you may want to add this value to the registry to increase the file size that the cache can hold. |
| *MemCacheSize* | (DWORD) 0 – 2500 MB | Approximately one-half of available physical memory | The default value is adjusted dynamically every 60 seconds. |
| *ObjectCacheTTL* | (DWORD) 0 – 4,294,967,295 (unlimited) | 30 seconds | Setting this value to 0xFFFFFFFF allows the cached object to remain in the cache until it is overwritten. This would be appropriate if your server has enough system memory and is handling data that is relatively static. |
| *ActivityPeriod* | (DWORD) 0 – 4,294,967,295 (unlimited) | 10 seconds | If this value is set to 0, IIS always caches files. |

**Dynamic Response Cache**
Web applications built using ASP.NET and ISAPI can cache responses and have subsequent requests for content served directly from the cache (rather than regenerating the content). This can be a huge performance gain for servers, because the caching can save significant CPU cycles and significantly improve overall response time.

By default, the kernel-mode cache can store dynamic content responses. Applications must be written (or configured by developers) to take advantage of this cache. For information on the ASP.NET output cache see the section entitled Kernel-Mode Response Cache, and for information on how to get ISAPI applications to leverage the kernel-mode response cache, see the section entitled Caching Responses from Dynamic Content in the Kernel.

**Logging**
From an implementation perspective, logging has also changed for IIS 6.0. The log formats have not changed, nor the behavior in general, but the Internet log files are now created by HTTP.sys in kernel mode, rather than by user-mode IIS components. HTTP.sys knows what type of log file to create and which fields are necessary, based on information passed down from the IIS components in user mode. The configuration repository for logging stays consistent from previous releases of IIS; log file parameters are stored in the IIS metabase.

**New Format in IIS 6.0 Centralized Binary Logging**
In addition to the existing log file formats of W3C and IIS NCSA, IIS 6.0 introduces a new format, centralized binary logging. Centralized means that IIS will override all of the existing log file settings on all sites on the server, and generate one log file for all of the sites or applications on that server[22]. The log file will be generated in a binary format which is optimized to eliminate variable length strings where possible. The centralized binary logging file also has its fields aligned on the right boundaries to maximize the I/O performance.

| **Note** |
|---|
| The centralized binary log file does not log all of the fields available in the W3C format. This cannot be changed or overridden, so if you have a need for a specific field that is not included, the binary log file is not an option for you. For more information about enabling centralized binary logging, see Centralized Binary Logging in the IIS 6.0 Help. |

*Why is centralized binary logging interesting?*

There are a few reasons why you would consider using centralized binary logging. For servers that have thousands of sites, having a log file per site can be a significant bottleneck because disk head access latency can affect the overall throughput of the server. (The server struggles to write blocks widely distributed over a physical medium). With centralized binary logging, disk head access time is improved; the single log file for the server can be placed on a separate device and the disk head of that device gets to write sequential blocks to the device.

The centralized binary logging format is optimized, and offers performance gains over the standard text-based logging format, while being smaller in size.

When IIS is logging in standard text mode, a 64K log buffer per log file is allocated for performance reasons[23]. If you have 20,000 sites on a server, by default, you might be allocating over 1 GB of RAM for the log file buffers. However, IIS 6.0 does allow for the log file buffer size to be decreased, so that if you have 20,000 sites on a server, you might peak at a little over 200 MB for the log file buffers[24]. When using centralized binary logging, you have exactly one 64K log buffer for the whole server, regardless of how many sites you put on that server. Therefore, there can be a significant memory benefit when using centralized binary logging.

| **Note** |
| --- |
| The term might is used, because the memory allocation numbers quoted above can happen if the server receives hits across all of the sites in a short period of time. Log buffers are timed out depending on how frequently log records are written to the buffer. |

Having a single log file can simplify the log file management on a server. You manage one log file rather than tens, hundreds, or thousands of log files per server.

Microsoft has released Log Parser v2.0 that can read the binary log file and generate W3C log files for the individual sites contained in the binary log file. With Log Parser, the binary log file can be copied to another computer for processing. Therefore, there is no need to invest the CPU cycles to do the formatting of the log file at run time, those CPU cycles can be better used on a different, non-production computer for log file interrogation and reporting. For more information, see the LogParser section of this document.

Logging is an area that has differing functional importance, depending on the scenario. For an ISP scenario, logging is a key mechanism in determining how many hits of usage a site has had for billing and capacity analysis purposes. For an average commerce application, the logs would not be analyzed constantly, but rather stored to help track down issues or do site analysis if necessary. On the basis of differing usage scenarios, operational staff have some options to consider when determining how logging will work for their particular implementation, and have some flexibility on how they can affect its overall performance.

The first thing to remember is that the best performance option is to disable logging; this relieves processing cycles dedicated to generating log files, and saves physical disk space and extra I/O associated with persisting the log files. This is done through the Enable Logging option on the Web Site Tab on the web site properties dialog.

Note, however, for many implementations this is not an option.

**Performance**

From a performance perspective, the most costly part of logging is the long, variable-length strings. These strings lead to inefficiencies in log buffer use and more physical bytes to write to the disk. The main perpetrators in the W3C log, are the User-Agent and UserName fields, and a few other fields.

If you are running a server that has a high throughput rate (peaks of thousands of requests / sec), you can get a throughput boost by just taking out the variable-length text strings.

**Straight W3C Optimization over W3C Default**

The data in the table below is from tests involving the execution of an ASP.NET page on a 900-MHz, 8-processor server. One test was done with the default logging settings; the second test was done with the User-Agent field of the W3C format disabled (Optimized W3C).

| | **Default W3C Logging** | **Optimized W3C** |
| --- | --- | --- |
| Requests / Sec | 5716 | 6115 |
| CPU % | 98.49 | 98.93 |

| | | |
|---|---|---|
| Log file Size (MB) | 263.63 | 165.02 |

The difference in taking out this one field is quite noticeable, throughput increased by approximately 7 percent and the resulting log file was 37 percent smaller than the original. Note that the User-Agent field value in this case was the string that represented Microsoft Internet Explorer 6.0 (approximately 70 characters).

In this same scenario, the centralized binary log file performs very similarly to the optimized W3C~7 percent more throughput than the default, and a smaller resultant log file. The reason that the binary log file is similar to the optimized W3C log file, in this case, is that 100 percent of the requests in the test actually executed the page (in user mode). This means that the binary log file has to write the Uniform Resource Identifier (URI) into every line item in the binary log file. When a 4K file was cached and served from the kernel, the binary log file increased server throughput by 2 percent over the optimized W3C format. The log file size was 7 percent smaller than the optimized W3C log file.

**Parallelism in Logging**

You should also consider, when configuring a log file subsystem on a server, logs for busy sites should be separated onto separate devices with separate controllers. The goal is to split out the I/O for busy sites so that log buffer writes for one busy site are not queuing behind log buffer writes for a different site, with both being written to the same device. If this is the case, the overall logging process has a much higher latency.

To determine the busy sites on a server, there are a few parameters you can monitor. In System Monitor, select the **Web Service** object. In **select counters from list**, select the **Total Method Requests** item. Then, in **select instances from list**, select the sites you want to inquire on.

The **Total Method Requests** item displays how many requests the site has had since IIS was started. From here you can quickly deduce which sites are the busy sites and start planning the log file subsystem accordingly.

**Heap Fragmentation**

One phenomenon that causes significant performance and scalability issues in ASP/COM+ and ISAPI applications is a condition called heap fragmentation. Heap fragmentation is frequently misdiagnosed as a memory leak when the system seems to be running out of memory resources and long-running applications slow down. Heap fragmentation is not a memory leaka developer could spend weeks trying to hunt down memory leaks in the application and not find any, and yet this condition would still exist. Heap fragmentation occurs when a memory heap (which is used for memory allocations) is expanded and continues to grow until it hits a system addressing barrier. What usually causes this heap fragmentation is memory requests of wildly differing sizes being made to the same memory heap.

When a memory heap has a request for a small memory block (say 30 bytes) followed by another request for a large memory block (say 100 KB), it will satisfy both and the heap will have a 30-byte allocation followed by a 100-KB allocation next to it. Now imagine the 30 bytes were de-allocated and a new request for 2,048 bytes came in. The system cannot allocate the 2,048 bytes in the space where the 30 bytes were; it must go to the end of the heap, possibly extend its size[25], and satisfy the allocation there.

Imagine a scenario where over a long period of time, Windows has to keep growing the heap. For each allocation request, the heap manager has a lot of searching to do, and it will eventually get to a point where it cannot allocate any more addresses (2-GB address limit). The heap manager will then fail the heap extension and pass back the memory allocation error to the application.

The way to rectify this situation is to form buckets for allocations across a number of heaps and only allocate the same size for each heap. This will resolve the heap fragmentation issue. Another alternative is to use Low Fragmentation Heap,[26] Microsofts implementation of the above, to rectify heap fragmentation issues.

**What Can I do as an Administrator?**

The discussion so far has focused on the cause and the remedy, unfortunately, there isn't much to address the symptoms and keep an application going when it is slowing down and exhausting memory resources. With IIS 6.0, an administrator can keep this condition at bay without having to baby-sit the server or take it offline.

To detect an application that is fragmenting the heap, you can use System Monitor. The signature of an application that is fragmenting the heap is the worker process in which it is executing increases its use of virtual memory over time, while its private bytes stay static. You can detect this issue by logging the *Process(w3wp#)->Virtual Bytes and Process(w3wp#)->Private Bytes* for the specific worker process instance you are interested in.

Once the issue is detected, keeping heap fragmentation at bay until developers can rectify the situation, means setting the recycle parameters on the application pool to recycle once the application exceeds acceptable virtual memory use limits.

**Note**

Avoid setting the virtual memory limit recycle setting on ASP.NET applications; the ASP.NET runtime takes care of heap fragmentation intrinsically.

The Recycle tab of the Application Pool dialog (Figure 4) enables recycling on Virtual memory use and Real memory use (private bytes).




**Figure 5: Parameter to control applications that display heap fragmentation**

**Multiprocessor Servers**

This section contains considerations and areas of focus for situations when Windows Server 2003 is running on large multiprocessor servers.

**Localized Loads**

When running on a large multiprocessor server[27], it is important to understand the physical hardware bus architecture and the costs of having a workload distributed across processors. Many large multiprocessor computers are built using a building block approach. A 32-processor computer can have a four-CPU base unit, with an interconnection architecture that links eight of these four-CPU building blocks to an I/O back plain, a common RAM area, etc. Note, however, that the interconnections between the building block processing unit (and the caches around the building block unit), impose a cost when breached.

In the above diagram, a workload that processes for a while on processor 0, then context-switches to processor 23, incurs a much heftier tax than if the workload context switched to processor 1. By going to processor 23, the hardware has to facilitate the transfer and cache-coherence of any data the workload was referencing in the caches surrounding processors 0-3, to the caches surrounding processors 20-23 before it can be re-scheduled to run. Often, depending on the specific hardware architecture, the process of synchronizing the processor caches can also stall other processors in the two building block units to ensure that the caches are not randomly polluted before the transfer operation completes. Therefore, the transfer also has an impact on workloads outside the initial workload.

**Application Silos**

Windows Server 2003 allows an administrator to configure a server to minimize the impact of the above processing locality issue when running Web services or Web applications. This can be done within a single instance of the operating system.

In order to preserve locality on a large multiprocessor, IIS 6.0 offers the ability to configure an application pool to have affinity for a specific processor. In the above example, an administrator can set up eight application pools to have affinity with the building block processors. The below illustrates what this code would look like:

```
<IIsApplicationPool   Location ="/LM/W3SVC/AppPools/BuildingBlock0"  //Procs 0-3
SMPAffinitized="TRUE"
SMPProcessorAffinityMask28="15"
>
<IIsApplicationPool   Location ="/LM/W3SVC/AppPools/BuildingBlock1"  //Procs 4-7
SMPAffinitized="TRUE"
SMPProcessorAffinityMask="240"
>
```

**Figure 6 IIS 6.0 metabase configuration sample for processor silos** [29]

By assigning an application to an application pool, that application will then automatically maintain its locality. Items of work assigned to that application will not go outside the building block and pollute other caches or minimize the necessary operating system housekeeping.

Note that this is not an optimal setup for all scenarios, because an application could require more than four CPUs to process, and in the above setup, would be limited to four. Additionally, the above setup does not address reliability aspects at all. You may want to configure a separate application pool for an application purely to ensure that it is isolated from other applications. Of course, you can define an application pool for one application and establish affinity between that application pool and a processor.

The other consideration here is networking. If your application is hosted in an application pool locked to processors 20-23, having your entire network I/O for the application come in on processor 0 is still going to introduce cross-processor transfers. Most of these large multiprocessor servers allow configurable I/O so you can wire in the network card(s) processing loads on behalf of the application to one of the physical processors in the processor group where the application is being executed. Using the above example, if the application was accepting load on two gigabit network adapters, those network adapters would ideally be wired into processors 20 and 21. You can also use the Network Interrupt Filter (available with the Windows Server 2003 Resource Guide companion CD) to bind a specific network card to a specific processor.

**Software Locks (Resource Contention)**

A potential issue to be aware of when putting an application on a large multiprocessor server is inherent software locking internal to the custom application itself. A software lock is when two independent pieces of work contend over a certain resource during the application run time. The first one to ask for the resource will gain access to it, and the second item will have to wait until the first one has finished using the resource before it can gain access.

The net result of this contention can be that when the specific application has more processors to execute on, it may not display any additional throughput, because all of the work is being funneled through a single resource.

A classic sign that you might have a contention problem in an application is a high system performance counter System-> Context Switches/Sec. When multiple work items hit the contention point, the typical Windows operating system response is to swap out the thread on which the item is executing, from the processor it is executing on, and a new thread will be swapped in to the processor (context switch). If all the work items execute for a short time and hit the contended resource, you will likely see high context switch rates[30].

*How can you alleviate software locks?*

The best way to fix an application that has high contention, unfortunately, is to get a software engineer, skilled in performance tuning, to analyze the common usage scenarios (common code paths for the application). The engineer would assess the locks[31] (contention points) at run time and re-architect those contention points so they offer greater parallelism.

**WebGardens**

Web gardens are designed to help alleviate the effect of an application that has serious software contention issues without needing an application redevelopment exercise. The approach is relatively simple; a Web garden alleviates the contention on a resource by indirectly creating multiple instances of the resource and evenly distributing the load throughout them. The creating of a new instance is controlled through the application pool configuration.

To create a Web garden, you set the number of worker processes property in the application pool configuration. The default number of worker processes for each application pool is one. By changing that default, you create a Web garden with however many worker processes you select[32].

IIS then distributes the load to the specific instances of the application. In a Web garden, the requests for a specific TCP/IP connection are all sent to one worker process, but every new connection is sent to the next worker process in a round robin fashion. This has the effect of localizing workload related to a specific client, but spreading out the workload across all of the worker processes of the Web garden. By spreading out the workload, you also alleviate some of the contention on the shared resources.

Using a Web garden consisting of four worker processes, the following pattern would be achieved:

- Requests from connection 1 go to Worker Process 1
- Requests from connection 2 go to Worker Process 2
- Requests from connection 3 go to Worker Process 3

- Requests from connection 4 go to Worker Process 4
- Requests from connection 5 go to Worker Process 1

This pattern is represented in Figure 5.




**Figure 7: Assignment of HTTP connections to Web garden process members**

**WebGardens and Affinity**
A Web garden can help relieve software contention without the use of hard affinity. It is possible, however, to have an application pool with affinity to a set of processors, and then run multiple worker process instances. All worker process instances of the application pool will share the same affinity to processors.

**SSL Implementation**
The SSL implementation on Windows Server 2003 is different architecturally to SSL in previous Windows Server releases. The good news is that the design of the SSL implementation on Windows Server 2003 is self tuning and aware of the incoming load.

A frequently asked question is How much will HTTPS (SSL) cost me, compared to straight HTTP? The following test was devised to demonstrate the cost of SSL. The test hardware used was an 8-processor, 900-MHz server with 2 gigabit network cards. The SSL certificate key length was 1024 bit, the content used was an 8K static file, with six requests sent per connection. Each new connection performed a separate SSL handshake. The results of the test are below:

| | **Without SSL** | **With SSL (1024-bit key)** | **Difference** |
|---|---|---|---|
| Requests / Sec | 9368.30 | 2462.10 | ~ -74% |
| CPU | 47.81 | 51.72 | ~ 8% |

The test above was purposely done to consume around 50% CPU to show what could be achieved in both configurations from a throughput perspective, at a CPU range that would be acceptable in a number of operational environments. The conclusion, for this load type, is that the throughput for HTTPS is roughly a quarter of what the straight HTTP case is. NOTE: The throughput will vary depending on the size of the data, the number of full SSL handshakes occurring, and the speed of the CPU.

**SSL Accelerator cards**
An SSL accelerator card can be a very effective acquisition to offload the CPU-intensive computations associated with SSL. SSL accelerator cards are now relatively cheap, and can be a great investment for improving the SSL throughput on a server.

A test was performed to gauge the throughput-enhancing capabilities of an SSL accelerator card with a new-generation, very fast CPU computer. The target server was a 1-processor, 3-Ghz HT (2 logical processors), with a 100-megabit network card and 512 MB of RAM.

The test characteristics were six requests to a piece of content that returned a 2K response per SSL connection, with a full SSL handshake after every six requests. The SSL public key length was 1024 bits. The results are below:

| | SSL in Software | SSL Accelerator Card | Difference |
|---|---|---|---|
| Requests / Sec | 1281.39 | 3075.20 | +139.99% |
| CPU % | 100.00 | 76.90 | - 23.10% |

As represented in the table above, the impact of an SSL accelerator card was quite significant to this load. Even when a high-speed processor is used, the SSL accelerator card outperforms substantially. The SSL accelerator card could have done even better, in terms of throughput, but its network interface was bottlenecked in the test.

| **Note** |
|---|
| This test is not a real world test. It is a test designed to show the best-case effect of SSL accelerator cards. |

**Remote (or Centralized) Content**

Many implementations prefer to have one copy of content stored on a Network Attached Storage (NAS) device or file server and have the Web server or application server access the remote content over a private network between it and the source content (Figure 6).




**Figure 8: Remote content scenario**

The main advantages with this scenario are that one copy of content is present and managed in a farm of servers. This might not seem like a big administrative chore, but when the server population reaches over 10 servers, this can quickly become an administrative concern.

Windows Server 2003 is a very good platform for this physical separation of the content from the server that processes the request and the server that stores the content. However, there are some performance considerations to be made, as the latency between when a piece of content is requested and returned grows in this scenario.

**Caching**

| **Important** |
|---|
| There is a very subtle point here. Given that the IIS static file handler is the component that does the file attribute checks on remote content, all requests for remote content need to be funneled through the static file handler to ensure the content being served is not stale (changed on the remote device). |

As the static file handler is a user-mode component, the implication here is that *all* requests for static content (hosted on a remote server) are *not* cached in the kernel-mode response cache. Rather, the user-mode static-file cache is used to cache these files.

If the device that is actually storing the files is a Windows file server, IIS 6.0 can be configured to cache the files in the kernel (see below for details).

There is typically a bigger impact on static content than dynamic content in the remote content scenario. Dynamic content is generally loaded in memory, or compiled and cached. Therefore, one of the big focus areas for IIS 6.0 was the rework of the caching algorithms for remote static content.

**Last-Modified Caching**
The user-mode static-file cache has implemented a new caching algorithm where, by default, instead of asking for change notification for each directory structure (i.e., the remote device tells all of the servers that a directory has changed), the user-mode static file cache simply asks the file system for the last-modified date and time on the cached file. If the file is new, or the file has changed, the cache entry is updated with the new content and served. If the file is the same as before, the cached version of the file is sent out. As a further optimization, IIS adds a staleness check to avoid having to do this last-modified check on a per request basis. Therefore, the last-modified check occurs only every five seconds[33].

**Use of Directory Change Notification for Remote Content**
As mentioned above, Directory Change Notification is the more expensive way to configure remote content (if that content resides on a Windows file server), because the change notifications need to be active and using redirector resources. However, by using this mechanism, remote content can be stored in the IIS kernel-mode response cache and IIS will be informed by the content server that the content has changed and will refresh the kernel-mode response cache.

It is recommended that you test this implementation before using it. If a server is connected to large numbers of UNC paths, this can reduce the servers ability to get change notifications.

To enable the Directory Change Notification, set the following registry key:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\InetInfo\Parameters\ DoDirMonitoringForUnc (DWORD). 1=on 0=off.**

**Implications of Switching to Change NotificationBased Caching**
Several significant tradeoffs need to be considered when deciding which caching mechanism to use. The most reliable and secure method is to use the default, last-modified time-based caching. It is reliable, because it does not rely on the change notification functionality of the file server, and will scale well on a wide application structure (where your Web server has lots of applications pointing at different file structures). Last-modified time-based caching is more secure because share level permission changes will be honored (the cache will access the file often, which automatically triggers the authorization check against the current file system Access Control List (ACL) and the current share ACL).

However, if you are using a file server that reliably reports change notifications *and* you have a small number of sites or virtual directories, using change notification is significantly faster. For example, if you have a single Web site with a couple of applications that include all your content, stored on a Windows .NET file server, it is recommended that you use the change notification method of caching.

**General Performance Tools**
This section introduces some basic tools that can be used to diagnose how applications are performing, system performance and a level of information from the operating system that can identify the workloads and state of the operating system.

**Event Tracing Enhancements**
Event tracing provides a mechanism for managing data that is useful to performance monitors, capacity planning tools, and other applications that analyze information relating to system resource utilization. Event tracing is enabled by running new command-line tools that control event trace log scheduling, collection, and analysis on local and remote computers. Log files can be taken from the server for analysis. Event tracing is a very high-performance implementation of tracing through core Windows components. A benefit of event tracing for Windows for Web and application server loads is that an administrator can enable event tracing on a production server and have it generate detailed statistics and information on the running load, without a dramatic impact on the performance of the production load.

Listed below is some of the information that can be extracted at run time from a server:

- Most-requested URLs
- Slowest URLs
- URLs with the most CPU usage
- URLs with the most bytes sent
- Clients with the most requests

- Clients with the slowest responses
- Statistics on all the running processes (kernel CPU, user CPU, number of threads, etc.)
- Disk statistics
- Files causing most disk I/Os

For more information about event tracing, see Tracerpt in the Windows Server 2003 product documentation.

**System Performance Monitor**
System Monitor is one of the key tools that display the internal system information of Windows at run time. It is a primary tool to use in performance analysis of the Windows platform and should be used as the master source of information about Windows Server internals when doing performance analysis.

The objects that are of particular interest, in the context of this paper, are counters that belong under the following objects:

- .NET CLR LocksAndThreads
- ASP pages
- ASP.NET v1 applications
- Network interface
- Physical disk
- Process
- System
- Web service
- Web service cache

In addition to offering live information about what is happening in the system, System Monitor also offers the ability to log performance information and generate alerts.

**Windows Task Manager**
The Windows Task Manager, another useful tool, provides similar information to System Monitor, but displays it in a summary fashion. Additionally, the Windows Task Manager only displays Windows system information. System Monitor is extensible and will therefore load many third-party libraries when calibrating performance information; however, it does not load any third-party code.

The Windows Task Manager gives an administrator a low-cost way to see the process internals, how busy network cards are, etc.

**LogParser v2.0**
LogParser is a versatile tool that makes it possible to run SQL-like queries against different input sources, and have the results written to the screen or to different output targets.

With LogParser, it is possible to:

- Quickly search for data and patterns in files of various formats, including IIS log files, Windows event log files, generic comma separated value (CSV) files, W3C files, and text files.
- Create formatted reports and XML files containing data retrieved from different sources.
- Export data to SQL tables. You can export entire files or filter the data to obtain only relevant entries.

- Convert data from one log file format to another.

LogParser is composed of three main components:

1. An input source generates records from the specified source. The following input sources are implemented:

    - All the IIS log file formats (W3C, IIS, IISMSID, NCSA, centralized binary logging, and ODBC), and the UrlScan and HTTP Error log formats
    - The NT event Log
    - W3C log files (e.g. personal firewall, WMS, etc.)
    - Generic CSV files
    - Directory structure information
    - Generic text files

2. The SQL engine processes these records; filtering, grouping, and ordering them according to the specified SQL-like query.

3. An output target writes these processed records to the specified target; the following output targets are implemented:

    - Native log format (a generic format to display results primarily to the screen)
    - Generic W3C log format text files
    - XML files
    - IIS log format text files
    - Direct output to SQL tables
    - Text files formatted according to user-defined templates
    - Generic CSV format text files

In SQL terms, LogParser treats an Input Source as a relational table, whose field types and names depend on the particular Input Source selected.

The LogParser can be found at: http://go.microsoft.com/fwlink/?LinkID=100882.

# Dedicated Application Scenario (Web and Application Server Infrastructure - Performance and Scalability)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

The Dedicated Application scenario describes when a server is predominantly used for serving a single, large, specific application. Generally, this application serves some critical business purpose and typically has multiple servers set up as copies of one another with some software splitting out the load and distributing it throughout the physical servers running the application.

| **Important** |
| --- |
| While many customers prefer this kind of administration approach where a single application is placed per server, this is not the preferred Microsoft architecture for application servers. Generally, we find that this approach leads to inefficient use of computing resources; most of the server deployment done in this design uses grossly under-utilized server hardware. Worse, this deployment style leads to higher costs in a data center, given TCO is directly mapped to the number of operating system instances being managed, regardless if they are running on the same hardware or virtualized in a host-operating system (VMWare). |

Windows Server 2003 has been redesigned to inherently support many applications running side-by-side on the same hardware and in the same instance of the operating system, and if possible, should be configured in such a way as to minimize TCO costs. See the documents referred to in the More Information section for greater detail.

A defining characteristic of the Dedicated Application scenario is that all settings of the application environment, the application itself, and the operating system, can be tuned to get better performance out of the application. This is relevant in that there is no application sociability with other applications that requires consideration when tuning the application.

**Considerations**
This sections covers relevant considerations for the Dedicated Application scenario.

**Application and Content Types**
**Managed Code Applications**
A managed code application refers to applications that are built onto the Microsoft Common Language Runtime (CLR). The CLR offers many services that allow application developers to build applications in a very efficient manner.

Examples are applications built in languages such as Microsoft Visual C# and VB.NET, and using the Microsoft Visual Studio .NET Tools.

**ASP.NET Web Applications (ASPX)**
ASP.NET Web applications are managed code applications that tend to be built as a thin client. For instance, a browser sends a request to a server, and the server platform performs business processing, application navigation, and generates the user interface in HTML. The HTML is then sent back to the client browser and the browser renders the interface for the application. The browser then accepts the next user request and sends a new request to the server.

In this scenario, the vast majority of the processing is being done on the server platform; therefore, it is vital that the request cost is minimized dramatically (in order to get substantial throughput on a server). It is also important that applications are further analyzed for scalability issues.

One of the tools available for additional scalability can be the Web garden feature, which is discussed in detail in the Software Locks (Resource Contention) section.

*When Should I consider the use of a Web garden rather than a single worker process?*

Whenever rich applications are being built, there is a chance that the application developer can inadvertently introduce software contention in an application. See the Software Locks (Resource Contention) section for a description of this effect.

To illustrate this effect and the assistance a Web garden can offer, an ASP.NET page was written where every request had to navigate a linked list with 2000 objects in it to manipulate a particular item in the list. When one request was accessing the list, another request had to wait. This application was placed on an 8-processor, 900-MHz server and had 2,600 simulated clients throw load at the applications. Each client had a think time of 50 milliseconds between each request. The results are below:

| | Single Worker Process | 8 Worker Process Web Garden | Difference |
|---|---|---|---|
| Requests / Sec | 2626.57 | 2961.20 | +12.7% |
| TTLB (milliseconds) | 49.19 | 43.63 | -11.3% |
| CPU % | 92.48 | 100.00 | +7.52% |
| Context Switches | 6808.19 | 3204.87 | -52.9% |

As can be seen, when pushing the server as hard as possible, the test run with one worker process was only able to get up to 92% CPU utilization. This is an indicator of a lock contention problem, no matter how much load you throw at a server, you cannot get it to 100% CPU utilization. With the addition of extra load, IIS would have started throwing back server too busy status codes when the kernel queues became full, and requests would be rejected.

The 8-processor Web garden distributes the 2,600 connections to the 8 worker processes serving the application, 325 connections to each worker process. Because all of the clients were sending in the same load, the requests were evenly distributed across the 8 worker processes in the Web garden. Although still bad, the really bad contention (queuing) behind one lock was split out into 8 locks with smaller queues behind them, thus the 12% gain in performance when the same request load was attempted.

It is important to note that a Web garden is not a silver bullet solution. There is a substantial memory cost to initializing a Web garden. The non-Web garden case had a real memory usage of approximately 35 MB (typical memory usage of a worker process running managed code), while the 8 worker process Web garden had a memory usage multiplied by 8; close to 280 MB used for the application. On a large, 8-processor server that is dedicated to one application, this might be acceptable, but it is a necessary tradeoff that must be considered.

If the actual problem is not software contention within a process, a Web garden will likely have a negative effect on throughput. If, for example, the contention problem for the application is a lock to gain access to a database page on a remote database server, an increase in worker processes is unlikely to have a positive affect. In fact, the additional threads per process are likely to increase overall context switch rates and decrease overall performance. The bottom line is not to assume that a Web garden will help. There are instances where it will, but its utility is in the nature of the contention problem. The only way to know is to test the Web garden configuration in a stress environment with expected loads before going into production.

**ASP.NET Threading Settings and High Latency Workloads**

If an application accesses slow resources, like a mainframe doing large, complicated queries, the threads of the ASP.NET subsystem can get blocked and, for a busy application, ASP.NET itself can run out of threads. If all ASP.NET threads are blocked, performance suffers, as there is no ability for ASP.NET to pick up new work. The work gets queued until other requests complete, and the situation ends up affecting user response times.

It is possible to override the number of threads ASP.NET is using to process requests. The parameter is found in the ASP.NET configuration system and can be changed at the root of the configuration system (the machine.config file), or specifically for the one instance of the application (the Web.config file in the virtual directory of the application). The threads parameter is in the following section of the configuration:

```
<httpRuntime>
minFreeThreads="[count]"
</httpRuntime>
```

**ASP.NET Applications and Queuing**

This section identifies how ASP.NET requests are queued in Windows Server 2003. The pipeline is complicated, but useful to understand when performing performance tuning changes.

1. HTTP requests come in off the network and are received by HTTP.sys.

2. After looking at the request, HTTP.sys places the request in a queue for an application pool.

3. The worker process servicing that application pool picks up the request and makes a decision on what type of processing environment that request needs to be executed. For ASP.NET Web applications, the worker process hands them to the ASP.NET handler.

4. The ASP.NET handler places them on a queue, and ASP.NET threads pick up individual requests and process them.

As you can see, there are a few places where a request could end up being en-queued, depending on how capable of

processing additional requests the subsequent layers are. The queue referred to in step 2 is controlled in the IIS 6.0 application pool dialog box (Figure 7).




**Figure 9: Kernel-Mode Queue Dialog (application pool properties)**

The queue referred to in step 4 is controlled through the machine.config / Web.config configuration system. For busy servers, it might be worthwhile to try increasing this parameter. The parameter is:

```
<httpRuntime>
appRequestQueueLimit=[count]
</httpRuntime>
```

The default queue length is set to 100.

| **Note** |
| --- |
| This parameter also applies to ASMX (Web services). |

**ASP.NET Web Services (ASMX)**

Rich client devices such as cell phones, PocketPCs, tablet devices, etc. have the ability to run rich applications, which communicate over public networks, using a Web service style of application integration. With this in mind, rich client applications connecting to Web services over HTTP have a few subtle differences to the traditional HTTP browser client. These differences can affect their performance.

If you were to trace the network packets associated with a Web service request over a physical network, and compare it to a standard ASPX Web application request made by a browser, you would notice a few differences:

1. The typical size of packet a browser would send is quite small (~400 - 500 bytes). It consists of the TCP/IP headers, an HTTP GET verb, the requested URI (/downloads/downloads.aspx), a few request headers that the browser uses to communicate content length, and a few other attributes of the request.

   Given the fact that browsers typically display data and/or graphics, the browser request usually gets back a large response (text/HTML, tens of KB of data) with embedded references to other large content (.gif/.jpeg) files. The browser then requests these individually.

   A Web service behaves a little differently. When a Web service client is communicating to a Web service server, the client is actually invoking a component, therefore the client end is supplying data (a set of parameters) to call the server program. The server program processes the request, generates a response, and sends the response back to the client. Therefore, with a Web service running over the HTTP protocol, you have a situation where the client is sending an HTTP POST (rather than a HTTP GET), and the client is sending an entity body that contains an XML serialized set of parameters, data structure, and other attributes of the request as part of the POST[34]. Therefore, the client side typically sends larger requests and does not typically get back huge responses.

2. The physical packets put on the wire, and the interaction between client and server, will differ. When a browser-style application makes a HTTP GET request, the request is in one packet, the server responds with a number of packets containing the response to the GET (the response is typically much larger than what can fit in one frame: 1,500 bytes).

   As the Web service client is potentially going to be sending large amounts of data to the server, it behaves differently. Given that Web service clients are likely to be non-computer devices, running on high latency networks with the potential of a per-byte charge on the data sent and received, the Web service stack sends a packet with the HTTP POST, URI, standard headers, and a special header: the Expect100Continue header. This is the equivalent of the client pinging the server to make sure it is worthwhile to send the bulk of its call. If the server responds with 100Continue it is telling the client that everything is OK and the client should send the body of its call. The server then picks up the request, processes the request, and sends back a response. Figure 8 depicts the Web service interaction.




**Figure 10 Expect100Continue packet interaction**

These different behaviors can really make the Web service client interaction efficient, because the client does not spend noticeable amounts of time sending requests that would eventually fail and frustrate the user.

However, when running on an internal network, where cost per byte might not be a concern, the additional latency introduced by the default Expect100Continue behavior, can slow down an applications responsiveness and the overall throughput of a server. As an example, a Web service test was created to measure the difference between the Expect100Continue communication mechanism and the same request done as a single HTTP POST. The table below summarizes the results for one Web service invocation:

| | Default (Expect100Continue) Behavior | All in one HTTP POST | Difference |
|---|---|---|---|
| Total Data (bytes) | 13329 | 11733 | -13.6% |
| Total Ethernet Frames | 31 | 21 | -47.6% |

The data numbers themselves are not necessarily a concern on a fast network, the bigger issue is the latency associated and potentially experienced at the client side, especially given that the TCP/IP client stack will, at times, invoke its nagling[35] behavior and delay small packets from being sent, in the hope that the application will fill the buffer of a packet more and make the overhead of transmission smaller.

If there is a perceived slowness from users (perceived high latency in applications), the above discussion is worthy of investigation. First, a network sniff on the client side of the application should be performed, and the parameters manipulated accordingly.

**How to Change Expect100Continue and Nagling Behavior**

Both the Expect100Continue behavior and nagling behavior can be changed on a per server or per application basis in Windows Server 2003. By default, both behaviors are on, that is, the value in the configuration system is true. To change the parameters for a system, you need to edit the following file:

%SYSTEMROOT%\Microsoft.NET\Framework\vVersion of the framework\CONFIG\machine.config

In the following section you need to add the properties and override them.

```
<system.net>

<settings>

<servicePointManager expect100Continue="true"|"false" useNagleAlgorithm="true"|"false"/>

</settings>
```

</system.net>

To change for a specific application, you modify the app.config file in the same directory as the application. The same change as above is made.

| **Note** |
| --- |
| Expect100Continue and useNagleAlgorithm are changed on the client side of the Web serviceor the consumer of the Web servicenot the server side. The server will behave accordingly to whatever behavior the client displays when interacting with the server. |

**.NET Remoting vs. Web Services**
Applications will often have many component parts. For scalability reasons, these component parts are sometimes separated onto different servers. This predicates a need for some form of object-based communication mechanism to glue the component parts of an application together. In the managed code world, two possible options are .NET Remoting and the use of Web services.

There are some functional differences between the two, which might force the use of one over the other. Assuming no functional restrictions, the use of .NET Remoting, using TCP/IP as a transport and making inter-object calls using the binary formatter, is the most performant option. This is discussed in depth (along with operational and functional tradeoffs) in the document:

http://msdn.microsoft.com/library/en-us/dnbda/html/bdadotnetarch14.asp.

**Unmanaged Applications**
Unmanaged applications are those that do not run on the Windows .NET Framework. This is the previous generation of applications (circa Windows NT 4.0 and Windows 2000), which are comprised of binaries compiled directly from Microsoft Visual C++ and other compilers, and also includes applications that are hosted under the VBScript and other script engines (interpreted applications). The tools used to build unmanaged applications include the following, among others:

- Visual C++ 6.0 and earlier
- Visual Basic 6.0 and earlier
- ASP (Script engine)
- PowerBuilder
- Microsoft Visual J++

The unmanaged code environment is much less forgiving for an application developer, because features such as automatic garbage collection, memory coalescence, pointers, and other productivity language features, are not available.

**ASP Scripts**
ASP technology (.asp) is the predecessor to the ASP.NET page (.aspx). A little known fact is that ASP pages are actually compiled. The high-level process for executing an ASP page is:

1. The ASP page (raw) is read by the ASP subsystem in IIS. That page is converted into what is called a template (the ASP page with all included code inserted and syntactically put together, and read for a script compiler to execute).
2. The template is then sent to the relevant script engine (VBScript or JScript) for compilation and execution.

The ASP subsystem caches both templates and script engines. Like any other cache, this avoids processing costs associated with executing a page again. For example, overriding the *ASPScriptFileCacheSize* metabase parameter, allows ASP to cache more of the templates referred to above, saving the cost of converting the raw script into a template.

Overriding the *AspScriptEngineCacheMax* metabase parameter allows ASP to store more instances of VBScript and JScript script engines in memory. If these engines had been used in the recent past, to execute a page for a different user, a further gain would have been realized, since the script engine does not need to recompile the page. The in-memory, already-compiled page is executed.

| Note |
| --- |
| A single script engine can only execute one page at a time, therefore, if 10 concurrent users all requested the same page at the same time, the system would have 10 concurrent copies of the script engine in memory. |

**Persisted ASP Template Cache**
In IIS 6.0, a further enhancement for ASP templates was undertaken. Templates can be written to disk to avoid the overhead of keeping them in memory and having to regenerate them.

If a site consists of numerous ASP pages, this cache de-allocates the oldest templates from memory to free space for new ones. If one of these ASP files gets requested again, the ASP engine loads the template instead of loading the ASP file and spending additional CPU processing.

**ASP/COM+ Applications**
Many ASP applications instantiate and call into COM+ components. The idea is that the ASP page provides HTML formatting, navigation, and can call into COM+ application components, which contain business logic. An additional advantage of using COM+ components, is that the COM+ components can be re-used in client and server-based applications, without modification.

Threading Model

One of the key things to be aware of in tuning an ASP/COM+ application on Windows Server 2003 is the threading model. The major applicable threading models are Apartment, Free, and Both. Apartment threaded components are COM objects that cannot function with two separate callers inside the component simultaneously; only one caller can be in the context of a method at a time. The converse is true for free threaded components. A free threaded component internally protects its data structures so that any client can call in at any time. Both threaded components can exist in other modes and will typically inherit the mode of objects that are calling into them.

As a typical guideline, COM components developed in Visual Basic 6.0 or Microsoft Visual J++, are all apartment threaded, while components built using Visual C++ 6.0 can be any of the threading models. These components should be inspected for appropriate threading.

The COM+ environment in which ASP executes is, by default, running in the single-threaded apartment (STA). Therefore, COM objects that are instantiated are placed in an apartment separate from ASP code execution. This isn't generally a bad thing, except that for execution control to transaction to the COM+ object, the operating system must stop a thread and perform a context switch. This can become a performance issue when it occurs thousands of times a second.

On Windows Server 2003, a new feature for ASP has been developed for ASP/COM+ applications that only involve COM+ components that are both or free threaded. Windows Server 2003 allows ASP scripts themselves to run in the multi-threaded apartment (MTA). This means that the ASP script itself will call directly into a free or both threaded COM+ object without having to cause a context switch and add threads to the system. This helps the scalability and overall system performance.

The most important consideration here is that the components called by the ASP page must be free or both threaded. For example, an ASP script that calls into the **ADODB.RecordSet** object can do so, without problem, under the MTA, since the **ADODB** object has a threading model of both.

After checking the components an application uses for its threading model and ensuring theyre either both or free threaded, an administrator can enable all applications on the server to run in the multi-threaded apartment by performing the following:

**cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AspExecuteInMTA 1**

Administrators can also run the above command in the specific application directories (for example, W3SVC/1/AspExecuteInMTA) to make the setting for a particular application.

Synchronous Execution

Given the nature of ASP, all execution is done synchronously, therefore, when an ASP application accesses resources off the server or resources that have a natural high latency, ASP can quickly run out of threads. Once ASP exhausts its thread allotment, the response time of a system can be drastically affected.

This is an area where an administrator or operations person should have knowledge of the application. Guidelines for how to handle highly synchronous applications can be found in the section entitled High Latency ASP Pages.

**Large Multiprocessor Single Application Scalability**
For the Dedicated Application scenario, the scale up of a single application is very similar to the considerations for scale up of applications on Windows platforms in general. Please see the Multiprocessor section for more information.

# Dense Hosting Scenario (Web and Application Server Infrastructure - Performance and Scalability)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This scenario is the opposite of the Dedicated Application scenario. The objective for the Dense Hosting scenario is to get as many different applications or sites on a physical server as possible, with only one operating system instance to manage, to minimize hardware and management (traditional TCO) costs. This is by far the most efficient way to run. Any form of virtualization and operating system partitioning[36] in this scenario is giving away precious processor cycles and RAM resources, which is contrary to the objective of cheap, efficient processing.

This form of dense hosting is prevalent at Internet Service Providers (ISPs), and can also be found in some enterprise environments where a central IT group identifies many non-critical applications executing on separate servers and brings them into a central environment. Note the term non-critical application: hundreds or thousands of applications can be hosted on a single piece of hardware, but should a hardware failure occur, many applications would experience some form of downtime while the servers hardware was being rectified. In an enterprise environment, this could be an extremely cost-effective way of providing service to a number of Tier 3 applications, where a 99.999 up time SLA (Service Level Agreement) is not required.

Windows Server 2003 has been tested with a 20,000+ defined applications or sites per operating system instance. This means that the server can easily support that number of applications defined. It does not mean that if these applications all have high resource needs, it will support them. If one of the 20,000+ application is not gated in its resource usage, it can consume 100 percent of the resources on the server, not allowing the other applications to function efficiently.

**Considerations**

This section contains relevant considerations for the Dense Hosting Scenario.

**Application Pools Shared vs. Independent**

One of the central tenets to be observed in this scenario is the level of application isolation. Note that the fundamental unit of isolation in Windows is the process. The process object in the Windows kernel offers operating system and processor instruction protection from one process to another. If one application faults in process A, another application running in process B will not suffer due to application As fault[37].

In IIS, the fundamental unit of isolation is the application pool. Administrators need to determine whether all applications/sites need to be put in their own application pool or can some applications share an application pool with other applications.

If you choose to have multiple applications sharing the same application pool, there is a chance that, at some stage, some of them will experience downtime when one application has an issue.

**Running Different Application Types Side-by-Side or Within the Same Process Space**

Windows Server 2003 offers the ability to mix different kinds of application types together in the same application pool. This means that you can take an existing ASP/COM+ application and run it in the same application pool as an ASP.NET (or managed) application. This also means that you can host native, compiled applications (ISAPI or Win32) in the same application pool as these other application types.

Note that running many different application types in the same process can save system resources (memory and CPU processing), as there is a fixed cost per process for these resources in the Windows operating system. Imagine you had two ASP.NET and one classic ASP/COM+ applications. By pooling these three applications together in the same application pool, you get the following memory footprint:

| All Applications Pooled in Single Process | Each Application Per Process |
|---|---|
| Base Process / IIS fundamentals | Base Process / IIS fundamentals |
| ASP.NET/CLR Runtime | ASP.NET/CLR Runtime |
| ASP Runtime | ASP.NET Application 1 |
| COM+ Runtime | Base Process / IIS fundamentals |
| ASP.NET Application 1 | ASP.NET/CLR Runtime |

| ASP.NET Application 2 | ASP.NET Application 1 |
|---|---|
| ASP/COM+ Application | Base Process / IIS Fundamentals<br>ASP.NET/CLR Runtime<br>ASP.NET Application 2<br>Base Process / IIS Fundamentals<br>ASP Runtime<br>COM+ Runtime<br>ASP/COM+ Application |

As you can see, to run the same three applications in separate application pools, you incur a substantially different set of costs in the isolated case vs. the shared case. But, the reliability tradeoff can be a factor in the decision of which application environments you configure in a shared or isolated way.

From an administrative perspective, the decision can be based on a number of different factors:

- **End customer.** You might decide that each business unit needs to run all of their applications in the same application pool because an application failure would likely be the fault of one of their applications. This would put the emphasis on each business unit to correct their own applications so that they do not incur outage themselves. In this case, the other business units running on the server would not incur outage, because they would be running in different application pools. For an ISP scenario, replace the term business unit above with customer, so that all applications from one customer run together in one application pool.

- **Known, problematic applications.** You might start off with a policy to keep all applications running together in the same pool to minimize the resources utilized, and when applications start to display issues (or failures), you could partition out the problematic applications. It is quite possible to segregate a specific virtual directory containing a particular application into its own application pool, with the rest of the related applications executing in a separate application pool.

- **Everything isolated until the limit.** In this scenario, all applications/sites are put in their own isolated application pools. Then, when the server resource utilization exceeds the administrative policy defined for servers in that environment, you consider which applications/sites could be pooled together. An effective strategy at that stage could be to start pooling together ASP.NET (or managed code) applications into the same application pool, or start pooling ASP-only applications into one application pool.

**All Application Pools Running with Unique Identity**

Another benefit of the IIS 6.0 re-architecture is that the administrator gets to assign the base identity of a worker process. This does not mean that custom code running in the worker process will run with this identity. IIS will still impersonate the caller when calling custom code (ASP, ASP.NET, ISAPI) if the caller is authenticated, or, if the caller is anonymous, IIS will run the custom code under the IUSR_<*computername*> identity. The application pool identity will be used when the IIS code is running in the worker process.

The default identity is the intrinsic account Network Service, which has rights approximating a member of the Users group of a server, but this identity can be set on a per application pool basis (Figure 9).




**Figure 11: Worker process identity tab on application pool dialog**

One consideration that needs to be made clear is that there is a limit to the number of concurrently running desktops in Windows Server. In order to scale this to hundreds or thousands of different identities running concurrently, the following registry key can be set to allow the different identities to share the same desktop.

**HKEY_LOCAL_MACHINE\System\Currentcontrolset\Services\w3svc\UseSharedWPDesktop to be set to (DWORD)1.**

| **Note** |
| --- |
| Sharing a desktop can be a security issue if critical data is being stored or manipulated. Physical windows and other graphical resources are shared on a per-desktop basis. Therefore, the security risk of sharing a desktop is dependant on what data and resources are being stored in desktop objects. For example, if sensitive data is being stored on a window title bar, as an intermediary for an application, other processes that have access to that desktop can gain access to the window title bar, and (if they have enough context) can get at that information. |

**Number of Worker Processes**

Depending on the isolation policy for your application or sites, as discussed above, Windows Server 2003 offers some additional tuning parameters to assist in effective resource management (Windows Server 2003 tries to only have applications running if they are processing actual load).

As discussed in the TCP/IP section, physical TCP/IP connections are managed by HTTP.sys, that is, all connections are held by the System process: PID 4. The effect of this is that, if Windows Server 2003 has 20,000+ applications defined on it, it can be listening on the endpoints and can accept TCP/IP connections for any one of those sites, regardless of whether the actual application is physically initialized. In fact, Windows Server 2003 can accept hundreds of connections without even having one physical worker process running.

The other side of this equation is that Windows Server 2003 can idle out the running application if it has been idle, without dropping TCP/IP connections that have sent requests to that application. This is a very useful feature for administrators looking to trim resource usage on a server.

**Demand Start and Idle Time-out**

Windows Server 2003 always performs a demand start for applications. When IIS 6.0 initializes, no worker processes are running, they only start when a client request for the application pool is received. The default idle time-out for application pools is 20 minutes. This means that if a set of applications running in a specific application pool has not received one request in 20 minutes, IIS 6.0 will spin down the worker process and give those resources back to the operating system.

Administrators can make this setting more aggressive. For applications that have very infrequent use, the idle time-out setting can be set to five or less minutes. Idle time-out can be found on the Performance tab of the Application Pool properties dialog box, or can be set through the adsutil command-line utility:

for a specific application pool:

**cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/IdleTimeout 5**

or, for all application pools:

**cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/IdleTimeout 5**

**Process Gating**

Some administrators prefer to simplify the administrative model and have all applications or sites totally isolated from each other. This configuration can greatly simplify administration and overall management; however, a downside with this approach can be forecasting the specific traffic pattern and usage of all applications (unless you have a small number of total application pools). This means that there can be a potential issue in having an excessive[38] number of concurrently-running processes in the system at a given time.

Therefore, Windows Server 2003 offers a process gating feature, which allows an administrator to set a high water mark for the number of active processes in the system at any given moment. The administrator can simply reset this value to be appropriate for the type of hardware the administrator has (the number of concurrent processes Windows Server 2003 can support is mostly limited by the amount of physical RAM in a system).

**How does process gating work?**

When a request is sent to an application pool that has not been started (or has idled out), IIS 6.0 makes a check of the number of concurrently-running processes. If the total number of processes, including the new one, would go over the threshold, the worker process is not started and a 503 Service Unavailable error response is sent to the request originator.

The process gating feature is not available through the IIS UI management tool. It is set in the IIS 6.0 metabase only. To override the process gating feature to have a maximum of 100 concurrently active application pools, run the following command at the Windows command prompt:

**cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/DemandStartThreshold 100**

to remove the override to get IIS 6.0 back to its standard behavior, run the following command:

**cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs delete W3SVC/DemandStartThreshold**

| **Note** |
| --- |
| The process gating feature is not on by default. This means that IIS will not gate processes by default and it is possible to get as many worker processes started as there are application pools. Further, if application pools are configured to run as Web gardens, it is possible to get even more concurrent worker processes than the total number of application pools (dependant on the number of worker processes set per Web garden). |

**IIS Logging Considerations with Thousands of Sites per Server**

On IIS 6.0, whenever a site is defined, W3C logging is established for that site. In order to not affect performance, every log file generated gets a 64K buffer assigned to avoid writing log entries to disk too often. This way, many log records are buffered and all written at one time.

When thousands (or tens of thousands) of sites are defined on a server, the memory use due to these log buffers can be substantial (64K per site). Therefore, to conserve memory, there is a way to control the size of the log buffers. Obviously, when the log buffers are decreased in size, the frequency of I/O to the disk is increased, but for sites that get infrequent hits, this may not be a massive impact on performance.

The parameter to manipulate the size of *all* log buffers on Windows Server 2003 is:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Http\Parameters\LogBufferSize (DWORD)**

with a range of: 12,288 (12K) 65,536 (64K). The value the parameter is set to must be divisible evenly by 4,096, otherwise it will be discarded and the 64K default will be used.

| **Note** |
| --- |
| Log buffer size cannot be changed for specific sites; the change is global for all sites on a server. |

**Centralized Binary Logging**

When running thousands of sites on a server experiencing high request rates, the centralized binary logging feature is very useful. Centralized binary logging alleviates a disk I/O bottleneck; it makes the log file size smaller and saves CPU in formatting

the log file information.

From a functional perspective, there are some restrictions on the centralized binary log file format (in terms of the fields logged), which cannot be changed. For more details on the binary logging implementation and comparison data, see the New Format in IIS 6.0 Centralized Binary Logging section.

**Connection Management Discussion**
Running a server with many sites can cause the server to hold onto many connections. The reason for this is that the majority of clients today are browsers. Most browsers send HTTP requests over the HTTP 1.1 protocol with a Connection: KeepAlive header in the request. This directive tells the server that once the browser->server connection is established, it should remain open. The reason for keeping the connection open is to forgo the connection establishment and latency for subsequent requests.

However, from a server perspective, having many connections established (tens of thousands per server possible) means that the server is investing some CPU in managing those connections, and the network infrastructure has an additional load in sustaining them. For these reasons, many very large implementations turn off keep alive connections at the server end, which means that exactly one request is sent per connection. Browsers and other clients need to go through the session establishment process for every request.

Administrators should gain an understanding of how their production servers behave with the KeepAlive header off as well as with it on. The resources to monitor are CPU, memory, network utilization, and latency at the client end. The right decision is based on operational policy for the organization.

For information about how to disable and enable HTTP Keep-Alive connections, see Enabling HTTP Keep-Alives and AllowKeepAlive in the IIS 6.0 Help.

**IP Addresses or Host-headers**
When placing many sites on Windows Server 2003, it is important to have a strategy for how sites are to be addressed. If the server is to support thousands of sites, the preferred mechanism to address them is via host-headers, rather than TCP/IP addresses or different ports.

For a detailed discussion of the considerations behind the way sites are addressed, see the TCP/IP Configuration and Connections section.

**Remote Content Considerations and Tuning**
Before you set up an IIS server that will have hundreds or thousands of sites on it, you must first decide where the content is going to be located. IIS 6.0 supports both local and remote content models. Local means that the physical files that make up each site are located on a disk attached to the system. Remote means that the content files are physically located on a remote file system (file server, NAS device, etc.)

The advantage to having remote content in this environment is that it becomes even more straightforward to reorganize which server is handling the requests for a specific site. If there is a need for reorganization, the server configuration for the site and DNS needs updating. The physical content is untouched and unmoved. If the content was local, it would require moving, and backing up that content might also require reconfiguration.

Even though the remote content case looks good, there are some performance and scalability considerations when using content on a remote share such as latency on content and access patterns. The Remote (or Centralized) Content section also introduces a discussion regarding the caches and parameters to control for the caches.

**Latency**
When content is hosted on a local server, the average access time of a piece of content is usually measured in milliseconds. When a content request is going across a network (even a very fast one), you might get to the content in milliseconds, or it might take seconds. This can have a dramatic impact on the responsiveness perceived by the user. Of course, once IIS 6.0 has seen the content, it will try to optimize its access patterns to make subsequent accesses much more efficient.

**Data Volume**
Similar to IISs ability to saturate several gigabit network cards, with the data being sent back to many thousands of clients, it is also possible to saturate the network between an IIS server and a NAS (file server or other computer), and hold the raw content. If a network card on either the Web/application server or shared content server is saturated, the overall performance of all the servers accessing that central server can be affected. This must be planned for in your design.

**Questions to Answer When Testing a Remote, Centralized Content Scenario**

- What is the likely profile of requests over a period of time to the Web/application servers? (LogParser v2.0 can use existing, typical IIS logs to help deduce a reasonable answer to this question.)

- What is the networking between my Web/application server and my central store?
- How many Web/application servers hit my central store at once?
- How will a server cope that is in a cold state,[39] just added to production? How does the server cope when all of the content is local? What happens when all of its content is on a remote server?

**Errors with UNC Content Under High Load**

If the Web server and the remote file server are not configured properly for high stress conditions, administrators might notice some errors. The following errors can show up on browsers, in the system event log on the IIS server/remote file server, and in the IIS logs.

- *Not enough server storage is available to process this command. Error code: 0x8007046.*
- *Windows cannot find the network path. Verify that the network path is correct and the destination computer is not busy or turned off. If Windows still cannot find the network path, contact your network administrator. Error code: 0x80070033.*
- *The network location cannot be reached. For information about network troubleshooting, see Windows Help. Error code: 0x800704CF.*

**MaxUserPort Setting**

The last two errors can be returned when the IIS server has a shortage of outbound TCP ports while making the UNC connection to the remote file server/NAD device. The Web/application server will then return HTTP 500 errors for all requests, until TCP ports are available again. To fix this problem, you can use the MaxUserPort registry setting (on the Web/application server) to increase the number of available ports.

For instance, if there are a large number of concurrent authenticated users (Basic, Digest, Kerberos/NTLM, Microsoft .NET Passport) a different Server Message Block (SMB) connection is opened for each individual user, from the Web/application server to the remote file server. Every SMB connection uses ports. By default, you are limited to ports from 1,024 to 5,000 for outbound TCP connections (a little less than 4,000 ports available), so you may need to increase this value to a higher number, such as 10,000 or 20,000, to allow for additional ports. Be aware that this may cause problems with other applications creating sockets on these higher ports. The following MSDN document discusses this further: http://www.microsoft.com/windows2000/techinfo/howitworks/communications/networkbasics/tcpip_implement.asp.

**Windows-based Remote Content Store**

This section only applies when a Windows Server is being used to house the content that several Web/application servers are accessing over Windows UNC file shares.

The first error message above (*Not enough server storage is available to process this command. Error code: 0x8007046A.*), requires a check of the Windows file server settings. Usually, that IIS error accompanies event log messages like the examples below:

- The server was unable to allocate from the system non-paged pool because the pool was empty.
- The server was unable to allocate a work item <n> times in the last 60 seconds.

Normally, those messages mean that the file server was set to use too much non-paged system memory. This can be quickly verified by opening Windows Task Manager, going to the Performance tab, and looking at the non-paged value under Kernel Memory frame. You should throttle it back and limit the amount of kernel memory the Server service should use.

When troubleshooting this class of errors, you will need to check these performance counters on the Web/application server:

- Server\Files Open
- Server\Server Sessions
- Server\Work Item Shortages

The last class of errors above can typically be fixed by tuning the following registry entries related to SMB[40] configuration, namely: **MaxCmds**, **MaxMpxCt**, **MaxWorkItems**.

Between two Windows servers, for each user, one SMB connection is used; however, within that SMB connection, there are a finite number of work items. Work items can be consumed in a variety of ways and for varying amounts of time. For example, doing an operation against a file (for example CreateFile or GetFileAttributes) only takes up an I/O work item for a short amount of time, but asking for a change notification on a directory structure will consume a work item for as long as the connection is established.

Here are some general details about work item configuration:

- Only one SMB connection is created per Web server / file server connection, per user.
- Each change notification instance uses a work item until the connection is closed.
- Each request from the Web/application server to the file server is going to temporarily use a work item.
- The registry settings at this location: **HKLM\System\CurrentControlSet\Services\LanmanServer\Parameters** can be updated on the remote file server, and the registry settings at: **HKLM\System\CurrentControlSet\Services\LanmanWorkstation\Parameters** can be updated on the Web/application server.
- For these registry values to take effect, you must stop and start the server service on the remote file server and the workstation service on the Web/application server. The servers do *not* require a reboot.
- Each work item typically takes up about 20K of non-paged pool memory. Although certain KB articles may specify setting the **MaxWorkItems** value to 16,000 this could consume up to around 320 MB of non-paged pool memory, which could starve the system. Always set this number lower to increase the reliability of your system.
- The rule of thumb for calculating **MaxCmds** is as follows: (# of distinct physical directories that IIS needs to monitor for change notifications) * (1 (if static files exist) + 1 (if ASP content exists) + 1 (if ASP.NET content exists)) + 50 (for concurrent default/regular File I/O).

| **Note** |
| --- |
| This calculation merely hints at the recommended setting. You may need to bump your settings higher depending on concurrent traffic. |

The following are the default and maximum registry values for some SMB settings on Windows Server 2003 (same as on Windows 2000):

| Registry Setting | Default Value | Maximum Value |
| --- | --- | --- |
| **MaxCmds** | 50 | 65535 |
| **MaxMpxCt** | 50 | 65535 (Windows 2000 SP2+; max 125 on W2Kpre-SP2) |
| **MaxWorkItems** | scaled on memory size, usually between 1K 4K | 65535 |

**ASP.NET Web Applications and Web Services in Dense Hosting Environments**

The richness of the managed code environment also means that the memory cost of running that environment is bigger than the cost of running applications in an unmanaged environment. When running managed code applications in a hosted environment, the first consideration is the deal the customer signed up for. If part of that deal is to have a very high level of reliability, it would be pragmatic to put that application in its own application pool.

If applications do not have stringent isolation requirements, it is fine to have a strategy of putting managed code applications into the same application pool. In this case, ASP.NET will place these applications into their own application domain within the worker process regardless; however, the risk is that if they are sharing application pool with an unmanaged component that happens to cause a failure, all applications in that application pool are be affected. Additionally, maintenance on any application in the application pool that requires a recycle will also temporarily affect the responsiveness of all other applications in that

application pool.

It is very important to understand what the maximum number of processes and memory footprint is on a server. This information can help an administrator size the server adequately so that the server has resource when many of the applications are initialized. This can be done by selectively pinging each possible application pool on a server and seeing what the impact is on memory, CPU utilization, and disk I/O overall responsiveness. When a comfortable threshold is reached, an administrator can consider using the (worker) Process Gating parameter[41] to lock down the maximum number of concurrent worker processes on a system. See the section entitled Number of Worker Processes for more information about Process gating.

# Server Consolidation Scenario (Web and Application Server Infrastructure - Performance and Scalability)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

The Server Consolidation scenario is a special case of dense hosting in that the goal of this scenario is to put several unrelated Web applications or Web services on the same server. The big distinguishing factor is that in the Server Consolidation scenario, typically, the applications have much higher reliability requirements. Additionally, the applications will typically have higher request rates and there will be fewer applications per server than the Dense Hosting scenario.

Server consolidation is becoming more and more popular; it is quite common to find companies having many servers running a single application. Even though many of these applications are mission critical, they do not have high request rates or large processor or memory needs. In fact, they barely use the resources available to them on the servers they are running. The management cost of these servers can be improved proportional to the number of operating system instances being managed. Therefore, by consolidating these applications on a single server, a company decreases management costs and sometimes gets better operational management of the remaining operating systems.

| **Note** |
| --- |
| The following discussion will focus on the homogeneous consolidation of Web applications and services, not the heterogeneous consolidation of Web applications with e-mail, database, etc. This discussion does not focus on the versioning and side-by-side considerations of the server consolidation scenario. |

**Considerations**

This section contains considerations relevant for tuning Windows Server 2003 in an (Web application) Consolidation scenario.

**Independent Application Pools and Independent IIS Sites for Applications**

This is a very likely configuration option in a Server Consolidation scenario. Given the reliability needs of applications, having them in separate application pools isolates failures and outages due to other applications. Applications can be managed separately; if they have bugs, they only have an impact on themselves, they can be put out of service for maintenance independently, without affecting other applications still online.

Having different applications or Web services running in their own application pool is also useful from a performance tuning perspective. Many settings are based on the application pool setting properties.

**Independent Sites**

The other setup configuration option, besides separate application pools, for a server consolidation scenario, is to have separate applications (or Web services) placed into separate IIS sites (Figure 10). There are additional settings and properties that can be controlled which are set at the site level. Examples include compression, bandwidth throttling, and total allowable connections to an application.

Creating an independent site for each application is simple. Each application must be accessed through a separate TCP/IP address, server port, or DNS name. The preferred approach is to use a separate DNS name per application, as this is much more efficient than using a separate IP address. Using a separate server port makes accessing applications a little more tedious (the port must be appended to all requests) when users are using a browser to access applications. If the application in question is a Web service, this is not as great a concern, since the separate port is usually ed away by a rich client or other program calling into the Web service.




**Figure 12: Site, application, and Web service isolation capabilities**

**Resource Control**

The application pool construct allows many settings to be applied per application. Many of these settings are resource controlled, but they also have performance impacts.

**CPU**

This is a primary resource control performance tuning parameter. The CPU utilization control available on the application pool object controls what percentage of the total CPU of the server that specific application gets.

The CPU control on the application pool object is the average amount of CPU the application should consume over unit time. So, if the CPU control on an application pool is set to CPU Utilization: 20%, for Time: 4 minutes, and an application consumes 40 percent for 2 minutes, the action for the application pool is fired. That means that if the application pool is set to Shutdown (should the CPU utilization be exceeded), the application pool will be shut down for 4 minutes, until the next period determined by the CPU occurs. Then, the application pool will be reactivated. If the action is set to No action, the excessive utilization will log an event in the system event log.

It is also possible to control the CPU utilization of a specific application in a server consolidation scenario, by configuring affinity between a processor and an application. For a discussion of affinity and large multiprocessor servers, see the Multiprocessor section (above). By establishing affinity between a specific application and CPU, you automatically limit the CPU available to an application. For example, if you are running on a 16-processor server, and the application has affinity with 4 processors, the maximum CPU that application has access to is 25 percent[42].

| **Note** |
| --- |
| Notice that the default IIS capabilities do not allow true throttling of the CPU on an application pool. For throttling, see the Windows System Resource Manager Tool discussion below. |

**Memory**

Memory utilization is an important performance parameter of an application. IIS 6.0 allows for the control of the memory utilization on a specific application, by limiting the Private memory and Virtual memory threshold settings. If these parameters are set too low, and the application exceeds the threshold frequently, the application will be recycled every few minutes, which will have a performance impact, given the system overhead of spinning up and warming up new processes and libraries.

In a server consolidation scenario, it is a good idea to have some excessive capacity on the server where you first place an application. Give the application an above-average amount of resources and monitor it for a few days, under normal usage. This will give an impression of the resource needs of the application. Then, select the resources you expect to constrain on the application and constrain them one by one. All the while, continue measuring the impact of the change, by assessing if the average response time of the application has increased. If so, you may have set one of the resources too aggressively for the application.

**Identity**

This is another property set on a per application pool basis. If set as a domain account, there will be slightly higher latency when starting processes, given the additional latency involved with communicating with a domain controller to log on the base process identity. This doesn't mean that the identity will be accessed per request, only at process startup.

**Logging**

This is an important consideration in the Server Consolidation scenario. The decision is between having a single log file for all the applications on the server or having a separate log file per application.

Having a single log file per server is efficient if it is getting many requests (high throughput rate), because the centralized binary log file allows administrators to place the log file on a separate device dedicated to logging. That style of logging is most efficient: the disk subsystem gets to optimize its writes for the one, sequentially-written log file. The downside with the centralized binary log file is the need to use the LogParser utility to separate out log entries for specific applications (all applications on the server have their log items in the same file). For more information, see the New Format in IIS 6.0 Centralized Binary Logging section (above).

By contrast, leaving the standard logging on will generate a separate log file for every site on the server, but this could also mean high disk latencies in a disk I/O subsystem. The determination is largely dependant on the operational policy, the needs of logging in the environment, and the request rate experienced by the server.

To better understand the behavior and performance consideration of the Logging subsystem, see the Logging section (above).

**Data Bandwidth Throttling**
This can also be a key consideration when supporting many important applications on a single, physical piece of hardware. It is good practice to understand the total bandwidth available on the server and to restrict applications in terms of the maximum bandwidth they can use on a server.

This is not a reservation scheme; an application is not guaranteed that amount of bandwidth, each application (or site) is limited in terms of how much bandwidth they can consume. Further, just because the network devices in a server have published data rates, does not mean it is possible to achieve that data rate given the limitation of the number of physical cards on a root bus, and the speed of the bus. For a deeper discussion of the capabilities and considerations of the networking subsystem of an Internet server, see the Networking section.

**Windows System Resource Manager (WSRM)**
WSRM also allows you to manage CPU and memory utilization on a per application basis, but it provides a mechanism for implementing rich application policies and integrating them with a schedule. This means that an administrator can run multiple applications on a server with greater safety.

WSRM policies can be applied according to a time/date schedule. This allows administrators to free up CPU and memory for maintenance applications during non-peak hours, and for mission critical applications during peak hours.

The WSRM accounting feature allows administrators to generate, store, view, and export resource utilization reports for systems management, and service level agreement (SLA) tracking and billing purposes.

WSRM enhances the base IIS 6.0 by providing the ability to throttle the CPU use of a specific application pool. With WSRM, you can enable a situation where an IIS 6.0 application pool uses up to 20 percent of the CPU, regardless of the load coming in.

**Tuning Parameters for Independent Applications**
By placing separate applications in independent application pools, many of the control and tuning parameters are available to administrators when setting up the Server Consolidation scenario. IIS 6.0 allows many different Web application types to run side-by-side. For example, it is possible to take an ISAPI application, ASP/COM+ application, ASP.NET Web application, ASP.NET Web service, and a third-party application environment[44], and run them all concurrently, side-by-side, on IIS 6.0.

Note, however, that for some of the Microsoft application technologies, additional configuration schemes contain some of the performance tuning parameters. For example, the ASP.NET technology also has a series of text files which contain configuration settings. The root configuration file for ASP.NET is the machine.config file (found in **%SYSTEMROOT%\Microsoft.NET\Framework\vvvv\CONFIG**). When an ASP.NET application is run, the properties from this file are read and overrides are applied in the application directories (to modify the machine.config file). Some of the parameters are thread-related and connection-related. This document references machine.config parameters; these files might require settings in addition to IIS 6.0 specific parameters.

Information about the ASP.NET configuration system inheritance can also be found at:

http://msdn.microsoft.com/library/en-us/cpguide/html/cpconconfigurationinheritance.asp.

**Large Multiprocessor Servers**
When consolidating important or mission-critical applications, it is not uncommon to use a large multi-processor server, as there is a greater capacity potential on that server and the hardware will also typically have additional reliability features built into it.

The Multiprocessor section (above) provides a detailed look at the performance, troubleshooting, and setup of large multiprocessor servers.

# Summary (Web and Application Server Infrastructure - Performance and Scalability)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Windows Server 2003 offers significant gains on performance and scalability, over its predecessor operating systems, when used in Web and application server scenarios. Additionally, Windows Server 2003 has many new features for the configuration, management, and tracing of applications running over the HTTP protocol.

**More Information**

1 The following resources contain in-depth information about IIS.

- Technical Overview of Internet Information Services (IIS) 6.0 at http://www.microsoft.com/windowsserver2003/techinfo/overview/iis.mspx
- Microsoft Server Consolidation Kit at http://www.microsoft.com/windowsserversystem/articles/consolidation/default.mspx
- What's New in IIS 6.0 at http://www.microsoft.com/windowsserver2003/evaluation/overview/technologies/iis.mspx
- IIS 6.0 Help at http://www.microsoft.com/technet/prodtechnol/WindowsServer2003/Library/IIS/848968f3-baa0-46f9-b1e6-ef81dd09b015.mspx

For the latest information about Windows Server, see the Windows Server Web site at http://www.microsoft.com/windowsserver2003/default.mspx.

| **Note** |
| --- |
| IIS 6.0 is wholly contained and ships as a component of Windows Server 2003. IIS 5.0 is wholly contained and ships as a component of Windows 2000. It is not possible to interchange IIS 5.0 onto Windows Server 2003 or IIS 6.0 onto Windows 2000, as both versions of IIS are tightly integrated into their respective operating system release. |

2 **worker process:** all Web application processing, including loading of ISAPI filters and extensions as well as authentication and authorization, is done by a new Web service DLL. This DLL is loaded into one or more host processes called *worker processes*, or *instances of W3wp.exe*. For more information about how worker processes interact with IIS 6.0, see the section entitled *Worker Process Isolation Mode*, following this section.

3 **application pool:** an application pool is a grouping of one or more URLs served by one or more worker processes. An application pool corresponds to one request queue within HTTP.sys and one or more worker processes. An application pool can serve requests for one or more unique Web applications. These applications are assigned to the application pool based on their URL. There can be multiple application pools operating at the same time. For more information about application pools, see the section entitled *Worker Process Isolation Mode*, following this section.

4 The IIS metabase is an XML file stored at: **%SYSTEMROOT%\system32\inetsrv\metabase.xml**.

5 8-processor, 900-MHz server (Saber class), 3 gigabit network cards.

6 TTFB/TTFB = Time to first byte/Time to Last byte response time latency from the client perspective.

7 See RFC2616 for a description of HTTP response headers (values, ranges, and purpose.)

8 If possible from a functionality perspective.

9 Test performed on Proliant DL360 2-processor (900-MHz) computer, running Windows Server 2003, Enterprise Edition (build 3700); 1 BCM5701 gigabit Ethernet adapter.

[10] Only ASP.NET pages with VaryByParam=none are cached in the kernel.

[11] CLR = Common Language Runtime, the underpinning of ASP.NET.

[12] This is CPU cycles measured on the physical processor. This number does not relate to the physical number of instructions counted in Microsoft intermediate language (MSIL). This is the resultant processor spins (stalls), cacheline flushes, etc. to get that handful of MSIL instructions executed on a busy system.

[13] A low context switch rate is less then 5,000 context switches per second, per processor. So, an 8-processor server with an overall context switch rate of 1,500 context switches per second would qualify as having a low context switch rate.

[14] When performance testing the Unisys ES/7000 platform (32 processor slots), the IIS test team frequently used 8 gigabit adapters per one operating system instance (32 processors). In some test cases, the 8 gigabit cards were still the bottleneck.

[15] Web browsers and other clients can ask IIS to send back a compressed response by including the Accept-Encoding: gzip, deflate request header in the initial content request.

[16] HT stands for Hyper-Threading technology from Intel where a single processor core is represented as 2 virtual processors in the operating system.

[17] The HTTP.sys cache will only cache HTTP GET requests that are anonymous. This includes unauthenticated SSL encrypted responses.

[18] PAE = Physical Address Extensions. An extended addressing scheme for x86 based systems that allows access to greater than 4 GB of RAM.

[19] This location is accurate for the initial release of Windows Server 2003. Microsoft reserves the right to change this location and parameters in future releases or service packs.

[20] Bottleneck is CPU.

[21] Bottleneck is gigabit network card.

[22] IIS will preserve existing per-site log-file settings when the CBL mode is selected, so if you want to revert back, you do not lose any log-file configuration.

[23] IIS will fill a 64K log buffer before physically writing it to disk.

[24] The registry key **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Http\LogBufferSize**, controls the size of the log buffers allocated by HTTP.sys. When this registry key is not present, HTTP.sys will allocate 64K for each log buffer. When this registry key is present, HTTP.sys will validate and use the size nominated as the log buffer. The valid range is 12K-64K.

[25] Windows will reserve large sections of virtual addresses from which to sub-allocate requests, when memory requests come in. This extension doesn't happen on a per-allocation request basis.

[26] See: http://msdn.microsoft.com/library/en-us/memory/base/low_fragmentation_heap.asp

[27] For the purposes of this discussion, a large multiprocessor server refers to a server with more than four physical CPU cores. Note, this is an arbitrary distinction I have made to discuss concepts in this paper. It is not uncommon to consider a large multiprocessor server having greater than 16 CPU cores.

[28] **SMPProcessorAffinityMask** is the decimal conversion of the binary number representing the processors to be used. For example, if an administrator wanted to establish affinity between an application pool and processors 4-7 of a multiprocessor server, the binary number representing that would be 1111 0000. Converting this binary number to decimal yields a value of 240. Therefore, the setting for that application pool would be **SMPPrcoessorAffinityMask=240**.

[29] SMP affinity can only be enabled/disabled through manipulation of the **SMPAffinitized** and **SMPProcessorAffinityMask** properties in the IIS metabase. These values can be manipulated through direct modification of the XML or by using the adsutil utility.

[30] Greater than 5,000 context switches per second, per processor (i.e., more than 40,000 / sec on an 8-processor server) is

generally considered a high context switch rate. Note, however, that context switch rates should also be considered on the request throughput at the time of calibration. For example, if a server is processing 50,000 requests / sec, a context switch rate of 40,000 is quite good (less than one per request).

[31] Example of tools a developer would use to assess locks at run time would be Microsoft Visual Studio .NET 2003 Analyzer or Intel VTune.

[32] An application must be able to be multi-instanced for Web gardens to work.

[33] This five-second interval can be configured; it can be overridden by adding the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\ Services\Inetinfo\Parameters\FileAttributeCheckThreshold (REG_DWORD)** registry key. This value will take effect when a worker process is recycled, so it can be changed in production on the fly.

[34] Because the Web service security standard is built out, some of the data items sent across, per method call, will be the large public/private key encryption values and many other SOAP envelope headers.

[35] Nagling is where a TCP/IP stack only has a small number of bytes to send, so instead of adding all of the TCP headers to a very small payload, TCP/IP will wait for 200 milliseconds in the hope that the application will give it more data to send. If no more data is passed down, TCP will send the small packet. This phenomenon has a higher probability of occurring in Web service interactions, given the split of the packets.

[36] VMWare is an example of an operating system virtualization that consumes significant amounts of system resources on a server, while still maintaining another operating system image.

[37] Assuming the two applications did not have a dependency on one another.

[38] Excessive is a relative term. Having 50 concurrently-running application pools on a server with 128 MB of physical RAM is excessive; having 50 concurrently-running applications on the exact same server with over 1 GB of RAM, would not be excessive.

[39] Cold in this context means a server that has just had IIS initialized, with all the IIS caches empty and no custom components initialized.

[40] SMB stands for Server Message Block protocol. This is the standard protocol used by Windows clients and file servers to share files and printers.

[41] The Process Gating parameter can be found at the W3SVC\MaxProcesses location in the IIS 6.0 metabase.

[42] Given that HTTP.sys does a lot of caching and performs some of the processing for an application, and the CPU use of HTTP.sys is not attributed to an application pool, CPU usage will not be perfect. For example, if an application has affinity to two processors on a 4-processor server and that application is driven as hard as possible, the overall processor utilization will be slightly over 50 percent, given the processing that HTTP.sys does.

[43] Only available for Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.

[44] Examples include Macromedia ColdFusion, IBM Websphere, and BEA WebLogic. NOTE: At time of writing, third-party vendors had not approved their application environments on Windows Server 2003.

# Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

By Brian Komar and David Cross

Microsoft Corporation

Microsoft Windows XP Professional and Windows Server 2003 provide an integrated, public key infrastructure (PKI) that enables you to securely exchange information across the Internet, extranets, intranets, and applications. This white paper provides a technical reference and planning guide for PKI administrators who wish to perform PKI cross-certification, deploy bridge Certification Authorities (CAs), and understand how to implement qualified subordination in Windows Server 2003.

**In This White Paper**

Acknowledgements (Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003)

Introduction (Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003)

Overview (Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003)

Understanding Constraints

Qualified Subordination Deployment Scenarios

Walkthrough (Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003)

Appendix A – Policy.Inf

Appendix B – CAPolicy.inf

Appendix C – CMC Dump of a Qualified Subordination Request

# Acknowledgements (Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

David Cross, Program Manager, Microsoft Corporation.

Trevor Freeman, Program Manager, Microsoft Corporation

Darren Canavor, Software Test Engineer, Microsoft Corporation

Sergio Dutra, Software Design Engineer, Microsoft Corporation

Krish Shenoy, Software Design Engineer, Microsoft Corporation

Jan De Clercq, Senior Consultant, Compaq Global Services

# Introduction (Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Qualified subordination is the process of cross-certifying Certification Authority (CA) hierarchies using basic, policy, naming, and application constraints to limit which certificates are accepted from partner CA hierarchies or a secondary hierarchy within the same organization. True cross-certification of CA hierarchies in a Windows 2000 network was not possible. The only available alternative was to define Certificate Trust Lists (CTLs) that trusted specific CAs and restricted certificate usage. By using qualified subordination, a CA administrator can clearly define which certificates issued by a partners PKI are trusted by the CA administrators organization. Qualified subordination also provides methods for compartmentalizing and controlling certificate issuance within an organization according to policy guidelines. Examples of both scenarios will be explained in this white paper.

For the purposes of this white paper, the term qualified subordination refers to a cross-certification between two CAs that implements basic constraints, name constraints, application constraints, policy constraints, or a mix of the four constraints. This limits which certificates are trusted from a partner or secondary CA hierarchy according to the rules and definitions defined in RFC 2459 and subsequently RFC 3280.

#### Scope

The scope of this white paper is to describe the ways that qualified subordination is used to govern the relationship between multiple organizations PKI hierarchies. The white paper describes the various constraints that can be implemented to define the relationship between PKI hierarchies, presents qualified subordination scenarios, and provides walkthroughs of the qualified subordination process.

#### Terms Used in This White Paper

**Application Constraints** A constraint that limits what purposes a certificate can be used for in a qualified subordination configuration. A presented certificate must contain the required application constraint to be accepted by the partner organization.

**Authority Information Access (AIA)** A certificate extension that contains URL locations where the issuing CAs certificate can be retrieved. The AIA extension can contain HTTP, FTP, LDAP, or FILE URLs.

**Authority Key Identifier (AKI)** A certificate extension used by the certificate chaining engine to determine what certificate was used to sign a presented certificate. The AKI can contain the issuer name and serial number, public key information, or no information at all. By matching the information in a certificates AKI extension to a CA certificates Subject Key Identifier (SKI) extension, a certificate chain can be built.

**CaPolicy.inf** A configuration file stored in the %SystemRoot% folder that defines configuration settings for CAs when they are installed and when the CAs certificate is renewed.

**CRL Distribution Point (CDP)** A certificate extension that indicates where the certificate revocation list for a CA can be retrieved. This extension can contain multiple HTTP, FTP, File, or LDAP URLs for the retrieval of the CRL.

**Certificate Trust List (CTL)** A method of restricting certificates chaining to a designated CA for limited time periods or usages. It is used more prevalently in a Windows 2000 network. In a Windows Server 2003 environment, qualified subordination is the preferred method for restricting certificate usage between organizations.

**Certificate Revocation List (CRL)** A digitally signed list issued by a CA that contains a list of certificates issued by the CA that have been revoked. The listing includes the serial number of the certificate, the date that the certificate was revoked, and the revocation reason. Applications can perform CRL checking to determine a presented certificates revocation status.

**Cross-Certification** The process of issuing subordinate CA certificates for existing CAs that link two root CAs.

**Cross-Certification Authority Certificate** A certificate issued by one CA for another CA's signing key pair (that is, for another CA's public verification key).

**Name Constraint** A constraint that limits what names are permitted or excluded in certificate requests submitted to a CA.

**Issuance Policy Constraint** A constraint that defines what issuance practices must be followed for certificates to be trusted by your organization. Issuance policy object identifiers (OIDs) in your organization are mapped to the matching object identifiers in a partner organization, so that object identifiers in presented certificates are recognized by your PKI.

**Policy.inf** A configuration file that defines the constraints that are applied to a CA certificate when qualified subordination is defined.

**Public Key Infrastructure (PKI)** A PKI provides an organization with the ability to securely exchange data over a public network using public-key cryptography. A PKI consists of CAs that issue digital certificates, directories that store the certificates (including Active Directory in Windows 2000 and Windows Server 2003), and X.509 certificates that are issued to security entities on the network. The PKI provides validation of certificate-based credentials and ensures that the credentials are not revoked, corrupted, or modified.

**Qualified Subordination** The process of configuring cross-certification with basic constraints, name constraints, application constraints, and issuance policy constraints to govern what certificates are trusted from a partner organization.

# Overview (Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Qualified subordination allows an organization to extend its PKI trust hierarchy to other organizations and secondary hierarchies within the same organization. Qualified subordination is also sometimes referred to as cross-certification. While extension of PKI trust was possible in a Windows 2000 network using CTLs, CTLs were limited in how restrictions could be defined between the two organizations. Qualified subordination provides a more flexible and manageable trust mechanism.

When implemented using qualified subordination, each qualified subordinate CA can have rules defined that do the following:

- Define the namespaces for which your PKI hierarchy will issue and accept trusted certificates
- Specify the acceptable uses of certificates issued by a qualified subordinate CA
- Define what issuance practices must be followed for a certificate issued by the qualified subordinate CA to be considered valid
- Create a managed trust between separate certification hierarchies

Windows Server 2003 Enterprise Edition provides the necessary tools to configure qualified subordination between CAs, so that two organizations can define how certificates will be trusted between the organizations. These tools include the following:

- **Version 2 certificate template** Allows CA administrators to modify certificate templates to meet their business purposes. Qualified subordination requires the use of version 2 certificate templates to include policy constraints. Modifications to version 2 templates can include:
    - Creating a new certificate template by duplicating and renaming an existing template
    - Modifying template properties such as certificate validity period, renewal period, cryptographic service provider (CSP), key size, and key archival settings
    - Establishing and applying enrollment policies, issuance policies, and application policies
- **Cross-Certification Authority Certificate** A version 2 certificate template that is issued from one CA to another CA establishing qualified subordination between the two certificates. The certificate used to sign the cross-CA certificate enforces the constraints defined in the certificate.
- **Qualified Subordination Signing Certificate** A version 2 certificate template that must be manually created; it contains the Qualified Subordination application policy object identifier. This certificate template validates that the holder is approved to sign cross-certification authority certificates.
- **Certutil.exe** A command-line program that is installed as part of Certificate Services. It is used to dump and display CA configuration information, configure Certificate Services, back up and restore CA components, and verify certificates, key pairs, and certificate chains.
- **Certreq.exe** A command-line program that is installed as part of Certificate Services. It is used to request certificates from a CA. In qualified subordination, certreq.exe is used to request the cross-certification authority certificate at the issuing CA.

# Understanding Constraints

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The key to qualified subordination is restricting what certificates issued by the qualified subordinate CA, or by CAs that chain through the qualified subordinate CA, are trusted by your organization. This is accomplished by defining basic constraints, name constraints, issuance policy constraints, or application policy constraints in one of two ways:

- Constraints can be defined during the installation of the CA. By including the sections defined below in the CAPolicy.inf file, the basic constraints, name constraints, issuance policy constraints, and application policy constraints are defined at the CA during the installation of the CA or during the certificate renewal process at a CA.

- Constraints can also be defined during the cross-certificate request process by defining the constraints in the Policy.inf file. The Policy.inf file defines the constraints implemented in the cross-certification certificate used to define the qualified subordination.

Appendix A contains an example of a Policy.inf file. Appendix B contains a sample of the CAPolicy.inf file.

| **Note** |
| --- |
| Constraints are defined in section 4.2 "Standard Certificate Extensions" in RFC 2459. |

| **Note** |
| --- |
| For more information about how the certificate chaining engine uses constraint information when building certificate chains, see Certificate Revocation and Status Checking (http://go.microsoft.com/fwlink/?LinkID=27081). |

**Basic Constraints**

Basic constraints allow a CA administrator to limit path length for a certificate chain. You can specify a basic constraint that defines the maximum number of CAs that can exist below the CA where the basic constraint is assigned. For example, if you define a path length of zero, the CA can only issue end-entity certificates and is not permitted to issue SubCA certificates.

Figure 1 shows a CA where a basic constraint is applied that limits the path length to two.



**Figure 1: Applying Basic Constraints**

The basic constraint ensures that the maximum path length for certificates that chain to the CorpCA that are trusted is two deep. This means that only two more layers of CAs can exist below the CA where the basic constraint is defined. In this example, the AsiaCA is restricted so that it can only issue trusted end-entity certificates. If a certificate issued to the JapanCA or any certificate issued by the JapanCA is presented, the chain would fail chain verification, because the certificate would result in a chain with a path length greater than two. Whenever a CA issues a certificate to another CA, it can reduce the path length, but never increase the permitted length.

| **Note** |
| --- |
| Basic constraints are fully defined in section 4.2.1.10 of RFC 2459. |

**Best Practice** A best practice that can be applied in regards to basic constraints is to only restrict the path length in a subordinate CA certificate, not a root CA certificate. The reason this is more commonly employed is that a change in the root requires a complete redeployment of the hierarchy and the root certificate should a change in the path length be desired. Therefore, in the previous example, the RegionCA certificate would be issued with a PathLength = 1 and have no length defined in the root CA certificate.

**Name Constraints**

Name constraints allow you to designate which namespaces are either permitted or excluded for certificates issued by CA where qualified subordination is defined. When a request is received at the qualified CA, the names present in the subject and the subject alternate name fields are compared to the configured name constraints to determine whether the namespace is permitted or excluded. In other words, the CA enforces all constraints defined in the CA certificate.

| **Note** |
| --- |
| If the name presented in the request is not present in the list of constraints, the request will be rejected by the qualified CA. All names in a certificate request must be within permitted namespaces and none may be within an excluded namespace for the certificate to be issued. |

Name constraints allow you to control what namespaces are managed by each CA in your organization and those that you trust from other organizations. When you deploy a qualified subordinate CA, you must carefully consider what namespaces you wish to allow the CA to issue, and more importantly in some cases, what namespaces you wish to prevent the qualified subordinate CA from supporting.

For example, when you configure qualified subordination between your organization and a partner organization, you would typically not want your partners CA infrastructure to issue certificates containing names in your organizations namespace. Remember that a valid certificate that contains a UPN for your organization is mapped to the user account with the UPN as an attribute. It does not matter which CA issues the certificate, only that the certificate chains to a CA certificate in the NTAuth store. The use of name constraints can ensure that your namespace, and all recognized formats of your namespace, are excluded for certificates issued by your partners CA hierarchy.

In the case where both permitted and excluded name constraints exist, the excluded name constraints will always take precedence. For example, if you create a permitted DNS name constraint for the namespace .microsoft.com and an excluded DNS name constraint for .subdomain.microsoft.com, all certificates for subdomain.microsoft.com will be rejected, even though the microsoft.com namespace is permitted.

| **Important** |
| --- |
| The default policy of a Microsoft client validating a certificate with name constraints asserted is that all names have to be explicitly permitted. For example, if a name constraint does not specify the e-mail name as a permitted type, and a certificate request contains an e-mail name, the request will be rejected. It is possible to relax this policy and implicitly allow names not defined in the name constraint extension by configuring the CA policy in the registry. |

The certificate request is only accepted if each name in the certificate request matches at least one of the permitted name constraints configured at the qualified subordinate CA. In other words, if the certificate request contains the requestors name in both an LDAP distinguished name format and in a UPN format, both names must match permitted name constraints. If one of the presented names does not match, the certificate request fails.

**Name Constraint Processing**

When a certificate request is presented to a subordinate CA under qualified subordination rules, all name forms in the certificate request must be within the permitted namespaces. In addition, the subject names must not be within an excluded

namespace. The permitted and excluded name constraints are processed separately.

The following rules are followed by the qualified subordinate CA:

- A certificate request is successful if all names in the certificate request succeed in matching the corresponding permitted name constraints.
- A certificate request is unsuccessful if any names in the certificate request succeed in matching the corresponding excluded name constraint.
- A certificate request containing names that cannot be matched with either the permitted or excluded names is considered excluded and fails.
- Comparison of the name identifying the subject entity with the qualified subordinate CA's name constraints is performed using the following rules:
    - The comparison is not case-sensitive if the name or part of the name is encoded with an IA5 string. If it is encoded as a Unicode or UTF8 string, the comparison will be a binary match of the content.
    - Name constraints are applied to the Subject name field and any existing Subject Alternate Name extensions.
    - A qualified subordinate CA's name constraints cannot permit names excluded by its parent CA. For example, if the parent CA excludes the DNS name .microsoft.com, then its qualified subordinate CA cannot permit the DNS sub-domain .example.microsoft.com. The name constraints are checked by the CA to ensure that its chain is valid by calling CryptoAPI.
    - Name constraints apply to all names contained in subject and subject alternate name extensions in an end certificate. Each name in the subject or subject alternate name extensions must match at least one of the name constraints listed for that name type. A subject name or subject alternate name that does not match a listed name type is rejected.
    - Constraints apply only when the namespace types specified as name constraints are present in the certificate request. If no namespace of the specified types is in the certificate request, the certificate is acceptable. For example, a CA is configured with the following name constraints:

```
[NameConstraintsExtension]
Include = NameConstraintsPermitted
Exclude = NameConstraintsExcluded
Critical = True

[NameConstraintsPermitted]
DirectoryName = "DC=Microsoft, DC=Com"
email = @microsoft.com
UPN = .microsoft.com
UPN = @microsoft.com

[NameConstraintsExcluded]
```

In this configuration, if the CA receives a request with the e-mail name of user1@northwindtraders.com, the request would be rejected. Likewise, the request would be rejected if the subject alternate name contained user1@northwindtraders.com while the subject name contained CN=user1,OU=Contractors,DC=Microsoft,DC=Com. Remember that all subject names must match the name constraints for a certificate to be issued. The default policy as mentioned previously requires all names to be defined unless otherwise configured on the certification authority.

| **Note** |
| --- |
| If you do not specify permitted or excluded names, you are creating a wildcard (*) name constraint for those name constraint formats. |

When the name constraint processing is completed, the validation process will result in one of the following outcomes:

- **Permitted** The end certificate contains a name that is listed as permitted in an issuers name constraints extension.
- **Not permitted** The end certificate contains a name that is not listed as permitted in an issuers name constraints extension.
- **Excluded** The end certificate contains a name that is listed as excluded in an issuers name constraints extension.
- **Not Defined** The issuer certificate does not list a constraint for a specific name type (such as Directory Name or IP Address).

| **Note** |
| --- |
| Windows Server 2003 CAs will have constraints set to permit all namespaces for a particular name type, unless specified otherwise in the name constraints extension in the Policy.inf file. |

**Applying Name Constraints**

The following naming and addressing format are natively supported for use in Policy.inf when defining name constraints:

- Relative distinguished name
- DNS domain name
- Universal Resource Identifiers (URI)
- E-mail name and User principal name (UPN)
- IP address

Additional name constraint forms may be enforced using the Other Name forms, which are identified by name and object identifier that are UTF8 or ASN.1 encoded. Other Names can be used to represent names that are outside of the standard name types such as Rfc822Name, DNS, or X.500 Directory Name. It consists of an object identifier and a binary blob. The most common Other Name is the Universal Principal Name (UPN) Name, which has the constant OID_NT_PRINCIPAL_NAME as the object identifier. The version 2 "Domain E-mail Replication" template certificates and the Version 1 "Domain controller" certificates have a subjectaltname extension that contains the constant OID_NTDS_REPLICATION as other name type and the GUID representation of the DC as the blob. In Windows 2000, only the UPN was present in the other name constraint. However, now in Windows Server 2003, the other name constraint has been expanded to support even the domain controller GUIDs. The following is the syntax where the first part is the object identifier, followed by a qualifier identifying the encoding of the following data as UTF8 or octet or ASN.1:

OtherName=1.2.3.4.99.100,{utf8}ssss

OtherName=1.2.3.4.99.101,{octet}ABCD

OtherName=1.2.3.4.99.102,"{asn}BAgAAQIDBAUGBw=="

OtherName=1.3.6.1.4.1.311.25.1

One scenario where this may be useful is when you have smartcard logon across forests, wherein the other domain has Windows 2000 domain controllers. In this case, you can provide an additional name constraint as "OtherName=1.3.6.1.4.1.311.25.1" to allow all Domain e-mail replication certificates and restrict the namespace using the DNS name constraint. It is technically possible to restrict the GUID of a domain controller as another name but that would not be a generic constraint but a specific constraint to only one domain controller.

**Configuration**

The name constraints are configured in one of two locations. When creating a new CA, you can define name constraints for that CA by configuring CAPolicy.inf to impose name constraints. Likewise, if you are creating a qualified subordinate CA

certificate, you would define name constraints in the Policy.inf file. In both cases, the following syntax is used:

```
[NameConstraintsExtension]
Include = NameConstraintsPermitted
Exclude = NameConstraintsExcluded
Critical = TrUe

[NameConstraintsPermitted]
DNS = ""
email=""
UPN=""

[NameConstraintsExcluded]
DNS = .nwtraders.com
email = @nwtraders.com
UPN = .nwtraders.com
UPN = @nwtraders.com
URI = ftp://.nwtraders.com
DIRECTORYNAME = "DC=NWtraders, DC=com"
```

| **Note** |
| --- |
| The critical = True in the [NameConstraintsExtension] indicates that this extension is marked as critical. If the validating computer cannot parse the extension, it must reject the certificate chain. |

**Relative Distinguished Name Constraints**

Relative distinguished names are used to identify the names of objects stored in directories such as Active Directory or X.500 directories. Using Relative distinguished name constraints allows you to restrict a qualified subordinate CA to issue certificates only to specific users or computers in Active Directory by using the objects Relative distinguished name. The distinguished name can be very specific, where the permitted or excluded Relative distinguished name references a specific object. Or, the distinguished name can be a wildcard value that references all the objects within a specific container, domain, or organizational unit.

The name constraint must be included in either the Policy.inf file used to create a qualified subordinate CA or in the CAPolicy.inf file used when installing a new subordinate CA. Table 1 shows some examples of LDAP name constraints and defines what objects the constraints refer to.

**Table 1 Name Constraint Examples**

| Relative Distinguished Name Constraint | Includes |
| --- | --- |
| `DC=example,DC=microsoft,DC=com` | All objects in the example.microsoft.com domain |
| `CN=Users,DC=example,DC=Microsoft,DC=com` | All objects in the Users container in the example.microsoft.com domain |
| `CN=Brian Komar,CN=Users,DC=microsoft,DC=com` | The user object Brian Komar in the Users container in the microsoft.com domain |
| `OU=Marketing,DC=Microsoft,DC=Com` | All objects in the Marketing OU of the Microsoft.com domain |
| `CN=Servers,CN=Sites,CN=Configuration,DC=Microsoft,DC=Com` | All domain controllers in the microsoft.com forest |
| `OU=Domain Controllers, DC=Redmond,DC=Microsoft,DC=Com` | All domain controllers in the Redmond.microsoft.Com domain. |

The range of characters in the allowed name constraint is defined by the relative distinguished name (also known as RDN) type. Most relative distinguished names (such as O=, OU=,CN=, and so on) support international (UTF8) characters, where relative distinguished names that include (DC=,C=,e-mail, and so on) support IA5 character types. A Relative name constraint will appear in either the permitted or excluded name constraints using the following format:

```
DIRECTORYNAME = "DC=NWtraders, DC=com"
```

**DNS Name Constraints**

DNS name constraints compare any DNS names contained in certificate requests to the list of permitted and excluded DNS name constraints. As with LDAP distinguished name constraints, the name in the certificate request must match one of the permitted names for the certificate to be issued.

The DNS domain names accepted by the subordinate CA must follow the standard DNS naming conventions as specified in RFCs 1034 and 1035, and are listed in either the Policy.inf or CAPolicy.inf file used. Therefore, the DNS name constraint cannot contain international character sets.

When configuring a DNS name constraint, you can either designate a specific DNS host name, for example, host.example.microsoft.com, or you can designate a DNS namespace. For example, the DNS name constraint .example.microsoft.com indicates all hosts that have a DNS name that ends with example.microsoft.com. The preceding period (.) indicates that any DNS name that ends with example.microsoft.com matches the name constraint.

**Note**

This does not include a host that does not have a period preceding example.microsoft.com. For example, the host named ourexample.microsoft.com would not result in a match when compared to the name constraint .example.microsoft.com.

Table 2 demonstrates the results for submitted DNS names against configured DNS name constraints.

**Table 2 Evaluating DNS Name Constraints**

| Submitted DNS Name | DNS Name Constraint | Evaluation |
|---|---|---|
| `www.example.microsoft.com` | `.example.microsoft.com` | Match |
| `www.south.example.microsoft.com` | `.example.microsoft.com` | Match |
| `example.microsoft.com` | `.example.microsoft.com` | No Match |
| `Ourexample.microsoft.com` | `example.microsoft.com` | No Match |
| `www.example.microsoft.com` | `example.microsoft.com` | Match |

**Important**

Evaluate your DNS constraints carefully. Remember that all excluded name constraints take precedence over a permitted name constraint. For example, if you permitted .example.microsoft.com and excluded .microsoft.com, a request from www.example.microsoft.com will be rejected. This DNS name matches both the excluded and permitted name constraints. Because there is a match on the excluded name constraint, the permitted name constraint is not observed.

A DNS name constraint will appear in either the permitted or excluded name constraints using the following format:

```
DNS = .nwtraders.com
```

**Uniform Resource Identifier (URI) Constraints**

Uniform Resource Identifiers (URIs) identify resources on the Internet that use identifiers such as URL, FTP, HTTP, telnet, mailto, news, and gopher. The URI naming conventions supported for name constraints must follow the syntax specified in RFC 2396.

When a qualified subordinate CA validates a URI name, the protocol element in the URI is ignored. For example, if the submitted URI is http://www.example.microsoft.com and the URI name constraint is URL=ftp://.example.microsoft.com, the actual comparison performed is between the host name www.example.microsoft.com and the name constraint .example.microsoft.com. In this case, the result would be a match, even though the protocols differ.

As with DNS constraints, a URI constraint can reference either a host or a domain. If a preceding period is included in the URI constraint, the domain name must include the entire suffix defined in the name constraint. For example, if the URI constraint is .example.microsoft.com, then matches are achieved for west.example.microsoft.com and for east.example.microsoft.com.

The most frequent application of URI name constraints is the validation of certificate requests for Web Server certificates, where the Web servers URL is submitted in the certificate request. For example, if the Web Server certificate request has a name of http://www.microsoft.com and the qualified subordinate CA's permitted URI name constraint is URL=http://www.microsoft.com, then this URI constraint would permit the certificate request for the Web Server certificate.

A URI name constraint will appear in either the permitted or excluded name constraints using the following format:

```
URL = "http://.microsoft.com"
```

**RFC 822, E-mail, and UPN Constraints**

RFC 822 and e-mail constraints must follow RFC 822 naming conventions and use IA5 encoding of character strings. User Principal Names (UPNs) use the same syntax but are encoded using UTF8. Although RFC 822 names, UPNs, and E-mail addresses share the same syntax, you must provide separate examples in the name constraints listing of Policy.inf to differentiate between requests using e-mail addresses and requests using UPN names.

When defining RFC 822, E-mail, or UPN constraints, you can specify constraints for individual addresses or for addresses that end with a specific domain.

For example, to create a name constraint for the e-mail address for a Joe Smith at Northwind Traders, you could create a E-mail address constraint for jsmith@nwtraders.com. To create a name constraint for all e-mail users at Northwind Traders, you would set the e-mail constraint to be @nwtraders.com. By not placing any text to the left of the ampersand, the e-mail constraint becomes a wildcard representing all e-mail addresses in the nwtraders.com domain.

Table 3 shows some examples of RFC822, E-mail, and UPN constraints and how they evaluate.

**Table 3 RFC 822, E-mail, and UPN Constraint Evaluation**

| Submitted Address | UPN/E-Mail Constraint | Evaluation |
|---|---|---|
| jsmith@nwtraders.com | jsmith@nwtraders.com | Match |
| bjsmith@nwtraders.com | jsmith@nwtraders.com | Match |
| jsmith@nwtraders.com | @nwtraders.com | Match |
| jbsmith@nwtraders.com | bjsmith@nwtraders.com | No Match |

**Note**

UPN name constraints should always be entered as two separate entries in the list of name constraints. One entry should include the ampersand character (such as @nwtraders.com); the second entry should replace the ampersand with a period (.), so this entry would be .nwtraders.com. This format allows for the possibility of having a UPN of subdomain.nwtraders.msft.

E-mail and UPN name constraints will appear in either the permitted or excluded name constraints using the following format:

```
email = @nwtraders.com
UPN = .nwtraders.com
UPN = @nwtraders.com
```

**IP Address Constraints**

IP address constraints allow you to specify either specific IP addresses or ranges of IP addresses that a CA can successfully receive certificates from a qualified subordinate CA. The format of the IP addresses in the constraints must follow either RFC 791 (for IPv4 addresses) or RFC 2460.

When a certificate request is received by a qualified subordinate CA, the source IP address in the certificate request is compared to the IP address constraints to determine if the IP address is permitted or excluded by the name constraints.

To specify a range of IP addresses, you must specify both the IP address and associated subnet mask in each constraint. For example, if you wanted to create an IP address constraint for the 192.168.3.0 network, the constraint would be 192.168.3.0/255.255.255.0.

An IP address constraint will appear in either the permitted or excluded name constraints using the following format:

```
IPADDRESS = 192.168.3.0/255.255.255.0
IPADDRESS = 192.168.2.254/255.255.255.255
IPADDRESS = 172.16.8.0/255.255.248.0
```

The first constraint includes all computers in the 192.168.3.0 network. The second constraint is a host-specific constraint that only includes the host with IP address 192.168.2.254. The final constraint is a Classless Internet Domain Routing (CIDR) example that includes all hosts with addresses between 172.16.8.1 and 172.16.15.254.

**Issuance Policy**

Issuance policies are used to identify the extent to which your organization trusts the identity presented in a certificate issued

by another organizations CA hierarchy. For example, an issuance policy may describe that you only trust certificates that were issued during a face-to-face meeting with a network administrator, such as the issuance of a smartcard certificate. Each issuance policy is described by an object identifier. The inclusion of an issuance policy object identifier in an issued certificate indicates that the certificate was issued meeting the issuance requirements associated with the issuance policy object identifier.

| **Note** |
| --- |
| Issuance policy is the Microsoft term for Certificate Policies defined in part 4.2.1.5 in RFC 2459 and is contained in the certificatePolicies extension of a certificate. |

For a specific certificate template, you can define one or more issuance policy object identifiers that are included in issued certificates. There is a special issuance policy, the *All issuance Policies* object identifier, commonly reserved for CA certificates, that indicates the issuance policy contains all other issuance policies.

Windows Server 2003 includes four predefined issuance policies:

- **All Issuance (2.5.29.32.0)** The all issuance Policy indicates that the issuance policy contains all other issuance policies. Typically, this object identifier is only assigned to CA certificates.

- **Low Assurance (1.3.6.1.4.1.311.21.8.x.y.z.1.400)** The low assurance object identifier is used to represent certificates that are issued with no additional security requirements. The *x.y.z* portion of the object identifier is a randomly generated numeric sequence that is unique for each Windows Server 2003 forest.

- **Medium Assurance (1.3.6.1.4.1.311.21.8.x.y.z.1.401)** The medium assurance object identifier is used to represent certificates that may have additional security requirements for issuance. For example, a smartcard certificate that is issued in a face-to-face meeting with a smartcard issuer could be considered a medium assurance certificate and contain the medium assurance object identifier. The *x.y.z* portion of the object identifier is a randomly generated numeric sequence that is unique for each Windows Server 2003 forest.

- **High Assurance (1.3.6.1.4.1.311.21.8.x.y.z.1.402)** The high assurance object identifier is used to represent certificates that are issued with the utmost security. For example, a key recovery agent certificate allows the holder to recover private key material from a Windows Server 2003 Enterprise Edition CA. The issuance of a key recovery agent certificate may require additional background checks and a digital signature from a designated approver. The *x.y.z* portion of the object identifier is a randomly generated numeric sequence that is unique for each Windows Server 2003 forest.

- **Best Practice** Technically, Issuance Policies can be added to any version 2 template that is issued from a Windows Server 2003 certification authority running Enterprise Edition. However, it is a best practice to apply issuance policies in the Policy.inf file for the qualified subordinate CA to ensure consistent policy for the CA. This is to avoid issuing templates with different policies unintentionally.

| **Note** |
| --- |
| A root CA automatically implies that all issuance policies are defined. For understanding how policy is evaluated in certificate chains, see the Certificate Status and Revocation white paper at http://www.microsoft.com/technet/security/topics/crypto/tshtcrl.mspx |

If your organization has existing object identifiers issued, they may be used for this purpose. In addition, you can create your own object identifiers to represent custom issuance policies. For example, two organizations involved in a purchaser/seller relationship may define custom object identifiers to represent digital signature certificates for specific purchase amounts. For example, object identifiers may be defined for purchase between $100,000 and $500,000, and another object identifier for purchases greater than $500,000. These object identifiers could then be used by applications to recognize whether a person had the appropriate signing authority for a specific volume purchase.

```
[PolicyStatementExtension]
Policies = HighAssurancePolicy, MediumAssurancePolicy, LowAssurancePolicy
CRITICAL = FALSE

[HighAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.402

[MediumAssurancePolicy]
```

```
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.401

[LowAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.400
```

**Note**

Issuance policy extensions are only recognized by Windows XP and Windows Server 2003 clients. If the extension is marked critical, CAPI passes the extensions to the application. It is up to the calling application to understand the extension when it calls with a specific policy provider. The behavior will vary by application. In the case of Outlook on Windows 2000, a critical issuance policy will cause Outlook to generate an error as it does not understand the extension. But on Windows XP, the issuance policy is recognized and does not cause an error.

**Policy Mapping**

When qualified subordination is configured between two CAs that use issuance policy, the object identifiers between the two organizations must be mapped. Policy mapping establishes what object identifiers in one organization are equivalent to object identifiers defined in a different organization's CA hierarchy. The policy mapping allows the object identifiers generated in certificates in one organization to be recognized in another organization through the creation of cross-certification certificates.

Policy mapping ensures that only authorized object identifiers from a partner organization are allowed in certificates issued by the partner organization. The policy mapping will associate the partner organization's object identifier with an object identifier defined in your organization's PKI hierarchy.

Figure 2 shows how policy mapping is defined in the Policy.inf file when cross certificates are issued to allow certificates to be used between the two CA hierarchies.




**Figure 2: Policy Mapping**

In this figure, the Northwind Traders is configured with an issuance policy named MillionDollar that has been assigned the object identifier 1.3.6.4.1.311.21.8.1.124. This issuance policy is mapped to the issuance policy named BigOrder configured at the Contoso Consulting CA.

If a certificate is issued by the Northwind Traders CA with an issuance policy other than the MillionDollar policy, the certificate is rejected by the Contoso CA because the policy does not map to an approved issuance policy object identifier.

When cross certificates are created to link the two organizations' CA hierarchies, the code snippets shown in the figure must be included in the Policy.inf files used to generate the cross-certification certificate at each CA. The Policy.inf files exist at the *issuer* CA. The object identifier that is mapped is the object identifier from the *subject* CA that is the object identifier for the matching issuance policy at the partner organization.

Before you perform the policy mapping, you must configure the issuer CA with the issuance policies and associated object IDs. This is done by configuring the CAPolicy.inf file and either installing the CA or renewing the CA certificate using the configured CAPolicy.inf file with the required object identifiers as shown in the following code sample:

```
[PolicyStatementExtension]
Policies = HighAssurancePolicy, MediumAssurancePolicy, LowAssurancePolicy
CRITICAL = FALSE

[HighAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.402

[MediumAssurancePolicy]
```

```
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.401

[LowAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.400
```

In this example, the High, Medium, and Low Assurance issuance policies are defined at the subject CA. The object identifiers can be retrieved from the Certificate Templates console by right-clicking Certificate Templates in the console tree and selecting View Object Identifiers from the pop-up menu.

| **Note** |
| --- |
| Appendix B contains a sample CAPolicy.inf file. |

The renewal of the CA certificate using the CAPolicy.inf file ensures that the issuance policies are now defined at the CA. At this point, the cross certificates can be generated that perform the policy mapping between the two organizations. The following process is used to perform the qualified subordination between the two CA hierarchies:

1. Determine the organization with which you want to establish a trust relationship. If more than two organizations exist in the relationship, the trust relationships must be defined in pairs.

2. Through discussions with the partner organization, establish the equivalence between the assurance levels used by the two organizations involved in the trust relationship.

3. Exchange the issuance policy object identifiers defined for each assurance level used in both organizations involved in the trust relationship.

4. Map the issuance policy object identifiers in the separate organizations in the Policy.inf file at each CA and define their policy constraints in the CA certificate request for the qualified subordinate CA you are installing for your organization.

5. Install the qualified subordinate CA with the policies, policy mappings, and policy constraints you wish to deploy in your organization.

When a user or computer in the organization where the *issuing* CA exists receives a file signed by a user in the organization where the *subject* CA exists, the certificate will be chained to the root CA in the *issuing* CA using the qualified subordination path. If the object identifier for the issuance policy in the certificate maps to the required issuance policy object ID in your organization, the certificate chain is considered valid. If the user or computer in the issuing CAs organization receives a certificate with a non-mapped object identifier, the certificate chain is assigned a lesser value by the validation process and is typically rejected by the application.

**Constraining Policy Mapping to Prevent Unexpected Trust**

You can further constrain issuance policy by defining parameters that define how the issuance policy defined in the qualified subordination affects other CAs below the qualified subordinate CA. There are two settings that can define the relationship:

- **Require explicit policy** Specifies the number of certificates that can exist in the hierarchy below the certificate with the policy constraint extension. For example, if the explicit policy is configured with a value of three, this means that the defined issuance policy must exist for three layers of the hierarchy with the CA, where the qualified subordination is defined being the first level.

- **Inhibit policy mapping** Specifies the number of additional certificates that can appear in the path before policy mapping is no longer permitted. An inhibit policy mapping value of three only restricts the policy mapping to three levels of CAs below the qualified subordinate CA.

These settings exist to prevent unexpected trust relationships. For example, Figure 3 shows the same qualified subordination configured previously between Northwind Traders and Contoso Consulting with the addition of a new CA at Fabrikam Corporation.




**Figure 3: Preventing Unexpected Trust**

If qualified subordination is configured between the Contoso Consulting CA and the Fabrikam Corp CA, mapping the BigOrder issuance policy to the SpecialSig issuance policy, it is possible that a certificate issued by the Fabrikam Corp CA could be accepted by the Northwind Traders CA. This is because the object identifier for SpecialSig would be mapped first to the BigOrder object identifier. Then, the BigOrder object identifier would be mapped to the MillionDollar object identifier.

This behavior is prevented by configuring the following Policy Constraint Extensions at the Northwind Traders CA.

```
[PolicyConstraintsExtension]
RequireExplicitPolicy = 1
InhibitPolicyMapping = 1
```

By assigning InhibitPolicyMapping a value of one, only certificates issued by the ContosoCA are recognized through the MillionDollar object identifier mapping defined at the Northwind Traders CA. A certificate issued by the Fabrikam Corp CA with the SpecialSig object identifier will fail the InhibitPolicyMapping =1 setting defined at the Northwind Traders CA. This prevents the certificate issued by the Fabrikam Corp CA from being accepted by the Northwind Traders CA.

**Policy Qualifiers**

You can provide additional information about the Issuance policies implemented at a CA by configuring policy qualifiers. Policy qualifiers provide information directly, or provide links to information, that describe the purpose of the issuance policy. The following code shows both formats of the policy qualifier:

```
[LegalPolicy]
OID = 1.3.6.1.4.1.311.21.43
Notice = "Legal policy statement text"
URL = "http://www.example.microsoft.com/policy/isspolicy.asp"
```

When viewing the certificate in an application, a person can follow the URL and read the description of the issuance policy in a Web browser. If a text message is used as a policy qualifier, the application will display the message.

**Best Practice:** It is recommended that URLs be implemented for certificate practice statements, rather than to use notice text. This provides several key benefits.

- It allows modification of the practice statement without having to re-issue the CA certificate.
- URLs do not dramatically increase the size of the certificate, unlike notice text.

**Application Policy**

Application policies included in certificates allow applications to determine if a certificate can be used when authenticating a user, encrypting data, or signing a device driver. An application can be coded to only accept certificates that contain specific application policies. When the application receives signed information from a user, the application will review the certificate associated with the private key used to sign the information and verify that the certificate chain has the required object identifier as a valid application policy.

Application policies are similar to the Enhanced Key Usage (EKU) extension in a certificate as both use one or more object

identifiers to prescribe how the public key in a certificate must be used. EKU is still supported by Windows XP and the Windows Server 2003 family to provide for PKIs that use this extension, but application policies are a replacement. The following rules are used when combining EKU and Application Policy extensions:

- If a certificate is presented that contains an EKU extension, but does not contain a separate Application Policy extension, the EKU extension will be treated as an Application Policy extension by Windows XP and Windows Server 2003 clients.
- If a certificate has an extension containing an application policy and also has an EKU extension, the EKU extension is ignored.
- If a certificate has an Application Policy extension and an EKU property, the effective policy for the certificate is the common policy between the EKU property object identifiers and the application policy object identifiers.

| **Note** |
| --- |
| If issuing certificates that include both Application Policy and EKU extensions, ensure that the two extensions are identical in their assignment of object identifiers so that the two extensions do not conflict with each other. This ensures consistent application of policy, when either Application Policy or EKU extensions are used. This is especially important in issuing cross certificates or in using templates that do not have predefined application policies. |

Application policies may be defined in qualified subordination certificates to limit what purposes certificates from a trusted organization can be used for. For example, you can create an application policy that only allows certificates with the Secure Mail object identifier to be recognized if they chain through the qualified subordinate CA.

When application policy is defined for a CA, the object identifiers associated with the application policy are included in all issued certificates.

| **Note** |
| --- |
| The All Applications object identifier indicates that the application policy includes all application policies. This application policy is normally reserved for certificates issued to CAs. For understanding how policy is evaluated in certificate chains, see the Certificate Status and Revocation white paper at http://www.microsoft.com/technet/security/topics/crypto/tshtcrl.mspx |

The following Application Policy object identifiers are recognized by Windows Server 2003.

| | |
| --- | --- |
| • Server Authentication (1.3.6.1.5.5.7.3.1) | • Client Authentication (1.3.6.1.5.5.7.3.2) |
| • Code Signing (1.3.6.1.5.5.7.3.3) | • Secure E-mail (1.3.6.1.5.5.7.3.4) |
| • IP Security End System (1.3.6.1.5.5.7.3.5) | • IP Security Tunnel Endpoint (1.3.6.1.5.5.7.3.6) |
| • IP Security User (1.3.6.1.5.5.7.3.7) | • Time Stamping (1.3.6.1.5.5.7.3.8) |
| • IP Security IKE Intermediate (1.3.6.1.5.5.8.2.2) | • All Application Policies (1.3.6.1.4.1.311.10.12.1) |
| • Microsoft Trust List Signing (1.3.6.1.4.1.311.10.3.1) | • Qualified Subordination (1.3.6.1.4.1.311.10.3.10) |
| • Key Recovery (1.3.6.1.4.1.311.10.3.11) | • Document Signing (1.3.6.1.4.1.311.10.3.12) |

**Note**

A list of available object identifiers can be viewed in the Certificate Templates console by right-clicking Certificate Templates in the console tree and selecting View Object Identifiers from the pop-up menu. Alternatively, custom application policy object identifiers can be defined and used in custom-developed applications.

To define application policies in a Policy.inf file, the following sections must be created:

```
[ApplicationPolicyStatementExtension]
Policies = AppEmailPolicy, AppCodeSignPolicy, AppAuthPolicy
CRITICAL = FALSE

[AppEmailPolicy]
OID = 1.3.6.1.5.5.7.3.4 ; Secure Email

[AppCodeSignPolicy]
OID = 11.3.6.1.5.5.7.3.3 ; Code Signing

[AppClAuthPolicy]
OID = 1.3.6.1.5.5.7.3.2 ; Client Authentication
```

The [ApplicationPolicyStatementExtension] section defines all application policy setting sections that exist in the Policy.inf file. In this case, three application policy sections are defined. For each [*AppPolicy*] section, the object identifier associated with the application policy is defined.

If a custom application policy object identifier is defined, an [ApplicationPolicyMappingsExtension] section must be defined. This section uses the same format, where the local object identifier is mapped to the object identifier used by the other organization participating in the qualified subordination as shown in the following code sample:

```
[ApplicationPolicyMappingsExtension]
1.3.6.1.4.1.311.21.64 = 1.2.3.4.98
1.3.6.1.4.1.311.21.65 = 1.2.3.4.100
critical = true
```

As with issuance policies, you can also configure limitations on how deep in the hierarchy explicit policy must exist and whether the defined application policies can be mapped to other application policies. These settings are defined in the [ApplicationPolicyConstraintsExtension] section as follows:

```
[ApplicationPolicyConstraintsExtension]
; consists of two optional DWORDs
; They refer to the depth of the CA hierarchy that requires explicit policy
; and inhibits Policy Mapping
RequireExplicitPolicy  = 6
InhibitPolicyMapping = 10
```

# Qualified Subordination Deployment Scenarios

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

This section looks at some of the more common certificate services deployments that use qualified subordination, including the following scenarios:

- Using qualified subordination to restrict certificate issuance to specific namespaces
- Configuring qualified subordination between two organizations; specifically two configurations:
  - Allowing trust between all CAs in the two organizations
  - Limiting trust to specific CAs in the CA hierarchy
- Configuring qualified subordination using a bridge CA

This section covers both the design of the scenarios and how to deploy the scenarios.

| **Note** |
| --- |
| This section looks at the design process of qualified subordination. The Walkthrough section provides the actual deployment steps involved when defining qualified subordination between two CA hierarchies. |

**Restricting Namespaces**

The first example is a common scenario in large organizations where certificate services administration is decentralized. In a decentralized environment, it may be desired to restrict certificate issuance to only specific objects in Active Directory. Use of name constraints is an effective way to ensure specific CAs can only issue to a well-defined subset of users framed by a namespace.

Assume that Contoso uses the forest structure in Figure 4 to provide decentralization of administration to regional subsidiaries.




**Figure 4: Contoso.coms Forest and Domain Structure**

Within Contoso.com, several CAs have been deployed for certificate services deployments. Figure 5 shows the current CA hierarchy used by Contoso.com.




**Figure 5: A Proposed CA Hierarchy for Contoso Corporation**

In this example, the CA named RegionCA exists to allow decentralized management of certificate deployment for Europe and Asia. The EuropeCA and AsiaCA are installed on member servers in the europe.contoso.com and asia.contoso.com domains. Qualified subordination can be used to restrict the issuance of certificates by the AsiaCA to only security principals (users, computers, services) located in the asia.contoso.com domain.

To do this, the AsiaCA must be installed with a CAPolicy.inf file in the %SystemRoot% folder with the following settings:

```
[NameConstraintsExtension]
Include = NameConstraintsPermitted
Exclude = NameConstraintsExcluded
Critical = True

[NameConstraintsPermitted]
DNS = .asia.contoso.com
email = @asia. contoso.com
UPN = .asia.contoso.com
UPN = @asia.contoso.com
DIRECTORYNAME = "DC=asia,DC=contoso,DC=com

[NameConstraintsExcluded]
```

These name constraints will only allow security principals from the asia.contoso.com domain to receive certificates issued by the AsiaCA certification authority. If a request is received from a security principal from a different namespace, the certificate request will be denied due to the configured name constraints.

| **Note** |
| --- |
| Alternatively, as mentioned earlier when defining the types of constraints, the AsiaCA can obtain a subordinate CA certificate from the RegionCA that includes the defined name constraints. Either method would provide the necessary name constraints to issued certificates. |

For example, Figure 6 shows the name constraints in place for AsiaCA restricting certificate issuance to the asia.contoso.com namespace.




**Figure 6: Applying Name Constraints at AsiaCA**

In this example, the certificate request by user3@asia.contoso.com is approved because Noels UPN is included in the permitted name constraints. In comparison, the request by user5@europe.contoso.com is denied because the UPN Europe.contoso.com is not listed in the permitted name constraints and is considered to be an excluded name constraint.

Enterprise CAs enforce name constraints during certificate requests based on the authenticated security principal performing the certificate request. The credentials of the security principal are used to extract the names of the user from Active Directory and these names are compared with the defined name constraints. In this case, user3 would be logged on to the network and performing the successful certificate request.

**Creating Trust Between Two Organizations**

The second scenario where cross-certification may be used is to allow certificates to be used and trusted between two organizations. Before the cross-certification is performed between the two organizations, the first thing that must be decided is what certificate usage will be trusted between the organizations. For example, you could require that certificates be trusted from the other organization(s) only for secure e-mail. Or, you could allow secure e-mail, client authentication, and server authentication.

This must be decided and included in the Policy.inf file when the qualified subordination is configured between CAs in the two organizations.

| **Note** |
| --- |
| For more information about how to perform the qualified subordination, see the Walkthrough section. |

When configuring the cross-certification between the two organizations, you also must decide whether all CAs in the other organizations CA hierarchy are to be trusted or only a specific grouping of CAs in the hierarchy.

| **Important** |
| --- |
| Cross-certification of a CA hierarchy may create very long certificate chains that can increase the size of network traffic for IPSEC, Secure Mail (s/mime), SSL, and so on. Use caution so as to not increase the chain length and overall certificate chain byte size to exceed an application limit. |

**Implementing Qualified Subordination to Trust All CAs**

In the first example, qualified subordination will be configured so that all CAs in the two organizations are trusted by the other organization. Figure 7 shows the two CA hierarchies.




**Figure 7: Two Non-Connected CA Hierarchies**

In this scenario, there is a business need for user 1 and user 2 to exchange secure e-mail using Outlook and S/MIME. Due to increased cooperation between Northwind Traders and Contoso Consulting, it is decided that qualified subordination be used to permit secure e-mail certificates issued by either organization to be trusted by both Northwind Traders and Contoso Consulting. In addition, the secure e-mail certificates can be issued by any CA in the individual CA hierarchies.

To accomplish this, two different methods can be used:

- Cross-certification between root Cas

  In this design, the root CAs issue cross-certification certificates to the other hierarchy's root CA. This is the easiest method to cross-certify organizations and is the easiest to understand when troubleshooting qualified subordination problems.

- Cross-certification between subordinate and Root Cas

  Many organizations are unwilling to issue certificates from their root CA to other organization's root CA. In this case, the cross-certification certificates can be issued by a subordinate CA to the root CA in the partner organization's CA hierarchy.

Both methods allow for complete trust between the CA hierarchies, subject to the constraints defined in the Policy.inf files. The decision factors will be based on whether your organization is willing to issue a cross-certification certificate to a partner organization's root CA from your root CA. Another consideration is that Root CAs generally publish their CRL less frequently than subordinate CAs. If there is a possibility that the cross-certification authority certificate may be revoked, it is preferable to issue the certificate from a subordinate CA than from a root CA. A certificate issued from a subordinate CA will have less lag time between the time the certificate is revoked and the revocation being recognized by all clients.

**Cross-Certification Between Root CAs**

If your organization is willing to perform cross-certification between the root CAs, the configuration in Figure 8 can be used.




**Figure 8: Configuring Complete Trust Between Root CAs**

In this example, the root CA for each organization is cross-certified with the root CA in the other organization. To create this configuration, the qualified subordination must be performed twice:

1. The PartnerCA at Contoso Consulting must be cross-certified using qualified subordination with the CorpCA at Northwind Traders.

2. The CorpCA at Northwind Traders must be cross-certified using qualified subordination with the PartnerCA at Contoso Consulting.

The qualified subordination will require that the following be in place:

- A Policy.inf is configured at each root CA. The Policy.inf must define any name constraints, issuance constraints, and application constraints defined between the two organizations and is used during the cross-CA certificate request as described in the Walkthrough section.

- The user account that is performing the qualified subordination must obtain a certificate with the Qualified Subordination application policy object identifier to sign the qualified subordination request between the two organizations. The signing certificate must be obtained from the CorpCA or any CA that is trusted on the CorpCA and PartnerCA computers for the first request and from the PartnerCA or any CA that is trusted on the CorpCA and PartnerCA computers for the second request.

The configuration shown in Figure 8 allows certificates issued by either CA to chain to a root that is trusted by each respective organization. For example, the certificate issued to User1 will validate up to the CorpCA if verified by a security principal at Northwind Traders as follows:




If the same certificate is validated by a computer in the Contoso Consulting organization, the certificate will validate up to PartnerCA.




Similarly, the certificate issued to User2 will chain to a trusted root in each organization. If the certificate is validated by a computer at Northwind Traders, the certificate chain will be as follows:




However, the same certificate would validate differently and use the PartnerCA root CA as the trust anchor when validated by a computer in the Contoso Consulting organization:

| **Note** |
| --- |
| In all cross-certification designs, a single certificate will chain to different roots depending on which organization's computers are validating the certificate. In both examples, a certificate validated by computers at Contoso Consulting will chain to the PartnerCA root, while certificates validated by computers at Northwind Traders will chain to the CorpCA root. |

The advantage of this design is that it also allows trust of new CAs installed into the CA hierarchies at either organization. For example, Figure 9 shows the addition of a new CA at Northwind Traders.




**Figure 9: Adding an Additional CA to the Hierarchy**

In this case, the certificate issued to User4 is trusted in both organizations. In Northwind Traders, the certificate chains to the CorpCA root CA.




While at Contoso Consulting, the certificate chains to the trusted PartnerCA.



**Cross-Certification Between Subordinate and Root CAs**

The same level of trust can be established between two organizations by issuing the cross-certification authority certificate from a subordinate CA in your CA hierarchy. This strategy is recommended if your organization's security policy prohibits the issuance of crossCA certificates by the root CA, your organization is simply unwilling to issue cross-certification certificates at the root CA, or there is a possibility that the cross-certification authority certificate may be revoked, and the CRL publication period for the root CA is a long interval. In these cases, the configuration shown in Figure 10 can be used to create complete trust between the two CA hierarchies, subject to defined qualified subordination constraints.




**Figure 10: Configuring Complete Trust Between Subordinate and Root CAs**

In this example, the Root CA for each organization is cross-certified with the subordinate enterprise CA in the other organization. To create this configuration, the qualified subordination must be performed twice:

1. The PartnerCA at Contoso Consulting must be cross-certified using qualified subordination with the IssuingCA at Northwind Traders.

2. The CorpCA at Northwind Traders must be cross-certified using qualified subordination with the PolicyCA at Contoso Consulting.

The qualified subordination will require that the following be in place:

- A Policy.inf is configured at each subordinate CA. The Policy.inf must define any name constraints, issuance constraints, and application constraints defined between the two organizations and is used during the cross-CA certificate request as described in the Walkthrough section.

- The user account that is performing the qualified subordination must obtain a certificate with the Qualified Subordination application policy OBJECT IDENTIFIER to sign the qualified subordination request between the two organizations. The signing certificate must be obtained from the IssuingCA or any CA that is trusted on the CorpCA and PartnerCA computers for the first request and from the PolicyCA or any CA that is trusted on the CorpCA and PartnerCA computers for the second request.

The following configuration allows certificates issued by either CA to chain to a root that is trusted by each respective organization. For example, the certificate issued to User1 will validate up to the CorpCA if verified by a security principal at Northwind Traders.




If the same certificate is validated by a computer in the Contoso Consulting organization, the certificate will validate up to PartnerCA.

Similarly, the certificate issued to User2 will chain to a trusted root in each organization. If the certificate is validated by a computer at Northwind Traders, the certificate chain will be as follows:

While, the same certificate would validate up to the PartnerCA root CA when validated by a computer in the Contoso Consulting organization.

**Note**

In all cross-certification designs, a single certificate will chain to different roots depending on which organization's computers are validating the certificate. In both examples, a certificate validated by computers at Contoso Consulting will chain to the PartnerCA root, while certificates validated by computers at Northwind Traders will chain to the CorpCA root.

As with the model where the cross-certification is established between root CAs, this design also allows trust of new CAs installed into the CA hierarchies at either organization. Figure 11 shows the addition of a new CA at Northwind Traders.




**Figure 11: Adding an Additional CA to the Hierarchy**

In this case, the certificate issued to User4 is trusted in both organizations. In Northwind Traders, the certificate chains to the CorpCA root CA.

While at Contoso Consulting, the certificate chains to the trusted PartnerCA.

**Implementing Qualified Subordination to Limit Trust to Specific CAs**

In some situations, two organizations may wish to restrict the cross-certification to specific CAs or to specific portions of the CA hierarchy, for example, where an organization uses outside consulting services. The organization (Northwind Traders) and consulting firm (Contoso) may wish to restrict the cross-certification to specific CAs in the hierarchy.

Figure 12 shows the scenario where the qualified subordination is performed to limit the trust between specific CAs in the organization.




**Figure 12: Configuring a Limited Trust Relationship**

With this configuration, the certificates issued by the IssuingCA and the PolicyCA will validate up to a trusted root CA in both organizations.

At Northwind Traders, the certificate issued to user1 will chain as follows:

The certificate issued to user2 will chain as follows:




At Contoso Consulting, the certificates will chain with the PartnerCA being the trusted Root CA. For the user1 certificate, the following chain will be selected by a validation process:




For the user2 certificate, the validation process will also select the chain that chains to the PartnerCA root CA.

The major difference between this qualified subordination design and the design that allowed total trust between CAs in the two hierarchies is best shown by investigating an additional CA added to the Northwind Traders CA hierarchy as shown in Figure 13.




**Figure 13: Adding an Additional CA to the Hierarchy**

In this case, the ProjectCA was added to the Northwind Traders CA hierarchy. To a computer in the Northwind Traders network, the certificate issued to user4 is valid because it chains to a trusted root CA, the CorpCA.

When a computer from the Contoso Consulting organization is presented the user4 certificate, only a partial chain can be built and the chain does not terminate at a trusted root certificate.

**Types of Constraints**

Qualified subordination allows constraints to be applied to a cross-certification configuration. Rather than implicitly trusting all certificates issued by another organizations CA hierarchy, qualified subordination allows you to define application, naming, and issuance constraints.

For example, when qualified subordination is established between two organizations, some of the constraints they may wish to configure include:

- **Application policy** Rather than allowing all certificates to be trusted, specific application purposes can be defined for the issued crossCA certificate by defining what object identifiers are required in certificates from the other organization. For example, if an organization wishes to only accept S/MIME certificates from a partner organization, an application policy can be defined to only accept certificates with the Secure E-mail object identifier (1.3.6.1.5.5.7.3.4) in the application policy extension or in the EKU extension.

| **Note** |
| --- |
| Enhanced Key Usage is defined in part 4.2.1.13 in RFC 2459. |

- **Name Constraints** Whenever qualified subordination is used between two organizations, the organization issuing the CrossCA certificate should include a name constraint that excludes its namespace through the cross-certification. This prevents the partner organization from submitting a certificate using the other organizations namespace. For example, Northwind Traders would implement a name constraint that prevents Contoso from presenting certificates that are issued to a user from Northwind Traders. In other words, if a CA in the Contoso hierarchy issued a certificate with a subject name of user1@northwindtraders.com, the name constraint would exclude the nortwindtraders.com e-mail address and would invalidate the certificate.

| **Note** |
| --- |
| Name Constraints are defined in part 4.2.1.11 in RFC 2459. |

- **Policy Constraints** To ensure a level of trust for partner certificates, policy constraints (sometimes referred to as issuance policy) can be defined and must be included as attributes of presented certificates to be accepted by the partner organization. For example, if two partner organizations are involved in a buyer/seller relationship, a Million Dollar object ID can be defined. For purchases of 1,000,000 USD or more, the signing certificate must contain the Million Dollar object identifier. Because the object identifier must be defined in both organizations, the Policy.inf file must map the object identifier from one organization to the other organizations object identifier. Likewise, constraints can be applied to whether the object identifier can be mapped to other object identifiers, or whether the object identifier must be issued from a specific CA or limit how many layers below the cross-certified CA can issue a certificate containing the object identifier.

| **Note** |
| --- |
| Policy Constraints is defined in part 4.2.1.5 in RFC 2459 and are referred to as the Certificate Policies extension in the RFC. |

| **Important** |
| --- |

If different combinations of policy and name constraints are required for application policies, you may have to perform the qualified subordination multiple times. For example, if you are willing to accept any Secure E-mail certificates, but will only accept High and Medium Assurance object identifiers for Client and Server authentication certificates, two separate Policy.inf files must be defined and the qualified subordination process must be executed twice—once for each Policy.inf file. This results in two separate Cross-Certification Authority certificates being issued—one for Secure E-mail certificates and one for Client and Server Authentication certificates. The cross-certification certificates are ultimately trusted for different purposes. Applications will verify the policy object identifiers when making authentication or authorization decisions.

**Creating Trust Between Multiple Organizations**

When you require trust between more than two, three, or more organizations, it is sometimes easier to design the qualified subordination by using a Bridge CA. The Bridge CA will act as a link between the CA hierarchies in each organization. Each organization will validate certificates issued by other organizations that use chains that include the Bridge CA. The Bridge CA design also reduces the complexity involved when new organizations are included in the trust.

A Bridge CA has several advantages over configuring separate trust relationships between multiple organizations, including

- A Bridge CA reduces the complexity for defining trust between CA hierarchies when three or more CA hierarchies exist.
- It is easier to add an additional organization to an existing Bridge CA trust relationship. The only tasks that must be performed are as follows:
    1. Issue the Bridge CA a cross-certification certificate using qualified subordination from each organization.
    2. Issue a cross-certification certificate from the Bridge CA for the root CA of the new organization.
    3. Publish the Bridge CA issued certificate in all other organizations to complete the web of trust created by the BridgeCA.
- The Bridge CA can offer a method to manage trust between subsidiary organizations by a holding company.

**Important**

Even if a BridgeCA is implemented, there is still nothing preventing two organizations participating in the Bridge CA structure from using qualified subordination to define a separate relationship between their hierarchies. For example, while Secure Mail certificates may be validated through the Bridge CA for e-mail sent between Northwind Traders and Contoso Consulting, a separate qualified subordination may be established between CAs in the two organizations to define specific policy constraints that must exist in certificates for signing purchase orders for transactions greater than 1 million dollars between the two organizations. Depending on the certificates attributes, the validation process will select the appropriate certificate chain based on the certificates usage and the defined qualified subordinations between the two organizations.

A common design for the Bridge CA deployment is shown in Figure 14.




**Figure 14: Using a Bridge CA**

In this design, total trust is established between all CAs in the three organizations, subject to any name, issuance, and application policies defined in the qualified subordination.

To achieve this, qualified subordination is defined so that the Bridge CA receives a subCA certificate defining all application, naming, and policy constraints from the subordinate enterprise CA in each organization. In addition, the root CA in each organization receives a subCA certificate defining all application, naming, and policy constraints from the Bridge CA.

| **Note** |
| --- |
| Alternatively, the cross-certification can be performed between the root CA of each organization and the BridgeCA. Again, the decision will be based on the security policy of each organization using the bridge CA, whether the organizations are willing to issue a cross-certification certificate from their root CA to the bridge CA, and the CRL publication period of the root CAs. |

| **Important** |
| --- |
| Cross-certification of a CA hierarchy may create very long certificate chains that can increase the size of network traffic for IPSEC, Secure Mail (s/mime), SSL, and so on. Use caution so as not to increase the chain length and overall certificate chain byte size to exceed an application limit. |

**Viewing Certificates from the Three Organizations**

The result is complete trust between the organizations using the BridgeCA for the certificate purposes defined in the qualified subordination. As long as an issued certificate meets all defined application, name, and policy constraints, the certificate will be considered valid in any organization connected to the Bridge CA in this manner.

In this example, the user1 certificate would chain to the CorpCA trusted root for all computers at Northwind Traders as follows:

The user1 certificate would chain to the OrgCA trusted root for all computers at Fabrikam Inc. as follows:




If the same certificate is evaluated by the certificate chaining engine at a computer in the Contoso organization, the resulting certificate chain would be as follows:

Likewise, the certificates for user6 and user2 would chain to the trusted root for each of the three organizations, based on what organization the computer evaluating the certificate belongs to.

**Adding New CAs to Existing CA Hierarchies**

This design for Bridge CA deployment allows for the addition of new CAs to the hierarchies at any of the participating organizations, without having to modify the existing qualified subordination using the Bridge CA. For example, if the Project CA is added to the Northwind Traders CA hierarchy as shown in Figure 15, the certificate issued to user4 would be trusted by

computers in the Northwind Traders, Fabrikam Inc., and Contoso Consulting organizations, subject to application, naming, and policy constraints.




**Figure 15: Adding a New CA to the Northwind Traders CA Hierarchy**

In this case, the user4 certificate would chain to the CorpCA trusted root for the computers in the Northwind Traders organization.

The certificate would chain to the OrgCA trusted root for the computers in the Fabrikam Inc. organization.




Finally, the certificate would chain to the PartnerCA trusted root for the computers in the Contoso Consulting organization.




**Using Name Constraints with Bridge CAs**

To prevent CAs in the other organization from creating certificates with subject names from your organizations namespace, name constraints should be created to exclude your namespace from certificates trusted through the BridgeCA.

For example, for the cross-certification authority certificate issued to BridgeCA from the CorpCA, the following name constraints could be added to the Policy.inf file at the CorpCA:

```
[NameConstraintsExtension]
Include = NameConstraintsPermitted
Exclude = NameConstraintsExcluded
Critical = TrUe

[NameConstraintsPermitted]
DNS = ""
email=""
UPN=""

 [NameConstraintsExcluded]
DNS = .northwindtraders.com
email = @northwindtraders.com
UPN = .northwindtraders.com
UPN = @northwindtraders.com
DIRECTORYNAME = "DC=Northwindtraders,DC=com"
```

These name constraints will ensure that all namespaces except the NorthwindTraders namespace will be permitted.

| **Note** |
| --- |
| Additional name constraints could be defined for NameConstraintsPermitted to restrict what specific namespaces are allowed into the organization through the BridgeCA. |

**Policy Constraints and Policy Constraint Mapping**

The qualified subordination between the organizations may also require specific policy constraints to exist for trusted certificates. If a specified policy constraint is not included in a certificate, the application should not select the chain including that certificate.

The problem with policy constraints is that for each CA hierarchy, the policy constraint may have a different name and it will have a different object identifier. To make the policy constraints interoperate between the CA hierarchies, the policy constraints must be defined and mapped in the Policy.inf file.

For example, if the BridgeCA will function to only allow Low and Medium Assurance certificates to be passed between the organizations, the policy constraints must be defined in each cross-certification certificate issued between CAs.

For the cross-certification authority certificate issued to CorpCA by the BridgeCA, the Policy.inf file must contain the following lines to define the medium and low assurance object identifiers. These object identifiers may be obtained from the Certificate Templates console or from the IANA:

```
[PolicyStatementExtension]
Policies = MediumAssurance, LowAssurance
CRITICAL = FALSE

[MediumAssurance]
OID = 1.3.6.1.4.1.311.21.8.1608391590.1259233725.355412102.3300744578.1.401

[LowAssurance]
OID = 1.3.6.1.4.1.311.21.8.1608391590.1259233725.355412102.3300744578.1.400
```

In addition to defining the object identifiers for its CA hierarchy, the Policy.inf file must also map the Medium and Low Assurance object identifiers to the object identifiers used at the BridgeCA. This is accomplished through the [PolicyMappingsExtension] section of the Policy.inf file:

```
[PolicyMappingsExtension]

1.3.6.1.4.1.311.21.8.1608391590.1259233725.355412102.3300744578.1.401 =
1.3.6.1.4.1.311.21.8.2812900009.1258984576.2698790533.44255976.1990074901.2398624335

1.3.6.1.4.1.311.21.8.1608391590.1259233725.355412102.3300744578.1.400 =
1.3.6.1.4.1.311.21.8.2812900009.1258984576.2698790533.44255976.572751058.2135922791

critical = yEs
```

This section maps the object identifier ID for medium assurance to an object ID defined at the BridgeCA. The object identifier does not have to be for the MediumAssurance policy constraint. In fact, it can be for a user-defined policy object identifier, as the previous example shows.

To complete the policy mappings, the following lines must be included in the Policy.inf file used by the BridgeCA to receive a cross-certification authority certificate from the IssuingCA:

```
[PolicyStatementExtension]
Policies = SecLevel1, SecLevel2
CRITICAL = FALSE

[SecLevel1]
OID =
1.3.6.1.4.1.311.21.8.2812900009.1258984576.2698790533.44255976.1990074901.239862
433

[SecLevel2]
OID =
1.3.6.1.4.1.311.21.8.2812900009.1258984576.2698790533.44255976.572751058.2135922
791

[PolicyMappingsExtension]
1.3.6.1.4.1.311.21.8.2812900009.1258984576.2698790533.44255976.1990074901.239862
433 = 1.3.6.1.4.1.311.21.8.1608391590.1259233725.355412102.3300744578.1.401

1.3.6.1.4.1.311.21.8.2812900009.1258984576.2698790533.44255976.572751058.2135922
791 = 1.3.6.1.4.1.311.21.8.1608391590.1259233725.355412102.3300744578.1.400

critical = yEs
```

As you can see, custom object identifiers were defined for the policies named SecLevel1 and SevLevel2. In the

[PolicyMappingsExtension] section, the same object identifiers are mapped with the order reversed from the Policy.inf file used by the CorpCA to receive a cross-certification authority certificate from the BridgeCA.

Figure 16 shows another example of policy constraint mapping.




**Figure 16: Policy Constraint Mapping**

In this figure, the NorthwindCA is configured with a MillionDollar object identifier, 1.3.6.1.4.1.311.21.8.1.124. This object identifier is mapped to the BigOrder object identifier, 1.3.6.4.1.204.22.33.44, at the ContosoCA. Although the certificate issued to v-user1 contains a different issuance policy object identifier from the certificate issued to user2, the policy constraint mapping makes the object identifiers equivalent and the object identifiers are recognized by both CA hierarchies.

**Application Policy**

The final policy that could be applied in this scenario is the application policy. Application policy will limit which purposes certificates issued by another organizations CAs can be used for. By defining application policy, you can prevent certificates for specific usages from being accepted for use in your organization.

For example, if you only want to allow certificates for S/MIME, client authentication, and server authentication to be recognized through the BridgeCA, you can add the following sections to the Policy.inf file:

```
[ApplicationPolicyStatementExtension]
; list of user defined policies
Policies = AppOIDPolicy1, AppOIDPolicy2, AppOIDPolicy3

[AppOIDPolicy1]
OID = 1.3.6.1.5.5.7.3.4 ; Secure Email

[AppOIDPolicy2]
OID = 1.3.6.1.5.5.7.3.1 ; Server Authentication

[AppOIDPolicy3]
OID = 1.3.6.1.5.5.7.3.2 ; Client Authentication

CRITICAL = FALSE
```

These application policies would prevent the use of an EFS certificate between the organizations. An EFS certificate would not chain in a trusted manner because the defined application policies define an application policy for EFS encryption services.

**Note**

Remember that these application policies also are combined with name constraints and policy constraints when applied to a certificate presented from another CA hierarchy. If you require different combinations, you must request multiple cross-certification authority certificates using separate Policy.inf files for each request. For example, you may not wish to implement policy constraints for e-mail certificates, but want to ensure that low or medium assurance object identifiers exist in client and server authentication certificates. This would require two separate cross-certification authority certificate requests—one for Secure E-mail and one for Client and Server Authentication.

# Walkthrough (Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

This section provides detailed implementation steps involved in defining qualified subordination between two CA hierarchies.

**Creating an Offline Bridge CA**

The following process can be used to deploy an offline root CA for your organization. The following is required before you establish your offline Bridge CA:

- Configure a CAPolicy.inf file to define the following settings for the offline Bridge CA.
    - Authority Information Access locations
    - CRL Distribution Point locations
    - Certificate Practice Statement
    - All issuance policies required to cross the Bridge CA
- Determine the CAs distinguished name.
- Determine what Cryptographic Service Provider (CSP) will be used for certificate creation and signing.

**Configuring the CAPolicy.inf File**

The Offline Bridge CA must include the following sections in a CAPolicy.inf file stored in the %Systemroot% folder. This example shows an addition of three policy settings for the Offline Bridge CA. The OIDs for the three policies were obtained from the organization's Active Directory. Alternatively, the OIDs could be obtained from the IANA at http://www.iana.org/cgi-bin/enterprise.pl.

```
[Version]
Signature= "$Windows NT$"

 [certsrv_server]
CRLPeriod = weeks
CRLPeriodUnits = 26

[PolicyStatementExtension]
Policies = HighAssurancePolicy, MediumAssurancePolicy, LowAssurancePolicy
CRITICAL = FALSE

[HighAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.402

[MediumAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.401

[LowAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.400

[PolicyConstraintsExtension]
RequireExplicitPolicy = 0
InhibitPolicyMapping = 0

[BasicConstraintsExtension]
PathLength = 0
```

**Note**

R As mentioned previously, it is not a best practice to place policy OIDs in root CA certificates. It is also not a best practice to place CDP or AIA extensions in the root CA certificate as most applications do not check for root CA revocation.

**Installing Certificate Services**

The actual installation process of a root CA is performed using the following steps:

1. Ensure that the updated CAPolicy.inf file is saved to the %systemroot% folder.

   **Note**

   Appendix B contains an example of a CAPolicy.inf file.

2. On the Start menu, click Settings, and then click Control Panel.

3. On the Control Panel, double-click Add or Remove Programs.

4. In the Add or Remove Programs window, click Add/Remove Windows Components.

5. In the Windows Components Wizard, on the Windows Components page, select the Certificate Services check box.

6. In the Microsoft Certificate Services dialog box, click Yes to accept that you cannot rename the computer or change its domain membership.

7. On the Windows Components page, click Next.

8. In the CA Type page (Figure 17), click Stand-alone root CA, select the Use custom settings to generate the key pair and CA certificate check box, and then click Next.




**Figure 17: Defining the CA Type Configuration**

9. On the Public and Private Key Pair page (Figure 18), select the CSP and the Key length (for the private and public key pair), and then click Next.




**Figure 18: Defining the Public and Private Key Pair Properties**

10. On the CA Identifying Information page (Figure 19), enter the Common name and Validity period (for the Root CA certificate), and then click Next.




**Figure 19: Defining the Common Name of the CA and Setting the Validity Period**

| **Note** |
| --- |
| The Cryptographic Key is now generated based on the provided information in the preceding dialog boxes. Depending on the length of key defined on the Public and Private Key Pair page, this may take some time. |

11. On the Certificate Database Settings page (Figure 20), accept the default locations, and then click Next.




**Figure 20: Defining the Database and Log Storage Locations**

12. In the Microsoft Certificate Services dialog box, click Yes to stop the Internet Information Services.

    The Configuring Components dialog box appears as the installation proceeds.

13. If the Files Needed dialog box appears (Figure 21), in the Copy files from box, enter the CD-ROM drive path or the network share path where the distribution files are located, and then click OK.




**Figure 21: Providing the Location of the Installation Files**

14. In the Windows Components Wizard, click Finish.

15. In the Add or Remove Programs dialog box, click Close.

16. Close the Control Panel.

**Disabling Delta CRLs at the Root CA**

Typically, the root CA for a CA hierarchy is removed from the network for security purposes. Periodically, the CA must be accessed for the publication of updated CRLs. The goal of delta CRLs is to publish the newly revoked certificates at a more frequent interval than the publication period for the base CRL. By disabling delta CRLs, you reduce the number of times that the offline root CA must be accessed for CRL publication. It is typically not an issue that the delta CRLs are not published, as the number of certificates issued by the root CA is minimal.

| **Important** |
| --- |
| If cross-certification is performed between root CAs, it is possible that there may be a delay between the revocation of a former partner's cross-certification certificate and the recognition of the revocation by the clients accessing the base CRL. |

1. On the Start menu, point to Programs, point to Administrative Tools, and then click Certification Authority.

2. In the Certification Authority console, in the console tree, expand *CAName* (where *CAName* is the name of your offline Bridge CA).

3. In the console tree, right-click Revoked Certificates, and then click Properties.

4. In the Revoked Certificates Properties dialog box, click to clear the Publish Delta CRLs check box, and then click OK.

5. In the console tree, right-click Revoked Certificates, click All Tasks, and then click Publish to publish a new base CRL that does not reference a Delta CRL.

6. In the Publish CRL dialog box, click OK.

**Removing the Non-Required AIA and CDP Extensions**

Once the offline Bridge CA is installed, the CDP and AIA locations for issued certificates must be modified to ensure that the certificate chaining engine can find the CRL and certificate for the offline Bridge CA. Because the CA is removed from the network, these default AIA and CDP extensions do not reference valid locations for the publication points. By removing the offline root CA from the network, you effectively remove the default publication points.

To remove non-required CDP and AIA extensions

1. In the Certification Authority console tree, right-click CAName, and then click Properties.

2. In the CAName Properties dialog box, click the Extensions tab.

3. In the Select extension box, ensure that the box contains CRL Distribution Point (CDP).

4. In the list of CDP locations, select the CDP location beginning with ldap:///, and then click Remove.

5. In the Confirm removal dialog box, click Yes to verify the removal.

6. In the list of CDP locations, select the CDP location beginning with http://, and then click Remove.

7. In the Confirm removal dialog box, click Yes to verify the removal.

8. In the list of CDP locations, select the CDP location beginning with file://\\, and then click Remove.

9. In the Confirm removal dialog box, click Yes to verify the removal.

| **Important** |
| --- |
| Do not delete the CDP location referencing c:\windows\System32. This location is required by the CA to locally publish update Certificate Revocation lists. |

10. On the Extensions tab, select Authority Information Access in the Select extension box.

11. In the list of AIA locations, select the AIA location beginning with ldap:///, and then click Remove.

12. In the Confirm removal dialog box, click Yes to verify the removal.

13. In the list of AIA locations, select the AIA location beginning with http://, and then click Remove.

14. In the Confirm removal dialog box, click Yes to verify the removal.

15. In the list of AIA locations, select the AIA location beginning with file://\\, and then click Remove.

16. In the Confirm removal dialog box, click Yes to verify the removal.

| **Important** |
| --- |
| Do not delete the AIA location referencing c:\windows\System32. This location is required by the CA to locally publish renewed Certificates. |

17. Click OK.

18. In the Certification Authority dialog box, click Yes to restart Certificate Services.

**Adding Required AIA and CDP Extensions**

The offline Bridge CA must add required CDP and AIA extensions to allow Active Directory clients and other clients to find the CAs certificate and CRL. You must modify the extensions to include paths that are available when the offline Bridge CA is removed from the network.

The examples provided in this white paper are not mandatory, but should be considered when defining the AIA and CDP extensions for the offline root CA.

| **Note** |
| --- |
| The AIA and CDP extensions are required by the certificate chaining engine to build certificate chains for validation. For more information about the certificate validation process, see Certificate Revocation and Status Checking (http://go.microsoft.com/fwlink/?LinkID=27081). |

**Adding the LDAP CDP Extension**

The first CDP extension that needs to be added is the LDAP:// CDP extension. This CDP extension allows Active Directory and other LDAP clients to acquire the CRL from an LDAP service available to the Internet.

| **Note** |
| --- |
| In some Windows 2000 and Windows Server 2003 deployments, a separate Active Directory forest is deployed in the extranet to prevent external access to the corporate Active Directory. If a separate LDAP service is used in the extranet, the LDAP URL must reference this Active Directory service. Otherwise, the LDAP URL must reference a link that is accessible by all clients that require access to the CRL. |

To add the CRL to Active Directory, you will need the following information from the forest root domain of Active Directory that will maintain a copy of the offline Root CAs CRL:

```
ForestRootDomain   The LDAP representation of the forest root domain. For example
security.nwtraders com would be represented as DC=security,DC=Nwtraders,DC=com.
```

You will also require some information from the offline Bridge CA. Specifically, you require the following:

```
CAName:    The name assigned to the CA during installation of the CA.
CAMachineName:    The NetBIOS name of the computer hosting the CA.
```

You will need to add the two following LDAP paths to the CDP extensions list:

```
ldap:///CN=CAName,CN=CAMachineName,CN=CDP,CN=Public Key Services,
CN=Services,CN=Configuration,CAForestName?certificateRevocationList
?base?objectClass=cRLDistributionPoint

and

ldap://ldap.company.com/CN=CAName,CN=CAMachineName,CN=CDP,CN=Public Key Services,
CN=Services,CN=Configuration,CAForestName

where ldap.company.com is an URL that refers queries to TCP port 389 to the
server hosting the LDAP service, typically a Windows 2000 domain controller.
```

For example, for the following information:

- *CAName*: BridgeCA
- *CAMachineName*: BridgeComp
- *ForestRootDomain*: extranet.nwtraders.com

The following are the LDAP paths you must add to the list of CDP extensions:

```
ldap:///CN=BridgeCA,CN=BridgeComp,CN=CDP,CN=Public Key Services,
CN=Services,CN=Configuration,DC=extranet,DC=nwtraders,DC=com?certificateRevocati
onList?base?objectClass=cRLDistributionPoint
```

and

```
ldap://ldap.extranet.nwtraders.com/CN=BridgeCA,CN=BridgeComp,CN=CDP,CN=Public
Key Services,CN=Services,CN=Configuration,DC=extranet,DC=nwtraders,DC=com
```

| **Note** |
| --- |
| The LDAP:/// URL is a local LDAP path that binds to the current domain controller used by the client computer. This LDAP URL will not work with external clients, therefore, the second LDAP URL is added to provide the URL to the LDAP server for external clients. |

1. In the Certification Authority console tree, right-click CAName, and then click Properties.
2. In the CAName Properties dialog box, click the Extensions tab.
3. Ensure that the Select Extensions list contains CRL Distribution Points (CDP), and then click Add.
4. In the Add Location dialog box (Figure 22), in the Location box, type the LDAP path as described previously, and then click OK.




**Figure 22: Adding an LDAP CDP Extension**

| **Note** |
| --- |
| Consider creating the path in a text file beforehand and then copying and pasting the complete LDAP URL into the Location box. There is no control to allow editing of the LDAP path if it is incorrectly typed in the console. |

5. On the Extensions tab, select the LDAP CDP extension. Select the Include in all CRLS. Specifics where to publish in the Active Directory when publishing manually and Include in the CDP extension of issued certificates check boxes as shown

in Figure 23. The first option is recommended when publishing offline CA CRLs to Active Directory. The second option ensures that all issued certificates include the updated paths to the CRL.




**Figure 23: Configuring the Properties of the LDAP Extension**

```
Repeat the process using the LDAP URL
ldap://ldap.company.com/CN=CAName,CN=CAMachineName,CN=CDP,CN=Public Key Services,
CN=Services,CN=Configuration,CAForestName
```

6. Leave the dialog box open to allow the addition of an HTTP CDP extension.

| **Note** |
| --- |
| It is recommended to include an HTTP CDP extension to ensure that non-LDAP clients can access the offline CA's CRL. It also is easier to deploy the CRL to external clients using the HTTP protocol than the LDAP protocol. |

**Adding the HTTP CDP Extension**

Additionally, at least one CDP extension should be added for the HTTP protocol that references an available Web server hosting the CRL file for the offline CA. You should provide more than one HTTP location to provide fault tolerance in the event that the Web server referenced in the CDP extension is unavailable. The URL must reference a valid HTTP location to work correctly.

To add an HTTP CDP extension

1. In the CAName Properties dialog box, click the Extensions tab.

2. Ensure that the Select Extensions list contains CRL Distribution Points (CDP), and then click Add.

3. In the Add Location dialog box (Figure 24), in the Location box, type the HTTP path as described previously, and then click OK.




**Figure 24: Adding an HTTP CDP Extension**

| **Note** |
| --- |
| Consider creating the path in a text file beforehand and then copying and pasting the complete HTTP URL into the Location box. There is no control to allow editing of the HTTP path if it is incorrectly typed in the console. |

4. On the Extensions tab, select the HTTP CDP extension, and then select the Include in the CDP extension of issued certificates check box as shown in Figure 25.




**Figure 25: Configuring the Properties of the HTTP Extension**

5. In the CAName Properties dialog box, click OK.

6. In the Certification Authority dialog box, click Yes to restart Certificate Services.

**Adding an HTTP AIA Extension**

In addition to CDP extensions, an AIA extension should be added to allow Web clients to access the offline Bridge CAs certificate from a Web server.

To add an updated HTTP AIA extension

1. In the Certification Authority console tree, right-click CAName, and then click Properties.

2. In the CAName Properties dialog box, click the Extensions tab.

3. Ensure that the Select Extensions list contains Authority Information Access (AIA), and then click Add.

4. In the Add Location dialog box (Figure 26), in the Location box, type the HTTP path to the Root CA certificate, and then click OK.




**Figure 26: Adding an HTTP AIA Extension**

| **Note** |
| --- |
| Consider creating the path in a text file beforehand and then copying and pasting the complete HTTP URL into the Location box. There is no control to allow editing of the HTTP path if it is incorrectly typed in the console. |

5. On the Extensions tab, select the HTTP AIA extension, and then select the Include in the AIA extension of issued certificates check boxes as shown in Figure 27.




**Figure 27: Configuring the Properties of the LDAP Extension**

6. In the CAName Properties dialog box, click OK.

7. In the Certification Authority dialog box, click Yes to restart Certificate Services.

**Publishing CRLs and Certificates**
Once the CDP and AIA extensions are modified, you must publish the updated CRL referencing the new CDP locations. The previous version of the CRL still contains the default locations for the CDP and AIA extensions.

To publish the new base CRL

1. In the console tree, right-click Revoked Certificates, click All Tasks, and then click Publish to publish a new base CRL.

2. In the Publish CRL dialog box, click OK.

Once the CRL is updated, the new version must be published to the locations referenced in the CDP and AIA extensions. This requires manually copying the CRL and Certificate to the Web servers and injecting the CRL and certificate into Active Directory.

**Publishing the CRL and CA Certificate to a Web Server**
The CRL and CA Certificate can be retrieved from the \\CAMachineName\CertEnroll share and copied to a disk or to a network location. The CRL and certificate must be published to the location referenced in the CDP and AIA extensions of the CA.

For example, if the CDP extension references http://www.microsoft.com/public/rootca.crl, the CRL must be copied to this location so that a Web client can retrieve the latest version of the CRL.

Likewise, if the AIA extension references http://www.microsoft.com/public/rootca.crt, the CAs certificate file must be copied to this location so that a Web client can retrieve the CAs certificate.

By default, the IIS default Web site is stored in \inetpub\wwwroot in the same drive where the Windows Server 2003 operating system is installed. To create the paths referenced above, you can create a Public subfolder and copy the two files into the Public folder (\inetpub\wwwroot\public).

Alternatively, you can create a virtual folder and have /Public reference any folder on the Web servers disk.

Once the CRL and certificate are published, you should validate that the paths are correct in the issued certificates by connecting to the CRL and certificate using the CDP and AIA URLs.

**Publishing the CRL and CA Certificate into Active Directory**
The Certutil.exe utility is used to publish the CAs certificate and CRL into the Configuration naming context of Active Directory. The CA's certificate is published to two locations:

- CN=*CAName*,CN=Certification Authorities,CN=Public Key Services,CN=Services,

  CN=Configuration,DC=*ForestRootDomainDN*

- CN=*CAName*,CN=AIA,CN=Public Key Services,CN=Services,

  CN=Configuration, DC=*ForestRootDomainDN*

The CA's CRL is published to the following location in the Configuration naming context:

- CN=*CAName*,CN=*CAComputerName*,CN=CDP,CN=Public Key Services,CN=Services
  CN=Configuration,DC=*ForestRootDomainDN*

The Certutil.exe utility is only installed at a CA by default, but can be installed at any Windows XP client by installing the Windows Server 2003 Administration Pack (adminpak.msi).

Assuming that the CAs certificate is named RootCA.crt and that the CAs CRL is named RootCA.crl, and that the files are stored in the current directory, the following commands can be run to inject the CAs certificate and the CAs CRL into Active Directory.

To install the CAs CRL into Active Directory, use the following command:

```
certutil.exe -dspublish -f RootCA.crl
```

If the command is successful, you should see the following output in response to the command:

```
ldap:///CN=RootCA,CN=RackCluster3,CN=CDP,CN=Public Key Services,CN=Services
,CN=Configuration,DC=bkforest2,DC=nttest,DC=microsoft,DC=com?certificateRevocati
onList?base?objectClass=cRLDistributionPoint?certificateRevocationList

Base CRL added to DS store.

CertUtil: -dsPublish command completed successfully.
```

To install the CAs certificate into Active Directory, a variation of the Certutil –dspublish command is used:

```
Certutil –dspublish –f RootCA.crt RootCA
```

If successful, the command should produce the following output:

```
ldap:///CN=RootCA,CN=Certification Authorities,CN=Public Key Services,CN=Se
rvices,CN=Configuration,DC=bkforest2,DC=nttest,DC=microsoft,DC=com?cACertificate

Certificate added to DS store.

ldap:///CN=RootCA,CN=AIA,CN=Public Key Services,CN=Services,CN=Configuratio
n,DC=bkforest2,DC=nttest,DC=microsoft,DC=com?cACertificate

Certificate added to DS store.

CertUtil: -dsPublish command completed successfully.
```

| **Note** |
| --- |
| The certificate is published to both the Certification Authorities and AIA locations under Public Key Services. |

**Installing the Enterprise Subordinate CAs**

The enterprise subordinate CAs must be installed with the issuance policies defined for the CA certificate. The following processes describe the installation of the Enterprise Subordinate CA.

**Configuring the CAPolicy.inf File**

The Enterprise Subordinate CA requires that a CAPolicy.inf file be created and installed in the %systemroot% folder to apply the correct policies. If the Enterprise CA already exists, you can still use CAPolicy.inf and renew the CAs certificate to apply the modified settings. The CAPolicy.inf file must contain the following sections:

```
[RequestAttributes]
CertificateTemplate = SubCA

[PolicyStatementExtension]
Policies = HighAssurancePolicy, MediumAssurancePolicy, LowAssurancePolicy
CRITICAL = FALSE

[HighAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.402

[MediumAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.401

[LowAssurancePolicy]
OID = 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1.400
```

| **Note** |
| --- |
| The OIDS for the assurance policies must be modified to match the OIDS for the forest where the Enterprise Subordinate CAs exist. Each forest will generate a unique OID that is used for the high, medium, and low assurance issuance policies. In this example, the unique component of each OID is 1.3.6.1.4.1.311.21.8.2473717464.1095930238.502626717.506190032.1, You must change the OIDs shown in this example to match the unique OID representing your forest. |

To modify the CAPolicy.inf file

1. Copy the modified CaPolicy.inf file to the %systemroot% folder.

2. Open the CaPolicy.inf file using Notepad.

3. Click Start, and then click Run. In the Open box, type **certtmpl.msc**.

4. In the console tree, right-click Certificate Templates, and then click View Object Identifiers.

5. In the View Object Identifiers dialog box, select High Assurance, and then click Copy Object Identifier.

6. Paste the updated object identifier for High Assurance into the CaPolicy.inf file.

7. In the View Object Identifiers dialog box, select Medium Assurance, and then click Copy Object Identifier.

8. Paste the updated object identifier for Medium Assurance into the CaPolicy.inf file.

9. In the View Object Identifiers dialog box, select Low Assurance, and then click Copy Object Identifier.

10. Paste the updated object identifier for Low Assurance into the CaPolicy.inf file.

11. In the View Object Identifiers, click Close.

12. Close the Certificate Templates console.

**Installing Certificate Services**

Once the CAPolicy.inf file is configured, the installation of the Enterprise CA can begin. The process is made up of three steps: the initial installation of the enterprise subordinate CA, the offline certificate request process, and the installation of the SubCA certificate.

**Installing the Enterprise Subordinate CA**

To perform the initial installation of certificate services on the enterprise subordinate CA

1. On the Start menu, click Settings, and then click Control Panel.

2. In the Control Panel, double-click Add or Remove Programs.

3. In the Add or Remove Programs window, click Add/Remove Windows Components.

4. In the Windows Components Wizard, on the Windows Components page, select the Certificate Services check box.

5. In the Microsoft Certificate Services dialog box, click Yes to accept that you cannot rename the computer or change its domain membership.

6. In the Windows Components page, click Next.

7. In the CA Type page, click Enterprise Subordinate CA, select the Use custom settings to generate the key pair and CA certificate check box, and then click Next.

8. On the Public and Private Key Pair page, define the CSP, Hash algorithm, and Key length the CA will use for the private and public key pair (Figure 28), and then click Next.




**Figure 28: Defining the Public and Private Key Settings**

9. On the CA Identifying Information page (Figure 29), enter the Common name for this CA, and then click Next.




**Figure 29: Defining the CA Name**

10. On the Certificate Database Settings page (Figure 30), accept the default settings, and then click Next.




**Figure 30: Defining the Certificate Database and Log Locations**

11. On the CA Certificate Request page (Figure 31), select the Save the request to a file option button, accept the default Request file name, and then click Next.




**Figure 31: Configuring an Offline Certificate Request**

| **Note** |
| --- |
| This assumes that the root CA or parent CA is offline and unable to receive online requests. |

12. In the Microsoft Certificate Services dialog box, click Yes.

13. If the Files Needed dialog box appears, enter the path to the installation files, and then click OK.

| **Note** |
| --- |
| A dialog box appears indicating that the installation is incomplete and that the request file must be submitted to the parent CA. |

14. In the Microsoft Certificate Services dialog box, click OK.

15. In the Windows Components Wizard, click Finish.

**Signing the Offline Certificate Request**

To sign the offline certificate request (These steps are performed at the offline Root CA or the offline parent CA.)

1. Copy the request file to a disk.

2. Take the disk to the offline root CA.

3. Open the request file using Notepad. You should see content similar to the following example:


```
 -----BEGIN NEW CERTIFICATE REQUEST-----
MIICJTCCAY4CAQAwPjETMBEGCgmSJomT8ixkARkWA2NvbTETMBEGCgmSJomT8ixk
ARkWA2FiYzESMBAGA1UEAxMJSXNzdWluZ0NBMIGfMA0GCSqGSIb3DQEBAQUAA4GN
ADCBiQKBgQDD7j/MtDoqG0ZWdGSkF7h+taDOD5fB8JhTqIgx31maN+YQE288n7Vm
xtHH7a6Mo+hTRyNBr9gut3ZD4+CNETE+ek3SAsqu/7yXKPzlURlrniKWSAQ9kseO
9llLNFVAWwE8dwxR/taqMAKW1hxflub7p7qnL95eqLzzLfzPfqHwoQIDAQABoIGm
MCkGCisGAQQBgjcNAgMxGxYZNS4xLjM1OTAuMi5TZXJ2aWNlIFBhY2sgMTB5Bgkq
hkiG9w0BCQ4xbDBqMBAGCSsGAQQBgjcVAQQDAgEAMB0GA1UdDgQWBBS64o3oWvft
Zy8bMazXc1SZsyCx2jAZBgkrBgEEAYI3FAIEDB4KAFMAdQBiAEMAQTALBgNVHQ8E
BAMCAYYwDwYDVR0TAQH/BAUwAwEB/zANBgkqhkiG9w0BAQUFAAOBgQAl15B50lwN
AfsSgFKTuPELSalkjWmnn11ZPAmGCLHnhZ6yhQwonWntN3nRaU7F1+KfEnvoibb2
DM1x7SLoMEzQrWQ8sWneoBSCtD0Sdg24dIpWwxlnKgsImRlfGnlEQEd/VHTgyxSh
hrS1gQXYQYH0CT8giCZ8PwGsg1qr/8dMIg==
-----END NEW CERTIFICATE REQUEST-----
```


4. Copy the entire contents of the request file to the Windows Clipboard.

5. Start Internet Explorer.

6. Open the http://localhost/certsrv URL.

7. On the Welcome page, click the Request a certificate link.

8. On the Request a Certificate page, click advanced certificate request.

9. In the Advanced Certificate Request page, click Submit a certificate request by using a base-64-encoded CMC or PKCS#10 file, or submit a renewal request by using a base-64-encoded PKCS#7 file.

10. In the Submit a Certificate Request or Renewal Request page, paste the contents of the request file in the Saved Request box, and then click Submit.

11. The Certificate Pending page appears informing you that an administrator must issue the certificate. You should note the request ID for future reference.

12. Close Internet Explorer.

To issue the certificate, you must use the Certification Authority Console.

1. On the Start menu, point to Programs, point to Administrative Tools, and then click Certification Authority.

2. In the console tree, expand CAName (where *CAName* is the name of your Certification Authority), and then click the Pending Requests container.

3. In the details pane, right-click the pending request, identified by its request ID, for the previously submitted request, click All Tasks, and then click Issue.

4. Close the Certification Authority console.

To retrieve the issued certificate (These steps are performed at the offline root CA or offline parent CA.)

1. Start Internet Explorer.
2. Open the http://localhost/certsrv URL.
3. On the Welcome page, click the View the status of a pending certificate request link.
4. On the View the Status of a Pending Certificate Request page, click the Saved-Request Certificate (Date and Time) link.
5. On the Certificate Issued page, click the Download certificate chain link.
6. In the File Download dialog box, click Save.
7. In the Save As dialog box, in the file name box, type **a:\certnew.p7b**, and then click Save to save the PKCS#7 file to the disk.
8. If the Download complete dialog box appears, click Close.
9. Close Internet Explorer.

**Installing the Certificate at the Enterprise Subordinate CA**

The PKCS#7 file must now be installed at the enterprise subordinate CA using the following steps:

1. Insert the disk containing the PKCS#7 file in the disk drive.
2. On the Start menu, point to Programs, point to Administrative Tools, and then click Certification Authority.
3. In the console tree, right-click CAName (where *CAName* is the name of your enterprise subordinate CA), point to All Tasks, and then click Install CA Certificate.
4. In the Select file to complete CA Installation dialog box, in the File name box, type **a:\certnew.p7b**, and then click Open.
5. In the console tree, right-click CAName, point to All Tasks, and then click Start Service.

| **Note** |
| --- |
| It may be required to also install the offline root CA's CRL manually if the CRL cannot be accessed directly from the network. You can load the CRL file using the certutil.exe –addstore <*CRL file*> command. |

**Renewing an Existing Enterprise Subordinate CA**

1. Ensure that the modified CAPolicy.inf file is in the %SystemRoot% folder.

   This ensures that either the newly defined constraints are applied or that the existing policy constraints are retained.

2. On the Start menu, point to Programs, point to Administrative Tools, and then click Certification Authority.
3. In the console tree, right-click CAName, and then click Renew CA Certificate.
4. In the Install CA Certificate dialog box, click Yes to stop Certificate Services.

   | **Note** |
   | --- |

| |
|---|
| In addition to defining constraints, the key size, key lifetime, CDP locations, AIA locations, and other settings defined in the CAPolicy.inf file can be redefined at certificate renewal. |

5. In the Renew CA Certificate dialog box (Figure 32), click No so that the same public and private key pair is used, and then click OK.



**Figure 32: Selecting Whether to Generate a New Public and Private Key Pair**

6. In the CA Certificate Request dialog box (Figure 33), enter the Computer Name and Parent CA name for the issuing CA, and then click Cancel. This will place a *CAMachineName_CAName.req* file in the root folder of the volume where Windows Server 2003 is installed.



**Figure 33: Selecting the CA to Which the CA Will Be a Subordinate CA**

| **Note** |
|---|
| If the CA that issued the SubCAs certificate is online, you can click OK to attempt the certificate request online. |

7. Copy the *CAMachineName_CAName.req* file to a disk, and then follow the procedures for performing the offline request and certificate installation used for installing a new subordinate CA certificate.

**Verifying the Inclusion of All Issuance Policies**

Once the CA certificate is installed, the Enterprise CAs certificate should be inspected to ensure that the Issuance policies are included in the certificate.

1. On the Start menu, point to Programs, point to Administrative Tools, and then click Certification Authority.

2. In the console tree, right-click CAName (where *CAName* is the name of your CA), and then click Properties.

3. In the CAName Properties dialog box, on the General tab, click View Certificate.

4. On the General tab, ensure that the intended purposes list includes:

    - High Assurance
    - Medium Assurance
    - Low Assurance

5. If any are missing, or appear as OIDs rather than text, review the %systemroot%\CAPolicy.inf file and ensure that the correct OIDs for the domain are included in the file. An incorrect OID will result in the OID being listed here, rather than being translated to the display names.

6. In the Certificate dialog box, click OK.

7. In the *CAName* Properties dialog box, click OK.

8. Close the Certification Authority console.

**Creating a Qualified Subordination Signing Certificate at a Stand-alone CA**
A stand-alone CA does not use certificate templates. Because certificate templates are not supported by stand-alone CAs, the certificate request must inject the OID for qualified subordination during certificate enrollment. The following steps allow a qualified subordination certificate to be obtained at a stand-alone CA.

**Performing the Initial Certificate Request**
To perform the initial request for a qualified subordination signing certificate

1. At the stand-alone CA, ensure that you are logged on as the Administrator of the CA computer.

2. Start Internet Explorer.

3. Open the **http://localhost/certsrv** URL.

4. On the Welcome page, click Request a certificate.

5. On the Request a Certificate page, click advanced certificate request.

6. On the Advanced Certificate Request page, click Create and submit a request to this CA.

7. On the Advanced Certificate Request page, enter the following information in Identifying Information:

    - Name: **Administrator** (or the name of the user account currently in use)
    - E-mail: Leave blank
    - Company: **YourOrganization** (where *YourOrganization* is the name of your organization)
    - Department: **YourDepartment** (where *YourDepartment* is the name of your department)
    - City: **CityName** (where *CityName* is the name of your city)
    - State: **StateName** (where *StateName* is the name of your state or province)

- Country/Region: **CountryName** (where *CountryName* is the name of your country or region)

8. In Type of Certificate Needed, select Other from the list.
9. In the OID box, type **1.3.6.1.4.1.311.10.3.10**.
10. In Key Options, set the following options:

    - **Create a new key set**
    - CSP: **Microsoft Enhanced Cryptographic Provider v1.0** (or a different CSP if you are using a hardware CSP).
    - Key Usage: **Signature**
    - **Automatic Key container name**
    - Mark keys as exportable: **unchecked**
    - Enable strong private key protection: **unchecked**
    - Use local machine store: **unchecked**

11. In Additional Options, set the following options:

    - Request format: **CMC**
    - Hash Algorithm: **SHA-1**
    - Save request to a file: **unchecked**
    - Attributes: **empty**
    - Friendly name: **Qualified Subordination Signing**

12. Review the entries, and then click Submit.

    The Certificate Pending page appears.

**Issuing the Pending Certificate Request**

The certificate request is set to a pending status by default on a stand-alone CA. To issue the certificate, a CA administrator must use the Certification Authority console.

1. On the Start menu, point to Programs, point to Administrative Tools, and then click Certification Authority.
2. In the console tree, click the Pending Requests container.
3. In the details pane, right-click the certificate request, point to All Tasks, and then click Issue.
4. Close the Certification Authority console.

**Installing the Qualified Subordination Signing Certificate**

The final step is to install the issued certificate.

1. In Internet Explorer, open the **http://localhost/certsrv/** URL.

2. On the Welcome page, click View the status of a pending certificate request.

3. On the View the status of a pending certificate request page, click the User-EKU (1.3.6.4.1.311.10.3.10) Certificate (*Date and Time*) link.

4. On the Certificate Issued page, click Install this certificate.

   The Certificate Installed page should appear.

5. Close Internet Explorer.

**Verifying the Qualified Subordination Signing Certificate**

Once you have installed the qualified subordination certificate, you should view the certificate to ensure that it is correctly configured.

1. Open an empty MMC console.

2. On the File menu, click Add/Remove Snap-in.

3. In the Add/Remove Snap-in dialog box, click Add.

4. In the Add Standalone Snap-in dialog box, select Certificates, and then click Add.

5. In the Certificates snap-in dialog box, click My user account, and then click Finish.

6. In the Add Standalone Snap-in dialog box, click Close.

7. In the Add/Remove Snap-in dialog box, click OK.

8. In the console tree, expand **Certificates** – Current User, expand Personal, and then click Certificates.

9. In the details panel, double-click the certificate with the Template Friendly Name of Qualified Subordination Signing.

10. In the Certificate dialog box, on the General tab, ensure that the intended purposes only list Qualified Subordination.

11. In the Certificate dialog box, on the Details tab, ensure that the Key Usage indicates Digital Signature, Non-Repudiation (c0).

12. Click OK.

13. Close the console without saving changes.

**Creating a Qualified Subordination Signing Certificate at an Enterprise CA**

An Enterprise CA running on Windows Server 2003 Enterprise Edition can issue version 2.0 certificate templates. For qualified subordination, a version 2.0 template can be created to allow qualified subordination signing. If you are going to issue the Qualified Subordination Signing certificate from an Enterprise CA, the following steps will allow the creation and retrieval of a qualified subordination signing certificate at an enterprise CA.

| **Note** |
| --- |
| While the goal is the same as the previous section on creating a qualified subordination signing certificate, the method is different when issuing the certificate from an Enterprise CA. The decision on which CA type to use should be based on what CAs are available in the infrastructure. |

**Creating a Qualified Subordination version 2 Certificate Template**
To create a version 2 certificate template for Qualified Subordination Signing

1. Open an empty MMC console.

2. On the File menu, click Add/Remove Snap-in.

3. In the Add/Remove Snap-in dialog box, click Add.

4. In the Add Standalone Snap-in dialog box, select Certificate Templates, and then click Add.

5. In the Add Standalone Snap-in dialog box, click Close.

6. In the Add/Remove Snap-in dialog box, click OK.

7. In the console tree, select Certificate Templates.

8. In the details pane, right-click Enrollment Agent, and then click Duplicate Template.

9. In the Properties of New Template dialog box, on the General tab, in the Template display name box, type **Qualified Subordination Signing**.

10. On the Request Handling tab (Figure 34), make the following property changes:



**Figure 34: Setting Request Handling Attributes**

- Purpose: **Signature**

- CSPs: Clear the check boxes for **Microsoft Base Cryptographic Provider 1.0** and **Microsoft Base DSS Cryptographic Provider** and ensure that **Microsoft Enhanced Cryptographic Provider 1.0** is checked.

11. On the Subject Name tab, ensure that the Subject name format is set to Fully distinguished name and that only the User principal name (UPN) check box is selected.

12. On the Issuance Requirements tab (Figure 35), ensure that the CA certificate manager approval check box is selected. For Require the following for reenrollment, ensure that the Valid existing certificate option is set.




**Figure 35: Configuring Issuance Requirement Settings**

| **Note** |
| --- |
| Depending on the security policy of your organization, you may choose to use different issuance settings. For example, you may choose to require one or more authorized signatures for certificate issuance, or that the same approval process be used for certificate reenrollment. |

13. On the Extensions tab, in the Extensions included in this template box, select Application policies, and then click Edit.

14. In the Edit Application Policies Extension dialog box, in the Application policies box, select Certificate Request Agent, and then click Remove.

15. In the Edit Application Policies Extension dialog box, click Add.

16. In the Add Application Policy dialog box, in the Application Policies box, select Qualified Subordination, and then click OK.

| **Note** |
| --- |
| Alternatively, you could create a custom application policy with its own assigned OID. This will require a modification later in the process where you will configure the XXXX Certificate template to require your custom application policy for a signature. |

17. In the Edit Application Policies Extension dialog box, ensure that Qualified Subordination is the only policy that appears in the Application policies dialog box, and then click OK.

18. On the Security tab, ensure that only Domain Admins and Enterprise Admins have the Enroll permission for the certificate template. Alternatively, you can create a custom security group and only assign that security group the Enroll permission.

19. Click OK to apply all configuration changes to the certificate template.

**Modifying the Cross-Certification Authority version 2 Certificate Template**

Once the qualified subordination signing template is created, the cross-certification authority certificate template must be

modified.

1. In the Certificate Templates console, in the details pane, double-click the Cross Certification Authority certificate template.

2. In the Cross-Certification Authority Properties dialog box, on the Issuance Requirements tab (Figure 36), make the following changes:




**Figure 36: Modifying the Issuance Requirements for the Cross CA Certificate**

   - **This number of authorized signatures** A request can be configured to require more than one signature for issuance. You should require at least one signature.

   - **Application Policy** If a custom application policy was created, change the default application policy from Qualified Subordination to the logical name you created for your custom application policy.

3. Click OK.

4. Close the Certificate Templates console

| **Important** |
| --- |
| By default, the CrossCA template is defined to publish to Active Directory when issued. Ensure that the certificate is formatted and configured properly before issuing. Invalid or unintended constraints may cause unwanted behavior in the environment and pose security risks, if not properly prepared. |

**Configuring Windows Server 2003 Enterprise Edition to Issue the Certificate Templates**

Since only Windows Server 2003 Enterprise Edition supports version 2 certificate templates, it is required for the certificate deployment. The Windows Server 2003 Enterprise Edition must be configured to issue both the Qualified Subordination Signing and cross-certification authority certificate templates.

1. Log on as a CA Administrator at Windows Server 2003 Enterprise Edition running Certificate Services configured as an Enterprise CA.

2. On the Start menu, point to Programs, point to Administrative Tools, and then click Certification Authority.

3. In the console tree, expand CAName (where *CAName* is the logical name assigned to your CA).

4. In the console tree, right-click Certificate Templates, point to New, and then click Certificate Template to Issue.

5. In the Enable Certificate Templates dialog box (Figure 37), in the list of available templates, click Cross-Certification Authority, control-click Qualified Subordination Signing, and then click OK.




**Figure 37: Loading the Cross-Certification Authority and Qualified Subordination Signing Certificates**

6. In the details pane, ensure that the Cross-Certification Authority and Qualified Subordination Signing certificate appear.

7. Close the Certification Authority console.

**Acquiring the Qualified Subordination Signing Certificate**

Although you typically request certificates from an Enterprise CA using the Certificates MMC console, the requirement for CA Certificate manager approval in the Qualified Subordination Signing certificate template makes the use of the Certificate Services Web enrollment pages the better method to acquire the Qualified Subordination Signing certificate.

**Performing the Initial Certificate Request**

To perform the initial request for a qualified subordination signing certificate

1. Ensure that you are logged on as a user with the Enroll permissions for the Qualified Subordination Signing Certificate template.

2. Start Internet Explorer.

3. Open the **http://CAMachineName/certsrv** URL (where *CAMachineName* is the name of the Windows Server 2003 Enterprise Edition configured to issue the Qualified Subordination Signing Certificates).

4. On the Welcome page, click Request a certificate.

5. On the Request a Certificate page, click advanced certificate request.

6. On the Advanced Certificate Request page, click Create and submit a request to this CA.

7. On the Advanced Certificate Request page, in the Certificate Template section, select Qualified Subordination Signing from the list.

8. In the Key Options section, set the following options:

   - **Create a new key set**

- CSP: **Microsoft Enhanced Cryptographic Provider v1.0** (or a different CSP if you are using a hardware CSP).
- Key Usage: **Signature**
- **Automatic Key container name**
- Mark keys as exportable: **unchecked**
- Enable strong private key protection: **unchecked**
- Use local machine store: **unchecked**

9. In the Additional Options section, set the following options:

   - Request format: **CMC**
   - Hash Algorithm: **SHA-1**
   - Save request to a file: **unchecked**
   - Attributes: **empty**
   - Friendly name: **Qualified Subordination Signing**

10. Review the entries, and then click Submit.

    The Certificate Pending page appears.

**Issuing the Pending Certificate Request**

The certificate request is set to a pending status by default on a stand-alone CA. To issue the certificate, a CA administrator must use the Certification Authority console.

1. Ensure that you are logged on to the Windows Server 2003 Enterprise Edition as a CA Certificate Manager.
2. On the Start menu, point to Programs, point to Administrative Tools, and then click Certification Authority.
3. In the console tree, click the Pending Requests container.
4. In the details pane, right-click the certificate request for a Qualified Subordination Signing certificate, point to All Tasks, and then click Issue.
5. Close the Certification Authority console.

**Installing the Qualified Subordination Signing Certificate**

The final step is to install the issued certificate.

1. At the computer where the original certificate request was performed, log on as the user who made the Qualified Subordination Signing certificate request.
2. In Internet Explorer, open the **http://CAMachineName/certsrv/** URL (where *CAMachineName* is the name of the Windows Server 2003 Enterprise Edition configured to issue the Qualified Subordination Signing Certificates).
3. On the Welcome page, click View the status of a pending certificate request.

4. On the View the status of a pending certificate request page, click the QualifiedSubordinationSigning Certificate (*Date and Time*) link.

5. On the Certificate Issued page, click Install this certificate.

   The **Certificate Installed** page should appear.

6. Close Internet Explorer.

**Verifying the Qualified Subordination Signing Certificate**

Once you have installed the qualified subordination certificate, you should view the certificate to ensure that it is correctly configured.

1. Open an empty MMC console.

2. On the File menu, click Add/Remove Snap-in.

3. In the Add/Remove Snap-in dialog box, click Add.

4. In the Add Standalone Snap-in dialog box, select Certificates, and then click Add.

5. In the Certificates snap-in dialog box, click My user account, and then click Finish.

6. In the Add Standalone Snap-in dialog box, click Close.

7. In the Add/Remove Snap-in dialog box, click OK.

8. In the console tree, expand Certificates – Current User, expand Personal, and then click Certificates.

9. In the details panel, double-click the certificate with the template Friendly Name of Qualified Subordination Signing.

10. In the Certificate dialog box, on the General tab, ensure that the intended purposes only list Qualified Subordination.

| **Note** |
| --- |
| If you used a custom application policy, then the name of this application policy and its OID will appear instead of Qualified Subordination. |

11. In the Certificate dialog box, on the Details tab, ensure that the Key Usage indicates Digital Signature (80).

12. Click OK.

13. Close the console without saving the changes.

**Performing Cross-Certification**

Once the qualified subordination signing certificate is successfully installed, you can now request the Cross-Certification Authority certificate from the other CA. In this example, the following CAs shown in Figure 38 will be used.




**Figure 38: Performing the Cross-Certification**

**Gathering the Necessary Files**

The first cross-certification request is to be run from the F1CA certification authority. To perform the qualified subordination, the following is required at the F1CA before starting the process:

- The CA certificate from F2Root.
- The Policy.inf file configured at the F1CA defining the naming constraints, policy constraints, and application constraints defined for the qualified subordination.
- The user performing the qualified subordination must have a qualified subordination signing certificate issued by F1CA.

**Generating the Cross-Certification Authority Certificate Request**

Once these items are obtained, the following procedure is used to obtain a Cross-Certification Authority certificate for F2Root from F1CA. The request is performed using Certreq.exe, a command-line tool used from the Windows Server 2003 administrative tools pack that requests certificates from the command line.

| **Important** |
| --- |
| The constraints defined in Policy.inf are enforced at the signing CA where you generate the request, not at the CA whose CA certificate you use during the request. It is also important to generate the cross-certificate request in the domain in which the cross-cert will be issued. |

1. Copy the F2Root certificate and the Policy.inf file into a folder on the F1CA computer.
2. Open a command prompt window.
3. Make the folder containing the F2Root certificate and the Policy.inf file the current directory.
4. Run certreq.exe –policy. This command constructs a cross-certification or qualified subordination request from an existing CA certificate or from an existing request.

   | **Note** |
   | --- |
   | If creating a cross-certificate request from a stand-alone CA to an Enterprise CA, the following command must be specified to include the proper CrossCA template name in the template. Example: |

   ```
   certreq -policy -attrib CertificateTemplate:CrossCA  <CatoXcertify> <inf
   file> <CMCoutFile>
   ```

5. In the Open Request File dialog box (Figure 39), in the Files of Type box, select Certificate Files (*.cer, *.crt, *.der), select the requesting CAs certificate (RackCluster3_F2Root in this example), and then click Open.




**Figure 39: Selecting the Target CAs Certificate**

6. In the Open Inf File dialog box (Figure 40), select the configured Policy.inf file, and then click Open.




**Figure 40: Selecting the Policy.inf File**

| **Note** |
| --- |
| You do not have to name the file Policy.inf. In this example, the file was named xrossForest2.inf to aid in determining which Policy.inf file was required. |

7. In the Certificate List dialog box (Figure 41), select the Qualified Subordination certificate that you requested earlier, and then click OK.




**Figure 41: Selecting the Qualified Subordination Signing Certificate**

8. In the Save Request dialog box (Figure 42), in the File name box, type a name describing the Cross-Certification Authority certificate request, and then click Save.




**Figure 42: Saving the Request File**

**Processing the Cross-Certification Authority Certificate Request**

At the F1CA, the final step is to process the Cross-Certification Authority certificate request file.

1. In Administrative Tools, open Certification Authority.

2. In the console tree, right-click CAName (where *CAName* is the name of your CA), point to All Tasks, and then click Submit New Request.

3. In the Open Request File dialog box (Figure 43), select the request file that you created in the previous process, and then click Save.




**Figure 43: Choosing the Request File**

4. In the Save Certificate dialog box, indicate the name of the Certificate File, and then click Save.

   The certificate file is automatically published into Active Directory, so that the clients in Forest1 can validate certificates issued by the CAs in Forest2.

**Verifying the Existence of the Cross Certificate**

The final step is verifying that the certificate has been successfully published into Active Directory. This process uses Certutil.exe to verify the existence of the cross certificate. The following walkthrough validates the certificate with subject name ASIA SA Root CA that was issued by the Microsoft Intranet CA.

1. Open a command prompt window.

2. Type **certutil -viewstore "CN=<CAName>,CN=AIA,CN=Public Key Services,CN=Services,CN=Configuration,DC=<ForestRootDN>?crossCertificatePair**

   Where:

   - *<CAName>* The name of the CA that is issued the cross-certification authority certification. This is the subject of the cross-certification authority certificate.

   - *<ForestRootDN>* The name of the forest root LDAP distinguished name. This is the forest in which the cross-certification authority certificate exists.

   In this example, the following command was used:

```
certutil -viewstore "CN=ASIA SA Root CA,CN=AIA,CN=Public Key
Services,CN=Services,CN=Configuration,DC=corp,DC=microsoft,DC=com?
crossCertificatePair
```

| **Note** |
| --- |
| If the subject name of the CA being cross-certified does not have CN= in the subject name, the Windows Server 2003 CA will generate a hash of the binary name of the CA and publish the cross certificate to that name. For example, to perform the command in the previous step in this case, the syntax would be similar to the following where the common name is the hash of the CA name: |

```
certutil -viewstore
CN=12d2cb85c5e62f9aa7591e2ddaa44c987de5abbc,CN=AIA,CN=Public Key
Services,CN=Services,CN=Configuration,DC=ntdev,DC=corp,DC=microsoft,DC=com
```

3. In the View Certificate Store window (Figure 44), all cross-certification authority certificates with the <CAName> in the subject name are listed.




**Figure 44: The Cross-Certification Authority Certificates List**

| **Note** |
| --- |
| Multiple cross-certification authority certificates can exist when the cross-certification authority certificate is renewed when the previous certificate expires. |

4. In the View Certificate Store dialog box, select the cross-certification authority certificate that you wish to view, and then click View Certificate.

5. In the Certificate dialog box, click the Certification Path tab.

6. On the Certification Path tab (Figure 45), ensure that the certification path is the correct path. In this example, the chain is correct, as it shows the ASIA SA Root CA chaining to the Microsoft Corporate Root Authority.




**Figure 45: Validating the Certification Path of the Cross-Certification Authority Certificate**

7. In the Certificate dialog box, click OK.

8. In the View Certificate Store dialog box, click OK.

9. Close the command prompt window.

When cross-certifying with some third-party certification authorities, you may have to specify where to retrieve additional cross certificates associated with the issuer of a presented certificate. Typically, a certificate will have a cross-certificate URL encoded in the certificate, but in some cases, you may need to specify additional URLs.

To add the additional URLs

1. Open a blank MMC console.

2. On the File menu, click the Add/Remove snap-in.

3. In the Add/Remove snap-in dialog box, click Add.

4. In the Add Standalone Snap-in dialog box, select Certificates, and then click Add.

5. In the Certificates snap-in dialog box, select either My user account, Service account, or computer account, based on whether the certificate that must chain using a cross certificate is issued to the computer, a service, or to your user account, and then click Finish.

| **Note** |
| --- |
| If you are not an Administrator of the local computer, you will automatically load the Certificates console focused on My user account. |

6. In the console tree, expand Certificates – Current User, expand Personal, and then click Certificates.

7. In the details pane, right-click the certificate that must be cross-certified, and then click Properties.

8. In the Certificate Properties dialog box, click the Cross-Certificates tab.

9. In the Cross-Certificates tab (Figure 46), select the Specify additional Cross-Certificate download locations check box. In addition, you can now do the following:




**Figure 46: Defining Manual Cross-Certificate Download Locations**

- Define the Cross-Certificate download interval. The default is every eight hours.
- Add URLs for Cross-Certificate download locations. The URL can be an LDAP, HTTP, FTP, or FILE location.

**Completing Cross-Certification**

The cross-certification must be performed in the opposite direction to ensure that all certificates issued in both organizations are trusted. This requires that the same process be performed at F2CA using the F1Root certificate.

**Completing the Cross-Certification When a Bridge CA Is Implemented**

An additional process must be performed when a Bridge CA is used. The certificates that are issued by the BridgeCA must be published in all forests that are connected to the BridgeCA to allow all certificates to be recognized by the organizations using the bridge.

1. From the BridgeCA, copy all issued CrossCA certificates to a common share point.
2. At the first forest connected to the bridge, run **Certutil –dspublish –f cert1.crt CrossCA** (where *cert1.crt* is the first crossCA cert).
3. Repeat the process for all certificates issued by the Cross CA, at all forests connected to the Bridge CA.

# Appendix A – Policy.Inf

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

```
[Version]
Signature= "$Windows NT$"
; ============================================================
; Request Attributes
; top level section
; ============================================================
[RequestAttributes]
CertificateTemplate = CrossCA
AttributeName1 = AttributeValue1
AttributeName2 = AttributeValue2
; ============================================================
; NameConstraintsExcluded Name Constraints Extension
; szOID_NAME_CONSTRAINTS   2.5.29.30
; top level section
; ============================================================
[NameConstraintsExtension]
Include = NameConstraintsPermitted
Exclude = NameConstraintsExcluded
Critical = TrUe
[NameConstraintsPermitted]
; list of user defined permitted DNS names
; the numeric second and third arguments are optional
; when present, the second argument is the minimum depth
; when present, the third argument is the maximum depth
; NOTE: Crypto APIs fail to process cert chains when the minimum or maximum
; depth is specified!
DNS = foo@domain.com
DNS = domain1.domain.com
email=me@you.com
UPN=user@domain.com
; the first is an IP address, the second is an IP address mask
IPADDRESS=255.255.18.172,255.255.255.0
ipaddress=::255.255.18.172,::255.255.255.0
ipaddress=1234:5678:9abc:def0:3210:7654:ba98:fedc,1234:5678:9abc:def0:3210:7654:ba98:fe
dc
ipaddress=::5678:9abc:def0:3210:7654:ba98:fedc,1234:5678:9abc:def0:3210:7654:ba98:fedc
ipaddress=1234::def0:3210:7654:ba98:fedc,1234:5678:9abc:def0:3210:7654:ba98:fedc
ipaddress=1234:5678:9abc:def0:3210:7654:ba98::,1234:5678:9abc:def0:3210:7654:ba9:fedc
ipaddress=1234:5678:9abc:def0:3210:7654::,1234:5678:9abc:def0:3210:7654:ba98:fedc
url=http://localhost/certsrv/default.html
url=file:\\localhost\certsrv\default.html
DIRECTORYNAME = "cn=mycn,ou=myou,s=mystate,c=us"
[NameConstraintsExcluded]
; list of user defined excluded DNS names
DNS = domain.com
; ============================================================
; Policy (CPS) Extension
; szOID_CERT_POLICIES   2.5.29.32
; top level section
; ============================================================
[PolicyStatementExtension]
; list of user defined policies
Policies = LegalPolicy, LimitedUsePolicy, ExtraPolicy, OIDPolicy
CRITICAL = FALSE
[LegalPolicy]
; each policy has one OID, and zero or more Notice and URL keys
OID = 1.3.6.1.4.1.311.21.43
Notice = "Legal policy statement text"
[LimitedUsePolicy]
OID = 1.3.6.1.4.1.311.21.47
URL = "http://http.site.com/some where/default.asp"
URL = "ftp://ftp.site.com/some where else/default.asp"
Notice = "Limited use policy statement text."
URL = "ldap://ldap.site.com/some where else again/default.asp"
```

```
[ExtraPolicy]
OID = 1.3.6.1.4.1.311.21.53
URL = http://extra.site.com/Extra Policy/default.asp
[oidpolicy]
OID = 1.3.6.1.4.1.311.21.55
; ============================================================
; Policy Mapping Extension
; szOID_POLICY_MAPPINGS   2.5.29.33
; top level section
; ============================================================
[PolicyMappingsExtension]
; list of user defined policy mappings
; first OID is Issuer Domain Policy OID, second is Subject Domain Policy OID
; each entry maps one foreign policy OID to local
1.3.6.1.4.1.311.21.53 = 1.2.3.4.87
1.3.6.1.4.1.311.21.54 = 1.2.3.4.89
critical = yEs
; ============================================================
; Policy Constraints Extension
; szOID_POLICY_CONSTRAINTS   2.5.29.36
; top level section
; ============================================================
[PolicyConstraintsExtension]
; consists of two optional DWORDs
; They refer to the depth of the CA hierarchy that requires explicit policy
; and inhibits Policy Mapping
RequireExplicitPolicy  = 3
InhibitPolicyMapping = 5
; ============================================================
; Application Policy (CPS) Extension
; szOID_APPLICATION_CERT_POLICIES   1.3.6.1.4.1.311.21.10
; top level section
; ============================================================
[ApplicationPolicyStatementExtension]
; list of user defined policies
Policies = AppLegalPolicy, AppLimitedUsePolicy, AppExtraPolicy, AppOIDPolicy
CRITICAL = FALSE
[AppLegalPolicy]
; each policy has one OID, and zero or more Notice and URL keys
OID = 1.3.6.1.4.1.311.21.54
Notice = "Application Legal policy statement text"
[AppLimitedUsePolicy]
OID = 1.3.6.1.4.1.311.21.58
URL = "http://http.site.com/application some where/default.asp"
URL = "ftp://ftp.site.com/application some where else/default.asp"
Notice = "Application Limited use policy statement text."
URL = "ldap://ldap.site.com/application some where else again/default.asp"
[AppExtraPolicy]
OID = 1.3.6.1.4.1.311.21.64
URL = http://extra.site.com/Application Extra Policy/default.asp
[Appoidpolicy]
OID = 1.3.6.1.4.1.311.21.66
; ============================================================
; Application Policy Mapping Extension
; szOID_APPLICATION_POLICY_MAPPINGS   1.3.6.1.4.1.311.21.11
; top level section
; ============================================================
[ApplicationPolicyMappingsExtension]
; list of user defined application policy mappings
; first OID is Issuer Domain Policy OID, second is Subject Domain Policy OID
; each entry maps one foreign policy OID to local
1.3.6.1.4.1.311.21.64 = 1.2.3.4.98
1.3.6.1.4.1.311.21.65 = 1.2.3.4.100
critical = trUE
; ============================================================
; Application Policy Constraints Extension
; szOID_APPLICATION_POLICY_CONSTRAINTS   1.3.6.1.4.1.311.21.12
; top level section
; ============================================================
[ApplicationPolicyConstraintsExtension]
; consists of two optional DWORDs
; They refer to the depth of the CA hierarchy that requires explicit policy
```

```
; and inhibits Policy Mapping
RequireExplicitPolicy  = 6
InhibitPolicyMapping = 10
; ============================================================
; Basic Constraints Extension
; szOID_BASIC_CONSTRAINTS2   2.5.29.19
; top level section
; ============================================================
[BasicConstraintsExtension]
; Subject Type is not supported  always set to CA
; maximum subordinate CA path length
PathLength = 3
[EnhancedKeyUsageExtension]
;OID = 1.3.6.1.4.1.311.21.6   ; szOID_KP_KEY_RECOVERY_AGENT
;OID = 1.3.6.1.4.1.311.10.3.9   ; szOID_ROOT_LIST_SIGNER
;OID = 1.3.6.1.4.1.311.10.3.1   ; szOID_KP_CTL_USAGE_SIGNING
; The following match the [ApplicationPolicyStatementExtension] section:
OID = 1.3.6.1.4.1.311.21.54
OID = 1.3.6.1.4.1.311.21.58
OID = 1.3.6.1.4.1.311.21.64
OID = 1.3.6.1.4.1.311.21.66
CriticAL = False
; ============================================================
; Cross Certificate Distribution Points Extension
; szOID_CROSS_CERT_DIST_POINTS   1.3.6.1.4.1.311.10.9.1
; top level section
; ============================================================
[CrossCertificateDistributionPointsExtension]
SyncDeltaTime = 24
URL = http://%1/Public/My CA.crt
URL = ftp://foo.com/Public/MyCA.crt
URL = file:////%1\Public\My CA.crt
CriticAL = false
```

# Appendix B – CAPolicy.inf

Updated: May 15, 2013

Applies To: Windows Server 2003 with SP1

The CAPolicy.inf contains settings that can be used to modify the default installation of the Certification Authority role of Active Directory Certification Service (AD CS). The file is also used when renewing the CA certificate. A CAPolicy.inf file is not required to install AD CS or renew a CA certificate. The file is only needed to modify default settings.

Once you have created your CAPolicy.inf file, you must copy it into the %windir% folder (such as the C:\Windows) of your server before you install AD CS or renew the CA certificate.

**CAPolicy.inf**

The following code illustrates the text of a sample CAPolicy.inf. You can create and save such a file with any text editor, such as Notepad.exe. Ensure that you save the file with the .inf file extension and not .txt or some other file extension. The name of the file must be CAPolicy.inf and must be in the %windir% folder, or it will not work.

```
[Version]
Signature= "$Windows NT$"
;[CAPolicy]
[PolicyStatementExtension]
Policies = LegalPolicy, LimitedUsePolicy, ExtraPolicy, OIDPolicy, EmptyPolicy
Critical = 0
[LegalPolicy]
OID = 1.3.6.1.4.1.311.21.43
Notice = "Legal policy statement text."
[LimitedUsePolicy]
OID = 1.3.6.1.4.1.311.21.47
URL = "http://http.site.com/some where/default.asp"
URL = "ftp://ftp.site.com/some where else/default.asp"
Notice = "Limited use policy statement text."
URL = "ldap://ldap.site.com/some where else again/default.asp"
[ExtraPolicy]
OID = 1.3.6.1.4.1.311.21.53
URL = http://extra.site.com/Extra Policy/default.asp
[oidpolicy]
OID = 1.3.6.1.4.1.311.21.55
[emptypolicy]
; For CRLDistributionPoint, AuthorityInformationAccess and
; CrossCertificateDistributionPointsExtension URLs:
;
; #define wszFCSAPARM_SERVERDNSNAME              L"%1"
; #define wszFCSAPARM_SERVERSHORTNAME            L"%2"
; #define wszFCSAPARM_SANITIZEDCANAME            L"%3"
; #define wszFCSAPARM_CERTFILENAMESUFFIX         L"%4"
; #define wszFCSAPARM_DOMAINDN                   L"%5"
; #define wszFCSAPARM_CONFIGDN                   L"%6"
; #define wszFCSAPARM_SANITIZEDCANAMEHASH        L"%7"
; #define wszFCSAPARM_CRLFILENAMESUFFIX          L"%8"
; #define wszFCSAPARM_CRLDELTAFILENAMESUFFIX     L"%9"
; #define wszFCSAPARM_DSCRLATTRIBUTE             L"%10"
; #define wszFCSAPARM_DSCACERTATTRIBUTE          L"%11"
; #define wszFCSAPARM_DSUSERCERTATTRIBUTE        L"%12"
; #define wszFCSAPARM_DSKRACERTATTRIBUTE         L"%13"
; #define wszFCSAPARM_DSCROSSCERTPAIRATTRIBUTE   L"%14"
;
; Setup APIs replace all %<number>% sequences with various directory paths.
; %3%8%9 in the first URL below presents two opportunities for string
; replacement with a directory path.  To avoid this, use two percent signs
; to escape the setup API string replacement.
;
; URLs with spaces or commas must be quoted to avoid INF parsing problems
;
; default CDP registry URLs:
;
; D:\WINDOWS\System32\CertSrv\CertEnroll\%3%8%9.crl
; ldap:///CN=%7%8,CN=%2,CN=CDP,CN=Public Key Services,CN=Services,%6%10
; http://%1/CertEnroll/%3%8%9.crl
```

```
; file://\\%1\CertEnroll\%3%8%9.crl
[AuthorityInformationAccess]
URL = http://%1/Public/My CA.crt
URL = ftp://foo.com/Public/MyCA.crt
URL = file://\\%1\Public\My CA.crt
CriticAL = falSe
[CRLDistributionPoint]
URL = http://%1/Public/My CA.crl
URL = ftp://%1/Public/MyCA.crl
URL = file://\\%1\Public\My CA.crl
CriticAL = trUe
[CrossCertificateDistributionPointsExtension]
SyncDeltaTime = 600 ; in seconds
URL = http://%1/Public/My CCDP.crl
URL = ftp://%1/Public/MyCCDP.crl
URL = file://\\%1\Public\My CCDP.crl
CriticAL = yeS
[EnhancedKeyUsageExtension]
OID = 1.3.6.1.4.1.311.21.6 ; szOID_KP_KEY_RECOVERY_AGENT
OID = 1.3.6.1.4.1.311.10.3.9 ; szOID_ROOT_LIST_SIGNER
OID = 1.3.6.1.4.1.311.10.3.1 ; szOID_KP_CTL_USAGE_SIGNING
CriticAL = False
[basicconstraintsextension]
pathlength = 13
criticaL=falsE
[certsrv_server]
renewalkeylength=2048
RenewalValidityPeriodUnits=0x18
RenewalValidityPeriod=years
CRLPeriod = days
CRLPeriodUnits = 2
CRLDeltaPeriod = hours
CRLDeltaPeriodUnits = 4
```

| **Note** |
| --- |
| Name constraints are not supported in the CAPolicy.inf |

**See Also**

**Other Resources**

Windows Server 2008 R2 CAPolicy.inf Syntax
Certification Authority Guidance
AD CS Step by Step Guide: Two Tier PKI Hierarchy Deployment
Test Lab Guide: Deploying an AD CS Two-Tier PKI Hierarchy

# Appendix C – CMC Dump of a Qualified Subordination Request

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

```
PKCS7/CMS Message:
  CMSG_SIGNED(2)
  CMSG_SIGNED_DATA_CMS_VERSION(3)
  Content Type: 1.3.6.1.5.5.7.12.2 CMC Data
PKCS7 Message Content:
================ Begin Nesting Level 1 ================
CMS Certificate Request:
Tagged Attributes: 2
  Body Part Id: 2
  1.3.6.1.5.5.7.7.8 CMC Extensions
  Value[0]:
    Data Reference: 0
    Cert Reference[0]: 1
  Extensions: 7
    2.5.29.36: Flags = 0, Length = 8
    Policy Constraints
        Required Explicit Policy Skip Certs=0
        Inhibit Policy Mapping Skip Certs=0
    2.5.29.32: Flags = 0, Length = 10
    Certificate Policies
429.195.0: 0x80070002 (WIN32: 2): LDAPFlags
        [1]Certificate Policy:
             Policy Identifier=Corporate High Assurance
    1.3.6.1.4.1.311.21.10: Flags = 0, Length = 6c
    Application Policies
        [1]Application Certificate Policy:
             Policy Identifier=Client Authentication
        [2]Application Certificate Policy:
             Policy Identifier=Smart Card Logon
        [3]Application Certificate Policy:
             Policy Identifier=Corporate RAS
        [4]Application Certificate Policy:
             Policy Identifier=Private Key Archival
        [5]Application Certificate Policy:
             Policy Identifier=Key Recovery Agent
        [6]Application Certificate Policy:
             Policy Identifier=Encrypting File System
        [7]Application Certificate Policy:
             Policy Identifier=Secure Email
        [8]Application Certificate Policy:
             Policy Identifier=Certificate Request Agent
    2.5.29.30: Flags = 0, Length = bf
    Name Constraints
        Permitted
             [1]Subtrees (0..Max):
                  Other Name:
                       Principal Name=.asia.northwindtraders.com
             [2]Subtrees (0..Max):
                  RFC822 Name=@northwindtraders.com
             [3]Subtrees (0..Max):
                  RFC822 Name=.northwindtraders.com
             [4]Subtrees (0..Max):
                  DNS Name=.asia.northwindtraders.com
             [5]Subtrees (0..Max):
                  Directory Address:
                       DC=ASIA
                       DC=Northwindtraders
                       DC=com
             [6]Subtrees (0..Max):
                  URL=
             [7]Subtrees (0..Max):
                  IP Address=
        Excluded=None
    1.3.6.1.4.1.311.20.2: Flags = 0, Length = c
```

```
    Certificate Template Name
        SubCA
    2.5.29.15: Flags = 0, Length = 4
    Key Usage
        Digital Signature, Certificate Signing, Off-line CRL Signing, CRL Signing (86)
    2.5.29.19: Flags = 1(Critical), Length = 5
    Basic Constraints
        Subject Type=CA
        Path Length Constraint=None
  Body Part Id: 3
  1.3.6.1.5.5.7.7.18 Reg Info
  Value[0]:
    CertificateTemplate: SubCA
Tagged Requests: 1
  CMC_TAGGED_CERT_REQUEST_CHOICE:
  Body Part Id: 1
================ Begin Nesting Level 2 ================
Element 0:
PKCS10 Certificate Request:
Version: 1
Subject:
    CN=ITG ASIA Corp CA 1
    DC=asia
    DC=Northwindtraders
    DC=com
Public Key Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.1 RSA
    Algorithm Parameters:
    05 00
Public Key Length: 2048 bits
Public Key: UnusedBits = 0
    0000  30 82 01 0a 02 82 01 01  00 c4 92 eb 3d e3 70 52
    0010  23 9f 9e a0 6c 9e 1e 26  43 7e 9c a3 d1 82 56 88
    0020  5f df 2d a3 c6 f2 0a ae  25 8b 4e c8 7c 2b a2 4a
    0030  72 49 ff 48 46 d9 59 6b  9e 1e 76 1a ff a9 1b 29
    0040  30 4f a7 00 0f 73 3d 16  6b 4c 57 cd 2b c5 3d 78
    0050  82 81 4a 90 26 b7 8b d4  b1 c4 08 ea 77 2a c2 f8
    0060  8e e9 93 98 47 21 96 8e  f9 9d ac bc 5f 01 f9 09
    0070  12 b6 73 70 9a 2e 35 1c  51 d0 74 54 ee 46 7e 92
    0080  03 5e d4 86 10 86 02 8b  8c 38 7e 76 10 55 0b 92
    0090  1d 85 b9 46 d7 eb c2 42  3d a4 3d 84 d7 1f dd 93
    00a0  30 ae 96 57 76 05 5d 2f  6e d0 7f 17 21 c2 87 1b
    00b0  82 0c 02 da 10 87 48 ec  c6 ba 45 45 75 22 3f 9a
    00c0  f8 1f c6 10 05 08 01 d5  fa 56 25 a3 19 2c da e0
    00d0  74 f6 43 9a c1 4d ed b6  9e 83 91 35 d0 c9 c3 6b
    00e0  72 2f b0 3c fd 05 27 35  7b ea 8b 9d 48 83 96 59
    00f0  bf b9 d3 80 b1 14 71 8e  75 e1 c9 da 69 86 4e cc
    0100  9a 00 01 83 f5 0f 4b 2e  55 02 03 01 00 01
Request Attributes: 2
  2 attributes:
  Attribute[0]: 1.3.6.1.4.1.311.13.2.3 (OS Version)
    Value[0][0]:
        5.1.3541.2.
  Attribute[1]: 1.2.840.113549.1.9.14 (Certificate Extensions)
    Value[1][0]:
    Unknown Attribute type
Certificate Extensions: 6
    1.3.6.1.4.1.311.21.1: Flags = 0, Length = 3
    CA Version
        V1.0
    1.3.6.1.4.1.311.21.2: Flags = 0, Length = 16
    Previous CA Certificate Hash
        d3 e5 cc ef 88 53 0d 13 b7 ae a2 7b 19 5f 57 5e 33 62 b0 ef
    2.5.29.14: Flags = 0, Length = 16
    Subject Key Identifier
        d3 b6 48 1a c0 76 07 ba 35 2b 1c 90 8b bc 1f 2b d3 b9 4d f8
    1.3.6.1.4.1.311.20.2: Flags = 0, Length = c
    Certificate Template Name
        SubCA
    2.5.29.15: Flags = 0, Length = 4
    Key Usage
        Digital Signature, Certificate Signing, Off-line CRL Signing, CRL Signing (86)
```

```
    2.5.29.19: Flags = 1(Critical), Length = 5
    Basic Constraints
        Subject Type=CA
        Path Length Constraint=None
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.5 sha1RSA
    Algorithm Parameters:
    05 00
Signature: UnusedBits=0
    0000  a1 4b fa 5b ec 6c 4d e8  e4 a4 2d 7f 4b b9 65 cb
    0010  61 1e d5 85 49 e9 07 9c  c5 8b 9b e2 a4 b5 18 26
    0020  76 33 d9 3d 66 2e 32 2c  e8 a5 23 65 9f 07 0a d6
    0030  d3 d0 3f 1a ac ef 75 2d  53 0d 79 13 90 6f e3 61
    0040  ab bb ad 31 e9 46 31 b4  78 33 3e 7d 4f d4 87 3c
    0050  cf 75 0a 03 51 f0 f4 f3  15 25 b8 d8 5f bb d4 78
    0060  5d a9 39 34 1d d7 f9 8a  8d 3c 4e 1f ea 79 23 ce
    0070  85 42 8f 36 c2 24 d2 9f  90 37 93 80 d4 f2 76 74
    0080  ad 65 a0 7e 83 fc 83 21  ea e5 d8 c9 5f 02 ea d0
    0090  9f 50 96 3a b5 c3 7f 85  9b b8 fc cd 68 c5 27 c5
    00a0  99 d6 5f df 8f 8b 82 7b  0f 21 f0 3d 9f 34 0c d8
    00b0  ec 2b b1 a9 55 c7 01 2d  9e f0 89 76 d3 ed d6 33
    00c0  55 8a 9a c7 8d 52 3e b1  5d da 35 61 28 f7 07 73
    00d0  57 52 80 ac c0 31 ad 9e  81 49 01 1f 48 1f f4 95
    00e0  9f 39 4b 1a be 33 c6 4b  cf 67 04 5b aa 94 59 4d
    00f0  d4 95 42 6f 25 d3 64 fc  d6 c8 e7 ca e3 b3 0a ad
Signature matches Public Key
Key Id Hash(sha1): d3 b6 48 1a c0 76 07 ba 35 2b 1c 90 8b bc 1f 2b d3 b9 4d f8
----------------  End Nesting Level 2  ----------------
Tagged Content Info: 0
Tagged Other Messages: 0
----------------  End Nesting Level 1  ----------------
Signer Count: 2
Signer Info[0]:
NULL signature verifies
CMSG_SIGNER_INFO_PKCS_1_5_VERSION(1)
CERT_ID_ISSUER_SERIAL_NUMBER(1)
    Serial Number: 00
    Issuer: OID.1.3.6.1.4.1.311.21.9=Dummy Signer
Hash Algorithm:
    Algorithm ObjectId: 1.3.14.3.2.26 sha1
    Algorithm Parameters: NULL
Encrypted Hash Algorithm:
    Algorithm ObjectId: 1.3.6.1.5.5.7.6.2 NO_SIGN
    Algorithm Parameters: NULL
Encrypted Hash:
    0000  75 80 82 26 ff 11 77 5b  92 52 ce 2e a2 8e a2 32
    0010  98 a7 1a a0
Authenticated Attributes[0]:
  3 attributes:
  Attribute[0]: 1.2.840.113549.1.9.3 (Content Type)
    Value[0][0]:
    Unknown Attribute type
    1.3.6.1.5.5.7.12.2 CMC Data
  Attribute[1]: 1.2.840.113549.1.9.4 (Message Digest)
    Value[1][0]:
    Unknown Attribute type
    Message Digest:
        5a 74 12 86 78 3a ab 5b 17 85 6f 4d 44 ea a2 74 2c 86 c1 1f
  Attribute[2]: 1.3.6.1.4.1.311.21.20 (Client Information)
    Value[2][0]:
    Unknown Attribute type
    Client Id: = 4
    XECI_CERTREQ -- 4
    User: ASIA\user2
    Machine: user2.asia.northwindtraders.com
    Process: certreq
Unauthenticated Attributes[0]:
  0 attributes:
Computed Hash: 75 80 82 26 ff 11 77 5b 92 52 ce 2e a2 8e a2 32 98 a7 1a a0
Signing Certificate Index: 1
-------- CERT_CHAIN_CONTEXT --------
ChainContext.dwInfoStatus = CERT_TRUST_HAS_PREFERRED_ISSUER (0x100)
```

```
CertContext.dwRevocationFreshnessTime: 27 Days, 4 Hours, 47 Minutes, 46 Seconds
SimpleChain.dwInfoStatus = CERT_TRUST_HAS_PREFERRED_ISSUER (0x100)
SimpleChain.dwRevocationFreshnessTime: 27 Days, 4 Hours, 47 Minutes, 46 Seconds
CertContext[0][0]: dwInfoStatus=101 dwErrorStatus=0
  Issuer: CN=W2K User Enrollment CA, OU=Test, O=NT Distributed Systems, L=Redmond,
S=WA, C=US, E=user1@northwindtraders.com
  Subject: CN=Darren Canavor
  Serial: 61321aaf00000000061e
  Template: EnrollmentAgent
  54 8e 6f e3 d6 88 31 29 b6 f0 fb be bd aa 91 12 76 dd 51 a3
  Element.dwInfoStatus = CERT_TRUST_HAS_EXACT_MATCH_ISSUER (0x1)
  Element.dwInfoStatus = CERT_TRUST_HAS_PREFERRED_ISSUER (0x100)
419.2199.0: 0x80070002 (WIN32: 2)
    CRL 0:
    Issuer: CN=W2K User Enrollment CA, OU=Test, O=NT Distributed Systems,
L=Redmond, S=WA, C=US, E=user1@northwindtraders.com
    4b f3 03 53 8b d0 82 f5 bf aa b7 d1 ee fc aa 3e 12 74 b5 fe
  Application[0] = 1.3.6.1.4.1.311.20.2.1 Certificate Request Agent
CertContext[0][1]: dwInfoStatus=101 dwErrorStatus=0
  Issuer: E=user1@northwindtraders.com, CN=ASIA W2K PCA, OU=Test, O=NT
Distributed Systems, L=Redmond, S=WA, C=US
  Subject: CN=W2K User Enrollment CA, OU=Test, O=NT Distributed Systems,
L=Redmond, S=WA, C=US, E=user1@northwindtraders.com
  Serial: 4860d04700020000bc47
  Template: SubCA
  fa 4b b5 e1 a6 9f 8d e9 1d 69 4b f4 42 9f 76 0b ef a9 c8 d9
  Element.dwInfoStatus = CERT_TRUST_HAS_EXACT_MATCH_ISSUER (0x1)
  Element.dwInfoStatus = CERT_TRUST_HAS_PREFERRED_ISSUER (0x100)
419.2199.0: 0x80070002 (WIN32: 2)
    CRL 0:
    Issuer: E=user1@northwindtraders.com, CN=ASIA W2K PCA, OU=Test, O=NT
Distributed Systems, L=Redmond, S=WA, C=US
    a7 a0 41 a8 90 71 0f 02 60 b1 28 bf 47 3b 4e 48 20 24 58 74
CertContext[0][2]: dwInfoStatus=102 dwErrorStatus=0
  Issuer: CN=ASIA SA Root CA, OU=Asia, O=Northwindtraders, C=US
  Subject: E=user1@northwindtraders.com, CN=ASIA W2K PCA, OU=Test, O=NT
Distributed Systems, L=Redmond, S=WA, C=US
  Serial: 1f54dfa100000000001d
  Template: SubCA
  61 0b 95 b6 06 ba 14 4c ae 89 24 9d 83 fd 06 49 9b ca 82 60
  Element.dwInfoStatus = CERT_TRUST_HAS_KEY_MATCH_ISSUER (0x2)
  Element.dwInfoStatus = CERT_TRUST_HAS_PREFERRED_ISSUER (0x100)
    CRL 49:
    Issuer: CN=ASIA SA Root CA, OU=Asia, O=Northwindtraders, C=US
    78 9e e3 0f 30 ed 2f d5 6e ec b1 9b 59 93 9b b9 b3 36 bb 8e
CertContext[0][3]: dwInfoStatus=10c dwErrorStatus=0
  Issuer: CN=ASIA SA Root CA, OU=Asia, O=Northwindtraders, C=US
  Subject: CN=ASIA SA Root CA, OU=Asia, O=Northwindtraders, C=US
  Serial: 212566f75e7584b8478f7b59b4a9e212
  Template: CA
  9e 90 bb 26 24 e4 da dc 63 11 b8 18 2d af ad 39 56 81 66 51
  Element.dwInfoStatus = CERT_TRUST_HAS_NAME_MATCH_ISSUER (0x4)
  Element.dwInfoStatus = CERT_TRUST_IS_SELF_SIGNED (0x8)
  Element.dwInfoStatus = CERT_TRUST_HAS_PREFERRED_ISSUER (0x100)
    CRL 49:
    Issuer: CN=ASIA SA Root CA, OU=Asia, O=Northwindtraders, C=US
    78 9e e3 0f 30 ed 2f d5 6e ec b1 9b 59 93 9b b9 b3 36 bb 8e
Exclude leaf cert:
  ea 88 66 a8 e8 c0 2d 50 c6 b0 21 a8 4d fb 87 2d 0b 8a da 83
Full chain:
  7c c0 2c 86 ba 49 40 95 45 4c 0c 7f e1 f7 07 d3 88 f1 8d d4
------------------------------------
Verified Issuance Policies: None
Verified Application Policies:
    1.3.6.1.4.1.311.20.2.1 Certificate Request Agent
Signer Info[1]:
Signature matches Public Key
CMSG_SIGNER_INFO_PKCS_1_5_VERSION(1)
CERT_ID_ISSUER_SERIAL_NUMBER(1)
    Serial Number: 61321aaf00000000061e
    Issuer: CN=W2K User Enrollment CA, OU=Test, O=NT Distributed Systems,
L=Redmond, S=WA, C=US, E=user1@northwindtraders.com
```

```
    Subject: CN=Darren Canavor
Hash Algorithm:
    Algorithm ObjectId: 1.3.14.3.2.26 sha1
    Algorithm Parameters: NULL
Encrypted Hash Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.1 RSA
    Algorithm Parameters: NULL
Encrypted Hash:
    0000  91 d8 40 e2 fc d7 9f dd  da 2a 16 ed 6e b4 62 39
    0010  08 c9 0b 08 c6 7a 19 8e  f4 7a af ee a0 c8 e5 5a
    0020  54 90 d3 bb f8 89 cf e8  3f e4 7a 33 45 1e 6b 09
    0030  29 a2 4a 3d e0 28 fe d8  45 15 59 67 74 f4 ab 03
    0040  82 d8 89 11 e6 bd 1a 5f  3b 73 02 a4 f8 be 9f f9
    0050  d2 65 cc 2a b1 47 d4 d1  ce 8f 1d 51 be 5e 5b 92
    0060  a7 79 da 80 4e 5e e5 72  3c 76 84 61 34 d4 42 f2
    0070  da 4d 4b 17 ec 34 53 9b  2c 86 71 60 82 47 54 1e
Authenticated Attributes[1]:
  3 attributes:
  Attribute[0]: 1.2.840.113549.1.9.3 (Content Type)
    Value[0][0]:
    Unknown Attribute type
    1.3.6.1.5.5.7.12.2 CMC Data
  Attribute[1]: 1.2.840.113549.1.9.4 (Message Digest)
    Value[1][0]:
    Unknown Attribute type
    Message Digest:
        5a 74 12 86 78 3a ab 5b 17 85 6f 4d 44 ea a2 74 2c 86 c1 1f
  Attribute[2]: 1.3.6.1.4.1.311.21.20 (Client Information)
    Value[2][0]:
    Unknown Attribute type
    Client Id: = 4
    XECI_CERTREQ -- 4
    User: ASIA\user2
    Machine: user2.asia.northwindtraders.com
    Process: certreq
Unauthenticated Attributes[1]:
  0 attributes:
Computed Hash: 75 80 82 26 ff 11 77 5b 92 52 ce 2e a2 8e a2 32 98 a7 1a a0
No Recipient
Certificates:
================ Begin Nesting Level 1 ================
Element 0:
X509 Certificate:
Version: 3
Serial Number: 4860d04700020000bc47
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.5 sha1RSA
    Algorithm Parameters:
    05 00
Issuer:
    E=user1@northwindtraders.com
    CN=ASIA W2K PCA
    OU=Test
    O=NT Distributed Systems
    L=Redmond
    S=WA
    C=US
NotBefore: 6/26/2001 1:03 PM
NotAfter: 6/7/2002 11:41 AM
Subject:
    CN=W2K User Enrollment CA
    OU=Test
    O=NT Distributed Systems
    L=Redmond
    S=WA
    C=US
    E=user1@northwindtraders.com
Public Key Algorithm:
    Algorithm ObjectId: 1.2.840.10040.4.1 DSA
    Algorithm Parameters:
    0000  30 82 01 1e 02 81 81 00  a1 13 9a 69 07 1f 39 de
    0010  e8 f2 6f 24 0b 80 91 67  a1 f5 b4 d0 e5 24 65 3c
```

```
    0020  2e f1 48 0c ef 3e e9 5d  1c fe 9a 47 1b 3e d3 41
    0030  ba c4 c5 c1 6e 55 9c b2  c4 dc 0b 9a 1a a1 12 f6
    0040  45 a5 32 f6 8d d3 50 79  e1 f0 1e 75 70 6b aa a8
    0050  75 fb 99 bd 74 75 a8 b1  05 b0 a6 fd 4e 20 fa 9d
    0060  9d 5e 79 51 9b 19 b0 c3  62 dd c2 a3 c1 ad af f6
    0070  00 6b f7 70 3a 7e 22 e1  a9 d1 df 2f fc 53 d5 04
    0080  95 b2 6c 7a 0c c4 52 5f  02 15 00 ee 57 07 f1 1f
    0090  44 e7 75 0b d2 a0 f9 65  0a ec b8 dc 7c ad 6d 02
    00a0  81 80 38 5d bc e4 06 1c  6c 16 70 54 7c 3a 65 d0
    00b0  f3 bb 08 83 90 d0 b1 1b  ea 53 90 23 8b b7 2e e2
    00c0  a0 16 b7 11 41 31 20 f2  2c 56 a9 f3 8d 2b e8 74
    00d0  32 c0 7e f4 90 a1 0f 30  c1 5e df e3 c7 a4 20 90
    00e0  73 6a 02 bb eb 46 31 bd  29 70 45 e7 d7 43 22 86
    00f0  55 33 e5 b9 d7 ac 4f 0b  d4 53 5f ec 9c ae 34 0c
    0100  14 35 7e 7f ad 0c 2d 50  4c ea 7d 47 34 1a 19 0b
    0110  63 a3 1a 4a 3a 4d b8 4a  8a 7d b1 36 48 64 d0 f7
    0120  e2 41
Public Key Length: 1024 bits
Public Key: UnusedBits = 0
    0000  02 81 80 6f be 7e 8d a2  4c d2 4a c2 bb dc 71 9c
    0010  85 d5 d3 0c e7 df d0 8e  a5 85 2b d2 5b a4 a1 0d
    0020  a8 55 1b d6 4b 04 2d 56  f8 0a a7 78 8b 1f d1 73
    0030  b7 3e 2a af 1e 21 13 c6  4e 98 ce 88 4c 34 60 d1
    0040  4c a4 80 4e 1c 76 ad 8e  dd 60 6c f0 22 55 47 95
    0050  09 3b 93 75 51 11 eb 7c  74 4a a1 72 2c cf d4 28
    0060  ef 60 f0 8a 18 eb 4a 19  24 93 c0 27 3f af 55 98
    0070  d6 1b 69 63 4a a6 7b f7  69 92 77 4c 28 60 f8 97
    0080  6c a2 d0
Certificate Extensions: 7
    2.5.29.14: Flags = 0, Length = 16
    Subject Key Identifier
        0a 2f e7 66 09 39 b6 c3 9d 96 9a 49 ff 75 73 fd 72 80 92 45
    1.3.6.1.4.1.311.20.2: Flags = 0, Length = c
    Certificate Template Name
        SubCA
    2.5.29.15: Flags = 0, Length = 4
    Key Usage
        Certificate Signing, Off-line CRL Signing, CRL Signing (06)
    2.5.29.19: Flags = 1(Critical), Length = 5
    Basic Constraints
        Subject Type=CA
        Path Length Constraint=None
    2.5.29.35: Flags = 0, Length = 76
    Authority Key Identifier
        KeyID=7a fc 16 de 56 19 08 a3 39 21 5d 55 0f f2 57 be 8f 5e c7 7f
        Certificate Issuer:
             Directory Address:
                  CN=ASIA SA Root CA
                  OU=Asia
                  O=Northwindtraders
                  C=US
        Certificate SerialNumber=1f 54 df a1 00 00 00 00 00 1d
    2.5.29.31: Flags = 0, Length = 11c
    CRL Distribution Points
        [1]CRL Distribution Point
             Distribution Point Name:
                  Full Name:
URL=ldap:///CN=ASIA%20W2K%20PCA,CN=W2KPCA,CN=CDP,CN=Public%20Key%20Services,CN=S
ervices,CN=Configuration,DC=asia,DC=Northwindtraders,DC=com?certificateRevocatio
nList?base?objectclass=cRLDistributionPoint
        [2]CRL Distribution Point
             Distribution Point Name:
                  Full Name:
URL=http://w2kpca.asia.northwindtraders.com/CertEnroll/ASIA%20W2K%20PCA.crl
    1.3.6.1.5.5.7.1.1: Flags = 0, Length = 130
    Authority Information Access
        [1]Authority Info Access
             Access Method=Certification Authority Issuer (1.3.6.1.5.5.7.48.2)
             Alternative Name:
RL=ldap:///CN=ASIA%20W2K%20PCA,CN=AIA,CN=Public%20Key%20Services,CN=Services,CN=
Configuration,DC=asia,DC=Northwindtraders,DC=com?cACertificate?base?objectclass=
certificationAuthority
```

```
        [2]Authority Info Access
             Access Method=Certification Authority Issuer (1.3.6.1.5.5.7.48.2)
             Alternative Name:
URL=http://w2kpca.asia.northwindtraders.com/CertEnroll/W2KPCA.asia.northwindtrad
ers.com_ASIA%20W2K%20PCA(2).crt
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.5 sha1RSA
    Algorithm Parameters:
    05 00
Signature: UnusedBits=0
    0000  70 a2 0e c9 94 30 81 83  43 33 1b b3 6a 83 ae a5
    0010  2a 1f fc c9 d8 28 ae 7e  b3 a2 fb c2 f0 be 8a 3b
    0020  31 e9 bc e2 80 1c 9d 5a  5c 08 db 79 df 7f 88 5e
    0030  7c e2 90 fe 9b e5 83 20  b5 7b 2b 2c 86 bb 22 c4
    0040  4e 9e 09 3c 6a c1 c0 35  57 85 82 a3 07 b4 ca e5
    0050  d8 48 38 33 ec b2 f2 11  56 c4 83 2c 27 2e 9b 19
    0060  7e 9c 62 5a 8f 10 78 16  3c 36 c3 19 fc 7f 63 38
    0070  b8 eb ab 24 7a 45 86 ec  25 b0 b0 63 b0 2e 98 91
    0080  54 13 6d c3 eb d7 c3 42  10 64 49 19 cf 20 d5 55
    0090  18 bb a6 fd 5b ff 12 5a  88 62 22 12 5a 02 1e de
    00a0  e2 21 e8 fa d2 83 6e 13  fd 55 3b ca 8a 56 27 8d
    00b0  28 79 3c 15 df 58 79 4c  ce fb d9 44 d4 fc 7f 6b
    00c0  92 6b 67 3c e6 29 b2 ed  6a 30 0f 89 75 ab 9e 04
    00d0  6b 31 ec e0 79 76 c3 51  cd 91 1e 13 cd 1e 06 8d
    00e0  ce c8 c9 9b cb 14 23 88  ae e0 c3 1f 18 56 ae 55
    00f0  a6 15 c5 95 18 61 5e 65  b6 24 9e c8 ca 87 fe 20
Non-root Certificate
Key Id Hash(sha1): 0a 2f e7 66 09 39 b6 c3 9d 96 9a 49 ff 75 73 fd 72 80 92 45
Cert Hash(md5): d4 8e c5 46 fa 21 77 2a 32 f0 8c 28 78 9a a2 92
Cert Hash(sha1): fa 4b b5 e1 a6 9f 8d e9 1d 69 4b f4 42 9f 76 0b ef a9 c8 d9
----------------  End Nesting Level 1  ----------------
================ Begin Nesting Level 1 ================
Element 1:
X509 Certificate:
Version: 3
Serial Number: 61321aaf00000000061e
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.10040.4.3 sha1DSA
    Algorithm Parameters: NULL
Issuer:
    CN=W2K User Enrollment CA
    OU=Test
    O=NT Distributed Systems
    L=Redmond
    S=WA
    C=US
    E=user1@northwindtraders.com
NotBefore: 10/29/2001 6:16 PM
NotAfter: 6/7/2002 11:41 AM
Subject:
    CN=Darren Canavor
Public Key Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.1 RSA
    Algorithm Parameters:
    05 00
Public Key Length: 1024 bits
Public Key: UnusedBits = 0
    0000  30 81 89 02 81 81 00 eb  8c 61 aa 06 26 3a 0f 74
    0010  bb 03 53 94 8e 8b d7 c9  1c b8 af d9 71 12 3a 07
    0020  96 fb 15 dd 1f ca d1 14  9e 47 ae aa 79 7d f0 ba
    0030  71 3c 02 d7 1a 53 2e fe  56 23 be 64 ee f4 7f 7e
    0040  02 68 22 ab f1 8d 94 f6  4f ee e2 45 f6 0e 5f 34
    0050  dd c9 60 32 f0 fd 55 6b  4f 3d 5a 8d c3 21 97 ba
    0060  a2 6b af 40 b1 ba 59 de  27 15 e4 e4 e3 2f 9f 84
    0070  22 92 29 25 88 42 a5 c9  90 84 2e 46 86 32 21 99
    0080  1f 52 98 5d 79 d7 eb 02  03 01 00 01
Certificate Extensions: 8
    2.5.29.15: Flags = 0, Length = 4
    Key Usage
        Digital Signature (80)
    2.5.29.14: Flags = 0, Length = 16
    Subject Key Identifier
```

```
        3e 09 89 24 b9 c6 f9 68 34 e8 02 1c f5 25 cd 96 6e 13 68 d0
    1.3.6.1.4.1.311.20.2: Flags = 0, Length = 20
    Certificate Template Name
        EnrollmentAgent
    2.5.29.35: Flags = 0, Length = ca
    Authority Key Identifier
        KeyID=0a 2f e7 66 09 39 b6 c3 9d 96 9a 49 ff 75 73 fd 72 80 92 45
        Certificate Issuer:
             Directory Address:
                  E=user1@northwindtraders.com
                  CN=ASIA W2K PCA
                  OU=Test
                  O=NT Distributed Systems
                  L=Redmond
                  S=WA
                  C=US
        Certificate SerialNumber=48 60 d0 47 00 02 00 00 bc 47
    2.5.29.31: Flags = 0, Length = 136
    CRL Distribution Points
        [1]CRL Distribution Point
             Distribution Point Name:
                  Full Name:
URL=ldap:///CN=W2K%20User%20Enrollment%20CA,CN=w2keobca,CN=CDP,CN=Public%20Key%2
0Services,CN=Services,CN=Configuration,DC=asia,DC=Northwindtraders,DC=com?certif
icateRevocationList?base?objectclass=cRLDistributionPoint
        [2]CRL Distribution Point
             Distribution Point Name:
                  Full Name:
URL=http://w2keobca.asia.northwindtraders.com/CertEnroll/W2K%20User%20Enrollment
%20CA.crl
    1.3.6.1.5.5.7.1.1: Flags = 0, Length = 147
    Authority Information Access
        [1]Authority Info Access
             Access Method=Certification Authority Issuer (1.3.6.1.5.5.7.48.2)
             Alternative Name:
URL=ldap:///CN=W2K%20User%20Enrollment%20CA,CN=AIA,CN=Public%20Key%20Services,CN
=Services,CN=Configuration,DC=asia,DC=Northwindtraders,DC=com?cACertificate?base
?objectclass=certificationAuthority
        [2]Authority Info Access
             Access Method=Certification Authority Issuer (1.3.6.1.5.5.7.48.2)
             Alternative Name:
URL=http://w2keobca.asia.northwindtraders.com/CertEnroll/w2keobca.asia.northwind
traders.com_W2K%20User%20Enrollment%20CA.crt
    2.5.29.37: Flags = 0, Length = e
    Enhanced Key Usage
        Certificate Request Agent (1.3.6.1.4.1.311.20.2.1)
    2.5.29.17: Flags = 0, Length = 52
    Subject Alternative Name
        Other Name:
             Principal Name=user2@asia.northwindtraders.com
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.10040.4.3 sha1DSA
    Algorithm Parameters: NULL
Signature: UnusedBits=0
    0000  c2 c8 2e 86 c9 32 ca 80  5c f3 ba 09 08 14 fc 01
    0010  ab 87 1c 34 14 02 32 aa  7c af d0 36 8b df ac 24
    0020  da 2c a6 7e 21 fc f8 49  da 80 00 15 02 2d 30
Non-root Certificate
Key Id Hash(sha1): 3e 09 89 24 b9 c6 f9 68 34 e8 02 1c f5 25 cd 96 6e 13 68 d0
Cert Hash(md5): 59 6d 1c 02 87 8f 91 08 cb 33 82 c1 d2 4a f8 1c
Cert Hash(sha1): 54 8e 6f e3 d6 88 31 29 b6 f0 fb be bd aa 91 12 76 dd 51 a3
----------------  End Nesting Level 1  ----------------
================ Begin Nesting Level 1 ================
Element 2:
X509 Certificate:
Version: 3
Serial Number: 1f54dfa100000000001d
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.5 sha1RSA
    Algorithm Parameters:
    05 00
Issuer:
```

```
    CN=ASIA SA Root CA
    OU=Asia
    O=Northwindtraders
    C=US
NotBefore: 6/7/2001 11:31 AM
NotAfter: 6/7/2002 11:41 AM
Subject:
    E=user1@northwindtraders.com
    CN=ASIA W2K PCA
    OU=Test
    O=NT Distributed Systems
    L=Redmond
    S=WA
    C=US
Public Key Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.1 RSA
    Algorithm Parameters:
    05 00
Public Key Length: 2048 bits
Public Key: UnusedBits = 0
    0000  30 82 01 0a 02 82 01 01  00 c2 0b 38 8e 77 00 c5
    0010  ec 23 62 2f 1f 5a bd f9  72 5f f4 71 25 87 38 ca
    0020  1d 8c 1c 07 cb 02 4e 12  43 7a d9 fe db 40 0f b1
    0030  ca 98 72 b0 d9 fb 10 37  42 84 72 24 0f 26 36 f8
    0040  ab a7 5b 44 a1 0f d7 18  f5 57 d6 e3 79 36 c5 cc
    0050  bc af 47 ae b7 5e 0c 0f  3d 76 e7 06 84 af c1 2e
    0060  99 90 e4 82 3f 20 d4 d6  bf 0b 37 9e 7f 31 e4 38
    0070  b9 32 46 0b 94 0f b3 b9  55 cb f5 03 c9 c4 3b 1c
    0080  6a 3b 78 77 06 71 f1 16  ed f9 a6 4b 94 35 00 a4
    0090  9a 13 fe fc 7b 7a 8f cd  6e 9b 87 8d de 19 8e 06
    00a0  88 ce b5 04 e4 fd 2a 50  a7 1e d5 7a d2 80 f1 e5
    00b0  3f 08 2e 55 5e 05 57 97  0e d6 13 d8 6c 16 7d 5e
    00c0  10 65 4e 2a 44 cc 5d f9  3d 52 9c d1 1e 15 e0 4d
    00d0  a4 ec a1 0f 2f 5a e6 29  d5 4e 45 04 09 fc 45 0e
    00e0  11 0f d0 fa d5 8d 0c 41  0d fd 79 69 e2 2a 09 f7
    00f0  92 cd 2d fe 4d 61 13 b9  b9 f4 06 fb 78 9a e6 7c
    0100  19 e4 1f 22 64 81 89 c7  29 02 03 01 00 01
Certificate Extensions: 8
    1.3.6.1.4.1.311.21.1: Flags = 0, Length = 3
    CA Version
        V2.0
    2.5.29.14: Flags = 0, Length = 16
    Subject Key Identifier
        7a fc 16 de 56 19 08 a3 39 21 5d 55 0f f2 57 be 8f 5e c7 7f
    1.3.6.1.4.1.311.20.2: Flags = 0, Length = c
    Certificate Template Name
        SubCA
    2.5.29.15: Flags = 0, Length = 4
    Key Usage
        Digital Signature, Non-Repudiation, Certificate Signing, Off-line CRL
Signing, CRL Signing (c6)
    2.5.29.19: Flags = 1(Critical), Length = 5
    Basic Constraints
        Subject Type=CA
        Path Length Constraint=None
    2.5.29.35: Flags = 0, Length = 18
    Authority Key Identifier
        KeyID=77 c9 74 69 2c 39 fe 38 65 f4 87 05 58 08 ce bd ba 97 da 10
    2.5.29.31: Flags = 0, Length = 131
    CRL Distribution Points
        [1]CRL Distribution Point
             Distribution Point Name:
                  Full Name:
URL=http://whicasarootca.asia.northwindtraders.com/CertEnroll/ASIA%20SA%20Root%20CA.crl
URL=ldap:///CN=ASIA%20SA%20Root%20CA,CN=whicasarootca,CN=CDP,CN=Public%20Key%20S
ervices,CN=Services,CN=Configuration,DC=asia,DC=Northwindtraders,DC=com?certific
ateRevocationList?base?objectClass=cRLDistributionPoint
    1.3.6.1.5.5.7.1.1: Flags = 0, Length = 145
    Authority Information Access
        [1]Authority Info Access
             Access Method=Certification Authority Issuer (1.3.6.1.5.5.7.48.2)
             Alternative Name:
```

```
URL=http://whicasarootca.asia.northwindtraders.com/CertEnroll/whicasarootca.asia
.northwindtraders.com_ASIA%20SA%20Root%20CA.crt
        [2]Authority Info Access
             Access Method=Certification Authority Issuer (1.3.6.1.5.5.7.48.2)
             Alternative Name:
URL=ldap:///CN=ASIA%20SA%20Root%20CA,CN=AIA,CN=Public%20Key%20Services,CN=Servic
es,CN=Configuration,DC=asia,DC=Northwindtraders,DC=com?cACertificate?base?object
Class=certificationAuthority
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.5 sha1RSA
    Algorithm Parameters:
    05 00
Signature: UnusedBits=0
    0000  1f 1b 35 ad 28 ce 75 25  8b 26 18 19 8a 38 60 1c
    0010  95 f6 bf d1 fb de 61 76  ba 24 71 97 f6 1d 48 92
    0020  df 11 36 f8 40 de 58 20  b1 6a 55 ac 27 f9 8b f7
    0030  c2 b6 ca 76 18 8a 47 69  39 28 e0 fd 81 98 3d 07
    0040  df 6f 01 12 76 c3 5b 2a  9b 42 d9 b5 9c 40 fd 15
    0050  0b 4a 9c 5f 88 17 f7 3a  b8 42 90 58 19 88 10 4d
    0060  4f 53 cf d8 29 1b 3c 5b  c9 c0 f2 ad 13 61 b0 e7
    0070  70 b5 25 df 15 0c 36 2a  50 95 b6 8f b7 1d d5 6e
Non-root Certificate
Key Id Hash(sha1): 7a fc 16 de 56 19 08 a3 39 21 5d 55 0f f2 57 be 8f 5e c7 7f
Cert Hash(md5): f2 bf 51 9f 3a d7 37 ec 03 20 79 b5 69 17 c4 26
Cert Hash(sha1): 61 0b 95 b6 06 ba 14 4c ae 89 24 9d 83 fd 06 49 9b ca 82 60
----------------  End Nesting Level 1  ----------------
================ Begin Nesting Level 1 ================
Element 3:
X509 Certificate:
Version: 3
Serial Number: 212566f75e7584b8478f7b59b4a9e212
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.5 sha1RSA
    Algorithm Parameters:
    05 00
Issuer:
    CN=ASIA SA Root CA
    OU=Asia
    O=Northwindtraders
    C=US
NotBefore: 9/20/2000 1:24 PM
NotAfter: 9/20/2002 1:33 PM
Subject:
    CN=ASIA SA Root CA
    OU=Asia
    O=Northwindtraders
    C=US
Public Key Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.1 RSA
    Algorithm Parameters:
    05 00
Public Key Length: 1024 bits
Public Key: UnusedBits = 0
    0000  30 81 89 02 81 81 00 c7  29 b7 6c 1b 49 f7 77 a9
    0010  f5 83 3d 78 5b 6b 25 29  85 03 c3 46 e8 eb 71 4c
    0020  a4 4b 2f 2a 2b 5c c6 0d  53 32 ec 76 8c ef 19 67
    0030  52 67 09 73 6e f0 13 6a  4c eb ce b8 ae aa ae d0
    0040  81 a0 73 26 f4 b4 3a af  32 03 3b 61 a9 fd 23 05
    0050  0c ac 1a f4 c7 d4 b1 e2  7a 8d db 98 21 45 38 e5
    0060  2d 1a f7 dd 24 66 c4 32  f4 db f1 c4 f4 cb 10 20
    0070  3c 9e ce af 45 99 b5 ae  fb 7f f0 11 50 d5 96 bf
    0080  a8 3b 4c d5 14 85 ed 02  03 01 00 01
Certificate Extensions: 6
    1.3.6.1.4.1.311.20.2: Flags = 0, Length = 6
    Certificate Template Name
        CA
    2.5.29.15: Flags = 0, Length = 4
    Key Usage
        Non-Repudiation, Certificate Signing, Off-line CRL Signing, CRL Signing (46)
    2.5.29.19: Flags = 1(Critical), Length = 5
    Basic Constraints
        Subject Type=CA
```

```
        Path Length Constraint=None
    2.5.29.14: Flags = 0, Length = 16
    Subject Key Identifier
        77 c9 74 69 2c 39 fe 38 65 f4 87 05 58 08 ce bd ba 97 da 10
    2.5.29.31: Flags = 0, Length = ae
    CRL Distribution Points
        [1]CRL Distribution Point
             Distribution Point Name:
                  Full Name:
URL=http://whicasarootca.asia.northwindtraders.com/CertEnroll/ASIA%20SA%20Root%20CA.crl
URL=file://\\whicasarootca.asia.northwindtraders.com\CertEnroll\ASIA%20SA%20Root%20CA.c
rl
    1.3.6.1.4.1.311.21.1: Flags = 0, Length = 3
    CA Version
        V0.0
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.5 sha1RSA
    Algorithm Parameters:
    05 00
Signature: UnusedBits=0
    0000  9d 64 9e 14 7c 07 32 06  f7 86 8b a6 fc b9 52 74
    0010  31 35 ab 30 98 ee b5 d7  7d 1c 8a 3d f7 a4 89 e2
    0020  2c f2 cc f9 ad 93 66 29  95 42 a8 77 a8 1b 7c 1c
    0030  4a 4b 25 b1 68 3f 1e db  47 2c e6 46 dd fd c9 b3
    0040  28 8a 55 14 c1 a6 64 9d  64 46 90 82 9a 73 55 85
    0050  2e 6e 5d ff 19 2c 95 18  fa a1 dc e3 b8 54 bc 9a
    0060  c3 3c 1b a7 e0 51 b1 90  3d a7 3b de e3 e8 55 a0
    0070  54 40 a4 90 04 37 ff f6  ac a8 cd 24 6f e1 f9 08
Signature matches Public Key
Root Certificate: Subject matches Issuer
Key Id Hash(sha1): 77 c9 74 69 2c 39 fe 38 65 f4 87 05 58 08 ce bd ba 97 da 10
Cert Hash(md5): 51 66 26 77 89 a4 3d 07 f7 62 56 d2 de 0e d8 f6
Cert Hash(sha1): 9e 90 bb 26 24 e4 da dc 63 11 b8 18 2d af ad 39 56 81 66 51
----------------  End Nesting Level 1  ----------------
No CRLs
CertUtil: -dump command completed successfully.
```

This is a preliminary document and may be changed substantially prior to final commercial release of the software described herein. The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

2003 Microsoft Corporation. All rights reserved. Microsoft, Windows, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Other product and company names mentioned herein may be the trademarks of their respective owners.

Microsoft Corporation One Microsoft Way Redmond, WA 98052-6399 USA

# Planning and Implementing Federated Forests in Windows Server 2003

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This paper describes basic concepts that involve forest trusts, common scenarios where you can deploy forest trusts, and the technologies that are included with Windows Server 2003 that enable multiple forest deployments.

**In This White Paper**

Introduction

Scope

Overview

Scope

Multiple Forest Deployment Scenarios

Technologies for Federating Multiple Forests

Considerations for Deploying Forest Trust

Enabling Multiple Forest Scenarios in Windows Server 2003

Walkthroughs

Appendix A: List of Ports

Appendix B: Terms That Are Used in this White Paper

# Introduction

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The concept of the Active Directory forest was introduced in versions of Microsoft Windows 2000, and it was presumed that most organizations would deploy a single forest that spanned the entire organization. However, many organizations deploy multiple forests. In addition, collaboration might be required across forests in different organizations. When this occurs, establishing an external trust relationship requires an enormous amount of management and does not use the newest technological advances.

The concept of the forest trust is introduced in Windows Server 2003 to make multiple forest deployments easier. The forest trust allows administrators to federate two Active Directory forests with a single trust relationship to provide a seamless authentication and authorization experience across the forests. In this white paper, common scenarios where you can deploy forest trusts, basic concepts that involve forest trusts, and the technologies that are included with Windows Server 2003 that enable these multiple forest deployments are described.

| **Note** |
| --- |
| This document refers to features that are included with Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. These features are not included on computers that are running the Windows Server 2003, Web Edition operating system. |

# Scope

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The scope of this white paper is restricted to concepts about authentication, authorization, and trust management across multiple forests. Although the basic reasons why you may want to deploy multiple forests are explained, this white paper is not intended to be a multiple forest deployment decision guide. Also, this white paper does not discuss how to synchronize Microsoft Outlook address book contacts or Microsoft Exchange Server public folders. This paper does, however, provide information about other technologies, such as Microsoft Metadirectory Services (MMS) and Exchange Server public folder replication tools.

# Overview

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

In Windows 2000, the concept of an Active Directory forest that integrates traditional Windows NT Server 4.0 domains is introduced. This integration allows you to enable rich collaboration between domains in the forest. Applications can also use this collaboration. For example, users can search Exchange Server global address lists to locate other users in the forest so that the users can send mail that is based on a variety of attributes.

When this integration was introduced, many of the Windows NT Server 4.0 boundaries disappeared. Domains were no longer completely isolated from one another. They became interdependent of each other, and moving to a single forest required a certain level of consent and trust between the administrators of different domains.

Where this level of trust cannot be achieved, multiple forests must be deployed. In addition, there are scenarios (such as mergers and acquisitions or collaboration across enterprises), where multiple forest deployment already exists. Because of this, a move to a single forest environment is not possible due to the complexity or a lack of trust between organizations.

Windows 2000 Server does not completely address the multiple forest deployment scenario. To enable authentication, administrators must set up trusts between every domain in one forest to every other domain in the other forest. This could lead to so many trusts that they may be difficult to manage. Also, the trust across forests is unlimited, which means that you can either trust every user in the other forest or you can decide to not trust any user. A more limited approach is needed in the scenario across enterprises, where you may want to allow only a specific set of users to authenticate from the other corporation. Trusts across enterprises most likely also have to span firewalls. Windows 2000 trusts have limited support for spanning firewalls.

In the Windows Server 2003 family, the concept of the forest trust is introduced. This concept enables trust between two forests, so that all of the domains in each forest are part of the trust. The problem of unlimited trust is solved by using the **Selective Authentication** option, which allows only specific users or sets of users to authenticate across the trust. Managing trust across network firewalls is also easier because the administrators can control the specific remote procedure call (RPC) ports that need to be open for the trust to work.

# Multiple Forest Deployment Scenarios

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

You may need multiple forests for various reasons. You may need to have different schemas, political divisions between resource owners, lack of trust in another organization's administrative infrastructure, mergers and acquisitions, geographically distributed business units, and so on.

There are two generic reasons why you may need to deploy multiple forests.

- You may need to deploy multiple forests for administrative autonomy. Separate departments do not trust each others administrators (for political or security reasons) or cannot agree on a common change control policy (for directory schema and configuration).

  **-and-**

- You may need to deploy multiple forests for asset isolation. Governmental agencies and business units that have become conglomerated through mergers and acquisitions sometimes maintain completely separate computing infrastructures for security or budgetary reasons; application service providers create separate infrastructures for outsourcing large business customers. The "Design Considerations for Delegation of Administration in Active Directory" white paper provides the following reason for why you may want to use multiple forests:

  Because data stored in Active Directory and on computers joined to Active Directory cannot be isolated from the service administrators of the directory, the only way for an organization to achieve complete data isolation is to create a separate forest.

  You can view the Microsoft TechNet "Design Considerations for Delegation of Administration in Active Directory" white paper at http://go.microsoft.com/fwlink/?LinkID=91034.

**Multiple Forests in the Same Corporation**

In a single corporation, the preferred approach is to deploy a single forest. However, for the reasons that are listed in the previous section, you may need to deploy multiple forests. In most scenarios, the forests should be connected by using dedicated leased lines or virtual private networks (VPNs), which allow the necessary data traffic to flow between the forests. If there are packet filters or firewalls, they usually have an "allow all, deny explicit" approach. The administrators for the two forests may be the same people, or they may be different people, but they may want to treat users from the other forest in the same way that they treat users from their own forests.

Consider the following example. There are three forests in the Contoso, Inc. Corporation: contoso.com, plant.contoso.com, and hr.contoso.com. The Human Resources (HR) department has its own forest because it has a program that requires schema changes. The IT department of the HR organization wants to retain control over these changes without going through the central IT department. In addition, the hr.contoso.com forest has a test domain that the contoso.com forest wants to exclude from the trust relationships.

The Fabrikam manufacturing plant also needs to run a program that uses Active Directory. In this scenario, they needed a separate forest because the domain controllers have to be placed on the floor at the manufacturing plant; everyone in the plant will have physical access to the domain controllers.



**Figure 1: Multiple Forests in the Same Corporation**

**Multiple Forests in Different Corporations**

Each organization will have its own forests. However, in a business-to-business (B2B) scenario, collaboration between a forest in one organization and a forest in another organization is necessary. In this scenario, trust also needs to be established between these two forests. However, the administrators in the each forest want only the users in the other forest to have access to a limited set of resources. The two forests might be connected through dedicated leased lines or through VPN connections, but only limited traffic is allowed between the two forests. Because of this, the firewalls typically have rules that deny all traffic and allow only the traffic that is explicitly allowed.

In this scenario, Contoso.com and Fabrikam.com have partnered together to jointly market their toothpaste and toothbrush combination product. However, they would like to restrict access to each others project files to the marketing teams from the other side because much of the collaboration is related to marketing. To keep employees aware of the general progress, both companies want to open their internal news Web site to all of the employees in the other corporation.



**Figure 2: Multiple Forests in Different Corporations**

**Perimeter Network Scenario**

An organization might need to publish resources in the perimeter network (also known as DMZ, demilitarized zone, and screened subnet) and run programs that need Active Directory. However, deploying a separate domain of the internal forest on the perimeter network is not recommended, because it becomes an easy staging point from which to compromise the entire internal forest if the external domain is compromised. Because of this, you must deploy a separate forest in the perimeter network if someone wants all of the administrative and security benefits of managing the perimeter network servers in a domain.

In this scenario, the administrators of contoso.com want to have a separate forest, called dmz.contoso.com. They have a Web application that is running in the perimeter network that uses Windows Live ID authentication so that users can gain access to a portal. The administrators also want to allow users in the internal forest to publish resources to the perimeter network forest. Because of this, the users in the internal forest need to be able to gain access to the resources that are in the perimeter network by using their internal forest accounts.




**Figure 3: Perimeter Network Scenario**

# Technologies for Federating Multiple Forests

Updated: July 8, 2014

Applies To: Windows 7, Windows 8, Windows 8.1, Windows Server 2003 with SP1, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Vista, Windows XP

This section describes the technologies that are included in Windows Server (beginning with Windows Server 2003) that address the problems that are associated with multiple forest deployments.

**Forest Trust**

A forest trust makes federated forests possible in Windows Server. A forest trust enables a transitive trust between all of the domains in two forests. Microsoft recommends a forest trust be created between forests rather than an external trust. A forest trust ensures that Kerberos is used whenever possible. Kerberos provides better security and scalability over NTLM.

A forest trust is established between the root domain of both forests. Forest trusts can be one-way or two-way trusts. In Figure 4, Forest A and Forest B have a bidirectional trust between each other. Because of this, users in Forest A can authenticate to resources in Forest B and users in Forest B can authenticate to resources in Forest A. In addition to authentication, the trust enables authorization so that the resource owners in each forest can add principals from the other forest to discretionary access control lists (DACLs) and groups.




**Figure 4: Forest Trust**

**Forest Trust Requirements**

Before you can create a forest trust, you need to verify that all domain controllers in both forests are running Windows Server 2003. You also need to verify that you have the correct Domain Name System (DNS) infrastructure in place and that you have established the appropriate functionality level for Active Directory. This means that the forest has to be raised to the Windows Server 2003 functional level and you cannot add any more Windows 2000 or Windows NT Server 4.0 domain controllers to the forest.

To create a forest trust, the person who will create the trust in Fabrikam and Contoso must have Enterprise Admin privileges in the corresponding forest. Each trust is assigned a password that the administrators in both forests must know. An administrator who has Enterprise Admin privileges in both forests can create the trusts in both Fabrikam and Contoso at one time and, in this scenario, a password that is cryptographically random is automatically generated and written for both forests.

Both Windows NT Server 4.0 and Windows 2000 clients can use the forest trust for authentication. However, Windows NT Server 4.0 clients support only NTLM and they do not support user principal names (UPNs) for logging on to the network.

**External Trusts vs. Forest Trusts**

External trusts are used to enable trusts between two domains that are in different forests. A forest trust is similar to an external trust between the root domains of both forests, but it takes the trust one step further to encompass all of the domains. The basic trust information (for both external trust and forest trust) is represented by a trusted domain object (TDO). The TDO for an external trust contains the name of the other domain with which the trust is established, its domain security identifier (SID), the trust direction, the type of trust, and some other attributes.

When there is a forest trust, the trust is marked as "Forest." The TDO also contains an additional **Forest Trust Information** attribute. The **Forest Trust Information** attribute contains information about all of the domains in the remote forest, the tree names, and any alternate name suffixes. This information is used to route authentication and lookups to the remote forest when needed. This information is propagated to the global catalogs so that any domain controller can look it up. The TDOs are stored in the Domain container that is in Active Directory.

The tree names and alternate name suffixes are stored in TopLevelName (or TLN) records. The DNS name, network basic

input/output system (NetBIOS) name, and SID for every domain is stored in DomainInfo records. Because of this, the forest trust contains enough information to reasonably determine which services or users can reside in the other forest.

The following table shows the key differences between both external and forest trusts.

**Table 1 External vs. Forest Trusts**

| | **External trust** | **Forest trust** |
|---|---|---|
| **TDO** | Yes | Yes |
| **Top level names** | No | Yes |
| **DomainInfo (for every domain)** | No | Yes |
| **Kerberos V5 support** | The following prerequisites must be in place to allow Kerberos to function over an external trust:<br>• The trust has to be created using the fully qualified domain name (FQDN). Kerberos referral fails if the FQDN is missing from the TDO. Windows Server 2003 **Add Trust** wizard does not create trusts with Windows 2000 and newer domains without DNS name resolution. For more information see, DNS and NetBIOS Name Resolution to Create External, Realm, and Forest Trusts (http://go.microsoft.com/fwlink/?LinkID=198352)<br>• User name syntax is UPN and the UPN suffix is resolvable to a DC in DNS (implicit UPN)<br>• UDP 389, UDP/TCP 88, and UDP/TCP 464 (password change requests) ports are open for the domain controllers in the user domain.<br>• The server name in the trusting resource domain has to be the FQDN, and the domain suffix of the server name has to match the AD DS domain's DNS FQDN.<br>• Applications building Service Principal Names need to build three-part SPNs using the service/server@realm syntax<br>• Since it is an external trust, the cross-forest targets only cover the domains that have *direct* trusts. Only these DCs should be used to target Kerberos requests or Kerberos Forest Searches.<br>• There may be additional scenarios, unknown today, where Kerberos over an external trust may fail. As mentioned, Microsoft recommends forest trusts where possible. Forest trusts ensure that Kerberos is used whenever possible. Kerberos provides better security and scalability over NTLM. | Yes |
| **Object picker browsing** | Yes | Yes |
| **NTLM support** | Yes<br>**Note**<br>Interactive logon across external trusts will attempt Kerberos. On Windows XP and Windows Server 2003, NTLM will be tried if Kerberos fails. Windows Vista and newer operating systems will not allow fallback to NTLM for interactive logon over external trusts. | Yes |

| | | |
|---|---|---|
| **UPN logons** | Implicit and explicit UPNs where the suffix matches the AD DS DNS domain name.<br>**Note**<br>Certificate logons may fail if the UPN suffix in the certificate does not match the AD DS domain name. | Implicit and explicit UPNs |

**Routing of Kerberos Authentication**

After you establish the forest trust, Kerberos and NTLM authentication are routed accordingly to the other forest so that they can validate user credentials.

With Kerberos, users receive a Ticket Granting Ticket (TGT) when they log on to the computer. When users need to be authenticated to a resource, the TGT is presented to their domain controller, and then the domain controller returns a service ticket to authenticate. This is accomplished through a Kerberos ticket granting service (TGS) request. If the resource is in another domain, the users' domain controller cannot issue the service ticket. Instead, the domain controller issues a referral for the user to contact the resource domain controller if it has a direct trust to it. If there is no direct trust, the domain controller issues a referral to the next domain controller that is in the trust chain. The final resource domain controller then issues the service ticket to the user. If the resource is in the same forest, the user's domain controller can determine if the resource is in the same forest by querying the global catalog.

If a resource is in another forest that is connected by a forest trust, the global catalog does not have information about all of the resources that are in the other forest. However, the global catalog can determine which forest contains this resource by searching the forest trust information for all of the forest TDOs. After the information is determined, the referrals are issued so that the user reaches the local forest root. The local forest root then refers the user to the root of the trusting forest. After the referrals reach the trusting forest, the user is then authoritatively referred to the resource domain controller to obtain the service ticket.

If a user from a trusted forest is logging onto a computer in the trusting forest, the Kerberos Authentication Service (AS) requests are routed in a similar manner. However, when this occurs, the user does not have to traverse all of the domains that are in the trust path; the user is instead directly referred to the root domains and then to a domain controller that contains the users account that is in the other forest.




**Figure 5: Authentication Process to Another Forest**

Figure 5 shows how the client uses Kerberos to authenticate to a service in the other forest:

1. Clients contacts the domain controller (DC1) to request a service ticket to use the service.

2. The domain controller cannot find the service in its domain, so the domain controller queries a global catalog. After examining the forest trust information, the global catalog sends a hint that the service is probably in the other forest, and the domain controller issues a referral to its parent.

3. The client then contacts the parent root domain controller (DC2) to request a service ticket to use the service. The root domain controller (DC2) sends the client a referral to the root domain controller (DC3) for the other forest.

4. The client then contacts the root domain controller (DC3) that is in the other forest to request a service ticket to use the service. In this example, the root domain controller (DC3) determines that the service is in its own forest after querying the global catalog in its forest, and then the root domain controller (DC3) refers the client to the next domain controller

(DC4) in the trust path.

5. The client then contacts the next domain controller (DC4) to which it is referred (the service domain controller) to request a service ticket. Because this domain controller determines that the service is in its domains, it issues a service ticket to the client.

6. The client then presents this service ticket to the server in a Kerberos KERB_AP_REQ message. The server responds, which establishes a secure communication channel.

**How the Object Picker Tool Works**

You can use the Object Picker command-line tool to add users or groups to DACLs on resources and as members of other groups. For example, if the administrator in the trusting forest needs to add a user from the trusted forest to a group in the trusting forest, the administrator opens the Active Directory Users and Computers Microsoft Management Console (MMC) Snap-in, clicks the desired group, and then clicks **Add members**. The Object Picker tool then prompts the administrator to type the name of the remote forest user. First, Object Picker tries to locate the domain in which the user is located. Because the user is from another forest, the local global catalog cannot find the user name in its forest. The global catalog then examines the forest trust information to determine if the requested object might be in a different forest. If it is possible that the object might be in a different forest, a referral is sent to Object Picker to contact a global catalog in the other forest. The Object Picker tool then contacts the global catalog in the remote forest to authoritatively determine which remote forest domain holds the queried object. After this is determined, Object Picker connects to a domain controller in that domain, and then performs a Lightweight Directory Access Protocol (LDAP) search to retrieve certain attributes of the object, such as its security identifier (SID). After this information is obtained, you can set the DACL or modify group membership.

Object Picker support for selecting objects from a trusted forest is available on computers that are running the Microsoft Windows XP operating system or an operating system that was released later than Windows XP. Note that you have to type the full UPN of the user or group when you use Object Picker to select objects from a trusted forest. Object Picker browsing, which allows you to use wildcard characters in searches, is available on computers that are running Windows Server 2003.

Note that you must have credentials in the remote forest if you want to look up a user from a remote forest. If there is a two-way trust, then the remote forest also trusts the local forest and the default local forest credentials are used. If the local forest trusts only the remote forest, the user in the local forest is prompted for credentials in the remote forest by the Object Picker tool.

| **Note** |
| --- |
| Group membership rules prevent both Universal and Global groups from containing members from another forest. |

A best practice for Active Directory is to have a global group in each domain that contains users in that domain and to add these global groups to a domain local group that is in the resource domain.

**SID Filtering**

When authentication requests come from a trusted forest, you trust that forest to have properly validated the users' credentials. However, you also receive authorization information in the form of a list of security identifiers (SIDs), along with the authentication request. If any of the SIDs correspond to the Enterprise Administrator's SID (or to a similar privileged group or user SID) of the trusting forest, then the user from the trusted forest receives elevated rights to resources. To protect against the possibility of the trusted forest assigning counterfeit SIDs, the operating system always filters the SIDs to verify that they are relative to one of the domain SIDs in the trusted forest. This procedure ensures that the trusted forest issues authorization information only for domains for which it is authoritative. For example, when the domain controller of the root domain of the fabrikam.com forest receives the request from the client, it filters the SIDs in the request to verify that the client is not passing any unauthorized SIDs (see step 4 in Figure 5). SID filtering is automatically enabled across all trust relationships.

SID filtering on forest trusts breaks SID history across forests. If a user from the fabrikam.com forest migrates to the contoso.com forest, the SID history of the users contoso.com account contains the SID of the user's user account in the fabrikam.com forest. When the user tries to gain access to resources that are in the fabrikam.com forest by using the user's contoso.com account, the fabrikam.com account SID is presented with the SIDs from contoso.com. The fabrikam.com root domain domain controller determines that the fabrikam.com account SID that is presented by the user is not authorized to be presented by a user account in the contoso.com forest and filters it out. This breaks SID history because the user does not have access to the user's previous resources.

If the trusted forest administrator can be trusted not to spoof unauthorized SIDs, you can allow SID History to traverse the forest trust. To do this, use the **netdom** command-line tool (Netdom.exe) with the **/enablesidhistory:yes** option. To allow users to access shared resources in their original domain over an external trust, ensure that SID filter quarantining is disabled.

To do this, use the **netdom** tool with the **/quarantine:no** option. The Netdom tool is included with the Windows Server 2003 support tools. You can install Windows Server 2003 support tools from the Support\Tools folder on the Windows Server 2003 CD. To do this, right-click the Suptools.msi file in the Support\Tools folder, and then click **Install**. As an alternative, you can install support tools from the Microsoft Download Center. To download the support tools, see Windows Server 2003 Service Pack 2 32-bit Support Tools (http://go.microsoft.com/fwlink/?LinkId=100114).

For more information about SID filtering and SID spoofing, see Microsoft Help and Support (http://go.microsoft.com/fwlink/?LinkID=69288).

#### Selective Authentication

The **Selective Authentication** option is conceptually similar to the network firewall in the way that it works. In the same way that a network firewall allows only certain network access requests through the firewall, the **Selective Authentication** option allows only specific authentication requests. For more information about how to enable the **Selective Authentication** option, see the "Walkthroughs" section in this document.

#### The Problem of Authenticating Users from a Trusted Forest

When users authenticate from a trusted forest, they receive the Authenticated Users SID in their token. Many of the default rights for users in a forest are granted through the Authenticated Users SID. Because the Authenticated Users group is a computed group and its SID is added on the server to which the user authenticates, you cannot change the membership of the group. Because of this, the users from the other forest receive some default rights to all of the resources that are in the trusting forest that are accessible by the Authenticated Users group after you set up a forest trust. You may not want to do this if the trust is set up only to allow access to a small subset of users who are in the trusted forest.

#### The Solution

The **Selective Authentication** option provides a method that you can use to grant limited rights for authentication requests that come across a trust. When you enable the **Selective Authentication** option, all authentication is examined on the service domain controller. The service domain controller verifies that the user is explicitly allowed to authenticate to the resource before allowing the authentication request through. Because of this, you need to specify which users who come across the trust can authenticate to which resources in the domain when you enable the **Selective Authentication** option across a trust. You can do this if you set up the "Allowed to Authenticate" control access right on an object for that particular user or group from the other forest or domain. When a user authenticates across a trust with the **Selective Authentication** option enabled, a special "Other Organization" SID is added to the users authorization data. The presence of this SID triggers a verification on the service domain to ensure that the user is allowed to authenticate to the particular service. After the user is authenticated, the server to which the user authenticates adds another SID, the "This Organization" SID. (The "This Organization" SID is added if the "Other Organization" SID is not already present.) For example, if the user did not authenticate across a trust with the **Selective Authentication** option enabled, only then is the "This Organization" token present. Because of this, only one of these special SIDs can be present in an authenticated user's context.

#### Trust Port Configurations

Trusts need to be deployed across various network boundaries. Because of this, the trusts may have to span one or more firewalls. When this is required, you can either tunnel the traffic across a firewall or open specific ports in the firewall to allow the trust traffic to pass through. In Windows Server 2003, this procedure is simplified through configurable remote procedure call (RPC) ports for the main trust services. The two main configurable RPC ports are:

- The Local Security Authority (LSA) RPC port. This port is used for trust creation and other access to the LSA policy database: **TCP/IPPort** entry in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** registry key.

  **-and-**

- The Netlogon RPC port. This port is used for NTLM and secure channel: **DCTcpipPort** entry in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters** registry key.

  To view a full description of the ports that you can open for different scenarios, see the "List of Ports" section in Appendix A in this document.

# Considerations for Deploying Forest Trust

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section discusses the various considerations and best practices to keep in mind when you deploy forest trusts.

**UPN Considerations**
When you deploy forest trusts, you can use explicit UPN suffixes across multiple forests. You could not use explicit UPN suffixes across multiple forests with external trusts. Administrators who want to use explicit UPN suffixes should carefully plan their UPN structure across forests. UPN suffixes across forests should be unique and should not overlap.

You can use overlapping UPN suffixes, such as plant.contoso.com and contoso.com, but you must also use TopLevelName exclusion (TLNEx) records. This is because the TopLevelName record for contoso.com (UPN suffixes and tree names of trusted or trusting forests are stored as TopLevelName records in the local forests) claims all of the namespaces that are under contoso.com, including plant.contoso.com. Because of this, there is a conflict about where to route a request for a name, such as server.plant.contoso.com. To ensure that this request is sent to the correct forest, all of the forests that have a trust to contoso.com should add a TopLevelName exclusion record for plant.contoso.com to that trust. The TopLevelName exclusion records specify that requests for names that are under plant.contoso.com should not be routed to that forest. TopLevelName exclusion records add an additional level of complexity, so it is recommended that administrators choose UPN suffixes that do not overlap.

Shared UPN suffixes, such as two forests that both claim contoso.com, are not permitted when you use forest trusts. Administrators who want to use shared UPN suffixes only for logons in their forest (users can use the UPN suffix only to logon to the same forest that their account is in) have to remove the definition of the suffix across the forest and define only the UPN attribute on the user object. This means that you cannot set the UPN suffix on the User object through the Active Directory Users and Computers MMC Snap-in user interface (UI). Because of this, you have to programmatically write the attribute to the user object. When you remove the definition of the UPN suffix across the forest by using the Active Directory Users and Computers MMC Snap-in, the UPN suffix is not propagated to another forest that shares the suffix and, therefore, does not result in a collision. (Note that a collision is when two forests clam the same TLN.) If a collision does occur, the forest trust is disabled. Note that you cannot set an exclusion record if both forests claim the same namespace (for example, if two forests both claim contoso.com). With this approach, you cannot log on across forests with the shared UPN. Because of this, you should use this approach only when the occurrence of logons to computers that are in another forest (a forest other than the forest in which the user account is located) are minimal. Note that you can seamlessly gain access to resources in another forest after you log onto a computer in the forest to which you are a member by using the shared UPN suffix.

**Considerations for Restricting the Scope of a Forest Trust**
There are two methods that you can use to restrict which users are permitted to gain access to the forest from the trusted forest:

- You can disable the corresponding DomainInfo record for the domain or the TopLevelName record for the tree in the UI. This method is useful when only a small part of the other forest is not trusted. Note that only authentication requests from users in that domain are disabled when you disable a DomainInfo record. When you disable a DomainInfo record, authentication requests are not disabled if those authentication requests are received from users who are in the local forest if those users want to gain access to resources that are in the disabled domain.

  **-or-**

- You can use the **Selective Authentication** option. This option is useful only when you want to allow access to a small number of users. If there is a large number of users who want to gain access to a large set of resources, it can be difficult to manage the requests to grant the Allowed to Authenticate right to users. Note that only users who have Write_DAC permissions to the resource object in Active Directory can set this permission. This permission is restricted to both the Domain Admins and delegated administrators who created the object. Note that when a user joins a computer to the domain by using the Add workstation to domain privilege, the user does not receive the Write_DAC permission and, because of this, the user cannot grant users who are in other forests the Allowed to Authenticate right on the computer that the user joined.

**Smart Card Logon Considerations**
For smart card logons to work across forests, users must have a public key infrastructure (PKI) trust and a forest trust. This means that the forest to which the user is logging on must trust the Certification Authority (CA) that issued the users smart

card certificate for smart card logon. In addition, the computer to which the user is logging on must trust the computer certificate for the domain controller of the users domain for server authentication. You can do this by either making the CA a trusted root CA, or by qualified subordination.

**User Experience Changes in Moving to Forest Trusts from External Trusts**

There may be certain changes in the user experience after the move to forest trusts. To minimize these changes and ensure that there is a consistent experience, you may want to upgrade to the latest service pack on computers that are running Windows XP, Windows 2000, or Windows NT Server 4.0; you may also want to deploy the appropriate hotfixes. The following list describes the impact of the service packs and hotfixes:

- Windows logons and authentication: When users log on to computers that are joined to a forest other than the forest where their user account is located, they do not see their account domain listed in the logon dialog box. The user has to type their user principal name (UPN) in the dialog box. This could be the implicit UPN suffix (*user_ID@DNS_domain_name.com*) or a UPN suffix that is explicitly defined by the administrator, such as someone@fabrikam.com. If you want to log on with the Windows NT Server 4.0 style name, you should apply the corresponding hotfix to the computer. Note that domains that are in other forests are not displayed in the logon box; you have to type the name as *domain\user* in the logon box. If your users are not currently using UPNs to log on, have them use the Windows NT Server 4.0 style so that the user logon experience is consistent.

- Authorization: On computers that are not running either Windows Server 2003 or Windows XP Service Pack 2 (SP2), you cannot browse principals in another forest to add them to DACLs and groups. Instead, you can type the UPN or Windows NT Server 4.0 name for the principal that is in the other forest. On computers that are running either Windows Server 2003 or Windows XP SP2, you can browse and search principals that are in the other forest. You must run either Windows 2000 Service Pack 4 (or the hotfix that is specified below) or Windows XP to resolve principals in the Object Picker tool that are in the other forest by using UPN names or Windows NT Server 4.0 names. On computers that are running Windows NT Server 4.0, you can type the Windows NT Server 4.0 style name for the user who is in the other forest. On computers that are running Windows 95, Windows 98, or Windows Millennium Edition (Me), you can choose the domain from which you want to browse for users, and then you can select the user. Microsoft Exchange Server 5.5 administrators and Microsoft SQL Server 2000 administrators have to explicitly type the Windows NT Server 4.0 style name when they want to select Windows users so that they can add the users, associate the users with a mailbox, and then enable SQL server logon. Note that you cannot browse for the users who are in the other forest.

The following tables describe the general user experience when they move to forest trust from external trust.

**Logons and Authentication**

| | **Windows 95, Windows 98, Windows Me** | **Windows NT Server 4.0** | **Windows 2000** | **Windows XP** | **Windows Server 2003** |
|---|---|---|---|---|---|
| **Logon box** | No change | Add the hotfix that enables the Windows NT Server 4.0 style name | Logon with the UPN name and add the hotfix that enables the Windows NT Server 4.0 style name | Logon with the UPN name and add the hotfix that enables the Windows NT Server 4.0 style name | Logon with the UPN name and add the hotfix that enables the Windows NT Server 4.0 style name |
| **Shares (UI)** | No change | No change | No change | No change | No change |
| **Shares (command prompt)** | No change | No change | No change | No change | No change |
| **Microsoft Internet Explorer** | No change | No change | No change | No change | No change |
| **Outlook** | No change | No change | No change | No change | No change |

| | | | | | |
|---|---|---|---|---|---|
| **Microsoft Internet Information Services (IIS)** | Not available | No change | No change | Not available | No change |
| **Exchange Server 5.5** | Not available | No change | No change | Not available | Not available |
| **Exchange 2000 Server** | Not available | Not available | No change | Not available | Not available |
| **Microsoft Exchange Server 2003** | Not available | Not available | No change | Not available | No change |
| **Domain DFS in a trusted forest** | Not available | Not available | Not available | Not available | No change |
| **Windows SharePoint Services** | Not available | Not available | No change | Not available | No change |
| **Microsoft Office SharePoint Server 2007** | Not available | Not available | No change | Not available | No change |
| **Microsoft SQL Server 7.0** | Not available | No change | No change | Not available | Not available |
| **SQL Server 2000** | Not available | No change | No change | Not available | No change |
| **Group Policy** | Not available | Not available | No change | No change | No change |

**Lookup and Authorization**

| | **Windows 95, Windows 98, Windows Me** | **Windows NT Server 4.0** | **Windows 2000** | **Windows XP** | **Windows Server 2003** |
|---|---|---|---|---|---|
| **Object Picker tool** | Not available | No change | Upgrade to the latest service packs and apply the latest hotfixes, and then type the UPN or Windows NT Server 4.0 name | Type the UPN or Windows NT Server 4.0 name, upgrade to SP2, and then browse | No change |
| **Shares (UI)** | Need to change the domain from which the users can be selected | No change | No change | No change | No change |
| **Shares (command prompt)** | Not available | No change | No change | No change | No change |
| **IIS** | Not available | No change | No change | Not available | No change |

| | | | | | |
|---|---|---|---|---|---|
| **Exchange Server 5.5** | Not available | Need to type names | Need to type names | Not available | Not available |
| **Exchange 2000 Server** | Not available | Not available | No change | Not available | Not available |
| **Exchange Titanium** | Not available | Not available | No change | Not available | No change |
| **Windows SharePoint Services** | Not available | N/A | No change | Not available | No change |
| **Microsoft Office SharePoint Server 2007** | Not available | N/A | No change when you upgrade to the latest service packs and apply the latest hotfixes, or when you apply the SharePoint Portal Server hotfix | Not available | No change |
| **SQL Server 7.0** | Not available | No change | No change | Not available | Not available |
| **SQL Server 2000** | Not available | Need to type names | Need to type names | Not available | No change |

# Enabling Multiple Forest Scenarios in Windows Server 2003

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The Windows Server 2003 technologies that are discussed in the previous section each help to enable a specific functionality across multiple forests. Note that these technologies are limited to trust management, authentication, and authorization problems.

You must use other technologies and tools to address other parts of the problem, such as synchronization of address books, user migration, DNS configuration, and so on. The following section describes how you can use each of the technologies that are discussed in this paper to enable the scenarios for multiple forests that were discussed earlier in this document. It also provides pointers to the other tools that you must use to solve the other parts of the problem.

**Multiple Forests That Are in the Same Corporation**

The first step that you should use when you want to deploy multiple forests in the same corporation is to ensure that you can perform DNS lookups from one forest to look up the domain controllers that are in the other forest. If the forests already use the same DNS infrastructure, then you do not need to perform any of the additional work. If the forests do not use the same DNS infrastructure, you can either merge the DNS infrastructures or you can use conditional forwarding.

After you configure DNS, set up a bidirectional forest trust between the contoso.com and hr.contoso.com forests. Because plant.contoso.com does not contain any user accounts that need to gain access to resources that are in the other forest, a trust that is one way is set up by the administrator to each of the other forests. You should not enable the **Selective Authentication** option in this scenario because there are a considerable number of users in each of the forests who need to gain access to resources that are in the other forests, and because the administrators for each forest feel more comfortable treating the users who are in the other forest as authenticated users. In addition, the users who are in the contoso.com forest should be able to look up users who are in the hr.contoso.com forest when they use the Exchange Server address books for sending e-mail messages. You can do this by using Microsoft Metadirectory Services (MMS) to synchronize the address books for both forests.

Because the hr.contoso.com forest trusts the plant.contoso.com forest for the plant.contoso.com TopLevelName and because the hr.contoso.com forest also trusts the contoso.com forest for the contoso.com TopLevelName, there is a namespace collision because the TopLevelName record for contoso.com also includes plant.contoso.com. To remove this collision, you must exclude the plant.contoso.com namespace from the contoso.com forest trust. You can do this by setting a TopLevelName exclusion record in the FTInfo record for the hr.contoso.com forest for its trust to contoso.com. Similarly, you must set up another exclusion record in the plant.contoso.com forest to exclude the hr.contoso.com namespace from the trust to the contoso.com forest. However, you do not need to set up an exclusion record in contoso.com because the plant.contoso.com and hr.contoso.com namespaces do not overlap.

In addition, the administrator for the contoso.com forest also does not want to trust the test.hr.contoso.com domain, which is a test domain in the hr.contoso.com forest. Because of this, the administrator can disable the DomainInfo record for the test.hr.contoso.com domain that is in the trust record for hr.contoso.com. When you do this, you ensure that the authentication that is performed by the test.hr.contoso.com domain is not accepted by contoso.com. Note that it is easier for you to use the **Explicit Deny** option that is in the forest trust, instead of the **Selective Authentication** option that is in this scenario because you want to prevent authentication for all principals that are in a specific domain that is in a trusted forest to any resource that is in the trusting forest. Note also that you disable authentication requests that are from users in a domain when you disable the DomainInfo record for a particular domain that is in a trusted forest. When you do this, authentication to resources that are in that domain are not disabled.




**Figure 6: Multiple Forests That Are in the Same Corporation Deployment**

**Multiple Forests That Are in Different Corporations**

Unique challenges are presented when you want to deploy multiple forests that are in different corporations. Not only do the trusts have to traverse firewalls, but you have to ensure that all of the resources are well managed so that users from another forest do not accidentally gain access to resources that they are not supposed to see.

In this scenario, Fabrikam and Contoso establish a forest trust. Each company also enables the **Selective Authentication** option. When the administrators enable the **Selective Authentication** option, the authentications that occur across the trust do not work automatically. The administrators must explicitly enable authentications across the trust. The administrator for the Contoso forest then sets a DACL on the marketing group's file server computer account so that members of the marketing group who are in the Fabrikam forest can authenticate to it. After this procedure, the administrator sets another DACL on the service account for the Web server so that all users who have the Authenticated Users SID can authenticate to it. The administrator of the Contoso forest configures the contoso.com forest in the same way that the Fabrikam administrator configured fabrikam.com. Note that one of these forests may be an extranet forest.

The administrators for both Contoso and Fabrikam create the trust independently by using the Netdom tool (Netdom.exe). After the administrators create the trusts, they try to validate the trust so that the Netlogon fixed port and the RPC Endpointmapper port can be opened (between the domain controllers in each domain) for this operation. The administrators can close the ports after the trust is validated. The administrators for both Contoso and Fabrikam then create the following firewall rules.

| Rule name | Port | Source | Destination |
|---|---|---|---|
| **Kerberos authentication inbound** | 88 (UDP, TCPinbound) | Other corporation's subnet | This corporation's domain controllers |
| **Kerberos authentication outbound** | 88 (UDP,TCPoutbound) | This corporation's subnet | Other corporation's domain controllers |
| **Inbound DACL lookups** | 88 (UDPinbound), 389 (TCP,UDPinbound), 135 (TCPinbound), NTDS fixed port (TCPinbound), Netlogon fixed port (TCPinbound) | Other corporation's subnet | This corporation's domain controllers |
| **Outbound DACL lookups** | 88 (UDPoutbound), 389 (TCP,UDPoutbound), 135 (TCPoutbound), NTDS fixed port (TCPoutbound), Netlogon fixed port (TCPoutbound) | This corporation's subnet | Other corporation's domain controllers |

Firewall rules enable authentication and lookups to work between both the corporations. Note that the administrators can also choose to have an Internet Protocol Security (IPSec) tunnel for the traffic that is described in the previous table by opening only the ports for IPSec. Most of the default Windows services use the negotiate package which first uses Kerberos for authentication and then uses NTLM when certain authentications do not work. If you also need NTLM authentication, you need to open the appropriate ports. These ports are specified in the "List of Ports" section in Appendix A in this document.




**Figure 7: Multiple Forests That Are in Different CorporationsDeployment**

**Perimeter Network Scenario**

In a perimeter network (also known as a DMZ, demilitarized zone, and screened subnet) forest, a trust that is one way from the perimeter network forest to the contoso.com corporate forest is established. This allows the users in the internal forest to be able to gain access to the resources in the perimeter network forest with their internal forest accounts. In this scenario, a two-way trust is not recommended, because compromise of the perimeter network forest will then provide a malicious user with an entry point into the corporate forest. Because of this, the administrator creates the following firewall rules.

| Rule name | Port | Source | Destination |
|---|---|---|---|
| **Trust creation** | 88 (UDP,TCPoutbound), 389 (UDP, TCPoutbound), 445 (TCPoutbound) | Other corporation's subnet | This corporation's domain controllers |
| **Kerberos authentication outbound** | 88 (UDP,TCPoutbound) | This corporation's subnet | Perimeter network domain controllers |
| **Inbound DACL lookups** | 88 (UDP,TCPinbound), 389 (TCP,UDPinbound), 135 (TCPinbound)<br><br>NTDS fixed port (TCPinbound), Netlogon fixed port (TCPinbound) | Perimeter network subnet | Internal domain controllers |

In this table:

- The "trust creation" rule needs to be applied by the administrator only when the trust is initially created. This rule allows trust creation for both domains from the internal forest and is immediately disabled after you create the trust.
- The "Kerberos authentication outbound rule" is always applied by the administrator to allow internal users to authenticate to the perimeter network forest.
- The "inbound DACL lookups" rule is applied by the administrator only when there needs to a change to the domain local groups in the perimeter network to add either global or universal groups from the internal forest. These changes should be minimal.




**Figure 8: Perimeter Network ScenarioDeployment**

# Walkthroughs

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The following sections provided the detailed implementation steps that you can use to set up federated forests.

**Walkthrough Setup**
The structure of the forests that are described in this walkthrough are detailed in the following figure.




**Figure 9: Walkthrough Infrastructure**

Each domain has one domain controller and the names of the domain controllers are CORP-DC-01, NW-DC-01, and MARKETING-DC-01. Both the CORP-DC-01 and NW-DC-01 domain controllers host the Active Directory-integrated DNS server for their respective forests.

**Creating a Forest Trust**
Before you begin this process, ensure that all of the domain controllers are running Windows Server 2003. To create a forest trust, use the methods in this section to perform the following tasks:

- Prepare both forests for the trust
    - Configure DNS
    - Raise the forest and all of the domains to the Windows Server 2003 functional level
- Create the forest trust on contoso.com
- Create the forest trust on nwtrader.com

**Configure DNS**
This section provides the steps that you can use to configure DNS for a forest trust.

**Scenario**
The administrators for both the Northwind Traders and Contoso corporations want to establish network connectivity between the two forests by configuring DNS. The Active Directory-integrated DNS zones for each forest are located in the respective root domains.

To configure DNS:

1. Log on to the corp.contoso.com domain with administrative privileges.
2. Configure DNS:
    - Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **DNS**.
    - Right-click **CORP-DC-01** and then click **Properties**.

- Click **New** on the **Forwarders** tab, type **nwtraders.com**, and then click **OK**.
- Type the IP address of the Northwind Traders DNS server (for example, type **10.1.1.2**), and then click **Add**.

3. Verify connectivity:

   - Click **Start**, click **Run**, type **cmd** in the **Open** box, and then press ENTER.
   - Type **ping nwtraders.com** at a command prompt, and then press ENTER.
   - You receive a reply.

**Raise All of the Domains and Forests to the Windows Server 2003 Functional Level**

This section describes how to raise the forest functional level to Windows Server 2003, which is required for forest trust to be established. To do this, you must first raise the domain functional level for each domain that is in the forest to the Windows Server 2003 functional level. After you raise all of the domains in the forest to the Windows Server 2003 functional level, you can raise the forest to the Windows Server 2003 functional level.

**Scenario**

The administrators for the Northwind Traders and Contoso corporations want to raise their forests to the Windows Server 2003 functional level to establish a forest trust.

To raise all of the domains and the forest to the Windows Server 2003 functional level:

**On the domain controller in the corp.contoso.com domain:**

1. Set the corp.contoso.com domain functional level to Windows Server 2003 by using the Active Directory Domains and Trusts MMC Snap-in:

   - Log on to the corp.contoso.com domain with administrative privileges.
   - Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.
   - Right-click **corp.contoso.com** and then click **Raise Domain Functional Level**.
   - In the **Select an available domain functionality** list, click **Windows Server 2003**, click **Raise**, click **OK**, and then click **OK**.

2. Set the forest functional level to Windows Server 2003 by using the Active Directory Domains and Trusts MMC Snap-in:

   - Log on to the corp.contoso.com domain with administrative privileges.
   - Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.
   - Right-click **marketing.contoso.com**.
   - Click **Raise Domain Functional Level**.
   - In the **Select an available domain functionality** list, click **Windows Server 2003**, click **Raise**, click **OK**, and then click **OK**.

3. Set the forest functional level to Windows Server 2003 by using the Active Directory Domains and Trusts MMC Snap-in:

- Log on to the corp.contoso.com domain with administrative privileges.
- Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.
- Right-click **Active Directory Domains and Trusts** and then click **Raise Forest Functional Level**.
- In the **Select an available domain functionality** list, click **Windows Server 2003**, click **Raise**, click **OK**, and then click **OK**.

**On a domain controller in the nwtraders.com domain:**

1. Set the nwtraders.com domain functional level to Windows Server 2003 by using the Active Directory Domains and Trusts MMC Snap-in:

   - Log on to the nwtraders.com domain with administrative privileges.
   - Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.
   - Right-click **Nwtraders.com** and then click **Raise Domain Functional Level**.
   - In the **Select an available domain functionality** list, click **Windows Server 2003**, click **Raise**, click **OK**, and then click **OK**.

2. Set the forest functional level to Windows Server 2003 by using the Active Directory Domains and Trusts MMC Snap-in:

   - Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.
   - Right-click **Active Directory Domains and Trusts** and then click **Raise Forest Functional Level**.
   - In the **Select an available domain functionality** list, click **Windows Server 2003**, click **Raise**, click **OK**, and then click **OK**.

**Implementing Trusts Across Forests**

This section describes how to implement forest trusts.

**Enabling Authentication and Authorization Across the Forest**

The administrators for both the Northwind Traders and Contoso corporations want to enable authentication and authorization across both forests. Both of the companies want users in their forest to be able to seamlessly gain access to resources in the other forest, without having to type a new user name and password. Because of this, the companies must establish a two-way trust across the forest.

To enable authentication and authorization across the forest:

- Verify connectivity to the Nwtraders.com domain:

  - Log on to the corp.contoso.com domain.
  - Click **Start**, click **Run**, type **cmd** in the **Open** box, and then press ENTER.
  - Type **ping nwtraders.com** at a command prompt, and then press ENTER.
  - You receive a reply.

- Create the forest trust:
    - Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.
    - Right-click **corp.contoso.com** and then click **Properties**.
    - Click the **Trusts** tab, click **New Trust**, and then click **Next** in the Trust Creation Wizard.
    - In the **Name** box, type the name of the forest to which you want to configure the trust (in this sample scenario, type **nwtraders.com**), and then click **Next**.
    - Click **Forest Trust**, and then click **Next**. If **Forest Trust** is not an option, verify that you raised the forest functional level to Windows Server 2003 by reviewing the steps in the previous section.
    - Click **Two Way**, and then click **Next**.
    - Click both the **This Domain** domain and the **Specified Domain** domain, and then click **Next**.
    - In the credentials dialog box for the nwtraders.com domain, type both the user name (**administrator**) and password (**nwtraders**), and then click **Next**.

      Note that you must have trust-creation privileges in the domain to create the trust.
    - Click **Allow authentication for all resources in the local forest**, and then click **Next**.
    - Click **Allow authentication for all resources in the Nwtraders.com forest**, and then click **Next**.

      Note that the **Selective Authentication** option for both sides of the trust is disabled when you do this. You will enable the **Selective Authentication** option in the next section.
    - Review the changes that are listed, and then click **Next** to approve the changes.
    - Click **Yes, confirm outgoing trust**, and then click **Next**.
    - Click **Yes, confirm outgoing trust**, and then click **Next**.
    - When the dialog box that lists the name suffixes that you want to route is displayed, do not make any changes. Click **Next**, click **Finish**, and then click **OK**.
- Create a file share in the Marketing.contoso.com domain, and then assign permissions to the share:
    - On the Marketing-DC-01 server, create a folder named Marketingshare, create a Sampletext.txt file with some text by using a text editor (such as Notepad), and then save the Sampletext.txt file in the C:\Marketingshare folder.
    - Right-click the **MarketingShare** folder and then click **Sharing and Security**.
    - Click **Share this folder**, and then click **Permissions**.
    - Click **Add** in the **Group or user names** box, type **administrator@nwtraders.com**, and then click **OK**.
    - Click **administrator@nwtraders.com** in the **Group or user names** box, and then click to select all of the check boxes in both the **Change** and **Read** boxes.
    - Click **Everyone** in the **Group or user names** box, and then click **Remove**.

Note that you cannot grant permissions by adding the user directly to the DACL file share when you use this procedure; however, you can create a domain local group to grant permission to the share and add the remote forest groups to this domain local group. You will directly add the users to the DACL in this section. More information about group membership rules is provided in the following section.

- Verify that you can gain access to the Marketing-DC-01 domain and the SampleText.txt file that you created:
  - Log on to the NW-DC-01 server with administrative privileges.
  - Click **Start**, click **Run**, type **\\Marketing-DC-01\Marketingshare** in the **Open** box, and then press ENTER.
  - Double-click the **Sampletext.txt** file to confirm that you can open and read the file. If you cannot open the file, verify that the permissions are properly assigned.
  - Create a Sampletext2.txt file in a text editor, such as Notepad, and then save the file to the \\Marketing-DC-01\Marketingshare folder to verify that you can save a file to the share.

**Adding an Alternate Name Suffix**

The administrators for the contoso.com forest want to add an alternate name suffix that corresponds to the users e-mail names. This allows the users to log on to the network with their e-mail names, such as user@e-mail.contoso.com. The suffix must be recognized in the other forest for the users to be able to use it.

Use the following steps to add an alternate name suffix, to add the name suffix to the user's UPN name, and then to enable routing to the alternate name suffix:

- Add an alternate name suffix to the corp.contoso.com domain:
  - Click **Start**, point to **All Programs**, and then click **Administrative Tools**.
  - Right-click **Active Directory Domains and Trusts** and then click **Properties**.
  - Type **e-mail.contoso.com** in the **Alternative UPN suffixes** box, and then click **Add**.
- Add the alternate name suffix to a users UPN in the corp.contoso.com domain:
  - Create a user with a **testuser** user login name in the corp.contoso.com domain. For more information about how to do this, see the Windows 2000 Server home page (http://go.microsoft.com/fwlink/?LinkId=145492).
  - In the Active Directory Users and Computers tool, right-click **testuser**, and then click **Properties**.
  - Click the **Account** tab, type **testuser** in **User logon name** box, and then click **e-mail.contoso.com** in the **Domain** box.
- Enable routing to the alternate name suffix in the nwtraders domain:
  - Log on to the nwtraders.com domain with administrative privileges.
  - Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.
  - Right-click **nwtraders.com** in the left pane, and then click **Properties**.
  - Click the **Trusts** tab, click **corp.contoso.com** in the **Domains trusted by this domain (outgoing trusts)** box, and then click **Properties**.
  - Click the **Name Suffix Routing** tab, click ***.e-mail.contoso.com** in the **Name Suffixes** box that is in the corp.contoso.com forest.

- Click **Enable**, click **OK**, and then click **OK**.

**Excluding a Domain from a Trust Relationship**

The administrators for the nwtraders.com forest no longer want a trust relationship with the marketing.contoso.com subdomain, so the administrators want to exclude that domain from the trust relationship.

To exclude a domain from a trust relationship:

- Create a file share in the nwtraders.com domain:
    - Log on to the nwtraders.com domain with administrative privileges.
    - Click **Start**, point to **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
    - Click the %SystemRoot% folder (the folder where Windows is installed) in the left pane, right-click an empty space in the right pane, point to **New**, click **Folder**, and then type **nwshare** for the folder name.
    - Right-click the **Nwshare** folder that you created in the right pane, and then click **Sharing and Security**.
    - Click **Share this folder**, click **Permissions**, and then click **Add**.
    - Type **administrator@marketing.contoso.com** in the **Enter the object names to select** box, and then click **OK**.
    - Click **Administrator@marketing.contoso.com** in the **Group or user names** box, click **Change** in the **Allow** column in the **Permissions for Administrator@marketing.contoso.com** box, click **Read** in the **Allow** column in the **Permissions for Administrator@marketing.contoso.com** box, and then click **OK**.
    - Click **Everyone** in the **Group or user names** box, click **Remove**, and then click **OK**.
    - Confirm that you can gain access the Nwshare folder from the marketing.contoso.com share while you are logged on to the computer with administrative privileges. To do this, Click **Start**, click **Run**, type **\\NW-DC-01\Nwshare** in the **Open** box, and then press ENTER.
- Disable the DomainInfo record for marketing.contoso.com in the nwtraders.com domain:
    - Log on to the nwtraders.com domain with administrative privileges.
    - Click **Start**, point to **All Programs**, point to **Administrator Tools**, and then click **Active Directory Domains and Trusts**.
    - Right-click **nwtraders.com** in the left pane, and then click **Properties**.
    - Click the **Trusts** tab, click **corp.contoso.com** in the **Domains trusted by this domain (outgoing trusts)** box, and then click **Properties**.
    - Click the **Name Suffix Routing** tab, click **marketing.contoso.com**, and then click **Disable**.

**Setting Up a TopLevelName Exclusion Record**

The administrators for the contoso.com forest want to set up a trust relationship with the nwtraders.com forest for all of the namespaces that are in nwtraders.com, except for plant.nwtraders.com. Note that the plant.nwtraders.com forest is a hypothetical forest that is not used in this sample scenario because it is in a different location and because the administrators for the contoso.com forest want to set up a trust relationship with it.

To set up a TopLevelName exclusion record:

- Log on to the corp.contoso.com domain controller with administrative privileges.
- Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.
- Right-click **corp.contoso.com** in the left pane, and then click **Properties**.
- Click the **Trust** tab, right-click **Nwtraders.com** in the **Domains trusted by this domain (outgoing trusts)** box, and then click **Properties**.
- Click the **Name Suffixes to exclude from routing to nwtraders.com** tab, click **Add**, type ***.plant.nwtraders.com**, click **OK**, and then click **OK**.

**The Selective Authentication Option**

The administrators for the Contoso company have decided that they cannot allow every user in the Northwind Traders forest to authenticate in their forest because of security constraints. The Contoso company administrators want to allow only members of the Enterprise Admins group from the other forest to be able to authenticate, but only to the marketing.contoso.com domain.

To enable the Selective Authentication option:

- Turn on the **Selective Authentication** option in corp.contoso.com to enable only selective authentication from nwtraders.com
    - Confirm that you are logged on to the corp.contoso.com domain with administrative privileges.
    - Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.
    - Right-click **corp.contoso.com** in the left pane, and then click **Properties**.
    - Click the **Trusts** tab, right-click **Nwtraders.com** in the **Domains trusted by this domain (outgoing trusts)** box, and then click **Properties**.
    - Click the **Authentication** tab, click **Allow authentication only to selected resources in the local forest**, click **OK**, and then click **OK**.
- Create a file share in the marketing.contoso.com domain, and then assign permissions to the share:
    - On the Marketing-DC-01 computer, click **Start**, point to **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
    - Click **Local Disk (C:)** in the left pane, right-click a blank area in the right pane, point to **New**, click **Folder**, and then type **Testfolder** for the name of the new folder.
    - Double-click the new **Testfolder** folder in the right pane to open the folder, right-click a blank area, point to **New**, click **Text Document**, and then type **Testdoc.txt** for the name of the document.
    - Right-click the **Testfolder** folder in the left pane, and then click **Sharing and Security**.
    - Click **Share this folder**, click **Permissions**, click **Add**, and then type **administrator@nwtraders.com**.
    - Click **Administrator@nwtraders.com** in the **Group or user names** box, click **Change** in the **Allow** column in the **Permissions for Administrator@ nwtraders.com** box, click **Read** in the **Allow** column in the **Permissions for Administrator@ nwtraders.com** box, and then click **OK**.
    - Click **Everyone** in the **Group or user names** box, and then click **Remove**.

- Verify that you cannot gain access to Marketing.contoso.com from Nwtraders.com:
    - Log on to the NW-DC-01 computer with administrative privileges.
    - Click **Start**, click **Run**, type **\\marketing-dc-01\marketingshare** in the **Open** box, and then press ENTER.

      You should not be able to gain access to the share because you enabled the **Selective Authentication** option. If you can gain access to the share, verify that the permissions are properly configured.
- Enable the **Selective Authentication** for Marketing-DC-01:
    - Log on to the marketing.contoso.com computer with administrative privileges.
    - Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
    - Click **Advanced Features** on the **View** menu, and then click **Domain Controllers** in the left pane.
    - Right-click **MARKETING-DC-01** in the right pane, and then click **Properties**.
    - Click the **Security** tab, click **Add**, type **administrator@nwtraders.com**, and then click **OK**.
    - Click **Administrator@nwtraders.com** in the **Group or user names** box, and then click to select the **Allowed to Authenticate** check box in the **Allow** column that is in the **Permissions for Administrator@nwtraders.com** box.

      After you do this, the administrator@nwtraders.com user can authenticate to the NW-DC-01 computer.
- Verify that you can gain access to marketing.contoso.com from nwtraders.com:
    - Log on to the NW-DC-01 computer with administrative privileges.
    - Click **Start**, click **Run**, type **\\marketing-dc-01\marketingshare** in the **Open** box, and then press ENTER.

      You can now gain access to the share.

**Configuring ISA Server 2000 to Enable Forest Trust Across Firewalls**

This section describes how to configure Microsoft Internet Security and Acceleration (ISA) Server 2000 to enable trust and authentication across the firewall.

**Setting Up the Trust**

It is recommended that you set up the trust for both forests from the internal computer because this method requires connections from only the internal network of the firewall to the outside when you set up the trust. To do this, start the ISA Server Firewall client to allow traffic from the internal domain controller to the external domain controller on both the LDAP (389 UDP and TCP) and Microsoft SMB (445 TCP) ports. You need to do this only to set up the forest trusts, so you can close the ISA Server Firewall client after you set up the trust relationship. You also need to add another rule to allow Kerberos (88 UDP) traffic. You must keep this rule so that authentication requests can pass through the firewall. To use this procedure on the ISA Server SP1 Firewall client:

1. Click **Start**, point to **All Programs**, point to **Microsoft ISA Server**, and then click **ISA Management**.
2. In the console pane, click to expand **Servers and Arrays**, click to expand your server, click to expand **Access Policy**, click **Protocol Rules**, and then, in the results pane, click **Create a Protocol Rule**.
3. Type a name for the rules, such as **allow xforest setup**, and then click **Next**.

4. Click **Allow**, and then click **Next**.

5. In the **Apply this rule to list** box, click **Selected protocols**, click to clear the **Show only selected protocols** check box, click to select the **Kerberos-Sec(UDP)** check box, and then click to select the **LDAP** check box, and then click **Next**.

6. In the **Use this schedule** box, click **Always**, and then click **Next**.

7. In the **Apply this rule to requests from** box, click **Any request**, click **Next**, and then click **Finish**.

8. Create protocol definitions for both Microsoft-DS and LDAP UDP, which are not defined by default:

    - After you create the new rule in step 1 through step 7, double-click the rule in the results pane, click the **Protocol** tab, and then click **New**.

    - In the **Properties** dialog box for the rule, type a name for your rules, such as **LDAP UDP**.

    - In the **Port number** box, type **389**.

    - In the **Protocol type** box, click **UDP**.

    - In the **Direction** box, click **Receive Send**, and then click **OK**.

9. Add a protocol definition for Microsoft-DS:

    - In the results pane, double-click the rule that you created in step 1 through step 7, click the **Protocol** tab, and then click **New**.

    - Type a name for your rules, such as **microsoft-ds outbound**.

    - In the **Port number** box, type **445**.

    - In the **Protocol type** box, click **TCP**.

    - In the **Direction** box, click **Outbound**, click **OK**, and then click **OK**.

10. After you create the rule, restart the firewall service to start the rule:

    - In the ISA Management Console, click to expand **Servers and Arrays**, click your server, and then click **Monitoring**, and then click **Services**.

    - In the results pane, click the **Firewall** service, click **Stop a Service**, and then click **Yes**.

    - In the results pane, click the **Firewall** service, click **Start a Service**, and then click **Yes**.

11. Set up the trust across the forest from the internal computer by following the steps that are in the "Implementing Trusts Across Forests" section earlier in this document.

12. Remove the trust setup rules because you no longer need them.

13. In the console tree, click to expand **Servers and Arrays**, click your server, click **Access policy**, and then click **Protocol rules**.

14. Double-click **Allow xforest setup**, and then click to clear the **Enable** check box.

15. Use the procedure in step 8 to restart the firewall service so that the rule is removed.

**Trust Validation and Network Logons**

If you want trust validation and network logons to shares on the external computer, you must open up certain ports. In protocol rules, you must open the LDAP UDP, Kerberos sec UDP, Any RPC server, DCE endpoint resolution outbound, and Microsoft-DS outbound ports. In addition, you must also publish both the LDAP UDP and Any RPC server ports on the internal forest domain controller. To do this:

1. Click **Start**, point to **All Programs**, point to **Microsoft ISA Server**, and then click **ISA Management**.

2. In the console tree, click to expand **Servers and Arrays**, click to expand your server, click to expand **Access policy**, and then click **Protocol rules**.

3. Click **Create a Protocol Rule**.

4. Type a name for the rules (such as **allow xforest validation**), click **Next**, click **Allow**, and then click **Next**.

5. In the **Apply this rule to** list, click **Selected protocols**, and then click to clear the **Show only selected protocols** check box.

6. Click to select all of the following check boxes:

   - **Any RPC server**
   - **Kerberos-Sec(UDP)**
   - **LDAP UDP**
   - **MS-DS outbound**

7. Confirm that the check boxes for the four definitions that you created (DCE Endpoint resolution, LDAP TCP inbound, LDAP UDP inbound, and Kerberos UDP inbound) are selected, confirm that the remaining check boxes are cleared, and then click **Next**.

8. Choose a schedule for your rule, and then click **Next**.

9. Choose to either make this rule applicable to only a specific client set or applicable to all of the clients, click **Next**, and then click **Finish**.

10. The protocol definitions for DCE endpoint resolution, LDAP TCP inbound, LDAP UDP inbound, and Kerberos UDP inbound are not available by default, so you must create them. Note that both LDAP UDP inbound and Kerberos UDP inbound are used only in the next section. To create the protocol definitions:

    - Add a protocol definition for DCE Endpoint resolution:

      In the results pane, double-click the rule, click the **Protocol** tab, and then click **New**.

      Type a name for your rules, such as **dce endpoint resolution outbound**.

      In the **Port number** box, type **135**.

      In the **Protocol type** box, click **TCP**.

      In the **Direction** box, click **Outbound**, and then click **OK**.

- Add a protocol definition for LDAP TCP inbound:

  In the results pane, double-click the rule, click the **Protocol** tab, and then click **New**.

  Type a name for your rules, such as **LDAP TCP inbound**.

  In the **Port number** box, type **389**.

  In the **Protocol type** box, click **TCP**

  In the **Direction** box, click **Inbound**, and then click **OK**.

- Add a protocol definition for LDAP UDP inbound:

  In the results pane, double-click the rule, click the **Protocol** tab, and then click **New**.

  Type a name for your rules, such as **LDAP UDP inbound**.

  In the **Port number** box, type **389**

  In the **Protocol type** box, click **UDP**.

  In the **Direction** box, click **Receive Send**, and then click **OK**.

- Add a protocol definition for Kerberos UDP inbound:

  In the results pane, double-click the rule, click the **Protocol** tab, and then click **New**.

  Type a name for your rules, such as **kerberos UDP inbound**.

  In the **Port number** box, type **88**

  In the **Protocol type** box, click **UDP**.

  In the **Direction** box, click **Receive Send**, and then click **OK**.

11. Publish the ports for the inbound connection:

- In the console tree, click to expand **Publishing**, and then click **Server Publishing Rules**.

- Click **Publish a Server**, type a name for the server (such as **LDAP UDP**), and then click **Next**.

- Type the IP address for the server that you want to publish, type the IP address for the external interface through which the external server will contact the internal server, and then click **Next**.

- Click the appropriate protocol (such as **LDAP UDP inbound**), and then click **Next**.

- Click **Specific computers**, and then click **Next**.

- Click **Add**, add a client set for the IP address for the external server, and then click **OK**.

- Click the client set, click **OK**, click **Next**, and then click **Finish**.

  Repeat above steps to publish the server for the "Any RPC server" rule.

**Object Picker: Allowing a Service Administrator to Set Up DACLs**
For Object Picker to work, you must publish Any RPC server, LDAP UDP inbound, LDAP TCP inbound, and Kerberos UDP inbound. When you do this, you can look up and add objects in the internal forest to DACLs that are in the external forest.

**Separately Setting Up the Trust in the Forests**
This method is not recommended and is only advisable if you cannot create the trust in both forests from the internal server. This involves publishing ports on the internal domain controller to allow the forest to make connections to it. You can create the trust on the forest that is on the internal side of the firewall by using the procedures that are in the previous section (set up the trust in both forests). To create the trust on the forest on the external side of the firewall, publish the following ports on the internal server domain controllers:

- Microsoft-DS (port 445)
- LDAP TCP (port 389)
- LDAP UDP (port 389)
- Kerberos-sec UDP (port 88)
- You can publish a server for each of these ports if you use the steps that are in the previous section.

Note that because the local Direct Host SMB service uses port 445 on the firewall by default, you must disable the service first. To do this, either set or create the registry key DWORD value to zero (0) in the following registry key, and then restart the computer:

**HKEY_LOCAL_MACHINE\SYSTEM\CCS\Services\NetBT\Parameters\SMBDeviceEnabled**

After you do this, the ISA Server firewall client can bind to port 445 and forward packets to the internal server.

**Error Messages**
You may receive the following error messages when you try to set up trusts across firewalls:

- **Specified domain is not a Windows domain name.**

  You may receive this error message if one domain controller cannot locate the other domain controller. DNS is not properly set up.

- **Parameter is invalid.**

  The LDAP UDP port is not open.

- **Access Denied.**

  The Kerberos port is not open, or the remote password is incorrect.

- **RPC server unavailable.**

  The Microsoft-DS port is not open.

- **Cannot query the forest functionality level. The forest is not operational.**

  The LDAP TCP port is not open.

- **The Endpointmapper does not have any more end points.**

  Either the NTDS or Netlogon port is not open.

  The Kerberos port is not open or the remote password is incorrect.

- **Unspecified error.**

Verify that the Endpointmapper port is open.

**Removing the Forest Trust**

The administrators for the Contoso company no longer need the trust across the forest.

In this scenario, you will remove the forest trusts. To remove the forest trusts:

1. Log on to the corp.contoso.com domain with administrative privileges.

2. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Domains and Trusts**.

3. Right-click **corp.contoso.com** in the left pane, and then click **Properties**.

4. Click the **Trusts** tab, right-click **nwtraders.com** in the **Domains trusted by this domain (outgoing trusts)** box, and then click **Remove**.

5. Click **Yes, remove the trust from the local domain and the other domain**.

6. Type **Administrator** in the **User name** box, and then type the password in the **Password** box.

7. Click **Yes**, and then choose the option to remove the trust.

Repeat step 4 through step 7 to remove the incoming trust in the **Domains that trust this domain (incoming trusts)** box.

# Appendix A: List of Ports

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The following tables shows the list of ports that you must open before you set up trusts.

| Scenario | Outbound Ports | Inbound Ports | From - To |
|---|---|---|---|
| Trust setup on both sides from the internal forest | LDAP (389 UDP and TCP)<br>Microsoft SMB (445 TCP)<br>Kerberos (88 UDP)<br>Endpoint resolution portmapper (135 TCP) Netlogon fixed port | | Internal domain domain controllers External domain domain controllers (all ports) |
| Trust validation from the internal forest domain controller to the external forest domain controller (outgoing trust only) | LDAP (389 UDP)<br>Microsoft SMB (445 TCP)<br>Endpoint resolution portmapper (135 TCP) Netlogon fixed port | | Internal domain domain controllers External domain domain controllers (all ports) |
| Object picker on the external forest to add objects that are in internal forest to groups and DACLs | | LDAP (389 UDP and TCP)<br>Windows NT Server 4.0 directory service fixed port<br>Netlogon fixed port<br>Kerberos (88 UDP)<br>Endpoint resolution portmapper (135 TCP) | External server Internal domain PDCs (Kerberos)<br>External domain domain controllers Internal domain domain controllers (Netlogon) |
| Setup trust on the external forest from the external forest | | LDAP (389 UDP and TCP)<br>Microsoft SMB (445 TCP)<br>Kerberos (88 UDP) | External domain domain controllers Internal domain domain controllers (all ports) |
| Kerberos authentication (internal forest client to external forest) | Kerberos (88 UDP) | | Internal client External domain domain controllers (all ports) |
| NTLM authentication (internal forest client to external forest) | | Endpoint resolution portmapper (135 TCP) Netlogon fixed port | External domain domain controllers Internal domain domain controllers (all ports) |

| | | | |
|---|---|---|---|
| Domain join from internal computer to external domain | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution portmapper (135 TCP) Netlogon fixed port<br><br>Windows NT Server 4.0 directory service fixed port | | Internal client External domain domain controllers (all ports) |

Configure the following keys to specify the services that you want to run on a fixed port:

1. LSA(Local Security Authority) RPC port (same as NTDS fixed port) used for trust creation and other access to the LSA Policy database TCP/IPPortentry in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** registry key.

2. Netlogon RPC port used for NTLM, Trust channel DCTcpipPort entry in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters** registry key.

# Appendix B: Terms That Are Used in this White Paper

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The following terms are used in this white paper:

- **External trust**: A single trust link between two domains that are in separate forests. The external trust establishes nontransitive trust from domain to the other. These trusts have to be explicitly created and maintained by the administrator. The trusts can be one-way or two-way trusts.

- **Internal Trust**: A single trust link between two domains in the same forest. The internal trust establishes transitive trust between the two domains. For example, if Domain A trusts Domain B and Domain B trusts Domain C, then Domain A trusts Domain C. These trusts are created automatically when a new domain is added to the forest and are required for the forest to work. These trusts are always two-way trusts.

- **Forest trust**: A single trust link between the root domains of two forests. The forest trust establishes transitive trust between all of the domains in each forest. For example, if Forest A trusts Forest B, then all of the domains in Forest A trust all of the domains in Forest B through the forest trust. However, this trust is not transitive across forests. For example, if Forest A trusts Forest B and Forest B trusts Forest C, then Forest A does not automatically trust Forest C. These trusts can be one-way or two-way trusts.

- **Trusted domain object (TDO)**: An object present in Active Directory that represents the trust between two domains or forests.

- **Forest TDO**: An existing trust definition object that has a bit that is set to indicate forest trust.

- **Namespace**: A grouping of related names in which names can be used to symbolically represent another type of information (for example, security principals) and for which specific rules are established to determine which names belong in that namespaces. For example, a namespace might be "all names ending with contoso.com except for the names that end with a.contoso.com or b.contoso.com." Forests manage many namespaces that frequently look similar but are used to identify different types of objects, such as user principal names (UPNs), service principal names (SPNs), and domain names.

- **Top-level name (TLN)**: A name that is used to identify a namespace. For example, corp.fabrikam.com is the TLN that identifies the namespace that contains all of the names under corp.fabrikam.com, such as server.corp.fabrikam.com, server1.corp.fabrikam.com, server2.corp.fabrikam.com, and so on.

- **Top-level name exclusion (TLNEx)**: A record that excludes a part of the namespace. For example, a TLNEx of test.fabrikam.com on the fabrikam.com TLN specifies that the namespace consists of all of the names that are under fabrikam.com except for the names that are under test.fabrikam.com.

- **Forest trust information (FTInfo)**: The set of namespaces (identified by their TLNs) that a trusted forest claims to manage, annotated with a field that indicates whether each claim is actually trusted by the trusting forest. FTInfo is stored on the forest TDO and is used for routing authentication and authorization requests to trusted forests.

- **User principal name (UPN)**: A variation of a users name that looks like an e-mail name, but can be used for logon. The syntax is *user_name@string*.

- **Implicit UPN**: The implicit UPN is always of the *user_id@DNS_domain_name.com* form. This UPN is always associated with the users account, even if an explicit UPN is not defined.

- **Explicit UPN**: The explicit UPN is always of the *string@any_string* form, where both *string* and *any_string* are explicitly defined by the administrator.

- **UPN suffix**: A string that indicates the namespace from which a UPN is derived. A UPN suffix includes everything on the

right side of the at sign (@) in the full UPN. Everything on the left side of the at sign is considered part of the user name, even if the text contains another at sign. UPN suffixes can be either flat or hierarchical in structure, although they are typically used to identify flat namespaces.

- **Service principal name (SPN)**: A multicomponent name that is used to identify a service that is associated with a computer account. SPNs are typically used to request Kerberos service tickets. The syntax is *service_type/instance_name* [:*port_number*] [/*service_name* [@*domain*] ]. For example, host/server.contoso.com, MSSQLSvc/server.contoso.com:1433.

- **SPN suffix**: A string that indicates the namespace from which a component of an SPN is derived. An SPN suffix may consist of the trailing substring of the *instance_name* or *service_name* component. It may include all of the *domain* component, except for the at sign (@).

- **Security identifier (SID)**: A unique identifier that is associated with every security principal, such as S-1-5-123232343434-544. The SID consists of a domain identifier and a relative identifier (RID).

# ADFS Design Guide

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) in the Microsoft® Windows Server® 2003 R2 operating system helps administrators meet federated identity management challenges. It does this by making it possible for organizations to securely share a user's identity information within an organization and across federated organizations—without creating and maintaining external trusts or forest trusts between those organizations. With ADFS, an administrator in an organization can control resources that users in that organization can access—both within that organization and at partner organizations. An administrator can also use ADFS to configure resources that users in other organizations can access. ADFS provides users with a Web-based, single-sign-on (SSO) experience when they access extranet Web sites or sites on the Internet that are accessible through federation partnerships.

For more information about how ADFS works and how to set up ADFS in a test lab, see the following resources:

- Appendix B: Reviewing Key ADFS Concepts
- Overview of Active Directory Federation Services (ADFS) in Windows Server 2003 R2 (http://go.microsoft.com/fwlink/?LinkId=54650)
- Step-by-Step Guide for Active Directory Federation Services (http://go.microsoft.com/fwlink/?LinkId=49531)

**About this guide**

This guide provides recommendations to help you plan a new deployment of ADFS, based on the requirements of your organization and the particular design that you want to create. This guide is intended for use by an infrastructure specialist or system architect. It highlights your main decision points as you plan your ADFS deployment. Before you read this guide, you should have a good understanding of how ADFS works on a functional level. You should also have a good understanding of the organizational requirements that will be reflected in your ADFS design.

This guide describes a set of deployment goals that are based on three primary ADFS designs, and the guide helps you decide the most appropriate design for your environment. You can use these deployment goals to form one of the following comprehensive ADFS designs or a custom design that meets the needs of your environment:

- Federated Web SSO to support business-to-business (B2B) scenarios and to support collaboration between business units with independent forests
- Federated Web SSO with Forest Trust to support business-to-employee (B2E) scenarios
- Web SSO to support customer access to applications in business-to-consumer (B2C) scenarios

For each design, you will find guidelines for gathering required data about your environment. You can then use these guidelines to plan and design your ADFS deployment. After you read this guide and finish gathering, documenting, and mapping your organization's requirements, you will have the information necessary to begin deploying ADFS using the guidance in the ADFS Deployment Guide.

**See Also**

**Concepts**

Understanding the ADFS Design Process
Identifying Your ADFS Deployment Goals
Mapping Your Deployment Goals to an ADFS Design
Evaluating ADFS Design Examples
Planning Partner Organization Deployments
Designing a Federated Application Strategy
Planning ADFS-Enabled Web Server Placement
Planning Federation Server Placement
Planning Federation Server Proxy Placement
Planning for ADFS Capacity
Finding Additional ADFS Resources

Appendix A: Reviewing ADFS Requirements
Appendix B: Reviewing Key ADFS Concepts
Appendix C: Documenting Your ADFS Design

# Understanding the ADFS Design Process

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

The following topics outline the Active Directory Federation Services (ADFS) design process for planning a new ADFS deployment that will meet the needs of your organization.

- Identifying Your ADFS Deployment Goals
- Mapping Your Deployment Goals to an ADFS Design
- Evaluating ADFS Design Examples
- Appendix C: Documenting Your ADFS Design

After you identify your deployment goals and map them to an ADFS design, you can begin documenting your design, based on the processes that are described in the following topics:

- Appendix A: Reviewing ADFS Requirements
- Planning Partner Organization Deployments
- Designing a Federated Application Strategy
- Planning ADFS-Enabled Web Server Placement
- Planning Federation Server Placement
- Planning Federation Server Proxy Placement
- Planning for ADFS Capacity

# Identifying Your ADFS Deployment Goals

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Correctly identifying your Active Directory Federation Services (ADFS) deployment goals is essential for the success of your ADFS design project. Depending on the size of your organization and the level of involvement that you want to provide for the information technology (IT) staff in any partner organizations, form a project team that can clearly articulate real-world deployment issues in a vision statement. Make sure that the members of this team understand the direction in which your deployment project must move in order to reach your ADFS deployment goals.

When you write your vision statement, identify, clarify, and refine your deployment goals. Prioritize and, possibly, combine your deployment goals so that you can design and deploy ADFS by using an iterative approach. You can take advantage of existing, documented, and predefined ADFS deployment goals that are relevant to the ADFS designs and develop a working solution for your scenarios.

The following table lists the three main tasks for articulating, refining, and subsequently documenting your ADFS deployment goals.

| **Deployment goal tasks** | **Reference links** |
|---|---|
| Evaluate predefined ADFS deployment goals that are provided in this section of the guide, and combine one or more goals to reach your organizational objectives. | • Provide federated access for your employees on the corporate network<br>• Provide federated access for your remote employees on the Internet<br>• Provide federated access for your hosted applications<br>• Provide single-sign-on access for customers to your hosted applications |
| Map one goal or a combination of any of the predefined ADFS deployment goals to an existing ADFS design. | • Mapping Your Deployment Goals to an ADFS Design |
| Document your deployment goals and other important details for your ADFS design. | • Appendix C: Documenting Your ADFS Design |

# Provide federated access for your employees on the corporate network

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

When you are the account partner administrator and you have a deployment goal to provide federated access for employees on the corporate network:

- Employees who are logged on to an Active Directory forest in the corporate network can use single sign-on (SSO) to access multiple applications, which are secured by Active Directory Federation Services (ADFS), when the applications are in a different organization. For more information, see Federated Web SSO design.

  For example, A. Datum Corporation may want corporate network employees to have federated access to applications that are hosted in Trey Research.

- Employees who are logged on to an Active Directory forest in the corporate network can use SSO to access multiple applications, which are secured by ADFS, in the perimeter network in your own organization. For more information, see Federated Web SSO with Forest Trust design.

  For example, A. Datum Corporation may want corporate network employees to have federated access to applications that are hosted in the A. Datum Corporation perimeter network.

- Information in the Active Directory account store can be populated into the employees' ADFS tokens.

The following components are required for this deployment goal:

- **Active Directory:** Active Directory contains the employees' user accounts that are used to generate ADFS tokens. Information, such as groups and attributes, is populated into ADFS tokens as group claims and custom claims. For more information about Active Directory, see Appendix B: Reviewing Key ADFS Concepts.

  | **Note** |
  | --- |
  | Active Directory Application Mode (ADAM) may also be used to contain the identities are used to generate ADFS tokens. However, ADAM is typically used for this purpose to host customer accounts on the perimeter network. |

- **Corporate DNS:** This implementation of Domain Name System (DNS) contains a simple address (A) resource record so that intranet clients can locate the account federation server. It may host other DNS records that are also required in the corporate network. For more information, see Name resolution requirements for federation servers.

- **Account federation server:** The account federation server is joined to a domain in the account partner forest. It authenticates employee user accounts and generates ADFS tokens. The client computer for the employee performs Windows Integrated authentication against the account federation server to generate an ADFS token. For more information, see Review the role of the federation server in the account partner organization.

  The following users can be authenticated by the account federation server:

  - Employees with user accounts in this domain
  - Employees with user accounts anywhere in this forest
  - Employees with user accounts anywhere in forests that are trusted by this forest (through a Windows trust)

- **Employee:** An employee accesses an ADFS-secured Web application through a supported Web browser while he or she is logged on to the corporate network. The employee's client computer on the corporate network communicates directly with the federation server for authentication.

The following illustration shows each of the required components for this ADFS deployment goal.



# Provide federated access for your remote employees on the Internet

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This deployment goal builds on the deployment goal that is described in Provide federated access for your employees on the corporate network. It also makes it possible for remote employees to obtain Active Directory Federation Services (ADFS) tokens from the account federation server. After it obtains the tokens, the remote employee's client computer can use the ADFS tokens to gain federated access to ADFS-secured applications that are hosted in another organization and allow employees to access resources in their own organization.

For more information, see Federated Web SSO design.

For example, A. Datum Corporation may want remote employees to have federated access to ADFS-secured applications that are hosted in Trey Research, without requiring A. Datum employees to be on the A. Datum corporate network.

In addition to the foundational components that are described in Provide federated access for your employees on the corporate network (which are shaded in the following figure), the following components are required for this deployment goal:

- **Account federation server proxy:** Employees that access the federated application from the Internet can use this ADFS component to perform authentication. By default, this component performs forms authentication, but it can also perform basic authentication. You may also configure this component to perform Secure Sockets Layer (SSL) client authentication if users at your organization have certificates to present. For more information, see Where to place a federation server proxy.

- **Perimeter DNS:** This implementation of Domain Name System (DNS) provides the host names for the perimeter network. For more information about how to configure perimeter DNS for a federation server proxy, see Name resolution requirements for federation server proxies.

- **Remote employee:** The remote employee accesses an ADFS-secured Web application through a supported Web browser, using valid credentials from the corporate network, while the employee is offsite using the Internet. The employee's client computer in the remote location communicates directly with the federation server proxy to generate a token and authenticate to the application.

The following illustration shows each of the required components for this ADFS deployment goal.



# Provide federated access for your hosted applications

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

When you are the resource partner administrator and you have a deployment goal to provide federated access to an application that resides in your organization (the resource partner organization), federated users both in your organization and in organizations that have configured a federation trust to your organization can access the Active Directory Federation Services (ADFS)-secured application that is hosted by your organization. For more information, see Federated Web SSO design and Federated Web SSO with Forest Trust design.

The following components are required for this deployment goal:

- **Active Directory:** The resource federation server must be joined to an Active Directory domain. If Windows NT token–based applications are supported, the domain also serves as the store that contains the resource accounts or resource groups. Claims-aware applications do not require local accounts in Active Directory. For more information about resource accounts and resource groups, see Determine your resource account mapping method.

- **Perimeter DNS:** This implementation of Domain Name System (DNS) contains a simple host (A) resource record so that clients can locate the resource federation server and the ADFS-enabled Web server. The DNS server may host other DNS records that are also required in the perimeter network. For more information, see Name resolution requirements for federation servers and Name resolution requirements for ADFS-enabled Web servers.

- **Resource federation server:** The resource federation server validates ADFS tokens that the account partners send. Account partner discovery is performed through this federation server. For more information, see Review the role of the federation server in the resource partner organization.

- **ADFS-enabled Web server:** The ADFS-enabled Web server can host a claims-aware application or a Windows NT token–based application. (The following illustration shows a claims-aware application.) The ADFS Web Agent confirms that it receives valid ADFS tokens from federated users before it allows access to the protected Web site. For more information, see When to create an ADFS-enabled Web server.

The following illustration shows each of the required components for this ADFS deployment goal.



# Provide single-sign-on access for customers to your hosted applications

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

When your ADFS deployment goal is to provide single-sign-on (SSO) access for customer accounts to hosted applications that are secured by Active Directory Federation Services (ADFS):

- Customers who are logged on to the Active Directory Application Mode (ADAM) account store, which is hosted in your perimeter network, can access multiple ADFS-secured applications, which are also hosted in your perimeter network, by logging on one time from client computers that are located on the Internet.

  In other words, when you host customer accounts to enable access to applications in your perimeter network, customers that you host in an account store can access one or more applications in the perimeter network simply by logging on once to the Federation Service. For more information, see Web SSO design.

- Information in the ADAM account store can be populated into customers' ADFS tokens.

The following components are required for this ADFS deployment goal:

- **Active Directory:** An Active Directory domain is required only for the resource federation server. It is not used to host customer accounts.

- **ADAM:** ADAM is used to contain the customer accounts that will be used to generate ADFS tokens. For more information about Active Directory or ADAM, see Appendix B: Reviewing Key ADFS Concepts.

  | **Note** |
  | --- |
  | Active Directory may also be used to contain customer accounts that will be used to generate ADFS tokens. |

- **Account/resource federation server:** This federation server serves in both the account role and the resource role. The account/resource federation server is configured so that the Federation Service includes values for both an application and an account store—in this case, ADAM—that contains the customer accounts. For more information, see Review the role of the federation server in the account partner organization and Review the role of the federation server in the resource partner organization.

- **ADFS-enabled Web server:** The ADFS-enabled Web server can host a claims-aware application or a Windows NT token–based application. The ADFS Web Agent confirms that it receives valid ADFS tokens from customer accounts before it allows access to the protected Web site. For more information, see When to create an ADFS-enabled Web server.

- **Customer:** While on the Internet, the customer accesses an ADFS-secured Web application through a supported Web browser. The customer client computer on the Internet communicates directly with the federation server for authentication.

The following illustration shows each of the required components for this ADFS deployment goal. In this case, because Active Directory is used only to support the federation server's requirement to be joined to a domain, it is shaded in this illustration.


Active Directory
Account/ resource federation server
Perimeter DNS
ADAM
ADFS-enabled Web server
Perimeter Network
Firewall
Customer
Internet

# Mapping Your Deployment Goals to an ADFS Design

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

After you finish reviewing the existing Active Directory Federation Services (ADFS) deployment goals and you determine which goals are related to your specific deployment, you can map those goals to a specific ADFS design. For more information about ADFS predefined deployment goals, see Identifying Your ADFS Deployment Goals.

Use the following table to determine which ADFS design maps to the appropriate combination of ADFS deployment goals for your organization. This table refers only to the three primary ADFS designs as described in this guide. However, you can create a hybrid or custom ADFS design by using any combination of the ADFS deployment goals to meet the needs of your organization.

| **ADFS deployment goal** | **Web SSO design** | **Federated Web SSO design** | **Federated Web SSO with Forest Trust design** |
|---|---|---|---|
| Provide federated access for your employees on the corporate network | No | Yes, in the account partner | Yes, in the account partner |
| Provide federated access for your remote employees on the Internet | No | Yes, optional in the account partner | Yes, optional in the account partner |
| Provide single-sign-on access for customers to your hosted applications | Yes | No | No |
| Provide federated access for your hosted applications | No | Yes, in the resource partner | Yes, in the resource partner |

# Web SSO design

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

In the Web Single-Sign-On (SSO) design in Active Directory Federation Services (ADFS), users must authenticate only once to access multiple ADFS-secured applications. In this design all users are external, and no federation trust exists because there are no partners. Typically, you deploy this design when you want to provide customer access to one or more ADFS-secured applications over the Internet, as shown in the following illustration.




With the Web SSO design, an organization that typically hosts an ADFS-secured application in a perimeter network can maintain a separate store of customer accounts in the perimeter network, which makes it easier to isolate customer accounts and employee accounts.

You can manage the local accounts for customers in the perimeter network by using either Active Directory or Active Directory Application Mode (ADAM) as the account store.

This design coincides with the deployment goal to provide SSO access for customers to your hosted applications. For more information, see Provide single-sign-on access for customers to your hosted applications.

To learn more about the flow of ADFS communications in this design, see Web SSO example.

For a list of detailed tasks that you can use to plan and deploy your Web SSO design, see Checklist: Implementing a Web SSO Design.

# Federated Web SSO design

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

The Federated Web Single-Sign-On (SSO) design in Active Directory Federation Services (ADFS) involves secure communication that spans multiple firewalls, perimeter networks, and name resolution servers—in addition to the entire Internet routing infrastructure.

Typically, this design is used when two organizations agree to create a federation trust relationship to allow users in one organization (the account partner organization) to access Web-based applications, which are secured by ADFS, in the other organization (the resource partner organization).

In other words, a federation trust relationship is the embodiment of a business-level agreement or partnership between two organizations. As shown in the following illustration, you can establish a federation trust relationship between two businesses, which results in an end-to-end federation scenario.




The one-way arrow in the illustration signifies the direction of the federation trust, which—like the direction of Windows trusts—always points to the account side of the forest. This means that authentication flows from the account partner organization to the resource partner organization.

In this Federated Web SSO design, two federation servers (one in A. Datum Corporation and the other in Trey Research) route authentication requests from user accounts in A. Datum Corporation to Web-based applications in Trey Research.

| **Note** |
| --- |
| For additional security, you can use federation server proxies to relay requests to federation servers that are not directly accessible from the Internet. |

In this example, A. Datum Corporation is the identity or account provider. The A. Datum Corporation portion of the Federated Web SSO design combines the following ADFS deployment goals:

- Provide federated access for your employees on the corporate network
- Provide federated access for your remote employees on the Internet

Trey Research is the resource provider. The Trey Research portion of the Federated Web SSO design achieves the following ADFS deployment goal:

- Provide federated access for your hosted applications

To learn more about the flow of ADFS communications in this design, see Federated Web SSO example.

For a list of detailed tasks that you can use to plan and deploy the Federated Web SSO design, see Checklist: Implementing a Federated Web SSO Design.

# Federated Web SSO with Forest Trust design

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

The Federated Web Single-Sign-On (SSO) with Forest Trust design in Active Directory Federation Services (ADFS) combines two Active Directory forests in a single organization, as shown in the following illustration.




Typically, you use this design when you want to provide employees on the corporate network and remote employees with federated access to ADFS-secured applications in the perimeter network, while using each employee's standard corporate domain credentials.

The one-way federation trust arrow in the illustration signifies the direction of the trust, which—like the direction of Windows trusts—always points to the account side of the forest. This means that authentication flows from the corporate network to the perimeter network.

Because a forest trust exists between the perimeter network and the corporate network, employee user accounts that are in the corporate network may be used to access the application, which eliminates the need for resource accounts or resource groups. A Windows NT token–based application requires that a user or group exists so that the ADFS token can be mapped into it. However, using Active Directory in the corporate network enables you to deploy the application without user accounts in the perimeter network.

| **Note** |
| --- |
| If a trust is not in place between the corporate network and the perimeter network and the application in the perimeter network is a Windows NT token–based application, resource accounts or groups must exist in the perimeter network. |

In this design, the single A. Datum Corporation organization combines the following ADFS deployment goals:

- Provide federated access for your employees on the corporate network
- Provide federated access for your remote employees on the Internet
- Provide federated access for your hosted applications

To learn more about the flow of ADFS communications in this design, see Federated Web SSO with Forest Trust example.

For a list of detailed tasks that you can use to plan and deploy the Federated Web SSO with Forest Trust design, see Checklist: Implementing a Federated Web SSO with Forest Trust Design.

# Evaluating ADFS Design Examples

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

The following Active Directory Federation Services (ADFS) design examples illustrate how you can use ADFS to enable federated identity and access management. You can use these topics to evaluate how ADFS communications work across all three ADFS designs and to determine which design or combination of predefined ADFS deployment goals best suits the goals of your organization.

- Web SSO example
- Federated Web SSO example
- Federated Web SSO with Forest Trust example

# Web SSO example

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

In this example, the fictitious company Adventure Works is an online retailer. The company sells products directly to customers over the Internet. The perimeter network hosts the Purchasing application and the Customer Service Web application, which are both claims-aware applications. Internet customer accounts and passwords are managed in Active Directory Application Mode (ADAM).

**Message flow for customer remote access**

The ADFS-enabled Web server that hosts the Purchasing and Customer Service applications is located in the perimeter network forest. Customers perform Active Directory Federation Services (ADFS) authentication for these applications by using the resource federation server for the ADFS-enabled Web server.

**Client application request**

The following illustration and corresponding steps provide a detailed description of the client application request process in ADFS using Transport Layer Security / Secure Sockets Layer (TLS/SSL).




1. The customer uses her Web browser to open the application on the ADFS-enabled Web server.

2. The ADFS-enabled Web server refuses the request because there is no ADFS authentication cookie, and the ADFS-enabled Web server redirects the client browser to the logon Web page on the resource federation server.

3. The client browser requests the logon Web page from the resource federation server.

4. The resource federation server redirects the client browser to its logon Web page.

**Authenticating the user**

The following illustration and corresponding steps continue to describe the client application request process in the previous section. Unless it is otherwise noted, all traffic uses TLS/SSL.




1. The Web page of the resource federation server prompts the client for user credentials.

2. The resource federation server does the following:

   - Validates the client's user credentials and retrieves attributes from ADAM using Lightweight Directory Access Protocol (LDAP).

   - Builds the security token for the ADFS-enabled Web server application.

   - Builds the ADFS authentication cookie.

3. The resource federation server redirects the Web browser to send the POST request to the ADFS-enabled Web server:

   - POST with security token in body and Java script to activate.

   - The ADFS authentication cookie is written to the Web browser.

4. The Web browser sends the POST request to the ADFS-enabled Web server.

5. The ADFS-enabled Web server redirects the Web browser to the Uniform Resource Locator (URL) of the application:

   - The ADFS-enabled Web server validates the security token.

   - Builds the new ADFS authentication cookie.

   - The ADFS authentication cookie is written to the Web browser.

6. The Web browser requests the original application URL from the ADFS-enabled Web server with the ADFS authentication cookie.

7. The application authorizes the user’s request, based on attributes from the security token.

The client browser requests additional application URLs from the ADFS-enabled Web server with its ADFS authentication

cookie that is created by the Web server.

# Federated Web SSO example

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

In this scenario, the fictitious company A. Datum Corporation has employees that work at their corporate offices and mobile employees that work at their homes. As with any other company, office supplies must be ordered from other retailers for the A. Datum employees. When users need to be authenticated, the account partner discovery page on the federation server in Trey Research always returns a Uniform Resource Locator (URL) for the external Federation Service Proxy.

| **Note** |
| --- |
| Domain Name System (DNS) servers in the perimeter network must be configured to resolve the host name of the Federation Service URL to the federation server proxy for employees that authenticate from home or on the road. |

As shown in the following illustration, employee sessions may originate from the corporate network for an office employee or from the Internet for a remote employee. Office employees are authenticated transparently to the federation server through their desktop logon sessions. Mobile employees are authenticated through passwords and forms authentication, or if it is configured, through client Secure Sockets Layer (SSL).




All the ADFS-enabled Web servers in the perimeter network of Trey Research are exposed directly to the Internet, without a federation server proxy. They are protected only by a firewall that screens out non-Web traffic. In a production environment, a business like Trey Research would probably deploy proxy servers in front of its Web farm servers. In this example, a proxy is omitted for the sake of clarity because it complicates the message flow with extra steps.

**Message flow for federation through internal access**

The ADFS-enabled Web server that is hosted by the online retailer is located in the perimeter network. When the office employees connect to the ADFS-enabled Web server directly, they are redirected to the default logon URL at the resource federation server proxy. Then, the employees must use home realm discovery to select the Federation Service endpoint URL. Corporate network DNS servers resolve that to the IP address of the internal account federation server.

The office employee authenticates by using his currently logged-on desktop session credentials and integrated authentication at an internal account federation server in the corporate network. The account federation server and Active Directory in the corporate network are used to validate the office employee's credentials and obtain attributes for building a Security Assertion Markup Language (SAML) security token.

**Client application request**

The following illustration and corresponding steps provide a description of the client application request process in ADFS using Transport Layer Security / Secure Sockets Layer (TLS/SSL).




1. The office employee uses his Web browser to open an application on the ADFS-enabled Web server.

2. The ADFS-enabled Web server refuses the request because there is no ADFS authentication cookie. The Web server redirects the client browser to the logon Web page on the resource federation server.

3. The client browser requests to sign-in from the resource federation server.

**Authenticating the user**

The following illustration and corresponding steps provide a description of the authentication process in ADFS. Unless it is otherwise noted, all traffic uses TLS/SSL.




1. The logon Web page on the resource federation server prompts the user for account partner discovery.

2. The resource federation server redirects the client browser to the logon Web page on the account federation server proxy.

3. The client browser requests the logon Web page from the account federation server proxy:

   - Internal DNS servers resolve the account federation server proxy URL to the CNAME of the account federation server.

   - Windows Integrated authentication occurs transparently.

4. The account federation server does the following:

- Validates user credentials and gets attributes from Active Directory in the corporate network forest using Lightweight Directory Access Protocol (LDAP).
- Builds the security token for the resource federation server.
- Builds the ADFS authentication cookie.

5. The account federation server redirects the Web browser to send the POST request to the resource federation server:
   - The POST request includes the security token in the body and Java script to activate.
   - The ADFS authentication cookie is written to the browser.
6. The client browser sends a POST request to the resource federation server.
7. The resource federation server redirects the Web browser to send the POST request to the ADFS-enabled Web server:
   - Builds the security token for the Web application.
   - Builds the new ADFS authentication cookie.
   - The POST request includes the security token in the body and Java script to activate.
   - The ADFS authentication cookie is written to the browser.
8. The Web browser sends the POST request to the ADFS-enabled Web server.
9. The ADFS-enabled Web server redirects the client to its application URL:
   - ADFS validates the security token.
   - Builds the new ADFS authentication cookie.
   - The ADFS authentication cookie is written to the browser.
10. The client browser uses the ADFS authentication cookie to request the original application URL from the ADFS-enabled Web server.
11. The ADFS-enabled Web server application authorizes the user's request based on attributes from the security token.

| **Note** |
| --- |
| If the resource application is an application that uses traditional Windows-based authorization such as a Windows NT token-based application, a resource account or resource group is required. |

The Web browser requests additional application URLs from the ADFS-enabled Web server with its ADFS authentication cookie.

**Message flow for federation through remote access**
When mobile employees connect to the ADFS-enabled Web server directly, they are redirected to the default logon URL at the resource federation server. Then, they must use home realm discovery to select the account Federation Service endpoint URL.

A mobile employee may then enter his credentials through a Web page that is displayed by the browser.

**Client application request**

The following illustration and corresponding steps provide a description of the client application request process in ADFS using TLS/SSL.




1. The remote employee uses the Web browser to open the application on the ADFS-enabled Web server.

2. The ADFS-enabled Web server refuses the request because there is no ADFS authentication cookie. The ADFS-enabled Web server redirects the client browser to sign-in on the resource federation server.

3. The client browser requests the logon Web page from the resource federation server.

4. The Web page on the resource federation server prompts the user for account partner discovery.

5. The resource federation server redirects the client browser to the logon Web page on the account federation server proxy.

6. The Web browser requests the logon Web page from the account federation server proxy.

**Authenticating the user**

The following illustration and corresponding steps provide a description of the user authentication process in ADFS. Unless it is otherwise noted, all traffic uses TLS/SSL.




1. The Web page on the account federation server proxy prompts the client for user credentials.

2. The account federation server proxy requests a security token from the account federation server by using Secure

Hypertext Transfer Protocol (HTTPS).

3. The account federation server does the following:

    - Validates user credentials and gets attributes from Active Directory in the corporate network forest using LDAP.
    - Builds the security token for the resource federation server.
    - Builds the ADFS authentication cookie.

4. The account federation server sends the security token and the ADFS authentication cookie to the account federation server proxy using Web methods and HTTPS.

5. The account federation server proxy redirects the Web browser to send the POST request to the resource federation server:

    - The POST request includes the security token in the body and Java script to activate.
    - The ADFS authentication cookie is written to the browser.

6. The client browser sends the POST request to the resource federation server.

7. The resource federation server redirects the Web browser to send the POST request to the ADFS-enabled Web server:

    - The POST request includes the security token in the body and Java script to activate.
    - Builds the new ADFS authentication cookie
    - The ADFS authentication cookie is written to the browser.

8. The client browser sends the POST request to the ADFS-enabled Web server.

9. The ADFS-enabled Web server redirects the Web browser to its application URL:

    - ADFS validates the security token.
    - Builds the new ADFS authentication cookie.
    - The ADFS authentication cookie is written to the browser.

10. The Web browser requests the original application URL for the ADFS-enabled Web server with the ADFS authentication, ADFS-enabled Web server cookie.

11. The Web application authorizes the user's request, based on attributes in the security token.

| **Note** |
| --- |
| If the resource application is an application that uses traditional Windows-based authorization, a resource account or resource group is required. |

The Web browser uses its ADFS authentication cookie to request additional application URLs from the ADFS-enabled Web

server.

# Federated Web SSO with Forest Trust example

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

In this scenario, the fictitious company Adventure Works is both a retail distributor and a wholesale distributor. The company sells products directly through Adventure Works salesmen, service departments, and sporting goods stores. Adventure Works obtains its entire inventory from other vendors. It does not manufacture the products it sells.

Adventure Works has two separate Active Directory forests. The forest in the perimeter network hosts the online inventory, purchasing, and customer service Web applications. The Inventory application that Adventure Works uses is critical to its operations. Therefore, this application is protected with strong security measures that require Windows impersonation and use access control lists (ACLs). The Purchasing application and the Customer Service application are claims-aware applications.

Traveling salesmen access perimeter network applications by using their employee accounts, which are managed in Active Directory in the corporate network forest.

As shown in the following illustration, a one-way trust from the perimeter network to the forest in the corporate network (an external trust for Microsoft® Windows® 2000 Server or a forest trust for Microsoft Windows Server® 2003) is created. The required ports are open in the firewall between the corporate network and the perimeter network to support trust and authentication as well as Web application traffic.




Because a one-way trust is already in place, Adventure Works decides to take advantage of it to enable applications to create access tokens for employees and have single sign-on (SSO) in the perimeter network.

Employee sessions may originate from the corporate network, for headquarters staff such as the Plant Manager, or from the Internet for mobile staff, such as traveling salesmen. This means that two types of authentication mechanisms are used:

- Intranet employees use integrated authentication.
- Mobile employees use Transport Layer Security / Secure Sockets Layer (TLS/SSL) client authentication or forms authentication.

Regardless of the type of user account or authentication mechanism, Active Directory Federation Services (ADFS) authentication always generates a security token, which is delivered to the application for the purpose of authorizing user requests. Because the inventory application is used only by employees who have Active Directory accounts, the ADFS Web Agent on the ADFS-enabled Web server also creates a Windows NT token for the user and the token can be impersonated by the application.

All the ADFS-enabled Web servers in the perimeter network are exposed directly to the Internet without a proxy server. They are protected only by a firewall that screens out non-Web traffic. In a production environment, Adventure Works would probably deploy proxy servers in front of its Web farm servers. In this example, however, a proxy server is omitted for the sake of clarity because it complicates the flow with extra steps.

**Message flow for traveling salesman remote access**

The ADFS-enabled Web server that hosts the Purchasing application is located in the Active Directory forest in the perimeter network. A traveling salesman authenticates to the application through an account federation server proxy. However, the default Federation Service for the ADFS-enabled Web server is the resource Federation Service. Therefore, when multiple account stores or multiple partners are configured, account partner discovery is required to redirect employees to the account federation server proxy. The account federation server and the corporate network Active Directory service validate the traveling salesman's credentials and obtain attributes for building a security token.

**Client application request**

The following illustration and corresponding steps provide a detailed description of the client application request process in ADFS using TLS/SSL.




1. The traveling salesman uses his Web browser to open the purchasing application on the ADFS-enabled Web server.

2. The ADFS-enabled Web server refuses the request because there is no ADFS authentication cookie. The Web server redirects the client browser to the logon Web page on the resource federation server.

3. The client browser requests the logon Web page from the resource federation server.

4. Depending on whether there are multiple account stores or multiple partners configured, the Web page on the resource federation server prompts the user for account partner discovery.

5. The resource federation server redirects the client browser to the logon Web page on the account federation server proxy.

6. The client browser requests the logon Web page from the account federation server proxy.

**Authenticating the user requesting access to the application**

The following illustration and corresponding steps describe how the traveling salesman is authenticated after requesting access to the Order Entry application. Unless it is otherwise noted, all traffic uses TLS/SSL.




1. The Web page on the account federation server proxy prompts the user for user credentials.

2. The account federation server proxy requests a security token from the account federation server by using Secure Hypertext Transfer Protocol (HTTPS).

3. The account federation server does the following:

    - Maps the user certificate and retrieves any required Lightweight Directory Access Protocol (LDAP) attributes from Active Directory in the corporate network using LDAP.

    - Builds the security token for the Web application.

    - Builds the ADFS authentication cookie.

4. The account federation server sends the authentication cookie and security token to the account federation server proxy using Simple Object Access Protocol (SOAP) and HTTPS.

5. The account federation server proxy redirects the Web browser to send the POST request to the resource federation server:

    - POST with security token in the body and Java script to activate.

    - The ADFS authentication cookie is written to the browser.

6. The client browser sends the POST request to the resource federation server.

7. The resource federation server does the following:

    - Builds the new security token for the Web application.

    - Builds the new ADFS authentication cookie.

The resource federation server redirects the Web browser to send the POST request to the ADFS-enabled Web server so that:

- POST with security token in the body and Java script to activate.
- The ADFS authentication cookie is written to browser.

8. The client browser sends the POST request to the ADFS-enabled Web server.
9. The ADFS-enabled Web server redirects the Web browser to its application Uniform Resource Locator (URL), and then:

    - The ADFS-enabled Web server component validates the security token.
    - Builds the new ADFS authentication cookie.
    - The ADFS authentication cookie is written to the browser.

10. The Web browser requests the original application URL from the ADFS-enabled Web server with the ADFS authentication cookie.
11. The ADFS-enabled Web server application:

    - Builds an impersonation Windows NT token from the security token.
    - Uses this Windows NT token to impersonate and to access check.

The Web browser requests additional application URLs from the ADFS-enabled Web server with its ADFS authentication cookie.

**Message flow for warehouse manager internal access**
The ADFS-enabled Web server that hosts the Inventory application is located in the perimeter network forest. The warehouse manager authenticates by using currently logged on desktop session credentials and integrated authentication at the account federation server in the corporate network. The account federation server and Active Directory in the corporate network validate the credentials and obtain attributes for building a security token.

When corporate network users connect to the ADFS-enabled Web server directly, they are redirected to the resource federation server. Then, as in the case of the traveling salesman, they are redirected through Domain Name System (DNS) servers to the canonical name (CNAME) of the account federation server.

**Client application request**
The following illustration and corresponding steps provide a detailed description of the client application request process in ADFS using TLS/SSL.




1. The warehouse manager uses a Web browser to open the application on the ADFS-enabled Web server.

2. The ADFS-enabled Web server refuses the request because there is an ADFS authentication cookie. The ADFS-enabled Web server redirects the client browser to the logon Web page on the (resource federation server).

3. The client browser requests the logon Web page from the resource federation server.

4. Depending on whether there are multiple account stores or multiple partners configured, the Web page on the resource federation server prompts the user for account partner discovery.

5. The client browser is sent directly to the account federation server.

**Authenticating the user**

The following illustration and corresponding steps provide a detailed description of the client application request process in ADFS. Unless it is otherwise noted, all traffic uses TLS/SSL.




1. The account federation server does the following:

    - Validates user credentials and gets attributes from Active Directory in the corporate network forest using LDAP.
    - Builds the security token for the Web application.
    - Builds the ADFS authentication cookie.

2. The account federation server redirects the Web browser to send the POST request to the resource federation server:

    - POST with security token in body and Java script to activate.
    - The ADFS authentication cookie is written to the browser.

3. The client browser sends the POST request to the resource federation server.

4. The resource federation server:

    - Builds the new security token for the Web application.
    - Builds the new ADFS authentication cookie.

    The resource federation server then redirects the Web browser to send the POST request to the ADFS-enabled Web server.

    - POST with security token in body and Java script to activate.
    - The ADFS authentication cookie is written to the browser.

5. The client browser sends the POST request to the ADFS-enabled Web server.

6. The ADFS-enabled Web server redirects the Web browser to its application URL:

    - ADFS validates the security token.
    - Builds the new ADFS authentication cookie.
    - The ADFS authentication cookie is written to the browser.

7. The client browser requests the original application URL from the ADFS-enabled Web server with the ADFS authentication cookie.

8. The Web application:

    - Builds an impersonation Windows NT token from the security token.
    - Uses this Windows NT token to impersonate and to access check.

The Web browser requests additional application URLs from the ADFS-enabled Web server with its ADFS authentication cookie.

# Planning Partner Organization Deployments

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

When you plan for cross-organizational (federation-based) collaboration using Active Directory Federation Services (ADFS), you first determine if your organization will host a Web resource to be accessed by other organizations across the Internet or if you will provide access to the Web resource for employees in your organization. This determination affects how you deploy ADFS, and it is fundamental in the planning of your ADFS infrastructure.

| **Note** |
| --- |
| Make sure that the role that your organization plays in the federation agreement is clearly understood by all parties. |

For the federation designs Federated Web SSO design and Federated Web SSO with Forest Trust design, ADFS uses terms such as "account partner" and "resource partner" to help differentiate the organization that hosts the accounts (the account partner) from the organization that hosts the Web-based resources (the resource partner).

"Account partner" and "resource partner" terminology is not applicable in the Web SSO design because one federation server acts in both the account partner role and the resource partner role, and there is no identification of two distinct partners in the Federation Service.

The following sections explain some of the ADFS partner organization concepts. They also contain links to topics in the Active Directory Federation Services Deployment Guide that contain information about setting up and configuring account partners and resource partners based on your ADFS deployment goals.

**Planning the account partner deployment**

An account partner represents the organization in the federation trust relationship that physically stores user accounts in either an Active Directory store or an Active Directory Application Mode (ADAM) store. The Federation Service in the account partner organization authenticates local users and creates security tokens that are used by the resource partner in making authorization decisions.

In scenarios in which you need to provide your users with access to multiple federated applications—when each application is hosted by a different organization—you can configure the account Federation Service so that you can deploy multiple resource partners.

For more information about how to set up and configure an account partner organization, see Checklist: Configuring the account partner organization.

**Planning the resource partner deployment**

The resource partner organization represents the organization whose ADFS-enabled Web servers are protected by the resource-side Federation Service. The Federation Service at the resource partner uses the security tokens that are produced by the account partner to make authorization decisions for ADFS-enabled Web servers that are located in the resource partner.

In scenarios in which you need to provide access to federated applications to many different users—when some users reside in different organizations—you can configure the resource Federation Service so that you can deploy multiple account partners.

For more information about how to set up and configure a resource partner in the account Federation Service, see Checklist: Configuring the resource partner organization.

# Review how ADFS may affect privacy

Updated: November 8, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) provides Web single-sign-on (SSO) technologies to authenticate a user to multiple Web applications over the life of a single online session. When you build an application using ADFS in Windows Server 2003 R2, your application may affect your end users' privacy.

For example, your application may explicitly collect user contact information, or it may request or send information over the Internet to your Web site or partner Web sites. If you embed Microsoft technology in your application, that technology may have its own behavior that may affect privacy.

**HTTP requests**

**Messages**

Messages contain Security Assertion Markup Language (SAML) tokens that may include application-defined personal data, depending on the claims that an administrator chooses to include in the token. The information that is contained in these SAML tokens is signed by the Federation Service's token-signing certificate, or, in some cases, the information is protected by a Kerberos session key. Messages are secured by the use of Secure Sockets Layer (SSL).

**Communication between the Federation Service and the application**

Although messages between Federation Services and Web applications transmit message routing information and body information, messages are typically secure by default. In addition, although messages contain application-specific data, such as application-defined personal information, this data is secured during transmission by SSL.

**Communication between the user and the application**

Query string parameters in the messages between a user and an application are subject to inspection by the account federation servers and the resource federation servers. Because an account federation server may be in a different organization, we recommend that you use a programming method that helps to ensure user privacy when you build Web applications.

A mitigation for this problem is to program an application to use HTTP POSTs, such as ASP.NET form submits, to communicate information that a user supplies. However, this method may have unreliable results because unauthenticated POSTs to applications that are protected by ADFS lose the POST body during ADFS authentication. (Unauthenticated POSTs include initial POSTs or POSTs that occur immediately after a token expires.) In spite of the potential for unreliable results, using an HTTP POST may prevent sensitive user data from being disclosed.

**Requested URL**

During the process of authenticating a user, the user is redirected to his or her account organization to perform authentication. As part of this request, the Uniform Resource Locator (URL) of the original requested page at the resource organization is included in encoded form. An account organization may decode this value to determine the first page at the resource that is accessed by the user.

**Cookies**

**Authentication**

After a token is verified at a Federation Service or Web application, an authentication cookie is issued to the client browser. Each time that a client must be authenticated, this cookie is used by the Federation Service, and, if the cookie is valid, the user is not required to enter his or her credentials again. The contents of this cookie are encrypted by SSL during transmission.

Application-defined personal information may be contained in this cookie, depending on the claims that an administrator includes in the token, and this personal information may be written to the disk of the Web browser. The contents of authentication cookies are signed, but they are not encrypted when they are written to disk.

All authentication cookies are session cookies that typically exist only in memory. These cookies are removed when the browser session ends or the ADFS sign-out procedure is invoked.

**Home-realm discovery**

During the process of home-realm discovery at a resource Federation Service or Federation Service Proxy, the client is prompted to select its account realm before being redirected to authenticate. After the account provider is selected, a persistent cookie is written to the disk of the Web browser computer that identifies the account provider. This cookie has a configurable lifetime, with a default of 30 days.

**Auditing**

When you enable the **Audit object access** policy setting and you configure ADFS for auditing, ADFS audits authentication

events. This includes the authentication of the user and any evidence, such as tokens, that the user presents. Application-defined personal information may be included in the security log, based on the configuration of the Federation Service that issues the security token. When auditing is enabled, both claim names and values are written to the security log. When the **Limit the auditing of this claim** check box is selected for a specific claim, only the claim's name is written to the security log.

#### Enhanced identity privacy

**Enable enhanced identity privacy** is an optional setting that you can configure for a resource partner in the trust policy. If the **Enable enhanced identity privacy** option is enabled, this setting hashes the user-name portion of outgoing user principal name (UPN) claims and e-mail claims. A random value is substituted for the common name.

The purpose of this feature is to prevent:

- The resource partner from correlating identity claims to personally identifiable user information.
- Collusion between partners in correlating identity claims to personally identifiable user information.
- Simple dictionary attacks to determine identity.

#### How enhanced identity privacy works

If you enable the **Enable enhanced identity privacy** setting, the user's identity is transformed by the account Federation Service before it is sent to the resource realm. The account Federation Service does this by hashing a combination of the privacy key (for salting), the resource partner Uniform Resource Identifier (URI), and the identity claim itself, as well as the privacy key (for salting), which reduces the possibility of dictionary attacks. When these three elements are combined, a unique hash per partner is created so that the resulting transformed identity is the same when the following values are the same:

- Privacy key
- Resource partner URI
- Identity claim (UPN or e-mail)

The hashing algorithm is case sensitive. Therefore, the hashed results are different if the case of the URI or the user-name portion of the UPN claim or e-mail claim is different.

#### The IT pro experience

#### Event logs

By default, ADFS records activities and events that are associated with messages to the event log. Event log entries are written that include data from messages and activities that are associated with those messages. This may include application-defined personal information, depending on the logging level and the claims that an administrator includes in the token.

#### Troubleshooting logs

You may configure ADFS to log messages that pass through the ADFS system, along with the activities and events that are associated with these messages. This feature is turned off by default. When logging is enabled, a log is written to the configured location. The log includes the messages and activities and the events that are associated with those messages. This may include application-defined personal information, depending on the logging level and the claims that an administrator includes in the token.

# Review the role of claims in the account partner organization

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

An account partner issues tokens containing claims to its users. The claims are built from the account store so that users can access Web-based applications in the resource partner. The following table describes the claim options in the account partner.

| **Claim option** | **Description** |
|---|---|
| Claim extractions | Claim extractions map a user or group in an account store to an organization claim. The account store can be either Active Directory or Active Directory Application Mode (ADAM). |
| Organization claims | Organization claims are used by the federation server—in this case, the account federation server. Claims passing through a federation server are mapped into and out of the organization claim set. These claims are then transformed, or mapped, into outgoing claims. This is the core set of claims that the organization uses for mapping. |
| Outgoing claim mappings | Outgoing claim mappings map organization claims to outgoing claims. The claim names that you configure here are determined by an agreement with your resource partner on a common namespace. |
| Outgoing claims | Outgoing claims are included in the security token of a user. They are generated by the account partner, and they are sent to the resource partner. Organization claims on the federation server of the account partner are mapped to outgoing claims that are then sent to the resource federation server. |

# Review the role of claims in the resource partner organization

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

A resource partner receives a token that contains claims, verifies that the token came from a trusted account partner, and makes the claims available to a Web-based application for authorization purposes. The following table describes claim options in the resource partner.

| Claim option | Description |
|---|---|
| Incoming claims | Incoming claims are received by the resource partner. They are generated by the account federation server as outgoing claims. The claim names that you configure here are determined by an agreement with your account partner on a common namespace. |
| Incoming claim mappings | Incoming claim mappings map incoming claims to organization claims. |
| Organization claims | The resource federation server transforms, or maps, incoming claims to organization claims. Organization claims are used by the federation server—in this case, the resource federation server. This is the core set of claims that the organization uses for mapping. |
| Claim transformation modules | You can use the Active Directory Federation Services (ADFS) user interface (UI) on your federation servers to change the name of a claim during the mapping process; for example, Administrators may become Admins. This way, you can share claims with partners and express claims in a common namespace that does not necessarily match the way that you define your own organization claims.<br><br>There may be situations, though, in which simple mapping of claims may not be sufficient for your scenario. In these situations, you can develop a claim transformation module that modifies claim names and values as they pass through the federation server. For example, the claim transformation module might convert a claim containing a monetary value into another currency based on an exchange rate. Or, the module might look into a sales database and determine the discount level for a partner.<br><br>If you determine that building a claim transformation module is necessary for your scenario, see the ADFS software development kit (SDK) on the MSDN Web site for additional information about how the claim transformation module functions. To find the ADFS SDK, search for "ADFS SDK" on the MSDN home page (http://go.microsoft.com/fwlink/?LinkId=16188). |

# Review the role of the federation server in the account partner organization

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

A federation server in the account partner is used to log on local user accounts in either an Active Directory store or an Active Directory Application Mode (ADAM) store. A federation server also issues initial security tokens that the local user accounts can use to access Web-based applications that are hosted in the resource partner. In addition, a federation server in the account partner issues cookies to users to maintain their logon status. These cookies include claims for those users. The cookies enable single-sign-on (SSO) capabilities so that users do not have to enter credentials each time that the users visit different Web-based applications in the resource partner.

| **Note** |
| --- |
| To create a federation server from a computer in the account partner organization, you must first join the computer to any domain in the forest where the federation server will be used to authenticate users from that forest. |

# Review the role of the federation server in the resource partner organization

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

A federation server in the resource partner validates the security tokens that are issued by the federation server in the account partner. A federation server in the resource partner also issues security tokens that are sent to the Web-based applications in the resource partner. In addition, a federation server in the resource partner issues cookies to the user accounts. The cookies come from the account partner. These cookies enable single-sign-on (SSO) capabilities so that users do not have to log on again at the federation server in the account partner when the users attempt to access different Web-based applications in the resource partner.

| **Note** |
| --- |
| To turn a computer into a federation server in the resource partner organization, you must first join the computer to any Active Directory domain in the resource partner organization. However, in scenarios in which the resource partner will use an ADFS-enabled Web server to host a Windows NT token–based application, you must join the federation server to a domain in the same forest as the ADFS-enabled Web server. |

In the Web SSO design, at least one federation server must be installed in the protected network. In the Federated Web SSO design and the Federated Web SSO with Forest Trust design, there must be at least one federation server in the account partner and at least one federation server in the resource partner.

# Review the role of the federation server proxy in the account partner organization

Updated: November 15, 2007

Applies To: Windows Server 2003 R2

The role of the federation server proxy in the perimeter network of the account partner is to collect authentication credentials from a client that logs on over the Internet and to pass those credentials to the federation server, which is located inside the corporate network of the account partner. (The account for the client is stored in the account partner.) The federation server issues a security token to the federation server proxy, which then sends the token to the client. The security token is used to provide access for that client to a specific resource partner.

The federation server proxy uses a default client logon Web form (clientlogon.aspx) to collect password-based credentials through forms-based authentication. However, you can customize this form to accept other supported types of authentication, such as Secure Socket Layer (SSL) client authentication. For more information about how to customize this page, see Customizing Client Logon and Home Realm Discovery Pages (http://go.microsoft.com/fwlink/?LinkId=104275). A federation server proxy does not accept credentials through Windows Integrated authentication.

| **Caution** |
| --- |
| Exposing a federation server proxy on the account partner extranet will make the client logon Web form accessible by anyone with Internet access. This can potentially leave your organization vulnerable to some password-based attacks, such as dictionary or brute force attacks that can trigger account lockouts for those user accounts that are stored in the corporate Active Directory directory service. |

To summarize, a federation server proxy in the account partner acts as a proxy for client logons to a federation server that is located in the corporate network. The federation server proxy also facilitates the distribution of security tokens to Internet clients that are destined for resource partners.

# Review the role of the federation server proxy in the resource partner organization

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

The role of the federation server proxy in the perimeter network of the resource partner is to perform account partner discovery for Internet clients and to redirect security tokens.

- **Account partner discovery.** An Internet client must identify which account partner will authenticate it. The client finds the account partner by using an account partner discovery Web form (discoverclientrealm.aspx), which is stored on the federation server proxy in the resource partner. If more than one account partner is configured in the trust policy in the Active Directory Federation Services snap-in, a drop-down menu appears with all the available account partners that are visible to Internet clients who access the account partner discovery Web form. You can change how the account partner discovery Web form is presented to clients by customizing the discoverclientrealm.aspx file.

- **Redirect security tokens.** The federation server proxy in the account partner sends the security tokens to the resource partner. The resource federation server proxy accepts these tokens and passes them on to the federation server in the resource partner. The resource federation server then issues a security token that is bound for a specific resource Web server. The resource federation server proxy returns the token to the client.

When it is necessary to help reduce the amount of hardware and number of required certificates, the federation server proxy can be located on the same computer as the ADFS-enabled Web server.

To summarize, a resource federation server proxy facilitates the federated logon process by redirecting clients to a federation server that can authenticate the clients. A resource federation server proxy also acts as a proxy for client security tokens to resource federation servers.

# When to enable Windows trusts

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) has built-in support for federation in organizations where a forest trust or external trust exists between two forests that represent the same organization. External trusts and forest trusts are also known as Windows trusts.

| **Note** |
| --- |
| External trusts and forest trusts are not required for ADFS to function. |

When Windows trusts have been established and you have decided to deploy a Federated Web Single-Sign-On (SSO) with Forest Trust design in your organization, you can configure ADFS to work over Windows trusts.

To successfully implement the Federation Web SSO with Forest Trust design and configure ADFS to support the Windows trust model, do all of the following:

- Verify that a forest trust exists between two Windows Server 2003 forests or that an external trust exists between Windows Server 2003 or Windows 2000 Server domains in each forest (where the resource forest trusts the account forest). For more information about Windows trusts, see Administering Domain and Forest Trusts (http://go.microsoft.com/fwlink/?LinkId=63917).

- Configure the resource Federation Service to enable the **Use Windows trusts relationship** option in the properties of the account partner. For more information, see Configure an account partner to use Windows trust.

- Configure the account Federation Service to enable the **Use Windows trust relationship for this partner** option in the properties of the resource partner. For more information, see Configure a resource partner to use Windows trust.

| **Note** |
| --- |
| Both the resource Federation Service administrator and the account Federation Service administrator must enable the Windows trust option so that the Federated Web SSO with Forest Trust design can function correctly. |

# Prepare client computers for federation

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

The easiest way for an administrator in the account forest to prepare client computers for access to Active Directory Federation Services (ADFS) federated applications is to use Group Policy. Group Policy provides a convenient way for you to push specific certificates and settings that are required for federation down to all the client computers that will be used to access federated applications.

So that your client computers can seamlessly access federated applications without certificate prompts or trusted site–related prompts, we recommend that you first prepare each client computer before you deploy ADFS broadly in your organization. Consider using Group Policy to:

- Configure Internet Explorer on each client computer to trust the account federation server.

  For more information, see Configure client computers to trust the account federation server.

- Install the appropriate account federation server, resource federation server, and ADFS-enabled Web server Secure Sockets Layer (SSL) certificates (or equivalent certificates that chain to a trusted root) on each client computer.

  For more information, see Distribute certificates to client computers using Group Policy.

# Designing a Federated Application Strategy

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

An important part of designing a new Active Directory Federation Services (ADFS) infrastructure is determining the full set of applications that will participate in the federation and which organizational partners will be the recipients of those application resources. Before you design a federated application strategy, consider the following questions:

- Will your organization host the federated application or applications?
- Will you be deploying a claims-aware application or a Windows NT token-based application?
- Will users on the corporate network require access to the federated application through Windows Integrated authentication?
- Are all of the ADFS-enabled Web servers that host federated applications running the Windows Server 2003 R2 operating system and Internet Information Services (IIS) 6.0?
- Who will the federated application provides resources for?
- Will the federated application be used by users in your perimeter network? If so, will Windows Integrated authentication be required?
- Will you use a Public Key Infrastructure (PKI) or the Kerberos authentication protocol to sign and protect tokens?

Answering these questions will help you plan a solid design. It will also assist you in creating a federated application strategy that is cost effective and resource efficient. For more information about designing the most appropriate federated application strategy for your organization, see the following topics:

- Provide federated access for your hosted applications
- Provide federated access for your employees on the corporate network
- Provide federated access for your remote employees on the Internet
- Provide single-sign-on access for customers to your hosted applications
- Review the role of ADFS Web Agents
- Identify the type of federated application to deploy
- When to use Authorization Manager
- Determine your security token protection method
- Determine your resource account mapping method

# Review the role of ADFS Web Agents

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Active Directory Federation Service (ADFS) Web Agents are Internet Server Application Programming Interface (ISAPI) extensions. They run on Internet Information Services (IIS) 6.0 and Windows Server 2003 R2, and they manage security tokens and authentication cookies for the Web server. An ADFS Web Agent intercepts incoming client Uniform Resource Locator (URL) requests for a protected resource and ensures that a valid authentication token is presented. If no valid authentication token is presented, the ADFS Web Agent redirects the client to the appropriate federation server or federation server proxy to begin the authentication and claims-creation processes.

| **Note** |
| --- |
| Several non-Microsoft Web agents are available that you can use to run federated applications on platforms other than IIS 6.0. These Web agents are based on a collaboratively developed specification for generating and consuming security tokens that fit a prescribed schema. This specification is called the WS-Federation Passive Requestor Interoperability Profile (WS-F PRP). |

If a valid authentication token is presented, the ADFS Web Agent provides the authorization data in the security token to the target application. An ADFS Web Agent can also write Hypertext Transfer Protocol (HTTP) cookies in client browsers to facilitate Web-based single sign-on (SSO).

ADFS Web Agents provide authorization support for two distinct types of Web applications:

- Applications that use Windows NT token–based authorization mechanisms, for example, access control lists (ACLs)
- Claims-aware ASP.NET 2.0 applications that can authorize against claims in the security token, using either ASP.NET roles or Authorization Manager

For more information about these types of federated applications, see Identify the type of federated application to deploy.

An ADFS Web Agent comprises two separate components:

- The ADFS Web Agent Authentication Service
- The ADFS Web Agent ISAPI Extension

The following sections describe these components.

**ADFS Web Agent Authentication Service**

The ADFS Web Agent Authentication Service validates incoming tokens and cookies. For this service to function correctly, the service account must be granted the TrustedComputingBase (TCB) Privilege. By default, this service runs as Local System.

| **Note** |
| --- |
| If you have a farmed Windows NT token-based application and you have configured the option **Domain service account** (with Kerberos) for the applications security token protection method, you must configure the ADFS Web Agent Authentication Service on each ADFS-enabled Web server that participates in the farm to run as the same domain service account that has been granted the TCB privilege. |

The ADFS Web Agent Authentication Service can generate tokens using either Service-for-User (S4U) or the ADFS authentication package. The IIS application pool is not required to run as Local System.

The ADFS Web Agent Authentication Service has remote procedure call (RPC) interfaces that may be called only from the local computer with local remote procedure call (LRPC), not RPC. This service returns a Windows NT access token if it is given an ADFS security token or an ADFS cookie.

**ADFS Web Agent ISAPI Extension**

The ADFS Web Agent ISAPI Extension is an IIS extension that enables ADFS authentication for Windows NT token-based

applications when those applications have had the ADFS Web Agent enabled through IIS.

In IIS Manager you can use the options on the **ADFS Web Agent** tabs on the **Web Sites** folder, **Default Web Site**, or on any additional Web site or virtual directory property pages to enable and configure a Windows NT token–based application for federation.

The ADFS Web Agent properties in the following table are inheritable. These properties are required on an IIS resource if the ISAPI extension is going to support WS-F PRP.

| Properties | Description |
|---|---|
| Federation Service URL | The URL of the Federation Service. This URL is required so that it may be queried for trust information by all the Web sites and federated applications on a Web server where the ADFS Web Agent is enabled. |
| Cookie path | The path that is specified when the authentication cookie is written. The cookie path identifies the location in an IIS Web site virtual directory to which cookies will be sent in response to an application request by an ADFS client.<br><br>If no cookie path is specified, the path defaults to a location that depends on whether a cookie domain is set, as follows:<br><br>• A cookie domain is set: / below the domain that applies for the requested address on the basis of the cookie domain.<br><br>• No cookie domain is set: / below the domain that is specified in the return URL.<br><br>If you specify a cookie path (for example, /test01), the cookie is sent for requests under that path (for example, /test01, /test01/index.html, /test01/test05/, and so on).<br><br>**Note**<br>The cookie path is case sensitive. It must match exactly with the name of the virtual directory that the federated application uses. |
| Cookie domain | The domain for which the cookie is valid. You can use the cookie domain setting to share an application at a higher level than the domain level that is specified in the return URL. In this way, you can expand the scope of requests for which a cookie will be sent.<br><br>If you do not configure a cookie domain, cookies are sent only in response to requests in which the domain that is specified matches the domain in the return URL. For example, if no cookie domain is set and the domain in the return URL is Sales.Adatum.com, cookies are sent only in response to requests in which the request URL matches Sales.Adatum.com. However, if you set Adatum.com as the cookie domain, cookies are sent in response to requests for Sales.Adatum.com plus requests for any other domain with the suffix Adatum.com. For example, cookies are also sent for Northwest.Adatum.com. |
| Return URL | The return URL that the token from the Federation Service comes back to after authentication at the Federation Service. The return URL identifies an application that is authenticated by ADFS to both the ADFS Web Agent and the Federation Service.<br><br>If the return URL for a federated application changes, you must change its value in IIS (for the Windows NT token–based Web Agent) or in the web.config file (for the claims-aware Web Agent).<br><br>**Note**<br>The **return URL** must match the **Application URL** that you configure in the properties of the application in the resource Federation Service. Therefore, if you change one URL, you must change the other. |

# Identify the type of federated application to deploy

Updated: January 10, 2008

Applies To: Windows Server 2003 R2

As part of the process of designing your deployment, identify the type of federated application that will be secured by Active Directory Federation Services (ADFS). All applications that are secured by Windows Server 2003 R2 ADFS Web Agents must be running under Internet Information Services (IIS) 6.0.

| **Note** |
| --- |
| Applications that require special client software or that are used for server-to-server communication are not good candidates for ADFS authentication. |

So that they can use ADFS Web Agents, each application must be at least one of the following application types:

- **Claims-aware application:** Claims-aware applications are Microsoft ASP.NET 2.0 applications that are written so that they can query ADFS security token claims. When a claims-aware application is presented with a valid token, the application can make authorization decisions based on the claims in the token.

- **Windows NT token–based application:** Windows NT token–based applications use Windows-based authorization mechanisms. These applications require Windows NT security tokens to make authorization decisions. ADFS can enable the transition of an ADFS security token to an impersonation-level Windows NT access token.

The following table shows the differences in requirements for each type of federated application.

| **Requirements and other factors** | **Claims-aware application** | **Windows NT token–based application** |
| --- | --- | --- |
| Requires the application to be written using ASP.NET 2.0 | Yes | No |
| Requires a user account in Active Directory | No | Yes |
| Requires a local Active Directory user store (to generate a user context for the application) | No | Yes |
| Requires Windows-based authorization mechanisms (such as access control lists (ACLs)) | No | Yes |
| Requires resource accounts or groups | No | Yes |
| Supported with ASP.NET 1.1 | No | Yes |
| Requires a web.config file | Yes | No |
| Can handle claims within the application code | Yes | No |
| Can use Authorization Manager for access control | Yes | No |

**Claims-aware applications**

When you set up a claims-aware application, all the information about a given identity is contained in the token that is presented by the application. The application may store additional information that links to the identity that is presented in the token, but a user account in Active Directory is not required to access a claims-aware application.

Federated applications that are claims aware should be developed specifically to use the claims that are produced by ADFS. When an application is presented with a valid token, the claims-aware application can make authorization decisions based on

the claims in the token.

For more information about how to deploy a claims-aware application, see Checklist: Installing a claims-aware application.

| Note |
| --- |
| When you configure the Web server to host only claims-aware applications, you do not have to provide values on any of the **ADFS Web Agent** tabs in the property sheets in IIS Manager. |

**Integrating with Authorization Manager**

Claims-aware applications can take advantage of role-based access control applications, such as Authorization Manager. Authorization Manager provides a flexible framework for integrating role–based access control into applications. Administrators who configure claims-aware applications can use Authorization Manager to provide access through assigned user roles that relate to job functions.

Authorization Manager applications maintain authorization policy in the form of authorization stores in Active Directory or XML files. The stores apply authorization policy at run time. For more information about ADFS and Authorization Manager, see When to use Authorization Manager.

**Windows NT token–based applications**

When you set up a Windows NT token–based application, you can map claims in incoming ADFS tokens to either user accounts or groups in the local Active Directory user store in the resource partner. The user accounts or groups are also known as resource accounts or resource groups. The application then uses the resource accounts or groups to perform authorization. For more information about resource accounts, see When to use resource accounts.

Applications that are protected by Windows Integrated authentication can—in some cases—be enabled as Windows NT token–based applications.

For more information about how to deploy a Windows NT token–based application, see Checklist: Installing a Windows NT token-based application.

| Note |
| --- |
| Because the ADFS Web Agent for Windows NT token–based applications implements authentication functions as an Internet Server Application Programming Interface (ISAPI) extension, an application that requires authentication in an ISAPI filter cannot be protected by ADFS. |

**Using SharePoint products as federated applications**

With ADFS in Windows Server 2003 R2, you can use either Microsoft® Windows® SharePoint® Services, Microsoft Office SharePoint Portal Server 2003, or Office SharePoint Server 2007 as supported federated applications. However, the level of support that ADFS provides depends largely on which SharePoint application that you will deploy.

| Important |
| --- |
| Before you deploy a specific SharePoint product for use with ADFS in your production environment, you should first read the following sections and the linked content to determine which SharePoint product best suits the needs of your organization when you use it as a federated application. |

**Using ADFS with previous versions of SharePoint products**

Windows SharePoint Services 2.0 and SharePoint Portal Server 2003 are considered to be Windows NT token–based applications as a result of their reliance on basic Windows-based authorization mechanisms. Because these SharePoint products depend on Windows-based access control, certain functionality will not be available when either Windows SharePoint Services 2.0 or SharePoint Portal Server 2003 are used as federated applications.

For more information about which functionality is supported in previous SharePoint products, see article 912492 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=107034).

For more information about how to configure Windows SharePoint Services 2.0 or SharePoint Portal Server 2003 as federated applications in a test lab environment, see the Step-by-Step Guide for Active Directory Federation Services (http://go.microsoft.com/fwlink/?LinkId=49531). This guide also provides instructions for removing unsupported SharePoint functionality that is not provided in the Knowledge Base article.

| Note |
| --- |

Previous versions of SharePoint products require the exclusive use of the default Web site in IIS.

**Using ADFS with current versions of SharePoint products**

Windows SharePoint Services 3.0 and Office SharePoint Server 2007 products can be configured as claims-aware applications as a result of their support for membership and role provider authorization mechanisms. Because these versions of SharePoint products can use membership and role-based access control, they do not require that resource accounts be created or used in the resource partner organization to take advantage of the single-sign-on (SSO) capabilities that are integrated into ADFS. This means that you can configure Windows SharePoint Services 3.0 and Office SharePoint Server 2007 as claims-aware applications for use with ADFS by using the SharePoint site administration page to configure membership and role-based access control permissions directly. However, for ADFS in Windows Server 2003 R2 to work smoothly with these SharePoint products, Service Pack 2 (SP2) must be installed on the ADFS-enabled Web server where the SharePoint product is installed.

For more information, see article 920764 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=107036).

For more information about how to configure Office SharePoint Server 2007 as a federated application, see Configure Web SSO authentication by using ADFS (http://go.microsoft.com/fwlink/?LinkID=84805).

**Note**

Although current versions of SharePoint products do not require the exclusive use of the default Web site in IIS, Windows SharePoint Services 3.0 and Office SharePoint Server 2007 do require the exclusive use of port 80. For this reason, installing these SharePoint products will automatically disable the default Web site in IIS. You can, however, enable the default Web site at a later time after you reassign the port number to something other than port 80.

# When to use Authorization Manager

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Authorization Manager is a component of Windows Server 2003 that provides role-based access control (RBAC) infrastructure for applications. When Authorization Manager is used with Active Directory Federation Services (ADFS), it offers the following advantages:

- **Administrative efficiency:** Administrators can use Authorization Manager to map ADFS claims to specific application roles to more easily control and enforce corporate access policy, rather than have corporate access policy built directly into a claims-aware application.

- **Developer flexibility:** Developers can create claims-aware applications that take advantage of the application authorization framework of Authorization Manager. The applications can then use RBAC policy and claims to make authorization decisions.

Federated claims-aware applications must be written specifically to take advantage of RBAC. For information about how to map ADFS claims to Authorization Manager roles, see Deploying Applications Using Windows Authorization Manager (http://go.microsoft.com/fwlink/?LinkId=77376).

# Determine your security token protection method

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You can use the Active Directory Federation Services snap-in to modify the way in which a federated application accepts security tokens from a resource federation server. This can be important when you have an application that requires a specific security-token-signing method, such as Public Key Infrastructure (PKI) or the Kerberos authentication protocol.

Security tokens that an ADFS-enabled Web server receives are signed and protected against tampering by resource federation servers using one of the following methods:

- **PKI:** The PKI method signs security tokens with the resource federation server's token-signing certificate. By default, the security-token-protection method in the Active Directory Federation Services snap-in is PKI.

- **Domain service account**: The domain service account method relies on the Kerberos authentication protocol to sign tokens using an assigned service principal name (SPN).

Federated applications that use PKI for token-signing have the following advantages over applications that use the domain service account method:

- PKI does not require the additional administrative overhead that is associated with manual configuration of service accounts for farmed federated applications.

- PKI can be used in scenarios in which you want an ADFS-enabled Web server, which hosts a federated claims-aware application, to not be joined to a domain (in a workgroup) or to be joined to a domain that is in a different domain than the domain that contains the resource federation server.

- PKI can be used to sign tokens in environments where Active Directory is not present, whereas the Kerberos protocol (which the domain service account method uses), relies on the Key Distribution Center (KDC) of Active Directory domain controllers to sign tokens.

Federated applications that use domain service account method (with the Kerberos authentication protocol) for token signing have the following advantages over PKI:

- The Kerberos protocol does not have the problem of certificate expiration because the expiration of Kerberos tickets is handled automatically.

- The Kerberos protocol provides better signing performance because it uses symmetric key operations, which are faster than the PKI asymmetric key operations.

For more information about how to configure either of these methods, see Configure the security token protection method for a federated application.

**PKI method**

The PKI method is the default value in the properties of the application in the Federation Service. This method is applied automatically when you create a new entry for a Windows NT token–based application or a claims-aware application. PKI is the default security-token-protection method because it does not require the additional configuration steps that the domain service account method requires.

When you use PKI as the security-token-protection method, the resource federation server signs only the tokens that are destined for the federated application (using the federation server's PKI-issued token-signing certificate) to protect the application from tampering.

**Domain service account method**

The domain service account security-token-protection method is an optional value in the properties of the application in the Federation Service. This method is identified by a service principal name (SPN). When a token is transferred using this method, the token contains a binary Kerberos V5 signature for the configured service account. This signature protects the token from

tampering.

If you select the domain service account method and you are using a single ADFS-enabled Web server to host your federation application, no additional service account configuration is necessary. This is because the domain service account on ADFS-enabled Web servers defaults to HOST\<*computername*>. For example, if you select the domain service account method in the properties of the application in the Active Directory Federation Services snap-in and the computer name of your ADFS-enabled Web server is ws.treyresearch.net, you type *HOST\ws.treyresearch.net*. Typing HOST\ws is also acceptable, but the full DNS name of the computer is preferred.

However, if you will be using the domain service account method for a federated application that is hosted on an ADFS-enabled Web server farm (where two or more servers run the appropriate Web Agent and they are used to load-balance the same federated application), there are additional considerations. The following section describes those considerations.

**Using the domain service account method in an ADFS-enabled Web server farm environment**
When you farm a federated application across multiple ADFS-enabled Web servers, you configure each ADFS-enabled Web server in the farm so that it points to the same SPN that is configured in the domain service account settings of the resource Federation Service. Complete the following steps to successfully configure your environment for the domain service account security-token-protection method.

**1. Create the service account**
If the domain service account method is selected in the Federation Service, you must type the SPN for a service account in the properties of the application. Before you can enter the SPN for a service account, you must first create a dedicated service account. The Kerberos protocol needs this service account to function properly and to allow pass-through authentication to the ADFS-enabled Web server. This service account is a dedicated account (that is, it should be used only as a service account) that you create in the Active Directory forest in the resource partner organization.

| **Note** |
| --- |
| To ensure that this service account is not interrupted as a result of domain password change requirements, consider configuring the account so that its password never expires. In the properties of the user account, select the **Password never expires** check box. |

**2. Set the SPN for the new service account**
Because the domain service account runs as a domain account, you must configure the SPN for this account in the domain. You configure the SPN with the Setspn command-line tool. Using a computer that is joined to the same domain where the service account resides, install Windows Server 2003 Support Tools to obtain Setspn.exe.

To install the Support Tools, on your Windows Server 2003 CD-ROM, go to the Support\Tools directory, double-click **Suptools.msi**, and follow the setup instructions. After you install the Support Tools, open a command-line prompt, type the following command, and then press ENTER:

**setspn -a <http/<** *farmclusterdnsname* **>** **<** *serviceaccountname* **>**

For example, in the scenario in which all ADFS-enabled Web servers are clustered under the DNS host name http://ws.treyresearch.net and the service account name that is assigned to the domain service account is adfswsfarm, type the following command:

**setspn -a** http/ws.treyresearch.net **adfswsfarm**

You need to complete this task only once for this account.

**3. Configure the Federation Service**
Now you configure the resource Federation Service to use the same ADFS-enabled Web server SPN that you configured previously, but without specifying the service account name. You can configure this SPN value in the **Domain service account** box for the application in the Active Directory Federation Services snap-in.

For example, if you select the domain service account method and your ADFS-enabled Web server has a computer name of ws.treyresearch.net, type the SPN value of http/ws.treyresearch.net. Also, the format host/ws.treyresearch.net is also acceptable in this field.

**4. Configure the service account on the ADFS-enabled Web server**
Now that the Federation Service contains the domain service account value of the ADFS-enabled Web server SPN, configure each ADFS-enabled Web server in the farm to also use the identity of the service account. However, the location where you configure the service account for each ADFS-enabled Web server is different, depending on the type of federated application that you are deploying in the farm. The following table indicates where in the user interface (UI) you specify the service account on each ADFS-enabled Web server, depending on the type of application that you are deploying in the farm.

| Federated application type | Configure the following: |
| --- | --- |
| Claims-aware application | The IIS application pool identity that is used by the federated application |
| Windows NT token–based application | The ADFS Web Agent Authentication Service |

You can use the following steps to configure each of your ADFS-enabled Web servers—based on the type of federated application—to use the service account that is required for Kerberos token-signing in the domain service account security-token-protection method.

**Claims-aware applications**

Complete the following steps on each of the ADFS-enabled Web servers in the farm when you want to use claims-aware applications with the Kerberos authentication protocol:

1. Open IIS Manager.
2. In the console tree, double-click **Application Pools**, right-click the specific **AppPool** that your federated application is using, and then click **Properties**.
3. On the **Identity** tab, click **Configurable**, and then type the service account name and password.
4. Click **OK** to exit the property page.
5. Repeat these steps on each ADFS-enabled Web server in the farm.

**Windows NT token–based applications**

Complete the following prerequisite tasks on each of the ADFS-enabled Web servers in the farm to configure the service account so that it will function correctly with the ADFS Web Agent Authentication Service:

- Grant the service account the TrustedComputingBase (TCB) privilege. For more information about the TCB privilege, see Review the role of ADFS Web Agents.
- Grant the service account the Audit privilege so that you can view all logged security audit events.
- Add the service account to the local IIS_WPG group.

Complete the following steps on each of the ADFS-enabled Web servers in the farm when you want to use Windows NT token–based applications with the Kerberos authentication protocol.

1. Open the Services snap-in.
2. Right-click **ADFS Web Agent Authentication Service**, and then click **Properties**.
3. On the **Log On** tab, click **This account**, and then type the service account name and password.
4. Click **OK** to exit the property page.
5. Repeat these steps on each ADFS-enabled Web server in the farm.

# Determine your resource account mapping method

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

In the past, Windows NT token–based applications could be used only by Windows users from the local forest or in trusting forests, that is, by users who could log on to the computer with Windows authentication techniques. By using Active Directory Federation Services (ADFS) and resource account mapping methods, you can extend access limits for Windows NT token-based applications, even across organizational boundaries in nontrusted forests.

In ADFS, resource account mapping is the act of mapping a federated user or a group of federated users to a security principal in Active Directory in the resource partner organization so that standard Windows authorization mechanisms can be applied to that security principal on the ADFS-enabled Web server. Resource account mapping is required when you federate Windows NT token–based applications because the Windows NT token–based Web Agent must reference an Active Directory security principal in the resource partner forest to build the Windows NT access token and thereby enforce access control permissions on the application.

| **Note** |
| --- |
| Claims-aware applications do not require resource account mapping because they use the ASP.NET membership and roles model, which means that they can inherently consume user principal names (UPNs) and group claims directly from the ADFS security token. |

The federation server in the resource partner organization supports any combination of the following three resource account mapping methods:

- Resource account
- Resource group
- Group-to-UPN mapping

You can use the information in the following table to help you determine which of these resource account mapping methods best suits your administrative needs.

| **Resource account mapping method** | **Description** | **Advantages** | **Disadvantages** |
| --- | --- | --- | --- |
| Resource account | A single security principal—usually a user account—created in Active Directory that is used to map to a single federated user | • Detailed user access control<br>• Verbose auditing results | High administrative overhead. Requires that account provisioning (creation, maintenance, and deletion) in Active Directory in the resource partner be tightly coupled with Active Directory in the account partner |

| | | | |
|---|---|---|---|
| Resource group | A single security group that is created in Active Directory that incoming group claims (ADFS group claims from the account partner) are mapped to. You can create more than one resource group. | • Eliminates the need to create and manage separate resource accounts for each federated user<br>• Uses standard access control lists (ACLs) to control authorization<br>• Highly scalable: one resource group can support an unlimited number of federated users that are mapped to it<br>• Accurate and accountable auditing results: The resource federation server provides a mechanism by which auditing can record events that uniquely identify each federated user. | The resource organization cannot control access to individual users, and it has to trust the account organization regarding group memberships |
| Group-to-UPN mapping | A group of federated users represented by the UPN of a user account that is created in the resource forest | • Eliminates the need to create and manage separate resource accounts for each federated user<br>• Highly scalable: one resource group can support a nearly unlimited number of federated users that are mapped to it | Inaccurate auditing results. Multiple user accounts are mapped to a single resource account that corresponds to a single UPN. Therefore, auditing cannot distinguish between the different accounts that were mapped. |

Because resource accounts, resource groups, and group-to-UPN mappings are all used solely for resource mapping, these methods are not used when the following conditions are true:

- The ADFS-enabled Web server and the resource federation server are joined to a domain in the forest or in a trusting forest where the identity resides.
- A one-way, cross-forest trust exists from the resource partner forest to the account partner forest.

In these situations, the federated user needing access to the Windows NT token–based application can access the resources directly in the resource forest by using the Windows trust.

The following topics are also relevant to resource account mapping methods:

- When to use resource accounts
- Select the optimal resource account option
- When to use resource groups
- When to use group-to-UPN mapping

# When to use resource accounts

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Resource accounts are simply user accounts that are stored in a resource partner forest. What makes a resource account different from a standard user account is that the sole purpose of a resource account is to represent the security context of a single identity that does not originate from the resource forest or from any trusting forests.

Because a federated identity is not created from the resource forest, it does not have a security identifier (SID) that identifies it as a security principal to standard, Windows-based authorization controls (such as access control lists (ACLs)) on servers in the resource forest. Therefore, the federated identity is prevented from gaining access to any Web-based federated resources in the resource partner, such as Windows NT token–based applications, that use ACLs.

Active Directory Federation Services (ADFS) requires a way for a federated identity to correspond or map to an existing security principal that does, in fact, have a SID that represents the resource forest. ADFS uses resource accounts to map federated identities to existing security principals.

ADFS creates a security token using the resource account. It then presents the SID of the resource account to any access control mechanism in the resource partner forest or a trusting forest. At this point—because there is a one-to-one mapping of the federated identity to the resource account—the federated identity can be treated like any other user account by Web server administrators.

| **Note** |
| --- |
| A federated identity, in ADFS terms, is known as a federated user. |

Consider using the resource account mapping method when:

- You are deploying a Windows NT token–based application as your federated application, and it requires access control settings.

- You want to provide more detailed levels of access to sensitive or private data.

  For example, the user AlanSh must have a different set of rights than the rest of the Information Technology (IT) staff. In this case, the fact that a user account is AlanSh is given preference in terms of access rights over the fact that he is a member of the group that represents IT staff so that he can get the allowed level of privilege.

So that federated users can access Windows NT token–based applications, resource accounts must be security principals in Active Directory. You cannot use Active Directory Application Mode (ADAM) user accounts as resource accounts in the resource partner forest. However, you can use ADAM user accounts in the account partner forest. You can create a corresponding resource account in the resource partner forest to represent each ADAM user account.

| **Note** |
| --- |
| A resource account is not intended to be accessed directly by the federated user whose user account it represents. Therefore, the federated user does not have to know about the resource account because he or she does not have to log on with this account directly. Also, the password for the resource account does not have to match the password of the federated user account. |

**Best practices for using resource accounts**

The following practices are recommended for using resource accounts to map to a federated user:

- As a security best practice, add all resource accounts to a security group. Then, add this group to the **Deny log on locally** and **Deny log on through Terminal Services** user rights in Group Policy.

  | **Note** |
  | --- |
  | If a resource account is disabled, resource account mapping will not function because a Windows NT token cannot be created for that account. |

- Create a single organizational unit (OU)—or perhaps one OU per account partner—where you can place all of the resource accounts for a particular partner so that you can more easily manage or delegate administrative control through Group Policy.

- To prevent possible interruption of service for federated users as a result of password restrictions, configure each account (in the properties of the user account) so that the password never expires. Use the **Password never expires** setting in Group Policy.

- To optimize how resource accounts are processed in the resource Federation Service, review the advantages and disadvantages of various resource account options. For more information, see Select the optimal resource account option.

- To decrease administrative overhead, write your own provisioning software with a claim transformation module to create a resource account automatically, based on the incoming claims. For more information, see Configure a claims transform module.

**Comparing resource accounts to resource groups**

A single resource account has a one-to-one mapping to a federated user—and only that user. A resource group, however, has a one-to-many mapping scheme. You can use a single resource group to map to one or more federated users that are members of a security group in the account partner forest. For more information about resource groups, see When to use resource groups.

Although both resource account mapping methods provide similar functionality, meaning that both methods provide federated users with a means to access resources without possessing their own resource forest SID, a resource group can usually provide a more well-rounded set of benefits than its resource account counterpart. This is mostly attributable to the amount of administrative resources that you can consume by regularly creating, managing, and pruning individual resource accounts.

As shown in the following illustration, the choice of a correct resource account mapping method for your deployment is important, because it can, over time, severely affect the level of administrative overhead in your organization.




**Preparing the resource forest for resource accounts**

If you decide to use the resource account mapping method, you must make some minor changes to Active Directory. You must also make some minor access control changes to the ADFS-enabled Web server. As shown in the following illustration and described in the following steps, these changes include the following:

- Adding the appropriate user principal name (UPN) suffixes to the resource forest

- Creating each resource account using the appropriate UPN suffix

- Assigning the appropriate permissions to the resource accounts in the Windows NT token–based application




| **Note** |
| --- |
| No modifications to the account forest are necessary to configure resource account mapping. |

**1. Add a UPN suffix to the resource forest for each account partner**
The first step in successfully routing authentication requests that are sent from federated users in the forest of the account partner organization to Windows NT token–based applications that are hosted on an ADFS-enabled Web server in the forest of the resource partner organization is to add a UPN suffix to the resource partner forest for each account partner forest that contains federated users. For example, if you have two account forests named Adatum.com and Tailspintoys.com and the resource forest is named Treyresearch.net, you add Adatum.com and Tailspintoys.com as UPN suffixes in the Treyresearch.net forest.

For more information about adding a UPN suffix to a forest, see Add user principal name suffixes (http://go.microsoft.com/fwlink/?LinkId=78171).

**2. Create the resource accounts using the account partner UPN suffix**
After you add a UPN suffix in the resource forest to mirror the UPN suffix of the account forest, you create individual resource accounts and assign the newly added UPN suffix to them. The UPN suffix in the Security Assertion Markup Language (SAML) tokens that are created by the account partner must match the UPN suffix of the associated resource accounts for successful access to federated Windows NT token–based applications. For more information about creating a resource account, see Create a resource account in the resource partner forest.

For example, say that Alan Shen is an employee of A. Datum Corporation. He regularly logs in to the network using his UPN, AlanSh@Adatum.com. The administrator for Trey Research has been given the task of creating a resource account for the AlanSh account. In this example, Adatum.com is the name of the account partner forest, and Treyresearch.net is the name of the resource partner forest. The administrator of Treyresearch.net uses the Active Directory Users and Computers snap-in to create a new user account. In the **New Object - User** dialog box, he types **AlanSh** in the **User logon name** field. He then selects the newly created UPN suffix **@adatum.com** in the drop-down list and creates the user account.

**3. Assign access permissions to the resource accounts**

After you create your resource accounts, you configure your Windows NT token–based application on the ADFS-enabled Web server with the appropriate access permissions for those accounts. Depending on the application, you can configure access permissions using standard ACLs or using the authorization controls that are built directly into some Windows NT token–based applications.

For example, if you want to configure an ADFS-enabled Web server to federate a Microsoft® Windows® SharePoint® Services site across organizations, you assign the resource account permissions using the **Site Settings** link on the Windows SharePoint Services site. For more information about ADFS and Windows SharePoint Services support, see Identify the type of federated application to deploy.

# Select the optimal resource account option

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) provides controls that the resource administrator—that is, the ADFS administrator in the resource partner forest—can use to define how resource accounts are used for a particular account partner. The following table describes resource account options and the circumstances in which a resource administrator might use them.

| Resource account option | Description | Can be used when the resource administrator wants to: |
|---|---|---|
| **Resource accounts exist for all users** | Specifies that a resource account is configured for each user from the account partner that needs access to the resource. In this case, incoming group claims are not mapped to resource groups even if resource groups are configured. | • Maintain complete control over access that is granted to specific account partner users. This option is most suitable for managing a small number of account partner users. |
| **Resource accounts exist for some users (prefer resource account)** | Specifies that resource groups are used for some user accounts. This means that some users may have individual resource accounts created for them, while other users may be configured to use resource groups.<br><br>When this option is selected, ADFS first looks for resource accounts that match the user principal name (UPN) claim or e-mail claim that is specified in the incoming token. ADFS uses these resource accounts if they exist. Otherwise, if the token has a group claim that is mapped to a resource group, ADFS uses the resource group. | • Prefer a specific set of account partner users. In addition, this option allows the account partner to manage its users through groups.<br><br>• Override the rights of specific users in a group.<br><br>• Provide specific users who belong in a privileged group in the account partner with more-detailed access to the resource. |
| **Resource accounts exist for some users (prefer groups in token)** | This is the default setting. Specifies that ADFS can use its logic to determine if each incoming token maps to a resource group or if it looks for a resource account.<br><br>When this option is selected, ADFS first looks in the token for incoming group claims that it can map to a resource group. If ADFS finds the incoming group claims, it uses the resource group. If there is no incoming group claim, ADFS looks for a resource account to use. | • Prefer delegating account management to an account administrator, including management of a specific set of account partner users.<br><br>• Provide access for specific users who are not in the privileged group in the account organization when they need access to the resource. For example, a resource is accessible only to executives; however, the account administrator needs access to parts of the resource.<br><br>• Move delegation of account management to the account administrator. |
| **No resource accounts exist for this account partner** | Specifies that one or more resource groups will be used for all users in this account partner. This means that every token that is issued from this account partner will be required to contain one or more group claims that map to one or more resource groups in the resource partner forest. | • Completely delegate account management to the account organization. |

For more information about how you can modify these resource account options, see Configure resource account options.

**Comparing resource account options**

Before you select any other option besides the default option, compare the various advantages and disadvantages that can result from using another resource account option, as described in the following table.

| Resource account option | Advantages | Disadvantages |
|---|---|---|
| **Resource accounts exist for all users** | • Access control list (ACL) model—the resource must be assigned permissions for the security identifier (SID) of the resource account or against SIDs for groups that the resource account is a member of. | • Administrative cost—high for large-scale account management at the resource partner. Costs can rise, depending on whether the account partner uses privacy enhancements. |
| **Resource accounts exist for some users (prefer resource account)** | • ACL model—the resource must be assigned permissions for the SID of the existing resource accounts, against SIDs for groups that the resource account is a member of, and also the SIDs of resource groups that correspond to groups in the token.<br>• Administrative cost—scalable, depending on the number of exception cases. | • Performance—slow as a result of account lookup<br>• Administrative cost—costs can rise, depending on whether the account partner uses privacy enhancements. |
| **Resource accounts exist for some users (prefer groups in token)** | • ACL model—the resource must be assigned permissions for the SID of the existing resource accounts, against SIDs for groups that the resource account is a member of, and also the SIDs of resource groups that correspond to groups in the token.<br>• Performance—In most cases, there is better performance. Exceptions are cases in which user account lookup occurs.<br>• Administrative cost—scalable, depending on the number of exception cases. | • Administrative cost—costs can rise, depending on whether the account partner uses privacy enhancements. |
| **No resource accounts exist for this account partner** | • ACL Model—ACL using groups.<br>• Performance—Better.<br>• Administrative cost—significantly less administrative costs are associated with account partners using privacy enhancements. | • Administrative cost—low, depending on the number of exception cases |

# When to use resource groups

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

In its simplest form, a resource group is a security group in the resource partner forest. Depending on the needs of your organization, you can create a resource group in Active Directory as a domain local group, a global group, or a universal group. When it is used with Active Directory Federation Services (ADFS), however, a resource group becomes a highly effective mapping method that administrators can use to easily map multiple federated users to a single security group.

Resource groups are considered to be a low-cost alternative to resource accounts because they do not require the administrative overhead that is associated with creating and maintaining resource accounts. Consider using the resource group mapping method when:

- You must eliminate or reduce the use of resource accounts. For more information, see When to use resource accounts.
- A Windows NT token–based application will be deployed as a federated application, and it requires per-group authorization.

**Mapping federated users to a resource group**

After federated users have been mapped to a resource group, ADFS-enabled Web servers can make authorization decisions to Windows NT token–based applications based on the access permissions that are assigned to the security identifier (SID) for the resource group. Federated users can be mapped to multiple resource groups for role based access control.

Windows NT token–based applications require access control mechanisms that are based on a valid SID that originates from the resource forest. When a Security Assertion Markup Language (SAML) token from an account partner organization is presented to the Federation Service in the resource partner organization, the resource federation server maps the claims in the SAML token to organization claims that are defined in its own organization, and it forms an organizational token. A component of the Federation Service referred to as the mapping module maps claims using the rules that you provide in the properties of the account partner in the resource Federation Service. If the organizational token contains any group claims that are associated with resource groups, the resource group information is also added to the SAML token. When this organizational token is presented to the Web application, the ADFS Web Agent for Windows NT token–based applications has an authentication package that can create tokens for federated users with group memberships as the resource groups in their token.

The following illustration shows the mapping module and how organization claims that have been previously mapped to a resource group carry the SID of the resource group into the new SAML token for the federated user. This illustration also shows how you can map multiple account security groups to a single resource group.




**How resource groups are applied to the SAML token**

The following illustration and corresponding steps show what happens when resource groups are applied to the SAML token.




1. The SAML token is sent by the account partner, and it contains the incoming group claim payload, along with any other claims that are assigned to Nate.

2. The mapping module in the Federation Service extracts Nate's claims from the SAML token and associates those claims with organization claims that reside in the Federation Service.

3. If the administrator has mapped an organization claim in the Federation Service to a resource group, the SID for that

resource group is added to the new SAML token that the resource federation server creates.

4. The new SAML token for Nate, which now includes the SID of the resource group, is then sent to the ADFS-enabled Web server through a client redirect so that the Web server can perform an access check against the access control lists (ACLs) that are associated with the Windows NT token–based application.

**How resource groups provide access control**

The following illustration and corresponding steps show what happens when resource groups provide access control.




1. A separate SAML token is created for each federated user and then sent by an account federation server to the resource federation server. Each token contains incoming group claims along with any other claims that are assigned to that federated user.

2. The mapping module in the Federation Service extracts the claims and maps them to organization claims that reside in the Federation Service.

3. If resource group SIDs are mapped to organization claims that are used by the incoming claims, these SIDs are added to the new SAML token that is created by the resource federation server. These SIDs are then sent in the token to the ADFS-enabled Web server.

4. The SAML token is used by the ADFS-enabled Web server to perform an access check on the Windows NT token–based application.

5. The access check determines whether or not the client has permission to perform the current operation with the application.

6. The ADFS-enabled Web server returns a page to the client, depending on the outcome of the access check.

# When to use group-to-UPN mapping

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Group-to-UPN (user principal name) claim mapping is one resource account mapping method that you can use in the resource Federation Service when claims are incoming from an account partner. In this context, "group" is the account partner security group and all the incoming claims that come from its federated member accounts; "UPN" represents a single user account in the resource partner forest. Group-to-UPN mappings map the claims of multiple federated user accounts, which are members of a security group in the account partner, to a single UPN that represents a single user account in the resource partner.

You can use the Active Directory Federation Services snap-in to create an ordered list of group-to-UPN claim mappings. The order of the group-to-UPN mappings is specified in the trust policy for the Federation Service. A group-to-UPN mapping list might be as follows:

1. Dev to devuser@adatum.com
2. Test to testuser@adatum.com
3. PM to pmuser@adatum.com

For example, if an incoming claim set contains (Common name=John Smith, Group=[Dev]), the organization claim set contains (Common name=John Smith, UPN=devuser@adatum.com). Because the list is ordered, a claim set of (Common name=John Smith, Group=[Dev,PM]) results in (Common name=John Smith, UPN=devuser@adatum.com).

**Comparing group-to-UPN mappings to resource groups**

When you use group-to-UPN mappings, audits that pertain to any user with a particular group claim contain the same mapped UPN as the user name. This means that the group-to-UPN mapping method cannot produce audits that identify which specific federated user was attempting an auditable action. This auditing limitation makes the group-to-UPN mapping method different from the resource group mapping method.

### Recommendation

Group-to-UPN mappings act in similar ways to resource groups. However, because group-to-UPN mappings lack the capability to record audits in more detail, we recommend that you consider deploying either the resource account mapping method or the resource group mapping method instead. For more information about how to implement a mapping method that records detailed audits, see When to use resource groups.

# Planning ADFS-Enabled Web Server Placement

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

An Active Directory Federation Services (ADFS)–enabled Web server relies on ADFS Web Agents to check if incoming requests need to be authenticated, and if so, the server directs the requests to a resource federation server to perform the actual authentication. The ADFS Web Agents also parse the tokens and cookies that are issued by the federation server to determine if access to the given application should be granted.

You must place at least one ADFS-enabled Web server in the resource partner organization. See the following topics for details about determining when and where to create and place an ADFS-enabled Web server:

- When to create an ADFS-enabled Web server
- Where to place an ADFS-enabled Web server
- When to create an ADFS-enabled Web server farm

| **Note** |
| --- |
| Although this information may help with your placement planning for ADFS-enabled Web servers, it does not explain how to determine the proper number of ADFS-enabled Web servers or the server hardware requirements for each ADFS design. For more information about determining the proper number of ADFS-enabled Web servers for each design, see Planning for ADFS-enabled Web server capacity. |

For examples of where ADFS-enabled Web servers can be placed in any of the three primary ADFS design scenarios, see Mapping Your Deployment Goals to an ADFS Design.

**See Also**
**Concepts**
Certificate requirements for ADFS-enabled Web servers
Name resolution requirements for ADFS-enabled Web servers

# When to create an ADFS-enabled Web server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Create at least one Active Directory Federation Services (ADFS)–enabled Web server in the resource partner organization when you deploy any of the following ADFS designs:

- Web SSO design
- Federated Web SSO design
- Federated Web SSO with Forest Trust design

The quickest way to get your federated applications up and running is to install and configure them on a single ADFS-enabled Web server. This way, you can set up a small-scale installation without performing the additional steps necessary to set up an ADFS-enabled Web server farm.

For more information about deploying federated applications, see Designing a Federated Application Strategy.

**Why an ADFS-enabled Web server is required**

An ADFS-enabled Web server provides the appropriate Web Agent software (ADFS Web Agents)—either claims-aware Web Agents or Windows NT token–based Web Agents—that are necessary for authenticating and authorizing federated access to locally hosted, Web-based applications. ADFS-enabled Web servers use these Web Agents to consume security tokens and authentication cookies—to either allow or deny a user access to the protected application—taking into consideration application-specific access control settings. Web Agents enforce application-based access control requirements by creating a security context in which the application can make the appropriate authorization decision.

For the ADFS-enabled Web server to know what tokens to accept, it must have a relationship with a federation server. This relationship is necessary so that all security tokens that are presented to the Web Agent (and destined for the application) are signed by that federation server (or any of the federation servers that represent that Federation Service). A signed security token indicates that the federation server has successfully verified the authenticity of the federated user.

To summarize, ADFS-enabled Web servers are a critical component of the ADFS infrastructure. ADFS Web Agents on these servers confirm that the incoming security tokens are signed by a valid federation server before they send federated access requests to the protected application.

# Where to place an ADFS-enabled Web server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

As is typical with perimeter networks, an intranet-facing firewall is established between the perimeter network and the corporate network, and an Internet-facing firewall is often established between the perimeter network and the Internet. In this situation, an Active Directory Federation Services (ADFS)–enabled Web server is placed between both of these firewalls on the perimeter network so that users coming from the Internet can access it.

In contrast, federation servers are typically placed inside the corporate network for security purposes. In scenarios in which you want to reduce the number of servers or public certificates in your ADFS deployment, you can make your ADFS-enabled Web server a federation server or federation server proxy by installing either the Federation Service or Federation Service Proxy component. However, placing a federation server on the perimeter network exposes that server to Internet clients, and this is not a security best practice. For more information, see Where to place a federation server.

**Configuring your firewall servers for an ADFS-enabled Web server**

So that ADFS-enabled Web servers can communicate directly with federation servers, both the intranet-facing firewall server and the Internet-facing firewall server must be configured to allow Secure Hypertext Transfer Protocol (HTTPS) traffic, usually over port 443. HTTPS configuration is required because ADFS-enabled Web servers communicate with federation servers—and with clients—over HTTPS. ADFS relies on HTTPS to provide channel security.

In addition, intranet-facing and Internet-facing firewall servers, such as servers running Internet Security and Acceleration (ISA) Server, use a process known as server publishing to distribute Internet client requests to the appropriate corporate federation servers and ADFS-enabled Web servers. Consequently, you must manually create a server publishing rule on any intranet and Internet ISA Server computers that publish the ADFS-enabled Web server URL (http://ws.treyresearch.net).

For general information about how to configure ISA Server to publish a server, see Create a secure Web publishing rule (http://go.microsoft.com/fwlink/?LinkId=74605).

**Joining an ADFS-enabled Web server to a domain**

When they host a Windows NT token-based application, ADFS-enabled Web servers must be joined to the same domain that the resource federation server belongs to. Because claims-aware applications do not require a Windows NT token for authorization, servers that host only claims-aware applications do not have to be joined to any domain.

# When to create an ADFS-enabled Web server farm

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You create an Active Directory Federation Services (ADFS)-enabled Web server farm when you want to balance the load of incoming federated access requests that are made to one or more protected applications. The obvious benefits that can be obtained from a Web server farm are fault tolerance for the hosted applications and a possible increase in client-side browser performance. To client computers, the Web server farm performs like a single Web server servicing a highly scalable federated application.

If you are thinking about creating a Web server farm solely to increase performance, consider whether it might be better to optimize the hardware or software performance of a single Web server before adding the additional costs and administrative overhead that are associated with synchronizing data between two or more Web servers. For more information about how to measure whether your ADFS-enabled Web server is performing optimally, see Planning for ADFS-enabled Web server capacity.

You can use a network load-balancing service, such as the Microsoft Network Load Balancing (NLB) service that is included in Windows Server 2003 operating systems, to create and manage your ADFS-enabled Web server farm. The number of Web servers that is required to support your federated applications is determined largely by the type and complexity of the applications and their client connection state (stateless as opposed to stateful) requirements.

Guidelines for creating server farms for federated applications using NLB are similar to the guidelines for creating server farms for nonfederated Web applications running on Internet Information Services (IIS) 6.0. For more information about using nonfederated Web applications with NLB, see Identifying Applications That Benefit from NLB (http://go.microsoft.com/fwlink/?LinkId=74610).

When they are used together, the recommendations in the NLB deployment documentation and in the topic Planning for ADFS-enabled Web server capacity should provide sufficient data for you to determine whether your federated applications can benefit from the improved scalability and availability that an ADFS-enabled Web server farm environment provides.

**Configuring servers in the farm**

Creating an ADFS-enabled Web server farm involves more than just placing the Web servers in a resource partner organization and then configuring NLB clustering. The following table provides additional guidance for identically configuring each of the ADFS-enabled Web servers in a farm.

| To configure the ... | To ... | See ... |
|---|---|---|
| Claims-aware Web Agent in the web.config file for the protected application | Point to the same Federation Service Uniform Resource Locator (URL) for each ADFS-enabled Web server in the farm | Set the Federation Service URL for a claims-aware application |
| Windows NT token–based Web Agent in the properties of the Web Sites folder in IIS 6.0 | Point to the same Federation Service URL for each ADFS-enabled Web server in the farm | Set the Federation Service URL for a Windows NT token-based application |
| Server authentication certificate | Export the private key so that the same server authentication certificate can be assigned to each ADFS-enabled Web server in the farm<br><br>Most certificates that are issued by certification authorities (CA)s can be used on multiple computers without first exporting them. If this is the case in your scenario, you do not have to perform this procedure.<br><br>**Note**<br>There is no requirement to use the same certificate on all ADFS-enabled Web servers as long as all Web servers in the farm have a certificate that is issued by the same CA and each of the certificates have a matching "subject name" field. | Export the private key portion of a server authentication certificate |

| | | |
|---|---|---|
| Server authentication certificate | Install on the appropriate Web site or virtual directory where your federated application will reside<br><br>(For an example of how to do this using the default Web site, see the following link.) | Import a server authentication certificate to the default Web site |

For additional details about how to configure an ADFS-enabled Web server farm, see Checklist: Installing an ADFS-enabled Web server.

# Certificate requirements for ADFS-enabled Web servers

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Each Web server that hosts an Active Directory Federation Services (ADFS) Web Agent requires a Secure Sockets Layer (SSL) server authentication certificate to communicate securely with Web clients. Publicly issued certificates are recommended for SSL server authentication certificates. However, if you are deploying the ADFS Web Single Sign-On (SSO) design, using either a public or corporate certification authority (CA) to obtain your server authentication certificate is sufficient.

| **Note** |
| --- |
| Token-signing certificates and SSL client authentication certificates are not necessary for ADFS-enabled Web servers. |

If you will be hosting additional ADFS components, such as the Federation Service or the Federation Service Proxy, on an already established ADFS-enabled Web server, it is not necessary to obtain additional server authentication certificates for each of those components. The ADFS Web Agent, the Federation Service, and the Federation Service Proxy can use a single server authentication certificate simultaneously. For more information about hosting multiple ADFS components on an ADFS-enabled Web server, see Where to place an ADFS-enabled Web server.

You can request and install server authentication certificates through the Microsoft Management Console (MMC) snap-in for Internet Information Services (IIS). For more general information about using SSL certificates, see Configuring Secure Sockets Layer (http://go.microsoft.com/fwlink/?linkid=62785) and Obtaining Server Certificates (http://go.microsoft.com/fwlink/?linkid=62479).

**Certificate requirements for an ADFS-enabled Web server farm**

In an ADFS-enabled Web server farm scenario, Web servers must obtain server authentication certificates in one of the following ways for ADFS to work:

- **Share the same certificate:** Web servers can share the same server authentication certificate across the farm. To share the same certificate across the Web servers, export the private key of that certificate and install it on the appropriate Web site for each Web server.

  For more information, see Export the private key portion of a server authentication certificate and Import a server authentication certificate to the default Web site.

- **Obtain individual certificates:** If you decide to obtain separate server authentication certificates for each Web server in a farm, you must ensure that the subject names for each of the individual server authentication certificates match. The subject name value for a server authentication certificate is used to identify the computer that the certificate represents.

  | **Note** |
  | --- |
  | Certification authorities (CAs) such as Microsoft Certificate Services create the subject name from the common name (CN) of the requestor that is obtained in Active Directory. |

For more information about configuring an ADFS-enable Web server farm, see When to create an ADFS-enabled Web server farm.

# Name resolution requirements for ADFS-enabled Web servers

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Before a Web browser client can contact an Active Directory Federation Services (ADFS)–enabled Web server over the Internet, the Web browser must first use Domain Name System (DNS) to resolve the Web server's host name to the actual IP address for an ADFS-enabled Web server that is located in a perimeter network. The following standard DNS tree-walking processes accomplish name resolution to federated applications:

1. The browser client contacts a top-level-domain DNS server on the Internet to resolve the fully qualified domain name (FQDN) of the target Web site that was typed into the browser.

2. The top-level-domain DNS server resolves the FQDN by providing the client with the IP address for the DNS server that is authoritative for the DNS domain that is specified in the Web site address.

   | **Note** |
   | --- |
   | In the federated world of ADFS, this DNS server is the DNS server that is located in the perimeter network of the resource partner organization. |

3. Using values that are stored in preconfigured host address (A) resource records, the perimeter DNS server resolves the target FQDN to the IP address of the ADFS-enabled Web server and then provides that information back to the client.

For more information about how the DNS process works, see How DNS Works (http://go.microsoft.com/fwlink/?LinkId=74637).

**Configuring perimeter DNS**

DNS is required for successful name resolution across the Internet to an ADFS-enabled Web server. DNS must be configured for a new host record that will resolve the IP address of the Web server cluster (if the Web servers are farmed) to a single Web server IP address and DNS host name. You configure DNS in the perimeter network of the resource partner.

In the following illustration, you can see how to configure the perimeter DNS so that it contains a single host (A) record for ws (ws.treyresearch.net) and so that it points to the IP address of the ADFS-enabled Web server cluster in the perimeter network. In this scenario, Network Load Balancing (NLB) provides a single, cluster FQDN name and a single, cluster IP address for an existing ADFS-enabled Web server farm.




For more information about how to configure a cluster IP address or a cluster FQDN using Microsoft NLB technology, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?LinkId=74651).

For more information about how to configure perimeter DNS, see Add a host (A) record to perimeter DNS for an ADFS-enabled Web server.

# Planning Federation Server Placement

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

The most critical component of an Active Directory Federation Services (ADFS) deployment is the federation server. Therefore, it is important that you plan your federation server placement strategy carefully, including when and where to deploy federation servers. The information in the following topics can help you determine when and where to create a federation server and whether to use that federation server in either the account partner role, the resource partner role, or both.

- Review the role of the federation server in the account partner organization
- Review the role of the federation server in the resource partner organization
- When to create a federation server
- Where to place a federation server
- When to create a federation server farm
- Certificate requirements for federation servers
- Name resolution requirements for federation servers

| **Note** |
| --- |
| Although this information might help with your placement planning for federation servers, it does not explain how to determine the proper number of federation servers and the hardware requirements for each ADFS design. For more information about determining the proper number of federation servers for each design, see Planning for federation server capacity. |

For examples of how a federation server can be placed in any of the three primary ADFS design scenarios, see Mapping Your Deployment Goals to an ADFS Design.

# When to create a federation server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You create federation servers in your organization whenever you want to deploy any of the following Active Directory Federation Services (ADFS) designs:

- Web SSO design
- Federated Web SSO design
- Federated Web SSO with Forest Trust design

When you create a federation server, you provide a means by which your organization can engage in Web single-sign-on (SSO)–based communication with another organization (that also has at least one federation server) and, when necessary, with the employees in your own organization (who need access over the Internet).

The role that a federation server plays in your organization depends on whether you place the federation server in the account partner organization or in the resource partner organization. For example, when a federation server is placed in the corporate network of the account partner, its role is to authenticate the user credentials of browser clients and send security tokens to the clients. When a federation server is placed in the corporate network of the resource partner, its role is to authenticate users, based on a security token that is issued by a federation server in the resource partner organization, or its role is to redirect token requests from configured Web applications to the account partner organization that the client belongs to.

For more information about account partner and resource partner organizations, see Planning Partner Organization Deployments.

| **Note** |
| --- |
| For the Federated Web SSO and Federated Web SSO with Forest Trust designs, there must be at least one federation server in the account partner and at least one federation server in the resource partner. |

If necessary, an organization that deploys a Federated Web SSO or Federated Web SSO with Forest Trust design can configure a single federation server so that it acts in both the account partner role and in the resource partner role. In this case, the federation server may produce Security Assertion Markup Language (SAML) tokens, based on user accounts in its own organization, or reroute token requests to the organization, based on where the users' accounts reside.

In the Web SSO design, where the user account exists in the same organization as the resource, it is not necessary to configure federation servers for either an account partner or a resource partner. The Federation Service, in this case, produces SAML tokens for the user that can be directly used on the resource.

**Differences between federation servers and federation server proxies**

Federation servers can serve out Web pages for sign-in, policy, authentication, and discovery in the same way that a federation server proxy does. The primary differences between a federation server and a federation server proxy have to do with what operations a federation server can perform that a federation server proxy cannot perform.

The following are the operations that only a federation server can perform:

- The federation server performs the cryptographic operations that produce the token. Although federation server proxies cannot produce tokens, they can be used to route or redirect the tokens to clients and, when necessary, back to the federation servers. For more information about using federation server proxies, see When to create a federation server proxy.
- Federation servers support the use of Windows Integrated authentication for clients on the corporate network, where federation server proxies do not. For more information about using Windows Integrated authentication with federation servers, see When to create a federation server farm.

**See Also**

**Concepts**

Review the role of the federation server in the account partner organization
Review the role of the federation server in the resource partner organization

# Where to place a federation server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

As a security best practice, place federation servers in front of a firewall and connect them to the corporate network to prevent exposure from the Internet. This is important because federation servers have full authorization to grant security tokens. Therefore, they should have the same protection as a domain controller. If a federation server is compromised, a malicious user has the ability to issue full access tokens to all Web applications and federation servers that are protected by Active Directory Federation Services (ADFS) in all resource partner organizations.

| **Note** |
| --- |
| As a security best practice, avoid having your federation servers directly accessible on the Internet. Consider giving your federation servers direct Internet access only when you are setting up a test lab environment or when your organization does not have a perimeter network. |

For typical corporate networks, an intranet-facing firewall is established between the corporate network and the perimeter network, and an Internet-facing firewall is often established between the perimeter network and the Internet. In this situation, the federation server sits inside the corporate network, and it is not directly accessible by Internet clients.

| **Note** |
| --- |
| Clients that are connected to the corporate network can communicate directly with the federation server through Windows Integrated authentication. |

A federation server proxy should be placed in the perimeter network before you configure your firewall servers for use with ADFS. For more information, see Where to place a federation server proxy.

**Configuring your firewall servers for a federation server**

So that the federation servers can communicate directly with federation server proxies, the intranet firewall server must be configured to allow Secure Hypertext Transfer Protocol (HTTPS) traffic from the federation server proxy to the federation server. This is a requirement because the intranet firewall server must publish the federation server using port 443 so that the federation server proxy in the perimeter network can access the federation server.

| **Note** |
| --- |
| All communications to and from clients also occur over HTTPS. |

In addition, the intranet-facing firewall server, such as a server running Internet Security and Acceleration (ISA) Server, uses a process known as server publishing to distribute Internet client requests to the appropriate corporate federation servers. This means that you must manually create a server publishing rule on the intranet server running ISA Server that publishes the clustered federation server Uniform Resource Locator (URL), for example, http://fs.adatum.com.

For more information about how to configure server publishing in a perimeter network, see Where to place a federation server proxy. For information about how to configure ISA Server to publish a server, see Create a Secure Web Publishing Rule (http://go.microsoft.com/fwlink/?LinkId=75182).

If you are using the Federated Web SSO with Forest Trust design, be sure to open the ports that are required for setting up a one-way, forest trust. For more information about network ports that trusts use, see Domain and Forest Trust Tools and Settings (http://go.microsoft.com/fwlink/?LinkId=79059).

# When to create a federation server farm

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Consider installing additional federation servers when you have a larger Active Directory Federation Services (ADFS) deployment and you want to provide fault tolerance, load balancing, or scalability to your organization's Federation Service. The act of creating two or more federation servers in the same network, configuring each of them to use the same trust policy file, and adding the public key of each server's token-signing certificates (verification certificates) to the trust policy creates a federation server farm.

Before federation servers can be grouped as a farm, they must first be clustered so that requests arriving at a single fully qualified domain name (FQDN) are routed to the various federation servers in the server farm. You can create the server cluster by deploying Network Load Balancing (NLB) inside the corporate network. This guide assumes that NLB has been configured appropriately to cluster each of the federation servers in the farm.

For more information about how to configure a cluster FQDN using Microsoft NLB technology, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?LinkID=74651).

**Best practices for deploying a federation server farm**

The following best practices are recommended for deploying a federation server in a production environment:

- If you will be deploying multiple federation servers at the same time or you know that you will be adding more servers to the farm over time, consider creating a server image of an existing federation server in the farm and then installing from that image when you need to create additional federation servers quickly.

  | **Note** |
  | --- |
  | If you do decide to use the server image method for deploying additional federation servers, you do not need to complete the tasks in Checklist: Installing a federation server every time that you want to add a new server to the farm. |

- Use NLB or some other form of clustering to allocate a single IP address for many federation server computers.

- Reserve a static IP address for each federation server in the farm and, depending on your Domain Name System (DNS) configuration, insert an exclusion for each IP address in Dynamic Host Configuration Protocol (DHCP). Microsoft NLB technology requires that each server that participates in the NLB cluster be assigned a static IP address.

- Avoid editing the trust policy from multiple federation servers at the same time.

**Configuring federation servers for a farm**

The following table describes the tasks that must be completed so that each federation server can participate in a farmed environment.

| Task | Description |
| --- | --- |

| | |
|---|---|
| Store the trust policy file on a protected network shared folder | A federation server farm consists of two or more federation servers that share the same trust policy information and token-signing certificates. If you currently have one federation server and you will be adding additional federation servers, you will have to move the location of the trustpolicy.xml file to a protected network shared folder so that it can be accessed by all new federation servers that participate in the farm.<br><br>**Note**<br>For farmed scenarios, it is important that the trust policy file be shared on a computer that does not also participate as a federation server in that farm. Microsoft NLB does not allow any of the computers that participate in a farm to communicate with one another.<br><br>After the trustpolicy.xml file has been placed in a shared folder, you protect that share with the appropriate permissions. This means that for each new federation server to successfully share a trust policy file, you must provide at least Read-only access permissions to each of the machine accounts on every federation server in the farm. The administrator of the Federation Service will be able to modify the trust policy file even though the machine accounts have Read-only permissions.<br><br>**Note**<br>Ensure that the identity of the ADFSAppPool on every federation server that participates in the farm has Read access to the shared trust policy file. |
| Obtain and share certificates | A single server authentication certificate can be obtained from a public certification authority (CA)—for example, Verisign—and then configured so that all federation servers share the same private key portion of the certificate. For more information about how to share the same certificate, see Checklist: Configuring certificates for a federation server.<br><br>Unique token-signing certificates can be acquired for each federation server or you can obtain one certificate and share it among all other federation servers in the farm. For more information, see Certificate requirements for federation servers. |
| Point to the trust policy file location | The new federation server must point to the same trust policy file that is used by other federation servers in the farm so that the new server can participate in the farm. The federation server uses the values that are specified in the Federation Service endpoint **Uniform Resource Locator (URL)** and Federation Service **Uniform Resource Identifier (URI)** fields, in the trust policy properties, to identify precisely which ADFS instance, or Federation Service, it belongs to.<br><br>**Note**<br>After you point the new federation server to the existing trust policy file location successfully (which is typically done at setup time), you do not have to alter the existing URL and URI values that are specified in the properties of the trust policy because this information is shared and it applies to every federation server in the farm. |
| Add the verification certificate to the trust policy | The public key of the token-signing certificate that is assigned to the new federation server must be added to the trust policy. The public key of the token-signing certificate is also referred to as the verification certificate.<br><br>This certificate is automatically added to the trust policy when you select an existing trust policy during the installation of the Federation Service or when you select a different token-signing certificate using the Active Directory Federation Services snap-in (click **Yes** when you are prompted to add it to the trust policy).<br><br>In this scenario, you do not have to manually add a verification certificate file to the trust policy. |
| Add Common View State Element to the web.config file | Add the `<pages enableViewStateMac="true"/>` element under the `<system.web>` section of the web.config file on each of the federation servers that participate in the farm. You must configure this element in the web.config file to post form data and to prevent view state validation failures.<br><br>For more information about this element, see How to Configure MachineKey in ASP.NET 2.0 (http://go.microsoft.com/fwlink/?LinkId=79067). |

**Enable Windows Integrated authentication in the account partner**

You must complete the following tasks in the account partner organization when you want to allow clients on the corporate

network to authenticate to any of the federation servers in the farm using Windows Integrated authentication:

1. Create a dedicated user/service account in the Active Directory forest that is located in the account partner organization. This account is necessary for the Kerberos authentication protocol to work in a farm scenario and to allow pass-through authentication on each of the federation servers. Use this account only for the purposes of the federation server farm. To ensure that this service account's function is not interrupted as a result of domain password change requirements, configure the account (in the properties of the user account) so that the password never expires. (Select the **Password never expires** check box.)

   | **Note** |
   | --- |
   | Using the Network Service account for this dedicated account will result in random failures when access is attempted through Windows Integrated authentication as a result of Kerberos tickets not validating from one server to another. |

2. On each federation server in the farm, configure the ADFSAppPool to run as the service account. In IIS Manager, navigate to the properties of **ADFSAppPool** (under the **Application Pools** folder), click the **Identity** tab, click **Configurable**, and type the new service account name and password. Repeat this task for each account federation server in the farm.

3. Because the application pool identity for ADFSAppPool is running as a domain user/service account, you must configure the service principal name (SPN) for that account in the domain using the Setspn.exe command-line tool. To obtain Setspn.exe, install the Windows Server 2003 Support Tools on a computer that is joined to the same domain where the user/service account resides. To install Support Tools, use your Windows Server 2003 CD. Navigate to the Support\Tools directory, double-click **Suptools.msi**, and follow the setup instructions. After you install the Support Tools, open a command prompt, and then type the following command:

   **setspn -a <http/<**farmclusterdnsname**> <**serviceaccountname**>**

   For example, in a scenario in which all federation servers are clustered under the DNS host name http://fs.adatum.com and the service account name that is assigned to the ADFSAppPool is named adfsaccountfarm, type the command as follows:

   **setspn -a http:/fs.adatum.com adfsaccountfarm**

   It is necessary to complete this task only once for this account.

4. After the ADFSAppPool identity is changed to the service account, set the access control lists (ACLs) on the trust policy file to allow Read access to this new account so that ADFSAppPool can read the trust policy.

# Certificate requirements for federation servers

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

In any Active Directory Federation Services (ADFS) design, various certificates must be used to secure communication and facilitate user authentications that are made between Internet clients and federation servers. Each federation server is required to have a server authentication certificate and a token-signing certificate before it can participate in ADFS communications. The trust policy requires an associated certificate, known as a verification certificate, which is the public key portion of the token-signing certificate. The following table describes each of these certificate types.

| Certificate type | Description |
|---|---|
| Token-signing certificate | A token-signing certificate is an X509 certificate. Its associated public/private key pair is used by federation servers to digitally sign all security tokens that they produce. |
| Verification certificate | A verification certificate represents the public-key portion of the token-signing certificate. It is stored in the trust policy and used by the federation server in one organization to verify that incoming security tokens have been issued by valid federation servers in this organization's farm and in other organizations. |
| Server authentication certificate | Federation servers use a server authentication certificate to secure Web services traffic for communication with Web clients or with federation server proxies. |

**Note**

Client authentication certificates are not necessary for federation servers.

**Determining your certification authority strategy**

Because the verification certificate represents the public key portion of the token-signing certificate, only the token-signing certificate and the server authentication certificate have to be issued by either a public certification authority (CA)—for example, Verisign—or a corporate CA.

ADFS does not require that token-signing and server authentication certificates be issued by a CA. However, we recommend that you not use self-signed certificates for server authentication certificates.

**Important**

Use of self-signed, Secure Sockets Layer (SSL) certificates in a production environment can allow a malicious user in an account partner to take control of federation servers in a resource partner organization. This security risk exists because self-signed certificates are root certificates. They must be added to the trusted root store of another federation server (for example, the resource federation server), which can leave that server vulnerable to attack.

After you receive a token-signing certificate and server authentication certificate from a CA, make sure that both certificates are imported into the personal certificate store of the local computer. You can import both certificates to the personal store with the Certificates snap-in in Microsoft Management Console (MMC). For more information, see Import a certificate (http://go.microsoft.com/fwlink/?linkid=20040).

As an alternative to using the Certificates snap-in, you can also import the server authentication certificate with the IIS Manager snap-in at the time that you assign the server authentication certificate to the default Web site. For more information, see Import a server authentication certificate to the default Web site.

**Note**

Before you install the Federation Service on the computer that will become the federation server, make sure that both certificates are in the Local Computer personal certificate store and the server authentication certificate is assigned to the default Web site. For more information about the order of the tasks that are required to set up a federation server, see Checklist: Installing a federation server.

Depending on your security and budget requirements, carefully consider which of your certificates will be obtained by a public CA or a corporate CA. The following figure shows the recommended CA issuers for a given certificate type. This recommendation reflects a best practice approach regarding security and cost.




**Token-signing certificates**

Federation servers require the use of token-signing certificates to prevent attackers from altering or counterfeiting security tokens in an attempt to gain unauthorized access to federated resources. Every token-signing certificate contains cryptographic private keys and public keys that are used to digitally sign (by means of the private key) a security token. Later, after they are received by a partner federation server, these keys validate the authenticity (by means of the public key) of the encrypted security token.

Because each security token is digitally signed by the account partner, the resource partner can verify that the security token was in fact issued by the account partner and that it was not modified. Digital signatures are verified with verification certificates. After the signature is verified, the resource federation server generates its own security token for its organization and signs the security token with its own token-signing certificate.

The Web server in the resource partner uses the public key of the token-signing certificate to verify that the security token is signed by the resource federation server. The Web server then allows the appropriate access to the client.

Digital signatures on security tokens are also used in the account partner when there is more than one federation server (for example, in a farm deployment). In this case the digital signatures verify the origin and integrity of the security tokens that are issued by other federation servers in the account partner.

Token-signing certificates must meet the following requirements to work with ADFS:

- The token-signing certificate must include digital signature key usage. Examples of certificate types that include this key usage are server authentication certificates or code-signing certificates. Enhanced key usage (EKU) is not required for token-signing certificates.

- The token-signing chain must be composed of only X509 certificates, and it must chain to a client-accessible, trusted root CA.

- For a token-signing certificate to successfully sign a security token, the token-signing certificate must contain a private key.

1. The ADFSAppPool identity must have access to the token-signing certificate's private key in the personal store of the local computer.

| **Note** |
| --- |
| This is taken care of by Setup. You can also use the Active Directory Federation Services snap-in to ensure this access if you subsequently change the token-signing certificate. |

- The trust policy includes the entire token-signing certificate chain in the trustpolicy.xml file. This operation takes place automatically. You do not have to configure it manually.

| Note |
| --- |
| It is a public key infrastructure (PKI) best practice to not share the private key for multiple purposes. Therefore, do not use the server authentication certificate that you installed on the federation server as the token-signing certificate. |

**Deployment considerations for token-signing certificates**

When you deploy the first federation server in a new ADFS installation, you must obtain a token-signing certificate and install it in the local computer personal certificate store on that federation server. You can obtain a token-signing certificate by requesting one from an enterprise CA or a public CA or by creating a self-signed certificate.

Token-signing certificates are installed differently, depending on how you create the server farm. There are two server farm options that you can consider when you obtain token-signing certificates for your deployment:

- A private key from one token-signing certificate is shared among all the federation servers in a farm.

  In a federation server farm environment, we recommend that all federation servers share (or reuse) the same token-signing certificate. You can install a single token-signing certificate from a CA on a federation server and then export the private key, as long as the issued certificate is marked as exportable.

  As shown in the following figure, the private key from a single token-signing certificate can be shared to all the federation servers in a farm. This option—compared to the following "unique token-signing certificate" option—reduces costs if you plan to obtain a token-signing certificate from a public CA.

- There is a unique token-signing certificate for each federation server in a farm.

  When multiple unique certificates are used throughout your farm, each server in that farm signs tokens with its own unique private key.

  As shown in the following figure, you can obtain a separate token-signing certificate for every single federation server in the farm. This option is more expensive if you plan to obtain your token-signing certificates from a public CA.




For information about installing token-signing certificates when you use Microsoft Certificate Services as your enterprise CA, see Submit an advanced certificate request via the Web to a Windows Server 2003 CA (http://go.microsoft.com/fwlink/?linkid=64020). For information about installing a token-signing certificate from a public CA, contact your public CA. For information about creating self-signed certificates, see Create a self-signed, code-signing certificate (http://go.microsoft.com/fwlink/?linkid=68171).

**Verification certificates**

Verification certificates verify that a security token was issued by a valid federation server and that the token was not modified. Verification certificates are the public-key portion of your own token-signing certificate as well as the certificates of other federation servers that you trust. As shown in the following figure, verification certificates are exported (from the token-signing certificate) and added to a verification certificate list in the trust policy at the time when you first assign a federation server to an existing trust policy. This process takes place automatically during the installation of the Federation Service. You do not have to add the certificate manually.




To verify that a security token was issued by a given federation server and not modified, the federation server must have a verification certificate for the federation server that issued the security token. For example, if fs.adatum.com (a federation server in the account partner) creates and signs a security token and then sends it to fs.treyresearch.net (a federation server in the resource partner), fs.treyresearch.net must see the verification certificate for fs.adatum.com in the verification certificate list that is stored in its own trust policy.

If the certificate has been issued by a CA, we require that:

1. The certificate revocation lists (CRLs) of the certificate be accessible to resource partners and Web agents that trust the federation server.

2. The root CA certificate be trusted by the resource partners and Web agents that trust the federation server.

**Server authentication certificates**

Federation servers require the use of server authentication certificates so that clients can establish the server's identity because the federation server presents the client with a server authentication certificate that discloses its source. In this way, a client can verify that the data that is transmitted is usable only by the organization that is identified by the certificate.

Server authentication certificates must meet the following requirements to work with ADFS:

- The server authentication certificate must include the server authentication EKU extension.

- The CRLs must be accessible for all the certificates in the chain from the server authentication certificate to the root CA certificate. The root CA must also be trusted by any federation server proxies and Web agents that trust this federation

server.

- The subject name that is used in the server authentication certificate must match the Domain Name System (DNS) name of the Federation Service endpoint Uniform Resource Locator (URL) in the trust policy.

**Deployment considerations for server authentication certificates**

Configure server authentication certificates so that all federation servers use the same certificate. If you are deploying either of the two Federated Web SSO designs, we recommend that your server authentication certificate be issued by a public CA. You can request and install these certificates through the Internet Information Services (IIS) snap-in.

You can use self-signed server authentication certificates successfully on federation servers in a test lab environment. However, for a production environment, we recommend that you obtain server authentication certificates from a public CA. The following list provides reasons why you should not use self-signed server authentication certificates for a live deployment:

- A self-signed Secure Sockets Layer (SSL) certificate must be added to the trusted root store on each of the federation servers in the resource partner organization. While this alone does not enable an attacker to compromise a resource federation server, trusting self-signed certificates does increase the attack surface of a computer, and it can lead to security vulnerabilities if the certificate signer is not trustworthy.

- It creates a bad user experience. Clients will receive Security Alert prompts when they try to access federated resources that display the message "The security certificate was issued by a company you have not chosen to trust." This is expected behavior, because the self-signed certificate is not trusted.

  | **Note** |
  | --- |
  | If necessary, you can work around this condition by using Group Policy to manually push down the self-signed certificate to the trusted root store on each client computer that will attempt to access an ADFS-enabled site. |

- CAs provide additional certificate-based features, such as private key archive, renewal, and revocation, that are not provided by self-signed certificates.

**Certificate revocation lists**

If any certificate that you use has CRLs, the server with the configured certificate must be able to contact the server that distributes the CRLs. The type of CRL determines what ports are used.

# Name resolution requirements for federation servers

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

When clients on the corporate network attempt to access an Active Directory Federation Services (ADFS)-secured application, they must first authenticate to a federation server. One way to authenticate is to have the corporate network clients access a local federation server through Windows Integrated authentication.

**Configure corporate DNS**

So that successful name resolution through Windows Integrated authentication on local federation servers can occur, Domain Name System (DNS) in the corporate network of the account partner must be configured for a new host record that will resolve the fully qualified domain name (FQDN) host name of the federation server to the IP address of the federation server cluster.

In the following illustration, you can see how this task is accomplished for a given scenario. In this scenario, Microsoft Network Load Balancing (NLB) provides a single cluster FQDN name and a single cluster IP address for an existing federation server farm.




For information about how to configure a cluster IP address or cluster FQDN using NLB, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?LinkId=75282).

For information about how to configure corporate DNS for a federation server, see Add a host (A) record to corporate DNS for a federation server.

For information about how to configure federation server proxies in the perimeter network, see Name resolution requirements for federation server proxies.

# Planning Federation Server Proxy Placement

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

After you gather all of the information that you will use to design your Active Directory Federation Services (ADFS) infrastructure and after you plan your federation server and Web server strategy, you can plan when and where to place federation server proxies in your new design. The information in the following topics can help you determine when and where to place a federation server proxy and whether to use either the account partner role or the resource partner role.

- When to create a federation server proxy
- Review the role of the federation server proxy in the account partner organization
- Review the role of the federation server proxy in the resource partner organization
- Where to place a federation server proxy
- When to create a federation server proxy farm

| **Note** |
| --- |
| Although this information might help with your placement planning for federation server proxies, it does not explain how to determine the proper number of proxies and the proxy hardware requirements for each ADFS design. For more information about determining the proper number of federation server proxies for each design, see Planning for federation server proxy capacity. |

For examples of how a federation server proxy can be placed in any of the three primary ADFS design scenarios, see Mapping Your Deployment Goals to an ADFS Design.

**See Also**
**Concepts**
Certificate requirements for federation server proxies
Name resolution requirements for federation server proxies

# When to create a federation server proxy

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Creating a federation server proxy in your organization adds additional security layers to your Active Directory Federation Services (ADFS) deployment. Consider deploying a federation server proxy in your organization's perimeter network when you want to:

- Prevent external clients from directly accessing your federation servers. By deploying a federation server proxy in your perimeter network, you effectively isolate your federation servers so that they can be accessed only by clients that are logged in to the corporate network and through federation server proxies, which act on behalf of external clients. Federation server proxies do not have access to the keys that are used to produce tokens.

- Provide a convenient way to differentiate the sign-in experience for users who are coming from the Internet as opposed to users who are coming from your corporate network (using Windows Integrated authentication). A federation server proxy collects credentials or home realm details from Internet clients by using the logon (clientlogon.aspx), logout (signout.aspx), and account partner discovery (discoverclientrealm.aspx) pages that are stored on the federation server proxy.

  In contrast, clients coming from the corporate network encounter a different experience, based on the configuration of the federation server. The corporate network federation server is often configured for Windows Integrated authentication, which provides a seamless sign-in experience for users on the corporate network.

The role that a federation server proxy plays in your organization depends on whether you place the federation server proxy in the account partner organization or in the resource partner organization. For example, when a federation server proxy is placed in the perimeter network of the account partner, its role is to collect the user credential information from browser clients. When a federation server proxy is placed in the perimeter network of the resource partner, it relays security token requests to the resource Federation Service and produces organizational security tokens in response to the security tokens that are provided by its account partners.

# Where to place a federation server proxy

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You place federation server proxies in a perimeter network to provide a protection layer from malicious users coming from the Internet. Federation server proxies are ideal for the perimeter network environment because they do not have access to the keys that are used to create tokens but they can efficiently route incoming requests to federation servers that are authorized to produce those tokens.

| **Note** |
| --- |
| It is not necessary to place a federation server proxy inside the corporate network for either the account partner or the resource partner because clients that are connected to the corporate network can communicate directly with the federation server. In this scenario, the federation server also provides federation server proxy functionality for clients coming from the corporate network. |

As is typical with perimeter networks, an intranet-facing firewall is established between the perimeter network and the corporate network, and an Internet-facing firewall is often established between the perimeter network and the Internet. In this scenario, the federation server proxy sits between both of these firewalls on the perimeter network.

**Configuring your firewall servers for a federation server proxy**

For the federation server proxy redirection process to be successful, all firewall servers must be configured to allow Secure Hypertext Transfer Protocol (HTTPS) traffic. The use of HTTPS is required because the firewall servers must publish the federation server proxy, using port 443, so that the federation server proxy in the perimeter network can access the federation server in the corporate network.

| **Note** |
| --- |
| All communications made to and from clients also occur over HTTPS. |

In addition, the Internet-facing firewall server, such as a computer running Internet Security and Acceleration (ISA) Server, uses a process known as server publishing to distribute Internet client requests to the appropriate perimeter and corporate network servers, such as federation server proxies or federation servers.

Server publishing rules determine how server publishing works—essentially, filtering all incoming and outgoing requests through the ISA Server computer. Server publishing rules map incoming client requests to the appropriate servers behind the ISA Server computer. For information about how to configure ISA Server to publish a server, see Create a Secure Web Publishing Rule (http://go.microsoft.com/fwlink/?LinkId=75182).

In the federated world of Active Directory Federation Services (ADFS), these client requests are typically made to a specific Uniform Resource Locator (URL), for example, a federation server (http://fs.adatum.com). Because these client requests are incoming from the Internet, the Internet-facing firewall server must be configured to publish the federation server URL for each federation server proxy that is deployed in the perimeter network.

**Configuring ISA Server to allow SSL**

To facilitate secure ADFS communications, you must configure ISA Server to allow Secure Sockets Layer (SSL) communications between:

- **Federation server proxies and federation servers.** An SSL channel is required for all communications between federation servers and federation server proxies. Therefore, it is a requirement that you configure ISA Server to allow an SSL connection between the corporate network and the perimeter network.

- **Clients, federation servers, and federation server proxies.** So that communications can occur between clients and federation servers or between clients and federation server proxies, you can place a computer running ISA Server in front of the federation server or federation server proxy.

  If your organization performs SSL client authentication on the federation server or federation server proxy, when you place a computer running ISA Server in front of the federation server or federation server proxy, the server must be configured for pass-through of the SSL connection because the SSL connection must terminate at the federation server or federation server proxy.

If your organization does not perform SSL client authentication on the federation server or federation server proxy, an additional option is to terminate the SSL connection at the computer running ISA Server and then re-establish an SSL connection to the federation server or federation server proxy.

| **Note** |
| --- |
| The federation server or federation server proxy requires that the connection be secured by SSL to protect the contents of the security token. |

# When to create a federation server proxy farm

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Consider installing additional federation server proxies when you have a large Active Directory Federation Services (ADFS) deployment and you want to provide fault tolerance, load-balancing, and scalability for the Federation Service Proxy. The act of creating two or more federation server proxies in the same perimeter network—configuring each of them to protect the same Federation Service—and then adding each of the servers' client authentication certificates to the trust policy creates a federation server proxy farm.

Before all the federation server proxies can function together as a farm, you must first cluster them under one IP address and one Domain Name System (DNS) fully qualified domain name (FQDN). You can cluster the servers by deploying Network Load Balancing (NLB) inside the perimeter network. The tasks in the following table require NLB to be configured appropriately to cluster the federation server proxies in the farm.

For more information about how to configure an FQDN for a cluster using Microsoft NLB technology, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?linkid=74651).

**Configuring federation server proxies for a farm**

The following table describes the tasks that must be completed so that each federation server proxy can participate in a farm.

| Task | Description |
|---|---|
| Point to the Federation Services Uniform Resource Locator (URL) | When you create the federation server proxies, you must type the same Federation Service DNS host name in the **Federation Services URL** field for all of the federation server proxies that participate in the farm. The federation server proxy uses the URL that makes up this DNS host name to determine which Federation Service it contacts.<br><br>For more information, see Install the Federation Service Proxy component of ADFS. |
| Obtain and share certificates | You can obtain a server authentication certificate from a public certification authority (CA)—for example, Verisign—and then configure the certificate so that all federation server proxies share the same private key portion of the same certificate on the default Web site for each federation server proxy. To share the certificate, you must install the same server authentication certificate on the default Web site for each federation server proxy. For more information, see Import a server authentication certificate to the default Web site.<br><br>When you obtain a client authentication certificate for a federation server proxy, you can share that certificate across all federation server proxies in the farm or you can obtain separate certificates for each of the federation server proxies in the farm.<br><br>**Note**<br>The Trust Policy user interface (UI) in the Active Directory Federation Services snap-in refers to client authentication certificates as **Federation Service Proxy (FSP) certificates**.<br><br>For more information, see Certificate requirements for federation server proxies. |
| Add the federation server proxy certificate to the trust policy | You must add the public key portion of the client authentication certificate for the federation server proxy to the trust policy on a federation server with which the federation server proxy communicates so that the Federation Service can authenticate the federation server proxy.<br><br>For more information, see Add a Federation Service Proxy (FSP) certificate to the trust policy (http://go.microsoft.com/fwlink/?LinkId=76073). |

For more information about adding new federation server proxies to create a federation server proxy farm, see Adding a New Federation Server Proxy (http://go.microsoft.com/fwlink/?LinkId=76068).

# Certificate requirements for federation server proxies

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Servers that are running the Federation Service Proxy component of Active Directory Federation Services (ADFS) are required to use the following types of certificates:

- **Secure Sockets Layer (SSL) server authentication certificates:** Federation server proxies use SSL server authentication certificates to secure Web server traffic communication with Web clients. Federation server proxies are usually exposed to computers on the Internet that are not included in your enterprise public key infrastructure (PKI). Therefore, use a server authentication certificate that is issued by a public (third-party) certification authority (CA), for example, Verisign.

  When you have a federation server proxy farm, all federation server proxy computers must use the same server authentication certificate. (For more information, see When to create a federation server proxy farm.) It is important to verify that the subject name in the server authentication certificate matches the Domain Name System (DNS) name of the Federation Service endpoint Uniform Resource Locator (URL) in the trust policy.

  For general information about using SSL certificates, see Configuring Secure Sockets Layer (http://go.microsoft.com/fwlink/?linkid=62785) and Obtaining Server Certificates (http://go.microsoft.com/fwlink/?linkid=62479).

- **SSL client authentication certificates:** Each federation server proxy uses a client authentication certificate to authenticate to the Federation Service. You can use any certificate with client authentication extended key usage (EKU) that chains to a trusted root CA on the federation server as a client authentication certificate for the federation server proxy. In addition, you must explicitly add the client authentication certificate to the trust policy. However, only the federation server proxy stores the private key that is associated with the federation server proxy client authentication certificate. You can install a client authentication certificate by connecting to an enterprise CA or by creating a self-signed certificate.

  | **Important** |
  | --- |
  | Do not use a certificate that was issued by your enterprise CA for client authentication of an Active Directory user (especially a domain administrator) because the private key is stored on the federation server proxy. Storing a private key on the federation server proxy allows an administrator or a successful attacker to assume the identity that the certificate represents. |

  For general information about installing client authentication certificates when you use Microsoft Certificate Services as your enterprise CA, see "Submit an advanced certificate request via the Web to a Windows Server 2003 CA" (http://go.microsoft.com/fwlink/?linkid=64020).

| **Note** |
| --- |
| Token-signing certificates do not have to be issued for federation server proxies. |

If any certificate that you use has certificate revocation lists (CRLs), the server with the configured certificate must be able to contact the server that distributes the CRLs. The type of CRL determines what ports are used.

# Name resolution requirements for federation server proxies

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

When Internet clients attempt to access an application that is secured by Active Directory Federation Services, the clients must first authenticate to the federation server. In most cases, the federation server is usually not directly accessible from the Internet. Therefore, Internet clients must be redirected to the federation server proxy instead. You can accomplish successful redirection by adding the appropriate Domain Name System (DNS) records to your DNS zone or zones that face the Internet.

The method that you use to redirect Internet clients to the federation server proxy depends on how you configure the DNS zone in your perimeter network or a DNS zone that you control on the Internet. Federation server proxies are intended for use in a perimeter network. They redirect Internet client requests to federation servers successfully only when DNS has been configured properly in all of the Internet-facing zones that you control. Therefore, the configuration of your Internet-facing zones—whether you have a DNS zone serving only the perimeter network or a DNS zone serving both the perimeter network and internet clients—is important.

This topic describes the steps that you can use to configure name resolution when you place a federation server proxy in your perimeter network. To determine which steps to follow, first determine which of the following DNS scenarios most closely matches the DNS infrastructure in the perimeter network of your organization. Then, follow the steps for that scenario.

**DNS zone serving only the perimeter network**
In this scenario, your organization has one or two DNS zones in the perimeter network, and your organization does not control any DNS zones on the Internet. Successful name resolution for a federation server proxy in the DNS zone serving only the perimeter network scenario depends on the following conditions:

- The federation server proxy must have a setting in the hosts file to resolve the fully qualified domain name (FQDN) of the Federation Server endpoint Uniform Resource Locator (URL) to an IP address of a federation server or federation server cluster.

- DNS in the perimeter network of the account partner must be configured so that the FQDN of the Federation Server endpoint URL resolves to the IP address of the federation server proxy.

The following illustration and corresponding steps shows how each of these conditions is achieved for a given example. In this illustration, Microsoft Network Load Balancing (NLB) technology provides a single cluster FQDN and a single cluster IP address for an existing federation server farm.




For more information about configuring a cluster IP address or a cluster FQDN using NLB, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?LinkId=75282).

**1. Configure the hosts file on the federation server proxy**

Because DNS in the perimeter network is configured to resolve all requests for fs.adatum.com to the account federation server proxy, the account federation server proxy has an entry in its local hosts file to resolve fs.adatum.com to the IP address of the actual account federation server (or cluster DNS name for the federation server farm) that is connected to the corporate network. This makes it possible for the account federation server proxy to resolve the host name fs.adatum.com to the account federation server rather than to itself—as would occur if it attempted to look up fs.adatum.com using perimeter DNS—so that the federation server proxy can communicate with the federation server.

**2. Configure perimeter DNS**

Because there is only a single Federation Service endpoint URL that clients are directed to—whether they are on an intranet or the Internet—clients on the Internet that use the perimeter DNS server must resolve the FQDN for the account federation server (fs.adatum.com) to the IP address of the account federation server proxy on the perimeter network. So that it can forward clients on to the account federation server proxy when they attempt to resolve fs.adatum.com, perimeter DNS contains a limited corp.adatum.com DNS zone with a single host (A) record for fs (fs.adatum.com) and the IP address of the account federation server proxy on the perimeter network.

For more information about how to modify the hosts file of the federation server proxy and configure DNS in the perimeter network, see Configure name resolution for a federation server proxy in a DNS zone serving only the perimeter network.

**DNS zone serving both the perimeter network and Internet clients**

In this scenario, your organization controls the DNS zone in the perimeter network and at least one DNS zone on the Internet. Successful name resolution for a federation server proxy in this scenario depends on the following conditions:

- DNS in the Internet zone of the account partner must be configured so that the FQDN of the Federation Server endpoint URL resolves to the IP address of the federation server proxy in the perimeter network.

- DNS in the perimeter network of the account partner must be configured so that the FQDN of the Federation Server endpoint URL resolves to the IP address of the federation server in the corporate network.

The following illustration and corresponding steps shows how each of these conditions is achieved for a given example.




**1. Configure perimeter DNS**

For this scenario, because it is assumed that you will configure the Internet DNS zone that you control to resolve requests that are made for a specific endpoint URL (that is, fs.adatum.com) to the federation server proxy in the perimeter network, you must also configure the zone in the perimeter DNS to forward these requests to the federation server in the corporate network.

So that clients can be forwarded to the account federation server when they attempt to resolve fs.adatum.com, perimeter DNS is configured with a single host (A) record for fs (fs.adatum.com) and the IP address of the account federation server on the corporate network. This makes it possible for the account federation server proxy to resolve the host name fs.adatum.com to the account federation server rather than to itself—as would occur if it attempted to look up fs.adatum.com using Internet DNS—so that the federation server proxy can communicate with the federation server.

**2. Configure Internet DNS**

For name resolution to be successful in this scenario, all requests from clients on the Internet to fs.adatum.com must be resolved by the Internet DNS zone that you control. Consequently, you must configure your Internet DNS zone to forward client requests for fs.adatum.com to the IP address of the account federation server proxy in the perimeter network.

For more information about how to modify the perimeter network and Internet DNS zones, see Configure name resolution for a federation server proxy in a DNS zone serving both the perimeter network and Internet clients.

# Planning for ADFS Capacity

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Capacity planning for Active Directory Federation Services (ADFS) is the process of planning for growth and forecasting peak usage periods to meet specific ADFS server role and federated application requirements. It involves analyzing results from extensive performance testing to establish the ideal resource use and transaction throughput for an ADFS server role or federated application under specific load conditions.

This section describes performance and scalability guidelines for specific ADFS server roles: federation server, federation server proxy, and ADFS-enabled Web servers. These guidelines are based on lab testing that was performed by the ADFS product team at Microsoft. The purpose of this section is to help you:

- Closely estimate the hardware needs for your organization's specific ADFS deployment.
- Get recommendations about ADFS server role and federated application capacity based on an analysis of actual test result data.

| **Note** |
| --- |
| Recommendations are identified with a **Recommendations** heading throughout this section so that you can quickly locate them and assess how they may affect your deployment decisions. |

- Understand how capacity requirements and scaling techniques can affect peak load conditions that might be handled by specific ADFS server roles and federated applications.
- Accurately project the expected peak usage for sign-in requests, plan for growth, and ensure that your ADFS deployment is capable of handling that expected peak usage.

This section does not contain formulas for determining the minimal or optimal number of federation servers, federation server proxies, or ADFS-enabled Web servers for a specific ADFS design.

To understand the full ADFS capacity planning story, see the following topics:

- Planning for federation server capacity
- Planning for federation server proxy capacity
- Planning for ADFS-enabled Web server capacity

**Terminology**

The following table describes important terms that are used in this capacity planning section of the Active Directory Federation Services Design Guide.

**Account federation server**

The federation server in the account partner organization. The account federation server issues security tokens to users based on user authentication. The server authenticates the user, pulls the relevant attributes and group membership information out of the account store, and generates and signs a security token to return to the user—either to be used in its own organization or to be sent to a partner organization.

**Federated application**

An application that is ADFS-enabled, meaning that it can be accessed by federated users.

**Resource federation server**

The federation server in the resource partner organization. The resource federation server typically issues security tokens to users based on a security token that is issued by an account federation server. The server receives the security token, verifies the signature, transforms the organizational claims based on its trust policy, generates a new security token based on information in the incoming security token, and signs the new token to return to the user and ultimately to the Web

application.

**Scaling out**
A design approach that adds additional servers to your deployment to more evenly distribute application processing load across multiple computers. Scaling out adds more servers in the anticipation of further growth, and it provides flexibility so that a server that participates in a Web farm can be taken offline for upgrades with relatively little impact on the cluster.

**Scaling up**
A design approach that upgrades an existing system or individual hardware components of a single system, such as CPU or memory, in your deployment. When a system reaches the maximum limit for the number of CPUs (or other hardware components) as it is scaled up, the only scaling option that remains is scaling out.

#### Hardware and software

This section describes the hardware and software that the ADFS product team used to perform its tests. The team used the following computer hardware and software configuration to gather performance and scalability data in tests of the federation server and ADFS-enabled Web server.

- HP Proliant DL560 G1
- Four Intel Xeon CPUs, each CPU running at 2.2 gigahertz (GHz) with hyper-threading enabled
- 4 gigabytes (GB) of memory
- Windows Server 2003 R2, Enterprise Edition

The team used the following computer hardware and software configuration to gather performance and scalability data for the federation server proxy tests.

- Dell Optiplex GX620
- One Intel Pentium 4 CPU, running at 3.4 GHz
- 2 GB of memory
- Windows Server 2003 R2, Enterprise Edition

#### Measuring ADFS server capacity

Typically, the hardware components that affect server performance and scalability are the CPU, memory, the disk, and network adapters. Fortunately, each of the ADFS components requires very little demand on memory and disk space. Network connectivity is an obvious requirement. Therefore, load tests that are performed on federation servers, federation server proxies, and ADFS-enabled Web servers concentrate on three primary areas for measuring server capacity:

- **Sign-in requests per second:** The number of sign-in requests that are processed per second on federation servers. This measurement can help you determine how many simultaneous users can sign in to a given server. You can use this measurement in conjunction with the CPU consumption measurement to understand this measurement's effect on performance.
- **CPU consumption:** The percentage by which CPU capacity is measured. This measurement can help you determine the overall CPU load that occurred based on the number of incoming sign-in requests per second.
- **Application requests per second:** The number of client requests to a federated application that are processed per second on an ADFS-enabled Web server. This measurement can help you measure the performance impact of client requests on both the federated application and the ADFS Web Agent.

# Planning for federation server capacity

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Capacity planning for federation servers helps you estimate the following:

- The appropriate hardware requirements for each federation server
- The number of federation servers to place in each organization

Federation servers issue security tokens to users. These tokens are presented to a relying party for consumption. Federation servers issue security tokens after authenticating a user or after receiving a security token that was previously issued by a partner Federation Service. A security token is requested from a Federation Service when users initially sign in to federated applications or when their security tokens expire while they are accessing federated applications.

Federation servers are designed to accommodate high-availability server farm configurations that use Microsoft Network Load Balancing (NLB) technology. Federation servers in a farm configuration can service requests independently, without accessing any common farm components for each request. Therefore, there is little overhead involved in scaling out a federation server deployment.

**Recommendations**

- For mission-critical or high-availability deployments, we recommend that you create a small federation server farm in each partner organization, with at least two federation servers per farm, to provide fault tolerance.
- With the need for high availability and the ease of scaling out federation servers, scaling out is the recommended method for handling high numbers of requests per second for a particular Federation Service. Scaling up beyond the base configuration in this guide is unlikely to produce significant capacity handling gains.

**Estimating hardware requirements**

Fortunately, memory and disk space requirements for federation servers are modest, and they are not likely to be a driving factor in hardware decisions. Memory and disk space requirements for federation servers depend largely on the type of entry —such as claims, applications, or partners—that is made to the trust policy and the number of entries for each type of entry.

As the number of claim, application, and partner entries in the trust policy grow, so does the need for more memory. The following table contains general size requirements for each type of trust policy entry as tested by the Microsoft Active Directory Federation Services (ADFS) product team.

**Memory and disk size requirements for trust policy entries**

| Trust policy entry type | Number of entries | Disk size requirements | Memory size requirements |
|---|---|---|---|
| Applications | 10 | 25 kilobytes (KB) | 75 KB |
| Claims | 100 | 13 KB | 40 KB |
| Partners* | 10 | 40 KB | 120 KB |

**Note**

*Each partner in this configuration contained five claim mappings. More claim mappings increase the size of memory that is necessary to represent the partner entry.

These requirements vary based on the configured content of each claim, application, and partner. Consider them to be general guidelines.

It is important to note that, using a minimal trust policy, the memory usage of the federation server is approximately 35 megabytes (MB). Therefore, add all calculated memory requirements on top of this base memory footprint. Whenever the trust policy is changed, a new copy of the trust policy is loaded into memory while the old copy remains active. After the new copy is loaded in its entirety, the old copy is removed. Therefore, when the trust policy is changed, the memory requirements of the service will briefly double.

### Recommendation

- If your ADFS deployment involves a particularly large trust policy, double your calculations of the memory requirements to account for trust policy changes.

#### Example

You decide that your organization needs a trust policy that contains 45,000 organization group claims, 9,000 applications, and 30 account partners with 5 group claim mappings per partner. Based on actual testing by the ADFS product team, the trust policy file size is about 28 MB with this number of entries. Memory usage for this trust policy is approximately 130 MB (adding about 95 MB to the minimum federation server memory footprint of 35 MB). When you change the trust policy, memory usage is approximately 200 MB. A server with 1 GB of RAM should be able to handle this trust policy configuration, in addition to any other operating system requirements.

Capacity recommendations for federation servers can vary considerably, depending on whether the federation server is acting in the account federation server role or resource federation server role.

#### Account federation server capacity considerations

An account federation server issues security tokens to users based on user authentication. The server authenticates a user, pulls the relevant attributes and group membership out of the account store, creates and signs a security token, and then sends the security token to the user. The security token is ultimately consumed by a resource federation server. Account federation server capacity is measured in terms of the peak number of these user sign-ins.

When you estimate the expected peak number of user sign-ins per second, one important factor is the token-caching lifetime. Consider adjusting the default values for the token-caching lifetime in the resource Federation Service for the Web application. A longer token-caching lifetime results in fewer overall sign-ins and token requests at the account federation server. We do not recommend token cache lifetimes that exceed the token expiration time of the account federation server.

For more information about how to adjust values for the token-caching lifetime, see Configuring ADFS Servers for Troubleshooting (http://go.microsoft.com/fwlink/?LinkID=74970).

#### Gathering log data from the resource partner

Determining accurate capacity requirements for account federation servers can be challenging because the applications that require security tokens from an account federation server are often located in another organization.

If your organization does not have access to the application that currently logs application sign-ins, you may have to request estimates or logs from the organization that does have access to the application so that you can more accurately assess your capacity requirements.

After you have all the application logs, analyze them for peak numbers of user sign-ins. Estimate the peak number of sign-ins per second.

#### Determining peak sign-in capacity

The following table shows test lab results for the hardware that was used by the ADFS product team at Microsoft. The table indicates the peak number of sign-in requests that were made to an account federation server, along with their impact on CPU consumption.

**Peak sign-ins per second for the account federation server**

| Scenario | Sign-in requests per second | CPU consumption, percent |
|---|---|---|
| Token issuance with Windows Integrated authentication on the account federation server | 94 | 95.93 |

You can use the numbers in this table to help calculate how many account federation servers you must have to handle the expected peak sign-ins per second. Administrators for most production deployments are interested in lower CPU usage. They can scale out their hardware accordingly.

### Recommendation

- If the expected value for peak sign-in requests per second for each server in your farm exceeds the peak value in the previous table, we recommend that you add additional federation servers until the expected number of peak sign-ins per second is well under this peak value for each server in the farm.

#### Example

Your organization has 36,000 employees that arrive between 8:00 A.M. and 9:00 A.M. on Monday mornings. The employees

typically access a federated Web application to display their 401(k) information soon after arriving to work. It is important to analyze actual logs of the 401(k) application access to determine peak usage. Analysis of the 401(k) application logs indicates peak usage of 200 sign-ins per second between 8:55 A.M. and 9:00 A.M. The account federation server is configured to authenticate employees using Windows Integrated authentication. If 200 sign-ins per second is the peak usage of the overall system, the previous table indicates that a load-balanced federation server farm, containing at least three federation servers with hardware similar to the hardware that was used in the ADFS product team tests, can adequately handle the load and guarantee high availability, while leaving room for future growth in usage.

**Resource federation server capacity considerations**

A resource federation server typically issues security tokens to users based on the security tokens that it receives from an account federation server. After the resource federation server receives a security token, it verifies the signature, maps the claims based on its trust policy, creates a new security token, signs the new security token, and then sends token to the user. Ultimately, the security token is consumed by the federated application.

If your federation server is acting in an account role as well as a resource role, use the capacity planning numbers in the "Account federation server capacity considerations" section because user authentication is more expensive than validating security tokens.

| **Note** |
| --- |
| Information that is provided here about resource federation server capacity applies only to the resource partner side of the Federated Web SSO design. |

**Gathering log data**

To estimate your resource federation server's peak number of sign-ins per second more precisely, gather and analyze the historical logs of the applications to which users will be signing in. After you have all the application logs, analyze them for peak numbers of user sign-ins.

An important factor that influences the peak number of sign-ins is the configuration of token-caching lifetimes at the resource Federation Service and the Web application. Longer token-caching lifetimes result in less sign-ins at the Federation Service. Each time a new browser session is initiated, a new sign-in occurs.

**Determining peak sign-in capacity**

The following table shows test lab results that the ADFS product team used to measure peak capacity of sign-in requests to a resource federation server, along with their impact on CPU consumption.

**Peak resource partner sign-ins per second for the resource federation server**

| Scenario | Sign-in requests per second | CPU consumption, percent |
| --- | --- | --- |
| Token issuance for an account partner security token on the federation server | 112 | 97.56 |

You can use the numbers in this table to help calculate how many resource federation servers you need to handle the expected peak sign-ins per second.

### Recommendation

- If the expected value for peak sign-in requests per second for each server in the farm exceeds the peak value shown in the previous table, we recommend that you add additional federation servers until the expected number of peak sign-ins per second is well under the value in this table for each server in the farm.

**Example**

Your organization hosts three federated Web applications for account partners to access. Before the ADFS deployment, account partner users were required to log in separately to each application. Analysis of the logs indicates that sign-ins peaked at 30, 50, and 60 sign-ins per second at each application, respectively. Further analysis shows that each application experiences its peak usage at the same time each day. The aggregated number of sign-ins per second is 140. While the single sign-on (SSO) that is provided by the Federation Service reduces the sign-ins that users perceive, the Federation Service must still issue a security token for each application. Therefore, 140 sign-ins per second is a safe number to use as an estimate. The data in the previous table indicates that 140 sign-ins per second can be handled easily by a high-availability deployment of a small, load-balanced federation server farm containing at least two federation servers with hardware similar to the hardware that was used in the ADFS product team tests—while leaving room for future growth in usage.

# Planning for federation server proxy capacity

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Capacity planning for federation server proxies helps you estimate the following:

- The appropriate hardware requirements for each federation server proxy
- The number of federation servers and federation server proxies to place in each organization

Federation server proxies provide security tokens from a protected federation server in the corporate network to federated users. A federation server proxy does not actually sign tokens or access trust policy. Therefore, the hardware requirements for the federation server proxy are usually lower than the hardware requirements for a federation server. Because every request to a federation server proxy results in a request to a federation server or federation server farm, capacity planning for federation servers and federation server proxies must be performed in parallel.

Estimating the peak sign-ins per second for the federation server proxy requires an understanding of the usage patterns of the federated users that will be signing in through the federation server proxy. In many deployments, the federated users who sign in using the federation server proxy are located on the Internet. You can estimate the peak sign-ins per second by looking at the usage patterns of these federated users on the existing Web applications that will be protected by Active Directory Federation Services (ADFS).

**Determining peak sign-in capacity**

The following table shows test lab results that the ADFS product team used to measure the peak capacity of sign-in requests to an account federation server proxy. The table also shows the impact on CPU consumption for the federation server proxy and the federation server. All sign-ins against the account federation server proxy for the tests used forms-based authentication. Secure Sockets Layer (SSL) client authentication was not tested. Because each request to the federation server proxy generates a request from the proxy to the federation server, the CPU consumption of the federation server is shown here as well.

**Peak sign-ins per second for the account federation server proxy and the account federation server**

| Scenario | Sign-in requests per second | CPU consumption, percent |
|---|---|---|
| Token issuance with forms-based authentication on the federation server proxy | 110 | 93.08 (account federation server proxy)<br><br>93.18 (account federation server) |

**Example**

You have 7200 Internet customers that use a federated purchasing application (named "Widgetbrowsing") to shop for widgets that your company makes. Previous analysis of the application logs indicates that Widgetbrowsing peaks at 50 requests per second every day. Because this is an Internet scenario, users sign in through forms-based authentication on a federation server proxy. The previous table indicates that one ADFS server, similar to the server that the Microsoft ADFS product team used for testing, can easily handle the 50-request-per-second load. However, the business depends on customers being able to sign in and browse widgets to purchase. This situation demands high availability. Therefore, you deploy a small federation server farm, containing at least two federation servers, and a small federation server proxy farm, containing at least two federation server proxies to balance the load and to prevent a single point of failure. Depending on usage, you might be able to scale down your federation server proxies slightly because your organizational load is much less than the peak capacity of the server that the ADFS product team used in its tests.

# Planning for ADFS-enabled Web server capacity

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Capacity planning for Active Directory Federation Services (ADFS)–enabled Web servers helps you estimate the following:

- The appropriate hardware requirements for each ADFS-enabled Web server
- The number of ADFS-enabled Web servers to place in the resource partner organization

A federated application is a Web application that has been configured to be authenticated through the ADFS Web Agent on a Web server. When a Web server is running the ADFS Web Agent, it is referred to as an ADFS-enabled Web server. You can use an ADFS-enabled Web server to federate Windows NT token-based applications or claims-aware applications.

When a federated user obtains a security token and presents it to the ADFS-enabled Web server, the Web server receives the security token, validates the signature on that token, writes an authentication cookie, and then redirects the user back to the originally requested Uniform Resource Locator (URL). As the user browses the application, the ADFS-enabled Web server also makes information about the user in the security token available to the application.

Because Web applications vary widely in their performance characteristics, capacity planning for ADFS-enabled Web servers includes testing a federated application under load while the ADFS Web Agent is enabled. Using the tables below as a guide, you can see how configuring an application to use the ADFS Web Agent affects the performance of an application.

The base application does nothing beyond displaying the user name. Secure Sockets Layer (SSL) is used throughout all scenarios to keep the scenarios consistent. Make sure that you assess the impact of using SSL before using the tables as a guide to the impact on the ADFS-enabled Web server.

The following table contains capacity numbers for the base, nonfederated application with SSL turned on. You can use these numbers as a baseline to compare to an existing SSL application that is not configured for federation.

**Capacity of a simple, nonfederated application**

| Scenario description | Requests handled per second | CPU consumption, percent |
|---|---|---|
| SSL-protected ASP.NET application using no authentication | 2,068 | 91.25 |
| SSL-protected ASP.NET application using Windows Integrated authentication | 504 | 99.79 |

Using the nonfederated application table above as a baseline, you can use the capacity numbers in the following federation application table to estimate the impact on the ADFS-enabled Web server when it runs one or more federated applications.

The first item in the following table is the capacity for initial sign-in to the Web application. During initial sign-in, the ADFS-enabled Web server validates the token, issues a cookie, and redirects to the originally accessed URL. The requests of the user after initial sign-in are labeled “continued access” in the table. The capacity numbers for continued access are categorized by the type of application (claims-aware or Windows NT token-based). Because the mix of requests during any one second is a combination of initial sign-ins and continued access requests, the actual performance falls between the numbers for initial sign-in and continued access. The following table contains capacity numbers for the same application that is described in the previous table, except that this table includes values that are based on the ADFS Web Agent being enabled.

**Capacity of a simple, federated application**

| Scenario description | Requests handled per second | CPU consumption, percent |
|---|---|---|
| Initial sign-in: authentication cookie issuance, given a security token from the Federation Service | 585 | 96.41 |
| Continued access to the Web application using the claims-based Web Agent, given the authentication cookie | 949 | 98.38 |

| | | |
|---|---|---|
| Continued access to the Web application using the Windows NT token–based Web Agent, given the authentication cookie | 712 | 98.79 |

**Example**

You are configuring a Windows NT token-based application to use the ADFS Web Agent so that users in the account partner can access the application. The current deployment of the application uses Windows Integrated authentication over SSL, and it can handle 250 requests per second. Looking at the table that describes the nonfederated application, the deployment of your application has approximately half the capacity of the sample application in the ADFS product team's tests before the ADFS Web Agent is enabled. It is reasonable to project for planning purposes that the capacity after the ADFS Web Agent is enabled will be approximately half the value that is specified in the table that describes the nonfederated application. Given this projection, you can project that performance to be between 290 requests per second (approximately half of the 585 initial sign-ins per second capacity) and 440 requests per second (approximately half of the 885 security identifier (SID)-based logon, continued access requests-per-second capacity), depending on the mix of initial sign-in requests and continued access requests in any given second.

# Finding Additional ADFS Resources

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You can find the following documentation about Active Directory Federation Services (ADFS) on the Windows Server 2003 R2 TechCenter Web Site:

- Overview of ADFS (http://go.microsoft.com/fwlink/?LinkId=63491)
- Step-by-Step Guide for Active Directory Federation Services (http://go.microsoft.com/fwlink/?LinkId=49531)
- ADFS Deployment Guide
- ADFS Design Guide
- ADFS Operations Guide (http://go.microsoft.com/fwlink/?LinkId=78683)
- Active Directory Federation Services (Product Help) (http://go.microsoft.com/fwlink/?LinkId=57765)
- Active Directory Federation Services (portal page) (http://go.microsoft.com/fwlink/?LinkId=79542)

For other ADFS-related information, see the following Web resources:

- Automate Information Access with Identity Management (http://go.microsoft.com/fwlink/?LinkId=78692)
- Web Services and Other Distributed Technologies (http://go.microsoft.com/fwlink/?LinkId=44189)
- Web Services Protocol Workshops (http://go.microsoft.com/fwlink/?LinkId=44190)
- Web Services Specifications (http://go.microsoft.com/fwlink/?LinkId=44191)
- Web Services Interoperability Organization (WS-I) (http://go.microsoft.com/fwlink/?LinkId=34328)

# Appendix A: Reviewing ADFS Requirements

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

So that the organizational partners in your Active Directory Federation Services (ADFS) deployment can collaborate successfully, you must first make sure that your corporate network infrastructure is configured to support ADFS requirements for accounts, name resolution, and certificates. ADFS has the following types of requirements:

- Hardware requirements
- Software requirements
- Browser requirements
- Network requirements
- Account store requirements
- Authentication requirements

The following sections describe these requirements.

**Hardware requirements**

The following minimum and recommended hardware requirements apply to the federation server, federation server proxy, and Web server computers.

| Hardware requirement | Minimum requirement | Recommended requirement |
|---|---|---|
| CPU speed | 133 megahertz (MHz) | 550 MHz |
| RAM | 128 megabytes (MB) | 256 MB |
| Disk space | 10 MB | 100 MB |

**Note**

For additional hardware considerations that are essential for planning the physical topology of your ADFS servers, see Planning for ADFS Capacity.

**Software requirements**

ADFS relies on server functionality that is built into the Windows Server 2003 R2 operating system. The Federation Service, Federation Service Proxy, and ADFS Web Agent components cannot run on other operating systems. The following table identifies the software requirements for each server hosting a specific ADFS component.

| Server hosting the Federation Service | Server hosting the Federation Service Proxy | Server hosting the ADFS Web Agent |
|---|---|---|
| Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Datacenter Edition | Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Datacenter Edition | Windows Server 2003 R2, Standard Edition; Windows Server 2003 R2, Enterprise Edition; or Windows Server 2003 R2, Datacenter Edition |
| Internet Information Services (IIS) | IIS | IIS |
| Microsoft ASP.NET 2.0 | Microsoft ASP.NET 2.0 | Microsoft ASP.NET 2.0 |
| Microsoft .NET Framework 2.0 | Microsoft .NET Framework 2.0 | Microsoft .NET Framework 2.0 |

| | | |
|---|---|---|
| A default Web site that is configured with Transport Layer Security / Secure Sockets Layer (TLS/SSL) | A default Web site that is configured with TLS/SSL | At least one Web site in IIS must be configured with TLS/SSL so that federated users can access Web-based applications that are hosted on the Web server. |

| **Note** |
|---|
| The Federation Service and Federation Service Proxy components cannot coexist on the same computer. |

**Browser requirements**

Although any current Web browser with JScript enabled can work as an ADFS client, only Internet Explorer 6, Internet Explorer 5 or 5.5, Mozilla Firefox, and Safari on Apple Macintosh have been tested by Microsoft. For performance reasons, we highly recommend that JScript be enabled. Cookies must be enabled, or at least trusted, for the federation servers and Web applications that are being accessed.

The ADFS product team at Microsoft successfully tested the browser and operating system configurations in the following table.

| Browser | Windows 2003 | Windows XP | Windows 2000 | Windows 98 | Macintosh OS X |
|---|---|---|---|---|---|
| Internet Explorer 6.0 | X | X | X | X | N/A |
| MSN 9.0 | X | X | X | N/A | N/A |
| Netscape 8.0 | X | X | X | X | X |
| Netscape 7.2 | X | X | X | X | X |
| Netscape 4.8 | X | X | X | X | N/A |
| FireFox 1.0.7 | X | X | X | X | N/A |
| Safari | N/A | N/A | N/A | N/A | X |

**Cookies**

ADFS creates authentication cookies that must be stored on client computers to provide single-sign-on (SSO) functionality. Therefore, the client browser must be configured to accept cookies. Authentication cookies are always Secure Hypertext Transfer Protocol (HTTPS) session cookies that are written for the originating server. If the client browser is not configured to allow these cookies, ADFS cannot function correctly.

The authentication cookie is signed but not encrypted, which is one reason why use of TLS/SSL is mandatory in ADFS.

**Network requirements**

Configuring the following network services appropriately is critical for a successful deployment of ADFS in your organization.

**TCP/IP network connectivity**

For ADFS to function, TCP/IP network connectivity must exist between the client; a domain controller; and the computers that host the Federation Service, the Federation Service Proxy (when it is used), and the ADFS Web Agent.

**DNS**

The primary network service that is critical to the operation of ADFS, other than Active Directory, is Domain Name System (DNS). When DNS is deployed, users can use friendly computer names that are easy to remember to connect to computers and other resources on IP networks.

Windows Server 2003 uses DNS for name resolution instead of the Windows Internet Name Service (WINS) network basic input/output system (NetBIOS) name resolution that was used in Windows NT 4.0-based networks. It is still possible to use WINS for applications that require it. However, Active Directory and ADFS require DNS name resolution.

The process of configuring DNS to support ADFS varies, depending on whether:

- Your organization already has an existing DNS infrastructure. In most scenarios, DNS is already configured throughout your network so that Web browser clients in your corporate network have access to the Internet. Because Internet access and name resolution are requirements of ADFS, this infrastructure is assumed to be in place for your ADFS deployment.

- You intend to add a federated server to your corporate network. For the purpose of authenticating users in the corporate network, internal DNS servers in the corporate network forest must be configured to return the CNAME of the internal server that is running the Federation Service. For more information, see Name resolution requirements for federation servers.

- You intend to add a federated server proxy to your perimeter network. When you want to authenticate user accounts that are located in the corporate network of your account partner organization, the internal DNS servers in the corporate network forest must be configured to return the canonical name (CNAME) of the internal federation server proxy. For information about how to configure DNS to accommodate the addition of federation server proxies, see Name resolution requirements for federation server proxies.

  Name resolution requirements are also necessary for ADFS-enabled Web servers. For more information, see Name resolution requirements for ADFS-enabled Web servers.

- You are setting up DNS for a test lab environment. If you plan to use ADFS in a test lab environment where no single root DNS server is authoritative, it is likely that you will have to set up DNS forwarders so that queries to names between two or more forests will be forwarded appropriately. For general information about how to set up an ADFS test lab environment, see Step-by-Step Guide for Active Directory Federation Services (http://go.microsoft.com/fwlink/?LinkId=49531).

**Account store requirements**

ADFS requires at least one account store to be used for authenticating users and extracting security claims for those users. ADFS supports two types of account stores: Active Directory and Active Directory Application Mode (ADAM).

Account store requirements depend on whether your organization is acting as the account partner (hosting the federated users) or the resource partner (hosting the federation application). Use the following table to determine the best account store type to use for your partner role in the federation.

| Account store type | Can be used by account partners who: | Is required by resource partners who: |
|---|---|---|
| Active Directory | Want to allow locally stored identities to access federated applications (both claims-aware applications and Windows NT token–based applications) in a resource partner | Need to grant access permissions for Windows NT token–based applications to federated users in an account partner. Each federated user must be mapped to a resource account or resource group in a local Active Directory domain.<br>**Note**<br>If federated users need access to only claims-aware applications, Active Directory is not required in the resource partner. |
| ADAM | Want to allow locally stored identities to access federated applications (both claims-aware applications and Windows NT token–based applications) located in a resource partner | Not required. ADAM cannot be used to map federated users to local resource accounts or resource groups. |

**ADAM**

For ADFS to operate successfully, computers that host the ADAM account store must be running Windows® 2000 Server with Service Pack 4 (SP4) and critical updates, Windows Server 2003 with Service Pack 1 (SP1), or Windows Server 2003 R2.

**Active Directory**

For ADFS to operate successfully, Active Directory domain controllers in either the account partner organization or the resource partner organization must be running Windows 2000 Server SP4 with critical updates, Windows Server 2003 SP1, or Windows Server 2003 R2.

**Important**

Because ADFS requires the installation of Internet Information Services (IIS), we strongly recommend that you not install any ADFS components on a domain controller in a production environment.

**Schema requirements**
ADFS does not require schema changes or functional-level modifications to Active Directory. To ensure that ADAM works with ADFS, install the version of ADAM that comes with Windows Server 2003 R2.

**Functional-level requirements**
ADFS does not require Active Directory functional-level modifications to operate successfully. However, if you are using Windows NT token–based applications and you want a token to be generated using Kerberos Service-for-User (S4U), the domain functional level must be Windows 2000 native or Windows Server 2003.

**Windows NT token–based application requirements**
As indicated in the previous table, Windows NT token–based applications require an Active Directory local user store. Incoming ADFS tokens can be mapped into either user accounts or groups, which the application then uses to perform authentication.

**Authentication requirements**
ADFS integrates naturally with existing Windows authentication, for example, Kerberos authentication, NTLM, smart cards, and X.509 v3 client-side certificates. Federation servers use standard Kerberos authentication to authenticate the user against the domain. Clients can authenticate by using forms-based, smart card, and Windows Integrated authentication, depending on how you configure authentication.

What ADFS can do is require SSL client authentication at an account federation server proxy. You can configure a corporate network-connected account federation server to require SSL client authentication. However, typically you configure the account federation server for Windows integrated authentication to provide a seamless Web single-sign-on (SSO) experience for corporate network users. In this situation, ADFS has no control over what credentials the user employs for Windows desktop logon.

**Smart card logon**
Although ADFS can enforce the type of credentials that it uses for authentication (passwords, SSL client authentication, or Windows Integrated authentication), it does not directly enforce authentication with smart cards. Therefore, ADFS does not provide a client-side user interface (UI) to obtain smart-card personal identification number (PIN) credentials. This is because Windows-based clients intentionally do not provide user credential details to federation servers or ADFS-protected Web servers.

**Smart card authentication**
Smart card authentication uses the Kerberos protocol to authenticate to an account federation server. ADFS cannot be extended to add new authentication methods. The certificate in the smart card is not required to chain up to a trusted root on the client computer. Use of a smart card–based certificate with ADFS requires the following conditions:

- The reader and cryptographic service provider (CSP) for the smart card must work on the computer where the browser is located.
- The smart card certificate must chain up to a trusted root on the account federation server and the account federation server proxy.
- The certificate must map to user account in Active Directory by either of the following methods:
    - The certificate subject name corresponds to the Lightweight Directory Access Protocol (LDAP) distinguished name of a user account in Active Directory.
    - The certificate subject altname extension has the user principal name (UPN) of a user account in Active Directory.

# Appendix B: Reviewing Key ADFS Concepts

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This appendix briefly explains the following key Active Directory Federation Services (ADFS) concepts:

- Certificate sources
- Claim definitions
- Claim transformation module
- Account stores
- Port configurations

**Certificate sources**
There are three sources for the certificates that you use with ADFS:

- A certification authority (CA) that is deployed in your organization
- A third-party CA
- Self-signed certificates

Each certificate source has advantages and disadvantages that are described in the following sections.

**CA in your organization**
When you implement a CA in your organization, you create a robust infrastructure that provides life-cycle management, renewal, trust management, and revocation for all the certificates in your organization, at the cost of having to deploy additional infrastructure.

**Third-party CA**
Certificates that are generated by a third-party CA have many of the benefits of implementing a CA in your organization. However, third-party certificates must be purchased. Of the certificates that are used in an ADFS deployment, the server authentication certificates for the various ADFS components are frequently third-party certificates. This makes it possible for client browsers to be configured in advance to trust the certificates.

**Self-signed certificates**
Self-signed certificates do not require the presence a CA. You must configure these certificates explicitly in certain locations on the server as trusted certificates. It is more difficult to establish an infrastructure for certificate life-cycle management, renewal, trust management, and revocation with self-signed certificates.

The various types of certificates that ADFS uses can come from any combination of these three sources. For example, you may find that your server authentication certificates are obtained from a third-party CA, but your token-signing certificate is obtained from your organization's CA.

For more information about certificates, see Public Key Infrastructure for Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=19936).

**Claim definitions**
Claims are statements that a federation server makes about a user. They are used by Web applications to make authorization decisions. Claims originate from either a local account store or an account partner.

In ADFS, there are several claim types:

- **Identity:** An identity claim defines the user or security principal that the claim set belongs to. ADFS supports three types of identity claims:

    - User principal name (UPN)
    - E-mail
    - Common name

- **Group:** This claim type indicates membership in a group or role. For example, you might define the following set of group claims: [Developer, Tester, Program Manager]. Each group claim is a separate unit of administration for claim population and mapping. It is useful to think of the value of a group claim as a Boolean value indicating membership.

- **Custom:** This claim type is a name-value pair. For example, a custom claim may indicate that a user's employee ID number is 1234.

In ADFS, there are several types of groupings of claims:

- **Organization claims:** These are claims that serve as the normalized set of claims within an organization. All internal Federation Service actions are performed on the organization claim set.

- **Incoming claims:** These are claims that a specific account partner sends to your organization.

- **Outgoing claims:** These are claims that your organization sends to a specific resource partner.

When you design an ADFS deployment, it is necessary to understand how the claims will be extracted from the account store into the organization claim set, transformed at both the account federation server and the resource federation server, and then ultimately presented to the Web application. You can use the tables in Appendix C: Documenting Your ADFS Design to define the claims that will be used by your federation servers.

The flow of claims through an ADFS deployment is depicted in the following illustration.




For more information about the role that claims play based on the organization they are being used in, see the following topics:

- Review the role of claims in the account partner organization

- Review the role of claims in the resource partner organization

**Claim transformation module**
You can use the Active Directory Federation Services snap-in on your federation servers to map claim names between the various transitions that are highlighted in the previous figure. You can use this UI to change the name of a claim during the mapping; for example, Administrators may become Admins. This way, you can share claims with partners and express claims in a common namespace that does not necessarily match the way that you have defined your own organization claims.

There may be situations, though, in which a simple mapping of claims may not be sufficient for your scenario. For these situations, you can develop a claim transformation module that modifies claim names and values as they pass through the federation server. For example, the claim transformation module might convert a claim containing a monetary value in one currency into another currency based on an exchange rate. Or, the claim transformation module might look into a sales database and determine the discount level for a partner.

If you determine that building a claim transformation module is necessary for your scenario, see the ADFS software development kit (SDK) on the MSDN Web site for additional information about how the claim transformation module functions. To find the ADFS SDK, search for "ADFS SDK" on the MSDN Home page (http://go.microsoft.com/fwlink/?LinkId=16188).

**Account stores**
ADFS requires an account store to authenticate users. The Federation Service uses Lightweight Directory Access Protocol (LDAP) to communicate with an account store. The account store also generates a claim set so that applications can make authorization decisions. In most situations the account store that ADFS will use has already been deployed and is populated with users.

The location of the user account store and the location from which users authenticate determine how you design ADFS to support the user identities. ADFS can use either Active Directory or Active Directory Application Mode (ADAM) as an account store. Depending on where the account store is located and where users will access the application (in an intranet or on the Internet), you can use one of the following options for the location of the account store:

- **Provide federated access for your employees on the corporate network:** Your organization's users access an ADFS-secured application (either your own application or a partner's application) when the users are logged on to the corporate intranet. For more information, see Provide federated access for your employees on the corporate network.

- **Provide federated access for your remote employees on the Internet:** Your organization's users access an ADFS-secured application (either your own application or a partner's application) when the users are logged on to the corporate intranet and when they log on from the Internet. For more information, see Provide federated access for your remote employees on the Internet.

- **Provide single sign-on access for customers to your hosted applications:** Your organization hosts external user accounts that must access an ADFS-secured application at your organization. For more information, see Provide single-sign-on access for customers to your hosted applications.

Depending on the requirements of your organization, you can combine several of these account store options to complete the design of your ADFS deployment.

| **Note** |
| --- |
| The Active Directory domain functional level may be set to either Windows 2000 mixed, Windows 2000 native, or Windows Server 2003. |

**Using ADAM**
The following list describes things you should know about using ADAM with ADFS:

- ADAM does not support SSL client authentication because it is implemented using a secure channel.

- ADFS cannot authenticate ADAM accounts that use parentheses as part of the account name. Accounts that have an open parenthesis in the user name cause an LDAP search failure as a result of the user name forming an invalid LDAP filter.

- We strongly recommend that the traffic between the ADAM server and the federation server be protected by TLS/SSL or other means, such as Internet Protocol security (IPsec).

- ADAM is typically used as an account store in the Web Single-Sign-On (SSO) design when you are deploying only

claims-aware applications (not Windows NT token–based applications).

- ADAM cannot be used in the resource partner when you are deploying a Windows NT token-based application that requires resource account mapping methods. For more information about resource account mapping methods, see Determine your resource account mapping method.

**Application user store requirements**

Windows NT token–based applications require the presence of an Active Directory local user store to generate a user context for the application. Claims in incoming ADFS tokens can be mapped into either user accounts or groups within this local user store. The user accounts or groups are then used by the application to perform authentication.

Claims-aware applications do not require a local user store. All information about a given identity is contained in the token that is presented by the application. The application may store additional information that links to the identity that is presented in the token, but a user account in Active Directory is not required.

**Port configurations**

Because ADFS serves Web browser clients over Secure Hypertext Transfer Protocol (HTTPS), connectivity through HTTPS must be available to the federation servers and federation server proxies. The default port for HTTPS is port 443, but other port numbers may be configured depending on your IIS configuration. Your firewalls between clients and federation servers or federation server proxies must be configured to allow HTTPS traffic.

Just as clients need HTTPS connectivity to the federation server, the federation server proxy requires HTTPS connectivity to the federation server.

If any certificate that you use has certificate revocation lists (CRLs), the server with the configured certificate must be able to contact the server that distributes the CRLs. The type of CRL determines that ports that are used.

For information about the port requirements for the Federated Web SSO with Forest Trust design, see How Domain and Forest Trusts Work (http://go.microsoft.com/fwlink/?LinkId=35356).

# Appendix C: Documenting Your ADFS Design

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You can use the following tables to document the various details of your Active Directory Federation Service (ADFS) design. Make sure that the role your organization plays in the federation agreement is clearly understood by all parties:

- If your organization is a resource provider, determine the application types and the organization claims for the organization, as well as the incoming claims for each account partner. In addition, if the resource that you are providing is a Windows NT token–based application, determine the resource accounts and groups (also known as shadow accounts or proxy groups) that will be mapped.

- If your organization is an account provider or identity provider, determine the account stores and the claim extractions for the organizations, as well as the outgoing claims for each resource partner.

- If your organization is both an account provider and a resource provider, document the requirements in the tables in all the following sections.

**Deployment goals**

Understanding the ADFS functionality that you want to enable can help you select the appropriate goals for your deployment. For each of the areas of functionality in the following table, specify whether or not your scenario requires them.

| Functionality | Yes/No |
|---|---|
| Provide federated access for your hosted applications | |
| Provide federated access for your employees on the corporate network | |
| Provide federated access for your remote employees on the Internet | |
| Provide single-sign-on access for customers to your hosted applications | |

The following table is an example of documented deployment goals.

| Functionality | Yes/No |
|---|---|
| Provide federated access for your hosted applications | Yes |
| Provide federated access for your employees on the corporate network | Yes |
| Provide federated access for your remote employees on the Internet | No |
| Provide single-sign-on access for customers to your hosted applications | No |

**Resource applications**

If your organization is hosting an application or multiple applications, use the following table to document the applications and application types that will be part of your ADFS deployment.

| Application name | Application type |
|---|---|
| | |
| | |
| | |
| | |

The following table is an example of documented resource application requirements.

| Application name | Application type |
|---|---|
| Purchasing Portal | Windows NT token–based |

| | |
|---|---|
| Ordering Application | Claims-aware |
| Sales Reports Application | Windows NT token–based |

**Account stores**

If your organization is hosting account stores, use the following table to document the account stores that will be used to access the applications.

| Account store | Account store type (internal, partner, hosted) |
|---|---|
| | |
| | |
| | |
| | |

The following table is an example of documented account store requirements.

| Account store | Account store type (internal, partner, hosted) |
|---|---|
| Corporate Active Directory | Internal account store (corporate network access) |
| Trey Research Employees | Federation account partner |
| Consolidated Messenger Customers | Hosted account store |

**Organization claims**

Organization claims are the normalized set of claims on the federation server. Use the following table to document the organization claims and claim types on your federation server.

| Organization claim | Claim type (identity, group, custom) |
|---|---|
| | |
| | |
| | |
| | |

The following table is an example of documented organization claim requirements.

| Organization claim | Claim type (identity, group, custom) |
|---|---|
| Administrators | Group |
| Purchasers | Group |
| Power Purchaser | Group |
| PurchaseLimit | Custom |
| EmployeeID | Custom |

**Claim extractions**

Claim extractions map a user or group in an account store to an organization claim. The account store can be Active Directory or Active Directory Application Mode (ADAM). If your organization is an account partner, use the following table to document the Active Directory users or groups for claim extractions and their corresponding organization claims.

| Active Directory user or group | Organization claim |
|---|---|
| | |
| | |
| | |
| | |

The following table is an example of documented claim extraction requirements.

| Active Directory user or group | Organization claim |
|---|---|
| Purchase Administrators | Administrators (Group) |
| Sales Managers (Group) | Purchasers (Group) |
| EmployeeID (Attribute) | EmployeeID (Custom) |
| John Smith (User) | Power Purchaser (Group) |

**Outgoing claims**

Organization claims on the federation server of the account partner are mapped to outgoing claims that are sent to the resource federation server. If your organization is an account partner, use the following table to document the organization claims and their corresponding outgoing claims.

| Organization claim | Outgoing claim |
|---|---|
| | |
| | |
| | |
| | |

| Note |
|---|
| Organization claims and outgoing claims can have the same names if it is not necessary for the claim names to be different. |

The following table is an example of documented outgoing claim requirements.

| Organization claim | Outgoing claim |
|---|---|
| Administrators | Admins |
| Purchasers | Allowed Purchasers |
| Power Purchaser | Power Purchaser |
| PurchaseLimit | PurchaseLimit |
| EmployeeID | EmployeeID |

**Incoming claims**

Incoming claims are received by the resource federation server from the account federation server. When incoming claims are received by the resource federation server, they are mapped to organization claims on the resource federation server. If your organization is a resource partner, use the following table to document the incoming claims and their corresponding organization claims.

| Incoming claim | Organization claim |
|---|---|
| | |
| | |
| | |
| | |

The following table is an example of documented incoming claim requirements.

| Incoming claim | Organization claim |
|---|---|
| Admins | Purchase Admins |
| Allowed Purchasers | Allowed Purchasers |
| Power Purchaser | Power Purchaser |
| PurchaseLimit | PurchaseLimit |

| | |
|---|---|
| EmployeeID | Employee Identity |

| **Note** |
|---|
| Incoming claims and organization claims can have the same names if it is not necessary for the claim names to be different. |

**Windows NT token–based application users and groups**

When the resource application is a Windows NT token–based application, the organization claims on the resource federation server must be mapped to either a user or a group in Active Directory. If your organization is a resource partner that hosts a Windows NT token–based application, use the following table to document the organization claims and the Active Directory users or groups that the claims must map to.

| Organization claim | Active Directory user or group |
|---|---|
| | |
| | |
| | |
| | |

The following table is an example of documented requirements for users or groups that use a Windows NT token–based application.

| Organization claim | Active Directory user or group |
|---|---|
| Purchase Admins (group) | Purchase Admins (group) |
| Allowed Purchasers (group) | Purchasers (group) |
| Power Purchaser (group) | Power Purchaser (user) |

# Designing Distributed File Systems

Updated: May 4, 2009

Applies To: Windows Server 2003 R2

Welcome to the design guide for the Distributed File System solution in the Microsoft® Windows Server™ 2003 R2 operating system. This preliminary guide contains design recommendations for two scenarios, data publication and data collection, in which DFS Namespaces and DFS Replication are commonly used. (Additional recommendations and scenarios will be available as they are developed.) This guide is intended for IT planners and architects who are evaluating these technologies or creating a distributed file system design for their organizations.

If you are not familiar with DFS Namespaces and DFS Replication in Windows Server 2003 R2, we recommend that you read the document titled "Overview of the Distributed File System Solution in Windows Server 2003 R2" available on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55315). The overview document describes the benefits of DFS Replication and the improvements it offers over File Replication service (FRS). It also describes the DFS Namespaces enhancements exposed in Windows Server 2003 R2. These enhancements, introduced as updated application programming interfaces (APIs) in Windows Server 2003 with Service Pack 1 (SP1), provide easier management and more flexibility for namespaces used in branch offices.

**Distributed File System Scenarios and Features**

Before you begin your design, it is helpful to understand the scenarios for which these technologies were designed and the basic features that can be configured for DFS Namespaces and DFS Replication in Windows Server 2003 R2. As you review the scenarios, keep in mind that there are no dependencies between DFS Namespaces and DFS Replication. Both services can be used independently of one another, but when used together they can help you achieve more powerful end-to-end scenarios of high availability and WAN load balancing.

**Recommended Scenarios for Using Distributed File System**

DFS Namespaces and DFS Replication can be used together to implement the following scenarios.

**Data Publication**

DFS Namespaces and DFS Replication in Windows Server 2003 R2 can be used to publish documents, software, and line-of-business data to users throughout an organization. In this scenario, data is distributed to multiple servers using DFS Replication, and user access to the data is simplified and made highly available using DFS Namespaces.

The following figure illustrates how DFS Replication can be used to replicate data in a branch office environment where data originates on one or more hub servers in a hub site or data center and replicates to servers in branch offices.




Because DFS Replication is designed for slow WAN links, DFS Replication is ideal for distributing files to branch offices in remote locations. Remote differential compression (RDC) helps reduce the amount of network bandwidth used for replication, and DFS Replication can resume replicating any partially replicated files that result from WAN interruptions. Some additional benefits of using DFS Replication for data distribution are as follows:

- To reduce the WAN traffic necessary to replicate new files that originate on the hub server, cross-file RDC can be used to identify files that are similar to the file that needs to be replicated. This is useful when a file exists on the hub server but not the branch server. Instead of replicating the entire file, DFS Replication can use portions of files that are similar to the replicating file to minimize amount of data transferred over the WAN. (Requirements for cross-file RDC are described in "DFS Replication Features" later in this guide.)

- To further reduce replication traffic needed for initial replication of new data, you can prestage the branch servers by using a restored backup. DFS Replication can use RDC and cross-file RDC to reduce the bandwidth required to replicate any new files or portions of changed files.

- You can schedule replication to occur during off-hours, and you can set bandwidth throttling to regulate how much bandwidth is used during replication.

- DFS Replication is very efficient in terms of load on the hub server when replicating to a large number of branch servers, because DFS Replication stages the file to be replicated (that is, prepares the file for replication with RDC hashes) once and re-uses the staged file to replicate to all partners.

DFS Replication provides numerous additional benefits not described here. For a complete list of these benefits, see the document titled "Overview of the Distributed File System Solution in Windows Server 2003 R2" available on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55315).

Although DFS Replication alone is sufficient to distribute data, using DFS Namespaces allows you to simplify how users access the distributed data and increases the availability of the data. When you create a namespace, you group shared folders located on different servers and present them to users as a virtual tree of folders. Any folder in the namespace can be hosted by multiple servers, each of which holds a replica (kept in sync by DFS Replication) of the published data in that folder. When browsing the namespace, users see a single folder and are not aware that the folder is hosted by multiple servers. The underlying servers are completely transparent to the users, who access the data using a UNC path, such as \\Contoso.com\Software\Products\Microsoft\Office.

There are several aspects of DFS Namespaces that work well in branch office environments. For example:

- Client computers access servers in their own Active Directory site first, if a server is present in the site. You can optionally restrict client computers so that they access only servers in their own site.

- If a server fails, clients can fail over to another server in the same site (if one exists). If no same-site servers exist, clients can fail over to a server that has the lowest connection cost (as defined in the Active Directory Sites and Services snap-in).

- After the local server is restored, you can configure the namespace so that clients fail back to the local server. (Note that this feature requires that you install a hotfix on the clients.)

The following figure illustrates two features of DFS Namespaces, client failover and target priority, as they might be used in an environment that has a hub site and two branch offices. Assume that during normal operations, the client at the branch office accesses the local server using a referral from the namespace server. If the local branch server fails, and target priority is configured, failover would occur as follows:







1. The client in the branch site attempts to access a folder target on the local server, but the server is unavailable.

2. The client attempts to fail over to one of the two servers in the hub site. If you want clients to always fail over to a particular server in the hub site, you can configure that hub server's target priority as first among targets of equal cost. The client attempts to fail over to this preferred hub server first.

3. If the preferred hub server is unavailable, the client attempts to access the other hub server, which might be a "hot standby" server for the hub site. You can configure this hub server's priority as last among all targets of equal cost.

| **Note** |
| --- |
| As described earlier, you can configure the namespace so that clients in branch offices will fail back to the branch server after it is restored. |

**Data Collection**

The data collection scenario helps address the need to eliminate the use of tape backup in branch offices. To accomplish this, DFS Replication is used to replicate data from a server in a branch office to a server in a hub site or data center. Administrators at the hub site can use backup software to back up the branch server's data from a hub server, eliminating the often error-prone process of having end users performing the backups at branch offices that are not staffed by trained IT personnel. Centralizing backups also helps reduce hardware and operational costs. The following figure illustrates this scenario.




Thanks to RDC, DFS Replication replicates only the differences (or changes) between the two servers; as a result, bandwidth use during replication is minimized, an important consideration for branch offices that use low-bandwidth WAN connections to the hub office. In addition, bandwidth throttling can be used to regulate the amount of bandwidth used during replication, providing you with more control over WAN traffic.

When DFS Replication is used in conjunction with DFS Namespaces, you can configure a namespace so that branch clients always connect to the branch server. If the branch server becomes unavailable, branch clients fail over to the hub server. And, if client failback is configured, branch clients will fail back to the branch server after it is restored. (Client failback requires a hotfix to be installed on clients. For more information, see "Client compatibility for DFS Namespaces" later in this guide.)

The following figure illustrates the client failback process, which is described in more detail after the figure.




In the previous figure, client failover and failback works as follows:

1. The local branch server or fails or is otherwise unavailable due to network problems.

2. The client fails over to the server in the hub site. The client will access this server until the client's referral expires, the client computer is rebooted, or the client's referral cache is cleared. (A referral is an ordered list of servers that host the shared folders associated with a folder in the namespace.)

3. The branch server is restored.

4. The client requests a new referral after the referral expires, the client is rebooted, or the client's referral cache is cleared. (This step is not related to the branch server's restoration.)

5. After receiving a new referral, the client fails back to the restored branch server.

Depending on your environment, you might consider using Data Protection Manager instead of DFS Replication, especially if

you have no other need for replication. We recommend using Data Protection Manager if your connections are greater than 512 kilobits per second (Kbps), all servers are in a single domain, and your organization cannot afford additional storage in the data center as needed by DFS Replication to collect the data and then protect it by Data Protection Manager. When you use DFS Replication to collect data on the hub server and Data Protection Manager to protect the data at the data center, the data in effect is duplicated on the hub server and the Data Protection Manager server.

**Features in Distributed File System**

To help you design and effectively use a distributed file system in your organization, Distributed File System provides a number of configurable settings and features in Windows Server 2003 R2. These settings and features are briefly described in the following sections.

**DFS Replication Settings and Features**

The following settings and features in DFS Replication can be customized or enabled as necessary to design a DFS Replication solution for your organization. For a complete list of DFS Replication features and benefits, see "Overview of the Distributed File System Solution in Windows Server 2003 R2" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55315).

**RDC**

Remote differential compression (RDC) is a protocol that can be used to efficiently update files 64 kilobytes (KB) or larger over a limited-bandwidth network. RDC detects insertions, removals, re-arrangements of data in files regardless of file type, enabling DFS Replication to replicate only the changes when files are updated. To compute the deltas to replicate, RDC typically works on an older version of the file with the same name that exists at the appropriate location in the replicated folder tree on the receiving member. RDC can also use a file with the same name in the Conflict and Deleted folder.

RDC is not used on files smaller than 64 KB; in this case, the file is compressed before it is replicated. You can also disable RDC on connections that are in a LAN where network bandwidth is not contended.

**Cross-file RDC**

On servers running Windows Server 2003 R2, Enterprise Edition or Datacenter Edition, an additional function of RDC, known as cross-file RDC, can be used to further reduce bandwidth usage. Cross-file RDC uses a heuristic to identify files that are similar to the file that needs to be replicated and uses the similar files as candidates for RDC. Cross-file RDC is useful when a file exists on the sending member and not the receiving member, but similar files exist on the receiving member. Instead of replicating the entire file, DFS Replication can use portions of files that are similar to the replicating file to minimize amount of data transferred over the WAN. Cross-file RDC can use multiple files as candidate files for RDC seed data.

**Replication schedule and bandwidth throttling**

DFS Replication supports replication scheduling and bandwidth throttling in 15-minute increments during a 7-day period. When specifying a replication window, you choose the replication start and stop times as well as the bandwidth to use during that window. The settings for bandwidth usage range from 16 kilobits per second (Kbps) to 256 megabits per second (Mbps) as well as full (unlimited) bandwidth. You can configure a default schedule and bandwidth that applies to all connections between members and optionally create a custom schedule and bandwidth for individual connections.

Because members of a replication group are often located in different time zones, it is important to consider the time zones of the sending and receiving members when you set the schedule. The receiving member initiates replication by interpreting the schedule either in Coordinated Universal Time (UTC) or in the receiving member's local time, depending on which setting you choose. You can choose this setting for the replication group schedule and for custom schedules on individual connections.

**Replication filters**

You can configure file and subfolder filters to prevent files and subfolders from replicating. Both types of filters are set on a per-replicated folder basis. For more information about the types of files that are filtered by default, see "Review DFS Replication Requirements" later in this guide.

**Staging folder**

DFS Replication uses staging folders to act as caches for new and changed files to be replicated from sending members to receiving members. Each replicated folder uses its own staging folder, and each staging folder has a configurable quota. The quota, which governs when files are purged based on high and low watermarks, must be carefully set based on each replicated folder's replication activity and the disk space available of the server. Guidelines for sizing the staging folder are described in each scenario; general guidelines are also described in "Review additional guidelines and considerations for DFS Replication" later in this guide.

**Conflict and Deleted folder**

DFS Replication uses a last writer wins method for determining which version of a file to keep when a file is modified on two or more members and each member has not seen the other's version. The losing file is stored in the Conflict and Deleted folder on the member that resolves the conflict. The Conflict and Deleted folder can also be used to store files that are deleted from replicated folders. Each Conflict and Deleted folder has a quota that governs when files are purged for cleanup purposes.

**Disabled memberships**

A *membership* defines the relationship between each replicated folder/member pair. Each membership has a status, either enabled or disabled. If you do not want a replicated folder to be replicated to certain members, you can disable the memberships for those members. Doing so allows you to replicate folders to only a subset of replication group members.

**DFS Namespaces Settings and Features**
The following settings and features in DFS Namespaces can be customized or enabled as necessary to design a DFS Namespaces solution for your organization. For a complete list of DFS Namespaces features and benefits, see "Overview of the Distributed File System Solution in Windows Server 2003 R2" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55315).

**Referral ordering**
A referral is an ordered list of targets, transparent to the user, that a client receives from a domain controller or namespace server when the user accesses the namespace root or a folder with targets in the namespace. The client caches the referral for a configurable period of time.

Targets in the client's Active Directory site are listed first in a referral. (Targets given the target priority "first among all targets" will be listed before targets in the client's site.) The order in which targets outside of the client's site appear in a referral is determined by one of the following referral ordering methods:

- Lowest cost
- Random order
- Exclude targets outside of the client's site

You can set referral ordering on the namespace root, and the ordering method applies to all folders with targets in the namespace. You can also override the namespace root's ordering method for individual folders with targets.

**Failover and failback**
Client failover in DFS Namespaces is the process in which clients attempt to access another target server in a referral after one of the servers fails or is removed from the namespace. Client failback is an optional feature that enables a client to fail back to a preferred, local server after it is restored.

Failback only occurs when a client has failed over to a more expensive server (in terms of site link cost) than the server that is restored. If the restored server has the same cost as the server that the client is currently connected to, failback does not occur to the restored server. For example, if there are two servers (Server 1 and Server 2) in the client's site, and Server 1 fails while the client is connected to it, the client will fail over to Server 2. However, the client will not fail back to Server 1 when it is restored, because Server 1 has the same site link cost as Server 2.

| **Note** |
| --- |
| You can view site link costs by using the Active Directory Sites and Services snap-in. |

For details about the hotfixes and operating systems that the clients and namespace servers must run for failback to work successfully, see "Client compatibility for DFS Namespaces" later in this guide.

**Target priority**
You can assign a priority to individual targets for a given namespace root or folder. This priority determines how the target is ordered in a referral. The options are:

- First among all targets
- Last among all targets
- First among targets of equal cost
- Last among targets of equal cost

It is important to note that setting target priority on a target will result in that target always being present in a referral, even in cases where you set the **Exclude targets outside of the client's site** option on the folder associated with the target.

**Redundant domain-based namespace servers**

A domain-based namespace can be hosted by multiple namespace servers to increase the availability of the namespace. Putting a namespace server in remote or branch offices also allows clients to contact a namespace server and receive referrals without having to cross expensive WAN connections.

**Root scalability mode**

To maintain a consistent domain-based namespace across namespace servers, it is necessary for namespace servers to periodically poll Active Directory to obtain the most current namespace data. If your organization will use more than 16 namespace servers to host a single namespace, we recommend that you enable root scalability mode. When this mode is enabled, namespace servers running Windows Server 2003 do not send change notification messages to other namespaces servers when the namespace changes, nor do they poll the PDC emulator every hour. Instead, they poll their closest domain controller every hour to discover updates to the namespace. (Regardless of whether root scalability mode is enabled, changes to the namespace are always made on the PDC emulator.)

Although traffic to the PDC emulator is reduced in root scalability mode, the trade-off is that namespace servers running Windows Server 2003 must wait for the hourly polling before they discover changes to the namespace. Therefore, if the namespace changes frequently, the namespace servers might provide out-of-date referrals until they poll for changes.

**Feature Requirements for Distributed File System**

DFS Namespaces and DFS Replication are two separate technologies and therefore have separate requirements. These requirements are described in the following sections.

| **Note** |
| --- |
| There are no dependencies between DFS Namespaces and DFS Replication. Both services can be used independently of one another. |

**DFS Replication Requirements**

The basic requirements for DFS Replication are:

- You must first update the Active Directory® schema to install the Active Directory classes and attributes used by DFS Replication. These schema changes are provided on disc 2 of the Windows Server 2003 R2 operating system installation discs. This schema can be applied to domain controllers running Windows Server 2003 R2, Windows Server 2003, and Windows® 2000 Server.

- All servers that participate in replication must run Windows Server 2003 R2.

- All members of a replication group must be in the same forest. You cannot enable replication across servers in different forests.

- Cross-file RDC is available only when the sending or receiving replication partner (in a pair of replicating servers) is running Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Datacenter Edition. If a pair of replicating servers are both running Windows Server 2003 R2, Standard Edition, or Windows Server 2003 R2, Web Edition, cross-file RDC is not used for replication between those members.

The full list of DFS Replication requirements is described in "DFS Replication Requirements" later in this guide. For more information about extending the schema, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55329).

**DFS Namespaces Requirements**

The requirements for DFS Namespaces vary for each setting or feature. The following table provides a high-level overview of the DFS Namespaces requirements for each configurable DFS Namespaces setting or feature. More in-depth requirements are described in "DFS Namespaces Requirements" later in this guide.

| Setting or Feature | Active Directory Requirements | Namespace Server Requirements | Client Requirements |
| --- | --- | --- | --- |

| | | | |
|---|---|---|---|
| Lowest cost referral ordering | • Sites must be set up in Active Directory. (Otherwise, all targets have the same cost.)<br>• Domain controllers must run Windows Server 2003 to provide domain-based root referrals based on cost.<br>• The domain controller acting as the Intersite Topology Generator (ISTG) in each site must run Windows Server 2003 to calculate site costs. | Namespace servers must run Windows Server 2003 or Windows Server 2003 R2 to provide referrals based on cost. | N/A |
| Client failback | N/A | Namespace servers must run Windows Server 2003 with SP1 or Windows Server 2003 R2. | Clients must run either:<br>• Windows® XP with Service Pack 2 and the Windows XP Client Failback hotfix.<br>• Windows Server 2003 SP1 and the Windows Server 2003 Client Failback hotfix. |
| Target priority | Domain controllers must run Windows Server 2003 with SP1 or Windows Server 2003 R2 to provide domain-based root referrals based on target priority. | Namespace servers must run Windows Server 2003 with SP1 or Windows Server 2003 R2 to provide referrals based on target priority. | N/A |
| Redundant domain-based namespace servers[1] | • Active Directory is required for domain-based namespaces. | Namespace servers must run Windows 2000 Server or Windows Server 2003. Namespace servers can be member servers or domain controllers. | N/A |
| Root scalability mode (domain-based namespaces only) | N/A | Namespace servers must run Windows Server 2003. | N/A |

[1]By default, the Standard Edition of Windows Server 2003 R2 and Windows Server 2003 supports one namespace per server. However, if you install the hotfix described in article 903651 in the Microsoft Knowledge Base, you can create multiple domain-based namespaces on a server running Standard Edition.

**Review Settings and Features Used in Each Distributed File System Scenario**

The following table compares the recommendations for the configurable settings and features for DFS Namespaces and DFS Replication, based on each scenario covered in this guide. More detailed information about these settings and features is available in the respective scenarios described later in this guide.

| Setting or Feature | Data Distribution | Data Collection |
|---|---|---|
| DFS Replication | Use to replicate data that originates on the hub server to the branch servers. | Use to replicate data that originates on the branch servers to the hub servers. |

| | | |
|---|---|---|
| Primary member | Choose the hub server as the primary member. | Choose the branch server as the primary member. |
| RDC | As new files are published, cross-file RDC can be used to replicate new files, based on their similarity to existing files. Cross-file RDC requires the hub or branch server to run Windows Server 2003 R2, Enterprise Edition or Datacenter Edition. | RDC is recommended so that only changes (deltas) are replicated from the branch servers to the hub server. (RDC is enabled by default.) |
| Replication schedule and bandwidth throttling | Use a staggered schedule to provide better scale-out on the hub server. | Varies; use a nightly replication schedule if you need to conserve bandwidth during the day. If you are using a namespace for failover purposes, consider whether to use a 24x7 schedule so that the hub server will be as consistent as possible with the branch server, if a branch client fails over to a hub server. |
| Replication filters | Use the default filters. | Add filters to exclude files you do not want to replicate or back up, such as .pst files, .mp3 files, and so forth. |
| Staging folders | Size the staging folder quota by evaluating the length of the replication window and the amount of data to replicate. The goal is to avoid having to restage files for subsequent replication windows. | To prevent frequent staging folder cleanups while replicating, size the staging folder quota so that it can accommodate staging files for all changes that occur during the day in order. |
| Conflict and Deleted folders | File conflicts are unlikely because data originates at the hub and is read-only for users. Therefore, file conflicts do not need to be considered in the quota size.<br><br>Use the **Move deleted files to Conflict and Deleted folder** feature if you want to restore deleted files. (This feature is enabled by default.) | File conflicts can occur if write permissions are set on the hub and branch servers. In this case, the quota size of the Conflict and Deleted folder should be chosen according to how much data you anticipate having to retrieve due to conflicts.<br><br>Using the **Move deleted files to Conflict and Deleted folder** is optional. If you have set up Shadow Copies of Shared folders, you might not need this option, though turning it on will save WAN bandwidth when a user accidentally deletes a file and then immediately restores a previous version. |
| Namespace | Use to simplify the way clients access distributed data; provides redundancy. | Optional in this scenario; used primarily for client failover and failback purposes. |
| Referral ordering | Use the **Exclude targets outside of client's site** referral ordering option and **Last among all targets** target priority to prevent clients from failing over to servers in other branch sites. | Use default ordering method (lowest cost) if a namespace is used. |
| Client failback | Optional; enable if you want supported clients to fail back to a local server. | Optional; enable if you want supported branch clients to fail back to the branch server. |
| Target priority | Use the **Exclude targets outside of client's site** referral ordering option and **Last among all targets** target priority to prevent clients from failing over to servers in other branch sites. | If the branch server is in its own site (as defined in Active Directory), target priority is not needed. Otherwise, use target priority to specify how branch and hub servers appear in referrals. |
| Redundant domain-based namespace servers | Use to increase the availability of domain-based namespaces; place namespace servers in the same sites as clients (or in sites connected by low-cost connections). | Use redundant namespace servers (one in the branch and one in the hub) if you want the namespace to continue to function when the WAN is unavailable. |

| | | |
|---|---|---|
| Root scalability mode (domain-based namespaces only) | Use if you plan to have more than 16 servers host a common namespace. | Use if you have a namespace server in each branch, the namespace servers host a common namespace, and you have more than 16 branches. |

**Distributed File System Design Process**
The following steps outline the general process for designing distributed file systems using DFS Namespaces and DFS Replication. Although the recommendations for each scenario differ, these steps generally apply to each scenario in this guide.

1. Identify data to replicate
2. Make initial namespace decisions
3. Design the replication topology
4. Plan for high availability and business continuity
5. Plan for delegation
6. Design the namespace hierarchy and functionality
7. Design replication schedules and bandwidth throttling
8. Review performance and optimization guidelines
9. Plan for DFS Replication deployment

**Plan for Data Publication**
The following sections describe the steps for planning for data publication.

**Identify Data to Replicate**
To begin planning for data publication, identify the servers and folders that contain the data you want to replicate to other servers. You will create a single replication group that contains multiple replicated folders if any of the following statements are true:

- The data to be replicated is located on multiple volumes or cannot be put into one folder tree.
- You want to replicate one or more folders to the full set of replication group members, or to a subset of replication group members.

| **Note** |
|---|
| If you want to replicate a folder to a subset of replication group members, ensure that the folder is not nested under another folder that you plan to replicate. For example, if you have a folder C:\Data that contains C:\Data\Software and C:\Data\Templates, but you want the Templates folder to be replicated to a subset of servers, do not use C:\Data as the replicated folder path. Instead, create a replicated folder for C:\Data\Software and another replicated folder for C:\Data\Tempates. |

Using a single replication group allows you to use the same topology for all published data, which simplifies the initial setup and management of replication. Each replicated folder can be replicated to all members or to a subset of members. You will need to create multiple replication groups if:

- Some replicated folders require different topologies.
- Some replicated folders require unique schedules and bandwidth throttling settings.

If the same data already exists on multiple servers that will be part of a replication group, identify the server that contains the most up-to-date version, if necessary. You will choose this server as the primary member when you configure replication. The primary member's data is considered authoritative during initial replication and will win any conflict resolution, even for files that are more up-to-date on the non-primary members.

The following table describes how prestaged files are handled during initial replication.

| File on Primary Member | File on Non-Primary Member | Result |
|---|---|---|
| File A.doc is identical to File A.doc on the non-primary member. | File A.doc is identical to File A.doc on the primary member. | The file is not replicated to the non-primary member. Minimal metadata is replicated, however, to update the DFS Replication database on the non-primary member. |
| File B.doc is more up-to-date than the version of File B.doc on the non-primary member. | File B.doc is outdated compared to the version of File B.doc on the primary member. | The primary member's version of File B.doc is considered authoritative. The version of File B.doc on the non-primary member is moved to the Conflict and Deleted folder. File B.doc from the primary member is replicated to the non-primary member. RDC and cross-file RDC can be used to replicate portions of the file to the non-primary member. |
| File C.doc does not exist on the primary member. | File C.doc exists on the non-primary member. | File C.doc on the non-primary member will be moved to the member's Preexisting folder at the end of initial replication. |
| File D.doc is outdated compared to the version of File D.doc on the non-primary member. | File D.doc is more up-to-date than the version of File D.doc on the primary member. | The primary member's version of File D.doc is considered authoritative. The version of File D.doc on the non-primary member is moved to the Conflict and Deleted folder. File D.doc from the primary member is replicated to the non-primary member. RDC and cross-file RDC can be used to replicate portions of the file to the non-primary member. |
| File E.doc does not exist on the primary member. | File E.doc is created on the non-primary member while initial replication is taking place. | File E.doc is replicated to the primary member after initial replication completes. |
| File G.doc is identical to File G.doc on the non-primary member. | File G.doc is deleted on the non-primary member while initial replication is taking place. | If File G.doc from the primary has not replicated to the non-primary member before the delete occurs, the delete does not replicate.<br><br>Otherwise, the delete replicates because the delete occurs on the primary member's version. |

Guidelines for managing preexisting files are described in "Plan for DFS Replication Deployment" later in this guide.

**Make Initial Namespace Decisions**

The data publication scenario works well when replicated folders are published in a namespace. Before you design a namespace, you need to answer some basic questions. The following flowchart will help you decide the type of namespace to create, based on your responses to the questions in the sections that follow.




**What type of namespace do you want to create?**

The size of the namespace can help you determine the type of namespace to choose. Essentially, the size of a namespace is based on the number of namespace folders that each correspond to a target. Typically you will create a namespace folder for each replicated folder, so the number of replicated folders you plan to create can help you decide the size and type of your namespace. If you plan to have fewer than 5,000 namespace folders in a namespace, we recommend using a domain-based namespace if you have deployed Active Directory. If you will create more than 5,000 namespace folders in a namespace, or Active Directory is not deployed in your organization, create a stand-alone namespace, or create multiple domain-based namespaces, each with fewer than 5,000 folders with targets. For more information about these size recommendations, see "DFS Namespaces size recommendations" later in this guide.

| **Note** |
| --- |
| You can also create "cascaded" namespaces where one namespace points to another. For details, see the question "What are the rules for using interlinks?" in the DFS FAQ on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=39465). |

**Does the namespace need to be fault-tolerant?**

A namespace hosted on a single, non-clustered server is not accessible if the server fails. Therefore, you must decide whether you want to make the namespace fault-tolerant, which prevents the situation in which the target servers are up and running but users cannot access them by using the namespace.

Each type of namespace (stand-alone or domain-based) is made fault-tolerant in a different way.

- To make a stand-alone namespace fault-tolerant, you create it on a server cluster using the Cluster Administrator snap-in.

- To make a domain-based namespace fault-tolerant, you need at least two namespace servers and two domain controllers. (A client that attempts to access a domain-based namespace will first contact a domain controller for a referral. If a domain controller is not available, the client cannot access the namespace.)

If you plan to host folder targets on the same server cluster where the stand-alone namespace is created, you will not be able to use DFS Replication on these folder targets because DFS Replication does not support server clusters.

**Is it OK for all branch clients to access a server in a central location to receive referrals?**

Decide whether it is acceptable for branch clients to access a domain controller (to request domain-based root referrals) or a namespace server (to request stand-alone root referrals and folder referrals) in a site other than the client's site, such as the hub site. Although clients will cache these referrals for a configurable amount of time, the referrals are purged if the client is rebooted, the referral expires, or the client's referral cache is flushed. The client must then contact the namespace server again for a referral and possibly a domain controller as well if the namespace is a domain-based namespace. (For more information about the referral process, see the DFS Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=36988).

If you do not want clients to request referrals from a remote domain controller or namespace server, you must place domain controllers and namespace servers in each branch office as follows:

- For domain-based namespaces, place a domain controller and namespace server in the branch office.
- For stand-alone namespaces, place a namespace server in the branch office.

For both cases, the namespace server can be the same server as an existing domain controller or file server in the client's site, so it might not be necessary to deploy new hardware. Placing these servers in the client's site will also ensure that the namespace is available if the WAN connection fails. If you don't have a namespace server (and a domain controller if the namespace is domain-based) in the client's site, the client cannot access the namespace if the WAN is down. The tradeoff for putting these servers in the client's site is that they must be monitored to ensure that they are up and running and that the namespace and domain controller roles are healthy.

**Design the Replication Topology**

The replication topology is a framework of replication paths between members. A bi-directional replication path consists of two one-way connections that replicate data in the opposite direction. For ease in recovery, we recommend always using two one-way connections (as opposed to creating a single one-way connection) between each member, and we recommend using shared folder permissions to prevent changes from being made on branch servers.

To publish data, you will likely use a hub-and-spoke topology, where one or more hub servers are located in data centers, and servers in branch offices will connect to one or more hub servers. To prevent the hub servers from becoming overloaded, we recommend that fewer than 100 spoke members replicate with the hub server at any given time. If you need more than 100 spoke members to replicate with a hub server, set up a staggered replication schedule to balance the replication load of the hub server.

For recommendations regarding the number of connections per member, see "DFS Replication limits" later in this guide.

**Plan for Backups**

Backups are essential for a highly available distributed file system, because the ultimate recovery method is to restore from backup. Therefore, you must make plans to:

- **Regularly back up the namespace**. This involves exporting the namespace configuration using the Windows Support Tool Dfsutil.exe. The recovery process involves creating a new namespace root and then importing the namespace configuration using Dfsutil.exe.
- **Regularly back up the replicated data**. Because data originates from the hub servers in this scenario, you can perform the backups at the hub site.
- **Regularly back up Active Directory**. Configuration information for DFS Replication is stored in Active Directory. Therefore, this information will be backed up as part of your regular Active Directory backup process.
- **Create an inventory of DFS Replication settings**. You can use the Dfsradmin.exe command-line tool to create a list of all replication groups, replicated folders, and their respective attributes.

**Plan for Delegation**

The permissions required to create and manage namespaces and replication groups vary according to the type of task. Although in some cases membership in the Domain Admins group is required for these tasks by default, it is possible to delegate the ability to perform almost every task associated with namespaces and DFS Replication. Therefore, it is important to determine who will need to perform these tasks so that a member of the Domain Admins group can delegate permissions as appropriate.

Specifically, you will need to determine who will perform the following tasks:

- Create stand-alone namespaces
- Create domain-based namespaces
- Manage existing namespaces
- Create replication groups and enable replication on folder targets in a namespace

- Manage replication groups

For more information about delegating the ability to create and manage namespaces and replication groups, see "DFS Replication security requirements and delegation" and "DFS Namespaces security requirements and delegation" later in this guide.

**Design the Namespace**
The following sections will help you design the namespace hierarchy, choose the referral ordering method and target priority, and configure client failback.

**Namespace hierarchy**
To design the namespace hierarchy, you choose the name of the namespace (also called the root name), the names of the folders that will appear beneath the root, and the hierarchy of the folders. The root and folder names should reflect not only your organization's needs but also the type of data you plan to distribute.

The following basic guidelines will help you choose namespace and folder names:

- A namespace name is the point beyond the server name or domain name that is at the top of the hierarchy of the logical namespace. Standardized and meaningful names at this level are very important, especially if you have more than one namespace in a domain, because the namespace name is where users enter the namespace.

- The folder names and hierarchy of a namespace must be as clear as possible to the users so that they do not follow the wrong path and have to backtrack. To minimize the number of times that users traverse a wrong path and avoid the delay that users can experience while a connection is set up with an undesired target, develop a meaningful naming scheme for folders in the namespace.

- The namespace does not have to map to the logical organization of files on the file systems; instead, the namespace should map to your organization's business needs.

- The namespace should be independent of geography. For example, creating a namespace path such as \\Contoso.com\Washington\Applications does not make sense even if the Washington users see only a subset of applications. Disable the membership of replicated folders for given servers to control this for publication.

If you plan to create a namespace with a large number of folders, the naming scheme and hierarchy is especially important so users do not have to scan a long list of folders to find the folder they are looking for. Using folders without targets (essentially namespace subfolders) will help you build a deeper hierarchy so that users can make a choice from a small number of top-level folders.

To flatten the namespace and reduce the number of folders that the user sees when browsing the namespace, consider using access-based folder enumeration, first introduced in Windows Server 2003 SP1, to hide folders that the user does not have access to. For more information about this feature, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55319).

**Referral ordering and target priority**
When a client computer attempts to access a namespace, a domain controller or namespace server provides a referral to the client. The referral contains a list of target servers that are sorted according to the currently configured ordering method and target priority. When a client accesses the namespace root or a folder in the namespace, the client attempts to access the first target at the top of the referral; the client moves to the next target if the prior target was not available.

The lowest cost ordering method is selected by default. In this method, targets are ordered as follows (assuming no target priority settings override the default behavior for this method):

1. Targets in the same Active Directory site as the client are listed in random order at the top of the referral.

2. Next, targets outside of the client’s site are listed in order of lowest cost to highest cost. Referrals with the same cost are grouped together and within each group the targets are listed in random order.

If you want to prevent branch clients from failing over to a branch server at a different branch site, select the **Exclude targets outside of the client site** ordering method for each folder with targets, and then set target priority on each hub server's folder target by selecting the **Last among all targets** target priority. The result of selecting these two options is as follows:

- The **Exclude targets outside of the client site** setting ensures that only targets within the client's site will be included in referrals.

- The **Last among all targets** setting overrides the referral ordering method by including the hub server in the referral, even if the hub server is not in the client's site. (If multiple hub servers are used as folder targets for a given folder, those hub servers will appear last in the referral and be sorted in order of lowest cost after the other targets.)

If you plan to use lowest cost referral ordering and target priority, be aware that domain controllers and namespace servers must run the operating systems described in "DFS Namespaces Requirements" earlier in this guide. If domain controllers or namespace servers are running Windows 2000 Server, they cannot provide referrals based on cost or priority. Referrals from these servers will use random referral ordering as follows (assuming no target priority settings override the default behavior for this method):

1. Targets in the same site as the client are listed in random order at the top of the referral.

2. Next, targets outside of the client's site are listed in random order. If no same-site target servers are available, the client computer is referred to a random target server no matter how expensive the connection is or how distant the target is.

The lowest cost ordering method works properly for all targets only if the **Bridge all site links** option in Active Directory is enabled. (This option, as well as site link costs, are available in the Active Directory Sites and Services snap-in.) An Intersite Topology Generator that is running Windows Server 2003 relies on the **Bridge all site links** option being enabled to generate the intersite cost matrix that the Distributed File System service requires for its site-costing functionality.

If this option is turned off, the Distributed File System service only computes cost for sites that have a direct site link from the branch location to the other sites. All sites that do not contain a direct site link will have the maximum possible cost. For example, assume that the topology between a branch sites and three regional data centers is configured as shown in the following figure:




In this figure, there is a site link between the branch site and the regional data center site 1. No site link is configured between the branch site and the other regional data centers. When a client in the branch site receives a referral, the targets are ordered as follows:

1. The server in the branch site. (Because the server is in the same site as the client, its cost is 0.)

2. The server in regional data center site 1. (This server is listed second because a site link exists between this data center and the branch office.)

3. The next 2 data centers, in random order. (These servers are listed in random order because the Distributed File System service cannot determine their site costs.)

If the **Bridge all site links** option is enabled, the servers in a referral are listed in the following order:

1. The server in the branch site.
2. The server in regional data center site 1. (Cost = 10)
3. The server in regional data center site 2. (Cost = 30)
4. The server in regional data center site 3. (Cost = 50)

**Client failback**

Client failover in DFS Namespaces is the process in which clients attempt to access another server in a referral after one of the servers fails or is removed from the namespace. This behavior can be undesirable, though, if a client fails over to a hub server and continues to access the hub server even after the branch server is restored. If you want clients to fail back to a preferred local server when it is restored, plan to enable the **Clients fail back to preferred targets** option for the root. Folders with targets will also use failback if this option is selected for the root.

**Design Replication Schedules and Bandwidth Throttling**

The replication schedule determines the days and times at which replication occurs. You can configure the schedule in 15-minute intervals on a 7-day schedule. When designing a replication schedule, you can choose between two types of schedules:

- **Replication group schedule**. This schedule applies to all connections in the replication group except connections that have a custom connection schedule.
- **Custom connection schedule**. This is a unique schedule that is applied to an individual connection.

When using DFS Replication for data distribution, you will likely set up replication windows so that replication occurs at night or other periods of low network activity, as appropriate for each branch office. If you plan to schedule replication to occur during a replication window, we recommend that the replication window, combined with the bandwidth throttling setting, allows DFS Replication to replicate a higher than average or expected peak of the number of new or changed files within the window. The amount of data that can replicate during the window will depend on the replication throughput, which is determined by a number of factors:

- The number and size of new and changed files
- The bandwidth throttling settings
- The speed of the network
- The ability to compress changes using RDC and whether cross-file RDC is used
- The size of the staging folder quota
- The speed of the disk subsystem
- Whether you have optimized the servers by placing the replicated folder and staging folders on separate disks

DFS Replication efficiency using compression will vary based on the type of files. For example, text files compress very well; files that are already compressed, such as Windows Installer .msi files, will not compress as much. RDC is also very effective against modified files, and cross-file RDC works well when files are new but similar to existing files on the receiving member.

| **Note** |
| --- |
| Files with the following extensions are not compressed: WMA, WMV, ZIP, JPG, MPG, MPEG, M1V, MP2, MP3, MPA, CAB, WAV, SND, AU, ASF, WM, AVI, Z, GZ, TGZ, and FRX. |

To determine the replication rate, perform testing in a lab environment that resembles your production environment. Use the built-in diagnostic reporting (in the DFS Management snap-in) or Dfsrdiag.exe to watch the replication backlog; the backlog

count should go to zero before the replication window closes. If a backlog remains after the replication window closes, then some servers will have stale or missing data.

If the amount of data changes exceeds what DFS Replication can replicate in a given period of time, you need to change one of these factors. You might also want to consider allowing a longer replication window, such as over a weekend, for the rare cases for when a large set of data is modified during the week. Creating an extended replication window during the weekend will allow the backlogged files accumulated during the week to be replicated.

Because members can exist in different time zones, it is also important to consider how the schedule is affected by time zones, and whether daylight savings time is in effect. Replication is always initiated by the receiving member; therefore, the schedule reflects the time at which a receiving member initiates replication with a sending member. The receiving members can interpret the schedule in one of the following ways:

- **Universal Coordinated Time (UTC)**. This option causes the receiving member to treat the schedule as an absolute clock. For example, a schedule that begins at 0800 UTC is the same for any location, regardless of time zone or whether daylight savings time is in effect for a receiving member. For example, assume that you set replication to begin at 0800 UTC. A receiving member in Eastern Standard Time would begin replicating at 3:00 A.M. local time (UTC - 5), and a receiving member in Rome would begin replicating at 9:00 A.M. local time (UTC + 1). Note that the UTC offset shifts when daylight savings time is in effect for a particular location.

- **Local time of receiving member**. This option causes the receiving member to use its local time to start and stop replication. Local time is determined by the time zone and daylight savings time status of the receiving member. For example, a schedule that begins at 8:00 A.M. will cause every receiving member to begin replicating when the local time is 8:00 A.M. Note that daylight savings time does not cause the schedule to shift. If replication starts at 9 A.M. before daylight savings time, replication will still start at 9 A.M. when daylight savings time is in effect.

When the schedule is open, replication occurs as files are changed, created, or deleted. When the schedule closes, replication stops, regardless of whether all changed or new files have replicated.

**Review Performance and Optimization Guidelines**
The following sections discuss performance and optimization guidelines for data publication scenarios.

**Optimize the staging folder quota and replication throughput**
In a typical data publication scenario, a large amount of data, such as one or more software programs, originates at a hub site and is replicated out to branch servers, typically on a staggered schedule. Because the initial amount of data to be replicated is large, and because the data will be replicated to many servers, it is important to consider how replication throughput is impacted by the size of the staging folder quota on the hub servers. The quota should be based on the following guidelines:

**Size the quota to avoid restaging files**
DFS Replication uses staging folders to act as caches for new and changed files to be replicated from sending members to receiving members. The sending member begins staging a file when it receives a request from the receiving member. The process involves reading the file from the replicated folder and building a compressed representation of the file in the staging folder. This is the *staged file*. After being constructed, the staged file is sent to the receiving member; if remote differential compression [RDC] is used, only a fraction of the staging file might be replicated. The receiving member downloads the data and builds the file in its staging folder. After the file has completed downloading on the receiving member, DFS Replication decompresses the file and installs it into the replicated folder.

Each replicated folder has its own staging folder, which by default is located under the local path of the replicated folder in the DfsrPrivate\Staging folder. The default size of each staging folder is 4,096 MB. This is not a hard limit, however. It is only a quota that is used to govern cleanup and excessive usage based on high and low watermarks (90 percent and 60 percent of staging folder size, respectively). For example, when the staging folder reaches 90 percent of the configured quota, the oldest staged files are purged until the staging folder reaches 60 percent of the configured quota.

For best performance, size the staging folder quota on the hub servers to avoid restaging files. The reason for this is as follows: if a file needs to be replicated to one branch server, and later that file needs to be replicated to another branch server (due to staggered replication schedules), one of two results can occur:

- If the staging quota is large, the file is likely already staged and does not need to be restaged. This will reduce the CPU and disk I/O needed to replicate the file to subsequent branch servers.

- If the staging quota is low, the staging file is more likely to have been purged as described earlier, and the file will need to be restaged before it is replicated. This will increase the CPU and disk I/O needed to replicate the file to subsequent branch servers.

If you plan to use staggered replication schedules, it is important to size the staging folder quota so that it is large enough to store the amount of data that will be replicated during the replication window. For example, if you plan to drop 2 gigabytes (GB) of files into the replicated folder, and you want that data to replicate within the replication window (assuming the window is long enough), set the staging quota size to at least 2 GB. That way, the full 2 GB of data will still be in the staging folder when the next replication window opens on the staggered schedule. (Note that DFS Replication compresses the staging files, so a 2 GB quota might not be necessary.)

RDC also performs better when files are kept in the staging folder. The reason for this performance increase is as follows: When a file is changed and subsequently staged in the staging folder, RDC logically "breaks" the file into portions and uses an algorithm to generate a checksum value for each portion. By comparing the checksums on the sending and receiving members, RDC identifies the mismatched checksums, which indicate that the portion has changed, and then replicates only the changed portions.

These checksums are stored as an alternate data stream on the staging file. As long as the file is kept in the staging folder, the checksums do not need to be regenerated if the file is replicated again. Generating the checksum value does require CPU overhead, which is why it is good to minimize checksum regeneration by keeping staged files. However, after the file is replicated and installed into in the replicated folder on the receiving member, the checksum is removed from the file.

| **Note** |
| --- |
| The staging file and the RDC checksum will be invalidated if the file has changed on the sending member, but this is not a frequent occurrence in the publication scenario. |

**Size the quota to avoid having the staging folder reach or exceed the configured quota**

If large files are being replicated, make sure the staging folder quota on the hub server is approximately ten times larger than the largest nine files. To understand this guideline, consider the following functions of DFS Replication:

- The staging folder is purged when it reaches 90 percent of the configured quota.
- If the staging folder size is below 100 percent of the configured quota, DFS Replication will typically replicate nine files at a time. (The 9 files consist of 5 serving threads and 4 download threads, assuming the server is participating in both data publication and data collection.)

A decrease in replication throughput can occur if the staging folder reaches 100 percent or more of its configured quota, such as when a large file needs to be replicated. In this case, the file will be staged (along with other previously staged files), which causes two events to occur:

- Staging folder cleanup is triggered, because the staging folder has exceeded 90 percent of its configured quota.
- DFS Replication stops replicating nine files at a time and instead replicates one file at a time until the staging folder is less than 90 percent of the configured quota.

To avoid this potential decrease in replication throughput, ensure that the each of the nine largest files will consume no more than 10 percent of the staging folder quota.

**Optimize the Conflict and Deleted folder quota**

Although conflicts are unlikely in a data publication scenario, you might want to use the Conflict and Deleted folder to store accidental deletions. The benefit of doing so is that when you restore the files on one member, RDC will be used on the receiving members to reconstruct the deleted files and put them back into the replicated folder without replicating them in full over the WAN, thus saving network bandwidth.

When sizing the Conflict and Deleted folder quota for deleted files, consider how much data might be accidentally deleted; in a publication scenario, this could be a large amount of data if an entire tree of program files, for example, is accidentally deleted.

**Optimize the amount of time that clients cache referrals**

For clients that are not running Windows XP with SP2 or Windows Server 2003 with SP1, the cache duration for a referral determines the earliest time that a client will request a new referral, but only if the existing referral expires before it is accessed again. Clients that use a cached referral will renew the Time to Live of the referral and thus use the referral indefinitely until the client’s referral cache is cleared or the client is restarted.

This behavior has changed for clients running Windows XP with SP2 or Windows Server 2003 with SP1. Specifically, the Time to Live value is not reset each time a client accesses a target using a cached referral. Instead, the referral expires after the Time to Live value lapses. This change has several effects:

- Clients running Windows XP with SP2 or Windows Server 2003 with SP1 will request referrals more frequently than other clients, which can cause moderately increased load on the domain-based namespace servers and domain controllers.

- Because they request new referrals more frequently, clients running Windows XP with SP2 or Windows Server 2003 with SP1 will discover namespace updates more quickly than other clients.

By default, clients cache referrals for namespace roots for 300 seconds (5 minutes) and referrals for folders with targets for 1800 seconds (30 minutes). In branch office environments, adjust these values using the following guidelines:

- If a namespace server and domain controller exists in the branch office, use the default cache duration for both namespace roots and folders with targets. You do not need to adjust these values because the clients in the branch do not need to access the WAN to request referrals.

- If a namespace server or domain controller does not exist in the branch office, and clients must access the WAN to request a referral, increase the cache duration. Doing so serves two purposes: clients will request referrals less frequently, and clients can continue to access previously visited portions of the namespace even when the WAN is down (assuming the targets servers are accessible). When increasing this value, however, keep in mind that the cache duration also affects when client failback is triggered. The longer the cache value, the longer clients will continue to access an out-of-site server if a same-site server is later restored.

**Plan for DFS Replication Deployment**

In the section "Identify Data to Replicate," you identified the servers that contain the data you want to replicate, the servers to participate in replication, the servers to be prestaged, and the server that contains the most up-to-date version of the data. Ensure that these servers will run Windows Server 2003 R2 and that the DFS Replication service is installed on these servers prior to your deployment.

You must also ensure that enough hard disk is available on each member. You need to account for the amount of data in each replicated folder, potential growth of the replicated data, and the quota sizes of the staging and Conflict and Deleted folders that are created for each replicated folder.

If time is not a concern, you can enable replication using highly throttled bandwidth and let the data replicate over a number of days or as long as necessary for initial replication to complete.

After initial replication completes, plan to review the Preexisting folder on each non-primary member so that you can delete the files to reclaim disk space (these files are never purged by DFS Replication) or move the files back into the replicated folder tree, if desired. You should also review the Conflict and Deleted folder manifest on non-primary members to identify any conflicting files (files with same name but different data as compared to the primary member) that you want to retrieve. It is important to review the conflicted or deleted files because they will be purged when the high watermark (90 percent of staging folder quota) is reached. (You can view a log of conflict files, as well as their original file names, by viewing the ConflictandDeletedManifest.xml file in the DfsrPrivate folder under the local path of the replicated folder.)

If you plan to publish replicated folders in a namespace, we recommend that you disable referrals to the non-primary members' folder targets until initial replication is complete. This will prevent users from accessing the servers until the data has completely replicated, at which point you can enable referrals to those folder targets.

**Plan for Data Collection**

The following sections describe the steps for planning for data collection.

**Identify Data to Replicate**

To begin planning for data collection, identify the servers in branch offices that contain the data that you want to replicate to a server in the hub site. Typically you will set up a replication group for each hub server/branch server pair.

Setting up individual replication groups (for each branch office) has the following benefits over setting up a single replication group to contain all branch servers and all branch-related replicated folders:

- You do not need to disable memberships to prevent one branch server's data from being replicated to another branch server. Replication performance and DFS Management snap-in performance are also improved when you do not have a

large number of disabled memberships.

- You can choose a replication schedule and bandwidth throttling settings best suited for each branch office. For example, time zone differences can come into play if the branch offices are geographically distributed, so the schedule for one branch office might not work for another. By using individual replication groups for each branch server/hub server pair, you can adjust the replication window as appropriate for each branch office.

- Because delegation is done at the replication group level, you have better control in delegating management permissions.

It is also important to consider how NTFS permissions work in the data collection scenario. Specifically, the NTFS permissions for a replicated folder on the branch server will be applied to the same replicated folder on the hub server. Because the type of data collected is typically user data, you must ensure that users can change files from only one server to avoid change conflicts. Publishing the replicated folders in a namespace can ensure that branch clients access the replicated folder in their own site first and fail over to the hub server only if the branch server is unavailable. If you do not plan to use a namespace, ensure that shared folder permissions on the hub server prevent users from changing files there.

Because files are replicated only after they are closed, using DFS Replication for data collection is not advised for databases or any other types of files that are held open for long periods of time. If the files to be replicated include .pst files, we recommend that you configure Outlook® to periodically to release locks on .pst files. Depending on the version of Outlook, you might also need to install an Outlook hotfix. If Outlook is not configured properly, sharing violations will occur and the .pst files will not be replicated because they are held open. The recommended procedures are:

- **Outlook 2000**. Refer to Knowledge Base article 222328 for information about decreasing the timeout period. This article is available in the Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55321).

- **Outlook 2003**. Contact Microsoft Product Support Services to obtain the post-SP1 hotfix package described in Knowledge Base article 839647, available on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55324). After you install the fix, follow the same process as for Outlook 2000 except use the following registry key: HKEY_CURRENT_USER\Software\Microsoft\Office\11.0\Outlook\PST key

If there are types of files that you do not want to back up and therefore want to exclude from replication, you can set up replication filters.

**Make Initial Namespace Decisions**

Although a namespace is not required for this scenario, using a namespace can provide the following benefits:

- Users can navigate the logical namespace without having to know the physical server names or shared folders hosting the data. Because of this, you can physically move data to another server without having to reconfigure applications and shortcuts and without having to re-educate users about where they can find their data.

- A namespace can provide increased data availability. If the branch server fails, clients can fail over to the hub server. However, if you use a namespace for failover purposes, you need to make sure that the hub server's version of the data is kept relatively up-to-date with respect to the branch server. If you plan to replicate data only at night (when backups are performed), you must decide whether it is acceptable for users to access the data from the previous day on the hub server.

If you do not need branch clients to fail over to the hub server if the branch server fails, you do not need to deploy a namespace. If you need failover and plan to deploy a namespace, you need to consider the following questions.

**How many namespaces do you want to create?**

For data collection, we recommend that you set up one namespace for each branch server/hub server pair if you do not want branch clients to browse to other branch servers using the namespace. If data collaboration is desired across branches, you can add the branch server targets and hub server targets to the same namespace.

| **Important** |
| --- |
| Conflicts can arise in data collaboration scenarios because DFS Replication has no distributed file locking. Therefore, using DFS Replication for data collaboration is suitable only if users in one branch office are passive readers while users in the other branch office are active writers that modify the data. |

For redundancy, the branch server and hub server can both host the namespace. Namespaces can be hosted on member servers and domain controllers, so you do not need to deploy additional hardware to host the namespace if a server already exists in the branch office. If namespace redundancy is not required, you can host the namespace on the branch server only.

If branch office servers will host more than one domain-based namespace, you must install the hotfix described in article 903651 in the Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55325) to enable servers running Windows Server 2003, Standard Edition, to host more than one domain-based namespace.

**What type of namespace do you want to create?**
In the data collection scenario, the namespace will likely be small if it is created for a specific branch office. Therefore, we recommend that you create a domain-based namespace if the following conditions are true:

- The branch server has fewer than 5,000 replicated folders. Each replicated folder will correlate to a folder in the namespace
- You have Active Directory deployed in your organization.

If you choose to create a domain-based namespace, you need to determine whether it is acceptable for clients to request referrals from namespace servers outside of the branch site and whether you want to make the namespace fault-tolerant. Both of these issues are covered in the following sections.

**Does the namespace need to be fault-tolerant?**
A namespace hosted on a single, non-clustered server is not accessible if the server fails. Therefore, you must decide whether you want to make the namespace fault-tolerant, which prevents the situation in which the target servers are up and running but users cannot access them by using the namespace.

To make a domain-based namespace fault-tolerant, you need at least two namespace servers and two domain controllers. (A client that attempts to access a domain-based namespace will first contact a domain controller for a referral. If a domain controller is not available, the client cannot access the namespace.)

If it acceptable for users to access the target servers directly instead of by using the namespace, you do not need a fault-tolerant namespace.

**Is it OK for all clients to access a server in a central location to receive referrals?**
Decide whether it is acceptable for clients to access a domain controller (to request domain-based root referrals) or a namespace server (to request folder referrals) in a site other than the client's site, such as the hub site. Although clients will cache these referrals for a configurable amount of time, the referrals are purged if the client is rebooted, the referral expires, or the client's referral cache is flushed. The client must then contact the domain controller or namespace server again for a new referral. If you do not want clients to request referrals from a remote domain controller or namespace server, you must place a domain controller, namespace server, or both in the client's site. (The namespace server can be the same server as an existing domain controller or file server in the client's site.) Doing so will also ensure that the namespace is available if the WAN connection fails; the tradeoff is that these servers must be monitored to ensure that they are up and running.

**Design the Replication Topology**
If you use the New Replication Wizard to set up a replication group for each branch server, the wizard sets up two one-way connections between the hub server and the branch server. This means that changes made at the hub server will be replicated to the branch server and vice versa, which could result in file conflicts if users are accessing both servers. Setting shared folder permissions can help prevent conflicts from occurring. Or, if you are using a namespace, you can disable referrals to the hub server so that clients are directed only to the branch server. If the branch server fails, you will need to re-enable the referrals to the hub server.

Because the hub server will be part of multiple replication groups (and will therefore host multiple replicated folders), you must take into account the fact that each replicated folder will have its own staging folder and Conflict and Deleted folder, and each of these folders has its own quota. Therefore, you need to ensure that the hub server has adequate disk space to store files in the staging and Conflict and Deleted folders.

**Plan for Backups**
Backups are essential for a highly available distributed file system, because the ultimate recovery method is to restore from backup. Therefore, you must make plans to:

- **Regularly back up the namespace**. This involves exporting the namespace configuration using the Windows Support Tool Dfsutil.exe. The recovery process involves creating a new namespace root and then importing the namespace configuration using Dfsutil.exe.

- **Regularly back up the replicated data**. Because data from the branch offices is replicated to the hub servers in this scenario, you can perform the backups at the hub site.

- **Regularly back up Active Directory**. Configuration information for DFS Replication is stored in Active Directory. Therefore, this information will be backed up as part of your regular Active Directory backup process.

- **Create an inventory of DFS Replication settings**. You can use the Dfsradmin.exe command-line tool to create a list of all replication groups, replicated folders, and their respective attributes.

**Plan for Delegation**

The permissions required to create and manage namespaces and replication groups vary according to the type of task. Although in some cases membership in the Domain Admins group is required for these tasks by default, it is possible to delegate the ability to perform almost every task associated with namespaces and DFS Replication. Therefore, it is important to determine who will need to perform these tasks so that a member of the Domain Admins group can delegate permissions as appropriate.

Specifically, you will need to determine who will perform the following tasks:

- Create stand-alone namespaces

- Create domain-based namespaces

- Manage existing namespaces

- Create replication groups and enable replication on folder targets in a namespace

- Manage replication groups

For more information about delegating the ability to create and manage namespaces and replication groups, see "DFS Replication security requirements and delegation" and "DFS Namespaces security requirements and delegation" later in this guide.

**Design the Namespace**

If you plan to create a namespace for each branch office, the namespace hierarchy can be relatively simple. The name of the namespace root can match the branch office's name or purpose, and each namespace folder can represent a replicated folder available at the branch.

Assuming the branch clients are in same site as the branch server, clients will always connect to the branch server first. This is because the namespace referral will always list same-site servers first (assuming no target priority settings override this behavior). If the branch server is unavailable, the branch clients will fail over to the hub servers. To enhance the namespace functionality, consider setting the following options.

- **Client failback**. If you want clients to fail back from the hub server to the branch server (after the branch server is restored), enable client failback for the namespace root; this setting will be inherited by all folders with targets, which represent replicated folders.

- **Target priority**. If you have two hub servers, the branch client will fail over to either hub server without preference for one or the other. If you prefer clients to fail over to a particular hub server, select the option **Last among all targets** on the non-preferred hub server's root target.

**Design Replication Schedules and Bandwidth Throttling**

In the data collection scenario, you will typically replicate at night during a replication window. Waiting to replicate changes at night is beneficial because, as changes are made throughout the day, they are not immediately replicated like they would be if the schedule were open 24 hours a day. Instead, the changes are collected (or dampened), and thus the changes are only replicated once when the replication window opens.

To determine the length of the replication window, you will need to analyze the amount of data that would need to be replicated as a result of the day's changes and take into account the time needed to back up the changed files. If the replication

window is longer than the amount of time needed to replicate the files, you can adjust the bandwidth so that less bandwidth is used during the window, or you can leave the leave the bandwidth as-is to deal with unexpected surges in replication traffic.

The downside to replicating only at night is that files on the hub server are not up-to-date during the day. If the branch server fails, users will not be able to access the most recent files on the hub server. This can be an issue if you are using a namespace and users transparently fail over to the hub server. If the hub server must be up-to-date, you will need to replicate continuously.

If the hub server will host a number of branch servers' data, consider staggering the replication schedules to reduce the load on the hub server. Staggering might also be necessary if branch offices are located in different time zones. Because replication is always initiated by the receiving member, the schedule reflects the time at which a receiving member initiates replication with a sending member. Therefore, adjust the schedules accordingly to ensure that replication occurs at off-hours for both hub servers and branch servers.

**Review Performance and Optimization Guidelines**

In the data collection scenario, changes are made throughout the day on branch servers, and those changes are queued in the staging folder for replication at night. Therefore, it is important to size the staging folder quota to accommodate the changes that will be replicated when the replication window opens. Doing so ensures that files remain in the staging folder and are not purged during the day, which can affect replication throughput if the files need to be restaged before they are replicated. Having files in the staging folder also increases RDC performance.

If large files are being replicated, make sure the staging folder quotas on the branch servers are approximately ten times larger than the largest nine files. The reasons behind this recommendation are described in the data publication scenario earlier in this guide.

| **Note** |
| --- |
| To determine the efficiency of RDC and DFS Replication, generate a diagnostic report against the branch servers. |

**Plan for DFS Replication Deployment**

If branch files are not prestaged on the hub server, initial replication can take a long time to complete. If time is not important but limiting bandwidth used is important, set the bandwidth throttling to a low setting and let the data trickle down from the branch servers to the hub server. To determine whether the hub server has received the initial files (at which point you can begin backing up the data at the hub server), check the backlog size reported in the diagnostic report, use Dfsrdiag.exe to check backlog, or manually compare the number of files on the hub and branch servers.

If you need initial replication to complete more quickly, you can prestage the data on the hub server by backing up the branch servers' data to media. Ship the media to the hub site and copy the files to the replicated folder on the hub server. Then, when you set up replication, only the changes are replicated from the branch server to the hub server.

| **Important** |
| --- |
| When you use the New Replication Group Wizard to set up a replication group for data collection, the wizard assumes that the branch server has the most up-to-date data and uses the branch server as the primary member. This means that the branch server's data will be authoritative if any file conflicts are encountered on the hub server. If for some reason you set up a multipurpose replication group using this wizard, be sure to choose the branch server as the primary member. |

For more information about prestaging, see the "Plan for DFS Replication Deployment" section in the data publication scenario earlier in this guide.

**Plan for Monitoring**

There are two methods for monitoring DFS Replication: the built-in diagnostic reporting (in the DFS Management snap-in) and the DFS Replication MOM Pack. The DFS Replication MOM Pack is designed for monitoring 50 or more servers. The diagnostic report is suitable monitoring up to 50 servers; if used for monitoring more than 25 servers, the diagnostic report might take a long time to generate, and the resulting file can be quite large and slow to open in a Web browser.

**Review DFS Replication Requirements**

The following sections describe the various requirements of DFS Replication that apply to all scenarios described in this guide.

**Active Directory requirements for DFS Replication**

- Active Directory must be deployed in the organization.
- The Active Directory schema must be updated to include the new DFS Replication classes and attributes. These schema

changes are provided on the second Windows Server 2003 R2 operating system CD. This schema can be applied to domain controllers running Windows 2000 Server, Windows Server 2003, and Windows Server 2003 R2.

- All members of a replication group must be in the same forest. You cannot enable replication across servers in different forests.

**Operating system requirements for using and managing DFS Replication**

The servers that will participate in DFS Replication must run Windows Server 2003 R2. After you install Windows Server 2003 R2, you must install the DFS Replication Service on each server that will take part in replication, and you must install the DFS Management snap-in on one server to manage replication.

You can also manage DFS Namespaces and DFS Replication from a computer running Windows XP with Service Pack 2 (SP2) by installing the Administration Tools Packs for Windows Server 2003 R2. When this pack is installed, the DFS Management snap-in is available as part of the File Server Management snap-in. For more information about installing this pack, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55225).

**DFS Replication compatibility**

- **File systems**. Replicated folders must be stored on NTFS volumes.
- **Server clusters**.On server clusters, replicated folders should be located in the local storage of a node, because the DFS Replication service is not designed to work in a coordinated way with cluster components, and the service will not fail over to another node.
- **Files not replicated**. DFS Replication does not replicate the following types of files:
    - NTFS mounted drives within the local path of the replicated folder. (However, the local path of a replicated folder can be at or under a mounted drive.)
    - Files encrypted by using encrypting file system (EFS).
    - Any reparse points except those associated with DFS Namespaces. If a file has a reparse point used for Hierarchical Storage Management (HSM) or Single Instance Store (SIS), DFS Replication replicates the underlying file but not the reparse point.
    - Files on which the temporary attribute has been set.
    - Files with case-sensitive names created by using a UNIX-compatible application and saved to a Network File System (NFS) shared folder. (NTFS supports the Portable Operating System Interface [POSIX] standard, which provides the ability to create two files that have the same name but different capitalization.) For example, if a user creates two files, FILE.doc and File.doc, on an NFS shared folder, DFS Replication will replicate the first but not the second.
- **Quota software**. If you plan to use disk quotas, follow these guidelines:
    - Disable any quotas on the DfsrPrivate folder that is automatically created under the local path of each replicated folder. Doing so ensures that disk quotas do not prevent DFS Replication from using the staging folder and Conflict and Deleted folder as configured.
    - Set quotas identically on all members.
    - Disable quotas on the \System Volume Information\DFSR folder that is in the root of the system volume and all volumes that contain replicated folders.
    - Watch for disk full events. Disk full events can be caused by quotas on certain folders even when there is sufficient space on the volume.
    - For specific guidelines for using the quotas in the File Server Resource Manager, see "Interoperability guidelines

for File Server Resource Manager" later in this guide.

- **Ultrasound and Sonar**. The Ultrasound and Sonar monitoring tools designed for File Replication service (FRS) are not compatible with DFS Replication. Use the built-in diagnostic reporting or the DFS Replication MOM pack to monitor DFS Replication.

- **Firewalls**. DFS Replication might not work across firewalls because it uses the RPC dynamic endpoint mapper. Additionally, configuring DFS Replication using the DFS Management snap-in does not work when a firewall is enabled. You must define a port exception or disable the firewall, or you can use Dfsrdiag.exe to set the static RPC port for DFS Replication. For more information, see the DFS Replication Operations Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55327).

- **SYSVOL**. DFS Replication is not supported for SYSVOL replication in Windows Server 2003 R2. Do not attempt to configure DFS Replication on SYSVOL by disabling FRS and setting up a replication group for SYSVOL. Continue to use FRS for SYSVOL replication on domain controllers running Windows Server 2003 R2. FRS and DFS Replication can coexist on the same member server or domain controller.

- **Antivirus software**. Antivirus software must be compatible with DFS Replication; contact your antivirus software vendor to check for compatibility. (With RDC, excessive replication does not occur due to incompatible virus software, but there is overhead associated with the RDC protocol.)

**DFS Replication security requirements and delegation**

The following table describes the groups that can perform basic DFS Replication tasks by default and the method for delegating the ability to perform these tasks.

| Task | Users or Groups That Can Perform This Task By Default | Delegation Method |
|---|---|---|
| Create a replication group or enable DFS Replication on a folder that has folder targets | Domain Admins group in the domain where the replication group will be created. | Right-click the **Replication** node in the console tree, and then click **Delegate Management Permissions**. |
| Administer a replication group | Domain Admins group in the domain where the replication group is configured, or the creator of the replication group. | Right-click the replication group in the console tree, and then click **Delegate Management Permissions**. |
| Add a server to a replication group[1, 2] | If the server is a member server, the user must be a member of the local Administrators group of the server to add.<br><br>If the server is a domain controller, the user must be a member of the Domain Admins group in the domain where the server is located. | Add the user to local Administrators group of the member server to add, or add the user to the Domain Admins group of the domain controller to add. |

[1] Assumes that the user has been delegated the ability to administer the replication group.

[2] The server to be added must be online.

If you plan to delegate the ability to create and administer replication groups, note the following two important considerations:

- If you delegate to a user or group the ability to create replication groups, and you later remove the user or group from the delegation list, there is no change to the security settings on existing replication groups.

- If you delegate to a user or group the ability to administer a specific replication group, and you later remove the user or group from the delegation list, there is no change to the security settings on any existing configuration data. For example, if the user who is being removed had created a connection in the replication group, then the user would still have permissions to edit that connection because the user is the owner of the Active Directory object that contains the configuration information for the connection.

**DFS Replication scalability limits**
We recommend that your deployments approach but not exceed the following tested limits:

- Each server can be a member of up to 256 replication groups.
- Each replication group can contain up to 256 replicated folders.
- Each server can have up to 256 connections (for example, 128 incoming connections and 128 outgoing connections).
- On each server, the number of replication groups multiplied by the number of replicated folders multiplied by the number of simultaneously replicating connections must be kept to 1024 or fewer. (If the replication schedule is staggered, you do not need to count the connections that are not replicating due to a closed schedule.)
- A replication group can contain up to 256 members.
- A volume can contain up to 8 million replicated files, and a server can contain up to 1 terabyte of replicated files. These are tested numbers and are recommended guidelines for performance and scalability reasons.

**Review additional guidelines and considerations for DFS Replication**
The following sections contain guidelines and considerations that apply to all DFS Replication scenarios.

**Additional information about DFS Replication staging folders**
Staging folders are an important part of the replication process. You should take the following considerations into account when sizing the staging folder quotas for your deployment.

- Staging cleanup can be triggered before the high watermarks are hit if a disk full condition is detected on a volume while the DFS Replication service is performing certain replication activities. The special cleanup will attempt to trim in half all staging folders and all Conflict and Deleted folders on the volume that encountered the disk full condition.
- Staging folders are managed on a per-replicated folder basis. If a server has many replicated folders, especially if there are several located on the same volume, staged files in the staging folders might consume a large portion of the volume and lead to disk full conditions that end users might encounter. To improve replication performance, you can dedicate a separate volume for staging folders. However, if the volume that hosts the staging folders becomes unavailable, replication will stop on all replicated folders that use the affected volume for staging.
- Staging cleanup is not necessarily bad if there are only two members and they are replicating with each other. In addition, during initial replication, frequent staging folder cleanup is expected. However, if a hub server is replicating to many partners, frequent staging cleanups are undesirable and you should consider increasing the size of the staging folder quota to avoid cleanups.

**Replicating the root of the volume**
The DFS Replication service supports configuring a volume root, such as C:\, as the local path of a replicated folder, but this configuration is not recommended for the following reasons:

- DFS Replication replicates metadata (including ACLs and attributes) that are set on the local path of the replicated folder. Volume roots are special because they always implicitly have both the hidden and system attribute bits set. If replication partners are configured to have the corresponding local path of the replicated folder to be some place other than a volume root, that replicated folder will have system and hidden attributes set.
- System folders are typically placed in the volume root. DFS Replication is designed to ignore these folders, but there might be interoperability issues with third-party applications.

**Choosing the replication group domain in cross-domain configurations**
DFS Replication supports having replicated folders on members in different domains within the same forest. When you set up DFS Replication in a multiple-domain environment, you must choose a domain in which to hold the Global Settings objects in Active Directory. Ideally you want to place the replication group in a domain that minimizes the total site link cost for each member server to the closest domain controller in the replication group's domain. Doing so reduces the cost when the DFS Replication service on each member polls Active Directory as follows:

- The DFS Replication service polls Active Directory at regular intervals (once per hour by default) to get the current configuration. The service will access the LocalSettings object, which is under the member server's computer object in Active Directory, and the objects under the GlobalSettings object.

- The DFS Replication service will also perform lightweight polling (every 5 minutes) by polling only the computer's LocalSettings object. Lightweight polling can be disabled.

- Any change detected by a lightweight poll will trigger a full poll.

**Choosing between multiple replication groups, multiple replicated folders, or multiple folders under a replicated folder**
DFS Replication offers flexibility in replicating folders among servers. For example, if you have three folders to replicate, you can:

- Create one replication group, and replicated folder, and put the three folders into the replicated folder.

- Create one replication group and three replicated folders that correspond to each of the folders you want to replicate.

- Create three replication groups, each with a single replicated folder that corresponds to one of the folders you want to replicate.

To help you choose among these options, consider the following issues:

- If you have a fast network, consider creating a single replicated folder with multiple subfolders instead of creating a replicated folder for each folder to replicate. Using multiple replicated folders increases the number of concurrent downloads, resulting in higher throughput, but the trade off is that this configuration can potentially cause heavy disk I/O usage for staging file creation, staging folder cleanup, and so forth. Creating a single replicated folder with many subfolders is a way to throttle the amount of data replicated concurrently and thus minimize disk I/O.

- Replicated folders can be disabled and enabled, allowing for selective memberships.

- Replicated folders have dedicated staging folders and Conflict and Deleted folders. Management of these folders increases as the number of replicated folders increases.

- Delegation, topology, and replication schedules can be made on a per-replication group basis.

- Performance wise, if there is the same number of replicated folders, it should not matter if they are spread across one replication group or multiple replication groups.

- There can be performance benefits to spreading the replicated folders across multiple volumes. To do this, you must use multiple replicated folders, not a single replicated folder.

- Essentially, you should consider management first and performance second when deciding the number of replication groups and replicated folders to create. Creating one replication group with multiple replicated folders might be the best balance. If the folders are relatively small, using a single replicated folder (with multiple subfolders) is fine.

- For high-priority content, creating a separate replication group will provide the flexibility of detailed control of scheduling as well as the option to force replication on a specific connection.

**Interoperability guidelines for File Server Resource Manager**
If you plan to use DFS Replication on volumes that are monitored using File Server Resource Manager, note the following guidelines:

- If you are using auto quotas on the local path of the replicated folder, disable or delete the quota that is automatically created for DfsrPrivate and its subfolders. (You might need to right-click the **Quotas** node and then click **Refresh** to cause the DfsrPrivate quota to appear in the list.)

- Do not enable quotas on the local path of the replicated folder, because doing so will also count the DfsrPrivate

subfolder. Because quotas cannot be set to ignore a particular subfolder tree, you must enable quotas on individual folders within the local path of the replicated folder.

**Review DFS Namespaces Requirements**
**Operating system and server requirements for DFS Namespaces**

- Servers where namespace management tasks are performed using the DFS Management snap-in must run Windows Server 2003 R2 or Windows XP (if the Administration Tools Packs for Windows Server 2003 R2 is installed). For more information about installing this pack, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55225).

- To support new namespace features (client failback, target priority, and enhanced delegation), all servers that host namespaces must run Windows Server 2003 with SP1 or Windows Server 2003 R2.

- To support new namespace features, all domain controllers must run Windows Server 2003 with SP1 or Windows Server 2003 R2.

- Namespaces must be created on NTFS volumes.

It is possible to use DFS Namespaces when domain controllers and namespace servers run a mix of Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 without SP1, and Windows 2000 Server, but some functionality is disabled or available inconsistently, depending on the operating systems on the servers. Some examples of mixed-mode behavior are as follows:

- If domain controllers or namespace servers are running Windows Server 2003 without SP1, they cannot provide referrals that support target priority or client failback.

- If domain controllers or namespace servers are running Windows 2000 Server, they cannot provide referrals that support target server priority or client failback, nor can they order targets by lowest cost in referrals. Additional configuration is required to enable these namespace servers and domain controllers to detect the site of each target server in the namespace. For details, see the answer to the question "What are the issues to consider when I use multiple servers running Windows 2000 Server and Windows Server 2003 to host a domain-based DFS root?" in the DFS FAQ on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=39465).

- If the DFS Management snap-in connects to a namespace server that is not running Windows Server 2003 R2 or Windows Server 2003 with SP1, none of the new configuration settings (such as client failback and target priority) can be enabled. Renaming or moving namespace folders will not work, and delegation will not be effective.

**Client compatibility for DFS Namespaces**
Clients that access namespaces must run one of the following operating systems:

- R2 versions of Windows Server 2003

- Windows Server 2003

- Windows Storage Server 2003

- Windows XP

- Windows Preinstallation Environment (Windows PE) (Can access stand-alone namespaces, but cannot access domain-based namespaces.)

- Windows 2000 Server family

- Windows 2000 Professional

- Windows NT® Server 4.0 with Service Pack 6a

- Windows NT Workstation 4.0 with Service Pack 6a

Additionally, clients that will be configured for client failback must run the following operating systems and hotfixes:

- Windows XP with Service Pack 2 and the Windows XP Client Failback hotfix.
- Windows Server 2003 with SP1 and the Windows Server 2003 Client Failback hotfix.

For more information about the Client Failback hotfix, see article 898900 in the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=53202).

**DFS Namespaces size recommendations**

- For stand-alone namespaces, we recommend a maximum of 50,000 folders with targets. This recommendation is based on startup times for the Distributed File System service for large namespaces. For deployments that will approach 50,000 folders with targets, we recommend that you test startup times on your own hardware to determine if startup time is acceptable for service level agreements (SLAs) in your organization.
- For domain-based namespaces, we recommend keeping the namespace object in Active Directory® to 5 megabytes (MB) or less (equivalent to about 5,000 folders with targets). This recommendation is aimed at reducing the time and bandwidth needed to replicate the namespace object using Active Directory replication each time the namespace changes.

| **Note** |
| --- |
| Folders without folder targets do not count against these limits. |

**DFS Namespaces security requirements and delegation**

The following table describes the groups that can perform basic namespace tasks by default, and the method for delegating the ability to perform these tasks.

| **Task** | **Groups That Can Perform This Task By Default** | **Delegation Method** |
| --- | --- | --- |
| Create a domain-based namespace | Domain Admins group in the domain where the namespace is configured | Right-click the **Namespace** node in the console tree, and then click **Delegate Management Permissions**. You must also add the user to the Local Administrators group on the namespace server. |
| Add a namespace server to a domain-based namespace | Domain Admins group in the domain where the namespace is configured | Right-click the domain-based namespace in the console tree, and then click **Delegate Management Permissions**. You must also add the user to the Local Administrators group on the namespace server to be added. |
| Manage a domain-based namespace | Local Administrators group on each namespace server | Right-click the domain-based namespace in the console tree, and then click **Delegate Management Permissions**. |
| Create a stand-alone namespace | Local Administrators group on the namespace server | Add the user to the local Administrators group on the namespace server. |
| Manage a stand-alone namespace* | Local Administrators group on the namespace server | Right-click the stand-alone namespace in the console tree, and then click **Delegate Management Permissions**. |

| | | |
|---|---|---|
| Create a replication group or enable DFS Replication on a folder | Domain Admins group in the domain where the namespace is configured | Right-click the **Replication** node in the console tree, and then click **Delegate Management Permissions**. |

**Note**

Delegating management permissions to manage a stand-alone namespace does not grant the user the ability to view and manage security using the **Delegation** tab. This issue occurs because the DFS Management snap-in cannot retrieve the discretionary access control lists (DACLs) for the stand-alone namespace from the registry. To enable the snap-in to display delegation information, you must follow the steps in article 314837 in the Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=46803).

# Windows Server 2003 Deployment Guide

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The *Microsoft Windows Server 2003 Deployment Guide* provides guidelines and recommended processes for designing and deploying Windows Server 2003 family technologies to meet your business needs and IT goals.

This guide includes the following books:

Planning, Testing, and Piloting Deployment Projects

Deploying Network Services

Designing and Deploying Directory and Security Services

Planning Server Deployments

Automating and Customizing Installations

Designing a Managed Environment

This guide includes the following technical references:

Windows Server 2003 Resource Kit Registry Reference

Windows Server 2003 Performance Counters Reference

You must install the following resources from the *Windows Server 2003 Deployment Kit* companion CD or download them from the Microsoft Web site:

Resource Kit Tools for Windows Server 2003 (http://go.microsoft.com/fwlink/?linkid=20334)

Deploying Internet Information Services (IIS) 6.0, in the *Internet Information Services (IIS) 6.0 Resource Kit* (http://go.microsoft.com/fwlink/?LinkId=22803)

**Note**

- The *Windows Server 2003 Deployment Kit* is only available in English, regardless of the language of your operating system. If you are running the kit on a non-English operating system or an operating system with a Multilingual User Interface Pack (MUI), you will see some English content mixed with non-English content in the Help and Support Center table of contents.

# Automating and Customizing Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Welcome to *Automating and Customizing Installations* of the *Microsoft Windows Server 2003 Deployment Kit*.

This book provides comprehensive information about planning, designing, and implementing automated installations in medium and large organizations. Options range from automated installations of a basic operating system to complex installations of a customized operating system and applications. The technologies and tools discussed in this book include: unattended installation, image-based installation with the System Preparation (Sysprep) tool, and Remote Installation Services (RIS). IT professionals can use the guidelines discussed in this book to create a functional specification that describes how to automate the installation of Windows Server 2003 and Windows XP Professional.

This book contains the following chapters:

Choosing an Automated Installation Method

Designing Unattended Installations

Designing Image-based Installations with Sysprep

Designing RIS Installations

Migrating User State

# Choosing an Automated Installation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Automated installations are faster, easier, less expensive, and more consistent than having users or IT professionals install the operating system manually. You can design and deploy automated installations by using one of three automated installation methods that are included with the Microsoft® Windows® Server 2003 family of operating systems. You can determine which method to use by evaluating your available resources, the existing or planned infrastructure, and the requirements of the configurations you plan to deploy.

**In This Chapter**


Overview of Choosing an Automated Installation Method

Choosing a Method Based on Clean Installations and Upgrades

Choosing a Method Based on Software Considerations

Choosing a Method Based on Network and Hardware Configurations

Choosing a Method Based on Directory Services Considerations

Additional Resources for Choosing an Automated Installation Method


**Related Information**

- For more information about designing unattended installations, see "Designing Unattended Installations" in this book.
- For more information about designing Remote Installation Services (RIS) installations, see "Designing RIS Installations" in this book.
- For more information about designing Sysprep-based installations, see "Designing Image-based Installations with Sysprep" in this book.

# Overview of Choosing an Automated Installation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are three automated deployment methods you can use to perform automated operating system installations: Remote Installation Services (RIS), the System Preparation tool (Sysprep.exe), and the Unattended Setup tool (Winnt32.exe). You can use the automated installation tools included in the Windows Server 2003 family to automate and customize your corporate client or server operating system deployments. For organizations with many computers, automating installations is more efficient and cost-effective than using the interactive Setup program.

You need to design the client and server configurations that you want to deploy in your organization before you perform automated installations of the Windows Server 2003 or Microsoft® Windows® XP operating system. This includes designing the configuration of all networking, directory services, and security components. You need this client and server design information to customize your automated installation, as well as to help you decide which method is best to use. To choose which automated installation method is the best, you need to have access to complete information about the network topology, directory services, hardware inventory, and software inventory of your organization, and the time required to perform automated deployments within that infrastructure. You need to weigh all these considerations carefully and consider the benefits and limitations of each installation method, even if one of the considerations seems to point you definitively toward one of the methods. You might also determine that one method of automated installation is best for one set of circumstances in your organization, but that other installation circumstances in your organization are better suited for another method.

The scope of your deployment plan might also have an impact on your choice of deployment tools and methods. If you are planning a very large client-side remote installation on thousands of computers in one central location, consider the impact on network availability. If you are planning a number of smaller deployments in geographically remote locations, consider the methods you will need to use to distribute and install the operating systems and reference images.

When you have completed the tasks in this chapter you will be able to choose the automated installation method best suited for your organization. You can then begin designing your automated installation according to the guidelines in the appropriate design chapter for the method you have chosen.

Considerations for choosing an automated installation method that are discussed in this chapter apply only to deployments and rollouts; they do not apply to ongoing operational tasks such as reinstallation after a hard disk failure or reinstallation caused by software or hardware failure.

# Process for Choosing an Automated Installation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To choose the best automated installation method for your situation, you need to systematically evaluate a variety of different aspects of your installation circumstances. These aspects include the logistics of the actual installation, the hardware and software involved in the installation, and the network and IT infrastructure of your organization. Use the hardware and software inventory and deployment plans of your organization as the source for this information. For more information about inventories, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

Figure 1.1 shows the process for choosing an automated installation method.

**Figure 1.1 Choosing Your Automated Installation Method**




The order of the tasks outlined in this chapter is designed to help you narrow your choices early in the process and fine-tune your decision toward the end of the process. Although your choice might appear clear early in the process, it is important that you carefully examine all of the factors that might affect your decision to be certain that you have not overlooked an important factor.

For a job aid to assist you in choosing an automated installation method, see "Choosing a Method for Automated Installation" (ACIOV_01.xls) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Choosing a Method for Automated Installation" on the Web at http://www.microsoft.com/reskit). Answering the questions listed in the job aid, as you read the information in this chapter, can help you determine the best automated installation method for your environment. This job aid is designed to be used online.

# Fundamentals of Automated Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can automate and customize installations by using answer files, scripts, and batch files that configure the operating system and applications automatically by using several installation tools provided with the Windows Server 2003 family.

**Basic Concepts of Automated and Customized Installations**

An *automated installation* runs with minimal or no user interaction. This provides a faster, more consistent, and trouble-free installation. The automated installation tools use two basic methods to accomplish an automated installation:

- An *image–based installation* is a method of copying, or cloning, a preconfigured operating system and software applications from a master computer onto destination clients and servers. For the purposes of this chapter, the term *image-based installation* refers to installations using Sysprep or the Remote Installation Preparation Wizard (Riprep.exe) installation tool.

- An *answer file–based installation* uses a text file that contains setup instructions. These instructions include:

  - Answers to the questions that Windows Setup normally presents during an installation.

  - Instructions for configuring operating system settings.

  - Instructions for installing applications without user intervention.

  For the purposes of this chapter, the term *answer file–based* installations refers to installations using the Unattend and Remote Installation Services Setup (Risetup.exe) installation tools.

A *custom installation* is an operating system installation that is modified to support specific hardware and software configurations and meets specific organizational and user needs. You can customize an automated installation by using the configuration and setting design information that you have determined for your Windows Server 2003 family and Windows XP deployment, including applications, additional language support, service packs, and device drivers. You customize an automated installation by:

- Modifying the answer file to provide the Setup program with specific answers and instructions.

- Adding custom files, applications, and programs to the distribution folder.

- Modifying the configuration of the master computer.

**Windows Server 2003 Automated Installation Tools**

Three automated installation tools are included with the Windows Server 2003 family. Each is described briefly in the following sections.

### Note

- You can start a destination computer by using a Windows Preinstallation Environment (Windows PE) CD, and then using the **diskpart** command to partition a disk and the **format** command to format a disk. WinPE is a bootable operating system that provides limited operating system functionality for performing preinstallation tasks. Windows PE is only available if you have purchased Enterprise Agreement 6.0, Enterprise Subscription Agreement 6.0, or Select License 6.0 with Software Assurance (SA). For more information about Windows PE and Windows PE licensing plans, see the Windows Preinstallation Environment link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Remote Installation Services**

With RIS, you can design a destination computer–initiated automated installation scenario to deploy clean, preconfigured, file system image–based or script-based installations on multiple client computers from remote master server computers over a

network connection. Using RIS, you can create and store reference images on a server; the destination computers initiate the installation process. RIS is the only method you can use to install an operating system without the need for an administrator to physically visit each computer to initiate the installation.

You can use two components of RIS to perform remote installations in different ways: Remote Installation Services Setup (Risetup.exe) and Remote Installation Preparation Wizard (Riprep.exe). The following summarizes the differences:

- **Risetup.exe.** You use this component to set up the RIS server and create a distribution folder for the operating system and software files for the installation.

- **Riprep.exe.** You can use Riprep.exe to create a customized image of an operating system such as Microsoft® Windows® XP Professional. Use Riprep to prepare an image from an existing operating system installation on a master computer and replicate that image to an available RIS server on your network. The image can include the operating system with default parameters applied, or the operating system with a preconfigured desktop, locally installed applications, and drivers.

RIS is a service that requires a dedicated server. You must also configure your network and domain to use RIS. For more information see "Designing RIS Installations" in this book.

**Sysprep**

You can use the System Preparation tool (Sysprep) to prepare a master computer for disk imaging after performing the initial setup steps on that computer. Sysprep assigns a unique security identifier (SID) to each destination computer the first time the computer is rebooted. Using Sysprep is the fastest way to set up a computer. Applying a Sysprep disk image to a destination computer takes just a few minutes.

With a third-party disk-imaging tool, you can copy the contents of the hard disk (a Sysprep disk image) of a master computer onto removable media. You can use this disk image to quickly install exact copies of the master computer onto the hard disk of destination computers in your organization.

Sysprep (Sysprep.exe) is in the Deploy.cab file in the \Support\Tools folder on the Windows XP Professional or Windows Server 2003 operating system CD. For more information see "Designing Image-based Installations with Sysprep" in this book.

**Unattended Installation**

Unattended installation (Unattended Setup) uses an answer file to automate the answers to the questions Windows Setup normally presents to the user during the installation. An answer file can also contain instructions for configuring operating system settings and installing applications without user intervention. You can, by means of a distribution share or media, distribute the answer file and any device drivers and other files that are required to customize the installation. Unattended installations take longer to perform because all of the files are copied to the destination computer.

Unattended installation consists of two command-line tools:

- **Winnt32.exe**, used when starting your installation from the Microsoft® Windows® 95, Windows® 98, Windows NT®, or Windows® 2000 operating systems.

- **Winnt.exe**, used when starting your installation from the Microsoft® Windows® 3.1, Microsoft® Windows® for Workgroups, and MS-DOS® operating systems.

These tools are in the \i386 folder on the Windows XP Professional or Windows Server 2003 operating system CD. For more information, see "Designing Unattended Installations" in this book.

# Choosing a Method Based on Clean Installations and Upgrades

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The decision to perform either clean operating system installations or upgrades is important when determining the best automated installation method to use. When moving to a new operating system, most organizations choose to perform clean installations. This helps them to maintain uniformity and, for client computers, to reset the corporate-installed base. If, however, you have older line-of-business applications or peripherals that you plan to continue using in your organization after moving to the new operating system, you might need to perform an upgrade to retain the ability to use those applications and device drivers.

If you are deploying clean installations of Windows XP Professional or Windows Server 2003, you can use any of the Windows Server 2003 automated installation methods. Before you perform a clean installation or an upgrade, test the installation of your older software and device drivers to be sure they work. For more information about performing upgrades, clean installations, and upgrade paths, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

Figure 1.2 illustrates the place of this step in the process of choosing an automated installation method.

**Figure 1.2 Choosing a Method Based on Clean Installations and Upgrades**




**Installation Tools for Upgrading the Operating System**

You cannot use RIS or Sysprep to upgrade an operating system. The only automated installation method you can use to perform upgrades is unattended installation using Winnt32.exe. You cannot use Winnt.exe to perform an upgrade.

Because some registry settings and system files are retained when you perform an upgrade, you need to thoroughly test your upgrade scenarios in your test lab before rolling out the installation to the production environment. Testing the upgrade can help avoid unexpected loss of data or configurations. For more information about designing an unattended installation, see "Designing Unattended Installations" in this book.

**Installation Tools for Performing Clean Installations**

You can use any of the Windows Server 2003 automated installation tools to perform clean installations of the operating system. If you are deploying clients and you want to retain user settings and data before using an automated installation method to deploy clean installations, consider using the User State Migration Tool (USMT). The guidelines outlined in the remainder of this chapter help you determine which of the Windows Server 2003 automated installation tools is best for a clean installation in your organization.

**Note**

- The USMT tool is included on the Windows Server 2003 operating system CD in the \ValueAdd\Msft\USMT folder.

# Choosing a Method Based on Software Considerations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Software considerations that affect your automated installation method include whether you plan to deploy server or client operating systems, what types of applications you plan to deploy along with the operating system, and the configuration of servers you plan to deploy. Figure 1.3 illustrates the place of this step in the process of choosing an automated installation method.

**Figure 1.3 Choosing a Method Based on Software Considerations**



# Choosing a Method Based on Operating System

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An important consideration when choosing your automated installation method is whether you are planning to deploy client or server operating systems. Answer file–based methods provide an opportunity to do the type of fine-tuned configuration necessary for deploying servers, and image-based methods are ideal for quickly deploying a common desktop environment to many client computers.

**Installing on Client Computers**
An image-based automated installation method is ideal for quickly deploying standard configurations to client computers. The device drivers for most Plug and Play devices for standard desktop and portable computers are included in Windows XP Professional, so no additional configuration is needed, even if these devices vary in your organization.

However, inventory your hardware to be certain that your device drivers are included in the installation. If you have a large number of desktop or portable computers that have a variety of specialized device drivers or drivers that are not included with Windows XP, an answer file–based method provides a way to reconfigure the installation with the appropriate drivers.

**Note**

- If you are using RIS to add a large number of client computers to an existing environment, include capacity planning in your deployment design to ensure an adequate level of service availability during the installation process.

**Installing on Servers**
Different considerations apply to installations on individual servers that have varying roles throughout the organization and installations on members of a server farm. Consequently, you might choose different installation methods for individual servers and server farms.

**Installing on individual servers**
Because installing a server operating system often involves customizing configurations, especially when deploying several different roles in a single installation, an answer file–based automated installation method is ideal.

**Installing on server farms**
Server farms, especially load-balanced server farms, often require identical configuration of the servers in the farm. Using Sysprep to set up the servers quickly with the same configuration can be the most efficient choice in this situation. If you are using Windows Network Load Balancing (NLB) as your load-balancing solution, however, you must script the installation and configuration of NLB after installing the operating system by using Sysprep.

# Choosing a Method Based on Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are planning to install applications together with the operating system, considerations that might affect your choice of automated installation methods include the compatibility of the application with the installation method.

**Testing Application Installation Compatibility for Image-Based Installations**

If you are planning an image-based installation and plan to install applications with the operating system, you need to thoroughly test the installation. Although most applications should copy correctly, some configurations, settings, or other aspects of the application might cause an application to behave unexpectedly after an image-based installation. If your applications do not install properly on an image, you can install by using Unattended Setup instead.

# Choosing a Method Based on Server Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are installing a server operating system, your plans for configuring the operating system affect the installation method you choose.

**Planning for Certificate Services**
You can use an answer file–based automated installation method to install and configure Certificate Services as part of the installation. However, when you use an image-based installation method, you must install and configure Certificate Services after the installation is complete.

**Planning for the Cluster Service**
You can use an answer file–based automated installation method to install and configure the Cluster service as part of the installation. However, when you use an image-based installation method, you must install and configure the Cluster service after the installation is complete.

**Planning for Domain Controllers**
As part of an answer file–based automated installation, you can include a script that starts the Active Directory Installation Wizard (Dcpromo.exe) to configure a server as a domain controller. However, when you use an image-based installation method, you must configure servers as domain controllers after the installation is complete.

**Planning for Internet Information Services**
The configuration settings for Internet Information Services (IIS) are built into the answer files used with Unattend and Risetup; therefore these methods are the most efficient way to deploy IIS servers. However, image-based installation of an IIS server is fully supported.

# Choosing a Method Based on Network and Hardware Configurations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are several network topology and hardware considerations that affect your automated installation. Figure 1.4 illustrates the place of this step in the process of choosing an automated installation method.

**Figure 1.4 Choosing a Method Based on Network and Hardware Configurations**



# Choosing a Method Based on Network Topology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Network bandwidth and existing network protocols are important factors in deciding which automated installation method to use. For example, if you do not have a high-bandwidth connection to a network server, a method that uses media such as a CD-ROM or DVD for the installation is usually more appropriate than using RIS for automated installations.

**Examining Network Connectivity**

To use RIS or to perform an installation by using Sysprep or Unattend from a share, you need to have reliable, high-bandwidth network connections in place. RIS requires that a TCP/IP network be in place. This is not a requirement for Unattend and Sysprep.

If the destination computers for your automated installation are connected to the network by low-bandwidth connections, such as when you have clients located in remote locations, an automated installation method that you install from a disk, such as Sysprep or Unattend, is better than a RIS-based installation. RIS requires that a robust network be in place. Installations performed with Sysprep or Unattend take place locally on the computer. Nothing needs to travel across the network.

### Note

- To use RIS for automated installation, you need a network card that supports Pre-Boot eXecution Environment (PXE) technology. Wireless network cards and many token ring network cards do not support PXE. For more information, see "Choosing a Method Based on Hardware Inventory" later in this chapter.

**Examining IP Address Allocation**

After an image created with Sysprep or Riprep is copied onto a destination computer, you must configure static IP address settings. When a disk image is copied onto a destination computer, all of the network adapters on the destination computer are initialized to the default settings, which include dynamic allocation of IP addresses. For this reason, an answer file–based installation method might be more convenient to use when you need to configure static IP addresses. For more information about how Sysprep affects network settings, see article Q271369, "Statically-Entered TCP/IP Settings Are Not Present After Sysprep" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

RIS requires the presence of an active Dynamic Host Configuration Protocol (DHCP) server on the same network as the client computers. The remote boot–enabled client computers receive an IP address from the DHCP server prior to contacting the RIS server. Routers in a multi-subnet network are configured to forward the DHCP packets between the client and the RIS server. In addition, on routed networks, DHCP packets must be relayed to the RIS server as well as the DHCP server. For more information, see "Designing RIS Installations" in this book.

# Choosing a Method Based on Hardware Inventory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When choosing your automated installation method, a number of considerations about the hardware of both destination and master computers might affect your choice of methods. These include both compatibility and configuration considerations. In general, if you are deploying to a homogeneous hardware base, an image-based automated deployment method is optimal. However, if you are deploying to a heterogeneous hardware base, such as older hardware with varying drivers that are not included with the operating system that you are deploying, an answer file–based automated deployment method is optimal.

**Examining HAL Compatibility**

You can only perform an image-based installation (using Sysprep or Riprep) if the hardware abstraction layer (HAL) on the disk image is compatible with the hardware on the destination computer. For example, if the master computer on which you run Sysprep or Riprep has an Advanced Configuration and Power Interface (ACPI) HAL, then the destination computers you designate to receive operating system images generated from that master computer must also have ACPI HALs. In some cases, you can upgrade the HAL that is on a disk image to suit the HAL requirements of a destination computer, but you must be certain that the HAL is compatible for this type of upgrade. If the HALs of the master computer and the destination computers are not compatible, an answer file–based installation method might be more convenient.

**Evaluating Support for PXE**

To initiate a RIS-based operating system installation, a RIS client must first perform a remote network boot by connecting to a RIS server over the network. To make a remote boot possible, both the network adapter and ROM BIOS of the destination computer need to support PXE.

It is possible to emulate PXE support by using a Peripheral Component Interconnect (PCI)–based network adapter that boots from a RIS boot floppy. The RIS boot floppy is a startup disk that simulates the PXE startup process for computers that lack a remote boot–enabled BIOS. The Remote Boot Floppy Generator tool (Rbfg.exe), which is a part of RIS and is located on the RIS server, allows you to generate RIS boot floppy disks for use with RIS clients that are not PXE-enabled. Rbfg.exe supports a limited number of PCI cards. It is not possible to add support for a card that is not supported by the disk. You can obtain a list of supported network cards by clicking the **Adapter List** button in the rbfg.exe application.

If the network adapter and ROM BIOS of the destination computer do not support PXE, either Unattended Setup or Sysprep might be a more convenient choice for automated installation.

**Evaluating Mass Storage Controllers**

Examine the mass storage controllers in your organization. If you have mass storage controllers that are not listed in any device information (.inf) file, such as Machine.inf, Scsi.inf, Pnpscsi.inf, or Mshdc.inf, you need to specially configure the Mini-Setup stage of an image-based installation. In this case, the overhead involved might indicate that using Unattended Setup or Risetup is a better choice.

# Choosing a Method Based on Directory Services Considerations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When choosing your automated installation method, you need to consider the directory service your organization has in place. Figure 1.5 illustrates the place of this step in the process of choosing an automated installation method.

**Figure 1.5 Choosing a Method Based on Directory Services Considerations**




**Examining Plans for Directory Services**

If you plan to use RIS for automated installations, you must be using the Active Directory® directory service. RIS relies on Active Directory for security and computer account placement. In addition, RIS uses Active Directory to identify RIS clients and RIS servers.

**Identifying Applications That Are Dependent on Active Directory**

Identify any applications that you plan to include with the automated installation that are dependent on Active Directory, such as client applications that access human resources or proprietary data. These applications cannot be installed and configured on a Sysprep image. You must install and configure these applications after the disk image is copied onto the destination computer and the computer is restarted. In this case, an Unattended Setup installation might be a better choice because Active Directory–dependent applications can be included with the rest of the installation.

**Examining Plans for Domain Controllers**

Special considerations apply if you intend to create domain controllers by using an automated installation method. You cannot configure a Sysprep master computer as a domain controller. You need to first configure a master computer as a stand-alone server, and then install Active Directory by using the Active Directory Installation Wizard (Dcpromo.exe) after the disk image is copied onto a destination computer. However, you can script Dcpromo.exe with an answer file, and you can use the GuiRunOnce entry in your answer file to automatically start it at the end of an Unattended Setup. This is, therefore, a more efficient choice for installing preconfigured domain controllers.

# Additional Resources for Choosing an Automated Installation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing Unattended Installations" in this book.
- "Designing RIS Installations" in this book.
- "Designing Image-based Installations with Sysprep" in this book.
- The Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about hardware and software that is compatible with the Windows Server 2003 family and Windows XP.

**Related Tools**

- Deploy.cab

  The Deploy.cab file contains the Sysprep tool and the Setup Manager tool, along with other related automated installation tools. The Deploy.cab file also contains the Microsoft Windows Preinstallation Reference file (Ref.chm), which has information about answer file sections, keys, and values; and the *Microsoft Windows Corporate Deployment Tools User's Guide*, which has detailed information about performing an automated installation. The Deploy.cab file is in the \Support\Tools folder on the Windows XP Professional and Windows Server 2003 operating system CDs. You can use Windows Explorer or run Extract.exe to extract and view the Ref.chm file.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Planning for unattended Setup" in Help and Support Center for Windows Server 2003.
- "Remote Installation Services" in Help and Support Center for more information about installing and managing RIS for Windows Server 2003.

**Related Job Aids**

"Choosing an Automated Installation Method" (ACIOV_01.xls) (or see "Choosing an Automated Installation Method" on the Web at http://www.microsoft.com/reskit).

# Designing Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Unattended installation is the most flexible and versatile automated method of deploying the Microsoft® Windows® XP Professional operating system and the Microsoft® Windows® Server 2003 operating system. You can automate both clean installations and upgrades, and you can automate the installation and configuration of a wide range of system settings and system components without running batch files or scripts. Unattended installations also speed up the deployment process, minimize user involvement during installation, and ensure consistency throughout your organization, which lowers support costs.

**In This Chapter**


Overview of Unattended Installation

Evaluating Hardware and Software for Unattended Installations

Deciding Whether to Perform an Upgrade or a Clean Installation

Choosing a Distribution Method

Designing Preinstallation Tasks for Unattended Installations

Designing Answer File and Setup Settings for Unattended Installations

Creating Startup Media, Answer Files, and Distribution Shares

Performing Unattended Installations

Additional Resources for Designing Unattended Installations


**Related Information**

- For more information about automated installations, see "Choosing an Automated Installation Method" in this book.
- For more information about image-based installations, see "Designing Image-based Installations with Sysprep" in this book.
- For more information about Remote Installation Services (RIS), see "Designing RIS Installations" in this book.

# Overview of Unattended Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Unattended installations are commonly used to perform bulk installations with minimal user intervention. Unattended installations are particularly useful if you are:

- Upgrading a Windows operating system to Windows XP Professional or Microsoft® Windows Server 2003, Standard Edition; Microsoft® Windows Server 2003, Web Edition; or Microsoft® Windows Server 2003, Enterprise Edition operating systems.
- Performing automated installations on computers that have heterogeneous hardware configurations.
- Performing automated installations on specific types of servers, such as domain controllers, remote access servers, and servers that run Certificate Services or the Cluster service.
- Configuring a wide range of operating system settings during an automated installation without using batch files and scripts.

In addition to these deployment solutions, unattended installation is a useful method of creating master installations for image-based and RIS installations.

As a deployment solution, unattended installation requires substantial up-front planning and design. This chapter is designed to help IT professionals in medium and large organizations plan and design an unattended installation. It is assumed that you have already read "Choosing an Automated Installation Method" in this book. It is also assumed that you have designed the client and server configurations that you want to deploy in your organization. This includes designing the configuration of all networking, directory services, and security components. You will use this client and server design information throughout this chapter to customize your unattended installation. When you finish the planning and design work described in this chapter, you will be ready to perform an unattended installation.

**Note**

- You cannot perform an unattended installation of Windows Server 2003, Datacenter Edition.

# Unattended Installation Design Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Planning and designing an unattended installation involves a design team and a deployment team. The design team is responsible for assessing your current environment, deciding whether to upgrade or to perform clean installations, and designing the overall deployment process, including the distribution mechanism, preinstallation tasks, and automated installation and post-installation tasks. The deployment team is responsible for implementing all design decisions, including creating answer files, distribution shares, and startup media. The steps in this decision-making process are shown in Figure 2.1.

**Figure 2.1 Designing Unattended Installations**



# Unattended Installation Fundamentals

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To perform an unattended installation, you first create an *answer file* — a text file that contains answers to the questions that Windows Setup normally prompts you for during an installation. An answer file also can contain instructions for configuring operating system settings and installing applications without user intervention. After you configure your answer file, you typically create a *distribution share* — a folder that contains the Windows XP Professional or Windows Server 2003 installation files, as well as any device drivers or other files that are required to customize the installation. A distribution share uses a hierarchical folder structure that is similar to the one used on computers running Windows Server 2003 or Windows XP Professional, and is typically stored on a server to which your destination computers can connect and retrieve copies of the files during an unattended installation.

**Note**

- You do not need to use a distribution share to perform an unattended installation; you can use an operating system CD instead of a distribution share.

After you have created an answer file and a distribution share, you are ready to start an unattended installation on a destination computer. To do this, you must run one of two Windows Setup programs: Winnt.exe or Winnt32.exe. Winnt.exe runs on 16-bit operating systems, including Microsoft® MS-DOS®, Microsoft® Windows 3.1, and Microsoft® Windows® for Workgroups operating systems. Winnt32.exe is used on 32-bit operating systems, including Microsoft® Windows® 95, Microsoft® Windows® 98, Microsoft® Windows® Millennium Edition, Microsoft® Windows NT®, Microsoft® Windows® 2000, Windows XP Professional, and Windows Server 2003 operating systems. Usually, you start the destination computer with a floppy disk that has been formatted as an MS-DOS startup disk, an operating system CD, or the existing operating system that is on the computer's hard disk. Depending on the operating system that is running on the destination computer, you then run either Winnt.exe or Winnt32.exe.

When you run Winnt.exe or Winnt32.exe to perform an unattended installation, you specify various command-line parameters. For example, you specify the name of the answer file you want Setup to use, and the location of the distribution share that contains the installation files. You also can specify various options, including whether to use Dynamic Update, whether to install Emergency Management Services, or whether to install the Recovery Console. Setup then runs and carries out all of the instructions specified in the answer file.

#### New in Windows Server 2003 and Windows XP Professional

Several answer file headings and entries are new for Windows Server 2003 and Windows XP Professional. In addition, some headings and entries have been modified, and some headings and entries found in earlier versions of Windows no longer apply to Windows Server 2003 and Windows XP Professional. For more information about answer file changes, see "Changes in Answer Files" in *Microsoft® Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

#### Unattended Installation Terms and Definitions

The following key terms are associated with unattended installations.

**Setup Manager** A program that creates answer files and distribution shares for unattended installations. Setup Manager (Setupmgr.exe) is included in the Deploy.cab file in the Support folder on the Windows XP Professional and Windows Server 2003 operating system CDs.

**Unattend.txt** The default name of the answer file that you use to automate Windows Setup during an unattended installation. Unattend.txt contains headings and parameters that instruct Setup to perform various configuration tasks.

**Winnt.sif** The name you give Unattend.txt when you perform an unattended installation by using the operating system CD instead of a distribution share.

**Cmdlines.txt** A configurable text file that you use to customize an unattended installation. Cmdlines.txt contains a list of commands that run synchronously after Setup finishes, but before a computer restarts. Cmdlines.txt can exist on the destination computer's hard disk or on a floppy disk, and must be specified in the [Unattended] section of Unattend.txt or Winnt.sif.

**[GUIRunOnce]** A section in your answer file that is used to customize an unattended installation. The [GUIRunOnce] section contains a list of commands that run synchronously after a destination computer is started for the first time and a user logs on.

**Mini-Setup** A wizard that is a subset of Windows Setup. Mini-Setup provides prompts for user-specific information, configures operating system settings, and detects new hardware. You can automate Mini-Setup by using a Sysprep.inf answer file.

**File-copy mode** The first of the three stages of Setup, where the Windows program files and any additional files specified are copied to the computer's hard disk.

**Text Mode** The second of the three stages of Setup, during which Setup determines the basic hardware of the computer (CPU, motherboard, hard disk controllers, file systems, and memory), installs the base operating system necessary to continue, and creates specified folders.

**GUI mode** The third of the three stages of Setup, during which Setup configures the computer's hardware and network settings, prompts you to provide an Administrator password, and allows you to customize the installation.

# Evaluating Hardware and Software for Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you start designing an unattended installation, you must determine whether the hardware and software on your destination computers is compatible with the operating system you are deploying, and whether you have the appropriate device drivers for an unattended installation. You can use your hardware and software inventory to identify the hardware and software that is installed on your destination computers. For more information about creating a hardware or software inventory, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

Figure 2.2 shows the steps you need to follow to evaluate hardware and software for unattended installations.

**Figure 2.2 Evaluating Hardware and Software**



# Evaluating Hardware and Software Compatibility

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must evaluate hardware and software compatibility before you perform an unattended installation because incompatibilities can cause an unattended installation to fail. For example, if Setup attempts to install a device or an application that is incompatible during an unattended installation, Setup cannot display a warning dialog box or prompt a technician for alternative installation instructions. Instead, Setup simply stops running, and the installation fails. Therefore, you must replace or upgrade incompatible hardware and software before you perform an unattended installation.

**Evaluating Hardware Compatibility**

You can use several resources and tools to verify hardware compatibility.

**Windows Catalog**

The Windows Catalog contains a list of software and hardware products that are designed for, or are compatible with, Windows XP Professional and Windows Server 2003. You can search the catalog by manufacturer, product type, product name, or model. If you do not see a product in the Windows Catalog, it does not mean the product will not work with Windows XP or Windows Server 2003 — check with the product's manufacturer to determine whether the product works with these versions of Windows. For more information, see the Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Windows Upgrade Advisor and Windows XP Upgrade Advisor**

Windows Upgrade Advisor and Windows XP Upgrade Advisor are tools that check your system hardware and software to see whether they are ready to be upgraded to Windows XP Professional or Windows Server 2003. Although the tools assume you are upgrading to Windows XP Professional or Windows Server 2003, you also can use them to identify software and hardware that are not compatible during a clean installation of Windows XP Professional or Windows Server 2003. To download either of the tools, see the Windows Upgrade Advisor link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. You also can run either tool by using the **/checkupgradeonly** parameter with Winnt32.exe. Winnt32.exe is included in the i386 folder on any Windows XP Professional or Windows Server 2003 operating system CD.

**Evaluating Application Compatibility**

In addition to Windows Catalog and Windows Upgrade Advisor, you can use the Application Compatibility Toolkit to verify software compatibility. The Application Compatibility Toolkit contains documents and tools to help you diagnose and resolve application compatibility issues with Windows XP Professional and Windows Server 2003. For more information about application compatibility, see "Planning and Testing for Application Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit. For more information about the Application Compatibility Toolkit, see article Q294895, "Description of the Application Compatibility Toolkit 2.0 for Windows XP," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Identifying Supplemental Device Drivers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most device drivers are provided on the Windows XP Professional and Windows Server 2003 operating system CDs. However, some devices require special device drivers, which you must obtain from the hardware vendor. Some devices that require special drivers include:

- Plug and Play devices for which there are no device drivers on the operating system CD. Sound cards and video adapters on portable computers are common examples.
- Mass storage controllers, particularly small computer system interface (SCSI) controllers and redundant array of independent disks (RAID) controllers.
- Hardware abstraction layers (HALs) created by a third party.
- Legacy (non–Plug and Play) devices for which there are no device drivers on the operating system CD.

When you obtain device drivers from a third party, make sure that Microsoft has digitally signed the device drivers. Signing ensures that the device driver has undergone a certain level of testing, and that the device driver cannot be altered or overwritten by another program's installation process. Device driver files can include: system (.sys) files, dynamic-link library (.dll) files, information (.inf) files, and text setup (Txtsetup.oem and Txtsetup.sif) files.

In addition to making sure that you have the appropriate device drivers for your hardware, you need to make sure that the device drivers are available to Setup during the unattended installation process. If they are not, Setup might stop running and the unattended installation might fail. To make sure device drivers are available to Setup, you need to save the device drivers in a distribution share so they are copied to the destination computer during an unattended installation, or you need to save the device drivers on the destination computer's hard disk before you initiate the installation.

For a worksheet to help you record information about supplemental device drivers, see "Unattended Installation Worksheet" (ACIUI_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Unattended Installation Worksheet" on the Web at http://www.microsoft.com/reskit). For more information about device-driver signing and how Windows selects device drivers, see "Using Signed Drivers" in *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Deciding Whether to Perform an Upgrade or a Clean Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Unlike other automated installation methods, such as image-based and RIS installations, an unattended installation can be configured to perform a clean installation or an upgrade of an existing operating system. To determine which installation method is more appropriate, you need to evaluate possible upgrade paths, and then evaluate the advantages and disadvantages of upgrades and clean installations. If you determine that upgrading is appropriate, you then identify upgrade issues that you need to address during the planning and design phase of an unattended installation. Figure 2.3 shows the steps you need to follow to determine whether to perform an upgrade or a clean installation.

**Figure 2.3 Determining Whether to Perform an Upgrade or a Clean Installation**



# Evaluating Possible Upgrade Paths

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If the computer on which you want to install Windows Server 2003 or Windows XP Professional is running Windows 95, Windows 3.1, or Microsoft® Windows NT® Server version 3.51, you need to perform a clean installation. In addition, you need to choose clean installation if you are using Windows 98, Windows Millennium Edition, Microsoft® Windows NT ®Workstation version 4.0, Microsoft® Windows® 2000 Professional, or Windows XP Professional operating systems and you want to install Windows Server 2003. In most other cases, you can choose to upgrade. Table 2.1 shows the possible upgrade paths for Windows XP Professional and Windows Server 2003.

**Table 2.1 Upgrade Paths for Windows XP Professional and Windows Server 2003**

| Existing Operating System | Upgrade to Windows XP Professional | Upgrade to Windows Server 2003, Standard Edition | Upgrade to Windows Server 2003, Enterprise Edition | Upgrade to Windows Server 2003, Web Edition |
|---|---|---|---|---|
| Windows NT Server 3.51 | | | | |
| Microsoft® Windows NT ®Server version 4.0 (with Service Pack 5) | | ● | ● | |
| Microsoft® Windows NT® version 4.0, Terminal Server Edition | | ● | ● | |
| Microsoft® Windows NT® version 4.0, Enterprise Edition | | | ● | |
| Microsoft® Windows® 2000 Server | | ● | ● | |
| Microsoft® Windows® 2000 Advanced Server | | | ● | |
| Windows Server 2003, Standard Edition | | | ● | |
| Windows 3.1 | | | | |
| Windows 95 | | | | |
| Windows 98 | ● | | | |
| Windows Millennium Edition | ● | | | |
| Windows NT Workstation 4.0 | ● | | | |
| Windows 2000 Professional | ● | | | |
| Microsoft® Windows XP Home Edition | ● | | | |

There are three exceptions to the options listed in Table 2.1:

- You cannot upgrade from one localized version of a Windows operating system to another localized version of a Windows operating system. For example, you cannot upgrade from the Japanese version of Windows 2000 Server to the English version of Windows Server 2003, Standard Edition.

- You can upgrade a computer that has the Multilingual User Interface Pack (MUI Pack) only to the English version of the operating system.

- You can upgrade from Windows NT 4.0 only if you have Service Pack 5 or later installed. It is not possible to upgrade a computer that has Windows NT 4.0 with Service Pack 4 or earlier.

# Evaluating Differences Between an Upgrade and a Clean Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

During an upgrade, existing user settings are retained, as are installed applications and application settings. During a clean installation, the operating system files are installed in a new folder, and you must reinstall all of your applications and reconfigure user preferences, such as desktop and application settings.

The biggest benefit of an upgrade is that you can accommodate heterogeneous hardware and software configurations without having to customize individual computers in your organization. Also, because an upgrade does not affect applications, files, or settings, you do not have to spend time configuring computers or installing applications during a rollout, which speeds up the deployment process. Another benefit is that you do not need to migrate user data before an upgrade.

The biggest benefit of a clean installation is that all of your systems can be deployed with the same configuration. If you use the same answer file for all your systems during a clean installation, all applications, files, and settings are reset the same way, which means all of the desktops or servers in your organization can be standardized. In this way, you can avoid many of the support problems that are caused by irregular or inconsistent configurations.

You can use the following guidelines to determine whether to perform an upgrade or a clean installation.

Choose a clean installation if:

- No operating system is installed on the destination computer.
- The installed operating system cannot be upgraded to Windows XP Professional or Windows Server 2003.
- The computer has a multiple-boot configuration that needs to support the current operating system and either Windows XP Professional or Windows Server 2003.
- You are planning on implementing a managed environment through Group Policy, the Active Directory® directory service, or other means, but you have not yet implemented a managed environment. In this case, clean installations are desirable because they ensure that you have a standard configuration on which to implement your managed environment.
- You want to reset the desktop or server configuration in your organization to a consistent, known standard.
- You are purchasing new hardware or software as part of your deployment.

Choose to upgrade if:

- You already have a Windows operating system that is suitable for upgrading, and your IT department centrally manages the computers in your organization.
- You want to use existing hardware and software, and you do not want to reconfigure user settings, operating system settings, or application settings.
- You need to retain hardware or software settings for compatibility reasons.

You cannot use a distribution share to perform an unattended upgrade installation. You must use the product CD. Also, the Windows Setup program reads only a limited number of answer file sections and entries during an unattended upgrade installation. For more information, see "Performing an Unattended Upgrade Installation" later in this chapter.

**Important**

- Installing multiple operating systems on the same partition is not supported, and doing so can prevent one or both operating systems from working properly.

For a worksheet to help you record information about your installation, see "Unattended Installation Worksheet" (ACIUI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Unattended Installation Worksheet" on the Web at http://www.microsoft.com/reskit).

# Choosing a Distribution Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

*Distribution* refers to the way you store and transfer system files, device drivers, and software program files during an unattended installation. There are two types of distribution methods:

- You can store files in a distribution share on a network share, and then transfer the files across the network to each of the computers on which you are performing an unattended installation. This method works best when your destination computers are connected to a server by a reliable, high-bandwidth network.

- You can store files on media — such as a CD, a DVD, or a floppy disk — and then transfer the files locally from the media to each of the computers on which you are performing an unattended installation. Typically, when you use this method, you use the operating system CD in conjunction with a floppy disk. The operating system CD contains the system files and device drivers; the floppy disk contains the answer file that is necessary to perform an unattended installation.

Figure 2.4 shows the steps you need to follow to choose a distribution method for unattended installations.

**Figure 2.4 Choosing a Distribution Method**



# Evaluating Distribution Methods

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Usually, unattended installations are performed with a distribution share. However, there are instances where a distribution share is not appropriate and you need to use media, rather than a distribution share, to perform an unattended installation. The guidelines in this section help you evaluate which distribution method is more appropriate for your organization's needs.

# Using a Distribution Share to Perform an Unattended Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A *distribution share* contains the Windows XP Professional or Windows Server 2003 installation files, as well as any device drivers or other files that are required to customize the installation. The distribution share is structured hierarchically and resembles the folder structure on a computer running Windows Server 2003 or Windows XP Professional. The folder typically resides on a server to which your destination computers can connect. You typically create one distribution share for each operating system you are deploying.

By using a distribution share to perform an unattended installation, you can perform the following tasks:

**Copy folders and files to destination computers** A distribution share can contain special folders and files that are copied to a computer's hard disk during an unattended installation. For example, if you want to create a Scripts folder on drive C of your destination computers, you can add the Scripts folder to your distribution share and it will automatically be copied to your destination computers during installation.

**Copy device drivers to destination computers** A distribution share contains folders for storing mass storage device drivers, HALs, and Plug and Play device drivers. The contents of these folders are copied to a computer's hard disk during an unattended installation.

**Copy a Sysprep folder to destination computers** A distribution share can contain a Sysprep folder, which can be used to store Sysprep files and the Sysprep.inf file. This is useful if you are performing image-based installations and you want to use unattended installation to create the master installations. In this case, you can add a Sysprep folder to the distribution share, and the Sysprep folder will automatically be created on the master installation when you perform an unattended installation.

Use a distribution share to perform your unattended installation if you want to:

- **Increase consistency.** When you use a distribution share, you create a single source for system files, device drivers, and customized installation files that are copied to each destination computer, which ensures that each unattended installation is consistent throughout your organization.

- **Reduce administrative costs.** All system files, device drivers, and customized installation files are stored in a central location, which reduces the cost of updating and changing system files and device drivers. For example, if you need to upgrade an existing device driver or add new device drivers, you only have to make the change in the distribution share. You do not need to create new answer files or create new floppy disks or CDs.

- **Control access to your installation files.** You can secure a distribution share by using file and folder permissions, which lets you specify the users and groups who can gain access to your installation files. By default, in Windows Server 2003, the Everyone group is granted read-only permissions when a folder is shared. Therefore, when you create a distribution share you need to grant write permissions to all of the users who are responsible for configuring the distribution share.

In addition, distribution shares are useful if you are performing other types of automated installations, such as image-based or RIS installations. In these cases, you can use unattended installation with distribution shares to create consistent master installations that are easily modified and updated.

Despite the advantages, there are several disadvantages that might preclude you from using a distribution share. If any of the following conditions are true, you probably do not want to use a distribution share to perform unattended installations:

- **Slow network connections.** Accessing a distribution share across a slow network connection, such as a dial-up connection or a slow wide area network (WAN) connection is overly time-consuming and impractical. It is recommended that you do not use a distribution share if your destination computers use a slow network connection to reach a distribution share.

- **Minimal file server capacity.** Distribution shares are typically stored on file servers. If your file servers do not have sufficient capacity to store all of your distribution shares, or if your file servers cannot handle an increase in file throughput during your rollout, then you cannot use a distribution share. The minimum disk space required for a

distribution share for Windows Server 2003, Standard Edition is 650 megabytes (MB).

- **Minimal number of unattended installations.** You need to create, test, and maintain distribution shares. If you are performing only a few unattended installations, it is probably not cost-effective or efficient to use distribution shares. However, if you are performing image-based or RIS installations, and you are using unattended installations to create only a few master installations, distribution shares are useful and cost-effective because a distribution share ensures consistency among master installations and makes it easy to make configuration changes.

If you choose to use distribution shares, you can save the distribution shares on multiple servers. This allows Setup to copy files simultaneously from several servers, thereby speeding up the file copy stage of Setup. In addition, you can use Distributed File System (DFS) to increase the availability of your distribution shares. For more information about copying installation files from multiple servers, see "Choosing Winnt32.exe Parameters" later in this chapter. For more information about DFS, see "Designing and Deploying File Servers" in *Planning Server Deployments* of this kit.

# Using Media to Perform an Unattended Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Media distribution is useful if you are deploying computers in remote locations that do not have high-speed network connections, or in locations that do not have a local IT department available to set up and perform the unattended installation. Media distribution is also useful as an alternative installation method when network congestion limits your ability to access a distribution share, or when a destination computer has a malfunctioning network adapter.

To use media distribution instead of a distribution share, you use the operating system CD to start the destination computer. Then, you insert a floppy disk containing an answer file into the floppy disk drive. Setup reads the answer file, copies the appropriate installation files from the operating system CD to the destination computer's hard disk, and then configures the destination computer based on the parameters specified in the answer file.

Using an operating system CD to perform an unattended installation has the following requirements:

- The destination computer must have a bootable CD-ROM drive, and the boot-order sequence in the destination computer's BIOS must list the CD-ROM drive before the hard disk.
- The destination computer must have a floppy disk drive.
- The answer file must be named Winnt.sif, and it must be saved on a floppy disk.
- The installation files must all be present on the operating system CD; you cannot access supplemental device drivers or files that are not on the CD.
- The answer file must have a [Data] section, and the entries within the [Data] section must be configured appropriately for an unattended installation that uses an operating system CD.

In addition, using an operating system CD to perform an unattended installation has the following limitations:

- You cannot perform an upgrade, only a clean installation. To perform an upgrade, you must have an operating system running on the destination computer, which you do not have when you start a computer with the operating system CD.
- You cannot use a uniqueness database file (.udf) to modify the answer file. A .udf file overrides answer file settings, and is typically used to set the computer name during an unattended installation. For example, you might have a .udf file that contains a list of predetermined computer names for your organization. Instead of creating numerous answer files, each containing a unique computer name, you can configure the answer file so it obtains a computer name from the .udf file.
- You cannot use Dynamic Update to add updated installation files and device driver files during setup.
- You cannot stage the installation of applications. For more information about staging the installation of applications, see "Preinstalling Applications" in *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

Despite the requirements and limitations, there are some advantages to performing an unattended installation with the operating system CD. Namely, you can configure the answer file to do the following:

- **Configure disks.** By using the Repartition parameter in the [Unattended] section, you can delete all partitions on the destination computer and create a new partition that is formatted with the NTFS file system.
- **Skip Mini-Setup.** By using the UnattendSwitch parameter in the [Unattended] section, you can prevent the Mini-Setup program from running when the destination computer is started for the first time after unattended installation is completed.
- **Force BIOS startup.** By using the UseBIOSToBoot parameter in the [Data] section, you can force a destination computer

to use the BIOS to start, even though Setup detects that it is more appropriate to use a device miniport driver to start the computer. The current generation of hardware uses the BIOS to start the computer, so this entry is rarely required. However, on computers with large drives that support extended int13 BIOS calls, this might not be the default behavior. Using the BIOS starts computers faster by eliminating possible delays caused by a miniport driver. Do not use this entry unless you are sure that the BIOS supports the extended int13 functions.

For more information about configuring an answer file for unattended installations, see "Designing Automated Installation Tasks" later in this chapter, and *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Designing the Distribution Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you choose a distribution method, you need to design the distribution process. If you use a distribution share to perform unattended installations, this includes designing the distribution share and determining which files and folders in the distribution share need to be renamed. If you use an operating system CD to perform unattended installations, this includes designing the startup process.

# Designing a Distribution Share

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To design and configure a distribution share, you need to:

- **Structure the distribution share.** This includes identifying and structuring all of the supplemental files and folders that you want copied to the hard disk of each destination computer, such as device drivers, applications, and scripts.

- **Identify files that need to be renamed.** This includes files in the distribution share that need to be renamed after they are copied to a destination computer.

Three job aids are available to assist you with the design and configuration steps discussed in this section:

- For a worksheet to help you record information about the structure of your distribution share, see "Distribution Share Worksheet" (ACIUI_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Distribution Share Worksheet" on the Web at http://www.microsoft.com/reskit).

- For a worksheet to help you record information about the files and folders in your distribution share that you need to rename, see "Renamed Files and Folders Worksheet" (ACIUI_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Renamed Files and Folders Worksheet" on the Web at http://www.microsoft.com/reskit).

- For a worksheet to help you record general information about your unattended installation, including the location and name of your distribution share, see "Unattended Installation Worksheet" (ACIUI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Unattended Installation Worksheet" on the Web at http://www.microsoft.com/reskit).

**Structuring a Distribution Share**

A distribution share consists of a single, top-level folder, named i386, and several subfolders. Figure 2.5 shows the structure of a distribution share.

**Figure 2.5 Distribution Share Structure**



**i386**
This is the distribution share. You create it at the root of the distribution server (the server on which the distribution share is located). You can use any name for this folder, but the name must not contain more than eight characters. You can create this folder by copying the contents of the i386 folder on a Windows Server 2003 or Windows XP Professional operating system CD.

**$OEM$**
You create the $OEM$ subfolder directly beneath the i386 folder. During setup, you can automatically copy folders, standard 8.3 format files, and any tools needed for your automated installation process to the $OEM$ subfolder.

If you want to create the $OEM$ subfolder outside the distribution share, you can use the OEMFILESPATH parameter in the answer file. For more information about answer file parameters and syntax, see *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

The $OEM$ subfolder can contain the optional file Cmdlines.txt, which contains a list of commands to be run at the end of the graphical user interface (GUI)-mode stage of Setup. These commands can be used to install additional services or applications that you want to include with your installation. For more information about the Cmdlines.txt file, see "Designing Automated Post-Installation Tasks" later in this chapter.

As long as Setup finds the $OEM$ subfolder in the root of the distribution share, Setup copies all of the files found in this folder to the temporary directory on the destination computer that is created during the text mode stage of Setup.

**$OEM$\Textmode**
The $OEM$\Textmode subfolder contains new or updated files for installing mass storage device drivers and HALs. These files can include OEM HALs, drivers for SCSI devices, and Txtsetup.oem, which directs the loading and installing of these components. Be sure to include the Txtsetup.oem file. All files placed in the $OEM$\Textmode subfolder (HALs, drivers, and Txtsetup.oem) must be listed in the [OEMBootFiles] section of the answer file.

**$OEM$\$$**
The $OEM$\$$ subfolder is equivalent to the systemroot or windir environment variables. The subfolder contains additional files that you want copied to the various subfolders of the Windows Server 2003 or Windows XP Professional installation directory. The structure of this subfolder must match the structure of a standard Windows Server 2003 or Windows XP Professional installation, where $OEM$\$$ matches systemroot or windir (for example, C:\Windows), $OEM$\$$\Fonts matches *windir*\Fonts, and so on. Each subfolder must contain the files that need to be copied to the corresponding system folder on the destination computer.

**$OEM$\$$\Help**
A subfolder that contains the OEM Help files to be copied to the *systemroot*\Help folder during setup.

**$OEM$\$$\System32**
A subfolder that contains files to be copied to the *systemroot*\System32 folder during setup.

**$OEM$\$1**
This folder points to the drive on which Windows Server 2003 or Windows XP Professional is installed; $1 is equivalent to the systemdrive environment variable. For example, if you are installing Windows Server 2003 on drive D, $OEM$\$1 points to drive D.

**$OEM$\$1\Pnpdrvrs**
You can use the $OEM$\$1*Pnpdrvrs* subfolder to place new or updated Plug and Play device drivers in your distribution shares. These folders are copied to the *systemdrive**Pnpdrvrs* location on the destination computer. Adding the OemPnPDriversPath parameter to your answer file directs Windows Server 2003 or Windows XP Professional to look (both during and after setup) for new or updated Plug and Play drivers in the folders that you created, in addition to those originally included with the system. Note that you can replace *Pnpdrvrs* with a name of your own that is eight or fewer characters long.

**$OEM$\$1\Sysprep**
The $OEM$\$1\Sysprep subfolder is optional. This subfolder contains the files that you need to run Sysprep. For information about these files, see "Designing Image-based Installations with Sysprep" in this book.

**$OEM$\drive_letter**
During the text mode stage of Setup, the structure and contents of each $OEM$*drive_letter* subfolder is copied to the root of the corresponding drive on the destination computer. For example, files that you place in the $OEM$\D subfolder are copied to the root of drive D. You can also create subfolders within these subfolders. For example, $OEM$\E\Misc causes Setup to create a subfolder called Misc on drive E.

**Tip**

- If you are using MS-DOS to start the installation, and your MS-DOS-based tools cannot copy folders with path names longer than 64 characters, you can use short file names for the folders and then use $$Rename.txt to rename them later.

**Converting Short File Names to Long File Names by Using $$Rename.txt**

If you are using Winnt.exe to perform an unattended installation, your distribution share can only contain files and folders that have short file names. This is because Winnt.exe only runs on MS-DOS, and MS-DOS can only process files and folders that use the 8.3 naming convention. The 8.3 naming convention allows only eight characters to the left of the decimal point and three characters to the right of the decimal point. Because of the short file name limitation, you need to shorten the names of files and folders so you can put them into your distribution share. You can convert these short file and folder names back to long file and folder names during setup by using the $$Rename.txt file.

$$Rename.txt lists all of the files and folders in a specific folder that need to be renamed. Each folder that contains short file names that need to be renamed must contain a separate version of $$Rename.txt. If you are using Winn32.exe to perform an unattended installation, short file names are automatically converted to long file names during setup.

To use $$Rename.txt, put the file in a folder that contains files and folders that need to be converted. Setup automatically looks for $$Rename.txt; if it finds a $$Rename.txt file, Setup renames the files in that folder. The syntax for $$Rename.txt is:

```
[section_name_1]
short_name_1 = " long_name_1"
short_name_2 = " long_name_2"
short_name_x = " long_name_x"
.
.
.
[section_name_2]
short_name_1 = " long_name_1"
short_name_2 = " long_name_2"
short_name_x = " long_name_x"
.
.
.
```

Where:

- *section_name_x* is the path to the folder that contains the files and folders. A section does not need to be named; it can have a backslash (\) as a name, which indicates that the section contains the names of the files or subfolders that are in the root of the drive.

- *short_name_x* is the name of the file or folder within this named folder that needs to be renamed. The name *must not* be enclosed in quotation marks.

- *long_name_x* is the new name of the file or folder. If the name contains spaces or commas, it *must* be enclosed in quotation marks.

# Designing the Media Distribution Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you plan to use an operating system CD to perform an unattended installation, you do not need to design a media distribution process. The installation files are already saved on the operating system CD, and you cannot change the structure of the operating system CD or add files to it. However, you do need to determine how to configure settings in the [Data] section of the answer file. Table 2.2 contains a description of these settings and guidelines to help you configure them. For a worksheet to help you record answer file settings, see "Answer File Settings Worksheet" (ACIUI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Answer File Settings Worksheet" on the Web at http://www.microsoft.com/reskit).

**Table 2.2 Answer File Settings for Installations That Use the Operating System CD**

| Answer File Setting ([Section]/Entry) | Comment |
|---|---|
| [Data]<br>AutoPartition = 1 | Installs Windows on the first available partition that has adequate space for a Windows installation and does not already contain an installed version of Windows. Either omit the AutoPartition entry from your answer file or set the value of AutoPartition to 1. If AutoPartition is set to 1, the **/tempdrive** parameter of Winnt32.exe is ignored during setup. If you do not set the value, text mode stage of Setup installs Windows on the partition where $WinIN_NTnt$.~lsLS is located. |
| [Data]<br>MsDosInitiated = 0 | Informs the Windows Setup Loader that an unattended installation is running directly from the operating system CD. You must set this value to 0. If you do not set the value to 0, setup fails at the beginning of the GUI mode stage of Setup. |
| [Data]<br>UnattendedInstall = Yes | Informs the Windows Setup Loader that an unattended installation is running directly from the operating system CD. You must set this value to Yes if you are installing Windows from the operating system CD. |
| [Data]<br>UseBIOSToBoot = 0 or 1 | Specifies whether Setup uses the BIOS to start the computer, even though Setup might detect that it is best to use a device miniport driver to start the computer. If you want to use the BIOS to start the computer, you must set this value to 1. The default setting is 0. Do not use this entry unless you are sure that the BIOS supports the extended int13 functions. |
| [Unattended]<br>Repartition = Yes or No | Specifies whether to delete all partitions on the first drive of the destination computer and to reformat the drive with NTFS. The default setting is No. Change this to Yes if you want the installation to manually run the GUI mode stage of Setup (the text mode stage of Setup is automated, but the GUI mode stage of Setup is not). If you leave the default setting, and Setup detects a Windows installation on the first drive of the destination computer, Setup will stop during the text mode stage of Setup and ask you whether you want to delete the existing Windows installation. |
| [Unattended]<br>UnattendSwitch = Yes or No | Specifies whether Setup skips Mini-Setup when installing Windows XP Professional or Windows Server 2003. Use only with Winnt.exe. The default setting is No. |

For more information about answer file settings, see *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Designing Preinstallation Tasks for Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you determine which distribution method to use, you need to identify your preinstallation tasks. Preinstallation tasks are performed before you begin an unattended installation. You might not need to perform any preinstallation tasks; you will need to perform some preinstallation tasks if you want to:

- Migrate user state before you perform an unattended installation. When you migrate user state, you save user settings and user data on some type of external media so you can restore the settings and data after the unattended installation is complete. You will likely not need to do this if you use folder redirection and roaming profiles to store user data and settings on a server. For more information about roaming profiles, see "Using roaming user profiles" in Help and Support Center for Windows Server 2003. For more information about folder redirection, see "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit.

- Change the size or format of the system partition before you perform an unattended installation.

- Use Dynamic Update to download updated installation files and device drivers with your unattended installation.

Figure 2.6 shows the design steps you need to follow to design your preinstallation tasks.

**Figure 2.6 Designing Preinstallation Tasks**



# Creating a User State Migration Plan for Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You will need to create a user state migration plan if any of your destination computers contain any of the following items that you want to restore after installation is complete:

- User data that you want to be available to the end user. User data includes such things as documents, e-mail messages, spreadsheets, and databases.
- User settings, such as desktop settings, shortcuts, and Internet Explorer Favorites.
- Application settings such as application-specific keyboard shortcuts, spell-checking options, and default file locations.

At a minimum, your user state migration plan must do the following:

- Identify the data you want to migrate, including user data, user settings, and application settings.
- Describe how to collect and restore the data.
- Determine where to store the data while you perform the unattended installation.
- Create a schedule for migrating the data on each of your destination computers.

Microsoft provides two tools for migrating user data and settings. The tool you use depends on your environment:

- **Files and Settings Transfer Wizard.** Designed for home users and small office users, the wizard is also useful in a corporate network environment for employees who get a new computer and need to migrate their own files and settings without the support of an IT department or Helpdesk.
- **User State Migration Tool.** Designed for IT administrators who perform large deployments of Windows XP Professional in a corporate environment, the User State Migration Tool provides the same functionality as the wizard, but on a large scale targeted at migrating multiple users. The User State Migration Tool gives administrators command-line precision for customizing specific settings, such as unique modifications to the registry. To download a free version of the tool, see the User State Migration Tool link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about migrating user data and settings, see "Migrating User State" in this book. Also see the articles "User State Migration in Windows XP," "Step-by-Step Guide to Migrating Files and Settings," "Deploying Windows XP Part I: Planning," and "Deploying Windows XP Part II: Implementing." To find these articles, see the Microsoft TechNet link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Creating a Disk Configuration Plan for Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to create a disk configuration plan if you need to do any of the following before you perform the installation:

- Change the size of the system partition on your destination computers.
- Repartition and reformat the system partition on your destination computers, and you are not using an operating system CD to start your destination computers.
- Create and format extra partitions on your destination computers.

You do not need to create a disk configuration plan in the following situations:

- You want to extend the system partition, create and format extra partitions during or after the installation, or convert an existing system partition to NTFS. These tasks do not require substantial analysis and planning, and are relatively easy to perform by configuring answer file settings or running commands or scripts from the answer file.
- You are using an operating system CD to perform an unattended installation, and you want to repartition and format the system partition on destination computers. In this case, you can use the Repartition entry in the [Unattended] section of the unattended installation answer file to repartition and format the system partition before the unattended installation begins.

**Configuring Disk Settings**
Begin your disk configuration plan by choosing the tool that most suits your needs:

**MS-DOS or Windows 98 disk configuration tools**
You can start a destination computer by using an MS-DOS or a Windows 98 boot disk, and then use the **fdisk** and **format** commands to partition and format the hard disk before you perform an unattended installation. This works only if you want to format your disks with the FAT or FAT32 file systems. If you want your hard disks formatted with NTFS, you will have to use the **convert** command to convert the FAT or FAT32 file system to NTFS after you have installed the operating system onto the destination computer, or you will have to use the **oformat** command. For more information about disk tools, including commands for configuring disks, see "Helpful Command Line Tools" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD. For more information about disk partitions and file systems, see the *Server Management Guide* of the *Microsoft® Windows® Server 2003 Resource Kit* (or see the Server Management Guide on the Web at http://www.microsoft.com/reskit).

**Third-party disk configuration tools**
Some third-party disk management programs provide a bootable floppy disk or CD that allows you to partition and format hard disks. If you use a third-party program to partition or format a disk, be sure that the third-party program creates partitions that are compatible with NTFS 3.1, which is the version of NTFS that is used in Windows XP Professional and Windows Server 2003 operating systems.

**Windows Preinstallation Environment**
You can start a destination computer by using a Microsoft® Windows® Preinstallation Environment (Windows PE) CD, and then using the **diskpart** command to partition the hard disk and the **format** command to format the hard disk. Windows PE is a bootable operating system that provides limited operating system functionality for performing preinstallation tasks. Windows PE is available only if you have purchased Enterprise Agreement 6.0, Enterprise Subscription Agreement 6.0, or Select License 6.0 with Software Assurance (SA). For more information about Windows PE and Windows PE licensing plans, see the Windows Preinstallation Environment link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Each method of configuring disk settings has advantages and disadvantages. You need to determine which method is better suited to your organization and your deployment needs. For more information about configuring disks, see "Disk Management overview" in Help and Support Center for Windows Server 2003.

**Components of a Disk Configuration Plan**

After you determine which method to use to configure disk settings, you need to create your disk configuration plan. At a minimum, your disk configuration plan must identify:

- **Disk configuration settings.** Disk configuration settings include the number of partitions, partition sizes, and file system formats for the destination computer. Disk configuration settings are based on several factors, including disk sizes, disk types, backup capabilities, and user needs. Analyze these factors in your disk configuration plan to determine the right disk configuration for your organization.

- **Procedures for configuring disk settings.** Your disk configuration plan must describe every step of the disk configuration process, including how to start a destination computer and how to run the partitioning or formatting tools.

- **Tools that you use to configure disk settings.** Disk configuration tools include the **format**, **fdisk**, and **diskpart** commands. Your disk configuration plan must describe all the tools you will use to configure disk settings, including the tools you will use to start a destination computer and to partition, format, and check a disk.

# Planning for Dynamic Update

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use Dynamic Update to update installation files and device drivers that are used by Setup during an unattended installation. Dynamic Update does not replace Windows Update; it downloads only a small subset of Windows Update files and device driver files that prevent critical errors from occurring during the setup process. The files that Dynamic Update downloads include:

- **Updated installation files.** This can include system files, in-box device drivers, Setup information (.inf) files required during upgrades, DLL files used by the Winnt32.exe Setup program, and file assemblies (.asm files). Dynamic Update downloads only replacements for existing installation files. It does not add new installation files to the setup process.

- **New device drivers.** This can include new device drivers that are critical to the setup process and are not on the operating system CD. New device driver files are not replacements for in-box device drivers. Replacements for in-box device drivers are considered updated installation files.

You can use Dynamic Update only if you are installing on destination computers that have an existing connection to your network or the Internet. For example, you can use Dynamic Update if your destination computer is running Windows 2000 and it is connected to your network when you run Windows Setup. In addition, you can use Dynamic Update only with Winnt32.exe; you cannot use Dynamic Update with Winnt.exe. By default, Dynamic Update is disabled during an unattended installation of Windows Server 2003 and Windows XP Professional.

In addition, if you are upgrading a computer that is running Windows 95 with Internet Explorer 4.01, you need to upgrade to Internet Explorer 5 or a higher version of Internet Explorer. The version of Secure Sockets Layer (SSL) in Internet Explorer 4.01 is not compatible with Dynamic Update, and will cause Dynamic Update to fail.

**Delivering Dynamic Update files to destination computers**

You can deliver Dynamic Update files to destination computers two ways: you can download Dynamic Update files across the Internet from the Windows Update Web site, or you can download Dynamic Update files across your corporate network from a shared folder that you create on a server in your organization. The latter method is better suited for large-scale corporate deployments because it ensures consistency among your destination computers.

When you download Dynamic Update files from a server in your organization, you can guarantee that the same set of Dynamic Update files are downloaded onto each of your destination computers because you have full control of the Dynamic Update files that are on the server. When you download Dynamic Update files across the Internet from the Windows Update Web site, you might introduce inconsistencies among your destination computers because the Windows Update Web site is periodically updated, and you cannot control when this occurs. In addition, downloading Dynamic Update files from a server in your organization eliminates many security issues that can arise when destination computers are connected to the Internet.

**Note**

- This book assumes you are delivering Dynamic Update files across your corporate network from a server in your organization. This book does not describe how to configure an unattended installation so that Dynamic Update files are delivered across the Internet.

**Preparing to use Dynamic Update**

To use Dynamic Update, you need to determine the following:

- Which files to deliver to your destination computers.

- How to prepare the files for delivery.

- How to configure answer files settings or Winnt32.exe settings.

For a worksheet to help you record information about your Dynamic Update design, see "Dynamic Update Worksheet" (ACIUI_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Dynamic Update Worksheet" on the Web at http://www.microsoft.com/reskit).

# Identifying and Downloading Dynamic Update Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to download two types of files to prepare for Dynamic Update: Dynamic Update packages, which contain updated installation files; and device driver .cab files, which contain new device driver files that are not present on the operating system CD. Device driver files must be downloaded individually; they are not assembled in packages.

**Downloading Dynamic Update Packages**
You can download Dynamic Update packages from the Microsoft Download Center. To do this, go to the Microsoft Download Center, specify the operating system that you are deploying, and then search for the keywords "dynamic update." Download the most current Dynamic Update package (the Dynamic Update package with the highest version number). To use the Microsoft Download Center, see the Microsoft Download Center link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For information about downloading the most recent Dynamic Update package for Windows XP Professional, see article Q311220, "Description of the Dynamic Update Feature in Windows XP Setup," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. For information about downloading the most recent Dynamic Update package for Windows Server 2003, search the Microsoft Knowledge Base using the keywords "dynamic update."

**Downloading device driver files**
You can download new device driver .cab files from the Windows Update corporate Web site. Before you do this, you need to identify all of the devices on your destination computers for which there are no device drivers on the operating system CD. To identify these devices, use an operating system CD to manually install the operating systems you are deploying on a representative sampling of your destination computers. You can then use Device Manager to check the device status for each device. Devices that do not have device drivers are denoted with a yellow question mark and a yellow exclamation mark in Device Manager. To download device driver .cab files, see the Windows Update link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Important**

- Use Dynamic Update to install missing device drivers only for devices that are critical for Setup, such as hard disk controllers, mice, video adapters, and keyboards. Do not use Dynamic Update to install missing device drivers for peripheral devices, such as scanners, cameras, and printers. To install missing device drivers for peripheral devices, or devices that are not critical for Setup, use the OemPnPDriversPath entry in the [Unattended] section of your answer file.

# Preparing Dynamic Update Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you download the Dynamic Update packages and the device driver .cab files, you need to prepare these files for delivery to your destination computers. Perform the following tasks to prepare Dynamic Update files.

**Extract the Dynamic Update packages**
Run the executable Dynamic Update package to extract the Dynamic Update files. This creates separate folders for each operating system. For example, if you run the executable Dynamic Update package for Windows XP, this creates two folders: a Windows XP Professional folder named IP, and a Windows XP Home Edition folder named IC. If you run the executable Dynamic Update package for Windows Server 2003, this creates three folders in your new folder: a Windows Server 2003, Standard Edition folder named Standard; a Windows Server 2003, Enterprise Edition folder named Enterprise; and a Windows Server 2003, Web Edition folder named Web. Each folder can contain one or more of the following files: Winnt32.cab, Updates.cab, Upginf.cab, Duasms.cab. Delete the folders for any operating systems that you are not installing. For example, if you are not installing Windows XP Home Edition, delete the IC folder.

**Copy the .cab files to new folders**
Copy the .cab files for each operating system, and each of the new device driver .cab files for each operating system, to a new folder. For example, copy the .cab files in the IP folder to a folder named DU_XPPro, and copy any new device driver files you downloaded for Windows XP Professional to the DU_XPPro folder.

**Use the /duprepare parameter to prepare each folder**
Run Winnt32.exe with the **/duprepare** parameter on each of the new folders. This prepares each folder for Dynamic Update. During the preparation process, Winnt32.exe creates one or more of the following subfolders: Duasms, Dudrvs, Updates, Upginfs, and Winnt32. Winnt32.exe also copies the contents of the .cab files to one or more of these subfolders.

For more information about preparing files for Dynamic Update, see article Q312110, "How to Deploy the Windows XP Dynamic Update Package," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Configure a Shared Folder for Dynamic Update files**
To distribute Dynamic Update files across your corporate network, you need to create a shared folder on a server. The shared folder must be available to all of your destination computers. For example, do not create the shared folder on a server that is in a restricted subnet. Typically, you create the shared folder on the same server that contains your distribution shares.

In addition, you need to assign permissions to the shared folder. If you create the shared folder on the same server that contains your distribution shares, the shared folder should have the same permissions as your distribution shares. Permissions ensure that only authorized users can access the shared folder. If the shared folder is not secure, a malicious user could tamper with the Dynamic Update files. For example, a malicious user could replace the .cab files in the shared folder with .cab files that contain a virus. For more information about permissions, see "Best practices for permissions and user rights" in Help and Support Center for Windows Server 2003.

**Copy the prepared Dynamic Update Files to the shared folder**
After you create the shared folder on a server, copy each of the prepared Dynamic Update folders to it. The "Dynamic Update Worksheet" (ACIUI_4.doc) shows the structure of the shared folder, and provides spaces for you to record the names of the folders containing your Dynamic Update files. See "Dynamic Update Worksheet" (ACIUI_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Dynamic Update Worksheet" on the Web at http://www.microsoft.com/reskit).You will need the names of these folders to design your answer file settings and Setup program settings later in this chapter.

# Configuring Answer File and Winnt32.exe Settings for Dynamic Update

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Dynamic Update is disabled by default when you perform an unattended installation with an answer file. Dynamic Update is enabled by default when you perform an unattended upgrade installation by running **Winnt32.exe /unattend**.

To use Dynamic Update during an unattended installation with an answer file, you need to add the following entry to the [Unattended] section of your answer file:

```
DUDisable=No
```

This entry tells Setup to enable Dynamic Update. By default, this entry is set to Yes, which means Dynamic Update is disabled.

**Important**

- Do not use the **/dudisable** parameter with Winnt32.exe if you enable Dynamic Update in your answer file. The **/dudisable** parameter takes precedence over the DUDisable entry in your answer file: that is, using the **/dudisable** parameter disables Dynamic Update regardless of how you configure your answer file.

In addition, you need to specify the location of the shared folder that contains the prepared Dynamic Update files. You can do this by creating the following entry in the [Unattended] section of your answer file:

```
DUShare =
path_to_prepared_dynamic_update_files
```

You can also do this by using the following command-line parameter with Winnt32.exe:

```
/DUShare:
path_to_prepared_dynamic_update_files
```

In both cases, *path_to_prepared_dynamic_update_files* is the path to the folder containing the Dynamic Update files that you downloaded, prepared, and saved on a secure shared folder in your organization. This path is different for each operating system you are installing.

In addition, when you perform an unattended installation, you can change the way Setup responds to Dynamic Update errors by using the DuStopOnError parameter in the [Unattended] section of your answer file. Dynamic Update errors include any failure to process Dynamic Update files, or the inability to connect to Windows Update. By default, the Dynamic Update process stops when an error is detected. You can change this by adding the following entry to your answer file:

```
DuStopOnError=No
```

For a worksheet to help you record answer file information, see "Answer File Settings Worksheet" (ACIUI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Answer File Settings Worksheet" on the Web at http://www.microsoft.com/reskit). For more information about answer file settings and Winnt32.exe parameters, see *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Designing Answer File and Setup Settings for Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By creating an answer file and by using Setup command line parameters, you can automate the following tasks that occur during and after an unattended installation:

- **Software installation and configuration.** You can automatically install and configure client and server applications. You also can install and configure Windows components, services, and applications.
- **Hardware installation and configuration.** You can automatically update device drivers and configure device settings.
- **Computer configuration.** You can automatically configure computer settings, such as domain membership, computer name, network protocols, display settings, and system services. You also can configure server roles, such as installing Active Directory.

Try to automate as many installation and post-installation tasks as possible during an unattended installation. By automating installation and post-installation tasks, you can:

- Reduce the number of errors caused by technicians, administrators, and end users during your deployment.
- Ensure consistency throughout your organization, which reduces support costs after deployment.
- Increase productivity by requiring little or no end-user interaction during your deployment.
- Update or modify your installation process without having to educate or retrain end users, technicians, or administrators.

To automate installation and post-installation tasks, you need to create and configure an answer file. In addition, you need to choose which command line parameters to use when you start the Setup program.

The answer file for unattended installations of Windows Server 2003 and Windows XP Professional is usually named Unattend.txt, but you can name it anything you want. However, if you start a destination computer from an operating system CD and the answer file is on a floppy disk, you must name the answer file Winnt.sif. Setup will then detect the answer file on the disk without user input.

The answer file supplies Setup with answers to all the questions that you are asked during a standard, interactive installation. The answer file also contains information about your installation and configuration requirements. In addition, an answer file tells Setup how to interact with the distribution shares and files that you have created (or the installation files on the operating system CD, if that is what you are using to install the product).

An answer file contains multiple sections — some mandatory and some optional — that you can modify. Section names are delineated by brackets (for example, [Unattended]). Every section has one or more entries that contain installation and configuration information. Entries are delineated by an entry name, an equal sign, and a value (for example, ComputerName = Computer1). The entry name represents a specific computer setting or action; the value represents the unique way you want the setting configured or the action performed. For more information about answer file sections and entries, see "Unattend.txt" in the "Reference" section of the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

Setup parameters are specified at the command line when you run Winnt.exe or Winnt32.exe.

Some installation and configuration tasks can be performed by configuring either Setup parameters or answer file settings. For example, you can implement Dynamic Update by configuring answer file settings or by configuring Setup parameters when you run Winnt32.exe. In both cases, the functionality is the same. For more information about Setup, see "Using Winnt.exe to Run Setup" and "Using Winnt32.exe to Run Setup" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

Figure 2.7 shows the design process you need to follow to automate tasks during an unattended installation.

**Figure 2.7 Designing Answer Files and Setup Parameters**



# Designing Automated Installation Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To design automated installation tasks for an unattended installation, you must identify:

- The installation tasks you want to automate.
- The settings you need to configure for the answer file.

Installation tasks are usually described in your component design or your client and server configuration design. For example, your network topology should provide information about the network protocol settings you need to configure during an installation. Likewise, your file server design should provide information about the disk settings you need to configure during an installation.

Answer file settings correspond to the installation tasks you need to perform. For example, if one of your installation tasks involves configuring network adapter settings, you need to identify the answer file sections and entries that configure network adapter settings, and then determine the proper values to assign to the answer file entries.

**Identifying Automated Installation Tasks**

Use Table 2.3 to determine which installation tasks you can automate by configuring an answer file. Try to automate as many installation tasks as possible with the answer file.

**Table 2.3 Tasks You Can Automate with an Unattended Installation Answer File**

| Installation Task | Comments |
|---|---|
| Hardware installation and configuration | This includes installing and configuring mass storage controllers that are required at startup, such as SCSI hard disks and Plug and Play devices that are not included on the operating system CD. |
| Setup configuration | This includes partitioning and formatting hard disks prior to setup, and configuring Setup options that are usually configured by end users during GUI mode and text mode stage of Setup. This also includes configuring upgrade options, uninstall options, and other settings that affect the way Setup runs. |
| Operating system configuration | This includes configuring power management, telephony, and display settings, and regional options. This also includes configuring error reporting, Windows file protection, remote assistance, system restore, licensing, and shell settings. |
| Internet Explorer configuration | This includes configuring Internet Explorer options, such as favorites, proxy server settings, branding, and default Home and Search pages. Also includes configuring Internet Explorer Enhanced Security Configuration settings.* |
| Networking configuration | This includes configuring Internet Connection Sharing (ICS), Windows Firewall, and domain membership settings. This also includes installation and configuration of protocols, network adapters, and networking services and components. |
| Services configuration | This includes configuring Internet Information Services (IIS), Certificate Services, Remote Installation Services, Terminal Server, fax service, and Simple Network Management Protocol (SNMP) service. This also includes installing and configuring a domain controller by using Active Directory. |
| Optional Windows components and services installation | This includes all Windows components listed in Add or Remove Programs in Control Panel, such as accessories, games, media services, and Indexing Service. |
| Software application installation and configuration | This includes Windows Installer (.msi) packages and staged software. Software installation must run in quiet mode, which means the installation must be fully automated and cannot rely on user interaction. Usually, when you run an installation program in quiet mode, you must provide an answer file. |

| | |
|---|---|
| Running programs, scripts, and batch files | Programs, scripts, and batch files must be fully automated and cannot rely on user interaction, which means you must provide an answer file for any programs, scripts, or batch files you are running, and you must be able to run the programs, scripts, or batch files in quiet mode. |

* For more information about Internet Explorer Enhanced Security Configuration settings, see "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003. For more information about answer file settings related to Internet Explorer Enhanced Security Configuration, see the Readme.txt file in Deploy.cab. Deploy.cab is in the Support folder on the Windows Server 2003 operating system CD.

**Identifying Answer File Settings**

Use Table 2.4 to find the specific answer file section that corresponds to each installation task. For a Word document to assist you in recording your answer file settings, see "Answer File Settings Worksheet" (ACIUI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Answer File Settings Worksheet" on the Web at http://www.microsoft.com/reskit).

**Table 2.4 Installation Tasks and Corresponding Answer File Sections**

| To Automate This Task | Configure These Answer File Sections |
|---|---|
| Hardware installation and configuration | [MassStorageDrivers]<br>[OEMBootFiles]<br>[Unattended] |
| Setup configuration | [Data]<br>[GuiUnattended]<br>[Unattended]<br>[Win9xUpg] |
| Operating system configuration | [Display]<br>[LicenseFilePrintData]<br>[PCHealth]<br>[RegionalSettings]<br>[Shell]<br>[SystemFileProtection]<br>[SystemRestore]<br>[TapiLocation]<br>[UserData] |
| Internet Explorer configuration | [Branding]<br>[FavoritesEx]<br>[Proxy]<br>[URL]<br>[Components]<br>[IEHardening] |

| | |
|---|---|
| Networking configuration | [Homenet]<br>[Identification]<br>[MS_AppleTalk parameters]<br>[MS_ATMArps parameters]<br>[MS_ATMLANE parameters]<br>[MS_ATMUni parameters]<br>[MS_L2TP parameters]<br>[MS_MSClient parameters]<br>[MS_NetMon parameters]<br>[MS_NWClient parameters]<br>[MS_NWIPX parameters]<br>[MS_NwSapAgent parameters]<br>[MS_PPTP parameters]<br>[MS_Psched parameters]<br>[MS_RAS parameters]<br>[MS_RasSrv parameters]<br>[MS_Server parameters]<br>[MS_TCPIP parameters]<br>[MS_WLBS parameters]<br>[NetAdapters]<br>[NetBindings]<br>[NetClients]<br>[NetOptionalComponents]<br>[NetProtocols]<br>[NetServices]<br>[Networking] |
| Services configuration | [DCInstall]<br>[Fax]<br>[InternetServer]<br>[OsChooser]<br>[RemoteInstall]<br>[SNMP]<br>[TerminalServices]<br>[CertSrv_Client]<br>[CertSrv_Server] |
| Optional Windows components and services installation | [Components] |

| | |
|---|---|
| Software application installation and configuration | [GuiRunOnce]<br><br>[SetupParams] |
| Running programs, scripts, and batch files | [GuiRunOnce]<br><br>[SetupParams] |

Most answer file settings are optional; however, if you want a fully automated unattended installation, you must configure the following sections and entries in your answer file. If you do not provide values for these entries in your answer file, Setup will prompt the end user (or whoever is performing the unattended installation) for the values:

**[GuiUnattended]** You must specify values for AdminPassword and TimeZone. The value for AdminPassword cannot begin with an asterisk (*). Using a password that begins with an asterisk can cause the password to be set to a null value.

**[Identification]** You must specify values for JoinDomain, DomainAdmin, and DomainAdminPassword.

**[LicenseFilePrintData]** You must specify values for AutoMode and AutoUsers if you are installing a Windows Server 2003 product.

**[Networking]** If your destination computer requires network connectivity, you must specify values for various network protocol entries.

**[Unattended]** You must specify values for UnattendMode and TargetPath.

**[UserData]** You must specify values for FullName and ComputerName.

For more information about specific answer file settings — including procedural and reference information about creating, formatting, and configuring an answer file for unattended installation — see *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Designing Automated Post-Installation Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Post-installation tasks include any installation and configuration tasks that need to be performed after the operating system is installed. Post-installation tasks are usually described in your component design or your client and server configuration design. You usually perform these tasks by running a command, program, script, or batch file after Setup is finished running. For example, your client networking design might include information about network settings that need to be configured with the Netset.exe command line tool. Likewise, your client configuration design might include information about mapped network drive settings that need to be configured with the **net use** command or printers that need to be installed with the Rundll32.exe program.

You can automate these types of tasks after the operating system is installed by using a Cmdlines.txt file or by using the [GuiRunOnce] section in your answer file. Both methods allow you to run commands, programs, scripts, and batch files.

To design automated post-installation tasks for an unattended installation, you must:

- Identify the tasks you want to perform after the operating system is installed.
- Design a method for automating the tasks you want to perform.

# Identifying Automated Post-Installation Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use an answer file to automate only a limited number of installation and configuration tasks during installation. Many installation and configuration tasks must be performed after the operating system is installed and configured. Testing your unattended installations is the best way to determine whether an installation or configuration task must be performed after the operating system is installed. However, the following installation and configuration tasks always must be performed after the operating system is installed:

- Tasks that cannot be performed by setting an answer file entry. You cannot add any other sections or entries to an answer file that is used to perform an unattended installation. For a worksheet that contains a list of all possible answer file sections and entries, see "Answer File Settings Worksheet" (ACIUI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Answer File Settings Worksheet" on the Web at http://www.microsoft.com/reskit).

- Tasks that rely on Active Directory directory service. For example, if a software installation program registers information in Active Directory or requires information from Active Directory, you must run the installation program after the operating system is installed and the computer is joined to a domain.

- Tasks that can be performed only while a user is logged on. For example, some software installation programs create shortcuts on the **Start** menu and the desktop. If you want these shortcuts applied to a specific user profile, then you need to run the installation program after the operating system is installed and the user is logged on.

You can automate post-installation tasks only if you can run the command, program, script, or batch file in quiet mode, or if you can suppress all user prompts by supplying an answer file for the command, program, script, or batch file. For example, if you delete folders after you install the operating system, you need to use the **/q** parameter with the **rmdir** command. Likewise, if you install Microsoft Word after you install the operating system, you need to create a Setup.ini file that provides configuration information to the Word Setup program.

# Choosing a Method for Automating Post-Installation Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can perform automated post-installation tasks two ways: you can use a Cmdlines.txt file to run commands, programs, scripts, or batch files just after Setup finishes but before the computer restarts, or you can configure the [GuiRunOnce] section in an answer file to run commands, programs, scripts, and batch files after the computer restarts and a user logs on. Use the following guidelines to determine which method to use.

**Using Cmdlines.txt**
Use Cmdlines.txt when:

- You are running commands, programs, scripts, or batch files from the $OEM$ folder on a distribution share.
- You want to install applications or perform configuration tasks during GUI mode stage of Setup.
- You want to install applications or perform configuration tasks under the Local System security account.
- You do not need network connectivity to perform the installation or configuration task.
- You are not using Windows Installer packages (.msi files) to install applications.
- You want to install applications or perform configuration tasks while no user is logged on.

**Using [GuiRunOnce]**
Use [GuiRunOnce] when:

- You need access to hard drives, CD-ROM drives, shared folders on the network, or other storage devices.
- You want to install applications or perform configuration tasks under a specific user account.
- You need network connectivity to perform an installation or configuration task.
- You are using Windows Installer packages (.msi files) to install applications.
- You need to control the order in which programs, scripts, or batch files run.

**Important**

- You cannot use Cmdlines.txt if you are using an operating system CD and a Winnt.sif file to perform an unattended installation. You can use Cmdlines.txt only if you are installing from a distribution share.

For information about configuring Cmdlines.txt and [GuiRunOnce] to run programs, scripts, and batch files, see "Configuring Cmdlines.txt to Perform Tasks" and "Configuring [GuiRunOnce] to Perform Tasks" later in this chapter.

# Configuring Cmdlines.txt to Perform Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Cmdlines.txt file contains a list of commands, programs, scripts, or batch files that are executed at the end of GUI mode stage Setup. You can configure Cmdlines.txt to run multiple commands, programs, scripts, or batch files. The commands, programs, scripts, and batch files run synchronously; for example, a command waits for the previous command to finish running before it starts. This means that you can control the order in which commands, programs, scripts, and batch files run. By default, Cmdlines.txt is not created, so you must manually create the file and save it in the $OEM$ folder in your distribution share.

The syntax for Cmdlines.txt is as follows:

```
[Commands]
"command_1""command_2"
      .
      .
"command_x"
```

Where *command_1*, *command_2*, and *command_x* refer to the commands, programs, scripts, or batch files that you want to run when GUI mode stage of Setup is complete. Note that all commands must be enclosed in quotation marks. Also, if you are using the command line (Cmd.exe) to run commands, programs, scripts, or batch files, then you need to use the **/c** parameter with the **cmd** command. For example, to create a new folder named Test on drive C, you type:

```
" cmd /c mkdir c:\Test"
```

To use Cmdlines.txt, you need to:

- Add the following section, entry, and value to your answer file:

  ```
  [Unattended]
  OemPreinstall = Yes
  ```

- Configure your Cmdlines.txt file, and save it in the $OEM$ folder in your distribution share.

- Copy all of the programs, scripts, and batch files that are listed in your Cmdlines.txt file to the $OEM$ folder in your distribution share.

For a worksheet to assist you in recording the commands, programs, scripts, and batch files that you want to run by using Cmdlines.txt, see "Answer File Settings Worksheet" (ACIUI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Answer File Settings Worksheet" on the Web at http://www.microsoft.com/reskit).

# Configuring [GuiRunOnce] to Perform Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The [GuiRunOnce] section of an answer file contains a list of commands, programs, scripts, or batch files that run the first time a user logs on to the computer after the operating system is installed. You can configure [GuiRunOnce] to run multiple commands, programs, scripts, or batch files. The commands, programs, scripts, and batch files run synchronously, which means each command, program, script, or batch file runs to completion before the next one starts running. This lets you control the order in which tasks are performed.

The syntax for [GuiRunOnce] is:

```
[GuiRunOnce]
"command_1""command_2"
      .
      .
"command_x"
```

Where *command_1*, *command_2*, and *command_x* refer to the commands, programs, scripts, or batch files that you want to run after the operating system is installed and a user logs on. Note that all commands must be enclosed in quotation marks.

To use [GuiRunOnce], you need to:

- Add the [GuiRunOnce] section and the corresponding commands to your answer file.
- Copy all of the programs, scripts, and batch files that are listed under [GuiRunOnce] to the $OEM$ folder in your distribution share.

For a worksheet to assist you in recording the commands, programs, scripts, and batch files that you want to run by using [GuiRunOnce], see "Answer File Settings Worksheet" (ACIUI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Answer File Settings Worksheet" on the Web at http://www.microsoft.com/reskit).

If you use the [GuiRunOnce] section to install software, you need to adhere to the following guidelines:

**You cannot run installation programs that require a restart**
You can automate installation tasks only if you can prevent the installation program from restarting the computer. When a computer restarts, all remaining entries in the [GuiRunOnce] section are lost. If the system restarts before completing entries listed in the [GuiRunOnce] section, the remaining items will not run. Therefore, you need to suppress restarts. If you cannot suppress a restart within the installation program, you can try to repackage the application into a Windows Installer package. For more information about Windows Installer packaging, see the Windows Installer documentation link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**You cannot run installation programs that rely on the Windows Explorer shell**
The Windows Explorer shell is not loaded when the operating system starts running commands, programs, scripts, or batch files that are listed in the [GuiRunOnce] section of an answer file. In some cases, you can get an updated installer program from the application vendor that does not rely on the Windows Explorer shell. If you cannot, you might be able to repackage the application as a Windows Installer package (.msi file).

**You might need to run an installation program from a batch file so you can control the installation process with the /wait parameter**
Installation programs often start and stop several different processes. In some cases, when you are installing multiple applications, this can inadvertently cause the next installation program listed in the [GuiRunOnce] section to start before the previous installation program is finished running. When this occurs, the second installation program usually fails. To prevent this, you can run the installation programs from batch files by using the **start** command with the **/wait** parameter. This forces each installation program to run to completion before the next command listed in [GuiRunOnce] runs. For more information about using the **start** command to run installation programs in batch files, in Help and Support Center for Windows Server 2003, under **Support Tasks**, click **Tools**, click "Command-line reference A-Z", and then click Start.

# Designing Setup Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An unattended installation relies on one of two Setup programs: a 16-bit Setup program called Winnt.exe, or a 32-bit Setup program called Winnt32.exe. You can run Winnt.exe on a computer that is running MS-DOS, Windows 3.1, or Windows for Workgroups. You can run Winnt32.exe on a computer that is running Windows 95, Windows 98, Windows Millennium Edition, Windows NT, Windows 2000, Windows XP, and the Windows Server 2003 family.Winnt.exe and Winnt32.exe can be found on a Windows XP Professional or Windows Server 2003 product CD. For a worksheet to assist you in recording Setup options for each of your unattended installations, see "Unattended Installation Worksheet" (ACIUI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Unattended Installation Worksheet" on the Web at http://www.microsoft.com/reskit).

# Choosing Winnt.exe Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can start Setup by using the Winnt.exe command line tool from MS-DOS, Windows for Workgroups, or Windows 3.1. Winnt.exe has the following syntax:

```
winnt [/s:[sourcepath]] [/t:[tempdrive]] [/u:[answer_file]] [/udf:id
[,UDB_file]] [/r:folder] [/rx:folder] [/e:command] [/a]
```

Use the following parameter descriptions to determine which parameters to use for your unattended installation. For a worksheet to assist you in recording these parameters, see "Unattended Installation Worksheet" (ACIUI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Unattended Installation Worksheet" on the Web at http://www.microsoft.com/reskit).

**/s**
Specifies the source location of the Windows operating system files. The location must be a full path of the form *drive_letter***:***path_to_folder* or *server**path_to_shared_folder*. The default is the current folder. Use this parameter if you are using a distribution share.

**/t**
Directs Setup to place temporary files on the specified drive and to install Windows on that drive. If you do not specify a location, Setup attempts to locate a drive for you.

**/u**
Performs an unattended installation with an answer file. The answer file provides answers to some or all of the prompts that the end user normally responds to during setup. If you use **/u**, you must also use **/s**.

**/udf**
Indicates an identifier (*id*) that Setup uses to specify how a uniqueness database file (.udf) modifies an answer file (see **/u**). The .udf file overrides values in the answer file, and the identifier determines which values in the .udf file are used. If no .udf file is specified, Setup prompts the user to insert a disk that contains the .udf file.

**/r**
Specifies an optional folder to create during setup. For example, you might create a folder named Drivers on the C drive, and then copy device driver files to the folder after the installation completes. The folder remains after Setup finishes.

**/rx**
Specifies an optional folder to create and to copy the contents of another folder to during setup. The folder is deleted after the Setup program stops running. For example, you might create a folder named Scripts on the C drive, and then copy scripts that you want to run during the installation to that folder. After the installation completes, the scripts and the folder are deleted.

**/e**
This parameter specifies a command to run after the GUI mode stage of Setup.

**/a**
This parameter enables accessibility options.

You cannot use Winnt.exe to perform an upgrade. You can perform only clean installations when you use Winnt.exe. In addition, you must use the 8.3 naming convention to name all of the files and folders in your distribution share. To adhere to the 8.3 naming convention, you cannot have more than 8 characters before the decimal point, and you must have a three-character extension after the decimal point. You need to use a $$Rename.txt file to convert 8.3 file names to long file names during setup. For more information about renaming files and folders during setup, see "Designing a Distribution Share" earlier in this chapter.

For more information about Winnt.exe parameters, see "Winnt.exe Command Line Options" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Choosing Winnt32.exe Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can start Setup by using the Winnt32.exe command line tool from Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000, Windows XP Professional, and the Windows Server 2003 family. You cannot run Winnt32.exe on an Itanium-based computer from the extensible firmware interface (EFI), and no Winnt32.efi is available.

Winnt32.exe has the following syntax:

```
winnt32 [/checkupgradeonly] [/cmd:command_line] [/cmdcons] [/copydir:
{i386|ia64}\folder_name] [/copysource:folder_name] [/debug[level]:[filename]]
[/dudisable] [/duprepare:pathname] [/dushare:pathname] [/emsport:
{com1|com2|usebiossettings|off}] [/emsbaudrate:baudrate] [/m:folder_name]
[/makelocalsource] [/noreboot] [/s:sourcepath] [/syspart:drive_letter:]
[/tempdrive:drive_letter:] [/udf:id [,UDB_file]] [/unattend[num]:[answer_file]]
```

Use the following parameter descriptions to determine which parameters to use for your unattended installation. For a worksheet to assist you in recording these parameters, see "Unattended Installation Worksheet" (ACIUI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Unattended Installation Worksheet" on the Web at http://www.microsoft.com/reskit).

**/checkupgradeonly**
Checks your computer for upgrade and installation compatibility with Windows XP Professional and Windows Server 2003.

If you use this option with **/unattend**, no user input is required. Otherwise, the results are displayed, and you can save them with a specified file name. The default file name is Upgrade.txt in the systemroot folder.

**/cmd**
Instructs Setup to execute a specific command before beginning the final stage of the setup process. This occurs after your computer has restarted and after the Setup program has collected the necessary configuration information, but before the setup process finishes.

**/cmdcons**
Installs the Recovery Console as a startup option on a functioning x86-based computer. The Recovery Console is a command-line interface from which you can perform tasks such as starting and stopping services and accessing the local drive (including drives formatted with NTFS). You can use the Recovery Console only after Setup finishes.

**/copydir**
Creates one or more folders in the folder where the Windows files are installed.

**/copysource**
Creates one or more temporary folders in the folder where the Windows files are installed.

**/debug**
Creates a debug log at the level specified, for example, **/debug4:Debug.log**. The default log file is %SYSTEMROOT%\Winnt32.log, and the default debug level is 2. The log levels are: 0, representing severe errors; 1, representing errors; 2, representing warnings; 3 representing information; and 4, representing detailed information for debugging. Each level includes the levels under it.

**/dudisable**
Prevents Dynamic Update from running. Without Dynamic Update, Setup runs only with the original Setup files. This option disables Dynamic Update even if you set DUDisable equal to No in the [Unattended] section of your answer file.

**/duprepare**
Prepares a distribution share so that it can be used with Dynamic Update files downloaded from the Windows Update Web site. This distribution share can then be used for installing Windows for multiple clients.

**/dushare**
Specifies a distribution share on which you previously downloaded Dynamic Update files (updated files for use with Setup) from the Windows Update Web site, and on which you previously ran **/duprepare:***pathname*. When run on a client, specifies

that the client installation uses the updated files on the distribution share specified in *pathname*.

**/emsport**
Enables or disables Emergency Management Services during setup and after a member of the Windows Server 2003 family of operating systems has been installed. With Emergency Management Services, you can remotely manage a server in emergency situations that normally require a local keyboard, mouse, and monitor, such as when the network is unavailable or the server does not function properly. Emergency Management Services has specific hardware requirements, and is available only for products in the Windows Server 2003 family. For more information about Emergency Management Services, see "Planning for Remote Server Management" in *Planning Server Deployments* of this kit.

**/emsbaudrate**
For x86-based computers, this parameter specifies the baud rate for Emergency Management Services. (The option is not applicable for Itanium-based computers.) Must be used with **/emsport:com1** or **/emsport:com2** or else **/emsbaudrate** is ignored.

**/m**
Specifies that Setup copies replacement files from an alternative location. Instructs Setup to look in the alternative location first, and, if files are present, to use them instead of the files from the default location.

**/makelocalsource**
Instructs Setup to copy all installation source files to your local hard disk. Use **/makelocalsource** when installing from a CD to provide installation files when the CD is not available later in the installation.

**/noreboot**
Instructs Setup not to restart the computer after the file-copy stage of Setup finishes, so that you can execute another command.

**/s**
Specifies the location of the Windows files. To simultaneously copy files from multiple servers, type the **/s:***sourcepath* option multiple times (up to a maximum of eight).

**/syspart**
On an x86-based computer, this parameter specifies that you can copy Setup startup files to a hard disk, mark the disk as active, and then install the disk onto another computer. When you start the computer onto which you have installed the disk, it automatically starts with the next phase of Setup. You must always use the **/tempdrive** parameter with the **/syspart** parameter.

You can start Winnt32.exe with the **/syspart** option on an x86-based computer running Windows NT 4.0, Windows 2000, or Windows XP Professional. The computer cannot be running Windows 95, Windows 98, or Windows Millennium Edition.

**/tempdrive**
Directs Setup to place temporary files on the specified partition. For a new installation, Windows is installed on the specified partition. For an upgrade, the **/tempdrive** option affects the placement of temporary files only; the operating system is upgraded in the partition from which you run Winnt32.exe.

**/udf**
Indicates an identifier (*id*) that Setup uses to specify how a uniqueness database file (.udf) modifies an answer file (see the **/unattend** option).

If you start from the Windows Server 2003 operating system CD and run an unattended setup, you cannot use the **/udf** command-line option for Winnt32.exe.

**/unattend**
Upgrades your previous version of Windows 98, Windows Millennium Edition, Windows NT 4.0, or Windows 2000 in unattended mode (without user input). Setup downloads the Dynamic Update files from Windows Update and includes these files in the installation. All user settings are taken from the previous installation, so no user intervention is required during setup.

**/unattend[num]:[answer_file]**
Performs a fresh installation of Windows in unattended mode using the specified answer file. Setup downloads the Dynamic Update files from the Windows Update Web site and includes these files in the installation. The specified *answer_file* provides Setup with your custom specifications.

For more information about Winnt32.exe parameters, see "Winnt32.exe Command Line Options" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on

the Windows Server 2003 operating system CD.

# Creating Startup Media, Answer Files, and Distribution Shares

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you design your unattended installation, you need to implement your design. Usually, your deployment team is responsible for implementing your design decisions. The worksheets that are referenced in this chapter contain sufficient design information for your deployment team to perform these tasks.

Figure 2.8 shows the steps you must follow to create startup media, answer files, and distribution shares for an unattended installation.

**Figure 2.8 Creating Startup Media, Answer Files, and Distribution Shares**



# Creating Startup Media for Destination Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are performing clean installations, and you want to configure the hard disks on your destination computers, you might need to create startup media to start the destination computer. This section provides guidelines for choosing and creating appropriate startup media.

If you are performing upgrades, you can skip this section because you cannot use startup media to start destination computers. You must start the destination computers by using the operating system that is installed on those computers.

Likewise, if you are performing clean installations on destination computers that are running Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000, or Windows XP Professional, and you do not want to configure hard disks on your destination computers, then you do not need to create startup media. You can perform an unattended installation by starting the destination computers with the operating system that is already installed on those computers.

# Choosing the Startup Media

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Choosing startup media is a two-step process. First, you determine what type of startup media your hardware supports. Not every organization can support CD or DVD startup media. Next, you determine whether one particular type of startup media is more appropriate than another, based on the way you are performing your unattended installations.

**Evaluating Hardware Support for Startup Media**
Follow these steps to determine which type of startup media your organization can support.

1. Evaluate your hardware inventory for floppy disk support.

   To use a floppy disk as startup media, every destination computer must have a floppy drive, and the boot-order sequence in every BIOS must list the floppy drive.

2. Evaluate your hardware inventory for CD or DVD support.

   To use a CD or DVD as startup media, all of your destination computers must have bootable CD or DVD drives. Some older CD drives and many DVD drives are not bootable devices. In addition, the boot-order sequence in the BIOS of each computer must include the CD or DVD drive. Some older BIOSes do not let you add the CD or DVD drive to the boot-order sequence.

3. Evaluate your CD or DVD writeable device.

   To create your own CD or DVD startup media, you must have the proper hardware, software, and instructions to create bootable CDs or DVDs. Microsoft does not provide tools for creating CD or DVD startup media; however, several manufacturers provide the hardware, software, and system files that you need to create bootable CDs or DVDs.

If your organization supports only floppy disk startup media, you are ready to create your startup media. For more information about creating startup media, see "Creating the Startup Media" later in this chapter.

**Choosing Startup Media**
If you have the proper hardware and software to create CD or DVD startup media, and the destination computers in your organization support CD or DVD startup media, you need to determine which type of startup media is most suitable for your unattended installation.

Use a floppy disk to start your destination computers if:

- Your design requirements allow you to use Winnt.exe to start Windows Setup. Winnt.exe runs only under MS-DOS, and has limited functionality compared to Winnt32.exe. For more information about Winnt.exe and Winnt32.exe, see "Performing Unattended Installations" later in this chapter, and the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

- You want to create partitions or format disks before you install the operating system onto the destination computer, and your disk configuration tools can run under MS-DOS.

- You do not want to use Dynamic Update to download and install critical updates and device drivers from the Internet. You can use Dynamic Update to download and install critical updates and device drivers from a server within your organization, but only if you start your destination computers with a floppy disk that provides network support.

Use a CD or DVD to start your destination computers if:

- Your design requires Winnt32.exe to start Windows Setup. Winnt32.exe is a 32-bit program that has greater functionality than Winnt.exe. However, Winnt32.exe can be used only if you start a destination computer with the operating system CD; or with Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000 (32-bit edition),

Windows XP Professional (32-bit edition), or Windows Server 2003 (32-bit edition). You cannot run Winnt32.exe on a destination computer that is running Windows Advanced Server, Limited Edition (64-bit edition). For more information about Winnt.exe and Winnt32.exe, see "Performing Unattended Installations" later in this chapter and the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

- You want to create partitions or format disks before you install the operating system onto you destination computer, and your disk configuration tools can run under a 32-bit operating system.

# Creating the Startup Media

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Startup media contains the system files and device drivers that are necessary to start a computer so that the primary hard disk is accessible, but not in use. Startup media might also contain network adapter and network drivers, CD and DVD device drivers, disk configuration tools, and scripts or batch files. The startup media that you choose to use depends mostly on personal preference and your organization's capabilities; however, your startup media must:

- Provide network support if you need to access a distribution share that is on a shared folder on the network, or if you need to access Dynamic Update files on a shared folder on the network.
- Provide CD or DVD device support if you need to access a distribution share on a CD or DVD, or if you need to access disk configuration tools on a CD or DVD.
- Support any tools you need to configure disks before you perform your unattended installation.
- Be able to run the Setup program that your design requires — that is, either Winnt.exe or Winnt32.exe.

You can use the following methods to create startup media:

**Create a TCP/IP boot disk** You can use a Windows NT Server 4.0 operating system CD to create startup media if you need network support. You must create a separate disk for each network adapter. For more information, see the TCP/IP Boot Disk link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Create a network boot disk by using Windows 2000** You can use the Network Client Administrator and a computer running Windows 2000 to create startup media if you need network support. You must have a Windows NT Server 4.0 operating system CD. For more information, see article Q252448, "How to Create an MS-DOS Network Startup Disk in Windows 2000" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Create a network boot disk by adding NDIS drivers to an MS-DOS boot disk** You can use this method if you need network support. You must create a separate disk for each network adapter. You must have an MS-DOS boot disk that was created by using the Network Client Administrator, which is included in the \Clients folder on the Windows NT Server 4.0 operating system CD. For more information, see articles Q142857, "How to Create a Network Installation Boot Disk," and Q128800, "How to Provide Additional NDIS2 Drivers for Network Client 3.0," in the Microsoft Knowledge Base. To find these articles, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Create an MS-DOS boot disk** If you are running Windows XP Professional or Windows Server 2003, you can create an MS-DOS boot disk by using the following procedure: In Windows Explorer, right-click a floppy drive, click **Format** on the shortcut menu, and then select the **Create an MS-DOS startup disk** check box.

**Create a bootable CD or DVD** You can use your writable CD or DVD device to create bootable CDs or DVDs. For more information about creating bootable CDs or DVDs, see the documentation that came with your CD or DVD drive or the documentation for the software that you use to create CDs or DVDs.

# Creating Answer Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An answer file contains the operating system settings that you want to configure during an unattended installation. Before you create an answer file, you need to design your answer file settings. For a worksheet to assist you in recording the settings, see "Answer File Settings Worksheet" (ACIUI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Answer File Settings Worksheet" on the Web at http://www.microsoft.com/reskit).

There are two ways to create an answer file:

- You can use Setup Manager (Setupmgr.exe) to create a minimal answer file that contains critical sections and entries necessary for an unattended installation. Setup Manager is an interactive tool that prompts you for configuration settings, and then builds an answer file based on your responses.

- You can use a text editor, such as Notepad, to manually build an answer file by adding sections and entries to a text file. Using a text editor to build an answer file is often faster and easier than using Setup Manager; however, text editors are more error prone because they do not check answer file syntax.

Typically, you use Setup Manager to build a new answer file, and then use Notepad to manually configure the answer file. You can use the "Answer File Settings Worksheet" (ACIUI_5.doc) to help you configure your answer file.

# Creating an Answer File with Setup Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use Setup Manager to create an answer file, and configure a limited number of answer file settings. Setup Manager is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD. You can run Setup Manager only on Windows XP Professional or on Windows Server 2003.

**To create an answer file by installing and running Setup Manager**

1. From the Deploy.cab file in the \Support\Tools folder of the Windows Server 2003 operating system CD, copy the Setupmgr.exe file to your hard disk.

2. At the command prompt, use the **cd** command to change your current folder to the folder that contains Setupmgr.exe, type **setupmgr**, and then press Enter.

3. Follow the instructions that appear on your screen.

Setup Manager prompts you for the configuration settings listed in Table 2.5. To configure other settings, you must manually edit the answer file as a text file.

**Table 2.5 Answer File Settings That Can Be Configured with Setup Manager**

| Setup Manager Page | Answer File | Section | Entry |
|---|---|---|---|
| User Interaction Page | Unattend.txt | [Unattended] | UnattendMode |
| Name and Organization Page | Unattend.txt | [UserData] | FullName<br>OrgName |
| Display Settings Page | Unattend.txt | [Display] | BitsPerPel<br>XResolution<br>YResolution<br>VRefresh |
| Time Zone Page | Unattend.txt | [GuiUnattended] | TimeZone |
| Product Key Page | Unattend.txt | [UserData] | ProductKey |
| Computer Names Page | Unattend.txt | [UserData] | ComputerName |
| Administrator Password Page | Unattend.txt | [GuiUnattended] | AdminPassword*<br>EncryptedAdminPassword |
| Networking Computers Page | Unattend.txt | [Networking]<br>[NetAdapters]<br>[NetClients]<br>[NetProtocols] | |
| Workgroup or Domain Page | Unattend.txt | [Identification] | |
| Windows Components Page | Unattend.txt | [Components] | |
| Telephony Page | Unattend.txt | [TapiLocation] | |

| | | | |
|---|---|---|---|
| Regional and Language Options Page | Unattend.txt | [RegionalSettings] | |
| Languages Page | Unattend.txt | [RegionalSettings] | Language |
| Browser and Shell Settings Page | Unattend.txt | [FavoritesEx] | |
| Installation Folder Page | Unattend.txt | [Unattended] | TargetPath |
| Install Printers Page | Unattend.txt | [GuiRunOnce] | |
| Run Once Page | Unattend.txt | [GuiRunOnce] | |
| Additional Commands Page | Cmdlines.txt | [Commands] | |

* The value for AdminPassword cannot begin with an asterisk (*). Using a password that begins with an asterisk can cause the password to be set to a null value.

For more information about Setup Manager, see "Using Setup Manager" and "Setup Manager Settings" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Creating an Answer File Manually

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can create or modify an answer file manually by using a text editor, such as Notepad. If you created an answer file by using Setup Manager, the answer file already has several sections and entries. You can add more sections and entries by simply typing the section, entry, and value. However, you can use only valid sections, entries, and values. You cannot create sections, entries or values that are not listed in the "Answer File Settings Worksheet" (ACIUI_5.doc).

Use the following guidelines when you manually create or modify an answer file:

- Sections are always enclosed in square brackets (for example, [Unattended]).
- Sections and entries are not case sensitive.
- Entries must have valid values. For example, "on" is not equivalent to "enable," and "off" is not equivalent to "disable." Invalid sections and entries can generate errors, or cause Setup to fail.
- Each entry must have a value.

In addition, you do not need to add every section or entry to the answer file. Setup ignores missing sections and entries.

# Creating Distribution Shares

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To install Windows XP Professional or Windows Server 2003 on multiple computers over a network, you must create at least one set of distribution shares. The distribution shares typically reside on a server to which each of your destination computers can connect. You can use the same set of distribution shares with different answer files to create different system configurations. Even if you intend to use disk imaging as your installation method, building your master installations with distribution shares provides consistent implementations for a variety of system types. In addition, you can use distribution shares to update future images by editing the files in the distribution shares or by modifying the answer files to generate updated images without having to rebuild each of your master computers.

Before you can create your distribution shares, you need to design them. For a worksheet to assist you in recording the design requirements, see "Distribution Share Worksheet" (ACIUI_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Distribution Share Worksheet" on the Web at http://www.microsoft.com/reskit). You can then use the Distribution Share Worksheet to help you create your distribution shares.

There are two ways to create a distribution share:

- You can use Setup Manager to create a minimal distribution share that contains the system files and device drivers necessary for an unattended installation. Setup Manager is an interactive tool that prompts you for configuration settings, and then builds a distribution share based on your responses.

- You can manually copy files and folders to a distribution share. However, this method is vulnerable to mistakes, and it is recommended that you use it to make only minor changes to your distribution shares.

Typically, you use Setup Manager to create your basic distribution shares. Then, you manually add files and folders to the distribution shares based on your design requirements.

**To create a distribution share by installing and running Setup Manager**

1. From the Deploy.cab file in the \Support\Tools folder of the operating system CD, copy the Setupmgr.exe file to your hard disk.

2. At the command prompt, use the **cd** command to change your current folder to the folder that contains Setupmgr.exe, type **setupmgr**, and then press Enter.

3. Follow the instructions that appear on your screen.

Use the following procedure to manually create a distribution share.

**To manually create a distribution share**

1. Create a shared folder on a server and give it a name that describes the operating system you are installing. For example, if you are creating a distribution share for Windows XP Professional installations, you might name the folder XP_Pro. Be sure each of your destination computers can connect to the shared folder. This folder is your distribution share.

2. Set permissions on the shared folder so that only authorized users can access the folder.

   Authorized users are those users who perform unattended installations in your organization.

3. Copy the contents of the i386 folder on your operating system CD to the i386 folder on your server.

4. In the shared folder, create a folder named $OEM$.

   The $OEM$ folder is the top level folder for all supplemental installation files — such as device drivers, utilities, programs, and scripts — that you want copied to you destination computers.

5. Copy all supplemental files to the $OEM$ folder.

   Be sure to follow the structure shown in Figure 2.5. You can also use your "Distribution Share Worksheet" (ACIUI_2.doc) to help you identify the structure and contents of the $OEM$ folder.

# Performing Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are several ways to perform an unattended installation:

- You can start a destination computer with the operating system CD, and use a Winnt.sif answer file to automate a clean installation of the operating system.
- You can start a destination computer with an MS-DOS startup disk, and then run Winnt.exe with an answer file to perform an automated clean installation.
- You can start a destination computer with a 32-bit operating system — for example, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000, Windows XP Professional, or Windows Server 2003 — and then run Winnt32.exe with an answer file to perform an automated clean installation.
- You can start a destination computer with a 32-bit operating system — for example, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000, Windows XP Professional, or Windows Server 2003 — and then run Winnt32.exe with an answer file to perform an automated upgrade.

To determine which unattended installation method to use, use the information you recorded in the "Unattended Installation Worksheet" (ACIUI_1.doc). This worksheet contains your overall design requirements, such as upgrade and clean installation requirements, startup method, and distribution requirements.

Figure 2.9 shows when you perform an unattended installation.

**Figure 2.9 Performing Unattended Installations**



# Performing a Clean Unattended Installation with an Operating System CD

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Performing a Clean Unattended Installation with an Operating System CD**

Before you perform an unattended installation with an operating system CD, verify that:

- Your answer file is named Winnt.sif, and that it is saved on a floppy disk.
- The BIOS settings in your destination computer list the CD-ROM drive as the first startup device.
- Your answer file contains the following entries, which are required if you perform an unattended installation with an operating system CD:

```
[Data]
MsDosInitiated=0
UnattendedInstall=Yes
[Unattended]
OemPreinstall=No
UnattendSwitch=Yes
```

- Make sure the destination computer is connected to the network (if necessary), and that all peripheral devices, such as printers, scanners, and cameras, are connected to the computer.

After you perform these verification steps, you are ready to perform the unattended installation.

**To perform an unattended installation with an operating system CD**

1. Perform the user state migration tasks that are discussed in your user state migration plan (if you have one).
2. Insert the Windows XP Professional or Windows Server 2003 operating system CD into the CD-ROM drive of the destination computer.
3. Start the destination computer.
4. Press any key when you see "**Press any key to boot from the CD**" appear on your screen. This message appears during text mode stage of Setup. This message does not appear in a dialog box.
5. Insert the floppy disk containing the Winnt.sif file into the floppy disk drive of the destination computer.

After you complete this procedure, Setup starts, reads the Winnt.sif answer file, and installs the operating system with the configuration settings specified in your answer file.

# Performing a Clean Unattended Installation with an MS-DOS Startup Disk

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you perform an unattended installation with an MS-DOS startup disk, verify that:

- The destination computer is connected to the network if you are using a distribution share, and any peripheral devices, such as scanners, printers, and cameras, are connected to the destination computer.
- Your MS-DOS startup disk contains the appropriate device drivers for network connectivity, DVD drives, and any other peripherals that your unattended installation requires. For example, if you are installing from a distribution share on a server, the destination computer needs network connectivity. Likewise, if you are installing from the operating system CD, the destination computer needs to load the drivers for the DVD drive.
- The BIOS settings in your destination computer list the floppy disk drive as the first or second startup device.
- You have the proper permissions to access the distribution share, if you are installing from a distribution share.
- Your answer file is saved on the MS-DOS startup disk or in your distribution share.

After you perform these verification steps, you are ready to perform the unattended installation.

**To perform an unattended installation with an MS-DOS startup disk**

1. Perform the tasks described in your user state migration plan (if you have one).
2. Perform the tasks described in your disk configuration plan (if you have one).
3. Insert the MS-DOS startup disk into the floppy disk drive, and make sure you do not have a CD in the CD-ROM drive.
4. Start the destination computer.
5. If you are installing from a distribution share that is on a CD, or if you are installing from an operating system CD, insert the CD into the CD-ROM drive.
6. At the command prompt, run Winnt.exe with the parameters listed on the "Unattended Installation Worksheet (ACIUI_1.doc).

After you complete this procedure, Setup starts, reads your answer file, and installs the operating system with the configuration settings specified in your answer file.

# Performing a Clean Unattended Installation with a 32-bit Operating System

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you perform an unattended installation with a 32-bit operating system, verify that:

- The destination computer is connected to the network, and that you have the proper permissions to access your distribution share if you are installing from a distribution share.
- Your answer file is saved on a floppy disk or in your distribution share.
- You have the proper permissions to access the distribution share, if you are installing from a distribution share.

After you perform these verification steps, you are ready to perform the unattended installation.

**To perform a clean unattended installation with a 32-bit operating system**

1. Perform the tasks described in your user state migration plan (if you have one).
2. Perform the tasks described in your disk configuration plan (if you have one).
3. Start the destination computer.
4. If you are installing from a distribution share that is on a CD, or if you are installing from an operating system CD, insert the CD into the CD-ROM drive.
5. At the command prompt, run Winnt32.exe with the parameters listed on the "Unattended Installation Worksheet (ACIUI_1.doc).

After you complete this procedure, Setup starts, reads your answer file, and installs the operating system with the configuration settings specified in your answer file.

# Performing an Unattended Upgrade Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you perform an upgrade unattended installation, verify that:

- Your answer file is saved on a floppy disk.
- You have the product CD for the operating system you are installing.
- If you are upgrading from Windows 98 or Windows Millennium Edition to Windows XP Professional, you have the following entries in your answer file:

```
[Unattended]
Win9xUpgrade=Yes
```

- If you are upgrading a server to Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Web Edition, you have the following entries in your answer file:

```
[Unattended]
NtUpgrade=Yes
```

- You do not have the following entries in your answer file. Theses entries cannot be used if you are performing an upgrade unattended installation:

```
[Unattended]
OemPreinstall=Yes
```

**Note**

- Setup only reads the following entries in an answer file during an unattended upgrade installation: Win9xUpgrade, NtUpgrade, OemPreinstall, ProductKey, AutoActivate, DuDisable, DuShare, and DuStopOnError.

After you perform these verification steps, you are ready to perform the unattended installation.

**To perform an unattended upgrade installation**

1. Perform the tasks described in your user state migration plan (if you have one).
2. Start the destination computer.
3. Insert the floppy disk containing your answer file into the floppy drive.
4. Insert the product CD into the CD-ROM drive.
5. At the command prompt, run Winnt32.exe with the parameters listed on the "Unattended Installation Worksheet (ACIUI_1.doc).

After you complete this procedure, Setup starts, reads your answer file, and installs the operating system with the configuration settings specified in your answer file.

# Additional Resources for Designing Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Choosing an Automated Installation Method" in this book for more information about planning Sysprep installations.
- "Designing Image-based Installations with Sysprep" in this book for more information about answer files, distribution shares, and unattended and automated installations.
- "Migrating User State" in this book for more information about migrating user data and settings.
- "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit for more information about folder redirection.
- "Planning for Remote Server Management" in *Planning Server Deployments* of this kit for more information about Emergency Management Services.
- The *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit) for more information about disk partitions and file systems.
- *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm) for more information about using Sysprep. Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.
- The Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Windows Update link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Windows Preinstallation Environment link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about Windows PE and Windows PE licensing plans.
- The Microsoft Platform SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources to download the Core SDK which contains information about configuring a .theme file.
- The TCP/IP Boot Disk link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about creating a TCP/IP boot disk for distributing disk images across a network.
- The Microsoft TechNet articles "User State Migration in Windows XP," "Step-by-Step Guide to Migrating Files and Settings," "Deploying Windows XP Part I: Planning," and "Deploying Windows XP Part II: Implementing" for more information about migrating user data and settings.
- Article Q294895, "Description of the Application Compatibility Toolkit 2.0 for Windows XP," in the Microsoft Knowledge Base.
- Article Q216573, "How Windows Determines ACPI Compatibility," and article Q298898, "How to Determine the Hardware Abstraction Layer (HAL) That Is Used in Windows XP," in the Microsoft Knowledge Base for more information about determining the type of HAL that is installed on a computer.
- Article Q252448, "How to Create an MS-DOS Network Startup Disk in Windows 2000," in the Microsoft Knowledge Base for more information about creating a network boot disk by using a Windows NT Server 4.0 operating system CD.

- Article Q167685, "How to Create an El Torito Bootable CD-ROM," in the Microsoft Knowledge Base for more information about using the El Torito specification to create a bootable CD.

- Articles Q142857, "How to Create a Network Installation Boot Disk," and Q128800, "How to Provide Additional NDIS2 Drivers for Network Client 3.0," in the Microsoft Knowledge Base for more information about creating a network boot disk by adding NDIS and NDIS2 drivers to an MS-DOS boot disk.

**Related Tools**

- Upgrade Advisor

  Use Upgrade Advisor to identify incompatible software and hardware on a destination computer before you perform a clean installation or an upgrade. To download the Upgrade Advisor tools, see the Windows Upgrade Advisor link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. You also can run the Upgrade Advisor tools by using the **/checkupgradeonly** parameter with the Winnt32.exe tool. The Winnt32.exe tool is included in the i386 folder on the Windows XP Professional operating system CD and on the Windows Server 2003 operating system CD.

- User State Migration Tool

  Use the User State Migration tool to save user settings and data before you perform an unattended installation. To download a free version of the User State Migration Tool, see the User State Migration Tool link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Microsoft Baseline Security Analyzer

  Use the Microsoft Baseline Security Analyzer to identify security vulnerabilities that require further configuration after you perform an unattended installation. See article Q320454, "Microsoft Baseline Security Analyzer (MBSA) Version 1.0 Is Available" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Using roaming user profiles" in Help and Support Center for Windows Server 2003, for more information about roaming profiles

- "Disk Management overview" in Help and Support Center for Windows Server 2003, for more information about configuring disks.

- "Best practices for permissions and user rights" in Help and Support Center for Windows Server 2003, for more information about permissions.

- "Start" in Help and Support Center for Windows Server 2003, for more information about using the **start** command to run installation programs in batch files. To find "Start" in Help and Support Center for Windows Server 2003, look under **Support Tasks**, click **Tools**, click **Command-line reference A-Z**, and then click **Start**.

**Related Job Aids**

- "Unattended Installation Worksheet" (ACIUI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Unattended Installation Worksheet" on the Web at http://www.microsoft.com/reskit) for a worksheet to help you design your unattended installations.

- "Distribution Share Worksheet" (ACIUI_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Distribution Share Worksheet" on the Web at http://www.microsoft.com/reskit) for a worksheet to help you record

information about your distribution shares.

- "Renamed Files and Folders Worksheet" (ACIUI_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Renamed Files and Folders Worksheet" on the Web at http://www.microsoft.com/reskit) for a worksheet to help you record information about your renamed files and folders.

- "Dynamic Update Worksheet" (ACIUI_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Dynamic Update Worksheet" on the Web at http://www.microsoft.com/reskit) for a worksheet to help you record information about your Dynamic Update design.

- "Answer File Settings Worksheet" (ACIUI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Answer File Settings Worksheet" on the Web at http://www.microsoft.com/reskit) for a worksheet to help you record your answer file settings.

# Designing Image-based Installations with Sysprep

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the System Preparation tool (Sysprep), included with the Microsoft® Windows® operating system, to perform image-based installations when you want to install identical operating systems and software configurations on multiple computers as quickly as possible. By carefully planning and designing your image-based installation, you can accommodate hardware and software differences among computers, minimize end-user interaction during installation, and reduce the number of images you have to manage.

**In This Chapter**


Overview of Image-based Installations

Identifying Inventory Requirements for Image-based Installations

Defining Disk Images

Designing the Image Delivery Process

Designing Preinstallation Tasks for Image-based Installations

Designing Automated Setup Tasks

Creating Disk Images

Creating Startup Media for a Destination Computer

Deploying Disk Images

Additional Resources for Designing Image-based Installations with Sysprep


**Related Information**

- For more information about planning Sysprep installations, see "Choosing an Automated Installation Method" in this book.
- For more information about unattended installations, see "Designing Unattended Installations" in this book.

# Overview of Image-based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

*Image-based installation* is a method of copying, or cloning, preconfigured operating systems and software applications onto clients and servers. You can perform image-based installations of the Microsoft® Windows® XP Professional operating system and the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Web Edition; and Windows® Server 2003, Enterprise Edition operating systems by using Sysprep and a third-party disk-imaging program.

Image-based installation is a suitable deployment method if you need to:

- Install identical operating systems and software configurations on multiple computers.
- Install an operating system and software configuration as quickly as possible.
- Perform clean installations of an operating system, rather than upgrade an existing installation.
- Minimize end-user interaction and post-installation tasks.
- Install operating systems on computers that have similar hardware and compatible hardware abstraction layer (HAL).

In addition to these deployment solutions, you can customize an image-based installation to accommodate different hardware and software requirements; this allows you to use one disk image to deploy several different hardware and software configurations.

As a deployment solution, image-based installation requires substantial up-front planning and design. This chapter is designed to help IT professionals in medium and large organizations plan and design an image-based installation using Sysprep and a third-party disk-imaging program. It is assumed that you have already designed the client and server configurations that you want to deploy in your organization. This includes designing the configuration of all networking, directory services, and security components. You will use this client and server design information throughout this chapter to customize your image-based installation.

This chapter does not discuss image-based installation from an operations standpoint. In other words, the planning process and design decisions that are discussed in this chapter apply only to corporate deployments and rollouts; they do not apply to ongoing operational tasks such as reinstallation after a hard disk crash or reinstallation due to software or hardware failure.

**Note**

- For a list of the job aids available to assist you in deploying image-based installations with Sysprep, see "Additional Resources for Designing Image-based Installations with Sysprep" later in this chapter.

# Image-based Installation Design Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing an image-based installation involves a design team and a deployment team. The design team is responsible for assessing your current environment to identify inventory requirements, defining disk image requirements, and designing the overall deployment process, including the image delivery process, preinstallation tasks, and automated setup tasks. The deployment team is responsible for implementing all design decisions, including creating disk images and startup media and deploying the disk images. Figure 3.1 shows the process for designing image-based installations.

**Figure 3.1 Designing Image-based Installations**



# Image-based Installation Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To perform an image-based installation, you first set up a *master installation* — a computer with the operating system, software applications, and configuration settings that you want to install onto the destination computers in your organization. Then you run Sysprep, which prepares the master installation so that you can create a *disk image* (that is, a functionally identical replica of its disk) that can be copied onto multiple computers. Next, you use a third-party disk-imaging program to create the disk image of the master installation. Finally, you copy the disk image onto your destination computers.

You need two tools to perform an image-based installation: Sysprep, which can be found on any Windows XP Professional or Windows Server 2003 operating system CD; and a third-party disk-imaging program, which you must purchase from a third-party vendor. The Sysprep tool consists of three separate programs: Sysprep.exe, Setupcl.exe, and Factory.exe. However, you only run Sysprep.exe; Setupcl.exe and Factory.exe are secondary programs that Sysprep.exe runs as needed. To obtain Sysprep, open the Support\Tools folder on any Windows XP Professional or Windows Server 2003 operating system CD, and then open Deploy.cab. For more information about using Sysprep, see the *Microsoft® Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

You run Sysprep on the master computer before you create an image of the master computer's hard disk. Sysprep configures various operating system settings on the master computer to ensure that every copy of the master computer's disk image is unique when you distribute it to a destination computer. Specifically, Sysprep configures a master installation so that unique security identifiers (SIDs) are generated on each destination computer. Sysprep also configures the master computer's disk image so that every destination computer starts in a special setup mode known as *Mini-Setup*. After you copy a disk image onto a destination computer, Mini-Setup runs the first time you start the destination computer.

**Note**

- You can run Sysprep from the command line or from the graphical user interface (GUI). When you run Sysprep from the command line, you can use various command-line parameters to control the way Sysprep runs. When you run Sysprep from the GUI, you use check boxes and command buttons to control the way Sysprep runs. This chapter assumes you are running Sysprep from the command line.

You use the third-party disk-imaging program to create an image of the master computer's hard disk. You also use the disk-imaging program to copy the disk image from the master computer onto a shared folder or a CD, and from the shared folder or CD onto a destination computer.

#### Requirements for Image-based Installation

You can use image-based installation to deploy operating systems and software applications to desktop computers, portable computers, and servers. However, image-based installation is dependent on several factors and can be used only when certain conditions are met. The following are some of the key conditions.

**Clean installation only** You can only use image-based installation to install a clean version of the operating system and clean versions of software applications. You cannot use image-based installation to upgrade an operating system or software configuration.

**Limited server configuration** Some server components must be installed and configured after the image-based installation is complete. These include Certificate Services, Cluster service, and any software that is dependent on the Active Directory® directory service. These also include any application or service that stores the computer name or the computer SID and cannot recover if the computer name or SID changes.

**HAL compatibility** You can only perform an image-based installation if the HAL on the disk image is compatible with the hardware on the destination computer. In some cases, Windows XP Professional and Windows Server 2003 automatically upgrade the HAL that is on a disk image to suit the HAL requirements of a destination computer, but this can only occur if the HAL on the disk image meets several requirements.

**Special domain controller installation process** You cannot deploy preconfigured domain controllers by using image-based installation. However, you can configure a domain controller by first deploying a member server and then automatically running a script that runs Dcpromo.exe, the Active Directory Installation Wizard.

**Limited configuration of some security settings** You cannot use image-based installation to deploy computers that

contain any files that are encrypted using Encrypting File System (EFS). In addition, you cannot use image-based installation to deploy systems that have already been configured with NTFS security settings, such as file and folder permissions, unless the disk-imaging program supports the NTFS file system. You can use a script to configure these settings after the image-based installation is complete.

#### Terms and Definitions

The following key terms are associated with image-based installation and Sysprep.

**Mini-Setup** A wizard that is a subset of Windows Setup. Mini-Setup provides prompts for user-specific information, configures operating system settings, and detects new hardware. You can automate Mini-Setup by using Sysprep.inf.

**Factory mode** A network-enabled state that allows you to perform installation and configuration tasks before you prepare the computer for final delivery to an end user. To use Factory mode, you must make sure Factory.exe is in the same folder as Sysprep.exe and Setupcl.exe. To put a computer into Factory mode, use the **-factory** parameter when you run Sysprep.

**Sysprep.inf** An answer file that you can use to automate Mini-Setup, configure system settings, and perform installation tasks. For example, you can configure Sysprep.inf to automatically set display options, join the computer to a domain, or set telephony options. You can also configure Sysprep.inf to run scripts, programs, or commands after Mini-Setup runs. The Sysprep.inf file must exist on the hard disk of the destination computer (in the *systemdrive*\Sysprep folder).

**Winbom.ini** An answer file that you can use to automate tasks when a computer is started in Factory mode. The Winbom.ini file can exist in one of several locations, including: the hard disk of the destination computer, a removable disk, or a CD.

**Cmdlines.txt** A configurable text file that is used to customize an image-based installation. Cmdlines.txt contains a list of commands that run synchronously after Mini-Setup runs, but before a computer restarts. Cmdlines.txt must exist in the s*ystemdrive*\Sysprep\$OEM$ folder on the destination computer's hard disk, and the path to Cmdlines.txt must be specified by the InstallFilesPath parameter, which is in the [Unattended] section of Sysprep.inf.

**GuiRunOnce** A section in Sysprep.inf that is used to customize an image-based installation. The [GuiRunOnce] section contains a list of commands that run synchronously after a destination computer is started for the first time and a user logs on.

#### New in Sysprep

Sysprep has several new features that are useful for image-based installations in corporate environments. Table 3.1 summarizes the new features.

**Table 3.1 New Sysprep Features**

| Feature | Description |
|---|---|
| Cancel restart support | A Sysprep parameter that prevents a computer from restarting after you run Sysprep. This parameter is mainly used for testing, especially to check if the registry was modified correctly after you perform installation tasks. |
| Countdown timer setting for product activation | A Sysprep parameter that prevents a reset of the countdown timer for product activation. By default, the countdown timer for product activation is reset when you run Sysprep. This parameter is useful if you activate a computer before you deliver it to an end user. This setting is not relevant if you have a volume license. |
| Mass storage support | A Sysprep parameter (**-bmsd**) and an answer file entry (BuildMassStorageSection) that instructs Sysprep to build a list of drivers for mass storage controllers. This prevents you from having to enter device driver information manually in the Sysprep answer file, if an image supports more than one type of mass storage controller. |
| Device driver cleanup support | A Sysprep parameter that clears unused mass storage drivers added by the [SysprepMassStorage] section of Sysprep.inf, and removes placeholder devices created by Plug and Play. |
| Audit support | A Sysprep parameter that lets you verify software and hardware installation without generating new SIDs or processing any items in the Factory mode answer file (Winbom.ini). You can only use audit support with the new Factory mode feature. |

| | |
|---|---|
| Factory mode | A Sysprep parameter that restarts a computer in a network-enabled state without running Mini-Setup. Factory mode is useful for updating drivers, running Plug and Play enumeration, installing applications, testing, configuring the computer with customer data, or making other configuration changes before you deliver a computer to an end user. The Factory mode answer file, Winbom.ini, allows you to automate many installation tasks. |
| Forced shutdown support | A Sysprep parameter that forces a computer to shut down after you run Sysprep. This parameter is useful if a computer has an Advanced Configuration and Power Interface (ACPI) BIOS and it does not shut down properly when you run Sysprep. |
| Reseal support | A Sysprep parameter that clears the Event Viewer logs and prepares the computer for delivery to the customer. Typically, you use the **-reseal** parameter after you perform installation and auditing tasks in Factory mode. |

In addition, the Sysprep answer file (Sysprep.inf) has several changes that affect the way you perform an image-based installation. For more information about the changes in Sysprep.inf, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Identifying Inventory Requirements for Image-based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can plan and design a deployment that uses image-based installation, you need to update your hardware and software inventories to identify hardware and software that can affect the way you perform an image-based installation and verify that all hardware and software is compatible with the new operating system. If you do not take this hardware and software into account while you are planning and designing an image-based installation, the installation might fail.

The hardware and software described in this section must be listed in your hardware and software inventories. If you do not already have hardware and software inventories, you must create them before you can plan and design your image-based installation. For more information about creating a hardware or software inventory, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

Figure 3.2 shows the process that you must follow to ensure that your hardware and software inventories are up-to-date and contain the information you need to design an image-based installation.

**Figure 3.2 Identifying Inventory Requirements**



# Identifying Hardware That Impacts Image-based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Identifying Hardware That Impacts Image-based Installations**

Most Plug and Play peripheral devices, such as sound cards, network adapters, modems, and video cards, do not impact image-based installations. You do not need to inventory these types of devices because they are automatically detected, installed, and configured after you copy a disk image onto a destination computer. You do, however, need to identify several other types of peripheral devices and hardware components, including the following:

- HALs
- Mass storage controllers
- Minimum available hard disk space
- Portable computer devices
- Vendor-specific devices
- Legacy devices

**HALs**

Identify how many HALs your organization uses. You can use image-based installation only if any of the following are true:

- The HAL on the master computer is identical to the HAL on the destination computer.
- The master computer has a uniprocessor or multiprocessor Advanced Programmable Interrupt Controller (APIC) HAL, and the destination computer has a uniprocessor or multiprocessor APIC HAL.
- The master computer has a uniprocessor or multiprocessor Advanced Configuration and Power Interface (ACPI) APIC HAL, and the destination computer has a uniprocessor or multiprocessor ACPI APIC HAL.

Table 3.2 lists the types of HALs that Windows XP Professional and Windows Server 2003 support.

**Table 3.2 HALs Compatible with Windows Server 2003 and Windows XP Professional**

| This HAL | Can Be Used on These Computers |
|---|---|
| Non-ACPI Programmable Interrupt Controller (PIC) HAL (Hal.dll) | • Non-ACPI PIC computers<br>• Non-ACPI APIC uniprocessor and multiprocessor computers<br>• ACPI PIC computers<br>• ACPI APIC uniprocessor and multiprocessor computers |
| Non-ACPI APIC uniprocessor HAL (Halapic.dll) | • Non-ACPI APIC uniprocessor computers<br>• ACPI APIC uniprocessor computers |

| | |
|---|---|
| Non-ACPI APIC multiprocessor HAL (Halmps.dll) | • Non-ACPI APIC multiprocessor computers<br>• Non-ACPI APIC uniprocessor computers |
| ACPI PIC HAL (Halacpi.dll) | • ACPI PIC computers |
| ACPI APIC uniprocessor HAL (Halaacpi.dll) | • ACPI APIC uniprocessor computers |
| ACPI APIC multiprocessor HAL (Halmacpi.dll) | • ACPI multiprocessor computers<br>• ACPI uniprocessor computers |

The type of HAL that is installed on a computer is often dependent on the BIOS. Before you determine the type of HAL a computer needs, make sure that the BIOS is current. For example, a computer might have ACPI-compatible peripherals, but if the BIOS is old and is not ACPI-compatible, the computer could still have a non-ACPI HAL because Setup installs the HAL based on the capabilities of the BIOS. For more information about ACPI-compatible HALs, see article Q216573, "How Windows Determines ACPI Compatibility," in the Microsoft Knowledge Base. For more information about determining the type of HAL that is installed on a computer, see article Q298898, "How to Determine the Hardware Abstraction Layer (HAL) That Is Used in Windows XP," in the Microsoft Knowledge Base. To find these articles, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**To determine the type of HAL that is installed on a computer**

1. In Windows Explorer, open the *Systemroot*\System32 folder.
2. Right-click **Hal.dll**, and then click **Properties** on the shortcut menu.
3. On the **Version** tab, in the **Item name** list, click **Original file name**.
4. Use the file name of the HAL and Table 3.2 to determine the type of HAL that is installed on the computer.

You cannot use the command prompt or the Microsoft® MS-DOS® operating system to determine the type of HAL that is installed on a computer.

For more information about HAL compatibility and image-based installations, see "Reducing the Number of Master Images for Computers with Multiprocessors" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

**Mass storage controllers**

You might need to identify certain types of mass storage controllers that are used in your organization. In the past, you had to create a separate disk image for each mass storage controller. This is no longer true with Windows XP Professional and Windows Server 2003; however, if you have a type of mass storage controller that is not listed in any of the device information (.inf) files that ship with Windows Server 2003 or Windows XP Professional — Machine.inf, Scsi.inf, Pnpscsi.inf, or Mshdc.inf — you need to use the following information when you design automated setup tasks for the Mini-Setup stage of an image-based installation:

- The hard disk controller's description, as specified in its .inf file (for example, Intel 82371AB/EB PCI Bus Master IDE Controller).
- The hard disk controller's Plug and Play ID, as specified in its .inf file (for example, PCI\VEN_8086&DEV_7111).
- The file name of the hard disk controller's .inf file.
- Driver file names for the hard disk controller. This includes the following files: *Driver*.sys, *Driver*.inf, *Driver*.dll, *Driver*.cat, and Txtsetup.oem, where *Driver* is the name of the device driver. Some drivers, such as small computer system interface (SCSI) miniport drivers, might not have a .dll file.

- The name of the tag file (also known as a disk tag), whose presence on a floppy disk or CD tells the driver installation program that the floppy disk or CD containing the device drivers is inserted into the floppy disk drive or CD-ROM drive. The name of the tag file is specified in the hard disk controller's Txtsetup.oem file.

For a worksheet to help you record information about your mass storage controllers, see "Mass Storage Controller Worksheet" (ACISYS_2.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Mass Storage Controller Worksheet" on the Web at http://www.microsoft.com/reskit). For more information about mass storage controller compatibility and image-based installations, see "Reducing the Number of Master Images for Computers with Different Mass Storage Controllers" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

**Minimum available hard disk space**
Identify the smallest hard disk, or the smallest partition, that you plan to distribute each disk image to. This is important from a design standpoint because your disk image must be smaller than the minimum space that is available on a destination computer. By making your disk image smaller than the smallest disk or smallest partition in your organization, you make the image more versatile.

Although most disk-imaging programs can extend or shrink a partition to fit the size of the disk image, using a disk-imaging program to do so is not recommended if the disks are formatted with NTFS. Instead of having a third-party disk-imaging program extend a partition, you can have Windows extend it. This will ensure that NTFS is not compromised. For more information about extending the size of a partition, see "Automating Tasks Before Mini-Setup" later in this chapter.

**Portable computer devices**
Identify any special devices that are installed on portable computers. Some devices are compatible only with portable computers and cannot be installed on desktop computers. For example, if you create a disk image of a portable computer that has an inboard (built-in) pointing device, such as a trackpad, and you distribute the image to a desktop computer, the desktop computer might not have any support for the mouse or keyboard during the Mini-Setup phase of an image-based installation. To prevent this behavior, create a separate disk image for portable computers.

The primary portable devices that you need to identify are:

- DVD, CD-RW, and CD-ROM drives that require vendor-specific or third-party drivers and codecs.
- Human input devices, such as trackpads and track sticks, that require vendor-specific or third-party drivers.
- Inboard or motherboard-resident devices — such as display adapters, network adapters, modems, infrared ports, and sound cards — that require vendor-specific or third-party drivers.

**Vendor-specific devices**
Identify special devices that require vendor-specific device drivers or third-party device drivers that are not available with Windows XP Professional or Windows Server 2003. Examples of these devices include smart card readers, redundant array of independent disks (RAID) controllers, flash disk devices, and IEEE 1394 bus host controllers. You might need to create a separate disk image that contains these device drivers, or you might need to install these devices after you copy a disk image onto a destination computer.

**Legacy devices**
Identify all legacy devices that are installed and used in your organization. Legacy devices are devices that do not support Plug and Play and might require manual installation and configuration after a disk image is copied onto a computer. Legacy devices do not necessarily require you to create separate disk images, but they can force you to alter the way you perform an image-based installation. For example, you can create a disk image for computers that have only Plug and Play devices, but still use that disk image on computers that have non–Plug and Play devices. For those computers that have non–Plug and Play devices, you might have to run a script after Mini-Setup to configure the non–Plug and Play device settings. For more information, see "Automating Tasks After Mini-Setup" later in this chapter.

# Identifying Software That Impacts Image-based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Several types of software can affect the way you perform an image-based installation. Some applications cannot be installed and configured on a disk image; they must be installed after the disk image is copied onto a destination computer. Some applications can only be installed on portable computers and cannot be installed on desktop computers, which can force you to create separate disk images for portable and desktop computers. To determine whether your software impacts your image-based installation, identify which of the following types of software you need to install.

**Core applications**
Identify the core applications that you want installed on every client and server computer. For client computers, this typically includes an office productivity suite, which includes such applications as an e-mail client, a spreadsheet, and a word processor. For server computers, this typically includes tools for maintenance and operations, such as performance-monitoring applications, remote management programs, and backup programs. Virus-detection programs are also core applications because they are usually installed on all computers.

Core applications are typically installed and configured on the disk image. If there are any computers that you do not want to install core applications onto, or any computers require different configuration settings for the core applications, note this in your software inventory.

**Line-of-business applications**
Identify all of the line-of-business applications that are used in your organization, and identify which groups use them. Accounting programs, specialized database programs, and investment modeling programs are examples of line-of-business applications. You might want to create a separate disk image for certain groups if they use line-of-business applications that require substantial configuration or take a long time to install.

**Applications that depend on Active Directory**
Identify all applications that are dependent on Active Directory. These applications cannot be installed and configured on a disk image: you must install and configure these applications after the disk image is copied onto a destination computer. An application is dependent on Active Directory if it uses any data from Active Directory or writes any data to Active Directory when the application is installed or when the application is run. You do not need to identify applications that are built into the operating system, such as snap-ins, optional components, or system tools.

**Third-party tools**
Identify all third-party tools that are specific to a certain computer or hardware configuration. For example, many computer manufacturers have a suite of diagnostic tools that are designed for their specific computers. Likewise, portable computers commonly have a suite of hardware-specific tools that let you configure power options and other features. You might need to install these tools after the disk image is copied onto destination computers, or you might want to create a separate disk image for the computers that require these tools.

**Service packs, hotfixes, and other software updates**
Identify all service packs, hotfixes, and other software updates that are installed in your organization. Be sure you record the revision number and the revision date of the service pack, hotfix, or other software update.

Having this information in your software inventory makes subsequent design steps easier for you to perform. If you do not have a software inventory, or you need to update your software inventory, keep these tips in mind:

- Be sure to allow plenty of time for preparing your software inventory. Customer data shows that you and your administrative staff might spend considerably more time identifying the software that is used in your organization than you estimate.

- Do not rely on end-user feedback to create a software inventory. End users often do not know what programs they use, because some programs do not have a user interface and some programs start without end-user interaction.

- Create a database for your software inventory and keep the database updated. You can use the database to plan, design, and track rollouts of service packs, hotfixes, and other software updates.

# Verifying Software and Hardware Compatibility

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As with any operating system installation, you need to make sure that your software and hardware are compatible with the new version of the operating system. You can use any of the following tools to check hardware and software compatibility.

**Windows Catalog**
Windows Catalog contains a list of software and hardware products that are designed for, or are compatible with, Windows XP Professional. You can search the catalog by manufacturer, product type, product name, or model. If you do not see a product in the Windows Catalog, it does not mean the product will not work with the Microsoft® Windows® XP family of operating systems — check with the product's manufacturer to determine whether the product works with Windows XP. To use the Windows Catalog, see the Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Upgrade Advisor**
Windows XP Upgrade Advisor and Windows Upgrade Advisor are tools that check your system hardware and software to see whether they are ready to be upgraded to Windows XP Professional or Windows Server 2003. Although these tools are designed for use in upgrading to Windows XP Professional or Windows Server 2003, you can use them to identify software and hardware that is not compatible with a clean installation of Windows XP Professional or Windows Server 2003. To download the Upgrade Advisor tools, see the Windows Upgrade Advisor link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. You can also run the Upgrade Advisor tools by using the **/checkupgradeonly** parameter with the Winnt32.exe tool. The Winnt32.exe tool is included in the I386 folder on the Windows XP Professional and Windows Server 2003 operating system CD.

**Application Compatibility Toolkit**
The Application Compatibility Toolkit contains documents and tools to help you diagnose and resolve application compatibility issues. For more information about the Application Compatibility Toolkit, see article Q294895, "Description of the Application Compatibility Toolkit 2.0 for Windows XP," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Defining Disk Images

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A disk image contains the operating system, software applications, and configuration settings that you want to copy onto a group of computers. Most corporate rollouts, regardless of size, require more than one disk image. However, it is a good idea to minimize the number of disk images your organization uses. Creating and maintaining a disk image is time-consuming and requires ongoing maintenance as your organization's hardware and software needs change. Having fewer disk images reduces the total cost of ownership and simplifies the deployment process.

Several factors influence how many disk images you need.

**Operating system versions** You need to create a separate disk image for each version of the operating system you are deploying. For example, if you are deploying Windows XP Professional, Windows Server 2003, Standard Edition, and Windows Server 2003, Web Edition, you will need at least three disk images.

**Hardware** You might have to create additional disk images if the destination computers have different peripheral devices or hardware configurations. For example, you cannot copy a disk image that contains an ACPI HAL onto a computer that requires a non-ACPI HAL. In this case, you have to create separate images for an ACPI HAL and a non-ACPI HAL. Portable computers are another example of hardware that commonly requires a separate disk image.

**Software** You might have to create additional disk images if you are deploying different software configurations and you do not want to install and configure software after the image-based installation is complete. You might also have to create additional disk images if you need to install proprietary line-of-business applications or special tools onto a specific group of computers (for example, portable computers).

**Operating system and software settings** You might want to create additional disk images for computers that require special operating system or software settings. For example, if you want to configure special local policy settings for a group of computers, and you do not want to do this by using a script after the image-based installation is complete, you can create a separate disk image that includes the special policy settings for that group of computers.

In addition to reducing the number of disk images you need to maintain, try to reduce the size of your disk images. This reduces the time it takes to transfer disk images across a network, and it reduces the time it takes to create disk images of your master computers. For more information about reducing the size of your disk images, see "Optimizing Your Images" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

For a worksheet to help you in defining your disk images, see "Disk Image Worksheet" (ACISYS_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Disk Image Worksheet" on the Web at http://www.microsoft.com/reskit). You will need a separate copy of the worksheet for each disk image. You will also need the information from your software and hardware inventories to define your disk images.

Figure 3.3 shows the design steps you need to follow to define your disk images.

**Figure 3.3 Defining Disk Images**


Identify inventory requirements
Define disk images
Design the image delivery process
Design preinstallation tasks
Design automated setup tasks
Create disk images
Create startup media for destination computers
Deploy disk images
Evaluate operating system differences
ACISYS_1.doc
Evaluate hardware differences
ACISYS_1.doc
Evaluate software differences
ACISYS_1.doc
Evaluate operating system and software settings
ACISYS_1.doc

# Evaluating Operating System Differences

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To start defining your images, create a copy of the job aid "Disk Image Worksheet" (ACISYS_1.doc) for each operating system you are deploying. Under "Operating System Installed on This Image," enter the product name, the full version number of the operating system, and any service packs, security updates, or fixes that you want to include in each disk image. If you are testing preliminary installations of the operating systems you plan to deploy — and you have installed all of the service packs, security updates, and fixes on your test computers — you can get a comprehensive listing of this information by running the **systeminfo** command at the command prompt. If you do not know which service packs, security updates, or fixes are available for the operating systems you plan to deploy, you can use Windows Update to determine this information. For more information about Windows Update, see the Windows Update link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Evaluating Hardware Differences

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You might need to create separate disk images if your computers have different hardware configurations. To determine whether you need additional disk images based on your hardware configurations, you need to evaluate the following categories of hardware.

**HALs**
You can only copy a disk image onto a destination computer if the HAL on the disk image is compatible with the HAL on the destination computer. Compatible HALs are those that are identical or those that can be updated during an image-based installation. If the HAL on your disk image is not compatible with the HAL that is required on a destination computer, you need to create a separate disk image for the HAL that is required on the destination computer. Table 3.2 can help you determine which HALs are compatible.

**Caution**

- Copying a disk image to a destination computer that has an incompatible HAL can cause the destination computer to restart continuously, become unresponsive, or generate Stop errors.

Using the information in your hardware inventory, you can determine whether you need more than one disk image for each operating system based on your HAL requirements. For example, if you are installing Windows Server 2003, Standard Edition on 50 computers, 10 of which require a non-ACPI HAL and 40 of which require an ACPI HAL, you need to create two disk images of Windows Server 2003, Standard Edition: one image for the ACPI HAL and one image for the non-ACPI HAL.

Use the following guidelines to determine how many disk images you need based on the HALs used in your organization:

- PIC and APIC HALs are not compatible: You need one disk image for computers that require PIC HALs and one disk image for computers that require APIC HALs.

- ACPI and non-ACPI HALs are not compatible: You need one disk image for computers that require ACPI HALs and one disk image for computers that require non-ACPI HALs.

- Non-ACPI APIC uniprocessor HALs and non-ACPI APIC multiprocessor HALs are compatible: You can use a single disk image for all computers that require these HALs.

- ACPI APIC uniprocessor HALs and ACPI APIC multiprocessor HALs are compatible: You can use a single disk image for all computers that require these HALs.

- ACPI PIC HALs are not compatible with any other HAL: You need a separate disk image for computers that require ACPI PIC HALs.

In addition, you might need to have a separate disk image if any of your computers support hyper-threading. Hyper-threading enables multithreaded software applications to execute threads in parallel within each processor. On a hyper-threading enabled system, Windows XP and the Windows Server 2003 family function as they would on a multi-processor system, even when only a single physical processor is installed.

Windows automatically uses the hyper-threading capabilities of the processor if the following conditions are met:

- The computer hardware supports hyper-threading, and this functionality is enabled in the computer's BIOS.

- Hyper-threading functionality is installed in the computer processor.

- The computer is using an ACPI Uniprocessor HAL.

- Windows detects one or more processors or enabled threads.

When the listed conditions are met, Windows automatically updates the HAL to the ACPI Multiprocessor HAL and installs an additional processor.

When a disk image is copied to a hyper-threading-enabled system, and the HAL is incompatible, the computer might not boot correctly. If the computer does boot correctly, it might not be able to take advantage of hyper-threading technology.

To ensure that the HALs on your disk images are compatible with hyper-threading enabled systems, you must create the master installation on a computer that has one of the following:

- An ACPI hyper-threading enabled HAL.
- An ACPI APIC uniprocessor HAL.
- An ACPI multiprocessor HAL.

If you do not create the master installation on one of these types of computers, Windows is not able to update the system to use multiple processors. A Sysprep image made on an ACPI-compliant multiprocessor computer will run in a multiprocessor configuration even if support for hyper-threading is turned off in the BIOS.

In each copy of the job aid "Disk Image Worksheet " (ACISYS_1.doc), under "Hardware Installed on this Image," enter the type of HAL that will be included on the disk image. You might have to create more copies of the worksheet if you are deploying several different types of HALs with the same operating system and the HALs are incompatible. Be sure to include the file name of the HAL in the worksheet. HAL file names and descriptions are listed in Table 3.2.

**Portable computer devices**
You might also need to create a separate disk image for portable computers. Portable computers often require vendor-specific or hardware-specific device drivers. Frequently, these specialized device drivers are not compatible with desktop computers. For example, if you configure a disk image of a portable computer that has an inboard pointing device, and then copy the disk image onto a desktop computer, the desktop computer might not have mouse or keyboard support.

If the portable computers in your organization require special device drivers, consider creating separate disk images for your portable computers. Having separate disk images for your portable computers prevents device conflicts and ensures that the appropriate device drivers are installed on both your portable and desktop computers.

In each copy of the worksheet under "Hardware Installed on this Image," record the names of the portable devices the disk image supports and the names of the device driver files that are associated with the devices. You might have to create more copies of the worksheet if you decide to create separate disk images for your portable computers.

**Other devices**
You do not need to create separate disk images for legacy (non–Plug and Play) hardware or hardware that requires vendor-specific device drivers; however, you might want to do this if you have a large number of computers that require the same special device drivers and have the same hardware configuration. For example, you might want to create a separate disk image for file servers that have RAID storage devices. Frequently, you have to install vendor-specific drivers for RAID devices, and you have to use a vendor-specific utility to configure them. Having a separate disk image lets you copy a fully configured and optimized system without having to perform any configuration tasks after deployment.

You can use your hardware inventory to identify legacy devices or other devices that require vendor-specific device drivers. In each copy of the worksheet under "Hardware Installed on this Image," record the names of any legacy devices and vendor-specific devices a disk image supports. Also include the names of all files that are associated with the device drivers.

**Mass storage controllers**
You do not need a separate disk image for each mass storage controller. However, you need to make sure that the appropriate device drivers for a mass storage controller are available on a disk image or on a floppy disk. If a mass storage controller is listed in the device information files that ship with Windows Server 2003 and Windows XP Professional — Machine.inf, Scsi.inf, Pnpscsi.inf, or Mshdc.inf — then the device drivers for it will be available on the disk image. If the mass storage controller is not listed in one of these .inf files, you need to record the name of the mass storage controller under "Hardware Installed on this Image." If the mass storage controller is not listed in the "Mass Storage Controller Worksheet" (ACISYS_2.doc), you also need to record the following information for the mass storage controller under "Hardware Installed on this Image": the Plug and Play device ID, the names of the device driver files, the tag file or disk tag, and the name of the .inf file.

# Evaluating Software Differences

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You do not need to create a separate disk image for every computer that has a different software configuration. You can automatically install and configure most software applications after you copy a disk image onto a destination computer (for more information about automatically installing applications at the end of an image-based installation, see "Automating Tasks After Mini-Setup" later in this chapter). This reduces the number of disk images you have to manage, and it makes it easier for you to modify a software configuration as the needs of a group change.

Still, you might want to create a separate disk image based on a specific software configuration when the following conditions exist:

- A group of computers requires software applications or tools that conflict with other software programs. For example, you might want to create a separate disk image for portable computers that require the same vendor-specific programs, such as power-management utilities, DVD codecs, or diagnostic tools.

- A group of computers requires software applications that cannot be automatically installed and configured after the disk image is copied onto a destination computer. For example, you might want to create a separate disk image for Web servers that all run the same suite of third-party data analysis and monitoring applications. This ensures that all of your Web servers have a consistent configuration, and it eliminates the need to manually install and configure third-party applications after you copy a disk image onto a Web server.

- A group of computers requires that the same unique software configuration be installed frequently. For example, you might want to create separate disk images for trade-show kiosks or computers that are used for training purposes, because these computers require frequent reinstallation.

Regardless of whether or not you create separate disk images for specific software configurations, you need to determine which software applications you want to include on a disk image and which software applications you want to install after the disk image is copied onto a destination computer. You can use your software inventory in conjunction with the following design guidelines to define a software configuration for each disk image.

**Identify core applications** Identify the applications that you want installed on every computer, then install and configure these applications on the disk image. Record the names of the programs in each worksheet under "Software Installed on This Image."

**Identify service packs, hotfixes, and patches** It is a good idea to install and configure service packs, hotfixes, and patches on the disk image. Service packs often take a long time to install; it is more efficient to install them on the disk image. Record the names and versions of service packs, hotfixes, and patches in each worksheet under "Operating System Installed on This Image." Keep in mind that if you install any applications after the disk image is copied onto a destination computer, you might have to install service packs, hotfixes, and patches after the disk image is copied onto a destination computer. Record this information in each worksheet under "Software Installed After Disk Copy."

**Identify applications that cannot be installed on a disk image** You must install and configure some applications after you copy a disk image to a destination computer and Mini-Setup has finished running. This includes: programs that depend on Active Directory, such as Message Queuing (also known as MSMQ); special server applications, such as Certificate Services and Cluster service; and special applications that you want installed only on certain computers or certain groups of computers. Record the names of these programs in each copy of the worksheet under "Software Installed After Disk Copy."

Do not include information about configuration settings in the worksheet at this point. You will evaluate and record configuration settings later in this chapter.

**Note**

- The Sysdiff.exe tool is not available for Windows XP Professional or Windows Server 2003. If you used Sysdiff.exe to deploy Microsoft® Windows® 2000 or Windows NT® version 4.0 operating systems, and you need a tool with similar functionality, you need to find another tool. For more information about other resources that are similar to Sysdiff.exe, see article Q298389, "Sysdiff.exe Deployment Tool Is Not Included in Windows XP," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Evaluating Operating System and Software Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You do not usually need to create a separate disk image for each computer or group of computers that has different operating system settings or different software settings. You can configure most operating system and software settings on the master computer before you create the disk image. However, if you need to configure settings that are unique, you can automatically configure the settings during Mini-Setup while the destination computer is being set up, or after Mini-Setup is complete and the destination computer has been restarted.

Examples of operating system and software settings include:

- Local policy settings, such as Group Policy Administrative Template settings.
- Control Panel settings, such as power options, sound scheme settings, system startup and recovery options, system performance settings, and accessibility options.
- Internet Explorer settings, such as the default home page, security and privacy settings, and connection settings.
- Optional Windows components settings, such as network monitoring tools, Remote Storage, and Services for Macintosh.
- Services settings, such as startup type, logon accounts, and recovery actions.
- Desktop settings, such as desktop shortcuts and folder options.
- Microsoft® Word or Microsoft® Excel settings, such as view, edit, save, and spelling options.

You can record operating system settings in each copy of the worksheet under "Operating System Installed on This Image" and under "Operating System Settings Configured After Disk Copy." Likewise, you can record software configuration settings in each worksheet under "Software Installed on This Image" and under "Software Installed After Disk Copy."

Most operating system settings and software settings can be configured on the master computer before you create the disk image. However, you must configure the following settings after the disk image is copied onto a destination computer.

**Domain controller settings** You cannot configure a master computer as a domain controller. You first configure a master computer as a member server, and then configure it as a domain controller after the disk image is copied onto a destination computer. You use Dcpromo.exe (also known as the Install Active Directory Wizard) to configure a server as a domain controller.

**Plug and Play hardware settings** You must configure settings for Plug and Play hardware after a disk image is copied onto a destination computer. For example, you cannot configure power management settings for a specific device, such as enabling wake-on-LAN settings for a network adapter, or configuring link speed or duplex settings for a network adapter. Sysprep configures the master computer so that Plug and Play devices are detected, installed, and reconfigured with default settings the first time you start a computer after a disk image is copied onto it.

**Static IP address settings** You must configure static IP address settings after a disk image is copied onto a destination computer (by using Sysprep.inf). When a disk image is copied onto a destination computer, all of the network adapters on a destination computer are initialized to the default settings, which include dynamic allocation of IP addresses. For more information about how Sysprep affects network settings, see article Q271369, "Statically-Entered TCP/IP Settings Are Not Present After Sysprep," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Encrypting File System settings** You must configure EFS settings on files and folders after a disk image is copied onto a destination computer. If you run Sysprep on a hard disk that contains encrypted files or folders, the data in those files and folders will become completely unreadable and unrecoverable.

**Policy settings** You can configure local Group Policy Administrative Template settings on a master computer. This is because Administrative Template settings are stored in the registry, and Sysprep does not change this part of the registry. Any changes you make to Administrative Template settings will appear on the disk image. You can also configure other local Group Policy settings if the local Group Policy object is not linked to a site, domain, or organizational unit Group Policy object. This is

because Sysprep does not change the local Group Policy object as long as local Group Policy settings are not overridden by site, domain, or organization unit Group Policy settings. If Group Policy settings are not overridden by site, domain, or organizational unit Group Policy settings, any changes you make to local Group Policy settings will appear on the disk image. If the local Group Policy settings are overridden by site, domain, or organizational unit Group Policy settings, and you configure these settings on a master computer, the settings will be overridden on each destination computer.

**Mini-Setup settings** Several types of operating system settings are configured during Mini-Setup. Examples of these settings include: telephony settings, licensing information, computer name, administrative password, and domain membership settings. You cannot configure Mini-Setup settings on the master computer before you create the disk image. For more information about Mini-Setup settings, see "Designing Automated Setup Tasks" later in this chapter.

# Designing the Image Delivery Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

*Image delivery* is the process of creating, managing, and distributing disk images. The image delivery process begins after you configure your master computer, and it ends after you copy a disk image onto a destination computer. To design an effective image delivery process, perform the following tasks:

- Choose a disk-imaging program, which you will use to create and manage disk images.
- Choose an image-distribution method, which you will use to store and transfer disk images to destination computers.

Your disk-imaging program must be compatible with the operating system and file system you are deploying, and your distribution method must be compatible with your organization's networking and hardware capabilities.

Figure 3.4 shows the steps to follow in designing your image delivery process.

**Figure 3.4 Designing the Image Delivery Process**



# Choosing a Disk-Imaging Program

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft does not provide disk-imaging software. You must purchase a third-party disk-imaging program to create a disk image of a master computer's hard disk.

Not all disk-imaging programs are compatible with Windows Server 2003 and Windows XP Professional. When you evaluate disk-imaging programs, make sure you choose a program that supports the following Windows Server 2003 and Windows XP Professional features:

- **Long file names.** Be sure your disk-imaging program supports long file names. (Long file names can be up to 255 characters and can contain spaces, multiple periods, and special characters that are not allowed in MS-DOS file names.) Most commercial third-party disk-imaging programs can handle long file names, but some shareware and freeware disk-imaging programs cannot.

- **NTFS 3.1.** Be sure that your disk-imaging program supports NTFS 3.1, which is the version of NTFS used by Windows Server 2003 and Windows XP Professional. Although many disk-imaging programs support NTFS, these programs do not necessarily support the new features in NTFS 3.1, such as the clean shutdown flag.

In addition to these required features, consider choosing a disk-imaging program that supports the following optional features:

- **Network share support.** Some disk-imaging programs can copy disk images to and from network shares. This feature is essential if you distribute disk images across a network.

- **CDR-RW support.** Some disk-imaging programs can write the disk image directly to a writable CD. This feature is useful if you distribute disk images on CDs.

- **Large-file support (also known as file splitting or disk spanning).** Some disk-imaging programs can copy an image onto multiple CDs or other media. This is useful because a typical disk image of Windows Server 2003 or Windows XP Professional does not fit on one CD.

- **Stand-alone support.** Some disk-imaging programs provide a mechanism for booting a computer that is not connected to a network, and then copying an image from removable media without using a network connection. This is useful if you distribute your disk images on CD or DVD.

- **Multicast image deployment.** Some disk-imaging programs have a multicast server feature that lets you simultaneously copy a disk image onto multiple computers over a network connection. This is useful for large-scale rollouts where you want to automate and control the disk copy process.

- **Image management.** Some disk-imaging programs have image-management features that let you view, add, and remove files and folders from a disk image. This is useful for updating a disk image without having to reconfigure a master computer and create a new disk image.

Some disk-imaging programs can create, resize, or extend a partition before you copy a disk image onto a destination computer. Although these features might be useful, not all disk-imaging programs can perform these tasks: in fact, some programs might cause a STOP 0x7B error (INACESSIBLE_BOOT_DEVICE). If you want to create a partition on a destination computer's hard disk before you perform an image-based installation, you need to be sure the disk-imaging program is compatible with the file systems used by Windows Server 2003 and Windows XP Professional. If you want to resize or extend a partition before you copy a disk image onto a destination computer, use the ExtendOemPartition parameter in the Sysprep.inf file.

For more information about Stop 0x7B errors, see article Q257813, "Using Sysprep May Result in 'Stop 0x7B (Inaccessible Boot Device)' on Some Computers," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. For more information about using the ExtendOemPartition parameter, see "Automating Tasks Before Mini-Setup" later in this chapter.

**Note**

- If you are deploying a 64-bit edition of Windows XP or a 64-bit version of the Windows Server 2003 family, you must use a 64-bit disk-imaging program.

# Choosing an Image Distribution Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

*Image distribution* refers to the way you store a disk image and the way you transfer a disk image to a destination computer. You can distribute disk images two ways:

- You can store images on a network share and then distribute the images across the network to destination computers. This is referred to as *network distribution*.

- You can store images on media such as a CD or a DVD, and then distribute the images from the media to destination computers. This is referred to as *media distribution*.

You need to determine which distribution method to use for each of your disk images. You will likely have to use both distribution methods for some disk images. For example, you might use network distribution to roll out Windows XP Professional in a corporate office that has a high-speed local area network (LAN), but you might also use media distribution to roll out the same image in a branch office that has a slow and unreliable network connection.

# Distributing Disk Images Across a Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To distribute disk images across a network, you need:

- High-speed network connectivity.
- Adequate file server capacity.
- A disk-imaging program that supports network distribution.
- A network boot disk.

Network distribution might also require additional administrative overhead, such as network configuration and troubleshooting, file server configuration and management, and security configuration. For example, you might have to configure network settings or troubleshoot network issues if a destination computer cannot access the network. Likewise, you might have to add storage capacity to your file servers and address performance issues to ensure that your file servers are optimally configured for handling disk images. You might also have to configure permissions, security policies, or user rights on your file servers so that unauthorized users do not download or copy your disk images.

**High-speed network connectivity**
You must have a network connection to every destination computer that you are deploying. Ethernet LANs and Token Ring LANs are well-suited for distributing disk images across a network. Wide area networks (WANs) are generally not fast enough, unless the LAN segments that make up the WAN are connected with a fast T-Carrier service (T2 or higher). Digital subscriber line (DSL), cable modem, Integrated Services Digital Network (ISDN), and dial-up modem connections are not suitable for network distribution.

Table 3.3 shows disk image transfer times based on connection type and network speed. Image transfer times are based on optimum network speeds only and are calculated for a 2.5 gigabyte (GB) disk image. File server performance is not factored into the disk image transfer times. You can use Table 3.3 as a rough guide to help you determine whether your network is a suitable for network distribution.

**Table 3.3 Disk Image Transfer Times Based on Connection Type and Network Speed**

| Connection Type | Network Speed | Transfer Time (2.5 GB Disk Image) |
|---|---|---|
| Fast Ethernet | 100 Megabits per second (Mbps) | 3 minutes, 25 seconds |
| Fast Token Ring | 16 Mbps | 21 minutes, 22 seconds |
| Ethernet | 10 Mbps | 34 minutes, 9 seconds |
| T2 | 6.312 Mbps | 54 minutes, 6 seconds |
| Token Ring | 4 Mbps | 1 hour, 25 minutes |
| T1 | 1.544 Mbps | 3 hours, 41 minutes |

**Adequate file server capacity**
You must have a file server configuration that can handle large file transfers. Several factors determine whether a file server is adequate for large file transfers. The disk type (integrated device electronics [IDE] or SCSI), disk access speed, network adapter settings, disk rotation speed, bus speed, and protocol type can all influence the performance of a file server. Many hard disk manufacturers provide applications that measure your disk performance.

**Disk-imaging program that supports network distribution**
You must have a third-party disk-imaging program that supports network deployment or transfer of disk images. Many disk-imaging programs can copy a disk image directly to a network share. Others can only copy a disk image onto a hard disk on the same computer you are imaging, which means you must manually copy the disk image to a network share. Some

programs also provide network deployment features, such as a multicast feature that you can use to deploy images simultaneously to multiple destination computers, and a subnet selection feature that you can use to distribute images to selected subnets. Although these features are not required for network distribution, they can make an image-based deployment faster and easier.

**Network boot disk**

You must have a network boot disk (floppy or CD) in order to transfer disk images across the network. You use the boot disk to start the destination computer you are deploying and connect the destination computer to a network. Some third-party disk-imaging programs provide a network boot disk (floppy). You can also create one yourself. For more information about creating a network boot disk, see "Creating Startup Media" later in this chapter.

# Distributing Disk Images by Using Media

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To distribute images on media, you need CD or DVD recording hardware, a disk-imaging program that supports media distribution, a file-splitting or disk-spanning program, and a boot disk with CD or DVD support.

Media distribution might also require additional administrative overhead. The most common administrative tasks associated with media distribution include: configuring and troubleshooting CD-ROM drives, maintaining and updating disk images, and managing security. For example, you might have to configure or troubleshoot CD-ROM or DVD drives if your boot disk fails to load the appropriate CD-ROM or DVD device drivers on a destination computer. Likewise, you might have to record new CDs or DVDs (and, for security purposes, destroy old CDs and DVDs) every time you make a change to your disk image. You will also spend administrative time ensuring that the CDs and DVDs are physically secure and not available to unauthorized users.

**CD or DVD recording hardware**
You must have a CD or DVD recorder to distribute images on media. You can use any type of CD or DVD recording device (for example, CD-R or CD-RW). However, you must make sure that the CD-ROM drives in your destination computers can read the media you create.

**Disk-imaging program that supports media distribution**
You must have a disk-imaging program that allows you to copy a disk image directly onto a hard disk, CD, or DVD on the same computer. Some disk-imaging programs do not support stand-alone disk image creation (for example, disk-to-disk or disk-to-CD). This feature is necessary if your master computers are not connected to a network, or you want to create distribution media immediately after you create a disk image.

**File-splitting or disk-spanning program**
Most disk images of Windows XP Professional or Windows Server 2003 will not fit on a single CD, so you will need a file-splitting or disk-spanning tool that splits a disk image into several files. Some disk-imaging programs provide this functionality, but most do not. To find vendors and shareware Web sites that offer file-splitting or disk-spanning programs, search the Web by using the keywords **file splitting** or **disk spanning**.

**Boot disk with CD or DVD support**
You must have a boot disk to start the destination computer. The boot disk can be the CD or DVD that contains the disk image, or it can be a separate floppy disk. The boot disk must also include the device drivers for the CD-ROM or DVD drive that is in the destination computer. Some third-party disk-imaging programs provide a network boot disk (floppy). You can also create one yourself. For more information about creating a network boot disk, see "Creating Startup Media" later in this chapter.

# Comparing Disk Image Distribution Methods

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Each method of distributing disk images has advantages and disadvantages. Table 3.4 summarizes the advantages and disadvantages of each distribution method. You can use Table 3.4 to identify which distribution method is best suited for each of your disk images and your organization.

**Table 3.4 Comparison of Disk Image Distribution Methods**

| Feature | Network Distribution | Media Distribution |
|---|---|---|
| Can be used to deploy disk images in remote offices that do not have fast network connections. | No | Yes |
| Can be used to deploy disk images to computers that do not have network connectivity. | No | Yes |
| Requires a file server with adequate capacity to store disk images. | Yes | No |
| Requires software-based security to prevent unauthorized access to disk images (permissions, user rights). | Yes | No |
| Requires physical security to prevent unauthorized access to disk images (locks on doors, locks on office desks). | No, but it is a good idea. | Yes |
| Accommodates disk images of any size without special file-splitting or disk-spanning software. | Yes | No |
| Requires CD or DVD recording hardware and media. | No | Yes |

You can use the following guidelines to choose a distribution method.

Choose network distribution if all of the following statements are true:

- You are deploying computers that are connected to a fast network (> 4 Mbps).
- You have a file server with sufficient capacity to store all of your disk images.
- You have a disk-imaging program that supports network distribution of disk images.

Choose media distribution if all of the following statements are true:

- You are deploying computers that are connected to a slow network, or you are deploying computers that are not connected to a network.
- You have CD or DVD recording hardware.
- You have a disk-imaging program that supports disk-to-disk or disk-to-CD copying.

# Designing Preinstallation Tasks for Image-based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you design the image delivery process, you need to identify and plan your preinstallation tasks. Preinstallation tasks are performed before you copy a disk image onto a destination computer. You might not need to perform any preinstallation tasks; you will need to perform some preinstallation tasks if you want to:

- Migrate user state before you copy a disk image onto a destination computer. You will likely not need to do this if you use folder redirection and roaming profiles to store user data and settings on a server.

- Change the size or format of the system partition before you copy a disk image onto a destination computer. You do not need to do this if all of your hard disks are already partitioned and formatted the way you want them.

Figure 3.5 shows the steps to follow in designing your preinstallation tasks.

**Figure 3.5 Designing Preinstallation Tasks**



# Creating a User State Migration Plan for Image-based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You will need to create a user state migration plan if any of your destination computers contains any of the following items that you want to restore after installation is complete:

- User data that you want to be available to the end user. User data includes such things as documents, e-mail messages, spreadsheets, and databases.
- User settings such as desktop settings, shortcuts, and Internet Explorer Favorites.
- Application settings such as application-specific keyboard shortcuts, spell-checking options, and default file locations.

At a minimum, your user state migration plan must do the following:

- Identify the data you want to migrate, including user data, user settings, and application settings.
- Determine where to store the data while you perform the image-based installation.
- Create a schedule for migrating data on each of your destination computers.
- Describe how to collect and restore the data.

Microsoft provides two tools for migrating user data and settings. Each tool is designed for different types of users and environments.

- **Files and Settings Transfer Wizard.** Designed for home users and small office users, the wizard is also useful in a corporate network environment for employees who get a new computer and need to migrate their own files and settings without the support of an IT department or Help desk.
- **User State Migration Tool.** Designed for IT administrators who perform large deployments of Windows XP Professional in a corporate environment, the User State Migration Tool provides the same functionality as the wizard, but on a large scale targeted at migrating multiple users. The User State Migration Tool gives administrators the flexibility of a command line approach to customizing specific settings, such as registry entries. To download a free version of the tool, see the User State Migration Tool link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about migrating user data and settings, see "Migrating User State" in this book. For more information about using USMT, see the User State Migration Tool on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Creating a Disk Configuration Plan for Image-based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to create a disk configuration plan if any of the following are true:

- You want to change the size of the system partition on your destination computers before you perform an image-based installation.
- You want to reformat the system partition on your destination computers before you perform an image-based installation.
- You want to create and format extra partitions on your destination computers before you perform an image-based installation.

You do not need to create a disk configuration plan if you configure disk settings after a disk image is copied onto a destination computer. It is relatively easy to delete, create, and format extra partitions — or extend an existing partition — and these tasks do not require substantial analysis and planning. You can automate these types of disk configuration tasks during an image-based installation by configuring parameters in Winbom.ini and Sysprep.inf.

**Note**

- You cannot put the system partition and boot partition on separate partitions if you use Sysprep to prepare a master computer for disk imaging: The system partition and boot partition must be on the same partition.

**Configuring Disk Settings**

There are three ways to configure disk settings before you copy a disk image onto a destination computer.

**Use MS-DOS or Windows 98 disk configuration tools** You can start a destination computer by using a boot disk for MS-DOS or a boot disk for the Microsoft® Windows® 98 operating system, and then use the **fdisk** and **format** commands to partition and format the disk before you copy a disk image onto the destination computer. This method only works if you want to format your disks with the file allocation table (FAT) or FAT32 file system. If you want your hard disks formatted with NTFS, you will have to run the **convert** command by using the [GuiRunOnce] section in Sysprep.inf to convert the FAT or FAT32 file system to NTFS after you have copied the disk image onto the destination computer. For more information about disk partitions and file systems, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Use third-party disk configuration tools** Some third-party disk-imaging programs and disk management programs provide a bootable floppy disk or CD that allows you to partition and format a hard disk before you copy the disk image onto the destination computer. If you use a third-party program to partition or format a disk, be sure that the third-party program creates partitions that are compatible with NTFS 3.1, which is the version of NTFS that is used in the Windows XP Professional and Windows Server 2003 operating systems.

**Use the Windows Preinstallation Environment** You can start a destination computer by using a Windows Preinstallation Environment (Windows PE) CD, and then using the **diskpart** command to partition a disk and the **format** command to format a disk. Windows PE is a bootable operating system that provides limited operating system functionality for performing preinstallation tasks. Windows PE is only available if you have purchased Enterprise Agreement 6.0, Enterprise Subscription Agreement 6.0, or Select License 6.0 with Software Assurance (SA). For more information about Windows PE and Windows PE licensing plans, see the Windows Preinstallation Environment link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. For more information about the **diskpart** command, see "DiskPart Scripting" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

Each method of configuring disk settings has advantages and disadvantages. You need to determine which method is better suited to your organization and your deployment needs.

**Components of a Disk Configuration Plan**

After you determine which method to use to configure disk settings, you need to create your disk configuration plan. At a minimum, your disk configuration plan must identify:

- **Disk configuration settings.** Disk configuration settings include the number of partitions, partition sizes, and file system formats. Disk configuration settings are based on several factors, including disk sizes, disk types, backup capabilities, and user needs. Analyze these factors in your disk configuration plan to determine the right disk configuration for your organization.

- **Procedures for configuring disk settings.** Your disk configuration plan must describe every step of the disk configuration process, including how to start a destination computer and how to run the partitioning or formatting tools.

- **Tools that you use to configure disk settings.** Disk configuration tools include the **format**, **fdisk**, and **diskpart** commands. Your disk configuration plan must describe all the tools you will use to configure disk settings, including the tools you will use to start a destination computer and to partition, format, and check a disk.

**Important**

- You need to rewrite the disk signature if your destination computer's partition is smaller than the new image partition you are copying. Failure to rewrite the disk signature in this case can prevent you from copying a disk image to a disk drive. For example, if you have a target computer with two 5 GB partitions and you need to install an image that is 8 GB, you will need to rewrite the disk signature. You can accomplish this with some disk-imaging programs, but Microsoft does not support doing so. You can also accomplish it with **fdisk -mbr**. The **format** command will not rewrite the disk signature.

# Designing Automated Setup Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can automate the following installation tasks after a disk image is copied onto a destination computer:

- **Mini-Setup.** You can automate Mini-Setup, which is a subset of Windows Setup, by using Sysprep.inf. Mini-Setup runs automatically after you copy a disk image onto a destination computer and an end user starts the destination computer. The primary role of Mini-Setup is to gather user-specific information, and detect and install hardware.

- **Software installation and configuration.** You can automatically install and configure client and server applications. You can also install and configure Windows components.

- **Hardware installation and configuration.** You can automatically update device drivers and configure device settings.

- **Computer configuration.** You can automatically configure computer settings, such as network protocols, display settings, and system services. You can also configure server roles; for example, you can promote servers to domain controllers.

Try to automate as many installation tasks as possible. By automating installation tasks, you can:

- Lower the number of errors caused by technicians, administrators, and end users during your deployment.

- Ensure consistency throughout your organization, which reduces support costs after you complete your deployment.

- Increase productivity by requiring little or no end-user interaction during your deployment.

- Update or modify your installation process without having to educate or retrain end users, technicians, or administrators.

To automate installation tasks, you need to use several types of configuration files, including information (.inf) files, answer (.txt) files, and initialization (.ini) files. Configuration files contain information that is used to answer end-user prompts and configure computer settings before, during, and after Mini-Setup. Configuration files can also contain instructions for running programs, scripts, or commands before, during, and after Mini-Setup. The following configuration files are used to automate installation tasks.

**Sysprep.inf** You can use this file to automate Mini-Setup, configure computer settings, and automatically run programs, scripts, or commands during and after Mini-Setup. When you copy a disk image onto a destination computer, the destination computer automatically searches for Sysprep.inf the first time it is started.

**Cmdlines.txt** You can use this file to automatically run programs, scripts, or commands during Mini-Setup. You must configure the InstallFilesPath parameter in Sysprep.inf to use the Cmdlines.txt file.

**Unattend.txt** You can use this file to automate a program so that no user input is required. For example, an Unattend.txt file can be used with Dcpromo.exe (the Active Directory Installation Wizard) to automatically configure a server as a domain controller.

**Winbom.ini** You can use this file to install software, update device drivers, and configure operating system settings after you copy a disk image onto a destination computer but before you prepare the computer for final delivery to an end user. You can also use this file to run auditing scripts or programs, which can help you identify device errors and verify that your applications and drivers are installed properly.

Figure 3.6 shows the process you need to follow to design an automated image-based installation.

**Figure 3.6 Designing Automated Setup Tasks**


Identify inventory requirements
Define disk images
Design the image delivery process
Design preinstallation tasks
Design automated setup tasks
Automate tasks before Mini-Setup
Automate tasks during Mini-Setup
Automate tasks after Mini-Setup
Create disk images
Create startup media for destination computers
Deploy disk images

# Automating Tasks Before Mini-Setup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can automate some installation and configuration tasks before Mini-Setup runs by using a special Sysprep feature known as *Factory mode*. Factory mode is a network-enabled state that uses an answer file to automate installation and configuration tasks before you prepare the computer for final delivery to an end user. Factory mode is commonly used in a manufacturing environment where every computer requires some customization prior to final delivery to an end user; however, it is also useful for corporate deployments. Use Factory mode if you want to customize individual computers or groups of computers after you have copied a disk image onto a destination computer but before Mini-Setup runs.

To automate installation tasks in Factory mode, you need to configure a Winbom.ini file (short for "Windows bill of materials"). Winbom.ini is the answer file for Factory mode. Using Wimbom.ini is a completely automated process. You prepare a master computer by using the **-factory** parameter with Sysprep, and then create a disk image of the master computer. You then copy the image onto a destination computer. The first time the destination computer starts, it starts in Factory mode and automatically searches for Winbom.ini. The computer then automatically performs the installation and configuration tasks you specified in Winbom.ini. When the destination computer finishes all of the tasks listed in Winbom.ini, you must run Sysprep with the -**reseal** parameter, which prepares the computer for delivery to an end user. To do this, use the Reseal and ResealMode entries in the [Factory] section of Winbom.ini. Factory.exe must be present in the Sysprep\I386\$OEM$ folder with Sysprep.exe and Setupcl.exe in order for Factory mode to work.

**Note**

- You cannot manually install and configure software and system components when a computer is in Factory mode. You must use a Winbom.ini file to install and configure software and system components when a computer is in Factory mode.

For corporate deployments, where you typically do not need to create a unique Winbom.ini file for every computer, you can create Winbom.ini files manually by using a text editor.

Use Factory mode and a Winbom.ini file to automate installation tasks before Mini-Setup if you need:

- **Network connectivity**. You can access resources in shared folders, such as data files, device drivers, and applications.
- **Choice of security context**. You can choose a user account under which to run automated installation tasks.
- **Staged installation support**.You can stage the installation of software, which is the fastest method of installing and configuring applications after you copy a disk image onto a destination computer. For more information about staging the installation of applications, see "Preinstalling Applications" in the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.
- **Windows Installer support**.You can use Windows Installer (.msi) packages to install programs.
- **Disk configuration support**.You can create partitions and format disks on a destination computer, but only if you use Factory mode with the Windows Preinstallation Environment.
- **Synchronous and asynchronous program execution**.You can run programs, scripts, and batch files synchronously by using the [OEMRunOnce] section of Winbom.ini, or asynchronously by using the [OEMRun] section of Winbom.ini. During *synchronous execution*, each program does not run until the previous program finishes running. During *asynchronous execution*, programs start one after the other without waiting for the previous program to finish running.
- **Faster uptime for end users**.You can reduce the number of installation tasks that need to be performed after a computer is delivered to an end user.
- **Auditing support**.You can perform auditing tasks after a destination computer has been started in Factory mode.
- **Device driver installation support**. You can use the [PnPDrivers] and [PnPDriverUpdate] sections of Winbom.ini to

connect to a server and download device drivers. By comparison, when you use the [OEMPnPDriversPath] section of Sysprep.inf, you can only copy device drivers from the local hard disk.

There are some limitations to automating installation tasks before Mini-Setup runs. You cannot use Factory mode and a Winbom.ini file to do the following:

- Perform installation and configuration tasks that publish information in Active Directory.
- Install and configure Cluster service and domain controllers.
- Perform installation and configuration tasks that rely on the computer name or the computer's SIDs.

To design automated installation tasks that occur before Mini-Setup, you must identify:

- The installation tasks you want to automate.
- The configuration files you need to use.
- The settings you need to configure for each configuration file.

**Identifying Automated Installation Tasks You Can Perform Before Mini-Setup**

You can use each copy of the "Disk Image Worksheet" (ACISYS_1.doc) to identify the installation tasks you need to perform after each of your disk images is copied onto a destination computer. If you have not created a worksheet for each of your disk images, identify the installation and configuration tasks that need to be performed after each of your disk images is copied onto a destination computer.

Next, use Table 3.5 to determine which installation tasks to automate before Mini-Setup. Try to automate as many installation tasks as possible.

**Table 3.5 Installation Tasks You Can Automate Before Mini-Setup**

| Installation Task | Comments |
|---|---|
| Install Windows components | These include all Windows components listed in Add or Remove Programs in Control Panel, such as accessories, games, media services, and networking services. |
| Install and configure software | This includes Windows Installer (.msi) packages as well as staged software. Software installation must run in quiet mode, which means the installation must be fully automated and cannot rely on user interaction. Usually, when you run an installation program in quiet mode, you must provide an answer file. |
| Configure computer settings | These include power management and display settings. |
| Run programs, scripts, and batch files | Programs, scripts, and batch files must be fully automated and cannot rely on user interaction, which means you must provide an answer file for any programs, scripts, or batch files you are running, and you must be able to run the programs, scripts, or batch files in quiet mode. |
| Update device drivers | Device driver files are copied onto the destination computer and the device driver location is added to the device path, which is stored in the registry. (The device path tells the Plug and Play module where drivers are stored). The Plug and Play feature installs the drivers the next time the computer starts. |
| Configure shell settings | These include desktop themes, Windows Messenger, and the appearance of the Start menu. |
| Enable an Internet connection firewall | Enables the firewall feature in Windows XP and Windows Server 2003. |

| | |
|---|---|
| Configure the list of most frequently opened or accessed applications | The list of most frequently opened or accessed applications can be configured for specific users. |
| Update registry and file settings | Uses directives based on Windows .inf file processing standards. For information about .inf file sections and directives, see the Driver Development Kits link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. |
| Configure user accounts | This includes user name, .NET Passport, password, and description settings. |
| Update files | This includes renaming, deleting, or copying files in addition to updating the contents of .inf and .ini files. |
| Extend a primary partition | Extends the primary partition that the Windows operating system is installed on. |
| Configure Internet Explorer Enhanced Security Configuration | This includes configuring Internet Explorer Enhanced Security Configuration for members of the User and Guests groups, and the Administrators group.* |

* For more information about Internet Explorer Enhanced Security Configuration settings, see "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003. For more information about answer file settings related to Internet Explorer Enhanced Security Configuration, see the Readme.txt file in Deploy.cab. Deploy.cab is in the Support folder on the Windows Server 2003 operating system CD.

You can perform other automated installation tasks with a Winbom.ini file if you are using Factory mode with the Windows Preinstallation Environment. For more information about these tasks, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

**Identifying Configuration Files to Use Before Mini-Setup**

You might have to configure several answer files or configuration files if you automate installation tasks before Mini-Setup. Table 3.6 describes these answer files and configuration files, and explains where you need to save them.

**Table 3.6 Configuration Files Used to Automate Tasks Before Mini-Setup**

| Configuration File | Description | Where to Save the Configuration File |
|---|---|---|
| Winbom.ini | Answer file for Factory mode. | Any of the following locations:<br>• The root of all removable media drives, including CD-ROM drives and floppy disk drives.<br>• The same location as Factory.exe (usually the *systemdrive*\Sysprep folder).<br>• The root of *systemdrive*. |
| *filename*.txt* | Answer file for programs or scripts that run during Factory mode. This includes answer files for software installation (setup) programs. | Any location you specify in the [OEMRun] or [OEMRunOnce] section of Winbom.ini. |
| *filename*.theme* | Configuration file for desktop themes. | Any location you specify in the [Shell] section of Winbom.ini. |

* Where *filename* can be any valid file name you choose.

For more information about how to configure an answer file for a program or script, see the documentation for the program

or script. For more information about configuring a .theme file, see the Core Software Development Kit (SDK). To download the Core SDK, see the Microsoft Platform SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Identifying Configuration File Settings for Winbom.ini**

Use Table 3.7 to match an installation task with a specific section in a Winbom.ini file.

**Table 3.7 Installation Tasks and Corresponding Winbom.ini Section Names**

| To Automate This Task | Configure These Sections of Winbom.ini |
|---|---|
| Install Windows components | [Components] |
| Install and configure software | [OEMRun], [OEMRunOnce] |
| Configure computer settings | [ComputerSettings] |
| Run programs, scripts, and batch files | [OEMRun], [OEMRunOnce] |
| Update device drivers | [PnPDrivers], [PnPDriversUpdate] |
| Configure shell settings | [Shell] |
| Enable an Internet connection firewall | [SetupHomenet] |
| Configure the list of most frequently used applications | [StartMenuMFUlist] |
| Update registry and file settings | [UpdateSystem] |
| Configure user accounts | [UserAccounts] |
| Update files | [UpdateSystem] |
| Extend a primary partition | [ComputerSettings] |

For more information about specific configuration file settings, including procedural and reference information about creating, formatting, and configuring a Winbom.ini file, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Automating Tasks During Mini-Setup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can automatically update device drivers and configure user preferences and computer settings during Mini-Setup by using a Sysprep.inf file. Sysprep.inf is most commonly used as the answer file for Mini-Setup, but you can use it to perform other automated tasks during Mini-Setup. For example, if you use Sysprep.inf in conjunction with a Cmdlines.txt file, you can install software and run programs, scripts, and batch files just after Mini-Setup is finished but before a destination computer shuts down or restarts.

When you prepare a disk image with Sysprep and then copy the image onto a destination computer, the destination computer automatically searches for Sysprep.inf the first time it starts and automatically performs the installation and configuration tasks you specify in Sysprep.inf. When the destination computer finishes all of the tasks listed in Sysprep.inf, it restarts and is ready for use by an end user (unless there are automated tasks that need to be performed after Mini-Setup runs — in that case, the tasks are performed, the computer restarts, and then it is ready for use by an end-user).

Use a Sysprep.inf file to automate installation tasks during Mini-Setup if you need:

- **Automated Mini-Setup.** You can automate end-user prompts during Mini-Setup, which reduces the amount of end user interaction.
- **Page file regeneration.** You can delete and recreate the page file on the disk image, which is useful if the amount of RAM on the destination computer is different from the amount of RAM that was on the master computer.
- **Mass storage controller support.** You can predefine driver information for mass storage controllers so that Windows loads the correct mass storage controller driver on a destination computer.
- **Cmdlines.txt support.** You can use this file to install programs or run programs, scripts, and batch files when Mini-Setup is finished but before the computer restarts. All tasks listed in Cmdlines.txt run under the Local System security account because there is no logged-on user.

Automating tasks during Mini-Setup has some limitations. You cannot automate tasks during Mini-Setup if:

- You need network connectivity. You cannot access network resources, such as shared folders, during Mini-Setup. This is also true for any tasks you automate by using a Cmdlines.txt file.
- You need to perform a task under a specific user account. All tasks are performed under the Local System security account during Mini-Setup. This is also true for any tasks you automate by using a Cmdlines.txt file.
- You want to install software by using Windows Installer packages. You cannot use the Cmdlines.txt file to install .msi packages.

You can create a Sysprep.inf file either by using Setup Manager or by using a text editor such as Notepad. Setup Manager steps you through the Mini-Setup process and records your answers in a Sysprep.inf file. Setup Manager is included in the Deploy.cab file in the Support folder on the *Windows Server 2003* operating system CD. For more information about using Setup Manager, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

To design automated installation tasks that occur during Mini-Setup, you must identify:

- The installation tasks you want to automate.
- The configuration files you need to use.
- The settings you need to configure for each configuration file.

**Identifying Automated Installation Tasks You Can Perform During Mini-Setup**

You can use each copy of the "Disk Image Worksheet" (ACISYS_1.doc) to identify the installation tasks you need to perform after each of your disk images is copied onto a destination computer. If you have not created a worksheet for each of your disk images, identify the installation and configuration tasks that need to be performed after each of your disk images is copied onto a destination computer.

Next, use Table 3.8 to determine which installation tasks to automate during Mini-Setup. If possible, automate tasks before Mini-Setup by using Factory mode and a Winbom.ini file.

**Table 3.8 Installation Tasks You Can Automate During Mini-Setup**

| Installation Task | Comments |
|---|---|
| Extend a primary partition | Extends the primary partition that the Windows operating system is installed on. This task is performed during the Mini-Setup phase of installation, but just before Mini-Setup runs. |
| Update device drivers and device path | Runs Plug and Play at the end of Mini-Setup to re-enumerate all the installed drivers and install any updated drivers that are found in the device path. The device path tells the Plug and Play module where device drivers are located, and is stored in the registry. You can also change the device path. |
| Regenerate a page file | Deletes and regenerates a page file based on the amount of RAM in the destination computer. This task is performed during the Mini-Setup phase of installation, but just before Mini-Setup runs. |
| Install drivers for mass storage controllers | This task is performed during the Mini-Setup phase of installation, but just before Mini-Setup runs. |
| Configure user settings | These settings include computer name, user name, organization name, and product key. If you automate this task, end users are not prompted for this information during Mini-Setup. |
| Configure regional options | These options include locale and language settings. If you automate this task, end users are not prompted for this information during Mini-Setup. |
| Set date and time | If you automate this task, end users are not prompted for this information during Mini-Setup. |
| Set server licensing mode | If you automate this task, end users are not prompted for this information during Mini-Setup. |
| Configure display settings | These settings include color depth, resolution, and refresh rate. If you automate this task, end users are not prompted for this information during Mini-Setup. |
| Configure telephony settings | These include area code, country code, dial tone type, and number to dial for an outside line. If you automate this task, end users are not prompted for this information during Mini-Setup. |
| Configure computer settings | These include computer name, organizational unit membership, domain or workgroup membership, and administrator password. If you automate this task, end users are not prompted for this information during Mini-Setup. |
| Configure network settings | These include installation of optional networking components and configuration of network services, protocols, and network adapters. If you automate this task, end users are not prompted for this information during Mini-Setup. |
| Install software | If you use Cmdlines.txt to install software, you cannot use .msi packages and you cannot access network resources. In addition, installation tasks are run under the Local System security account because there is no logged-on user. |
| Run programs, scripts, and batch files | You cannot access network resources if programs, scripts, or batch files are run with Cmdlines.txt. In addition, programs, scripts, and batch files are run under the Local System security account because there is no logged-on user. |

#### Identifying Configuration Files to Use During Mini-Setup

You might have to configure several answer files or configuration files if you automate installation tasks during Mini-Setup. Table 3.9 describes these answer files and configuration files, and explains where you need to save them.

**Table 3.9 Configuration Files Used to Automate Tasks During Mini-Setup**

| Configuration File | Description | Where to Save the Configuration File |
|---|---|---|
| Sysprep.inf | Primarily the answer file for Mini-Setup, but can also be used to configure computer settings. | Must be saved in the same location as Sysprep.exe and Setupcl.exe, which are in the *systemdrive*\Sysprep folder on the destination computer. |
| Cmdlines.txt | Configuration file for running programs, scripts, or batch files during Mini-Setup. | Must exist in the *systemdrive*\Sysprep\$OEM$ folder on the destination computer's hard disk. |
| *filename*.txt* | Answer file for programs or scripts that run during Mini-Setup. This includes answer files for software installation (setup) programs. | Same location as the program or script that you need to run without user intervention during Mini-Setup. |

* Where *filename* can be any valid file name you choose.

You can use different Sysprep.inf files with a single disk image by copying Sysprep.inf to the *systemdrive*\Sysprep folder after you copy the disk image to a destination computer. The easiest way to do this is to start the destination computer in Factory mode, and then copy the appropriate Sysprep.inf file from a shared folder to the *systemdrive*\Sysprep folder. You can also use a disk-imaging application to copy files onto a disk image after the disk image has been created, although not all disk-imaging programs provide this functionality.

**Note**

- You can also save Sysprep.inf on a floppy disk with Sysprep.exe and Setupcl.exe; however, the floppy disk controller on the master computer must be identical to the floppy disk controller on the destination computer. If the floppy disk controllers are different, the Setup program will not find the floppy disk controller on the destination computer, and the automated tasks and settings specified in Sysprep.inf will fail.

For more information about using Sysprep.inf and Cmdlines.txt files, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm), which is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD. For more information about how to configure an answer file for a program or script, see the documentation for the program or script.

#### Identifying Configuration File Settings for Sysprep.inf

Use Table 3.10 to match an installation task with a specific section and parameter in a Sysprep.inf file.

**Table 3.10 Installation Tasks and Corresponding Sysprep.inf Section Names**

| To Automate This Task | Configure These Sections in Sysprep.inf |
|---|---|
| Extend a primary partition | [Unattended] |
| Update device drivers and change device path | [Unattended] |
| Regenerate a page file | [Unattended] |
| Install drivers for mass storage controllers | [Sysprep] and [SysprepMassStorage] |
| Configure user settings | [UserData] |
| Configure regional options | [GuiUnattended] and [RegionalSettings] |
| Set date and time | [GuiUnattended] |
| Set server licensing mode | [LicenseFilePrintData] |
| Configure display settings | [Display] |

| | |
|---|---|
| Configure telephony settings | [TapiLocation] |
| Configure computer settings | [Networking] and [Identification] |
| Configure network settings | [Networking] |
| Install software | [GuiRunOnce] |
| Run programs, scripts, and batch files. | [GuiRunOnce] |

**Note**

- You can also use the [GuiRunOnce] section in a Sysprep.inf file to install software and run programs, scripts, and batch files. All tasks that are listed in the [GuiRunOnce] section occur after Mini-Setup finishes and the computer restarts. For more information about the [GuiRunOnce] section, see "Automating Tasks After Mini-Setup" later in this chapter.

# Automating Tasks After Mini-Setup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can automatically install software and run commands, programs, scripts, and batch files after Mini-Setup runs and the computer restarts by using the [GuiRunOnce] section in Sysprep.inf. All of the tasks listed in the [GuiRunOnce] section of a Sysprep.inf file are stored in the following registry subkey:

HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Runonce

When a computer is started for the first time after Mini-Setup, the commands listed in the [GuiRunOnce] section are read from the registry and executed.

Automate tasks by using the [GuiRunOnce] section in Sysprep.inf if you need:

- **Choice of security context**. You can install software and run programs, scripts, and batch files under the context of an end user or a local Administrator account.
- **Network access**. You can access network resources, such as shared folders and drives.
- **Server configuration support**. You can install Cluster service and configure domain controllers.
- **Active Directory support**. You can access Active Directory, depending on the security context of the logged-on user.

On the other hand, automating tasks by using the [GuiRunOnce] section in Sysprep.inf has one key disadvantage: longer startup time for end users. End users must wait for the tasks specified in the [GuiRunOnce] section to run before they can access their computers.

You can configure the [GuiRunOnce] section in a Sysprep.inf file either by using Setup Manager or by using a text editor such as Notepad. Setup Manager prompts you for commands, programs, scripts, or batch files that you want to run after Mini-Setup runs and a computer restarts. Your answers are saved in the [GuiRunOnce] section of a Sysprep.inf file. Setup Manager is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD. For more information about using Setup Manager, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

To design automated installation tasks that occur after Mini-Setup, you must identify:

- The installation tasks you want to automate.
- The configuration files you need to use.
- The settings you need to configure for each configuration file.

#### Identifying Automated Installation Tasks You Can Perform After Mini-Setup

You can use each copy of the "Disk Imaging Worksheet" (ACISYS_1.doc) to identify the installation tasks you need to perform after each of your disk images is copied onto a destination computer. If you have not created a worksheet for each of your disk images, identify the installation and configuration tasks that need to be performed after each of your disk images is copied onto a destination computer.

Next, use Table 3.11 to determine which installation tasks to automate after Mini-Setup. If possible, try to automate tasks before Mini-Setup by using Factory mode and a Winbom.ini file.

**Table 3.11 Installation Tasks You Can Automate After Mini-Setup**

| Installation Task | Comments |
|---|---|
| Install software | Software installation tasks are run under either the end user or local Administrator account. |

| | |
|---|---|
| Run commands, programs, scripts, and batch files | Commands, programs, scripts, and batch files are run under the security context of the end user or the local Administrator account. |

**Identifying Configuration Files to Use After Mini-Setup**

You might have to configure several answer files or configuration files if you automate installation tasks after Mini-Setup. Table 3.12 describes these answer files and configuration files, and explains where you need to save them.

**Table 3.12 Configuration Files Used to Automate Tasks After Mini-Setup**

| Configuration File | Description | Where to Save the Configuration File |
|---|---|---|
| Sysprep.inf | Primarily the answer file for Mini-Setup, but can also be used to run programs, scripts, or batch files after Mini-Setup runs and a destination computer restarts. To do this, you must include the programs, scripts, or batch files in the [GuiRunOnce] section of Sysprep.inf. | Must be saved in the same location as Sysprep.exe and Setupcl.exe, which are in the *systemdrive*\Sysprep folder on the destination computer. |
| *filename*.txt* | Answer file for programs or scripts that run after Mini-Setup and a computer restarts. This includes answer files for software installation (setup) programs. | Same location as the program or script that you need to run without user intervention after Mini-Setup. |

* Where *filename* can be any valid file name you choose.

You can use different Sysprep.inf files with a single disk image by copying Sysprep.inf to the *systemdrive*\Sysprep folder after you copy the disk image to a destination computer. The easiest way to do this is to start the destination computer in Factory mode, and then copy the appropriate Sysprep.inf file from a shared folder to the *systemdrive*\Sysprep folder. You can also use a disk-imaging application to copy files onto a disk image after the disk image has been created, although not all disk-imaging programs provide this functionality.

**Note**

- You can also save Sysprep.inf on a floppy disk with Sysprep.exe and Setupcl.exe; however, the floppy disk controller on the master computer must be identical to the floppy disk controller on the destination computer. If the floppy disk controllers are different, the Setup program will not find the floppy disk controller on the destination computer, and the automated tasks and settings specified in Sysprep.inf will fail.

For more information about using Sysprep.inf and Cmdlines.txt files, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm), which is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD. For more information about how to configure an answer file for a program or script, see the documentation for the program or script.

**Identifying Configuration File Settings to Use After Mini-Setup**

To run programs, scripts, or batch files after Mini-Setup runs and a destination computer restarts, you must use the [GuiRunOnce] section of Sysprep.inf. Programs, scripts, and batch files that are run using the [GuiRunOnce]section run in the context of the currently logged-on end user. If the end user does not have the necessary user rights to run the program, script, or batch file completely, the application fails. Because programs, scripts, and batch files are run in the context of a logged-on end user rather than as a service, the registry entries that the application creates are written for the current end user rather than the default user. (Default user registry settings are propagated to new end users.) If you want any settings and updates to show specifically for the logged-on end user only, using the [GuiRunOnce] section is appropriate. Otherwise, you can use Cmdlines.txt to run applications because it runs programs, scripts, and batch files under the context of the Local System security account.

For more information about configuring the [GuiRunOnce] section in Sysprep.inf, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Creating Disk Images

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Creating disk images is a straightforward process that involves three key tasks:

- Build a master installation for each of your disk images. A master installation is built on a master computer. Building a master installation includes installing and configuring the operating system and any software you want to include on your disk image.
- Prepare each master installation by using Sysprep. This includes configuring and running Sysprep on each master computer.
- Create a disk image of each master installation by using the disk-imaging program. This includes saving each disk image to a permanent storage location.

Figure 3.7 shows the process you must follow to create disk images. None of the steps are optional: you must perform each step in the order shown.

**Figure 3.7 Creating Disk Images**



# Building Master Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To build a master installation, you need to configure disk settings, install the operating system and software, and configure operating system and software settings. You can automate some or all of these tasks, or you can perform them manually on each of your master computers.

**Configuring Disk Settings on a Master Installation**
To configure disk settings on a master computer, you need to perform the following tasks.

**Determine minimum disk size**
Use the "Disk Image Worksheet" (ACISYS_1.doc) or your hardware inventory to determine the minimum available disk space for each disk image. The minimum available disk space for a disk image will be either the smallest hard disk or the smallest system partition among the destination computers.

**Determine partition size and file system format**
Use the job aid "Disk Image Worksheet" (ACISYS_1.doc) to determine the partition size and file system format you want to use on each disk image. The size of the system partition on each master installation must be equal to or less than the minimum hard disk or partition size on your destination computers. If you create a partition on the master installation that is smaller than the minimum hard disk or partition size on your destination computers, you can use the ExtendOemPartition parameter in Sysprep.inf to extend the partition on the destination computer after you copy the disk image onto it. There are several reasons to create a master installation on a partition that is smaller than the minimum hard disk or partition size on your destination computers:

- You reduce installation time if you are distributing disk images across a network.
- You reduce the number of CDs or DVDs you need for each disk image if you are distributing disk images on media.
- You reduce the amount of file server space you need to store your disk images.
- You reduce the amount of time it takes to create disk images.

**Format and partition the hard disk on the master computer**
Use Microsoft tools, such as the **fdisk**, **diskpart**, and **format** commands, to create and format partitions. Also, be sure to create your master installation on drive C.

**Installing and Configuring a Master Installation**
Use the "Disk Image Worksheet" to determine the operating system and software you need to install, and the settings you need to configure, for each of your master installations. For a master installation, you can configure the operating system and software using one of three methods.

**Manual installation and configuration method**
You can manually install the operating system by using Windows Setup, and then manually install software and configure system settings. There are two ways you can do this:

- Start the master computer by using the Windows Server 2003 operating system CD. Windows Setup will start automatically. When you finish installing the operating system, you can then install and configure software applications.
- Start the master computer by using an MS-DOS startup disk, and then start Windows Setup by running Winnt.exe, which is located in the I386 folder on the operating system CD. When you finish installing the operating system, you can then install and configure software applications.

**Semi-automated installation and configuration method**
You can use an answer file to automate Windows Setup, and then manually install and configure software. This automated method of installing the operating system is known as unattended installation. You can perform an unattended installation by starting the master computer with an operating system CD, and then automating Windows Setup by using a Winnt.sif answer file. You can also perform an unattended installation by starting the master computer with an MS-DOS startup disk, and then automating Windows Setup by using an Unattend.txt answer file.

**Fully automated installation and configuration method**

To fully automate the installation and configuration of the operating system and software, you can use answer files in conjunction with configuration sets that reside on a distribution share. A *configuration set* contains device drivers, software files, answer files, and configuration settings that are required to build a master installation. A *distribution share* is a shared folder that contains all of your configuration sets. You can use Setup Manager to create configuration sets.

Table 3.13 compares the three methods of building master installations.

**Table 3.13 How Manual and Automated Methods of Building Master Installations Differ**

| Benefit/Requirement | Manual Method | Semi-automated Method | Automated Method |
|---|---|---|---|
| Stores master installation files in a central location. | No.<br><br>Master installation files, such as device drivers and answer files, are transferred directly from operating system CDs, file shares, and floppy disks to the master computers. | No.<br><br>Master installation files, such as device drivers and answer files, are transferred directly from operating system CDs, file shares, and floppy disks to the master computers. | Yes.<br><br>Master installation files are transferred from operating system CDs, file shares, and floppy disks to a centralized distribution share. |
| Requires a fast network connection between the file server and master computers. | No.<br><br>Network connectivity is only necessary if you need to access installation files that are not on a CD or floppy disk. | No.<br><br>Network connectivity is only necessary if you need to access installation files that are not on a CD or floppy disk. | Yes.<br><br>Network connectivity is necessary to transfer installation files from the distribution share on a file server to a master computer. |
| Simplifies updating and modifying disk images. | No.<br><br>Each master installation must be updated or modified individually before new disk images can be created. | Yes.<br><br>Master installations can be updated by modifying answer files, and then automatically installed on master computers before new disk images are created. | Yes.<br><br>Master installations are updated at a single location and then automatically installed on master computers before new disk images are created. |
| Simplifies testing. | No.<br><br>Errors must be fixed on each master computer. | Yes.<br><br>Errors can be fixed in answer files and then each master computer can be automatically updated. | Yes.<br><br>Errors can be fixed on the distribution share and then each master computer can be automatically updated. |
| Requires record-keeping to track installation and configuration information for each disk image. | Yes.<br><br>You must keep a record of the installation and configuration procedures you performed for each disk image. | Some.<br><br>The answer files track operating system installation and configuration information, but you still need to record information about the software that is installed on each disk image. | No.<br><br>The answer files in each configuration set provide installation and configuration information for each disk image. |
| Complements unattended installations. | No.<br><br>Manually installing and configuring master installations does not affect your unattended installations. | Somewhat.<br><br>You can use the answer files to perform your unattended installations. | Yes.<br><br>The distribution share you use to create your master installations can be used to perform your unattended installations. |

| | | | |
|---|---|---|---|
| Ensures consistency every time you make a change to a master installation. | No.<br><br>Changes are made individually to each master installation by a technician or administrator, which increases the potential for inconsistency and errors. | Somewhat.<br><br>Changes to the operating system are made through answer files, which lessens the potential for errors; however, changes to software are made individually to each master installation by a technician or administrator, which increases the potential for inconsistency and errors. | Yes.<br><br>Changes are made to configuration sets on a centralized distribution share, which lessens the potential for inconsistency and errors. |

You can use the following guidelines to determine which method to use to build your master installations.

Choose manual installation and configuration if any of the following are true:

- You are creating no more than three disk images.
- You seldom upgrade computers and you seldom perform unattended installations.
- You seldom update the configuration of your disk images.
- You have limited network bandwidth in your organization.
- You have limited file server capacity in your organization.

Choose semi-automated or automated installation and configuration if any of the following are true:

- You are creating more than three disk images.
- You frequently upgrade computers or perform unattended installations.
- You frequently change the configuration of your disk images.

For more information about installing and configuring a master installation, and for more information about using Setup Manager to create configuration sets and distribution shares, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD. For more information about answer files, distribution shares, and unattended and automated installations, see "Designing Unattended Installations" in this book.

**Important**

- You must perform all installation and configuration tasks under the local Administrator account on your master installation. This ensures that all configuration settings are stored in the same user profile and that all global configuration settings are stored in the Default User or All Users user profiles.

# Preparing Master Installations by Running Sysprep

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can create your disk images, you must prepare each of your master installations. The preparation process cleans up, configures, audits, and prepares a master installation so an image of its disk can be created and then distributed to destination computers.

To prepare your master installations for imaging:

- Identify the cleanup, configuration, and auditing tasks you need to perform.
- Choose the settings you need to configure when you run Sysprep.

# Identifying Cleanup, Configuration, and Auditing Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You perform cleanup, configuration, and auditing tasks on each of your master installations before you run Sysprep. For a typical image-based installation, you perform the following cleanup, configuration, and auditing tasks in the order in which they are presented. Perform the following tasks under the local Administrator account, except where noted otherwise.

1. Delete files and folders that you do not want end users to see, such as:

   - Files and folders that you used to build the master installation, such as tools, documents, and scripts.
   - Temporary Internet files, including cookies.
   - Temporary user files, which can include items in *systemroot*\Temp, *systemdrive*\Temp, and the folder used by the TEMP environment variable.
   - Files and folders in the Recycle Bin.
   - Files and folders in My Documents.

2. Create a folder for the Sysprep tool and configuration files.

   To do this, create a folder named Sysprep in the root directory of your master installation (for example, C:\Sysprep). Next, extract the Sysprep tool from Deploy.cab, which is located in the Support folder on the Windows Server 2003 operating system CD and the Windows XP Professional operating system CD. The Sysprep tool comprises three files: Sysprep.exe, Setupcl.exe, and Factory.exe. The Sysprep folder is deleted when a destination computer is restarted after Mini-Setup.

3. Remove the master computer from the domain and add it to a workgroup.

   Sysprep cannot completely finish running if the master computer is joined to a domain. If the master computer is joined to a domain, Sysprep will automatically remove it from the domain; however, the preferred method is to remove the master computer from the domain before you run Sysprep. Do not restart the computer when prompted to do so.

4. Run auditing and diagnostic tools.

   Run all auditing and diagnostic tools, such as Disk Defragmenter, Check Disk, and virus detection tools. Be sure to delete any temporary files created by the auditing and diagnostic tools. Also run the Microsoft Baseline Security Analyzer (Mbsa.exe for the graphical user interface version; Mbsacli.exe for the command-line version). For more information about the Microsoft Baseline Security Analyzer, see article Q320454, "Microsoft Baseline Security Analyzer (MBSA) Version 1.0 Is Available," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

   When you finish running auditing and diagnostic tools, restart the computer and log on under the local Administrator account.

5. Perform final cleanup tasks.

   Clear the Event Viewer log files and clear the Administrator account password. You can set the password in the Sysprep.inf file, or you can let end users choose a password; if a password is present when you run Sysprep, every computer will have the same Administrator password and you will not be able to change it during Mini-Setup. Finally, empty the Recycle Bin and clear the Start menu list, which includes command, program, document, and Internet Explorer history.

# Choosing Sysprep Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Sysprep prepares a master installation for disk imaging. You can run Sysprep at the command line or within the Windows graphical user interface. You need to put the Sysprep files in the Sysprep folder on each of your master installations or on a floppy disk.

**Running Sysprep**
You run Sysprep just before you create a disk image of a master installation. This ensures that any changes Sysprep makes are present on the disk image, which in turn ensures that the changes are present on every destination computer onto which you copy the disk image.

When you run Sysprep without specifying any parameters, Sysprep:

- Searches for Sysprep.inf, and, if the file is found, temporarily stores the path to Sysprep.inf in the registry.
- Determines whether a master computer is a member of a domain, and, if it is, removes the master computer from the domain.
- Copies Setupcl.exe to *systemroot*\System32, and then runs Setupcl.exe, which resets SIDs.
- Removes all network adapters (except legacy network adapters), which removes all network settings such as DNS and IP configuration settings.
- Configures the registry so that Mini-Setup runs the next time a destination computer is started.
- Issues a shutdown command so a disk image of the master installation can be created.

You can run Sysprep on a master installation without specifying any parameters if:

- You do not want to perform any auditing or testing after a disk image is copied onto a destination computer and before it is delivered to an end user.
- You are not performing automated installation and configuration tasks by using a Winbom.ini file.
- You do not want to install or configure software, device drivers, or system components after a disk image is copied onto a destination computer and before it is delivered to an end user.
- You do not need to enumerate non–Plug and Play devices the first time a destination computer starts.
- Your master computer shuts down properly after you run Sysprep. Some computers do not shut down after you run Sysprep; if this is the case, you must use the **-forceshutdown** parameter with Sysprep.
- You do not want to reset the grace period for Windows Product Activation, nor clear the critical devices database, nor run Sysprep without generating new SIDs.

If you cannot run Sysprep with its default settings, you need to specify optional parameters. You can use the following guidelines to help you configure Sysprep.

**Using the -activated parameter**
Use the **-activated** parameter if you activate your destination computers in Factory mode. For more information about Windows Product Activation and Sysprep, see article Q299840, "How to Use Sysprep with Windows Product Activation or Volume License Media to Deploy Windows XP," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. The **-activated** parameter is not applicable if you have a volume license.

**Using the -audit parameter**
Use the **-audit** parameter to audit or test a computer in Factory mode. If you use this parameter, you must clear the event logs and delete all files that you created while you were auditing or testing. You cannot use the **-audit** parameter with any other Sysprep parameters.

**Using the -bmsd parameter**
Use the **-bmsd** parameter to populate the [SysprepMassStorage] section of Sysprep.inf with the Plug and Play IDs of mass-storage devices specified in Machine.inf, Scsi.inf, Pnpscsi.inf, and Mshdc.inf. Sysprep only builds the list of mass-storage devices; it does not install these devices in the critical device database or complete any other processing.

You can only use this parameter if the [SysprepMassStorage] section exists in Sysprep.inf, but does not contain any entries. You do not need to add the BuildMassStorageSection parameter to the [Sysprep] section in Sysprep.inf when you use the **-bmsd** parameter. In addition, you cannot use the **-bmsd** parameter with any other Sysprep parameters.

**Using the -clean parameter**
Use the **-clean** parameter to delete device drivers for mass storage controllers that are loaded but not physically present on a computer. You can only use the **-clean** parameter if you used the [SysprepMassStorage] section in Sysprep.inf and the **-bmsd** parameter to load device drivers for mass storage controllers. You typically run Sysprep with the **-clean** parameter in a Cmdlines.txt file. You cannot use the **-clean** parameter with any other Sysprep parameters.

**Using the -factory parameter**
Use the **-factory** parameter to perform installation and configuration tasks — such as installing, configuring, auditing, or testing software and system components — before you prepare a computer for delivery to an end user. You must run Sysprep again on the destination computer when you are finished performing installation and configuration tasks in Factory mode. To do this, use the Reseal and ResealMode entries in the [Factory] section of Winbom.ini.

**Using the -forceshutdown parameter**
Use the **-forceshutdown** parameter if a computer with an ACPI BIOS does not shut down after you run Sysprep.

**Using the -noreboot parameter**
Use the **-noreboot** parameter to test installation and configuration changes in a nonproduction environment. When you run Sysprep with this parameter, Sysprep performs all tasks without shutting down or restarting the computer.

**Using the -nosidgen parameter**
Use the **-nosidgen** parameter if you are not duplicating the computer on which you are running Sysprep.

**Using the -pnp parameter**
Use the **-pnp** parameter only if legacy (non–Plug and Play) hardware is not being detected properly. The **-pnp** parameter can only be used to install legacy hardware, such as COM ports, and cannot be used to install unsigned device drivers. In addition, a destination computer can take up to 20 minutes to start when you use the **-pnp** parameter. This is because the **-pnp** parameter forces a computer to enumerate every device.

**Using the -quiet parameter**
Use the **-quiet** parameter to run Sysprep without displaying onscreen confirmation messages. This is useful if you are automating Sysprep. For example, if you plan to run Sysprep immediately following an unattended Setup, add **sysprep -quiet** to the [GuiRunOnce] section of the Unattend.txt file.

**Using the -reboot parameter**
Use the **-reboot** parameter to force a computer to automatically reboot and then start Mini-Setup, or Factory mode, as specified. This is useful when you want to audit the system and verify that the first-run experience is operating correctly.

**Using the -reseal parameter**
Use the **-reseal** parameter to prepare a destination computer for final delivery to an end user after you have performed installation and configuration tasks in Factory mode. This parameter clears the Event Viewer logs and configures the registry so that Mini-Setup is set to start at the next boot. If you run the command **sysprep -factory**, you must seal the installation as the last step in your preinstallation process, either by running the command **sysprep -reseal** or by clicking the **Reseal** button in the **Sysprep** dialog box.

For more information about Sysprep parameters, see the *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Creating Disk Images of Master Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you run Sysprep on a master installation, you can create the disk image by using a third-party program. Microsoft does not provide a disk-imaging program.

Disk imaging typically involves the following steps:

1. Start the master computer by using a floppy disk, CD, or DVD. The third-party disk-imaging product includes a startup disk or CD that contains the imaging software.

2. Run the third-party disk-imaging program to create an image of the master installation.

3. Save the image in a shared folder, or write the image directly to a CD or DVD.

4. Shut down the master computer.

The disk-imaging process might vary depending on the disk-imaging software you use. Refer to the documentation that came with your disk-imaging software to design your disk-imaging process.

**Caution**

- After you use Sysprep to prepare a master installation, you must not reset SIDs or perform other system preparation tasks by using third-party disk-imaging programs. Doing so after you run Sysprep can damage your master installation and make your disk image unusable. Furthermore, resetting SIDs by using a third-party tool is not supported.

# Creating Startup Media for a Destination Computer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can copy a disk image onto a destination computer, you need to start the destination computer from some type of startup media, such as a floppy disk, CD, or DVD. Some disk-imaging programs provide startup media; some do not. If your disk-imaging program provides startup media, use the media and the instructions that came with it to start your destination computers. If your disk-imaging program does not provide startup media, you need to create the startup media yourself.

Figure 3.8 shows the process you follow to create startup media.

**Figure 3.8 Creating Startup Media for Destination Computers**



# Choosing Startup Media

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Choosing startup media is a multistep process. First, you determine what type of startup media your hardware supports. Not every computer can support CD or DVD startup media. Next, you determine whether one particular type of startup media is more appropriate than another, based on the way you are performing your image-based installations.

**Evaluating Hardware Support for Startup Media**
Follow these steps to determine what type of startup media your organization can support.

1. Evaluate your hardware inventory for floppy disk support.

   To use a floppy disk as startup media, every destination computer must have a floppy disk drive and the boot-order sequence in every BIOS must list the floppy disk drive.

2. Evaluate your CD or DVD writable device.

   To use a CD or DVD as startup media, you must have the proper hardware, software, and instructions to create bootable CDs or DVDs. Microsoft does not provide any tools for creating CD or DVD startup media; however, several manufacturers provide the hardware, software, and system files that you need to create bootable CDs or DVDs.

3. Evaluate your hardware inventory for CD or DVD support.

   To use a CD or DVD as startup media, all of your destination computers must have bootable CD-ROM or DVD drives. Some older CD-ROM drives and many DVD drives are not bootable devices. In addition, the boot-order sequence in the BIOS of each computer must include the CD-ROM or DVD drive. Some older BIOSes do not let you add the CD-ROM or DVD drive to the boot-order sequence.

If your organization supports only floppy disk startup media, you are ready to create your startup media. For more information about creating startup media, see "Creating Startup Media" later in this chapter.

**Identifying Which Startup Media to Use for Image-based Installation**
If you have the proper hardware and software to create CD or DVD startup media, and the destination computers in your organization support CD or DVD startup media, you need to determine which type of startup media is best for your image-based installation. You can use the following guidelines to determine which type of startup media to choose.

Use a floppy disk to start your destination computers if:

- You are distributing disk images across a network.
- You do not need to create partitions or format disks before you copy the disk image onto the destination computer, or your disk configuration tools do not fit on a floppy disk.

Use a CD or DVD to start your destination computers if:

- You are distributing disk images on media (CD or DVD). You can put the disk image on the same media that you use to start the destination computer.
- You are configuring disk settings on destination computers before the disk image is copied onto the destination computer. You can put your disk configuration tools, scripts, and batch files on the same media that you use to start the destination computer.

# Creating Startup Media

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Startup media contains the system files and device drivers that are necessary to start a computer so that the primary hard disk is accessible but not in use. Startup media might also contain network adapter and network drivers, CD and DVD device drivers, disk configuration tools, and scripts or batch files. The method you choose depends mostly on personal preference and your organization's capabilities; however, there are a few guidelines to follow:

- Your startup media must provide network support if you are distributing disk images across a network.
- Your startup media must provide CD or DVD device support if you are distributing disk images on media and you are using a floppy disk as your startup media.
- Your startup media must support the tools you need to copy a disk image from a storage location onto a destination computer. For example, if your startup media is an MS-DOS boot disk, you need to use MS-DOS tools to copy the disk image onto the destination computer.

You can use the following methods to create startup media.

**Create a TCP/IP boot disk** You can use an operating system CD for the Microsoft® Windows NT® Server version 4.0 operating system to create startup media if you are distributing disk images across a network. You must create a separate disk for each network adapter. For more information, see the TCP/IP Boot Disk link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Create a network boot disk by using Windows 2000** You can use the Network Client Administrator and a computer running the Windows 2000 operating system to create startup media if you are distributing disk images across a network. You must have a Windows NT Server 4.0 operating system CD. For more information, see article Q252448, "How to Create an MS-DOS Network Startup Disk in Windows 2000," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Create a network boot disk by adding NDIS drivers to an MS-DOS boot disk** You can use this method if you are distributing disk images across a network. You must create a separate disk for each network adapter. You must have an MS-DOS boot disk that was created by using the Network Client Administrator, which is included in the \Clients folder on the Windows NT Server 4.0 operating system CD. For more information, see articles Q142857, "How to Create a Network Installation Boot Disk," and Q128800, "How to Provide Additional NDIS2 Drivers for Network Client 3.0," in the Microsoft Knowledge Base. To find these articles, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Create an MS-DOS boot disk** You can create an MS-DOS boot disk by right-clicking a floppy disk drive in Windows Explorer, clicking **Format** on the shortcut menu, and then selecting the **Create an MS-DOS startup disk** check box.

**Create a bootable CD or DVD** You can use your writable CD or DVD device to create bootable CDs or DVDs. You can also create a bootable CD according to the El Torito specification. For more information about using the El Torito specification to create a bootable CD, see article Q167685, "How to Create an El Torito Bootable CD-ROM," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Deploying Disk Images

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The final step in your imaged-based installation, shown in Figure 3.9, is to deploy disk images. To deploy disk images, implement the procedures or tasks described in your user state migration plan (if you have one). Then, start the destination computer by using your startup media, and implement the procedures or tasks described in your disk configuration plan (if you have one). When you are finished implementing your disk configuration plan, use your disk-imaging program to copy the disk image to the destination computer, and then shut down the destination computer. Finally, if you are not using Factory mode to configure the destination computer, deliver the destination computer to the end user. If you are using Factory mode, start the computer, let Factory mode perform the configuration tasks specified in Winbom.ini, and then deliver the computer to the end user.

**Figure 3.9 Deploying Disk Images**



# Additional Resources for Designing Image-based Installations with Sysprep

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Choosing an Automated Installation Method" in this book for more information about planning Sysprep installations.
- "Designing Unattended Installations" in this book for more information about answer files, distribution shares, and unattended and automated installations.
- "Migrating User State" in this book for more information about migrating user data and settings.
- The *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit) for more information about disk partitions and file systems.
- *Microsoft Windows Corporate Deployment Tools User's Guide* (Deploy.chm) for more information about using Sysprep. Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.
- The Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about using the Windows Catalog.
- The Windows Update link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about Windows Update.
- The Windows Preinstallation Environment link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about Windows PE and Windows PE licensing plans.
- The Microsoft Platform SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources to download the Core SDK, which contains information about configuring a .theme file.
- The TCP/IP Boot Disk link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about creating a TCP/IP boot disk for distributing disk images across a network.
- Article Q216573, "How Windows Determines ACPI Compatibility," and article Q298898, "How to Determine the Hardware Abstraction Layer (HAL) That Is Used in Windows XP," in the Microsoft Knowledge Base for more information about determining the type of HAL that is installed on a computer. To find these articles, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- Article Q294895, "Description of the Application Compatibility Toolkit 2.0 for Windows XP," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- Article Q298389, "Sysdiff.exe Deployment Tool Is Not Included in Windows XP," in the Microsoft Knowledge Base for more information about other resources that are similar to Sysdiff.exe. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- Article Q271369, "Statically-Entered TCP/IP Settings Are Not Present After Sysprep," in the Microsoft Knowledge Base for more information about how Sysprep affects network settings. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Article Q257813, "Using Sysprep May Result in 'Stop 0x7B (Inaccessible Boot Device)' on Some Computers," in the Microsoft Knowledge Base for more information about Stop 0x7B errors. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Article Q299840, "How to Use Sysprep with Windows Product Activation or Volume License Media to Deploy Windows XP," in the Microsoft Knowledge Base for more information about Windows Product Activation and Sysprep. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Article Q252448, "How to Create an MS-DOS Network Startup Disk in Windows 2000," in the Microsoft Knowledge Base for more information about creating a network boot disk by using a Windows NT Server 4.0 operating system CD. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Article Q167685, "How to Create an El Torito Bootable CD-ROM," in the Microsoft Knowledge Base for more information about using the El Torito specification to create a bootable CD. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Articles Q142857, "How to Create a Network Installation Boot Disk," and Q128800, "How to Provide Additional NDIS2 Drivers for Network Client 3.0," in the Microsoft Knowledge Base for more information about creating a network boot disk by adding NDIS and NDIS2 drivers to an MS-DOS boot disk. To find these articles, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Tools**

- Sysprep.exe, Setupcl.exe, and Factory.exe

  Use Sysprep.exe, Setupcl.exe, and Factory.exe to prepare a hard disk for disk imaging. To obtain Sysprep, open the Support\Tools folder on any Windows XP Professional or Windows Server 2003 operating system CD, and then open Deploy.cab. You can also see the Windows Downloads link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Windows Upgrade Advisor

  Use Windows Upgrade Advisor to identify incompatible software and hardware on a destination computer before you perform an image-based installation. To download the Upgrade Advisor tools, see the Windows Upgrade Advisor link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. You can also run the Windows Upgrade Advisor tools by using the **/checkupgradeonly** parameter with the Winnt32.exe tool. The Winnt32.exe tool is included in the I386 folder on the Windows XP Professional and Windows Server 2003 operating system CD.

- User State Migration Tool

  Use the User State Migration tool to save user settings and data before you perform an image-based installation. To download a free version of the User State Migration tool, see the User State Migration Tool link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Microsoft Baseline Security Analyzer

  Use the Microsoft Baseline Security Analyzer to identify security vulnerabilities that require further configuration after you perform an image-based installation. See article Q320454, "Microsoft Baseline Security Analyzer (MBSA) Version 1.0 Is Available," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Job Aids**

- "Disk Image Worksheet" (ACISYS_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Disk

Image Worksheet" on the Web at http://www.microsoft.com/reskit) for a worksheet to help you define your disk images.

- "Mass Storage Controller Worksheet" (ACISYS_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Mass Storage Controller Worksheet" on the Web at http://www.microsoft.com/reskit) for a worksheet to help you record information about your mass storage controllers.

# Designing RIS Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use Remote Installation Services (RIS) to simplify the installation of operating systems on client computers in your organization by performing remote image- and CD-based installations on client computers with no operating system or on failed computers that need restoration of the operating system. RIS enables you to create, maintain, and quickly install identical operating system and software configurations on multiple remote client computers with a predefined level of end-user interaction.

**In this Chapter**


Overview of the RIS Deployment Process

Planning RIS Installations

Designing RIS-based Installations

Configuring and Deploying RIS

Deploying an Operating System

Additional Resources for Designing RIS Installations


**Related Information**

- For information about Sysprep installations, see "Designing Image-based Installations with Sysprep" in this book.
- For information about unattended installations, see "Designing Unattended Installations" in this book.

# Overview of the RIS Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

RIS enables you to support on-demand image-based or script-based clean operating system installations over a network connection from a RIS server to a RIS client. RIS is included in Microsoft® Windows® Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition operating systems. RIS allows you to standardize client operating system installations, control the end-user installation experience, and choose the software distribution media you use. RIS supports large-scale deployments of Windows Server 2003 and Microsoft® Windows® XP Professional and can also serve as an operations and recovery tool.

RIS uses Pre-boot eXecution Environment (PXE) technology to enable client computers without an operating system to boot remotely to a RIS server that performs installation of a supported operating system over a TCP/IP network connection. You can deploy one or more RIS servers to accommodate your client operating system needs, but each client must have compatible hardware, which includes a BIOS and network adapter that support the remote-boot process. You can create different sets of RIS images to accommodate various configurations of different groups of client computers. You can also use Group Policy settings to limit the installation options that RIS presents to clients. You can use RIS to provide interactive operating system installations that require user input, or fully-automated installations that require no user input other than logon credentials.

The process steps described in this chapter can help IT professionals in medium and large organizations successfully plan, design, and execute a RIS-based operating system deployment. Job aids that are available to assist you in deploying RIS are listed in "Additional Resources for Designing RIS Installations" later in this chapter.

It is recommended that you use RIS when you have a large number of clients that need clean installations of an operating system and when you have an idea of the software configurations you want to deploy in your organization. To deploy RIS, your network infrastructure must be able to support RIS-based installations. Also, DNS and Dynamic Host Configuration Protocol (DHCP) servers must be running on the network and you must have the Active Directory® directory service installed. In addition, you must be familiar with the Windows XP Professional and Windows Server 2003 setup process.

Once you deploy RIS, you can automate large-scale operating system installations, customized for various predefined client and server configurations, in either the interactive or fully-automated mode. You can deliver standard desktops to clients using predefined master images, including customized application installation and configuration. You can also provide installations based on CD-type images from a distribution share. In addition, you can secure client installations by setting user access rights and configuring Group Policy settings that control user installation options.

# Process for Deploying RIS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The RIS deployment process consists of three phases including planning, design, and configuration and deployment. You will need a design team to handle the planning and design phases and a deployment team for the configuration and deployment phase.

Figure 4.1 illustrates the process steps for deploying RIS in your organization, followed by the deployment of an operating system.

**Figure 4.1 Deploying RIS**




**RIS Deployment Teams**

Design team personnel might consist of high-level system architects who make planning decisions and designers who are upper-tier administrators that choose the appropriate methods for carrying out the planning decisions. Deployment team personnel might consist of lower-tier administrators and technicians who implement the design decisions.

Some of the responsibilities of the design team include tasks such as evaluating the current environment, assessing RIS network security, and evaluating image requirements, in addition to designing the overall installation configuration, the supporting infrastructures, the RIS server configuration, and a RIS test environment.

Some of the responsibilities of the deployment team include tasks such as configuring the network infrastructure to support RIS, creating and configuring RIS servers, creating images, customizing answer files, creating client boot disks, and designing the Client Installation Wizard (CIW) process. The deployment team also deploys the operating systems.

**RIS Technology Background**

RIS was introduced in Microsoft® Windows® 2000 to allow server-based installation of an operating system onto client computers that do not currently contain one. Improvements to RIS in the Windows Server 2003 family are summarized in the following section.

**New in Windows Server 2003**

With the release of Windows Server 2003, RIS now supports the following new capabilities:

- Deployment of Microsoft® Windows® 2000 Professional, Microsoft® Windows® 2000 Server, Microsoft® Windows® 2000 Advanced Server, Windows XP Professional, and the Windows Server 2003 family operating systems.
- Automation of the CIW using the Autoenter feature.
- Enhanced cross-domain functionality.
- Increased security by adding a masked double-prompt administrator password.
- Automatic DHCP authorization with Risetup.exe.
- Auto-detection of the target system Hardware Abstraction Layer (HAL) type to allow filtering of images from the CIW.
- Support for the Recovery Console and support for Microsoft® Windows® Preinstallation Environment.

- Support for Microsoft® Windows® XP 64-Bit Edition Version 2003 and the 64-bit versions of the Windows Server 2003 family.
- Support for the Uniqueness Database in .sif files.
- Support for Secure Domain Join.
- Support for NTLM version 2 (NTLMv2)
- Support for encrytped local administrator password entries.

**RIS Components**

RIS consists of several components that facilitate the remote installation of client operating systems. To create the RIS server configuration, you must install the **Remote Installation Services** Windows component from **Add or Remove Programs** in **Control Panel**. This component configures and starts the following services:

**Remote Installation Services (Binlsvc)** This service detects PXE-initiated DHCP requests from RIS clients and facilitates a response to those requests. Remote Installation also directs clients to files on the RIS server that initiate the installation process and then services CIW requests. In addition, Remote Installation checks Active Directory to verify client credentials, determines if a client can be serviced, and confirms whether to create a new computer account object or reset an existing account on behalf of the client. Also, if a client that is prestaged in Active Directory has settings specifying that a particular RIS server must answer the client, then Remote Installation facilitates the response to that client from the specified RIS server.

**Note**

- Remote Installation Services was formerly known as the Boot Information Negotiation Layer (BINL) service in Microsoft® Windows® 2000 and Windows XP Professional.

**Trivial File Transfer Protocol Daemon (TFTPD)** A RIS server uses TFTP to download the CIW and the initial files needed to start the remote installation process on the client computer.

**Note**

- On a RIS server, TFTP is called a *daemon* or service (TFTPD) while on the client side it is referred to as a *protocol* (TFTP).

The first file that downloads is Startrom.com, which is a small bootstrap program that displays the **Press F12 for Network Boot** prompt to the client. If the user presses F12 within 3 seconds, the CIW downloads to the client so the installation process can begin. The file Startrom.com is located on your RIS server in the directory path *ServerName*\RemoteInstall\OSChooser\i386\.

**Note**

- For installations of Windows XP 64-Bit Edition Version 2003, the first file downloaded is Oschoice.efi It is not necessary to press F12 for these installations.

**Single Instance Store (SIS) Service** SIS consists of an NTFS file system filter driver and a groveler agent that interacts with RIS images. The SIS service reduces the hard disk storage requirements for RIS images. SIS does this by monitoring the RIS server partition for duplicate files. Whenever the groveler agent finds a duplicate file, SIS copies the original file into a directory and an NTFS reparse point containing the current location, size, and attributes of the original file. This way, SIS retains only a single instance of the file while replacing duplicate files with links to the single instance. This enables SIS to store the duplicate files it finds in RIS images and reduce disk space usage on your RIS server.

**Caution**

- When backing up a RIS server, you must use an SIS-aware backup solution. Failure to use an SIS-aware backup solution while backing up a RIS server consumes unnecessary disk space while performing a restore operation and might result in some of your files not being restored. The backup program included in Windows Server 2003 is SIS-aware.

**Remote Boot and Installation Setup Processes**

RIS uses PXE technology to allow RIS clients without an operating system to initiate the boot sequence from their network adapters, thus facilitating operating system installations from remote network locations. To initiate the remote boot process and set up a RIS-based operating system installation, PXE interacts with the Dynamic Host Configuration Protocol (DHCP), the Remote Installation services, and the TFTPD, as shown in Figure 4.2.

**Figure 4.2 RIS Installation Configuration**




When you start a new PXE-enabled RIS client computer, the following sequence of events occurs:

1. The client computer initiates the communication by sending a DHCP Discover broadcast on its subnet. A DHCP server with an active scope for that subnet will issue an IP address to the client.

2. All Remote Installation servers that receive the client's DHCP Discover broadcast extract (from the PXE data portion of the packet) the UUID of the client that is requesting service. The RIS server then queries its preferred domain controller to search for this UUID in all prestaged computer accounts in Active Directory.

   If the domain controller does not find the UUID in the local domain, the RIS server queries the global catalog to locate the client computer account. If the UUID is found in either location, the client computer is recognized as a known client; otherwise, it is an unknown client. If the client is unknown, it will only receive an answer from a RIS server that is configured to answer unknown clients, provided one exists on the network.

3. If the client is known, all available RIS servers query the domain to determine whether the prestaged client computer account has a setting that specifies that only a particular RIS server can answer the client.

   If this is the case, then only the designated RIS server answers the service request, and other RIS servers simply notify the client of the particular RIS server configured to answer it. If the client computer account does not have a setting that requires it to be answered by a particular server, any RIS server can answer the request. However, the client only receives service from the first RIS server it contacts.

4. The user receives a prompt to press the F12 key to initiate a network service boot request from the RIS server.

5. Using the TFTP daemon (service), the contacted RIS server downloads the CIW to the RIS client, along with all client dialog boxes contained within the CIW.

6. The CIW prompts the user to log on with a valid user name, password, and domain name.

7. The user receives an offering of operating system images hosted on the RIS server for installation on the client computer.

   The list of operating system images offered to the user is based on the user's credentials or security group membership.

**PXE Specifications**

The published PXE specification defines the remote boot process and also establishes the PXE compliance standards for

hardware manufacturers and other vendors. RIS uses PXE environment extensions to DHCP, an industry-supported technology, to allow workstations to do the following:

- Boot remotely using their network adapters to access bootstrap code from a network location.
- Install an operating system from a remote source to a client's local hard disk.

The PXE environment is built upon Internet protocols and services that are widely used in the computer industry. This includes TCP/IP, DHCP, and TFTP. The PXE extensions to the DHCP protocol enable information to be sent to network-bootable systems and also allow these systems to locate remote installation services. For more information about PXE and the protocols required to support network booting, see the Preboot Execution Environment (PXE) Specification link on the Web Resources page at: http://www.microsoft.com/windows/reskits/webresources.

**Note**

- Network adapters that meet the PXE .99n specification should work correctly with RIS.

#### RIS Components

The following RIS components enable you to install, configure, and implement RIS in your organization.

**Remote Installation Services** An optional Windows component that you can install with Windows Server 2003 or you can add it at any time after the operating system installation. Services that install with RIS include Remote Installation, TFTPD, and the SIS Groveler.

**Remote Installation Services Setup (Risetup.exe)** You use this component to initially set up the RIS server and create at least one CD-based operating system image. You can initiate the setup process from the **Start** menu of your RIS server. By selecting **Remote Installation Services Setup** from the **Administrative Tools** group, a wizard starts and does the following:

- Requests preliminary information, including the installation folder name and the path to the operating system installation files.
- Copies Windows installation files.
- Updates the CIW screens.
- Creates a default answer file (Ristndrd.sif).
- Starts RIS services.
- Authorizes the DHCP server.

**Note**

- Risetup is also used to create any additional CD-based operating system images after the initial installation is created.

**Remote Installation Preparation Wizard (Riprep.exe)** Riprep.exe allows you to create a customized image of an operating system such as Windows XP Professional. To create an image means that you create a replica of a hard disk that you can install on other computers in your organization. You use Riprep to image an existing Windows XP Professional operating system installation on a master computer and replicate that image to an available RIS server on your network. The image can include the operating system with default parameters applied, or the operating system with a preconfigured desktop, locally-installed applications, and drivers.

**Remote Boot Floppy Generator (Rbfg.exe)** Rbfg.exe allows you to create remote boot floppy disks for some RIS clients that are not PXE-enabled, so that these clients can emulate the remote boot process and install an operating system over the network using RIS. However, for non PXE-enabled RIS clients to use the remote boot floppy disk, they must each have a supported Peripheral Component Interconnect (PCI) network adapter.

**Client Installation Wizard (OSChooser)** The OSChooser is the client-side service of the CIW. It is a text-based program downloaded by the RIS server that allows the client to communicate with the RIS server during setup of the installation

process. Remote Installation is the server-side component that sends a default set of CIW screens to guide the client through the remote installation process. Remote boot-enabled clients use the CIW to log on and select from operating system installation options. You can customize these setup screens to meet the needs of your organization.

**Active Directory Users and Computers Extension for RIS (Dsa.msc)** When you create the RIS server, the Active Directory Users and Computers extension installs on the RIS server. The extension provides a **Remote Install** tab within the computer account **Properties** dialog box of each RIS server that allows you to administer the RIS server. You can start this extension by specifying the Microsoft Management Console (MMC) snap-in Dsa.msc in the **Run** dialog box or you can start it from the command line.

You can administer RIS locally or through a Terminal Services session on another network computer. You can also administer RIS from a computer running Windows XP Professional if you install the Adminpak.msi on that computer.

#### RIS Tasks

Table 4.1 describes some of the tasks that you might perform while using RIS, the corresponding RIS components you would use, and which users can perform the tasks.

**Table 4.1 RIS Components, Tasks, and Users**

| Task | RIS Component | User |
|---|---|---|
| Install RIS | Remote Installation Services Windows component | Server administrator |
| Complete RIS server installation | Remote Installation Services Setup (Risetup.exe) | Server administrator |
| Configure Group Policy settings related to RIS | Active Directory Users and Computers RIS Extension (Dsa .msc) | Server administrator |
| Create operating system images, including application and desktop configurations, and install on RIS servers | Remote Installation Preparation Wizard (Riprep.exe) | Desktop administrator |
| Create boot floppy disk for non PXE–enabled client computers to install operating systems using RIS | Remote boot floppy generator (Rbfg.exe) | Desktop administrator |
| Provide log on and selection of operating system images to RIS clients | Client Installation Wizard (OSChooser.exe) | End user |

#### RIS Technology Limitations

You can use RIS technology to install operating systems, with or without software applications, to portable and desktop computers in your organization, which include member servers, stand-alone servers, and domain controllers. However, limitations to the scope of RIS-based operating system installations include the following:

**Clean Installs** You can only use RIS to provide a clean version of an operating system, with or without software applications. You cannot use RIS to upgrade an operating system or software configuration.

**Server Components** If you use RIS to install a server operating system, you might not be able to include all the server components you want to provide with the RIS image. For example, some server components require that you install and configure them only after the RIS-based installation is complete. This can include components such as Certificate Services, Cluster service, or software that is dependent on Active Directory.

**Domain controllers** You cannot install a preconfigured domain controller using a RIS image. However, you can use RIS to install a stand-alone server and then configure the server as a domain controller by running the Active Directory Installation Wizard.

**Encryption and security settings** You cannot use RIS to deploy files that are encrypted with a system such as the Encrypting File System (EFS). Also, you cannot use RIS to deploy systems with preconfigured user-level security settings such as file and folder permissions. To configure these settings, you can run a script after completing your RIS-based installation.

**Wireless networks** Wireless networks do not support pre-booting computers using PXE technology.

**Multihomed computers** Multihomed RIS servers are supported if the network adapters service multiple separate subnets or if all network adapters service the same subnet. In both cases the RIS server must also be the DHCP server. The DHCP server must have active scopes for each subnet serviced and must be authorized for each IP address on the network adapters being

serviced.

**Supported operating systems** RIS has certain limitations depending on the operating system that you are installing. For more information about operating systems supported by RIS, see "Operating systems supported by Remote Installation Services" in Help and Support Center for Windows Server 2003.

# Planning RIS Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Careful planning is critical to ensure successful RIS-based installations. Effective planning minimizes the time and effort required to support large-scale operating system installations across your organization.

In large organizations that support hundreds or even thousands of desktop computers, it is expensive and inefficient to manually install every operating system and respond to each dialog box that Setup displays. In this type of environment, the best approach is to use RIS to automate and customize the operating system installation process for your clients. An appropriate plan for custom RIS-based operating system installations lets you:

- Accommodate differing software and hardware configurations and varying user needs.
- Control the interaction level of end users during installations.
- Minimize the number of operating system images you need to manage.

Typically you would use a volume license for bulk rollouts of Windows XP or Windows Server 2003. If you would like to use individual product keys for each installation, you need to use the Windows Management Instrumentation (WMI) Windows Product Activation (WPA) provider. For more information about the WMI WPA provider, see the at http://www.microsoft.com/reskit, or see the MSDN Scripting Clinic link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

When planning your RIS-based operating system installation, you must first gather the data you need to choose the appropriate configurations for your RIS servers and clients. The steps in this process are illustrated in Figure 4.3.

**Figure 4.3 Planning RIS Installations**




For additional information about Best Practices to assist your planning process, click the Index button in Help and Support for Windows Server 2003 and in the keyword box type **Remote Installation Services**, then select **Best Practices**.

# Identifying Client Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have existing client computers on your network that you are considering for RIS-based operating system installations, you need to evaluate the hardware and software configurations of these clients during the planning phase. If you need to obtain new client computers to receive RIS-based operating system installations, you can acquire them from an OEM/Solution provider. For both new and existing clients, you need to determine if your clients satisfy the requirements for RIS-based installations.

**Important**

- You can only use RIS to perform clean operating system installations. You cannot use RIS to perform an operating system upgrade.

To identify requirements for RIS client computers on which you want to install an operating system by using a RIS-based installation, you need to:

- Evaluate whether your client computers meet the minimum hardware requirements for the operating system you intend to install.
- Determine if your client computers utilize the same HAL as the master computer (for Riprep images only).
- Evaluate the remote boot capabilities of RIS clients (whether they are PXE-enabled or if they require a RIS boot floppy disk).
- Audit existing client computers so you can inventory software and hardware configurations.
- Decide if you want to prestage RIS clients in Active Directory for more secure RIS-based operating system installations.
- Evaluate operating system configurations.

You can perform these tasks in any order. For more information about each task, see the sections that follow.

# Evaluating RIS Client Hardware

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In general, RIS client computers must meet the requirements for the operating system that you install on them. Whether you are deploying a client or server operating system, it is still considered a client installation with respect to RIS. However, hardware requirements differ between computers that host client operating system software and those that host server operating system software.

For more information about Remote Installation Services system requirements, see "Remote Installation Services system requirements" in Help and Support Center for Windows Server 2003.

To verify whether your existing client hardware is compatible with Windows XP, see the Windows Catalog link on the Web Resources page at: http://www.microsoft.com/windows/reskits/webresources.

You can also verify hardware compatibility by consulting the following:

- **Hardware Compatibility List**. A list of software and hardware products that are compatible with Windows 2000 and earlier Windows operating systems. See the Hardware Compatibility List link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- **Windows XP Professional Upgrade Center**. A Microsoft Web site that provides information for verifying whether your system hardware and software is compatible with Windows XP Professional. See the Windows XP Professional Upgrade Center link on the Web Resources page at: http://www.microsoft.com/windows/reskits/webresources.

For more information about verifying software and hardware compatibility see "Designing Image-based Installations with Sysprep" in this book.

For this part of your planning process, use job aid "Planning for RIS Clients" (ACIRIS_01.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Clients" on the Web at http://www.microsoft.com/reskit) to indicate whether client computers require hardware upgrades. Also indicate the domains or organizational units in which clients must receive an upgrade and the personnel you want to assign to the task.

# Determining RIS Client HAL Types

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you want to create an image-based RIS installation using the tool Riprep.exe, you first need to determine if your RIS clients have a HAL that is compatible with the master computer from which you create your image. For example, if the master computer where you run Riprep.exe has an Advanced Configuration and Power Interface (ACPI) HAL, then the client computers you designate to receive operating system images generated from that master computer must also have an ACPI HAL. The HAL type is indicated by the original file name of the file Hal.dll.

You need to ascertain how many different HAL types exist in your organization. This determines how many different master installations you will need with HALs that are compatible with the client HALs in your organization. To verify the HAL type on your client computers, you can do one of the following:

- Use a management tool such as Microsoft® Systems Management Server (SMS) to obtain your client inventory, from which you can determine the HAL types.
- View the properties of Hal.dll to determine the HAL types.

To view the properties of Hal.dll on a client computer:

**To obtain the HAL type on a client computer**

1. In **Windows Explorer**, open the *systemroot*\System32 folder.
2. Right-click Hal.dll, and then click **Properties** on the shortcut menu.
3. In the **Itemname** list on the **Version** tab, click **Original File name**.

   The original file name that displays in the **Value** list (such as Halacpi.dll or Hal.dll) indicates the HAL type.

For more information about HAL, see "Designing Image-based Installations with Sysprep" in this book, to determine the type of HAL that is installed on the computer.

You can install RIS-based operating system images if any of the following conditions are true regarding HALs:

- The master and destination computer HALs are identical.
- The master and destination computers both have either uniprocessor or multiprocessor Advanced Programmable Interrupt Controller (APIC) HALs.
- The master and destination computers both have either uniprocessor or multiprocessor Advanced Configuration and Power Interface (ACPI) HALs.

For a listing of the HAL types that Windows XP Professional and Windows Server 2003 support, along with additional information about determining HAL types and identifying hardware that impacts image-based installations, see "Designing Image-based Installations with Sysprep" in this book.

**Tip**

- By default the HAL auto-detect feature of Riprep.exe causes the CIW to filter images based on the HALs of the client computers. CIW only lists images with compatible HAL types. You must generate the operating system images from master computers that have HALs compatible with those of the RIS clients.

For this part of your planning process, use job aid "Planning for RIS Clients" (ACIRIS_01.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Clients" on the Web at http://www.microsoft.com/reskit) to indicate the number of different HAL types in your organization and whether you require multiple master installations to match the

differing HAL types of client computers. Also specify the domains or organizational units where the clients exist.

# Evaluating Remote Boot Capabilities of RIS Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To initiate a RIS-based operating system installation, a RIS client must first perform a remote (network) boot by connecting to a RIS server over the network. When the RIS client locates and downloads the boot files located on the RIS server, CIW displays on the client and prompts the client to logon and begin the installation process. The best way to facilitate the remote boot is to use a PXE-enabled RIS client, which means that both the network adapter and BIOS of the RIS client support PXE.

However, if a RIS client does not have a PXE-enabled network adapter and a supporting BIOS, you can emulate PXE support by using a PCI-based network adapter that boots from a RIS boot floppy disk. The RIS boot floppy disk is a startup disk that simulates the PXE startup process for computers that lack a remote boot-enabled BIOS. To be able to use a RIS boot floppy disk, RIS client computers must meet the minimum processor, Random Access Memory (RAM), and hard drive specifications referenced in "Evaluating RIS Client Hardware" earlier in this chapter. By using the RIS boot floppy disk to emulate PXE support, you can enable RIS-based operating system installations on non-PXE-compliant client systems. The Remote Boot Floppy Generator (RBFG) utility allows you to generate RIS boot floppy disks for use with RIS clients that are not PXE-enabled.

**Note**

- The remote boot floppy generator does not support Windows XP 64-Bit Edition Version 2003 or the 64-bit versions of the Windows Server 2003 family.

**Verifying the RIS Client Remote Boot Configuration**

You need to verify whether your RIS clients are PXE-enabled. To do this, you can check the documentation and historical records that came with the system. If you do not have this information, you need to do the following:

- Verify your RIS clients have PCI, Mini-PCI, or CardBus type network adapters. RIS clients can only perform a remote boot from these types of network adapters.

- Verify that the BIOS of your RIS clients is capable of using the network adapter as the primary boot device. A ROM BIOS that is at least version .99n satisfies this requirement.

  **Note**

  - Most computers that conform to the Net PC or PC98 specifications have a PXE remote boot-enabled network adapter and remote boot-enabled BIOS.

You might be able to obtain this information using a management system such as SMS or directly from inspection of the client computers.

However, it might be easier for you to use a remote script to determine whether the BIOS of client computers in a specified domain supports PXE-enabled remote booting. To do this, you can use the BIOS Information script, which you can find from the Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources.). To return the BIOS information, run this script at the command line and specify the Active Directory domain name and the **getallbios** command. The returned information verifies if the BIOS of client computers supports PCI adapters and selectable booting. If so, then the BIOS supports remote booting from a network adapter.

**Note**

- The BIOS information returned by the BIOS information script does not conclusively determine if your system is completely PXE-enabled. However, if the information indicates that the BIOS does not support PCI network adapters or does not have a selectable boot capability, then it is a certainty that the BIOS does not support PXE-enabled remote booting.

The BIOS information script uses Active Directory Service Interfaces (ADSI) to query the Active Directory "Computers" container to obtain the computer objects for which the BIOS information is returned. The script also uses Windows Management Instrumentation (WMI) technology to query each computer and return the specific BIOS information to the command line. You can also redirect this output to a text file that you specify. For the script to function properly, you must have

WMI installed on both the computer running the script and the computers that you query with the script. In addition, you must have ADSI installed on the computer running the script.

Network adapter manufacturers can embed the PXE-based remote boot code on a chip as part of the network adapter. Some manufacturers create a version of PXE ROM code as part of the client system BIOS to support the PXE environment specification.

For RIS clients that are not PXE-enabled, you need to determine if they can use a RIS boot floppy disk. To use a RIS boot floppy disk, these clients must have a PCI-type network adapter because RIS boot floppy disks do not support Personal Computer Memory Card International Association (PCMCIA), CardBus, ISA, USB, or other non-PCI network adapters. You can generate the RIS boot floppy disks for these clients by running the Rbfg.exe utility. This utility is located on the RIS server in the following directory location:

*RISServerName*\RemoteInstall\Admin\i386\Rbfg.exe

**Note**

- To view a list of supported PCI network adapters supported by Remote Boot Floppy Generator, run Rbfg.exe and click the Adapter List button on the displayed dialog box.

The RIS boot floppy disk also enables you to use RIS to install operating systems on portable computers. However, portable computers often use PCMCIA network adapters that PXE does not support, so you cannot use these adapters with RIS-based operating system installations. Alternatively, you can place the portable computer in a docking station that contains a PCI network adapter and use a RIS boot floppy disk to facilitate the installation. You cannot add network adapters to the RIS boot floppy disk you create with Rbfg.exe.

**Note**

- You should not use the same docking station to install RIS for multiple portable computers.

For RIS clients that require the use of a RIS boot floppy disk, the boot sequence in the BIOS must be set so that booting the floppy disk occurs first. For PXE-enabled RIS clients, the BIOS must be set to allow booting from the network adapter.

**Important**

- You can use these BIOS settings when you implement interactive installations. However, if you intend to use automated installations, you must ensure that the BIOS boot configuration is set to use the hard disk as the first boot device. For more information about automated installations, see "Fully-Automated Installation Design Background" later in this chapter.

For more information about the PXE environment, including its security context, see "PXE Specifications" earlier in this chapter.

For this part of your planning process, use job aid "Planning for RIS Clients" (ACIRIS_01.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Clients" on the Web at http://www.microsoft.com/reskit) to record whether client computers require:

- Upgrades to support PXE.
- The use of RIS boot floppies.
- BIOS upgrades to support selectable booting.

Also specify the domains where the clients exist and the personnel you want to assign to the upgrade tasks.

# Auditing Existing Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you control your existing client computers using Systems Management Server or a similar management product, you can wipe these computers clean in preparation for a new RIS-based installation. However, before you do this, it is advisable to audit these client computers by:

- Conducting an inventory.
- Backing up all user data and settings.
- Obtaining client UUIDs.

In this part of your planning process, use job aid "Planning for RIS Clients" (ACIRIS_01.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Clients" on the Web at http://www.microsoft.com/reskit) to indicate your decisions to conduct an inventory, backup user data, and obtain client computer UUIDs. You can also specify the domains or organizational units to audit and the personnel you want to assign to inventory and backup tasks.

**Conducting an Inventory**
You conduct an inventory to determine such things as the number of existing clients, the types of existing desktop configurations in your organization, hardware configurations, software and hardware compatibility, and existing applications. For more information about creating hardware and software inventories, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

**Inventory-related factors that impact RIS installation planning**
The following are some issues that can arise when conducting an inventory and how they can affect the RIS installation planning process. Your primary concern is how these issues affect the number of images you need to create.

**Number of existing clients** The number of clients that have differing requirements can give you an indication of the number of RIS images you might need to create. Also, the number of clients you have determines the tool you need to use for automating the installation process. If you have a large number of clients, RIS is a good choice for the task. If you have a small number of clients, you might be better off using the Winnt.exe or Winnt32.exe programs in unattended mode. For more information about unattended installations, see "Designing Unattended Installations" in this book.

**Existing desktop configurations** The types of existing desktop configurations in your organization provide an indication of the number of different RIS images and answer files you might need to create. Alternatively, you can apply a standard desktop to all clients in your organization, which reduces the number of RIS images and answer files you need to maintain.

**Hardware configuration types** The types of existing client hardware configurations you have affects the tool you use to create installation images, the hardware device drivers you add to a distribution image, and how you configure answer files.

For example, if you want to use Riprep to create file-system-based images of an operating system, the master computer and all destination client computers must have identical HAL types.

Also, if you have hardware not supported by the Mini-Setup process, which occurs after image installation, this affects how you create RIS-based installations. For example, if you have a SCSI hard disk device that is not supported during the text mode phase of Mini-Setup, you need to add the driver files for the SCSI device — along with its Textsetup.oem file — to the distribution folder that you create using Risetup. You must also modify the [MassStorageDrivers] section of the RIS default answer file to add values for the appropriate driver entries.

For more information about Mini-Setup, see "Riprep Image Design Background" later in this chapter.

**Software and hardware compatibility** You need to ensure that your existing client computers meet the software and hardware requirements that support installation of the operating system images you plan to make available to them.

For information about hardware compatibility, see the Windows Catalog link on the Web Resources page at: http://www.microsoft.com/windows/reskits/webresources

To verify software compatibility for your existing clients, you can use the **Application Compatibility Toolkit**, which contains documents and tools to help you diagnose and resolve application compatibility issues. For more information about verifying

software and hardware compatibility, see "Designing Image-based Installations with Sysprep" in this book.

**Application configurations** The application configurations that exist on client computers indicate what you need to include in the Riprep images you create or what you need to provide on the distribution folder you create using Risetup.

#### Special hardware components

As part of planning your RIS-based installation, you must inventory special peripherals, hardware devices, and software configurations because these components can affect the number of images you need to create. A special component might be an existing hardware device that Mini-Setup (which follows image installation) does not support. If you do not take these components into account in the planning phase, the RIS-based installation that you implement might fail.

Some of the special hardware components you might need to add to your inventory include:

- Mass storage controllers
- Portable computer devices
- Vendor-specific devices
- Legacy devices

**Note**

- Most Plug and Play peripheral devices, such as sound cards, network adapters, modems, and video cards, have no impact on RIS-based installations. It is therefore unnecessary to inventory these devices because Setup automatically detects, installs, and configures them after the RIS image is copied to the destination computer.

For more information about inventories for special hardware and software components, see "Designing Image-based Installations with Sysprep" in this book.

#### Special software components

The installation requirements of some applications might cause you to alter the way you perform a RIS-based installation or they might call for you to create separate images. For example, some applications require that you install them after the RIS installation is complete, rather than installing and configuring them in a master installation (for a Riprep image). Also, you might be able to install some applications only on portable computers but not on desktops. In this situation, you probably need to create separate RIS images.

Some of the special software components you might need to inventory include:

**Core software applications** These applications might include office productivity suites and anti-virus software. You typically install and configure core applications in your master installation, although you can add applications to a distribution share that you create with Risetup images. If you have any computers on which you do not want to install core applications or any computers that require special configuration settings for these applications, note these in your software inventory. You probably need to create separate images to accommodate these differences.

**Line-of-business applications** These applications might include accounting, database, or investment modeling programs. You need to identify the groups in your organization that require line-of-business applications because you might want to create separate RIS images for specific groups, especially if they require substantial configuration or take a long time to install.

**Applications with Active Directory dependencies** You need to identify any applications that depend on Active Directory. You cannot include these types of applications in a RIS image because you can only install and configure them after installation of a RIS image is complete.

**Third-party utilities and tools** These applications might include the diagnostic tools of a computer manufacturer or a suite of hardware-specific utilities for a portable computer that allow you to configure power options and other features. You might need to install these utilities and tools after installation of a RIS image is complete. You might also want to create a separate RIS image for the computers that require these utilities and tools.

**Service packs, hotfixes, and other software updates** Identify all service packs, hotfixes, and other software updates that are installed in your organization. Be sure you record the revision number and the revision date of the service pack, hotfix, or other software update in your inventory.

#### Migrating User State

You will need to create a user state migration plan if any of your destination computers contain any of the following items that you want to restore after installation is complete:

- User data that you want to be available to the end user. User data includes such things as documents, e-mail messages, spreadsheets, and databases.
- User settings such as desktop settings, shortcuts, and Microsoft® Internet Explorer Favorites.
- Application settings such as application-specific keyboard shortcuts, spell-checking options, and default file locations.

At a minimum, your user state migration plan must do the following:

- Identify the data you want to migrate, including user data, user settings, and application settings.
- Determine where to store the data while you perform the image-based installation.
- Create a schedule for migrating data on each of your destination computers.
- Describe how to collect and restore the data.

Microsoft provides two tools for migrating user data and settings. Each tool is designed for different types of users and environments.

- **Files and Settings Transfer Wizard.** Designed for home users and small office users. The wizard is also useful in a corporate network environment for employees who get a new computer and need to migrate their own files and settings without the support of an IT department or Help desk.
- **User State Migration Tool.** Designed for IT administrators who perform large deployments of Windows XP Professional in a corporate environment, the User State Migration Tool provides the same functionality as the wizard, but on a large scale targeted at migrating multiple users. The User State Migration Tool gives administrators command-line precision for customizing specific settings, such as unique modifications to the registry. The User State Migration Tool is located in the \valueadd\msft\usmt\ folder on the Windows Server 2003 CD.

For more information about migrating user data and settings, see "Migrating User State" in this book. Also see the articles "User State Migration in Windows XP," "Step-by-Step Guide to Migrating Files and Settings," "Deploying Windows XP Part I: Planning," and "Deploying Windows XP Part II: Implementing." To find these articles, see the Microsoft TechNet link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources.

**Note**

- You might not need to backup user data if you use folder redirection and roaming profiles to store user data and configuration settings on a server. For an overview of Group Policy and information about using folder redirection and roaming user profiles, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

#### Obtaining Client Computer UUIDs

During the auditing process, you can also obtain the UUIDs for your existing client computers. You need the IDs for existing clients that you are planning to prestage in Active Directory to enhance the security of RIS installations. For more information about obtaining UUIDs for prestaging your client computers, see "Evaluating the RIS Client Prestaging Process" later in this section.

Some of your clients, such as those which are not PXE-enabled, might not provide a UUID. For these clients, you need to use the RIS boot floppy disk. After using the RIS boot floppy disk to initiate startup on these clients, Remote Installation automatically generates UUIDs and computer accounts for these clients based on the MAC address of the network adapter.

#### Evaluating the Creation of UUIDs for Non PXE-Enabled Clients

When a PXE-enabled client prestaged in Active Directory connects to a RIS server, the UUID of the client computer is passed to the RIS server. The UUID is a unique 32 character hexadecimal or 16-byte number, which is stored with the computer account

object that you create in Active Directory. To generate a UUID for these computers, the Remote Installation Services uses the media access control (MAC) address of the network adapter, which is 12 characters long, and prepends 20 zeroes to create a unique 32 character hexadecimal or 16-byte number. After this occurs, Remote Installation creates a new computer account object in Active Directory and associates this unique 32-bit number with the account. Because the MAC address of a network adapter is unique on the network, so is the 32 character hexadecimal or 16-byte identifying number and the computer account as well.

If you want to enhance security for non PXE-enabled clients, you can prestage computer accounts in Active Directory for these clients using a prestaging script. The prestaging script uses UUIDs generated by a BIOS information script, which retrieves UUIDs from PXE-enabled client computers, suggests possible UUIDS for non PXE-enabled client computers, and stores this information in an Excel spreadsheet. This script creates UUIDs by using a process similar to the Remote Installation method To find these scripts, see the Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources. For more information about obtaining UUIDs and prestaging computer accounts, see "Evaluating the RIS Client Prestaging Process" later in this chapter.

For this part of your planning process, use job aid "Planning for RIS Clients" (ACIRIS_01.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Clients" on the Web at http://www.microsoft.com/reskit) to indicate whether you intend to allow Remote Installation to suggest possible UUIDs and computer accounts or if you want to create UUIDs for non PXE-enabled clients by using the BIOS information script.

# Evaluating the RIS Client Prestaging Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Prestaging computer accounts in Active Directory means that you create computer account objects in Active Directory prior to client computer startup, using the UUIDs of the client computers to configure the netbootMachineFilePath attribute in each computer object. You also configure the user accounts that will use the client machines by providing them with read, write, and set/change password permissions on the computer account objects. When these clients boot to a RIS server, they send their UUID to the RIS server. The server then checks Active Directory for a UUID that matches the UUID that the client sends to the RIS server. If one is found, the RIS server accepts the request for service from that client. By using the prestaging process, you can greatly enhance security by causing your RIS server to recognize specific clients only.

If you are not concerned with the security of servicing RIS client requests for operating system installations and you are not planning to provide automated installations, you can bypass the prestaging process. If you decide to prestage your client computers in Active Directory to enhance security or provide automated installations, you need to obtain the UUIDs for client computers so you can provide them during the prestaging process. Prestaging clients in Active Directory assures that the RIS server recognizes service requests from these clients, while ignoring all others (if you configure the RIS server to do so). For more information about prestaging client computers to enhance security, see "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit).

In many cases, you can get an Excel spreadsheet with the UUID information from the OEM supplier of the client computers. If you do not have UUID information from the OEM, you can use the BIOS information script. To find this script, see the Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources.). You can use this script to automate the process of obtaining UUIDs of all client computers in the default Active Directory Computers container.

To obtain UUIDs by using the BIOS information script, you must run the script at the command line and use the **getalluuids** command. The script uses ADSI and WMI technologies to return valid UUIDs that you can use to prestage client computers in Active Directory. The script provides usage instructions that explain all the input arguments and commands you must specify.

Alternatively, you can also use SMS to obtain the UUID for a computer or group of computers. To use SMS to identify the UUID of a computer or group of computers, see the documentation provided with SMS.

Once you have the UUIDs for your client computers, you can use them to prestage clients by creating new computer accounts in Active Directory. For procedures to prestage RIS clients using the Active Directory snap-in on a RIS server, see "Remote Installation Services administration overview" in Help and Support Center for Windows Server 2003.

You can automate the prestaging process by using the prestaging script; to find this script, see the Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources.). If you have UUID listings on an OEM spreadsheet, you can adapt this information as input data to the script, but you must use the exact same Excel spreadsheet format that BIOS information script creates. Otherwise, use the BIOS information script with the /ExcelPath: option to print the UUIDs to an Excel spreadsheet for the data you need as input to the prestaging script. For more information about designing Active Directory support, including prestaging RIS clients in Active Directory and automating the prestaging process by using scripts, see "Designing the Active Directory Infrastructure" later in this chapter.

For this part of your planning process, use job aid "Planning for RIS Clients" (ACIRIS_01.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Clients" on the Web at http://www.microsoft.com/reskit) to indicate your decision to prestage client computers in Active Directory. Also specify the method you intend to use to obtain the UUIDs and the personnel assigned to the task.

# Evaluating Operating System Configurations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To determine the configurations in which to deploy operating systems to your RIS clients, you need to determine what operating system image configurations you need in your organization. To do this, you need to assess user needs for the following:

- Operating systems
- Desktop configurations
- Drivers and applications
- Computer workstation types, such as desktop or portable computer
- User requirements
- Server components

Completing this evaluation also helps you define the type of installations you require, such as:

- Single-image or multiple-image
- Risetup and Riprep images
- Interactive or fully automated modes

In this part of your planning process, use job aid "Planning for RIS Clients" (ACIRIS_01.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Clients" on the Web at http://www.microsoft.com/reskit) to define your operating system configurations and installation types.

**Operating Systems**
The primary decision you need to make is how many different client operating system installation choices you need to provide to your RIS clients. RIS can support installations of Windows XP Professional, the Windows Server 2003 Family, Windows 2000 Professional, and Windows 2000 Server. For each operating system you provide, you need to generate a separate image. This can be a file-system-based image generated from your master computer with Riprep, or a CD-based image that you create in a distribution share on the RIS server using Risetup. In both cases, you might also want to create different versions of each image to include certain applications, tools, and drivers for specific RIS clients, or special desktops in the case of Riprep images. The methods for creating custom images are different for Riprep.exe and Risetup.exe. Also keep in mind that each different operating system configuration you choose is another image you must create and maintain.

For information about Riprep and creating master computer installations see "Design a Riprep-Based Installation" later in this chapter. For information about Risetup and creating a CD-based image for a distribution share, see "Design a Risetup-Based Installation" later in this chapter.

**Desktop Configurations**
The number of desktop configurations you need depends on the desktop configurations you require for new client computers and the types of existing client desktops you already have (see "Auditing Existing Clients" earlier in this chapter). You can apply a new standard desktop configuration or an existing desktop configuration type to client computers throughout your organization. However, the number of different desktops you plan on making available to RIS clients, along with other components, dictates the number of different operating system images you need to generate and maintain. To make a particular desktop configuration available to RIS clients, you must install an operating system on the master computer, configure the desktop as you want it, and run Riprep.exe on the master computer to create the image and store it on the RIS server. The image on the server then contains the desktop configuration that you configured on the master computer.

**Note**

- You cannot customize desktops using Risetup images.

**Drivers and Applications**
If you need to provide custom drivers, applications, or support files to your RIS clients, you can use Riprep.exe to generate a file-system-based image from a master computer that you configure with the drivers, applications, and Help or support files you need. The most efficient way to provide specific application configurations to users is to prestage a large number of applications on the master computer and create different answer files for each application configuration. For more information about creating your master computer configuration, see "Configuring a Master Installation" later in this chapter.

You can also provide custom drivers and applications to your clients with CD-based images you create using Risetup.exe. After creating the Risetup image on the RIS server, you can do this by:

- Populating the distribution share with the specific drivers and applications your users need — this can include hardware device drivers, such as drivers for storage and Plug and Play devices.

- Modifying the answer file that the installation uses.

**Workstation Types**
The types of computers you want to receive RIS-based operating system installations, such as desktops or portables, can impact the operating system configurations you provide. You might want to provide different applications, drivers, and desktop configurations to each of these clients. For example, you might need to provide remote access and a unique suite of applications for your portable computer configurations. You might also need to provide display resolution settings for portable computers.

**User Requirements**
You need to identify the requirements of your users so you can define the types of users you have. Obtaining this information is important because it helps you decide how to customize your RIS-based installation design. For example, if a particular group of users needs a specific application, you must add it to the distribution share on your RIS server or configure it on the master computer from which you create a Riprep image.

You can classify user types based on criteria such as the following:

- Computer knowledge — such as beginner, intermediate, or advanced.

- Location — such as on-site, roaming, or remote.

- Job function — such as marketing, research, or customer service.

- Job category — such as manager, project lead, or individual contributor.

As an example, suppose that you classify certain users into types according to computer knowledge. This has an impact on what operating system installation choices you make available to these users. You want to allow less knowledgeable, task-oriented users to make few or no installation choices, while more advanced users might have several installation options from which to choose. To accommodate these differences, you control the configuration of the CIW using Group Policy settings to allow or disallow specific operating system installation options.

User requirements can also include such things as local account passwords, language needs, regional considerations, desktop configurations, and applications such as line-of-business, spreadsheets, and word processing.

**Server Components**
The components you need to provide with a Windows Server 2003 installation can vary, depending on the member server roles you need to provide. If you have several member server component configurations, you can associate multiple answer files with a single CD-based image of Windows Server 2003. For each CD-based image that you create using Risetup.exe, RIS creates a default answer file called Ristndrd.sif. By providing a modified version of this answer file to represent each member server configuration, you can offer a variety of unattended Windows Server 2003 installation configurations from the same source image on the RIS server. This way, you do not have to create additional images to provide different member server component configurations.

You can provide the server components by adding them to the distribution share you create after running Risetup.exe. You can add applications such as IIS, Telnet Services, or Exchange. You then need to configure various parameters for the operating

system and server components using the unattended answer file associated with the Risetup image on the distribution share.

# Evaluating RIS Server Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To create a RIS server in your organization, you need to add the Remote Installation Services component to a computer running Windows Server 2003. You can do this in **Control Panel** under **Add or Remove Programs**. You must also configure your RIS server using Risetup.exe and you must create at least one operating system image on the server.

You create a RIS server in the deployment phase of your RIS deployment process, but you need to do some preliminary analysis in the planning stage to assure good performance. As part of the planning process, you need to evaluate the following elements to ensure that your RIS server meets the necessary specifications:

- Hardware requirements
- Software requirements
- Server placement
- Server performance

In this part of your planning process, use job aid "Planning for RIS Servers" (ACIRIS_02.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Servers" on the Web at http://www.microsoft.com/reskit) to record your planning decisions. At this point, you can specify which personnel you want to perform the task of creating RIS servers and initial installation images.

# Evaluating RIS Server Hardware Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In general, RIS servers must meet the requirements of the product version of the Windows Server 2003 family you install. However, you are encouraged to increase those requirements to more efficiently support RIS image deployment in your organization.

You need to have at least two disk partitions available on the RIS server, one for booting the server operating system and another to contain the directory structure for the client operating system images. You need to allocate an entire separate partition to the RIS server directory tree because you cannot install RIS on the same drive as the system volume. Also, you need to format the RIS partition as NTFS.

The partition containing the images must be large enough to store one or more operating system images, depending on your requirements. A CD-based image for Windows XP Professional is 650 MB–700 MB in size. Most organizations store more than one operating system image to meet the needs of different clients.

For up-to-date hardware compatibility information, see the Windows Catalog link on the Web Resources page at: http://www.microsoft.com/windows/reskits/webresources.

For this part of your planning process, use job aid "Planning for RIS Servers" (ACIRIS_02.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Servers" on the Web at http://www.microsoft.com/reskit) to record whether you require an upgrade to RIS server hardware. Also indicate the personnel you want to assign to the task.

# Assessing RIS Server Software Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Several software components are necessary to support the functioning of a RIS server in your network. You need to verify whether the Active Directory, DNS, DHCP services exist on your network before deploying RIS in your organization. You can find all of these services, including the RIS component, on a server running Windows Server 2003. However, you might also have the DNS or DHCP services provided by a third party. In any case, to assure that you do not have any performance degradation in your production environment, install these services on different servers and configure each server for the specific role of that service.

**Note**

- You can install all of these software components on a single RIS server when you create your test environment, as discussed in "Designing a Test RIS Environment" later in this chapter.

**Remote Installation Services (RIS)**
RIS is an optional component of Windows Server 2003 that provides the services that enable you to automate the remote installation of an operating system, such as Windows XP Professional, on client computers in your organization.

**Domain Name System (DNS)**
RIS servers rely on DNS to locate the required Active Directory servers to facilitate domain operations. If you use Windows Server 2003 DNS, you have the benefit of dynamic updates for your DNS server. However, it is not a requirement to use Windows Server 2003 DNS for RIS to function. Whichever DNS server you use, it must support the SRV RR record type and the dynamic update protocol specified in RFCs 2052 and 2136, respectively. For more information about DNS, see "Deploying DNS" in *Designing Network Services* of this kit.

**Dynamic Host Configuration Protocol (DHCP)**
RIS servers require a DHCP server on the network which is authorized and has an activated scope. Remote boot-enabled clients must receive an IP address from a DHCP server before they can contact a RIS server to request an operating system installation. You can install Windows Server 2003 DHCP or you can use existing DHCP services provided with Windows 2000 Server. In addition, you can use a third party DHCP server. For more information about Dynamic Host Configuration Protocol (DHCP), see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

**Active Directory**
You must install RIS on a computer running Windows Server 2003 in an Active Directory domain. For best results, configure this computer as a member server. Although you can install RIS on a domain controller, the heavy traffic load generated by RIS can impact the performance of the domain controller.

RIS uses Active Directory to locate RIS clients and other RIS servers. You can administer the RIS server from the Active Directory Users and Computers snap-in (dsa.msc) located on the RIS server. For more information about Active Directory, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

For this part of your planning process, use job aid "Planning for RIS Servers" (ACIRIS_02.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Servers" on the Web at http://www.microsoft.com/reskit) to indicate whether you need to verify the existence of the software components required to support RIS, along with the personnel you want to assign to this task.

# Assessing RIS Server Placement

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the test environment that you use to test RIS features, RIS performance, and the installation process, RIS server placement is not critical. For example, you can place the RIS server and other supporting components, such as Active Directory, DNS, and DHCP all on a single server. For more information about creating a RIS test environment, see "Designing a Test RIS Environment" later in this chapter.

In the production environment of a large organization, an all-in-one configuration is not recommended because high RIS traffic volumes can cause significant performance degradation in other services on the server that hosts RIS. This can occur when Server Message Block (SMB) traffic from the server to numerous clients in the setup process precludes other traffic on the network. This results in inhibiting DHCP traffic, new PXE requests, other TCP/IP network traffic, and even some Active Directory replication modes. When the performance of DHCP degrades, this can slow the network or even bring it down. For these reasons, avoid configuring production domain controllers with multiple roles that include RIS, Active Directory, DNS, and DHCP.

To assure adequate performance in a large organization, consider using separate computers in your production environment for the following components:

- Domain controller with Active Directory and DHCP
- DNS
- RIS server

In relatively small organizations, placing your DHCP and RIS servers on the same computer is a common practice and works well in most situations. However, you need to be aware that whenever your RIS server approaches its maximum capacity to handle PXE requests, you might begin to see a failure in the ability of your DHCP server to service client DHCP requests. For further information about combining DHCP and RIS, see "DHCP and RIS Server Considerations" later in this section.

Also observe the following guidelines when placing your RIS server(s) in the production environment:

- Do not place RIS on the same computer that is running Exchange Server or Microsoft® SQL Server™.

  The high traffic levels of RIS can degrade the performance of these products, and vice versa.

- You cannot host RIS on a computer in a wireless network.

- Do integrate RIS into networks with preexisting third party remote installation servers.

  RIS technology allows the coexistence of remote installation servers from multiple vendors on the same physical network. When you set RIS servers to ignore boot requests from unknown clients, you can introduce them on a network without interfering with preexisting remote installation servers that use the same remote boot protocols.

**DHCP and RIS Server Considerations**

Because PXE-enabled RIS clients use the DHCP discovery mechanism to obtain a network (IP) address and to locate RIS servers, the relationship of RIS to DHCP in your organization can play a key role in determining your RIS server placement strategy.

In simple environments, a common solution is to add RIS to each DHCP server in use. When you use a combination Windows Server 2003 DHCP/RIS server approach, this reduces the number of initial network packets that RIS clients send to the DHCP and RIS servers. This also increases the server's initial response time. In addition, combining a Windows Server 2003 DHCP and RIS server provides simultaneous answering of client requests. This creates a simplified form of load balancing, because it takes advantage of existing groupings of client computers associated with the DHCP server. This configuration also simplifies troubleshooting and administrative procedures.

**Important**

- If the RIS server and DHCP server coexist on the same computer and the RIS server becomes too busy to answer PXE

requests, then the DHCP server also becomes too busy and cannot answer IP address requests. However, this issue might be significant only in relatively large organizations.

Because RIS servers can generate large network loads, they often require high-end hardware and usually must be located near the clients they service. By contrast, a DHCP server generates far less traffic, does not typically require high-end server hardware, and is often centralized rather than near client computers. In a centralized configuration, you might find it impractical to simply add RIS to your existing DHCP servers. In such cases, consider adding RIS services to existing software installation point servers, because they have planning and placement requirements similar to RIS servers.

Also, because the PXE-based remote boot process does not provide a way to determine from which RIS server a client receives service, you need to control which RIS server answers specific clients. This is a primary issue when RIS servers are separate from DHCP servers, or when DHCP servers that are not running Windows Server 2003 are in use. In this situation, the location of RIS clients can have an impact on how you configure RIS server selection and load balancing, and subsequently on where you place a RIS server.

**RIS Server Selection and Load Balancing**

By default, when a PXE-enabled RIS client broadcasts a request for service, all RIS servers receiving the request initiate a reply. The first RIS server on the network from which the client receives a response is the one that provides service to the client. However, combined DHCP/RIS servers have response priority over servers hosting these applications separately. Although this provides a simple form of load balancing among multiple RIS servers available to clients, it is better to balance the load by explicitly restricting which servers can respond to specific clients. This also allows you to prevent specific clients from using RIS servers that you do not intend them to use.

To control server selection, you can physically control network routing so that DHCP discovery broadcasts are forwarded only when appropriate. For example, you might want to forward DHCP broadcasts when a server local to a client is busy or down when the client requests service. By using forwarding to control the routing to a DHCP/RIS server, you can allow answers from only those RIS servers that you configure to receive forwarded client requests.

Another way to control server selection is to configure clients that you prestage in Active Directory to only use a specific RIS server. When a RIS server answers a service request from a client, the server checks the Active Directory forest for the UUID configured in the client computer account to verify whether there is a match to the UUID submitted with the service request. If the RIS server finds a matching computer account, it also checks the account to confirm whether it is configured to use a specific RIS server. If the account designates that only a specific RIS server can provide service, then this is the server that services the client request. The RIS server initially processing the client request then points the client to the RIS server that actually provides the requested RIS services. This is known as a *server referral*. You can use this mechanism as a simple way to control which RIS servers offer operating system installation services to specific client computers.

For information about combining prestaging and referral capabilities to enhance flexibility and security, see "RIS Server Configuration Design Tasks" later in this chapter.

At this point in your planning process, use job aid "Planning for RIS Servers" (ACIRIS_02.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Servers" on the Web at http://www.microsoft.com/reskit) to specify the server/software configuration you want to use. Also record the method you intend to use to load balance client service requests to RIS servers (either by configuring routing or using RIS referral servers) and the personnel you want to assign to the task.

**Network Load Considerations**

Because RIS servers install operating system images on client computers, the amount of traffic the server produces is similar to that of other servers performing as software installation points on your network. However, the amount of RIS server traffic is more predictable than traffic from a general purpose software installation point that provides applications and regular updates. For example, RIS traffic increases when many users are loading images, during deployment of a new operating system or when you add new computers to the network, and decreases after the initial installations are complete.

To accommodate the periodic high traffic volumes that a RIS server generates, you need to place the RIS server in a location that minimizes its impact across your network. In general, place a RIS server near the client computers it services. This localizes the traffic and reduces its impact during times when multiple image downloads occur. If you have an environment in which re-installations occur frequently, you might consider segmenting the physical network to isolate the RIS server and dedicate it to client installations on that segment.

If you have an environment in which you perform a large number of operating system pre-installations before delivery to clients, consider implementing a RIS-based preinstallation lab. A lab such as this allows you to process a high volume of computers by using high-speed networking to reduce installation times. This also avoids any impact on network performance, because you do not have to place your RIS server on the network at all.

For this part of your planning process, use job aid "Planning for RIS Servers" (ACIRIS_02.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Servers" on the Web at http://www.microsoft.com/reskit) to indicate whether you plan to:

- Add your RIS servers to existing software installation points, or DHCP servers.
- Create a preinstallation lab for RIS servers.

If you choose to create a preinstallation lab, record the personnel that you assign to the task.

# Planning RIS Server Performance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To plan for the needs of your RIS server, you can assess RIS server performance in your test environment prior to deploying RIS in your network. The results you obtain provide some initial performance indications that can assist you in deciding where to make necessary improvements to the RIS server configuration.

To determine if you need improvements to enhance RIS server performance, you might first analyze the areas in which potential performance bottlenecks can occur. Doing this in your RIS test environment is beneficial for a baseline performance assessment, because the environment uses multiple server roles on the computer hosting RIS. However, for more definitive assessments, you can monitor performance in your actual environment during times of high demand on your RIS servers.

For more information about the RIS test environment, see "Creating a RIS Test Environment" later in this chapter.

Because RIS primarily uses a file-copy process, RIS server performance factors are common to those of other file input/output (I/O) intensive servers, such as Web, file, and print servers. These performance factors include server throughput, disk throughput, and network speed. When a RIS server or its network connection overloads, the result is increased installation time on client computers and TFTP time-outs during the initial file-copy phase.

For this part of your planning process, use job aid "Planning for RIS Servers" (ACIRIS_02.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Servers" on the Web at http://www.microsoft.com/reskit) to indicate whether you plan to use the test or production environment to obtain RIS server performance indications.

**Performance Bottlenecks**

A *bottleneck* is the part of a computer system that restricts workflow. Generally, a bottleneck is caused by over-consumption of a specific resource. Some causes of bottlenecks include:

- Too many active processes that need access to RAM.
- A processor running at a high percentage of utilization.
- A disk controller or drive that is slow when accessing data.

To monitor for bottlenecks, you can use the Performance tool provided with the Windows Server 2003 operating system to determine which system components are having an adverse effect on total I/O throughput. You can locate the Performance tool on a Windows Server 2003 from the **Start** menu in the **Administrative Tools** group. When you start the tool in MMC, a Help menu item is available to assist you in using the tool.

The components you should consider monitoring include the following:

**Memory.** The best indicator of a memory bottleneck is a sustained high rate of hard page faults. Hard page faults occur when the data a program needs is not found in the working memory visible to the program or elsewhere in physical memory, and must therefore be retrieved from the disk. An acceptable range for this parameter is 0–20 hard page faults per second.

**Processor.** Processor activity is especially important for server-based applications. Two of the most common causes of CPU bottlenecks are CPU-bound applications and excessive interrupts generated by inadequate disk or network subsystem components. Consider processor use in excess of 75 percent to be a bottleneck.

**Disk subsystem.** Just as processor usage is the bottleneck for server applications, the disk subsystem is often the bottleneck for file I/O performance. Because a RIS server is file I/O intensive, disk activity is a primary area to analyze. With the Performance tool, you can evaluate the disk activity of your RIS server in terms of total I/O throughput of the server and percentage use of the disk subsystem.

Table 4.2 describes the Performance tool objects and counters you can use to assess performance of critical systems, including memory, processor, and the disk subsystem.

**Table 4.2 System Monitor Objects and Counters**

| Object | Counters |
|---|---|

| | |
|---|---|
| Memory | Pages/sec |
| Processor | % Processor Time |
| Physical Disk | % Disk Time, Average Disk Queue Length |

**Tip**

- You can also use the Disk Bytes/Transfer and Disk Bytes/Second counters for the Physical Disk object. High values for these counters indicate efficient use of the disk subsystem, despite heavy disk loads.

Based on the performance monitoring results you obtain, consider making improvements in one or more of the following areas to optimize performance of your RIS servers:

- Increase processor speed.
- Increase the data transmission rate on the network.
- Decrease processor utilization, which can include shutting down unnecessary services on the computer hosting RIS.
- Improve disk utilization.

  Consider providing a disk configuration that uses multiple SCSI II (minimum) hard drives with Redundant Array of Independent Disks (RAID) 5 disk striping with parity.
- Increase the amount of installed RAM.

For this part of your planning process, use job aid "Planning for RIS Servers" (ACIRIS_02.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Servers" on the Web at http://www.microsoft.com/reskit) to indicate the components you plan to test (memory, processor, disk subsystem), the potential upgrades you anticipate, and the personnel you want to assign to these tasks.

# Assessing Master Computer Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To plan for your master computer configuration, you need to assess the following:

- Hardware requirements
- Operating system image requirements
- Placement of the master computer

**Master Computer Hardware Requirements**

The minimum hardware requirements for your master computer are similar to those of client computers because the master must be capable of supporting the installation of the client's target operating system. This might be an operating system such as Windows XP Professional or Windows Server 2003. For more information about client computer hardware requirements, see "Evaluating RIS Client Hardware" earlier in this chapter.

Because the master computer never participates in remote installations across the network, it is unnecessary for the master computer to be PXE-enabled. However, the master computer must meet the following requirements:

- The HAL must be the same as that of client computers.
- There can be only one disk partition and it must not contain any encrypted files.
- The disk partition size must be less than or equal to the client partition size.

The master computer disk partition contains the operating system and applications you intend to provide to clients. The size of this partition determines the minimum disk size required on client computers that receive the image you create from the master computer using Riprep.exe. During image creation, if a client computer hard disk contains more than one partition, RIS sends a message to the client desktop stating that only the system partition (the one with the Windows folder) will be copied from the master installation image. If the client's boot partition and system partition are different, the computer cannot be imaged.

For this part of your planning process, use job aid "Planning the Master Computer Configuration" (ACIRIS_03.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning the Master Computer Configuration" on the Web at http://www.microsoft.com/reskit) to indicate whether your master computer requires a hardware upgrade and if you need to verify that the master computer has the following:

- A HAL matching that of client computers receiving a Riprep image.
- A single disk partition the same size (or smaller) than the client computer disk partition, and with no encrypted files.

**Master Operating System Image Requirements**

To create installation images of an operating system, run the Riprep wizard on the master computer. When you are ready to create your image, run Riprep.exe from the master computer by specifying the following in the **Run** dialog box:

**\\RISServerName\Reminst\Admin\i386\Riprep.exe**

The master computer contains the operating system, locally-installed applications, and any configured system settings that represent a standard client configuration that you want to deploy. Plan to carefully test each configuration before running Riprep.exe to create the installation image. After Riprep.exe replicates the image to the RIS server, you cannot alter the image configuration without installing the image, making the alterations, and re-running the Riprep wizard to create a new image.

**Note**

- Overwriting an existing Riprep image on the RIS server is not recommended or supported. It is recommended that you delete the old image and make a new one.

It is likely that you will need to create multiple operating system images from your master computer. When this is the case, consider installing the operating systems you intend to image by using unattended installations. This way, you can create and maintain a separate answer file for each image. This makes it easier to recall the original configurations if you need to make changes. For more information about using unattended installations for your master computer, see "Configuring a Master Installation" later in this chapter.

**Note**

- At least one Risetup or CD-based image of the master computer operating system must exist on the RIS server.

For this part of your planning process, use job aid "Planning the Master Computer Configuration" (ACIRIS_03.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning the Master Computer Configuration" on the Web at http://www.microsoft.com/reskit) to indicate whether you plan to:

- Install different operating systems on the master computer using unattended answer files.
- Test the master installation prior to running Riprep.exe.

Earlier, in job aid "Planning for RIS Clients" (ACIRIS_01.doc), you recorded the operating system images you plan to make available to clients, along with application, driver, and desktop configurations.

**Master Computer Placement**

If you are placing a RIS server on the network, you might consider placing the master computer in close proximity to the RIS server because you might need to administer the images to create permissions configurations for clients. For more information about specifying security permissions on Riprep images, see "Evaluating Security for Operating System Images" later in this chapter. If you set up a lab for preinstalling and configuring operating system images on client computers prior to delivery to the client, you can place your master computer in any suitable location in the lab.

For this part of your planning process, use job aid "Planning the Master Computer Configuration" (ACIRIS_03.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning the Master Computer Configuration" on the Web at http://www.microsoft.com/reskit) to indicate the location for your master computer, either near your RIS server on the network, or in a lab.

# Assess Existing Network Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The impact that RIS-based operating system installations have on your network varies with the demand for RIS services. When large scale installations are in progress, your network carries a heavy load because multiple simultaneous operating system image downloads create significant increases in network traffic. For this reason, it is necessary to assess whether your network can support the speed and traffic volumes that permit RIS-based operating system installations in a reasonable amount of time.

To facilitate the assessment of your network, consider drawing a Visio diagram of the network. You can consult Active Directory and your DHCP scopes when determining the organization of the subnets and domains for your diagram.

Use the following guidelines to verify network requirements to support RIS-based operating system installations:

- **Network Topology.** Determine if your network topology can support a data transmission rate of at least 10 megabits per second (Mbps), but preferably 100 Mbps. To support the minimum recommended transmission rate, you need a 10 Mbps Ethernet network with supporting transmission media. For a 100 Mbps transmission rate, you need a minimum of a Fast Ethernet network with Category 5 UTP cable.

  **Note**

  - You cannot use RIS on a wireless network.

  You also need to determine if the network adapters of your RIS servers and clients can support the minimum data transmission rate.

- **Number of Servers.** Estimate how many RIS servers you need, based on the number of RIS clients in your organization. A single RIS server can use network pipes and allow 75 simultaneous instances of a Riprep or Risetup image to the equivalent number of clients. If you attempt to use a single RIS server to provide service to more than 75 clients, all subsequent PXE client requests for RIS service are ignored. If clients reboot after a TFTP time-out initiated by the RIS server, they are unable to obtain a client connection to that RIS server.

  **Note**

  - The limit of 75 clients per server is based on Microsoft internal testing. Your results might be different depending on your network topology.

- **Firewall**. Assess your need for a firewall. Because PXE does not provide any inherent security mechanisms, ensure that you have a correctly configured firewall for your network to prevent unauthorized access to your RIS server. For more information about RIS server security, see "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit).

- **Disk Image Transfer Time**. You must have a network connection to every RIS client computer. Ethernet local area networks (LANs) and Token Ring LANs are well suited for distributing disk images across a network. Wide area networks (WANs) are generally not fast enough, unless the LAN segments that make up the WAN are connected with a fast T-carrier service (T2 or higher). Digital subscriber line (DSL), cable modem, Integrated Services Digital Network (ISDN), and dial-up modem connections are not suitable for network distribution of RIS images.

  Table 4.3shows connection speeds and image transfer times for various network connections. Image transfer times are based on optimum network speeds only and are calculated for a 2.5 GB disk image. File server performance is not factored into the disk image transfer times. You can use Table 4.3 as a rough guide to help you determine whether your network is a suitable for RIS-based installations.

  **Table 4.3 Approximate Image Transfer Times for RIS-Based Installations**

| Connection Type | Network Speed | Transfer Time (2.5 GB Disk Image) |
|---|---|---|
| Fast Ethernet | 100 Mbps | 3 minutes, 25 seconds |
| Fast Token Ring | 16 Mbps | 21 minutes, 22 seconds |
| Ethernet | 10 Mbps | 34 minutes, 9 seconds |
| T2 | 6.312 Mbps | 54 minutes, 6 seconds |
| Token Ring | 4 Mbps | 1 hour, 25 minutes |
| T1 | 1.544 Mbps | 3 hours, 41 minutes |

- **Network Component Requirements**. Your network configuration must include a Windows Server 2003–based server, the DHCP service, DNS, and Active Directory. It is unnecessary for the RIS server to be the sole DNS/DHCP server or the domain controller.

  **Note**

  - It might be necessary to manually add a record to an existing DNS server to allow it to locate a Windows Server 2003–based server where Active Directory is located.

For this part of your planning process, use job aid "Planning the RIS Network Configuration" (ACIRIS_04.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning the RIS Network Configuration" on the Web at http://www.microsoft.com/reskit) to indicate your decisions to:

- Render a network structure diagram.
- Upgrade your network components, including installation and configuration of a firewall.

In the job aid, also record your image transfer time requirements and include the personnel you want to assign to upgrading or other tasks.

Other network-related factors you should evaluate include:

- The network installation point for RIS servers.
- The redirection of RIS client requests.
- The role of routers in forwarding client DHCP requests.

# Evaluating Network Installation Points

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to determine suitable local or remote installation points on the network for each RIS server. It might make sense to have your RIS servers each placed locally in a separate LAN. For example, you might want to isolate the installation traffic by using bridges. If that is the case, make sure all components that redirect network traffic, such as a bridges, switches, hubs, and repeaters, can all support the minimum data transmission rate.

Also, it might be feasible to determine the locations of RIS servers based on the DHCP scopes in your organization. If that is the case, you might place your RIS servers on the DHCP servers throughout your enterprise. For more information about combining RIS and DHCP servers, see "Assessing RIS Server Placement" earlier in this chapter.

If you have a small LAN containing a single subnet and no router, a single RIS server can provide service to all PXE-enabled client computers as long as installation traffic does not exceed network bandwidth and server resource limitations. For routed environments, you can configure prestaged clients with settings that direct them to the nearest RIS server in proximity for service. Also consider having high-speed communication links between routers. In a branch office site connected by a slow WAN communication link, you might place a RIS server at the branch site to avoid overloading the WAN link. Because WAN connections are typically slow, avoid having RIS-based operating system installations cross this type of connection. If they do, RIS server traffic can quickly cause a bottleneck and produce unpredictable results.

For this part of your planning process, use job aid "Planning the RIS Network Configuration" (ACIRIS_04.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning the RIS Network Configuration" on the Web at http://www.microsoft.com/reskit) to indicate the following:

- The number of RIS server installation points you estimate.
- Whether you need to install RIS servers on existing software installation points or on existing DHCP servers.
- Whether RIS servers will be placed either locally or remotely (cross-domain locations with respect to RIS client locations on the network).
- The type of internetworking connection you plan to use for RIS servers.

# Redirecting RIS Client Requests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can perform RIS-based operating system installations across routers and domains. You can also perform RIS installations across Active Directory forests, providing that you configure trusts across forests. If you need RIS-based installations to cross routers in your organization, you must determine if your routers must be upgraded to support the data transmission rate. Also, for PXE-enabled clients to contact RIS servers located across routers, you must configure the RIS server IP address in the router IP helper tables. By configuring the router in this manner, a PXE-enabled client that broadcasts an IP address request is redirected to the specific RIS server listed in the DHCP scope. As a result, the client does not need to broadcast another DHCP discover packet to locate a RIS server.

DHCP allows you to configure options 60, 66, and 67 for directing PXE-enabled clients to a RIS server without having to update routers with the IP address of the RIS server. However, for reliable RIS service, avoid this configuration because it is known to not function properly for PXE 1.0 and RIS boot floppy disk clients. This configuration also has other drawbacks.

For more information about using options to redirect PXE-enabled clients, see article Q259670, "Using Dynamic Host Configuration Protocol Options 60, 66, 67 to Direct PXE Clients to RIS Servers May Fail" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For this part of your planning process, use job aid "Planning the RIS Network Configuration" (ACIRIS_04.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning the RIS Network Configuration" on the Web at http://www.microsoft.com/reskit) to indicate whether you plan to update your router address tables with the IP addresses of RIS servers.

# Forwarding Client DHCP Requests through Routers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because client service requests are based on the DHCP discovery process, configuring your network to support RIS-based operating system installations across routers has the same requirements as configuring your network to support DHCP across routers.

Routers that you configure to forward DHCP broadcasts also automatically forward client service requests, however, you must ensure that the requests are forwarded to the proper RIS servers in addition to any DHCP servers. Depending on the model in use and the specific configuration, your router might support DHCP broadcast forwarding to a subnet, a specific host, or another router interface. If you use Windows Server 2003 DHCP but you place your RIS servers on separate computers, or if you use a third-party DHCP service, you must ensure that the routers forward DHCP broadcasts to both the DHCP and RIS servers. Otherwise, the client does not receive a reply to its remote boot request.

**Caution**

- You need to enable DHCP on all routers along all router hops between RIS servers and clients.

For this part of your planning process, use job aid "Planning the RIS Network Configuration" (ACIRIS_04.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning the RIS Network Configuration" on the Web at http://www.microsoft.com/reskit) to indicate whether you plan to configure your routers to forward DHCP broadcast traffic.

# Planning RIS Network Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

During RIS-based operating system installations, you need to maintain network security. The elements you need to consider when planning security for your RIS network include those that relate directly to the network, the client, server authorization, and administrative tasks. To plan for securing your RIS server on the network, address the following issues:

- Security risks of your PXE environment.
- NTLM authentication protocol level needed to log on securely over the network.
- Security for non-prestaged RIS clients.
- Enhancement of network security by using prestaged RIS clients.
- Restriction of client installation options.
- Control of the user interaction level during installation.
- Security for operating system images.
- Security for RIS server authorizations.
- Planning security for RIS administrative tasks.

For a job aid to record your planning decisions for RIS server security, see "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit).

# Assessing the Security of the PXE Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because of the design of PXE architecture, the PXE environment can introduce some inherent security risks in a network containing a RIS server, as follows:

- PXE has no provisions to detect or prevent unauthorized installations on PXE-enabled client computers from an unknown server. Any server that establishes a connection with a PXE-enabled client can perform an installation on the client computer.

- PXE has no provisions to prevent packet spoofing. As a result, an attacker could send malicious packets to integrate into the client installation.

- PXE cannot prevent unknown PXE-enabled computers on the network from receiving a remote operating system installation from a RIS server.

  This last risk is offset by the fact that RIS provides service only to users who log on with valid user credentials. In addition, if you prestage your client computers in Active Directory and configure your RIS server to only respond to known clients, a PXE-enabled intruder gaining access to your network cannot receive an operating system installation or any information about your client computer configurations.

To minimize the potential for successful attacks on your PXE-enabled clients, plan to take the following steps to ensure that unauthorized users cannot connect to them:

- Install and configure a firewall on your network.

- Implement safeguards, such as auditing and monitoring, to detect intrusions on your network.

- Secure physical access to your network.

- Enforce a strict password policy throughout your network.

For this part of your security planning process, use job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to indicate the steps you choose to take to secure PXE-enabled clients. Record this information under the section "PXE Environment Security."

# Evaluating the NTLM Authentication Level

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As part of planning for RIS server security, you need to evaluate which level of the NTLM challenge/response authentication protocol you require in your network. RIS can use either of two versions of NTLM to support RIS client network logons, including NTLM (the first version) and NTLMv2. NTLMv2 is inherently more secure than NTLM because of the way it handles encryption keys.

The NTLM version you choose affects the authentication protocol level that clients use, the level at which the protocol negotiates session security, and the authentication level that servers accept. For more information about choosing the most appropriate LAN Manager authentication level in a network that includes RIS, see "Setting the LAN Manager Authentication Level on a network that includes RIS" in Help and Support Center for Windows Server 2003.

When determining the most appropriate version of NTLM for your network, consider the following:

- **The network logon security level you need**. If you choose the highest level of security by using the **Send NTLMv2 response only\refuse LM & NTLM** option, then only NTLMv2 is used. However, when this is the case, you must ensure that all computers involved in the authentication process are running software that supports NTLMv2. If you choose the lower security level by using the **Send NTLM response only** option, NTLMv2 is used wherever possible and NTLM is used only when authenticating computers do not support NTLMv2.

- **The various operating systems you are running**. The NTLM version you choose can affect the ability of Windows 2000 Server, Windows 2000 Professional, Windows Server 2003, and Windows XP Professional computers to communicate with Windows NT 4.0 and earlier clients over the network. For example, computers running versions of Windows NT 4.0 earlier than SP4 do not support NTLMv2 and computers running Windows 95 or Windows 98 do not support any NTLM version.

For this part of your security planning process, use job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to indicate the NTLM authentication level you want to use, along with the platforms in use that support the NTLM level you choose. Record this information under the section "NTLM Authentication Level."

# Assessing Security for Non-Prestaged Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Providing RIS-based operating system installations to RIS clients that are not prestaged could pose a security risk. To service these clients, you must configure your RIS server to respond to all clients that request service. In this situation, the RIS server does not discriminate between authorized and unauthorized clients making service requests. This could expose your network to malicious clients.

For this part of your planning process, use job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to record your choice to configure non-prestaged RIS clients with the right to join the domain or check the box indicating that you will allow Remote Installation to create computer accounts.

# Planning for Network Security Enhancement Using Prestaged Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can enhance the security of a network that contains a RIS server by prestaging client computer accounts in Active Directory. By using client computer accounts prestaged in Active Directory and configuring your RIS server to respond only to these known clients, you ensure that unauthorized clients do not receive an operating system installation. You also make sure that the prestaged clients are serviced only by authorized RIS servers.

To prestage client computer accounts in Active Directory, you must obtain the UUID for the client computer and specify it when you create the client computer account. For more information about the requirements for prestaging client computers, see "Evaluating the RIS Client Prestaging Process" earlier in this chapter. For more information about how to prestage client computers, see "Evaluating the RIS Client Prestaging Process" earlier in this chapter and "Designing the Active Directory Infrastructure" later in this chapter.

If you want to optimize security using prestaged RIS clients, plan to do the following:

- Obtain the UUIDs for client computers and prestage client computer accounts in Active Directory.
- Configure users of prestaged client computers with read, write, and set or change password permissions on the prestaged computer account objects.
- Configure your RIS server to only respond to known (prestaged) clients by setting options in RIS server **Properties**.

For more information about how RIS servers respond to prestaged clients, see "RIS Server Configuration Design Tasks" later in this chapter.

For this part of your security planning process, use job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to indicate your choices to:

- Enhance security by prestaging client computers in Active Directory.
- Obtain UUIDs for client computers.
- Configure your RIS server to respond to known clients.
- Set user permissions on prestaged computer accounts to enhance security.
- Select the Active Directory domain or organizational unit to which these decisions apply.

# Assessing Security Benefits of Restricting Client Installation Options

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To enhance security, you can place restrictions on the client installation process by modifying Group Policy with specific RIS installation options. Group Policy applies to sites, domains, and organizational units. If you have personnel designated to deal with Group Policy issues in your organization, you can flag this as a task they need to perform.

You can use the Group Policy Object Editor MMC snap-in to alter the choices the CIW displays to a particular user or user group. You can configure these choices in the Default Domain Policy or you can create new group policies for specific groups of users that require certain installation options.

If you want to enhance security using Group Policy to modify how the CIW displays installation options to the client, plan to use some of the following RIS-specific Group Policy options:

- **Automatic Setup**. Accommodates an automatic setup process using predefined computer names and locations within Active Directory for client computer accounts. Include this option for client computers you are prestaging in Active Directory to enhance network security or for non-prestaged clients for which you predefine a computer naming format on your RIS server.

  **Note**

  - Under this option, if a UUID for a client is not found in Active Directory, the client computer receives a name based on the automatic computer naming format you configure in RIS server **Properties**. Also, the computer account is created in the location you specify in RIS server **Properties**.

- **Custom Setup**. Allows users to define a unique name for their computer and specify where to create the computer account within Active Directory. Include this option for clients you are not prestaging in Active Directory and for client computers that are to be set up by you or someone else during installation. This is a less secure configuration because the RIS server must be configured to recognize any client requesting service. For more information about defining CIW setup options in Group Policy, see "CIW Design Tasks" later in this chapter.

- **Restart Setup**. Allows users to restart an operating system installation attempt if it fails prior to completion. It is best to include this option only with prestaged clients because it is less secure to make multiple installation attempts available to unknown clients.

- **Tools**. Allows users to access tools, including the Recovery Console, from the CIW. Depending on which ISV and OEM tools are installed in your RIS server RemoteInstall share, you might want to limit which RIS clients can access them.

You also need to evaluate whether you want to apply the Group Policy as the default domain policy, or if you need to create new Group Policy objects for particular user groups. These choices are closely associated with how you configure the RIS deployment mode and the CIW. For more information about defining Group Policy, see "Designing the RIS Deployment Mode" later in this chapter.

For this part of your security planning process, use the "Security Enhancement with Group Policy and User Interaction Level Control" section of job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to record your decision to use Group Policy to enhance security. Also indicate if you want to use either the default domain policy or new Group Policy objects. If you plan to use new GPOs, then indicate the user groups where they apply.

# Assessing Security Benefits of Controlling the User Interaction Level

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can enhance the level of security in your network by managing client interaction with RIS-based operating system installations. You can predetermine how much user interaction occurs by modifying the CIW configuration. If you want to enhance security by controlling the user interaction level, you might plan to modify the CIW in the following ways:

- Add or remove entire CIW screens.
- Add or remove individual options within CIW screens.
- Provide additional text or instructions to users.
- Create new screens that prompt the user for specific information that you use to control image installation.

For example, you might add a new screen that prompts users to provide additional security information. Also consider that for any entry that you can specify in an answer file, you can create OSC variables in an answer file to capture information that users input in response to CIW prompts. By implementing modifications to the CIW, you can limit the information presented to potential unauthorized clients or require them to provide specific information that ensures they are valid RIS clients. For more information about defining and modifying the CIW configuration, see "CIW Design Tasks" later in this chapter.

You might also consider whether different user groups should receive different modifications to increase the security of RIS installations.

For this part of your security planning process, use the "Security Enhancement with Group Policy and User Interaction Level Control" section in job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to record your choice to use the CIW to control the user interaction level. Also indicate the types of controls you want to use and the user groups to which they will be applied.

# Evaluating Security for Operating System Images

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As part of planning RIS installation security, you need to evaluate how you intend to control which operating system images you make available to specific RIS clients. If you want to control which clients can access a particular operating system image, plan to do the following:

- Configure user permissions on each image you create, to define which users can install a particular image from the RIS server.
- Remove specific users from the access control list (ACL) on the operating system image folder on your RIS server to prevent these users from viewing (and therefore accessing) the image.

By setting permissions in the ACL of the answer file associated with an operating system image, you can prevent certain users from installing the image. By this means, you can also configure which users can install the image. If you do not set specific permissions on the answer file, then all users can install the image. If you remove a user account (or the group account to which it belongs) from the ACL on the operating system image folder, you disable a user's ability to view an image.

If you intend to use the default answer file with your Riprep or Risetup images, you need to set permissions on the Ristndrd.sif file. Otherwise, you need to set permissions on any custom answer files you create and associate with operating system images you want to configure for access control.

**Note**

- To enable a RIS user to view and subsequently install an operating system image from the CIW, you need to provide Read permission on both the answer file and the operating system image folder on your RIS server.

For more information about making images available to RIS clients, click the Index button in Help and Support for Windows Server 2003 and in the keyword box type **Remote Installation Services**, then select **Best Practices**.

For this part of your security planning process, use the "Operating System Image Security" section of job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to indicate your decision to control access to operating system images by modifying the ACLs of the default or custom answer files. Also indicate whether you want to disable the user capability of viewing and installing an image and the users/groups to which this decision applies.

# Assessing RIS Server Authorization Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a RIS server attempts to start on the network, Active Directory checks the RIS servers IP address against a list of authorized RIS servers. If a match is found, the RIS server is authorized to provide service on the network, otherwise, the RIS server is not authorized and cannot answer client service requests.

Part of your planning process for RIS server security involves assessing how you plan to authorize RIS servers on your network. You must authorize every RIS server in Active Directory to prevent unauthorized servers from servicing RIS clients on your network. The factors to consider when assessing the means for authorizing your RIS servers include:

- Who you designate to perform RIS server authorizations.
- Which computer you use to perform authorizations.
- How you perform the authorization of RIS servers.

**Understanding RIS Authorizers**

The person who authorizes RIS servers must be logged on as a member of the Enterprise Admins group. You can perform this task as an Administrator, but you might also consider delegating this task to qualified personnel to whom you give Administrative credentials. You might create a special security group to handle this task and add it to the Enterprise Admins group.

For this part of your security planning process, use the "RIS Server Authorization" section in job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to indicate your decision to assign the task of authorization to specific personnel. Also record whether you plan to create a special security group for RIS server authorizers or add the user accounts of authorization personnel to the Enterprise Admins group.

**Understanding RIS Authorization Locations**

You can authorize a RIS server from Active Directory Users and Computers MMC snap-in extension (Dsa.msc) on the RIS server itself or you can do so through a server running Windows Server 2003 or Windows XP Professional Remote Desktop session to the RIS server. All the RIS administrative tools, such as the Active Directory extension, are included when you create a RIS server on a computer running Windows Server 2003.

### Note

- In Windows 2000, you can administer a RIS server remotely using a Terminal Session in administration mode.

Alternatively, you can authorize a RIS server from a computer running Windows XP Professional. However, to do this you will need to install the Administrative Tools package on the computer running Windows XP Professional. You can install this package using the adminpak.msi application which is located in the System32 directory of computers running Windows Server 2003.

For this part of your security planning process, use the "RIS Server Authorization" section in job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to indicate the authorization location for your RIS servers.

**Understanding RIS Authorization Methods**

The methods you can use to perform authorization of a RIS server include using the **Verify** function in RIS server **Properties**, running Risetup at the command line with the /Check argument, or using the DHCP snap-in on a computer running Windows XP Professional. To use the **Verify** function or Risetup at the command line, you must be logged on at the RIS server and belong to the Enterprise Admins group. To use the DHCP snap-in on a computer running Windows XP Professional, you need to install the Administrative Tools package on that computer using the adminpak.msi application.

Note that you use this same DHCP snap-in to authorize DHCP servers. Therefore, if you install RIS on a DHCP server, which is already authorized in Active Directory, it is unnecessary to re-authorize the RIS server.

### Note

- The authorization process does not depend on how you combine or separate RIS and DHCP, nor on whether or not you use Windows Server 2003 DHCP.

For this part of your security planning process, use the "RIS Server Authorization" section in job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to indicate the authorization method you plan to use for your RIS servers.

# Planning Security for RIS Administrative Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To secure RIS administrative tasks, you need to decide whether you are planning to delegate tasks and then consider the best way to accomplish this securely. Also, you need to seriously consider securing the use of administrator credentials when performing administrative tasks.

For security reasons, when Windows XP Professional Service Pack 1 (SP1) or Windows Server 2003 is installed from the RIS server, the Administrator account is disabled as soon as the client computer is joined to the domain. Also, the Domain Admin group is added to the local computer when it is joined to a domain. If you want to prevent the Administrator account from being disabled when the client computer is joined to a domain, remove the entry DisableAdminAccountOnDomainJoin from the .sif file for that RISetup image.

**Assessing Delegation of RIS Administrative Tasks**

If you plan to delegate any RIS administrative tasks, you need to decide how to do this while maintaining security in your network. The best way to delegate RIS administrative tasks is to use existing security groups or define new ones for which you configure the appropriate permissions to perform specific RIS administrative tasks. This allows you to delegate tasks such as managing client installation images, managing prestaged computer accounts, and authorizing and configuring RIS servers.

For example, to install a RIS server and authorize it to Active Directory, the installer must be a member of the Enterprise Admins group. Others, who are responsible for configuring RIS servers and creating installation images, can have user accounts in Enterprise Admins or in another administrative group such as Domain Admins. This ensures they can perform all RIS configuration tasks.

If you have people in your organization who manage accounts and permissions, but do not configure RIS servers or create client installation images, you might make them members of the Account Operators group. You can then grant them folder permissions on the RIS server to perform their management tasks, rather than making them members of the Domain Admins or Enterprise Admins groups. This approach conforms to the best practices principle of granting permissions only where needed.

To set up a new security group for RIS tasks, create an administrative group in Active Directory, add qualified administrative personnel to the group, and then designate the appropriate permissions for RIS tasks. You can set permissions on the RIS server computer account object in Active Directory using the **Remote Install** tab in your RIS server **Properties**.

For more information about permission requirements for RIS tasks, see "Set permissions for administrators who manage client installation images for RIS" in Help and Support Center for Windows Server 2003.

For this part of your security planning process, use the "RIS Administrative Task Security" section of job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to record whether you want to delegate RIS administrative tasks and whether you want to use new or existing security groups. You can also specify the personnel to which you delegate the tasks.

**Assessing Security for RIS Administrative Tasks**

To minimize security risks, consider not logging on to your computer with administrative credentials to perform RIS administrative tasks. Instead, you and other RIS administrators can log on with a domain user account and use the **Run as** command to accomplish your administrative tasks. For this reason, consider creating alternate user accounts for all your RIS administrators. In addition, strongly consider training all your RIS administrators in the use of the **Run as** command, so they can perform RIS administrative tasks securely.

**Run as** enables you to run various programs and wizards under your administrator account and security context while you are logged on with a different account, such as that of a domain user. This allows you to expose your administrative context only for the specific program you are running and only for the duration of program execution. For more information about the security risks associated with logging on to the network as an administrator, see "Groups and Default Security Settings" in Help and Support Center for Windows Server 2003.

**Note**

- As a domain administrator, you should seriously consider using **Run as** to accomplish administrative tasks securely. For example, if you run your computer with domain administrator credentials, your Active Directory domain and forest are susceptible to Trojan horses and other attacks that target the logon sequence.

You can access the **Run as** command by using the command line or the user interface:

- **User interface**. In the Windows user interface, you can right-click the executable program (.exe), Control Panel (.cpl) item, or MMC (.msc) console you want to run, then select **Run as** and provide a user account and password.

- **Command-line**. You can use the **runas** command to provide the same capabilities as the "Run as" command in the user interface. For **runas** usage instructions, type the following syntax at the command line (cmd.exe):

```
runas ?
```

For this part of your security planning process, use the "RIS Administrative Task Security" section of job aid "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit) to record whether you want to:

- Use **Run as** to secure administrative tasks.

- Create alternate domain user accounts or a special group account for RIS Administrators who will use the **Run as** command.

- Train RIS administrators in the use of "Run as."

# Designing RIS-based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

At this point, you can pass your planning job aids to the design team to serve as inputs to the design process. The tasks you must accomplish at this stage of your RIS deployment are illustrated in Figure 4.4.

**Figure 4.4 Designing RIS-based Installations**



# Designing the RIS Installation Type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can design RIS-based installations using either Riprep or Risetup or both. Riprep installations are based on one or more file system images. Risetup installations are based on one or more distribution folders that contain a CD-like structure for installing the operating system. You can host both Riprep and Risetup images on the same RIS server.

# Design a Riprep-Based Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are designing a Risetup-based installation, you can skip this section and proceed to "Design a Risetup-Based Installation."

Use a Riprep-generated image if you want to distribute an image of a fully-configured workstation complete with applications. A Riprep image is essentially a file system image that is located on a remote RIS server. It is similar to the hard disk-images you create using a third-party disk-imaging tool and the Windows System Preparation tool (Sysprep).

You create Riprep images by running the Riprep wizard (Riprep.exe) on a master computer which has the operating system configuration, applications and settings, and desktop customizations you want to deploy to client computers in your organization.

Riprep images are most useful for cloning a standard operating system configuration to clients. Riprep images generally require more disk space on your RIS server than Risetup images because they usually include preconfigured applications and tools. However, they install faster than equivalent size Risetup images.

# Riprep Image Design Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To create a Riprep-based installation, you must first set up a *master installation*. This is the reference computer that contains the operating system, software applications, and configuration settings you plan to install on destination computers in your organization. After you configure the master installation, you run Riprep, which is on the server at the following location: *servername*\reminst\admin\i386\riprep. This converts the master installation into a remote installation image — a functionally identical replica of the master computer disk — that you can install on multiple destination computers. Riprep also replicates the image to a RIS server where it is available for installation on remote-boot-enabled client computers. Clients who request installation of an operating system can access Riprep-based images on a remote RIS server if you configure them to do so.

The best way to install the operating system on your master computer is to use RIS with an unattended installation. For more information about setting up a master installation, see "Configuring a Master Installation" later in this chapter. However, you can also install the operating system locally using the appropriate operating system CD. If you do this, use the disk partitioning utility found on the Windows Server 2003 installation CD and use the text mode setup to clear the disk partition and ensure a clean installation.

After installing the operating system on the master computer, you can install any applications needed by your clients, including line-of-business applications. Before running Riprep to create an image, it is prudent to test the master installation to verify proper configuration and functioning.

Riprep configures various operating system settings on the master computer to ensure that every copy of the master computer's disk image is unique when you install it on destination computers. This includes resetting the security identifiers (SIDs) and ACLs. Riprep also configures the master installation image so that, after the initial installation of the image, every destination computer starts in a special setup mode known as Mini-Setup.

To create a Riprep image from a master installation and store it on a RIS server, your master computer must meet the requirements described in "Assessing Master Computer Requirements" earlier in this chapter. In addition, the following apply:

- You must have at least one Risetup image stored on the RIS server that matches the operating system on the master computer, from where you create the Riprep image.

  The default answer file (Riprep.sif) that Riprep generates for the image refers the client computer to the RIS server to obtain drivers that start the text-mode portion of the CIW during installation of the operating system image.

- The Risetup image on the RIS server must use the same language and have the same first two designations in the version number. For example, a 5.1.2600.0 image will work for a 5.1.2600.1106 version of the same SKU. Only the first two numbers and the SKU type are checked (5.1 Professional) to verify that they are the same as the master computer.

- During image installation, Mini-Setup uses Plug and Play to detect hardware differences between the master and destination computers. Therefore, identical hardware is not needed except for the HAL type. You can only perform a RIS-based installation if the HAL in the RIS (Riprep) image is compatible with the HAL on the destination computer.

- The Riprep wizard only supports preparation and replication of images from the system partition (C:\ *system partition*) on the master computer. The master computer must have a partition no larger than the partition on the destination computer.

In addition, RIS uses the unattended installation process, which means it uses text files, known as setup information (.sif) files, to store the configuration settings for installation images. Each time the client uses the CIW to choose an operating system image, the installation processes the information contained in the .sif file. By modifying the contents of a .sif file associated with an installation image, you can predefine the configuration settings for that image. For more information about software configurations and Risetup images, including modifying answer files, see "Riprep Image Design Tasks" later in this chapter. For more information about defining the CIW configuration, see "CIW Design Tasks" later in this chapter.

# Riprep Image Design Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Riprep Image Design Tasks**
When designing a Riprep-based installation, your primary tasks are to define how to create and use the master installation. Accomplishing these tasks involves choosing the following:

- The operating systems you want to image.
- The software applications you want to install and configure in the master installation.
- The special hardware drivers you want to include in the master installation.
- Operating system and software settings you want to provide.
- Whether to reuse the master installation to create new images with differing operating system and application configurations. It is not recommended that you reuse the master installation more than two times because Windows Product Activation can only be reset three times.

For a job aid to record your design decisions for Riprep images, see "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit).

**Operating System Images**
To begin defining your operating system images, make a copy of job aid "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit) for each operating system you want to deploy. Under "Operating System Image," create an identifier for each image and enter it under "Image ID Number." Also enter the product name and the full version number of the operating system for each Riprep image. For Windows Server 2003 and Windows XP Professional operating systems, you can find the version number by viewing the properties of the file Ntdll.dll located in the i386 folder on the operating system CD.

You can locate the version number on the **Version** tab of the **Properties** dialog box. The version number is in the format 5.1.XXXX.Y, where XXXX is the build number. Be sure you include the build number in the worksheet.

**Software Configurations for the Master Installation**
Your primary task after installing a master computer operating system is to define which software applications to include in your master installation image. You can use your software inventory in conjunction with the following design guidelines to define a software configuration for each Riprep image you need to create.

**Identify core applications**
Choose the applications you want to include in your master installation image for the client computers and record the application names in job aid "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit) under "Software to Install with This Image."

**Identify service packs, hotfixes, and patches**
It is a good idea to configure and include service packs, hotfixes, and patches in your master installation images. Because service packs can take a long time to manually install, it is more efficient to include them in the master installation image. Record the names and versions of service packs, hotfixes, and patches under "Operating System Image" in the appropriate copy of job aid "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit) for each operating system image. Keep in mind that if you intend to install any applications after copying a disk image to a destination computer, you might also have to install service packs, hotfixes, and patches after copying the disk image to the destination computer. If you do, record this information under "Software to Install After Image Copy" in each copy you use of job aid "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit).

**Identify applications you cannot install on a master installation image**

In some cases, you must install and configure certain applications only after you copy a disk image to a destination computer. You cannot add applications such as these to a Riprep image. The following are examples of these types of applications:

- Programs that depend on Active Directory, such as Message Queuing (also known as MSMQ) or Microsoft Exchange.
- Special server applications, such as Certificate Services and Cluster service.
- Special applications you want to install only on certain computers.

When you have applications you need to add and configure after installing an image, you might be able to use a Risetup-based installation with an answer file that calls Cmdlines.txt or uses the [GuiRunOnce] section. For more information about designing a Risetup-based installation, see "Design a Risetup-Based Installation" later in this chapter.

If these methods do not work for your application configuration, you can use Sysprep rather than RIS. With Sysprep, you can stage a group of applications on your master computer and create separate answer files that manage the installation of specific application configurations. For more information about using Sysprep for image-based installations, see "Designing Image-based Installations with Sysprep" in this book.

**Software Application Configurations and Separate Images**

When creating Riprep images, it is necessary to create a separate image for each group of computers on which you want a different software configuration. You might need to create separate images when conditions such as the following exist:

- A group of computers requires software applications, utilities, or tools that conflict with software programs required for other computers. For example, you might need to create a separate image for portable computers that require the same vendor-specific programs, such as power-management utilities, DVD codecs, or diagnostic tools.
- A group of computers requires software applications that you cannot automatically install and configure after copying the image to destination computers. For example, you might want to create a separate image for Web servers that all run the same suite of third-party data analysis and monitoring applications. This ensures that all of your Web servers have a consistent configuration and also eliminates the need to manually install and configure third-party applications after you copy an image to a Web server.
- A group of computers requires frequent installation of the same unique software configuration. For example, you might want to create separate disk images for trade-show kiosks or computers that you use for training purposes, because these computers require frequent reinstallation.

For these cases, you can create separate images containing unique application sets. Use job aid "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit) to record your application configurations under "Software to Install with This Image."

**Hardware Configurations and Separate Images**

It is unnecessary to provide separate images for groups of computers on which hardware is detected by Plug and Play. In this case, you do not need to have the same hardware on your destination computers as you configured in your master installation image. However, certain groups of computers in your organization might have special hardware that Plug and Play does not detect. To ensure that the drivers for this hardware are properly detected, you need to add the drivers and any support files to your Riprep image. You can only perform a RIS-based installation if the HAL in the RIS (Riprep) image is compatible with the HAL on the destination computer.

For this part of your Riprep image design process, you need to identify all the unique hardware driver and support files that you need to provide to clients. For each unique configuration, you must create a separate master installation image. You can record the hardware driver and support file information in job aid "Defining Riprep Images" (ACIRIS_06.doc) (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit) on the *Windows Server 2003 Deployment Kit* companion CD under "Hardware Drivers to Install with This Image."

**Defining the path to special hardware drivers in Riprep answer files**

When Plug and Play runs, it enumerates the hardware to get the appropriate IDs and then tries to match the existing hardware with the correct drivers. To find the device driver search path, Plug and Play checks a registry entry that stores the default path to the location of all drivers. The registry entry is **Device Path** in:

HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion

The default device path is %windir%\System32\Drivers\. If you have special or legacy hardware that requires drivers that do not exist in the default device path, you can use your riprep.sif answer file to add a new device search path to the registry entry. This enables Mini-Setup (and each subsequent startup sequence of client computers) to find the special hardware drivers.

To add the new device search path to the registry entry, you must specify the entry **OemPnPDriversPath** in the [Unattended] section of the riprep.sif answer file along with the path to the folder containing the drivers. You can include your special drivers, along with any catalog or .inf files, on your system root partition in a folder with a name such as \Drivers.

Riprep prepends %SYSTEMDRIVE% to the folder name, therefore you must ensure that you place your \Drivers folder in the system root partition. Record the special hardware device search path in job aid "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit) under "Special Hardware Driver Path."

**Defining paths to multiple hardware drivers**
If you have different driver types, you can create subfolders under the Drivers folder, which Mini-Setup will automatically search to find all drivers. For example, if you create an Audio subfolder and a Video subfolder, the OemPnpDriversPath entry in your Riprep.sif answer file in the [Unattended] section is as follows:

```
OemPnPDriversPath=Drivers\Audio; Drivers\Video
```

**Operating System and Software Settings**
You can configure most operating system and software settings on the master installation before running Riprep. Types of settings you can configure include the following:

- Local policy settings, such as Group Policy Administrative Template settings.
- Control Panel settings, such as power options, sound scheme settings, system startup and recovery options, system performance settings, and accessibility options.
- Internet Explorer settings, such as the default home page, security and privacy settings, and connection settings.
- Optional Windows components settings, such as network monitoring tools, Remote Storage, and Services for NetWare.
- Services settings, such as startup type, logon accounts, and recovery actions.
- Desktop settings, such as desktop shortcuts, folder options, and fonts.
- Display appearance and settings.
- Microsoft Word or Microsoft Excel settings, such as view, edit, save, and spelling options.

For this part of your Riprep image design process, use the "Operating System Settings" section in job aid "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit) to record the settings you intend to include for each master installation image you create.

**Reusing the Master Installation and Windows Product Activation**
If you intend to use Riprep images for RIS installations, it is best to use volume license (VL) media for installing your master computer operating system. After running Riprep on a master installation you create with a VL media source, you can reuse the same installation after rebooting and running Mini-Setup, as indicated by the Riprep wizard. At this time, you can install and configure additional applications and run the Riprep wizard on the master installation again. However, if you are unable to use VL media for some reason, ensure that you always clean-install operating systems on the master computer using CD-type images from RIS, or by using a non-VL media CD-ROM. After installing the master computer operating system, you can install and configure applications and run Riprep again.

If you do not have VL media for the Windows XP Professional or Windows Server 2003 operating systems, from which you intend to make Riprep images, you also need to be concerned about Windows Product Activation (WPA). WPA includes an activation timer that is reset each time you create a Riprep image from a master installation configured with a non-VL source. The activation timer controls the grace period during which users must activate their operating systems after installing the Riprep image by using RIS.

**Note**

- The activation timer reset behaves similarly on Riprep images that you *restore* for the purpose of creating a new image.

If you used non-VL media and you attempt to run Riprep on the same master installation more than three times, a warning displays indicating that the activation timer failed to reset. As a result, RIS-based installations of master images of the fourth generation or greater do not have the grace period reset. Therefore, if users install these images after the grace period, they are required to activate Windows at the first logon. If they do not, they are automatically logged out.

**Note**

- The WPA issues described here do not apply to flat images that you create with Risetup.

For this part of your Riprep image design process, use the "Operating System Image" section of job aid "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit) to record the media type you intend to use to generate Riprep images.

# Riprep Image Design and User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When designing a Riprep image, be aware that the profile of the user that installs applications and makes configuration changes in the master installation has an impact on the users of client computers where you install the Riprep images. This is a primary concern for the functioning and availability of applications and configurations that are not compliant with Windows Server 2003. For example, some applications might rely upon per-user configurations that are specific to the profile of the user (usually the Administrator) installing the applications prior to running Riprep, rather than to all users of the client computer where the Riprep image is installed.

For this part of the Riprep image design process, use the "User Profiles and Applications" section in job aid "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit) to record any applications included with your Riprep image that are not compatible with Windows Server 2003. Also include the name of the user profile under which you intend to install applications and make configuration changes. For more information about the effect of user profiles on Riprep images, see "Testing Riprep Images and User Profiles" later in this chapter.

# Design a Risetup-Based Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Risetup allows you to create a CD-type image to be used with RIS installations on client computers. When designing a Risetup-based installation, use the job aid "Planning for RIS Clients" (ACIRIS_01.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Clients" on the Web at http://www.microsoft.com/reskit). as input to the design processes outlined in this section.

Use a Risetup-generated image if you want to distribute the network equivalent of CD-based installation functionality. A Risetup image is a replica of an operating system CD file structure, located across the network on a remote RIS server.

You create Risetup images by running the Risetup wizard on a RIS server, while using an operating system CD to create the image. When using Risetup images, you cannot provide a fully-configured clone of an operating system with applications and desktop customizations, as you can with Riprep images. However, you can add applications and drivers to the distribution folder where the Risetup images are located and use answer files to install the applications and specify the location of drivers.

# Risetup Image Design Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Installing a Risetup image is similar to setting up a workstation directly from a CD, however, the source files are located across the network on a RIS server. Figure 4.5 shows the directory structure of your RIS server under the RemoteInstall folder where Risetup images are stored. You can define the name of the folder <imagename> where the images are located.

**Figure 4.5 Directory Location for Risetup Images**




When you install remote installation services on your RIS server, it automatically creates one Risetup image of the server operating system and stores it under the Images folder within the RIS directory structure, as shown in Figure 4.5. This image is available to remote boot enabled clients. Clients that request installation of an operating system can access Risetup images on a remote RIS server, if you configure them to do so.

You can make additional Risetup images using operating system CDs for Windows 2000 Professional, Windows 2000 Server, and Windows 2000 Advanced Server, in addition to Windows XP Professional and Windows Server 2003. To create a Risetup image, place the CD in the CD drive of your RIS server and run the Risetup wizard from the **Images** tab of RIS server **Properties**. You can also specify the following command string at the command line interface to start the Risetup wizard:

```
risetup -add
```

After creating a Risetup image, you can add applications and device drivers located in the RIS server directory structure, as described in "Riprep Image Design Tasks" later in this chapter.

**Note**

- The RemoteInstall folder, which is the parent folder of your RIS server directory structure, is shared as Reminst, so that images, applications, and drivers are available to RIS clients on this distribution share.

You can create and associate multiple answer files with Risetup images, which allows you to customize the applications and drivers you want to install with each image. However, you cannot include preconfigured application or desktop configurations with a Risetup image. Also, Risetup images take longer to install than an equivalent size Riprep image.

# Risetup Image Design Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Risetup Image Design Tasks**
When designing a Risetup-based installation, your primary tasks are to choose the following:

- The operating systems you want to image.
- The software applications you want to include on your distribution share.
- The special hardware drivers you want to include on your distribution share.
- The operating system configuration settings or components you want to provide.

For a job aid to record your design decisions for Risetup images, see "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit).

**Operating Systems**
To begin defining your operating system images, make a copy of job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit) for each operating system you want to deploy. Under "Operating System Image," create an identifier for each image and enter it under "Image ID Number." Also enter the product name and the full version number of the operating system for each Risetup image.

For Windows Server 2003 and Windows XP Professional operating systems, you can find the version number by viewing the properties of the file Ntdll.dll located in the i386 folder on the operating system CD. You can locate the version number on the **Version** tab of the **Properties** dialog box. The version number is in the format 5.1.XXXX.Y, where XXXX is the build number. Be sure you include the build number in the job aid.

**Software Configurations and Risetup Images**
In many cases, you can use a modified Ristndrd.sif answer file to automatically install and configure software applications after installing a Risetup image. The answer file provides the information necessary to do this at the end of Mini-Setup or after Mini-Setup finishes, which follows initial image installation on the client. To enable application installation at these times, you need to create an \$OEM$ folder at the same level as the i386 folder for the appropriate image in your RIS directory structure, and place your applications and related support files in a folder you create with a name such as Applications. The path to such a folder is as follows:

RemoteInstall\Setup*Language*\Images*ImageName*\$OEM$\$1\Applications

By staging your applications in this location and then using answer files to install the applications on the client, you can minimize the number of images you need to manage and simplify the modification of software configurations as group needs change. However, you need to maintain separate answer files that define the unique application installation configurations. The methods available to define these configurations in your answer files include:

- Pointing to Cmdlines.txt
- Specifying the appropriate entries in the [GuiRunOnce] section

Using these methods enables you to maintain a single image from which you derive multiple application (or driver) configurations that you can provide as installation options to meet varying client needs. You provide these options to the client by associating each answer file with the Risetup image and then configuring access control entries (ACEs) on each answer file to define which clients can receive specific image-answer file sets. For more information about defining operating system installation options, see "Interactive Installation Design Tasks" later in this chapter.

**Note**

- For application installation, you can also use Group Policy to configure computer or user deployments of applications.

For more information about Group Policy-based software deployment, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

If you intend to stage your applications, use job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit) to record your pool of applications under "Applications Staged in \$OEM$ Folder." Also, record the staging path to your applications under "Staging Path for Applications," using the name of the new folder you plan to create to contain them. In addition, use the "Answer File Versions" section of job aid ACIRIS_07.doc to record your answer file names and software configuration for each answer file. In job aid ACIRIS_07.doc, you can identify your Risetup answer files with any suitable names, such as Ristndrd01.sif and Ristndrd 02.sif.

To create the answer files, you can copy and modify the default Risetup answer file (Ristndrd.sif) from the following location after you create a Risetup image:

RemoteInstall\Setup*Language*\Images*ImageName*\i386\Templates\

To create new answer files or to modify existing answer files, you can use Setupmgr.exe or a text editor such as Notepad.

For further information about configuring answer files, see the Deploy.cab file in the \Support\Tools folder on your Windows XP Professional or Windows Server 2003 operating system CD. This file contains a complete reference to the entries you can add or modify for the [Unattended] section. However, be aware that not all sections and entries you find in Deploy.cab apply to RIS.

**Tip**

- Consider testing your answer file configurations with an actual operating system installation in your RIS test environment. This ensures proper functioning before performing RIS installations in your production environment.

**Hardware Configurations and Risetup Images**

It is unnecessary to create separate answer files that specify installation of hardware drivers for groups of computers on which hardware is automatically detected by Plug and Play. However, you might have special hardware on certain groups of computers in your organization that Plug and Play does not detect. To ensure that the drivers for this hardware are properly detected, you need to add the drivers and any support files to your RIS distribution share in an \$OEM$ folder you create, and point to this location in your answer files. For storing your driver files, you can create a subfolder with a name such as "Drivers" in the following path within your RIS directory structure:

RemoteInstall\Setup*Language*\Images*ImageName*\$OEM$\$1\Drivers

**Note**

- You can also structure the Drivers folder with subfolders; Mini-Setup searches all the subfolders when attempting to detect drivers for your hardware.

For this part of your Risetup image design process, you need to identify all the unique hardware driver and support files that you want to provide to clients. For each unique configuration, you must create a separate answer file that provides the commands to install the specific drivers. You can record the hardware driver and support file information in job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit) under "Hardware Drivers Staged in \$OEM$\$1\Drivers Folder." You can record the answer file names with the hardware driver configurations that they install under "Answer File Versions" in job aid "Defining Risetup Images" (ACIRIS_07.doc).

**Note**

- If you have a large number of special drivers, you can include them all in your Drivers folder. After installing the Risetup image on a group of client computers, you can have your answer file call Cmdlines.txt, which you can preconfigure with commands to remove the unused hardware drivers for that user group.

**Defining the path to special hardware drivers in Risetup answer files**

For special or legacy hardware that requires drivers that Plug and Play cannot detect, you can use your Ristndrd.sif answer file to add a new device search path. This path enables Mini-Setup (and each subsequent Plug and Play detection sequence on

client computers) to find the special hardware drivers.

As with Riprep images, you can specify the path to special hardware drivers by including the registry entry **OemPnpDriversPath** in your Ristndrd.sif answer file in the [Unattended] section as follows:

OemPnPDriversPath=Drivers\Audio; Drivers\Video

For more information about defining the path to special hardware drivers in Riprep answer files, see "Riprep Image Design Tasks" earlier in this chapter.

If you have multiple hardware configurations for different user groups that use the same Risetup image, you can create different answer files with different search paths to the required hardware drivers. In each answer file, add the [Unattended] section and OemPnpDriversPath entries to specify the different paths.

For this part of the image design process, identify the path(s) to your special hardware drivers using the new folder(s) you plan to create, and record this information in job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit) under "Staging Paths for Hardware Drivers."

**Installing Applications and Drivers**

After you install a Risetup image on a client computer, you can install special applications and drivers by calling the Cmdlines.txt file at the end of Mini-Setup or by using [GuiRunOnce] to provide commands that execute after Mini-Setup runs. To set up for application or driver installation, you need to create Cmdlines.txt and configure it with the commands you want to run, or use the [GuiRunOnce] section in your Ristndrd.sif answer file to run the commands that perform the installation.

**Using Cmdlines.txt**

Consider using Cmdlines.txt in your Risetup image design when you want to do the following:

- Install applications from the \$OEM$ folder on your RIS distribution share.
- Install applications at the end of Mini-Setup.
- Install an application that is designed for installation by one user, and you need to replicate user-specific information to all users of the computer.

However, before you decide to incorporate Cmdlines.txt into your Risetup image design, consider the following limitations:

- **Logon context.** When Cmdlines.txt is called, it runs as a system service.

  This means there is no logged-on user and no network connectivity. All user-specific information is written to the default user portion of the registry, and all subsequently created users inherit those registry settings.

- **Application support files.** You must place the files necessary for an application or utility to run in your RIS distribution share.
- **Quiet mode.** When you install applications by using Cmdlines.txt, you must install the application with an unattended setup so the user does not need to respond to messages about the application.
- **Limited applications.** You cannot install Windows Installer (.msi) applications with Cmdlines.txt. To install these types of applications, use [GuiRunOnce] instead.

For this part of the image design process, check the "Use Cmdlines.txt" check box in job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit), if you intend to use Cmdlines.txt to install applications or drivers.

**Creating Cmdlines.txt commands**

To use Cmdlines.txt for installing components such as applications and drivers, you first need to create the file and then add the specific command strings that start the installations. You create the Cmdlines.txt file with a text editor such as Notepad (or Setupmgr.exe). In the Cmdlines.txt file, use the following syntax to specify the commands you want to run to install applications or drivers:

```
[Commands]
```

```
"command_1"
"command_2"
     . . .
"command_x"
```

The command strings enclosed in quotation marks specify the commands that execute in order when Cmdlines.txt is called. For example, you could create a command string to run the commands in an .inf file as follows:

```
"%windir%\System32\Rundll32.exe ApplicationName.inf"
```

When you finish creating Cmdlines.txt, you need to place it in the following path in the RIS server directory structure:

RemoteInstall\Setup*Language*\Images*ImageName*\$OEM$\$1*Subfolder*

The folder *Subfolder* might be named Drivers or Applications, as described earlier in this section in the discussion about software configurations and Risetup images. You also need to place the application(s) you want to install (along with any support files) in the \$OEM$\$1\Application subfolder and the drivers you want to install along with support files in the \$OEM$\$1\Drivers subfolder. In addition, you must specify the following value in the [Unattended] section in your Ristndrd.sif answer file to point to the location of the applications, special drivers, and Cmdlines.txt:

```
InstallFilesPath=RemoteInstall\Setup\Language\Images\ImageName\$OEM$\$1\Subfolder
```

After you properly configure all these parameters, Cmdlines.txt will run automatically at the end of Mini-Setup and process each command that it contains to install the applications you specify.

For this part of your Risetup image design process, record the commands you intend to run in Cmdlines.txt under the "Cmdlines.txt Commands" section of job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit).

**Using GUIRunOnce**

Consider using [GuiRunOnce] in your Risetup image design when you want to do the following:

- Install applications or drivers from any source, including network shares, hard disks, or CD-ROM drives.
- Install applications or drivers, but you cannot use Cmdlines.txt; this is the case if you want to install applications that use the Windows Installer (.msi).
- Log on as a user and not automatically replicate application registry settings to future users.
- Require a user to be logged on to the computer when the application or driver installs.
- Control the order in which applications or drivers install.

However, before you decide to incorporate [GuiRunOnce] into your Risetup image design, consider the following restrictions:

- **Enabled log-on function.** [GuiRunOnce] requires enabling the log-on function.

  This requires you to set up the automatic logon process by adding the AutoLogon=Yes value in the [Unattended] section in your Ristndrd.sif answer file.

- **Login context.** [GuiRunOnce] commands run in the security context of the currently logged-in user.

  If the logged-in user does not have the permissions necessary to run the command completely, then the application fails to install. Because of running in this context rather than as a service, the registry entries that the application creates are for the current user rather than the default user. If you want settings to register only for the currently logged-in user, then [GuiRunOnce] might be appropriate. Otherwise, use Cmdlines.txt to run commands that install applications, because it runs as a system service.

- **Application restarts.** [GuiRunOnce] commands that install applications that force a restart might not complete unless you can suppress the restart.

  When a computer restarts, all entries in the [GuiRunOnce] section are lost, meaning that these commands do not execute. If the application does not provide a way to suppress restarts, you might be able to repackage the application into a

Windows Installer package (.msi) using Veritas WinInstall.

For further information about the Windows Installer, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

- **Shell not loaded.** [GuiRunOnce] does not function if an application requires installation of the Windows Explorer shell, because the shell is not loaded at the time GuiRunOnce commands execute.

- **Use of /Wait switch.** [GuiRunOnce] commands often require you to use of the /Wait switch to avoid application setup failures caused by multiple instances of the installation mechanism running.

  While Setup is running, initiating a second process and closing an active one might cause the next routine listed in RunOnce registry subkey to start. When this occurs, more than one instance of the installation mechanism is running, which usually causes the second process to fail.

For this part of the image design process, check the "Use GUIRunOnce" check box in job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit), if you intend to use [GUIRunOnce] to install applications or drivers.

**Creating GUIRunOnce commands**

The [GuiRunOnce] section of your answer file contains commands that execute the first time users log on to their computers after Mini-Setup completes. To set up this process, each command under [GuiRunOnce] creates an entry in the registry subkey Runonce when your answer file is parsed during image installation. Runonce is located in:

HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion.

When the user logs on after Mini-Setup finishes, this registry subkey is checked for any commands that should be executed.

For example, you might start the wizard that installs Active Directory by specifying the following command string in the [GuiRunOnce] section in your Ristndrd.sif answer file:

```
[GuiRunOnce]
"%windir%\system32\dcpromo.exe"  /answer:answer_file
```

Note that you must specify each command line in quotation marks under [GUIRunOnce]. Also, the dcpromo command requires an answer file, specified on the command line, to provide input to the wizard.

For this part of your Risetup image design process, record the commands you intend to run under the "[GuiRunOnce] Commands" section of job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit).

For additional information about GUIRunOnce, including controlling application installation, see the Deploy.chm help file in the Support\Tools folder on your Windows XP Professional or Windows Server 2003 operating system CD.

**Important**

- Deploy.chm and Ref.chm must be in the same directory for Deploy.chm to provide the most complete RIS help information. You can open these files on the operating system CD, or open them on your computer by installing Deploy.cab in a folder on your computer.

**Operating System Settings**

You can configure many operating system settings in your Ristndrd.sif answer file. For example, you can use the following entries in the [Components] section in your answer file to determine whether the components associated with these keys will be installed with the operating system:

- Calc

- Certsrv

- Deskpaper

- Fax

- MousePoint
- MSWordpad
- Paint

For this part of your Risetup image design process, record the components or configuration settings you plan to provide using Ristndrd.sif answer files in "Answer File Versions" under "Components/Configuration Settings" in job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit).

For further information about operating system settings you can configure in Ristndrd.sif answer files, see the Reference/Unattend.txt section of the Deploy.chm help file in the \Support\Tools folder on your Windows XP Professional operating system CD.

# Designing the RIS Deployment Mode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can design your RIS-based operating system deployment to use either the interactive or fully-automated mode. Interactive mode requires the user to respond to predefined options presented during installation at the client computer; the fully-automated mode requires no user intervention other than turning on the client computer and providing logon credentials.

You might want to provide interactive installations for more knowledgeable users who can easily provide the user input you require. For users that are less knowledgeable, you can provide fully-automated installations that require little or no input. This ensures that installation on these clients goes smoothly and does not require your presence at the client computer.

If you want to provide a combination of interactive and fully automated installations to RIS clients, you can prestage your client computer accounts in Active Directory by using the prestaging script along with an input file that specifies which Startrom.com boot file each prestaged client must use. To find the prestaging script, see the Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources.).

**Note**

- The RIS deployment mode design process is closely associated with the CIW design process. For further details, see the job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit).

# Interactive Installation Design Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An interactive installation requires users at client computers to press the F12 key to initiate a network service boot. This occurs automatically when you use the default Startrom.com boot file for RIS installations. After the client initiates the network service boot and obtains an IP address from DHCP, the CIW downloads to the RIS client. At this point, the information gathering and option choosing process of the CIW begins.

You can configure the CIW to display numerous required user input fields on a menu of installation options. You can also configure the CIW to present specific operating system installation options, such as a list of Risetup or Riprep images. These capabilities enable you to control the user input necessary to enable the installation and to designate who can access predefined installation configurations.

For example, you might allow one group of users to have access to all predefined operating system installation options, but restrict another group of users to a single installation option. You can achieve this type of control over the installation process by using Group Policy and ACL configuration.

# Interactive Installation Design Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When designing an interactive installation, your primary tasks are to choose the following:

- The information you require users to input to the CIW.
- The operating system installation options you want to make available to different users via the CIW.
- Which users receive specific RIS setup options that you configure in Group Policy.

For a job aid to record your design decisions for an interactive installation, see "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit).

**Defining User Input**

You can configure the user information that is required as input to the CIW. For example, the first screen presented by the CIW is the Welcome screen. You can modify this screen to include multilanguage choices or additional information such as a company message. The next screen is the Logon screen, which requires the user to provide logon credentials, including user name, password, and domain name. You can also build custom CIW screens to prompt the user for specific input information. For more information about defining the CIW configuration, including customizing CIW screens with language options and input prompts, see "CIW Design Tasks" later in this chapter.

At this point in your interactive installation design process, you might need to further analyze your user input requirements. If this is the case, you can wait until you review the section "Designing the CIW Process" later in this chapter before recording your design decisions. Otherwise, use the "Required User Input" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) to record the choices you make for required user input. You can also specify the users or user groups from which you require the input, under "Applicable User Groups."

**Defining Operating System Installation Options**

You can define which operating system images the CIW allows specific users to install, and you can limit the images the CIW displays to specific users by configuring ACLs. By providing ACEs on the answer files associated with a RIS installation image and by providing ACEs on the RIS server operating system image folder, you can determine which clients get to view and install a particular image.

For example, to enable a particular user group to install a RIS image, you must configure the group with Read permissions on the answer file and Read permissions on applicable RIS server operating system image folder. This causes the operating system images associated with that answer file to display in the CIW as installation options for the user group. If you specify Read permissions only on the answer file and not on the image folder, the option to install the image displays in the CIW, however, the image does not install due to the lack of Read permissions on the image folder. When configuring ACEs on the operating system image folder on your RIS server, use the following directory path:

*RISServerName*\RemoteInstall\Setup*Language*\Images*ImageName*\

You can also set an answer file permission configuration to Deny Read access to a particular user group. This causes all operating system images associated with that answer file to *not* display in the CIW as installation options for that user group.

When enabling access to operating system images, you will probably want to simplify administration. Therefore, the best approach might be to enable Read permissions for all users on the image folder and then use answer files to provide Read permissions to specific user groups. However, note that by configuring ACEs on the image folder, you can explicitly prevent certain users and groups from being able to view an installation option for a particular operating system image located on your RIS server.

For this part of your interactive installation design process, use job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit) to record your choices for the operating system installation options you want to provide, along with the user groups that can or cannot receive them. In this job aid, you can specify the following under the "Operating System Installation Access" section:

- Answer file name.
- Operating system image associated with the answer file.
- Users or groups of users granted or denied access to operating system images based on the answer file ACL.
- Users or user groups denied access to specific images based on the image folder ACL.

If you need to record this information for multiple operating system images, make an additional copy of the job aid for each operating system image. You might also make more copies of this job aid for recording information that applies to automated installations.

**Defining Installation Options with Group Policy**

You can use Group Policy to configure certain installation options presented by the CIW. With Group Policy, you can create different Group Policy objects to apply to specific users and configure each policy with settings such as **Automatic Setup**, **Custom Setup**, **Restart Setup**, and **Tools**. Each of these options has three different settings associated with it: Enabled, Disabled, and Not Configured.

For example, you could create a separate Group Policy object that provides the **Automatic Setup** option. This causes the CIW to present the appropriate options for clients you have prestaged in Active Directory or for non-prestaged clients that receive predefined computer account names and locations from your RIS server. For non-prestaged clients, you can also create a separate Group Policy object that configures the **Custom Setup** option. This option allows the client to override the automatic computer naming format and causes the CIW to present options to these clients that allow them to specify computer names and the location for their computer accounts in Active Directory.

You can also enable the **Restart Setup** option to allow certain user groups to restart a failed operating system installation attempt if the failure occurs during the text-mode portion of the setup process. This might be useful for security purposes. Also, you can enable the Tools option from Group Policy to allow certain user groups to have access to maintenance and troubleshooting tools for client computers, including the Recovery Console.

**Note**

- If you specify **Not Configured** for a particular RIS setting in a Group Policy object that you define, the default Domain Group Policy setting for the associated option applies.

You can also use the RIS-related Group Policy capabilities to improve security in your network. For more information about the security implications of Group Policy as well as ACL settings, see "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit). For more information about defining CIW setup options in Group Policy, including Group Policy options, see "CIW Design Tasks" later in this chapter.

For this part of your interactive installation design process, use the "Group Policy Configuration" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit) to record your choices for RIS-related Group Policy settings. In this job aid, you can record Group Policy names, the configuration options you want to set, and the user groups to which they apply.

# Fully-Automated Installation Design Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can design your RIS installation to occur in the fully-automated mode, so that no user input is necessary other than logon credentials. To automate installations on remote boot enabled clients, you can use the alternate RIS boot file Startrom.n12 to automatically start downloading the CIW after the RIS client successfully accesses a RIS server. When this occurs, the user initiating the RIS installation does not receive a prompt to press the F12 key for a PXE boot. Instead, the fully automated installation of Risetup or Riprep images proceeds silently after the user logs on.

To enable a fully automated RIS-based installation, it is necessary to substitute the Startrom.n12 boot file for the default Startrom.com boot file. There are two different ways of doing this, depending on the configuration you want:

**Rename the files**
To configure all clients serviced by a RIS server with an automated installation, rename the startup boot files as follows:

- Change Startrom.com to Startrom.bak
- Change Startrom.n12 to Startrom.com

These files are located in the following directory location on your RIS server:

RemoteInstall\OSChooser\i386

You can change the names of the startup boot files manually or you can do this remotely from any computer in the RIS domain by using the change boot file script. To find this script, see the Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources.). You can also run this script locally on a RIS server. By specifying the "automated" or "interactive" command arguments with the script, you can toggle the names of the boot files back and forth to support the default interactive mode or the automated mode, as appropriate. Once you configure the names of the boot files, all clients contacting the RIS server receive the same installation mode.

**Assign Startrom.n12 to specific clients**
If you want to configure automated installations on certain clients only, you can specifically assign them the Startrom.n12 file as the startup file. However, you only have this level of control by using the prestaging script. This script prestages RIS client computer accounts in Active Directory using an Excel spreadsheet generated by the BIOS information script to provide input data. In the spreadsheet, you can add information to specific data cells to configure which startup file each prestaged RIS client should receive. To find these scripts, see the Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources.).

**Important**

- If you change the name of the startup boot files on a RIS server and you also prestage client computers with specific startup file configurations, client computers might not receive the correct startup file. To avoid this problem, either configure startup files with the prestaging script or change the name of the startup file so all clients contacting the RIS server receive the same boot file. Do not use both methods simultaneously.

For more information about designing Active Directory support, including prestaging by using scripts, see "Designing the Active Directory Infrastructure" later in this chapter.

Also, when using Startrom.n12 as the startup file, a repetitive installation loop can occur in some situations. This can happen with PXE-enabled client computers that have their BIOS boot configuration set with the network adapter as the first boot device. It can also happen with non-PXE-enabled clients that have their BIOS boot configuration set with the floppy disk drive as the first boot device.

When you use the Startrom.n12 file to automate image installations, it causes client computers that are powering up to immediately initiate a PXE boot, whether from the RIS boot floppy disk or the network adapter — providing that the BIOS boot sequence is set to allow this. When using Startrom.n12, if the boot sequence is set so that computers boot from a PXE-enabled device or the RIS boot floppy disk, every reboot initiates a RIS installation. By contrast, the RIS default startup file Startrom.com provides the client with the option of performing a remote network boot to RIS, which allows the client to avoid the installation loop.

There are two ways to circumvent this problem with automated installations:

- Use the preferred method of altering the BIOS boot configuration of RIS clients so they always boot first from the hard disk and then from the device that boots to RIS.

- Rename the startup files to the default configuration before clients complete the text mode portion of the setup process (after image download), at which time a reboot occurs. For more information about defining new CIW screens and OSC variables see "CIW Design Tasks" later in this chapter.

If you change the boot sequence in the BIOS of RIS client computers so that the hard disk is set as the first boot device, you must also disable the master boot record (MBR) on the boot partition of the hard disk to intentionally cause the first boot attempt to fail. This way, the next boot device in the sequence is the device that boots to RIS over the network, using either the RIS boot floppy disk or a PXE-enabled network adapter. When the text mode portion of the setup process is complete, the disabled disk partition is rebuilt and automatically reactivated to prevent further PXE booting.

To support automated installations on RIS clients, you need to disable the boot partition on each RIS client you want to receive an automated installation. To do this, you can use a tool such as Diskpart.exe, which you can find in the System32 directory of a computer running Windows Server 2003 or Windows XP Professional Service Pack 1. The syntax to run Diskpart.exe is as follows:

```
Diskpart.exe /S:disablepart.txt
```

The file disablepart.txt, which provides input to Diskpart.exe, contains the following commands:

```
SEL DIS 0
SEL PAR 1
INACTIVE
EXIT
```

Diskpartshould be run on every active partition on the boot drive of all client computers receiving an automated installation. To simplify administration, you might be able to run this program from a client management system such as SMS.

**Note**

- Clients that are configured for an automated installation by using the prestaging script can still have the installation loop problem, if their BIOS boot configuration is not set with the hard disk as first boot device, as described earlier.

# Fully-Automated Installation Design Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When designing a fully automated installation, your primary tasks involve choosing the following:

- The operating system image to configure in the CIW for automated clients.
- The client computers that receive the fully automated installation configuration.
- The configurations for RIS Group Policy options.
- The boot configuration of client computers.

For a job aid to record your design decisions for a fully-automated installation, see "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit).

**Tip**

- Job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) allows you to record design parameters for either automated or interactive installations. When designing an automated installation, simply skip the sections that apply to interactive installations.

**Defining the Operating System Image**

If you want to fully automate the operating system image installation process, it makes sense to provide only one operating system installation option to clients that receive the automated installation. When there is only one operating system image that users are configured to receive, that image is selected automatically and the CIW does not display the operating system choices screen.

For an automated installation, you make a single image available to select users or user groups the same way you make multiple images available for interactive installations: by configuring permissions on answer files and the operating system image folder on your RIS server.

For example, you might have a user group with users that all require an automated installation of a common operating system. For this group, you create an answer file and associate it with a single operating system image. You then set Read permissions on the answer file and image folder to allow each user in the group to receive installation of the operating system.

For this part of your automated installation design process, use the "Operating System Installation Access" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit) to record the operating system image and the user groups that you want to receive it. Also, record the name of the applicable answer file that gives the user group permission to access the operating system image.

You might want to defer recording this information until you design the CIW process. For more information about defining and configuring CIW operating system installation options, see "Interactive Installation Design Tasks" earlier in this chapter.

**Choosing Clients for Automated Installation**

To provide an automated installation to all clients that contact a RIS server, you can rename the Startrom.n12 file in the \Remoteinstall\OSChooser\i386\ folder on your RIS server to Startrom.com. However, this method does not apply to RIS referral servers since these servers only provide clients with referrals to other RIS servers, which in turn supply the appropriate boot files. For more information about RIS referral servers see "RIS Server Configuration Design Tasks" later in this chapter.

For example, you might have a particular group of clients, configured to be serviced by a single RIS server, for which you want to provide automated installations. In this situation, renaming the Startrom.n12 file to Startrom.com is the best solution.

If you need to configure only certain clients for automated installations, you must explicitly specify which clients receive Startrom.n12 and which clients receive Startrom.com. You can only do this if you prestage clients in Active Directory and specify the startup files for each prestaged computer account as part of the input file to the prestaging script. This way, you can

provide an automated installation to any client serviced by any RIS server that is configured to respond to known clients. To find the prestaging script, see the Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources.).

For this part of your automated installation design process, use the "Automated Installations" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit) to record the following information:

- The RIS server name.
- The default startup file configuration you intend to use on your RIS server.
- Whether you prestage clients and specify startup files.
- The users or groups you want to receive an automated installation.

**Defining the Group Policy Configuration**

When providing automated installations, you want users to be able to start their computers, log on, and initiate the installation process without providing any further input. To assure that this occurs, you can configure the RIS-related Group Policy options so that clients receive an installation based on the **Automatic Setup** option. This means that clients do not have to specify a computer account name and an Active Directory location during the CIW process. This option works well for clients that have prestaged computer accounts in Active Directory and for non-prestaged clients that receive a computer account name and location based on preconfigured RIS server settings.

It does not matter if you prestage by creating computer accounts from the Active Directory extension on your RIS server, or if you prestage by using the prestaging script. In either case, you can configure these clients with the **Automatic Setup** option to support automated installations. However, when prestaging from the Active Directory extension, you are limited to providing automated installations to all clients configured to be serviced by the RIS server, because the extension does not provide any options for specifying the startup file.

The approach that gives you the most flexibility is to prestage clients by using the prestaging script, specify the startup file each client uses, and set the Group Policy option to **Automatic Setup**. You can create as many Group Policy objects as you need to configure the user groups that will receive an automated installation.

**Note**

- When designing a fully-automated installation, you might want to disable the **Restart Setup** and **Tools** options in Group Policy settings

For example, you might design a fully-automated installation that requires implementing a process such as the following:

- Create a Group Policy object and set the **Automatic Setup** option.
- Apply the Group Policy to a particular user group for which you want automated installations to occur.
- Prestage computer accounts for the users in that group and specify the use of Startrom.n12 file for those users.
- Create an answer file for the group and configure permissions to view a single operating system image.

For this part of your automated installation design process, use the "Group Policy Configuration" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit) to record the Group Policy name, Group Policy configuration, and the user group(s) to which it applies.

You might want to defer recording this information until you design the CIW process. For more information about configuring Group Policy setup options in the CIW, see "CIW Design Tasks" later in this chapter.

**Defining the Boot Configuration**

To ensure that your automated installations function as expected, you need to define the proper boot configuration by making

the appropriate choices:

**Startup file** Choose which method you intend to use for the startup file configuration:

- Change the names of the startup files on a RIS server to enable all clients serviced by the RIS server to perform a remote network boot of an automated installation.

- Prestage automated installation clients in Active Directory and configure startup files for each client.

You might have already recorded these parameters in job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit) when reviewing the "Choosing Clients for Automated Installation" section, earlier in this chapter. If not, record them now.

**Boot sequence** Choose the client computers on which to configure the appropriate BIOS boot sequence. For automated installations, you must set the boot sequence to use the hard disk as the first boot device and the network adapter (or the floppy disk drive in the case of non-PXE enabled clients) as the second device.

**Hard drive** Choose the client computers for which you must disable all active partitions on the boot drive, using the Microsoft Diskpart.exe tool.

For this part of your automated installation design process, use the "Automated Installations" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) to record the BIOS boot sequence you define for specific clients.

# Designing the CIW Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After a RIS client computer successfully executes a remote boot to a RIS server, the CIW displays on the client and directs the user to the correct choices for installing an operating system. You can customize the CIW process by configuring which setup and operating system installation options the CIW presents to the user. Also, you can customize CIW screens to support your organization's internal requirements by including items such as Help Desk contact information and other pointers. The CIW supports installations of both Risetup and Riprep operating system images on the client computer in either the interactive or fully-automated mode.

A primary consideration when designing the CIW process is the sophistication level of the users who perform the remote installations, because this influences your choice between deploying interactive and automated installations. In turn, the design of the CIW process is affected by the installation type you choose. If you have not done so already, review "Designing the RIS Deployment Mode" earlier in this chapter and document the appropriate design decisions using the suggested job aids.

# CIW Design Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The CIW is a text-based tool that guides the user through the remote operating system installation process. The CIW is the first user interface that displays on the client computer after the user installing the operating system begins the remote boot process.

After a RIS client successfully connects to a RIS server for remote installation of an operating system, a startup boot file downloads from the RIS server to the client. In most cases, this is the default startup file Startrom.com. When the default startup file downloads, it prompts the user performing the installation to press the F12 key, at which time the CIW downloads to the client via TFTP. However, the RIS client can also download startrom.n12 or startrom.n12 renamed to startrom.com for an automated installation, if you appropriately configure the RIS client and RIS server. For more information about the remote boot and installation setup processes, see "Process for Deploying RIS" earlier in this chapter.

The CIW displays whether you configure the user to receive an interactive installation (initiated by pressing the F12 key) or an automated installation. However, in fully automated installations, you typically do not provide any setup or other installation options to the user, which minimizes the number of default CIW screens that are presented.

**Important**

- You must not remove the Welcome, Logon, or Summary screens from the CIW configuration, and you must have at least one screen that has the <meta server action=dnreset> tag.

**CIW Default Configuration**

The default configuration of the CIW process provides basic guidance for installing an operating system by using RIS. You can use the default configuration or you can modify it to accommodate your unique requirements. The following is a brief summary of processes that occur in the default CIW configuration.

When the CIW first downloads to the client computer, the Welcome screen displays. After the user responds to the welcome, the CIW displays the Logon screen which prompts the user to log on to the network with credentials that include the user account, password, and logon domain. After Active Directory validates the user's logon credentials, RIS checks certain Group Policy options to verify the installation configuration for the user. The CIW then presents the specific setup options the user is configured to receive.

After the user specifies the appropriate information for setup options, the CIW displays operating system choices, which can include selections for both Risetup and Riprep images. Following selection of an operating system image and display of the Caution and CIW Summary screens, the installation process begins.

**CIW Default Screens and User Interaction**

When you install RIS and run Risetup.exe, a default set of CIW screens is installed on your RIS server. These screens guide RIS clients through the part of setup that requires user interaction. Use job aid "The CIW Process" (ACIRIS_11.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "The CIW Process" on the Web at http://www.microsoft.com/reskit). You can use this information as a reference when considering the details of your CIW design process.

**CIW Screen Functions**

The CIW screens that install with your RIS server consist of text files with an .osc extension. These screens consist of a default set of files and several example files that show how to enhance CIW functionality. The examples include files that display a multilanguage Welcome screen and a screen that prompts the user for the administrator password.

The default CIW screens provide the basic functionality you need to perform remote operating system installations on client computers in your organization. You can choose to use the default screens or you can build your own customized screens. This section describes the functionality provided by the CIW screens and various modifications you can make to them.

**Welcome screen**

The Welcome screen is a file named Welcome.osc. This screen is the first to display to the user installing the operating system by using RIS. You can modify this screen to include the following:

- Custom information, such as a company-specific message or other preinstallation guidance, which you can add using OSChooser Markup Language (OSCML) tags.

- Custom language options.

For an example of a multilanguage Welcome screen, see the file MultiIng.osc in the following directory path on your RIS server:

RemoteInstall\OSChooser*language*\

For more information about defining a multilanguage CIW process, see "CIW Design Tasks" later in this chapter.

**Logon screen**
The Logon screen is a file named Login.osc. This screen requires the user to log on to the network with valid user credentials consisting of user name, password, and domain name. After the user successfully logs on to the network, RIS uses the user credentials and the Group Policy options applied to the user to determine which setup options the user is configured to receive. If the logon attempt is unsuccessful, the CIW prompts the user to log on again.

You can modify this file to cause a new screen to display that prompts the user to enter the administrator password. For an example of how you can do this, see the file Logined.osc in the following directory path on your RIS server:

RemoteInstall\OSChooser*language*\

### Caution

- If you customize Login.osc, be careful not to modify any of the values within the following OSML tags:
- <INPUT NAME="NTLMV2Enabled" VALUE=%NTLMV2Enabled% MAXLENGTH=255 type=VARIABLE>
- <INPUT NAME="ServerUTCFileTime" VALUE=%ServerUTCFileTime% MAXLENGTH=255 type=VARIABLE>
- Changing any of these values causes you to lose NTLM v2 support.

**Setup options screen**
The Setup Options screen is a file named Choice.osc. This screen displays setup options to the user based on the Group Policy configuration applied to the user. You configure Group Policy setup options by accessing the Group Policy Object Editor from the Active Directory extension on your RIS server. Setup options consist of Automatic Setup, Custom Setup, Restart Setup, and Tools. For more information about Group Policy configurations and defining CIW setup options in Group Policy, see "CIW Design Tasks" later in this chapter.

**Error screen**
The Error screen is a file named Dupauto.osc. This screen displays to the user if RIS finds a duplicate UUID for the client computer in Active Directory. The screen instructs the user to contact the network administrator.

**Operating system choice screen**
The Operating System Choice screen is a file named Oschoice.osc. This screen displays a list of operating system images on the RIS server that are available to the logged-on user. If there is only one possible operating system image that the user is configured to install, then that image is selected and the user does not see this screen.

You can add operating system image choices to this screen by setting permissions on the answer files associated with the images you want to add. This causes additional operating system choices to automatically display to the users who have permission to install them, providing that the users are also configured to be serviced by the RIS server that hosts the images.

### Note

- To make operating system images available to RIS clients, you must also configure Read permissions for the image folder on your RIS server.

You can choose to automate the operating system installation choice or you can provide a selection of images to the user.

**Caution screen**
The Caution screen is a file named Warning.osc. This screen displays a warning message to users indicating that an operating system will be installed on their computers. The message also indicates that this action causes the hard disk to be repartitioned and formatted and that all existing data will be lost.

**Summary screen**
The Summary screen is a file named Install.osc. This screen displays information gathered by the CIW, including the following:

- Computer name
- Computer UUID
- RIS server hosting the installation

By this point in the process, the RIS server has created a computer account object in Active Directory for the client computer. RIS can now readily identify this computer and its associated installation settings, should an operating system need to be reinstalled. The installation now begins if the user presses any key on the keyboard.

**Other error screens**

There are additional screens in the \OSChooser subdirectory on your RIS server that can display when an error occurs. For example, an error might occur if the user enters an incorrect user name or password. When RIS encounters an error, it retrieves the error code (such as 20008 for a login error) and changes it to its hexadecimal equivalent value (00004e28) and appends .osc to the result to derive an error screen file name such as 00004e28.osc. RIS then checks the language directory in use for a screen file with that name. If RIS finds an error screen with a matching name, that screen displays. If RIS cannot find a matching error screen name, an internally generated error message displays. You cannot customize this internally generated error message, however, you can customize the CIW error screens using OSCML tags, just as you can for other CIW screens.

**Custom screens**

You can customize existing screens or create new screens for the CIW using OSCML tags. For new and existing screens, you can use OSC variables to capture user input. For more information about OSC variables, see the discussion about defining new CIW screens and OSC variables later in this section, and see job aid "Reserved OSC Variables" (ACIRIS_12.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Reserved OSC Variables" on the Web at http://www.microsoft.com/reskit). For more information about OSCML tags, see job aid "OSCML and Client Installation Wizard Variables" (ACIRIS_13.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "OSCML and Client Installation Wizard Variables" on the Web at http://www.microsoft.com/reskit).

# CIW Design Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When designing the CIW process, your primary tasks are to choose the following

- The operating system installation options you want to present to clients.
- The setup options you want to provide to clients.
- The CIW configuration you want to provide.

For a job aid to record your design decisions for your CIW process, see "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit).

**Defining CIW Operating System Installation Options**

The operating system installation options that you provide in the CIW enable users to receive the correct operating systems on their computers. By setting explicit user or group security permissions on the answer file associated with specific operating system images on your RIS server, you can control which installation images the CIW displays to users or user groups. You can restrict users and user groups to viewing a single image in the CIW or you can make all operating system images on your RIS server available to them.

Each Risetup or Riprep image that you add to your RIS server has an associated \Templates directory that contains a default answer file (Ristndrd.sif or Riprep.sif) and any additional answer files that you create and associate with that image. These are the answer files on which you configure permissions to cause the CIW to display operating system image installation options on the OSChoice screen. However, you must set these permissions from your RIS server **Properties** dialog box, rather than setting them directly on the answer files in the \Templates directory.

Because selecting individual users for specific image access is time consuming, consider using security groups when applying permissions on answer files. This way, any new users that you add to a particular group automatically receive the correct operating system image installation options.

**Caution**

- The default security permissions allow the **Everyone** group access to operating system images displayed in the CIW. To restrict access to operating system images, remove the **Everyone** group, and add only the users that you want to access the images.

For this part of your CIW design process, it is unnecessary to record your operating system installation access configuration options if you specified them earlier in "Defining the Operating System Image" using job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit). However, if you did not do so, decide on them now and record them under the "Operating System Installation Access" section in the job aid.

**Defining CIW Setup Options in Group Policy**

The setup options that you configure in Group Policy help guide users to the correct choices for installing their operating systems. When determining the setup options you want to provide to users, you need to consider the sophistication level of the users who will be installing using RIS. You can provide advanced users with options that allow them to:

- Specify a name for their computer accounts and a location in Active Directory, if you enable the **Custom Setup** option.
- Restart a failed installation attempt, if you enable the **Restart Setup** option.
- Access maintenance and troubleshooting tools, if you enable the **Tools** option.

For less knowledgeable users, you might want to prestage client computers in Active Directory, automate installations for those computers, and enable the **Automatic Setup** option in the Group Policy object that users at the client computer receive.

If the Group Policy object applied to the users is configured with the **Automatic Setup** option, then the CIW does not present computer account setup options to the users. In this case, the CIW only displays the following installation options:

- Those that offer a choice of operating systems the user can install.
- Those that exist on any other custom screens you create.

**Tip**

- You can configure the CIW to not present any setup or operating system installation options if you configure a fully-automated installation. You can also configure both **Automatic** and **Custom Setup** options to simultaneously provide for less knowledgeable users and also allow advanced users to enter their own input.

You can apply a setup option configuration for clients by using the default domain Group Policy settings, or you can create new Group Policy objects that you customize and apply to specific user security groups. If you apply a new Group Policy object that has any of the setup options set to "Not Configured", the settings for the corresponding options in the default domain policy apply.

For this part of your CIW design process, it is unnecessary to record your Group Policy configuration options if you specified them earlier in "Designing for the RIS Deployment Mode" using job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit). However, if you did not do so, record your Group Policy setup option design choices after evaluating the information in the following paragraphs.

**Automatic setup**
This option provides the simplest way to install an operating system on the client computer. When the Group Policy object applied to the client is configured with this option, RIS searches Active Directory for a computer account object with a UUID that matches the UUID of the client computer. If RIS finds a match, the client computer is named according to the prestaged computer account name. If RIS does not find a match, the client computer account is named and located according to the automatic naming format and Active Directory location that you preconfigured on the RIS server.

**Custom setup**
This option allows users to override the automatic computer naming process and the default location specification in Active Directory for client computer account objects. However, if the user leaves either the computer name or location field blank, RIS uses the automatic naming and location process.

The **Custom Setup** option is similar to the **Automatic Setup** option, although you can use **Custom Setup** to set up the client computer for other users. Because the default configuration of a RIS-based installation is to automatically create computer account names based on the user who logs on to the client computer, it is impractical to use the **Automatic Setup** option when you need to set up another user's computer. For example, Help Desk personnel can use the **Custom Setup** option to preinstall an operating system on a client computer, prior to delivery to the client.

If the name and location the user enters under the **Custom Setup** option and the client UUID match the name, location, and UUID of an existing computer account object, the existing computer account object is reused. If only one of the fields between name, location, and UUID match, a duplicate name or duplicate UUID error screen displays. However, because the user can bypass the error screen, do not use the **Custom Setup** option for prestaged client computers in installations that users perform directly.

**Restart setup**
This option allows users to restart a failed installation attempt, which might occur if the installation process fails or network connectivity is disrupted during the text-mode stage of the setup process. This part of setup is where the CIW process runs prior to the image copy stage. If you provide this option, a **Restart Setup** command is available to users the next time they restart their computers. If the user executes the command, the operating system installation process restarts using the information provided in the previous installation attempt. Display of this command is controlled by the client-specific temporary answer file that RIS generates and then deletes before running Mini-Setup.

**Tools**
This option allows users to access maintenance and troubleshooting tools for use prior to completing the operating system installation. Such tools include a system flash BIOS update tool, diagnostic tools, or other tools provided by third parties.

**Defining the CIW Configuration**
You can choose to use the default CIW configuration or create a custom configuration. If the default CIW configuration is adequate for your purposes, you can use it as a quick way to provide basic installation guidance to your clients. If necessary,

review the functionalities described in "CIW Screen Functions" earlier in this chapter before making your decision.

For a customized configuration, you can build new screens by renaming existing screens and modifying them with OSCML tags. You can also modify existing screens by adding certain OSCML tags that allow you to provide additional information. In addition, you can construct screen prompts for custom user input and use OSC variables in your answer files to capture the input information. You can also use one of the sample .osc files as a new CIW screen to display to users.

If you anticipate network problems, you can even create waiting screens to allow you to debug network delays that might exist between clients and the RIS server, or between the RIS server and the domain controller. For more information, see the discussion about defining new CIW screens and OSC variables later in this section.

**Note**

- To modify CIW screens (.osc files), you can use a text editor such as the Notepad application.

For this part of your CIW design process, decide if you want to use the default CIW configuration. If you think you will need custom screens or modifications to existing screens, choose a custom CIW configuration. To record your decision, use the "CIW Configuration" section in job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit).

If you choose the default configuration, you can skip the remainder of this section. If you choose to define a custom CIW configuration, you need to decide if you want to do one or more of the following:

- Add new screens to the CIW process, including wait screens.
- Modify existing screens with custom prompts or information displays.
- Remove any screens from the CIW process.

**Defining new CIW screens and OSC variables**

You can customize the CIW process by creating additional screens that require the user to provide custom information to complete the installation. This can assist you when creating interactive deployments because it gives you some flexibility in the way you gather user input.

For example, you can build a new screen from an existing screen by modifying its contents using OSCML tags. You can then add user input prompts to the new CIW screen. You also use an OSCML tag to define which screen the CIW displays next, so you can insert a new screen into the CIW process at the correct point. This tag has the following format:

```
<FORM ACTION =”SCREENAME”>
```

When you add input prompts, you must also have a method to capture the user input information. To do this, you create OSC variables in the image answer file. When the CIW runs, RIS replaces the OSC variables in the answer file with the values the user enters in the CIW screen input prompt. RIS then uses these values in the temporary answer file it creates for the client installation configuration.

You can create an OSC variable for any answer file value that works with an unattended setup process in addition to those specifically designed for RIS. However, you must specify OSC variables in your answer file by enclosing them in percent signs so RIS can identify and parse them correctly. For example, you could create prompts on a custom CIW screen that ask users to specify the X and Y resolution and the refresh rate for their video display. To do this, you must include the following OSCML tags in the new screen and specify input name strings such as the following:

```
<INPUT NAME="X-res" >
<INPUT NAME="Y-res" >
<INPUT NAME="Refresh" >
```

Then, in your answer file, you need to specify the OSC variables that capture the user input. You can do this by setting the appropriate unattend values equal to the corresponding OSC variables in the unattended [Display] section, as follows:

```
[Display]
XResolution=%X-res%
YResolution=%Y-res%
VRefresh=%Refresh%
```

You can use up to 64 unique OSC variables per client session; however, there are certain variables that are reserved for RIS use

only. For a list of these variables, see job aid "Reserved OSC Variables" (ACIRIS_12.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Reserved OSC Variables" on the Web at http://www.microsoft.com/reskit). Also, you cannot define OSC variables prior to the time the user logs on to the network, with exception of the %language% variable, which the user can only set before logging on. This allows users to select the proper language in which to proceed in the installation. For more information about OSCML tags, see job aid "OSCML and Client Installation Wizard Variables" (ACIRIS_13.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "OSCML and Client Installation Wizard Variables" on the Web at http://www.microsoft.com/reskit).

The logon process from a PXE client to a Windows Server 2003 RIS server uses NTLM v2 by default to encrypt the user name and password sent between the client and RIS server. After authentication, the client and server communicate using SMB. For more information about NTLM v2, see "Evaluating the NTLM Authentication Level" earlier in this chapter.

From this point forward in the CIW process, you can create any OSC variables you need for custom or existing screens, as long as you associate them with sections in the Unattend.txt answer file. For more information about sections in the Unattend.txt answer file, see "Unattended.txt" in the *Microsoft® Windows Server 2003 Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

Also, if you observe significant time delays between CIW screens, you might consider inserting waiting screens at the appropriate point of the CIW process. One way you might observe where delays are occurring is to configure waiting screens in a test CIW configuration and perform a test installation in your production network environment. If there are no delays, the waiting screens display only for a moment and then the next CIW screen appears. Otherwise, the waiting screens display for the length of the network delay. Once you know where the delays are occurring, you can include waiting screens in your user CIW process.

For example, users might experience a delay between the time they log on and the time they receive the operating system choice screen Choice.osc. To improve the user experience, you can place a waiting screen between Login.osc and Choice.osc with a statement that asks the user to please wait while their logon credentials are verified. Also, slow network access to the domain controller might occur when creating computer account objects in Active Directory. To accommodate the delay, you might add a waiting screen between Choice.osc and Autodup.osc asking the user to wait while creating their computer account.

Lastly, when you insert a new or modified screen into the CIW process on a RIS server, you enable all users who are configured to use the RIS server to view the new screen. Also, you must include changes to the CIW process on each RIS server because there is no capability to synchronize changes to .osc files across multiple RIS servers.

**Note**

- You might want to test your CIW process in your RIS test environment before making changes to the CIW process on RIS servers in your production environment

For this part of your CIW design process, use the "CIW Configuration" section in job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit) to record the following information:

- New screens to be added, including wait screens.
- Insertion points in the CIW process for new screens.
- New user input prompts.
- OSC variables to create for new user prompts and the associated answer file entries.

Also, use the "Operating System Installation Access" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) to indicate which answer files require configuration with OSC variables.

**Defining modifications to existing CIW screens**

You can modify existing screens with the same techniques that you use for customizing new CIW screens. This means you use a combination of OSCML tags in your .osc files and OSC variables in your image answer file (in cases where you are gathering custom user input). However, remember that you cannot create OSC variables in screens that display before the Logon screen, with exception of the %language% variable, which you can use in the Welcome screen configuration.

When you need to provide custom information in a CIW screen, you can use the following OSCML tag to add the text:

```
<BODY>
Text you want to provide.
</BODY>
```

For this part of your CIW design process, use the "CIW Configuration" section in job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) to record the following information:

- Names of the existing screens you want to modify.
- Input prompts or information you want to provide.
- OSC variables and answer file entries you want to use.

Also indicate the answer files that require configuring with OSC variables. Use the "Operating System Installation Access" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) to record this information.

**Defining screens to remove**

There might be specific screens you want to remove from the CIW process. This could include the Setup Options, Error, Operating System Choice, and Caution screens. You might want to do this for RIS clients that receive an automated installation, or if you decide you don't want a particular screen to display. However, be aware that you must not remove the Welcome, Logon, or Summary screens.

If you want to remove a screen from the CIW process, you can change the target of the appropriate OSCML tag so that it no longer calls the screen. These OSCML tags exist in each CIW screen and they point to the next .osc file the CIW displays. For example, the following OSCML tag in the Oschoice.osc file calls the Warning.osc screen:

```
<FORM ACTION =”WARNING”>
```

Therefore, if you want to skip the Warning screen, you can change the target screen of this tag to the Summary screen.

For this part of your CIW design process, use the "CIW Configuration" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit) to identify the screens you want to remove from the CIW process.

At this point, you should have gathered enough information to identify the new CIW screen sequence. Record this information in job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) under the "New CIW Sequence" section along with the RIS servers to which you want the new configuration to apply. Also, if you want to test the new CIW configuration in your RIS test environment, select the appropriate check box in the "New CIW Sequence" section.

**Defining a multilanguage CIW process**

If you need to support multiple languages, you can use a single RIS server to provide service to clients that install different language versions of the Windows Server 2003 or Windows XP Professional operating systems. You can modify the Welcome screen to include the language options you want the RIS server to support. You can find the Welcome.osc file in the \OSChooser directory on your RIS server. Also in this directory is the sample Multilng.osc file that shows how to configure multilanguage support.

The CIW screens that display following the Welcome screen are obtained by RIS from the specific language folder in the \OSChooser directory. The list of operating system images presented to the user is restricted to images for the language selected. For each language version you want to make available to users, you must provide a set of CIW screens in a separate language folder in the \OSChooser directory. The \OSChooser\English version is provided by default.

Certain language restrictions apply to the CIW; these include not providing support for the following:

- Non-101 key keyboards.
- Non OEM fonts.
- Multi-Byte Character Set (MBCS)/Unicode character sets.

These restrictions apply to data used within the CIW, including computer, domain, and directory or file names such as answer file names. It also applies to any example or descriptive text you display to users and to the text that users input to the CIW, such as user names, passwords, and domain names. Because of these restrictions, you must ensure these user input strings do

not contain any non-ASCII characters, since they cannot be used within the CIW. Furthermore, even though you can use a different set of CIW screen files for each language you want to make available, in many cases it is not possible to create screens that are fully localized to a specific language using available character sets.

Also, to support Riprep images in different languages, a Risetup (CD-type) image of the language must also exist on your RIS server. This is necessary to supply files such as device drivers that are needed during installation of the Riprep image, but that did not exist on the master installation from which you derived the Riprep image.

For this part of your CIW design process, record your decision to provide a multilanguage CIW process under "CIW Configuration" in job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit). Also indicate your choices for the language options you want to provide using the "Required User Input" section of job aid "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc), if you did not do so in "Defining User Input" earlier in this chapter.

# Designing the RIS Server Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can design your RIS server configuration to accommodate smaller localized networks as well as full scale corporate networks. In smaller networks of 100 or less client computers, you can minimize the number of RIS servers you need to service client requests for operating system installations. However, in larger network environments, you need to carefully consider the following:

- Where to place RIS servers on the network so as to minimize the impact of RIS traffic.
- Where RIS clients are located in proximity to the RIS servers that service them.
- How many clients you intend to service.
- How you distribute different operating system images to various user groups.
- What security methods you apply to ensure secure operating system installations.
- How you configure your Active Directory infrastructure to support RIS.

To accommodate full-scale corporate environments, you will need multiple RIS servers across your network, preferably using a combination of referral servers that accept, process, and forward client requests, and install servers, which provide the client with boot files, CIW screens, and the actual image download.

# RIS Server Configuration Design Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The way you design your RIS server configuration directly impacts its performance. For example, where you place your RIS servers on the network makes a difference because RIS servers generate heavy traffic during periods when clients are installing operating system images. Also, the number of RIS servers you have plays a role in installation performance because there is a limit to how many clients each RIS server can handle before time-outs occur during client service requests. In addition, by making use of multiple distribution points to provide different operating system images to RIS clients, you can mitigate network traffic and provide faster installation times to clients.

Another factor that impacts your RIS server configuration is the way you implement RIS server security. In corporate environments, you need to design a RIS server configuration that provides secure responses to clients requesting service. To do this, you need to set specific RIS server properties, provide security for non-prestaged clients, and secure the operating system images that you distribute to clients. You can also include prestaging RIS client computer accounts in Active Directory as part of your design, to maximize the security of RIS-based operating system installations.

# RIS Server Configuration Design Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When designing your RIS server configuration, your primary tasks are to define the following:

- Network deployment configuration and the supporting Active Directory infrastructure
- RIS server properties and other RIS configuration parameters.
- RIS security configuration.

For a job aid to record your design decisions for your RIS server configuration, see "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit).

**Note**

- Although your Group Policy settings are part of your RIS server configuration, it is unnecessary to design them here because you should have already made those design decisions in "Designing the RIS Deployment Mode" earlier in this chapter, and recorded them in job aid "Designing for the RIS Deployment Mode"(ACIRIS_08.doc).

**Designing the RIS Network Deployment Configuration**

RIS servers are dependent on your network configuration: the way you deploy and manage your RIS servers on the network determines how they perform. Depending on how you place and configure your RIS servers, one operating system image can support multiple Active Directory sites, domains, and organizational units, or you can provide multiple customized images that you distribute to clients from strategically placed RIS servers.

Because each RIS server can only handle a limited number of simultaneous client installations, you might consider load balancing client service requests by using a RIS referral server. Figure 4.6 shows a basic RIS configuration unit that illustrates the relationship between PXE-enabled remote boot clients, a RIS referral server, and RIS install servers on the network that provide service to clients.

**Figure 4.6 RIS Server Network Deployment**




In Figure 4.6, a PXE-enabled remote boot client requests the remote installation of an operating system. The request is passed to the RIS referral server, which is configured with the **Do not respond to unknown clients** option. This allows only prestaged clients to be acknowledged by the RIS referral server. The RIS referral server checks Active Directory to verify whether the client has a prestaged computer account and if it is configured to receive service from a specific RIS install server. If it finds a prestaged computer account and a designated RIS install server, the RIS referral server passes the request to the appropriate RIS install server (RIS Install Server 3) in Figure 4.6. The client then downloads the CIW and begins the installation process.

RIS Install Servers 1, 2, and 3 are install servers that only provide operating system installations and do not respond to initial client requests for service. Conversely, the referral server does not provide image support, but does answer initial client service

requests.

Figure 4.6 shows how RIS referral and install server configurations can work in an enterprise setting. In this configuration, you can apply tight control to which clients can access which RIS servers. This enables you to load-balance client service requests to ensure that each RIS server is not overloaded. You have this capability because you can specify which RIS server services which clients when you prestage client computer accounts in Active Directory. When you do this, be sure not to configure more than 75 clients per RIS server if you expect heavy service request traffic from clients. Alternatively, you can implement a simpler solution by configuring all RIS servers to respond and provide service to all RIS clients, however, this foregoes the additional security gained by using prestaged RIS clients.

To design a RIS server network deployment that includes configuration units such as the one depicted in Figure 4.6, begin by deciding the following:

- The number of RIS servers you require (including both RIS image and RIS referral servers).
- Where you will place RIS servers.
- How you will distribute RIS server images to clients.

**Defining the number of RIS servers**

The number of RIS servers you need is largely dependent upon how many RIS clients you need to support. You might need multiple RIS servers to support the clients in a large organization or only one RIS server if you are deploying Windows XP Professional on a small LAN or network segment.

The number of RIS servers you will need is impacted by the demand for new, upgrade, or custom operating system installations. As a result, you will need to determine your needs prior to deploying a standard desktop configuration of Windows XP Professional or other operating systems to your clients. Once you determine your needs, you can calculate how many RIS servers to deploy. You can base your estimate on the following metric for best case scenarios: one RIS server can send multiple operating system images over the network for up to 75 clients simultaneously.

The speed of your network and the hardware you use on your RIS server to support image distribution can also have a bearing on how many RIS servers you need. If you have slower network connections or RIS server hardware with marginal capabilities, you will need more RIS servers to handle client service requests to avoid network traffic bottlenecks during periods when RIS servers are active. If you follow the hardware recommendations specified in "Evaluating RIS Server Hardware Requirements" earlier in this chapter, you will be able to maintain support for the maximum number of clients per RIS server.

For load balancing and security reasons, consider using prestaged clients with a RIS referral and install server configuration. If you decide in favor of this configuration, then you must also determine the number of RIS referral servers you need to use. A RIS install server should be in close proximity to the clients it services, but a RIS referral server can pass client service requests to RIS install servers that are located across routers and domains. This is possible as long as the routers are enabled to pass DHCP traffic and there is a trust relationship between domains. As a general guideline for calculating how many RIS referral servers you will need, you can use a metric of one RIS referral server for every three RIS install servers.

For this part of your RIS server configuration design process, use the "RIS Network Deployment Configuration" section in job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to record the total number of clients you need to support and the total number of RIS servers you need to provide image services. Also include the total number of RIS referral servers that you will need.

**Defining RIS server placement**

The primary issues concerning RIS server placement involve where you physically locate the server and where you place it in your Active Directory infrastructure. For more information about designing your Active Directory infrastructure, see "Designing the Active Directory Infrastructure" later in this chapter.

As a general guideline, place RIS servers in close physical proximity to the client computers they service rather than making connections across a WAN link. However, it might be necessary for your clients to locate a RIS server across a router or domain. When this is the case, the router must be configured to pass DHCP packet traffic and there must also be a trust relationship between domains. When considering RIS server placement in your network, you might also consult your DHCP scopes to analyze your domain structure.

In large organizations, do not place your RIS server on a DHCP server. This avoids potential failures in DHCP service if the RIS server becomes overloaded with client service requests. For more information about RIS server placement on the network, see "Assessing RIS Server Placement" earlier in this chapter.

Other placement issues are associated with the type of network connection you use to integrate RIS servers into your environment. Slow connections to RIS servers can hinder the speed of your entire network during periods when RIS is active. Inappropriate RIS server hardware that cannot support network demands can do the same thing. As a practical example, if your organization has branch offices, it is best to place a RIS server in each branch rather than attempting to have clients connect to a RIS server across a slow WAN connection.

For this part of your RIS server configuration design process, use the "RIS Network Deployment Configuration" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to record the following:

- The network location or site name.
- The names of RIS install servers that provide service to specific clients.
- Whether you need to enable DHCP on routers for cross-domain client service requests.
- Whether you need to establish cross domain trusts.
- The names of your RIS referral servers and the Active Directory domains/subnets which they support.

**Defining the distribution of RIS server images**

Depending on the size of your network and the number of clients you have, you might need to create a scheme for managing the distribution of multiple operating system images from different RIS servers to ensure quick installations across the network. You can do this by using multiple RIS servers that provide custom operating systems installations to specific clients. To provide specific operating system images to clients from designated RIS servers, you will need to do the following:

- Create the operating system images you want on each RIS server using Risetup.exe or Riprep.exe.
- Create unique answer files and associate them with specific operating system images on each RIS server.
- Set security permissions on the answer files to configure which users or user groups can access the images.

You can also create unique versions of the CIW process with custom .osc files on each RIS server to manage how you identify and distribute images associated with each RIS server. By distributing operating system images from different RIS servers in this manner, you can mitigate network traffic and accelerate the installation process for designated RIS clients.

For this part of your RIS server configuration design process, use the "RIS Network Deployment Configuration" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to record whether you intend to use multiple RIS servers to handle the distribution of a single operating system or multiple operating systems. If you choose multiple operating systems, record the operating system image names, the RIS servers that will host them, and whether you want to use a corresponding custom CIW process on each RIS server.

For more information about creating operating system images, see "Designing the RIS Installation Type" earlier in this chapter. For more information about customizing the CIW process see "CIW Design Tasks" earlier in this chapter.

**Designing RIS Server Properties**

How you configure your RIS server properties has an impact on RIS server performance and function. The properties you can set on a RIS server are located in the RIS server **Properties** dialog box. You can access this dialog box by using the Active Directory extension on your RIS server. To open the dialog box, right-click the RIS server computer account object in Active Directory and select **Properties** to access the **Remote Install** tab. From here you have access to RIS server properties, which includes the following options:

- **Client support.** Consists of options that allow you to determine which clients the RIS server responds to.
- **Computer naming format.** Consists of various options that determine how computer account objects will be named.
- **Computer account location.** Consists of options that determine where the computer account objects will be placed in Active Directory.

To determine the most appropriate settings to use for RIS server properties in your organization, use Table 4.4 as a guide.

**Table 4.4 RIS Server Property Settings**

| Use This Setting | When |
|---|---|
| **Client support options** | You need to configure the way RIS servers respond to clients requesting installation service. For more information about client support options, see the discussion about designing RIS server security in Designing RIS Server Security later in this document. |
| **Respond to client computers requesting service** | You want a RIS server to acknowledge all clients requesting service, including prestaged and non-prestaged clients, to whom the server makes its operating system images available. Use when maximum security is unnecessary or when you are setting up a RIS referral server. |
| **Do not respond to unknown client computers** | You want a RIS server to acknowledge only clients with prestaged computer accounts in Active Directory, to whom the server makes its operating system images available. Use when you want to maximize the security applied to RIS clients so unauthorized clients cannot receive an operating system installation. |
| **Client computer naming format options** | You configure the Automatic Setup option in Group Policy, so you can apply the computer naming format to non-prestaged clients and to Custom Setup clients that do not provide input for computer name and Active Directory location. |
| **User name** | You want to name the client computer requesting RIS service based on the user name of the operating system installer. This is the default setting. |
| **NP plus MAC address** | You want to name the client computer requesting RIS service based on the media access control (MAC) address of the client network adapter. |
| **Custom naming scheme** | You want to name the client computer requesting RIS service based on a custom naming format that you specify. |
| **Other name variations** | You want to name the client computer requesting RIS service based on name variations such as first name, last name, initial, and so on. |
| **Client account location options**: | You want to define the default Active Directory container for all client computer accounts prior to installation. |
| **Default directory service location** | You want to specify that the client computer account object is created in the Computers container by default when the client joins the domain. Use when you want the client computer to become a member of the same domain as the RIS server handling the client installation process. |
| **Same location as that of the user setting up the client computer** | You want to specify that the client computer account object is created within the same Active Directory container as the user account of the user setting up the computer, for example, in the Users container. |
| **The following directory service location** | You want to predetermine where client computer account objects are created in Active Directory. Use when you want to configure an account location for all client computers installed from a RIS server. |

**Tip**

- If a prestaged client exists in a forest separate from the RIS forest and RIS is configured to not respond to unknown clients, this client will not be answered by RIS. You can fix this by configuring RIS to answer unknown clients and specifying the directory service location in the correct forest where computer accounts are created. Do this using the New Clients tab of the Remote Install dialog box on the RIS server.

For this part of your RIS server configuration design process, use the "RIS Server Properties" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to record the configuration settings you choose from Table 4.4.

**Defining Other RIS Server Configuration Parameters**

The **Remote Install** tab also provides you with access to other dialog boxes that allow you to do the following:

- Associate new answer files with existing images.
- Set security permissions on answer files.
- Add new Risetup images to the RIS server.
- Remove tools or view properties of tools provided by third parties.
- Set security permissions on the RIS server computer account object in Active Directory.

From the **Remote Install** tab, you can also browse Active Directory to do such things as display the UUIDs of all your RIS clients along with the RIS servers designated to service them.

**Defining answer file associations**

By clicking the **Advanced Settings** button in RIS server **Properties**, you can define answer file associations on your RIS server. For example, from the **Images** tab, you can associate answer files with existing operating system images. This allows you to provide custom operating system installations based on answer files that you create and tailor for specific user needs. After you associate the answer file with an image, you can set permissions on the answer file to enable specific users to access the image associated with it.

**Note**

- In **Advanced Settings** in RIS server **Properties**, setting permissions on an item under **Descriptions** on the **Images** tab sets permissions on answer files associated with images rather than on the images themselves.

You should already have recorded the design decisions that specify which answer files you associate with RIS installation images and the user groups that you permit or deny access to these files. These tasks are part of designing the RIS deployment mode and the CIW process.

**Choosing additional Risetup images to host on RIS servers**

From the **Image** tab of RIS server **Properties**, you can add new Risetup images to your RIS server based on an operating system CD that you provide. If you click the **Add** button on the **Images** tab, a dialog box displays with an option that starts the Risetup Wizard. The design decisions about which Risetup images you intend to host on your RIS server(s), made in "Riprep Image Design Tasks" earlier in this chapter, were recorded using job aid "Defining Risetup Images" (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit).

**Choosing to remove tools**

From the **Tools** tab, you can remove tools or view the properties of system maintenance and troubleshooting tools provided by third parties. You cannot add tools to your RIS server from the **Tools** dialog box. Only independent software vendors (ISVs) or original equipment manufacturers (OEMs) can provide system maintenance and troubleshooting tools to administrators, technical support staff, and users of client computers. ISVs and OEMs use a custom setup program to add their tools to the \RemoteInstall directory on a RIS server.

Your RIS server configuration design might involve removing certain tools from your RIS server so that they are not available to clients. However, note that you can achieve the same objective by using Group Policy settings for specific user groups rather than by deleting the tool entirely. You cannot retrieve a tool once you delete it, except by the OEM reinstalling it. Record which tools you want to delete in the "Other RIS Server Configuration Parameters" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit).

**Choosing to delegate RIS administrative tasks**

If you decide to delegate administration of your RIS server, you can set permissions on your RIS server computer account in Active Directory from the **Security** tab of RIS server **Properties**. The decision to delegate RIS administrative tasks is addressed in the discussion about assessing delegation of RIS administrative tasks in "Planning Security for RIS Administrative Tasks" earlier in this chapter. If you did record your decision in job aid "Planning RIS Server Security" (ACIRIS_05.doc) earlier, record the information now in the "RIS Administrative Task Security" section of the job aid. See "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit).

**Designing RIS Server Security**

Most RIS server security issues are addressed in "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit). The security design issue details that must now be completed include choosing how to do the following:

- Provide secure responses from your RIS server to clients, with load balancing.
- Provide security for non-prestaged RIS clients.
- Optimize network security for RIS services.
- Provide authorization for your RIS servers.

**Designing secure RIS server responses and load balancing**

To control how a RIS server responds to remote boot-enabled clients that request service, set **Client support** options on the RIS server **Properties** dialog box. Available settings consist of the following:

- **Respond to client computers requesting service**. The RIS server responds to all clients requesting service. This is the least secure setting because the RIS server does not distinguish between authorized and unauthorized clients.
- **Do not respond to unknown client computers**. The RIS server only responds to clients that have a prestaged computer account object in Active Directory. This is the most secure setting for your network because it enables you to limit access to only authorized clients that are prestaged in Active Directory.

If you configure a RIS server with the **Respond to all clients requesting service** option, you designate that server to handle all client requests for RIS services. In this configuration, you have less security with respect to unknown and possibly unauthorized clients accessing the RIS server. However, you can enhance security by configuring the RIS server to only respond to prestaged clients using the **Do not respond to unknown client computers** option.

In addition, if you prestage all computer accounts and use the RIS referral and install server configuration described in "Designing the RIS Network Deployment Configuration," you can provide load balancing for client service requests by:

- Dedicating RIS servers as referral servers that acknowledge all initial prestaged client service requests and then provide referrals to the appropriate RIS install servers.
- Using specific RIS install servers to handle service requests from designated clients.

Figure 4.7 illustrates how a referral server responds to non-prestaged and prestaged RIS clients.

**Figure 4.7 Securing Client Request Responses and Achieving Load Balancing With RIS Servers**




In Figure 4.7, only Server B is configured as a referral server because it is the only one that can respond to initial client requests for RIS services. It is also configured to only respond to prestaged or "known" clients. Because Client 1 and Client 3 are prestaged and configured to obtain service from a specific RIS server, they receive replies from Server B that refer them to either Server A or Server C.

In Figure 4.7, Servers A and C cannot reply to initial client service requests, but only provide operating system installation services to Client 1 and Client 3 through referrals from Server B. Client 2 is not recognized by Server B because it is not prestaged and therefore cannot receive service from any RIS server.

If you configure Server B to not use the **Do not respond to unknown client computers** option, then Server B itself replies to service requests from Client 2 and offers itself as the remote boot server. Server B functions this way because it is configured to respond to all clients requesting service (Respond = Yes in Figure 4.7).

If you have not already done so, use the "RIS Server Properties" and "RIS Network Deployment Configuration" sections of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to record the **Client support** options you choose and whether you want to use the RIS referral and install server configuration.

**Designing security for non-prestaged RIS clients**

To improve the security of non-prestaged RIS clients, you can control which valid users can create computer accounts in Active Directory during installation. You do this by using the Active Directory **Delegation** feature to preassign the right to join computers to the domain. This automatically provides the user with the Create/Delete Computer Objects permission. You can also do this by explicitly adding the Create Computer Objects and Delete Computer Objects permissions to the user within the **Computers** container of the appropriate domain or organizational unit in Active Directory.

By pre-assigning prestaged client computers with the right to join a domain, you enable users to turn on their systems, connect to a RIS server, log on with their domain accounts, and perform an unassisted installation of an operating system image — all without compromising the security of your network.

For this part of your RIS server security design process, use the "RIS Server Security" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to indicate whether you want to secure non-prestaged RIS clients by giving them the right to join a domain using the Active Directory **Delegation** feature.

**Designing an optimal security configuration with prestaged clients**

You can optimize RIS server security by using prestaged clients. After you prestage computer accounts in Active Directory, configure your RIS server to only respond to these prestaged clients. To further enhance security, you can configure your users with read, write, and reset or change password permissions on the prestaged computer account objects.

For this part of your RIS server security design process, use the "RIS Server Security" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to record your decision to enhance the security of prestaged clients by setting user permissions on the prestaged computer accounts. Also indicate the user groups you want to receive these permissions.

**Designing the RIS server authorization method**

To ensure that your RIS clients are serviced by known RIS servers on the network, you must authorize each RIS server. This ensures that the RIS server is recognized in Active Directory.

The easiest way to authorize RIS on a computer running Windows Server 2003 is to use the **Verify Server** feature on the **Remote Install** tab of the RIS server **Properties** dialog box. You can also type the following command at the command line:

```
Risetup /Check
```

If you intend to delegate this task to specific personnel, they must be part of the Enterprise Admins security group or another group that you configure with this permission in order to access and configure a RIS server.

Alternatively, you can authorize a RIS server to Active Directory by using the **Authorize** function in the **Manage Authorized Servers** dialog box in the Windows Server 2003, Windows XP, or Windows 2000 DHCP snap-in.

To use the DHCP snap-in to authorize the RIS server, it is unnecessary to install the DHCP service. You can use this snap-in if the Administrative Tools package is installed on a computer running Windows XP Professional or Windows Server 2003, from which you can authorize the RIS server. You can install this package by running the adminpak.msi installer — located in the System32 directory of a computer running Windows Server 2003 — on the computer running Windows XP Professional.

You should not attempt to install Windows Server 2003 DHCP on a RIS server just to obtain the snap-in. To service RIS clients, any combined Windows Server 2003 DHCP/RIS server must have a fully functional DHCP service with defined and active scopes. This is because the Windows Server 2003 DHCP service on a combined server is aware that RIS is also present. If a client requests DHCP and remote boot services in its DHCP discovery broadcast, DHCP issues a single reply containing the

specific details on DHCP and remote booting for that server. If the Windows Server 2003 DHCP service is not answering clients properly, the server does not generate a remote boot reply to clients requesting service.

For this part of your RIS server security design process use the "RIS Server Security" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to record the following information:

- The names of RIS server authorization personnel, who are either included in the Enterprise Admins group or in a separate RIS authorizers group that has appropriate permissions.
- The RIS server authorization location.
- The RIS server authorization method.
- Whether you need to install the Administrative Tools package on a computer running Windows XP Professional.

If you have multiple RIS servers, you might simplify things by using a common location and authorization method for each one. For example, you can choose to authorize all RIS servers from a remote administration session by using the **Verify** button in RIS server **Properties**.

# Designing the Active Directory Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The successful implementation of RIS in your environment requires you to carefully analyze your Active Directory architectural design. The logical structure of Active Directory is separate and distinct from the physical network structure. You use physical structures to configure and manage network traffic. You use the logical structure of Active Directory to organize your network resources.

The core units of logical structure in Active Directory are forests, trees, domains, and organizational units. Forests consist of multiple trees which in turn consist of multiple domains that share a contiguous namespace. For any given domain, Active Directory provides organizational units, which are containers that you can use to organize users, computers, and resources into logical administrative groups. This can assist you when defining your RIS server location in Active Directory.

The core units of physical structure associated with Active Directory are sites, which consist of one or more Internet Protocol (IP) subnets connected by a high-speed link. Sites map the physical structure of a network; domains map the logical structure of your organization. Active Directory allows multiple domains in a single site and multiple sites in a single domain.

For further information about Active Directory planning and deployment, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

**Choosing the Active Directory Location for RIS Servers**

Where you place your RIS servers in Active Directory depends in part on how many clients you need to provide with RIS services. The Active Directory location you choose might also depend on your existing infrastructure. If you have a domain containing subnets that each have 75 clients or less, you might create an infrastructure with organizational units for each subnet to be serviced by a single RIS install server. Otherwise, for domains with 75 clients or less, you can use a single RIS install server to provide service to domain clients. Also, to simplify administrative organization, you can choose to create a logical grouping of your RIS servers by placing them all in the same organizational unit.

If it is not possible to locate your clients in close physical proximity to your RIS server, you might locate a RIS server at a particular site and allow RIS clients to connect to it remotely to receive remote installation services. If RIS servers and clients are at different Active Directory sites on separate IP subnets in a common domain, you must connect them with a high speed links, such as a fiber optic backbone. This ensures adequate installation times and minimal traffic congestion during periods when RIS is active.

For this part of your Active Directory infrastructure design process, use the "RIS Network Deployment Configuration" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to record the Active Directory location of each RIS server.

**Designing Active Directory Support**

To design the Active Directory configuration that supports RIS, you need to define the following:

- Any new Active Directory security groups you want to provide for RIS administrators.
- The details of prestaging RIS clients in Active Directory.

**Defining new Active Directory security groups**

If you decided to delegate RIS administrative tasks in job aid "Planning RIS Server Security" (ACIRIS_05.doc), you need to create a new group in Active Directory for RIS administrators. For more information about delegation issues see "Planning Security for RIS Administrative Tasks" earlier in this chapter. If you want to designate more than one group with each handling different tasks, you need to create multiple security groups. After you create the groups, you need to set the appropriate permissions to allow performance of assigned tasks.

For this part of your Active Directory design, use the "Designing Active Directory Support" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to finalize your decision to create Active Directory security groups for RIS administrators or add administrative personnel to the Enterprise Admins group. If you decide to create new groups, record the names of the groups and the personnel you want to add to them.

**Defining prestaging details**

You can prestage client computer accounts either manually using the Active Directory snap-in or with the prestaging script from the Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources.). If you have a small number of clients, it might be sufficient to use the snap-in. If you use the snap-in, however, you can only configure the computer name, UUID, and the RIS server you choose to support the client. You cannot specify which startup file is designated for each client, as you might want to do when configuring some clients with automated installations and others with interactive installations. However, you can designate the startup file by using the prestaging script.

While the primary use of the prestaging script is to automate the prestaging process, you can also use it to automate the configuration of startup boot files for client use. Automating this process helps reduce administrative efforts in a large environment. However, for the prestaging script to work properly, you must run it from within the domain where you want to prestage clients, and the computer from where you run it must have ADSI installed.

The prestaging script uses an Excel spreadsheet created by the BIOS information script as input data, as described in "Evaluating the RIS Client Prestaging Process" earlier in this chapter. You run the BIOS information script to automate the process of obtaining the UUIDs of existing client computers on your network for prestaging these computer accounts in Active Directory.

If you have an OEM spreadsheet with the UUIDs of new client computers, you can add this information to the second column of the Excel spreadsheet generated by the BIOS information scirpt. The OEM UUIDs that you add to the spreadsheet must each be a 32-bit hexadecimal number in raw byte order format as follows:

1534A67812B41C34123F12365E432D16

**Notes**

- When you prestage manually using the Active Directory snap-in, you can use either the raw byte or *pretty print* format. Pretty print format includes curly braces and spaces, as follows:

- {12345678-1234-1234-1234-15E4160B15F2}

When you add OEM UUIDs to the spreadsheet, you must also add other information, including the new computer account name, location, domain\user, description, and startup boot file path. See the prestaging script for more information. The startup boot file path is the path to the RIS server location where the boot files are located, for example:

*RISServername*\REMINST*OSChooser*\i386\Startrom.n12

In the spreadsheet, you can specify which startup file you want for each client, by using either the Startrom.n12 or Startrom.com boot files. However, the prestaging script also provides options that allow you to set all clients to either boot file, to accommodate groups of clients that you configure with interactive or automated installations. When you choose to use these options, you must specify the appropriate action command, the RIS server name, the image name, and the path to a fully-configured input spreadsheet file. In this case, the script does not read data from the cells in Startup File Path column of the spreadsheet, but applies the value you enter at the command line to each client computer account listed in the spreadsheet. Values are **automate**, to configure the client with Startrom.n12 and **interactive**, to configure the client with Startrom.com.

The prestaging script contains usage instructions that explain how to run the script and the commands or input arguments you must provide. The script also provides header information that explains the details of the Excel spreadsheet file format. Whether you prestage by script or manually, you still must acquire the UUIDs for your client computers. For more information about methods to acquire the UUIDs for client computers, see "Evaluating the RIS Client Prestaging Process" earlier in this chapter.

For this part of your Active Directory design, use the "Designing Active Directory Support" section of job aid "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit) to record your decision to prestage client computers in Active Directory either manually or using the prestaging script. If you decide to prestage, also record the name of the input Excel file that the script requires and the personnel who will create and configure this file. You can also specify the method you will use to obtain UUIDs.

# Designing a Test RIS Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before implementing a full scale RIS deployment in your organization, you can create a RIS test environment to facilitate preliminary RIS configuration and testing. After confirming that the test environment produces the desired results, you can begin your RIS rollout. For more information about creating a test environment, see "Designing a Test Environment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

The RIS test environment enables you to create and configure your RIS server, generate operating system images, and perform test deployments of RIS images to clients outside your production environment. You can also use the RIS test environment to run tests on uniquely tailored answer files and custom CIW configurations to assure proper functioning prior to introducing them in your network. You might also use your test environment to do a preliminary test run of the BIOS information, prestaging and boot file name scripts to become familiar with them before running them on the network. In addition, you can use the RIS test environment to acquire baseline performance indications for your RIS server. For more information about RIS server performance, see "Planning RIS Server Performance" earlier in this chapter.

To create the RIS test environment, you need a minimum of two computers. You can set up one computer as a domain controller with multiple roles, including Active Directory, DNS, DHCP, and RIS, and a second computer that serves as a master computer from which you create custom file system images using Riprep. After you configure one or more images and place them on the RIS server, you can use the second computer as a client on which you install the image. Also, if you want to create CD-type operating system images in your test environment and you have the operating system CDs, you can run Risetup on the RIS server.

Alternatively, you can use a three-computer configuration such as the following:

- One domain controller with Active Directory, DNS, and DHCP.
- One member server on which you install RIS.
- One combined master and client computer running the operating system you want to image, such as Windows XP.

If you have additional computers, you can break up the combined master and client configuration and provide separate computers for each. This simplifies things because you can avoid repeated installations of operating systems on the master computer after performing test installations on the client with RIS.

**Important**

- In your production environment, strongly consider not having your domain controller share multiple roles. Instead, you can configure separate computers with unique roles such as domain controller with DHCP, member server with Active Directory and DNS, and member server with RIS.

For this part of your design process, use job aid "Designing a RIS Test Environment" (ACIRIS_10.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing a RIS Test Environment" on the Web at http://www.microsoft.com/reskit) to record the configuration of your test environment and the tasks you intend to carry out there, including any of the following:

- Configuring your test RIS server.
- Testing-run scripts for prestaging, getting client UUIDs, and changing startup file names.
- Creating operating system images using Riprep and Risetup.
- Testing answer files and CIW configurations in trial deployments.
- Analyzing RIS server performance.

# Configuring and Deploying RIS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

At this point, you can pass your design job aids to the deployment team to serve as inputs to the RIS configuration and deployment process. Figure 4.8 illustrates the order of the configuration and deployment tasks to accomplish at this stage.

**Figure 4.8 Configuring and Deploying RIS**




For specific procedures that support the RIS configuration and deployment process, see "Remote Installation Services" in Help and Support Center for Windows Server 2003.

# Creating a RIS Test Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To create a RIS test environment for running preliminary tests, you need to accomplish the following tasks:

1. Install the appropriate hardware on your RIS server to support your test environment and then install the Windows Server 2003 operating system.
2. Install RIS on your test RIS server. (This automatically creates a CD-type Risetup image.)
3. Create any additional Risetup images using operating system CDs.
4. Install the appropriate hardware on a computer you designate as a domain controller for your RIS test environment and then install the Windows Server 2003 operating system.
5. Install Active Directory by running dcpromo.exe at the command line or by running the Active Directory Installation Wizard.
6. Configure and activate your DHCP scope to provide IP addresses on the test network.
7. Install the appropriate hardware and the Windows XP Professional operating system on a computer you designate in the role of master computer and RIS client.
8. Connect the domain controller, RIS server, master computer and RIS client to the test network using linear bus topology and Ethernet components.
9. Configure Active Directory on the domain controller with user and computer accounts and join the RIS server and client to the test domain. (You can prestage the client computer account using the UUID of the client computer.)
10. Install and configure applications on the master computer, set operating system parameters, and add any special drivers required for a Riprep image.
11. Create the Riprep image on your RIS server by running Riprep.exe from the master computer.
12. Create any custom answer files (including permissions) and CIW configurations.
13. Configure your test RIS server with client support options, a computer naming format, and the location where computer accounts are generated.
14. Reinstall an operating system on the master computer, which now acts as the client.
15. If you configure the RIS server for automated installations, set the boot sequence in the BIOS of client computers to boot from the hard disk first and the network boot device second.
16. Disable any active boot partitions on the client computer hard disk using a tool such as Diskpart.exe.

To run preliminary tests, perform the following tasks:

1. Test-run the scripts for obtaining client UUIDs, prestaging clients in Active Directory, and changing startup file names.
2. Test custom answer file, CIW, and boot configurations in trial operating system deployments.
3. Analyze RIS server performance.

After you achieve the results you expect in your RIS test environment, proceed to the configuration tasks for your production environment.

# Configuring Networking Support

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to configure your domain controller to support the RIS server that provides operating system installation services. You also need to configure your network infrastructure to support RIS-based operating system installations.

**Configuring Domain Controller Support**

To configure your domain controller, you need to accomplish the following:

1. Install the appropriate hardware on your domain controller to support the RIS environment and then install the Windows Server 2003 operating system.

2. Ensure that you have the software required to support RIS, including:

   - DHCP installed and activated on your domain controller, unless you are using a third-party application on a member server.

   - A configured DNS server to provide name resolution services for the domain.

3. Add a RIS installation account to your domain to enable autologon. Use the Active Directory Users and Computers snap-in to create this account.

**Network Infrastructure Configuration**

To configure your network infrastructure, you need to accomplish the following:

1. Make any necessary changes to your network topology to provide support for the minimum data transmission rate of 10 Mbps (100 Mbps is highly recommended). See "Assess Existing Network Infrastructure" earlier in this chapter.

2. Ensure that you have the necessary components in your network to support RIS, including a domain controller running Windows Server 2003 with DHCP and Active Directory services, and a DNS server. See "Assessing RIS Server Software Requirements" earlier in this chapter.

3. Install your RIS servers at suitable network installation points See "Evaluating Network Installation Points" earlier in this chapter.

4. Set the NTLM authentication level you require on the network. See "Evaluating the NTLM Authentication Level" earlier in this chapter.

5. Set up standard security measures for your network, including auditing and monitoring, securing physical access to the network, and enforcing a strict password policy. See "Assessing the Security of the PXE Environment" earlier in this chapter.

# Configuring Production Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure production clients in preparation for RIS-based operating system installations, you need to accomplish the following:

1. Install new client computers that use the correct hardware to support RIS installations in the appropriate locations on your network. See "Evaluating RIS Client Hardware" earlier in this chapter.

2. Install the correct hardware on existing clients to support RIS installations. See "Evaluating RIS Client Hardware" earlier in this chapter.

3. Update any existing client computers that have incorrect HAL types for the Riprep images they are to receive. For more information about verifying the RIS client remote boot configuration, see "Evaluating Remote Boot Capabilities of RIS Clients" earlier in this chapter.

4. Run the BIOS information script to determine client BIOS compatibility with network adapter booting. For more information about verifying the RIS client remote boot configuration, see "Evaluating Remote Boot Capabilities of RIS Clients" earlier in this chapter.

   **Note**

   - Before running the BIOS information script, you must ensure that the computer running the script has ADSI and WMI installed and that client computers have WMI installed

5. Create RIS boot floppy disks for non-PXE enabled clients. For more information about verifying the RIS client remote boot configuration, see "Evaluating Remote Boot Capabilities of RIS Clients" earlier in this chapter.

6. Migrate user state. For more information about migrating user state, see "Auditing Existing Clients" earlier in this chapter.

7. Obtain client computer UUIDs, using either Systems Management Server, OEM listings, or run the BIOS information script. For more information about obtaining client computer UUIDs, see "Evaluating the RIS Client Prestaging Process" earlier in this chapter.

8. For PXE-enabled clients receiving an automated installation, configure the boot sequence in the BIOS to boot from the hard disk first and the network adapter second.

9. For non-PXE enabled clients receiving an automated installation, configure the boot sequence in the BIOS to boot from the hard disk first and the floppy disk second.

10. Disable all active boot partitions on the hard disk of all client computers receiving an automated installation by using a tool such as Diskpart.exe with the /S:disablepart.txt argument.

11. Erase the hard disks of client computers using a management application such as Systems Management Server.

For more information about defining the boot configuration, also see "Fully-Automated Installation Design Tasks" earlier in this chapter.

# Creating a Production RIS Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To create a production RIS server to provide operating system installation services on the network, you need to accomplish the following:

1. Install the hardware necessary on each server to support RIS installations. For more information about RIS server hardware requirements, see "Evaluating RIS Server Hardware Requirements" earlier in this chapter.

2. Install the Windows Server 2003 operating system on each RIS server. For more information about RIS server software requirements, see "Assessing RIS Server Software Requirements" earlier in this chapter.

3. Install RIS on each RIS server. For more information about RIS components and deploying RIS, see "Process for Deploying RIS" earlier in this chapter.

4. Create any additional Risetup images for the RIS servers using operating system CDs. For more information about choosing additional Risetup images to host on RIS servers, see "RIS Server Configuration Design Tasks" earlier in this chapter.

   **Note**

   - You can join the RIS server to the domain where it will provide service after you configure domain controller support.

# Configuring a Master Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure a master computer from which you generate Riprep images, you need to accomplish the following:

1. Install the appropriate hardware and the Windows XP Professional operating system on a computer you designate as the master computer and join it to the domain where the RIS server is located. For more information about master computer hardware requirements, see "Assessing Master Computer Requirements" earlier in this chapter.

2. Install and configure applications on the master computer, set operating system parameters, and add any special drivers required for a custom Riprep image. For more information about designing Riprep images, see "Design a Riprep-Based Installation" earlier in this chapter.

3. Test your Riprep images to determine the impact on user profiles. For more information about how Riprep image design affects user profiles, see "Riprep Image Design and User Profiles" earlier in this chapter.

4. Create the Riprep image on your RIS server by running Riprep.exe from the following location: *risserver*\reminst\admin\i386\riprep.exe.

5. Configure permissions on the answer files for each Riprep image you create and on the operating system image folders, to allow users to access the images. For more information about setting permissions on answer files, see "Evaluating Security for Operating System Images" earlier in this chapter.

**Note**

- If you set the local administrator password on a Riprep image, passwords entered in the CIW are ignored.

# Installing the Master Computer Operating System

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you install the operating system on the master computer, use the unattended method. By using unattended installs, it is easier to create and catalog unique configurations of operating systems that include different components such as drivers and applications. For each operating system installation, you create a separate answer file that identifies the components of the master installation configuration. The answer file facilitates the unattended installation and also serves as a record of your master operating system installation configuration, which is easily stored and can be used to re-create the original configuration.

Another benefit of using an unattended installation for your master computer configuration is that you can include a command in the GUIRunOnce section of the answer file to automatically start Riprep.exe when the unattended installation completes. By integrating Riprep into the unattended installation, you automatically create the master installation image on your RIS server.

For example, creating a master installation for a file server using an unattended installation with Riprep involves the following steps:

1. Create an answer file and configure it with entries and values for components that are appropriate for your master operating system configuration.

2. Specify the command to run Riprep.exe in the GUIRunOnce section of the answer file by using the following entry:

   [GuiRunOnce]

   "\\risserver\reminst\admin\i386\riprep.exe"

3. At a command prompt on the server, use the following syntax to run the unattended installation:

```
Winnt32.exe /unattend: answerfilepath
```

For more information about creating answer files for unattended installations, see "Answer files" in *Microsoft Windows Server 2003 Corporate DeploymentWindows Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

# Configuring the Master Computer Operating System

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the Master Computer Operating System**
To configure the master computer, you need to accomplish the following:

1. Create the desktop configuration.
2. Add hardware devices.
3. Set passwords.
4. Create an optional command file.
5. Set language and regional options.
6. Add and configure applications

You must also clean up the installation by doing the following:

**Reset history settings** Reset all history settings in the operating system and all applications that are installed on the master installation. This includes the most frequently used applications list and the history list in Internet Explorer.

**Delete files and folders** Delete all files and folders that you do not want end users to see. This might include:

- Files and folders that you used to build the master installation, such as tools, documents, and scripts.
- Temporary Internet files, including cookies.
- Files and folders in My Documents.

**Configure the user profile for the Default User** Configure the user profile for the Default User so that all of the installation and configuration tasks you performed are available to end users. To do this, you first create a local user account, and then add the account to the local Administrators group. Next, log on to the computer using the new user account, and copy the Administrator user profile to the Default User profile. Do not log off.

# Testing Riprep Images and User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When creating Riprep images, you need to understand how user profiles affect the changes you make to a master installation and the impact on users who log on to computers that receive the Riprep image.

Applications that are compliant with the operating system that you are preparing to deploy properly separate user-specific and computer-specific configuration settings and data. Installing such applications in your Riprep master installation makes them available to all users of client computers on which you install the Riprep image. Applications that are not compliant with Windows Server 2003 might rely on per-user configurations that are specific to the profile of the user (typically the Administrator) who actually installs the applications prior to running Riprep, rather than computer-specific configurations. These configurations remain specific to the installing user, which can result in the application or configuration setting not being available or not functioning properly for the users of the client computer where you install the Riprep image. Also, some non-application configuration changes, such as the wallpaper specified for a user's desktop, are applied only to the current user's profile by default, which means they are not available to users of the client computer where you install the Riprep image.

To ensure that the application or configuration settings of your Riprep image work properly with the implementation of user profiles in your organization, you need to test them thoroughly. To test a configuration, you can make a change while logged on as the Administrator, log off, and then log on with a user account that is representative of your organization. If the change you made applies to the user account with which you log on, then it should also apply to all users who log on to systems installed with a Riprep image containing the same change. Complete your testing by creating a Riprep image and installing it on a client computer. Log on to the computer with a different user account and verify that the change applies and is fully functional.

# Running the Riprep Wizard on the Master Computer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you are ready to create your image, you run Riprep.exe from the master computer by specifying the following in the **Run** dialog box:

```
\\RISServerName\Reminst\Admin\i386\Riprep.exe
```

By default, Riprep-based images do not perform full Plug and Play enumeration during operating system installation on the client. If you want full Plug and Play enumeration to occur, you must start Riprep.exe with the **/pnp** option, as follows:

```
\\RISServerName\Reminst\Admin\i386\Riprep.exe /pnp
```

After you run this command, full Plug and Play enumeration occurs. If you want to turn off full Plug and Play enumeration, you must recreate the image.

**Note**

- You can also run Riprep.exe by entering these commands at the command line.

**Replicating Images to Other RIS Servers**

RIS does not provide a mechanism for replicating operating system images from one RIS server to another. However, you can use third-party tools to replicate operating system images. If you use a third-party tool, make sure that the replication mechanism supports alternate file streams, file maintenance attributes, extended attributes, and security settings of the source images.

# Configuring Answer File and Image Folder Permissions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Each time you run Riprep.exe on a master installation, you create a default answer file that is specifically configured for that master image. To make this image available to your users, you must set Read permissions on the answer file associated with it. Using the **Images** tab in RIS server **Properties**, set the ACL on each answer file to specify which users can receive the master installation image. Note that you also need to set Read permissions on the operating system image folder located in the following directory path on the RIS server:

*RISServer*\RemoteInstall\Setup*LanguageFolder*\Images*OSImageFolder*\

Setting permissions on the image folder ensures that specified users will have access to the operating system image.

# Building a Master Distribution Share Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you create your RIS server on a computer running Windows Server 2003, you also automatically create an initial CD-type image of Windows XP Professional or Windows Server 2003. However, you can also create additional CD-type images. To create a CD-type image of an operating system on a RIS server from which you can build custom distribution share installations, you need to accomplish the following:

1. Log on to the RIS server.
2. Insert a CD of the operating system you want to image in the CD-ROM drive of the master computer.
3. Locate your RIS server by using the Active Directory extension.
4. Use the **Add** button on the **Images** tab of your RIS server **Properties** dialog boxto run the **Remote Installation Services Setup Wizard** (Risetup), which locates the operating system source files and creates the directory structure of the image.
5. Install applications and drivers in your distribution share.
6. Configure new answer files that create unique operating system installations by using Cmdlines.txt and the entry [GUIRunOnce] unattended.

For more information about designing a Risetup image, including configuring custom answer files, see "Design a Risetup-Based Installation" earlier in this chapter.

# Configuring the RIS Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By configuring your RIS server, you determine the way it responds to RIS clients requesting service. If your design includes a RIS referral server, your configuration options are slightly different.

**RIS Server Configuration**

To configure your RIS server, you need to accomplish the following tasks:

1. Authorize the RIS server to Active Directory using the **Verify** button on the **Remote Install** tab of your RIS server **Properties** dialog box.

2. Set client response options that determine whether the RIS server responds to all clients requesting service or only to those that have prestaged computer account objects in Active Directory.

3. Set the computer naming policy.

4. Set the location where new managed computer account objects will be created in Active Directory.

5. Set permissions for RIS users to create computer accounts using the **Delegation** feature of Active Directory to preassign them the right to join computers to the domain. Alternatively, do this by explicitly adding the Create Computer Objects and Delete Computer Objects permissions to the client within the **Computers** container of Active Directory.

6. If you are prestaging client computers in Active Directory, add prestaged computer accounts to the domain by using the Active Directory extension on your RIS server or automate the process by running the prestaging script from any domain computer that has ADSI and WMI installed.

7. Using the **Images** tab of your RIS server **Properties** dialog box, associate any additional answer files you have with Riprep and Risetup images that are located on the RIS server.

**Note**

- Any configuration changes you make from the Active Directory extension on your RIS server might not be immediately replicated to other domain controllers in Active Directory. The time required to replicate changes depends on your network topology.

**RIS Referral Server Configuration**

To configure a RIS server as a referral server, you need to accomplish the following tasks:

1. Install the hardware required to support the Windows Server 2003 operating system on the computer you designate as a RIS referral server.

2. Install RIS as described above and at least one CD image.

3. Authorize the RIS referral server to Active Directory using the **Verify** button in your RIS server **Properties** dialog box.

4. Set the **Do not respond to unknown client computers** option in the RIS server **Properties** dialog box.

**Note**

- The RIS referral server only acts in the capacity of making referrals if you prestage clients in Active Directory and configure them to use specific RIS install servers.

# Creating the CIW Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are using the default CIW configuration, you do not need to perform any of the following configuration tasks. Otherwise, to create a custom CIW configuration, you need to accomplish the following:

1. If you have not done so already, configure the operating system installation options you want to present to clients by setting answer file ACLs to permit or deny client access to operating system images in the CIW.

2. If you have not done so already, configure permissions for the operating system image folder on the RIS server to explicitly deny access to specific user groups, if applicable.

3. Configure the setup options you want to provide to RIS clients by configuring RIS-related settings in the default domain Group Policy or in any new Group Policy objects that you create for specific user groups.

4. Modify existing screens and use OSCML tags to create custom input prompts.

5. Add any new screens you want and use OSCML tags to create custom input prompts or information displays.

6. Configure answer files with OSC variables to capture user input.

7. Remove screens from the CIW by modifying the <FORM ACTION="screenname"> tags with different screen names, as required.

8. Add any language options to the CIW by modifying the Welcome screen.

For an OSCML tag reference, see job aid "OSCML and Client Installation Wizard Variables" (ACIRIS_13.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "OSCML and Client Installation Wizard Variables" on the Web at http://www.microsoft.com/reskit). For more information about Group Policy settings that affect the CIW in addition to using OSCML tags and OSC variables, see "CIW Design Tasks" earlier in this chapter.

# Deploying an Operating System

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Once you have completed configuring and deploying RIS in your production environment, you are ready to perform operating system deployments in your production environment. Figure 4.9 illustrates the deployment options at this stage of the process.

**Figure 4.9 Deploying an Operating System**




To deploy an operating system by using RIS requires that client computers initiate a remote network boot. To do this, you must use one of the following methods:

- Use PXE-enabled client computers to boot from their network cards to a remote RIS server.
- Use a RIS boot floppy disk to emulate the PXE process to boot supported non PXE-enabled client computers from a remote RIS server.

# Using a Network Boot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To perform deployment of an operating system image hosted on a RIS server using a network boot from a PXE-enabled client, you need to accomplish the following:

1. Power up the client computer.

2. If the client is receiving an interactive deployment, wait for the client to receive a prompt to press the F12 key for a network boot to a RIS server. The prompt appears after the client receives an IP address from the DHCP service.

3. Press the F12 key to initiate the network boot.

4. If the client is receiving an automated deployment, wait for the CIW to begin downloading. The download begins after the client receives an IP address from the DHCP service. The F12 prompt is bypassed for automated deployments.

5. Follow the CIW screens after TFTP downloads the CIW files to the client.

# Using a RIS Boot Floppy Disk

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To perform deployment of an operating system image hosted on a RIS server using a network boot from a RIS boot floppy disk client, you need to accomplish the following:

1. Insert the RIS boot floppy disk into the client computer floppy disk drive.

2. Accomplish the tasks listed in the previous section "Using a Network Boot."

# Additional Resources for Designing RIS Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this chapter.

**Related Information**

- "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit for information about Active Directory planning and deployment.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for an introduction to TCP/IP, or for more information about Dynamic Host Configuration Protocol (DHCP).
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about the Windows Installer, and about using Group Policy to configure computer and user deployments of applications.
- "Migrating User State" in this book for information about migrating user data and settings.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Operating systems supported by Remote Installation Services" in Help and Support Center for Windows Server 2003 for more information about operating systems supported by RIS.
- "Remote Installation Services system requirements" in Help and Support Center for Windows Server 2003 for more information about Remote Installation Services system requirements.
- "Remote Installation Services administration overview" in Help and Support Center for Windows Server 2003 for procedures to prestage RIS clients using the Active Directory snap-in on a RIS server.
- "Setting the LAN Manager Authentication Level on a network that includes RIS" in Help and Support Center for Windows Server 2003 for more information about choosing the most appropriate LAN Manager authentication level in a network that includes RIS.
- "Set permissions for administrators who manage client installation images for RIS" in Help and Support Center for Windows Server 2003 for more information about permission requirements for RIS tasks.
- "Remote Installation Services" in Help and Support Center for Windows Server 2003 for specific procedures that support the RIS configuration and deployment process.

**Related Job Aids**

- "Planning for RIS Clients" (ACIRIS_01.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Clients" on the Web at http://www.microsoft.com/reskit).
- "Planning for RIS Servers" (ACIRIS_02.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning for RIS Servers" on the Web at http://www.microsoft.com/reskit).
- "Planning the Master Computer Configuration" (ACIRIS_03.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning the Master Computer Configuration" on the Web at http://www.microsoft.com/reskit).

- "Planning the RIS Network Configuration" (ACIRIS_04.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning the RIS Network Configuration" on the Web at http://www.microsoft.com/reskit).

- "Planning RIS Server Security" (ACIRIS_05.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning RIS Server Security" on the Web at http://www.microsoft.com/reskit).

- "Defining Riprep Images" (ACIRIS_06.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Riprep Images" on the Web at http://www.microsoft.com/reskit).

- "Defining Risetup Images (ACIRIS_07.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Defining Risetup Images" on the Web at http://www.microsoft.com/reskit).

- "Designing the RIS Deployment Mode and CIW Process" (ACIRIS_08.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Deployment Mode and CIW Process" on the Web at http://www.microsoft.com/reskit).

- "Designing the RIS Server Configuration" (ACIRIS_09.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing the RIS Server Configuration" on the Web at http://www.microsoft.com/reskit).

- "Designing a RIS Test Environment" (ACIRIS_10.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing a RIS Test Environment" on the Web at http://www.microsoft.com/reskit).

- "The CIW Process" (ACIRIS_11.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "The CIW Process" on the Web at http://www.microsoft.com/reskit).

- "Reserved OSC Variables" (ACIRIS_12.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Reserved OSC Variables" on the Web at http://www.microsoft.com/reskit).

- "OSCML and Client Installation Wizard Variables" (ACIRIS_13.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "OSCML and Client Installation Wizard Variables" on the Web at http://www.microsoft.com/reskit).

- The Remote Installation Scripts link on the Web Resources page http://www.microsoft.com/windows/reskits/webresources for the BIOS information script, the prestaging script, and the boot file name script.

# Migrating User State

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you move client computers to the Microsoft® Windows® XP operating system from earlier versions of Windows, it is important to save and then restore user data and settings. This process is known as *migrating user state*. By carefully planning and implementing user state migration, you help conserve IT staff time, preserve important data, scale the migration as needed, and minimize costs while maintaining user productivity and workplace morale.

**In This Chapter**

Overview of User State Migration

Choosing the User State Collection Method

Identifying Migration Content

Creating a Detailed Migration Plan

Testing Your Migration Process

Additional Resources for Migrating User State

**Related Information**

- For information about using Remote Installation Services (RIS), see "Designing RIS Installations" in this book.
- For information about scripting in a Microsoft® Windows® Server 2003 environment, see the at http://www.microsoft.com/reskit, or see the MSDN Scripting Clinic link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Overview of User State Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Any time that you perform a new installation of Windows XP on a client workstation, you should migrate user state to ease users into the new system and maintain user productivity. *User state* consists of *user data* — the files that users create and need to do their jobs — along with *user settings* containing application-specific and user-specific information. Additionally, application settings supply the user with links, menus, and other information that can be essential for their productivity.

If user state is not migrated, an organization can accrue costs as users spend production time reconfiguring their applications and other settings. Organizations must evaluate the cost/benefits ratios for migrating various types of items. They must understand the security issues related to migration and be sure to educate users about what to expect before and after the migration.

The way that you choose to deploy Windows XP affects your user state migration plan. Ideally, an organization can perform either a parallel or a wipe-and-load deployment, restoring collected user state to a clean environment.

The method that you use to collect and restore user state is critical to the success and efficiency of your user state migration. To avoid the high migration cost of a strictly manual migration, an organization can:

- Partially script the migration, leaving nonstandard items to the discretion of the individual user or IT staff.
- Use migration tools that automate the migration of common settings but allow customization.
- Create its own custom tools.

The degree to which an organization should automate user state migration depends on these and a variety of other factors, including the number of users to be migrated and how widely dispersed they are; how centralized the organization's IT effort is; the degree to which users share a common desktop, folder hierarchy, and computing requirements; the IT expertise available to assist in and support the migration; and whether the deployment involves simultaneous domain migration.

Before you begin planning a user state migration, identify the computers on which you will deploy Windows XP, and determine the appropriate deployment method for each computer. When you complete this process, you will be ready to deploy Windows XP, with a complete, tested plan for migrating user state during the system deployment and a schedule for the migration.

# Process of User State Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User State Migration Process**

Proper planning is essential for a successful user state migration. Before creating a detailed migration plan, identify the best methods for collecting, storing, and restoring user state data and decide which user data and settings to migrate. After making these decisions, prepare a migration plan that addresses storage and security requirements; potential registry, drive, and domain changes; scheduling; and the education of users. Then test your migration process before embarking on a large-scale migration. Figure 5.1 shows the process for planning a user state migration.

**Figure 5.1 Migrating User State**



# Tools Used for the Migration Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some large organizations develop their own migration tool. This frequently provides an excellent migration experience for the user, because the tool is customized to the specific environment. Such a tool can capture and restore user settings and files. Adjustments can be made quickly when new applications are deployed or bugs are found.

Typically, this option requires a significant investment in development time. If your organization does not have personnel who can create a migration tool, the cost of hiring programmers to create one might be prohibitive. Even if you do have programmers on staff, compare the cost of tool development with the migration costs of other available methods.

Many tools currently available from vendors can collect and restore most necessary settings and are extensible to include additional application settings. These tools often provide multiple types of rules to specify which files to migrate. However, while fairly thorough, such tools are not targeted to your specific environment, and their initial cost can be high.

Microsoft provides two tools designed specifically for migrating user state in a Windows environment:

- The Files and Settings Transfer Wizard
- The User State Migration Tool (USMT)

Both tools automate the migration of basic application, operating system, and user settings, as well as user data, and both support customization.

**Files and Settings Transfer Wizard**
The Files and Settings Transfer Wizard is a Windows XP accessory, available in System Tools. (On the Program menu, point to All Programs, Accessories, System Tools, and then click Files and Settings Transfer Wizard.) The wizard enables users to migrate personal display properties, folder and taskbar options, and Internet browser and mail settings, as well as specific files or entire folders (such as My Documents, My Pictures, and Favorites) from their old computer to their new one without any manual configuration.

Designed for home users and small office users, the Files and Settings Transfer Wizard is also useful in a corporate network environment for employees who get a new computer and need to migrate their own files and settings without the support of an IT department or Help desk. For information about using the Files and Settings Transfer Wizard, see Help and Support Center for Windows XP.

**User State Migration Tool (USMT)**
Designed for IT administrators who are performing large deployments of the Microsoft® Windows® XP Professional operating system in a corporate environment, USMT provides the same functionality as the Files and Settings Transfer Wizard, but on a large scale targeted at migrating multiple users. USMT enables administrators to precisely configure unique settings, such as making user-specific modifications to the registry. The tool is included on the Windows Server 2003 operating system CD in the \ValueAdd\Msft\USMT folder.

USMT uses the following files in collecting and migrating user data and settings:

- Scanstate.exe collects user state.
- Loadstate.exe restores user state.
- Migapp.inf determines which application settings are migrated.
- Migsys.inf determines which operating system settings are migrated.
- Miguser.inf determines which user settings are migrated.
- Sysfiles.inf defines files that *must not* be migrated despite any other rules. These are operating system files that will conflict with the newer version of the files in Windows XP. The SysFiles.inf file should not be modified except to add more files to the list of files that never migrate under any circumstances.

These files are shipped with Windows XP in the ValueAdd\Msft\USMT folder.

Table 5.1 and Table 5.2 list the file types, folders, settings, and system components that are migrated by default using USMT. (See also the Inf Commands.doc file included on the Windows Server 2003 operating system CD in the \ValueAdd\Msft\USMT folder.)

**Table 5.1 File Types and Folders Migrated by Default by USMT**

| File Types Migrated | Folders Migrated |
|---|---|
| .ch3 | Desktop |
| .csv | My Documents |
| .dif | My Pictures |
| .doc | Favorites |
| .dot | Cookies |
| .dqy | |
| .iqy | |
| .mcw | |
| .oqy | |
| .pot | |
| .ppa | |
| .pps | |
| .ppt | |
| .pre | |
| .rqy | |
| .rtf | |
| .scd | |
| .sh3 | |
| .slk | |
| .txt | |
| .wk? | |
| .wpd | |
| .wps | |
| .wq1 | |
| .wri | |
| .xl? | |

**Table 5.2 Settings and System Components Migrated by Default by USMT**

| Settings and System Components Migrated |
|---|

| |
|---|
| Accessibility options |
| Classic Desktop settings |
| Dial-up connections |
| Display properties |
| Folder options |
| Fonts |
| Microsoft® Internet Explorer settings |
| Localization/International settings |
| Microsoft® Office settings |
| Mouse and keyboard settings |
| Network drives and printers |
| Microsoft® Outlook® settings and store |
| Microsoft® Outlook® Express settings and store |
| Phone and modem options |
| Regional options |
| Screen saver selection (not users' personal screen saver files) |
| Shortcuts (shell tools, network items, and so forth) |
| Sounds and audio devices settings |
| User certificates (personal, e-mail, Internet Explorer security, and so forth) |
| Taskbar settings |

USMT offers multiple customization options for including various file types and settings in the user state migration. Administrators should expect to customize the default set of data and settings. Customization should be performed by technical personnel with knowledge of the registry.

This chapter describes how to plan and test a user state migration, but does not describe how to use the USMT tool. For code samples to assist you in customizing the .inf files used with USMT and the Files and Settings Transfer Wizard, see the file Inf Commands.doc included on the Windows Server 2003 operating system CD in the folder \ValueAdd\Msft\USMT.

For more information about using USMT, see the User State Migration Tool link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Choosing the User State Collection Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in planning your user state migration is to determine the best way to collect user state in your environment (Figure 5.2). Four user state migration methods are available for collecting and restoring user state:

- Manual migration
- Scripted-manual migration
- Centralized automation
- User-driven migration

**Figure 5.2 Choosing a User State Collection Method**




The method by which you deploy Windows XP affects which user state migration method you should choose. Table 5.3 explains how each system deployment method affects the environment into which the user state will be migrated. A clean environment reduces management and support requirements.

**Table 5.3 Effects of Windows XP Deployment Methods on User State Migration**

| Deployment Method | Effects on User State Migration |
|---|---|
| **Wipe-and-load**. The computer's hard drive is reformatted before installing Windows XP. This is the recommended deployment method when the existing hardware is sufficient to run Windows XP, because it provides a clean platform on which to restore applications and settings. | Presents a completely clean environment in which to restore user state. |
| **Parallel deployment**. The original computer is replaced with a new computer running Windows XP. This is the recommended deployment strategy when the old computer has insufficient hardware capability to run Windows XP. You can keep the original computer running until you are sure that the new computer is completely functional. | Presents a completely clean environment in which to restore user state. Commonly used when deployment of Windows XP is timed with computer replacement, as in lease rollover. |

| | |
|---|---|
| **Operating system upgrade.** The original computer is upgraded using the Upgrade option during the setup phase of Windows XP deployment. This leaves the user's files, folders, settings, and installed applications intact. | Does not provide a clean environment, thereby increasing support and management costs. The migration of System Policy, registry settings, files, drivers, DLLs, and folder hierarchies can cause problems and nonstandard installations. Not recommended in production environments. |

When determining the best user state collection method for your situation, weigh these factors:

- The size of your organization
- The number of users to be migrated
- The level of desktop management already in place
- The uniformity of file locations on workstations
- The type of technical personnel available to assist in the migration
- The amount of time you can dedicate to the migration process

Each migration method is particularly well suited to specific scenarios. It is likely that a large corporate deployment will involve several of these scenarios and employ a mixture of migration methods.

# Manual Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a manual migration, an onsite technician personally attends each computer, typically performing these tasks:

- Ensures that the user's computer is ready for migration (for example, checks to see whether all important files are in the folders that are being migrated).
- Collects the user's state by running either USMT (the Scanstate.exe command-line tool) or the Files and Settings Transfer Wizard.
- Deploys Windows XP either by providing a new computer running Windows XP or by doing a wipe-and-load deployment of Windows XP. (Remote Installation Services [RIS] provides a convenient way to deploy a common Windows XP image.)
- Restores the user state by again running either USMT (the Loadstate.exe command-line tool) or Files and Settings Transfer Wizard. The same tool that is used to collect user state must be used to restore it.
- Is available to help with any issues while the user checks to make sure that everything has been migrated properly.

Because technical labor costs in manually collecting state data can be very high, it is often beneficial to combine manual collection with the use of automated scripts.

Table 5.4 summarizes advantages and disadvantages of manually migrating user state. Because of the noted disadvantages, a strictly manual approach is not recommended.

**Table 5.4 Advantages and Disadvantages of Manual User State Migration**

| Advantages | Disadvantages |
|---|---|
| • A technician is available to deal with unexpected problems.<br>• Users are reassured by having a person to ask questions of during the migration. | • Expensive because of high technical labor costs.<br>• Slow because a technician must visit each computer individually.<br>• Higher chance of human error than with automated methods.<br>• Does not scale to distributed or remote office scenarios. |

# Supplementing a Manual Migration with Scripting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By supplementing a manual deployment with scripting, you can reduce the costs of an exclusively manual approach while providing the flexibility to handle special situations.

Scripting speeds up the process of migrating user state on individual computers, enabling a technician to migrate user state on multiple computers in the same physical area simultaneously. This greatly reduces the likelihood of human error, yet maintains the advantage of having someone onsite to deal with unexpected problems. During scripted-manual migrations, you also can use less skilled technicians for the onsite phase of the migration.

Table 5.5 summarizes the advantages and disadvantages of using a scripted-manual approach for migrating user state. Scripted-manual migration is the recommended approach for a parallel deployment.

**Table 5.5 Advantages and Disadvantages of Scripted-Manual User State Migration**

| Advantages | Disadvantages |
|---|---|
| • A technician is available to deal with unexpected problems.<br>• Users are reassured by having a person to ask questions of during the migration.<br>• Lower chance of human error than with a purely manual migration method. | • Requires that a technician be onsite, with physical access to each computer that is being migrated.<br>• Requires that script files be created.<br>• Does not scale to distributed or remote office scenarios. |

The manual-scripted migration process for a wipe-and-load deployment, in which the computer's hard drive is reformatted before the new operating system is installed, is slightly different from the process for a parallel deployment.

**Migration Process in a Wipe-and-Load Deployment**

In a wipe-and-load Windows XP deployment, use this combination of scripted and manual steps:

1. Create scripts to collect and restore user state.
2. Have a technician perform the following tasks at each computer:
   a. Run the script for collecting user state.
   b. Format the computer's hard drive and run RIS to install the new operating system image.
   c. Log on as the administrator, and run the restoration script.
3. Have the computer's user log on and check for proper restoration of data and settings.

**Migration Process in a Parallel Deployment**

Perform these steps to migrate user state in a parallel deployment:

1. Create two scripts for each computer that is to be replaced, one for collecting user state and the other for restoring user state.
2. On the computer that will be replaced, run the script for collecting user state.
3. On the new computer:

a. Install Windows XP.

b. Log on as the administrator, and run the script for restoring user state.

4. Have the computer's user log on to the new computer and check for proper restoration of data and settings.

# Centralized Automation of User State Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With centralized automation, you can extend the efficiencies available through the scripted-manual method for migrating user state. To centralize automation of user state migration, you refine the user state collection and restoration scripts to such a degree that no onsite input from a technician is required. IT technicians can deploy the scripts to targeted computers from a remote location.

Centralized automation enables enormous cost savings and provides a common migration experience corporation-wide. Centralized automation does not work well in parallel deployments because of the complexity in determining target and destination computer addresses for a large number of computers, but it is ideal for wipe-and-load deployments.

Table 5.6 summarizes the advantages and disadvantages of using centralized automation to migrate user state. Centralized automation is the ideal solution for wipe-and-load deployments.

**Table 5.6 Advantages and Disadvantages of Centrally Automating User State Migration**

| Advantages | Disadvantages |
|---|---|
| • Allows simultaneous migration of user state for large numbers of computers (limited only by network bandwidth and server storage).<br>• Produces results that can be replicated.<br>• Produces a common user experience.<br>• Scales well to distributed or remote office scenarios. | • Requires that script files be created.<br>• Does not work well in parallel deployments. |

The key challenges in the centralized automation of user state migration are:

- Targeting and deploying scripts so that they run on the user's computer in the appropriate context.
- Associating user state with a specific computer.
- Writing scripts that will create a temporary store for each user's state and then restore that state on the destination computer.
- Automatically deploying a Windows XP image on remote computers without user intervention.

**Targeting and deploying scripts to run in the appropriate context** Follow these rules when targeting and deploying automated scripts during user state migration:

- The collection script must run under the user's logon account.
- The restoration script must run securely in the administrator's context, without user intervention at the remote computer during log on.
- The computer's user should not be using the computer when the script is run.
- No applications can be running when the script is run.

Several options are available for automatically running the script at a specific time and under in the appropriate context. It is best to use a management solution such as Microsoft® Systems Management Server (SMS) for this. SMS provides advanced targeting options, contains software deployment structural components, and can target packages to run at specific times in specific contexts. Other options include logoff scripts, e-mail that includes the script (which automatically shuts down the mail

client), or deployment automation delivered by way of a Web site.

For more information about Systems Management Server, see the SMS Product Information link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Associating user state with a computer** When deploying a standard image to a series of computers, plan how to discover which user's state to restore to which computer. Two common methods for identifying computers are by the media access control (MAC) address for the network adapter or by the serial number of the processor.

For example, you can write a script that collects both the network adapter MAC address of the computer and the logon name of the user. This data pair is stored in a central database along with the mapping of the user state storage location. When restoring user state, the script looks up the network adapter MAC address to find the user's logon name and user state storage location, and restores the user state to the appropriate computer.

**Storing and restoring the user state** In centralized automation, you can use the same scripts that collect and restore user state in the scripted-manual method, with these adjustments:

- The collection script must be able to create a separate subdirectory for storing each user's state during the migration. Appending the user's logon name to the root storage path (for example, \\State*Username*) is a good solution. If the user has multiple computers, use both the computer name and the user's logon name (for example, \\State*Username**Computername*).

- The restoration script must read the user state storage path from the central database that the collection script wrote to and restore the user state from the appropriate storage location.

**Automating Windows XP image installation** Microsoft offers several options for deploying operating systems. For information about your options for automating the deployment of a Windows XP image, see "Choosing an Automated Installation Method" in this book.

# User-Driven State Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Perform a user-driven user state migration when no central management is in place for a deployment, or when users connect to the organization's network from a remote location.

In a user-driven migration, have the user run the Files and Settings Transfer Wizard to collect and restore user state. The user can use the same customized .inf files that you might use with USMT in other types of migrations. With the Files and Settings Transfer Wizard, the user can decide which applications and components to migrate (although you set the defaults), and can add and remove files and folders from the set to be migrated. Be prepared to offer these users some training on using the wizard and the .inf files.

Table 5.7 summarizes the advantages and disadvantages of a user-driven migration of user state. User-driven migration is the recommended approach for nonmanaged environments and remote users.

**Table 5.7 Advantages and Disadvantages of User-Driven User State Migration**

| Advantages | Disadvantages |
|---|---|
| • Controlled tool better enables users to drive their own migration.<br>• Can use the data captured using customized .inf files. | • IT staff are not involved in the actual migration; relies on the user's judgment.<br>• No standard storage location.<br>• Minimal control over what is migrated, resulting in a less standardized desktop. |

# Identifying Migration Content

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After determining how to perform the migration, identify which user files and folders to migrate, and then identify key user settings to migrate, as shown in Figure 5.3. Consider which data and settings are worth migrating and which do not provide sufficient benefits to justify the cost of migration.

**Figure 5.3 Identifying Migration Content**




Part of identifying migration content is recognizing that you must provide a transitional storage place for everything that you migrate. Be sure that you have enough storage capacity to store the aggregate content during the migration. For information about estimating required storage capacity, see "Determining Storage Requirements" later in this chapter.

# Identify User Data to Migrate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When identifying what data to migrate, consider these questions:

- Do users save their files in a single folder or in files scattered across a hard disk?
- Which data files do the users work with regularly?

One way to determine which user data to migrate is to identify folders to migrate based on known locations. These can be locations that the system is aware of, such as the My Documents folder and Favorites, or locations that the organization-specifies, such as \EngineeringDrafts or C:\Data.

Another way to determine which user data to migrate is to identify the applications that the users use and then look for files with corresponding file types. Organizations commonly use an e-mail package and productivity suite such as Microsoft Office. These applications typically use specified file name extensions. For example, Microsoft® Word primarily uses the .doc file name extension. However, Word also uses file types such as templates (.dot files) and hypertext files (.htm files).

If you use this method to identify files to migrate, create a list of important file types based on applications that your organization uses. A good starting point for identifying the file types to migrate is to look at the registered file types on the standardized desktop image that you will install. The registered file types are listed in **Folder Options**.

**To view a list of registered file types**

1. Double-click the **My Computer** icon on the desktop.
2. On the **Tools** menu, click **Folder Options**.
3. Click the **File Types** tab to display the registered file types.

**Important**

- Do not attempt to migrate the applications associated with the files. Instead, reinstall the applications from a software distribution point, or include them in the standard desktop image.

# Identifying User Settings to Migrate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Consider the following questions when identifying which user settings to migrate:

- Are you moving toward a more managed environment? If so, which settings will users be able to change in this new environment?
- Which settings do the users need to get their work done?
- Which settings make the work environment comfortable for users, allowing them to be more productive?
- Which settings will reduce help desk calls after the migration?

# Identifying Key Settings for User Productivity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

List the important settings that the user needs to become productive immediately after the migration. These settings might include an e-mail server and account, a remote access connection, an Internet connection, and accessibility features. One good place to find the relevant settings is in your organization's system configuration handbook for new users.

Locating application-specific settings can be time-consuming, because various applications store settings in different locations. Therefore, limit your list to settings that the user must have to maintain productivity.

Some applications provide tools that scan the registry and then display settings and their storage location in a format that is easy to read. For other applications, you must compare registry entries before and after an installation to trace the settings.

Typically, the user settings for an application are stored in the registry in the subkey HKEY_CURRENT_USER\SOFTWARE*Companyname**Application*. You can use the Sysdiff.exe tool to compare images of the registry before and after migration. Sysdiff.exe finds updated registry entries. The Sysdiff.exe tool is documented and available for download from the Resource Kit Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

# Evaluating Costs vs. Benefits of Migrating Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is not always cost-effective to collect and restore all user-specific settings in the registry. In deciding which settings to migrate, weigh the cost of lost productivity while users recreate their settings against the IT costs of migrating them.

**To determine the cost-effectiveness of collecting settings:**

1. Perform the following calculations:

    a. Determine the number of users whose user state will be migrated.

    b. Multiply that number by the average time that it takes a user to reconfigure his or her settings.

    c. Multiply the result by the users' average hourly wage.

2. Compare this cost with the costs involved in tracing, collecting, and restoring the settings.

Table 5.8 shows user cost calculations for an enterprise with 5,000 users who have an average hourly wage of $20, based on an estimated reconfiguration time of 2 hours per user. The total user costs for not migrating user settings are compared with estimated IT costs for migrating the settings, shown in the final column.

**Table 5.8 Sample Cost Comparison for Migrating vs. Not Migrating User Settings**

| Number of Users | User Reconfiguration Time | Average Hourly Wage | Total User Costs (No Migration) | Estimated IT Costs for Migration |
|---|---|---|---|---|
| 5,000 | 2 hours | $20.00 | $200,000.00 | $50,000 |

Less measurable, but equally important, is the time that IT professionals and the Help desk staff spend helping users reconfigure desktops when their personal settings are not migrated

Not migrating settings also can lead to lost productivity and decreased morale. Familiar desktop settings help users learn about the new system more easily, and reduce the potential for help desk calls. While the value of migrating some personal settings might not be immediately apparent, it is worthwhile to capture the settings that make new computers familiar and comfortable to users — for example, desktop settings, display settings, and folder options. If such settings are not migrated, productivity can decrease while users adjust these settings to suit them.

# Creating a Detailed Migration Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After selecting a collection method and identifying the content to migrate, review the technical details involved in a successful migration and devise a detailed migration plan that addresses storage and security issues, resolves changes in registry settings from earlier versions of Windows, and includes adjustments for concurrent domain migration if applicable. In your migration plan, include a migration schedule and detailed instructions for users. Figure 5.4 shows the tasks involved in creating a detailed plan for user state migration.

**Figure 5.4 Creating a Detailed Migration Plan**



# Resolving Storage and Data Issues

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because the process of migrating user state consists of moving data, the first step in preparing for your migration is to identify and resolve storage and data concerns associated with the move. Determine how much temporary storage space is required for the data that is being moved. Verify that your selection of user data and settings to be migrated is complete as well as cost-effective. Plan how to resolve any file name conflicts or other file relocation issues that might arise during the migration.

# Determining Storage Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Determine how much disk space is required for an temporary storage location for user state. Base your calculations on the volume of e-mail, personal documents, and system settings for each user. The best way to estimate these is to survey a few average desktops to estimate the size of average stores in your environment.

**Important**

- Allow a minimum buffer of 20 percent additional space in the temporary storage location. To enhance performance, locate the temporary store on high-speed drives. Ensure that the temporary storage of user state is the only task that the store performs and that it has an optimized (high-speed) network connection.

**E-mail** If users deal with a large volume of e-mail or keep e-mail on their local computers instead of on a mail server, the e-mail can take up as much disk space as all other user files combined. (This is not a factor if the e-mail is stored on a server only.) Prior to migrating user data, make sure that users who store e-mail locally synchronize with their mail server.

**User documents** The types of documents that an organization uses can make a substantial difference in storage requirements. For example, an architectural firm that predominantly uses Computer-Aided Design (CAD) files needs much more space than does a law firm dealing primarily with word processing documents. If your users already store many documents on file servers through such mechanisms as Folder Redirection, and they will have access to these locations after the migration, you do not need to migrate those documents.

**User system settings** Typically, 5 megabytes (MB) of storage is adequate to save a user's registry settings. This requirement can fluctuate based on the number of applications installed over the lifetime of the computer, but it is rare for the user-specific portion of the registry to exceed 5 MB.

# Reviewing Data Collection and Restoration Selections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are using the USMT, in most cases you will customize the .inf files included with the tool to limit the operating system and application settings that are migrated and to include additional file types and folders.

When customizing the .inf files, it is important to keep a backup copy of the original files and to thoroughly test your customizations. To yield the best results, keep the rules in the .inf file as simple as possible.

For information about the default files and settings that USMT migrates, see "Tools Used for the Migration Process" earlier in this chapter.

For instructions and code samples to assist you in customizing the .inf files used by USMT, see the User State Migration Tool link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Addressing Issues with File Relocation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you change the computer hard disk configuration during migration, you might not be able to restore files to the same drive or directory structure from which they were collected. For example, if you replace two small drives with one large drive, the large drive will not be available to receive the collected user data. In this case, you must relocate the files.

A relocated file might be written to a folder that already contains a file with the same name, causing a name conflict. USMT handles this problem by appending "(1)" to the original filename, and incrementing that number for each new file with the same name. For example, if two files by the name of Example.doc were written to a directory that already contained an Example.doc file, the relocated files would be named Example(1).doc and Example(2).doc.

One way to avoid file name collisions when you move files is to duplicate as much of the original path as possible in the new location. For example, if the full path and file name of the original file was D:\EngineeringDrafts\Example.doc, and the new root location is C:\Documents and Settings*Username*\My Documents, create the new path and file name C:\Documents and Settings*Username*\My Documents\EngineeringDrafts\Example.doc.

# Identifying Security Concerns for Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Maintaining security during and after user state migration is a significant issue. In particular, take into consideration these issues:

- Typically, the Access Control Lists (ACLs) associated with files and folders are not migrated, so the ACLs must be restored or recreated.
- Encrypted File System (EFS) information is not migrated, and encryption that occurs during migration affects who can read the files in their new destinations.
- In some organizations, it is deemed critical to secure user state during a migration.

# Restoration of Lost Access Control Lists (ACLs)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In planning user state migration, it is best to assume that access control lists will not migrate during your user state migration. Several factors affect the migration of ACLs:

- The USMT tool and the Files and Settings Transfer Wizard do not migrate ACLs — instead, default ACLs are assigned to each folder that is created on the destination computer.

- If users are changing domains during a migration, there is a good chance that the original ACLs will not work unless you use a tool such as SIDHistory as part of the user state migration process. For information about managing access control lists during a domain migration, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

- When you migrate a Windows NT workstation that uses an NTFS file system drive, ACLs for individual files often do not migrate with the files. Instead, the files inherit the default ACLs of the folder into which they are copied.

# Managing Data Encryption During Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Encrypted File System (EFS) certificate data is not migrated when you use either USMT or the Files and Settings Transfer Wizard. The two tools treat encryption differently during a user state migration:

- The Files and Settings Transfer Wizard decrypts encrypted files during migration, and does not encrypt the files when it writes them to the destination computer (unless writing them to a folder that is encrypted).

- USMT decrypts encrypted files during migration, but if the temporary store is encrypted, the file will be encrypted under the user's credentials (because Scanstate.exe is run in the user's context). In addition, if the destination folder for the migrated file is encrypted, the restored file might be encrypted and, because the file will have been written under the administrator's credentials, the administrator, not the user, will be able to read the file.

In general, assume that files are not protected by encryption during a user state migration. Furthermore, because EFS certificates are not migrated, if a file does get encrypted during the migration, the user will not be able to read the file unless the EFS certificate is recovered from the network. For information about performing this type of operation, see "Encrypting File System" at http://go.microsoft.com/fwlink/?linkid=4514.

# Keeping User State Secure During Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In some organizations, keeping the user's state secure from the IT technician who is performing the migration is a potential issue.

If the IT technician's access to user state is a security concern for you, take these steps:

- Have the user drive the migration using either USMT or a scripted-manual method. Under the scripted-manual method, the user must be able to restore user state by logging on as the administrator.

- When securing the state in the temporary store, make sure that while the root folder might allow full user access, the individual user folders only allow access for IT staff and the owner of the folder.

- To protect data as it traverses the network, use Internet Protocol security (IPSec) or other network security protocols to secure these transfers.

# Translating and Relocating Registry Entries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because the name or location of some registry entries for the operating system has been changed in later versions of Windows, many registry values must be translated during migration, and others must be relocated within the registry. This is also true with different versions of some applications. For example, copying the subkey HKEY_CURRENT_USER\Control Panel\Desktop\WindowsMetrics during migration causes problems, because entries such as **IconFont** are not translated correctly.

USMT automatically translates and relocates the operating system settings for the user state that it migrates. To prevent problems with custom settings, either do not migrate entries that are unique to a specific version of an application and cause problems, or use the renaming and relocating capabilities of USMT to adjust the entries. You cannot use USMT to translate registry values: The best solution for this type of change is to write custom code.

You can use tools such as Sysdiff.exe to compare before and after images of the registry. This tool helps in finding registry changes between versions of an application or between the same version of the application running on different versions of Windows. Sysdiff.exe is documented and available for download from the Resource Kit Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Windows XP registry restrictions**

The enhanced security of Windows XP can mean that registry settings that were accessible in the Microsoft® Windows® 95 or Microsoft® Windows® 98 operating systems are no longer accessible.

In Windows XP, a user with no administrative permissions has write access to only three locations (depending on default security settings): the HKEY_CURRENT_USER registry subtree, the User Profile, and a shared documents location. Users cannot change settings outside those locations without administrative permissions.

If an application writes settings outside HKEY_CURRENT_USER, users will not be able to run the application after migration. You should deal with these applications on a case-by-case basis. Sometimes it is acceptable to change the access rights to a part of the registry; at other times, this can grant to the user unacceptable access to the registry. The best solution is to work with the software vendor or in-house developer to determine migration requirements when introducing a revised version of an application.

# Adapting a Plan for Domain Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain migration introduces additional issues for user state migration. Command-line parameters available in USMT make it easy to change the domain name and the user name during a migration. However, these switches do not resolve issues related to security identifiers (SIDs) during domain migration.

The Active Directory Migration Tool (ADMT), which is used to migrate to the Microsoft® Windows® 2000 Active Directory® directory service, addresses issues that USMT does not handle. This tool can help you diagnose possible problems before starting migration operations.

ADMT provides the following capabilities:

- Using task-based wizards, you can migrate users, groups, and computers; set correct file permissions; and migrate mailboxes for Microsoft® Exchange Server.
- Using the tool's Reporting feature, you can assess the impact of the migration both before and after move operations.
- The tool also provides support for parallel domains, so that you can maintain your existing domains under the Microsoft® Windows NT® version 4.0 operating system while you deploy a later version of the Windows operating system (Microsoft® Windows® 2000; Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; or Windows® Server 2003, Datacenter Edition).

For information about issues related to the migration of domain user objects and how domain migration affects the migration of user state, see the ADMT Cookbook link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Scheduling the Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Consider the impact of the migration on network bandwidth. The user state traverses the network as it is being moved to a temporary storage location on a network server and again when moved to the destination computer.

To minimize the impact of your migration on users and your network, follow these general guidelines:

- Minimize the impact on the network while other people need to use it by scheduling migrations during off-peak hours.
- Trying to move too many users at a time risks network collisions. Determine the optimal number of users to migrate at a time.
- Minimize the use of network capacity by locating storage servers close to the clients (for example, on the same subnet).
- Keep in mind the simple impact of migration on your normal flow of business. Work with the teams that you are migrating to ensure that the migration will not jeopardize any crucial projects. Determine whether teams need to be migrated as a group.

# Educating Users

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To minimize productivity loss and support costs, prior to migration, set user expectations to match the results that you obtained during pilot testing. Set expectations early and clearly to reduce user frustration and Help desk calls.

Provide a schedule that indicates when each user's computer will be migrated, as well as clear guidelines that tell exactly what the user needs to do to prepare for the migration. Your migration plan must include backing up user files, verifying that applications that use synchronization mechanisms (such as e-mail) are also backed up, and preparing for changes to the desktop.

**Preparing files and folders for migration** The period just before migration is a good time for users to get their files into a stable state. For example, if version control software is in use, make sure that all users check in all files that they have checked out. If users are supposed to save all in-progress documents in a specific network folder, make sure that the users save all relevant files to that folder.

If the user state collection process will retrieve data only within a known folder (for example, My Documents), have users move all of their important documents to that folder. If the folder that contains all of these files is a network share, no migration of the files is needed as long as the user has access to that share from the new system.

**Preparing e-mail and other applications that must be synchronized** It is recommended that users send all pending e-mail prior to migration. Along with e-mail, My Briefcase, Microsoft Outlook, Microsoft® Notes, and any other application or feature that uses synchronization must be synchronized prior to the migration.

**Preparing users for changes to the desktop** The more closely that the users' new environment mirrors their previous one, the less support they will need, and the sooner they can resume productivity. If the migration involves changes to the desktop, prepare the users for these changes.

If you will not migrate specific settings, tell users in advance which settings they will need to reenter and which related files they might need to migrate (for example, a personal photograph used as a background on the user's desktop).

# Testing Your Migration Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before rolling the migration out to a large number of computers, test the process in a controlled laboratory setting, and then run a pilot test. Figure 5.5 shows the tasks in the testing process.

**Figure 5.5 Testing Your Migration Process**




Even with thorough testing before the migration, you can often improve the process by making adjustments as the migration proceeds. It is best to move groups of users in phases. Migrate one group or team and make sure the migration is successful before starting the next group. This gives you a chance to modify your plan as needed between groups.

# Performing Lab Tests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the lab test, match your test environment closely to your production network:

- Use the same type of server in your test environment as in your real environment. If you will migrate users from an old domain to a new domain, include both the old and new types of servers in your test environment.

- Run a test migration from at least one computer running each operating system from which you will migrate. For example, if you need to migrate user state from some computers running Microsoft® Windows® 98, some running the Microsoft® Windows NT® Workstation version 4.0 operating system, and some running the Microsoft® Windows® 2000 Professional operating system, test at least one computer running each operating system.

- Make backups of the data residing on the computers from which you are migrating user state so that you can easily reproduce any problems that you encounter. If you adjust a custom script to solve a problem, it is hard to know whether the change solved the problem if you cannot reproduce the problem.

# Performing a Pilot Test for Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After thoroughly testing your user state collection, operating system deployment, and user state restoration processes, conduct a pilot test on a small group of users in a production environment.

In the pilot test, concentrate on these areas:

- Make sure that all data and settings migrate as expected.
- Note the storage space requirements for the pilot data and adjust your initial calculations accordingly.
- If unexpected problems arise, address them before going further.

As you did during lab testing, make backups of the data on your source computers so that, for testing purposes, you can easily reproduce any problems that you encounter.

Only after you are fully satisfied with the success of your pilot test should you begin a full migration.

# Additional Resources for Migrating User State

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Performing a Pilot Test for Your Migration

These resources contain additional information related to user state migration for Windows XP.

**Related Information**

- "Designing RIS Installations" in this book for information about using Remote Installation Services (RIS).
- For information about scripting in a Windows Server 2003 environment, see the at http://www.microsoft.com/reskit, or see the MSDN Scripting Clinic link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The file Inf Commands.doc, included on the Windows Server 2003 operating system CD in the folder \ValueAdd\Msft\USMT.

**Related Tools**

- User State Migration Tool (USMT)

  This command-line tool is used to collect a user's documents and settings before an operating system migration to Windows XP from an earlier version of Windows and to restore them after the installation. For a reference to the commands and syntax employed in the .inf files that are used to customize the selection of files, settings, and registry entries migrated by USMT, see the User State Migration Tool link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Files and Settings Transfer Wizard

  The Files and Settings Transfer Wizard enables users to migrate personal display properties, folder and taskbar options, and Internet browser and mail settings, as well as specific files or entire folders, such as My Documents, My Pictures, and Favorites, without manual configuration. For more information about the Files and Settings Transfer Wizard, see Help and Support Center for Windows XP.

- Sysdiff.exe

  Sysdiff.exe is an automated installation tool that enables you to pre-install applications, including applications that do not support scripted installation, as part of an automated setup. For more information about Sysdiff.exe, see the Sysdiff.exe link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Windows 2000 Active Directory Migration Tool (ADMT)

  The ADMT is used to migrate to Windows 2000 Active Directory, providing a task-based wizard that is used to migrate users, groups, and computers; to set correct file permissions; and to migrate Microsoft Exchange Server mailboxes. For information about ADMT, see the Active Directory Migration Tool (Windows 2000) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Designing a Managed Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Welcome to *Designing a Managed Environment* of the *Microsoft Windows Server 2003 Deployment Kit*.

This book provides the foundation you need to understand and deploy Group Policy and its supporting technologies, included within Microsoft Windows Server 2003. The guidance in this book will help you reduce the cost of managing a diverse set of organizational and computing needs, and to centrally manage domain-wide and forest-wide security, servers, computers, and users.

This book contains the following chapters:


Planning a Managed Environment

Designing a Group Policy Infrastructure

Staging Group Policy Deployments

Deploying Security Policy

Migration of User State

Implementing User State Management

Deploying a Managed Software Environment

Deploying a Simple Managed Environment

# Planning a Managed Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many organizations are moving towards a highly managed computing environment based on a configuration management infrastructure that is designed to reduce the cost of managing a diverse set of organizational needs. The Microsoft® Windows® Server 2003 operating system includes a set of configuration management technologies — collectively known as IntelliMirror® — that you can use to centrally manage configurations for users and computers.

IntelliMirror and Group Policy combine the advantages of centralized computing with the performance and flexibility of distributed computing. This book explains how to design and deploy a centrally managed distributed computing environment.

**In This Chapter**


Overview of Deploying a Managed Environment Process

Developing a Deployment Plan

Defining the Scope of Your Configuration Management Project

Assessing Your Current Environment

Determining Your Requirements for Configuration Management Technologies

Designing Managed Configurations

Designing Your Organizational Unit Structure

Testing Your Configuration Management Design

Staging and Deploying Your Design to the Production Environment

Additional Resources for Planning a Managed Environment


**Related Information**

- For more information about designing and deploying your Active Directory® directory service infrastructure, see "Planning an Active Directory Deployment Project" and "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

- For more information about Group Policy, see the *Windows Security Collection* of the Windows Server 2003 Technical Reference.

# Overview of Deploying a Managed Environment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IntelliMirror technologies included in Microsoft® Windows® 2000, Microsoft® Windows® XP, and Windows Server 2003 operating systems provide the capabilities for directory-based configuration management. By using IntelliMirror on the server and the client, you can ensure that a user's data, applications, and settings remain constant throughout the user's environment. You also use Group Policy to centralize the process for deploying and managing security for servers running Windows 2000 and Windows Server 2003 and clients running Windows 2000 and Microsoft® Windows® XP Professional. Together, Active Directory and Group Policy provide the foundation for IntelliMirror. Based on the Group Policy settings you specify, IntelliMirror can deploy, recover, restore, and replace user's data, software, and personal settings. You can set policy definitions once and rely on the system to apply the policy settings without further administrative intervention. By using IntelliMirror, you can centralize and simplify the management of users, servers, desktops, and security.

The term *configuration management* as used in this guide refers to a subset of technologies in Windows Server 2003 (and Windows 2000) that administrators can use as part of their overall infrastructure for change and configuration management. By using the IntelliMirror management technologies provided in Windows 2000 and later, you can design configurations for servers, desktops, and users.

Microsoft has developed guidelines you can use to help you effectively design, develop, deploy, operate, and support solutions built on Microsoft technologies. The guidelines are organized into two integrated frameworks, Microsoft Operations Framework (MOF) and Microsoft Solutions Framework (MSF). The guidelines include white papers, operations guides, assessment tools, best practices, case studies, templates, support tools, and services. MOF provides technical guidance that enables organizations to achieve system reliability, availability, supportability, and manageability of IT solutions. MOF addresses the people, process, technology, and management issues pertaining to operating complex, distributed, heterogeneous IT environments.

For more information about MOF, see the Microsoft Operations Framework (MOF) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. For more information about Microsoft Solutions Framework, see the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Advantages of Using IntelliMirror**

IntelliMirror refers to the ability to provide users with consistent access to their applications, application settings, roaming user profiles, and user data, from any managed computer – even when they are disconnected from the network. IntelliMirror is delivered via a set of Windows features that enable IT administrators to implement standard computing environments for groups of users and computers.

IntelliMirror can significantly boost user productivity and satisfaction by doing the following:

- Allowing users to continue working efficiently in intermittently connected or disconnected scenarios by enabling uninterrupted access to user and configuration data under these conditions.
- Delivering a consistent computing environment to users from any computer when their desktop or laptop computer is unavailable or in scenarios where users are not assigned a specific computer.
- Minimizing data loss by enabling centralized backup of user data and configuration files by the IT organization.
- Minimizing user downtime by enabling automated installation and repair of applications.

Implementing IntelliMirror also boosts administrator efficiency and reduces IT costs by doing the following:

- Eliminating the need to manually configure user settings, install applications, or transfer user files to provide users access to their computing environments on any computer.
- Enabling scenarios where users don't have an assigned computer but log on to any available computer in a pool of computers. This helps reduce hardware and administration costs.
- Easing the IT task of implementing centralized backup of user files while satisfying need for these files to be available on

the user's computer.

- Reducing support costs by using Windows Installer to automatically repair broken application installations.

IntelliMirror is implemented by means of a set of Windows features, including Active Directory, Group Policy, Software Installation, Windows Installer, Folder Redirection, Offline Folders, and Roaming User Profiles.

Deploying IntelliMirror technologies for configuration management provides the following advantages.

**User data management** IntelliMirror facilitates the centralized backup and restore of user data. Storing the users data on servers simplifies data backup and security and ensures that data is protected and highly available. Users can access their data from any computer on the network. User data management relies on Active Directory, Group Policy, Folder Redirection, Offline Files, and disk quotas.

**User settings management** Using the user settings management capabilities, you can manage user settings and preferences (such as user-defined screen colors and display fonts) and make them available to users regardless of which networked computer the users log on to. You can centrally define managed configurations for users and mirror user settings to the network. User settings management relies on Active Directory, Group Policy, Roaming User Profiles, Folder Redirection, Offline Files, and Synchronization Manager.

**Security settings** Using Group Policy, you can manage the following types of security options for users, clients, servers, and domain controllers: Internet Protocol security (IPSec), security settings, software restrictions policies, and wireless network policies.

**Group Policy–based software installation** Using Group Policy–based software installation, you can manage software installation throughout an application's life cycle. You can centrally manage software installation, repairs, updates, and removal.

**Internet Explorer policy settings** Using Microsoft® Internet Explorer policy settings you can customize the appearance of the browser, define connection settings, and define custom Universal Resource Locators (URLs), security settings, and program associations.

**Scripts** You can specify scripts to run when the computer starts and shuts down and when users log on or log off.

**Remote setup and configuration** Using Remote Installation Services (RIS), you can simplify operating system installations by remotely setting up and configuring operating system installations on computers throughout the enterprise.

To ensure the success of your managed environment deployment, you must create a configuration management plan that defines the resources to be managed and the processes that you are implementing.

You start your deployment plan by defining your business objectives for change and configuration management and determining the scope of your configuration management deployment in terms of the users and resources you need to manage. In the first phase, you also need to perform an assessment of your current environment and determine the requirements for configuration management.

Next, determine which IntelliMirror configuration management technologies meet your business requirements, and then design managed configurations for the various types of users and computers in your organization based on those requirements.

Before you deploy configuration management technologies into your production environment, you must fully test your designs in a controlled test environment, and then review your test results to determine whether the systems meet the design requirements. It is strongly recommended that you stage your deployment of Group Policy and related technologies. For more information about staging, see "Staging Group Policy Deployments" in this book.

After you complete the testing of your design in a controlled environment, you can deploy to production. To learn about using the Windows Server 2003 Group Policy Management Console (GPMC) to facilitate testing and staging your deployments of Group Policy, see "Designing a Group Policy Infrastructure" in this book.

# IntelliMirror Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying IntelliMirror configuration management technologies entails several processes. You begin your deployment plan by establishing your business objectives and identifying the resources to be managed, and then you define the scope of your project and evaluate your current environment in terms of the resources you need to manage and your requirements for configuration management technologies. Based on these requirements, you design managed configurations, determine the scope of application of your solution, test, and then stage and deploy your solution.

**What this Guide Contains**

Planning how you organize users and computers in Active Directory containers and how you apply Group Policy settings to manage them is an integral part of IntelliMirror. To help guide your design decisions for creating managed desktop configurations, you need a clear understanding of the appropriate technologies to use and the processes involved in implementing a desktop management solution that meets your business needs.

Although it is anticipated that you will read only those chapters in this guide that apply to your particular deployment needs, be aware that the IntelliMirror technologies rely on Group Policy and Active Directory for their delivery. Therefore, it is recommended that you read "Designing a Group Policy Infrastructure" in addition to the chapters that pertain to your particular deployment.

Each chapter is intended to assist you in planning, designing, testing, deploying, and implementing managed desktop configurations using IntelliMirror technologies.

**Designing a Group Policy Infrastructure**

Provides information about the process of defining your Group Policy objectives and designing a Group Policy infrastructure. Read "Designing a Group Policy Infrastructure" before you deploy Group Policy–based solutions that are described in other chapters in this book.

**Staging Group Policy Deployments**

Discusses the processes for creating and testing your Group Policy deployments in a controlled test environment. For example, "Staging Group Policy Deployments" explains how to create a test domain that mirrors a production domain by using the sample scripts provided with the Group Policy Management Console. You can perform incremental policy changes in the test environment, verify your changes, and then use the **Import** functionality of GPMC and migration tables to migrate the Group Policy objects (GPOs) you created to the production domain.

**Deploying Security Policy**

Discusses incorporating the Windows 2003 Server Group Policy security features into your overall security strategy to protect your corporate environment. "Deploying Security Policy" explains how to evaluate existing security policies, determine the level of risk acceptable for your environment, and deploy security policies.

**Migrating User State**

Discuses how to use the user state migration tools included in Windows Server 2003 to save and restore user data and settings when you move client computers to the Windows XP operating system from earlier versions of Windows. "Migration of User State" guides you through the processes of planning and implementing user state migration to conserve IT staff time, preserve important data, and minimize costs while maintaining user productivity.

**Implementing User State Management**

Discusses the processes involved in implementing user data and user settings management technologies to set up, manage, and control the availability of user data and personal settings across your network. IntelliMirror allows data and settings to be available to users even when their computers are disconnected from the network. "Implementing User State Management" helps you plan, test, and configure managed desktop configurations by using Roaming User Profiles, Folder Redirection, Offline Files and synchronization, and disk quotas and user profile quotas.

**Deploying a Managed Software Environment**

Discusses the process of deploying a managed software environment using the software installation extension of Group Policy. By using Group Policy–based software installation, you create a controlled environment that provides on-demand software installation and automatic repair of applications. "Deploying a Managed Software Environment" explains how to use Group Policy–based software deployment to manage software throughout all phases of the software administration life cycle.

**Deploying a Simple Managed Environment**

Discusses the process of combining the IntelliMirror technologies to create a simple managed environment. "Deploying a Simple Managed Environment" guides you through the steps required to deploy a complete IntelliMirror solution, which you

can use used as an initial production environment for a small organization, or a test environment for a larger organization.

**Summary of Job Aids**

As supplement to this book, you can use job aids, such as worksheets and checklists, which are designed to help you collect data for planning your configuration management deployments. Each worksheet is created for the type of information you need for planning, such as questions to answer, points to consider, and tables to help you organize your information. You can modify the designated job aids to develop your configuration management designs as appropriate for your organization's needs. For worksheets to assist you with the deployment processes discussed in this book, see "Additional Resources for Planning a Managed Environment," later in this chapter.

# Introduction to IntelliMirror Configuration Management Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you use Windows XP and Windows 2000 clients, and Microsoft® Windows® 2000 Server and Windows Server 2003 networks with Active Directory installed, you can take full advantage of IntelliMirror and Group Policy management features. Active Directory and Group Policy provide the foundation for IntelliMirror. Based on the Group Policy settings you specify, IntelliMirror can deploy, recover, restore, and replace user's data, software, and personal settings.

Table 1.1 lists the IntelliMirror core features, the technologies that enable these features, and the advantages of using IntelliMirror.

**Table 1.1 IntelliMirror Features, Advantages, and Technologies**

| IntelliMirror Feature | Advantages | Technologies Used |
|---|---|---|
| User Data Management | Data is protected and highly available. Users can access their data from any computer on the network. For example, if you redirect specific user data folders, such as the **My Documents** folder, to a network location, and then making this location available to the users for offline use, users can access to their data wherever they log on.<br><br>You can also manage configuration settings for Internet Explorer by using Group Policy. | Active Directory<br><br>Group Policy<br><br>Offline Files<br><br>Folder Redirection<br><br>Synchronization Manager<br><br>Disk Quotas<br><br>Enhancements to the Windows shell |
| User Settings Management | Users get their preferred desktop configuration from any computer on the network. Their preferences and settings for the desktop and applications are available wherever they log on. | Active Directory<br><br>Group Policy<br><br>Offline Files<br><br>Roaming User Profiles<br><br>Enhancements to the Windows shell |
| Computer Settings Management | You can use Group Policy to define options for computers, including security and network settings. For example, you can set options such as remote assistance, system restore, Windows file protection, and Terminal Services. | Active Directory Users and Computers snap-in<br><br>Group Policy |
| Software Installation | You can centrally manage software installations, updates, repairs and removal. | Active Directory<br><br>Group Policy<br><br>Windows Installer |

# IntelliMirror Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory and Group Policy together provide the foundation for IntelliMirror. The following technologies enable the IntelliMirror features.

**Active Directory**
Windows–based directory service stores information about objects on a network and makes this information available to administrators and users. By using Active Directory, you can view and manage network objects on their network from a single location, and users can access permitted network resources by using a single logon. The *Active Directory Users and Computers*Microsoft Management Console (MMC) snap-in is the recommended tool for managing Active Directory objects, including organizational units (OUs), users, contacts, groups, computers, printers, and shared file objects.

**Group Policy**
The infrastructure within Active Directory that enables directory-based configuration management of user and computer settings on computers running Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems and Windows 2000 and Windows XP Professional operating systems. The Group Policy settings that you create are contained in a Group Policy object (GPO). To create a GPO, use the Group Policy Object Editor snap-in for the MMC, which can be started from the Group Policy Management Console MMC snap-in (GPMC).

The GPMC tool provides unified management of all aspects of Group Policy across an enterprise, including cross-forest management. Using GPMC, you can manage all GPOs, Windows Management Instrumentation (WMI) filters, and permissions in your network.

GPMC consists of the following: a set of scriptable interfaces for managing Group Policy, the Group Policy Modeling Wizard for planning Group Policy deployments prior to implementing them in the production environment, the Group Policy Results Wizard for viewing GPO interaction and for troubleshooting Group Policy deployments, and a new MMC snap-in that gives you the ability to manage Group Policy across your organization through a single user interface.

For more information about Group Policy and GPMC, see "Designing a Group Policy Infrastructure" in this book.

You can download GPMC from the Microsoft Download Center; see the Group Policy Management Console link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. You can use GPMC to manage both Windows 2000– and Windows Server 2003–based domains with Active Directory. By using GPMC to link a GPO to selected Active Directory system containers — sites, domains, and OUs — you can apply the GPO's policy settings to the users and computers in those Active Directory containers.

You can use Group Policy to configure policy settings for the following:

**Registry-based policy settings** The Administrative Templates extension contains all registry-based policy settings, including those for the Windows 2000 or Windows Server 2003 operating systems and their components as well as any registry-based policy settings provided by applications. You use these policies to mandate registry settings that control the behavior and appearance of the desktop, the operating system components, and applications that provide registry-based policy.

**Security settings** You can set security options for computers and users within the scope of a Group Policy object. You can define local computer, domain, and network security settings.

**Scripts** You can use scripts to automate tasks when the computer starts and shuts down, and when the user logs on and logs off. You can use any language supported by Windows Script Host. These include the Microsoft Visual Basic® development system, Scripting Edition (VBScript), JavaScript, Perl, and MS-DOS®-style batch files (.bat and .cmd).

**Folder Redirection** You can redirect special folders on Windows 2000 Professional and Windows XP from their default user profile location to an alternate location on the network. These special folders include **My Documents**, **My Pictures**, **Application Data**, **Desktop**, and the **Start** menu.

**Software installation** You can centrally manage software in your organization. You can assign and publish software to users and assign software to computers.

**Microsoft Internet Explorer Maintenance** You can manage and customize Internet Explorer on computers running Windows 2000, Windows XP, and Windows Server 2003. You can also export settings for Microsoft® Windows® 95, Windows® 98, and Windows NT® version 4.0 clients (the settings are exported into an .ins and .cab file format for those

platforms). Administrators can set options for browser user interface, connections, Uniform Resource Locators (URLs), proxy settings, security zones, and the **Favorites** folder and the **Links** bar.

**Remote Installation Services** Remote Installation Services (RIS) is used to control the behavior of the Remote Operating System Installation feature as displayed to client computers.

**User profiles**
A user profile describes the desktop computing configuration for a specific user, including the user's environment and preferred settings. The user profile is created when a user first logs on to a computer running Windows XP, Windows 2000, Windows NT, or Windows Server 2003. A user profile consists of a group of settings and files, which defines the environment that the system loads when a user logs on. It includes all the user-specific configuration settings, such as application settings, screen colors, network connections, printer connections, and mouse settings.

The following user profiles are available in Windows Server 2003, Windows XP, and Windows® 2000 Server and Microsoft® Windows® 2000 Professional:

- Local user profile. This is created the first time that a user logs on to a computer. The local user profile is stored on a computer's hard disk. Any changes made to the local user profile are specific to the computer on which the changes are made.

- Roaming user profile. A copy of the local profile is copied to and stored on a server share. This profile is downloaded every time that a user logs on to any computer on the network, and any changes made to a roaming user profile are synchronized with the server copy when the user logs off.

- Mandatory user profile. A type of profile that administrators can use to specify particular settings for users. Only system administrators can make changes to mandatory user profiles. Changes that a user makes to desktop settings are lost when the user logs off.

**Offline Files**
Offline Files complements Folder Redirection. By using Offline Files, users can disconnect from the network and continue working on the files as if they were still connected. When the computer is offline, the files and folders appear in the same directory as they do online — as if they still resided in the same location on the network. The user can then edit files even when not connected to the network. The next time the user connects to the network, the offline changes are synchronized with the network share.

**Folder Redirection**
By using Folder Redirection, users and administrators can redirect the path of a folder to a new location. The new location can be a folder on the local computer or a directory on a network share. Users can work with documents on a server as if the documents were based on the local drive. For example, you can redirect the **My Documents** folder (typically stored on the local hard disk) to a network location. The documents in the folder are available to the user from any computer on the network.

**Synchronization Manager**
When using Offline Files, users can use Synchronization Manager to synchronize all network resources. Users can set Synchronization Manager to automatically synchronize some or all resources. For example, users can set certain files and folders to be synchronized every time they log on or off the network. Synchronization Manager quickly scans the system for any changes, and if it detects changes, *only* the resources that have changed are updated, speeds up the synchronization process.

**Disk quotas**
You can set disk quotas to track and manage the amount of disk space used on the servers containing users' redirected folders. You can specify a disk quota limit and a disk quota warning level. If a user exceeds their designated quota limit, the user is prevented from storing additional files on the volume without first clearing some disk space. You can configure the disk quota system to log a system event when the users reach their quota warning level.

**Software Update Services**
SUS is a server component that you can install on a server running Windows 2000 or Windows Server 2003 to allow small and medium enterprises to bring critical updates from Windows Update inside their firewalls for distribution to computers running Windows 2000, Windows XP, and Windows Server 2003. You can use SUS to download the latest operating system patches to an intranet server, test the patches in your operating environment, select the patches you want to deploy to specific computers, and then deploy the patches.

**User State Migration tool**

If you are performing a large deployment of the Windows XP Professional operating system, use the User State Migration tool (USMT) to automate the migration of user state information. By using USMT, you can migrate users' personal display properties, folder and taskbar options, Internet browser and mail settings, as well as specific files or entire folders (such as **My Documents**, **My Pictures**, and **Favorites**) from their old computer to their new one. The USMT tool is included on the Windows Server 2003 CD in the \ValueAdd\Msft\USMT folder.

For more information about migrating user state information, see Migration of User State in this book. For more information about using USMT, see the User State Migration Tool link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Remote Installation Services**

You can use RIS to install the base operating system on a new computer or to replace a system that has failed. By using RIS you can do such installations or replacements without visiting the individual computer. Clients must support remote booting by means of the Pre-Boot eXecution Environment (PXE) ROM. For client computers that do not have a remote boot–enabled ROM, you can use a remote-startup floppy disk; these clients must have a Peripheral Component Interconnect (PCI)–based network adapter.

You can centrally set client configuration options for Remote Installation Services by using Group Policy. The remote operating system installation feature uses Active Directory, Group Policy, Dynamic Host Configuration Protocol (DHCP), and RIS.

To facilitate computer replacement, you can use RIS to install the operating system, Group Policy–based software installation to recover applications, Roaming User Profiles to restore user profiles, and Folder Redirection to manage files centrally.

You can use IntelliMirror features separately or you can combine them depending on your business and organizational requirements.

# IntelliMirror Implementation Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The examples that follow use some typical events at any organization to illustrate how the implementation of the IntelliMirror configuration management features in Windows Server 2003 impacts the computing environment. Example situations illustrate how the features and technologies of IntelliMirror can address typical desktop management needs.

# Setting up the New Employee's Computer Example

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A new user logs on to a new computer and finds documents and shortcuts already on the desktop. These shortcuts link to common files, data, and URLs that are useful to all users. Some examples of documents and shortcuts are the employee handbook or shortcuts to the user's departmental guidelines and procedures or to the intranet.

Also, if your computers come to you without an operating system, or if you have your own customized operating system installation, you can completely automate the installation of the client operating system for the new employee by using RIS. No technician is required to visit the computer. For more information about using RIS, see the Remote Installation Services (RIS) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Implementing the Setting up of a New Employee's Computer Example**

A default domain user profile and Group Policy are used to configure the new user's environment based on the user's job requirements. The administrator created a customized default domain profile that applies to all new domain users the first time they log on. When the users log on, they receive the customized settings from this profile. The advantage for the administrator of using a default domain profile is that all new users start from a base configuration that the administrator has chosen. Then, as the user personalizes the desktop settings, these settings are saved in the user's profile, which is stored either locally, or in a roaming user profile in a predetermined location on the network. By implementing roaming user profiles, you can provide the user with the required business information and settings whenever needed.

# Logging on the First Time Example

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A new employee logs on for the first time and sees that required software, such as Microsoft® Word, on the **Start** menu. When the new user selects Word on the **Start** menu, or double-clicks a Word document, Word is installed on demand.

**Implementing a First Time Logon Example**
Software Installation is implemented primarily through the use of Group Policy and the Windows Installer service. Based on the user's Active Directory location and the Group Policy object applied, the user is assigned Word.

IntelliMirror ensures that the necessary and correct items assigned to the user appear in the **Start** menu. Using Group Policy, the administrator can assign or publish differently configured versions of the same application to different groups of users. When the user starts Word or opens an employee handbook in Word for the first time, the Windows Installer service checks to see if the application is installed on the local computer. If not, Windows Installer downloads and installs the necessary files for Word to run and sets up the necessary local user and computer settings. Although Windows Installer continues to check each time the application is invoked, it does not attempt to reinstall or repair the application unless necessary files are missing.

# Making Offline Files Accessible to a Portable Computer Example

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A portable-computer user working at the office creates a number of documents and saves them to the **My Documents** folder. After saving the documents, the user logs off, disconnects from the network, and takes the portable computer on a trip. While on the trip and off the network, the user continues to edit the documents saved earlier in **My Documents**.

**Implementing Making Offline Files Accessible to a Portable Computer Example**
In this situation, the documents are simultaneously saved to a network location and cached to the local computer in a process that is transparent to the user. This action takes place because the network folder is configured to be available offline. This configuration creates a copy of the network folder's contents on the local computer. The user can then access the data offline. By combining Offline Files with Folder Redirection, you can keep user data files backed up and secure on a centrally managed server. If a folder is both redirected and set to offline, that folder receives the benefits of being secure on a server drive, accessible by any computer the user logs on to, and remains available on the user's computer even if the network becomes inaccessible.

**Returning to the Corporate Network**
In this example, the user who uses a portable computer returns to the office and logs on to the network. Because the user made changes to files while working offline, a reminder ballon appears over the notification area showing that the changed files are being synchronized with the network files. In this situation, IntelliMirror technology identifies that the data in **My Documents** has changed and automatically updates the version on the network.

User Data Management and User Settings Management technologies are used to allow users to work on files offline and automatically update network versions of those files when they later reconnect to the network. Offline Files allows users to work on network files when not actually connected to the network. The Synchronization Manager coordinates synchronization of any changes between the offline version of a file and the network version.

Synchronization Manager helps you manage the multiuser use of network files. If multiple users modify the same network file, IntelliMirror notifies the users about the conflict and offers several resolution methods. The user can save the network version, the local version, or both versions. If both versions are to be kept, the user is asked for a new file name to store one of the versions.

# Setting up a Shared Computer Environment Example

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Users work in a department, such as a call center or IT support environment, where they use different computers from day to day. A user is working on a document late one night when the shift ends. The user saves the document and logs off the computer. The next day the user logs on to a different computer and connects to the network. The desktop has the same appearance and configuration as the computer used the previous night. The user resumes working on the same document that was saved in the **My Documents** folder the previous night.

**Implementing Setting Up a Shared Computer Environment Example**

A user's desktop was configured to use Roaming User Profiles so that a copy of the user's working environment was stored on a network server. When the user logged onto the computer, the user's existing preferences, shortcuts, and documents were copied to the local computer. The user was able to continue working as if using the original computer. A variation of this example is when one uses Roaming User Profiles in conjunction with Folder Redirection. This approach saves downloading time because the redirected folders do not need to be copied down to each computer. Users can have the same work environment and access to the same documents from any computer on the network. Changes that the user makes on one computer are synchronized with the other computer the next time the user logs on.

# Replacing a Computer Example

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The computer that a user is working on suddenly has a complete hardware failure. The user calls the technical support, and a new computer with only the Windows XP operating system installed arrives. Without further technical assistance, the user plugs in the new computer, connects it to the network, starts it, and can immediately log on. The user finds that the desktop has the same configuration as the computer it replaced — the same color scheme, screensaver, application icons, shortcuts, and preferences. More importantly, all the user's data files have been restored.

**Implementing the Replacement of a Computer Example**

IntelliMirror helps the user quickly get a computer installed and running with a minimum of support because data and settings are stored independently of any specific computer. By using Roaming User Profiles, Folder Redirection, and Group Policy, you can make the user's data, settings, and applications available wherever the user logs on to the network. In this example, the solution is not limited to getting the user a new computer. Instead, the user could also move to another computer in the office because all of the user's data, settings, and environment are mirrored on the network.

The features of IntelliMirror can be used separately or combined to address the range of needs, from minor configuration changes and updates to complete disaster recovery. This example only addresses IntelliMirror features. In this case, the support department shipped a computer that was preloaded with Windows XP Professional. However, by using RIS, it is also possible to send out computer hardware that has not been preloaded or configured. In that case, RIS can install Windows XP Professional after the computer is on site.

# Implementing a Customized Home Page URL for all Domain Users Example

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You are an administrator and you want to stipulate that the home page in Internet Explorer is configured to use the URL address of your organization's intranet home page for all domain users.

**Implementing the Customized Home Page Example**
In this example, a GPO is created, and policy settings are defined to use a customized home page URL. The GPO is linked at the domain level to ensure it applies to all users in that domain.

# Managing Desktops in Non-Active Directory Environments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The availability of configuration management tools and features differs depending on whether your network operates exclusively in an Active Directory environment or in another network environment.

In a non–Active Directory environment, you can use other tools, such as Microsoft® Systems Management Server (SMS) for managing software distribution, the Internet Explorer Administration Kit for managing Internet Explorer settings, and System Policy (for Windows 95, Windows 98, and Windows NT 4.0) for managing registry-based settings. In addition, each local computer running Windows 2000, Windows XP Professional, Microsoft® Windows® XP 64-Bit Edition, or a Windows Server 2003 operating system has exactly one local Group Policy object (LGPO), even if it does not participate in a domain. Although it is possible to set a variety of settings by using the LGPO, the System Policy scales more easily to a large number of clients. The LGPO can be useful if you only need to apply certain settings to a small number of clients running Windows XP in a Windows NT 4.0 domain. Note that a local GPO does not support Folder Redirection or Group Policy–based Software Installation, and some security settings are not available in local GPOs.

Table 1.2 summarizes how desktop management tools differ between Active Directory and non–Active Directory environments.

**Table 1.2 Desktop Management Tools and Features in Active Directory and Non–Active Directory Environments**

| Management Task | Active Directory | Non–Active Directory |
|---|---|---|
| Configure registry-based settings for computers and users | Administrative templates deployed by using Group Policy<br>Administrative templates, deployed using local Group Policy object (LGPO) | System Policy<br>LGPO |
| Manage local, domain, and network security | Security settings deployed by using Group Policy<br>Security Settings deployed LGPO | LGPO |
| Centrally install, update, and remove software | SMS<br>Group Policy–based software distribution | SMS |
| Manage Internet Explorer configuration settings after deployment | Internet Explorer Maintenance in the Group Policy Object Editor snap-in<br>Internet Explorer Maintenance deployed by using LGPO<br>Internet Explorer Administration Kit (IEAK) | LGPO<br>IEAK |
| Apply scripts during user logon/logoff and computer startup/shutdown | Logon/logoff and startup/shutdown scripts can be centrally configured in Group Policy or independently by using the LGPO | LGPO |
| Centrally manage users' folders and files on the network | Folder Redirection in conjunction with Offline Files and Folders | System Policy<br>Manipulation of registry settings |
| Centrally manage user settings on the network | Roaming User Profiles | Roaming User Profiles (for Windows domains) |

# Developing a Deployment Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To use the IntelliMirror technologies for managing configurations effectively, you must develop a plan that defines the resources to be managed and the management processes to be implemented. To ensure the success of your deployment of the technologies for managing configurations, you need to do the following:

- Analyze your organization's current configuration management process. When you develop your configuration management infrastructure, determine how your organization currently manages its computing environment.

- Determine goals for your managed environment deployment. As you begin your design process, establish the criteria that define the success of your managed environment project. The business requirements of the various organizational groups of users may differ, and success, therefore, means different things to each group. For example, a system administrator might consider a project successful if it results in a reduction in the ten most common support calls, although for a department manager, success includes increased user productivity. It is important that you set achievable and measurable goals, and that you review them periodically to track your progress.

- Design your solution for configuration management by using the IntelliMirror technologies to meet your business requirements.

- Follow a process to implement a configuration management infrastructure. Analyze your business requirements, service level agreements, security, network, and IT requirements, and then determine the appropriate technologies to use to meet your requirements. By following the guidelines in this book for designing your configuration management infrastructure, you can establish the approach that best suits your organization's needs.

- Create a configuration management team. The majority of change and configuration management deployments are likely to have cross-functional boundaries, so as part of preparing your deployment, it is important to consult with various functional teams in your organization and ensure they participate during analysis, design, test, and implementation phases as appropriate.

  A recommended option is to use the Microsoft Solutions Framework (MSF) team model. In the MSF model, small multidisciplinary teams are created, and team members share responsibilities to focus on the project.

  You can align your team with the MSF model by including the six roles in the change and configuration management design team.

  - *Product Management*. Ensures the goals of the project are met.

  - *Program Management*. Facilitates the team and manages resources.

  - *Development*. Designers and implementers of the CCM infrastructure.

  - *Testing*. Ensures all issues are known before the release of the design.

  - *Logistics Management*. Treats physical requirements of the deployment.

  - *User Education*. Designs user and team training materials.

For more information about the Microsoft Solutions Framework Team Model, see the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Defining the Scope of Your Configuration Management Project

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to identify both the organizational and the technical scope for your configuration management deployment. The primary purpose is to deploy a configuration management infrastructure to meet the needs of a defined set of users and to manage a clearly defined computing environment. Defining the scope is important both to set user expectations and concrete goals for the project.

**Organizational Scope**
To understand the configuration management requirements of an organization, you need to identify the organizational structure and determine organizational goals.

You can create a map of your organization to show a representation of its logical hierarchy (an organizational chart). This chart is useful as an input to the design process as it defines natural boundaries within the business, which might function as separation points in the use of IntelliMirror technologies. The breakdown of the organization's resources includes all levels of job roles. Use the job roles as input when you categorize users by job type.

**Resources to be Managed**
The technical scope for deploying a configuration management infrastructure involves all computing resources to be managed, including hardware and software.

It is useful to develop a diagram of the network to serve as a logical representation of servers and services within the network. Work with the network administrators and the Active Directory deployment team to obtain this information. The diagram needs to illustrate the following information:

- Location of domain controllers and member servers
- Site configuration
- Link speed within and across sites
- Servers that provide specific services, such as Domain Naming Services (DNS) and DHCP
- Interfaces with infrastructure elements outside the scope of configuration management, such as firewalls, networking equipment, and systems running other operating systems

Use the network diagram to indicate which components of the computing environment are to be managed by the configuration management infrastructure and which components are to remain unmanaged.

# Assessing Your Current Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assessing your current environment is the first step toward identifying how IntelliMirror can meet your desktop management needs. To identify the areas where IntelliMirror can provide the most benefit for your organization, examine how you perform your desktop management tasks. For example, evaluate how your IT staff handles common desktop management situations such as the following:

- New applications that need to be made available to users as quickly as possible.
- Users who move from one location to another and must retain access to their primary applications and data even on a different computer.
- Users who travel frequently and have intermittent network access, sometimes over slow links.
- New employees who require a newly configured computer that includes all their required applications.
- Users who receive new computers and need to have an operating system installed, along with their applications and data.

Use this process as a starting point to define the tasks your IT group performs for configuration management. You need to determine:

- How your organization currently performs the tasks associated with change and configuration management
- How your organization currently manages desktops

After you determine how your organization manages desktops, you can create a flowchart or table to identify the most common and frequently occurring desktop management tasks that your IT department performs.

# Evaluating Your Desktop Management Processes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Analyzing your current configuration management processes helps you to identify ways you can improve service and reduce your total cost of ownership (TCO).

You need to assess your desktop environment and evaluate the current methods your organization uses for installing operating system, applications and service packs. You also need to evaluate the backup processes in place and determine the level of security to provide for desktops.

For a worksheet to help you evaluate your desktop environment, see "Worksheet A.4 Assessing Your Current Desktop Environment" (DMEUSE_4.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Worksheet A.4 Assessing Your Current Desktop Environment" on the Web at http://www.microsoft.com/reskit).

Large organizations typically support hundreds of software applications as well as multiple versions of operating systems. To help reduce the cost associated with client computing, you can implement standards for corporate-wide software such as virus protection, e-mail, word processing, and spreadsheets — and then retire obsolete or unnecessary software.

To develop your client application standards, examine the operating system types and the versions your organization has installed, the commercial applications your organization uses (such as word processing software), and the line-of-business applications that your corporation has developed for tasks such as managing clients or filling orders. For a worksheet to help you assess your organization's software standards, see "Worksheet A.5 Evaluating Software Standards" (DMEUSE_5.doc) on the Windows Server 2003 Deployment Kit companion CD (or see "Worksheet A.5 Evaluating Software Standards" on the Web at http://www.microsoft.com/reskit).

# Evaluating Computer and User Roles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is recommended that you configure a user's environment according to the user's job or role in the organization. You can then specify the security and other configuration settings for different types of computers such as member servers and desktops. The Group Policy settings you define for each group of users and computers should be based on their business requirements.

**Domain controllers and member servers**
If you have different types of member servers, such as Microsoft® Exchange 2000 Server, terminal servers, or file and print servers, it is likely that you need to define and create different Group Policy objects to configure the settings on each type of server. For example, when managing terminal servers, you need to ensure that you apply user settings Group Policy based on the *computer* that the user logs on to. To achieve this, consider using the Group Policy setting for loopback policy processing. (The **User Group Policyloopback processing mode** policy setting is an advanced option that applies user settings based on the location of the computer object.)

To simplify troubleshooting and problem solving, place the domain controllers and the different types of member servers into separate organizational units and allocate Group Policy objects accordingly.

**Desktops**
Different computers require different policy settings based on their roles. To support these differences, you need to create Group Policy objects that configure each computer according to its role.

To use IntelliMirror to create configurations for computers, you need to consider things such as whether a given computer is configured to allow multiple users to log on. You also need to determine if the computer retains or discards user profiles between user sessions and if the computer provides local storage space for copies of offline files.

# Categorizing Users by Job Types and Location

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Categorizing users in your organization according to their job requirements, computer experience, and location is an important part of the process of developing Group Policy objects for managing the various desktop configurations users require.

Several user types can be used to describe the majority of job roles in most organizations. Every user has specific job-based needs for data management, and users often have multiple roles. For example, a high-performance user might also perform the tasks of a knowledge user.

For each user category, consider what type of configuration options to allow users to make on their computers. Determine if users should run with User or Power User privileges, whether users can install software, or make desktop configuration changes. For example, you might allow high-performance users to change the desktop colors, and you probably want to prevent data-entry workers from making any desktop configuration changes.

**High-performance users**
High-performance users — engineers, graphic artists, and software developers, for example — rely on information technology (IT) to do jobs that are often project- or process-driven. They typically require highly specialized applications.

It is recommended that such users store their data on corporate servers even though they might also need access to local storage on a disk for performance. However, local data is not backed up. Users need permanent access to their data regardless of their location.

**Knowledge workers**
Knowledge users require IT to collect data from many sources, convert the data into information, and communicate the information to support decision-making transactions. Knowledge workers' jobs are typically project-driven, requiring both personal productivity and specialized applications. For example, this category includes executive staff, financial analysts, consultants, researchers, and reporters.

Knowledge workers mainly store data on corporate servers. Users tend to share their working data with others in their department or with those working on the same project.

**Process workers**
Process workers — in customer service, claims and loan processing, for example — use IT to add value to a process. These users perform repetitive process-driven tasks, which require personal productivity and your organization's customized applications.

These users' data is typically stored on corporate servers, and they might have limited write-access to the local computer. In some cases, it might be appropriate to use a fully locked-down desktop.

**Data-entry workers**
Data-entry jobs — processing airline reservations, making order entries, and typing transcriptions — require IT to transcribe data from one medium to another and to add value to data by making it available for other uses.

Data-entry users have no access to disk drives, CD-ROM drives, or the local file system, and their data is stored on corporate servers. Users might share workstations with users on other shifts or with temporary workers.

In addition to categorizing users according to their IT service requirements, you need to identify them by the locations from which they access the corporate network to perform their jobs. Location categories are largely independent of the user types already defined in this section.

**Stationary users**
Stationary users access the corporate network by using local area network (LAN) links. Their computers have permanent LAN connections and high-speed links to connect to the users' data on local servers. The nature and volume of the users' work is fairly predictable. Therefore, their needs for IT services can be easily anticipated.

**Roaming users**
Roaming users access the corporate network by using LAN links. They have permanent LAN connections when working locally, but if they roam between sites, they might have restricted network bandwidth back to some servers. They need to access their data from multiple workstations from many different areas in the same physical location.

**Mobile users**

Mobile users need to access the network at different times and locations by dial-up connections, varying LAN connections, or across a wide area network (WAN) link. Therefore, network services must be accessible at any time.

The following characteristics apply to mobile users:

- Their computers are often connected by slow or intermittent network links.
- The bandwidth, quality, and consistency of their network connections are highly variable.
- Users data and settings need to be saved locally when the users work offline. The users' data and settings might be synchronized to a file server.
- The availability of different types of services depends on whether the users are connected to the corporate network and on the speed and reliability of their connections.

**Remote users**

Remote users are occasionally connected by slow or intermittent network links, but they connect by using the same link each time. They need access to the corporate network at their convenience. Therefore, network services must be accessible at any time. Users need to save data and settings locally, which might not be synchronized to a file server.

Two other important considerations to take into account when analyzing your users and your organization for desktop management are security and corporate policy. You need to consider how you can protect your desktops from possible security threats. Also, you need to consider how your desktop management plans might affect the corporate culture. For more information, see "Determining Corporate Policies" later in this chapter.

# Evaluating Your Security Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The security requirements for managed desktops in your organization are an essential part your configuration management deployment. Examine both the internal and the external security threats that might exist. Internal threats range from accidental damage that your users might cause to their desktops to intentionally malicious actions. External threats include viruses, malicious users, or attackers. Some things to consider as you plan for desktop security include:

- What user security level is appropriate for users to have on their computers?

  Three fundamental levels of security are granted to users by membership in one of these groups: Users, Power Users, and Administrators. Membership in the Users group gives the most protection from a number of external threats, such as viruses, and it limits the damage that a user can accidentally or intentionally cause to their computers. However, user level permissions have the most incompatibility problems with older applications. Take particular care before you give users privileged access to computers that they share with other employees.

- What type of systems do the workstations need to interoperate with?

  They might need to interoperate with Windows NT 4.0 servers, UNIX server message block (SMB) servers, or other types of servers. Interoperability with older systems means that some security you might use in a pure Windows Server 2003 environment must be relaxed.

- Do users have to provide any level of support on their own computers, or do they have to configure their own computers?

  Users who use portable computers and provide their own support might require administrator rights on their computers. Other high-performance users, such as developers, might also need administrative rights.

A series of security templates is supplied with Windows Server 2003 that you can use to manage security configurations. For more information about deploying Windows Server 2003 security, see "Deploying Security Policy" in this book or see "Planning a Secure Environment" in *Designing and Deploying Directory and Security Services* of this kit.

# Determining Corporate Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Desktop management has both political and technical consequences. Depending on the current corporate policy and the culture of your organization, implementing restrictive desktop configurations can either be straightforward or complicated. Often, the more centralized the current IT structure and standards are, the easier it is to gain user acceptance of more restrictive desktop configurations. For example, organizations that have implemented system policies in Windows NT 4.0 to configure common managed desktops might find that users accept quite restrictive policies in Windows 2000 and later. However, organizations moving from an unmanaged Windows NT 4.0 environment (or from a Windows 95 or Windows 98 environment) might encounter resistance to increased restrictions. This can delay or challenge your efforts to impose new or restrictive managed configurations.

This does not mean that an organization with a distributed administration and policy structure cannot benefit from IntelliMirror desktop management. The majority of IntelliMirror features are not readily visible to the user, and many — such as Offline files and application publishing, for example — are enabling rather than restricting technologies. You can actually save time and resources by taking time to carefully examine the benefits of imposing a particular restriction and weigh it against the potential costs of resistance. Where the benefits outweigh the costs, you might want to think about a user education program (as part of the migration training) so that users understand the motivation behind the policy settings.

For more information educating users, see "Preparing Users for Deployment" later in this chapter.

# Assessing Network and Storage Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The availability of network bandwidth can affect how Group Policy settings are applied. It is also important to have sufficient network bandwidth available between servers and workstations when you deploy Roaming User Profiles, Offline Files, and Folder Redirection. Collecting information about your network infrastructure and network traffic patterns helps you plan your deployment of IntelliMirror technologies.

To plan for disk space allocation, you need to determine the disk storage that the users and computers in your organization require.

**Evaluate Your Network Infrastructure**

Evaluate how your network infrastructure meets your user-demand for bandwidth. If some of your users connect to your network over slow or intermittent communication links, the placement of your distribution points and other servers assumes even greater importance in your planning. Determine if any areas of your network are problematic. Where remote users are highly managed, you might justify the cost of upgrading the communication link. By default, when Windows Server 2003 detects a slow link, it does not apply all Group Policy objects to the remote user or computer. For example, security settings and administrative templates are applied, but software installation and scripts are not. The reliability and speed of your local- and wide-area communication links influence your placement of servers and the management functionality you are able to offer to users.

Offline Files, Folder Redirection, and Roaming User Profiles all transfer data files across your network. Software installation sends applications. By testing and piloting, you can best determine the amount of data that passes over your network as a result of implementing these technologies.

The central networking questions you need to answer when deploying IntelliMirror include:

- What changes do you need to make in the physical, logical, and site topology of your network to support these technologies?
- Does your network have the correct protocols to support these technologies?
- Is your network bandwidth within and between sites sufficient to support IntelliMirror features?

You also need to collect the following data:

- The number of clients at each location served. Correlate the placement of servers, routers, and domain controllers to the clients serviced by each device. Remember to include mobile clients.
- Configuration information about network devices such as modems, routers, and hubs. Note whether your routers are configured to pass on broadcast packets. This information is useful when deciding where to place software distribution points.

**Evaluate Your Network Traffic Patterns**

Because IntelliMirror operations involve sending information across your network, knowing the pattern and load of network traffic is essential for planning. Collect the following data about your network traffic patterns:

- Bandwidth demands such as peak utilization and percentage utilization.
- Usage patterns. Establish whether all users log on and off at the same time, or whether user logon time is spread more evenly.

You can estimate the network load generated by software installation by noting:

- Whether installation of the software is mandatory (assigned) or optional (published), and whether the assignment is to a user or to a computer.

- The size of the software package that is installed.
- The placement of the software distribution point in relation to the targeted users.

If you assign software to a computer, the software installation takes place when the computer is restarted. When you perform large deployments to groups of computers, the workload on your network increases substantially when the users start their computers. Consider performing such deployments so that you minimize disruptions to users.

Assigning software to be available on-demand to a user places a shortcut to the software on the user's **Start** menu, but the software is not actually installed until the user starts the application. You can also use the **Install this application at logon** option to install software the next time the user starts the computer, or after the user logs off and then back on. To some degree the load is spread out as users install the software whenever they require it. However, in the case of large-scale deployments to many users, a significant number of users might install the assigned software within the same time period. You can spread out demand on bandwidth over time by publishing software to users so that each user decides when or if to install the package.

You also need to consider the impact of Roaming User Profiles and Folder Redirection on the network loads. If possible, try to minimize the bandwidth impact of synchronizing user data and profiles to network file servers by placing the servers so that data traveling back and forth is localized to the network segment instead of traveling over your entire network or across slow links.

**Evaluate Storage Requirements**

When implementing IntelliMirror technologies, you need to allocate server storage space to hold user data and settings. Allocating server storage space for user data and user settings is complicated by the variety of files (both in number and size) that each user stores on a network server. You need to determine the following:

- Type of user you are supporting
- Typical desktop environment in your organization
- IntelliMirror technologies you are deploying

Recording this information helps you estimate how much server storage space you need to allocate for user data and profiles.

Although you can use Disk Quotas to limit the amount of server disk space available to each user, you can also choose to set a limit but not enforce it. In that case, Windows Server 2003 logs a system event whenever a user uses more than the specified amount of disk space. You can then examine the logs to estimate the amount of disk space that you need for storing user data and settings.

To calculate your server capacity needs when you use Folder Redirection, begin by listing the folders you need to redirect to the network. For example, you can redirect any of the following folders for each user: **My Documents** folder, **Desktop**, **Application Data** folder, and the **Start** menu. Then determine how much local disk space your users' data consumes in the specified folders. You can gather this information by writing scripts (VBScript, WMI queries, or Perl, for example). Or you can use other inventory tools to gather file size data, filtering for the specified directories and known data files extensions such as .doc and .xls. Use this information to estimate the server storage space needed for Folder Redirection.

If possible, minimize the impact on bandwidth of synchronizing user data and profiles to network file servers. To do this, place the servers so that the data traveling back and forth is localized to the network segment, instead of traveling over your entire network or across slow links. Track and control the total number of users saving data files to each server, and track how much data is being stored.

You can control the disk space used to store user data and settings by imposing disk quotas or controlling the size of a user's profile. You can monitor and enforce the size of a user's profile by setting profile quotas. If a user's profile is larger than the size you specify, the user is prevented from logging off until the user reduces the amount of disk space used by that profile. Testing and piloting IntelliMirror features on a sample group of users can provide valuable information about the average and median profile size of your users.

If you use Roaming User Profiles while redirecting the **My Documents** folder of your roaming users, you might not want to impose profile quotas. The reason is that files such as custom dictionaries (size outside user's control) are written to the user profile. Imposing profile quotas can frustrate users who cannot control the size or number of files in their profile.

**Evaluating Service Level Agreements**

Many IT organizations find that the best way to maintain service level agreements with various departments within an organization is to structure services by levels. For example, an organization might define service level agreements (SLAs) that specify the maximum amount of time allowed for the logon screen in Windows 2000 to appear (when the workstation is turned on) or for the user to have access to the desktop after successfully logging on.

You might categorize your services by basic and custom services. *Basic services* are a predetermined set of services you agree to provide to a business unit for a specified cost. You can have several levels of basic services, each with its own assigned cost. Basic services might include standard hardware and software configurations for desktop computers.

*Custom services* are additional specialized services that a business unit might request. When a business unit requests custom services, you have a specified charge-back rate for them. Custom services might include special software configurations created for a user with unique needs.

# Determining Your Requirements for Configuration Management Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To ensure that your design supports the administrative requirements of your organization, begin by determining your organization's current administrative practices and the existence of any administrative boundaries, technical or political, that might impact the design. To determine your specific administrative requirements, you can ask a number of questions about the current configuration management tasks and methods, document the responses, and collect data. This will help you define the goals of your configuration management infrastructure.

Evaluate the administrative tasks you currently perform when installing operating systems, managing user settings and data, computers, and software distribution.

**Operating System Installations**
Assess how your organization manages operating system installations and upgrades, including the following:

- Number of operating system installations that are performed per month
- Frequency of operating system upgrades
- Time to install or upgrade a client computer
- Automation processes used to reduce the time required for installations or upgrades

**User Settings and Data Management**
Managing user settings requires that you assess the levels of control users need to have over their environment and how user profiles are stored. If a policy is currently in place, determine the business requirements it meets.

When you examine the organizational requirements for user data and settings management, determine whether a different strategy can improve user productivity and reduce the need for IT intervention.

Managing user data requires that you consider which configuration management technologies can be implemented to improve the process of managing user data. To do this, document your requirements for the following:

- Backup procedures for user data
- Types of users who regularly use more than one computer and ways that improved data access can enhance their job performance
- Amount of time and money spent recovering and recreating data when hardware or software problems cause a computer failure
- Potential amount of lost revenue to the organization if a catastrophic loss of user data occurs
- Need for users to take server-based data with them off site. Mobile and remote users have different requirements due to the transient nature of their network connections and distance from support staff.

**Computer Settings Management**
Analyze the administrative tasks you currently perform to manage computer settings, such as security and network configuration settings. You need to ensure that you provide appropriate, secure computer configurations based on the business requirements of your organization. You should evaluate issues such as user authentication, access to resources, and network settings.

**Authentication and access to resources** Determine how users are currently authenticated to the network and their computers, what resources users are permitted to access, and to which groups they belong. By using Group Policy, you can define security settings to manage multiple computers. You can create a security policy by importing a security template to a

Group Policy object. Security settings include options for administering the following: account policies, local policies, Event log settings, membership in restricted groups, startup and permissions for system services, and registry and file system access control list (ACLs) permissions.

You can also specify policy settings for IPSec, software restriction policies, public key policies, and wireless network (IEEE 802.11) configurations.

**Administrative rights requirements** Determine whether the users have to configure and support their own computers. For example, users who have to provide their own support (such as users who use portable computers) might require administrator rights on their computers. Other high-performance users, such as developers, might also need to use administrative rights.

**Interoperability in mixed server environments** Determine the type of systems with which the users' workstations need to interoperate, such as Windows NT 4.0 servers, UNIX server message block servers, or other types of servers. To support interoperability with older systems means that some security settings you might use in a pure Windows 2000 environment must be relaxed.

**Network settings** Determine whether you need to set up specific network settings on a per-computer basis — DNS or proxy settings, for example.

**Scripts** Determine the types of scripts to use when a computer starts or shuts down. You can use both VBScript (.vbs) and JScript (.js) scripts.

**Internet Explorer maintenance** To manage Internet Explorer settings for computers, you can use Group Policy for the following: security zones, proxy settings, controlling the installation of Internet Explorer components by users, and displaying or hiding the Internet Explorer splash screen.

**Users privileges on the local computer** Establish whether a user, or set of users, should be a member of the Users, Power Users, or Administrators group.

**Note**

- Typically, for security purposes, it is *not* recommended that users log on to their computers with administrative credentials. If users need to perform administrative tasks on their computer, the users can use the **Run as** command to perform such tasks.

#### Software Distribution

Evaluate the current methods your organization uses for deploying and managing software throughout the software lifecycle. Include the following considerations:

- How applications are deployed, whether the methods used are adequate, and whether the software distribution is too wide in its scope.
- How applications are used, and the common support issues that arise.
- How often applications are upgraded, and how upgrades are performed.
- How applications are removed, and whether full removal is achieved.

For more information about managing configurations by using IntelliMirror technologies, see the following chapters in this book:

- "Designing a Group Policy Infrastructure" for information about deploying Group Policy to manage groups of users and computers.
- "Implementing User State Management" for information about managing user data and user settings.
- "Deploying Security Policy" for information about managing security settings.
- "Deploying a Managed Software Environment" for information about using Group Policy to deploy software.

After you identify your business needs, you can determine which features are most useful for your organization.

# Identifying the IntelliMirror Technologies That Meet Your Configuration Management Needs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following IntelliMirror technologies to improve your configuration management processes:

- User data management
- User settings management
- Computer settings management
- Software management
- Remote operating system installations

**User Data Management**

Ensuring that data remains available is a leading concern for most organizations. What happens to user data when a hard disk fails? Who ensures that users back up their files on a timely basis? Too often, user data backups are not performed, and important files are lost if the user's hard disk fails.

Other concerns include the availability of user's data, whether or not users have access to their data if they move to a different computer on the network, or if they are only intermittently connected to the network. Using the appropriate IntelliMirror desktop management technologies, you can ensure that users can access their data from any computer wherever they log on, whether online or offline. You can back up user data centrally and provide fast computer replacement in disaster recovery situations.

By using IntelliMirror to manage user data, you provide the following advantages:

- You can provide improved protection of user data by ensuring that local data is also redirected or copied to a network share, providing a central location for administrator-managed backups. This capability helps to enforce corporate directives such as to place all important data on servers.
- You can ensure that the most up-to-date versions of a user's data reside on both the local computer and on the server. Local caching maintains data on the local computer even when it is not connected to the network. This ensures that data is available to the user, even when the user is working offline.
- When a user roams to another computer on the network, the user's data can follow. This provides increased accessibility because users can use any computer on the network to access their data.

Implementing user data management relies on some or all of the following technologies:

- Active Directory
- Group Policy
- Roaming User Profiles
- Folder Redirection
- Offline Files
- DFS (Distributed File System)

- EFS (Encrypting File System)
- Disk quotas

For more information about how to implement user data management technologies, see "Implementing User State Management" in this book.

**User Settings Management**

In most organizations, new and existing users who change computers often need help from the IT department initially to configure their computers. By using the IntelliMirror technologies to manage user settings, you can centrally define computing environments for groups of users and computers so that users automatically get the correct configurations for their jobs. You can also restore user settings if a computer fails to ensure that a user's desktop settings follow the user if they roam to another computer.

By managing user state, you can accomplish the following:

- Reduce support calls by providing a preconfigured desktop environment appropriate for the user's job.
- Save time and cost for computer replacement.
- Help users work more efficiently by providing a consistent, secure, preconfigured desktop environment, no matter where they work.

The settings you can manage include:

- Desktop configurations
- Security settings
- Language settings
- Application settings
- Scripts (for when a computer starts or shuts down and when a user logs on or logs off)

Administrators can also redirect any of the special folders in a user profile to a network share. Then the same user profiles are available wherever a user logs on.

User settings, like user data, can follow the user regardless of where that user logs on. Administrators use settings to customize and control users' computing environments and to grant or deny the users the ability to customize their own computing environments. These settings can be applied to both users and computers. When users have permission, they often customize the style and default settings of their computing environment to suit their needs and work habits. Settings contain three basic types of information: user and administrative information, temporary information, and data that is specific to the local computer. For example:

- User settings include items such as Internet Explorer favorites, quick links, cookies, and the Microsoft® Outlook® Express personal Web address book or background bitmap.
- Temporary information includes items such as the user's personal Internet Explorer cache.
- Local computer settings include items such as the folders and files marked for offline use.

You use the following technologies to implement user settings management:

- Active Directory
- Group Policy

- Offline Files
- Synchronization manager
- Folder Redirection
- Roaming User Profiles

For more information about how to implement User Settings Management, see "Implementing User State Management" in this book.

**Computer Settings Management**
Group Policy provides numerous settings to help you manage the computers and servers in your organization. By using Group Policy, you can manage the following types of settings for computers:

**Security settings**
You configure security settings on computers to protect resources on the computers or the network.

**Account policies** Defined on computers, they determine how user accounts interact with the computer or the domain. This includes password policy for domain or local accounts, account lockout policy, and Kerberos policy.

**Local policies** Include audit policy, user rights policy, and security options.

**Public key policies** Includes Encrypting File System policy settings.

**Software restriction policies** Used to protect your computing environment by identifying and specifying which applications are allowed to run on computers.

**Internet Protocol security** Provides protection against private network or Internet attacks. IPSec uses cryptographic security services to ensure private, secure data transmission over Internet Protocol (IP) networks.

**Scripts**
You can specify scripts to run when the computer starts up or shuts down.

**Windows components**
You can specify policy settings for NetMeeting, Internet Explorer, Task Scheduler, Terminal Services, Windows Installer, and Windows Messenger.

**System**
You can specify policy settings for **User Profiles**, **Scripts** (to determine how scripts run), **Logon**, **Disk Quotas**, **Net Logon**, **Group Policy** (to specify how Group Policy should be processed), **Remote Assistance**, **Error Reporting**, **Windows File Protection**, **Remote Procedure Call**, and **Windows Time Service**.

**Network**
You can specify policy settings for **DNSClient**, **Offline Files**, **Network Connections**, **QoSPacket Scheduler**, and **SNMP**.

**Printers**
You can specify policy settings for printers such as allowing printers to be published to or deleted from Active Directory or Web-based printing.

**Software Management**
User productivity is enhanced when users have all of the software applications that enable them to perform their jobs efficiently. It is also important that IT tracks applications that are no longer being used, or are out of date, and makes sure those are phased out. The IT group determines when to stop supporting software that is no longer useful. In some cases, the best solution is to remove the obsolete application instead of risk compatibility issues and other problems that can result from its continued use.

Application management tasks can be extremely labor intensive. That is why many organizations want to automate them for large groups or even for all client computers at one time.

Microsoft provides several software deployment solutions. By evaluating your business objectives for deploying software, the types of requirements your solution must meet, and the types of client operating systems in your organization, you can determine the most appropriate method for deploying software.

**Group Policy–based software installation**

You can use Group Policy–based software installation to install software applications when a computer is started, when the user logs on, or on demand. This approach is suitable for small and medium organizations that have deployed Active Directory, and whose client computers are Windows 2000 Professional or later. You can also use Group Policy–based software installation to upgrade deployed applications or remove earlier applications that are no longer required. You can ensure that a user cannot install any software from local media, such as a CD-ROM, or disk, or other unapproved applications.

The Group Policy–based software installation feature also provides for the following situations:

- If a user inadvertently deletes files from an application, the application repairs itself.
- If a user moves from one computer to another, the software remains available to the user.
- If an application is not installed on a user's computer, and the user tries to open a document associated with that application, the application is automatically installed and the document opens.

You can apply software installation policy settings to users or computers in your Active Directory structure. You can also use Group Policy to set software restrictions to help protect the computer environment from questionable or unknown software.

Group Policy–based software installation does have some basic limitations, including the inability to schedule installation, manage network bandwidth, or provide feedback on the status of the installation. Consider using Systems Management Server if you need to provide scheduled installations, manage network use, perform hardware and software inventory, or monitor installation status.

For more information about deploying software by using Group Policy–based software installation, see "Deploying a Managed Software Environment" in this book.

**Systems Management Server**

SMS is appropriate for medium and large organizations that use client computers running Windows 2000 or earlier. SMS does not require Active Directory. SMS provides advanced capabilities such as inventory-based targeting, status reporting, server- and client-side scheduling, multisite facilities, complex targeting, centralized hardware and software inventory, remote diagnostic tools, software metering, software distribution-point population and maintenance, support for Windows 95, Windows 98, Windows NT 4.0, Windows 2000, and Windows XP clients, and enhanced software deployment features.

For more information about SMS, see the Microsoft Systems Management Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Remote Operating System Installations**

You can use the Remote Installation feature to perform a new installation of Windows on Preboot eXecution Environment (PXE) remote boot-enabled client computers throughout your organization. An administrator does not have to visit the new computer to install a new operating system and core applications. You can provide a customized, fully automated installation process from a remote source. When the computer is turned on, the user presses F12 to initiate the operating system installation process. The computer then starts from a network server that supports RIS.

You can use RIS servers using RIsetup and RIPrep to deploy all editions of Windows 2000 (except Microsoft® Windows® 2000 Datacenter Server), Windows XP Professional, and all editions of Windows Server 2003 (except Windows Server 2003, Datacenter Edition). You can also use RIS servers using RIsetup to deploy Windows XP 64-Bit Edition and the 64-bit versions of Windows Server 2003.

You use the following technologies to implement Remote Installation:

- Active Directory
- Group Policy
- DNS
- DHCP
- RIS

# Designing Managed Configurations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many help desk calls occur after users attempt to upgrade and install nonstandard hardware and software. By creating standards, and designing and deploying managed configurations in your organization, you can specify which applications users can access and which features are available, based on users' job types, services provided by the IT department, and the needs of your corporate environment.

You can limit unnecessary access to system functionality to decrease the number of problems users might inadvertently cause. For example, you can use Group Policy to prevent users from accessing various storage devices such as floppy disk drives, hard disks, or CD-ROMs. By using security policy or ACLs, you can also secure objects, such as system files and the registry, so that your users cannot gain access to them.

Implementing standard configurations can result in increased user productivity by reducing the incidence of computer-related problems. Also, because standard configurations are easier to troubleshoot or replace, they bring about a reduction in support costs.

# Mapping Managed Configurations to User and Computer Roles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you categorize users by their job requirements and locations, and determine the types of servers and clients in your organization, you can determine how best to combine these categories with other organizational considerations to create managed configurations for your organization.

Managed configurations can be mapped to the user categories as described in "Categorizing Users by Job Types and Location" earlier in this chapter. You can design configurations to address the specific business requirements of your organization.

The following examples illustrate how you can use IntelliMirror features to create managed configurations.

**Highly Managed**
Use the highly managed configuration for process workers whose work requires highly restricted configurations with only a few applications. This configuration is suitable for marketing, processing claims and loans, and serving customers.

A highly managed desktop has the following characteristics:

- Users working different shifts can share the computer. Each user needs a unique logon account.
- Users can customize a limited number of application-specific settings.
- Users can access their data from any computer.
- User data is stored on server shares, and users do not store data locally

Table 1.3 lists the features that you can use to design a highly managed configuration.

**Table 1.3 Features Used for a Highly Managed Configuration**

| Feature | Comments |
|---|---|
| Roaming User Profiles | User settings are available at any computer on the network to which the users log on.<br><br>You can remove cached copies of the roaming user profile when the users log off by using Group Policy. |
| Folder Redirection | Redirect **My Documents** to a network share. |
| Offline Files | Designated files are cached locally to ensure that in event of network disconnection, the users can continue to work on these important files. |
| Software installation | Core applications are installed on all laptops (these applications are assigned).<br><br>Optional applications are available for users to install locally (these applications are published). |
| Disk quotas | You can use disk quotas on the servers that contain redirected folders so that individual users cannot use excessive amounts of disk space. |
| Software installation | You can assign core applications to users.<br><br>Optional applications are available for users to install locally (these applications are published). |
| Security context and template | User or Power User.<br><br>The secure workstation security template is used as the basis for this configuration. |

**Lightly Managed**
Use for high-performance users who require a lot of control over their computers, in an organization where highly managed desktops are not acceptable to users, or where desktop management is highly delegated.

The lightly managed desktop has a minimal set of restrictions that help reduce desktop support costs and user down time:

- Users can customize most settings that affect them but are prevented from making unauthorized system changes.
- Users can log on to any computer on the network and access their data.
- User data is saved on server shares and is not stored locally.

Table 1.4 lists the features that you can use for creating a lightly managed configuration.

**Table 1.4 Features Used for Configuring a Lightly Managed Configuration**

| Feature | Comments |
|---|---|
| Roaming User Profiles | Users get the same settings on any computer they use. |
| Redirected Folders and Offline Files | Any files that you redirect must also be set as offline files, so that the users can continue to work on their files even when their computer is not connected to the network. |
| Software Installation | Core applications are installed on all laptops (these applications are assigned).<br><br>Optional applications are available for users to install locally (these applications are published). |
| Disk quotas | Use disk quotas on the servers that contain redirected folders so that individual users cannot exhaust available shared disk space. |
| Security context and template | User or Power User.<br><br>The secure workstation security template is used as the basis for this configuration. |

**Mobile User**

A mobile user configuration is appropriate for managing *mobile users* — traveling users who often use portable computers. Mobile users typically log on to the same computer, and they connect by both high speed and low speed.

The following characteristics apply to mobile-user desktops:

- Can be configured so that users have access to user data whether the computer is connected to or not connected to the network.
- Can save data locally or on network servers.
- Can be configured so that users can disconnect from the network without logging off or shutting down, and to have data files synchronized automatically.

Table 1.5 summarizes the configuration management features you can use to create a mobile user configuration.

**Table 1.5 Features Used for a Mobile User Configuration**

| Feature | Explanation |
|---|---|

| | |
|---|---|
| User profiles | Use the following criteria to determine which type of profile to use:<br>• If the user regularly connects to the network by using a fast link and if you want to back up user state or they use multiple computers, use a roaming user profile.<br>• If the user rarely connects by a fast link, use a local profile.<br>• If the user roams to LAN-connected computers in the domain and also has a portable computer, use a roaming user profile. |
| Folder Redirection | Redirect the **My Documents** folder so that users can access centrally stored data and documents from anywhere. Redirected folders are automatically made available offline, to provide access when users are not connected to the network. |
| Ability for user to customize | Permit customizations within certain guidelines so that users can personalize their work environment but are prevented from making changes to critical system settings. |
| Software installation | Core applications are installed on all laptops. Administrators assign these applications.<br>Optional applications are available for users to install locally (these applications are published). |
| Group Policy settings | Use Group Policy settings to create the managed environment. |
| Security context and template | User or Power User.<br>The highly secure workstation security template is used as the basis for this configuration. |

**Multiuser**

The multiuser desktop is appropriate for environments such as a university laboratory, public computing center, or a library where users might be allowed to save some customizations, such as desktop wallpaper and color scheme preferences, but are prevented from changing hardware or connection settings.

A multiuser configuration has the following characteristics:

- The system is mostlyrestricted, but some personal settings are allowed.
- Users can log on and use a configured roaming profile.
- Users share this computer with other users either by having a unique logon account or by using a Guest account.
- User data is saved on server shares, and users do not store data locally.

Table 1.6 lists the features you can use for a multiuser configuration.

**Table 1.6 Features of a Multi-User Configuration**

| Feature | Comments |
|---|---|
| Multiple users | Users share this computer. Each user has a unique logon account. |
| Roaming User Profiles | Makes user settings available from any computer and enables administrators to easily replace computers without losing their configuration. When the user logs off, the local cached version of the profile is removed to preserve disk space. |
| Folder Redirection | User data is saved on server shares and Group Policy is set to prevent users from storing data locally. Redirect **My Documents**, **Desktop**, and **Application Data**. |

| | |
|---|---|
| Ability for user to customize | Most of the system is locked down, but some personal settings are available. |
| Assigned applications | Core applications common to all users are assigned to the computer. Other applications are available for on-demand install by means of user assignment. |
| Published applications | Applications are available for users to install from **Add or Remove Programs** in **Control Panel**. |
| Group Policy settings | Group Policy settings are used to create the managed environment. |
| Security context and template | User.<br><br>Based on the highly secure workstation security template. |

**Kiosk**

The term *kiosk* in this context refers to a public workstation that runs only one application and one user account, runs unattended, and automatically logs on. Users are unknown to the kiosk owner and do not provide logon credentials. A kiosk workstation is highly secure, simple to operate. Users can not change the default settings.

Use the kiosk desktop in a public area where multiple users access the computer or where you want to prevent users from making any customizations. For example, the kiosk is frequently used in airports where passengers check in and view their flight information.

The following characteristics apply to the kiosk desktop:

- The system is highly restricted by applying policy settings.
- Users cannot customize the installed applications.
- Users cannot save data to the computer locally or to the network.
- The computer can be in a stand-alone environment without any network connectivity.
- Users cannot add or remove applications.
- Users are anonymous, and all users share the same user account.

Table 1.7 lists the features you can use for a kiosk configuration.

**Table 1.7 Features of a Kiosk Configuration**

| Feature | Comments |
|---|---|
| User profile type | Local. |
| Folder Redirection | Not used. |
| Ability for user to customize | No user customizations are permitted. Use policy settings to prevent users from accessing the **Taskbar** and **Start** menu. |
| Assigned applications | One application is assigned to the computer. |
| Security context and template | User.<br><br>The highly secure workstation security template is used as the basis for this configuration. |

**Task Station**

Use the *task station* desktop — an entry terminal for orders on a manufacturing floor or in a call center, for example — for data

entry workers when you need dedicated computers to run a single application.

A task station configuration has the following characteristics:

- The computer is dedicated to running a single application.
- Users on different shifts often share computers.
- Each user has a unique logon account.
- Many users roam between multiple computers that run the same single application.
- User data is saved on server shares and can be stored locally.

Table 1.8 lists the features you can use to configure a task station desktop.

**Table 1.8 Features of a Task Station Configuration**

| Feature | Comments |
|---|---|
| Multiple users | Users share this computer. Each user has a unique logon account. |
| Roaming User Profiles | Makes user settings available from any computer and enables administrators to easily replace computers without losing their configuration. When the user logs off, the local cached version of the profile is removed to preserve disk space. |
| Folder Redirection | Redirect **My Documents** and **Application Data**. |
| Ability for user to customize | No user customizations are permitted. You can also use policy settings to prevent users from accessing the **Taskbar** and **Start** menu. |
| Assigned applications | Typically one application is assigned to the computer. No published applications are made available for users to install. |
| Security context and template | User.<br><br>The highly secure workstation security template is used as a basis for this configuration. |

**Configurations for Computer Replacement**

To simplify computer replacement and to minimize interruption to users, you can store the user data and settings independently of any specific computer. By using Roaming User Profiles and Folder Redirection, you can ensure that the users' data, settings, and applications are available wherever the users log on.

You can also simplify setting up a new managed computer on your network by using RIS to create standardized operating system configurations. You can create a customized image of a Windows XP Professional or Windows 2000 Professional desktop from a source computer. You can save that desktop image to the RIS server, and then use that preconfigured image to set up multiple desktops. You can only include the operating system in the image or a pre-configured desktop image that includes both the operating system and a standard, locally installed desktop application.

# Designing Your Organizational Unit Structure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you plan your configuration management solution, ensure that you design an OU structure that facilitates the management of Group Policy.

The OU hierarchy does not need to mirror your organization's departmental hierarchy. Create every OU to have a defined purpose, such as delegation of authority or application of Group Policy. Business needs must drive the OU hierarchy. By delegating administrative authority, you can designate groups of users to have control over the users and computers or other objects in an OU. An OU is the smallest Active Directory container to which you can assign Group Policy settings.

**Note**

- **Redirusr.exe** (for user accounts) and **redircomp.exe** (for computer accounts) are two new tools included with Windows Server 2003 to assist with the application of Group Policy to new user and computer accounts. These tools are located in %windir%\system32. New user and computer accounts are created in the CN=Users and CN=Computers containers by default. It is not possible to apply Group Policy directly to these containers. By running Redirusr.exe and Redircomp.exe once for each domain, the domain administrator can specify OUs into which all new user and computer accounts are placed at the time of creation. This allows administrators to manage these unassigned accounts by using Group Policy *before* the administrators assign them to the OU in which they are finally placed. It is recommended that the OUs used for new user and computer accounts be highly restricted by means of linked GPOs to increase security around new accounts.

For more information about redirecting the Users and Computers containers, see article Q324949, "Redirecting the Users and Computers Containers in Windows Server 2003 Domains," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about the **redirusr.exe** and **redircomp.exe** tools, see the Redirecting Users and Computers link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Application of Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can link GPOs to sites, domains, and OUs to implement Group Policy settings as broadly or as narrowly in the organization as necessary. Keep in mind how Group Policy is applied when you determine the scope of application of Group Policy objects:

- The policy settings in Group Policy objects are inherited, cumulative, and apply to all users and computers in an Active Directory container.
- Group Policy objects are processed in the following order: local GPO, site, domain, and OU.
- By default, Group Policy inheritance is evaluated starting with the Active Directory container farthest from the computer or user object. The Active Directory container closest to the computer or user overrides Group Policy set in a higher-level Active Directory container unless you set the **No Override** option for that GPO.
- If you link more than one GPO to an Active Directory container, the GPO processing order (priority) is as follows: the GPO highest in the **Group Policy Object Links** list, displayed in the Group Policy page of the Active Directory container's **Properties** page, has precedence by default. If you set the **No Override** option in one or more of the GPOs, the highest GPO that is set to **No Override** takes precedence.

For information about creating an Active Directory structure see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

For more information about defining the scope of application of Group Policy, see "Designing a Group Policy Infrastructure" in this book.

# Testing Your Configuration Management Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before deploying your management solutions to a wide base, fully test your design in a test lab environment. Minimally, your test environment consists of at least two domain controllers, a member server, two or more workstations, and possibly a mobile computer connected by means of a slow link. If you are testing software installation through Group Policy, include one or more servers set up as software distribution points. By setting up a test-to-production environment deployment process and using features of the Group Policy Management Console, you can ensure that you provide a reliable and consistent configuration management solution.

Document the testing network as well as all steps required to set it up. If new hardware, such as a new server, is being added to your organization's network, use this same hardware in your test deployment if possible.

To minimize variables and to ensure that testing does not interfere with your organization's network services, keep the testing network on its own, isolated LAN.

**Note**

- If there is a network server that you cannot simulate, such as a database service, you might need to temporarily connect the testing network to the organization's network to verify that the deployment interacts with this service as expected. Use this approach *only* if you have no other options available for verifying the interaction of the deployment with the service.

After completing tests in a controlled environment, select a group of users to pilot your configuration. Keep the users to a manageable number. A pilot can expose unexpected problems on a small scale so that you can resolve them before deploying on a large scale. Verify that the deployed technology is operating as expected. If you perform an iterated deployment, deploy and test it in phases, and then emphasize the testing of the final configuration.

You can perform two types of testing:

- Proof of concept testing
- Pilot testing

When conducting such tests, be prepared to iterate if necessary. Depending on the scope of any design updates, you might want to repeat the proof of concept or pilot stage before proceeding further.

# Conducting Proof of Concept Testing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Proof of concept testing provides the opportunity to try out a design in a controlled environment and to identify any potential major problems or challenges prior to conducting a pilot with users.

**Use a Test Environment** Conducting a proof of concept test requires a test environment that can be configured to simulate the intended configuration for your production environment. Typically, this is a small-scale simulation that is also large enough to demonstrate the core structures and policy settings that you have designed.

Primary aspects of conducting a proof of concept test are to give a representative sample of users and administrators access to the test environment and to confirm that the design meets their requirements and expectations.

**Involve the Design Team** Ensure that all members of the design team have the opportunity to review the Proof of Concept and are confident that it meets the design objectives.

# Conducting Pilot Testing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can conduct a pilot test to try out a design with a controlled number of users. This enables you to capture refinements that have not been identified during the design phase, and to highlight any issues that have not previously been seen during the proof of concept phase. In some circumstances, you can use pilot testing to check system performance and resilience under real-user conditions. Carrying out a pilot test involves tasks, starting with selecting the appropriate users.

**Selecting the appropriate users** Typically, it is good practice to select users who have volunteered to participate in the pilot, instead of imposing the test on them. Volunteer users are more likely to be fully involved in the pilot, and take the time to test the configurations.

**Providing pre-pilot training** In any pilot, it is essential to provide adequate user training before deploying the system, and to set expectations about what is required from the users during the test phase. For example, you need to tell users who to contact if they encounter problems and how to provide feedback.

**Deploying the pilot** When deploying the pilot, it is good practice to involve any implementation team that is required to implement the final systems. The team can then identify issues in advance and plan how to solve them.

**Reviewing and refining** After conducting the pilot, perform a review of the process to ascertain whether the systems met the design requirements. This helps you identify any improvements you need to consider for inclusion in the final design.

For more information about how to stage, test, and deploy your configuration management deployments, see "Staging Group Policy Deployments" in this book.

# Preparing Users for Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Desktop management changes the way many users interact with their computers. It might also change how they access their data, where their settings are stored, how they install software, and which software they can access. As you begin to implement new desktop management standards, it is important to develop a plan to educate and support your users.

If you create standard desktops, you need to let users know well in advance what changes are planned and how you plan to implement those changes. In addition, you need to train users, as well as IT and support staff, to use the new features. Resources that users can refer to for information when questions arise, such as a corporate Web site or newsletter, can help to ease the transition.

You also need to think about how users might perceive your new management processes. For example, a common phrase many IT professionals use to describe a highly managed desktop is "locked down." Although system administrators and IT professionals understand this phrase, many users might not like the idea of the IT department locking down their desktops. As you train your users, emphasize how a managed environment provides value. For example, users are likely to appreciate the convenience of automatic software installation and repair.

You also need to determine how you will provide user support during the transition. As you begin to implement new features, determine what percentage of your staff will support upgraded users and the escalation procedures the support staff will use. Make sure to document the problems that users encounter as you implement different parts of your plan. That information helps you support the next set of users.

# Staging and Deploying Your Design to the Production Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The importance of staging your deployments of Group Policy can not be overemphasized. After you have performed incremental policy changes in the test environment and verified your changes, you can begin to migrate the Group Policy objects you created to your production domain by using the migrate functionality and the migration tables included in GPMC.

As you begin to implement desktop management services in your organization, look for the desktop management technology that provides the greatest benefit to your organization. Implement that service before you begin to implement others. This gives you the opportunity to test one set of technologies relative to the technologies already in place before you implement another set.

For example, you might decide that protecting and backing up user data would provide the greatest benefit to your organization. You might then choose to implement Folder Redirection in your company so that all user data is centrally stored on a server, facilitating backup. After you deploy Folder Redirection to your users, you can reevaluate whether additional IntelliMirror technologies such as Roaming User Profiles or Software Installation and Maintenance is appropriate to implement next.

By following the processes described in "Staging Group Policy Deployments" in this book, you can ensure that you provide a predictable, consistent configuration management solution.

# Additional Resources for Planning a Managed Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Deploying a Simple Managed Environment" in this book for an example of how IntelliMirror technologies are used to create a simple managed environment for your users.
- "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about Group Policy.
- "Designing RIS Installations" in *Automating and Customizing Installations* of this kit.
- "Designing a Group Policy Infrastructure" in this book.
- "Staging Group Policy Deployments" in this book.
- "Deploying Security Policy" in this book.
- "Migration of User State" in this book.
- "Implementing User State Management" in this book.
- "Deploying a Managed Software Environment" in this book.
- "Deploying a Simple Managed Environment" in this book.

**Related Job Aids**

- "Worksheet A.1 Identifying Primary Client Support Tasks" (DMEUSE_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.1 Identifying Primary Client Support Tasks" on the Web at http://www.microsoft.com/reskit).
- "Worksheet A.2 Assessing User Data Management Tasks" (DMEUSE_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.2 Assessing User Data Management Tasks" on the Web at http://www.microsoft.com/reskit).
- "Worksheet A.3 Assessing User Settings Management Tasks" (DMEUSE_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.3 Assessing User Settings Management Tasks" on the Web at http://www.microsoft.com/reskit).
- "Worksheet A.4 Assessing Your Current Desktop Environment" (DMEUSE_4.doc) on the Windows Server 2003 Deployment Kit companion CD (or see "Worksheet A.4 Assessing Your Current Desktop Environment" on the Web at http://www.microsoft.com/reskit).
- "Worksheet A.5 Evaluating Software Standards" (DMEUSE_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.5 Evaluating Software Standards" on the Web at http://www.microsoft.com/reskit).
- "Worksheet A.6 Analyzing Your Users by Job Function" (DMEUSE_6.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.6 Analyzing Your Users by Job Function" on the Web at

http://www.microsoft.com/reskit).

- "Worksheet A.7 Classifying Your Users" (DMEUSE_7.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.7 Classifying Your Users" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.8 Managing Different Types of Users" (DMEUSE_8.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.8 Managing Different Types of Users" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.9 Evaluating your Security Requirements" (DMEUSE_9.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.9 Evaluating Your Security Requirements" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.10 Considering Common Configurations" (DMEUSE_10.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.10 Considering Common Configurations" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.11 Training Needs for Various Workers" (DMEUSE_11.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.11 Training Needs for Various Workers" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.12 Updating Network Documentation" (DMEUSE_12.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.12 Updating Network Documentation" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.13 Documenting Servers (Example)" (DMEUSE_13.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.13 Documenting Servers (Example" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.14 Documenting Client Computers" (DMEUSE_14.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.14 Documenting Client Computers" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.15 Documenting Clients (Example)" (DMEUSE_15.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.15 Documenting Clients (Example)" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.16 Evaluating Network Infrastructure" (DMEUSE_16.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.16 Evaluating Network Infrastructure" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.17 Documenting Network Infrastructure" (DMEUSE_17.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.17 Documenting Network Infrastructure" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.18 Evaluating Network Traffic Patterns" (DMEUSE_18.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.18 Evaluating Network Traffic Patterns" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.19 Monitoring Network Performance with System Monitor" (DMEUSE_19.doc) on the Windows Server 2003 Deployment Kit companion CD (or see "Worksheet A.19 Monitoring Network Performance with System Monitor" on the Web at http://www.microsoft.com/reskit).

# Designing a Group Policy Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft® Windows® Server 2003 Group Policy enables administrators to manage configurations for groups of computers and users, including options for registry-based policy settings, security settings, software deployment, scripts, folder redirection, Remote Installation Services, and Microsoft® Internet Explorer maintenance. By using Group Policy, you can significantly reduce an organization's total cost of ownership. Because of factors such as the large number of policy settings available, the interaction between multiple policies, and inheritance options, Group Policy design can be complex. By carefully planning, designing, and testing a solution based on your organization's business requirements, you can provide the standardized functionality, security, and management control that your organization needs.

**In This Chapter**

Overview of Group Policy

Planning Your Group Policy Design

Designing Your Group Policy Model

Deploying Group Policy

Maintaining Group Policy

Additional Resources for Group Policy Infrastructure

**Related Information**

- For more information about the Active Directory® directory service, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

- For more information about security in Windows Server 2003, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

- For more information about managing Group Policy, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Overview of Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy enables Active Directory–based change and configuration management of user and computer settings on computers running a member of the Microsoft® Windows® Server 2003 or Microsoft Windows® 2000 families of operating systems, or the Microsoft® Windows® XP Professional operating system. You use Group Policy to define configurations for groups of users and computers, including policy settings for registry-based policies, software installation, scripts, folder redirection, Remote Installation Services, Internet Explorer maintenance, and security. You can also use Group Policy to help manage server computers, by configuring many server-specific operational and security settings.

The Group Policy settings that you create are contained in a Group Policy object (GPO). To create a GPO, use the Group Policy Management Console (GPMC). To edit a new GPO, use the Group Policy Object Editor snap-in for the Microsoft Management Console (MMC), which you can start from GPMC. By using GPMC to link a GPO to selected Active Directory system containers — sites, domains, and organizational units (OUs) — you apply the policy settings in the GPO to the users and computers in those Active Directory containers.

To guide your Group Policy design decisions, you need a clear understanding of your organization's business needs, service level agreements, and security, network, and IT requirements. By analyzing your current environment and users' requirements, defining the business objectives you want to meet by using Group Policy, and following this chapter's guidelines for designing a Group Policy infrastructure, you can establish the approach that best supports your organization's needs.

To see example standard desktop configurations and the actual policy settings used for those configurations, see the Group Policy scenarios in the whitepaper at the Implementing Common Desktop Management Scenarios link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. For a list of these sample configurations, see Table 2.3 later in this chapter.

IntelliMirror refers to the ability to provide users with consistent access to their applications, application settings, roaming user profiles, and user data, from any managed computer – even when they are disconnected from the network. IntelliMirror is delivered by means of a set of Windows features that enable IT administrators to implement standard computing environments for groups of users and computers.

IntelliMirror can significantly boost user productivity and satisfaction by doing the following:

- Allowing users to continue working efficiently in intermittently connected or disconnected scenarios by enabling uninterrupted access to user and configuration data under these conditions.
- Delivering a consistent computing environment to users from any computer when their desktop or laptop computer is unavailable or in scenarios where users are not assigned a specific computer.
- Minimizing data loss by enabling centralized backup of user data and configuration files by the IT organization.
- Minimizing user downtime by enabling automated installation and repair of applications.

Implementing IntelliMirror also boosts administrator efficiency and reduces IT costs by doing the following:

- Eliminating the need to manually configure user settings, install applications, or transfer user files to provide users access to their computing environments on any computer.
- Enabling scenarios where users do not have an assigned computer but log in to any available computer in a pool of computers. This helps reduce hardware and administration costs.
- Easing the IT task of implementing centralized backup of user files while satisfying the need for these files to be available on the user's computer.
- Reducing support costs by using Windows Installer to automatically repair broken application installations.
- Enabling rapid deployment of security settings to ensure resources on the network are secure.

Windows features that implement IntelliMirror include Active Directory, Group Policy, Software Installation, Windows Installer, Folder Redirection, Offline Folders, and Roaming User Profiles.

# Process for Implementing a Group Policy Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process for implementing a Group Policy solution entails planning, designing, deploying, and maintaining the solution. These steps are illustrated in Figure 2.1.

**Figure 2.1 Implementing Group Policy**




When you plan your Group Policy design, make sure that you design the OU structure to ease Group Policy manageability and to comply with service level agreements. Establish good operational procedures for working with GPOs, such as those defined in this chapter, make sure that you understand Group Policy interoperability issues, and determine whether or not you plan to use Group Policy for software deployment.

During the design phase, define the scope of application of Group Policy, determine the policy settings that are applicable to all corporate users, classify users and computers based on their roles and locations, and then plan desktop configurations based on the user and computer requirements.

The deployment phase begins by staging in a test environment. This includes creating standard desktop configurations, filtering the scope of application of GPOs, specifying exceptions to default inheritance of Group Policy, delegating administration of Group Policy, evaluating effective policy settings by using Group Policy Modeling (formerly known as Resultant Set of Policy planning mode), and evaluating the results by using Group Policy Results (formerly known as Resultant Set of Policy logging mode).

Staging is critical. Thoroughly test your Group Policy implementation in a test environment before deploying to your production environment. After you complete staging and testing, migrate your GPO to your production environment by using GPMC. Consider an iterative implementation of Group Policy: rather than deploying 100 new Group Policy settings, initially stage and then deploy only a few settings to validate that the Group Policy infrastructure is working well.

Finally, prepare to maintain Group Policy by establishing control procedures for working with and troubleshooting GPOs by using GPMC.

# What You Need Before Designing Your Group Policy Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before designing your Group Policy implementation, you need to understand your current organizational environment, and you need to take preparatory steps in the following areas:

**Active Directory** Ensure that the Active Directory OU design for all domains in the forest supports the application of Group Policy. For more information about Active Directory OU design, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

**Group Policy Management Console (GPMC)** Download and install the Group Policy Management Console, which consists of scripting interfaces and a Microsoft Management Console snap-in.

**Networking** Make sure that your network meets the requirements for change and configuration management technologies. Because Group Policy works with fully qualified domain names, you must have DNS running in your forest in order to correctly process Group Policy; you cannot use NETBIOS only. Also, because client or destination computers must be able to contact your network's domain controllers, do not turn off the ICMP protocol. If destination computers cannot ping the domain controllers, Group Policy processing will fail.

**Security** Obtain a list of the security groups currently in use in your domain. Work closely with the security administrators as you delegate responsibility for organizational-unit administration and create designs that require security-group filtering. For more information about filtering GPOs, see "Applying GPOs to Selected Groups (Filtering)" in "Defining the Scope of Application of Group Policy" later in this chapter.

**IT requirements** Obtain a list of the administrative owners and corporate administrative standards for the domains and OUs in your domain to develop a good delegation plan and to ensure that Group Policy is properly inherited.

**Note**

- Turning off the ICMP protocol will cause Group Policy processing to fail. Turning off Read access on Active Directory containers that are in the hierarchy of a user or computer object will cause Group Policy processing for that object to fail.

#### Administrative Requirements for Group Policy

To use Group Policy, your organization must be using Active Directory and the destination desktop and server computers must be running Windows 2000 Professional, Windows 2000 Server, Windows XP Professional, or Windows Server 2003. You can manage server computers as well as client computers by using Group Policy; Group Policy offers many settings specific to server computers.

Using GPMC will greatly improve the manageability of your Group Policy deployment and enable you to take full advantage of the power of Group Policy by providing an enhanced and simplified Group Policy management interface.

By default, only domain administrators or enterprise administrators can create and link GPOs, but you can delegate this task to other users. For more information about administrative requirements for Group Policy, see "Delegating Administration of Group Policy" later in this chapter.

# New in Windows Server 2003: Group Policy Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

GPMC is a new tool that unifies management of all aspects of Group Policy across multiple forests in an enterprise. GPMC allows you to manage all GPOs, Windows Management Instrumentation (WMI) filters, and Group Policy-related permissions in your network. Think of GPMC as your primary access point to Group Policy, with all the Group Policy management tools available from the GPMC interface. The information presented in this book is based on using GPMC for Group Policy deployment and ongoing management.

GPMC consists of a set of scriptable interfaces for managing Group Policy and an MMC-based user interface (UI). The UI integrates all previous Group Policy tools into a unified Group Policy-management console. GPMC runs on 32-bit computers that are running a member of the Windows Server 2003 family operating system or Windows XP Professional with Service Pack 1 and the Microsoft® .NET Framework. This tool can manage both Windows Server 2003 and Windows 2000 Active Directory–based domains.

GPMC provides the following:

- A new user interface that integrates existing Group Policy functionality currently accessible by using various tools such as the Active Directory Users and Computers snap-in, the Active Directory Sites and Services snap-in, the Delegation of Control Wizard, the RSoP snap-in, the Delegation Wizard, and the ACL editor. The UI also simplifies inheritance and enforcement of GPOs.
- Access to the Group Policy Object Editor (previously known as the Group Policy MMC snap-in).
- Importing and exporting GPOs.
- Copying and pasting GPOs.
- Backing up and restoring GPOs.
- Searching for existing GPOs.
- Integration of RSoP capabilities:
    - Group Policy Modeling. Allows you to simulate RSoP data for planning Group Policy deployments prior to implementing them in the production environment.
    - Group Policy Results. Allows you to get RSoP data for viewing GPO interaction and for troubleshooting Group Policy deployments.
- Support for migration tables to facilitate the importing and copying of GPOs across domains and across forests.
- Reporting GPO settings and RSoP data in HTML reports that you can save and print.
- Scripting all operations that are available within the tool. You cannot, however, use scripts to edit individual policy settings in a GPO.

**Note**

- To help you get started, the GPMC installation includes sample scripts that use COM interfaces. The sample scripts are installed in the folder \Program Files\GPMC\Scripts\.

GPMC deployment and troubleshooting operations are described throughout this chapter. For detailed, step-by-step information about using GPMC to deploy and manage your Group Policy infrastructure, see Help in GPMC. Full details of the scripting interfaces are documented in the Group Policy Management Console Software Development Kit (SDK), which is

located at program files\gpmc\scripts\gpmc.chm on any computer where you install GPMC. The GPMC SDK is also available in the Microsoft® Platform SDK. For more information and to download SDKs, see the Microsoft Platform SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Before you begin planning your Group Policy design, install GPMC. It is available as a download from the Microsoft Web site. See the Group Policy Management Console link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Planning Your Group Policy Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you plan your Group Policy design, ensure that your Active Directory design supports the application of Group Policy. Then you need to clearly define your objectives for deploying Group Policy. Specifically, understand any service-level agreements and administrative issues that pertain to Group Policy and consider your business requirements and how Group Policy can help you achieve them. Finally, incorporate any operational, interoperability and software installation considerations into your plan. Figure 2.2 illustrates the steps in the Group Policy planning process.

**Figure 2.2 Group Policy Planning**



# Designing an OU Structure that Supports Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In an Active Directory environment, you assign Group Policy settings by linking GPOs to sites, domains, or organizational units (OUs). Typically, most GPOs are assigned at the organizational unit level, so be sure your OU structure supports your Group Policy-based client-management strategy. You might also apply some Group Policy settings at the domain level, particularly those such as password policies, which only take effect if applied at the domain level. Very few policy settings are likely to be applied at the site level. A well-designed OU structure, reflecting the administrative structure of your organization and taking advantage of GPO inheritance, simplifies the application of Group Policy. For example, it can prevent needing to duplicate certain policies so that the policies can be applied to different parts of the organization, or having to link the same GPO to multiple Active Directory containers to achieve your objectives. If possible, create OUs to delegate administrative authority as well as to help implement Group Policy.

OU design requires balancing requirements for delegating administrative rights – independent of Group Policy needs – and the need to scope the application of Group Policy. The following OU design recommendations address delegation and scope issues:

**Delegating administrative authority** You can create OUs within a domain and delegate administrative control for specific OUs to particular users or groups. Your OU structure might be affected by requirements to delegate administrative authority. For more information about planning for delegation of Active Directory administrative authority, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

**Applying Group Policy** An OU is the lowest-level Active Directory container to which you can assign Group Policy settings.

Think primarily about the objects you want to manage when you approach the design of an OU structure. You might want to create a structure that has OUs organized by workstations, servers, and users near the top level. Depending on your administrative model, you might consider geographically based OUs either as children or parents of the other OUs, and then duplicate the structure for each location to avoid replicating across different sites. Add OUs below these only if doing so makes the application of Group Policy clearer, or if you need to delegate administration below these levels.

By using a structure in which OUs contain homogeneous objects, such as either user or computer objects but not both, you can easily disable those sections of a GPO that do not apply to a particular type of object. This approach to OU design, illustrated in Figure 2.3, reduces complexity and improves the speed at which Group Policy is applied. Keep in mind that GPOs linked to the higher layers of the OU structure are inherited by default, which reduces the need to duplicate GPOs or to link a GPO to multiple containers.

Note that the default Users and Computers containers cannot have Group Policy applied to them until you use the new Redirusr.exe and Redircomp.exe tools. When designing your Active Directory structure, the most important considerations are ease of administration and delegation.

**Figure 2.3 Example OU Structure**




**Applying Group Policy to New User and Computer Accounts**

New user and computer accounts are created in the CN=Users and CN=Computers containers by default. It is not possible to apply Group Policy directly to these containers, although they inherit GPOs linked to the domain. Redirusr.exe (for user accounts) and Redircomp.exe (for computer accounts) are two new tools included with Windows Server 2003 that enable you to change the default location where new user and computer accounts are created so you can more easily scope GPOs directly

to newly created user and computer objects. These tools are located in %windir%\system32. By running Redirusr.exe and Redircomp.exe once for each domain, the domain administrator can specify the OUs into which all new user and computer accounts are placed at the time of creation. This allows administrators to manage these unassigned accounts by using Group Policy before the administrators assign them to the OU in which they are finally placed. You might want to consider restricting the OUs used for new user and computer accounts by using Group Policy to increase security around these accounts.

For more information about redirecting users and computers, see article 324949, "Redirecting the Users and Computers Containers in Windows Server 2003 Domains," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about the redirusr.exe and redircomp.exe tools, see the Redirecting Users and Computers link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Sites and Replication Considerations**

As you determine which policy settings are appropriate, be aware of the physical aspects of Active Directory, which include the geographical location of sites, the physical placement of domain controllers, and the speed of replication.

GPOs are stored in both Active Directory and in the Sysvol folder on each domain controller. These locations have different replication mechanisms. Use the Resource Kit tool Gpotool.exe to help diagnose problems when you suspect that a GPO might not have replicated across domain controllers. For more information about Gptool.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Windows Resource Kit Tools Help.

Domain controller placement is an issue when slow links, typically to clients at remote sites, are involved. If the network link speeds between a client and the authenticating domain controller fall below the default slow-link threshold of 500 kilobits per second, only the administrative template (registry-based) settings, the new Wireless Policy extension, and security settings are applied by default. All other Group Policy settings, including software distribution and folder redirection, are not applied by default. You can, however, modify this behavior by using Group Policy.

You can change the slow link threshold by using the **Group Policy Slow Link Detection** policy for both the user and computer aspects of a GPO. If necessary, you can also adjust which Group Policy extensions are processed below the slow-link threshold. However, it might be more appropriate to place a local domain controller at a remote location to serve your management needs.

# Complying with Service Level Agreements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some IT groups use service level agreements to specify how services should operate. For example, a service-level agreement might stipulate the maximum length of time required for computer startup and logon, how long users can use the computer after they log on, and so on. Service-level agreements often set standards for service responsiveness. For example, a service level agreement might define the amount of time allowed for a user to receive a new software application or gain access to a previously disabled feature. Issues that can affect service responsiveness are the site and replication topology, the positioning of domain controllers, and the location of Group Policy administrators.

To reduce the amount of time required to process a GPO, consider using one of the following options:

- If a GPO contains only computer or user settings, disable the portion of the policy that does not apply. The destination computer does not scan the portions of a GPO that you disable, which reduces processing time.

- When possible, combine smaller GPOs to form a consolidated GPO. This reduces the number of GPOs that apply to a user or computer. Applying fewer GPOs to a user or computer can reduce startup or logon times and make it easier to troubleshoot the policy structure.

- The changes you make to GPOs are replicated to domain controllers and result in new downloads to client or destination computers. If you have large or complex GPOs that require frequent changes, consider creating a new GPO that contains only the sections that you update regularly. Test this approach to determine whether the benefits you get by minimizing the impact on the network and improving the destination computer's processing time outweigh the increased troubleshooting potential by making the GPO structure more complex.

**Note**

- You should implement a Group Policy change control process and log any changes made to GPOs. This can help you troubleshoot and correct problems with Group Policy objects. Doing so also helps comply with service level agreements that require keeping logs.

# Defining Your Group Policy Objectives

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you plan the deployment of Group Policy, identify your specific business requirements and how Group Policy can help achieve them. You can then determine the most appropriate policy settings and configuration options to meet your requirements.

The objectives for each Group Policy implementation vary depending on user location, job needs, computer experience, and corporate security requirements. For example, in some cases, you might remove functionality from users' computers to prevent them from modifying system configuration files (which might disrupt computer performance), or you might remove applications that are not essential for users to perform their jobs. In other cases, you might use Group Policy to configure operating system options, specify Internet Explorer maintenance settings, or establish a security policy.

Having a clear understanding of your current organizational environment and requirements helps you design a plan that best meets your organization's requirements. Collecting information about the types of users (such as process workers and data entry workers) and existing and planned computer configurations is essential. Based on this information, you can define your Group Policy objectives.

**Evaluating Existing Corporate Practices**

To help you identify the appropriate Group Policy settings to use, begin by evaluating current practices in your corporate environment, including such things as:

- User requirements for various types of users.
- Current IT roles, such as the various administrative duties divided amongst administrator groups.
- Existing corporate security policies.
- Other security requirements for your server and client computers.
- Software distribution model.
- Network configuration.
- Data storage locations and procedures.
- Current management of users and computers.

**Defining Group Policy Objectives**

Next, as part of defining the goals for Group Policy, determine the following:

- Purpose of each GPO
- Owner of each GPO – the person who requested the policy and who is responsible for it
- Number of GPOs to use
- Appropriate container to link each GPO (site, domain, or OU)
- Types of policy settings contained in each GPO, and the appropriate policy settings for users and computers
- When to set exceptions to the default processing order for Group Policy
- When to set filtering options for Group Policy

- The software applications to install and their locations
- What network shares to use for redirecting folders
- The location of logon, logoff, startup, and shutdown scripts to execute

# Establishing Group Policy Operational Guidelines

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you design and implement your Group Policy solution, it is also important to plan for the ongoing administration of Group Policy.

Establishing administrative procedures to track and manage GPOs can ensure that all changes are implemented in a prescribed manner.

To simplify and regulate ongoing management of Group Policy, it is recommended that administrators:

- Always stage Group Policy deployments using the following pre-deployment process:

    1. Use Group Policy Modeling to understand how a new GPO will interoperate with existing GPOs.
    2. Deploy new GPOs in a test environment modeled after your production environment.
    3. Use Group Policy Results to understand which GPO settings actually are applied in your test environment.

- Use GPMC to make backups of your GPOs on a regular basis.
- Use GPMC to manage Group Policy across the organization.
- Do not modify the default domain policy or default domain controller policy unless necessary. Instead, create a new GPO at the domain level and set it to override the default settings in the default policies.
- Define a meaningful naming convention for GPOs that clearly identifies the purpose of each GPO.
- Designate only one administrator per GPO. This prevents one administrator's work from being overwritten by another's.

Windows Server 2003 and GPMC allow you to delegate permission to edit and link GPOs to different groups of administrators. Without adequate GPO control procedures in place, delegated administrators can duplicate GPO settings, or create GPOs that conflict with settings set by another administrator or that are not in accordance with corporate standards. Such conflicts might adversely affect the users' desktop environment, generate increased support calls, and make troubleshooting GPOs more difficult.

# Identifying Interoperability Issues

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to consider possible interoperability issues when planning a Group Policy implementation in a mixed environment. Group Policy only applies to computers running Windows 2000, Windows XP Professional, or Windows Server 2003. You can only use GPMC on a computer running Windows XP Professional or Windows Server 2003 in an Active Directory network.

Windows Server 2003 and Windows XP Professional include many new Group Policy settings that are not used on Windows 2000. However, even if the client and server computers in your organization mostly run Windows 2000, and you have any Windows Server 2003 – based computers, you should use the Windows Server 2003 administrative templates (.adm files) because they are the latest .adm files and so are the most inclusive. If you apply a GPO with newer settings to a previous operating system that does not support the setting, it will not cause a problem. Destination computers that are running Windows 2000 or Windows XP Professional will simply ignore settings supported only in Windows Server 2003. To determine which settings apply to which operating systems, in the description for the setting, see the **Supported on** information, which explains which operating systems can read the setting.

If you are planning to deploy Group Policy in mixed environments, take the following Group Policy processing behavior into account.

- If the destination computer is running Microsoft® Windows NT® version 4.0, Microsoft Windows® 95, or Microsoft Windows® 98, it uses System Policy rather than Group Policy. System Policy is a Windows NT 4.0-style policy based on registry settings specified by using the System Policy Editor, Poledit.exe.

- If the destination computer is running Windows 2000, Windows XP Professional, or Windows Server 2003, and the computer account and the account for the logged-on user are both located in a Windows 2000 or Windows Server 2003 domain, both the computer and the user portions of a GPO are processed. If either the logged-on user account or the computer account is located in a Windows NT 4.0 domain, System Policy is processed for those accounts located in the domain.

Identifying and troubleshooting problems in a mixed environment where both System Policy and Group Policy apply can be difficult and time-consuming. When possible, move both the computer and the user account into a Windows 2000 or Windows Server 2003 domain. For more information about migrating user and computer accounts from a Windows NT 4.0 domain to Windows 2000 or Windows Server 2003 domain, see Overview of Upgrading Windows NT 4.0 Domains in *Designing and Deploying Directory and Security Services* of this kit.

# Determining When Group Policy Changes are Applied

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Changes to Group Policy settings might not be immediately available on users' desktops because changes to the Group Policy object must first replicate to the appropriate domain controller. In addition, clients use a 90-minute refresh period (randomized by up to approximately 30 minutes) for the retrieval of Group Policy. Therefore, it is rare for a changed Group Policy setting to apply immediately. Components of a GPO are stored in both Active Directory and on the Sysvol folder of domain controllers. Replication of a GPO to other domain controllers occurs by two independent mechanisms:

- Replication in Active Directory is controlled by Active Directory's built-in replication system. By default, this typically takes less than a minute between domain controllers within the same site, if these domain controllers are part of a fresh-installed (not upgraded) forest running on Windows Server 2003–based domain controllers, or if the upgraded forest's Functional Level is set to Windows Server 2003. However, in environments such as a partially upgraded forest that contains domain controllers running Windows 2000 and Windows Server 2003, a typical replication might take up to 15 minutes. This process can be even slower if your network is slower than a LAN.

- Replication of the Sysvol folder is controlled by the File Replication service (FRS). Within sites, replication occurs every 15 minutes. If the domain controllers are in different sites, the replication process occurs at set intervals based on site topology and schedule; the lowest interval is 15 minutes. In a Windows Server 2003 LAN environment, replication typically takes less than a minute.

**Note**

- If it is critical to immediately apply a change to a specific group of users or computers in a specific site, you can connect to the domain controller closest to these objects, and then make the configuration change on that domain controller so those users get the updated policy first.

**Policy Refresh Interval**

The primary mechanisms for refreshing Group Policy are startup and logon. Group Policy is also refreshed on a regular basis. The policy refresh interval affects how quickly changes to Group Policy objects are applied. By default, clients and servers running Windows 2000, clients running Windows XP Professional, and destination servers running Windows Server 2003 check for changes to Group Policy objects every 90 minutes by using a randomized offset of up to 30 minutes. Domain controllers running Windows 2000 Server or Windows Server 2003 check for computer policy changes every five minutes. Although this polling frequency can be changed (by using one of these policy settings: **Group Policy Refresh Interval for Computers,Group Policy Refresh Interval for Domain Controllers**, or **Group Policy refresh Interval for Users**), shortening the frequency between refreshes is not recommended because of the potential increase in network traffic and the additional load placed on the domain controllers.

**Triggering a Group Policy Refresh**

If necessary, you can trigger a policy refresh manually from a local computer without waiting for the automatic background refresh. To do this, you can type `gpupdate` at the command line to refresh the user or computer policy settings. You cannot trigger a policy refresh by using GPMC.

The **gpupdate** command triggers a background policy refresh on the local computer from which the command is run. The **gpupdate** command is used in Windows Server 2003 and Windows XP environments. In Windows 2000, use the **secedit /refreshpolicy** command.

For more information about the **gpupdate** command, see "Changing the Group Policy Refresh Interval" later in this chapter.

**Note**

- Some policy settings, such as folder redirection and the assignment of software applications, require the user to log off and log on again before they take effect. Software applications assigned to computers are installed only when the computer is restarted. See Help and Support Center for Windows Server 2003 for specific settings for information about when those settings take effect.

# Identifying Issues Pertaining to Software Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although Group Policy can successfully install software applications – especially in small- or medium-sized organizations – you need to determine if it is the best solution for your needs. Using Group Policy, assigned applications are installed or updated only when the computer is restarted or when the user logs on.

Using SMS for software deployment provides enterprise-level functionality that is not available with Group Policy-based software deployment, such as inventory-based targeting, status reporting, and scheduling. For this reason, you might use Group Policy to configure the desktop, set system security and access permissions, and use SMS to deliver software applications. This approach provides bandwidth control by scheduling application installation outside core business hours.

Your choice of tools depends on your requirements, your environment, and whether or not you need the additional functionality and security that SMS provides.

For more information about the Software Installation extension of Group Policy, see "Deploying a Managed Software Environment" in this book.

# Designing Your Group Policy Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your primary objective is to design the GPO structure based on your business requirements. Keeping in mind the computers and users in your organization, determine which policy settings must be enforced across the organization, as well as which policy settings are applicable to all users or computers. Also determine which settings to use to configure computers or users according to type, function, or job role. Then group these different types of policy settings into GPOs and link them to the appropriate Active Directory containers.

Also, keep in mind the Group Policy inheritance model and how precedence is determined. By default, options set in GPOs linked to higher levels of Active Directory containers —sites, domains and OUs — are inherited by all containers at lower levels. However, inherited policy can be overridden by a GPO that is linked at a lower level. For example, you might use a GPO linked at a high level for assigning standard desktop wallpaper, but want a certain OU to get different wallpaper. To do so, you can link a second GPO to that specific lower-level OU. Because lower-level GPOs apply last, the second GPO will override the domain-level GPO and provide that specific lower-level OU with a different set of Group Policy settings. However, you can modify this default inheritance behavior by using **Block Inheritance** and **Enforced**.

Figure 2.4 illustrates the steps detailed in this section.

**Figure 2.4 Group Policy Design Model**




The following guidelines can help tailor your Group Policy design to the needs of your organization:

- Determine if there are any policy settings that must always be enforced for particular groups of users or computers. Create GPOs that contain these settings,link them to the appropriate site, domain, or OU, and designate these links as **Enforced** (formerly known as **No Override**). By setting this option, you enforce a higher-level GPO's settings by preventing GPOs in lower-level Active Directory containers from overriding them. For example, if you define a specific GPO at the domain level and specify that it is enforced, the policies that the GPO contains apply to all OUs under that domain; GPOs linked to the lower-level OUs cannot override that domain Group Policy.

  **Note**

  - Use the **Enforced** and **Block Policy Inheritance** features sparingly. Routine use of these features can make it difficult to troubleshoot policy because it is not immediately clear to administrators of other GPOs why certain settings do or do not apply.

- Decide which policy settings are applicable to the entire organization and consider linking these to the domain. You can also use GPMC to copy GPOs or import GPO settings, thereby creating identical GPOs in different domains.

- Link the GPOs to the OU structure (or site), and then use Security Groups to selectively apply these GPOs to particular users or computers.

- Classify the types of computers and the roles or job function of users in your organization, group them into OUs, create GPOs to configure the environment for each as needed, and then link the GPOs to those OUs.

- Prepare a staging environment to test your Group Policy-based management strategy before deploying GPOs into your production environment. Think of this phase as staging your deployment. This is a crucial step toward ensuring that your Group Policy deployment will meet your management goals. This process is fully detailed in "Staging Group Policy Deployments," in this book.

# Defining the Scope of Application of Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To define the scope of application of Group Policy objects, consider these three main questions:

- Where will your GPOs be linked?
- What security filtering on the GPOs will you use?
- What WMI filters will be applied?

Also, remember that by default, Group Policy objects are inherited, cumulative, and affect all computers and users in an Active Directory container and its children. They are processed in the following order: local GPO (LGPO), site, domain, and then OU, with the last processed overriding the earlier GPOs. The default inheritance method is to evaluate Group Policy starting with the Active Directory container farthest from the computer or user object. The Active Directory container closest to the computer or user overrides Group Policy set in a higher-level Active Directory container unless you set the **Enforced** (**No Override**) option for that GPO link or if the **Block Policy inheritance** setting has been applied to the domain or OU. The LGPO is processed first, so settings from GPOs linked to Active Directory containers override the local settings. For more information about LGPOs, see the Windows Security Collection in the Windows Server 2003 Technical Reference on the Web (http://go.microsoft.com/fwlink/?linkid=4571).

Another issue is that although you can link more than one GPO to an Active Directory container, you need to be aware of the processing order (priority). The GPO link with the lowest link order in the **Group Policy Object Links** list (displayed in the **Linked Group Policy Objects** tab in GPMC) has precedence by default. However, if one or more GPO links have the **Enforced** option set, the highest GPO link set to **Enforced** takes precedence.

Stated briefly, **Enforced** is a link property, **Block Policy Inheritance** is a container property, and **Enforced** takes precedence over **Block Policy Inheritance**. In addition, you can disable settings on the GPO itself in four other ways: A GPO can be disabled; and a GPO can have its computer settings disabled, its user settings disabled, or all of its settings disabled.

GPMC greatly simplifies these tasks by allowing you to view GPO inheritance across your organization and manage links from one MMC console. Figure 2.5 shows Group Policy inheritance as displayed in GPMC.

**Figure 2.5 Group Policy Management Linking and Inheritance**




**Note**

- To view full details of inheritance and precedence for GPO links to a domain, site, or OU, you must have Read permissions on the domain, site, or OU containing the GPO links as well as on the GPOs. If you have Read access to the domain, site, or OU, but not on one of the GPOs linked there, it will appear as **Inaccessible GPO**, and you will not be able to read the name or other information for that GPO.

# Determining the Number of Group Policy Objects

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The number of Group Policy objects you need depends on your approach to design, the complexity of the environment, your objectives, and the scope of the project. If you have a multi-domain forest, or multiple forests, you might find that the number of GPOs required in each domain is different. Domains supporting highly complex business environments with a diverse user population typically require more GPOs than smaller, simpler domains.

As the number of GPOs required to support an organization increases, so can the workload of Group Policy administrators. There are steps you can take to ease the administration of Group Policy. In general, you should group settings that apply to a given set of users or computers and are managed by a common set of administrators into a single GPO. Further, if various groups of users or computers have common requirements, and only a few of the groups need incremental changes, consider applying the common requirements to all these groups of users or computers by using a single GPO linked high in the Active Directory structure, and then add additional GPOs, which apply only the incremental changes, at the relevant OU. This might not always be possible or practical, so you might need to make exceptions to this guideline. If so, be sure to keep track of them. Note that a maximum of 999 GPOs is supported for processing GPOs on any one user or computer. If you exceed the maximum, no GPOs will be processed. This limitation affects only the number of GPOs that can be applied at the same time; it does not affect the number of GPOs you can create and store in a domain.

Consider that the number of GPOs applied to a computer affects startup time, and the number of GPOs applied to a user affects the amount of time needed to log on to the network. The greater the number of Group Policy objects that are linked to a user — especially the greater the number of settings within those GPOs — the longer it takes to process them whenever a user logs on. During the logon process, each GPO from the user's site, domain, and OU hierarchy is applied, provided both the Read and Apply Group Policy permissions are set for the user. In GPMC, the Read and Apply Group Policy permissions are managed as a single unit called Security Filtering.

If you use Security Filtering and you remove the Apply Group Policy permission for a given user or group, also remove the Read permission, unless you need that user to have read access for some reason. (If you are using GPMC, you need not worry about this, because GPMC does this for you automatically.) If the Apply Group Policy permission is not set, but the Read permission is, the GPO is still inspected (although not applied) by any user or computer that is in the OU hierarchy where the GPO is linked. This inspection process increases logon time slightly.

Always test your Group Policy solution in a test environment to ensure that the policy settings you define do not unacceptably prolong the time it takes to display the logon screen, and that they comply with desktop service level agreements. During this staging period, log on with a test account to gauge the net effect of several GPOs on objects in your environment.

# Linking GPOs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To apply the settings of a GPO to the users and computers of a domain, site, or OU, you need to add a link to that GPO. You can add one or more GPO links to each domain, site, or OU by using GPMC. Keep in mind that creating and linking GPOs is a sensitive privilege that should be delegated only to administrators who are trusted and understand Group Policy.

**Linking GPOs to the Site**

If you have a number of policy settings to apply to computers in a particular physical location only — certain network or proxy configuration settings, for example — these settings might be appropriate for inclusion in a site-based policy. Because domains and sites are independent, it is possible that computers in the site might need to cross domains to link the GPO to the site. In this case, make sure there is good connectivity.

If, however, the settings do not clearly correspond to computers in a single site, it is better to assign the GPO to the domain or OU structure rather than to the site.

**Linking GPOs to the Domain**

Link GPOs to the domain if you want them to apply to all users and computers in the domain. For example, security administrators often implement domain-based GPOs to enforce corporate standards. They might want to create these GPOs with the GPMC **Enforce** option enabled to guarantee that no other administrator can override these settings.

**Important**

- If you need to modify some of the settings contained in the **Default Domain Policy GPO**, it is recommended that you create a new GPO for this purpose, link it to the domain, and set the **Enforce** option. In general, do not modify this or the **Default Domain Controller Policy GPO**. If you do, be sure to back up these and any other GPOs in your network by using GPMC to ensure you can restore them.

As the name suggests, the **Default Domain Policy GPO** is also linked to the domain. The **Default Domain Policy GPO** is created when the first domain controller in the domain is installed and the administrator logs on for the first time. This GPO contains the domain-wide account policy settings, Password Policy, Account Lockout Policy, and Kerberos Policy, which is enforced by the domain controller computers in the domain. All domain controllers retrieve the values of these account policy settings from the **Default Domain Policy GPO**. In order to apply account policies to domain accounts, these policy settings must be deployed in a GPO linked to the domain, and it is recommended that you set these settings in the Default Domain Policy. If you set account policies at a lower level, such as an OU, the settings only affect local accounts (non-domain accounts) on computers in that OU and its children.

Before making any changes to the default GPOs, be sure to back up the GPO using GPMC. If for some reason there is a problem with the changes to the default GPOs and you cannot revert back to the previous or initial states, you can use the Dcgpofix.exe tool to recreate the default policies in their initial state.

Dcgpofix.exe is a command-line tool that completely restores the Default Domain Policy GPO and Default Domain Controller GPO to their original states in the event of a disaster where you cannot use GPMC. Dcgpofix.exe restores only the policy settings that are contained in the default GPOs at the time they are generated. The only Group Policy extensions that include policy settings in the default GPOs are RIS, Security, and EFS. Dcgpofix.exe does not restore other GPOs that administrators create; it is only intended for disaster recovery of the default GPOs.

Note that Dcgpofix.exe does not save any information created through applications, such as SMS or Exchange. The Dcgpofix.exe tool is included with Windows Server 2003 and only works in a Windows Server 2003 domain.

Dcgpofix.exe is located in the C:\Windows\Repair folder. The syntax for Dcgpofix.exe is as follows:

```
DCGPOFix [/Target: Domain | DC | BOTH]
```

Table 2.1 describes the options you can use with the command line parameter /Target: when using the Dcgpofix.exe tool.

**Table 2.1 Dcgpofix.exe Options for Using the /Target Parameter**

| /Target option: | Description of option |
|---|---|

| | |
|---|---|
| DOMAIN | Specifies that the Default Domain Policy should be recreated. |
| DC | Specifies that the Default Domain Controllers Policy should be recreated. |
| BOTH | Specifies that both the Default Domain Policy and the Default Domain Controllers Policy should be recreated. |

For more information about Dcgpofix.exe, in Help and Support Center for Windows Server 2003 click **Tools**, and then click Command-line reference A-Z

**Linking GPOs to the OU Structure**

Most GPOs are normally linked to the OU structure because this provides the most flexibility and manageability:

- You can move users and computers into and out of OUs.
- OUs can be rearranged if necessary.
- You can work with smaller groups of users who have common administrative requirements.
- You can organize users and computers based on which administrators manage them.

Organizing GPOs into user- and computer-oriented GPOs can help make your Group Policy environment easier to understand and can simplify troubleshooting. However, separating the user and computer components into separate GPOs might require more GPOs. You can compensate for this by adjusting the **GPO Status** to disable the user or computer configuration portions of the GPO that do not apply and to reduce the time required to apply a given GPO.

**Changing the GPO Link Order**

Within each domain, site, and OU, the link order controls the order in which GPOs are applied. To change the precedence of a link, you can change the link order, moving each link up or down in the list to the appropriate location. Links with the lowest number have higher precedence for a given site, domain, or OU. For example, if you add six GPO links and later decide that you want the last one that you added to have the highest precedence, you can adjust the link order of the GPO link so it has link order of 1. To change the link order for GPO links for a domain, OU, or site, use GPMC.

# Using Security Filtering to Apply GPOs to Selected Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, a GPO affects all users and computers contained in the linked site, domain, or OU. However, you can use Security Filtering on a GPO to modify its effect to apply only to a specific user or the members of a security group by modifying the permissions on the GPO. By combining Security Filtering with appropriate placement in OUs, you can target any given set of users.

In order for a GPO to apply to a given user or computer, that user or computer must have both **Read** and **Apply Group Policy** permissions on the GPO. By default, Authenticated Users have both **Apply Group Policy** and **Read** permissions set to **Allow**. Both of these permissions are managed together as a single unit by using Security Filtering in GPMC.

To set the permissions for a given GPO, in the GPMC console tree, expand **Group Policy Objects** in the forest and domain containing that GPO. Click the GPO, and in the details pane, on the **Scope** tab, under **Security Filtering**, remove **Authenticated Users**, and then click **Add** to add a new group.

To filter a GPO by a particular security group, add that group to the Security Filtering section on the **Scope** tab of a GPO in GPMC.

For example, if you want only a subset of users within an OU to receive a GPO, remove the Authenticated Users from Security Filtering. Instead, add a new group with Security Filtering permissions that contains the subset of users who are to receive the GPO. Only members of this group that are within the site, domain, or OU where the GPO is linked receive the GPO; members of the group in other sites, domains, or OUs do not receive the GPO.

You might want to prevent certain Group Policy settings from applying to the Administrator group. To accomplish this, you can do one of the following:

- Create a separate OU for administrators and keep this OU out of the hierarchy to which you apply most of your management. In this way administrators do not receive most of the settings that that you provide for managed users. If this separate OU is a direct child of the domain, the only possible settings administrators receive are settings from GPOs linked either to the domain or the site. Typically, only generic, broadly applicable settings are linked here, so it might be acceptable to have administrators receive these settings. If this is not what you intend, you can set the **Block Inheritance** option on the administrators' OU.

- Have administrators use separate administrative accounts for use only when they perform administrative tasks. When not performing administrative tasks, they would still be managed.

- Use Security Filtering in GPMC so that only non-administrators will receive the settings.

# Applying WMI Filters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use WMI filters to control the application of GPOs. Each GPO can be linked to one WMI filter, however, the same WMI filter can be linked to multiple GPOs. Before you can link a WMI filter to a GPO, you must create the filter. The WMI filter is evaluated on the destination computer (running either Windows XP or Windows Server 2003) during processing of Group Policy. The GPO will only apply if the WMI filter evaluates to TRUE. On Windows 2000–based computers, the WMI filter is ignored and the GPO is always applied.

Using GPMC, you can perform the following operations for WMI filters: create and delete, link and unlink, copy and paste, import and export, and view and edit attributes.

WMI filters can only be used if at least one domain controller in the domain is running Windows Server 2003, or if you have run ADPrep with the /Domainprep option in that domain. If not, the **WMI Filtering** section on the **Scope** tab for GPOs and the **WMI Filters** container under the domain will not be present. See Figure 2.6 to help you identify the items discussed in this section.

**Figure 2.6 GPMC Tab Scope**




**Setting WMI Filtering Options**

WMI exposes management data from a destination computer, such as hardware and software inventory, settings, and configuration information, including data from the registry, drivers, the file system, Active Directory, SNMP, Windows Installer, and networking. Administrators can create WMI filters, which consist of one or more queries based on this data — to control whether the GPO is applied. The filter is evaluated on the destination computer (which must be running either Windows XP Professional or Windows Server 2003). If the WMI filter evaluates to true, the GPO is applied to that destination computer; if the filter evaluates to false, the GPO is not applied. On Windows 2000–based client or server targets, WMI filters are ignored, and the GPO is always applied. In the absence of any WMI Filter, the GPO is always applied.

Administrators can use WMI filters to target Group Policy based on a variety of objects and other parameters. Table 2.2 illustrates example query criteria that might be specified for WMI filters.

**Table 2.2 Sample WMI Filters**

| WMI Data Queried | Sample Query Criteria |
|---|---|
| Services | Computers with the DHCP service running |

| | |
|---|---|
| Registry | Computers that have a specified registry key or entry populated |
| Windows Event Log | Computers that reported an audit event in the last five minutes |
| Operating system version | Computers running Windows Server 2003 and later |
| Hardware inventory | Computers with a Pentium III processor |
| Hardware configuration | Computers with network adapters on in level 3 |
| Service associations | Computers that have any service dependent on SNA service |
| Ping | Computers that can ping Server1 in less than 100 milliseconds |

A WMI filter consists of one or more WMI Query Language (WQL) queries. The WMI filter applies to every setting in the GPO, so administrators must create separate GPOs if they have different filtering requirements for different settings. The WMI filters are evaluated on the destination computer after the list of potential GPOs is determined and filtered based on security group membership. For detailed information about WMI filters, see the Windows Security Collection in the Windows Server 2003 Technical Reference on the Web (http://go.microsoft.com/fwlink/?linkid=4571).

Although you can do limited inventory-based targeting for software deployment by combining Group Policy-based software deployment with WMI filters, this is not recommended as a general practice for the following reasons:

- Each GPO can only have one WMI Filter. If applications have different inventory requirements, you need multiple WMI filters and therefore multiple GPOs. Increasing the number of GPOs impacts startup and logon times and also increases management overhead.

- WMI filters can take significant time to evaluate, so they can slow down logon and startup time. The amount of time depends on the construction of the query.

**Note**

- It is recommended that you use WMI filters primarily for exception management. They can be powerful solutions for targeting GPOs to specific users and computers, but because WMI filters are evaluated every time Group Policy is processed, they increase startup and logon time. Also, there is no time-out for WMI filters. Use them only when necessary.

**Example WMI Filters**

WMI filters are most useful as tools for exception management. By filtering for particular criteria, you can target particular GPOs to only specific users and computers. The following are WMI filters that illustrate this technique.

**Targeting Based on Operating System**

In this example, an administrator wants to deploy an enterprise monitoring policy, but wants to target only Windows XP Professional–based computers. The administrator can create a WMI filter such as the following:

```
Select * from Win32_OperatingSystem where Caption = " Microsoft Windows XP
Professional"
```

Most WMI filters use the **Root\CimV2** namespace, and this option is populated by default in the GPMC user interface.

Because WMI filters are ignored on Windows 2000-based computers, a filtered GPO will always be applied on them. However, you can work around this by using two GPOs and giving the one with Windows 2000 settings higher precedence (using link order). Then use a WMI filter for that Windows 2000 GPO, and only apply it if the operating system is Windows 2000, not Windows XP Professional. The Windows 2000-based computer will receive the Windows 2000 GPO and will override the settings in the Windows XP Professional GPO. The Windows XP Professional client will receive all the settings in the Windows XP Professional GPO.

**Targeting Based on Hardware Inventory**

In this example, an administrator wants to distribute a new network connection manager tool only to desktops that have modems. The administrator can deploy the package by using the following WMI filter to target those desktops:

```
Select * from Win32_POTSModem Where Name = " MyModem"
```

If you use Group Policy with a WMI filter, remember that the WMI filter applies to all settings in the GPO. If you have different

requirements for different deployments, you need to use different GPOs, each with its own WMI filter.

**Targeting Based on Configuration**
In this example, an administrator does not want to apply a GPO on computers that have multicasting turned on. The administrator can use the following filter to identify those computers which have multicast support:

```
Select * from Win32_NetworkProtocol where SupportsMulticasting = true
```

**Targeting Based on Amount of Disk Space and File System Type**
In this example, an administrator wants to target computers that have more than 10 megabytes (MB) of available space on the C, D, or E partition. The partitions must be located on one or more local fixed disks and they must be running the NTFS file system. The administrator can use the following filter to identify computers that meet these criteria:

```
SELECT * FROM Win32_LogicalDisk WHERE (Name = " C:"  OR Name = " D:"  OR Name = " E:" )
AND DriveType = 3 AND FreeSpace > 10485760 AND FileSystem = " NTFS"
```

In the preceding example, DriveType value = 3 represents a local disk and FreeSpace units are in Bytes (10 MB = 10,485,760 Bytes).

**To Create a WMI Filter**

1. In the GPMC console tree, right-click **WMI Filters** in the forest and domain in which you want to add a WMI filter.

2. Click **New**.

3. In the **New WMI Filter** dialog box, type a name for the new WMI Filter in the **Name** box, and then type a description of the filter in the **Description** box.

4. Click **Add**.

5. In the **WMI Query** dialog box, either leave the default namespace (root\CIMv2) or enter another namespace by doing one of the following:

    - In the **Namespace** box, type the name of the namespace that you want to use for the WMI query. The default is **Root\CimV2.** In most cases, you do not need to change this.

    - Click **Browse**, select a namespace from the list, and then click **OK**.

6. Type a WMI query in the **Query** box, and then click **OK**.

7. To add more queries, repeat steps 4 through 6 to add each query.

8. After you add all the queries, click **Save**.

The WMI filter is now available to be linked.

# Using Group Policy Inheritance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is often useful to define a corporate-standard GPO. As used here, *corporate standard* refers to policy settings that apply to a broad set of users in an organization. An example of where defining a corporate standard GPO might be appropriate is a business requirement that states: "Only specially authorized users can access the command prompt or the registry editor." Group Policy inheritance can help you apply these corporate standards while customizing settings for different groups of users.

One way to do this is to set the policy settings **Prevent access to the command prompt**, and **Prevent access to registry editing tools** in a GPO, such as the Standard User Policy GPO that is linked to an OU, such as the User Accounts OU. By default this will apply these settings to all users in that OU. Then create a GPO, such as an Administrator User Policy GPO, which explicitly allows administrators access to the command prompt and registry editing tools. Link the GPO to the Administrators OU, which overrides the settings configured in the Standard User Policy GPO. This approach is illustrated in Figure 2.7.

**Figure 2.7 Standard User Policy GPO**




If another group of users requires access to the command prompt, but not the registry, you can create another GPO that allows them to do so. Access to the registry editing tools is still denied because the new GPO does not override the registry tools setting made in the Standard User Policy GPO. Typically, a corporate standard GPO includes more settings and configuration options than those shown in the preceding illustration. For example, corporate standard Group Policy objects are typically used to achieve the following:

- Remove all potentially harmful and nonessential functionality for users.
- Define access permissions, security settings, and file system and registry permissions for member servers and workstations.

Typically, GPOs are assigned to the OU structure instead of the domain or site. If you structure your OU model around users, workstations, and servers, it is easier to identify and configure corporate standard settings. You can also disable either the user or computer portions of policy that do not apply, making Group Policy easier to manage.

When you set default values for security-related settings such as restricted group membership, file system access permissions, and registry access permissions, it is important to understand that these settings work on a last-writer-wins principle, and that the settings are not merged. The following example demonstrates this principle.

**Example: Last-Writer-Wins Principle**
An administrator creates a Default Workstations GPO that defines the membership of the local Power Users group as the Technical Support and Help Desk groups. The Business Banking group wants to add the Business Banking Support group to this list and creates a new Default Workstations GPO to do so. Unless the new GPO specifies that allthree groups are members of Power Users, only the Business Banking Support group has Power User rights on affected workstations.

# Deploying Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before deploying your Group Policy solution, administrators need to be familiar with the procedures for working with GPOs, including creating GPOs, importing settings, backing up and restoring GPOs, editing and linking GPOs, setting exceptions to default inheritance of GPOs, filtering the application of GPOs, delegation of administration, and using Group Policy Modeling for planning and Group Policy Results for evaluating GPO application. Figure 2.8 illustrates the place of this step in the process of designing a Group Policy infrastructure.

**Figure 2.8 Deploying Group Policy**




Always fully test your GPOs in safe (nonproduction) environments prior to production deployment. The more you plan, design, and test GPOs prior to deployment, the easier it is to create, implement, and maintain an optimal Group Policy solution. The importance of testing and pilot deployments in this context cannot be overemphasized. Your tests should closely simulate your production environment.

A design is not complete until you test and validate all its significant variations and your deployment strategy. Thorough testing of your GPO implementation strategy is not possible until you configure your GPOs using specific settings, such as security settings, and desktop and data management. Do this for each group of users and computers in the network. Use your test environment to develop, test, and validate specific GPOs. Take full advantage of the GPMC Modeling Wizard and the Results Wizard.

Consider an iterative implementation of Group Policy. That is, rather than deploying 100 new Group Policy settings, stage and then initially deploy only a few settings to validate that the Group Policy infrastructure is working well.

For more information about this phase, see "Staging Group Policy Deployments" in this book. This chapter details the information you need to deploy Group Policy into your production environment, but never deploy an untested GPO. For step-by-step details about creating a new GPO, search Help and Support Center for Windows Server 2003 using the keywords "Group Policy Object Editor." For step-by-step details about using GPMC to perform other GPO tasks such as importing settings from one GPO into another, see Help in GPMC.

Table 2.3 lists the six sample standard desktop configurations that are described in "Planning a Managed Environment" in this book. Examining these standard configurations can help you configure your Group Policy deployment. Each file is a GPO with sample standard desktop configurations you can copy into your test environment. The actual settings for each GPO are listed in the corresponding HTML reports.

**Table 2.3 Sample Standard Desktop Configurations**

| Standard Desktop Configuration | Scenario File Name |
|---|---|
| Lightly managed desktop | LightlyManagedDesktop.htm |
| Mobile users desktop | MobileUser.htm |
| Multi-user desktop | MultiUserDesktop.htm |

| | |
|---|---|
| Highly managed desktop | Appstation.htm |
| TaskStation desktop | Taskstation.htm |
| Kiosk | Kiosk.htm |

For information about using the sample GPOs, see the whitepaper at the Implementing Common Desktop Management Scenarios link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Creating and Working with GPOs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because changes to a GPO take place immediately, keep the GPO unlinked from its production location (site, domain, or OU) until you have fully tested it in a test environment. While you are developing the GPO, keep it either unlinked or linked to a test OU.

For more information about testing your Group Policy configurations prior to deployment, see "Staging Group Policy Deployments" in this book. This section details the process of creating and deploying GPOs.

The following procedures detail creating GPOs by using GPMC:

**To create an unlinked Group Policy object**

1. Expand the GPMC console tree until you see **Group Policy Objects** in the container domain for which you want to create a new GPO.

2. Right-click **Group Policy Objects**.

3. Click **New**.

4. In the **New GPO** dialog box, specify a name for the new GPO, and then click **OK**.

**To edit a GPO**

1. In the GPMC console tree, right-click the GPO you want to edit, and then click **Edit**. This starts the Group Policy Object Editor MMC snap-in.

2. In the Group Policy Object Editor console tree, expand the items to view their policies in the details pane.

3. In the details pane, double-click the names of the settings to open their **Properties** dialog boxes and change their settings. Note that some settings, such as the settings for deploying a new software installation package, use unique user interfaces.

The primary mechanism for applying the settings in a GPO to users and computers is by linking the GPO to a container in Active Directory. GPOs can be linked to three types of containers in Active Directory: sites, domains, and organizational units. A GPO can be linked to multiple Active Directory containers.

**To link a GPO**

GPOs are stored on a per-domain basis. For example, if you link a GPO to an OU, the GPO is not actually located in that OU. A GPO is a per domain object that can be linked anywhere in the forest. The UI in GPMC helps clarify the distinction between links and actual GPOs. In GPMC, you can link a GPO to Active Directory containers using any of the following methods:

- Right-click a site, domain, or OU item, and then click **Link an existing GPO here**. This option is equivalent to choosing **Add** on the **Group Policy** tab that was available in the Active Directory Users and Computers MMC snap-in, prior to installing GPMC. This requires that the GPO already exist in the domain.

- Drag a GPO from under the **Group Policy objects** item to the OU. This drag-and-drop functionality works only within the same domain.

You can also use the GPMC user interface to simultaneously create a new GPO and link it at the same time in one step, as described below.

**To create and link a GPO**

- Right-click a domain or OU item, and then click **Create and Link a GPO here**. This option is equivalent to clicking **New**

on the **Group Policy** tab that was available in the Active Directory Users and Computers snap-in, prior to installing GPMC. In the **New GPO** dialog box, type a name for the new GPO, and then click **OK**. Although this operation is presented in GPMC as one action to the user, there are actually two steps taking place. First, a GPO is created in the domain, and second, the new GPO is linked to the domain or OU.

To create a GPO and link it to a site, you must first create the GPO in the domain, and then link it.

**To unlink a GPO from a site, domain, or OU**

1. In the GPMC console tree, right-click the GPO you want to unlink.
2. In the details pane, click the **Scope** tab.
3. In the **Links** section, right-click the Active Directory object.
4. Select **Delete Link** from the drop-down list.

### Note

- Deleting a link is different than deleting a GPO. It you delete only the link, the GPO still exists, as do any other existing links to that GPO. However, if you delete a GPO, you will be prompted to delete the GPO and all links to it in this domain. This does not delete links to the GPO from other domains. Be sure to remove links to the GPO in other domains before deleting this GPO in this domain.

**Disabling the User or Computer Configuration Property of a GPO**

If you are creating a GPO to set only user-related policy settings, you can disable the Computer Configuration portion of the GPO. Doing this slightly reduces computer startup time because the Computer GPO does not have to be evaluated to determine if any policies exist. If you are configuring only computer-related policy settings, turn off the User Configuration portion of the GPO.

See Figure 2.9 to help you identify the GPMC items referred to in the procedure that follows.

**Figure 2.9 Setting GPO Status**




**To disable the User or Computer Configuration settings in a GPO**

1. In the GPMC console tree, click the GPO you want to modify.
2. Click the **Details** tab, or right-click the GPO.
3. In the **GPO Status** drop-down box, select one of these choices:

- **All settings disabled**
- **Computer settings disabled**
- **Enabled (default)**
- **User settings disabled**

**Special Considerations for Site-linked GPOs**

GPOs linked to sites might be appropriate to use for setting policy for proxy settings and network-related settings. Any GPO that is linked to a site container is applied to all computers in that site, regardless of which domain in the forest the computer belongs to. This has the following implications:

- Ensure that the computers do not access a site Group Policy object across a WAN link, which would lead to significant performance issues.
- By default, to manage site GPOs, you need to be either an Enterprise Admin, or the domain admin of the forest root domain.
- Active Directory service data-replication between domain controllers in different sites occurs less frequently than replication between domain controllers in the same site, and occurs during scheduled periods only. Between sites, FRS replication occurs spontaneously, and is not determined by the site link replication schedule; this is not an issue within sites. The directory service replication schedule and frequency are properties of the site links that connect sites. The default inter-site replication frequency is three hours. To change it, go to the appropriate site link, go to the IP link, and change the replication frequency or schedule as needed.

Changing either the replication frequency or schedule can significantly affect policy. For example, assume that you have replication set to three hours or longer, and you create a GPO and link it to an OU in a domain that spans several sites. You will likely need to wait several hours before all users in that OU receive the GPO.

If most of the users in an OU are in a remote location, and you have a domain controller in that site, you can work around inter-site replication latency by performing all Group Policy operations on a domain controller in that site.

# Using Loopback Processing to Configure User Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The **User Group Policyloopback processing mode** policy setting is an advanced option that is intended to keep the configuration of the computer the same regardless of who logs on. This option is appropriate in certain closely managed environments, such as servers, terminal servers, classrooms, public kiosks, and reception areas. Setting the loopback processing mode policy setting applies the same user settings for any user who logs onto the computer, based on the computer.

When you apply Group Policy objects to users, normally the same set of user policy settings applies to those users when they log on to any computer. By enabling the loopback processing policy setting in a GPO, you can configure user policy settings based on the computer that they log on to. Those settings are applied regardless of which user logs on. When you use this option, you must ensure that both the computer and user portions of the GPO are enabled.

You can set the loopback policy in the Group Policy Object Editor snap-in by using the **User Group Policy loopback processing mode** policy setting under **Computer Settings\Administrative settings\System\Group Policy**. Two options are available:

**Merge mode** In this mode, the list of GPOs for the user is gathered during the logon process. Then, the list of GPOs for the computer is gathered. Next, the list of GPOs for the computer is added to the end of the GPOs for the user. As a result, the computer's GPOs have higher precedence than the user's GPOs.

**Replace mode** In this mode, the list of GPOs for the user is not gathered. Instead, only the list of GPOs based on the computer object is used. The User Configuration settings from this list are applied to the user.

# Delegating Administration of Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your Group Policy design will probably call for delegating certain Group Policy administrative tasks. Determining to what degree to centralize or distribute administrative control of Group Policy is one of the most important factors to consider when assessing the needs of your organization. In organizations that use a centralized administration model, an IT group provides services, makes decisions, and sets standards for the entire company. In organizations that use a distributed administration model, each business unit manages its own IT group.

You can delegate the following Group Policy tasks:

- Creating GPOs
- Managing individual GPOs (for example, granting Edit or Read access to a GPO)
- Performing the following tasks on sites, domains, and OUs:
    - Managing Group Policy links for a given site, domain, or OU
    - Performing Group Policy Modeling analyses for objects in that container (not applicable for sites)
    - Reading Group Policy Results data for objects in that container (not applicable for sites)
- Creating WMI filters
- Managing and editing individual WMI filters

Based on your organization's administrative model, you need to determine which aspects of configuration management can best be handled at the site, domain, and OU levels. You also need to determine how responsibilities at each site, domain, and OU level might be further subdivided among the available administrators or administrative groups at each level.

When deciding whether to delegate authority at the site, domain, or OU level, remember the following points:

- Authority delegated at the domain level affects all objects in the domain, if the permission is set to inherit to all child containers.
- Authority delegated at the OU level can affect either that OU only, or that OU and its child OUs.
- Managing permissions is easier and more efficient if you assign control at the highest OU level possible.
- Authority delegated at the site level is likely to span domains and can influence objects in domains other than the domain where the GPO is located.

Following are descriptions of how to use GPMC to perform these delegation tasks.

**Delegating management of individual GPOs**
Using GPMC, you can easily grant additional users permissions on a GPO. GPMC manages permissions at the task level. There are five levels of allowed permissions on a GPO: Read, Edit, and Edit/Delete/Modify Security, Read (from Security Filtering), and Custom. These permission levels correspond to a fixed set of low-level permissions. Table 2.4 shows the corresponding low-level permissions for each option.

**Table 2.4 GPO Permission Options and Low Level Permissions**

| GPO Permission Option | Low Level Permissions |
|---|---|
| Read | Allow Read Access on the GPO. |
| Read (from Security Filtering) | This setting cannot be set directly, but appears if the user has Read and Apply Group Policy permissions to the GPO, which is set using Security Filtering on the **Scope** tab of the GPO. |
| Edit settings | Allow Read, Write, Create Child Objects, Delete Child Objects. |
| Edit, delete, and modify security | Allow Read, Write, Create Child Objects, Delete Child Objects, Delete, Modify Permissions, and Modify Owner. This essentially grants full control on the GPO, except that the "Apply Group Policy" permission is not set. |
| Custom | Any other combinations of rights, such denying permissions, appear as **Custom** permissions. You cannot set custom rights by clicking **Add**. They can only be set by using the ACL editor directly, which can be started by clicking the **Advanced** button. |

You can click **Add** to grant users permissions on a GPO. This starts the object picker so you can find the desired user or group to set the permission level. You can then set the permission level by selecting the **Read**, **Edit**, or **Edit, Delete, Modify Security** permissions.

Note that the **Apply Group Policy** permission, which is used for Security Filtering, cannot be set using the **Delegation** tab. Because setting **Apply Group Policy** is used for scoping the GPO, this permission is managed on the **Scope** tab of the GPMC user interface. When you grant a user Security Filtering on the **Scope** tab, you are actually setting both the **Read** and **Apply Group Policy** permissions.

You can grant additional groups and users permissions on a GPO by using GPMC, as described above. Table 2.5 lists the default security permission settings for a GPO.

**Table 2.5 Default Security Permissions for GPOs.**

| Security Group | Permissions |
|---|---|
| Authenticated Users | Read (from Security Filtering) |
| Enterprise Domain Controllers | Read |
| Domain Administrators<br><br>Enterprise Administrators<br><br>Creator Owner<br><br>SYSTEM | Edit settings, delete, modify security |

**Note**

- Because Administrators are also part of the Authenticated Users group, they have the Apply Group Policy ACE set to Allow by default; as a result, policy settings apply to them as well if they are located in the container where the GPO is linked.

**Delegating Group Policy-Related Tasks on Sites, Domains, and OUs**

You can manage three Group Policy tasks on a per-container basis in Active Directory:

- Linking GPOs to an Active Directory container (site, domain, or OU)
- Performing Group Policy Modeling analysis for objects in that container (domains and OUs)
- Reading Group Policy Results data for objects in that container (domains and OUs)

By default, Domain Administrators have GPO linking permission for domains and OUs, and Enterprise Administrators and

Domain Administrators of the forest root domain can manage links to sites. You can delegate permissions to additional groups and users by using GPMC.

By default, access to Group Policy Modeling and remote access to Group Policy Results data is restricted to Enterprise Administrators and Domain Administrators. Organizations can delegate access to these tasks to lower-level administrators by setting these permissions in GPMC.

The following procedures detail Group Policy delegation tasks on Active Directory containers.

**To delegate Group Policy administrative tasks on a container**

1. To delegate Group Policy-related permission on a site, domain, or OU, click the appropriate container in the GPMC console.

2. In the right pane for the site, domain, or OU, click the **Delegation** tab.

3. In the drop-down list box, select the desired permission you want to manage: **Link GPOs**, **Perform Group Policy Modeling analyses**, or **Read Group Policy Results data**. Note that GP Modeling and GP Results are not available for sites.

4. To add new groups, use the **Add** button.

5. To modify the **Applies To** setting for an existing permission, right-click the user or group in the list and then select either **This container only** or **This container and all child containers**.

6. To remove an existing group or user from having the specified permission, select the user or group from the list and click the **Remove** button. Only domain administrators have permission to do this.

7. To add or remove custom permissions, click **Advanced** at the bottom-right of the details pane and select the object whose permissions you want to change. Note that setting custom permissions is *not* recommended.

**Delegating Creation of GPOs**

The ability to create GPOs in a domain is a permission that is managed on a per-domain basis. By default, only Domain Administrators, Enterprise Administrators, Group Policy Creator Owners, and SYSTEM can create new Group Policy objects. If the domain administrator wants a non-administrator or non-administrative group to be able to create GPOs, that user or group can be added to the Group Policy Creator Owners security group. Alternatively, you can use the **Delegation** tab on the Group Policy Objects container in GPMC to delegate creation of GPOs. When a non-administrator who is a member of the Group Policy Creator Owners group creates a GPO, that user becomes the creator owner of the GPO and can edit the GPO and modify permissions on the GPO. However, members of the Group Policy Creator Owners group cannot link GPOs to containers unless they have been separately delegated the right to do so on a particular site, domain, or OU. Being a member of the Group Policy Creator Owners group gives the non-administrator full control of only those GPOs that the user creates. Group Policy Creator Owner members do not have permissions for GPOs that they do not create.

**Note**

- When an administrator creates a GPO, the Domain Administrators group becomes the Creator Owner of the Group Policy object. By default, Domain Administrators can edit all GPOs in the domain.

The right to link GPOs is delegated separately from the right to create GPOs and the right to edit GPOs. Be sure to delegate both rights to those groups you want to be able to create and link GPOs. By default, non-Domain Admins cannot manage links, and this prevents them from being able to use GPMC to create and link a GPO. However, non-Domain Admins can create an unlinked GPO if they are members of the **Group Policy Creator Owners** group. After a non-Domain Admin creates an unlinked GPO, the Domain Admin or someone else who has been delegated permissions to link GPOs an a container can link the GPO as appropriate.

Creation of GPOs can be delegated to any group or user. There are two methods of granting a group or user this permission:

- Add the group or user to the Group Policy Creator Owners group. This was the only method available prior to GPMC.

- Explicitly grant the group or user permission to create GPOs. This method is newly available with GPMC.

You can manage this permission by using the **Delegation** tab on the Group Policy objects container for a given domain in GPMC. This tab shows the groups that have permission to create GPOs in the domain, including the Group Policy Creator Owners group. From this tab, you can modify the membership of existing groups that have this permission, or add new groups.

Because the Group Policy Creator Owners group is a domain global group, it cannot contain members from outside the domain. Being able to grant users permissions to create GPOs without using Group Policy Creator Owners facilitates delegating GPO creation to users outside the domain. Without GPMC, this task cannot be delegated to members outside the domain.

If you require that users outside the domain have the ability to create GPOs, create a new domain local group in the domain (for example, "GPCO – External"), grant that group GPO creation permissions in the domain, and then add domain global groups from external domains to that group. For users and groups in the domain, you should continue to use the Group Policy Creator Owners group to grant GPO-creation permissions.

Adding a user to the membership of Group Policy Creator Owners and granting the user GPO-creation permissions directly using the new method available in GPMC are identical in terms of permissions.

**Delegating creation of WMI filters**

WMI filters are a new feature in Windows Server 2003 and Windows XP. WMI filters are created in the **WMI Filters** container in GPMC. There are two levels of permission for creating WMI filters:

- **Creator Owner**: Allows the user to create new WMI filters in the domain, but does not grant permissions on WMI filters created by other users.

- **Full Control**: Allows the user to create WMI filters, and grants full control on all WMI filters in the domain, including new filters created after users are granted this permission.

To delegate these permissions, use the **Add** button on the **Delegation** tab of the **WMI Filters** pane.

An administrator can **Add**, **Remove**, and view **Properties** for WMI Filter delegations from the **Delegation** tab. Selecting **Add** prompts for a user or group before selecting the permission level (**Creator Owner** or **Full Control**) to assign to the user or group. Selecting **Remove** prompts for confirmation that the delegation should be removed. Selecting **Properties** displays the user or group properties for that object.

**Delegating permissions on individual WMI filters**

GPMC allows you to delegate permissions on individual WMI filters. There are two levels of permissions that can be granted to a user or group on an individual WMI filter:

- **Edit**: Allows the user or group to edit the selected WMI filter.

- **Full Control**: Allows the user or group to edit, delete, and modify security on the selected WMI filter.

These permissions are managed by using the **Delegation** tab of a WMI filter

The **Delegation** tab shows the users and groups that have permissions on the WMI filter, their permission levels, and whether the permission is inherited from a parent container. Buttons on this tab let you add users and groups to the delegation list for the WMI filter, or remove them from this list.

Note that all users have **Read** access to all WMI filters. GPMC does not allow this permission to be removed. If the **Read** permission were removed, this can cause Group Policy processing on the destination computer to fail.

# Defining Group Policy Operational Procedures

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To facilitate future management of Group Policy, you should develop operational procedures to ensure that changes to GPOs are made in an authorized and controlled manner. In particular, make sure that all new GPOs and changes in existing GPOs are properly staged before deployment to your production environment. You should also create regular backups of your GPOs.

In some organizations, different teams might be responsible for managing different aspects of Group Policy. For example, a software deployment team is typically concerned with the **Software Installation (Users)** and **Software Installation (Computers)** sections of Group Policy. The remaining settings, relating to items such as scripts and Folder Redirection, are unlikely to be of interest to this team.

To reduce complexity and minimize the likelihood of introducing errors, consider creating separate GPOs for different groups of administrators.

Alternatively, you might restrict administrators to the parts of Group Policy they are authorized to change. You can use the **Restricted/Permitted Snap-ins\Extension snap-ins** policy setting to restrict the snap-ins that administrators can access. This setting is available in the Group Policy Object Editor snap-in under **User Configuration\Administrative Templates\Windows Components\Microsoft Management Console**. The **Restricted/Permitted Snap-ins\Extension snap-ins** policy pertains to the UI that is accessible by using the Group Policy Object Editor snap-in. Remember that some teams may need access to more than one type of extension snap-in. To start the Group Policy Object Editor, right-click a GPO in GPMC and then click **Edit**.

**Note**

- The MMC policy settings only affect the UI that is accessible by using MMC; if Group Policy is edited by using a programmatic means, any GPO settings can be edited, including important Sysvol aspects.

For more information about these and other Group Policy settings, double-click the setting in the details pane of the Group Policy Object Editor, and then click the **Explain** tab in the policy **Properties** dialog box. Note that this information is always available by simply clicking the setting if you have **Extended View** enabled, which is the default behavior.

# Specifying a Domain Controller for Editing Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In each domain, GPMC uses the same domain controller for all operations in that domain. This includes all operations on the GPOs that are located in that domain, as well as all other objects in that domain, such as OUs and security groups. When you open the Group Policy Object Editor from GPMC, it uses the same domain controller.

GPMC also uses the same domain controller for all operations on sites. This domain controller is used to read and write information about what links to GPOs exist on any given site, but information regarding the GPOs themselves is obtained from the domain controllers of the domains that host the GPOs.

By default, when you add a new domain to the console, GPMC uses the PDC emulator in that domain for operations in that domain. For managing sites, GPMC uses the PDC emulator in the user's domain by default.

The choice of domain controllers is important for administrators to consider to avoid replication conflicts. This is especially important because GPO data is located both in Active Directory and in Sysvol, which rely on independent replication mechanisms to replicate GPO data to the various domain controllers in the domain. If two administrators simultaneously edit the same GPO on different domain controllers, it is possible for the changes written by one administrator to be overwritten by another administrator, depending on replication latency.

To avoid this, GPMC uses the Operations Master token for the PDC emulator in each domain as the default. This helps ensure that all administrators are using the same domain controller and guards against data loss. However, it might not always be desirable for an administrator to use the PDC to edit GPOs. For example, if the administrator is located in a remote site, or if the majority of the users or computers targeted by the GPO are in a remote site, the administrator might choose to target a domain controller at the remote location. For example, if you are an administrator in Japan and the PDC emulator is in New York, it might be inconvenient to rely on a WAN link to access the New York PDC emulator.

**Important**

- If multiple administrators manage a common GPO, all administrators should use the same domain controller when editing a particular GPO in order to avoid collisions in the File Replication service (FRS).

Use the **Change Domain Controller** function to specify the domain controller to be used for a given domain or for all sites in a forest. In each case, you have four options:

- **The domain controller with the Operations Master token for the PDC emulator** (the default option)
- **Any available domain controller**
- **Any available domain controller running Windows Server 2003 or later**.
- **This domain controller:** Select a specific domain controller to be used.

The selected option is used each time that you open a saved console, until you change the option.

This preference is saved in the .msc file and is used when you open that .msc file. It is generally not recommended that you use the **Any available domain controller** option unless you are performing read-only operations.

# Group Policy and Slow Links

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some Group Policy is not applied when the connection speed fall below specified thresholds. Therefore, when your Group Policy solution calls for applying policy over slow links or by using remote access, you need to consider policy settings for slow link detection.

Although slow links and remote access are related, there are distinctions, and Group Policy processing varies for each. Having a computer connected to a LAN does not necessarily imply a fast link, nor does a remote access connection imply a slow link. The default value for what Group Policy considers a slow link is any rate slower than 500 Kilobits per second (Kbps). You can change this threshold using Group Policy. Group policy uses the following algorithm to measure link speed.

**How Group Policy measures link speed**

1. Ping the server with 0 bytes of data and time the number of milliseconds. This value is time#1. If it is less than 10 ms, exit (assume a fast link).

2. Ping the server with 2 kilobytes (KB) of uncompressible data, and time the number of milliseconds. This value is time#2. The algorithm uses a compressed .jpg file to ping the server.

3. DELTA = time#2 - time#1. This removes the overhead of session setup, with the result being equal to the time to move 2 KB of data.

4. Calculate Delta three times, adding to TOTAL each DELTA value. Use the following calculations:

5. TOTAL/3 = Average of DELTA in milliseconds.

6. 2 * (2 KB) * (1000 ms/sec) / DELTA Average ms = X

7. X = (4000 KB/sec) / DELTA Average

8. Z Kbps = ((4000 KB) / DELTA Average) *(8 bits/byte)

9. Z Kbps = 32000 kbps/Delta Average.

   Two KB of data have moved in each direction (this is represented by the leading factor two on the left side in step six) through each modem, Ethernet card, or other device in the loop once. The resulting Z value is evaluated against the policy setting. A default of less than 500 Kbps is considered a slow link; faster than 500 Kbps is a fast link.

Administrators can use a policy setting to define a slow link for the purpose of applying Group Policy, as described in the following sections.

# Specifying Group Policy for Slow Link Detection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Administrators can partially control which Group Policy extensions are processed over a slow link. By default, when processing over a slow link, not all components of Group Policy are processed.

Table 2.6 shows the default settings for processing Group Policy over slow links.

**Table 2.6 Default Settings for Processing Group Policy over Slow Links**

| Setting | Default |
|---|---|
| Security Settings | ON (cannot be turned off) |
| IP Security | ON |
| EFS | ON |
| Software Restriction Policies | ON |
| Wireless | ON |
| Administrative Templates | ON (cannot be turned off) |
| Software Installation | OFF |
| Scripts | OFF |
| Folder Redirection | OFF |
| IE maintenance | ON |

Administrators can use a Group Policy setting to define a slow link for the purposes of applying and updating Group Policy. The default value defines a rate slower than 500 Kbps as a slow link.

To specify settings for Group Policy slow link detection for computers, use the **Group Policy slow link detection policy** setting in the **Computer Configuration\Administrative Templates\System\Group Policy** item of the Group Policy Object Editor. The unit of measurement for connection speed is Kbps.

To set this for users, use the **Group Policy slow link detection policy** setting in **User Configuration\Administrative Templates\System\Group Policy**.

For User Profiles, the **Slow network connection timeout for user profiles** setting is located in the **Computer Configuration\Administrative Templates\System\User Profiles** item. This setting allows users to ping the server and to check the performance of the file system. This is because user profiles can be stored anywhere, and the server might or might not have IP support. Therefore, the user profile first tries to ping the server. If the server does not have IP support, it falls back to measuring the performance of the file system. You must specify connection speeds in both Kbps and milliseconds when you set this policy.

**Notes**

- If the **Do not detect slow network connections** policy setting is enabled, the **Slow network connection timeout for user profiles** policy setting is ignored.

- If **Delete cached copies of roaming profiles** is enabled, there is no local copy of the roaming profile to load when the system detects a slow connection.

**Setting Computer Policy for Client-Side Extensions for Processing Across Slow Links**
Group Policy is implemented almost entirely as a series of client-side extensions, such as security, administrative templates, and folder redirection. There is a computer policy that allows configuring slow-link behavior for each client-side extension. You can use these policy settings to specify the behavior of client-side extensions when processing Group Policy. There is a maximum of three options for each policy setting. The **Allow processing across a slow network connection** policy option controls processing policy settings across slow links. The other two options can be used to specify that policy should not be processed in the background, or that policy be updated and reapplied even if policy settings have not changed. For more information about policy for client-side extensions, see "Specifying Group Policy for Slow Link Detection" earlier in this chapter.

Some extensions move large amounts of data, so processing across a slow link can affect performance. By default, only the administrative templates and security-related settings are processed over a slow link.

You can set the options for processing the following:

- IP Security policy
- EFS recovery policy
- Internet Explorer Maintenance policy
- Scripts policy
- Folder Redirection policy

**Group Policy and Remote Access Connections**
Processing of Group Policy over a remote access connection differs from processing over a slow link. Group Policy is applied during a remote access connection as follows:

- When users click to select the **Logon using dial-up connection** check box at the logon prompt, both user and computer Group Policy are applied if the computer is a member of the domain that the remote access server belongs to or trusts. However, computer-based software installation settings are not processed, nor are computer-based startup scripts executed, because computer policy is normally processed before the logon screen appears. However, in the case of using a dial-up connection, the application of computer policy is done as a background refresh during the logon process.
- When the processing of cached credentials is completed and a remote access connection is established, Group Policy is not applied, except during a background refresh.

Group Policy is not applied to computers that are members of a workgroup, because computer policy is never applied to computers that are in a workgroup.

# Controlling Client-Side Extensions by Using Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some of the Group Policy components include client-side extensions (typically implemented as .dll files) that are responsible for processing and applying Group Policy settings at the destination computers.

For each client-side extension, the GPO processing order is obtained from a list of GPOs, which is determined by the Group Policy engine during processing. Each client-side extension processes the resulting list of GPOs. For details about the processing of Group Policy, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

A computer policy exists to control the behavior of each of the Group Policy client-side extensions. Each policy includes up to three options, and some include more specific configuration options. You can set computer policies for client-side extensions by opening the **Computer Configuration/Administrative Templates/System/Group Policy** item of the Group Policy Object Editor.

You can set the following computer policy options:

- **Allow processing across a slow network connection**. Some extensions move large amounts of data, so processing across a slow link can decrease performance. By default, only the administrative templates and security settings are processed over a slow link. You can set this policy to mandate that other client-side extensions are also processed across a slow link. To control what is considered a slow link, use the **Group Policy slow link detection** setting. See "Specifying Group Policy for Slow Link Detection" earlier in this chapter.

- **Do not apply during periodic background processing**. Computer policy is applied at boot time and again every 90 minutes. User policy is applied when the user logs on to the computer and in the background approximately every 90 minutes after that. The **Do not apply during periodic background processing** option gives you the ability to override this behavior and prevent Group Policy from running in the background.

  **Note**

  - The Software Installation and Folder Redirection extensions process Group Policy only at startup and when the user logs on to the network because of the risks in processing these policies in the background, when users might have applications and files open

- **Process even if the Group Policy objects have not changed**. If the GPOs on the server do not change, it is not usually necessary to continually reapply them to the destination computer except to override possible local changes. Because local administrators might be able to modify the parts of the registry where Group Policy settings are stored, you might want to reapply these settings as needed during the logon process or during periodic background processing to return the computer to the desired state.

  For example, assume that Group Policy defines a specific set of security options for a file. Then a user who has administrative credentials logs on and changes those security options. The Group Policy administrator might want to set the policy option to process Group Policy even if the GPOs have not changed so that the security options specified in Group Policy are reapplied the next time policy is refreshed. The same considerations apply to applications: with this option set, if Group Policy installs an application, but the user removes the application or deletes its icon, the application is re-advertised the next time the user logs on to the computer.

By default, Security Policy settings delivered by Group Policy are applied every 16 hours (960 minutes) even if a GPO has not changed. It is possible to change this default period by using the registry entry **MaxNoGPOListChangesInterval** in the following subkey:

HKLM\Software\Microsoft\Windows NT\CurrentVersion\Winlogon \GPExtentions\{82...},

The data type of this entry is REG_DWORD and the value is number of minutes.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing

settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

# Minimizing the Growth of DC System Volume Folders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The policy settings information in GPOs is stored in two locations: in Active Directory and in the Sysvol folder of domain controllers. The Active Directory container is also known as a *Group Policy container*, and the Sysvol folder contains the *Group Policy template*. The Group Policy container contains attributes that are used to deploy GPOs to the domain, to OUs, and to sites within the domain. The Group Policy container also contains a link to the Group Policy template, where most Group Policy settings are stored.

Information stored in the Sysvol folder includes security settings, script files, information for deploying applications that are available for Group Policy Software Installation, and Administrative Template–based Group Policy settings. Administrative Templates (.adm files) provide Group-Policy setting information for the items that appear under the Administrative Templates item in the Group Policy Object Editor.

Limiting the size of stored Group Policy information is important largely because the two storage locations use different replication mechanisms, which can cause replication conflicts. Replication conflicts are discussed in "Specifying a Domain Controller for Editing Group Policy," earlier in this chapter.

By default, when you open the Group Policy Object Editor snap-in to edit a new or existing GPO (or when you click **Edit** after right-clicking a GPO in GPMC), the administrative computer's .adm files for that GPO are sent to the domain controller and placed in a folder associated with the new GPO. The .adm files stored are the ones that are included with the operating system for the administrative computer (such as System.adm, Inetres.adm, Conf.adm, and Wmplayer.adm), as well as any .adm files that are added or removed by the administrator by using the **Add/Remove Templates** shortcut menu.

When an existing GPO is edited, a comparison is made of the timestamps of the local .adm files and the versions in the Sysvol folder. If the local timestamp is more recent, those files are sent to the domain controller where they are stored in the Sysvol GPO folder.

#### Benefits of Storing .Adm Files in the Sysvol Folder

Storing copies of .adm files in the Sysvol folder provides a consistent experience for administrators across computers. For example, if a specific set of .adm files is selected when you edit a GPO using the Group Policy Object Editor on one computer (by using **Add/Remove Templates**), these same .adm files are also included when the GPO is edited on another computer. The selected .adm files need not be installed on the second computer, because the files are retrieved from the domain controller. The list of selected .adm files is stored in the Sysvol GPO folder in the **Admfiles.ini** file.

#### Drawbacks of Storing .Adm Files in the Sysvol Folder

Each GPO stores its own copy of the .adm files selected for it. Because some .adm files can be large — more than 1Mb— it is possible that the Sysvol GPO folder can grow large, especially when you have a large number of GPOs.

An additional problem occurs if administrators use different language versions of the operating system. In this case, the .adm with the newest timestamp is uploaded to the Sysvol folder on the domain controller, regardless of language. This is important because non-U.S. English .adm files typically have a later timestamp than the U.S. English versions. This occurs because simply opening **Administrative Templates** in the Group Policy Object Editor on a non-U.S. English administrative computer results in the Group Policy templates being updated with the local language version of the .adm files, even if you do not actually edit any settings.

#### Group Policy Settings to Control Treatment of .Adm Files

To provide administrators with control how .adm files are treated, you can use the following policy settings: **Always use local ADM files for Group Policy Editor** and **Turn off automatic updates of ADM files**.

#### Always use local ADM files for Group Policy Object Editor

This policy setting forces the Group Policy Object Editor to use the local computer .adm files in the *systemroot*\inf directory. To access this Group Policy setting, in the Group Policy Object Editor snap-in, navigate to **Computer Configuration\Administrative Templates\System\Group Policy**. This setting is not available in the Windows 2000 family.

This setting does not alter the default treatment of .adm files when you edit a newly-created GPO: The local .adm files are always used and are copied to the Sysvol GPO folder. When you edit an existing GPO, any .adm files stored in the GPO folder on the Sysvol are ignored. When this setting is used, the Group Policy Object Editor always uses only the .adm files from the local computer (including any custom .adm files). If an .adm file is not found, the corresponding Group Policy settings are not available. This Group Policy setting can be used to ensure that language-specific versions of the .adm files on the administrative computer are used, rather than any other language version stored in the GPO folder.

**Turn off automatic update of .Adm files**

Enabling this policy setting prevents the system from automatically updating the .adm files in the Sysvol folder when you open the Group Policy Object Editor. This policy setting is available in the Group Policy Object Editor snap-in under **User Configuration\Administrative Templates\System\Group Policy**.

When this Group Policy setting is enabled and you edit a new GPO, no .adm files are copied from the local computer to the Sysvol GPO folder on the domain controller. Similarly, when you edit an existing GPO, no local .adm files are sent to the domain controller. This means that the Group Policy Object Editor presents only the .adm files found on the local computer, and no .adm files are copied to the domain controller.

# Changing the Group Policy Refresh Interval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Administrators can change the default refresh policy interval setting by using one of these policy settings: **Group Policy Refresh Interval for Computers,Group Policy Refresh Interval for Domain Controllers**, or **Group Policy refresh Interval for Users.** By using these settings, you can stipulate an update rate from zero to 64,800 minutes (45 days). You can also set the policy to **Turn off background refresh of Group Policy**.

**Important**

- When you set the refresh interval to 0 minutes, the computer tries to update Group Policy every seven seconds. Because such updates might interfere with users' work and increase network traffic, very short update intervals are appropriate only in test environments.

**Group Policy Refresh Interval for Computers**
This setting specifies how often Group Policy for computers is updated in the background. It specifies a background update rate only for Group Policy settings under Computer Configuration. Computer Group Policy is updated in the background every 90 minutes by default, with a random offset of 0 to 30 minutes. In addition to background updates, computer Group Policy is always updated when the system starts. This policy setting is available in the Group Policy Object Editor under **Computer Configuration\Administrative Templates\System\Group Policy**.

**Group Policy Refresh Interval for Domain Controllers**
This setting specifies how often Group Policy is updated in the background on domain controllers. By default, Group Policy on domain controllers is updated every five minutes. This policy setting is available in the Group Policy Object Editor under **Computer Configuration\Administrative Templates\System\Group Policy**.

**Group Policy Refresh Interval for Users**
This setting specifies how frequently Group Policy is updated in the background only for the Group Policy settings in the User Configuration folder. In addition to background updates, Group Policy for users is always updated when users log on. This policy is in the **User Configuration\Administrative Templates\System\Group Policy** item.

**Turn off background refresh of Group Policy**
This policy prevents Group Policy settings from being updated while the computer is in use. It applies to Group Policy for computers, users, and domain controllers. This policy setting is available in **Computer Configuration\Administrative Templates\System\Group Policy** item.

**Running Command Line Options to Refresh Policy**
From a given computer, you can refresh the settings that are deployed to that computer by using the Gpupdate.exe tool. One reason this is useful is to ensure that security settings are enforced. Table 2.7 describes parameters for Gpupdate.exe. The Gpupdate.exe tool is used in Windows Server 2003 and Windows XP environments and replaces the /refreshpolicy option previously used with the secedit command. For Windows 2000, continue to use the secedit /refreshpolicy command.

The Gpudate.exe tool uses the following syntax:

```
gpupdate [/target:{computer|user}] [/force] [/wait:value] [/logoff] [/boot]
```

**Table 2.7 Gpudate.exe Parameters**

| Parameter | Description |
|---|---|
| **/target**: {computer\|user} | Depending on what target you specify, Gpupdate.exe processes the computer settings, the current user settings, or both. By default, both the computer and the user settings are processed. |
| **/force** | Reapplies all settings and ignores processing optimizations. |
| **/wait:** *value* | Specifies the number of seconds that policy processing waits to finish. The default is 600 seconds. A value of 0 means no wait; -1 means wait indefinitely. |

| | |
|---|---|
| **/logoff** | Logs off after the policy refresh completes. This is required for Group Policy client-side extensions that do not process on a background refresh cycle but do process when the user logs on, such as user Software Installation and Folder Redirection. This option has no effect if there are no extensions called that require the user to log off. |
| **/boot** | Restarts the computer after the policy refresh completes. This is required for those Group Policy client-side extensions that do not process on a background refresh cycle but do process when the computer starts up, such as computer Software Installation, for example. This option has no effect if there are no extensions called that require the computer to be restarted. |
| **/?** | Displays Help at the command prompt. |

# Using Group Policy Modeling and Group Policy Results to Evaluate Group Policy Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before deploying your Group Policy solution, it is critical that you assess it to determine the effects of applying the various policy settings that you select, individually and in combination. The primary mechanism for assessing your Group Policy deployment is to create a staging environment and log on using a test account. This is the best way to understand the impact and interaction of all the applied GPO settings. Staging your Group Policy deployment is critical for creating a successful managed environment. For more information, see "Staging Group Policy Deployments" in this book.

For Active Directory networks with at least one Windows Server 2003 domain controller, you can use Group Policy Modeling in GPMC to simulate the deployment of GPOs to any destination computer running Windows 2000 Server or Professional, Windows XP Professional, or Windows Server 2003. The primary tool for viewing the actual application of GPOs is by using Group Policy Results in GPMC.

Group Policy Modeling was previously called Resultant Set of Policy (RSoP) planning mode, and Group Policy Results was previously called RSoP logging mode.

**Using Group Policy Modeling to Simulate Resultant Set of Policy**

The built-in Group Policy Modeling Wizard calculates the simulated net effect of GPOs. Group Policy Modeling can also simulate such things as security group membership, WMI filter evaluation, and the effects of moving user or computer objects to a different Active Directory container. The simulation is performed by a service that runs on domain controllers running Windows Server 2003. These calculated settings are reported in HTML and are displayed in GPMC on the **Settings** tab in the details pane for the selected GPO. To expand and contract the settings under each item, click **hide** or **show all** so that you can see all the settings, or only a few. To perform a Group Policy Modeling analysis, you must have at least one domain controller running Windows Server 2003, and you must have the **Perform Group Policy Modeling analyses** permission on the domain or organizational unit that contains the objects on which you want to run the query.

To run the wizard, right-click **Group Policy Modeling** (or an Active Directory container), and then click **Group Policy Modeling Wizard.** If you run it from an Active Directory container, the wizard fills in the **Container** fields for user and computer with the LDAP distinguished name of that container.

When you have answered all the questions in the wizard, your answers are displayed as if they were from a single GPO. They are also saved as a query represented by a new item under the **Group Policy Modeling** item. The display also shows which GPO is responsible for each setting, under the heading **Winning GPO**. You can also see more detailed precedence information (for example, which GPOs attempted to set the settings, but did not succeed). To do so, right-click the item, and then click **Advanced View**. This starts the traditional RSoP snap-in. Each setting has a **Precedence** tab.

Keep in mind that modeling does not include evaluating any LGPOs. Because of this, in some cases you might see a difference between the simulation and the actual results.

To save the results of the modeling, right-click the query, and then click **Save Report**.

**Using Group Policy Results to determine Resultant Set of Policy**

Use the Group Policy Results Wizard to see what Group Policy settings are actually in effect for a user or computer by gathering RSoP data from the destination computer. In contrast to Group Policy Modeling, Group Policy Results reveals the actual Group Policy settings that were applied to the destination computer. The target must be running Windows XP Professional or later.

The settings are reported in HTML and are displayed in a GPMC browser window on the **Summary** and **Settings** tabs in the details pane for the selected GPO. You can expand and contract the settings under each item by clicking **hide** or **show all** so that you can see all the settings, or only a few. To remotely access Group Policy Results data for a user or computer, you must have the **Remotely access Group Policy Results data** permission on the domain or organizational unit that contains the user or computer, or you must be a member of a local Administrator's group on the appropriate computer and must have network connectivity to the destination computer. To delegate Group Policy Results, you need the Windows 2003 Server schema in your Active Directory. To update your schema, run **ADPrep /forestprep** on the domain controller with the schema operations master role.

To run the wizard, right-click the **Group Policy Results** container**,** and then click **Group Policy Results Wizard**.

When you have answered all the questions in the wizard, GPMC creates a report that shows the resultant set of policy for the user and computer you entered in the wizard. The display shows which GPO is responsible for each setting on the **Settings** tab, under the heading **Winning GPO**.

You can save the results by right-clicking the query and choosing **Save Report**.

**Using Gpresult.exe to Evaluate Policy Settings**

You can run Gpresult.exe on the command line of any remote computer within the scope of your management to get the same data you can get by using GPMC Group Policy Results. By default, Gpresult.exe returns settings in effect on the computer on which it runs.

For Windows Server 2003 and Windows XP Professional, Gpresult.exe uses the following syntax:

```
gpresult [/scomputer [/udomain\user/ppassword]] [/userTargetUserName] [/scope {user|computer}] [/v] [/z]
```

Table 2.8 describes the parameters for Gpresult.exe.

**Table 2.8 Gpresult.exe Parameters**

| Parameter | Description |
|---|---|
| **/s** *computer* | Specifies the name or IP address of a remote computer. (Do not use backslashes.) The default is the local computer. |
| **/u** *domain\user* | Runs the command using the account permissions of the user that is specified by User or Domain\User. The default is the permissions of the current logged-on user on the computer that issues the command. |
| **/p** *password* | Specifies the password of the user account that is specified in the /u parameter. |
| **/user** *TargetUserName* | Specifies the user name of the user whose RSoP data is to be displayed. |
| **/scope** {user\|computer**}** | Displays either user or computer results. Valid values for the /scope parameter are user or computer. If you omit the /scope parameter, gpresult displays both user and computer settings. |
| **/v** | Specifies that the output display verbose policy information. |
| **/z** | Specifies that the output display all available information about Group Policy. Because this parameter produces more information than the /v parameter, redirect output to a text file when you use this parameter (for example, gpresult /z >policy.txt). |
| **/?** | Displays help at the command prompt. |

**To run Gpresult.exe on your computer**

1. Click **Start**, click **Run**, type **cmd**, and then press ENTER.
2. Type **`gpresult /z>gp.txt`** to write the output to a text file that is named Gp.txt.
3. Type **`notepad gp.txt`**to open the text file.

Note that using Group Policy Results is recommended over using Gpresult.exe. GPResult.exe is primarily useful if you desire command-line *presentation* of resultant set of policy data. If you require command-line access to RSoP data but not command-line presentation of the data, you can simply run a script based on the GPMC interfaces to generate an RSOP report in the form of an .htm file.

# Backing up, Restoring, Migrating, and Copying GPOs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

GPMC provides mechanisms for backing up, restoring, migrating, and copying existing GPOs. This is very important for maintaining your Group Policy deployments in event of error or disaster. It helps you avoid having to manually recreate lost or damaged GPOs and then go through the planning, testing, and deployment phases again. Part of your ongoing Group Policy operations plan should include regular backups of all GPOs. Inform all Group Policy administrators about how to use GPMC to restore GPOs.

GPMC also provides for copying and importing GPOs, both from the same domain and across domains. You can use GPMC to migrate an existing GPO, for example, from an existing domain into a newly deployed domain. You can either copy GPOs or import settings from one GPO into another GPO. This can save you a lot of time and trouble by enabling you to re-use the contents of existing GPOs. Copying GPOs enables you to move straight from the staging phase to production, provided you have the proper trust between the environments. Import allows you to transfer settings from a backed-up GPO into an existing GPO, and is especially useful in situations where there is no trust between the source and destination domains. If you want to reuse existing GPOs, copying also allows you to conveniently move GPOs from one production environment to another.

**Using GPMC to Work with GPOs**

To create GPO backups, you must have at least **Read** access to the GPOs and **Write** access to the folder in which the backups are stored. See Figure 2.10 to help you identify the items referred to in the procedures that follow.

**Figure 2.10 Backup Options in the Group Policy Management Console User Interface**




**Using GPMC to Back up GPOs**

The backup operation backs up a live GPO to the file system. The location of the backup can be any folder to which you have write access. After backing-up GPOs, you must use GPMC to display and manipulate the contents of your backup folder, either by using the GPMC UI or programmatically by using a script. Do not interact with backed-up GPOs directly through the file system. Once backed up, archived GPOs can be processed by the **Import** and **Restore** operations.

Note that you can back up multiple instances of the same GPO to the same location — GPMC uniquely identifies each backup instance and provides mechanisms to allow you to pick which instance of the archived GPO you want to work with. For example, you can choose to display only the most recent backups when viewing the contents of a backup folder through GPMC. This can be useful when you make backups of a GPO after changing it, and later need to restore a previous version of that GPO.

**To back up all GPOs in a domain**

1. In the GPMC console tree, expand the forest or domain that contains the GPOs you want to back up.

2. Right-click the **Group Policy Objects** container.

3. In the context menu, click **Back Up All**.

4. In the **Backup Group Policy Object** dialog box, enter the path to the location at which you want to store the GPO backups. Alternatively, you can click **Browse**, locate the folder in which you want to store the GPO backups, and then click **OK**.

5. Type a description for the GPOs that you want to back up, and then click **Backup**.

6. After the operation completes, click **OK**.

#### To back up a specific GPO

1. In the **Group Policy Objects** container, right-click the GPO you want to back up.

2. In the context menu, click **Back Up**.

3. In the **Backup Group Policy Object** dialog box, enter the path to the location at which you want to store the GPO backup. Alternatively, you can click **Browse**, locate the folder in which you want to store the GPO backup, and then click **OK**.

4. Type a description for the GPO that you want to back up, and then click **Backup**.

5. After the operation completes, click **OK**.

#### To manage GPO backups

1. In the GPMC console tree, expand the domain that contains the GPOs that you want to back up.

2. Right-click **Group Policy Objects** container.

3. In the **Manage Backups** dialog box, click **Browse**, locate the folder that contains your GPO backups, and then click **OK**.

### Important

- You should secure backed-up GPOs by ensuring that only authorized administrators have permission to access the folder to which you are saving GPOs. Use good security on the file system where they are backed up.

#### Using GPMC to Restore GPOs

You can also restore GPOs. This operation takes a backed-up GPO and restores it to the same domain from which it was backed up. You cannot restore a GPO from backup into a domain different from the GPO's original domain.

#### To restore a previous version of an existing GPO

1. In the GPMC console tree, expand the forest or domain that contains the GPOs you want to restore.

2. Expand the **Group Policy Objects** container, right-click the GPO you want to restore to a previous version, and then click **Restore from Backup**.

3. When the Restore Group Policy Object Wizard appears, follow the instructions and provide the appropriate information about the backed-up GPO that you want to restore, and then click **Finish**.

4. After the Restore Group Policy Object Wizard completes the restore operation, click **OK**.

#### To restore a deleted GPO

1. Right-click the **Group Policy Objects** container and then select the **Manage Backups** option from the context menu.

2. In the **Manage Backups** dialog box, click **Browse,** and then locate the file system that contains your backed-up GPOs.

3. Select the GPO that you want to restore, and then click **Restore**.

4. When you are prompted to confirm the restore operation, click **OK**.

### Note

- Settings that are stored outside the GPOs, such as WMI filters and IPSec policies are not backed up or restored during these processes. However, links to WMI filters and IPSec are stored in the GPO, and these are backed up as part of the GPO. During restore, the links are preserved if the underlying object still exists in Active Directory.

**Using GPMC to Copy GPOs and Import GPO Settings**

GPMC allows you to copy GPOs, both in the same domain and across domains, as well as import Group Policy settings from one GPO to another. Perform these operations as part of your staging process prior to deployment in your production environment. These operations are also useful for migrating GPOs from one production environment to another.

Although the collection of settings which comprises a GPO is logically a single entity, the data for a single GPO is stored in multiple locations and in a variety of formats; some data is contained in Active Directory and other data is stored in the Sysvol folder on domain controllers. This means that you cannot simply copy GPOs by copying a folder from one computer to another. However, GPMC provides built-in support that allows you to do this safely and relatively simply.

A copy operation copies an existing, live GPO to the desired destination domain. A new GPO is always created as part of this process. The destination domain can be any trusted domain in which you have the rights to create new GPOs. Simply add the desired forests and domains in GPMC and use GPMC to copy and paste (or drag and drop) the desired GPOs from one domain to another. To copy a GPO, you must have permission to create GPOs in the destination domain.

An additional option available when copying GPOs is to copy the Discretionary Access Control List (DACL) on the GPO in addition to the settings within the GPO. This is useful for ensuring that the new GPO that is created as part of the copy operation has the same security filtering and delegation options as the original GPO.

Importing a GPO allows you to transfer settings from a backed-up GPO to an existing GPO. Importing a GPO transfers only the GPO settings; it does not modify the existing security or links on the destination GPO. Importing a GPO is useful for migrating GPOs across untrusted environments, because you only need access to the backed-up GPO, not the actual GPO. Because an import operation only modifies settings, Edit permissions on the destination GPO are sufficient to perform the operation.

When copying or importing a GPO, you have the option of specifying a migration table if the GPO contains security principals or UNC paths that might need to be updated to new values in the target domain. Use the Migration Table Editor to create and edit migration tables; migration tables are described in the next section, "Using GPMC to Migrate GPOs."

**To copy a GPO**

1. In the GPMC console tree, right-click the GPO that you want to copy, and then click **Copy**.

    - To place the copy of the GPO in the same domain as the source GPO, right-click the **Group Policy Objects** container, and then click **Paste**.

    - To place the copy of the GPO in a different domain (either in the same or a different forest), expand the destination domain, right-click the **Group Policy Objects** container, and then click **Paste**.

2. If you are copying within a domain, click **Use the default DACL for new GPOs** or **Preserve the existing DACL**, and then click **OK**.

3. If you are copying to or from another domain, answer all the questions in the cross-domain copying wizard that appears, and then click **Finish**.

**To import settings from a backed up GPO into a GPO**

1. In the GPMC console tree, expand the domain that contains the GPO into which you want to import settings.

2. Expand the **Group Policy Objects** container, right-click the GPO, and then click **Import Settings**.

3. When the Import Settings Wizard appears, follow the instructions and provide the appropriate information about the backed up GPO that contains the settings you want to import, and then click **Finish**.

**Using Migration Tables**

Because some data in a GPO is domain-specific and might not be valid if copied directly to another domain, GPMC introduces *migration tables*. Migration tables are stored with the file name extension .migtable, and are actually XML files.

A migration table is a file that maps references to users, groups, computers, and UNC paths in the source GPO to new values in the destination GPO. The migration table consists of one or more mapping entries. Each mapping entry consists of a source type, source reference, and destination reference. If you specify a migration table when performing an import or copy, each reference to the source entry is replaced with the destination entry when the settings are written into the destination GPO. To use a migration table, the destination references specified in the migration table must already exist.

To examine a sample migration table, see **program files\GPMC\Scripts\SampleMigrationTable.migtable** on a computer where GPMC is installed. You do not need to know XML to create or edit migration tables; GPMC provides the Migration Table Editor for manipulating migration tables.

You can use migration tables to update security principals and UNC paths to new values as part of the import or copy operation.

The following items can contain security principals and can be modified using a migration table.

- Security policy settings of the following types:
    - User rights assignment
    - Restricted groups
    - System services
    - File system
    - Registry
- Advanced folder redirection settings
- The GPO DACL, if it is preserved during a copy operation
- The DACL on software installation objects, which is only preserved if the option to copy the GPO DACL is specified

Also, the following items can contain UNC paths, which might need to be updated to new values as part of the import or copy operation, because servers in the original domain might not be accessible from the domain to which the GPO is being migrated:

- Folder redirection Group Policy settings
- Software installation Group Policy settings
- References to scripts (such as for logon and startup scripts) that are stored outside the GPO. The script itself is not copied as part of the GPO copy or import operation, unless the script is stored inside the source GPO.

A migration table is a simple table that specifies a mapping between a source value and a destination value. Figure 2.11 shows a migration table in the Migration Table Editor in GPMC.

**Figure 2.11 Migration Table Editor**




The purpose of this table is to perform, during the copy or import operation, conversion of the references in a GPO to new references that will work in the target domain.

For more information about using migration tables, see "Staging Group Policy Deployments," in this book, and see the white papers available from the Administering Group Policy with GPMC link and the Migrating GPOs Across Domains by Using GPMC link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Maintaining Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After deployment, your Group Policy implementation might need routine maintenance and modification as your organization and its needs evolve and as your experience with Group Policy grows. By establishing control procedures for creating, linking, editing, importing settings into, backing up, and restoring Group Policy objects, you can minimize help desk and support calls arising from poorly planned Group Policy deployments. You can also simplify troubleshooting GPOs and help lower the total cost of ownership for computers in your network.

By establishing GPO control mechanisms, you can create GPOs that:

- Conform to corporate standards.
- Ensure that policy settings do not conflict with those set by others.

Figure 2.12 illustrates the place of this step in the process of designing a Group Policy infrastructure.

**Figure 2.12 Maintaining Group Policy**




To assist with troubleshooting GPOs, you can use the GPMC Group Policy Results Wizard to identify possible Group Policy deployment errors. For more information about this tool, see "Using Group Policy Modeling and Group Policy Results to Evaluate Group Policy Settings" earlier in this chapter. You can also use the GPMC Group Policy Modeling Wizard to evaluate the consequences of new Group Policy settings prior to deploying them to your production environment.

Whenever you deploy new technology solutions, such as wireless networking, you need to revisit your Group Policy configurations to ensure compatibility with new technology. To help manage various technologies, Group Policy offers settings such as those for Wireless Network IEEE (802.11) Policies options in **Security Settings**, Terminal Services in **Administrative Templates**, and settings for many other technologies.

Modifying Group Policy settings can have significant consequences. When performing Group Policy maintenance, you need to take reasonable precautions to test proposed changes and evaluate their effects in a staging environment prior to deployment.

# Group Policy Considerations for Renaming a Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain names are a critical part of the proper functioning of a Group Policy implementation. In the Windows Server 2003 family, administrators can rename a domain. This is accomplished by using the Rename Domain Tools available for download from the Windows Server 2003 Domain Rename Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. The Domain Rename tools provide a secure and supported methodology to rename one or more domains (as well as application directory partitions) in a deployed Active Directory forest.

**Important**

- Be sure to back up all your GPOs using GPMC after domain rename is complete. Backups made prior to domain rename cannot be used for **Restore** after a domain rename.

Renaming one or more domains is a complex process that requires thorough planning and understanding of the domain rename procedures. You must also modify any affected GPOs so that they work correctly. To modify the GPOs, use the Gpfixup.exe tool, which is included with the Domain Rename tools. Gpfixup.exe repairs GPOs as well as GPO references in each renamed domain. It is necessary to repair the GPOs and the Group Policy links after a domain rename operation in order to update the old domain name embedded in these GPOs and their links.

**Important**

- Renaming a domain is a complex process that requires thorough planning and understanding of the domain rename procedures. For more information about the domain renaming process, see the Windows Server 2003 Domain Rename Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Using Scripts to Manage Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

GPMC provides a comprehensive set of COM interfaces that you can use to script many of the operations supported by the snap-in. To help you get started, the GPMC installation includes sample scripts that use these interfaces. The sample scripts are installed in the \Program Files\GPMC\Scripts\ folder. Full details about the scripting interfaces are documented in the Group Policy Management Console SDK, is installed as Gpmc.chm in Program Files\GPMC\Scripts\ on any computer where you have installed GPMC. The GPMC SDK is also located in the Platform SDK. For more information about scripting GPMC operations, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

The sample scripts form the basis for a scripting toolkit that you can use to solve specific administrative problems. For example, you can run queries to find all GPOs in a domain that have duplicate names, or to generate a list of all GPOs in a domain whose settings are disabled or partially disabled. The scripts also illustrate some key scripting objects and methods to provide an overview of the many administrative tasks that you can accomplish with the Group Policy Management snap-in. You can modify and distribute the sample scripts in accordance with the terms of the End User License Agreement.

All of the scripts run from the command line. Running a script with the **/?** switch displays the usage for that script.

The sample scripts echo output to the command window and must be executed by using Cscript.exe. If Cscript.exe is not your default scripting host, you will need to explicitly specify Cscript.exe on the command line. For example, type **`d: \Program Files\GPMC\Scripts>cscript ListAllGPOs.wsf`**.

To make Cscript the default scripting host, type **`cscript //h:cscript`** at the command line.

Many of the sample scripts rely on a library of common helper functions contained in the file Lib_CommonGPMCFunctions.js. If you copy these scripts to another location, you must also copy this library file to that location in order for the script samples to work.

# Additional Resources for Group Policy Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this chapter.

**Related Information**

- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about Group Policy.
- "Deploying Security Policy" in this book.
- "Deploying a Managed Software Environment" in this book.
- "Staging Group Policy Deployments," in this book.
- "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.
- "Designing and Deploying File Servers" in *Planning Server Deployments* of this kit.
- "Planning an Active Directory Deployment Project" in *Designing and Deploying Directory and Security Services* of this kit.
- "Designing and Deploying Directory and Security Services"of this kit.
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for information about Active Directory.
- The Implementing Common Desktop Management Scenarios link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about using the sample Group Policy objects included on the Windows Server 2003 Deployment Kit companion CD.
- The Best Practice Active Directory Design for Managing Windows Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Group Policy Management Console White Paper link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Migrating GPOs Across Domains by Using GPMC link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, click to select the **Search in title only** check box.

- "Group Policy" in Help and Support Center for Windows Server 2003.
- Help in GPMC for detailed information about using GPMC to help deploy Group Policy and troubleshoot your deployment.
- Help for specific Group Policy settings in the default **Extended** view in Group Policy Object Editor (select a Group Policy setting to see detailed information for that setting).
- Command-line reference A-Z in Help and Support Center for Windows Server 2003, for more information about command-line tools such as Dcgpofix.exe.

- Type **hh system.chm** at the command line to see the full description of Group Policy settings. If you open this Help file on a Windows XP Professional–based computer, you will only see settings for Windows 2000 and Windows XP Professional; if you open this Help file on a computer running Windows 2003 Server, you will see settings for Windows 2000, Windows XP Professional, and Windows 2003 Server.

# Staging Group Policy Deployments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft® Windows® Server 2003 Group Policy provides powerful capabilities for deploying configuration changes across the enterprise. As with any other change within the enterprise, Group Policy deployments and ongoing updates require careful planning and testing to ensure a highly available and secure infrastructure. By using features included in the Group Policy Management Console (GPMC), you can create a test/staging/production deployment process that ensures predictability and consistency during Group Policy deployments.

**In This Chapter**


Overview of Group Policy Staging

Creating the Staging Environment

Synchronizing the Staging and Production Environments

Testing Group Policy in the Staging Environment

Preparing for Deployment to Production

Deploying Staged GPOs to the Production Environment

Additional Resources for Staging Group Policy Deployments


**Related Information**

- For more information about Group Policy and GPMC features and capabilities, see "Designing a Group Policy Infrastructure" in this book.
- For more information about using IntelliMirror to create a managed environment "Planning a Managed Environment" in this book.

# Overview of Group Policy Staging

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy is a powerful tool for configuring Microsoft® Windows® 2000, Windows® XP Professional, and Windows Server 2003 operating systems across the enterprise. This ability to affect configurations across hundreds or even thousands of computers necessitates good change management practices to ensure that the changes made to a Group Policy object (GPO) produce the expected results for the intended targets — users and computers — in infrastructures that use the Active Directory® directory service.

Most enterprises have change management processes in place to ensure that any new configurations or deployments of production systems go through rigorous testing in a non-production environment prior to moving them into the production environment.

In many change management processes, enterprises differentiate between a test environment, which is used to test changes, and a staging environment, which is a pristine environment that resembles production and is the last stop for a change before it is deployed to production. In this chapter, the terms test and staging are used interchangeably, without differentiating between them as physical environments. You can, however, use the techniques described in this chapter to create separate test and staging environments if your change management processes require them.

This kind of change management process is equally important for Group Policy changes because Group Policy is capable of affecting everything from registry settings to security settings to deployed software on a workstation or server. In addition to the many configuration settings that Group Policy accommodates, GPOs can be linked to a number of different scopes, and their effect can be filtered by users, computers or security groups. The ability to stage GPOs in a pre-production test environment and then test the various ways that they can be deployed prior to committing them to live users and computers is critical to ensure reliable, robust operation of your Windows-based infrastructure.

The creation of a staging environment is critical to any successful deployment of Group Policy within your Active Directory–based infrastructure. There are several options that you can choose from to create such an environment. These options are enabled by using features within the Group Policy Management Console (GPMC). GPMC runs on 32-bit computers running a member of the Windows Server 2003 family operating system or Windows XP Professional with Service Pack 1 (SP1). GPMC consists of a set of scriptable interfaces for managing Group Policy and an MMC-based user interface (UI). The UI integrates all previous Group Policy tools into a unified Group Policy-management console. GPMC can manage both Windows Server 2003 and Windows 2000 Active Directory–based domains, although some features are available only in Active Directory environments running Windows Server 2003.

GPMC console-based features and accompanying Windows Script Host (WSH) scripts give you the ability to create a staging environment that mimics your production environment. You can then use the staging environment to test new or changed GPOs. Once those GPOs are validated, you can use GPMC to migrate them to your production domains.

# Group Policy Staging Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process for staging Group Policy involves creating a staging environment that mimics the production environment, testing new Group Policy settings in the staging environment, and then deploying those settings in the production environment. The specific deployment approach you use depends on the configuration of your staging environment. Figure 3.1 shows the steps in the Group Policy staging process.

**Figure 3.1 Staging Group Policy**

Initially, assembling a staging environment for Group Policy is simply a matter of identifying available hardware that can be used in creating an infrastructure that is similar to your production environment, and then setting up the appropriate logical structure. You can then use tools in GPMC to import production Group Policy into the staging environment. Once you have created the environment, testing Group Policy involves implementing changes and measuring their effect on test users and computers that mimic production users and computers. Once you have validated your changes, you can again use tools in GPMC to migrate changed or new Group Policy settings to your production environment.

On an ongoing basis you need to maintain Group Policy and continue to evaluate changes. Consequently, you need to keep the staging environment synchronized with the production environment over time. You can use GPMC tools such as the supplied scripts and the console-based Group Policy backup, copy and import features to maintain the staging environment over time.

# GPMC Staging Technology Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

GPMC is a new tool for managing Group Policy. It includes several features for creating and maintaining Group Policy, including the Group Policy Modeling Wizard for planning Group Policy deployments, the Group Policy Results Wizard for viewing GPO interaction and for troubleshooting, and the ability to use a single MMC interface to manage Group Policy across your organization, including importing and exporting, copying and backing up and restoring GPOs. For detailed information about GPMC features and capabilities, see Help in GPMC, and see Designing a Group Policy Infrastructure in this book. For staging Group Policy, the most important GPMC tools are those that allow you to stage and migrate GPOs between forests and domains.

**Backup and Import**

GPMC provides the ability to back up one or more GPOs. These backups can then be used to restore an individual GPO to its previous state (using the restore option), or backups can be imported into an existing GPO, overwriting any previous policy settings. The restore operation is used only to restore a GPO into the same domain from which it was backed up. By contrast, the import operation is used in cases where the backup was made from any GPO in the same domain, a different domain or even in a different untrusted forest, such as a test forest isolated from the production forest. Note that while both restore and import operate on previously backed-up GPOs, restore has some additional capabilities that are specific to a restore operation. You will use the backup and import operations and the copy operation to stage and migrate GPOs into your production environment. Figure 3.2 illustrates the import operation. In this case, GPO X in a test forest contains a number of security principals who are assigned the Log on locally user right. This GPO is backed up and then imported into the production forest. During the import operation, the original security principals are mapped to new ones that exist in the production domain.

**Figure 3.2 GPO Import Operation Across Forests**




**Copy**

Using the copy capability in GPMC, you can right-click a GPO, copy it from one domain, and paste it into a new domain. In a copy operation, when you copy a GPO into a new domain, a new GPO is created. This differs from the import operation, which erases and then overwrites an existing GPO. Copy operations require that the destination domain is trusted by the source domain. In addition, the administrator performing the copy operation must have rights to read the source GPO in order for the copy operation to succeed. With both import and copy operations, GPMC supports the ability to perform security principal and UNC mapping between references to those objects in the source and destination GPOs. Figure 3.3 illustrates a copy operation. In this case a GPO is migrated from Domain B to Domain C and some of its associated security principals are mapped to new principals on Domain C.

**Figure 3.3 Copy Operation Between Domains in a Production Forest**




**Migration Tables**

GPOs can contain references to security principals and UNC paths as part of a policy setting. For example, within security settings policies, you can control who can start and stop a particular Windows service by using the Group Policy Object Editor to specify a user or user group for that policy. Figure 3.4 illustrates the security settings that can be applied to the Messenger service. In this case, these security settings can be mapped from security principals in the staging environment to security principals in a production environment by using migration tables.

**Figure 3.4 Security Principals on the Messenger Service**




In addition, a GPO itself has an associated Discretionary Access Control List (DACL) that is used to control which computers, users, or groups process a GPO and which users can create, modify, and edit the GPO. The security principals associated with the DACL on a GPO can also be taken into consideration when the GPO is deployed from one domain to another.

Migration tables also support mapping of UNC paths, which might exist in Software Installation, Folder Redirection, or Scripts policy. To deal with any differences in these paths between the test and production environments, you can use migration tables to replace server and share names as you migrate Group Policy.

If a GPO created in another domain or forest is migrated to your production environment, you need a way of modifying the associated security principal references to reflect the references found in the production domain. GPMC provides a Migration Table Editor (MTE) that you can use to create a mapping file for security principals and UNC paths. MTE creates an XML format

file with a .migtable extension; this file specifies the source and destination security principals or UNC paths for a GPO migration. For more information about MTE, see Creating Migration Tables later in this chapter.

# Creating the Staging Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in staging and deploying Group Policy is the creation of the staging environment. This involves building a test infrastructure that mirrors that of production and allows you to test new or changed Group Policy settings without affecting production users and computers. Figure 3.5 illustrates this step in the staging process.

**Figure 3.5 Creating the Staging Environment**




At this point, you need to make decisions about the placement of your staging environment and its trust relationships to your production environment. You can choose to create:

- A staging domain within the production forest.
- A staging forest with no trusts to the production forest.
- A staging forest with trusts to the production forest.

Each option has its benefits and drawbacks, as described in Table 3.1.

**Table 3.1 Choosing a Staging Approach**

| Approach | Advantages | Disadvantages |
|---|---|---|

| | | |
|---|---|---|
| Staging Domain within Production Forest | • Can use GPMC copy operation to move GPOs between staging and production environments.<br>• Can leverage existing production infrastructure services (for example, DNS, DHCP).<br>• Might require less hardware to implement than a completely isolated environment that requires supporting infrastructure.<br>• Easier to keep synchronized with production environment because all settings and services are in the same forest.<br>• Might require less use of migration tables if migrating from domain to domain in the production forest (for example, some security principals can be re-used regardless of domain). | • Might not be sufficiently isolated from production environment to ensure testing does not affect it (for example, site-linked GPOs cannot easily be tested because sites span domains within a forest).<br>• Might prove restrictive if changes to the environment are required for testing. |
| Staging Forest with no trusts to Production Forest | • Completely isolated from production environment; provides maximum protection from test GPOs affecting production computers and users.<br>• No security overlap between staging and production; administrators in staging and production forests need not have access to both forests.<br>• Provides flexibility; administrators can experiment freely with settings and configurations without affecting the production environment. | • Difficult to keep synchronized with production forest.<br>• Without trusts, movement of data and settings between forests is more cumbersome.<br>• Migration Tables are required to move GPOs that contain security principals or UNC paths from staging to production.<br>• Cannot use GPMC copy operation to migrate GPOs; must use GPMC import. |
| Staging Forest with trusts to Production Forest | • Can use GPMC copy operation to move GPOs between staging and production environments.<br>• Somewhat isolated from production environment.<br>• Provides flexibility; administrators can experiment freely with settings and configurations without affecting the production environment.<br>• Might not need migration tables to map UNC paths, since all paths could be available via current trusts. | • Difficult to keep synchronized with production forest.<br>• Trusts between staging and production environments allow users in one environment to access resources in the other.<br>• Migration Tables required to move GPOs that contain security principals from the staging environment to the production environment. |

Weigh the advantages and disadvantages described in Table 3.1 when making your choice of a staging approach. Ultimately, your choice depends on your own unique requirements. Once you have made your choice, you are ready to determine the hardware requirements for the staging environment.

**Hardware Requirements**

Regardless of the approach you choose, it will be necessary to dedicate some additional hardware to the construction of your staging environment. How much hardware you need depends upon the kinds of testing you need to do and how specific your Group Policy testing requirements are. For example, if your production forest has workstations located across slow network links from your domain controllers, this fact can affect the application of Group Policy because some Group Policy settings are not applied across slow links. It is important that your test environment reflect this situation for you to get an accurate picture of how your production environment will be impacted by changes in Group Policy. GPMC can help in such a situation, for example, by providing the capability to model the impact of slow links on GPO application. However, you might not be able to fully mirror your production environment unless you dedicate sufficient systems and network hardware to the staging environment. Your goal is to produce a testing and staging environment that reflects what computers and users in your production environment will see when new or changed GPOs are applied.

**Preparing the Staging Environment**

Once you have chosen a staging approach and set up your hardware, install Windows Server 2003 (or Windows 2000 Server) and Active Directory on your staging servers in preparation for synchronizing the configuration of production and staging environments. In most cases, you should ensure that the staging environment is running at the same operating system, service pack and hot fix levels as your production environment. This is important to ensure consistent test results. In addition, ensure that the supporting infrastructure, such as Directory Name Service (DNS), Distributed File System (DFS), and related services are also configured as in the production environment. DNS especially is critical to proper processing of GPOs. If you decide to use a staging approach that places a staging domain in your production forest, then you can use your existing production DNS infrastructure for name services.

### Important

- You can use GPMC to manage both Windows Server 2003 and Windows 2000 Active Directory–based domains, but you must run GPMC on a computer running a member of the Windows Server 2003 family operating system or Windows XP Professional SP1 and the .Net Framework.

If you build a separate forest for staging, then you need to address the issue of name services integration. Name services might include either DNS or Windows Internet Name Service (WINS), depending on the types of trusts you created. You might need to create a separate DNS infrastructure for your staging environment. This is particularly true if you are using Active Directory–integrated DNS in your production forest, because Active Directory–integrated zones cannot support dynamic registration of clients from foreign forests. If you plan to create trusts between your staging and production forests, the name services infrastructure in each forest must be aware of the other. Note that if both your staging and production forests are running at Windows Server 2003 forest functionality level, and you want to create inter-forest trusts between the two forests, then DNS integration is required. However, if you are running at Windows 2000 forest functionality level, which requires the use of down-level trusts, you must have the WINS service available and integrated between the two forests to create the required trusts.

Once your staging environment is fully configured with the base elements required to deploy Group Policy, the next step is to synchronize the staging and production environments.

# Synchronizing the Staging and Production Environments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Once you have created a basic staging infrastructure that reflects your production environment, you need to ensure that all GPO settings and security are identical between the two environments. Synchronization also requires ensuring that a sufficient representation of OUs, users, computers and groups exist in both environments, because you need to be able to test GPO links and the effects of security group filtering as they would exist in the production environment. Figure 3.6 illustrates the synchronization step in the staging process.

**Figure 3.6 Synchronizing the Staging and Production Environments**




The goal of any test environment is to ensure that it mirrors the production environment as closely as possible. You can use GPMC tools to help both with the initial synchronization and to keep the test environment synchronized with the production environment over time. The first set of tools you will use are two Windows Script Host (WSH) scripts that are installed with GPMC. These two scripts, CreateXMLFromEnvironment.wsf and CreateEnvironmentFromXML.wsf, are installed in the Scripts folder in the folder Program Files\GPMC\Scripts. The script CreateXMLFromEnvironment.wsf is run against a production domain, stores all policy-related information in an XML format file, and creates backups of the GPOs it finds in the production domain. Note that this script works only against a single domain at a time, not against an entire forest. The script CreateEnvironmentFromXML.wsf, uses the XML format file and any backup GPOs created by CreateXMLFromEnvironment.wsf to re-create GPOs and other objects from the production domain in a staging domain. Table 3.2 describes the objects and settings that CreateXMLFromEnvironment.wsf captures, and shows additional objects you can capture by using command-line options when running the script.

**Table 3.2 Objects Captured by CreateXMLFromEnvironment.wsf**

| Object Type | Captured by Script | Additional Command-line options |
|---|---|---|
| All GPOs and GPO settings for the domain or OU | Yes | To get GPO settings, you must provide a Template Path by using the /TemplatePath option to specify a file system location to store backed up GPOs. If no template path is specified, GPOs are not backed up.<br>You can exclude permissions on GPOs by using the /ExcludePermissions option. |
| Organizational Units | Yes | You can capture only a portion of the OU tree by using the /StartingOU switch, along with the DN-style path to an OU. |

| | | |
|---|---|---|
| GPO Links and link attributes (for example, disabled, block inheritance) | Yes, except that links on site objects are not captured | None |
| Policy-related Permissions | Yes | You can exclude permissions by using the /ExcludePermissions option. |
| WMI Filters | Yes | None |
| Users | Optional | User accounts are not captured unless you use the /IncludeUsers option. |
| Security Groups | Yes | By default, only security groups defined in OUs are captured by the script. You can extend this to include all groups in the Users container and in the domain root by using the /IncludeAllGroups option. |
| Computers | No | None |
| Sites | No | None |

There are a few things to bear in mind when using the script CreateXMLFromEnvironment.wsf. First, if you use the /IncludeUsers option to capture user objects, when those objects are re-created in the staging domain, you will need to supply a password for each user captured. You can do this by manually editing the resulting XML file and adding a password for each user. (For an example of the syntax, see the file SampleEnvironment.xml in the Program Files\GPMC\Scripts folder on any computer where you have installed GPMC.) Alternatively, if you have any users that do not have passwords specified in the XML file, the CreateEnvironment script will prompt you to supply a password. All users that do not have passwords specified in the XML file will be created with this password. Also note that the script does not capture computers. This is because computer objects in Active Directory correspond to physical hardware resources, and those might differ between the production and staging environments. Finally, the script captures neither sites nor GPO links on sites. Because sites can span multiple domains and can have an impact on Active Directory replication, it is best to re-create these objects, and GPO links on them, manually in your staging environment.

# Example: Creating an XML Format File from a Production Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assume that your production domain is called Contoso.com, and you want to export Group Policy and related information to create a new staging domain for GPO testing. In this example, assume that you want to capture GPOs from the entire domain and include user accounts as well as groups. To export the information you need, complete the following tasks:

1. Ensure that you have sufficient permissions on the production domain to extract the necessary data. You must have the rights to read all objects that you are capturing, including GPOs, OUs, users, and groups (and their memberships).

2. Create a folder to store the XML format file that describes the information collected by the script.

3. Create a folder to store backups of the GPOs that are extracted by the script.

4. Run the script CreateXMLFromEnvironment.wsf from the Scripts folder in the GPMC installation folder. You must precede the script name with the command `cscript` if cscript.exe is not your default WSH engine. For this example, type the following from the command line:

   ```
   Cscript "%programfiles%\gpmc\scripts\CreateXmlFromEnvironment.wsf".\production.xml
   /Domain:contoso.com /DC:contoso-dc1 /TemplatePath:.\GPObackups /IncludeUsers
   ```

   This command creates the XML format file Production.xml in the folder where the script is run. The backed-up GPOs are created in a sub-folder of the current folder called GPObackups. Placing a backslash (\) in front of the production.xml and GPObackups paths causes the script to use a relative path, and create the XML file and backup GPO folders in the current directory from which the script is run. Using a relative path makes it easier to copy the XML and backups to different locations from which they can be restored.

   The script starts its capture at the domain level, Contoso.com. You can also run the script at an OU level, in which case you would use the /StartingOU option in addition to the /Domain option. If you exclude the /Domain option, the current domain is assumed. The /DC option tells the script to use the domain controller contoso-dc1, and the /TemplatePath option specifies that the backups of all of the GPOs that are captured are stored in the folder GPOBackups. Finally, the /IncludeUsers option ensures that user accounts are captured by the script as well.

   **Caution**

   - You can open and edit the XML format files produced by the script CreateXMLFromEnvironment.wsf in a text editor or any XML editor. Be aware, however, that XML-formatted files must adhere to a specific syntax. If you change that syntax, you might affect the ability of the script CreateEnvironmentFromXML.wsf to read the input file.

Once you have captured the production environment by running the script CreateXMLFromEnvironment.wsf, you need to run the script CreateEnvironmentFromXML.wsf, using the .XML format file output by CreateXMLFromEnvironment.wsf as input. You must run the script CreateEnvironmentFromXML.wsf, from within the staging domain, or you can run this script from a computer that is not in the staging domain if you already have trusts to the staging domain in place.

# Importing Production GPOs into the Staging Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The script CreateEnvironmentFromXML.wsf has a number of different options that you can use to qualify the creation of GPOs in your staging environment. In the simplest form, just supply an XML format file created from the production domain to the script and optionally direct the operation of the script to a DC in your staging domain. The script creates GPOs and related objects in the staging domain that correspond to the data that was captured from the production domain. If you need to modify this process, the script provides a number of command-line options:

- **Undo.** This option removes all objects (GPOs, GPO permissions, OUs, WMI filters, users and groups) specified by the XML format file from the staging environment. This option is useful if you need to reverse changes you made to your staging domain.

- **ExcludeSettings.** This option creates GPOs in the destination domain, but with no settings. Use it when you do not actually want to import the settings in any Group Policy objects, but rather just want to create any OUs, users, and user groups that might have been captured.

- **ExcludePermissions.** This option causes the script to ignore any policy-related permissions contained in the XML format file. Instead, when the new GPOs and other objects are created in the staging environment, they are created with the default permissions.

- **MigrationTable.** This option lets you specify a .migtable file that you create using the MTE to specify mapping of security principals and UNC paths in your production environment GPO settings to the appropriate security principals and UNC paths in the staging environment.

- **ImportDefaultGPOs.** This option imports settings into the default domain policy and the default domain controllers policy if settings for these GPOs are specified in the XML file. If this option is not specified, these GPOs will not be modified.

- **CreateUsersEnabled.** This option creates user accounts as enabled instead of disabled.

- **PasswordForUsers.** This option allows you to specify the password to use for any users that do not have passwords specified in the XML file. The same password will be used for all users that do not already have passwords specified in the XML file.

- **Q.** The Q option runs the script in quiet mode, provided that all necessary parameters have been supplied on the command line. Without this option, you are warned that this script should only be used for creating staging environments, and if necessary, you will be prompted to supply a password for any users that do not have passwords defined in the XML file.

# Example: Populating the Staging Domain from the XML Format File

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assume that your staging environment is the domain test.contoso.com, and is in the same forest as the production domain captured earlier in this chapter. Even if the staging domain is not in the same forest as the production domain, the steps for populating the staging domain are the same, but different mapping of security principals using migration tables might be required. To populate the staging domain, complete the following tasks:

1. Ensure that you are running the script CreateEnvironmentFromXML.wsf with sufficient permissions in the staging domain. You should run the script as a user as a member of Domain Admins or have equivalent access in the domain.

2. Ensure that you have access to the XML format file and backup GPOs that were created in the production domain by running CreateXMLFromEnvironment.wsf.

   Note that when you run CreateEnvironmentFromXML.wsf, you are only referencing the XML format file (not the location of the backup GPOs) in the command-line options. That file includes the paths to the backup GPO files. Consequently, when you specify the XML file to CreateEnvironmentFromXML.wsf, the script uses any backup GPO files in the folder that was specified when the script CreateXMLFromEnvironment.wsf was run. If you ran CreateXMLFromEnvironment.wsf using the command as shown in "Example: Creating an XML Format File from a Production Environment" earlier in this chapter, then the XML file will indicate that the backups are in a sub-folder of the current folder. If you did not use a relative path when running CreateXMLFromEnvironment.wsf, there are three ways to ensure that CreateEnvironmentFromXML.wsf, can find the required files:

   - Copy the specified folder structure from the location where it was created to an identical path on the local computer from which you run CreateEnvironmentFromXML.wsf.

   - Specify a network share rather than a local drive when you first create the XML format file (the share must be also accessible from the location where you run CreateEnvironmentFromXML.wsf).

   - Edit the XML format file to change the path entries to point to a different location for the backup GPO files.

3. Run CreateEnvironmentFromXML.wsf from the Scripts folder in the GPMC installation folder. You must precede the script name with the command `cscript` if cscript.exe is not your default WSH engine. For this example, type the following from the command line:

   ```
   Cscript CreateEnvironmentFromXml.wsf /xml:c:\staging\production.xml
   /Domain:test.contoso.com /DC:test-dc1
   ```

   The script generates a warning that the script is intended for creating staging environments only, and then prompts you to enter a password for user objects. If you use the **/Q** option and supply the password using the **PasswordForUsers** option when you run this script, these messages are not presented. If you confirm that you want to proceed, the script provides status as it processes the XML file and GPOs. You can then confirm that all steps completed correctly by using the Active Directory Users and Computers MMC snap-in and GPMC to verify that users, groups, and GPOs were successfully created.

# Maintaining Synchronization of Staging and Production

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The scripts CreateXMLFromEnvironment.wsf and CreateEnvironmentFromXML.wsf are used to create an initial staging environment from your production environment. But maintenance of Group Policy, including testing new and changed GPOs, is ongoing; how do you keep your staging environment synchronized with the production environment on a continuing basis? These two scripts provide an all or nothing method for populating GPOs — they are not specific enough to capture and import only specific GPOs.

The backup and import functions in GPMC give you the ability to selectively synchronize specific GPOs between your production and staging environments. You can use the backup capability to create a backup of the settings and security of a production GPO. You can then import the backup over an existing GPO in your staging domain, thereby synchronizing it with the production GPO. For more information about backing up and importing GPOs, see "Deployment Examples" later in this chapter.

# Testing Group Policy in the Staging Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have created your staging environment and synchronized Group Policy with your production environment, you can begin to test planned Group Policy changes. The best mechanism for testing Group Policy is by using a combination of the Results and Modeling tools provided with GPMC and exercising real user accounts and computers in the test environment to process actual GPOs.

The Group Policy Results tool is useful when you have applied new GPO settings to a computer and user, and need to verify that all of the expected settings were actually applied. Group Policy Modeling can be used to determine the effects of changing the location of a user or computer within the Active Directory namespace, changing the group membership of a user or computer, or to observe the effects of a slow link or loopback policy. The Group Policy Modeling tool lets you test the effects of a change without actually making the change, while Group Policy Results tool tells you what actually happened. Group Policy Results runs on the target computer, so you must have access to that computer. Group Policy Modeling runs on a Windows Server 2003 Domain Controller, so there must be one available to run the modeling process. Note that with Group Policy Modeling, you can model policy settings on computers running Windows 2000 and Windows XP Professional as well as Windows Server 2003. Bear in mind that the Group Policy Modeling tool simulates processing of policy, while the Group Policy Results tool shows the effects of policies actually processed.

Figure 3.7 illustrates this step in the process.

**Figure 3.7 Testing Group Policy in the Staging Environment**




**Testing by Logging on as Test User**

The first and best method for testing Group Policy is to make the actual changes to your staging domain GPOs and then test the results by logging on to workstations with test user accounts to observe the effect of the changes. In this way you can see first hand how the user has been affected by the changes.

**Testing by Using Group Policy Results**

You can use the Group Policy Results Wizard in GPMC to get detailed reports of which policies are applied to users and computers. You can then make any needed changes in your test GPOs accordingly. The Group Policy Results Wizard is used after all Group Policy is processed for a given user and computer to inform you as to what settings were applied. The results are gathered by querying the WMI-instrumented Group Policy logging facility on a Windows XP–based computer or Windows Server 2003–based server that processed Group Policy. The wizard thus returns the settings that were actually applied rather than expected settings.

For more information about the Group Policy Results Wizard, see "Designing a Group Policy Infrastructure" in this book.

**Testing by Using Group Policy Modeling**

The second method for testing Group Policy is to use the Group Policy Modeling Wizard in GPMC to model changes to your environment before you actually make them. Group Policy Modeling lets you perform hypothetical tests on user and computer objects prior to a production rollout to see how Group Policy settings would be applied if you made changes such as moving the user or computer objects to a different OU, changing their security group membership, or changing the effective WMI filters. Be aware, however, that results obtained using Modeling are modeled rather than actual policy settings. Therefore, once you have modeled the scenario that meets your needs, it is always best to use the GP Results wizard to verify the expected settings.

Because Group Policy Modeling does not let you specify proposed changes to settings in a GPO, you need to make the proposed changes to your staging GPOs and then run the Group Policy Modeling Wizard for a given OU, user, or computer to determine the resultant set of policy.

Group Policy Modeling also gives you the ability to model Group Policy behavior when your computers are processing policy across a slow network link, which can affect which Group Policy extensions are processed. For example, if a computer connects to a domain controller over a slow network link (defined by default as any rate slower than 500 kilobits per second (Kbps)) then Group Policy extensions such as Software Installation and Folder Redirection are not processed. Group Policy Modeling can simulate a slow link speed and use it to determine what the effective policy settings will be for the user and computer being modeled. In addition, Group Policy Modeling supports testing the effects of Group Policy loopback processing. With loopback processing enabled, the same settings are applied to a computer regardless of the user who logs on to it. Note that you must specify that you want to model loopback processing within the Modeling wizard; loopback processing is not modeled by default.

You can specify slow-link detection, loopback processing, or both when using the Group Policy Modeling Wizard. For loopback processing, you can choose whether to replace or merge user-specific policy. The replace mode replaces all of a user's normal policy settings with those defined in the user configuration of the GPOs that apply to the computer object (the loopback settings). Merge mode merges the user's normal policy settings and the loopback settings. In the case where a policy item in the user's normal policy conflicts with the loopback settings, the loopback settings are applied.

Note that the Group Policy Modeling process actually runs on a Windows Server 2003 domain controller in your Active Directory domain. By contrast, the Group Policy Results Wizard runs at the Windows XP–based workstation or Windows Server 2003–based server that is processing Group Policy. Group Policy Results uses the RSoP WMI provider to generate information about policy processing. Group Policy Modeling relies on the Windows Server 2003 Resultant Set of Policy Provider service to perform its analysis.

For more information about the Group Policy Modeling Wizard, see "Designing a Group Policy Infrastructure" in this book.

# Preparing for Deployment to Production

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Once you are satisfied that your changes to Group Policy have been thoroughly tested in the staging environment, you are almost ready to deploy the new or changed GPOs in your production environment. Before you can do that, however, you need to assess whether you will need to map security principals or UNC paths contained in your GPOs to different values as part of the migration. This step is illustrated in Figure 3.8.

**Figure 3.8 Preparing for Deployment to Production**




**Determining Your Migration Mapping Requirements**

Your staging environment might be a test domain in production, a separate but trusted test forest, or a separate test forest that is not trusted. In each case, you will probably have to create and use a migration table as you deploy new or changed GPOs in your production environment. Migration tables satisfy three different types of mapping requirements:

- You need to map an Access Control Entry (ACE) on one or more GPOs to different security principals as you migrate the GPOs to the production environment. The ACEs on a GPO describe which users, computers and computer groups will process that GPO, and which users or user groups can view and edit settings in or delete the GPO.

- You need to map security principals within security policy settings or Folder Redirection settings defined in one or more GPOs. Specifically, policies such as User Rights Assignment, Restricted Groups, File System, Registry, or System Services allow you to specify particular users or groups who can access or configure those resources. The Security Identifier (SID) for that user or group is stored in the GPO and must be modified to reflect production domain users or groups when the GPO is migrated.

- You need to map UNC paths when you have defined software installation policy, folder redirection or scripts policies that reference UNC paths. For example, you might have a GPO that references a script stored in an external path, such as the Netlogon share on a remote server. This path might need to be mapped to a different path when the GPO is migrated. UNC paths are usually specific to a given environment, and might need to be changed when you migrate the GPO to your production environment.

If any of these three conditions above is true, you will need to create a migration table that can be used to map the values in your test GPOs to the correct values in the production domain when they are migrated

# Creating Migration Tables

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Migration tables are XML format files that specify how to map security principals and UNC paths to their production domain equivalents. Migration table files have a default file name extension of .migtable. The Migration Table Editor (MTE) is a GPMC tool that can be used to create and edit migration tables. It can be accessed in either of two ways:

- You can start MTE and create or edit a migration table during a GPMC copy or import operation. In this case, MTE starts in a separate window that allows you to create a new migration table or edit an existing one.

- You can start MTE in stand-alone mode (independently of an import or copy operation) and create or edit the migration table in advance of migrating GPOs to your production environment.

You can also create migration tables by using a supplied script, as described later in this section.

An advantage of creating the migration table in advance is that you can be sure that the migration settings you define are exactly what you want before beginning the deployment. Therefore, when you are ready to move your test GPOs into production, you should first create one or more migration tables for the GPO(s) that you need to migrate. Note that a single migration table can be used for more than one GPO. You might use a single migration table that covers every possible security principal and UNC path combination for a given migration from a staging-to-a production domain. In that case you can simply apply the same migration table to every GPO that is deployed from the staging domain to the production domain, and those principals and paths that match will be correctly mapped.

**Using the Stand-alone Migration Table Editor**
To start MTE in stand-alone mode, run Mtedit.exe from the GPMC installation folder. MTE starts with a blank migration table that you can populate manually by typing entries into the grid, or you can auto-populate the table by using one of the auto-populate methods.

**Auto-populating the migration table**
The easiest way to start creating a migration table is to use one of the auto-populate features, accessed from the **Tools** menu in MTE. You can auto-populate from both backup GPOs and from live GPOs. To auto-populate a migration table, you need to perform the following tasks:

1. Choose to auto-populate the table from live GPOs or from backup GPOs.

   When you are ready to migrate a GPO in your staging environment into your production environment, you can use **Populate From GPO** against the live GPO in the staging environment to start the migration table. The process for auto-populating the table from a backup GPO is the same, except that you must provide a path to the backup GPO. In that case, if you have more than one backed up GPO, a list is displayed from which you can choose.

   Note that you can select multiple GPOs or backup GPOs when auto-populating a single migration table. This allows you to use a single migration table for all GPOs in a domain.

2. Choose whether to include security principals from the DACL on the GPO.

   When you auto-populate a migration table, you can select the option to include security principals from the DACL on the GPO. If you select this option, security principals on the GPO's DACL are included in the table along with security principals referenced in the GPO settings. Duplicate source security principals are not repeated in the migration table. MTE supports a number of different object types that can be mapped, as described in Table 3.3.

   **Table 3.3 Object Types Supported in the Migration Table**

| Object Type | Used to Map... |
|---|---|
| User | Individual user accounts. |

| | |
|---|---|
| Domain Global Group | Domain global groups. |
| Domain Local Group | Domain local groups. |
| Universal Group | Universal groups. |
| Computer | Computer names. For example, an individual computer can be given read and apply Group Policy permissions on a GPO. |
| UNC Path | UNC paths used in software installation policy. |
| Free Text or SID | Undetermined security principals. For example, you might reference security principals in a GPO by name rather than by SID (typed as "administrators" instead of "DomainX\Administrators"); or it might not be possible to resolve security principals to determine the type.<br><br>This can occur because, for example, when setting restricted group security policy, you can type in the name of the group rather than resolve it against a live domain. In this case, the group name is stored in the Group Policy object as its name rather than its corresponding SID. MTE considers such a security principal as Free Text or SID.<br><br>In addition, you can type raw SIDs into the MTE. In that case, because the object type is not known by MTE, it must be specified as Free Text or SID. |

3. Modify the **Destination Name** for each security principal and UNC path.

   After you have populated the migration table, you can choose to modify the **Destination Name** field for each record. The default **Destination Name** value is **Same As Source**, which means that the same security principal or UNC path will be used in the destination GPO as the source. In this case, the value is copied without modification and the mapping accomplishes no changes. Typically you will need to change this field for one or more source entries when deploying a GPO from a test to a production environment. To change the destination field, you can either type in an entry or right-click the field and make a choice from the menu.

   The two menu options available are **Browse** and **Set Destination**. Choosing **Browse** allows you to use the object picker to select a security principal in any trusted domain. If you choose **Set Destination**, you can choose one of three options:

   - **No Destination.** If you specify **No Destination**, the security principal is not included in the destination GPO when migrated. This option is not available for UNC path entries.

   - **Map by Relative Name.** If you specify **Map by Relative Name**, the security principal name is assumed to already exist in the destination domain, and that destination name will be used for the mapping. For example, if the source name is **Domain Admins** for the test.contoso.com domain and you are migrating the GPO into the contoso.com domain, then the name Domain Admins@test.contoso.com will be mapped to Domain Admins@contoso.com. The group must already exist in the target domain for the actual import or copy operation to succeed. This option is not available for UNC path entries.

   - **Same As Source.** If you specify **Same As Source**, the same security principal is used in both the source and destination GPOs. Essentially, the security entry is left as-is. Note that this option is only practical if you are migrating from a test domain in the same forest as the production domain, or if you are migrating from a test domain in a different forest that trusts the production forest. The requirement for a source name to map successfully is that it can be resolved by users and computers in the production forest.

   There are some restrictions on the options available for the destination name. UNC paths only support the **Same As Source** option, or you can manually enter a different UNC path. Security Principals designated as Free Text or SID do not support **Map by Relative Name**.

It is also important to note that you will receive a warning if you map from one group type to another. For example, if you have a source principal that is a Domain Global Group and you select a Domain Local Group as the destination, you will be warned that the destination name is of a different type from the source. If you then try to validate the file, the validation process fails, but you can still use the migration table to perform a migration. Note that the migration table does not support mapping to a built-in security group such as the Administrators group.

If you need to delete a row from MTE, select the desired row, then right-click the row and choose **Delete** from the menu.

4. Validate the migration table.

   Before saving the migration table, it is best to validate the file. To do this, from the **Tools** menu, choose **Validate**. The validation process determines if the XML format of the resulting file is valid and verifies that the destination names are valid from a migration perspective. For example, if you enter a UNC path for the destination, and the path does not exist, the validation process will return a warning. Specifically, the validation process:

   - Validates the existence of destination security principals and UNC paths.
   - Checks that source entries with UNC paths do not have destinations of Map By Relative Name or No Destination, which are not supported.
   - Checks that the type of each destination entry in the table matches the real type in Active Directory.

     If you are entering data manually, the validation process is especially important to ensure that an entry error does not prevent a successful migration. Note that a validation of the mapping file might fail because the user editing the file does not have the ability to resolve the security principals or UNC paths specified in the file. However, that does not mean that the file will not work as expected during a migration, provided that the user who performs the migration can resolve the security principal and UNC names. Validation messages will indicate whether there is a syntax error in the table or whether the validator simply cannot resolve a security principal name or UNC path. In the case of a name resolution failure, ensure that you will have sufficient access to both source and destination resources during the actual migration.

5. When you are finished editing the table, save the resulting .migtable file by choosing **Save** from the **File** menu.

**Entering Migration Table Entries Manually**

If you choose not to use the Auto-Populate feature, or if you need to enter some data manually, take care to adhere to the proper formats in order for the migration table to be valid. Table 3.4 shows the proper form for each object type supported in the migration table. Note that these formats are required in both source and destination fields.

**Table 3.4 Required Formats for Migration Objects**

| Object Type | Required Format |
|---|---|
| User | *a. UPN - User@UPN suffix*<br>*b. SAM – NetBIOS domain name\user*<br>*c. DNS – DNS domain name\user*<br>For example, PilarA@contoso.com, contoso\PilarA, or contoso.com\PilarA. |
| Domain Global Group | *a. UPN - Group@UPN suffix*<br>*b. SAM – NetBIOS domain name\group*<br>*c. DNS – DNS domain name\group*<br>For example, Domain Admins@contoso.com, contoso\Domain Admins or Contoso.com\Domain Admins. |

| | |
|---|---|
| Domain Local Group | *a. UPN - Group@UPN suffix*<br><br>*b. SAM – NetBIOS domain name\group*<br><br>*c. DNS – DNS domain name\group*<br><br>For example, Administrators@contoso.com, contoso\Administrators, Contoso.com\Administrators. |
| Universal Group | *a. UPN - Group@UPN suffix*<br><br>*b. SAM – NetBIOS domain name\group*<br><br>*c. DNS – DNS domain name\group*<br><br>For example, Enterprise Admins@contoso.com, contoso\Enterprise Admins, or contoso.com\Enterprise Admins. |
| Computer | *a. UPN – Computer Name$r@UPN suffix*<br><br>*b. SAM – NetBIOS domain name\computer name$*<br><br>*c. DNS – DNS domain name\Computer name$*<br><br>For example, server1$@contoso.com, contoso\server1$, or contoso.com\server1$. The $ indicates the computer's hidden computer account. |
| UNC Path | *servername**sharename*\. For example, \\server1\packages. |
| Free Text or SID | Either a string or the string representation of a SID. For example, "PilarA" or "S-1-5-21-1473733259-1489586486-3363071491-1005". SIDs cannot be specified in the destination field. |

**Creating a Migration Table by Using a Script**

If you need to automate the process of creating migration tables, you can use a supplied script, CreateMigrationTable.wsf. You can also use this script in place of MTE to generate the initial migration table, and then use MTE to modify the table. The script CreateMigrationTable.wsf is located in the Scripts folder in the GPMC installation folder. The script supports auto-populating a migration table using either a live GPO or a backup GPO location. You can also have the script read from all GPOs within a domain. In that case, all possible security principals found in GPOs in the staging domain are put into the migration table, and that single migration table can be used for any GPO deployment from that staging domain to a production domain.

To use the script CreateMigrationTable.wsf, from the command line, change directories to the Scripts folder and run the script, preceded by the Cscript.exe command. Note that the script always includes the security principals that are part of the DACL on the GPO, unlike MTE, which gives you the option to include them. The script also has an option to set the destination name to **Map by Relative Name** rather than the default **Same As Source**. You use the /MapByName argument to implement relative naming. The following command illustrates how the script can be used. In this command, a GPO named Finance OU Desktop Policy is located in a staging domain named staging.contoso.com. This command auto-populates the migration table called FinanceStaging.migtable from the live GPO:

```
Cscript.exe CreateMigrationTable.wsf c:\migtables\FinanceStaging.migtable
/GPO: "Finance OU Desktop Policy" /domain:staging.contoso.com
```

To create a migration table from the backup of this GPO instead of from the live copy, simply add the /BackupLocation option to the command above, and provide a folder path that contains the backup copy of the GPO. Note that if you use the /BackupLocation option and there is more than one backup GPO located in that folder path, all available backed up GPOs will be used to populate the migration table.

**Final Preparations for Deployment**

As a final step before your production deployment, you should back up your staging GPOs. A backup is required if you are using a GPO import to perform your migration from staging to production. This method is required in cases where your staging environment is in a forest that is separate from and not trusted by your production domain, or where you need to update an existing GPO that already exists in your production environment. You can use GPMC to back up one or more GPOs or, you can use the script BackupGPO.wsf to back up a single GPO, or use the script BackupAllGPOs.wsf to back up all GPOs in the staging domain. To back up a GPO by using GPMC, in the console tree, right-click the GPO and select **Backup** from the menu.

To back up a GPO by using BackupGPO.wsf, run the script from the Scripts folder in the GPMC installation folder. The following command-line syntax backs up the GPO Finance OU Workstation Security Policy in the domain staging.contoso.com to the folder c:\gpobacks:

```
Cscript.exe backupgpo.wsf "Finance OU Workstation Security Policy"
c:\gpobacks /comment:"Backup prior to prod" /domain:staging.contoso.com
```

The preceding syntax includes a comment that indicates the purpose of the backup.

# Deploying Staged GPOs to the Production Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Once you have built your staging environment, synchronized it to your production environment, tested new and changed GPOs, and created migration tables, you are ready to perform the actual production deployment. Figure 3.9 shows this final step in the process.

**Figure 3.9 Deploying Staged GPOs to the Production Environment**



# Deployment Precautions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To ensure uninterrupted service to your users, it is a good idea to observe some precautions when migrating staged GPOs to your production environment. Although migrating new GPOs is typically a quick process that does not adversely impact production users or computers, it is prudent to avoid making such a change until the least possible number of users will be affected. Typically this might be during off hours, when users are not active on the network.

Remember that when a GPO is updated, the update is performed first against the Domain Controller that is currently targeted by GPMC for a particular domain. If you are using GPMC to perform the migration, you can click the **Domains** item in the console tree to see which DC is currently being used for each domain under management. To change the DC, right-click the domain name and choose the **Change Domain Controller** from the menu before migrating your changes.

**GPO Replication**

Keep in mind that GPO changes propagate according to your Active Directory replication and Sysvol File Replication service topologies, and therefore might take an extended period of time to replicate to all locations in a worldwide Active Directory deployment. Also keep in mind that a GPO is composed of two parts — the portion that is stored and replicates as part of Active Directory, and the portion that is stored and replicates as part of Sysvol. Because these are two separate objects that need to replicate across your network, both need to be synchronized before the new GPO is applied.

You can view the replication status on a given DC by using GPMC. From the Group Policy Objects node in GPMC, click a GPO to check, and then select the **Details** tab in the details pane. If the GPO is synchronized on that DC, the Active Directory and Sysvol version numbers will be identical for user and computer configuration. However, the user version numbers do not need to match the computer version numbers.

**Requirements for Performing the Deployment**

The primary requirement to keep in mind as you prepare to deploy your staged GPOs to your production environment is whether you have sufficient permissions on the destination GPOs. You typically need only read access to the source domain to complete a deployment. Depending on the configuration of your staging environment, you might need to take some specific steps prior to migration. If you are performing a copy operation, you will need to have sufficient permissions to create a new GPO in the destination domain. If you are importing a backup GPO, you will need to be able to read the backup files, wherever they might be located, and then have sufficient permissions to modify an existing GPO in the destination domain that is the target of the import operation. Finally, you should ensure that the migration table that you created for each GPO that requires one is stored where it is accessible to you while performing the migration. The following checklist summarizes the items to verify before running the migration:.

- **For a copy operation:** ensure that the destination domain is trusted by the source domain and that you have GPO Creation permissions on the destination domain.

  You can confirm GPO Create permissions on a domain by using GPMC. Click the **Group Policy Objects** node and, in the details pane, select the **Delegation** tab to see which users or groups can create new GPOs in the domain.

- **For an import operation:** ensure that you have access to the backup GPO files and that you have GPO Edit Settings permission on the destination GPO.

- **If you are using a migration table (.migtable):** ensure that you have access to the file from GPMC.

# Deployment Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following two examples illustrate deploying GPOs from staging to production environments. In the first example, the staging domain is located in the same forest as the production domain. In the second example, the staging domain is in a separate forest that is not trusted by the production domain. If you use a separate staging forest that is trusted by the production domain, the steps are the same as in the first example, where the staging domain is part of the production forest.

**Staging to a Production Domain in a Single Forest or from a Trusted Staging Forest**

When the staging domain is part of your production forest or you have a separate staging forest that is trusted by your production domain, your deployment method depends on whether the GPO is new or changed. If the GPO is new and does not exist in the production domain, use the copy method to deploy the new GPO. If you are deploying an update to an existing GPO, then you must to use the import method to update the production GPO's settings with those from the backup staging GPO.

In this example, you will deploy a new GPO named **Sales OU Workstation Security Policy** from the staging domain to the production domain by using GPMC. Figure 3.10 illustrates the staging and production domain configuration and shows the accompanying migration table.

**Figure 3.10 Migration from Staging to Production Domain**




| Source Name | Source Type | Destination Name |
|---|---|---|
| Sales Admins@staging.contoso.com | Domain Global Group | Sales Admins@contoso.com |
| Authenticated Users | Unknown Group | <Same As Source> |
| WORKSTATION1$@staging.contoso.com | Computer | <None> |
| Domain Admins@staging.contoso.com | Domain Global Group | <Map by Relative name> |
| Administrators | Unknown Group | <Same As Source> |
| Enterprise Admins@staging.contoso.com | Domain Global Group | <Map by Relative name> |
| Sales Users@staging.contoso.com | Domain Global Group | Sales Users@contoso.com |
| sales users | Unknown Group | Sales Users@contoso.com |
| \\server1\packages | UNC Path | \\prodserver\packages |

Before beginning the deployment load both the source and destination domains in GPMC. If you are copying from a separate trusted forest, open both forests in GPMC.

**To deploy a new GPO using the copy method**

1. In GPMC, in the staging domain, right-click the GPO that you plan to migrate and choose **Copy** from the menu.

2. In GPMC, in the production domain, right-click **Group Policy Objects** and choose **Paste** from the menu or drag-and-drop the source GPO onto the Group Policy Objects node of the destination domain. The Cross-Domain Copying Wizard starts.

3. On the **Cross-Domain Copying Wizard** page, click **Next**, select the option to Preserve or migrate the permissions from the original GPOs, and then click **Next**.

   If you choose the first option, **Use the default permissions for new GPOs**, this GPO will receive the default permissions that would be applied to any new GPO in the production domain. For this example, you want to use the second option, **Preserve or migrate the permissions from the original GPOs**, because that option lets you use a migration table to map the DACL on the staging GPO to its production equivalents.

   The wizard scans the source GPO to determine any security principal or UNC path mapping requirements. After these are found you can advance the wizard to the next step.

4. On the **Migrating References** page of the **Cross-Domain Copying Wizard**, select the option to either copy the security principals exactly as they exist in the source GPO to the destination or to use a migration table to map references and click **Next**.

The wizard gives you two choices for performing the migration. The first choice, **Copying them identically from the source**, leaves all security principals and UNC paths in the new GPO exactly as they are in the source. Because you are migrating a new GPO from the staging environment to the production environment, you need to choose **Using this migration table to map them to new values in the new GPOs** instead. This option lets you choose a migration table to use as part of the deployment. You further have the option to **Use migration table exclusively...** If you check this option the wizard will attempt to map all security principals and UNC paths using the migration table you specify. If a security principal or UNC path that exists in the source GPO is not present in the migration table, then the entire migration fails. This is useful to ensure you have accounted for all security principals and UNC paths in your migration table. After you make your choice move to the next step in the wizard, which is the summary screen prior to migration.

5. On the **Completing the Cross-Domain Copying Wizard** page, you can see your chosen migrations options prior to performing the actual migration. This page shows the source GPO you have chosen, the option you chose for migrating permissions and the migration table (if any) that you will be using. If you need to change an option, you can select **Back** to go back through the wizard and change your choices.

6. When you click **Finish**, the migration of the staging GPO begins. Keep in mind that the new GPO is being created in the production domain but will not yet be linked to any container objects.

7. When the wizard completes the copy operation, right-click the Active Directory site, domain, or OU to which you want to link the copied GPO, and then select **Link an Existing GPO** from the menu. In the **Select GPO** dialog box, select the GPO you just copied.

   After you link the new GPO, and replication is complete, the GPO is live in the production domain.

**Using a Script to Perform a Copy Deployment**

You can also perform a copy deployment by using the script CopyGPO.wsf, which is installed in the Scripts folder in the GPMC installation folder. This script copies a GPO between staging and production domains in a single command. To perform the same copy operation as in the foregoing procedure, use the following command:

```
Cscript CopyGPO.wsf "Sales OU Workstation Security Policy" "Sales OU
Workstation Security Policy" /SourceDomain:staging.contoso.com /TargetDomain:contoso.com
/SourceDC:staging-dc1 /TargetDC:prod-DC1 /migrationtable:c:\migtables\SalestoProd.migtable
/CopyACL
```

The first two arguments in this command specify the same name for both the source and target GPO. The next four arguments specify the source and target domain names and a domain controller in each domain. The /migrationtable argument specifies the migration table to use and the /CopyACL argument is used to preserve the DACL from the source GPO and use the specified migration table to map the source DACLs to their production domain equivalents.

**Deploying to a Production Domain from an Untrusted Staging Forest**

If you are deploying a GPO from a staging forest that is not trusted by the production forest, the only choice for deployment is an import operation. You can also use an import to deploy an update to an existing GPO in the production domain even if a trust relationship exists between the staging and production domains.

**Import Operation Prerequisites**

Before performing the deployment in this example, there are some prerequisites that you should be aware of:

- If you are deploying a new GPO by using GPMC, you need to create a new, empty GPO in your production domain that can act as a target for the import operation. Remember that the GPMC import operation works by importing the settings from a backup GPO into an existing destination GPO. However, you can also use the script ImportGPO.wsf to create a new GPO automatically, as part of the import process.

- Before beginning the import, make sure you back up the GPOs from your staging domain that you plan to deploy to production. This is necessary because the import operation uses backup GPOs rather than live GPOs.

- If you are using the GPMC UI rather than a script to perform the import, you have the opportunity to back up the current production GPO before completing the import. You should always back up an existing production GPO before deploying a new version in case there are problems with the deployment. In that event, you can perform a restore operation from GPMC to restore the previous version of the GPO.

When these prerequisites have been satisfied , you are ready to deploy a new GPO into the production environment using the

import method.

**To deploy a new GPO to the production domain using the Import method**

1. In GPMC, in the production domain, right-click the GPO to be updated and select **Import Settings** from the menu.

   The Import Settings Wizard starts.

2. In the **Import Settings Wizard** on the **Backup GPO** page, select **Backup** to back up the existing production GPO prior to performing the import, and then click **Next**. When you select the **Backup** button, you're given the option of choosing a location to store your backup and a name to reference it by. Select the Backup button to start the backup process.

   Once the wizard finishes backing up your production GPO, from the **Backup location** page, you need to choose the folder that contains the backup of the staging GPO that you want to import. You must have access to the folder where you backed up your staging GPOs. If your backups were made on a server in your staging forest, you might need to map a drive to that folder from the workstation where you are running the import operation, using credentials from the staging forest.

3. After you select the backup folder, choose Next to present the **Source GPO page**. This page presents a the list of backup GPOs, from which you can choose the staging GPO you want to import. Once you make the choice, click **Next**.

4. On the **Migrating References** page, you will see the same options for migrating security principals as in the **Cross-Domain Copying Wizard** used in the previous copy method. You can choose to either use the existing security principal information contained in the staging GPO or apply a migration table against it. Because you are deploying a GPO from staging domain that does not have a trust relationship with the production domain, you must use a migration table to migrate security principal and UNC path information. Otherwise, the security principals and UNC paths referenced in the untrusted forest cannot be resolved by the production domain. To select a migration table to use, choose the **Using this migration table to map them in the destination GPO** option and then select a path to the migration table you created for this migration. You can select the **Use migration table exclusively...** option to only import the GPO if all security principals found in the backed up version are accounted for in the migration table.

5. Select **Next** to view the **Completing the Import Settings Wizard** page which presents a summary of your chosen migration options. You can select **Back** if you need to go back to change some of the options prior to the import. Select **Finish** to start the import operation.

   If you created the production GPO from scratch to perform this import, you will need to link the new GPO to the appropriate container object. To link the GPO, in GPMC, in the production domain, right click the Active Directory site, domain, or OU to which you want to link the imported GPO and choose Link an Existing GPO from the menu. Once you link the new GPO and replication is complete, the GPO is live in the production domain.

**Using a Script to Perform an Import Deployment**

You can also perform an import deployment by using the script ImportGPO.wsf, which is installed in the Scripts folder in the GPMC installation folder. This script lets you import a backup GPO into your production domain. If the target GPO does not yet exist, the script also lets you create a new GPO to receive the import as part of the process. To perform the same import operation as described in the previous procedure, type the following command:

```
Cscript ImportGPO.wsf c:\gpobacks "Sales OU Workstation Security Policy" "Sales OU Workstation Security Policy" /CreateIfNeeded /MigrationTable:c:\migtables\salesprod.migtable /Domain:contoso.com
```

The first argument in this command specifies the location of the backup GPO files. The second argument specifies the name of the backed up GPO to import from (you can instead provide the Backup ID, which is a 128-bit GUID value generated by the backup utility to uniquely identify the backup). The third argument specifies the name of the destination GPO to import into. The /CreateIfNeeded argument indicates that if the destination GPO does not yet exist, it should be created before performing the import. The /MigrationTable argument specifies the path and name of the migration table file. The /Domain argument provides the DNS name of the destination domain.

**Rollback**

In the event that you have a problem with a GPO after you deploy it from the staging environment to the production environment, the best way to roll back the deployment is to use the backup GPO you created in the previous deployment steps to restore the original GPO. You can also use the RestoreGPO.wsf script to perform the restore process. As part of your deployment, it is a good idea to create a set of scripts that can perform an automated rollback of all of your changes using RestoreGPO.wsf. In the event that you need to perform a rollback, the script is ready and available to use with minimal user disruption.

# Additional Resources for Staging Group Policy Deployments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this chapter.

**Related Information**

- "Designing a Group Policy Infrastructure" in this book for detailed information about Group Policy and GPMC features and capabilities.
- "Planning a Managed Environment" in this book for more detailed information about using IntelliMirror to create a managed environment.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, click to select the **Search in title only** check box.

- "Group Policy" in Help and Support Center for Windows Server 2003.
- Help in GPMC for detailed information about using GPMC to help deploy Group Policy and troubleshoot your deployment.

# Deploying Security Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Administrators who deploy enterprise networks must design and implement many aspects of security. As part of your security strategy to manage servers and workstations, you can deploy Group Policy with the Active Directory® directory service to specify options for Internet Protocol security (IPSec), security settings, software restrictions policies, and wireless network policies.

**In This Chapter**

Overview of Security Policy Deployment

Designing Security Policy

Configuring Security Policy

Additional Resources for Deploying Security Policy

**Related Information**

- For more information about security in the Microsoft® Windows® Server 2003 operating systems, see the security chapters in *Designing and Deploying Directory and Security Services* of this kit.
- For more information about Internet Authentication Service (IAS) and wireless networks, see "Deploying IAS" and "Deploying a Wireless LAN" in *Deploying Network Services* of this kit.
- For more information about IPSec, see "Deploying IPSec" in *Deploying Network Services Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Overview of Security Policy Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your organization's data is subject to various types of risks, including user errors and malicious attacks. Attackers can gain access to the system, rendering systems useless or disrupting services. Attackers can also modify, delete, or steal information. As you analyze the potential security threats and how to handle those threats, consider how much time, money, and effort you want to invest in developing security strategies and controls for your organization.

Group Policy is one of the main tools that you can use to specify security policy for servers, workstations, and domain controllers. By using Group Policy, you can create and apply security policies to simplify and centralize the process for configuring and managing security for servers running Microsoft® Windows® 2000 Server and Windows Server 2003 and workstations running Microsoft® Windows® XP Professional and Microsoft® Windows® 2000 Professional.

You can use Group Policy to specify configuration options for IPSec, security settings (such as user rights, password policies, file and registry access control lists [ACLs], and service startup modes), software restriction policies, and wireless network configurations.

Before you deploy security policies by using Group Policy in an Active Directoryenvironment, you must complete your deployment design for the following areas:

- **IT requirements.** Understand how domains and organizational units (OUs) are created and managed in your organization and obtain a list of the administrative owners.

- **Active Directory and security.** You must have Active Directory to use Group Policy. While you are deploying Active Directory, design the OUs in your organization to delegate administration authority and to implement Group Policy. Create a security plan that includes authentication, password requirements, resource security (such as file security), auditing, and security group membership. For more information about Active Directory and creating a security plan, see *Designing and Deploying Directory and Security Services* of this kit.

# Security Policy Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To provide security for their organizations, administrators must plan and develop a strategy that protects the availability, integrity, and confidentiality of the organization's data. Implementing security settings by using Group Policy ensures that any changes that you make to a Group Policy object (GPO) setting affects all servers and workstations in the OU or domain to which the GPO is linked.

**Figure 4.1 Planning a Security Strategy**



# Security Policy Concepts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy is the infrastructure in Active Directory that enables centralized management of user and computer settings. Use Group Policy to define security configurations for groups of users and computers, including the following security settings:

- Account policies (password policy, account lockout policy, and Kerberos policy)
- Local policies (user rights assignment, audit policy, and security options)
- IPSec policies
- Software restriction policies
- Wireless network configurations
- File and registry ACLs
- Service startup modes
- Public key policies

For more information about individual security policies, see "Security Settings" in Help and Support Center for Windows Server 2003.

The Group Policy settings that you create are contained in a GPO. By associating a GPO with selected Active Directory system containers — sites, domains, and OUs — you can apply the GPO's policy settings to the users and computers in the Active Directory containers. Although some security settings affect user accounts, most settings are controlled by computer settings that must be applied to computers accounts; only software restriction policies and public key policies can be applied to user accounts.

For more information about Group Policy design, see "Designing a Group Policy Infrastructure" in this book. For information about the mechanics of Group Policy, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Group Policy Management Console

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Group Policy Management Console (GPMC) is a tool that permits you to manage Group Policy for multiple domains and sites in one or more forests. This chapter assumes that you are using GPMC for security policy deployment and management.

GPMC is not included with Windows Server 2003. To obtain GPMC, see the Group Policy Management Console (GPMC) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Starting GPMC**
To start GPMC, click **Start**, click **Programs**, click **Administrative Tools**, and then click **Group Policy Management**. GPMC can be used on computers that are running Windows XP Professional and Windows Server 2003. The procedures in this chapter use GPMC.

For more information about using GPMC, see Group Policy Management Console Help. To view GPMC Help, open GPMC, right-click **Group Policy Management**, and then click **Help**.

# Security Configuration Manager Tool Set

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows 2000 and Windows Server 2003 provide an extensive set of security technologies and management tools, including the Security Configuration Manager tools, support for IPSec, the Kerberos V5 authentication protocol, public key infrastructure (PKI), Encrypting File System (EFS), and smart cards.

The Security Configuration Manageris a set of tools that includes:

- **Security templates.** Use these .inf files to create customized security policies. You can import a predefined security template to a Group Policy object, and then customize it. For more information about security templates, see "Security Templates" in Help and Support Center for Windows Server 2003.

- **Security Configuration and Analysis.** Use this tool with a security template to analyze or configure computer security settings for the local computer. For more information about security configuration and analysis, see "Security Configuration and Analysis" in Help and Support Center for Windows Server 2003.

- **Group Policy Object Editor.** Use this tool to manage security settings in a Group Policy object for a site, domain, or OU. The Group Policy Object Editor is not part of Security Configuration Manager. However, the **Security Settings** node, which is found in the Group Policy Object Editor, is part of Security Configuration Manager. For more information about security settings, see "Security Settings extension to Group Policy" in Help and Support Center for Windows Server 2003.

- **Local security policy.** Use this tool to manage security settings on the local computer, particularly for clients who are running Windows 2000 or Windows XP Professional on Microsoft® Windows NT® or non-Microsoft networks. For more information about local security policy, see "Local Security Policy" in Help and Support Center for Windows Server 2003.

- **Secedit.exe.** Use this command-line tool to automate security configuration tasks. For more information about Secedit, see "Secedit" in Help and Support Center for Windows Server 2003.

# Designing Security Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use GPMC to centralize the process of deploying and managing Group Policy–based security for servers running Windows 2000 Server and Windows Server 2003 and clients running Windows 2000 Professional and Windows XP Professional. Use Group Policy–based security policies to deploy and manage the following security areas:

- **Security settings.** These settings are used to define values for various security-relevant operating system parameters, such as password policy, user rights assignment, audit policy, registry values, file and registry ACLs, and service startup modes.

- **IPSec policies.** These policiesare used to configure IPSec services for authenticating or encrypting network traffic. An IPSec policy consists of a set of security rules, and each security rule consists of an IP filter with an action.

- **Software restriction policies.** These policies are used to help protect computers from code that is not trusted by identifying and specifying which applications are permitted to run.

- **Wireless network policies.** These policiesare used to configure settings for the Wireless Configuration Service, a user-mode service that operates on each of the IEEE 802.11 wireless network adapters that are installed on a computer.

**Figure 4.2 Designing Security Policy**




**Important**

- As with all Group Policy settings, you must fully test your implementation before you deploy your security settings to your production environment. For more information about Group Policy staging and testing, see "Staging Group Policy Deployments" in this book and "Designing a Test Environment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

# Designing IPSec Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use IPSec policies to filter, authenticate, or encrypt network traffic. An IPSec policy consists of a set of security rules. Each security rule consists of an IP filter with a filter action to permit or block traffic or to negotiate security. If you define a filter action to negotiate security and encrypt a specific type of traffic, you must configure additional settings, such as an authentication method for establishing trust between IPSec peers. IPSec can be configured in many different ways; therefore, you must understand IPSec in detail and test your policy configurations in lab environments before you attempt to use policies in a production environment.

For example, an IPSec policy that secures end-to-end traffic cannot be assigned to a single computer; however, an IPSec policy that only uses permit or block filtering *can* be assigned to a single computer. You must fully understand IPSec before you assign an IPSec policy. If an IPSec policy is not correctly configured and assigned, it might inadvertently block all communication to a computer.

| **Important** |
| --- |
| The IPSec filter lists, actions, and default policies included with Windows Server 2003 are not appropriate to use under any circumstances. They are intended *only* as examples to demonstrate behavior. Administrators must design and build their own policies, and customize the policies for their particular situation and security requirements. |

For more information about IPSec policy, see "Deploying IPSec" in *Deploying Network Services* of this kit, and the *Networking Collection* of the*Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit). For more information about how to use IPSec and Group Policy, see "Assign or unassign IPSec policy in Group Policy" and "Creating, modifying, and assigning IPSec policies" in Help and Support Center for Windows Server 2003.

Assigning an IPSec policy to a GPO records a pointer to the IPSec policy that is inside the GPO attribute ipsecOwnersReference. The GPO itself contains only a Lightweight Directory Access Protocol (LDAP) distinguished name (also known as DN) reference to the IPSec policy. Group Policy is used only to deliver the policy assignment to the computer's IPSec service. The computer's IPSec service then retrieves the IPSec policy from Active Directory, maintains a current cache of the policy locally, and keeps it current by using a polling interval that is specified in the IPSec policy itself. Because the IPSec policy itself is not stored inside the GPO, its settings can be assigned to and shared by many GPOs. Consider the following characteristics when you plan for the behavior and management of Group Policy for IPSec.

**Plan IPSec Policy to Fit Your Active Directory Structure**

By default, in Windows Server 2003, Active Directory restricts Read permissions on the IP Security Policies container to a greater degree than in Windows 2000. If you are deploying a new installation of Windows Server 2003 Active Directory, be aware that IPSec policies cannot be read by computers in child domains, even though the GPO can be read by computers in the child domain. The domain administrator must explicitly allow permissions for computers in child domains to read the IPSec policy from the parent domain.

For clean Windows Server 2003 installations of Active Directory, the Group Policy Creator Owners administrative group does *not* have permission by default to create or modify IPSec policies. By default, only members of the Domain Admins group have this permission, and the Group Policy Creator Owners group has read-only permission. Upgrades of Windows 2000 Active Directory domains to Windows Server 2003 domains do not change permissions on existing IPSec policy objects.

Domain-based IPSec policy objects are stored in the IP Security Policies container in Active Directory, which is separate from the GPOs to which IPSec policies are applied. The domain administrator must grant permissions to the IP Security Policies container for other delegated administrators to administer IPSec policies. Standard delegation tools cannot be used to delegate permissions to administer IPSec policies. Instead, domain administrators must use Active Directory Service Interfaces (ADSI) Edit tool for this purpose.

ADSIEdit is a Microsoft Management Console (MMC) snap-in that domain administrators can use to edit objects in the Active Directory database. When domain administrators delegate permissions to others to administer IPSec policies, the delegated administrators must have Full Control permissions to all IPSec policy objects in the IP Security Policies container.

After an IPSec administrator creates an IPSec policy, a member of the Group Policy Creator Owners group or other delegated owner of the GPO can assign the IPSec policy to the appropriate GPOs.

**Caution**

To modify access control lists (ACLs) on the IP Security Policies container, domain administrators should use the tools that are provided with the Windows 2000 or Windows Server 2003 operating system CD. When domain administrators modify ACLs on the IP Security Policies container, they must specify that all modifications are also propagated to all child objects. Incorrect modification of the ACLs on the IP Security Policies container, failure to propagate modifications to the child objects of this container, or both, can result in the failure of IPSec policy settings to be properly applied.

ADSIEdit is one of the Windows Support tools included on the Windows 2000 and Windows Server 2003 operating system CDs. You can install these tools by running the Support Tools Setup program, Suptools.msi, from the \Support\Tools folder.

For more information about permissions for the IP Security Policies container, see article 329194, "IPSec Policy Permissions in Windows 2000 and Windows Server 2003" in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkID=106552.

Unlike most Group Policy settings, multiple IPSec policies that are assigned in different GPOs are not merged. The last GPO in the Active Directory directory tree (closest to the computer object) that contains an IPSec policy assignment is the one that takes effect. Because the actual IPSec policy that is applied on the computer depends on the network adapter configuration of that computer, using the IP Security Monitor snap-in or Netsh IPSec context (the **netsh ipsec static show** and **netsh ipsec dynamic show** commands) is the only way to view the detailed IPSec policy settings as they are applied on the computer.

**Note**

Netsh is a command-line tool for configuring networking components on the local computer or on remote computers running Windows 2000, Windows XP Professional, or Windows Server 2003. The Netsh IPSec context is only available on Windows Server 2003. For more information about using the Netsh IPSec context, see "Netsh commands for Internet Protocol security (IPSec)" in Help and Support Center for Windows Server 2003.

Group Policy inheritance in Active Directory cannot be blocked for a specific GPO without affecting other settings in the GPO. If you must control IPSec policy inheritance, create a new GPO that is dedicated exclusively to deploying and assigning IPSec policy.

For recommendations on the uses of IPSec, see "Special IPSec considerations" in Help and Support Center for Windows Server 2003.

**Ensure That Your New IPSec Policies Are Applied**

When you deploy an IPSec policy by using Group Policy, Group Policy (Winlogon) polling detects changes in policy assignments within the GPOs. Additionally, during the Group Policy poll, IPSec checks whether its policy has changed regardless of whether the GPO has changed. By using the **gpupdate /target:computer** command, it detects that an IPSec policy is newly assigned in a GPO and causes the IPSec service to apply that policy. Additionally, it causes IPSec to detect if the IPSec policy has changed even though the GPO may not have changed. If an assigned IPSec policy has not changed, then no IPSec policy changes are applied to the computer. Using the **gpupdate /target:computer/force** command causes the IPSec policy agent to reload the assigned IPSec policy regardless of whether the GPO or the IPSec policy has changed.

IPSec also uses its own polling mechanism to detect a change in an IPSec policy that is already assigned (for example, a filter list change). This polling mechanism provides compatibility with previous versions of Windows, but it is no longer necessary because IPSec detects changes to IPSec policies during the Group Policy interval. However, it might be necessary to configure the IPSec polling interval to be shorter than the Group Policy polling interval during IPSec policy change rollout. Configuring the IPSec polling interval permits you to remove any changes that might have been made if problems occur or if you are using an IPSec domain policy to respond to security incidents. For more information, see article 813878, "How to Block Specific Network Protocols and Ports by Using IPSec," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

To completely delete and reload all existing IPSec policy configurations that affect communication, you must stop and then restart the IPSec service by using the **net stop policyagent** and **net start policyagent** commands. If your remote access service is configured to use Layer Two Tunneling Protocol (L2TP)/IPSec, you must restart the Routing and Remote Access service after the newly restarted IPSec service is running. Also, if an L2TP/IPSec virtual private network (VPN) client connection is connected when the IPSec service is restarted, the VPN connection is interrupted and must be reconnected.

IPSec policies cannot be applied by using security templates, and they cannot be analyzed by the Security Configuration Manager.

**IPSec Differences When Using the Group Policy Management Console (GPMC)**

Because the IPSec policy itself is not stored inside the GPO, note the following differences when using GPMC to manage GPOs that have IPSec policies applied:

- You can use the GPMC **Backup** and **Restore** capabilities to store information about which IPSec policies are assigned to specific Group Policy objects. However, because the IPSec policies themselves aren't stored in GPOs, you must use the **Export Policies** and **Import Policies** commands of IP Security Policy Management snap-in to back up and restore the IPSec policies themselves.

- GPMC **Delegation of rights** and **Security Filtering** permissions only apply to the GPO, not to the IPSec policy that is assigned in the GPO. Thus delegation of edit rights within GPMC only allows a user to assign or unassign an existing IPSec policy in the specific GPO, but only if the user also has read access rights to the IPSec policy. Delegation of rights to create, edit, or delete IPSec policies must be done on the IPSec policies container.

- GPMC **Import Settings** can only import an IPSec policy assignment and cannot be used to import IPSec policies into a GPO. The IP Security Policy Management snap-in provides export and import capabilities for the IPSec policy store. When using GPMC to import or copy a GPO into another domain or forest, the IPSec policy assignment is invalidated for the new GPO. The administrator must assign an IPSec policy in that domain or forest to the GPO.

- The GPMC Group Policy Results wizard is used to show which GPOs are applied to a computer, which includes the IPSec policy assignments. To find out which IPSec policy is assigned to a specific computer, after running the wizard, right-click the computer node, and then select **Advanced View**.

Additionally, on computers running Windows XP, IPSec does not provide Resultant Set of Policy (RSoP) information. GPMC Group Policy Results shows the GPO being processed, but does not show which IPSec policy is assigned. Use **netdiag /test:ipsec** to view the assigned IPSec policy.

| **Note** |
| --- |
| Additional IPSec design and deployment information is not included in this chapter. For more information about deploying IPSec, see "Deploying IPSec" in *Deploying Network Services* of this kit. For more information about how to create IPSec policies, see "Creating, modifying, and assigning IPSec policies" in Help and Support Center for Windows Server 2003. |

# Designing Security Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses the recommendations and considerations for using security templates and implementing security settings. Before implementation make sure to consider the following design elements. For more information about importing security templates for domain controllers, servers, or workstations, see "Importing Security Templates and Modifying Security Settings in a GPO" later in this chapter.

**Use the NTFS file system**
You cannot secure Windows-based computers that are installed on file allocation table (FAT) file systems. All the security design and deployment recommendations in this chapter assume that you use the NTFS file system on all computers that you need to secure.

**Complete all policy settings before applying them**
Make sure that all the settings of your policy are in place before you apply the policy to any GPOs. If you apply the policy in parts, and a user refreshes the policy before all the parts are in effect, this can adversely affect the user's computer when configuring software restriction policies.

**Avoid editing the Default Domain GPO and the Default Domain Controllers GPO**
The Default Domain GPO and Default Domain Controllers GPO are vital to the health of any domain. The Default Domain GPO provides the basic domain encryption key, and if that policy is removed or deleted, users cannot unencrypt their files. There are two exceptions:

- Editing the Default Domain GPO to define account policies, including password policies, account lockout policies, and Kerberos policies.
- Editing the Default Domain Controllers GPO to define user rights assignment and audit policy for the domain controllers OU.

**Set domain account policy in the Default Domain GPO**
When you set account policies (including password policies, account lockout policies, and Kerberos policies) in Active Directory, there can only be *one* domain account policy throughout all servers, workstations, and domain controllers in the domain. The policy is the account policy that is applied at the root domain of a domain tree. Although account policies affect user accounts, the policies are defined on computers.

**Do not link to a GPO in another domain**
Avoid linking to a GPO in another domain because this can degrade performance.

**Consider precedence of policy application when multiple GPOs are applied**
Because a computer can have more than one GPO applied to it, security settings can conflict. From highest to lowest, the settings apply in the following order of precedence: OU, domain, site, and local computer. For example, policies that are defined in Active Directory at the OU, domain, or site level always override the local security policy for a computer if there is a conflict. If the same computer is a member of an OU, the organizational unit policy overrides all other settings. If the computer is a member of nested OUs, the OU that immediately contains the computer takes precedence.

For nested OUs, the GPO that is linked closest to the OU takes precedence. The precedence rules that apply to restricted groups also apply to other Group Policy settings when there is a conflict. For example, if a computer is a member of OU A, which is nested in OU B, and both OUs define Power Users as a restricted group, then the definition of Power Users according to OU A takes precedence on the computer.

To find out which policies are currently applied to a specific computer, use the **Group Policy Results** node in the GPMC snap-in.

**Plan Event log size and wrapping according to business and security requirements**
Define Event log size and overwrite of events logs (also known as *log wrapping*) to match the business and security requirements of your organization's security plan. Implement these Event log settings at the site, domain, or OU level to take advantage of Group Policy settings.

**Understand the use of Restricted Groups**

If you create a Restricted Groups policy for a group, any users and groups that are not specified as members of the group within the policy are removed from the group. For example, if you create a Restricted Groups policy for the local Administrators group, and the newly created policy specifies only the Domain Admins group as members, all other members of the local Administrators group (including any local accounts) are removed from the local Administrators group when the policy is applied. Note that if the Restricted Groups members are defined in more than one GPO, only the members that are defined in the GPO with the highest precedence are applied. This also applies to the groups that the group can be a member of.

**Be aware that security settings can persist**
For Windows Server 2003 and Windows XP, security settings might persist even if the setting is no longer defined in the GPO that originally applied it. This occurs under the following conditions:

- The setting was not defined for the local computer at the time that the policy setting was applied.
- The setting is for a registry object.
- The setting is for a file system object.

Windows 2000 security settings may persist even if the setting is no longer defined in the GPO that originally applied it. This occurs under the following conditions:

- The setting has not been defined for the local computer at the time that the policy setting was applied.
- The setting is for a registry object.
- The setting is for a file system object.
- The setting is for a service.
- The setting is for a Restricted Groups policy.
- The setting is an Event log setting.

All settings that are applied through local policy or a GPO are stored in a local database on your computer. Whenever a security setting is modified, the computer saves the security setting value to the local database. The database retains a history of all the settings that have been applied to the computer. If a policy defines a security setting and then no longer defines that setting, the setting reverts to the previous value in the database. If a previous value does not exist in the database, the setting remains defined as is. This behavior is sometimes called *tattooing*.

Any other settings that persist maintain the values that are applied through the policy until that setting is set to a different value.

# Using Custom Security Templates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Security templates are files that represent a security configuration. These files can be imported to a GPO, applied to a local computer, or used to analyze security. To edit individual security settings on a domain, site, or OU, administrators use the Group Policy Object Editor.

You can view the templates as text files or you can use the Security Templates snap-in to view the settings in each template. Security templates can be transported, imported, and exported. All of the security attributes in Table 4.1 can be contained in a security template.

**Table 4.1 Security Template Attributes**

| Security Attribute | Description |
|---|---|
| Account Policies | Password policy, account lockout policy, and Kerberos policy |
| Local Policies | Audit policy, user rights assignment, and other security options, including numerous security-related registry values |
| Event Log | Application, system, and security Event log settings |
| Restricted Groups | Membership of security-sensitive groups |
| System Services | Startup and permissions for system services |
| Registry | Permissions and auditing entries for registry keys |
| File System | Permissions and auditing entries for folders and files |

# Selecting Predefined Security Templates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 includes a set of predefined security templates that you can use to create customized security policies. You can customize the templates by using the Security Templates snap-in. It is recommended that you save one of the predefined templates under another name, and then modify it to meet the needs of your organization. You can use a template to configure either a single computer or an entire domain. You can also use a security template — with the Security Configuration and Analysis snap-in — as a baseline for analyzing the local computer for potential security problems or policy violations.

The default storage location of the predefined security templates is %systemroot%\Security\Templates. Only administrators have permissions to modify the templates in this folder; users without administrative permissions can save their templates to another folder. This folder includes the following templates.

**Setup Security.inf**
Setup security.inf is a computer-specific template that represents the default security settings that are applied when the operating system is installed, including the file permissions for the root of the system drive.

The Setup security.inf template is created for each computer during installation. It can vary from computer to computer and is based on whether the installation was a clean installation or an upgrade. This template can be used on servers and client computers; it cannot be applied to domain controllers. You can apply portions of this template for disaster recovery purposes.

Do not apply Setup security.inf using Group Policy. It contains a large amount of data and can seriously degrade performance if it is applied by using Group Policy. Degraded performance can occur because policy is periodically refreshed, and a large amount of data then moves through the domain.

It is recommended that you apply the Setup security.inf template in parts by using the Secedit command-line tool. For more information about Secedit, see "Secedit" in Help and Support Center for Windows Server 2003.

**DC Security.inf**
This template is created when Active Directory is installed onto a server. It reflects file, registry, and system service default security settings. Reapplying this template resets these settings to default values. Note that this template might overwrite permissions on new files, registry keys, and system services that are created by other applications. It can be applied by using the Security Configuration and Analysis snap-in or the Secedit command-line tool.

**Compatws.inf**
Default permissions for workstations and servers are primarily granted to three local groups: Administrators, Power Users, and Users. Of the three, the Administrators group has the most permissions, while the Users group has the least. Because of this, you can significantly improve security, reliability, and the total cost of system ownership by:

- Making sure that end users are members of the Users group.
- Deploying applications that can be run successfully by members of the Users group.

Members of the Users group can successfully run applications that are a part of the Windows Logo Program. However, members of the Users group might not be able to run applications that do not meet the requirements of the program. If other applications must be supported, there are two options:

- Permit members of the Users group to be members of the Power Users group.
- Relax the default permissions that are granted to the Users group.

Because Power Users have additional permissions such as creating users, groups, printers, and shares, some administrators prefer to relax the default User permissions instead of permitting members of the Users group to be members of the Power Users group. This is precisely what the Compatible template is for. The Compatible template changes the default file and registry permissions that are granted to the Users group in a way that is consistent with the requirements of most applications that do not belong to the Windows Logo Program. Additionally, because it is assumed that the administrator who is applying the Compatible template does not want members of the Users group to be Power Users, the Compatible template also

removes all members of the Power Users group. For more information, see "Default security settings for groups" in Help and Support Center for Windows Server 2003.

For more information about the Windows Logo Program for Software, see the Windows Logo Program for Software link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Do not apply the Compatible template to domain controllers. For example, do not import the Compatible template to the Default Domain GPO or the Default Domain Controllers GPO.

**Secure*.inf**
The Secure templates define enhanced security settings that are least likely to impact application compatibility. For example, the Secure templates define stronger password, lockout, and audit settings.

The Secure templates prevent anonymous users, such as users from untrusted domains, from:

- Enumerating account names and shares.
- Performing SID-to-name or name-to-SID translations. If this is allowed, anonymous users can request user names (such as the administrator's user name) based on the user's security ID (SID) as well as requesting the SID for a user based on the user name.

The Secure templates enable server-side server message block (SMB) packet signing. By default, SMB packet signing is disabled for member servers. Because client-side SMB packet signing is enabled by default, SMB packet signing is always negotiated when workstations and servers are operating at the Secure level.

Additionally, the Secure templates limit the use of LAN Manager and NTLM authentication protocols by configuring clients to send only NTLM version 2 (NTLMv2) responses and by configuring servers to refuse LAN Manager responses.

Considerations for applying Securews.inf to a member computer include:

- All domain controllers that contain the accounts of all users that log on to the client must run Windows NT 4.0 Service Pack 4 (SP4) or later, Windows 2000, or Windows Server 2003.
- If the member computer is joined to a domain that contains domain controllers that are running Windows NT 4.0, the clocks on both the domain controllers and the member computers must be set within 30 minutes of each other.

If a client is configured with Securews.inf, it cannot connect to:

- Servers that only use the LAN Manager authentication protocol or that run Windows NT 4.0 Service Pack 3 (SP3) or earlier using a local account that is defined on the target server.
- Servers running Windows 2000 SP3 and earlier or Windows NT 4.0 using a local account that is defined on the target server unless the clock on the target server is set within 30 minutes of the clock on the client.
- Computers running Windows XP by using a local account that is defined on the target computer unless the clock on the target computer is set within 36 hours of the clock on the client.
- Servers running LAN Manager that are running in share-level security mode.

If a server is configured with Securews.inf:

- A user with a local account on the server cannot connect to the server by using that local account from a client computer that is running only LAN Manager.
- For Windows 2000 SP3 and earlier, a client using a local account on the server that is also configured to use NTLMv2 authentication cannot connect unless the clocks on the two computers are set within 30 minutes of each other.

Likewise, if a domain controller is configured with Securedc.inf, a user who has an account in that domain cannot connect to any member server from a client computer that is running only LAN Manager using that domain account.

You can configure computers running Windows NT SP4, Windows 2000, Windows XP, or Windows Server 2003 to send only NTLMv2 responses by doing at least one of the following:

- Specify this preference in the **Network security: LAN Manager Authentication Level security** option.
- Set HKLM\System\CurrentControlSet\Control\LSA\LMCompatibilityLevel to 3 or higher.

  **Caution**

  - Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

Computers that run LAN Manager include Windows for Workgroups and Microsoft® Windows® 95 and Microsoft® Windows® 98 platforms that do not have the Active Directory Client Extensions Pack installed. If the Active Directory Client Extensions Pack is installed on Windows 95 or Windows 98, those clients can use NTLMv2. Microsoft® Windows® Millennium Edition supports NTLMv2 without additional modification.

**Hisec*.inf**

The Highly Secure templates are supersets of the Secure templates that require greater levels of encryption and signing for authentication and for the data that flows over secure channels and between SMB clients and servers. For example, the Secure templates cause servers to refuse LAN Manager responses; the Highly Secure templates cause servers to refuse both LAN Manager and NTLM responses. The Secure template enables server-side SMB packet signing; the Highly Secure template requires it.

The Highly Secure templates limit the use of cached logon data, such as data stored by Winlogon and Stored User Names and Passwords.

Hisecws.inf uses Restricted Group settings to:

- Remove all members of the Power Users group.
- Make sure that only the Domain Admins group and the local Administrator account are members of the local Administrators group.

Hisecws.inf defines these group restrictions under the assumption that users only use applications that belong to the Windows Logo Program. When using those applications, neither the Compatible template nor the Power Users group is needed. Instead, users can run those applications successfully under the secure context of a member of the Users group. The Users group is defined by the default security settings of the file system and registry. Members of the Administrators group can still use any application.

Additionally, the Highly Secure templates require strong encryption and signing for the secure channel data that constitutes domain-to-member and domain-to-domain trust relationships. Thus, consideration for applying Hisecws.inf to a member computer include:

- All of the domain controllers that contain the accounts of all users that can log on to the client must be running Windows NT 4.0 SP4 or later, Windows 2000, or Windows Server 2003.
- All of the domain controllers for the domain that the client is joined to must run Windows 2000 or Windows Server 2003.

Clients that are configured with Hisecws.inf cannot connect to:

- Computers that only run LAN Manager or computers running Windows NT 4.0 SP3 or earlier using a local account defined on the target server.
- Servers running Windows 2000 SP3 and earlier or Windows NT 4.0 SP4 using a local account that is defined on the target server unless the clock on the target server is set within 30 minutes of the clock on the client.

- Computers running Windows XP using a local account that is defined on the target computer unless the clock on the target computer is set within 36 hours of the clock on the client.

- LAN Manager servers that are operating in share-level security mode.

To apply Hisecdc.inf to a domain controller, all of the domain controllers in all trusted or trusting domains must run Windows 2000 or Windows Server 2003 operating systems.

If a server is configured with Hisecws.inf:

- A user who has a local account on that server cannot connect to that server from a client that does not support NTLMv2.

- A client that has a local account on that server cannot connect to that server unless the client is configured to send NTLMv2 responses.

- All clients that want to use SMB to connect to the server must enable client-side SMB packet signing. By default, all clients running Windows 2000 and Windows XP operating systems enable client-side SMB packet signing.

If a domain controller is configured with Hisecdc.inf:

- If a user attempts to connect to member servers by using a domain account in the same domain, the connection will fail if the user's client only uses the LAN Manager authentication protocol.

- LDAP clients cannot bind with the Active Directory LDAP server unless data signing is negotiated. LDAP BIND requests using ldap_simple_bind or ldap_simple_bind_s are rejected. By default, all Microsoft LDAP clients included with Windows XP request data signing if Transport Layer Security/Secure Sockets Layer (TLS/SSL) is not already being used. If TLS/SSL is being used, data signing is negotiated.

- A user who has an account in that domain cannot connect to member servers by using that domain account unless:

    - Both the client and target server are running Windows 2000, Windows XP, or Windows Server 2003, and both can use Kerberos-based authentication instead of LAN Manager-based authentication.

    - The client is configured to send NTLMv2 responses.

**Rootsec.inf**
Rootsec.inf specifies the root permissions. By default, Rootsec.inf defines these permissions for the root of the system drive. This template can be used to reapply the root directory permissions if they are inadvertently changed, or the template can be modified to apply the same root permissions to other volumes. As specified, the template does not overwrite explicit permissions that are defined on child objects; it propagates only the permissions that are inherited by child objects.

**Notssid.inf**
The default file system and registry ACLs on servers assign permissions to a Terminal Server security identifier (SID). The Terminal Server SID is used only when Terminal Server is running in application compatibility mode. If Terminal Server is not being used, this template can be applied to remove the Terminal Server SIDs that are not necessary from the file system and registry locations. However, removing the access control entry (ACE) for the Terminal Server SID from these default file system and registry locations does not increase the security of the system. Instead of removing the Terminal Server SID, run Terminal Server in Full Security mode. When you run Terminal Server in Full Security mode, the Terminal Server SID is not used.

# Design Recommendations for Using Predefined Security Templates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Server 2003 security templates are for computers that use the default security settings. These templates incrementally modify the default security settings if they are on the computer. They do not install the default security settings before performing the modifications.

Consider the following recommendations before you use the predefined security templates.

**Read the security template descriptions before you select which template to use**
To successfully deploy the security templates, you need to understand which template is appropriate for your computer.

**Use the Setup security.inf template only on the local computer**
Do not apply the Setup security.inf template using Group Policy. Apply it only to the local computer by using the Security Configuration and Analysis snap-in or the Secedit.exe command-line tool. The Setup security.inf template is modified during installation and is created specifically for each computer. It might vary from computer to computer. This template contains a large amount of data and can degrade performance if it is applied by using Group Policy. This occurs because Group Policy is refreshed periodically, and a large amount of data is moving through the network each time the policy is refreshed. The advantage of using the Secedit.exe command-line tool to apply Setup security.inf is that the command-line tool permits you to configure subareas of the default settings. For example, by using Secedit.exe, you can apply only the default file system ACLs without also resetting the user rights and registry ACLs.

**Apply workstation, server, and domain controller templates appropriately**
Apply templates of the form *ws.inf *only* to workstation or server computers; do not apply *ws.inf to domain controllers. Likewise, apply templates of the form *dc.inf *only* to domain controller class machines; do not apply *dc.inf to workstations or servers. If you apply a predefined template at the domain root level, it applies to all computers in the domain by default. For example, account policies (password policies, account lockout policies, and Kerberos policies) are always defined at the domain level, but local policies are subject to precedence rules.

**Use Group Policy to apply templates to groups of computers**
You can import a security template to a Group Policy object to make sure that any computers where the Group Policy object is applied automatically receive the template's security settings when the Group Policy settings are refreshed.

**Use the appropriate tools to apply templates to local computers**
Configure individual computers by using the Security Configuration and Analysis snap-in, the Secedit.exe command-line tool, or by importing the template into the local security policy. Configure groups of computers by importing a template into the Group Policy Object Editor.

For more information about importing security templates for domain controllers, servers, or workstations, see "Import a security template to a Group Policy object" in Help and Support Center for Windows Server 2003.

# Refreshing Group Policy to Activate New Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If any changes are made to the settings, all Group Policy settings are refreshed automatically every 90 minutes on a workstation or server and every five minutes on a domain controller. Additionally, security settings are refreshed every 16 hours, regardless of whether any changes have been made.

Security settings are also refreshed when you restart the computer, and you can trigger a Group Policy object refresh to apply a newly deployed GPO, including security settings. You can do so by using the **gpupdate** command. The **gpupdate** commandreplaces the **/refreshpolicy** option that was previously used for the **secedit** command.

For more information about the **gpupdate** command including full syntax, see "Gpupdate" in Help and Support Center for Windows Server 2003.

# Designing Software Restriction Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Software restriction policies permit you to identify software that is running on computers in your domain and to control its ability to run. By using software restriction policies, you can:

- Control which programs can run on your system. For example, you can control which types of programs can be accepted as e-mail attachments if you are concerned about users being sent e-mail viruses.
- Run only digitally signed scripts.
- Permit users to run only specific files on multiuser computers. For example, you might want to lock down specific executable files in specific directories on a terminal server.
- Decide who can add trusted publishers to a computer.
- Control whether software restriction policies affect all users or only specific users who use a computer.
- Prevent any files from running on a local computer.

Consider the following design recommendations before you create rules and apply software restriction policies:

- Decide if you want the policy to apply to multiple computers or users in a domain or OU or if you want it only to apply to the local computer. If you want the policy to apply to many computers or users in a domain or other Active Directory container, use a GPO. If you want the policy to apply only to the local computer, use the Local Security Policy snap-in.
- Decide if you want the policy to apply to users or computers. Applying a policy to users affects the users regardless of which computer they use to log on to the network. Applying a policy to computers affects the computers regardless of which users log on.
- Decide what the default security level of the policies will be. The default security level can either be **Disallowed**, which does not permit software to run, or **Unrestricted**, which permits software to run with the full permissions of the user who is logged on. After you have decided what the default security level will be, you can begin creating rules. The Rules types are:
    - **Hash rules.** Software is identified by its hash.
    - **Certificate rules.** Software is identified by its signing certificate.
    - **Path rules.** Software is identified by its file path or registry path.
    - **Internet zone rules.** Software is identified by using a zone that is specified by Microsoft® Internet Explorer. Zone rules apply only to Windows Installer packages.

For more information about deploying software restriction policies, see:

- "Deploying a Managed Software Environment" in this book.
- The Software Restriction Policies link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- "Software Restriction Policies" in Help and Support Center for Windows Server 2003.

# Designing Wireless Network Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Administrators can define policy settings to help secure wireless network configurations for both the IEEE 802.11 configuration and the IEEE 802.1X authentication engine. If you use Group Policy for wireless networking, the policy settings that are defined in the GPO that affects a specific computer take precedence over the user-defined settings. The type of network that you use also affects precedence. You can design wireless network policies for three types of networks:

- **Access point (infrastructure).** In this type of network, wireless devices with radio network adapters, such as a portable computer or personal digital assistant, connect to wireless access points.

- **Computer-to-computer (ad hoc).** In this type of network, wireless devices connect to each other directly instead of through wireless access points.

- **Any available network access point preferred.** When you select this option, a connection to an access point wireless network is always attempted first, if one is available. If an access point network is not available, a connection to a computer-to-computer wireless network is attempted.

The wireless networking technologies that are included in Windows XP Professional and Windows Server 2003 support the IEEE 802.11 and IEEE 802.1X standards for wireless communications. For complete information about deploying these technologies, see "Deploying a Wireless LAN" in *Deploying Network Services* of this kit. For step-by-step information about creating wireless network policies, see "Wireless Networking" in Help and Support Center for Windows Server 2003.

# Design Considerations for Wireless Network Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Consider the following issues that pertain to authentication methods and wireless network policies:

- Computer authentication is recommended. By default, authentication is set to **Enabled**.
- The access point must support the authentication method that you select. For example, the access point must support 802.1X. If you choose EAP-TLS, all computers must support it (for example, a RADIUS server must support EAP-TLS).
- Your servers and wireless clients must support the authentication method you plan to deploy. Whether you choose EAP-TLS or PEAP as the authentication method over 802.1X, both your RADIUS server and your wireless clients need to support it.
- It is recommended that you permit certificate autoenrollment for users and computer when you use EAP-TLS. For more information, see "Checklist: Configuring certificate autoenrollment" in Help and Support Center for Windows Server 2003.
- The wireless network configuration settings that are defined in GPOs take precedence over user-defined settings. The only exception to this is the list of preferred networks, where the policy-defined list is merged with the user-defined list. For more information, see "Creating Wireless Network Policies" later in this chapter.
- If a domain policy for wireless configuration exists, the local user (whether the user is an administrator or non-administrator) cannot remove or disable the domain policy.
- When a Group Policy change occurs, the Wireless Configuration service breaks the current association *if and only if* the new policy takes precedence (for example, a visible network is now a more preferred network according to the policy's list of preferred networks). In all other cases, the association does not change.
- If a GPO that contains wireless network policies is deleted, the Wireless Configuration service clears its policy cache, initiates and processes a soft reset, and then reverts to the user-configured settings.

# Understanding Wireless Network Policy Precedence

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If wireless network configuration settings have been defined both on the local computer (by the user) and in a GPO that affects that computer, the wireless network settings are merged, and the user cannot change the Group Policy wireless network settings. One exception to this rule is the Wired Equivalent Privacy (WEP) key. The user also cannot change the order in which the settings are applied. The wireless networks can be access point (infrastructure) or computer-to-computer (ad hoc) networks. The merging process and order of precedence occur according to the following rules:

- Infrastructure networks *always* have higher precedence than ad hoc networks.
- Group Policy overrides user-defined policy, and the wireless network policy configurations have the highest precedence of their respective group of configurations (infrastructure or ad hoc).

For example, an administrator might define a GPO with the following wireless configuration settings, and then select **The key is provided for me automatically** option:

- Service Set Identifier (SSID): any
- Network type: either infrastructure or ad hoc
- Authentication mode: either open or shared
- Encryption: WEP

In this case, the user can clear the **The key is provided for me automatically** check box,and type in different key information because WEP key configuration can be changed locally.

**Example: Preferred Network Precedence**
This example illustrates how the merging and ordering process occurs. The client computer is a member of WirelessClientsOU. A GPO named WirelessConfigGPO is assigned to WirelessClientsOU,and the GPO defines the following list of preferred networks:

- Infrastructure network Ip1
- Infrastructure network Ip2
- Ad hoc network Ap1
- Ad hoc network Ap2

The client computer has this list of user-defined preferred networks (local):

- Infrastructure network Iu1, which is user defined and the same network as Ip1.
- Infrastructure network Iu2, which is user defined and the same network as Ip2.
- Infrastructure network Iu3.
- Ad hoc network Au2, which is the same network as Ap2.
- Ad hoc network Au3.

When the GPO is applied to that client, the resulting list of preferred networks is merged in the following order: Ip1, Ip2, Iu3, Ap1, Ap2, and Au3. This occurs because infrastructure networks take precedence over ad hoc networks and because settings in the OU take precedence over local settings. For example, Ip1 and Iu1 are the same network with different configurations; Ip1 is configured by a Group Policy, and Iu1 is configured locally.

For more information, see "Define Wireless Network Policies on a Client Computer" and "Define Active Directory–based wireless network policies" in Help and Support Center for Windows Server 2003.

**Multiple GPOs**

If a computer is subject to multiple GPOs that contain wireless network policies, the wireless settings that are defined in the GPO associated with the Active Directory container closest to the computer object takes precedence. Those settings will override the settings that are assigned to a higher level Active Directory container. In this case, the settings are not merged. For example, a client is a member of OU1 in the Redmond domain, and the following GPOs are being used:

- **GPOA.** This GPO defines wireless network policies and is assigned to the Redmond domain.
- **GPO1.** This GPO contains wireless policy settings and is assigned to OU1. The client is a member of this OU.

When the GPO list for this client is processed, the GPO1 wireless network settings take precedence and are in effect for this computer. The settings are not merged.

# Measuring Your Design Using the GPMC

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Windows Server 2003, you can use the Group Policy Results Wizard in the GPMC to view information about the following settings defined in your wireless network policies:

- The preferred list of wireless networks to which clients can connect
- The wireless network association settings
- The IEEE 802.1X authentication and settings

The Group Policy Results wizard replaces the Resultant Set of Policy (RSoP) – Logging Mode capabilities. For step-by-step information about using the Group Policy Results wizard, see "How to create a new Group Policy Results query" in GPMC Help.

# Configuring Security Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use Group Policy to manage all of the computers in your organization from one computer. By using GPMC, you can create and edit policies that affect workstations, servers, and domain controllers throughout your organization. Figure 4.3 shows which security settings to configure.

**Figure 4.3 Configuring Security Policy**



# Configuring Security Settings for Sites, Domains, and OUs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using the Group Policy Object Editor, you can modify individual security settings in a GPO, or you can import a template to a GPO that is linked to an OU, domain, or site. If you want to change the default security settings on your computers, there are several ways to configure security settings on multiple computers.

Before you make any changes, make sure that you understand the default security settings. For more information, see Default security settings for groups and Default groups in Help and Support Center for Windows Server 2003. For information about the default setting for each policy, see Security Setting Descriptions in Help and Support Center for Windows Server 2003.

# Importing Security Templates and Modifying Security Settings in a GPO

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using Group Policy Object Editor and a security template, you can create a security policy for your computers. You can use GPMC to navigate to any GPO in the forest, and then open the Group Policy Object Editor. The Group Policy Object Editor permits you to import security templates and edit security settings. The security template is a single location that contains the full range of security settings. After you have imported the security template, you can edit individual policies.

**To import a security template for a domain or OU**

1. Open GPMC.
2. In the console tree, expand the domain or OU that you want to manage, right-click the Group Policy object that you want to edit, and then click **Edit**.
3. In the Group Policy Object Editor console tree, click **Computer Configuration**, click **Windows Settings**, right-click **Security Settings**, and then select **Import Policy**.
4. Click the security template that you want to import, and then click **Open**.

For step-by-step information about configuring security templates, see "Security Templates" in Help and Support Center for Windows Server 2003.

After you have imported a template to the GPO, you can manually modify the security settings in the GPO.

**To modify security settings**

1. Open GPMC.
2. In the console tree, expand the domain or OU that you want to manage, right-click the Group Policy object that you want to edit, and then click **Edit**.
3. In the Group Policy Object Editor console tree, click **Computer Configuration**, click **Windows Settings**, and then click **Security Settings**.
4. Do one of the following:
   - To edit **Password Policy**, **Account Lockout Policy**, or **Kerberos Policy**, click **Account policies.** Before making any changes to these policies make sure you understand how account policies are applied. For more information, see "Modifying Account Policies in the Default Domain GPO" later in this section.
   - To edit **Audit Policy**, **User Rights Assignment**, or **Security Options**, click **Local Policies**.
   - To edit Event log settings, click **Event Log**.
5. In the details pane, double-click the security setting that you want to modify.
6. If the security settings have not been defined, you can click to select the **Define these policy settings** check box.
7. Modify the security settings, and then click **OK**.

For step-by-step information about configuring security templates, see "Security Templates" in Help and Support Center for Windows Server 2003.

# Modifying Account Policies in the Default Domain GPO

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Account policies include password policy, account lockout policy, and Kerberos policy. Although they affect user accounts, account policies are defined on computers. In Windows Server 2003 and Windows 2000, there is only one set of account policies for each domain, and these policies are defined at the domain level.

- The account policies for the domain affect all users who log on to the domain, regardless of any other account policies that are defined in a linked GPO. However, if a user logs on to a member computer by using a local computer account, the user is subject to the usual order of precedence as described in "Designing Security Settings" earlier in this chapter. Domain controllers do not have local accounts and always receive account policies from the domain.

  - If you want to modify the account policies at the domain level, it is recommended that you modify the default domain GPO. The account policies for domains include:

  - **Password policy settings for domains.** The most common way to authenticate a user's identity is by using passwords. After a user has been identified and authenticated, the user can perform any tasks or access any resource for which he or she is authorized. Strong passwords generally enhance security for users. Using strong passwords helps avoid the threat of an unauthorized user guessing a weak password (also known as *cracking*) and acquiring the credentials of the compromised user account (also known as *spoofing*). This is especially true for administrative accounts because an unauthorized user might obtain administrative credentials and gain elevated privileges. For more information, see "Strong passwords" in Help and Support Center for Windows Server 2003.

  - **Account lockout policy settings for domains.** More than a few unsuccessful password tries during logon might represent an attacker's attempt to determine an account password by trial and error. Windows Server 2003 and Windows 2000 tracks the number of logon attempts, and it can be configured to respond to this type of attack by disabling the account for a preset period of time. This is known as an account lockout.

  - **Kerberos policy settings for domains.** Kerberos V5 authentication protocol provides the default mechanism for authentication services and the authorization data that is necessary for a user to gain access and perform a task on a resource. By reducing the lifetime of Kerberos authentication tickets, you reduce the risk of having a legitimate user's credentials stolen and successfully used by an attacker. However, authorization overhead is increased.

In addition, there are five Security Option settings that behave in the same way as account policies. These are:

- Network Security: Force logoff when logon hours expire
- Accounts: Administrator account status
- Accounts: Guest account status
- Accounts: Rename administrator account
- Accounts: Rename guest account

For more information about the individual account policies, see "Account Policies" in Help and Support Center for Windows Server 2003.

# Modifying Local Policies in the Default Domain Controllers GPO

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can establish local policies to secure your domain controllers. Local policies are divided into a number of settings categories, including:

- **User rights assignment.** User rights permit users to log on and perform specific administrative or operations tasks on your domain controllers. Make sure that the appropriate user rights are assigned to users in the domain so that the users can perform their intended functions without compromising the security of the domain controllers.

  **Important**

  - Make changes to the user rights assignment policy settings directly in the Default Domain Controllers GPO for application compatibility reasons. Not doing so can cause unexpected behavior if server applications are later installed on the domain controller.

- **Audit policy settings.** When you enable auditing on your domain controllers, the number of events that is recorded in the security event log increases. As a result, the maximum size of the security event log must be increased. As a part of your typical operations tasks, archive the security and system event logs regularly and frequently before they fill up. If the logs fill up, you might miss events.

  **Important**

  - Make changes to the audit policy settings directly in the Default Domain Controllers GPO for application compatibility reasons. Not doing so can cause unexpected behavior if server applications are later installed on the domain controller.

- **Security options policy settings.** The domain controller security options policy settings affect Active Directory–related security configuration settings and network, file system, and user logon security configuration settings.

For more information about securing domain controllers, see the Securing Active Directory Installations link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Implementing Software Restriction Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using software restriction policies, you can protect your computer or network by identifying and specifying the software that is permitted to run. The first step is to set the security level for a computer to **Unrestricted** or **Disallowed**. If the default level of security is set to **Unrestricted**, users can run all programs, except for the programs that you restrict by using rules. If the default security level is **Disallowed**, no programs can run on the computer, except the programs that meet the requirements of the rules. After you set a default security level to **Unrestricted** or **Disallowed**, you can make exceptions to the default security level by creating software restriction policy rules.

**Rule Precedence**

Rules are evaluated in a specific order. The rules that more specifically match a program take precedence over rules that more generally match a program. From highest to lowest, rule precedence is as follows:

- Hash rule
- Certificate rule
- Path rule
- Internet zone rule
- Default security level

**Example of Rule Precedence with a Default Security Level of Unrestricted**

Table 4.2 shows an example of rules that are configured by using software restriction policies on a computer that has a default security level of **Unrestricted**.

**Table 4.2 Example of Software Restriction Policy Rules**

| Rule | Type of Rule | Description | Setting |
|---|---|---|---|
| Rule 1 | Hash rule | Hash of Pagefileconfig.vbs | **Disallowed** |
| Rule 2 | Certificate rule | IT management certificate | **Unrestricted** |
| Rule 3 | Path rule | %WINDIR%\System32*.VBS | **Unrestricted** |
| Rule 4 | Path rule | *.VBS | **Disallowed** |
| Rule 5 | Path rule | %WINDIR% | **Unrestricted** |

Table 4.3 shows examples of how the software restriction policy rules in Table 4.2 determine which programs users can run.

**Table 4.3 Example of the Application of Software Restriction Policy Rules**

| Program Started | Applied Rules | Outcome |
|---|---|---|

| | | |
|---|---|---|
| C:\WINDOWS\SYSTEM32\EventQuery.vbs | Rule 3 is applied because EventQuery.vbs is a .vbs file in the System32 folder.<br><br>Rule 4 is applied because EventQuery.vbs has a .vbs extension.<br><br>Rule 5 is applied because EventQuery.vbs is stored in a subfolder of the Windows folder. | Rule 3 is the most specific match for EventQuery.vbs. Because Rule 3 has a security level of **Unrestricted**, EventQuery.vbs is permitted to run. |
| C:\WINDOWS\SYSTEM32\Pagefileconfig.vbs | Rule 1 is applied because the hash in the rule matches the hash of Pagefileconfig.vbs.<br><br>Rule 3 is applied because Pagefileconfig.vbs is a .vbs file in the System32 folder.<br><br>Rule 4 is applied because Pagefileconfig.vbs has a .vbs extension.<br><br>Rule 5 is applied because Pagefileconfig.vbs is stored in a subfolder of the Windows folder. | Rule 1 is the most specific match for Pagefileconfig.vbs. Because Rule 1 has a security level of **Disallowed**, Pagefileconfig.vbs is not permitted to run. |
| \\LOGIN_SRV\Scripts\CustomerScript1.vbs | Rule 2 is applied because CustomerScript1.vbs is digitally signed by the certificate that belongs to the customer's IT management group.<br><br>Rule 4 is applied because CustomerScript1.vbs has a .vbs extension. | Rule 2 is the most specific match for CustomerScript1.vbs. Because Rule 2 has a security level of **Unrestricted**, CustomerScript1.vbs is permitted to run. |
| C:\Documents and Settings\user1\LOVE-LETTER-FOR-YOU.TXT.VBS | Rule 4 is applied because LOVE-LETTER-FOR-YOU.TXT.VBS has a .vbs extension. | Rule 4 is the most specific match for LOVE-LETTER-FOR-YOU.TXT.VBS. Because the Rule 4 has a security level of **Disallowed**, LOVE-LETTER-FOR-YOU.TXT.VBS is not permitted. |

**DLL Checking**

Most programs require multiple .dll files to run. By default, software restriction policy rules are not enforced against .dll files. This is the recommended option for most users for the following reasons:

- Setting a program's main executable file as **Disallowed** in the Software Restriction Policies snap-in prevents the program from running. Therefore, it is typically not necessary to set all of the related .dll files to **Disallowed**.

- Checking .dll files results in performance degradation on the system. If a user runs 10 programs during his or her logon session, the software restriction policy is evaluated 10 times. If .dll checking is turned on, the software restriction policy is evaluated for each .dll file that is loaded with each program. For example, if each program uses 20 .dll files, 10 .exe files are checked, and then 200 .dll files are checked. The software restriction policy is evaluated 210 times.

- If the default security level is set to **Disallowed**, the main executable file has to be identified so that it can run, and all of the executable file's .dll files must be identified.

To activate .dll checking, in the **Enforcement Properties** policy, set **Apply software restriction policies to the following** to **All software files**. For more information about .dll checking, see "Apply software restriction policies to DLLs" in Help and Support Center for Windows Server 2003.

**Exclude Administrators**

You might want to prevent most users from running certain programs by setting them as **Disallowed** but permit administrators to run any programs. For example, a user might have a shared computer that multiple users connect to by using Terminal Server. You might want the users to run specific applications on the computer, while permitting members of the Local Administrators group to run any programs. To exclude local administrators, you can set the option to apply software restriction policies to **All users except local administrators**.

### Note

- Setting the software restriction policies to exclude local administrators is only valid for Group Policy objects that are applied to computers.

For more information about software restriction policies, see "Prevent software restriction policies from applying to local administrators" and "Software Restriction Policies" in Help and Support Center for Windows Server 2003, and "Deploying a Managed Software Environment" in this book.

# Configuring Wireless Network Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Wireless network settings can be configured locally, by users on client computers, or centrally. To enhance the deployment and administration of wireless networks, you can use Group Policy to centrally create, modify, and assign wireless network policies for Active Directory clients. When you use Group Policy to define wireless network policies, you can configure wireless network connection settings, enable IEEE 802.1X authentication for wireless network connections, and specify the preferred wireless networks that clients can connect to. When you create and configure wireless policies, you have the options that are described in Table 4.4. For more information about configuring wireless policies, see "Define Active Directory-based wireless network policies" in Help and Support Center for Windows Server 2003.

Table 4.4 Configuration Settings for General Policy

| Options | Comments |
|---|---|
| **Name** | Name of the policy. Use a unique and descriptive name of up to 255 characters that easily identifies the policy. |
| **Check for policy changes every** | Specifies in minutes how often to poll Active Directory for changes to this policy. Applies only to computers that are members of an Active Directory domain. The default is 180 minutes. |
| **Network to access**<br>• Any available network (access point preferred)<br>• Access point (infrastructure) networks only<br>• Computer-to-computer (ad hoc) networks only | Specifies the types of IEEE 802.11 wireless networks that are available for clients to try to connect to. |
| **Use Windows to configure my wireless network settings** | Specifies whether client settings are automatically configured for IEEE 802.11 wireless network connections. |
| **Automatically connect to non-preferred networks** | Specifies whether clients can try to connect to any available IEEE 802.11 wireless networks that are within range. |

**Network authentication services**

The IEEE 802.11–supported network authentication services provide open system and shared key authentication. Open system authentication permits any wireless device to associate with an access point. Shared key authentication requires a network key to be used. For security reasons shared key authentication is not recommended. Instead, open system authentication used in conjunction with 802.1X authentication is recommended.

**Network keys**

When you enable WEP, you can require that a network key be used for encryption. You can specify a key (by typing a key in the **Network key** text box when you configure the wireless connection). If you specify a key, you can also provide its location in the **Key index** text box (on the **Properties** page for **Wireless Network Connections**). Table 4.5 includes descriptions of the configuration settings for requiring network keys.

Table 4.5 Configuration Settings for Preferred Networks

| Options | Comments |
|---|---|

| | |
|---|---|
| **Networks** | Lists the IEEE 802.11 wireless networks to which clients can try to connect. Use the **Move Up** and **Move Down** buttons to prioritize the list. Use the **Add** button to add a new wireless network. You can also edit properties of a network by using the **Edit** button, or use the **Remove** option to remove an entry from the list. |
| **Network name (SSID)** | Specifies the name for the specified wireless network. Under the IEEE 802.11 standard, the network name is also known as the Service Set Identifier (SSID). |
| **Description** | Provides a description for the specified wireless network. Use a unique description of up to 255 characters. |
| **Wireless network key (WEP)**<br>• Data encryption (WEP enabled)<br>• Network authentication (Shared mode)<br>• The key is provided automatically | **Data Encryption (WEP enabled)** specifies that a network key is used to encrypt the data that is sent over the network.<br>• **Network authentication (Shared mode)** specifies that a network key be used for authentication to the wireless network.<br>**The key is provided automatically** specifies that a network key is automatically provided for clients. |
| **This is a computer-to-computer (ad hoc) network; wireless access points are not used** | Specifies whether this preferred network is a computer-to-computer ad hoc network. If this check box is not selected, this network is an access point (infrastructure) network. |

**IEEE 802.1X authentication**

To provide user and computer identification, centralized authentication, and dynamic key management, you can enable IEEE 802.1X authentication.

You can use Group Policy to create a wireless configuration policy to configure IEEE 802.11 and IEEE 802.1X values. Table 4.6 and Table 4.7 list the wireless network policy settings that you can specify.

**Table 4.6 Wireless Network (IEEE 802.11) Policy Settings**

| Options | Comments |
|---|---|
| **Enable network access control using IEEE 802.1X** | Use 802.1X authentication when you connect to an 802.11 wireless network. |
| **EAPOL-Start message**<br>• Do not transmit<br>• Transmit<br>• Transmit per IEEE 802.1X | Specifies how Extensible Authentication Protocol over LAN (EAPOL)-start messages are transmitted. |

**Table 4.7 Wireless Network (IEEE 802.1X) Authentication Settings**

| Options | Comments |
|---|---|

| | |
|---|---|
| **Parameters (seconds):**<br>• Max start<br>• Held period<br>• Start period<br>• Authentication period | Default **Max start** value is 3 seconds.<br>Default **Held period** is 60 seconds.<br>Default **Start period** is 60 seconds.<br>Default **Authentication period** is 30 seconds. |
| **EAP type:**<br>• Smart card or other certificate<br>• Protected Extensible Authentication Protocol (PEAP) | Click **Settings** to specify the options to use when connecting, including: using a smart card or certificate on the computer; validating server certificate; specifying which servers to connect to; Trusted Root Certification Authorities; viewing certificates; and selecting and configuring an authentication method. |
| **Authenticate as guest when user or computer information is unavailable** | Specifies whether clients attempt authentication to the wireless network as guests when user or computer information is not available. |
| **Authenticate as computer when computer information is available** | Specifies whether client computers must attempt authentication to the wireless network when a user is not logged on. The default setting is **Enabled**. |
| **Computer authentication:**<br>• With user authentication<br>• With user re-authentication<br>• Computer only | It is recommended that you select **With user re-authentication**. When this option is selected, authentication is performed by using the computer credentials when users are not logged on to the computer. After a user logs on to the computer, authentication is performed by using the user credentials. When a user logs off of the computer, authentication is performed by using the computer credentials. |

# Creating Wireless Network Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can define wireless network policies for your organization by using the Group Policy Object Editor snap-in.

**To access Wireless Network (IEEE 802.11) Policies**

1. Open GPMC.
2. Right-click the GPO that you want to edit, and then click **Edit**.
3. In the Group Policy Object Editor console tree, click **Computer Configuration**, click **Windows Settings**, and then click **Security Settings**.
4. Right-click **Wireless Network (IEEE 802.11) Policies on Active Directory**, and then click **Create Wireless Policies**. The Wireless Policy Wizard starts.

# Defining Wireless Configuration Options for Preferred Networks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using the **Properties** page for your wireless configuration policy, you can define a list of preferred networks to use. You can use the **General** tab to specify how often to check for policy changes, which networks to access, whether to disable Zero Configuration, or automatically connect to non-preferred networks.

**To define preferred wireless networks**

1. Open GPMC.
2. In the console tree, expand the domain or OU that you want to manage, right-click the Group Policy object that you want to edit, and then click **Edit**.
3. In the Group Policy Object Editor console tree, click **Computer Configuration**, click **Windows Settings**, and then click **Security Settings**.
4. Click **Wireless Network (IEEE 802.11) Policies**, right-click the wireless network policy that you want to modify, and then click **Properties**.
5. Click the **Preferred Networks** tab, and then click **Add**.
6. Click the **Network Properties** tab, and then in the **Name** box, type a unique name.
7. In the **Description** box, type a description of the wireless network, such as the type of network and whether WEP and IEEE 802.1X authentication are enabled.
8. In the **Wireless network key (WEP)** box, specify whether a network key is used for encryption and authentication, and whether a network key is provided automatically. The options are:
   - **Data encryption (WEP enabled).** Select this option to require that a network key be used for encryption.
   - **Network authentication (Shared mode).** Select this option to require that a network key be used for authentication. If this option is not selected, a network key is not required for authentication, and the network is operating in open system mode.
   - **The key is provided automatically.** Select this option to specify whether a network key is automatically provided for clients (for example, whether a network key is provided for wireless network adapters).
9. To specify that the network is a computer-to-computer (ad hoc) network, click to select the **This is a computer-to-computer (ad hoc) network; wireless access points are not used** check box.

**To define 802.1X authentication**

1. Open GPMC.
2. In the console tree, expand the domain or OU that you want to manage, right-click the Group Policy object that you want to edit, and then click **Edit**.
3. In the Group Policy Object Editor console tree, click **Computer Configuration**, click **Windows Settings**, and then click **Security Settings**.

4. Click **Wireless Network (IEEE 802.11) Policies**, right-click the wireless network policy that you want to modify, and then click **Properties**.

5. On the **Preferred Networks** tab, under **Networks**, click the wireless network for which you want to define IEEE 802.1X authentication.

6. On the **IEEE 802.1X** tab, check the **Enable network access control using IEEE 802.1X** check boxtoenable IEEE 802.1X authentication for this wireless network. This is the default setting. To disable IEEE 802.1X authentication for this wireless network, clear the **Enable network access control using IEEE 802.1X** check box.

7. Specify whether to transmit EAPOL-start message packets and how to transmit them.

8. Specify EAPOL-Start message packet parameters.

9. In the **EAP type** box, click the EAP type that you want to use with this wireless network.

10. In the **Certificate type** box, select one of the following options:

    - **Smart card.** Permits clients to use the certificate that resides on their smart card for authentication.

    - **Certificate on this computer.** Permits clients to use the certificate that resides in the certificate store on their computer for authentication.

11. To verify that the server certificates that are presented to client computers are still valid, select the **Validate server certificate** check box.

12. To specify whether client computers must try authentication to the network, select one of the following check boxes:

    - **Authenticate as guest when user or computer information is unavailable.** Specifies that the computer must attempt authentication to the network if user information or computer information is not available.

    - **Authenticate as computer when computer information is available.** Specifies that the computer attempts authentication to the network if a user is not logged on. After you select this check box, specify how the computer attempts authentication.

For detailed information about using the wireless network policies, see "Define Wireless Network Policies on a Client Computer" in Help and Support Center for Windows Server 2003.

# Testing Security Policies and Rolling Out a Pilot Project

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy your security policy solution to your organization, be sure to test the configuration in a test environment and follow security policy best practices that are outlined in this chapter and in Help and Support Center for Windows Server 2003. If you have problems with your configuration, it is much easier to correct them in a non-production environment than after you deploy the configuration to your servers and clients. For more information about staging and testing Group Policy, see "Staging Group Policy Deployments" in this book.

**Test customized security templates before you apply them**
Do not apply the predefined security templates to production computers without first testing your customized version to make sure that the appropriate level of application functionality is maintained for the network and computer architecture.

**Make sure that automatic Startup services start correctly**
If you set the Startup service to start automatically, you must perform adequate testing to verify that the services can start without user intervention. For more information about setting the Startup service to start automatically, see "Configure how a service is started" in Help and Support Center for Windows Server 2003.

For more information about testing, see "Designing a Test Environment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

# Additional Resources for Deploying Security Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- *Designing and Deploying Directory and Security Services* of this kit for more information about distributed security.
- "Deploying IPSec" in *Deploying Network Services* of this kit for more information about Internet Protocol security (IPSec).
- "Deploying a Wireless LAN" in *Deploying Network Services* of this kit for more information about wireless networks.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about Group Policy.
- "Deploying IAS" in *Deploying Network Services* of this kit for more information about Internet Authentication Service (IAS).
- The TechNet Security link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about Windows Server 2003 security.

**Related Job Aids**

- "Worksheet A.27 Using Custom Security Templates" (DMEUSE_27.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.27 Using Custom Security Templates" on the Web at http://www.microsoft.com/reskit).
- "Worksheet A.28 Selecting Predefined Security Templates" (DMEUSE_28.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.28 Selecting Predefined Security Templates" on the Web at http://www.microsoft.com/reskit).
- "Worksheet A.29 Designing Security Policy" (DMEUSE_29.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.29 Designing Security Policy" on the Web at http://www.microsoft.com/reskit).

**Related Tools**

- Group Policy Management Console

  To download the GPMC tool, see the Group Policy Management Console link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Security Settings Concepts" in Help and Support Center for Windows Server 2003 for more information about using the security settings and security templates.
- "Internet Protocol Security (IPSec)" in Help and Support Center for Windows Server 2003 for more information about using IPSec.

- "Internet Authentication Service" in Help and Support Center for Windows Server 2003 for more information about using IAS.
- "Software Restriction Policies" in Help and Support Center for Windows Server 2003 for more information about using software restriction policies.
- "Creating, modifying, and assigning Active Directory-based wireless network policies" in Help and Support Center for Windows Server 2003 for more information about how to set up security policies for wireless networking.
- "Assign or unassign IPSec policy in Group Policy" in Help and Support Center for Windows Server 2003 for more information about using IPSec with Group Policy.
- "Gpupdate" in Help and Support Center for Windows Server 2003 for more information about the **gpupdate** command including full syntax.
- "Security Templates" in Help and Support Center for Windows Server 2003.
- "Security Configuration and Analysis" in Help and Support Center for Windows Server 2003.
- "Security Settings Concepts" in Help and Support Center for Windows Server 2003.
- "Local Security Policy" in Help and Support Center for Windows Server 2003.
- "Secedit" in Help and Support Center for Windows Server 2003.
- "Creating, modifying, and assigning IPSec policies" in Help and Support Center for Windows Server 2003.
- "Customize a predefined security template" in Help and Support Center for Windows Server 2003.
- "Predefined security templates" in Help and Support Center for Windows Server 2003.
- "Import a security template to a Group Policy object" in Help and Support Center for Windows Server 2003.
- "Define Wireless Network Policies on a Client Computer" Help and Support Center for Windows Server 2003.
- "Configure how a service is started" in Help and Support Center for Windows Server 2003.

# Migration of User State

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you move client computers to the Microsoft® Windows® XP operating system from earlier versions of Windows, it is important to save and then restore user data and settings. This process is known as *migrating user state*. By carefully planning and implementing user state migration, you help conserve IT staff time, preserve important data, scale the migration as needed, and minimize costs while maintaining user productivity and workplace morale.

**In This Chapter**

Overview of Migrating User State

Choosing a Method for Collecting User State

Identifying Migration Content for Your Migration

Creating a Detailed Plan for User State Migration

Testing Your Process for Migration

Additional Resources for User State Migration

**Related Information**

- For information about using Remote Installation Services (RIS), see "Designing RIS Installations" in *Automating and Customizing Installations* of this kit.

- For information about scripting in a Microsoft® Windows® Server 2003 environment, see the at http://www.microsoft.com/reskit, or see the MSDN Scripting Clinic link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Overview of Migrating User State

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Any time that you perform a new installation of Windows XP on a client workstation, you should migrate user state to ease users into the new system and maintain user productivity. *User state* consists of *user data* — the files that users create and need to do their jobs — along with *user settings* containing application-specific and user-specific information. Additionally, application settings supply the user with links, menus, and other information that can be essential for their productivity.

If user state is not migrated, an organization can accrue costs as users spend production time reconfiguring their applications and other settings. Organizations must evaluate the cost/benefits ratios for migrating various types of items. They must understand the security issues related to migration and be sure to educate users about what to expect before and after the migration.

The way that you choose to deploy Windows XP affects your user state migration plan. Ideally, an organization can perform either a parallel or a wipe-and-load deployment, restoring collected user state to a clean environment.

The method that you use to collect and restore user state is critical to the success and efficiency of your user state migration. To avoid the high migration cost of a strictly manual migration, an organization can:

- Partially script the migration, leaving nonstandard items to the discretion of the individual user or IT staff.
- Use migration tools that automate the migration of common settings but allow customization.
- Create its own custom tools.

The degree to which an organization should automate user state migration depends on these and a variety of other factors, including the number of users to be migrated and how widely dispersed they are; how centralized the organization's IT effort is; the degree to which users share a common desktop, folder hierarchy, and computing requirements; the IT expertise available to assist in and support the migration; and whether the deployment involves simultaneous domain migration.

Before you begin planning a user state migration, identify the computers on which you will deploy Windows XP, and determine the appropriate deployment method for each computer. When you complete this process, you will be ready to deploy Windows XP, with a complete, tested plan for migrating user state during the system deployment and a schedule for the migration. For worksheets to assist you with the deployment of migrating user state, see "Additional Resources for User State Migration," later in this chapter.

# User State Migration Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Proper planning is essential for a successful user state migration. Before creating a detailed migration plan, identify the best methods for collecting, storing, and restoring user state data and decide which user data and settings to migrate. After making these decisions, prepare a migration plan that addresses storage and security requirements; potential registry, drive, and domain changes; scheduling; and the education of users. Then test your migration process before embarking on a large-scale migration. Figure **6**.1 shows the process for planning a user state migration.

**Figure 6.1 Migrating User State**



# Tools Used in the Migration Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some large organizations develop their own migration tool. This frequently provides an excellent migration experience for the user, because the tool is customized to the specific environment. Such a tool can capture and restore user settings and files. Adjustments can be made quickly when new applications are deployed or bugs are found.

Typically, this option requires a significant investment in development time. If your organization does not have personnel who can create a migration tool, the cost of hiring programmers to create one might be prohibitive. Even if you do have programmers on staff, compare the cost of tool development with the migration costs of other available methods.

Many tools currently available from vendors can collect and restore most necessary settings and are extensible to include additional application settings. These tools often provide multiple types of rules to specify which files to migrate. However, while fairly thorough, such tools are not targeted to your specific environment, and their initial cost can be high.

Microsoft provides two tools designed specifically for migrating user state in a Windows environment:

- The Files and Settings Transfer Wizard
- The User State Migration Tool (USMT)

Both tools automate the migration of basic application, operating system, and user settings, as well as user data, and both support customization.

**Files and Settings Transfer Wizard**

The Files and Settings Transfer Wizard is a Windows XP accessory, available in **Systems Tools**. (On the **Start** menu, point to **All Programs**, **Accessories**, **System Tools**, and then click **Files and Settings Transfer Wizard**.) The wizard enables users to migrate personal display properties, folder and taskbar options, and Internet browser and mail settings, as well as specific files or entire folders (such as **My Documents**, **My Pictures**, and **Favorites**) from their old computer to their new one without any manual configuration.

Designed for home users and small office users, the Files and Settings Transfer Wizard is also useful in a corporate network environment, for employees who get a new computer and need to migrate their own files and settings without the support of an IT department or Help desk. For information about using the Files and Settings Transfer Wizard, see Help and Support Center for Windows XP.

**User State Migration Tool (USMT)**

Designed for IT administrators who are performing large deployments of the Microsoft® Windows® XP Professional operating system in a corporate environment, USMT provides the same functionality as the Files and Settings Transfer Wizard, but on a large scale targeted at migrating multiple users. USMT gives administrators precision in configuring unique settings, such as making user-specific modifications to the registry. The tool is included on the Microsoft® Windows® Server 2003 operating system CD in the \ValueAdd\Msft\USMT folder.

USMT uses the following files in collecting and migrating user data and settings:

- Scanstate.exe collects user state.
- Loadstate.exe restores user state.
- Migapp.inf determines which application settings are migrated.
- Migsys.inf determines which operating system settings are migrated.
- Miguser.inf determines which user settings are migrated.
- Sysfiles.inf defines files that *must not* be migrated despite any other rules. These are operating system files that will conflict with the newer version of the files in Windows XP. The SysFiles.inf file should not be modified except to add more files to the list of files that never migrate under any circumstances.

These files are shipped with Windows XP in the ValueAdd\Msft\USMT folder.

Table 6.1 and Table 6.2 list the file types, folders, settings, and system components that are migrated by default using USMT. (See also the Inf Commands.doc file included on the Windows Server 2003 operating system CD in the \ValueAdd\Msft\USMT folder.)

**Table 6.1 File Types and Folders Migrated by Default by USMT**

| File Types Migrated | Folders Migrated |
|---|---|
| .ch3 | Desktop |
| .csv | My Documents |
| .dif | My Pictures |
| .doc | Favorites |
| .dot | Cookies |
| .dqy | |
| .iqy | |
| .mcw | |
| .oqv | |
| .pot | |
| .ppa | |
| .pps | |
| .ppt | |
| .pre | |
| .rqy | |
| .rtf | |
| .scd | |
| .sh3 | |
| .slk | |
| .txt | |
| .wk? | |
| .wpd | |
| .wps | |
| .wq1 | |
| .wri | |
| .xl? | |

**Table 6.2 Settings and System Components Migrated by Default by USMT**

| Settings and System Components Migrated |
|---|

| |
|---|
| Accessibility options |
| Classic Desktop settings |
| Dial-up connections |
| Display properties |
| Folder options |
| Fonts |
| Microsoft® Internet Explorer settings |
| Localization/International settings |
| Microsoft® Office settings |
| Mouse and keyboard settings |
| Network drives and printers |
| Microsoft® Outlook® settings and store |
| Microsoft® Outlook® Express settings and store |
| Phone and modem options |
| Regional options |
| Screen saver selection (not users' personal screen saver files) |
| Shortcuts (shell tools, network items, and so forth) |
| Sounds and audio devices settings |
| User certificates (personal, e-mail, Microsoft® Internet Explorer security, and so forth) |
| Taskbar settings |

USMT offers multiple customization options for including various file types and settings in the user state migration. Administrators should expect to customize the default set of data and settings. Customization should be performed by technical personnel with knowledge of the registry.

This chapter describes how to plan and test a user state migration, but does not describe how to use the USMT tool. For code samples to assist you in customizing the .inf files used with USMT and the Files and Settings Transfer Wizard, see the file Inf Commands.doc included on the Windows Server 2003 operating system CD in the folder \ValueAdd\Msft\USMT. For more information about using USMT, see the User State Migration Tool Reference Materials link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources/.

# Choosing a Method for Collecting User State

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in planning your user state migration is to determine the best way to collect user state in your environment (Figure **6**.2). Four user state migration methods are available for collecting and restoring user state:

- Manual migration
- Scripted-manual migration
- Centralized automation
- User-driven migration

**Figure 6.2 Choosing a User State Collection Method**




The method by which you deploy Windows XP affects which user state migration method you should choose. Table 6.3 explains how each system deployment method affects the environment into which the user state will be migrated. A clean environment reduces management and support requirements.

**Table 6.3 Effects of Windows XP Deployment Methods on User State Migration**

| Deployment Method | Effects on User State Migration |
|---|---|
| **Wipe-and-load.** The computer's hard drive is reformatted before installing Windows XP. This is the recommended deployment method when the existing hardware is sufficient to run Windows XP, because it provides a clean platform on which to restore applications and settings. | Presents a completely clean environment in which to restore user state. |
| **Parallel deployment.** The original computer is replaced with a new computer running Windows XP. This is the recommended deployment strategy when the old computer has insufficient hardware capability to run Windows XP. You can keep the original computer running until you are sure that the new computer is completely functional. | Presents a completely clean environment in which to restore user state. Commonly used when deployment of Windows XP is timed with computer replacement, as in lease rollover. |

| | |
|---|---|
| **Operating system upgrade.** The original computer is upgraded using the Upgrade option during the setup phase of Windows XP deployment. This leaves the user's files, folders, settings, and installed applications intact. | Does not provide a clean environment, thereby increasing support and management costs. The migration of System Policy, registry settings, files, drivers, DLLs, and folder hierarchies can cause problems and nonstandard installations. Not recommended in production environments. |

When determining the best user state collection method for your situation, weigh these factors:

- The size of your organization
- The number of users to be migrated
- The level of desktop management already in place
- The uniformity of file locations on workstations
- The type of technical personnel available to assist in the migration
- The amount of time you can dedicate to the migration process

Each migration method is particularly well suited to specific scenarios. It is likely that a large corporate deployment will involve several of these scenarios and employ a mixture of migration methods.

# Manual Migration of User State

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a manual migration, an onsite technician personally attends each computer, typically performing these tasks:

- Ensures that the user's computer is ready for migration (for example, checks to see whether all important files are in the folders that are being migrated).
- Collects the user state by running either USMT (the Scanstate.exe command-line tool) or the Files and Settings Transfer wizard.
- Deploys Windows XP either by providing a new computer running Windows XP or by doing a wipe-and-load deployment of Windows XP. (Remote Installation Services [RIS] provides a convenient way to deploy a common Windows XP image.)
- Restores the user state by again running either USMT (the Loadstate.exe command-line tool) or Files and Settings Transfer wizard. The same tool that is used to collect user state must be used to restore it.
- Is available to help with any issues while the user checks to make sure that everything has been migrated properly.

Because technical labor costs in manually collecting state data can be very high, it is often beneficial to combine manual collection with the use of automated scripts.

Table 6.4 summarizes advantages and disadvantages of manually migrating user state. Because of the noted disadvantages, a strictly manual approach is not recommended.

**Table 6.4 Advantages and Disadvantages of Manual User State Migration**

| Advantages | Disadvantages |
|---|---|
| • A technician is available to deal with unexpected problems.<br>• Users are reassured by having a person to ask questions of during the migration. | • Expensive because of high technical labor costs.<br>• Slow because a technician must visit each computer individually.<br>• Higher chance of human error than with automated methods.<br>• Does not scale to distributed or remote office scenarios. |

# Scripted-Manual Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By supplementing a manual deployment with scripting, you can reduce the costs of an exclusively manual approach while providing the flexibility to handle special situations.

Scripting speeds up the process of migrating user state on individual computers, enabling a technician to migrate user state on multiple computers in the same physical area simultaneously. This greatly reduces the likelihood of human error, yet maintains the advantage of having someone onsite to deal with unexpected problems. During scripted-manual migrations, you also can use less skilled technicians for the onsite phase of the migration.

Table 6.5 summarizes the advantages and disadvantages of using a scripted-manual approach for migrating user state. Scripted-manual migration is the recommended approach for a parallel deployment.

**Table 6.5 Advantages and Disadvantages of Scripted-Manual User State Migration**

| Advantages | Disadvantages |
|---|---|
| • A technician is available to deal with unexpected problems.<br>• Users are reassured by having a person to ask questions of during the migration.<br>• Lower chance of human error than with a purely manual migration method. | • Requires that a technician be onsite, with physical access to each computer that is being migrated.<br>• Requires that script files be created.<br>• Does not scale to distributed or remote office scenarios. |

The manual-scripted migration process for a wipe-and-load deployment, in which the computer's hard drive is reformatted before the new operating system is installed, is slightly different from the process for a parallel deployment.

**Migration Process in a Wipe-and-Load Deployment**

In a wipe-and-load Windows XP deployment, use this combination of scripted and manual steps:

1. Create scripts to collect and restore user state.
2. Have a technician perform the following tasks at each computer:
   a. Run the script for collecting user state.
   b. Format the computer's hard drive and run RIS to install the new operating system image.
   c. Log on as the administrator, and run the restoration script.
3. Have the computer's user log on and check for proper restoration of data and settings.

**Migration Process in a Parallel Deployment**

Perform these steps to migrate user state in a parallel deployment:

1. Create two scripts for each computer that is to be replaced, one for collecting user state and the other for restoring user state.
2. On the computer that will be replaced, run the script for collecting user state.
3. On the new computer:

a. Install Windows XP.

b. Log on as the administrator, and run the script for restoring user state.

4. Have the computer's user log on to the new computer and check for proper restoration of data and settings.

# Centralized Automation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With centralized automation, you can extend the efficiencies available through the scripted-manual method for migrating user state. To centralize automation of user state migration, you refine the user state collection and restoration scripts to such a degree that no onsite input from a technician is required. IT technicians can deploy the scripts to targeted computers from a remote location.

Centralized automation enables enormous cost savings and provides a common migration experience corporation-wide. Centralized automation does not work well in parallel deployments because of the complexity in determining target and destination computer addresses for a large number of computers, but it is ideal for wipe-and-load deployments.

Table 6.6 summarizes the advantages and disadvantages of using centralized automation to migrate user state. Centralized automation is the ideal solution for wipe-and-load deployments.

**Table 6.6 Advantages and Disadvantages of Centrally Automating User State Migration**

| Advantages | Disadvantages |
|---|---|
| • Allows simultaneous migration of user state for large numbers of computers (limited only by network bandwidth and server storage).<br>• Produces results that can be replicated.<br>• Produces a common user experience.<br>• Scales well to distributed or remote office scenarios. | • Requires that script files be created.<br>• Does not work well in parallel deployments. |

The key challenges in the centralized automation of user state migration are:

- Targeting and deploying scripts so that they run on the user's computer in the appropriate context.
- Associating user state with a specific computer.
- Creating scripts that create a temporary store for each user's state and then access that store to restore the state on the destination computer.
- Automatically deploying the Windows XP image with no one attending the remote computer to accept the RIS image.

**Targeting and deploying scripts to run in the appropriate context** Follow these rules when targeting and deploying automated scripts during user state migration:

- The collection script must run under the user's logon account.
- The restoration script must run securely in the administrator's context with no one attending to the remote computer to log on.
- The computer's user should not be using the computer when the script is run.
- No applications can be running when the script is run.

Several options are available for automatically running the script at a specific time and under in the appropriate context. It is best to use a management solution such as Microsoft® Systems Management Server (SMS) for this. SMS provides advanced targeting options, contains software deployment structural components, and can target packages to run at specific times in specific contexts. Other options include logoff scripts, e-mail that includes the script (which automatically shuts down the mail

client), or deployment automation delivered by way of a Web site.

For more information about Systems Management Server, see the SMS Product Information link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Associating user state with a computer** When deploying a standard image to a series of computers, plan how to discover which user's state to restore to which computer. Two common methods for identifying computers are by the media access control (MAC) address for the network adapter or by the serial number of the processor.

For example, you can write a script that collects both the network adapter MAC address of the computer and the logon name of the user. This data pair is stored in a central database along with the mapping of the user state storage location. When restoring user state, the script looks up the network adapter MAC address to find the user's logon name and user state storage location, and restores the user state to the appropriate computer.

**Storing and restoring the user state** In centralized automation, you can use the same scripts that collect and restore user state in the scripted-manual method, with these adjustments:

- The collection script must be able to create a separate subdirectory for storing each user's state during the migration. Appending the user's logon name to the root storage path (for example, \\State*Username*) is a good solution. If the user has multiple computers, use both the computer name and the user's logon name (for example, \\State*Username**Computername*).

- The restoration script must read the user state storage path from the central database that the collection script wrote to and restore the user state from the appropriate storage location.

**Automating Windows XP image installation** Microsoft offers several options for deploying operating systems. For information about your options for automating the deployment of a Windows XP image, see "Choosing an Automated Installation Method" in *Automating and Customizing Installations* of this kit.

# User-Driven Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Perform a user-driven user state migration when no central management is in place for a deployment, or when users connect to the organization's network from a remote location.

In a user-driven migration, have the user run the Files and Settings Transfer Wizard to collect and restore user state. The user can use the same customized .inf files that you might use with USMT in other types of migrations. With the Files and Settings Transfer Wizard, the user can decide which applications and components to migrate (although you set the defaults), and can add and remove files and folders from the set to be migrated. Be prepared to offer these users some training on using the wizard and the .inf files.

Table 6.7 summarizes the advantages and disadvantages of a user-driven migration of user state. User-driven migration is the recommended approach for nonmanaged environments and remote users.

**Table 6.7 Advantages and Disadvantages of User-Driven User State Migration**

| Advantages | Disadvantages |
|---|---|
| • Controlled tool better enables users to drive their own migration.<br>• Can use the data captured using customized .inf files. | • IT staff are not involved in the actual migration; relies on the user's judgment.<br>• No standard storage location.<br>• Minimal control over what is migrated, resulting in a less standardized desktop. |

# Identifying Migration Content for Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After determining how to perform the migration, identify which user files and folders to migrate, and then identify key user settings to migrate, as shown in Figure 6.3. Consider which data and settings are worth migrating and which do not provide sufficient benefits to justify the cost of migration.

**Figure 6.3 Identifying Migration Content**




Part of identifying migration content is recognizing that you must provide a transitional storage place for everything that you migrate. Be sure that you have enough room to store the aggregate content during the migration. For information about estimating required storage capacity, see "Determining Storage Requirements for User State" later in this chapter.

# Identifying User Data to Migrate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When identifying what data to migrate, consider these questions:

- Do users save their files in a single folder or in files scattered across a hard disk?
- Which data files do the users work with regularly?

One way to determine which user data to migrate is to identify folders to migrate based on known locations. These can be locations that the system is aware of, such as the My Documents folder and Favorites, or locations that the organization-specifies, such as \EngineeringDrafts or C:\Data.

Another way is to identify the applications that the users use and then look for files with corresponding file types. Organizations commonly use an e-mail package and productivity suite such as Microsoft Office. These applications typically use specified file name extensions. For example, Microsoft® Word primarily uses the .doc file name extension. However, Word also uses file types such as templates (.dot files) and hypertext files (.htm files).

If you use this method to identify files to migration, create a list of important file types based on applications that your organization uses. A good starting point for identifying the file types to migrate is to look at the registered file types on the standardized desktop image that you will install. The registered file types are listed in **Folder Options**.

**To view a list of registered file types**

1. Double-click the **My Computer** icon on the desktop.
2. On the **Tools** menu, click **Folder Options**.
3. Click the **File Types** tab to display the registered file types.

**Important**

- Do not attempt to migrate the applications associated with the files. Instead, reinstall the applications from a software distribution point, or include them in the standard desktop image.

# Identifying the User Settings to Migrate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Consider the following questions when identifying which user settings to migrate:

- Are you moving toward a more managed environment? If so, which settings will users be able to change in this new environment?
- Which settings do the users need to get their work done?
- Which settings make the work environment comfortable for users, allowing them to be more productive?
- Which settings will reduce support desk calls after the migration?

# Identifying Key Settings for Users' Productivity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

List the important settings that the user needs to become productive immediately after the migration. These settings might include an e-mail server and account, a remote access connection, an Internet connection, and accessibility features. One good place to find the relevant settings is in your organization's system configuration handbook for new users.

Locating application-specific settings can be time-consuming, because various applications store settings in different locations. Therefore, limit your list to settings that the user must have to maintain productivity.

Some applications provide tools that scan the registry and then display settings and their storage location in a format that is easy to read. For other applications, you must compare registry entries before and after an installation to trace the settings.

Typically, the user settings for an application are stored in the registry in the subkey HKEY_CURRENT_USER\SOFTWARE*Companyname**Application*. To compare before and after images of the registry in order to find the updated entries, you can use the Sysdiff.exe tool, documented and available for download from the Resource Kit Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources/.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference on the *Microsoft® Windows®* Server 2003 Deployment Kit companion CD or at http://www.microsoft.com/reskit.

# Evaluating Costs and Benefits of Migrating Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is not always cost-effective to collect and restore all user-specific settings in the registry. In deciding which settings to migrate, weigh the cost of lost productivity while users recreate their settings against the IT costs of migrating them.

**To determine the cost-effectiveness of collecting settings:**

1. Perform the following calculations:

    a. Determine the number of users whose user state will be migrated.

    b. Multiply that number by the average time that it takes a user to reconfigure his or her settings.

    c. Multiply the result by the users' average hourly wage.

2. Compare this cost with the costs involved in tracing, collecting, and restoring the settings.

Table 6.8 shows user cost calculations for an enterprise with 5,000 users who have an average hourly wage of $20, based on an estimated reconfiguration time of 2 hours per user. The total user costs for not migrating user settings are compared with estimated IT costs for migrating the settings, shown in the final column.

**Table 6.8 Sample Cost Comparison for Migrating vs. Not Migrating User Settings**

| Number of Users | User Reconfiguration Time | Average Hourly Wage | Total User Costs (No Migration) | Estimated IT Costs for Migration |
|---|---|---|---|---|
| 5,000 | 2 hours | $20.00 | $200,000.00 | $50,000 |

Less measurable, but equally important, is the time that IT professionals and the Help desk staff spend in helping users to reconfigure desktops when their personal settings are not migrated.

Not migrating settings also can lead to lost productivity and decreased morale. Familiar desktop settings help users learn about the new system more easily, and reduce the potential for help desk calls. While the value of migrating some personal settings might not be immediately apparent, it is worthwhile to capture the settings that make new computers familiar and comfortable to users — for example, desktop settings, display settings, and folder options. If such settings are not migrated, productivity decreases while users adjust these settings to the way that they prefer them.

# Creating a Detailed Plan for User State Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After selecting a collection method and identifying the content to migrate, review the technical details involved in a successful migration and devise a detailed migration plan that addresses storage and security issues, resolves changes in registry settings from earlier versions of Windows, and includes adjustments for concurrent domain migration if applicable. Include a migration schedule and detailed instructions for users in your migration plan. Figure 6.4 shows the tasks involved in creating a detailed plan for user state migration.

**Figure 6.4 Creating a Detailed Migration Plan**



# Resolving Issues with Storage and Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because the process of migrating user state consists of moving data, the first step in preparing for your migration is to identify and resolve storage and data concerns associated with the move. Determine how much intermediate storage space is required for the data that is being moved. Verify that your selection of user data and settings to be migrated is complete as well as cost-effective. Plan how to resolve any file name conflicts or other file relocation issues that might arise during the migration.

# Determining Storage Requirements for User State

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Determine how much disk space is required for an intermediate storage location for user state. Base your calculations on the volume of e-mail, personal documents, and system settings for each user. The best way to estimate these is to survey a few average desktops to estimate the size of average stores in your environment.

**Important**

- Allow a minimum buffer of 20 percent additional space in the intermediate storage location. To enhance performance, locate the intermediate store on high-speed drives. Ensure that the intermediate storage of user state is the only task that the store performs and that it has an optimized (high-speed) network connection.

**E-mail** If users deal with a large volume of e-mail or keep e-mail on their local computers instead of on a mail server, the e-mail can take up as much disk space as all other user files combined. (This is not a factor if the e-mail is stored on a server only.) Prior to migrating user data, make sure that users who store e-mail locally synchronize with their mail server.

**User documents** The types of documents that an organization uses can make a substantial difference in storage requirements. For example, an architectural firm that predominantly uses Computer-Aided Design (CAD) files needs much more space than does a law firm dealing primarily with word processing documents. If your users already store many documents on file servers through such mechanisms as Folder Redirection, and they will have access to these locations after the migration, you do not need to migrate those documents.

**User system settings** Typically, 5 MB of storage is adequate to save a user's registry settings. This requirement can fluctuate based on the number of applications installed over the lifetime of the computer, but it is rare for the user-specific portion of the registry to exceed 5 MB.

# Reviewing Selections for Data Collection and Restoration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are using the User State Migration Tool (USMT), in most cases you will customize the .inf files included with the tool to limit the operating system and application settings that are migrated and to include additional file types and folders.

When customizing the .inf files, it is important to keep a backup copy of the original files and to thoroughly test your customizations. To yield the best results, keep the rules in the .inf file as simple as possible.

For information about the default files and settings that USMT migrates, see "Tools Used in the Migration Process" earlier in this chapter. For more information about using USMT, see the User State Migration Tool link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources/.

# Dealing with File Relocation Issues

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you change the computer hard disk configuration during migration, you might not be able to restore files to the same drive or directory structure from which they were collected. For example, if you replace two small drives with one large drive, the second drive will not be available to receive the collected user data. In this case, you must relocate the files.

A relocated file might be written to a folder that already contains a file with the same name, causing a name conflict. USMT handles this problem by appending "(1)" to the original filename, and incrementing that number for each new file with the same name. For example, if two files by the name of Example.doc were written to a directory that already contained an Example.doc file, the relocated files would be named Example(1).doc and Example(2).doc.

One way to avoid file name collisions when you move files is to duplicate as much of the original path as possible in the new location. For example, if the full path and file name of the original file was D:\EngineeringDrafts\Example.doc, and the new root location is C:\Documents and Settings*Username*\My Documents, create the new path and file name C:\Documents and Settings*Username*\My Documents\EngineeringDrafts\Example.doc.

# Identifying Security Concerns

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Maintaining security during and after user state migration is a significant issue. In particular, take into consideration these issues:

- Typically, the Access Control Lists (ACLs) associated with files and folders are not migrated, so the ACLs must be restored or recreated.
- Encrypting File System (EFS) information is not migrated, and encryption that occurs during migration affects who can read the files in their new destinations.
- In some organizations, it is deemed critical to secure user state during a migration.

# Restoring Lost Access Control Lists (ACLs)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In planning user state migration, it is best to assume that access control lists will not migrate during your user state migration. Several factors affect the migration of ACLs:

- The USMT tool and the Files and Settings Transfer Wizard do not migrate ACLs — instead, default ACLs are assigned to each folder that is created on the destination computer.

- If users are changing domains during a migration, there is a good chance that the original ACLs will not work unless you use a tool such as SIDHistory as part of the user state migration process. For information about managing access control lists during a domain migration, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

- When you migrate a Windows NT workstation that uses an NTFS file system drive, ACLs for individual files often do not migrate with the files. Instead, the files inherit the default ACLs of the folder into which they are copied.

# Managing Data Encryption During Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Encrypting File System (EFS) certificate data is not migrated when you use either USMT or the Files and Settings Transfer Wizard. The two tools treat encryption differently during a user state migration:

- The Files and Settings Transfer Wizard decrypts encrypted files during migration, and does not encrypt the files when it writes them to the destination computer (unless writing them to a folder that is encrypted).

- USMT decrypts encrypted files during migration, but if the temporary store is encrypted, the file will be encrypted under the user's credentials (since Scanstate.exe is run in the user's context). In addition, if the destination folder for the migrated file is encrypted, the restored file might be encrypted and, because the file will have been written under the administrator's credentials, the administrator, not the user, will be able to read the file.

In general, assume that files are not protected by encryption during a user state migration. Furthermore, because EFS certificates are not migrated, if a file does get encrypted during the migration, the user will not be able to read the file unless the EFS certificate is recovered from the network. For information about performing this type of operation, see "Encrypting File System" in *Microsoft® Windows® XP Professional Resource Kit Documentation* (or see "Encrypting File System" on the Web at http://www.microsoft.com/reskit).

# Securing User State During Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In some organizations, keeping the user's state secure from the IT technician who is performing the migration is a potential issue. In general, of course, IT technicians should be trusted people.

If the IT technician's access to user state is a security concern for you, take these steps:

- Have the user drive the migration using either USMT or a scripted-manual method. Under the scripted-manual method, the user must be able to restore user state by logging on as the administrator.

- When securing the state in the temporary store, make sure that while the root folder might allow full user access, the individual user folders only allow access for IT staff and the owner of the folder.

- To protect data as it traverses the network, use Internet Protocol security (IPSec) or other network security protocols to secure these transfers.

# Translating and Relocating Registry Entries for Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because the name or location of some registry entries for the operating system has been changed in later versions of Windows, many registry values must be translated during migration, and others must be relocated within the registry. This is also true with different versions of some applications. For example, copying the subkey HKEY_CURRENT_USER\Control Panel\Desktop\WindowsMetrics during migration causes problems, because entries such as **IconFont** are not translated correctly.

USMT automatically translates and relocates the operating system settings for the user state that it migrates. To prevent problems with custom settings, either do not migrate entries that are unique to a specific version of an application and cause problems, or use the renaming and relocating capabilities of USMT to adjust the entries. You cannot use USMT to translate registry values: The best solution for this type of change is to write custom code.

You can use tools such as Sysdiff.exe to compare before and after images of the registry. This tool helps in finding registry changes between versions of an application or between the same version of the application running on different versions of Windows. Sysdiff.exe is documented and available for download from the Resource Kit Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources/.

**Windows XP registry restrictions**

The enhanced security of Windows XP can mean that registry settings that were accessible in the Microsoft® Windows® 95 or Microsoft® Windows® 98 operating systems are no longer accessible.

In Windows XP, a user with no administrative permissions has write access to only three locations (depending on default security settings): the HKEY_CURRENT_USER registry subtree, the User Profile, and a shared documents location. Users cannot change settings outside those locations without administrative permissions.

If an application writes settings outside HKEY_CURRENT_USER, users will not be able to run the application after migration. You should deal with these applications on a case-by-case basis. Sometimes it is acceptable to change the access rights to a part of the registry; at other times, this can grant to the user unacceptable access to the registry. The best solution is to work with the software vendor or in-house developer to determine migration requirements when introducing a revised version of an application.

# Adapt Your Plan for Domain Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain migration introduces additional issues for user state migration. Command-line parameters available in USMT make it easy to change the domain name and the user name during a migration. However, these switches do not resolve issues related to security identifiers (SIDs) during domain migration.

The Active Directory Migration Tool (ADMT) addresses issues that USMT does not handle. This tool can help you to diagnose possible problems before starting migration operations.

ADMT provides the following capabilities:

- Using task-based wizards, you can migrate users, groups, and computers; set correct file permissions; and migrate mailboxes for Microsoft® Exchange Server.
- Using the tool's Reporting feature, you can assess the impact of the migration both before and after move operations.
- The tool also provides support for parallel domains, so that you can maintain your existing domains under the Microsoft® Windows NT® version 4.0 operating system while you deploy a later version of the Windows operating system (Microsoft® Windows® 2000; Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; or Windows® Server 2003, Datacenter Edition).

For information about issues related to the migration of domain user objects and how domain migration affects the migration of user state, see the ADMT Cookbook link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Scheduling Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Consider the impact of the migration on network bandwidth. The user state traverses the network as it is being moved to a temporary storage location on a network server and again when moved to the destination computer.

To minimize the impact of your migration on users and your network, follow these general guidelines:

- Minimize the impact on the network while other people need to use it by scheduling migrations during off-peak hours.
- Trying to move too many users at a time risks network collisions. Determine the optimal number of users to migrate at a time.
- Minimize the use of network capacity by locating storage servers close to the clients (for example, on the same subnet).
- Keep in mind the simple impact of migration on your normal flow of business. Work with the teams that you are migrating to ensure that the migration will not jeopardize any crucial projects. Determine whether teams need to be migrated as a group.

# Educating Users About Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To minimize productivity loss and support costs, prior to migration, set user expectations to match the results that you obtained during pilot testing. Set expectations early and clearly to reduce user frustration and Help desk calls.

Provide a schedule that indicates when each user's computer will be migrated, as well as clear guidelines that tell exactly what the user needs to do to prepare for the migration. Your migration plan must include backing up user files, verifying that applications that use synchronization mechanisms (such as e-mail) are also backed up, and preparing for changes to the desktop.

**Preparing files and folders for migration** The period just before migration is a good time for users to get their files into a stable state. For example, if version control software is in use, make sure that all users check in all files that they have checked out. If users are supposed to save all in-progress documents in a specific network folder, make sure that the users save all relevant files to that folder.

If the user state collection process will retrieve data only within a known folder (for example, My Documents), have users move all of their important documents to that folder. If the folder that contains all of these files is a network share, no migration of the files is needed as long as the user has access to that share from the new system.

**Preparing e-mail and other applications that must be synchronized** It is recommended that users send all pending e-mail prior to migration. Along with e-mail, My Briefcase, Microsoft Outlook, Microsoft® Notes, and any other application or feature that uses synchronization must be synchronized prior to the migration.

**Preparing users for changes to the desktop** The more closely that the users' new environment mirrors their previous one, the less support they will need, and the sooner they can resume productivity. If the migration involves changes to the desktop, prepare the users for these changes.

If you will not migrate specific settings, tell users in advance which settings they will need to reenter and which related files they might need to migrate (for example, a personal photograph used as a background on the user's desktop).

If you will migrate your users to the new Windows XP Luna desktop, let them know to expect the change. If you are using USMT, you can optionally migrate their "classic desktop" if you need to delay this change until the users have had training in how to use the new Luna desktop.

# Testing Your Process for Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before rolling the migration out to a large number of computers, test the process in a controlled laboratory setting, and then run a pilot test. Figure 6.5 shows the tasks in the testing process.

**Figure 6.5 Testing Your Migration Process**




Even with thorough testing before the migration, you often can improve the process by making adjustments as the migration proceeds. It is best to move groups of users in phases. Migrate one group or team and make sure the migration is successful before starting the next group. This gives you a chance to modify your plan as needed between groups.

# Performance of Lab Tests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the lab test, match your test environment closely to your production network:

- Use the same type of server in your test environment as in your real environment. If you will migrate users from an old domain to a new domain, include both the old and new types of servers in your test environment.

- Run a test migration from at least one computer running each operating system from which you will migrate. For example, if you need to migrate user state from some computers running Microsoft® Windows® 98, some running the Microsoft® Windows NT® Workstation version 4.0 operating system, and some running the Microsoft® Windows® 2000 Professional operating system, test at least one computer running each operating system.

- Make backups of the data residing on the computers from which you are migrating user state so that you can easily reproduce any problems that you encounter. If you adjust a custom script to solve a problem, it is hard to know whether the change solved the problem if you cannot reproduce the problem.

# Performing a Pilot Test

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After thoroughly testing your user state collection, operating system deployment, and user state restoration processes, conduct a pilot test on a small group of users in a production environment.

In the pilot test, concentrate on these areas:

- Make sure that all data and settings migrate as expected.
- Note the storage space requirements for the pilot data and adjust your initial calculations accordingly.
- If unexpected problems arise, address them before going further.

As you did during lab testing, make backups of the data on your source computers so that, for testing purposes, you can easily reproduce any problems that you encounter.

Only after you are fully satisfied with the success of your pilot test should you begin a full migration.

# Additional Resources for User State Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to user state migration for Windows XP.

**Related Information**

- "Designing RIS Installations" in *Automating and Customizing Installations* of this kit for information about using Remote Installation Services (RIS).
- The at http://www.microsoft.com/reskit for information about scripting in a Windows Server 2003 environment.
- The MSDN Scripting Clinic link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about scripting in a Windows Server 2003 environment.
- The file Inf Commands.doc, included on the Windows Server 2003 operating system CD in the folder \ValueAdd\Msft\USMT.

**Related Tools**

- User State Migration Tool (USMT)

  This command-line tool is used to collect a user's documents and settings before an operating system migration to Windows XP from an earlier version of Windows and to restore them after the installation. For a reference to the commands and syntax employed in the .inf files that are used to customize the selection of files, settings, and registry entries migrated by USMT, see the User State Migration Tool Reference Materials link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Files and Settings Transfer Wizard

  The Files and Settings Transfer Wizard enables users to migrate personal display properties, folder and taskbar options, and Internet browser and mail settings, as well as specific files or entire folders, such as My Documents, My Pictures, and Favorites, without manual configuration. For more information about the Files and Settings Transfer Wizard, see Help and Support Center for Windows XP.

- Sysdiff.exe

  Sysdiff.exe is an automated installation tool that enables you to pre-install applications, including applications that do not support scripted installation, as part of an automated setup. For more information about Sysdiff.exe, see the Sysdiff.exe link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Windows 2000 Active Directory Migration Tool (ADMT)

  The Active Directory Migration Tool (ADMT) is used to migrate to the Windows 2000 Active Directory® directory service, providing a task-based wizard that is used to migrate users, groups, and computers; to set correct file permissions; and to migrate Microsoft Exchange Server mailboxes. For information about ADMT, see the Active Directory Migration Tool (Windows 2000) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Job Aids**

- "Worksheet A.34 Identifying Data and Settings for Migration" (DMEUSE_34.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.34 Identifying Data and Settings for Migration" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.35 Evaluating Costs vs. Benefits of Migrating Settings" (DMEUSE_35.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Worksheet A.35 Evaluating Costs vs. Benefits of Migrating Settings" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.36 Choosing an Intermediate Data Store" (DMEUSE_36.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.36 Choosing an Intermediate Data Store" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.37 Collecting User State Information" (DMEUSE_37.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.37 Collecting User State Information" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.38 Preparing for the Migration" (DMEUSE_38.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.38 Preparing for the Migration" on the Web at http://www.microsoft.com/reskit).

# Implementing User State Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By implementing the user data management and user settings management features that IntelliMirror® management technologies offer, you can centrally manage and control desktops, settings, and user data for large or small groups throughout your organization. Users can also move from one workstation to another and still have access to their own data, settings, and familiar desktop.

For information about deploying these technologies on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

**In This Chapter**

Implementing User State Management Overview

Planning for User State Management

Designing Configurations for Managed Desktops

Preparing for Deployment

Configuring User State Management Features

Additional Resources for Implementing User State Management

**Related Information**

- For more information about managing and reconfiguring desktops, managing user data and settings, and the infrastructure and mechanics of Group Policy, see the Windows Security Collection in the Windows Server 2003 Technical Reference on the Web (http://go.microsoft.com/fwlink/?linkid=4571).

- For more information about implementing Group Policy, see "Designing a Group Policy Infrastructure" in this book.

# Implementing User State Management Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To ensure that users' data, applications, and settings remain constant throughout the users' environments, you can use IntelliMirror technologies on the server and the clients. Active Directory® directory service and Group Policy provide the foundation for IntelliMirror. Using Group Policy to define managed desktop configurations enables IntelliMirror to deploy, recover, restore, and replace users' data, software, and personal settings.

By deploying IntelliMirror technologies to manage users' data and settings, you provide the following advantages:

- Users' data is protected, up to date, and highly accessible. Users can access their data from any computer on the network. To enable these capabilities, you can redirect specific user data folders, such as the My Documentsfolder, to a network location, and then make this location available to users for offline use.

- Users' preferences and settings for the desktop and applications are available wherever the users log on. By creating a preconfigured, customized desktop, based on the users' jobs, you can provide a consistent desktop environment and system settings, and you can reduce support calls and minimize the time and cost that is associated with replacing computers.

- You can also centrally manage software installations, updates, repairs, upgrades, and software removal by performing a Group Policy–based software installation.

You start your deployment plan by defining your objectives for managing user data and settings, and then you assess the needs and resources of your organization. You can then determine which IntelliMirror technologies meet the needs of your users and your organization. After you design the managed desktop configurations for the various types of users and computers in your organization based on their requirements, you must test your designs in a controlled environment and instruct your users about the technologies that you are deploying.

# User State Management Implementation Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process of implementing user state management typically entails four phases. Begin by developing a plan for managing user data and settings. Before beginning your design, you must fully understand your organization in terms of the users' requirements, computer roles, the current network environment, your organization's security requirements, and corporate policies. Then design your managed desktop configurations based on the requirements of the various types of users and computers in your organization. The next phase consists of thoroughly testing the configurations before you deploy them and preparing your users by informing them about the technologies you will use to configure and manage their data and settings.

Figure 7.1 illustrates how to structure your approach to implementing user state management.

**Figure 7.1 Process for Implementing User State Management**



# IntelliMirror Technology Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many users must work with network-based files, even when their computers are not connected to the network. IntelliMirror technologies, such as Folder Redirection and Offline Files, make data and settings available to users even when the network is not. This increased availability is a result of storing user data and settings on network servers, while maintaining local copies of selected files and folders.

The IntelliMirror management technologies that are included in Microsoft® Windows® 2000, Microsoft® Windows® XP, and Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition, operating systems provide directory-based change and configuration management capabilities. By using IntelliMirror technologies on the server and the client, a user's data, applications, and settings remain consistent throughout the user's environment. When you use Windows XP and Windows 2000 clients, and Windows 2000 Server, and Windows Server 2003 networks with Active Directory, you can take full advantage of IntelliMirror and Group Policy management features. Active Directory and Group Policy provide the foundation for IntelliMirror. Based on the Group Policy settings you specify, IntelliMirror can deploy, recover, restore, and replace a user's data, software, and personal settings.

For more information about Group Policy, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit). For more information about Active Directory, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

Table 7.1 describes the advantages of using the core technologies in IntelliMirror that support user state management.

**Table 7.1 IntelliMirror Features and Technologies for User State Management**

| IntelliMirror Feature | Advantages | Technologies |
|---|---|---|
| User data management | Data is protected and highly available. Users can access their data from any computer on the network. | Active Directory<br>Group Policy<br>Offline Files<br>Folder Redirection<br>Synchronization Manager<br>Disk Quotas<br>Enhancements to the Windows shell |
| User settings management | Users get their preferred desktop configuration from any computer on the network. User preferences and settings for the desktop and applications are available wherever the user logs on. | Active Directory<br>Group Policy<br>Offline Files<br>Roaming User Profiles<br>Enhancements to the Windows shell |
| Computer settings management | Administrators can use Group Policy to define the appearance and behavior of the desktop. For example, you can set options for the programs that appear on the desktop and options for the Start menu. | Active Directory Users and Computers snap-in<br>Group Policy |

| | | |
|---|---|---|
| Software installation and maintenance | Administrators can centrally manage software installations, updates, repairs, and removal. | Active Directory<br>Group Policy<br>Software Installation (an extension of Group Policy Object Editor)<br>Windows Installer |

**Active Directory**
Windows–based directory service, Active Directory, stores information about objects on a network and makes this information available to administrators and users. By using Active Directory, you can view and manage network objects on the network from a single location, and users can access permitted network resources by using a single logon. Active Directory Users and Computers Microsoft Management Console (MMC) snap-in is the recommended tool for managing Active Directory objects, including organizational units (OUs), users, contacts, groups, computers, printers, and shared-file objects.

To manage sites and services, use the Active Directory Sites and Services MMC snap-in. To administer domains and trusts, use the Active Directory Domains and Trusts snap-in.

**Group Policy**
The infrastructure within Active Directory that enables directory-based configuration management of user and computer settings on computers running Windows® Server 2003, the Windows® 2000 family, and the Microsoft® Windows XP Professional operating systems. By using Group Policy, you can define configurations for groups of users and computers, including policy settings for Windows Server 2003 registry-based policies, software installation, scripts, folder redirection, Remote Installation Services, Microsoft® Internet Explorer maintenance, and security.

The Group Policy settings that you create are contained in a Group Policy object (GPO). To create a GPO, use the Group Policy Management Console MMC snap-in (GPMC). To edit policy settings in GPOs, use the Group Policy Object Editor, which can be started from the GPMC. By using GPMC to link a GPO to selected Active Directory system containers — sites, domains, and OUs — you apply the policy settings in the GPO to the users and computers in those Active Directory containers.

For more information about Group Policy and GPMC, see "Designing a Group Policy Infrastructure" in this book. To download GPMC from the Microsoft Download Center; see the Group Policy Management Console link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**User Data and User Settings**
Before you deploy IntelliMirror technologies, it is useful to understand the distinctions between user data and user settings. *User data* describes the files that a user creates and uses; user data belongs to the user. Examples of user data include word processing documents, spreadsheets, or graphics files. Examples of user-accessed data that is *not* considered to be user data include database records that exist in a corporate database and documents that are shared by many users but that are not exclusively owned by any single user.

In a managed environment, user data is stored in files in the **My Documents** folder. Administrators can use Folder Redirection to redirect the paths of the following special folders to a network location to back up and protect the data: My Documents, My Pictures, Application Data, Desktop, and Start Menu.

*User settings* are the configuration choices — stored by the operating system or applications — that the user applies to a desktop or to applications. Typically, settings include variables such as the customized toolbar settings in an application, icon arrangement and color scheme of the desktop, mouse pointers, and language options. User settings are stored in the registry, the Application Data folder, on the Desktop, and on the Start Menu in the user's user profile.

**User Profiles**
A user profile includes a user's unique settings, such as printer connections, desktop icons, mouse settings, folder settings, and the special folders that can be redirected. (Table 7.2 lists the contents of each user profile folder.) A user profile is automatically created the first time that a user logs on.

A user profile consists of a *registry hive* and a *set of profile folders*.

**Registry hive** NTuser.dat in file form. It is loaded by the system at logon and mapped to the registry subtree HKEY_CURRENT_USER. NTuser.dat stores the user's registry-based preferences and configuration.

**Set of profile folders** Stored in the file system. User profile files are stored in the file system in %SYSTEMDRIVE%\Documents and Settings(for example, C:\Documents and Settings), in per-user folders. The user profile folder is a container for applications and other system components to populate with subfolders and per-user data, such as

documents and configuration files. Windows Explorer uses the user profile folders extensively for the user's desktop, the **Start** menu, and the **My Documents** folder.

**Table 7.2 Contents of the User Profile Folder**

| User profile folder | Contents |
|---|---|
| Application Data* | Program-specific data (for example, a custom dictionary). Program vendors decide what data to store in this user profile folder. |
| Cookies | User information and preferences. |
| Desktop | Desktop items, including files, shortcuts, and folders. |
| Favorites | Shortcuts to favorite locations on the Internet. |
| Local Settings* | Application data, history, and temporary files. The Local Settings folder and its subfolders do *not* roam with the roaming user profiles. |
| My Documents | User documents and subfolders. |
| My Recent Documents | Shortcuts to the most recently used documents and most accessed folders. |
| NetHood* | Shortcuts to My Network Places items. |
| PrintHood* | Shortcuts to printer folder items. |
| SendTo | Shortcuts to document-handling utilities. |
| Start Menu | Shortcuts to program items. |
| Templates* | User template items. |

* These folders are hidden by default.

The three types of user profiles are *local, roaming*, and *mandatory*.

**Local user profile** The default profile type that resides only on the computer at which the user is logged on. A local profile is created whenever a user first logs on to a computer. Local profiles are highly recommended for users who never connect over fast links (such as remote users) or those who need their user settings to roam to whichever computer they use to log on.

**Roaming User Profile** Ideal for users who use different computers on the network and who need to have their customized settings and data available to them at each computer they use. A roaming user profile is copied to a specified server at logoff. When the user logs on to another computer in the network, the roaming user profile is copied from the server to that computer. When the user logs off the second computer, the profile is copied back to the server, thus maintaining the most recent version of the profile on the server.

**A mandatory user profile** A profile that you create to provide specific settings for users. Any changes that users make to their desktop while they are logged on are lost when they log off. With Windows Server 2003, the recommended tool for establishing strict control over workstations is Group Policy, not mandatory profiles, because mandatory profiles are less manageable and more likely to create administrative problems.

#### Folder Redirection

Use Folder Redirection to relocate specific user folders to centrally managed shared folders on the network. You can redirect **My Documents**, **My Pictures**, **Application Data**, **Desktop**, and **StartMenu**. Users can then work with the files in those folders from any computer on the network, and the folders receive the benefits of centralized security and backup.

**My Documents** The standard folder in which user data is stored. By redirecting **My Documents** to a shared network server, the user can access all the documents from any computer. Also, important user data can be more easily backed up as part of routine system administration, requiring no action on the part of the user.

**My Pictures** The default location for pictures and images in Windows 2000. Typically, **My Pictures** is contained in the **My Documents** folder. By default, if **My Documents** is redirected, **My Pictures** is also redirected. Alternatively, **My Pictures** can be redirected independently of **My Documents**. However, it is recommended that you allow **My Pictures** to remain within

**My Documents**, unless you need to separate these folders for a specific reason such as server scalability, for example.

**Application Data** A folder in the user's profile where applications often place large amounts of data (such as a custom dictionary). By redirecting the Application Data folder, users with roaming profiles can still access files in their Application Data folder without downloading those files every time they log on.

**Desktop** A folder that includes files and shortcuts.

**Start Menu** Windows Server 2003 allows Folder Redirection to redirect the Start Menu folder.

#### Offline Files

Offline Files is a feature that complements Folder Redirection and lets users disconnect from the network and work as if they were still connected. When the computer is offline, the files and folders appear in the same directory as they did online — as if they still resided in the same location on the network. This allows the user to edit files when they are disconnected from the network. The next time the user connects to the network, the offline changes are synchronized with the shared folder on the network.

Offline Files is a stand-alone technology. You do not need to pair it with Folder Redirection. However, using the technologies together works well. To ensure that users can access their files even when the network is unavailable, it is recommended that you use Offline Files along with Folder Redirection, especially if you are redirecting the Application Data folder.

#### Synchronization Manager

File synchronization ensures that local copies of offline files match their network counterparts. When using Offline Files, users can synchronize all network resources by using the Synchronization Manager. You can set the Synchronization Manager to automatically synchronize some or all resources. For example, users can set certain files and folders to be synchronized every time they log on or off the network. The Synchronization Manager quickly scans the system for any changes, and if it detects changes, the resources are updated. Only resources that have changed are updated.

#### User Profile Quotas

To specify the maximum disk space that a user profile can occupy, use the **Limit Profile Size** Group Policy setting. This setting is in the **User Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor. When you enable this policy setting, you can set options that determine the maximum profile size, specify whether to include registry files when calculating the profile size, and define the user notification options. Profile quotas apply to all types of profiles.

For roaming user profiles, it is recommended that you use profile quotas to limit profile size, not disk quotas. If you use disk quotas to manage profile size, the users find out they have exceeded their quota *after* they log off. If you use profile quotas, the Proquota.exe tool displays a message to the users when they exceed their profile quota. Users can then manually select and delete files from one or more profile folders by using Windows Explorer or **My Computer**. Proquota.exe prevents a user from logging off from the network until the profile size is decreased. This tool is enabled when you set the **Limit Profile Size** policy setting.

#### Disk Quotas

Administrators can use disk quotas to track and manage disk space usage on the servers that contain users' redirected folders. Disk quotas apply to NTFS volumes. You can specify a disk quota limit and a disk quota warning level. For example, you can set a disk quota limit to 500 megabytes (MB) and the disk quota warning level to 450 MB. In this case, the user can store no more than 500 MB of files on the volume. You can configure the disk quota system to log a system event if the user stores more than 450 MB of files on the volume.

#### Remote OS Installation Feature

Administrators can enable remote installation of Windows 2000– and Windows 2003–based operating systems and desktop images on new or replacement computers By using Remote Installation Services, you can direct client computers to a Remote Installation Services server and install automated, customized versions of the operating system. Administrators can use Group Policy to manage Remote Installation Services by centrally setting client configuration options. The remote OS installation feature uses Active Directory, Group Policy, the Dynamic Host Configuration Protocol (DHCP) service, and Remote Installation Services.

To facilitate computer replacement, you can use Remote Installation Services to install the operating system, Group Policy–based software installation to recover applications, Roaming User Profiles to restore user profiles, and Folder Redirection to manage files centrally.

For more information about managing user profiles, Folder Redirection, Offline Files, and Synchronization Manager for Windows XP in a Windows 2000 environment, see the User Data and Settings Management link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Planning for User State Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

During the planning stage of your implementation of user state management, you need to assess your current environment. Assessment includes determining your users' business needs and the types of computer roles in your organization; estimating disk storage requirements; and evaluating network bandwidth. Based on this information and the requirements of your organization — security, for example — you can define managed desktop configurations for the various user and computer groups. Figure 7.2 illustrates the planning process.

**Figure 7.2 Planning for User State Management**



# Determining User Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, larger organizations have a wide range of user needs. In addition to knowing your organization's needs, you must understand the work needs of users before you design and deploy services that define the users' computing environment.

When you plan a deployment, determine the following:

- The number of users to be accommodated. Allocating resources, such as disk space and network bandwidth, varies depending on the numbers of users.
- The primary applications and computing services in use, such as databases and e-mail. Databases are stored in shared folders on the network and do not move with users. However, users typically have user settings that must be available to them wherever the users log on to. E-mail applications might store users' e-mail folders in network shares or on the local workstation. In either case, these folders must remain available.
- Whether mobile users switch from one workstation to another or use a single portable computer. Users who switch between workstations benefit from roaming user profiles more than those who work exclusively on a single computer.
- The security requirements of different users, applications, and data files. Security involves access controls and backup requirements.

To develop a plan for managing user data and settings, categorize users according to their job types and computer experience. Then, determine whether users need to be able to work from any computer on the network and retain access to network-based files when they are not connected to the network.

For more information about user categories, see "Planning a Managed Environment" in this book.

For job aids to assist you with the deployment processes discussed in this chapter, see Additional Resources for Implementing User State Management later in this chapter.

# Evaluating Computer Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For the purpose of managing computers, determine whether a specified computer offers the following services:

- Supports logging on by different users
- Retains or discards user profiles between user sessions
- Provides local data store for offline file copies

# Determining Storage Requirements for Users, Local Workstation, and Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To plan for disk space allocation, you need to determine the typical disk storage requirements of the users and computers in your organization.

**User Storage Requirements**
User settings typically require about 10–20 MB of storage space per user, and they are simple to plan for. User data size is typically more varied, and planning its management involves multiple factors.

**Local Workstation Storage**
The IntelliMirror technologies provide a wide range of options for configuring your data storage. The options you select have associated costs in hardware and computing overhead as follows:

- Some Group Policy configurations leave copies of roaming user profiles on workstations. This reduces logon delays, but as the number of users who share a workstation increases, the disk space that is devoted to user profile storage increases. The performance improvement typically outweighs the additional workstation storage requirements, but you have to plan for residual profile storage space.

- Redirecting large folders (My Documents, for example) out of the user's profile reduces the local disk space that is required by the user's profile. Redirecting My Documents to a shared folder on the network and leaving My Pictures in My Documents is the best option for most users.

- Implementing Offline Files on a workstation requires space for the offline file cache and the Offline Files database. Manual selection of specific files or folders for offline caching improves reliability of access to critical information while ensuring offline caches are no larger than necessary.

**Server Disk Storage**
Estimating the amount of server storage space to allocate for user data and profiles depends in part on understanding local workstation disk storage requirements. First, determine how much disk space is currently occupied by users' data and, separately, by users' profiles.

You can use the Diruse.exe command-line tool that is included with the Windows Support Tools to gather disk usage information, such as how much disk space each folder uses in a directory tree on a disk volume. For information about how to use the **Diruse** command, type `diruse/?` at the command line. You can install the Windows Support Tools from the Windows Server 2003 CD. In the \Support\Tools folder, double-click **Suptools.msi** and follow the instructions that are displayed on the screen.

To prevent large folders from being unnecessarily copied across the network along with the rest of the roaming profile when a user logs on, redirect My Documents and My Pictures. This reduces the roaming profile size substantially because redirected folders are not stored or copied as part of the roaming profile.

To exclude additional folders from the user's profile, you can use the **Exclude directories in roaming profile** policy setting. This setting is in the **User Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor. This policy blocks specified folders from roaming.

You also need to estimate the amount of server disk storage that you need for user data and profiles.

**Estimating server disk space for user data and profiles example**
You have 5,000 users who can benefit from Redirected Folders. Of these, 4,000 users also need roaming user profiles. You have determined that the average amount of disk space used for each user's personal files is 80 MB. You have also viewed the technologies available for roaming user profiles and found that a typical roaming user profile uses 15 MB. In this situation, you need approximately 400 GB disk space for user data and 60 GB disk space to store roaming user profiles. To provide a small margin, add 50 percent to these numbers, for a total disk storage requirement of approximately 600 GB for user data and 75 GB for user profiles.

Investigate profile sizes in your organization to determine average and maximum sizes for your own users' profiles. A typical

disk-quota setting for users is about 200 MB. Due to the wide variation in disk-space requirements of users, disk quotas often require a lot of adjustment. For more information about establishing disk quotas, see "Configuring Disk Quotas on NTFS Volumes" later in this chapter.

You can control the maximum size of a user's profile by using Group Policy. If a user profile exceeds the allocated size, the user must reduce the size of the profile to log off. Restrictions for profile size apply to both local and roaming profiles. For more information, see "Setting Quotas on User Profiles" later in this chapter.

# Assessing Network Bandwidth Availability and Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The availability of network bandwidth can affect how Group Policy settings are applied. By default, some policies do not process across a slow network connection. If the network link speeds between a client and the authenticating domain controller fall below the default slow-link threshold of 500 kilobits per second (Kbps), only the administrative template (registry-based) settings and security settings are applied. By default, all other Group Policy settings, including Software Installation and Folder Redirection, are not applied. When the available bandwidth between the client and the domain controller falls below this preset threshold, the client is said to be on a slow link.

If necessary, you can modify the default slow link behavior by using the **Group Policy slow link detection** policy setting for both the user and computer aspects of a GPO. The **Group Policy slow link detection** policy setting is available in both **Computer Configuration** and **User Configuration**, under the **Administrative Templates\System** node. You can also adjust the Group Policy extensions that are processed below the slow link threshold. However, depending on your situation, it might be more appropriate to place a local domain controller at a remote location to serve your management requirements.

For more information about slow links, see "Designing a Group Policy Infrastructure" in this book. For more information about Group Policy mechanics and infrastructure, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

Is it important to have sufficient network bandwidth available between servers and workstations when you deploy Roaming User Profiles, Offline Files, and Folder Redirection. It is also recommended that the servers to which workstations connect for this data are on a fast network link. Check your network configuration for ways to minimize network routing hops when accessing frequently needed data. Keeping the needed data and the user on the same subnet improves performance.

When evaluating network bandwidth for Folder Redirection, be aware that slow link detection is based on link speed between the workstation and the domain controller to which the workstation is connected. If the link to the domain controller is fast but the link to a shared folder that contains redirected folders is slow, users might experience slow performance when accessing redirected files. Selecting these redirected files for offline caching does not improve file access speed in this case because the slow link is not detected. As a result, the Offline Files cache is not brought online.

Logging on the first time after having a folder configured for redirection might be slow for the user because all the files are being copied from the user's local drive to the shared folder on the network. The time this takes depends on the amount of data being copied, the local computer's disk I/O speed, the available network bandwidth, and server performance. You can minimize logon time by doing the following:

- Locate the validating domain controller in close proximity to the workstations that it needs to validate.
- Simplify or remove existing user logon scripts. Group Policy can replace many conventional logon script settings. Group Policy creates additional download time at logon; however, this is insignificant if Group Policy objects (GPOs) are judiciously used (keeping in mind that the number of GPOs influences logon time). If Group Policy is used for new folder redirection or software installation, logon time for those actions increases.
- Keep roaming user profiles as small as possible. Folder Redirection and profile quotas are common tools to manage roaming user profile size. For more information about managing roaming user profile size, see "Recommendations for Folder Redirection and "Setting Quotas on User Profiles," later in this chapter.

Understanding network traffic can help you design your network and servers to balance network loads. In some cases, it is cost effective to add another server or subnet to redistribute load; in other cases, making simple changes to server configuration can produce similar benefits.

# Designing Configurations for Managed Desktops

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use IntelliMirror and Group Policy to design well-managed desktops that enhance user and computer efficiency. To take full advantage of these benefits, it is recommended that you define and set up default user configurations.

A managed desktop configuration must be carefully adapted to the target users' applications, tasks, and locations. It can also increase productivity by preventing users from making system changes that can cause downtime. By establishing managed configurations, you make troubleshooting or replacement easier and reduce the cost of support.

IntelliMirror and Group Policy are designed for environments where administrators need to centralize tasks such as the following:

- Managing new users
- Managing mobile users
- Creating lightly managed desktops
- Creating highly managed desktops
- Managing multi-user desktops
- Replacing computers

Figure 7.3 describes the process for deploying managed desktop configurations that you select.

**Figure 7.3 Designing Managed Desktop Configurations**




For more information about user and computer types and their management, see "Planning a Managed Environment" in this book.

# Creating a New User Desktop Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

IntelliMirror, Group Policy, Windows Installer, and Remote Installation Services significantly streamline adding new users and their computers to your network. You can use these technologies to add a new managed user.

**Example** A new user logs on to a new computer and finds shortcuts to documents on the desktop. These shortcuts link to common files, data, and URLs such as the employee handbook, the company intranet, and appropriate departmental guidelines and procedures. Settings such as desktop options, application configurations, and Internet settings are configured to your organization's standard. As users customize their personal desktop environment (within boundaries you define), their changes are added to the initial environment. For instance, the user might change the screen resolution for better visibility or add shortcuts to the desktop.

In the preceding example, a default domain profile and Group Policy are used to configure the new user's environment, based on job requirements. The advantage of using a default domain profile is that all new users start from a common, administrator-defined configuration in an existing domain structure. You create a customized domain profile that applies to all new domain users the first time they log on, and they receive the customized settings from this profile. Then, as a user personalizes desktop settings and items, the new settings are saved in the user's profile that is stored locally, or in the case of a roaming user profile, in a predetermined location on the network. By implementing a default domain profile in conjunction with roaming user profiles, you provide users with the necessary business information as a starting point and also allow them to access their settings whenever and wherever needed. Finally, you can use Folder Redirection to redirect a user's My Documents folder to a network location, so that the user's documents are securely stored on a network server, which can be backed up regularly.

The administrator uses the Software Installation extension of Group Policy to assign Microsoft® Word to a user or a specific group of users. The new users log on for the first time and see that the software required to do their jobs is listed on the Start menu. When the user selects Microsoft Word from the Start menu, or double-clicks a Word document, Windows Installer verifies that Word is installed on the local computer. If it is not, Windows Installer downloads and installs the necessary files for Word to run and sets up the necessary local user and computer settings for on-demand installation.

For more information about Software Installation, see "Deploying a Managed Software Environment" in this book. For more information about Windows Installer, see *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

# Designing Managed Configurations for Mobile Users

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many organizations have *mobile users* — traveling employees who often use portable computers. Typically, mobile users log on to the same computer. Sometimes they connect by using a high-speed line and sometimes by using a low-speed (or dial-up) line. Some mobile users never have a fast connection. Mobile users fall into two main categories:

- Users who spend the majority of time away from the office or have no fixed office. Typically, these users connect by using slow links, although they might have occasional LAN access to their logon server, data servers, and application-delivery servers.

- Users who spend most of their time in an office but occasionally work at home or in another location. The majority of their network access is at LAN-speed, but they occasionally connect by using remote access or remote network links.

Despite the apparent differences between these two types of users, you can often accommodate them with a single configuration. However, you might want to create a slightly different GPO for users who spend the majority of their time out of the office.

Mobile users are often expected to provide much of their own computer support because on-site support is not available. For this reason, you might want to elevate those users' permissions so that they can install printers, for example.

However, you can also restrict mobile users from making system changes that might damage or disable their computers. For example, you might restrict mobile users from altering certain Internet Explorer settings or adding unapproved hardware devices. If those users need access to some MMC administration snap-ins, you can make a restricted set of those tools available to them.

Mobile users expect transparent access to the most critical parts of their data and settings even if their portable computer is not connected to the network. For example, they roam to desktop computers while their portable computer is in use to read mail while they are in a remote office. Finally, mobile users frequently disconnect their portable computer from the network without logging off and shutting down. This is more likely to happen when they use the hibernate and standby features of Windows XP Professional.

IntelliMirror provides several tools that simplify managing mobile users. By using the management tools of user data and settings, users can work on files offline and automatically update network versions of those files when they reconnect to the network. By using the Offline Files feature, users can work on network files when they are not connected to the network. Synchronization Manager coordinates synchronization of any changes between the offline version of a file and the network version.

**Note**

- If users are likely to disconnect from the network without logging off, it is recommended that you set Offline Files to periodically synchronize in the background. If Offline Files is set to synchronize only when users log off, users' files might not be current. You might also want to educate users to manually synchronize their data before disconnecting from the network to ensure that all files are current.

Synchronization Manager also helps manage multi-user network files. If multiple users modify the same network file, Synchronization Manager notifies the users about the conflict and offers several resolution methods. The users can save the network version, their local version, or both versions. If both are to be kept, the user is asked for a new file name to store one of the versions so that uniqueness is maintained.

Software Installation for mobile users requires some additional planning. You can make sure that all important software components, defined by you or the user, are completely installed initially. This allows users to access the necessary software even when they are not connected to the network. That means that prior to these users leaving the office, you must ensure that all relevant features within the application are installed locally and are not just advertised. For example, make sure the spelling checker for Microsoft® Office is locally installed so that the user does not trigger on-demand installation of this feature while offline.

It is not recommended that you publish software for mobile users who connect over slow links. Additionally, when mobile users connect over a slow link, user-assigned software effectively behaves the same as if you published it for these users. If you

set the Group Policy slow-link detection setting to the default in the user interface, the software is not installed on demand. However, in the Group Policy setting for slow-link detection, you can define the connection speed that you consider to be a slow link.

**Note**

- It is recommended you treat any link that is slower than local area network (LAN) speed as a slow link.

If you determine that it is appropriate for mobile users to download software from a remote location and they experience difficulty staying connected when downloading the software, you can verify that the connection speed and Group Policy settings are set appropriately in the **Group Policy slow link detection** setting in **Computer Configuration/Administrative Templates/System/Group Policy** or **User Configuration/Administrative Templates/System/Group Policy**.

Typically, a mobile user has a single portable computer and does not roam between portable computers (unless the computer is replaced). However, roaming user profiles give some protection against the failure or loss of a portable computer and allow a user to roam to a desktop computer to connect to a fast network. If the mobile user is *not* often connected to a fast network, it is better not to use roaming user profiles.

Data accessed by the mobile user typically falls into one or more of the following categories:

- Data that resides on a network server and that users want to access while not connected to the network. Users typically own this data (for example, their home directory), but shared data can also be stored on the local computer.

- Data that resides only on the network server (either not needed offline or volatile shared data that is inappropriate for storing offline).

- Data that resides only on the portable computer hard disk. Examples are policy manuals or other read-only items or large document sets that are needed offline by the user but the performance overhead of synchronizing precludes storing them on a file server. (In this case, a suitable backup mechanism is definitely needed.) Other examples are large database files or other data items that have their own synchronization mechanism, such as the offline storage feature in Microsoft® Outlook® messaging and collaboration client.

Table 7.3 summarizes desktop management features that you can use to create a mobile user configuration.

**Table 7.3 Features of a Mobile User Configuration**

| Feature | Explanation |
|---|---|
| Folder Redirection | Use to redirect the My Documents folder. This allows users to access centrally stored data and documents from anywhere. Redirected folders are automatically made available offline, to provide access when users are not connected to the network. |
| Ability for user to customize | Permit customization within certain guidelines. You can allow users to personalize their work environment but still prevent them from making changes to critical system settings. |
| Software Installation | Core applications are installed on all laptops (these applications are assigned).<br>Optional applications are available for users to install locally (these applications are published). |
| Group Policy Settings | Use Group Policy settings to create the managed environment. |

# Recommendations for Configuring Mobile Users Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, a mobile user works offsite, connects to the corporate network over a slow link from various locations, and occasionally connects to the network by using a fast link. This user also uses a desktop computer at the home office, so it is preferable that the user's portable computer and desktop computer have the same settings. To provide a consistent user experience, follow the guidelines in this section for configuring user profiles.

When configuring multiple mobile users, set up an OU for mobile computers and apply the same settings to all users in that OU.

**User Profile**

The following suggestions can guide you in planning user profile configurations for mobile users:

- If the users regularly connect to the network by using a fast link, and you want to back up their state or if they use multiple computers, use a roaming user profile.
- If the users rarely connect by a fast link, use a local profile.
- If users roam to LAN-connected computers in the domain and also have portable computers, use a roaming user profile. For portable computers, enable the **Only allow local user profiles** policy setting which is available in the **Computer Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor. Note that a Computer Configuration Group Policy setting takes precedence over a User Configuration setting, so the user receives the user setting on a desktop computer and the computer setting on the portable computer.

**Roaming User Profiles**

By default, roaming user profiles do *not* roam over slow links. It is recommended that you configure roaming user profiles only for mobile users that log on *frequently* over a fast link, so that their profiles are saved to the server. This provides backup of the user state. However, if the user *never* logs on to the network over a fast link, or does so infrequently, it is *not* useful to provide roaming user profiles for that user.

**Folder Redirection, Offline Files and Slow Links**

If mobile users frequently connect at high speed, Folder Redirection and Offline Files are beneficial for backup and synchronization. By default, Folder Redirection and Offline Files are *not* applied over slow links. You can override the default by using the **Group Policy Slow Link Detection** policy setting. This setting is in the **User Configuration\Administrative Templates\System\Group Policy** node of the Group Policy Object Editor.

### Important

- If you change the default setting for Group Policy slow-link detection, all group policies are affected, and all Group Policy settings are applied over a slow link, including policies that install software. This is not appropriate for most users.

# Designing Highly Managed Desktops

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, a highly managed desktop is the most appropriate workstation configuration for users in an office environment. Users' personal files are cached locally, so they are available if a server becomes unavailable. Authoritative versions of the files are retained on central servers, so backup and physical security are easier to protect. The user's personal configuration and customizations are stored on central servers, so replacing a workstation does not alter a user's computing environment. Also, users can receive their own personal computing environment on another user's workstation. After the visiting user logs off, and the regular user logs on, the regular user's personal computing environment is restored. To implement this process, you must configure Roaming User Profiles, Folder Redirection, Offline Files, and server disk quotas.

To create highly managed desktops, perform the following tasks:

- Set up Folder Redirection to redirect My Documents to a shared folder on the network.

- Set up Roaming User Profiles for users who use multiple computers so that their settings follow them to any computer they use. For more information about implementing roaming user profiles, see "Using roaming user profiles" in Help and Support Center for Windows Server 2003.

- Set up Offline Files to minimize vulnerability to network status inherent in redirected folders. Designated files are cached locally to ensure that in the event of network disconnection, users can continue to work on these important files. For more information about using Offline Files to minimize vulnerability to network status, see "Configuring Offline Files" later in this chapter.

- Set up Disk Quotas on the servers that contain redirected folders so that individual users cannot use too much disk space.

- To remove cached versions of roaming profiles when the users log off, apply the **Delete cached copies of roaming profiles** policy setting. This policy setting is available in the **Computer Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor.

- Make applications available to users by using Group Policy–based software management to assign or publish applications to users. You can also use Microsoft Systems Management Server to manage software if your organization is medium or large, if your users run operating systems earlier than Microsoft® Windows® 2000 Professional, or if you require more advanced capabilities for planning, distributing, and tracking software.

  For more information about managing software, see "Deploying a Managed Software Environment" in this book.

  For more information about Systems Management Server, see the SMS Product Information link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Designing Lightly Managed Desktops

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Lightly managed workstations are often a good solution for advanced users, such as software developers, who require a lot of control over their computers. To implement a lightly managed workstation, perform the following tasks:

- Set up Roaming User Profiles for users, so that their settings follow them to any computer they use.
- Set up Folder Redirection and Offline Files. Any files that are redirected must be set as Offline Files, so that the users can continue to work on their important files in the event of network disconnection.
- Set up Disk Quotas on the servers that contain redirected folders so that individual users cannot exhaust available shared disk space.

# Designing Managed Multi-User Desktops

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A multi-user desktop is a managed desktop that many users can use. In this configuration, users can partially configure the desktop. The multi-user desktop is ideal for computers that are set up for public access, such as computers in a library or school. The multi-user desktop gets heavy traffic and must be reliable and unbreakable. It must also be flexible enough to allow some customization. For example, students might need access to customized applications for instructional purposes and need to install applications that the network administrator publishes. You might also want to allow users to change their desktop wallpaper and color scheme.

However, you need to control the configuration of hardware and connection settings to maintain security. Multi-user computers often require certain tools, such as word processing software, spreadsheet software, or a development studio.

With the multi-user desktop configuration, users can do the following tasks:

- Modify Internet Explorer and the desktop
- Run assigned or published applications
- Configure some Control Panel options

However, users cannot:

- Use the **Run** command on the **Start** menuor at a command prompt.
- Add, remove, or modify hardware devices.

In the multi-user environment, a user is unlikely to return to the same computer. Therefore, local copies of roaming user profiles that are cached on the computer are removed after the user logs off if the roaming user profile settings are successfully synchronized back to the server. Roaming user profiles use the My Documents and Application Data folders that are redirected to a network folder. However, users can log on even if their network profile is not available. In this case, the user receives a temporary profile that is based on the default profile.

The multi-user computer is assigned a set of core applications that is available to all users who log on to that particular computer. In addition, a wide variety of applications can be assigned to or published for users. Users cannot install from a disk, a CD-ROM, or an Internet location. To conserve disk space on the workstation, most applications must be configured to run from a network server. Start menu shortcuts and registry-based settings are configured when the user selects an application to install, but most of the files that are associated with an application remain on the server. The shared folders that store the applications can be configured for automatic caching for programs so that application files are cached at the workstation on first use.

Table 7.4 shows the desktop management features that are used to create a multi-user computing environment.

**Table 7.4 Features of a Multi-User Desktop Configuration**

| Feature | Specifics | Explanation |
|---|---|---|
| Multiple users | Per-user logon accounts | Users share this computer at different times. Each user has a unique logon account. |
| Roaming User Profiles | Yes | User settings are available from any computer and administrators can easily replace computers without losing their configuration. When the user logs off, the local cached version of the profile is removed to preserve disk space. |

| | | |
|---|---|---|
| Folder Redirection | My Documents, Desktop, and Application Data folders | User data is saved in shared folders on the network, and Group Policy prevents users from storing data locally. |
| Ability for user to customize | Some | Most of the system is locked down, but some personal settings are available. |
| Assigned applications | Multiple | Core applications that are common to all users are assigned to the computer. Other applications are available for on-demand installation by means of user assignment. |
| Published applications | Multiple | Applications are available for users to install from **Add or Remove Programs** in Control Panel. |
| Group Policy settings | Yes | Group Policy settings are used to create the managed environment. |

# Designing Managed Desktops for Replacing Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a user receives a new or different computer, it can cause a time-consuming interruption in productivity. It is important that users regain productivity in the shortest possible time and with a minimum of support. This can be accomplished by storing user data and settings independently of any specific computer. By using the Group Policy features Roaming User Profiles and Folder Redirection, you can ensure that the user's data, settings, and applications are available wherever the user logs on to the network.

To further simplify setting up a new managed computer on your network, use Remote Installation Services to create standardized operating system configurations. Remote Installation Services allows you to create a customized image of a Windows XP Professional or Windows 2000 Professional desktop from a source computer. Then, you can save that desktop image to the Remote Installation Services server. The image can include the operating system alone or a preconfigured desktop image, including the operating system and a standard, locally installed desktop application. You can use that preconfigured image to set up multiple desktops, saving valuable time. Create as many standard desktop images as you need to meet the needs of all types of users in your organization. For more information about using Remote Installation Services, see the Remote Installation Services (RIS) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

These IntelliMirror technologies can work together when a computer fails, as the following example illustrates.

- **Example**. A user's computer suddenly undergoes a complete hardware failure. The user calls the internal support line. Presently, a new computer, loaded only with the Windows XP Professional operating system, arrives. Without waiting for technical assistance, the user plugs in the new computer, connects it to the network, starts it, and can immediately log on.

  Because roaming user profiles are enabled, the user finds that the desktop takes on the same configuration as the computer it replaced — the same color scheme, screensaver, and all the application icons, shortcuts, and favorites are present. Because Folder Redirection and Software Installation are enabled, the user can access data files on the server after the appropriate applications are automatically installed.

# Preparing for Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can minimize expense, effort, and user inconvenience by implementing your deployments in phases. You can further prepare for user satisfaction by communicating clearly to your users the deployment plan and schedule.

The IntelliMirror tools for managing computing environments for users are powerful, and they can have unwanted consequences if applied incorrectly. Therefore, it is essential that you thoroughly test all your deployments of Group Policy and User Data and Settings before you deploy them on a large scale. Figure 7.4 highlights the process of preparing for deployment of your design.

**Figure 7.4 Preparing for Deployment**



# Performing Controlled Testing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you choose the IntelliMirror technologies to deploy, set up a test environment to thoroughly test your configurations.

For testing, you typically need a domain controller, a file server, two or more workstations, and possibly a mobile computer that is connected by a slow link. These requirements might increase depending on the scale at which you want to test your proposed configuration. Build on these recommendations until you have included all of the requirements for your organization.

# Piloting a Limited Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you complete tests in a controlled environment, you are ready to conduct a pilot deployment to expose and resolve unexpected problems before you deploy on a large scale. Keep the group of users that you select for your pilot deployment to a manageable number (10 to 20).

On the first day of your pilot deployment, modify one or two user accounts from a single group. On the second day, after resolving any new problems, apply the configurations to 10 or 20 users in that group. Expand the pilot deployment gradually.

# Preparing Users

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The deployment of user data management and user settings management policies is more noticeable to users than the deployment of many other technologies because those settings directly affect the look and use of users' desktops. Users who know what to expect from a deployment are better prepared to deal with temporary inconveniences and alterations in their work environment. Pilot users can help you inform other users to help them adapt to changes in their work environment. Having pilot users share their understanding and expertise throughout your user community improves overall user satisfaction and reduces the number and cost of technical support calls.

# Teaching Users about Roaming Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, roaming user profiles require very little user training. Users who connect to the network over a LAN appreciate seeing the same desktop settings when they log on to any computer. Users need only a few common concepts to understand roaming profiles.

**Logging on to and off from more than one computer**
When a user with a roaming user profile uses multiple computers simultaneously, the settings from the last time a user logs off from the computer are preserved. This means that if a user is logged on to two computers simultaneously, the user should log off *last* from the computer whose configuration the user wants to preserve. *Only* the last copy of the user profile is preserved.

**Roaming user profiles on a fast network link only**
By default, roaming user profiles roam over a fast network link only. Users cannot receive their roaming user profiles over a slow link, as a message explains to them. If users do not need to see the message, you can use Group Policy to disable it, or you can set the **Timeout for dialog boxes** policy setting in the Group Policy Object Editor to 1 to make it less noticeable. (This setting is in the **Computer Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor.)

Educate users to keep their profile size to a minimum. For example, they can save shortcuts to documents on their desktop instead of saving the actual document. If you use the **Limit profile size** policy setting to manage profile sizes, teach users how to respond to the messages they receive when they exceed that limitation. If you force users to reduce their profile size before logging off, show them how to do so safely. The **Limit profile size** policy setting is available in the **User Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor.

# Teaching Users about Folder Redirection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If users are accustomed to seeing their documents in the My Documents folder, when you deploy Folder Redirection, most applications (such as Microsoft Word) appear and use the files in the same manner as they did before you redirected the folders.

**Folder Redirection and remote access to another domain**
When users connect to the network by using remote access and their computer accounts are not a part of the domain to which they are connecting, they need to know the full Universal Naming Convention (UNC) path to the shared folder that contains their files. For example: *Servername*\Redirected Folders\%*username*%\My Documents.

The logon process might access one or more files from a redirected My Documents folder during a dial-up session.

# Teaching Users about Offline Files and Synchronization Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When files are modified offline by more than one user, file conflicts can occur. Offline files are referenced to a specific version of the corresponding file on the server. When a file is modified offline, it replaces the server version of that file when the user reconnects to the network. Synchronization Manager checks the server version of the file to determine if it has changed since it was last synchronized.

If the server version of the file has not changed, the local copy of the modified file is copied to the server. If the server version of the file has changed since it was last synchronized, Synchronization Manager displays the **Resolve File Conflicts** dialog box. The user is offered the following options for dealing with multiple versions of the file.

**Keep both versions** Both versions of the file are saved to the server. The user name and a version number are appended to the local file name. For example, Report.txt becomes Report (John Doe v1).txt.

**Keep only the version on my computer** The local file replaces the server file.

**Keep only the network version** The server file replaces the local file.

**View: Network Version** Open and view the server version of the file.

**View: Local version** Open and view the local version of the file.

Typically, it is a good practice is to minimize the number of users who can modify a file that is available offline. This practice prevents data loss and multiple file versions. The user can use Windows NTFS permissions to specify who can access the files. Grant *Change* permissions to users who are allowed to modify the file offline and *Read* permissions to users who need the data to be available while they are offline. Permissions still apply when users are offline.

Tell users that when the status of their network connection changes, an **Offline Files** icon appears in the notification area, and a message is displayed over the icon to notify users of the change. Users can continue to work with their files without interruption. For more information about the status of their connection, they can click the **Offline Files** icon in the notification area.

# Teaching Users about Disk and Profile Quotas

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you implement disk or profile quotas, tell your users what to expect if they exceed their quotas and how to resolve the situation. Tell them how to request a quota increase. Make sure users know how to determine how much disk space they have used and how much they have left.

When a user tries to exceed a quota on a disk volume, the disk volume appears to be full. Any files that the user cannot save to that disk volume can be saved to an alternate disk volume to which they have access.

A user who exceeds a profile quota cannot log off until the profile size is reduced to the quota allowance. The user must remove files from the roaming user profile until the combined size of the redirected folders in the user's roaming profile drops below the profile quota.

# Configuring User State Management Features

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

While the Folder Redirection, Roaming User Profiles, and Offline Files technologies are complementary, they are functionally independent. You can apply them either in combination or one at a time. The preferred method is to use these technologies in Active Directory environments with Group Policy because this approach provides the full benefits of IntelliMirror.

Using Active Directory and Group Policy to implement the IntelliMirror technologies helps you accomplish the following tasks and responsibilities:

- Centrally create and manage desktop configurations for all users and computers in your organization.
- Allow users to access their files from any location at any time by using Roaming User Profiles and Folder Redirection in combination with Offline Files.
- Manage how software is deployed and installed on computers to ensure that users have the software that they need to perform their jobs.
- Manage and enforce centralized data storage. This helps you keep important corporate data backed up.
- Replace computers efficiently by using the following technologies:
    - Remote Installation Services and Group Policy–based software installation to replace applications
    - Roaming User Profiles to recover user profiles
    - Folder Redirection to centrally store files

As you determine which technologies to deploy, consider how the deployment of one technology, such as Roaming User Profiles interacts with another, such as Folder Redirection. For example, redirected folders are not copied back and forth with Roaming User Profiles. Implementing Folder Redirection before implementing Roaming User Profiles provides two benefits:

- Converting users profiles from local to roaming causes new logon delays because folders that were previously local are copied from a shared folder on the network during logon.
- Redirecting **My Documents** before implementing Roaming User Profiles ensures that the **My Documents** folder is never copied back and forth during a logon session. This also ensures that copies of **My Documents** that are left in profiles from previous logon sessions do not remain on workstation hard drives.

Figure 7.5 shows the processes you use to implement the user state management features that you select.

**Figure 7.5 Implementing User State Management Features**




Implementing new technologies in phases has a significant advantage: Testing is simpler. Because you have introduced fewer variables in each phase, it is easier to trace unexpected results back to their source.

It is important to include in your plan adequate time for testing the technologies that you intend to implement so that you can prevent problems. For more information about testing, see "Preparing for Deployment" earlier in this chapter.

Implementing the IntelliMirror deployment plan entails selecting group policies for GPOs and then applying those Group Policy objects to selected groups of users or computers. For more information about working with GPOs, see "Designing a Group Policy Infrastructure" in this book and the Group Policy topics in Help and Support Center for Windows Server 2003.

For information about managing desktops in environments that do not use Active Directory, see "Desktop Strategies for Computers Running Windows 2000 and Windows XP Without Active Directory," later in this chapter.

# Implementing Roaming User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can reduce setup time, improve consistency, and simplify manageability of your user accounts by creating a default profile for each group of users, defined by their organizational function or role. For example, field engineers who primarily connect from remote locations and accountants who work exclusively on site are likely to benefit from different policy settings. You can start by customizing one profile to fit the general needs of users in a domain and then set the customized profile as the domain default user profile. Thereafter, you need only modify your default profile as needed for each group.

For information about deploying Roaming User Profiles on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

# Recommendations for Roaming User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following guidelines and recommendations for customizing roaming profile settings for your deployment include considerations for users roaming between different versions of the Windows operating systems.

For information about deploying Roaming User Profiles on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

**Considerations for users roaming between different versions of the Windows operating systems**
The following guidelines apply to roaming users who roam between computers running any of the following operating systems: Windows Server 2003, Windows 2000, and Windows XP:

- If you can avoid roaming between versions of the operating system, it is preferable that you do so. There is nothing inherent in roaming that causes problems, but the data that the applications put in the profile might cause unintended side effects on other versions of the operating system.
- Install the same application versions on all client computers.
- Make sure that the applications are installed using the same path and drive on all client computers.
- Make sure that the different versions of the operating system are installed on the same %systemdrive% and in the same %windir% location on all clients.
- If users roam between Microsoft® Windows® NT 4.0–based clients and Windows XP–, Windows 2000–, or Windows Server 2003–based clients, consider setting the profile path during the installation of Windows XP, Windows 2000, or Windows Server 2003. Differences in the default profile path (%windir%\Profiles versus %systemdrive%\Documents and Settings) may cause problems for users that roam between Windows NT 4.0–based clients and Windows XP–, Windows 2000–, or Windows Server 2003–based clients. To minimize the chance of problems, make sure the path to the profile is the same on all clients.

**Redirect My Documents out of user's roaming profile**
To reduce logon delay, use Folder Redirection to relocate large folders, such as **My Documents**, to shared folders on the network. Redirected folders are not copied as part of the Roaming User Profile during logon or logoff.

**Do not use Encrypting File System with roaming user profiles**
Encrypting File System (EFS) is *not* compatible with files in roaming user profiles. If you use EFS to encrypt profile folders or files, the user's profile does not roam.

**Do not place roaming user profiles on a server volume that uses NTFS compression**
It is recommended that you do not place user home folders and roaming profiles on a volume that uses NTFS compression. The frequent file creation and deletion can cause a compressed volume to become fragmented.

**Combine Roaming User Profiles with Folder Redirection**
When a user gets a roaming user profile, the data and settings that are stored in the user's local profile directories become mobile. The profile data is stored on a server and copied to a workstation when the user logs on. Storing this data on servers simplifies data backup and security. However, because large folders, such as **My Documents**, are copied, there can be logon delays. You can avoid these delays by using Folder Redirection in combination with Roaming User Profiles.

When a user profile folder is redirected to a shared folder on the network, the folder is no longer part of the user profile. While other profile data is copied to the local workstation when the user logs on, the redirected folder remains on the server and is accessed over the network.

**Do not set disk quotas too low for users with roaming profiles**
For users who have roaming user profiles, set the disk quota limits high enough to allow the system to create a temporary duplicate copy of the user's profile. During the synchronization process, a temporary profile is created in the user's context. This process debits the user's quota. If a user's disk quotas are set too low, roaming profile synchronization might fail.

If possible, use profile quotas to manage disk space and apply disk quotas *only* if additional control is required.

**Do not use Offline Folders on roaming user profile shared folders**

To prevent synchronization problems, make sure that you turn off Offline Folders for shared folders where roaming user profiles are stored. If you do not turn off Offline Folders for a user's profile, you might experience synchronization problems because both Offline Folders and Roaming User Profiles try to synchronize the files in a user's profile. This synchronization conflict does not affect using Offline Folders with redirected files, such as **My Documents**.

# Group Policy Recommendations for Roaming User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following guidelines for setting Group Policy for roaming user profiles to help to ensure optimum functionality.

**Reducing disk space use, profile size, and logon processing**

To save disk space, you can use the **Delete cached copies of roaming profiles** policy setting to specify that locally cached versions of roaming user profiles be deleted when the user logs off. This setting is in the **Computer Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor snap-in.

To manage the user profile size, use the **Limit profile size** policy setting. This setting is in the **User Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor snap-in. You can set the maximum size of the roaming user profile and determine the system's response when the limit is reached. When you set this policy, the Proquota.exe program is enabled. This program alerts users when their user profile exceeds a predefined limit, and it prevents users from logging off until the size of the profile files is reduced.

To reduce both profile size and logon processing time, you can use the **Exclude directories in roaming profile** setting to designate certain folders to exclude from the user's profile. This setting is in the **User Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor.

**Turning off the fast logon enhancement**

With the fast logon enhancement in Windows XP, when users change from a local to a roaming profile, two logons on each computer are required for profile changes to be registered. This is because the user always logs on with cached credentials. Therefore, it takes one logon for the network to detect that the user has become a roaming user and a second logon to apply the new settings.

To ensure the best user experience, enable the **Always wait for the network at computer startup and logon** policy setting. This setting is in the **Computer Configuration\Administrative Templates\System\Logon** node of the Group Policy Object Editor.

**Limiting use of Group Policy loopback processing if you use roaming profiles**

The **User Group Policyloopback processing mode** policy setting is an advanced option that is intended to keep the desktop configuration the same regardless of who logs on. This option is appropriate for use in certain closely managed environments, such as classrooms, public kiosks, reception areas, servers, and terminal servers. When you enable the loopback processing policy setting in a GPO, you can configure user-based policy settings within that GPO, and those settings are applied regardless of which user logs on to the computer, meaning that those settings override individual user settings. When you use the **User Group Policyloopback processing mode** policy setting, you must ensure that both the computer and user portions of the GPO are enabled. The **User Group Policyloopback processing mode** policy setting is available in the **Computer Configuration\Administrative Templates\System\Group Policy** node of the Group Policy Object Editor.

By default, the user's Group Policy objects determine which user settings apply. When you enable this policy setting, you also specify either the *Replace mode* or the *Merge mode,* which determines the policies that are applied.

**Replace mode** The user settings that are defined in the computer's Group Policy objects replace the user settings that are typically applied to the user.

**Merge mode** The user settings that are defined in the computer's Group Policy objects and the user settings that are typically applied to the user are combined. If the settings conflict, the user settings in the computer's Group Policy objects take precedence over the user's policy settings.

If you disable or do not configure the **User Group Policyloopback processing mode** setting, the user's Group Policy objects determine which user settings apply.

**Caution**

- Use caution when using loopback policy processing and roaming profiles, especially if roaming users use two or more computers that run different versions of the Windows operating system, including Windows Server 2003, Windows 2000, Windows XP, and Windows NT 4.0. You might see some settings persist in the registry because applications can store policy settings in the HKEY_CURRENT_USER\Software\Policies subkey regardless of the operating system version. Windows NT 4.0 also stores some Internet Explorer policy settings in the HKEY_CURRENT_USER\Software\Microsoft\windows\currentversion\explorer\policies subkey, and those settings persist

in the registry. Windows Server 2003, Windows 2000, and Windows XP clear these registry subkeys each time before re-applying the current Group Policy object settings. However, Windows NT 4.0 does *not* clear the registry subkeys, which causes some settings to persist if you roam from a computer that runs Windows Server 2003, Windows 2000, or Windows XP.

**Additional Policy Settings That Are Related to User Profiles**

Tables 7.5 and 7.6 list additional Group Policy settings that you can set for user profiles. These policy settings are available in **Group Policy Object Editor** inthe **User Configuration\Administrative Templates\System\User Profiles** and **Computer Configuration\Administrative Templates\System\ User Profiles** respectively. For more information about the policy settings, click the **Explain** tab on each policy's **Properties** page.

**Table 7.5 Policy Settings for User Profiles: User Configuration**

| Policy Setting | Description |
|---|---|
| Connect home directory to root of the share | Restores the definitions of the %HOMESHARE% and %HOMEPATH% environment variables to those that are used in Windows NT 4.0 and earlier. If you enable this setting, the system uses the Windows NT 4.0 definitions. Along with %HOMEDRIVE%, these variables define the home directory of a user profile. The home directory is a persistent mapping of a drive letter on the local computer to a local or remote directory. |

**Table 7.6 Policy Settings for User Profiles: Computer Configuration**

| Policy Setting | Description |
|---|---|
| Delete cached copies of roaming profiles | Determines whether the system saves a copy of a user's roaming profile on the local computer's hard disk drive when the user logs off. This policy and the related policies in this folder define a strategy for managing user profiles that reside on remote servers. Specifically, these policies indicate to the system how to respond when a remote profile is slow to load. |
| Slow network connection timeout for user profiles | Defines a slow connection for roaming user profiles. If the server on which the user's roaming user profile resides takes longer to respond than the thresholds that are set by this policy permit, the system considers the connection to the profile to be slow. This policy and related policies in this folder together determine how the system responds when roaming user profiles are slow to load. |
| Add the Administrators security group to the roaming user profile share | This setting adds the Administrator security group to the roaming user profile shared folder. After an administrator has configured a user's roaming profile, the profile is created at the user's next login at the location that is specified by the administrator.<br><br>For Windows 2000 and Windows XP operating systems, the default file permissions for the newly generated profile are full control, or read and write access for the user, and no file access for the administrators group. By configuring this setting, you can alter this behavior. If you enable this setting, the administrator group is given full control to the user's profile folder. |
| Prevent roaming profile changes from propagating to the server | This policy determines if the changes a user makes to their roaming profile are merged with the server copy of their profile. If this policy is set, at logon users receive their roaming profile, but any changes users make to their profile are not merged to the users' roaming profile at logoff. |
| Only allow local profiles | This setting determines if roaming user profiles are available on a particular computer. By default, when roaming profile users log on to a computer, their roaming profile is copied to the local computer. If they have previously logged on to this computer, the roaming profile is merged with the local profile. When the users log off this computer, the local copy of their profile, including any changes they have made, is merged with the server copy of their profile.<br><br>If you enable this setting, the following occurs on the affected computer: At first logon, the user receives a new local profile, rather than the roaming profile. At logoff, changes are saved to the local profile. All subsequent logons use the local profile. |

| | |
|---|---|
| Prompt user when slow link is detected | Notifies users when their roaming profiles are slow to load, letting a user decide whether to use a local copy or to wait for the roaming user profile.<br><br>If you disable this policy or do not configure it, the system does not notify the user when a roaming user profile is slow to load. The system loads the local copy of the profile. If you enable the **Wait for remote user profile** policy, the system loads the remote copy without prompting the user. |
| Maximum retries to unload and update user profile | Applies to Windows 2000 only. This setting determines how many times the system tries to unload and update the registry portion of a user profile. When the number of trials that is specified by this policy is exhausted, the system stops trying. As a result, the user profile might not be current, and local and roaming user profiles might not match. |
| Do not detect slow network connections | Disables the slow link detection feature. *Slow link detection* measures the speed of the connection between a user's computer and the remote server that stores the roaming user profile. When the system detects a slow link, the related policies in this folder tell the system how to respond. |
| Wait for remote user profile | Directs the system to wait for the remote copy of the roaming user profile to load even when loading is slow. The system waits for the remote copy when the user is notified about a slow connection but does not respond within the time allowed. |
| Timeout for dialog boxes | Determines how long the system waits for a user response before it uses a default value. The default value is used when the user does not respond to messages.<br><br>Use this setting to override the system's default value of 30 seconds. |

# Security Recommendations for Roaming User Profiles Shared Folders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to ensure that access permissions are set appropriately on shared folders that contain user profile folders and to secure the servers in which the users' data is stored. To provide enhanced security, host the roaming profile shared folders on servers running Windows 2000 or later, use NTFS on the volumes containing the users' data, and grant share access permissions as follows.

For information about deploying Roaming User Profiles on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

**Granting profile share permissions**

A common error in user profiles is permissions that are incorrectly set. To ensure that permissions are set correctly, use the following guidelines:

- When you create the shared folders for roaming user profiles, limit access to the folder to only users who need access.
- Because a roaming profile contains personal information, such as the user's documents and EFS certificates, it is important to ensure that roaming user profiles are secure. Here are some ways you can enhance the security of roaming user profiles:
    - Restrict the shared folder to only users who need access. Create a security group for users who have profiles on a particular shared folder, and then limit access to only those users.
    - When you create the shared folder, hide the folder by putting a dollar sign ($) after the share name. This hides the folder from casual browsers and hides the folder in **My Network Places**.
    - Unless you need special permissions on the profile folder, do not create profile folders in advance for the user. Instead, allow the system to create them.
    - Assign users the minimum permissions that are required as described in Tables 7.7, 7.8, and 7.9. These tables list the required NTFS and share level server message block (SMB) permissions for roaming user profile shares and folders.

**Table 7.7 NTFS Permissions for Roaming Profile Parent Folder**

| User Account | Minimum Permissions Required |
|---|---|
| Creator Owner | Full Control, Subfolders and Files Only |
| Administrator | None |
| Security group of users needing to put data on share | List Folder/Read Data, Create Folders/Append Data - This Folder Only |
| Everyone | No permissions |
| Local System | Full Control, This Folder, Subfolders and Files |

**Table 7.8 Share level (SMB) Permissions for Roaming Profile Share**

| User Account | Default Permissions | Minimum Permissions Required |
|---|---|---|
| Everyone | Read only | No permissions |
| Security group of users needing to put data on share | N/A | Full Control |

**Table 7.9 NTFS Permissions for Each User's Roaming Profile Folder**

| User Account | Default Permissions | Minimum Permissions Required |
|---|---|---|
| %Username% | Full Control, Owner of Folder | Full Control, Owner of Folder |
| Local System | Full Control | Full Control |
| Administrators | No Permissions* | No Permissions |
| Everyone | No Permissions | No Permissions |

* No permissions is the default unless the **Add the Administrator security group to the roaming user profile share** policy setting is set, in which case the Administrators group has full control. (The **Add the Administrator security group to the roaming user profile share** policy setting requires Windows 2000 Service Pack 2 or later).

#### Hosting profile shares on servers running Windows 2000 or Windows Server 2003

A user's roaming profile contains personal information that is copied to and from the client computer and the server that hosts the roaming profile; therefore, it is important to ensure that the data is protected as it travels over the network.

The major potential threats to the privacy and integrity of a user's data come from malicious users intercepting and tampering with data as it passes over the network, or the server hosting the user's data.

Several features of Windows 2000 and Windows Server 2003 can help to secure a user's data:

- **Kerberos.** Standard on all versions of Windows 2000–based servers, Kerberos ensures the highest level of security to network resources. While NTLM authenticates the client only, Kerberos authenticates the server and the client. When NTLM is used, the client does not detect whether the server is valid. This is particularly important if the client exchanges personal files with the server, as is the case with roaming profiles. Kerberos provides better security than NTLM and is not available on Windows NT 4.0 or earlier operating systems.

- **IP Security Protocol (IPSec).** IPSec provides network-level authentication, data integrity, and encryption to ensure that roamed data is safe from the following:

  - Data modification while en route
  - Interception, viewing, or copying
  - Access by unauthenticated parties

  For more information about IPSec, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

- **Server Message Block Signing.** The SMB authentication protocol supports message authentication. This prevents active message and "man-in-the-middle" attacks. SMB signing provides this authentication by placing a digital signature into each SMB, which is then verified by both the client and the server. To use SMB signing, you must either enable it or require it on both the SMB Service client and the SMB Service server.

**Note**

- SMB signing imposes a performance penalty even though it does not consume any more network bandwidth; it does use more CPU cycles on the client and server.

#### Using the NTFS File System for Volumes Containing User Data

For the most secure configuration, always configure servers that host roaming profiles to use NTFS. Unlike a file allocation table (FAT), NTFS supports discretionary access control lists (DACLs) and system access control lists (SACLs) which determine who can perform operations on a file and what events trigger logging of actions performed on a file.

# Configuring Roaming User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you create a roaming user profile, you need to create each user account. Then, log on to a server as an administrator to create a network share to store the roaming user profiles, designate the groups of users to receive the roaming user profiles, and grant all users Full Control permissions.

Use the following procedures when you create and manage roaming user profiles.

For information about deploying Roaming User Profiles on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

**Creating Roaming User Profiles**

To perform the following procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. For enhanced security, consider using the **Runas** command to perform this procedure.

**To create a roaming user profile**

1. Open **Active Directory Users and Computers**.
2. Click the domain and the OU where the user account resides.
3. Right-click the user account for which to set a roaming profile, and then click **Properties**.
4. Click the **Profile** tab, and then type the profile path information in **Profile path**. (Use the full path in each user account. For example, type ***\\Server\ShareName\UserName***.)

Another way to populate the profile path is to use an Active Directory® Service Interfaces (ADSI) script. ADSI provides a single set of interfaces for managing resources on the network. You can use ADSI in combination with Microsoft® Visual Basic® Scripting Edition (VBScript) or JScript scripts to manage Active Directory resources such as users and services.

For information about ADSI and ADSI scripts, see the Microsoft Platform SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Changing User Profile Type from Local to Roaming**

Typically, a large organization has many users with local profiles. For ease of management, you might want to change many of the local profiles to roaming profiles. Moving user's data and settings from the workstation to a server reduces the user's dependence on the workstation's availability, simplifies user data management, and allows centralized account management.

**To create a roaming user profile for a user that has a local profile**

1. Open **Active Directory Users and Computers**.
2. Click the domain and the OU where the user account resides.
3. Right-click the appropriate user account for which to set a roaming profile, and then click **Properties.**
4. Click the **Profile** tab, and type the profile path information in **Profile path** (for example, type **\\Server\ShareName\UserName**).

### Note

- To change a user's local profile to a roaming profile for a user who uses multiple computers simultaneously, the user must log off last from the computer that has the profile that the user wants to use.

**Disabling Roaming User Profiles on Certain Computers**

You can prevent computers from receiving roaming profiles by enabling the **Only allow local user profiles** policy setting, which blocks roaming profiles from being used on a computer. By default, when roaming profile users log on to a computer, the user's roaming profile is copied to the local computer. If the user has previously logged on to this computer, the roaming profile is merged with the local profile. Similarly, when the user logs off from this computer, the local copy of the profile, including any changes the user made, is merged with the server copy of the profile.

If you enable the **Only allow local user profiles** policy setting, the following occurs on the affected computer: When the user first logs on, the user receives a new local profile instead of the roaming profile. At logoff, changes are saved to the local profile. All subsequent logons use the local profile.

If you enable both the **Prevent Roaming Profile changes from propagating to the server** setting and the **Only allow local user profiles** setting, roaming profiles are disabled for that computer. These policy settings are in the **Computer Configuration\Administrative Templates\System\User Profiles** node.

#### Creating Accounts That Possess roaming user profiles

You can save time and reduce the chances for error by scripting many repetitive tasks, such as creating user accounts. A script to automate the creation of user profiles for roaming user might look something like the sample script Listing 7.1, which shows a script for creating user accounts that have roaming profiles.

**Listing 7.1 Creating User Accounts That Have Roaming User Profiles**

```
set Args = Wscript.ArgumentsouName = Args(0)
usrName = Args(1)
RUProot = Args(2)

RUPpath = RUProot & " \"  & usrName

'Get the domain
Set dse = GetObject(" LDAP://RootDSE" )
Set domain = GetObject( " LDAP://"  & dse.Get(" defaultNamingContext" ))

set ou = domain.GetObject(" organizationalUnit" , " OU="  & ouName )

wscript.echo " Creating user in "  & ou.Name

set usr = ou.Create(" user" , " cn="  & usrName )
usr.Put " samAccountName" , usrName
usr.Put " userPrincipalName" , usrName
usr.Put " Profilepath" , RUPpath

usr.SetInfo

wscript.echo "  User "  & usrName & "  was created successfully in "  & ou.Name & "
with a RUP Path of: "  & RUPpath

```

Every Windows Server 2003 user has a profile. If the operating system does not have a profile to apply to the user when the user logs on, a new local profile is created for the user, based on the defaults in place. Windows Server 2003 applies a generic user profile format by default.

#### Configuring a Default Profile

You can create a default profile to ensure that all users within a domain receive an identical profile the first time they log on. This option simplifies administrative control over the users' desktops and settings.

To create a default user profile, you must be logged on as Administrator or a member of the Administrators group. Create a default profile for all new user accounts in a domain. Include any domain-specific customizations that you want in the profile. To create subsequent profiles, you can create a new user account as a template.

Before creating a new user account to use as a new user's profile template, perform the following tasks:

1. Log on to the domain as the new user, and then customize the desktop if appropriate.

2. Optionally, install and configure any applications to be shared by user accounts made from this template.

3. Log off, and then log on as the administrator.

For more information about creating a new user account, see "Create a new user account" in Help and Support Center for Windows Server 2003.

**To configure a new user account to use as a new user's profile template**

1. After you create a new user account template, in **Control Panel**, click **System**.
2. On the **Advanced** tab, under **User Profiles**, click **Settings**.
3. Under **Profiles stored on this computer**, select the user that you created in step 1, and then click **Copy To**.
4. To create the default user profile for the domain, type the path to NETLOGON\Default User on the domain controller.
5. In the **Copy To** dialog box, under **Permitted to use**, click **Change**.
6. In the **Select User or Group** dialog box, enter the object name to select, and then type: **Everyone**.

**Troubleshooting: Creating a Log File for User Profiles**

User profiles log events in the Application event log. To aid in troubleshooting, administrators can also create detailed log files by using the following procedure.

### Caution

- Do not edit the registry unless you have no alternative. The registry editor, regedit.exe, bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you edit the registry, make sure to back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

**To create a detailed log file for user profiles**

1. In the **Run** dialog box, type **regedit,** and then click **OK**.
2. Locate the following subkey: HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Winlogon.
3. Create a new entry named **UserEnvDebugLevel** of data type **REG_DWORD**, and set its value to 0x30002.

The log file is stored in this location: %windir%\Debug\Usermode\Userenv.log.

# Implementing Folder Redirection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Folder Redirection is simple to configure and manage. While Folder Redirection reduces the disk space that is required on the client, it requires network shares to accommodate the redirected folders. By combining Folder Redirection and roaming user profiles, you can decrease logon and logoff times for roaming and mobile users. Besides the improved availability and backup benefits of having the data on the network, users also have performance gains with low-speed network connections and subsequent logon sessions.

By using Folder Redirection in conjunction with roaming user profiles, you can also provide fast computer replacement. If a user's computer needs to be replaced, the data that a user requires can quickly be re-established on a replacement computer.

For information about deploying Folder Redirection on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

# Recommendations for Folder Redirection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To get the best performance from Folder Redirection, it is recommended that you: create the root share on the server and let the system create the users' folders, synchronize files at logoff when you use Folder Redirection with Offline Files, and use the following guidelines for redirecting **My Documents**, and setting options for Offline Files.

For information about deploying Folder Redirection on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

#### Letting the system create folders for each user

For optimal performance of the Folder Redirection feature, it is strongly recommended that you create *only* the root share on the server, and then let the system create the folders for each user. If you must create the folders for users, ensure that you correctly assign permissions. For more information about assigning permissions see "Security Recommendations for Folder Redirection" later in this chapter.

**Important**

- If you must create folders for users, make sure that you set the correct permissions. Then, clear the **Grant exclusive rights to** check box on the **Settings** tab of the **Folder RedirectionProperties** page. If you do not clear this check box, Folder Redirection first checks preexisting folders to determine if the user is the owner. If the administrator previously created the folder, the check fails, and redirection is cancelled. Folder Redirection logs an event in the Application event log indicating that redirection failed and that the new directories for the redirected folder cannot be created due to not being able to assign a security ID as the owner of the folder (Event ID 101).

#### Accepting the default settings for Folder Redirection

If you are storing roaming user profiles on the server where you have enabled Offline Files, Folder Redirection ensures that Offline Files are set to synchronize when the user logs on and logs off.

For Windows 2000 and earlier clients, if you are using Offline Files in conjunction with Folder Redirection and roaming user profiles, it is recommended that you leave the default setting of synchronizing Offline Files at logoff enabled to ensure best performance of these features. This is because if a shared folder is unavailable, Offline Files considers the entire server to be unavailable until the offline cache is manually synchronized. Roaming profiles are *not* synchronized with the server while Offline Files treats the server as being unavailable. For clients that run Windows XP or later, whether the default synchronization setting is selected or not, roaming profiles continue to roam even if Offline Files has marked the server as being offline.

#### Using offline files settings

Using offline files settings on a server share where the user data is stored is especially useful for users of portable computers. It is recommended that you use Folder Redirection in conjunction with Offline Files. Table 7.10 lists the recommendations for Offline Files.

**Table 7.10 Recommendations for Configuration of Offline Files**

| Redirected Folder | Recommended Offline Files Settings |
|---|---|
| My Documents | **All files and programs that users open from the share will be automatically available offline** *or* **Only the files and programs that users specify will be available offline** (if you want users to manually designate files and folders to be available offline). |
| My Pictures | **All files and programs that users open from the shared folder will be automatically available offline** *or* **Only the files and programs that users specify will be available offline.** |
| Application Data | **All files and programs that users open from the share will be automatically available offline.** |
| Desktop | **All files and programs that users open from the share will be automatically available offline** if the desktop is read only. |

| | |
|---|---|
| Start Menu | **All files and programs that users open from the share will be automatically available offline**. |

**Redirecting My Documents**

The following suggestions for redirecting the My Documents folder are appropriate for most deployments and can provide a faster and simpler deployment.

- Redirect My Documents to a network share.
- Allow Folder Redirection to create folders for you. When setting redirection policy for a group, use the path to the share, such as *server**share*. Folder Redirection then appends the user name and the folder name when the policy is applied.
- Allow Folder Redirection to perform all the moving of folders and files when you select a folder for redirection or change the target network share to which you redirect the folder. The Folder Redirection client not only moves files to the appropriate network share, but it also sets proper folder security and renames entries in the Offline Files cache database so that they continue to link to the correct target folders and files. Any files pinned by the user in the Offline File Cache stay pinned.
- Combine Folder Redirection with Offline Files to provide the user access to My Documents, even when the user's workstation is temporarily disconnected from the network. This is particularly useful for people who use portable computers. For more information, see "Make a file or folder available offline" in Help and Support Center for Windows Server 2003.
- Include redirected folders, particularly **My Documents**, in routine server backups. Performing backups of user data that is located on network shares is simpler and more reliable because it requires no action on the part of the user or interaction with the workstation.
- Use Group Policy to set profile quotas and disk quotas to establish limits on the disk space that is used by users' data and settings.
- Redirect user-specific data from the hard disk that holds the operating system files. This data can be redirected to a different hard disk on the user's local computer or to a network share. This simplifies system maintenance by separating system files from user files.
- Centralize storage on large shares to reduce workstation hardware and maintenance costs. Pooling disk space more than offsets the cost of increasing server disk capacity.
- Do not use the **Redirect to home folder** policy setting unless you have already deployed home folders in your organization.
- Leave the **My Pictures** folder located in the **My Documents** folder.

When you redirect the **My Documents**, the Recycle Bin size for **My Documents** defaults to a percentage of the size of the server partition where the redirected **My Documents** resides. You can manually change this size in 1 percent increments. Because a Recycle Bin can grow large, encourage users to empty their Recycle Bins periodically.

**Note**

- You can grant users exclusive rights to their redirected folders. Select the **Grant the user exclusive rights to My Documents** check box under the **Settings** tab in each folder's **Properties** dialog box to grant full control over the folder to the user and the local system only.

**Redirect folders to home folders**

Typically, it is recommended that you do *not* redirect to a home directory unless you have already deployed home directories in your organization. However, if you have home directories and want to transition your users to use **My Documents** while maintaining compatibility with the home directory environment, you can redirect a user's **My Documents** folder to the user's home folder. The **Redirect to home folder** policy setting is intended *only* for organizations in which home folders are already in place. Redirect only the **My Documents** folder to the home folder. For this type of redirection, the client computer must run

one of the following operating systems: Windows XP Professional, Microsoft® Windows XP, 64-Bit Edition, or Windows Server 2003. This redirection option does not work for clients that run Windows NT, Windows 2000, or Microsoft® Windows® XP Home Edition.

When a folder is redirected to the home folder, security and ownership are not checked, and permissions are not changed. Folder Redirection behaves as though the administrator has set directory security correctly. This relaxed security is the reason that redirection to the home folder is not recommended if the home folder structure is not already in place, and you have not updated your configuration.

Typically, folder redirection fails if a user is not the owner of the folder to which the **My Documents** folder is redirected. Because redirection to the home folder is intended for an earlier environment, Folder Redirection does not check for proper folder ownership. Instead, ownership check is left to the administrator.

Users must have the home folder property set correctly on their user object in Active Directory. The client computer gets the path for the user's home folder from the user object in Active Directory when the user logs on. User accounts that have redirected folders must have this path set correctly, or Folder Redirection fails.

For more information about creating home folders for profiles, see "Add a home folder to a profile" in Help and Support Center for Windows Server 2003.

**Combining Folder Redirection with Offline Files**
The Offline Files technology applies to network shares or mapped drives that contain documents or data that a user might want to use offline. Folder Redirection and Offline Files are functionally independent but complementary in operation.

Folder Redirection provides user access to redirected folders that have been relocated to network shares. Offline Files provides the user with access to any specified folder or file when the network share is unavailable, or offline.

**Redirecting the Start Menu folder**
Folder Redirection of the **Start** Menu folder is available in Windows XP, Windows XP 64-Bit Edition, or Windows Server 2003 operating systems. **Start** Menu redirection is treated differently from other redirected folders in that the contents of the user's **Start** Menu are not copied to the redirected location. It is assumed that a redirected **Start** Menu has been previously created by an administrator and that all users share the same **Start** Menu. As a best practice for Windows XP–based computers, do *not* use Folder Redirection to redirect the **Start** Menu folder; instead, use Group Policy to control what appears on the **Start** Menu.

**Redirecting application data and use of Outlook**
If a user redirects the **Application Data** folder and runs multiple instances of the Microsoft Outlook messaging and collaboration client on different computers, including an instance of Microsoft® Outlook 2000, mail opening performance is delayed. Outlook 2000 continually keeps the Outcmd.dat file open. (This file stores information about toolbar customizations that users make in Outlook.) When another instance of Outlook tries to access Outcmd.dat, it is unable to access it because Outlook 2000 has locked the file. The second copy of Outlook repeatedly tries to access Outcmd.dat, causing a delay when the user tries to open or reply to messages.

**Note**

- Outlook *2002*, included in Microsoft® Office XP, does *not* hold the Outcmd.dat file open. This behavior occurs *only* when Outlook 2000 is running on one of the computers.

**Encrypted file system considerations**
Folder Redirection has implications for encrypted files that are located in redirected folders.

- Files redirected to a server can be encrypted by Encrypting File System (EFS) *only* if an administrator has designated the remote server as *trusted for delegation*. Administrators can establish a service or computer as trusted for delegation to allow that service or computer to complete delegated authentication, receive a ticket for the user who makes the request, and then access information for that user.

- Encrypted files are decrypted before being transmitted over the network. File encryption only protects the files while they reside on the disk.

- The Offline Files cache cannot be encrypted on Windows 2000 Professional.

For more information about delegating trust and enabling delegated authentication, see "Designing an Authentication Strategy" in *Designing and Deploying Directory and Security Services* of this kit.

# Group Policy Recommendations for Folder Redirection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following Group Policy guidelinesto ensure that offline files are fully synchronized and available when users work offline, and to specify the behavior of redirected folders and their contents when a Group Policy object no longer applies.

**Enabling the Synchronize all offline files before logging off Policy Setting**
Always enable the **Synchronize all offline files before logging off** policy setting to ensure that offline files are fully synchronized and that all the files in the users' redirected folder are available when users work offline. If this setting is not enabled, the operating system only performs a quick synchronization, and only recently used files are cached. The **Synchronize all offline files before logging off** policy setting is available in the **Computer Configuration\Administrative Templates\Network\Offline Files** node of the Group Policy Object Editor.

**Removing Group Policy**
It is important to consider the behavior of your Folder Redirection policy settings when Group Policy is removed. To specify Group Policy removal options, use the selected folder's **Properties** page and select the **Settings** tab. You can access this page under **User Configuration\Windows Settings\Folder Redirection** of the Group Policy Object Editor by right-clicking a folder (for example, **Application Data**), clicking **Properties,** and then selecting the **Settings** tab to specify the policy removal options to use.

Table 7.11 indicates what happens to redirected folders and their contents when a Group Policy object no longer applies.

**Table 7.11 Summary of Folder Properties Page Settings for Redirected Folders**

| Setting: Move the contents of special folder to the new location option | Policy Removal Option | Result of Removal |
|---|---|---|
| Enabled | **Redirect the folder back to the user profile location when policy is removed** | • The special folder returns to its user profile location.<br>• The folder contents are copied back to the user profile location.<br>• The contents are not deleted from the redirected location.<br>• The user can continue to access the contents, but only on the local computer. |
| Disabled | **Redirect the folder back to the user profile location when policy is removed** | • The special folder returns to its user profile location.<br>**Note:** In this case, the folder contents are not copied or moved to the user profile location. As a result, the user can no longer see the contents. |
| Either Enabled or Disabled | **Leave the folder in the new location when policy is removed** | • The special folder remains at its redirected location.<br>• The contents remain at the redirected location.<br>• The user continues to have access to the contents at the redirected folder. |

**Important**

- Changing the redirection option to **No Administrative policy specified***does not* redirect the folder to the local profile,

this option means that Folder Redirection is not configured — *if a folder was previously redirected it continues to be redirected to the previous location*. If you want to return the folder to the local user profile, use the **Redirect the folder back to the user profile location when policy is removed** option.

# Security Recommendations for Folder Redirection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following guidelines when you create the shares for redirected folders to ensure you set access permissions appropriately, and to help provide the most secure configuration.

For information about deploying Folder Redirection on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

**Restricting access to the share**

Redirected folders contain personal information such as documents and EFS certificates so it is important to protect this data.

- Create a security group for users who have redirected folders on a particular share and limit access only to those users
- Create a hidden share by putting a dollar sign ($) after the share name. The share is not visible in the network neighborhood.
- Grant users the minimum permissions that are required to access the data.

**Assigning permissions for root folder, shares, and user's redirected folder**

Tables 7.12, 7.13, and 7.14 show the permissions for the folder redirection root, share, and the users' redirected folders.

**Table 7.12 NTFS Permissions for Folder Redirection Root Folder**

| User Account | Minimum Permissions Required |
|---|---|
| Creator Owner | Full Control, Subfolders and Files Only |
| Administrator | None |
| Security group of users that need to put data on share | List Folder/Read Data, Create Folders/Append Data - This Folder Only |
| Everyone | No Permissions |
| Local System | Full Control, This Folder, Subfolders and Files |

**Table 7.13 Share level (SMB) Permissions for Folder Redirection Share**

| User Account | Default Permissions | Minimum permissions required |
|---|---|---|
| Everyone | Full Control | No permissions |
| Security group of users that need to put data on share. | N A | Full Control |

**Table 7.14 NTFS Permissions for Users' Redirected Folders**

| User Account | Default Permissions | Minimum permissions required |
|---|---|---|
| %Username% | Full Control, Owner of Folder | Full Control, Owner of Folder |
| Local System | Full Control | Full Control |
| Administrators | No permissions | No permissions |
| Everyone | No permissions | No permissions |

**Host redirected file shares on servers running Windows 2000 or Windows Server 2003**

To provide the best protection as data is transmitted over the network, ensure that you set up the redirected folders shares on servers running Windows 2000 and later. The Kerberos, IPSec, and SMB signing security features of Windows 2000 and

Windows Server 2003 help protect the users' data.

**Using the NTFS file system for user data volumes**
Always configure the servers hosting redirected files to use NTFS to provide the most secure configuration.

**Do not rely on EFS to encrypt users' files when transmitted over the network**
When you use EFS to encrypt files on a remote server, the data is encrypted *only* while it is stored on the disk, not when it is transmitted over the network. The exceptions to this are when your system includes IPSec or Web Distributed Authoring and Versioning (WebDAV). IPSec encrypts data while it is transported over a TCP/IP network. If the file is encrypted before being copied or moved to a WebDAV folder on a server, it remains encrypted during the transmission and while it is stored on the server.

**Encrypting the Offline Files cache**
While access control lists (ACLs) protect the Offline Files cache on NTFS partitions by default, encrypting the cache enhances security on a local computer. By default, the cache on the local computer is *not* encrypted, so any encrypted files that are cached from the network are not encrypted on the local computer. This might pose a security risk in some environments.

When you enable encryption, all files in the Offline Files cache are encrypted, including existing files and any files that you add later. The cached copy on the local computer is affected, but the associated network copy is not.

You can encrypt the cache in one of two ways:

- By using Group Policy to enable the **Encrypt the offline files cache** policy setting. This setting is in the **Computer Configuration\Administrative Templates\Network\Offline Files** node in the Group Policy Object Editor snap-in.

- Manually, by clicking **Folder Options** on the **Tools** menu in **Windows Explorer**. Click the **Offline Files** tab, and then select the **Encrypt offline files to secure data** check box.

### Note

- Encryption of the Offline File cache is only available in Windows XP and Windows Server 2003; it is *not* possible to encrypt the cache on Windows 2000–based computers.

For information about encrypting the Offline Files cache for Windows XP, see the How to Encrypt Offline Files link on the Web resources page at http://www.microsoft.com/windows/reskits/webresources. For information about encrypting files for Windows 2000, see the Encrypting File System for Windows 2000 link on the Web resources page.

# Configuring Folder Redirection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the Group Policy Object Editor snap-in to redirect special folders to network locations.

For information about deploying Folder Redirection on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

The primary tasks that pertain to configuring Folder Redirection include:

- Creating or selecting a GPO for Folder Redirection.
- Selecting **Basic** or **Advanced** Folder Redirection options. **Advanced** redirection allows you to apply the redirection to users that belong in a specified security group.
- Identifying the target folder location. You can redirect folders to any of the following locations:
  - The root of a network file share
  - A folder on a network share
  - A local folder
  - A home directory (only applicable if you have existing home directories in your environment)
- Setting permissions for root folders and users' folders.

**Note**

- If you allow Folder Redirection to create the redirected folders on a specified network, the folders that are created in this way have proper permissions assigned to them. If you create the folders manually, you must make sure that permissions are properly assigned.

**Redirecting Special Folders**

In most cases, you should redirect special folders to the root directory instead of to a specific location because this appends the correct paths for you, including %username%.

**To redirect special folders to the root directory**

1. Open a GPO that is linked to the site, domain, or OU that contains the users whose special folders you want to redirect.
2. In the console tree, under User **Configuration\Windows Settings\Folder Redirection**, double-click **Folder Redirection** to display the special folder that you want to redirect.
3. Right-click the special folder that you want to redirect (such as **Desktop** or **My Documents**), and then click **Properties**.
4. Click the **Target** tab, and then in the **Settings** box, select **Basic - Redirect everyone's folder to the same location**.
5. Under **Target folder location**, select **Create a folder for each user under the root path**.
6. In the **Root Path** box, type a Universal Naming Convention (UNC) path, such as ***\\servername\sharename***, and then click **OK**.
7. In the **Properties** dialog box for the special folder, click **OK**.

The user name and folder name are appended to the UNC path automatically.

**To redirect special folders to a specific path**

1. Open a GPO that is linked to the site, domain, or OU that contains the users whose special folders you want to redirect.

2. In the console tree, under User Configuration\Windows Settings\Folder Redirection, double-click Folder Redirection to display the special folder that you want to redirect.

3. Right-click the special folder that you want to redirect, and then click **Properties**.

4. In the **Setting** box on the **Target** tab, select **Basic - Redirect everyone's folder to the same location**.

5. Under **Target folder location**, select **Redirect to the following location**. In the **Root Path** box, type an UNC path, such as \\server\share, or type a locally valid path, such as C:\somefolder, and then click **OK**.

If necessary, you can return redirected folders to their original default locations by using the following procedure.

**To redirect special folders to the local profile location**

1. Open a GPO that is linked to the site, domain, or OU that contains the users whose special folders you want to redirect.

2. In the console tree, double-click Folder Redirection to display the special folder that you want to redirect.

3. Right-click the special folder that you want to redirect (such as Desktop or My Documents), and then click **Properties**.

4. In the **Setting** box on the **Target** tab, select **Basic - Redirect everyone's folder to the same location**.

5. Under **Target folder location**, select **Redirect to the local user profile location**, and then click **OK**.

**Redirecting to the home directory**

It is recommended that you do *not* redirect to a home directory unless you have already deployed home directories in your organization. However, if you have already deployed home directories and you want to maintain compatibility with your existing home directory environment, you can redirect My Documents to the home directory.

**To redirect My Documents to the home directory**

1. Open a GPO that is linked to the site, domain, or OU that contains the users whose My Documents folders you want to redirect.

2. In the console tree, double-click Folder Redirection to display My Documents.

3. Right-click My Documents, and then click **Properties**.

4. In the **Setting** box, on the **Target** tab, select **Basic - Redirect everyone's folder to the same location**.

5. Under **Target folder location**, select **Redirect to the user's home directory**, and then click **OK**.

### Note

- Domain administrators have full control over the users' My Documents folder when you redirect My Documents to the home directory. This is the case even if you select the **Grant the user exclusive rights to My Documents** option on the **Settings** tab in the **My Documents Properties** dialog box.

**Redirecting special folders based on membership in security groups**

You can redirect special folders for use by members of security groups to the root directory, a specific path, the local profile directory, or to the home directory.

**To redirect special folders to the root directory based on security group membership**

1. Open the GPO that contains the users whose special folders you want to redirect.

2. In the console tree, double-click Folder Redirection to display the special folder that you want to redirect (for example, Desktop or My Documents), right-click the special folder, and then click **Properties**.

3. On the **Target** tab, in the **Setting** box, select **Advanced - Specify locations for various user groups**, and then click **Add**.

4. In the **Specify Group and Location** dialog box, under **Security Group Membership**, type the name of the security group, or click **Browse** to search for it.

5. Under **Target Folder Location**, select **Create a folder for each user under the root path**.

6. In **Root Path**, type the UNC path to use for redirection (for example, \\servername\sharename), and then click **OK**.

7. In the special folder's Properties dialog box, click OK.

For more information about Folder Redirection, in Help and Support Center for Windows Server 2003, on the toolbar, click the **Index** button, in **Type in the keyword to find**, type **folder redirection**, and then double-click the topic that you want in the index list.

**Troubleshooting: Creating a Log File for Folder Redirection**
Folder Redirection logs events in the Application Event Viewer log. To aid troubleshooting, administrators can also create detailed log files for Folder Redirection.

### Caution

- Do not edit the registry unless you have no alternative. The registry editor, regedit.exe, bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you edit the registry, make sure to back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

**To enable logging for Folder Redirection**

1. In the **Run** dialog box, type **regedit**, and then click **OK**.

2. Locate the following subkey: HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Diagnostics.

3. Create a new entry called FdeployDebugLevel of data type Reg_DWORD, and set its value to 0x0f.

The log file is created in %windir%\Debug\Usermode\Fdeploy.log.

# Implementing Offline Files and Synchronization

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By deploying the Offline Files technology, you allow network users to access files that are stored on network shares, even when the client computer is not connected to the network. A mobile user that accesses the share while disconnected from the network can still browse, read, and edit files because the files are cached on the client computer. You can initiate synchronization manually or set Synchronization Manager to control the synchronization of offline files with the network.

For information about deploying Offline Files on newer versions of Windows, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles.

The Offline Files technology provides either a *full* or *quick* synchronization. A full synchronization ensures that the most current version of every network file that has been made available offline is available. A quick synchronization is much faster than a full synchronization, and it ensures that a complete version of each cached file is available. The quick synchronization does not ensure that each cached file is the most current version of every network file that has been made available offline. The quick synchronization ensures that each file in the offline cache is a complete. For more information about Offline Files, file synchronization, and managing user data and settings, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Recommendations for Using Offline Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using Offline Files, you can manage user data more efficiently while enhancing users' access to their files and folders. When combined with Folder Redirection, Offline Files improves availability; when combined with Synchronization, Offline Files increases the reliability of content.

**Files and Folders You Can Make Available Offline**

You can make available for offline use any shared files or folders on a Microsoft network. You can make files available offline from any computer that supports Server Message Block-based file and printer sharing, including Microsoft® Windows® 95, Microsoft® Windows® 98, Windows NT 4.0, and the Windows Server 2003 family.

You can set up shared folders so that every network file the user opens in that shared folder is available offline. Alternatively, you can specify which network files are made available offline. Table 7.15 shows typical Offline Files caching settings for each redirected folder.

**Note**

- If you make a shortcut to a file available offline, that file is made available offline. However, if you make a shortcut to a *folder* available offline, the contents of that folder are not available offline.

**Table 7.15 Offline Files Settings**

| Special Folder | Offline File Settings |
|---|---|
| My Documents | **All files and programs that users open from the share will be automatically available offline** *or* **Only the files and programs that users specify will be available offline** to allow users to designate files and folders for offline use. |
| My Pictures | **All files and programs that users open from the share will be automatically available offline** *or* **Only the files and programs that users specify will be available offline** so users can designate files and folders for offline use. |
| Application data | **All files and programs that users open from the share will be automatically available offline**. |
| Desktop | **All files and programs that users open from the share will be automatically available offline** if the desktop is Read Only. |
| Start Menu | **All files and programs that users open from the share will be automatically available offline**. |

**Offline Files Caching Options**

In the Windows Server 2003 family, Offline Files is, by default, not enabled. When enabled, the following three types of Offline Files caching are available:

- **Only the files that users specify will be available offline** (previously called **Manual Caching for Documents**). Provides offline access to only those files on a network share that have been manually selected. Manual caching for documents is the default option when you set up a shared folder to use offline.

- **All files and programs that users open from the share will be automatically available offline** (previously called **Automatic Caching for Documents** and **Automatic Caching for Programs**). Selecting this option makes every file that someone opens from your shared folder available offline. However, only those files that have been opened are available offline.

- **Files or programs from the share will not be available offline**. Selecting this option prevents users from storing files offline.

In Windows XP and Windows Server 2003, the default behavior is to cache all redirected folders automatically. Even if you specify manual caching for a redirected folder, the folder is still subject to automatic caching. To disable this default behavior, enable the **Do not automatically make redirected folders available offline** policy setting, available in **User Configuration\Administrative Templates\Network\Offline Files**. If you enable this policy setting, the users must manually select the files to be made available offline.

#### Combining Offline Files with EFS

When you use Offline Files, you can specify that offline files are encrypted. If you choose to encrypt offline files, only the files on your local computer are encrypted. The files on the network are not encrypted, nor are the files encrypted as they traverse the network.

The **Encrypt offline files to secure data** setting is disabled under the following conditions:

- You are not an administrator on the computer.
- The local drive is not NTFS or does not support encryption.
- A system administrator has implemented an encryption policy for Offline Caching.
- The Offline Files technology is not available in Windows XP Home Edition.
- In the Windows Server 2003 family, Offline Caching is not enabled by default.

#### Deploying Offline Files on Terminal Services Clients

Multi-user Terminal Services disables Offline Files on computers that run Windows Server 2003. Terminal Services on Windows XP Professional is always single user, but Terminal Services on servers is multi-user by default. To use Terminal Services and Offline Files together on servers, set the value of the **AllowMultipleTSSessions** registry entry to 0. This entry is found in the following registry subkey**:**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

### Caution

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

#### Offline Files and Synchronization Manager

The copies of offline files that are stored in the Offline Files data store are kept up to date by Synchronization Manager. The Synchronization Manager tool is invoked when users with roaming user profiles log on, log off, or access selected files or folders.

#### Caching Offline Files

Offline Files notifies users when the status of the network connection changes. A message notifies users when they are offline, and the users can continue working with the files. When they work offline, users can still browse network drives and shared folders that are set up for Offline Files. A red X appears over any disconnected network drives that are not currently available. For more information about their network connection status, users can click the **Offline Files** icon in the notification area.

When the network connection is restored, any changes the user made while working offline are updated to the network. If other users have made changes to the same file, options are offered to the users for resolving the version conflict. For more information, see "Handling file conflicts" in Help and Support Center for Windows Server 2003.

#### Forcing Synchronization

A user who is concerned about losing network connectivity can force synchronization to ensure that local copies of network files are current.

You can use Synchronization Manager to perform periodic synchronizations to ensure that users' offline files are kept current.

### Note

- To keep offline files current with minimal interruption to users, schedule offline files to be synchronized whenever users' computers are idle.

**Synchronizing Offline Files**

If you configure Offline Files, you need to decide the appropriate time to synchronize offline files. You also have a choice of two methods, full or quick. For a *full synchronization*, enable the **Synchronize all offline files before logging off** policy setting. Disabling this policy causes a*quick synchronization* to be performed when the user logs off.

A full synchronization ensures that you have the most current versions of every network file that has been made available offline. A quick synchronization ensures that you have complete versions of all your offline files, even if they are not the most current versions.

In Windows Server 2003, the **Synchronize all offline files when logging on** policy setting mirrors the **Synchronize all offline files before logging off** policy setting. Both policy settings are available in both User Configuration and Computer Configuration under **Administrative Templates\Network\Offline Files**. For more information, see the **Explain** tab on the **Properties** page of each policy.

By default, the Offline Files feature is not enabled in the Windows Server 2003 family because any Terminal Services mode that allows multiple, concurrent user logon sessions is not compatible with Offline Files. Offline Files is enabled by default in Windows XP Professional and Windows 2000 Professional.

To maximize the benefits of Offline Files, the Synchronization Manager service frequently synchronizes server and client copies of any offline file. However, synchronization makes demands on file system and network resources.

# Configuring Group Policy for Offline Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Group Policy for Offline Files**

Tables 7.16 and 7.17 list the Offline Files policy settings that administrators can use. The settings are available in **User Configuration\Administrative Templates\Network\Offline Files** and in **Computer Configuration\Administrative Templates\Network\Offline Files** in the Group Policy Object Editor snap-in. For more information about policy settings, see the **Explain** tab on the **Properties** page of each policy.

**Note**

- Most of the policy settings exist in both the User Configuration and the Computer Configuration nodes.

**Table 7.16 User Configuration Policy Settings for Offline Files**

| Policy Setting | Description |
|---|---|
| **Prohibit user configuration of Offline Files** [1] | Prevents users from enabling, disabling, or changing the configuration of Offline Files.<br><br>This setting removes the **Offline Files** tab from the **Folder Options** dialog box. It also removes the **Settings** item from the **Offline Files** shortcut menu and disables the **Settings** button on the **Offline FilesStatus** dialog box. As a result, users cannot view or change the options on the **Offline Files** tab or in the **Offline Files** dialog box.<br><br>This is a comprehensive policy setting that locks down the configuration that you establish by using other settings in this folder. |
| **Synchronize all Offline Files when logging on** [1] | Determines whether offline files are fully synchronized when users log on.<br><br>This setting also disables the **Synchronize all offline files before logging on** option on the **Offline Files** tab. This prevents users from changing the option while a setting controls it. |
| **Synchronize all Offline Files before logging off** [1] | Determines whether offline files are fully synchronized when users log off.<br><br>This setting also disables the **Synchronize all offline files before logging off** option on the **Offline Files** tab. This prevents users from trying to change the option while a setting controls it. |
| **Synchronize Offline Files before suspend** | Determines whether offline files are synchronized before a computer is suspended.<br><br>If you enable this setting, offline files are synchronized whenever the computer is suspended. Setting the synchronization action to **Quick** ensures only that all files in the cache are complete. Setting the synchronization action to **Full** ensures that all cached files and folders are up to date with the most current version.<br><br>If the computer is suspended by closing the display on a portable computer, synchronization is not performed. If multiple users are logged on to the computer at the time the computer is suspended, synchronization is not performed. |
| **Action on server disconnect** [1] | Determines whether network files remain available if the computer is suddenly disconnected from the server that is hosting the files.<br><br>This setting also disables the **When a network connection is lost** option on the **Offline Files** tab, which prevents users from trying to change the option while a setting controls it.<br><br>Also, see the **Non-default server disconnect actions** policy setting in this table. |

| | |
|---|---|
| **Non-default server disconnect actions** [2] | Determines how computers respond when they are disconnected from particular offline file servers.<br><br>This setting overrides the default response, a user-specified response, and the response that is specified in the **Action on server disconnect** policy setting. This setting does not prevent users from setting custom actions on the Offline Files tab. However, users cannot change any custom actions that were established by means of this setting. |
| **Remove 'Make Available Offline'** [1] | Prevents users from making network files and folders available offline.<br><br>This setting removes the **Make Available Offline** option from the **File** menu and from all shortcut menus in Windows Explorer. As a result, users cannot designate files to be saved on their computer for offline use. However, this setting does not prevent the system from saving local copies of files that reside on network shares that are designated for automatic caching. |
| **Prevent use of Offline Files folder** [1] | Disables the Offline Files folder.<br><br>This setting disables the **View Files** button on the **Offline Files** tab. As a result, users cannot use the Offline Files folder to view or open copies of network files that are stored on their computer. Also, they cannot use the folder to view characteristics of offline files, such as their server status, type, or location.<br><br>This setting does not prevent users from working offline or from saving local copies of files available offline. Also, it does not prevent users from using other programs, such as Windows Explorer, to view their offline files. |
| **Administratively assigned offline files** [3] | Lists network files and folders that are always available for offline use.<br><br>This setting makes the specified files and folders available offline to users of the computer. |
| **Turn off reminder balloons** [1] | Hides or displays reminder balloons, and prevents users from changing the setting.<br><br>Reminder balloons appear above the **Offline Files** icon in the notification area to notify users when they have lost connection to a networked file and are working on a local copy of the file. Users can then decide how to proceed. |
| **Reminder balloon frequency** [1] | Determines how often reminder balloon updates appear.<br><br>Reminder balloons appear when the user's connection to a network file is lost or reconnected, and they are updated periodically. By default, the first reminder for an event is displayed for 30 seconds. Then, updates appear every 60 minutes and are displayed for 15 seconds. You can use this setting to change the interval between updates. |
| **Initial reminder balloon lifetime** [1] | Determines how long the first reminder balloon for a network status change is displayed. |
| **Reminder balloon lifetime** [1] | Determines how long updated reminder balloons are displayed. |
| **Event logging level** [1] | Determines which events the Offline Files feature records in the event log.<br><br>Offline Files records events in the Application event log when it detects errors. By default, Offline Files records an event only when the offline files storage cache is corrupted. However, you can use this setting to specify additional events that you want Offline Files to record. |
| **Prohibit 'Make Available Offline' for these files and folders** [4] | Prohibits specific network files and folders from being made available for offline use.<br><br>This setting does not prevent files from being automatically cached if the network share is configured for automatic caching. It only affects the availability of the **Make Available Offline** menu option in the user interface.<br><br>If the **Remove 'Make Available Offline'** setting is enabled, this setting has no effect. |

| | |
|---|---|
| **Do not automatically make redirected folders available offline** | By default, makes all redirected shell folders, such as **My Documents**, **Desktop**, **Start Menu**, and **Application Data**, are available offline.<br><br>This setting allows you to change this behavior so that redirected shell folders are not automatically available for offline use. However, users can still choose to make files and folders available offline.<br><br>Do not enable this setting unless you are certain that users do not need access to all of their redirected files in the event that the network or the server holding the redirected files becomes unavailable.<br><br>This setting does not prevent files from being automatically cached if the network share is configured for automatic caching, nor does it affect the availability of the **Make Available Offline** menu option in the user interface. |

[1] This setting appears in the Computer Configuration and User Configuration folders. If both settings are configured, the setting in Computer Configuration takes precedence over the setting in User Configuration.

[2] This setting appears in the Computer Configuration and User Configuration folders. If both settings are configured for a particular server, the setting in Computer Configuration takes precedence over the setting in User Configuration. Both Computer and User configurations take precedence over a user's setting.

[3] This setting appears in the Computer Configuration and User Configuration folders. If both settings are configured, the settings are combined and all specified files are available for offline use.

[4] This setting appears in the Computer Configuration and User Configuration folders. If both settings are configured, the settings are combined, and the specified files are *not* available for offline use.

**Note**

- Most Offline Files policy settings in the User Configurationnode, listed in Table 7.16, also appear in the Computer Configurationnode. Table 7.17 lists *only* those settings that are unique to the Computer Configuration node. For more information about the settings, open the **Explain** tab on the **Properties** page of each policy.

**Table 7.17 Computer Configuration Policy Settings for Offline Files**

| Policy Setting | Description |
|---|---|
| **Allow or disallow use of the Offline Files feature** | Determines whether the Offline Files feature is enabled.<br><br>This setting also disables the **Enable Offline Files** option on the **Offline Files** tab. This prevents users from changing the option while a setting controls it.<br><br>By default, Offline Files is enabled in Windows 2000 Professional and is disabled in Windows 2000 Server. |
| **Default cache size** | Limits the percentage of the computer's disk space that can be used to store automatically cached offline files.<br><br>This setting also disables the **Amount of disk space to use for temporary offline files** option on the **Offline Files** tab. This prevents users from changing the option while a setting controls it.<br><br>Automatic caching can be set on any network share. When a user opens a file on the share, the system automatically stores a copy of the file on the user's computer.<br><br>This setting does not limit the disk space available for files that users make available offline manually. |
| **Files not cached** | Lists types of files that cannot be used offline.<br><br>This setting lets you exclude certain types of files from automatic and manual caching for offline use. The system does not cache files of the type that is specified in this setting even when they reside on a network share that is configured for automatic caching. Also, if users try to make a file of this type available offline, the operation will fail and the following message is displayed in the **Synchronization Manager Progress** dialog box: "Files of this type cannot be made available offline." |

| | |
|---|---|
| **At logoff, delete local copy of users offline files** | Deletes local copies of the user's offline files when the user logs off.<br><br>This setting specifies that automatically and manually cached offline files are retained only while the user is logged on to the computer. When the user logs off, the system deletes all local copies of offline files.<br><br>**Caution:** Files are not synchronized before they are deleted. Any changes to local files since the last synchronization are lost. |
| **Subfolders always available offline** | Makes subfolders available offline whenever their parent folder is made available offline.<br><br>This setting automatically extends the **Make available offline** policy setting to all new and existing subfolders of a folder. Users do not have the option of excluding subfolders. |
| **Encrypt the Offline Files cache** | Determines whether offline files are encrypted.<br><br>Offline files reside on a user's hard disk, not the network, and they are stored in a local cache on the computer. Encrypting this cache enhances security on a local computer. If the cache on the local computer is not encrypted, any encrypted files cached from the network will not be encrypted on the local computer. This might pose a security risk in some environments.<br><br>By default, this cache is protected on NTFS partitions by ACLs. |
| **Configure Slow link speed** | Configures the threshold value at which Offline Files considers a network connection to be slow.<br><br>Any network speed below this value is considered to be slow. |

# Security Recommendations for Shared Resources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Note**

- When you create a new shared resource, offline access is permitted by default; this means that secure shared resources can be stored offline on potentially unsecured computers. You must ensure that you set appropriate permissions by using share permissions or access control on the NTFS file system.

**Assigning permissions to groups, not to user accounts**
Assigning permissions to groups simplifies the management of shared resources because you can then add users to or remove them from the groups without having to reassign permissions. To deny all access to a shared resource, deny the Full Control permission.

**Assigning the most restrictive permissions that still allow users to perform required tasks**
For example, if users only need to read information in a folder, and they will never delete, create, or change files, assign the Read permission.

**Centralizing the administration of shared resources**
On computers running Windows XP Professional that are connected to a domain, allow access to shared resources by using domain user accounts, instead of by using local user accounts. This centralizes the administration of share permissions.

**Denying permissions**
Avoid explicitly denying permissions to a shared resource unless you want to override specific permissions that are already assigned.

**Using NTFS file system permissions or access control shared resources**
Use NTFS file system permissions and access control if users log on locally to access shared resources (such a Terminal Services). Share permissions apply only to users who access shared resources over the network; they do not apply to users who log on locally.

**Organizing resources for objects with similar security requirements**
Organize resources so that objects with the same security requirements are located in the same folder. For example, if users require Read permissions for several application folders, store the application folders in the same parent folder. Then, share the parent folder, instead of each individual application folder. If you need to change the location of an application, you might need to reinstall it.

**Limiting membership in Administrators group**
To enable administrators to manage application software and to control user rights, limit membership in the Administrators group, and assign Full Control permissions to that group.

**Creating a strong password**
To keep drives secure, make sure that you use a strong password for all accounts. A strong password has the following characteristics:

- Is at least seven characters long.
- Does not contain your user name, real name, or company name.
- Does not contain a complete dictionary word.
- Is significantly different from previous passwords and does not contain incremental numbers or letters — such as Password1, Password2, Password3 or PasswordA, PasswordB, and so on.
- Contains characters from each of the following four groups: upper case letters, lower case letters, numerals, and symbols. An example of a strong password is J*p2leO4>F.

You can use Group Policy to enforce password complexity requirements. You can configure the **Password must meet complexity requirements** security setting by selecting the appropriate GPO and expanding the console tree: **Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy**. For more information about password policy, see "Apply or modify password policy" in Help and Support Center for Windows Server 2003.

**Using a firewall**

To protect shared resources from unauthorized Internet access, use a firewall. For more information about making shared resources more secure, see "Securing shared resources" in Help and Support Center for Windows Server 2003.

# Configuring Offline Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Offline Files technology applies to any mapped or mounted drive that contains documents that users might want to use offline. You set up and configure Offline Files independently of the Folder Redirection snap-in extension. The default state (Enabled or Disabled) of the Offline Files feature depends on the operating system version that you are running, as indicated in Table 7.18.

**Table 7.18 Default State of the Offline Files Feature**

| Operating System Version | Offline Files State |
|---|---|
| Windows 2000 Professional | Enabled |
| Windows XP Professional | Enabled |
| Windows 2000 Server | Disabled |
| Windows Server 2003 family products | Disabled. The Offline Files feature is not compatible with Terminal Services. |

**Enabling the Offline Files Feature for Windows Server 2003 Computers**

By enabling the Offline Files features in Windows Server 2003, you can set up a computer to use Offline Files or make file or folder available offline.

**To enable Offline Files on a server running Windows Server 2003**

1. Right-click **My Computer**, and then click **Properties.**
2. In the **System Properties** dialog box, click the **Remote** tab.
3. Under **Remote Desktop**, clear the **Allow users to connect remotely to this computer** check box, and then click **OK**.

**To set up a computer to use Offline Files**

1. In **My Computer**, on the **Tools** menu, click **Folder Options**.
2. On the **Offline Files** tab, select **Enable Offline Files**.

   -Or-

   To do a quick synchronization, clear the **Synchronize all offline files before logging off** check box.

After you set up a computer to use Offline Files, select the files and folders that you want to make available offline.

**To make a file or folder available offline**

1. Open **My Computer**, and then right-click the shared file or folder that you want to make available offline.
2. On the **File** menu, click **Make available offline**.
3. Follow the instructions in the Offline Files Wizard.

**Configuring Offline Settings**

Offline Files stores a version of the shared resources on the client computer in the file system *cache*, a reserved portion of disk space, which the client computer can access even when it is disconnected from the network. You can set caching options for

each shared resource to specify which files or programs will be available offline.

**To configure offline settings for shared files or folders**

1. Open Computer Management.

2. In the console tree, navigate to System Tools, Shared Folders, and then Shares.

3. In the details pane, right-click the shared file or folder for which you want to configure offline settings, and then click **Properties**.

4. On the **General** tab, click **Offline Settings**.

5. In the **Offline Settings** dialog box, select one of the following options, and then click **OK**.

   - **Only the files and programs that users specify will be available offline**. (The default option that gives users control over which files are available offline.) .

   - **All files and programs that users open from the share will be automatically available offline**. (Allows all the files that users open from the shared folder to be automatically available offline.) If you select the **Optimized for performance** check box, all programs are automatically cached so that they can run locally. This option is especially useful for file servers that host applications because it reduces network traffic and improves server scalability.

   - **Files or programs from the share will not be available offline**. This option prevents users from storing files offline.

**Note**

- When you use the option for automatic caching of all files and programs, be sure to restrict permissions on the files and programs that are contained in the shared folder to read-only access.

After you configure Offline Files, configure Synchronization Manager to periodically ensure that the Offline Files cache has a complete and correct copy of every file that is designated for offline use.

On client computers running Windows XP Professional or Windows Server 2003, redirected folders are automatically added to Offline Files storage. In Windows 2000, you can enable this feature, but it is not enabled by default.

You can make individual files from a shared network folder or an entire shared network folder available offline. If you create a shortcut to a *file* that is available offline, the file is also available offline. If you create a shortcut to a *folder* that is available offline, the contents of that folder are *not* available offline.

For a procedure to store offline files and folders in a network share, see "Set client-side caching for a File Share resource" in Help and Support Center for Windows Server 2003.

**Tip**

- Use the Group Policy setting **Administratively assigned offline files** to set multiple folders at the same time for offline use.

# Configuring Synchronization Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Synchronization Manager maintains offline files that are stored in the Offline Files cache on the user's workstation. You can configure Synchronization Manager to initiate Offline File synchronization in response to one or more of the following situations:

- When the user logs on or off the computer
- At specific intervals while the computer is idle
- At scheduled times

**Synchronizing Files and Folders**

Use the **Synchronization Manager** wizard to specify the items to synchronize and to define options for the synchronization process.

**To open Synchronization Manager**

- In the **Run** dialog box, type **mobsync**, and then click **OK**.

**To synchronize offline items**

1. Open Synchronization Manager.
2. In the **Select the check box for any items you want to synchronize** box, under **Name**, select the check box for each offline item that you want to synchronize, and then click **Synchronize**.

**To synchronize individual files, folders, or Web pages**

1. Open Windows Explorer, My Computer, or Internet Explorer.
2. Select the item that you want to synchronize, and then on the **Tools** menu, click **Synchronize**.

### To synchronize offline items when you log on to or log off from a computer

1. Open Synchronization Manager, click **Setup**, and then click the **Logon/Logoff** tab.
2. In **When I am using this network connection**, click the network connection that you want to use.
3. In **Synchronize the following checked items**, select the check box for the offline items that you want to synchronize.
4. Under **Automatically synchronize the selected items**, do one of the following:
    - To synchronize items when you log on to your computer, select the **When I log on to my computer** check box.
    - To synchronize items when you log off of your computer, select the **When I log off my computer** check box.
    - To request permission before synchronizing your offline items, select the **Ask me before synchronizing the items** check box.

**To schedule offline items to synchronize**

1. Open **Synchronization Manager**,click **Setup**, and then click the **Scheduled** tab.

2. Click **Add** to start the Scheduled Synchronization Wizard, and then follow the instructions in the wizard.

**To synchronize offline items while your computer is idle**

1. Open Synchronization Manager, click **Setup**, and then click the **On Idle** tab.

2. Click the network connection that you want to use.

3. In **When I am using this network connection**, click the network connection that you want to use.

4. In **Synchronize the following checked items**, select the check box for the offline items that you want to synchronize.

5. Select the **Synchronize the selected items while my computer is idle** check box.

# Assigning Profile and Disk Quotas

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, the disk space that users need varies, depending on users' roles.

Consolidating user disk space on network shares is more efficient and cost-effective than providing large amounts of storage at each workstation. However, you need to manage the shares. Profile and disk quotas can assist you in monitoring and managing network shares.

#### Profile Quotas and Local Drives

Profile quotas do not restrict redirected folders because they are not stored in the user's profile. If you want to reduce disk space requirements of workstations, do not save profiles locally. To establish quotas, apply profile quotas instead of disk quotas to local disk drives.

#### Combining Redirection of My Documents with Roaming User Profiles

If you combine Folder Redirection of the **My Documents** folders with Roaming User Profiles, do *not* apply disk quotas to the volume in which roaming profiles are stored. Instead, use a profile quota to manage the size of the profile, and use a disk quota to manage the size of the **My Documents** and **My Pictures** folders.

### Note

- If you are combining Folder Redirection of the **My Documents** folder and roaming user profiles, it is recommended that you do *not* use disk quotas on the profile folders. The operating system and applications write items to the user profile folder on the user's behalf. This includes items such as Custom.dic and Favorites.

#### Limiting Profile Size

To manage the profile quota size, use the **Limit profile size** policy setting in the Group Policy Object Editor snap-in. This policy sets the maximum size of each roaming user profile and determines the system's response when a roaming user profile reaches the maximum size. The **Limit profile size** policy setting is available in the **User Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor.

#### Removing Cached Copies of Roaming Profile to Minimize Disk Space Usage

If disk size on a multi-user computer is a concern (for example, a public computer that numerous users use), you can use the **Delete cached copies of roaming profiles** policy, which removes cached versions of the profile when a user logs off. The **Delete cached copies of roaming profiles** policy setting is available in the **Computer Configuration\Administrative Templates\System\User Profiles** node of the Group Policy Object Editor.

# Setting Quotas on User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To manage the size of a user profile, you can set a maximum profile quota or remove local copies of a user's roaming profile from the workstation when a user logs off.

**To set a maximum profile quota**

1. Open the Group Policy Object Editor snap-in.
2. In **User Configuration\Administrative Templates\System\User Profiles**, double-click the **Limit profile size** policy setting.
3. In the **Limit profile size Properties** page, click **Enable**, specify the options you want to use, and then click **OK**.

**To remove local copies of a user's roaming profile from the workstation when a user logs off**

1. Open Group Policy Object Editor.
2. In the **Computer Configuration\Administrative Templates\System\User Profiles** folder, double-click the Delete cached copies of roaming profiles policy setting.
3. Click **Enable**, and then click **OK**.

Deleting cached profiles when the user logs off adds time to the logon process; however, it is recommended for workstations in locations that are not secure. User profile data needs to be treated as confidential. Ideally, such data is only retained on computers that are physically secure.

# Over-Allocating Disk Space

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some tasks, such as moving large folders within a disk volume, require a user to temporarily allocate a large amount of disk space. This temporary space is more readily available if it is pulled from a large pool of space. Therefore, you need to allocate more space for each user than they need for long-term storage.

Determination of the total quantity of shared disk space to be made available to users is influenced by historical requirements, planned deployments of new servers and user accounts, and your hardware purchasing budget. Typically, the median disk usage is much higher than the average disk usage due to the varying allocations of disk space in a large group of users.

**Requirements for Managing Disk Quotas**

Disk quotas track and control disk space usage for user data that is stored on the network. You can set disk quotas by using Group Policy. When disk quotas are enabled in Group Policy, the settings affect all NTFS volumes on all computers running Windows 2000, Windows XP Professional, or Windows Server 2003 to which the GPO applies. Alternatively, you can manually set disk quotas on individual servers.

To administer disk quotas, you must be a member of the Administrators group. You can set quotas on local volumes, network volumes, and removable drives that are formatted with NTFS. The network volumes must be shared from the volume's root directory, and removable drives must be shared.

**Recommendations for Using Disk Quotas**

Apply disk quota limits according to the actual disk space your users need.

1. Try to estimate how much disk space various types of users will require and then categorize users according to their disk space requirements (for example, users who work with art programs might require a large amount of disk space).

2. Next, structure your disk volumes by those categories, and then use disk quotas to limit the amount of disk space allowed for the users on each volume.

3. Also, consider metadata. Each file stored on the volume can use up to 64 kilobytes (KB) of NTFS metadata that is not applied to a user's quota limit. To avoid running out of disk space, ensure that you provide sufficient disk space to accommodate this metadata.

**Setting default quota limits**

Set moderately restrictive default limits for all user accounts, and then modify the limits to allow more disk space to users who work with large files. You might want to increase disk quota limits for a few user accounts rather than to force some users to work with a quota limit that does not meet their needs. User policy settings occupy at least 2 MB of disk space, so you should set the default limits for all user accounts above 2 MB.

**Setting limits that permit growth in disk space use**

Set the quota level to allow for increasing use of disk space. Increase quota when necessary as a percentage of the current quota setting.

**Group Policy for Disk Quotas**

Table 7.19 lists the Group Policy settings for disk quotas. These settings are available in the **Computer Configuration\Administrative Templates\System\Disk Quotas** node.

Table 7.19 Disk Quota Policy Settings in Disk Quotas

| Policy Setting | Description |
|---|---|
| **Enable disk quotas** | Enables and disables disk quota management on all NTFS volumes of the computer and prevents users from changing the setting. |
| **Enforce disk quota limit** | Determines whether disk quota limits are enforced and prevents users from changing the setting. |
| **Default quota limit and warning level** | Specifies the default disk quota limit and warning level for new users of the volume. |

| | |
|---|---|
| **Log event when quota limit exceeded** | Determines whether the system records an event in the local Application log when users reach their disk quota limit on a volume, and prevents users from changing the logging setting. |
| **Log event when quota warning level exceeded** | Determines whether the system records an event in the Application log when users reach their disk quota warning level on a volume. |
| **Apply policy to removable media** | Extends the disk quota policies in this folder to NTFS file system volumes on removable media. |

**Group Policy for Disk Quota Processing**

Administrators can use Group Policy to specify how often disk quotas are applied. The **Disk quota policy processing** policy setting is available under **Computer Configuration\Administrative Templates\System\Group Policy**. This policy applies to all policy settings that use the disk quota component of Group Policy, such as those in **Computer Configuration\Administrative Templates\System\Disk Quotas**.

The **Quota policy processing** policy setting includes the following options:

- **Allow processing across a slow network connection.** Updates the policies even when the update is being transmitted across a slow link. Updates across slow connections can cause significant delays.

- **Process even if the Group Policy objects have not changed.** Updates and reapplies the policies even if the policies have not changed. Many policy implementations specify that they are updated only when changed. However, you might want to update unchanged policies, such as reapplying a setting in case a user has changed it.

- **Do not apply during periodic background processing.** Prevents the system from updating affected policies in the background while the computer is in use.

  **Caution**

  - Background updates can disrupt the user, cause a program to stop or operate abnormally, and, in rare cases, damage data.

Administrators and users with rights to create computer objects can circumvent quota restrictions. This is because they can create any number of computer accounts, each with the default quota limit.

# Configuring Disk Quotas on NTFS Volumes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To enable disk quotas, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

**To enable disk quotas**

1. Open My Computer, right-click the disk volume for which you want to enable disk quotas, and then click **Properties**.
2. On the **Quota** tab, click the **Enable quota management** check box.
3. Specify the quota limits, thresholds, and event logging options that you want to use, and then click **OK**.

Table 7.20 shows the available options for disk quotas.

**Table 7.20 Quota Management Options**

| Option | Description |
|---|---|
| **Deny disk space to users exceeding quota limit** | Users who exceed their quota limit receive an "insufficient disk space" error message from Windows and cannot write additional data to the volume without first deleting or moving some existing files from it. Individual programs determine their own error handling for this condition. To the program, it appears that the volume is full. If you clear this check box, users can exceed their quota limit. Enabling quotas and not limiting disk space use are useful when you do not want to deny users access to a volume but want to track disk space use on a per-user basis. You can also specify whether to log an event when users exceed either their quota warning level or their quota limit. |
| **Limit disk space to** | Enter the amount of disk space that new users of the volume are allowed to use and the amount of disk space that must be used before an event is written to the system log. Administrators can view these events in Event Viewer. You can use decimal values (for example, 20.5). For the disk space and warning levels, select the appropriate units from the drop-down list (for example, KB, MB, or GB). |
| **Log event when a user exceeds their quota limit** | If quotas are enabled, an event is written to the system log on the local computer whenever users exceed their quota limit. Administrators can view these events in Event Viewer and filter for disk event types. By default, quota events are written hourly to the system log on the local computer. You can change the interval at which quota events are written to the system log on the local computer using the fsutil behavior command. Fsutil is a command-line tool that you can use to perform tasks related to FAT and NTFS file systems. |
| **Log event when a user exceeds their warning level** | If quotas are enabled, an event is written to the system log on the local computer whenever users exceed their quota warning level. Administrators can view these events in Event Viewer, filtering for disk event types. By default, quota events are written hourly to the system log on the local computer. You can change the interval at which quota events are written to the system log on the local computer by using the **Fsutil behavior** command. |

For more information about using disk quotas, see "Disk Quotas" in Help and Support Center for Windows Server 2003. For more information about the command-line tool Fsutil, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Command-line reference A-Z.

# Desktop Strategies for Computers Running Windows 2000 and Windows XP Without Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have computers running Windows 2000 Professional and Windows XP Professional in a domain that does *not* include Active Directory, you can manage desktops locally by implementing the following features:

**Roaming User and Logon Scripts** When using either a Windows NT 4.0 domain or Active Directory, both roaming user profiles and logon scripts are configured on the user object.

**Folder Redirection** You can redirect special folders to alternate locations, either to a local or network location by modifying the values of the registry entries in the following registry subkey:
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

**Redirecting folders by using logon scripts** Group Policy is the recommended method for redirecting user's folders. You can also achieve similar results by using logon scripts to set values of the registry entries in the subkey **User Shell Folders**. This approach provides functionality that is similar to Folder Redirection.

In Windows NT 4.0 environments, you can use System Policy to set the appropriate values. However, when you use System Policy to set values, the registry settings persist. And you do not get the advantages of using Group Policy to set paths such as automatic moving of files when the path changes.

**Internet Explorer Maintenance** Instead of using Group Policy to control Internet Explorer settings, administrators can use the Internet Explorer Administration Kit (IEAK) to apply settings to Internet Explorer clients using auto-configuration packages. For more information about IEAK, see the Microsoft Internet Explorer Administration Kit (IEAK) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**System Policy (for registry-based policy)** Domain-based Group Policy processing requires that the User and Computer objects be located in Active Directory. If the User or Computer objects are located in a Windows NT 4.0 domain, then Windows NT 4.0 System Policy is processed for whichever of these objects is located in that domain — this could be the Computer or User object, or both. System Policy is defined as the policy mechanism that is used natively in Windows NT 4.0; it is a set of registry settings that together define the computer resources that are available to a group of users or an individual.

**Local Group Policy object** You can set settings in the local Group Policy object (LGPO) for any computer, whether or not it participates in an Active Directory domain. Although System Policy scales more easily to a large number of clients, the LGPO can be useful if you only need to apply certain settings to a small number of Windows XP Professional or Windows 2000–based clients in a Windows NT 4.0 or other domain.

The LGPO is located at *systemroot*\System32\GroupPolicy. Not all Group Policy extensions are available for the local GPO. Each Group Policy extension snap-in queries the Group Policy engine to get the GPO type, and then determines whether the GPO is to be displayed. To set the LGPO, use the Group Policy snap-in focused on the local computer. You can use a LGPO to set policy for these Group Policy extensions only: scripts, security settings, Administrative Templates, and Internet Explorer maintenance; you cannot set policy for Software Installation, Folder Redirection, and Remote Installation Services.

# Additional Resources for Implementing User State Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- The *Distributed Services Guide* of the *Windows Server 2003 Technical Reference* (or see the Distributed Services Guide on the Web at http://www.microsoft.com/reskit) for more information about Group Policy and advanced Group Policy.
- "Deploying a Managed Software Environment" in this book.
- The User Data and Settings Management link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about managing user data and settings for Windows XP in a Windows 2000 environment.
- The Group Policy Management Console link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about Group Policy management tools.

**Related Tools**

- Diruse.exe

  You can use the command-line tool Diruse.exe to gather disk usage information for compressed files and folders in NTFS volumes, such as how much disk space each folder uses in a directory tree on a disk volume. For more information about Diruse.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Windows Support Tools.

- Fsutil

  You can use the **Fsutil** command-line tool to perform many tasks related to file allocation table (FAT) and NTFS file systems, such as managing reparse points, managing sparse files, or dismounting a volume. **Fsutil** is a powerful tool that should only be used by administrators who have a thorough knowledge of Windows XP or the Windows Server 2003 family of operating systems. For more information about **Fsutil**, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Command-line reference A-Z.

**Related Job Aids**

- "Worksheet A.39 Determining User Requirements" (DMEUSE_39.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.39 Determining User Requirements" on the Web at http://www.microsoft.com/reskit).
- "Worksheet A.40 Determining Server Disk Space Requirements" (DMEUSE_40.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.40 Determining Server Disk Space Requirements" on the Web at http://www.microsoft.com/reskit).
- "Worksheet A.41 Considerations for Mobile Users" (DMEUSE_41.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.41 Considerations for Mobile Users" on the Web at http://www.microsoft.com/reskit).

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Using roaming user profiles" in Help and Support Center for Windows Server 2003 for more information about implementing roaming user profiles.
- Group Policy topics in Help and Support Center for Windows Server 2003 for more information about working with GPOs.
- "Create a new user account" in Help and Support Center for Windows Server 2003.
- "Make a file or folder available offline" in Help and Support Center for Windows Server 2003.
- "Add a home folder to a profile" in Help and Support Center for Windows Server 2003.
- "Handling file conflicts" in Help and Support Center for Windows Server 2003.
- "Apply or modify password policy" in Help and Support Center for Windows Server 2003.
- "Securing shared resources" in Help and Support Center for Windows Server 2003 for more information about making shared resources more secure.
- "Set client-side caching for a File Share resource" in Help and Support Center for Windows Server 2003 for a procedure to store offline files and folders in a network share.
- "Disk Quotas" in Help and Support Center for Windows Server 2003.

# Deploying a Managed Software Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using Group Policy–based software management, you can centrally deploy, install, and manage applications throughout an organization. From a central location, you can also perform routine maintenance tasks such as upgrading, patching, and removing applications without going to individual workstations. You can also configure Software Restriction Policies to prevent users from running unknown or dangerous programs.

**In This Chapter**

Deploying a Managed Software Environment Overview

Preparing Applications for Deployment

Deploying Applications in a Managed Environment

Migrating Applications to a Managed Environment

Patching, Upgrading, and Removing Applications

Troubleshooting Software Deployment

Additional Resources for Deploying a Managed Software Environment

**Related Information**

- For information about deploying Group Policy and using the new Group Policy Management console (GPMC) MMC snap-in, see "Designing a Group Policy Infrastructure" in this book.

- For more information about managing Group Policy–based software deployment, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Deploying a Managed Software Environment Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Software administrators face ongoing challenges for deploying and maintaining software. They must ensure that users have reliable access to necessary software while preventing unknown or dangerous software from being installed or run. Users can, for example, unintentionally turn a desktop environment into an administrator's nightmare by deleting files and programs or by downloading unsupported software from the Internet or from removable media such as a CD-ROM.

The software installation extension of Microsoft® Windows® Server 2003 Group Policy enables you to create a controlled environment, providing on-demand software installation and automatic repair of applications. Users benefit from reliable access to the applications that they need to perform their jobs on any computer they use on their network.

For worksheets to assist you with the deployment processes discussed in this book, see Additional Resources for Deploying a Managed Software Environment, later in this chapter.

# Deploying a Managed Software Environment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The software administrator needs to manage software throughout all phases of the software administration life cycle. Managing this life cycle can shorten the time it takes to deploy software, and can increase the productivity of users. Figure 8.1 illustrates the process for deploying a managed software environment by using Group Policy-based software deployment.

**Figure 8.1 Deploying Managed Software**



# Assessing Microsoft Software Management Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft offers several software solutions for your networked users. Your organization might already have objectives and requirements for a software installation and management product. Before you plan your deployment, you must verify those objectives to be certain that you use the appropriate technologies for software deployment.

Group Policy, which is built-in to Microsoft® Windows® 2000 and later operating systems, offers a convenient method for distributing software in your Active Directory® directory service environment, especially if you are already using Group Policy for other purposes, such as securing your client and server computers. However, a Group Policy-based software installation has some basic limitations, including difficulties with scheduling installation, consistently managing network bandwidth, and providing feedback on the status of the installation. If you need to carefully schedule installations, manage network use, perform hardware and software inventory, or monitor installation status, consider using Microsoft® Systems Management Server (SMS). For more information about SMS, see the Microsoft Systems Management Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources/.

Using the right solutions can benefit your organization by giving you a centralized, efficient means to perform routine tasks such as updating software. Table 8.1 compares the various software management technologies.

**Software installation extension of Group Policy** You can use the software installation extension of Group Policy to deploy and manage software if your organization is small or medium in size, and the following conditions exist:

- You have deployed Active Directory.
- You have determined that Group Policy provides the management features your organization requires.
- You have a solid base of client computers running Microsoft® Windows® 2000 Professional or Windows® XP Professional and member servers running Microsoft® Windows® 2000 Server and Windows Server 2003.

**Note**

- Your servers do not have to run Windows Server 2003 for you to use Group Policy.

Group Policy can also serve the needs of large enterprises that use other software installation solutions, such as SMS, from the top level across the organization. Consider using Group Policy for distributing software within various groups, such as individual divisions, where you might not need the advanced capabilities of SMS.

**Software Update Services** You can use Microsoft® Software Update Services (SUS) to quickly acquire and distribute critical Windows patches to computers in your organization. By using SUS, you can choose which of the latest critical or security patches to download, test them in a company-standard operating environment, and then efficiently deploy the patches to the appropriate computers running the Automatic Updates client.

**Systems Management Server** You can use Systems Management Server (SMS) if any of the following conditions exist:

- Your organization is medium or large in size.
- Your users are running operating systems earlier than Windows 2000 Professional.
- You require more advanced capabilities for planning, scheduling, distributing, and tracking software.

The advanced capabilities of SMS include such features as inventory-based targeting, status reporting, server-side and client-side scheduling, multisite facilities, centralized hardware and software inventory, remote diagnostic tools, software metering, software distribution-point population and maintenance, support for Microsoft® Windows® 95, Windows® 98, Windows NT® 4.0, Windows 2000, and Windows® XP clients, and enhanced software deployment features. Additionally, SMS does not require Active Directory. For more information about SMS, see the Microsoft Systems Management Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Terminal Services** You can use Microsoft® Terminal Services if you have Windows-based desktop applications that require

frequent updates, and the users who require those applications are in remote locations and have low bandwidth. When used as a terminal server, a server becomes a Windows application server. This allows the user to run Windows-based applications remotely on the server while only the mouse, keyboard, and display data are transmitted to the local computer. By using Terminal Services, you can offer your users software as a remote service instead of as a local installation package. For more information about using Microsoft® Terminal Services, see "Hosting Applications with Terminal Server" in *Planning Server Deployments* in this kit.

**Table 8.1 Comparing Software Management Technologies**

| Management Function | Group Policy | SMS | Terminal Services | SUS |
|---|---|---|---|---|
| Patch and upgrade Windows XP, Windows Server 2003, and Windows 2000 | N/A | Yes<br>When using SMS for software management, also use it to patch your Windows systems instead of SUS. | Although Terminal Services does not automate patching, you can use it to remotely log on and apply patches. | Windows patches only (no upgrade) |
| Consistent user environment (persistence of data, software, and settings) | Yes | Software only | Yes | N/A |
| Disaster recovery for applications in Windows 2000 and Windows XP | Yes | Yes | N/A | N/A |
| Inventory, advanced deployment, troubleshooting, and diagnostic tools | Limited | Yes | Limited | None |
| Manage environments that are not Active Directory-based | No | Yes | Yes | Yes (Windows patches only) |

Although all these Microsoft management technologies provide important software distribution capabilities, SMS is the preferred Microsoft software distribution solution for medium-sized, and especially for enterprise-sized, organizations. SMS provides advanced features for deploying and managing software, Windows patches, and critical updates. If you use SMS as your software management solution, use the SMS Feature Pack, instead of SUS, to distribute patches and critical updates to your clients. However, SUS, used with the Automatic Updates client, is the recommended solution for distributing Windows patches in conjunction with Group Policy–based software distribution.

Although there are certain instances where you would choose one software deployment method over another, you can also use many of these Microsoft technologies together, depending on your needs. For more information about using these Microsoft software deployment methods to provide a combined solution, see the Application Deployment link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Group Policy Software Deployment Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To deploy software using Group Policy, you must have an Active Directory–based domain and Windows 2000 or Windows Server 2003 domain controllers. Also, the clients must run Windows 2000 Professional or Windows XP Professional. By using other Windows Server 2003 features and technologies, such as those described in Table 8.2, you can take full advantage of Group Policy-based software deployment.

**Table 8.2 Essential Tools and Components for Deploying Software in a Managed Environment**

| Component or Tool | General Description | Combined with Group Policy Software Installation Extension |
|---|---|---|
| Active Directory | A hierarchical collection of objects including domains, sites, OUs, users, computers, and printers that allow an organization to manage these resources. | Provides the scope of management mechanism to locate users and computers.<br>Stores software deployment information through Group Policy. |
| Group Policy | An administrative tool for defining and controlling the way programs, network resources, and the operating system work for users and computers in an organization. In an Active Directory environment, you apply Group Policy to users or computers on the basis of their membership to sites, domains, or OUs. | Enables you to deploy applications in a Group Policy object (GPO) associated with one or more Active Directory containers, such as sites, domains, or OUs. Use the software installation extension of the Group Policy Object Editor Microsoft Management Console (MMC) snap-in to deploy applications. |
| Windows Installer | A service based on an operating system, which provides software installation services using a standard package format. You can use Windows Installer to manage the installation, modification, upgrade, and removal of software applications. | Installs, modifies, upgrades, and removes software applications. |
| Software installation extension of the Group Policy Object Editor MMC snap-in | An extension of the Group Policy Object Editor MMC snap-in that includes a user interface that allows administrators to deploy and manage software. | Communicates with Active Directory, GPOs, and Windows Installer to assign or publish applications as follows:<br>• Assigns software to users. Installs user-assigned applications entirely the first time the user logs on after deployment, or allows users to install certain components or features of an application as needed.<br>• Assigns software to computers.Installs an application the next time the computer starts. The application is available for all the users on that computer.<br>• Publishes applications for users only: Users can choose to install the software from a list of published applications located in **Add or Remove Programs** in Control Panel. |

| | | |
|---|---|---|
| Group Policy Management Console (GPMC) | A new tool that consists of an MMC snap-in and command-line tools. This tool unifies management of all aspects of Group Policy across an enterprise. GPMC allows you to manage all GPOs, Windows Management Instrumentation (WMI) filters, and permissions on your network. | Group Policy Modeling (formerly known as RSoP planning) allows you to run hypothetical scenarios to verify software configurations under various sites, domains, and OUs. Provides printable HTML reports.<br><br>Group Policy Results(formerly known as RSoP logging)verifies which software applications are properly installed for a specific group of users or computers. It also pinpoints the causes of unintended removal or damage to software. Provides HTML printable reports. |
| Add or Remove Programs | A user interface in Control Panel of Windows XP Professional and Windows 2000 Professional. **Add or RemovePrograms** lets users manage software on their own computers. | Lists both published and assigned applications so that users can install, modify, and remove software from their desktop computers. |
| Command line and Graphical User Interface (GUI) tools or scripts | These include GPResult.exe, GPOTool.exe, GPUpdate.exe, ReplMon.exe, NetDiag.exe, InstallShield, and the new Group Policy Management MMC snap-in. Some are installed by default; others must be installed manually. | Helps you manage, optimize, or troubleshoot Group Policy-based software deployment. |

The software installation extension of Group Policy allows you to centrally manage the installation of software on all client computers in your organization. You do this either by assigning applications to users or computers, or by publishing applications for users.

Assign software on a per-user or per-computer basis when you do not want to give users the choice to install or remove the software. For example, if a user accidentally removes a user-assigned application by using **Add or Remove Programs**, the software installation extension of Group Policy automatically reapplies the advertisement information after the user logs on or the computer restarts, and the software is reinstalled the next time a user selects it. It is not possible for a user to delete a computer-assigned application. In most cases, packages that you assign to users or computers include applications that are essential but do not create congestion between the clients and the software distribution points.

If you use Group Policy-based software deployment, you can publish software for users only (not available for computers). When you publish software for users, you give them the opportunity to decide if and when they want to install it. They can install the software from a list of published applications in **Add or Remove Programs** in Control Panel. For example, not everyone in the organization requires software for project management. Therefore, a software administrator is likely to publish a project management package for only those users who require it. Managers who require the software can then choose to install it.

Users can always see both assigned and published applications in **Addor Remove Programs**.

For more information about assigning software to users and computers and publishing software for users, see "Assigning and Publishing Software" later in this chapter.

# Preparing Applications for Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Installer is the engine for deploying software packages when you use the software installation extension of Group Policy. It is a core service of the Windows operating system, which enables software administrators to manage the state of software applications. The process of preparing applications involves determining the appropriate packaging method, as indicated in Figure 8.2.

**Figure 8.2 Preparing Applications for Deployment**




Windows Installer technology uses the following two components to help you install and manage software:

- A Windows Installer package (an .msi file), which is a database containing information that describes the installed state of an application. The Windows Installer package performs the installation, modification, and removal of the software.
- An application programming interface (API) that allows applications and management tools to interact with Windows Installer to install or remove additional features of the application after the initial installation is completed.

The managed state of an application includes installation, modification, upgrade, or removal. Windows Installer provides you with consistent and reliable methods to customize installations, update and upgrade applications, and resolve configuration problems. It can assist you as follows:

- Helps you manage the state of software during and after installation.
- Defines a standard format for application setup and tracks components such as groups of files, registry entries, shortcuts, and other aspects of the application that must be managed together.
- Detects whether software is installed correctly or whether a program file is missing. It can immediately reinstall the damaged or missing files.
- Repairs applications and ensures that they are installed or removed without overwriting or deleting files that are required by another application.
- Enables you to modify an installation by adding or removing features after the installation.
- Enables you to deploy 32-bit and 64-bit Windows applications (Version 2.0).

**Note**

- Versions of Windows Installer earlier than version 2.0 can install and manage 32-bit Windows Installer packages only on 32-bit operating systems. Windows Installer version 2.0 supports all of the setup functionality that is available in earlier Windows Installer versions.

# Determining the Preparation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you start the software distribution process, determine which preparation method to use for each new application, patch, or upgrade that you plan to deploy. The following methods are available:

- **Use native Windows Installer packages.** If the application you are installing includes a built-in .msi package you can either deploy the software as-is, or customize it further. For information about using transforms for customizing Windows Installer packages, see "Packaging Software for Deployment" later in this chapter.

- **Reauthor the setup program to include a native .msi.** This method is not recommended except for expert Windows Installer package authors. However, when used, this method is most appropriate if the application is relatively simple, or you have thorough knowledge of its structure and of the application setup on the Windows platform that uses Windows Installer. For more information about reauthoring applications for Windows Installer, see "Packaging Software for Deployment" later in this chapter.

- **Create a Software Installation Settings (.zap) file.** A .zap file is a text file similar to an .ini file, which contains instructions that allow Windows to publish an application (Setup.exe) for users to install by using **Add or Remove Programs** in **Control Panel**. To publish applications that do not install by using Windows Installer, you must create a .zap file, copy the .zap file to the software distribution point servers, and then use the Group Policy–based software deployment to publish the application for the users. You cannot use .zap files for assigned applications.

For information about using other installing programs, see article 231747 in the Microsoft Knowledge Base.

Repackaging applications into an .msi format has limitations and application manufacturers typically do not support it. Therefore, consider repackaging as a last resort, except for applications that are developed specifically for your organization. For information about repackaging applications for Windows Installer, see the Windows Security Collection in the Windows Server 2003 Technical Reference on the Web (http://go.microsoft.com/fwlink/?linkid=4571).

Table 8.3 describes the advantages and disadvantages of the various packaging methods.

**Note**

- Packages that are created on Windows 2000 by using Veritas WinInstall LE (the repackaging software that is included with Windows 2000) work on target computers running Windows XP. However, Veritas WinInstall LE does not run on Windows XP and is not intended for repackaging new .msi files.

**Table 8.3 Comparing Packaging Methods**

| Advantages and Disadvantages | Natively Author or Reauthor Windows Installer-Based Packages | Repackaging | .zap Files |
|---|---|---|---|

| | | | |
|---|---|---|---|
| Advantages | Can benefit from all Windows Installer features, including just-in-time feature installation, feature repair, and installation with elevated permissions.<br><br>Can be run without user intervention.<br><br>Can be assigned or published.<br><br>Does not require user to be a local administrator to install.<br><br>Can automatically repair itself if key files are damaged or missing.<br><br>If application includes native .msi package, easy to deploy.<br><br>Can roll back an unsuccessful installation, modification, repair, or removal. | Same benefits as natively authoring. | Easy to create and fast to deploy.<br><br>Display application in **AddorRemove Programs**. |
| Disadvantages | Time-consuming to build and test. | Can easily fail if repackaging not performed on a clean computer.<br><br>Time-consuming to build and test.<br><br>Control over installation less detailed than natively authoring. | Run existing setup, requiring user intervention.<br><br>Do not benefit from Windows Installer features.<br><br>Might require user to have local administrator permissions to install.<br><br>Cannot automatically install software or features on demand.<br><br>Cannot automatically repair if key files are damaged or missing.<br><br>Cannot roll back an unsuccessful operation. |

**Important**

- Before you reauthor or repackage applications in your organization:

    1. Check with the application manufacturer or internal development group to see if either has or is planning to develop a native Windows Installer implementation of the program.

    2. Weigh the benefits of reauthoring and repackaging against the costs involved. Keep in mind that both are advanced operations. There are tools for each operation that can help software package developers create the final Windows Installer package, but the procedures are still quite resource-intensive and can be costly.

# Packaging Software for Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Each time you use the software installation extension of Group Policy to deploy an application, patch, or upgrade package, you must first prepare the application for Windows Installer. If the application does not include a natively authored Windows Installer package (.msi), you must obtain one or create a .zap file.

Each Windows Installer package file contains a database that stores all the instructions and data that are required to manage the state of a program. For example, an .msi file of an application might contain instructions for installing the application when an earlier version of the application has been installed. The .msi file might also contain instructions for installing the software on a computer where that application has never been installed.

To package software, you might only have to perform an administrative installation to prepare the application for later installation by client users or computers from a network location. However, you might want to customize the application deployment so that you can create a transform that modifies the .msi. Also, if your application does not include a Windows Installer package, packaging can be much more complex. In that case, you must determine how long your organization expects to use the application. Your decision might lead you to reauthor the Setup.exe file of the application.

Figure 8.3 illustrates the process for packaging software for deployment by using Windows Installer and the software installation extension of Group Policy.

**Figure 8.3 Packaging Applications for Deployment**




**Using Native Windows Installer Packages**

Many software authors develop applications to include native Windows Installer packages. A native Windows Installer package contains a single product, such as Microsoft® Office 2000, which can be made up of several features, such as Microsoft® Word, Microsoft® Excel, and Microsoft® PowerPoint®. However, you can configure the software to install features individually. When the user selects an uninstalled feature, the feature is installed. Each feature (Word, Excel, PowerPoint, and so on) contains components, such as a thesaurus, spelling checker, or an additional user-interface language. When the user selects a feature or component that is not installed, the feature or component is automatically installed. Automatic installation of selected features saves network bandwidth during the initial installation. It also gives users only the features and components that they need to do their jobs. Automatic installation can also save space on users' local disk drives. However, this type of installation delays the availability of a feature when the user first selects it.

Windows Installer packages ensure that an accidentally deleted file, such as Winword.exe, is reinstalled the next time the user tries to start Word because Windows Installer detects and reinstalls missing files.

After you obtain a native Windows Installer package, perform one of the following tasks if the application you plan to deploy includes a native Windows Installer package:

- If the package requires no further customization, copy it to the designated shared folders on the software distribution point servers.

- If you want to modify the .msi package use transforms to customize it.

For information about copying packages to the software distribution servers, see "Deploying Applications in a Managed Environment" later in this chapter.

**Customizing Windows Installer Packages by Using Transforms**
A *transform* (.mst file) is a collection of specified changes that you apply to a base Windows Installer package file at the time of deployment. After you package the software in the Windows Installer package format, you can use transforms to customize the software for your organization. At this phase, the modular design of Windows Installer packages simplifies deployment. When you apply transforms to an .msi file, Windows Installer can dynamically add or modify data in the installation database to customize the installation of the application.

Office 2000, for example, provides the Office 2000 Custom Installation wizard, which you can use to build transforms. You can create a transform for managing the configuration of Office 2000 that you deploy to users in your organization. Other tools in the Windows Installer SDK, or other non-Microsoft tools, can also help you to create transforms for Windows Installer packages that do not include their own custom installation tools.

**Purposes for Transforms**
Transforms tailor the installation of an application. Although they are optional, you can use transforms for a variety of purposes including encapsulating and customizing.

**Customizing**
Customizing can involve configuring installations so that a particular set of features from a specified software application, or suite of software applications, is installed locally on the computer.

You can also use transforms to add new features to an existing application's package. For example, you can add custom Excel templates for the Finance group. However, if the templates change frequently, it might be a better practice to package them, and then assign or publish them as a separate package.

**Encapsulating**
You can encapsulate numerous customizations of a base package that are required by different groups. For example, in organizations where the Finance and Marketing departments require different installations of a product, you can make the product's base package available to everyone at one software distribution point. Then you can distribute the appropriate customizations separately as transforms to each group of users.

### Important

- Store transforms at the same software distribution point and in the same shared folders as the Windows Installer package that they customize or the installation will fail.

**Customizing a Native Windows Installer Package Example**
An organization maintains a large number of applications. Some of these applications contain native Windows Installer packages that require some customization, and many of the applications require modification to prepare them for Windows Installer. The Finance and Marketing departments of the organization require different installations of a custom-designed customer database application that includes a native Windows Installer package (.msi). The Finance department requires financial information about customer accounts, and the Marketing department requires customer purchase history to determine the best use for marketing resources.

The software administrator made the application's base package available to everyone from one software distribution point, and then distributed the appropriate customizations to each group of users separately as transforms. To achieve this, she did the following tasks:

1. Copied the original .msi to the designated software distribution point servers.

2. Created two separate transform files (.mst), each containing the customizations for the Finance and Marketing departments.

3. Copied the .mst files to the same shared folders on the software distribution point server where the original .msi files are located.

4. Created a <a ID="wPopUp" href="help=glossary.hlp TOPIC=gpo1">GPO </a>and assigned the application to the users in both the Finance and Marketing departments.

When users log off and log on again, the application is published on their desktops, and it is also listed under **Programs** on the **Start** menu.

You associate a transform with a Windows Installer package by using a GPO when you first deploy the package.

**To associate a transform with a Windows Installer package by using a GPO**

1. Select the software installation extension of the Group Policy object.
2. In the right pane of the Group Policy Object Editor, right-click the managed software.
3. Click **Properties**.
4. In the appropriate fields on the **Modifications** tab, specify the modifications or transforms that are associated with the package.

### Important

- Do not change the manufacturer's product code for the package when you create a transform to customize a native Windows Installer package. The software installation extension of Group Policy treats the new code as a new product. This can cause a failed installation and void your licensing agreement.

**Modifying a Transform**

To modify a transform by using the software installation extension of Group Policy, you can create an upgraded relationship between the application and the old transform, or between the application and the new transform. An appropriate time to do is this when you create a new transform or when you distribute new software that requires a transform.

**To modify a transform**

1. Create a new .mst file.
2. Remove the existing associated .mst files by using the software installation extension of the Group Policy Object Editor.
3. Associate the new .mst file by using the original .msi.

### Important

- Remember that a transform is applied at the time of assignment or publication, not at the time of installation. Verify that the **Modifications** tab of the package properties dialog box is properly configured *before* you click **OK**. If you do not do this, and you deploy an incorrectly transformed package, you must either remove the software and redeploy it, or upgrade the software to a correctly transformed version.

**After Modifying a Transform**

After you create or modify a transform, you must deploy the new transform.

**To deploy a new or modified transform**

1. Copy a new .msi file and a new .mst file to the designated shared folders on the software distribution point servers.
2. Delete the existing package object from the software installation extension of the Group Policy Object Editor.
3. Associate the new .mst file with the original .msi file.
4. Deploy the new package and the transform at the same time.

For information about copying software to the software distribution point servers, see "Deploying Applications in a Managed Environment" later in this chapter.

**Reauthoring Applications for Windows Installer**
When you reauthor an application, you create an application that adheres to the Windows Installer format. You are essentially redeveloping the setup portion of the application to take full advantage of the advanced capabilities of Windows Installer.

If you plan to reauthor an application that does not include a Windows Installer package, you must have the following:

- All executable files, dynamic-link library (DLL) files, and other resources. For all but the simplest applications, you need the source code to understand the logic and actions of the original setup program.
- An understanding of the application and the registry entries, shortcuts, and other information that are needed for it to run correctly.
- An authoring tool that supports creating Windows Installer packages.

There are some authoring tools available to help developers create new Windows Installer packages, but the procedures can be resource-intensive and costly. If you determine that the application will play a key role in the future of the company, it is important to weigh the benefits of reauthoring the application with the costs of reauthoring the application.

**Tools to Help You Create Windows Installer Packages**
Several tool vendors supply .msi package–authoring tools that developers can use to create Windows Installer packages. Such authoring tools include, but are not limited to:

- Microsoft® Visual Studio® Installer
- Microsoft® Visual Studio® .NET
- Commercial installers for Windows

**Tip**

- To develop any new or custom applications for your organization, use Windows Installer technology during the design phase. By designing the application with Windows Installer in mind from the outset, you can take full advantage of Windows Installer capabilities.

If you plan to reauthor an application to include a Windows Installer package, see the Windows Installer Software Developer's Kit (SDK) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources

For more information about the Microsoft® Visual Studio® Installer authoring tools for installing software, see the Microsoft Visual Studio Installer link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources

**After Reauthoring Applications for Windows Installer**
After you have reauthored an application to include a Windows Installer package, the next step is to perform one of the following tasks:

- If the package requires no further customization, copy it to the designated shared folders on the software distribution point servers.
- If you want to customize the .msi package for your users, use transforms (.mst files) to modify it.

For information about copying packages to the software distribution point servers, see "Deploying a Managed Software Environment" later in this chapter.

**Reauthoring an Application for Windows Installer Example**
Administrators have developed a simplified custom application so that users can arrange their own business travel. Because the application was developed internally, the organization has all the files for the software, and the developers understand how

the software must be installed. IT management determined that this application is ideal for reauthoring packages for Windows Installer.

To reauthor the business travel application, the administrator performed the following tasks:

1. Compared the benefits of having the application with the costs to reauthor it.
2. Identified the developers to perform the task of reauthoring the application.
3. Chose the appropriate reauthoring tool.
4. Reauthored the application to include a native Windows Installer package.
5. Placed the native .msi on the software distribution point.
6. Created a GPO and published the application for a particular group of users.

Users who need the business travel application can go to **Add or Remove Programs** in Control Panel, and then download the software from a published list of applications.

**Creating .zap Files**

Applications that do not use the .msi file format for the Windows Installer Service can be set up for distribution by creating a text file that has a .zap file extension. This method is not as flexible as .msi package files. If you have many applications that do not contain native Windows Installer packages, and you know that your organization plans to discontinue these applications, you can create software installation settings (.zap) files for the installation executable (such as Setup.exe or install.exe) files. Additionally, if you use custom applications that do not have Windows Installer support, but you plan to use them in the long term, .zap files might be your only choice. When you create .zap files, you do not benefit from the capabilities of Windows Installer. By creating .zap files, you wrap 32-bit or 64-bit Setup.exe files into a .zap file format that the software installation extension of Group Policy recognizes. This method allows you to publish the applications for users to install by using **Add or Remove Programs**.

Because these applications do not use Windows Installer setup programs, they do not do the following:

- Use elevated permissions for installation.
- Install a feature on the first use of the feature.
- Roll back an unsuccessful operation, such as install, modify, repair, or removal.
- Detect a broken state and automatically repair it.
- Implement customized installations (transforms).

**Suggestions for Working with .zap Files**

When you work with .zap files, consider the following:

- While applications that are installed by using .zap files run their original setup programs, they do not run with the elevated permissions that Windows Installer packages have. If installing the application requires administrative permissions, only users who have those permissions can install it.
- Because .zap files are typically created by using text editors, the files might have a .zap.txt file name extension. Make sure that the file name extension of a .zap file is only .zap (without the .txt extension). Also, make sure that any software installation file that you distribute by using GPOs does not end in .txt.
- If you have 64-bit clients, test 32-bit .zap applications to verify that you can install and run them on 64-bit clients. This is important because more .zap applications fail on 64-bit clients than on 32-bit clients.
- Unless you change the default behavior, 32-bit .zap applications are deployed so that they are not listed in **Add or Remove Programs** on 64-bit clients.

**After Creating .zap Files**
After you create a .zap file, copy it to the designated shared folders of the software distribution point servers. You can then use the software installation extension of Group Policy to publish the application in **Add or Remove Programs** in Control Panel. This makes the application readily available to users. You cannot customize .zap files by using transforms.

**Creating a .zap File Example**
The software administrators of the organization have identified a group of users who use an Excel template that is only compatible with Microsoft® Excel 97. Although all other users in the organization plan to upgrade to Microsoft® Office XP, this group of users must continue using Excel 97 until the end of this particular project. Excel 97 does not include a native .msi package file. After this project is completed, this organization has no further need for Excel 97.

To make sure that these users have Excel 97, the administrator performed the following tasks:

1. Wrapped the setup program of Excel 97 into a .zap file.

2. Copied it to the designated software distribution point servers.

3. Created a GPO and published the software to that particular group of users.

The users who need Excel 97 can now go to **Add or Remove Programs** in Control Panel, and then download the software from a published list of applications.

To publish Excel 97 by using the existing Excel 97 setup program, only the three path-information lines are needed in a .zap file. The following example represents the .zap file that the software administrator created for the purpose of deploying Excel 97 to the users on this project.

### Path information

```
<Application>
FriendlyName = " Microsoft Excel 97"
SetupCommand="" \\server\share\Excel 97\setup.exe""
```

The path and the name of the .exe file are enclosed in quotation marks in the Application section. If there are no command-line arguments, the .exe file path and name must be enclosed in two sets of quotation marks. For example:

### Absolute path

```
SetupCommand=" \\server\share\long folder\setup.exe"  /argument
SetupCommand="" \\server\share\long folder\setup.exe""
```

### Relative path

```
SetupCommand="setup.exe"  /argument
SetupCommand=""setup.exe""
```

### Note

- When you create your own .zap file, modify the information in the preceding sample .zap file according to the application that you are managing and the location of your software distribution point.

For more information about publishing .zap file-packaged applications, see "Targeting Software to Users and Computers" later in this chapter.

**Packaging 64-Bit Applications**
In the Windows Server 2003 family of servers, the software installation extension of Group Policy and Windows Installer 2.0 continue to support and protect the investment that you have made in 32-bit applications. Additionally, Microsoft® Windows® Server 2003, Enterprise Edition and Windows® Server 2003, Datacenter Edition introduce support for 64-bit application installation.

Windows Installer version 2.0 installs three types of Windows Installer packages on a computer running a 64-bit operating system:

- 32-bit packages that contain only 32-bit components

- 64-bit packages containing some 64-bit components and some 32-bit components
- 64-bit packages containing only 64-bit components

**Note**

- You cannot publish or install a 64-bit application on a 32-bit operating system.

You can use the Software Installation extension of Group Policy to allow or disallow installation of 32-bit applications to 64-bit clients.

**To allow or disallow installation of 32-bit applications on 64-bit clients**

1. Open the Group Policy Object Editor to the Software Installation item.
2. Right-click the managed software, and then select **Properties**.
3. On the **Deployment** tab, click **Advanced**.

You might decide not to deploy a 32-bit package to a 64-bit system if that application functions poorly or not at all on 64-bit systems. Make your configuration determinations as you test the applications and system combinations in your test lab and during your pilot phase. For more information about piloting the deployment of applications, see "Conducting a Pilot for Software Deployment" later in this chapter. For guidelines about testing the applications in your organization, see "Planning and Testing for Application Deployment" in *Planning, Testing, and Piloting Deployment Projects* in this kit.

**Important**

- Do not deploy a 32-bit application that uses the same Windows Installer 2.0 Product ID as a 64-bit application. Instead of creating such Windows Installer 2.0 packages, create a separate Product ID for the 64-bit version of a product. Windows Installer does not support different architecture packages that have the same product code.

If you incorrectly configure a 64-bit package for 32-bit clients (or use the wrong Product ID), Group Policy tries to install it, but then removes the package at logon. It then installs the package again the next time a user logs on, and then removes it again, creating network traffic and preventing users from using their computers.

**Repackaging Applications for Windows Installer**

When you cannot reauthor a package to use Windows Installer, you might want to repackage it. Repackaging an application for Windows Installer involves taking a snapshot of a clean computer (including the registry settings, files, and system settings), installing the software, and then taking a post-installation snapshot of the computer. The repackaging software detects the difference between the two snapshots, and then creates the necessary instructions to reproduce the installation. If any registry changes, files changes, or system setting changes occur during the capture process, they are included in the installation. You use repackaging when you do not have control over DLL files, source code, and registry entries, or for applications about which you do not have in-depth knowledge.

Use this method only as a last resort when you need to repackage an application into an .msi. It is easy to underestimate the cost of repackaging in terms of labor hours. Also, users often set their expectations too high for the reliability of repackaged applications. Repackaging requires a thorough knowledge of the application's installation program and of the Windows Installer setup on the Windows platform.

Success with repackaging is affected by the state of the computer where you perform the repackaging. For best results, always perform a repackaging by using a clean computer. For the purpose of repackaging, a clean computer is defined as a computer that has only the operating system and operating system service packs installed before you run the repackaging software. Because of this limitation, and other issues, repackaging is not recommended.

Repackaging is not a function or a feature of Windows Installer. As with reauthoring applications, several vendors provide tools to enable administrators to repackage applications for a variety of needs. The same vendors who provide tools to reauthor applications can also help you repackage them.

When you repackage an application, you replace the existing components, such as DLLs, .ini files, registry settings, and

shortcuts, and then you create a path for Windows Installer to find these items at installation time.

Packages that you create on a computer running Windows 2000 by using Veritas WinInstall LE, will work on target computers running Windows XP. However, the Veritas WinInstall LE program itself does not run on Windows XP. For more information about repackaging applications on newer computers, contact your repackaging-software vendor.

**Important**

- Microsoft supports authoring and customizing applications that natively use Windows Installer for installation and maintenance. However, Microsoft does not provide support for applications that are repackaged as .msi files. Contact the application manufacturer for this type of support.

- As with other repackaging techniques, repackaging applications into the Windows Installer format has limitations and might not be supported by the application manufacturer. For more information, contact each application manufacturer.

If you have an application that does not include a native Windows Installer package, and you decide to repackage it in the .msi format, see the Windows Installer SDK for detailed documentation about the Windows Installer package format. For more information about the authoring tools for deploying a managed software environment, see the product links, as listed by product name, on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

After you repackage an application for Windows Installer, copy it to the designated shared folders of the software distribution point servers. You cannot customize repackaged applications by using transforms.

# Deploying Applications in a Managed Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you prepare the application for deployment, you can deploy the software by using Group Policy or another software distribution technology. Each time you deploy applications, patches, or updates in your test environment, and then to the rest of the organization, you must prepare the software for Windows Installer. You are then ready to perform the tasks as listed in Figure 8.4.

**Figure 8.4 Deploying Applications**




Deploying applications is the process of setting up and managing distribution points or server shares where users have access to the packaged software and can install it on their computers. Before you can begin distributing applications, you must clearly understand the current arrangement of servers and client computers on your network, the software requirements of users, and their locations on the network. Also, plan for who will install and manage the software.

Microsoft Distributed File System (DFS) and File Replication service (FRS) can enhance the security and availability of your distribution points. For more information about DFS, FRS, and Microsoft file-server technologies, and how to best deploy these technologies, see "Designing and Deploying File Servers" in *Planning Server Deployments* of this kit. This is an important topic to understand before you deploy your Group Policy–based software distribution solution. Read that chapter before you configure your distribution point servers.

Whether you decide to use DFS, FRS, or another method to configure your distribution point servers, you can distribute software to users by using the software installation extension of Group Policy.

Network traffic can become excessive when clients simultaneously download software. To prepare for this load, perform the following tasks before you place software distribution point servers on your network:

- Identify the users' software requirements and where these users are located.
- Examine the network infrastructure:
    - Measure network capacity and bandwidth.
    - Identify the slow or intermittently connected network links.
    - Identify performance issues between clients and software distribution point servers.
- Plan for securing your software distribution point servers.

- Evaluate using DFS and FRS as supporting technologies for Group Policy-based software distribution.

#### Identifying Software Requirements and Locations of Users

When you know what software users require and where those users are located on the network, place the software distribution point servers close to those users (on the same subnet, for example). If you have users who connect over slow links, consider the impact to the network performance if you allow them to download software over those slow links.

**Note**

- If you are deploying the same application to clients at two locations, deploy it by using two targeted GPOs instead of one. One of the GPOs applies to all clients at one location and uses a UNC path to a source point at that location. The other GPO applies to all clients at the other location and uses a UNC path to a source point at that location. GPOs can be linked to Active Directory sites, which can be a valuable option in this context.

#### Examining the Network Infrastructure

When you are familiar with the network infrastructure in your organization, you can determine which method for managing software distribution point servers is best for your organization. If you use DFS and FRS, you must understand your network infrastructure to design your DFS namespace, replication topology, and replication schedule. For more information about using DFS and FRS, see "Designing and Deploying File Servers" in *Planning Server Deployments*.

#### Measuring network capacity and bandwidth

You can measure network capacity by determining the number of connections that the server establishes and maintains. Available bandwidth varies widely, depending on the transmission capacity of the link, the server configuration, the server workload, and competing traffic. The capacity of a single server changes during operation in response to demand and competition for shared network resources.

With the fluctuation of server capacity in mind, you must place software distribution point servers as close as possible to the users who need to access them (preferably on local networks). If you implement DFS, you have three options for configuring DFS to select targets: default, restricted, or least expensive.

**Default target selection** If the last (or only), target in an Active Directory site fails or is taken offline, DFS directs clients to another target in the same site, if a target is available. Clients are directed to a random target if no same-site targets are available. DFS does not consider bandwidth cost or speed when choosing the random target.

**Restricted same-site target selection** You can limit client access to only those targets that are in the same site as the client. When you use this option, plan to have at least one target (or two targets, for fault tolerance) in every site. If no same-site targets exist, clients in that site are denied access to the data in the namespace.

**Least-expensive target selection** You can make DFS use an alternate target that is based on connection cost if no same-site targets are available. Windows Server 2003 uses the site and cost information in Active Directory to determine whether sites are linked by inexpensive high-speed links, or by expensive WAN links.

To avoid installation failure, test the replication of application components (the executable files, the Windows Installer packages, and any transforms) from the location of the administrator who is initiating the deployment to all the software distribution points near intended recipients. If an administrator plans to work remotely, be sure to test this configuration. For more information about deploying file servers, see "Designing and Deploying File Servers" in *Planning Server Deployments*.

**Important**

- To ensure reliable transfer and usable connection speed, place frequently accessed data in the same sites as the primary users of the information. You can increase availability by placing multiple servers in each site, as needed.

#### Identifying slow-link connections

Many remote and mobile users connect to the network by using slow-link connections. The software installation extension of Group Policy can move large amounts of data, so processing across a slow link can affect performance significantly. By default, Group Policy–based software deployment does not operate over slow network or dial-up connections. However, you can allow users to install published applications over a slow link if absolutely necessary. However, you cannot configure Group Policy for automatic assignment, upgrade, or removal of applications over slow links.

**Note**

- A remote access connection is not necessarily a slow link, nor is a local area network (LAN) necessarily a fast link. By

default, the fast or slow status of a link is based on a test ping to the server. If it takes less than 2000 milliseconds (two seconds), then it is considered a fast link. If it takes longer than 2000 milliseconds it is considered a slow link. You can set this value by using the Group Policy setting, **Group Policy slow link detection**.

#### Evaluating Strategies for Connecting Remote Users

If any users connect primarily over slow links, it is essential that you define an appropriate software deployment strategy for them. While keeping the application management objective in the forefront, you can address this challenge by various means:

**Terminal Services** Run the applications on a server running Terminal Services. This is currently the best method for providing remote computers with access to Windows–based applications that are deployed by using Group Policy–based software installation. The terminal server becomes a single access point that allows multiple users to have access to a desktop where they can download applications, run programs, save files, and use network resources without an administrator having to manage multiple, remote desktops.

When using Terminal Services, administrators install applications on a per-computer basis, meaning that applications are available to any user who has access to the terminal server. Users can gain access to a terminal server over any Transmission Control Protocol/Internet Protocol (TCP/IP) connection including Remote Access, Ethernet, the Internet, wireless, wide area network (WAN), or virtual private network (VPN). The user experience is limited only by the characteristics of the weakest link in the connection. The TCP/IP deployment in the data center governs the security of the link.

For more information about using Terminal Services to deploy applications, see "Hosting Applications with Terminal Server" in *Planning Server Deployments* of this kit.

**Direct Installation** Have remote and mobile user bring their portable computers to the office so that you can install the software. If you use this method, use the **Install this application at logon** option to install software in its entirety automatically the next time the user logs on to the network. For more information about this option see "Assigning and Publishing Software" later in this chapter. This option will not interfere if mobile users manually install the needed software whenever they have access to a fast link.

**Add or Remove Programs** Educate remote users to install applications periodically from **Add or Remove Programs** in Control Panel. Any assigned application behaves like a published application when users gain access to it over a slow link. Having remote users installing applications might not be to the best method because of the time and bandwidth required, but it lets users decide when to begin the lengthy installations. Keep in mind that, if many users share a slow link, they might block the connection. However, if only a few users share a slow link, this method might provide a practical solution.

**SMS** Use SMS to deploy software to users who can connect only over slow network links. For example, SMS supports a CD-ROM distribution method by which you can install a distribution package on a CD-ROM, and then send it to mobile or remote users. For more information about providing connectivity for remote and mobile users by using SMS, see the Microsoft Systems Management Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Do not use WebDAV to publish software on a Web server. This method can allow an unauthorized user to run a script under the security context of the WebDAV thread, and gain access to your network. See Microsoft Security Bulletin MS01-022 for details.

#### Identifying performance issues

To understand the performance issues between the client computers and the software distribution point servers in your organization, perform these tasks (not necessarily in this order) in your test lab:

- Determine if the installation of the software is assigned (required) or published (optional), and if the assignment is to users or computers. For example, assigning an application to a large number of users or computers has a heavy impact on the network when a high number of users log on the same day and automatic, simultaneous software downloads start.

- If you publish the application, users can download software as-needed. In effect, publishing staggers the downloading of the software, which decreases congestion on the network.

- Determine how many settings you can include in a GPO and still maintain the balance between settings and logon times. Having many GPOs often increases logon times; logon times are especially affected by the number of settings in those GPOs. Having fewer GPOs is easier to manage but limits flexibility. Testing these configurations in the lab can help you to determine the optimum balance for your organization. It is recommended that you use GPMC to test, stage, and deploy your GPOs.

- Consider the size of each software application that you plan to distribute. Obviously, the larger the application, the larger the impact on the network, especially if you assign the application to a large group of users or computers at one time.

- Consider the placement of the software distribution point server in relation to the targeted users: How long does it take for a particular package to go from a software distribution point server to a client? Test the most common client and server access circumstances, the least common client and server access circumstances, and worst-case situations.

- Plan for the number and frequency of the software deployments. These factors affect the amount of server disk space that is required and the subsequent network load. For example, when many users install a small package simultaneously, a large amount of data suddenly moves across the network. This can have a greater impact on your network than a large package that is rarely installed.

- Consider how many users access a package at one time. In this instance, determine if you will use the **Install this application at logon** option (located in the software installation extension of Group Policy), or if you will use the default option and allow the users to install the application components on an as-needed basis. The following are two instances where large numbers of users might access a package at the same time:

    - **New software package placed on the software distribution point.** In this case, expect a large number of initial installations whether you publish or assign the software. Stagger the installation so that all users do not download packages at the same time on the same day. For example, you can break up the users into small groups and assign a certain application to 500 users one week, and then assign the same application to 500 different users the following week. This minimizes the effect that software installation has on your network bandwidth.

    - **Many new employees starting on the same day.** If this occurs, consider the network overhead for these users to log on and receive assigned packages. Also, consider the number of published packages that these new employees must gain access to over a specified time. After you determine the results in your situation, make decisions that protect your servers from becoming overloaded by providing sufficient software distribution points that are part of a DFS namespace. Note that SMS uses scheduling to avoid this problem.

Test potential issues in a lab to determine what might happen on the network if a specified number of users try to gain access to a specified software distribution point simultaneously to install a 10 megabyte (MB) application, a 50 MB application, a 100 MB application, or a larger application.

After you have determined your network capacity, you have identified slow link connections, and you have identified and resolved performance issues between client computers and software distribution point servers, you must configure the software distribution point servers.

# Determining Requirements for Software Distribution Point Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The method that works better for deploying and managing software distribution point servers depends on the objectives of your organization. You can use either of the following methods:

- Set up a universal naming convention (UNC) path to a server share.
- Use DFS.

**Setting up a UNC Path to a Server Share**
By using the UNC names, a user or application can specify the physical server and share names to gain access to file information. For example: *Server**Share**Path**File_name*. You can use a UNC path to allow direct access to a shared file by mapping to a network drive, where the drive letter denotes *Server**Share*. You can also perform a deep **net use** to navigate beyond the redirected drive.

However, as networks grow and as organizations begin using existing storage for new purposes, mapping a single drive letter to individual shares becomes inefficient. Also, despite the fact that users and applications can refer to UNC names directly, the increasing number of places users must go to retrieve data can be overwhelming.

**Using DFS to Manage Your Software Distribution Point Servers**
DFS provides fault tolerance for your software distribution points by mapping a given logical name to shared folders on multiple file servers. This way, software remains available for installation, regardless of whether one of the physical servers where the software deployment files reside becomes unavailable. DFS also improves storage scalability because you can deploy additional or higher-performance file servers and present the storage on the new computers as new directories in an existing namespace.

When you use DFS in combination with Group Policy–based software deployment, you benefit from its location independence and load-sharing abilities. These features simplify management and optimize the installation for users. Instead of allowing all users to install software from a single server, and taxing the server, you can design a DFS namespace to distribute network traffic across multiple servers.

For more information about DFS, EFS, and creating a replication topology and schedule, see "Designing and Deploying File Servers" in *Planning Server Deployments*.

# Configuring Software Distribution Point Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you place the software distribution point servers, you must configure each server for the software applications that are accessible to users. Configurations include granting the appropriate permissions so that users (for user-assigned applications) and computers (for computer-assigned applications) can gain access to the software.

After you complete the deployment, you cannot change the path to another software distribution point without redeploying the software package. Redeploying software can be quite intrusive to users because it automatically removes the application on computers where the software is already installed. If the software is assigned, it is reinstalled automatically; if it is published, it must be manually selected again for installation by using **Addor Remove Programs** or by opening an appropriate file type.

**Important**

- When you create a DFS software distribution point for a new application, you must specify a different DFS path, regardless of whether the target software is on the same share. You can then retire software distribution points for the old software by taking down targets of the DFS tree while all the new software distribution points of the new application remain available.

To set up a software distribution point server, copy the software and your Windows Installer packages to the designated servers. Perform an administrative install of the software before you distribute it to clients and servers. To perform an administrative install, run the application's Setup.exe from a command line against the application's CD-ROM. Set the **/a** option to the relevant share in the form of *Server Name**Share Name*. This unpacks all the files from the .cab files, and then drops all the packaged files on the target share. You do not have to log on to the software distribution point server. However, any user who performs an administrative install must have administrative access to the destination share.

Typically, you copy the Windows Installer packages to an intermediate location and then run **Setup /a** from there to the destination that you want. Alternatively, you can run **Setup /a** from and to the software distribution share itself. The drawback of this method is that it unpacks the .cab files, and then leaves them on the same share as the Windows Installer package. This method uses space on the share. However, many administrators do this intentionally so that the .cab files are always available from the destination share. This is a matter of preference, or ease of use, and does not affect performance.

When you set up a software distribution point server, be sure to maintain the security of your network by performing the following tasks:

1. Restrict physical access to the server to the administrators who must physically touch the server.

2. Consider using digital signing. For more information about digital certificates, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

3. Create the folders for the software on the server that you designate as the software distribution point, and then make the folders network shares. For example: *Server_Name**Share_Name*

4. Copy the Windows Installer packages, application executable files, .mst files, .msp files and .zap files (if used) to the appropriate shared folders.

5. Set appropriate permissions on the folders.

If you allow high-level permissions to a very broad group, you increase the chance that your servers might be tampered with. Software distribution point servers can become security hazards in your organization if they are loaded with harmful or damaged software packages. Allow users to read from the software distribution point. However, only grant Full Control to designated software deployment administrators.

Set the following discretionary access control list (DACL) permissions:

- **Everyone:** Read

- **Administrator:** Full Control, Change, and Read. Grant these permissions only to software deployment administrators

who must add or modify installation packages on the server. Also, limit access permissions to only those permissions that the administrators need to do their jobs. Note that administrators who have edit rights on a GPO effectively acquire local administrator rights on all computers that receive that GPO to install and remove software, although they do not become local administrators.

**Note**

- You must have software licenses for software that is written by independent software vendors and distributed by software distribution points. It is your responsibility to limit the number of users who gain access to software through software distribution points to the number of licenses that you own. It is also your responsibility to verify that you are working within the licensing requirements that are provided by each independent software manufacturer.

Whether you install applications by using SMS or Group Policy, it is important to understand the implications of software license agreements. Some applications are licensed per seat, and others are licensed per user or per site. The license agreement for an application might influence your choice of installation method and your decision about whether to have Group Policy assign the application to the computer, or whether to assign or publish the application to the user.

**Targeting Software**

After you distribute prepared applications to the appropriate software distribution point servers, target each application to a list of intended recipients. Recipients can be specified users, groups of users (such as all members of a certain defined OU), computers, or any combination of these.

Typically, administrators define the hierarchical structure of the sites, domains, and OUs for Active Directory. They do this based on an assessment of their organization's business objectives, administrative requirements, and user needs. When you perform your assessment to target software to users and computers, you narrow your scope to the software-based needs of all users in the organization and the administrative requirements for managing the software. It is essential that you consider the Active Directory structure when determining your software deployment strategy.

It is also important that you evaluate each application that you plan to deploy, from the perspective of both the users who need the application and the administrators who will manage it.

**To assess the administrative requirements and business needs of the users in your organization**

1. Identify what software is mandatory for your entire organization and what software is appropriate for a particular job or business unit.

2. Identify the users who must have specific language versions of the software.

3. Determine whether to assign the software to users or to computers, or to publish for users. If you assign software to users, decide whether you want the software to install fully after computer startup, or if the on-demand approach works better for your users and your network infrastructure requirements. See "Assigning and Publishing Software" later in this chapter.

4. Identify remote software installation requirements.

5. Determine whether you must customize the software for different parts of the company.

6. Determine how often you will have to update user templates or custom files.

There are other requirements that determine targeting specifications, such as users who have existing versions of a package or minimum hardware requirements. Special considerations for targeting include the dynamic nature of any organization and the resulting addition and removal of resources from a target set over time.

To help you assess the administrative and user requirements in your organization, see "Planning a Managed Environment" in this book. To learn how to create and then link a GPO to a site, to a domain, or to an OU, see "Designing a Group Policy Infrastructure" in this book.

# Targeting Software to Users and Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin the process of targeting software to users and computers, you must know how to create a GPO and how to link it to a site, to a domain, or to an OU.

Group Policy is an extremely flexible tool that you can use to provide users with controlled access to the specific software that they need to do their jobs. You can use Group Policy to control various levels of software administration for users and computers in a specified site, domain, or OU.

To create GPOs for software deployment, you must find commonality across applications, users, and administration. In a simple situation, a set of users who are members of the same OU must have access to a specified set of applications that are administered by a single group of administrators. In this situation, you can put the applications in a single GPO, link that GPO to the applicable OU, and then assign or publish the applications for the group of users.

**Determining the Number of Software Installation GPOs**
To minimize the number of GPOs that you must have to install the software applications in the environment, it is recommended that you set the GPOs to publish a number of related applications. For example, if you determine that when users install Microsoft® Project, they typically also must have Microsoft® Visio® drawing and diagramming software. Therefore, it makes sense to combine these two applications in a single GPO. By logically grouping applications, you can reduce overall administration and management costs.

An alternative to grouping applications is to create a separate GPO for each application you want to manage by using Group Policy. By creating separate GPOs, you have more flexibility because the application of GPOs is more specific. However, the number of GPOs that are assigned to each client is higher. More GPOs leads to slower start and logon and times. However, the number of settings in each GPO is the most important factor because processing each .adm file requires a set amount of time, regardless of the number of settings that the file contains.

In small organizations, it is not uncommon to see more than 100 different applications. In large multinational organizations, the number can be significantly higher. It can be difficult and time-consuming to manage all these applications by using a separate GPO for each application. However, if it ever becomes necessary to delete or disable an application installation Group Policy, deploying the applications in separate GPOs can prevent massive removals and reinstallations of the applications if the applications fall out of management scope (for example, if the client computer or user is moved to another organizational unit). Administrators must balance the benefits of this approach from the point of view of both administrators and users.

**Recommendations for Creating GPOs for Software Deployment**
When you create GPOs for software deployment, it is recommended that you do the following:

- Create separate GPOs for administrators, based on their roles, to simplify administrative delegations. However, these administrators must communicate with each other about the applications that they are deploying to avoid affecting each other's installations and creating unnecessary network traffic.
- Use GPOs to match applications to users and the tasks they perform.
- Assign or publish a package only once per GPO. For example, assign Office XP to the computers that are affected by a GPO, but do not assign or publish it to users who are affected by the same GPO.
- Minimize the number of GPOs that you create because the overhead for processing each GPO increases user logon time. Network performance is much better if you use a single GPO that deploys 100 applications than if you use 100 targeted GPOs. Take advantage of GPO inheritance to distribute the application throughout your organization. If necessary, use security descriptors, such as Access Control Entries on the GPO for finer control over who receives the software.

**Note**

- The depth of OUs has very little effect on client processing time per GPO. However, a deep OU hierarchy can significantly complicate the troubleshooting process. For more information about inheritance and security descriptors, see "Designing a Group Policy Infrastructure" in this book.

- Place GPOs that assign or publish software applications as high as possible in the organizational unit structure. This makes it easy to give all users or computers in an organization access to an application. It also reduces administration because you can use a single GPO to assign or publish an application, instead of re-creating that object in multiple containers that are lower in Active Directory.

- Plan ahead for removal when you initially deploy the software. If you want the application to be removed when a GPO is no longer applicable, select the **Uninstall this application when it falls out of the scope of management** option. You can configure this option on the **Deployment** tab in the **Properties** dialog box. Right-click the managed software in Group Policy, and then click **Properties**.

- Specify application categories. When you use categories, it is easier for users to find applications in **Add or Remove Programs**. You can configure this option on the **Categories** tab in the **Properties** dialog box. Right-click the managed software in Group Policy, and then click **Properties**.

**Important**

- Security Group filtering is not recommended because the final effect of filtering is difficult to predict before deployment or to analyze afterward. Although using this method might seem simple at first — one GPO can be used for everyone, with certain settings excepted for certain groups — it is difficult in practice.

**The impact of advertisement scripts on startup and logon time** Evaluate the size of each advertisement script of every application that you assign. The size of the script can vary dramatically, from 1 kilobyte (KB) to 200 KB (typical) to 700 KB (very large application suites). Windows XP Professional advertises all assigned applications each time a user logs on. This does not generate network traffic, but it does require the client computer to process Windows Installer each time a user logs on. The first time that applications are assigned, the target must load and process the advertisement scripts for each new application. After this, each time the target user logs on, Windows Installer reprocesses the downloaded scripts for every assigned application to verify that these applications are properly configured. Typically, published applications do not affect user logon time. Therefore, it is not a problem to publish applications to a large number of users, regardless of whether they need the applications right away. The exception is when you publish an upgrade to an application that is already published or assigned. The operating system retrieves this publication data from Active Directory during the logon process. The effect this has on users is not as significant as the effect of assigning nonessential applications. It is best to assign only necessary software.

You can see the size of application management files by opening the applications folder or the scripts folder in Sysvol. These folders are located in the user folder or the computer folder of your domain controller. For example:

*Server Name*\Sysvol*Domain Name*\Policies*GPO GUID*\User\applications

-or-

*Server Name*\Sysvol*Domain Name*\Policies*GPO GUID*\Computer\Scripts

**Caution**

- Never edit or open the files in these Sysvol folders. Doing so can damage the application scripts and affect every client in the GPO. Always perform all necessary edits by using Group Policy Object Editor.

**Scaling Group Policy to Meet Your Needs**

The way you use Group Policy to manage software depends on your network infrastructure, administrative hierarchy, and the business requirements of your organization. Many large organizations have a central IT administration that has delegated authorization. Others have either strictly centralized or decentralized IT groups. Regardless of how your IT administration is organized, you must communicate with software administrators at all levels of the domain hierarchy so that all software administrators know what software each other assigns, publishes, or removes at each level. The administrative tasks that other software administrators perform at other levels of the domain hierarchy can affect how you scale Group Policy to deploy and manage software at your level.

**Applying Group Policy high in the domain hierarchy**

When you apply Group Policy settings high in the domain structure, you can reach more users than if you apply settings at a lower OU level. You also decrease application administration for the same application. You can filter user groups lower in the Active Directory hierarchy to avoid giving all users all the applications that you deploy by using Group Policy.

If an application is assigned at the domain level, all users in the domain receive that application. Administrators at the domain

level manage applications that are deployed at that level. By managing applications at the domain level, you relieve administrators in the OU levels beneath it from having to manage them.

Figure 8.5 illustrates how you can assign software to users and computers by applying GPOs at the domain level, and how that software assignment is inherited by the associated OUs by default.

**Figure 8.5 Assigning Application-Based GPOs at the Domain Level**




**Important**

- Administrators cannot control installation order. Although the Group Policy user interface displays the list of applications alphabetically, installation order is based on the order in which the applications are added to the GPO. However, you can create launch conditions in the .msi package so that a particular software package is not installed until a condition is met. For example, you might specify that a dependent package is not installed unless the application it requires to operate correctly is installed.

#### Targeting Software to Multinational Users

While you are planning your Group Policy deployment to meet the needs of your users, note the needs of multinational users. Multinational corporations share information and collaborate on a worldwide basis. Before Windows 2000, the operating-system support and the files generated with an application in one language were not necessarily compatible with a file that was created by the same application in a different language. However, in Windows Server 2003, you no longer have to apply, track, and maintain a different service pack and tool and application set for each localized version. You can now assign or publish the same software to every user or computer in the company, regardless of location.

Table 8.4 shows the conditions that the software installation extension of Group Policy must verify to determine whether to install an application on a computer that runs Windows 2000 or a later version of the operating system.

**Table 8.4 Application Installation Conditions to Verify**

| Condition | State of the Package |
|---|---|
| The language for non-Unicode applications is any value and the product language is either neutral or English. | The package is either installed or advertised. |
| The **Ignore Language** parameter is set in the software installation extension of Group Policy for the managed package. | The package is either advertised or installed, whether or not the system locale and product language match. |

The **Ignore Language** parameter can be configured by clicking the **Advanced** button on the **Deployment** tab, which you can open by right-clicking the managed software in Group Policy, and then selecting **Properties**.

**Caution**

- If you deploy two versions of the same application that have different product languages (such as an English version and a German version), and they are in the same GPO, be sure to give the application a different product code for each language.

# Assigning and Publishing Software

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because you can publish software for users, assign software to users, or assign software to computers, you can establish a workable combination of those three options to meet your software management goals. The following is a comparison of these methods.

**Publishing software for users**
Typically, after you publish a software package to users in a site, domain, or OU, the users can use **Add or Remove Programs** to install the software. An exception is when you publish an application in a new GPO, and you must simultaneously link the GPO to the users in a site, domain, or OU. If you link a GPO and deploy the software at the same time, you must refresh the Group Policy before the application appears in **Add or Remove Programs**. Additionally, the application can be installed by opening an associated document if the application is deployed to do that (if **Auto-Install** is selected).

The user can remove the software, and then later choose to reinstall it, by using **Add or Remove Programs**.

**Assigning software to users**
There are three methods for assigning software: assign to users on-demand, assign to users, or assign to computers.

**Important**

- Check software license agreements before you assign applications to users because assigning software can result in an application being installed on multiple computers. Issues might occur, regardless of whether you use the policy setting option **Remove the application if it falls out of the scope of management.**

**Assigning software to be available on demand** After you assign a software package to users in a site, domain, or OU, the software is advertised on the desktop. The application becomes available to the user the next time the user logs on (if application's GPO applies to that user). The application is fully installed by the user from the **Start** menu, from **Add or Remove Programs**, from a desktop shortcut, or by opening a document (on demand) that has a file name extension that is associated with the application.

The user can remove the software, and then later choose to reinstall it as they did previously. By using Group Policy, you make sure that assigned applications that are available on-demand are available, regardless of whether users remove them, and that the applications are available again the next time the user logs on or starts the computer.

**Assigning software to users** After you assign a software package to users in a site, domain, or OU, you can use the **Install this application at logon** option to install the whole application the next time the computer starts, or after the user logs off and then logs on again. The application is also immediately available in **Add or Remove Programs**.

The user can remove the software, and then later choose to reinstall it as they did previously.

**Note**

- Some applications that you have published might not appear in **Add or Remove Programs** in a domain that has multiple domain controllers until the changes have replicated to all domain controllers in the domain.

**Assigning software to computers** After you assign a software package to computers in a site, domain, or OU, the software is installed the next time the computer restarts or the user logs on.

Only the local or network administrator can remove the software, though a user can repair the software.

**Caution**

- To avoid installation errors and reduce network traffic, do not assign or publish a Windows Installer package more than once in the same GPO.

**Assigning Software to Users and Computers**
Assign software to users or computers for either of the following reasons:

- To make a particular application available to all users of one computer, assign that application to the computer.

- To make mission-critical software available to users or computers at all times, assign the application to the users or computers.

### Note

- If you assign many applications instead of publishing them, you can cause congestion between client computers and the software distribution point servers. Use DFS to distribute the server load among multiple servers.

**Assigning standard software**

Typically, packages that you assign to users or computers are essential. Therefore, the applications on your standard software list are good candidates for assignment to users or computers. The easiest method for assigning standard software to a large number of users in your organization is to apply the GPO at the highest level of the domain hierarchy, as shown in the following example.

**Assigning Software to Computers and Users Example**

It is standard for all users at a corporation to receive virus-protection software and e-mail. The software administrator creates two GPOs and assigns the two software packages. She assigns the virus protection software to all computers, and the e-mail application to all users. After the GPOs are created, she applies the GPOs at the domain level of the Active Directory structure so that all members in the domain receive the software assignments.

It is recommended that you assign virus-scanning software to all computers in the organization because this software must function for every user of each computer. Some organizations consider e-mail application to be mission-critical, but some e-mail packages are very large. Installing large packages over an already congested or slow network link can negatively affect network bandwidth. If this is not an issue for you, and you want all users to have e-mail, you can assign the e-mail package to everyone in the organization. After you configure software assignment in the appropriate GPO, apply the GPO that is associated with standard software applications to the root domain.

**To prepare for assignment of software to users and computers**

1. Determine the size of each application. Very large applications might not be appropriate for automatic installation. For example, a product such as Office 2000 can take a long time to install. Make sure that your deployment plan includes an analysis of how much traffic your network can handle.

2. Assign applications only to the users who require them.

3. Determine whether you can include some of the common applications in a Remote Installation Protocol preparation (RIPrep) image, or other automated image technology. RIPrep can reduce software installation time during the logon process or at initial selection of the application. RIPrep and Remote Installation Services (RIS) are excellent methods for creating these images. For more information about RIPrep and RIS, see "Designing RIS Installations" in *Automating and Customizing Installations* of this kit.

**Configuring a complete application Installation**

When you assign an application to a user, you have the option to install the whole application the first time the user logs on after deployment, or you can configure to install the application on demand. You can select the **auto-install by file activation** option on the **Deployment** tab in the **Properties** dialog box. Right-click the managed software in Group Policy, and then click **Properties**.

You can configure the installation to occur the first time the user logs on after deployment. This method ensures that the user has the whole application available when it is needed. However, this method also requires a longer logon time while the application is being installed. Without this method, a portable-computer user who is not connected to the network might discover that an essential feature is not available. By using this method, you provide a less confusing experience for users who might think that an application is installed, only to find that clicking the shortcut triggers an installation.

For applications that you can customize, such as Office XP, you can make all components of the application available at installation. This approach to installation is a Windows Installer authoring function, not a Group Policy software deployment function. For example, the author of the installation package can select to make features such as a spelling checker available on first installation. This increases the installation time somewhat, but it also provides all needed features on first use. Performing a complete application installation is a good method to use for mobile users who are not connected to the network most of the

time. When the user requires the spelling checker, it is already installed on the computer.

**Note**

- By default, Group Policy allows you to configure a user-assigned application that has a staggered, on-demand installation. By using Windows Server 2003, you can turn off the default installation and install the entire application at once. This mirrors the behavior of computer-assigned application installation.

**Enabling users to install applications and features on demand**

When you configure Group Policy so that users can install only the features (such as the spelling checker) or components of a product as they use them, you avoid wasting client disk space to store features that users do not need or use. Additionally, this method helps to prevent network congestion that is caused by users downloading large applications. The core application is not installed until the user activates the application on the computer by one of three ways: selecting the application from the **Start** menu, clicking s shortcut on the **Desktop**, or by activating a document of a file type that is associated with the application. After the core application is installed, the user can install features of the product as needed.

The following installation process is typical for user-assigned applications intended for on-demand applications:

1. The user logs on to a computer running Windows 2000 (or a later version of the operating system).
2. The application management service process advertises applications on the user's desktop or on the **Start** menu.
3. The user invokes the needed software from either the **Desktop** or the **Start** menu, or by selecting a file that has a file name extension for an assigned application. This action starts Windows Installer.
4. Windows Installer installs the requested Windows Installer package from the distribution point.
5. Windows Explorer starts the application.

**Publishing Software for Users**

The benefit of publishing software, instead of assigning it, is that it requires less management when change occurs in the Active Directory structure. Typically, you publish applications that are nonessential for the users. When you publish software for a user, it does not initially appear to be installed on the computer. There is no Windows Installer advertisement information about the software on the computer in the registry, on the desktop, or on **Start** menu as a shortcut. On an as-needed basis, the user installs the published software by using **Add or Remove Programs** in **Control Panel**. Users can also install the published application by selecting a file that has a file name extension for an application.

**To publish software for users in your organization**

1. Determine the size of each application. Some products take a long time to install, so consider if it is more appropriate to assign the application, instead.
2. Determine whether you can publish certain applications to all users (without restriction) in your administrative area.
3. Create a table of applications. This table includes the locations from which users can install the application files.
4. Publish all .zap files. You cannot assign .zap files.

When you publish applications, users do not need to remember server share names or locations for installing software. In Windows XP and Windows Server 2003, when a user clicks **Add or Remove Programs** in Control Panel, and then clicks **Add New Programs**, a list appears that provides available software categories. In these specific categories, the user can see a list of the software that is published for that user name. Users can install only the software you have published for them.

**Note**

- Because users might be accustomed to installing software from a designated share on your network, it is important that you educate users about installing and removing published software by using **Add or Remove Programs** in Control Panel**.**

For more information about using **Add or Remove Programs** to install software on a client computer, see "Making Software Available to Users and Computers" later in this chapter. For more information about creating software categories, see "Categorizing Applications" later in this chapter.

**Note**

- Files that have a .zap file name extension can only be published in **Add or Remove Programs**.

After the user installs a published application, it behaves like an assigned application until the user removes the application by using **Add or Remove Programs**, or until the software administrator removes the application.

**Publishing an Application Example**

Employees of an organization use a custom application that includes a corporate organizational chart and an employee locator map. This application is not essential to everyone because it does not directly affect the job they perform at the company. However, employees can save time locating coworkers by using the application. Therefore, most employees will use it occasionally.

The software administrators decided to publish this application for all users and to apply the GPO at the highest level of the domain hierarchy. A user who wants to gain access to this application can install it by using **Add or Remove Programs** in Control Panel.

**Publishing Software for Large and Small Groups of Users**

If you have a lightly managed IT environment, you can publish an application to all users at the domain-level without restriction, and then specify a category for the application, such as Sales. In this situation, you can expect users in the Sales department to install the software in the Sales category. However, this does not prevent unauthorized users from installing the software from the Sales category.

To prevent users from installing certain software, you can either assign software to the targets that need it instead of publishing it, or you can create different GPOs. You can also turn on loopback processing, a Group Policy setting that allows you to configure user-based policy settings in a GPO so that those settings are applied regardless of who logs on to the computer. For more information about using loopback processing, see "Designing a Group Policy Infrastructure" in this book.

To publish an application to smaller groups at lower levels of the domain infrastructure, plan for more administrative management than for small groups of users at a higher level. This kind of fine-tuning requires more GPOs or filtering.

# Categorizing Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can organize assigned and published applications in logical categories to make it easier for users to locate the applications that they need in **Add or Remove Programs**. Use the software installation extension of Group Policy to create and modify categories of software that you want to appear in **Add or Remove Programs**. If you have a lot of software to manage, make sure that users can easily locate the applications that they need to install by creating categories that define your organizational structure, job functions, and type of software.

**Note**

- Windows Server 2003 does not provide predefined software categories. You must define any categories that you want.

When you create software categories, do the following:

**Reflect your organizational structure** For example, when users in a department use a common set of applications, you can create a category named "Finance," which includes accounting applications such as Excel and Microsoft® Access.

**Reflect job functions** For example, you can create a category named "Project Managers," which includes applications such as Project and Excel.

**Classify applications by the type of software** For example, a category named "Presentation Tools" might include graphics programs that your organization supports, such as Microsoft® Publisher and PowerPoint.

You can also define broad categories, such as line-of-business tools or site-licensed applications.

The categories that you establish are per domain, not per GPO. You need to define them only once for the whole domain. Only users to whom you have assigned or published software can see or install the categories containing that software. To avoid creating duplicate categories, allow only one administrator to define categories for the entire organization.

**Creating Categories Example**

A software administrator at an organization created the Finance category for users in the Finance department. Only the users in Finance can view and install applications in this category because they are members of the GPO for which those applications are published. Correspondingly, the users in the Finance OU cannot view or install applications from the Administration, Sales, or Shipping categories because they are not members of the GPO for which those applications are published.

# Conducting a Pilot for Software Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Whether you assign or publish software, you need to conduct a pilot deployment using a small group of users to test and evaluates the software before you deploy it to the rest of the organization.

You can manage a software evaluation in any of the following ways:

- Evaluate the software outside the corporate environment. For example, set up a laboratory or test network environment.
- Create a GPO to manage the evaluation, and then assign or publish the software to users who are managed by that GPO.
- Edit the security settings on an existing GPO, or on the assigned or published package, to control who can install the package for evaluation.
- Manage the state of the software by switching any of the variables, such as assigned or published, visible or hidden, or auto-install set or not set. You can manage these parameters for each package by using the property pages of the software installation extension of the Group Policy Object Editor snap-in.

**Deploying Software for Users**

Although a single solution cannot fit all situations, the following situations illustrate commonalities that you are likely to encounter during a pilot that includes user who have different needs.

**Stationary workers**

In most organizations, users have their own computers and stay in one location. Typically the computers they use are desktop computers that have a consistent, high-speed connection to the network.

Typically, administrators make these clients their base model and define a standard operating environment. This environment includes standards for how to install and configure the operating system and how to use Group Policy to further configure and manage the environment. This environment also determines the software that is needed.

If these computers have the necessary hardware to use the Remote Installation Services (RIS) feature of Windows Server 2003, the administrator can include the necessary software with the operating system image. The software is then copied to the computer that runs the operating system.

For computers that do not have the necessary hardware to use Remote Installation Services, you can use the software installation extension of Group Policy to assign software to the computers. Do this by editing the GPO that manages the computers, and by using the software installation extension in **Computer Configuration** of the Group Policy namespace.

**Roaming users**

In many organizations, some users move or *roam* from one location to another. It is important to understand that, although the roaming users log on to different computers to do their jobs, these computers are typically connected by a high-speed connection or a LAN connection. Additionally, managed applications follow the roaming users to any computers they log on to. For example, when a user logs on to computer A and installs software, and then goes to another location and logs on to computer B, the user can see that the software that is installed on computer A is also installed on computer B. Any changes made to that software for the user or by the user on one computer propagate to all other computers to which the user logs on.

**Note**

- Because Group Policy software deployment settings are not applied over a slow link by default, roaming users who connect to their organization's network by slow links might not see changes to their software. You can change this default behavior in Group Policy so that applications install over slow links. You can also change the network speed threshold that Group Policy uses to decide whether a link is considered "slow." For more information about slow links and Group Policy, see "Designing a Group Policy Infrastructure" in this book.

You can assign applications either to the users or to the computers. For example, if all roaming users of a certain type use a certain application, it makes sense to assign that application to the computers so that it is already installed on the computers that the roaming users will use. Sometimes you might have to assign the software to the users, instead. In this case, when the

various roaming users log on to the computer, they see the applications that are assigned to them because the application is advertised. The application is only installed for the users who actually run the application.

**Mobile users**
Some users travel extensively to do their jobs. For example, sales personnel often spend more time at the offices of customers than at their own offices. Mobile users are different from roaming users in that they typically work from portable computers away from an office. Although mobile users log on to the same computer, their computers sometimes connect by a high-speed link and sometimes by a low-speed link.

One challenge in assigning applications to mobile users is that source files might not be available. Use the following recommendations if your organization has mobile users:

- When connected over a slow link, user assignment effectively behaves the same as publishing software to users. If Group Policy slow-link processing is set to the default in the user interface, the software is not installed on demand. However, users can go to **Add or Remove Programs** to install the assigned software.

- If users experience difficulty staying connected when they install software, verify that the connection speed and Group Policy settings are set appropriately. You can define the connection speed that is considered to be a slow link.

- Verify that all important software components that are defined by you for the user are installed initially. This allows a user who is not connected to the network to have access to necessary software components.

**Shared computers**
In many organizations, users share computers, such as bank tellers who work at different times and might use a different counter and computer each shift. In these environments, the software is often task-based. Although users change, the software does not. Also, the software might track who is logged on. You can manage these users or computers from a single GPO by grouping them appropriately, and then using the software installation extension of Group Policy to assign software to the computers by using the **Computer Configuration** of the GPO namespace. The software is then available for every user of that computer.

Note that when new software is assigned to a computer, it is installed when the computer restarts. If computers restart between shifts, the new software installation or upgrade might affect the total startup time of the computer. This increase in startup time occurs only if new software is assigned or the existing software is upgraded.

Another example where computers are shared is a computer lab or classroom where users share computers for a short period of time. This situation is different from the previous shared computer situation because each user might use the same applications as the previous user, or different applications. However, the computers do not move. In this case, using an Active Directory site to manage software makes sense, although grouping the computers into a single OU can also work.

Choose the method that gives you the correct level of control for applying Group Policy. Depending on your requirements, you might decide to assign software to the computer. This can work well if the software is written to keep user information (such as configuration information and saved files) separate from software information (such as executable files). Another way to manage this environment is to assign software so that users have access to the software that each needs for their training.

**Tip**

- To rebuild a shared environment quickly and efficiently, use RIS.

# Making Software Available to Users and Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can make software available to users in Control Panel**.** By using **Add or Remove Programs**, users can manage software on their own computers. However, you can control what software is available to users in **Add or Remove Programs** by using Group Policy settings.

**Add or Remove Programs** includes an active Web link for each application, which provides users with the support information they need to install certain applications. However, you can overwrite the default link by using the software installation extension of Group Policy. The support link then corresponds to your internal product support resources.

You can also have this Web link point to a support page that includes information such as an FAQ about a specified application, a help desk article about using the application, or instructions for requesting support. This can save time for both users and help desk personnel.

**Linking to an Internal Resource Example**
When deploying Office, an administrator at a corporation replaced the default product support link with an internal link to the organization's product support resource. This allowed users to request assistance for an Office issue from an internal resource instead of by going outside the organization for product support.

**Deploying Applications in a Managed Environment Example**
The administrators at an organization use Group Policy to deploy and manage software. Each time they deploy software in their organization, they prepare the software for Windows Installer, and then they distribute, target, and install the software.

The sales personnel of the organization require two new sales-based applications: a sales database and an application for order entry. The sales personnel include outside sales representatives, inside sales representatives, sales management, and clerical sales staff. These users are dispersed across four OUs in the domain infrastructure. Each OU is named for the five sales office locations: Seattle Sales, San Jose Sales, Boston Sales, Atlanta Sales, and Dublin Sales. Each OU contains up to 150 users.

The software administrator plans to publish the two new sales applications for all sales personnel. The following is the process the software administrator follows to deploy the new sales applications. The software administrator performs the following tasks:

1. Obtains the Windows Installer Packages. The new sales software comes complete from the manufacturer with Windows Installer packages (SalesDB.msi and OrderEntry.msi). Each application is less than 50 MB.

   **Note**

   - Because the author of the application provided the .msi files, the administrator did not need to reauthor or repackage the software.

2. Creates a test environment in the lab to verify that the installation works as planned.

3. Places the software distribution point servers as close as possible to the users. In this case, the IT administrator uses the existing software distribution points in each of the U.S. sales offices because they are only 20 percent used. For the Dublin, Ireland, sales office, the IT administrator sets up a new software distribution point file server because none currently exists.

4. Creates the file shares on the software distribution points, and then copies the packages to the shares.

5. Uses DFS to manage the network traffic. The administrator configures DFS to manage the network traffic during software installation time because DFS offers load-sharing among servers and increases availability by distributing the same data across multiple servers. This helps balance the software installation traffic and protects against bottlenecks.

6. Creates a new GPO named Sales Personnel, specifying that these two applications be published to the members of the sales office organizational units. Before linking the GPO, the administrator runs GPMC Modeling to further test the deployment.

7. Links the GPO.

8. Breaks up the software deployment into four phases to minimize network traffic. However, the administrator does not want to create four GPOs. The deployment is performed Monday through Thursday for up to 150 users per day by using e-mail. On Monday the administrator sends half of the users in Atlanta and half of the users in Boston an e-mail message telling them that the new applications, "Sales Database" and "Order Entry," are now available for installation under the **Sales Applications** category in **Add or Remove Programs**. The following day, the administrator contacts the other half of the users in the Atlanta and Boston sales offices. The administrator follows the same process for all sales offices until all sales personnel are informed about the new software.

While the deployment is progressing, IT administration supports both the previously installed sales applications and the newly deployed sales applications.

After deployment, the administrator allows a grace period for users to upgrade to the new version. The administrator informs the users that the old sales database and order entry applications will expire 30 days after the new software is fully deployed. This encourages all sales personnel to install and use the new software in a timely manner.

# Migrating Applications to a Managed Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you decide to migrate your users and computers from an unmanaged environment to a managed environment, you must also choose a method to migrate your unmanaged applications. Figure 8.6 shows this phase in the Group Policy-based application-deployment process.

**Figure 8.6 Migrating Applications Stage of Deployment**




The two most common methods of migration are the following:

- Deploy applications to Windows NT 4.0 for transition to a managed Windows Server 2003 environment. Use this method if you are using a version of Windows that is earlier than Windows 2000.

- Use RIS to deploy software, and use the software installation extension of Group Policy to manage it. Use this method if you run Windows 2000 Server or Windows Server 2003. This method is recommended for deploying the operating system and standard applications to a large group of users or computers.

In either case, you can use the software installation extension of Group Policy to manage your applications.

**Deploying Applications for Transition to a Managed Windows Server 2003 Environment**
If you have plans to migrate to Windows Server 2003, you can deploy applications in a Windows NT 4.0 environment now, and then migrate to the Group Policy–based, managed Windows Server 2003 environment later. By planning ahead for a migration from Windows NT 4.0 to Windows 2000 Server or Windows Server 2003, you do not have to remove and reinstall applications later.

**Planning for Future Migration Example**
A software administrator at an organization wants to deploy Office 2000 in a Windows NT 4.0 environment. Management wants to deploy the Windows 2000 Server or Windows Server 2003 operating system, and then use Group Policy to manage Office 2000. With that in mind, the administrator anticipates moving computers that have the unmanaged Office 2000 to the managed environment. This is to make sure that the administrator will not have to remove and then reinstall Office 2000.

Use a Windows Installer package, a transform (if you must have one) for the package, and a software distribution point that is exactly the same as the Windows Installer package and transform combination that you plan to use when it is managed.

You can create the software distribution points now for use after you successfully transition to the new software.

**To make sure that the future transition is successful**

1. Create the software distribution points, and then install one or more Windows Installer package and transforms in the

same shared folder.

2. Install the software by using the Windows Installer with the following parameters:

```
\>msiexec /I \\servername\share\<software.msi> TRANSFORMS = \\servername\share\
<software.mst> /qb
```

In this example, Software.msi is the Windows Installer package that you are installing, and Software.mst is the transform that you want applied at deployment time.

After Windows 2000 Server or Windows Server 2003 is in place, including the Active Directory and Group Policy infrastructure, you can perform the following tasks:

1. Assign the software in the appropriate GPO, using the same software distribution point, Windows Installer package, and transform.

2. Move the computer object into an Active Directory container with the associated GPO, or link the GPO to the Active Directory container with the computer object.

**Important**

- Use the original Windows Installer package, the same transform, and the same software distribution point and shared folder when you redeploy the application in the managed environment. In this case, the software will not be removed and then reinstalled. Instead, the software installation extension of Group Policy detects that the software is the same and does only what is necessary to continue to manage the software in the new environment.

**Using RIS to Deploy Software and Group Policy to Manage Applications**

You can use RIS to deploy the operating systems and specific applications to client computers, and then use the software installation extension of Group Policy to manage the applications.

When an organization first acquires new computers, IT administrators spend a lot of time preparing them for the users. Many highly managed organizations format the hard drives of new computers to configuring them to their organization's standard configuration.

To save time, you can use RIS to create a customized image of a Windows XP Professional desktop on a source computer. Then you can save that desktop image to the RIS server. That image can include only the operating system, or it can be a preconfigured desktop image that includes the operating system and standard locally-installed desktop applications. You can use that preconfigured image to set up multiple desktops. You can create as many desktop images as required to meet the needs of all types of users in your organization.

You can rapidly and efficiently deploy the operating system and standard applications by using RIS, and then bring the software into a state where you can manage it by using the software installation extension of Group Policy.

For more information about using RIS to deploy software, see "Designing RIS Installations" in *Automating and Customizing Installations*.

**Migrating Software to a Managed Environment Example**

An organization continuously acquires new computers for new users and to replace old computers. The administrator saves valuable time by installing and customizing the operating system and standard software, which in this case includes Office XP. The administrator uses RIS to deploy the software, and then uses Group Policy to manage the applications after deployment. To accomplish this, the administrator performed the following tasks:

1. Set up and configures a RIS server.

2. Created a RIPrep image.

3. Used the Group Policy snap-in to create a GPO to manage the computers that are associated with the RIPrep image. Because this is an effort to standardize all desktops in the organization, this GPO applies to all computers in the organization.

4. Placed and sets up a software distribution point file server for Office XP.

   **Note**

   - Office XP includes a native Windows Installer package (.msi). Therefore, the administrator did not have to obtain a package for this application.

5. Used the Custom Installation wizard (in Office XP) to customize Office XP based on the users' needs. This produced a transform.

6. Used Group Policy–based software deployment to assign Office XP to the computers in the organization-wide GPO.

7. Installed Windows XP Professional on a new source computer, and then configured the operating system. The source computer does not need to have the same hardware, but it must use the same Hardware Abstraction Layer (HAL) as the computers that the image is installed on.

8. Added the computer to the Active Directory container where it remains when it is deployed. This container holds the GPO that has Office XP assigned to the computer associated with it.

9. Restarted the source computer. When the computer restarted, the software that the administrator assigned to the computer (by using the software installation extension) installed.

After installing the software on the RIS server, the administrator used the RIPrep tool of RIS to build a desktop image of the source computer that has Windows XP Professional and Office XP installed, and then put this image on the RIS server.

When the image becomes available, a user who receives a new, PXE-enabled, client computer that supports the RIS server only has to connect to the network (connect a cable from the network card to the hub), connect the keyboard, the mouse, and the monitor, and turn on the computer.

The client computer locates the RIS server, and then downloads the Windows XP Professional operating system and Office XP. When the computer restarts after installing Windows XP Professional remotely, Windows Installer determines that Office XP is already installed on the computer, and then updates only the advertisement information. The advertisement update takes only a few seconds.

The user logs on to the computer, and then selects the first Office XP application. This causes Windows Installer to start. Although Office XP is already installed, Windows Installer properly separates the installation from the user configuration. To complete the small amount of configuration that is required for each user, Windows Installer starts each time a new user starts Office XP. This happens in Office XP regardless of whether Office XP is assigned to the user, assigned to the computer, published for the user, or installed by using of RIS.

# Patching, Upgrading, and Removing Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The purpose of deploying patches, updates, and upgrades is similar to deploying new applications, to bring computers up-to-date. When software becomes obsolete and can no longer be updated or patched, you must remove the software.

The tasks involved in maintaining software that has been deployed are described in Figure 8.7.

**Figure 8.7 Maintaining Applications After Deployment**




It is important to have a decision-making process for determining when to patch, upgrade to new versions of managed applications, or retire old versions. The process of deploying, patching, upgrading, and retiring is known as the *software life cycle*. Figure 8.8 illustrates a sample software life cycle.

**Figure 8.8 Sample Software Life Cycle**


Deploying software in a Windows 2000 or later environment?
No
Yes
Migrate software to a managed environment
Deploy applications in a managed environment
Prepare version 1 for deployment
Deploy version 1
Patch Needed?
Yes
No
Prepare patch A for deployment
Deploy patch A
Prepare version 2
New Users?
Yes
No
Prepare version 2 for new deployment
Deploy version 2
Prepare version 2 for upgrade
Deploy upgrade to version 2
Retire version 1 that has patch A applied

# Updating Installed Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order to update previously installed applications, you obtain updates (.msp files) from the software manufacturers or from the internal developers of the original program. You can update an existing application without removing the product. This approach preserves the customizations of the installation and can lower the cost of making the change. The update might change only a few bytes of a single application file. It is more efficient to distribute those few bytes than to remove and redeploy the whole product.

**Note**

- An update might change all of the files and registry keys in a product.

After you apply the update, it is cached on the user's computer, this allows the user to:

- Perform any installation on demand.
- Reinstall the application.
- Repair the application.
- Remove the application.

The updating information in this chapter applies to those organizations that have deployed software by using the software installation extension of Group Policy.

A update package (.msp) does not include a database as a regular installation package does. Instead, it contains, at a minimum, one database transform that adds update information to the database of its target installation .msi package. For an application to perform maintenance operations, such as adding, removing, or repairing the installation, the package codes for both the installed application and the source must match.

**Important**

- To remove an update after applying it, you must remove the entire application, and then reinstall it without the update. There is no way to roll back the changes made by an update.
- Do not deploy an update to an application that might add 64-bit components to a 32-bit application. This can cause failures on 32-bit clients because they cannot run the 64-bit component.

For more information about applying updates by using command-line options, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection).

For step-by-step instructions for updating applications, see Help and Support Center for Windows Server 2003.

# Upgrading Installed Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you determine the type of software renewal method you plan to use, Group Policy simplifies the upgrade process.

Before you update the software in your organization, you must determine which of three types of software renewal methods is appropriate for your situation:

- Small Upgrade
- Minor Upgrade
- Major Upgrade

Table 8.5 provides a quick view of the significant differences among small updates, minor upgrades, and major upgrades.

**Table 8.5 Comparing Upgrade Types**

| Type of Renewal | Description | ProductCode Property | ProductVersion Property |
|---|---|---|---|
| Small update | An update to one or two .msi or application files that is too small to warrant changing the product version. The package code in the Revision Number Summary Property does change. This can be shipped as a patch package or as a full product installation package. | No change | No change |
| Minor upgrade | A small update that makes large enough changes to warrant changing the product version, but not the product code. By changing the product version, you prevent downgrades. You also enable package sequencing. This can be shipped as a patch package or as a full-product installation package. | No change | Change |
| Major upgrade | A comprehensive update of the product warranting a change in the product code. Essentially, it is a new installation; optionally, it is an application removal. This can be shipped as a patch package (.msp) or as a full product installation package (.msi); only possible by using Windows Installer version 1.1 or later. | Change | Change |

**Small Updates and Minor Upgrades**

Small updates and minor upgrades (see Table 8.6) are essentially a reinstallation that uses a new package. Regardless of the type of software renewal you choose, the .msi package code changes to make sure that the new package is used.

**Table 8.6 Comparing Small Updates with Minor Upgrades**

| Situation | Renewal Method |
|---|---|
| A software renewal comes from the application manufacturer, and it does not include changes to the product version or product code. | A small update to the original package |
| A software renewal from the manufacturer changes the product version, but not the product code. | A minor upgrade to the original package |

The only difference between small updates and minor upgrades is if you must differentiate between product versions.

**Major Upgrades**

During a major upgrade, Windows Installer searches the user's computer for applications that are related to the pending upgrade. When it locates one, Windows Installer retrieves the version of the installed application from the registry. Windows Installer then uses information in the upgrade version's database to determine whether to upgrade the installed application.

During a major upgrade, the installed version of an application might be removed, or might remain and coexist on the system

with the newer version. This behavior, and its exposure to administrators, depends on the implementation of the application's setup program. For example, multiple versions of an application that are installed on the same computer might prevent the version that you install from functioning properly.

If any of the following conditions and application changes applies to your organization, you must perform a major upgrade:

- Coexisting installations of the original and updated products on the same computer
- A change in an .msi file name
- A change in the component code of an existing component
- An addition or removal of a component from an existing feature
- A change that causes an existing feature to become a child of an existing feature
- Removal of an existing child feature from its parent feature

You do not need to change the product code when you add a new child feature that consists entirely of new components that are added to an existing feature.

**To perform an update or upgrade**

1. Provide adequate notice to the users that the software will be updated or upgraded by a specified date.
2. Obtain the Windows Installer package file.
3. Locate the software distribution point server and share where the original package resides.
4. Apply the update or upgrade to the original package.
5. Inform the users that the updated software is now available.

It is important that you support both the newly installed update or upgrade and the software that you plan to retire soon. Be sure to maintain the previous distribution point until all users have upgraded. Eventually, you must remove the retired version of the software. For information about removing outdated software, see "Patching, Upgrading, and Removing Applications" later in this chapter.

**Upgrading Software by Using Windows Installer and Group Policy**

During a minor upgrade or a major upgrade, Windows Installer searches for upgradeable products by querying the **Upgrade** table of the upgrade package. The newer version of the product is installed. If Windows Installer finds an older version of the product, it removes the old version. The author of the application's setup can choose to remove the old version, and then install the new version. The maintenance mode and removal do not trigger these actions because "remove existing versions of an application" is now automatic.

You can use the software installation extension of Group Policy to manually create upgrade relationships between the new package and the packages that the application replaces. This includes making a formal upgrade relationship between two similar products from completely different vendors. Again, you can either replace one vendor's application with another, or you can upgrade a repackaged application. Of course, it is recommended that you pilot or test an upgrade before putting it into production.

For examples of upgrades that are deployed by using the software installation extension of Group Policy, see "Patching, Upgrading, and Removing Software Examples" later in this chapter.

For a more information about configuring upgrades, and for a procedure for upgrading applications by using the software installation extension of Group Policy, see Help and Support Center for Windows Server 2003.

If you plan to deploy repackaged applications by using the software installation extension of Group Policy, you must allow extra time for unexpected situations that require you to create upgrade relationships manually or to create a script to remove unwanted files.

**Manually creating upgrade relationships**

The Windows Installer package for a repackaged application does not have declared upgrade relationships. You must manually create upgrade relationships. You can configure these relationships by using the software installation extension of Group Policy.

**To manually create an upgraded relationship**

1. In the Software Installation section of the Group Policy object, right-click the managed software.

2. Click **Properties**.

3. On the **Upgrades** tab, specify the packages that the selected package will upgrade and the packages in the current GPO that will upgrade this package.

4. Click **Add** to select the packages.

**Writing a script to remove files**

During an upgrade, it might not be possible to completely remove a repackaged application. The removal of a repackaged application might leave components on the desktop, regardless of whether the component is shared or needed. You can create a script to remove the remaining files.

# Removing Installed Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you no longer support a software version, plan for its removal. First, be sure to give adequate notice to the software users before you remove the application.

You can either force the removal of the software or make it optional, depending on, for example, whether the old software is compatible with newly deployed software. These two choices, available from the software installation extension of Group Policy, affect the removal of software.

**Forced removal** If the software is assigned to a computer, it is removed the next time the computer restarts. If it is assigned to a user, it is removed the next time the user logs on.

**Optional removal** You can stop managing the software without forcing its physical removal from the computers of users who use it. Users who currently have the software installed can continue to use it until they choose to remove it themselves.

Figure 8.9 shows how to open the **Remove Software** options menu dialog box. Right-click the software installation extension of Group Policy to view this menu.

**Figure 8.9 Locating the Remove Software Option Menu**




**Note**

- To save time later, you can plan for eventual removal when you initially deploy the software. If you want the application to be removed when a GPO is no longer applicable, select the **Uninstall this application when it falls out of the scope of management** option. For more information about this option, see "Targeting Software to Users and Computers" earlier in this chapter.

See the following section for an example of removing software by using the software installation extension of Group Policy.

# Restricting Software Access and Protecting Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although a business might use a set of programs that it knows and trusts, and might train its administrators and Help Desk personnel to support those programs, administrators lose control as soon as users begin running unknown code. The increasing role of the Internet in business puts your network at a greater risk than ever before.

Such user-installed programs can conflict with other installed programs, change critical configuration data, or introduce a virus. Users often cannot make safe or informed choices about what software to run because viruses intentionally conceal the malicious purpose of the program. Also, the problems that are associated with running unknown code can increase support costs substantially because they lead to more system maintenance, more help desk time, and lost user productivity.

Software restriction policies, new with Windows XP and Windows Server 2003, provide a policy-driven mechanism that enables you to identify the programs that are running on computers in your domain, and to control their ability to run. By using software restriction policies, you can:

- Control what programs run on your system. For example, you can apply a rule that does not allow certain file types to run in the mail attachment directory of your e-mail program if you are concerned about users receiving viruses through e-mail.
- Run only digitally signed scripts.
- Allow users to run only specific files on multiuser computers. For example, if you have multiple users who use a single computer, you can set up software restriction policies and Access Control List settings so that users cannot make changes to computers.
- Decide who can add trusted publishers to a computer.
- Control whether software restriction policies affect all users or only certain users who use a computer.
- Prevent any files from running on a local computer. For example, if you are aware of a known virus, you can disallow a hash of that virus so that the computers in your domain cannot run that program.

A software restriction policy consists of a rule, or set of rules, that determines what programs are allowed to run, and any exceptions to the rule. Use the Group Policy Object Editor to create software restriction policies. Go to the User or Computer configuration, as appropriate. Under **Windows Settings**, click **Software Restriction Policies**. Here you can configure the parameters for **Security Levels**, set **Additional Rules**, specify **Designated File Types Properties**, and identify **Trusted PublishersProperties**.

The purpose of a rule is to identify one or more software applications, and to specify whether the application is allowed to run. Creating rules is mainly identifying software that is an exception to the default rule. You can include descriptive text with each rule to help communicate why the rule was created. A software restriction policy supports four rules to identify software.

**Hash rule**
A cryptographic fingerprint of the file, also called a *message digest*. When you create a hash rule for a program, Software Restriction Policies calculates a hash of the program, and then stores the hash securely. When a user tries to open a program, a hash of the program is compared to existing hash rules for Software Restriction Policies. The hash of a program is always the same, regardless of the location of the program on the user's computer. However, if a program is altered in any way (by applying a hotfix, for example), its hash also changes, and it no longer matches the hash in the Software Restriction Policies hash rule.

For example, you can create a hash rule, and then set the security level to **Disallowed** to prevent users from running a certain file. A file can be renamed or moved to another folder and still result in the same hash. However, if any changes are made to the file itself, they also change its hash value and allow it to bypass restrictions.

**Certificate rule**
A software publisher certificate used to digitally sign a file. For example, a company can require that all scripts and Microsoft® ActiveX® controls be signed with a particular set of publisher certificates. Certificates that are used in a certificate rule can be

issued from a commercial Certification Authority (CA) such as VeriSign, a Windows 2000 or Windows Server 2003 public key infrastructure (PKI), or by a self-signed certificate. Certificate rules are a strong means to identify software because the certificate matches the files, regardless of name or location, by using signed hashes that are contained in the signature of the signed file.

**Path rule**
The local or universal naming convention (UNC) path to where the file is stored. Path rules can include the following:

- Environment variables. Path rules are evaluated in the client environment. Therefore, using an environment variable, such as %Windir%, allows a rule to adapt to a user's environment.

- Wildcard characters such as and asterisk (***)** or question mark (**?)**. By including a wildcard character, a rule can match all files of the specified type. For example, a rule such as *.vbs matches all Microsoft® Visual Basic® script files.

- Registry paths. You can create a path rule that looks up these registry keys for applications that store paths to their installation folders or application directories in the system registry. The registry path is formatted as follows:

  **%[*Registry Subtree* (Hive) *Name*]\[*Registry Key Name*]\[*Registry Entry* (Value) *Name*]%**

  **Note**

  - A registry path rule suffix cannot contain a backslash (\) character immediately following the last percent sign (%) in the rule.

The following conditions apply:

1. The registry path must be enclosed by percent signs (%).

2. The registry setting must be a **REG_SZ** or **REG_EXPAND_SZ** type. If the registry value contains environment variables, these will be expanded when the policy is evaluated.

3. Do not use HKLM as an abbreviation for **HKEY_LOCAL_MACHINE**, or HKCU as an abbreviation for **HKEY_CURRENT_USER**.

4. A registry path rule can also include a suffix path.

   For example, you can use:

   **%HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\Shell Folders\Cache%OLK***

   **Note**

   - When you set a path rule, verify the Access Control List (ACL) entries on the path. If users have Write access to a path, they can modify its contents. For example, if you allow Write access to C:\Program Files, any power user on the computer can copy software to the Program Files folder.

**Zone rule**
An Internet zone rule identifies software by the Microsoft® Internet Explorer zone from which it is downloaded. These zones are: Internet, Intranet, Restricted Sites, Trusted Sites, and My Computer. This rule applies only to Windows Installer (.msi) packages. It does not apply to software that is downloaded by using Internet Explorer.

When a user tries to install a piece of code, Windows Installer queries the software restriction policy to determine the level at which the code is allowed to run. Software restriction policies integrate with the operating system and common scripting runtimes to control the running of software, not just to hide access to applications by removing them from the **Start** menu or to hide the **Run** command. Software restriction policies go beyond this by removing the common access points for software.

Essentially, software restriction policies protect against the various Trojan horse and worm viruses that propagate through e-mail and over the Internet. Software restriction policies are a powerful way to identify software so that it can be classified as

**Unrestricted** or **Disallowed**. After you identify programs, you can apply a policy to either restrict the software or to let it run.

Table 8.7 shows some recommended situations for applying available software restriction policy rules.

**Table 8.7 Situations for Applying Each Rule**

| Sample Task | Situation to Apply Rule |
|---|---|
| Allow or disallow a specific version of a program. | **Hash rule**<br><br>Browse to file to create hash. |
| Identify a program that is always installed in the same place. | **Path rule with environment variables**<br><br>%Program Files%\Internet Explorer\iexplore.exe |
| Identify an anti-virus program that can be installed anywhere on a client computer. | **Registry path rule**<br><br>%HKEY_LOCAL_MACHINE\SOFTWARE\ ComputerAssociates\InoculateIT\6.0\Path\HOME% |
| Identify a set of scripts on the central server. | **Path rule**<br><br>*Server Name**Share Name* |
| Identify a set of scripts on a set of servers named DC01, DC02, and DC03. | **Path rule with wildcards**<br><br>\\DC??*Share Name* |
| Disallow all .vbs files except those in a login script directory. | **Path rule with wildcards**<br><br>*.VBS set to Disallowed<br><br>\\LOGIN_SRV*Share Name**.VBS set to Unrestricted |
| Disallow a file that is installed by a virus that is always named Flcss.exe. | **Path rule**<br><br>Flcss.exe set to **Disallowed** |
| Identify a set of scripts that can be run anywhere. | **Certificate rule**<br><br><Certificate used to digitally sign the scripts> |
| Allow software to be installed from trusted Internet zone sites. | **Internet zone rule**<br><br>Trusted Sites set to **Unrestricted** |

The following guidelines can help you use software restriction policies effectively:

**Create a separate GPO for software restriction policies** When you create a separate GPO for your software restriction policy settings, you can turn them off in an emergency without affecting the rest of your security or other settings.

**Never modify the default domain policy** When you do not edit the default domain policy, you can always reapply it.

**Never link to a software restriction policy in another domain** Linking to a GPO in another domain can result in poor performance, or it might not work at all. This is true for any GPO, not just for software restriction policies in a GPO.

**Test your software restriction policy first** Typographical errors in rules can result in a software restriction policy that does not perform as you expect. Be sure to test your policy on a test computer first.

**Turn off a software restriction policy while editing it** When you edit a software restriction policy, each change you make updates the GPO. If a computer tries to synchronize Group Policy settings during an editing session, the user might receive the software restriction policy before you finish editing it. If this happens, the user's computer might be adversely affected.

**Verify that trusted virus scanners are not restricted** Most antivirus software has a real-time scanner program that starts when the user logs in. This program scans all files that are accessible by the user to locate possible virus contamination. Make sure that your rules allow your virus-scanning programs to run. Software restriction policies are not a replacement for antivirus software.

**Filter user policies based on membership in security groups** You can decide which users do not have a Group Policy object applied to them by denying them the **Apply Group Policy** or **Read** permissions on that policy. You require both of these permissions to apply Group Policy.

**Use software restriction policies in conjunction with access control settings** Prepare carefully if using Disallowed as the default setting. Many applications start other applications to perform certain tasks. Verify that all major tasks in your applications are covered by your rules. If your computer needs to run logon scripts, make sure that you have created a path rule that allows them to run.

**Restart in Safe Mode if you experience problems with applied policies** Software restriction policies do not apply in Safe Mode. If you accidentally lock down a workstation by applying software restriction policies, restart in Safe Mode, log on as a local administrator, modify the policy, run Gpupdate.exe, restart the computer, and then log on the usual way.

**Always test policy on a test computer before applying it to any other computers** Do not apply rules at the **Disallow** level without the proper testing. Restrictions on certain files can seriously affect the operation of your computer or network.

Rules are evaluated in a specific order. If two rules are similar or conflict, the more restrictive rule takes precedence, and the rules that more specifically match a program take precedence over rules that loosely match a program. In the case of multiple matching path rules, the most specific matching rule takes precedence. Because each rule has an associated GUID, two identical rules have two GUIDs.

**Note**

- Environment variables are not protected by Discretionary Access Control Lists (DACLs). Therefore, if users can start a command prompt, they can redefine an environment variable to a path of their choosing.

The overall rule precedence is as follows:

1. Hash rule
2. Certificate rule
3. Path rule
4. Internet zone rule
5. Default rule

When multiple path rules match, the most specific matching rule has precedence. The following is a set of paths, listed from highest precedence (a more specific match) to lowest precedence (a more general match):

- Drive:\Folder1\Folder2\FileName.Extension
- Drive:\Folder1\Folder2*.Extension
- *.Extension
- Drive:\Folder1\Folder2\
- Drive:\Folder1\

# Patching, Upgrading, and Removing Software Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An essential part of the software life cycle involves patching, updating, upgrading, and removing software. This section describes how to use Group Policy–based software deployment to simplify managing this life cycle.

**Patching an Installed Application Example**

The address of corporate headquarters for the organization has changed and the software administrators want to deploy a Microsoft® Word version 2002 fax and letterhead template patch to replace the previous fax and letterhead template files. Except for the sales offices, the OU structure is designed around the organization's departmental structure. With this in mind, the IT administrator must selectively deploy the fax and letterhead templates to the departmental OUs and the Seattle-based sales office. The sales offices in San Jose, Boston, and Atlanta do not need to receive the patch. The administrator performed the following tasks to patch some previously deployed, computer-assigned software:

1. Copied the patched (.msp) file to the applicable software distribution points. This replaced only the previously deployed fax and letterhead files for Word 2002 on each file server.

2. Opened the software installation extension in the GPO. The GPO includes the software installation extension that contains the application deployment properties of Word 2002.

3. Applied the patch to the source image by running **msiexec /a <package.msi> /p <patch.msp>**.

4. Redeployed the fax and letterhead files. In the GPO, the administrator right-clicked **Word 2002**, selected **All Tasks**, and then selected **Redeploy Application** from the available menu.

5. Prompted users by e-mail to restart their computers to make the computers recognize the patched software. After the users restarted their computers, the GPO applied the patched file, and the users had access to the new fax and letterhead templates.

**Performing an In-Place Upgrade Example**

For years, everyone at the organization has used a custom application specific to company business. This software is mission-critical for all users. Therefore, it will continue to be assigned to everyone in the organization. Recently, the company recruited developers to rewrite the application to include a Windows Installer package. The software administrator wants to deploy the new .msi package, and has chosen to perform an in-place upgrade. The administrator used the following process to upgrade the software:

1. Copied the new .msi package upgrade to the software distribution points.

2. Created a GPO that published the new software to all users.

3. Advised all users to update their software by the established date, with a reasonable grace period to accomplish the task.

4. After verifying that all users had upgraded, changed the software installation configuration in the GPO to assign the new version to all users.

5. Sent e-mail prompting the users to restart their computers.

After the users restarted their computers, the new Group Policy setting applied, and the users had access to the new custom application.

**Performing an Upgrade to Remove and Replace Software Example**

The organization signed a global licensing agreement with a new antivirus software provider. Until then, all computers in the organization had *X* antivirus software installed. By using the software installation extension of Group Policy, IT administrators safely replaced *X* antivirus software with *Y* antivirus software.

Antivirus software is critical to any network operating environment. Therefore, IT administrators assigned *Y* antivirus software to all computers in the organization. The IT administrator used the following process to replace *X* antivirus software with *Y* antivirus software:

1. Opened the software installation extension in the GPO. The software installation extension contained the application deployment properties of both X antivirus software and Y antivirus software.

2. Created an upgrade relationship between X antivirus software and Y antivirus software. On the Upgrade tab of the Y antivirus software Properties, the administrator selected the Required upgrade for existing packages check box. This makes all users of X antivirus software receive the upgrade to Y antivirus software.

3. Sent e-mail to prompt users to restart their computers.

   After the users restarted their computers, the computers recognized the new *Y* antivirus software.

**Removing Software Example**

The organization terminated a site license agreement with a certain software distributor for a marketing application that was previously published to users in the Marketing department. Software administrators of the organization needed to verify the removal of all copies of the unlicensed software. The software administrator used the following procedure to remove the unlicensed software in a GPO:

1. Opened the software installation extension in the GPO, which included the deployment properties of the application to be removed.

2. Right-clicked the application that the administrator wanted to remove.

3. Selected **All Tasks**, and then selected **Remove**.

4. In the **Remove** dialog box, selected the **Immediately Uninstall** check box.

After completing the removal procedure, the administrator did not prompt users to restart their computers. Instead, users were allowed to restart their computers at various times, in the course of their own work routines. That approach naturally reduced network load. However, use care at peak times, such as Monday mornings, when many people arrive at work at the same time. If the software is assigned to the users, their typical log off and log on process performs the same function.

After the users restarted their computers, the GPO was applied, and the users no longer had access to the unlicensed software.

# Troubleshooting Software Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The new Group Policy management snap-in (GPMC) provides tools for preventing and resolving Group Policy–based software installation issues. Also, Microsoft® Windows® Server 2003, Standard Edition includes command-line RSoP tools that you can use to prevent and troubleshoot problems that might occur when you deploy Group Policy–based software (process indicated in Figure 8.10). GPMC and RSoP help you to do the following:

- Assess how Group Policy affects software deployment, updates, and removal.
- Identify and diagnose Group Policy setting failures.
- Prevent conflicts among applications that share components.
- Protect computers from unknown or infected code.
- GPMC also provides printable HTML reports for modeling your Group Policy design before deployment, and for logging its results after deployment.

**Figure 8.10 Using GPMC to Troubleshoot Your Application Deployment**




To prevent problems, use the Group Policy Modeling feature of GPMC to determine how software deployment, update, and removal configurations affect a target group of users and computers before you deploy, update, or remove software. To troubleshoot post-deployment problems, use the Group Policy Results feature of GPMC to see what Group Policy settings are actually in effect for a user or computer.

For more information about using GPMC to prevent and resolve Group Policy deployment issues, see "Designing a Group Policy Infrastructure" in this book. For complete step-by-step information about using GPMC, see GPMC online Help.

# Using GPMC Group Policy Modeling to Evaluate Group Policy Settings Before Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use Group Policy Modeling to calculate the net effect of several GPOs before deployment. These calculated settings are reported in HTML and are displayed in a GPMC browser window on the **Settings** tab in the **Details** pane for the selected GPO. You can expand and contract the settings under each item by clicking **hide** or **show** so that you can see all the settings, or only a few. To create a Group Policy Modeling query, you must have the **Perform Group Policy Modeling analyses** permission on the domain or OU that contains the objects on which you want to run the query.

To run the wizard, right-click **Group Policy Modeling** (or an Active Directory container), and then click **Group Policy Modeling Wizard.** If you run the wizard from an Active Directory container, the wizard fills in the **Container** fields for user and computer with the LDAP distinguished name of that container.

When you have answered all the questions in the wizard, your answers are saved as a query that is represented by a new item under the **Group Policy Modeling** item. Your answers appear as if they were from a single GPO. However, the display does show which GPO is responsible for each setting, under the heading **Winning GPO**.

To save the results of the modeling, right-click the query, and then click **Save Report**. You can also print the results by right-clicking the query, and then clicking **Print**.

# Using GPMC Group Policy Results to Evaluate Group Policy Settings After Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use Group Policy Results to see what Group Policy settings are actually in effect for a user or computer. Use this GPMC feature in your staging environment before you deploy managed software in your production environment. The settings are reported in HTML and appear in a GPMC browser window on the **Settings** tab in the **Details** pane for the selected GPO. You can expand and contract the settings under each item by clicking **hide** or **show** so that you can see all the settings, or only a few. To access Group Policy Results data for a user or computer, you must have the **Remotely access Group Policy Results data** permission on the domain or organizational unit that contains the user or computer, or you must be a member of a local Administrator's group on the appropriate computer

To run the wizard, right-click **Group Policy Results**, and then click **Group Policy Results Wizard**.

When you have answered all the questions in the wizard, your answers are saved as a query that is represented by a new item under the **Group Policy Results** item. Your answers appear as if they were from a single GPO. However, the display does show which GPO is responsible for each setting, under the heading **Winning GPO**.

To save the results, right-click the query, and then click **Save Report**. You can also print the results by right-clicking the query, and then clicking **Print**.

**Preventing and Resolving Software Deployment Issues Examples**
You can use the Group Policy Modeling feature of GPMC to prevent or resolve software deployment problems. Perhaps more importantly, you can use Software Restriction Policies to control the running of code and prevent future problems.

**Example one**
By using Group Policy Modeling in GPMC, an administrator simulated moving a user from one OU to another to see what effect it might have on that user. Each of these OUs had two separate GPOs applied to it. By seeing the results of the simulated move, the administrator discovered that the user's spreadsheet application had been downgraded to an earlier version. Additionally, the user no longer had access to a word processor application.

Before the actual move, the administrator of the existing OU contacted the administrator of the target OU to request that the most recent version of the spreadsheet and word processor applications be made available to the migrating user in the target OU. The target administrator did this by creating a security filter, and then making the user a member of either one or the other, and then filtering.

**Example two**
An administrator wanted to examine the result of deploying a software upgrade of an existing application that was assigned to a particular group of users. When the administrator ran the scenario by using Group Policy Modeling in GPMC, the following unexpected result occurred: The GPO that contained the package was not visible to users in the Sales and Marketing departments. By examining the HTML results of the Group Policy Modeling, the administrator verified that another GPO was taking precedence over the GPO that contained the new upgrade. In this case, the administrator configured the GPO that contained the new upgrade to override the GPO that was already in place.

For more information about using GPMC to prevent software deployment issues, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**Troubleshooting by Using GPMC Modeling Examples**
The following situations demonstrate ways to use GPMC Modeling at the organization level.

**Example one**
A help desk administrator used logging mode to locate and determine the reason that a user received the wrong version of an application. The user received a less-powerful version of the required software.

The administrator used the Group Policy Modeling in GPMC to verify which version was installed. The HTML report showed that the client computer did not have the correct language version.

After diagnosing the problem, the administrator published the correct version of the application. The user could then install it by using **Add or Remove Programs**.

**Example two**

A sales employee from Seattle accepted a new position in the Boston Human Resources department. The standard applications that were available in Seattle were Microsoft® Exchange E-mail, Office, a sales database, and an order entry application. In Boston, the employee discovered that he did not have access to the Human Resources database, so he contacted the help desk for assistance.

The help desk administrator started Group Policy Modeling, and then viewed the software installation extension part of the HTML report. The report showed which applications appeared in **Add or Remove Programs** on the client computer. The **Origin** field in the list of available applications showed that all the applications were coming from the GPO that was linked to the Seattle OU.

Because GPOs are set at the OU level, the administrator moved the user from the Seattle OU to the Boston OU. After that, the user could install the Human Resources database by using **Add or Remove Programs**.

For more information about using GPMC for troubleshooting software deployment issues, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**Blocking a Malicious Script by Using Software Restriction Policies Example**

The organization wanted to protect itself from script-based viruses. The LoveLetter virus, technically called a worm, had cost the organization considerable expense. This worm has over 80 variants. The LoveLetter worm, written in the language Visual Basic Script, appears as LOVE-LETTER-FOR-YOU.TXT.VBS.

Administrators at the organization created a software restriction policy to block this worm from running by explicitly blocking LOVE-LETTER-FOR-YOU.TXT.VBS. They used a hash rule to prevent this script from running regardless of whether the file name changed. In the past, the organization had used VB Script files for systems management and logon scripts. By blocking all .vbs files from running, they would have protected the organization. However, they also would have penalized it because users could not have use VB Scripts for legitimate purposes.

For information about how to obtain a certificate and digitally sign files to increase the level of security in your environment, see Help and Support Center for Windows Server 2003.

# Additional Resources for Deploying a Managed Software Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Choosing an Automated Installation Method" in the *Automating and Customizing Installations Guide* in this kit, for more information about deploying and installing operating systems.
- "Designing a Group Policy Infrastructure" in this book.
- The Windows Security Collection in the Windows Server 2003 Technical Reference on the Web (http://go.microsoft.com/fwlink/?linkid=4571) for more information about Active Directory, Group Policy, and Windows Installer.
- *Designing and Deploying Directory and Security Services* of this kit for more information about Active Directory.
- "Designing RIS Installations" in *Automating and Customizing Installations* of this kit for more information about Remote Installation Services.
- "Planning Deployments" in *Microsoft® Windows® XP Professional Resource Kit Documentation* (or see "Planning Deployments" on the Web at http://www.microsoft.com/reskit).
- "Designing the Active Directory Logical Structure," in *Designing and Deploying Directory and Security Services s* of this kit.
- "Designing and Deploying File Servers" in *Planning Server Deployments* of this kit.
- "Planning a Secure Environment" in *Designing and Deploying Directory and Security Services* of this kit.
- *Microsoft® Systems Management Server Administrator's Guide* for more information about SMS.
- The Group Policy Management Console link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for a free download of GPMC.
- The Microsoft Systems Management Server link on the Web Resources page athttp://www.microsoft.com/windows/reskits/webresources.
- Windows Installer Software Development Kit (SDK) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Tools**

- Group Policy Management console (GPMC).

  The all-in-one Group Policy management tool. For a free download, see the Group Policy Management Console link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Job Aids**

- "Worksheet A.42 Assessing Software Management Tasks" (DMEUSE.42 doc) on the *Microsoft® Windows® Server 2003*

*Deployment Kit* companion CD (or see "Worksheet A.42 Assessing Software Management Tasks" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.43 Packaging Software" (DMEUSE.43 doc) on the *Windows® Server 2003 Deployment Kit* companion CD (or see "Worksheet A.43 Packaging Software" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.44 Preparing Your Network for Deploying Applications" (DMEUSE_44.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.44 Preparing Your Network for Deploying Applications" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.45 Identifying Performance Issues Between Client Computers and Software Distribution Point Servers" (DMEUSE_45.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.45 Identifying Performance Issues Between Client Computers and Software Distribution Point Servers" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.46 Evaluating Strategies for Connecting Remote Users" (DMEUSE_46.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.46 Evaluating Strategies for Connecting Remote Users" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.47 Assessing Administrative Requirements and Business Needs of the Users in Your Organization" (DMEUSE_47.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.47 Assessing Administrative Requirements and Business Needs of the Users in Your Organization" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.48 Assigning and Publishing Software" (DMEUSE_48.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.48 Assigning and Publishing Software" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.49 Using GPO Back Up Log Sheet" (DMEUSE_49.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.49 Using GPO Back Up Log Sheet" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.50 Distributing Software Using Group Policy" (DMEUSE_50.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.50 Distributing Software Using Group Policy" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.51 Targeting Software to Multilingual Users" (DMEUSE_51.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.51 Targeting Software to Multilingual Users" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.52 Creating a Pilot Plan for Software Installation" (DMEUSE_52.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.52 Creating a Pilot Plan for Software Installation" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.53 Identifying Software Installation Issues" (DMEUSE_53.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.53 Identifying Software Installation Issues" on the Web at http://www.microsoft.com/reskit).

- "Worksheet A.54 Identifying Software Maintenance Issues" (DMEUSE_54.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Worksheet A.54 Identifying Software Maintenance Issues" on the Web at http://www.microsoft.com/reskit).

# Deploying a Simple Managed Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft® IntelliMirror® management technologies and Group Policy combine the advantages of centralized computing with the performance and flexibility of distributed computing. This chapter incorporates the use of these technologies, described throughout this book, in a cohesive plan for creating a simple managed environment.

**In This Chapter**


Overview of Deploying a Simple Managed Environment

Deploying the Network Infrastructure

Deploying the Logical Infrastructure

Preparing Shared Folders

Installing Software

Setting Policies to Manage User Data and Settings

Testing and Deploying Your Configuration

Additional Resources for Deploying a Simple Managed Environment


**Related Information**

- For more information about how to use IntelliMirror technologies to create a managed environment, see "Planning a Managed Environment" in this book.
- For more information about Group Policy, see "Designing a Group Policy Infrastructure" in this book.
- For more information about user profiles, folder redirection, disk quotas, and offline files, see "Implementing User State Management" in this book.
- For more information about using Group Policy to deploy application software to your users, see "Deploying a Managed Software Environment" in this book.

# Overview of Deploying a Simple Managed Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This chapter provides a blueprint for deploying a simple managed environment. While it provides a complete end-to-end management solution for a small computing network, larger organizations can also use this chapter as a starting point for their configuration, and scale out as needed.

You can leverage the information in this chapter in two ways.

**Create a small production environment** If you are the system administrator of a small organization that wants to deploy a new network based on servers running the Microsoft® Windows® Server 2003 operating system, use the instructions here to build your initial production environment from the ground up, adding client computers and additional servers as needed. Be sure to follow all of the testing procedures included in "Testing and Deploying Your Configuration" later in this chapter, before inviting users to start their production work in the environment.

**Create a test environment** If you are part of a large organization's IT management team, use the guidelines and procedures in this chapter to build a test network for testing various implementations of Active Directory® directory service and Group Policy in Windows Server 2003. After the structure is in place and your work has been tested, staged, and deployed to your production network, you can continue to use this small network for modeling and testing new Group Policies.

# Process for Deploying a Simple Managed Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By performing the procedures in this chapter, you will deploy a domain controller (DC) within your intranet, create a logical infrastructure, assign and publish application software, and set up policies to manage your users and computers. You will use Microsoft Software Update Services (SUS) to review, approve, and test critical operating system updates before deploying them to your clients, so that you are secure in the knowledge that the clients are fully protected from the latest security risks. The steps for deploying a simple managed environment are illustrated in Figure 9.1.

**Figure 9.1 Deploying a Simple Managed Network**



# Chapter Assumptions and Sample Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A sample configuration has been devised to illustrate the various steps in creating a simple managed environment. This sample consists of a fictional organizational unit (OU) structure. You might want to go through each step using the sample data before implementing your own design; or, you might prefer to work through this guide, replacing the samples with your actual OUs and user accounts.

**Organizational Units in the Sample Configuration**

The chapter steps you through the creation of the OU structure and the users described in Table 9.1.

**Table 9.1 Example of an OU Structure for a Simple Managed Environment**

| User Name | User Requirements | Security Group | Organizational Unit | Computer Type |
|---|---|---|---|---|
| Florian Voss | Domain Administrator; roaming profile; offline files; redirected folders | Administrators<br><br>RUPUsers | IT | Mobile |
| Hung-Fu Ting<br><br>Christie Moon | No roaming profile; offline files; redirected folders; needs access to the Microsoft® Office productivity suite only. | | Bookkeeping | Desktop |
| Nicole Holliday<br><br>Tzipi Butnaru | Roaming profiles; offline files; redirected folders; needs access to Microsoft Office only. | RUPUsers | Sales | Mobile |

The following common characteristics and groupings are important to remember when applying policies (discussed in "Setting Policies to Manage User Data and Settings" later in this chapter):

- All users have redirected folders and offline files, and need access to Microsoft® Office applications.
- All users except the members of the Bookkeeping OU have roaming user profiles (RUP).
- The users in the IT and Sales OUs have mobile computing needs, but they almost always connect to the local area network (LAN) over a fast link.
- The users in the IT OU need access to the Windows Server 2003 Administrative Tools Pack.

This fictitious OU exists in a fictitious company named A. Datum Corporation with the domain name of adatum.com.

**Hardware Configuration and Assumptions**

Certain hardware and network assumptions are made in this chapter, which might not be true of your own environment. Read this chapter with the understanding that you need to reject assumptions if they do not apply to your organization.

The hardware configuration outlined here is intended to supply the basic needed components for a simple environment. The infrastructure includes one mobile client computer and one desktop client computer to demonstrate management techniques for these two common types of client computers. One server is deployed, which takes on several roles: Active Directory domain controller, Domain Name System (DNS) server, secure software distribution point using Distributed File System (DFS), and Software Update Services (SUS) distribution server. This computer also contains all of the application software, user state, and user data files for the organization.

Organizations frequently deploy a second server to host applications, SUS, and user data and settings. With this configuration, the DC can maintain the highest possible performance in its primary role. However, for a small organization, one server usually can fulfill all of these roles.

It is highly recommended that, if possible, you deploy a second DC to provide backup and redundancy. With two domain controllers, if one DC becomes unavailable, users still can log on and access their data using the second domain controller. By default, if no DC is available, only users with cached credentials can log on to the network.

Having a second DC already in place provides a much faster method of restoring service to users than having to create a new DC and populate its information from an external backup source. The second DC contains a complete and up-to-date replica of the entire Active Directory database; if the hard disk or data is corrupted on one DC, the second DC ensures that you will not have to rebuild your OU structure or repopulate your Active Directory database. The addition of a second DC should be a primary upgrade consideration for scaling your initial network and keeping it highly available.

If you choose to add a second DC to your network, consider the following:

- Replication affects network traffic. Place both DCs on the same network segment for best replication performance.
- A client computer can retrieve Group Policy objects (GPOs) from a DC to which a recent change has not yet been replicated.
- In a simple environment, it is recommended that you leave the operations master roles on one DC.

**Note**

- For complete information about deploying multiple domain controllers in an Active Directory environment, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

Because this chapter's goal is to explain how to use change and configuration management techniques to manage a simple environment, the configuration that is presented does not include a second domain controller. The following assumptions and configuration are used in this chapter:

- Your network infrastructure does not include a secondary server or DC.
- The DC has one physical drive with two partitions. The first partition is logical drive C, and it will contain the Windows Server 2003 operating system. The second partition is logical drive D, and it will contain all software and user data that is not part of the operating system.
- The CD-ROM for the DC is logical drive E.

**Minimum Hardware and Software Configurations for a Simple Managed Environment**

To begin building your simple managed environment, use the minimum hardware and software configurations listed in Table 9.2. Your hardware list will expand to fit your organization's computing needs.

**Table 9.2 Minimum Hardware and Software Configuration for a Simple Managed Environment**

| Device | Requirements |
|---|---|
| Domain controller running Windows Server 2003 | See the Windows Server 2003 System Requirements link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. |
| One desktop client computer running the Microsoft® Windows® XP Professional operating system | See the Windows XP Professional System Requirements link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. |
| One mobile client computer running Windows XP Professional | See the Windows XP Professional System Requirements link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. |
| Network Address Translator (NAT)-enabled router | Standard 5-port router recommended. |
| Broadband (cable or DSL) modem | No requirements. |

# Deploying the Network Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before implementing any management techniques, you must deploy your network infrastructure. In the chapter example, a Network Address Translation (NAT)router provides Dynamic Host Configuration Protocol (DHCP) addresses to all computers on the network, including the domain controller. Microsoft highly recommends the use of static IP addresses for domain controllers; however, in the sample configuration, the router does not continually release and renew leases, so the dynamic IP assigned to the domain controller is stable and essentially complies with that recommendation. If you want to assign a static IP address to your DC you can do so; however, this requires additional configuration that this chapter does not cover.

Figure 9.2 illustrates the steps involved in setting up the network infrastructure.

**Figure 9.2 Deploying the Network Infrastructure**



# Setting Up the Physical Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The components that you use in setting up your physical network will vary depending upon the equipment that you already have in place, your organization's specific needs, and the purpose of this network — that is, whether you are building a test LAN or an initial production LAN. The configuration documented here is that of a basic small network, which can be easily scaled to fit your computing needs. The router used here is a standard 5-port NAT router with a built-in firewall. Specific router instructions are not included, as those depend on the router that you have purchased.

**To configure your router**

Follow the directions in the documentation for your router to configure the router to these specifications:

- Ensure that port 53 on the router is enabled to support DNS. (This is the default state of many routers.)
- If the router is wireless, enable 128-bit WEP security.
- Set a strong administrator password on the router.
- Use the instructions that you received with your router to configure the router to receive its IP configuration from your ISP using DHCP (this is the default state for many routers).

**To configure your LAN-router connection**

1. Connect the LAN cable from the computer that is to be the DC to an available port on the NAT router.
2. Connect the LAN cable from the broadband modem to the WAN port on the NAT router.
3. Turn on the router and the modem.

Your network should be similar to the one illustrated in Figure 9.3.

**Figure 9.3 LAN Router Connection for a Simple Managed Environment**



# Deploying a Domain Controller Inside the Intranet

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must deploy a domain controller (DC) in order to establish a top-level domain and forest on your private LAN.

For more information about domains and forests, see "Designing a Group Policy Infrastructure" in this book, and "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

To deploy the DC, perform the following tasks:

- Install and configure the DC.
- Synchronize the timeserver for the domain controller with an external source.
- Format additional drive partitions.

# Install and Configure the Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To install and configure the domain controller, you will perform the following tasks:

- Install the Windows Server 2003 operating system.
- Install Active Directory on the domain controller, and configure the server role.
- Configure DNS.
- Install the Application Server role (Internet Information Services [IIS], ASP.NET). This step is only necessary for servers hosting Software Update Services (SUS) and is not a core requirement for a DC.

**To install Windows Server 2003**

1. Boot from your Windows Server 2003 operating system CD-ROM. Follow the instructions in the documentation for Windows Server 2003 to install the operating system on the computer that is to be your domain controller. Create disk partitions with the following properties.

| Partition | Use | File System | Size |
|---|---|---|---|
| Primary | Operating System | NTFS | At least 2 gigabytes (GB) |
| Secondary | Applications | NTFS | Remaining space |

**Note**

- If your LAN includes a second server, you can choose to create only one partition on the domain controller's hard drive, to store the operating system, and use the other server for storing additional software and data.

2. During Windows Setup, enter the following values:

- **Computer Name**: Enter **DC01**.
- **Administrator Password**: Enter a strong password.

**Important**

- Computer security requires the use of a strong password for your administrator account. A strong password has from 7 through 14 characters, and contains letters (both uppercase and lowercase), numerals, and symbols (all other characters, such as $%*&). The password should contain at least one symbol character in the second through sixth positions.
- **Network settings**: Select typical settings.
- When prompted about whether this computer is part of a Workgroup or Computer Domain, select **Workgroup** and accept the default name of Workgroup.

3. After the computer restarts, log on as Administrator.

4. Click **Start**, point to **All Programs**, and click **Activate Windows**. Follow the prompts to activate and register your copy of Windows Server 2003 through the Internet.

   If you cannot access the Internet, refer to your router and modem instructions for troubleshooting assistance.

**To configure the server as a domain controller**

1. Click **Start**, and click **Manage Your Server**. Select **Custom Configuration**. Click **Add or remove a role**, and then click **Next**. Wait for the wizard to review the computer's current configuration.

2. Select the **Domain Controller (Active Directory)** role. Proceed to run the Active Directory Installation Wizard. Use the following values as you are prompted for them:

   a. Select **Domain controller for a new domain**.

   b. Select **Domain in a new forest**.

   c. Enter your domain name (in the sample configuration, this is **adatum.com**).

   d. Accept the default values for **Domain NetBIOS name**, **Database folder**, **Log folder**, and **SYSVOL folder location**.

      Because DNS has not yet been installed on this server, the **DNS Registration Diagnostics** will indicate that none of the DNS servers used by this computer responded within the timeout interval.

   e. Select **Install and configure the DNS server on this computer, and set this computer to use this DNS server as its preferred DNS server**.

   f. Select **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**.

   g. In the **Directory Services Restore Mode Administrator Password** field, enter a strong password.

      The wizard will notify you that the computer has a dynamically assigned IP address. Typically you would not assign a dynamic IP address to a domain controller. However, this configuration is acceptable for this simple network in which the router is used as the DHCP server.

   h. When the **Local Area Connection Properties** page displays, click **Cancel**.

   i. When the wizard finishes configuring Active Directory, select **Restart Now**. After the computer has restarted, click **Finish**.

**To configure the server as a DNS server**

1. From the **Manage Your Server** screen, click **Manage this DNS server**.

2. Right-click **DC01**, click **Configure a DNS Server**, and then click **Next**.

3. Select **Create aForward Lookup Zone**.

4. Select **This server maintains the zone**. Type your domain name for the zone; for example, **adatum.com**.

5. Select **Allow dynamic updates**.

6. Select **Yes, forward queries to DNS servers with the following IP addresses**, and type the IP address of the NAT router.

7. Exit the **Manage DNS Server** snap-in.

   **Note**

   - You will receive a message that the forward lookup zone cannot be added to the server, because the zone already exists. This is because the zone was created when the DNS server role was initially configured. This message does not indicate an error condition.

**To install Application Server (IIS, ASP.NET)**

IIS is a necessary component for SUS. Use Manage Your Server to install IIS.

1. From the **Manage Your Server** screen, click **Add or remove a role**, and then click **Next**.

2. Select **Application Server (IIS, ASP.NET)**, and run the **IIS Installation** wizard. Do not load **FrontPage Server Extensions**, and do not **Enable ASP.NET** unless you need them for applications beyond the scope of this chapter.

**To confirm that you have completed installing server roles**

1. In **Manage Your Server**, verify that your server has been configured with the following roles:

   - Application Server (for IIS)
   - Domain Controller (Active Directory)
   - DNS server

2. Close the **Manage** y**our Server** program.

   **Tip**

   - With Windows Server 2003, remote administration capability is built into the domain controller server role. Use the Remote Desktop for Administration program from any other computer to administer this server. As an extra security measure, it is recommended that you require smart card security for anyone who logs onto the DC remotely. For information about smart cards, see "Deploying Smart Cards" in *Designing and Deploying Directory and Security Services* of this kit. For more information about using Remote Desktop, see "Remote Desktop for Administration" in Help and Support Center for Windows Server 2003.

# Synchronize the Time Server for the Domain Controller with an External Source

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, the primary domain controller (PDC) emulator gets its time from the BIOS clock. In a network with a single DC, that DC automatically has this role.

The PDC emulator establishes the time and date settings for all computers within its domain. If the time is not accurately set in the PDC emulator's BIOS, all computers in the domain have incorrect time and date settings.

To prevent this, you can synchronize the domain controller with an external time source such as the time servers provided by the National Institute of Standards and Technology (NIST). For a list of the names and IP addresses of NIST time servers for your area, see the National Institute of Standards and Technology (NIST) Internet Time Servers link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Note**

- Be aware that the Network Time Protocol (NTP) is unauthenticated, and unencrypted, and it is possible for an intruder to spoof the time root source, causing the wrong time to be set on the DC. You can avoid this possibility by using IPSec to secure the transmission, by accessing the time root source by its IP address rather than its fully qualified domain name, or by purchasing an NTP-capable hardware clock for your DC time synchronization.

**To synchronize the domain controller with an external time source**

1. Click **Start,** and then click **Command Prompt**.

2. In the **Command Prompt** window, type the following line, where ***peers*** is a comma-separated list of IP addresses of the appropriate time sources, and press ENTER:

   **w32tm /config /manualpeerlist: *peers* /syncfromflags:MANUAL**

   The time sources you choose depend on your time zone. For example, if your domain controller is located in the Pacific Time zone, this line might read:

   **w32tm /config /manualpeerlist:131.107.1.10 /syncfromflags:MANUAL**

   In this example, the IP address of the timeserver is used instead of the fully qualified domain name for security purposes.

3. Press ENTER. You should get a message that the command completed successfully.

4. Type **w32tm /config /update**

5. Press ENTER. You should get a message that the command completed successfully.

   W32time uses a variable poll interval based on the quality of timesync with the server. On DCs, this interval defaults to between 64 and 1024 seconds.

6. To immediately synchronize with the external time server, type **w32tm /resync** and press ENTER. You should get a message that the command completed successfully.

7. Type **Exit** and press ENTER.

# Format Additional Drive Partitions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must format additional drive partitions on the domain controller before you can use them to store software and user data.

**To format drive partitions**

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Computer Management**.

2. Expand **Storage**, and click **Disk Management**.

3. In the right pane, right-click the volume (**D:**), and click **Format**.

4. Type a volume label that reflects the contents of this partition (such as **Apps**). Accept the other formatting defaults (**NTFS**, **default allocation unit size**), and click **OK**.

5. When formatting is complete, close the **Computer Management** snap-in.

# Deploying Windows XP Professional

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For the simple managed environment described in this chapter, client computer operating systems are installed manually. Alternatively, you can automate these installations as well as configure and deploy desktop images to specified computers throughout your domain. For more information about those capabilities, see "Designing Unattended Installations" in *Automating and Customizing Installations* of this kit.

The process of deploying Windows XP Professional includes these tasks:

- Deploying the client computers and installing the operating system
- Establishing the client's IP addressing information
- Joining the client computer to the adatum.com domain
- Activating Windows and installing updates

**Note**

- If the computer on which you are installing Windows XP has small computer system interface (SCSI) drives, and you receive the error message that "Setup did not find any hard disk drives installed in your system," see article Q234054, "Error Message: Setup Did Not Find Any Hard Disk Drives Installed in Your System," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**To deploy the client computers and install the operating system**

1. Connect the client computers' Ethernet connections to available ports on the router.
2. Follow the instructions in the documentation for Windows XP Professional to install the operating system on each client computer. Enter a strong password for the Administrator account. Use the default values for most settings, entering the following values when prompted.

| Setting | Desktop Computer Value | Laptop Computer Value |
|---|---|---|
| **File System** | NTFS | NTFS |
| **User Name** | Christie Moon | Tzipi Butnaru |
| **Org Name** | A. Datum Corporation | A. Datum Corporation |
| **Computer Name** | Desktop01 | Laptop01 |
| **Network Settings** | Typical | Typical |
| **Domain/workgroup** | Workgroup | Workgroup |

**To establish the client's IP addressing information**

It is important for the client's network configuration to contain the proper addressing information. To determine its addressing, perform the following steps.

1. Click **Start**, and click **Command Prompt**.
2. Type **ipconfig /all**, and press ENTER.

3. Make a note of the IP address, the default gateway, and the first listed DNS server.

   The default gateway IP address must match the IP address of the NAT router. The first DNS server IP address must match that of DC01. If either address does not match the appropriate device, follow steps 4 through 7 to manually set the proper addresses.

4. Click Start, point to **Control Panel**, click **Network and Internet Connections**, and then click **Network Connections**.

5. Double-click **Local Area Connection**. Click **Properties**, select **Internet Protocol (TCP/IP),** and click **Properties**.

6. Select **Use the following DNS server addresses**, and type the IP address of DC01 into this field.

7. Click **Advanced**. Under **Default gateways**, click **Add** and enter the IP address of your NAT router. Click **OK**, and close all windows.

   **Tip**

   - Before joining the computer to the domain, use the TCP/IP **ping** command at the command prompt to test the connection to the IP addresses of DC01 and the NAT router in order to ensure that you can communicate to both devices.

**To join the client computer to the adatum.com domain**

1. Click **Start**, click **My Computer**, and in the left pane, click **View system information**.

2. Click **Computer Name**, and click **Change**.

3. Click **Domain**, and type **adatum.com** as the domain name. When prompted for the name and password of a user who is authorized to join this computer to the domain, type **Administrator** and enter the strong password that you supplied for DC01.

4. You will receive a **Welcome to the adatum.com domain** message. Close all windows and reboot the client computer.

**To activate Windows and install updates**

1. Log on to the client computer as its local Administrator.

2. Click **Start**, point to **All Programs**, and click **Activate Windows**. Follow the prompts to activate and register your copy of Windows XP Professional through the Internet.

3. Click **Start**, point to **All Programs**, and click **Windows Update**. Click **Pick updates to install**, then click **Review and install updates**.

4. Select and install **Windows XP SP1**. Follow the prompts to download and run the Windows XP Service Pack 1 Setup Wizard. Select the option **Do Not Archive Files**. After the service pack has been loaded, you are prompted to reboot. Do so and log on again as the local Administrator of this computer.

5. Click **Start,** point to **All Programs**, and click **Windows Update**. Click **Pick updates to install**. Remove all critical updates and service packs from the right pane (you will use Software Update Services to download approved critical updates to clients).

6. In the left pane, click **Windows XP**. In the right pane, click the **Add** button for **Microsoft .NET Framework**. Click **Review and install updates**. Follow the prompts to download and install **Microsoft .NET Framework.**

This completes the deployment of Windows XP to the client computers. The network infrastructure that you have created

should be similar to the one in Figure 9.4.

**Figure 9.4 Network Infrastructure with Clients Deployed**



# Deploying the Logical Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can manage your environment, you must create an Active Directory structure. This structure includes organizational units (OUs) for users and computers, as well as security groups for assigning file share permissions. Use Group Policy Management Console (GPMC) and the Active Directory Computers and Users snap-in to develop this structure. Figure 9.5 shows the process for deploying the logical infrastructure for your simple managed environment.

**Important**

- For all subsequent steps, you must log on to the domain controller with a domain administrator account unless otherwise instructed.

**Figure 9.5 Deploying the Logical Infrastructure**



# Installing Group Policy Management Console (GPMC)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy Management Console (GPMC) unifies management of all aspects of Group Policy. Using GPMC, you can manage all Group Policy objects (GPOs), Windows Management Instrumentation (WMI) filters, and permissions on your network.

GPMC is available for download from the Microsoft Web site. For instructions for downloading and installing GPMC, see the Group Policy Management Console link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Tip**

- To access the GPMC download site, add Microsoft.com to your trusted sites in Microsoft® Internet Explorer.

**To run GPMC**

- To open the GPMC snap-in, click **Start**, point to **Administrative Tools**, and click **Group Policy Management**.

# Creating a Basic OU Structure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to create OUs for the users and computers you want to manage within your domain. This chapter uses a simple OU structure for a simple environment. If your environment requires a more complex OU structure, refer to the following documentation for planning information:

- "Designing a Group Policy Infrastructure" in this book.
- "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

**To create the OU structure**

1. Using GPMC, expand the forest node, the domains node, and the domain **adatum.com**.
2. Right-click **adatum.com**,and click **Active Directory Users and Computers**.
3. In the **Active Directory Users and Computers** MMC, expand **adatum.com**.
4. Right-click **adatum.com**, point to **New**, and click **Organizational Unit**.
5. For the new OU name, type **IT**.
6. Repeat steps 4 and 5 to create OUs named Bookkeeping and Sales.
7. Right-click the OU named IT, point to **New**, and click **User**.
8. Type the user information for Florian Voss, as shown in Table 9.1. Set the **User logon name** as Florian, and click **Next**.
9. When prompted for a **User Password**, type in a strong password. Make sure that **User must change password at next logon** is not checked.
10. Repeat steps 7 through 9 to create the user accounts within the appropriate OUs, as shown in Table 9.1.
11. With **adatum.com** highlighted, click **Builtin** in the left pane, and double-click **Administrators**.
12. Click **Members**, and click **Add**. In the **object names** text box, type **Florian**,and click **Check Names** to resolve the entry to Florian Voss. Click **OK**.
13. In the right pane, double-click **Domain Admins**.
14. Click **Members** and repeat step 12 to add Florian to the **Domain Admins** security group. Click **OK** twice, and click **adatum.com** in the left pane.
15. Right-click **adatum.com**, point to **New**, and click **Group**.
16. Type the group name **RUPUsers**. Make sure that the group type is **Security**, and that the group scope is **Global**.
17. Double-click the security group **RUPUsers**.
18. Click the **Members** tab, and click **Add**.
19. Click **Advanced**, and click **Locations**.

20. Expand **adatum**, click **Sales**, and click **OK**.

21. To select all users in the Sales OU, click **Find Now**. To add the users to the RUPUsers security group, click **OK**.

22. Repeat steps 17 through 19 to add all of the IT OU members to the RUPUsers security group.

23. Close the **Active Directory Users and Computers** snap-in, and minimize GPMC.

Although not addressed in this chapter's procedures, it is useful to understand that new user and computer accounts are created in the CN=Users and CN=Computers containers by default. It is not possible to apply Group Policy directly to these containers, although they inherit GPOs linked to the domain. Redirusr.exe (for user accounts) and Redircomp.exe (for computer accounts) are two new tools included with Windows Server 2003 that enable you to change the default location where new user and computer accounts are created so you can more easily scope GPOs directly to newly created user and computer objects. These tools are located in %windir%\system32.

By running Redirusr.exe and Redircomp.exe once for each domain, the domain administrator can specify the OUs into which all new user and computer accounts are placed at the time of creation. This allows administrators to manage these unassigned accounts by using Group Policy before the administrators assign them to the OU in which they are finally placed. You might want to consider restricting the OUs used for new user and computer accounts using Group Policy to increase security around new accounts.

For more information about redirecting users and computers, see article 324949, "Redirecting the Users and Computers Containers in Windows Server 2003 Domains," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For information about Redirusr.exe and Redircomp.exe, see the Redirecting Users and Computers link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For additional information about these tools, see "Designing a Group Policy Infrastructure" in this book.

# Deploying Security Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is beyond the scope of this chapter to prescribe the security measures that you should take to protect your computing environment. Some things to consider are:

- Protect your physical equipment. For example, you might want to keep your domain controller computer inside a locked room or cabinet.

- Secure data as it traverses the network. If you are using Encrypting File System (EFS) to encrypt files, keep in mind that encrypted data is not encrypted when in transit over the network, but only when stored on disk. EFS does not work with roaming user profiles; if users encrypt files that are stored in their roaming profile, they will receive an error message when the profile is saved to the server. (You can, however, use IPSec in combination with EFS to protect the data when it is in transit.)

  For more information about securing data as it traverses a network, see "Deploying Security Policy" in this book and "Deploying IPSec" in *Deploying Network Services* of this kit.

- If you have a wireless network, be sure to enable 128-bit WEP protection. Consider deploying Group Policy settings specifically designed to protect your data in a wireless environment; for more information, see "Deploying Security Policy" in this book.

For more information about Group Policy and security policies, see "Designing a Group Policy Infrastructure" and "Deploying Security Policy" in this book. For best practices and considerations for securing user data and settings, see "Implementing User State Management" in this book. For more information about how to secure your Windows computing environment, see "Planning a Secure Environment" in *Designing and Deploying Directory and Security Services* of this kit.

It is recommended that you deploy your selected security measures before continuing with the tasks in this chapter. If you are not ready to do so, you can proceed with the tasks, but be sure to secure your system before you deploy in a production environment.

# Preparing Shared Folders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The next step in deploying your managed environment is to create a DFS root share to provide virtual views of the file shares. DFS is deployed to create scalable and flexible software distribution points. Next, you will prepare shared folders that will contain application software and user data (roaming user profiles and redirected folders). These steps are illustrated in Figure 9.6.

**Figure 9.6 Preparing Shared Folders**



# Creating a DFS Root Share

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Even though this chapter pertains to a very simple environment with only one server, setting up DFS when you initially install and deploy your applications can save you a lot of time later on when your organization grows. You can add additional and distributed storage capacity to the DFS namespace without changing the way your users access the data.

When creating a DFS namespace, you create either a stand-alone DFS root or a domain-based DFS root. Stand-alone DFS is an appropriate choice in cases where:

- Your organization does not use Active Directory.
- You need to create a DFS namespace and you are not part of the Domain Admins group, or company policy prevents you from delegating authority to manage a domain-based DFS namespace.
- You need to create a single namespace with more than 5,000 links. (If you can divide your links among two or more namespaces, domain-based DFS is an option.)
- You want to ensure the availability of the namespace by clustering the Distributed File System.

The chapter example sets up the DFS shares as domain-based because:

- The anticipated number of links is less than 5,000 links.
- No clustering hardware is deployed.
- Administration of DFS is centralized.
- The organization is using Active Directory, with the anticipation of using File Replication service (FRS) for replication once additional servers are added to the configuration.

For more information about DFS and FRS, see "Designing and Deploying File Servers" in *Planning Server Deployments* of this kit.

**To set up DFS**

1. In Windows Explorer, navigate to **D:**.
2. On the **File** menu, click **New**, click **Folder**, type **UserApps,** and press ENTER.
3. Right-click the **UserApps** folder, and then click **Sharing and Security**.
4. Select **Share this folder,** and then accept the default share name of **UserApps**.
5. Click **Permissions**, and with **Everyone** highlighted in the top pane, click **Remove.**
6. Click **Add**, and then type: **auth**. Click **Check Names**, and your field entry resolves to the **Authenticated Users** built-insecurity group with the default permission of **Read**. Click **OK**.
7. Click **Offline Settings**.
8. Select **Files or programs from the share will not be available offline**, and then click **OK**.
9. Minimize Windows Explorer.

10. On the **Start** menu, point to **Programs**, point to **Administrative Tools**, and then click **Distributed File System**.

11. In the left pane, right-click **Distributed File System**, and then click **New Root**.

12. In the **New root wizard**, click **Next**, select **Domain root**, and then accept **adatum.com** as the domain to host the root.

13. In the **server name** field, type **DC01**.

14. In the **Root name** field, type **Public**, and click **Next**.

15. In the **Folder to share** field, type **D:\Public**, and then click **Yes** to create the folder.

16. Click **Finish** to complete the wizard.

17. In the left pane, right-click **\\adatum.com\Public**.

18. Click **New Link**, and for the link name, type **UserApps**. In **Path to target**, type **\\dc01\UserApps**.

    This path can now be accessed as \\adatum.com\Public\UserApps\.

19. Close the **Distributed File System** snap-in.

Completion of these steps creates a DFS root and link for use as a software installation point. You can add folders and applications in UserApps, or you can add other links to other physical locations that contain software packages you want to deploy using Group Policy. Doing this insulates the system from future changes to the underlying file servers and shares.

# Creating Server Shares for Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For the chapter example, two applications are installed: Microsoft Office and the Windows Server 2003 Administrative Tools Pack. Microsoft Office will be assigned to all computers in the organization; the Administrative Tools Pack will be published to the IT OU only. For more information about assigning and publishing software using Group Policy, see "Deploying a Managed Software Environment" in this book.

**Important**

- Although it is not specifically covered in this chapter, it is highly recommended that you purchase and assign virus protection software to all computers in the domain.

**To prepare a server share for Microsoft Office**

Create a folder for Microsoft Office in the UserApps share. It inherits the file permissions from the parent folder.

1. Maximize Windows Explorer.
2. Navigate to D:\\UserApps and on the **File menu**, click **New**, click **Folder**, and type **MSOffice**. Press ENTER.

**To prepare a server share for the Administrative Tools Pack**

Create a folder for system administrator tools in the UserApps share. (These steps assume you are already in Windows Explorer on DC01.)

1. In **D:\UserApps**, on the **File menu**, click **New**, click **Folder**, and then type **AdminTools**. Press ENTER.
2. Right-click the **AdminTools** folder, and then click **Sharing and Security**.
3. Click the **Security** tab.
4. Click **Advanced**, and then clear the **Allow inheritable permissions from the parent to propagate to the object and all child objects** check box.
5. Click **Remove** to confirm the action, and then click **OK**.
6. Only the **Administrators (ADATUM\Administrators)** group is displayed in the top pane. Select this group, click **Full Control** in the bottom pane, and then click **OK**.

The preceding procedure allows full control for the **Administrators** group and allows no access for any other users.

# Creating Server Shares for User Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some of the users in the chapter example OU have roaming user profiles, and all of them have redirected folders. You must create server shares on DC01 and apply the appropriate permissions.

**To prepare a server share for roaming user profiles**

1. On \\DC01, click **Start**, and then click **Windows Explorer**.
2. Expand **My Computer**, and click drive **D:**.
3. Click **File**, click **New**, click **Folder**, and type **RUP** into the **New Folder** text box. Press ENTER.
4. Make the folder a network share. To set user permissions:
   a. Right-click the **RUP** folder, and click **Sharing and Security**.
   b. Select **Share this folder**. For the Share name, type **RUP$**.(The **$** hides the share so that it is not visible in **My Network Places**.)
   c. Click **Permissions**.
   d. Click **Add**, and type **RUP** in the **Object Names** text box. Click **Check Names** to resolve the entry to the **RUPUsers** security group. Click **OK**.
   e. With **RUPUsers** highlighted in the top pane, click **Allow** next to **Full Control** in the bottom pane. Click **OK**.
   f. Click **Offline Settings**.
   g. Select **Files or programs from the share will not be available offline**. (If you do not turn off Offline Folders for a user's profile, you can experience synchronization problems as both Offline Folders and Roaming Profiles try to synchronize the files in a user's profile.)
   h. Click the **Security** tab. Click **Add** and type **RUP** in the **Object Names** text box. Click **Check Names** and it will resolve to the **RUPUsers** security group. Click **OK**.
   i. With **RUPUsers** highlighted in the top pane, click **Allow** next to **Full Control** in the bottom pane. Click **OK**.

**To prepare a server share for redirected folders**

Create a network file share on DC01 to contain all the individual user folders that will be redirected. These steps assume you are already in **Windows Explorer** and have drive D expanded.

1. Click **File**, click **New**, click **Folder**, and type **Redir** into the **New Folder** text box. Press ENTER.
2. Make the folder a network share. To set user permissions:
   a. Right-click the **Redir** folder, and click **Sharing and Security**.
   b. Select **Share this folder.** Create the share name of **Redir$**. Click **Permissions**, and with **Everyone** highlighted in the top pane, click **Remove** in the bottom pane.
   c. Click **Add**, and type **auth** in the **Object Names** text box. Click **Check Names**, and the entry will resolve to the **Authenticated Users** built-insecurity group.

d. With **Authenticated Users** highlighted in the top pane, click **Allow** next to **Full Control** in the bottom pane. Click **OK**.

e. Click the **Security** tab. Click **Add**, and type **auth** in the **Object Names** text box. Click **Check Names**, and the entry will resolve to the built-in **Authenticated Users** security group. Click **OK**.

f. With **Authenticated Users** highlighted in the top pane, click **Allow** next to **Full Control** in the bottom pane.

g. With **Everyone** highlighted in the top pane, click **Remove** in the bottom pane. Click **OK**.

Following the preceding steps, you have created server file shares for roaming user profiles and redirected folders. You have set the permissions appropriately for the users who will need access to the shares, while denying access to other users.

# Installing Software

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Two types of software are installed on the domain controller: application software for users and Software Update Services (SUS), an administrative software package for reviewing and distributing critical and security operating system updates. Figure 9.7 illustrates the steps you will perform to install this software.

**Figure 9.7 Installing Software**



# Installing MS Office and the Administrative Tools Package

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By performing the previous procedures, you created server shares and set up DFS software distribution points. You must now perform an administrative installation of the software so that it can properly load onto client computers.

**To perform an administrative installation of Microsoft Office**

1. On the domain controller, insert the installation CD for Microsoft Office. When the Microsoft Office splash screen displays, close it.
2. Click **Start**, and then click **Command Prompt**.
3. To perform an administrative installation of the software, type **E:\setup /a**.
4. Enter the product key for your copy of Microsoft Office.
5. When prompted for the location in which to install Microsoft Office, navigate to **d:\UserApps\MSOffice**.
6. When the installation of MS Office is complete, insert the Windows Server 2003 operating system CD into the CD-ROM drive. When the splash screen is displayed, close it.

**To install the Windows Server 2003 Administrative Tools Pack**

- At the command line, type the following entry, and press ENTER:

  **COPY E:\I386\ADMINPAK.MSI D:\USERAPPS\ADMINTOOLS**

All software is ready to assign and publish using Group Policy (in "Setting Policies to Manage User Data and Settings" later in this chapter).

### Note

- Be sure that you purchase the appropriate number of licenses for the software that you install. In many cases, you need to purchase a multiuser version of an application software package in order to perform an administrative install. Check with your software vendor for more information.

# Installing Software Update Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Software Update Services (SUS) helps you collect, approve, and distribute critical operating system updates to resolve security vulnerabilities and stability issues. In this chapter, SUS is used to distribute such updates to the Windows XP-based clients in the domain. For more information about deploying SUS, see the Software Update Services Deployment White Paper link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Perform the following procedure to install SUS version 1.0 with Service Pack 1 (SP1) on DC01.

**Note**

- To use SUS on a DC, you must install SUS version 1.0 with SP1 or later.

**To install SUS 1.0 SP1**

1. To download the SUS 1.0 SP1 server component, click the Software Update Services Download Site link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, and follow the instructions.

2. When prompted to specify whether to open or save the file, click **Open**.

   SUS will be downloaded and installed into a folder named \SUS on the largest drive partition available on DC01.

3. Choose the **Typical** setup.

   After the wizard installs SUS, it provides you with a URL for administering SUS on DC01 (at http://DC01/SUSAdmin) and opens that Web page in Internet Explorer.

4. In the left pane of the SUS Administration Page, click **Set Options**.

5. Select the **Synchronize directly from the Microsoft Windows Update servers** option.

6. Select **Automatically approve new versions of previously approved updates**.

7. Accept the default for saving updates to a local folder, and click **Clear All** to clear all languages. Then select only those languages that you need in your environment. For the chapter example, check **English**. When you finish selecting languages, click **Apply**.

8. To download updates from the Microsoft Windows Update Server, in the left pane, click **Synchronize server** and then click **Synchronize Now**.

9. After downloading the updates, sort them by **Platform**. Select all Windows XP updates, and click **Approve**.

10. Click **Synchronize server** again, and click **Synchronization Schedule**. Set SUS to synchronize updates with the Microsoft Windows Update servers at 23:00 on Sunday.

11. Close the SUS Administration Page window.

You will use SUS later to download critical updates to the client computers.

**Caution**

- This configuration is appropriate when you are administering SUS locally; if you need to administer SUS remotely, you should enable secure Hypertext Transfer Protocol Secure/Secure Sockets Layer (HTTPS/SSL).

# Setting Policies to Manage User Data and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy settings determine how user state is managed in the environment described in this chapter. For information about managing user state with different types of users and different requirements, and for a fuller understanding of the range of options available for doing so, read "Implementing User State Management" in this book.

If you have questions about any specific GPO setting, you can click the **Explain** tab when that setting is being edited. The **Explain** tab provides a brief description of the setting and how to anticipate its applied result.

Figure 9.8 shows the steps involved in this part of your configuration.

**Figure 9.8 Setting Policies to Manage User Data and Settings**



# Setting Policies for All Users

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the chapter example, all users have redirected folders and offline files, and they all need access to Microsoft Office applications. Windows Update on the client computers must be configured for SUS. As an administrator, you also want to limit the amount of shared disk space that any one user can absorb. Use the following sets of procedures while logged on to DC01 to set up this part of your managed environment.

**To create a GPO for all users**

1. If GPMC is not already running on DC01, click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Group Policy Management**.

2. In GPMC, expand **adatum.com**.

3. Expand **Group Policy Objects**. Right-click **Group Policy Objects**, and click **New**. When prompted for the **New GPO** name, type **All Users**, and click **OK**.

**To set up folder redirection**

1. Right-click the **All Users** GPO, and click **Edit**.

2. Expand **User Configuration**, expand **Windows Settings**, click and expand **Folder Redirection**. Icons for the personal folders that can be redirected are displayed.

3. Right-click **My Documents** for redirection. (Note that when you redirect **My Documents**, all of its subfolders [**My Music**, **My Pictures**, and so forth] are automatically redirected as well; this is preferred.)

4. Click **Properties**, and in **Settings**, select **Basic - Redirect everyone's folders to the same location**.

5. Accept the default **Target folder location** of **Create a folder for each user under the root path**. Under the **Root Path** text box, type **\\dc01\redir$**. Click **OK**.

   This creates a unique folder for each user under **D:\Redir** on DC01.

**To configure offline files**

1. Under **User Configuration** in the left pane, expand **Administrative Templates**, expand **Network**, and click **Offline Files**.

2. In the right pane, double-click **Do not automatically make redirected folders available offline**. Click **Disabled** and click **OK**. This setting ensures that redirected folders are always available when the user is offline.

3. In the right pane, double-click **Prohibit user configuration of Offline Files**. Click **Enabled**, and click **OK**.

   This setting prevents users from disabling the offline file cache for My Documents, as well as preventing them from creating additional file caches. The use of this setting is discretionary.

4. In the right pane, double-click **Non-default server disconnect actions**. Click **Enabled**, and click **Show**.

5. In the **Show Contents** window, click **Add**, and type **DC01** as the name of the item to be added. Type **0** (zero) for the value, and click **OK** three times.

**To configure Automatic Updates (Windows Update) for SUS**

1. Expand **Computer Configuration**, expand **Administrative Templates**, expand **Windows Components**, and click **Windows Update**.

2. In the right pane, double-click **Configure Automatic Updates**, and click **Enabled**.

3. Select the following options, and click **OK**:

   - On the **Configure automatic updating** dropdown menu, select **4 - Automatically download and schedule the installation**. This choice is appropriate because no users are logging onto their computers as local administrators.

   - On the **Scheduled install day** dropdown menu, select **every Monday**.

   - On the **Scheduled install time** dropdown menu, select **21:00**. Click **OK**.

   When this time and date arrive, updates will be downloaded and installed to the target computers.

4. In the right pane, double-click **Specify intranet Microsoft update service location**. Click **Enabled**, and in the **Set the intranet update service** text box, type **http://dc01/sus**. Type the same value in the **intranet statistics server** text box, and click **OK**.

**To assign Microsoft Office to all authenticated users**

1. Under **User Configuration** in the left pane, expand **Software Settings**.

2. Right-click **Software Installation**, point to **New**, and click **Package**.

3. In the **File Open** dialog box, type **\\adatum.com\public\userapps\msoffice**, click the appropriate.msi file,and click **Open**. (The name of the .msi file is dependent on the version of Office that you are installing.)

   **Note**

   - It is important to use the UNC name of the share as described in step 3. If you navigate to the local location in lieu of entering the DFS Root share name and path, users will not be able to access the application.

4. Select **Assigned**, and click **OK**.

5. In the left pane, click **Software Installation**.

   In the right pane, Microsoft Office is displayed as an assigned deployment.

   **Note**

   - Make sure that you have the proper number of software licenses for the number of users who will use the application software.

**To set disk quotas**

You can use disk quotas to manage disk space usage on the servers containing users' data. You can specify a disk quota limit and a disk quota warning level. For example, you can set a user's disk quota limit to 200 megabytes (MB), and the disk quota warning level to 150 MB. In this case, the user can store no more than 200 MB of files on the volume. If the user stores more than 150 MB of files on the volume, a system event is logged.

For more information about disk quotas, see "Implementing User State Management" in this book.

1. Expand **Computer Configuration**, expand **Administrative Templates**, and expand **System**. Click **Disk Quotas**.

2. In the right pane, double-click **Enable disk quotas**. Click **Enabled**, and click **OK**.

3. In the right pane, double-click **Enforce disk quota limit**. Click **Enabled** and click **OK**.

4. In the right pane, double-click **Default quota limit and warning level**. Click **Enabled**. Set the default quota limit to 200 MB. Scroll down and set the **Default warning level** to 150 MB, and click **OK**.

5. In the right pane, double-click **Log event when quota limit exceeded**. Click **Enabled**, and click **OK**.

6. In the right pane, double-click **Log event when quota warning level exceeded**. Click **Enabled**, and click **OK**.

7. In the right pane, double-click **Apply policy to removable media**. Click **Disabled**, and click **OK**.

8. Close the GPO editor.

9. If you want to view the settings that you just applied to this GPO, in GPMC, click the **All Users** GPO, click **Settings**, then click **show all**.

**Warning**

- Be careful when setting disk quotas for users who have roaming profiles. If you set those users' disk quotas too low, they can run out of profile space and be unable to log off. The 200 MB storage limit is an estimate of the amount of storage that the typical user needs. You should analyze your own user data patterns before setting this limit in your production environment.

# Setting Policies for Users with Roaming User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the chapter example, all users except for the members of the Bookkeeping OU have roaming user profiles. The GPO that you create here will be linked to those users in a subsequent procedure.

**To create roaming user profile paths**

1. In GPMC, right-click **adatum.com**, and click **Active Directory Users and Computers**.
2. Expand **adatum.com**, and double-click **RUPUsers** in the right pane.
3. Click the **Members** tab, double-click **Florian Voss**, and click the **Profile** tab.
4. In the **Profile path** text box, type **\\dc01\rup$\%username%** to create a directory called Florian in the RUP server share. Click **OK**.
5. Repeat steps 3 and 4 for the other users in the RUPUsers security group.

   When each user logs on, the system will create a folder in the RUP share containing his or her roaming user profile. Each folder is identified by the user's logon name.
6. Close the **Active Directory Users and Computers** snap-in.

**To create a better user experience with roaming user profiles**

1. In GPMC, right-click **Group Policy Objects**, and click **New**. When prompted for the **New GPO** name, type **Enhance RUP**, and click **OK**.
2. Right-click the **Enhance RUP** GPO and click **Edit**.
3. Expand **Computer Configuration**, expand **Administrative Templates**, expand **System**, and click **Logon**.
4. Double-click **Always wait for the network at computer startup and logon** in the right pane.

   This GPO setting improves logon time for users with roaming user profiles.
5. Click **Enabled**, and click **OK**.

   You cannot encrypt files that are part of a roaming user profile, so a Group Policy setting is used to disable EFS for RUP users. For more information about EFS and other security options, see "Deploying Security Policies" in this chapter.
6. Under **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Public Key Policies**, and right-click **Encrypting File System**.
7. Click **Properties**. Clear the check box for **Allow users to encrypt files using Encrypting File System (EFS)**. Click **OK**, and exit the Group Policy editor.
8. Close the GPO editor.

   **Note**

   - This step disables EFS for all files on the computer, not just the files in the redirected My Documents folder. Because this example assumes that all user data is kept in My Documents, this setting in this case is appropriate. If your

users save data to other areas on their disk and they want to encrypt those files using EFS, you should not enable this setting. If files within these users' roaming profiles are inadvertently encrypted with EFS, the user will receive an error message when the system tries to save their profile to the server.

# Setting Policies for the Users in the IT OU to Access the Administrative Tools Pack

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a previous section, you created a server share with the appropriate permissions for the Windows Server 2003 Administrative Tools Pack. The final step to make this software available to administrative users is to publish it using Group Policy.

**To publish the Administrative Tools Pack**

1. In GPMC, right-click **Group Policy Objects**, and click **New**. When prompted for the **New GPO** name, type **AdminTools**, and click **OK**.

2. Right-click the **Admintools** GPO, and click **Edit**.

3. Expand **User Configuration**, and expand **Software Settings**.

4. Right-click **Software Installation**, point to **New**, and click **Package**.

5. In the **File Open** dialog box, type **\\adatum.com\public\userapps\admintools\adminpak.msi**, and click **Open**.

   **Note**

   - It is important to use the UNC name of the share as described in step 5. If you navigate to the local location in lieu of entering the DFS Root share name and path, users will not be able to access the application.

6. Select **Published**, and click **OK**.

7. To verify that the Administrative Tools Pack was published successfully, click **Software Installation** in the left pane.

   The Administrative Tools package displays in the right pane as a published deployment.

8. Close the GPO editor.

# Setting Policies for Mobile Users

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Several users in the example organization are considered mobile users: they use laptop computers and access the LAN from a variety of locations.

For more information about user types and best practices, see "Implementing User State Management" in this book.

**To create a pleasant mobile user experience**

These Group Policy settings help ensure that your mobile users have the best user experience when logging on to a managed environment.

1. In GPMC, right-click **Group Policy Objects**, and click **New**. When prompted for the **New GPO** name, type **Mobile Users**, and click **OK**.

2. Right-click the **Mobile Users** GPO, and click **Edit**.

3. Expand **Computer Configuration**, expand **Administrative Templates**, expand **System**, and click **User Profiles**.

4. In the right pane, double-click **Do not detectslow network connections**. Click **Disabled**, and click **OK**.

   This GPO setting blocks the user from disabling slow-link detection.

5. In the right pane, double-click **Timeout for dialog boxes**. Click **Enabled**, and then click **OK** to accept the default value (30 seconds).

   This GPO setting causes a message to display briefly, saying that the user is not receiving their RUP because the user is logged on over a slow link.

6. In the left pane, click **Scripts**. In the right pane, double-click **Maximum wait time for Group Policy scripts**. Click **Enabled**, and accept the default wait time of 600 seconds.

   This setting prevents the user from having to wait for longer than the specified maximum wait time for a logon script to complete execution over a slow link.

7. Close the GPO Editor.

# Testing and Deploying Your Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most organizations have change management processes in place to ensure that any new configurations or deployments made to production systems go through rigorous testing within a non-production test environment prior to migration to production. This kind of change management process is especially important for Group Policy changes, because Group Policy can affect everything from registry settings to security settings to deployed software on a workstation or server. In addition to the many configuration settings that Group Policy accommodates, Group Policy objects can be linked to a number of different scopes, and their effect can be filtered by the use of user or computer security groups.

The ability to test Group Policy objects in a pre-production test environment, and to test the various ways that they can be deployed prior to committing them to live users and computers, is critical to maintaining your Windows infrastructure.

You should consider the testing phase to be essential to any successful Group Policy deployment. For complete information about testing and staging your group policies, see "Staging Group Policy Deployments" in this book.

Figure 9.9 illustrates the steps that you will take to test and deploy your configuration.

**Figure 9.9 Testing and Deploying Your Configuration**



# Linking GPOs and Using Group Policy Modeling to View the Results

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Now that you have created your GPOs, you can use GPMC to link them to your OUs and simulate the results.

**To link a Group Policy object**

1. From within GPMC, expand and right-click the **Bookkeeping** OU.
2. Click **Link an Existing GPO**. Select the GPO named **All Users**, and click **OK**.

   In the left pane, the **All Users** GPO is displayed beneath **Bookkeeping**.

   **Tip**

   - You can simultaneously link multiple GPOs to an object by holding down the CTRL key while selecting GPOs.

**To use Group Policy Modeling to simulate the application of GPOs**

1. To simulate the results, right-click **Group Policy Modeling**, and click **Group Policy Modeling Wizard**.
2. On the **User and Computer Selection** page, click the **User** option button, and click **Browse**.
3. Expand **adatum**, click **Bookkeeping**, and click **OK**. Click **Advanced**, and click **OK** to select all users in this group.
4. Select **Skip to the final page of this wizard without collecting additional data**, and click **Next**.

   The **Summary of Selections** page displays the criteria that the wizard will use to process the simulation.
5. Click **Next**, and then click **Finish**.

   The Group Policy modeling for **adatum.com/Bookkeeping** displays on the right side of the screen.
6. To view all settings, click **show all**.

   As you scroll down the page, you will see that the All Users GPO has been applied to the Bookkeeping OU.

   Continue to link GPOs as specified in Table 9.3.

   **Table 9.3 Group Policy Objects Linked to OUs**

| OU | Linked GPOs |
|---|---|
| Bookkeeping | All Users |
| IT | All Users, Enhance RUP, Mobile Users, Admin Tools |
| Sales | All Users, Enhance RUP, Mobile Users |

7. When you have finished, the linked GPOs displayed in the left pane of GPMC should look like those in Figure 9.10.

   **Figure 9.10 GPMC Modeling Results**


Group Policy Management
File
Action
View
Window
Help
Group Policy Management
Forest: adatum.com
Domains
adatum.com
Default Domain Policy
Bookkeeping
All Users
Domain Controllers
IT
AdminTools
All Users
Enhance RUP
Mobile Users
Sales
All Users
Enhance RUP
Mobile Users
Group Policy Objects
AdminTools
All Users
Default Domain Controllers Policy
Default Domain Policy
Enhance RUP
Mobile Users
WMI Filters
Sites
Group Policy Modeling
Group Policy Results
Forest: adatum.com
Contents
Name
Domains
Sites
Group Policy ...
Group Policy ...
4 object(s)

# Testing Computers and Users

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must thoroughly evaluate your results before releasing this configuration to production users. Use this series of tests to validate your settings.

**To force an immediate update of linked GPOs, and review the results using Group Policy Results**
By default, clients running Windows XP Professional, and servers running Windows Server 2003, check for changes to Group Policy objects every 90 minutes, plus or minus a random offset. Domain controllers running Windows Server 2003 check for computer policy changes every 5 minutes.

You can manually trigger a policy refresh by using the **gpupdate** command or by rebooting the client computers. The **gpupdate** command triggers a background policy refresh on the computer from which the command is run. For more information about the **gpupdate** command, see "Designing a Group Policy Infrastructure" in this book.

1. To force the linked GPOs to immediately take effect, reboot each client computer.

2. In GPMC on DC01, right-click **Group Policy Results**, and click **Group Policy Results Wizard**. On the **Computer Selection** page, click **Another computer** and type **Laptop01** into the name text box.On the User Selection page, click **ADATUM\Florian**. The Group Policy Results report in the right pane displays the combined settings for this computer and the selected user. You can save this report to disk or print it for your records.

3. Repeat step 2 to view the resultant set of policies for other computer and user combinations.

**To test computer and user configurations**
Log on to each client computer as at least one user from each OU who would use that type of computer, as shown in Table 9.4 and Table 9.5. Perform the following steps.

1. Click **Start** and point to **All Programs**. The Office suite of programs should be listed.

2. Select Microsoft Excel, Microsoft Word, or other Office program. The software application will automatically install the first time the user accesses it. Thereafter, the program loads and runs when the user selects it.

3. When logged on as Florian, install the Administrative Tools Pack by using **Add or Remove Programs** in **Control Panel.** Those tools will then be available for Florian and other administrators from the **Start\Programs\Administrative Tools** menu.

   The ability of users to install and run these applications is a confirmation that the assigned and published software programs are working correctly.

4. Check the Event Logs on each computer to make sure that no errors have occurred.

5. To test roaming user profiles, use one of the following methods:

   - Log on, log off, log on again, and then access the share \\dc01\RUP$\%username%. The user's profile should have been copied to that share at logoff.

   - Make a simple change to a setting (for example, change the desktop background). Then log off and log on to another computer to make sure that the setting roams.

6. To test folder redirection, right-click **My Documents**, and click **Properties**. The **Targetfolder location** should point to the server share, not to the local computer.

   If everything is installed and configured properly, users should have the settings and access to the applications listed in

Table 9.4 and Table 9.5.

**Table 9.4 Results for Users Logged On to Desktop01**

| Logon As | OU | Group Policy Settings | Applications |
|---|---|---|---|
| Hung-Fu | Bookkeeping | Offline Files<br>Redirected Folders | Microsoft Office |
| Christie | Bookkeeping | Offline Files<br>Redirected Folders | Microsoft Office |

**Table 9.5 Results for Users Logged On to Laptop01**

| Logon As | OU | Group Policy Settings | Applications |
|---|---|---|---|
| Nicole | Sales | Roaming User Profiles<br>Offline Files<br>Redirected Folders<br>Mobile User Enhancements | Microsoft Office |
| Tzipi | Sales | Roaming User Profiles<br>Offline Files<br>Redirected Folders<br>Mobile User Enhancements | Microsoft Office |
| Florian | IT | Roaming User Profiles<br>Offline Files<br>Redirected Folders<br>Mobile User Enhancements | Microsoft Office<br>Windows Server 2003 Administrative Tools Pack |

# Backing Up the Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The importance of safeguarding your data by backing it up onto removable media cannot be overstated. Obtain a tape backup, writeable CD-ROM drive, or other suitable backup device for backing up your server and user data.

In the chapter example, user data is redirected to DC01, so backups of individual client machines are not necessary. Your own environment might require the ability to back up client hard drives as well as servers. After determining your backup needs, you can use the Backup program included in Windows Server 2003, or choose a third-party software program, to back up your data.

In addition to a backup device and software, you need a good backup plan. Many organizations perform a full server backup weekly, and then perform incremental backups (backing up only those files that have changed) each day.

Rotate your backup media. Using the backup scheme described in the preceding paragraph, you will need at least one set of backup media for each day of the week as well as a separate set of media for the fully weekly backups. Keep at least three separate sets on hand for rotation of the backups. In other words, use backup set 1 for one week; backup set 2 for the second week, and so forth.

Once a month, make a full system backup and data backup, and store the backup media offsite. In addition, make a system backup of your Active Directory database and the Sysvol folder on your DC — you will need this backup if you ever need to recover your Active Directory or Group Policy information. All of these backups should be considered permanent system archives, and not included in your rotation plan.

If you use tape cartridges for backup, re-tension the tapes after several uses. Replace tapes after six months of use. Tape erodes.

Always store your backup media offsite in a secure location — perhaps a safety deposit box or another office. In the event of disaster (flood or fire, for example), you want to be sure that you can recover all of your data.

In the chapter example, DC01 hosts Active Directory and all of its associated objects, as well as SUS, application software, RUPs, and redirected folders. Therefore, the information that must be backed up is stored in multiple locations on the computer. The subsequent procedures for backing up DC01 give you the basic commands for performing the various types of backup.

**To back up the entire contents of DC01**

1. Click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and click **Backup**.
2. Select **Backup Files and Settings**.
3. Select **All Information on this Computer**.
4. Designate where to save the data, click **Next**, and then click **Finish**.

**To back up Active Directory and Group Policy**

1. Click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and click **Backup**.
2. Select **Backup Files and Settings**.
3. Select **Let me choose what to back up**.
4. Navigate through drive C, and select the **System Volume** folder.
5. Scroll down, and click **System State**.
6. Designate where to save the data, click **Next**, and then click **Finish**.

These steps back up the current state of the domain controller (including Active directory settings and data) and Group Policy settings.

**To back up GPOs**

It is very useful to make a backup of your GPOs, in the event that you need to restore only those settings, or want to copy them to another forest.

1. Click **Start**, and then click **GPMC**.

2. Expand **Forests**, expand **Domains**, and expand **adatum.com**.

3. Right-click **Group Policy Objects**, and click **Back Up All**.

4. Click **Browse**, and browse to the preferred storage location.

   This might be a secure area on your server that will then be backed up onto tape, or you might back the Group Policy settings up directly on removable media.

5. Click **Backup**.

   Note that you can also back up GPOs individually.

   Another option for backing up GPOs in a domain is to use the BackupAllGPOs.wsf script, located in the %programfiles%\gpmc\scripts folder. The syntax for using the script requires that you provide the backup location and (optionally) a comment. For example, to back up GPOs in the adatum.com domain, you would enter the following command:

   ```
   cscript BackupAllGPOs.wsf \\dc01\GPO-Backups /comment:"Weekly backup"
   ```

**To back up data only**

1. Click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and click **Backup**.

2. Select **Backup Files and Settings**.

3. Select **Let me choose what to back up**.

4. Navigate through drive C, and select **Redir**.

   This folder contains all the redirected folders from the users. In the chapter example, all data is stored in these folders, so by backing up **Redir**, you are making a full data backup.

5. Designate where to save the data, click **Next**, and then click **Finish**.

# Deploying Your Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you are completely satisfied with the results of your configuration, and have taken the time to back up the server, proceed to deploying the managed solution in your production environment. Be prepared to answer user questions regarding RUP and redirected folders. Explain disk quotas to users, so that they know what to do if they need additional shared storage space.

If the LAN that you created will be your new production environment, start migrating your users over to the LAN. If you have previous Windows-based clients and want to migrate existing user settings to the new environment, see "Migration of User State" in this book. You can scale this production network as your organization grows.

If this LAN will be a test environment, you can now back up your GPOs and restore them to your production environment. You can use the knowledge that you gained here to create new management solutions in that environment.

Deployment of a management solution is never a one-stop process. As your environment grows and changes, you will need to make adjustments in accommodation. The environment that this chapter develops is intended to give you the basic structure that will ease system management responsibilities as well as scale out to future demands.

# Additional Resources for Deploying a Simple Managed Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Additional Resources**

- "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit for information about designing and deploying your Active Directory infrastructure.
- "Deploying DNS" in *Deploying Network Services* of this kit for information about DNS.
- "Planning a Managed Environment" in this book for information about using IntelliMirror technologies to create a managed environment.
- "Deploying a Managed Software Environment" for information about using using Group Policy to deploy application software to your users.
- "Designing a Group Policy Infrastructure" in this book for information about Group Policy.
- "Staging Group Policy Deployments" in this book for information about staging and testing your group policies.
- "Deploying Security Policy" in this book and "Planning a Secure Environment" in *Designing and Deploying Directory and Security Services* of this kit for information about security.
- "Deploying IPSec" in *Deploying Network Services* of this kit for information about securing your data as it travels across the network.
- "Overview of Deploying IIS 6.0" in *Deploying Internet Information Services (IIS) 6.0* of this kit for information about Internet Information Services (IIS).
- "Designing and Deploying File Servers" in *Planning Server Deployments* of this kit for information about file servers, DFS, and FRS.
- "Implementing User State Management" in this book for information about user profiles, folder redirection, disc quotas, and offline files.
- "Designing RIS Installations" in *Automating and Customizing Installations* of this kit for information about using RIS to deploy client operating systems and desktop images.

# Deploying Network Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Welcome to *Deploying Network Services* of the *Microsoft Windows Server 2003 Deployment Kit*.

You can use the deployment guidelines in this book to build a secure core network infrastructure that uses Transmission Control Protocol/Internet Protocol (TCP/IP), Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), Windows Internet Name Service (WINS), Microsoft® Internet Security and Acceleration (ISA) Server, Windows Firewall, and Internet Protocol security (IPSec).

You can also use the deployment guidelines in this book to extend your core network infrastructure by deploying a network access infrastructure supporting a variety of network access methods, including dial-up, virtual private network (VPN), wireless, and switch access. You can provide centralized authentication, authorization, auditing, and accounting for all of these network access types by using Remote Authentication Dial-In User Service (RADIUS).

This book contains the following chapters:

Part I: Core Network Infrastructure

Designing a TCP/IP Network

Deploying DHCP

Deploying Domain Name System (DNS)

Deploying WINS

Deploying ISA Server

Deploying IPSec

Deploying Windows Firewall

Part II: Network Access Infrastructure

Deploying Internet Authentication Service (IAS)

Deploying Dial-up and VPN Remote Access Servers

Deploying Remote Access Clients Using Connection Manager

Connecting Remote Sites

Deploying a Wireless LAN

# Designing a TCP/IP Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The TCP/IP protocol suite defines industry standard networking protocols for data networks, including the Internet. Determining the best design and implementation of your TCP/IP network ensures optimal reliability, availability, scalability, security, and performance for your enterprise. You can also start to explore the next generation of the Internet layer protocol of the TCP/IP protocol suite — IP version 6 (IPv6) — by introducing Microsoft® Windows® Server 2003 IPv6 into part of your IPv4 network.

**In This Chapter**


Overview of Designing a TCP/IP Network

Planning the IP-Based Infrastructure

Developing Routing Strategies

Designing an IP Addressing Scheme

Planning an IP Configuration Strategy

Planning Security

Improving Availability

Planning IP Multicasting

Introducing IPv6 on Your Network

Testing Your Design

Additional Resources for Designing a TCP/IP Network


**Related Information**

- For more information about IP configuration strategies using Dynamic Host Configuration Protocol (DHCP), see "Deploying DHCP" in this book.
- For more information about using Domain Name System (DNS) for name resolution, see "Deploying Domain Name System (DNS)" in this book.
- For more information about using Windows Internet Name Service (WINS) for name resolution in networks that support clients running Microsoft® Windows NT®, see "Deploying WINS" in this book.

# Overview of Designing a TCP/IP Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing your IP deployment includes deciding how you want to implement IP in a new environment, or — for most organizations — examining your existing infrastructure and deciding what to change. Windows Server 2003 TCP/IP, the most widely used networking protocol, can connect different types of systems, provide a framework for client/server applications, and give users access to the Internet. TCP/IP is included in the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; and Windows® Server 2003, Web Edition operating systems.

Before you start the TCP/IP design process, inventory your hardware and software and create or update a map of your network topology. Preparing an inventory and network map can save time and help you focus on the design decisions you want to address. After you review your existing network, you might upgrade several servers to Windows Server 2003 in order to take advantage of end-to-end support for TCP/IP, or you might decide to redesign your entire network to improve its efficiency and prepare for the future of IP networking. Determine which design tasks are relevant to your environment, and then decide what changes you want to make to your network. For more information about creating a hardware and software inventory and a network topology map, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

To start the TCP/IP design process, you must make a number of design decisions about your network infrastructure. For enterprise-wide scalability, you might decide to plan your IP infrastructure based on a hierarchical network design model. You must also choose between hardware and software-based routers, and decide where to use static routing or dynamic routing protocols. You must carefully design a structured model for IP address assignment that fits your current networking environment and that accommodates expected growth. Your model can use either public or private addresses, or you can use a combination of public and private addresses.

In addition, consider security issues for an IP network, including where best to use Internet Protocol security (IPSec) and which options are appropriate for securing your perimeter network. For higher availability and load balancing, you can include redundancy in your network design. Decide whether you need to use technology enhancements such as IP multicast to optimize server workload and network bandwidth. You might start deploying IPv6 on certain network servers or clients, and, if so, decide how you want to implement IPv6/IPv4 coexistence.

After you develop your network design, you can use the remaining chapters in this book as a guide for deploying core features, such as DHCP, DNS, and WINS, as well as optional technologies, such as support for mobile or home users, connecting remote sites, or deploying wireless solutions.

# Process for Designing a TCP/IP Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 1.1 shows the design stages involved in deploying TCP/IP. Although the figure lists the stages sequentially, you must consider each topic in relation to the others rather than as a linear step-by-step process.

**Figure 1.1 Designing a TCP/IP Network**



# Windows Server 2003 TCP/IP Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 TCP/IP enables enterprise networking and connectivity on computers running Windows Server 2003, Microsoft® Windows® XP, Windows® 2000, Windows NT®, Windows® Millennium Edition, Windows® 98, and Windows® 95.

# Benefits of Windows Server 2003 TCP/IP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Using TCP/IP in a Windows Server 2003 configuration offers the following advantages:

- **Enables the most widely used network protocol.** Windows Server 2003 TCP/IP is a complete, standards-based implementation of the most widely accepted networking protocol in the world. IP is routable, scalable, and efficient. IP forms the basis for the Internet, and it is also used as the primary network technology on most major enterprise networks in production today. You can configure computers running Windows Server 2003 with TCP/IP to perform nearly any role that a networked computer requires.

- **Connects dissimilar systems.** Although all modern networking operating systems offer TCP/IP support, Windows Server 2003 TCP/IP provides the best platform for connecting Windows–based systems to earlier Windows systems and to non-Windows systems. Most standard connectivity utilities are available in Windows Server 2003 TCP/IP, including the File Transfer Protocol (FTP) program, the Line Printer (LPR) program, and Telnet, a terminal emulation protocol.

- **Provides client/server framework.** Windows Server 2003 TCP/IP provides a cross-platform client/server framework that is robust, scalable, and secure. Windows Server 2003 TCP/IP offers the Windows Sockets programming interface, which is ideal for developing client/server applications that can run on Windows Sockets–compliant TCP/IP protocol implementations from other vendors.

- **Provides access to the Internet.** Windows Server 2003 TCP/IP can provide users with a method of gaining access to the Internet. A computer running Windows Server 2003 can be configured to serve as an Internet Web site, it can function in a variety of other roles as an Internet client or server, and it can use nearly all of the Internet-related software available today.

# Planning the IP-Based Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To create or expand an enterprise network, you can choose from many design models, including a network infrastructure model based on the *three-tier design model*. This model, a hierarchical network design model described by Cisco Systems, Inc. and other networking vendors, is widely used as a reference in the design of enterprise networks.

Figure 1.2 shows the tasks involved in creating a three-tier TCP/IP infrastructure.

**Figure 1.2 Planning the IP-Based Infrastructure**




The modular nature of a hierarchical model such as the three-tier model can simplify deployment, capacity planning, and troubleshooting in a large internetwork. In this design model, the tiers represent the logical layers of functionality within the network. In some cases, network devices serve only one function; in other cases, the same device may function within two or more tiers.

The three tiers of this hierarchical model are referred to as the *core*, *distribution*, and *access tiers*. Figure 1.3 illustrates the relationship between network devices operating within each tier.

**Figure 1.3 Three-Tier Network Design Model**




**Designing the Access Tier**

The *access tier* is the layer in which users connect to the rest of the network, including individual workstations and workgroup servers. The access tier usually includes a relatively large number of low- to medium-speed access ports, whereas the distribution and core tiers usually contain fewer, but higher-speed network ports. Design the access tier with efficiency and

economy in mind, and balance the number and types of access ports to keep the volume of access requests within the capacity of the higher layers.

**Designing the Distribution Tier**

The *distribution tier* distributes network traffic between related access layers, and separates the locally destined traffic from the network traffic destined for other tiers through the core.

Network security and access control policies are often implemented within this tier. Network devices in this layer can incorporate technologies such as firewalls and address translators.

The distribution tier is often the layer in which you define subnets; through the definition of subnets, distribution devices often function as routers. Decisions about routing methods and routing protocols affect the scalability and performance of the network in this tier.

A server network in the distribution layer might house critical network services and centralized application servers. Computers running Windows Server 2003 can be used there to run the Active Directory® directory service, DNS, DHCP, and other core infrastructure services.

**Designing the Core Tier**

The *core tier* facilitates the efficient transfer of data between interconnected distribution tiers. The core tier typically functions as the high-speed backbone of the enterprise network. This tier can include one or more building-wide or campus-wide backbone local area networks (LANs), metropolitan area network (MAN) backbones, and high-speed regional wide area network (WAN) backbones.

The primary design goal for the core is reliable, high-speed network performance. As a general rule, locate any feature that might affect the reliability or performance of this tier in an access or distribution tier instead.

Select highly reliable network equipment for the core tier, and design a fault-tolerant core system whenever possible. Many products meet these criteria, and most major network vendors offer complete solutions to meet the requirements of the core tier.

For more information about designing a three-tier network model, see "Additional Resources for Designing a TCP/IP Network" later in this chapter.

# Developing Routing Strategies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After planning your network infrastructure based on your design model, plan how to implement routing. Figure 1.4 shows the tasks involved in developing a unicast routing strategy. For information about IP multicast routing, see "Planning IP Multicasting" later in this chapter.

**Figure 1.4 Developing a Routing Strategy**




To plan an effective routing solution for your environment, you must understand the differences between hardware routers and software routers; static routing and dynamic routing; and distance vector routing protocols and link state routing protocols.

# Choosing Hardware or Software Routing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A *router* is a device that holds information about the state of its own network interfaces and contains a list of possible sources and destinations for network traffic. The router directs incoming and outgoing packets based on that information. By projecting network traffic and routing needs based on the number and types of hardware devices and applications used in your environment, you can better decide whether to use a dedicated hardware router, a software-based router, or a combination of both. Generally, dedicated hardware routers handle heavier routing demands best, and less expensive software-based routers are sufficient to handle lighter routing loads.

A software-based routing solution, such as the Windows Server 2003 Routing and Remote Access service, can be ideal on a small, segmented network with relatively light traffic between subnets. Conversely, enterprise network environments that have a large number of network segments and a wide range of performance requirements might need a variety of hardware-based routers to perform different roles throughout the network.

# Choosing Static or Dynamic Routing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Routing can be either static or dynamic, depending on how routing information is generated and maintained:

- In *static routing*, routing information is entered manually by an administrator and remains constant throughout the router's operation.

- In *dynamic routing*, a router is configured to automatically generate routing information and share the information with neighboring routers.

You must decide where best to implement each type of routing.

# Static Routing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In static routing, a network administrator enters static routes in the routing table manually by indicating:

- The network ID, consisting of a destination IP address and a subnet mask.
- The IP address of a neighboring router (the next *hop*).
- The router interface through which to forward the packets to the destination.

Static routing has significant drawbacks. Because a network administrator defines a static route, errors are more likely than with a dynamically assigned route. A simple typographical error can create chaos on the network. An even greater problem is the inability of a static route to adapt to topology changes. When the topology changes, the administrator might have to make changes to the routing tables on every static router. This does not scale well on a large internetwork.

However, static routing can be effective when used in combination with dynamic routing. Instead of using static routing exclusively, you can use a static route as the redundant backup for a dynamically configured route. In addition, you might use dynamic routing for most paths but configure a few static paths where you want the network traffic to follow a particular route. For example, you might configure routers to force traffic over a given path to a high-bandwidth link.

# Dynamic Routing Protocols

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Conceptually, the dynamic routing method has two parts: the routing protocol that is used between neighboring routers to convey information about their network environment, and the routing algorithm that determines paths through that network. The protocol defines the method used to share the information externally, whereas the algorithm is the method used to process the information internally.

The routing tables on dynamic routers are updated automatically based on the exchange of routing information with other routers. The most common dynamic routing protocols are:

- Distance vector routing protocols
- Link state routing protocols

Understanding how these protocols work enables you to choose the type of dynamic routing that best suits your network needs.

**Distance Vector Routing Protocols**

A *distance vector* routing protocol advertises the number of hops to a network destination (the *distance*) and the direction in which a packet can reach a network destination (the *vector*). The distance vector algorithm, also known as the *Bellman-Ford algorithm*, enables a router to pass route updates to its neighbors at regularly scheduled intervals. Each neighbor then adds its own distance value and forwards the routing information on to its immediate neighbors. The result of this process is a table containing the cumulative distance to each network destination.

Distance vector routing protocols, the earliest dynamic routing protocols, are an improvement over static routing, but have some limitations. When the topology of the internetwork changes, distance vector routing protocols can take several minutes to detect the change and make the appropriate corrections.

One advantage of distance vector routing protocols is simplicity. Distance vector routing protocols are easy to configure and administer. They are well suited for small networks with relatively low performance requirements.

Most distance vector routing protocols use a hop count as a routing metric. A *routing metric* is a number associated with a route that a router uses to select the best of several matching routes in the IP routing table. The *hop count* is the number of routers that a packet must cross to reach a destination.

Routing Information Protocol (RIP) is the best known and most widely used of the distance vector routing protocols. RIP version 1 (RIP v1), which is now outmoded, was the first routing protocol accepted as a standard for TCP/IP. RIP version 2 (RIP v2) provides authentication support, multicast announcing, and better support for classless networks. The Windows Server 2003 Routing and Remote Access service supports both RIP v1 and RIP v2 (for IPv4 only).

Using RIP, the maximum hop count from the first router to the destination is 15. Any destination greater than 15 hops away is considered unreachable. This limits the diameter of a RIP internetwork to 15. However, if you place your routers in a hierarchical structure, 15 hops can cover a large number of destinations.

**Link State Routing Protocols**

Link state routing protocols address some of the limitations of distance vector routing protocols. For example, link state routing protocols provide faster convergence than do distance vector routing protocols. Convergence is the process by which routers update routing tables after a change in network topology — the change is replicated to all routers that need to know about it. Although link state routing protocols are more reliable and require less bandwidth than do distance vector routing protocols, they are also more complex, more memory-intensive, and place a greater load on the CPU.

Unlike distance vector routing protocols, which broadcast updates to all routers at regularly scheduled intervals, link state routing protocols provide updates only when a network link changes state. When such an event occurs, a notification in the form of a link state advertisement is sent throughout the network.

The Windows Server 2003 Routing and Remote Access service supports the Open Shortest Path First (OSPF) protocol, the best known and most widely used link state routing protocol. OSPF is an open standard developed by the Internet Engineering Task Force (IETF) as an alternative to RIP. OSPF compiles a complete topological database of the internetwork. The shortest path first (SPF) algorithm, also known as the *Djikstra algorithm*, is used to compute the least-cost path to each destination. Whereas RIP

calculates cost on the basis of hop count only, OSPF can calculate cost on the basis of metrics such as link speed and reliability in addition to hop count.

Unlike RIP, OSPF can support an internetwork diameter of 65,535 (assuming that each link is assigned a cost of 1). OSPF transmits multicast frames, reducing CPU usage on a LAN. You can hierarchically subdivide OSPF networks into areas, reducing router memory overhead and CPU overhead.

Like RIP v2, OSPF supports variable length subnet masks (VLSM) and noncontiguous subnets. For information about variable length subnet masks and noncontiguous subnets, see "Creating a Structured Address Assignment Model" later in this chapter.

**Selecting the Appropriate Routing Protocol**

Select a routing protocol based on the following considerations:

- For a small, simple network that is not expected to grow, use a simpler distance vector routing protocol like RIP v2. For a large, complex internetwork, use a newer, more sophisticated link state routing protocol like OSPF.

- Use RIP v2 or OSPF if you need to support variable length subnet masks. Although the outdated RIP v1 is still widely used in private networks, it does not support VLSM and thus is not well suited for enterprise networks. For more information about VLSM, see "Planning Variable Length Subnet Masks (VLSM)" later in this chapter.

# Designing an IP Addressing Scheme

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before assigning addresses, design an IP addressing scheme that meets the requirements of your networking infrastructure. Figure 1.5 shows the tasks involved in designing your IP addressing system, including planning your address assignment model, address allocation, and public or private addressing. Most organizations choose to use classless IP addressing, classless IP routing protocols, and route summarization.

**Figure 1.5 Designing an IP Addressing Scheme**




For information about IP multicast addressing, see "Planning IP Multicasting" later in this chapter.

# Creating a Structured Address Assignment Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can ease the burden of enterprise internetwork administration by designing a structured address assignment model. A structured address assignment model makes troubleshooting easier and more systematic and helps you interpret network maps and locate specific devices. It also simplifies the use of network management software. For enterprise scalability, assign address blocks hierarchically.

The structured address assignment model reflects more than just hierarchical concerns. To maximize network stability and scalability, assign a block of addresses based on a physical network rather than on membership within a department or team, to avoid complications when you move a workstation to a new location. For more information about address allocation as it relates to your IP addressing scheme, see "Choosing an Address Allocation Method" later in this chapter.

As a general rule, assign static addresses to routers and servers, and assign dynamic addresses to workstations. This scheme minimizes manual addressing, reducing the chances of address duplication and stabilizing the network's addressing structure. You can assign meaningful numbers when using static addresses; for example, reserve host addresses in the low or high portion of the range, and manually assign these addresses to routers or servers.

To design a structured model for assigning addresses:

- Plan classless IP addressing.
- Plan classless routing.
- Use route summarization.
- Plan variable length subnet masks (VLSM).
- Plan supernetting and classless interdomain routing (CIDR).

# Planning Classless IP Addressing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Classless IP addressing makes traditional classful IP addressing methods — restricted to the standard IP address classes in their default formats — out of date for enterprise networks. Of the five address classes, Class A, B, and C addresses, collectively known as *IPv4 unicast addresses*, are assigned to specific devices on an IPv4 network. Class D addresses, known as *multicast addresses*, are used for IP multicasting (simultaneously sending a message to more than one network destination). Class E addresses are reserved for experimental purposes.

To be able to use subnetting or supernetting, you must first understand the default formats of the unicast addresses. Unicast addresses have the following formats:

- All 32-bit IPv4 addresses contain four octets of 8 bits each, often represented as four decimal numbers separated by dots (known as *dotted decimal notation*).
- In Class A addresses, the first byte, or octet, represents the network ID, and the three remaining bytes are used for node addresses.
- In Class B addresses, the first 2 bytes represent the network ID, and the last 2 bytes are used for nodes.
- In Class C addresses, the first 3 bytes are used for the network ID, and the final byte is used for nodes.

Without some means of subdividing class-designated networks, all available IP addresses would have been depleted long ago. Classless IP addressing, which allows subnetting, was developed to handle this problem.

**Determining the Number of Subnets and Hosts**

To better use the address space, instead of using the unicast addresses in their default formats, you can use *subnet addressing*, which lets you "borrow" additional bits from the host part of the address to divide the network into subnets. In subnetting, the *subnet mask* consists of the octets assigned to the network plus the bits added for the subnet. You can use subnet mask notation to indicate these leftmost contiguous bits.

For example, for a Class B address, which has a default subnet mask of 255.255.0.0, you might allocate an additional 8 bits for subnets. That is, for a Class B address such as 131.107.65.37, you can use the following subnet mask, shown in both decimal and binary notation.

| Subnet Mask in Decimal Notation | Subnet Mask in Binary Notation |
|---|---|
| 255.255.255.0 | 11111111 11111111 11111111 00000000 |

By using 8 host bits for subnetting, you obtain 256 (that is, $2^8$) subnetted network IDs (subnets), supporting as many as 254 hosts per subnet. The number of hosts per subnet is 254 because 8 bits ($2^8$ minus 2) are reserved for the host ID. You subtract 2 because subnetting rules exclude the host IDs consisting of all ones or all zeros.

An alternative to subnet mask notation is the *network prefix length notation*. A *network prefix* is shorthand for a subnet mask, expressing the number of high-order bits that constitute the subnetted network ID portion of the address in the format *<IP address>/<# of bits>*, where *# of bits* defines the network/subnet part of the IP address, and the remaining bits represent the host ID portion of the address.

The following is the network prefix length notation for the Class B address in the previous example:

131.107.65.37/24

The bit notation "/24" refers to the number of high-order bits set to 1 in the binary notation for the subnet mask, leaving 8 bits for hosts (the eight bits set to 0).

**Note**

- IPv6 supports only network prefix length notation. It does not support dotted decimal subnet masks. For more information about IPv6, see "Introducing IPv6 on Your Network" later in this chapter.

By contrast, if you anticipate needing only 32 subnets rather than 256, each of the 32 subnets can support as many as 2,046 hosts ($2^{11}$ minus 2). That subnet mask has the following decimal and binary notations.

| Subnet Mask in Decimal Notation | Subnet Mask in Binary Notation |
|---|---|
| 255.255.248.0 | 11111111 11111111 11111000 00000000 |

The following network prefix length notation indicates the 21 bits needed to create as many as 32 subnets:

131.107.65.37/21.

Again, "/21" indicates the number of high-order bits set to 1 in binary notation, leaving 11 bits (the 11 zeros) for the host ID portion of the address.

To determine the appropriate number of subnets versus hosts for your organization's network, consider the following:

- **More subnets.** Allocating more host bits for subnetting supports more subnets but fewer hosts per subnet.
- **More hosts.** Allocating fewer host bits for subnetting supports more hosts per subnet, but limits the growth in the number of subnets.

For an introduction to TCP/IP, including information about subnetting, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Planning Classless Routing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Organizations today typically implement classless routing solutions. With classful routing protocols, IP hosts and routers recognize only the network address designated by the standard address classes. An IP host device or a router using a classful protocol such as RIP v1 cannot recognize subnets.

Classless routing protocols extend the standard Class A, B, or C IP addressing scheme by using a subnet mask or mask length to indicate how routers must interpret an IP network ID. Classless routing protocols include the subnet mask along with the IP address when advertising routing information. Subnet masks representing the network ID are not restricted to those defined by the address classes, but can contain a variable number of high-order bits. Such subnet mask flexibility enables you to group several networks as a single entry in a routing table, significantly reducing routing overhead. In addition to RIP v2 and OSPF, described earlier, classless routing protocols include Border Gateway Protocol version 4 (BGP4) and Intermediate System to Intermediate System (IS-IS).

If your network contains routers that support only RIP v1 and you want to upgrade from classful to classless routing, upgrade the RIP v1 routers to support RIP v2 or use another protocol such as OSPF. For example, you might use VLSM to implement subnets of different sizes or CIDR to implement supernetting. (VLSM and CIDR are described later in this chapter.)

**Planning Classless Noncontiguous Subnets**

One reason that classful routing is out of date is that classful routing protocols cannot reliably handle noncontiguous subnets of a subnetted class-based network ID. As mentioned earlier, classful routing protocols recognize only those networks indicated by an address class. Because classful protocols do not transmit subnet mask or prefix length information, noncontiguous subnets, when summarized by a classful routing protocol, can have the same class-based network ID.

**Noncontiguous subnets with classful routing**

Noncontiguous subnets occur when another network with a different network ID separates subnets of a classful network. For example, the two routers in Figure 1.6 separate two subnets that each use the base prefix 10.0.0.0/8, which is a Class A private network. A segment of another class-based network connects the two routers. (For more information about private addresses, see "Choosing Public or Private Addresses" later in this chapter.)

**Figure 1.6 Classful Routing Not Appropriate for Noncontiguous Subnets**




Each router in Figure 1.6 must use a subnet mask to look up a match in the routing table. Because a classful address, by definition, has only its class-based default subnet mask, the router uses the network mask that corresponds to the class of the subnet ID when advertising the route for the subnet. With classful routing, each of the routers in Figure 1.6 summarizes and advertises the class-based network ID of 10.0.0.0/8, resulting in two routes to 10.0.0.0/8, each of which might have a different metric. Therefore, a packet meant for one subnet could be incorrectly routed to the other subnet. In the figure, the arrows represent the routes advertised by the routers.

**Noncontiguous subnets with classless routing**

Figure 1.7 also shows an unrelated network connecting two noncontiguous subnets. In this example, using classless routing, the locations on the noncontiguous subnets are unambiguous because the classless protocol includes a subnet mask when advertising the route. Routers in the intermediate network can distinguish between the two noncontiguous subnets.

**Figure 1.7 Classless Routing Appropriate for Noncontiguous Subnets**



# Using Route Summarization

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With *route summarization,* or *aggregation*, in a hierarchical routing infrastructure, one route in a routing table represents many routes. A routing table entry for the highest level (the network) is also the route used for subnets and sub-subnets. In contrast, in a flat routing infrastructure, the routing table on every router in the network contains an entry for each network segment. When you use flat routing, the network IDs have no network/subnet structure and cannot be summarized. RIP-based Internet Packet Exchange (IPX) internetworks use flat network addressing and have a flat routing infrastructure.

Using route summarization, you can contain topology changes occurring in one area of the network within that area. Route summarization simplifies routing tables and reduces the exchange of routing information, but it requires more planning than does a flat routing infrastructure.

To support route summarization, your IP addressing scheme must meet the following requirements:

- Classless routing protocols (those including subnet mask or prefix length information along with the IP address) must be used.
- All IP addresses used in route summarization must share identical high-order bits.
- The length of the prefix can be any number of bits up to 32 (for IPv4).

# Planning Variable Length Subnet Masks (VLSM)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Variable length subnet masks (VLSMs) allow you to use different prefix lengths at different locations so that subnets of different sizes can coexist on the same network. Instead of using one subnet mask throughout the network, you apply several masks to the same address space, producing subnets of different sizes. For example, given the Class B network ID of 131.107.0.0, you can configure one subnet with as many as 32,766 hosts, 15 subnets with as many as 2,046 hosts, and 8 subnets with as many as 254 hosts.

**Tip**

- When using VLSM, do not accidentally overlap blocks of addresses. If possible, start with equal-size subnets and then subdivide them.

VLSM also can be used when a point-to-point WAN link connects two routers. One way to handle such a WAN link is to create a small subnet consisting of only two addresses. Without VLSM, you might divide a Class C network ID into an equal number of two-address subnets. If only one WAN link is in use, all the subnets but one serve no purpose, wasting 252 addresses.

Alternatively, you can divide the Class C network into 16 workgroup subnets of 14 nodes each by using a prefix length of 28 bits (or, in subnet mask terms, 255.255.255.240). By using VLSM, you can then subdivide one of those 16 subnets into 8 smaller subnets, each supporting only 2 nodes. You can use one of the 8 subnets for your existing WAN link and reserve the remaining 7 subnets for similar links that you might need in the future. To accomplish this act of sub-subnetting by using VLSM, use a prefix length of 30 bits (or, in subnet mask terms, 255.255.255.252).

Figure 1.8 shows variable length subnetting for two-host WAN subnets.

**Figure 1.8 Variable Length Subnetting of 131.107.106.0**




If your network includes numerous WAN links, each with its own subnet, this approach can require significant administrative overhead. If you do not use route summarization, each subnet requires another entry in the routing table, increasing the overhead of the routing process.

Some routers support unnumbered connections; a link with unnumbered connections does not require its own subnet.

# Planning Supernetting and Classless Interdomain Routing (CIDR)

Updated: April 13, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Similar to the way that subnetting allows you to divide class-based networks into smaller subnets by "borrowing" bits from the host part of the address, *supernetting* allows you to combine contiguous subnets into larger supernets by "borrowing" bits from the network part of the address. For example, rather than allocate a Class B network ID to an organization that has 2,000 hosts, the Internet Assigned Numbers Authority (IANA) might allocate a range of eight Class C network IDs. Each Class C network ID accommodates 254 hosts, for a total of 2,032 host IDs.

Although this technique helps conserve Class B network IDs, it creates a new problem. Using conventional routing techniques, the routers on the Internet must, in this example, have eight Class C network ID entries in their routing tables to route IP packets to the organization. To prevent Internet routers from becoming overwhelmed with routes, a technique called *Classless Interdomain Routing (CIDR)*, which the Internet uses to summarize routes, collapses multiple network ID entries into a single entry. In this example, CIDR collapses the network IDs that correspond to the eight Class C network IDs allocated to that organization into one entry.

A supernetted subnet mask conveys the starting network ID and the number of Class C network IDs allocated. The following tables demonstrate how eight Class C network IDs are allocated. Table 1.1 indicates the contiguous allocation of eight Class C network IDs, starting with network ID 220.78.168.0. Note that the first 21 bits (underlined) are the same for the starting network ID and the ending network ID. The last 3 bits of the third octet, which are borrowed from the network ID, range from 000 through 111. In decimal notation, the range is 0 through 7, or 8 total contiguous subnets, which are combined into one supernet.

**Table 1.1 Supernetted Block of Addresses**

| | Network ID (decimal) | Network ID (Binary) |
|---|---|---|
| **Starting Network ID** | 220.78.168.0 | 11011100 01001110 10101000 00000000 |
| **Ending Network ID** | 220.78.175.0 | 11011100 01001110 10101111 00000000 |

A block of supernetted addresses, such as those in Table 1.2, is known as a *CIDR block*. Table 1.2 indicates the single CIDR entry that appears in the routing table. This entry represents all eight Class C network IDs that are allocated to the example organization.

**Table 1.2 CIDR Routing Table Entry**

| Network ID | Subnet Mask | Subnet Mask (Binary) |
|---|---|---|
| 220.78.168.0 | 255.255.248.0 | 11111111 11111111 11111000 0000000 |

In network prefix length notation, the CIDR block is 220.78.168.0/21.

RIP v2, OSPF, and BGP4, which can exchange routing information in the form of [*Network ID*, *Network Mask*] pairs, support CIDR.

# Choosing an Address Allocation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Choose an address allocation method that best fits your structured address model. Addressing by topology is recommended. However, you can choose one or more of the following methods:

- **Random address allocation.** Under a random addressing structure, you can assign blocks of addresses randomly. Random address allocation might be the most frequently used address allocation method, but it is the least desirable. For a small network where no significant growth is anticipated, this approach might be appropriate. However, if the network does grow, random address allocation can cause extra work for network administrators. Summarizing the random collection of routes might be difficult or impossible. This method can cause stability problems, with numerous routes being advertised to the core tier.

- **Addressing by organization chart.** To base your address structure on your organization chart, you create subnets based on a pool of addresses preassigned to a department or team. If, for example, you designate the Sales department as 10.2.0.0/16, the address 10.2.1.0/24 might be the subnet for the sales team at one site and 10.2.2.0/24 might be the subnet for the sales team at another site. To the extent that contiguous subnets remain unassigned, this address allocation method offers limited possibilities for route summarization, but, as a rule, this kind of addressing scheme does not scale well.

- **Addressing by geographical region.** When you base your address structure on location, a greater degree of summarization is possible. However, as the internetwork of a geographically diverse organization continues to grow, fewer routes are available for summarization.

- **Addressing by topology.** By basing your address structure on topology, you can ensure that summarization takes place and that an internetwork remains scalable and stable. Addressing by topology makes the addressing structure router-centric, enhancing efficiency.

# Choosing Public or Private Addresses

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you use a direct (routed) connection to the Internet, you must use public addresses. If you use an indirect connection such as a proxy server or Network Address Translator (NAT), use private addresses. If your organization is not connected to the Internet, use private addresses (rather than "unauthorized" addresses) so that if you later connect to the Internet using an indirect connection, you do not need to change addresses already in use.

If you connect to the Internet by using an Internet service provider (ISP), the ISP might provide only private addresses. The ISP itself uses public addresses to connect to the Internet.

**Public Addresses**
IANA assigns public addresses and guarantees them to be globally unique on the Internet. In addition, routes are programmed into the routers on the Internet so that traffic can reach those assigned public addresses. That is why public addresses can be reached on the Internet.

**Private Addresses**
Private addresses are a predefined set of IPv4 addresses that the designers of the Internet provided for those hosts within an organization that do not require direct access to the Internet. These addresses do not duplicate already assigned public addresses. RFC 1918, "Address Allocation for Private Internets," defines the following three private address blocks:

- **10.0.0.0/8.** The 10.0.0.0/8 private network is a Class A network ID that supports the following range of valid IP addresses: 10.0.0.1 through 10.255.255.254. The 10.0.0.0/8 private network has 24 host bits that a private organization can use for any subnetting scheme within the organization.

- **172.16.0.0/12.** The 172.16.0.0/12 private network can be interpreted either as a block of 16 Class B network IDs or as a 20-bit assignable address space (20 host bits) that can be used for any subnetting scheme within the private organization. The 172.16.0.0/12 private network supports the following range of valid IP addresses: 172.16.0.1 through 172.31.255.254.

- **192.168.0.0/16.** The 192.168.0.0/16 private network can be interpreted either as a block of 256 Class C network IDs or as a 16-bit assignable address space (16 host bits) that can be used for any subnetting scheme within the private organization. The 192.168.0.0/16 private network supports the following range of valid IP addresses: 192.168.0.1 through 192.168.255.254.

Because IANA never assigns IP addresses in the private address space as public addresses, routes for private addresses never exist on the Internet routers. Any number of organizations can repeatedly use the private address space, which helps to prevent the depletion of public addresses.

Private addresses cannot be reached on the Internet. Therefore, Internet traffic from a host that has a private address must either send its requests to an application layer gateway (such as a proxy server), which has a valid public address, or have its private address translated into a valid public address by a NAT before it is sent over the Internet.

For an introduction to TCP/IP and more information about public and private addresses, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

**Unauthorized Addresses**
Network administrators of private networks who have no plans to connect to the Internet can choose any IP addresses they want, even public addresses that IANA has assigned to other organizations. Such potentially duplicate addresses are known as *unauthorized* (or *illegal*) *addresses*. Later, if the organization decides to connect directly to the Internet after all, its current addressing scheme might include addresses that IANA has assigned to other organizations. You cannot connect to the Internet by using unauthorized addresses.

Do not use unauthorized addresses if even the slightest possibility exists of ever establishing a connection between your network and the Internet. On some future date, discovering that you need to quickly replace the IP addresses of all the nodes on a large private network can require considerable time and interrupt network operation.

**Network Address Translation**
*Network address translation*, defined in RFC 3022, is the translation process performed by an IP router functioning as a

network address translator (NAT). A NAT translates IP addresses from private network addresses used inside an organization to public addresses used outside the organization. Typically, a NAT-enabled router connects an internal corporate network with the Internet and builds a table that maps the connections between hosts inside the network and hosts outside on the Internet.

You can use NAT to map multiple internal private addresses to a single external public IP address. For example, a small business might obtain an ISP–allocated public IP address for each computer on its network. By using NAT, however, the business could use private addressing internally and have NAT map its private addresses to one or more public IP addresses that the ISP allocates.

NAT makes it more difficult for external users to attack systems on a private network. NAT also allows several nodes on the private network, each with its own private address, to share a smaller number of scarcer public addresses to access the Internet. However, although NAT allows you to reuse the private address space, it does not support standards-based network layer security or the correct mapping of all higher layer protocols. One purpose for the large number of addresses made available with the introduction of IPv6 is to make address conservation techniques such as NAT unnecessary.

Windows Server 2003 also supports IPSec NAT traversal (NAT-T), which allows nodes located behind a NAT (that is, they use private addresses) to use Encapsulating Security Payload (ESP) to protect traffic. This capability allows the creation of Layer Two Tunneling Protocol with IPSec (L2TP/IPSec) connections from remote access clients and routers located behind NATs.

For more information about unicast IP routing, including technical information about the NAT routing protocol component of the Routing and Remote Access service, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Planning an IP Configuration Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Every computer on an IP network must have a unique IP address. As noted earlier, using static addressing for clients is time-consuming and prone to error. To provide an alternative for IPv4, the IETF developed the Dynamic Host Configuration Protocol (DHCP), based on the earlier bootstrap protocol (BOOTP) standard. Figure 1.9 shows the stage in the TCP/IP design process during which you decide what to use for IP configuration. Most organizations choose to use DHCP for IPv4.

**Figure 1.9 Planning an IP Configuration Strategy**




Although BOOTP and DHCP hosts can interoperate, DHCP is easier to configure. BOOTP requires maintenance by a network administrator, whereas DHCP requires minimal maintenance after the initial installation and configuration.

The DHCP standard, defined in RFC 2131, defines a DHCP server as any computer running the DHCP service. Compared with static addressing, DHCP simplifies IP address management because the DHCP server automatically allocates IP addresses and related TCP/IP configuration settings to DHCP-enabled clients on the network. This is especially useful on a network with frequent configuration changes — for example, in an organization that has a large number of mobile users.

The DHCP server dynamically assigns specific addresses from a manually designated range of addresses called a *scope*. By using scopes, you can dynamically assign addresses to clients on the network no matter where the clients are located or how often they move.

#### DHCP Integration with DNS and WINS

The DHCP implementation in Windows Server 2003 is closely linked to name resolution services such as the Domain Name System (DNS) service and the Windows Internet Name Service (WINS). Network administrators benefit from combining all three when planning a deployment.

If you use DHCP servers for Windows-based network clients, you must use a name resolution service. In addition to name resolution, Windows Server 2003 networks use DNS to support Active Directory. Domain-based networks supporting clients running Windows NT version 4.0 or earlier or NetBIOS applications must use WINS servers. Networks supporting a combination of clients running Windows XP, Windows 2000, Windows Server 2003, and Windows NT 4.0 must implement both WINS and DNS.

#### DHCP, APIPA, and IP Address Allocation

DHCP clients receive IP addresses as follows:

- **Dynamic allocation — from DHCP server.** After you configure DHCP, the DHCP server automatically assigns an IP address from a specified scope to a client for a finite period of time called a *lease*. Most clients receive a dynamic IP address.

- **Static allocation — from DHCP server.** For a specific computer (such as a DHCP, DNS, or WINS server, or a print server, firewall, or router), you can manually configure the TCP/IP properties, including the IP address, the DNS and WINS parameters, and default gateway information. For the static clients to be on the same subnet as other, dynamically allocated computers, the static IP addresses must be within the scope or subnet defined for dynamic address allocation. You can use the DHCP snap-in to set an exclusion range to prevent the DHCP server from dynamically allocating the static IP address.

- **Client reservation — from DHCP server.** By using the DHCP snap-in, you can also reserve a specific IP address for permanent use by a given DHCP client.

- **Automatic allocation — APIPA.** In the absence of a DHCP server, Automatic Private IP Addressing (APIPA) lets a workstation configure itself with an address in the range from 169.254.0.1 to 169.254.255.254. Computers using APIPA addresses can communicate only with other computers that are also using APIPA addresses within a single subnet. In this case, a computer has an IP address but cannot connect outside the subnet. APIPA regularly checks for the presence of a DHCP server; if it detects one, it yields to the DHCP service, which then assigns a dynamic address to replace the APIPA address. APIPA is designed primarily for simple networks with only one subnet, such as small or home-based networks. On a larger network, APIPA can be useful for identifying problems with DHCP: when a client uses an APIPA address, this indicates that a DHCP server has not been found.

- **Alternate configuration — user configured.** In the absence of a DHCP server, alternate configuration lets a computer use an IP address configured manually by the user. Alternate configuration is designed for a computer that is used on more than one network, such as a laptop used both at the office and at home. The user can specify an IP address on the computer's TCP/IP properties **Alternate Configuration** tab if at least one of the networks (for example, the home office) does not have a DHCP server and APIPA addressing is not wanted. If alternate configuration is not configured and no DHCP server is found, TCP/IP uses APIPA by default.

For more information about developing a DHCP strategy, see "Deploying DHCP" in this book.

# Planning Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

IP does not have a default security mechanism. Without security, both public and private IP networks are susceptible to unauthorized monitoring and access. To prevent these types of security breach, develop a security strategy for your IP deployment in tandem with your overall network security plan.

Ways that you can enhance security when deploying IP include:

- **Securing IP packets.** Provide end-to-end security by securing IP packets, which requires that you not use address translation (unless both peers support IPSec NAT-T and use ESP to protect traffic). IPSec is the most efficient way to provide a secure data stream.

- **Deploying a perimeter network.** Use a perimeter network to help secure your internal network from intrusion. Several options are available for doing this.

Figure 1.10 shows the tasks involved in incorporating IPSec and a perimeter network in your IP security plan.

**Figure 1.10 Planning IP Security**



# Using IPSec

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Effective integration with IPSec is becoming increasingly important to the secure deployment of IP in an enterprise internetwork. IPSec is a framework of open standards for ensuring private, secure communications over IP networks through the use of cryptographic security services. The implementation of IPSec that runs on Windows Server 2003, Windows XP, and Windows 2000 is based on standards developed by the IETF IPSec working group.

IPSec provides a comprehensive technology for securing networks. However, the larger your organization, the more planning and engineering are required to implement IPSec. Assess the relative importance of your information resources — domain controllers, mail servers, and financial servers may rank high among the resources you want to protect. Include confidentiality considerations in your assessment. For example, many organizations might target Human Resources information for IPSec protection. After identifying the critical information resources to secure, configure IPSec policies as appropriate on those computers.

Windows Server 2003 uses the IPSec protocol suite to protect data traffic as it crosses a network. Although file encryption and required passwords protect information stored on network resources, they do not protect information as it moves across a network.

By implementing IPSec, you can secure the following types of data:

- Data that moves across the part of your intranet that external users do not access.
- Data that moves across the part of your intranet that can be accessed by external users who have appropriate permissions.
- Data that moves across the Internet.
- Data that moves across an extranet.

IPSec security protects the content of IP packets from both active and passive attacks. In an active attack, a hacker modifies existing data or adds false data. In a passive attack, an intruder reads data.

IPSec secures communication through the following methods:

- **Peer authentication.** IPSec verifies the identity of each computer. Each peer sends security credentials that are verified by the peer at the other end of the connection. Windows Server 2003 IPSec provides multiple methods of peer authentication.
- **Data origin authentication.** By incorporating a cryptographic checksum calculated with a shared secret key with each packet of protected data, IPSec can verify that the packet must have been sent by a peer that has knowledge of the secret key.
- **Confidentiality (data encryption).** IPSec offers confidentiality by encrypting data before transmission, ensuring that the data cannot be read during transmission — even if an attacker monitors or intercepts the packet. IPSec encryption is applied at the IP network layer, which makes it transparent to applications that use TCP or User Datagram Protocol (UDP) for network communication.
- **Integrity.** IPSec protects data from unauthorized modification in transit, ensuring that the information received is exactly the same as the information sent.
- **Anti-replay.** IPSec ensures that any attacker who might intercept data cannot reuse or replay that data to establish a session or to illegally gain information or access to resources.

Deploying IPSec requires careful planning. For more information about deploying IPSec, see "Deploying IPSec" in this book. For more technical information about IPSec, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Using a Perimeter Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A perimeter network protects your intranet or enterprise LAN from intrusion by controlling access from the Internet or other large network. The perimeter network (also known as a *demilitarized zone* or *DMZ*) is bounded by firewalls. A firewall is not a single component, but rather a system or combination of systems that enforces a boundary between two or more networks.

Figure 1.11 shows a perimeter network bounded by firewalls placed between a private network and the Internet in order to secure the private network.

**Figure 1.11 Perimeter Network Securing an Internal Network**




Organizations vary in their use of firewalls for providing security. IP packet filtering offers weak security, is cumbersome to manage, and is easily defeated. Application gateways are more secure than packet filters and easier to manage because they pertain only to a few specific applications, such as a particular e-mail system. Circuit gateways are most effective when the user of a network application is of greater concern than the data being passed by that application. The proxy server — the recommended solution — is a comprehensive security tool that includes an application gateway, safe access for anonymous users, and other services.

**IP packet filtering** You can configure packet filtering, the earliest implementation of firewall technology, to accept or deny specific types of packets. Packet headers are examined for source and destination addresses, TCP and UDP port numbers, and other information. Packet filtering is a limited technology that works best in clear security environments where, for example, everything outside the perimeter network is not trusted and everything inside is. You cannot use IP packet filtering when IP packet payloads are encrypted because the port numbers are encrypted and therefore cannot be examined.

In recent years, various vendors have improved on the packet filtering method by adding intelligent decision-making features to the packet-filtering core, thus creating a new form of packet filtering called *stateful protocol inspection*.

**Application gateways** Used when the actual content of an application is of greatest concern, application gateways do not adapt easily to changes in technology. However, unlike IP packet filtering, application gateways can be used in conjunction with encryption.

**Circuit gateways** As tunnels connecting specific processes or systems on each side of a firewall, circuit gateways are best employed in situations where the person using an application is potentially a greater risk than the information that the application carries. The circuit gateway differs from a packet filter in its capability for connecting to an out-of-band application scheme that can add additional information.

**Proxy servers** Proxy servers are comprehensive security tools that include firewall and application gateway functionality to manage Internet traffic to and from a private intranet. Proxy servers also provide document caching and access control. A proxy server can improve performance by caching and directly supplying frequently requested data such as a popular Web page. A proxy server also can filter and discard requests that the owner does not consider appropriate, such as requests for unauthorized access to proprietary files.

Take advantage of those firewall security features that can help you. Position a perimeter network in your network topology at a point where all traffic from outside the corporate network must pass through the perimeter that the external firewall maintains. You can fine-tune access control for the firewall to meet your needs and can configure firewalls to report all attempts at unauthorized access.

# Improving Availability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

*Availability* refers to how much time the network is operational. Planning well for availability improves both your network's mean time between failures (MTBF) and its mean time to recovery (MTTR) after a network failure.

To improve availability in your IP network design, you must know your organization's availability requirements. For some organizations, unanticipated down time is simply an irritating inconvenience. In other environments, unanticipated down time could mean financial disaster, drastic loss of credibility, or, as in health care or law enforcement, a risk to safety.

Figure 1.12 shows the process for improving availability on your network.

**Figure 1.12 Improving Availability**




Each method for improving availability places different demands on the design of your network. As the risk of down time to your operation increases, build more redundancy into your design, both in hardware and routing. Similarly, as the consequences of failure increase, make your network more resilient by increasing the amount of stress it can handle before it loses functionality.

**Implementing Redundancy**

Single points of failure, such as devices, links, and interfaces, can make a network vulnerable. If one such point fails, it isolates users from services and, in the worst case, causes entire sections of the network to fail. For a purely hierarchical network — one based on summarization and controlled access between tiers — every device and link is a point of failure.

Redundancy provides alternative paths around points of failure. In a purely redundant network, each individual device, link, and interface is dispensable. No single device, link, or interface can isolate users or cause the network to fail.

In most production environments, neither a purely hierarchical nor a purely redundant network is practical. You must balance the efficiency of a hierarchical network with the safety net of redundancy.

**Implementing Secondary Paths**

After deploying multiple devices to eliminate single points of failure, configure secondary paths to take advantage of the multiple devices. A *secondary path*, or *backup path*, consists of the interconnecting devices and the links between them that duplicate the devices and links in the primary path. For example, you can configure multiple routers to provide redundancy.

A redundant design uses the secondary path to maintain network connectivity when any of the primary path's devices or links fails. Be sure to test any secondary paths on a regular basis. Do not assume that they will work. If possible, ensure that the switch from the primary path to the secondary path occurs transparently. For mission-critical applications, automatic failover is mandatory.

**Using Load Balancing**

In addition to its safety net function, redundancy plays a second valuable role. By properly configuring two or more paths that connect the same source and destination networks, you can significantly improve throughput by providing load balancing. Load balancing evenly divides the flow of traffic among parallel links.

Most routing protocols based on open standards support load balancing across paths that the protocol determines to be equally favorable to the destination. In addition, some vendors' proprietary routing protocols support load balancing where the costs of the paths (their relative favorability to the destination in terms of shortest distance, number of hops, and other criteria) are not considered equal.

For more information about network load balancing, see "Designing Network Load Balancing" in *Planning Server Deployments* of this kit.

# Planning IP Multicasting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With IP multicasting, one device can send a single data stream that the network replicates only as necessary so that multiple devices receive the data. Because of the minimal overhead required to create the data stream and the low overhead on the network, multicast communication is particularly suitable for multiple-user multimedia applications such as video conferencing, distance learning, and collaborative computing. You can also use multicast traffic to discover resources on the internetwork and to support datacasting applications such as file distribution or database synchronization.

Using the IP multicast components of the Windows Server 2003 TCP/IP protocol and the Routing and Remote Access service, you can send and receive IP multicast traffic from multicast-enabled portions of your intranet or the Internet and from remote access clients. You can use IP multicast to optimize server loading and network bandwidth.

Figure 1.13 shows the tasks involved in planning IP multicasting.

**Figure 1.13 Planning IP Multicasting**




In *multicast routing*, routers communicate multicast group membership information to each other using multicast routing protocols, and forward data across the internetwork. *Multicast forwarding* refers to the process of forwarding multicast traffic to networks on which other multicast devices are listening. The multicast-capable portion of the Internet is referred to as the Internet *multicast backbone*, or *MBone*.

All computers running Windows Server 2003 can both send and receive IP multicast traffic. Windows Server 2003 TCP/IP can listen for IPv4 multicast traffic and use a multicast forwarding table to determine where to forward incoming multicast traffic.

Figure 1.14 shows one common configuration of IP multicast components. For examples of a number of supported multicast configurations, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

**Figure 1.14 IP Multicast Components**


IGMP proxy
mode interface
IGMP router
mode interface
IP multicast-enabled
router running a
multicast routing protocol
Multicast
source
Computer running
Windows Server 2003 Routing
and Remote Access service
IP multicast-enabled
router running a
multicast routing protocol
Windows XP
Host
Enterprise Intranet

# Planning MADCAP Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Multicast Address Dynamic Client Allocation Protocol (MADCAP), built on a client/server model, enables a computer to request an IP multicast address from one or more multicast address allocation servers, known as MADCAP servers. If a client sends a message and does not receive a response, it can retransmit its request.

MADCAP as defined in RFC 2730, "Multicast Address Dynamic Client Allocation Protocol (MADCAP)," differs substantially and is separate from DHCP. However, the Windows Server 2003 DHCP service combines support for both the DHCP and MADCAP protocols for IPv4. Although MADCAP is packaged in the DHCP service, the DHCP and MADCAP services are independent of each other. A DHCP client might or might not be a MADCAP client, and a MADCAP client might or might not be a DHCP client.

**MADCAP Without DHCP**
To use the DHCP service to deploy MADCAP servers independently of DHCP servers, create one or more multicast scopes, but do not create other scopes or superscopes. The MADCAP server also functions as a DHCP server only if you configure other scopes or superscopes.

**MADCAP Security**
The IPSec protocol meets MADCAP requirements for client/server identification and integrity protection as described in RFC 2730, and requires no modifications to the MADCAP protocol. Therefore, when you require strong security, use IPSec to protect all of the unicast messages of the MADCAP protocol.

For more information about MADCAP, including how to use IPSec in conjunction with MADCAP, see RFC 2730, "Multicast Address Dynamic Client Allocation Protocol (MADCAP)."

# Planning IP Multicast-Enabled Routers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To implement IP multicasting on a multiple-router intranet, you must install routers enabled for multicast routing and configured with one or more multicast routing protocols.

Windows Server 2003 does not provide any multicast routing protocols. To provide multicast forwarding within a single-router intranet or when connecting a single-router intranet to the Internet, you can configure the Internet Group Management Protocol (IGMP) routing protocol component of the Routing and Remote Access service with interfaces set to IGMP router mode and IGMP proxy mode. The IGMP routing protocol component exchanges and updates information in the IP multicast forwarding table about host membership in specific groups.

The IGMP routing protocol is not a multicast routing protocol. To support efficient multicast forwarding and routing on a multiple-router intranet, you must also install IP multicast-enabled routers that use one or more multicast routing protocols. Multicast routers use multicast routing protocols to communicate multicast group information with each other.

**Note**

- You can configure the IGMP router mode and IGMP proxy mode interfaces to provide multicast forwarding support in multiple-router intranets, but doing so is not efficient and is therefore not recommended or supported.

Although Windows Server 2003 does not include any multicast routing protocols, the Routing and Remote Access service is an extensible platform that can support multicast routing protocols. Multicast routing protocols include Protocol-Independent Multicast (PIM) in both Sparse Mode (PIM-SM) and Dense Mode (PIM-DM), Multicast Extensions to OSPF (MOSPF), and the Distance Vector Multicast Routing Protocol (DVMRP). Your choice of multicast routing protocol will depend on the size and type of network and the distribution of multicast group members.

- **Protocol-Independent Multicast (PIM).** The PIM protocol routes to multicast groups whose members span wide-area and interdomain internetworks. PIM functions independently of any unicast routing protocol. A multicast group that uses PIM can declare itself sparse or dense, using either Sparse Mode or Dense Mode:

    - **Protocol-Independent Multicast Sparse Mode (PIM-SM)**, the most widely used multicast routing protocol, is designed for multicast groups whose members are distributed sparsely across a large region. PIM-SM can operate in a LAN environment but is most efficient in a WAN environment. Using a dense-mode protocol for a multicast group whose members are distributed thinly can cause unnecessary transmission and router storage of data packets or membership report information. This overhead might be acceptable where multicast group members are populated densely, but it is inefficient for a sparse mode multicast group. In sparse mode, routers must explicitly join and leave multicast groups, which eliminates unnecessary traffic and storage.

    - **Protocol-Independent Multicast Dense Mode (PIM-DM)** is a dense-mode multicast routing protocol designed for multicast groups whose members are distributed thickly over an area where bandwidth is plentiful. PIM-DM is interoperable with the sparse mode, PIM-SM. PIM-DM does not scale well.

- **Multicast Extensions to OSPF (MOSPF).** The MOSPF protocol, an extension of OSPF, is also a dense-mode multicast routing protocol. MOSPF employs a unicast routing protocol that requires that each router in a network be aware of all available links. MOSPF is intended for use on a single organization's network, and does not scale well. MOSPF requires OSPF as its accompanying unicast routing protocol. It can sometimes put a heavy load on router CPU bandwidth.

- **Distance Vector Multicast Routing Protocol (DVMRP).** The original IPv4 multicast routing protocol, DVMRP runs over multicast-capable LANs such as Ethernet. DVMRP can also tunnel IP multicast packets as unicast packets through routers with no multicast capability. DVMRP is a dense-mode multicast routing protocol that does not scale well.

# Configuring IGMP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To support IPv4 multicast applications on a single-router intranet or when connecting a single-router intranet to the Internet, you can use the Routing and Remote Access service on one or more computers running Windows Server 2003, add the IGMP routing protocol component on each server, and configure the server's outbound interface for IGMP router mode and its inbound interface for IGMP proxy mode. If your multicast applications cross the Internet, the outbound interface is the intranet interface and the inbound interface is the Internet interface.

- **IGMP router mode on the outbound interface.** In Windows Server 2003, an outbound interface running in IGMP router mode listens for IGMP Membership Report messages and tracks group membership. Enable IGMP router mode on the interfaces to listening multicast hosts. The TCP/IP protocol and the IGMP routing protocol component for interfaces running in IGMP router mode forward multicast traffic.

- **IGMP proxy mode on the inbound interface.** IGMP proxy mode is designed to pass IGMP Membership Report messages within a single-router intranet or from a single-router intranet to the MBone. (As explained earlier, in a multiple-router intranet, you must install routers that use one or more multicast routing protocols.) With IGMP proxy mode enabled on the inbound interface, hosts can receive multicast traffic from multicast sources and can send multicast traffic to other hosts.

Within a single-router intranet, or when connecting a single-router intranet to the Internet, you do not need routers running multicast routing protocols. However, within a multiple-router intranet that uses multicast routers running multicast routing protocols, you can still use the Routing and Remote Access service as a multicast forwarding router on the periphery of your intranet.

RFC 1112, "Host Extensions for IP Multicasting," defines address and host extensions for IP hosts that support multicasting, and defines IGMP Version 1. RFC 2236, "Internet Group Management Protocol (IGMP), Version 2," defines IGMP Version 2. Windows Server 2003 supports IGMP Version 3, described in the Internet Draft "Internet Group Management Protocol, Version 3." Under IGMP Version 3, hosts can specify interest in receiving multicast traffic from specified sources or from all but a specific set of sources.

# Configuring IP Multicast Scopes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Multicast addressing supports *dynamic membership*, under which individual computers can join or leave a multicast group at any time. Group membership is not limited by size, and computers are not restricted to membership in any single group.

On all IP networks, each computer must first be configured with its own unicast IP address. After assigning this unicast address, you can configure the computer to support a multicast address. A multicast group of computers shares the same multicast IP address. IPv4 multicast addresses range, in dotted decimal notation, from 224.0.0.0 through 239.255.255.255 (224.0.0.0/4). Such a multicast group also uses a MAC-layer multicast address, which allows all devices to filter unsolicited multicast traffic at the link layer. Ethernet addresses reserved for multicasting range from 01-00-5E-00-00-00 through 01-00-5E-7F-FF-FF.

Typically, you specify IP address ranges for multicast scopes on your MADCAP server in the following ways:

- **Administrative scoping** is designed for multicast IP addresses used privately on your intranet. You use the 239.192.0.0 range of the multicast (Class D) address space with a subnet mask of 255.252.0.0. This is known as the IPv4 Organization Local Scope. It provides 262,144 group addresses ($2^{18}$) for use in all subnets on your network. For more information about administrative scoping, see RFC 2365, "Administratively Scoped IP Multicast."

- **Global scoping** is designed for multicast IP addresses used on the Internet. You use the 233.0.0.0 range of the multicast address space. Global addresses are allocated in the following way:

  - IANA (or another network registry) allocates and reserves the first 8 bits of the range (the "233" portion).

  - The next 16 bits are based on your Autonomous System (AS) number. For information about obtaining your existing AS number or acquiring a new one, see "Using multicast scopes" in Help and Support Center for Windows Server 2003.

  - The last 8 bits provide the IP address range from which to configure any multicast scopes for group addresses that you want to use publicly on the Internet. Use a subnet mask of 255.255.255.0.

  For more information about global scoping, see RFC 3180, "GLOP Addressing in 233/8." For more information about AS numbering, see RFC 1930, "Guidelines for Creation, Selection, and Registration of an Autonomous System (AS)."

# Configuring Client Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

On participating clients, install and configure the appropriate MADCAP-aware hardware and software. For example, for video conferencing, install video conferencing software and a video camera, sound card, and audio headset.

Standards for the multicast transmission of a data stream between the subnets of an internetwork include RFC 1112, "Host Extensions for IP Multicasting"; RFC 2236, "Internet Group Management Protocol, Version 2"; and the Internet Draft "Internet Group Management Protocol, Version 3." Such standards instruct routers how and where to route multicast traffic.

# Introducing IPv6 on Your Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In addition to the IPv4 stack installed by default, Windows Server 2003 and Windows XP include an IPv6 protocol stack that you can use to test IPv6, to explore IPv6-enabled applications, and to prepare for possible eventual migration to a native IPv6 infrastructure.

It is expected that IPv4 and IPv6 will coexist on enterprise networks for a number of years. Depending on their needs, some organizations might continue to use IPv4 exclusively, some will migrate slowly while running both IPv4 and IPv6 in the interim, and some will maintain IPv4 in one or more sections of their organization and implement IPv6 in other sections.

To ensure that your organization makes best use of IPv6 capabilities with the least administrative overhead, include a plan for introducing IPv6 into the design for your TCP/IP network. To prepare to introduce IPv6, you must explore the new functionality introduced by IPv6, plan IPv6 addressing, plan how to route IPv6 traffic over an existing IPv4 infrastructure or an IPv6 infrastructure, decide whether to deploy DNS dynamic update, and decide whether to deploy PortProxy to enable IPv4 applications (where possible) for IPv6. Figure 1.15 shows each task in the planning process.

**Figure 1.15 Introducing IPv6 on Your Network**



# Exploring IPv6

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 includes an IPv6 stack, in addition to the IPv4 stack, which you can use to explore the capabilities of IPv6, test new applications and network technologies, and plan the first steps toward the wider adoption of IPv6 on your network.

The current version of the Internet Protocol — IP version 4, known as *IPv4* — dates from 1981 and has not changed substantially since it was introduced in RFC 791, "Internet Protocol." Although IPv4 proved to be remarkably robust and enduring, in the early 1990s the Internet Engineering Task Force (IETF) began to develop a suite of protocols and standards — IPv6 — to better address the demands of modern networking. Two of the most important of these protocols are RFC 2460, "Internet Protocol, Version 6 (IPv6) Specification," which defines IPv6, and RFC 2463, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification," which specifies a set of ICMP messages for use with IPv6.

Before considering the design choices that you must make when introducing IPv6 on your network, you must become familiar with some of the basics about IPv6, including:

- IPv6 features.
- Supported features, server applications, and application programming interfaces (APIs).
- Supported IPv6 tools.
- Types of nodes.

**IPv6 Features**

The IPv6 protocol includes the following features and improvements over IPv4:

- **New header format.** The IPv6 header is designed to minimize overhead. Although the IPv6 address field is four times as long as the address field in IPv4, the IPv6 header is only twice as large as the IPv4 header overall. The more efficient header design enables faster processing at intermediate routers. Because IPv6 headers are not interoperable with IPv4 headers, and the IPv6 protocol is not backward compatible with IPv4. A host or router must use an implementation of both IPv4 and IPv6 in order to recognize and process both header formats.
- **Large address space.** IPv6 provides 128-bit IP addresses, in contrast with the 32-bit IPv4 IP addresses. The address space is designed to accommodate a vast number of interconnected devices on any network, and its structure is designed to reduce the number of routing table entries in IPv6 routers.
- **Hierarchical addressing and routing infrastructure.** IPv6 global addresses are designed to facilitate a hierarchical routing infrastructure that is based on the common occurrence of multiple levels of ISPs. It is anticipated that the routing tables for backbone routers on the IPv6 Internet will be much smaller and, as a result, will be processed much more efficiently.
- **Automatic address configuration.** IPv6 simplifies address configuration and renumbering by enabling automatic address configuration for all hosts. Host interfaces automatically learn their addresses through interactions with local IPv6 routers. They can learn new addresses on the fly, making network renumbering much simpler than in IPv4.
- **Integrated network security.** Support for IPSec is an IPv6 protocol suite requirement.
- **Better support for Qualify of Service (QoS).** The IPv6 header contains a new field that can be used to determine how to identify and prioritize traffic. Because the traffic type can be identified within the IPv6 header, support for QoS is available even when IPSec encryption is in use.
- **New protocol for neighboring node interaction.** The *IPv6 Neighbor Discovery protocol* is a series of Internet Control Message Protocols for IPv6 messages (ICMPv6) that manage the interaction of nodes on the same link. Neighbor Discovery replaces broadcast-based Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and ICMPv4 Redirect messages with efficient multicast and unicast Neighbor Discovery messages.

- **Extensibility.** IPv6 can be easily extended to incorporate new features by adding extension headers after the IPv6 header. The size of IPv6 extension headers is limited only by the size of the IPv6 packets.

**Supported Features, Server Applications, and APIs**

Windows Server 2003 supports IPv6 functionality for a wide range of services. Table 1.3 shows which IPv6 features Windows Server 2003 IPv6 supports.

**Table 1.3 IPv6 Features Supported by Windows Server 2003 IPv6**

| IPv6 Feature | Supported by Windows Server 2003 IPv6 |
|---|---|
| Installation<br><br>(Use **Add** protocol GUI, or use the **Netsh** command-line tool) | Yes |
| Uninstallation<br><br>(Use **Remove** protocol GUI, or use the **Netsh** command-line tool) | Yes |
| Dual IPv6/IPv4 stack | Yes |
| 6to4 | Yes |
| ISATAP | Yes |
| 6over4 (manual) | Yes |
| IPv6 NAT Traversal<br><br>(also referred to as Teredo) | No |
| DNS over IPv6<br><br>(also referred to as DNS AAAA records) | Yes |
| Linklocal Multicast Name Resolution (LLMNR) | No |
| DNS dynamic update | Yes |
| DHCP | No |
| TCP PortProxy | Yes |
| Remote Desktop | No |
| Remote Assistance | No |
| IPv6 Management Information Base (MIB) for Simple Network Management Protocol (SNMP) | Yes |
| Microsoft Network Monitor version 2 (Netmon) | Yes |
| Visual Studio .NET (VS.NET) | Yes |
| IPSec authentication | Yes |
| IPSec encryption | No |

Table 1.4 shows which server applications Windows Server 2003 IPv6 supports.

**Table 1.4 Server Applications Supported by Windows Server 2003 IPv6**

| Server Applications | Supported by Windows Server 2003 IPv6 |
|---|---|

| | |
|---|---|
| File sharing, printer sharing | Yes |
| Windows Media Server | Yes |
| Internet Information Services (IIS) 6.0 (HTTP only) | Yes |
| Telnet server | Yes |
| FTP server | No |
| Active Directory | No |
| Microsoft® Exchange Server | No |
| SQL Server™ | No |

Windows Server 2003 IPv6 also supports Internet Explorer. However, it does not include support for literal addresses.

In addition, the following APIs support Windows Server 2003 IPv6:

- .NET Framework
- Windows Sockets 2 (Winsock2) API
- Remote procedure call (RPC)
- Distributed Component Object Model (DCOM)
- Windows Internet (WinINet) API (does not include support for literal addresses)
- Windows HTTP Services (WinHTTP)
- HTTP.sys
- IP Helper API (IPHLPAPI) module
- Debuggers

**Supported IPv6 Tools**

Windows Server 2003 IPv6 supports the following tools.

- Ping
- Tracert
- Pathping
- Ipconfig
- Route
- Netsh (Use `netsh interface IPv6` commands)
- Netstat
- Nslookup
- Telnet client

- FTP client

For more information about these TCP/IP tools and commands, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

**Types of Nodes**

To understand IPv6 tunneling technologies, such as 6to4 and ISATAP (described later), you must understand the types of nodes that might be involved. Table 1.5 shows IPv4 and IPv6 node types.

**Table 1.5 IPv4 and IPv6 Node Types**

| Node Type | Description |
|---|---|
| IPv4-only node | A device that can communicate only with IPv4 nodes and applications and that does not support IPv6. |
| IPv6-only node | A device that can communicate only with IPv6 nodes and that does not support IPv4. |
| IPv6/IPv4 node | A device that implements both IPv4 and IPv6 and that can communicate with either IPv6 or IPv4 nodes and applications. |
| IPv4 node | Any device that supports IPv4. Both IPv4-only and IPv6/IPv4 nodes are IPv4 nodes. |
| IPv6 node | Any device that supports IPv6. Both IPv6-only and IPv6/IPv4 nodes are IPv6 nodes. |

For more information about the different node types, see RFC 2893, "Transition Mechanisms for IPv6 Hosts and Routers."

# Planning IPv6 Addressing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To plan an efficient IPv6 addressing strategy, you must understand how IPv6 addressing works. IPv6 addressing is a major departure from IPv4 addressing. The most obvious difference is that IPv4 uses 4-byte source and destination addresses, typically expressed in the familiar dotted-decimal notation, whereas IPv6 uses 16-byte addresses, typically expressed in *colon-hexadecimal* notation. Colon-hexadecimal notation uses eight 4-digit hexadecimal numbers, with colons separating the 16-bit blocks (the 4-digit numbers).

To manage addresses more easily, IPv6 suppresses leading zeros and compresses a single contiguous all-zero 16-bit block, representing the contiguous block with two colons (::) (known as *double-colon compression*). Table 1.6 shows the effects of suppressing leading zeros and double-colon compression on the notation for an IPv6 address.

**Table 1.6 Leading Zero Suppression and All-Zero Contiguous Block Compression**

| IPv6 Address Notation | IPv6 Address |
|---|---|
| IPv6 address | FEC0:0000:0000:0000:02AA:00FF:FE3F:2A1C |
| IPv6 address with leading zeros suppressed | FEC0:0:0:0:2AA:FF:FE3F:2A1C |
| IPv6 address with leading zeros suppressed and an all-zero contiguous block compressed | FEC0::2AA:FF:FE3F:2A1C |

The 16 bytes, or 128 bits, provided in the IPv6 address space potentially supports $2^{128}$ addresses. However, the purpose of this large address space is not only to provide an inexhaustible supply of addresses, but also to enable a hierarchical routing infrastructure that can be summarized. IPv6 addressing is designed to minimize the size of routing tables and to reduce routing complexity.

IPv6 supports address configuration both in the presence of a DHCP server, known as *stateful address configuration*, and in the absence of a DHCP server, known as *stateless address configuration*. Stateless address configuration introduces the use of *link-local addresses*, whereby hosts on the same link automatically configure themselves with IPv6 addresses for that link and can use those addresses to communicate with the other hosts on the same link. If one or more local routers exist, hosts can use router discovery to automatically determine the routers' addresses and can then communicate with IPv6 hosts beyond the local link.

As in IPv4, the high-order bits in an IPv6 address identify the type of address. In IPv6, the high-order bits are known as the *Format Prefix (FP)*. IPv6 does not use subnet masks to specify the network ID. Instead, it uses only prefix notation.

# IPv6 Address Types

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

IPv6 has three types of addresses, which can be categorized by type and scope:

- **Unicast addresses.** A packet is delivered to one interface.
- **Multicast addresses.** A packet is delivered to multiple interfaces.
- **Anycast addresses.** A packet is delivered to the nearest of multiple interfaces (in terms of routing distance).

IPv6 does not use broadcast messages.

Unicast and anycast addresses in IPv6 have the following scopes (for multicast addresses, the scope is built into the address structure):

- **Link-local.** The scope is the local link (nodes on the same subnet).
- **Site-local.** The scope is the organization (private site addressing).
- **Global.** The scope is global (IPv6 Internet addresses).

In addition, IPv6 has special addresses such as the loopback address. The scope of a special address depends on the type of special address.

Much of the IPv6 address space is unassigned.

**Unicast IPv6 Addresses**
IPv6 has several major unicast address types.

**Unicast global addresses**
IPv6 unicast global addresses are similar to IPv4 public addresses. Also known as *aggregatable global unicast addresses*, global addresses are globally routable. The structure of an IPv6 unicast global address creates the three-level topology shown in the following illustration.




Table 1.7 explains each field in a unicast global address.

**Table 1.7 Fields in a Unicast Global Address**

| Field | Description |
|---|---|
| **001** | Identifies the address as an IPv6 unicast global address. |
| **Top Level Aggregation Identifier (TLA ID)** | Identifies the highest level in the routing hierarchy. TLA IDs are administered by IANA, which allocates them to local Internet registries, which then allocate a given TLA ID to a global ISP. |
| **Res** | Reserved for future use (to expand either the TLA ID or the NLA ID). |
| **Next Level Aggregation Identifier (NLA ID)** | Identifies a specific customer site. |
| **Site Level Aggregation Identifier (SLA ID)** | Enables as many as 65,536 ($2^{16}$) subnets within an individual organization's site. The SLA ID is assigned within the site; an ISP cannot change this part of the address. |

| | |
|---|---|
| **Interface ID** | Identifies the interface of a node on a specific subnet. |

**Unicast site-local addresses**

IPv6 unicast site-local addresses are similar to IPv4 private addresses. The scope of a site-local address is the internetwork of an organization's site. (You can use both global addresses and site-local addresses in your network.) The prefix for site-local addresses is FEC0::/48.

The following illustration shows the structure of a site-local address.




The initial 48 fixed bits are followed by a 16-bit **Subnet ID** field, which provides as many as 65,536 subnets in a flat subnet structure. Alternatively, you can subdivide the high-order bits of the **Subnet ID** field to create a hierarchical routing infrastructure. The last field is a 64-bit **Interface ID** field that identifies the interface of a node on a specific subnet.

### Note

- Global addresses and site-local addresses share the same structure after the first 48 bits — the 16-bit SLA ID of a global address and the 16-bit Subnet ID of a site-local address both identify the subnets of an organization's site. Because of this, you can assign a specific subnet number to identify a subnet that is used for both global and site-local unicast addresses.

**Unicast link-local addresses (FE80::/64)**

IPv6 unicast link-local addresses are similar to IPv4 APIPA addresses used by computers running Microsoft Windows. Hosts on the same link (the same subnet) use these automatically configured addresses to communicate with each other. Neighbor Discovery provides address resolution. The prefix for link-local addresses is FE80::/64. The following illustration shows the structure of a link-local address.

**Unicast unspecified address**

The IPv6 unicast unspecified address is equivalent to the IPv4 unspecified address of 0.0.0.0. The IPv6 unspecified address is 0:0:0:0:0:0:0:0:, or a double colon (::).

**Unicast loopback address**

The IPv6 unicast loopback address is equivalent to the IPv4 loopback address, 127.0.0.1. The IPv6 loopback address is 0:0:0:0:0:0:0:1, or ::1.

**Unicast 6to4 addresses (2002::/16)**

IPv6 uses 6to4 addresses to communicate between two IPv6/IPv4 nodes over the IPv4 Internet. A 6to4 address combines the prefix 2002::/16 with the 32 bits of the public IPv4 address of the node to create a 48-bit prefix — 2002:*WWXX*:*YYZZ*::/48, where *WWXX*:*YYZZ* is the colon-hexadecimal representation of *w.x.y.z*, a public IPv4 address. Therefore, the IPv4 address 157.60.91.123 translates into a 6to4 address prefix of 2002:9D3C:5B7B::/48.

The following illustration shows the structure of a 6to4 address.




However, this is often written using the hexadecimal prefix: 2002:*WWXX*:*YYZZ*:*SLA ID*:*Interface ID*.

The following example shows how the *WWXX*:*YYZZ* portion of the address is translated from colon-hexadecimal notation to dotted-decimal notation. In this example, 9D3C:5B7B translates to 157.60.91.123, as illustrated in the following example.

| Notation Type | Use a calculator to convert each constituent number from one notation type to the other |
|---|---|
| Colon-hexadecimal | 9D 3C 5B 7B |
| Dotted-decimal | 157 60 91 123 |

For more information about 6to4 tunneling, see "Routing IPv6 Traffic over an IPv4 Infrastructure" later in this chapter.

**Unicast ISATAP addresses**

IPv6 uses ISATAP addresses to communicate between two IPv6/IPv4 nodes over an IPv4 intranet. An ISATAP address combines a 64-bit unicast link-local, site-local, or global prefix (a global prefix might be a 6 to 4 prefix) with a 64-bit suffix constructed of the ISATAP identifier 0:5EFE, followed by the IPv4 address assigned to an interface of the host. The prefix is known as the *subnet prefix*. Although a 6to4 address can incorporate only a public IPv4 address, an ISATAP address can incorporate either a public or a private IPv4 address.

The following illustration shows the structure of an ISATAP address.

| Subnet prefix <A link-local, site-local, or global prefix> (64 bits) | 0000:5EFE (32 bits) | *WWXX:YYZZ* (32 bits) |
|---|---|---|

Table 1.8 shows an example of each type of ISATAP address.

**Table 1.8 Examples of ISATAP addresses**

| Type of ISATAP Address | ISATAP Address |
|---|---|
| With link-local prefix | FE80::5EFE:131.107.129.8* |
| With site-local prefix | FEC0::1111:0:5EFE:131.107.129.8* |
| With global prefix | 3FFE:1A05:510:1111:0:5EFE:131.107.129.8* |
| With global 6to4 prefix | 2002:9D36:1:2:0:5EFE:131.107.129.8* |

*Alternatively, the IPv4 address (in this example, 131.107.129.8) can be written in hexadecimal (in this example, 836B:8108).

By default, the IPv6 protocol for Windows XP and members of Windows Server 2003 automatically configures the ISATAP address of FE80::5EFE:*w.x.y.z* for each IPv4 address that is assigned to the node. This link-local ISATAP address allows two hosts to communicate over an IPv4 network by using each other's ISATAP address.

For more information about ISATAP tunneling, see "Routing IPv6 Traffic over an IPv4 Infrastructure" later in this chapter.

**Multicast IPv6 Addresses**

IPv6 multicast addresses are similar to IPv4 multicast addresses. Packets addressed to a multicast address are delivered to all interfaces that the address identifies.

The following illustration shows the structure of an IPv6 multicast address.

| 1111 1111 (8 bits) | Flags (4 bits) | Scope (4 bits) | Group ID (112 bits) |
|---|---|---|---|

Table 1.9 explains each field in an IP multicast address. The prefix for multicast addresses is FF00::/8.

**Table 1.9 Fields in a Multicast Address**

| Field | Description |
|---|---|
| **1111 1111** | Identifies the address as an IP multicast address. |
| **Flags** | Currently, the only defined flag is the **Transient** (**T**) flag. Set to zero, the T flag identifies the address as a permanently assigned multicast address. Set to **1**, it identifies a transient address. |
| **Scope** | Indicates the scope of the multicast traffic, such as interface-local, link-local, site-local, organization-local, or global scope. |
| **Group ID** | identifies the multicast group. |

**Multicast solicited node address**

The IPv6 multicast solicited node address is used for efficient address resolution. The IPv4 ARP Request frame is sent to the MAC-level broadcast, which disturbs all nodes on the network segment. The multicast solicited node address combines the prefix FF02::1:FF00:0/104 with the last 24 bits of the IPv6 address being resolved. IPv6 uses the solicited node multicast address for the Neighbor Solicitation message (the IPv6 equivalent to the ARP Request frame) that resolves an IPv6 address to

its link-layer address, disturbing few nodes during the address resolution process.

**Anycast IPv6 Addresses**

Anycast IPv6 addresses are similar to but more efficient than the anycast addresses in IPv4, which are used primarily by large ISPs. Anycast addresses use the unicast address space but function differently from other unicast addresses. IPv6 uses anycast addresses to identify multiple interfaces. IPv6 delivers packets addressed to an anycast address to the nearest interface that the address identifies. In contrast to a multicast address, where delivery is from one to many, an anycast address delivery is from one to one-of-many. Currently, anycast addresses are assigned only to routers and are used only as destination addresses.

# IPv6 Addresses Assigned to Hosts and Routers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An IPv6 host, including those with only one interface, typically has multiple IPv6 addresses. By default, link-local addresses are automatically configured for each interface on each IPv6 host or router. To communicate with non-neighboring nodes, a host must also be configured with unicast site-local or global addresses. A host obtains these additional addresses either from router advertisements or by manual assignment. Use commands in the **netsh interface ipv6** context to manually configure IPv6 addresses.

In IPv6, hosts and routers are typically assigned the following addresses:

- **Unicast addresses**:
  - A link-local address for each interface
  - A site-local address for each interface
  - One or more global addresses for each interface
  - The loopback address for the loopback interface
- **Multicast addresses** (to listen for multicast traffic):
  - The interface-local scope all-nodes address (FF01::1)
  - The link-local scope all-nodes address (FF02::1)
  - The solicited node address for each unicast address on each interface
  - The multicast address for each joined group on each interface

In addition, IPv6 routers also have the following addresses:

- **Multicast addresses**:
  - The interface-local scope all-routers address (FF01::2)
  - The link-local scope all-routers address (FF02::2)
  - The site-local scope all-routers address (FF05::2)
- **Anycast addresses**:
  - A subnet-router anycast address for each subnet
  - Optional — Additional anycast addresses

Table 1.10 summarizes the major differences between IPv6 and IPv4 addresses.

**Table 1.10 Differences Between IPv4 Addressing and IPv6 Addressing**

| IPv4 Address | IPv6 Address |
|---|---|

| | |
|---|---|
| Internet address classes | N/A |
| Multicast addresses (224.0.0.0/4) | IPv6 multicast addresses (FF00::/8) |
| Broadcast addresses | N/A |
| Unspecified address is 0.0.0.0 | Unspecified address is :: |
| Loopback address is 127.0.0.1 | Loopback address is ::1 |
| Public IP addresses | Aggregatable global unicast addresses |
| Private IP addresses | Site-local addresses (FEC0::/48) |
| Autoconfigured addresses | Link-local addresses (FE80::/64) |
| Dotted decimal notation | Colon hexadecimal format |
| Subnet mask or prefix length notation | Prefix length notation only |
| A resource records | AAAA resource records |

# Routing IPv6 Traffic over an IPv4 Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An eventual successful transition to IPv6 requires interim coexistence of IPv6 nodes in today's predominantly IPv4 environment. To support this, IPv6 packets are automatically tunneled over IPv4 routing infrastructures, enabling IPv6 clients to communicate with each other by using 6to4 or ISATAP addresses and tunneling IPv6 packets across IPv4 networks. For information about automatic tunneling of IPv6 packets, see RFC 2893, "Transition Mechanisms for IPv6 Hosts and Routers."

Support for IPv6 automatic tunneling technologies in Windows XP and Windows Server 2003 includes:

- 6to4, to provide automatic intersite tunnels across the IPv4 Internet.
- ISATAP, to provide automatic intrasite tunnels.

A computer running Windows XP or Windows Server 2003 can automatically configure itself for 6to4 and ISATAP tunneling. The IPv6 Helper service, included with the IPv6 protocol for Windows XP and Windows Server 2003, provides support for 6to4 hosts and 6to4 routers. Use **netsh interface IPv6 isatap** context commands to configure the IPv6 Helper service. In addition, you can configure a computer running Windows XP or Windows Server 2003 as a 6to4 router by enabling the Internet Connection Sharing (ICS) feature on the interface that is connected to the Internet.

Both 6to4 and ISATAP encapsulate an IPv6 packet within an IPv4 header. However, they send the packet across an IPv4 infrastructure in different ways:

- **6to4 uses the IPv6 prefix.** 6to4 uses a public IPv4 address to create the 64-bit subnet identifier portion for an IPv6 address. For example, 131.107.71.152 becomes 2002:836B:4798::/48.
- **ISATAP uses the IPv6 interface ID.** ISATAP uses a locally assigned IPv4 address (public or private) to create a 64-bit interface identifier. For example, 172.31.71.152 becomes ::0:5EFE:172.31.71.152.

In both cases, IPv4 addresses that are embedded in portions of the IPv6 address provide the information to determine the source and destination addresses in the encapsulating IPv4 header.

By deploying 6to4 or ISATAP, you can integrate IPv6 traffic into your IPv4 network environment. Understanding examples of each automatic tunneling technology can help you decide whether to deploy 6to4, ISATAP, or both as you introduce IPv6 on your network.

**Note**

- For an introduction to IPv6, including information about router-to-router, host-to-router, router-to-host, and host-to-host tunneling configurations that underlie 6to4 and ISATAP tunneling, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

#### Using 6to4 for IPv6 Traffic Between Subnets or Between Sites

6to4 is an address assignment and router-to-router automatic tunneling technology that is described in RFC 3056, "Connection of IPv6 Domains via IPv4 Clouds." To facilitate the introduction of IPv6 in current IPv4 environments, IPv6 is designed so that you can use 6to4 to handle traffic between IPv6 nodes without obtaining an IPv6 global address prefix from an IPv6 ISP, and without a direct connection to the IPv6 Internet.

Figure 1.16 shows one way to use 6to4 to handle the following types of traffic:

- **Direct 6to4 host communication within a site (no tunnel).** A 6to4 host can communicate directly with another 6to4 host within the same site. A *6to4 host* is an IPv6 host that is configured with at least one 6to4 address (a global address with the 2002::/16 prefix). Host A and Host B in Figure 1.16 use the local 6to4 router to communicate with each other.
- **Tunnel across the IPv4 Internet by using a 6to4 router.** A 6to4 host can communicate with a non-local 6to4 host by using a tunnel from a local 6to4 router across an IPv4 network (such as the Internet) to a 6to4 router at the destination site. The first 6to4 router encapsulates the packet in an IPv4 header; the receiving 6to4 router removes the IPv4 header

and then forwards the IPv6 packet to the destination 6to4 host. During the first and last stages of the packet's transmission — from the sending 6to4 host to its 6to4 router, and from the recipient 6to4 router to the destination 6to4 host — the IPv6 routing infrastructure in place at each site is used. In Figure 1.16, 6to4 Host A (or 6to4 Host B) sends its packet to 6to4 Router 1, which tunnels it across the IPv4 Internet to 6to4 Router 2, which then forwards the packet to 6to4 Host C.

- **Tunnel across the IPv4 Internet to the IPv6 Internet by using a 6to4 router and a 6to4 relay.** A 6to4 host on an IPv4 network can communicate with an IPv6-only host on the IPv6 Internet by using a tunnel from a local 6to4 router across the IPv4 Internet to a 6to4 relay that then forwards the packet across the IPv6 Internet to the recipient IPv6-only host. In this case, it is the 6to4 relay that removes the IPv4 header and forwards the IPv6 packet to the recipient IPv6-only host. In Figure 1.16, Host A (or Host B) sends its packet to 6to4 Router 1, which tunnels it across the IPv4 Internet to the 6to4 relay, which then forwards the packet to 6to4 Host D.

**Figure 1.16 Using 6to4 to Route IPv6 Packets**




In Figure 1.16, 6to4 Router 2 represents a computer running Windows XP with ICS enabled. The private interface of the ICS computer connects to a single-subnet intranet, and the ICS computer's public interface connects to the IPv4 Internet. The private interface of an ICS computer always uses the private IPv4 address 192.168.0.1.

**Using ISATAP for IPv6 Traffic Between Subnets**

Intrasite Automatic Tunnel Addressing Protocol (ISATAP) is an address assignment and automatic tunneling technology that is described in the Internet Draft "Intrasite Automatic Tunnel Addressing Protocol (ISATAP)." ISATAP enables unicast communication between IPv6/IPv4 nodes in an IPv4 intranet.

ISATAP derives an interface identifier (the last 64 bits of an IPv6 address) from any IPv4 address assigned to the node, either public or private. The ISATAP address format supports configuration of global addresses (including 6to4), site-local addresses, and link-local addresses.

Figure 1.17 shows two IPv6/IPv4 hosts communicating over an IPv4 network by using each other's automatically configured link-local ISATAP address.

**Figure 1.17 Using Link-Local ISATAP Addresses to Route IPv6 Packets on an IPv4 Network**




IPv6/IPv4 hosts can also communicate with non-local IPv6/IPv4 hosts by using ISATAP-derived global addresses, and by using an ISATAP router to tunnel packets through an IPv4 infrastructure. Under the IPv6 protocol that Windows XP and Windows Server 2003 support, you can use either of the following methods to configure the intranet IPv4 address of an ISATAP router:

- **Name resolution (preferred).** For computers running Windows XP (SP1 or later) or Windows Server 2003, automatic resolution of the name ISATAP to an IPv4 address. To ensure successful name resolution, name the computer used as the ISATAP router ISATAP. A computer running Windows XP or Windows Server 2003 then automatically registers the appropriate records in DNS and WINS. For computers running Windows XP (earlier than SP1), the name resolved is _ISATAP.

- **Netsh commands for Interface IPv6.** Manual configuration by using commands in the Netsh Interface IPv6 context.

An ISATAP host sends an IPv4-encapsulated Router Solicitation message to a configured ISATAP router. The ISATAP router responds with an IPv4-encapsulated unicast Router Advertisement message that contains prefixes for use in autoconfiguring ISATAP-based addresses. This additional configuration is needed only when the host's subnet does not contain an IPv6 router.

The example in Figure 1.18 shows how two ISATAP hosts that use 6to4 prefixes can communicate across the Internet even though each site is using the 192.168.0.0/16 private address space.

**Figure 1.18 Using 6to4 and ISATAP to Route IPv6 Packets Across the IPv4 Internet**




### Note

- Hosts running Windows XP or Windows Server 2003 determine whether to use 6to4, ISATAP, or both depending on their IPv4 configuration.

# Configuring DNS for IPv6/IPv4 Coexistence

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Through *DNS dynamic update*, DNS client computers register and dynamically update their resource records with a DNS server whenever an IP address changes. This reduces the need to manually administer zone files, especially for clients that frequently move or change locations and that use DHCP to obtain an IP address.

In an IPv4 environment, by default the DNS Client service on computers running Windows 2000, Windows XP, or Windows Server 2003 dynamically updates host (A) resource records (RRs) in DNS. If all hosts on your network run those operating systems, DNS dynamic updates are automatic.

However, on hosts that do not support dynamic update, you must either enable dynamic update or manually add or update their DNS records. The same is true on a network to which IPv6 has been introduced: hosts that do not support dynamic update must have dynamic update enabled or must have DNS records added manually. IPv6 has the additional requirement that IPv6 nodes use a new type of address resource record, known as *AAAA (quad-A) resource records*, to resolve a fully qualified domain name to an IPv6 address. (Four "A"s are used for the name of these resource records because 128-bit IPv6 addresses are four times as large as 32-bit IPv4 addresses.)

Systems that support IPv6 use the same domain names as the domain names used in IPv4 but have both IPv6 and IPv4 addresses registered in DNS. The DNS Server service in Windows Server 2003 and Windows 2000 support processing for DNS IPv6 host records as defined in RFC 1886, "DNS Extensions to Support IP Version 6."

An IPv6 host sends DNS name queries to the DNS server to resolve host names to IPv6 addresses. The AAAA resource records stored on the DNS server provide the mapping from a host name to its IPv6 address.

DNS traffic is also supported over IPv6 for both client and server. The client and server are configured for IPv6 over DNS using anycast or unicast DNS server IP addresses. For more information, see "IPv6 configuration items" in Help and Support Center for Windows Server 2003.

Because IPv6 addresses are too long to remember easily, you can populate your DNS servers with IPv6 address resource records to support IPv6 name-to-address resolutions and optionally with pointer resource records to support IPv6 address-to-name resolutions:

- **Address Resource Records.** To successfully resolve names to addresses, the DNS infrastructure must contain the following resource records, populated either manually or dynamically:
  - A resource records for the IPv4 addresses of IPv4 nodes.
  - AAAA resource records for the IPv6 addresses of IPv6 nodes. The following is an example of a AAAA resource record:

    ```
    host1.microsoft.com     IN     AAAA     FEC0::2AA:FF:FE3F:2A1C
    ```

- **Pointer (PTR) Resource Records (optional; not recommended).** The DNS infrastructure can also contain the following resource records, populated either manually or dynamically, to resolve addresses to host names in reverse queries:
  - PTR records in the IN-ADDR.ARPA domain for the IPv4 addresses of IPv4 nodes.
  - PTR records in the IP6.ARPA domain for the IPv6 addresses of IPv6 nodes. (Recall that RFC 3152 specifies that IP6.INT be phased out and replaced by IP6.ARPA.) The IP6.INT domain was created specifically for IPv6 reverse queries. To create the namespace for reverse queries, each hexadecimal digit in the 32-digit IPv6 address (zero compression and double-colon compression notation cannot be used) becomes a separate level in inverse order in the reverse domain hierarchy. Therefore, the reverse lookup domain name for the address FEC0::2AA:FF:FE3F:2A1C is:

    C.1.A.2.F.3.E.F.F.F.0.0.A.A.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.C.E.F.IP6.INT

Integrating PTR resource record support into your DNS infrastructure is not recommended.

For name-to-address resolution, after the querying node obtains the set of addresses corresponding to the name, that node must determine the best set of addresses to use as the source and destination for outbound packets.

While name-to-address resolution is fairly straightforward in an IPv4-only environment, it becomes more complex in an environment in which IPv4 and IPv6 coexist. In the mixed IPv6/IPv4 scenario, a DNS query can return both IPv4 and IPv6 addresses. The querying host is configured with at least one IPv4 address and, typically, multiple IPv6 addresses. Determining the type of address (IPv4 versus IPv6), and then the scope of the address (for IPv4, public versus private; for IPv6, link-local versus site-local versus global versus coexistence), for both the source and the destination addresses is complex.

Two algorithms, one to select the source address and another to select the destination address, specify default behavior for IPv6 implementations. These algorithms do not override choices made by applications or upper-layer protocols, nor do they preclude the development of more advanced mechanisms for address selection. The two algorithms include an optional mechanism that lets you override the default behavior. In dual-stack implementations, the destination address selection algorithm considers both IPv4 and IPv6 addresses, and determines whether it prefers IPv6 addresses over IPv4 addresses, or vice-versa.

For more information about default address selection rules for IPv6, including the source address selection algorithm and the destination address selection algorithm, see the Internet Draft "Default Address Selection for IPv6."

For an introduction to IPv6 and more information about Windows Server 2003 IPv6, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit), or see the IPv6 link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Enabling IPv4 Applications for IPv6

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the PortProxy service as an application-layer gateway for nodes or applications that do not support IPv6. PortProxy facilitates the communication between nodes or applications that cannot connect using a common address type, Internet layer protocol (IPv4 or IPv6), and TCP port. The primary purpose of the service is to allow IPv6 nodes to communicate with IPv4 TCP applications.

PortProxy relays TCP traffic from IPv4 to either IPv4 or IPv6, or from IPv6 to either IPv6 or IPv4. In the context of IPv6/IPv4 coexistence or migration, use the PortProxy service to enable any of the following scenarios:

- An IPv6 node accessing an IPv4-only application that is running on an IPv4 node.
- An IPv4-only node accessing an IPv6-only node.
- An IPv6-only node accessing an IPv4-only node.

The Netsh commands for Interface Portproxy provide a command-line tool for administering servers that act as proxies between IPv4 and IPv6 networks and applications. For more information about how to use the Netsh Interface PortProxy commands, see the Netsh command-line help, or see "Netsh commands for Interface Port Proxy" in Help and Support Center for Windows Server 2003.

**Note**

- The PortProxy service transmits only TCP traffic for application-layer protocols that do not embed address or port information in the TCP segment. For example, the File Transfer Protocol (FTP), which embeds addresses when using the FTP Port command, does not work across a PortProxy computer. Unlike NAT, the PortProxy service does not include an equivalent to NAT editors.

# Testing Your Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After acquiring any new hardware and software that your network design requires, systematically measure the new solution against your organization's business and technical goals. Testing your design before deploying it in a production environment ensures that those goals are met with minimum impact.

Predeployment testing lets you assess the performance characteristics of network devices and technologies. Testing also helps you identify deployment-related risks, and instills confidence in the deployment process throughout your organization.

Figure 1.19 shows the process for testing a TCP/IP network design.

**Figure 1.19 Testing Your Network Design**




**Reviewing Industry Tests**

Vendors, trade journals, and independent test labs extensively test devices and other network solutions. You might find their published results useful for validating or rejecting assumptions. Keep in mind that most lab tests are component tests rather than system tests and can fail to measure how a particular network design might impact the performance of the specific device or technology.

**Using Network Testing Tools**

Use the following types of tools to test your network design:

- Modeling and simulation tools
- Network management and monitoring tools

**Modeling and simulation tools**

Use statistical analysis and modeling techniques to simulate a mathematical model of a network. By creating a model, you can isolate potential performance problems before you actually deploy any part of an IP network. In most cases, these tools do not measure actual traffic behavior, so evaluate the results with this limitation in mind.

**Network management and monitoring tools**

Typically, you use network management and monitoring tools after deploying a network. However, these tools can also help you test your IP network design in a lab. You can use a number of effective commercially available network management applications to identify problems and potential problems on your test network.

Many of these applications run on dedicated network management stations (NMSs) and communicate with internetworking devices using Simple Network Management Protocol (SNMP) or Remote Monitoring (RMON). By using data supplied by an SNMP or RMON Management Information Base (MIB) located on the devices, a network management application can isolate performance problems in a proposed network design.

Windows Server 2003 includes the Network Monitor tool (Netmon.exe), a protocol analyzer that you can use to monitor a new network design. Network Monitor captures and displays packets, analyzing their traffic patterns, rate of broadcast, errors, utilization, and other aspects of their behavior.

The Network Monitor component that ships with Windows Server 2003 can capture frames that are sent to or from the computer on which Network Monitor is installed. To capture frames that are sent to or from a remote computer, you can use the Network Monitor component that ships with Microsoft® Systems Management Server (SMS), which can capture frames sent to or from any computer on which the Network Monitor driver is installed.

For more information about the Network Monitor component, see Help and Support Center for Windows Server 2003.

# Additional Resources for Designing a TCP/IP Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this chapter.

**Related Information**

- "Deploying IPSec" in this book for more information about using Internet Protocol security (IPSec).
- "Deploying ISA Server" in this book for more information about deploying Network Address Translation (NAT).
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about TCP/IP, IPSec, and IPv6 in Windows Server 2003.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for technical information about unicast IP routing, including the NAT routing protocol component of the Routing and Remote Access service.
- "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit for more information about inventorying your network hardware and software and creating a map of your network topology.
- *Cisco Internetwork Design* by Matthew Birkner, 2000, Indianapolis, IN: Cisco Press for more information about the three-tier network design model.
- *Top-Down Network Design* by Priscilla Oppenheimer, 1999, Indianapolis, IN: Cisco Press/Macmillan Technical Publishing for more information about the three-tier network design model.
- *Understanding IPv6* by Joseph Davies, 2002, Redmond, WA: Microsoft Press.
- *Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference* by Joseph Davies and Thomas Lee, 2002, Redmond, WA: Microsoft Press.
- *Routing in the Internet (2nd Edition)* by Christian Huitema, 2000, Upper Saddle River, NJ: Prentice Hall PTR.
- *Interconnections (2nd Edition)* by Radia Perlman, 2000, Reading, MA: Addison-Wesley.

**Related Tools**

- Netsh commands for Interface IPv6

  You can use the Netsh commands for Interface IPv6 to manage configuration of the IPv6 protocol. For more information about how to use the Netsh commands for Interface IPv6, see the Netsh command-line help or see "Netsh commands for Interface IPv6" in the Help and Support Center for Windows Server 2003.

- Netsh commands for Interface Portproxy

  The Netsh commands for Interface Portproxy provide a command-line tool for administering servers that act as proxies between IPv4 and IPv6 networks and applications. For more information about how to use the Netsh Interface PortProxy commands, see the Netsh command-line help or see "Netsh commands for Interface PortProxy" in Help and Support Center for Windows Server 2003.

- Ipsec6.exe

  For experimenting with IPSec for IPv6, you can use the Ipsec6 tool to configure IPSec policies and security associations in

an IPv6 environment. For more information about Ipsec6, see "IPv6 utilities" in Help and Support Center for Windows Server 2003.

- Network Monitor (Netmon.exe)

  The Network Monitor tool (Netmon.exe) is a protocol analyzer that you can use to monitor a new network design. For more information about Netmon.exe, see "Network Monitor" in Help and Support Center for Windows Server 2003.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Using multicast scopes" in Help and Support Center for Windows Server 2003.

- "Netsh commands for Interface Portproxy" in Help and Support Center for Windows Server 2003.

# Deploying DHCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Dynamic Host Configuration Protocol (DHCP) in the Microsoft® Windows® Server 2003 family of operating systems enables centralized automatic management of IP addresses and other TCP/IP settings for network clients. You can reduce administrative overhead in your organization by designing and implementing a reliable and scalable DHCP solution.

**In This Chapter**


Overview of DHCP Deployment

Creating Your DHCP Server Design

Integrating DHCP with Other Services

Defining Scopes

Implementing Your DHCP Solution

Example DHCP Implementation

Additional Resources for Deploying DHCP


**Related Information**

- For more information about Dynamic Host Configuration Protocol (DHCP), see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

- For more information about integrating DHCP with Domain Name System (DNS), see "Deploying Domain Name System (DNS)" in this book.

- For more information about integrating DHCP with Windows Internet Name Service (WINS), see "Deploying WINS" in this book.

# Overview of DHCP Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Valid IP addresses and other network options must be configured for all computers and other devices on the network, such as printers, in a corporate network. Manually configuring this information is a time-consuming process that adds significant costs to an organization and is susceptible to user error.

Windows Server 2003 Dynamic Host Configuration Protocol (DHCP) reduces the complexity and administrative overhead involved in managing network client IP addressing and configuration. DHCP allows you to assign IP addresses to network clients automatically and dynamically, as needed, and to automate, centralize, and simplify IP address and option configuration and distribution across your network. This protects against common configuration errors that occur when values are entered manually at each computer and helps to prevent address conflicts.

By using DHCP options, you can configure DHCP servers to supply a full range of configuration values when assigning a DHCP lease, allowing you to configure a large number of computers at one time, and to change configuration as necessary.

Different types of organizations can benefit from the automation and centralization that Windows Server 2003 DHCP provides, including:

- Organizations in which administrators are responsible for configuring IP addresses and options for a large number of devices.
- Organizations that include a large population of mobile users or network configurations that frequently change.
- Organizations that already use versions of DHCP earlier than Microsoft® Windows® 2000, which can benefit from the added functionality of Windows Server 2003 DHCP by upgrading their DHCP deployment.

If you are deploying DHCP in a new Windows Server 2003 environment, begin your deployment process by creating a design for your DHCP servers. When you have completed all of the design steps for your DHCP infrastructure, you can implement your DHCP solution by configuring your DHCP servers. If you are upgrading your existing DHCP infrastructure to Windows Server 2003, begin by reviewing your current DHCP design and modifying it if required, and then migrate your existing DHCP databases.

# DHCP Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying or upgrading DHCP in your organization involves several important planning, design, and implementation steps. Figure 2.1 shows the process for deploying DHCP.

**Figure 2.1 Deploying DHCP**



# Creating Your DHCP Server Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to create a DHCP server design that meets the needs of your organization in terms of functionality, availability, interoperability, and total cost of ownership (TCO).

Figure 2.2 shows the process for creating a DHCP server design.

**Figure 2.2 Designing Your DHCP Server**



# Upgrading Your Existing DHCP Server Hardware

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Determine whether your current hardware and software, including routers, switches, and other servers and clients, support Windows Server 2003 DHCP. Windows Server 2003 DHCP servers support Windows DHCP clients and third-party operating systems that use DHCP and comply with RFC 2131. For more information about the operating systems that support DHCP, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

If your current systems support Windows Server 2003, but are close to the end of their expected lifecycle, consider upgrading your hardware at the same time that you upgrade to Windows Server 2003. Upgrading DHCP servers running Microsoft® Windows NT® Server version 4.0 or earlier to Windows Server 2003 allows you take advantage of benefits related to the Active Directory directory service, such as integrated secure dynamic updates of the DNS database.

For information about hardware life expectancy, contact your hardware vendor or refer to any internal metrics that your organization might have developed. For information about hardware compatibility, see the Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about performing a hardware inventory, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

# Determining DHCP Server Locations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To determine where to locate your DHCP servers, consider whether you are deploying a distributed, centralized, or combined DHCP infrastructure.

For a distributed DHCP infrastructure, locate a DHCP server on each subnet. Because distributed infrastructures use a DHCP server on each subnet, they require a greater number of servers than centralized networks. For example, a network that includes 30 subnets and that is using a true distributed topology requires at least 30 DHCP servers, and possibly more to provide for redundancy.

In a centralized DHCP infrastructure, DHCP servers are placed in a central location. A centralized DHCP topology requires the deployment of DHCP/bootstrap protocol (BOOTP) relay agents. Additional hardware resources are not generally required for DHCP relay agents; in most cases, the routers that are positioned between each subnet can assume this role, as defined in RFC 1542. If the routers cannot relay DHCP messages, configure a computer running Windows Server 2003 to act as a DHCP/BOOTP relay agent. For more information about configuring relay agents, see "Enabling DHCP Support for Multiple Subnets" later in this chapter.

Combining both distributed and centralized DHCP infrastructures provides the maximum efficiency for your network. In a combined DHCP infrastructure, the locations for DHCP servers are based on the physical characteristics of the local area network (LAN) or wide area network (WAN) infrastructure, and not the logical groupings defined by the Active Directory logical structure. DHCP servers are not required for every subnet if the connecting routers support DHCP/BOOTP relay agents. You can administer Windows Server 2003 DHCP servers remotely from a computer running Windows Server 2003 and Microsoft Management Console (MMC) DHCP snap-in. You can also administer a DHCP server remotely at the command line by using Netsh commands for DHCP, or you can remotely administer a DHCP server from a computer running Microsoft® Windows® XP Professional that has the Windows Server 2003 Administration Tools Pack installed. You must have the correct level of security permissions in order to administer a DHCP server. For more information about using the Windows Server 2003 Administration Tools Pack, see "Windows Server 2003 Administration Tools PackWindows Server 2003 Administration Tools Pack" in Help and Support Center for Windows Server 2003.

# Optimizing DHCP Server Performance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can optimize the performance of DHCP servers in your organization by doing the following:

- Extending the duration of the IP address lease.
- Improving DHCP server hardware, specifically upgrading to a faster hard drive or adding random access memory (RAM).

**Extending the IP Address Lease and Renewal Duration**

The volume of traffic on your network can have a negative impact on DHCP performance. For example, a subnet that relies on a DHCP server at a remote location on the WAN might experience poor performance at start of day, when users turn on computers and a large load of requests might be sent over the network.

DHCP traffic does not use significant network bandwidth during periods of normal usage; however, the following two phases of DHCP client configuration generate some network traffic load:

- IP address lease
- IP address renewal

When a DHCP client initializes TCP/IP or renews its address lease, it acquires an IP address from the DHCP server. This process results in an exchange between the DHCP client and the DHCP server, which typically consists of four packets, each containing a maximum of 4 kilobytes (KBs). For more information about the DHCP exchange, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

You can reduce the amount of network traffic generated by DHCP IP address lease and IP address renewal by extending the lease duration. Before you extend the lease duration, you must take into consideration other factors in your network, such as ratio of clients to available IP addresses, or clients that frequently lease addresses on more than one subnet, such as laptops that move frequently. If you have a relatively stable network and many more available IP addresses than DHCP clients, increasing the lease duration reduces network traffic, because DHCP messages are sent less frequently. If, however, you have a limited number of IP addresses available to your DHCP clients, or a network that changes frequently, a longer lease duration might cause you to run out of IP addresses because IP addresses are not returned to the address pool and made available to other DHCP clients until the lease expires. For more information about extending the lease duration, see "Determining Lease Duration" later in this chapter.

**Improving DHCP Server Hardware**

You can optimize DHCP performance in your system by optimizing individual DHCP server performance. Windows Server 2003 includes performance monitoring tools that you can use to test and monitor your servers.

For more information about performance monitoring tools, see "Performance Monitoring Tools" in Help and Support Center for Windows Server 2003.

The primary factors that impact DHCP server performance include:

- The speed of the server disk drives.
- The amount of RAM installed in the DHCP server computer.

The greatest volume of disk usage occurs when the service is started and when the database is backed up. When planning your DHCP server hardware specifications, evaluate the average time required for disk access and for disk read/write operations. If necessary, maximize DHCP server performance by increasing RAM and purchasing high-speed disk drives for the servers.

# Enabling DHCP Support for Multiple Subnets

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have multiple subnets in your network, and do not have a DHCP server on every subnet, determine whether your current routers relay DHCP/BOOTP messages.

If your routers cannot be used for DHCP/BOOTP relay, set up a DHCP/BOOTP relay agent on at least one computer running Windows Server 2003 on each subnet. The DHCP/BOOTP relay agent relays DHCP and BOOTP message traffic between the DHCP-enabled clients on the local network and a remote DHCP server located on another physical network by using the IP address of the remote DHCP server.

Figure 2.3 shows a simple, routed network in which the router acts as a DHCP relay agent.

**Figure 2.3 Subnets Configured to Use a DHCP Relay Agent**




If your routers cannot be used for DHCP/BOOTP relay and you choose not to configure DHCP/BOOTP relay agents, you must configure your network so that a DHCP server has a network adapter on each subnet it serves. You can accomplish this by either placing a DHCP server on each subnet, or by multihoming DHCP servers. This distributed configuration does not provide fault tolerance. If a DHCP server becomes unavailable, DHCP clients on the subnet cannot receive IP addresses and options.

**Note**

- The DHCP Relay Agent service is available only on computers running Windows Server 2003, Microsoft® Windows® 2000, or Windows NT 4.0. To use the DHCP Relay Agent routing protocol, the Routing and Remote Access service must be installed and enabled.

For more information about installing and configuring DHCP relay agents, see "Configure the DHCP Relay Agent" in Help and Support Center for Windows Server 2003. For more information about the DHCP Relay Agent service, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) or see Help and Support Center for Windows Server 2003.

# Determining How Many DHCP Servers to Deploy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A single DHCP server can serve an almost unlimited number of clients. However, factors such as the size and layout of your network, the IP address class selected for use, and the volume of traffic on your network often make this impractical. You can deploy multiple DHCP servers to reduce the volume of DHCP-related traffic across your network and create faster response times for DHCP messages. Deploying multiple DHCP servers also creates fault tolerance on your network. If you choose to deploy more than one DHCP server, it is important to weigh the benefits of increased response times against the costs required for additional hardware.

When deciding how many DHCP servers you need, consider:

- The location of DHCP-enabled clients on your network.

- The transmission speeds between the segments for which DHCP service is provided.

  If you have slower WAN or dial-up links, place a DHCP server on both sides of these links to improve DHCP response times for local clients.

- The network traffic that DHCP produces, as well as your current network traffic.

  If your current volume of network traffic is high, consider deploying multiple DHCP servers to reduce the volume of DHCP requests traveling across the network. Make sure to account for periods when network traffic is heaviest, such as the beginning of the day, when many users turn on their computers at the same time.

- Disk space requirements.

  It is important to consider the database size when choosing your hardware. Each lease requires approximately 600 bytes per lease for the database, plus 1200 bytes for backup (600 bytes for the backup and 600 bytes for the temporary directory). In addition, the audit logs require approximately 500 bytes per lease transaction and are stored for seven days.

  **Tip**

  - In general, allow at least 50-70 MB for the audit logs, however the number of lease transactions depends on the number of leases as well as the lease duration.

  To figure out how much hard disk space is required, first multiply the number of leases by 600 bytes, then multiply the estimated number of lease transactions by 500 bytes, and add these two results. The sum is the minimum amount of disk space required by the DHCP server.

  For example, a DHCP server with 10,000 leases and lease duration of one week requires approximately 18 MB to store the leases and the backup (6 MB for the database, 6 MB for the backup, and 6 MB for the temporary database). The audit logs would require an absolute minimum of 10 MB: 5 MB (500 bytes x 10,000 leases) for startup, and 5 MB when the leases renew halfway through the week. If the number of leases increases or if lease time is shortened, this requirement will increase. A company might allocate 100 MB for audit logs to allow for flexibility in adding leases or reducing lease duration, as well as dealing with any peak-load events.

# Optimizing DHCP Availability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A highly available solution must account for all possible points of failure, including server failures, WAN link interruptions, and router failures. You can increase the fault tolerance and availability of your design by using one or more of the following, depending on your needs and hardware cost considerations:

- Split-scope configurations
- Clustered DHCP servers
- Using standby servers

# Using Split-Scope Configurations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can increase fault tolerance by splitting DHCP scopes between multiple DHCP servers. With a split-scope configuration, if one server becomes unavailable, the other server can take its place and continue to lease new IP addresses or renew existing clients. Splitting DHCP scopes also helps to balance server loads.

When splitting the IP address pool of a scope between two servers, assign the same scope to both servers, and exclude opposite portions of the address range. You also need to make identical reservations at both DHCP servers, so that either server can assign the reserved IP address, ensuring that the intended device receives the address that is reserved for its use.

Figure 2.4 shows a network that is using a split-scope configuration.

**Figure 2.4 Split-Scope Configuration**




In Figure 2.4, DHCP Server 1 has 80 percent of the addresses in the scope and DHCP Server 2 has 20 percent of the addresses in the scope. Splitting a scope between servers in this way, which is commonly referred to as the "80/20 rule," often relies on the proximity of the DHCP servers to the clients it serves. For example, when a DHCP client that is on the same subnet as DHCP Server 1 sends out a DHCP Discover packet, it takes longer for DHCP messages from clients to reach the DHCP Server 2 than DHCP Server 1, because DHCP Server 2 is on the other side of a router from the DHCP client. You can also configure a delay on the DHCP relay agent to ensure the local DHCP server has adequate time to respond. Because DHCP clients always accept the lease from the DHCP server that sends the first response, clients normally obtain leases from DHCP Server 1. If DHCP Server 1 goes offline for any reason, clients accept leases from DHCP Server 2.

# Using Clustered DHCP Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 DHCP can use Windows Clustering, which allows two or more servers to be managed as a single system. You can increase DHCP (or multicast address dynamic client allocation protocol [MADCAP]) scalability, availability, and reliability by using the Cluster service to deploy a DHCP server cluster.

By using clustering support for DHCP, you can implement a local method of DHCP server failover. In this way, you can achieve greater fault tolerance and minimize disruptions and work stoppages. Windows Clustering can automatically detect the failure of an application or server and restart the application on or transfer the server role to an alternate server. Users experience only a brief break in service.

Windows Clustering creates a virtual DHCP server so that if one of the clustered nodes fails, the namespace and all of the services contained in that node are automatically transferred to a second node. No changes are visible to the client, which sees the same IP address for the clustered DHCP servers.

Use Windows Clustering alone to create a fault-tolerant design that makes efficient use of available IP addresses. To further enhance DHCP fault tolerance and availability, combine DHCP server clustering with a remote failover configuration — such as a split-scope configuration across different segments of your network. Although combining server clustering with a split-scope configuration increases DHCP availability, you must consider whether the benefits to your organization outweigh the hardware costs involved.

Figure 2.5 shows an example of clustered DHCP servers. DHCP Server 1 is the active DHCP server, and DHCP Server 2 is the backup DHCP server.

**Figure 2.5 Clustered DHCP Servers**




**Example: Using a Split-Scope Configuration in Combination with a DHCP Server Cluster to Enhance Availability and Fault Tolerance**
A company has its main corporate office in North America, and two European offices in Milan and Seville. The North American DHCP server cluster includes two servers, which are configured as nodes of the server cluster.

Twenty percent of the available IP leases for the Milan and Seville sites are configured on the North American server cluster. The remaining 80 percent of the available IP leases are configured on the local Milan and Seville DHCP servers.

This configuration provides the following levels of fault tolerance for the European sites:

1. Under normal circumstances clients in Milan and Seville request IP lease assignments from the local DHCP server, which contains 80 percent of the available IP addresses for the subnet.

2. If either the Milan or Seville DHCP server is slow or unavailable, European DHCP clients use the North American DHCP server, which contains 20 percent of the available IP addresses for the Milan and Seville subnets. The relay agent between the European offices and the North American office is configured to delay the DHCP messages from the European offices, allowing the local (Milan or Seville) DHCP server enough time to respond.

3. Because the North American DHCP server is configured as a server cluster, if one node of the cluster becomes unavailable, Windows Clustering automatically brings the other node online.

# Using Standby Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A standby server and its scopes are not activated for use under normal conditions, and are activated by the administrator only when needed, such as when a DHCP server fails or is taken offline for an extended period of time. Standby servers require manual administration to ensure failover transition, and therefore might not be as effective as other failover methods, such as split scopes and clustered servers.

To use a standby configuration, configure an additional DHCP server to server as a backup if the primary server goes offline. You can either configure the standby server to be identical to your primary DHCP server or configure the standby server with unused scopes to temporarily replace the primary DHCP server.

If you are configuring the standby server with the identical scope to your primary DHCP server, you must implement server-side address conflict detection to prevent the assigning of duplicate addresses.

**Important**

- Because server-side conflict detection uses Address Resolution Protocol (ARP) and Internet Control Message Protocol (ICMP) messages to detect conflicts, Internet Connection Firewall (ICF) or other firewalls that are installed on clients on your network might interfere with conflict detection.

For more information about backing up your DHCP servers, see "Backing up the DHCP databaseBacking up the DHCP database" or "Netsh commands for DHCP" in Help and Support Center for Windows Server 2003.

# Integrating DHCP with Other Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you use DHCP servers for Microsoft network clients, you must use a name resolution service. Networks that support clients running Windows 2000, Microsoft® Windows® XP Professional, and Windows Server 2003 use the DNS service to support name resolution. Networks that support clients running versions of the operating system earlier than Windows 2000 must use a form of NetBIOS name resolution, such as WINS. Networks that support both types of clients must implement both WINS and DNS servers.

Windows Server 2003 DHCP provides support for both DNS dynamic updates and secure DNS dynamic updates. DHCP and DNS work together to perform dynamic updates and work with Active Directory to perform secure DNS dynamic updates. DHCP also works with Active Directory to prevent unauthorized DHCP servers from running on the network.

Figure 2.6 shows the process for integrating DHCP with other services.

**Figure 2.6 Integrating DHCP with Other Services**

For detailed instructions on using this template, right click on the blank page with all objects de-selected. Choose the menu items you want to view information about listed at the bottom of the list of menu items.



# Configuring Dynamic Update and Secure Dynamic Update

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Server 2003 DHCP Server service can be configured to perform DNS dynamic updates and secure DNS dynamic updates for DHCP clients, which eliminates the need for administrators to update DNS records manually when a client's IP address changes. Clients running Windows 2000, Windows XP, or Windows Server 2003 can also perform dynamic updates.

Clients running versions of Windows earlier than Windows 2000 do not support DNS dynamic update. To enable the DHCP server to perform DNS dynamic updates on behalf of these clients, use the default client preference settings. Clients using WINS for name resolution cannot make an explicit request for DNS dynamic update protocol preference. For these clients, the DHCP service can be configured to update both the PTR and the A resource records.

By itself, dynamic update is not secure; any client can modify DNS records. When secure dynamic update is configured, the authoritative name server accepts updates only from clients and servers that are authorized to make dynamic updates to the appropriate objects in Active Directory. Secure dynamic update is available only on Active Directory–integrated zones. To configure secure dynamic updates, you can use the Windows Server 2003 secure dynamic update feature.

Secure dynamic update protects zones and resource records from being modified by unauthorized users by enabling you to specify the users and groups that can modify zones and resource records. By default, Windows Server 2003, Windows XP Professional, and Windows 2000 clients attempt unsecured dynamic updates first. If that request fails, they attempt secure updates.

When using multiple DHCP servers and secure dynamic updates, add each of the DHCP servers as members of the DnsUpdateProxy global security group so that any DHCP server can perform a secure dynamic update for any record. Otherwise, when a DHCP server performs a secure dynamic update for a record, that DHCP server is the only computer that can update the record.

**To configure dynamic update for DHCP clients and servers**

1. In the DHCP snap-in, select and right-click the DHCP server you want to configure, and then click **Properties**.
2. In the *server name***Properties** dialog box, click the **DNS** tab.
3. On the **DNS** tab, select the **Enable DNS dynamic updates according to the settings below** check box.
4. On the **DNS** tab, select the dynamic update method you want: either always updating DNS A and PTR, or only updating the records when requested by the DHCP client.

Use the DNS snap-in to enable secure dynamic update. For more information about dynamic update and secure dynamic update, see "Deploying Domain Name System (DNS)" in this book and in Help and Support Center for Windows Server 2003.

**Important**

- If DHCP will perform DNS dynamic updates, do not install it on a domain controller. Instead, install DHCP on a member server. When DHCP is installed on a domain controller and is configured to perform dynamic updates on behalf of clients in DNS zones that are configured to allow only secure dynamic update, specify a user account to update the DNS records. For more information about installing DHCP, see "Checklist: Installing a DHCP server" in Help and Support Center for Windows Server 2003.

# Authorizing DHCP Servers in Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An unauthorized DHCP server on a network can cause a variety of problems, such as the leasing of incorrect IP addresses and options. To protect against this type of problem, when a Windows 2000 or Windows Server 2003 domain member DHCP server attempts to start on the network, it first queries Active Directory. The DHCP server compares its IP address and server name to the list of authorized DHCP servers. If either the server name or IP address is found on the list of authorized DHCP servers, the server is authorized as a DHCP server. If no match is found, the server is not authorized in Active Directory and does not respond to DHCP traffic. The process of authorizing DHCP servers is useful for only Windows 2000–based or Windows Server 2003–based DHCP servers. This process cannot be used for DHCP servers running Windows NT Server, or servers running non-Windows-based DHCP services. Only a member of the Enterprise Admins group can authorize or unauthorize a DHCP server in Active Directory.

**Important**

- You must be logged in as an enterprise administrator to authorize a DHCP server.

**To authorize a DHCP server in Active Directory**

1. In the DHCP snap-in, right-click **DHCP**.
2. Select **Manage authorized servers**.
3. In the **Manage Authorized Servers** dialog box, click **Authorize**.
4. In the **Authorize DHCP Server** dialog box, type the name or IP address of the DHCP server, and then click **OK**.

**Note**

- Detection of unauthorized DHCP servers requires the deployment of Active Directory and the DHCP service running on Windows 2000 or Windows Server 2003. Other DHCP servers do not attempt to determine whether they are authorized by Active Directory before offering IP address leases.

# Defining Scopes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before DHCP clients can use a DHCP server for dynamic TCP/IP configuration, you must define and activate scopes for your DHCP clients. A *scope* is the full, consecutive range of possible IP addresses for a subnet. The IP addresses in a scope define a single subnet on which DHCP services are offered. DHCP servers use scopes to manage network IP address distribution and the configuration of DHCP options.

Figure 2.7 shows the process for defining scopes.

**Figure 2.7 Defining Scopes**



# Creating Scopes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must create a DHCP scope for each subnet in your network. Each subnet has a DHCP scope, with a single continuous range of IP addresses.

Before you create scopes, you must install DHCP on your server. For more information about installing DHCP, see "Checklist: Installing a DHCP server" in Help and Support Center for Windows Server 2003.

You can use the DHCP MMC snap-in to create a new scope on your DHCP server.

**To create a DHCP scope**

1. In the DHCP snap-in, select the name of the DHCP server.

2. Select **Action**, andthen select **New Scope**. This opens the **New Scope Wizard**. Complete the New Scope Wizardby configuring the following properties:

    - A scope name and description of the scope.

    - A consecutive range of possible IP addresses.

    - A unique subnet mask. For more information about defining subnet masks, see "Designing a TCP/IP Network" in this book.

    - The IP addresses that are to be excluded from the scope.

    - Lease duration values.

    - Options.

For a worksheet to assist you in completing the New Scope Wizard, see the "Scope Data Collection Worksheet" (DNSDHC_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see the "Scope Data Collection Worksheet" on the Web at http://www.microsoft.com/reskit).

After you create the scope, you can set reservations and configure any additional DHCP options. For information about setting reservations, see "Creating Reservations" later in this chapter. For information about configuring additional DHCP options, see "Configuring DHCP Options" later in this chapter.

# Setting Exclusion Ranges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To prevent address conflicts, the scopes that you define must exclude the IP addresses of devices that you statically configure, such as DHCP servers. By setting exclusion ranges, an administrator can exclude IP address ranges within a scope so that those addresses are not offered to DHCP clients. When you create a new scope, immediately exclude addresses of the existing statically configured computers. Excluded IP addresses can be active on your network, but only when these addresses are manually configured or distributed as reserved IP addresses. For more information about reservations, see "Creating Reservations" later in this chapter.

You can set exclusion ranges in the **Add Exclusions** page of the **New Scope Wizard**. For more information about using the New Scope Wizard, see "Creating Scopes" earlier in this chapter.

**To set an exclusion range after a scope is created**

1. In the DHCP snap-in, expand the scope you want to configure.
2. Select and right-click **Address Pool** under the appropriate scope.
3. Select **New Exclusion Range**.
4. In the **Add Exclusion** dialog box, type the starting and ending IP addresses of the exclusion range, and then click **Add**.

# Determining Lease Duration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a scope is created, the default lease duration is set to eight days. However, because lease renewal is an ongoing process that can affect the performance of DHCP clients and your network, you can increase or decrease the lease duration to fit your specific needs. Determine what segments of your network have specific lease duration requirements, and decide how best to modify lease duration settings to improve DHCP performance on your network.

You can set the lease duration in the **Lease Duration** page of the **New Scope Wizard**. For more information about using the New Scope Wizard, see "Creating Scopes" earlier in this chapter.

**To set the lease duration after a scope is created**

1. In the **DHCP** snap-in, select and right-click the scope you want to configure.
2. Select **Properties**.
3. In the **Lease duration for DHCP clients** box, adjust the lease time for the scope.

**Increasing the Default Lease Duration**

You can increase the lease duration in a scope to reduce network traffic. Increase the lease duration only if that segment of your network has a large number of IP addresses available and a configuration that rarely changes. Increasing lease duration reduces the rate at which IP addresses are reclaimed when changes occur.

In a more stable environment, you can use a long lease, such as several months. This ensures both that addresses are ultimately recovered, and that DHCP-related network traffic is kept to a minimum.

**Caution**

- Use caution when configuring unlimited lease durations. Even stable environments have a certain amount of client turnover. At a minimum, roving computers might be added and removed, desktop computers might be moved from one office to another, and network adapters might be replaced. If a client with an infinite lease is removed from the network, the DHCP server is not notified, and the IP address cannot be reused.

**Reducing the Default Lease Duration**

Reduce the lease duration for segments of your network that have any of the following:

- A limited number of IP addresses available.
- Client configurations that change frequently.
- Clients that relocate often; for example, because they connect to conference rooms or wireless networks.

Although reducing the lease duration creates more DHCP-related network traffic, it increases the rate at which addresses are returned to the available address pool for reassignment. With an average volume of DHCP request traffic, Windows Server 2003 DHCP has a four-hour default grace period after which an expired lease can be reused. This means that an address is marked for deletion four hours after the lease expires, regardless of lease duration. When the volume of DHCP-related traffic is heavy and no leases are available to service lease requests, DHCP immediately instantiates a cleanup cycle, which reclaims any leases marked for deletion. By default, the cleanup cycle occurs every 60 minutes. You can adjust the duration of the default grace period after which an expired lease is marked for deletion by editing the following key in the registry:

```
HKLM\System\CurrentControlSet\Services\DHCPServer\Parameters\LeaseExtension
```

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up

first.

# Configuring DHCP Options

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DHCP uses options to pass additional IP settings to DHCP clients on a network. Examples of DHCP options include:

- The default gateway IP address
- The DNS server IP address
- The DNS domain name

You can configure DHCP options for specific values and enable them for assignment and distribution to DHCP clients based on server, scope, class, or reserved client levels. For example, you can enable the vendor class option Release on Shutdown for any laptops on your network to allow IP addresses assigned to mobile clients to be reincorporated into the address pool more quickly.

You can configure options for an entire server, a scope, or for a single reserved client. The most specific options (reserved client) take precedence over the least specific options (server). Values configured manually on a client override any DHCP options of any type and of any level.

Using the **New Scope Wizard**, you can configure some scope-level options, including router (default gateway), domain name, DNS servers, and WINS servers. You can also configure options at the server, scope, and reserved-client levels in the DHCP snap-in.

**To configure server-level options**

1. In the DHCP snap-in, expand the server for which you want to configure options.
2. Right-click **Server Options**, and then click **ConfigureOptions**.
3. In the **Server Options** dialog box, select the options you want to configure.
4. In the **Data Entry** section of the **Server Options** dialog box, type the option parameters, and then click **OK**.

**To configure scope-level options**

1. In the DHCP snap-in, expand the scope for which you want to configure options.
2. Right-click **Scope Options**, and then click **Configure Options**.
3. In the **Scope Options** dialog box, select the options you want to configure.
4. In the **Data Entry** section of the **Scope Options** dialog box, type the option parameters, and then click **OK**.

**To configure options for a reserved client**

1. In the DHCP snap-in, expand the scope that holds the reservation for which you want to configure options, and then expand **Reservations**.
2. Right-click the reservation for which you want to configure options, and then click **Configure Options**.
3. In the **Reservation Options** dialog box, select the options you want to configure.
4. In the **Data Entry** section of the **Reservation Options** dialog box, type the option parameters, and then click **OK**.

For more information about configuring reservations, see "Creating Reservations" later in this chapter.

Many option types are predefined in Windows Server 2003 DHCP. Other standard DHCP option types can be added as needed to support DHCP client software that recognizes or requires them. Windows Server 2003 DHCP supports all DHCP options, including those defined in RFC 2132, although most DHCP clients use or support only a small subset of the available option types.

In general, use the following guidelines when configuring DHCP options for clients on your network:

- Add or define new, custom option types only if you have new software or applications that require a nonstandard DHCP option.
- If your network is large, be conservative and selective when assigning global options. These options apply to all clients of a DHCP server, unless more specific options are specified.
- Use scope-level options for most options that clients are assigned. Setting options at the scope level allows you to take scope-related differences into account, such as different client needs or the use of a different DNS server from other scopes in the network.
- Use class-level options if you have a large network or diverse groups of clients that are able to support membership in option classes.
- Use reserved client options only for clients that have special requirements, for example, if your intranet has a DNS server that performs forwarding for resolving Internet DNS names not authoritatively managed on your network. In this case, you need to add the IP address of an external DNS server on your DNS server computer. You can configure your DNS server as a reserved client in DHCP and set this address as another reserved client option.

# Configuring Option Classes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 DHCP includes vendor-defined and user-defined option classes. Use DHCP option classes to configure the parameters necessary for network clients to meet the special requirements of custom applications. Equipment from multiple vendors on a network can also use different option numbers for different functions. The option types used to support vendor-defined classes — the vendor class identifier and the vendor-specific option — are defined in the Internet DHCP options standard reference, RFC 2132.

You can add and configure vendor-defined classes for submanaging DHCP options that are assigned to clients identified by vendor type. You can also add and configure user-defined classes for submanaging DHCP options that are assigned to clients identified by a common need for a similar DHCP option configuration.

After you configure specific user-defined and vendor-defined options classes, you must configure scopes to assign the option classes to clients.

**Vendor-Defined Classes**
DHCP clients can use vendor-defined classes to identify the client's vendor type and configuration to the DHCP server when obtaining a lease. For a client to identify its vendor class during the lease process, the client needs to include the vendor class ID option (option code 60) when it requests or selects a lease from a DHCP server.

When vendor options are specified, the server performs the following additional steps to provide a lease to the client:

1. The server checks to see that the vendor class identified by the client request is a recognized class defined on the server.

2. If the vendor class is recognized, the server checks to see whether any additional DHCP options are configured for this class in the active scope.

3. If the vendor class is not recognized, the server ignores the vendor class identified in the client request, and returns options allocated to the default vendor class (includes all DHCP Standard Options).

4. If the scope contains options configured specifically for use with clients in this vendor-defined class, the server returns those options and uses the vendor-specific option type (option code 43) as part of its acknowledgment message.

In most cases, the default vendor class — DHCP Standard Options — provides a default grouping for any Windows Server 2003 DHCP clients or other DHCP clients that do not specify a vendor class ID. In some cases, you might define additional vendor classes for other DHCP clients, such as printers or some types of UNIX clients. When you add other vendor classes for these purposes, be sure that the vendor class identifier that you use to configure the class at the server matches the identifier used by clients for your third-party vendor.

**User-Defined Classes**
User-defined classes allow DHCP clients to specify what type of client they are, such as a remote access client or a desktop computer. For Windows Server 2003 clients, you can define specific user class identifiers to relate information about a client's software configuration, its physical location in a building, or its user preferences. If user-defined option classes are not specified, default settings are assigned.

# Configuring DHCP Option Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a DHCP server actively provides option parameters, clients receive and use the associated values in their local TCP/IP configurations for the period of leased configuration.

By default, Microsoft-based DHCP clients can recognize and use two categories of option parameters: information options and protocol options.

**Information Options**
Use the MMC DHCP snap-in to *explicitly* configure information options and any associated values provided to clients. These options are not required and can be assigned at your discretion. Use information options to assign values, such as DNS servers, WINS servers, and domain name.

**Windows Server 2003 Protocol Options**
You *implicitly* configure values for protocol options based on properties configured at either the applicable server or one of its scopes. These options are always included in DHCP client/server messages, as they are a required part of protocol design. For example, Lease Duration/Time is a protocol option.

In most cases, the actual values included in these protocol options are provided to clients based on property settings for the applicable DHCP server. Depending on the needs of your clients, you can also use the DHCP snap-in to configure these options individually for defined scopes, identifying members of a specified user or vendor class, or for a single reserved client.

# Creating Reservations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For clients that require a constant IP address, you can either manually configure a static IP address, or assign a reservation on the DHCP server. *Reservations* are permanent lease assignments that are used to ensure that a specified client on a subnet can always use the same IP address. You can use DHCP reservations for hosts that require a consistent IP address, but do not need to be statically configured. Reserved IP addresses differ from statically configured IP addresses in one significant manner: when network parameters are changed at the DHCP server, the device configured with a reserved IP address receives the new network parameters when the device requests renewal of its lease. To change network parameters on a device configured with a static IP address, the changes must be made manually to the device.

Determine the clients for which you need to manually configure an IP address — such as DHCP server, DNS servers, WINS servers, routers, and domain controllers — and which clients can receive addresses from DHCP. Keep in mind that for clients for which you manually configure static IP addresses, you must insert all configuration parameters that the client requires in order to interact with the network. This includes IP addresses, DNS and WINS parameters, and default gateway information. Clients that have reserved IP addresses always have the same IP address, but still receive updated configuration information from the DHCP server.

You might want to assign network printers and certain servers DHCP reservations to ensure that they always have the same IP address, but continue to receive updated configuration information from the DHCP server. For example, create reservations for servers that must always have the same IP address, such as:

- Windows Internet Name Service (WINS) and Domain Name System (DNS) servers
- Print servers that use TCP/IP print services
- Firewalls
- Routers

DHCP-enabled clients receive any available options, such as DNS server or router (default gateway), from the DHCP server when they renew their leases. If these devices are manually configured, an administrator must reconfigure each device individually when a change occurs.

**To create a reservation**

1. In the DHCP snap-in, expand the scope for which you want to create a reservation.
2. Select and right-click **Reservations**, and then click **New Reservation**.
3. In the **NewReservation** dialog box, enter the **Reservation name**, **IP address**, **MAC address**, and **Description** of the reservation.
4. Select the appropriate **Supported types**: DHCP only, BOOTP only, or Both.
5. Click **Add**.

# Creating Superscopes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A *superscope* is an administrative grouping of scopes that can support multiple logical IP subnets on the same physical subnet. Superscopes contain a list of member scopes that can be activated together. You cannot configure scope-level properties on superscopes; you must configure these on the member scopes. A superscope allows a DHCP server to provide leases from more than one scope to clients on a single physical network.

You can use superscopes to resolve DHCP service issues for the following situations:

- DHCP clients are located on a single physical network segment that includes multiple logical IP subnets.
- Multiple DHCP servers manage separate logical subnets on the same physical subnet.
- The available address pool for an active scope is nearly depleted and more computers must be added to the physical network segment.
- Clients are migrating to a new scope.
- You need to support DHCP clients on a network that has multiple logical subnets in one physical subnet on the other side of a BOOTP/DHCP relay agent.

Before you create a superscope, you must use the DHCP MMC snap-in to define at least one scope to be included in the superscope. Scopes added to a superscope are called *member scopes*. You can add additional member scopes either from the superscope menu, or from the individual scope menus.

**To create a new superscope**

1. In the DHCP snap-in, create at least one scope to be included in the superscope. For information about creating scopes, see "Creating Scopes" earlier in this chapter.
2. Select and right-click the DHCP server, and then select **New Superscope**. This opens the **New Superscope Wizard**.
3. On the **Superscope Name** page of the **New Superscope Wizard**, type a name for the superscope.
4. On the **Select Scopes** page of the **New Superscope Wizard**, inthe **Available Scopes** list, select one or more scopes to include in the superscope.

**To add scopes to an existing superscope**

1. In the DHCP snap-in, select and right-click the superscope, and then select **New Scope**. This opens the **New Scope Wizard**.
2. Complete the **New Scope Wizard**. For information about creating scopes, see "Creating Scopes" earlier in this chapter.

   –or–

3. Create a new scope, right-click the new scope, and then select **Add to Superscope**.
4. In the **Add Scope** *name* **to a Superscope** dialog box, in the **Available superscopes** list, select the superscope.

**Note**

- Deleting the superscope does not delete the member scopes.

# Configuring Multicast Scopes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 DHCP service offers MADCAP support in the form of multicast scopes. MADCAP supports dynamic assignment and configuration of IP multicast addresses on TCP/IP-based networks. Multicast scopes provide ranges of Class D IP addresses, which are reserved for multicast operation, by using directed transmission from one point to multiple points. With the exception of DHCP-assignable options, which multicast scopes do not support, you can configure a multicast scope in the same way that you configure a regular DHCP scope.

Multicast IP addresses allow multiple clients to receive data that a server sends to a single IP address, enabling point-to-multipoint communication. This type of transmission is often used for streaming media transmissions, such as video conferencing.

**Important**

- In all TCP/IP networks, each computer requires a unique primary computer IP address from one of the standard address classes used for building the network (Class A, B, or C range). You must assign this required primary computer IP address before you can configure a computer to support and use secondary IP addresses such as multicast IP addresses.

Although the Windows Server 2003 DHCP service supports both DHCP and MADCAP, the services function independently; clients that do not obtain IP addresses from the DHCP service can still obtain MADCAP addresses from the DHCP service. DHCP scopes are used to allocate IP address ranges from Class A, B, or C addressing schemes, which enable unicast for point-to-point communication between networked computers. MADCAP scopes allocate Class D IP addresses to enable point-to-multipoint communication.

**To configure a MADCAP scope**

1. In the DHCP snap-in, select and right-click the DHCP server you want to configure. The **New Multicast Scope Wizard** appears.

2. In the **New Multicast Scope Wizard**, type a name and description for this multicast scope.

3. Set the multicast IP address range and Time to Live (TTL).

4. Add any exclusion ranges and the lease duration, then activate the multicast scope.

**Tip**

- Clients that use MADCAP must be configured to use the MADCAP API. For more information about writing or programming applications that use the MADCAP API, see the MSDN Online link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Removing DHCP Scopes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remove DHCP scopes when a subnet is no longer in use or when you need to renumber your network to use a different IP address range.

Do not remove a scope while it has active leases. Before you remove a scope, deactivate the scope until all client leases expire or all lease renewal requests are denied. When you have confirmed that the scope no longer contains active leases, you can remove it by using the DHCP snap-in. For more information about deactivating scopes, see "Deactivate a scope"Deactivate a scope in Help and Support Center for Windows Server 2003.

# Implementing Your DHCP Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have planned your DHCP solution and made the necessary configurations, you must implement your new or updated DHCP solution in your production environment. The DHCP implementation process involves the following steps:

- Configuring your DHCP clients.
- Migrating existing DHCP servers, if necessary.
- Testing your DHCP solution.

Figure 2.8 shows the process for implementing your DHCP solution.

**Figure 2.8 Implementing Your DHCP Solution**



# Installing DHCP on Your Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Install DHCP on your server either by using Manage Your Server or by adding DHCP from Add/Remove Windows Components. If you are migrating an existing DHCP server to a new Windows Server 2003–based server, install DHCP on the Windows Server 2003–based server before continuing with your migration. For more information about installing DHCP on your server, see "DHCP server role: Configuring a DHCP server," "Installing a DHCP serverInstalling a DHCP server," and "Checklist: Installing a DHCP server" in Help and Support Center for Windows Server 2003.

# Migrating Existing DHCP Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you already have a working DHCP infrastructure on your network, you can migrate your DHCP database and server configuration from your existing servers to new Windows Server 2003–based servers. Migrating can save you time because you do not need to reconfigure the scope, lease, and option information on the new server.

Before migrating the database, create a backup of your working configuration, and test the migration procedures in a lab environment.

The following procedures migrate lease, scope, and option information, including reservations and exclusions. If you have changed any registry settings from their defaults on the existing server, you must manually make these changes on the new Windows Server 2003–based server for them to take effect.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Referenceon the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

# Exporting DHCP Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When migrating from an existing DHCP server running Windows NT 4.0 or Windows 2000 Server, the first step is to export the current DHCP settings. Use the following procedure to export the DHCP database from an existing Windows NT 4.0–based server or a Windows 2000–based server.

**To export the DHCP server settings from a Windows NT 4.0–based server or Windows 2000–based server**

1. Run DHCPExim.exe. This opens the **DhcpExim** dialog box.

   To download and install DHCPExim.exe, see the Resource Kit Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

2. In the **DhcpExim** dialog box, select **Export configuration of the local service to a file**.

3. In the **DHCPEXIM Export To File** dialog box, enter a file name and location to save the file, and then click **OK**.

   The target location the file is shown in the **DhcpExim Export** dialog box.

4. In the **DhcpExim Export** dialog box, select all of the scopes on the list to migrate all settings on the server. Alternatively, you can only the specific scopes you want to export if you are not migrating the entire server. Select the **Disable the selected scopes on local machine before export** check box to disable the scopes on this server before exporting, and then click **Export**.

   **Tip**

   - This step might take several minutes to complete, and there is no dialog box to indicate the progress of the export. You can view the process running on the **Processes** tab of the **Windows Task Manager**.

5. Click **OK** when a message appears that says "The operation was completed successfully."

6. Copy the exported file to a location where you can access it from the new Windows Server 2003–based server.

# Importing DHCP Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After exporting the DHCP settings and moving the exported file to a location where you can access it, you are ready to use the Netsh command-line tool to import the DHCP settings to the new DHCP server running Windows Server 2003. You might want to configure the new DHCP server to use the IP address that was formerly used on the old DHCP server. Before you do this, make sure to disconnect the old server from the network in order to avoid any IP address conflicts.

**To import the DHCP server settings to Windows Server 2003–based DHCP server**

1. Install DHCP on the Windows Server 2003–based server. DHCP must be installed before you can import the file.

2. At the command prompt, type:

   ```
   netsh DHCP server import <path of export file> all
   ```

   This imports the configurations of all scopes that you exported from the server running Windows NT 4.0 or Windows 2000.

For more information about the Netsh command-line tool, see "Netsh" in Help and Support Center for Windows Server 2003.

# Configuring DHCP Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When configuring your DHCP clients, it is important to consider the types of computers on each subnet of the network. For example, you might configure one subnet with a shorter lease duration for portable computer users. To configure a Windows client to obtain an IP address from the DHCP server, select **Obtain an IP address automatically** from the **Internet Protocol (TCP/IP) Properties** window.

In addition to Windows-based clients, you can configure DHCP for UNIX, Linux, or Macintosh clients. Refer to the documentation for these operating systems to configure the clients for DHCP.

**Note**

- You can override DHCP information by configuring the individual client. Any information entered manually into a client's TCP/IP configuration overrides dynamic settings.

# Providing Support for Remote Access Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The remote access server can be configured to use DHCP to lease IP addresses in blocks of 10 from the DHCP server and store them in the registry. The remote access server leases additional addresses in blocks of 10 as needed.

When passing the leases on to remote access clients, the remote access server discards the options configured for the scope. However, remote access clients receive WINS or DNS entries from the remote access server. This information is not taken from the DHCP options of the lease that the remote access client is using; it is taken directly from the settings of the remote access server.

You can use a DHCP relay agent to provide DHCP scope options to remote access clients. In this situation, the remote access client continues to receive an IP address from the remote access server, but uses DHCPInform packets to obtain WINS addresses, DNS addresses, domain names, or other DHCP options.

**Note**

- DNS and WINS addresses obtained from DHCPInform packets override DNS and WINS addresses obtained from the remote access server.

For more information about using DHCP for remote access clients, see "Using Routing and Remote Access servers with DHCP" in Help and Support Center for Windows Server 2003 and article Q216805, "RAS Server Behavior When Configured to Use DHCP to Assign IP Addresses," in the Microsoft Knowledge Base. For more information about configuring a DHCP relay agent to provide DHCP options to remote access clients, see article Q232703, "How to Use DHCP to Provide RAS Clients with DHCP Options," in the Microsoft Knowledge Base. To find these articles, see the Microsoft Knowledge Base link on the Web Resources page at www.microsoft.com/windows/reskits/webresources.

# Configuring Support for BOOTP Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 provides support for dynamic BOOTP, which allows you to provide BOOTP support without making a reservation for each BOOTP client.

You can use DHCP to configure and support BOOTP clients. BOOTP clients can be configured to receive boot file information, or IP address information. You can also use DHCP to configure options for your BOOTP clients.

**Configuring DHCP to Provide Boot File Information**
To configure the DHCP service to provide boot file information to BOOTP clients, create BOOTP entries for each client-specific platform in the BOOTP table on the DHCP server.

Information stored in the BOOTP table is returned to any BOOTP clients on the network that broadcast a BOOTP request message. If the BOOTP table includes at least one BOOTP entry, the DHCP service replies to BOOTP client requests. If no BOOTP entries are configured, the BOOTP client gets a lease, but no options or other information are passed to the BOOTP client.

The reply message returned by the DHCP service indicates the name of a Trivial File Transfer Protocol (TFTP) server on the network and the location of the boot file. The client then contacts the TFTP server to retrieve the boot image file.

BOOTP clients download the image file from a TFTP server. Because Windows Server 2003 does not provide a full TFTP file service, you might need to use a third-party TFTP server to support BOOTP clients that must start from an image file (usually diskless workstations).

**Configuring DHCP to Provide IP Address Information**
Windows Server 2003 allows you to designate a pool of addresses from which IP addresses for BOOTP clients are dynamically assigned, similar to the way that a scope is used for DHCP clients. After verifying that a specified lease time has elapsed and that the BOOTP client IP address is not still in use, the DHCP service can reclaim addresses used in the dynamic BOOTP address pool.

To include BOOTP client support in your DHCP scope design, designate the DHCP scope as a BOOTP address pool or as both a DHCP and BOOTP address pool on the **Advanced** tab of the **Scope Properties** window in the DHCP snap-in. For more information about configuring a BOOTP address pool, see "Add dynamic BOOTP client support to a scope" in Help and Support Center for Windows Server 2003.

**Configuring Options for BOOTP Clients**
To retrieve all options, the client must specify option 55 (the Options Request List parameter) in the BOOTP request. Windows Server 2003 DHCP servers return as many options to BOOTP clients as can fit in a single datagram response. Because BOOTP only allocates 300 bytes for options, any options that exceed the 300 bytes are not sent. Therefore, when configuring options on a server that will service BOOTP clients, be aware of the size limitations for BOOTP response packets.

For more information about BOOTP options, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

### Important

- DHCP options can apply to both DHCP and BOOTP clients. Therefore, you must configure your scopes to ensure that DHCP options are applied correctly.

# Testing Your DHCP Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The ideal environment for validating a deployment plan is a test laboratory that is constructed to simulate your production environment. If you do not have a DHCP infrastructure in place, create a new DHCP infrastructure in the pilot lab. If you already have a DHCP infrastructure in place, recreate your existing infrastructure in the pilot lab and then test your migration and your updated solution.

Use the following guidelines when testing your DHCP solution in the pilot lab:

- Ensure that your test servers come from the same vendor and are configured in the same way as the servers that exist in your production environment. Represent all client computers that use DHCP in your organization in your pilot lab.

- Duplicate the network configuration of your organization in the pilot lab. If the network uses both Ethernet and Token Ring, you need to include both configurations in the test lab. If possible, set up a separate Windows Server 2003 domain for the lab so that you can monitor the performance of the domain controllers without interference from other network activity. If you are planning to deploy DHCP on a WAN, include routers in your lab design and use a link simulator to simulate network latency.

- Deploy a typical set of applications that run together in your production environment on the test server. This allows you to identify any interoperability issues that might arise when users run different applications simultaneously.

- Test your DHCP solution in burst traffic conditions. Make sure that all services on your servers continue to be available during burst traffic conditions, such as following a large-scale power outage, when many clients might restart simultaneously.

For more information about testing your design, see "Designing a Test Environment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

Testing your DHCP solution in a pilot lab allows you to identify problems with your infrastructure, system configuration, or software. After you have tested your configuration thoroughly in a pilot lab, you can roll out your DHCP solution in your production environment.

# Example DHCP Implementation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This DHCP example implementation depicts how a fictional company uses the Windows Server 2003–based DHCP servers to streamline and automate administration and assignment of IP addresses and other client configuration information. Though your network configuration might differ from this example, you can apply the basic concepts.

The company has two offices, a main administrative office, and branch offices 25 miles away. The company employs 600 people, many of whom travel frequently between the two offices with portable computers.

The main office is housed on two floors, with 75 desktop computer users per floor. Because many of these employees also have portable computers, conference rooms are configured with wireless networking. Wireless networking also provides ease of network connectivity to employees from the branch offices who are attending meetings at the main office.

**Connectivity and Routing**

A digital subscriber line (DSL) account at a rate of 1.1 megabits per second (Mbps) is established for each of the two company sites, providing a dedicated connection to the Internet through a DSL modem. Additional network adapters are installed on the server computers at both sites, one network adapter per subnet.

The DSL-connected server in the main office has two additional network adapters, one for Subnet A and one for Subnet B. On Subnet A, network cable runs from the network adapter to a hub. On wireless Subnet B, media network cable runs from the network adapter to a hub, then to wireless access points stationed in conference rooms.

The DSL-connected server computer in the branch office has four additional network adapters, one each for subnets C through F. On wireless subnet C, network cable runs from the network adapter to a hub, then to wireless access points stationed in conference rooms. On subnets D, E, and F, network cable runs from the network adapters to hubs that extend to each individual subnet.

To provide secure, private, encrypted communication between the two sites over DSL, both DSL-connected servers are also configured as virtual private network (VPN) servers.

To allow messages from both DHCP servers and DHCP clients to cross from site to site and between subnets at each site, the Routing and Remote Access service on the DSL-connected servers have the DHCP relay agent routing protocol installed and configured.

For more information about connecting remote sites, see "Connecting Remote Sites" in this book.

**Transmission Security Between Sites**

Because the company wants the highest level of security available for the VPN connection that allows the two sites to communicate using their 1.1 Mbps DSL connections over the Internet, they install a Windows Server 2003 certification authority (CA) and establish a Layer Two Tunneling Protocol (L2TP) connection, relying on Internet Protocol security (IPSec) for encryption services. The combination of L2TP and IPSec is known as L2TP/IPSec. With certificate-based authentication and an L2TP/IPSec connection, the company is using the strongest form of authentication in Windows Server 2003. For more information about using a certification authority, see "Designing an Authentication Strategy" in the *Designing and Deploying Directory and Security Services* book of this kit. For more information about using IPSec, see "Deploying IPSec" in this book.

**Active Directory Domain Structure**

The company uses Active Directory to list IP addresses of DHCP servers authorized for operation on the network. If an unauthorized DHCP server running Windows 2000 or Windows Server 2003 is started on the network, it determines its authorization status from the directory service. If the server determines that it is not authorized, it stops functioning in the domain as a DHCP server, and cannot be used to provide IP address leases to clients. There are two sites to define in Active Directory, one at the main administrative site and one at the branch office site. Creating two sites optimizes the exchange of directory information and facilitates administration by centralizing resources such as configuration information.

# Subnets and IP Addressing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although the company employs approximately 600 people, most employees use both a desktop computer and a portable computer. Because the portable computers function primarily via wireless connections to the network, one subnet at each of the two company sites is configured as a wireless subnet.

In all, there are six subnets on the company network, with 600 desktop clients and approximately 425 portable clients divided among the subnets as shown in Table 2.1. All subnets on the network are configured with the subnet mask 255.255.255.0. All six subnets are serviced by DHCP, with one DHCP server located at each site.

The main office houses 150 employees dispersed over two floors of the building. All users also have portable computers that they frequently take to meetings in conference rooms. All of conference rooms are configured for wireless networking.

The company has divided the main office into two Class C subnets:

- Subnet A — Floors 1 and 2
- Subnet B — Wireless subnet with wireless access points installed in conference rooms

The branch office houses 450 employees, dispersed throughout a three-story building. Over 200 users also have portable computers that they frequently take to meetings in conference rooms, all of which are configured for wireless networking.

The company has divided the branch office into four Class C subnets:

- Subnet C — Wireless subnet for conference rooms
- Subnet D — Floor 3
- Subnet E — Floor 2
- Subnet F — Floor 1

**Table 2.1 Address Range and Clients per Subnet**

| Subnet | Address Range | Approximate Number of Clients |
|---|---|---|
| A | 192.168.0.1/24 to 192.168.0.254/24 | 150 |
| B | 192.168.1.1/24 to 192.168.1.254/24 | Up to 225 |
| C | 192.168.2.1/24 to 192.168.2.254/24 | Up to 225 |
| D | 192.168.3.1/24 to 192.168.3.254/24 | 150 |
| E | 192.168.4.1/24 to 192.168.4.254/24 | 150 |
| F | 192.168.5.1/24 to 192.168.5.254/24 | 150 |

Before creating and activating scopes on the DHCP server, the company plans IP address ranges, exclusion ranges, and reservations (where applicable) for each subnet. The company uses Class C IP address ranges on every subnet, as each Class C address range provides 254 IP addresses when the subnet mask is defined as 255.255.255.0.

Because the client count per non-wireless subnet is 150 or fewer clients, 254 IP addresses per subnet allow the company to provide static address assignments for any servers that require them, and dynamic assignments to all clients on the network — with plenty of IP addresses remaining to provide for future network expansion.

The wireless subnets are configured with a lease duration of 15 minutes. Because of the short lease time, 254 IP addresses per subnet is a sufficient amount for the wireless subnets, even though these subnets experience substantial traffic, with portable computers joining and leaving the subnet at high volume during a typical day. These short leases expire soon after the portable computer is disconnected from the network, and the IP address used by that computer becomes available for lease to other DHCP clients as they connect to the wireless subnet.

# Exclusion Ranges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some network devices need to use statically assigned IP addresses rather than addresses dynamically assigned through DHCP. For example, DHCP servers must have statically configured IP addresses. Also, some devices (such as legacy network printers) do not support DHCP.

For the devices that need static IP assignments, the company creates an exclusion range from each IP address range. Creating one or more exclusion ranges prevents the DHCP server from assigning a client lease with any address in the exclusion range, thereby protecting it for use as a static IP address and preventing address conflicts between statically configured devices and dynamically configured devices.

Although any addresses in the address range can be excluded, the company chooses to exclude the first 20 addresses from each address range for non-wireless subnets, and the first five IP addresses from each address range for wireless subnets. The company uses additional exclusion ranges to configure load balancing and fault tolerance using the 80/20 rule. For more information about DHCP scopes, see "Scope Configuration" later in this chapter.

# Address Pools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After the address range and exclusion ranges are defined, the remaining addresses form the available address pool within the scope. These addresses are eligible for dynamic assignment by the server to DHCP clients on the network. Table 2.2 shows the address pool for each subnet prior to adding the exclusion ranges used to apply the 80/20 rule.

**Table 2.2 DHCP Address Pools per Subnet**

| Subnet | Address Pool |
|---|---|
| A | 192.168.0.21/24 to 192.168.0.254/24 |
| B (wireless) | 192.168.1.6/24 to 192.168.1.254/24 |
| C (wireless) | 192.168.2.6/24 to 192.168.2.254/24 |
| D | 192.168.3.21/24 to 192.168.3.254/24 |
| E | 192.168.4.21/24 to 192.168.4.254/24 |
| F | 192.168.5.21/24 to 192.168.5.254/24 |

For more information about DHCP scopes and the 80/20 rule, see "Scope Configuration" later in this chapter.

# Reservations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The company uses IP address reservations for file and print servers on their network. Reservations are used to create a permanent IP address lease assignment by the DHCP server. Reservations ensure that a specified hardware device on the subnet can always use the same IP address.

When using the 80/20 rule and splitting a scope's IP address pool between two servers for load balancing and fault tolerance, identical reservations must be made at both DHCP servers. When reservations are made at both servers, neither server assigns the reserved IP address to another client, assuring that the intended device receives the address reserved for its use.

Table 2.3 shows two example address reservations. For more information about the 80/20 rule and for an example of how these reservations are created in specific scopes at each DHCP server, see "Scope Configuration" later in this chapter.

**Table 2.3 Example Address Reservations**

| Device | Subnet | Reserved IP Address |
|---|---|---|
| Application server | A | 192.168.0.21/24 |
| File server | D | 192.168.3.68/24 |

**Note**

- Reservations can be created using any IP address in the scope's address range, even if the IP address is also within an exclusion range. Because of this design, when the 80/20 rule is implemented and some addresses in the scope are excluded (80 percent at one server, 20 percent at the other), reservations still function properly.

# Message Routing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The company uses DHCP relay agents to relay DHCP messages between subnets and sites. To support and use DHCP across multiple subnets, routers connecting each subnet should comply with DHCP/BOOTP relay agent capabilities described in RFC 1542. To cut the cost of expensive network hardware such as routers, the company uses the Windows Server 2003 Routing and Remote Access service including, DHCP relay agents, to forward DHCP/BOOTP messages between subnets. Because the VPN servers act as routers for network traffic, the DHCP relay agents are configured on the VPN servers.

For more information about installing and configuring DHCP relay agents, see "Configure the DHCP Relay Agent" in Help and Support Center for Windows Server 2003.

# Installing DHCP and Authorizing Servers in Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The company completes the following steps at both sites:

1. Configures the servers.
2. Installs the DHCP Server service.
3. Opens the DHCP snap-in from the MMC.
4. Authorizes the DHCP server in Active Directory.
5. Adds DHCP administrators to the DHCP Administrators group or the DHCP Users group, depending on their user rights.
6. Creates, configures, and activates one scope for each subnet.

After creating scopes at the DHCP servers, the company defines a site in Active Directory for the branch office. Each of the six subnets is then associated with the appropriate site. When configuring your network, however, you can configure all sites in Active Directory before creating DHCP scopes.

# Scope Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using the 80/20 split-scope configuration for fault tolerance and availability, scopes for all six subnets on the company network are defined on both DHCP servers. Exclusion ranges are used to allocate available addresses per scope, per server, as follows:

- The main office DHCP server is configured with 80 percent of the IP addresses available for lease to clients in each scope serving subnets A and B and 20 percent of the IP addresses available for lease to clients located at the branch office (subnets C through F).

- The branch office DHCP server has 80 percent of all addresses in all scopes available for lease to clients in the branch office (subnets C through F) and 20 percent of all addresses in all scopes available for lease to clients located in the main office (subnets A and B).

Thus, if either server suffers a hard-disk failure or other failure, the alternate server is available to assign and renew leases on all subnets.

To achieve the 80/20 rule, each Class C IP address range of 254 IP addresses available in each non-wireless scope is divided in the following manner:

- 20 IP addresses for static assignments.

- 187 IP addresses, or 80 percent of the addresses for lease, in the address pool of the DHCP server at the same site.

- 47 IP addresses, or 20 percent of the addresses for lease, in the address pool of the DHCP server at the other site.

The wireless scopes are divided in the following manner:

- 5 IP addresses for static assignments.

- 203 IP addresses, or 80 percent of the addresses for lease, in the address pool of the DHCP server at the same site.

- 51 IP addresses, or 20 percent of the addresses for lease, in the address pool of the DHCP server at the other site.

Table 2.4 shows the address pools and exclusion ranges configured on the main office DHCP server.

**Table 2.4 Scope Configurations on the DHCP Server at the Main Office**

| Scope Name | Address Range | Exclusion Ranges | Address Pool |
|---|---|---|---|
| A | 192.168.0.1 to 192.168.0.254 | 192.168.0.1 to 192.168.0.20, 192.168.0.21 to 192.168.0.67 | 192.168.0.68 to 192.168.0.254 |
| B (wireless) | 192.168.1.1 to 192.168.1.254 | 192.168.1.1 to 192.168.1.5, 192.168.1.204 to 192.168.1.254 | 192.168.1.6 to 192.168.1.203 |
| C (wireless) | 192.168.2.1 to 192.168.2.254 | 192.168.2.1 to 192.168.2.203 | 192.168.2.204 to 192.168.2.254 |
| D | 192.168.3.1 to 192.168.3.254 | 192.168.3.1 to 192.168.3.20, 192.168.3.68 to 192.168.3.254 | 192.168.3.21 to 192.168.3.67 |

| | | | |
|---|---|---|---|
| E | 192.168.4.1 to 192.168.4.254 | 192.168.4.1 to 192.168.4.20, 192.168.4.68 to 192.168.4.254 | 192.168.4.21 to 192.168.4.67 |
| F | 192.168.5.1 to 192.168.5.254 | 192.168.5.1 to 192.168.5.20, 192.168.5.68 to 192.168.5.254 | 192.168.5.21 to 192.168.5.67 |

Table 2.5 shows the address pools and exclusion ranges configured on the branch office DHCP server.

**Table 2.5 Scope Configurations on the DHCP Server at the Branch Office**

| Scope Name | Address Range | Exclusion Ranges | Address Pool |
|---|---|---|---|
| A | 192.168.0.1 to 192.168.0.254 | 192.168.0.1 to 192.168.0.20, 192.168.0.68 to 192.168.0.254 | 192.168.0.21 to 192.168.0.67 |
| B (wireless) | 192.168.1.1 to 192.168.1.254 | 192.168.1.1 to 192.168.1.203 | 192.168.1.204 to 192.168.1.254 |
| C (wireless) | 192.168.2.1 to 192.168.2.254 | 192.168.2.1 to 192.168.2.5, 192.168.2.204 to 192.168.2.254 | 192.168.2.6 to 192.168.2.203 |
| D | 192.168.3.1 to 192.168.3.254 | 192.168.3.1 to 192.168.3.20, 192.168.3.21 to 192.168.3.67 | 192.168.3.68 to 192.168.3.254 |
| E | 192.168.4.1 to 192.168.4.254 | 192.168.4.1 to 192.168.4.20, 192.168.4.21 to 192.168.4.67 | 192.168.4.68 to 192.168.4.254 |
| F | 192.168.5.1 to 192.168.5.254 | 192.168.5.1 to 192.168.5.20, 192.168.5.21 to 192.168.5.67 | 192.168.5.68 to 192.168.5.254 |

**Lease Duration**

Subnet B in the main office and Subnet C in the branch office are both wireless subnets. Because wireless clients (portable computers and other portable devices) are connected to and disconnected from the network in large numbers and for short intervals throughout the average day, lease duration on these two subnets is set for 15 minutes. These short lease times help to ensure that the maximum number of IP addresses are available in the scope as clients connect to the network.

Lease time for all other (nonwireless) subnets is eight days.

**Scope Options**

Each scope is configured with option 249, classless static routes. Using classless static routes, each DHCP client can be easily configured with the route to any destination on the network, and the subnet mask can be specified. Because each scope represents a physical subnet, the scope can be viewed as the start location for any message that is to be sent by a client to another subnet. The parameters used to configure option 249 are Destination, Mask, and Router. One or more static routes can be configured with option 249; the company provides all DHCP enabled clients on the network with routes to all other subnets using option 249.

This option is not configured as a server option because it maps routes between subnets, so no one set of values for the required parameters of Destination, Mask, and Router is always correct. For example, subnets A and D each use a router (that is, a VPN server configured with the Routing and Remote Access service and DHCP Relay Agent service enabled) to communicate with each other. Of course, the routers they use are different, and the destination is different in each case.

**Server Options**

The DHCP standard options shown in Table 2.6 are configured as server options at the main office DHCP server.

**Table 2.6 Example DHCP Options**

| Option Number | Description | Value |
|---|---|---|

| | | |
|---|---|---|
| 006 | DNS servers | 192.168.0.3, 192.168.3.3 |
| 044 | WINS/NBNS servers | 192.168.0.4, 192.168.3.4 |
| 132 | Enable NBT hostname resolution | Byte: 0x1 (1=on) |
| 133 | Enable gethostbyname() WINS resolution | Byte: 0x1 (1=on) |

**Other Settings**

Scopes on both DHCP servers are configured with the same reservations, lease durations, scope options, and server options.

# Additional Resources for Deploying DHCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about Dynamic Host Configuration Protocol (DHCP).
- "Dynamic Host Configuration Protocol" in the *TCP/IP Core Networking Guide* of the *Microsoft® Windows® 2000 Server Resource Kit*.
- "Deploying Domain Name System (DNS)" in this book.
- "Deploying IPSec" in this book.
- "Designing an Authentication Strategy" in the *Designing and Deploying Directory and Security Services* book of this kit for more information about using a certification authority.
- *DHCP for Windows 2000*, by Neall Alcott, 2001, Sebastopol, CA: O'Reilly & Associates.
- RFC 1542: *Clarifications and Extensions for the Bootstrap Protocol*.
- RFC 2131: *Dynamic Host Configuration Protocol*.
- RFC 2132: *DHCP Options and BOOTP Vendor Extensions*.

**Related Tools**

- DHCPExim tool

  For more information about DHCPExim, see the Resource Kit Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Netsh

  For more information about Netsh, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Command-line reference A-Z.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Checklist: Installing a DHCP server" in Help and Support Center for Windows Server 2003.
- "Checklist: Installing a DHCP Service resource" in Help and Support Center for Windows Server 2003.
- "Checklist: Installing a MADCAP server" in Help and Support Center for Windows Server 2003.
- "Cluster support for DHCP servers" in Help and Support Center for Windows Server 2003.
- "Configure the DHCP Relay Agent" in Help and Support Center for Windows Server 2003.

- "Netsh" and "Netsh commands for DHCP" in Help and Support Center for Windows Server 2003.
- "Add dynamic BOOTP client support to a scope" in Help and Support Center for Windows Server 2003.
- "Using Routing and Remote Access servers with DHCP" in Help and Support Center for Windows Server 2003.

# Deploying Domain Name System (DNS)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft® Windows® Server 2003 Domain Name System (DNS) provides efficient name resolution and interoperability with standards-based technologies. Deploying DNS in your client/server infrastructure enables resources on a TCP/IP network to locate other resources on the network by using host name-to-IP address resolution and IP address-to-host name resolution. The Active Directory® directory service requires DNS for locating network resources.

**In This Chapter**


Overview of DNS Deployment

Examining Your Current Environment

Designing a DNS Namespace

Designing a DNS Server Infrastructure

Designing DNS Zones

Configuring and Managing DNS Clients

Securing Your DNS Infrastructure

Integrating DNS with Other Windows Server 2003 Services

Implementing Windows Server 2003 DNS

Additional Resources for Deploying DNS


**Related Information**

- For more information about DNS, the Windows Server 2003 DNS Server service, and Windows Server 2003 DNS Client service, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Overview of DNS Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DNS is the primary method for name resolution in the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems (collectively referred to as "Windows Server 2003" in this chapter). DNS is also a requirement for deploying Active Directory, but Active Directory is not a requirement for deploying DNS. However, integrating DNS with Active Directory enables DNS servers to take advantage of the security, performance, and fault tolerance capabilities of Active Directory.

If you are planning to deploy DNS to support Active Directory, plan your DNS namespace in conjunction with planning your Active Directory logical structure. For more information about designing the Active Directory logical structure, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

# Process for Deploying DNS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying DNS involves planning and designing your DNS infrastructure, including the DNS namespace, DNS server placement, DNS zones, and DNS client configuration. In addition, if you are integrating DNS with Active Directory, you must plan the level of integration and identify your security, scalability, and performance requirements. Figure 3.1 shows the DNS deployment process.

**Figure 3.1 Deploying DNS**



# Concepts about DNS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 DNS is based on Requests For Comments (RFCs) standards developed by the Internet Engineering Task Force (IETF) and is therefore interoperable with other standards-compliant DNS implementations. DNS uses a distributed database that implements a hierarchical naming system. This naming system enables an organization to expand its presence on the Internet and enables the creation of names that are unique both on the Internet and on private TCP/IP-based intranets.

By using DNS, any computer on the Internet can look up the name of any other computer in the Internet namespace. Computers running Windows Server 2003 and Microsoft® Windows® 2000 also use DNS to locate domain controllers and other servers running Active Directory.

**DNS Roles**

Deploying a DNS infrastructure involves design, implementation, and maintenance tasks. The individuals who are responsible for these tasks include DNS designers and the DNS administrators. Before you begin designing your DNS deployment, it is helpful to identify the individuals in your organization who are responsible for these roles. Table 3.1 lists the responsibilities of the DNS designer and DNS administrator roles.

**Table 3.1 DNS Roles**

| Role | Responsibility |
|---|---|
| DNS designer | • Designing the DNS namespace<br>• Placing DNS servers and zones within the DNS namespace<br>• Creating a secure DNS infrastructure<br>• Designing DNS integration with Active Directory |
| DNS administrator | • Deploying, configuring, and managing the DNS infrastructure<br>• Managing Active Directory integration |

**DNS designer role**

If you are deploying DNS to support Active Directory in an environment that does not already have a DNS infrastructure, the DNS designer is responsible for the DNS integration with the entire Active Directory forest. The DNS designer works closely with the DNS administrator for Active Directory.

If you are deploying DNS to support Active Directory in an environment that has an existing DNS infrastructure, the DNS designer works with the DNS administrator for Active Directory to delegate the forest root DNS name to Active Directory. The Active Directory forest administrator delegates management of DNS to a DNS administrator.

**DNS administrator role**

DNS administrators manage and maintain the DNS namespace, DNS servers, DNS clients, DNS zones, and zone propagation. DNS administrators are also responsible for maintaining network security by anticipating and mitigating new security threats. In addition, DNS administrators are responsible for DNS integration with other Windows Server 2003 services.

**New in Windows Server 2003**

Windows Server 2003 DNS includes several new features, including:

- **Conditional forwarding.** Conditional forwarding enables a DNS server to forward DNS queries based on the DNS domain name in the query. For more information about conditional forwarding, see Help and Support Center for Windows Server 2003.

- **DNS application directory partitions.** DNS application directory partitions enable you to set the replication scope for Active Directory–integrated DNS data. By limiting the scope of replication traffic to a subset of the servers running Active

Directory in your forest, you can reduce replication traffic.

- **DNSSEC.** DNS provides basic support for the DNS Security Extensions (DNSSEC) protocol as defined in RFC 2535: *Domain Name System Security Extensions*. For more information about DNSSEC, see Help and Support Center for Windows Server 2003.

- **EDNS0.** Extension Mechanisms for DNS (EDNS0) enable DNS requestors to advertise the size of their UDP packets and facilitate the transfer of packets larger than 512 octets, the original DNS limit for UDP packet size. For more information about EDNS0, see Help and Support Center for Windows Server 2003.

**Tools for Deploying DNS**

Windows Server 2003 includes a number of tools to assist you in deploying a DNS infrastructure.

**Netdiag.exe**

The Netdiag.exe tool assists you in isolating networking and connectivity problems. Netdiag.exe performs a series of tests that you can use to determine the state of your network client. For more information about Netdiag.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Windows Support Tools.

**Nslookup.exe**

You can use the Nslookup.exe command-line tool to perform query testing of the DNS domain namespace and to diagnose problems with DNS servers.

**Dnscmd.exe**

You can use the Dnscmd.exe command-line tool to perform administrative tasks on the DNS server the same as you can by using the DNS Microsoft Management Console (MMC) snap-in.

**DNSLint**

DNSLint is a command-line tool that you can use to address some common DNS name resolution issues, such as lame delegation and DNS record verification. DNSLint is in the Support.cab file in the \Support\Tools folder on the Windows Server 2003 operating system CD. You can install DNSLint by running Suptools.msi.

**Terms and Definitions**

The following are some important DNS-related terms.

**Authoritative DNS server** A DNS server that hosts a primary or secondary copy of zone data. Each zone has at least one authoritative DNS server.

**Conditional forwarding** A DNS query setting that enables a DNS server to route a request for a particular name to another DNS server by specifying a name and IP address. For example, a DNS server in contoso.com can be configured to forward queries for names in treyresearch.com to a DNS server hosting the treyresearch.com zone.

**Delegation** The process of using resource records to provide pointers from parent zones to child zones in a namespace hierarchy. This enable DNS servers in a parent zone to route queries to DNS servers in a child zone for names within their branch of the DNS namespace. Each delegation corresponds to at least one zone.

**DNS client resolver** A service that runs on client computers and sends DNS queries to a DNS server. Some resolvers use a cache to improve name resolution performance.

**DNS namespace** The hierarchical naming structure of the domain tree. Each domain label that is used in a fully qualified domain name (FQDN) indicates a node or branch in the domain tree. For example, host1.contoso.com is an FQDN that represents the node host1, under the node Contoso, under the node com, under the DNS root.

**DNS server** A computer that hosts DNS zone data, resolves DNS queries, and caches the query responses.

**Domain tree** In DNS, the inverted hierarchical tree structure that is used to index domain names within a namespace. Domain trees are similar in purpose and concept to the directory trees used by computer filing systems for disk storage.

**Public namespace** A namespace on the Internet, such as www.microsoft.com, that can be accessed by any connected device. Beneath the top-level domains, the Internet Corporation for Assigned Names and Numbers (ICANN), the Internet Assigned Numbers Authority (IANA), and other Internet naming authorities delegate domains to organizations such as Internet Service Providers (ISPs), which in turn delegate subdomains to their customers or host zones for their customers. For more information about public namespaces, see the Internet Assigned Numbers Authority (IANA) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Forward lookup zone** An authoritative DNS zone that is primarily used to resolve network resource names to IP addresses.

**Fully qualified domain name (FQDN)** A DNS name that uniquely identifies a node in a DNS namespace. The FQDN of a computer is a concatenation of the computer name (for example, client1) and the primary DNS suffix of the computer (for example, contoso.com), and a terminating dot (for example, contoso.com.).

**Internal namespace** A namespace internal to an organization to which it can control access. Organizations can use the internal namespace to shield the names and IP addresses of its internal computers from the Internet. A single organization might have multiple internal namespaces. Organizations can create their own root servers and any subdomains as needed. The internal namespace can coexist with an external namespace.

**Iterative query** A query made by a client to a DNS server for an authoritative answer that can be provided by the server without generating additional server-side queries to other DNS servers.

**Primary DNS server** A DNS server that hosts read-write copies of zone data, has a DNS database of resource records, and resolves DNS queries.

**Secondary DNS server** A DNS server that hosts a read-only copy of zone data. A secondary DNS server periodically checks for changes made to the zone on its configured primary DNS server, and performs full or incremental zone transfers, as needed.

**Recursive query** A query made by either a client or a DNS server on behalf of a client, the response to which can be an authoritative answer or a referral to another server. Recursive queries continue until the DNS server receives an authoritative answer for the queried name. By default, recursion is enabled for Windows Server 2003 DNS.

**Resource record (RR)** A DNS database structure containing name information for a particular zone. For example, an address (A) resource record can map the IP address 172.16.10.10 to the name DNSserverone.contoso.com or a namespace (NS) resource record can map the name contoso.com to the server name DNS1.contoso.com. Most of the basic RR types are defined in RFC 1035: *Domain Names — Implementation and Specification*, but additional RR types are defined in other RFCs.

**Reverse lookup zone** An authoritative DNS zone that is primarily used to resolve IP addresses to network resource names.

**Stub zone** A partial copy of a zone that can be hosted by a DNS server and used to resolve recursive or iterative queries. Stub zones contain the Start of Authority (SOA) resource records of the zone, the DNS resource records that list the zone's authoritative servers, and the glue address (A) resource records that are required for contacting the zone's authoritative servers. Stub zones are used to reduce the number of DNS queries on a network, and to decrease the network load on the primary DNS servers hosting a particular name.

**Zone** In a DNS database, a contiguous portion of the domain tree that is administered as a single separate entity by a DNS server. The zone contains resource records for all of the names within the zone.

**Zone file** A file that consists of the DNS database resource records that define the zone. DNS data that is Active Directory–integrated is not stored in zone files because the data is stored in Active Directory. However, DNS data that is not Active Directory–integrated is stored in zone files.

**Zone transfer** The process of copying the contents of the zone file located on a primary DNS server to a secondary DNS server. Using zone transfer provides fault tolerance by synchronizing the zone file in a primary DNS server with the zone file in a secondary DNS server. The secondary DNS server can continue performing name resolution if the primary DNS server fails.

# Examining Your Current Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy Windows Server 2003 DNS, you must assess your current environment to determine the DNS needs and constraints of your organization. After that, create a Windows Server 2003 DNS deployment plan to match those needs and constraints. Figure 3.2 shows the process for examining your current environment.

**Figure 3.2 Examining Your Current Environment**



# Determining Internet Status

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you want devices outside your private network to be able to resolve names belonging to your internal namespace, your IP addresses and DNS domain names must be registered with an Internet registration authority (Internet registrar). Internet registrars are organizations that are responsible for:

- Assigning IP addresses.
- Registering DNS domain names.
- Keeping public records of registered IP addresses and domain names.

If your network is currently or will be connected to the Internet, you must ensure that the domain name of your organization is valid and registered to you.

# Identifying the DNS Data Host

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Determine who will host your DNS data. You can either host your DNS data on your own DNS servers or you can have an external ISP host your DNS data. By hosting your own DNS data, you have complete control of network resource allocation and security. Use an ISP if you have a small network that will not be expanding rapidly in the near future.

If you decide to use an ISP to host your DNS data, ensure that the DNS infrastructure of the ISP can support your DNS deployment.

Even if you decide to host your own DNS data, you might consider using an ISP to resolve names outside your private network. For example, the company contoso.com might decide to host all names belonging to the internal namespace corp.contoso.com on its own DNS servers while using an ISP to enable its employees to resolve external Web addresses such as www.microsoft.com.

# Analyzing Your Network Topology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Analyze your existing network topology and plan your DNS namespace while accounting for the service and administrative goals of your organization.

Allow for anticipated expansion of the number of nodes in your DNS hierarchy by including domain name placeholders between the domain names that you are initially deploying. Adding domain name placeholders enables you to avoid having to redesign your DNS infrastructure to accommodate additional domain names.

Anticipate the possibility of changes to your business model by assigning domain names that can be used in a modified context. For example, instead of using the domain name *accounting.contoso.com*, you might use *finance.contoso.com*, in anticipation of the possibility of expanding into additional financial services beyond accounting.

**Diagram Your Existing DNS Infrastructure**

If you are upgrading to Windows Server 2003, rolling out a new DNS deployment, or integrating Windows Server 2003 DNS with Active Directory, you might not need to make any changes to your existing DNS infrastructure. However, if you are migrating from a third-party DNS infrastructure or integrating Windows Server 2003 DNS with an existing third-party DNS infrastructure, create a diagram of your existing DNS infrastructure, including domains, servers, and clients. Use this diagram to assist you in deciding whether to make changes to your current DNS infrastructure when you deploy Windows Server 2003 DNS.

**Identify Your Security Policies**

Identify and document your organization's security policies before you begin to design and deploy your DNS infrastructure. In this way, you can ensure that your DNS servers, zones, and resource records support these policies. For more information about security policies, see "Deploying Security Policy" in *Designing a Managed Environment* of this kit.

# Designing a DNS Namespace

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy a DNS infrastructure, the DNS designer in your organization must design a DNS namespace. You can design an external namespace that is visible to Internet users and computers, or you can design an internal namespace that is accessible only to users and computers that are within the internal network. After your DNS namespace has been deployed, DNS administrators are responsible for managing and maintaining the DNS namespace. Figure 3.3 shows the process for designing a DNS namespace.

**Figure 3.3 Designing a DNS Namespace**



# Identifying Your DNS Namespace Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in designing a DNS namespace is to determine whether you need a new namespace for your organization, or whether you can retain an existing Windows or third-party DNS namespace.

**Important**

- You must plan your DNS namespace in conjunction with planning your Active Directory logical structure. For more information about designing the Active Directory logical structure, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

Table 3.2 summarizes the DNS namespace design requirements for each possible scenario.

**Table 3.2 DNS Namespace Design Requirements**

| Scenario | Design Requirements |
|---|---|
| You are upgrading an existing DNS infrastructure from a version of Windows earlier than Windows Server 2003. | DNS namespace design can remain the same. |
| You are upgrading from a third-party DNS infrastructure that uses DNS software that adheres to standard DNS domain naming guidelines. | DNS namespace design can remain the same. |
| Your existing DNS software does not conform to standard DNS domain naming guidelines. | Bring your existing DNS namespace design into compliance with DNS domain naming guidelines before deploying a Windows Server 2003 DNS namespace. |
| You are integrating Windows Server 2003 DNS into an existing third-party DNS software that adheres to standard DNS domain naming guidelines. | Integrate Windows Server 2003 DNS with your current DNS infrastructure. You do not need to change the namespace design of the third-party DNS infrastructure or your existing namespace. |
| You are deploying a new Windows Server 2003 DNS infrastructure. | Design a logical naming convention for your DNS namespace based on DNS domain naming guidelines. |
| You are deploying Windows Server 2003 DNS to support Active Directory. | Create a DNS namespace design that is based on your Active Directory naming convention. |
| You are modifying your existing DNS namespace to support Active Directory, but you do not want to redesign your DNS namespace. | Ensure that Active Directory domain names match your existing DNS names. This enables you to deploy the highest level of security by using the simplest management techniques. |

# Creating Internal and External Domains

Updated: September 14, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Organizations that require an Internet presence as well as an internal namespace must deploy both an internal and an external DNS namespace and manage each namespace separately. You can create a mixed internal and external DNS namespace in one of two ways:

- By making the internal domain a subdomain of the external domain.
- By using different names for the internal and external domains.

| **Note** |
| --- |
| You can also use the same name for the internal domain and the external domain. However, this method is not recommended. It creates name resolution problems because it introduces DNS names that are not unique. This method requires additional configuration to enable optimized performance. |

Select the configuration design option that best meets the needs of your organization. Table 3.3 lists the design options for deploying a mixed internal and external namespace and the level of management complexity for each option, along with an example to illustrate each option.

**Table 3.3 Mixed Internal and External DNS Namespace Design Options**

| Design Option | Management Complexity | Example |
| --- | --- | --- |
| The internal domain is a subdomain of the external domain. Microsoft strongly recommends this option. For more information, see Using an Internal Subdomain. | Easy to deploy and administer. | An organization with an external namespace contoso.com uses the internal namespace corp.contoso.com. |
| The internal and external domain names are different from each other. For more information, see Using Different Internal and External Domain Names. | More complicated than previous option. | An organization uses contoso.com for its external namespace, and corp.internal for its internal namespace. |

# Using an Internal Subdomain

Updated: September 14, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The recommended configuration option for a mixed internal and external DNS namespace is to make your internal domain a subdomain of your external domain. For example, an organization that has an external namespace domain name of contoso.com might use the internal namespace domain name corp.contoso.com. Using an internal domain that is a subdomain of an external domain:

- Requires you to register only one name with an Internet name authority even if you later decide to make part of your internal namespace publicly accessible.
- Ensures that all of your internal domain names are globally unique.
- Simplifies administration by enabling you to administer internal and external domains separately.

You can use your internal subdomain as a parent for additional child domains that you create to manage divisions within your company.

For example, to implement an internal subdomain of the external domain contoso.com:

1. Configure the external DNS server with the namespace contoso.com. This server contains a static zone with only records for servers that are to be available publicly on the Internet. These records would typically include such servers as www.contoso.com, ftp.contoso.com, and so on.
2. Configure the internal DNS server with the namespace corp.contoso.com. Note that while the internal namespace is a subdomain of the external namespace, the internal zone is not delegated from the external zone. That is, the external server does not have a delegation record.
3. Register all hosts in the organization's internal network in the corp.contoso.com namespace, either in the corp.contso.com domain or in child domains within that namespace. For example, a server for the sales department might have the fully qualified domain name of salesvr1.sales.corp.contoso.com.
4. Disable dynamic updates on the external server because no computers in the company's internal network should be registered in the external namespace.
5. Configure the internal DNS server to forward Internet queries to enable internal hosts to resolve external (Internet) names. You can do this in one of two ways:

   - Configure the internal DNS server to forward to the external server, which enables recursion on the external server. In addition, you can configure the internal server as a secondary server to the external server. This enhances the security of the internal network by ensuring that the internal server never makes a query directly on the Internet.
   - Configure the internal DNS server to forward queries to the Internet service provider's DNS server. This method enhances security on the external server because it makes it possible to disable recursion on the external server.

   Instead of using forwarders, the internal DNS server may be configured simply to use the standard Internet root hints to resolve Internet names. The DNS server is configured with the list of standard root hints when the DNS server role is installed. This option is the least secure, however, because it requires the internal DNS server to be able to access DNS servers on the Internet.

**See Also**
**Concepts**
Namespace planning for DNS

Securing DNS deployment

**Other Resources**

http://go.microsoft.com/fwlink/?LinkId=99540

http://go.microsoft.com/fwlink/?LinkId=99541

# Using Different Internal and External Domain Names

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If it is not possible for you to configure your internal domain as a subdomain of your external domain, use a stand-alone internal domain. This way, your internal and external domain names are unrelated. For example, an organization that uses the domain name contoso.com for their external namespace uses the name corp.internal for their internal namespace.

The advantage to this approach is that it provides you with a unique internal domain name. The disadvantage is that this configuration requires you to manage two separate namespaces. Also, using a stand-alone internal domain that is unrelated to your external domain might create confusion for users because the namespaces do not reflect a relationship between resources within and outside of your network. In addition, you might have to register two DNS names with an Internet name authority if you want to make the internal domain publicly accessible.

# Deciding Whether to Deploy an Internal DNS Root

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have a large distributed network and a complex DNS namespace, it is best to use an internal DNS root that is isolated from public networks. Using an internal DNS root streamlines the administration of your DNS namespace by enabling you to administer your DNS infrastructure as if the entire namespace consists of the DNS data within your network.

If you use an internal DNS root, a private DNS root zone is hosted on a DNS server on your internal network. This private DNS root zone is not exposed to the Internet. Just as the DNS root zone contains delegations to all of the top-level domain names on the Internet, such as .com, .net, and .org, a private root zone contains delegations to all of the top-level domain names on your network. The DNS server that hosts the private root zone in your namespace is considered to be authoritative for all of the names in the internal DNS namespace.

Using an internal DNS root provides the following benefits:

- **Simplicity.** If your network spans multiple locations, an internal DNS root might be the best method for administering DNS data in a network.

- **Secure name resolution.** With an internal DNS root, DNS clients and servers on your network never contact the Internet to resolve internal names. In this way, the DNS data for your network is not broadcast over the Internet. You can enable name resolution for any name in another namespace by adding a delegation from your root zone. For example, if your computers need access to resources in a partner organization, you can add a delegation from your root zone to the top level of the DNS namespace of the partner organization.

**Important**

- Do not reuse names that exist on the Internet in your internal namespace. If you repeat Internet DNS names on your intranet, it can result in name resolution errors.

If name resolution is required by computers that do not support software proxy, or by computers that support only LATs, then you cannot use an internal root for your DNS namespace. In this case, you must configure one or more internal DNS servers to forward queries that cannot be resolved locally to the Internet.

Table 3.4 lists the types of client proxy capabilities and whether you can use an internal DNS root for each type.

**Table 3.4 Client Proxy Capabilities**

| Proxy Capability | Microsoft Software with Corresponding Proxy Capabilities | Forwards Queries | Can You Use an Internal Root? |
|---|---|---|---|
| No Proxy | Generic Telnet | ● | |
| Local Address Table (LAT) | Winsock Proxy (WSP) 1.*x* and later<br>Microsoft® Internet Security and Acceleration (ISA) Server 2000 and later | ● | |
| Name Exclusion List | WSP 1.*x* and later<br>Internet Security and Acceleration (ISA) Server 2000 and later, and all versions of Microsoft® Internet Explorer | ● | ● |
| Proxy Auto-configuration (PAC) File | WSP 2.*x*, Internet Security and Acceleration Server (ISA) Server 2000 and later, Internet Explorer 3.01 and later | ● | ● |

# Configuring Name Resolution for Disjointed Namespaces

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you need to create or merge two DNS namespaces when you deploy Windows Server 2003 DNS, this can result in *disjointed namespaces* — a DNS infrastructure that includes two or more top-level DNS domain names. To configure internal name resolution for multiple DNS top-level domains, you must do one of the following:

- If you have an internal DNS root, add delegations for each top-level DNS zone to the internal DNS root zone.
- If you want to reduce cross-domain DNS query traffic, configure the DNS servers that host the DNS zones in the first and second namespaces to host secondary zones for the DNS zones in each other's namespaces. In this configuration, the DNS servers that host the DNS zones in each namespace are aware of the DNS servers in the other namespace. This solution requires increased storage space for hosting secondary copies of zones in different namespaces, and generates increased zone transfer traffic.
- If storage capacity on DNS servers is a consideration, configure the DNS servers that host the DNS zones in one namespace to forward name resolution queries in a second namespace to the DNS servers that are hosting the DNS zones for the second namespace. Then configure the DNS servers that host the DNS zones in the second namespace to forward name resolution queries in the first namespace to the DNS servers that are hosting the DNS zones for the first namespace. You can use Windows Server 2003 DNS conditional forwarders for this configuration.

You can also use Windows Server 2003 DNS stub zones to facilitate DNS data distribution between separate namespaces. For more information about conditional forwarders and stub zones, see Help and Support Center for Windows Server 2003 and the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Integrating a Windows Server 2003 DNS Infrastructure Into an Existing DNS Namespace

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 DNS is standards-compliant and interoperates with other implementations of DNS, including Microsoft® Windows NT® version 4.0, BIND 9.1.0, BIND 8.2, BIND 8.1.2, and BIND 4.9.7. The complexity of your integration process depends, in part, on the DNS features that you need to support. If the computers in your DNS infrastructure are running versions of DNS that support the same features, then integrating the Windows Server 2003 DNS infrastructure is a simple process. If the computers in your DNS infrastructure are running versions of DNS that do not support the same DNS features, then the integration process is more complex.

Table 3.5 compares feature support in Windows Server 2003 DNS and other implementations of DNS.

**Table 3.5 Feature Support in Different Implementations of DNS**

| Feature | Windows Server 2003 | Windows 2000 | Windows NT 4.0 | BIND 9 | BIND 8.2 | BIND 8.1.2 | BIND 4.9.7 |
|---|---|---|---|---|---|---|---|
| Supports RFC 2782: *A DNS RR for specifying the location of services (DNS SRV)* | ● | ● | ● | ● | ● | ● | ● |
| Dynamic update | ● | ● | | ● | ● | ● | |
| Secure dynamic update based on the GSS-Transaction signature (TSIG) algorithm | ● | ● | | | | | |
| WINS and WINS-R records | ● | ● | ● | | | | |
| Incremental zone transfer | ● | ● | | ● | ● | | |
| UTF-8 character encoding | ● | ● | | | | | |
| DNS MMC snap-in | ● | ● | | | | | |
| Dnscmd.exe | ● | ● | | | | | |
| Active Directory–integrated zones | ● | ● | | | | | |
| Storage of zones in the DNS application directory partition | ● | | | | | | |
| Aging and scavenging of obsolete records | ● | ● | | | | | |
| Stub zones | ● | ● | | ● | | | |
| Conditional forwarding | ● | | | ● | | | |

# Creating DNS Domain Names and Computer Names

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy your Windows Server 2003 DNS infrastructure, you must create a naming convention for your DNS Internet and internal domains and the DNS computers on your network. To create a DNS naming convention, you must establish the following:

- An Internet DNS domain name, if your organization is connected to the Internet.
- An internal DNS domain name for your organization.
- A DNS resource naming convention.

In addition, you must determine whether you need to support NetBIOS names in your organization.

# Creating an Internet DNS Domain Name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are deploying a new Windows Server 2003 DNS infrastructure that is connected to the Internet, you must create an Internet DNS domain name for your organization. Because all of the nodes in your network that require name resolution are assigned a DNS name that includes the Internet DNS domain name for your organization, it is important to select an Internet DNS domain name that is short and easy to remember. Because DNS is hierarchical, DNS domain names grow when you add subdomains to your organization. Short domain names result in computer names that are easy to remember, facilitating resource access.

A DNS namespace that is connected to the Internet must be a subdomain of a top-level or second-level domain of the Internet DNS namespace. If you are deploying a new Windows Server 2003 DNS namespace, you must select a top-level Internet DNS domain in which to register your Internet DNS domain name. For example, you can register your domain as a subdomain of .com, .org, or .net, or as a subdomain of the domain name that is assigned to your country/region, such as .au (Australia), .fr (France) or .ca (Canada).

When you have selected your Internet DNS domain name and identified the top-level domain for which your DNS domain is a subdomain, complete the following steps to register your DNS domain name:

1. Search the Internet to confirm that the DNS domain name that you selected for your organization is not registered to another organization. If the DNS domain name that you selected is owned by another organization, you can attempt to buy it from that organization, or select a different DNS domain name.

2. Configure at least one authoritative DNS server to host the DNS zone for your domain name. This DNS server might be located on your network or on the network of your ISP.

3. Register your DNS domain name with an Internet registrar, and supply the registrar with the DNS name and IP address of at least one DNS server that is authoritative for your DNS domain name. For a list of Internet registrars, see the ICANN link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

The Internet domain name registration process varies according to the design of your DNS namespace. Table 3.6 lists the domain names that you need to register for each type of DNS namespace design.

**Table 3.6 Internet DNS Domain Name Registration**

| Namespace Design | Domain Name Registration | Example |
|---|---|---|
| The internal domain name is a subdomain of the external domain. | Register only the external domain name. | The domain name contoso.com is used for the external namespace.<br>The domain name corp.contoso.com is used for the internal namespace. |
| The internal and external domain names are different from each other. | Register the external domain name, and then, if you want the internal domain to be publicly accessible, also register the internal domain name. | The domain name contoso.com is used for the external namespace.<br>The domain name corp.contoso.com is used for the internal namespace. |

When you register your DNS domain name, the Internet registrar creates a delegation in the DNS zone that is authoritative for the top-level domain that you selected. This is the top-level domain for the DNS servers that are authoritative for your organization's Internet DNS domain name.

**Note**

- If a domain name that you want to register is not available in one top-level domain, such as .com, and you register the same domain name in another top-level domain, such as .net, then people who are searching for your domain name on the Internet might assume that computers and services in the wrong top-level domain belong to your company.

# Creating Internal DNS Domain Names

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When creating names for your internal domains, use the following guidelines:

- If your organization has an Internet presence, use names relative to your registered Internet DNS domain name. For example, if you have registered the Internet DNS domain name contoso.com for your organization, use a DNS domain name such as corp.contoso.com for your intranet domain name.
- Do not use the name of an existing corporation or product as your domain name.
- Do not use top-level Internet domain names, such as .com, .net, .org, .us, .fr, .gr, on your intranet. Using top-level Internet domain names on your intranet can result in name resolution errors for computers on your network that are connected to the Internet.
- Do not use acronyms or abbreviations for domain names. The business units that these acronyms represent can be difficult for users to recognize.
- Do not use business unit or division names for domain names. Business units and other divisions change periodically and the domain names can become obsolete or misleading.
- Do not use geographic names that are difficult to spell and remember.
- Avoid extending your DNS domain name hierarchy more than five levels from the internal or DNS root domain. Limiting the extent of the domain name hierarchy reduces administrative costs.

If you are deploying DNS in a private network and do not plan to create an external namespace, it is recommended that you register the DNS domain name that you create for your internal domain. If you do not register the name and later attempt to use it on the Internet, or connect to a network that is connected to the Internet, you might find that the name is unavailable.

# Creating DNS Computer Names

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to develop a practical DNS computer naming convention for your network. This enables users to remember the names of computers on public and private networks easily, and therefore facilitates access to resources on your network.

The DNS computer name creation process varies according to whether you are creating a new DNS infrastructure, integrating your DNS infrastructure with an existing third-party infrastructure, or upgrading an existing DNS infrastructure.

**Creating Computer Names in a New Windows Server 2003 DNS Infrastructure**
Use the following guidelines when creating names for the DNS computers in your new Windows Server 2003 DNS infrastructure:

- Select computer names that are easy for users to remember.
- Identify the owner of a computer in the computer name. For example, john-doe-1 indicates that John Doe uses the computer.
- Select names that describe the purpose of the computer. For example, a file server named past-accounts-1 indicates that the file server stores information related to past accounts.
- Do not use character case to convey the owner or purpose of a computer. DNS is not case-sensitive.
- If you are deploying DNS to support Active Directory, match the Active Directory domain name to the primary DNS suffix of the computer name. For more information about designing the Active Directory logical structure, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.
- Use unique names for all computers in your organization. Do not assign the same computer name to different computers in different DNS domains.
- Use ASCII characters to ensure interoperability with computers running versions of Windows earlier than Windows 2000. For DNS computer names, use only the characters listed in RFC 1123:*Requirements for Internet Hosts — Application and Support*, which include A–Z, a–z, 0–9, and the hyphen (-). Windows Server 2003 DNS supports almost any UTF-8 character in a name; however, do not use extended ASCII or UTF-8 characters unless all of the DNS servers in your environment support them.

**Note**

- Windows Server 2003 DNS is configured to use UTF-8 name checking by default.

**Creating Computer Names in an Integrated DNS Infrastructure**
If you are integrating Windows Server 2003 DNS with an existing third-party DNS infrastructure, you do not need to make any changes to your third-party DNS host names. If you are migrating to Windows Server 2003 DNS from a third-party DNS infrastructure, you must ensure that the host names that are used in the third-party DNS infrastructure conform to the DNS Internet naming standards.

If you are integrating or migrating an existing public DNS infrastructure that is connected to the Internet into your existing DNS infrastructure, you do not need to make any changes to the DNS domain names of your infrastructure.

**Creating Computer Names When Upgrading a DNS Infrastructure**
If you are upgrading to Windows Server 2003 DNS from Windows NT 4.0, you do not need to change your DNS host names; however, you might need to convert any NetBIOS names to DNS names. DNS names must conform to the DNS standard defined by RFC 1123.

Table 3.7 lists the different character sets that are supported by standard DNS, Windows Server 2003 DNS, and NetBIOS.

**Table 3.7 Character Set Restrictions**

| Character Set Restriction | Standard DNS (Including Windows NT 4.0) | DNS in Windows 2000 and Windows Server 2003 | NetBIOS |
|---|---|---|---|
| Characters permitted | Supports RFC 1123, which permits A–Z, a–z, 0–9, and the hyphen (-). | Supports RFC 1123 and UTF-8. On a per-server basis, You can configure the Windows 2000 DNS Server service to allow or disallow the use of UTF-8 characters on your DNS server. | Not allowed: Unicode characters, numbers, white space, symbols: / \ [ ] : \| < > + = ; , ? and *) |
| Maximum host name and FQDN length. | 63 octets per label.<br>255 bytes per FQDN (254 bytes for the FQDN plus one byte for the terminating dot). | The same as standard DNS with the addition of UTF-8 support. Character count is insufficient to determine size because some UTF-8 characters exceed one octet in length. Domain controllers are limited to 155 bytes for an FQDN. | 16 bytes in length. |

**Important**

- Names encoded in UTF-8 format must not exceed the limits defined in RFC 2181: *Clarifications to the DNS Specification*, which specifies a maximum of 63 octets per label and 255 octets per name.

# Determining if You Need to Support NetBIOS Names

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

During a domain upgrade to Windows Server 2003, you might need to support NetBIOS on your network if your domain includes clients that are running versions of Windows earlier than Windows 2000. For example, if your network is multi-segmented, WINS is required to create the NetBIOS browse list. Without WINS, the network must rely on Active Directory for browsing resources. This can have a significant impact on clients that are running applications that require NetBIOS support, even if the client operating system does not require NetBIOS support. When WINS is installed, performance monitor counters for WINS are also installed. Use these WINS performance monitor counters to determine how many queries WINS is receiving, and how many names WINS is resolving. This information will help you to determine whether it is necessary to support NetBIOS names on the network.

Windows Server 2003 DNS is compatible with WINS; therefore, in a mixed networking environment, you can use a combination of DNS and WINS. Windows NT 4.0–based clients can register in both Windows 2000 WINS and Windows Server 2003 WINS. Also, computers running either the Microsoft® Windows® 2000 Professional or Windows® XP Professional operating systems can register in Windows NT 4.0 WINS. To maintain backward compatibility, each computer is given a NetBIOS name that must be unique in the domain to which the computer belongs.

Preserving existing NetBIOS names can be difficult because NetBIOS names have a broader character set than DNS names. One solution is to replace NetBIOS names with DNS names to ensure that the names adhere to existing DNS naming standards. This is not possible for organizations that support computers running versions of Windows earlier than Windows 2000.

RFC 2181: *Clarifications to the DNS Specification* expands the character set that is allowed in DNS names to include any binary string. The binary strings do not have to be interpreted as ASCII. Windows 2000 and Windows Server 2003 support UTF-8 character encoding (RFC 2044). UTF-8 is a superset of ASCII and a translation of the UCS-2 (or Unicode) character encoding. The UTF-8 character set enables you to transition from Windows NT 4.0 NetBIOS names to Windows 2000 and Windows Server 2003 DNS names

By default, multibyte UTF-8 name checking is used. This provides the greatest tolerance when the DNS service processes characters. This is the preferred name-checking method for most DNS servers that are not providing name resolution services for Internet hosts.

**Important**

- Windows Server 2003 and Windows 2000 DNS support NetBIOS and UTF-8 characters for computer names. Other versions of DNS only support the characters permitted in RFC 1123. Therefore, only use NetBIOS and UTF-8 character sets when you are certain that Windows Server 2003 or Windows 2000 DNS is the method used for name resolution. Names that are intended to be visible on the Internet must contain ASCII-only characters, as recommended in RFC 1123.

# Creating Subdomains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are deploying DNS on a large enterprise network, or if you expect your network to expand to include additional subnets and sites, consider distributing the management of portions of your DNS namespace to the administrators for the different subnets and sites in your network. To distribute the management of your DNS namespace, create subdomains of your initial DNS domain and delegate the authority for these subdomains to DNS servers located on different subnets or sites. In this way, you can create any number of separate and autonomous entities within a DNS namespace, each of which is authoritative for a portion of the overall namespace.

# Example: Merging DNS Namespaces

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso Corporation merged with Trey Research Corporation. Before the merger, each corporation used internal domains that were subdomains of their external domains. The Contoso Corporation used a private root to simplify their DNS server administration. The Trey Research Corporation forwarded queries to the Internet, rather than using a private root.

The external namespace of the newly merged corporation contains the zones contoso.com and treyresearch.com. Each zone in the external namespace contains the DNS resource records that the companies want to expose to the Internet. The internal namespace contains the internal zones, *corp.contoso.com* and *corp.treyresearch.com*.

The Contoso division and the Trey Research division each use a different method to support name resolution for names in their namespace. The Contoso division uses the name contoso.com externally and corp.contoso.com internally. The internal root servers host the root zone. Internal servers also host the zone, corp.contoso.com. The name contoso.com is registered with an Internet name authority.

To ensure that every client within the organization can resolve every name in the newly merged organization, the private root zone contains a delegation to the zone for the top level of the merged organization's internal namespace, corp.treyresearch.com.

To resolve internal and external names, every DNS client must submit all queries to either the internal DNS servers or to a proxy server. Figure 3.4 shows this configuration.

**Figure 3.4 Name Resolution in the Contoso Division**




Based on this configuration, internal clients can query for names in the following ways:

- **Query internal DNS servers for internal names.** The internal DNS servers resolve the query. If a DNS server that receives a query does not contain the requested data in its zones or cache, it uses root hints to contact the internal root DNS servers.

- **Query a proxy server for names on the Internet.** The proxy server forwards the query to DNS servers on the Internet. The DNS servers on the Internet resolve the query.

- **Query internal DNS servers for names in the Trey Research division.** Because the root servers contain a delegation to the top level of the DNS namespace of the Trey Research division, the internal DNS servers recursively resolve the query by contacting the DNS servers in the Trey Research division.

External clients:

- **Cannot query for internal names.** This limitation helps secure the internal network.

- **Query DNS servers on the Internet for names in the contoso.com external namespace.** The DNS servers on the Internet resolve the query.

The Trey Research division uses the name treyresearch.com externally and the name corp.treyresearch.com internally. The server InternalDNS.treyresearch.com hosts the corp.treyresearch.com zone. The Trey Research division does not have a private root.

To simplify management of clients and DNS servers, Trey Research division administrators decided to use conditional forwarding. Administrators configured the DNS server InternalDNS.treyresearch.com to forward queries in the following manner:

- The server forwards all queries destined for the Contoso division to a DNS server for the Contoso division. For example, the server forwards queries destined for corp.contoso.com to InternalDNS.contoso.com.

- At the same time, the server forwards all other queries destined for contoso.com to a DNS server on the Internet.

Figure 3.5 shows this configuration.

**Figure 3.5 Conditional Forwarding in the Trey Research Division**



# Designing a DNS Server Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DNS servers store information about the DNS namespace and use the information to answer queries from DNS clients. The size of the DNS zone data, how many DNS clients you have, and where these clients are physically located all impact your DNS server topology.

The DNS designer in your organization designs DNS servers that enable you to create an effective DNS data distribution and update topology while minimizing query and zone transfer network traffic. The DNS administrators in your organization manage and maintain your DNS servers. Figure 3.6 shows the process for designing DNS servers.

**Figure 3.6 Designing a DNS Server Infrastructure**



# Allocating Hardware Resources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A typical recommendation for DNS server hardware includes the following:

- Single-processor computers with 400 megahertz (MHz) Pentium II CPUs.
- 256 megabytes (MB) of RAM for each processor.
- At least 4 gigabytes (GB) of available hard disk space.
- A network adapter.

Using faster CPUs, more RAM, and larger hard drives improves the scalability and performance of your DNS servers. DNS servers use approximately 100 bytes of RAM for each resource record. Using this figure, you can calculate how much memory you need.

# Determining the Number of Required DNS Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To reduce administrative overhead, use the minimum number of DNS servers required. Be sure to make at least two DNS servers authoritative for each zone to enable fault tolerance and load sharing.

Add additional DNS servers in order to:

- Provide redundancy when your namespace design requires greater DNS availability.
- Improve query response time when better DNS performance is required.
- Reduce WAN traffic for remote locations.

Use the following guidelines to determine the number of DNS servers that you need to deploy:

- If the ratio of DNS servers to clients is very low and you are experiencing significant name resolution delays, add additional DNS servers to host secondary or Active Directory–integrated zones. Use your anticipated number of queries and dynamic updates per second to determine the number of DNS servers that you need. The Windows Server 2003 DNS Server service is capable of responding to more than 10,000 queries per second on a Pentium III microprocessor running at 700 MHz.

  For information about capacity planning, see "Allocating Hardware Resources" earlier in this chapter.

- If you delegate zones, add additional DNS servers to handle the delegated zones. Note that you do not need to delegate zones when you have multiple zones. You can host all zones on the same server or servers. One DNS server running Windows Server 2003 can host 20,000 small zones.
- If you plan to host Active Directory–integrated zones, you must place these zones on Windows 2000–based or Windows Server 2003–based domain controller.
- If high-volume traffic is a consideration in your environment, add additional DNS servers to balance the workload. Although DNS helps reduce broadcast traffic between local subnets, it does create some traffic between servers and clients, particularly in complex routed environments. In addition, although the DNS service supports incremental zone transfers (IXFRs) and clients and servers can cache recently used names, traffic considerations can still remain an issue, depending on available bandwidth. This is especially true when using short Dynamic Host Configuration Protocol (DHCP) leases, which require more frequent dynamic updates.
- If you have a high number of client nodes on a single subnet, placing more than one DNS server on the subnet allows for backup and failover in the event that the primary DNS server stops responding.

If your DNS design includes primary and secondary zones and you run a large number of secondary servers for a zone, the primary DNS server can become overloaded when the secondary servers poll to ensure that their zone data is current. You can solve this problem in one of three ways:

- Use some of the secondary DNS servers as primary servers for the zone. Other secondary servers can poll and request zone updates from these primary servers.
- Increase the refresh interval so that the secondary servers poll less frequently. Note, however, that a longer refresh interval might cause your secondary zones to be outdated more often.

# Determining DNS Server Placement

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The placement of your DNS servers and the number of DNS servers that you deploy affects the availability of DNS. It is important to ensure that you plan the placement of your DNS servers to allow for DNS availability and Active Directory availability.

**Placing DNS Servers for Availability**
To ensure that DNS is always available, make sure that your DNS infrastructure does not include any single points of failure. To improve fault tolerance and load sharing have clients point to a primary and alternate DNS server. In a LAN configuration, place the pair of authoritative DNS servers on separate subnets. In a WAN configuration, place the pair of authoritative DNS servers on different networks, and then ensure that at least one DNS server is available for each network. This configuration removes routers as potential points of failure. Whenever possible, distribute your DNS servers across different geographic locations to enable communications to continue in the event of a natural disaster.

If you identify single points of failure in your network, determine whether they affect only DNS or all network services. If a router goes down and your clients cannot access any network services, then DNS failure is not an issue. If a router goes down and local DNS servers are unavailable but other network services are available, then your clients cannot access required network resources because they cannot look up DNS names.

If you have an Internet presence, DNS must be working properly for Internet clients to access your Web servers, send mail, and locate other services; therefore, it is recommended that you run a secondary DNS server offsite. If you have a business relationship with an organization on the Internet, either business partners or ISPs, they might agree to run a secondary server for you; however, ensure that the data on the organization's server is secured against Internet attackers.

To ensure that DNS is available if your offsite primary DNS servers are down, consider deploying a secondary DNS server offsite.

For more information about how to place DNS servers to maximize Active Directory availability, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

# Using Forwarding

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If a DNS server does not have the data to resolve a query in its cache or in its zone data, it forwards the query to another DNS server, known as a *forwarder*. Forwarders are ordinary DNS servers and require no special configuration; a DNS server is called a forwarder because it is the recipient of a query forwarded by another DNS server.

Use forwarding for off-site or Internet traffic. For example, a branch office DNS server can forward all off-site traffic to a forwarder at the company headquarters, and an internal DNS server can forward all Internet traffic to a forwarder on the external network. To ensure fault tolerance, forward queries to more than one forwarder.

Forwarders can increase network security by minimizing the list of DNS servers that communicate across a firewall.

You can use conditional forwarding to more precisely control the name resolution process. Conditional forwarding enables you to designate specific forwarders for specific DNS names. You can use conditional forwarding to resolve the following:

- Queries for names in off-site internal domains
- Queries for names in other namespaces

**Using Conditional Forwarding to Query for Names in Off-Site Internal Domains**

In Windows Server 2003 DNS, non-root servers resolve names for which they are not authoritative, do not have a delegation, and do not have in their cache by doing one of the following:

- Querying a root server.
- Forwarding queries to a forwarder.

Both of these methods generate additional network traffic. For example, a non-root server in Site A is configured to forward queries to a forwarder in Site B, and it must resolve a name in a zone hosted by a server in Site C. Because the non-root server can forward queries only to Site B, it cannot directly query the server in Site C. Instead, it forwards the query to the forwarder in Site B, and the forwarder queries the server in Site C.

When you use conditional forwarding, you can configure your DNS servers to forward queries to different servers based on the domain name specified in the query. This eliminates steps in the forwarding chain and reduces network traffic. When conditional forwarding is applied, the server in Site A can forward queries to forwarders in Site B or Site C, as appropriate.

For example, the computers in the Seville site need to query computers in the Hong Kong site. Both sites use a common DNS root server, DNS3.corp.fabrikam.com, located in Seville.

Before the Contoso Corporation upgraded to Windows Server 2003, the server in Seville forwarded all queries that it could not resolve to its parent server, DNS1.corp.contoso.com, in Seattle. When the server in Seville queried for names in the Hong Kong site, the server in Seville first forwarded those queries to Seattle.

After upgrading to Windows Server 2003, administrators configured the DNS server in Seville to forward queries destined for the Hong Kong site directly to a server in that site, instead of first detouring to Seattle, as shown in Figure 3.7.

**Figure 3.7 Conditional Forwarding to an Off-Site Server**




Administrators configured DNS3.corp.fabrikam.com to forward any queries for corp.treyresearch.com to DNS5.corp.treyresearch.com or DNS6.corp.treyresearch.com. DNS3.corp.fabrikam.com forwards all other queries to DNS1.corp.contoso.com or DNS2.corp.contoso.com.

For more information about conditional forwarding in Windows Server 2003 DNS, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

**Using Conditional Forwarding to Query for Names in Other Namespaces**

If your internal network does not have a private root and your users need access to other namespaces, such as a network belonging to a partner company, use conditional forwarding to enable servers to query for names in other namespaces. Conditional forwarding in Windows Server 2003 DNS eliminates the need for secondary zones by configuring DNS servers to forward queries to different servers based on the domain name.

For example, the Contoso Corporation includes two namespaces: Contoso and Trey Research. Computers in each division need access to the other namespace. In addition, computers in both divisions need access to computers in the Supplier private namespace.

Before upgrading to Windows Server 2003, the Trey Research division created secondary zones to ensure that computers in both the Contoso and Trey Research namespace can resolve names in the Contoso, Trey Research, and Supplier namespaces. After upgrading to Windows Server 2003, the Trey Research division deleted its secondary zones and configured conditional forwarding instead.

# Upgrading DNS Servers to Windows Server 2003 DNS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The procedure you need to follow to upgrade DNS servers to Windows Server 2003 depends on whether you want to support Active Directory or not. If you are upgrading to Windows Server 2003 DNS and might not support Active Directory, for information about upgrading your existing DNS servers or migrating third-party DNS servers, see "Migrating servers" in Help and Support Center for Windows Server 2003. Migration involves the following:

- Plan your migration schedule to ensure that your DNS clients have access to a DNS server at all times.
- Back up your existing configuration.
- Migrate data from existing DNS servers to Windows Server 2003 DNS.

If you are upgrading your DNS servers to support Active Directory, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

After you have upgraded or migrated your servers, test them to ensure that they are resolving correctly. For more information about DNS troubleshooting and testing DNS server performance, see "Monitor Servers" in Help and Support Center for Windows Server 2003, and the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Designing DNS Zones

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Each zone type that is available in Windows Server 2003 DNS has a specific purpose. The DNS designer in your organization selects the type of zones to deploy based on the practical purpose of each zone. The DNS administrators in your organization manage and maintain your DNS zones. Figure 3.8 shows the process for designing DNS zones.

**Figure 3.8 Designing DNS Zones**



# Choosing a Zone Type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Design zones to correspond to your network administration infrastructure. If a site in your network is administered locally, deploy a zone for the subdomain. If a department has a subdomain, but no administrator, keep the subdomain in the parent zone. Decide whether or not to store your zones in Active Directory. Active Directory distributes data using a multimaster replication model that provides more security than standard DNS. With the exception of secondary zones, you can store all zone types in Active Directory because all other zones are considered primary zones. When designing DNS zones, host each zone on more than one DNS server.

Decide which type of zone to use, based on your domain structure. For each zone type, with the exception of secondary zones, decide whether to deploy file-based zones or Active Directory–integrated zones.

**Primary Zones**
Deploy primary zones that correspond to your planned DNS domain names. You cannot store both an Active Directory–integrated and a file-based primary copy of the same zone on the same DNS server.

**Secondary Zones**
Add secondary zones if you do not have an Active Directory infrastructure. If you do have an Active Directory infrastructure, use secondary zones on DNS servers that are not serving as domain controllers. A secondary zone contains a complete copy of a zone. Therefore, use secondary zones to improve zone availability at remote sites if you do not want zone data propagated across a WAN link by means of Active Directory replication.

**Stub Zones**
A *stub zone* is a copy of a zone that contains only the original zone's start of authority (SOA) resource record, the name server (NS) resource records listing the authoritative servers for the zone, and the glue address (A) resource records that are needed to identify these authoritative servers.

A DNS server that is hosting a stub zone is configured with the IP address of the authoritative server from which it loads. DNS servers can use stub zones for both iterative and recursive queries. When a DNS server hosting a stub zone receives a recursive query for a computer name in the zone to which the stub zone refers, the DNS server uses the IP address to query the authoritative server, or, if the query is iterative, returns a referral to the DNS servers listed in the stub zone.

Stub zones are updated at regular intervals, determined by the refresh interval of the SOA resource record for the stub zone. When a DNS server loads a stub zone, it queries the zone's primary servers for SOA resource records, NS resource records at the zone's root, and glue address (A) resource records. The DNS server attempts to update its resource records at the end of the SOA resource record's refresh interval. To update its records, the DNS server queries the primary servers for the resource records listed earlier.

You can use stub zones to ensure that the DNS server that is authoritative for a parent zone automatically receives updates about the DNS servers that are authoritative for a child zone. To do this, add the stub zone to the server that is hosting the parent zone. Stub zones can be either file-based or Active Directory–integrated. If you use Active Directory–integrated stub zones, you can configure them on one computer and let Active Directory replication propagate them to other DNS servers running on domain controllers.

Although conditional forwarding is the recommended method for making your servers aware of other namespaces, you can also use stub zones for this. For more information about using stub zones, see Help and Support Center for Windows Server 2003.

### Note

- Only DNS servers running Windows Server 2003 and BIND 9 support stub zones.

**Stub Zones and Conditional Forwarding**
Stub zones and conditional forwarding are Windows Server 2003 DNS features that enable you to control the routing of DNS traffic over a network. These features enable a DNS server to respond to a query by doing one of the following:

- Providing a referral to another DNS server.
- Forwarding the query to another DNS server.

A stub zone enables a DNS server that is hosting a parent zone to be aware of the names and IP addresses of DNS servers that are authoritative for a child zone, even if the DNS server does not have a complete copy of the child zone. In addition, when a stub zone is used, the DNS server does not have to send queries to the DNS root servers. If the stub zone for a child zone is hosted on the same DNS server as the parent zone, the DNS server that is hosting the stub zone receives a list of all new authoritative DNS servers for the child zone when it requests an update from the stub zone's primary server. In this way, the DNS server that is hosting the parent zone maintains a current list of the authoritative DNS servers for the child zone as the authoritative DNS servers are added and removed.

Use conditional forwarding if you want DNS servers in one network to perform name resolution for DNS clients in another network. You can configure DNS servers in separate networks to forward queries to each other without querying DNS servers on the Internet. If DNS servers in separate networks forward DNS client names to each other, the DNS servers cache this information. This enables you to create a direct point of contact between DNS servers in each network and reduces the need for recursion.

If you are using a stub zone and you have a firewall between DNS servers in the networks, then DNS servers on the query/resolution path must have port 53 open. However, if you are using conditional forwarding and you have a firewall between DNS servers in each of the networks, the requirement to have port 53 open only applies to the two DNS servers on either side of the firewall.

**Active Directory–Integrated Zones**
If your DNS topology includes Active Directory, use Active Directory–integrated zones. Active Directory–integrated zones enable you to store zone data in the Active Directory database. Zone information about any primary DNS server within an Active Directory–integrated zone is always replicated.

Because DNS replication is single-master, a primary DNS server in a standard primary DNS zone can be a single point of failure. In an Active Directory–integrated zone, a primary DNS server cannot be a single point of failure because Active Directory uses multimaster replication. Updates that are made to any domain controller are replicated to all domain controllers and the zone information about any primary DNS server within an Active Directory–integrated zone is always replicated. Active Directory–integrated zones:

- Enable you to secure zones by using secure dynamic update.
- Provide increased fault tolerance. Every Active Directory–integrated zone can be replicated to all domain controllers within the Active Directory domain or forest. All DNS servers running on these domain controllers can act as primary servers for the zone and accept dynamic updates.
- Enable replication that propagates changed data only, compresses replicated data, and reduces network traffic.

If you have an Active Directory infrastructure, you can only use Active Directory–integrated zones on Active Directory domain controllers. If you are using Active Directory–integrated zones, you must decide whether or not to store Active Directory–integrated zones in the application directory partition.

You can combine Active Directory–integrated zones and file-based zones in the same design. For example, if the DNS server that is authoritative for the private root zone is running on an operating system other than Windows Server 2003 or Windows 2000, it cannot act as an Active Directory domain controller. Therefore, you must use file-based zones on that server. However, you can delegate this zone to any domain controller running either Windows Server 2003 or Windows 2000.

**Storing Active Directory–Integrated Zones in Application Directory Partitions**
Windows Server 2003 Active Directory enables you to configure an application directory partition that limits the scope of replication. Data stored in an application directory partition is replicated to a subset of domain controllers. This subset is determined by the replication scope of the data. In the default configuration of Windows Server 2003 Active Directory, DNS application directory partitions are present only on the domain controllers that run the DNS Server service. By storing Active Directory–integrated zones in an application directory partition, you can reduce the number of objects that are stored in the global catalog, and you can reduce the amount of replication traffic within a domain.

In contrast, Active Directory–integrated zones that are stored in domain directory partitions are replicated to all domain controllers in the domain. Storing Active Directory–integrated zones in an application directory partition allows replication of DNS data to domain controllers anywhere in the same Active Directory forest.

When you are setting up your Active Directory environment and installing the first Windows Server 2003 domain controller in the forest, if you install DNS, two Windows Server 2003 DNS application directory partitions are created by default. A forest-

wide DNS application directory partition called **ForestDNSZones** will be created, and for each domain in the forest, a domain-wide DNS application directory partition called **DomainDNS Zones** will be created.

# Choosing a Propagation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you decide which zone each DNS server hosts, decide how to propagate the zones among the servers. Propagated zones provide higher availability, improve query response time, and reduce network traffic produced by name queries. However, propagated zones require storage space and increase network traffic. If your network is distributed and managed at different sites, use subdomains for these sites. If you do not have a distributed network, avoid using subdomains when possible.

In Windows Server 2003, zones are propagated by means of file-based zone transfer or Active Directory replication. If you use file-based zones, file-based zone transfer is the method of propagation. If you have Windows Server 2003 and Windows 2000 Active Directory–integrated zones, use Active Directory replication.

**File-Based Zone Transfer**
Windows Server 2003 and Windows 2000 DNS support both incremental and full zone transfer of file-based zones. Incremental zone transfer is the default method, but if this method is not supported by a third-party DNS server that is involved in the transfer, DNS servers running Windows Server 2003 and Windows 2000 transfer the full zone.

Incremental zone transfer, described in RFC 1995:*Incremental Zone Transfer in DNS*, provides better use of available network bandwidth. Rather than sending the entire contents of the zone file, the primary server only transfers the incremental changes in the zone. This reduces the impact of DNS zone transfers on network traffic. Without incremental zone transfers, the primary server transfers the entire zone file to the secondary server every time a DNS zone is updated.

Windows Server 2003 DNS uses full zone transfer when zones must be transferred to DNS servers that do not support incremental zone transfer, such as DNS servers running on Windows NT 4.0 or earlier versions of BIND 8.

**Active Directory Replication**
Active Directory replication propagates zone changes between domain controllers. Replication processing differs from DNS full zone transfers, in which the DNS server transfers the entire zone. Replication processing also differs from incremental zone transfers, in which the server transfers all changes made since the last change.

Active Directory zone replication provides the following additional benefits:

- Network traffic is reduced because the domain controllers only send the final result of all changes.
- When a zone is stored in Active Directory, replication occurs automatically. No additional configuration is required.
- When Active Directory zone replication occurs between sites, zone data that is greater than the default transfer size is automatically compressed before it is transferred. This compression decreases the network traffic load.

After careful analysis, you can partition and delegate your DNS zones based on what is required for providing efficient and fault-tolerant name service to each location or site.

If you are using Active Directory–integrated zones in a Windows Server 2003 domain, you must select an Active Directory–integrated zone replication scope. When selecting a replication scope, note that network traffic increases as you broaden the replication scope. For example, if you choose to replicate Active Directory–integrated DNS zone data to all DNS servers in the forest, this produces greater network traffic than replicating the DNS zone data to all DNS servers in a single Active Directory domain in that forest. Balance your need to minimize replication traffic against your need to minimize zone query traffic. The DNS administrators in your organization are responsible for managing zone replication.

Table 3.8 lists the replication options for Active Directory–integrated zone data.

**Table 3.8 Replication Options for Active Directory–Integrated Zone Data**

| Option | Description | When to Use |
|---|---|---|

| | | |
|---|---|---|
| All DNS servers in the Active Directory forest | The zone data replicates to all the DNS servers running on Windows Server 2003–based domain controllers in all domains of the Active Directory forest. | You want the broadest scope of replication. This option generally produces the most zone replication traffic. Note that you can choose this option only if all DNS servers hosting an Active Directory–integrated copy of this zone run Windows Server 2003. |
| All DNS servers in a specified Active Directory domain | The zone data replicates to all DNS servers running on Windows Server 2003–based domain controllers in the specified Active Directory domain. This option is the default setting for Active Directory–integrated DNS zone replication.<br><br>(The specified Active Directory domain is the domain hosted by the domain controller on which the DNS server hosting the zone is running.) | You do not need the zone to be replicated throughout the forest and you want to limit zone replication traffic. This option produces less zone replication traffic than replicating the zone to all DNS servers in the forest or to all domain controllers in the domain. If you choose this option, the zone data does not replicate to DNS servers running on Windows 2000–based domain controllers. |
| All domain controllers in the Active Directory domain | The zone data replicates to all domain controllers in the specified Active Directory domain, whether or not the DNS Server service runs on the domain controllers in the domain. | You host an Active Directory–integrated copy of this zone on a DNS server running on a Windows 2000–based domain controller. |
| All domain controllers specified in the replication scope of a DNS application directory partition | The zone data replicates to all the domain controllers specified in the replication scope of the DNS application directory partition. | You want to customize zone replication scope for your organization. With this option, you can minimize zone replication traffic while maximizing functionality. However, this option requires more administrative overhead. You can choose this option only if all DNS servers hosting an Active Directory–integrated copy of this zone run Windows Server 2003. |

# Migrating Zones to Windows Server 2003 DNS Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can migrate zones to DNS servers running Windows Server 2003 in one of two ways:

- By using zone transfer.
- By copying the zone files.

If you copy the zone files, you must manually verify the integrity of the zones. Regardless of the method that you use to migrate zones, you must decide whether to take the original DNS server offline, or to use it as a secondary server. If you determine that the original third-party DNS server causes interoperability problems on your network, or if you need to use that server hardware for another purpose, take the server offline. Otherwise, keep the server on you network to provide backup for your primary DNS server running Windows Server 2003.

For more information about using zone transfer, see "Initiate a zone transfer at a secondary server" in Help and Support Center for Windows Server 2003.

# Configuring and Managing DNS Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you configure DNS clients, you must specify a list of DNS servers for clients to use when resolving DNS names. You can also specify a DNS suffix search list to be used by the clients when performing DNS query searches for short, unqualified domain names.

Figure 3.9 shows the process for configuring and managing DNS clients.

**Figure 3.9 Configuring and Managing DNS Clients**



# Configuring Client DNS Server Lists and Suffix Search Lists

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure your clients' DNS server lists and suffix search list by including at least two DNS server IP addresses on the clients and domain controllers: the IP address for a preferred server and the IP address for an alternate server. Use a server running in the local site for the preferred server. The alternate server can be running in either a local or a remote site.

The DNS suffix search list is populated based on the primary DNS suffix of the client and any connection-specific DNS suffixes. The client uses these suffixes to try to resolve unqualified names. You can modify the DNS suffix search list manually, or by using Group Policy. Limit the size of your suffix search list if you can, because a large suffix search list increases network traffic.

# Using Group Policy to Simplify Client Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 includes a new set of Group Policy settings to simplify the rollout of Windows Server 2003 DNS clients. You can use them to set your suffix search lists, dynamic update configuration, and many other settings. As with all Group Policy settings, you can specify different settings based on site, domain, or OU.

For more information about these Group Policy settings, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Securing Your DNS Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because DNS was designed to be an open protocol, DNS data can be vulnerable to security attacks. Windows Server 2003 DNS provides improved security features to decrease this security issue. The DNS designer in your organization is responsible for creating a secure DNS infrastructure. The DNS administrators in your organization are responsible for maintaining network security by anticipating and mitigating new security threats.

Figure 3.10 shows the process for securing your DNS infrastructure.

**Figure 3.10 Securing Your DNS Infrastructure**



# Identifying DNS Security Threats

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A DNS infrastructure is vulnerable to a number of types of security threats.

**Footprinting** The process of building a diagram, or footprint, of a DNS infrastructure by capturing DNS zone data such as domain names, computer names, and IP addresses for sensitive network resources. DNS domain and computer names often indicate the function or location of domains and computers.

**Denial-of-service attack** An attack in which the attacker attempts to deny the availability of network services by flooding one or more DNS servers in the network with recursive queries. When a DNS server is flooded with queries, its CPU usage eventually reaches its maximum and the DNS Server service becomes unavailable. Without a fully operating DNS server on the network, network services that use DNS are unavailable to network users.

**Data modification** The use of valid IP addresses in IP packets that an attacker has created to destroy data or conduct other attacks. Data modification is typically attempted on a DNS infrastructure that has already been foot printed. If the attack is successful, the packets appear to be coming from a valid IP address on the network. This is commonly called IP spoofing. With a valid IP address (an IP address within the IP address range of a subnet), an attacker can gain access to the network.

**Redirection** An attack in which an attacker is able to redirect queries for DNS names to servers that are under the control of the attacker. One method of redirection involves the attempt to pollute the DNS cache of a DNS server with erroneous DNS data that might direct future queries to servers that are under the control of an attacker. For example, if a query is made to example.contoso.com and a referral answer provides a record for a name that is outside of the contoso.com domain, the DNS server uses the cached data to resolve a query for the external name. Redirection can be accomplished when an attacker has writable access to DNS data, such as with non-secure dynamic updates.

For more information about common types of attacks, developing a security policy, and evaluating your level of risk, see "Designing an Authentication Strategy" and "Designing a Resource Authorization Strategy" in *Designing and Deploying Directory and Security Services* of this kit.

# Developing a DNS Security Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your DNS data is compromised, attackers can gain information about your network that can be used to compromise other services. For example, attackers can harm your organization in the following ways:

- By using zone transfer, attackers can retrieve a list of all the hosts and their IP addresses in your network.
- By using denial-of-service attacks, attackers can prevent e-mail from being delivered to and from your network, and they can prevent your Web server from being visible.
- If attackers can change your zone data, they can set up fake Web servers, or cause e-mail to be redirected to their servers.

Your risk of attack varies depending on your exposure to the Internet. For a DNS server in a private network that uses a private namespace, a private addressing scheme, and an effective firewall, the risk of attack is lower and the possibility of discovering the intruder is greater. For a DNS server that is exposed to the Internet, the risk is higher.

Developing a DNS security policy involves:

- Deciding what access your clients need, what tradeoffs you want to make between security and performance, and what data you most want to protect.
- Familiarizing yourself with the security issues common to internal and external DNS servers.
- Studying your name resolution traffic to see which clients can query which servers.

You can choose to adopt a low-level, mid-level, or high-level DNS security policy.

**Low-Level DNS Security Policy**

Low-level security does not require any additional configuration of your DNS deployment. Use this level of DNS security in a network environment in which you are not concerned about the integrity of your DNS data, or in a private network in which no external connectivity is possible. A low-level security policy includes the following characteristics:

- All DNS servers in your network perform standard DNS resolution.
- All DNS servers are configured with root hints that point to the root servers for the Internet.
- All DNS servers permit zone transfers to any server.
- All DNS servers are configured to listen on all of their IP addresses.
- **Secure cache against pollution** is disabled on all DNS servers.
- Dynamic update is allowed for all DNS zones.
- User Datagram Protocol (UDP) and TCP/IP port 53 is open on the firewall for your network for both source and destination addresses.

**Mid-Level DNS Security Policy**

Mid-level DNS security consists of the DNS security features that are available without running DNS servers on domain controllers and storing DNS zones in Active Directory. A mid-level security policy includes the following characteristics:

- The DNS infrastructure of your organization has limited exposure to the Internet.

- All DNS servers are configured to use forwarders to point to a specific list of internal DNS servers when they cannot resolve names locally.
- All DNS servers limit zone transfers to servers listed in the NS records in their zones.
- DNS servers are configured to listen on specified IP addresses.
- **Secure cache against pollution** is enabled on all DNS servers.
- Secure dynamic update is allowed for all DNS zones.
- Internal DNS servers communicate with external DNS servers through the firewall with a limited list of allowed source and destination addresses.
- External DNS servers in front of your firewall are configured with root hints pointing to the root servers for the Internet.
- All Internet name resolution is performed by using proxy servers and gateways.

**High-Level DNS Security Policy**

High-level DNS security uses the same configuration as mid-level security and also uses the security features available when the DNS Server service is running on a domain controller and DNS zones are stored in Active Directory. Also, high-level security completely eliminates DNS communication with the Internet. This is not a typical configuration, but it is recommended whenever Internet connectivity is not required. High-level security policy includes the following characteristics:

- The DNS infrastructure of your organization has no Internet communication by means of internal DNS servers.
- Your network uses an internal DNS root and namespace, where all authority for DNS zones is internal.
- DNS servers that are configured with forwarders use internal DNS server IP addresses only.
- All DNS servers limit zone transfers to specified IP addresses.
- DNS servers are configured to listen on specified IP addresses.
- **Secure cache against pollution** is enabled on all DNS servers.
- Internal DNS servers are configured with root hints that point to the internal DNS servers hosting the root zone for your internal namespace.
- Secure dynamic update is configured for all DNS zones except for the top-level and root zones, which do not allow dynamic updates at all.
- All DNS servers are running on domain controllers. An access control list (ACL) is configured on the DNS Server service to allow only specific individuals to perform administrative tasks on DNS servers.
- All DNS zones are stored in Active Directory. An ACL is configured to allow only specific individuals to create, delete, or modify DNS zones.
- ACLs are configured on DNS resource records to allow only specific individuals to create, delete, or modify DNS data.

**Note**

- Windows Server 2003 DNS does not support the use of DACLs on zones to control which clients or users can send queries to the DNS server.

**Cache Pollution Protection**

When cache pollution protection is enabled, the DNS server disregards DNS resource records that originate from DNS servers that are not authoritative for the resource records. Cache pollution protection is a significant security enhancement; however, when cache pollution protection is enabled, the number of DNS queries can increase.

In Windows Server 2003 DNS, cache pollution protection is enabled by default. You can disable cache pollution protection to reduce the number of DNS queries; however, to ensure the security of your system, it is strongly recommended that you leave cache pollution protection enabled on your DNS servers.

For information about cache pollution protection, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Securing DNS Servers That Are Exposed to the Internet

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DNS servers that are exposed to the Internet are especially vulnerable to attack. You can secure your DNS servers that are exposed to the Internet by doing the following:

- Place the DNS server on a perimeter network instead of your internal network.

  For more information about perimeter networks, see "Deploying ISA Server" in this book.

  Use one DNS server for publicly accessed services inside your perimeter network and a separate DNS server for your private internal network. This reduces the risk of exposing your private namespace, which can expose sensitive names and IP addresses to Internet-based users. It also increases performance because it decreases the number of resource records on the DNS server.

- Add a secondary server on another subnet or network, or on an ISP. This protects you against denial-of-service attacks.

- Eliminate single points of failure by securing your routers and DNS servers, and distributing your DNS servers geographically. Add secondary copies of your zones to at least one offsite DNS server.

- Encrypt zone replication traffic by using Internet Protocol security (IPSec) or virtual private network (VPN) tunnels to hide the names and IP addresses from Internet-based users.

- Configure firewalls to enforce packet filtering for UDP and TCP port 53.

- Restrict the list of DNS servers that are allowed to initiate a zone transfer on the DNS server. Do this for each zone in your network.

- Monitor the DNS logs and monitor your external DNS servers by using Event Viewer.

# Securing Internal DNS Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Internal DNS servers are less vulnerable to attack than external DNS servers, but you still need to protect them. To secure your internal DNS servers:

- Eliminate any single point of failure. Note, however, that DNS redundancy cannot help you if your clients cannot access any network services. Think about where the clients of each DNS zone are located, and how they resolve names if the DNS server is compromised and unable to answer queries.
- Prevent unauthorized access to your servers. Allow only secure dynamic update for your zones and limit the list of DNS servers that are allowed to obtain a zone transfer.
- Monitor the DNS logs and monitor your internal DNS servers by using Event Viewer. Monitoring your logs and your server can help you detect unauthorized modifications to your DNS server or zone files.
- Implement Active Directory–integrated zones with secure dynamic update.

# Securing Dynamically Updated DNS Zones

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use Active Directory–integrated zones and configure them for secure dynamic update. Secure dynamic update resolves the security risks associated with using dynamic update. Because dynamic update allows any computer to modify any record, an attacker can modify zone data, then impersonate existing servers.

For example, if you install the Web server, web.contoso.com, and it registers its IP address in DNS by using dynamic update, an attacker can install a second Web server, also name it web.contoso.com, and use dynamic update to modify the corresponding IP address in the DNS record. In this way, the attacker can impersonate the original Web server and capture secure information.

To prevent server impersonation, implement secure dynamic update. By using secure dynamic update, only the computers and users specified in an access control list (ACL) can modify objects within a zone.

If your security policy demands stricter security, modify these settings to further restrict access. Restrict access by computer, group, or user account, and assign permissions for the entire DNS zone and for the individual DNS names within the zone.

For more information about securing dynamically updated DNS zones, see the *Networking Guide* of the *Windows Server 2003 Resource Kit* (or see the Networking Guide on the Web at http://www.microsoft.com/reskit).

# Securing DNS Zone Replication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Zone replication can occur either by means of zone transfer or as part of Active Directory replication. If you do not secure zone replication, you run the risk of exposing the names and IP addresses of your computers to attackers. You can secure DNS zone replication by doing the following:

- Using Active Directory replication.
- Encrypting zone replication sent over public networks such as the Internet.
- Restricting zone transfer to authorized servers.

**Using Active Directory Replication**

Replicating zones as part of Active Directory replication provides the following security benefits:

- Active Directory replication traffic is encrypted; therefore zone replication traffic is encrypted automatically.
- The Active Directory domain controllers that perform replication are mutually authenticated, and impersonation is not possible.

**Note**

- Use Active Directory–integrated zones whenever possible, because they are replicated as part of Active Directory replication, which is more secure than file-based zone transfer.

**Encrypting Replication Traffic Sent Over Public Networks**

Encrypt all replication traffic sent over public networks by using IPSec or VPN tunnels. When encrypting replication traffic sent over public networks:

- Use the strongest level of encryption or VPN tunnel authentication that your servers can support.
- Use the Windows Server 2003 Routing and Remote Access service to create the IPSec or VPN tunnel.

**Restricting Zone Transfer to Authorized Servers**

If you have secondary servers and you replicate your zone data by using zone transfer, configure your DNS servers to specify the secondary servers that are authorized to receive zone transfers. This prevents an attacker from using zone transfer to download zone data. If you are using Active Directory–integrated zones instead, configure your servers to disallow zone transfer.

# Integrating DNS with Other Windows Server 2003 Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you deploy Windows Server 2003 DNS, it is important to integrate the DNS service with other Windows Server 2003 services, such as DHCP and WINS. DNS administrators are responsible for integrating DNS with WINS and DHCP. Figure 3.11 shows the process for integrating Windows Server 2003 DNS with other Windows Server 2003 services.

**Figure 3.11 Integrating DNS with Other Windows Server 2003 Services**



# Integrating DNS with DHCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 DNS supports DHCP by means of the dynamic update of DNS zones. By integrating DHCP and DNS in a DNS deployment, you can provide your network resources with dynamic addressing information stored in DNS. To enable this integration, you can use the Windows Server 2003 DHCP service.

The dynamic update standard, specified in RFC 2136: *Dynamic Updates in the Domain Name System (DNS UPDATE)*, automatically updates DNS records. Both Windows Server 2003 and Windows 2000 support dynamic update, and both clients and DHCP servers can send dynamic updates when their IP addresses change.

Dynamic update enables a DHCP server to register address (A) and pointer (PTR) resource records on behalf of a DHCP client by using DHCP Client FQDN option 81. Option 81 enables the DHCP client to provide its FQDN to the DHCP server. The DHCP client also provides instructions to the DHCP server describing how to process DNS dynamic updates on behalf of the DHCP client.

The DHCP server can dynamically update DNS A and PTR records on behalf of DHCP clients that are not capable of sending option 81 to the DHCP server. You can also configure the DHCP server to discard client A and PTR records when the DHCP client lease is deleted. This reduces the time needed to manage these records manually and provides support for DHCP clients that cannot perform dynamic updates. In addition, dynamic update simplifies the setup of Active Directory by enabling domain controllers to dynamically register SRV resource records.

If the DHCP server is configured to perform DNS dynamic updates, it performs one of the following actions:

- The DHCP server updates resource records at the request of the client. The client requests the DHCP server to update the DNS PTR record on behalf of the client, and the client registers A.

- The DHCP server updates DNS A and PTR records regardless of whether the client requests this action or not.

By itself, dynamic update is not secure because any client can modify DNS records. To secure dynamic updates, you can use the secure dynamic update feature provided in Windows Server 2003. To delete outdated records, you can use the DNS server aging and scavenging feature.

# Integrating DNS with WINS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you upgrade to Windows Server 2003 DNS from an earlier version of Windows, you might need to continue support an existing WINS infrastructure. Windows Server 2003 DNS enables you to support an existing WINS deployment by allowing you to configure a DNS server to query a WINS server as a DNS zone setting.

WINS provides dynamic NetBIOS name resolution. If your organization supports clients and applications that use WINS for NetBIOS name resolution, you need to continue to support WINS. If some of your clients are registered in WINS and other clients need to resolve their names but are not capable of NetBIOS name resolution, you can use WINS lookup to enable your DNS server to look up names in the WINS namespace.

This feature is particularly useful if some of your clients that require NetBIOS name resolution cannot use WINS or if some of your clients cannot register with DNS (for example, clients that run the Microsoft® Windows® 95 or Windows® 98 operating system). Use WINS referral if some of your DNS servers do not support the resource records used for WINS lookup and WINS reverse lookup.

**WINS Lookup and WINS Reverse Lookup**

By configuring your DNS server for WINS lookup, you can direct DNS to query WINS for name resolution, so that DNS clients can look up the names and IP addresses of WINS clients. To direct DNS to query WINS for name resolution, add a WINS lookup record to the authoritative zone. An authoritative DNS server checks that zone when it receives a query for a name. If the DNS server does not find the name in the authoritative zone, but the zone contains a WINS lookup record, the DNS server queries the WINS server. If the name is registered with WINS, the WINS server returns the associated record to the DNS server.

The DNS server then compiles and returns the corresponding DNS record in response to the original DNS request. DNS clients do not need to know whether a client is registered with WINS or DNS, and they do not need to query the WINS server.

You can also configure your DNS server for WINS reverse lookups. Reverse lookups work slightly differently. When an authoritative DNS server is queried for a nonexistent PTR record, and the authoritative zone contains the WINS-R record, the DNS server uses a NetBIOS node adapter status lookup.

**Note**

- For fault tolerance, you can specify multiple WINS servers in the WINS lookup record. The server that is running the Windows 2000 or Windows Server 2003 DNS Server service tries to locate the name by searching the WINS servers in the order specified by the list.

**Configuring WINS Referral**

Computers that are running third-party implementations of DNS do not support the records used for WINS lookup and WINS reverse lookup. If you attempt to use a combination of Microsoft and third-party DNS servers to host a zone containing these records, the mixture can cause data errors or failed zone transfers at the third-party DNS servers.

If you have such a combination, you can use WINS referral to create and delegate a special WINS zone that refers DNS lookups to WINS. This zone does not perform any registrations or updates. Next, you configure all DNS clients to append the WINS referral zone name to unqualified queries. That way, the client can query both DNS and WINS at the same time, using a DNS query. To simplify administration, you can use DHCP or Group Policy to configure the clients to perform the configuration. Deploying this configuration overrides the default DNS client resolver behavior, requiring you to finish populating the suffix search order with combinations of suffixes, such as the primary DNS suffix, the primary DNS suffix devolved, and connection specific suffixes.

For more information about DHCP, see "Deploying DHCP" in this book. For more information about Group Policy, see "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit.

**Note**

- The WINS zone must be hosted on a DNS server that is running Windows Server 2003 or Windows 2000 and must not be propagated to third-party DNS servers. Third-party DNS servers do not support WINS resource records and might not be able to host the zone.

# Implementing Windows Server 2003 DNS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have tested your configuration in a pilot lab, you can implement your changes in your production environment. Figure 3.12 shows the process for implementing Windows Server 2003 DNS.

**Figure 3.12 Implementing Windows Server 2003 DNS**



# Preparing to Deploy Windows Server 2003 DNS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Prepare your environment for Windows Server 2003 DNS deployment by ensuring that you have reliable backups of anything that you plan to change, including servers and zones. Test your backups before you proceed with your deployment. In addition, create a recovery plan for contingencies such as data loss, server failure, and failure of network connections.

Before implementing your DNS deployment, ensure that routing links between the servers that you plan to deploy are in place and are working correctly. Depending on how your DNS infrastructure is configured, your DNS servers might need to query the following:

- Root servers.
- Forwarders.
- Servers hosting parent zones.
- Servers hosting child zones.
- Servers on an external network or the Internet.

If you expect clients to query for names on the Internet and you plan to use a proxy server, ensure that the proxy server is in place and that a proxy-client/firewall-client is installed on the client. In addition, ensure that the Web client configuration is set in a Conseil Européen pour la Recherche Nucléaire (CERN)–compliant Internet browser. Microsoft Internet Explorer is an example of a CERN–compliant Internet browser.

# Setting up the DNS Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you install DNS, ensure that your computer is named correctly and that you can ping other computers in the network that your DNS servers might need to query. Because clients locate DNS servers by IP address, assign a static IP address to each DNS server.

You can set up your DNS server in one of four ways:

- Install DNS on a server by using the Active Directory Installation Wizard to install Active Directory.

  The Active Directory Installation Wizard automatically creates an Active Directory–integrated copy of the forward lookup zone that corresponds to the name of the Active Directory domain, and configures the zone for secure dynamic update. In addition, the wizard creates the standard reverse lookup zones recommended by the DNS RFCs.

  You can start the Active Directory Installation Wizard by running **dcpromo.exe** at a command prompt. For more information about Active Directory installation and removal, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

- Before or after you install Active Directory on the server, you can use the Add or Remove Programs tool to install the DNS Server service and then run the Configure DNS Server Wizard to configure your zones. As with the Active Directory Installation Wizard, the Configure DNS Server Wizard creates the standard reverse lookup zones recommended by the DNS RFCs, and either configures the server as a root server or initializes the root hints.

- You can use the command-line tool Dnscmd.exe to configure the DNS server.

- You can use Microsoft® Visual Basic® Scripting Edition (VBScript) or other scripting languages through the Windows Management Instrumentation (WMI) provider packaged with Windows Server 2003.

For more information about these setup options and for information about Windows Server 2003 DNS, including how the Active Directory Installation Wizard and the Configure DNS Server Wizard determine whether or not to initialize the root hints, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at www.microsoft.com/reskit).

# Setting up Zones

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you install DNS by using the Active Directory Installation Wizard, the wizard creates DNS zones that correspond to the Active Directory domains that you specify. If the zones that you specified during the zone planning phase of your deployment do not already exist, create them now. Note that the default DNS installation by the Active Directory Installation Wizard includes secure dynamic update and an Active Directory–integrated zone. If this is not the configuration you want to deploy, change the default settings.

If the zone that the wizard creates is not the type of zone that you want, change it now.

**Note**

- Converting any zone to an Active Directory–integrated zone can increase the use of DNS server resources and network resources. This is because converting a zone can trigger Active Directory replication.

If you want to push updates to secondary DNS servers for a zone, configure **DNS notify** at the primary DNS server.

For more information about how to add and remove zones, see Help and Support Center for Windows Server 2003.

# Configuring Forwarding

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If any of your servers need to forward queries to any other server, configure forwarding on the servers that must forward queries. If you want your server to forward queries to different servers depending on the DNS suffix specified in the query, configure conditional forwarding appropriately.

For more information about conditional forwarding, see "Using Forwarding" earlier in this chapter, and see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Configuring the DNS Server for Dynamic Update

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Depending on how you deploy your server and zones, your zones might already be configured for dynamic update or secure dynamic update. If the zones are not configured as you intend, make any necessary changes. You can configure your zones to perform dynamic updates, secure dynamic updates, or no updates. However, configuring your zones to perform unsecured dynamic updates is a security risk and is not recommended.

For information about how to configure dynamic updates, see "Dynamic update" in Help and Support Center for Windows Server 2003. For information about how to allow only secure dynamic updates, see "Allow only secure dynamic updates" in Help and Support Center for Windows Server 2003.

# Configuring Aging and Scavenging

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With dynamic update, whenever a computer joins the network and registers with DHCP, the DNS server automatically adds resource records to the zone. However, in some cases, the DNS server does not automatically delete them, and they can become outdated. Outdated resource records use disk space on the server, and a server might use an outdated resource record to answer a query. As a result, DNS server performance suffers. To solve these problems, the Windows Server 2003 DNS Server service can scavenge outdated records by searching the database for records that are obsolete and deleting them.

You can configure aging and scavenging from DNS Manager or by using **Dnscmd.exe**.

For more information about configuring aging and scavenging in Windows Server 2003 DNS, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Configuring Zone Transfer and Replication Scope

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you use file-based primary or secondary zones, configure each primary server for a zone to allow zone transfers to each secondary server for a zone. Next, configure each secondary server for a zone with a list of primary servers for that zone.

If you use Active Directory replication in a Windows Server 2003 domain, configure the zone replication scope as described in Table 3.8.

If the DNS server is a domain controller, you can change the zone type to Active Directory–integrated. However, if the DNS server is not a domain controller, this option is not available. Active Directory–integrated zone data is stored and replicated as part of the Active Directory database.

For more information about DNS zone storage and replication in Active Directory, enlisting a DNS server in a DNS application directory partition, removing a DNS server from a DNS application directory partition, or changing zone replication scope, click the **Index** button in Help and Support for Windows Server 2003, and then in the keyword box, type **DNS zones**.

# Verifying that the DNS Server is Operating Correctly

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you install and configure the DNS server, verify that it is operating correctly. Use the monitoring features of the DNS MMC snap-in, such as simple or recursive query testing. You can also examine the event log or use the DNSLint Windows Server 2003 command-line support tool to test DNS servers for problems with delegations and missing Active Directory replication records. In addition, you can use the Nslookup.exe command-line tool to attempt queries. To access the monitoring features of the DNS MMC snap-in, click **Properties** on the **Action** menu, and then click the **Monitoring** tab.

For more information about DNS troubleshooting and verifying DNS server operation, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Setting Up DNS Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Set up each DNS client with the following:

- Host name and DNS suffix.
- Preferred and alternate DNS servers.
- Optionally, Proxy Auto-Configuration file or name exclusion list (if a proxy client).

You can use any of the following methods to set up your DNS clients:

- Use the TCP/IP settings of the client.
- Use Group Policy to configure groups of clients.
- Use the DHCP Server service to configure some client settings automatically.

For more information about how to install and configure DNS clients, see "Configuring DNS client settings" in Help and Support Center for Windows Server 2003.

# Using Command-Line Tools to Deploy DNS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the command-line tool Dnscmd.exe to perform most of the tasks that you can perform from the DNS MMC snap-in. By using Dnscmd.exe, you can create, delete, and view zones and records and reset server and zone properties. You can also perform routine administration operations, such as creating or updating a zone, reloading the zone, refreshing the zone, writing the zone back to a file or Active Directory, pausing and resuming the zone, clearing the cache, adding records to root hints, stopping and starting the DNS service, and viewing statistics.

You can also use Dnscmd.exe to write batch file scripts and for remote administration. For more information about Dnscmd.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click **Windows Support Tools**. For information about installing and using the Windows Server 2003 Support Tools and Support Tools Help, see the Sreadme.doc file in the \Support\Tools folder on the Windows Server 2003 operating system CD.

You can use the Nslookup.exe command-line tool to perform query testing of the DNS domain namespace and to display configuration information.

# Additional Resources for Deploying DNS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing a Resource Authorization Strategy" in *Designing and Deploying Directory and Security Services* of this kit for information about establishing security policies.
- "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit for information about how to deploy DNS specifically for Active Directory.
- "Designing Security Policy" in *Designing a Managed Environment* of this kit for more information about security policies.
- "Designing an Authentication Strategy" in *Designing and Deploying Directory and Security Services* of this kit.
- "Deploying ISA Server" in this book for more information about perimeter networks.
- "Deploying DHCP" in this book.
- "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about the DNS Server service and DNS troubleshooting.
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information, about Active Directory installation and removal.
- RFC 1035: *Domain Names — Implementation and Specification*.
- *DNS and BIND*, 4th ed., by Paul Albitz and Cricket Liu, 2001, Sebastopol, CA: O'Reilly & Associates for more information about DNS.
- *Windows 2000 TCP/IP Protocols and Services*, by Thomas Lee and Joseph Davies, 2000, Redmond, Washington: Microsoft Press for more information about the DNS wire protocol.
- The Internet Engineering Task Force (IETF) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about Request for Comments (RFC) documents and IETF Internet-Drafts.

**Related Tools**

For information about installing and using the Windows Server 2003 Support Tools and Support Tools Help, see the file Sreadme.doc in the \Support\Tools folder of the Windows Server 2003 operating system CD.

- **Dnscmd.exe**

  You can use the Dnscmd.exe command-line tool to perform most of the tasks that you can perform from the DNS MMC snap-in.

- **DNSLint**

  DNSLint is a command-line tool that you can use to address some common DNS name resolution issues, such as lame

delegation, DNS record verification, and verifying DNS records that are used for Active Directory replication.

- **Netdiag.exe**

  Netdiag.exe helps you to isolate networking and connectivity problems by performing a series of tests to determine the state of your network client and whether it is functional.

- **Nslookup.exe**

  You can use the Nslookup.exe command-line tool to submit DNS queries and display the results of the queries.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Migrating servers" in Help and Support Center for Windows Server 2003 for information about upgrading your existing DNS servers or migrating third-party DNS servers.

- "Monitor Servers" in Help and Support Center for Windows Server 2003 for more information about testing DNS server performance.

- "Initiate a zone transfer at a secondary server" in Help and Support Center for Windows Server 2003 for more information about using zone transfer.

- "Dynamic update" in Help and Support Center for Windows Server 2003 for information about how to configure dynamic updates.

- "Allow only secure dynamic updates" in Help and Support Center for Windows Server 2003 for information about how to allow only secure dynamic updates.

- "Configuring DNS client settings" in Help and Support Center for Windows Server 2003 for more information about how to install and configure DNS clients.

# Deploying WINS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Internet Name Service (WINS) in the Microsoft® Windows® Server 2003 operating system allows large organizations to accomplish NetBIOS name resolution with high availability, security, and performance. The following sections describe the WINS deployment process, including how to design and customize a secure replication strategy. WINS migration information and examples are also provided.

**In This Chapter**


Overview of WINS Deployment

Building Your WINS Server Strategy

Designing Your WINS Replication Strategy

Securing Your WINS Solution

Integrating WINS with Other Services

Implementing Your WINS Solution

Additional Resources for Deploying WINS


**Related Information**

- For more information about Windows Internet Name Service (WINS), see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

- For more information about planning and designing your Domain Name System (DNS) network, see "Deploying Domain Name System (DNS)" in this book.

# Overview of WINS Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

WINS provides a dynamic solution for network basic input/output system (NetBIOS) name resolution in enterprise networks. Although most large networks currently have a WINS infrastructure, some still rely on other methods of NetBIOS name resolution, such as the Lmhosts file. If your organization does not currently use WINS, and intends to continue operating with Microsoft® Windows® 95, Windows 98, Windows Millennium Edition, or Microsoft® Windows NT® version 4.0, consider implementing WINS when you deploy Windows Server 2003 in order to automate NetBIOS name resolution. Certain applications, such as Microsoft® Exchange Server, also rely on NetBIOS name resolution. Therefore, even if all of your computers are running Microsoft® Windows® 2000, Windows XP, or Windows Server 2003, you might still require NetBIOS name resolution based on the applications running in your environment.

If you are upgrading your current WINS servers to Windows Server 2003, determine if your existing hardware is compatible with Windows Server 2003, then migrate your WINS solution to Windows Server 2003.

By deploying WINS, you provide NetBIOS name resolution for clients on your network. WINS implements a distributed database for NetBIOS names and their corresponding addresses. WINS clients register their names at a local WINS server, and the WINS servers replicate the entries to the other WINS servers. This ensures the uniqueness of NetBIOS names and makes local name resolution possible.

# WINS Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying WINS involves building a server strategy, designing a replication strategy, securing your WINS solution, integrating WINS with other services, and implementing your WINS solution. Figure 4.1 shows the general WINS deployment process.

**Figure 4.1 Deploying WINS**

# WINS Technology Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Smaller, non-routed networks can be configured as broadcast nodes, also known as B-nodes, accomplishing NetBIOS name registration and resolution by using broadcast packets. A non-WINS solution is viable where the broadcast domain is small and the resulting broadcast traffic is low. However, the traffic generated by broadcasts can overload a large network. In addition, some routers do not allow broadcast messages to pass through, so this method of name resolution is not an option for most enterprise networks. Although you can also use the static Lmhosts file for NetBIOS name resolution, manually editing the file with each name or IP address change can be time-consuming and prone to administrative error. Also, it is not a viable solution in a Dynamic Host Configuration Protocol (DHCP) environment. These more complex environments require a non-broadcast-based solution, which WINS provides by using unicast NetBIOS name registration and resolution.

WINS client support allows you to specify up to 12 WINS servers for redundancy. Different configurations, or node types, are available through WINS. The node type determines the method or methods that are used for NetBIOS name resolution. WINS supports the following node types, as shown in Table 4.1.

**Table 4.1 NetBIOS Node Types**

| Node Type | Resolution Method |
|---|---|
| B-node | IP broadcast messages register and resolve NetBIOS names to IP addresses. Windows 2000–based and Microsoft® Windows® XP–based computers use modified B-node name resolution. If the broadcast fails to resolve the name, an lmhosts file is used. |
| P-node | Point-to-point communication with a NetBIOS name server, such as WINS, to register and resolve computer names to IP addresses. |
| M-node | A mix of B-node and P-node communication to register and resolve NetBIOS names. M-node first uses broadcast resolution, and then attempts a server query if necessary. |
| H-node | A hybrid of B-node and P-node. An H-node computer attempts to query a server first and uses broadcasts only if direct queries fail. Windows 2000 and Windows XP–based computers are configured to use H-node by default. |

**New Features for Windows Server 2003**

The following improvements to the Windows Internet Name Service (WINS) have been made in the Windows Server 2003 family:

**Filtering records**

Improved filtering and new search functions help you locate records by showing only those records that fit the criteria you specify. These functions are particularly useful in analyzing very large WINS databases. You can use multiple criteria to perform advanced searches for WINS database records. This improved filtering capability allows you to combine filters for customized and precise query results. Available filters include: record owner, record type, NetBIOS name, and IP address with or without subnet mask.

Because you can now store query results in the cache of the memory on your local computer, the performance of subsequent queries is increased, and network traffic is reduced.

**Accepting replication partners**

When determining a replication strategy for your organization, you can define a list that controls the source of incoming name records during pull replication between WINS servers. In addition to blocking name records from specific replication partners, you can also choose to accept only name records owned by specific WINS servers during replication, excluding the name records of all servers that are not on the list.

# Building Your WINS Server Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When building your WINS server strategy, account for any existing hardware that you might need to upgrade, how many WINS servers are needed for your design, and how your server strategy increases WINS availability and optimizes WINS performance. Figure 4.2 shows the process for building your WINS server strategy.

**Figure 4.2 Building Your WINS Server Strategy**



# Reviewing WINS Hardware

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Determine whether your current WINS server hardware is sufficient to upgrade to Windows Server 2003. You might need to upgrade your server hardware for optimal WINS performance. A dual-processor WINS server increases performance about 25 percent, and a dedicated disk drive measurably improves WINS server name replication response time.

When selecting your hardware, consider the following performance guidelines:

- 67Use high-performance disk hardware. WINS causes frequent and intense activity on server hard disks.
- Consider using a redundant array of independent disks (RAID)-based solution, which improves disk access time.
- When evaluating the performance of a server, include WINS to ensure the server can handle its demanding use of central processing unit (CPU), memory, and disk input/output (I/O). Monitor server usage to determine whether WINS server hardware needs to be upgraded.

For a current list of compatible hardware, see the Hardware Compatibility List (HCL) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about determining hardware compatibility, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

# Determining How Many WINS Servers to Deploy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The number of WINS servers needed and the locations of each server depend on the number of WINS clients per server and the network topology.

The number of users each server can support depends on usage patterns, data storage, and the processing capabilities of the server. A WINS server can typically register 1,500 names per minute or answer 4,500 queries per minute. This means that a single WINS server can adequately service up to 10,000 clients.

You might install additional WINS servers in locations separated by slow, or pay-by-usage wide area network (WAN) links. Set conservative client counts for a WINS server to minimize client query response times. Allow room in your design for peak-load conditions, such as large-scale power outages that force many computers to restart simultaneously, thereby bombarding the WINS servers with registration requests.

# Designing WINS for High Availability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Any design that requires high availability must include more than one WINS server. Consider all possible points of failure, including servers, WAN links, and routers. These factors, along with the business goals of the organization, determine the required level of WINS redundancy and fault tolerance.

To ensure that you are planning a fault-tolerant WINS design, ask the following question for each server on your network: What happens to WINS if this server shuts down or if clients cannot reach it?

To help answer the question, consider two common situations in which a WINS server might fail to perform its role on a network:

- A hardware or power failure requires downtime for server repair or maintenance.
- A network link or router failure isolates a WINS server from clients.

To prepare for both of these situations:

- Consider the length of time a WINS server might be out of service on your network, factoring in both planned and unplanned outages.
- Consider what happens to your WINS clients if their primary WINS server shuts down. By maintaining and assigning secondary WINS servers for clients, you can reduce the impact of a single WINS server being offline.

For each client, specify the servers for WINS lookup and the node type. When designing your WINS client support strategy for maximum availability, do the following:

- Specify multiple WINS servers for clients to provide server redundancy.
- For fault tolerance in the case of link failure, point clients to a local WINS server as their primary WINS server, and a remote WINS hub as their secondary WINS server. Ideally, the secondary WINS server is located in a separate building and on a separate power grid from the primary WINS server.
- Consider using an Lmhosts file to provide secondary name resolution in the event of a WINS failure.

  While Lmhosts files are not a recommended solution, in rare circumstances they can provide an effective temporary solution. Lmhosts files must be tightly managed because changes in the NetBIOS environment do not automatically update in static name files.

  For more information about the Lmhosts file, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Using Multiple Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To provide additional fault tolerance, configure a secondary (or backup) WINS server. Although WINS replication architecture benefits from employing a minimum number of WINS servers, employing a secondary WINS server improves the availability of your design. This solution balances performance and availability against cost and manageability.

When using two WINS servers to provide redundancy and load balancing, configure the replication relationship between these servers as a pull or push partnership. When you use replication, both servers contain the same WINS database information.

When a WINS server is configured as a pull partner, it periodically queries the partner server to determine if any updates are available. Use pull partnerships:

- Over lower-speed WAN or congested local area network (LAN) connections.
- To reduce replication traffic by consolidating WINS database updates.
- To perform WINS database updates at scheduled intervals.

When a WINS server is configured as a push partner, the WINS server notifies the partner server that updates are available for replication. Use push partnerships:

- Over LAN or higher-speed WAN connections.
- When the network traffic created by frequent WINS replication updates is not a consideration.
- To ensure WINS database updates are received as soon as possible.

The availability that is provided by WINS replication is appropriate for solving availability issues at local and remote locations. Adding a WINS server to a remote location ensures WINS availability in the event that a WAN link or router fails. For more information about replication strategies, see "Designing Your WINS Replication Strategy" later in this chapter.

# Using Windows Clustering

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Clustering provides a higher level of fault tolerance but consumes additional system resources. If your business goals require a WINS design that provides the highest availability, use server clusters as provided by Windows Clustering. By configuring WINS on multiple servers belonging to the same cluster, you:

- Share a common WINS database.
- Provide immediate failover in the event of failure.
- Restore failed servers sooner, because database resynchronization is not required between the cluster nodes.

**Note**

- If you choose to cluster your WINS servers, be sure to equip the servers with a hard disk containing high-speed I/O that is dedicated to WINS. This can speed up the database response and ensure clustering efficiency.

**Example: A Company Uses a Cluster to Simplify their WINS Design**

A large corporation uses a server cluster to provide infrastructure services, including WINS. Prior to implementing the server cluster, the company had a large and complicated Windows NT 4.0–based WINS replication topology. To maintain consistency and provide accurate information to clients, WINS client records were replicated to all WINS servers.

To simplify the replication matrix, provide redundancy, and more efficiently manage the WINS traffic load, a server cluster is used as the WINS replication hub. Applications and services running on nodes in the cluster are exposed to users and workstations as virtual servers. Figure 4.3 shows the replication matrix before the WINS cluster implementation.

**Figure 4.3 WINS Topology Pre-Clustering**




Figure 4.4 shows the new simplified replication matrix using a server cluster.

**Figure 4.4 WINS Topology Post-Clustering**



Windows Clustering only solves local availability issues. Windows Server 2003–based servers that belong to the same cluster require persistent, high-speed connections between all servers in the cluster.

**Note**

- Before adding WINS to a set of clustered servers, be sure to consider both the advantages and disadvantages of doing so. In many cases, the overall number of WINS servers is small, so clustering WINS is not necessary — replication makes WINS fault tolerant. Instead, configure your WINS clients with the address of a secondary WINS server to ensure uninterrupted service.

For more information about server clusters, see "Designing and Deploying Server Clusters" in *Planning Server Deployments* of this kit.

**Caution**

- WINS does not support rolling upgrades from Windows 2000 to Windows Server 2003 in a server cluster. You can upgrade and failover to Windows Server 2003. However, when WINS is brought online on Windows Server 2003, it cannot fail back to the Windows 2000 node.

# Optimizing WINS Performance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although WINS is designed to help reduce broadcast traffic between local subnets, it creates some traffic between servers and clients. This is particularly important if you use WINS on routed TCP/IP networks.

To optimize performance, begin by estimating the amount of network traffic between WINS clients and WINS servers under normal conditions. Estimate and monitor the following:

- NetBIOS names commonly registered by WINS clients.
- WINS registration and renewal caused by daily startup of clients.
- Mobile users and their effect when moving within a routed network.
- The effects of slower links, such as WAN links and their effect on replication performance and convergence.

**Reducing Response Time**

Reducing the response time of WINS improves performance, with the greatest visibility to users and management. As a result, a design that reduces the response time of WINS is highly successful.

The performance of your WINS design largely depends on other network traffic. For example, a subnet that relies on a WINS server elsewhere on the WAN might experience poor performance during peak hours when network usage is high. Increase the NetBIOS name registration renewal period, which defaults at six days, to reduce client-to-server renewal traffic. This setting must be changed on the WINS server.

Obtain reliable figures on the number of locations and hosts that your WINS design must support. When planning for WINS client traffic on large, routed networks, estimate and monitor the effect of name query, registration, and response traffic routed between subnets. Name requests and responses that occur at the daily startup of computers must pass through the traffic queues on the routers and might cause delays at peak times.

**Consolidating Multiple Subnets**

When you have multiple subnets in a small remote office, consider consolidating the office to one subnet address.

You can do this using asynchronous transfer mode (ATM) switching or a virtual private network (VPN) configuration. By consolidating to one subnet address, you can configure clients to use local broadcasts to resolve names before attempting to contact a WINS server across the WAN. Changing the client to M-node allows it to broadcast locally for resources before contacting a WINS server for NetBIOS name resolution. This can help to reduce the overall amount of WINS-associated traffic, especially WAN traffic.

Use DHCP scope option 046, WINS/NBT Node Type, to configure your WINS clients as M-node clients. For more information about configuring DHCP options at the DHCP server, see "Assign a scope-based optionAssign a scope-based option" in Help and Support Center for Windows Server 2003.

**Configuring Burst Handling**

Burst handling supports a high volume of WINS client name registration. When a large number of WINS clients simultaneously try to register their NetBIOS names, the WINS server can become saturated. In burst handling mode, the WINS server responds positively to clients that submit a registration request before the WINS server has processed and entered these updates in the WINS server database. The WINS server immediately sends a relatively short, random Time to Live (TTL) lease length to all WINS clients. The short TTL lease length forces WINS clients to reregister after the excessive WINS registration traffic subsides, therefore decreasing the load on the network and varying the delay interval to distribute the load over time.

Using the WINS MMC snap-in, you can configure the level of burst handling for the server, which modifies the size of the burst queue.

**To configure burst handling**

1. In the WINS MMC snap-in, right-click the appropriate WINS server.

2. Select the **Advanced** tab from the *server name* properties dialog box.

3. In **Enable Burst Handling**, select **Low (300)**, **Medium (500)**, **High (1000)**, or **Custom (50-5000)** as the burst queue size.

**Load Balancing with Redundant WINS Databases**

A WINS implementation design provides higher performance by specifying that multiple WINS servers contain replicas of WINS databases. These redundant servers improve performance by providing load balancing.

Use load balancing with redundant WINS databases when:

- The length of time to perform WINS functions is unacceptably long.
- The connections between the WINS servers support the additional WINS replication traffic.
- The traffic generated by WINS clients accessing a WINS server in another location saturates a WAN link.
- The cost of adding a server is not prohibitive.

# Designing Your WINS Replication Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A good replication design is essential to your WINS availability and performance. Designs encompassing multiple WINS servers distribute NetBIOS name resolution across LAN and WAN environments, confining WINS client traffic to localized areas. To ensure consistent, network-wide name resolution, WINS servers must replicate their local entries to other servers. For more information about a WINS replication strategy, see the examples later in this section.

Figure 4.5 shows the process for designing your WINS replication strategy.

**Figure 4.5 Designing Your WINS Replication Strategy**




Before configuring replication, carefully design and review your WINS replication topology. For WANs, this planning can be critical to the success of your deployment and use of WINS.

WINS provides the following choices when you are configuring replication:

- You can manually configure WINS replication for a WAN environment.
- For larger networks, you can configure WINS to replicate within a LAN environment.
- In smaller or bounded LAN installations, consider enabling and using WINS automatic partner configuration for simplified setup of WINS replication.
- In larger or global installations, you might have to configure WINS across untrusted Windows NT domains.

If your network uses only two WINS servers, configure them as push/pull replication partners to each other. When configuring replication partners, avoid push-only or pull-only servers except where necessary to accommodate slow links. In general, push/pull replication is the most simple and effective way to ensure full WINS replication between partners. This also ensures that the primary and secondary WINS servers for any particular WINS client are push/pull partners of each other, a requirement for proper WINS functioning in the event of a failure of the primary server of the client.

In most cases, the hub-and-spoke model provides a simple and effective design for organizations that require complete

convergence with minimal administrative intervention. For example, this model works well for organizations with centralized headquarters or a corporate data center (the hub) and several branch offices (the spokes). Also, a second or redundant hub (that is, a second WINS server in the central location) can increase the fault tolerance for WINS.

In some large enterprise WINS networks, limited replication partnering can effectively support replication over slow network links. However, when you plan limited WINS replication, ensure that each server has at least one replication partner. Furthermore, balance each slow link that employs a unidirectional link by a unidirectional link elsewhere in the network that carries updated entries in the opposite direction.

# Specifying Automatic Partner Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure a WINS server to automatically configure other WINS server computers as replication partners. By using this automatic partner configuration, other WINS servers are discovered when they join the network and are added as replication partners.

When using automatic partner configuration, each WINS server announces its presence on the network by using periodic multicasts. These announcements are sent as Internet Group Management Protocol (IGMP) messages for the multicast group address of 224.0.1.24, which is reserved for WINS server use.

Automatic partner configuration is typically useful in small networks, such as single subnet LAN environments. However, you can use automatic partner configuration in routed networks. For WINS multicast support in routed networks, the forwarding of multicast traffic is made possible by configuring routers for each subnet to forward traffic to the WINS multicast group address of. 224.0.1.24.

Because periodic multicast announcements between WINS servers can add traffic to your network, automatic partner configuration is recommended only if you have a small number of installed WINS servers (typically, three or fewer).

Automatic partner configuration monitors multicast announcements from other WINS servers, and performs the following configuration steps:

- Adds the IP addresses for the discovered servers to its list of replication partner servers.
- Configures the discovered servers as push/pull partners.
- Configures pull replication at two-hour intervals with the discovered servers.

If a remote server is discovered and added as a partner by means of multicasting, it is removed as a replication partner when WINS shuts down properly. To have automatic partner information persist when WINS restarts, you must manually configure the partners.

To manually configure replication with other WINS servers, use the WINS Microsoft Management Console (MMC) snap-in or the Netsh command-line tool to specify roles for each partner and any related information.

For more information about the Netsh command-line tool, see "Netsh" and "Netsh commands for WINS" in Help and Support Center for Windows Server 2003.

# Determining Replication Partners

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Choosing whether to configure a WINS server as a push partner, pull partner, or push/pull partner depends on several considerations, including the specific configuration of servers at your site, whether the partner is across a WAN, and how important it is to distribute changes immediately throughout the network.

In the hub-and-spoke configuration, you can configure one WINS server as the central server and all other WINS servers as push/pull partners of this central server. Such a configuration ensures that the WINS database on each server contains addresses for every node on the WAN. Figure 4.6 shows replication using a hub-and-spoke topology.

**Figure 4.6 WINS Replication in a Hub-and-Spoke Topology**




You can select other configurations for replication partner configurations to meet the particular needs of your site. For example, Figure 4.7 shows replication in a T network topology, in which Server1 has only Server2 as a partner, but Server2 has three partners. So Server1 gets all the replicated information from Server2, but Server2 gets information from Server1, Server3, and Server4.

**Figure 4.7 Replication in a T Network Topology**




If Server2 needs to perform pull replications with Server3, make sure it is a push partner of Server3. If Server2 needs to push replications to Server3, configure it as a pull partner of Server3. Determine whether to configure WINS servers as either pull or push partners, and set partner preferences for each server.

# Determining Convergence Time

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The time needed to replicate a new entry in a WINS database, from the WINS server that owns the entry to all other WINS servers on the network is defined as *convergence time*. When planning for WINS servers, you must decide what is acceptable as the convergence time for your network; the longer the replication path, the longer the convergence time.

Name query requests can succeed before the convergence time elapses, resulting in earlier replication of the new entry. This happens when:

- The entries replicate over a shorter path than the worst-case path.
- **The Number of changes in version ID before replication** threshold expires in the push replication settings before the **Replication Interval** expires in the pull replication settings in the **Replication Partners Properties** dialog box in the WINS snap-in.

# Configuring Replication Across WANs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When configuring WINS replication across WANs, the two most important issues are:

- Whether your WINS replication occurs over slower WAN links.
- The length of time required for all replicated changes in the WINS database to converge and achieve consistency on the network.

The frequency of WINS database replication between WINS servers is a major design issue. The WINS server database must be replicated frequently enough to prevent the downtime of a single WINS server from affecting the reliability of the mapping information in other WINS servers. However, the time interval between replications cannot be so small that it interferes with network throughput.

Network topology can influence your decision on replication frequency. For example, if your network has multiple hubs connected by relatively slow WAN links, you can configure WINS database replication between WINS servers on the slow links to occur less frequently than replication on the LAN or on fast WAN links. This reduces traffic across the slow link and reduces contention between replication traffic and WINS client name queries.

After determining the replication strategy that works best for your organization, map the strategy to your physical network. For example, if you have chosen a hub-and-spoke strategy, indicate on your network topology map which sites have the "hub" server, and which have the "spoke" servers. Also indicate whether the replication is push/pull, push-only, or pull-only.

Document the configurations of each WINS server, including the hardware configuration, IP address, replication configuration, and replication partners.

For more information about WIN configuration across WANs, see "Configuring WINS replication" in Help and Support Center for Windows Server 2003.

# Configuring Replication Across LANs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When configuring WINS replication across LANs, the issues are similar to those that occur in WAN environments, although less critical.

Because the data throughput of the underlying network links for LANs is much greater than for WANs, it might be acceptable to increase the frequency of WINS database replication by specifying push and pull parameters for LAN-based replication partners. For push/pull partners, you can do this by decreasing the **Number of changes in version ID before replication** and **Replication interval** settings from what you use for WAN-based partners on slower links.

For example, between LAN-based replication partners it often works to enable WINS to use a persistent connection between the servers. Without a persistent connection, the normal update count threshold defaults to a minimum of 20. You can specify a smaller update count with a persistent connection.

Next, you can specify a much smaller number, such as a value of one to three in the **Number of changes in version ID before replication** setting before WINS sends a push replication trigger to the other partner. For pull partners, you might also consider setting the **Replication interval** setting to a value in minutes, instead of hours.

As in WAN replication planning, the WINS server database must replicate frequently enough to prevent the downtime of a single WINS server from affecting the reliability of the mapping information in other WINS servers. However, the time interval between replications cannot be so small that it interferes with network throughput.

In environments with a large amount of network traffic it is a good idea to use a network monitoring tool, such as Network Monitor, to help measure and determine how to optimize your WINS replication strategy.

For more information about WINS configuration across LANs, see "Configuring WINS replication" in Help and Support Center for Windows Server 2003. For more information about the Network Monitor tool, see "Network Monitor" in Help and Support Center for Windows Server 2003.

# Configuring Replication Between Untrusted Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is possible to set up WINS replication between one or more WINS servers in domains that do not have a trust relationship. You can do this without a valid user account in the untrusting domain. To configure replication, an administrator for each WINS server must use the WINS snap-in or Netsh commands to manually configure each server to permit this replication.

**Important**

- If you require replication across a firewall, keep in mind that WINS replication occurs over TCP port 42. Therefore, this port must not be blocked on any network device between two WINS replication partners.

For more information about WINS configuration across domains that do not have trust relationships, see "Configuring WINS replication" in Help and Support Center for Windows Server 2003. For more information about domain trusts, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Mapping the Replication Architecture to the Physical Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After determining the replication strategy that works best for your organization, map the strategy to your physical network. For example, if you have chosen a hub-and-spoke strategy, indicate on your network topology map which sites will have the "hub" server, and which will have the "spoke" servers. Also indicate whether the replication is push/pull, push-only, or pull-only.

Document the configurations of each WINS server, including the hardware configuration, IP address, replication configuration, and replication partners.

The convergence time for the system is the sum of the two longest convergence times to the hub. For example, in an organization that has five WINS servers (WINS-A through WINS-E), if WINS-B and WINS-D replicate with WINS-A (the hub) every 30 minutes, and WINS-C and WINS-E replicate with the hub every 4 hours, the convergence time is 8 hours.

The following examples show three different types of replication.

**Example 1: Deploying WINS Over a Large Number of Branch Offices**

In this example, a medium-sized company has two main sites: a New York and a Los Angeles office with 500 computers in each office, connected through high-speed links. The company also has more than 160 small branch offices, including local sales offices. To save on the costs of the links, some branches act as concentrators for a region. Figure 4.8 shows a WINS server placement strategy for an organization with many small branch offices.

**Figure 4.8 Deploying WINS Over a Large Number of Branch Offices**

In most cases, the branches do not have local WINS servers — there is simply no need for a separate server for each branch. Instead, the company adds regional WINS servers when the costs of registration and query traffic increase above the cost of deploying the additional server. When the link to a regional WINS server fails, local names can still be resolved by the broadcast mechanism.

The regional WINS servers are not required for this configuration to function correctly, but they do provide a cost optimization. The company's network administrators avoid deploying the regional servers whenever possible because the added servers increase the convergence time. Administrators configure regional WINS servers as replication partners of the WINS servers in the main sites.

Clients in the main site are configured with the IP address of their local WINS server as primary, and the IP address of the WINS server in the other main site as secondary. Clients in the regional branches are configured with the IP address of the regional WINS server as primary, and the address of the closest main site WINS server as secondary. The replication interval is set to 15 minutes between sites; therefore, all computers are reachable within 15 minutes of an address registration or change.

**Example 2: Deploying WINS with a Concentrated User Base**

Figure 4.9 shows the network configuration of another example company that is very different. The network serves a larger company with three sites, Philadelphia, Seattle, and Houston, each with 5,000 users. The number of users justifies two WINS servers at each site.

**Figure 4.9 Deploying WINS Over a Few Large Sites**




The clients are configured with a local primary and secondary WINS server. Half of the clients have one local WINS server as primary and the other as secondary. The other half has exactly the opposite configuration. This balances the registration and query load over both WINS servers, and it provides a backup for maintenance purposes and in case of a server failure.

The local WINS servers use a very short replication interval of 10 minutes; therefore, all computers within the same site are reachable within 10 minutes of an address registration or change. The replication interval between the sites can be longer — about 30 minutes — because most users work with resources within their site.

**Example 3: A Large Hub-and-Spoke Design**

Figure 4.10 shows an extremely large WINS implementation, serving more than 100,000 nodes. In a configuration with so many WINS servers, a common error is to create many push/pull relationships for redundancy. This can lead to a system that, while functional, is overly complex and difficult to understand and troubleshoot.

**Figure 4.10 Large-Scale WINS Deployment Using Hub Topology**




Four major hubs are located in Seattle, San Francisco, Chicago, and Los Angeles. These hubs serve as secondary WINS servers for their regions while connecting the four geographic locations. All primary WINS servers are configured as push/pull partners with their hubs, and the hubs are configured as push/pull partners with other hubs.

The primary WINS servers replicate with the hubs every 15 minutes, and the hub-to-hub replication interval is 30 minutes. The convergence time of the WINS system is the time it takes for a client registration to be replicated to all WINS servers.

In this case the longest convergence time would be 1.5 hours from a Seattle primary server to a Chicago primary server. The total convergence time can be calculated by adding up the maximum time between:

- Seattle primary to Seattle secondary, 15 minutes
- Seattle secondary to San Francisco secondary, 30 minutes
- San Francisco secondary to Chicago secondary, 30 minutes
- Chicago secondary to Chicago primary, 15 minutes

However, the convergence time might be longer for WINS servers connected across slow links. It is probably not necessary for the servers in Paris or Berlin to replicate every 15 minutes. You might configure them to replicate every two hours or even every 24 hours, depending on the volatility of names in the WINS system.

This network contains low redundancy. If the link between Seattle and Los Angeles is down, replication still occurs through San Francisco. If, for example, the Seattle hub fails, the Seattle area can no longer replicate with the rest of the WINS system. Network connectivity, however, is still functional — all WINS servers contain the entire WINS database, and name resolution functions normally. All that is lost are changes to the WINS system that occurred since the Seattle hub failed. A Seattle user cannot resolve the name of a file server in Chicago that comes online after the Seattle hub fails. When the hub returns to service, all changes to the WINS database are replicated normally.

# Securing Your WINS Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In many WINS implementations, WINS replication occurs across public networks, such as the Internet. Replicating the NetBIOS names and IP addresses of all hosts within the organization over these public networks creates a security risk, which you can mitigate by using VPN tunnels or placing servers within a perimeter network. Figure 4.11 shows where you perform this step in the process of deploying your WINS solution.

**Figure 4.11 Securing WINS During the Deployment Process**



# Securing WINS Traffic with Tunnels

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

All WINS replication traffic sent over public networks should be encrypted. Encrypt the replication traffic by using Internet Protocol security (IPSec) or VPN tunnels. When choosing to encrypt replication traffic by using IPSec or VPN tunnels, do the following to further increase security:

- Use the strongest level of encryption.
- Use the Routing and Remote Access service to provide the IPSec or VPN tunnel.
- Use Kerberos V5 or other certificate-based authentication for secure communication channels.

For more information about deploying IPSec, see "Deploying IPSec" in this book. For more information about virtual private networks and the Routing and Remote Access service, see "Deploying Dial-up and VPN Remote Access Servers" in this book. For more information about enabling Kerberos V5 authentication, see "Enabling Kerberos V5 authentication" in Help and Support Center for Windows Server 2003.

# Running WINS on a Perimeter Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Place WINS servers in a perimeter network when you must send WINS traffic over a public network to avoid exposing intranet NetBIOS names and WINS data. This placement protects corporate resources while providing NetBIOS name resolution to external clients that need access to these resources.

**Caution**

- If you require replication from the WINS server in the perimeter network to a WINS server within the intranet, in the WINS snap-in, select **Replicate Only with Partners** in the **Replication Partners Properties** dialog box on both the WINS servers. Also consider using only pull replication from the intranet servers. To maintain security, encrypt all replication traffic across the inner firewall using IPSec or VPN tunnels.

# Integrating WINS with Other Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most network administrators deploying WINS also plan a strategy for DNS and DHCP servers, because WINS is so closely linked to DNS and DHCP. Figure 4.12 shows when you perform this step in the process of deploying your WINS solution.

**Figure 4.12 Integrating WINS During the Deployment Process**



# Integrating WINS with DNS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If most of your clients use NetBIOS and your servers are running Windows 2000 or Windows Server 2003 DNS, enable WINS lookup on your DNS servers. When WINS lookup is enabled on DNS servers, WINS resolves any names that DNS resolution does not find. DNS does not support the WINS forward lookup and WINS-R reverse lookup records in versions of Windows earlier than Windows 2000. For information about enabling WINS lookup, see "Deploying Domain Name System (DNS)" in this book.

**Note**

- For a smooth integration with DNS, do not use extended characters in NetBIOS names, especially the underscore (_) and the period (**.**). Consult with your DNS administrator when determining NetBIOS naming standards.

If all of your network computers are running Windows 2000, Windows XP, or Windows Server 2003 and you are not supporting any applications that require NetBIOS names, you might consider establishing DNS as your only method of name resolution. However, before you consider decommissioning your WINS servers, identify any computers or applications that rely on NetBIOS, and determine the impact of removing NetBIOS. You might find that a critical application relies on NetBIOS (with no alternative currently available) in which case, you must continue to use WINS. For example, certain applications, such as Microsoft® Systems Management Server (SMS) and Microsoft® BackOffice® client/server mail configurations using Exchange Server, might require NetBIOS naming.

For more information about DNS, see "Deploying Domain Name System (DNS)" in this book or see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Integrating WINS with DHCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When using DHCP and WINS together on your network, use additional DHCP scope options to assign WINS node types and to identify WINS primary and secondary servers for DHCP clients.

Computers with static IP addresses can be problematic and their initial registration record in WINS becomes tombstoned if they are not periodically stopped and restarted. You can have a more reliable and manageable network by creating DHCP reservations for these computers. These reservations ensure that the computer gets the same IP address from the DHCP server for each request.

When you configure a network to use both DHCP and WINS, set the DHCP lease period to be equal to or greater than the WINS renewal period. This prevents a situation in which the WINS server fails to notice that a DHCP client has released a DHCP-assigned IP address. Specifically, the client cannot send a WINS renewal request because it did not renew its IP address. If another computer is assigned that IP address before the WINS server notes the change, the WINS server mistakenly directs requests for the address to the new client. This is significant only if you do not use the default lease lengths for both services, and lease durations were changed for either DHCP or WINS individually.

For more information about deploying DHCP, see "Deploying DHCP" in this book.

# Implementing Your WINS Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For best performance, avoid deploying WINS on heavily loaded servers, or on servers that perform other tasks that might reduce performance of the hard disk, memory, and processors. If you do host more than one service on a WINS server, consider how each service might impact the others. For example, hosting two name services (such as WINS and DNS) on the same server can result in slow performance for each service due to intensive hard disk access. Or a WINS server hosting a print service might be slow to register names if it is handling several large print jobs simultaneously. It might be appropriate to run multiple services on the same computer in a branch office because servers in branch offices are usually not under heavy use.

Figure 4.13 shows the process for implementing your WINS solution.

**Figure 4.13 Implementing Your WINS Solution**



# Migrating WINS to Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

1. Before migrating from legacy WINS servers, make sure your existing WINS infrastructure is appropriate for your current needs. For example, if you have recently upgraded most desktop computers in your organization to Windows 2000 or Windows XP, or if you have recently stopped using an application that relies heavily on WINS, your current WINS structure might be too robust for your current needs, and might not be structured in the most efficient way possible. In a case such as this, start the deployment from the design phase, rather than migrating the existing database to new servers.

2. Follow these steps when migrating your WINS database from Windows NT 4.0 or Windows 2000 to Windows Server 2003:

3. Install the WINS service.

   This can be installed either during or after installing Windows Server 2003.

4. Configure the WINS service.

   Verify that the server is pointing to itself for WINS. You can do this by viewing the TCP/IP properties of your network adapter.

5. Convert the WINS database for use on the Windows Server 2003–based server.

   This conversion might occur automatically from existing Windows NT 4.0–based or Windows 2000–based servers. If not, follow these steps:

   a. At the command prompt, type **net stop wins** on both the existing and new servers.

   b. Copy the contents of the %SystemRoot%\System32\Wins folder from the existing server to the new Windows Server 2003–based server.

   c. At the command prompt, type **net start wins** on both servers.

   **Note**

This process can take 30 minutes or more to complete depending on the size of the database. Do not stop the process until it is finished. It is normal for Jetconv.exe to require heavy CPU usage during the conversion.

During the conversion process, you might be prompted for additional files from the Windows Server 2003 operating system CD.

**To access WINS conversion files**

1. Copy the Edb500.dl_ file from the I386 folder on the CD-ROM to the %SystemRoot%\System32 folder on the server.

2. At the command prompt, type **expand edb500.dl_ edb500.dll** to expand the Edb500.dl_ file on the server.

3. At the command prompt, type **net start wins** to finish the conversion process.

4. Verify that the WINS database is shown in the WINS snap-in on the server.

# Testing Your WINS Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After completing your WINS design, test it in a lab to find potential problems before implementing your design on your production network. As you roll out your design, test your network to ensure it is working as expected.

The best time to discover potential problems with your design is in a test lab prior to your full implementation. When preparing your test lab, be sure to:

- Use a server computer from the same vendor and with the same configuration as the servers that will be used for the actual WINS servers. Set up a representative sample of the computers in your organization to be tested as WINS clients.
- If you are planning to deploy WINS over a WAN, design your lab with routers and use a link simulator to simulate network latency.
- Deploy a typical set of applications together on the WINS test server. This step is vital in determining any compatibility issues that might arise when users run different applications simultaneously.

For more information about planning a test environment, see "Designing a Test Environment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

# Evaluating the Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After implementing your WINS design, evaluate your deployment to ensure that it complies with your design and meets your organization's business goals.

**To assess the availability of your design** Stage a simulated failure to ensure that functionality, security, and performance are maintained.

**To evaluate WINS service availability** Disable or disconnect each WINS server that is a part of a redundant WINS design. Provide procedures detailing how to restore synchronization of WINS databases after a failed server is reactivated or repaired.

**To evaluate WINS security** Initiate WINS replication, and examine the data transmissions between the locations to ensure that the WINS replication traffic is encrypted.

# Additional Resources for Deploying WINS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more information about WINS, refer to the following sources:

**Related Information**

- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about Windows Internet Name Service (WINS), Windows Server 2003 DNS, or the Lmhosts file.
- "Deploying Domain Name System (DNS)" in this book for information about enabling WINS lookup or about planning and designing your DNS network.
- "Designing and Deploying Server Clusters" in the *Planning Server Deployments* book of this kit.
- "Deploying DHCP" in this book.
- "Deploying Dial-up and VPN Remote Access Servers" in this book for more information about virtual private networks and the Routing and Remote Access service.
- "Deploying IPSec" in this book.
- "Designing a Test Environment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about domain and forest trusts.
- RFC 1001: *Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods*

**Related Tools**

- For more information about the Network Monitor tool, see "Network Monitor" in Help and Support Center for Windows Server 2003.
- For more information about the Netsh command-line tool, see "Netsh" in Help and Support Center for Windows Server 2003.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "WINS" in Help and Support Center for Windows Server 2003.
- "Netsh commands for WINS" in Help and Support Center for Windows Server 2003.
- "Configuring WINS replication" in Help and Support Center for Windows Server 2003 for more information about WINS configuration across WANs, LANs, or untrusted domains.
- "Enabling Kerberos V5 authentication" in Help and Support Center for Windows Server 2003.

# Deploying ISA Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft® Internet Security and Acceleration (ISA) Server is an integrated firewall and Internet caching server. Deploying ISA Server saves network bandwidth by providing faster Web access for users. ISA Server secures your network, allowing you to implement your business security policy by configuring a broad set of rules that specify which sites, protocols, and content can be passed through the firewall.

**In This Chapter**


Overview of ISA Server

Determining the Role of ISA Server

Evaluating Interoperability Issues

Designing for High Availability

Securing the Design

Implementing Your ISA Server Solution

Additional Resources for Deploying ISA Server


**Related Information**

- For more information about creating a virtual private network (VPN), see "Deploying Dial-up and VPN Remote Access Servers" in this book.

# Overview of ISA Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ISA Server offers a complete Internet connectivity deployment solution as a firewall and a Web caching server.

**Important**

- ISA Server (the successor to Proxy Server 2.0) is a separate product from the Microsoft® Windows® Server 2003 operating system. You must install ISA Server Service Pack (SP1) for ISA Server immediately after installing ISA Server for it to be interoperable with Windows Server 2003.

ISA Server monitors requests and responses between the Internet and internal client computers; ISA Server also controls which computers on the Internet the internal clients can access.

ISA Server offers many security and caching options, including:

- **Enhanced security with multilayer firewall and integrated intrusion detection**.

  ISA is a complete firewall product that provides packet filtering, stateful packet inspection and application-level awareness for many common protocols, such as Simple Mail Transfer Protocol (SMTP) and Domain Name System (DNS). You can create access policies based on user-level information, IP addresses, or Web content.

- **Secure publishing**.

  You can use ISA Server as a reverse cache server to define a secure publishing policy. This protects internal publishing servers and makes them safely accessible to Internet clients.

- **ISA Server can cache Web content, potentially reducing your external bandwidth requirements**.

  You can configure the cache to contain content that your organization uses frequently or that your Internet clients access.

- **ISA Server is extensible**.

  You can program the COM interface in ISA Server using high-level programming languages or scripting languages. Developers can implement application and web filters to extend the core firewall functionality. You can use the ISA Server management interface to develop additional administration tools and management scripts.

# ISA Server Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process for deploying ISA Server includes determining the deployment mode (cache, firewall, or integrated) you need to use, examining the integration between ISA Server and other components in your network, and choosing the best method for securing your network design.

Figure 5.1 shows the process for deploying ISA Server.

**Figure 5.1 Deploying Your ISA Server**



# Determining the Role of ISA Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ISA Server can act as a firewall, a Web caching server, or both in your network. Figure 5.2 shows the process for determining the role of your ISA server.

**Figure 5.2 Determining the Role of ISA Server**




You can install ISA Server in firewall, cache, or integrated mode.

- In **firewall mode**, you can secure network communication by configuring rules and access policies that control communication between your internal network and the Internet. You can also publish internal servers.

- In **cache mode**, you can improve network performance and save bandwidth by storing frequently accessed content closer to the user. You can also route requests from internal users to the appropriate Web server and publish Web servers in cache mode.

- In **integrated mode**, all cache and firewall features are available. You can configure a policy to meet both cache performance and security requirements.

Use Table 5.1 to determine which mode of ISA Server installation is most appropriate in your network.

**Table 5.1 Determining the ISA Server Installation Mode**

| Goal | Mode |
|---|---|
| Secure your connection to the Internet, connect remote offices, or implement secure extranets. | Firewall |
| Increase performance of your Internet connection. | Cache |
| Secure your connection to the Internet and increase the performance of your Internet connection. | Integrated |

# Implementing ISA Server in Firewall Mode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ISA Server in *firewall mode* acts as a secure gateway between the Internet and internal clients. By configuring the access policies, you can prevent unauthorized access and malicious content from entering the network, as well as restrict outbound traffic.

In Figure 5.3, ISA Server is deployed in firewall mode and as a publisher of internal services, including e-mail and Web services.

**Figure 5.3 ISA Server in Firewall Mode**




All inbound traffic requiring access to the Web or e-mail servers must pass through the firewall first. Likewise, ISA Server can also limit Internet access to specified clients. In this example, ISA Server is acting as a dedicated firewall controlling access to the internal network.

# Implementing ISA Server in Cache Mode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ISA Server in *cache mode* accelerates Web access performance by caching Internet content locally. ISA can provide access control for Web content, both in forward cache mode and reverse cache mode.

**Forward Cache Mode**

When a client in the internal network requests a Web page, ISA Server in *forward cache mode* checks if the content is cached locally. If so, the request is not forwarded to the Internet, and the forward cache server returns the Web pages to the client. If the Web page is not stored locally, ISA Server (acting on behalf of the client) retrieves the Web page from the Internet. ISA Server then saves that Web page in the local cache. The next time a client requests that page, ISA Server can fulfill the request without going to the Internet to retrieve the page. This results in using less bandwidth on the Internet connection. Figure 5.4 illustrates a forward cache mode configuration.

**Figure 5.4 ISA Server in Forward Cache Mode**




**Reverse Cache (Web Publishing) Mode**

You can also configure ISA Server in reverse cache mode. ISA Server in *reverse cache mode* caches content provided to the Internet from the Web server. When an Internet client requests a Web page, the request is sent to the ISA Server–based computer first. If the page is stored there locally, there is no need to retrieve the page from the Web server. This increases performance for Internet clients accessing the Web site from the Internet, as well as increasing security for the server. Figure 5.5 illustrates a reverse cache mode configuration.

**Figure 5.5 ISA Server in Reverse Cache Mode**



# Implementing ISA Server in Integrated Mode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ISA Server in *integrated mode* provides a firewall solution and acts as a Web cache server simultaneously by allowing both services to coexist on the same server.

Before implementing both the firewall and Web cache on the same server, consider the following points:

- Purchasing less equipment can minimize costs.
- Centralizing the management of both resources on a single computer can simplify administration.
- Implementing both services on a single computer presents a single point of failure for both services. If that computer goes offline, multiple services are taken offline.

# Determining Client Types

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ISA Server supports the following types of clients.

**Web Proxy client**
Makes all Internet requests to the ISA Server "Outgoing Web Requests" listener. Most often, this is a Web browser that is either configured manually by the user, or is configured automatically by using Group Policy or a configuration script. ISA restricts user-based Web access controls to Web Proxy clients only.

**SecureNAT client**
Provides IP-based security, but does not allow for user-level authentication. To configure a SecureNAT client, you only have to set the default route between the client and the ISA Server default internal IP address. Because a SecureNAT client uses no other configuration, any computer that uses TCP/IP can be a SecureNAT client.

**Firewall client**
Restricts access on a per-user, per-application basis for outbound access for requests that use Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). To configure a firewall client, you must install the Firewall Client software on each client computer. You can install the Firewall Client software from the shared folder *ISA Server Name*\mspclnt on the ISA Server–based computer.

You can only install the Firewall Client software on computers running Microsoft® Windows® 95 Service Release 2, Windows NT 4.0, Windows® 98, Windows® Millennium Edition, Windows® 2000 Professional, Windows XP Professional, Windows® XP 64-Bit Edition, or the Windows Server 2003 family.

Before you deploy or configure client software assess your organizational needs, determine which applications and services your internal clients require, and assess how you plan to publish servers. Finally, map these needs to the client types supported by ISA Server.

Use Table 5.2 to determine which clients to deploy on your network.

**Table 5.2 Determining Which Clients to Deploy**

| Goal | Client | Reason |
|---|---|---|
| Improve the performance of Web requests for internal clients.<br><br>Combine user-level and content controls to Web access. | Web Proxy | Web Proxy clients do not require any software to be installed but does require specific configuration. |
| Avoid deploying client software or configuring client computers. | SecureNAT | SecureNAT clients do not require any software or specific configuration. |
| Improve Web performance in an environment with non-Microsoft operating systems. | SecureNAT | SecureNAT client requests pass transparently to the ISA Server firewall service and then to the caching service. |
| Publish servers that are located on your internal network. | SecureNAT | You can publish Internal servers as SecureNAT clients, which eliminates the need for creating special configuration settings on the publishing server. |
| Allow Internet access only for authenticated users. | Firewall or Web Proxy | You can configure user-based access policy rules for firewall clients. There is also an option to require authentication with Web Proxy clients. |

# Evaluating Interoperability Issues

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can make ISA Server interoperable with other network services and computers, you need to evaluate possible implementation and integration issues within your configuration, as shown in Figure 5.6.

**Figure 5.6 Evaluating Interoperability Issues**



# Running Other Services with ISA Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can run ISA Server on a computer with other services running already, such as e-mail or Web servers. However, it is not recommended to place other services on the firewall. For improved security and protection, place services behind the firewall, as shown in Figure 5.7. For example, if your organization needs to deploy ISA Server to connect branch offices to the Internet in integrated mode, ISA Server can be installed on existing e-mail or Web servers to simplify deployment and minimize the cost.

**Figure 5.7 Running Other Services with ISA Server**



# Implementing ISA Server in a Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ISA Server with ISA Server SP1 can be installed as a stand-alone server or as a member of an array in a Microsoft® Windows® 2000 Server or Windows Server 2003 domain. When you install ISA Server as a stand-alone server, the configuration information is saved to the local registry. When you install ISA Server as a member of an array in a Windows 2000 or Windows Server 2003 domain, the ISA Server schema is installed into the Active Directory® directory service. You need to install Active Directory on the ISA Server domain to use ISA Server arrays.

ISA Server can be installed as a stand-alone server in a Windows NT Server 4.0 domain. No special configuration is required. You can use arrays of ISA Server computers to connect and secure Windows NT 4.0 domain users and clients to the Internet. However, the array must be set up on a separate Windows 2000 or Windows Server 2003 domain. Then you can establish a trust relationship from the domain containing the ISA Server computer to the Windows NT 4.0 domain.

# Using ISA Server with Other Network Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You might have run the Routing and Remote Access service in Windows 2000 or Windows Server 2003 to make network services and computers available to remote clients. ISA Server enables remote connectivity and extends routing and remote access functionality by providing additional security features. ISA Server packet filtering replaces the packet filtering functionality in the Routing and Remote Access service. In addition, ISA Server uses the dial-up connections that you previously configured for routing and remote access.

Similarly, you might have used Internet Connection Sharing (ICS) or network address translation (NAT) to access the Internet. ISA Server provides the connectivity enabled by NAT or ICS, while adding sophisticated security and caching features.

**Caution**

- You cannot combine ICS or the NAT that is included with the Routing and Remote Access service for Windows Server 2003 with ISA Server or else you will lose firewall functionality. Also, IIS cannot run on the same computer as ISA Server without special configuration. For more information about configuring IIS for use with ISA Server, see the ISA Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, and then search for the keyword *socket pooling*.

# Designing for High Availability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before deploying ISA Server, estimate your Internet connectivity requirements. Use the following sections as guidelines for planning the size and type of your servers, and whether to deploy them as an array. Figure 5.8 shows the process for designing for high availability.

**Figure 5.8 Designing for High Availability**



# Performing Capacity Planning

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The flowchart in Figure 5.9 outlines the capacity planning process for ISA Server.

**Figure 5.9 ISA Server Capacity Planning Process**




Use the following guidelines for capacity planning:

- Ensure the minimum hardware requirements deploying ISA Server are met.
- Decide whether to install ISA Server as a firewall. If yes, ensure that the minimum firewall requirements are met.
- Decide whether or not to install ISA Server as a Web cache server. If yes, ensure that the minimum hardware requirements are met.

The following list describes the minimum hardware requirements for installing ISA Server:

- A computer with a 300 megahertz (MHz) or higher Pentium II-compatible CPU.
- The computer must be running either a member of the Windows 2000 Server family, or a member of the Windows Server 2003 family. If you are installing ISA Server on a computer running Windows Server 2003, you must also install ISA Server SP1.

  **Note**

  - Using the latest service pack is always recommended.
- 256 megabytes (MB) of memory.
- 20 MB of available hard disk space.
- A network adapter to communicate with the internal network.
- One local hard disk partition that is formatted with the NTFS file system.
- To implement the array and enterprise-level policy configuration, you must also run Active Directory.
- If you are using ISA Server in firewall or integrated mode, two network adapters are required.
- If your ISA Server is also supporting other services, such as e-mail or Web services, additional resources might be

required.

Use Table 5.3 to determine the type of computers to use and whether you require an array.

**Table 5.3 Hardware Requirements for Different Network Loads**

| Hits per Second | Minimum Hardware Required | RAM |
|---|---|---|
| Less than 500 | One computer, Pentium II, 300 MHz processor | 256 MB |
| 500 to 900 | One computer, Pentium III, 550 MHz processor | 256 MB |
| More than 900 | One computer, Pentium III, 550 MHz processor, for each 800 hits per second increment | 256 MB per server |

**Note**

- You can also use Performance Monitor to identify bottlenecks and determine whether to add more servers.

If multiple computers are required to handle the network load, consider setting up an array of ISA Server–based computers. Arrays allow Web cache routing across a group of ISA Server–based computers. For more information, see "Configuring ISA Server in an Array" later in this chapter.

**Firewall Requirements**

Table 5.4 lists hardware requirements and network connections based on expected throughput for firewall clients accessing content on the Internet.

**Table 5.4 Hardware Requirements for Firewall**

| Throughput Requirements | Minimum Hardware Required | Internet Connection Type |
|---|---|---|
| 36 Kilobits per second (Kbps) to 1 Megabits per second (Mbps) | One computer, Pentium II, 300 MHz processor | POTS modem, cable modem, or xDSL |
| 384 Kbps to 1.5 Megabits per second (Mbps) | One computer, Pentium II, 300 MHz processor | T1 |
| 3 Mbps - 44 Mbps | One computer, Pentium III, 550 MHz processor | T3 or faster |
| More than 44 Mbps | One computer, Pentium III, 550 MHz, for each 50 MB/second required | OC3 or faster |

**Forward Caching Requirements**

You can deploy ISA Server as a *forward-caching server*, which maintains a centralized cache of frequently-requested Internet content. In this case, consider how many users might access the Internet.

Table 5.5 lists the hardware requirements for using ISA Server in forward cache mode.

**Table 5.5 Hardware Requirements for Forward Caching**

| Internet Users | Minimum Hardware Required | RAM | Disk Space for Caching |
|---|---|---|---|
| Up to 500 | One computer, Pentium II, 300 MHz processor | 256 MB | 2–4 Gigabytes (GB) |
| 500 -1,000 | One computer, two Pentium III, 550 MHz processors | 256 MB | 10 GB |
| More than 1,000 | Two computers, Pentium III, 550 MHz processors | 256 MB for each server | 10 GB for each server |

If your user-base exceeds 1,000 users, you can use hardware with faster processors and more memory, or you can add more ISA Server installations.

**Reverse Caching Requirements**

You can deploy ISA Server as a *reverse-caching server* to fulfill Web requests from the Internet to your network. For example,

you might place an ISA Server computer between the Internet and an organization's Web server that is hosting a commercial Web business or providing access to business partners. In that case, you need to consider how often external clients might request content from the publishing servers.

Table 5.6 lists hardware requirements for ISA Server in reverse cache mode, based on the number of hits per second from Internet users.

**Table 5.6 Hardware Requirements for Reverse Caching**

| Hits Per Second | Minimum Hardware Required |
|---|---|
| Fewer than 100 | One computer, Pentium II, 300 MHz processor |
| 101 to 250 | One computer, Pentium III, 450 MHz processor |
| More than 250 | One computer, Pentium III, 550 MHz processor for each 250 hits per second. You can use Performance Monitor to determine bottlenecks, and then add more servers or more powerful hardware, as necessary. |

Memory requirements depend on the size of the cacheable content that you are publishing, and the working set of the content. Ideally, all cacheable content should fit into the available memory. By default, the ISA Web Proxy service uses half of the available server memory for RAM caching. For example, if the Web site you are publishing has 250 MB of cacheable content, then your ISA server computer should have at least twice this much available RAM before the Web Proxy service starts.

**Adding Computers**

In some cases, you need to decide whether to add an additional ISA Server–based computer or to improve the performance of the existing computer by adding an additional processor. Each option has different advantages.

When you add a new computer and create an array of ISA Server–based computers, you set up a fault-tolerant system. If one computer fails, the other continues to function. On the other hand, adding a computer means that you have to purchase and manage additional hardware and any software that is installed on the computer.

# Designing for Scalability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When designing for scalability, consider differences between Microsoft® Internet Security and Acceleration (ISA) Server Standard Edition and Microsoft® Internet Security and Acceleration (ISA) Server Enterprise Edition, such as:

- ISA Server Standard Edition supports only a single computer configuration, and therefore, cannot be used in an array.
- ISA Server Enterprise Edition can be configured either in a single computer configuration or in an array.

# Configuring ISA Server in an Array

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Computers running ISA Server Enterprise Edition can be grouped together in arrays. An *array* is a group of ISA Server–based computers used to perform Web cache routing. Arrays allow a group of ISA Server–based computers to be treated and managed as a single, logical entity. An array installation also provides increased performance and bandwidth savings. Grouping your ISA Server–based computers in an array allows your client requests to be distributed among multiple servers, thereby improving response time for clients.

All the servers in an array share a common configuration. This saves management time because the array is configured once and the configuration is applied to all the servers in the array. Furthermore, you can apply an enterprise policy to an array. This allows centralized management for all the arrays in the enterprise. A unique array policy can be applied to each array in the enterprise.

It is recommended that you consider installing ISA Server as an array even if there is only one server. The advantages to this include the ability to easily add an additional server to the array in the future and the ability to use the advanced array management features.

**Note**

- All array members must be in the same domain and in the same site.

Table 5.7 compares ISA Server features as a stand-alone server and in an array configuration.

**Table 5.7 Comparing Features of ISA Server as a Stand-Alone Server or as an Array**

| ISA Server Stand-Alone Server | ISA Server Array |
|---|---|
| Can be installed in a Windows NT 4.0 domain. | Requires Active Directory. |
| Cannot use array or enterprise policies. | Uses both enterprise- and array-level policies. |
| Installs from either ISA Server Standard or ISA Server Enterprise Edition. | Installs from ISA Server Enterprise Edition only. |

# ISA Server and DNS Round Robin

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Firewall and Web proxy clients can achieve fault tolerance when two or more computers running ISA Server are used together with a Domain Name System (DNS) server.

You can use DNS to assign the same name to all the ISA Server–based computers in a cluster. With this configuration, when a client requests an object from the ISA Server–based computer specifying the DNS name, the DNS server resolves the name to one of the computers running ISA Server in the array in a round robin fashion. This increases fault tolerance through redundancy and improves performance through the use of multiple computers answering client requests.

**Note**

- For DNS round robin to work for an ISA array, the duplicated resource records must all use the array name.

Figure 5.10 shows the DNS server receiving a request from the clients and forwarding the request to the computers running ISA Server in a round-robin configuration.

**Figure 5.10 DNS Round Robin**



# Securing the Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ISA Server secures your connection to the Internet or to remote sites and extranets. Security decisions for implementing ISA Server are discussed in the following sections. Figure 5.11 shows the process for securing the design.

**Figure 5.11 Securing the Design**



# Connecting Remote Sites Using ISA Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Using ISA Server to connect remote offices includes the following benefits:

- The ability to connect remote offices together through the Internet using a virtual private network (VPN).
- Hierarchical caching can also be implemented across the wide area network (WAN) on ISA Server–based computers.

# Securing Network Perimeters with ISA Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A perimeter network, also known as a screened subnet, is a network that is set up separately from an organization's private network and the Internet. The perimeter network allows external users access to the specific servers located in the perimeter network, while preventing access to the internal network. In addition, an organization might allow very limited access from computers in the perimeter networks to computers in the internal network.

A perimeter network is commonly used for deploying the e-mail and Web servers. The perimeter network can be set up using either of the following configurations:

- Back-to-back perimeter network configuration with two ISA Server–based computers on either side of the perimeter network.

- Three-homed ISA Server–based computer with both the perimeter and internal network protected by the same computer.

# Designing a Back-to-Back Perimeter Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a back-to-back perimeter network configuration, two ISA Server–based computers are located on either side of the perimeter network. Figure 5.12 shows a back-to-back perimeter network configuration.

**Figure 5.12 Back-to-Back Perimeter Network**




Both ISA Server–based computers are set up in integrated or firewall mode. This configuration reduces the risk of compromise by requiring anyone attempting to access the internal network from the Internet to access both systems to reach the internal network.

Perform the following steps to make the servers on the perimeter network available to Internet clients:

1. Configure the local address table (LAT) on the ISA Server–based computer that is connected to the internal network to include the IP addresses of the computers in the internal network.

2. Configure the LAT on the ISA Server–based computer connected to the Internet to include the IP address of the ISA Server–based computer connected to the internal network, and the IP addresses of all the publishing servers in the perimeter network.

3. Create a Web publishing rule on the ISA Server–based computer connected to the Internet to publish the Web server.

4. Create a server publishing rule on the ISA Server–based computer connected to the Internet to publish the e-mail server. Configure the server publishing rule to apply to the e-mail server.

5. Create a Web publishing rule to publish the Web server, and configure the rule to redirect requests to the hosted site.

6. With this back-to-back perimeter network design, selected traffic can access the e-mail or Web server without accessing the internal network. This example publishes the e-mail and the Web servers without exposing the internal network to the Internet.

# Designing a Three-Homed Perimeter Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a three-homed perimeter network, a single ISA Server–based computer is set up with three network adapters:

- The first network adapter connects to clients on the internal network.
- The second network adapter connects to the servers located in the perimeter network.
- The third network adapter connects to the Internet.

Figure 5.13 illustrates the three-homed perimeter network configuration.

**Figure 5.13 Three-Homed Perimeter Network**




Perform the following configuration steps for the three-homed ISA Server perimeter network:

- Configure the LAT to include all of the addresses on the internal network. The LAT should not include the addresses on the perimeter network.
- Enable packet filtering and IP routing.
- Create IP packet filters for each of the servers in the perimeter network. For each IP packet filter, the local computer should be specified as the IP address of the server on the perimeter network.

# Using ISA Server in Extranets

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An extranet is a private network that is configured for use outside your internal network. The extranet is installed to support selected partners who require access to your network. ISA Server supports the installation of extranets through the built-in capability of VPNs. Figure 5.14 shows ISA Server within an extranet design.

**Figure 5.14 ISA Server in Extranets**



# Implementing Your ISA Server Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After determining the roles you want for ISA Server and completing your ISA Server design, you will implement your ISA Server solution. Figure 5.15 shows when to implement your ISA Server solution.

**Figure 5.15 Implementing Your ISA Server Solution**




Figure 5.16 shows the process for implementing the ISA Server solution in your network.

**Figure 5.16 Details of Implementing Your ISA Server Solution**




Use the following steps to deploy ISA Server in your network:

1. Ensure that you are running either Windows 2000 Server or the Windows Server 2003 family. If necessary, upgrade the operating system before continuing with the deployment.

2. Ensure that the latest service pack in installed. If necessary, install the service pack before proceeding with the deployment.

3. Ensure that the network adapters are installed and working correctly; make sure you have a valid connection to the Internet.

    - You can choose to connect your network to the Internet through either a persistent connection (T1, T3, xDSL, or cable modem) or a dial-up connection. If you choose a direct connection, you need to set up a network adapter that connects the ISA Server–based computer to the Internet.

    - When you set TCP/IP properties for the external network adapter, consult with your ISP for the correct settings. Specifically, you need the IP address, subnet mask, default gateway, and IP addresses for the DNS servers to use in DNS name searches. In some cases, your ISP might be using Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) for dynamic assignment of client addresses.

    - Typically, ISA Server has only one IP default gateway. You should configure the IP address of the default gateway on the external — not internal — network adapter. Leave the **Default Gateway** setting for the internal network adapter blank.

    - When setting TCP/IP properties for any internal network adapter, you should enter a permanently reserved IP address for the ISA Server–based computer and an appropriate subnet mask for your internal network. Addressing that is assigned by DHCP should not be used for the internal network adapter, because DHCP might reset the default gateway you selected for the ISA Server–based computer. The external network adapter can be DHCP-enabled, including the default gateway and DNS settings, or these settings can be statically defined.

4. If you are installing ISA Server Enterprise Edition you need to initialize the enterprise. You can run the Enterprise Initialization Tool in ISA Server startup.

    - If this is the first time you are installing ISA Server as an array member, you should run the ISA Server Enterprise Initialization Tool.

    - If you are installing a stand-alone server, or if you have previously installed ISA Server in your enterprise as an array member, you do not need to run the Enterprise Initialization Tool, and you can select **Install ISA Server**.

    **Important**

    - In order to install the ISA Server schema to Active Directory, you must be a member of both the Enterprise Administrators and Schema Administrators groups.

5. Install ISA Server. Setup asks for the following information:

    - **Installation options**.

      You can select a **Typical** installation, **Full** installation, or **Custom** installation.

    - **Array selection**.

      If you previously initialized the enterprise, you can select which array to join. If you did not initialize the enterprise, then ISA Server is installed as a stand-alone server.

    - **Mode**.

      You can install ISA Server in firewall mode, integrated mode, or cache mode.

    - **Cache configuration**.

If you install ISA Server in integrated or cache mode, you need to configure which cache drives to use and the size of the cache.

- **LAT configuration**.

  If you install ISA Server in integrated or firewall mode, you need to configure the address ranges to include in the LAT.

6. Before building your LAT, ensure that you enter all subnets correctly.

   - The LAT is automatically constructed from the Windows Server 2003 routing table. If the computer is connected to a routed internal network and you are unsure of the routing topology of your network or how to add static routes, you can manually construct the LAT to contain the range(s) of IP addresses that your internal clients use.

   - Because a default gateway cannot be set on the internal interface of the ISA Server–based computer, you need to create static routes for your internal network at a later time to achieve full connectivity. You can do this using the **route** command from the command prompt.

   - A LAT that is configured correctly ensures that ISA Server can determine which network adapter to use, in order to access different portions of your internal network. If the LAT is not configured correctly, a client request for an internal IP address might be incorrectly routed to the Internet or redirected through the firewall service.

7. After installation, ISA Server effectively blocks all communication between your internal network and the Internet. No communication can occur until you configure an access policy with the protocol, site, and content rules specifically allowing access. Similarly, you must configure publishing rules if you want to allow Internet clients access to computers on your internal network.

   If you installed ISA Server as an array member, then an enterprise policy might be applied to the array. In this case, ISA Server might allow communication if the enterprise policy is configured appropriately.

8. Install the latest service pack for ISA Server.

# Additional Resources for Deploying ISA Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Deploying Dial-up and VPN Remote Access Servers" in this book for more information about creating a virtual private network (VPN).
- The Migrating from Proxy Server 2.0 link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about migrating from Proxy Server 2.0 to ISA Server.
- The Internet Engineering Task Force link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about RFC documents and Internet Engineering Task Force (IETF) Internet-Drafts.
- The ISA Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about extending ISA Server.
- The Hardware Compatibility List (HCL) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about hardware compatibility.

# Deploying IPSec

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Internet Protocol security (IPSec) in the Microsoft® Windows® Server 2003 operating system protects networks from active and passive attacks by securing IP packets through the use of packet filtering, cryptography, and the enforcement of trusted communication. IPSec is useful in host-to-host, virtual private network (VPN), site-to-site (also known as gateway-to-gateway or router-to-router), and secure server scenarios, and can be managed by using Group Policy or scripted by using command-line tools.

**In This Chapter**


Overview of IPSec Deployment

Determining Your IPSec Needs

Designing an IPSec Policy

Assigning IPSec Policies

Deploying Your IPSec Solution

Additional Resources for Deploying IPSec


**Related Information**

- For information about Windows Server 2003 IPSec, see the see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Overview of IPSec Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Private networks face various external and internal threats. You can use specific tools to counter specific threats. For example, antivirus systems can protect desktops and servers to a limited extent against known attacks.

IPSec can be used to secure all communications between parties — for example, by using authentication and encryption — without requiring any further modifications to applications or protocols. If you need to ensure that messages are not modified in transit or that they are unreadable to network intruders, IPSec provides a solution you can use to achieve these ends. Although IPSec can be applied to a variety of situations, this chapter describes solutions that are useful for most, but not all environments.

Although businesses typically separate their internal network from the Internet by using firewalls that block traffic sent to specified ports and protocols, internal corporate networks have become so complex that it is difficult to protect all mission-critical data at all times from attackers. Yet business applications fully depend on network access, so using a firewall to block incoming traffic by port is not always practical; it is difficult to know which ports require protection and which ports carry essential communication. IPSec provides a much stronger access model based on enforcing trusted communication.

Designed by the Internet Engineering Task Force (IETF), Internet Protocol security (IPSec) is a framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services.

IPSec is not a full-featured host firewall. However, it does provide the ability to centrally manage policies that can permit, block, or negotiate security for unicast IP traffic, based on specific addresses, protocols, and ports. Microsoft Internet Connection Firewall (ICF) is designed as a locally managed basic host firewall. The primary difference between IPSec and ICF is that IPSec provides complex static filtering based on IP addresses, while ICF provides stateful filtering for all addresses on a network interface. It is recommended that you use ICF when you want to implement a firewall for a network interface that can be accessed through the Internet. It is recommended that you use IPSec when you want to secure traffic on the wire, or when you need to allow access only to a group of trusted computers.

IPSec is included in Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; Windows® Server 2003, Web Edition operating systems; and the 64-bit editions of Windows Server 2003. Portions of IPSec and related services in these platforms were jointly developed by Microsoft and Cisco Systems, Inc. This chapter primarily focuses on design decisions and planning required to use IPSec for end-to-end secure networking scenarios for Windows Server 2003 and Microsoft® Windows® XP operating systems.

To best deploy IPSec in a typical network environment, you also need Active Directory® directory service domains and Group Policy in place.

Focus on the following areas when planning an IPSec deployment:

- Deciding where and how you must secure computers on your network by grouping computers in Active Directory organizational units (OUs).

- Deciding how tight to make the security by what policy you assign.

# IPSec Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before deploying IPSec on your network, make decisions about which computers to secure, where the security should be implemented, and how tight the security must be. Then test your deployment plan in a lab, and finally deploy IPSec on your network.

Use the steps shown in Figure 6.1 to design the deployment process of IPSec on your network.

**Figure 6.1 Deploying IPSec**



# Windows Server 2003 IPSec Concepts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the Windows Server 2003 implementation of IPSec to compensate for the limited protections provided by applications for network traffic, or as a network-layer foundation of a defense-in-depth strategy. Do not use IPSec as a replacement for other user and application security controls, because it cannot protect against attacks from within established and trusted communication paths. Your authentication strategy must be well defined and implemented for the potential security provided by IPSec to be realized, because authentication verifies the identity and trust of the computer at the other end of the connection. For example, if IPSec is required to encrypt all traffic only to trusted domain member computers, but there are weak controls on how computers can join the domain, the IPSec-protected traffic might be exposed to a trusted man-in-the-middle attack. Similarly, if attackers can join their computers to the domain, they might obtain trusted, encrypted access to IPSec-protected servers. If computers cannot be properly authenticated, IPSec does not provide security.

Table 6.1 outlines the specific security needs that IPSec addresses.

**Table 6.1 IPSec Security Solutions**

| Security Concern | IPSec Solution |
|---|---|
| Network attacks using specific protocols or ports. | Two methods:<br>• IPSec traffic blocking, to block specified traffic.<br>• IPSec policy filter lists can allow traffic from only trusted senders over particular protocols to particular addresses and ports. |
| Eavesdropping. Also known as *sniffing*. Eavesdropping occurs when an attacker uses a frame capture program (also known as a sniffer) to view the data that is placed on the network. | IPSec uses the Encapsulating Security Payload (ESP) protocol to encrypt data by using Triple Data Encryption Standard (3DES) or Data Encryption Standard (DES) so that IP packets cannot be read if intercepted in transit. |
| Data modification. If an intruder can access the information, they might be able to change it in such a way that the recipients cannot detect that any change has occurred. | IPSec uses a cryptographic checksum that incorporates a secret key to provide data integrity. Although a packet can be modified, the checksum cannot be updated for the proper value without knowledge of the secret key. An invalid checksum indicates to the receiving computer that the packet was modified in transit. |
| Identity spoofing. Intruders falsely using an IP address to compromise packet filter security. | IPSec uses Kerberos, public key certificates, or preshared key authentication to verify the identity of computer systems before the application level communication can take place. |
| Denial-of-service attack. By flooding a computer or network with malicious, malformed, or useless information, an attacker can disrupt normal operations. | IPSec uses IP packet filtering to determine whether communication is allowed, secured, or blocked, according to specified IP address ranges, protocols, or TCP or User Datagram Protocol (UDP) ports. IPSec cannot be targeted by denial-of-service attacks. |

When communication is secured between computers, IPSec authenticates computer identities and negotiates security options. Security options include aspects such as encryption strength or acceptable authentication methods. Negotiating security options ensures that communication partners agree how to secure the connection. By requiring IPSec for all network access to a server, you can ensure that only trusted computers can access the server. Active Directory can be used to deliver Group Policy security settings for all clients in a domain, as well as for groups of servers in an OU. Thus you can create, distribute, and enforce access rules for groups of computers for any of the following:

- Computers that are members of a specific domain

- A specific group of computers with common needs or roles in an OU
- A specific computer

**Compatibility**

Operating systems older than Microsoft® Windows® 2000 do not provide built-in support for IPSec. These include Microsoft® Windows® 98, Windows® Millennium Edition, and Microsoft® Windows NT®. If you have computers running these operating systems in your environment, make sure they are not required to use IPSec because the enforcement of IPSec-secured communications denies them access to resources.

Several features in the Windows Server 2003 implementation of IPSec are not provided in Windows 2000 or in Microsoft® Windows® XP Professional. If you plan to apply IPSec policies that use any of the new features that are available only in the Windows Server 2003 implementation of IPSec, do not use the Windows XP or Windows 2000 IP Security Policy Management snap-ins to manage these policies. Using those operating systems to manage Windows Server 2003 policies will result in the new features in the Windows Server 2003 policies being lost. IPSec policies that do not use the new Windows Server 2003 features, however, can be applied to computers running any of the three operating systems.

Remote management and monitoring of IPSec is only supported between computers running the same version of the Windows operating system. For remote management and monitoring of IPSec, it is recommended that you use Remote Desktop Connection.

For central management of IPSec policy on computers running Windows XP, Windows 2000, and Windows Server 2003, use Active Directory.

The IPSec internal infrastructure components were significantly modified for Windows Server 2003. Policies created on computers running Windows XP, Windows 2000, and Windows Server 2003 can be shared between any of those operating systems, but many tools have changed for creating and managing those policies and some tools are not installed by default. In Windows Server 2003, use the IP Security Policy Management snap-in, the Netsh IPSec context, and the Resultant Set of Policy (RSoP) snap-in for management and the IP Security Monitor snap-in for monitoring. Netdiag.exe, which is included in the Support Tools folder of the Windows Server 2003 operating system CD, no longer includes IPSec functionality. IPSecmon.exe, which was included in Windows 2000, is not included with Windows Server 2003 (it has been replaced by the new IP Security Monitor snap-in). IPSecpol.exe, which was included with the *Microsoft® Windows® 2000 Server Resource Kit*, is not included with Windows Server 2003. IPSeccmd.exe, which was included in the Support Tools folder of the Windows XP operating system CD, is not included with Windows Server 2003. For Windows Server 2003, the Netsh IPSec context replaces and enhances the functionality provided by Netdiag.exe and IPSeccmd.exe.

**New in Windows Server 2003**

IPSec provides new features that enhance security, scalability, availability, ease of deployment, and administration. Although the additions to Windows Server 2003 provide greater security, they are not always compatible with earlier versions of Windows. Ensure that you understand which technologies work with clients in your environment, and then determine whether to upgrade clients, postpone the use of specific IPSec technologies, or a mix of both. Then test your plan in a lab before deploying it in your production environment.

The new technologies highlighted in this section focus on those pertinent to deployment.

**Filters update IP configuration of IPSec partners** You can use the IP Security Policy Management snap-in or Netsh to configure the source or the destination address fields that the target computer interprets as the addresses for the Dynamic Host Configuration Protocol (DHCP) server, the Domain Name System (DNS) servers, the Windows Internet Name Service (WINS) servers, and the default gateway. As a result, IPSec policies can now automatically accommodate changes in the IP configuration of the source or destination, by using either DHCP or static IP configurations. Enhanced subnet definitions are also supported by Windows Server 2003 IPSec.

**NAT traversal support has been added** Windows Server 2003 IPSec meets new Internet Engineering Task Force (IETF) draft version 2 specifications, "UDP Encapsulation of IPSec Packets" and "Negotiation of NAT-Traversal in the IKE," which describe methods by which ESP-protected IPSec traffic can pass through a network address translator (also known as a NAT). However, some applications might not work when their traffic is first protected with IPSec ESP and then passed through a network address translator. The Internet Key Exchange (IKE) protocol automatically detects network address translators and uses UDP ESP encapsulation to send all user data by using UDP port 4500. Using the Authentication Header (AH) protocol across network address translators is not supported by IPSec NAT traversal (NAT-T).

**Netsh enables command-line support for IPSec** Replacing IPSecpol.exe and IPSeccmd.exe, the new IPSec context of the Netsh command-line tool can be used to automate IPSec configuration, deploy IPSec remotely or across many computers, and perform other tasks. For example, domain logon scripts, startup scripts, and other files can be used to configure IPSec policy.

For more information, see "Netsh commands for Internet Protocol Security (IPSec)" in Help and Support Center for Windows Server 2003.

**Default exemptions have been removed** IPSec provides greater protection because filtering exemptions that existed in Windows XP and Windows 2000 have been removed from Windows Server 2003. In Windows XP and Windows 2000, all broadcast, multicast, Internet Security Association and Key Management Protocol (ISAKMP), Kerberos, and Resource Reservation Protocol (RSVP) traffic was exempt from IPSec filtering. In Windows Server 2003, only ISAKMP traffic is exempt from IPSec filtering. Exemptions can be added by using the IP Security Policy Management snap-in or the Netsh IPSec context. The Kerberos protocol is no longer a default exemption; if you want to enable Kerberos authentication, you must create filters in the IPSec policy that explicitly allow all traffic to each domain controller IP address in the domain member's domain. If you created policies for Windows XP or Windows 2000, you might need to configure explicit permit filters to maintain the expected behaviors of the IPSec policy on servers running Windows Server 2003. For more information, see "Understanding IP Filters, Filter Actions, and Filter Lists" later in this chapter.

**RSoP support has been added** RSoP provides information that can help assess and resolve unexpected IPSec behavior resulting from Group Policy conflicts. For more information, see "Using RSoPd to Analyze IPsec Policy Assignments" later in this chapter.

Other improvements that enhance security and IPSec policy management include:

- 2048-bit Diffie-Hellman keys, which provide very strong cryptographic key generation to maximize 3DES encryption strength.
- Persistent IPSec policy, which enhances security during computer startup. Additionally, persistent IPSec policy provides a method for local administrators to force local IPSec policy settings to be applied when Active Directory-based policy is assigned. For more information about persistent IPSec policy, see "Assigning IPSec Policies Locally" later in this chapter.
- Stateful filtering of network traffic during computer startup, which permits only the following traffic during computer startup: the outbound traffic that the computer initiates during startup, the inbound traffic that is sent in response to the outbound traffic, and DHCP traffic.
- The IP Security Monitor snap-in, which displays more details about IPSec communication than IPsecmon.exe (used for monitoring IPSec in Windows 2000). These additional details and performance enhancements enable effective monitoring of IPSec-secured servers.

For more information about new features of Windows Server 2003, see "New features for IPSec" in Help and Support Center for Windows Server 2003.

#### Terms and Definitions

This section provides brief definitions of core IPSec and IPSec-related terms. For more information, see the sources mentioned in "Additional Resources for Deploying IPSec" later in this chapter, or Help and Support Center for Windows Server 2003.

**Authentication** The process for verifying that a peer or object is who or what it claims to be. Examples include confirming the source of information, such as verifying a digital signature, or verifying the identity of a computer.

**Authentication Header (AH)** An IPSec protocol that provides data origin authentication, data integrity, and anti-replay for the entire packet (the IP header and the data payload carried in the packet, except fields in the IP header that are allowed to change in transit). AH can be used alone, in combination with ESP, or in IPSec tunnel mode. Windows AH functionality complies with RFC 2402.

**Cryptography** The process of keeping messages and data secure. Cryptography enables and ensures data confidentiality, data integrity, authentication (peer and data origin), and nonrepudiation.

**Default response** An IPSec rule that is used to ensure that the computer responds to requests for secure communication. If an active policy does not have a rule defined for a computer that is requesting secure communication, the default response rule is applied and security is negotiated. The default response rule cannot be deleted, but it can be deactivated. It is activated by default for all policies.

**Encapsulating Security Payload (ESP)** An IPSec protocol that provides data confidentiality, data origin authentication, data integrity, and anti-replay for the ESP payload. ESP can be used alone, in combination with AH, or in IPSec tunnel mode. Windows ESP functionality complies with RFC 2406.

**Filter action** The part of an IPSec rule that defines the security requirements for the data transmission. A filter action can be

configured to permit traffic, block traffic, or negotiate secure communications with IPSec. If security negotiation is selected, you must also configure security methods and their order, whether to accept initial incoming unsecured traffic, whether to allow unsecured communication with computers that do not support IPSec, and whether to use perfect forward secrecy (PFS).

**Group Policy** The infrastructure within Active Directory that enables directory-based change and configuration management of user and computer settings, including security and user data. You use Group Policy to define configurations for groups of users and computers. The Group Policy settings that you create are contained in a Group Policy object (GPO). By associating a GPO with selected Active Directory system containers — domains, sites, and OUs — you can assign the GPO's policy settings to the users and computers in those Active Directory containers. Although Group Policy applies to both users and computers, IPSec policy is a computer configuration Group Policy setting that is only applied to computers.

**Group Policy Object (GPO)** A collection of Group Policy settings that are stored at the domain level that affect users and computers contained in domains, sites, and OUs. In addition, each computer has one local GPO. To create an individual GPO, use the Group Policy Object Editor snap-in. Although GPOs can apply to both users and computers, IPSec policy objects can only be applied to computers.

**IPSec authentication** The process of verifying trust in an identity. IKE uses mutual authentication between computers to establish trusted communications, and requires the use of one of the following authentication methods: Kerberos V5, a computer X.509 v3 public key infrastructure (PKI) certificate, or a preshared key. The two communication endpoints must have at least one common authentication method, or communication fails.

**IPSec filter** A specification in an IPSec rule that is used to match IP packets to filter actions (permit, block, or negotiate security). An IP filter list contains one or multiple filters, and it can be shared among different IPSec policies. For example, one filter list can contain the list of each domain controller IP address in a domain.

**IPSec policy** A collection of one or more IPSec rules that determine IPSec behavior.

**IPSec rule** A statement in an IPSec policy that associates a filter list with a filter action, an authentication method, and an IPSec mode. An IPSec rule is typically configured for a specific purpose (for example, "Block all inbound traffic from the Internet to TCP port 135"). Many IPSec rules can be defined in a single IPSec policy.

**IPSec transport mode** A mode of IPSec that is used to protect host-to-host communications. Transport mode provides security between computers that are on the same local area network (LAN) or connected by private wide area network (WAN) links.

**IPSec tunnel mode** A mode of IPSec that is used to protect site-to-site (gateway-to-gateway) traffic between networks, such as site-to-site networking through the Internet. The sending gateway encapsulates the entire IP datagram by creating a new IP packet that is then protected by one of the IPSec protocols. Windows Server 2003 supports IPSec tunnel mode for configurations where Layer Two Tunneling Protocol (L2TP) cannot be used. For more information about recommended uses of tunnel mode, see "Using IPSec in Tunnel Mode" later in this chapter.

**Netsh** A command-line scripting tool for configuring networking components on the local computer or on remote computers running Windows 2000, Windows XP Professional, or Windows Server 2003. The Netsh IPSec context is only available on Windows Server 2003, although policies created with Netsh can be used for computers running other versions of Windows that support IPSec.

**Perfect forward secrecy (PFS)** A mechanism that determines whether the existing keying material for a master key can be used to derive a new session key. Session key PFS performs a new Diffie-Hellman key exchange to generate new master key keying material instead of using master key keying material to derive more than one session key.

**Tunnel endpoint** The tunneling computer that is closest to the IP traffic destination, as specified by the associated IP filter list. Two rules are required to describe an IPSec tunnel. For the outbound traffic rule, the tunnel endpoint is the IP address or subnet of the IPSec peer on the other end of the tunnel. For the inbound traffic rule, the tunnel endpoint is an IP address or subnet configured on the local computer.

# Determining Your IPSec Needs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in deploying IPSec on your network is to determine which computers need to be secured. Some data on your network might need higher levels of security than others. Although adding IPSec can protect your network and data, it can also reduce performance, and some computers might not support IPSec. As shown in Figure 6.2, you must decide which data and which computers on your network you need to secure and at what level of security.

**Figure 6.2 Determining Your IPSec Needs**




When beginning your design process, make sure to have specifications about your current network environment available for use. Your hardware and software inventory and a map of network topology are essential tools for the design phase. Because IPSec policies are based on which computers and data you want to secure, where that data is stored within the network, and the operating systems you need to support, this information lets you know where you can deploy IPSec. For more information about creating those documents, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

Some network environments are well suited to IPSec as a security solution and others are not. This section lists some common examples.

**Recommended IPSec Uses**

IPSec is recommended for the following uses:

- **Packet filtering.** IPSec provides limited firewall capabilities for end systems. You can use IPSec with the NAT/Basic Firewall component of the Routing and Remote Access service to permit or block inbound or outbound traffic. You can also use IPSec with ICF, which provides stateful filtering. However, to ensure proper IKE management of IPSec security associations (SAs), you must configure ICF to permit ISAKMP UDP port 500 traffic.

- **Securing host-to-host traffic on specific paths.** You can use IPSec to provide protection for traffic between servers or other static IP addresses or subnets. For example, IPSec can secure traffic between domain controllers in different sites or between Web servers and database servers.

- **Securing traffic to servers.** You can require IPSec protection for all client computers that access a server. In addition, you can set restrictions on which computers are allowed to connect to a server running Windows Server 2003.

- **L2TP/IPSec for VPN connections.** You can use the combination of the L2TP and IPSec (L2TP/IPSec) for all VPN scenarios. This does not require the configuration and deployment of IPSec policies.

- **Site-to-site (gateway-to-gateway) tunneling.** You can use IPSec in tunnel mode for site-to-site (gateway-to-gateway) tunnels when you need interoperability with third-party routers, gateways, or end systems that do not support L2TP/IPSec or Point-to-Point Tunneling Protocol (PPTP) connections.

**Note**

- Because IPSec depends on IP addresses for establishing secure connections, you cannot specify dynamic IP addresses. It is often necessary for a server to have a static IP address in IPSec policy filters. In large network deployments and in some mobile user cases, using dynamic IP addresses at both ends of the connection can increase the complexity of IPSec policy design.

**IPSec Uses That Are Not Recommended**

IPSec can reduce processing performance and increase network bandwidth consumption. Additionally, IPSec policies can be quite complex to configure and manage. Finally, the use of IPSec can introduce application compatibility issues. For these reasons, IPSec is not recommended for the following uses:

- **Securing communication between domain members and their domain controllers.** In addition to reduced network performance, using IPSec for this scenario is not recommended because of the complexity of the IPSec policy configuration and management required.

- **Securing all traffic in a network.** In addition to reduced network performance, using IPSec for this scenario is not recommended because:

    - IPSec cannot negotiate security for multicast and broadcast traffic.

    - Traffic from real-time communications, applications that require Internet Control Message Protocol (ICMP), and peer-to-peer applications might be incompatible with IPSec.

    - Network management functions that must inspect the TCP, UDP, and protocol headers are less effective or cannot function at all, due to IPSec encapsulation or encryption of IP payloads.

Additionally, the IPSec protocol and implementation have characteristics that require special consideration in the following scenarios:

- **Securing traffic over wireless 802.11 networks.** You can use IPSec transport mode to protect traffic sent over 802.11 networks. However, IPSec is not the recommended solution for providing security for corporate 802.11 wireless LAN networks. Instead, it is recommended that you use 802.11 Wired Equivalent Privacy (WEP) encryption and IEEE 802.1X authentication. Support for IPSec, configuration management, and trust are required on client computers and servers. Because many computers on a network do not support IPSec or are not managed, it is not appropriate to use IPSec alone to protect all 802.11 corporate wireless LAN traffic. Additionally, IPSec tunnel mode policies are not optimized for mobile clients with dynamic IP addresses, nor does IPSec tunnel mode support dynamic address assignment or user authentication, which is needed for remote access VPN scenarios. To secure remote access traffic to organization networks when that traffic is sent over public wireless networks that are connected to the Internet, use L2TP/IPSec VPN connections.

- **Using IPSec in tunnel mode for remote access VPN connections.** Using IPSec in tunnel mode is not recommended for remote access VPN scenarios. Instead, use L2TP/IPSec or PPTP for remote access VPN connections.

# Using IPSec in Transport Mode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Table 6.2 provides a summary of when IPSec transport mode is appropriate to use.

**Table 6.2 IPSec Transport Mode Uses**

| IPSec Scenario | IPSec Transport Mode Usage |
|---|---|
| Require packet filtering | Although IPSec does not provide full firewall functionality, it can be used to statically permit or block traffic based on source and destination address combinations, and based on the IP protocol and TCP and UDP ports. Some of the functions found in standard firewalls that IPSec does not provide include stateful inspection, application protocol awareness, intrusion inspection, and packet logging. Although IPSec lacks some features of firewalls, the packet blocking and filtering it provides can be effective in limiting the spread of viruses or in thwarting specific attacks known to use specific ports. It can also be used to prevent specific applications and services from being used on the network. |
| Require end-to-end security | This is the easiest way to secure traffic with IPSec, typically between servers or clients and servers. End-to-end security creates a secure channel for trustworthy communication. |

You can also combine these two IPSec methods to enhance the network-level security of your infrastructure. Both methods are detailed later in this chapter.

# IPSec Packet Filtering

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use Windows Server 2003 IPSec to secure specific servers in your enterprise. IPSec can be configured to permit or block specific types of traffic based on source and destination address combinations and specific protocols and specific ports. You can use this feature of IPSec to block well-known ports of software so that even if a server is infected, it cannot be used to infect other computers or allow access via the well-known port.

For example, nearly all the systems illustrated in Figure 6.3 can benefit from packet filtering to restrict communication to only specific addresses and ports.

**Figure 6.3 Filtering Packets by Using IPSec**




By blocking communication to specific ports of a server, it is more difficult for an intruder to access the server or for the server to be used to attack other computers. You can strengthen security by using IPSec filtering to control exactly the type of communication that is allowed between systems. For example, as illustrated in Figure 6.3:

- The internal network domain administrator can assign a domain-based IPSec policy to block all traffic from the perimeter network.
- A perimeter network domain administrator can assign a domain-based IPSec policy to block all traffic to the internal network.
- The administrator of the computer running Microsoft® SQL Server™ on the internal network can create an exception to the domain-based IPSec policy to permit SQL protocol traffic to the Web application server on the perimeter network.
- The administrator of the Web application server on the perimeter network can create an exception to the domain-based policy to permit SQL traffic to the computer running SQL Server on the internal network.
- The administrator of the Web application server on the perimeter network can also block all traffic from the Internet, except requests to TCP port 80 (HTTP) and TCP port 443 (SSL), which are used by Web services. This provides additional security against traffic allowed in from the Internet in case the firewall was mis-configured or compromised by an attacker.
- The domain administrator can block all traffic to the management computer, but allow traffic to the perimeter network.

You can also use IPSec with the NAT/Basic Firewall component of the Routing and Remote Access service or IP packet filtering to enhance IPSec permit or block filtering of inbound or outbound traffic.

# End-to-End Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

IPSec end-to-end security establishes trust and security from a unicast source IP address to a unicast destination IP address (peer-to-peer). You can use this IPSec capability for secured server scenarios in which only trusted computers are allowed access to a server. In secured server scenarios, you can use IPSec to control access to all applications and services running on the server, and then either encrypt or only simply authenticate all application network traffic. The computers (IPSec peers) are authenticated through the use of Kerberos V5, certificate-based public key authentication, or — only if necessary — preshared key authentication.

**Peer-to-Peer Communication**

As shown in Figure 6.4, only the sending and receiving computers need to be aware of IPSec. Each handles security at its respective end and assumes that the medium over which the communication takes place is not secure. The two computers can be located near each other, as on a single network segment, or across the Internet. Computers or network elements that route data from source to destination are not required to support IPSec. A firewall between IPSec peers must be configured to forward IPSec traffic on UDP source and destination port 500, IP protocol 50 (ESP), or IP protocol 51 (AH). If network address translation is taking place between two computers, they must both support IETF IPSec NAT-T specification draft version 2. A firewall must be configured to forward traffic on UDP source and destination port 4500.

**Figure 6.4 Peer-to-Peer Security in IPSec**




End-to-end security is typically used to secure data sent over connections that transmit sensitive data or that might be susceptible to eavesdropping, such as:

- Servers within a corporate network.
- Servers over the Internet. For example, end-to-end security can secure traffic used by business-to-business application systems, protect file and database content as it is distributed to your business partners, or encrypt Remote Authentication Dial-In User Service (RADIUS) traffic between RADIUS clients, proxies, and servers.
- Front-end servers in a data center or perimeter network and servers inside the corporate internal network. There are two methods for securing data in this scenario:
    - End-to-end IPSec transport protection of traffic between the servers. Use IPSec to guarantee authenticity and unmodified traffic between servers, other computers using static IP addresses, or subnets. For example, IPSec can secure traffic between domain controllers, between routers connecting remote sites or forests, or between Web servers and database servers. You can use encryption if the traffic does not require inspection by an intrusion detection system between the servers.
    - End-to-gateway IPSec protection of traffic, so that the traffic is secured with IPSec, but can be fully inspected or modified by a firewall. Depending on the gateway, you can use IPSec transport mode or IPSec tunnel mode.

For any VPN scenarios, you can use the combination of L2TP and IPSec (L2TP/IPSec). For compatibility with other routers and gateways, or end-systems that do not support L2TP/IPSec or PPTP connections, you can use IPSec in tunnel mode for site-to-site (gateway-to-gateway) tunnels.

**Securing an Application Server**

Figure 6.5 shows end-to-end IPSec in transport mode securing an application server.

**Figure 6.5 Using IPSec to Secure an Application Server**




In this scenario, an application server in an internal corporate network must communicate with clients running Windows 2000 or Windows XP Professional; a WINS server, DNS server, and DHCP server; Active Directory domain controllers; and a third-party data backup server. Typically, an administrator uses Group Policy to configure and assign an Active Directory-based IPSec policy to an application server. In this scenario, the IPSec policy specifies that traffic is secured or permitted for different data paths, as described below.

**Secured traffic**

The users on the client computers are company employees who access the application server to view their personal payroll information and performance review scores. Because the traffic between these clients and the application server involves highly sensitive data, and because the administrator wants the server to only communicate with other domain members, he or she uses an IPSec policy that requires ESP encryption and communication only with trusted computers in the Active Directory domain.

Additionally, in Windows Server 2003, a specific group of computers can be authorized to use IPSec when either Kerberos V5 or certificates are used for IKE authentication. This capability is known as *certificate-to-account mapping*. When IPSec certificate-to-account mapping is used, an administrator can control which computers are authorized to use IPSec by configuring Group Policy security settings and assigning either the **Access this computer from the network** or the **Deny access to this computer from the network** logon right to individual or multiple computers, as needed. For more information, see "Public Key Certificate" later in this chapter.

**Permitted traffic**

Traffic is permitted for the following data paths:

- Traffic between the WINS server, DNS server, DHCP server, and the application server is permitted because WINS servers, DNS servers, and DHCP servers must typically communicate with computers that run on a wide range of operating systems, some of which might not support IPSec.

- Traffic between Active Directory domain controllers and the application server is permitted, because using IPSec to secure communication between domain members and their domain controllers is not a recommended usage due to the complexity of the IPSec policy configuration and management required in Active Directory.

- Traffic between the third-party data backup server and the application server is permitted because the third-party backup server does not support IPSec.

**Example: Contoso's Use of End-to-End Security**

The fictitious Contoso Corporation implemented a mix of IPSec solutions to meet its end-to-end security needs. Using IPSec end-to-end security makes sense for Contoso because it has branches throughout the world that need to securely exchange data; messages need to pass securely and smoothly across the Internet. Also, Contoso might need to prevent network access to its servers from attackers who might penetrate their network, either through the perimeter network or by directly using a network port in any Contoso office.

Contoso considered using IPSec end-to-end security on all communications, but quickly realized that this was unnecessary and

instead chose to authenticate network access to a select set of computers, and to encrypt only its most sensitive communications. Contoso uses IPSec on traffic sent between the following computers or locations:

- Servers within their network, such as between their Distributed File System (DFS) shares and all clients. Contoso used Kerberos authentication for these connections.

- Contoso headquarters servers and partner sites at which Contoso products are manufactured. In the case of the partner connections, Contoso also restricted which partner computers can connect to their corporate network. In this case, Contoso used public key certificate authentication because they shared common root certification authorities (CAs) with their partners. Their partners obtained computer certificates from third-party PKI providers. Contoso allowed only IPSec traffic to pass through the firewall when communicating with partner networks. Contoso mapped computer certificates for partner networks to a shadow computer account in Active Directory, and specified that only computer accounts from the partner networks were allowed IPSec access to their servers. To do this, they configured the **Access this computer from the network** logon right on each of their servers and assigned this right to computer accounts in the partner networks.

- Web server access to the data center database on cluster servers. Contoso used certificate authentication to enable IPSec-protected connections because the Web servers did not have a security domain trust with the data center cluster servers. The IPSec policy on the data center servers allowed only inbound SQL traffic from the Web servers. To allow IPSec-protected access to Web servers that are managed by a contractor and that do not run the Windows operating system, Contoso used preshared key authentication.

By protecting these connections, Contoso protected high-value communications that were critical to their business success. If they had failed to properly protect order information, competitors might have gained valuable data, or orders might have been canceled or modified by attackers.

# Using IPSec in Tunnel Mode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Tunnel mode is primarily used for interoperability with gateways or end systems that do not support L2TP/IPSec or PPTP VPN site-to-site connections. Table 6.3 provides a summary of scenarios in which IPSec tunnel mode is appropriate to use.

**Table 6.3 IPSec Tunnel Mode Usage**

| IPSec Scenario | IPSec Tunnel Mode Usage |
|---|---|
| Establish gateway-to-gateway tunnels between sites, when the gateways or end systems do not support L2TP/IPSec VPN connections. | Use only when necessary for interoperability with gateways or end systems that do not support L2TP/IPSec VPN connections. When L2TP/IPSec VPN connections are supported, use L2TP/IPSec (IPSec transport mode) instead. When IPSec in tunnel mode is used for gateway-to-gateway tunnels, static IP addresses are required for the IPSec tunnel endpoints.<br><br>For an illustration of this scenario, see Figure 6.6. |
| Protect traffic end-to-end, but one endpoint of the communication does not support IPSec. | Use an IPSec tunnel from one endpoint to a firewall or IPSec-enabled router nearest to the endpoint that does not support IPSec. This scenario requires static IP addresses for both tunnel endpoints.<br><br>For an illustration of this scenario, see Figure 6.7. |
| Encrypt traffic end-to-end between two computers, but a third-party firewall or network intrusion detection system requires that traffic be decrypted. | Use one IPSec tunnel to reach the third-party firewall or the network intrusion detection system. Use another tunnel from the third-party firewall or network intrusion detection system to reach the destination. This scenario requires static IP addresses for all four tunnel endpoints.<br><br>For an illustration of this scenario, see Figure 6.8. |

**Gateway-to-Gateway Traffic Must be Secured**

In this scenario, traffic is being sent between a client computer in a vendor site (Site A) and a File Transfer Protocol (FTP) server at the corporate headquarters site (Site B). Although an FTP server is used for this scenario, the traffic can be any unicast IP traffic. The vendor uses a third-party gateway, while corporate headquarters uses a gateway running Windows Server 2003. An IPSec tunnel is used to secure traffic between the third-party gateway and the gateway running Windows Server 2003. Figure 6.6 shows IPSec in tunnel mode for gateway-to-gateway tunneling between sites.

**Figure 6.6 Gateway-to-Gateway Tunneling Between Sites**




**Note**

- For ease of deployment, it is recommended that the gateway running Windows Server 2003 be a stand-alone computer and that you remotely manage this computer by using Remote Desktop Connection. If this computer is joined to a domain, the computer must be able to access Active Directory locally, not through an IPSec tunnel.

**One Endpoint Does Not Support IPSec**

In this scenario, an IPSec tunnel is used to secure traffic between a computer running Windows Server 2003 in a perimeter network and an IPSec-enabled router (for example, a Cisco IOS router). Traffic on the path between the router and the third-party server in the internal corporate network is not secured, because the third-party server does not support IPSec. Note that this is not a VPN remote access scenario, because there is no dynamic address assignment, nor is user authentication used to establish the tunnel. The tunnel rules can be defined to protect traffic between the computer running Windows Server 2003 and either the internal network subnet or the specific IP address of the third-party server. Figure 6.7 shows IPSec in tunnel mode to provide end-to-end security of traffic when one endpoint of the communication does not support IPSec.

**Figure 6.7 One Endpoint Does Not Support IPSec**




**Traffic Must Be Decrypted for Third-Party Firewall Inspection**

In many cases, IPSec transport mode can provide secure end-to-end communication through firewalls. If traffic must be inspected at the firewall, you can use IPSec AH or ESP with null encryption in transport mode to maintain end-to-end security. However, if the traffic must be encrypted and inspected, you can use IPSec tunnel mode to secure traffic to the inspection point.

Figure 6.8 shows IPSec in tunnel mode encrypting traffic end-to-end between two computers: a line of business (LOB) application server running Windows Server 2003 in a perimeter network, and a computer running SQL Server that functions as a data store for the application server in an internal corporate network.

**Figure 6.8 Traffic Must Be Decrypted for Third-Party Firewall Inspection**




In this scenario, IPSec tunnels are used to secure traffic between the application server and the third-party firewall's external interface and between the third-party firewall's internal interface and the computer running SQL Server. Traffic is decrypted for firewall inspection and then reencrypted when it is forwarded to the IPSec peer. The tunnel rules can be defined to protect traffic between the application server and either the internal network subnet or the specific IP address of the computer running SQL Server.

**Note**

- This scenario cannot work if a computer running Microsoft® Internet Security and Acceleration (ISA) Server is used as the firewall. Use IPsec transport mode on the side where the ISA interface IP address is being used as the source or destination IP address.

For more information about using IPSec in tunnel mode, including detailed configuration procedures, see article Q252735, "How to Configure IPSec Tunneling in Windows 2000," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Weighing IPSec Tradeoffs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Whether you use packet filtering or end-to-end security, deploying IPSec provides greater security at the expense of network and processing performance and compatibility with other services, tools, and features. Use the information in this section to determine the impact on your environment before deploying IPSec.

For example, you might decide that information sent between some departments, such as Accounting and Legal, require higher security than information sent to or from other departments, such as Production. You might also decide that all private information transmitted across the Internet must be secured. Finally, you might decide that you only want trusted computers to have network access to your servers. When planning your deployment, be aware that deploying IPSec on your network increases the overall cost of network and system administration. The exact cost cannot be predicted because it depends on many factors, including the level of application traffic. However, if you stage IPSec deployment in a lab, you can arrive at a good estimate. Do not deploy IPSec if the security it provides is not required. Consider whether the added cost is worth the added protection.

**Slower connection establishment**
IPSec is not recommended for scenarios in which clients make infrequent, short-lived connections. IPSec increases the time required to initially establish authenticated connections. Depending on policy design, network roundtrip time, and the load on the systems required to establish the connection, negotiating security initially using IKE typically increases connection time by one to three seconds. Make sure that you want to trade the one- to three-second increase in connection time for the added security. Typically, after IPSec SAs are established, no further delays occur. Because IPSec is not integrated with TCP, TCP connections can be established and ended rapidly after IPSec SAs are established.

**Reduced computing performance**
Although the Windows Server 2003 IPSec filtering engine is enhanced to provide significantly faster packet processing, filtering and encryption of each packet consumes computing power. As the volume of IPSec traffic — including policies, rules, and filters — increases, the load placed on the network and on computer CPUs increases also. For example, policies with 1,000 or more filters can cause a small load increase. Filters with specific source and destination IP addresses can be processed much more quickly than filters that specify a subnet or **Any IP Address**. IPSec implements a cache algorithm that can accelerate the processing of packets that match a filter. Packets that do not match any filter require the greatest amount of time to process.

If your systems already have high loads on their CPUs, consider adding more computers or using IPSec offload network adapters before deploying IPSec. IPSec offload network adapters accelerate cryptographic operations used to secure IPSec packets. However, these adapters do not accelerate filter processing time. IPSec-enabled network adapters can be used with Windows 2000, Windows XP, and Windows Server 2003. Such network adapters typically complete IPSec operations at such a high rate that there is minimal change in transmission speeds, resulting in minimal performance degradation. However, such adapters generally have a maximum number of SAs that they can support, after which additional SAs are processed by the Windows IPSec components.

**Thwarting some networking inspection technologies**
ESP and AH encapsulate IP payloads by adding a security header to every packet. Existing network management and diagnostic tools typically cannot interpret IPSec-protected packets. The tools that might be impacted include:

- Router-based access control lists.
- Firewall access controls.
- Network traffic analysis and reporting tools.
- Quality of Service (QoS).
- Third-party load balancing devices. Windows Server 2003 IPSec supports the Windows Server 2003 Network Load Balancing service for host-based load balancing, which might provide an alternate solution.
- Network address translation. Windows Server 2003 supports IPSec NAT-T, which is used to allow IPSec ESP-protected traffic to traverse one or more network address translators.

- Stateful filtering controls on the direction of traffic flow. IKE and IPSec communication typically function best when traffic flows in both directions. IPSec does not easily allow for connections in only one direction.

- Network-based intrusion detection.

These functions require that network packets be inspected, modified, or both. If IPSec protocols are not used to encrypt packets, these functions can be enhanced to inspect inside the IPSec protocol header. If IPSec encryption is required for security, however, it might be difficult to maintain full network management functionality.

**Increased packet size reduces throughput, increases network utilization**
IPSec protection adds overhead to IP packets, reducing effective throughput and increasing network utilization. Although the IETF design and Microsoft implementation of IPSec attempt to minimize both impacts, existing networks and servers might be operating at maximum capacity without IPSec. Before deploying IPSec, test several different IPSec policy configurations, because each policy configuration is likely to result in different performance impacts on clients and servers. It might be necessary to increase the available network bandwidth or CPU power, or install IPSec offload adapters to compensate for the increased overhead of IPSec.

**Application incompatibility with IPSec NAT-T**
Windows Server 2003 supports IPSec NAT-T as described by version 2 of the IPSec NAT-T Internet drafts, but some applications might not work when their traffic is first protected with IPSec and then passed through a network address translator. Test how your planned implementation of IPSec interacts with network address translators in a lab before deploying IPSec in your environment.

**Loss of connectivity during cluster node addition or removal**
IPSec establishes and maintains cryptographic state between computers. Many clustering and load-balancing services use the same IP address for all nodes in a cluster, which creates incompatibilities with IPSec. Windows Server 2003 IPSec has proprietary extensions that allow it to work with the Windows Server 2003 Network Load Balancing service and Windows Cluster Service. However, support for these extensions does not exist in the current Windows 2000 and Windows XP IPSec client implementations, so you might experience some loss of connectivity when you add or remove cluster nodes.

**Requiring IPSec for communication between Active Directory domain members and domain controllers might block connections**
IPSec is based on the authentication of computers on a network; therefore, before a computer can send IPSec-protected data, it must be authenticated. The Active Directory security domain provides this authentication using the Kerberos protocol. Accordingly, when IKE uses Kerberos to authenticate, the Kerberos protocol and other dependent protocols (DNS, UDP LDAP and ICMP) are used for communication with domain controllers. Additionally, Active Directory–based IPSec policy settings are typically applied to domain members through Group Policy. As a result, if IPSec is required from domain members to the domain controllers, authentication traffic will be blocked and IPSec communications will fail. In addition, no other authenticated connections can be made using other protocols, and no IPSec other policy settings can be applied to that domain member through Group Policy. For these reasons, using IPSec for communications between domain members and domain controllers is not supported.

**ICMP-based functionality might fail**
When designing an IPSec policy to secure or block ICMP traffic, keep in mind that such a policy might cause services and tools that rely on ICMP to measure network response times to produce misleading results. For example, Windows 2000 and Windows XP domain members have an IP/DNS-compatible locator, which measures ICMP response times for domain controller load balancing within a site. Also, Group Policy measures ICMP response times to determine if the connection to the domain controller is a slow link. Also, tools that depend on ICMP, such as Tracert, might not work when ICMP traffic is protected by IPSec.

**Multicast and broadcast traffic might fail**
In Windows Server 2003, you can configure IPSec filters to block multicast and broadcast traffic. Therefore, any application that uses multicast or broadcast might fail if the computer has an IPSec policy that blocks such traffic.

# Examples: IPSec in a Corporate Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The fictitious Contoso Corporation evaluated their security needs when considering whether to implement IPSec. Although there are many tools to counter threats in a computing environment, this example focuses on security concerns and solutions that relate to IPSec. To begin this process, administrators created a map of the hardware and software in their environment. Notable aspects of an environment include domain membership, operating system versions, computer roles, and the location of computers on the network.

Some computers have relatively low security needs or are very difficult to manage. Although it might be advantageous to secure all data, using IPSec demands administrative resources on both the client side and the server side of the communication. Accordingly, using ICF and anti-virus tools protects computers that are difficult to manage or that have relatively low security needs. Contoso relied on firewall functionality provided by a perimeter network, and a proxy server for Internet access from their internal network. Contoso specified a default client configuration consisting of computers running Windows 2000 and Windows XP. Client computers are configured as follows:

- The client computers and users have accounts in the corporate Active Directory security domain. An enterprise CA is installed on a server running Windows Server 2003, to provide automatic certificate enrollment for user and computers.
- An antivirus program is used to examine all incoming and outgoing files and e-mail.
- Encrypting File System (EFS) encrypts sensitive documents on the client computers' hard disk.
- Services and applications use Kerberos authentication when it is available.
- All domain members are assigned Group Policy settings that do the following:
    - Automatically enroll with the CA to obtain a computer certificate. This certificate can be used to gain remote access to the corporate network when the computer is connected to the Internet for L2TP/IPsec client VPN connections. It is also used for IPSec when Kerberos authentication is not available.
    - ICF is enabled by default for all network connections on portable computers running Windows XP. However, domain administrators disable ICF on portable computers when those computers are connected to the corporate network. ICF is disabled so that users on can share files and allow inbound access to other applications on their desktop.
    - Enforce an IPSec policy to negotiate security with specific IPSec-protected servers and to block ports that are used by known viruses and prohibited applications.

Three examples of how Contoso used IPSec on their servers are as follows:

- The Contoso IT group used a set of servers to receive order information from customers, and decided that they needed to ensure that the orders were received only from their partners and were not modified in transit. Because data authenticity and integrity were important, but confidentiality was not, they required AH protection for customers to connect with these computers. IPSec ESP with null encryption was allowed for customers that must traverse a network address translator to make the IPSec-protected connection to Contoso. Contoso allowed customers to use certificates from a specific list of third-party PKI providers (public PKI root CAs).
- The Contoso Human Resources department stored information about employee salaries and medical claims. The department carefully controlled access with user security settings. But it was also necessary to meet health care regulations to maintain the privacy of this traffic as it flowed over the network. Further, Contoso wanted a defense-in-depth strategy to protect against attackers who might gain remote access to their servers over the network, through the use of compromised employee user IDs and passwords. IPSec allows Contoso to restrict client network access only to computers that are members of the domain. If necessary, access can be further restricted to a specific group of computers in the domain. Accordingly, to secure their servers, Contoso used an IPSec policy to require ESP encryption for

domain members to connect.

Because Windows Server 2003 IPSec can detect the presence of network address translators and automatically use UDP headers to allow IPSec traffic to traverse the network address translators, Contoso was able to discontinue using PPTP for remote access VPN connections. For enhanced security, administrators switched all remote users to L2TP/IPSec VPN connections using L2TP/IPSec NAT-T-capable clients, newly available from Microsoft. For more information, see the Virtual Private Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.Use the job aid "Designing an IPSec Policy" (DNSIPS_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing an IPSec Policy" on the Web at http://www.microsoft.com/reskit) to record information about your environment. Use a separate copy of the worksheet for each policy you create to meet the security needs of each group of servers.

# Designing an IPSec Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to create a set of IPSec policies that matches the needs of your environment. Organizations that have consistent security needs and a more simple network structure can create fewer and less complex policies to meet their goals. Other organizations with more stringent security needs and more complex environments require a greater number of policies and potentially more rules in their policies.

Use the steps shown in Figure 6.9 to help you decide how restrictively secure you want to make your IPSec policies.

**Figure 6.9 Designing IPSec Policies**




When designing an IPSec policy, you need to decide whether to apply an IPSec policy to protect traffic just for specific paths (between specific IP addresses) or to provide general protection for all traffic sent to and from a specific computer. If you want to use IPSec to provide general protection for all traffic sent to and from a specific computer, configure a policy that is based on blocking all traffic and then permitting exceptions, or permitting all traffic and then blocking particular ports. Blocking all traffic provides greater security, but this approach also requires a detailed analysis of the computer's communications, to ensure that required traffic is not blocked.

Consider a simple scenario that requires a simple IPSec policy strategy. All computers belong to a single Active Directory domain. The environment is composed entirely of servers running Windows Server 2003 and clients running Windows XP. All servers (except domain controllers) request confidentiality by using ESP, and all clients support IPSec. In such a case, only two policies are required: one for clients, and one for member servers. These policies can be applied by using Group Policy at the OU level in Active Directory. Because a different IPSec policy is applied to servers, the server computer accounts are grouped into their own OU. A GPO is created for the OU, and then the IPSec server policy is assigned to that GPO.

Because of the operating systems that the computers run and because all of the computers are domain members, both the IPSec server and client policies require that Kerberos be used for mutual authentication. To ensure that domain controllers do not receive any IPSec policy, the predefined domain security group for domain controllers is denied Read access to the GPO that assigns IPSec client policy at the domain level.

Conversely, many more policies can be required in a more complicated scenario. A single IPSec policy can contain many rules, each rule tailored for a specific type of traffic. Many different IPSec policies might need to be designed, to meet the security requirements different computers. Factors that can increase the number of policies required in your environment include:

- **Computer roles.** Different servers performing in different roles require different types of security, including some combination of packet filtering, confidentiality, integrity, and tunneling. Meeting the security needs of an environment that includes many servers with different roles and requirements can require hundreds of different configurations per server. Furthermore, servers might serve multiple roles, which complicates role-based configurations.

- **Sensitivity of data sent over the network.** Different levels of encryption are required to accommodate different security needs. For example, 3DES, which provides greater security than DES, can be used to protect transmission of personnel files stored on a file server.

- **Computer operating systems.** Some operating systems automatically support IPSec transport mode, but some require special client software to be installed. Depending on their operating system, computers might receive IPSec policy through different means: some by Group Policy, others only through locally installed policies. Additionally, Windows Server 2003 IPSec includes several new features that are not supported by Windows 2000 or Windows XP.

- **Domain memberships.** Some clients and servers belong to domains in different forests, some to UNIX realms, and some are members of workgroups. The trusts and authentication methods in place between such forests, realms, and workgroups affect which IPSec configurations are available.

- **Domain relationships.** As above, what IPSec configurations are available and what authentication is possible depends on whether domains are in the same forest or different forests, and on whether those forests trust each other.

If you have a complex environment, use IPSec only where it is truly needed. Although you want to provide an appropriate level of security, use as few policies as possible to minimize the complexity of your system. A simpler system is less likely to produce problems and is also easier to troubleshoot if it does.

After identifying what network segments, communications, and computers you want to secure by using IPSec, you need to specify the settings and rules that make up an IPSec policy. The settings and rules you specify also determine how strictly security is enforced by your IPSec policies. Not only must you put policies in place to meet your security requirements, but you must also ensure that computers on your network have compatible policies that can allow them to negotiate SAs.

#### General IPSec Policy Settings

General IPSec policy settings must be specified whether you want the policy to provide packet filtering or end-to-end security. Make sure to manage all IPSec policies in a controlled way from design, through testing, and into production. To ensure proper management, include version numbers in policy names. General IPSec policy settings are shown in Table 6.4.

**Table 6.4 General IPSec Policy Settings**

| Setting | Description | Example | For More Information |
|---|---|---|---|
| Name | The name and version number of the policy.<br><br>If the policy will be imported into a Windows 2000 Active Directory store or a local policy store, limit the name to 62 characters. | ContosoDefDomain v1.23 | "Add, edit, or remove IPSec policies" in Help and Support Center for Windows Server 2003 |
| Description | Optional text that describes the policy, includes the name of its administrative owner, and that can aid in management of this system. | *Administrative owner_name*. This policy blocks certain ports and requests integrity to specific server IP addresses. | "Add, edit, or remove IPSec policies" in Help and Support Center for Windows Server 2003 |
| Policy change poll interval | Specifies the period of time in minutes between polling by the IPSec Policy Agent for changes to existing applied policies.<br><br>IPSec polling does not detect changes made to domain or OU membership or the assigning or unassigning of policies in a GPO. These changes are detected by the Winlogon service every 90 minutes, by default. | Default of 180 minutes accepted | "Add, edit, or remove IPSec policies" in Help and Support Center for Windows Server 2003 |
| Key exchange settings | Determines how keys are created, and how frequently (in seconds) IKE negotiates an ISAKMP SA.<br><br>In the IP Security Monitor snap-in, relevant statistics are displayed under **Main Mode\IKE Policies**. | Default settings accepted | "Configure key exchange settings" in Help and Support Center for Windows Server 2003 |

| | | | |
|---|---|---|---|
| Key exchange methods | Determines how identities are protected when keys are exchanged.<br><br>You can specify which algorithms are used, including Message Digest 5 (MD5) and Secure Hash Algorithm (SHA1) for integrity, DES and 3DES for confidentiality, and the length of the master key. An ordered list of security settings is also specified, so that several settings can be offered during negotiation with the IPSec peer. | Default methods accepted | "Create key exchange security methods" in Help and Support Center for Windows Server 2003 |

**Note**

- Computers running Windows Server 2003 and Windows XP support the 3DES and DES algorithms and do not require installation of any additional components. However, computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use 3DES. If a computer running Windows 2000 is assigned a policy that uses 3DES encryption, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the security method defaults to the weaker DES algorithm. To ensure at least some level of privacy for communication, make sure to allow DES as a fallback option whenever a 3DES setting is applied to a group of computers in case some of them cannot support 3DES.

**IPSec Rules**

IPSec rules determine which traffic is affected by an IPSec policy and which actions take place when that type of traffic is encountered. Table 6.5 describes the contents of IPSec rules that two computers use to establish a secure, authenticated channel.

**Table 6.5 Settings of IPSec Rules**

| Setting | Description | Setting Example | For More Information |
|---|---|---|---|
| Filter list | Specifies a named list of filters. Each filter in the filter list specifies the types of traffic to which the filter action is applied. Filters can be defined to match specific IP protocols, source and destination TCP and UDP ports, and source and destination IP addresses.<br><br>The filter list name might include the version number, the last update time, and the administrative owner. Each computer discards the filter list name during policy processing.<br><br>The description or name for each filter is maintained on each computer during policy processing. Make sure to name each filter. | Source Address: My IP Address<br><br>Destination Address: 172.16.0.4<br><br>Protocol: TCP<br><br>Source Port: Any<br><br>Dest Port: 1434<br><br>Mirrored: Yes<br><br>Name: Me to sqlsvr3 TCP * 1434 | "Filter list" in Help and Support Center for Windows Server 2003 |
| Filter action | Specifies whether a packet is permitted, blocked, or secured. If packets are to be secured, specifies how they are secured. A list of security methods specifies the security protocol, cryptographic algorithm, and session key regeneration frequency. | Request Security | "Filter action" in Help and Support Center for Windows Server 2003 |
| Authentication methods | One or more authentication methods, which are specified in order of preference. Available options are Kerberos V5, certificate, or preshared key. | Kerberos V5 | "Authentication methods" in Help and Support Center for Windows Server 2003 |

| | | | |
|---|---|---|---|
| Tunnel endpoint | Specifies whether to use tunnel mode and, if so, the tunnel's endpoint. | 172.16.0.5 | "Tunnel endpoint" in Help and Support Center for Windows Server 2003 |
| Connection type | Specifies whether the rule applies to LAN connections, remote access connections, or both. | LAN | "Connection type" in Help and Support Center for Windows Server 2003 |

For more information about IPSec policy rules, see "IPSec Policy Rules" in Help and Support Center for Windows Server 2003.

# Understanding Default IPSec Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 includes three default IPSec policies that are provided as examples only. Do not use any part of the examples as templates to edit or change when creating your own IPSec policies. Instead, design new custom IPSec policies for operational use. The example policies will be overwritten during operating system upgrades and when IPSec policies are imported (when the import files contain other definitions of the same example policies). The three default IPSec policies are as follows:

- **Client (Respond Only).** This default policy contains one rule, the default response rule. The default response rule secures communication only upon request by another computer. This policy does not attempt to negotiate security for any other traffic.

- **Server (Request Security).** This default policy contains two rules: the default response rule and a second rule that allows initial incoming communication to be unsecured. The second rule then negotiates security for all outbound unicast IP traffic (security is not negotiated for multicast or broadcast traffic). The filter action for the second rule allows IKE to fall back to unsecured communication when required. This policy can be combined with the Client (Respond Only) policy when you want traffic secured by IPSec when possible, yet allow unsecured communication with computers that are not IPSec-enabled. If IKE receives a response from an IPSec-enabled client, but the IKE security negotiation fails, the communication is blocked. In this case, IKE cannot fall back to unsecured communication.

- **Secure Server (Require Security).** This default policy has two rules: the default response rule and a rule that allows the initial inbound communication request to be unsecured, but requires that all outbound communication be secured. The filter action for the second rule does not allow IKE to fall back to unsecured communication. If the IKE security negotiation fails, the outbound traffic is discarded and the communication is blocked. This policy requires that all connections be secured with IPSec. Any clients that are not IPSec-enabled cannot establish connections.

Both the Server (Request Security) and Secure Server (Require Security) default policies allow inbound unsecured communication, so that they can be used with the Client (Respond Only) policy. IPSec policies that are designed for operational use are typically customized to combine these behaviors as needed for a specific computer or network environment. The most secure IPSec configuration is not be represented by the default policies.

The most secure configuration is one that negotiates security and neither receives unsecured traffic nor allows communication with non-IPSec-aware computers. In such a configuration, the client computer policy must have a filter in a rule to initiate secured communication when applications on the client computer attempt outbound connections to the server.

**Predefined Filter Lists**

There are two predefined example filter lists included with Windows Server 2003:

- **All ICMP Traffic.** This filter affects any ICMP traffic (IP protocol 1) that is sent and received by using the unicast IP address of any network interface on a computer and all other computers. In Windows Server 2003, this filter matches outbound IP broadcast and multicast traffic when that traffic also uses the ICMP protocol.

- **All IP Traffic.** This filter affects all IP traffic that is sent and received using the unicast IP address of any network interface on the computer to any destination IP address. Because inbound broadcast and multicast traffic use a multicast or broadcast type for the destination address, the inbound traffic is exempt. All ISAKMP traffic sent over UDP port 500, which is required for establishing IPSec-secured communication, is also exempted from IPSec filtering. For more information, see "Understanding IP Filters, Filter Actions, and Filter Lists" later in this chapter.

**Predefined Filter Actions**

There are three predefined example filter actions included with Windows Server 2003:

- **Permit.** All traffic that matches the associated filters in the filter list is permitted. This permit behavior is static, not stateful.

- **Request Security.** All inbound traffic that matches the associated filters in the filter list is allowed unsecured. All

outbound traffic that matches the associated filters in the filter list negotiates security. Traffic is secured if the computer at the other end of the connection supports IPSec with a complementary filter action and filter in its policy. If security negotiation fails to elicit a response from the peer within three seconds, a security association for plaintext traffic is created (this association is known as a *soft SA*). If the peer responds within three seconds and the security negotiation fails, the communication is blocked. After IPSec SAs are established, all traffic that matches the associated filters is secured in both directions between the two computer IP addresses. This behavior is designed so that servers can request security for all clients, but fall back to unsecured communication with computers that are not IPSec-enabled.

- **Require Security.** As with the Request Security filter action, all inbound traffic that matches the associated filters in the filter list is allowed unsecured, and all outbound traffic that matches the associated filters in the filter list negotiates security. Unlike Request Security, if IKE fails to receive a response and the security negotiation fails, the communication is blocked. After IPSec SAs are established, all traffic that matches the associated filters is secured in both directions between the two computer IP addresses.

Each of these three filter actions allows inbound unsecured communication. This is not the most secure behavior, nor can it be used to initiate secured communication in some trust environments. These filter actions allow the Server (Request Security) and Secure Server (Require Security) example policies to be used with the Client (Respond Only) example policy. If Kerberos authentication is used, this behavior is appropriate for servers that only communicate with clients in mutually trusted domains. For servers that are in one-way domain trust environments, the client computer policy must have a filter in a rule to initiate secured communication with the server.

It is not appropriate to allow inbound unsecured communication for servers that have an Internet interface, because it can make the servers vulnerable to a denial-of-service attack. Internet attackers can easily spoof IP packets to flood the server, causing a denial of service as the server attempts many IKE negotiations with non-existent source IP addresses. The most secure filter action is one that requires security for all inbound and outbound traffic. To create this filter action, verify that the **Accept unsecured communication, but always respond using IPsec** and **Allow unsecured communication with non-IPsec aware computers** check boxes are cleared.

# Understanding IP Filters, Filter Actions, and Filter Lists

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

IP filters are used to specify the type of traffic that is permitted, blocked, or secured by an IPSec policy. For more information about specifying IP filter lists, see "Add, edit, or remove IP filter lists" in Help and Support Center for Windows Server 2003.

# Specifying Default Exemptions to IPSec Filtering

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default in Windows 2000 and Windows XP, broadcast, multicast, Kerberos, RSVP, and ISAKMP traffic is exempt from IPSec filtering. In Windows Server 2003, the default filtering exemptions have been removed for Kerberos, RSVP, and multicast and broadcast traffic, but remain for ISAKMP traffic, and inbound multicast and broadcast traffic.

To modify the default filtering behavior for Windows Server 2003 IPSec, you can use the Netsh IPSec context or modify the registry.

To modify the default filtering behavior by using the Netsh IPSec context, use the following command:

**netsh ipsec dynamic set config ipsecexempt value={0|1|2|3}**

Depending on which exemptions you want, specify the values as follows:

- A value of **0** specifies that multicast, broadcast, RSVP, Kerberos, and ISAKMP traffic are exempt from IPSec filtering. This is the default filtering behavior for Windows 2000 and Windows XP.

  Use this setting only if required for compatibility with Windows 2000 and Windows XP. However, if Kerberos traffic is exempted from filtering, an attacker can bypass other IPSec filters by using either UDP or TCP source port 88 to access any open port. Many port scan tools will not detect this because these tools do not allow setting the source port to 88 when checking for open ports.

- A value of **1** specifies that Kerberos and RSVP traffic are not exempt from IPSec filtering (multicast, broadcast, and ISAKMP traffic are exempt).

- A value of **2**specifies that multicast and broadcast traffic are not exempt from IPSec filtering (RSVP, Kerberos, and ISAKMP traffic are exempt).

- A value of **3** specifies that only ISAKMP traffic is exempt from IPSec filtering. This is the default filtering behavior for Windows Server 2003.

If you change the value for this setting, you must restart the computer for the new value to take effect.

To modify the default filtering behavior by using the registry, do the following:

1. Under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IPSEC**, add a new DWORD entry named **NoDefaultExempt**.

2. Assign this entry any value from **0**through **3**.

3. Restart the computer.

The filtering behaviors for each value are equivalent to those noted above for the **netsh ipsec dynamic set config ipsecexempt value=***x* command.

| ⚠**Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference at http://go.microsoft.com/fwlink/?LinkID=33702. |

The following table summarizes the equivalent filters that are implemented if all default exemptions to IPSec filtering are enabled (that is, if **NoDefaultExempt** is **0**). When the IP address is specified, the subnet mask is 255.255.255.255. When the IP address is **Any**, the subnet mask is 0.0.0.0.

**Table 6.6 Equivalent Filters for NoDefaultExempt=0**

| Source Address | Destination Address | Protocol | Source Port | Destination Port | Filter Action |
|---|---|---|---|---|---|
| My IP Address | Any IP Address | UDP | Any | 88 | Permit |
| Any IP Address | My IP Address | UDP | 88 | Any | Permit |
| Any IP Address | My IP Address | UDP | Any | 88 | Permit |
| My IP Address | Any IP Address | UDP | 88 | Any | Permit |
| My IP Address | Any IP Address | TCP | Any | 88 | Permit |
| Any IP Address | My IP Address | TCP | 88 | Any | Permit |
| Any IP Address | My IP Address | TCP | Any | 88 | Permit |
| My IP Address | Any IP Address | TCP | 88 | Any | Permit |
| My IP Address | Any IP Address | UDP | 500 | 500 (1) | Permit |
| Any IP Address | My IP Address | UDP | 500 | 500 | Permit |
| My IP Address | Peer IP Address | UDP | 4500 | 4500 (2) | Permit |
| Peer IP Address | My IP Address | UDP | 4500 | 4500 | Permit |
| My IP Address | Any | 46 (RSVP) | | | Permit |
| Any IP Address | My IP Address | 46 (RSVP) | | | Permit |
| Any IP Address | <multicast> (3) | | | | Permit |
| My IP Address | <multicast> | | | | Permit |
| Any IP Address | <broadcast> (4) | | | | Permit |
| My IP Address | <broadcast> | | | | Permit |
| <All IPv6 protocol traffic> (5) | | | | Permit | |

(1) In order for IPSec transport mode to be negotiated through an IPSec tunnel mode SA, ISAKMP traffic is not exempted if it needs to pass through the IPSec tunnel first.

(2) When IPSec NAT-T is performed, the filter exemption for UDP port 4500 is automatically generated, based on the source and destination IP addresses used during the initial part of the IKE negotiation on UDP port 500. This dynamic permit filter for port 4500 is displayed in the IP Security Monitor snap-in, under **Quick Mode\Specific Filters**, and in the output for the **netsh ipsec dynamic show qmfilter** command.

(3) Multicast traffic is defined as the class D range, with a destination address range of 224.0.0.0 with a 240.0.0.0 subnet mask, which corresponds to the range of addresses from 224.0.0.0 to 239.255.255.255.

(4) Broadcast traffic is defined as a destination address of 255.255.255.255 (the limited broadcast address), or as having the host ID portion of the IP address set to all 1's (the subnet broadcast address).

(5) IPSec does not support filtering for IP version 6 (IPv6) packets, except when IPv6 packets are encapsulated with an IPv4 header.

Windows Server 2003 IPSec does not support specific filters for broadcast protocols or ports, nor does it support multicast

groups, protocols, or ports. Because IPSec does not negotiate security for multicast and broadcast traffic, these types of traffic are dropped if they match a filter with a corresponding filter action to negotiate security. A filter with a source address of **Any IP Address** and a destination address of **Any IP Address** can block or permit all multicast and broadcast traffic. By default (and if the **NoDefaultExempt** registry key is set to a value of **2** or **3**), outbound multicast or broadcast traffic will be matched against a filter with a source address of **My IP Address** and a destination address of **Any IP Address**. More specific unicast IP address filters that block, permit, or negotiate security for unicast IP traffic should be configured in the same IPSec policy to achieve appropriate security.

For more information about viewing or modifying filter settings, see "Add, edit, or remove filter actions" and "Select a filter action for a rule" in Help and Support Center for Windows Server 2003.

# Linking Filter Actions and Filter Lists

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Filter lists and filter actions are linked together to form a rule in an IPSec policy. Although filter lists can be reused in different policies, they cannot be reused in the same policy. Filter actions can be reused by rules in the same policy, and they can be shared among different policies.

You can manage IPSec policy in one of two ways:

- Create a new policy and define the set of rules for the policy, adding filter lists and filter actions as required.

  In this method, you create an IPSec policy first and then add and configure rules. Filter lists (specifying traffic types) and filter actions (specifying how the traffic is treated) are added during rule creation.

- Create the set of filter lists and filter actions, and then create the policies and add rules that combine the filter lists with filter actions.

  In this method, you configure the filter lists and the filter actions. Next, you create IPSec policies and add rules that combine the appropriate filter list with the appropriate filter action. Additionally, you specify authentication methods, connection types, and tunnel settings.

# Configuring Firewalls

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The most secure firewall configuration is one in which the firewall permits only IKE and IPSec traffic to flow between the specific IP addresses of the peers. However, if these addresses are not static, or if there are many addresses, a less secure configuration might be required to permit IPSec and IKE traffic to flow between subnets.

When a firewall or filtering router exists between IPSec peers, it must be configured to forward IPSec traffic on UDP source and destination port 500, IP protocol 50 (ESP), or IP protocol 51 (AH). If you are using IPSec NAT-T, the firewall or filtering router must also be configured to forward IPSec traffic on UDP source and destination port 4500.

First, to permit IPSec traffic on UDP source and destination port 500, use the following settings to create a firewall filter called **Permit ISAKMP traffic on UDP port 500**:

- Source address = *Specific_IP_address*
- Destination address = *Specific_IP_address*
- Protocol = **UDP**
- Source port = **500**
- Destination port = **500**

If you are using IPSec NAT-T, to permit traffic on UDP source and destination port 4500, use the following settings to create a firewall filter called **Permit ISAKMP traffic on UDP port 4500**:

- Source address = *Specific_IP_address*
- Destination address = *Specific_IP_address*
- Protocol = **UDP**
- Source port = **4500**
- Destination port = **4500**

The firewall filter must also permit or track fragments for ISAKMP. IKE with certificate or Kerberos authentication requires ISAKMP packets to be fragmented because the IKE protocol uses UDP. ISAKMP messages that are larger than the local interface MTU are automatically fragmented by IP. If only certificate authentication is used, Windows Server 2003 implements a method to avoid IKE message fragmentation. When certificate authentication is used for communication between computers running Windows Server 2003 IPSec and Windows XP IPSec or Windows 2000 IPSec, fragmentation is required.

You must also allow IKE to be initiated from either a source or destination IP address. RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*, specifies that the IKE protocol must be able to negotiate security in either direction. Stateful filtering that allows only one computer to initiate IKE to a responder typically times out and deletes the stateful inbound filter in the firewall. As a result, IKE cannot rekey IPSec security associations, and IPSec connectivity is lost.

During the security negotiation, IKE detects whether IPSec NAT-T is required, and, if so, IKE automatically uses UDP port 4500 as the source and destination port. However, the exact firewall configuration that is required to permit this traffic depends on the location of the network address translator relative to the firewall. Typically, the network address translator modifies the source address and the source port of the IKE UDP header.

To permit IPSec traffic on IP protocol 50 (ESP) or IP protocol 51 (AH), use the following setting to create a firewall filter called **Permit IPSec traffic on ESP or AH protocol (50 or 51)**:

- Source address = *Specific_IP_address*

- Destination address = *Specific_IP_address*

- Protocol = **50** or **51**

When TCP traffic is protected by IPSec transport mode or tunnel mode, it should not require fragmentation. But UDP traffic that is protected by IPSec might require fragmentation. Therefore, you might need to enable fragment tracking for protocol 50 (ESP) or protocol 51 (AH).

- ESP and AH traffic should be allowed to flow in either direction. The upper layer protocols such as TCP will determine when IPSec packets are generated. However, ESP fully encapsulates the TCP header, so stateful filters might not be able to inspect the TCP header to track the TCP connection state and determine when to close the hole. IPSec SA state is controlled by IKE in the encrypted quick mode negotiation. Accordingly, IPSec SA creation and deletion can occur in either direction and cannot be detected by inspection of the IKE negotiation packets. Stateful filtering on the IPSec protocols might impact connectivity. To assess the possible impact, ensure that you test how effectively IPSec protects your application traffic with the appropriate firewall or filtering router.

After the firewall is opened to allow the IKE and IPSec protocols, the firewall might not be able to inspect packets to control which traffic is secured by IPSec. IPSec policy filters determine which traffic IPSec can secure, so if you want only a specific protocol to flow between two peers, you must create IPSec filters that enforce this behavior. Port-specific filters can control the direction in which connections are made. Therefore, if a computer is in a more trusted network (inside the firewall) and you want IPSec to secure traffic only over certain protocols and ports on that computer, do not enable the default response rule in the IPSec policy for that computer. If the default response rule is enabled and an attacker compromises the remote computer, the attacker might be able to modify the IPSec policy to negotiate security for all traffic through the firewall.

# Choosing the IPSec Protocol

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 IPSec protects data sent across the network by encapsulating the original payload with either an IPSec header (ESP or AH) for transport mode or an IPSec header and an additional IP header for tunnel mode. An IPSec protocol is often used to convert many different TCP and UDP ports used by applications and system services into one stream of secure traffic for the purpose of traversing an unsecured network. IPSec policy filters can control what protocols and ports are protected by IPSec.

Depending on which protocol is used, the entire original packet can be encrypted, encapsulated, or both. There are two IPSec protocols, Authentication Header (AH) and Encapsulating Security Payload (ESP).

**AH**

AH provides data origin authentication, data integrity, and anti-replay protection for the entire packet (both the IP header and the data payload carried in the packet), except for the fields in the IP header that are allowed to change in transit. It does not provide data confidentiality, which means that it does not encrypt the data. The data is readable, but protected from modification.

**ESP**

ESP provides data origin authentication, data integrity, anti-replay protection, and the option of confidentiality for the IP payload only. ESP in transport mode does not protect the entire packet with a cryptographic checksum. The IP header is not protected. ESP in tunnel mode encapsulates and secures the entire original packet, including the IP header (but not the outer IP header). Windows Server 2003 supports IPSec NAT-T, and ESP is the only protocol that can be used to traverse network address translators. The IKE protocol automatically detects the presence of a network address translator and proposes UDP-ESP encapsulation, rather than ESP, to allow IPSec traffic to pass through the network address translator. If only AH is specified as a security method in the filter action, the security negotiation will fail. ESP supports the option of null encryption when network address translator traversal, and not encryption, is needed.

Use Table 6.7 to choose which protocol to use.

**Table 6.7 Choosing IPSec Protocol Types**

| Protocol | Requirement | Usage |
|---|---|---|
| AH | The data and the header need to be protected from modification and authenticated, but remain readable. | Use for data integrity in situations where data is not secret but must be authenticated — for example, where access is enforced by IPSec to trusted computers only, or where network intrusion detection, QoS, or firewall filtering requires traffic inspection. |
| ESP | Only the data needs to be protected by encryption so it is unreadable, but the IP addressing can be left unprotected. | Use when data must be kept secret, such as file sharing, database traffic, RADIUS protocol data, or internal Web applications that have not been adequately secured by SSL. |
| Both AH and ESP | The header and data, respectively, need to be protected while data is encrypted. | Use for the highest security. However, there are very few circumstances in which the packet must be so strongly protected. When possible, use ESP alone instead. |

Using both AH and ESP is the only way to both protect the IP header and encrypt the data. However, this level of protection is rarely used because of the increased overhead that AH would incur for packets that are already adequately protected by ESP. ESP protects everything but the IP header, and modifying the IP header does not provide a valuable target for attackers. Generally, the only valuable information in the header is the addresses, and these cannot be spoofed effectively because ESP guarantees data origin authentication for the packets. In addition, some IPSec hardware offload adapters do not support the use of AH and ESP on the same packet. Hardware offload network adapters can accelerate IPSec processing by performing hardware offload of IPSec cryptographic functions. If you are using such offload adapters, determine the protocol support that they provide before selecting an IPSec protocol to use.

# Selecting IPSec Authentication Methods

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Peer authentication is the process of ensuring that an IPSec peer is who it claims to be. By using peer authentication, IPSec can determine whether or not to communicate with another computer before the communication begins. Windows Server 2003 IPSec performs peer authentication, but requires only mutual trust of the identities exchanged. It does not verify that the identity that is received is authorized to use a particular IP address.

Table 6.8 describes the uses of the authentication methods.

**Table 6.8 Choosing Authentication Methods**

| Security Requirement | Authentication Method | Examples |
|---|---|---|
| Communication within a Windows Server 2003 or Windows 2000 domain, or between trusted Windows Server 2003 or Windows 2000 domains. | Kerberos V5 | Clients accessing a Human Resources database<br>A Web server in a perimeter network connecting to a computer running SQL Server on an internal network |
| Communication outside of your domain or across the Internet where Kerberos V5 is not available but access to a CA is available. | Public key certificate | Partner organizations using a Web-based CA to access resources on your private network |
| Communication with systems that do not support the Kerberos V5 protocol and do not have access to a CA. | Preshared key | Windows 98 or Macintosh clients using third-party IPSec implementations<br>UNIX servers on a Windows network |

Windows Server 2003 IPSec supports three methods of peer authentication so that computers running different operating systems or that exist in different environments can find a common method when negotiating communication with a peer. An IPSec policy rule associates each IP address in a filter with an authentication method list, so that IKE can determine which authentication method list to use with each IP address. During IKE negotiation, the IKE initiator proposes a list of authentication methods to the IKE responder. The responder must use the source IP address of the initiator to identify which filter controls the IKE negotiation. The authentication method list that corresponds to the filter in the responder's IPSec policy is used to select one authentication method from the initiator's list. The responder then replies to inform the initiator of the agreed upon authentication method. If the selected authentication method fails, IKE does not provide a method for retrying a different authentication method.

# Kerberos Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Kerberos V5 is the default authentication standard in Windows Server 2003 and Windows 2000 domains. This method of authentication can be used by any computer in the domain or a trusted domain. Also, because the Kerberos protocol is no longer a default exemption, if you want to enable Kerberos authentication, you must create filters in the IPSec policy that explicitly allow all traffic to your domain controllers.

For IKE Kerberos authentication to be used successfully in Windows 2000 and Windows Server 2003 cross-forest trust configurations, you must use FQDNs when configuring external trusts. In addition, you must configure the IPSec policy on both computers to allow an IKE initiator to communicate to any domain controller in the forest domain hierarchy, so that the initiator can obtain a ticket from a domain controller in the responder's domain.

It is recommended that you use Kerberos authentication in a two-way (mutual) domain and forest trust environment. Although one-way domain and forest trusts are supported, it is not recommended to use Kerberos in this type of trust environment. If you do, you must design the IPSec policy to ensure that the IKE initiator can obtain a ticket from a domain controller in the responder's domain. Also, traffic might be lost if IKE fails to rekey main mode negotiations in the opposite direction. Because an IKE main mode negotiation is performed on demand when quick mode negotiations are rekeyed, when configuring lifetimes in kilobytes (KBs) for the associated filter action, use as high a value as possible to minimize rekey authentication failures. If you are using ESP DES or 3DES for encryption, you can set the lifetime value to as high as 200,000 KB (200 megabytes). If you set a higher value, however, the risk of a sophisticated attacker gaining knowledge of the 56-bit encryption keys increases. If you are using ESP or AH for integrity and authenticity, you can set the value to as high as 2 gigabytes because the key sizes are much larger for MD5 (128-bit) and SHA1 (160-bit).

By default, when Kerberos authentication is used, the **Access this computer from the network** or **Deny this computer access from the network**logon right (defined in Group Policy) is evaluated. This evaluation is only performed by the IKE main mode responder (the computer that receives the ticket and must determine whether accept it). Typically, clients obtain Kerberos tickets to access a server. Likewise, the IPSec policy must contain filters that will trigger the IKE negotiation from the IKE initiator, the client that is requesting IPSec access to the server. The server can then use these logon rights to restrict access to certain client computer security groups. However, if IKE rekeys main mode negotiations, the server does not evaluate these rights.

# Public Key Certificate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Windows 2000 and Windows Server 2003, you can use Certificate Services to automatically manage computer certificates for IPSec throughout the certificate lifecycle. Certificate Services is integrated with Active Directory and Group Policy, and it simplifies certificate deployment by enabling certificate auto-enrollment and renewal and by providing configurable certificate templates. In addition, by publishing the computer certificate as an attribute of the domain computer account, Certificate Services allows you to use IPSec to restrict network access to a server.

Use a public key certificate in situations that include access to corporate resources, external business partner communications, or computers that do not run the Kerberos V5 authentication protocol. This requires that at least one trusted root CA is configured on your network and that client computers have an associated computer certificate.

IPSec supports the use of a variety of third-party X.509 public key infrastructure (PKI) systems, in addition to Windows 2000 Server or Windows Server 2003 Certificate Services. Windows Server 2003 IKE has basic compatibility with several certificate systems, including those offered by Microsoft, Entrust, VeriSign, and Netscape. If you are using a third-party PKI system, the PKI system must be able to issue certificates to computers and store their certificates in the Windows Cryptographic Application Programming Interface (CryptoAPI) computer certificate store.

**Note**

- Certificates obtained from Certificate Services with the advanced option set for **Enable strong private key protection** do not work for IKE authentication because a personal identification number (PIN) cannot be entered to access the private key during IKE negotiation.

**General certificate authentication considerations**

IKE's use of certificate authentication is compatible with many different PKI architectures, and IKE places relatively few requirements on the contents of a certificate. Typically, computers that have a common trusted root, or whose certificates can chain through a cross-certification trust relationship, can successfully use IKE authentication. To use certificates for IKE authentication, you define an ordered list of acceptable root CA names in the authentication method. This list controls the certificates that IKE can select and the certificates that IKE will select. If IKE authentication fails, you cannot retry the authentication using a different method. For this reason, before you apply an IPSec policy that can use certificates for authentication, make sure that all target computers have the correct root CA certificates and relevant cross-certificates in the CryptoAPI PKI stores, as well as valid computer certificates. Additionally, to ensure that certificate authentication works as intended, test your PKI infrastructure with various IPSec policy configurations before deployment.

The following sections describe the IKE certificate selection and acceptance process. If you decide to use certificates for IKE authentication, understanding this process and its requirements is integral to ensuring proper deployment.

**IKE certificate selection process**

When IKE negotiates to use certificates for authentication, the following process is used to select a computer certificate:

1. The list of trusted roots is prepared. This list is the list of the CA root names provided by the peer in the Certificate Request Payloads (CRPs), and it matches the CA root names configured in the list of trusted roots in the appropriate authentication method of the IPSec policy. If there are no matching CA root names, all trusted CA root names from the appropriate authentication method are used.

2. IKE searches the computer store for an IPSec certificate that chains to any of the trusted CA roots identified in step 1. An IPSec certificate contains an Enhanced Key Usage (EKU) attribute with a value equal to the IP security IKE intermediate object identifier (OID) 1.3.6.1.5.5.8.2.2.

3. For each certificate chain found, the following checks are performed to verify that:

    - The certificate chain does not have any trust errors.

    - The certificate chain is not a root-only chain.

- The computer certificate has a private key.
- The computer certificate has an RSA type public/private key pair.
- The computer certificate has a public key length that is greater than 512 bits.
- The computer certificate has a Digital Signature key usage.
- The computer certificate is not a CA signing certificate that is used to issue certificates.
- The certificate chain passes certificate revocation list (CRL) checking, which is performed by default or if the value of the **StrongCrlCheck** registry key is set to **1** or **2**. For more information about CRL checking, see "CRL Checking" later in this chapter.

If all of these checks succeed, IKE selects the certificate chain to be sent to the IPSec peer. If any of these checks fails, IKE continues to search for another IPSec type certificate, using the same list of root CA names.

4. If a valid computer certificate chain is not located, IKE retries the process, from step 2. Although the same list of root CA names is used, IKE does not search for an IPSec type certificate.

5. If a valid computer certificate chain is still not found and if the list of root CA names in step 1 is a subset of the names allowed by the local IPSec policy, IKE retries, from step 2. This time, IKE uses the entire list of root CA names allowed by the local authentication method.

   This step is required for successful authentication when cross-certificates are used to establish trust relationships.

6. After IKE selects a computer certificate, it includes all intermediate certificates in the chain up to the root, except for the root CA certificate. A certificate chain in PKCS#7 format is then sent to the IPSec peer. If there are no intermediate CAs, only the computer certificate is sent.

   If a computer certificate cannot be selected, the authentication fails.

### Note

- If IKE negotiates with another computer running Windows Server 2003, or with other Microsoft IKE implementations that use IPSec NAT-T (such as Microsoft L2TP/IPsec VPN Client), a special method to avoid fragmentation of ISAKMP UDP packets might be implemented. Otherwise, the ISAKMP message that contains the certificate chain will likely be fragmented as the packet is transmitted.

#### IKE certificate acceptance process

1. IKE receives the peer's certificates or certificate chains and verifies that the peer's certificates chain up to any of the root CAs in the appropriate authentication method of the local IPSec policy.

2. For each peer certificate chain, the following checks are performed to verify that:

   - The computer certificate Subject Name or Subject AltName is consistent with the peer's ID field passed in the IKE negotiation.
   - The computer certificate chain does not have any trust errors. If there is a trust error, the peer authentication fails.

3. If the two checks in step 2 succeed, the following checks are performed for the peer certificate chain to verify that:

   - The certificate chain passes CRL checking, which is performed by default or if the value of the **StrongCrlCheck**

registry key is set to **1** or **2**.

- The computer certificate has an RSA type public/private key pair.
- The computer certificate has a public key length that is greater than 512 bits.
- The computer certificate has a Digital Signature key usage.

If any of these checks fails, the peer authentication fails.

4. If certificate-to-account mapping is enabled in the IPSec policy for the certificate root CA of the peer, IKE calls the Windows secure channel (Schannel) APIs to perform the mapping. Schannel completes the mapping and builds an access token for the computer account. This access token is automatically evaluated against the **Access this computer from the network** or the **Deny this computer access from the network** logon right defined in Group Policy Security settings.

   If the logon right evaluation fails, the peer authentication fails.

**Editing Default Certificate Templates to Include the Subject Name**

The IP Security Monitor snap-in, Event Viewer, and most IPSec diagnostics report the Subject Name attribute of the IPSec peer certificate. In Windows Server 2003 Certificate Services, the default version 2 certificate template for IPSec certificates does not include the computer name in the **Subject Name** field (the version 1 certificate templates cannot be modified). To aid in deployment and troubleshooting, edit the default IPSec certificate template to include either the FQDN or the common name of the computer in the **Subject Name** field. To do this, use the following procedure (note that you must be logged on as a member of the Enterprise Admins group or the root domain's Domain Admins group in Active Directory):

**To edit the default IPSec certificate template**

1. Open the Certificate Templates snap-in (Certtmp.msc).
2. Right-click the certificate template that you want to modify, and then click **Properties**.
3. On the **Subject Name** tab, click **Build from this Active Directory Information**, and then click **Common Name** or **Fully Distinguished Name**.
4. Enter the common name or FQDN as required.

If you choose to enter the common name of the computer, the name in the **Subject Name** field of the certificate will appear as CN=< FQDN>. If any FQDN is greater than 64 characters, select **Fully Distinguished Name**, not **Common Name**. After you complete this procedure, all subsequently issued computer certificates will contain the FQDN in the **Subject Name** field of the certificate. If you have already completed enrollment for computer certificates by using the default certificate template, create a new certificate template and configure it to supersede the existing template. Doing this will automatically update all computer certificates so that the blank name in the **Subject Name** field is replaced with the FQDN of the computer. If the existing certificate template is not superseded, each computer will have two computer certificates that can be used for IKE authentication. IKE can choose either one.

**Auditing IKE Negotiation Successes and Failures**

You can view the success or failure of IKE negotiations in the Event Viewer security log. To view these events, enable success or failure auditing for the **Audit logon events** audit policy for your domain or local computer. If auditing is enabled for IKE events, and IKE authentication fails, event 547 is recorded. If IKE authentication succeeds, event 541 is recorded.

When you enable success or failure auditing for the **Audit logon events** audit policy, IPSec records the success or failure of each main mode and quick mode negotiation and the establishment and termination of each negotiation as separate events. In many cases, it might be necessary to enable success auditing for the **Audit logon events** audit policy, to track user activity. Keep in mind, however, that enabling this type of auditing can cause the security log to fill with IKE events. You can disable auditing of IKE events by modifying the registry.

To disable auditing of IKE events in the security log, do the following:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Audit\DisableIKEAudits** registry setting to a value of **1**.

   The **DisableIKEAudits key** does not exist by default and must be created.

2. Do one of the following:

   - Restart the computer
   - Stop, and then restart the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

   **Note**

Stopping and restarting the IPSec service can disconnect all of the computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. If you restart the IPSec service immediately, TCP-based communication might resume, due to the retransmit behavior of TCP, after new IKE and IPSec SAs are established.

**CRL Checking**

By default, in Windows XP and Windows Server 2003, IPSec CRLs are automatically checked during IKE certificate authentication, but a fully successful CRL check is not required for the certificate to be accepted. However, if enhanced security is required, a fully successful CRL check is also required. CRL checking can cause delays in authentication or unnecessary failures, and some third-party PKI systems might not support it. You can disable IPSec CRL checking or specify a stronger level of IPSec CRL checking by using the Netsh IPSec context or by modifying the registry.

To disable IPSec CRL checking or specify a different level of IPSec CRL checking by using the Netsh IPSec context, use the following command:

```
netsh ipsec dynamic set config strongcrlcheck value={0 | 1 | 2}
```

**To disable IPSec CRL checking or specify a different level of IPSec CRL checking**

1. Under **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent**, add a new **Oakley** key, with a DWORD entry named **StrongCRLCheck.**

2. Assign this entry any value from **0** through **2**, where:

   - A value of **0** disables CRL checking.
   - A value of **1** enables the default level of CRL checking for Windows Server 2003 IKE. When this level of CRL checking is performed, certificate validation fails only if the certificate is revoked.
   - A value of **2** enables strong CRL checking, which means that certification validation fails if any error is encountered during CRL processing.

3. Do one of the following:

   - Restart the computer.
   - Stop, and then restart the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

Note that IPSec CRL checking does not guarantee that certificate validation fails immediately when a certificate is revoked. There is a delay between the time that the revoked certificate is placed on an updated and published CRL and the time when the computer that performs the IPSec CRL checking retrieves this CRL. The computer does not retrieve a new CRL until the current CRL has expired or until the next time the CRL is scheduled to be published. By default, IKE requests that Crypto API (CAPI) waits 15 seconds to complete the CRL retrieval. If the CRL cannot be retrieved at that time, IKE either ignores the error (if the value of the **StrongCRLCheck** registry key is set to **1**, or it causes authentication to fail (if the value of **StrongCRLCheck** is set to **2**). CRLs are cached in memory and in \Documents and Settings*UserName*\Local Settings\Temporary Internet Files by CAPI. Because CRLs persist across computer restarts, if a CRL cache problem occurs, restarting the computer does not resolve the problem. For more information about CRLs, see the Troubleshooting Certificate Status and Revocation link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Certificate-to-Account Mapping**

In Windows Server 2003, a specific group of computers can be authorized to use IPSec when either Kerberos V5 or certificates are used for IKE authentication. This capability enables much stronger peer authentication and IPSec to be used to restrict network access to a server. When you enable IPSec certificate-to-account mapping, the IKE protocol associates (that is, maps) a computer certificate to a computer account in an Active Directory domain or forest, and then retrieves an access token, which includes the list of computer security groups. This process ensures that the certificate offered by the IPSec peer corresponds to an active computer account in the domain, and that the certificate is one that should be used by that computer.

Certificate-to-account mapping can only be used for computer accounts that are in the same forest as the computer performing the mapping. This provides much stronger authentication than simply accepting any valid certificate chain. For example, you can use this capability to restrict access to computers that are only within the same forest. Certificate-to-account mapping, however, does not ensure that a specific trusted computer is being allowed IPSec access.

Certificate-to-account mapping is especially useful if the certificates come from a PKI that is not integrated with your Active Directory deployment, such as if business partners obtain their certificates from third-party providers. You can configure the IPSec policy authentication method to map certificates to a domain computer account for a specific root CA. You can also map all certificates from an issuing CA to one computer account. This allows IKE certificate authentication to be used to limit which forests are allowed IPsec access in an environment where many forests exist and each perform autoenrollment under a single internal root CA. If the certificate-to-account mapping process is not completed properly, authentication will fail and IPSec-protected connections will be blocked. For more information about establishing certificate-to-account mapping, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**Excluding the CA Name from Certificate Requests**

If you use certificate authentication to establish trust between IPSec peers, you can also use Windows Server 2003 to exclude CA names from certificate requests. Excluding the CA name prevents a malicious user from learning sensitive information about the trust relationships of a computer, such as the name of the company that owns the computer and the domain membership of the computer (if an internal PKI is being used). Although excluding the CA name from certificate requests enhances security, computers with multiple certificates from different roots might require the CA root names to select the correct certificate. Also, some third-party IKE implementations might not respond to a certificate request that does not include a CA name. For these reasons, excluding the CA name from certificate requests might cause IKE certificate authentication to fail in certain cases.

# Preshared Key

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are not using Kerberos V5 authentication and do not have access to a CA, a preshared key can be used. For example, a stand-alone computer on a network that does not connect to the Internet might need to use a preshared key, because neither Kerberos authentication through the computer's domain account nor access to a CA on the Internet are available.

A *preshared key* is a shared secret key that has been agreed upon by administrators who wish to secure the computer's communications by using IPSec. Administrators must manually configure their systems to use the same preshared key.

**Important**

- Microsoft does not recommend the use of preshared key authentication, because the authentication key is stored in plaintext format in the system registry and hex-encoded in Active Directory–based IPSec policy. Well-known methods can enable attackers with access to these data stores to discover weak preshared key values.

Use preshared key authentication only where no stronger method can be used. Using Kerberos or certificate-based authentication is recommended to avoid security risks associated with preshared key authentication.

If you must use preshared key authentication, use only local or persistent IPSec policy, a 25-character or longer random key value, and a different preshared key for each IP address pair. These practices result in different security rules for each destination, and ensure that a compromised preshared key compromises only those computers that share the key. For more information about local or persistent IPSec policies, see "Assigning IPSec Policies Locally" later in this chapter.

# IKE Authentication Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use Windows Server 2003 IPSec policies and IKE authentication to limit network access to trusted computers. In most scenarios, successful IKE authentication results in successful network access to a computer. IKE is not aware of the identity or the public key that is expected from the peer. Therefore, if the certificate private key or domain password for a computer were compromised, an attacker might be able to use that computer to successfully authenticate and gain access to another IPSec-protected computer or to conduct trusted man-in-the-middle attacks on IPSec communication. To ensure that IPSec provides the appropriate network access controls, carefully consider the following:

- How computers are joined to an Active Directory security domain.
- How trusts between domains and forests are controlled.
- How computers obtain a certificate from a trusted root CA (or an issuing CA).
- How PKI trust is controlled (for example, how cross certificates are handled).
- How to reduce the number of trusted computers.

# Assigning IPSec Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As a domain administrator, you can configure IPSec policies to meet the security requirements of a user, group, application, domain, site, or global enterprise from a domain controller. IPSec policy can also be implemented in a non-Windows 2000-based domain environment by using local IPSec policies.

Follow the steps shown in Figure 6.10 to assign IPSec policies during your deployment.

**Figure 6.10 Assigning IPSec Policies**




Make sure that you test these IPSec policies before you actually assign them in your production environment. For more information about testing IPSec policies, see "Testing Your Policies in a Test Lab" later in this chapter.

First determine whether to use Active Directory to apply IPSec policy to clients. If you decide to do so, make sure that you understand how Group Policy inheritance affects the way in which the IPSec policies are applied and how IPSec policy precedence differs from standard Group Policy inheritance.

# Understanding the New Default IP Security Policies Container Permissions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain-based IPSec objects are stored in the IP Security Policies container in Active Directory, which is separate from the Group Policy objects (GPOs) to which the IPSec policies are applied. By default in Windows Server 2003, Active Directory restricts Read permissions on the IP Security Policies container more than in Windows 2000. In Windows Server 2003, only members of the Group Policy Creator Owners and the Domain Computers groups have Read permissions on the IP Security Policies container. The domain administrator must grant permissions to the IP Security Policies container for other delegated administrators to administer IPSec policies.

If you perform a new installation of Windows Server 2003 and create a domain controller, the domain controller will have the following new default permissions on the IP Security Policies container:

- Owner: Domain Admins
- Group: Domain Admins
- Allow domain computers: Read only
- Allow Group Policy Creator Owners: Read only
- Allow Domain Admins: Full Control

When you run the Adprep tool (available in the I386 directory in the Windows Server 2003 product CD) to prepare Windows 2000 domains and forests for an upgrade to Windows Server 2003, it will change the permissions on the IP Security Policies container to these new settings, unless you have already changed the default permissions in Windows. When you create new objects in the IP Security Policies container, those objects inherit the permissions of that container. Upgrading from Windows 2000 will not modify the permissions on existing IPSec policy objects.

Due to the more restrictive permissions in Windows Server 2003, a member of the Domain Admins group must explicitly allow the assignment of an IPSec policy to child domains within a forest by enabling child domain computers to read the IP Security Policies container in the directory. For these computers to retrieve domain-based IPSec policy from Active Directory, the Local System account for each computer must have Read permissions on the IP Security Policies container. If computer accounts in child domains must read a parent domain-based IPSec policy, you must modify the permissions on the IP Security Policies container to allow this.

Make sure to carefully control Modify access to the IP Security Policies container. Do not assign specific permissions for individual IPSec policies. The IPSec policies are a collection of related directory objects, some of which can be shared between policies. Accordingly, you should control permissions on the IP Security Policies container itself.

An IPSec policy administrator typically has Write access to all IPSec policies. You can restrict who can create and modify GPOs so that only authorized individuals can assign a domain-based IPSec policy. Make sure to investigate these permissions and set them as needed for your environment.

Standard delegation tools cannot be used to delegate permissions to administer IPSec policies. Instead, domain administrators must use the Active Directory Service Interfaces (ADSI) Edit tool for this purpose. ADSIEdit is a Microsoft Management Console (MMC) snap-in that domain administrators can use to edit objects in the Active Directory database. When domain administrators delegate permissions to others to administer IPSec policies, the delegated administrators must have Full Control permissions to all IPSec policy objects in the IP Security Policies container. After an IPSec administrator creates an IPSec policy, a member of the Group Policy Owner Creators group or other delegated owner of the GPO can assign the IPSec policy to the appropriate GPOs.

**Caution**

To modify access control lists (ACLs) on the IP Security Policies container, domain administrators should use the tools that are provided with the Windows 2000 or Windows Server 2003 operating system CD. When domain administrators modify ACLs on the IP Security Policies container, they must specify that all modifications are also propagated to all child objects. Incorrect modification of the ACLs on the IP Security Policies container, failure to propagate modifications to the child objects of this container, or both, can result in the failure of IPSec policy settings to be properly applied.

ADSIEdit is one of the Windows Support tools included on the Windows 2000 and Windows Server 2003 operating system CDs. You can install these tools by running the Support Tools Setup program, Suptools.msi, from the \Support\Tools folder.

For more information about permissions for the IPSec Policies container, see article 329194, "IPSec Policy Permissions in Windows 2000 and Windows Server 2003" in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkID=106552.

# Understanding IPSec Policy Precedence

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To know how to apply IPSec policy in a domain environment, you must understand IPSec policy precedence. Unlike most Group Policy settings, which are cumulative, only one IPSec policy can be assigned to a computer at a time. Therefore, if there are multiple IPSec policies assigned at different levels, the last one applied is the one that takes effect. IPSec policy uses the same precedence sequence as other Group Policy settings, which is from lowest to highest, as follows:

- **Local GPO.** Each computer has one local GPO. For a computer that is not a member of a domain, this is the only place where IPSec policy can be assigned. Although you can assign an IPSec policy by editing the local GPO, you can also assign a local policy directly in the IP Security Policy Management snap-in, outside of the Group Policy context, or by using the Netsh IPSec context. When IPSec policy assignments are made outside of the GPO context, the GPO cannot display the local IPSec policy that is assigned.

- **Site.** IPSec policies are not often assigned at the site level because all computers within a site must have the same security needs, which is unlikely. Furthermore, if the computer moves to another subnet — such as when a user travels to another location with a laptop that uses DHCP — different policies are applied, which results in different security behaviors.

- **Domain.** Simple IPSec policies are often assigned at the domain level and then superseded by more specific IPSec policies, as required on various OUs.

- **OU.** Specific IPSec policies are assigned to groups of computers. This is the last policy applied under normal conditions, and, therefore, the policy takes precedence. If an OU is nested within another OU, the IPSec policy assigned to the nested OU takes precedence.

Note that this is the default order in which policies of different types are applied. This order can be overridden by using a number of options, including **Enforced**, **Block Policy Inheritance**, and **Loopback** processing. For more information about the results of these options, see "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit.

**Note**

- If you create a persistent policy, this policy adds to or overrides the local or Active Directory policy and remains in effect regardless of whether other policies are applied. For more information about persistent IPSec policy, see "Assigning IPSec Policies Locally" later in this chapter.

# Selecting a Method for Assigning IPSec Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The next step is to decide how to distribute and assign IPSec policies. This might include storing and assigning IPSec policies from Active Directory as part of Group Policy, storing and assigning them from the local computer, or assigning them from a remote computer.

The primary method of building the IPSec policy is using the graphical user interface provided by the IP Security Policy Management snap-in. You can use the IP Security Policy Management snap-in to create, modify, and activate IPSec policies, and then assign them to an OU in Active Directory that contains the server using the Group Policy Object Editor snap-in. Use this for situations where larger numbers of computers need to be managed in a consistent fashion. A Group Policy–based management solution addresses this situation well.

If a computer is not a member of a Windows 2000 domain or a Windows Server 2003 domain, it cannot retrieve IPSec policy from Active Directory. The IPSec policy configuration for a server can be distributed two ways that do not use Active Directory:

- As a Netsh IPSec script that is included as a startup script for the computer.
- As an IPSec file that can be imported from another computer by using the IP Security Policy Management snap-in or Netsh if the computer is running Windows Server 2003. Use this method when computers need to secure their communications, but there are few enough of them that applying policies to OUs is inconvenient. After the policy is imported, you can use the IP Security Policy Management snap-in to assign the policy to the computer.

If you use either of these two methods, make sure you use strict version and change control processes to ensure that the policy file cannot be altered after it was created. If the policy is accessed by using Lightweight Directory Access Protocol (LDAP), the IPSec policy configuration data can be authenticated and encrypted.

Using explicit credentials for remote management and monitoring of IPSec is not supported. Instead, the IP Security Policy Management and IP Security Monitor snap-ins use the credentials that are provided during the desktop logon process to authenticate to a remote computer.

# Assigning Domain-based, OU-Level, and Local IPSec Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

IPSec policies can be applied to local computers, domains, sites, organizational units, or any Group Policy object in Active Directory. The following sections describe deployment considerations for domain-based, OU-level, and local IPSec policies.

# Assigning Domain-based IPSec Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you determine where you want to protect data and create IPSec policies for each of those areas, you must determine which IPSec policy to assign to GPOs.

A GPO defines access, configuration, and usage settings for accounts and resources. When an IPSec policy is assigned to a GPO, the IPSec policy is propagated to any computer accounts that are affected by the GPO.

Keep these factors in mind when selecting GPOs for IPSec policy assignment:

- IPSec policies can be assigned to the GPO of a site, domain, or OU. However, only a single local or domain-based IPSec policy can be assigned to a specific computer.
- Like Group Policy settings, the assignment precedence for IPSec policies, from lowest to highest, is local, site, domain, and OU. In addition, persistent IPSec policy has the highest precedence of all, even though it is stored on the local computer.
- Unlike Group Policy settings, IPSec policies from different OUs are never merged.
- The highest level of the Active Directory hierarchy is typically used to broadly assign IPSec policies, to reduce the amount of configuration and administration required.
- For domain-based IPSec policy, limit the number of rules to 10 or less, even though Windows Server 2003 supports over 1500 rules per policy.
- Create and apply an IPSec policy at the domain level to provide a baseline of IPSec protection. For example, the default domain policy GPO can often be used to assign IPSec policy for all clients in a domain.
- A new IPSec policy can be tested as a local policy on a server or a client first before it is uploaded into the domain. Be sure to adequately test the impact of new IPSec policies before assigning them in the domain.
- When configuring the **Negotiate security** filter action, select the **Allow unsecured communication with non-IPSec aware computers** check box during the rollout phase, so that communication will not be blocked after a computer starts using IPSec policy. After verifying that all computers have IPSec policy and are working correctly, clear this check box to require only IPSec-secured communications.
- Use the **Export Policies** and **Import Policies** commands in the IP Security Policy Management console to back up and restore all of the IPSec policy objects in the IP Security Policies container in Active Directory. The **Export Policies** command exports all IPSec policy objects from the policy store into one .ipsec file. The **Import Policies** command imports all IPSec policy objects in the .ipsec file into the destination policy store. Note that importing IPSec policy into Active Directory will overwrite existing IPSec policy objects.

**Caution**

- When importing or editing IPsec policy in Active Directory, do not close the IP Security Policy Management snap-in (or exit the Netsh environment, if you are using the **netsh ipsec static importpolicy** command) before all IPSec policy data has been written to Active Directory. If the IP Security Policy Management snap-in or Netsh Ipsec cannot finish writing all of the policy data into Active Directory, IPSec policy corruption might result.
- If you detect IPSec policy corruption, you can try to reimport the IPSec policy file. In some cases, you must use either LDP.exe or ADSIedit.exe to delete the IPSec policy objects. (LDP.exe is a Support Tool that was included in the Support Tools folder of the Windows 2000 operating system CD, and is included in the Support Tools folder of the Windows XP and Windows Server 2003 operating system CDs). The IPSec policy objects must be deleted so that a new IPSec policy import operation can be successfully completed. If you are managing IPSec policy remotely over slow links, transfer the

IPSec policies in .ipsec export files by using a file copy technique first. Then use Remote Desktop Connection to connect to the remote server and perform the import operation quickly.

# Assigning OU-Level IPSec Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After assigning a domain-based IPSec policy as a baseline for IPSec protection, you can assign other IPSec policies to specific OUs to provide tighter IPSec filtering in sensitive areas of your network. By default, the IPSec policies assigned to a GPO at the OU level override the baseline policy for the domain.

For example, you can organize the computer accounts in the domain to keep domain clients in one OU and IPSec-protected servers in another OU, so that client policy will not affect servers.

**Note**

- Assigning an IPSec policy to a GPO (such as an OU) records a pointer to the IPSec policy inside the GPO. Thus Group Policy only detects changes in IPSec policy assignments, not changes made within an IPSec policy after it is assigned to a GPO. The IPSec service detects changes in the related IPSec policy. You can specify the frequency of this detection process, known as the *IPSec polling interval,* by setting a value on the **General** tab of the properties of a policy by using the IP Security Policy Management snap-in.

It is recommended that you test your new IPSec policies as a local policy on a server or a client first before assigning them into the domain.

# Assigning IPSec Policies Locally

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Each computer running Windows Server 2003 has one local GPO, which is also known as the *local computer policy*. When this local GPO is used, Group Policy settings can be stored on individual computers regardless of whether they are members of an Active Directory domain. The local GPO can be overridden by GPOs assigned to sites, domains, or OUs in an Active Directory environment that have higher precedence. On a network without an Active Directory domain (that is, a domain that does not have a domain controller running Windows 2000 or Windows Server 2003), the local GPO settings determine IPSec behavior because they are not overridden by other GPOs.

Local policy assignment is a way to enable IPSec for computers that are not members of a domain.

You can also create and assign *persistent IPSec policy*, which secures a computer even if a local IPSec policy or an Active Directory–based IPSec policy cannot be applied. This policy adds to or overrides the local or Active Directory policy, and remains in effect regardless of whether other policies are applied or not. Persistent IPSec policies enhance security by providing a secure transition from computer startup to IPsec policy enforcement. Persistent policy also provides backup security in the event of an IPSec policy corruption, or if errors occur during the application of local or domain-based IPSec policy. To configure persistent policies, you must use the **`netsh ipsec static set store location=persistent`** command.

When designing persistent IPSec policy, it is important to consider the potential impact of persistent policy on remote management. If local or domain-based IPSec policy is not applied and the persistent IPSec policy is the only policy that is applied, attempts to remotely diagnose an issue might be blocked by the persistent IPSec policy. To allow for remote management in case troubleshooting is required, it is recommended that you create appropriate permit filters when configuring persistent IPSec policy.

To enable remote management, add the following exemption by using the Netsh IPSec context command:

```
        netsh ipsec dynamic set config bootexemptions tcp:0:3389:inbound
UDP:0:68:inbound
```

This command specifies two things: that the destination port 3389 for inbound TCP connections is permitted, thereby enabling clients to use Remote Desktop Connection or Remote Assistance; and that destination port 68 for inbound UDP connections is permitted, thereby preserving DHCP functionality.

If an error occurs when persistent IPSec policy is applied, the IPSec driver will block all traffic except for that which matches any specific permit filters that you configure (by using the **netsh ipsec dynamic set config bootexemptions** command), and DHCP.

For more information about assigning local IPSec policies, see "Creating, modifying, and assigning IPSec policies" in Help and Support Center for Windows Server 2003.

# Understanding IPSec Protection During Computer Startup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 IPSec provides several options for protection during computer startup. To provide maximum protection against attacks during computer startup, it is highly recommended that you configure and assign a persistent IPSec policy. If you do not configure a persistent policy, the IPSec driver cannot enforce IPSec policy until domain-based or local IPSec policy is retrieved and applied.

# Understanding IPSec Driver Startup Modes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To understand how IPSec policies are processed and applied, it is important to understand the modes in which the IPSec driver operates. The IPSec driver can perform in any of three following computer startup modes:

- **Permit.** In this mode, the IPSec driver permits all inbound and outbound traffic. After persistent, local, or domain-based IPSec policy is applied, the IPSec driver no longer operates in this mode.

- **Block.** In this mode, the IPSec driver blocks all inbound and outbound traffic until persistent policy is applied, except for traffic that matches any specific permit filters that you configure (by using the **netsh ipsec dynamic set config bootexemptions** command), and DHCP traffic.

- **Stateful.** In this mode, the IPSec driver permits all outbound traffic initiated by the computer during startup. Inbound traffic that is sent in response to the outbound traffic is permitted, specific to the peer IP address, protocol, and source and destination ports. The IPSec driver also permits inbound traffic that matches any specific filters that you configure (by using the **netsh ipsec dynamic set config bootexemptions** command), and DHCP traffic. All other inbound unicast, broadcast, and multicast packets are dropped. The stateful inbound permit filters are discarded after the IPSec service starts and sets persistent IPSec policy.

If a persistent, local, or domain-based IPSec policy has been assigned to a computer, the IPSec Policy Agent sets the stateful mode for the IPSec driver by default. You can change the default startup mode for the IPSec driver by using the Netsh IPSec context or by modifying the registry.

To change the default startup mode of the IPSec driver by using the Netsh IPSec context, use the following command:

```
netsh ipsec dynamic set config bootmode value={stateful | block | permit}
```

**To change the default startup mode of the IPSec driver by modifying the registry**

1. Under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IPSEC**, add a new DWORD entry named **OperationMode**

2. Assign this new registry entry any value from **0, 1, or 3** where:

   - A value of **0** specifies that the IPSec driver operate in Permit mode, to allow all inbound and outbound traffic.

   - A value of **1** specifies that the IPSec driver operate in Block mode, except to allow traffic that matches any specific permit filters that you configure by using Netsh, and DHCP traffic.

   - A value of **2** is reserved.

   - A value of **3** specifies that the IPSec driver operate in Stateful mode and permits inbound traffic that matches any specific filters that you configure (by using the **netsh ipsec dynamic set config bootexemptions** command), and DHCP traffic.

3. Restart the computer.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit*

companion CD or at http://www.microsoft.com/reskit.

# Using IPSec in Safe Mode with Networking

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a computer is started in Safe Mode with Networking (for example, if you are performing Directory Service Restore Mode operations), the IPSec service does not start. As a result, no persistent, local, or domain-based IPSec policy can be applied, and IKE cannot negotiate security for network traffic. However, if an IPSec policy has been assigned to a computer, IPSec still can provide some level of filtering protection. If the IPSec service startup type is set to Automatic, the IPSec driver remains in computer startup mode (which is Stateful mode, by default). When the IPSec driver operates in this mode, you can configure specific permit filters to allow inbound traffic over specific protocols and ports.

If you must change the IPSec driver mode to Permit to allow all inbound connections, you must either set the **OperationMode** registry key to a value of **0** or change the IPSec service startup type to Manual or Disabled, and then restart the computer in Safe Mode with Networking again. Note that in this case, you cannot use the **netsh ipsec dynamic set config bootmode value=** command to change the IPSec driver startup type,because Netsh IPSec requires the IPSec service to be running.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

# Using Netsh Scripts to Assign IPSec Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Windows Server 2003, the Netsh IPSec tool provides a scriptable command-line method of building an IPSec policy. This is useful in cases such as when a computer is not a domain member or is running an older version of Windows. Netsh IPSec can create a persistent policy or a local policy, both of which are stored in the local computer registry, or it can create a policy that is stored in Active Directory.

The IPSec context of the Netsh tool can also dynamically insert new IPSec rules into the run-time system. This "dynamic mode" IPSec policy is part of the run-time state and is not stored; therefore, it is lost when the IPSec service is stopped either administratively or during a restart.

The IPSec internal infrastructure components were significantly modified for Windows Server 2003 such that the Netdiag.exe, IPSecpol.exe, and IPSeccmd.exe tools from earlier Windows releases cannot run properly. You do not need to import policies created by these tools during an upgrade – policies in place before an upgrade to Windows Server 2003 continue to function after the upgrade has been completed. In all cases, the user or process that sets the IPSec policy must be running as a local or domain administrator.

For more information about importing IPSec policies, see "Creating, modifying, and assigning IPSec policies" in Help and Support Center for Windows Server 2003.

Windows 2000, Windows XP, and Windows Server 2003 do not have published programmatic APIs in the Microsoft Windows Platform Software Development Kit for IPSec policy. Command-line scripting using the Netsh IPSec tool is the only method of managing policy in an automated fashion.

# Coordinating IPSec Client and Server Policies in a Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Both clients and servers need IPSec policy before secure communications can be enabled. Because the initial application of IPSec policy and policy changes can briefly interrupt connectivity, make sure to schedule these tasks during low use periods as much as possible. When policies are changed, the IPSec system might be forced to delete existing IPSec SAs, so that new SAs can be negotiated according to the new IPSec policy. This can cause a temporary loss of communication.

Remember that domain IPSec policy will override local IPSec policy. If a local administrator has secured a server with local IPSec policy, applying a domain policy might make the server less secure or break its secure communication. If persistent policy was used, the domain policy cannot override it, but can enhance that baseline security.

# Deploying Your IPSec Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After scoping your needs, building IPSec policies, and determining your strategy for assigning the policies to specific OUs, test the IPSec policies in a lab environment and conduct a pilot project before rolling them out for production use.

To ensure that IPSec policy functions as expected and provides the appropriate level of security, test specific IPSec policy configurations on clients and servers in a lab environment, and then conduct pilot or beta tests in a limited operational environment before conducting a full-scale deployment. Test both functional behavior and the policy assignment steps required to ensure that one end of the communication does not receive a policy that blocks communication until the other end receives a compatible policy.

Follow the steps shown in Figure 6.11 to deploy IPSec.

**Figure 6.11 Deploying Your IPSec Solution**



# Testing Your Policies in a Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you roll out IPSec into your production environment, test your IPSec policies by following these steps:

1. Configure a lab in a controlled environment separate from your production network. Make the lab simulate your production environment as closely as possible. For example, you might want to include a perimeter network, a data center server, firewalls or filtering routers, clients with various operating systems, application test configurations, multiple domains with various OUs, and a network with varying connection speeds. Also, make sure to activate those Active Directory services that are running in your production environment.

2. Use the RSoP component of Group Policy to verify that the IPSec policies are being applied the way in which you assigned them, and that the policy you expect takes precedence if multiple policies are applied.

3. Implement your policies on the server and clients to test communication of all software and services. Use the IP Security Monitor snap-in or the **netsh ipsec dynamic show mmfilter** and **netsh ipsec dynamic show qmfilter** commands to verify that specific filters are being applied as expected. In the IP Security Monitor snap-in, the **Specific Filters** nodes display quick mode and main mode filters configured for persistent, local, or domain-based IPSec policies. However, default exemptions to IPSec filtering are not displayed.

4. Look for communications that have failed. Check if any network traffic was blocked by any of the IPSec polices that you have implemented.

5. Simulate a load on the network that is similar to what your network would experience, to test for performance impact.

6. Complete penetration tests to verify that the filters are configured properly and behave as you expect. Verify that the exemptions to IPSec filtering are appropriate.

Use the IP Security Monitor snap-in to gather information you can use to identify problems and optimize performance where IPSec is deployed. For example, you can view details about IPSec policies and filters, statistics about performance, and SAs.

For more information about the IP Security Monitor snap-in, see "Viewing details about active IPSec policies in IP Security Monitor" in Help and Support Center for Windows Server 2003. For more information about setting up a test environment, see "Designing a Test Environment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

# Using RSoPd to Analyze IPsec Policy Assignments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use RSoP to analyze IPSec policy assignments. RSoP is a Group Policy–related snap-in that you can use to view IPSec policy assignments for a computer or for members of a Group Policy container.

To view IPSec policy assignments in RSoP, you must first open the RSoP snap-in, and then run a query. RSoP allows you to troubleshoot policy precedence issues and determine the exact set of policies that are being used by IPSec clients.

For more information about RSoP, see "Using Resultant Set of Policy to view IPSec policy assignments" in Help and Support Center for Windows Server 2003. For more information about IPSec policy precedence, see "Understanding IPSec Policy Precedence" earlier in this chapter.

# Testing Your Policies in a Pilot Project

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Verifying successful IPSec operation is an important, and often neglected, step in deploying IPSec. Incorrectly configuring IPSec and related Windows networking components can cause network traffic to be unintentionally blocked or sent unprotected. The specifics of how to verify successful operation depend on the particulars of your deployment. You might need to engage the Help desk and conduct training in the technology, as well as create FAQs or self-help documents and procedures, so that you can quickly isolate problems.

Testing successful IPSec operations is uniquely different from testing other networking components. You can often test for successful operation of a system by performing tests to see if the applications behave as expected. However, applications in the overall system can perform properly while not protected by IPSec. Testing IPSec involves testing that the applications work, and that IPSec is in fact restricting access and performing protection.

It is also important to test IPSec policy management procedures and key IT operations processes, such as:

- Importing and exporting IPSec policy
- Remotely managing IPSec policy
- Initially assigning IPSec policy to different types of computers
- Changing IPSec policy on each computer
- Backup and restore operations
- Server and network monitoring
- Deploying new domain controllers
- Unassigning IPSec policy
- Adding subnets to the network
- Server scaling
- Cross-platform IPSec policy compatibility

Finally, coordinate with your network operations team to determine how procedures for responding to internal network attacks might impact the use of IPSec in your network.

For more information about planning and rolling out a pilot project, see "Designing a Pilot Project" in *Planning, Testing, and Piloting Deployment Projects* in this kit.

# Deploying IPSec Across Your Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The general steps for assigning IPSec policy by using Group Policy are:

1. Create the IPSec policy in the directory. If you have a computer dedicated to using local IPSec policy, you can use the **netsh ipsec static importpolicy** command or the IP Security Policy Management snap-in.

2. Create a new GPO or identify an existing GPO to use to deliver the IPSec policy assignment. Remember that only members of the Domain Admins group can modify the GPO and assign IPSec policy.

3. Assign the GPO to the domain level or an appropriate OU level, and then resolve any IPSec policy precedence issues. Group Policy objects, which include IPSec policies, can be configured to not be updated over slow links. IPSec policy should *always* be implemented, regardless of whether the updates must use slow links. If you prevent IPSec policy updates over slow links, communication failures might occur.

4. Plan the rollout of the IPSec policy on both the client and server side. Determine whether to use Group Policy, the IP Security Policy Management snap-in, executable files, or some combination to deploy this technology. If you are going to secure a server, make sure to enable client-side IPSec policy first, allow time for all clients to retrieve the policy, then enable the server-side policy. When configuring the **Negotiate security** filter action, select the **Allow unsecured communication with non-IPSec aware computers** check box, so that IKE can fall back to unsecured communication if a remote computer has not yet retrieved IPSec policy.

5. Ensure that IPSec authentication methods are ready before rolling out IPSec policy. All certificate autoenrollment problems must be investigated and resolved before assigning IPSec policy. Any cross-certificate trusts must be established and deployed to all systems first. Likewise, if Kerberos is being used, ensure that all domain trusts between clients and servers are mutually trusted and that Kerberos traffic is permitted. After IPSec policy is applied, if IKE authentication fails, communication will be blocked. For more information about certificate autoenrollment, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

# Additional Resources for Deploying IPSec

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this chapter.

**Related Information**

- "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit for more information about creating hardware and software inventories and mapping your existing network before deploying new features.
- "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit.
- "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit for more information about certificate auto-enrollment.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about virtual private networking.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about Internet Protocol security (IPSec).
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about Certificate Services.
- For more information about IPSec, see the Windows IPSec Web site link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- For more information about VPN, see the Virtual Private Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Job Aids**

- "Designing an IPSec Policy" (DNSIPS_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Designing an IPSec Policy" on the Web at http://www.microsoft.com/reskit).

**Related Tools**

- Netsh

  Netsh is a command-line scripting tool included with Windows Server 2003 that allows you to, either locally or remotely, display or modify the network configuration of a computer that is currently running. For more information about Netsh, see "Netsh" in Help and Support Center for Windows Server 2003.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center for Windows Server 2003, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Establishing an IPSec security plan"
- "Special IPSec considerations"
- "Creating, modifying, and assigning IPSec policies"
- "Add, edit, or remove IPSec policies"

- "Start the IP Security Policy Management snap-in"
- "Assign or unassign IPSec policy in Group Policy"
- "Assign or unassign IPSec policy on a computer"
- "Configure key exchange settings" and "Create key exchange security methods"
- "New features for IPSec"
- "IPSec Policy Rules"
- "Filter list"
- "Filter action"
- "Add, edit, or remove filter actions"
- "Authentication methods"
- "Add, edit, or remove IP filter lists"
- "Default IPSec policies"
- "Viewing details about active IPSec policies in IP Security Monitor"
- "Netsh commands for Internet Protocol Security (IPSec)"
- "Using Resultant Set of Policy to view IPSec policy assignments"

# Deploying Internet Authentication Service (IAS)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Servers running the Internet Authentication Service (IAS) component of the Microsoft® Windows® Server 2003 operating system perform centralized authentication, authorization, auditing, and accounting for many types of network access, including dial-up, virtual private network (VPN), wireless, and 802.1X authenticating switch access. IAS is the Microsoft implementation of the Remote Authentication Dial-In User Service (RADIUS) protocol. A number of design, implementation, and deployment issues must be considered when rolling out a scalable and robust IAS solution.

Information about deploying remote access clients can be found in other chapters in this book.

**In This Chapter**

Overview of IAS Deployment

Designing IAS

Designing an Optimized IAS Solution

Creating a Remote Access Policy Strategy

Securing Your Remote Access Strategy

Implementing Your IAS Solution

Additional Resources for Remote Access

**Related Information**

- For information about Windows Server 2003 Internet Authentication Service (IAS), see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

- For information about Windows Server 2003 Routing and Remote Access, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

- For information about Windows Server 2003 Routing and Remote Access deployment, see "Deploying Dial-up and VPN Remote Access Servers," "Deploying Remote Access Clients Using Connection Manager," and "Connecting Remote Sites" in this book.

# Overview of IAS Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use IAS to provide authentication and authorization for dial-up, VPN, wireless, and authenticating switch access to your network. For example, organizations that outsource network access and perform joint ventures with other organizations require the authentication of user accounts from outside of the private network. In addition, organizations that provide outsourcing services, such as Internet service providers (ISPs), require remote user connection accounting so that they can charge subscribers.

Windows Server 2003 IAS enables you to centralize authorization, authentication, and accounting for remote access clients, enhancing the security of your network. Windows Server 2003 IAS works with other standards-based implementations of the Remote Authentication Dial-In User Service (RADIUS) protocol, so that you can use it with any standards-compliant RADIUS client, server, or proxy server.

Windows Server 2003 IAS is included in Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition. IAS is not provided with Microsoft® Windows® Server 2003, Web Edition. In addition, Windows Server 2003, Standard Edition, limits some IAS features. For more information, see "Concepts for IAS" later in this chapter.

Windows Server 2003 IAS provides the following solutions for organizations that require secure network access:

- Compatibility with RADIUS servers and clients from any vendor that meets the specifications outlined in RFCs 2865, 2866, and 2869.
- Integration with the Active Directory® directory service. IAS allows you to take advantage of Active Directory for user authentication, authorization, and client configuration, thus reducing management costs.
- Use of standards-based strong authentication methods for network access.
- Management of network access that is outsourced to an ISP. IAS allows you to create an agreement between your organization and the ISP in which the ISP can charge a roaming user's department for that employee's network usage. In this way, each employee does not need to submit an expense statement or create a roaming account to connect to the corporate network remotely.
- Dynamic key management for wireless access points as a means to increase network security.

# IAS Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process of deploying IAS involves deciding what role IAS will play in your network infrastructure, configuring connection request and remote access policies to allow users to access the network, and securing and optimizing the design. Figure 7.1 shows the IAS deployment process.

**Figure 7.1 Deploying IAS**



# Concepts for IAS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

IAS implements the Internet Engineering Task Force (IETF) standard Remote Authentication Dial-In User Service (RADIUS) protocol, specified in RFCs 2865, 2866, and 2869, which enables the use of a homogeneous or heterogeneous network of dial-up, VPN, wireless, or authenticating switch equipment. When a remote client tries to connect to an access server configured to use the RADIUS protocol, the access server sends the connection request to the IAS server by using the RADIUS protocol. When an IAS server is a member of an Active Directory® domain, IAS uses the directory service as its user account database and is part of a single sign-on solution. The same set of credentials is used for network access control (authenticating and authorizing access to a network) and to log on to an Active Directory domain.

In addition, the IAS server can accept or reject the request based on conditions that you specify in the remote access policies. *Remote access policies* are an ordered set of rules that define how connections are either authorized or rejected. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting. Access to the network resources can be controlled by applying policies to users or groups of users. For more information about using remote access policies to grant access, see "Remote Access Policies" in Help and Support Center for Windows Server 2003.

When using IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range. However, when using Windows Server 2003, Standard Edition, you can configure IAS with a maximum of 50 RADIUS clients and a maximum of two remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query.

RADIUS is a client/server protocol that requires a RADIUS client and a RADIUS server to provide network access. An access server or a RADIUS proxy is a *RADIUS client*, and the computer making the determination of authentication and authorization is a *RADIUS server*.

Figure 7.2 shows a typical IAS architecture. An access client contacts an IAS RADIUS proxy located at an ISP by using a local telephone connection. The IAS proxy examines the user name, which contains two elements – the identification of the user account name and the identification of the user account location (also known as a *realm*). Based on the realm portion of the user name in the connection request, the IAS RADIUS proxy forwards the connection request to a RADIUS server on a private network, which authenticates and authorizes the connection attempt with the Active Directory user accounts database and user account properties.

**Figure 7.2 IAS Architecture**




For more information about Internet Authentication Service (IAS), see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) and "Internet Authentication Service" in Help and Support Center for Windows Server 2003.

For more information about realm names, see "Realm names" in Help and Support Center for Windows Server 2003.

A general understanding of the following topics is also essential to proper IAS deployment:

- Remote access.
- The Windows Server 2003 implementation of remote access.
- Any network access mechanism your users require, such as dial-up, VPN, wireless, or authenticating switch access.
- Network security.

- Active Directory.
- Authentication.
- Accounting.

For more information about any of these topics, see the Windows Server 2003 Technical Reference and Help and Support Center for Windows Server 2003.

**New in Windows Server 2003**
Windows Server 2003 IAS includes the following new features:

- **Support for RADIUS proxy**. With RADIUS proxy support, you can use IAS as a RADIUS message router or forwarder between access servers and other IAS servers. Based on attributes in the incoming RADIUS message, the RADIUS proxy forwards the message to a specific RADIUS server or client.
- **Support for mapping network authentication and authorization for IAS proxy**. The proxy component of IAS supports the ability to separate the authentication and authorization of connection requests from access servers. The IAS proxy can forward password-based user credentials to an external RADIUS server for authentication, and perform authorization against a user account in an Active Directory domain and a locally configured remote access policy. Alternate user authentication databases can be used but connection authorization and restrictions are still determined through local administration. You can configure the proxy component with the **Remote-RADIUS-to-Windows-User-Mapping** attribute in the advanced properties of a connection request policy. For more information, see "Mapping network authentication and authorization" and "Connection request policies" in Help and Support Center for Windows Server 2003.
- **Support for IEEE 802.1X wireless and authenticating switches**. IAS provides authentication, authorization, and accounting for connections that use the link-layer standard IEEE 802.1X for wireless and authenticating switch access. For more information about configuring IAS for wireless access authentication, see "Checklists: Configuring IAS for Wireless Access" and "Wireless access" in Help and Support Center for Windows Server 2003.
- **Support for Protected Extensible Authentication Protocol (PEAP) for 802.11 wireless clients**. Protected Extensible Authentication Protocol (PEAP) provides protection for clients and authenticators (IAS or RADIUS servers) that are using Extensible Authentication Protocol (EAP). The next generation of EAP, PEAP uses Transport Layer Security (TLS) to create end-to-end communication between client and authenticator after the identity of the authenticator is verified. For more information, see "PEAP" in Help and Support Center for Windows Server 2003.
- **Enhanced EAP configuration for remote access policies**. In Microsoft® Windows® 2000 Server, you can select only a single EAP type for a remote access policy. In Windows Server 2003 IAS, this limitation is removed. For example, you can select different computer certificates for VPN connections and EAP-TLS authentication for wireless connections, or you can select multiple EAP types for wireless connections in the circumstance where some of your wireless clients use EAP-TLS authentication and some of them use PEAP with Microsoft Challenge Handshake Authentication Protocol version 2 (EAP-MS-CHAPv2). For more information, see "EAP" and "Configure PEAP and EAP methods" in Help and Support Center for Windows Server 2003.
- **Support for IAS Network Access Quarantine Control**. IAS Network Access Quarantine Control provides phased network access, which restricts the access of remote clients to quarantine mode until each client is either verified as meeting or configured according to organization network access policy. After the client computer configuration is verified as meeting organization network policy, the quarantine restrictions, which consist of Quarantine IP-Filters and Session Timers, are removed and standard remote access policy is applied to the connection. For more information, see "IAS Network Access Quarantine Control" in Help and Support Center for Windows Server 2003.
- **Support for logging to MSDE 2000 and SQL Server 2000 databases**. You can use an XML-compliant database, such as Microsoft® SQL Server™ 2000 and SQL Server Desktop Engine (MSDE 2000), to log user authentication requests, periodic data, and accounting requests received from one or more access servers. For more information, see "SQL Server database logging" in Help and Support Center for Windows Server 2003.
- **Support for ignoring the dial-in properties of user accounts**. You can configure a RADIUS attribute on the profile

properties of a remote access policy to ignore the dial-in properties of user accounts. To support multiple types of connections for which IAS provides authentication and authorization, it might be necessary to disable the processing of user account dial-in properties. This can be done to support scenarios in which specific dial-in properties are not required. For more information, see "New features for IAS" in Help and Support Center for Windows Server 2003.

- **Support for configuring RADIUS clients by IP address range**. For IAS in Windows 2000, you must specify a RADIUS client by IP address or by Domain Name System (DNS) name. In addition, you must configure each RADIUS client separately, even if you have a number of RADIUS clients on the same subnet. While this is not an issue for typical dial-in or VPN access server configurations, numerous wireless access points can be placed on the same subnet, creating a circumstance where use of an IP address range simplifies configuration and administration. In Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, IAS allows you to specify a RADIUS client by using an IP address range. All of the RADIUS clients in the range must use the same configuration and shared secret. For more information, see "Configure RADIUS Clients" in Help and Support Center for Windows Server 2003.

- **Support for computer authentication**. In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, Active Directory and IAS support the authentication of computer accounts by using standard user authentication methods such as Point-to-Point Protocol (PPP). This allows a computer and its computer certificate to be authenticated for wireless or authenticating switch access clients.

- **Support for checking user certificate purposes**. To enforce the use of specific types of user certificates for specific connection types, you can configure IAS to check the purposes (also known as *object identifiers* or*OIDs*) of certificates in their Enhanced Key Usage (EKU) extensions. You can configure a list of object identifiers that are required to be present in the user certificate. For more information, see "Network access authentication and certificates" and "Add RADIUS attributes to a remote access policy" in Help and Support Center for Windows Server 2003.

- **Improved attribute manipulation**. In Windows 2000, you can use IAS to manipulate the contents of the User-Name RADIUS attribute. Using connection request policies in IAS for Windows Server 2003, you can manipulate the User-Name, Called-Station-ID, and Calling-Station-ID RADIUS attributes. For more information, see "Connection request policies" in Help and Support Center for Windows Server 2003.

- **Support for the Authentication Type remote access policy condition**. You can create remote access policies by using the Authentication Type condition in IAS for Windows Server 2003. You can use the Authentication Type condition to specify connection constraints that are based on the authentication protocol or method that is used by the access client. For more information, see "Elements of a remote access policy" in Help and Support Center for Windows Server 2003.

- **Improved support for the Class attribute**. In Windows 2000, IAS automatically generates a value for the Class attribute and appends it to the existing value of the Class attribute received in the RADIUS request message. The result is the Class attribute in the RADIUS response message. In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, you can disable the automatic generation of a value for the Class attribute by using the **Generate-Class-Attribute** setting on the **Advanced** tab in the properties of a remote access policy profile. Automatic generation of a value for the Class attribute is disabled by default. Instead of appending the generated value of the Class attribute to the existing Class attribute, IAS creates a separate Class attribute. The RADIUS response message contains both the original Class attribute and the second Class attribute that is generated by IAS. For more information, see "Add RADIUS attributes to a remote access policy" in Help and Support Center for Windows Server 2003.

#### IAS Terms and Definitions

This section provides brief definitions of important IAS terms. For more information, see the sources mentioned in "Additional Resources for Remote Access" later in this chapter, or Help and Support Center for Windows Server 2003.

**Internet Authentication Service (IAS)** A component of Windows Server 2003 that performs centralized authentication, authorization, auditing, and accounting for a variety of different types of network access, including dial-up, VPN, wireless, and authenticating switch access.

**Remote Authentication Dial-In User Service (RADIUS) protocol** A protocol, specified in RFCs 2138 and 2139, that enables use of a homogeneous or heterogeneous network of dial-up or VPN equipment.

**RADIUS server** A server that implements the RADIUS protocol. A RADIUS server performs authentication, authorization, and

accounting on behalf of a RADIUS client.

**RADIUS client** A network access server that uses a RADIUS server to perform authentication, authorization, and accounting for its access clients.

**RADIUS proxy** A server that forwards incoming RADIUS messages to specific RADIUS servers for additional processing, based on RADIUS attributes in the incoming RADIUS message. An IAS server can act as a RADIUS server for some requests and a RADIUS proxy for others.

**Authentication** The process of performing identity verification of the entities that communicate over a network; for example, the process that verifies the identity of a user who logs on to a computer either locally, at a computer's keyboard, or remotely, by means of a network connection.

**Authorization** The process that determines what a user is permitted to do on a computer system or network. For remote access or demand-dial routing connections, the verification that the connection attempt is allowed. Authorization occurs after successful authentication. With IAS, you can manage authorization using remote access policies and Active Directory user account properties.

**Auditing** The process that tracks the activities of users by recording selected types of events in the security log of a computer.

**Accounting** A method of tracking account activity information such as logon and logoff records, to help maintain records for billing purposes. With IAS, you can also track authentication information, such as each accept, reject, and automatic account lockout.

# Designing IAS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After taking an inventory of your network environment, the first step in designing an IAS solution is to determine the role of the IAS server. For example, determine whether you need the IAS server to authenticate the connection request that it receives, forward the request to another IAS server for authentication, or perform a mixture of both functions depending on context. Finally, an important step in the design process is to configure IAS to work with different types of clients.

Figure 7.3 shows the process for designing IAS.

**Figure 7.3 Designing IAS**



# Inventory Your Current Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When beginning your design process, make sure to have specifications about your current network environment available for use. Specifically, your hardware and software inventory and a map of network topology can be very helpful in reducing time spent during the design phase. For more information about creating those documents, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

# Determine the Role of the IAS Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure your IAS server to act as a RADIUS server, a RADIUS proxy, or both, depending on where you want network access requests to be authenticated.

**RADIUS server**
If you want your IAS server to authenticate the connection requests that it receives, rather than forwarding connection requests to another IAS server, use the IAS server as a RADIUS server. For example, if your access servers connect directly to your network, then the IAS server is configured as a RADIUS server to authenticate the connection.

Figure 7.4 shows an IAS server configured as a RADIUS server. An access client connects to an access server. The access server sends a connection request to an IAS RADIUS server located on the corporate network, which authenticates and authorizes the connection attempt.

**Figure 7.4 IAS Configured as a RADIUS Server**




**RADIUS proxy**
If you want an IAS server to forward connection requests to another IAS server, use IAS as a RADIUS proxy. Use the RADIUS proxy capabilities in the following situations.

**Using IAS proxy at a third-party ISP**
You are an ISP providing outsourced network connection services to multiple customers. Your network access servers send connection requests to the IAS RADIUS proxy. Based on the realm portion of the user name in the connection request, the IAS RADIUS proxy forwards the connection request to a RADIUS server maintained by the customer that can authenticate and authorize the connection attempt.

Figure 7.5 shows an IAS server configured as a RADIUS proxy. An access client contacts an access server at an ISP. The ISP access server sends a connection request to an IAS RADIUS proxy. Based on the realm portion of the user name in the connection request, the IAS RADIUS proxy forwards the connection request to a RADIUS server located on the corporate network, which authenticates and authorizes the connection attempt.

**Figure 7.5 IAS Configured as a RADIUS Proxy at a Third-Party ISP**




**Using IAS proxy with multiple forests**
You have multiple forests and want to perform cross-forest authentication with Extensible Authentication Protocol-Transport Layer Security (EAP-TLS). Rather than configuring your access servers to send their connection requests to an IAS RADIUS server, configure them to send their connection requests to an IAS RADIUS proxy.

Figure 7.6 shows an IAS server configured as a RADIUS proxy forwarding RADIUS messages to RADIUS servers in multiple forests. The IAS RADIUS proxy uses the domain name portion of the user name and forwards the request to an IAS server in each forest.

**Figure 7.6 IAS as a RADIUS Proxy with Multiple Forests**




**Using IAS proxy for load balancing**

You want to increase the capacity for connection requests. In this case, rather than configure your access servers to attempt to load balance across multiple RADIUS servers, configure them to send their connection requests to an IAS RADIUS proxy. The IAS RADIUS proxy can load balance across multiple RADIUS servers and scale up to large numbers of RADIUS clients and authentications per second.

Figure 7.7 shows an access server forwarding a request to a RADIUS proxy to load balance to multiple RADIUS servers. The remote client connects to a RADIUS client, such as an access server. The access server sends the authentication request to the RADIUS proxy, which load balances the request across different IAS servers.

**Figure 7.7 Load Balancing**




**IAS server and proxy**

If you need your IAS server to authenticate some requests and forward other requests, use IAS as both a RADIUS server and a RADIUS proxy. For example, if you are performing cross-forest authentication, use your IAS server as a RADIUS server to authenticate users in the same forest, and use it as a RADIUS proxy to forward authentication requests to another IAS server for users in another forest.

Figure 7.8 shows an IAS server configured to be both a RADIUS server and a RADIUS proxy. The remote client connects to an IAS server configured as both a RADIUS server and a RADIUS proxy. Based on the realm portion of the access client user name, the IAS server determines whether to authenticate the request directly or forward the authentication request on to another IAS server in a different forest.

**Figure 7.8 IAS Configured as Both a RADIUS Server and a RADIUS Proxy**



# Design IAS as a RADIUS Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing IAS as a RADIUS server involves some or all of the following tasks:

- Determining the IAS server domain membership
- Planning for RADIUS clients
- Planning authentication
- Determining compatibility with third-party access servers
- Adding RADIUS or vendor-specific attributes (VSAs) to the remote access policy
- Installing a backup IAS server

For more information about configuring IAS as a RADIUS server, see "Deploy IAS as a RADIUS Server" later in this chapter and "IAS as a RADIUS server" in Help and Support Center for Windows Server 2003.

**Determine IAS Server Domain Membership**

You must determine which domain the IAS server computer will belong to. In multiple-domain environments, an IAS server can authenticate user credentials for user accounts in the domain in which it is a member and for user accounts in all domains that have a two-way trust relationship with the domain in which it is a member. For increased performance, install IAS on the domain controllers in the domain that will authenticate the users. If you install IAS on a domain controller, you do not have to add the domain controller computer account to the RAS and IAS Servers group of the domain.

For more information about configuring IAS servers, see "Implementing Your IAS Solution" later in this chapter.

**Plan for RADIUS Clients**

A RADIUS client can be an access server (for example, a dial-up or VPN server, a wireless access point, or an Ethernet authenticating switch) or a RADIUS proxy. IAS supports all access servers and RADIUS proxies that comply with RFC 2865, *Remote Authentication Dial-in User Service (RADIUS)*.

Configure each access server or RADIUS proxy that sends RADIUS request messages to the IAS server as a RADIUS client. For each RADIUS client, you can configure a friendly name, an IP address or DNS name, the client vendor, the shared secret, and whether the RADIUS **Message-Authenticator** attribute is used.

You can specify IP addresses or DNS names for RADIUS clients. In most cases, it is better to specify RADIUS clients with IP addresses. When you use IP addresses, IAS is not required to resolve host names at startup and starts more quickly as a result. This is especially beneficial if your network contains a large number of RADIUS clients. You can use DNS names to specify RADIUS clients when you require something other than administrative flexibility (for example, the ability to map multiple RADIUS client IP addresses to a single DNS name).

If you are using Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, you can specify a RADIUS client by using an IP address range. All of the RADIUS clients in the range must use the same configuration and shared secret. This configuration is useful, for example, if you place many wireless access points on the same subnet.

The address range for RADIUS clients is expressed in the network prefix length notation *w.x.y.z/p*, where *w.x.y.z* is the dotted decimal notation of the address prefix and *p* is the prefix length (the number of high order bits that define the network prefix). This is also known as *Classless Inter-Domain Routing (CIDR) notation*. An example is 192.168.21.0/24. For more information, see "Configure RADIUS Clients" in Help and Support Center for Windows Server 2003.

**Plan Authentication**

When designing network authentication solutions, protecting the authentication channel can prevent potential security attacks.

Most password-based authentication methods, such as CHAP and MS-CHAP, do not provide privacy-protected channels to guard against offline dictionary attacks by hackers that intercept authentication traffic on the network. Because of this, ensure that password-based authentication methods are always deployed with protection from IPSec or PEAP.

**EAP-TLS**
Certificate-based authentication methods are much more secure than password-based authentication methods. Use certificate-based authentication methods, such as PEAP and EAP, in all possible circumstances because they protect the authentication channel and provide strong security. EAP-TLS is designed, in part, to protect against spoofing or other attacks and can be deployed without protection from IPSec or PEAP. EAP-TLS requires a public key infrastructure (PKI) and is an authentication method that can be used with all connection types supported by IAS (wireless, authenticating switch, VPN, and dial-up).

By using Group Policy in Windows Server 2003, you can easily distribute certificates to clients and servers with auto-enrollment. For the maximum strength user credentials, deploy EAP-TLS with a smart cards. Smart cards provide maximum strength with two-factor authentication. Certificates installed in the computer certificate store (without smart cards) offer single-factor authentication.

**PEAP**
PEAP uses Secure Sockets Layer (SSL) technology to privacy-protect authentication communications, and to key the encryption of link layer network connections. You can deploy PEAP with either EAP-MS-CHAPv2 or EAP-TLS as the authentication type. PEAP-EAP-MS-CHAPv2 is a secure password authentication method recommended for wireless deployments. However, PEAP is not available for VPN or dial-up connections.

PEAP provides protection for the EAP method negotiation that occurs between client and server through a TLS channel. This helps prevent an attacker from injecting packets between the client and the access server to cause the negotiation of a less secure EAP method.

Clients using PEAP as the authentication method have the ability to authenticate the IAS or RADIUS server. Because the server also authenticates the client, mutual authentication occurs.

PEAP fast reconnect, which reduces the delay in time between an authentication request by a client and the response by the IAS or RADIUS server, allows wireless clients to move between wireless access points without repeated requests for authentication. This reduces resource requirements for both client and server, and allows users to move between access points without reentering credentials.

If you deploy PEAP, do not deploy the same authentication type both inside of the PEAP Transport Layer Security (TLS) channel and outside of the PEAP TLS channel. For example, do not deploy both PEAP-EAP-TLS and EAP-TLS on the same network. You can deploy different authentication types inside and outside the TLS channel. For example, you can deploy PEAP-EAP-MS-CHAPv2 and EAP-TLS on the same network.

**Authentication and PPP and PPTP Connections**
For dial-up Point-to-Point Protocol (PPP) or Point-to-Point Tunneling Protocol (PPTP) VPN connections, it is recommended that you use EAP-TLS with smart cards or certificates as the authentication method.

For L2TP/IPSec VPN connections, you can use MS-CHAPv2 as the authentication method. Internet Protocol security (IPSec) uses computer certificates to establish a secure channel before authentication proceeds, providing the necessary protection for authentication and other communication.

When planning authentication:

- For wireless connections, you can configure all the connection properties on the client (including authentication methods) using Windows Server 2003 Group Policy. For example, you can configure wireless connections to specific networks to require use of EAP-TLS.

- You can use Connection Manager Administration Kit (CMAK) to create a Connection Manager profile for installation on client computers. With Connection Manager, you can manage the client connection properties (including authentication methods) used for access to your network.

- Configure remote access policies at the IAS server to only allow the authentication methods you want to allow per connection type, such as dial-up or VPN. .

For more information, see "Authentication methods" and "Public Key Infrastructure" in Help and Support Center for Windows Server 2003.

**Add RADIUS or VSA Attributes to Remote Access Policy**
If you plan to return additional RADIUS attributes or VSAs with the responses to RADIUS requests, you must add the RADIUS attributes or VSAs to the appropriate remote access policy.

**Install Backup IAS Servers**

To provide fault tolerance for RADIUS-based authentication and accounting, you must always use at least two IAS servers working together in the same location. One IAS server is used as the primary RADIUS server, and the other is used as a backup. RADIUS proxies are configured to use both IAS servers. When the primary IAS server becomes unavailable, the RADIUS proxies automatically use the backup IAS server instead.

# Design IAS as a RADIUS Proxy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing IAS as a RADIUS proxy involves some or all of the following tasks:

- Planning the connection request policy
- Adding RADIUS and VSA attributes
- Planning for load balancing and failure detection
- Installing backup IAS proxies

For more information about configuring IAS as a RADIUS proxy, see "Deploy IAS as a RADIUS Proxy" later in this chapter and "Deploy IAS as a RADIUS Proxy" in Help and Support Center for Windows Server 2003.

**Plan the connection request policy**
The default connection request policy **Use Windows authentication for all users** is configured for IAS when it is used as a RADIUS server. To create a connection request policy to use IAS as a RADIUS proxy, complete the following steps:

1. Create a remote RADIUS server group on the domain that will authenticate the users.
2. Create a connection request policy that forwards authentication requests to the remote RADIUS server group.
3. Either delete the default connection request policy, or set the new connection request policy first in the order before the default connection request policy so that it is evaluated first.

**Add RADIUS attributes and VSAs**
If you plan to return additional RADIUS attributes and VSAs with RADIUS requests, you must add the RADIUS attributes and VSAs to the appropriate connection request policy.

**Plan for load balancing and failure detection**
When you configure multiple servers in a remote RADIUS server group, you can configure settings that determine how the IAS proxy server balances the load of authentication and accounting requests over the RADIUS servers in the group. You can use additional settings to configure IAS to detect and recover from the failure of a remote RADIUS server group member.

For more information about load balancing for remote RADIUS server group members, see "Configure the load balancing properties of a group member" in Help and Support Center for Windows Server 2003.

**Install backup RADIUS proxies**
To provide fault tolerance for RADIUS-based authentication and accounting, you must always use at least two IAS proxy servers. One IAS server is used as the primary RADIUS proxy, and the other is used as a backup. RADIUS clients (access servers or other RADIUS proxies) are configured on both IAS proxy servers. In addition, configure both the primary and backup IAS proxy servers on each RADIUS client. When the primary IAS proxy becomes unavailable, the access servers automatically use the backup RADIUS server instead.

For more information about synchronizing the configuration of multiple IAS servers, see "Managing multiple IAS servers" in Help and Support Center for Windows Server 2003.

# Designing Support for Client Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have designed your IAS infrastructure and determined what role IAS will play in your network, you must consider what types of network and remote access clients you will need to support. For each type, your IAS server configuration will differ, and you will need to make different planning decisions.

Windows Server 2003 IAS supports the following forms of network access:

- Dial-up remote
- VPN
- Wireless
- Authenticating switch

**Note**

- You can use the Routing and Remote Access service included with Windows Server 2003 or the Microsoft® Windows® 2000 operating system to provide network access. However, you can also use third-party network access servers with Windows Server 2003 IAS. To ensure that the third-party access server works with Windows Server 2003 IAS, check with the manufacturer to confirm that the product supports RFCs 2865, 2866, and 2869.

# Designing Support for Dial-up Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 has extensive support for remote access technology to connect dial-up and other types of remote clients to corporate networks or to the Internet. For information about how to plan, design, and implement Windows Server 2003 components (such as Routing and Remote Access Service and Connection Manager) in an enterprise environment, see "Deploying Remote Access Clients Using Connection Manager" in this book.

IAS supports both voluntary and compulsory tunneling. Table 7.1 describes the differences between both types of tunneling and when you might use each type.

**Table 7.1 Comparison of Voluntary and Compulsory Tunneling**

| Tunnel Type | When to Use It | Which IAS Components to Use |
|---|---|---|
| Voluntary tunneling | Use this option if your clients need to choose their tunneling location themselves.<br><br>For example, the user dials in to an ISP. At the client's request, the ISP creates a tunnel to the corporation. The user can alternatively request a tunnel to somewhere else, such as the Internet. | The ISP uses IAS as a RADIUS proxy to forward the request to the corporate IAS server.<br><br>The corporation uses IAS as a RADIUS server to authorize and authenticate the request.<br><br>Either the ISP or the corporation can use a third-party RADIUS server. |
| Compulsory tunneling | Use this option if you want to use one tunnel for many clients.<br><br>For example, if a corporation has clients in geographically dispersed locations, it can contract with an ISP to deploy regional tunnel servers. Any client can dial in to any of the tunnel servers, which then creates a compulsory tunnel to the corporation. Compulsory tunneling is typically used for dial-up clients, but can also be used for VPN clients. | The ISP uses IAS as a RADIUS server to authorize the request.<br><br>The corporation uses IAS as a RADIUS server to authorize and authenticate the request.<br><br>Either the ISP or the corporation can use a third-party RADIUS server. |

**Voluntary Tunneling**

In voluntary tunneling, during the authorization phase, the corporate IAS server restricts the client connection to use a specificVPN protocol (PPTP or L2TP/IPSec) if an administrator has configured remote access policies to restrict connections to those using the specified protocol. The designated protocol must be installed on the client or the connection attempt is rejected. After the access request is authenticated and authorized, the client establishes a VPN tunnel to the corporate access server.

A user or client computer can issue a VPN request to configure and create a voluntary tunnel. In this case, the user's computer is a tunnel endpoint that acts as the tunnel client. Voluntary tunneling occurs when a workstation or router uses tunneling client software to create a VPN connection to the target tunnel server. In order to accomplish this, the appropriate tunneling protocol must be installed on the client computer. In a dial-up situation, which is the most common use, the client must establish a dial-up connection to the internetwork before the client can set up a tunnel. A good example of this is the dial-up

Internet user, who must dial an ISP and obtain an Internet connection before a tunnel over the Internet is created.

Voluntary tunneling is not different from other types of network access, and IAS can be used for authentication, authorization, and accounting.

**Compulsory Tunneling**

Compulsory tunneling is the creation of a secure tunnel by another computer or network device on the client computer's behalf. Compulsory tunnels are configured and created automatically for users without their knowledge or intervention. With a compulsory tunnel, the user's computer is not a tunnel endpoint. Another device between the user's computer and the tunnel server is the tunnel endpoint, which acts as the tunnel client.

The computer or network device that provides the tunnel for the client computer is known as a front-end processor (FEP) in PPTP, an L2TP access concentrator (LAC) in L2TP, or an IP security gateway in IPSec. The term *FEP* is used to describe tunnel creation functionality, regardless of the protocol used. To perform its function, the FEP must have the appropriate tunneling protocol installed and must be capable of establishing the tunnel when the client computer attempts a connection. In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000, the Routing and Remote Access service cannot be used as a FEP.

An organization can contract with an ISP to deploy a nationwide set of FEPs. These FEPs can establish tunnels across the Internet to a VPN server that is connected to the organization private network, thereby consolidating calls from geographically diverse locations into a single Internet connection at the organization network.

There are two types of compulsory tunneling. In the first type, the tunnel is created before the access client is authenticated. Based on the realm name or the caller ID of the access client, the FEP sends an Access-Request message to an IAS server. The IAS server sends back an immediate Access-Accept message with RADIUS attributes for the tunnel creation without performing authentication and authorization. After the tunnel is created, the access client authenticates against the tunnel server.

In the second type of compulsory tunneling, the tunnel is created after the access client is authenticated by the FEP. In this case, the FEP sends the Access-Request message with the client credentials to an IAS server. The IAS server authenticates and authorizes the connection attempt and returns RADIUS attributes in the Access-Accept message, which specify to the NAS how to initiate a tunnel to a VPN server. The tunnel endpoint (the VPN server at which the tunnel is terminated), can be changed on the basis of conditions in a remote access policy. For example, the tunnel endpoint can be changed on the basis of the user name or the user account group membership. Controlling compulsory tunnels with remote access policies provides more flexibility than static tunneling (which requires a dedicated access server) or realm-based tunneling (which requires all users in a specific realm to use the same tunnel settings).

For more information, see "IAS and tunnels" in Help and Support Center for Windows Server 2003.

# Designing Support for VPN Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Depending on the credentials of the authenticating user, you can configure IAS to direct the traffic from a VPN client through a tunnel to specific parts of the enterprise network. Both voluntary and compulsory tunneling can be used for outsourced VPN access.

**Special Considerations for VPN Access**
Consider the following access server compatibility issues:

- IAS can provide compulsory tunneling connection attributes for network access servers (NASs) that support compulsory tunneling. However, the Windows Server 2003 Routing and Remote Access service does not support compulsory tunneling.

- VPN servers running the Microsoft® Windows NT® version 4.0 operating system must also be running Routing and Remote Access Service (RRAS) if you want to configure them as RADIUS clients.

For more information about Internet Authentication Service, including the authorization and authentication process used in voluntary and compulsory tunneling, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

For information about the RADIUS attributes used with compulsory tunneling, see "Compulsory tunnels" in Help and Support Center for Windows Server 2003.

# Designing Support for Wireless Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Wireless networking technology is becoming more widespread with the adoption of industry standards such as IEEE 802.11 and 802.11b. Wireless networking allows a user to roam around a building or campus and automatically connect to the network after the user is in proximity to a wireless access point (wireless AP).

You can use IAS to support wireless access in the following scenarios:

- Wireless LAN access
- Outsourced wireless access

**Countering Wireless Security Risks**

While providing convenience, wireless networking technologies and wireless APs represent the following security risks:

- Anyone who has a compatible wireless network adapter can gain access to the network.
- Wireless networking signals use radio waves to send and receive information. Anyone within an appropriate distance to a wireless access point can detect and receive all data sent to and from the wireless access point.

IAS provides enhanced security for wireless access. To counter the first security risk, you can set up the wireless access point as a RADIUS client, and then configure it to send access requests and accounting messages to a central RADIUS server running IAS.

To counter the second security risk, you can encrypt the data sent between the wireless devices and the wireless access points. The authentication method used by the wireless client must be able to use encryption keys.

**Special Considerations for Wireless Access**

If you will be supporting wireless access, include the following elements in your design:

- An authentication mechanism. With wireless access, you can use Protected Extensible Authentication Protocol (PEAP), EAP-TLS, or unauthenticated wireless access. For more information about PEAP, see "PEAP" in Help and Support Center for Windows Server 2003.
- Certificates and a certification authority, if you are deploying authentication methods that use certificates on both the client and server, such as EAP-TLS. For more information, see "Integrate IAS with the Certificate Infrastructure" later in this chapter.
- The **Wireless-IEEE 802.11** and **Wireless-Other** port types for the **NAS-Port-Type** condition of remote access policies. By using these port types, you can create a separate remote access policy that contains connection parameters and encryption settings specifically designed for wireless devices.
- The **Ignore-User-Dialin-Properties** attribute in the profile settings of a remote access policy. The dial-in properties of the user account are designed for clients dialing into an access server, not for clients connecting to a wireless port or authenticating switch. You can disable them on the remote access policy. For more information about profile settings, see Dial-in properties of a user account in Help and Support Center for Windows Server 2003.

# Designing Support for Authenticating Switch Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A network switch filters the traffic received on each port of the switch and allows better traffic management by segmenting a network. A typical switch, however, sends and receives packets to any node that is connected to it. This can create a security risk, especially for switch ports in a conference room of an organization, which might be accessed by visitors or employees of partner organizations.

**Securing Authenticating Switch Access with IAS**

To prevent physical layer access to the network, use authenticating switches as RADIUS clients and use the industry-standard RADIUS protocol to send access requests and accounting messages to a central RADIUS server. The RADIUS server has access to a user account database and a set of rules for determining whether to grant authorization.

If you need to grant access to different types of users for different parts of your network, you can use authenticating switch access. For example, a corporate conference room has a virtual local area network (VLAN) for visitors. A visitor logging on without presenting credentials is granted access to the conference room VLAN. At the same time, an employee logging on with appropriate credentials is granted access to the entire corporate network. Thus, administrators can use IAS with authenticating switch access to secure parts of the network from malicious users and from inexperienced users who might inadvertently cause errors in network configuration. You can also use VLANS with IAS for wireless access configurations.

Because the data sent between the node and the authenticating switch physically travels on a dedicated wire, encryption of the data is not required or typically implemented.

**Special Considerations for Authenticating Switch Access**

If you will be supporting authenticating switch access, include the following elements in your design:

- An authentication mechanism. You can use EAP-TLS or secure password-based authentication by using PEAP-EAP-MS-CHAPv2. For more information, see "Select Authentication Protocols" later in this chapter.

- Certificates and a certification authority, if you are deploying authentication methods that use certificates on both the client and server, such as EAP-TLS. For more information, see "Integrate IAS with the Certificate Infrastructure" later in this chapter.

- The use of the **Ethernet** port type for the **NAS-Port-Type** condition of remote access policies. By using this port type, you can create a separate remote access policy that contains connection parameters specifically designed for nodes connecting to the network through an authenticating switch. For more information, see "Elements of a remote access policy" in Help and Support Center for Windows Server 2003.

- The **Ignore-User-Dialin-Properties** attribute in the profile settings of a remote access policy. The dial-in properties of the user account are designed for clients dialing into an access server, not for clients dialing into a wireless port or authenticating switch. You can disable them on the remote access policy. For more information, see "Dial-in properties of a user account" in Help and Support Center for Windows Server 2003.

# Designing an Optimized IAS Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Optimize your IAS design by planning how to scale your IAS servers, whether or not to add IAS servers, where to place your IAS servers, and other steps as illustrated in Figure 7.9.

**Figure 7.9 Designing an Optimized IAS Solution**




IAS RADIUS clients and servers require minimal management and administration. However, over time, changes in the number of access clients, changes in WAN technology, and other factors can reduce the performance of IAS.

You can optimize IAS performance by positioning your IAS servers strategically. Use the following guidelines when deciding where to position your IAS servers:

- Locate IAS servers in the same domain with the server that provides remote user account authentication.
- Locate IAS on a domain controller and store the user account database in Active Directory.

In addition, the following factors can negatively impact IAS performance:

- The current load of the domain controller.
- The resolution of user principal names, resulting in an additional remote procedure call (RPC) query against the computer that contains the global catalog.
- EAP-based authentication methods, involving multiple challenge-response exchanges.
- The type of hardware in use.
- Network latency between:
    - The IAS server and the domain controller.
    - The IAS server and the computer that contains the global catalog.

- The IAS server and the access server.

You can optimize the performance of an IAS solution by scaling IAS to meet increasing demands in your organization and by including more than one RADIUS client and server in your network design.

# Scale an IAS Server to Optimize Performance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

IAS servers can accommodate a large volume of accounts and authentications per second. For this reason, they perform well in large organizations. A single IAS server can often meet the dial-up and VPN connection needs for an entire network. However, if you have more complex network access needs, you might need to scale even a single IAS server. For example, certificate-based authentication can generate a large volume of authentication traffic that is transmitted every time a client logs onto the network. If you use certificate-based authentication, you must scale IAS to meet all of your Internet authentication needs.

For more information about scaling IAS, see "Deploying IAS as a RADIUS Proxy" in Help and Support Center for Windows Server 2003.

# Add IAS Servers to Optimize Availability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to prevent an IAS server from becoming a single point of failure.

To prevent an IAS server from becoming a single point of failure, deploy IAS servers either in pairs (a primary and a backup IAS server) to provide fault tolerance, or in multiples, to balance the load of a large volume of authentication and accounting requests. Add at least two additional IAS servers per forest.

The level of availability that you need to provide depends on the following:

- How likely it is that a network link will fail
- How important it is to provide uninterrupted authentication to your clients at all times

Balance your need for availability against your need to minimize the costs of the hardware investment. Keep in mind that you can add IAS to existing domain controllers.

When you deploy more than one IAS server, make sure to design your IAS servers to use RADIUS proxies for load balancing and synchronization of the server configurations.

# Use Load Balancing with Multiple IAS Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you add multiple IAS servers for performance reasons, you must load balance the servers. You can load balance and provide failover at each access server by configuring the access server to send queries to multiple RADIUS servers in a specified order of importance. You can also load balance by using IAS as a RADIUS proxy server that forwards connection requests to IAS servers in groups called remote RADIUS server groups. These configurations are useful for the following:

- Organizations that use EAP-TLS for authentication, which increases the load on RADIUS proxies and servers. For example, ISPs and corporations that want to provide wireless or authenticating switch access often use EAP-TLS.
- Organizations that need to sustain continuous service availability.
- ISPs that outsource VPN access for other organizations. The outsourced VPN services can generate a large volume of authentication traffic.

You can use RADIUS proxies to balance the load of a large volume of authentication traffic. Without RADIUS proxies, each network access server balances its RADIUS requests across multiple RADIUS servers and detects unavailable RADIUS servers.

When RADIUS proxies are in place, the load of authentication, authorization, and accounting traffic is distributed across all of the IAS servers in the organization. Additionally, there is a consistent scheme for failure detection and RADIUS server failover.

# Duplicate IAS Server Configurations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you deploy more than one IAS server to provide the same authentication, authorization, and accounting service to RADIUS clients and proxies, you must copy the configuration of one IAS server computer to the other IAS servers. To duplicate the configuration of one IAS server to multiple IAS servers, use the IAS snap-in. This duplication method is useful when the number of configuration changes is small or if you are duplicating the configuration to only a few IAS servers. You can use the snap-in to manage both local and remote IAS servers. If you make a configuration change to one IAS server, you must make the same configuration change to all of the IAS servers that provide the same service.

To duplicate one IAS server configuration when there are a large number of configuration changes or a larger number of IAS servers, you can copy the configuration of one IAS server to another IAS server in the following way:

- Make configuration changes on the primary IAS server.

- On the primary IAS server, use the **netsh aaaa dump** command to export the configuration of one IAS server to a Netsh script file. The **dump** command displays the configuration of the IAS database file (Ias.mdb) as a Netsh command script that you can use to duplicate the configuration of the server on which the command is executed. The Netsh command script contains the configuration of the IAS server, including the registry keys and database file (Ias.mdb), in a compressed text format as a large data block. This large data block is used by the **set config** command within the script to import the configuration of a saved data block into an existing IAS database on the same or another computer, which you can perform by using the **netsh exec** command. To save the Netsh command script to a file, type: **netsh aaaa show config >***Path\File.txt*

- On the target computers, use the **netsh exec** command to import the primary IAS server configuration to the other IAS servers.

By using these two Netsh commands, you can automate the process in a simple batch file or script for multiple IAS servers.

Use this method to manage RADIUS and remote access policy configurations in a large enterprise network.

The **netsh aaaa** commands also provide a way to export and import individual aspects of the IAS server configuration rather than the entire configuration. For example, you can export and import only the remote access policies, or you can export and import only the RADIUS clients configured on a server.

For more information, see "Netsh" and "Netsh commands for AAAA" in Help and Support Center for Windows Server 2003.

# Optimize RADIUS Client Performance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because RADIUS clients, such as VPN servers and dial-up access servers, manage access client connections, their performance directly affects access client performance. Providing highly available RADIUS clients ensures that access clients can reliably connect to the network and can always be authenticated by IAS.

You can optimize performance for access clients that are using all types of connections by doing any of the following:

- Upgrading the hardware of existing RADIUS clients. Choose this option if your hardware is outdated, and it is cost effective to upgrade it.
- Replacing existing RADIUS clients with higher-performance servers. Choose this option if you cannot upgrade the hardware of your existing resources.
- Adding RADIUS clients. Choose this option if your hardware is up-to-date, but you require more fault tolerance. Be sure to register redundant RADIUS clients with the IAS servers to ensure proper authentication and accounting.

You can optimize the performance of RADIUS clients that support dial-up connections by assigning remote access clients different primary and backup telephone numbers. This ensures that the remote access clients connect to different RADIUS clients. Choose this option if you require additional load balancing.

# Optimize Authentication and Accounting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

RADIUS servers provide authentication and accounting for RADIUS clients, and interact with the authentication servers. As a result, RADIUS server performance and availability impacts authentication and accounting performance.

To optimize authentication in your network, ensure that all redundant IAS servers use the same user account database, thereby ensuring that the user accounts are available for authentication. Also, specify that RADIUS clients use the redundant IAS servers to ensure proper authentication and accounting.

In larger organizations with complex forest or domain topologies, use IAS as a RADIUS proxy to forward authentication requests to remote RADIUS server groups. You can also designate remote RADIUS server groups to process only accounting requests, freeing the servers performing authentication from handling accounting traffic.

- To optimize authentication and accounting performance in your IAS design, take the following actions:

    - Run IAS on the same computer as the domain controller. This speeds IAS access to the Active Directory user accounts database when IAS is performing user authentication and authorization.

    - Run IAS on the same computer that contains the global catalog. If it is not possible to run IAS on the same computer as the domain controller or the computer that contains the global catalog, verify that you have an efficient domain and site topology, and place the IAS server on the same subnet as a domain controller or global catalog server.

    - Reduce the number of user accounts in each domain by redesigning your domain topology.

    - Add IAS proxy servers to load balance authentication and accounting between servers in remote RADIUS server groups.

    - Upgrade the hardware resources of the existing IAS servers.

    - Replace existing IAS servers with higher performance servers.

    - Reduce the level of detail recorded in IAS accounting. IAS accounting can record user authentication requests, accounting requests, and periodic data. Make sure you are logging only the amount of information you need to troubleshoot network access.

    - If you configure IAS accounting for SQL Server logging, install SQL Server Desktop Engine (MSDE 2000) on the IAS server, and log to MSDE 2000 instead of directly to SQL Server 2000 running on another computer. This configuration assists in preventing logging failure due to network hardware failure or heavy network traffic. Use a custom application, service, or component to periodically publish the accounting logs from the MSDE 2000 database on each IAS server to the master SQL Server 2000 database.

    - For wireless deployments, use PEAP-EAP-MS-CHAPv2 with fast reconnect. PEAP uses cached TLS keys during re-authentication with access points configured as RADIUS clients of a single IAS server. Cached authentication is critical for wireless deployments because wireless clients authenticate each time they move to and associate with a new access point. In addition to improving performance, PEAP fast reconnect significantly reduces the latency of authentication and the public key operation overhead on both the client and the RADIUS server.

    - Use the **MaxConcurrentApi** registry entry (**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\ Parameters**) to increase the number of multiplexed connections to the domain controller.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

For more information, see "Remote Access Logging" in Help and Support Center for Windows Server 2003.

# Optimize IAS Deployment in Large Organizations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can optimize IAS performance in large organizations by doing the following:

- If you are using remote access policies to restrict access for all but certain groups, create a universal group for all of the users to whom you want to allow access, and create a remote access policy that grants access for that universal group. If you have a large number of users on your network, create global groups within the universal group, and add the users to the global groups.

- Use IAS as a proxy server and configure connection request policies to distribute authentication requests to remote RADIUS server groups based on the realm name portion of the user name. In this manner you can load balance traffic based on domain membership, and authentication requests are sent to a remote RADIUS server group that resides in the same domain where the user account is located.

- Install IAS as a dedicated RADIUS server. If you choose not to run the IAS server on a domain controller with a global catalog, you can run it on a computer that has other services running on it, as long as the services are not resource-intensive.

In very large environments (such as an ISP with millions of remote access users and extremely heavy load conditions) that must process a large number of both authentication requests and accounting packets per second, you can optimize IAS performance by doing the following:

- Using a faster domain controller to yield better throughput. The number of authentications per second depends on the hardware used for the domain controller.

- Using separate IAS servers for authentication and accounting. IAS proxy servers can send all accounting requests to a specific remote RADIUS server group, while sending authentication requests to other groups. For more information, see "Configure accounting" in Help and Support Center for Windows Server 2003.

- Running the IAS server on a domain controller with a global catalog. Choose this option if you have a high-latency connection between your IAS server and your domain controller, or between your IAS server and your global catalog, but you do not have problems with your IAS performance.

- Increasing the number of concurrent authentication calls in progress at one time by using the **MaxConcurrentApi** registry entry. Keep in mind that if you assign too high a value to this registry entry, your IAS server can place an excessive load on your domain controller. Values from 2 to 5 provide the best performance.

  For more information about the **MaxConcurrentApi** registry entry, see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

# Creating a Remote Access Policy Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remote clients connect to the network by using remote access policies. Plan an authorization method for security, and then plan your user groups and user accounts. Next, create your remote access policies to centralize management. Figure 7.10 illustrates this process.

**Figure 7.10 Creating a Remote Access Policy Strategy**



# Establish a Remote Client Access Authorization Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use remote access policies to authorize remote client access either by group or by user.

If you authorize access by group, you can restrict network access based on an account's group membership. For example, you can create one group for wireless users and another group for dial-in users, and create separate remote access policies for each group by using the **Windows-Groups** attribute (condition) of the remote access policy.

To authorize access by group in a Windows Server 2003 or a Windows 2000 native domain, set the remote access permission on every user account to **Control access through Remote Access Policy**. Remote access permissions are then determined by the remote access permission setting on the remote access policy.

For more information, see "Remote Access Policies" in Help and Support Center for Windows Server 2003.

# Plan Remote Access Policy Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can manage remote access authorization by user or by group. By using the Active Directory Users and Computers snap-in, you can create groups to which specific IAS remote access policies are applied, allowing you to grant different types of remote access to different groups. The functional level of the domain impacts the type of group that you can use.

In Windows Server 2003 domains or Windows 2000 native-mode domains, you can use universal and global groups as follows:

- Universal groups can include groups and accounts from any Windows Server 2003 or Windows 2000 domain in the domain tree or forest and can be granted permissions in any domain in the domain tree or forest.

- Global groups can include groups and accounts only from the domain in which the group is defined and can be granted permissions in any domain in the forest.

In Windows Server 2003 interim domains or Windows 2000 mixed-mode domains, you can use only global groups. Domain local groups cannot be used because they are created in a single domain, are visible only in that domain, and can be assigned permissions only to resources within that domain.

For more information about remote access policies, see "Remote Access Policies" in Help and Support Center for Windows Server 2003.

For more information about functional levels, see "Domain and forest functionality" in Help and Support Center for Windows Server 2003.

# Plan the User Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, the user account for a stand-alone server or server running Active Directory contains a set of dial-in properties that are used when performing authorization (allowing or denying a connection attempt made by a user). On a stand-alone server, you can set the dial-in properties on the **Dial-in** tab in the user account in the **Local Users and Groups** dialog box. On a server running Active Directory, you can set the dial-in properties on the **Dial-in** tab in the user account in the **Active Directory Users and Computers** snap-in. On these servers, you cannot use the Windows NT 4.0 User Manager for Domains administrative tool.

In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition, you can configure a RADIUS attribute to ignore the dial-in properties of user and computer accounts in the profile properties of a remote access policy. To support multiple types of connections for which IAS provides authentication and authorization, it might be necessary to disable the processing of user account dial-in properties. This can be done to support scenarios in which specific dial-in properties are not required and is accomplished by configuring the **Ignore-User-Dialin-Properties** attribute on the **Advanced** tab of the profile settings for a remote access policy.

Computer accounts also have dial-in properties that are similar to user accounts. Therefore computers can be authenticated as if they were users when an IEEE 802.1X Ethernet client uses EAP-TLS and an installed computer certificate to authenticate itself to an Ethernet switch. Note that the **Ignore-User-Dialin-Properties** attribute disables the use of all dial-in properties for the user account. Specific dial-in properties cannot be individually disabled.

For more information about planning user accounts, see "Dial-in properties of a user account" in Help and Support Center for Windows Server 2003.

# Configure Remote Access Policies for Secure Remote Connections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Remote Access Policies**

You can create multiple remote access policies to accommodate the individual security requirements of your remote connections. For example, you can use remote access policies to grant or deny authorization according to the time of day and the day of the week, the Windows Server 2003 or Windows 2000 Server group to which the remote access user belongs, the type of connection being requested (dial-up networking or VPN connection), and so on. You can also configure settings within the remote access policies that limit the maximum session time, specify the authentication and encryption strengths, set Bandwidth Allocation Protocol (BAP) policies, and so forth.

Remote access policies are applied in the order in which they are listed. When a policy is found that contains conditions that match the connection attempt, access is granted or denied, regardless of the conditions specified by the other policies later in the list. Therefore, it is a good idea to list more specific remote access policies before general remote access policies.

If you do not implement any remote access policies, all connection attempts fail.

Before you configure remote access policies, you must make decisions about the following:

- Whether to use Network Access Quarantine Control for VPN and dial-up connections.
- Whether to use custom or common policies. When you use the New Remote Access Policy Wizard in the IAS snap-in, you can choose to create a common or a custom policy. For a common policy, you must configure an access method, whether to grant access permissions by user or by group, authentication methods, and levels of allowed encryption (depending on the access method selected). For a custom policy, you must configure a set of policy conditions, whether remote access permission for the policy is granted or denied, and remote access policy profile settings. For more information, see "Add a remote access policy" in Help and Support Center for Windows Server 2003.
- The groups and users to which the remote access policies apply.
- Whether the remote access policy grants or denies access to the users or the group.
- The restrictions that are placed on the users or the group.

For more information about Internet Authentication Service, including remote access policies, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Using Network Access Quarantine Control

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Network Access Quarantine Control provides phased network access for client computers that connect to the network by using VPN or dial-up remote access. By using connection request policies and remote access policies, you can apply the MS-Quarantine-IPFilter attribute, the MS-Quarantine-Session-Timeout attribute, or both attributes to restrict clients to quarantine mode. During quarantine mode, an administrator-provided network policy requirements script is run on the client computer. After the client computer configuration is either brought into or determined to be in accordance with your organization's network policy, quarantine mode is removed by the network access server, and standard remote access policy is applied to the connection.

You can implement Network Access Quarantine Control with one or more servers running Windows Server 2003 and the Routing and Remote Access service, one or more servers running Windows Server 2003 and IAS, a Connection Manager (CM) profile created with Connection Manager Administration Kit (CMAK), an administrator-provided script, and two additional components: the notifier component and the listener component. You can create your own notifier and listener components, or you can use Rqs.exe (a listener component) and Rqc.exe (a notifier component), which are available in the Windows Deployment and Resource Kits.

For more information, see "Deploying Dial-up and VPN Remote Access Servers" and "Deploying Remote Access Clients Using Connection Manager" in this book.

# Using Common vs. Custom Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must decide whether to use common or custom policies:

- A common policy includes typical settings for a particular access method.
- A custom policy includes a detailed configuration of a particular access method.

**Create Specifications for a Common Policy**
To create a common or custom policy, you must specify the following:

- An access method. The access method is used to configure the access server Port Type condition automatically. You can choose one of the following four access methods:
    - VPN access
    - Dial-up access
    - Wireless access
    - Ethernet switch access
- Whether to grant access permissions by user or by group. If you choose to grant access by group, the **Windows-Groups** condition is automatically set to the chosen groups.
- Authentication methods. For more information about how to configure authentication methods, see "Configure authentication" in Help and Support Center for Windows Server 2003.
- Levels of allowed encryption (depending on the access method chosen). For more information about how to configure encryption, see "Configure encryption" in Help and Support Center for Windows Server 2003.

When you create a common policy, the remote access permission is always set to **Grant remote access permission**.

**Create Specifications for a Custom Policy**
For each custom policy, create detailed specifications for the following elements:

- Conditions. Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt.
- Remote access permission. Specify whether the permission is granted or denied if the conditions of a remote access policy are met.
- Profile. Specify the remote access policy profile properties to set dial-in constraints and other restrictions.

For specific information about settings for each element, see Help and Support for Windows Server 2003 or the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

**Note**

- Not all network access servers send all of the RADIUS attributes. Consult the documentation for your network access server to see which attributes it sends.

**Conditions**
Specify the remote access conditions for each policy. Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt. In order for the connection attempt to match the policy, all conditions must match the settings of the connection attempt. Remote access policy profile settings are applied only if the connection attempt matches the policy. Thus, remote access policies are applied only if the connection attempt matches all of the conditions of the policy.

**Permission**
Specify whether the permission is granted or denied if the conditions of a remote access policy are met. You use the **Grant remote access permission** option or the **Deny remote access permission** option to set remote access permission for a policy.

During the authorization process, the dial-in properties of user accounts are evaluated before remote access policy is applied. If the dial-in properties for the user account are set to **Deny access**, the connection attempt is rejected and remote access policies are not evaluated. When dial-in properties for the user account are set to **Control access through Remote Access Policy**, the remote access policy alone determines whether the user is granted access.

When the dial-in properties for the user account are set to **Grant access**, remote access policies are evaluated next. In this circumstance it is possible for the user to be denied access by settings in the remote access policy. For example, if the remote access policy is configured to allow the user to connect only between the hours of 8 AM and 5 PM and the user is attempting to connect at 6 PM, the connection attempt fails due to the settings in the remote access policy.

**Profile**
Specify the remote access policy profile properties to set dial-in constraints and other restrictions. These properties are applied to a connection after the connection is authorized, whether the connection has been authorized through the user account permission setting or the remote access policy.

You can use these properties to specify the series of RADIUS attributes that are sent back to the RADIUS client by the IAS server, including any vendor-specific attributes you are using. For more information about VSAs, see "Configuring IAS for Compatibility with Third-Party Access Servers" later in this chapter.

### Note

- Elements of a remote access policy correspond to RADIUS attributes that are used during RADIUS-based authentication. For an IAS server, verify that the network access servers that you use are sending the RADIUS attributes that correspond to the configured remote access policy conditions and profile settings. If an access server does not send a RADIUS attribute that corresponds to a remote access policy condition or profile setting, then all RADIUS authentications from that access server are denied.

# Apply Policies to Users and Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are controlling remote client access by means of groups, you must determine which groups the policies that you create will apply to. If you are using nested groups, make sure that policies configured for the parent group do not conflict with policies configured for the nested group.

If you are controlling remote client access by means of the access-by-user administrative model, you must establish the criteria that you use to determine whether the policy is applied; for example, you can use the client name, client IP address, access server IP address, or any other condition attribute.

Two of the attributes, **NAS-Port-Type** and **Tunnel-Type**, deserve special mention because they are used to support VPN, wireless, and authenticating switch access clients:

- The **NAS-Port-Type** attribute. This attribute allows you to specify the type of physical port that is used by the access server originating the request.

- The **Tunnel-Type** attribute. This attribute allows you to specify the type of tunnel that the access server can create. You can specify either Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP).

For a detailed list of condition attributes, see "Connection request policies" in Help and Support Center for Windows Server 2003 and the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Determine Which Restrictions to Apply

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can specify dial-in conditions or any other profile property to restrict access. For a detailed list of profile properties that can be used to restrict access, see "Introduction to remote access policies" in Help and Support Center for Windows Server 2003 or the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

To achieve the highest level of security possible, begin by granting the least amount of access possible and increase access only when necessary.

# Configure Client-Specific Remote Access Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Different types of remote access clients need different settings in their remote access policies. The following sections describe design considerations for creating the different types of policies.

**Configure Remote Access Policies for VPN Clients**
For VPN clients, use the **Tunnel-Type** attribute to specify the type of tunnel that is created by the requesting client. Tunnel types include the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol (L2TP), which are used by remote access clients and demand-dial routers running Windows XP; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000. You can use this condition to specify profile settings, such as authentication methods or encryption strengths for a specific type of tunneling technology.

**Configure Remote Access Policies for Wireless Access Clients**
For wireless access clients, create a wireless remote access policy, and then add the condition **NAS-Port-Type** with the values **Wireless-IEEE 802.11** and **Wireless-Other**.

For more information, see "Wireless access example" in Help and Support Center for Windows Server 2003.

**Configure Remote Access Policies for Authenticating Switch Access Clients**
For authenticating switch access clients, specify the use of the **Ethernet** port type for the **NAS-Port-Type** attribute. By using this port type, you create a separate remote access policy that contains connection parameters specifically designed for authenticating switch nodes.

# Configure a Quarantine Remote Access Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When deploying Network Access Quarantine Control, create a new policy called Quarantine Remote Access Policy, and then choose whether to quarantine clients with IP filters, whether to apply a session timer when clients connect to VPN or dial-up servers, or whether to apply both attributes to the connections. Specify the use of the **MS-Quarantine-IPFilter** attribute if you want to apply IP filters that restrict client access to the network. Specify the use of the **MS-Quarantine-Session-Timeout** attribute if you want to set a restriction on the amount of time the client network policy requirements script has to notify the access server that it has run successfully. When you create a Quarantine Remote Access Policy, position it as the first remote access policy so that it is evaluated before other policies.

Quarantine IP filters and session timers are for use only with the Windows Server 2003 Routing and Remote Access service.

For more information, see "Deploying Dial-up and VPN Remote Access Servers" and "Deploying Remote Access Clients Using Connection Manager" in this book.

# Securing Your Remote Access Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use IAS authentication, authorization, and accounting to secure your remote access solutions. You can also implement security precautions to protect your IAS server and IAS-related traffic. Figure 7.11 illustrates this process.

**Figure 7.11 Securing Your Remote Access Strategy**



# Select Authentication Protocols

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

- Windows Server 2003 IAS can perform authentication on behalf of any access server that is configured as a RADIUS client. You can use one of a number of protocols to authenticate dial-up, VPN, wireless, and authenticating switch users before allowing them access to the network.

Before you deploy IAS, determine which authentication protocols you will use to authenticate remote access clients. Use the most secure protocols that your network access servers and clients can support. If you need a high degree of security, you can configure IAS to accept only a few very secure authentication protocols. Alternatively, if your organization requires more flexibility, you can configure IAS to accept less secure authentication protocols when attempts to use more secure authentication protocols are unsuccessful.

Table 7.2 lists the authentication protocols that IAS supports and summarizes the conditions for which each protocol is used and the requirements for each protocol.

**Table 7.2 Authentication Protocols That IAS Supports**

| Protocol | Type of Authentication | Protocol Characteristics | Protocol Requirements |
|---|---|---|---|
| Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) | Certificate-based | Works for dial-up, VPN, wireless access, and authenticating switch access.<br><br>Provides authentication by means of a registry-based user certificate or a smart card.<br><br>Provides mutual authentication.<br><br>Generates cryptographic keys, which are needed for wireless LAN connections and for dial-up and VPN connections that use Microsoft Point-to-Point Encryption (MPPE).<br><br>Enables uninterrupted transfer between wireless access points (user does not need to re-enter credentials when moving between access points).<br><br>Enables unauthenticated access for visitors. | IAS must be a member of a Windows 2000 or Windows Server 2003 domain.<br><br>Both client and server must support this protocol.<br><br>The IAS server must have a certificate installed in the certificate store. The certificate must contain the Server Authentication purpose in EKU extensions. The certificate must meet all other certificate requirements.<br><br>The client computer or user certificate must contain the Client Authentication purpose in EKU extensions. The certificate must meet all other certificate requirements. The certificate can be installed in the client computer certificate store or on a smart card. |

| | | | |
|---|---|---|---|
| Protected Extensible Authentication Protocol (PEAP) | Certificate and password-based (depending upon the selected authentication method) | Currently for 802.1X wireless and authenticating switch clients only.<br><br>PEAP does not specify an authentication method, but provides a secure "wrapper" for other EAP authentication protocols, such as EAP-MS-CHAPv2, that operate within the outer TLS-encrypted channel provided by PEAP.<br><br>Enables uninterrupted transfer between wireless access points (user does not need to re-enter credentials when moving between access points).<br><br>Does not allow unauthenticated access for visitors. | Both client and server must support this protocol.<br><br>The IAS server must have a certificate installed in the certificate store. The certificate must contain the Server Authentication purpose in EKU extensions. The certificate must meet all other certificate requirements. |
| Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) | Password-based | Provides mutual authentication.<br><br>Generates cryptographic keys, which are needed for dial-up and VPN connections that use MPPE.<br><br>Enables you to change passwords.<br><br>Works for dial-up and VPN access. | Both client and server must support this protocol. |
| EAP-MSCHAPv2 | Password-based | Provides mutual authentication.<br><br>Generates cryptographic keys, which are needed for dial-up and VPN connections that use MPPE.<br><br>Enables you to change passwords. | IAS must be a member of a Windows 2000 or Windows Server 2003 domain.<br><br>Both client and server must support this protocol. |
| Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) | Password-based | Provides encrypted authentication for Microsoft® Windows® 98, Microsoft® Windows® Millennium Edition operating system, or Microsoft Windows NT 4.0 (with the latest dial-up networking upgrade).<br><br>Generates cryptographic keys, which are needed for dial-up and VPN connections that use Microsoft Point-to-Point Encryption (MPPE).<br><br>Enables you to change passwords. | Both client and server must support this protocol. |
| Extensible Authentication Protocol-Message Digest 5 (EAP-MD5) | Password-based | Provides less security than EAP-TLS.<br><br>Does not generate cryptographic keys.<br><br>Works for dial-up and authenticating switch–based access, but not wireless or PPTP-based VPN access. | Requires a reversibly encrypted password to be stored in the account database (local Security Accounts Manager (SAM) or domain).<br><br>IAS can be a member of a Windows NT 4.0 domain, a Windows 2000 domain, or a Windows Server 2003 domain.<br><br>Both client and server must support this protocol. |

| | | | |
|---|---|---|---|
| Challenge Handshake Authentication Protocol (CHAP) | Password-based | Provides encrypted authentication for a combination of different operating systems, such as Macintosh or UNIX operating systems.<br><br>Does not generate cryptographic keying material. | Requires a reversibly encrypted password to be stored in the account database (local SAM or domain).<br><br>Both client and server must support this protocol and reversibly encrypted passwords. |
| Password Authentication Protocol (PAP) | Password-based | Provides unencrypted authentication. Use only if clients do not support other protocols.<br><br>Does not generate cryptographic keying material. | Both client and server must support this protocol. |
| Unauthenticated Access | | Grants access when the remote access client does not supply authentication credentials.<br><br>Does not generate cryptographic keying material. | Both client and server must support this protocol. |

Before the IAS server can access Active Directory–based domains to authenticate user credentials and user access account properties, the IAS server must be registered in those domains.

# Integrate IAS with the Certificate Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Whether you need a certificate infrastructure for IAS depends on whether you are using EAP-TLS as your authentication protocol. If you are using EAP-TLS, you need a certificate infrastructure for your clients. Otherwise, you do not.

A certificate infrastructure consists of the following components:

- One or more certificate servers
- An IAS server with a certificate
- Clients with certificates

For more information about how to design a certificate infrastructure, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

When planning how to distribute certificates to clients, decide the following:

- **Whether to use a computer certificate or a user certificate.** Computer certificates are available in Windows Server 2003 and Windows 2000. Use them for servers and for computers on a LAN that do not often move. User certificates are new for Windows Server 2003. Use them for roaming users, such as wireless users.
- **How to install the certificate.** After the certification authority (CA) is configured, you can install a certificate in three ways. Table 7.3 shows each method and when to use it.

**Table 7.3 Selecting a Certificate Installation Method**

| Installation Method | When to Use |
|---|---|
| By using Group Policy to configure auto-enrollment of computer certificates to computers in a Windows Server 2003 domain. | Use this method if you have large numbers of domain member clients that you need to enroll. In this case, setting up a Group Policy takes less time than manually obtaining certificates.<br>This method has the advantage that after auto-enrollment is configured and enabled, all domain member computers receive computer certificates when Group Policy is refreshed next, whether the refresh is triggered manually with the **gpupdate** command, or by logging on to the domain.<br>If you use auto-enrollment for user certificates, any user with a valid password can obtain a certificate. This makes your certificate authentication the same as password-based authentication. |
| By using Certificate Manager to obtain a computer certificate. | Use this method if you have only a few computers, such as IAS servers. |
| By using Microsoft Internet Explorer and Web-based enrollment. | Use this method if the client is not a member of the domain.<br>Use this method or smart cards for user certificates. |

For specific information about how to perform these steps, see "Computer certificates for certificate-based authentication" in Help and Support Center for Windows Server 2003.

For more information about certificate enrollment methods and domain membership, see "Network access authentication and certificates" in Help and Support Center for Windows Server 2003.

# Install Computer Certificates for IAS Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are using PEAP-EAP-MS-CHAPv2 or EAP-TLS, you must install a computer certificate on your IAS servers. That certificate must be issued from a CA that can follow a certificate chain to a root CA that is trusted by the access clients. Likewise, the IAS server must trust the root CA of the CA that issued the user or computer certificate to the access client.

You can install multiple computer certificates on the IAS servers and configure separate remote access policies to use different computer certificates. However, you can select only a single certificate for all remote access policies that specify authentication by using EAP-TLS.

The server certificate must also contain the Server Authentication purpose in Enhanced Key Usage extensions, and meet other certificate requirements for PEAP and EAP authentication.

To install a certificate on the IAS server, you can use Group Policy and auto-enrollment, the CA Web enrollment tool provided with Certificate Services for Windows Server 2003, or you can request a certificate by using the Certificates snap-in.

For more information about certificate requirements for PEAP and EAP, see "Network access authentication and certificates" in Help and Support Center for Windows Server 2003.

# Install Computer Certificates for Access Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In addition to adding a certificate to your IAS servers, you must add a certificate to any computer that uses EAP-TLS or PEAP-EAP-TLS authentication. The certificate must be issued from a certification authority (CA) that can follow a certificate chain to a root CA that is trusted by the IAS server.

The computer certificate must also contain the Client Authentication purpose in Enhanced Key Usage extensions and meet other certificate requirements for PEAP and EAP authentication.

For more information about creating a certificate infrastructure, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

# Secure the IAS RADIUS Server and RADIUS Proxy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to secure your IAS server. Regardless of whether you configure your IAS server as a RADIUS server or a RADIUS proxy, you must apply a number of basic security precautions.

**Use strong shared secrets**
RADIUS supports a simple password called a secret. Configure strong shared secrets to prevent dictionary attacks, and change them frequently. RADIUS secrets are combined with a 16-byte random number and then passed through a one-way Message Digest 5 (MD5) hash to create a 16-byte encryption value. The 16-byte encryption value is stored with the password entered by the remote access user.

Include RADIUS secrets in your remote access design when you are mutually authenticating RADIUS computers and you encrypt the remote user password. It is best to specify RADIUS secrets that are at least 16 characters in length and that include uppercase letters, lowercase letters, numbers, and punctuation.

**Use the Message-Authenticator attribute**
Use the **Message-Authenticator** attribute (also known as a digital signature or the signature attribute) for connection requests that use the PAP, CHAP, MS-CHAP, and MS-CHAPv2 authentication protocols. This attribute ensures that an incoming RADIUS Access-Request message was sent from a RADIUS client configured with the correct shared secret. You must enable the use of the **Message-Authenticator** attribute on both the IAS server (as part of the configuration of the RADIUS client) and the RADIUS client (the network access server or RADIUS proxy). Ensure that the RADIUS client supports the **Message-Authenticator** attribute before you enable the attribute. The **Message-Authenticator** attribute is always used with EAP, regardless of whether it is enabled on the IAS server and access server.

**Configure your Internet firewall**
In the most common configuration, the Internet firewall is situated on your perimeter network between your secure network and the Internet. The perimeter network (also known as a *screened subnet*) is an IP network segment that contains resources (such as Web and VPN servers) that are available to Internet users. In this configuration, the IAS server is an intranet resource that is connected to the perimeter network.

If your IAS server is on a perimeter network, configure your Internet firewall to allow RADIUS messages to pass between your IAS server and RADIUS clients on the Internet. You might need to configure an additional firewall that is placed between your perimeter network and your intranet, which allows traffic to flow between the IAS server on the perimeter network and domain controllers on the intranet.

If your IAS server is on the perimeter network, it might use either of the following to contact a domain controller on the intranet:

- An interface on the perimeter network and an interface on the intranet (IP routing is not enabled).
- A single interface on the perimeter network. In this configuration, IAS communicates with intranet domain controllers through another firewall that connects the perimeter network to the intranet.

For more information about Internet firewalls, see "Deploying ISA Server" in this book.

**Enable remote access account lockout**
Enable remote access account lockout to protect against online dictionary attacks. Remote access account lockout disables network access for user accounts after a configured number of failed connection attempts has been reached.

Remote access account lockout can also be used to prevent a malicious user from intentionally locking out a domain account by attempting to make multiple dial-up or VPN connections with the wrong password. You can set the number of failed attempts for remote access account lockout to a number that is lower than the number of logon retries for domain account lockout. By doing this, remote access account lockout occurs before domain account lockout, which prevents the domain account from being intentionally locked out.

For more information about account lockout, see "Remote access account lockout" in Help and Support Center for Windows Server 2003.

# Protect Against Access Server Vulnerabilities

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some access servers expose the network to unwanted visitors who can gain access to your network. For example, intruders might connect to an authenticating switch accessible from a conference room network port, set up their own access server to connect to your network, and access your resources. To protect against this form of attack, configure mutual authentication by using PEAP-EAP-MS-CHAPv2 or EAP-TLS as the authentication method for network access connections. For more information about configuring mutual authentication, see "PEAP" in Help and Support Center for Windows Server 2003.

Authenticating RADIUS clients and RADIUS servers also protects against this type of network attack. You can use three methods to authenticate RADIUS clients and RADIUS servers.

**RADIUS shared secrets**
Include RADIUS shared secrets in your network access design. Specify RADIUS secrets that are at least 22 characters in length and consist of a random sequence of uppercase letters, lowercase letters, numbers, and punctuation.

**Secure RADIUS traffic with IPSec**
Securing RADIUS traffic with Internet Protocol security (IPSec) provides you with the ability to secure RADIUS servers against unwanted traffic by filtering on specific network adapters (allowing or blocking specific protocols) and enabling you to choose source IP addresses from which traffic is allowed. For organizational units, you can create IPSec policies, which are stored in Active Directory, or you can create local policies on RADIUS servers, and then apply the local policies to specific computers. If you create IPSec policies for an organizational unit, the policy is applied by using Group Policy.

You can enable IPSec between IAS proxies and IAS servers, or between RADIUS clients and IAS servers.

For more information about securing RADIUS traffic with IPSec, see "Securing RADIUS traffic with IPSec" in Help and Support Center for Windows Server 2003.

**VPN tunnels**
All RADIUS computers that require authorization support VPNs.

# Implementing Your IAS Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying IAS involves configuring IAS as a RADIUS server or proxy, optimizing your IAS configuration to best meet your needs, and configuring compatibility with third-party access servers. Figure 7.12 illustrates this process.

**Figure 7.12 Implementing Your IAS Solution**



# Deploy IAS as a RADIUS Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying IAS as a RADIUS server on your network involves configuring Active Directory, configuring the primary and backup IAS servers, and configuring the access servers. Figure 7.13 illustrates this process.

**Figure 7.13 Deploying IAS as a RADIUS Server**



# Configure User Accounts and Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure user accounts and groups, do the following:**

1. Ensure that all users that are making remote access connections have a corresponding user account.

2. Set the remote access permission on user accounts to **Grant access** or **Deny access** to manage network access by user. Or, to manage network access by group, set the remote access permission on user accounts to **Control access through Remote Access Policy**.

3. Organize remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies.

4. If you are using CHAP, enable support for reversibly encrypted passwords in the appropriate domains. You can configure reversibly encrypted passwords by using Group Policy. For more information, see "Enable reversibly encrypted passwords in a domain" in Help and Support Center for Windows Server 2003.

5. If you are using certificate-based authentication, configure the domain in which IAS server computers will be members for the auto-enrollment of certificates. With Windows Server 2003, you can auto-enroll both user and computer certificates.

For more information about configuring user accounts for IAS, see "Dial-up and VPN remote access" in Help and Support Center for Windows Server 2003.

For more information about reversibly encrypted passwords, see "Enable reversibly encrypted passwords in a domain" in Help and Support Center for Windows Server 2003.

# Configure the Primary IAS Server on a Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the primary IAS server on a domain controller, do the following:**

1. On the domain controller, install IAS by using **Add/Remove Windows Components**. For more information, see "Install IAS" in Help and Support Center for Windows Server 2003.

2. Configure the IAS server to read the properties of user accounts in the domain. For more information, see "Enable the IAS server to read user accounts in Active Directory" in Help and Support Center for Windows Server 2003.

3. If the IAS server authenticates connection attempts for user accounts in other domains, use the Active Directory Domains and Trusts snap-in to verify that these domains have a two-way trust with the domain in which the IAS server is a member. Next, configure the IAS server to read the properties of user accounts in other domains by adding the IAS server to the **RAS and IAS Servers** security group on all domain controllers with user account databases to be accessed by the IAS server. For more information, see "Enable the IAS server to read user accounts in Active Directory" in Help and Support Center for Windows Server 2003.

4. Enable file logging for accounting and authentication events. You can log session information to text files in either IAS format or database-compatible format, or you can log to a server running SQL Server 2000 or later. In addition, you can configure which information you want to log. For more information, see "Remote Access Logging" in Help and Support Center for Windows Server 2003.

5. If needed, configure additional User Datagram Protocol (UDP) ports for authentication and accounting messages that are sent by RADIUS clients. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

6. Add the access servers as RADIUS clients of the IAS server. Verify that you are configuring the correct name or IP address and shared secrets. Enable the use of the **Message-Authenticator** attribute, but only when it is also supported by the RADIUS client.

7. Create remote access policies that reflect your network access usage scenarios.

8. If you have created new remote access policies, either delete the default remote access policies or move them so that they are the last policies to be evaluated.

For more information about configuring IAS, see "Dial-up and VPN remote access" in Help and Support Center for Windows Server 2003.

# Configure the Secondary IAS Server on a Different Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the secondary IAS server on a different domain controller, do the following:**

1. On the other domain controller in the same domain, install IAS by using **Add/Remove Windows Components**. For more information, see "Install IAS" in Help and Support Center for Windows Server 2003.

2. Configure the secondary IAS server to read the properties of user accounts in the domain by adding the IAS server to the **RAS and IAS Servers** security group. For more information, see "Enable the IAS server to read user accounts in Active Directory" in Help and Support Center for Windows Server 2003.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server is a member. Next, configure the secondary IAS server to read the properties of user accounts in other domains by adding the IAS server to the **RAS and IAS Servers** security group in those domains. For more information, see " Enable the IAS server to read user accounts in Active Directory" in Help and Support Center for Windows Server 2003.

4. Copy the configuration of the primary IAS server to the secondary IAS server.

For more information about copying IAS configuration, see "Copy the IAS configuration to another server" in Help and Support Center for Windows Server 2003.

# Configure RADIUS Clients for Use with IAS Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure each RADIUS client to use the primary and secondary IAS servers for authentication, authorization, and accounting of remote access connections, do the following:

- If the RADIUS client is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, use the Routing and Remote Access snap-in to configure the RADIUS client to use the primary and secondary IAS servers as RADIUS servers. For more information about configuring the RADIUS client, see "Use RADIUS authentication" and "Use RADIUS accounting" in Help and Support Center for Windows Server 2003.

- If the RADIUS client is a computer running Windows NT 4.0 and the Routing and Remote Access Service (RRAS), see Windows NT 4.0 Help for information about how to configure RRAS to use the primary and secondary IAS servers as RADIUS servers for RADIUS authentication.

- If the RADIUS client is a third-party network access server, see the documentation for the NAS to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS servers).

# Deploy IAS as a RADIUS Proxy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying IAS as a RADIUS proxy on your network involves configuring the primary and secondary IAS proxies, configuring the Internet firewalls, and configuring for compatibility with third-party access servers. Figure 7.14 illustrates this process.

**Figure 7.14 Deploying IAS as a RADIUS Proxy**



# Configure the Primary IAS Proxy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the primary IAS proxy in the perimeter network, do the following:**

1. On a computer running Windows Server 2003, Standard Edition or Windows Server 2003, Enterprise Edition in the perimeter network, install IAS by using **Add/Remove Windows Components**. For more information, see "Install IAS" in Help and Support Center for Windows Server 2003. The computer on which IAS is installed does not need to be dedicated to forwarding RADIUS messages. You can install IAS on a computer running other services, such as a DHCP server, a file server, or a DNS server.

2. If needed, configure additional UDP ports for RADIUS messages that are sent by the RADIUS proxies. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

3. Add the RADIUS proxies as RADIUS clients of the IAS server. Verify that you are configuring the correct name or IP address and shared secrets.

4. Create a remote RADIUS server group that contains the IAS servers in your organization.

5. Create a connection request policy that forwards RADIUS request messages based on the realm name of your organization. For more information about realm names, see "Realm names" in Help and Support Center for Windows Server 2003.

6. Use the New Connection Request Policy Wizard to create a connection request policy that forwards connection requests to a remote RADIUS server group and where the realm name matches the realm name of the user accounts in your organization. Clear the check box that removes the realm name for authentication. In the New Connection Request Policy Wizard, use the New Remote RADIUS Server Group Wizard to create a remote RADIUS server group with members that include the two IAS servers within your intranet.

7. Delete the default connection request policy named **Use Windows authentication for all users**.

For more information about configuring IAS proxies in the perimeter network, see "Outsourced dial and a proxy in the perimeter network" in Help and Support Center for Windows Server 2003.

# Configure the Secondary IAS Proxy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the secondary IAS proxy in the perimeter network, do the following:**

1. On another computer running Windows Server 2003, Standard Edition or Windows Server 2003, Enterprise Edition in the perimeter network, install IAS by using **Add/Remove Windows Components**. For more information, see "Install IAS" in Help and Support Center for Windows Server 2003.

2. Using Netsh, copy the configuration of the primary IAS proxy to the secondary IAS proxy in the perimeter network.

For more information about copying IAS configuration, see "Copy the IAS configuration to another server" in Help and Support Center for Windows Server 2003.

# Configure the Intranet and Internet Firewalls

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To support RADIUS traffic, you must take two steps. First, configure the Internet firewall to allow RADIUS traffic between the IAS proxies on the perimeter network and the RADIUS clients on the Internet. Then configure the intranet firewall to allow RADIUS traffic between the IAS proxies on the perimeter network and the IAS servers on the intranet.

**Filters on the Internet Interface**

Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1812 (0x714).

  This filter allows RADIUS authentication traffic from Internet-based RADIUS clients to the IAS server. This is the default UDP port used by IAS as defined in RFC 2856. If you are using a different port, substitute that port number for 1812, as used in this example.

- Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1813 (0x715).

  This filter allows RADIUS accounting traffic from Internet-based RADIUS clients to the IAS server. This is the default UDP port used by IAS as defined in RFC 2857. If you are using a different port, substitute that port number for 1813, as used in this example.

Configure the following output filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the IAS server's perimeter network interface and UDP source port of 1812 (0x714).

  This filter allows RADIUS authentication traffic from the IAS server to Internet-based RADIUS clients. This is the default UDP port used by IAS as defined in RFC 2856. If you are using a different port, substitute that port number for 1812, as used in this example.

- Source IP address of the IAS server's perimeter network interface and UDP source port of 1813 (0x715).

  This filter allows RADIUS accounting traffic from the IAS server to Internet-based RADIUS clients. This is the default UDP port used by IAS as defined in RFC 2857. If you are using a different port, substitute that port number for 1813, as used in this example.

**Filters on the Perimeter Network Interface**

Configure the following input filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the IAS server's perimeter network interface and UDP source port of 1812 (0x714).

  This filter allows RADIUS authentication traffic from the IAS server to Internet-based RADIUS clients. This is the default UDP port used by IAS as defined in RFC 2856. If you are using a different port, substitute that port number for 1812, as used in this example.

- Source IP address of the IAS server's perimeter network interface and UDP source port of 1813 (0x715).

  This filter allows RADIUS accounting traffic from the IAS server to Internet-based RADIUS clients. This is the default UDP port used by IAS as defined in RFC 2857. If you are using a different port, substitute that port number for 1813, as used in this example.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1812 (0x714).

This filter allows RADIUS authentication traffic from Internet-based RADIUS clients to the IAS server. This is the default UDP port used by IAS as defined in RFC 2856. If you are using a different port, substitute that port number for 1812, as used in this example.

- Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1813 (0x715).

  This filter allows RADIUS accounting traffic from Internet-based RADIUS clients to the IAS server. This is the default UDP port used by IAS as defined in RFC 2857. If you are using a different port, substitute that port number for 1813, as used in this example.

For added security, if you know the IP address of each RADIUS client sending the packets through the firewall, you can configure more specific filters for traffic between the IP address of the RADIUS client and the IP address of the IAS server on the perimeter network.

For more information about configuring packet filters, see "Manage Packet Filters" and "Apply packet filters for business partner extranet" in Help and Support Center for Windows Server 2003.

# Configure RADIUS Clients for Use with IAS Proxies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure each RADIUS client to use the primary and secondary IAS proxies, do the following:

- If the RADIUS client is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, use the Routing and Remote Access snap-in to configure the RADIUS client to use the primary and secondary IAS servers.

- If the RADIUS client is a computer running Windows NT 4.0 and the Routing and Remote Access Service (RRAS), see Windows NT 4.0 Help for information about how to configure RRAS to use the primary and secondary IAS proxies as RADIUS servers.

- If the RADIUS client is a third-party network access server, see the documentation for the NAS to determine how to configure it as a RADIUS client with two RADIUS servers.

For more information about configuring RADIUS accounting and authentication, see "Use RADIUS accounting" and "Use RADIUS authentication" in Help and Support Center for Windows Server 2003.

# Configuring IAS for Compatibility with Third-Party Access Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are using Windows Server 2003 IAS with a third-party access server, you might need to configure the IAS server to prevent it from sending the class attribute. Specifically, if the IAS server sends the class attribute, some network access servers require the use of certain vendor-specific attributes to be included in the class attribute.

Alternatively, you can configure vendor-specific attributes and custom attributes to enable compatibility with a third-party access server.

For more information about how to configure the class attribute, see "Add RADIUS attributes to a remote access policy" in Help and Support Center for Windows Server 2003. For more information about configuring IAS for compatibility with a third-party access server, see "IAS as a RADIUS server design considerations" in Help and Support Center for Windows Server 2003.

# Configure Vendor-Specific Attributes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

IAS allows administrators to associate specific RADIUS attributes with each remote access policy. These attributes are described in RFC 2865. In addition to the standard RADIUS attributes, some access server manufacturers use vendor-specific attributes (VSAs) to provide functionality that is not supported in standard attributes.

If your access server manufacturer uses VSAs, you might need to associate those VSAs with remote access protocols so that your access server processes authentication requests properly. IAS enables you to create or edit VSAs to take advantage of proprietary functionality supported by some access server vendors.

To provide access server–specific support, you can add vendor-specific attributes to remote access policies. IAS includes VSAs from some vendors in its dictionary. You cannot add other VSAs to the dictionary.

**To configure IAS to support vendor-specific attributes, complete the following steps:**

1. Determine which vendor-specific attributes the access server requires. The third-party access server manufacturer documentation provides information about the individual attributes that must be specified.

2. Determine which attributes are standard RADIUS attributes and which require the addition of vendor-specific RADIUS attributes (attribute type 26). For more information, see "Interpreting IAS-formatted log files" in Help and Support Center for Windows Server 2003.

3. Specify the attribute values for the VSAs in the IAS remote access policy. When you add VSAs to the remote access policy profile to support proprietary functionality for your access point, you must determine whether the VSA conforms to the format recommended in RFC 2865. If it does, you must specify a network access vendor by either name or vendor code, a vendor-assigned attribute number, the attribute format (that is, the type of data, such as string or hexadecimal), and the attribute value. If it does not, you must specify a network access vendor by either name or vendor code and a hexadecimal attribute value that represents the attribute data. For more information, see "Vendor-specific attribute overview" in Help and Support Center for Windows Server 2003.

# Configure Custom Attributes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IAS Software Development Kit (SDK), part of the Windows Server 2003 SDK, allows you to configure your IAS server to send custom attributes to the access server in addition to standard attributes. This enables you to build your own extension to the IAS snap-in.

For more information about configuring custom attributes, see the Software Development Kit (SDK) information in the MSDN Library link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Additional Resources for Remote Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about Internet Authentication Service.
- "Deploying Remote Access Clients Using Connection Manager" in this book.
- "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit for more information about how to design a certificate infrastructure.
- "Deploying Dial-up and VPN Remote Access Servers" in this book.
- "Deploying a Wireless LAN" in this book for information about deploying wireless access clients.
- RFC 2865: *Remote Authentication Dial In User Service (RADIUS)*.
- RFC 2866: *RADIUS Accounting*.
- RFC 2869: *RADIUS Extensions*.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Internet Authentication Service" in Help and Support Center for Windows Server 2003.
- "Remote Access Policies" in Help and Support Center for Windows Server 2003.
- "IAS as a RADIUS server" in Help and Support Center for Windows Server 2003.
- "Deploying IAS as a RADIUS Proxy" in Help and Support Center for Windows Server 2003.
- "Compulsory tunnels" in Help and Support Center for Windows Server 2003 for information about the RADIUS attributes used with compulsory tunneling.
- "Computer certificates for certificate-based authentication" in Help and Support Center for Windows Server 2003.
- "Dial-up and VPN remote access" in Help and Support Center for Windows Server 2003 for more information about configuring user accounts for IAS.
- "Copy the IAS configuration to another server" in Help and Support Center for Windows Server 2003 for more information about copying IAS configuration.
- "Outsourced dial and a proxy in the perimeter network" in Help and Support Center for Windows Server 2003 for more information about configuring IAS proxies in the perimeter network.
- "Add RADIUS attributes to a remote access policy" in Help and Support Center for Windows Server 2003 for more information about how to configure the class attribute.

- "Manage Packet Filters" in Help and Support Center for Windows Server 2003 for more information about configuring packet filters.
- "Use RADIUS accounting" and "Use RADIUS authentication" in Help and Support Center for Windows Server 2003 for more information about configuring RADIUS accounting and authentication.
- "Managing multiple IAS servers" in Help and Support Center for Windows Server 2003 for more information about synchronizing the configuration of multiple IAS servers.
- "Configure accounting" in Help and Support Center for Windows Server 2003 for more information about using separate IAS servers for authentication and accounting.
- "Configure authentication" in Help and Support Center for Windows Server 2003.
- "Configure encryption" in Help and Support Center for Windows Server 2003.
- "PEAP" in Help and Support Center for Windows Server 2003.
- "Network access authentication and certificates" in Help and Support Center for Windows Server 2003 for more information about certificate enrollment methods and domain membership

# Deploying Dial-up and VPN Remote Access Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the remote access and security technologies in the Microsoft® Windows® Server 2003 operating systems to provide remote users with secure and reliable access to your network resources. By using the Routing and Remote Access service in Windows Server 2003, you can design and deploy a dial-up solution or you can take advantage of the Internet by deploying a virtual private network (VPN) solution.

**In This Chapter**


Overview of Deploying Dial-up and VPN Remote Access Servers

Process for Deploying Dial-up and VPN Remote Access Servers

Choosing Dial-up or VPN

Designing a Remote Access Server Solution

Deploying a VPN Remote Access Server Solution

Deploying a Dial-up Remote Access Server Solution

Additional Resources for Deploying Dial-up and VPN Remote Access Servers


**Related Information**

- For information about designing and deploying a public key infrastructure (PKI), see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.
- For more information about deploying smart cards, see "Planning a Smart Card Deployment" in *Designing and Deploying Directory and Security Services* of this kit.
- For information about designing and deploying Internet Authentication Service (IAS), see "Deploying Internet Authentication Service (IAS)" in this book.

# Overview of Deploying Dial-up and VPN Remote Access Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To support users who require access to your network from remote locations, you can deploy a dial-up network, a VPN, or a combination of both. *Dial-up networking* enables remote users to dial in directly to a remote access server on your network using a phone line. A *virtual private network (VPN)* enables remote users who are connected to the Internet to establish a connection to a VPN server on your network.

In deciding which solution will best serve your organization, consider the relative cost-effectiveness of each solution and how well it meets your organization's requirements for security and availability. Also consider the network infrastructure of the intranet needed to support your remote access server design. Without proper design of the supporting infrastructure, remote access clients cannot obtain IP addresses and resolve intranet names, and packets cannot be forwarded between remote access clients and intranet resources.

Use the process described in this chapter to design and deploy a new remote access solution or to reexamine and improve your existing infrastructure. If you already have a dial-up or VPN infrastructure, you might benefit from replacing existing components because of anticipated obsolescence or failure of components, scalability limitations, or increased security requirements.

Before you begin work on your remote access server design, your organization should have deployed the following supporting technologies.

- Active Directory® directory service to store and manage information about network resources.
- A Public key infrastructure (PKI) to enable the use of certificate-based authentication.
- Internet Authentication Service (IAS) to provide authentication and authorization for dial-up and VPN network access. As a RADIUS server, Windows Server 2003 IAS performs authentication and authorization on behalf of any remote access server configured as a RADIUS client.

All editions of Windows Server 2003 support VPN connections. However, some limitations apply to the Microsoft® Windows® Server 2003, Web Edition, and Windows® Server 2003, Standard Edition, operating systems. On computers running either of these operating systems, you can create as many as 1,000 connections, using Point-to-Point Tunneling Protocol (PPTP) ports or Layer Two Tunneling Protocol (L2TP) ports. Windows Server 2003, Standard Edition, can accept 1,000 concurrent VPN connections; however, Windows Server 2003, Web Edition, accepts only one VPN connection at a time. For more information about features included in Windows Server 2003, Web Edition, see "Overview of Windows Server 2003, Web Edition" in Help and Support for Windows Server 2003.

# Process for Deploying Dial-up and VPN Remote Access Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The key to providing a secure and reliable remote access solution that meets your organization's needs is to carefully plan and test your remote access design after deciding whether to deploy a dial-up network, a VPN, or a combination of both. Deploying a VPN remote access server involves different tasks than are required to deploy a dial-up remote access server; however, often your remote access solution will involve elements of both. Figure 8.1 shows the process for designing and deploying dial-up and VPN remote access servers.

**Figure 8.1 Deploying Dial-up and VPN Remote Access Servers**



# Choosing Dial-up or VPN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With a secure network architecture based on Windows Server 2003 in place, the first step in designing a remote access server solution is deciding whether to provide network access to remote clients by using dial-up networking, a VPN solution, or a combination of both. Figure 8.2 shows the placement of this design decision in the process for designing and deploying dial-up and VPN remote access servers.

**Figure 8.2 Choosing Dial-up or VPN**




Each method for providing remote access has advantages and disadvantages that you must weigh based on the needs of your organization. A dial-up networking solution provides a secure data path over a circuit-switched connection, and it provides the convenience of direct dial-up connectivity to your network for mobile users. In contrast, a VPN solution, by using the Internet as a connection medium, saves the cost of long-distance phone service and hardware costs. To mitigate the public nature of the Internet, VPNs use a variety of security technologies, including tunneling, encryption, and authentication.

**Using Dial-up Networking for Remote Access**

In a dial-up networking solution, remote users call in to a remote access server on your network. Dial-up lines are inherently more private than a solution that uses a public network such as the Internet. However, with dial-up networking, your organization faces a large initial investment and continuing expenses throughout the life cycle of the solution. These expenses include:

- **Hardware purchase and installation.** Dial-up networking requires an initial investment in modems or other communication hardware, server hardware, and phone line installation.

- **Monthly phone costs.** Each phone line that is used for remote access increases the cost of dial-up networking. If you use toll-free numbers or the callback feature to defray long distance charges for your users, these costs can be substantial. Most businesses can arrange a bulk rate for long distance, which is preferable to reimbursing users individually at their more expensive residential rates.

- **Ongoing support.** The number of remote access users and the complexity of your remote access design significantly affect the ongoing support costs for dial-up networking. Support costs include network support engineers, testing equipment, training, and help desk personnel to support and manage the deployment. These costs represent the largest portion of your organization's investment.

Figure 8.3 shows an example of a simple dial-up remote access networking design.

**Figure 8.3 Dial-up Remote Access Design**




**Providing Remote Access over a VPN**

In a VPN solution for remote access, users connect to your corporate network over the Internet. VPNs use a combination of tunneling, authentication, and encryption technologies to create secure connections. To ensure the highest level of security for a VPN deployment, use Layer Two Tunneling Protocol with Internet Protocol security (L2TP/IPSec).

Many organizations with extensive remote access requirements implement a VPN solution. VPNs reduce remote access expenses by using the existing Internet infrastructure. You can use a VPN to partially or entirely replace your centralized, in-house, dial-up remote access infrastructure and legacy services.

VPNs offer two primary benefits:

- **Reduced costs.** Using the Internet as a connection medium saves long-distance phone expenses and requires less hardware than a dial-up networking solution does.

- **Sufficient security.** Authentication prevents unauthorized users from connecting. Strong encryption methods make it extremely difficult for an attacker to interpret the data sent across a VPN connection.

Figure 8.4 shows an example of a simple VPN remote access networking design.

**Figure 8.4 VPN Remote Access Design**




**Note**

- Regardless of the approach that you choose, you can increase manageability of your remote access server solution by using IAS to centralize VPN or dial-up networking authentication, authorization, and accounting. For the Microsoft® Windows® 2000 Server family, IAS is a RADIUS server; for the Windows Server 2003 family, IAS is a RADIUS server and proxy. For information about designing and deploying IAS, see "Deploying Internet Authentication Service (IAS)" in this book.

# Designing a Remote Access Server Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After choosing the remote access to deploy, begin the design process. The design for a remote access server solution encompasses hardware requirements, outsourcing options, placement of remote access servers on the network, routing for remote access clients, and a security plan. Before deploying your remote access server solution, optimize and test your design. Figure 8.5 shows the process for designing a remote access server solution.

**Figure 8.5 Designing a Remote Access Server Solution**



# Determining Hardware Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in designing a remote access solution is to determine hardware requirements. Determine the capacity required for your dial-up or VPN servers, as well as additional hardware required, such as modem banks for a dial-up solution. When selecting server hardware, consider central processing unit (CPU), RAM, and network hardware requirements. For optimal performance, run lab tests on each model of equipment to be used in your production environment.

# Determining Hardware Requirements for Dial-up Networking

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following guidelines when determining hardware requirements for your dial-up networking design:

- A dial-up remote access server must have a modem or a multiport adapter, and it must have access to an analog telephone line. Based on the maximum number of dial-up remote access clients that will dial in at one time, install modem bank equipment and phone lines that meet your organization's needs.

- Each server-side modem requires a serial port on the remote access server. For multiple modems (modem banks), use a multiport serial adapter or a high-density combination card. A multiport serial adapter can connect a large number of analog modems or Integrated Services Digital Network (ISDN) modems to one remote access server. A high-density combination card combines multiple modems and serial adapters in one device.

- For best performance, consider using intelligent communications adapters (multiport serial boards) to offload processing from the remote access server.

# Determining Hardware Requirements for VPN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following guidelines when determining network hardware requirements for your VPN design:

- For interfaces on the public network, use network adapters capable of IPSec hardware offload.
- Assuming that you have a 10/100 Ethernet infrastructure, set all devices to 100 Mbps Full Duplex.
- Connect interfaces on the private network directly to a high-capacity switch that also connects the data servers and routers that remote access clients will access frequently.

**CPU Requirements**

Use the following guidelines when determining CPU requirements for your VPN design:

- Processing inbound and outbound packets requires CPU cycles. By increasing the available processing power, you can increase throughput.
- Doubling the speed of a single processor is more effective than doubling the number of processors.
- In the case of multiprocessor platforms, binding one CPU to each network adapter can increase the efficiency of interrupt handling, freeing cycles and shrinking the performance gap between the use of a large number of less powerful CPUs and a few faster, more expensive CPUs.

**RAM Requirements**

Use the following guidelines when determining the RAM needed for VPN servers:

- Each active connection consumes a small block of nonpageable memory (approximately 40 KB). If you do not need to handle more than 1,000 concurrent calls from remote access users, 512 MB of RAM is adequate.
- If you require the capacity to handle more than 1,000 concurrent calls, for every 1,000 concurrent calls provide an extra 128 MB of RAM over recommended RAM capacity for the server, plus a base of 128 MB more for remote access and related services.

  For example, for a dedicated remote access server that will support as many as 2,000 simultaneous VPN calls, if the recommended RAM capacity for Windows Server 2003 is 256 MB, provide 768 MB of RAM:

  256 MB + (128 MB * 2) + (128 MB * 2)

**Performing Capacity Planning**

The two greatest potential performance constraints in your remote access server solution are the number of simultaneous connections and the overall data throughput.

The number of simultaneous connections that a VPN server can support is determined by the available nonpaged pool memory as well as other factors, such as the use of data compression. With compression, each connection uses more nonpaged pool memory and requires more processing. Turning off compression can improve performance.

The processing power that is available to a VPN server determines the server's data throughput capacity. Tunneling protocols also have an impact on data throughput. PPTP connections require less processing power than L2TP/IPSec connections do; however, L2TP/IPSec connections are the most secure. You can mitigate the impact of L2TP/IPSec on processing power by using IPSec hardware offload.

# Considering Outsourcing Options

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Outsourcing part or all of a remote access solution through a wholesale contract with an Internet service provider (ISP) can provide some organizations significant cost savings, minimizing both setup and operations costs. You can purchase a wholesale contract with an ISP that provides a guaranteed level of service for some or all components of your remote access solution. An ISP can provide access to an extensive network of connections across a wide geographical area for a fixed cost. Many ISPs provide a guaranteed level of service that rivals those of dial-up wide area network (WAN) connections.

If you are deploying a dial-up networking solution, consider outsourcing your deployment to avoid the long-distance charges that you are likely to accrue.

If you are deploying a VPN solution, an ISP can provide many of the components required to support VPN access, including:

- Network access servers (NASs) for access from various geographical access points.

  If you outsource support, ensure that the NASs that the ISP provides meet your access requirements. VPN users dial first into a NAS. Unless your outsourcing agreement specifies the support required from the ISP's NASs, including connection speeds, device support, and levels of service, the NASs can be the weak link in your VPN solution.

- RADIUS proxy servers for routing of access requests to the enterprise.

  You can either contract with the ISP to provide authentication services or have the ISP use a RADIUS proxy and send authentication requests to a RADIUS server that you manage.

- Phone book support for delivering the latest access numbers either to the enterprise or directly to the client.

# Placing Remote Access Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In deciding where to place remote access servers on your network, consider firewall placement and the placement of other network resources. Place remote access servers close to the network resources that remote access clients need. These resources might include a CA, a RADIUS server, a domain controller, or file and application servers.

In a dial-up remote access design, servers usually are placed behind the firewall. Because a VPN design involves Internet connectivity, server placement relative to the firewall is a greater issue.

If you are designing a VPN remote access solution, choose between two options for server placement, each with different design requirements:

- **VPN server behind the firewall.** The firewall is attached to the Internet, with the VPN server between the firewall and the intranet. This is the placement used in a perimeter network configuration, in which one firewall is positioned between the VPN server and the intranet, with another between the VPN server and the Internet.

- **VPN server in front of the firewall.** The VPN server is connected to the Internet, with the firewall between the VPN server and the intranet.

**VPN Server Behind the Firewall**

The most common configuration for a VPN remote access design is to locate the VPN server behind a firewall. In this configuration, the firewall is connected to the Internet, and the VPN server is an intranet resource that is connected to the perimeter network. The VPN server has an interface on both the perimeter network and the intranet. The Internet firewall (the firewall between the Internet and the VPN server) filters all traffic from Internet clients. The intranet firewall (the firewall between the VPN server and the intranet) filters intranet traffic from VPN clients.

Placing a VPN server behind the firewall requires the following configuration:

- Configure the Internet interface on the firewall with inbound and outbound filters that allow traffic to the VPN server. You can specify additional filters to allow traffic to the Web servers, File Transfer Protocol (FTP) servers, and other types of servers on the perimeter network.

- For an added layer of security, configure the perimeter network interface on the VPN server with PPTP or L2TP/IPSec packet filters.

**VPN Server in Front of the Firewall**

Another option is to place the VPN server in front of the firewall, directly connected to the Internet. For inbound traffic, the VPN server decrypts the tunneled data and forwards it to the firewall. The firewall acts as a filter for intranet traffic, and it can prevent access to specific resources, scan data for viruses, perform intrusion detection, and carry out other functions.

To place a VPN server in front of the firewall, you must configure inbound and outbound filters on the VPN server to allow only VPN traffic to and from the IP address of the VPN server's Internet interface.

### Note

- To enable access to services running on the VPN server, make sure that the network BIOS (NetBIOS) and DNS names of the VPN server's Internet interface are not registered in the intranet namespaces. This is the default behavior for Windows Server 2003 VPN servers.

# Determining Routing for VPN Remote Access Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because VPN remote access clients typically do not support routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), the clients often have very simple routing tables. Your remote access design must provide a method for routing packets from the remote access clients to the intranet and the Internet — in some cases simultaneously, depending on the needs of the remote user.

For example, a computer on a small local area network (LAN) has a route to its own subnet on the LAN, and it might have a default route to a gateway router on the LAN for all other traffic. Without routing protocol support, remote access clients cannot automatically determine the best route to a destination.

**Choosing Between Routing Approaches**

The preferred method for directing packets to a remote network is to create a default route on the remote access client that directs packets to the remote network (the default configuration for VPN remote access clients). Any packet that is not intended for the neighboring LAN segment is sent to the remote network. When a connection is made, the remote access client by default adds a default route to its routing table and increases the metric of the existing default route to ensure that the newest default route is used. The newest default route points to the new connection, which ensures that any packets that are not addressed to the local LAN segment are sent to the remote network.

Under this configuration, when a VPN client connects and creates a new default route, Internet sites that have been accessible are no longer accessible. This poses no problem for VPN remote access users who only require access to the organization's network. However, it is not acceptable for remote users who need access to the Internet while they are connected to the organization's network.

Based on whether or not the default route is added, the client has broad access either to Internet locations or to locations on the intranet, but not to both:

- If the default gateway on the remote network is not being used, Internet locations are reachable, but only intranet locations matching the address class of the assigned IP address can be reached.

- If the default gateway on the remote network is being used, all intranet locations are reachable, but only the IP address of the VPN server and locations available through other routes can be reached on the Internet.

For most VPN clients with an Internet connection, this does not present a problem, because the client is typically engaged in either intranet communication or Internet communication, but not both.

To work around this problem, instead of having the client create a new default route when a connection is made, you can configure the client's routing table with specific routes that direct packets to the organization's network over the VPN connection. While connected to the intranet, the client can obtain Internet access using the existing default route over the connection to the ISP. This configuration is known as *split tunneling*.

**Choosing Among Split Tunneling Options**

The VPN client can obtain the routes needed for split tunneling in several ways:

- If the VPN client has a configured connection without a default route, the client adds a route that it infers from the class of the IP address assigned to it for the current connection. For a simple target network, such as a small office, this one route is sufficient to allow packets to be routed to the target network. However, for a complex network, you need to configure multiple routes to successfully direct packets to the remote network.

- A client running the Microsoft® Windows® XP or Windows Server 2003 operating systems uses a DHCPINFORM message after the connection to obtain additional information about the connection, such as a DNS name or a set of routes for the target network. This additional information is only available if the DHCP server has been configured to provide the DHCP Classless Static Routes option, and if the remote access server has been configured to relay the DHCPINFORM message to the DHCP server.

- If the remote access client is managed using the Connection Manager component of Windows Server 2003, the network administrator can prepare a routing table for the remote access client and associate a post-connect action with the

managed connection to update the client's routing table when a connection is made. For more information about using Connection Manager, see "Deploying Remote Access Clients Using Connection Manager" in this book.

- If none of the above approaches meets your needs, the end user or network administrator can write a program or batch file that updates the routing table on the client with the necessary routes to the remote network.

**Security considerations for split tunneling**

Some network administrators consider split tunneling to be a needless security risk. Their concern is that an attacker might be able to compromise the network by enabling traffic to be routed between the Internet and the network using the remote access client. To prevent this, you can add packet filters to the remote access policy profile settings for the VPN connection to allow only inbound traffic that originates from remote access clients. The default remote access policy named **Connections to Microsoft Routing and Remote Access server** has these packet filters already configured.

**Obtaining IP Addresses for Remote Access Clients**

Use either of the following methods to obtain IP addresses to assign to remote access clients:

- Let Routing and Remote Access assign IP addresses obtained from a DHCP server to remote access clients.

- Configure a static pool of IP addresses to assign to remote access clients. You can specify either an on-subnet or off-subnet range of IP addresses for the static pool.

# Planning Security for a VPN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because a dial-up networking solution provides a secure data path over a circuit-switched connection, security is not a critical issue in your design for a dial-up remote access solution. In contrast, a VPN remote access solution routes data over a packet-switched connection that does not intrinsically provide the same level of security. Therefore, security is an important part of your VPN remote access server design.

The security of a VPN is based on the tunneling and authentication protocols that you use and the level of encryption that you apply to VPN connections. For the highest level of security, use a remote access VPN based on L2TP/IPSec with certificate-based IPSec authentication and Triple-DES for encryption. If you decide to use a PPTP-based VPN solution to reduce costs and improve manageability and interoperability, use Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) as the authentication protocol.

In designing security for your VPN remote access server solution, perform the following tasks:

- Select a VPN protocol.
- Select authentication protocols.
- Select the scope and level of encryption.
- If needed, plan a certificate infrastructure to support client authentication for remote access.
- Optionally, plan for Network Access Quarantine Control.
- Optionally, enhance security by using remote access account lockout.

**Note**

- You can increase the security and manageability of your remote access server solution by using IAS to centralize VPN or dial-up networking authentication, authorization, and accounting. In operating systems in the Windows 2000 Server family, IAS is an implementation of a RADIUS server; in Windows Server 2003, IAS is an implementation of a RADIUS server and proxy. For information about designing and deploying Internet Authentication Service (IAS), see "Deploying Internet Authentication Service (IAS)" in this book.

# Selecting a VPN Protocol

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Tunneling and authentication protocols, and the encryption levels applied to VPN connections, determine VPN security. L2TP/IPSec provides the highest level of security. For a VPN design, determine which VPN protocol best meets your requirements. Windows Server 2003 supports two VPN protocols: Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol with Internet Protocol security (L2TP/IPSec).

**PPTP**

PPTP uses Point-to-Point Protocol (PPP) user authentication methods and Microsoft Point-to-Point Encryption (MPPE) to encrypt IP traffic. When used with MS-CHAP v2 for password-based authentication and strong passwords, PPTP is a secure VPN technology. For stronger authentication for PPTP connections, you can implement a PKI using smart cards or certificates and Extensible Authentication Protocol — Transport Level Security (EAP-TLS).

PPTP is widely supported and easily deployed, and it works with most network address translators (NATs).

**L2TP/IPSec**

The more secure of the two VPN protocols, L2TP/IPSec uses PPP user authentication methods and IPSec encryption to encrypt IP traffic. This combination uses certificate-based computer identity authentication to create IPSec security associations in addition to PPP-based user authentication. L2TP/IPSec provides data integrity, data origin authentication, data confidentiality, and replay protection for each packet.

Support for L2TP/IPSec is provided with Windows Server 2003, as well as with Windows 2000 and Windows XP. To use L2TP/IPSec with the Microsoft® Windows® 98, Windows® Millennium Edition (Windows Me), or Windows NT® Workstation 4.0 operating system, download and install Microsoft L2TP/IPSec VPN Client (Mls2tp.exe). For information about Mls2tp.exe, see the Microsoft L2TP/IPSec VPN Client link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Table 8.1 summarizes the advantages and constraints associated with the use of the PPTP and L2TP/IPSec protocols.

**Table 8.1 Advantages and Constraints of the PPTP and L2TP/IPSec VPN Protocols**

| Factor | PPTP Advantages and Constraints | L2TP/IPSec Advantages and Constraints |
|---|---|---|
| Client operating systems supported | Supported on clients running Windows 2000, Windows XP, Windows Server 2003, Windows NT Workstation 4.0, Windows Me, or Windows 98. | Natively supported on clients running Windows 2000, Windows XP, or Windows Server 2003.<br><br>With Mls2tp.exe installed, supported on clients running Windows 98, Windows Me, or Windows NT Workstation 4.0. |
| Certificate support | For EAP-TLS authentication to issue computer certificates to the authenticating server and user certificates to all VPN clients or to issue smart cards to all users, PPTP requires a certificate infrastructure. | To issue computer certificates to the VPN server and all VPN clients, L2TP/IPSec requires a certificate infrastructure or a preshared key (PSK). |
| Security | Provides data confidentiality. (Captured packets cannot be interpreted without the encryption key.)<br><br>Does not provide data integrity (proof that the data was not modified in transit) or data origin authentication (proof that the data was sent by the authorized user).<br><br>To increase security, use MS-CHAP v2 as the authentication protocol with strong passwords. | Offers the highest level of security, providing data confidentiality, data integrity, data origin authentication, and replay protection. |

| | | |
|---|---|---|
| Performance | A VPN server supports more PPTP connections than L2TP/IPSec connections. | Because IPSec encryption is processing intensive, a VPN server supports fewer L2TP connections than PPTP connections. To support additional L2TP connections, increase CPU processing power or use offload network adapters. |
| NAT support | PPTP-based VPN clients can be located behind a NAT if the NAT includes an editor that can translate PPTP. | If you locate L2TP/IPSec–based clients or servers behind a NAT, both client and server must support IPSec NAT traversal (NAT-T). |

**NAT Requirements for VPN Protocols**

If you are using a NAT with your VPN remote access server solution, your security plan for remote access must include the required setup for placing VPN clients behind a NAT. The VPN protocol that you deploy affects the NAT requirements.

A *network address translator (NAT)* translates the IP addresses and Transmission Control Protocol / User Datagram Protocol (TCP/UDP) port numbers of packets that are forwarded between a private network and the Internet. The NAT on the private network can provide IP address configuration information to the other computers on the private network.

The NAT can act as a simplified DHCP server that allocates an IP address, a subnet mask, a default gateway, and the IP address of a DNS server. The NAT can become the DNS proxy for the computers on the private network. When the NAT receives name resolution requests from a computer on the private network, it forwards the request to a specified Internet-based DNS server and returns a response to the requesting computer on the private network.

**Using a NAT with PPTP connections** If a VPN client that uses a PPTP connection is behind a NAT, the NAT must include a NAT editor that can translate PPTP traffic. The NAT editor is required, because PPTP tunneled data has a Generic Routing Encapsulation (GRE) header rather than a TCP header or a UDP header. The NAT editor uses the Call ID field in the GRE header to identify the PPTP data stream and translate IP addresses and call IDs for PPTP data packets that are forwarded between a private network and the Internet.

The NAT/Basic Firewall routing protocol component of the Routing and Remote Access service includes a NAT editor for PPTP traffic.

**Using a NAT with L2TP connections** *IPSec NAT Traversal (NAT-T)* enables IPSec peers to communicate when behind a NAT. IPSec NAT-T provides UDP encapsulation of IPSec packets to enable Internet Key Exchange (IKE) and Encapsulating Security Payload (ESP)–protected traffic to pass through a NAT. IKE automatically detects that a NAT is present and uses User Datagram Protocol — Encapsulating Security Payload (UDP-ESP) encapsulation to enable ESP-protected IPSec traffic to pass through the NAT.

To use NAT-T, both the remote access VPN client and the remote access server must support IPSec NAT–T. IPSec NAT-T is supported by Windows Server 2003 and Microsoft L2TP/IPSec VPN Client.

# Selecting an Authentication Protocol

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because L2TP/IPSec user authentication occurs after the VPN client and the VPN server have established a secure channel of communication, your choice of authentication protocol has no effect on VPN security if you use L2TP/IPSec. However the use of MS-CHAP v2 and EAP-TLS is recommended.

To use encryption on a PPTP connection, you must use one of the following authentication protocols:

- MS-CHAP
- MS-CHAP v2
- EAP-TLS

**Authentication Protocols for PPTP VPN Connections**
PPTP VPN connections require the use of the MS-CHAP, MS-CHAP v2, or EAP-TLS authentication protocols. Only these three authentication protocols provide a mechanism to generate the same encryption key on both the VPN client and the VPN server. MPPE uses this encryption key as a basis for encrypting all PPTP data sent over the VPN connection.

**MS-CHAP and MS-CHAP v2** MS-CHAP and MS-CHAP v2 are password-based authentication protocols. In the absence of certificates or smart cards, use MS CHAP v2, a stronger authentication protocol than MS CHAP, which provides mutual authentication. With *mutual authentication*, the VPN server authenticates the VPN client, and the VPN client authenticates the VPN server.

**Note**

- If you must use a password-based authentication protocol, enforce the use of strong passwords on your network. A strong password has more than eight characters and a random mixture of uppercase and lowercase letters, numbers, and punctuation marks. For example, "f3L*q02~>xR3w#4o" is a strong password. In an Active Directory domain, use Group Policy settings to enforce strong user passwords.

**EAP-TLS** The EAP-TLS authentication protocol is designed for use with a certificate infrastructure and either certificates or smart cards. With EAP-TLS, the VPN client sends its user certificate for authentication, and the authenticating server for the VPN server sends a computer certificate for authentication. This is the strongest authentication method, because it does not rely on passwords. For more information about EAP-TLS authentication, see "Connecting Remote Sites" in this book.

You can use Certificate Services in Windows Server 2003 as the CA for your organization, or you can use a third-party CA when you deploy EAP-TLS as your authentication method. For information about certificate requirements with Certificate Services, see "Network access authentication and certificates" in Help and Support Center for Windows Server 2003.

Using a third-party CA requires the following setup:

- The certificate in the computer store of the authenticating server must contain the Server Authentication certificate purpose in Enhanced Key Usage (EKU) extensions. A certificate purpose is identified with an object identifier (OID). The object identifier for Server Authentication is 1.3.6.1.5.5.7.3.1.
- The Subject Alternative Name property of the computer certificate must contain the fully qualified domain name (FQDN) of the computer account of the authenticating server.
- The cryptographic service provider for the computer certificates on the authenticating server must support the secure channel (Schannel) security package. Without support for Schannel, the authenticating server cannot use the certificate, and the certificate will not be available for use in the remote access policy.
- The certificate installed on a remote access client that is running Windows Server 2003 must contain the Client Authentication certificate purpose (OID 1.3.6.1.5.5.7.3.2).

- The Subject Alternative Name property of the user certificate must contain the FQDN of the user account of the VPN client.

- Both the certificate in the computer store of the authenticating server and the user certificate of the remote access client must contain a private key.

**Authentication Protocols for L2TP/IPSec Connections**

For L2TP/IPSec connections, you can use any user authentication protocol, because the authentication occurs after the VPN client and VPN server have established a secure channel of communication. This is referred to as an *IPSec security association (SA)*. It is strongly recommended that you use either MS-CHAP v2 or EAP-TLS to provide the most secure user authentication that is available.

**Guidelines for Selecting Authentication Protocols**

Consider the following factors when choosing an authentication protocol for VPN connections:

- If you use smart cards or have a certificate infrastructure that issues user and computer certificates, use the EAP-TLS authentication protocol for both PPTP and L2TP connections. EAP-TLS is supported by VPN clients running Windows 2000, Windows XP, or Windows Server 2003.

- If you must use a password-based authentication protocol, use MS-CHAP v2, and enforce strong passwords using Group Policy. MS-CHAP v2 is supported by VPN clients running Windows XP, Windows Server 2003, Windows 2000, Windows NT Workstation 4.0 with Service Pack 4 (SP4) and later, Windows Millennium Edition, or Windows 98.

- Use the most secure protocols that your network access servers and clients can support. If you need a high level of security, configure the remote access server and the authenticating server to accept only a few very secure authentication protocols. Alternatively, if flexibility is more important than maintaining a high level of security, configure the authenticating server to accept less secure authentication protocols. For more information about designing and deploying IAS, see "Deploying Internet Authentication Service (IAS)" in this book.

# Selecting the Scope and Level of Encryption

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

On a VPN, you protect your data by encrypting it between the VPN client and the VPN server. Always use data encryption for VPN connections when private data is sent across a public network, which always presents a risk of interception. For VPN connections, Windows Server 2003 uses MPPE for PPTP connections and IPSec encryption for L2TP connections.

**Note**

- Nonencrypted PPTP connections (over which the PPP frame is sent in plaintext) and nonencrypted non-IPSec-based L2TP connections (over which the PPP frame is sent in plaintext) are not secure, and they are not recommended for VPN connections over the Internet.

To ensure successful encryption and decryption, the sender and the receiver must use a common encryption key. The length of the encryption key is an important security parameter, especially over public networks. To ensure the highest level of encryption, use the largest key size.

**Protection Provided by Link Encryption**

In *link encryption,* data is encrypted only on the link between the VPN client and the VPN server. A VPN connection has link encryption, regardless of the VPN protocol in use. PPTP connections use MPPE with MS-CHAP, MS-CHAP v2, or EAP-TLS authentication. For L2TP/IPSec connections, IPSec provides encryption on the link between the VPN client and the VPN server.

When data encryption is performed between the VPN client and the VPN server, you do not need to encrypt the data on the communication link between a dial-up client and its ISP. For example, a mobile user might use a dial-up networking connection to dial in to a local ISP. After the Internet connection is made, the user creates a VPN connection with the enterprise VPN server. Because the VPN connection is encrypted, no encryption is needed on the dial-up networking connection between the user and the ISP.

**Providing End-to-End Encryption**

For an additional layer of security, configure end-to-end encryption. *End-to-end encryption* encrypts the data between the source host and the destination host. After a VPN connection is made, IPSec can be used to provide end-to-end encryption. For an L2TP/IPSec connection, IPSec end-to-end encryption is used in addition to IPSec link encryption.

Table 8.2 shows which authentication methods support specific encryption requirements.

**Table 8.2 Encryption Support Provided Under CHAP, MS-CHAP, and EAP-TLS**

| Requirement | Authentication Protocols | Encryption Enforcement |
|---|---|---|
| Secured password with no data encryption | CHAP, MS-CHAP, MS-CHAP v2 | Optional encryption. (Connect even if no encryption.) |
| Secured password with MPPE data encryption | MS-CHAP, MS-CHAP v2 | Required encryption. (Disconnect if server declines.) |
| Smart card with no data encryption | EAP-TLS | Optional encryption. (Connect even if no encryption.) |
| Smart card with data encryption | EAP-TLS | Require encryption. (Disconnect if server declines.) |

Data encryption for L2TP connections relies on IPSec, which does not require any specific authentication protocol. IPSec enforces the encryption; if the server declines data encryption, the connection is denied.

The strength of link encryption is set through the remote access policies that govern PPTP and L2TP connections on the server. A *remote access policy* is a collection of conditions and settings that define authorization and access privileges for connection attempts. For IAS servers and servers running Routing and Remote Access, remote access policies are used to determine whether a connection attempt is accepted or rejected.

Table 8.3 shows the encryption support provided for PPTP and L2TP/IPSec connections by each level of encryption that is set in a remote access policy.

**Table 8.3 Encryption Required at Each Encryption Level for PPTP and L2TP/IPSec Connections**

| Encryption Level | PPTP Encryption Required | L2TP Encryption Required |
|---|---|---|
| **No Encryption** | No encryption required. | No encryption required. |
| **Basic** | MPPE 40-bit data encryption | IPSec 56-bit Data Encryption Standard (DES) |
| **Strong** | MPPE 56-bit data encryption | IPSec 56-bit DES |
| **Strongest** | MPPE 128-bit encryption | IPSec 168-bit Triple DES (3DES) |

For a procedure for setting the encryption level in a remote access policy, see "Configuring authentication and data encryption" in Help and Support Center for Windows Server 2003. For more information about using Windows Server 2003 remote access policies, see "Introduction to remote access policies" in Help and Support Center for Windows Server 2003.

# Planning a Certificate Infrastructure to Support Client Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The authentication methods that are deployed for remote access clients determine whether or not your remote access design requires a certificate infrastructure.

A certificate infrastructure is required for L2TP/IPSec-based VPN connections, because certificates are used to negotiate IPSec peer authentication. For PPTP-based VPN connections, a certificate infrastructure is required if either smart cards or certificates and EAP-TLS authentication are in use. Password-based authentication protocols, such as MS-CHAP v2, do not use certificates in authentication; therefore, a certificate infrastructure is not required.

Table 8.4 shows where you must install certificates to support remote access clients over L2TP/IPSec-based VPN connections and over PPTP-based connections using EAP-TLS. For EAP-TLS authentication, L2TP/IPSec-based VPN connections require one more certificate on the VPN client than PPTP-based VPN connections require.

**Table 8.4 Certificate Infrastructures Required for Remote Access Client Authentication**

| VPN/Authentication Protocol | Required Certificate Infrastructure |
|---|---|
| L2TP/IPSec-based VPN connection | • Install a computer certificate on the VPN server.<br>• Install a computer certificate on each VPN client. |
| PPTP-based VPN connection using smart cards and EAP-TLS | • Install a computer certificate on the authenticating server for the VPN server.<br>• Install a user certificate on each smart card. |
| PPTP-based VPN connection using registry-based user certificates and EAP-TLS | • Install a computer certificate on the authenticating server for the VPN server.<br>• Install a user certificate on each VPN client. |

If your PPTP-based VPN connections require a certificate infrastructure, install a computer certificate on the authenticating server for the VPN server. If you are using smart cards, install a user certificate on each smart card distributed to a VPN client user. If you are using registry-based user certificates with EAP-TLS authentication, install a user certificate on each VPN client.

For an L2TP/IPSec-based VPN connection, install a computer certificate on all VPN clients and on the VPN server. A certificate infrastructure is also required when you are using either smart cards or certificates and EAP-TLS for user authentication.

For more information about certificate requirements, see "Network access authentication and certificates" in Help and Support Center for Windows Server 2003.

For more information about deploying certificate services to support L2TP/IPSec, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

# Planning for Network Access Quarantine Control

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because typical remote access connections only validate the credentials of a remote access user, a remote access client that connects to a private network can access network resources even if the configuration of the remote access client does not comply with corporate network policies. You can implement Network Access Quarantine Control to delay normal remote access to a private network until the configuration of the remote access client has been examined and validated by a client-side script.

When a remote access client initiates a connection to a remote access server, the user is authenticated, and the remote access client is assigned an IP address. If Network Access Quarantine Control is in use, the connection is placed in quarantine mode until a client-side script is run on the remote access client and the configuration of the remote access client is validated against current network policies. While the remote access connection is in quarantine mode, network access is limited. When the remote access server is notified that the configuration of the remote access client is validated against current network policies, quarantine mode is removed, and the remote access client is granted normal remote access.

The components for Network Access Quarantine Control are included in the *Microsoft® Windows® Server 2003 Resource Kit Tools*. For instructions on setting up Network Access Quarantine Control, see "Configuring Network Access Quarantine Control" later in this chapter.

**Note**

- Network Access Quarantine Control is designed to prevent clients with unsafe configurations from attaching to a private network. It does not protect a private network from malicious users who have obtained a valid set of credentials.

**Processing a Connection Attempt Under Network Access Quarantine Control**

Under Network Access Quarantine Control, the user on a quarantine-compatible remote access client uses an installed Connection Manager profile to connect with a quarantine-compatible remote access server. The remote access client passes its authentication credentials to the remote access server. The Routing and Remote Access service sends a RADIUS Access-Request message to the IAS server. The IAS server validates the authentication credentials of the remote access client and, assuming valid credentials, checks its remote access policies. If the connection attempt matches the quarantine policy, the connection is accepted with quarantine attributes and the connection is placed in quarantine mode.

While the connection attempt is in quarantine mode, the remote access server implements a set of network restrictions for the connection. These network restrictions are configured in IAS. The IAS server sends a RADIUS Access-Accept message that contains the MS-Quarantine-IPFilter and MS-Quarantine-Session-Timeout attributes. The remote access client completes the remote access connection, obtaining an IP address and other configuration settings, and the Windows Server 2003 Routing and Remote Access service configures the MS-Quarantine-IPFilter and MS-Quarantine-Session-Timeout settings for the connection. At this point, the remote access client can only successfully send traffic that matches the quarantine filters. The client must notify the remote access server that the client has passed network compliance testing within the time limit (in seconds) specified by the MS-Quarantine-Session-Timeout attribute in the quarantine remote access policy.

**Note**

- The process described in this section incorporates the use of both the MS-Quarantine-IPFilter attribute and the MS-Quarantine-Session-Timeout attribute. Both attributes are optional.

The Connection Manager profile initiates a post-connect action, which runs the embedded client-side script. The script verifies that the remote access computer's configuration complies with network policy requirements. If the script runs successfully, the script runs the notification component, Rqc.exe, which notifies the remote access server that the remote access client complies with network policy.

The listener component on the remote access server, known as the Remote Access Quarantine Agent service (Rqs.exe), receives the notification. Routing and Remote Access removes the MS-Quarantine-IP Filter and MS-Quarantine-Session-Timeout settings from the connection, giving the remote access client normal access to the intranet.

**Components of Network Access Quarantine Control**

Figure 8.6 shows the components of Windows remote access for Network Access Quarantine Control.

**Figure 8.6 Components of Network Access Quarantine Control for Remote Access**




This configuration consists of the following components:

- Quarantine-compatible access clients
- Quarantine-compatible access server
- Quarantine-compatible RADIUS server
- Quarantine resources
- Accounts database
- Quarantine remote access policy

**Quarantine-compatible access clients**

The remote access client must be a computer running one of the following operating systems: Microsoft® Windows® XP Professional, Windows® XP Home Edition, Windows 2000, Windows Millennium Edition, Windows 98 Second Edition, or Windows Server 2003.

These versions of Windows support Connection Manager profiles that are created by the Connection Manager Administration Kit (CMAK) provided with Windows Server 2003 and contain:

- A post-connect action setting that runs a network policy requirements script.

  The post-connect action setting is configured when the CM profile is created with CMAK.

- A network policy requirements script

  The network policy requirements script performs validation checks on the remote access client computer to verify that it conforms to network policies. The script can be a custom executable file or as simple as a command file (also known as a batch file).

- A notifier component.

  When the script has run successfully and the connecting computer has satisfied all of the network policy requirements verified by the script, the script executes a notifier component (an executable file) with the appropriate parameters. The notifier component sends a message to the quarantine-compatible remote access server that indicates a successful execution of the script. You can use your own notifier component or you can use Rqc.exe, which is provided with the *Windows Server 2003 Resource Kit Tools*.

With these components installed, the remote access client computer uses the CM profile to perform its own network policy requirements check and indicate its success to the remote access server as part of the connection setup.

**Quarantine-compatible access server**
A quarantine-compatible remote access server requires the following components:

- A computer running Windows Server 2003 and the Routing and Remote Access service.

  Routing and Remote Access with Windows Server 2003 supports the use of a listener component and the MS-Quarantine-IPFilter and MS-Quarantine-Session-Timeout RADIUS vendor-specific attributes (VSAs), which are used to specify quarantine settings.

- A listener component.

  The listener component listens for messages from quarantine-compatible remote access clients that indicate that their scripts have run successfully. You can create your own custom listener component (matched with your own custom notifier component) or you can install the Remote Access Quarantine Agent service (Rqs.exe) from the *Windows Server 2003 Resource Kit Tools*.

With these components installed, the remote access server implements quarantine mode for connecting remote access clients and listens for notifier messages that the clients have satisfied network policy requirements and can be taken out of quarantine mode.

**Quarantine-compatible RADIUS server**
A quarantine-compatible RADIUS server requires a computer running Windows Server 2003 and IAS, which supports the configuration of the MS-Quarantine-IPFilter and MS-Quarantine-Session-Timeout RADIUS vendor-specific attributes (VSAs) to specify quarantine settings for the quarantine-compatible remote access server.

**Quarantine resources**
In quarantine mode, a remote access client must have access to the following resources:

- To perform name resolution, the client must have access to DNS servers.

- To obtain the latest version of the script, the client must have access to file servers with anonymous access allowed.

- To obtain instructions and components needed to bring the remote access client into compliance with network policies, the client must have access to Web servers with anonymous access allowed.

**Accounts database**
For Windows Server 2003-based networks, the Active Directory service is used as the accounts database to store user accounts and their dial-in properties.

**Quarantine remote access policy**
For Network Access Quarantine Control, you must configure a quarantine remote access policy with the appropriate conditions for remote access connections, and with profile settings that specify the MS-Quarantine-IPFilter and MS-Quarantine-Session-Timeout attributes (configured on the **Advanced** tab of the profile).

- The MS-Quarantine-IPFilter attribute is used to configure inbound and outbound packet filters to allow only the traffic generated by the notifier component. If you are using Rqc.exe, configure a single inbound packet filter to only allow traffic from TCP port 7250 and to TCP port 7250 (the default TCP port for Rqc.exe), and specify that all other traffic be discarded. Additional packet filters are needed in order for the quarantined remote access client to access the quarantine resources. These include filters that allow the remote access client to access DNS servers, file shares, and Web servers.

  The packet filters configured for the MS-Quarantine-IPFilter attribute provide the quarantine, or isolation, of the traffic of the remote access client until the notifier component on the remote access client indicates that the computer is in compliance with network policies.

- The MS-Quarantine-Session-Timeout attribute specifies how long the remote access server waits to receive the notification that the script has executed successfully before terminating the connection.

# Enhancing Security by Using Remote Access Account Lockout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To prevent dictionary attacks on remote server accounts, you can use remote access account lockout. When deciding whether or not to use remote access account lockout, remember that if you enable remote access account lockout, a malicious user can intentionally attempt multiple authentications for a user account to force the account to be locked out, thereby preventing the authorized user from being able to create a remote access connection.

Remote access account lockout is configured in the registry in Windows Server 2003. It is not related to the account lockout policy for domain or local user accounts.

To enable remote access account lockout, modify the following subkey in the registry on the server that authenticates remote access requests:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess \Parameters\AccountLockout

If the remote access server is configured for Windows authentication, modify the registry on that server. If the remote access server is configured for RADIUS authentication, and you are using IAS, modify the registry on the IAS server.

For more information about modifying the AccountLockout subkey, see "Configuring Remote Access Account Lockout for a VPN Solution" later in this chapter.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry reference on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

If your organization is using smart cards, the smart card manufacturer controls account lockout for personal identification numbers (PINs) that are not valid. Recovery from account lockout as a result of an invalid PIN might require smart card replacement.

For more information about remote access account lockout, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Optimizing Your Remote Access Server Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In addition to increasing performance by upgrading server hardware, consider increasing availability, security, and performance by incorporating the following elements in your remote access design:

- Redundant servers for increased availability
- Network Load Balancing for increased availability and performance

**Increasing Availability by Using Redundant Servers**
Remote access solutions with redundant servers can provide higher availability for remote access clients. If degradation of service is not a critical issue, you can use your primary remote access servers as backups for each other. If service degradation is not acceptable, provide redundancy by enlisting an extra server to provide failure protection.

If one or more user groups require high-priority access, consider using separate remote access servers for these user groups.

**Increasing Performance by Using Network Load Balancing**
By using Network Load Balancing, which is available in Windows Server 2003, you can increase VPN server performance and availability. Network Load Balancing distributes traffic from remote access VPN clients among multiple VPN servers.

Network Load Balancing also provides immediate failover if a VPN server fails. If a VPN server fails, client sessions handled by that server also fail. Clients are prompted to log on again, and their new session is handled by one of the remaining hosts.

To provide load balancing for VPN clients, use the default port rule in configuring all hosts, as follows:

- Set the port range to 0–65535 (the default). The default range covers all of the ports, so the port rule remains valid even if there is a change in the port numbers that you want to cover.
- Accept the default filtering mode, load weight/equal load distribution, and affinity settings.

For more information about using Network Load Balancing in a VPN scenario, see "Deploying Network Load Balancing" in *Planning Server Deployments* of this kit.

# Testing Your Remote Access Server Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you complete the design of your remote access server solution, test the design. Testing the design includes testing individual VPN client access to VPN servers, as well as comprehensive testing of the entire external connectivity design. If you are integrating multiple remote access solutions, test them together after testing them individually.

Because of vulnerabilities that exposure to the Internet might introduce, isolate your network perimeter from your intranet during testing. Do not integrate your network perimeter with your intranet until you are confident that you have addressed all security issues.

Be sure to test the ISP infrastructure, including the RADIUS proxy (if applicable), and a representative sampling of access points.

Testing is critical to the security of any external connectivity solution, and it is important for ensuring that the connections function as planned. Before testing, simulate both internal and external connections to prevent exposure and corruption of any part of your network.

**Tools for Testing a Remote Access Server Design**

The following tools are useful in testing a remote access server design:

- TCP/IP troubleshooting tools, including Netsh, Ping, Pathping, Route, and Tracert.
- Remote access logging in the Routing and Remote Access snap-in. The log includes authentication and accounting information.
- Event Viewer. This is an administrative tool for Windows Server 2003 that displays monitoring and troubleshooting information from Windows and other programs.
- Network Monitor. This is an optional networking component for Windows Server 2003 that allows a system administrator to capture and examine packets on a LAN and save the packets to a capture file.
- Remote Access Event Tracing. This is enabled through Netsh.

# Deploying a VPN Remote Access Server Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying a remote access server solution involves configuring the server as a VPN remote access server; configuring routing on the VPN server and the VPN clients; and implementing security — installing any certificates required for the authentication of remote clients, configuring encryption in requirements the remote access policy for VPN connections, and optionally configuring remote access account lockout. Figure 8.7 shows the process for deploying a VPN remote access server solution.

**Figure 8.7 Deploying a VPN Remote Access Server Solution**




**Important**

- The procedures for deploying a VPN remote access server solution assume that you have deployed Active Directory on the server, have a PKI in place, and have deployed an IAS server.

- For information about designing and deploying a PKI, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit. For information about designing and deploying IAS, see "Deploying Internet Authentication Service (IAS)" in this book. For more information about Active Directory, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

# Configuring a VPN Remote Access Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The configuration of a VPN remote access server involves the following tasks:

- Configure the server as a VPN remote access server.
- Configure TCP/IP on the server.
- Configure name resolution on the server.
- Configure packet filters for the server.
- Optionally, configure Network Access Quarantine Control.

# Configuring the Server as a VPN Remote Access Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure the server as a VPN remote access server, use the Configure Your Server Wizard and select **Remote access/VPN server** as the server role, or use the Rounting and Remote Access snap-in. For instructions on using the wizard, see "Remote access/VPN server role: Configuring a remote access/VPN server" in Help and Support Center for Windows Server 2003.

# Configuring TCP/IP on the VPN Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After configuring the server as a remote access server, configure the TCP/IP settings for the Internet or perimeter network interface and for the intranet interface.

**Note**

- Because of routing issues related to configuring TCP/IP automatically, it is recommended that you not configure a VPN server as a DHCP client. Instead, manually configure TCP/IP on the intranet interfaces of a VPN server. For a full discussion of the routing options for a VPN server, see "Configuring Routing on a VPN Server" later in this chapter.

Manually configure the Internet or perimeter network interface of the VPN server with a default gateway. Configure the TCP/IP settings with a public IP address, a subnet mask, and the default gateway of either the firewall (if the VPN server is connected to a perimeter network) or an ISP router (if the VPN server is connected directly to the Internet).

**To configure TCP/IP for the Internet or perimeter network interface**

1. In **Control Panel**, double-click **Network Connections**, and then double-click the network adapter for the Internet or perimeter network interface.

2. In the network adapter status dialog box (for example, **Local Area Connection Status**), click **Properties**.

3. SelectInternet Protocol (TCP/IP), and then click **Properties**.

4. On the **General** tab, configure the IP address, subnet mask, and default gateway.

   The IP address must be a public IP address assigned by an ISP. As an option, you can configure the VPN server with a private IP address but assign it a published static IP address by which it is known on the Internet. When packets are sent to and from the VPN server, a NAT that is positioned between the Internet and the VPN server translates the published IP address to the private IP address.

   When you configure a VPN connection, give your VPN servers names that can be resolved to IP addresses using DNS.

5. Click **Advanced** to display the **Advanced TCP/IP Settings** dialog box.

6. To prevent the VPN server from dynamically registering the public IP address of its Internet interface with an intranet DNS server, on the **DNS** tab, clear the **Register this connection's addresses in DNS** check box. This check box is cleared by default.

7. To prevent the VPN server from registering the public IP address of its Internet interface with intranet WINS servers, on the **WINS** tab, select the **Disable NetBIOS over TCP/IP** check box. This check box is selected by default.

When you configure TCP/IP for the VPN server's intranet interface, do not configure the default gateway on the intranet connection. This will prevent default route conflicts with the default route pointing to the Internet.

**To configure TCP/IP for the intranet interface**

1. In **Control Panel**, double-click **Network Connections,** and then double-click the network adapter for intranet interface.

2. In the network adapter status dialog box (for example, **Local Area Connection 2 Status**), click **Properties**.

3. Select **Internet Protocol (TCP/IP),** and then click **Properties**.

4. On the **General** tab, configure the IP address, subnet mask, and DNS server address.

To prevent default route conflicts with the default route pointing to the Internet, do not configure the default gateway on the intranet connection.

5. Click **Advanced** to display the **Advanced TCP/IP Settings** dialog box.

6. On the **WINS** tab, configure the IP addresses of your WINS servers.

# Configuring Name Resolution on a VPN Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you use Domain Name System (DNS) to resolve intranet host names or Windows Internet Name Service (WINS) to resolve intranet NetBIOS names, manually configure the VPN server with the IP addresses of the appropriate DNS and WINS servers.

During the PPP connection setup process, VPN clients receive the IP addresses of DNS and WINS servers. By default, the VPN clients inherit the DNS and WINS server IP addresses configured on the VPN server. However, VPN clients that are capable of sending a DHCPINFORM message (computers running Windows 2000, Windows XP, or Windows Server 2003) get their DNS and WINS server IP addresses from the DHCP server.

# Configuring Packet Filters for a VPN Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The firewall is configured with rules to filter the packets that a VPN server sends and receives and control intranet traffic to and from VPN clients, based on your network security policies. Packet filtering is based on the fields of inbound and outbound packets.

The Routing and Remote Access Server Setup Wizard for Windows Server 2003 and Windows Server 2000 Service Pack 2 (SP2) or later automatically configures the appropriate packet filters for VPN traffic. Alternatively, you can use the Routing and Remote Access snap-in to configure the packet filters.

The following sections summarize the packet filters that are required when the VPN server is placed behind a firewall or in front of a firewall. For procedures explaining how to enter the packet filter configurations, see "VPN servers and firewall configuration" in Help and Support Center for Windows Server 2003.

**Configuring Filters for a VPN Server Behind a Firewall**

If the VPN server is behind a firewall, you must configure packet filters for both an Internet interface and a perimeter network interface. In this configuration, the firewall is connected to the Internet, and the VPN server is an intranet resource that is connected to the perimeter network. The VPN server has an interface on both the perimeter network and the Internet.

**PPTP connections**

For a PPTP connection, configure the following packet filters on the Internet and perimeter network interfaces of the firewall.

**Internet interface of the firewall** On the firewall's Internet interface, configure the inbound and outbound filters in Table 8.5, specifying that all packets are dropped except those that are selected by the filters.

**Table 8.5 VPN Server Behind a Firewall: PPTP Filters on the Firewall's Internet Interface**

| | Filter | Action |
|---|---|---|
| Inbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**TCP destination port** = 1723 (0x6BB) | Allows PPTP tunnel maintenance traffic from the PPTP client to the PPTP server. |
| | **Destination IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 47 (0x2F) | Allows PPTP tunneled data from the PPTP client to the PPTP server. |
| | **Destination IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site (also known as router-to-router) VPN connection. If you allow all traffic to the VPN server from TCP port 1723, network attacks can emanate from sources on the Internet that use this port. You should only use this filter in conjunction with the PPTP filters that are also configured on the VPN server. |

| | | |
|---|---|---|
| Outbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Allows PPTP tunnel maintenance traffic from the PPTP server to the PPTP client. |
| | **Source IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 47 (0x2F) | Allows PPTP tunneled data from the PPTP server to the PPTP client. |
| | **Source IP address** = Perimeter network interface of VPN server<br><br>**TCP destination port** = 1723 (0x6BB) | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. If you allow all traffic from the VPN server to TCP port 1723, network attacks can emanate from sources on the Internet using this port. You should only use this filter in conjunction with the PPTP filters that are also configured on the VPN server. |

**Perimeter network interface of the firewall** On the firewall's perimeter network interface, configure the inbound and outbound filters in Table 8.6, specifying that all packets are dropped except those that are specified by the filters.

**Table 8.6 VPN Server Behind a Firewall: PPTP Filters on the Perimeter Network Interface**

| | Filter | Action |
|---|---|---|
| Inbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Allows PPTP tunnel maintenance traffic from the VPN server to the VPN client. |
| | **Source IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 47 (0x2F) | Allows PPTP tunneled data from the VPN server to the VPN client. |
| | **Source IP address** = Perimeter network interface of VPN server<br><br>**TCP destination port** = 1723 (0x6BB) | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. If you allow all traffic from the VPN server to TCP port 1723, network attacks can emanate from sources on the Internet using this port. |

| | | |
|---|---|---|
| Outbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Allows PPTP tunnel maintenance traffic from the PPTP client to the PPTP server. |
| | **Destination IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 47 (0x2F) | Allows PPTP tunneled data from the PPTP client to the PPTP server. |
| | **Destination IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. If you allow all traffic to the VPN server from TCP port 1723, network attacks can emanate from sources on the Internet using this port. |

**L2TP/IPSec connections**

For an L2TP/IPSec connection, configure the following packet filters on the Internet and perimeter network interfaces of the firewall.

**Internet interface of the firewall** On the firewall's Internet interface, configure the inbound and outbound filters in Table 8.7, specifying that all packets are dropped except those that are specified by the filters.

**Table 8.7 VPN Server Behind a Firewall: L2TP/IPSec Filters on the Firewall's Internet Interface**

| | Filter | Action |
|---|---|---|
| Inbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**UDP destination port** = 500 (0x1F4) | Allows IKE traffic to the VPN server. |
| | **Destination IP address** = Perimeter network interface of VPN server<br><br>**UDP destination port** = 4500 (0x1194) | Allows IPSec NAT-T traffic to the VPN server. |
| | **Destination IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 50 (0x32) | Allows IPSec ESP traffic to the VPN server. |
| Outbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**UDP source port** = 500 (0x1F4) | Allows IKE traffic from the VPN server. |
| | **Source IP address** = Perimeter network interface of VPN server<br><br>**UDP source port** = 4500 (0x1194) | Allows IPSec NAT-T traffic from the VPN server. |
| | **Source IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 50 (0x32) | Allows IPSec ESP traffic from the VPN server. |

No filters are required for L2TP traffic at UDP port 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

**Perimeter network interface of the firewall** On the firewall's perimeter network interface, configure the inbound and outbound filters in Table 8.8, specifying that all packets are dropped except those that are selected by the filters.

**Table 8.8 VPN Server Behind a Firewall: L2TP/IPSec Filters on the Firewall's Perimeter Network Interface**

| | Filter | Action |
|---|---|---|
| Inbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**UDP source port** = 500 (0x1F4) | Allows IKE traffic from the VPN server. |
| | **Source IP address** = Perimeter network interface of VPN server<br><br>**UDP source port** = 4500 (0x1194) | Allows IPSec NAT-T traffic from the VPN server. |
| | **Source IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 50 (0x32) | Allows IPSec ESP traffic from the VPN server. |
| Outbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**UDP destination port** = 500 (0x1F4) | Allows IKE traffic to the VPN server. |
| | **Destination IP address** = Perimeter network interface of VPN server<br><br>**UDP destination port** = 4500 (0x1194) | Allows IPSec NAT-T traffic to the VPN server. |
| | **Destination IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 50 (0x32) | Allows IPSec ESP traffic to the VPN server. |

**Configuring Filters for a VPN Server in Front of a Firewall**

When a VPN server is in front of a firewall and connected to the Internet, configure inbound and outbound packet filters on the VPN server to allow only VPN traffic to and from the IP address of the VPN server's Internet interface. Use this configuration if your VPN server is in a perimeter network, with one firewall positioned between the VPN server and the intranet and another between the VPN server and the Internet.

**PPTP connections**

For a PPTP connection, configure the VPN server with the inbound and outbound filters in Table 8.9, specifying that all packets be dropped except those that are specified by the filters. These filters are automatically configured when you:

- Rrun the Routing and Remote Access Server Setup Wizard and choose the **Remote access (dial-up or VPN)** option.
- Select the correct interface.
- Select the **Enable security on the selected interface by setting up packet filters** option on the VPN Connection page. This setting is enabled by default.

**Table 8.9 VPN Server in Front of a Firewall: Packet Filters for PPTP**

| | Filter | Action |
|---|---|---|
| Inbound | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**TCP destination port** = 1723 | Allows PPTP tunnel maintenance to the VPN server. |

| | | |
|---|---|---|
| | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**IP Protocol ID** = 47 | Allows PPTP tunneled data to the VPN server. |
| | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**TCP (established) source port** = 1723 | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. Accepts TCP traffic only when a VPN server initiates the TCP connection. |
| Outbound | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**TCP source port** = 1723 | Allows PPTP tunnel maintenance traffic from the VPN server. |
| | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**IP Protocol ID** = 47 | Allows PPTP tunneled data from the VPN server. |
| | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**TCP (established) destination port** = 1723 | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. Sends TCP traffic only when a VPN server initiates the TCP connection. |

**L2TP/IPSec connections**

For an L2TP/IPSec connection, configure the VPN server with the inbound and outbound filters in Table 8.10, specifying that all packets be dropped except those that are specified by the filters.

**Table 8.10 VPN Server in Front of a Firewall: Packet Filters for L2TP/IPSec**

| | Filter | Action |
|---|---|---|
| Inbound | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP destination port** = 500 | Allows IKE traffic to the VPN server. |

| | | |
|---|---|---|
| | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP destination port** = 1701 | Allows L2TP traffic from the VPN client to the VPN server. |
| | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP destination port** = 4500 | Allows IPSec NAT-T traffic from the VPN client to the VPN server. |
| Outbound | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP source port** = 500 | Allows IKE traffic from the VPN server. |
| | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP source port** = 1701 | Allows L2TP traffic from the VPN server to the VPN client. |
| | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP source port** = 4500 | Allows IPSec NAT-T traffic from the VPN server to the VPN client |

# Configuring Network Access Quarantine Control

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Perform the following steps to configure Network Access Quarantine Control:

1. Create quarantine resources.

   a. Configure a DNS server, a file server and share point for updated scripts, and a Web server as quarantine resources.

   b. Create Web pages containing network policy compliance instructions.

2. Create a *notification component* that notifies the remote access server that the remote access client complies with network policy requirements. The notification component sends a notifier message that indicates that a client-side script has run successfully on the remote access client, network policy requirements have been met, and the remote access connection quarantine restrictions can be removed. If you do not want to create your own notification component, you can use Rqc.exe in the *Windows Server 2003 Resource Kit Tools*.

3. Create a *client-side script* that validates the client configuration based on your network policy requirements. If all of the verification checks in the script are successful, the script executes the notification component with the appropriate parameters.

4. Create a *listener component* that receives the network policy compliance notification from the notification component. If you do not want to create a listener component, use Rqs.exe in the *Windows Server 2003 Resource Kit Tools*.

   **Note**

   - Rqs.exe and Rqc.exe use TCP port 7250 by default. When you create the quarantine policy, you must configure quarantine inbound filters to allow network traffic on TCP port 7250. Otherwise, Rqc.exe, which runs on client computers, cannot notify Rqs.exe that the client-side script has run successfully. If you specify another TCP port for Rqc.exe and Rqs.exe, you must configure the filter to allow traffic on that TCP port.

5. Create a quarantine Connection Manager profile, to be installed on all remote access clients that access servers participating in Network Access Quarantine Control. Only those remote access clients that have the quarantine Connection Manager profile installed can obtain a full-access connection.

   Use the Windows Server 2003 Connection Manager Administration Kit (CMAK) to create a profile with the following elements:

   - Specify a post-connect action to run the client-side script with the appropriate parameters.

   - Embed the client-side script and the notification component within the profile.

   For information about creating a Connection Manager profile using CMAK, see "Deploying Remote Access Clients Using Connection Manager" in this book.

6. Install the Quarantine Connection Manager profile on all remote access clients that access servers participating in Network Access Quarantine Control.

7. Use the New Remote Access Policy Wizard to create a *quarantine remote access policy* that restricts a remote access client's access while the client computer's configuration is verified against network policy requirements. The quarantine remote access policy can contain the following attributes:

- MS-Quarantine-IPFilter, to restrict a quarantined remote access client's access to only quarantine resources and the port designated for notification traffic.

- MS-Quarantine-Session-Timeout, to restrict the length of time during which a client can remain connected in quarantine mode before being disconnected.

To be quarantine-compatible, a remote access server must be running Windows Server 2003 and the Routing and Remote Access service. Routing and Remote Access with Windows Server 2003 supports the use of a listener component and the RADIUS vendor-specific attributes (VSAs) MS-Quarantine-IP Filter and MS-Quarantine-Session-Timeout, which are used to specify quarantine settings.

**Note**

- For an overview of Network Access Quarantine Control, see "Planning for Network Access Quarantine Control" earlier in this chapter.

# Configuring Routing for a VPN Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure routing on your intranet to allow computers to reach VPN clients. After either configuring static routing on the VPN server or configuring the server as a dynamic router, configure routing on the VPN clients.

# Configuring Routing on a VPN Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To enable a VPN server to correctly forward traffic to locations on your intranet, perform one of two routing configurations:

- Configure the server with static routes that summarize all possible IP addresses on the intranet.
- Configure the server with routing protocols that enable it to act as a dynamic router, automatically adding routes for intranet subnets to its routing table.

In a small, stable networking environment, static routing might be an appropriate choice for a VPN solution. However, in most corporate networking environments, the increased administrative overhead required to maintain static routes is prohibitive. The preferred method for a VPN solution is to configure the VPN server as a dynamic router.

**Configuring Static Routes on the Server**
If you manually configure IP address ranges for a static address pool on any of your VPN servers, and if any of the ranges is an off-subnet range, your intranet routing infrastructure must include routes representing the off-subnet address ranges. To provide the best summarization of address ranges for routes, choose your address ranges so that they can be expressed using a single prefix and subnet mask.

To ensure this, add static routes representing the off-subnet address ranges to the routers neighboring the VPN servers, and then use the routing protocol of your intranet to propagate the off-subnet routes to other routers. When you add the static routes to the neighboring routers, specify that the gateway or the next hop address is the intranet interface of the VPN server.

For information about adding static routes, see "Configuring the branch office network" in Help and Support Center for Windows Server 2003.

**Configuring the Server as a Dynamic Router**
If you are using RIP or OSPF, you can configure any VPN server that is using off-subnet address ranges as a RIP or OSPF router.

For OSPF, you must also configure the VPN server as an autonomous system boundary router (ASBR). For more information, see "OSPF design considerations" in Help and Support Center for Windows Server 2003.

If you use a routing protocol other than a RIP or OSPF, such as Interior Gateway Routing Protocol (IGRP), on the VPN server's neighboring intranet router, configure the interface connected to the subnet to which the VPN server is assigned for RIP or OSPF and configure all other interfaces for IGRP.

To configure the VPN server with an on-subnet address range, configure the VPN server to obtain IP addresses through DHCP or manually configure on-subnet address ranges.

For information about:

- Configuring the VPN server as a RIP router, see "Configure RIP for IP" in Help and Support Center for Windows Server 2003.
- Configuring the VPN server as an OSPF router, see "OSPF design considerations" and "Configure OSPF" in Help and Support Center for Windows Server 2003.

# Configuring Routing on a VPN Client

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, when a Windows-based VPN client makes a VPN connection, the VPN client automatically adds a new default route for the VPN connection and sets a higher metric for the existing default route. Because a new default route has been added, all Internet locations, except for the IP address of the tunnel server and locations based on other routes, are not reachable for the duration of the VPN connection.

Whether the default route is acceptable for the VPN connection depends on your remote access clients' needs (whether they need simultaneous access to both the intranet and the Internet) and security issues. For a full discussion of the routing options for VPN remote access clients, see "Determining Routing for VPN Remote Access Clients" earlier in this chapter.

Based on your design, implement one of the following routing options on the VPN client:

- If the remote access user does not require concurrent access to intranet and Internet resources, use the default gateway for the VPN connection.

- If the remote access user requires concurrent access to intranet and Internet resources over a VPN connection, choose one of the following options:

    - If you want to allow Internet access through the organization's intranet, use the default gateway for your VPN connection.

      Internet traffic between the VPN client and Internet hosts passes though firewalls or proxy servers as though the VPN client were physically connected to the organization's intranet. This method can affect performance, but it enables an organization to filter and monitor Internet access according to its network policies while the VPN client is connected to the organization network.

    - If the addressing within your intranet is based on a single class-based network ID, and the addresses assigned to VPN clients are from that single class-based network ID, prevent the use of the default gateway for your VPN connection.

    - If the addressing within your intranet is not based on a single class-based network ID, prevent the use of the default gateway for your VPN connection. Then, use one of the split tunneling methods described in "Determining Routing for VPN Remote Access Clients" earlier in this chapter.

**To prevent the VPN client from creating a new default route during a VPN connection**

1. In **Control Panel**, double-click **Network Connections,** and then double-click the name of the VPN connection.

2. In the **Connect** dialog box, click **Properties**.

3. In the properties dialog box for the VPN connection, click the **Networking** tab.

4. SelectInternet Protocol (TCP/IP), and then click **Properties**.

5. On the **General** tab, click **Advanced** to display the **Advanced TCP/IP Settings** dialog box.

6. To prevent a default route from being created during a VPN connection, on the **General** tab, clear the **Use default gateway on remote network** check box.

   No default route will be created for the connection. However, a route corresponding to the Internet address class of the assigned IP address will be created. For example, if the IP address assigned during the connection process is 10.0.12.119, the Windows Server 2003–based or Windows XP–based VPN client creates a route for the class-based network ID 10.0.0.0 with the subnet mask 255.0.0.0.

# Implementing Security for a VPN Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After configuring the VPN server as a remote access server and configuring routing on both the server and the remote access clients, implement security of a VPN solution by performing the following tasks:

- Install certificates for L2TP/IPSec VPN connections.
- Configure the appropriate level of data encryption.
- Optionally, configure remote access account lockout.

# Installing Certificates for VPN Connections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A certificate infrastructure is a requirement for L2TP/IPSec-based VPN connections. Certificates provide stronger authentication security than password-based authentication does.

To provide a certificate infrastructure for a VPN client that makes L2TP/IPSec connections:

1. Install a certificate in the Local Computer certificate store on the VPN server.

2. Install a user certificate in the Current User certificate store of each client.

   The certificate provides authentication for establishing IPSec security associations (SAs).

To provide a certificate infrastructure for user-level authentication with EAP-TLS:

1. Install a certificate on the authenticating server for the VPN server.

2. If you are not using smart cards, install a registry-based user certificate on each client.

   -Or-

   If you are using smart cards, install a certificate on each smart card distributed to a VPN client user.

Before you can install a certificate, a certification authority must be present and reachable. For a computer in a Windows Server 2003 domain, you can use auto-enrollment or the Certificates snap-in to install a certificate. Alternatively, you can install a certificate by using a Web browser to connect the VPN client to the CA Web enrollment agent. To install a certificate by using a CA Web enrollment agent, perform the following procedure:

**To use the CA Web enrollment tool to install a certificate on a VPN client**

1. Use a Web browser to connect the VPN client to the CA Web enrollment tool at http:**//***ServerName***/**certsrv, where *ServerName* is the name of the server hosting the CA.

2. Click Request a certificate, and then click Advanced Certificate Request.

3. Click **Create and submit a request to this CA** to display a Web form for entering certificate information.

4. Enter the required information on the Web form, and then click **Submit**.

5. Click **Install this certificate**.

For information about:

- Using the Certificates snap-in to install a certificate, see "Using Certificates" in Help and Support Center for Windows Server 2003.

- Using certificate autoenrollment to install a certificate, see "Certificate autoenrollment" in Help and Support Center for Windows Server 2003.

- Deploying smart cards, see "Planning a Smart Card Deployment" in *Designing and Deploying Directory and Security Services* of this kit.

# Configuring Encryption for a VPN Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the remote access policy that governs VPN connections on the server, set the appropriate encryption strengths for PPTP and L2TP/IPSec connections. For a procedure for entering encryption settings in a remote access policy, see "Remote access/VPN server role: Configuring a remote access/VPN server" in Help and Support Center for Windows Server 2003.

For PPTP-based VPN connections, specify one of the following encryption strengths:

- To support MPPE with a 40-bit key, select **Basic**.
- To support MPPE with a 56-bit key, select **Strong**.
- To support MPPE with a 128-bit key, select **Strongest**.

For L2TP/IPSec-based VPN connections, specify one of the following encryption strengths:

- To support 56-bit DES, select either **Basic** or **Strong**.
- To support 3DES encryption, select **Strongest**.

**Note**

- The **No Encryption** level, which allows connections that do not use data encryption, is not recommended.

For more information about using Windows Server 2003 remote access policies, see "Introduction to remote access policies" in Help and Support Center for Windows Server 2003.

# Configuring Remote Access Account Lockout for a VPN Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you will use remote access account lockout to prevent online dictionary attacks, enable remote access account lockout by modifying the AccountLockout subkey in registry on the server that authenticates remote access requests.

If the remote access server is configured for Windows authentication, modify the registry on that server. If the remote access server is configured for RADIUS authentication, and you are using IAS, modify the registry on the IAS server.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

The AccountLockout subkey can be found in the following subkey:

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

The AccountLockout subkey does not exist in the registry until you enable the Routing and Remote Access service or install the Internet Authentication Service.

To configure remote access account lockout, modify two entries in the AccountLockout subkey:

1. To enable account lockout, set **MaxDenials** to 1 or greater.

   **MaxDenials** sets the maximum number of failed attempts that can occur within the configured reset time before the account is locked out. By default, **MaxDenials** is set to zero, which disables account lockout.

2. To change the interval at which the failed attempts counter is reset, set the number of minutes in **ResetTime (mins)**.

   By default, the failed attempts counter is reset every 48 hours (a value of 0xb40, or 2,880 minutes). To modify this interval, enter the preferred number of minutes.

**Note**

- To manually reset a user account that has been locked out before the failed attempts counter is automatically reset, delete the following registry subkey, which corresponds to the user's account name: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout*domainname:username*.

For more information about remote access account lockout, see "Remote Access Technologies" in the *Networking Collection* of the *Windows Server 2003 Technical Reference*.

# Deploying a Dial-up Remote Access Server Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying a dial-up remote access server involves three major tasks: configuring a Windows Server 2003–based server as a dial-up remote access server, configuring the LAN adapter to provide the server with a connection to the intranet, and configuring the appropriate level of encryption strength in the dial-up remote access policy. Figure 8.8 shows the process for deploying a dial-up remote access server solution.

**Figure 8.8 Deploying a Dial-up Remote Access Server Solution**



# Configuring a Dial-up Remote Access Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To provide dial-up access to your organization's intranet, configure a computer running Windows Server 2003 as a dial-up remote access server.

Before configuring the server as a dial-up remote access server, you must enable the Routing and Remote Access service, which is installed automatically with Windows Server 2003. Use the Routing and Remote Access Server Setup Wizard. For instructions on using the wizard, see "Remote access/VPN server role: Configuring a remote access/VPN server " in Help and Support Center for Windows Server 2003.

**Note**

- You can optionally implement Network Access Quarantine Control to quarantine each new remote access connection until the configuration of the client computer can be verified against network policy restrictions. For more information, see "Planning for Network Access Quarantine Control" and "Configuring Network Access Quarantine Control" earlier in this chapter.

With Routing and Remote Access enabled, configure the properties of a dial-up remote access server by using the Routing and Remote Access snap-in.

**To configure a server for dial-up remote access**

1. Open the **Routing and Remote Access** snap-in.

2. In the console tree, right-click the server name, and then click **Properties**.

3. On the **General** tab of the **Server Properties** dialog box, verify that the **Remote access server** check box is selected.

4. On the **Security** tab, set up authentication for dial-up remote access clients:

    a. Click **Authentication Methods**, and in the dialog box select the authentication methods that the server will accept for dial-up connections.

    **Note**

The server is configured by default to accept certain authentication methods. You can use remote access policies to control which authentication methods to accept. For more information about using Windows Server 2003 remote access policies, see "Introduction to remote access policies" in Help and Support Center for Windows Server 2003.

1. Under **Authentication Provider** on the **Security** tab, select the authentication provider to use for dial-up networking clients.

2. Under **Accounting Provider**, select and configure the accounting provider to use for recording dial-up connection accounting information.

1. On the **IP** tab, set up routing for remote access clients:

    a. Verify that the **Enable IP routing** and A**llow IP-based remote access and demand-dial connections** check boxes are selected.

    b. If you are using DHCP to obtain IP addresses for remote access clients, select **Dynamic Host Configuration Protocol (DHCP)**.

    -Or-

Select **Static address pool**, and then configure ranges of IP addresses that are dynamically assigned to dial-up networking clients.

If the static IP address pool consists of ranges of IP addresses for a separate subnet, either enable an IP routing protocol on the remote access server or add static IP routes for each range to your IP routing infrastructure. If the routes are not added, remote access clients cannot receive traffic from resources on the intranet.

# Configuring a Dial-up Connection to the Intranet

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A LAN adapter provides the connection from a dial-up remote access server to the intranet. To enable this connection, you must configure TCP/IP on the LAN adapter and, on the dial-up remote access server, configure the modem ports for remote access.

**Configuring TCP/IP on the LAN Adapter**
Configure the following TCP/IP settings on the LAN adapter that provides the connection from the dial-up remote access server to the intranet:

- The IP address and subnet mask assigned by a network administrator.
- The default gateway of a local router.
- The IP addresses of DNS and WINS servers.

**Configuring a Connection to Dial-up Networking Clients**
To enable multiple dial-up clients to connect to the intranet simultaneously, the dial-up solution must have a modem bank connected to a telecommunications provider. The modem bank adapter includes drivers that you install on the dial-up remote access server.

**Configuring Dial-in Ports for Remote Access**
With the modem bank adapter drivers installed, the modem bank appears as a device with multiple modem ports. Use the Routing and Remote Access snap-in to configure all of the active modem bank ports on the server for remote access.

**To configure the ports of a device for remote access**

1. Open the **Routing and Remote Access** snap-in.
2. In the console tree, right-click **Ports**, and then click **Properties**.
3. In the **Ports Properties** dialog box, select the device that you want to configure, and then click **Configure**.
4. In the **Configure Device** dialog box, select the appropriate routing connection options.

# Configuring Encryption for a Dial-up Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the remote access policy that governs connections for remote access on the dial-up remote access server, use Routing and Remote Access to set the appropriate encryption level. For a procedure for entering encryption settings in a remote access policy, see "Configuring authentication and data encryption" in Help and Support Center for Windows Server 2003.

In the remote access policy for dial-up connections on the dial-up remote access server, choose one of the following encryption levels:

- To use MPPE with a 40-bit key, select **Basic**.
- To use MPPE with a 56-bit key, select **Strong**.
- To use MPPE with a 128-bit key, select **Strongest**.

For more information about using Windows Server 2003 remote access policies, see "Introduction to remote access policies" in Help and Support Center for Windows Server 2003.

# Additional Resources for Deploying Dial-up and VPN Remote Access Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.
- "Deploying Smart Cards" in *Designing and Deploying Directory and Security Services* of this kit.
- "Deploying Internet Authentication Service (IAS)" in this book.
- "Deploying Remote Access Clients Using Connection Manager" in this book.
- The *Internetworking Collection* of the *Windows Server 2003 Technical Reference* (or see the Internetworking Collection on the Web at http://www.microsoft.com/reskit) for more information about Routing and Remote Access.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Remote access/VPN server role: Configuring a remote access/VPN server" in Help and Support Center for Windows Server 2003.
- "Configure RIP for IP" in Help and Support Center for Windows Server 2003 for information about configuring the VPN server as an RIP router.
- "OSPF design considerations" and "Configure OSPF" in Help and Support Center for Windows Server 2003 for information about configuring a VPN server as an OSPF router.
- "OSPF design considerations" in Help and Support Center for Windows Server 2003 for information about configuring a VPN server as an autonomous system boundary router (ASBR).
- "Introduction to remote access policies" in Help and Support Center for Windows Server 2003.
- "Configuring Authentication and Data Encryption" in Help and Support Center for Windows Server 2003 for a procedure for setting the encryption level in a remote access policy.
- "VPN servers and firewall configuration" in Help and Support Center for Windows Server 2003 for procedures explaining how to configure packet filters on the firewall and the VPN server.
- "Configuring the branch office network" in Help and Support Center for Windows Server 2003.
- "Using Certificates" in Help and Support Center for Windows Server 2003 for information about using the Certificates snap-in to install a certificate.
- "Certificate autoenrollment" in Help and Support Center for Windows Server 2003 for information about using certificate autoenrollment to install a certificate.

# Deploying Remote Access Clients Using Connection Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use Connection Manager, a component of the Microsoft® Windows® Server 2003 operating systems, to provide customized remote access to your network through a dial-up or a virtual private network (VPN) connection. By deploying remote access clients with the Connection Manager family of programs, which includes the Connection Manager Administration Kit (CMAK), Connection Point Services (CPS), and the Connection Manager client, you can configure the remote access experience for your users. This configuration can include providing a phone book in which users can find the most convenient dial-up access number.

**In This Chapter**


Overview of Remote Access Client Deployment

Providing Connection Manager Phone Book Support

Customizing Connection Manager

Implementing Your Connection Manager Solution

Example: Deploying Remote Access Clients

Additional Resources for Deploying Remote Access Clients Using Connection Manager


**Related Information**

- For information about designing and deploying a remote access server solution, see "Deploying Dial-up and VPN Remote Access Servers" in this book.
- For information about designing and deploying Internet authentication, see "Deploying Internet Authentication Service (IAS)" in this book.
- For information about creating connections to branch offices and remote sites, see "Connecting Remote Sites" in this book.

**Note**

- For the purposes of this chapter, "phone book" and "Connection Manager phone book" refer to a list of Points of Presence (POPs) and other information configured by using the Connection Manager family of programs.

# Overview of Remote Access Client Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have users who travel frequently or need to access your network from home or other locations, Connection Manager provides a way for you to customize a self-installing service profile for your users. The Connection Manager client allows users to either connect to your network directly or to create a VPN connection from a remote location. Using this managed remote access solution reduces administration by providing a single connection client for all remote access users.

Before you deploy a remote access client solution, you must design and deploy your remote access and VPN servers and related infrastructure. You must also deploy an authentication service, such as Internet Authentication Service (IAS), to enable authentication, authorization, and accounting. Many of the decisions that you make when you deploy your Connection Manager service profiles are based on the decisions that you make when you design your servers.

# Deploying Connection Manager Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process for deploying Connection Manager includes designing phone books if you are configuring dial-up access, and optionally deploying phone book servers to provide phone book updates; customizing the Connection Manager service profile through the CMAKwizard; testing your remote access solution; and distributing your Connection Manager profile to users.

Figure 9.1 shows the process for deploying Connection Manager.

**Figure 9.1 Deploying Connection Manager**



# Remote Access Clients Background Information

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order for users to take advantage of an organization's remote access solution, each client must be configured to connect to the remote access dial-up or VPN servers. You can either use the native connection features in Windows to configure clients or use a managed client solution, such as Connection Manager and its components, to create and distribute a custom service profile.

# Native Connection Capabilities and Limitations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is possible for users to manually configure remote access connections using the native network connection capabilities in Windows. To connect to the remote access server using these native capabilities, the user configures the network settings on the client. These settings include:

- **Dial-up connections** The telephone number for your remote access server, user authentication method, encryption settings, and dialing scripts.

- **VPN connections** The host name or IP address for the VPN server, VPN type, user authentication method, and encryption settings.

The native connection capabilities are best suited for when there are few users connecting to the network. These connections are relatively simple to set up when there are a small number of clients; however, there are major disadvantages to this method when you are administering a large network with many remote access users, including:

- The procedure for manually configuring remote access clients varies between versions of Windows; therefore, you would need a separate set of procedures for each client operating system you support.

- Each client must be manually configured; either an administrator must configure each client individually, or the users must configure their own settings using operating system–specific instructions. Either approach can lead to a large resource drain in the IT department.

- If any telephone numbers change, either the administrator or the user must manually reconfigure the connection. For example, you might contract with a telecommunications supplier to provide multiple dial-up telephone numbers and worldwide access for users who travel. Similarly, you might choose to use VPN connections over Internet connections supplied by an Internet service provider (ISP), with multiple access numbers and worldwide Internet access. If any of these telephone numbers change, you need a way to notify the users.

Connection Manager provides a solution for these and other issues when you deploy a large number of remote access clients.

# Connection Manager Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Connection Manager family of programs is a set of optional components used to create a managed remote access solution. Connection Manager enables a network administrator to preconfigure remote access clients, add custom behavior and a custom appearance, and provide an updateable phone book that enables users to find the most convenient dial-up access number. The Connection Manager family of products includes:

- The Connection Manager client

  The Connection Manager client provides a simplified way of connecting to a remote network. Typically, the user only needs to enter a user name and password and select a phone number if applicable. The administrator configures all other settings before distributing the service profile.

- The Connection Manager Administration Kit (CMAK)

  CMAK allows the administrator to create and configure the service profile and creates a small, self-installing package. CMAK also allows the administrator to customize Connection Manager features such as branding, custom actions, and custom Help files, as well as enhanced security features.

- Connection Point Services (CPS)

  CPS allows you to create and maintain phone books. It consists of two parts:

  - Phone Book Administrator (PBA)

    PBA is a tool used to create and maintain phone book files, and to publish new or updated phone book files on the PBS server.

  - Phone Book Service (PBS)

    PBS distributes phone books to Connection Manager clients on request.

# Connection Methods

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remote users connect to networks by using one of two methods: they either connect with *direct dial*, where they connect directly by using dial-up lines, or they use VPNs to connect over the Internet. When using a VPN to connect, remote users who do not have a pre-existing connection to the Internet must use a *double-dial* configuration, where they first dial an ISP number to access the Internet and then establish the VPN connection. Connection Manager can make this double dial process look like a single connection attempt to the end user.

**Direct Dial**

Users who connect to your network by using direct dial call directly into your network, using the dial-up phone numbers that your organization provides to connect to remote access dial-up servers. You can easily manage a small number of users calling a small number of phone numbers. However, if a large number of users are dialing into your network, or if your network can be reached through many phone numbers, Connection Manager and CPS are useful for managing remote access.

**VPN**

Organizations that offer VPN access to their remote users approach this in either one of two ways:

- Assume that the users have their own connections to the Internet.
- Provide users with an easy method to dial up the Internet and establish a subsequent VPN connection to the corporate network.

An organization can also contract with an ISP to supply a national or worldwide collection of phone numbers for Internet access. Connection Manager provides a method to expose these numbers from the ISP in a phone book and automatically establish the VPN connection after the ISP connection is complete. For more information about working with an ISP, see "Providing Connection Manager Phone Book Support" in this chapter. For more information about double-dial connections, see "Example: Deploying Remote Access Clients" in this chapter.

# User Authentication Methods

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The user authentication method that you implement depends on the operating systems that your clients are running and the level of security that you require for your network. For example, you might require passwords, certificates, or smart cards for user authentication, depending on your organization's security needs. For more information about user authentication methods, see "Designing an Authentication Strategy" in *Designing and Deploying Directory and Security Services* of this kit. For more information about deploying smart cards, see "Deploying Smart Cards" in *Designing and Deploying Directory and Security Services* of this kit.

# Virtual Private Network (VPN) Tunneling Protocols

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For VPN connections, you can require Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP) connections, or you can allow users to attempt L2TP connections and use PPTP if they cannot connect using L2TP. If you require L2TP/IPSec authentication, you can use preshared keys or certificates. For more information about choosing a VPN strategy, see "Deploying Dial-up and VPN Remote Access Servers" in this book.

# Network Access Quarantine Control

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Network Access Quarantine Control, a new feature in the Windows Server 2003 family, delays normal remote access to a private network until the configuration of the remote access computer has been examined and validated by an administrator-provided script. Without Network Access Quarantine Control, only the credentials of the user are verified, and a user with the correct credentials can connect even if their configurations do not comply with corporate network policy. For example, a remote access user with valid credentials can connect to a network with a computer that does not have required antivirus software installed on it. Remote access client can use either a manually-configured connection or a Connection Manager profile. For more information about configuring Network Access Quarantine Control by using a Connection Manager profile, see "Incorporate Custom Actions" later in this chapter.

**Important**

- Network Access Quarantine Control allows an administrator to prevent computers with unsafe or undesirable configurations from connecting to a private network, not to protect a private network from malicious users who have obtained a valid set of credentials.

For more information about Network Access Quarantine Control, see "Deploying Dial-up and VPN Remote Access Servers" and "Deploying Internet Authentication Service (IAS)" in this book, and "IAS Network Access Quarantine Control" in Help and Support Center for Windows Server 2003.

# Providing Connection Manager Phone Book Support

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use CPS to create, publish, and update Connection Manager phone books, which include all information relating to a POP, such as area code, phone number, and user authentication methods supported. The Connection Manager phone book also includes various network settings that you configure when you run the CMAK wizard.

Figure 9.2 shows the process for providing phone book support.

**Figure 9.2 Providing Connection Manager Phone Book Support**




For more information about security considerations when using CPS, see "Security information for Connection Point Services" in Help and Support Center for Windows Server 2003.

# Using PBA to Create and Maintain Phone Books

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

PBA in CPS allows you to create and maintain phone books for use with Connection Manager. By using CPS, you can ensure your remote access users have the most recent phone numbers to connect to your network.

**Note**

- If you are providing a VPN-only profile you do not need to create a phone book and can skip to "Customizing Connection Manager" later in this chapter.

Before using PBA, you must install it on a computer running either Windows Server 2003 or Microsoft® Windows® XP Professional. You install PBA by running Pbainst.exe from the VALUEADD\MSFT\MGMT\PBA folder on the Windows Server 2003 family CD-ROM or Windows XP Professional CD-ROM.

For information about how to administer phone books from the command line, see "Administer phone books" in Help and Support Center for Windows Server 2003.

**Creating, Publishing, and Maintaining Phone Books**

A phone book contains one or more Points of Presence (POPs), each of which includes a telephone number to access the network or the Internet. Connection Manager phone books allow the user to have a complete POP list, so they have more than one phone number to choose from when traveling or if there is a problem connecting to a specific number.

Before you begin adding POPs to the phone book, you can create a list of regions to organize the POPs. *Regions* allow the user to see a filtered list within their country or dependency. This reduces the number of POPs the Connection Manager client displays at one time, so the user chooses from a list of POPs targeted to their region. Use the **Region Editor** in PBA to create a new list of regions, or import regions from an existing .pbr file.

After you have gathered all the POPs contained in the phone book, use PBA to create your new phone book.

When you create a new phone book, PBA creates a phone book file (.pbk) and a region file (.pbr). When you finish creating the phone book, you need to publish it so that it is available on the PBS server.

**To publish the phone book**

1. Start the FTP service on the PBS server.
2. Enable the account that was created on the PBS server for phone book posting.
3. Start PBA, click the **Tools** menu, and then click **Publish Phone Book**.
4. In the **Options** dialog box, type the server address, user name, and password used to post the phone book to the server running PBS.
5. Click **Create** to create the .cab files in the release directory.
6. Enter the address of the PBS server and click **Post** to post the new release to the PBS server.
7. Disable the account on the PBS server.
8. Stop the FTP service on the PBS server.

Use PBA to edit your phone book when adding or changing telephone numbers, and then use the same procedure to re-publish the updated phone book.

**Updating Phone Books**

To ensure the client always has the most recent version of the phone book, leave the **Automatically download phone book updates** check box selected on the **Phone Book** page of the CMAK wizard.

If you have configured automatic phone book updates, after the client connects to the network, the client sends a request to PBS that includes the name of the phone book file and the version currently installed on the client. The server then replies with either a full or an incremental update, if one is required.

# Using an Outsourced Phone Book

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To provide a wider range of phone numbers, your organization might outsource your phone book to a telecommunications company or an ISP. If your telecommunications company or ISP uses CPS, you can import that company's phone book.

When negotiating with the ISP or telecommunications company, you should ask that company to provide a service profile that contains all its phone book information, including an update URL, and then use CMAK to merge this profile into your service profile.

If the company does not provide you with a service profile, request the phone book .(pbk) and region .(pbr) files for the phone book, as well as the update URL. Use these files to create a service profile, and merge it into the service profile you will send to your users.

If your telecommunications company or ISP does not use CPS, obtain the phone numbers in another format, such as a text file, and import this file into your phone book. For more information about importing phone book files using PBS, see "Add, edit, or delete POPs by command line" in Help and Support Center for Windows Server 2003.

For more information about merging service profiles, see "Merging Service Profiles" later in this chapter.

# Hosting Phone Books on a PBS Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You do not need to run PBS on a high-end computer because any computer capable of running Windows Server 2003 and Internet Information Services (IIS) can run PBS.

Just before publishing, ensure that the PBS server is using IIS to run FTP and the Web server service. This is necessary for PBS servers to receive posts from PBA and to provide updated files to clients. For more information about IIS, "Internet Information Services (IIS) 6.0 overview" in Help and Support Center for Windows Server 2003.

**Caution**

- For security reasons, do not allow anonymous FTP access to your PBS server. Also, only run the FTP service when you are actually publishing or updating a phone book. For more information about CPS security, see "Security information for Connection Point Services" in Help and Support Center for Windows Server 2003.

PBS servers can be located on your perimeter network or in the perimeter network of the ISP if you outsource phone book support. Because the phone book updates after connecting to your network, locate the PBS server such that the client has access to it after a successful connection to your network. For more information about security in perimeter networks, see "Securing Network Perimeters with ISA Server in this book.

# Customizing Connection Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the CMAK wizard to create a custom service profile. You then distribute this .exe file to your users. When they double-click the file, it installs a service profile customized with the information you entered in the CMAK wizard. The Connection Manager client then allows users to dial in to your organization directly or to complete a VPN connection, based on the information provided in the CMAK wizard. Figure 9.3 shows steps in the process for customizing your service profile.

**Tip**

- You can install the CMAK wizard from **Management and Monitoring Tools** details in the Windows Components Wizard. For more information about installing the CMAK wizard, see "The Connection Manager Administration Kit Wizard" in Help and Support Center for Windows Server 2003.

**Figure 9.3 Customizing Connection Manager**




Before you run the CMAK wizard, make sure you know the following information, which is required to complete the wizard:

- The service name and a file name that you will use for the new profile and related files.
- A realm name, if your service requires it. A *realm name* is a prefix or suffix that Connection Manager automatically adds to the user name.
- Any existing service profiles that you plan to merge into the new profile. For more information about merging service profiles, see "Merging Service Profiles" later in this chapter.
- VPN Support information, including: the VPN server address(s) and whether or not the client will use the same passwords for the dial-up and VPN connection in a double-dial situation. For an example of configuring VPN Support information, see "Example: Deploying Remote Access Clients" later in this chapter. For information about configuring a VPN-only profile, see "Implementing VPN support" in Help and Support center for Windows Server 2003.
- VPN Entries network and security information. For more information about networks and security, see "Configuring Network and Security Settings" later in this chapter.
- Preshared key information, if needed for L2TP/IPSec VPN connections. If you are using a preshared key, encrypt it with a PIN that follows strong password rules. *Strong passwords* include a combination of uppercase and lowercase letters, numbers, and special characters so the password is protected from a dictionary attack or a database of popular passwords.
- The location of the phone book file that was created by PBA to include in this service profile and any text that should appear in the **More access numbers** box in the **Phone Book** dialog box.

- Phone book file name for downloading updates and the update URL to point to on the phone book server, if applicable.
- Dial-up Networking Entries and security information. For more information about networks and security, see "Configuring Network and Security Settings" later in this chapter.
- Routing table update information, if you are planning to implement *split-tunneling* where users can connect to both your internal network and the Internet simultaneously.
- Automatic Proxy Configuration settings, if you want Connection Manager to automatically update proxy settings for this connection.
- Custom Actions, which are any programs you want to start automatically before, during, or after users connect to your service. For more information about custom actions, see "Incorporate Custom Actions" later in this chapter.
- Branding information, including custom graphics, icons, menu items for the notification area shortcut, custom Help, and support information, if applicable. For more information about including branding information in your service profile, see "Branding Your Connection Manager Client" later in this chapter.
- Whether to include the latest version of Connection Manager with your service profile. This is a small file, so if you are not sure that all clients have the latest version of CM, include the latest version with your profile.
- A custom license agreement, if applicable. For more information about including a custom license agreement in your service profile, see "Branding Your Connection Manager Client" later in this chapter.
- Any additional files you want to include in this service profile.
- Any information you require for advanced customization, if applicable. For more information about advanced customization, see "Providing Advanced Customization" later in this chapter.

For a worksheet to assist you in completing the CMAK wizard, see "Preparation for Running the CMAK Wizard" (DNSRAC_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Preparation for Running the CMAK Wizard" on the Web at http://www.microsoft.com/reskit).

For more specifics about customizing Connection Manager using CMAK, see "Preparing to run the CMAK wizard" and "Connection Manager Administration Kit" in Help and Support Center for Windows Server 2003.

# Merging Service Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Merging service profiles is especially useful for incorporating information from multiple phone books, including service types, POP names, and access numbers. You can also use it to consolidate different dial-up access numbers that are covered by more than one of your phone books. By merging existing service profiles into a single top-level service profile, you can present several dissimilar networks as a single, cohesive service.

You can merge multiple profiles so that the *top-level profile*, which is the service profile you distribute to users, behaves as if it has a single phone book, which contains all the POPs defined in each of the component profiles. For example, if you outsource your phone book through multiple ISPs, and each ISP provides you with a service profile containing its phone book(s), you can merge the profiles together within a top-level service profile that contains dial-up numbers for all the POPs from all the providers.

For an example of merging service profiles, see "Example: Deploying Remote Access Clients" later in this chapter.

# Configuring Network and Security Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure network and security settings by editing the appropriate networking entry from the **VPN Entries** page and/or the **Dial-Up Networking Entries** page of the CMAK wizard. From the **Edit VPN Entry** and the **Edit Dial-Up Networking Entry** dialog boxes, you can customize general network properties, TCP/IP settings, and security settings.

**Configuring General Network Properties**
Use the general network properties to disable file and print sharing and enable clients to log on to a network. For a dial-up connection, you can also enter a dial-up networking script.

**Configuring TCP/IP Settings**
Use the TCP/IP settings to change the Domain Name System (DNS) and Windows Internet Name Service (WINS) client configuration by either allowing the server to assign these addresses or manually configuring the DNS and WINS addresses for this profile. You can also choose to make this connection the default gateway for the client and to use IP header compression.

**Configuring Security Settings**
Use the security settings to select the authentication method for VPN or dial-up users and the VPN strategy for VPN clients. You can choose to use basic or advanced security settings for all computers or use a combination. If you choose a combination of both basic and advanced security settings, Connection Manager uses the advanced security settings for clients running on the operating systems that support them and basic security settings for clients running on operating systems that do not support advanced security settings.

# Incorporate Custom Actions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Connection Manager has the ability to run custom actions at various points when establishing a connection. By providing custom actions, you can enhance the connection experience for your users. Use the CMAK wizard to include custom actions in your service profile to automatically start programs when users connect to your service. A *custom action* can be any batch file, executable file, or dynamic-link library (DLL). These custom actions can use programs that users have installed, or you can distribute the programs with your service profile.

Using the CMAK wizard, you can specify custom actions for each of the following points during the connection process.

- **Pre-initialization actions.** These actions occur immediately when the user starts the Connection Manager client.
- **Pre-connect actions.** These actions occur before the connection attempt.
- **Pre-dial actions.** These actions occur before every dialing attempt, including redials. (For dial-up connections only.)
- **Pre-tunnel actions.** These actions occur before tunneling. (For VPN connections only.)
- **Post-connect actions.** These actions occur immediately after the connection is established.
- **Disconnect actions.** These actions occur immediately after the user or server disconnects.
- **On cancel actions.** These actions occur whenever the user abandons a connection attempt.
- **On error actions.** These actions occur whenever the connections attempt fails due to an error.

You might want to use pre-connect actions to start an application before you connect, such as an e-mail program, or use a post-connect action to upload logs of connection activity or to download the latest virus signatures. An on error action could also be used to point the user to custom Help files for self-help information, potentially reducing help desk calls.

Several common custom actions are built into CMAK, such as:

- A post-connect action checks for phone book updates. This action is automatically included in your profile if you leave the **Automatically download phone book updates** check box selected on the **Phone Book** page of the CMAK wizard.
- A post-connect action obtains and installs routing tables for the target network. This action is automatically included in your profile if you enable the **Routing Table Update** feature.
- A post-connect action updates proxy settings of the client during the connection. This action is automatically included in your profile if you enable the **Automatic Proxy Configuration** feature.

The *Windows Server 2003 Technical Reference* also contains custom actions you can use to customize your profile:

**Profile update**
This includes the files Getcm.exe, which runs as a post-connect action that checks for and downloads an updated service profile, and Instcm.exe, which runs as a disconnect action that checks to see if an updated service profile has been downloaded and installs it.

**Certificate deployment**
This DLL (Cmgetcer.dll) allows Connection Manager to automatically obtain a certificate for L2TP/IPSec connections.

**Network Access Quarantine Control**

- This network policy requirements script runs as a post-connect action. The network policy requirements script performs validation checks on the remote access client computer to verify that it conforms to network policies. The script can be a

custom executable file or simple batch file.

When the script has run successfully and the connecting computer has satisfied all of the network policy requirements (as verified by the script), the script executes a notifier component (an executable) with the appropriate parameters. You can also configure your script to download the latest version of the script from a quarantine resource. If the script does not run successfully, it directs the remote access user to a quarantine resource such as an internal Web page, which describes how to install the components that are required for network policy compliance.

The notifier component sends a message to the quarantine-compatible remote access server that indicates a successful execution of the script. You can use your own notifier component or you can use Rqc.exe, which is provided on the *Windows Server 2003 Deployment Kit* companion CD. With these components installed, the remote access client computer uses the Connection Manager profile to perform its own network policy requirements check and indicate its success to the remote access server as part of the connection setup.

**Tip**

- Because Network Access Quarantine Control introduces a delay in obtaining normal remote access, applications that run immediately after the connection is complete might encounter problems. One way to minimize the delay is to separate your script into two scripts: one that runs as a pre-connect action and one that runs as a post-connect action.

For more information about Network Access Quarantine Control, see "IAS Network Access Quarantine Control" in Help and Support Center for Windows Server 2003, "Deploying Dial-up and VPN Remote Access Server" and "Deploying Internet Authentication Service (IAS)" in this book. For a sample notifier component, see the Windows SDK. For more information about the Windows SDK, see the Software Development Kit (SDK) information in the MSDN Library link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

In addition to these predefined custom actions, you can create your own custom actions to include in the service profile.

For security reasons, custom actions cannot be run when users log on to the network using dial-up networking unless certain registry keys have been set. For more information about custom actions, see "Incorporating custom actions" in Help and Support Center for Windows Server 2003.

# Branding Your Connection Manager Client

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Another important feature of CMAK is the ability to apply your own branding to the client by using your organization's graphics, icons, menu items for the notification area shortcut, Help, and license agreement. If you do not want to customize the appearance of the client, accept the defaults provided in CMAK.

**Using Custom Graphics and Icons**
You must specify whether to use the default bitmaps or your own graphics. Customize the user interface by including your company logo or another image that identifies your organization. You can replace the graphics in both the **Logon** dialog box and the **Phone Book** dialog box. The replacement graphic must be a bitmap (.bmp) file.

Provide custom icons for your service profile in the form of either one file containing icons in multiple sizes or one file for each icon in each size. Use the CMAK wizard to specify the icons.

**Customizing Connection Manager Help and License Agreement**
To offer your users customized Help and license agreement files, enter the appropriate file information in the CMAK wizard.

The default Help file is Cmmgr32.hlp. If you do not want to use the default Help, create your own .hlp file, and then use the CMAK wizard to replace the default Help with your .hlp file in the service profile. To include a license agreement, create a .txt file containing the agreement and enter the file name into the CMAK wizard when prompted. For more information about providing your users with custom Help, see "Providing custom Windows Help" in Help and Support Center for Windows Server 2003.

# Providing Advanced Customization

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The CMAK wizard guides you through most of the customization features that you need to build a service profile. However, you can provide additional customization by selecting the **AdvancedCustomization** check box in the **Ready to Build the Service Profile** page of the CMAK wizard.

You can also delete any section or key from the .cms or .cmp files by using the **Advanced Customization** page of the CMAK wizard.

**To delete a section**

1. On the **Advanced Customization** page of the CMAK wizard, in the **File name** box, select the appropriate file.
2. In the **Section name** box, select the section you want to delete.
3. Clear the **Key name** and **Value** boxes.
4. Click **Apply**.
5. Click **Yes** to confirm that you want to delete the entire section.

**To delete a key**

1. On the **Advanced Customization** page of the CMAK wizard, in the **File name** box, select the appropriate file.
2. In the **Section name** box, select the appropriate section.
3. In the **Key name** box, select the key you want to delete.
4. Clear the **Value** box.
5. Click **Apply**.
6. Click **Yes** to confirm that you want to delete the key.

**Caution**

- Use extreme care when deleting sections or keys from the .cms or .cmp files by using the **Advanced Customization** page of the CMAK wizard, particularly when you are editing an existing service profile.

For more information, including a comprehensive list of the service profile files and keys that you can customize through advanced customization, see "Advanced Customization" in Help and Support Center for Windows Server 2003.

The following procedures show four ways to use advanced customization to increase security for user connections.

**Tip**

- If the key names you want to customize do not appear in the drop down list, simply type them in the **Key name** text box.

**Removing the Save Password Option**

Edit the HideRememberPassword key to remove the **Save Password** check box from the Connection Manager user interface.

**To remove the Save Passwordcheck box from the Connection Manager user interface**

1. On the **Advanced Customization** page of the CMAK wizard, in the **File name** box, select *FileName***.cms**.

2. In the **Section name** box, select **Connection Manager**.

3. In the **Key name** box, type or select **HideRememberPassword**.

4. In the **Value** box, type **1**.

5. Click **Apply**.

   **Important**

   - In order for any advanced customization setting to be recorded, you must click **Apply** after entering each setting.

**Disabling ICS**
Edit the DisableICS key to disable Internet Connection Sharing (ICS) in Windows XP for this connection.

**To disable Internet Connection Sharing (ICS) from the Connection Manager user interface**

1. On the **Advanced Customization** page of the CMAK wizard, in the **File name** box, select *FileName***.cms**.

2. In the **Section name** box, select **Connection Manager**.

3. In the **Key name** box, type or select **DisableICS**.

4. In the **Value** box, type **1**.

5. Click **Apply**.

**Enabling ICF**
Edit the EnableICF key to turn on Internet Connection Firewall (ICF) in Windows XP for this connection.

**To enable ICF from the Connection Manager user interface**

1. On the **Advanced Customization** page of the CMAK wizard, in the **File name** box, select *FileName***.cms**.

2. In the **Section name** box, select **Connection Manager**.

3. In the **Key name** box, type or select **EnableICF**.

4. In the **Value** box, type **1**.

5. Click **Apply**.

**Hiding the Advanced Tab**
Edit the HideAdvancedTab key to hide the **Advanced** tab from the Connection Manager user interface in Windows XP for this connection. The **Advanced** tab is where users control ICF and ICS for the connection. You should only enable this key if you are using the DisableICS and EnableICF keys.

**To hide the Advanced tab from the Connection Manager user interface**

1. From the **Advanced Customization** page of the CMAK wizard, in the **File name** box, select *FileName***.cms**.

2. In the **Section name** box, select **Connection Manager**.

3. In the **Key name** box, type or select **HideAdvancedTab**.

4. In the **Value** box, type **1**.

5. Click **Apply**.

# Implementing Your Connection Manager Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you create your service profile(s), test your remote access solution in its entirety before distributing your service profile to the users.

Figure 9.4 shows the process for implementing your managed remote access client solution using Connection Manager.

**Figure 9.4 Implementing Your Connection Manager Solution**



# Testing Your Remote Access Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When testing your solution in a lab, recreate the actual user experience as closely as possible. It is likely that the client will be deployed on computers not directly under the control of your organization, such as a user's home computer. For this reason, it is necessary to test your service profile(s) using a standard set of applications and test on the same types of hardware that you will deploy to.

**Important**

- Test both the server and client portions of your remote access design.

Use the following guidelines to test your deployment:

- Load and run the client on each operating system you are supporting for remote access. If you are supporting both dial-up and VPN connections, test both of these types of connections on your test clients.
- Test custom actions in detail, using a standard set of application that users might have on their computers. Be sure your test is representative of the end-user experience.
- Make sure your phone books are updating. If phone books are not updating, check the URL that is pointing to the phone book server, check for firewall conflicts, and ensure that IIS is correctly configured on PBS server.
- Test your distribution method before announcing and rolling it out. For information about establishing a distribution method, see "Distributing Your Connection Manager Service Profiles" later in this chapter.

# Distributing Certificates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you configured the security settings of the VPN Entries in the CMAK wizard to use L2TP/IPSec, you might need to distribute certificates to your users. The certification authority (CA) is generally set up as a Web server. You can either have your CA on the Internet or on your intranet.

**Internet Enrollment**
With Internet enrollment, users go to a public Web site to obtain their certificates. Internet enrollment is useful if you are using a CA that is provided by another company.

**Intranet Enrollment**
If certificates are optional but recommended, users can obtain their certificates after connecting to your intranet. Configure the service profile to attempt authentication by using L2TP first. This setting allows the client to attempt a connection using L2TP; if L2TP is not available, the client connects using PPTP. When you configure this setting, the client will first attempt to connect using L2TP each time the client connects. By using this setting, clients can connect the first time by using PPTP and get a certificate. After receiving the certificate, subsequent connections will use L2TP.

You can configure the Connection Manager Certificate Deployment Tool, Cmgetcer.dll, as a custom action. This tool enables the client to get a certificate from the certification authority.

For more information about certification authorities and certificates, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

# Educating Users About Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When distributing your service profiles, you should also inform users of their responsibilities in protecting the organization's resources. Educate your users about potential threats and how to avoid them, including:

- Enable a personal firewall (such as ICF in Windows XP).
- Use strong passwords on their remote computers.
- Never save passwords for any connection.

  To prevent users from saving their password for this connection, disable the **Save Password** check box on the Connection Manager client. For more information, see "Providing Advanced Customization" earlier in this chapter.
- Lock their computers when they are not actively using them. They do this by password-protecting the screen saver or through the **Ctrl-Alt-Delete** dialog box.
- Do not share VPN connections or run a VPN connection from an ICS host. Sharing the VPN connection allows all computers on the ICS network — using the VPN connection — to access your organization's network and resources using the VPN connection's credentials.

# Distributing Your Connection Manager Service Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are several ways to distribute your service profile, each with costs and benefits. Choose one of the following methods, or provide more than one method to give your users a choice.

**Distributing Service Profiles on CD or Floppy Disk**
You can distribute CDs or floppy disks containing your self-installing Connection Manager package. Connection Manager and the service profile fit on a floppy disk. However if you want to include other programs, such as anti-virus software, you might need more space than a floppy provides so a CD is a better choice.

The benefit of distributing this way is that you can physically give a copy to all users or send them easily through the mail. However, this solution might be costly and has little inherent security.

**Distributing Service Profiles by E-mail**
You can send a service profile through e-mail to your users. If you choose to send the service profile through e-mail, ensure that users are able to receive .exe files, because not all e-mail systems allow executable files as attachments.

**Distributing Service Profiles by Download**
You can set up a Web site where users can download the service profile. Desktop users can download to a floppy disk, and portable-computer users can download directly to their computers from a Web site inside your network.

It is also possible to make the service profile available by download from a Web site over the Internet. However, identify any security risks to your organization before posting your service profile on an Internet site.

**Pre-installing Service Profiles**
You can install the service profile on each client individually. The benefit of this method is that users are not required to install anything themselves, which can reduce user frustration and calls to your help desk. However, this method requires administrator or help desk resources during the initial installation, which might be a large resource hit during the roll out phase of your deployment. This method is useful when there are a small number of client computers or when all of the client computers and devices are controlled by your organization.

**Combining Distribution Methods**
You can also use a combination of distribution methods. For example, a company could distribute the Connection Manager service profiles on CD to users who work from their own computers from remote locations, provide downloads for local employees who have portable computers, and pre-install the service profile on any new portable computers before distribution.

# Example: Deploying Remote Access Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A large company, Contoso Ltd., is redesigning remote access infrastructure. Contoso decides to use the Connection Manager family of products to provide managed remote access to their company network through both dial-up connections and VPN connections. The new VPN server allows both PPTP and L2TP/IPSec VPN connections.

Contoso contracts with an ISP, A. Datum Corporation, to provide bulk dial-up Internet access. Under the arrangement, A. Datum will provide single sign-on for users using their Contoso credentials. This is accomplished by using a Remote Authentication Dial-In User Service (RADIUS) proxy to forward connection requests to a Contoso RADIUS server using a realm name agreed to by both A. Datum and Contoso. For more information about deploying a RADIUS proxy and RADIUS server, see "Deploying Internet Authentication Service (IAS)" in this book.

Contoso has the following primary objectives for providing a remote access solution to its users:

- Allow local users to connect to the corporate intranet with direct dial using a local phone number, which dials directly into Contoso's remote access servers.
- Reduce costs by eliminating the need for toll free (1-800) dial-up access numbers for users traveling within the United States. The company has a contract with A. Datum to provide dial-up access numbers to the Internet, which will be used to carry VPN connections to the company.
- Provide all users with automatic phone book updates when dial-up access numbers change.

  Allow users to connect by making a VPN connection to the corporate Intranet over their existing connections to the Internet, such as digital subscriber line (DSL) and cable modem connections.

The company would also like to improve the connection experience for their users in the following ways:

- Provide a simplified method of setting up all types of connections on a variety of Windows operating systems.
- Provide a unified phone book for all access numbers.
- Provide a customized user interface for the connection client, including custom icons and graphics.
- Provide a single sign-on experience for double-dial VPN users by using a realm name and a RADIUS proxy.

The Connection Manager family of products provides Contoso with solutions to meet all of these goals.

# Contoso Prepares Phone Books

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before Contoso creates the Connection Manager service profile, they create the local phone book file (Contoso.pbk) and the region file (Contoso.pbr). These files contain the local phone numbers that allow users to dial directly into Contoso's corporate intranet.

The company also receives a phone book file (Adatum.pbk), a region file (Adatum.pbr), and an update URL (http://pbupdate.adatum.com) from A. Datum, the ISP. To incorporate the phone numbers from the ISP, Contoso creates a component service profile that it will merge into a top-level profile.

The company creates a phone book that includes all the direct-dial numbers for users to the company. Contoso performs the following steps to create the phone book:

1. Installs CPS and PBA. For more information about installing and running PBA, see "Providing Connection Manager Phone Book Support" earlier in this chapter.
2. Runs PBA.
3. Creates a new phone book named Contoso.
4. Uses the **Region Editor** to enter the regions for these numbers.
5. Adds POP entries and enters the information for each phone number.
6. Publishes the phone book.

To provide phone book updates, Contoso installs PBS on a computer running a member of the Windows Server 2003 family. The phone books from the ISP are already in the form of a phone book file (Adatum.pbk) and a region file (Adatum.pbr), so Contoso does not have to create any additional phone book files.

# Contoso Creates Service Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso uses the "Preparation for Running the CMAK Wizard" worksheet to collect all the information necessary to run the CMAK wizard. This information is used to create a component profile, which includes the phone numbers and the update URL provided by the ISP, and the top-level profile, which is the service profile the company will distribute to its users.

# Preparing to Create the Component Profile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A. Datum supplies Contoso with a phone book file (Adatum.pbk) and a region file (Adatum.pbr), as well as a phone book update URL. Contoso uses this information to create a component profile. Contoso also includes the realm name, @contoso.com, agreed to by both Contoso and A. Datum. The realm name identifies the Contoso user and allows A. Datum to forward the RADIUS messages to the RADIUS server for Contoso, as shown in Figure 9.5.

**Figure 9.5 Double-Dial Using RADIUS Proxy**




Figure 9.6 shows the worksheet that Contoso uses to create the component profile.

**Important**

- Figure 9.6 only shows the portions of the "Preparation for Running the CMAK Wizard" worksheet that Contoso customizes to create the component profile. To see the entire blank worksheet to help prepare to run the CMAK wizard, see "Preparation for Running the CMAK Wizard" (DNSRAC__1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Preparation for Running the CMAK Wizard" on the Web at http://www.microsoft.com/reskit).

**Figure 9.6 Preparation for Running the CMAK Wizard: Component Profile**

| Include? (Y/N) | Implementation Information | | | | | Ready? (Y/N) |
|---|---|---|---|---|---|---|
| Y | Service Profile Source *(Choose one)* | | | | | Y |
| | | X | Create a new service profile | | | |
| | | | Edit an existing service profile | | | |
| | | | Existing profile name: | | | |
| Y | Service Name<br>*Use up to 35 characters, including spaces. Do not use any of the following: *\|/:<>\|[]* | | | A. Datum phone numbers | | Y |
| | File name<br>*Enter the name of the files that CMAK creates. Use up to 8 alphanumeric characters.* | | | adatum | | |
| Y | Realm Name<br>*Specify a realm name if your service requires users to enter one and you want Connection Manager to add it automatically .* | | | | | Y |
| | *Choose one:* | | | | | |
| | | | Do not add a realm name to the user name | | | |
| | | X | Add a realm name to the user name | | | |
| | | | *Choose one:* | Prefix | X Suffix | |
| | | | Realm name *(Use the line below and include separator character):* | | | |
| | | | @Contoso.com | | | |
| Y | Phone Book<br>This .pbk file must have a corresponding .pbr region file. This name cannot be longer than 8 characters and cannot contain a file extension, spaces, or any of the following characters:!,;*=\/:?"<> | | | | | Y |
| | Phone book file: | | Adatum.pbk | | | |
| | Check this box if you want Connection Manager to automatically download phone book updates. | | | | | |
| | X | Automatically download phone book updates. | | | | |
| Y | Phone Book Updates<br>If you checked "Automatically download phone book updates," type the URL for the CPS server. | | | | | Y |
| | Connection Point Services server: | | http://pbupdate.adatum.com | | | |

# Preparing to Create the Top-Level Profile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso then creates its top-level profile, merging the component profile into the top-level profile, which is the service profile that Contoso will distribute to its users. Additional information about files included in this profile, such as CMProxy.txt and CMRoute.txt, is discussed in the following section.

Figure 9.7 through Figure 9.13 show the worksheet that Contoso uses to create the component profile and complete the CMAK wizard. Additional information about files included in this profile, such as CMProxy.txt and CMRoute.txt, is discussed in the following section.

**Important**

- Figure 9.7 through Figure 9.13 only show the portions of the "Preparation for Running the CMAK Wizard" worksheet that Contoso customizes to create the top-level profile. To see the entire blank worksheet to help you prepare to run the CMAK wizard, see "Preparation for Running the CMAK Wizard" (DNSRAC_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Preparation for Running the CMAK Wizard" on the Web at http://www.microsoft.com/reskit).

**Figure 9.7 Preparation for Running the CMAK Wizard: Top-Level Profile (page 1)**

| Include? (Y/N) | Implementation Information | | Ready? (Y/N) |
|---|---|---|---|
| Y | Service Profile Source *(Choose one)* | | Y |
| | [X] Create a new service profile | | |
| | [ ] Edit an existing service profile | | |
| | Existing profile name: | | |
| Y | Service Name<br>*Use up to 35 characters, including spaces. Do not use any of the following: *\|/:<>\|[]* | Remote access to Contoso | Y |
| | File name<br>*Enter the name of the files that CMAK creates. Use up to 8 alphanumeric characters.* | Contoso | |
| N | Realm Name<br>*Specify a realm name if your service requires users to enter one and you want Connection Manager to add it automatically.* | | Y |
| | *Choose one:* | | |
| | [X] Do not add a realm name to the user name | | |
| | [ ] Add a realm name to the user name | | |
| | *Choose one:* [ ] Prefix [ ] Suffix | | |
| | Realm name *(Use the line below and include separator character):* | | |
| | | | |
| Y | Merging Profile Information<br>*List any existing profiles to merge into the profile you are creating.* | | Y |
| | adatum | | |
| | | | |
| Y | VPN Support<br>Establish a VPN connection when the user dials a phone number from these phone books. *Check all that apply.* | | Y |
| | [ ] Phone books from this profile | | |
| | [X] Phonebooks from the merged profiles | | |
| Y | VPN Support: VPN Server name or IP address *(Choose one)* | | Y |
| | [X] Always use the same VPN server | | |
| | *VPN server name or IP address* | vpn.contoso.com | |
| | [ ] Allow the user to chose a VPN server before connecting | | |
| | *Text file containing VPN server names or IP addresses:* | | |
| | Check, if appropriate: | | |
| | [X] Use the same user name and password for VPN and dial-up connections | | |

**Figure 9.8 Preparation for Running the CMAK Wizard: Top-Level Profile (page 2)**

| Include? (Y/N) | Implementation Information | Ready? (Y/N) |
|---|---|---|
| Y | VPN Entries<br>VPN entries provide configuration data for one or more VPN servers. Provide the network and security configuration to enable a client to connect to a VPN server.<br>*Complete the following section for each VPN Entry. Note that default entries are already checked in the Wizard.* | Y |
| | General | |
| | [ ] Disable file and printer sharing | |
| | [X] Enable clients to log on to a network *(default)* | |
| | TCP/IP Settings | |
| | Client DNS and WINS configuration *(Choose one)* | |
| | [X] Server assigns addresses *(default)* | |
| | [ ] Use these specific addresses for VPN connections | |
| | *Specify all that apply:* | |
| | Primary DNS address | |
| | Secondary DNS address | |
| | Primary WINS address: | |
| | Secondary WINS address | |
| | Choose all that apply: | |
| | [X] Make this connection the default gateway of the client *(default)* | |
| | [X] Use IP header compression *(default)* | |
| | Security<br>Basic Security Settings are supported by all Microsoft operating systems that can use Connection Manager profiles. Advanced Security Settings are only supported by Microsoft® Windows® 2000, Windows XP, and the Windows Server 2003 family. *Choose to configure Basic Settings, Advanced Settings, or both (the default).* | |
| | [X] Basic Security Settings | |
| | User authentication method *(Choose one)* | |
| | [ ] Allow any method | |
| | [ ] Require a secured password | |
| | [X] Require a Microsoft secured password | |
| | [X] Require data encryption *(default)* | |
| | VPN Strategy *(Choose one)* | |
| | [ ] Always use PPTP *(default)* | |
| | [X] Use L2TP/IPSec if available | |
| | [ ] Use a pre-shared key when using L2TP/IPSec | |

**Figure 9.9 Preparation for Running the CMAK Wizard: Top-Level Profile (page 3)**

| Include? (Y/N) | Implementation Information | Ready? (Y/N) |
|---|---|---|
| | [X] Advanced Security Settings | |
| | Data encryption (*Choose one*) | |
| | [ ] No encryption | |
| | [X] Require encryption *(default)* | |
| | [ ] Option encryption | |
| | Login Security (*Choose one*) | |
| | [ ] Use Extensible Authentication Protocol (EAP) | |
| | *Choose one, if using EAP* | |
| | [ ] MD5- Challenge | |
| | [ ] Protected EAO (PEAP) (encryption enabled) | |
| | [ ] Secured password (EAP-MSCHAP v2) (encryption enabled) | |
| | [ ] Smart Card or other certificate (encryption enabled) | |
| | [X] Authentication Methods *(default)* | |
| | *Choose all that apply, if not using EAP* | |
| | [ ] Unencrypted password (PAP) | |
| | [ ] Shiva Password Authentication Protocol (SPAP) | |
| | [ ] Challenge Handshake Authentication Protocol (CHAP) | |
| | [ ] Microsoft CHAP (MS-CHAP) *(default)* | |
| | [X] Microsoft CHAP Version 2 (MS-CHAP v2) *(default)* | |
| | VPN Strategy (*Choose one*) | |
| | [ ] Only use PPTP | |
| | [ ] Try PPTP first *(default)* | |
| | [ ] Only use L2TP | |
| | [X] Try L2TP first | |
| | [ ] Use a pre-shared key when using L2Tp/IPSec | |
| Y | Phone Book<br>This .pbk file must have a corresponding .pbr region file. This name cannot be longer than 8 characters and cannot contain a file extension, spaces, or any of the following characters:!,;*=\/:?"<> | Y |
| | Phone book file: `Contoso.pbk` | |
| | Check this box if you want Connection Manager to automatically download phone book updates. | |
| | [X] Automatically download phone book updates. | |
| Y | Phone Book Updates<br>If you checked "Automatically download phone book updates," type the URL for the CPS server. | Y |
| | Connection Point Services server: `http://pbupdate.contoso.com` | |

**Figure 9.10 Preparation for Running the CMAK Wizard: Top-Level Profile (page 4)**

| Include? (Y/N) | Implementation Information | Ready? (Y/N) |
|---|---|---|
| Y | Dial-up Networking Entries<br>Dial-up networking entries provide configuration data for one or more phone numbers.<br>*Complete this section for each dial-up networking entry. Note that default entries are already checked in the Wizard.* | Y |
| | [ ] Disable file and printer sharing | |
| | [ ] Enable clients to log on to a network | |
| | Dial up networking script (if applicable) | |
| | TCP/IP Settings: | |
| | Client DNS and WINS configuration (*Choose one*) | |
| | [X] Server assigns addresses *(default)* | |
| | [ ] All clients use these addresses | |
| | *Specify all that apply:* | |
| | Primary DNS address | |
| | Secondary DNS address | |
| | Primary WINS address: | |
| | Secondary WINS address | |
| | *Choose all that apply:* | |
| | [X] Make this connection the default gateway of the client *(default)* | |
| | [X] Use IP header compression *(default)* | |
| | Security<br>Basic Security Settings are supported by all Microsoft operating systems that can use Connection Manager profiles. Advanced Security Settings are only supported by Microsoft® Windows® 2000, Windows XP, and the Windows Server 2003 family. *Choose to configure Basic Settings, Advanced Settings, or both (default).* | |
| | [X] Basic Security settings | |
| | User authentication method (*Choose one*) | |
| | [ ] Allow any method | |
| | [ ] Require a secured password *(default)* | |
| | [X] Require a Microsoft secured password | |
| | [X] Require data encryption | |
| | [X] Advanced Security Settings | |
| | Data encryption (*Choose one*) | |
| | [ ] No encryptions | |
| | [X] Require encryption | |
| | [ ] Option encryption | |

**Figure 9.11 Preparation for Running the CMAK Wizard: Top-Level Profile (page 5)**

| Include? (Y/N) | Implementation Information | Ready? (Y/N) |
|---|---|---|
| | Login security *(Choose one)* | |
| | [ ] Use Extensible Authentication Protocol (EAP). | |
| | *Choose one, if using EAP:* | |
| | [ ] MD5- Challenge | |
| | [ ] Protected EAP (PEAP) (encryption enabled) | |
| | [ ] Secured password (EAP-MSCHAP v2) (encryption enabled) | |
| | [ ] Smart Card or other certificate (encryption enabled) | |
| | [X] Authentication Methods *(default)* | |
| | *Choose all that apply, if not using EAP.* | |
| | [ ] Unencrypted password (PAP) | |
| | [ ] Shiva Password Authentication Protocol (SPAP) | |
| | [ ] Challenge Handshake Authentication Protocol (CHAP) | |
| | [ ] Microsoft CHAP (MS-CHAP) *(default)* | |
| | [X] Microsoft CHAP Version 2 (MS-CHAP v2) *(default)* | |
| Y | Routing Table Update<br>Specify routing table changes in your service profile. You can include a route file in your service profile, provide a URL to a route file, or both. *Choose one.* | Y |
| | [ ] Do not change the routing tables | |
| | [X] Define a routing table update | |
| | Route the file to include | |
| | URL to a route file `http://www.contoso.com/CM/CMroute.txt` | |
| | *Check, if appropriate:* | |
| | [ ] If this URL is unavailable, disconnect the client. | |
| Y | Automatic Proxy Configuration<br>Automatically configures Internet Explorer proxy settings for your users through a proxy server. *Check one.* | Y |
| | [ ] Do not configure proxy settings | |
| | [X] Automatically configure proxy settings | |
| | Proxy settings file: `CMproxy.txt` | |
| | *Check, if appropriate:* | |
| | [X] Restore previous proxy settings of the user, after disconnecting | |
| Y | Logon Bitmap | Y |
| | *Check one:* | |
| | [ ] Default graphic | |
| | [X] Custom graphic (.bmp format, 330 x 140 pixels) | |
| | File Name `CMLogon.bmp` | |

**Figure 9.12 Preparation for Running the CMAK Wizard: Top-Level Profile (page 6)**

| Include? (Y/N) | Implementation Information | | | Ready? (Y/N) |
|---|---|---|---|---|
| Y | Phone Book Bitmap | | | Y |
| | *Check one:* | | | |
| | | Default graphic | | |
| | X | Custom graphic (.bmp format, 114 x 309 pixels | | |
| | File Name | CMPhone.bmp | | |
| Y | Icons | | | Y |
| | *Check one:* | | | |
| | | Default icons | | |
| | X | Custom icons (.ico files) | | |
| | | *Check all that apply and enter the custom icon file names:* | | |
| | | Icon type | Icon File Name | |
| | X | Program icon (32 x 32 pixels) | CMProgIc.ico | |
| | X | Title bar icon (16 x 16 pixels) | CMTitleIc.ico | |
| | | Notification area icon (16 x 16 pixels) | | |
| Y | Support Information | | | Y |
| | This is the text you want to appear in the logon dialog box. For example: For customer service, call 1-800-555-5555. | | | |
| | Support Information | For help, please call 1-800-555-0100 | | |
| Y | Connection Manager Software | | | Y |
| | Check, if appropriate | | | |
| | X | Install Connection Manager 1.3 with this service profile. | | |
| Y | Ready to Build the Service Profile | | | Y |
| | You have completed all steps required to build your service profile. To further customize the information included in this profile, check the Advanced Customization box. | | | |
| | *Check, if appropriate.* | | | |
| | X | Advanced Customization | | |
| Y | Advanced Customization | | | Y |
| | For each file you want to modify, enter the file name, the section name, the name of the key that you want to modify, and the new value for the key. | | | |
| | First file to customize: | | | |
| | | Description | Remove "Save Password" option | |
| | | File name | Contoso.cms | |
| | | Section name | Connection Manager | |
| | | Key name | HideRememberPassword | |
| | | Value | 1 | |

**Figure 9.13 Preparation for Running the CMAK Wizard: Top-Level Profile (page 7)**

| Include? (Y/N) | Implementation Information | | | Ready? (Y/N) |
|---|---|---|---|---|
| | Second file to customize: | | | |
| | | Description | Disable ICS | |
| | | File name | Contoso.cms | |
| | | Section name | Connection Manager | |
| | | Key name | DisableICS | |
| | | Value | 1 | |
| | Third file to customize: | | | |
| | | Description | Enable ICF | |
| | | File name | Contoso.cms | |
| | | Section name | Connection Manager | |
| | | Key name | EnableICF | |
| | | Value | 1 | |
| | Fourth file to customize: | | | |
| | | Description | Hide Advanced Tab | |
| | | File name | Contoso.cms | |
| | | Section name | Connection Manager | |
| | | Key name | HideAdvancedTab | |
| | | Value | 1 | |

Based on the decisions to provide routing table updates, automatic proxy configuration, and custom graphics, Contoso creates

several files before running the CMAK wizard:

- A route update file, CMRoute.txt.

  This plain text file includes information required to add or delete routes in the following format:

  ```
  Command Destination maskNetmask Gatewaymetric Metric if Interface
  ```

  Certain parameters can contain the value of **default**. In those cases, the appropriate information from the client computer is used.

  The file CMRoute.txt contains the following text to make all locations in the address range 192.168.0.0/16 reachable through the VPN connection:

  ```
  ADD 192.168.0.0 MASK 255.255.0.0 default METRIC default IF default
  ```

  For more information about including routing table updates, see Including routing table updates in Help and Support Center for Windows Server 2003.

- A proxy setting file, CMProxy.txt.

  This plain text file includes information to ensure that the user has appropriate access to internal and external resources.

  The proxy setting file for Contoso, CMProxy.txt, includes the following information:

  ```
  [Automatic Proxy]
  AutoProxyEnable=1

  [Manual Proxy]
  ProxyEnable=1
  ProxyServer=Contosoproxy:80
  ProxyOverride=<local>
  ```

  For more information about using automatic proxy configuration, see Using automatic proxy configuration in Help and Support Center for Windows Server 2003.

- Bitmap (.bmp) files for each of the custom graphics.

  Contoso creates custom bitmap files for the logon bitmap (330 x 140 pixels) and the phone book bitmap (114 x 309 pixels).

- Icon (.ico) files for each of the custom icons.

  Contoso creates custom icon files for the program icon (32 x 32 pixels) and the title bar icon (16 x 16 pixels). Contoso leaves the default notification area icon to show the connection status in the notification area.

# Contoso Uses CMAK to Create the Service Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After completing these worksheets, the company completes the CMAK wizard using this information.

The company completes the following steps to run the CMAK wizard:

1. Install CMAK from the Windows Component Wizard.

2. Create the component profile based on the information gathered in the component profile worksheet (Figure 9.6).

3. Create the top-level profile, merging the component profile to include the phone book provided by the ISP. The top-level profile is completed by using the information gathered in the top-level profile worksheet (Figure 9.7 to Figure 9.13).

# Contoso Tests Its Remote Access Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After creating the top-level service profile, Contoso tests its entire remote access solution before rolling it out. To test both the client and server aspects of the remote access solution, Contoso installs the Connection Manager profile onto a computer running each operating system the company supports and runs each possible user scenario. The company tests for the following from the Connection Manager client of each operating system:

- A VPN connection to the company's VPN remote access server.
- A direct-dial connection to the company's dial-up remote access server. By reviewing the Connection Manager log file, Contoso also confirms that the phone books are updating from the update URL provided by the company.
- A double-dial connection to several phone numbers provided by the ISP to ensure that the RADIUS proxy is forwarding the RADIUS messages to the Contoso RADIUS server. By reviewing the Connection Manager log file, Contoso also confirms that the phone books are updating from the update URL provided by A. Datum.

# Distributing the Connection Manager Profile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After thoroughly testing the entire remote access solution, Contoso distributes the Connection Manager service profile Contoso.exe to users at the company. Contoso uses a combination of distribution methods: the service profile is made available for download from the company's corporate network, and the service profile is preinstalled onto all new portable computers before distributing them.

The service profile can be downloaded from inside the corporate network and either installed on portable computers or saved to floppy disk and installed later on the user's home computer.

After the user has Contoso.exe on their computer, the user installs and sets up Connection Manager by using the following procedure.

**To install Contoso.exe**

1. Double-click **Contoso.exe**.
2. Click **Yes** when prompted to install Remote access to Contoso.
3. Click **My use only** to make the connection available to only the intended user and not all users of that computer. The user can also choose to add a shortcut to the connection on the desktop, and then click **OK**.
4. After the installation is complete, the user selects a primary and backup phone number or chooses to connect over their existing Internet connection.

   **Tip**

   - The user can create additional settings for home, travel, or other locations. Users create these customized settings by using the **New** button on the **General** tab of the **Remote access to Contoso Properties** dialog box and access them by using the **Use settings for** list on the **Remote access to Contoso** dialog box.

**To select a phone number**

1. Click **Properties** in the **Remote access to Contoso** dialog box.
2. On the **General** tab of the **Remote access to Contoso Properties** dialog box, click **Dial a phone number to connect**.
3. Click **Phone Book** to specify a phone number and a backup number.

**To use an existing connection to the Internet**

1. Click **Properties** in the **Remote access to Contoso** dialog box.
2. On the **General** tab of the **Remote access to Contoso Properties** dialog box, click **I am already connected to the Internet**.
3. The user clicks **OK** to return to the **Remote access to Contoso** dialog box.
4. Finally, the user enters their credentials and connects to the Contoso intranet.

# Additional Resources for Deploying Remote Access Clients Using Connection Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Related Information**

- "Deploying Dial-up and VPN Remote Access Servers" in this book for information about designing and deploying a remote access server solution.
- "Deploying Internet Authentication Service (IAS)" in this book for information about designing and deploying Internet authentication.
- "Connecting Remote Sites" in this book for information about creating connections to branch offices and remote sites.
- The *Internetworking Guide* of the *Windows Server 2003 Resource Kit* (or see the Internetworking Guide on the Web at http://www.microsoft.com/reskit) for information about the Routing and Remote Access service and remote access issues, virtual private networks, and Connection Manager.
- "Deploying Internet Authentication Service (IAS)" in this book for more information about Network Access Quarantine Control.

**Related Tools**

- Certificate Deployment (Cmgetcer.dll)

  Cmgetcer.dll allows Connection Manager to automatically obtain a certificate for L2TP/IPSec connections. For more information about Cmgetcer.dll, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Windows Resource Kit Tools Help.

- Profile Update (Getcm.exe and Instcm.exe)

  Getcm.exe runs as a post-connect action that checks for and downloads an updated service profile; Instcm.exe runs as a disconnect action that checks to see if an updated service profile has been downloaded and installs it. For more information about Getcm.exe and Instcm.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Windows Resource Kit Tools Help.

- Remote Access Quarantine Support Tools (RQC.exe and RQS.exe)

  IAS Network Access Quarantine Control provides phased network access for remote client computers. The Remote Access Quarantine Agent service is included when RQS.exe is installed on a remote access server. When you create the Connection Manager profile, you can include an administrator-provided script and RQC.exe. For more information about RQC.exe and RQS.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Windows Resource Kit Tools Help.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Administer phone books" in Help and Support Center for Windows Server 2003 for more information about how to administer phone books from the command line.
- "Add, edit, or delete POPs by command line" in Help and Support Center for Windows Server 2003 for more information about importing phone book files by using PBS.

- "Preparing to run the CMAK wizard" and "The Connection Manager Administration Kit Wizard" in Help and Support Center for Windows Server 2003 for more specifics about customizing Connection Manager by using CMAK.

- "Incorporating custom actions" in Help and Support Center for Windows Server 2003 for more information about custom actions.

- "Advanced Customization" in Help and Support Center for Windows Server 2003 for more information about advanced customization, including a comprehensive list of the service profile files and keys that you can customize.

- "IAS Network Access Quarantine Control" in Help and Support Center for Windows Server 2003 for more information about Network Access Quarantine Control.

- "Security information for Connection Point Services" in Help and Support Center for Windows Server 2003 for more information about phone book security.

- "Set up FTP accounts for known users" in Help and Support Center for Windows Server 2003 for more information about setting up FTP accounts for CPS.

# Connecting Remote Sites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Routing and Remote Access service in the Microsoft® Windows® Server 2003 operating system supports remote site connectivity by using dial-up or virtual private network (VPN) connections. You can deploy either a temporary or persistent site-to-site (also known as router-to-router) connectivity solution in your network to achieve optimal scalability, security, and performance while reducing your connection costs.

**In This Chapter**


Overview of Remote Site Connectivity

Choosing the Remote Site Connection Type

Choosing Security Features

Integrating the Remote Site Connection into Your Network

Preparing for Server Configuration

Deploying a Site-to-Site Connection

Additional Resources for Connecting Remote Sites


**Related Information**

- For more information about name resolution, see "Deploying Domain Name System (DNS)" and "Deploying WINS" in this book.
- For more information about IP address assignment, see "Deploying DHCP" in this book.
- For more information about the Active Directory® directory service replication between sites, see "Designing the Site Topology" in *Designing and Deploying Directory and Security Services* of this kit.
- For more information about creating hardware and software inventories and mapping your existing network before deploying new features, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

# Overview of Remote Site Connectivity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many organizations have offices located in different geographical locations, requiring remote site connectivity. You can use the Windows Server 2003 Routing and Remote Access service to deploy a cost-effective and secure site-to-site solution.

Traditionally, organizations have used wide area network (WAN) site-to-site connection technologies, such as T-Carrier or Frame Relay, to connect remote sites across a private data network. However, these private lines are expensive. For example, the prices for T-Carrier services are based on both bandwidth and distance, which makes the connections relatively costly. In addition, T-Carrier typically requires a dedicated infrastructure, including a Channel Service Unit/Data Service Unit (CSU/DSU) and line-specific routers at each end of the connection.

In contrast, you can integrate the Windows Server 2003 Routing and Remote Access service solution into your organization's current network by using existing servers. With the site-to-site connections provided by the Routing and Remote Access service, you have two alternatives to conventional WAN links: a site-to-site dial-up connection or a site-to-site VPN connection. If you deploy a Routing and Remote Access solution to replace an existing WAN connection, or to implement a new connection, you can optimize cost savings by tailoring your connection type to your traffic volume. You can also customize security to fit your organization's requirements.

**Note**

- The Routing and Remote Access service can support both site-to-site connections between remote offices and remote access connections for individual computers. This chapter focuses on site-to-site connections.

Before you begin the design process to introduce a new site-to-site connection into your network, or modify an existing connection, inventory your existing hardware and software, and create or update a map of your current network. Updating your inventory and network configuration information facilitates both the design and the deployment phases. For a guide to conducting inventories and creating a network map, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

**Note**

- For a list of job aids that are available to assist you in deploying site-to-site connection technology, see "Additional Resources for Connecting Remote Sites" later in this chapter.

# Remote Site Connectivity Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To begin the design process for deploying a site-to-site connection, choose the remote site connection type and its configuration options that are most appropriate, and decide which security features to use to protect that connection. Next, decide how you want to integrate the remote site connection into your existing network infrastructure. Finally, prepare the servers that you plan to configure as routers. After you complete these design decisions, you are ready to deploy your remote site connection. The flowchart in Figure 10.1 outlines this process.

**Figure 10.1 Connecting Remote Sites**



# Remote Site Connectivity Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can design and deploy a remote site connection that is optimal for your organizational and network environment by using the connectivity, security, and network configuration options provided by the Routing and Remote Access service.

The following sections describe three typical remote site connection solutions — a Point-to-Point Tunneling Protocol (PPTP) VPN connection, a Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec) VPN connection, and a dial-up connection. For detailed information about each connection type and the range of available configuration and security options, see "Choosing the Remote Site Connection Type" and "Choosing Security Features" later in this chapter.

**PPTP VPN Solution**

Organizations with moderate to heavy traffic between a branch office and a main office and existing connections to the Internet might choose a PPTP–based site-to-site connection. In the example shown in Figure 10.2, a firewall creates a perimeter network at each end of the Internet tunnel. Windows Server 2003 also supports VPN functionality without the use of a firewall.

Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2), which provides a high-security protocol for password authentication, is a highly recommended method for authentication and encryption key generation for a site-to-site connection. Alternatively, you can use Extensible Authentication Protocol-Transport Layer Security Protocol (EAP-TLS), which provides an even stronger user-level authentication than the password-based MS-CHAP v2.

The main office site must have a permanent WAN link to its local Internet service provider (ISP), but the branch office site can use a dial-up WAN link to its local ISP. An on-demand connection that disconnects when the connection is idle ensures that the connection is not active when not in use. Reciprocal replication ensures that replication between domain controllers in separate sites takes place over the one-way initiated on-demand connection. Figure 10.2 depicts this solution.

**Figure 10.2 One-Way Initiated On-Demand Dial-up PPTP VPN Solution**




**L2TP/IPSec VPN Solution**

Organizations that need maximum security to support substantial two-way traffic between a large branch office and a headquarters office might choose an L2TP/IPSec VPN solution. A firewall creates a perimeter network at each end of the Internet tunnel. A persistent connection, enabled by a dedicated link to the ISP at both sites, allows around-the-clock traffic.

The recommended method for the computer-level authentication provided by L2TP/IPSec is the exchange of computer certificates by the calling and answering endpoints, which requires a certificate infrastructure provided by the certification authority (CA) server. EAP-TLS provides stronger user-level authentication than does the password-based MS-CHAP v2, because it requires a user certificate on the calling endpoint and a computer certificate on the answering endpoint. L2TP/IPSec uses IPSec as its encryption method. For a persistent connection, replication takes place across the site link at specified intervals — you do not need to configure reciprocal replication as you do for a one-way initiated on-demand connection. The example in Figure 10.3 depicts this solution.

**Figure 10.3 Persistent Two-Way Initiated L2TP/IPSec VPN Solution**




**Dial-up Solution**

Organizations with moderate traffic between a branch and a main office might choose to replace an existing leased WAN link with a dial-up WAN link or use a dial-up WAN link to create a new connection. Figure 10.4 shows an example of a site-to-site connection that uses an ISDN dial-up link. One common situation in which you might deploy a dial-up link is as a backup solution when a VPN link provides your primary connection. For more information about when to use a dial-up connection, see "Choosing a Dial-up or VPN Connection" later in this chapter.

A one-way initiated on-demand dial-up connection disconnects when the connection is idle for a specified period of time and thus provides efficient access for branch office users who need only intermittent access to the main office. A dial-up connection typically uses MS-CHAP v2as the user authentication method to authenticate the calling router together with Microsoft Point-to-Point Encryption (MPPE) for data encryption. Configuring reciprocal replication enables replication between domain controllers in separate sites over a one-way initiated on-demand connection. Figure 10.4 depicts this solution.

**Figure 10.4 One-Way Initiated On-Demand Dial-up Solution**



# Choosing the Remote Site Connection Type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first set of decisions that you need to make when designing remote site connectivity are about the connection type. The flowchart in Figure 10.5 shows the tasks required when choosing connection type options.

**Figure 10.5 Choosing a Remote Site Connection Type**



# Choosing a Dial-up or VPN Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When considering your connection type options, your first decision is whether to use a dial-up (non-VPN) connection or a VPN connection.

# Routing and Remote Access Dial-up Connections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Routing and Remote Access connections that run over a physical device (such as an ISDN adapter or analog modem) that is installed on both the calling and answering routers are known as *dial-up connections*. These connections differ from traditional WAN links in that they use existing telephone lines instead of leased lines. A dial-up connection differs from a VPN connection in that it does not cross the Internet.

Using existing phone lines can substantially reduce connection costs, especially where traffic volume between sites is low. Instead of paying for a permanent WAN connection 24 hours a day, you can configure the link to disconnect when no traffic crosses the connection for a specified period of time. For example, customers who use ISDN typically pay by the minute or by the byte, so configuring the call to hang up when the connection is idle reduces cost.

To keep data transmission secure, a dial-up connection uses private, dedicated lines or circuits across a carrier's network. In addition, the connection uses Point-to-Point Protocol (PPP) user authentication and MPPE for data encryption.

An organization might choose to deploy a dial-up connection if it has a small branch office that needs an occasional dial-up connection to a main office. For a dial-up link used as the main site-to-site connection, using ISDN is a viable option.

A larger organization that uses a VPN connection might also deploy an ISDN dial-up connection as a backup solution in case the Internet connection in either site is ever unavailable. You can use a Public Switched Telephone Network (PSTN) dial-up link as the backup if the primary link is Frame Relay with a committed information rate (CIR) of 56 Kbps. If the link for which you need a backup is faster than 56 Kbps CIR, using a PSTN dial-up connection as the backup is not feasible, because users accustomed to a faster connection will perceive the dial-up connection as too slow or inoperable. If the primary link is faster than 56 Kbps, instead use a synchronous circuit solution, such as ISDN or a leased-line circuit, for the backup.

In either of these situations, a dial-up connection provides a solution that is easy to deploy and cost-effective.

# Virtual Private Network Connections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Routing and Remote Access connections that cross the Internet (or other shared network) are known as virtual private network (VPN) connections. A VPN connection typically uses a physical link to a local ISP. A VPN-based answering router always uses a permanent WAN link to an ISP, such as a T-Carrier, Frame Relay, DSL, or cable modem link. The answering router's permanent link to the Internet ensures that the router is available whenever a calling router attempts to establish a connection. This link requires a static IP address, which is assigned by your local ISP (or by InterNIC), on the answering router's Internet-connected interface. However, a VPN-based calling router might use a permanent WAN link to an ISP, or it might use a temporary link (such as a modem or adapter) to the ISP.

When traffic volume is high or permanent connectivity across long distances is required, using a VPN connection to transport data across the Internet is more cost-effective than using a dedicated WAN line or a dial-up connection.

To keep data transmission secure, a VPN connection uses PPP user authentication, routes packets encapsulated in a secure tunnel across the Internet, and uses MPPE or IPSec encryption to protect the data portion of the packet. This virtual point-to-point connection emulates a dedicated, private, physical point-to-point connection but lets you replace long-distance WAN links with local WAN links to your nearby ISP. The ISP does not install any VPN software on its equipment, nor is any VPN software installed on the intermediate routers that a VPN data packet crosses.

If you choose a VPN connection in which the calling router uses a temporary link to connect to its local ISP, the temporary link can be either a dial-up link or a PPP over Ethernet (PPPoE) link. The calling router first establishes the dial-up or PPPoE link to the ISP and then establishes the VPN tunnel across the Internet to the answering router. Many broadband ISPs use PPPoE. PPPoE links to the ISP are faster than dial-up links.

An example of an organization that might choose to deploy a VPN connection is a large enterprise with several large satellite offices that each need secure connection for high-volume traffic across the Internet to a headquarters office.

**Choosing PPTP or L2TP/IPSec**

If you choose a VPN site-to-site connection, you must next decide whether to use PPTP or L2TP/IPSec as the VPN technology. Table 10.1 lists some of the factors you need to consider in order to determine whether to deploy a PPTP or an L2TP/IPSec solution. For more information about the security options (user authentication, certificates, and encryption) described briefly in the table, see "Choosing Security Features" later in this chapter.

**Table 10.1 Comparing a PPTP Solution with an L2TP/IPSec Solution**

| Factor | Using PPTP | Using L2TP/IPSec |
|---|---|---|
| Windows version | PPTP is supported by Windows Server 2003, as well as the Microsoft® Windows® 2000 Server operating system, and the Microsoft® Windows NT® version 4.0 operating system with the Routing and Remote Access Service (RRAS). Most third-party VPN routers also support PPTP. | L2TP/IPSec is supported by Windows Server 2003 and Windows 2000 Server (with the Routing and Remote Access service). Most third-party VPN routers also support L2TP/IPSec. |
| User authentication | EAP-TLS or MS-CHAP v2 is recommended. | EAP-TLS or MS-CHAP v2 is recommended. |

| | | |
|---|---|---|
| Certificates | PPTP requires certificates only when using EAP-TLS for user authentication, in which case a user certificate for the calling router and a computer certificate for the authenticating server of the answering router are required. | • Using EAP-TLS user authentication with L2TP/IPSec requires a user certificate for the calling router and a computer certificate for the authenticating server of the calling router.<br>• For computer-level authentication, L2TP/IPSec supports computer certificates or preshared keys as the authentication method for IPSec. Computer certificate authentication is recommended and requires a computer certificate on both the calling and answering routers. |
| Encryption | For a PPTP-based VPN connection, choosing either EAP-TLS or MS-CHAP v2 for user authentication provides MPPE for data encryption. MPPE provides data confidentiality (encryption); that is, captured packets cannot be interpreted without the encryption key. | L2TP/IPSec generates its encryption keys by using IPSec. IPSec provides data confidentiality (encryption), replay protection, data integrity, and data origin authentication. |
| NATs | In most cases, you can locate PPTP-based calling routers behind a network address translator (NAT), so you can configure a small or home office (SOHO) network to share a single connection to the Internet. Most NATs include a NAT editor that can accurately translate PPTP-tunneled data. | With Windows Server 2003–based calling or answering routers, you can use IPSec NAT traversal (NAT-T) to create L2TP/IPSec connections across NATs. Using NAT-T requires running Windows Server 2003 on both the calling and answering routers. With NAT-T, hosts that are hidden behind a NAT can use IPSec to connect to a remote site. |
| Ease of deployment | When using MS-CHAP v2 (or MS-CHAP) for user authentication, PPTP is cost-effective and easier to deploy than L2TP/IPSec with computer certificates. | • When you use computer certificates as the authentication method, L2TP/IPSec requires a certificate infrastructure and, therefore, requires more administrative overhead to deploy and maintain than does PPTP used with password-based MS-CHAP v2. For optimal security, using certificates is the recommended method.<br>• Using preshared keys as the authentication method, L2TP/IPSec requires less administrative overhead (for initial setup but not for long-term administration) than using L2TP/IPSec with certificates but more administrative overhead than using PPTP. To ensure security, do not use preshared keys.[1] |

1. For more information about the advantages of computer certificates over preshared keys, see "Choosing Computer Certificates or Preshared Keys for Computer-Level Authentication" later in this chapter.

**Using Both a PPTP Connection and an L2TP/IPSec Connection**

You can deploy both a PPTP solution and an L2TP/IPSec solution at the same time. By default, a VPN router running Windows Server 2003 simultaneously supports both connection types. A VPN router running Windows 2000 Server supports both connection types if the router is not behind a NAT; however, Windows 2000 Server supports only PPTP — but not L2TP/IPSec — across a NAT.

You might want to use PPTP for some site-to-site connections and L2TP/IPSec for others. Table 10.2 lists some situations in which you might use PPTP or L2TP/IPSec for different connections on the same network.

**Table 10.2 Using PPTP and L2TP/IPSec for Different VPN Connections on the Same Network**

| VPN Connection Type | Typical Uses |
|---|---|

| | |
|---|---|
| PPTP | • To connect from calling routers that are running Windows NT 4.0 with RRAS.<br>• To connect from routers that are running Windows Server 2003 or Windows 2000 Server that do not have an installed computer certificate.<br>• To establish a VPN connection when you want to place Windows 2000 Server–based routers behind a Routing and Remote Access NAT or a third-party NAT. |
| L2TP/IPSec | • To connect from calling routers running Windows Server 2003 or Windows 2000 Server that have an installed computer certificate.<br>• To enable hosts that are hidden behind a NAT (because they use private addresses) to use IPSec to connect to a remote site. This is possible because NAT-T, which is enabled by default on Windows Server 2003–based computers, can create L2TP/IPSec connections across a NAT.<br>• To provide the highest security solution available.<br>• To connect from calling routers running Windows Server 2003 that use preshared keys (not recommended for security reasons). |

If you use both a PPTP and an L2TP/IPSec solution, you can create separate remote access policies that define different connection parameters for each connection type. For example, if you have PPTP VPN connections to two branch offices whose calling routers run Windows NT 4.0 with RRAS, you might want to create a more restrictive remote access policy for those connections than the remote access policy that you create for an L2TP/IPSec VPN connection to a branch office whose calling routers run Windows Server 2003 and use computer-level authentication. For more information about creating remote access policies for a site-to-site connection, see "Configure a Remote Access Policy" later in this chapter.

# Choosing an On-Demand or Persistent Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure the calling router for any of the connection types — dial-up, PPTP VPN, or L2TP/IPSec VPN — with either an on-demand or a persistent connection. Table 10.3 describes and compares these connection type options.

**Table 10.3 Comparing On-Demand and Persistent Connections**

| Connection Type | Description | Use |
|---|---|---|
| On-demand connection | Establishes a connection when traffic is forwarded, and it terminates the connection when the link is not used for a specified period of time. | Use an on-demand connection if using the communications link incurs per-minute charges.<br><br>For an on-demand VPN connection, the initiating router can use either a permanent or a dial-up link to the Internet. The answering router must have a permanent link to the Internet to ensure that it is available when a calling router attempts to establish a connection. |
| Persistent connection | Sustains a connection for 24 hours a day. | Use a persistent connection in the following circumstances:<br>• When the cost of the connection is based on a flat fee, such as for a link to a local ISP for each site when sites are located in separate cities or for a connection between different sites within the same city.<br>• When data traffic is time-sensitive. For example, if you support mainframe terminal connectivity between sites, if the terminals must wait for an on-demand VPN connection to be activated, the connection attempt will time out before the session can be launched.<br><br>For a persistent VPN connection, both the calling and the answering router must use a permanent link to the Internet. |

For on-demand connections, to prevent the calling router from establishing unnecessary connections, you can use demand-dial filtering and dial-out hours:

- **Demand-dial filters.** To prevent a VPN calling router from initiating unnecessary connections, you can configure demand-dial filters to specify the types of IP traffic for which the router will or will not create a demand-dial connection. You can identify traffic to accept or reject based on source and destination addresses of incoming traffic and the protocol in use. It is recommended that you match the demand-dial filters to the IP packet filters configured on the demand-dial interface. If there is specific traffic that is not allowed across the demand-dial interface when it is connected, that same traffic should not be allowed to initiate a demand-dial connection using that interface. For example, if you have a packet filter that prevents ICMP traffic from being sent across the demand-dial interface, then you should configure a demand-dial filter to prevent ICMP traffic from initiating the demand-dial connection. For more information about matching demand-dial filters to IP packet filters, see "Integrate the VPN Server into a Perimeter Network" and "Configure IP Packet Filters and Demand-Dial Filters" later in this chapter.

- **Dial-out hours.** To prevent a dial-up or VPN calling router from initiating unnecessary connections, you can configure dial-out hours to specify the hours during which a calling router is either permitted to make a site-to-site connection or denied the connection. You can also configure remote access policies on the answering router to restrict the time periods when incoming demand-dial connections are allowed.

# Choosing a One-Way or Two-Way Initiated Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can set up a site-to-site connection so that it can be initiated from only one location, or you can configure a connection that can be initiated from either side of the dial-up or VPN connection. Table 10.4 describes and compares these options.

**Table 10.4 Comparing One-Way and Two-Way Initiated Connections**

| Connection Type | Description | Use |
|---|---|---|
| One-way initiated connection | A one-way initiated connection is one in which one site contains only an answering router and the other site contains only a calling router. | Use an on-demand or persistent one-way initiated connection when users at a branch office need to connect to a headquarters office but not vice versa.<br>Use a persistent one-way initiated connection when you have 10 or more branch offices and users at each site need to access the other sites. When you use a two-way initiated connection for 10 or more connections, performance is too slow. |
| Two-way initiated connection | A two-way initiated connection is one in which a router at each site can function as both the answering router and as the calling router. | Use a two-way initiated connection when users at both locations need to access resources or people in the other location and you have 10 or fewer connections (if you have 10 or more connections, use a one-way initiated persistent connection). |

For both a one-way and a two-way initiated connection, use the properties of the router in the Routing and Remote Access snap-in to configure both the calling and the answering router as a local area network (LAN) router and as a demand-dial router.

# Choosing Security Features

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As described in "Choosing the Remote Site Connection Type" earlier in this chapter, security is one of the factors that you weigh when you choose a connection type. For all connection types, securing site-to-site communications requires making decisions about several additional security features. The flowchart in Figure 10.6 shows the tasks required when choosing these additional security features.

**Figure 10.6 Choosing Security Features**



# Integrate the VPN Server into a Perimeter Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 supports VPN functionality without the use of a firewall. However, many organizations use firewalls to implement a perimeter network (also known as a DMZ, demilitarized zone, and screened subnet) to increase security by protecting their internal network from intrusion through the Internet.

Only servers that provide resources to users over the Internet, such as Proxy, Web, and FTP servers, are located in the perimeter network. Traffic between dial-up routers does not cross the Internet. Therefore, you do not need to locate dial-up routers in a perimeter network. However, placing VPN routers in a perimeter network helps ensure that communications between sites are protected.

**VPN Router Placement in Relation to Firewall**

If your organization already uses a perimeter network, you can add your VPN router to the existing set of servers on the perimeter network. If not, consider adding a perimeter network to your infrastructure when you deploy a VPN site-to-site connection.

How you configure firewall filters and the filters on the VPN router depends on the position of the VPN router relative to the firewall. Although it is possible to place the VPN router in front of the firewall (with the VPN router attached to the Internet), the more common — and recommended — configuration for a site-to-site connection is to place the VPN router behind the firewall (attaching the firewall to the Internet). When you place the VPN router behind the firewall, you configure the firewall with input and output filters on the firewall's Internet and perimeter network interfaces to restrict traffic to the VPN server. These filters are configured the same for a site-to-site VPN server as for a remote access VPN server. For a description of these filters and how they function, see the instructions for configuring packet filters for a VPN server in "Deploying Dial-up and VPN Remote Access Servers" in this book.

For more information about VPN servers and firewalls, including configuration of PPTP and L2TP/IPSec packet filters both for VPN servers behind the firewall and for VPN servers in front of the firewall, see "VPN servers and firewall configuration" in Help and Support Center for Windows Server 2003.

**Match IP Packet Filters to Demand-Dial Filters**

At the same time that you plan where to place your VPN router in relation to a firewall and how to configure PPTP and L2TP/IPSec IP packet filters on the firewall, also plan how to configure demand-dial filters in conjunction with the IP packet filters configured on the demand-dial interfaces. Although IP packet filters and demand-dial filters serve different purposes, Microsoft recommends that you configure them together.

You use demand-dial filters, which are applied *before* a connection is made, to specify which types of traffic are allowed to create a connection in the first place. You use IP packet filters, which are applied *after* a connection is made, to specify what traffic is allowed into and out of an interface after a connection is established. To prevent the demand-dial connection for traffic that will be discarded by the IP packet filters, you need to match your demand-dial and IP packet filters.

For more information about configuring IP packet filters to match your demand-dial filters, see Demand-dial routing design considerations in Help and Support Center for Windows Server 2003.

# Choosing the Authentication Provider

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Site-to-site connections between remote offices require an authentication and accounting provider for:

- Authentication of calling router credentials and authorization of the site-to-site connection.
- Accounting services that record the creation and termination of each site-to-site connection.

When a connection is attempted, the answering router authenticates the credentials of the calling router by using one of two authentication providers: Windows or RADIUS. Your choice of authentication provider is determined by whether your solution involves a site-to-site only connection or a combined site-to-site and remote access connection:

- **For a site-to-site only connection, choose Windows.** When you choose Windows as the authentication, authorization, and accounting provider, the same Windows authentication process that validates user credentials when a user logs on also validates the calling router.
- **For a combined site-to-site and remote access connection, choose either Windows or RADIUS.** If the same answering router will support both a site-to-site connection and remote access users (such as home or mobile users), you can use either Windows or Remote Authentication Dial-in User Service (RADIUS) as your authentication provider. Servers running the Internet Authentication Service (IAS) component of Windows Server 2003 provide an Internet standards–compliant RADIUS server and proxy.

If you have more than one answering router or other types of access servers (such as wireless access servers), you can use a single RADIUS server to provide centralized authentication, authorization, and accounting for all answering routers and access servers instead of administering each answering router and access server separately. To simplify administration for a combined site-to-site and remote access connection, you can use IAS to store both site-to-site and remote access information.

In an Active Directorydomain, it is recommended that you use the Windows Server 2003 version of IAS as your RADIUS server. The IAS RADIUS server is tightly integrated with Windows Server 2003, Active Directory, and the Routing and Remote Access service. When you use RADIUS authentication, you configure each of participating answering router as a RADIUS client. After you configure both the answering router and the IAS server, the answering router uses remote access policies stored on the IAS server instead of those on the answering router.

Although it is possible to use RADIUS as the authentication provider for a site-to-site only connection, you do not need RADIUS. Deploying an IAS server is unnecessary administrative overhead for a demand-dial connection that connects two sites but does not support remote access users.

The credentials that the calling router passes to the answering router for verification are those of a user account, either in Active Directory or on the local computer. Authorization is granted based on the dial-in properties that you specify in the user account and on remote access policies configured on the answering router (or on the RADIUS server). For more information, see "Choosing Router User Accounts and Groups" and "Choosing a Remote Access Policy Type" later in this chapter.

The authentication provider that you choose also functions as the authorization provider. However, the Routing and Remote Access service does not require that you use the same provider for authentication and authorization that you use for accounting. You can use Windows for authentication and RADIUS for accounting, or vice versa. However, if you have multiple answering routers that support remote access users, consider using RADIUS for integrated authentication, authorization, and accounting, and, if you use IAS, to manage remote access policies.

For more information about RADIUS authentication, see "Checklist: Configuring IAS for dial-up and VPN accessChecklist: Configuring IAS for dial-up and VPN access" in Help and Support Center for Windows Server 2003. For more information about deploying IAS to perform centralized authentication, authorization, and accounting, see "Deploying Internet Authentication Service (IAS)" in this book.

# Choosing an Authentication Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For site-to-site connections, you must implement user-level authentication, and, for greater security, you might decide to implement computer-level authentication as well:

- User-level authentication — for all site-to-site connection types.
- Computer-level authentication — for L2TP/IPSec connections only.

# Choosing EAP-TLS or MS-CHAP v2 for User-Level Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

All site-to-site connection technologies require user-level authentication. The "user" to be authenticated is the calling router. To prevent a nonauthorized router from making a connection, the Routing and Remote Access service requires that the calling router requesting the connection be authenticated by the answering router.

The Routing and Remote Access service can use any of several PPP authentication methods to provide user-level authentication. For optimal security, choose one of the two strongest recommended PPP user authentication methods — EAP-TLS or MS-CHAP v2.

For more information about the other PPP authentication protocols (in addition to EAP-TLS and MS-CHAP v2) supported by Windows Server 2003, Windows 2000 Server, and Windows NT Server 4.0, see "Authentication methods" in Help and Support Center for Windows Server 2003.

**Certificate-based EAP-TLS**

EAP-TLS is a more secure protocol than MS-CHAP v2 for user-level authentication on a dial-up, PPTP VPN, or L2TP/IPSec VPN connection. EAP-TLS requires a certificate infrastructure, because it uses a user certificate for the calling router and a computer certificate for the authenticating server of the answering router. The identity of the authenticating server depends on the authentication provider:

- If Windows provides authentication, as is the case for a site-to-site only connection, the computer certificate is installed on the answering router. The answering router must be joined to the Active Directory domain.
- If RADIUS provides authentication, as might be the case if your answering router also supports remote access users, the computer certificate is installed on the RADIUS server. If the RADIUS server is an IAS server, it must be joined to the Active Directory domain. In this case, the answering router does not need to belong to the domain.

EAP-TLS is recommended for all connection types both because it is the strongest method for user authentication and because it provides mutual user authentication between calling and answering routers. The calling router submits a user certificate for authentication, and the answering router (or RADIUS server) submits a computer certificate for authentication. With EAP-TLS, the user account that the answering router uses to authenticate and authorize the calling router must be an Active Directory user account.

If you plan to use certificate-based EAP-TLS, you can use a third-party CA if the conditions in Table 10.5 are met.

**Table 10.5 Requirements for Using a Third-Party CA**

| Certificates Installed on Authenticating Server (Answering Router or RADIUS Server) | Certificates Installed on Calling Router |
|---|---|

| | |
|---|---|
| • Computer certificates must be installed in the Local Computer certificate store.<br>• Computer certificates must have a corresponding private key.<br>• The cryptographic service provider for the certificates must support secure channel (Schannel). (If not, the authenticating server cannot use the certificate and the certificate cannot be selected from the properties of the Smart Card or Other Certificate EAP type on the Authentication tab in the properties of a profile for a remote access policy.)<br>• Computer certificates must contain the Server Authentication certificate purpose (also known as an enhanced key usage [EKU]). An EKU is identified using an object identifier (also known as an OID). The object identifier 1.3.6.1.5.5.7.3.1 represents Server Authentication.<br>• Computer certificates must contain the fully qualified domain name (FQDN) of the computer account of the authenticating server computer in the Subject Alternative Name property.<br>• The root CA certificates of the CAs that issued the calling router user certificates must be installed in the Trusted Root Certification Authorities certificate store. | • User certificates must have a corresponding private key.<br>• User certificates must contain the Client Authentication EKU (the 1.3.6.1.5.5.7.3.2 object identifier).<br>• User certificates must be installed in the Current User certificate store.<br>• User certificates must contain the universal principal name (UPN) of the user account in the Subject Alternative Name property.<br>• The root CA certificates of the CAs that issued the authenticating server computer certificates must be installed in the Trusted Root Certification Authorities store. |

For examples of how to deploy EAP-TLS certificate-based user authentication for site-to-site connections, see Deploying certificate-based authentication for demand-dial routing in Help and Support Center for Windows Server 2003. For more information about certificate types and requirements, see Certificate Services in Help and Support Center for Windows Server 2003, and see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

**Password-based MS-CHAP v2**

Like EAP-TLS, MS-CHAP v2 also provides mutual authentication between the calling router and the answering router. MS-CHAP v2 is a password-based authentication method in which the answering router uses either an Active Directory user account or a local user account to authenticate the calling router. MS-CHAP v2 is the recommended method for user authentication if a certificate infrastructure is not available.

If you use a local user account for MS-CHAP v2 authentication, the demand-dial routers do not need to join the Active Directory domain. Be sure to use strong passwords with MS-CHAP v2. In an Active Directory domain, you can use Group Policy settings to enforce the use of strong passwords.

**Advantages of EAP-TLS and MS-CHAP v2**

Although several weaker PPP authentication protocols are also available for user authentication, EAP-TLS and MS-CHAP v2 are the preferred methods because both provide the following benefits:

- **Mutual authentication.** With mutual authentication, the calling router is authenticated by the answering router, and the answering router is then authenticated by the calling router. This mutual authentication prevents an attacker from masquerading as the answering router, whereas other types of authentication verify only the identity of the calling router.

- **Encryption features.** In addition, for a dial-up or PPTP connection, EAP-TLS and MS-CHAP v2 provide the following encryption features. (L2TP/IPSec generates its encryption keys using IPSec and therefore does not need EAP-TLS or MS-CHAP v2 for encryption keys.)

    - **Stronger initial encryption keys.** MS-CHAP v2 uses both a unique session identifier and user credentials to generate encryption keys. EAP-TLS uses user certificates to verify user identity. (In contrast, MS-CHAP uses only the user name and password to create these keys, making it easier for attackers to determine the keys in use.)

    - **Different keys for sending and receiving data.** The calling router and answering router use separate keys to encrypt data and to send data. Using different keys makes it more difficult for attackers to determine which key is

used for encryption.

- **MPPE for encryption.** MPPE uses Rivest-Shamir-Adleman (RSA) RC4 stream ciphering with 40-bit, 56-bit, and 128-bit encryption keys. (RC4, developed by RSA Data Security, is a secret key cryptographic method that uses a variable-length key.) Encryption key strength is negotiated during connection establishment (as is authentication method). The client (calling router) and server (answering router) negotiate the strongest mutually allowable key size. If the answering router requires a larger key than that supported on the calling router, the answering router rejects the connection.

For more information about encryption for site-to-site connections, see "Choosing MPPE or IPSec Encryption" later in this chapter.

# Choosing Computer Certificates or Preshared Keys for Computer-Level Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The only site-to-site connection technology that provides computer-level authentication is an L2TP/IPSec VPN connection. Computer-level authentication occurs in one of two ways:

- **Computer certificates:** the exchange of computer certificates by the calling and answering routers. Computer-level authentication requires that you deploy a public key infrastructure (PKI). Although computer certificate authentication requires more administrative overhead for initial setup than does the use of preshared keys, it is the recommended method because it provides stronger computer authentication than the preshared keys method. Windows Server 2003 supports the automatic enrollment of certificates, which makes certificate deployment and management easier than using preshared keys over the long term.

- **Preshared keys:** the exchange of preshared keys during the establishment of the IPSec security association (SA). Support for preshared keys is new with Windows Server 2003, and it requires running Windows Server 2003 on both VPN routers. A preshared key is a text string that is configured on both the calling and the answering router. Because a preshared key is a weaker computer authentication method than certificate authentication, Microsoft recommends that you use preshared key authentication only during the time you are deploying a PKI to enable the use of certificates. Using preshared key authentication is not as secure as using computer certificates, but it requires less administrative overhead.

The IPSec Internet Key Exchange (IKE) protocol can use either certificate-based or preshared key authentication to negotiate security for the L2TP traffic.

For more information about creating a certificate infrastructure, see Certificate Services in Help and Support Center for Windows Server 2003, and see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit. For more information about key exchange, see Internet Key Exchange in Help and Support Center for Windows Server 2003.

# Choosing MPPE or IPSec Encryption

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For site-to-site connections, the connection type and the user authentication protocol that you choose to deploy determine the data encryption method. Table 10.6 shows the available options.

**Table 10.6 Choosing a Data Encryption Method**

| Connection Type | Recommended User Authentication Protocol | Encryption Method |
|---|---|---|
| Dial-up connection | EAP-TLS or MS-CHAP v2 | MPPE |
| PPTP connection | EAP-TLS or MS-CHAP v2 | MPPE |
| L2TP connection | EAP-TLS or MS-CHAP v2 | IPSec |

Understanding the following features can help you decide how you want to manage encryption:

- **Link encryption versus end-to-end encryption.** MPPE provides *link encryption*. Link encryption encrypts data as it passes between the calling and answering routers. In addition to providing computer-level authentication, IPSec provides *end-to-end* encryption for data that passes between the sending and receiving nodes.

- **Encryption method used if VPN connection type is Automatic.** If you configure a VPN connection for an **Automatic** server type (the default), the connection first tries to use PPTP and its associated MPPE encryption, and then it tries to use L2TP and its associated IPSec encryption. If you configure the VPN connection to connect to a PPTP server, only MPPE encryption is used. If you configure the VPN connection to connect to an L2TP server, only IPSec encryption is used.

- **No encryption needed for link to ISP.** For VPN connections, you do not need to use encryption for the link between your site and the ISP, because no data is transmitted during the process that establishes this connection. After the connection to the ISP is made, the data that passes between the calling and answering routers is encrypted as it passes through the VPN tunnel.

You configure MPPE and IPSec encryption strengths on the **Encryption** tab for the properties of a remote access policy. For information about how to configure encryption in a remote access policy for a site-to-site connection, see "Configure a Remote Access Policy" later in this chapter. For general information about configuring encryption, see Add a remote access policy and Remote Access Policies Examples in Help and Support Center for Windows Server 2003.

Configure either MPPE or IPSec to use one of the encryption keys as shown in Table 10.7.

**Table 10.7 Encryption Strength by Connection Type**

| Encryption Strength | Dial-up or PPTP | L2TP/IPSec |
|---|---|---|
| **Basic** | 40-bit MPPE | 56-bit DES |
| **Strong** | 56-bit MPPE | 56-bit DES |
| **Strongest** | 128-bit MPPE | 3DES (three 56-bit keys) |

**Note**

- Windows NT 4.0 with the 128-bit version of Service Pack 4 (SP4) can support 128-bit MPPE, but it does not support 56-bit MPPE. Therefore, any Windows operating system earlier than Windows NT 4.0 SP4 is not recommended, because security enhancements for MS-CHAP and MPPE are not included.

# Choosing Router User Accounts and Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After a calling router is authenticated by using either Windows or RADIUS as the authentication provider, it must be authorized: that is, it must be given permission to establish a connection with the answering router. You use two interrelated sets of components to authorize access by the calling router: user accounts and (optionally) groups, and remote access policies.

Before you can successfully configure either the router user accounts or remote access policies (both described later in this chapter), you need to understand the relationship between the two. Configuring the router user account includes the option of choosing whether to use remote access policies to grant the calling router access to the answering router. You can grant or deny permission for the calling router to access the answering router either at the user account level or at the remote access policies level. The permission specified in the user account overrides the permission specified in a remote access policy. However, if you choose the **Control access through Remote Access Policy** option in the user account that the answering router uses to authenticate the calling router, the remote access policy permission specified on Properties page for the remote access policy governs whether the user account of the calling router is granted or denied access. This option is available only for accounts on stand-alone routers or members of a native mode Active Directory domain. For more information about remote access policies, see "Choosing a Remote Access Policy Type" later in this chapter.

To allow or reject connection attempts according to a variety of criteria, you can specify several remote access options in the user account of the calling router and multiple additional options by using remote access policies. This level of precision enhances the security of your site-to-site connection by providing great flexibility in how you can control access to the answering router and its network resources.

You can configure router user accounts individually for each router or by adding router accounts to an Active Directory group.

# Configuring Router User Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Successfully configuring the router's user account includes the following tasks:

- Configure the router user account name to match the demand-dial interface name.
- Configure the authentication provider for the router.
- Choose an Active Directory user account or a local user account.
- Choose dial-in options for the router user account.
- Ensure that the calling router can establish a connection.

**Configure the Router User Account Name to Match the Demand-Dial Interface Name**

The name of a demand-dial interface configured on the answering router must match the name of the user account that the answering router uses to authenticate the calling router (unless, for a one-way initiated connection, you do not create a demand-dial interface on the answering router, instead configuring static routes in the calling router's user account). For example, if you have one headquarters site and multiple branch offices, on the headquarters answering router configure a separate demand-dial interface for the calling router at each branch-office site. This creates a one-to-one relationship between each calling router's user account name and the name of its corresponding demand-dial interface on the answering router.

**Caution**

- If the user account name does not match the name of the demand-dial interface, the answering router identifies the calling router as a remote access client rather than as a demand-dial router. In this case, the connection might not work correctly. For this reason, do not change the user account name without also changing the corresponding demand-dial interface name.

Table 10.8 shows an example configuration of the demand-dial interfaces and user accounts for a two-way initiated VPN connection. The example in Table 10.8 uses Active Directory user accounts for the routers. In some cases, you can alternatively use a local user account.

**Table 10.8 Example Configuration of Demand-Dial Interfaces and User Accounts for a Two-Way Initiated VPN Connection**

| Router | Demand-Dial Interface | User Account for Calling Router |
|---|---|---|
| Router located in Seattle<br><br>(Functions as both a calling and an answering router) | Interface Name: VPN_NewYork<br><br>Destination IP Address: 131.107.0.1<br><br>Credentials: VPN_Seattle<br><br>Password: PasswordOne | VPN_NewYork<br><br>PasswordTwo |
| Router located in New York<br><br>(Functions as both a calling and an answering router) | Interface Name: VPN_Seattle<br><br>Destination IP Address: 157.54.0.1<br><br>Credentials: VPN_NewYork<br><br>Password: PasswordTwo | VPN_Seattle<br><br>PasswordOne |

For a two-way initiated VPN connection, such as the example in Table 10.8, the name of the demand-dial interface for each router must match the user name in the remote router's user credentials so that each router knows that a requested connection is for a site-to-site connection rather than for a remote access connection. The destination address of the remote router is configured on the demand-dial interface so that each router knows what address to connect to for the site-to-site

connection. Each router uses the user account (the credentials and password of the remote router) to authenticate the remote router when it attempts to establish a connection.

Table 10.9 shows an example configuration of the demand-dial interfaces and user accounts for a one-way initiated dial-up connection in which two branch offices — Portland and Olympia — call the headquarters office in Seattle. In this case, because the connection is a one-way initiated connection, no demand-dial interface is required on the answering router. Therefore, the "rule" that the calling router's user account name must match the name of a demand-dial interface on the answering router does not apply. If you do not create a demand-dial interface on the answering router, you use the calling router's local or Active Directory user account to configure static routes that identify the network IDs of the calling router's site.

**Table 10.9 Example Configuration for a Dial-up One-Way Initiated Connection**

| Router | Demand-Dial Interface | Calling Router User Accounts |
|---|---|---|
| Answering router in Seattle | No demand-dial interface needed.<br><br>If you do not configure a demand-dial interface on the answering router, you must configure static routes specifying the network IDs of the calling router's site on the **Dial-in** tab of the calling router's user account. | DD_Olympia<br><br>PasswordOne<br><br>-and-<br><br>DD_Portland<br><br>PasswordTwo |
| Calling router in Olympia | Interface Name: DD_Seattle<br><br>Modem or adapter: Modem on COM1<br><br>Phone number: 555-0111<br><br>Credentials: DD_Olympia / PasswordOne | No user accounts needed, because the Seattle router does not call the routers in Olympia or Portland. |
| Calling router in Portland | Interface Name: DD_Seattle<br><br>Modem or adapter: Modem on COM2<br><br>Phone number: 555-0111<br><br>Credentials: DD_Portland / PasswordTwo | No user accounts needed, because the Seattle router does not call the routers in Olympia or Portland. |

**Configure the Authentication Provider for the Router**

The answering router uses the calling router's credentials to authenticate and authorize the calling router. As described earlier, if the answering router is configured for Windows authentication, as is the case for a site-to-site only connection, Windows security is used to authenticate the credentials of the calling router. If the answering router is configured for RADIUS authentication, as might be the case if the same answering router is used to support both a site-to-site connection and remote access users, the RADIUS server verifies the credentials of the calling router.

The user accounts and groups that Windows and IAS can use to authenticate and authorize site-to-site dial-up or VPN connection attempts are those available in Windows Server 2003 or Windows 2000 native-mode or mixed-mode Active Directory domains and Windows NT 4.0 domains.

**Choose an Active Directory User Account or a Local User Account**

The answering router must be able to access an Active Directory user account or a local user account to verify the calling router's credentials. In either case, because demand-dial routers make connections automatically, the user account that the answering router uses to authenticate and authorize the calling router must not prompt for human intervention.

**Active Directory user account**

If your demand-dial routers belong to an Active Directory domain, you use the Active Directory Users and Computers snap-in (rather than the Demand-Dial Interface Wizard) to create a user account for each calling router before you deploy a dial-up or VPN connection. When you create the router account, be sure to clear the **User must change password at next logon** check box and select the **Password never expires** check box.

If you use EAP-TLS user authentication with Windows as the authentication provider, the answering router must belong to the Active Directory domain, and you must create an Active Directory user account for the calling router. (If you use EAP-TLS user authentication with RADIUS as the authentication provider, you must join the IAS server to the Active Directory domain.)

If you use RADIUS authentication, the RADIUS server can use either an Active Directory user account or a local user account to

authenticate the calling router. If you use a Windows Server 2003 IAS server as your RADIUS server, Microsoft recommends that you use Active Directory accounts.

For information about how to create an Active Directory user account, see "Create a new user account" in Help and Support Center for Windows Server 2003.

**Local user account**

If you use Windows authentication and your demand-dial routers do not belong to an Active Directory domain, the answering router uses a local user account to authenticate the calling router. This user account is created locally on the answering router.

Instead of creating the user account for the calling router manually, you use the Demand-Dial Interface Wizard to create the account. When you run the Demand-Dial Interface Wizard on an answering router that is not joined to an Active Directory domain, the wizard adds a local account for the calling router when you select **Add a user account so a remote router can dial in** on the **Protocols and Security** page. The wizard clears the **User must change password at next logon** check box and selects the **Password never expires** check box on the user account object.

For information about using the Demand-Dial Interface Wizard, see "Configure the Routing and Remote Access Service and Demand-Dial Interfaces" later in this chapter, and see Add a demand-dial interface in Help and Support Center for Windows Server 2003.

If you use a local account and password, you must use MS-CHAP v2 rather than EAP-TLS for user authentication.

If you use RADIUS authentication, the RADIUS server can use a local user account to authenticate the calling router. If you use a Windows Server 2003 IAS server as your RADIUS server, however, using Active Directory accounts is recommended.

**Choose Dial-in Options for the Router User Account**

You use the **Dial-in** tab of the calling router's user account to specify the remote access permission as follows:

- **Allow access (for native-mode or mixed-mode domains).** Allows access when a connection is attempted. (When you use the Demand-Dial Interface Wizard to create a local account, the remote access permission is set to **Allow access** by default.) This option overrides the grant or deny remote access permission setting specified on the Properties page of any associated remote access policy. The remote access permission on the associated remote access policy is either **Deny remote access permission** or **Grant remote access permission**. Thus, for example, if you select **Allow access** on the router user account, and if you select **Deny remote access permission** on the remote access policy, the connection attempt succeeds.

- **Deny access.** Denies access when a connection is attempted. This option also overrides the remote access permission configured on the remote access policy.

- **Control access through Remote Access Policy (available only for native-mode domains).** With this option selected, the grant or deny remote access permission setting specified on the Properties page of any associated remote access policy is used. Thus, for example, if you select **Control access through Remote Access Policy** on the router user account, and if you select **Deny remote access permission** on the remote access policy, the connection attempt is blocked.

  **Note**

  - If the user account is in a Windows Server 2003 or Windows 2000 mixed-mode domain, the **Control access through Remote Access Policy** setting is not available, and you must set the remote access permission in each calling router's user account to **Allow access**.

Even when remote access permission is allowed, the connection attempt still can be denied on the basis of other user account properties, such as callback options configured on the **Dial-in** tab, or it can be denied on the basis of remote access policy conditions or profile settings. Authorization of the site-to-site connection is controlled by a combination of the account dial-in properties (which includes the account remote access permission), the policy remote access permission, and the profile constraints configured on the remote access policy.

You can configure the following settings on the user account **Dial-in** tab:

- **Verify Caller ID (for dial-up router only).** Typically, you use this option to increase security when accepting dial-in connections from telecommuters, but you can also use it for a calling router. Selecting this option requires that the calling router always call from the phone number that you type. Windows Server 2003 uses the caller ID functionality of

the modem and the phone connection to verify that the dial-up router is in fact using the correct number; if the number matches, the connection is allowed. This option is not available if the answering router is in a Windows NT 4.0 domain or a Windows Server 2003 or Windows 2000 mixed-mode domain.

**Note**

- For caller ID verification to work, the calling router, the answering router, and the phone system that connects them must all support caller ID. In addition to supporting the caller ID feature, the calling and answering routers must have the appropriate drivers to support the transmission of caller ID information to the Routing and Remote Access service.

- **Callback Options.** Typically, these options are used for managing individual dial-up remote access users rather than demand-dial calling router user accounts. However, you can also configure callback options for a dial-up demand-dial connection. For example, you can specify that the main office router call the branch office router back if the main office has fixed or low-cost long-distance rates and the branch office has higher long-distance rates.

- **Assign a static IP address (for dial-up or VPN routers).** Use this option to cause the remote router to assign a static IP address to this router when a connection is made. For more information about assigning IP addresses, see "IP Address Assignment for the Logical Interface" later in this chapter. You can increase security by specifying which IP address the router must use. This option is not available in a Windows NT 4.0 domain or a Windows Server 2003 or Windows 2000 mixed-mode domain.

  If the static IP address that you specify is not valid or is already in use by another calling router or by a remote access user, the answering router assigns a dynamic IP address to the calling router.

- **Apply static routes (for dial-up or VPN routers).** This option is designed for one-way initiated site-to-site connections. You select **Apply static routes** to add static IP routes (network IDs) for the calling router's site to the routing table of the answering router when a connection is made. This lets the answering router forward packets to all addresses on the calling router's intranet across the connection to the calling router. This option is not available if the answering router is in a Windows NT 4.0 domain or a Windows 2000 mixed-mode domain.

#### Ensure That the Calling Router Can Establish a Connection

To ensure that the router user account is not prevented from establishing a site-to-site connection, verify that the following options are configured as appropriate:

- When you create the calling router user account, clear the option **User must change password at next logon**. If you select **User must change password at next logon**, the site-to-site connection attempt fails, because changing the password while attempting to make the connection is an interactive process. For information about how to create a user account, see "Create a new user account" in Help and Support Center for Windows Server 2003.

- For a site-to-site connection, do not use the remote access account lockout feature. If the router user account is disabled (locked out by using the remote access account lockout feature), the connection attempt is rejected. For information about the remote access account lockout feature, see Remote access account lockout in Help and Support Center for Windows Server 2003.

- If you block the calling router from establishing a connection during certain times, make sure that the times specified are appropriate for your organization. If the router user account is not permitted to log on during the time specified for the site-to-site connection, the connection attempt is rejected. For information about how to use remote access policies to configure dial-in hours on an answering router in order to block incoming connections from a calling router at specific times of the day or week (such as at night or on weekends), see Configure dial-in constraints in Help and Support Center for Windows Server 2003.

# Configuring Router Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Adding demand-dial router user accounts to an Active Directory group reserved for demand-dial routers simplifies administration by letting you centrally manage the list of demand-dial routers on your network. If you decide to manage authorization by group rather than by each router's individual user account, you must set the remote access permission on each calling router's user account to either **Control access through Remote Access Policy** or **Allow access** and then create a remote access policy based on connection type and group membership.

For example, if multiple calling routers will use a VPN connection, you can create an Active Directory global group called VPN-Routers and add the user account of each calling router to that group.

Then, create a remote access policy with two conditions:

- Set **NAS-Port-Type** to **Virtual (VPN)**. A network access server (NAS) is a server that accepts point-to-point connections from a calling router, or other remote client, and then acts as a gateway to the network for the calling router; the NAS-Port-Type is the media type that a calling router uses to access the site of the answering router).

- Set **Windows-Group** to **VPN-Routers**.

Finally, configure the profile for the remote access policy, selecting an authentication method and encryption strength.

# Choosing a Remote Access Policy Type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use remote access policy settings to validate a variety of connection settings *before* a connection is authorized and to specify a variety of connection restrictions *after* the connection is authorized. For example, you can use remote access policies to allow or reject connection attempts based on membership in an Active Directory group or by time of day or day of week; to require a specific authentication method and encryption strength; or to limit the connection based on bandwidth.

The computer on which you create a remote access policy is determined by which authentication provider you use for site-to-site connections. If you use Windows authentication, you create and manage remote access policies on each answering router. If you use RADIUS authentication, as you might if your answering router supports both a site-to-site connection and home or mobile users, you can configure and manage all remote access policies centrally on the IAS server that provides RADIUS authentication for multiple answering routers.

If you use RADIUS as the authentication provider in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition, you can configure a RADIUS attribute to ignore the dial-in properties of a user account in the profile properties of a remote access policy. To support the multiple types of connections for which IAS provides authentication and authorization, you might need to disable the processing of user account dial-in properties. For more information about dial-in properties, see "Dial-in properties of a user accountDial-in properties of a user account" and "Add RADIUS attributes to a remote access policy" in Help and Support Center for Windows Server 2003.

You can use one of three types of remote access policies:

- **Common policy.** You can create a common policy, which uses typical settings for a particular access method.

- **Custom policy.** You can create a custom policy, which lets you specify a detailed configuration for a particular access method. If you want to manage authorization and connection parameters other than by group or by type of connection, you must configure custom remote access policies. For more information about each possible setting for remote access policy conditions and profiles for a custom policy, see "Elements of a remote access policyElements of a remote access policy" in Help and Support Center for Windows Server 2003.

- **Default policy.** If enforcing a high level of security is not important for your organization, you can use one of the existing default policies. Two default remote access policies are created when you enable and configure the Routing and Remote Access service on a demand-dial router or when you install IAS: The **Connections to Microsoft Routing and Remote Access server** policy and the **Connections to other access servers** policy. To use the **Connections to Microsoft Routing and Remote Access server** policy for a site-to-site connection, all you have to do is change the remote access permission on the policy's Properties page to **Grant remote accesspermission**.

For more information about using Windows Server 2003 remote access policies, see "Remote Access Policies" in Help and Support Center for Windows Server 2003. For more information about using remote access policies with an Internet Authentication Service (IAS) server, see "Deploying Internet Authentication Service (IAS)" in this book and see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# Integrating the Remote Site Connection into Your Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can successfully deploy a dial-up or VPN remote site connection on your network, you must decide how you want to integrate the connection into your existing network infrastructure. The flowchart in Figure 10.7 shows the tasks required to integrate the connection into your network.

**Figure 10.7 Integrating the Remote Site Connection into Your Network**



# Designing the Routing Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Integrating a remote site connection into your existing routing infrastructure requires that you decide which choices you want to make about the following tasks:

- Adding static routes
- Using routing protocols
- Servicing Internet traffic
- Enabling multicast connectivity between sites

# Adding Static Routes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In some cases, instead of using routing protocols to dynamically update routing tables, you must configure one or more static routes on the intranet interfaces and demand-dial interfaces of the demand-dial routers in your site-to-site deployment. A static route, which creates a specific path to a destination IP address in an IP network, is one of a set of routes in a routing table that are permanent until changed by a network administrator or by an automatically scheduled auto-static update.

The following topics provide the information that you need to manage static routes for a site-to-site connection:

- Static routes for a site-to-site connection
- Auto-static updates
- Using on-subnet or off-subnet address ranges

**Static Routes for a Site-to-Site Connection**

You might need to create one or more the following types of static routes for your site-to-site connection:

- **LAN interface at both sites.** On both the calling and the answering router, configure a static route or routes on the LAN interface that connects the router to the local intranet. Include a static route for each subnet on the local area network.

  Alternatively, you can use a routing protocol instead of configuring static routes. For more information about using routing protocols, see "Using Routing Protocols" later in this chapter.

- **Demand-dial interface for the remote site.** On the calling router, configure a static route or routes on the demand-dial interface that connects the router to the remote site. Include a static route for each subnet in the answering router's network that you want users to be able to access (or you can use the default route).

  Alternatively, for a persistent site-to-site connection only, you can enable a routing protocol on the demand-dial interface instead of configuring static routes. For more information about using routing protocols, see "Using Routing Protocols" later in this chapter.

- **Demand-dial interface for the local ISP.** On the calling router — for a VPN connection in which the branch office router uses a temporary link to a local ISP only — you must configure a static host route on the demand-dial interface that connects to the local ISP. The destination that you specify for this static host route is the IP address of the answering router's Internet-connected interface; this IP address is assigned to the answering router by its local ISP (or by InterNIC).

- **Router user account.** For a one-way connection in which the answering router is a standalone router or a member of a native-mode Active Directory domain, you can omit creating a demand-dial interface on the answering router. In this case, you must configure a static route or routes in the calling router's user account that identify the network IDs of the calling router's site.

For information about how to configure these static routes, see "Configure Static Routes" later in this chapter. For general information about the difference between using static routes or routing protocols, see the discussion of developing routing strategies in "Designing a TCP/IP Network" in this book.

**Auto-static Updates**

You can add the static routes that correspond to the network IDs available across a demand-dial interface either manually or by using auto-static updates. An auto-static update is a one-time, one-way transfer of routing information. In contrast to the periodic announcements issued by routing protocols, an administrator must either issue a command to initiate a manual auto-static update or must schedule auto-static updates by running the update as a scheduled task.

When instructed, a demand-dial interface that is configured for auto-static updates sends a request across an active connection to request all of the routes of the router on the other side of the connection. In response to the request, all of the routes of the requested router are automatically entered as static routes in the routing table of the requesting router.

**Using On-Subnet or Off-Subnet Address Ranges**
If any of the static address ranges that you configure in the IP properties of the answering router is an off-subnet address range, you must add routes to the routing infrastructure in order for the *logical interfaces* of calling routers to be reachable. During the PPP negotiation, each router typically assigns an IP address to the logical interface of the other router. When a site-to-site connection is made, each router sends traffic to the other router using the logical interface that corresponds to the dial-up, PPTP, or L2TP port of the connection. For more information about how each router assigns an IP address to the other, see "IP Address Assignment for the Logical Interface" later in this chapter.

The method used to ensure the reachability of the logical interfaces in a site-to-site connection depends on how you configure each router to obtain IP addresses for calling routers (and for remote access clients, if your network also supports them). You use either an on-subnet or an off-subnet address range for these IP addresses.

**On-subnet address range**
An on-subnet address range is an address range of the subnet to which the answering router is attached. An on-subnet address range provides the IP address for a logical interface whenever the router is configured to use Dynamic Host Configuration Protocol (DHCP) to obtain IP addresses, a DHCP server is available, and the manually configured pool (or pools) of IP addresses are within the range of addresses of the attached subnet. If you use an on-subnet address range, no additional routing configuration is required.

**Off-subnet address range**
An off-subnet address range is an address range that represents a different subnet than the subnet to which the router is attached. Off-subnet addressing uses a separate subnet address space that is unique to the intranet. An off-subnet address range provides the IP address for a logical interface whenever the router is manually configured with a pool of IP addresses for a separate subnet.

If you use an off-subnet address range, you must add the route or routes that summarize the off-subnet address range to the intranet routing infrastructure so that traffic destined to the logical interfaces of connected routers are forwarded from the originating node to the local dial-up or VPN router and then sent by it to the appropriate connected remote router. You can add the routes that summarize the off-subnet address range to the routing infrastructure by using one of the following methods:

- Add static routes for the off-subnet address range that point to the dial-up or VPN router's intranet interface to the neighboring router. Configure the neighboring router to propagate this static route to other routers in the site by using the dynamic routing protocol used in your site.

- If the dial-up or VPN router uses Open Shortest Path First (OSPF) and participates as a dynamic router, you must configure the router as an autonomous system boundary router (ASBR) so that the static routes of the off-subnet address range are propagated to the other OSPF routers in the site.

If your site consists of a single subnet, and you use an off-subnet address range, you must do one of the following:

- Configure each intranet host for a persistent route (or routes) that points to the dial-up or VPN router's intranet interface. The route (or routes) expresses the off-subnet address range.

  In the Routing and Remote Access snap-in, you must configure address ranges with a starting and ending address. To simplify the set of routes needed to express the off-subnet address ranges, express each range as an IP address with a subnet mask. For more information about using an IP address and a mask to express an address range, see Expressing an IP address range with a mask in Help and Support Center for Windows Server 2003.

- Configure each intranet host with the IP address of the intranet-connected interface of the dial-up or VPN router as its default gateway.

Therefore, if your site consists of a single subnet, it is more efficient to use an on-subnet address range.

# Using Routing Protocols

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you use the Routing Information Protocol version 1 or 2 (RIPv1 or RIPv2) or the OSPF routing protocol in your site, you can add and configure the RIP or OSPF routing protocol components of the Routing and Remote Access service so that a demand-dial router participates in the propagation of routing information as a dynamic router. As an efficient alternative to manually configuring static routes, you can use routing protocols on each LAN interface and on each demand-dial interface that is used for a persistent site-to-site connection. You must configure dynamic routing protocols for each new router so that it can participate in the dynamic routing architecture and share its subnets.

If you use a routing protocol other than RIP or OSPF, such as the Cisco proprietary Interior Gateway Routing Protocol (IGRP) or Enhanced Interior Gateway Routing Protocol (EIGRP), configure the neighboring Cisco router for RIP or OSPF on the interface connected to the subnet that contains the dial-up or VPN router, and then configure IGRP or EIGRP on all other interfaces. You must also configure the neighboring Cisco router to redistribute the IGRP or EIGRP routes into the OSPF or RIP routing tables.

Do *not* use routing protocols across an on-demand site-to-site connection. The periodic announcements that most routing protocols use to propagate routing information cause one demand-dial router to call another each time an announcement occurs. For example, RIPv1, by default, announces routing information every 30 seconds. If your router incurs a long-distance charge every 30 seconds, the phone bill that results defeats the cost savings that an on-demand link is designed to provide. Instead, add routes for network IDs that are available across the demand-dial interface as static routes to the routing tables of demand-dial routers.

# Servicing Internet Traffic

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Depending on the immediate task at hand, users in an organization that includes more than one site might want to access people or resources at the remote site, or they might want to access the Internet. You can use a demand-dial router to handle both types of traffic. If you do use the demand-dial router to enable users to access the Internet, you must decide whether you want users at a branch office to access the Internet through the main office or to access the Internet directly from the branch office.

**For security, route branch office Internet traffic through the main office**
If you deploy Microsoft® Internet Security and Acceleration (ISA) Server at the main office, and you want branch office Internet traffic to be protected by this system, configure branch office Internet traffic to go over the dial-up or VPN connection to the ISA server at the main office. ISA Server is an integrated firewall and Internet caching server that can also act as a VPN router to provide Internet access that is both secure and fast. For information about deploying ISA Server, see "Deploying ISA Server" in this book.

**For performance, route branch office Internet traffic directly to the Internet**
To give branch office users faster access to the Internet than is possible if the Internet traffic must travel to the main office and back, configure branch office Internet traffic to go directly out to the Internet through the demand-dial router. In addition, if you use an on-demand connection rather than a persistent connection, you must provide direct access to the Internet through the demand-dial router rather than over the link to the main office. Configuring client computers to allow users in the branch office to access the Internet directly through the demand-dial router is known as *split tunneling*.

For information about handling each alternative, see "Configure Internet Access Through the Calling Router" later in this chapter.

# Enabling Multicast Connectivity Between Sites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An increasing number of organizations use multicast communication for such multiple-user applications as video conferencing, collaborative computing, and distance learning. Computers running Windows Server 2003 can both send and receive IP multicast traffic. The Routing and Remote Access service includes the Windows Server 2003 Internet Group Management Protocol (IGMP) routing protocol component, which you can configure with IGMP router mode and IGMP proxy mode to enable the exchange of multicast packets between remote sites. The IGMP Router and Proxy routing protocol is not itself a multicast routing protocol.

When you make your demand-dial routers multicast-capable, they can listen for all multicast traffic over all site-to-site connections, they can listen for IGMP Membership Report messages, and they can update the TCP/IP multicast forwarding table so that routers can determine where to forward incoming multicast traffic.

For a site-to-site connection, use the Routing and Remote Access snap-in to enable IGMP with IGMP proxy mode on the demand-dial interface on both the calling router and the answering router, and to enable IGMP with IGMP router mode on all other interfaces internal to the site.

For more information about IP multicasting, see "Designing a TCP/IP Network" in this book and the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit). For more information about multiple supported multicast configurations, see the *Internetworking Collection* of the *Windows Server 2003 Technical Reference* (or see the Internetworking Collection on the Web at http://www.microsoft.com/reskit).

# Planning IP Address Assignment and Name Resolution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The DHCP and name resolution issues that you must consider when planning how to integrate a site-to-site connection into your existing network include:

- Client IP address assignment
- IP address assignment for the logical interface
- Using IP addresses to avoid name resolution issues
- Using name resolution when accessing other services on a VPN router

# Client IP Address Assignment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, when you have slower WAN links or when you have dial-up links, you place a DHCP server on each side of a remote site connection. Installing the DHCP service on a calling router or other server at a branch office lets branch office clients obtain IP addresses from that server.

However, a DHCP server located in one subnet can support other subnets. If no DHCP server is located at the branch office, when a client at the branch office requests an address, the calling router can forward the client's request across the dial-up or VPN link to the DHCP server in the main office, which then leases the client an address. For this to work, ensure the following:

- The calling router contains the DHCP Relay Agent (the DHCP Relay Agent component is added with the internal interface by default).
- The calling router is configured with the DHCP server's IP address.
- The site-to-site connection is a persistent connection.

For an on-demand connection, if the DHCP Relay Agent is installed on a branch office router, this can cause the router to dial the main office router (for a non-VPN dial-up connection) or the local ISP (for a VPN connection that uses a dial-up link to the ISP) every time a DHCP packet is sent by a computer in the branch office network. Therefore, if you deploy an on-demand connection, install the DHCP service in the branch office and uninstall the DCHP Relay Agent on the calling router.

For more information about deploying DHCP, see "Deploying DHCP" in this book.

# IP Address Assignment for the Logical Interface

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a calling router initiates a connection, the router creates a temporary logical interface (also known as a *virtual interface* or a *virtual network adapter*) and requests that the answering router assign an IP address to this logical interface. The answering router then creates a logical interface and requests an IP address for itself from the calling router. The logical interface on the calling router connects to the logical interface on the answering router to form the demand-dial connection. The IP address assignment for each demand-dial interface lasts for the duration of the connection. These IP addresses can be either private or public IP addresses. If both requests are successful, the logical interface on the PPP connection for each router is assigned an IP address from the other router. This is known as a *numbered connection*. In the absence of a numbered connection, a site-to-site connection can also use an unnumbered connection.

**Numbered Connection**

A numbered IP address assignment can occur in one of the following ways:

- **Dynamic IP addresses allocated by a DHCP server.** The answering router obtains the IP address to assign to a calling router from a DHCP server. This is the default method for IP address assignment. If no DHCP server is available, the router uses an address from the Automatic Private IP Addressing (APIPA) range 169.254.0.1–169.254.255.254. You must install a DHCP server on each network that contains an answering router.

- **Dynamic IP addresses allocated from a static address pool configured on the answering router.** The answering router obtains the IP address from a static pool of addresses. If you configure a static address pool, be sure to use only IP addresses that are not in a range that your DHCP server might assign to another computer and that are not already assigned to specific computers. The pool can be ranges of addresses that are a subset of addresses from the IP network to which the server is attached or from a separate subnet. As described earlier, if the static IP address pool address ranges represent a different subnet, ensure that routes to the address ranges exist on the routers of your intranet so that traffic to the logical interface of a calling router is forwarded to the remote server.

- **Static IP address specified in the calling router user account.** You configure a static IP address on the router user account **Dial-in** tab for both the calling and the answering router. In this case, when a calling router initiates a connection, creates a temporary logical interface, and requests that the answering router assign an IP address to this logical interface, the answering router assigns the IP address specified in the calling router's user account. The answering router then creates a logical interface and requests an IP address for itself from the calling router. The calling router assigns the IP address specified in the answering router's user account.

**Unnumbered Connection**

Site-to-site connections in Windows Server 2003 and Windows 2000 do not require a numbered connection (Windows NT 4.0 RRAS connections do require a numbered connection). If one of the routers rejects the request for an IP address during the connection establishment process between a calling and an answering router (because no addresses are available to assign, probably due to a misconfiguration), a connection is still established. In this case, the logical interface on the PPP connection does not have an assigned IP address. This is known as an *unnumbered connection*.

The routing protocols provided with Routing and Remote Access cannot operate over unnumbered connections. Therefore, you must use static routing if you use unnumbered connections.

# Using IP Addresses to Avoid Name Resolution Issues

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you configure a demand-dial interface for a dial-up calling router, you provide the phone number of the answering router with which this router establishes a connection. In this case, name resolution is not an issue.

For VPN connections, however, the corresponding entry on the demand-dial interface is the destination address of the answering router. In this case, name resolution can be an issue, because you can provide either the Domain Name System (DNS) name or the IP address of the answering router. Microsoft recommends that you provide the IP address rather than the answering router name when you configure a demand-dial interface so that resolution of the name to the IP address by means of a DNS or Windows Internet Name Service (WINS) server is not needed.

If you do use names for demand-dial interfaces, make sure that an appropriate DNS record exists on the DNS server that hosts the public namespace that is visible to Internet users and computers, or on the DNS server of your ISP, so that the DNS names of your answering routers can be resolved.

# Using Name Resolution When Accessing Other Services on a VPN Router

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In some organizations, you might want users to be able to access other services, such as file shares, on the answering VPN router. For this type of configuration, if you specify the DNS name rather than the IP address in the demand-dial interface, and the name resolves to the public IP address of the answering router, traffic sent to the services running on the router is sent in clear text (unencrypted) across the Internet. It is not encapsulated, encrypted, and sent using the VPN connection, which can compromise security.

A related problem is that if you configure packet filters on the answering router to allow only traffic over a VPN connection, all other traffic is discarded. Attempts to connect to services running on the answering router fail in this situation, because traffic attempting to connect to those services is not sent over the site-to-site VPN connection.

If the site DNS and WINS servers do not contain a record mapping the name of the VPN router to its public IP address, traffic to services running on the VPN router is always sent across the VPN connection. To ensure that the name of the VPN router is always resolved to the private or site IP address of the VPN router, disable DNS dynamic update and NetBIOS over TCP/IP (NetBT) on the Internet-connected interface (or interfaces) of the VPN router from the properties of the Internet connection in Network Connections as follows:

- **Prevent DNS name resolution.** To prevent a VPN router from dynamically registering the public IP address of its Internet interface on the site DNS servers, on the Internet interface of the router, configure the properties of **Internet Protocol (TCP/IP)** by clicking the **Advanced** button, selecting the **DNS** tab, and then clearing the **Register this connection's addresses in DNS** check box.

- **Prevent WINS name resolution.** To prevent a VPN router from dynamically registering the public IP address of its Internet interface on the site WINS servers, on the Internet interface of the router, configure the properties of **Internet Protocol (TCP/IP)** by clicking the **Advanced** button, selecting the **WINS** tab, and then selecting the option **Disable NetBIOS over TCP/IP**.

  **Note**

  - By default, the Routing and Remote Access Wizard clears **Register this connection's addresses in DNS** and selects **Disable NetBIOS over TCP/IP**. Be sure not to change these defaults if you want users to be able to access services such as file shares on the answering VPN router.

For more information about name resolution, see "Deploying Domain Name System (DNS)" and "Deploying WINS" in this book.

# Planning Active Directory Integration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Integrating a remote site connection into an Active Directory–based network requires you to decide which choices you want to make about the following tasks:

- Put a domain controller at the branch office.
- Use one domain to include geographically remote sites.
- Use scheduled replication or reciprocal replication.
- Join demand-dial routers to the Active Directory domain.

**Putting a Domain Controller at the Branch Office**
If you deploy a persistent site-to-site connection between a branch office and a main office, you might not need a domain controller at the branch office. Branch office users can access a domain controller in the main office when they log on to their computers or use other Active Directory services.

For an on-demand connection, install a domain controller at the remote site.

**Using One Domain to Include Geographically Remote Sites**
You can include a main office and a branch office in one Active Directory domain. However, geographically remote sites must not share the same address space and must have separate Active Directory sites. You must create a separate Active Directory site for the branch office and create a child object for the branch office, providing the appropriate network ID and subnet mask for the branch office site.

For more information about deploying Active Directory, see "Deploy Active Directory" later in this chapter.

**Using Scheduled Replication or Reciprocal Replication**
Typically, domain controllers have a constantly available connection so that all domain controllers obtain a steady flow of updated directory information. If you have domain controllers in sites that are connected by a site-to-site connection, you must ensure that replication takes place. Directory updates can be exchanged through a site-to-site connection in one of two ways:

- **Scheduled replication.** On a persistent site-to-site connection, you can schedule replication to take place at specified intervals.
- **Reciprocal replication.** For a one-way initiated on-demand connection in which no constantly available connection exists between domain controllers in the two sites, you must enable reciprocal replication. With reciprocal replication, all replication occurs simultaneously between the domain controllers in the two sites, and the connection is closed when replication is complete. Reciprocal replication maximizes the efficiency of directory information exchange while minimizing connection time and eliminating timeout errors that can occur if the main site domain controller requests changes from the branch site domain controller when the connection is not available. You can configure reciprocal replication on a site link or on a connection.

For more information, see "Configure Replication for Active Directory" later in this chapter.

**Joining Demand-Dial Routers to the Active Directory Domain**
In an Active Directory domain, you can choose either to join a demand-dial router computer to the domain or not to, based on the following factors:

- If your answering router uses Active Directory user accounts to authenticate and authorize a calling router, you must join the answering and calling routers to the domain.
- If you use EAP-TLS user authentication with Windows as the authentication provider, you must join the answering router (which is the authenticating server) to the Active Directory domain. If you use EAP-TLS user authentication with RADIUS as the authentication provider, you must join the IAS server (which is the authenticating server) to the Active Directory

domain. With RADIUS authentication, the answering router does not need to be joined to the domain. (Use Windows authentication for a site-to-site only connection. You might use RADIUS authentication for a site-to-site connection if the answering router also supports remote access users.)

- If you use L2TP/IPSec with computer certificates, the demand-dial routers are not required to join the Active Directory domain. However, the Windows Server 2003 PKI uses Active Directory, if it is installed, to store certificates, certificate revocation lists, and delta certificate revocation lists and to publish root CA certificates and cross-certificates. Using Active Directory makes this information easy to locate from anywhere on the network. (If you use L2TP/IPSec with preshared keys, you are not required to join the VPN routers to the Active Directory domain.)

For more information about planning and deploying Active Directory for an organization with multiple branch office sites, see the Active Directory Branch Office Planning Guide link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about deploying Active Directory domains and sites and Active Directory replication between sites, see "Designing the Site Topology" in *Designing and Deploying Directory and Security Services* of this kit.

# Preparing for Server Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have existing Windows NT 4.0 or Windows 2000 site-to-site connections, you must decide whether to upgrade your routers to Windows Server 2003 and migrate to the Windows Server 2003 Routing and Remote Access service, or whether to continue to support your current configuration. Before you install the Routing and Remote Access service on your demand-dial routers, you must decide how many routers you need, and you must make sure that the computers meet Windows Server 2003 requirements.

For more information about creating hardware and software inventories, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit. The flowchart in Figure 10.8 shows the tasks required when preparing to configure the router server.

**Figure 10.8 Preparing for Server Configuration**



# Migrating Routers from Windows NT 4.0 or Windows 2000

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have an existing site-to-site connection between remote offices using Windows NT 4.0–based or Windows 2000–based servers and plan to upgrade most of your network to Windows Server 2003, Windows Server 2003 can support your existing Routing and Remote Access or RRAS servers. Alternatively, you can take advantage of the new features in Windows Server 2003 by upgrading your demand-dial routers.

The following topics can help you decide whether to upgrade your demand-dial routers:

- New features
- Migrating router settings

**New Features**

In Windows NT Server 4.0, routing and remote access are separate services. In Windows 2000 Server and Windows Server 2003, these functions are combined in the single Routing and Remote Access service. Table 10.10 lists new features available in Windows Server 2003 and Windows Server 2000.

**Table 10.10 New Features for Dial-up or VPN Site-to-Site Connections Since Windows NT Server 4.0 RRAS**

| Windows Release | New Features |
|---|---|
| Windows 2000 Server | • **L2TP/IPSec.** An L2TP/IPSec VPN tunnel provides stronger security than a PPTP VPN tunnel.<br>• **Remote access policies.** This feature gives administrators more flexibility in setting remote access permissions and connection restraints.<br>• **MS-CHAP v2.** This type of password–based user authentication provides a stronger alternative to MS-CHAP. (MS-CHAP v2 is also available in Windows NT 4.0 SP4.)<br>• **EAP.** Support for EAP lets you use installable authentication methods, such as EAP-TLS certificate-based user authentication, which is stronger than password-based user authentication. |

| | |
|---|---|
| Windows Server 2003 | • **Improved wizards and snap-in.** The Routing and Remote Access and Demand-Dial Interface wizards and the Routing and Remote Access snap-in are easier to use.<br><br>• **Intranet and Internet-connected interface improvements.** By default, the Routing and Remote Access service now disables dynamic update on the intranet interface and disables dynamic update and NetBIOS over TCP/IP on the Internet-connected interface to ensure correct name resolution of the router and to ensure access to services running on the router.<br><br>• **Preshared keys.** Support for configuring preshared keys using the Routing and Remote Access snap-in provides an alternative to computer certificates in L2TP/IPSec authentication.<br><br>• **NAT/Basic Firewall configuration.** You can use Manage Your Server to configure the NAT/Basic Firewall component of Routing and Remote Access. NAT integration with static and dynamic packet filtering lets you configure NAT interfaces to work with Basic Firewall or with incoming or outgoing static packet filters.<br><br>• **NAT-T.** IPSec NAT traversal (NAT-T) lets you create L2TP/IPSec connections from a calling or answering router that is located behind one or more NATs.<br><br>• **PPPoE.** Support for PPPoE lets a small business use NAT/Basic Firewall and their broadband Internet connection to connect a branch office to their local ISP. Using PPPoE for an on-demand connection is faster than using a dial-up link to connect to the ISP. |

**Migrating Router Settings**

When you upgrade from Windows NT 4.0 or Windows 2000 to Windows Server 2003, you retain all IP-based routing configuration, including demand-dial, RIP, OSPF, and DHCP Relay Agent settings. However, Windows Server 2003 does not support the NetWare routing protocol Internetwork Packet Exchange (IPX). If you upgrade from Windows NT 4.0 to Windows Server 2003, and IPX settings are detected, you are provided the option not to upgrade after all.

# Planning Server Capacity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Table 10.11 provides capacity planning information that you can use to help determine how many demand-dial servers you need to deploy and how much data throughput your site-to-site connection can support.

**Table 10.11 Capacity Planning**

| Factor | Capacity |
|---|---|
| Number of connections | For two-way connections, one answering router supports 10 simultaneous calling router connections before performance begins to degrade.<br>• For one-way connections, one answering router supports 100 simultaneous calling router connections before performance begins to degrade. |
| Data throughput | The amount of data throughput that a site-to-site connection can support depends, in part, on what resources the users are using on the network. Other factors that affect data throughput include:<br>• **IPSec offload card.** You can increase data throughput for L2TP/IPSec by installing an IPSec offload card. Using an IPSec offload card and a dual processor lets a server process more than 50 Mbps of fully encrypted data. To install an IPSec offload card, follow the manufacturer's instructions.<br>• **Compression.** Using Microsoft Point-to-Point Compression (MPPC) decreases data throughput. To increase data throughput, turn off MPPC. You can turn off MPPC by using one of the following methods:<br>  • **To clear compression on a specific demand-dial interface.** In Routing and Remote Access, in the console tree, click **Network Interfaces**, right-click the demand-dial interface on which you want to clear compression, and then click **Properties**. On the demand-dial interface Properties page, click the **Networking** tab, click **Settings**, and then in the **PPP Settings** dialog box, clear the check box **Enable software compression**.<br>  • **To clear compression for all types of PPP connections.** In Routing and Remote Access, in the console tree, right- click the demand-dial server icon, click **Properties**, click the **PPP** tab, and then clear the check box **Software compression**. This disables compression both for site-to-site connections and for remote access connections. |

# Planning Server Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following information can help you plan how to set up your server before you deploy the remote site connection:

- Meeting server requirements
- Disabling unused services
- Planning physical and administrative security

# Meeting Server Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Table 10.12 lists the minimum hardware and software requirements for a demand-dial router.

**Table 10.12 Hardware and Software Requirements for a Demand-Dial Router**

| Component | Requirement |
|---|---|
| Processor | Pentium 233 MHz processor (550 MHz recommended) |
| Memory | 128 MB RAM (256 MB recommended) |
| Hard drive | 4 GB hard drive |
| LAN adapter | A network adapter connected to the intranet. The adapter must have a driver that displays the Designed for Windows logo. The server must have multiple network adapters or a single network adapter configured with multiple IP addresses. |
| WAN adapter | • **For a dial-up site-to-site connection:** An ISDN adapter, analog modem, or other physical device that is connected to the line that connects the two sites.<br>• **For a VPN site-to-site connection:** A network adapter that is connected to the Internet, either directly or through a perimeter network. Typically, this is a T-Carrier or Frame Relay adapter, or a DSL or cable modem. |
| Software | The Windows Server 2003 operating system. Windows Server 2003 includes the Routing and Remote Access service, which you must enable. |

Before you deploy a dial-up or VPN router, install and configure the required hardware and the appropriate drivers, and test whether each one functions properly. To determine if the hardware in your organization is certified and compatible, see the Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Disabling Unused Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Disabling services that you do not use on your demand-dial server has two advantages: It returns resources to the server that other components can use, and it makes your backend network more secure by shutting off services that are potential entry points for attackers trying to break into your network.

The following is a list of services that you might be able to disable on a server that you plan to use as a demand-dial router:

- Distributed File System (DFS)
- Distributed Transaction Coordinator
- Fax Services
- File Replication service (FRS)
- Indexing Service
- Internet Connection Sharing (ICS)
- Intersite Messaging
- Kerberos Key Distribution Center
- License Logging Service
- Print Spooler
- Task Scheduler
- Telnet
- Windows Installer

**Important**

- Do not disable the remote registry service. If you disable it, the demand-dial router cannot operate correctly.

# Planning Physical and Administrative Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To prevent the intentional or unintentional modification of your router configuration, use the following guidelines to secure your routers:

- Keep router computers physically secure from unauthorized users or intruders, for example, by placing them in a locked room.
- Delegate administration rights and permissions for the Routing and Remote Access service to a limited number of individuals.
- Assign administrator rights to an Active Directory group whose role is to administer remote site connections, and add only authorized users to that group. You can easily update changes to the group membership, whereas, if you do not use a group, you must administer each Administrator account separately on each Routing and Remote Access server.
- On both the calling router and the answering router, rename the local Administrator account and use strong passwords for the account.

# Deploying a Site-to-Site Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you make all the necessary design decisions, you can deploy a remote site connection for your organization. The flowchart in Figure 10.9 shows the tasks that can help you to successfully deploy a site-to-site connection.

**Figure 10.9 Deploying a Site-to-Site Connection**




If you want to try a test deployment to prepare for your production deployment, see "Perform a Test Deployment in Your Lab" later in this chapter. If you are ready to connect two sites in your production environment now, see "Deploy a Remote Site Connection" later in this chapter.

# Perform a Test Deployment in Your Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By reading the design sections earlier in this chapter, you can develop a good idea about which remote site connectivity features are appropriate for your organization. However, before you deploy any technology in your production network, it is a good practice to perform a test deployment in a lab first. Testing in a lab familiarizes you with how the technology works and gives you the opportunity to experiment with different features when alternatives are available.

When testing how to deploy a remote site connection, evaluate the following key issues in a lab setting before deciding how to deploy the technology in your production environment:

- Whether to deploy a dial-up connection, a PPTP VPN, or an L2TP/IPSec VPN.
- Whether to deploy an on-demand connection or a persistent connection.
- Whether to deploy a one-way initiated connection or a two-way initiated connection.
- If you plan to deploy a VPN connection, what type of perimeter network you want to use.
- Whether to use certificate-based EAP-TLS or password-based MS-CHAP v2 for user-level authentication.
- If you plan to deploy an L2TP/IPSec VPN connection (the only connection type that offers computer-level authentication), whether to use computer certificates or preshared keys.
- Whether the encryption method, such as a dial-up or PPTP VPN connection using MPPE encryption or an L2TP/IPSec VPN using IPSec for encryption, influences your decision about which connection type to use.
- Whether to use an Active Directory account (and join your routers to the Active Directory domain) or use a local account for your router user accounts.
- Which dial-up options to set in the router user account.
- Whether to use a default remote access policy, a common policy, or a custom policy.
- Which static IP addresses and which static routes are needed.
- Which routing protocol or protocols are needed.
- Whether your demand-dial routers will support Internet traffic.
- Whether to enable multicast connectivity between your sites.
- Whether to deploy a DHCP server in each site.
- Whether to deploy a domain controller in each site.

For a tool to assist you in evaluating some of the remote connectivity features presented in this list, see "Example: Contoso Connects Remote Sites" (DNSREM_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Example: Contoso Connects Remote Sites" on the Web at http://www.microsoft.com/reskit). This job aid shows you how to deploy a PPTP VPN and a dial-up connection in a lab environment.

For additional test lab deployment examples, including information about deployments that include certificates and L2TP/VPN connections, see "Routing Scenarios" and "Virtual Private Network Implementation Examples" in Help and Support Center for Windows Server 2003.

# Deploy a Remote Site Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To connect remote sites using a site-to-site link, you must decide which of the design options described earlier in this chapter you want to deploy. Use the "Choose Optional Tasks You Need" column in Figure 10.10 to mark which optional tasks you will perform for your deployment.

**Figure 10.10 Deployment Tasks for a Site-to-Site Connection Worksheet**

Site-to-Site Connection Deployment Steps

| Prepared By | | Date | |
|---|---|---|---|

Instructions: Perform the deployment tasks in the sequence shown here. Enter an X in the second column to choose the optional tasks needed for your deployment.

| Task # | Choose Optional Tasks You Need (X) | Task Description |
|---|---|---|
| 1 | | Deploy Active Directory.[1] |
| 2 | | Deploy a certificate infrastructure.[1] |
| 3 | | Deploy an IAS server for RADIUS authentication.[1] |
| 4 | | Configure Active Directory user accounts and groups.[1] |
| 5 | Required | Configure the WAN adapter (modem/interface). |
| 6 | Required | Configure the intranet connection. |
| 7 | | Join the router to the domain.[1] |
| 8 | | Place the router in your perimeter network.[1] |
| 9 | | For L2TP/IPSec, install computer certificates.[1] |
| 10 | | For EAP-TLS, install computer and user certificates.[1] |
| 11 | Required | Configure the Routing and Remote Access service and demand-dial interfaces. |
| 12 | | Configure remote access policies on the answering router or IAS server.[2] |
| 13 | Required | Configure a persistent connection or a disconnect interval. |
| 14 | Required | Configure static routes. |
| 15 | | Configure routing protocols.[1] |
| 16 | | Configure Internet access through the calling router.[1] |
| 17 | | Configure IP multicasting.[1] |
| 18 | | Configure the authentication provider.[2] |
| 19 | | Configure authentication methods.[2] |
| 20 | | Configure ports.[2] |
| 21 | | Configure dial-out or dial-in hours.[2] |
| 22 | | Configure IP packet filters and demand-dial filters.[1] |
| 23 | Required | Initiate the connection. |
| 24 | Required | Test connectivity. |
| 25 | | Configure replication.[1] |

[1] If you do not perform this task, this feature is not deployed for your site-to-site connection.
[2] If you do not perform this task, the default configuration is used for your site-to-site connection.

For a copy of Figure 10.10 that you can use to record which deployment options you need for your site-to-site deployment, see "Site-to-Site Connection Deployment Steps" (DNSREM_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Site-to-Site Connection Deployment Steps" at http://www.microsoft.com/reskit). You can also use this worksheet as a guide for the sequence in which you must perform the tasks.

# Deploy Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

How you manage Active Directory and domain controllers in a branch office varies depending on the size, complexity, and structure of your overall network. If you do *not* plan to deploy a domain controller in the branch office location, Microsoft recommends that you establish a persistent connection between the two sites. If you do plan to locate a domain controller in the remote site with which you want to establish a site-to-site connection, you must place it in a subnet separate from the computers at the main site and create a separate Active Directory site for the branch office subnet.

**To deploy a domain controller for a new branch office site**

**Note**

- This procedure assumes that both sites belong to the same Active Directory domain.

1. Install the Windows Server 2003 operating system on a computer at the main office, promote it to a domain controller, and then confirm that it successfully replicates with the existing domain controllers.

2. Log on as a member of either the Domain Administrators or Enterprise Administrators security group of the forest root domain, open Active Directory Sites and Services, right-click Sites, select New Site, type a name for the new branch office site, and then, under Link Name, select the appropriate site link (such as DEFAULTIPSITELINK).

3. Right-click the server object of the new branch office domain controller from its current location (which might be under Default-First-Site-Name), click Move, and then, in the Move Server dialog box, click the new site that you just created.

4. Right-click Subnets, click New, and then click Subnet to create a new child for the new remote site, providing the appropriate network ID and subnet mask.

   **Note**

   - If, until now, your domain has had only one site, you must create two child objects under **Subnets** — one for the main office site and one for the branch office site, each with its own network ID and subnet mask.

5. If necessary to optimize network performance, make the domain controller a global catalog.

6. Ship the new domain controller to the remote site.

7. Install a LAN adapter connected to the branch office intranet, and then, on the LAN adapter's **Internet Protocol (TCP/IP) Properties** page, configure an IP address, subnet mask, and gateway appropriate for the branch office subnet.

The domain controllers in the main and branch offices cannot replicate yet, because the demand-dial connection does not exist yet. For information about how to configure replication after a connection exists, see "Configure Replication for Active Directory" later in this chapter.

# Deploy a Certificate Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you plan to use EAP-TLS as your authentication protocol for user-level authentication on a dial-up, PPTP VPN, or L2TP/IPSec VPN connection, or if you plan to use L2TP/IPSec with computer certificates for your VPN connection, or if you plan to do both, you must have a certificate infrastructure in place in your network.

For information about creating a certificate infrastructure, see Certificate Services in Help and Support Center for Windows Server 2003, and see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

# Deploy an IAS Server for RADIUS Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For a site-to-site only connection, you use Windows authentication and do not need to deploy an IAS server. However, if you use the same answering router for both a site-to-site connection and a remote access connection that supports mobile or home users, you might decide to use RADIUS authentication instead. If you plan to use RADIUS authentication and Windows Server 2003 IAS, you must have an IAS server available in your network. Deploying an IAS server is the same for both dial-up and VPN site-to-site connections.

**To enable RADIUS authentication**

1. Install an IAS server. To ensure that RADIUS authentication and accounting services remain available, configure both a primary IAS server and one or more backup (secondary) IAS servers to provide redundancy and fault tolerance.

2. Register the IAS servers in the appropriate Active Directory domain.

3. Configure the primary IAS server with RADIUS clients corresponding to your answering routers.

4. Configure each answering router with the RADIUS servers of your primary and secondary RADIUS servers.

5. After you enable the Routing and Remote Access service, configure remote access policies that reflect your dial-up or VPN connection requirements on the primary IAS server. For more information, see "Configure the Routing and Remote Access Service and Demand-Dial Interfaces" and "Configure a Remote Access Policy" later in this chapter.

6. Configure logging methods for user authentication and accounting requests.

7. Copy the IAS configuration (including the remote access policies) from the primary IAS server to the secondary IAS server.

For more information about installing an IAS server and using it for RADIUS authentication, see "Deploying Internet Authentication Service (IAS)" in this book, and see Checklist: Configuring IAS for dial-up and VPN access in Help and Support Center for Windows Server 2003.

# Configure Active Directory User Accounts and Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Note**

- If you plan to use local accounts, do not perform these steps. Instead, use the Demand-Dial Interface Wizard to create a user account for the calling router locally on the answering router.

Each calling router must have a user account, which the answering router uses to authenticate the calling router. If you have more than one calling router and if you joined your routers to the Active Directory domain, you can add each router user account to an Active Directory group to simplify administration.

Use the following procedures to accomplish these tasks:

- Create Active Directory user accounts for routers.
- Add router user accounts to an Active Directory group.

**Create Active Directory User Accounts for Routers**

If you plan to use Active Directory user accounts for demand-dial routers, you manually create an Active Directory user account for each calling router (you can use the Demand-Dial Interface Wizard to create the user account only if you use a local user account).

**To create an Active Directory user account for a router**

1. Open the Active Directory Users and Computers snap-in, and create a user account for the calling router (for a two-way connection, create a user account for the calling router in both sites). The name of the account must match the name of a corresponding demand-dial interface on the remote router.
2. To ensure that connectivity occurs, clear the **User must change password at next logon** check box and select the **Password never expires** check box on the **Account** tab on the property sheet for the user account object.
3. On the user account **Dial-in** tab, select one of the following options:
   - **Allow access.** This option overrides the grant or deny remote access permission setting specified on the Properties page of any associated remote access policy.
   - **Control access through Remote Access Policy.** This option ensures that the grant or deny remote access permission setting specified on the Properties page of any associated remote access policy is used.

**Add Router User Accounts to an Active Directory Group**

If you use Active Directory user accounts for demand-dial routers, you can add the router user accounts to a group to simplify administering them and to simplify configuring remote access policies.

**To add router user accounts to a group**

1. Open the Active Directory Users and Computers snap-in, create an Active Directory group to contain the user accounts of the calling routers. Use an appropriate name, such as **BranchOfficeRouters**.
2. Add the user accounts of the calling routers to the group.

# Configure the WAN Adapter

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

On the demand-dial router at each site, use the following procedures, as appropriate, to configure the WAN adapter for a dial-up connection, the WAN adapter for a dedicated link to an ISP, or (if needed) to configure the DNS server IP address.

**To configure the WAN adapter for a dial-up connection**

1. On the demand-dial router at each site, install the WAN adapter. This WAN adapter is a modem, an ISDN adapter, or other physical device connected to one of the following:

    - For a non-VPN dial-up connection, the line that links the two sites.
    - For a VPN with a dial-up connection to the local ISP, the line that links the branch office router to its local ISP.

2. Provide the phone number and the authorized user name and password. Leave the **Domain** field blank. The phone number is one of the following:

    - For a non-VPN dial-up connection, the phone number of the answering router.
    - For a VPN that uses a modem, ISDN, or PPPoE dial-up connection to its local ISP, the phone number of the local ISP.

**To configure the WAN adapter for a dedicated link from a VPN router to an ISP**

1. On the demand-dial router at each site, install the Internet-connected WAN adapter. This WAN adapter is a DDS, T1, Fractional T1, or Frame Relay adapter.

2. On the WAN adapter's **Internet Protocol (TCP/IP) Properties** page, configure the following:

    a. The public IP address and subnet mask allocated to you by your ISP (or by the InterNIC).

    b. The default gateway of the firewall (if you plan to have the router connect to a perimeter network) or of the ISP router (if you plan to have the router connect directly to the Internet).

**To configure the DNS server IP address (if needed)**

- If you have a VPN connection, configure the DNS server IP address on the Internet-connected interface of the calling router if either of the following is true:

    - If you use the DNS name of the answering router (instead of its IP address), you must specify the DNS server IP address in the Internet-connected interface of the calling router. Because the Routing and Remote Access Wizard clears the **Register this connection's addresses in DNS** check box (on the **DNS** tab) and selects **Disable NetBIOS over TCP/IP** (on the **WINS** tab), you must configure the DNS server IP address on the VPN Internet-connected interface after you run the wizard.

    - If the branch office calling router will provide users access to the Internet (in addition to supporting site-to-site traffic to the main office), you must specify the DNS server IP address in the Internet-connected interface of the calling router. For more information about configuring access to the Internet, see "Configure Internet Access Through the Calling Router" later in this chapter.

# Configure the Intranet Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure the intranet interface on the demand-dial router at each site.

**To configure the connection to the local intranet**

1. Install a LAN adapter connected to the intranet.

2. On the LAN adapter's **Internet Protocol (TCP/IP) Properties** page, configure the following:

    a. The public or private IP address and subnet mask that your network administrator assigned.

    b. The IP addresses of the DNS and WINS name servers for your organization's intranet.

3. Typically, leave the default gateway value on the LAN adapter blank to avoid default route conflicts with the default route pointing to the Internet.

    **Note**

    - If a router is positioned between the demand-dial router that you are configuring and the intranet, the Routing and Remote Access Wizard might not be able to add the demand-dial router to the list of valid remote access servers in the Active Directory. One workaround to this problem is to temporarily configure the default gateway on the LAN adapter. To do this, enter a value for the default gateway, run the Routing and Remote Access and Demand-Dial Interface wizards (as described later), and then remove the default gateway from the LAN adapter.

Do not configure the intranet interface of the router to be a DHCP client. However, an answering router can still use DHCP to obtain IP addresses for calling routers, and vice versa.

# Join the Router to the Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you install the Windows Server 2003 operating system on the computers that you want to use as the calling and answering routers, you can join the routers to the Active Directory domain.

If you have a calling router that is located in the branch office, and you already have a domain controller in the branch office, join the router computer to the domain there.

If you do not have a domain controller at the branch office yet, you can install the Windows Server 2003 operating system on the computer that you want to use as the calling router while that computer is located in the main office, join the computer to the domain, and then ship it to the branch office.

**Note**

- Alternatively, you can set up the routers to dial in once, install Active Directory on the computer in the branch office that you want to use as the domain controller, replicate, and then join the routers to the domain.

# Place the Router in Your Perimeter Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For optimal security, place both the answering router and the calling router in a perimeter network at their respective sites. For more information about VPN routers and perimeter networks, see "VPN servers and firewall configuration" in Help and Support Center for Windows Server 2003.

# Install Computer Certificates for L2TP/IPSec

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you use an L2TP/IPSec site-to-site connection, you must install a computer certificate on both the answering router and on the calling router. You must have a certification authority (CA) in your network to issue these certificates.

You can install a computer certificate for L2TP/IPSec by using one of three methods:

- Configure the automatic enrollment of computer certificates in a Windows Server 2003 domain system container by using Group Policy.
- Use the Certificates snap-in to request a computer certificate.
- Use your Web browser to connect to the CA Web enrollments pages to request a certificate.

**Note**

- It is also possible to use a preshared key to provide authentication for IPSec security associations for an L2TP/IPSec connection. However, using computer certificates is the recommended method.

For information about how to create a certificate infrastructure and install computer certificates, see Certificate Services in Help and Support Center for Windows Server 2003, and see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit. For more information about configuring a preshared key, see Configure a pre-shared key for a demand-dial routing interface in Help and Support Center for Windows Server 2003.

# Install Computer and User Certificates for EAP-TLS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you use EAP-TLS as the user authentication method for your site-to-site connection, you must install a computer certificate on the authentication server of the answering router and a user certificate on the calling router. These steps are the same for a dial-up router and for a VPN router.

For information about how to deploy certificates required by EAP-TLS, see Deploying certificate-based authentication for demand-dial routing in Help and Support Center for Windows Server 2003.

# Configure the Routing and Remote Access Service and Demand-Dial Interfaces

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedures to enable the Routing and Remote Access service and to establish a site-to-site connection:

- Enable Routing and Remote Access.
- Configure the demand-dial interface for the remote site connection.
- Configure an additional demand-dial interface for a temporary ISP link.

**Enable Routing and Remote Access**

When you run the Routing and Remote Access Wizard to enable the Routing and Remote Access service, the choices you make are the same for dial-up routing and for VPN routing.

**To enable the Routing and Remote Access service**

**Note**

- You can skip step 1 if either of the following is true:
  - If this server uses local authentication or authenticates against a RADIUS server.
  - If you have administrative rights to add the computer account of the Routing and Remote Access server to the **RAS and IAS Servers** security group. The wizard automatically adds the computer to **RAS and IAS Servers**.

1. Enable the router as follows:

Ask your domain administrator to add the router's computer account to the **RAS and IAS Servers** security group for this domain by using the Active Directory Users and Computers snap-in or the **netsh ras add registeredserver** command.

If this router must access other domains, ask your domain administrator to add the router's computer account to the **RAS and IAS Servers** security group of the other domains.

Restart the router for the change to take effect immediately.

1. Open **Routing and Remote Access**, select the computer on which you want to enable the Routing and Remote Access service (probably the computer you are currently working on), and then, on the **Action** menu, select **Configure and Enable Routing and Remote Access** to start the Routing and Remote Access Wizard. Complete the wizard pages as shown in Table 10.13.

**Table 10.13 Enabling the Routing and Remote Access Service**

| Wizard Page | Action |
|---|---|
| Configuration | Select **Secure connection between two private networks**. |
| Demand-Dial Connections | Select **Yes** (to use demand-dial routing to access remote networks). |

| | |
|---|---|
| IP Address Assignment | Choose one of the following alternative options: Select **Automatically** to use DHCP if you want to assign addresses automatically without using a specified range of addresses. -or- Select **From a specified range of addresses** if you want to specify an address range (recommended):<br>1. On the **Address Range Assignment** screen, select **New**, and then type values for the following:<br>• **Starting address**<br>• **Ending address**<br>You can use public or private address ranges. Based on what you specify for the starting and ending addresses, the **Number of addresses for the IP address pool** field is prepopulated for you. **Note.** For example, for a two-way connection, you might specify the range 192.168.10.1–192.168.10.2 on the calling router and the range 192.168.0.220–192.168.0.221 on the answering router. In this case, if the calling router initiates the connection, the calling router assigns 192.168.10.1 to itself, and it assigns 192.168.10.2 to the answering router.<br>2. If the static IP address pool address range is an off-subnet address range, ensure that the routes to the address range exist in the routers of your intranet. |

When the Routing and Remote Access Wizard completes, you might see the message "Windows was unable to add this computer to the list of valid remote access servers in the Active Directory. Before you can use this computer as a remote access server, the domain administrator must complete this task." If you see this message, click **OK**. Later, after you complete the Demand-Dial Interface Wizard (described next), you will add the computer account to the RAS and IAS Servers security group.

**Configure the Demand-Dial Interface for the Remote Site Connection**

- The Demand-Dial Interface Wizard appears automatically after the Routing and Remote Access Wizard completes.

**To configure the demand-dial interface for a remote site connection**

- Complete the wizard pages for the Demand-Dial Interface Wizard as shown in Table 10.14.

**Table 10.14 Configuring the Demand-Dial Interface for a Remote Site Connection**

| Wizard Page | Action |
|---|---|
| Interface Name | Type a name for the remote router that matches the user account name that you created earlier for the remote router. |

| | |
|---|---|
| Connection Type | Choose one of the following alternative options:<br><br>**Connect using a modem, ISDN adapter, or other physical device**. Select this option to establish a device-to-device dial-up connection.<br><br>• On the **Select a device** screen, select the modem or adapter this interface will use from the prepopulated list.<br><br>• On the **Phone Number** screen, if this is a calling router, type the phone number of the router this interface will call. (If this is an answering router that is not also a calling router, you can leave this blank.)<br><br>-or-<br><br>**Connect using virtual private networking (VPN)**. Select this option to establish a VPN connection over the Internet.<br><br>• On the **VPN Type** screen, select one of the following:<br><br>  • **Automatic** (accepts either PPTP or L2TP connections)<br><br>  • **Point to Point Tunneling Protocol (PPTP)**<br><br>  • **Layer Two Tunneling Protocol (L2TP)**<br><br>• On the **Destination Address** screen, if this is a calling router, type the IP address of the remote router this interface will connect to. (If this is an answering router, you can leave this field blank.)<br><br>Do *not* select the third option, **Connect using PPP over Ethernet (PPPoE)**, because PPPoE is used to link to the local ISP, not to create a device-to-device dial-up link or a VPN tunnel. |
| Protocols and Security | • Select **Route IP packets on this interface** (the default).<br><br>• If this is an answering router that is not joined to an Active Directory domain, add a local account by selecting **Add a user account so a remote router can dial in**. This creates a local user account on the demand-dial router. (Do *not* select this option if you earlier created an Active Directory user account for the answering router to use to authenticate the calling router.) |
| Static Routes for Remote Networks | To add one or more static routes to define the permanent route between this network and the remote network, click **Add**, and then, in the **Static Route** dialog box, do the following:<br><br>• **Destination** — Type the network ID of the remote site.<br><br>• **Network Mask** — Type the subnet mask for the network ID of the remote site.<br><br>• **Metric** — Select an appropriate number for the metric. |
| Dial In Credentials (for an answering router) | Type and confirm a password for the local user account. **Note.** This page appears *only* if this is an answering router and if you chose **Add a user account so a remote router can dial in** on the **Protocols and Security** page earlier in the wizard (to use a local account rather than an Active Directory account for router authentication). Notice that the prepopulated **User name** provided is the same name as that used for the demand-dial interface. |

| | |
|---|---|
| Dial Out Credentials (for a calling router) | Specify the dial-out credentials that this interface will use to connect to the remote router:<br>• **User name** — Type the name of the user account for the calling router that matches the name of the corresponding demand-dial interface on the answering router.<br>• **Domain** — Type the domain name; typically, both sites belong to the same domain.<br>• **Password** and **Confirm Password** — Type the password.<br>**Note.** If this is an answering router that is not also a calling router, you do not need to provide this information; however, the wizard requires that you fill in this page, so type any name, domain, and password. |

If the Routing and Remote Access Wizard (which ran before the Demand-Dial Interface Wizard) was unable to add the computer to the list of valid remote access servers in Active Directory, you saw the error message "Windows was unable to add this computer to the list of valid remote access servers in the Active Directory. Before you can use this computer as a remote access server, the domain administrator must complete this task." To enable the computer to function as a remote access server, add the computer account for the router to the RAS and IAS Servers security group. For information about how to add a computer account to a group, see Add a computer account to a group in Help and Support Center for Windows Server 2003. If you did not see the error message indicating that the computer had not been added to the valid remote access servers in Active Directory, you do not need to perform this step.

After at least one demand-dial interface exists, you can run the Demand-Dial Interface Wizard at any time to add additional demand-dial interfaces by right-clicking **Network Interfaces** in console tree, and then clicking **New Demand-dial Interface**. You run the wizard again for the following reasons:

- To add other branch office sites, repeat the steps in this procedure for each additional demand-dial interface you want to create.
- To establish a temporary link to the local ISP at the branch office in order to create a demand-dial interface for that link, perform the steps as described in the next section.

#### Configure an Additional Demand-Dial Interface for a Temporary ISP Link

If this is a VPN connection, and you connect your branch office to its local ISP through a temporary link, you must run the Demand-Dial Interface Wizard a second time to create a demand-dial interface for this physical link to the ISP. This link to the ISP can be a dial-up link or a PPPoE link.

### Note

- If you are deploying a non-VPN dial-up link, or a VPN connection between two sites, each of which connects to its local ISP through a dedicated link, do not perform these steps. Instead, perform the steps in "Configure the Demand-Dial Interface for the Remote Site Connection" earlier in this chapter.

#### To configure a demand-dial interface for a temporary link to the ISP

- Open Routing and Remote Access, right-click **Network Interfaces**, click **New Demand-dial Interface**, and then complete the wizard pages for the Demand-Dial Interface Wizard as shown in Table 10.15.

**Table 10.15 Configuring an Additional Demand-Dial Interface for a Temporary Link to the ISP**

| Wizard Page | Action |
|---|---|
| Interface Name | Type an appropriate name, such as **Dial_ISP**. |

| | |
|---|---|
| Connection Type | Choose one of the following alternative options:<br><br>Select **Connect using a modem, ISDN adapter, or other physical device**. Select this option to create a dial-up link to your local ISP.<br><br>• On the **Select a device** screen, select the modem or adapter that this interface will use from the prepopulated list.<br><br>• On the **Phone Number** screen, type the phone number of your local ISP.<br><br>-or-<br><br>Select **Connect using PPP over Ethernet (PPPoE)**. Select this option to create a PPPoE link to your local ISP.<br><br>• On the **Service Name** screen, type the name of the service in the text box provided. (If you leave this text box blank, Windows will automatically detect and configure your service when you connect.)<br><br>Do *not* select the third option, **Connect using virtual private networking (VPN)**, because this demand-dial interface is for the link to the ISP, not for a VPN tunnel. |
| Protocols and Security | Select **Route IP packets on this interface** (do *not* select **Add a user account so a remote router can dial in**). |
| Static Routes for Remote Networks | To add a static host route for the IP address allocated to the answering router by the answering router's ISP (or by InterNIC):<br><br>• **Destination** — Type the IP address of the answering router's Internet-connected interface.<br><br>• **Network Mask** — Type 255.255.255.255<br><br>• **Metric** — Select an appropriate number for the metric. |
| Dial-In Credentials | This page does not appear. |
| Dial-Out Credentials | Specify the dial-out credentials that this interface will use to connect to the local ISP:<br><br>• **User name** — Type the name of the user account that has permission to access the local ISP (this is *not* the router user account).<br><br>• **Domain** — leave this field blank.<br><br>• **Password** and **Confirm password** — Type the password.<br><br>**Note.** Open Active Directory Users and Computers, and then, on the **Dial-in** tab of the user object's **Properties** page for the user account that has permission to access the local ISP, select **Allow access**. |

# Configure a Remote Access Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use a remote access policy to validate a variety of connection settings before a connection is authorized, and to specify a variety of connection restrictions after the connection is authorized.

**Configure the Default Policy or Create a New Policy**

Configuring the Routing and Remote Access service on a demand-dial router or installing IAS on a computer running Windows Server 2003 creates two default remote access policies. You can use the **Connections to Microsoft Routing and Remote Access server** default policy for your site-to-site connection. However, if you want more precise control over connection requirements than the default policy provides, you can create a common or a custom remote access policy.

**To enable the default policy**

**Note**

- Do not perform these steps if you plan to create a common or custom remote access policy, described next.

1. To enable the default policy, do one of the following:

   - If you use Windows authentication, on the answering router open Routing and Remote Access, and, if necessary, double-click Routing and Remote Access and the server name. (Use Windows authentication for a site-to-site only connection.)

   - If you use RADIUS authentication, on the IAS server open Internet Authentication Service, and, if necessary, double-click Internet Authentication Service. (Use either Windows or RADIUS authentication if the answering router for the site-to-site connection also supports remote access users.)

2. In the console tree, click **Remote Access Policies**. In the details pane, right-click the default policy **Connections to Microsoft Routing and Remote Access server**, and then click **Properties**.

3. Select **Grant remote access permission**. (The default selection is **Deny remote access permission**.)

**To add a common or custom remote access policy**

**Note**

- Do not perform these steps if you plan to use the default policy, described earlier.

1. To add a common or custom remote access policy, do one of the following:

   - If you use Windows authentication, open Routing and Remote Access, and, if necessary, double-click Routing and Remote Access and the server name.

   - If you use RADIUS authentication, open Internet Authentication Service, and, if necessary, double-click Internet Authentication Service.

2. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**. Use the New Remote Access Policy wizard to create a common policy, as shown in Table 10.16, or to create a custom policy, as shown in Table 10.17.

**Table 10.16 Creating a Common Remote Access Policy by Using the New Remote Access Policy Wizard**

| Wizard Page | Action |
|---|---|
| Policy Configuration Method | Select **Use the wizard to set up a typical policy for a common scenario**, and then type an appropriate name for the policy, such as **Authenticate BranchOfficeRouters**. |
| Access Method | Select **VPN** or **Dial-up**, as appropriate. |
| User or Group Access | Click **Group**, click **Add**, and then type the group name you created earlier, such as **BranchOfficeRouters**. |
| Authentication Methods | Either accept the default method, **MS-CHAP v2**, or choose **Extensible Authentication Protocol (EAP)** and specify its type (either **MD5-Challenge** or **Smart card or other certificate**). |
| Policy Encryption Level | Select **Strongest encryption**, and clear any other selections. |

**Table 10.17 Creating a Custom Remote Access Policy by Using the New Remote Access Policy Wizard**

| Wizard Page | Action |
|---|---|
| Policy Configuration Method | Select **Set up a custom policy**, and then type an appropriate name for the policy, such as **Authenticate BranchOfficeRouters**. |
| Policy Conditions | **If this is a dial-up (non-VPN) connection**:<br>1. Click **Add**.<br>2. Select **Windows-Groups**, click **Add** twice, and then specify the group name you created earlier (such as **BranchOfficeRouters**). Click **OK** twice to return to the **Policy Conditions** page.<br>3. Click **Add**, and select **NAS-Port-Type.** Click **Add**, and select the appropriate device type, such as **Async (Modem)**, **ISDN Async V.100**, **ISDN Async V.120**, or **ISDN Sync**. Then click **Add**.<br>4. Click **Add**, select **Authentication Type**, click **Add**, select either **MS-CHAP v2** or **EAP**, and then click **Add**.<br>5. Select and configure any other attributes for which you want to specify a setting.<br>-or-<br>**If this is a VPN connection**:<br>1. Click **Add**.<br>2. Select **Windows-Groups**, click **Add** twice, and then specify the group name you created earlier (such as **BranchOfficeRouters**). Click **OK** twice to return to the **Policy Conditions** page.<br>3. Click **Add**, select **NAS-Port-Type**, click **Add**, select **Virtual VPN**, and then click **Add**.<br>4. Click **Add**, select **Tunnel-Type**, click **Add**, select either **Point-to-Point Tunneling Protocol** or **Layer 2 Tunneling Protocol** (as appropriate), and then click **Add**.<br>5. Click **Add**, select **Authentication-Type**, select either **MS-CHAP v2** or **EAP**, and then click **Add**.<br>6. Select and configure any other attributes for which you want to specify a setting. |
| Permissions | Select **Grant remote access permission**. |

| | |
|---|---|
| Profile | If you want to change the defaults, click **Edit Profile**, and then make the desired changes. For example, click **Edit Profile**, select the **Encryption** tab, select **Strongest encryption**, and clear any other selections. |

# Configure a Persistent Connection or a Disconnect Interval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, the wizards configure a site-to-site connection to be an on-demand rather than a persistent connection, and the idle time before disconnecting the on-demand connection is set to **Never**. You can change these defaults.

**To configure a persistent connection or a disconnect interval**

1. On the calling router, open the Routing and Remote Access snap-in.

2. In the console tree, under the desired router, click **Network Interfaces**. Right-click the appropriate demand-dial interface in the details pane, and then click **Properties**.

3. Click the **Options** tab, and then, under **Connection type**, do one of the following:

   - If you want a connection that disconnects, accept the default **Demand dial** (on-demand) option, and then configure the **Idle time before hanging up** by using the drop-down list to select, for example, **10 minutes**.

   - If you want a connection that is always available, select **Persistent connection**.

4. Specify values for **Redial attempts** and **Average redial interval**.

# Configure Static Routes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure several types of static routes for a site-to-site connection. You can add static routes manually or, when requesting routes from the remote router, by using auto-static updates.

**Create Static Routes for a Site-to-Site Connection**

For a site-to-site connection, you can use Routing and Remote Access to create the following types of static routes.

- LAN interface. If you do not use routing protocols for the local intranet, you must create one or more static routes for locations on the local intranet.
- Demand-dial interface for the remote site connection. Although you were prompted to provide a static route for the network ID of the remote site when you ran the Demand-Dial Interface Wizard, you might need to recreate this static route in the future. Alternatively, for a persistent site-to-site connection only, you can enable a routing protocol instead of using static routes.
- Demand-dial interface for a link to a local ISP. If you created an additional demand-dial interface for a link to the local ISP, you were prompted to provide a static host route for the IP address of the answering router's Internet-connected interface when you ran the Demand-Dial Interface Wizard. You might need to recreate this static host route in the future.
- Router user account. For a one-way initiated connection in which the answering router is a standalone router or a member of a native-mode Active Directory domain, you can omit creating a demand-dial interface on the answering router, but, in this case only, you must create one or more static routes on the calling router's user account that identify the network IDs of the calling router's site.

Use the following procedures to accomplish these tasks:

- Create static routes on the LAN or demand-dial interfaces.
- Create static routes on the router user account.

**Create static routes on the LAN or demand-dial interfaces**

For each static route you create, fill out the **Static Route** dialog box as shown in Table 10.18. For information about how to add static routes, see Add a static route in Help and Support Center for Windows Server 2003.

**Table 10.18 Configuring the Static Route Dialog Box for a Site-to-Site Connection**

| Interface | Action |
|---|---|
| LAN interface<br><br>(Calling and answering routers) | Specify values for the following fields:<br>• **Interface.** Select **Local Area Connection**.<br>• **Destination.** Type the network ID of the local site.<br>• **Network mask.** Type the subnet mask of the local site.<br>• **Gateway.** Do *not* specify a default gateway on the LAN interface.<br>• **Metric.** Select a number representing the appropriate metric.<br><br>(The check box **Use this route to initiate demand-dial connections** is unavailable.) |

| | |
|---|---|
| Demand-dial interface for the remote site [1]<br><br>(Calling router only) | Specify values for the following fields:<br><br>• **Interface.** Select the demand-dial interface for the connection to the remote site.<br><br>• **Destination.** Type the network ID of the remote site. (Alternatively, you can use the default route.)<br><br>• **Network mask.** Type the subnet mask of the remote site.<br><br>• **Gateway.** This field is unavailable.<br><br>• **Metric.** Select a number representing the appropriate metric. To prevent this static route from causing problems with RIP or OSPF, give it a higher cost than the cost of the static route configured on the LAN interface.<br><br>Select the following check box:<br><br>• **Use this route to initiate demand-dial connections.** |
| Demand-dial interface for the local ISP (if any) [1]<br><br>(Calling router only) | Specify values for the following fields:<br><br>• **Interface.** Select the demand-dial interface for the link to the local ISP.<br><br>• **Destination.** Type the IP address (static host route) of the answering router's Internet-connected interface.<br><br>• **Network mask.** 255.255.255.255<br><br>• **Gateway.** This field is unavailable.<br><br>• **Metric.** Select a number representing the appropriate metric. To prevent this static route from causing problems with RIP or OSPF, give it a higher cost than the cost of the static route configured on the LAN interface.<br><br>Select the following check box:<br><br>• **Use this route to initiate demand-dial connections.** |

1. You might have already created one or more static routes for one or both of these demand-dial interfaces when you ran the Demand-Dial Interface Wizard.

**Create static routes on the router user account**

For a one-way connection, you can create a demand-dial interface on the answering router, but this is not required. If you do not create a demand-dial interface on the answering router, you must create a static route or routes that identify the network IDs of the calling router's site on the calling router's user account.

For information about how to configure static routes on either a local user account or on an Active Directory user account, see Configure static routes for a dial-in user in Help and Support Center for Windows Server 2003. When performing the steps in that Help topic, you are prompted to provide a value for **Destination**. Type a network ID of the calling router's site.

**Configure Auto-static Updates**

You can use auto-static updates to request all of the routes of the router on the other side of a site-to-site connection. Auto-static updates are supported when you use RIP for IP, but not OSPF.

For more information about how to manually configure auto-static updates, see Perform manual auto-static updates in Help and Support Center for Windows Server 2003.

For more information about how to schedule auto-static updates, see Perform scheduled auto-static updates in Help and

Support Center for Windows Server 2003.

# Configure Routing Protocols

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Instead of manually configuring static routes, you can use routing protocols for the following:

- On each LAN interface.
- On each demand-dial interface that is used for a persistent site-to-site connection.

If you use routing protocols, you must configure each new LAN interface or demand-dial interface to use the protocols, because, typically, each interface is connected to a different subnet.

**To configure routing protocols for a persistent site-to-site connection**

- For information about how to deploy the IP unicast routing protocols supported by Routing and Remote Access, see RIP for IP and OSPF in Help and Support Center for Windows Server 2003.

# Configure Internet Access Through the Calling Router

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

At a branch office, you can configure the calling router to grant users access to the Internet in addition to sending traffic to the main site over the site-to-site connection. Choose one of the following scenarios if you want to configure access to the Internet:

- Access the Internet through the main office — for greater security.
- Access the Internet directly — for faster performance.

**For Security, Access the Internet Through the Main Office**

To access the Internet through the main office, use the following steps to add a default (0.0.0.0/0.0.0.0) route to the demand-dial interface used for the dial-up or VPN connection. The default route ensures that all IP packets that cannot find specific routes on the private branch office network are sent to the Internet-connected interface of the demand-dial router at the main office. You might use this alternative if you use ISA Server at the main office.

**To use the calling router to access the Internet through the main office**

1. On the branch office demand-dial router, open the **Routing and Remote Access** snap-in.
2. In the console tree, expand the router you want to configure, expand **IP Routing**, right-click **Static Routes**, and then select **New Static Route**.
3. In the **Static Route** dialog box, configure the following:
   a. In the **Interface** box, select the demand-dial interface used for the dial-up or VPN connection to the main office.
   b. In the **Destination** box, type **0.0.0.0**
   c. In the **Network Mask** box, type **0.0.0.0**
   d. In the **Metric** box, accept the default value (**1**).
   e. Select **Use this route to initiate demand-dial connections**.
   f. Click **OK**.

**For Performance, Access the Internet Directly**

To access the Internet directly from the branch office, you have two options, depending on how the branch office connects to the local ISP:

- The branch office uses a dial-up connection to its local ISP.
- The branch office uses a demand-dial connection to its local ISP.

**Option 1: The branch office uses a dial-up connection to its local ISP**

This method, which is more common and requires configuring only one static route, assumes that the branch office uses a dial-up connection to the local ISP in conjunction with the demand-dial connection to the main office.

**To use the calling router to access the Internet directly if the branch office uses a dial-up connection to its local ISP**

1. On the branch office demand-dial router, open the **Routing and Remote Access** snap-in.
2. In the console tree, expand the router you want to configure, expand **IP Routing**, right-click **Static Routes**, and then

select **New Static Route**.

3. In the **Static Route** dialog box, configure the following:

    a. In the **Interface** box, select the demand-dial interface used for the dial-up or VPN connection to the main office.

    b. In the **Destination** box, type the network ID of the main office network.

    c. In the **Network Mask** box, type the network mask for the main office network ID.

    d. In the **Metric** box, accept the default value (**1**).

    e. Select **Use this route to initiate demand-dial connections**.

    f. Click **OK**.

**Option 2: The branch office uses a demand-dial connection to its local ISP**

This method, which is less common and requires configuring two static routes, assumes that the branch office uses a demand-dial connection to the local ISP in conjunction with the demand-dial connection to the main office.

**To use the calling router to access the Internet directly if the branch office uses a demand-dial connection to its local ISP**

1. On the branch office demand-dial router, open the **Routing and Remote Access** snap-in.

2. In the console tree, expand the router you want to configure, expand **IP Routing**, right-click **Static Routes**, and then select **New Static Route**.

3. In the **Static Route** dialog box, configure the following:

    a. In the **Interface** box, select the demand-dial interface used for connecting the branch office to its ISP.

    b. In the **Destination** box, type **0.0.0.0**.

    c. In the **Network Mask** box, type **0.0.0.0**.

    d. In the **Metric** box, accept the default value (**1**).

    e. Select **Use this route to initiate demand-dial connections**.

    f. Click **OK**.

4. In the console tree, right-click **Static Routes**, and then select **New Static Route**.

5. In the **Static Route** dialog box, configure the following:

    a. In the **Interface** box, select the demand-dial interface used for the dial-up or VPN connection to the main office.

    b. In the **Destination** box, type the network ID of the main office network.

    c. In the **Network Mask** box, type the network mask for the main office network ID.

    d. In the **Metric** box, accept the default value (**1**).

    e. Select **Use this route to initiate demand-dial connections**.

    f. Click **OK**.

# Configure IP Multicasting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure your site-to-site connection to support IP multicast applications.

**To configure demand-dial routers for IP multicasting**

1. For both the calling and the answering router, in the Routing and Remote Access snap-in, expand **IP Routing**, right-click **General**, click **New Routing Protocol**, and then click **IGMP Router and Proxy**.

2. Right-click **IGMP**, click **New Interface**, select the interface you want to enable, and then configure the interface as follows:

   a. Select **Enable IGMP** (the default), and then do the following:

      If this is the demand-dial interface, select **IGMP proxy**.

      If this is an intranet interface, select **IGMP router**.

   b. Typically, for **IGMP protocol version**, you can accept **Version 3** (the default).

3. If needed, modify the default IGMP router settings. In most cases, no changes are needed.

   For more information about how to modify these settings, if analysis of your network indicates that you do need to modify them, see Configure IGMP router settings in Help and Support Center for Windows Server 2003.

# Configure the Authentication Provider

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After the Routing and Remote Access and Demand-Dial Interface wizards complete, Windows authentication and Windows accounting are selected by default. You can change these defaults from Windows authentication and Windows accounting to RADIUS authentication and RADIUS accounting, or you can choose separate providers for authentication and accounting. For a deployment that supports only a site-to-site connection, use Windows authentication and Windows accounting. However, you can change these defaults if the same answering router will support both the site-to-site connection and remote access users, and you want to use RADIUS as either the authentication provider or the accounting provider.

Use the following procedures to accomplish these tasks:

- Configure the authentication provider on the answering router.
- Configure the accounting provider on the answering router.

**To configure the authentication provider on the answering router**

1. Configure either Windows authentication or RADIUS authentication.

   For information about how to configure Windows or RADIUS authentication, see Use Windows Authentication or "Use RADIUS authentication" in Help and Support Center for Windows Server 2003.

2. If you select RADIUS authentication, add the answering router as a RADIUS client on the IAS server.

   For information about how to add the answering router as a RADIUS client, see "Configure RADIUS Clients" in Help and Support Center for Windows Server 2003.

**To configure the accounting provider on the answering router**

1. Select either Windows accounting or RADIUS accounting.

   For information about how to select the accounting method you want to use, see Use Windows Accounting or "Use RADIUS accounting" in Help and Support Center for Windows Server 2003.

2. If you select RADIUS accounting, add the answering router as a RADIUS client on the IAS server.

   For information about how to add the answering router as a RADIUS client, see "Configure RADIUS Clients" in Help and Support Center for Windows Server 2003.

# Configure Authentication Methods

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, the answering router is configured to accept EAP-TLS, MS-CHAP v2, and MS-CHAP as the authentication methods. To increase security, use either EAP-TLS alone or EAP-TLS along with MS-CHAP v2. Alternatively, you can use MS-CHAP v2 with passwords for user authentication.

**To configure the authentication method on the answering router**

1. On the answering router, specify which authentication method or methods to accept.

   By default, the answering router is configured to accept EAP-TLS, MS-CHAP v2, and MS-CHAP.

   - To increase security, clear the MS-CHAP selection, and, if you plan to use EAP-TLS only, also clear the MS-CHAP v2 selection.

     -or-

   - To use MS-CHAP v2 with passwords for user authentication, clear the EAP-TLS selection.

   For information about how to add or clear authentication methods, see Enable authentication protocols in Help and Support Center for Windows Server 2003.

2. If you select EAP-TLS authentication, be sure to perform the steps in "Install Computer and User Certificates for EAP-TLS" earlier in this chapter.

# Configure Ports in Routing and Remote Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If needed, you can change the default values for **Ports** in Routing and Remote Access.

**To configure ports**

- Fill out the **Configure Device** dialog box as shown in Table 10.19. For information about how to configure ports for your site-to-site connection, see Enable routing on ports in Help and Support Center for Windows Server 2003.

**Table 10.19 Using the Configure Device Dialog Box to Configure Ports**

| Option | Action |
|---|---|
| **Remote access connections (inbound only)** | Select this option if you want to allow connections from individual remote access users in addition to connections from a remote router. |
| **Demand-dial routing connections (inbound and outbound)** | Select this option if you want to enable a site-to-site demand-dial connection. This is the default. |
| **Demand-dial routing connections (outbound only)** | This option is available only for a PPPoE connection to the Internet. Select this option if you want this router to function only as a calling router. |
| **Phone number for this device** | If you configure a value for **Called-Station-ID** in the remote access policy associated with this connection, you must manually type a value for this field:<br>• If this is a dial-up connection, type the phone number that was entered in the **Called-Station-ID** remote access policy attribute.<br>• If this is a VPN connection, type the IP address of the VPN server's Internet-connected interface (entered in the **Called-Station-ID** remote access policy attribute). |
| **Maximum ports** | For PPTP or L2TP/IPSec VPN routers only, change the default number of ports. For optimal scalability, use 10 or fewer ports. |

# Configure Dial-out or Dial-in Hours

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For on-demand connections, you can specify time limits for dial-out or dial-in hours.

**To configure dial-out hours on a calling router**

- Configure dial-out routers on the calling router to prevent connections at unwanted times.

  For information about how to use Routing and Remote Access to configure dial-out hours on the demand-dial interface of a calling router to block outgoing connections at specific times of the day or week (such as at night or on weekends), see Configure dial-out hours in Help and Support Center for Windows Server 2003.

**To configure dial-in hours on an answering router**

- Configure dial-in hours on the answering router to specify when the calling router can access the answering router.

  For information about how to use remote access policies to configure dial-in hours on an answering router to block incoming connections at specific times of the day or week (such as at night or on weekends), see Configure dial-in constraints in Help and Support Center for Windows Server 2003.

# Configure IP Packet Filters and Demand-Dial Filters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For an on-demand VPN connection, you can specify IP packet filters and demand-dial filters.

Use the following procedures to accomplish these tasks:

- Configure IP packet filters on the Internet interface
- Match IP demand-dial filters to IP packet filters on the demand-dial interface

**Configure IP Packet Filters on the Internet Interface**
You can configure PPTP or L2TP/IPSec input and output filters on the Internet-connected interface of a VPN router to allow only PPTP or only L2TP/IPSec traffic to travel between the two sites.

How you configure firewall filters and filters on the VPN router depends on the relative position of the VPN router and firewall. For information about configuring filters for a VPN site-to-site server, see "Deploying Dial-up and VPN Remote Access Servers" in this book, and see "VPN servers and firewall configuration," Add PPTP Filters, and Add L2TP over IPSec Filters in Help and Support Center for Windows Server 2003.

**Configure IP Demand-Dial Filters and Match Them to IP Packet Filters on the Demand-Dial Interface**
You can configure demand-dial filters to specify which types of traffic are allowed to create a site-to-site connection. By matching demand-dial filters to the IP packet filters, you can also prevent a calling router from establishing a demand-dial connection for traffic that IP packet filters are configured to discard.

For information about how to configure demand-dial filters and to match them to IP packet filters, see Configure demand-dial filters and Demand-dial routing design considerations in Help and Support Center for Windows Server 2003.

# Initiate the Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

First, ensure that both routers are enabled for dial-up access, and then initiate the connection.

**To confirm that a router has permission to initiate an on-demand connection**

1. Open Active Directory Users and Computers, and then on the **Dial-in** tab of the user account of the router's **Properties** page, confirm that either **Allow access** or **Control access through Remote Access Policy** is selected.

2. If **Control access through Remote Access Policy** is selected, open the Routing and Remote Access snap-in and confirm that the remote access permission setting on the Properties page of the associated Remote Access Policy is set to **Grant remote access permission**.

**To initiate the on-demand connection on the calling router**

- On the calling router, in the **Routing and Remote Access** snap-in, click **Network Interfaces**, right-click the demand-dial interface for which you want to initiate a connection, and then click **Connect**.

If the calling router uses a dial-up connection to the local ISP, the local ISP assigns the router a temporary IP address. You can confirm that this IP address exists by typing **ipconfig** at a command prompt.

# Test Connectivity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedures to test the remote site connection:

- **Test network connectivity.** To test network connectivity, ping a client computer at the main office from a client computer in the branch office.

- **Test application connectivity.** To test application connectivity, map a network drive to a client computer at the main office from a client computer in the branch office. For example, map a drive to *ClientComputerName*\c$.

- **Force replication between domain controllers in both sites.** If you have a domain controller in both sites, use the following procedure to test replication:

  1. Open Active Directory Users and Computers, and create a new user account called, for example, *TestReplication*.

  2. Use the Active Directory **Replicate Now** option to test that replication takes place over the site-to-site connection.

     For information about how to force replication between the two sites, see Force replication over a connection in Help and Support Center for Windows Server 2003.

  3. Confirm that the new user account exists on domain controllers located in both sites.

# Configure Replication for Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Replication**
If you installed an Active Directory domain controller in the branch office, you must ensure that replication takes place continually. How you do this depends on whether this is a persistent connection or an on-demand connection.

Use the following procedures to accomplish these tasks:

- Configure a replication interval for a persistent connection
- Configure reciprocal replication for a one-way initiated on-demand connection

**Configure a Replication Interval for a Persistent Connection**
If this is a persistent connection, you can schedule replication to occur after a specified interval.

Use the Active Directory Sites and Services snap-in to configure a replication interval. For information about how to specify a replication interval, see Configure site link replication frequency in Help and Support Center for Windows Server 2003.

**Configure Reciprocal Replication for a One-Way Initiated On-Demand Connection**
If this is a one-way initiated on-demand connection, you must configure reciprocal replication using the Active Directory Service Interfaces (ADSI) Edit tool.

Install ADSI Edit, a Windows Support tool, on a domain controller in the main office or in the branch office. For information about how to install Windows Support Tools, which include ADSI Edit, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Install Windows Support Tools.

**Caution**

- If you use the ADSI Edit snap-in and incorrectly modify the attributes of Active Directory objects, you can cause serious problems that might require you to reinstall Windows Server 2003. Microsoft cannot guarantee that problems resulting from the incorrect modification of Active Directory object attributes can be solved. Modify these attributes at your own risk.

**To enable reciprocal replication on a site link for a one-way initiated on-demand connection**

1. On a domain controller, type **adsiedit.msc** in the Run dialog box to open the ADSI Edit snap-in.
2. Expand the **Configuration** container, expand the **Sites** container, expand the **Inter-Site Transports** container, and then click **CN=IP**.
3. In the details pane, right-click the site link object for the sites for which you want to enable reciprocal replication, and then click **Properties**.
4. In the **Attributes** box, double-click **Options**.
5. In the **Integer Attribute Editor** dialog box, take one of the following actions:
   a. If the **Value** box displays **<not set>**, type **2**.
   b. If a value is displayed, convert the integer value to a binary value and use the binary OR operation to combine that value with the binary value **0010**, and then type the integer value of the result in the **Value** box.

      For a job aid that includes a table of examples showing how to perform this operation, see "Example: Contoso Connects Remote Sites" (DNSREM_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Example: Contoso Connects Remote Sites" on the Web at http://www.microsoft.com/reskit).

Alternatively, you can enable reciprocal replication on a connection, instead of on a site link.

**To confirm reciprocal replication over a one-way initiated on-demand connection**

1. On the branch office domain controller, create a user account called TestReplication.

2. On the calling router in the branch office, establish a connection to the answering router in the main office.

3. Wait 15 to 20 minutes, and then open Active Directory Users and Computers on a domain controller in the main office. If your connection is working correctly, the TestReplication user account will be listed.

4. To confirm that replication did in fact take place over the one-way initiated site-to-site connection, perform the following steps:

    a. On both the calling router and the answering router, type **ipconfig** at a command prompt. If, when you ran the Routing and Remote Access wizard, you chose the recommended option to configure IP address assignment from a specified range of addresses, the output of both **ipconfig** commands will include IP addresses from the specified address ranges.

    b. On the branch office domain controller, type **tracert** at a command prompt. If, when you ran the Routing and Remote Access wizard, you configured IP address assignment from a specified range of addresses, the output of the **tracert** command will include an IP address from the specified range that the answering router assigns to the calling router.

# Additional Resources for Connecting Remote Sites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this chapter.

**Related Information**

- "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit for more information about creating hardware and software inventories and mapping your existing network before deploying new features.
- "Deploying Domain Name System (DNS)" and "Deploying WINS" in this book for more information about name resolution.
- "Deploying DHCP" in this book for more information about IP address assignment.
- "Deploying Dial-up and VPN Remote Access Servers" and "Deploying Remote Access Clients Using Connection Manager" in this book for more information about using the Routing and Remote Access service to connect individual computers (such as computers of users working from home or mobile users) to a central office.
- "Deploying Internet Authentication Service (IAS)" in this book for more information about deploying IAS.
- "Designing the Site Topology" in *Designing and Deploying Directory and Security Services* of this kit for more information about Active Directory replication between sites.
- "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit for more information about creating a certificate infrastructure.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about Internet Authentication Service (IAS), Routing and Remote Access, demand-dial routing, virtual private networking, unicast IP routing, and Network Address Translation (NAT).
- The Virtual Private Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about virtual private networking.
- The Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about hardware compatibility.
- The Active Directory Branch Office Planning Guide link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about planning and deploying Active Directory for multiple branch office sites.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Setting up demand dial routing Setting Up Demand-Dial Routing" and "Deploying Router-to-Router VPNs" in Help and Support Center for Windows Server 2003 for more information about deploying site-to-site connections.
- "Routing Scenarios" and "Virtual Private Network Implementation Examples" in Help and Support Center for Windows Server 2003 for more examples of test lab deployments for site-to-site connections.

**Related Job Aids**

- "Example: Contoso Connects Remote Sites" (DNSREM_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Example: Contoso Connects Remote Sites" on the Web at http://www.microsoft.com/reskit).

- "Site-to-Site Connection Deployment Steps" (DNSREM_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Site-to-Site Connection Deployment Steps" on the Web at http://www.microsoft.com/reskit).

# Deploying a Wireless LAN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Wireless connectivity offers users a high degree of mobility and provides another networking option when traditional wired networks are impractical. The Microsoft® Windows® Server 2003 operating system provides the networking services needed to deploy a secure and manageable wireless local area network (WLAN) infrastructure within an enterprise environment. The deployment of WLANs — the primary wireless networking support for an enterprise environment provided by Windows Server 2003 and the Microsoft® Windows® XP operating system — is detailed in this chapter.

**In This Chapter**

Overview of Deploying a Wireless LAN

Adapting the Network Infrastructure for a WLAN

Designing a Wireless Network Access Solution

Implementing a WLAN Test Environment

Additional Resources for Deploying a Wireless LAN

**Related Information**

- For more information about the Internet Authentication Service (IAS), see "Deploying Internet Authentication Service (IAS)" in this book and the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

- For information about designing a certificate infrastructure, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

- For more information about implementing the Active Directory® directory service, see "Designing the Active Directory Infrastructure" in *Designing and Deploying Directory and Security Services*.

# Overview of Deploying a Wireless LAN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To provide authorization and authentication, automatic IP address assignment, and name resolution for wireless users, your networking infrastructure should include the following services:

- Active Directory directory service
- Remote Authentication Dial-In User Service (RADIUS) servers and proxies
- A certificate infrastructure, also known as a *public key infrastructure* (*PKI*)
- Dynamic Host Configuration Protocol (DHCP) services
- Domain Name System (DNS) services

These services together provide the security, availability, and scalability needed for an enterprise WLAN solution. Before you begin designing and deploying an enterprise WLAN, all of the supporting components and services in your networking infrastructure should be in place.

All of the components required for an enterprise WLAN solution are included with Windows Server 2003 and Windows XP. Windows Server 2003 includes DHCP, DNS, and Certificate Services, and support for RADIUS (through the Internet Authentication Service [IAS]), the IEEE 802.1X standard, and certificate authentication. Windows XP with wireless network adapters provides support for wireless devices such as laptops and personal digital assistants (PDAs), the IEEE 802.1X standard, and certificate authentication.

**Note**

- Though the WLAN components are included in the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems, there are differences in the levels of features supported and the capabilities. For information about differences in these services among the Windows Server 2003 operating systems, see Help and Support Center for Windows Server 2003.

After designing your enterprise WLAN and deploying the networking components and services needed for the WLAN, you will be able to maintain a secure and manageable wireless network by using supported features such as the Wireless Zero Configuration (WZC) service included in Windows XP and Windows Server 2003, RADIUS-based 802.1X authentication, and interoperability with other networking services.

# Process for Deploying a Wireless LAN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In deploying a wireless LAN, adapt your existing network infrastructure for a WLAN before designing the wireless network access solution — that is, deciding where to locate wireless access points (APs) and how to deploy them; designing wireless security and unauthenticated access; optionally designing a public space WLAN; and designing for better manageability. Before embarking on a full-scale WLAN deployment, implement a WLAN test environment and test your wireless networking solution. Figure 11.1 shows the major steps in the process for deploying a WLAN.

**Figure 11.1 Deploying a Wireless LAN**



# WLAN Technology Background

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The WLAN solution provided by Windows XP and Windows Server 2003 is based on IEEE standards 802.11 and 802.1X.

**IEEE 802.11** IEEE 802.11, the standard for WLANs, specifies a technology that operates in the 2.4 through 2.5 GHz Industrial, Scientific, and Medical (ISM) band and has a maximum bit rate of 2 megabits per second (Mbps). IEEE 802.11b supports two additional speeds, 5.5 Mbps and 11 Mbps, in the ISM band.

**Note**

- The latest IEEE standard, IEEE 802.11a, specifies a technology that operates in a 5.725 through 5.875 GHz band with a maximum bit rate of 54 Mbps.

The 802.11 standard defines an *access point* (AP) as a device that functions as a transparent bridge between the wireless clients and the existing wired network. The AP contains at least one interface to connect to the existing wired network, and transmitting equipment to connect with the wireless clients. The AP also contains IEEE 802.1D bridging software, thereby acting as a transparent bridge between wireless and wired data-link layers.

**IEEE 802.1X** The 802.1X standard defines port-based network access control to provide authenticated network access for Ethernet networks. This port-based network access control uses the physical characteristics of the switched LAN infrastructure to authenticate devices attached to a LAN port. Access to the port can be denied if the authentication process fails. Although this standard is designed for wired Ethernet networks, it applies to 802.11 WLANs as well.

# Adapting the Network Infrastructure for a WLAN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In adapting your network infrastructure for a WLAN, verify you have the required components; eliminate any potential single points of failure; and define IP addressing and subnets needed to support your wireless clients. Figure 11.2 shows the process for adapting your existing network infrastructure for a WLAN.

**Figure 11.2 Adapting the Network Infrastructure for a WLAN**



# Verifying Required Components

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To support a secure wireless solution, your existing network infrastructure must include the following components:

- Active Directory, to store account properties and validate password-based credentials.
- DHCP services, to provide automatic IP configuration to wireless clients.
- DNS services, to provide name resolution.
- RADIUS support, to provide centralized connection authentication, authorization, and accounting.
- A certificate infrastructure, also known as a PKI, to issue and validate the certificates required for Extensible Authentication Protocol–Transport Level Security (EAP-TLS) and Protected EAP (PEAP)–TLS authentication. TLS can use either smart cards or registry-based user certificates for authenticating the wireless client.
- For PEAP-Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) authentication, computer certificates for the RADIUS servers and root CA certificates of the issuing CAs on the wireless clients (if needed).

Windows Server 2003 provides all of these components, with some variations in the levels of features supported and capabilities in different editions of the operating system (Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition). For information about differences in these services among the different editions of Windows Server 2003, see Help and Support Center for Windows Server 2003.

After verifying that your existing network infrastructure includes the appropriate components, update the design to support your WLAN.

**Active Directory: Planning Groups and Group Policy for Wireless Access**
Active Directory contains the user and computer accounts that are used for authentication and authorization of wireless users. It also contains the Group Policy settings that govern wireless connections — for example, information regarding autoenrollment for the user and computer certificates that are installed on wireless clients, and the Wireless Network (IEEE 802.11) Policies settings that specify preferred networks, Wired Equivalent Privacy (WEP) settings, and IEEE 802.1X settings for wireless connections.

To plan for the configuration of Active Directory for your wireless clients, identify the user and computer accounts for wireless users, and add them to a group that will be used in conjunction with a remote access policy to manage wireless access. You must also determine how to set the remote access permission on the user and computer accounts.

**Note**

- For a native-mode domain, you can use universal groups and nested global groups. For example, you might create a universal group named WirelessUsers that contains global groups of wireless user and computer accounts for intranet access.

For information about designing group policies for a WLAN, see "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit.

For information about configuring Active Directory for a WLAN, see "Implementing a WLAN Test Environment" later in this chapter. For information about the design and deployment of Active Directory, see "Designing the Active Directory Infrastructure" in *Designing and Deploying Directory and Security Services*.

**Note**

- You can configure a computer running Windows Server 2003 or the Microsoft® Windows® 2000 Server operating system as an Active Directory domain controller. To configure a Windows 2000 Server–based computer as an Active Directory domain controller for wireless access, you must install Service Pack 3 (SP3) or later.

**DNS: Identifying Zones Where Wireless Clients Will Register**
In DNS, identify the DNS zones where the wireless computers will register DNS address records, and ensure that the zones are configured for dynamic updates. Optionally specify that the DNS zones are Active Directory–integrated zones, which provide secured updates, and that the DNS zones are updated by DHCP. For information about configuring DNS zones, see "Deploying Domain Name System (DNS)" in this book.

**DHCP: Creating Scopes and Leases for Wireless Clients**
On the DHCP server, you must configure the scope and lease duration differently for wireless clients than for clients connected to the wired LAN.

A DHCP *scope* is the full, consecutive range of possible IP addresses for a network, which typically defines a single physical subnet where DHCP services are offered. A DHCP *lease duration* is the specified time during which the IP addresses within the DHCP scope can be leased.

Define separate scopes for your wired subnets and your wireless subnets, so that you can configure the lease duration differently for the clients requiring DHCP services on your WLAN than for the clients on your wired LAN.

The default lease duration for a DHCP scope is eight days. This might be optimal for clients on your wired network. Because wireless clients do not release their addresses when the wireless user roams to a new subnet, you should decrease the lease duration substantially for wireless clients. By setting a shorter lease duration for wireless clients, you free IP addresses for reuse throughout the day instead of leaving the addresses unavailable for as long as eight days. When determining the optimal lease duration for the wireless clients in your environment, keep in mind the additional processing load that the shorter lease duration places on your DHCP server.

For more information about configuring scopes and lease durations on a DHCP server, see "Deploying DHCP" in this book.

**RADIUS: Verifying Support for 802.1X and EAP-TLS**
Verify that the RADIUS servers, such as IAS, support EAP-TLS and PEAP-MS-CHAP v2 authentication. Optionally, verify that the RADIUS servers support the use of the NAS-Port-Type attribute in their Access-Request messages. IAS uses the NAS-Port-Type attribute to identify the wireless connection types.

Computers running Windows Server 2003 or Windows 2000 Server IAS can be used for RADIUS support. A Windows 2000 IAS server must have Service Pack 3 (SP3) or later installed.

### Notes

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client by using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query.

- With IAS in Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

By using IAS as your RADIUS server and Active Directory as your directory service, you can provide a single sign-on solution. IAS uses Active Directory as its user accounts database. The same set of credentials is used to control network access (authenticating and authorizing access to a network), to log the user on to an Active Directory domain, and to control access to resources in the domain.

The RADIUS server must support the EAP-TLS authentication protocol, which is used in certificate-based security environments. In addition, EAP-TLS is required if you are using smart cards for network access authentication. The RADIUS server also must support the PEAP-MS-CHAP v2 authentication protocol, which is used in password-based security environments.

**Certificate Infrastructure: Verifying Support for EAP-TLS and PEAP-MS-CHAP v2**
To support EAP-TLS authentication, verify that the certificate infrastructure supports issuing user and computer certificates to wireless client computers and issuing computer certificates to your RADIUS servers.

To support PEAP-MS-CHAP v2 authentication, verify that your RADIUS servers have the appropriate computer certificates installed, and that each wireless client computer has the root certification authority (CA) certificates of the issuers of the

computer certificates of the RADIUS servers installed.

You can use a third-party CA to issue certificates for wireless access, as long as the certificates that you install fulfill the requirements for TLS authentication.

For more information about certificate requirements, see the discussion of deploying a certificate infrastructure in "Implementing a WLAN Test Environment" later in this chapter. For information about designing and deploying a certificate infrastructure, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services*.

# Eliminating Single Points of Failure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To ensure that wireless clients can continue to be authenticated on the network and can access resources and applications, eliminate single points of failure in your network infrastructure by including:

- Redundant services (such as Active Directory domain controllers) on separate subnets.
- Clustered DHCP services, in the event that one of the cluster nodes fails.
- DNS on all domain controllers, in the event that a DNS server fails.
- Redundant RADIUS servers and proxies, to provide fault tolerance for RADIUS-based authentication.
- Redundant switches and routers, in the event that a switch or router fails.
- Redundant network paths between switches and routers.

# Defining IP Addressing and Subnets

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Determine how many additional IP addresses your wireless clients will require, and whether or not to define additional subnets.

**To determine the number of additional IP addresses that you will need for wireless access:**

1. Calculate the number of additional IP addresses that wireless users will require:

    a. Determine the average number of wireless clients currently using your corporate network at any given time.

    b. Add to this the estimated number of additional concurrent wireless clients your network will need to support in the future.

2. Estimate the number of APs (and associated IP addresses) that you will need for wireless network access. For information about how to determine how many wireless APs to deploy, see "Designing Wireless AP Location" later in this chapter.

Based on the number of IP addresses that you will add to accommodate your WLAN, decide whether or not to add additional subnets.

Creating separate subnets for your wireless networking components offers many benefits, including:

- Wired network components do not have to draw from the same pool of existing IP addresses as your wireless clients.
- IP addresses for wireless clients are easier to identify, which assists in easier management and troubleshooting.
- Separate subnets give you increased control over DHCP lease times.
- You can associate each of your physical subnets (both wireless and wired) with sites within Active Directory, which enables you to assign network access policies to the specific subnets.
- If all APs are on the same subnet, you can provide seamless network-layer roaming for the wireless clients. *Network-layer roaming* allows a wireless client to associate with a new AP within the same subnet, in the same wireless network. When crossing subnets, applications that cannot handle a change of address, such as some e-mail applications, might fail.

**Note**

- Network-layer roaming is to be distinguished from *general roaming,* which allows a wireless client to associate with a new AP within the same wireless network. In network-layer roaming, the wireless client associates with a new AP on the same subnet, within the same wireless network.

**Example: An Enterprise Corporation Designs Subnets and IP Addressing for a WLAN**
IEEE 802.11 APs are designed with Ethernet ports and use TCP/IP as a networking protocol. Thus, an enterprise corporation designed their network so that the wireless APs in the building are all attached to the same separate subnet, which is connected to a router.

To avoid using IP addresses from existing subnets, they assigned all of the wireless components to a separate subnet. Because they used a separate subnet for wireless components, the wireless components did not adversely affect the available number of host addresses allocated on previously configured wired subnets.

To keep track of the allocation of IP addresses, they created the IP address numbering scheme shown in Table 11.1. The corporation adopted this numbering convention for all of their buildings that have wireless network connectivity.

**Table 11.1 Example IP Address Allocation for IP Subnet 172.16.50.0/24**

| IP Address | Device |
| --- | --- |
| 172.16.50.1 | Router |
| 172.16.50.2–172.16.50.10 | Servers (terminal, proxy, IAS, and so forth) |
| 172.16.50.11–172.16.50.*x* | Wireless APs |
| 172.16.50.*x*+*1*–172.16.50.254 | Wireless clients |

Under this addressing scheme, addresses were assigned in the following manner:

- Within the IP subnet 172.16.50.0/24, they assigned the router connecting to the rest of the network the first IP address of 172.16.50.1.
- They assigned other devices — such as terminal servers, proxy servers, and IAS servers — addresses from 172.16.50.2 through 172.16.50.10.
- They assigned the wireless APs sequential IP addresses starting with 11. To make it easier to keep track of the wireless APs, they assigned IP addresses that were 10 digits higher than the wireless AP number. For example, Wireless AP 1 was assigned 172.16.50.11, Wireless AP 2 was assigned 172.16.50.12, and so forth.

As is usual practice, they assigned static IP addresses to the APs and any servers on the subnet. To prevent the DHCP server from allocating a static IP address to a wireless client, they created a DHCP scope for the wireless subnet that did not include the assigned servers and APs (the scope range was 172.16.50.*x*+*1* through 172.16.50.254).

They located all APs in a building on the same subnet, which allowed network-layer roaming throughout the building. This also made IP addressing by DHCP servers more manageable. The DHCP server assigned wireless clients dynamic IP addresses.

# Designing a Wireless Network Access Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In designing your wireless network access solution, determine where to locate wireless APs and plan how to deploy the APs; design security for wireless access; if needed, design for unauthenticated access and for a public space WLAN; and, finally, optimize your design for manageability. Figure 11.3 shows the process for designing your wireless network access solution.

**Figure 11.3 Designing Your Wireless Network Access Solution**



# Designing Wireless AP Location

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After designing and verifying that the services needed for your network infrastructure to support an enterprise WLAN are in place, begin the design process for the location of the wireless APs.

To determine where to locate your wireless APs:

- Identify the areas of coverage for wireless users.
- Determine how many wireless APs to deploy.
- Determine where to place wireless APs.
- Select the channel frequencies for wireless APs.

# Identifying the Areas of Coverage for Wireless Users

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To identify the areas of coverage for wireless users:

1. Obtain or create scaled architectural drawings of each floor for each building, in each geographic location in your organization that has wireless users.

2. On the drawing, identify the offices, conferences rooms, lobbies, or other areas where you want to provide wireless coverage.

   It might be useful to enable wireless coverage for a building in its entirety rather than for specific locations within the building. For example, this can prevent connectivity problems that might result from undocking a laptop from an office for use in a different part of your building.

3. Indicate any devices that interfere with the wireless signals.

   Any device that operates on the same frequencies as your wireless devices (in the 2.4 through 2.5 GHz ISM range) might interfere with the wireless signals. Devices that operate on the same frequency include:

   - Existing Bluetooth-enabled devices
   - Microwave ovens
   - Some models of cordless telephones
   - Wireless video cameras
   - Medical equipment

4. Indicate any building construction materials that interfere with wireless signals.

   Metal objects used in the construction of a building can affect the wireless signal. For example, the following common objects interfere with signal propagation:

   - Support girders
   - Elevator shafts
   - Rebar reinforcement in concrete
   - Heating and air-conditioning ventilation ducts
   - Wire mesh that reinforces plaster or stucco in walls

**Note**

- *Radio frequency attenuation* (the reduction of signal strength), shielding, and reflection can affect how you deploy your APs. Refer to the manufacturer of your APs for information regarding the different scenarios that might increase the radio frequency attenuation. Testing software is available with most APs to check for signal strength, error rate, and data throughput. This can be very beneficial during the deployment of your APs.

# Determining How Many Wireless APs to Deploy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To determine how many wireless APs to deploy, following these guidelines:

- Include enough wireless APs to ensure that wireless users have sufficient signal strength from anywhere in the area of coverage.

  Wireless APs typically have an indoor range within a 150-foot radius. Include enough wireless APs to ensure signal overlap between the wireless APs.

- Determine the maximum number of simultaneous wireless users per coverage area.

- Estimate the data throughput that the average wireless user requires.

Add additional wireless APs to:

- Improve wireless client network bandwidth capacity.

- Increase the number of wireless users supported within a coverage area

  Based on the total data throughput of all users, determine the number of users that you can connect through a wireless AP. Obtain a clear picture of throughput before deploying the network or making changes. Some wireless vendors provide an 802.11 simulation tool, which you can use to model traffic in a network and view throughput levels under various conditions.

- Ensure redundancy, in the event that a wireless AP fails.

# Determining Where to Place Wireless APs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to locate the APs close enough together to provide ample wireless coverage but far enough apart to not interfere with each other and increase the error rate. The actual distance needed between any two APs depends upon the combination of the type of AP, the type of AP antenna, and the construction of the building, as well as on sources of signal degradation, shielding, and reflection. For specifications and guidelines for placing wireless APs, see the manufacturer's documentation for the APs and the antennas used with them.

Maintain the best average ratio of wireless clients to APs. The greater the number of wireless clients that are associated with the AP, the lower the effective data transmission rate. Too many wireless clients attempting to use the same AP degrade the effective throughput or available bandwidth for each wireless client. By adding APs, you can increase throughput. To increase the number of APs per wireless client, you must increase the number of APs in a given coverage area. You can move APs closer together up to a point before they start to interfere with each other.

To determine where to place the wireless APs:

1. On the architectural drawings, place wireless APs so that each wireless AP is no further than 300 feet from an adjacent wireless AP.

2. To test the wireless AP placement, perform a site survey:

   - Temporarily place wireless APs in the locations specified on the architectural drawings.

   - Using a laptop equipped with an 802.11 wireless adapter and site survey software (site survey software ships with most wireless adapters), determine the signal strength within each coverage area.

3. In coverage areas where signal strength is low, make any of the following adjustments:

   - Reposition existing APs to increase the signal strength for that coverage area.

   - Reposition or eliminate devices that interfere with signal strength (such as Bluetooth devices or microwave ovens).

   - If possible, reposition or eliminate metal obstructions that interfere with signal propagation (such as filing cabinets and appliances).

   - Add additional wireless APs to compensate for the weak signal strength.

     It is important to remember that radio frequency is three-dimensional. It can be conceptualized as a sphere of signal.

   - Purchase antennas to meet the requirements of your building infrastructure.

     For example, to eliminate interference between APs located on adjoining floors in your building, you can purchase directional antennas that flatten the signal (forming a donut-shaped signal distribution) to increase the horizontal range and decrease the vertical range.

4. Update the architectural drawings to reflect the final number and placement of the wireless APs.

**Example: An Enterprise Corporation Determines AP Location**
An enterprise corporation used information from the manufacturer of its APs and from internal testing to determine the best locations for their APs when using IEEE 802.11b. They determined that a 30-foot radius for each AP (with 60 feet between adjacent APs) would provide the best coverage without interference between the APs.

To assist them in designing AP locations, the IT staff created a plan view of the building, with one drawing per floor. The floor plans displayed the location, name, IP address, and channel information for each AP. Based on the information in the floor plans, the IT staff determined that they could locate their APs in the plenum area between the lowered ceiling and the next floor up, which simplified troubleshooting for the AP equipment.

# Selecting Channel Frequencies for Wireless APs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Direct communication between an 802.11 wireless network adapter and an AP occurs over a common channel corresponding to a frequency range in the S-Band ISM frequency range. You set the channel in the AP, and the wireless network adapter automatically tunes to the channel of the AP with the strongest signal. The wireless network adapter continues communication with the AP until the signal gets weak, at which time it attempts to locate another AP with a stronger signal.

To reduce interference between wireless APs, ensure that wireless APs with overlapping signals use unique channel frequencies. The 802.11b standard reserves 14 frequency channels for use with wireless APs. Within the United States, the Federal Communications Commission (FCC) allows channels 1 through 11. In most of Europe, you can use channels 1 through 13. In Japan, you have only one choice: channel 14.

Figure 11.4 shows the 11 802.11b frequency channels available in the United States. Notice that the 802.11b signals overlap with adjacent channel frequencies. As a result, you can only use three channels (in the United States, channels 1, 6, and 11) without causing interference between adjacent APs.

**Figure 11.4 Channel Overlap for 802.11b APs in the United States**




To select the channel frequencies for the wireless APs:

1. Identify any wireless networks owned by other organizations in the same building. Find out the placement of their wireless APs and the channel frequencies assigned to the APs.

   Radio waves travel through floors and ceilings, so APs located near each other on different floors need to be set to non-overlapping channels. If another organization located on a floor adjacent to your organization's offices has a wireless network, the wireless APs for that organization might interfere with the wireless APs in your network. Contact the other organization to determine the placement and frequencies of their wireless APs so that you can ensure that any of your own wireless APs that provide overlapping coverage use a different channel frequency.

2. Identify overlapping wireless signals on adjacent floors within your own organization.

3. After identifying overlapping coverage areas outside and within your organization, assign channel frequencies for your wireless APs.

   a. Assign channel 1 to the first wireless AP.

   b. Assign channels 6 and 11 to any wireless APs that overlap coverage areas with the first wireless AP, to ensure that those APs do not interfere with one another.

   c. Continue assigning channel frequencies to the wireless APs, ensuring that any two wireless APs with overlapping coverage are assigned different channel frequencies.

**Example: An Enterprise Corporation Determines IEEE 802.11b Channels**

An enterprise corporation occupies multiple floors in a building. Because of this, they had to pay attention to both the horizontal and vertical dimensions when determining which IEEE 802.11b channel to assign to each AP. For example, if a certain spot on the first floor used channel 1, they assigned channel 6 to the same location on the second floor, and assigned

channel 11 to the same location on the third floor. They did not use channel 1 again until the fourth floor.

Figure 11.5 illustrates the selection of channels for the wireless APs on a building floor. The wireless AP channels were selected to ensure that no two overlapping areas of coverage have the same channel (frequency).

**Figure 11.5 Example of 802.11b Channel Allocation**



# Planning the Deployment of APs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After determining where to locate APs, you need to make additional decisions about how you will deploy the APs. These decisions affect the types of equipment required, where the equipment should be located, and how the equipment will be installed

**Will you install the APs on the walls of your building or in the plenum area?**
If you place your APs in the plenum area, you must determine the best method for powering the APs. Consult with the AP manufacturer to determine how to meet the power requirements for the APs. Some wireless APs can receive electrical power through the Ethernet cable.

The plenum area is regulated by fire codes. Therefore, for plenum placement, you must purchase APs that are fire-rated.

**Will you preconfigure APs before installing them or configure them remotely?**
Preconfiguring the APs before installing them on location can speed up the deployment process and can save labor costs, because less skilled workers can perform the physical installation. You can preconfigure APs by using the console port (serial port), Telnet, or a Web server that is integrated with the AP. If you will use a terminal server for console configuration of APs, decide where to locate the terminal server. Regardless of whether you decide to preconfigure the APs, make sure that you can access them remotely, particularly if you deploy many APs. This enables you to configure or upgrade the APs remotely by using scripts.

The terminal server can be located on the same subnet as the APs, or elsewhere.

It is helpful to physically locate the terminal server in the same wiring closet as the hub or patch panel connecting the AP. With this arrangement, you can pull together the Ethernet and console (serial) wires at the same time. This is only possible if the wires are run to and from the same place. You can then complete the configuration of many APs from a central location.

**What model of antenna will you use for the APs?**
For example, in a building with multiple floors, an omnidirectional antenna — which propagates the signal equally in all directions except the vertical — might work best. For information about which type of antenna will work best for your WLAN deployment, see the documentation for your APs.

**What model of AP will you purchase?**
The AP must be able to be configured as a RADIUS client and must support Wired Equivalent Privacy (WEP) and 802.1X authentication.

**Example: An Enterprise Corporation Mounts APs in the Plenum Area**
An enterprise corporation learned that mounting its APs in the plenum area would resolve some of the issues raised during their pilot test of their wireless network deployment.

For instance, users often disconnected data and power cables from the APs in order to plug their portable computers into the network. The unavailable APs prompted many service calls to the Help desk.

In addition, the IT staff found that by mounting APs in the plenum area, they could install APs in doorways and halls, avoiding users' offices. Though the initial installation cost was higher, they believe the placement will pay off in the long term in decreased user interference during working hours.

The enterprise corporation worked closely with the manufacturer of the APs to ensure that all procedures were followed and that the wireless APs were fire-rated for plenum placement.

# Designing for Wireless Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In designing security for your wireless LAN, choose the appropriate level of basic security available under IEEE 802.11 and 802.1X. Then, to close inherent security risks associated with wireless networking, require authorization and authentication of wireless clients before they exchange data with the network attached to the wireless APs, and encrypt the data sent between wireless clients and APs.

**Choosing the Right Basic Security for Your WLAN**
IEEE 802.11 open system or shared key authentication does not scale appropriately for a large, infrastructure mode wireless network (corporate offices and public places, such as airports and malls). In a large enterprise environment, you should not deploy 802.11 without also deploying 802.1X and RADIUS support.

To ensure the highest level of security for a WLAN in a corporate enterprise environment, use 802.1X with EAP-TLS authentication, a PKI, and RADIUS. The wireless clients must support 802.1X in a WLAN deployment using EAP-TLS.

Microsoft 802.1X Authentication Client provides 802.1X support for computers running any of the following Windows operating systems: Microsoft® Windows® 2000, Windows® 98, and Windows NT® version 4.0 Workstation.

**Note**

- The Microsoft 802.1X Authentication Client is only available for Windows® 98 and Windows NT® version 4.0 Workstation to customers who have Microsoft Premier Support.

Windows XP provides 802.1X support and additional wireless support, including automatic wireless configuration.

**Closing Inherent Security Risks for WLANs**
While providing convenience, wireless networking technologies and wireless APs present security risks. In wireless networks, the signals can be intercepted because the data is broadcast using an antenna. Furthermore, if the signals are not encrypted, an eavesdropper outside the premises can view the packets sent on a wireless network.

Wireless networking technologies and wireless APs present two security risks:

- Any person with a compatible wireless network adapter can associate with your wireless APs and attach to your network.
- Because wireless networking signals use radio waves to send and receive information, anyone within a certain distance of a wireless AP can detect and receive all data sent to and from the wireless AP.

**Enforcing authorization and authentication**
To counter the first security risk, wireless APs must require authentication and authorization of the wireless client before data can be sent to, and received from, the network attached to the wireless AP.

The solution is to use the combination of an 802.1X-enabled wireless client, such as Windows XP; an IEEE 802.1X-enabled and RADIUS-capable wireless AP; and EAP-capable RADIUS servers such as Windows Server 2003 IAS.

With this combination, wireless APs can send connection requests and accounting messages to central RADIUS servers. The RADIUS servers have access to a user accounts database, such as Active Directory, and a set of rules for granting authorization, such as IAS remote access policies. The RADIUS server processes the wireless AP connection request, and either accepts or rejects it.

**Encrypting data**
To counter the second security risk, encrypt the data sent between the wireless clients and the wireless APs. The method of encryption defined by the IEEE 802.11b standard is Wired Equivalent Privacy (WEP). To provide per-session strong cryptographic keys for WEP encryption, use EAP-TLS, PEAP-TLS, or PEAP-MS-CHAP v2 as the authentication method. IAS has been enhanced to support both PEAP-MS-CHAP v2 and EAP-TLS.

EAP-TLS, as defined in RFC 2716, is the TLS authentication scheme as an EAP type. TLS is used in certificate-based security environments. EAP-TLS is a secure channel (SChannel) authentication protocol that provides for mutual authentication, integrity-protected cipher-suite negotiation, and key exchange between the two endpoints by means of public key cryptography.

PEAP-MS-CHAP v2 provides a secure wireless authentication solution for small businesses without requiring a certificate infrastructure (PKI) and the installation of a user or computer certificate on each wireless client. With PEAP, you can use a password-based authentication method to securely authenticate wireless connections. PEAP creates an encrypted channel before the password-based authentication occurs. Therefore, password-based authentication exchanges such as occur in MS-CHAP v2 are not subject to offline dictionary attacks.

**Note**

- PEAP with MS-CHAP v2 is provided with Windows XP Service Pack 1 (SP1) and later, Windows Server 2003, and Microsoft 802.1X Authentication Client.

# Designing Unauthenticated Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Unauthenticated EAP-TLS access can be useful in both corporate and public space environments.

In a corporate environment, this feature can be used to grant guest access to visitors such as consultants. The unauthenticated users are redirected to a specific virtual LAN (VLAN), which provides only limited network access, such as access to the Internet.

In a public space environment, a wireless Internet service provider (ISP) can use this feature to give potential subscribers access to a restricted VLAN with local information. When the person subscribes for Internet access, the ISP provides connectivity to the Internet.

EAP-TLS unauthenticated access provides a means to grant guest access for a wireless client that does not have a certificate installed. EAP-TLS supports one-way authorization or unauthenticated access when a client does not send credentials. If a network access client does not provide credentials, IAS determines whether unauthenticated access is enabled in the remote access policy that matched the connection attempt.

Windows Server 2003 and IAS support unauthenticated wireless connections. For more information about unauthenticated wireless access, see "Wireless access" in Help and Support Center for Windows Server 2003.

# Designing a Public Space WLAN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you plan to deploy a public space WLAN in a venue such as an airport or shopping mall, you need to design your WLAN to meet some additional requirements.

- Plan for a single wireless network infrastructure that multiple service providers can share and access.

  A single wireless network infrastructure eliminates radio frequency interference. Because of the finite number of non-overlapping channels available in 802.11b, multiple wireless network infrastructures in the same location cause interference among wireless APs with overlapping channel frequencies.

- Make sure that the APs support VLANs, the capability for beaconing multiple Service Set Identifiers (SSIDs, also known as *network names*), and the capability for binding each SSID to a separate VLAN.

  Enhanced APs are necessary in a public space WLAN deployment. VLAN support enables the AP to route the wireless client to the correct network path. The capability for beaconing multiple SSIDs enables multiple service providers to share the same wireless network infrastructure. After the wireless client associates with the correct SSID, the AP must bind that SSID to the correct VLAN in order to route the network traffic to the correct destination. The AP maintains a table that maps each SSID to its respective VLAN number. The public space WLAN also must allow non-802.1X wireless clients access. To support this, you must assign a VLAN number for all non-802.1X wireless clients. The VLAN number routes the non-802.1X clients to a VLAN that is configured to provide non-802.1X clients with 802.1X credentials.

- To provide security, you need an IEEE 802.1X and RADIUS-capable wireless AP, and an EAP-capable RADIUS server such as Windows Server 2003 IAS.

- You might need to provide billing and accounting for services provided to customers connecting through the public space WLAN.

  A public space WLAN must provide a means for charging for services provided, typically by an ISP, to customers connecting through the public space WLAN. An ISP can charge the customer for this service in several ways. It can bill for the total time connected, the quantity of data transferred, or a combination of the two methods.

  You can configure the same IAS server that is used for the authorization of wireless users to capture this connection data and save it to an accounting log file. The log file contains the connection time, the amount of data transferred during a session, and other data that can be used to produce billing records for ISP customers. Database exporting can convert the log files into a format that can be read and interpreted to provide the billing records. IAS for Windows Server 2003 can also be configured to send accounting information to a SQL server database.

  In addition, third-party software is available to create billing solutions.

- Provide sufficient bandwidth to support the volume of users likely to use a public space WLAN.

  In designing a public space WLAN, consider how many users need to connect simultaneously through each AP. For example, if you design for an average bandwidth of 56 kilobits per second (Kbps), approximating a 56K modem, more users will be able to associate with the network than if you design the average bandwidth to be more than 56 Kbps.

Figure 11.6 shows the infrastructure for a public space WLAN with 802.1X designed for an airport. This public space WLAN enables ISPs to provide Internet access for general and corporate users with wireless devices that are 802.11b-capable.

**Figure 11.6 Example of a Public Space WLAN Infrastructure in an Airport**




**Example: Public Space WLAN Access**
A user with a computer running Windows XP (with the Wireless Zero Configuration [WZC] service and IEEE 802.1X) starts his computer. The wireless adapter attempts to authenticate with an AP, but can only associate with VLAN 0 since it has no authorization for VLAN 1, the ISP network for the airport. When associated with VLAN 0, the wireless device is directed to the airport's ISP Web server.

The Web server queries the user about access to the Internet or another company's network. As a free service, the ISP's Web site provides local travel information, including arrival and departure times and restaurants. If the user chooses to set up an account, the ISP creates the account for billing purposes and provides the wireless user with a certificate to join VLAN 1.

The wireless user now has a certificate and can access VLAN 1. The user is authenticated using IAS, which simultaneously creates or appends a log file. The log file and new user account are both used for billing purposes. The wireless user is granted permission to access the Internet.

If the user decides to access his own corporate network across the Internet, a virtual private network (VPN) connection can be created from the wireless client to a VPN server in the perimeter network.

### Notes

- As an alternative to a VLAN, a public space wireless network can support IP filtering. This requires the use of APs that are capable of IP filtering and can be configured to restrict access to only the IP addresses for the ISP's certificate, DHCP, and Web servers. These servers provide the minimum connectivity and services that are required in order to obtain authenticated access.

- If an AP is associated with repeatedly when you use IP filtering for authentication, the AP can consume the allotted quantity of IP addresses that the DHCP server has set aside, preventing additional wireless clients from obtaining an IP address. Although the infrastructure for IP filtering is less costly, because IP filtering saves a switch and a server, IP filtering is less secure. For these reasons, it is better to use a VLAN than IP filtering for a public space wireless network.

# Designing for Manageability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For a large wireless deployment to be practical, it must be easy to manage. The combination of Windows XP and Windows Server 2003 allows for efficient management of your wireless network.

For optimal manageability of your wireless network, ensure that your wireless clients use Windows XP, which provides support for automatic switching between APs during roaming and support for zero configuration through the WZC. Although you can use other Windows operating systems with Microsoft 802.1X Authentication Client, they do not support zero configuration.

**Automatic Switching Between APs During Roaming**

Windows XP supports automatic switching between APs when roaming, autodetection of wireless networks, 802.1X, and automatic wireless configuration.

Windows XP has improved and built upon the wireless support for clients that Windows 2000 provides. In Windows 2000, *media sense capability* (the capability for detecting an attached network) is used to control the configuration of the network stack and inform the user and applications when the network is unavailable. With Windows XP, media sense capability is used to enhance the wireless roaming experience. This is done by detecting a move to a new AP and then forcing re-authentication and DHCP renewal to ensure appropriate network access during roaming. Windows XP in addition supports autodetection of a wireless network, and automatic wireless configuration with the Wireless Zero Configuration (WZC) service.

**Distributing Certificates Through Autoenrollment**

An IAS server, which acts as a RADIUS server and proxy, also supports EAP–TLS. Because both Windows XP and IAS in Windows Server 2003 support EAP-TLS, the combination gives you support for a strong authentication method and a per-session key management system. To ease deployment, you can distribute the computer and user certificates used for authentication through certificate autoenrollment. *Autoenrollment* is the automatic requesting and issuing of certificates based on Group Policy settings.

Table 11.2 lists the types of certificates (user, computer, or both) for which autoenrollment is supported with each combination of client and server operating system. When using any other clients, such as Windows NT or Windows 98, you must manually enroll each client or use a scripted solution.

**Table 11.2 Support for Autoenrollment of Certificates Provided in Windows**

| Client | Server | Computer | User |
|---|---|---|---|
| Windows XP or Windows Server 2003 | Windows Server 2003 | ● | |
| Windows XP or Windows Server 2003 | Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition | ● | ● |
| Windows XP or Windows Server 2003 | Windows 2000 Server | ● | |
| Windows 2000 | Windows Server 2003 | ● | |
| Windows 2000 | Windows 2000 Server | ● | |

For more information about the design and deployment of IAS, see "Deploying Internet Authentication Service (IAS)" in this book.

**Taking Advantage of Autoconfiguration**

The addition of the Wireless Zero Configuration (WZC) service in Windows XP improves the manageability of your wireless network. The WZC service dynamically selects the wireless network to which to attempt connection based either on your preferences or on default settings. When a more preferred wireless network becomes available, the WZC service automatically selects and connects to that network. If none of the preferred wireless networks is found nearby, the WZC service configures the wireless adapter so that there is no accidental connection until the wireless client roams within the range of a preferred network.

To improve the roaming experience by automating the process of configuring the network adapter to associate with an

available network, Microsoft partnered with 802.11 wireless adapter vendors. The wireless network adapter scans for available networks and passes them to Windows XP, which then configures the wireless network adapter with an available network. If you are not using a WZC-capable network adapter, you must configure the network adapter manually by using the configuration software that the manufacturer provides.

For improved manageability of your wireless network, ensure that your wireless clients are using Windows XP so the WZC service is available. Windows Server 2003 also provides this service, which is known in Windows Server 2003 as the Wireless Configuration service. The Microsoft 802.1X Authentication Client does not provide WZC and roaming support.

**Managing APs Remotely**

For better manageability, design your network so that you can manage your APs from a remote location. You can remotely manage APs by using the AP console port (serial port) and an asynchronous terminal server, a Telnet session, or a Web server that is integrated with the AP.

To configure an AP for network access by an asynchronous terminal server, use the two unused pairs of Ethernet cable to return the serial communication lines to the data closet where they can be connected to an asynchronous terminal server. This enables you to configure an AP remotely if necessary. If you arrange to switch power off remotely, you can also restart APs remotely when they are not responding to a signal from an Ethernet or console port.

Using these methods, you will be able to completely manage APs remotely except when an AP fails and must be repaired or replaced.

**Using Active Directory-based Wireless Network Policies**

To centrally manage the configuration of secure wireless connections for wireless client computers, you can create Active Directory-based wireless network policies that specify the types of networks that users can access, preferred networks, WEP settings, IEEE 802.X settings, and other settings for wireless connections. The settings are configured in Group Policy, in Wireless Network (IEEE 802.1) Policies. The wireless network policy is replicated to computers that are associated with the computer configuration Group Policy object. Users do not need to enter the configurations or select settings.

For example, you can configure the following items in the Wireless Network (IEEE 802.11) Policies settings:

- Types of networks that users can access

  For example, you might restrict users' access to an AP (infrastructure) network only, or to a computer-to-computer (ad hoc) network only.

- Network name (SSID)

- WEP settings

- Enabling of network access control using IEEE 802.1X

- Authentication methods and settings

The Wireless Network (IEEE 802.1) Policies settings are only supported by wireless clients running Windows XP (SP1 and later) and Windows Server 2003.

For information about:

- Opening the Group Policy Object Editor, see "Ways to open Group Policy Object Editor" in Help and Support Center for Windows Server 2003. (Click the **Index** button, and in the keyword box type **Group Policy Object Editor**; then select **opening**.)

- Adding and defining wireless network policies, see "Define Active Directory-based wireless network policies" in Help and Support Center for Windows Server 2003.

- Designing wireless network policies, see "Deploying Security Policy" in *Designing a Managed Environment*.

# Implementing a WLAN Test Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before embarking on a full-scale WLAN deployment on your enterprise network, perform a trial deployment on a WLAN test network in your lab to familiarize yourself with how the technology works and resolve any issues that arise before deploying the enterprise WLAN. After setting up your WLAN test environment, perform a limited test to ensure that all components are working together under a simplified certificate infrastructure. Then expand the test environment, testing your Group Policy design and a three-tier CA infrastructure. Figure 11.7 shows the process for implementing a WLAN test environment.

**Figure 11.7 Implementing Your WLAN Test Environment**



# Setting Up Your WLAN Test Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Keep your initial WLAN test deployment simple so that you can focus on successfully installing and configuring the essential components. To simplify your initial test:

- Test a basic, single-tier PKI.

  When you deploy your enterprise WLAN, it is recommended that you provide the extra security of a three-tiered certificate infrastructure in which the root CA is offline. For guidelines for designing a certificate infrastructure, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services*.

- Omit using Group Policy in your initial test.

  In your production deployment, you will want to use Group Policy to preconfigure and control the following Wireless Network (IEEE 802.11) Policies settings for Windows XP (SP1 and later) and Windows Server 2003 wireless clients:

  - Types of networks that users can access

    For example, you might restrict users' access to an AP (infrastructure) network only, or to a computer-to-computer (ad hoc) network only.

  - Network name (SSID)

  - Wireless network key (WEP) settings

  - Enabling of network access control using IEEE 802.1X

  - Authentication methods and settings

In your initial test deployment, install and configure the following components and services:

- Wireless APs
- Active Directory
- DNS and DHCP
- A single-tier PKI
- At least one RADIUS server (IAS)
- Wireless clients

After successfully deploying your initial WLAN test environment, add the use of Group Policy and deploy a WLAN with a three-tier PKI in your lab, and confirm that your WLAN still functions properly, to prepare for a successful production deployment.

# Configuring Wireless APs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure your wireless APs according to the manufacturer's specifications.

To provide secure wireless connectivity, the APs must support WEP and 802.1X authentication.

Because your WLAN uses 802.1X and RADIUS, you do not need to manually enter the WEP keys: They are generated automatically during the EAP-TLS and PEAP-MS-CHAP v2 authentication process. (Most APs include the option for manually entering the WEP keys on the AP, which you can do if needed for the test deployment.)

In configuring the AP, you enter a value for the RADIUS shared secret, which you will later enter on the IAS server when you add the AP as a RADIUS client on the IAS server.

# Configuring Active Directory for the WLAN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your test deployment must include at least one Active Directory–based domain controller.

Perform the following configuration tasks in Active Directory:

- Configure the remote access permission on wireless user and computer accounts.
- Create a group for wireless user and computer accounts. You will create a remote access policy for the group.
- Register the IAS server in Active Directory.

**Configuring the Remote Access Permission**

To grant wireless user accounts and wireless computer accounts permission to access the network remotely, set the remote access permission in the computer and user accounts.

**To configure the remote access permission on wireless user and computer accounts**

1. On a domain controller (or on a member server on which the Active Directory snap-ins are installed), open Active Directory Users and Computers.
2. To configure each wireless user account with the permission to access the network remotely, right-click the user object in the **Users** folder, click **Properties**, click the **Dial-in** tab, and then under **Remote Access Permission (Dial-in or VPN)**, select either **Control access through Remote Access Policy** (for native-mode domains) or **Allow access** (for mixed-mode domains).
3. To configure each computer account with the permission to create wireless connections, right-click the computer object in the **Computers** folder, click **Properties**, click the **Dial-in** tab, and then under **Remote Access Permission (Dial-in or VPN)**, select either **Control access through Remote Access Policy** (for native-mode domains) or **Allow access** (for mixed-mode domains).

**Create a Group for Wireless Users and Computers**

Create an Active Directory group to contain wireless users and computers. Later in the test deployment, you will create a group-based remote access policy for wireless connections and specify the group.

**To create a group for wireless users and computers**

1. Open Active Directory Users and Computers.
2. Create a group for wireless users. For the test deployment, accept the default group scope — that is, create a global group.

   For information about how to add a group, see "Creating user and group accounts" in Help and Support Center for Windows Server 2003.
3. Add each user account for wireless users and each computer account to be used for wireless access as a group member.

   For information about how to add a member to a group, see "Changing group memberships" in Help and Support Center for Windows Server 2003.

   **Tip**

   - If you are unable to add computer objects when adding members to a group, use the **Object Types** button in the **Select Users, Contacts, Computers, or Groups** dialog box to add computers to the types of objects that you can

add to a group.

**Register the IAS Server in Active Directory**

The next step in setting up your WLAN test environment is to use the IAS snap-in to register the IAS server in Active Directory. The following procedure registers the IAS server by using the Internet Authentication snap-in.

**To register the IAS server in Active Directory**

1. Open the Internet Authentication Service snap-in on the IAS computer.

2. Right-click **Internet Authentication Service (Local)** for the IAS server, and then click **Register Server in Active Directory**.

3. When asked if you want to authorize this computer to read users' dial-in properties for this domain, click **OK**.

   **Note**

   - This procedure registers the IAS server only in its member domain, which is all that you need for your test deployment. For your production deployment, you will need to register the IAS server in its member domain, trusted domains, and so on. For information, see "Enable the IAS server to read user accounts in Active Directory" in Help and Support Center for Windows Server 2003.

# Configuring DNS and DHCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To support wireless computers in your test lab, your test deployment must include one or more servers running the DNS and DHCP services.

Configure the DNS and DHCP services as follows:

1. On the DNS server:

    a. Ensure that the DNS zone in which the wireless computers will register DNS address records is configured for dynamic updates.

    b. Optionally, specify that the DNS zone is an Active Directory-integrated zone, which provides secured updates, and that the DNS zone is to be updated by DHCP.

2. On the DHCP server, configure a separate scope and lease duration for your wireless client computers.

    The DHCP scope should not include the static IP addresses of your APs or any of your servers.

    For more information about configuring DHCP for your wireless deployment, see "Adapting the Network Infrastructure for a WLAN" earlier in this chapter. For information about configuring scopes and lease durations on a DHCP server, see "Deploying DHCP" in this book.

# Deploying a Certificate Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For your initial test lab deployment, use a simple certificate infrastructure. To be able to integrate the certificate services with Active Directory in your test environment (and, later, to use Group Policy to provide easier management of wireless clients), you must install the CA as an enterprise CA. After installing your enterprise root CA, you can install a computer certificate on the IAS server and install user and computer certificates on your wireless computers.

To set up the certificate infrastructure for your initial test environment, perform the following tasks:

- Install a single-tier CA.
- Install a computer certificate on the IAS server.
- Install user and computer certificates on wireless computers.

**Installing a Single-Tier CA**

To keep your initial test deployment simple, install a single-tier CA.

**To install a single-tier CA in your test environment**

- Install the enterprise root CA either on the domain controller or on a separate member server in your test environment.

  You must be logged on as a member of both the Enterprise Admins group and the Domain Admins group for the root domain.

  For installation instructions, see "Install an enterprise root certification authority" in Help and Support Center for Windows Server 2003. For your test lab deployment, you do not need to add certificate templates to the CA or configure the CA to allow subjects to request a certificate based on a template.

**Installing a Computer Certificate on the IAS Server**

On the IAS server, install a computer certificate from the issuing CA, which, in the single-tier CA infrastructure that you will deploy in your WLAN test environment, is the enterprise root CA. For your test lab deployment, use the Certificates Request Wizard located in the Certificates snap-in to obtain a computer certificate.

Start by creating a Certificates console on the IAS server that contains the Certificates - Local Computer snap-in, which you will use to request the computer certificate.

**To install a computer certificate on the IAS Server**

1. Create a Certificates console on your IAS server that contains the Certificates - Local Computer snap-in. For the test lab deployment, name the console Certificates.

   For information about how to add a snap-in to manage certificates, see "Manage certificates for a computer" in Help and Support Center for Windows Server 2003. To perform this task, you must be a member of the Domain Admins group (or a member of the Administrators group on the local computer).

2. Use the Certificates console to request a computer certificate for the IAS server.

   To install a computer certificate, click **Certificates - Local Computer** in the console tree, and select **Computer** as the certificate type (unless your IAS server is also a domain controller, in which case your only option is to select **Domain Controller**). For more information about how to use the Certificates console to request a certificate, see "Request a certificate" in Help and Support Center for Windows Server 2003.

For more information about using the Certificates Request Wizard for installing computer certificates, in addition to two alternative methods, see "Computer certificates for certificate-based authentication" in Help and Support Center for Windows

Server 2003.

**Verifying that the computer certificates meet IAS requirements**

Each computer certificate installed on an IAS server must meet the following requirements:

- The certificate must be installed in the Local Computer certificate store.
- The cryptographic service provider for the certificate must support the secure channel (Schannel) security package. If not, the IAS server cannot use the certificate, and the certificate is not available for selection in the properties of the **Smart Card or other certificate** EAP type in the remote access policy.

The computer certificate for the IAS server must meet additional requirements. The following procedure tells how to verify each requirement.

**To verify that the computer certificate for the IAS server meets all requirements**

1. From the **Certificates** console, double-click the certificate to open it.
2. On the **General** tab, confirm that **You have a private key that corresponds to this certificate** appears.
3. On the **Details** tab, under **Field**, click **Enhanced Key Usage**, and then confirm that there is an object identifier for Server Authentication (1.3.6.1.5.5.7.3.1).
4. On the **Details** tab, under **Field**, click **Subject Alternative Name**, and then confirm that the fully qualified domain name (FQDN) of the computer account for the IAS server (for example, **DNS Name=*IASServerName.TestDomainName*.com**) appears.
5. On the **Certification Path** tab, confirm that a valid certification path appears and that the statement **This certificate is OK** appears.

**Verifying the root CA certificate**

The root CA certificate of the CA that will issue the wireless client computer and user certificates must be installed in the Trusted Root Certification Authorities certificate store. The following procedure tells how to verify this.

**To verify that the root CA is in the Trusted Root Certification Authorities store**

1. From the Certificates console, expand **Certificates - Local Computer**, expand **Trusted Root Certification Authorities**, and then click **Certificates**.
2. In the **Details** pane, confirm that the name of your test lab enterprise root CA appears in the **Issued To** list.

   If the root CA is not in the list, you might need to refresh the display. To do this, **click Action**, and then click **Refresh**.

**Installing User and Computer Certificates on Wireless Clients**

When EAP-TLS is in use, as in your test lab deployment, wireless clients should have both a computer certificate and a user certificate in order to be authenticated to the network. When PEAP-MS-CHAP-v2 is in use, the root CA certificates of the issuing CAs for the computer certificates on the RADIUS servers must be installed on the wireless clients. You can do this manually by importing the root CA certificate on each wireless client, or you can publish the root CA certificate using Group Policy.

For your test lab deployment, use the Certificate Request Wizard located in the Certificates snap-in on the wireless client computer to obtain both a computer certificate and a user certificate for each wireless computer in your test lab.

**Before you begin**

- Connect the wireless client directly to the wired network that contains the CA infrastructure.

  The connection is required in the test environment in order for the wireless client to receive computer and user certificates. In your enterprise environment, this step might not be necessary, depending upon how you decide to deploy certificates.

If you connect the wireless client to the wired network, you can install the user certificate on the wireless client by using the Certificates - Current User snap-in (as described in the procedure), by using autoenrollment, by submitting a certificate request over the Web, or by implementing a CAPICOM program or script. If you prefer not to make a temporary connection between the wireless client and the wired network, you can install the certificate from a floppy disk.

**Note**

- CAPICOM is a COM client, supporting Automation, that performs cryptographic functions (the CryptoAPI) using Microsoft® ActiveX® controls and COM objects.

**To install user and computer certificates on a wireless client**

1. Create a single console that contains two snap-ins, for managing computer certificates and user certificates. For the test deployment, name the console Certificates.

   a. Install the snap-in for computer accounts under the name **Certificates - Local Computer**.

      For information about how to install a snap-in for managing computer certificates, see "Manage certificates for a computer" in Help and Support Center for Windows Server 2003.

   b. Install the snap-in for user accounts under the name **Certificate - Current User**.

      For information about how to install a snap-in for managing user certificates, see "Manage certificates for your user account" in Help and Support Center for Windows Server 2003.

   To install both snap-ins, log on under a user account with administrative credentials for the local computer. (You can install the user certificates snap-in but not the computer certificates snap-in if you log on under a user account in the test domain.)

   **Note**

   - For the initial test deployment, to receive computer and user certificates, the wireless client must be connected directly to the wired network that has the CA infrastructure.

2. Use the **Certificates - Local Computer** snap-in to request a computer certificate for the wireless client.

   For instructions telling how to use the Certificates console to request a computer certificate, see "Request a certificate" in Help and Support Center for Windows Server 2003.

   The Help topic provides instructions for requesting a *user* certificate. To request a *computer* certificate, instead of clicking **Certificates - Current User** in the console tree, click **Certificates - Local Computer**. Then, when prompted for a certificate type, select **Computer**.

3. Use the **Certificates - Current User** snap-in to request a user certificate.

   For instructions telling how to use the Certificates console to request a user certificate, see "Request a certificate" in Help and Support Center for Windows Server 2003. When prompted for a certificate type, select **User**.

**Verifying that the certificates meet all requirements**

After installing the computer and user certificates, perform the following procedures to verify that the certificates meet all requirements for the client to perform properly over a wireless connection.

**To verify that the computer certificate for the wireless client meets requirements**

1. Verify that the computer certificate is installed in the Local Computer certificate store (required for EAP-TLS

authentication).

After verifying the correct certificate store, verify the certificate configuration.

2. From the **Certificates** console, double-click the certificate to open it.

3. On the **General** tab, confirm that the statement **You have a private key that corresponds to this certificate** appears.

4. On the **Details** tab, under Field :

    a. Click **Enhanced Key Usage**, and confirm that the object identifier for Client Authentication is 1.3.6.1.5.5.7.3.2.

    b. Click **Subject Alternative Name**, and confirm that the FQDN of the wireless computer account (for example, **DNS Name=*LaptopName.TestDomainName*.com**) appears.

5. On the **Certification Path** tab:

    a. Confirm that a valid certification path appears.

    b. Confirm that the statement **This certificate is OK** appears.

**To verify that the user certificate for the wireless client meets requirements**

1. Verify that the user certificate is installed in the Current User certificate store (required for EAP-TLS authentication).

2. From the **Certificates** console, double-click the certificate to open it

3. On the **General** tab, confirm that **You have a private key that corresponds to this certificate** appears.

4. On the **Details** tab, under **Field**, confirm the following items:

    a. Click **Enhanced Key Usage**, and confirm that the object identifier for Client Authentication is 1.3.6.1.5.5.7.3.2.

    b. Click **Subject Alternative Name**, and confirm that the universal principal name (UPN) of the user account (**PrincipalName=*WirelessUserName@TestDomainName*.com**, for example) appears.

5. On the **Certification Path** tab:

    a. Confirm that a valid certification path appears.

    b. Confirm that the statement **This certificate is OK** appears.

# Configuring the RADIUS (IAS) Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you configure your RADIUS server, open Active Directory Users and Computers and verify that your IAS server is a member of the RAS and IAS Servers group.

Configure your RADIUS (IAS) server by performing the following tasks:

1. Add each wireless AP to the IAS server as a RADIUS client.

2. Create a remote access policy for wireless clients.

**Adding APs as RADIUS Clients**

On the IAS server, add each wireless AP as a RADIUS client. You will need to type the RADIUS shared secret that you configured earlier on the wireless AP.

**To add a wireless AP as a RADIUS client on the IAS server**

1. Open the Internet Authentication Service snap-in.

2. In the console tree, right-click the **RADIUS Clients** folder, and then click **New RADIUS Client**.

3. In the **Friendly name** field, type a name for the AP.

4. In the **Client address (IP or DNS)** field, type the IP address of the wireless AP. Then click **Next**.

5. If the remote access policy for wireless users is designed for a specific model of wireless AP (for example, a remote access policy that contains vendor-specific attributes), in the **Client Vendor** list, select the manufacturer's name.

   If you do not know the manufacturer, accept the default value, **RADIUS Standard**.

6. In the **Shared secret** and **Confirm shared secret** fields, type the shared secret value that you assigned when you configured the AP.

**Creating a Remote Access Policy for Wireless Clients**

To give wireless users access to the network, create a remote access policy for wireless clients, and then configure that policy for the highest level of encryption. To use IAS, you must be logged on using an account that has administrative credentials.

**To add a remote access policy for wireless clients**

1. Open the Internet Authentication Service snap-in.

2. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**.

3. Complete the New Remote Access Policy Wizard using the information provided in Table 11.3. Accept default settings when no information is specified.

   **Table 11.3 Adding a Remote Access Policy for Wireless Users**

| Wizard Page | Action |
|---|---|
| Policy Configuration Method | For **Policy Name**, type an appropriate name, such as WLAN Test Policy. |

| | |
|---|---|
| Access Method | Select **Wireless**. |
| User or Group Access | Click **Group**, and then click **Add**.<br><br>In the **Select Groups** dialog box, type the name of the group that you created for wireless users, and then click **Check Names** to confirm that the name you typed is correct. |
| Authentication Methods | Select **Smart Card or other certificate**. |

**To configure encryption for the new remote access policy**

1. In the console tree of the Internet Authentication Service snap-in, right-click the newly created wireless access policy, and then click **Properties**.

2. Verify that **Grant remote access permission** is selected, and then click **Edit Profile**.

# Configuring the Wireless Adapter on Wireless Clients

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

On each wireless client, you can manually configure the wireless adapter to recognize the wireless network that the client must access in order to gain wireless connectivity. However, this is generally not necessary, because Windows XP will sense the wireless network and prompt the user via the notification bar. Once the user selects the network name, the network is automatically added to the preferred networks list.

**To manually configure the client computer's wireless adapter to recognize the AP**

1. On the wireless computer, log on under a user account with administrative credentials on the local computer, and open **Network Connections**.

2. Right-click the wireless adapter icon, and click **Properties** to display properties for the wireless network adapter.

3. To configure TCP/IP for the wireless adapter, on the **General** tab, type the static IP address for your DNS server.

4. On the **Wireless Networks** tab, confirm that the name of the wireless network that you added appears in the **Available Networks** list.

5. Disable the network adapter that connects the client computer to the wired network, and then disconnect the client computer from the wire that connects it to the network.

# Testing Your Limited WLAN Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To test the deployment of your wireless network, roam the entire coverage area for your wireless network, associating with one AP after another. Use your floor plan (with the APs marked on it) to mark the areas that provide adequate coverage and those that require more troubleshooting. You should be able to roam around the building, associating with one AP after another, and test applications.

As you roam through coverage areas, perform the following tests to ensure that your wireless network will provide strong, uninterrupted coverage for wireless clients:

- Use the client software that the adapter manufacturer provided for the wireless device to determine that the wireless client associates with the nearest AP.

  If the wireless client does not readily associate with the closest AP when you move from one AP's coverage area to the next, turn the network adapter's radio off and back on using software provided by the adapter manufacturer. This forces the wireless adapter to find the strongest signal, which usually is the closest AP.

  The wireless client's ability to associate with an AP is determined by the error rate of the data packets and the signal strength. If the coverage from the first AP is still strong, the wireless transceiver receives few bad packets and maintains its association. If the closest AP is failing to associate with the wireless client, restarting the radio of the wireless network adapter forces the wireless adapter to find the strongest signal, which usually is the closest AP.

- Check the statistics for error rates and signal strength to be sure that they are within limits. Check the AP for throughput to determine whether the data transfer rate is adequate.

The following troubleshooting tools also can be useful when testing and deploying your WLAN:

- Use the Wireless Monitor MMC snap-in, included with Windows Server 2003, to gather and view statistical and configuration information for wireless APs and the Windows Server 2003 wireless client.

- Use a spectrum analyzer to determine the location and strength of interfering signals as you move from one signal area to another. A *spectrum analyzer* measures radio frequency radiation from low to high frequencies across a frequency spectrum. These signals are plotted on a graph that shows their strength and frequency. If necessary, you can shield or move any devices that are causing interference.

- Use a *protocol analyzer* to document usage intervals and traffic load. You can use Network Monitor or third-party tools to capture 802.11 packets sent between a wireless client and a wireless AP. With a protocol analyzer, you can capture 802.11 packets, but cannot view the contents of the encrypted payloads.

# Expanding Your WLAN Test Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After successfully deploying and testing a simple wireless network, you can add more complex features — such as Group Policy settings to more easily deploy and manage wireless clients, and a three-tier CA infrastructure to provide greater security for your enterprise WLAN.

Each time that you add a new component or feature, test your new deployment before expanding your test deployment further.

**Configuring Group Policy Settings**
For your initial test deployment, you configured your wireless clients without creating the Active Directory-based wireless network policies that enable you to preconfigure and replicate the wireless client configuration to all wireless clients. Wireless network policies are created by configuring Wireless Network (IEEE 802.11) Policies settings in Group Policy.

In addition, you did not use Group Policy to configure autoenrollment, which enables you to install certificates for the wireless clients automatically.

Instead, you manually configured some of the wireless client settings and used the Certificates console on the client computer to request the computer certificate. (Alternatively, you could have used Web enrollment to request the user or computer certificate.)

However, in your production WLAN deployment, you will want to use Group Policy to provide easier deployment and management of wireless clients and to enable autoenrollment for the installation of the certificates. Before embarking on an enterprise deployment of your WLAN, configure and test Group Policy settings to enable these features.

**Note**

- To support automatic computer certificate allocation, the issuing CA must be an enterprise CA server running either Windows 2000 or Windows Server 2003. To support automatic user and computer certificate allocation, the issuing CA must be an enterprise CA server running either Windows Server 2003, Enterprise Edition or Windows Server 2003, Data Center Edition.

When you configure Group Policy settings to support your WLAN, decide whether you want to manage wireless connections through the domain or create a separate organizational unit (OU) for this purpose. Using an OU might be more efficient than entering Group Policy settings for the domain, which includes both wired and wireless clients.

**Note**

- If you need to force a Group Policy update on the wireless client during your testing, you can use Gpupdate command-line tool. For Gpupdate parameters, see "Command-line reference" in Help and Support Center for Windows Server 2003.

For more information about:

- Designing OUs, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services*.
- Designing group policies, see "Designing a Group Policy Infrastructure" in *Designing a Managed Environment*.
- Deploying autoenrollment, see "Planning for autoenrollment deployment" in Help and Support Center for Windows Server 2003.
- The appropriate way to open Group Policy Object Editor for a specific type of object, see "Ways to open the Group Policy Object Editor Ways to open Group Policy Object Editor" in Help and Support Center for Windows Server 2003. (Click the **Index** button, and in the keyword box type **Group Policy Object Editor**; then select **opening**.)
- Adding and defining wireless network policies, see "Define Active Directory-based Wireless Network Policies" in Help and

Support Center for Windows Server 2003.

**Installing a Three-Tier CA**

When you deploy your enterprise WLAN, it is recommended that you provide the extra security of a three-tier certificate infrastructure in which the root CA is offline. Therefore, after you finish deploying and testing your WLAN test environment with a single-tier CA, and then introducing Group Policies and retesting, it is a good practice to install a test version of the CA infrastructure that you plan to implement in your enterprise environment in your lab before doing so in your production environment.

For information about designing and deploying a certificate infrastructure, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services*.

# Additional Resources for Deploying a Wireless LAN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Expanding Your WLAN Test Deployment

**Related Information**

- "Deploying DHCP" in this book for more information about deploying a DHCP solution on your network.
- "Deploying Domain Name System (DNS)" in this book for more information about deploying DNS within your client/server infrastructure.
- "Deploying ISA Server" in this book for more information about perimeter networks.
- "Deploying Internet Authentication Service (IAS)" in this book and the *Networking Guide* of the *Windows Server 2003 Resource Kit* (or see the Networking Guide on the Web at http://www.microsoft.com/reskit) for more information about deploying an IAS infrastructure in your network.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Enable the IAS server to read user accounts in Active Directory" in Help and Support Center for Windows Server 2003.
- "Install an enterprise root certification authority" in Help and Support Center for Windows Server 2003.
- "Request a certificate" in Help and Support Center for Windows Server 2003 for information about using the Certificates console to request a certificate.
- "Computer certificates for certificate-based authentication" in Help and Support Center for Windows Server 2003 for information about using the Certificates Request Wizard and other methods for installing computer certificates.
- "Manage certificates for a computer" in Help and Support Center for Windows Server 2003.
- "Manage certificates for your user account" in Help and Support Center for Windows Server 2003 for instructions telling how to install the **Certificates - My User Account** snap-in.
- "Planning for autoenrollment deployment" in Help and Support Center for Windows Server 2003.
- "Define Active Directory-based wireless network policies" in Help and Support Center for Windows Server 2003.
- "Command-line reference" in Help and Support Center for Windows Server 2003 for information about forcing a Group Policy update on a wireless client.
- "Open Group Policy Object Editor" in Help and Support Center for Windows Server 2003. (Click the **Index** button and in the keyword box type **Group Policy Object Editor**; then select **opening**.)
- "Creating user and group accounts" in Help and Support Center for Windows Server 2003.
- "Changing group memberships" in Help and Support Center for Windows Server 2003.

# Deploying Windows Firewall

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To protect your network from the growing number of malicious attacks, you need to adopt a defense-in-depth strategy. A defense-in-depth strategy incorporates a variety of network security technologies into your security architecture and implements those technologies in a layered structure that extends from your perimeter network (outer layer) to every computer in your organization (inner layer). By providing host firewall protection on the innermost layer of your network, Windows Firewall, a new security component in Microsoft® Windows Server™ 2003 with Service Pack 1 (SP1), can be an effective part of your defense-in-depth security strategy.

**In This Chapter**

Overview of Windows Firewall Deployment

Step-by-Step Guide for Using Windows Firewall

Deploying Windows Firewall with Group Policy

Deploying Windows Firewall During an Unattended Installation

Deploying Windows Firewall with a Netfw.inf File

# Overview of Windows Firewall Deployment

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, Windows Firewall is turned off (disabled) in Windows Server 2003 with Service Pack 1 (SP1). In addition, when Windows Firewall is turned on, all unsolicited incoming Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic that uses UDP or TCP is blocked. Therefore, to deploy Windows Firewall you must turn on (enable) Windows Firewall and configure Windows Firewall settings so that unsolicited incoming traffic is allowed to reach the programs and services that are acting as servers, listeners, or peers.

The Security Configuration Wizard (SCW) is the recommended tool for deploying Windows Firewall in small, medium, and large organizations. SCW is an optional component that must be installed through **Add or Remove Programs** in Control Panel. SCW guides you through the process of creating a security policy, based on the roles performed by a given server. Once a policy is created, it can be edited or applied to one or more similarly configured servers. For more information, see Configuring Windows Firewall with SCW on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48116).

In addition to SCW, there are several other deployment tools that are suitable for small-scale and large-scale deployments of Windows Firewall.

**Small-Scale Deployments**
If you are deploying Windows Firewall to a small number of servers, you can use Windows Firewall in Control Panel to turn on Windows Firewall and configure Windows Firewall settings on a server-by-server basis. This deployment method is not efficient if you are deploying more than a few servers, and can result in inconsistencies among your server configurations. Therefore, this deployment method is recommended only for small organizations or for servers that require special configuration settings. For more information about deploying Windows Firewall by using Windows Firewall in Control Panel, see the Step-by-Step Guide for Using Windows Firewall.

**Large-Scale Deployments**
If you are deploying Windows Firewall to a large number of servers, you can use one of the following tools to automate the startup and configuration of Windows Firewall.

**Windows Firewall Group Policy Settings**
Windows Server 2003 with SP1 includes several new Group Policy settings that allow you to configure Windows Firewall using domain-based or local Group Policy. Using the Windows Firewall Group Policy settings is the recommended method for deploying Windows Firewall in organizations that use Active Directory. For more information about deploying Windows Firewall with Group Policy, see Deploying Windows Firewall with Group Policy.

**Unattended Installation Answer File**
Windows Server 2003 with SP1 includes several new answer file entries that allow you to enable or disable Windows Firewall and configure Windows Firewall settings. This deployment solution is recommended if your organization does not use Active Directory or Group Policy and you are rolling out slipstream installations of Windows Server 2003. For more information about using an answer file to deploy Windows Firewall, see Deploying Windows Firewall During an Unattended Installation.

**Netfw.inf Information File**
Windows Server 2003 with SP1 includes a Netfw.inf file that you can use to configure Windows Firewall settings while you are installing SP1 or while you are performing a slipstream installation of Windows Server 2003 with SP1. The Netfw.inf file is primarily used by original equipment manufacturers (OEMs) in a manufacturing environment, but it can also be used in a corporate environment during a large-scale rollout. For more information, see Deploying Windows Firewall with a Netfw.inf File.

# Step-by-Step Guide for Using Windows Firewall

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Set up Windows Firewall for the first time using a four-step process: review Windows Firewall limitations, turn on Windows Firewall, configure Windows Firewall exceptions, and configure advanced settings.

This step-by-step guide shows you how to start and configure Windows Firewall for the first time on a computer that is running Windows Server 2003 with Service Pack 1 (SP1).

**Steps for Using Windows Firewall**

Step 1: Review Windows Firewall limitations.

Step 2: Turn on Windows Firewall.

Step 3: Configure exceptions.

Step 4: Configure advanced settings.

Do not use this guide if you are:

- Using domain-based Group Policy to configure Windows Firewall.
- Using an automated installation technology to configure Windows Firewall during setup, such as a Netfw.inf file or an answer file.
- Using Security Configuration Wizard (SCW) to configure Windows Firewall.

**Note** SCW is the recommended method for starting and configuring Windows Firewall. For more information, see Configuring Windows Firewall with SCW on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48116).

**Step 1: Review Windows Firewall limitations.**

It is recommended that you use Windows Firewall on all of your servers; however, there are a few server configurations on which you cannot run Windows Firewall. To determine whether Windows Firewall is appropriate for your server configuration, see Known issues for managing resets, startup, and shutdown on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48117).

In addition, Windows Firewall is designed to be a supplemental security solution; it should be part of a security architecture that implements a variety of security technologies. For more information, see Best practices for managing Windows Firewall on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48221) and Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729).

You might not want to start Windows Firewall if a server requires you to open numerous ports or allow a large number of applications and services to receive unsolicited traffic. Because a significant volume of network traffic will be allowed to pass through Windows Firewall anyway, by disabling Windows Firewall, you eliminate the operational overhead associated with Windows Firewall configuration and maintenance. You also avoid any performance impact related to Windows Firewall. However, you should closely evaluate the design of any client or server that requires you to open numerous ports. Clients and servers that are configured for numerous roles or to provide numerous services can be a critical point of failure in your organization and might indicate poor infrastructure design.

**Step 2: Turn on Windows Firewall.**

Windows Firewall is turned off by default on Windows Server 2003. When you turn on Windows Firewall, you must also start the Windows Firewall/Internet Connection Sharing service, if it is not already running. If the Windows Firewall/Internet Connection Sharing service is not started, and you attempt to start Windows Firewall, a **Windows Firewall** dialog box will appear in the graphical user interface asking if you want to start the service.

To turn on Windows Firewall, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**To turn on Windows Firewall using Windows Firewall in Control Panel**

1. Open Control Panel, and double-click **Windows Firewall**.

   If a **Windows Firewall** dialog box displays a message asking if you want to turn on the Windows Firewall/Internet Connection Sharing service, click **OK**.

2. On the **General** tab, click **On**.

3. Click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and you see **For your security, some settings are controlled by Group Policy** at the top of the **General** tab, the setting might be managed by Group Policy. In this case, you should not use this step-by-step guide.

If all Windows Firewall settings appear dimmed, and you see **You must be a computer administrator to change these settings** at the top of the **General** tab, you do not have administrative rights to configure Windows Firewall. For more information about administrative rights, see Default local groups on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43150) and Default groups on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43151). For more information about turning on Windows Firewall, see Enabling and disabling Windows Firewall on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48118).

**Step 3: Configure exceptions.**
When you turn on Windows Firewall for the first time, all unsolicited incoming TCP/IP traffic is blocked on all network connections. This means that any programs or system services that are acting as servers, listeners, or peers will be unable to receive traffic through TCP and UDP ports. To allow programs and system services to receive unsolicited traffic through these ports, you must add the program or system service to the Windows Firewall exceptions list. In some cases, if you cannot add a program or system service to the exceptions list, you must determine which port or ports the program or system service uses and add the port or ports to the Windows Firewall exceptions list.

To configure exceptions, do the following:

**Use Windows Firewall notifications to add programs to the exceptions list**

By default, Windows Firewall displays a **Windows Security Alert** dialog box whenever a program attempts to listen for incoming traffic and the incoming traffic is blocked. If you are a member of the Administrators group, the **Windows Security Alert** dialog box will display the option to keep blocking the program or unblock the program. Unblocking a program adds the program to the exceptions list and allows unsolicited incoming traffic to reach the program.

To add programs to the exceptions list, do the following when you see a **Windows Security Alert** dialog box:

1. Verify that the program listed in **Name** is a program that you installed and that it is not a malicious program (malware) or spyware.

2. Hover your cursor over the program name to see the path and file name for the program's executable (.exe) file. Verify that the path and file name are correct.

3. If the program is legitimate (not malware or spyware) and you want it to receive unsolicited incoming traffic, click **Unblock**.

   If the program is a malicious program, click **Keep Blocking**. You should immediately remove any malicious programs from your computer.

   If you are unsure, but you think it might be a legitimate program, click **Ask Me Later**. Windows Firewall will continue to block the program, but will prompt you again later.

Windows Firewall displays a **Windows Security Alert** dialog box only when a program is running and attempting to listen for incoming traffic. If a program is not running or does not attempt to listen for incoming traffic, Windows Firewall does not display a **Windows Security Alert** dialog box. You might see several **Windows Security Alert** dialog boxes over the course of several minutes or several hours as programs and system services start up on your server. You should respond to each of these dialog boxes to be sure add all required programs to the Windows Firewall exceptions list.

For more information, see Managing Windows Firewall Notifications on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48222), Add a program to the exceptions list on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48119), and Add a port to the exceptions list on the Microsoft Web site

(http://go.microsoft.com/fwlink/?LinkId=48120).

**Use the Windows Firewall Settings Technical Reference to determine which programs and ports to add to the exceptions list**

Use the Windows Firewall Settings Technical Reference to find out how to configure program and port exceptions for your specific server configuration. The Windows Firewall Settings Technical Reference provides Windows Firewall configuration settings for server roles, system services, remote administration tools, and optional components.

| **Note** |
| --- |
| If you do not know which server roles or optional components are installed on your server, or you do not know which system services or remote administration tools your server uses, you should use SCW to configure Windows Firewall settings. |

For more information, see Windows Firewall Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43155).

For detailed instructions about adding a program to the exceptions list, see Add a program to the exceptions list on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48119).

For detailed instructions about adding a port to the exceptions list, see Add a port to the exceptions list on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48120).

**Step 4: Configure advanced settings.**
To optimally configure Windows Firewall for your server, you must understand and configure the following:

**Windows Firewall profiles**

Windows Firewall settings can be configured in two profiles: a domain profile and a standard profile. This step-by-step guide has helped you configure settings in one profile only. For more information, see Managing Windows Firewall profiles on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48121).

**Windows Firewall scope settings**

When you configure a program, port, or system service exception, you must also configure scope settings for the exception. Scope settings control from which addresses unsolicited traffic is allowed to originate. For more information about scope settings, see Configuring scope settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48122).

**ICMP settings**

By default, Windows Firewall blocks all incoming Internet Control Message Protocol (ICMP) traffic and some outgoing ICMP traffic. This can prevent you from using certain troubleshooting tools, including the **ping** command. For more information about ICMP settings, see Configuring ICMP settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48224).

**Log file settings**

Windows Firewall has a log file that you can use to troubleshoot and monitor Windows Firewall. By default, the log file is disabled. For more information, see Using the Security Log on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48229) and Interpreting the Windows Firewall Log on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48228).

**Windows Firewall exceptions**

Although you have added program, system service, and port exceptions to the exceptions list, it is likely you still need to add (or remove) exceptions for your server to function optimally. For more information about exceptions, see Managing Program, Port, and System Service Exceptions on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=43261).

For more information about Windows Firewall, see Windows Firewall Help on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=33577).

**See Also**
**Concepts**
Deploying Windows Firewall

# Deploying Windows Firewall with Group Policy

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For organizations that use Active Directory, the recommended method for deploying and managing Windows Firewall is to use the new Windows Firewall Group Policy settings. When you use Group Policy to configure Windows Firewall, administrators will be unable to use the **Netsh firewall** command or Windows Firewall in Control Panel to configure the Windows Firewall settings that are managed through Group Policy. In other words, any Windows Firewall settings that are managed through Group Policy appear dimmed and are not accessible through Windows Firewall in Control Panel and cannot be configured through the **Netsh firewall** command.

You can configure Windows Firewall Group Policy settings in either the domain profile or the standard profile. The domain profile settings are used whenever the computer is joined to a network that contains the domain controllers for the domain in which the computer's computer account resides. The standard profile settings are used whenever the computer is joined to a network that does not contain the domain controllers for the domain in which the computer's computer account resides. Both the domain profile and standard profile contain the same set of Windows Firewall settings.

For more information about the Windows Firewall Group Policy settings, see "Windows Firewall Tools and Settings" in the Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729). For more information about using Group Policy settings to deploy Windows Firewall in Windows XP with Service Pack 2 (SP2), see "Managing Windows XP Service Pack 2 Features Using Group Policy" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=31974). Although the article is written specifically for Windows XP with SP2, most of the concepts and information apply to deployments of Windows Firewall in Windows Server 2003 with SP1.

**See Also**
**Concepts**
Deploying Windows Firewall

# Deploying Windows Firewall During an Unattended Installation

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The answer file for unattended installations of Windows Server 2003 with Service Pack 1 (SP1) has been updated to include a section for configuring Windows Firewall. You can use the new Windows Firewall parameters to turn on Windows Firewall and configure Windows Firewall settings if you are deploying clean installations of Windows Server 2003 with SP1. You cannot use an answer file to perform an unattended installation if you are installing SP1 on a computer that already has Windows Server 2003 installed.

| **Note** |
| --- |
| If you use an unattended installation answer file to automate the configuration of Windows Firewall settings, the settings that you specify in the answer file are deleted and are not reapplied when you restore Windows Firewall default settings. However, if you use a Netfw.inf to configure Windows Firewall settings, the settings specified in Netfw.inf are reapplied to the computer when you restore Windows Firewall default settings. If you want to restore the Windows Firewall settings that you configured during an unattended installation with an answer file, you must configure the settings manually. |

For more information about using the Windows Firewall parameters in an unattended installation answer file, see "Windows Firewall Tools and Settings" in the Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729). For more information about the Windows Firewall parameters that are available for use in an unattended installation answer file, see "Unattended Installation Tools and Settings" in the Unattended Installation Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729).

**See Also**
**Concepts**
Deploying Windows Firewall

# Deploying Windows Firewall with a Netfw.inf File

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the Netfw.inf file to turn on Windows Firewall and configure Windows Firewall settings during an installation of Windows Server 2003 with Service Pack 1 (SP1). The Netfw.inf file is typically used by original equipment manufacturers (OEMs) in a manufacturing setting, but it can also be used in a corporate environment during a large-scale rollout of SP1; however, unattended installation is the preferred method for deploying Windows Firewall in a corporate environment.

| **Note** |
| --- |
| If you use a Netfw.inf file to automate the configuration of Windows Firewall settings, the settings that you specify in Netfw.inf are reapplied to the computer when you restore Windows Firewall default settings. Any Windows Firewall settings that you specify in an answer file, such as Unattend.txt, are deleted and are not restored. If you want to restore the Windows Firewall settings that you configured during installation with an answer file, you must configure the settings manually. |

For more information about using the Netfw.inf file to deploy Windows Firewall, see "Windows Firewall Tools and Settings" in the Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729).

**See Also**
**Concepts**
Deploying Windows Firewall

# Designing and Deploying Directory and Security Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Welcome to *Designing and Deploying Directory and Security Services* of the *Microsoft Windows Server 2003 Deployment Kit*.

Whether you are designing a new Active Directory logical structure, deploying Active Directory for the first time, upgrading an existing Windows environment to Windows Server 2003 Active Directory, or restructuring your current environment to a Windows Server 2003 Active Directory environment, Part I of this book will assist you in meeting all of the Active Directory design and deployment goals for your organization.

You also need to make important decisions early in your deployment regarding the design and deployment of Windows Server 2003 distributed security services for authentication, access control, and, increasingly, certificate use. Part II of this book will assist you in preparing the design and planning documents needed to create a sound foundation of distributed security services.

This book contains the following chapters:

Part I: Designing and Deploying Directory Services

Planning an Active Directory Deployment Project

Designing the Active Directory Logical Structure

Designing the Site Topology

Planning Domain Controller Capacity

Enabling Advanced Windows Server 2003 Active Directory Features

Deploying the Windows Server 2003 Forest Root Domain

Deploying Windows Server 2003 Regional Domains

Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory

Upgrading from Windows 2000 Domains to Windows Server 2003 Domains

Restructuring Windows NT 4.0 Domains to an Active Directory Forest

Restructuring Active Directory Domains Between Forests

Restructuring Active Directory Domains Within a Forest

Part II: Deploying Distributed Security Services

Planning a Secure Environment

Designing an Authentication Strategy

Designing a Resource Authorization Strategy

Designing a Public Key Infrastructure

Planning a Smart Card Deployment

# Planning an Active Directory Deployment Project

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you deploy the Microsoft® Windows® Server 2003 Active Directory® directory service in your environment, you can take advantage of the centralized, delegated administrative model and single sign-on capability that Active Directory provides. After you identify the current environment and deployment goals for your organization, you can create the Active Directory deployment strategy that meets your organization's needs. Testing the deployment in an isolated lab environment and refining the deployment in selected pilot areas of your production environment help to ensure a smooth deployment throughout your organization.

**In This Chapter**

Overview of Planning an Active Directory Deployment Project

Determining Your Active Directory Design and Deployment Strategy

Testing and Verifying the Deployment Process

Additional Resources for Planning an Active Directory Deployment Project

**Related Information**

- For more information about planning, testing, and piloting a deployment project, see "Designing a Test Environment" and "Planning and Testing for Application Deployment" in *Planning, Testing, and Piloting Deployment Projects* in this kit.
- For more information about deploying Windows Server 2003 Domain Name System (DNS), see "Deploying DNS" in *Deploying Network Services* in this kit.
- For more information about Group Policy, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Overview of Planning an Active Directory Deployment Project

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory in the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems allows organizations to simplify user and resource management while creating a scalable, secure, and manageable infrastructure. You can use Active Directory to manage your network infrastructure, including branch office, Microsoft® Exchange Server, and multiple forest environments.

Although the guidelines presented in this book are appropriate for almost all network operating system (NOS) management deployments, the guidelines have been tested and validated specifically for environments that contain fewer than 100,000 users and fewer than 1,000 sites, with network connections of a minimum of 28.8 kilobits per second (Kbps). If your environment does not meet these criteria, consider using a consulting firm that has experience deploying Active Directory in more complex environments.

Deploying Active Directory provides the following benefits to your organization:

- **Simplified administration and resource management.** You can delegate administration to all levels of an organization, and you can use Group Policy to centralize administration.
- **Increased network security and single sign-on for users.** Active Directory supports multiple authentication protocols and X.509 certificates, and provides support for smart cards.
- **Interoperability with other directory services.** Active Directory provides standards-based, open interfaces that interoperate with other directory services and applications, such as e-mail applications.
- **Features that reduce administration costs, increase security, and provide additional functionality.** Application directory partitions allow you to configure application-specific data replication settings on domain controllers. When you raise domain or forest functional levels to Windows Server 2003, you can do the following:
    - Rename domains and domain controllers
    - Establish two-way forest trusts
    - Restructure forests
    - Improve replication
    - Remove some limitations in environments with a large number of sites

Although the Windows Server 2003 Active Directory design and deployment strategies that are presented in this book are based on extensive lab and pilot-program testing and successful implementation in customer environments, you might have to customize your Active Directory design and deployment to better suit specific, complex environments. For more information about deploying Active Directory in a branch office environment, see the *Active Directory Branch Office Planning Guide*. For more information about deploying Active Directory in an Exchange environment, see *Best Practice Active Directory Design for Exchange 2000*. For more information about deploying Active Directory in a multiple forest environment, see *Multiple Forest Considerations*. To download these guides, see the Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, and then click "Planning & Deployment Guides."

This book also provides flowcharts, job aids, and deployment examples to help you optimize your Active Directory design and deployment process.

# Process for Planning an Active Directory Deployment Project

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To plan a Windows Server 2003 Active Directory deployment project, first determine your design and deployment strategy, and then test and verify your design and deployment. Figure 1.1 shows the process for planning your Active Directory deployment project.

**Figure 1.1 Planning an Active Directory Deployment Project**

# Active Directory Background Information

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you design and deploy Windows Server 2003 Active Directory, become familiar with the Active Directory deployment project cycle, as well as Active Directory–related terms that are required for the Windows Server 2003 Active Directory deployment process.

# Active Directory Deployment Project Cycle

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An Active Directory deployment project involves three phases: a design phase, a deployment phase, and an operations phase. During the design phase, the design team creates a design for the Active Directory logical structure that best meets the needs of each division in the organization that will use the directory service. After the design is approved, the deployment team tests the design in a lab environment and then implements the design in the production environment. Because testing is performed by the deployment team and potentially affects the design phase, it is an interim activity that overlaps both design and deployment. When the deployment is complete, the operations team is responsible for maintaining the directory service.

Lab testing and the implementation of a pilot program continue throughout the lifetime of the Active Directory deployment.

Figure 1.2 shows the relationship between the phases of the Active Directory project cycle relative to the lifetime of the deployment project.

**Figure 1.2 Relationship Between Active Directory Project Cycle Phases**



# Terms and Definitions

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following terms are important to understanding the Windows Server 2003 Active Directory deployment process.

**Active Directory domain**
An administrative unit in a computer network that, for management convenience, groups several capabilities, including:

- **Network-wide user identity.** Domains allow user identities to be created once and referenced on any computer that is joined to the forest in which the domain is located. Domain controllers that make up a domain are used to store user accounts and user credentials, such as passwords or certificates, securely.

- **Authentication.** Domain controllers provide authentication services for users and supply additional authorization data, such as user group memberships. These services can be used to control access to resources on the network.

- **Trust relationships.** Domains extend authentication services to users in other domains in their own forest by means of automatic bidirectional trusts, and to users in domains in other forests by means of either manually created external trusts or forest trusts.

- **Policy administration.** The domain is a scope of administrative policies, such as password complexity and password reuse rules.

- **Replication.** The domain defines a partition of the directory tree that provides data that is adequate to provide the required services and that is replicated between the domain controllers. In this way, all domain controllers are peers in a domain and are managed as a unit.

**Active Directory forest**
A collection of one or more Active Directory domains that share a common logical structure, directory schema, and network configuration, as well as automatic two-way transitive trust relationships. Each forest is a single instance of the directory and defines a security boundary.

**Active Directory functional level**
A setting in Windows Server 2003 Active Directory that enables advanced domain-wide or forest-wide Active Directory features.

**Migration**
The process of moving an object from a source domain to a target domain, while preserving or modifying characteristics of the object to make it accessible in the new domain.

**Domain restructure**
A migration process that involves changing the domain structure of a forest. A domain restructure can involve either consolidating or adding domains, and can take place between forests or within a forest.

**Domain consolidation**
A restructuring process that involves eliminating Microsoft® Windows NT® 4.0 domains or Active Directory domains by merging their contents with the contents of other domains.

**Domain upgrade**
The process of upgrading the directory service of a domain to a later version of the directory service. This includes upgrading the operating system on all domain controllers and raising the Active Directory functional level where applicable.

**In-place domain upgrade**
The process of upgrading the operating system on all domain controllers that are based on Windows NT 4.0 or on the Microsoft® Windows® 2000 operating system and raising the functional level of the domain if applicable, while leaving domain objects, such as users and groups, in place.

**Regional domain**
A child domain that is created based on a geographic region in order to optimize replication traffic.

# Determining Your Active Directory Design and Deployment Strategy

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you perform a high-level assessment of your current environment and determine your Active Directory deployment goals, you can determine the deployment strategy that works best for your environment. Figure 1.3 shows the steps for defining the Active Directory deployment process.

**Figure 1.3 Determining Your Design and Deployment Strategy**




The Active Directory deployment strategy that you apply varies according to your existing network configuration. For example, if your organization currently runs Windows 2000, you can simply upgrade your operating system to Windows Server 2003. If your organization currently runs Windows NT 4.0 or a non-Windows network operating system, however, you must design an Active Directory infrastructure before you upgrade to Windows Server 2003.

Your deployment process might involve restructuring existing domains, either within an Active Directory forest or between Active Directory forests. You might need to restructure your existing domains after you deploy Windows Server 2003 Active Directory or after organizational changes or corporate acquisitions. You can also restructure domains from a Windows NT 4.0 environment to an Active Directory forest in order to upgrade your production environment to Windows Server 2003.

Table 1.1 lists the possible starting points and goals for a Windows Server 2003 Active Directory deployment and the corresponding deployment steps and topics in this guide that apply to each.

**Table 1.1 Current Environment, Goals, and Corresponding Topics for Deploying Windows Server 2003 Active Directory**

| Environment | Deployment goals | Corresponding topics |
|---|---|---|
| New organization | Create forest, domain, Domain Name System (DNS), and OU design. | Designing the Active Directory Logical Structure (http://go.microsoft.com/fwlink/?LinkId=132452) |
| Create a site and site link design. | Designing the Site Topology (http://go.microsoft.com/fwlink/?LinkId=132453) | |
| Assess hardware requirements. | Planning Domain Controller Capacity (http://go.microsoft.com/fwlink/?LinkId=132455) | |
| Deploy the forest root domain. | Deploying the Windows Server 2003 Forest Root Domain (http://go.microsoft.com/fwlink/?LinkId=132457) | |

| | | |
|---|---|---|
| Deploy regional domains. | Deploying Windows Server 2003 Regional Domains (http://go.microsoft.com/fwlink/?LinkId=132461) | |
| Raise the domain and forest functional levels. | Enabling Advanced Windows Server 2003 Active Directory Features (http://go.microsoft.com/fwlink/?LinkId=132462) | |
| Windows NT 4.0 | Create forest, domain, DNS, and OU design. | Designing the Active Directory Logical Structure (http://go.microsoft.com/fwlink/?LinkId=132452) |
| Create a site and site link design. | Designing the Site Topology (http://go.microsoft.com/fwlink/?LinkID=132453) | |
| Assess hardware requirements. | Planning Domain Controller Capacity (http://go.microsoft.com/fwlink/?LinkId=132455) | |
| Deploy the forest root domain. | Deploying the Windows Server 2003 Forest Root Domain (http://go.microsoft.com/fwlink/?LinkId=132457) | |
| Deploy regional domains. | Deploying Windows Server 2003 Regional Domains (http://go.microsoft.com/fwlink/?LinkId=132461) | |
| Upgrade in-place Windows NT 4.0 domains that will remain part of your Active Directory domain structure. | Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory (http://go.microsoft.com/fwlink/?LinkId=132466) | |
| Restructure other Windows NT 4.0 domains. | Restructuring Windows NT 4.0 Domains to an Active Directory Forest (http://go.microsoft.com/fwlink/?LinkId=132467) | |
| Raise the domain and forest functional levels. | Enabling Advanced Windows Server 2003 Active Directory Features (http://go.microsoft.com/fwlink/?LinkId=132462) | |
| Windows 2000 | Upgrade Windows 2000 domain controllers. | Upgrading from Windows 2000 Domains to Windows Server 2003 Domains (http://go.microsoft.com/fwlink/?LinkId=132484) |
| Raise the domain and forest functional levels. | Enabling Advanced Windows Server 2003 Active Directory Features (http://go.microsoft.com/fwlink/?LinkId=132462) | |

Table 1.2 lists the goals and corresponding topics that apply to restructuring domains either within or between forests.

**Table 1.2 Goals and Corresponding Topics for Restructuring Active Directory Domains**

| Action | Deployment goals | Corresponding topics |
| --- | --- | --- |
| Restructure domains within a forest | Create forest, domain, DNS, and OU design. | Designing the Active Directory Logical Structure (http://go.microsoft.com/fwlink/?LinkId=132452) |
| Create a site and site link design. | Designing the Site Topology (http://go.microsoft.com/fwlink/?LinkID=132453) | |
| Use a tool such as the Active Directory Migration Tool (ADMT) to restructure domains within a forest. | Restructuring Active Directory Domains Within a Forest (http://go.microsoft.com/fwlink/?LinkId=132486) | |
| Restructure domains between forests | Create forest, domain, DNS, and OU design. | Designing the Active Directory Logical Structure (http://go.microsoft.com/fwlink/?LinkId=132452) |
| Create a site and site link design. | Designing the Site Topology (http://go.microsoft.com/fwlink/?LinkID=132453) | |
| Use a tool such as ADMT to restructure domains between forests. | Restructuring Active Directory Domains Between Forests (http://go.microsoft.com/fwlink/?LinkId=132488) | |

# Determining Your Active Directory Design Requirements

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your network environment is currently operating without a directory service, or if you need to modify your current Active Directory infrastructure, complete the design process for your Active Directory infrastructure. You must complete a comprehensive design of your Active Directory logical structure before you deploy Active Directory. Thoroughly preparing your Active Directory design is essential to a cost-effective deployment.

**Logical Structure Design**
Before you deploy Windows Server 2003 Active Directory, you must plan for and design the Active Directory logical structure for your environment. The Active Directory logical structure determines how your directory objects are organized, and provides an effective method for managing your network accounts and shared resources. When you design your Active Directory logical structure, you define a significant part of the network infrastructure of your organization.

To design the Active Directory logical structure, determine the number of forests that your organization requires, and then create designs for domains, DNS, and organizational units.

**Site Topology Design**
After you design the logical structure for your Active Directory infrastructure, you must design the site topology for your network. The site topology is a logical representation of your physical network. It contains information about the location of Active Directory sites, the Active Directory domain controllers within each site, and the site links that support Active Directory replication between sites.

**Domain Controller Capacity Planning**
To ensure efficient Active Directory performance, you must determine the appropriate number of domain controllers for each site and verify that they meet the hardware requirements for Windows Server 2003. Careful capacity planning for your domain controllers ensures that you do not underestimate hardware requirements, which can cause poor domain controller performance and application response time.

**Advanced Active Directory Features**
Functional levels in Windows Server 2003 Active Directory allow you to enable new features, such as improved group membership replication, deactivation and redefinition of attributes and classes in the schema, and forest trust relationships that require that all domain controllers within the participating domain or forest run Windows Server 2003. Part of the Active Directory design process involves identifying the domain and forest functional levels that your organization requires. To implement these Windows Server 2003 Active Directory features in your organization, you must first deploy Windows Server 2003 Active Directory and then raise the forest and domain to the appropriate functional level.

# Determining Your Active Directory Deployment Requirements

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The structure of your existing environment determines your strategy for deploying Windows Server 2003 Active Directory. If you are creating an Active Directory environment and you do not have an existing domain structure, you must complete your Active Directory design before you begin creating your Active Directory environment. Then you can deploy a new forest root domain and deploy the rest of your domain structure according to your design.

**Windows Server 2003 Forest Root**
To deploy Active Directory, you must first deploy a Windows Server 2003 forest root domain. To do this, you must configure DNS, deploy forest root domain controllers, configure the site topology for the forest root domain, and configure operations master roles.

**Windows Server 2003 Regional Domains**
If you are creating one or more new regional domains in a Windows Server 2003 forest, you must deploy each regional domain after you deploy your forest root domain. To do this, you must delegate a DNS zone and deploy domain controllers for each regional domain.

**Windows NT 4.0 Domain Upgrade to Windows Server 2003**
When you perform an in-place domain upgrade of Windows NT 4.0 domains, you can begin to use Active Directory without making any modifications to your existing domain structure.

Alternatively, if you do not want to retain your existing domain structure, you can restructure your Windows NT 4.0 domains to a Windows Server 2003 forest. For more information about restructuring your Windows NT 4.0 domains to a Windows Server 2003 forest, see "Determining Your Restructure Requirements" later in this chapter.

**Windows 2000 Domain Upgrade to Windows Server 2003**
Upgrading your Windows 2000 domains to Windows Server 2003 domains is an efficient, straightforward way to take advantage of additional Windows Server 2003 features and functionality. Upgrading from Windows 2000 to Windows Server 2003 requires minimal network configuration and has little impact on user operations.

# Determining Your Restructure Requirements

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As part of your Active Directory deployment, you might choose to restructure your environment. Before doing so, you must determine when and how you want to restructure your environment. Organizations with an existing Windows NT 4.0 domain structure might perform an in-place upgrade of some domains and restructure others. In addition, you might decide to reduce the complexity of your environment by either restructuring domains between forests or restructuring domains within a forest after you deploy Active Directory.

**Windows NT 4.0 Domain Restructure to a Windows Server 2003 Forest**
Because of its greater scalability, a Windows Server 2003 Active Directory environment requires fewer domains than a Windows NT 4.0 environment. Instead of performing an in-place upgrade of your Windows NT 4.0 domains, it might be more efficient to consolidate a number of smaller Windows NT 4.0 account and resource domains into a few, larger Active Directory domains.

**Interforest Active Directory Domain Restructure**
When you restructure domains *between* Windows Server 2003 forests, you can reduce the number of domains in your environment and, therefore, reduce administrative complexity and overhead. When you migrate objects between forests as part of the restructuring process, both the source and target domain environments exist simultaneously. This enables you to roll back to the source environment during the migration, if necessary.

**Intraforest Active Directory Domain Restructure**
When you restructure Windows Server 2003 domains *within* a Windows Server 2003 forest, you can consolidate your domain structure and, therefore, reduce administrative complexity and overhead. Unlike the process for restructuring Windows Server 2003 domains between forests, when you restructure domains within a forest, the migrated accounts no longer exist in the source domain.

Table 1.3 lists the differences between an interforest and an intraforest domain restructure.

**Table 1.3 Differences Between Interforest and Intraforest Domain Restructures**

| Migration Consideration | Interforest Restructure | Intraforest Restructure |
|---|---|---|
| **Object preservation** | Objects are cloned rather than migrated. The original object remains in the source location to maintain user access to resources. | Objects are migrated and no longer exist in the source location. |
| **SID history maintenance** | Maintaining SID history is optional. | SID history is required. |
| **Password retention** | Password retention is optional. | Passwords are always retained. |
| **Local profile migration** | You must use tools such as ADMT to migrate local profiles. | For workstations that run Windows 2000 and later, local profiles are migrated automatically because the user's GUID is preserved. However, you must use tools such as ADMT to migrate local profiles for workstations that run Windows NT 4.0 and earlier. |
| **Closed sets** | You do not need to migrate accounts in closed sets. | You must migrate accounts in closed sets. |

# Example: Establishing an Active Directory Deployment Strategy

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To illustrate the Active Directory deployment process, the chapters in this book show an example of how a fictitious company, Contoso Pharmaceuticals, deploys Active Directory in its environment. The Contoso environment consists of four domains, all of which are running Windows 2000 Active Directory. Figure 1.4 shows the current domain structure for the Contoso corporation.

**Figure 1.4 Contoso Corporation Domain Structure**




After reviewing its existing environment and identifying its deployment goals, Contoso established the following Active Directory deployment strategy:

- Upgrade Windows 2000 domains to Windows Server 2003 domains.
- Enable advanced Active Directory features by raising the domain and forest functional levels to Windows Server 2003.

After upgrading all Windows 2000 domains to Windows Server 2003 domains, Contoso will restructure the africa.concorp.contoso.com domain within the forest to consolidate it with the emea.concorp.contoso.com domain.

The Contoso corporation is acquiring a company called Trey Research, which is currently running a Windows NT 4.0–based environment, as shown in Figure 1.5.

**Figure 1.5 Current Environment for Trey Research**




Contoso established the following Active Directory deployment strategy for their Trey Research acquisition:

- Design the Active Directory logical structure to create forest, domain, DNS, and organizational unit designs for the new Windows Server 2003 environment.
- Design the site topology to create the required sites, site links, and site link bridges.
- Plan domain controller capacity to determine the hardware requirements for the new Windows Server 2003 environment.
- Deploy trccorp.treyresearch.net as the forest root domain.
- Deploy three regional domains. Different teams can create these domains simultaneously.
    - Upgrade the EAST domain to Windows Server 2003 to become east.trccorp.treyresearch.net.
    - Create two new Windows Server 2003 regional domains called asia.trccorp.treyresearch.net and

west.trccorp.treyresearch.net.

- Restructure the BOSTON, MAIL-APPS, PROD-APPS, and OFFICE-APPS domains to the east.trccorp.treyresearch.net Windows Server 2003 domain by using ADMT.

- Raise the domain and forest functional levels to Windows Server 2003.

Figure 1.6 shows the interim environment for Trey Research.

**Figure 1.6 Interim Environment for Trey Research**




At a later time, Contoso determined that a single domain would be more cost-effective for the Europe, Middle East, and Asia region, so the final step in the deployment process is to restructure asia.trccorp.treyresearch.net into the emea.concorp.contoso.com domain in the Contoso forest by using ADMT.

Figure 1.7 shows the domain structure for the Contoso corporation after the acquisition of Trey Research and the Windows Server 2003 Active Directory deployment process is complete.

**Figure 1.7 Final Environment for Contoso and Trey Research**



# Testing and Verifying the Deployment Process

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In any Active Directory deployment, you can minimize the impact on normal business operations by testing design assumptions and verifying the deployment process in a pilot program. As you create the first draft of your Active Directory design, begin to test and verify. Testing and verifying begin during the design phase and continues through the deployment and operations phase.

Figure 1.8 shows the process for testing and verifying your Active Directory design and deployment.

**Figure 1.8 Testing and Verifying the Active Directory Design and Deployment**



# Testing the Design and Deployment in a Lab Environment

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Lab testing is the first evaluation of the Active Directory design. During lab testing, you confirm the assumptions made by the design architects.

As the first draft of the Active Directory design approaches completion, begin testing specific design assumptions in the deployment process in a lab environment. By testing the deployment process in your lab, you can discover potential design problems that affect the deployment process and provide feedback to the design team to correct problems before the deployment.

Ensure that the test lab environment is isolated from the rest of your organization's production network and represents, on a small scale, the hardware and operating system configuration of the computers in your organization. Include enough domain controllers in the lab environment to support a representative sample of your site design, including intrasite and intersite replication partners, site links, and realistic replication intervals.

Include user and group accounts and other resources that are exclusively designated for testing. Ensure that your test environment provides access to test configurations of external services, such as mainframe or Internet access, as required. Retain the lab permanently to test new procedures and train the deployment team.

The deployment team can use the lab environment to learn the specifics of your deployment process and to gain familiarity with the deployment and migration tools that are used during the Active Directory deployment.

Typically, the design assumption tests and the deployment process tests are performed by different teams. Table 1.4 lists the lab tests and the team members who perform the tests in the lab.

**Table 1.4 Lab Tests and Corresponding Team Members**

| Test Process | Lab Tests | Team Members |
|---|---|---|
| Test design assumptions | Analyze Active Directory replication and site topology. | Design team, site topology owner, and deployment team. |
| Test application and desktop compatibility. | Design team. | |
| Test deployment process | Test disaster recovery. | Forest owner and deployment team. |
| Test account and resource migration. | Forest owner and deployment team. | |
| Evaluate delegation, administration, and management. | Forest owner. | |

# Testing Design Assumptions

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

During the design process, the design team makes assumptions that are incorporated into the Active Directory design, such as Active Directory replication and application compatibility. After the team completes a preliminary draft of the design, it must prove these assumptions in the lab environment.

To test the design assumptions in the lab environment, the design team must:

- Analyze Active Directory replication site topology.
- Develop a test plan, and then test application and desktop compatibility.

**Analyze Active Directory Replication and Site Topology**

The site topology design specifies the maximum replication latency. This is the length of time that is required to replicate changes throughout the forest. The design team must make sure that forest-wide replication latency is less than or equal to the maximum replication latency specified in the design. The team must perform a worst-case test that is based on the maximum number of hops that are assumed in the design. The team must observe the time that is required for replication convergence when a domain controller or communications link fails.

**To analyze Active Directory intersite replication site failover**

1. Identify the domain controllers that are responsible for intersite replication by using Active Directory Sites and Services.
2. Disconnect domain controllers or disable communications links that are used in intersite replication.
3. Allow the Knowledge Consistency Checker (KCC) to automatically configure new replication topology.
4. Identify the domain controllers that are now responsible for intersite replication.
5. Reconnect the domain controllers or enable communications links.
6. Verify that the intersite replication topology returns to the original state, as identified in step 1.

**Verify Application and Desktop Compatibility**

The design team must also determine the compatibility between applications, desktop operating systems, and Active Directory. Typically, the aspects of application testing that are affected by an Active Directory migration or upgrade include applications that run on servers and client computers, in addition to remote access usage.

Verify the application and desktop compatibility design assumptions by creating a list of all critical applications. Have design team members test each application to make sure that it operates correctly in a migrated environment.

When you verify application and desktop compatibility, verify that:

- Existing server applications, such as those that currently run on a Windows NT 4.0 backup domain controller (BDC), can run on Windows Server 2003–based member servers and domain controllers.

  For example, some server applications that run on BDCs take advantage of Shared Local Groups. To run these server applications on a Windows Server 2003–based domain controller, verify that the applications run correctly by using Active Directory domain local groups.

- Server applications that run on a mixture of Windows Server 2003–based and Windows NT 4.0–based servers can interoperate with one another.

  For example, verify that a Windows Server 2003–based server running Microsoft® SQL Server™ can interact with a Windows NT 4.0–based server running the same application.

- Existing desktop applications run correctly when the domain infrastructure is migrated to Windows Server 2003 Active Directory.

- Existing applications that use integrated Windows security run correctly when the domain infrastructure is migrated to Windows Server 2003 Active Directory.

If you find that a server application cannot be migrated to a Windows Server 2003–based domain controller, you can try to reinstall the application or a later version of the application on a Windows Server 2003–based member server. If the application cannot run on a server that runs Windows Server 2003, you can continue to run the application on the server that runs Windows NT 4.0 or Windows 2000.

Provide feedback to the design team that the server application's domain cannot be upgraded in-place or consolidated and must remain until a version of the application that can run on a Windows Server 2003–based domain controller is available. As a long-term deployment goal, transition any applications that currently run on domain controllers to member servers.

# Testing Deployment Processes

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a lab environment before the pilot program begins, the deployment team must test specific tasks that are essential to the Active Directory deployment process, such as testing account and resource migration from Windows NT 4.0 to Windows Server 2003 Active Directory.

To verify the deployment process in the lab environment:

- Test disaster recovery.
- Test account and resource migration.
- Evaluate delegation, administration, and management.

**Test Disaster Recovery**

Test disaster recovery in your lab environment to validate that users can log on within an acceptable response time until a failed domain controller is restored and to determine the time that is required to restore the failed domain controller.

To implement a disaster recovery process in your Active Directory deployment, back up the System State data on at least two domain controllers in the lab environment. After you back up the data, you need to test the validity of the backup tape and the restore process. Test the following scenarios:

- Perform a non-authoritative restore of the domain controller whose directory services database contains corrupted data.
- Perform an authoritative restore of a domain controller to restore Active Directory data that has been deleted.

Make sure that the tests represent the slowest connection speeds in your environment and the largest number of user accounts.

For example, when you determine the time that is required to restore a failed domain controller, make sure to test the restore of System State data from your backup for any domain controller that is the only one in a site that is connected with a data rate of 128 Kbps or less. In addition, test the restore of System State data from your backup for any domain controller in a domain that contains more than 20,000 user accounts.

When a domain controller is connected to other domain controllers with a data rate that is equal to or greater than 128 Kbps, test your process for installing Active Directory on a new domain controller and letting Active Directory replication repopulate the Active Directory database.

For more information about testing disaster recovery, see the Active Directory Disaster Recovery (.doc) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Test Account and Resource Migration**

To test the deployment process for account and resource migration, use the procedures in the chapter for the restructure process that you are planning. Organizations that are planning to restructure Windows NT 4.0 domains can also perform the following tests of their restructure process:

**To test the deployment process for account and resource migration**

1. In two or more production Windows NT 4.0 account domains, create new backup domain controllers (BDCs).
2. Remove the new BDCs from the production network.
3. Install the new BDCs in the lab environment.
4. Promote the new BDCs to primary domain controllers (PDCs).
5. Perform in-place upgrades and restructure the account domains in your lab.

6. Perform account and resource migrations by using a migration tool such as ADMT.

7. Verify that migrated accounts have access to resources and retain user profiles.

**Evaluate Delegation, Administration, and Management**
Evaluate the delegation, administration, and management processes by creating the organizational unit structure that is specified by your Active Directory design. Delegate control of organizational units to specific group accounts that are used for administration. Use these steps to verify the success of the delegation:

**To verify successful delegation of control of OUs to specific groups**

1. Log on as a user who belongs to the group account to which you delegated control.

2. Perform administration tasks on objects within the organizational unit (for example, modify the properties of a user in an account organizational unit).

3. Try, and subsequently fail, to perform administrative tasks on organizational units to which the administration group does not have delegated control.

# Verifying the Deployment in a Pilot Program

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the lab environment, you verify that the deployment process works *outside* of your production environment on accounts and resources that approximate your production environment. If your environment runs Windows 2000 Active Directory, use your existing pilot program to verify your Windows Server 2003 deployment. In the pilot deployment program, identify a *controlled* subset of the accounts (users, groups, and services) and resources that exist in the production environment. Perform the deployment process on the identified accounts and resources.

The goals of the pilot program include:

- Extend testing into a subset of the production environment.
- Provide a test environment for other design and deployment groups, such as Exchange 2000 deployment.
- Verify process and procedures for network and operating system infrastructure updates.
- Verify proper operation of application updates.
- Evaluate the impact of monitoring solutions on the network infrastructure and the servers that are being monitored.
- Discover any potential problems in the deployment process that are caused by complexities that could not be modeled in the lab environment.
- Revise the deployment process to correct any problems that you discovered before the production deployment.

In your pilot deployment, begin with users who are involved in the deployment project, and then include users who are representative of your user population.

**To create a pilot deployment program in your environment**

1. Create ***forest_root_domain*** (where *forest_root_domain* is the name of an empty Active Directory forest root domain that was created by appending "-test" to the same name as the production forest root domain).
2. Create ***regional_domain*** (where *regional_domain* is the regional domain name in the pilot program) by appending "-test" to a production regional domain.
3. Establish the appropriate trust relationships between ***regional_domain*** and ***winnt_domain*** (where *winnt_domain* is a Windows NT 4.0 account or resource domain).
4. Migrate selected accounts and resources from ***winnt_domain*** to ***regional_domain***.
5. Verify that users and administrators can minimally perform the same tasks that they performed before the migration (such as resource access, account administration, and resource administration).

**Important**

- When you migrate production users to the pilot, leave the user accounts enabled in the production and the pilot environments. By leaving the user accounts enabled in the production environment, you provide a fallback plan if any problems occur in the pilot environment.

# Example: Creating a Pilot Deployment Program for Trey Research

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso corporation and Trey Research created a pilot deployment program for the Active Directory deployment. Table 1.5 lists the names that they provided for their pilot deployment.

**Table 1.5 Example of a Pilot Deployment Program**

| Domain | Name |
|---|---|
| forest_root_domain | trccorp-test.treyresearch.net |
| regional_domain | east-test.trccorp-test.treyresearch.net |
| winnt_domain | BOSTON for account domain<br>OFFICE-APPS for resource domain |

Figure 1.9 illustrates the pilot deployment configuration.

**Figure 1.9 Pilot Deployment Configuration**



# Completing the Pilot Deployment Program

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you finish the pilot deployment program, retain the pilot deployment environment and use it as a staging environment for your production deployment. Continue to use the pilot forest to verify new deployment processes, such as adding new applications or schema extensions, installing operating systems, creating Group Policy objects, or restructuring organizational units.

As noted earlier, provide a fallback plan by leaving user accounts enabled in both the original production environment and the pilot deployment environment. Keep users in the pilot deployment environment to do their production work; do not migrate the accounts from the pilot deployment environment to the production forest. Instead, if you decide to move pilot users to another production forest, migrate them from their original production environment. This ensures that all users in the production forest have consistent accounts and makes troubleshooting easier.

# Example: Completing the Pilot Deployment Program for Trey Research

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 1.10 shows a comparison between the design of the Active Directory pilot forest and the production forest deployment.

**Figure 1.10 Comparison of the Pilot Forest and the Production Forest**



# Additional Resources for Planning an Active Directory Deployment Project

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing the Active Directory Logical Structure" in this book.
- "Designing the Site Topology" in this book.
- "Planning Domain Controller Capacity" in this book.
- "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.
- "Deploying the Windows Server 2003 Forest Root Domain" in this book.
- "Deploying Windows Server 2003 Regional Domains" in this book.
- "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.
- "Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" in this book.
- "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.
- "Restructuring Active Directory Domains Between Forests" in this book.
- "Restructuring Active Directory Domains Within a Forest" in this book.
- "Deploying DNS" in *Deploying Network Services* in this kit.
- The Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Click "Planning & Deployment Guides" to find additional links where you can download the following guides:
    - *Active Directory Branch Office Planning Guide*
    - *Active Directory Operations Guide*
    - *Best Practice Active Directory Design for Exchange 2000*
    - *Multiple Forest Considerations*
    - *Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations: Part I*

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Active Directory" in Help and Support Center for Windows Server 2003.
- "Install Windows Support Tools" under "Tools" in Help and Support Center for Windows Server 2003.

# Active Directory Maximum Limits - Scalability

Updated: March 12, 2015

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

This topic describes Active Directory scalability and other limitations, as well as recommendations that apply when you are designing or implementing an Active Directory infrastructure. These limitations include the following:


- Maximum Number of Objects
- Maximum Number of Security Identifiers
- Maximum Number of entries in Discretionary and Security Access Control Lists
- Group Memberships for Security Principals
- FQDN Length Limitations
- File Name and Path Length Limitations
- Additional Name Length Limitations
- Maximum Number of GPOs Applied
- Trust Limitations
- Maximum Number of Accounts per LDAP Transaction
- Recommended Maximum Number of Users in a Group
- Recommended Maximum Number of Domains in a Forest
- Recommended Maximum Number of Domain Controllers in a Domain
- Recommended Maximum Kerberos Settings


**Maximum Number of Objects**

Each domain controller in an Active Directory forest can create a little bit less than 2.15 billion objects during its lifetime.

Each Active Directory domain controller has a unique identifier that is specific to the individual domain controller. These identifiers, which are called Distinguished Name Tags (DNTs), are not replicated or otherwise visible to other domain controllers. The range of values for DNTs is from 0 through 2,147,483,393 ($2^{31}$ minus 255). As objects are created on a domain controller, a unique value is used. A DNT is not reused when an object is deleted. Therefore, domain controllers are limited to creating approximately 2 billion objects (including objects that are created through replication). This limit applies to the aggregate of all objects from all partitions (domain NC, configuration, schema, and any application directory partitions) that are hosted on the domain controller.

Because new domain controllers start with low initial DNT values (typically, anywhere from 100 up to 2,000), it may be possible to work around the domain controller lifetime creation limit—assuming, of course, that the domain is currently maintaining less than 2 billion objects. For example, if the lifetime creation limit is reached because approximately 2 billion objects are created, but 500 million objects are removed from the domain (for example, deleted and then permanently removed from the database through the garbage collection process), installing a new domain controller and allowing it to replicate the remaining objects from the existing domain controllers is a potential workaround. However, it is important that the new domain controller receives the objects through replication and that such domain controllers not be promoted with the Install from Media (IFM) option. Domain controllers that are installed with IFM inherit the DNT values from the domain controller that was used to create the IFM backup.

At the database level, the error that occurs when the DNT limit is reached is "Error: Add: Operations Error. <1> Server error: 000020EF: SvcErr: DSID-0208044C, problem 5012 (DIR_ERROR), data -1076."

| **Note** |
| --- |
| KB article 2548145 includes a hotfix to correct a problem where the size of the Active Directory database can increase rapidly on a domain controller that hosts the DNS server role, with most of the growth contributed by deleted DNS records. To remove those records, the garbage collection task processes a maximum of 5000 records in a single database transaction (also known as MAX_DUMPSTER_SIZE), and will continue to re-schedule itself until it is done. The hotfix does not change the garbage collection process; it only changes behavior in the DNS Server component and how it deals with deletions of dnsNode objects. For more information about garbage collection, see How the Active Directory–Data Store Really Works (Inside NTDS.dit)–Part 4. |

Beginning with Windows Server 2012, you can check a rootDSE operational attribute named **approximateHighestInternalObjectID** to determine the actual highest DNT in use on a domain controller. The following script can be used to retrieve this attribute value:

```
$ossupported = $false

# max possible DNTs
$maxdnts = [System.Math]::Pow(2, 31) - 255

# connect rootDSE
$root = [ADSI]"LDAP://rootDSE"

# get operational attribute dsaVersionString
$root.RefreshCache("dsaVersionString")

# get version string usable in System.Version
$osverstring = $root.dsaVersionString[0].Substring(0,
$root.dsaVersionString[0].IndexOf(" "))

try
{
    $osver = New-Object System.Version $osverstring

    # we need at least W2K12 DC to see the approximateHighestInternalObjectID exposed
    if ($osver.Major -gt 6)
    { $ossupported = $true }

    elseif ($osver.Major -eq 6)
    { $ossupported = ($osver.Minor -ge 2) }
}

catch
{ Write-Host "ERROR: Could not evaluate DC OsVer!"  }

if ($ossupported)
{
    # get operational attribute approximateHighestInternalObjectID
    $root.RefreshCache("approximateHighestInternalObjectID")

    Write-Host "Approx highest committed DNT:
$($root.approximateHighestInternalObjectID[0])"

    Write-Host "DNTs left for new assignments: $($maxdnts -
$root.approximateHighestInternalObjectID[0]) from $maxdnts"
}

else
{ Write-Host "approximateHighestInternalObjectID not exposed (DC OsVer: $osverstring)"
}
```

**Maximum Number of Security Identifiers**

Beginning with Windows Server 2012, the maximum number of security identifiers that can be created over the life of a domain is increased to 2,147,483,647 RIDs. This increase is part of a series of improvements made to how RIDS are issued and monitored. For more information, see Managing RID Issuance.

In Windows Server 2008 R2 and earlier operating systems, there is a limit of approximately 1 billion security identifiers (SIDs) over the life of a domain. This limit is due to the size of the global relative identifier (RID) pool of 30 bits that makes each SID (that is assigned to user, group, and computer accounts) in a domain unique. The actual limit is $2^{30}$ or 1,073,741,823 RIDs. Because RIDs are not reused—even if security principals are deleted—the maximum limit applies, even if there are less than 1 billion security principals in the domain.

| **Note** |
| --- |
| RIDs are assigned in blocks of 500 by default from the domain controller that holds the RID operations master role in each domain. If a domain controller is demoted, the unused RIDs that were allocated to the domain controller are not returned to the global RID pool and are therefore no longer available for use in the domain. |

Beginning in Windows Server 2012, an artificial ceiling is introduced when the number of available RIDs reaches within 10 percent of the limit for the global RID space. When within one percent of the artificial ceiling, domain controllers that request RID pools will log Directory-Services-SAM warning event 16656 to their System event log. For more information, see Managing RID Issuance.

In Windows Server 2008 R2 and earlier operating systems, when all the available RIDs are assigned for a domain, the Directory Service log in the Application and Service Logs of Event Viewer also displays Event ID 16644 from an event log source of the Security Accounts Manager (SAM) that reads "The maximum domain account identifier value has been reached. No further account-identifier pools can be allocated to domain controllers in this domain." If you run Dcdiag when all the available RIDs are assigned for a domain, you see the error message "The DS has corrupt data: rIDAvailablePool value is not valid."

A partial work-around to this limitation is to create an additional domain to hold accounts and then migrate accounts to the new domain. However, you must create a trust relationship to migrate accounts in advance of reaching the limit. Creating a trust requires the creation of a security principal, which is also known as a trust user account. For more information about this limit, see articles 316201 (http://go.microsoft.com/fwlink/?LinkID=115211) and 305475 (http://go.microsoft.com/fwlink/?LinkId=115212) in the Microsoft Knowledge Base.

| **Note** |
| --- |
| The Active Directory database does not set limits on the number of objects in a container, such as organizational units (OUs). You might experience limits when you work with multiple thousands of objects. These limits are configured to help provide a certain level of application or service availability. For example, the Active Directory Users and Computers snap-in is configured by default to display a maximum of 2,000 objects per container. You can adjust this value by using the **Filter Options** settings on the **View** menu. There are also adjustable Lightweight Directory Access Protocol (LDAP) policies that are set by default to improve domain controller performance. These policies are described in article 315071 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=135481). |

**Maximum Number of entries in Discretionary and Security Access Control Lists**

The limitation for the number of entries in a discretionary access control list (DACL) or a security access control list (SACL) of an Active Directory object using the ntSecurityDescriptor attribute comes from a limitation in the size of the access control list (ACL), which is 64K. Since access control entries (ACEs) vary in size, the actual number of entries (SIDs) is approximately 1,820. For additional details, see How Security Descriptors and Access Control Lists Work (http://go.microsoft.com/fwlink/?LinkId=214683).

**Group Memberships for Security Principals**

Security principals (that is, user, group, and computer accounts) can be members of a maximum of approximately 1,015 groups. This limitation is due to the size limit for the access token that is created for each security principal. The limitation is not affected by how the groups may or may not be nested. For more information, see article 328889 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=115213). For a detailed discussion of access token limitations, see Addressing Problems Due to Access Token Limitation (http://go.microsoft.com/fwlink/?LinkId=146571).

For more information about how a domain controller creates the data structure that is used for authorization decisions, see:

- [MS-PAC]: Privilege Attribute Certificate Data Structure (http://msdn.microsoft.com/en-us/library/cc237917(PROT.13).aspx)
- 3.3.5.3.2 Initial Population of the PAC (http://msdn.microsoft.com/en-us/library/cc233956(PROT.13).aspx)
- 3.3.5.4.3 Domain Local Group Membership (http://msdn.microsoft.com/en-us/library/cc233950(PROT.13).aspx)

**FQDN Length Limitations**
Fully qualified domain names (FQDNs) in Active Directory cannot exceed 64 characters in total length, including hyphens and periods (.). For example, the following host name has 65 characters; therefore, it is not valid in an Active Directory domain:

server10.branch-15.southaz.westernregion.northamerica.contoso.com

This is an important limitation to keep in mind when you name domains. This limitation is due to the MAX_PATH length of 260 characters that the Win32 application programming interfaces (APIs) define, in combination with the way in which Group Policy objects (GPOs) are stored in the SYSVOL share. For more information, see article 245809 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=115219). For more information about naming limitations, see article 909264 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=106629).

**File Name and Path Length Limitations**
The physical files that Active Directory components use, such as SYSVOL, database (NTDS.DIT), and log file paths, are constrained by the MAX_PATH length of 260 characters, as defined by the Win32 APIs. When you are determining where to place your SYSVOL and database files during Active Directory installation, avoid nested folder structures that make the full file path to the SYSVOL folder, database, and log files longer than 260 characters. For more information, see article 245809 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=115219).

**Additional Name Length Limitations**
There are additional limitations regarding name lengths in Active Directory. The following limits are described in article 909264 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=106629):

- NetBIOS computer and domain names are limited to 15 characters.
- Domain Name System (DNS) host names are limited to 24 characters.
- OU names are limited to 64 characters.

**Name Length Limits from the Schema**
Default limits on attribute names for Active Directory objects that are imposed by the schema include the following. These items provide examples of schema-limited name attributes:

- Display names are limited to 256 characters. For more information, see Display-Name Attribute (http://go.microsoft.com/fwlink/?LinkId=153705).
- Common names are limited to 64 characters. For more information, see Common-Name Attribute (http://go.microsoft.com/fwlink/?LinkId=153706).
- The **SAM-Account-Name** attribute (also known as the pre–Windows 2000 user logon name) is limited to 256 characters in the schema. However, for the purpose of backward compatibility the limit is 20 characters. For more information, see SAM-Account-Name Attribute (http://go.microsoft.com/fwlink/?LinkId=153707).

**Name Length Limitations for LDAP Simple Bind Operations**
During binds to the directory, simple LDAP bind operations limit the distinguished name (also known as DN) of the user to 255 total characters. If you attempt a simple LDAP bind with more than 255 characters, you might experience authentication errors, such as the following:

```
Error <49>: ldap_simple_bind_s() failed: Invalid Credentials
Server error: 80090308: LdapErr: DSID-0C0903AA, comment: AcceptSecurityContext error,
data 57, v1771
Error 0x80090308 The token supplied to the function is invalid
```

You can avoid this issue by ensuring that the applications, scripts, and utilities that attempt to bind to your directory use secure LDAP binds. You can also avoid this issue by reducing the depth of the OU structure or the length of the OU names. For example, the following distinguished name is 261 characters:

```
CN=BobKelly,OU=CorporateVicePresidents,OU=CorporateOfficers,OU=ViewOfPugetSoundOffices,
OU=TopFloor,OU=Building1557,OU=CorporateCampus,OU=Redmond,OU=Washington,OU=NorthWestern
,OU=UnitedStatesOfAmerica,OU=NorthAmerica,DC=BusinessGroup,DC=humongousinsurance,DC=com
```

If the OU named CorporateVicePresidents is shortened to CVP, the distinguished name for the user account BobKelly is only 242 characters.

**Maximum Number of GPOs Applied**
There is a limit of 999 Group Policy objects (GPOs) that you can apply to a user account or computer account. This does not mean that the total number of policy settings on the system is limited to 999. Rather, a single user or computer will not be able to process more than 999 GPOs. This limit exists for performance reasons.

**Trust Limitations**
Trust limitations arise from the number of trusted domain objects (TDOs), the length of trust paths, and the ability of clients to discover available trusts. Limitations that apply include the following:

- Kerberos clients can traverse a maximum of 10 trust links to locate a requested resource in another domain. If the trust path between the domains exceeds this limit, the attempt to access the domain fails.
- When a client searches out a trust path, the search is limited to the trusts that are established directly with a domain and the trusts that are transitive within a forest.
- Previous testing shows that the increased time to complete TDO-related operations, such as authentication across domains, deteriorates performance noticeably if the Active Directory implementation in an organization contains more than 2,400 TDOs.

For more information about trust limitations, see "Practical Limitations of Trusts" in How Domain and Forest Trusts Work (http://go.microsoft.com/fwlink/?LinkID=35356).

**Maximum Number of Accounts per LDAP Transaction**
When you write scripts or applications that perform LDAP transactions, the recommended limit is to perform no more than 5,000 operations per LDAP transaction. An LDAP transaction is a group of directory operations (such as add, delete, and modify) that are treated as one unit. If your script or application performs more than 5,000 operations in a single LDAP transaction, you are at risk of running into resource limits and an operational time-out. If that happens, all the operations (changes, additions, and modifications) in the transaction are rolled back, which means that you lose all those changes.

As an example, if you are using Active Directory Service Interfaces (ADSI) to write a script, the SetInfo method completes a transaction. For more information about ADSI Methods, see Active Directory Service Interfaces (http://go.microsoft.com/fwlink/?LinkID=4487).

As another example, when you use the System.DirectoryServices (S.DS) namespace in the Microsoft .NET Framework, the DirectoryEntry.CommitChanges method completes an LDAP transaction. For more information about the DirectoryEntry.CommitChanges method, see DirectoryEntry.CommitChanges () (http://go.microsoft.com/fwlink/?LinkId=115220).

| **Note** |
| --- |
| Regardless of the method that you use for LDAP transactions, you should plan to send less than 5,000 directory operations in a single transaction. To learn more about the LDAP data structure that commits changes, see LDAPMod (http://go.microsoft.com/fwlink/?LinkId=115221). |

**Recommended Maximum Number of Users in a Group**
For Windows 2000 Active Directory environments, the recommended maximum number of members in a group is 5,000. This recommendation is based on the number of concurrent atomic changes that can be committed in a single database transaction.

Starting with Windows Server 2003, the ability to replicate discrete changes to linked multivalued properties was introduced as a technology called Linked Value Replication (LVR). To enable LVR, you must increase the forest functional level to at least Windows Server 2003 interim. Increasing the forest functional level changes the way that group membership (and other linked multivalued attributes) is stored in the database and replicated between domain controllers. This allows the number of group memberships to exceed the former recommended limit of 5,000 for Windows 2000 or Windows Server 2003 at a forest functional level of Windows 2000.

So far, testing in this area has yet to reveal any new recommended limits to the number of members in a group or any other linked multivalued attribute. Production environments have been reported to exceed 4 million members, and Microsoft scalability testing reached 500 million members.

**Important**

Increasing the forest functional level to Windows Server 2003 interim or higher does not modify the way that existing group members are stored or replicated. To do that, you must remove the members that were added to the group before the forest functional level was increased to Windows Server 2003 and then add them back again to the appropriate groups. Any group members that you either add or remove after the forest functional level is increased will be LVR enabled, even if the group contains other members that are not LVR enabled.

For more information about linked attributes, see Linked Attributes (http://go.microsoft.com/fwlink/?LinkId=142909). For more information about the replication process, see How the Active Directory Replication Model Works (http://go.microsoft.com/fwlink/?LinkId=142908).

**Recommended Maximum Number of Domains in a Forest**

For Windows 2000 Server, the recommended maximum number of domains in a forest is 800. For Windows Server 2003, the recommended maximum number of domains when the forest functional level is set to Windows Server 2003 (also known as forest functional level 2) is 1,200. This restriction is a limitation of multivalued, nonlinked attributes in Windows Server 2003. For more information, see "Maximum Database Record Size" in How the Data Store Works (http://go.microsoft.com/fwlink/?LinkId=134791).

**Recommended Maximum Number of Domain Controllers in a Domain**

To ensure reliable recovery of SYSVOL, we recommend a limit of 1200 domain controllers per domain.

If any Active Directory domain in your network is expected to exceed 800 domain controllers and those domain controllers are hosting Active Directory–integrated Domain Name System (DNS) zones, review article 267855 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=115222).

For more information about FRS limitations, see the FRS Technical Reference (http://go.microsoft.com/fwlink/?LinkId=115302).

**Recommended Maximum Kerberos Settings**

The maximum recommended size for a Kerberos ticket is 48,000 bytes, which is configured through the **MaxTokenSize** REG_DWORD value in the registry (**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Lsa\Kerberos\Parameters**) or through Group Policy, as described in KB article 938118.

**Note**

The maximum allowed value of MaxTokenSize is 65,535 bytes. If you are using Kerberos for IPSEC key management, the limit of 65,536 bytes. However, because of HTTP's base64 encoding of authentication context tokens, we do not recommend that you set the maxTokenSize registry entry to a value larger than 48,000 bytes. Starting with Windows Server 2012, the default value of the MaxTokenSize registry entry is 48,000 bytes.

For additional information about Kerberos tickets, including error conditions that can occur when Kerberos ticket size limits are set too low or too high, see Additional Resources for Troubleshooting Kerberos (http://go.microsoft.com/fwlink/?LinkId=134740).

# Designing the Active Directory Logical Structure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing your Microsoft® Windows® Server 2003 Active Directory® directory service logical structure involves defining the structure of and relationships between the forests, domains, and organizational units that you plan to deploy. Carefully designing your Active Directory logical structure enables you to minimize the time and effort required to implement Active Directory in your organization. In this way, you can establish an efficient structure that best meets the administrative needs of your organization.

**In This Topic**


Overview of Designing the Active Directory Logical Structure

Identifying the Deployment Project Participants

Creating a Forest Design

Creating a Domain Design

Designing a DNS Infrastructure to Support Active Directory

Designing Organizational Units for Delegation of Administration

Additional Resources for Designing the Active Directory Logical Structure


**Related Information**

- For more information about Active Directory concepts, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

# Overview of Designing the Active Directory Logical Structure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory in Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition enables organizations to create a scalable, secure, and manageable infrastructure for user and resource management and to support directory-enabled applications.

A well-designed Active Directory logical structure provides the following benefits:

- Simplified management of Windows networks that contain large numbers of objects.
- A consolidated domain structure and reduced administration costs.
- The ability to delegate administrative control over resources as appropriate.
- Reduced impact on network bandwidth.
- Simplified resource sharing.
- Optimal search performance.
- Low total cost of ownership.

A well-designed Active Directory logical structure facilitates the efficient integration of features such as Group Policy, enabling desktop lockdown, software distribution, and user, group, workstation, and server administration, into your system. In addition, a carefully designed logical structure facilitates the integration of services such as Microsoft® Exchange 2000, public key infrastructure (PKI), and domain-based distributed file system (DFS).

If your organization is currently operating in a non-Microsoft operating system or Microsoft® Windows NT® version 4.0 operating system environment, designing an Active Directory logical structure before deploying Active Directory enables you to optimize your deployment process in order to best take advantage of Windows Server 2003 Active Directory features.

**Note**

- For a list of the job aids that are available to assist you in designing your Active Directory logical structure, see "Additional Resources for Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" later in this chapter.

# Process for Designing the Active Directory Logical Structure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To design the Active Directory logical structure, a design team must first identify the requirements of the organization, and, based on this information, decide where to place the forest and domain boundaries. Then, the design team decides how to configure the DNS environment to meet the needs of the forest. Finally, the team identifies the organizational unit (OU) structure that is required in order to delegate management of resources in your organization.

Figure 2.1 shows the process for designing the logical structure.

**Figure 2.1 Designing the Active Directory Logical Structure**



# Active Directory Logical Structure Background Information

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you design your Active Directory logical structure, it is important to understand the Active Directory logical model. Active Directory is a distributed database that stores and manages information about network resources, as well as application-specific data from directory enabled applications. Active Directory allows administrators to organize elements of a network (such as users, computers, devices, and so on) into a hierarchical containment structure. The top-level container is the forest. Within forests are domains, and within domains are organizational units. This is called the logical model because it is independent of the physical aspects of the deployment, such as the number of domain controllers required within each domain and network topology.

Figure 2.2 shows the relationship between forests, domains, and organizational units.

**Figure 2.2 Relationship Between Active Directory Forests, Domains, and OUs**




**Active Directory forest**

A forest is a collection of one or more Active Directory domains that share a common logical structure, global catalog, directory schema, and directory configuration, as well as automatic two-way transitive trust relationships. Each forest is a single instance of the directory and defines a security boundary.

**Security boundary** A boundary that defines a container for which no administrator external to the container can take control away from administrators within the container. For example, a forest is a security boundary. No administrators from outside the forest can control access to information inside the forest unless first given permission to do so by the administrators within the forest. By contrast a domain is not a security boundary because within a forest it is not possible for administrators from one domain to prevent a malicious administrator from another domain from accessing data in their domain.

**Active Directory domain**

A domain is a partition in an Active Directory forest. Partitioning data enables organizations to replicate data only to where it is needed. In this way, the directory can scale globally over a network that has limited available bandwidth. In addition, the domain supports a number of other core functions related to administration, including:

- Network-wide user identity. Domains allow user identities to be created once and referenced on any computer joined to the forest in which the domain is located. Domain controllers that make up a domain are used to store user accounts and user credentials such as passwords or certificates securely.

- Authentication. Domain controllers provide authentication services for users and supply additional authorization data, such as user group memberships, which can be used to control access to resources on the network.

- Trust relationships. Domains can extend authentication services to users in domains outside their own forest by means of trusts.

- Replication. The domain defines a partition of the directory containing sufficient data to provide domain services and replicates it between the domain controllers. In this way, all domain controllers are peers in a domain and are managed as a unit.

**Active Directory organizational units**
OUs can be used to form a hierarchy of containers within a domain. Organizational units are used to group objects for administrative purposes, such as the application of Group Policies or delegation of authority. Control over an OU and the objects within it is determined by the access control lists (ACLs) on the OU and on the objects in the OU.

**Delegation of authority** To facilitate the management of large numbers of objects, Active Directory supports the concept of delegation of authority. By means of delegation, owners can transfer full or limited authority over objects to other users or groups. Delegation is important because it helps to distribute the management of large numbers of objects across a number of people trusted to perform management tasks.

**Active Directory Logical Structure Design Process**
Designing your Active Directory logical structure involves defining the relationships between the containers in your directory. These relationships might be based on administrative requirements, such as delegation of authority, or they might be defined by operational requirements, such as the need to control replication. In general, forests are used as security boundaries, domains are used to control replication, and OUs are used to delegate administration.

Figure 2.3 shows the distribution of users in a fictitious multinational organization. In this example, a single Active Directory forest contains all of the users, computers, devices, and other entities within the organization. The organization started with a domain for their corporate headquarters (Corp) and then, to control replication across the entire world-wide network, the organization created three additional domains (North America, South America, and Europe) as partitions of the forest. To support further delegation, the organization subdivided the North America domain into three OUs: West, Central, and East.

**Figure 2.3 Delegation of Administration Within an Organization**



# Identifying the Deployment Project Participants

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in establishing an Active Directory deployment project is to establish the design and deployment project teams who will be responsible for managing the design phase and deployment phase of the Active Directory project cycle. In addition, you must identify the individuals and groups who will be responsible for owning and maintaining the directory after the deployment is completed. Figure 2.4 shows the process for identifying the Active Directory deployment project participants.

**Figure 2.4 Identifying Active Directory Project Participants**




The design and deployment teams that you establish consist of individuals who understand the various needs of the organization. This includes members of IT groups in your organization who are responsible for network engineering and network planning with respect to directory services. It is important that the project teams include forward-looking IT managers who are responsible for the future growth of your networking environment, as well as individuals who are familiar with the Active Directory deployment process.

It is also important to include in the project teams individuals, called *owners*, who will be responsible for maintaining the directory when the deployment process is complete. After Active Directory is deployed, the owners ensure that service levels are maintained and are responsible for the day-to-day maintenance of the directory. Owners are usually IT managers who oversee the groups or teams of network administrators that perform the hands-on work.

# Defining Project-Specific Roles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An important step in establishing the project teams is to identify the individuals who are to hold project-specific roles. These include the executive sponsor, the project architect, and the project manager. These individuals are responsible for running the Active Directory deployment project.

After you appoint the project architect and project manager, these individuals establish channels of communication throughout the organization, build project schedules, and identify the individuals who will be members of the project teams, beginning with the various owners.

**Executive sponsor**
Deploying an infrastructure such as Active Directory can have a wide-ranging impact on an organization. For this reason, it is important to have an executive sponsor who understands the business value of the deployment, supports the project at the executive level, and can help resolve conflicts across the organization.

**Project architect**
Each Active Directory deployment project requires a project architect to manage the Active Directory design and deployment decision-making process. The architect provides technical expertise to assist with the process of designing and deploying Active Directory.

### Note

- If no existing personnel in your organization have directory design experience, you might want to hire an outside consultant who is an expert in Active Directory design and deployment.

The responsibilities of the Active Directory project architect include:

- Owning the Active Directory design.
- Understanding and recording the rationale for key design decisions.
- Ensuring that the design meets the business needs of the organization.
- Establishing consensus between design, deployment, and operations teams.
- Understanding the needs of Active Directory–integrated applications, such as Exchange 2000, and the impact they have on the Active Directory design.

The final Active Directory design must reflect a combination of business goals and technical decisions. Therefore, the project architect must review design decisions to ensure that they are in alignment with business goals.

**Project manager**
The project manager facilitates cooperation across business units and between technology management groups. Ideally, the Active Directory deployment project manager is someone from within the organization who is familiar with the operational policies of the IT group and the design requirements for the groups that are preparing to deploy Active Directory. The project manager oversees the entire deployment project, beginning with design and continuing through implementation, and makes sure that the project stays on schedule and within budget. The responsibilities of the project manager include:

- Providing basic project planning, such as scheduling and budgeting.
- Driving progress on the Active Directory design and deployment project.
- Ensuring that the appropriate individuals are involved in each part of the design process.
- Serving as single point of contact for the Active Directory deployment project.

- Establishing communication between design, deployment, and operations teams.
- Establishing and maintaining communication with the executive sponsor throughout the deployment project.

# Establishing Owners and Administrators

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In an Active Directory deployment project, individuals who are *owners* are held accountable by management for making sure that deployment tasks are completed and that Active Directory design specifications meet the needs of the organization. Owners do not necessarily have access to or manipulate the directory infrastructure directly. *Administrators* are the individuals responsible for completing the required deployment tasks. Administrators have the network access and permissions necessary to manipulate the directory and its infrastructure.

The role of the owner is strategic and managerial. Owners are responsible for communicating to administrators the tasks required for the implementation of the Active Directory design, such as the creation of new domain controllers within the forest. The administrators are responsible for implementing the design on the network according to the design specifications.

In larger organizations, different individuals fill owner and administrator roles; however, in some smaller organizations, the same individual might act as both the owner and the administrator.

**Service and Data Owners**

Managing Active Directory on a daily basis involves two types of owners:

- Service owners. Service owners are responsible for the planning and long-term maintenance of the Active Directory infrastructure, and ensuring that the directory continues to function, and that the goals established in service level agreements are maintained.

- Data owners. Data owners are responsible for the maintenance of the information stored in the directory. This includes user and computer account management and management of local resources, such as member servers and workstations.

It is important to identify the Active Directory service and data owners early on so that they can participate in as much of the design process as possible. Because the service and data owners are responsible for the long-term maintenance of the directory after the deployment project is finished, it is important for these individuals to provide input as to organizational needs and to be familiar with how and why certain design decisions were made. Service owners include the forest owner, the Active Directory DNS owner, and the site topology owner. Data owners include OU owners.

**Service and Data Administrators**

The operation of Active Directory involves two types of administrators: service administrators and data administrators. Service administrators implement policy decisions made by service owners and handle the day-to-day tasks associated with maintaining the directory service and infrastructure. This includes managing the domain controllers that are hosting the directory service; managing other network services, such as DNS, that are required for Active Directory; controlling the configuration of forest-wide settings; and ensuring that the directory is always available.

Service administrators are also responsible for completing ongoing Active Directory deployment tasks that are required after the initial Windows Server 2003 Active Directory deployment process is complete. For example, as demands on the directory increase, service administrators create additional domain controllers, and establish or remove trusts between domains as needed. For this reason, the Active Directory deployment team needs to include service administrators.

You must be careful to assign service administrator roles only to trusted individuals in the organization. Because these individuals have the ability to modify the system files on domain controllers, they can change the behavior of Active Directory, enabling them to gain access to and modify any resource in the forest. You must ensure that the service administrators in your organization are individuals who are familiar with the operational and security policies that are in place on your network and who understand the need to enforce those policies.

Data administrators are users within a domain who are responsible for maintaining data that is stored in Active Directory, such as user and group accounts, and maintaining computers that are members of their domain. Data administrators control subsets of objects within the directory and have no control over the installation or configuration of the directory service.

Data administrator accounts are not provided by default. After the design team determines how resources are to be managed for the organization, domain owners must create data administrator accounts and delegate them the appropriate permissions based on the set of objects for which the administrators are to be responsible.

It is best to limit the number of service administrators in your organization to the minimum number required to ensure that

the infrastructure continues to function. The majority of administrative work can be completed by data administrators. Service administrators require a much wider skill set because they are responsible for maintaining the directory and the infrastructure that supports it. Data administrators only require the skills necessary to manage their portion of the directory. Dividing work assignments in this way results in cost savings for the organization because only a small number of administrators need to be trained to operate and maintain the entire directory and its infrastructure.

For example, a service administrator needs to understand how to add a domain to a forest. This includes how to install the software to convert a server into a domain controller and how to manipulate the DNS environment so that the domain controller can be merged seamlessly into the Active Directory environment. A data administrator only needs to know how to manage the specific data that they are responsible for, such as the creation new user accounts for any new employees in their department.

**Service and Data Owners for Active Directory**
Deploying Active Directory requires coordination and communication between many different groups involved in the operation of the network infrastructure. These groups should appoint service and data owners who are responsible for representing the various groups during the design and deployment process.

Once the deployment project is complete, these service and data owners continue to be responsible for the portion of the infrastructure managed by their group. In an Active Directory environment these owners are:

**Forest Owner**
The forest owner is typically a senior IT manager in the organization, who is responsible for the Active Directory deployment process and who is ultimately accountable for maintaining service delivery within the forest after the deployment is complete. The forest owner assigns individuals to fill the other ownership roles by identifying key personnel within the organization who are able to contribute necessary information about network infrastructure and administrative needs. The forest owner is responsible for:

- The deployment of the forest root domain to create the forest.
- The deployment of the first domain controller in each domain to create the domains required for the forest.
- The memberships of the service administrator groups in all domains of the forest.
- The creation of the design of the OU structure for each domain in the forest.
- The delegation of administrative authority to OU owners.
- Changes to the schema.
- Changes to forest-wide configuration settings.
- The implementation of certain Group Policies, including
    - Domain user account policies, such as password complexity and account lockout.
    - Policies that apply to domain controllers.
    - Any other Group Policies that are applied at the domain level.

The forest owner has authority over the entire forest. It is the forest owner's responsibility to set policy and to select the individuals who are service administrators. The forest owner is a service owner.

**Active Directory DNS Owner**
The Active Directory DNS owner is an individual who has a thorough understanding of the existing DNS infrastructure and the existing namespace of the organization.

The Active Directory DNS owner is responsible for:

- Serving as a liaison between the design team and the IT group that currently owns the DNS infrastructure.
- Providing the information about the existing DNS namespace of the organization, to assist in the creation of the new

Active Directory namespace.

- Working with the deployment team to make sure that the new DNS infrastructure is deployed according to the specifications of the design team and that it is working properly.

- Managing the Active Directory DNS infrastructure including the DNS services and DNS data.

The Active Directory DNS owner is a service owner.

**Site Topology Owner**
The site topology owner is familiar with the physical structure of the network of the organization, including the mapping of individual subnets, routers, and the areas of the network that are connected by means of slow links. The site topology owner is responsible for the following tasks:

- Understanding the physical network topology and how it affects Active Directory.

- Understanding how the Active Directory deployment will impact the network.

- Determining the sites that need to be created.

- Updating server objects for domain controllers when a subnet changes.

- Creating site links.

The site topology owner is a service owner.

**OU Owner**
The OU owner is responsible for managing data stored in the directory. This individual needs to be familiar with the operational and security policies that are in place on the network. OU owners can only perform tasks that have been delegated to them by the service administrators and they can only perform those tasks on the OUs to which they are assigned. Some tasks that the OU owner might be assigned include the following:

- Performing all account management tasks within their assigned OU.

- Managing workstations and member servers that are members of their assigned OU.

- Delegating authority to local administrators within their assigned OU.

The OU owner is a data owner.

# Building Project Teams

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory project teams are temporary groups that are responsible for completing Active Directory design and deployment tasks. When the Active Directory deployment project is complete, the owners assume responsibility for the directory and the project teams can disband.

The size of the project teams varies according to the size of the organization. In smaller organizations, a single person can cover multiple areas of responsibility on a project team and be involved in more than one phase of the deployment. Larger organizations might require larger teams with different individuals, or even different teams, covering the different areas of responsibility. The size of the teams is not important as long as all areas of responsibility are assigned and the design goals of the organization are met.

# Identifying Potential Forest Owners

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Identify the groups within your organization that own and control the resources necessary to provide directory services to users on the network. These groups are considered potential forest owners.

The separation of service and data administration in Active Directory makes it possible for the infrastructure IT group or groups of an organization to manage the directory service, while local administrators in each group manage the data that belongs to their own groups. Potential forest owners have the required authority over the network infrastructure to deploy and support Active Directory.

For organizations that have one centralized infrastructure IT group, the IT group is generally the forest owner and therefore the potential forest owner for any future deployments. Organizations that include a number of independent infrastructure IT groups have a number of potential forest owners. If your organization already has an Active Directory infrastructure in place, then any current forest owners are also potential forest owners for new deployments.

Select one of the potential forest owners to act as the forest owner for each forest that you are considering for deployment. These potential forest owners are responsible for working with the design team to determine whether or not their forest will actually be deployed or if an alternate course of action, such as joining another existing forest, is a better use of the available resources and still meets their needs. The forest owner or owners in your organization are members of the Active Directory design team.

# Establishing a Design Team

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory design team is responsible for gathering all of the information needed to make decisions about the Active Directory logical structure design. The responsibilities of the design team include:

- Determining how many forests and domains are required and what the relationships are between the forests and domains.
- Working with data owners to ensure that the design meets their security and administrative requirements.
- Working with the current network administrators to ensure that the current network infrastructure supports the design and the design will not adversely affect existing applications deployed on the network.
- Working with representatives of the security group of the organization to ensure that the design meets established security policies.
- Designing OU structures that permit appropriate levels of protection and the proper delegation of authority to the data owners.
- Working with the deployment team to test the design in a lab environment to ensure that it functions as planned, and modifying the design as needed to address any problems that occur.
- Creating a site topology design that meets the replication requirements of the forest while preventing overload of available bandwidth. For more in formation about designing the site topology, see "Designing the Site Topology" in this book
- Working with the deployment team to ensure that the design is implemented correctly.

The design team includes the following members:

- Potential forest owners
- Project architect
- Project manager
- Individuals who are responsible for establishing and maintaining security policies on the network

During the logical structure design process, the design team identifies the other owners. These individuals must start participating in the design process as soon as they are identified.

After the deployment project is handed off to the deployment team, the design team is responsible for overseeing the deployment process to ensure that the design is implemented correctly. The design team also makes changes to the design based on feedback from testing.

# Establishing a Deployment Team

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory deployment team is responsible for testing and implementing the Active Directory logical structure design. This involves the following tasks:

- Establishing a test environment that sufficiently emulates the network environment.
- Testing the design by implementing the proposed forest and domain structure in a lab environment to verify that it meets the goals of each role owner.
- Developing and testing any migration scenarios proposed by the design in a lab environment.
- Making sure that each owner signs off on the testing process to ensure that the correct design features are being tested.
- Testing the deployment operation in a pilot environment.

When the design and testing tasks are complete, the deployment team:

- Creates the forests and domains according to the Active Directory logical structure design.
- Creates the sites and site link objects as needed based on the site topology design.
- Ensures that the DNS infrastructure is configured to support Active Directory and that any new namespaces are integrated into the existing namespace of the organization.

The Active Directory deployment team includes the following members:

- Forest owner
- Active Directory DNS owner
- Site topology owner
- OU owners

The deployment team works with the service and data administrators during the deployment phase to ensure that members of the operations team are familiar with the new design. This helps to ensure a smooth transition of ownership when the deployment operation is completed. At the completion of the deployment process, the responsibility for maintaining the new Active Directory environment passes to the operations team.

# Document the Design and Deployment Teams

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Document the names of and contact information for the people who will participate in the design and deployment of Active Directory. Identify who will be responsible for each role on the design and deployment teams. Initially this list includes the potential forest owners, the project manager, and the project architect. When you determine the number of forests that you will deploy, you might need to create new design teams for additional forests. Note that you will need to update your documentation as team memberships change and as you identify the various Active Directory owners during the design process.

For a worksheet to assist you in documenting your Active Directory design and deployment team members, see "Design and Deployment Team Information" (DSSLOGI_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Design and Deployment Team Information" on the Web at http://www.microsoft.com/reskit).

# Example: Identifying Deployment Project Participants

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso Pharmaceuticals Corporation decided to deploy Windows Server 2003 Active Directory in their environment. They established a design team to begin the forest planning and deployment process. This team included IT personnel who were familiar with the existing network, all of the potential forest owners, the individuals who manage the corporate namespace, and the owners and administrators who will be responsible for the directory after the deployment project is complete. Some members of the Active Directory logical structure design team will be identified as the design process progresses. Figure 2.5 shows the worksheet that Contoso corporation completed to document the design team members who will begin the deployment process.

**Figure 2.5 Example of a Design and Deployment Team Information Worksheet**

| Design and Deployment Team Information | | | |
|---|---|---|---|
| Prepared By | Maria Hammond | Date | 2/18/2003 |
| Name | Role | E-mail | Phone |
| Amy Rusko | Executive Sponsor | amy@contoso.com | x832 |
| Chris Gray | Project Architect | chris@contoso.com | x659 |
| Maria Hammond | Project Manager | maria@contoso.com | x245 |
| | | | |
| | Forest Owner | | |
| | Active Directory DNS Owner | | |
| | Site Topology Owner | | |
| Instructions: Record the names and contact information for the key members of the design and deployment teams above. Use the space below to record the contact information for any additional personnel whom you want to participate in the project. Add more pages as necessary. | | | |
| Name | Role | E-mail | Phone |
| Mark Lee | Potential Forest Owner - Central IT Group | mark@contoso.com | x874 |
| Mary Baker | Potential Forest Owner - IT Pharmaceuticals Division | mary@contoso.com | x989 |
| Reed Koch | Potential Forest Owner - IT Nutritional Products Division | reed@contoso.com | x114 |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

# Creating a Forest Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Creating a forest design involves first identifying the groups within your organization that have the resources available to host an Active Directory forest, and then defining your forest design requirements. Finally, you need to determine the number of forests that you require in order to meet the needs of your organization.

Figure 2.6 shows the process for creating a forest design.

**Figure 2.6 Creating a Forest Design**



# Identifying Forest Design Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order to create a forest design for your organization, you must identify the business requirements that your directory structure needs to accommodate. This involves determining how much autonomy the groups in your organization need to manage their network resources, and whether each group needs to isolate their resources on the network from other groups.

Active Directory enables you to design a directory infrastructure that accommodates multiple groups within an organization that have unique management requirements, and to achieve structural and operational independence between groups as needed.

Groups in your organization might have some of the following types of requirements:

- **Organizational structure requirements.** Parts of an organization might participate in a shared infrastructure to save costs, but require the ability to operate independently from the rest of the organization. For example, a research group within a large organization might need to maintain control over all of their own research data.

- **Operational requirements.** One part of an organization might place unique constraints on the directory service configuration, availability, or security, or use applications that place unique constraints on the directory. Examples are found in:

    - Military organizations.

    - Hosting scenarios.

    - Organizations that maintain a directory that is available both internally and externally (publicly accessible by users on the Internet).

- **Legal requirements.** Some organizations have legal requirements to operate in a specific way, such as by restricting access to certain information as specified in a business contract. Failure to meet these requirements can result in loss of the contract and possibly legal action. These requirements commonly apply to the following types of organizations:

    - Financial institutions that need to maintain private client financial records.

    - Defense contractors that are working on classified military projects.

    - Government organizations that maintain top secret data.

Part of identifying your forest design requirements involves identifying the degree to which groups in your organization can trust the potential forest owners and their service administrators, and identifying the autonomy and isolation requirements for each group in your organization.

# Service Administrator Scope of Authority

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you choose to participate in an Active Directory forest, you must trust the forest owner and the service administrators. The forest owners are responsible for selecting and managing the service administrators; therefore, when you trust a forest owner, you also trust the service administrators that the forest owner manages. These service administrators have access to all of the resources in the forest. Before making the decision to participate in a forest, it is important to understand that the forest owner and the service administrators will have full access to your data. You cannot prevent this access.

All service administrators in a forest have full control over all data and services on all computers in the forest. Service administrators have the ability to do the following:

- Correct errors on ACLs of objects. This enables the service administrator to read, modify, or delete objects regardless of the ACLs that are set on those objects.
- Modify the system software on a domain controller to bypass normal security checks. This enables the service administrator to view or manipulate any object in the domain, regardless of the ACL on the object.
- Use the Restricted Groups security policy to grant any user or group administrative access to any computer joined to the domain. Service administrators can thus obtain control of any computer joined to the domain, regardless of the intentions of the computer owner.
- Reset passwords or change group memberships for users.
- Gain access to other domains in the forest by modifying the system software on a domain controller. Service administrators can affect the operation of any domain in the forest, view or manipulate forest configuration data, view or manipulate data stored in any domain, or view or manipulate data stored on any computer joined to the forest.

For this reason, groups that store data in OUs in the forest and that join computers to a forest must trust the service administrators. In order for a group to join a forest, it must choose to trust all service administrators in the forest. This involves ensuring that:

- The forest owner can be trusted to act in the interests of the group, and does not have reason to act maliciously against the group.
- The forest owner appropriately restricts physical access to domain controllers. Domain controllers within a forest cannot be isolated from one another. It is possible for an attacker who has physical access to a single domain controller to make offline changes to the directory database, and by so doing interfere with the operation of any domain in the forest, view or manipulate data stored anywhere in the forest, or view or manipulate data stored on any computer joined to the forest. For this reason, physical access to domain controllers must be restricted to trusted personnel.
- You understand and accept the potential risk that trusted service administrators can be coerced into compromising the security of the system.

Some groups might determine that the collaborative and cost-saving benefits of participating in a shared infrastructure outweigh the risks that service administrators will misuse or will be coerced into misusing their authority. These groups can share a forest and use OUs to delegate authority. However, other groups might not accept this risk because the consequences of a compromise in security are too severe. Those groups require separate forests.

# Autonomy vs. Isolation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can design your Active Directory logical structure to achieve either of the following:

- **Autonomy**.Autonomy involves independent but not exclusive control of a resource. When you achieve autonomy, administrators have the authority to manage resources independently; however, administrators with greater authority exist who also have control over those resources and can take control away if necessary. You can design your Active Directory logical structure to achieve the following types of autonomy:
    - **Service autonomy.** This type of autonomy involves control over all or part of service management.
    - **Data autonomy.** This type of autonomy involves control over all or part of the data stored in the directory or on member computers joined to the directory.
- **Isolation.** Isolation involves independent and exclusive control of a resource. When you achieve isolation, administrators have the authority to manage a resource independently and no other administrators can take control of the resource away. You can design your Active Directory logical structure to achieve the following types of isolation:
    - **Service isolation.** This type of isolation prevents administrators other than those specifically designated to control service management from controlling or interfering with service management.
    - **Data isolation.** This type of isolation prevents administrators other than those specifically designated to control or view data from controlling or viewing a subset of data in the directory or on member computers joined to the directory.

Administrators who require only autonomy accept that other administrators who have equal or greater administrative authority have equal or greater control over service or data management. Administrators who require isolation have exclusive control over service or data management. Creating a design to achieve autonomy is generally less expensive than creating a design to achieve isolation.

In Active Directory, administrators can delegate both service administration and data administration in order to achieve either autonomy or isolation between organizations. The combination of service management, data management, autonomy, and isolation requirements of an organization impact the Active Directory containers that are used to delegate administration.

# Isolation and Autonomy Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The number of forests that you need to deploy is based on the autonomy and isolation requirements of each group within your organization. In order to identify your forest design requirements, you must identify the autonomy and isolation requirements for all groups in your organization. Specifically, you must identify the need for data isolation, data autonomy, service isolation, and service autonomy. You must also identify areas of limited connectivity in your organization.

**Data Isolation**

Data isolation involves exclusive control over data by the group or organization that owns the data. It is important to note that service administrators have the ability to take control of a resource away from data administrators, and data administrators do not have the ability to prevent service administrators from accessing the resources that they control. Therefore, you cannot achieve data isolation when another group within the organization is responsible for service administration. If a group requires data isolation, that group must also assume responsibility for service administration.

Because data stored in Active Directory and on computers joined to Active Directory cannot be isolated from service administrators, the only way for a group within an organization to achieve complete data isolation is to create a separate forest for that data. Organizations for which the consequences of an attack by a malicious or coerced service administrator are substantial might choose to create a separate forest to achieve data isolation. Legal requirements typically create a need for this type of data isolation. For example:

- A financial institution is required by law to limit access to data that belongs to clients in a particular jurisdiction to users, computers, and administrators located in that jurisdiction. Although the institution trusts service administrators that work outside the protected area, if the access limitation is violated, the institution will no longer be able to do business in that jurisdiction. Therefore, the financial institution must isolate data from service administrators outside that jurisdiction. Note that encryption is not always an alternative to this solution. Encryption might not protect data from service administrators.

- A defense contractor is required by law to limit access to project data to a specified set of users. Although the contractor trusts service administrators who control computer systems related to other projects, a violation of this access limitation will cause the contractor to lose business.

  **Note**

  - If you have a data isolation requirement, then you must decide whether you need to isolate your data from service administrators or from data administrators and ordinary users. If your isolation requirement is based on isolation from data administrators and ordinary users, then you can use ACLs to isolate the data. For the purposes of this design process, isolation from data administrators and ordinary users is not considered a data isolation requirement.

**Data Autonomy**

Data autonomy involves the ability of a group or organization to manage its own data, including making administrative decisions about the data and performing any required administrative tasks, without the need for approval from another authority.

Data autonomy does not prevent service administrators in the forest from accessing the data. For example, a research group within a large organization might want to be able to manage their project-specific data themselves, but not need to secure the data from other administrators in the forest.

**Service Isolation**

Service isolation involves exclusive control of the Active Directory infrastructure. Groups that require service isolation require that no administrator outside of the group can interfere with the operation of the directory service.

Operational or legal requirements typically create a need for service isolation. For example:

- A manufacturing company has a critical application that controls equipment on the factory floor. Interruptions in the service on other parts of the network of the organization cannot be allowed to interfere with the operation of the factory floor.

- A hosting company provides service to multiple clients. Each client requires service isolation so that any service interruption that affects one client does not affect the other clients.

**Service Autonomy**
Service autonomy involves the ability to manage the infrastructure without a requirement for exclusive control. This means that a group wants to make changes to the infrastructure, such as adding or removing domains, modifying the DNS namespace, or modifying the schema, without the approval of the forest owner.

Service autonomy might be required for a group within an organization that wants to be able to control the service level of Active Directory by adding and removing domain controllers as needed, or that needs to be able to install directory enabled applications that require schema extensions.

**Limited Connectivity**
If a group within your organization owns networks that are separated by devices that restrict or limit connectivity between networks, such as firewalls and Network Address Translation (NAT) devices, this can impact your forest design. When you identify your forest design requirements, be sure to note the locations where you have limited network connectivity. This information is required to enable you to make decisions regarding the forest design.

# Documenting the Forest Design Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The design team must document the isolation and autonomy requirements for service and data administration for each group in the organization that intends to use Active Directory. The team must also note any areas of limited connectivity that might affect the deployment of Active Directory.

For a worksheet to assist you in documenting your forest design requirements, see "Forest Design Requirements" (DSSLOGI_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Forest Design Requirements" on the Web at http://www.microsoft.com/reskit).

**Example: Documenting the Forest Design Requirements**
The Pharmaceuticals group of Contoso Corporation needs to keep their data confidential; therefore, they require complete administrative control over their information. They need to write a justification so that the design team understands their administrative requirements.

Figure 2.7 shows the forest design requirements worksheet that the design team created in conjunction with the Pharmaceuticals group.

**Figure 2.7 Example of Forest Design Requirements Worksheet**

Forest Design Requirements

| | | | |
|---|---|---|---|
| Prepared By | Maria Hammond | Date | 2/19/03 |
| Group | Pharmaceuticals Division | | |

Instructions: Complete the indicated information for the contacts in the group.

| Name | Role | E-mail | Phone |
|---|---|---|---|
| Yan Li | Pharmaceuticals Division IT Manager | yan@contoso.com | x164 |
| Lori Kane | Pharmaceuticals Division IT Analyst | lorik@contoso.com | x743 |
| Jay Adams | Pharmaceuticals Division IT Analyst | jay@contoso.com | x386 |
| Toby Nixon | Pharmaceuticals Division IT Analyst | toby@contoso.com | x224 |

In the area below, list the design requirements. Add check marks to the columns indicating the isolation and autonomy requirements for your decisions.

| | DI | DA | SI | SA | LC |
|---|---|---|---|---|---|
| The Pharmaceuticals Division maintains confidential research data. Only members of the division should be granted access to the information stored in the forest. Administrative access must be controlled. | x | | | | |
| The computers containing research data are on a network that is separated from the organization's network by a firewall. | | | | | x |
| Pharmaceuticals Division administrators must have complete autonomy to manage the information while the data must be isolated from administrators outside of the division. | x | | | | |
| | | | | | |

DI=Data Isolation DA=Data Autonomy SI=Service Isolation SA=Service Autonomy
LC=Limited Connectivity

# Determining the Number of Forests Required

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order to determine the number of forests that you must deploy, you need to carefully identify and evaluate the isolation and autonomy requirements for each group in your organization and map those requirements to the appropriate forest design models.

Note that if you are coming from a Windows NT 4.0 environment, the decision process for creating your Active Directory logical structure design is different than the decision process for deploying Windows NT 4.0 Master User Domains (MUDs). It is not always appropriate for you to attempt to model your existing Windows NT 4.0 infrastructure when creating your Active Directory forest design.

When determining the number of forests to deploy for your organization, consider the following points:

- Isolation requirements limit your design choices. Therefore, if you identify isolation requirements, be sure that the groups actually require data isolation and that data autonomy is not sufficient for their needs. Ensure that the various groups in your organization clearly understand the concepts of isolation and autonomy.

- Negotiating the design can be a lengthy process, and it can be difficult for groups to come to agreement about ownership and utilization of available resources. Ensure that you allow enough time for the groups in your organization to conduct adequate research to identify their needs. Set firm deadlines for design decisions and get consensus from all parties on the established deadlines.

- Determining the number of forests to deploy involves balancing costs against benefits. A single forest model is the most cost-effective option and requires the least amount of administrative overhead. Although a group in the organization might prefer autonomous service operations, it might be more cost-effective for the organization to subscribe to service delivery from a centralized, trusted IT group, allowing the group to own data management without creating the added costs of service management. Balancing costs against benefits might require input from the executive sponsor.

  A single forest is the easiest configuration to manage and allows for maximum collaboration within the environment because:

  - All objects in a single forest are listed in the global catalog. Therefore, no synchronization across forests is required.

  - Management of a duplicate infrastructure is not required.

- Co-ownership of a single forest by two separate and autonomous IT organizations is not recommended. In the future, the goals of the two IT groups might change, so that they can no longer accept shared control.

- Outsourcing service administration to more than one outside partner is not recommended. Multinational organizations that have groups in different countries or regions might choose to outsource service administration to a different outside partner for each country or region. Because multiple outside partners cannot be isolated from one another, the actions of one partner can affect the service of the other, which makes it difficult to hold the partners accountable to their service level agreements.

- Only one instance of an Active Directory domain should exist at any time. Microsoft does not support cloning, splitting, or copying domain controllers from one domain in an attempt to establish a second instance of the same domain. For further information about this limitation, see the following section, "Restructuring Limitations."

# Forest Design Models

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can apply one of the following three forest design models in your Active Directory environment:

- Organizational forest model
- Resource forest model
- Restricted access forest model

It is likely that you will need to use a combination of these models in order to meet the needs of all the different groups in your organization.

**Organizational Forest Model**
In the organizational forest model, user accounts and resources are contained in the forest and managed independently. The organizational forest can be used to provide service autonomy, service isolation, or data isolation, if the forest is configured to prevent access to anyone outside the forest.

If users in an organizational forest need to access resources in other forests, or vice versa, trust can be established between one organizational forest and the other forests, as shown in Figure 2.8. This makes it possible for administrators to grant access to resources in the other forest.

**Figure 2.8 Organizational Forest Model**




Every Active Directory design includes at least one organizational forest.

**Resource Forest Model**
In the resource forest model, a separate forest is used to manage resources. Resource forests do not contain user accounts other than those required for service administration and those required to provide alternate access to the resources in that forest if the user accounts in the organizational forest become unavailable. Forest trusts are established so that users from other forests can access the resources contained in the resource forest, as shown in Figure 2.9.

**Figure 2.9 Resource Forest Model**




Resource forests provide service isolation that is used to protect areas of the network that need to maintain a state of high availability. For example, if your company includes a manufacturing facility that needs to continue to operate when there are problems on the rest of the network, you can create a separate resource forest for the manufacturing group.

**Restricted Access Forest Model**

In the restricted access forest model, a separate forest is created to contain user accounts and data that must be isolated from the rest of the organization. Restricted access forests provide data isolation in a situation for project data for which the consequences of compromise are severe. Figure 2.10 shows a restricted access forest model.

**Figure 2.10 Restricted Access Forest Model**




Users from other forests cannot be granted access to the restricted data because no trust exists. In this model, users have an account in an organizational forest for access to general organizational resources and a separate user account in the restricted access forest for access to the classified data. These users must have two separate workstations, one connected to the organizational forest and the other connected to the restricted access forest. This protects against the possibility that a service administrator from one forest can gain access to a workstation in the restricted forest.

In extreme cases, the restricted access forest might be maintained on a separate physical network. Organizations that work on classified government projects sometimes maintain restricted access forests on separate networks in order to meet security requirements.

# Mapping Design Requirements to Forest Design Models

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most groups in your organization can share a single organizational forest that is managed by a single IT group and that contains the user accounts and resources for all of the groups that share the forest. This shared forest, called the *initial organizational forest*, is the foundation of the forest design model for the organization.

Because the initial organizational forest can host multiple groups in the organization, the forest owner must establish service level agreements with each group so that all parties understand what is expected of them. This protects both the individual groups and the forest owner by establishing agreed-upon service expectations.

If not all of the groups in your organization can share a single organizational forest, you must expand your forest design to accommodate the needs of the different groups. This involves identifying the design requirements that apply to the groups, based on their needs for autonomy and isolation and whether or not they have a limited-connectivity network, and then identifying the forest model that you can use to accommodate those requirements. Use Table 2.1 to identify the forest design model scenario that best meets your needs based on the autonomy, isolation, and connectivity factors that you identified. After you identify your forest design scenario, determine whether you need to make any additional decisions in order to meet your design requirements.

**Note**

- If a factor is listed as N/A, it is not a consideration because other requirements also accommodate that factor.

**Table 2.1 Scenarios Based on Isolation and Autonomy Design Requirements**

| Limited Connectivity | Data Isolation | Data Autonomy | Service Isolation | Service Autonomy | Scenario |
|---|---|---|---|---|---|
| No | No | Yes | No | No | Scenario 1: Join an existing forest for data autonomy. |
| No | No | N/A | No | Yes | Scenario 2: Use an organizational forest or domain for service autonomy. |
| No | No | N/A | Yes | N/A | Scenario 3: Use an organizational or resource forest for service isolation. |
| N/A | Yes | N/A | N/A | N/A | Scenario 4: Use an organizational or restricted access forest for data isolation. |
| Yes | No | N/A | No | No | Scenario 5: Use an organizational forest or reconfigure the firewall for limited connectivity. |
| Yes | No | N/A | No | Yes | Scenario 6: Use an organizational forest or domain and reconfigure the firewall for service autonomy with limited connectivity. |
| Yes | No | N/A | Yes | N/A | Scenario 7: Use a resource forest and reconfigure the firewall for service isolation with limited connectivity. |

**Scenario 1: Join an Existing Forest for Data Autonomy**

You can meet a requirement for data autonomy simply by hosting the group in OUs in an existing organizational forest. Delegate control over the OUs to data administrators from that group to achieve data autonomy. For more information about delegating control by using OUs, see "Designing Organizational Units for Delegation of Administration" later in this chapter.

**Scenario 2: Organizational Forest or Domain for Service Autonomy**

If a group in your organization identifies service autonomy as a requirement, it is recommended that you first reconsider this requirement. Achieving service autonomy creates more management overhead and additional costs for the organization. Ensure that the requirement for service autonomy is not simply for convenience and that you can justify the costs involved in

meeting this requirement.

You can meet a requirement for service autonomy by doing one of the following:

- Creating an organizational forest. Place the users, groups, and computers for the group that requires service autonomy into a separate organizational forest. Assign an individual from that group to be the forest owner. If the group needs to access or share resources with other forests in the organization, then they can establish a trust between their organizational forest and the other forests.

- Using organizational domains. Place the users, groups, and computers in a separate domain in an existing organizational forest. This model provides for domain-level service autonomy only, and not full service autonomy or service or data isolation. Other groups in the forest must trust the service administrators of the new domain to the same degree that they trust the forest owner. For this reason, this approach is not recommended.

  For more information about using organizational domains, see "Using the Organizational Domain Forest Model" later in this chapter.

**Scenario 3: Organizational or Resource Forest for Service Isolation**
You can meet a requirement for service isolation by doing one of the following:

- Using an organizational forest. Place the users, groups, and computers for the group that requires service isolation into a separate organizational forest. Assign an individual from that group to be the forest owner. If the group needs to access or share resources with other forests in the organization, then they can establish a trust between their organizational forest and the other forests.

- Using a resource forest. Place resources and service accounts into a separate resource forest, keeping user accounts in an existing organizational forest. If necessary, alternate accounts can be created in the resource forest to access resources in the resource forest in the event that the organizational forest becomes unavailable. The alternate accounts must have the authority required to log on to the resource forest and maintain control of the resources until the organizational forest is back on line.

  Establish a trust between the resource and organizational forests so that the users can access the resources in the forest while using their regular user accounts. This configuration enables centralized management of user accounts, while allowing users to fall back to alternate accounts in the resource forest in the event that the organizational forest becomes unavailable.

Considerations for service isolation include the following:

- Forests created for service isolation can trust domains from other forests, but must not include users from other forests in their service administrators groups. If users from other forests are included in administrative groups in the isolated forest, the security of the isolated forest potentially can be compromised because the service administrators in the forest do not have exclusive control.

- As long as domain controllers are accessible on a network, they are subject to attacks (such as denial of service attacks) from computers on that network. You can do the following to protect against the possibility of an attack:

  - Host domain controllers only on networks that are considered secure.

  - Limit access to the network(s) hosting the domain controllers.

- Service isolation requires the creation of an additional forest. Evaluate whether the cost of maintaining the infrastructure to support the additional forest outweighs the costs associated with loss of access to resources due to an organizational forest being unavailable.

**Scenario 4: Organizational or Restricted Access Forest for Data Isolation**
You can achieve data isolation by doing one of the following:

- Using an organizational forest. Place the users, groups, and computers for the group that requires data isolation into a

separate organizational forest. Assign an individual from that group to be the forest owner. If the group needs to access or share resources with other forests in the organization, then establish a trust between the organizational forest and the other forests. Only the users who require access to the classified information exist in the new organizational forest. Users have one account that they use to access both classified data in their own forest and unclassified data in other forests by means of trust relationships.

- Using a restricted access forest. This is a separate forest that contains the restricted data and the user accounts that are used to access that data. Separate user accounts are maintained in the existing organizational forests that are used to access the unrestricted resources on the network. No trusts are created between the restricted access forest and other forests in the enterprise. You can further restrict the forest by deploying the forest on a separate physical network so that it cannot connect to other forests. If you deploy the forest on a separate network, users must have two workstations: one for accessing the restricted forest and one for accessing the nonrestricted areas of the network.

Considerations for creating forests for data isolation include the following:

- Organizational forests created for data isolation can trust domains from other forests, but users from other forests must not be included in any of the following:

    - Groups responsible for service management or groups that can manage the membership of service administrator groups.

    - Groups that have administrative control over computers that store protected data.

    - Groups that have access to protected data, or groups that are responsible for the management of user or group objects that have access to protected data.

  If users from another forest are included in any of these groups, then a compromise of the other forest might lead to a compromise of the isolated forest and to disclosure of protected data.

- Other forests can be configured to trust the organizational forest created for data isolation so that users in the isolated forest can access resources in other forests. However, users from the isolated forest must never interactively log on to workstations in the trusting forest. The computer in the trusting forest can potentially be compromised by a malicious administrator and can be used to capture the logon credentials of the user.

  **Note**

    - To prevent servers in a trusting forest from impersonating users from the isolated forest and then accessing resources in the isolated forest, the forest owner can disable delegated authentication or use the constrained delegation feature. For more information about delegated authentication and constrained delegation, see "Delegating authentication" on the Microsoft Web site.

- You might need to establish a firewall between the organizational forest and the other forests in the organization to limit users' access to information outside of their forest.

- Although creating a separate forest enables data isolation, as long as the domain controllers in the isolated forest and computers that host protected information are accessible on a network, they are subject to attacks launched from computers on that network. Organizations that decide that the risk of attack is too high, or that the consequence of an attack or security violation is too great, need to limit access to the network(s) hosting the domain controllers and computers hosting protected data by using technologies such as firewalls and Internet Protocol security (IPSec). In extreme cases, they might choose to maintain the protected data on an independent network that has no physical connection to any other network in the organization.

  **Note**

    - If any network connectivity exists between a restricted access forest and another network, then the possibility exists for data in the restricted area to be transmitted to the other network.

**Scenario 5: Organizational Forest or Reconfigured Firewall for Limited Connectivity**
To meet a limited connectivity requirement, you can do one of the following:

- Place users into an existing organizational forest and then open the firewall enough to allow Active Directory traffic to pass through. For more information about configuring firewalls for use with Active Directory, see the Active Directory in Networks Segmented by Firewalls link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Use an organizational forest. Place the users, groups, and computers for the group for which connectivity is limited into a separate organizational forest. Assign an individual from that group to be the forest owner. The organizational forest provides a separate environment on the other side of the firewall. The forest includes user accounts and resources that are managed within the forest, so that users do not need to go through the firewall to accomplish their daily tasks. Specific users or applications might have special needs that require the ability to pass through the firewall to contact other forests. You can address those needs individually by opening the appropriate interfaces in the firewall, including those necessary for any trusts to function. For more information about configuring firewalls for use with Active Directory, see the Active Directory in Networks Segmented by Firewalls link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Scenario 6: Organizational Forest or Domain and Reconfigured Firewall for Service Autonomy with Limited Connectivity**
If a group in your organization identifies service autonomy as a requirement, it is recommended that you first reconsider this requirement. Achieving service autonomy creates more management overhead and additional costs for the organization. Ensure that the requirement for service autonomy is not simply for convenience and that you can justify the costs involved in meeting this requirement.

If limited connectivity is an issue and you have a requirement for service autonomy, then you can do one of the following:

- Use an organizational forest. Place the users, groups, and computers for the group that requires service autonomy into a separate organizational forest. Assign an individual from that group to be the forest owner. The organizational forest provides a separate environment on the other side of the firewall. The forest includes user accounts and resources that are managed within the forest, so that users do not need to go through the firewall to accomplish their daily tasks. Specific users or applications might have special needs that require the ability to pass through the firewall to contact other forests. You can address those needs individually by opening the appropriate interfaces in the firewall, including those necessary for any trusts to function. For more information about configuring firewalls for use with Active Directory, see the Active Directory in Networks Segmented by Firewalls link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Place the users, groups, and computers in a separate domain in an existing organizational forest. This model provides for domain-level service autonomy only, and not full service autonomy or service or data isolation. Other groups in the forest must trust the service administrators of the new domain to the same degree that they trust the forest owner. For this reason, this approach is not recommended. For more information about using organizational domains, see "Using the Organizational Domain Forest Model" later in this chapter.

You also need to open the firewall enough to allow Active Directory traffic to pass through. For more information about configuring firewalls for use with Active Directory, see the Active Directory in Networks Segmented by Firewalls link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources..

**Scenario 7: Resource Forest and Reconfigured Firewall for Service Isolation with Limited Connectivity**
If limited connectivity is an issue and you have a requirement for service isolation, then you can do one of the following:

- Use an organizational forest. Place the users, groups, and computers for the group that requires service isolation into a separate organizational forest. Assign an individual from that group to be the forest owner. The organizational forest provides a separate environment on the other side of the firewall. The forest includes user accounts and resources that are managed within the forest, so that users do not need to go through the firewall to accomplish their daily tasks. Specific users or applications might have special needs that require the ability to pass through the firewall to contact other forests. You can address those needs individually by opening the appropriate interfaces in the firewall, including those necessary for any trusts to function. For more information about configuring firewalls for use with Active Directory, see the Active Directory in Networks Segmented by Firewalls link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Use a resource forest. Place resources and service accounts into a separate resource forest, keeping user accounts in an existing organizational forest. It might be necessary to create some alternate user accounts in the resource forest to maintain access to the resource forest in the event that the organizational forest becomes unavailable. The alternate accounts must have the authority required to log on to the resource forest and maintain control of the resources until the organizational forest is back on line.

  Establish a trust between the resource and organizational forests so that the users can access the resources in the forest while using their regular user accounts. This configuration enables centralized management of user accounts, while allowing users to fall back to alternate accounts in the resource forest in the event that the organizational forest becomes unavailable.

Considerations for service isolation include the following:

- Forests created for service isolation can trust domains from other forests, but must not include users from other forests in their service administrators groups. If users from other forests are included in administrative groups in the isolated forest, the security of the isolated forest potentially can be compromised because the service administrators in the forest do not have exclusive control.

- As long as domain controllers are accessible on a network, they are subject to attacks (such as denial of service attacks) from computers on that network. You can do the following to protect against the possibility of an attack:

  - Host domain controllers only on networks that are considered secure.

  - Limit access to the network(s) hosting the domain controllers.

- Service isolation requires the creation of an additional forest. Evaluate whether the cost of maintaining the infrastructure to support the additional forest outweighs the costs associated with loss of access to resources due to an organizational forest being unavailable.

  Specific users or applications might have special needs that require the ability to pass through the firewall to contact other forests. You can address those needs individually by opening the appropriate interfaces in the firewall, including those necessary for any trusts to function. For more information about configuring firewalls for use with Active Directory, see the Active Directory in Networks Segmented by Firewalls link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Using the Organizational Domain Forest Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the organizational domain forest model, several autonomous groups each own a domain within a forest. Each group controls domain-level service administration, which enables them to manage certain aspects of service management autonomously while the forest owner controls forest-level service management.

Because a domain is not a security boundary, it cannot be used for service or data isolation. For this reason, it is important to ensure that service or data isolation are not requirements for the groups in the organization. In an organizational domain forest model, a malicious service administrator in one domain can access any other domain within the forest.

If you have an existing Windows NT 4.0 environment that is based on organizational domains, it is not recommended that you reproduce this configuration in an Active Directory forest. The Active Directory forest does not have the same administrative and security boundaries that exist in a Windows NT 4.0 environment.

Figure 2.11 shows an organizational domain forest model.

**Figure 2.11 Organizational Domain Forest Model**




**Note**

- Because the organizational domain forest model is not recommended, additional design information is not provided.

**Domain-Level Service Autonomy**

The organizational domain forest model enables the delegation of authority for domain-level service management. Table 2.2 lists the types of service management that can be controlled at the domain level.

**Table 2.2 Types of Service Management That are Controlled at the Domain Level**

| Type of Service Management | Associated tasks |
|---|---|
| Management of domain controller operations | • Creating and removing domain controllers.<br>• Monitoring the functioning of domain controllers.<br>• Managing services that are running on domain controllers.<br>• Backing up and restoring the directory. |

| | |
|---|---|
| Configuration of domain-wide settings | • Creating domain and domain user account policies, such as password, Kerberos, and account lockout policies.<br>• Creating and applying domain-wide Group Policies. |
| Delegation of data-level administration | • Creating OUs and delegating administration.<br>• Repairing problems in the OU structure that OU owners do not have sufficient access rights to fix. |
| Management of external trusts | • Establishing trust relationships with domains outside the forest. |

Other types of service management, such as schema or replication topology management, are the responsibility of the forest owner.

**Domain Owner**

In an organizational domain forest model, *domain owners* are responsible for domain-level service management tasks. Domain owners have authority over the entire domain, as well as access to all other domains in the forest. For this reason, domain owners must be trusted individuals selected by the forest owner.

Delegate domain-level service management to a domain owner if the following conditions are met:

- All groups participating in the forest trust the new domain owner and the service management practices of the new domain.
- The new domain owner trusts the forest owner and all the other domain owners.
- All domain owners in the forest agree that the new domain owner has service administrator management and selection policies and practices that are equal to or more strict than their own.
- All domain owners in the forest agree that domain controllers managed by the new domain owner in the new domain are physically secure.

Keep in mind that if a forest owner delegates domain-level service management to a domain owner, then other groups might choose not to join that forest if they do not trust that domain owner.

All domain owners must be aware that if any of these conditions change in the future, it might become necessary to migrate the organizational domains into a multiple forest deployment.

# Documenting the Forest Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you map all of your design requirements to forest models and select the forest model that meets the needs of your organization, you need to document the proposed forest design. Include in your documentation the name of the group for which the forest is designed, the contact information for the forest owner, the type of forest for each forest that you include, and the requirements that each forest is designed to meet. This documentation will help the design team to ensure that all of the appropriate people are involved in the design process and to clarify the scope of the deployment project.

For a worksheet to assist you in documenting the proposed forest design, see "Proposed Forests" (DSSLOGI_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Proposed Forests" on the Web at http://www.microsoft.com/reskit).

At this point, you also need to establish and document the design and deployment teams for each forest. For a worksheet to assist you in documenting the design and deployment teams for each forest, see "Design and Deployment Team Information" (DSSLOGI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Design and Deployment Team Information" on the Web at http://www.microsoft.com/reskit).

# Example: Documenting the Forest Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso Corporation decided to use a single organizational forest for most of the organization. In addition, the Pharmaceuticals group needs to keep their data confidential so they want complete administrative control over their own organizational forest. The manufacturing group wants their own resource forest so that the service accounts used to operate the factory applications can continue to authenticate even if the main corporate forest goes offline. As a result, Contoso is deploying three forests for their forest design.

The design team documented the forests as shown in Figure 2.12.

**Figure 2.12 Example of a Completed Forest Design Worksheet**

Forest Design

| Prepared By | Maria Hammond | Date | 2/19/03 |
|---|---|---|---|

Instructions: Record the group names of any new proposed forest owners, the contact information for the representatives of those groups, the forest type, and the requirements that the forest is designed to meet.

| Group Name | Contact | Forest Type | Requirements |
|---|---|---|---|
| Central IT | Mark Lee<br>mark@contoso.com<br>X874 | Organizational | This is the initial forest created for the corporation. Many groups within the company use this forest. |
| Pharmaceuticals Division | Mary Baker<br>mary@contoso.com<br>x989 | Restricted Access | This group needs data isolation and they need to control who can access the project data for the group. |
| Manufacturing Division | Paul West<br>paul@contoso.com<br>X804 | Resource | The organization needs to maintain control of the manufacturing forest even if the corporate forest is unavailable. |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

# Creating a Domain Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The forest owner is responsible for creating a domain design for the forest. Creating a domain design involves examining the replication requirements and the existing capacity of your network infrastructure, and then building a domain structure that enables Active Directory to function in the most efficient way. Figure 2.13 shows the process for creating a domain design.

**Figure 2.13 Creating a Domain Design**




Domains are used to partition the directory so that the information in the directory can be distributed and managed efficiently throughout the enterprise. The goal for your domain design is to maximize the efficiency of the Active Directory replication topology while ensuring that replication does not use too much available network bandwidth and does not interfere with the daily operation of your network.

# Reviewing the Domain Models

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following factors impact the domain design model that you select:

- The amount of available capacity on your network that you are willing to allocate to Active Directory. The goal is to select a model that provides efficient replication of information with minimal impact on available network bandwidth.

- The number of users in your organization. If your organization includes a large number of users, deploying more than one domain enables you to partition your data and gives you more control over the amount of replication traffic that will pass through a given network connection. This makes it possible for you to control where data is replicated and reduce the load created by replication traffic on slow links in your network.

The simplest domain design is a single domain. In a single domain design, all information is replicated to all of the domain controllers. If necessary, however, you can deploy additional regional domains. This might come about if portions of the network infrastructure are connected by slow links and the forest owner wants to be sure that replication traffic does not exceed the capacity that has been allocated to Active Directory.

It is best to minimize the number of domains that you deploy in your forest. This reduces the overall complexity of the deployment and, as a result, reduces total cost of ownership. Table 2.3 lists the administrative costs associated with adding regional domains.

**Table 2.3 Administrative Costs Associated With Adding Regional Domains to a Forest**

| Cost | Implications |
|---|---|
| Management of multiple Service Administrator groups | Each domain has its own service administrator groups that need to be managed independently. The membership of these service administrator groups must be carefully controlled. |
| Maintaining consistency among Group Policies that are common to multiple domains | Group Policy settings that need to be applied forest-wide must be applied separately to each individual domain in the forest. |
| Maintaining consistency among access control and auditing settings that are common to multiple domains | Access control and auditing settings that need to be applied forest-wide must be applied separately to each individual domain in the forest. |
| Increased likelihood of objects moving between domains | The greater the number of domains, the greater the likelihood that users will need to move from one domain to another. This migration can potentially impact end users. |

# Single Domain Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A single domain model is the easiest to administer and the least expensive to maintain. It consists of a forest that contains a single domain. This domain is the forest root domain and it contains all of the user and group accounts in the forest.

A single domain forest model reduces administrative complexity by providing the following advantages:

- Any domain controller can authenticate any user in the forest.
- All domain controllers can be global catalogs; therefore, you do not need to plan for global catalog server placement.

In a single domain forest, all directory data is replicated to all geographic locations that host domain controllers. While this model is the easiest to manage, it also creates the most replication traffic of the two domain models. Partitioning the directory into multiple domains limits the replication of objects to specific geographic regions but results in more administrative overhead.

# Regional Domain Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

All object data within a domain is replicated to all domain controllers in that domain. For this reason, if your forest includes a large number of users that are distributed across different geographic locations connected by a wide area network (WAN), you might need to deploy regional domains to reduce replication traffic over the WAN links. Geographically based regional domains can be organized according to network WAN connectivity.

The regional domain model enables you to maintain a stable environment over time. Base the regions used to define domains in your model on stable elements such as continental boundaries. Domains based on other factors, such as groups within the organization, can change frequently and might require you to restructure your environment.

The regional domain model consists of a forest root domain and one or more regional domains, as shown in Figure 2.14.

**Figure 2.14 Regional Domain Model**




Creating a regional domain model design involves identifying which domain is the forest root domain and determining the number of additional domains that are required in order to meet your replication requirements.

**Note**

- If your organization includes groups that require data or service isolation from other groups in the organization, create a separate forest for those groups. Domains do not provide data or service isolation.

# Determining the Number of Domains Required

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Every forest starts with a single domain. The maximum number of users that a single domain forest can contain is based on the slowest link that must accommodate replication between domain controllers and the available bandwidth that you want to allocate to Active Directory. Table 2.4 lists the maximum recommended number of users that a domain can contain based on a single domain forest, the speed of the slowest link and the percentage of bandwidth you want to reserve for replication. This information applies to forests that contain a maximum of 100,000 users and that have a connectivity of 28.8 Kilobits per second (Kbps) or higher.

**Note**

- For recommendations that apply to forests that contain more than 100,000 users or connectivity of less than 28.8 Kbps, consult an experienced Active Directory designer.

**Table 2.4 Maximum Number of Users in a Single Domain**

| Slowest Link Connecting a Domain Controller (Kbps) | Maximum Number of Users if 1% Bandwidth Available | Maximum Number of Users if 5% Bandwidth Available | Maximum Number of Users if 10% Bandwidth Available |
|---|---|---|---|
| 28.8 | 10,000 | 25,000 | 40,000 |
| 32 | 10,000 | 25,000 | 50,000 |
| 56 | 10,000 | 50,000 | 100,000 |
| 64 | 10,000 | 50,000 | 100,000 |
| 128 | 25,000 | 100,000 | 100,000 |
| 256 | 50,000 | 100,000 | 100,000 |
| 512 | 80,000 | 100,000 | 100,000 |
| 1500 | 100,000 | 100,000 | 100,000 |

To use this table:

1. In the Slowest Link Connecting a Domain Controller column, locate the value that matches the speed of the slowest link across which Active Directory will replicate in your domain.

2. In the row that corresponds to your slowest link speed, locate the column that represents the percentage bandwidth you want to allocate to Active Directory. The value at that location is the maximum number of users that the domain in a single domain forest can contain.

The values in Table 2.4 are based on the replication traffic generated in an environment that has the following characteristics:

- New users join the forest at a rate of 20 percent per year.

- Users leave the forest at a rate of 15 percent per year.

- Each user is a member of five global groups and five universal groups.

- The ratio of users to computers is 1:1.

- Active Directory–integrated DNS is used.
- DNS scavenging is used.

**Note**

- The figures listed in Table 2.4 are approximations. The quantity of replication traffic depends largely on the number of changes made to the directory in a given amount of time. Confirm that your network can accommodate your replication traffic by testing the estimated quantity and rate of changes on your design in a lab before deploying your domains.

If you determine that the total number of users in your forest is less than the maximum number of users that your domain can contain, then you can use a single domain. Be sure to accommodate for planned future growth when you make this determination. If you determine that the total number of users in your forest is greater than the maximum number of users that your domain can contain, then you need to reserve a higher percentage of bandwidth for replication, increase your link speed, or divide your organization into regional domains.

# Dividing the Organization into Regional Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you cannot accommodate all of your users in a single domain, then you must select the regional domain model. Divide your organization into regions in a way that makes sense for your organization and your existing network. For example, you might create regions based on continental boundaries.

Keep in mind that because you need to create an Active Directory domain for each region that you establish, it is best to minimize the number of regions that you define for Active Directory. Although it is possible to include an unlimited number of domains in a forest, for manageability reasons it is recommended that a forest include no more than ten domains. You must establish the appropriate balance between optimizing your replication bandwidth and minimizing your administrative complexity when dividing your organization into regional domains.

First, determine the maximum number of users that your forest can host based on the slowest link in the forest across which domain controllers will replicate and the average amount of bandwidth you want to allocate to Active Directory replication. Table 2.5 lists the maximum recommended number of users that a forest can contain based on the speed of the slowest link, and the percentage of bandwidth you want to reserve for replication. This information applies to forests that contain a maximum of 100,000 users and that have a connectivity of 28.8 KBps or higher.

**Table 2.5 Maximum Number of Users in a Forest**

| Slowest Link Connecting a Domain Controller (KBps) | Maximum Number of Users if 1% Bandwidth Available | Maximum Number of Users if 5% Bandwidth Available | Maximum Number of Users if 10% Bandwidth Available |
|---|---|---|---|
| 28.8 | 10,000 | 50,000 | 75,000 |
| 32 | 10,000 | 50,000 | 75,000 |
| 56 | 10,000 | 75,000 | 100,000 |
| 64 | 25,000 | 75,000 | 100,000 |
| 128 | 50,000 | 100,000 | 100,000 |
| 256 | 75,000 | 100,000 | 100,000 |
| 512 | 100,000 | 100,000 | 100,000 |
| 1500 | 100,000 | 100,000 | 100,000 |

To use this table:

1. In the Slowest Link Connecting a Domain Controller column, locate the value that matches the speed of the slowest link across which Active Directory will replicate in your forest.

2. In the row that corresponds to your slowest link speed, locate the column that represents the percentage bandwidth that you want to allocate to Active Directory. The value at that location is the maximum number of users that your forest can host.

**Note**

- The figures listed in Table 2.5 are approximations. The quantity of replication traffic depends largely on the number of changes made to the directory in a given amount of time. Confirm that your network can accommodate your replication traffic by testing the estimated quantity and rate of changes on your design in a lab before deploying your domains.

The values in Table 2.5 are based on the following assumptions:

- All domain controllers are global catalog servers.
- New users join the forest at a rate of 20 percent per year.
- Users leave the forest at a rate of 15 percent per year.
- Users are members of five global groups and five universal groups.
- The ratio of users to computers is 1:1.
- Active Directory–integrated DNS is used.
- DNS scavenging is used.

If the maximum number of users that your forest can host is greater than the number of users that you need to host, then a single forest will work for your design. If you need to host more users than the maximum number that you identified, then you need to increase the minimum link speed, allocate a greater percentage of bandwidth for Active Directory, or deploy additional forests.

If you determine that a single forest will accommodate the number of users that you need to host, the next step is to determine the maximum number of users that each region can support based on the slowest link located in that region. Divide your forest into regions that make sense to you. Make sure that you base your decision on something that is not likely to change. For example, use continents instead of sales regions. These regions will be the basis of your domain structure when you have identified the maximum number of users.

Determine the number of users that need to be hosted in each region and then verify that they do not exceed the maximum allowed based on the slowest link speed and the bandwidth allocated to Active directory in that region. Table 2.6 lists the maximum recommended number of users that a regional domain can contain based on the speed of the slowest link and the percentage of bandwidth you want to reserve for replication. This information applies to forests that contain a maximum of 100,000 users and that have a connectivity of 28.8 KBps or higher.

**Table 2.6 Maximum Number of Users in a Region**

| Slowest Link Connecting a Domain Controller (KBps) | Maximum Number of Users if 1% Bandwidth Available | Maximum Number of Users if 5% Bandwidth Available | Maximum Number of Users if 10% Bandwidth Available |
|---|---|---|---|
| 28.8 | 10,000 | 18,000 | 40,000 |
| 32 | 10,000 | 20,000 | 50,000 |
| 56 | 10,000 | 40,000 | 100,000 |
| 64 | 10,000 | 50,000 | 100,000 |
| 128 | 15,000 | 100,000 | 100,000 |
| 256 | 30,000 | 100,000 | 100,000 |
| 512 | 80,000 | 100,000 | 100,000 |
| 1500 | 100,000 | 100,000 | 100,000 |

To use this table:

1. In the Slowest Link Connecting a Domain Controller column, locate the value that matches the speed of the slowest link across which Active Directory will replicate in your region.

2. In the row that corresponds to your slowest link speed, locate the column that represents the percentage bandwidth that you want to allocate to Active Directory. That value represents the maximum number of users that the region can host.

**Note**

- The figures listed in Table 2.6 are approximations. The quantity of replication traffic depends largely on the number of changes made to the directory in a given amount of time. Confirm that your network can accommodate your replication traffic by testing the estimated quantity and rate of changes on your design in a lab before deploying your domains.

The values in Table 2.6 are based on the following assumptions:

- All domain controllers are global catalog servers.
- New users join the forest at a rate of 20 percent per year.
- Users leave the forest at a rate of 15 percent per year.
- Users are members of five global groups and five universal groups.
- The ratio of users to computers is 1:1.
- Active Directory–integrated DNS is used.
- DNS scavenging is used.

Evaluate each proposed region and determine whether the maximum number of users in each region is less than the recommended maximum number of users that a domain can contain. If you determine that the region can host the number of users that you require, then you can create a domain for that region. If you determine that you cannot host that many users, consider dividing your design into smaller regions and recalculating the maximum number of users that can be hosted in each region. The other alternatives are to allocate more bandwidth or increase your link speed.

Although the total number of users that you can put in a domain in a multidomain forest is smaller than the number of users in the domain in a single domain forest, the overall number of users in the multidomain forest can be higher. The smaller number of users per domain in a multidomain forest accommodates the additional replication overhead created by maintaining the global catalog in that environment.

**Note**

- For recommendations that apply to forests that contain more than 100,000 users or connectivity of less than 28.8 KBps, consult an experienced Active Directory designer.

Figure 2.15 shows the regions that Contoso Corporation chose to use and the slowest link and number of users located in each region.

**Figure 2.15 Contoso Corporation Proposed Regions and the Slowest Link Speed in Each Region**


noam
(128 KBps, 40,000 users)
eur
(32 KBps, 15,000 users)
soam
(32 KBps, 5,000 users)

# Documenting the Regions Identified

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you divide your organization into regional domains, document the regions that you want represented and the number of users that will exist in each region. In addition, note the speed of the slowest links in each region that you will use for Active Directory replication. This information is used to determine whether additional domains or forests are required.

For a worksheet to assist you in documenting the regions you identified, see "Identifying Regions" (DSSLOGI_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Identifying Regions" on the Web at http://www.microsoft.com/reskit).

**Example: Documenting the Regions Identified**

Contoso Corporation has facilities all over the world. They want to create a single domain for their entire organization. However, they have 60,000 users and they do not want to allocate any more than 10 percent of the total bandwidth to Active Directory. Based on the slowest link in that Active Directory would use, they can only host up to 50,000 users. With 32 KBps links they can host up to 75,000 users in a single forest as long as they are willing to allocate up to ten percent of the bandwidth to Active Directory. This means that they can still use a single forest but they need to use multiple domains.

Based on the slowest links in each region, Contoso decided to represent their network by using three regions: North America, South America, and Europe. The slowest link in the North America region is 128 KBps. In South America and Europe, the slowest links are 32 KBps. North America has 40,000 users, Europe has 15,000 users, and South America has 5,000 users. Based on these link speeds and an allocation of up to ten percent of the bandwidth, the number of users in each domain is within the maximum limits.

Figure 2.16 shows the documentation that Contoso created for the regions that they identified.

**Figure 2.16 Example of a Completed Identifying Regions Worksheet**

Identifying Regions

| | | | |
|---|---|---|---|
| Prepared By | Olinda Turner | Date | 2/20/03 |

| | |
|---|---|
| Forest Owner | Olinda Turner |
| Slowest link in forest | 32 kbps Sao Paolo |
| Number of users in this forest | 60,000 |
| Forests: ☒ Single Forest ☐ Multiple Forests | |
| Model: ☐ Single Domain ☒ Regional Domains | |

Instructions: List each region that is included in the forest. For each region, note the number of users and the link speed of the slowest link that will be used for Active Directory replication. Use this information to determine the number of domains required.

| Name of Region | Slow Link | # of Users | Comment |
|---|---|---|---|
| North America | 128 kbps (multiple) | 40,000 | Allocate max of 10% bandwidth |
| South America | 32 kbps (Sao Paolo) | 5,000 | Allocate max of 10% bandwidth |
| Europe | 32 kbps | 15,000 | Allocate max of 10% bandwidth |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

# Determining Whether to Upgrade Existing or Deploy New Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Each domain in your design will either be a new domain or an existing domain that has been upgraded in place. Users from existing domains that you do not upgrade in place must be migrated into new domains.

Moving accounts between domains can impact end users. Before deciding whether to migrate users into a new domain or upgrade existing domains in place, evaluate the long-term administrative benefits of a new Active Directory domain against the cost of migrating users into the domain.

If you are upgrading from Windows NT 4.0 to Windows Server 2003, evaluate each Master User Domain (MUD) in your organization to determine whether they are candidates for upgrade to an Active Directory domain. If you are migrating from an operating system other than Windows NT 4.0, plan to create new Active Directory domains. If no Windows NT 4.0 MUD upgrade candidates exist in your current environment, plan to create new Active Directory domains and migrate your existing accounts and resources into them.

For more information about consolidating and upgrading Windows NT 4.0 domains, see "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.

# Evaluating Current Master User Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Create a list of existing Windows NT 4.0 master user domains (MUDs) and compare it to the regional domains in your domain design. For each regional domain in your design, decide whether you will deploy a new domain or upgrade of one of your existing MUDs in place.

When upgrading from a Windows NT 4.0 environment to Active Directory, upgrading MUDs in place is the least expensive and most efficient method. Upgrading domains in place has no impact on end users. Migrating users to new domains is a more complicated process that creates additional administrative costs.

**Important**

- A domain must be in Windows 2000 native mode or operating at the Windows Server 2003 functional level before it can be the target for user migrations from Windows NT 4.0. If you have a large number of backup domain controllers (BDCs) in remote locations or that require hardware upgrades, you might not be able to achieve this in a timely manner. If time is a consideration for your upgrade process, consider performing an in-place upgrade of a different domain, or creating a new domain to be the target of the migration.

If you identified more than one MUD upgrade candidate, you can do one of the following:

- Upgrade the largest MUD and migrate accounts in the remaining MUDs into that MUD.
- Define additional regions for each MUD upgrade candidate and perform an in-place upgrade of each MUD into a corresponding new region. This option involves additional long-term management overhead costs but saves the short-term administrative costs associated with migrating users between domains.

Figure 2.17 shows how Contoso Corporation chose to create their new domain design based on the geographic regions where they do business.

**Figure 2.17 Contoso Corporation Proposed Domains**



# Documenting Plans for New and Upgraded Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

List your proposed regions and then indicate whether the domain that you create for that region will be a new domain or a MUD that you upgrade. If the new domain is to be the result of an upgrade, record the name of the MUD that will be upgraded and the contact information for the current MUD owner.

For a worksheet to assist you in documenting your plans for new and upgraded domains, see "Domain Planning" (DSSLOGI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Planning" on the Web at http://www.microsoft.com/reskit). You can use the same worksheet to document your domain name choices and forest root domain information later in the design process.

**Example: Documenting Plans for New and Upgraded Domains**
Contoso Corporation has 60,000 users located throughout facilities all over the world. Based on how the offices are connected to the network, they divided their organization into three regions: North America, South America, and Europe. Their current deployment has four MUDs, one in North America, one in South America, and two in Europe. The design team decided to upgrade the North and South American MUDs to Windows Server 2003 Active Directory domains, and to create a new Active Directory domain in Europe and migrate the users from the two European MUDs into the new domain.

Figure 2.18 shows the documentation that the Contoso Corporation design team created to document their plans for new and upgraded domains.

**Figure 2.18 Example of a Domain Planning Worksheet That Includes Upgrade and Migration Plans**

# Domain Planning

| | | | |
|---|---|---|---|
| Prepared By | Olinda Turner | Date | 2/20/03 |
| Forest Owner | Olinda Turner | | |

Instructions: Record information about your existing domains and any new domains that you will create. First, list the domains and whether they are new domains or existing domains that will be upgraded. Second, list the DNS prefix and NetBIOS name for each domain. Finally, list the forest root information.

| Region | Origin | DNS Prefix | NetBIOS Name | Justification / Notes |
|---|---|---|---|---|
| North America | ☐ New domain<br>☒ Upgrade<br>From: noam<br>Owner: Mark Lee x874 | | | |
| South America | ☐ New domain<br>☒ Upgrade<br>From: soam<br>Owner: Fernando Caro x332 | | | |
| Europe | ☒ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | Will migrate Northern Europe and Southern Europe MUDS into this domain. |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |

| | | |
|---|---|---|
| Forest Root Type | ☐ Dedicated ☐ Regional | |
| Forest Root DNS Suffix | | ☐ Existing ☐ New |
| Forest Root DNS Prefix | | |

# Assigning Domain Names

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must assign a name to every domain in your plan. Active Directory domains have two types of names: DNS names and NetBIOS names. In general, both names are visible to end users. The DNS names of Active Directory domains include two parts, a prefix and a suffix. When creating domain names, first determine the DNS prefix. This is the first label in the DNS name of the domain. The suffix is determined when you select the name of the forest root domain. To select the DNS prefix for the domain name:

1. For each Windows NT 4.0 MUD that you choose to upgrade in place, decide whether the current NetBIOS name of the domain will become its DNS prefix. Review prefix naming rules. If the name is appropriate to represent the region and satisfies the prefix naming rules, it is recommended that you retain the current NetBIOS name. In this case, the NetBIOS name of the domain is the same as the DNS prefix of the domain. Table 2.7 lists the prefix naming rules for DNS names.

   **Table 2.7 Rules for Selecting a Prefix for a Registered DNS Name**

   | Rule | Explanation |
   | --- | --- |
   | Select a prefix that is not likely to become outdated. | Avoid names such as a business line or operating system that might change in the future. Geographical names are recommended. |
   | Select a prefix that includes Internet standard characters only. | A-Z, a-z, 0-9, and (-), but not entirely numeric. |
   | Include 15 characters or less in the prefix. | If you choose a prefix length of 15 characters or less, then the NetBIOS name is the same as the prefix. |

   If the current NetBIOS name of the domain is inappropriate to represent the region or fails to satisfy the prefix naming rules, select a new prefix. In this case, the NetBIOS name of the domain is different from the DNS prefix of the domain.

2. For each new domain that you deploy, select a prefix that is appropriate for the region and that satisfies prefix naming rules. It is recommended that the NetBIOS name of the domain be the same as the DNS prefix.

# Documenting Domain Names

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Document the DNS prefix and NetBIOS names that you select for each domain in your forest. You can add the DNS and NetBIOS name information to the worksheet that you created to document your plan for new and upgraded domains.

For a worksheet to assist you in documenting the proposed domain names see "Domain Planning" (DSSLOGI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Planning" on the Web at http://www.microsoft.com/reskit).

**Example: Documenting Domain Names**

Contoso Corporation decided that the DNS prefixes for the North America, South America and European Active Directory domains will be noam, soam, and eur respectively. They also decided to use the prefix as the NetBIOS name.

Figure 2.19 shows the worksheet in which the design team documented the domain names.

**Figure 2.19 Example of a Domain Planning Worksheet That Includes Domain DNS Prefixes and NetBIOS Name**

| Domain Planning | | | |
|---|---|---|---|
| Prepared By | Olinda Turner | Date | 2/20/03 |
| Forest Owner | Olinda Turner | | |

Instructions: Record information about your existing domains and any new domains that you will create. First, list the domains and whether they are new domains or existing domains that will be upgraded. Second, list the DNS prefix and NetBIOS name for each domain. Finally, list the forest root information.

| Region | Origin | Prefix | NetBIOS Name | Justification / Notes |
|---|---|---|---|---|
| North America | ☐ New domain<br>☒ Upgrade<br>From: noam<br>Owner: Mark Lee x874 | noam | noam | |
| South America | ☐ New domain<br>☒ Upgrade<br>From: soam<br>Owner: Fernando Caro x332 | soam | soam | |
| Europe | ☒ New domain<br>☐ Upgrade<br>From:<br>Owner: | eur | eur | Will migrate Northern Europe and Southern Europe MUDS into this domain. |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |

| | | |
|---|---|---|
| Forest Root Type | ☐ Dedicated ☐ Regional | |
| Forest Root DNS Suffix | | ☐ Existing ☐ New |
| Forest Root Prefix | | |

# Selecting the Forest Root Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first domain that you deploy in an Active Directory forest is called the forest root domain. This domain remains the forest root domain for the life cycle of the Active Directory deployment.

The forest root domain contains the Enterprise Admins and Schema Admins groups. These service administrator groups are used to manage forest-level operations, such as the addition and removal of domains and the implementation of changes to the schema.

Selecting the forest root domain involves determining whether one of the Active Directory domains in your domain design can function as the forest root domain, or whether you need to deploy a dedicated forest root domain.

# Choosing a Regional or Dedicated Forest Root Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are applying a single domain model, then the single domain functions as the forest root domain. If you are applying a multiple domain model, then you can choose to deploy a dedicated forest root domain, or select a regional domain to function as the forest root domain.

**Dedicated Forest Root Domain**

A *dedicated forest root* domain is a domain that is created specifically to function as the forest root. It does not contain any user accounts other than the service administrator accounts for the forest root domain, and it does not represent any region in your domain structure. All other domains in the forest are children of the dedicated forest root domain.

Using a dedicated forest root provides the following advantages:

- **Operational separation of forest service administrators from domain service administrators.** In a single domain environment, members of the Domain Admins or built-in Administrators groups can use standard tools and procedures to make themselves members of the Enterprise Admins and Schema Admins groups. In a forest that uses a dedicated forest root domain, members of the Domain Admins or built-in Administrators groups in the regional domains cannot make themselves members of the forest-level service administrator groups by using standard tools and procedures.

  **Warning**

  - Because a domain is not a security boundary, it is possible for a malicious service administrator, such as a member of the Domain Admins group, to use nonstandard tools and procedures to gain full access to any domain in the forest or to any computer in the forest. For example, service administrators in a nonroot domain can make themselves members of the Enterprise Admins or Schema Admins group.

- **Protection from operational changes in other domains.** A dedicated forest root domain does not represent a particular region in your domain structure. For this reason, it is not affected by reorganizations or other changes that result in the renaming or restructuring of domains.

- **Serves as a neutral root so that no region appears to be subordinate to another region.** Some organizations might prefer to avoid the appearance that one country/region is subordinate to another country/region in the namespace. When you use a dedicated forest root domain, all regional domains can be peers in the domain hierarchy.

In a multiple regional domain environment in which a dedicated forest root is used, the replication of the forest root domain has minimal impact on the network infrastructure. This is because the forest root only hosts the service administrator accounts. The majority of the user accounts in the forest and other domain-specific data is stored in the regional domains.

One disadvantage to using a dedicated forest root domain is that it creates additional management overhead to support the additional domain.

**Regional Domain as a Forest Root Domain**

If you choose not to deploy a dedicated forest root domain, then you must select a regional domain to function as the forest root domain. This domain is the parent domain of all the other regional domains and will be the first domain that you deploy. The forest root domain contains user accounts and is managed in the same way that the other regional domains are managed. The primary difference is that it also includes the Enterprise Admins and Schema Admins groups.

The advantage to selecting a regional domain to function as the forest root domain is that it does not create the additional management overhead that maintaining an additional domain creates.

Select an appropriate regional domain to be the forest root, such as the domain that represents your headquarters or the region that has the fastest network connections. If it is difficult for your organization to select a regional domain to be the forest root domain, you can choose to use a dedicated forest root model instead.

In a Windows Server 2003 environment, global availability of the forest root is not as important as it is in Windows 2000 because forest-wide application partitions automatically replicate the forest-wide locator record zone to all domain controllers that are running DNS. Any domain controller can be used to write updates to the forest-wide locator records zone. In a

Windows 2000 environment, DNS does not use a forest-wide application partition; therefore, it is recommended that a dedicated forest root be used to make the zone containing the writable copy of the forest-wide locator records highly available. For more information about DNS and the forest-wide application partition, see "Designing a DNS Infrastructure to Support Active Directory" later in this chapter.

**Forest Root Alternatives With Single Global Domain Model**

This model consists of a forest that contains a dedicated forest root domain and one additional global domain. The global domain contains all the user, computer, and group accounts for the entire organization. Figure 2.20 shows an example of a forest with a single global domain with and without a dedicated forest root.

**Figure 2.20 Single Global Domain Model With and Without Dedicated Root**




**Forest Root Alternatives With Multiple Regional Domains**

This model consists of a dedicated forest root domain or a regional domain that is designated to be the forest root, and multiple regional domains that are children of the forest root, as shown in Figure 2.21.

**Figure 2.21 Multiple Regional Domains With and Without a Dedicated Forest Root**




In the case of multiple regional domains, a third alternative is to make each regional domain a separate tree in a single forest.

While this design is fully supported, it is not recommended because of the complexity of the DNS deployment.

**Figure 2.22 Forest with Multiple Trees**



# Assigning the Forest Root Domain Name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The forest root domain name is also the name of the forest. The forest root name is a DNS name that consists of a prefix and a suffix in the form of *prefix.suffix*. For example, an organization might have the forest root name corp.contoso.com. In this example, corp is the prefix and contoso.com is the suffix.

Select the suffix from a list of existing names on your network. For the prefix, select a new name that has not been used on your network previously. By attaching a new prefix to an existing suffix, you create a unique namespace. Creating a new namespace for Active Directory ensures that any existing DNS infrastructure does not need to be modified to accommodate Active Directory.

**Selecting a Suffix**
To select a suffix for the forest root domain:

1. Contact the DNS owner for the organization for a list of registered DNS suffixes that are in use on the network that will host Active Directory. Note that the suffixes used on the internal network might be different than the suffixes used externally. For example, an organization might use contosopharma.com on the Internet and contoso.com on the internal corporate network.

2. Consult the DNS owner to select a suffix for use with Active Directory. If no suitable suffixes exist, register a new name with an Internet naming authority.

It is best to use DNS names that are registered with an Internet authority in the Active Directory namespace. Only registered names are guaranteed to be globally unique. If another organization later registers the same DNS domain name, or if your organization merges with, acquires, or is acquired by other company that uses the same DNS names, then the two infrastructures cannot interact with one another.

**Note**

- Using single label names or unregistered suffixes, such as .local, is not recommended.

**Selecting a Prefix**
If you chose a registered suffix that is already in use on the network, select a prefix for the forest root domain name by using the prefix rules in Table 2.8. Add a prefix that is not currently in use to create a new subordinate name. For example, if your DNS root name is contoso.com, then you can create the Active Directory forest root domain name concorp.contoso.com, if the namespace concorp.contoso.com is not already in use on the network. This new branch of the namespace will be dedicated to Active Directory and can be integrated easily with the existing DNS implementation. When selecting a prefix, consider the following:

- If you are performing an in-place upgrade of a Windows NT 4.0 MUD to create the forest root domain, then you might be able to use the NetBIOS name of the domain as the prefix. Determine whether the NetBIOS name is an appropriate name for a root domain and meets the prefix rules listed in Table 2.8.

- If you selected a regional domain to function as a forest root domain, you might need to select a new prefix for the domain. Because the forest root domain name affects all of the other domain names in the forest, a regionally-based name might not be appropriate. For example, if Contoso Corporation decided to use their North American domain, called noam.contoso.com, as their forest root, then the European domain name would be europe.noam.contoso.com. In this case, a better choice would be to select a new prefix, such as corp, for the forest root; in this way, the name of the European domain would be europe.corp.contoso.com.

If you are using a new suffix that is not currently in use on the network, you can use it as the forest root domain name without choosing an additional prefix.

Table 2.8 lists the rules for selecting a prefix for a registered DNS name.

**Table 2.8 Rules for Selecting a Prefix for a Registered DNS Name**

| Rule | Explanation |
| --- | --- |
| Select a prefix that is not likely to become outdated. | Avoid names such as a business line or operating system that might change in the future. Generic names such as corp or ds are recommended. |
| Select a prefix that includes Internet standard characters only. | A-Z, a-z, 0-9, and (-), but not entirely numeric. |
| Include 15 characters or less in the prefix. | If you choose a prefix length of 15 characters or less, then the NetBIOS name is the same as the prefix. |

It is important for the Active Directory DNS owner to work with the DNS owner for the organization to obtain ownership of the name that will be used for the Active Directory namespace. For more information about designing a DNS infrastructure to support Active Directory, see "Designing a DNS Infrastructure to Support Active Directory" later in this chapter.

# Documenting the Forest Root Domain Name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Document the DNS prefix and suffix that you select for the forest root domain. At this point identify which domain will be the forest root. You can add the forest root domain name information to the worksheet that you created to document your plan for new and upgraded domains and your domain names.

For a worksheet to assist you in documenting the forest root domain name, see "Domain Planning" (DSSLOGI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD ( or see "Domain Planning" on the Web at https://www.microsoft.com/reskit).

**Example: Documenting the Forest Root Domain Name**
Contoso Corporation decided to use a dedicated forest root domain for their Active Directory deployment. They selected the name corp.contoso.com for their forest. They will use their registered corporate name contoso.com as the suffix for this domain and add the prefix corp to represent this domain specifically. The forest root domain name will therefore be corp.contoso.com. They recorded this information in their Domain Planning worksheet so that the service administrators can reference it as needed when they deploy the forest root domain.

Figure 2.23 shows the worksheet in which the design team documented the forest root domain.

**Figure 2.23 Example of a Domain Planning Worksheet That Includes Forest Root Domain Name Information**

# Domain Planning

| | | | |
|---|---|---|---|
| Prepared By | Olinda Turner | Date | 2/20/03 |
| Forest Owner | Olinda Turner | | |

Instructions: Record information about your existing domains and any new domains that you will create. First, list the domains and whether they are new domains or existing domains that will be upgraded. Second, list the DNS prefix and NetBIOS name for each domain. Finally, list the forest root information.

| Region | Origin | Prefix | NetBIOS Name | Justification / Notes |
|---|---|---|---|---|
| North America | ☐ New domain<br>☒ Upgrade<br>From: noam<br>Owner: Mark Lee x874 | noam | noam | |
| South America | ☐ New domain<br>☒ Upgrade<br>From: soam<br>Owner: Fernando Caro x332 | soam | soam | |
| Europe | ☒ New domain<br>☐ Upgrade<br>From:<br>Owner: | eur | eur | Will migrate Northern Europe and Southern Europe MUDS into this domain. |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |
| | ☐ New domain<br>☐ Upgrade<br>From:<br>Owner: | | | |

| | | |
|---|---|---|
| Forest Root Type | ☒ Dedicated ☐ Regional | |
| Forest Root DNS Suffix | contoso.com | ☒ Existing ☐ New |
| Forest Root Prefix | corp | |

# Creating a Consolidation Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are currently operating in a Windows NT 4.0 environment, you must decide how to restructure the remaining MUDs that you are not upgrading in place and your existing resource domains into the Active Directory domains that you have designed. Because you can use OUs in an Active Directory environment to delegate control, MUDs and resource domains are no longer needed for the delegation of control to administrators. As part of your domain design, decide how you will restructure your MUDs and resource domains into the new Windows Server 2003 environment.

# Restructuring Windows NT 4.0 MUDs into Windows Server 2003 Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization is currently operating in a Windows NT 4.0 environment, you must plan to restructure all MUDs that you did not select for in-place upgrade into your new Active Directory domains. Keep the following points in mind when you consolidate a MUD into an Active Directory domain:

- You can configure OUs to preserve the organization and administrative structure of the Windows NT 4.0 source domain. For more information about configuring OUs, see "Designing Organizational Units for Delegation of Administration" later in this chapter.
- The migration process potentially impacts users involved in the migration.
- If you have a MUD that spans geographic locations, migrate all the users within that MUD to a single Active Directory domain. If you need to relocate some of those users to a different domain, you can do so by means of an intraforest migration after the initial migration process is complete.

For more information about restructuring Windows NT 4.0 account and resource domains to Windows Server 2003 Active Directory, see "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.

Figure 2.24 shows how the Contoso Corporation chose to restructure Windows NT 4.0 resource domains into OU structures within the new Active Directory regional domains in the same geographical location.

**Figure 2.24 Contoso Corporation Restructuring of Windows NT 4.0 MUDs into Active Directory Regional Domains**



# Documenting the Migration Plan for Windows NT 4.0 Master User Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you identify all remaining Windows NT 4.0 Master User Domains and determine the OUs in the Active Directory domains to which the accounts in these domains will be migrated, you need to document your migration plan.

For a worksheet to assist you in documenting your migration plan for Windows NT 4.0 MUDs, see "Migration Plan for Windows NT 4.0 Master User Domains" (DSSLOGI_6.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migration Plan for Windows NT 4.0 Master User Domains" on the Web at http://www.microsoft.com/reskit).

**Example: Documenting the Migration Plan for Windows NT 4.0 Master User Domains**
Contoso Corporation decided to upgrade the noam and soam domains to Active Directory domains. They chose to migrate the remaining MUDs in Europe into the new European domain eur.corp.contoso.com. They documented their migration plan in a worksheet.

Figure 2.25 shows the Windows NT 4.0 MUD migration plan documentation that the Contoso Corporation design team created.

**Figure 2.25 Example of a Migration Plan for Windows NT 4.0 Master User Domains Worksheet**

Migration Plan for Windows NT 4.0 Master User Domains

| | | | |
|---|---|---|---|
| Prepared By | Olinda Turner | Date | 2/28/03 |
| Forest Name | corp.contoso.com | | |
| Forest Owner | Olinda Turner | | |

Instructions: List the remaining Windows NT 4.0 master user domains (MUDs) that will be migrated to OUs, the target OU for each MUD, and the appropriate MUD and OU owner names.

| Active Directory Domain | MUD Name | MUD Owner | OU Name | OU Owner | Justification / Notes |
|---|---|---|---|---|---|
| eur | NorthEur | Roger Lengel x551 | Europe North | Roger Lengel x551 | Consolidate into the eur domain. |
| | SouthEur | Stefan Delmarco x996 | Europe South | Stefan Delmarco x996 | Consolidate into the eur domain. |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |

# Restructuring Windows NT 4.0 Resource Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can restructure your Windows NT 4.0 resource domains into the Active Directory regional domain that is in the same geographic location, so that resources and the users who need them are in the same location. Keep in mind that when you restructure Windows NT 4.0 resource domains to Windows Server 2003 Active Directory regional domains, you can configure OUs to preserve the organization and administrative structure of the resource domain.

Alternatively, you can maintain the Windows NT 4.0 resource domain until you restructure the domain at a later point or eliminate it by attrition. You might choose to preserve a Windows NT 4.0 resource domain if you are running applications on domain controllers that are not compatible with Windows Server 2003.

**Note**

- If your resource domain is dedicated to hosting Exchange 5.5, do not restructure it into a Windows Server 2003 Active Directory regional domain. For information about upgrading your Microsoft Exchange environment, see the Exchange Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Figure 2.26 shows how Contoso Corporation chose to restructure Windows NT 4.0 resource domains into OU structures within the new Active Directory regional domains in the same geographical location.

**Figure 2.26 Contoso Corporation Restructuring of Windows NT 4.0 Resource Domains into OUs in Active Directory Regional Domains**



# Documenting the Migration Plan for Windows NT 4.0 Resource Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you identify all existing Windows NT 4.0 resource domains and determine how you will migrate the domains to OUs in Active Directory regional domains in the forest, it is important to document this information.

For a worksheet to assist you documenting your plan for restructuring Windows NT 4.0 resource domains, see "Migration Plan for Windows NT 4.0 Resource Domains" (DSSLOGI_7.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migration Plan for Windows NT 4.0 Resource Domains" on the Web at http://www.microsoft.com/reskit).

**Example: Documenting the Migration Plan for Windows NT 4.0 Resource Domains**
Contoso Corporation needs to migrate their Windows NT 4.0 resource domains to OUs in their Active Directory regional domains. After they identified the source and target locations, they documented the target OU for each resource domain that is to be restructured. Figure 2.27 shows the Windows NT 4.0 resource domain migration plan documentation that the Contoso Corporation design team created.

**Figure 2.27 Example of a Migration Plan for Windows NT 4.0 Resource Domains Worksheet**

Migration Plan for Windows NT 4.0 Resource Domains

| | | | |
|---|---|---|---|
| Prepared By | Olinda Turner | Date | 2/28/03 |
| Forest Name | corp.contoso.com | | |
| Forest Owner | Olinda Turner | | |

Instructions: List the remaining Windows NT 4.0 resource domains that will be migrated to OUs, the target OU for each resource domain, and the appropriate resource domain and OU owner names.

| Active Directory Domain | Resource Domain Name | Resource Domain Owner | OU Name | OU Owner | Justification / Notes |
|---|---|---|---|---|---|
| noam | NoamMfg | Paul West X804 | Noam Manufacturing | Paul West x804 | |
| | Research | Mark Lee X874 | Research | Mark Lee x874 | |
| | HR | Mary Baker x989 | Human Resource | Mary Baker x989 | |
| soam | Mfg | Fernando Caro X332 | Soam Manufacturing | Fernando Caro x332 | |
| | Shipping | Fernando Caro X332 | Shipping | Fernando Caro x332 | |
| eur | EurMfg | Roger Lengel x551 | Europe Manufacturing | Roger Lengel x551 | |
| | Finance | Stefan Delmarco x996 | Finance | Stefan Delmarco x996 | |
| | Transport | Stefan Delmarco x996 | Europe Transportation | Stefan Delmarco x996 | |
| | | | | | |
| | | | | | |
| | | | | | |

# Designing a DNS Infrastructure to Support Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

After you create your Active Directory forest and domain designs, you must design a DNS infrastructure to support your Active Directory logical structure. DNS enables users to use friendly names that are easy to remember to connect to computers and other resources on IP networks.

Windows Server 2003 uses DNS for name resolution instead of the Windows Internet Name Service (WINS) NetBIOS name resolution method used in Windows NT 4.0–based networks. It is still possible to use WINS for applications that require it; however, Active Directory requires DNS.

The process for designing DNS to support Active Directory varies according to whether your organization already has an existing DNS service or whether you are deploying a new DNS service. Figure 2.28 shows the process for creating a DNS design for Active Directory.

**Figure 2.28 Designing a DNS Infrastructure to Support Active Directory**




Before you design your DNS infrastructure to support Active Directory, it can be helpful to read about the DNS hierarchy, the DNS name resolution process, and how DNS supports Active Directory. For information about the DNS hierarchy and name resolution process, see the DNS Technical Reference (http://go.microsoft.com/fwlink/?LinkID=48145). For information about how DNS supports Active Directory, see the DNS Support for Active Directory Technical Reference (http://go.microsoft.com/fwlink/?LinkID=48147).

# Basic DNS Concepts

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DNS is a distributed database that represents a *namespace*. The namespace contains all of the information needed for any client to look up any name. Any DNS server can answer queries about any name within its namespace. A DNS server answers queries in one of the following ways:

- If the answer is in its cache, it answers the query from the cache.
- If the answer is in a zone hosted by the DNS server, it answers the query from its zone. A *zone* is a portion of the DNS tree stored on a DNS server. When a DNS server hosts a zone, it is *authoritative* for the names in that zone, which means that the DNS server can answer queries for any name in the zone. For example, a server hosting the zone contoso.com can answer queries for any name in contoso.com.
- If the server cannot answer the query from its cache or zones, it queries other servers for the answer.

It is important to understand the core features of DNS, such as delegation, recursive name resolution, and Active Directory–integrated zones, because they have a direct impact on your Active Directory logical structure design.

For more information about DNS and Active Directory, see "DNS and Active Directory" later in this chapter.

# Delegation

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For a DNS server to be able to answer queries about any name, it must have a direct or indirect path to every zone in the namespace. These paths are created by means of *delegations*. A delegation is a record in a parent zone that lists a name server that is authoritative for the zone in the next level of the hierarchy. Delegations enable servers in one zone to refer clients to servers in other zones. Figure 2.29 shows an example of delegation.

**Figure 2.29 Delegation**




The DNS root server hosts the root zone, represented as a dot ( . ). The root zone contains a delegation to a zone in the next level of the hierarchy, the com zone. The delegation in the root zone tells the DNS root server that, to find the com zone, it must contact the Com server. Likewise, the delegation in the com zone tells the Com server that to find the contoso.com zone, it must contact the Contoso server.

**Note**

- A delegation uses two types of records. The name server (NS) resource record provides the name of an authoritative server. The host (A) resource provides the IP address of an authoritative server.

This system of zones and delegations creates a hierarchical tree that represents the DNS namespace. Each zone represents a layer in the hierarchy and each delegation represents a branch of the tree.

By using the hierarchy of zones and delegations, a DNS root server can find any name in the DNS namespace. The root zone includes delegations that lead directly or indirectly to all other zones in the hierarchy. Any server that can query the DNS root server can use the information in the delegations to find any name in the namespace.

# Recursive Name Resolution

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Recursive name resolution is the process by which a DNS server uses the hierarchy of zones and delegations to respond to queries for which it is not authoritative.

In some configurations, DNS servers include *root hints,* a list of names and IP addresses, which enable them to query the DNS root servers. In other configurations, servers forward all queries that they cannot answer to another server. Forwarding and root hints are both methods that DNS servers can use to resolve queries for which they are not authoritative.

**Resolving Names by Using Root Hints**

Root hints enable any DNS server to locate the DNS root servers. After a DNS server locates the DNS root server, it can resolve any query for that namespace. Figure 2.30 shows how DNS resolves a name by using root hints.

**Figure 2.30 Using Root Hints for Recursive Name Resolution**




In this example, the following events take place:

1. A client sends a recursive query to a DNS server to request the IP address that corresponds to the name ftp.contoso.com. A *recursive query* indicates that the client wants a definitive answer to its query. The response to the recursive query must be a valid address or a message indicating that the address cannot be found.

2. Because the DNS server is not authoritative for the name and does not have the answer in its cache, the DNS server uses root hints to find the IP address of the DNS root server.

3. The DNS server uses an iterative query to ask the DNS root server to resolve the name ftp.contoso.com. An *iterative query* indicates that the server will accept a referral to another server in place of a definitive answer to the query. Because the name ftp.contoso.com ends with the label com, the DNS root server returns a referral to the com server that hosts the com zone.

4. The DNS server uses an iterative query to ask the com server to resolve the name ftp.contoso.com. Because the name ftp.contoso.com ends with the name contoso.com, the com server returns a referral to the Contoso server that hosts the contoso.com zone.

5. The DNS server uses an iterative query to ask the Contoso server to resolve the name ftp.contoso.com. The Contoso server finds the answer in its zone data and returns the answer to the server.

6. The server then returns the result to the client.

**Resolving Names by Using Forwarding**

Forwarding enables you to route name resolution through specific servers instead of using root hints. Figure 2.31 shows how DNS resolves a name by using forwarding.

**Figure 2.31 Using Forwarding for Recursive Name Resolution**




In this example, the following events take place:

1. A client queries a DNS server for the name ftp.contoso.com.

2. The DNS server forwards the query to another DNS server, known as a *forwarder*.

3. Because the forwarder is not authoritative for the name and does not have the answer in its cache, it uses root hints to find the IP address of the DNS root server.

4. The forwarder uses an iterative query to ask the DNS root server to resolve the name ftp.contoso.com. Because the name ftp.contoso.com ends with the name com, the DNS root server returns a referral to the com server that hosts the com zone.

5. The forwarder uses an iterative query to ask the com server to resolve the name ftp.contoso.com. Because the name ftp.contoso.com ends with the name contoso.com, the com server returns a referral to the Contoso server that hosts the contoso.com zone.

6. The forwarder uses an iterative query to ask the Contoso server to resolve the name ftp.contoso.com. The Contoso server finds the answer in its zone files and returns the answer to the server.

7. The forwarder then returns the result to the original DNS server.

8. The original DNS server then returns the result to the client.

For more information about DNS, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

# DNS and Active Directory

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory uses Domain Name System (DNS) name resolution services to make it possible for clients to locate domain controllers and to make it possible for the domain controllers that host the directory service to communicate with each other.

Active Directory enables easy integration of the Active Directory namespace into an existing DNS namespace. Features such as Active Directory–integrated zones make it easier for you to deploy DNS by eliminating the need to set up secondary zones and configure zone transfers.

For information about how DNS supports Active Directory, see the DNS Support for Active Directory Technical Reference (http://go.microsoft.com/fwlink/?LinkID=48147) as well as "Deploying DNS" in Deploying Network Services in this kit.

| **Note** |
| --- |
| If you implement a disjoint namespace in which the Active Directory domain name differs from the primary DNS suffix that clients use, Active Directory integration with DNS is more complex. For more information, see Disjoint Namespace. |

# Domain Controller Location

Updated: October 19, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Clients use DNS to locate domain controllers to complete operations such as processing logon requests or searching the directory for published resources. Domain controllers register a variety of records in DNS to help clients and other computers locate them. These records are collectively referred to as the *locator records*.

Domain controllers also use DNS to locate other domain controllers and perform tasks such as replication. The process by which domain controllers locate other domain controllers is the same as the process by which clients locate domain controllers.

# Active Directory-Integrated Zones

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DNS servers running on domain controllers can store their zones in Active Directory. In this way, it is not necessary to configure a separate DNS replication topology that uses ordinary DNS zone transfers, because all zone data is replicated automatically by means of Active Directory replication. This simplifies the process of deploying DNS and provides the following advantages:

- Multiple masters are created for DNS replication. Therefore:
    - Any domain controller in the domain running the DNS server service can write updates to the Active Directory–integrated zones for the domain name for which they are authoritative. A separate DNS zone transfer topology is not needed.
- Secure dynamic updates are supported. Secure dynamic updates allow an administrator to control which computers update which names, and prevent unauthorized computers from overwriting existing names in DNS.

Windows Server 2003 DNS Active Directory stores zone data in application directory partitions. The domain partition was the only Active Directory storage option in Windows 2000 Server, and it is available in Windows Server 2003 DNS for backward compatibility. The following DNS-specific application directory partitions are created during Active Directory installation:

- A forest-wide application directory partition, called ForestDnsZones.
- Domain-wide application directory partitions for each domain in the forest, named DomainDnsZones.

For more information about how Active Directory stores DNS information in application partitions, see the DNS Technical Reference (http://go.microsoft.com/fwlink/?LinkID=48145) and Windows Server 2003 Help and Support.

# Computer Naming

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a computer running Windows 2000, Microsoft® Windows® XP, or Windows Server 2003 joins a domain, by default the computer assigns itself a name comprising the host name of the computer (Computer Name in the System Properties) and the DNS name of the Active Directory domain that the computer joined (Primary DNS Suffix in System Properties). The concatenation of the host name and DNS name of the domain is known as the fully qualified domain name (FQDN). For example, if a computer with host name Server1 joins the domain corp.contoso.com, the FQDN of the computer is server1.corp.contoso.com.

If a computer already has a different DNS domain name that was statically entered into a DNS zone or registered by an integrated DNS/DHCP service, the FQDN of the computer is distinct from the name that was registered previously. The computer can be referred to by either name.

# Disjoint Namespace

Updated: August 17, 2011

Applies To: Windows Server 2003, Windows Server 2008, Windows Server 2008 R2

A disjoint namespace occurs when one or more domain member computers have a primary Domain Name Service (DNS) suffix that does not match the DNS name of the Active Directory domain of which the computers are members. For example, a member computer that uses a primary DNS suffix of corp.fabrikam.com in an Active Directory domain named na.corp.fabrikam.com is using a disjoint namespace.

A disjoint namespace is more complex to administer, maintain, and troubleshoot than a contiguous namespace. In a contiguous namespace, the primary DNS suffix matches the Active Directory domain name. Network applications that are written to assume that the Active Directory namespace is identical to the primary DNS suffix for all domain member computers do not function properly in a disjoint namespace.

**Support for disjoint namespaces**

Domain member computers, including domain controllers, can function in a disjoint namespace. Domain member computers can register their host (A) resource record and IP version 6 (IPv6) host (AAAA) resource record in a disjoint DNS namespace. When domain member computers register their resource records in this way, domain controllers continue to register global and site-specific service (SRV) resource records in the DNS zone that is identical to the Active Directory domain name.

For example, assume that a domain controller for the Active Directory domain named na.corp.fabrikam.com that uses a primary DNS suffix of corp.fabrikam.com registers host (A) and IPv6 host (AAAA) resource records in the corp.fabrikam.com DNS zone. The domain controller continues to register global and site-specific service (SRV) resource records in the _msdcs.na.corp.fabrikam.com and na.corp.fabrikam.com DNS zones, which makes service location possible.

| **Important** |
| --- |
| Although Windows operating systems may support a disjoint namespace, applications that are written to assume that the primary DNS suffix is the same as the Active Directory domain suffix may not function in such an environment. For this reason, you should test all applications and their respective operating systems carefully before you deploy a disjoint namespace. |

A disjoint namespace should work (and is supported) in the following situations:

- When a forest with multiple Active Directory domains uses a single DNS namespace, which is also known as a DNS zone

  An example of this is a company that uses regional domains with names such as na.corp.fabrikam.com, sa.corp.fabrikam.com, and asia.corp.fabrikam.com and uses a single DNS namespace, such as corp.fabrikam.com.

- When a single Active Directory domain is split into separate DNS namespaces

  An example of this is a company with an Active Directory domain of corp.contoso.com that uses DNS zones such as hr.corp.contoso.com, production.corp.contoso.com, and it.corp.contoso.com.

A disjoint namespace does not work properly (and is not supported) in the following situations:

- A disjoint suffix used by domain members matches an Active Directory domain name in this or another forest. This breaks Kerberos name-suffix routing.

- The same disjoint suffix is used in another forest. This prevents routing these suffixes uniquely between forests.

- When a domain member certification authority (CA) server changes its fully qualified domain name (FQDN) so that it no longer use the same primary DNS suffix that is used by the domain controllers of the domain to which the CA server is a member. In this case, you may have problems validating certificates the CA server issued, depending on what DNS names are used in the CRL Distribution Points. But if you place a CA server in a stable disjoint namespace, it works properly and is supported.

**Considerations for disjoint namespaces**

The following considerations may help you decide if you should use a disjoint namespace.

**Application compatibility**
As previously mentioned, a disjoint namespace can cause problems for any applications and services that are written to assume that a computer primary DNS suffix is identical to the name of the domain name of which it is a member. Before you deploy a disjoint namespace, you must check applications for compatibility issues. Also, be sure to check the compatibility of all applications that you use when you perform your analysis. This includes applications from Microsoft and from other software developers.

**Advantages of disjoint namespaces**
Using a disjoint namespace can have the following advantages:

- Because the primary DNS suffix of a computer can indicate different information, you can manage the DNS namespace separately from the Active Directory domain name.

- You can separate the DNS namespace based on business structure or geographical location. For example, you can separate the namespace based on business unit names or physical location such as continent, country/region, or building.

**Disadvantages of disjoint namespaces**
Using a disjoint namespace can have the following disadvantages:

- You must create and manage separate DNS zones for each Active Directory domain in the forest that has member computers that use a disjoint namespace. (That is, it requires an additional and more complex configuration.)

- You must perform manual steps to modify and manage the Active Directory attribute that allows domain members to use specified, primary DNS suffixes.

- To optimize name resolution, you must perform manual steps to modify and maintain Group Policy to configure member computers with alternate primary DNS suffixes.

  | **Note** |
  | --- |
  | The Windows Internet Name Service (WINS) could be used to offset this disadvantage by resolving single-label names. For more information about WINS, see the WINS Technical Reference (http://go.microsoft.com/fwlink/?LinkId=102303). |

- When your environment requires multiple primary DNS suffixes, you must configure the DNS suffix search order for all of the Active Directory domains in the forest appropriately.

  To set the DNS suffix search order, you can use Group Policy objects or Dynamic Host Configuration Protocol (DHCP) Server service parameters. You can also modify the registry.

- You must carefully test all applications for compatibility issues.

For more information about steps that you can take to address these disadvantages, see Create a Disjoint Namespace (http://go.microsoft.com/fwlink/?LinkId=106638).

**Planning a namespace transition**
Before you modify a namespace, review the following considerations, which apply to transitions from contiguous namespaces to disjoint namespaces (or the reverse):

- Manually configured Service Principal Names (SPNs) may no longer match DNS names after a namespace change. This can cause authentication failures.

  For more information, see Service Logons Fail Due to Incorrectly Set SPNs (http://go.microsoft.com/fwlink/?LinkId=102304).

  - If you use Windows Server 2003–based computers with constrained delegation, those computers may require additional configuration to change SPNs. For more information, see article 936628 in the Microsoft Knowledge

Base (http://go.microsoft.com/fwlink/?LinkId=102306).

- If you want to delegate permissions to modify SPNs to subordinate administrators, see Delegating Authority to Modify SPNs (http://go.microsoft.com/fwlink/?LinkId=106639).

- If you use Lightweight Directory Access Protocol (LDAP) over Secure Sockets Layer (SSL) (known as LDAPS) with a CA in a deployment that has domain controllers that are configured in a disjoint namespace, you must use the appropriate Active Directory domain name and primary DNS suffix when you configure the LDAPS certificates.

  For more information about domain controller certificate requirements, see article 321051 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=102307).

| **Note** |
| --- |
| Domain controllers that use certificates for LDAPS may require you to redeploy their certificates. When you do so, domain controllers may not select an appropriate certificate until they are restarted. For more information about LDAPS authentication and a related update for Windows Server 2003, see article 932834 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=102308). |

**Planning for disjoint namespace deployments**

Take the following precautions if you deploy computers in an environment that has a disjoint namespace:

1. Notify all software vendors with whom you do business that they must test and support a disjoint namespace. Ask them to verify that they support their applications in environments that use disjoint namespaces.

2. Test all versions of operating systems and applications in disjoint namespace lab environments. When you do, follow these recommendations:

   a. Resolve all software issues before you deploy the software into your environment.

   b. When possible, participate in beta tests of operating systems and applications that you plan to deploy in disjoint namespaces.

3. Ensure that administrators and helpdesk staff are aware of the disjoint namespace and its impact.

4. Create a plan that makes it possible for you to transition from a disjoint namespace to a contiguous namespace, if necessary.

5. Evangelize the importance of disjoint namespace support with operating system and application providers.

# Assigning the Active Directory DNS Owner Role

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The forest owner assigns an Active Directory DNS owner for the forest. The Active Directory DNS owner for the forest is a person or group of people responsible for overseeing the deployment of the Active Directory DNS infrastructure and making sure that, if necessary, domain names are registered with the proper Internet authorities.

The Active Directory DNS owner is responsible for the Active Directory DNS design for the forest. If your organization is currently operating a DNS service, then the DNS designer for the existing DNS service works with the Active Directory DNS owner to delegate the forest root DNS name to DNS servers running on domain controllers.

The Active Directory DNS owner for the forest also maintains contact with the DHCP and DNS teams of the organization and coordinates the plans of the individual DNS owners of each domain in the forest (if any) with these groups. The DNS owner for the forest ensures that the DHCP and DNS teams are involved in the Active Directory DNS design process so that each group is aware of the DNS design plan and can provide input early on.

# Identifying the DNS Infrastructure Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you deploy Active Directory in your environment, you must have a DNS infrastructure available to support the directory service. If you already have an existing DNS infrastructure, you need to integrate the Active Directory namespace into that environment. If you do not have a DNS infrastructure in place, you need to design and deploy a new DNS infrastructure to support Active Directory.

For the purpose of this design process, an organization does not have an existing DNS infrastructure if the following are true:

- The organization does not have any existing DNS servers in the network infrastructure.
- The organization does not have any clients that access DNS servers. This means that the organization does not rely on an external source, such as a network service provider, for DNS services.

For more information about designing a new DNS infrastructure, see the Domain Name System (DNS) Center link on the Web Resources page.

# Integrating Active Directory into an Existing DNS Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization already has an existing DNS service, the Active Directory DNS owner must work with the DNS owner for your organization to integrate Active Directory into the existing infrastructure. This involves creating a DNS server and DNS client configuration.

# Creating a DNS Server Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When integrating Active Directory with an existing DNS namespace, it is recommended that you do the following:

- Install the DNS Server service on every domain controller in the forest. This provides fault tolerance in the event that one of the DNS servers is unavailable. In this way, domain controllers do not need to rely on other DNS servers for name resolution. This also simplifies the management environment because all domain controllers have a uniform configuration.
- Configure the Active Directory forest root domain controller to host the DNS zone for the Active Directory forest.
- Configure the domain controllers for each regional domain to host the DNS zones that correspond to their Active Directory domains.
- Configure the zone containing the Active Directory forest-wide locator records (_msdcs.*forestname* zone) to replicate to every DNS server in the forest by using the forest-wide DNS application directory partition.

**Note**

- Active Directory uses forest-wide locator records to enable replication partners find each other and to enable clients to find global catalog servers. Active Directory stores the forest-wide locator records in the _msdcs.*forest_name* zone. Because the information in the zone must be widely available, this zone is replicated to all DNS servers in the forest by means of the forest-wide DNS application directory partition.

The existing DNS structure remains intact. You do not need to migrate any servers or zones. You simply need to create a delegation to your Active Directory zones from your existing DNS hierarchy.

**Configuration When Using Root Hints**

Figure 2.32 illustrates the server configuration model for a network that uses root hints. The DNS root zone is hosted on a DNS server at the root of the existing internal DNS infrastructure.

**Figure 2.32 Existing DNS Service Configured to Use Root Hints**




In this configuration, the DNS servers are configured with root hints that indicate the location of the DNS root servers. When a DNS server needs to resolve a query and cannot resolve it from its cache or zone data, the DNS server queries the DNS servers

listed in the root hints.

**Configuration When Using Forwarding**

Figure 2.33 illustrates the server configuration model for a network that uses forwarding.

**Figure 2.33 Existing DNS Service Configured to Use Forwarding**




In this configuration, the regional DNS servers that are used to support Active Directory forward queries to the nearest DNS server for the Active Directory forest root, which in turn forwards queries to the standard DNS forwarder that is used by the existing DNS infrastructure.

# Creating the DNS Client Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure DNS on client computers, the Active Directory DNS owner must specify the computer naming scheme and how the clients will locate DNS servers. Table 2.9 lists the recommended configurations for these design elements.

**Table 2.9 DNS Configuration for Client Computers**

| Design element | Configuration |
|---|---|
| Computer naming | Use default naming. When a Windows 2000, Windows XP, or Windows Server 2003 computer joins a domain, the computer assigns itself a fully qualified domain name (FQDN) comprising the host name of the computer and the name of the Active Directory domain. |
| Client Resolver Configuration | Configure client computers to point to any DNS server on the network. |

**Note**

- Active Directory clients and domain controllers can dynamically register their DNS names even if they are not pointing to the DNS server that is authoritative for their name.

A computer might have a different existing DNS name if the organization previously statically registered the computer in DNS or the organization previously deployed an integrated DHCP solution. If your client computers already have a registered DNS name, when the domain to which they are joined is upgraded to Windows Server 2003 Active Directory, they will have two different names:

- The existing DNS name
- The new FQDN

Clients can still be located by either name. Any existing DNS, DHCP, or integrated DNS/DHCP solution is left intact. The new primary names are created automatically and updated by means of dynamic update. They are cleaned up automatically by means of scavenging.

If you want to take advantage of Kerberos authentication when connecting to a server running Windows 2000, Windows XP, or Windows Server 2003, you must make sure that the client connects to the server by using the primary name.

# Documenting Your DNS Infrastructure Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization has an existing DNS infrastructure, you need to make sure that you understand how your DNS infrastructure will interact with the Active Directory namespace.

For a worksheet to assist you in documenting your existing DNS infrastructure design, see "DNS Inventory" (DSSLOGI_8.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DNS Inventory" on the Web at http://www.microsoft.com/reskit).

Figure 2.34 shows the DNS infrastructure design documentation that Contoso Corporation created.

**Figure 2.34 Example of a DNS Inventory Worksheet**

| DNS Inventory | | | |
|---|---|---|---|
| Prepared By | Mike Danseglio | Date | 2/20/03 |
| Forest Name | corp.contoso.com | | |
| Forest Owner | Olinda Turner | | |
| Active Directory DNS Owner | Mike Danseglio | | |
| Instructions: Complete this worksheet if you have existing DNS service on your network. | | | |

| DNS Information | | | |
|---|---|---|---|
| Item | Name of Owner | Name of Team | Contact Info |
| DNS Owner | Lori Oviatt | IT Network Infrastructure | lori@contoso.com x633 |
| DHCP Owner | Mandy Vance | IT Network Infrastructure | mandy@contoso.com x992 |

| Recursive Name Resolution Method | |
|---|---|
| Root Hints | |
| Name | IP Address |
| contosoDNS1 | 131.107.7.23 |
| contosoDNS2 | 131.107.7.98 |
| | |
| | |

| Forwarding | | |
|---|---|---|
| Name | IP Address | Physical Location |
| | | |
| | | |
| | | |
| | | |

| Notes |
|---|
| |

# Designing Organizational Units for Delegation of Administration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Forest owners are responsible for creating organizational unit designs for their domains. Creating an OU design involves designing the OU structure, assigning the OU owner role, and creating account and resource OUs. Initially, design your OU structure to enable delegation of administration. When the OU design is complete, you can create additional OU structures for the application of Group Policy to the users and computers and to limit the visibility of objects. Figure 2.35 shows the process for creating an organizational unit design.

**Figure 2.35 Process for Creating an Organizational Unit Design**



# Reviewing Organizational Unit Design Concepts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The OU structure for a domain includes the following:

- A diagram of the OU hierarchy.
- A list of OUs.
- For each OU:
  - The purpose for the OU.
  - A list of users or groups that have control over the OU or the objects in the OU.
  - The type of control that users and groups have over the objects in the OU.

The OU hierarchy does not need to reflect the departmental hierarchy of the organization or group. OUs are created for a specific purpose, such as the delegation of administration, the application of Group Policy, or to limit the visibility of objects.

You can design your OU structure to delegate administration to individuals or groups within your organization that require the autonomy to manage their own resources and data. OUs represent administrative boundaries and enable you to control the scope of authority of data administrators.

For example, you can create an OU called ResourceOU and use it to store all the computer accounts that belong to the file and print servers managed by a group. Then you can configure security on the OU such that only data administrators in the group have access to the OU. This prevents data administrators in other groups from tampering with the file and print server accounts. Figure 2.36 shows an organizational unit that is created to enable delegation of administration.

**Figure 2.36 Creating an Organizational Unit to Delegate Administration**




You can further refine your OU structure by creating subtrees of OUs for specific purposes, such as the application of Group Policy, or to limit the visibility of protected objects so that only certain users can see them. For example, if you need to apply Group Policy to a select group of users or resources, you can add those users or resources to an OU and then apply the Group Policy to that OU. You can also use the OU hierarchy to enable further delegation of administrative control.

Figure 2.37 shows a subtree that was created inside the ResourceOU. The subtree includes two additional OUs: The Print Servers OU includes all the computer accounts of the print servers; the File Servers OU includes the computer accounts for all of the file servers. This subtree enables the application of separate Group Policies to each type of server managed in the ResourceOU.

**Figure 2.37 Creating an Organizational Unit Subtree for Application of Group Policy**




**Note**

- While there is no technical limit to the number of levels in your OU structure, for the purpose of manageability, it is recommended that you limit your OU structure to a depth of no more than 10 levels. There is no technical limit to the number of OUs on each level. Note that Active Directory–enabled applications might have restrictions on the number of characters used in the distinguished name (the full LDAP path to the object in the directory), or on the OU depth within the hierarchy.

The Active Directory organizational unit structure is not intended to be visible to end users. The organizational unit structure is an administrative tool for service and data administrators and is easy to change. Continue to review and update your OU structure design to reflect changes in your administrative structure and to support policy-based administration.

# Organizational Unit Owner Role

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The forest owner designates an OU owner for each OU that you design for the domain. OU owners are data managers who control a subtree of objects in Active Directory. OU owners can control how administration is delegated, and how policy is applied to objects within their OU. They can also create new subtrees and delegate administration of OUs within that subtree.

Because OU owners do not own or control the operation of the directory service, you can separate ownership and administration of the directory service from ownership and administration of objects, thereby reducing the number of service administrators who have high levels of access.

OUs provide administrative autonomy and the means to control visibility of objects in the directory. OUs provide isolation from other data administrators but they do not provide isolation from service administrators. Although OU owners have control over a subtree of objects, the forest owner retains full control over all subtrees. This enables the forest owner to correct mistakes, such as an error in an access control list (ACL), or to reclaim delegated subtrees when data administrators are terminated

# Delegating Administration by Using OU Objects

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use organizational units to delegate the administration of objects, such as users or computers, within the OU to a designated individual or group. To delegate administration by using an OU, place the individual or group to which you are delegating administrative rights into a group, place the set of objects to be controlled into an OU, and then delegate administrative tasks for the OU to that group.

Active Directory enables you to control the administrative tasks that can be delegated at a very detailed level; for example, you can assign one group full control of all objects in an OU; assign another group the rights only to create, delete, and manage user accounts in the OU; and assign a third group the right only to reset user account passwords. You can make these permissions inheritable so that they apply to not only a single OU, but also any OUs that are placed in subtrees of the OU.

Default OUs and containers are created during the installation of Active Directory and are controlled by service administrators. It is best if service administrators continue to control these containers. If you need to delegate control over objects in the directory, create additional OUs and place the objects in these OUs. Delegate control over these OUs to the appropriate data administrators. This makes it possible to delegate control over objects in the directory without changing the default control given to the service administrators.

The forest owner determines the level of authority that is delegated to an OU owner. This can range from the ability to create and manipulate objects within the OU to only being allowed to control a single attribute of a single type of object in the OU. Granting a user the ability to create an object in the OU implicitly grants that user the ability to manipulate any attribute of any object that the user creates. In addition, if the object that is created is a container, then the user implicitly has the ability to create and manipulate any objects that are placed in the container.

# Administration of Default Containers and OUs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Every Active Directory domain contains a standard set of OUs and containers that are created during the installation of Active Directory. These include the following:

- Domain container, which serves as the root container to the hierarchy.
- Builtin container, which holds the default service administrator accounts.
- Users container, which is the default location for new user accounts and groups created in the domain.
- Computers container, which is the default location for new computer accounts created in the domain.
- Domain Controllers OU, which is the default location for the computer accounts for domain controllers computer accounts.

The forest owner controls these default OUs and containers.

Figure 2.38 shows the default containers and OUs in a standard Active Directory domain. For more information about resource OUs, see "Creating Resource OUs" later in this chapter.

**Figure 2.38 Default Containers and OUs in an Active Directory Domain**

**Domain Container**

The domain container is the root container of the hierarchy of a domain. Changes to the policies or the ACL on this container can potentially have domain-wide impact. Do not delegate control of this container; it must be controlled by the service administrators.

**Users and Computers Containers**

When you perform an in-place domain upgrade from Windows NT 4.0, existing users and computers are automatically placed into the Users and Computers containers. If you are creating a new Active Directory domain, the Users and Computers containers are the default locations for all new user accounts and non-domain controller computer accounts in the domain.

If you need to delegate control over users or computers, do not modify the default settings on the Users and Computers containers. Instead, create new OUs as needed and move the user and computer objects from their default containers and into the new OUs. Delegate control over the new OUs as needed. It is not recommended that you modify who controls the default containers.

Also, you cannot apply Group Policy settings to the default Users and Computers containers. To apply Group Policy to users and computers, create new OUs and move the user and computer objects into those OUs. Apply the Group Policy setting to the new OUs.

Optionally, you can redirect the creation of objects that are placed in the default containers to be placed in containers of your choice. For more information about redirecting the default location of new objects, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.

**Well-known Users and Groups and Built-in Accounts**

By default, several well-known users and groups and built-in accounts are created in a new domain. It is recommended that management of these accounts remains under the control of the service administrators. Do not delegate management of these accounts to an individual who is not a service administrator. Table 2.10 lists the well-known users and groups and built-in accounts that need to remain under the control of the service administrators.

**Table 2.10 Accounts That Need to Remain Under the Control of the Service Administrators**

| Well-known Users and Groups | Built-in Accounts |
|---|---|
| Cert Publishers | Administrator |
| Domain Controllers | Guest |
| Group Policy Creator Owners | Guests |
| KRBTGT | Account Operators |
| Domain Guests | Administrators |
| Administrator | Backup Operators |
| Domain Admins | Incoming Forest Trust Builders |
| Schema Admins (forest root domain only) | Print Operators |
| Enterprise Admins (forest root domain only) | Pre-Windows 2000 Compatible Access |
| Domain Users | Server Operators |
| | Users |

**Domain Controller OU**

When domain controllers are added to the domain, their computer objects are automatically added to the Domain Controller OU. This OU has a default set of policies applied to it. To ensure that these policies are applied uniformly to all domain controllers, it is recommended that you do not move the computer objects of the domain controllers out of this OU. Failure to apply the default policies can cause a domain controller to fail to function properly.

By default, the service administrators control this OU. Do not delegate control of this OU to individuals other than the service administrators.

# Delegating Administration of Account and Resource OUs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Account OUs contain user, group, and computer objects. Forest owners must create an OU structure to manage these objects and then delegate control of the structure to the OU owner. Resource OUs contain resources and the accounts that are responsible for managing those resources. The forest owner is also responsible for creating an OU structure to manage these resources and delegating control of that structure to the OU owner.

**Delegating Administration of Account OUs**

Delegate an account OU structure to data administrators if they need to be able to create and modify user, group, and computer objects. The administrative capabilities of an account OU structure are similar to those of a Windows NT 4.0 Master User Domain. The account OU structure is a subtree of OUs for each account type that must be independently controlled. For example, the OU owner can delegate specific control to various data administrators over child OUs in an account OU for users, computers, groups, and service accounts, as shown in Figure 2.39.

**Figure 2.39 Account OU Structure**




Table 2.11 lists and describes the child OUs that you can create in an account OU structure.

**Table 2.11 Child OUs in the Account OU Structure**

| OU | Purpose |
|---|---|
| Users | Contains user accounts for non-administrative personnel. |
| Service Accounts | Some services that require access to network resources run as user accounts. This OU is created to separate service user accounts from the user accounts contained in the Users OU. Also, placing the different types of user accounts in separate OUs enables you to manage them according to their specific administrative requirements. |
| Computers | Contains accounts for computers other than domain controllers. |
| Groups | Contains groups of all types, except for administrative groups, which are managed separately. |
| Admins | Contains user and group accounts for data administrators in the account OU structure, to allow them to be managed separately from regular users. Enable auditing for this OU so that you can track changes to administrative users and groups. |

Figure 2.40 shows the administrative group design for an Account OU.

**Figure 2.40 Delegation Model for Account OUs**




The owner of the Account OU is the *acct_ou*_OU_admins group, which is comprised of data administrators. This group has full control of the *acct_ou*_OU subtree, and is responsible for creating the standard set of child OUs and the groups to manage them.

Groups that manage the child OUs are granted full control only over the specific class of objects that they are responsible for managing. For example, *acct_ou*_group_admins has control only over group objects.

Note that no separate administrative group manages the Admins OU; rather, it inherits ownership from its parent OU, so it is managed by *acct_ou*_OU_admins. The OU owner might choose to create additional administrative groups, however. For example, the OU owner might create the optional group *acct_ou*_helpdesk_admins in the Admins OU to control password resets.

**Administrative Group Types**

The types of groups that you use to delegate control within an OU structure are based on where the accounts are located relative to the OU structure that is to be managed. If the admin user accounts and the OU structure all exist within a single domain, then the groups that you create to use for delegation must be global groups. Figure 2.41 shows the delegation of administration within a single domain.

**Figure 2.41 Group Type for Delegation Within a Single Domain**




If your organization has a department that manages its own user accounts and exists in more than one region, you might have a group of data administrators who are responsible for managing account OUs in more than one domain. If the accounts of the data administrators all exist in a single domain and you have OU structures in multiple domains to which you need to delegate control, make those administrative accounts members of global groups and delegate control of the OU structures in each domain to those global groups, as shown in Figure 2.42

**Figure 2.42 Using Global Groups to Manage OUs in Multiple Domains**




If the data administrators accounts to which you delegate control of an OU structure come from multiple domains, you must use a universal group. Universal groups can contain users from different domains and can therefore be used to delegate control in multiple domains.

Figure 2.43 shows a configuration in which universal groups are used to manage OUs in multiple domains.

**Figure 2.43 Using Universal Groups to Manage OUs in Multiple Domains**




**Delegating Administration of Resource OUs**

Resource OUs are used to manage access to resources. The resource OU owner creates computer accounts for servers that are joined to the domain that include resources such as file shares, databases, and printers. The resource OU owner also creates groups to control access to those resources.

Figure 2.44 shows the two possible locations for the resource OU.

**Figure 2.44 Resource OU Placement in a Domain**




The resource OU can be located under the domain root or as a child OU of the corresponding account OU in the OU administrative hierarchy. If the domain spans across different countries or regions, or the domain owner is responsible for a large number of OUs, Windows NT 4.0 resource domain owners might prefer to control resource OUs that are subordinate to account OUs, to ensure that they have direct access to support from the account OU owner. Resource OUs do not have any standard child OUs. Computers and groups are placed directly in the resource OU.

The resource OU owner owns the objects within the OU, but does not own the OU container itself. Resource OU owners manage only computer and group objects; they cannot create other classes of objects within the OU, and they cannot create child OUs.

**Note**

- The creator or owner of an object has the ability to set the ACL on the object, regardless of the permissions that are inherited from the parent container. If a resource OU owner can reset the ACL on an OU, he or she can create any class of object in the OU, including users. For this reason, resource OU owners are not permitted to create OUs.

For each resource OU in the domain, create a global group to represent the data administrators who are responsible for managing the contents of the OU. This group has full control over the group and computer objects in the OU, but not over the OU container itself.

Figure 2.45 shows the administrative group design for a resource OU. The *res_ou*_OU_admin group manages its own membership and is located in the resource OU.

**Figure 2.45 Resource OU Administrative Group Design**




Placing the computer accounts into a resource OU gives the OU owner control over the account objects but does not make the OU owner an administrator of the computers. In an Active Directory domain, the Domain Admins group is, by default, placed in the local Administrators group on all computers. This means that service administrators have control over those computers. If resource OU owners require administrative control over the computers in their OU, the forest owner can apply a Restricted Groups Group Policy to make the resource OU owner a member of the Administrators group on the computers in that OU.

# Creating Account OUs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process that you use to create account OUs varies according to whether you are upgrading a Windows NT 4.0 domain in place or creating a new Windows Server 2003 Active Directory domain. Use one of the following approaches to create account OUs for your domain:

1. If you are upgrading a Windows NT 4.0 MUD in place to an Active Directory domain, create an account OU in the domain for the accounts that are to be upgraded. Leave the service administrator groups and user accounts in the default structure and move the remaining user and group accounts to the newly created account OU.

2. If you are restructuring Windows NT 4.0 MUDs into an Active Directory domain, create an account OU for each independently managed source MUD that is to be migrated into this domain. If the same IT group manages multiple source MUDs, you can migrate accounts from those domains into a single account OU.

3. If you are deploying a new Active Directory domain, create an account OU for the domain so that you can delegate control of the accounts in the domain.

# Creating Resource OUs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use one of the following approaches to create resource OUs for your domain:

If the domain is the target domain in a Windows NT 4.0 resource domain restructure:

1. Create an OU for each source resource domain that is managed independently and that is to be migrated into this domain.

    - If you are restructuring multiple resource domains that are managed by the same IT group, create a single resource OU and migrate the objects from all domains into that OU.

    - If the source resource domains are managed by former Master User Domain owners, you can migrate objects into the Computers and Groups child OUs of the corresponding account OU subtree, instead of creating a separate resource OU under the account OU.

2. Place the resource OU under the domain root or under an account OU, depending on whether the resource OU owner reports to the account OU owner or directly to the forest owner. The source resource domain owner becomes the resource OU owner.

If the domain is not the target of a resource domain restructure, create resource OUs as needed based on the requirements of each group for autonomy in the management of data and equipment.

# Documenting the OU Design for Each Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assemble a team to design the OU structure that you use to delegate control over resources within the forest. The forest owner might be involved in the design process and must sign off on the OU design. You might also involve at least one service administrator to ensure that the design is valid. Other design team participants might include the data administrators who will work on the OUs and the OU owners who will be responsible for managing them.

It is important to document your OU design by listing the names of the OUs that you plan to create, and for each OU, documenting the type of OU, the OU owner, the parent OU, if applicable, and the origin of that OU.

For a worksheet to assist you in documenting your OU design, see "Identifying OUs for Each Domain" (DSSLOGI_9.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Identifying OUs for Each Domain" on the Web at http://www.microsoft.com/reskit).

**Example: Documenting the OU Design for Each Domain**
The HR department of Contoso Corporation needs to delegate control over print server administration and user accounts for HR employees and management of the HR database servers to three separate data administrators in their group. They designed an OU structure based on these needs. To assist the deployment team in implementing the OU design, they documented their OU design plans.

Figure 2.46 shows the documentation that the HR department created for their OU design.

**Figure 2.46 Example of an Identifying OUs for Each Domain Worksheet**

# Identifying OUs for Each Domain

| Prepared By | Mark Lee | Date | 2/19/03 |
|---|---|---|---|

| | |
|---|---|
| Forest Name | humrsc.contoso.com |
| Forest Owner | Mark Lee |
| Domain Name | humrsc.contoso.com |
| Domain Owner (Optional) | N/A |

Instructions: Complete this worksheet for each domain in the forest.

| OU Name | Type | Owner | Parent OU | Origin |
|---|---|---|---|---|
| HR Printers | Resource | Bart Duncan<br>bart@contoso.com | N/A | New |
| HR Data | Resource | Paula Nartker<br>paula@contoso.com | N/A | New |
| HR Accounts | Account | William Vong<br>William@contoso.com | N/A | New |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

# Applying Group Policy to OUs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your initial OU structure can be expanded to apply Group Policy to specific sets of users and computers. For more information about applying Group Policy to OUs, see "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* in this kit.

# Additional Resources for Designing the Active Directory Logical Structure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- The *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit)for more information about Active Directory concepts.
- "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.
- The Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search under "Administration and Configuration Guides" and download the *Active Directory Operations Guide* for more information about Active Directory operations.
- The Active Directory in Networks Segmented by Firewalls link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about configuring firewalls for use with Active Directory.
- "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book for more information about consolidating and upgrading Windows NT 4.0 domains.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about DNS and the forest-wide application partition.
- "Deploying DNS" in *Deploying Network Services* in this kit for more information about DNS integration with Active Directory.
- "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* in this kit for more information about applying Group Policy to OUs.
- The Active Directory Branch Office Planning Guide link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for a complete guide to information about Active Directory branch office implementations.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Active Directory" in Help and Support Center for Windows Server 2003.

**Related Job Aids**

- "Design and Deployment Team Information" (DSSLOGI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Design and Deployment Team Information" on the Web at http://www.microsoft.com/reskit).
- "Forest Design Requirements" (DSSLOGI_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Forest Design Requirements" on the Web at http://www.microsoft.com/reskit).
- "Forest Design" (DSSLOGI_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Forest Design" on the Web at http://www.microsoft.com/reskit).

- "Identifying Regions" (DSSLOGI_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Identifying Regions" on the Web at http://www.microsoft.com/reskit).

- "Domain Planning" (DSSLOGI_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Planning" on the Web at http://www.microsoft.com/reskit).

- "Migration Plan for Windows NT 4.0 Master User Domains" (DSSLOGI_6.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migration Plan for Windows NT 4.0 Master User Domains" on the Web at http://www.microsoft.com/reskit).

- "Migration Plan for Windows NT 4.0 Resource Domains" (DSSLOGI_7.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migration Plan for Windows NT 4.0 Resource Domains" on the Web at http://www.microsoft.com/reskit).

- "DNS Inventory" (DSSLOGI_8.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DNS Inventory" on the Web at http://www.microsoft.com/reskit).

- "Identifying OUs for Each Domain" (DSSLOGI_9.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Identifying OUs for Each Domain" on the Web at http://www.microsoft.com/reskit).

# Designing the Site Topology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An Active Directory® directory service site topology is a logical representation of your physical network. Designing an Active Directory site topology involves planning for domain controller placement and designing sites, subnets, site links, and site link bridges to ensure efficient routing of query and replication traffic.

**In This Chapter**


Overview of Designing a Site Topology

Collecting Network Information

Planning Domain Controller Placement

Creating a Site Design

Creating a Site Link Design

Creating a Site Link Bridge Design

Additional Resources for Designing the Site Topology


**Related Information**

- For more information about designing the Active Directory logical structure and the DNS infrastructure needed to support Active Directory, see "Designing the Active Directory Logical Structure" in this book.
- For more information about determining domain controller hardware requirements, see "Planning Domain Controller Capacity" in this book.
- For more information about deploying a DNS infrastructure for name resolution on your network, see "Deploying DNS" in *Deploying Network Services* of this kit.
- For more information about Active Directory replication topology, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

# Overview of Designing a Site Topology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing a site topology helps you efficiently route client queries and Active Directory replication traffic. A well-designed site topology helps your organization achieve the following benefits:

- Minimize the cost of replicating Active Directory data.
- Minimize administrative efforts that are required to maintain the site topology.
- Schedule replication that enables locations with slow or dial-up network links to replicate Active Directory data during off-peak hours.
- Optimize the ability of client computers to locate the nearest resources, such as domain controllers and Distributed File System (DFS) servers, reducing network traffic over slow, wide area network (WAN) links, improving logon and logoff processes, and speeding up file download operations.

Before you begin to design your site topology, you must understand your physical network structure. In addition, you must first design your Active Directory logical structure, including the administrative hierarchy, forest plan, and domain plan for each forest. You must also complete your DNS infrastructure design for Active Directory. For more information about designing your Active Directory logical structure and DNS infrastructure, see "Designing the Active Directory Logical Structure" in this book.

After you complete your site topology design, you must verify that your domain controllers meet the hardware requirements for the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition, operating systems. You must also determine the appropriate number of domain controllers for each domain that is represented in each site. For more information about determining the appropriate number of domain controllers and their hardware requirements, see "Planning Domain Controller Capacity" in this book.

**Note**

- For a list of job aids that are available to assist you in designing the site topology, see "Additional Resources for Designing the Site Topology" later in this chapter.

# Process for Designing a Site Topology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing a site topology includes determining what locations require domain controllers and determining where you need to create sites, site links, and site link bridges. Figure 3.1 illustrates the site topology design process.

**Figure 3.1 Designing a Site Topology**



# Site Topology Design Background Information

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your site topology significantly affects the performance of your network and the ability of your users to access network resources. Before you begin to design your site topology, become familiar with the functions for sites in Microsoft® Windows Server 2003, the different network topologies that organizations commonly use, the role of the site topology owner, and some Active Directory replication concepts.

# Functions for Sites in Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 uses site information for many purposes, including routing replication, client affinity, system volume replication, DFS, and service location.

**Routing replication**
Active Directory uses a multimaster, store-and-forward method of replication. A domain controller communicates directory changes to a second domain controller, which then communicates to a third, and so on, until all domain controllers have received the change. To achieve the best balance between reducing replication latency and reducing traffic, site topology controls Active Directory replication by distinguishing between replication that occurs within a site and replication that occurs between sites.

Within sites, replication is optimized for speed — data updates trigger replication and the data is sent without the overhead required by data compression. Conversely, replication between sites is compressed to minimize the cost of transmission over WAN links. When replication occurs between sites, a single domain controller per domain at each site collects and stores the directory changes and communicates them at a scheduled time to a domain controller in another site.

**Client affinity**
Domain controllers use site information to inform Active Directory clients about domain controllers present within the closest site as the client. For example, consider a client in the Seattle site that does not know its site affiliation and contacts a domain controller from the Atlanta site. Based on the IP address of the client, the domain controller in Atlanta determines which site the client is actually from and sends the site information back to the client. The domain controller also informs the client whether the chosen domain controller is the closest one to it. The client caches the site information provided by the domain controller in Atlanta and queries for the site-specific service (SRV) resource record (a DNS resource record used to locate domain controllers for Active Directory) and thereby finds a domain controller within the same site.

By finding a domain controller in the same site, the client avoids communications over WAN links. If no domain controllers are located at the client site, a domain controller that has the lowest cost connections relative to other connected sites advertises itself (registers a site-specific SRV resource record in DNS) in the site that does not have a domain controller. The domain controllers that are published in DNS are those from the closest site as defined by the site topology. This process ensures that every site has a preferred domain controller for authentication.

For more information about the process of locating a domain controller, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

**SYSVOL replication**
The system volume (SYSVOL) is a collection of folders in the file system that exists on each domain controller in a domain. The SYSVOL folders provide a default Active Directory location for files that must be replicated throughout a domain, including Group Policy objects (GPO), startup and shutdown scripts, and logon and logoff scripts. Windows Server 2003 uses the File Replication service (FRS) to replicate changes made to the SYSVOL folders from one domain controller to other domain controllers. FRS replicates these changes according to the schedule that you create during your site topology design.

**DFS**
DFS uses site information to direct a client to the server that is hosting the requested data within the site. If DFS does not find a copy of the data within the same site as the client, DFS uses the site information in Active Directory to determine which file server that has DFS shared data is closest to the client.

**Service location**
By publishing services such as file and print services in Active Directory, you allow Active Directory clients to locate the requested service within the same or nearest site. Print services use the location attribute stored in Active Directory to let users browse for printers by location without knowing their precise location. For information about enabling printer location tracking, see "Enable printer location tracking" in Help and Support Center for Windows Server 2003. For more information about designing and deploying print servers, see "Designing and Deploying Print Servers" in *Planning Server Deployments* of this kit.

# Site Topology Owner Role

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The administrator who manages the site topology is known as the site topology owner. The site topology owner understands the conditions of the network between sites and has the authority to change settings in Active Directory to implement changes to the site topology. Changes to the site topology affect changes in the replication topology. The site topology owner's responsibilities include:

- Controlling changes to the site topology if network connectivity changes.
- Obtaining and maintaining information about network connections and routers from the network group. The site topology owner must maintain a list of subnet addresses, subnet masks, and the location to which each belongs. The site topology owner must also understand any issues about network speed and capacity that affect site topology in order to effectively set costs for site links.
- Moving Active Directory server objects representing domain controllers between sites if a domain controller's IP address changes to a different subnet in a different site, or if the subnet itself is assigned to a different site. In either case, the site topology owner must manually move the Active Directory server object of the domain controller to the new site.

# Network Topologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization's network topology reflects its business needs. In some organizations, business needs result in users working in a few large locations that are well connected to one another. Alternatively, other organizations have users in many small, satellite locations connected to one of a few, well-connected hub sites. Such network topologies usually are one of three general types: ring, hub and spoke, and complex. These network topologies are illustrated in Figure 3.2.

**Figure 3.2 Network Topologies**



# Active Directory Replication Concepts

Updated: October 27, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before designing site topology, become familiar with some Active Directory replication concepts.

**Connection object**
A connection object is an Active Directory object that represents a replication connection from one domain controller to another. A domain controller is a member of a single site and is represented in the site by a server object in Active Directory. Each server object has a child NTDS Settings object that represents the replicating domain controller in the site.

The connection object is a child of the NTDS Settings object on the destination server. For replication to occur between two domain controllers, the server object of one must have a connection object that represents inbound replication from the other. All replication connections for a domain controller are stored as connection objects under the NTDS Settings object. The connection object identifies the replication source server, contains a replication schedule, and specifies a replication transport.

The Knowledge Consistency Checker (KCC) creates connection objects automatically, but they can also be created manually. Connection objects created by the KCC appear in the Active Directory Sites and Services snap-in as **<automatically generated>** and are considered adequate under normal operating conditions. Connection objects created by an administrator are manually created connection objects. A manually created connection object is identified by the name assigned by the administrator when it was created. When you modify an **<automatically generated>** connection object, you convert it into an administratively modified connection object and the object appears in the form of a GUID. The KCC does not make changes to manual or modified connection objects.

**KCC**
The KCC is a built-in process that runs on all domain controllers and generates replication topology for the Active Directory forest. The KCC creates separate replication topologies depending on whether replication is occurring within a site (intrasite) or between sites (intersite). The KCC also dynamically adjusts the topology to accommodate new domain controllers, domain controllers moved to and from sites, changing costs and schedules, and domain controllers that are temporarily unavailable.

Within a site, the connections between domain controllers are always arranged in a bidirectional ring, with additional shortcut connections to reduce latency in large sites. On the other hand, the intersite topology is a layering of spanning trees, which means one intersite connection exists between any two sites for each directory partition and generally does not contain shortcut connections. For more information about spanning trees and Active Directory replication topology, see Active Directory Replication Topology Technical Reference in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=44137).

On each domain controller, the KCC creates replication routes by creating one-way inbound connection objects that define connections from other domain controllers. For domain controllers in the same site, the KCC creates connection objects automatically without administrative intervention. When you have more than one site, you configure site links between sites and a single KCC in each site automatically creates connections between sites as well.

Figure 3.3 shows two sites (Seattle and Los Angeles) with domain controllers that are all in the same domain. The arrows represent possible inbound connections that the KCC creates. Because all Active Directory updates are transferred in a ring within a site and redundant connections exist, all domain controllers can receive updates from all other domain controllers in the Seattle site, although domain controllers within a site do not necessarily replicate in both directions. For replication to occur between Seattle and Los Angeles, one domain controller in each site has a replication agreement with a domain controller in the other site. Between sites, these replication partners replicate in both directions according to configuration settings for a site link object that represents the physical WAN connecting the two sites. In Figure 3.3, domain controllers DC-3 and DC-4 are replication partners between the Seattle and Los Angeles sites.

**Figure 3.3 Intersite and Intrasite Replication Connections**




**Failover functionality**
Sites ensure that replication is routed around network failures and offline domain controllers. The KCC runs at specified intervals to adjust the replication topology for changes that occur in Active Directory, such as when new domain controllers are added and new sites are created. The KCC reviews the replication status of existing connections to determine if any connections are not working. If a connection is not working due to a failed domain controller, the KCC automatically builds temporary connections to other replication partners (if available) to ensure that replication occurs. If all the domain controllers in a site are unavailable, the KCC automatically creates replication connections between domain controllers from another site.

**Subnet**
A subnet is a segment of a TCP/IP network to which a set of logical IP addresses are assigned. Subnets group computers in a way that identifies their physical proximity on the network. Subnet objects in Active Directory identify the network addresses that are used to map computers to sites.

**Site**
Sites are Active Directory objects that represent one or more TCP/IP subnets with highly reliable and fast network connections. Site information allows administrators to configure Active Directory access and replication to optimize usage of the physical network. Site objects are associated with a set of subnets, and each domain controller in a forest is associated with an Active Directory site according to its IP address. Sites can host domain controllers from more than one domain, and a domain can be represented in more than one site.

**Site link**
Site links are Active Directory objects that represent logical paths that the KCC uses to establish a connection for Active Directory replication. A site link object represents a set of sites that can communicate at uniform cost through a specified intersite transport.

All sites contained within the site link are considered to be connected by means of the same network type. Sites must be manually linked to other sites by using site links so that domain controllers in one site can replicate directory changes from domain controllers in another site. Because site links do not correspond to the actual path taken by network packets on the physical network during replication, you do not need to create redundant site links to improve Active Directory replication efficiency.

When two sites are connected by a site link, the replication system automatically creates connections between specific domain controllers in each site that are called *bridgehead servers*. In Microsoft® Windows® 2000, intersite replication of the directory partitions (domain, configuration, and schema) between domain controllers in different sites is performed by the bridgehead servers (one per directory partition) in those sites. In Windows Server 2003, all domain controllers in a site that host the same directory partition are candidates for being selected as bridgehead servers. The replication connections created by the KCC are randomly distributed among all candidate bridgehead servers in a site to share the replication workload. By default, the randomized selection process takes place only once, when connection objects are first added to the site.

However, you can run the Active Directory Load Balancing tool (Adlb.exe) to rebalance the load each time a change occurs in the site topology or in the number of domain controllers the site. In addition, Adlb can stagger schedules so that the outbound replication load for each domain controller is spread out evenly across time. Consider using Adlb to balance replication traffic between the Windows Server 2003–based domain controllers when they are replicating to more than 20 other sites hosting the same domain. For more information about using Adlb.exe and managing environments that have 100 or more branch sites, see the Windows Server 2003 Active Directory Branch Office Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=28523).

**Site link bridge**
A site link bridge is an Active Directory object that represents a set of site links, all of whose sites can communicate by using a common transport. Site link bridges enable domain controllers that are not directly connected by means of a communication link to replicate with each other. Typically, a site link bridge corresponds to a router (or a set of routers) on an IP network.

By default, the KCC can form a transitive route through any and all site links that have some sites in common. If this behavior is disabled, each site link represents its own distinct and isolated network. Sets of site links that can be treated as a single route are expressed through a site link bridge. Each bridge represents an isolated communication environment for network traffic.

Site link bridges are a mechanism to logically represent transitive physical connectivity between sites. A site link bridge allows the KCC to use any combination of the included site links to determine the least expensive route to interconnect directory partitions held in those sites. The site link bridge does not provide actual connectivity to the domain controllers. If the site link bridge is removed, replication over the combined site links will continue until the KCC removes the links.

Site link bridges are only necessary if a site contains a domain controller hosting a directory partition that is not also hosted on a domain controller in an adjacent site, but a domain controller hosting that directory partition is located in one or more other sites in the forest. Adjacent sites are defined as any two or more sites included in a single site link.

A site link bridge creates a logical connection between two site links, providing a transitive path between two disconnected sites by using an interim site. For the purposes of the intersite topology generator (ISTG), the bridge implies physical connectivity by using the interim site. The bridge does not imply that a domain controller in the interim site will provide the replication path. However, this would be the case if the interim site contained a domain controller that hosted the directory partition to be replicated, in which case a site link bridge is not required.

The cost of each site link is added, creating a summed cost for the resulting path. The site link bridge would be used if the interim site did not contain a domain controller hosting the directory partition and a lower cost link does not exist. If the interim site contained a domain controller that hosts the directory partition, two disconnected sites would set up replication connections to the interim domain controller and not use the bridge.

**Site link transitivity**

By default all site links are transitive, or "bridged." When site links are bridged and the schedules overlap, the KCC creates replication connections that determine domain controller replication partners between sites, where the sites are not directly connected by site links but are connected transitively through a set of common sites. This means that you can connect any site to any other site through a combination of site links.

In general, for a fully routed network, you do not need to create any site link bridges unless you want to control the flow of replication changes. All site links for a specific transport implicitly belong to a single site link bridge for that transport. The default bridging for site links occurs automatically, and no Active Directory object represents that bridge. The **Bridge all site links** setting found in the properties of both the IP and Simple Mail Transfer Protocol (SMTP) intersite transport containers implements automatic site link bridging.

**Global catalog server**

A global catalog server is a domain controller that stores information about all objects in the forest so that applications can search Active Directory without referring to specific domain controllers that store the requested data. Like all domain controllers, a global catalog server stores full, writable replicas of the schema and configuration directory partitions and a full, writable replica of the domain directory partition for the domain that it is hosting. In addition, a global catalog server stores a partial, read-only replica of every other domain in the forest. Partial, read-only domain replicas contain every object in the domain but only a subset of the attributes (those attributes that are most commonly used for searching the object).

**Universal group membership caching**

Universal group membership caching allows the domain controller to cache universal group membership information for users. You can enable domain controllers that are running Windows Server 2003 to cache universal group memberships by using the Active Directory Sites and Services snap-in.

Enabling universal group membership caching eliminates the need for a global catalog server at every site in a domain, which minimizes network bandwidth usage because a domain controller does not need to replicate all of the objects located in the forest. It also reduces logon times because the authenticating domain controllers do not always need to access a global catalog to obtain universal group membership information. For more information about when to use universal group membership caching, see "Planning Global Catalog Server Placement" later in this chapter.

# Collecting Network Information

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in designing an effective Active Directory site topology is to consult your organization's networking group to collect information and communicate with them regularly about your physical network topology. Figure 3.4 shows the process for collecting information about your physical network.

**Figure 3.4 Collecting Network Information**



# Creating a Location Map

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Create a location map that represents the physical network infrastructure of your organization. On the location map, identify the geographic locations that contain groups of computers with internal connectivity of 10 megabits per second (Mbps) or higher (local area network [LAN]-speed or better).

Figure 3.5 shows an example of a location map for Trey Research.

**Figure 3.5 Trey Research Locations**



# Listing Communication Links and Available Bandwidth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After creating a location map, document the type of communication link, its link speed, and the available bandwidth between each location. Obtain a WAN topology from your networking group. For a list of common WAN circuit types and their bandwidth, see "Determining the Cost" later in this chapter. You need this information to create site links later in the site topology design process.

*Bandwidth* refers to the amount of data that you can transmit across a communication channel in a given amount of time. *Available bandwidth* refers to the amount of bandwidth actually available for use by Active Directory. You can obtain available bandwidth information from your network group or you can analyze traffic on each link by using a protocol analyzer such as Network Monitor, which is a component that ships with Windows Server 2003. For information about installing Network Monitor, see "Monitoring network traffic" on the Microsoft Web site.

Document each location and the other locations that are linked to it. In addition, record the type of communication link and its available bandwidth. For a worksheet to assist you in listing communication links and available bandwidth, see "Geographic Locations and Communication Links" (DSSTOPO_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Geographic Locations and Communication Links" on the Web at http://www.microsoft.com/reskit).

Figure 3.6 shows an example of a worksheet listing the locations for Trey Research, the communication links between each location, and the available bandwidth for each communication link.

**Figure 3.6 Example of a Geographic Locations and Communication Links Worksheet**

| Geographic Locations and Communication Links | | | |
|---|---|---|---|
| Prepared By | Site Topology Owner | Date | 1/1/03 |
| Forest Name | Trccorp.treyresearch.net | | |

| Hub Location | Satellite Location | Connection Type | Available Bandwidth |
|---|---|---|---|
| Seattle | Los Angeles | ISDN (128.8 Kbps) | No more than 56 Kbps |
| | Phoenix | ISDN (128.8 Kbps) | No more than 28.8 Kbps |
| | New York | T1 (1.544 Mbps) | No more than 150 Kbps, but with 450-millisecond latency |
| New York | Washington DC | ISDN (128.8 Kbps) | No more than 56 Kbps |
| | Boston | T1 (1.544 Mbps) | No more than 200 Kbps, but with 450-millisecond latency |

# Listing IP Subnets Within Each Location

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you document the communication links and the available bandwidth between each location, record the IP subnets within each location. If you do not already know the subnet mask and IP address within each location, consult your networking group.

Active Directory associates a workstation with a site by comparing the workstation's IP address with the subnets that are associated with each site. As you add domain controllers to a domain, Active Directory also examines their IP addresses and places them in the most appropriate site.

For a worksheet to assist you in listing the IP subnets within each location, see "Locations and Subnets" (DSSTOPO_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Locations and Subnets" on the Web at http://www.microsoft.com/reskit).

Figure 3.7 shows an example of a completed worksheet listing the IP subnets within each Trey Research location.

**Figure 3.7 Example of a Locations and Subnets Worksheet**

Locations and Subnets

| | | | |
|---|---|---|---|
| Prepared By | Site Topology Owner | Date | 1/1/03 |
| Forest Name | Trccorp.treyresearch.net | | |

| Location | IP Addresses | Subnet Mask |
|---|---|---|
| Seattle | 172.16.16.0<br>172.16.24.0 | 255.255.252.0<br>255.255.252.0 |
| Boston | 172.16.36.0<br>172.16.40.0 | 255.255.252.0<br>255.255.252.0 |
| New York | 172.16.28.0<br>172.16.32.0 | 255.255.252.0<br>255.255.252.0 |
| Los Angeles | 172.16.136.0<br>172.16.92.0 | 255.255.252.0<br>255.255.252.0 |
| Washington DC | 172.16.116.0<br>172.16.140.0 | 255.255.252.0<br>255.255.252.0 |
| Phoenix | 172.16.64.0 | 255.255.252.0 |

# Listing Domains and Number of Users for Each Location

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The number of users for each regional domain that is represented in a location is one of the factors that determine the placement of regional domain controllers and global catalog servers, which is the next step in the site topology design process. For example, plan to place a regional domain controller in a location that contains more than one hundred regional domain users so they can still log on to the domain if the WAN link fails.

Record the locations, the domains that are represented in each location, and the number of users for each domain that is represented in each location. For a worksheet to assist you in listing the domains and the number of users that are represented in each location, see "Domains and Users in Each Location" (DSSTOPO_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domains and Users in Each Location" on the Web at http://www.microsoft.com/reskit).

Figure 3.8 shows a completed worksheet listing the domains that are represented in each Trey Research location, and the number of users for each domain.

**Figure 3.8 Example of a Domains and Users in Each Location Worksheet**

Domains and Users in Each Location

| | | | |
|---|---|---|---|
| Prepared By | Site Topology Owner | Date | 1/1/03 |
| Forest Name | Trccorp.treyresearch.net | | |

| Location | Domains in Location | Users per Domain per Location |
|---|---|---|
| Seattle | trccorp.treyresearch.net | 20 |
| | west.trccorp.treyresearch.net | 15,000 |
| | east.trccorp.treyresearch.net | 200 |
| Los Angeles | west.trccorp.treyresearch.net | 5,000 |
| Washington DC | east.trccorp.treyresearch.net | 50 |
| New York | east.trccorp.treyresearch.net | 7,000 |
| | west.trccorp.treyresearch.net | 200 |
| Boston | east.trccorp.treyresearch.net | 80 |
| Phoenix | west.trccorp.treyresearch.net | 20 |

# Planning Domain Controller Placement

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have gathered all of the network information that will be used to design your site topology, plan where you want to place domain controllers, including regional domain controllers, forest root domain controllers, operations master role holders, and global catalog servers. Figure 3.9 shows the process for determining domain controller placement.

**Figure 3.9 Planning Domain Controller Placement**




Although this process helps you plan where you intend to place domain controllers, it does not address how you determine the proper number of domain controllers and the domain controller hardware requirements for each domain that is represented in each site. For more information about determining the proper number of domain controllers for each domain that is represented in each site, see "Planning Domain Controller Capacity" in this book.

# Planning Forest Root Domain Controller Placement

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Forest root domain controllers are needed to create trust paths for clients that need to access resources in domains other than their own. Place forest root domain controllers at locations that host datacenters and in hub locations. If users in a given location need to access resources from other domains in the same location, and the network availability between the datacenter and the user location is unreliable, then you can either add a forest root domain controller in the location or create a shortcut trust between the two domains. It is more cost efficient to create a shortcut trust between the domains unless you have other reasons to place a forest root domain controller in that location.

Shortcut trusts help to optimize authentication requests made from users located in either domain. For more information about how to create shortcut trusts between domains, see "Create a shortcut trust" in Help and Support Center for Windows Server 2003.

For a worksheet to assist you in documenting your forest root domain controller placement, see "Domain Controller Placement" (DSSTOPO_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Controller Placement" on the Web at http://www.microsoft.com/reskit). For an example of a completed Domain Controller Placement worksheet, see "Example: Determining Domain Controller Placement" later in this chapter.

You need to refer to this information when you create the forest root domain. For more information about deploying the forest root domain, see "Deploying the Windows Server 2003 Forest Root Domain" in this book.

# Planning Regional Domain Controller Placement

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning Regional Domain Controller Placement**

For cost efficiency, plan to place as few regional domain controllers as possible. First review the Geographic Locations and Communication Links worksheet to determine whether a location is a hub. Plan to place regional domain controllers for each domain that is represented in each hub location.

In addition, ensure the physical security of domain controllers in hub locations so that unauthorized personnel cannot access them. Do not place domain controllers in a location in which you cannot guarantee the physical security of the domain controller.

After you place regional domain controllers in all hub locations, evaluate the need for placing regional domain controllers at satellite locations. Eliminating unnecessary regional domain controllers from satellite locations reduces the support costs required to maintain a remote server infrastructure.

To authenticate client logons and access to local file servers, most organizations place regional domain controllers for all regional domains that are represented in a given location. However, you must consider many variables when evaluating whether a business location requires its clients to have local authentication or the clients can rely on authentication and query over a WAN link. Figure 3.10 shows how to determine whether to place domain controllers at satellite locations.

**Figure 3.10 Determining Whether to Place Domain Controllers at Satellite Locations**




**Domain Controller Physical Security**

Do not place a regional domain controller in a satellite location if you cannot ensure the physical security of the domain controller. A person who has physical access to a domain controller can attack the system by:

- Accessing physical disks by starting an alternate operating system on a domain controller.
- Removing (and possibly replacing) physical disks on a domain controller.
- Obtaining and manipulating a copy of a domain controller system state backup.

Add regional domain controllers only to locations in which you can guarantee their physical security. For more information about securing domain controllers, see Best Practice Guide for Securing Windows Server Active Directory Installations on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=28521).

**On-site Technical Expertise Availability**

Domain controllers need to be managed continuously for various reasons. Place a regional domain controller only in locations that include personnel who can administer the domain controller, or be sure that the domain controller can be managed

remotely.

**WAN Link Availability**
WAN links that experience frequent outages can cause significant productivity loss to users if the location does not include a domain controller that can authenticate the users. If your WAN link availability is not 100 percent, place a regional domain controller in locations where the users require the ability to log on when the WAN link is down.

**Authentication Availability**
Certain organizations, such as banks, require that users be authenticated at all times. Place a regional domain controller in a location where the WAN link availability is not 100 percent but users require authentication at all times.

**Logon Performance over WAN Link**
If your WAN link availability is 100 percent, then whether you place a domain controller at the location depends on the logon performance requirements over the WAN link. Factors that influence logon performance over the WAN include link speed and available bandwidth, number of users and usage profiles, and the amount of logon network traffic versus replication traffic.

**WAN link speed and bandwidth utilization**
The activities of a single user can congest a slow WAN link. Place a domain controller at a location if logon performance over the WAN link is unacceptable.

The average percentage of bandwidth utilization indicates how congested a network link is. If a network link has an average bandwidth utilization that is greater than an acceptable value, place a domain controller at that location.

**Number of users and usage profile**
The number of users and their usage profile at a given location can help determine whether you need to place regional domain controllers at that location. To avoid productivity loss in the event of a WAN link failure, place a regional domain controller at a location that has a staff of one hundred or more users.

The usage profile indicates how the users use the network resources. You need not place a domain controller in a location that contains only a few users who do not frequently access network resources.

**Logon network traffic vs. replication traffic**
If a domain controller is not available within the same location as the Active Directory client, then the client creates logon traffic on the network. The amount of logon network traffic that is created on the physical network is influenced by several factors, including group memberships, number and size of GPOs, logon scripts, and IntelliMirror® features such as offline folders, folder redirection, and roaming profiles.

On the other hand, a domain controller that is placed at a given location generates replication traffic on the network. The frequency and amount of updates made on the partitions hosted on the domain controllers influence the amount of replication traffic that is created on the network. The different types of updates that can be made on the partitions hosted on the domain controllers include adding or changing users and user attributes, changing passwords, and adding or changing global groups, printers, or volumes.

To determine if you need to place a regional domain controller at a location, compare the cost of logon traffic created by a location without a domain controller versus the cost of replication traffic created by placing a domain controller at the location.

For example, consider a network that has branch offices that are connected through slow links to the headquarters and in which domain controllers can easily be added. If the daily logon and directory lookup traffic of a few remote site users causes more network traffic than replicating all company data to the branch, consider adding a domain controller to the branch.

If reducing the cost of maintaining domain controllers is more important than network traffic, centralize the domain controllers for that domain and do not place any regional domain controllers at the location.

For a worksheet to assist you in documenting the placement of regional domain controllers and the number of users for each domain that is represented in each location, see "Domain Controller Placement" (DSSTOPO_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Controller Placement" on the Web at http://www.microsoft.com/reskit). For an example of a completed Domain Controller Placement worksheet, see "Example: Determining Domain Controller Placement" later in this chapter.

You need to refer to the information about locations in which you need to place regional domain controllers when you deploy regional domains. For more information about deploying regional domains, see "Deploying Windows Server 2003 Regional Domains" in this book.

# Planning Global Catalog Server Placement

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Global catalog placement requires planning except if you have a single-domain forest. In a single-domain forest, configure all domain controllers as global catalog servers. Because every domain controller stores the only domain directory partition in the forest, configuring each domain controller as a global catalog server does not require any additional disk space usage, CPU usage, or replication traffic. In a single-domain forest, all domain controllers act as "virtual global catalog servers" in that they can all respond to any authentication or service request. This special condition for single-domain forests is by design (authentication requests do not require contacting a global catalog server, as they do when there are multiple domains and a user can be a member of a universal group that exists in a different domain). However, only domain controllers that are designated as global catalog servers can respond to global catalog queries on the global catalog port 3268. To simplify administration in this scenario and to ensure consistent responses, designating all domain controllers as global catalog servers eliminates the concern about which domain controllers can respond to global catalog queries. Specifically, any time a user uses Start\Search\For People or Find Printers, these requests go only to the global catalog.

In multiple-domain forests, global catalog servers facilitate user logon requests and forest-wide searches. Figure 3.11 illustrates how to determine which locations require global catalog servers.

**Figure 3.11 Determining the Placement of Global Catalog Servers**




**Adding Global Catalog Servers Based on Application Requirements**

Certain applications, such as Microsoft® Exchange® 2000, Message Queuing (also known as MSMQ), and applications using Distributed COM (DCOM), do not deliver adequate response over latent WAN links and therefore need a highly available global catalog infrastructure to provide low query latency. Determine whether any applications that perform poorly over a slow WAN link are running in locations or whether the locations include Exchange servers. If your locations include applications that do not deliver optimum response over a WAN link, you must place a global catalog server at the location to reduce query latency.

**Adding Global Catalog Servers for a Large Number of Users**

Place global catalog servers at all locations that contain more than 100 users to reduce congestion of network WAN links and to prevent productivity loss in case of WAN link failure.

**Using Highly Available Bandwidth**

You do not need to place a global catalog at a location that does not include applications that require a global catalog server, includes less than 100 users, and is also connected to another location that includes a global catalog server by a WAN link that is 100 percent available for Active Directory. In this case, the users can access the global catalog server over the WAN link.

**Adding Global Catalog Servers for a Large Number of Roaming Users**
Roaming users need to contact the global catalog servers whenever they log on for the first time at any location. Place a global catalog at a location that is visited by a large number of roaming users if the logon time over the WAN link is unacceptable.

**Enabling Universal Group Membership Caching**
For locations that include less than 100 users and do not include a large number of roaming users or applications that require a global catalog server, you can deploy domain controllers that are running Windows Server 2003 and enable universal group membership caching. Ensure that the global catalog servers are not more than one replication hop from the domain controller on which universal group membership caching is enabled, so that universal group information in the cache can be refreshed.

For more information about how to enable universal group membership caching, see "Cache universal group memberships" in Help and Support Center for Windows Server 2003. For information about how universal group caching works, see "How the Global Catalog Works" in the Windows Server 2003 Technical Reference.

For a worksheet to assist you in documenting where you plan to place global catalog servers and domain controllers with universal group caching enabled, see "Domain Controller Placement" (DSSTOPO_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Controller Placement" on the Microsoft Web site http://go.microsoft.com/fwlink/?linkid=9568). For an example of a completed Domain Controller Placement worksheet, see "Example: Determining Domain Controller Placement" later in this chapter. You need to refer to the information about locations in which you need to place global catalog servers when you deploy the forest root domain and regional domains.

# Planning Operations Master Role Placement

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning Operations Master Role Placement**

Active Directory supports multimaster replication of directory data, which means any domain controller can accept directory changes and replicate the changes to all other domain controllers. However, certain changes, such as schema modifications, are impractical to perform in a multimaster fashion. For this reason certain domain controllers, known as *operations masters*, hold roles responsible for accepting requests for certain specific changes.

Three operations master roles exist in each domain:

- The primary domain controller (PDC) emulator processes all replication requests from Microsoft® Windows NT® 4.0 backup domain controllers (BDCs) and processes all password updates for clients that are not running Active Directory client software.
- The relative identifier (RID) master allocates RIDs to all domain controllers to ensure that all security principals have a unique identifier.
- The infrastructure master for a given domain maintains a list of the security principals from other domains that are members of groups within its domain.

In addition to the three domain-level operations master roles, two operations master roles exist in each forest:

- The schema master governs changes to the schema.
- The domain naming master adds and removes domains to and from the forest.

Place the domain controllers hosting these operations master roles in areas where network reliability is high, and ensure that the PDC emulator and the RID master are consistently available.

Operations master role holders are assigned automatically when the first domain controller in a given domain is created. The two forest-level roles (schema master and domain naming master) are assigned to the first domain controller created in a forest. Additionally, the three domain-level roles (RID master, infrastructure master, and PDC emulator) are assigned to the first domain controller created in a domain.

Place the first domain controller for a domain in a location that has the largest number of users for that domain. Designate a standby operations master for a domain controller that hosts the operations master roles. The standby operations master is a domain controller that you identify as the computer that assumes the operations master role if the original role holder fails. Ensure that the standby operations master is a direct replication partner of the actual operations master.

Place the domain controllers that host these operations master roles in areas where the network is reliable and where the operations masters can be accessed by all other domain controllers in the forest. For more information about how certain operations might be affected if an operations master role holder is not reachable, see "Operations Master Placement for Networks with Limited Connectivity" later in this section.

**Operations Master Placement for Single Domain Forest**

In addition to hosting the operations master roles, the first domain controller created in a forest also hosts the global catalog. In a single domain forest, the database of a global catalog server is identical to that of a domain controller. Therefore, configure all domain controllers as global catalog servers because it does not cause any additional workload for the domain controllers. In a single domain forest where all domain controllers are configured as global catalog servers, leave all operation master roles on the first domain controller that is created in the forest and use the second domain controller as a standby operations master.

**Operations Master Placement for Forest Root Domain**

In a forest hosting multiple domains, if all domain controllers in the forest root domain are also global catalog servers, leave all the operations master roles on the first domain controller. Use the second domain controller deployed in the forest as the standby operations master.

However, if all domain controllers in the forest root domain are not also global catalog servers, move all the operations master roles to the second domain controller deployed in the forest root domain and ensure that this domain controller is never configured as a global catalog server. This is because the first domain controller is always a global catalog server and the infrastructure master should not be placed on a domain controller that is also a global catalog server unless all domain controllers are global catalog servers. To simplify the environment, keep all the operations master roles on the second domain controller. Configure the third domain controller deployed in the forest root domain as the standby operations master and ensure that this domain controller will also never be configured as global catalog server.

**Operations Master Placement for Regional Child Domain**

The three domain-level roles are assigned to the first domain controller in the domain by default. If any domain controllers in the regional domain will not host the global catalog, leave the three-domain-level operations master roles on the first domain controller and ensure that the first domain controller is never configured as a global catalog server. Configure the second domain controller deployed in this domain to be the standby operations master.

**Planning the PDC emulator placement**

The PDC emulator acts as a Windows NT 4.0 PDC if the domain contains pre–Active Directory clients or if it contains Windows NT 4.0 BDCs. It processes password changes from clients and replicates updates to the BDCs. Only one domain controller acts as the PDC emulator in each domain in the forest.

Even if all the domain controllers are upgraded to Windows 2000 or Windows Server 2003 and the domain is operating at the Windows 2000 native functional level, the PDC emulator receives preferential replication of password changes performed by other domain controllers in the domain. If a password was recently changed, that change takes time to replicate to every domain controller in the domain. If logon authentication fails at another domain controller due to a bad password, that domain controller forwards the authentication request to the PDC emulator before rejecting the logon attempt.

Place the PDC emulator in a location that contains a large number of users from that domain for password forwarding operations if needed. In addition, ensure that the location is well connected to other locations to minimize replication latency.

Use the Domain Controller Placement worksheet to document the information about where you plan to place PDC emulators and the number of users for each domain that is represented in each location. For an example of a completed Domain Controller Placement worksheet, see "Example: Determining Domain Controller Placement" later in this chapter.

You need to refer to the information about locations in which you need to place PDC emulators when you deploy regional domains. For information about deploying regional domains, see "Deploying Windows Server 2003 Regional Domains" in this book.

**Requirements for infrastructure master placement**

The infrastructure master updates the names of security principals from other domains that are added to groups in its own domain. For example, if a user from one domain is a member of a group in a second domain and the user's name is changed in the first domain, then the second domain is not notified that the user's name must be updated in the group's membership list. Because domain controllers in one domain do not replicate security principals to domain controllers in another domain, the second domain never becomes aware of the change in the absence of the infrastructure master.

The infrastructure master constantly monitors group memberships, looking for security principals from other domains. If it finds one, it checks with the security principal's domain to verify that the information is updated. If the information is out of date, the infrastructure master performs the update and then replicates the change to the other domain controllers in its domain.

Two exceptions apply to this rule. First, if all domain controllers are global catalog servers, the domain controller that hosts the infrastructure master role is insignificant because global catalogs replicate the updated information regardless of the domain to which they belong. Second, if the forest has only one domain, the domain controller that hosts the infrastructure master role is insignificant because security principals from other domains do not exist.

Do not place the infrastructure master on a domain controller that is also a global catalog server. If the infrastructure master and global catalog are on the same domain controller, the infrastructure master will not function. The infrastructure master will never find data that is out of date, so it will never replicate any changes to the other domain controllers in the domain.

**Operations Master Placement for Networks with Limited Connectivity**

Be aware that if your environment does have a central location or hub site in which you can place operations master role holders, certain domain controller operations that depend on the availability of those operations master role holders might be affected.

For example, suppose that an organization creates sites A, B, C, and D. Site links exist between A and B, between B and C, and between C and D. Network connectivity exactly mirrors the network connectivity of the sites links. In this example, all operations master roles are placed in site A and the option to **Bridge all site links** is not selected.

Although this configuration results in successful replication between all of the sites, the operations master role functions have the following limitations:

- Domain controllers in sites C and D cannot access the PDC emulator in site A to update a password or to check it for a password that has been recently updated.
- Domain controllers in sites C and D cannot access the RID master in site A to obtain an initial RID pool after Active Directory installation and to refresh RID pools thereafter as they become depleted.
- Domain controllers in sites C and D cannot add or remove directory, Domain Name System (DNS), or custom application partitions.
- Domain controllers in sites C and D cannot make schema changes.

Document the information about operations master roles placement. You need to refer to this information when you create the forest root domain and regional domains. For more information about deploying the forest root domain, see "Deploying the Windows Server 2003 Forest Root Domain" in this book. For more information about deploying regional domains, see "Deploying Windows Server 2003 Regional Domains" in this book.

Use the Domain Controller Placement worksheet to document information about where you plan to place operations master role holders. For an example of a completed Domain Controller Placement worksheet, see "Example: Determining Domain Controller Placement" later in this chapter.

For a worksheet to assist you in planning operations master role placement, see "Domain Controller Placement" (DSSTOPO_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Controller Placement" on the Web at http://www.microsoft.com/reskit).

# Example: Determining Domain Controller Placement

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 3.12 shows an example of a completed Domain Controller Placement worksheet for Trey Research, which has offices located in Seattle, New York, Los Angeles, Boston, Phoenix, and Washington, DC. Users from three domains exist in the organization: the forest root domain, and the EAST and WEST domains. Seattle is the hub location for Western locations and New York is the hub location for Eastern locations. For each location, Trey Research documented the names of the domains, the number of users for each domain, and the type of domain controllers needed.

**Figure 3.12 Example of a Domain Controller Placement Worksheet**

| Domain Controller Placement | | | | |
|---|---|---|---|---|
| Prepared by | Site Topology Owner | | Date | 1/1/03 |
| Forest name | Trccorp.treyresearch.net | | | |

| Location | Domains in Location | Users per Domain per Location | DC Needed? (Yes/No) | Domain Controller Type |
|---|---|---|---|---|
| Seattle | trccorp.treyresearch.net | 20 | Yes | Forest root<br>Global catalog |
| | west.trccorp.treyresearch.net | 15,000 | Yes | PDC emulator (WEST)<br>Global catalog<br>WEST regional |
| | east.trccorp.treyresearch.net | 200 | Yes | EAST regional |
| Los Angeles | west.trccorp.treyresearch.net | 5,000 | Yes | WEST regional<br>Global catalog |
| Washington DC | east.trccorp.treyresearch.net | 50 | Yes | EAST regional<br>Universal group membership caching enabled |
| New York | trccorp.treyresearch.net | 10 | Yes | Forest root |
| | east.trccorp.treyresearch.net | 7,000 | Yes | PDC emulator (EAST)<br>Global catalog<br>EAST regional |
| | west.trccorp.treyresearch.net | 200 | Yes | WEST regional |
| Boston | east.trccorp.treyresearch.net | 80 | Yes | EAST regional<br>Universal group membership caching enabled |
| Phoenix | west.trccorp.treyresearch.net | 20 | No | |

The following decisions were made regarding placement of domain controllers at each location based on WAN link speed between locations, number of users per domain, the need for users to access resources across the forest, and logon performance over WAN links:

- The PDC emulator for the WEST domain is placed in Seattle because it includes the largest number of users from the WEST domain. The PDC emulator for the EAST domain is placed in New York because it includes the largest number of users from the EAST domain.

- Because Trey Research includes a large number of users that are distributed across different geographic locations connected by a wide area network (WAN), two regional domains are created (EAST and WEST) to reduce replication traffic over slow WAN links. Regional domain controllers hosting the WEST domain are placed in Seattle, Los Angeles, and New York. Regional domain controllers hosting the EAST domain are placed in New York, Boston, Seattle, and Washington, DC. At any given time, an average of 200 mobile users from the WEST domain travel from Seattle to the New York location. Therefore, regional domain controllers hosting the WEST domain are placed in New York so that these mobile users belonging to the WEST domain can log on locally at the New York location. Similarly, at any given time, an average of 200 mobile users from New York travel to the Seattle location. Therefore, regional domain controllers hosting the EAST domain are placed in Seattle so that mobile users belonging to the EAST domain can log on locally in Seattle.

  Because domain controllers from both EAST and WEST domains are placed into the Seattle and New York locations

respectively, the resulting network traffic for replication is similar to the replication traffic created if Trey Research had deployed a single domain forest. However, because there are additional locations (Los Angeles, Phoenix, Boston, and Washington, DC) that are connected to the Seattle and New York hub locations through slower network links, and none of the users from these satellite locations will travel to these hub locations, Trey Research will still save network bandwidth caused by replication.

Because Phoenix has only 20 users from the WEST domain and users at the Phoenix location have acceptable logon performance over the WAN link between Phoenix and Seattle, the organization decides not to place any regional domain controllers in Phoenix. Users in Boston require local authentication at all times but do not have 100 percent WAN link availability between New York and Boston. Therefore, a domain controller hosting the EAST domain is placed at the Boston location.

- The TRCCORP forest root domain controllers are placed in Seattle and New York because they are hub locations in the Trey Research forest. Shortcut trusts are created between the WEST domain and the EAST domain because users in Boston need to regularly access resources from the WEST domain.

- Global catalog servers are placed in Seattle, Los Angeles, and New York because each location hosts more than one hundred users. A domain controller with universal group caching enabled is placed in Boston instead of a global catalog server because that location has only 80 users and does not include any applications that require a global catalog server.

- Washington, DC hosts only 50 users from the EAST domain, and does not have 100 percent WAN link availability between New York. Additionally, no applications that require a global catalog server are running at this location. Because Washington, DC has less than 100 users and also no application that requires a global catalog server is running at this location, no global catalog server is placed in Washington, DC. The users in Washington, DC need local authentication at all times, hence a domain controller running Windows Server 2003 is configured for universal group caching and placed there to facilitate user logon requests.

# Creating a Site Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Creating a site design involves deciding which locations will become sites, creating site objects, creating subnet objects, and associating the subnets with sites. Figure 3.13 shows the steps that are required to create a site design.

**Figure 3.13 Creating a Site Design**




**Deciding Which Locations Will Become Sites**

Figure 3.14 illustrates how to decide which locations will become sites.

**Figure 3.14 Deciding Which Locations Will Become Sites**




Decide which locations to create sites for as follows:

- Create sites for all locations in which you plan to place domain controllers. Refer to the information documented in the Domain Controller Placement worksheet to identify locations that include domain controllers. For an example of a completed Domain Controller Placement worksheet, see "Example: Determining Domain Controller Placement" earlier in this chapter.

- Create sites for those locations that include servers that are running applications that require a site to be created. Certain applications, such as DFS, use site objects to locate the closest servers to clients.

- If a site is not required for a location, add the subnet of the location to a site for which the location has the maximum WAN speed and available bandwidth.

Document those locations that will become sites and the network addresses and subnet masks within each location. For a worksheet to assist you in documenting sites, see "Associating Subnets with Sites" (DSSTOPO_6.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Associating Subnets with Sites" on the Web at http://www.microsoft.com/reskit).

**Creating a Site Object Design**
For every location where you have decided to create sites, plan to create site objects in Active Directory. Document those locations that will become sites in the "Associating Subnets with Sites" worksheet.

For more information about how to create site objects, see "Create a site" in Help and Support Center for Windows Server 2003.

**Creating a Subnet Object Design**
For every IP subnet and subnet mask associated with each location, plan to create subnet objects in Active Directory representing all the IP addresses within the site. To obtain information about IP subnets within each location, refer to the Locations and Subnets worksheet. For an example of a completed Locations and Subnets worksheet, see "Listing IP Subnets Within Each Location" earlier in this chapter.

When creating an Active Directory subnet object, the information about network IP subnet and subnet mask is automatically translated into the network prefix length notation format *<IP address>/<# of bits>*. For example, the network IP address 172.16.4.0 with a subnet mask 255.255.252.0 appears as 172.16.4.0/22.

For more information about how to create subnet objects, see "Create a subnet" in Help and Support Center for Windows Server 2003.

Document the Active Directory subnet object associated with each location in the "Associating Subnets with Sites" worksheet.

**Associating Subnets with Sites**
Associate each subnet object with a site object by referring to the "Associating Subnets with Sites" worksheet to determine which subnet is to be associated with which site.

For more information about associating subnets with sites, see "Associate a subnet with a site" in Help and Support Center for Windows Server 2003.

**Example: Associating Subnets with Sites**
Figure 3.15 shows an example of a completed Associating Subnets with Sites worksheet for Trey Research, which has offices located in Seattle, Los Angeles, Boston, New York, Phoenix, and Washington, DC. Trey Research listed potential site candidates, the locations to be included in those sites, and the corresponding Active Directory subnets. Because the Phoenix location includes fewer than one hundred users, a domain controller is not placed at that location. The Phoenix subnets are associated with the Seattle site because Seattle is the nearest location that has a direct communication link to Phoenix.

**Figure 3.15 Example of Associating Subnets with Sites Worksheet**

| Associating Subnets with Sites | | | |
|---|---|---|---|
| Prepared By | Site Topology Owner | Date | 1/1/03 |
| Forest Name | Trccorp.treyresearch.net | | |

| Site Candidate | Location | Active Directory Subnets |
|---|---|---|
| Seattle | Seattle<br><br>Phoenix | 172.16.16.0/22<br>172.16.24.0/22<br><br>172.16.64.0/22 |
| Boston | Boston | 172.16.36.0/22<br>172.16.40.0/22 |
| New York | New York | 172.16.28.0/22<br>172.16.32.0/22 |
| Los Angeles | Los Angeles | 172.16.136.0/22<br>172.16.92.0/22 |
| Washington DC | Washington DC | 172.16.116.0/22<br>172.16.140.0/22 |

# Creating a Site Link Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Create a site link design in order to connect your sites with site links. Site links reflect the intersite connectivity and method used to transfer replication traffic. You must connect sites with site links so that domain controllers at each site can replicate Active Directory changes. Figure 3.16 shows the process for creating a site link design.

**Figure 3.16 Creating a Site Link Design**



# Connecting Sites with Site Links

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To connect sites with site links, identify the member sites that you want to connect with the site link, create a site link object in the respective **Inter-Site Transports** container, and then name the site link. After you create the site link, you can proceed to set the site link properties.

When creating site links, ensure that every site is included in a site link. In addition, ensure that all sites are connected to each other through other site links so that the changes can be replicated from domain controllers in any site to all other sites. If you fail to do this, then the KCC generates an error message in the Directory Service log in Event Viewer stating that the site topology is not connected.

Whenever you add sites to a newly created site link, determine if the site being added is a member of other site links and change the site link membership of the site if needed. For example, if you make a site a member of the default-first-site-link when you initially create the site, be sure to remove the site from the default-first-site-link after you add the site to a new site link. If you do not remove the site from the default-first-site-link, the KCC will make routing decisions based on the membership of both site links, which may result in incorrect routing.

To identify the member sites that you want to connect with a site link, use the list of locations and linked locations that you recorded in the "Geographic Locations and Communication Links" worksheet. If multiple sites have the same connectivity and availability to each other, you can connect them with the same site link. For an example of a completed Geographic Locations and Communication Links worksheet, see "Listing Communication Links and Available Bandwidth" earlier in this chapter.

The Inter-Site Transports container provides the means for mapping site links to the transport that the link uses. When you create a site link object, you create it in either the IP container (which associates the site link with the remote call procedure [RPC] over IP transport) or the Simple Mail Transfer Protocol (SMTP) container (which associates the site link with the SMTP transport). When you create a site link object in the respective **Inter-Site Transports** container, Active Directory uses RPC over IP to transfer both intersite and intrasite replication between domain controllers. To keep data secure while in transit, RPC over IP replication uses both the Kerberos authentication protocol and data encryption.

When a direct IP connection is not available, you can configure replication between sites to use SMTP. However, SMTP replication functionality is limited and requires an enterprise certification authority (CA). SMTP can only replicate the configuration, schema, and application directory partitions, and does not support the replication of domain directory partitions.

To name site links, use a consistent naming scheme, such as *name_of_site1-name_of_site2*. Record the list of sites, linked sites, and the names of the site links connecting these sites in a worksheet.

For a worksheet to assist you in recording site names and associated site link names, see "Sites and Associated Site Links" (DSSTOPO_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Sites and Associated Site Links" on the Web at http://www.microsoft.com/reskit).

Figure 3.17 shows an example of a completed "Sites and Associated Site Links" worksheet.

**Figure 3.17 Example of a Sites and Associated Site Links Worksheet**

| Sites and Associated Site Links | | | |
|---|---|---|---|
| Prepared By | Site Topology Owner | Date | 1/1/03 |
| Forest Name | Trccorp.treyresearch.net | | |

| Site Names | Linked Sites | Name of the Site Link | Link Type |
|---|---|---|---|
| Seattle | LosAngeles | LAX-SEA | ISDN (128.8 Kbps) |
| Seattle | NewYork | SEA-NYC | T1 (1.544 Mbps) |
| NewYork | WashingtonDC | NYC-WDC | T1 (1.544 Mbps) |
| NewYork | Boston | NYC-BOS | ISDN (128.8 Kbps) |

# Setting Site Link Properties

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Intersite replication occurs according to the properties of the connection objects. When the KCC creates connection objects, it derives the replication schedule from properties of the site link objects. Each site link object represents the WAN connection between two or more sites.

Setting site link object properties includes the following steps:

- Determining the cost that is associated with that replication path. The KCC uses cost to determine the least expensive route for replication between two sites that replicate the same directory partition.
- Determining the schedule that defines the times during which intersite replication can occur.
- Determining the replication interval that defines how frequently replication should occur during the times when replication is allowed as defined in the schedule.

# Determining the Cost

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You assign cost values to site links to favor inexpensive connections over expensive connections. Certain applications and services, such as Domain Controller Locator and DFS, also use cost information to locate nearest resources. Site link cost can be used to determine which domain controller is contacted by clients located in one site if the domain controller for the specified domain does not exist at that site. The client contacts the domain controller by using the site link that has the lowest cost assigned to it.

It is recommended that the cost value be defined on a site-wide basis. Cost is usually based not only on the total bandwidth of the link but also on the availability, latency, and monetary cost of the link.

To determine the costs to place on site links, perform these tasks:

- Document the connection speed for each site link. Refer to the Geographic Locations and Communication Links worksheet for information about the connection speed that you identified. For an example of a completed Geographic Locations and Communication Links worksheet, see "Listing Communication Links and Available Bandwidth" earlier in this chapter.

  Table 3.1 lists the speeds for different types of networks.

  **Table 3.1 Speeds for Different Network Types**

| Network Type | Speed |
|---|---|
| Very slow | 56 kilobits per second (Kbps) |
| Slow | 64 Kbps |
| ISDN | 64 Kbps or 128 Kbps |
| Frame relay | Variable rate, commonly between 56 Kbps and 1.5 megabits per second (Mbps) |
| T1 | 1.5 Mbps |
| T3 | 45 Mbps |
| 10BaseT | 10 Mbps |
| Asynchronous transfer mode (ATM) | Variable rate, commonly between 155 Mbps and 622 Mbps |
| 100BaseT | 100 Mbps |
| Gigabit Ethernet | 1 gigabit per second (Gbps) |

- Use Table 3.2 to calculate the cost of each site link based on WAN link speed. For WAN link speed that is not listed in the table, you can calculate a relative cost factor by dividing 1024 by the log of the available bandwith as measured in Kbps.

  **Table 3.2 Calculating Costs Based on WAN Link Speed**

| Available Bandwidth (Kbps) | Cost |
|---|---|
| 9.6 | 1042 |

| | |
|---|---|
| 19.2 | 798 |
| 38.4 | 644 |
| 56 | 586 |
| 64 | 567 |
| 128 | 486 |
| 256 | 425 |
| 512 | 378 |
| 1024 | 340 |
| 2048 | 309 |
| 4096 | 283 |

These costs do not reflect differences in reliability between network links. Set higher costs on any failure-prone network links so that you do not need to rely on those links for replication. By setting higher site links costs, you can control replication failover when a site link fails.

# Determining the Schedule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can control site link availability by setting a schedule on site links. When replication between two sites traverses multiple site links, the intersection of the replication schedules on all relevant links determines the connection schedule between the two sites.

To plan for setting the site link schedule, create two overlapping schedules between site links that contain domain controllers that need to directly replicate with each other. Use the default (100-percent available) schedule on those links unless you want to block replication traffic during peak hours. By blocking replication, you give priority to other traffic, but you also increase replication latency.

Domain controllers store time in Coordinated Universal Time (UTC). Time settings in site link object schedules conform to the local time of the site and computer on which the schedule is set. When a domain controller contacts a computer that is in a different site and time zone, the schedule on the domain controller displays the time setting according to the local time for the site of the computer.

# Determining the Interval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must set the site link replication interval property to indicate how frequently you want replication to occur during the times when the schedule allows replication. For example, if the schedule allows replication between 02:00 hours and 04:00 hours, and the replication interval is set for 30 minutes, replication can occur up to four times during the scheduled time. The default replication interval is 180 minutes, or 3 hours.

Consider the following criteria to determine how often replication occurs within the schedule window:

- A small interval decreases latency but increases the amount of WAN traffic.
- To keep domain directory partitions up to date, low latency is preferred.

With a store-and-forward replication strategy, determining just how long a directory update might take to be replicated to every domain controller is difficult. To provide a conservative estimate of maximum latency, perform these tasks:

- Create a table of all the sites on your network, as shown in Table 3.3.

  **Table 3.3 Maximum Latency (in Hours) of Replication Between Sites**

| | Seattle | Boston | Los Angeles | New York | Washington DC |
|---|---|---|---|---|---|
| **Seattle** | 0.25 | | | | |
| **Boston** | | 0.25 | | | |
| **Los Angeles** | | | 0.25 | | |
| **New York** | | | | 0.25 | |
| **Washington DC** | | | | | 0.25 |

  An average, worst-case latency within a site is estimated to be 15 minutes.

- From the replication schedule, determine the maximum replication latency possible on any site link connecting two hub sites.

  For example, if replication occurs between Seattle and New York every three hours, then the maximum delay for replication between these sites is 3 hours. If the replication delay between New York and Seattle is the longest scheduled delay among all hub sites, then the maximum latency between all hubs is 3 hours.

- For each hub site, create a table of the maximum latencies between the hub site and any of its satellite sites.

  For example, if replication occurs between New York and Washington, DC, every four hours and this is the longest replication delay between New York and any of its satellite sites, then the maximum latency between New York and its satellites is four hours.

- Combine these maximum latencies to determine the maximum latency for the entire network.

  For example, if the maximum latency between Seattle and its satellite site in Los Angeles is one day, the maximum replication latency for this set of links (Washington, DC-New York-Seattle-Los Angeles) is 31 hours (4 (Washington, DC-New York)+ 3 (New York-Seattle) + 24 (Seattle-Los Angeles)) as shown in Table 3.4.

  **Table 3.4 Maximum Latency (in Hours) of Replication Between Sites**

| | Seattle | Boston | Los Angeles | New York | Washington DC |
|---|---|---|---|---|---|
| **Seattle** | 0.25 | 4+3 | 24.00 | 3.00 | 4+3 |
| **Boston** | | 0.25 | 4+3+24 | 4.00 | 4.00 |
| **Los Angeles** | | | 0.25 | 24 + 3 | 24+3+4 |
| **New York** | | | | 0.25 | 4.00 |
| **Washington, DC** | | | | | 0.25 |

Review your maximum replication latency, and revise your replication schedule and interval if necessary.

# Example: Creating a Site Link Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As an example of a site link design, Figure 3.18 shows three domain controllers from the same domain at three different Trey Research sites: Los Angeles, New York, and Seattle. The LAX-SEA and SEA-NYC site link objects are created to permit replication between domain controllers at the three sites. The site link object LAX-SEA is assigned a cost of 586, whereas the site link object SEA-NYC is assigned a lower cost of 471 because the T1 communication link that connects the Seattle and New York sites has a more available bandwidth (150 Kbps) than the 56 Kbps available bandwidth between the Seattle and Los Angeles sites.

Active Directory replication topology for the example in Figure 3.18 is computed by the KCC according to a spanning tree algorithm. For intersite replication, a spanning tree algorithm determines how to connect all the sites that need to be connected with the minimum number of connections and the least cost. For the example shown in Figure 3.18, the cost of replication between domain controller DC-2 and domain controller DC-1 is 586, the cost of replication between domain controller DC-1 and domain controller DC-3 is 471. Because no direct communication link exists between the New York and Los Angeles sites, the cost of replication between these two sites is equal to the sum of the replication costs between the Los Angeles and Seattle sites and the Seattle and New York sites (471 + 586 = 1,057).

Based on these calculated costs, the domain controllers in these three sites can replicate the domain information by creating three replication connections: DC-2 replicates with DC-1 (cost =586), DC-3 replicates with DC-1 (cost = 471), and DC-2 replicates with DC-3 (cost = 1,057). Two replication connections between DC-1 and DC-2 and DC-1 and DC-3 are sufficient for replicating the domain information between the three sites. Hence, based on the spanning tree algorithm, connection objects are created between DC-1 and DC-2, and between DC-1 and DC-3.

No site link is created between the Los Angeles and New York sites by the site topology owner because no direct communication link exists between these two locations. Even if a third site link with a cost of 1,057 is created between the Los Angeles and New York sites, no additional connection objects will be created between domain controllers DC-2 and DC-3.

With transitivity enabled in this example, DC-2 at Los Angeles always replicates with DC-1 (cost = 586) in Seattle and not DC-3 (cost = 586 + 471) in New York because the LAX-SEA site link cost is more favorable. Replication connections between DC-2 and DC-3 will only be created if DC-1 at the Seattle site is unavailable.

**Figure 3.18 Example of Site Links**




However, with transitivity disabled in this example, because no direct communication link exists between the Los Angeles and New York sites, DC-2 will only replicate with DC-1 and never replicate with DC-3 even if DC-1 in Seattle fails.

With site link transitivity enabled, if DC-2 in the Los Angeles site and DC-3 in the New York site are from a different domain (domain A) than DC-1 in the Seattle site (domain B), and if the LAX-SEA site link has a schedule of 18:00 hours to 24:00 hours and the SEA-NYC site link has a schedule of 17:00 hours to 20:00 hours, data can be replicated directly from DC-2 to DC-3 between 18:00 hours and 20:00 hours, which is the intersection of the LAX-SEA site link replication schedule and the SEA-NYC site link replication schedule. If the LAX-SEA site link schedule does not overlap with the SEA-NYC site link schedule, then replication between DC-2 and DC-3 can never occur.

# Creating a Site Link Bridge Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A site link bridge connects two or more site links. Figure 3.19 shows the steps for creating a site link bridge design.

**Figure 3.19 Creating a Site Link Bridge Design**




A site link bridge enables transitivity between site links. Each site link in a bridge must have a site in common with another site link in the bridge. The Knowledge Consistency Checker (KCC) uses the information on each site link to compute the cost of replication between sites in one site link and sites in the other site links of the bridge. Without the presence of a common site between site links, the KCC also cannot establish direct connections between domain controllers in the sites that are connected by the same site link bridge.

By default, all site links are transitive and it is recommended to keep transitivity enabled by not changing the default value of **Bridge all site links** (enabled by default). However, you will need to disable **Bridge all site links** and complete a site link bridge design if:

- Your IP network is not fully routed. When you disable **Bridge all site links**, all site links are considered nontransitive, and you can create and configure site link bridge objects to model the actual routing behavior of your network.

  –Or–

  You need to control the replication flow of the changes made in Active Directory. By disabling **Bridge all site links** for the site link IP transport and configuring a site link bridge, the site link bridge becomes the equivalent of a disjointed network. All site links within the site link bridge can route transitively, but they do not route outside of the site link bridge.

For more information about how to use Active Directory Sites and Services to disable the **Bridge all site links** setting, see "Enable or disable site link bridges" in Help and Support Center for Windows Server 2003.

# Creating a Site Link Bridge Design for Disjointed Networks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your IP network is not fully routed, you must disable **Bridge all site links** for the IP transport and configure site link bridge objects. If your IP network is fully routed but you have too many routes that you categorically do not want the KCC to consider, you also need to disable **Bridge all site links** and manually create site link bridges.

For more information about how to use the Active Directory Sites and Services snap-in to disable the **Bridge all site links** setting, see "Enable or disable site link bridges" in Help and Support Center for Windows Server 2003.

# Creating a Site Link Bridge Design to Control Active Directory Replication Flow

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Two scenarios where you might want to control replication flow include controlling replication failover and controlling replication through a firewall.

**Controlling Replication Failover**

If your organization has a hub-and-spoke network topology, you generally do not want the satellite sites to create replication connections to other satellite sites if all domain controllers in the hub site fail. In such scenarios, you must disable **Bridge all site links** and create site link bridges such that replication connections are created between the satellite site and another hub site that is just one or two hops away from the satellite site.

**Example: Creating a site link bridge design to control replication failover**

Figure 3.20 illustrates an organization's hub-and-spoke network topology, consisting of two hub sites (A and B) and six satellite sites (C through H). The site links between all sites are named A-B, A-C, A-D, A-E, B-F, B-G, and B-H.

**Figure 3.20 Creating a Site Link Bridge to Control Replication Failover**




With **Bridge All Site Links** enabled, if DC-1 in hub site A fails, DC-3 can create replication connections between DC-2, DC-4, DC-5, DC-6, DC-7 or DC-8. The only intended connection for DC-3 in this case is with DC-2 in hub site B, as this connection will not congest network traffic. To ensure that DC-3 only replicates with DC-2 when DC-1 fails, you can disable site link transitivity and create the AC-AB site link bridge between the A-C and A-B site links. Similarly, you can create the AB-BH site link bridge between the A-B and B-H site links. With the AC-AB site link bridge, if DC-1 is unavailable, the replication connections from DC-3 are created only with DC-2 in hub site B. With the AB-BH site link bridge, if DC-2 is unavailable, the replication connections from DC-8 are created only with DC-1 in hub site A. Create similar site link bridges between all satellite sites so that the KCC failover occurs only to the domain controllers at the hub sites.

**Controlling Replication Through a Firewall**

If two domain controllers representing the same domain in two different sites are specifically allowed to communicate with each other only through a firewall, you can disable **Bridge all Site Links** and create site link bridges for sites on the same side of the firewall.

**Example: Creating a site link bridge design to control replication through a firewall**

Figure 3.21 illustrates a similar scenario as shown in Figure 3.20, but in this case a firewall separates hub site A and hub site B. IPSec is used to replicate through the firewall. The IPSec policy allows only DC-1 and DC-2 to communicate through the firewall by using IPSec. With transitivity enabled, if DC-1 fails, then DC-3 in satellite site C can create replication connections not only with DC-4 and DC-5 as intended, but it can try to create replication connections across the firewall with DC-2, DC-6, DC-7, or DC-8.

To ensure that DC-3 directly replicates only on one side of the firewall, you can disable transitivity and create the WEST site link bridge, which includes site links A-C, A-D, and A-E. Creating the WEST site link bridge ensures that if domain controller DC-1

fails, DC-3, DC-4, and DC-5 directly replicate only with each other and not with DC-2, DC-6, DC-7, or DC-8. Similarly, you can create the EAST site link bridge, which includes site links B-F, B-G, and B-H, to ensure that if DC-2 fails, then DC-6, DC-7, or DC-8 directly replicate only with each other.

**Figure 3.21 Creating Site Link Bridges to Control Replication Through a Firewall**




Therefore, if your network is separated by firewalls, it is recommended that you disable transitivity of site links and create site link bridges for the network on one side of the firewall. For information about managing replication through firewalls, see "Active Directory in Networks Segmented by Firewalls" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=37928).

# Additional Resources for Designing the Site Topology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing the Active Directory Logical Structure" in this book.
- "Planning Domain Controller Capacity" in this book.
- "Deploying the Windows Server 2003 Forest Root Domain" in this book.
- "Deploying Windows Server 2003 Regional Domains" in this book.
- "Deploying DNS" in *Deploying Network Services* of this kit.
- Windows Server 2003 Active Directory Branch Office Guide on the Web (http://go.microsoft.com/fwlink/?LinkId=46583).
- Solutions for Branch Offices on the Web (http://go.microsoft.com/fwlink/?LinkId=34284).

**Related Tools**

- Adlb.exe

  You can use ADLB to improve replication efficiency in forests set to Windows Server 2003 functional level. For more information about ADLB and to download this tool, see the Windows Server 2003 Active Directory Branch Office Guide on the Web (http://go.microsoft.com/fwlink/?LinkId=46583).

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Active Directory" in Help and Support Center for Windows Server 2003.
- "Enable printer location tracking" in Help and Support Center for Windows Server 2003.
- "Create a shortcut trust" in Help and Support Center for Windows Server 2003.
- “Cache universal group memberships” in Help and Support Center for Windows Server 2003.
- "Create a site" in Help and Support Center for Windows Server 2003.
- "Create a subnet" in Help and Support Center for Windows Server 2003.
- "Associate a subnet with a site" in Help and Support Center for Windows Server 2003.
- "Enable or disable site link bridges" in Help and Support Center for Windows Server 2003.

**Related Job Aids**

- "Geographic Locations and Communication Links" (DSSTOPO_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Geographic Locations and Communication Links" on the Web (http://go.microsoft.com/fwlink/?

linkid=9565).

- "Locations and Subnets" (DSSTOPO_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Locations and Subnets" on the Web (http://go.microsoft.com/fwlink/?linkid=9566).

- "Domains and Users in Each Location" (DSSTOPO_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domains and Users in Each Location" on the Web (http://go.microsoft.com/fwlink/?linkid=9567).

- "Domain Controller Placement" (DSSTOPO_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Controller Placement" on the Web (http://go.microsoft.com/fwlink/?linkid=9568).

- "Sites and Associated Site Links" (DSSTOPO_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Sites and Associated Site Links" on the Web (http://go.microsoft.com/fwlink/?linkid=9569).

- "Associating Subnets with Sites" (DSSTOPO_6.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Associating Subnets with Sites" on the Web (http://go.microsoft.com/fwlink/?linkid=10186).

# Planning Domain Controller Capacity

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Overview of Planning Domain Controller Capacity

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Process for Planning Domain Controller Capacity

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Background Information for Planning Domain Controller Capacity

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Collecting Site Topology Design Information

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Determining the Number of Domain Controllers

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Determining the Minimum Number of Domain Controllers Required

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Adding Domain Controllers to Support Replication Between Sites

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Assessing Disk Space and Memory Requirements

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Determining Required Disk Space

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Determining Minimum Disk Space Requirements

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Adding Disk Space for Global Catalog Servers

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Adding Disk Space for Application Directory Partitions

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Determining Required Memory Allocation

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Example: Assessing Disk Space and Memory Requirements

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Monitoring Domain Controller Performance

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Additional Resources for Planning Domain Controller Capacity

Updated: February 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity planning information for domain controllers has been updated and published at Capacity Planning for Active Directory Domain Services.

# Planning for Active Directory Forest Recovery

Updated: April 25, 2013

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

This guide contains best-practice recommendations for recovering an Active Directory® forest if forest-wide failure renders all domain controllers (DCs) in the forest incapable of functioning normally. The steps it contains serve as a template for your forest recovery plan, which you can customize for your particular environment. These steps apply to DCs that run Microsoft® Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003 operating systems.

| **Note** |
| --- |
| Procedures that are unique for DCs that run Windows Server 2003 are consolidated in Appendix D. |

**In this guide**


- New Features, Assumptions, and Prerequisites for Using This Guide for Planning Active Directory Forest Recovery
- Devising a Custom Forest Recovery Plan
- Recovering Your Active Directory Forest
- Appendix A: Forest Recovery Procedures
- Appendix B: Frequently Asked Questions
- Appendix C: Recovering a Single Domain within a Multidomain Forest
- Appendix D: Forest Recovery with Windows Server 2003 Domain Controllers
- Additional Resources


**Publication and revision history**

The following table summarizes the revision history for this guide, including its original publication on Microsoft TechNet.

| Date | Revision |
| --- | --- |
| April 2013 | Updated to reflect new features in Windows Server 2012. |
| May 2009 | Updated to reflect new features in Windows Server 2008. |
| October 2006 | Updated for Windows Server 2003 and published on TechNet |
| June 2002 | Original publication on Microsoft Download Center |

# New Features, Assumptions, and Prerequisites for Using This Guide for Planning Active Directory Forest Recovery

Updated: April 25, 2013

Applies To: Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

This topic describes how the new virtualized domain controller cloning feature in Windows Server 2012 improves the forest recovery process and other issues to review before you use the guidelines.

- Using virtualized domain controller cloning in Windows Server 2012 to expedite forest recovery
- Assumptions for Using This Guide for Planning Active Directory Forest Recovery
- Prerequisites for Using This Guide for Planning Active Directory Forest Recovery

**Using virtualized domain controller cloning in Windows Server 2012 to expedite forest recovery**
Virtualized domain controller (DC) cloning simplifies and expedites the process for installing additional virtualized DCs in a domain, especially in centralized locations such as datacenters where several DCs run on hypervisors. After you restore one virtual DC in each domain from backup, additional DCs in each domain can be rapidly brought online by using the virtualized DC cloning process. You can prepare the first virtualized DC that you recover, shut it down, and then copy that virtual hard disk as many times as is necessary in order to create cloned virtualized DCs to build out the domain.

The requirements for virtualized DC cloning are:

- The hypervisor must support VM-GenerationID. Hyper-V in Windows Server 2012 and Windows 8 is an example of a hypervisor that supports VM-GenerationID. Check with your hypervisor vendor if VM-GenerationID is supported.
- The virtualized DC that is used as a source for cloning must run Windows Server 2012 and be a member of the Cloneable Domain Controllers group.
- The PDC emulator must run Windows Server 2012. You can clone PDC emulator if it is virtualized.

For step-by-step instructions about how to perform virtualized DC cloning, see Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100). For details about how virtualized DC cloning works, see Virtualized Domain Controller Technical Reference (Level 300).

**Assumptions for Using This Guide for Planning Active Directory Forest Recovery**
First, this guide assumes that you have:

- Worked with Microsoft Support to determine the cause of the forest-wide failure. This guide does not suggest a cause of the failure or recommend any procedures to prevent the failure.
- Evaluated any possible remedies.
- Concluded, in consultation with Microsoft Support, that restoring the whole forest to its state before the failure occurred is the best way to recover from the failure. In many cases, forest recovery should be the last option.

Second, this guide assumes that you have followed the Microsoft best-practice recommendations for using Active Directory–integrated Domain Name System (DNS). Specifically, there should be an Active Directory–integrated DNS zone for each Active Directory domain. If this is not the case, you can still use the basic principles of this guide to perform forest recovery. However, you will need to take specific measures for DNS recovery based on your own environment. For more information about using Active Directory–integrated DNS, see Designing a DNS Infrastructure to Support Active Directory.

Although the objectives of this guide are to recover the forest and maintain or restore full DNS functionality, recovery can result in a DNS configuration that is changed from the configuration before the failure. After the forest is recovered, you can revert to the original DNS configuration. The recommendations in this guide do not describe how to configure DNS servers to perform name resolution of other portions of the corporate namespace where there are DNS zones that are not stored in

AD DS.

Finally, although this guide is intended as a generic guide for forest recovery, not all possible scenarios are covered. For instance, beginning with Windows Server 2008, there is a Server Core version, which is a full version of Windows Server but without a full GUI. Although it is certainly possible to recover a forest consisting of just DCs that run Server Core, this guide has no detailed instructions. However, based on the guidance discussed here you will be able to design the required command-line actions yourself.

**Prerequisites for Using This Guide for Planning Active Directory Forest Recovery**
Before you begin planning for recovery of an Active Directory forest, you should be familiar with the following:

- Fundamental Active Directory concepts
- The importance of operations master roles (also known as flexible single master operations or FSMO). These roles include the following:
    - Schema master
    - Domain naming master
    - Relative ID (RID) master
    - Primary domain controller (PDC) emulator master
    - Infrastructure master

In addition, you should have backed up and restored AD DS and SYSVOL in a lab environment on a regular basis. For more information, see Backing up the System State data and Performing a nonauthoritative restore of Active Directory Domain Services.

# Devising a Custom Forest Recovery Plan

Updated: April 25, 2013

Applies To: Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

Depending on your environment and business requirements, you might or might not need to perform all the steps described in this guide to perform a successful forest recovery. Given that this guide serves only as a template for forest recovery, it is vital that you devise a custom forest recovery plan that suits your environment and meets your business needs.

For example, in your forest recovery plan, you should have a detailed topology map of your forests. The map should list all the information about the DCs, such as their names, their roles and backup status, and the trust relationships between them. For a tool that you can use to create a topology map, see Microsoft Active Directory Topology Diagrammer.

You should practice your forest recovery plan at least once a year. Also, it is a good idea to perform a forest recovery drill when there are membership changes to the Enterprise Admins or Domain Admins group. This helps ensure that your information technology (IT) staff fully understands the forest recovery plan.

# Recovering Your Active Directory Forest

Updated: April 25, 2013

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

This section provides an overview of the recommended path for recovering a forest. The forest recovery steps are described in detail later.

The following list summarizes the recovery steps at a high level:

1. Identify the problem

   Work with IT and Microsoft Support to determine the scope of the problem and potential causes, and evaluate possible remedies with all business stakeholders. In many cases total forest recovery should be the last option.

2. Decide how to recover the forest

   After you determine that forest recovery is necessary, complete preliminary steps to prepare for it: determine the current forest structure, identify the functions that each DC performs, decide which DC to restore for each domain, and ensure that all writeable DCs are taken offline.

3. Perform initial recovery

   In isolation, recover one DC for each domain, clean them, and reconnect the domains. Reset privileged accounts, and rectify problems caused by security breaches in this phase.

4. Redeploy remaining DCs

   Redeploy the forest to return it to its state before the failure. This step will need to be adapted to your specific design and requirements. Virtualized domain controller cloning can help expedite this process.

5. Cleanup

   After functionality has been restored, reconfigure name resolution as needed, and get LOB applications working.

The following flowchart shows the recovery process.




The steps in this guide are designed to minimize the possibility of reintroducing dangerous data into the recovered forest. You might have to modify these steps to account for such factors as:

- Scalability
- Remote manageability

- Speed of recovery

However, modifications to these forest recovery steps can increase the risk of reintroducing dangerous data. For more information about possible modifications to these forest recovery steps, see What can I do to speed up recovery?

**Identify the problem**

When symptoms of a forest-wide failure appear, such as in event logs or other monitoring solutions, work with Microsoft Support to determine the cause of the failure, and evaluate any possible remedies.

Examples of forest-wide failures include the following:

- All DCs have been logically corrupted or physically damaged to a point that business continuity is impossible; for example, all business applications that depend on AD DS are nonfunctional.
- A rogue administrator has compromised the Active Directory environment.
- An attacker intentionally—or an administrator accidentally—runs a script that spreads data corruption across the forest.
- An attacker intentionally—or an administrator accidentally—extends the Active Directory schema with malicious or conflicting changes.
- An attacker has managed to install malicious software on DCs, and you have been advised by Microsoft Support to recover the forest from backup.

  **Important**

  This paper does not cover security recommendations about how to recover a forest that has been hacked or compromised. In general, it is recommended to follow Pass-the-Hash mitigation techniques to harden the environment. For more information, see Mitigating Pass-the-Hash (PtH) Attacks and Other Credential Theft Techniques.

- None of the DCs can replicate with their replication partners.
- Changes cannot be made to AD DS at any domain controller.
- New DCs cannot be installed in any domain.

**Decide how to recover the forest**

Recovering an entire Active Directory forest involves either restoring it from backup or reinstalling Active Directory Domain Services (AD DS) on every domain controller (DC) in the forest. Recovering the forest restores each domain in the forest to its state at the time of the last trusted backup. Consequently, the restore operation will result in the loss of at least the following Active Directory data:

- All objects (such as users and computers) that were added after the last trusted backup
- All updates that were made to existing objects since the last trusted backup
- All changes that were made to either the configuration partition or the schema partition in AD DS (such as schema changes) since the last trusted backup

For each domain in the forest, the password of a Domain Admin account must be known. Preferably, this is the password of the built-in Administrator account, which must not be disabled. You must also know the DSRM password to perform a system state restore of a DC. In general, it is a good practice to archive the Administrator account and DSRM password history in a safe place for as long as the backups are valid, that is, within the tombstone lifetime period or within the deleted object lifetime period if Active Directory Recycle Bin is enabled. You can also synchronize the DSRM password with a domain user account in order to make it easier to remember. For more information, see KB article 961320. Synchronizing the DSRM account must be done in advance of the forest recovery, as part of preparation.

**Note**

The Administrator account is a member of the built-in Administrators group by default, as are the Domain Admins and Enterprise Admins groups. This group has full control of all DCs in the domain.

**Determining which backups to use**

Back up at least two writeable DCs for each domain regularly so you have several backups to choose from. Note that you cannot use the backup of a read-only domain controller (RODC) to restore a writeable DC. We recommend that you restore the DCs by using backups that were taken a few days before the occurrence of the failure. In general, you must determine a tradeoff between the recentness and the safeness of the restored data. Choosing a more recent backup recovers more useful data, but it might increase the risk of reintroducing dangerous data into the restored forest.

Restoring system state backups depends on the original operating system and server of the backup. For example, you should not restore a system state backup to a different server. In this case, you may see the following warning:

"The specified backup is of a different server than the current one. We do not recommend performing a system state recovery with the backup to an alternate server because the server might become unusable. Are you sure you want to use this backup for recovering the current server?"

If you need to restore Active Directory to different hardware, create full server backups and plan to perform a full server recovery.

**Important**

Beginning with Windows Server 2008, it is not supported to restore system state backup to a new installation of Windows Server on new hardware or the same hardware. If Windows Server is reinstalled on the same hardware, as recommended later in this guide, then you can restore the domain controller in this order:

1. Perform a full server restore in order to restore the operating system and all files and applications.

2. Perform a system state restore using wbadmin.exe in order to mark SYSVOL as authoritative.

For more information, see Microsoft KB article 249694.

If the time of the occurrence of the failure is unknown, investigate further to identify backups that hold the last safe state of the forest. This approach is less desirable. Therefore, we strongly recommend that you keep detailed logs about the health state of AD DS on a daily basis so that, if there is a forest-wide failure, the approximate time of failure can be identified. You should also keep a local copy of backups to enable faster recovery.

If Active Directory Recycle Bin is enabled, the backup lifetime is equal to the **deletedObjectLifetime** value or the **tombstoneLifetime** value, whichever is less. For more information, see Active Directory Recycle Bin Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkId=178657).

As an alternative, you can also use the Active Directory database mounting tool (Dsamain.exe) and a Lightweight Directory Access Protocol (LDAP) tool, such as Ldp.exe or Active Directory Users and Computers, to identify which backup has the last safe state of the forest. The Active Directory database mounting tool, which is included in Windows Server 2008 and later Windows Server operating systems, exposes Active Directory data that is stored in backups or snapshots as an LDAP server. Then, you can use an LDAP tool to browse the data. This approach has the advantage of not requiring you to restart any DC in Directory Services Restore Mode (DSRM) to examine the contents of the backup of AD DS.

For more information about using the Active Directory database mounting tool, see the Active Directory Database Mounting Tool Step-by-Step Guide.

You can also use the **ntdsutil snapshot** command to create snapshots of the Active Directory database. By scheduling a task to periodically create snapshots, you can obtain additional copies of the Active Directory database over time. You can use these copies to better identify when the forest-wide failure occurred and then choose the best backup to restore. To create snapshots, use the version of **ntdsutil** that ships with Windows Server 2008 or the Remote Server Administration Tools (RSAT) for Windows Vista or later. The target DC can run any version of Windows Server. For more information about using the **ntdsutil snapshot** command, see Snapshot.

**Determining which domain controllers to restore**

Ease of the restore process is an important factor when deciding which domain controller to restore. It is recommended to have a dedicated DC for each domain that is the preferred DC for a restore. A dedicated restore DC makes it easier to reliably plan and execute the forest recovery because you use the same source configuration that was used to perform restore tests.

You can script the recovery, and not contend with different configurations, such as whether the DC holds operations master roles or not, or whether it is a GC or DNS server or not.

| **Note** |
| --- |
| While it is not recommended to restore an operations master role holder in the interest of simplicity, some organizations may choose to restore one for other advantages. For example restoring the RID master may help prevent problems with managing RIDs during the recovery. |

Choose a DC that best meets the following criteria:

- A DC that is writeable. This is mandatory.
- A DC running Windows Server 2012 as a virtual machine on a hypervisor that supports VM-GenerationID. This DC can be used as a source for cloning.
- A DC that is accessible, either physically or on a virtual network, and preferably located in a datacenter. This way, you can easily isolate it from the network during forest recovery.
- A DC that has a good full server backup. A good backup is a backup that can be restored successfully, was taken a few days before the failure, and contains as much useful data as possible.
- A DC that was a Domain Name System (DNS) server before the failure. This saves the time required to reinstall DNS.
- If you also use Windows Deployment Services, choose a DC that is not configured to use BitLocker Network Unlock. In this case, BitLocker Network Unlock is not supported to be used for the first DC that you restore from backup during a forest recovery.

  BitLocker Network Unlock as the *only* key protector *cannot* be used on DCs where you have deployed Windows Deployment Services (WDS) because doing so results in a scenario where the first DC requires Active Directory and WDS to be working in order to unlock. But before you restore the first DC, Active Directory is not yet available for WDS, so it cannot unlock.

  To determine if a DC is configured to use BitLocker Network Unlock, check that a Network Unlock certificate is identified in the following registry key:

  HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\SystemCertificates\FVE_NKP

Maintain security procedures when handling or restoring backup files that include Active Directory. The urgency that accompanies forest recovery can unintentionally lead to overlooking security best practices. For more information, see the section titled "Establishing Domain Controller Backup and Restore Strategies" in Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations: Part II.

**Identify the current forest structure and DC functions**

Determine the current forest structure by identifying all the domains in the forest. Make a list of all of the DCs in each domain, particularly the DCs that have backups, and virtualized DCs which can be a source for cloning. A list of DCs for the forest root domain will be the most important because you will recover this domain first. After you restore the forest root domain, you can obtain a list of the other domains, DCs, and the sites in the forest by using Active Directory snap-ins.

Prepare a table that shows the functions of each DC in the domain, as shown in the following example. This will help you revert back to the pre-failure configuration of the forest after recovery.

| DC name | Operating system | FSMO | GC | RODC | Backup | DNS | Server Core | VM | VM-GenID |
| --- | --- | --- | --- | --- | --- | --- | --- | --- | --- |
| DC_1 | Windows Server 2012 | Schema master, Domain naming master | Yes | No | Yes | No | No | Yes | Yes |
| DC_2 | Windows Server 2012 | None | Yes | No | Yes | Yes | No | Yes | Yes |

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| DC_3 | Windows Server 2012 | Infrastructure Master | No | No | No | Yes | Yes | Yes | Yes |
| DC_4 | Windows Server 2012 | PDC emulator, RID Master | Yes | No | No | No | No | Yes | No |
| DC_5 | Windows Server 2012 | None | No | No | Yes | Yes | No | Yes | Yes |
| RODC_1 | Windows Server 2008 R2 | None | Yes | Yes | Yes | Yes | Yes | Yes | No |
| RODC_2 | Windows Server 2008 | None | Yes | Yes | No | Yes | Yes | Yes | No |

For each domain in the forest, identify a single writeable DC that has a trusted backup of the Active Directory database for that domain. Use caution when you choose a backup to restore a DC. If the day and cause of the failure are approximately known, the general recommendation is to use a backup that was made a few days before that date.

In this example, there are four backup candidates: DC_1, DC_2, DC_4, and DC_5. Of these backup candidates, you restore only one. The recommended DC is DC_5 for the following reasons:

- It satisfies requirements for using it as a source for virtualized DC cloning, that is, it runs Windows Server 2012 as a virtual DC on a hypervisor that supports VM-GenerationID, runs software that is allowed to be cloned (or that can be removed if it is not able to be cloned). After the restore, the PDC emulator role will be seized to that server and it can be added to the Cloneable Domain Controllers group for the domain.

- It runs a full installation of Windows Server 2012. A DC that runs a Server Core installation can be less convenient as a target for recovery.

- It is a DNS server. Therefore, DNS does not have to be reinstalled.

**Note**

Because DC_5 is not a global catalog server, it also has an advantage in that the global catalog does not need to be removed after the restore. But whether or not the DC is also a global catalog server is not a decisive factor because beginning with Windows Server 2012, all DCs are global catalog servers by default, and removing and adding the global catalog after the restore is recommended as part of the forest recovery process in any case.

**Recover the forest in isolation**

The preferred scenario is to shut down all writeable DCs before the first restored DC is brought back into production. This ensures that any dangerous data does not replicate back into the recovered forest. It is particularly important to shut down all operations master role holders.

**Note**

There may be cases where you move the first DC that you plan to recover for each domain to an isolated network while allowing other DCs to remain online in order to minimize system downtime. For example, if you are recovering from a failed schema upgrade, you may choose to keep domain controllers running on the production network while you perform recovery steps in isolation.

If you are running virtualized DCs, you can move them to a virtual network that is isolated from the production network where you will perform recovery. Moving virtualized DCs to a separate network provides two benefits:

- Recovered DCs are prevented from reoccurrence of the problem that caused the forest recovery because they are isolated.

- Virtualized DC cloning can be performed on the separate network so that a critical number of DCs can be running and tested before they are brought back to the production network.

If you are running DCs on physical hardware, disconnect the network cable of the first DC that you plan to restore in the forest root domain. If possible, also disconnect the network cables of all other DCs. This prevents DCs from replicating, if they are accidentally started during the forest recovery process.

In a large forest that is spread across multiple locations, it can be difficult to guarantee that all writeable DCs are shut down. For this reason, the recovery steps—such as resetting the computer account and krbtgt account, in addition to metadata cleanup—are designed to ensure that the recovered writeable DCs do not replicate with dangerous writeable DCs (in case some are still online in the forest).

However, only by taking writeable DCs offline can you guarantee that replication does not occur. Therefore, whenever possible, you should deploy remote management technology that can help you to shut down and physically isolate the writeable DCs during forest recovery.

RODCs can continue to operate while writeable DCs are offline. No other DC will directly replicate any changes from any RODC—especially, no Schema or Configuration container changes—so they do not pose the same risk as writeable DCs during recovery. After all the writeable DCs are recovered and online, you should rebuild all the RODCs.

RODCs will continue to allow access to local resources that are cached in their respective sites while the recovery operations are going on in parallel. Local resources that are not cached on the RODC will have authentication requests forwarded to a writeable DC. These requests will fail because writeable DCs are offline. Some operations such as password changes will also not work until you recover writeable DCs.

If you are using a hub-and-spoke network architecture, you can concentrate first on recovering the writeable DCs in the hub sites. Later, you can rebuild the RODCs in remote sites.

**Perform initial recovery**
This section includes the following steps:

- Restore the first writeable domain controller in each domain
- Reconnect each restored writeable domain controller to a common network
- Add the global catalog to a domain controller in the forest root domain

**Restore the first writeable domain controller in each domain**
Beginning with a writeable DC in the forest root domain, complete the steps in this section in order to restore the first DC. The forest root domain is important because it stores the Schema Admins and Enterprise Admins groups. It also helps maintain the trust hierarchy in the forest. In addition, the forest root domain usually holds the DNS root server for the forest's DNS namespace. Consequently, the Active Directory–integrated DNS zone for that domain contains the alias (CNAME) resource records for all other DCs in the forest (which are required for replication) and the global catalog DNS resource records.

After you recover the forest root domain, repeat the same steps to recover the remaining domains in the forest. You can recover more than one domain simultaneously; however, always recover a parent domain before recovering a child to prevent any break in the trust hierarchy or DNS name resolution.

For each domain that you recover, restore only one writeable DC from backup. This is the most important part of the recovery because the DC must have a database that has not been influenced by whatever caused the forest to fail. It is important to have a trusted backup that is thoroughly tested before it is introduced into the production environment.

Then perform the following steps. Procedures for performing certain steps are in Appendix A: Forest Recovery Procedures.

1. If you plan to restore a physical server, ensure that the network cable of the target DC is not attached and therefore is not connected to the production network. For a virtual machine, you can remove the network adapter or use a network adapter that is attached to another network where you can test the recovery process while isolated from the production network.

2. Because this is the first writeable DC in the domain, you must perform a nonauthoritative restore of AD DS and an authoritative restore of SYSVOL. The restore operation must be completed by using an Active Directory-aware backup and restore application, such as Windows Server Backup (that is, you should not restore the DC by using unsupported methods such as restoring a VM snapshot).

   An authoritative restore of SYSVOL is required because replication of the SYSVOL replicated folder must be started after you recover from a disaster. All subsequent DCs that are added in the domain must resynchronize their SYSVOL folder with a copy of the folder that has been selected to be authoritative before the folder can be advertised.

   **Caution**

| |
|---|
| Perform an authoritative (or primary) restore operation of SYSVOL only for the first DC to be restored in the forest root domain. Incorrectly performing primary restore operations of the SYSVOL on other DCs leads to replication conflicts of SYSVOL data. |

There are two options perform a nonauthoritative restore of AD DS and an authoritative restore of SYSVOL:

- Perform a full server recovery and then then force an authoritative synchronization of SYSVOL. For detailed procedures, see Performing a full server recovery and Perform an authoritative synchronization of DFSR-replicated SYSVOL.

- Perform a full server recovery followed by a system state restore. This option requires that you create both types of backups in advance: a full server backup and a system state backup. For detailed procedures, see Performing a full server recovery and Performing a nonauthoritative restore of Active Directory Domain Services.

3. After you restore and restart the writeable DC, verify that the failure did not affect the data on the DC. If the DC data is damaged, then repeat step 2 with a different backup.

   If the restored domain controller hosts an operations master role, you may need to add the following registry entry to avoid AD DS being unavailable until it has completed replication of a writeable directory partition:

   **HKLM\System\CurrentControlSet\Services\NTDS\Parameters\Repl Perform Initial Synchronizations**

   Create the entry with the data type **REG_DWORD** and a value of **0**. After the forest is recovered completely, you can reset the value of this entry to **1**, which requires a domain controller that restarts and holds operations master roles to have successful AD DS inbound and outbound replication with its known replica partners before it advertises itself as domain controller and starts providing services to clients. For more information about initial synchronization requirements, see KB article 305476.

   Continue to the next steps only after you restore and verify the data and before you join this computer to the production network.

4. If you suspect that the forest-wide failure was related to network intrusion or malicious attack, reset the account passwords for all administrative accounts, including members of the Enterprise Admins, Domain Admins, Schema Admins, Server Operators, Account Operators groups, and so on. The reset of administrative account passwords should be completed before additional domain controllers are installed during the next phase of the forest recovery.

5. On the first restored DC in the forest root domain, seize all domain-wide and forest-wide operations master roles. Enterprise Admins and Schema Admins credentials are needed to seize forest-wide operations master roles.

   In each child domain, seize domain-wide operations master roles. Although you might retain the operations master roles on the restored DC only temporarily, seizing these roles assures you regarding which DC hosts them at this point in the forest recovery process. As part of your post-recovery process, you can redistribute the operations master roles as needed. For more information about seizing operations master roles, see Seizing an operations master role. For recommendations about where to place operations master roles, see What Are Operations Masters?.

6. Clean up metadata of all other writeable DCs in the forest root domain that you are not restoring from backup (all writeable DCs in the domain except for this first DC). If you use the version of Active Directory Users and Computers or Active Directory Sites and Services that is included with Windows Server 2008 or later or RSAT for Windows Vista or later, metadata cleanup is performed automatically when you delete a DC object. In addition, the server object and computer object for the deleted DC are also deleted automatically. For more information, see Cleaning metadata of removed writable DCs.

   Cleaning up metadata prevents possible duplication of NTDS-settings objects if AD DS is installed on a DC in a different site. Potentially, this could also save the Knowledge Consistency Checker (KCC) the process of creating replication links when the DCs themselves might not be present. Moreover, as part of metadata cleanup, DC Locator DNS resource

records for all other DCs in the domain will be deleted from DNS.

Until the metadata of all other DCs in the domain is removed, this DC, if it were a RID master before recovery, will not assume the RID master role and therefore will not be able to issue new RIDs. You might see event ID 16650 in the System log in Event Viewer indicating this failure, but you should see event ID 16648 indicating success a little while after you have cleaned the metadata.

7. If you have DNS zones that are stored in AD DS, ensure that the local DNS Server service is installed and running on the DC that you have restored. If this DC was not a DNS server before the forest failure, you must install and configure the DNS server.

| **Note** |
| --- |
| If the restored DC runs Windows Server 2008, you need to install the hotfix in KB article 975654 or connect the server to an isolated network temporarily in order to install DNS server. The hotfix is not required for any other versions of Windows Server. |

In the forest root domain, configure the restored DC with its own IP address (or a loopback address, such as 127.0.0.1) as its preferred DNS server. You can configure this setting in the TCP/IP properties of the local area network (LAN) adapter. This is the first DNS server in the forest. For more information, see Configure TCP/IP to use DNS.

In each child domain, configure the restored DC with the IP address of the first DNS server in the forest root domain as its preferred DNS server. You can configure this setting in the TCP/IP properties of the LAN adapter. For more information, see Configure TCP/IP to use DNS.

In the _msdcs and domain DNS zones, delete NS records of DCs that no longer exist after metadata cleanup. Check if the SRV records of the cleaned up DCs have been removed. To help speed up DNS SRV record removal, run:

```
nltest.exe /dsderegdns:server.domain.tld
```

8. Raise the value of the available RID pool by 100,000. For more information, see Raising the value of available RID pools. If you have reason to believe that raising the RID Pool by 100,000 is insufficient for your particular situation, you should determine the lowest increase that is still safe to use. RIDs are a finite resource that should not be used up needlessly.

If new security principals were created in the domain after the time of the backup that you use for the restore, these security principals might have access rights on certain objects. These security principals no longer exist after recovery because the recovery has reverted to the backup; however, their access rights might still exist. If the available RID pool is not raised after a restore, new user objects that are created after the forest recovery might obtain identical security IDs (SIDs) and could have access to those objects, which was not originally intended.

To illustrate, consider the example of the new employee named Amy that was mentioned in the introduction. The user object for Amy no longer exists after the restore operation because it was created after the backup that was used to restore the domain. However, any access rights that were assigned to that user object might persist after the restore operation. If the SID for that user object is reassigned to a new object after the restore operation, the new object would obtain those access rights.

9. Invalidate the current RID pool. The current RID pool is invalidated after a system state restore. But if a system state restore was not performed, the current RID pool needs to be invalidated to prevent the restored DC from re-issuing RIDs from the RID pool that was assigned at the time the backup was created. For more information, see Invalidating the current RID pool.

| **Note** |
| --- |
| The first time that you attempt to create an object with a SID after you invalidate the RID pool you will receive an error. The attempt to create an object triggers a request for a new RID pool. Retry of the operation succeeds because the new RID pool will be allocated. |

10. Reset the computer account password of this DC twice. For more information, see Resetting the computer account

password of the domain controller.

11. Reset the krbtgt password twice. For more information, see Resetting the krbtgt password.

    Because the krbtgt password history is two passwords, reset passwords twice to remove the original (prefailure) password from password history.

    | **Note** |
    | --- |
    | If the forest recovery is in response to a security breach, you may also reset the trust passwords. For more information, see Resetting a trust password on one side of the trust. |

12. If the forest has multiple domains and the restored DC was a global catalog server before the failure, clear the **Global catalog** check box in the NTDS Settings properties to remove the global catalog from the DC. The exception to this rule is the common case of a forest with just one domain. In this case, it is not required to remove the global catalog. For more information, see Removing the global catalog.

    By restoring a global catalog from a backup that is more recent than other backups that are used to restore DCs in other domains, you might introduce lingering objects. Consider the following example. In domain A, DC1 is restored from a backup that was taken at time T1. In domain B, DC2 is restored from a global catalog backup that was taken at time T2. Suppose T2 is more recent than T1, and some objects were created between T1 and T2. After these DCs are restored, DC2, which is a global catalog, holds newer data for domain A's partial replica than domain A holds itself. DC2, in this case, holds lingering objects because these objects are not present on DC1.

    The presence of lingering objects can lead to problems. For instance, e-mail messages might not be delivered to a user whose user object was moved between domains. After you bring the outdated DC or global catalog server back online, both instances of the user object appear in the global catalog. Both objects have the same e-mail address; therefore, e-mail messages cannot be delivered.

    A second problem is that a user account that no longer exists might still appear in the global address list. A third problem is that a universal group that no longer exists might still appear in a user's access token.

    If you did restore a DC that was a global catalog—either inadvertently or because that was the solitary backup that you trusted—we recommend that you prevent the occurrence of lingering objects by disabling the global catalog soon after the restore operation is complete. Disabling the global catalog flag will result in the computer losing all its partial replicas (partitions) and relegating itself to regular DC status.

13. Configure Windows Time Service. In the forest root domain, configure the PDC emulator to synchronize time from an external time source. For more information, see Configure the Windows Time service on the PDC emulator in the Forest Root Domain.

**Reconnect each restored writeable domain controller to a common network**

At this stage you should have one DC restored (and recovery steps performed) in the forest root domain and in each of the remaining domains. Join these DCs to a common network that is isolated from the rest of the environment and complete the following steps in order to validate forest health and replication.

| **Note** |
| --- |
| When you join the physical DCs to an isolated network, you may need to change their IP addresses. As a result, the IP addresses of DNS records will be wrong. Because a global catalog server is not available, secure dynamic updates for DNS will fail. Virtual DCs are more advantageous in this case because they can be joined to a new virtual network without changing their IP addresses. This is one reason why virtual DCs are recommended as the first domain controllers to be restored during forest recovery. |

After validation, Join the DCs to the production network and complete the steps to verify forest replication health.

- To fix name resolution, create DNS delegation records and configure DNS forwarding and root hints as needed. Run **repadmin /replsum** to check replication between DCs.

- If the restored DC's are not direct replication partners, replication recovery will be much faster by creating temporary connection objects between them.

- To validate metadata cleanup, run **Repadmin /viewlist *** for a list of all DCs in the forest. Run **Nltest /DCList:**<*domain*> for a list of all DCs in the domain.

- To check DC and DNS health, run DCDiag /v to report errors on all DCs in the forest.

**Add the global catalog to a domain controller in the forest root domain**
A global catalog is required for these and other reasons:

- To enable logons for users.

- To enable the Net Logon service running on the DCs in each child domain to register and remove records on the DNS server in the root domain.

Although it is preferred that the forest root DC become a global catalog, it is possible to elect any of the restored DCs to become a global catalog.

| **Note** |
| --- |
| A DC will not be advertised as a global catalog server until it has completed a full synchronization of all directory partitions in the forest. Therefore, the DC should be forced to replicate with each of the restored DCs in the forest.<br><br>Monitor the Directory Service event log in Event Viewer for event ID 1119, which indicates that this DC is a global catalog server, or verify the following registry key has a value of 1:<br><br>**HKLM\System\CurrentControlSet\Services\NTDS\Parameters\Global Catalog Promotion Complete** |

For more information, see Adding the global catalog.

At this stage you should have a stable forest, with one DC for each domain and one global catalog in the forest. You should make a new backup of each of the DCs that you have just restored. You can now begin to redeploy other DCs in the forest by installing AD DS.

**Redeploy remaining DCs**
The steps up to this point apply to all forests: find a valid backup for each domain, recover the domains in isolation, reconnect them, reset the global catalog, and clean up. In this next step you will redeploy the forest. The way to do this will greatly depend on your forest design, your service level agreements, site structure, available bandwidth, and numerous other factors. You will need to design your own redeployment plan based on the principles and suggestions in this section, in a way that is best suited to your business requirements.

The next step is to install AD DS on all DCs that were present before the forest recovery took place. If the DCs still exist, the AD DS service will need to be removed forcibly, or the DCs can be reinstalled. Any existing backups for these DCs cannot be reused, because the corresponding metadata has been removed during forest recovery. In an uncomplicated environment this redeployment process can be as simple as reconnecting the recovered DCs to the production network, and promoting new DCs as needed.

In a large enterprise faced with a worldwide infrastructure, a more sophisticated plan is needed. The first phase is usually to restore the AD as a service; this means to install strategically placed DCs such that all critical business divisions and applications can start working again. It may be acceptable for branch offices to temporarily have reduced performance as a result of this. As a second phase, all remaining and less critical DCs are redeployed.

There are two methods to install additional DCs, both of which can be automated:

- Cloning

  For virtualized environments that run Windows Server 2012, cloning is the fastest and simplest way to recover a large number of DCs. You can automate the recovery of all virtualized DCs in a domain after you restore a single virtualized DC from backup.

For more information about cloning and prerequisites, see Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100).

- Re-install AD DS by using Windows PowerShell on servers that run Windows Server 2012 (or Dcpromo.exe on servers that run earlier versions of Windows Server) or by using the user interface

  To expedite re-installing AD DS, you can use Install from Media (IFM) option to reduce replication traffic during the installation. For more information about using the **ntdsutil ifm** command to create installation media, see Installing AD DS from Media.

Consider the following additional points for each replica DC that is recovered in the forest by virtualized DC cloning or by installing AD DS (as opposed to restoring from backup):

- All software on a DC that is used as the source for cloning must be able to be cloned. Applications and services that cannot be cloned should be removed before cloning is initiated. If that is not possible, an alternative virtualized DC should be chosen as the source.

- If you clone additional virtualized DCs from the first virtualized DC to be restored, the source DC will need to be shut down while its VHDX file is copied. Then it will need to be running and available online when the clone virtual DCs are first started. If the downtime required by the shutdown is not acceptable for the first recovered DC, deploy an additional virtualized DC by installing AD DS to act as the source for cloning.

- There is no restriction on the host name of the cloned virtualized DC or the server on which you want to install AD DS. You can use a new host name or the host name that was in use previously. For more information about DNS host name syntax, see Creating DNS Computer Names (http://go.microsoft.com/fwlink/?LinkId=74564).

- Configure each server with the first DNS server in the forest (the first DC that was restored in the root domain) as the preferred DNS server in the TCP/IP properties of its network adapter. For more information, see Configure TCP/IP to use DNS.

- Redeploy all RODCs in the domain, either by virtualized DC cloning if several RODCs are deployed in a central location, or by the traditional method of rebuilding them by removing and reinstalling AD DS if they are deployed individually in isolated located locations such as branch offices.

  Rebuilding RODCs ensures that they do not contain any lingering objects and can help prevent replication conflicts from occurring later. When you remove AD DS from an RODC, *choose the option to retain DC metadata*. Using this option retains the krbtgt account for the RODC and retains the permissions for the delegated RODC administrator account and the Password Replication Policy (PRP), and prevents you from having to use Domain Admin credentials to remove and reinstall AD DS on an RODC. It also retains the DNS server and global catalog roles if they are installed on the RODC originally.

  When you rebuild DCs (RODCs or writeable DCs), there may be increased replication traffic during their reinstallation. To help reduce that impact, you can stagger the schedule of the RODC installations, and you can use the Install From Media (IFM) option. If you use the IFM option, run the **ntdsutil ifm** command on a writeable DC that you trust to be free of damaged data. This helps prevent possible corruption from appearing on the RODC after the AD DS reinstallation is complete. For more information about IFM, see Installing AD DS from Media.

  For more information about rebuilding RODCs, see RODC Removal and Reinstallation.

- If a DC was running the DNS Server service before the forest malfunction, install and configure the DNS Server service during the installation of AD DS. Otherwise, configure its former DNS clients with other DNS servers.

- If you require additional global catalogs to share authentication or query load for users or applications, you can either add the global catalog to the source virtualized DC before cloning or you can make a DC a global catalog server during the installation of AD DS.

**Cleanup**

Perform the following post recovery steps as needed:

- After the entire forest is recovered, you can revert to the original DNS configuration, including configuration of the preferred and alternate DNS servers on each of the DCs. After the DNS servers are configured as they were before the malfunction, their previous name resolution capabilities will be restored. Delete any DNS records for DCs that have not been recovered.

- Delete Windows Internet Name Service (WINS) records for all DCs that have not been recovered.

- You can transfer the operations master roles to other DCs in the domain or forest and add more global catalog servers based on the configuration before the failure.

- Because the entire forest is restored to a previous state, any objects (such as users and computers) that were added and all updates (such as password changes) that were made to existing objects after this point are lost. Therefore, you should re-create these missing objects and reapply the missing updates as appropriate.

- You might also need to restore outgoing trusts with external domains and forests, because these external trust relationships are not restored automatically from backups.

# Appendix A: Forest Recovery Procedures

Updated: April 25, 2013

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

This appendix contains procedures related to the forest recovery process described earlier in this guide. The procedures are applicable for Windows Server 2012, and are also applicable to Windows Server 2008 R2 and Windows Server 2008 with some minor exceptions. Procedures that include steps that vary for Windows Server 2003 are found in Appendix D: Forest Recovery with Windows Server 2003 Domain Controllers.

- Backing up a full server
- Backing up the System State data
- Performing a full server recovery
- Performing an authoritative synchronization of DFSR-replicated SYSVOL
- Performing a nonauthoritative restore of Active Directory Domain Services

  These steps explain how to perform an authoritative restore of SYSVOL at the same time.
- Configuring the DNS Server service
- Removing the global catalog
- Raising the value of available RID pools
- Invalidating the current RID pool
- Seizing an operations master role
- Cleaning metadata of removed writable domain controllers
- Resetting the computer account password of the domain controller
- Resetting the krbtgt password
- Resetting a trust password on one side of the trust
- Adding the global catalog
- Resources to verify replication is working

**Backing up a full server**

A full server backup is recommended to prepare for a forest recovery because it can be restored to different hardware or a different operating system instance. To perform a full server backup, perform the following procedure using Windows Server Backup. Windows Server Backup is not installed by default. In Windows Server 2012, install it by selecting **Windows Server Backup** on the **Select Features** page in the **Add Roles and Features Wizard** in **Server Manager**. For steps to install it in Windows Server 2008 and Windows Server 2008 R2, see Installing Windows Server Backup.

**To perform a full server backup using Windows Server Backup**

1. In Windows Server 2012, open **Server Manager**, click **Tools**, and then click **Windows Server Backup**.

   In Windows Server 2008 R2 and Windows Server 2008, click **Start**, point to **Administrative Tools**, and then click

**Windows Server Backup**.

2. If you are prompted, in the **User Account Control** dialog box, provide Backup Operator credentials, and then click **OK**.
3. In Windows Server 2012 only, first click **Local Backup**.
4. On the **Action** menu, click **Backup once**.
5. In the Backup Once Wizard, on the **Backup options** page, click **Different options**, and then click **Next**.
6. On the **Select backup configuration** page, click **Full server (recommended)**, and then click **Next**.
7. On the **Specify destination type** page, click **Local drives** or **Remote shared folder**, and then click **Next**.
8. On the **Specify destination type** page, choose the backup location as follows:
   - If you are backing up to a local volume, in **Backup destination**, select a drive, and then click **Next**.

     When you are prompted to exclude the destination volume from the list of items to be backed up, click **OK**.

   - If you are backing up to a remote shared folder, do the following:

     a. Type the path to the shared folder.

     b. Under **Access Control**, select **Do not inherit** or **Inherit** to determine access to the backup, and then click **Next**.

     c. In the **Provide user credentials for Backup** dialog box, provide the user name and password for a user who has write access to the shared folder, and then click **OK**.

9. On the **Confirmation** page, review your selections, and then click **Backup**.
10. After the Backup Once Wizard begins the backup, click **Close** at any time. The backup runs in the background and you can view backup progress at any time during the backup. The wizard closes automatically when the backup is complete.

**To perform a full server backup using Wbadmin.exe**

1. Open an elevated command prompt, type the following command and press ENTER:

```
wbadmin start backup -backuptarget:<Drive_letter_to store_backup>: -include:
<Drive_letter_to_include>: -quiet
```

**Backing up the System State data**

A System State backup must be restored to the same operating system instance and hardware. Therefore it is not as flexible during a forest recovery as a full server backup. But a system state backup can be used to perform a non-authoritative restore of AD DS and an authoritative restore of SYSVOL at the same time (using wbadmin.exe), which may be more convenient than the full server restore option.

To back up System State data, complete the following procedures:

**Backing up the System State data**

Use the following procedure to perform a system state backup on a DC by using Windows Server Backup or wbadmin.exe.

**To perform a system state backup using Windows Server Backup**

1. In Windows Server 2012, open **Server Manager**, click **Tools**, and then click **Windows Server Backup**.

   In Windows Server 2008 R2 and Windows Server 2008, click **Start**, point to **Administrative Tools**, and then click **Windows Server Backup**.

2. If you are prompted, in the **User Account Control** dialog box, provide Backup Operator credentials, and then click **OK**.
3. In Windows Server 2012 only, first click **Local Backup**.

4. On the **Action** menu, click **Backup once**.
5. In the Backup Once Wizard, on the **Backup options** page, click **Different options**, and then click **Next**.
6. On the **Select backup configuration** page, click **Custom**, and then click **Next**.
7. On the **Select backup items** page:

   In Windows Server 2012, click **Add Items**, click **System state**, then click **Next**.

   In Windows Server 2008 R2 and Windows Server 2008, select the volumes to include in the backup. If you select the **Enable system recovery** check box, all critical volumes are selected.

   As an alternative, you can clear that check box, select the individual volumes that you want to include, and then click **Next**.

   Your selection must include the volumes that store the operating system, Ntds.dit, and SYSVOL.

   | **Note** |
   | --- |
   | If you select a volume that hosts an operating system, all volumes that store system components are also selected. |

8. On the **Specify destination type** page, choose the backup location as follows:
   - If you are backing up to a local volume, in **Backup destination**, select a drive, and then click **Next**.

     When you are prompted to exclude the destination volume from the list of items to be backed up, click **OK**.

   - If you are backing up to a remote shared folder, do the following:

     a. Type the path to the shared folder.

     b. Under **Access Control**, select **Do not inherit** or **Inherit** to determine access to the backup, and then click **Next**.

     c. In the **Provide user credentials for Backup** dialog box, provide the user name and password for a user who has write access to the shared folder, and then click **OK**.

9. For Windows Server 2008 R2 and Windows Server 2008, on the **Specify advanced option** page, select **VSS copy backup** and then click **Next**,
10. On the **Confirmation** page, review your selections, and then click **Backup**.
11. After the Backup Once Wizard begins the backup, click **Close** at any time. The backup runs in the background and you can view backup progress at any time during the backup. The wizard closes automatically when the backup is complete.

**To perform a system state backup using Wbadmin.exe**

1. Open an elevated command prompt, type the following command and press ENTER:

```
wbadmin start systemstatebackup -backuptarget:<targetDrive>: -allCritical -quiet
```

**Performing a full server recovery**

Use the following procedure to perform a full server recovery for Windows Server 2012. A full server recovery is necessary if you are restoring to different hardware or a different operating system instance. The number drives on the target server needs to be equal to the number in the backup and they need to be the same size or greater.

The target server needs to be started from the operating system DVD in order to access the **Repair your computer** option. If the target DC is running in a VM on Hyper-V and the backup is stored on a network location, you must install a legacy network adapter.

After you perform a full server recovery, you need to separately perform an authoritative restore of SYSVOL, as described in the next section.

**To perform a full server recovery**

1. Start Windows Setup, specify the Language, Time and currency format, and keyboard options and click **Next**.
2. Click **Repair your computer**. If the backup is stored locally, skip to step 6. If the backup is stored on a network location, continue with step 3.
3. Click **Troubleshoot**, click **Command Prompt**.
4. Type the following command and press ENTER:

   ```
   wpeinit
   ```

5. To confirm the name of the network adapter, type:

   ```
   show interfaces
   ```

   Type the following commands and press ENTER after each command:

   ```
   netsh

   interface

   tcp

   ipv4

   set address "Name of Network Adapter" static IPv4 AddressSubnetMaskIPv4 Gateway
   Address 1
   ```

   For example:

   ```
   set address "Local Area Connection" static 192.168.1.2 255.0.0.0 192.168.1.1 1
   ```

   Type `quit` to return to a command prompt. Type `ipconfig /all` to verify the network adapter has an IP address and try to ping the IP address of the server that hosts the backup share to confirm connectivity. Close the command prompt when you are done.

6. Click **Troubleshoot**, click **System Image Recovery**, and click **Windows Server 2012**.
7. If you are restoring the most recent local backup, click **Use the latest available system image (recommended)** and click **Next** twice, click **Finish**, and click **Yes** to confirm the restore operation.

   If you are restoring a different backup, click **Select a system image** and click **Next**.

8. Click the name of a local backup file or click **Advanced** and click **Search for a system image over the network** to search for a backup over the network.
9. Type the UNC path to the backup share location (for example, \\server1\backups) and click **OK**. You can also type the IP address of the target server, such as \\192.168.1.3\backups.
10. Type credentials necessary to access the share and click OK.
11. Select the name of the backup file and click **Next**.
12. Select the drives in the backup file and click **Next**.
13. Click **Format and repartition disks** and click **Next**.
14. Click **Finish**, and click **Yes** to confirm that all disks will be restored.

**Performing an authoritative synchronization of DFSR-replicated SYSVOL**

There are different ways to perform an authoritative restore of SYSVOL. You can either edit the **msDFSR-Options** attribute or perform a system state restore using wbadmin –authsysvol. If you have the option to restore a system state backup (that is, you are restoring AD DS to the same hardware and operating system instance) then using wbadmin –authsysvol is simpler. But if you need to perform a bare metal restore, then you need to edit the **msDFSR-Options** attribute.

Use the following steps to perform an authoritative synchronization of SYSVOL (if it is replicated using DFSR) by editing the

**msDFSR-Options** attribute. If SYSVOL is replicated using FRS, see article 290762.

**To perform an authoritative synchronization of DFSR-replicated SYSVOL**

1. Open Active Directory Users and Computers.
2. Click **View**, and then select **Users, Contacts, Groups, and Computers as containers** and **Advanced Features**.
3. In the tree-view, click **Domain Controllers**, the name of the DC you restored, **DFSR-LocalSettings**, and then **Domain System Volume**.
4. In the Details pane, right-click **SYSVOL Subscription**, click **Properties**, and click **Attribute Editor**.
5. Click **msDFSR-Options**, click **Edit**, type **1**, and click **OK**
6. Click **OK** to close the Attribute Editor.

**Performing a nonauthoritative restore of Active Directory Domain Services**
To perform a nonauthoritative restore, complete the following procedure.

The following procedures use the Wbadmin.exe to perform a nonauthoritative restore of Active Directory or Active Directory Domain Services (AD DS). If you are using a different backup solution or if you intend to complete the authoritative restore of SYSVOL later in the forest recovery process, you can perform an authoritative restore of SYSVOL by using these alternative methods:

- If you are using File Replication Service (FRS) to replicate SYSVOL, follow the steps in article 290762 in the Microsoft Knowledge Base, using the **BurFlags** registry key to reinitialize FRS replica sets, or if necessary, article 315457 315457to rebuild the SYSVOL tree. To determine if SYSVOL is replicated by FRS, see Determining Whether a Domain Controller's SYSVOL Folder is Replicated by DFSR or FRS.

- If you are using Distributed File System (DFS) Replication to replicate SYSVOL, see Performing an authoritative synchronization of DFSR-replicated SYSVOL.

**Performing a nonauthoritative restore**
Use the following procedure to perform a nonauthoritative restore of AD DS and an authoritative restore of SYSVOL at the same time by using wbadmin.exe on a DC that runs Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008. The backup must explicitly include system state data; a full server backup that is used for full server recovery will not work. For more information about creating a system state backup, see Backing up the System State data.

**To perform a nonauthoritative restore of AD DS and authoritative restore of SYSVOL using wbadmin.exe**

- Include the **-authsysvol** switch in your recovery command, as shown in the following example:

  ```
  wbadmin start systemstaterecovery <otheroptions> -authsysvol
  ```

  For example:

  ```
  wbadmin start systemstaterecovery -version:11/20/2012-13:00 -authsysvol
  ```

**Configuring the DNS Server service**
If the DNS server role is not installed on the DC that you restore from backup, you must install and configure the DNS server.

**Install and configure the DNS Server service**
Complete this step for each restored DC that is not running as a DNS server after the restore is complete.

| **Note** |
| --- |
| If the DC that you restored from backup is running Windows Server 2008, you must connect the DC to an isolated network in order to install DNS server. Then connect each of the restored DNS servers to a mutually shared, isolated network. Run repadmin /replsum to verify that replication is functioning between the restored DNS servers. After you verify replication, you can connect the restored DCs to the production network If the DNS server role is already installed, you can apply a hotfix that makes it possible for a DNS server to start while the server is not connected to any network. You should slipstream the hotfix into the operating system installation image during your automated build processes. For more information about the hotfix, see Article 975654 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=184691). |

**To install and configure the DNS Server service using Server Manager**

1. Open Server Manager and start the **Add Roles Wizard**.
2. In the Add Roles Wizard, if the **Before You Begin** page appears, click **Next**.
3. In the **Roles** list, click **DNS Server**, and then click **Next**.
4. Read the information on the **DNS Server** page, and then click **Next**.
5. On the **Confirm Installation Options** page, verify that the DNS Server role will be installed, and then click **Install**.

   After the installation, complete the following steps to configure the DNS server.
6. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **DNS**.
7. Create DNS zones for the same DNS domain names that were hosted on the DNS servers before the critical malfunction. For more information, see Add a Forward Lookup Zone (http://go.microsoft.com/fwlink/?LinkId=74574).
8. Configure the DNS data as it existed before the critical malfunction. For example:
   - Configure DNS zones to be stored in AD DS. For more information, see Change the Zone Type (http://go.microsoft.com/fwlink/?LinkId=74579).
   - Configure the DNS zone that is authoritative for domain controller locator (DC Locator) resource records to allow secure dynamic update. For more information, see Allow Only Secure Dynamic Updates (http://go.microsoft.com/fwlink/?LinkId=74580).
9. Ensure that the parent DNS zone contains delegation resource records (name server (NS) and glue host (A) resource records) for the child zone that is hosted on this DNS server. For more information, see Create a Zone Delegation (http://go.microsoft.com/fwlink/?LinkId=74562).
10. After you configure DNS, you can speed up registration of the NETLOGON Records.

    | **Note** |
    | --- |
    | Secure dynamic updates only work when a global catalog server is available. |

    At the command prompt, type the following command, and then press ENTER:

    **net stop netlogon**
11. Type the following command, and then press ENTER:

    **net start netlogon**

**Removing the global catalog**

Use the following procedure to remove the global catalog from a DC.

Restoring a global catalog server from backup could result in the global catalog holding newer data for one of its partial replicas than the corresponding domain that is authoritative for that partial replica. In such a case, the newer data will not be removed from the global catalog and might even replicate to other global catalog servers. As a result, even if you did restore a DC that was a global catalog server, either inadvertently or because that was the solitary backup you trusted, you should remove the global catalog soon after the restore operation is complete. When the global catalog is removed, the computer removes all its partial replicas.

**To remove the global catalog using Active Directory Sites and Services**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
2. In the console tree, expand the **Sites** container, and then select the appropriate site that contains the target server.
3. Expand the **Servers** container, and then expand the *server* object for the DC from which you want to remove the global catalog.
4. Right-click **NTDS Settings**, and then click **Properties**.

5. Clear the **Global Catalog** check box.

**To remove the global catalog using Repadmin**

1. Open an elevated command prompt, type the following command, and press ENTER:

   ```
   repadmin.exe /options DC_NAME -IS_GC
   ```

**Raising the value of available RID pools**

Use the following procedure to raise the value of the relative ID (RID) pools that the RID operations master will allocate after that DC is restored. By raising the value of the available RID pools, you can ensure that no DC allocates a RID for a security principal that was created after the backup that was used to restore the domain.

Each domain has an object **CN=RID Manager$,CN=System,DC=**<*domain_name*>. This object has an attribute named **rIDAvailablePool**. This attribute value maintains the global RID space for an entire domain. The value is a large integer with upper and lower parts. The upper part defines the number of security principals that can be allocated for each domain (0x3FFFFFFF or just over 1 billion). The lower part is the number of RIDs that have been allocated in the domain.

| **Note** |
| --- |
| In Windows Server 2012, the number of security principals that can be allocated is increased to just over 2 billion. For more information, see Managing RID issuance. |

- Sample Value: 4611686014132422708
- Low Part: 2100 (beginning of the next RID pool to be allocated)
- Upper Part: 1073741823 (total number of RIDs that can be created in a domain)

When you increase the value of the large integer, you increase the value of the low part. For example, if you add 100,000 to the sample value of 4611686014132422708 for a sum of 4611686014132522708, the new low part is 102100. This indicates that the next RID pool that will be allocated by the RID master will begin with 102100 instead of 2100.

**To raise the value of available RID pools using adsiedit and the calculator**

1. At an elevated command prompt, type:

   ```
   adsiedit.msc
   ```

2. Connect do the Default Naming Context, and browse to the following distinguished name path: **CN=RID Manager$,CN=System,DC=**<*domain name*>.
3. Open the properties of CN=RID Manager$.
4. Select the attribute **rIDAvailablePool**, click **Edit**, and then copy the large integer value to the clipboard.
5. Start calculator, and from the **View** menu, select **Scientific Mode**.
6. Add 100,000 to the current value.
7. Using ctrl-c, or the **Copy** command from the **Edit** menu, copy the value to the clipboard.
8. In the edit dialog of adsiedit, paste this new value.
9. Click **OK** in the dialog, and **Apply** in the property sheet to update the **rIDAvailablePool** attribute.

**To raise the value of available RID pools using LDP**

1. At the command prompt, type the following command, and then press ENTER:

   **ldp**

2. Click **Connection**, click **Connect**, type the name of RID manager, and then click **OK**.

3. Click **Connection**, click **Bind**, type your administrative credentials, and then click **OK**.
4. Click **View**, click **Tree**, and then type the following distinguished name path:

   CN=RID Manager$,CN=System,DC=*domain name*

5. Click **Browse**, and then click **Modify**.
6. Add 100,000 to the current **rIDAvailablePool** value, and then type the sum into **Values**.
7. In **Dn**, type **cn=RID Manager$,cn=System,dc=**<*domain name*>.
8. In **Edit Entry Attribute**, type **rIDAvailablePool**.
9. Select **Replace** as the operation, and then click **Enter**.
10. Click **Run** to run the operation.
11. To validate the change, select the cn=RID Manager$,cn=System,dc=<domain name> object and verify then value of the **rIDAvailablePool** attribute.

**Invalidating the current RID pool**

Use the following procedure to us Windows PowerShell to invalidate the current RID pool on a domain controller. Windows PowerShell is enabled by default on Windows Server 2012 and Windows Server 2008 R2, but not Windows Server 2008 where it must be installed by using **Add Features**. It can be downloaded to run on Windows Server 2003.

To verify the command completed successfully, check for event ID 16654 (source is Directory-Services-SAM) in the System log in Event Viewer in Windows Server 2012. Earlier versions of Windows do not log this event.

| **Note** |
| --- |
| After you invalidate the RID pool, you will receive an error when you first attempt to create security principal (user, computer, or group). The attempt to create an object triggers a request for a new RID pool. Retry of the operation succeeds because the new RID pool will be allocated. |

**To invalidate the current RID pool**

1. Open an elevated Windows PowerShell session, run the following command and press ENTER:

```
$Domain = New-Object System.DirectoryServices.DirectoryEntry
$DomainSid = $Domain.objectSid
$RootDSE = New-Object System.DirectoryServices.DirectoryEntry("LDAP://RootDSE")
$RootDSE.UsePropertyCache = $false
$RootDSE.Put("invalidateRidPool", $DomainSid.Value)
$RootDSE.SetInfo()
```

**Seizing an operations master role**

Use the following procedure to seize an operations master role (also known as a flexible single master operations (FSMO) role). You can use Ntdsutil.exe, a command-line tool that is installed automatically on all DCs.

**To seize an operations master role**

1. At the command prompt, type the following command, and then press ENTER:

```
ntdsutil
```

2. At the **ntdsutil:** prompt, type the following command, and then press ENTER:

```
roles
```

3. At the **FSMO maintenance:** prompt, type the following command, and then press ENTER:

```
connections
```

4. At the **server connections:** prompt, type the following command, and then press ENTER:

```
Connect to server ServerFQDN
```

Where *ServerFQDN* is the fully qualified domain name (FQDN) of this DC, for example: **connect to server nycdc01.example.com**.

If *ServerFQDN* does not succeed, use the NetBIOS name of the DC.

5. At the **server connections:** prompt, type the following command, and then press ENTER:

```
quit
```

6. Depending on the role that you want to seize, at the **FSMO maintenance:** prompt, type the appropriate command as described in the following table, and then press ENTER.

| Role | Credentials | Command |
|---|---|---|
| Domain naming master | Enterprise Admins | **Seize naming master** |
| Schema master | Schema Admins | **Seize schema master** |
| Infrastructure master<br>**Note**<br>After you seize the infrastructure master role, you may receive an error later if you need to run Adprep /Rodcprep. For more information, see KB article 949257. | Domain Admins | **Seize infrastructure master** |
| PDC emulator master | Domain Admins | **Seize pdc** |
| RID master | Domain Admins | **Seize rid master** |

After you confirm the request, Active Directory or AD DS attempts to transfer the role. When the transfer fails, some error information appears, and Active Directory or AD DS proceeds with the seizure. After the seizure is complete, a list of the roles and the Lightweight Directory Access Protocol (LDAP) name of the server that currently holds each role appears. You can also run **Netdom Query FSMO** at an elevated command prompt to verify current role holders.

**Note**

If this computer was not a RID master before the failure and you attempt to seize the RID master role, the computer tries to synchronize with a replication partner before accepting this role. However, because this step is performed when the computer is isolated, it will not succeed in synchronizing with a partner. Therefore, a dialog box appears asking you whether you want to continue with the operation despite this computer not being able to synchronize with a partner. Click **Yes**.

**Cleaning metadata of removed writable domain controllers**

Metadata cleanup removes Active Directory data that identifies a DC to the replication system.

Use the following procedure to delete the DC objects for DCs that you plan to add back to the network by reinstalling AD DS.

If you are using the version of Active Directory Users and Computers or Active Directory Sites and Services that is included Remote Server Administration Tools (RSAT), metadata cleanup is performed automatically when you delete a DC object.

**Deleting a domain controller using Active Directory Users and Computers**

When you use the version of Active Directory Users and Computers or Active Directory Administrative Center in Remote Server Administration Tools (RSAT), metadata cleanup is performed automatically when you delete the DC object. The server object and the computer object are also deleted automatically.

As an alternative, you can also use Active Directory Sites and Services in RSAT to delete a DC object. If you use Active Directory Sites and Services, you must delete the associated server object and NTDS Settings object before you can delete the DC object.

To download RSAT:

- Remote Server Administration Tools for Windows 8
- Remote Server Administration Tools for Windows 7 with Service Pack 1 (SP1)
- Microsoft Remote Server Administration Tools for Windows Vista

The following procedure is the same for DCs that run either Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008. The target DC of the metadata cleanup operation can run any version of Windows Server.

**To delete a domain controller object using Active Directory Users and Computers in RSAT**

1. Click **Start**, click **Administrative Tools**, and then click **Active Directory Users and Computers**.
2. In the console tree, double-click the domain container, and then double-click the **Domain Controllers** organizational unit (OU).
3. In the details pane, right-click the DC that you want to delete, and then click **Delete**.
4. Click **Yes** to confirm the deletion. Select the **This Domain Controller is permanently offline and can no longer be demoted using the Active Directory Domain Services Installation Wizard (DCPROMO)** check box and click **Delete**.
5. If the DC was a global catalog server, click **Yes** confirm that the deletion.

**Resetting the computer account password of the domain controller**
Use the following procedure to reset the computer account password of the DC.

**To reset the computer account password of the domain controller**

1. At a command prompt, type the following command, and then press ENTER:

   ```
   netdom help resetpwd
   ```

2. Use the syntax that this command provides for using the Netdom command-line tool to reset the computer account password, for example:

   ```
   netdom resetpwd /server:domain controller name /userD:administrator /passwordd:*
   ```

   Where *domain controller name* is the local DC that you are recovering.

   | **Note** |
   | --- |
   | You should run this command twice. |

**Resetting the krbtgt password**
Use the following procedure to reset the krbtgt password for the domain. The following procedure applies writeable DCs, but not read-only domain controllers (RODCs).

| **Important** |
| --- |
| If you plan to recover RODCs online during the forest recovery, do not delete the krbtgt accounts for the RODCs. The krbtgt account for an RODC is listed in the format krbtgt_*number*.<br><br>If you use a customized password filter (such as passfilt.dll) on a DC, then you might receive an error when you try to reset the krbtgt password. For more information, including a workaround, see Microsoft Knowledge Base article 2549833 (http://support.microsoft.com/kb/2549833). |

**To reset the krbtgt password**

1. Click **Start**, point to **Control Panel**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
2. Click **View**, and then click **Advanced Features**.

3. In the console tree, double-click the domain container, and then click **Users**.
4. In the details pane, right-click the **krbtgt** user account, and then click **Reset Password**.
5. In **New password**, type a new password, retype the password in **Confirm password**, and then click **OK**. The password that you specify is not significant because the system will generate a strong password automatically independent of the password that you specify.

| **Note** |
| --- |
| You should perform this operation twice. The password history of the krbtgt account is two, meaning it includes the two most recent passwords. By resetting the password twice you effectively clear any old passwords from the history, so there is no way another DC will replicate with this DC by using an old password. |

**Resetting a trust password on one side of the trust**

If the forest recovery is related to a security breach, use the following procedure to reset a trust password on one side of the trust. This includes implicit trusts between child and parent domains as well as explicit trusts between this domain (the trusting domain) and another domain (the trusted domain).

Reset the password on only the trusting domain side of the trust, also known as the incoming trust (the side where this domain belongs). Then, use the same password on the trusted domain side of the trust, also known as the outgoing trust. Reset the password of the outgoing trust when you restore the first DC in each of the other (trusted) domains.

Resetting the trust password ensures that the DC does not replicate with potentially bad DCs outside its domain. By setting the same trust password while restoring the first DC in each of the domains, you ensure that this DC replicates with each of the recovered DCs. Subsequent DCs in the domain that are recovered by installing AD DS will automatically replicate these new passwords during the installation process.

**To reset a trust password on one side of the trust**

1. At a command prompt, type the following command, and then press ENTER:

   ```
   netdom experthelp trust
   ```

2. Use the syntax that this command provides for using the NetDom tool to reset the trust password.

   For example, if there are two domains in the forest—parent and child—and you are running this command on the restored DC in the parent domain, use the following command syntax:

   ```
   netdom trust parent domain name /domain:child domain name /resetOneSide
   /passwordT:password /userO:administrator /passwordO:*
   ```

   When you run this command in the child domain, use the following command syntax:

   ```
   netdom trust child domain name /domain:parent domain name /resetOneSide
   /password:password /userO:administrator /passwordO:*
   ```

| **Note** |
| --- |
| **passwordT** should be the same value on both sides of the trust. Run this command only once (unlike the **netdom resetpwd** command) because it automatically resets the password twice. |

**Adding the global catalog**

Use the following procedure to add the global catalog to a DC.

**To add the global catalog**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
2. In the console tree, expand the **Sites** container, and then select the appropriate site that contains the target server.
3. Expand the **Servers** container, and then expand the server object for the DC to which you want to add the global catalog.
4. Right-click **NTDS Settings**, and then click **Properties**.

5. Select the **Global Catalog** check box.

**To add the global catalog using Repadmin**

1. Open an elevated command prompt, type the following command, and press ENTER:

```
repadmin.exe /options DC_NAME +IS_GC
```

The following are ways to speed up the process of adding the global catalog to the DC in the root domain:

- Ideally, the DC in the root domain should be a replication partner of the restored DCs in the non-root domains. If so, confirm that the Knowledge Consistency Checker (KCC) has created the corresponding **repsFrom** object for the source DC and partition in the root DC. You can confirm this by running the **repadmin /showreps /v** command.

- If there is no **repsFrom** object created, create this object for the configuration partition. This way, the DC in the root domain can determine which DCs in the non-root domain have been deleted. You can do this with the following commands:

  ```
  repadmin /add ConfigurationNamingContext DestinationDomainController
  SourceDomainControllerCNAME
  ```

  ```
  repadmin /options DSA -Disable_NTDSCONN_XLATE
  ```

  The format for the *SourceDomainControllerCNAME* is:

  ```
  sourceDCGuid._msdcs.root domain
  ```

  For example, the repadmin /add command for the configuration partition of the contoso.com domain could be:

  ```
  repadmin /add cn=configuration,DC=contoso,DC=com DC01 937ef930-7356-43c8-88dc-
  8baaaa781cf6._msdcs.dDSP17A22.contoso.com
  ```

- If the **repsFrom** object is present, try to sync the DC in the root domain with the DC in the non-root domain as follows:

  ```
  Repadmin /sync DomainNamingContext DestinationDomainController
  SourceDomainControllerGUID
  ```

  Where *DestinationDomainController* is the DC in the root domain and *SourceDomainController* is the restored DC in the non-root domain.

- The root domain DNS server should have the alias (CNAME) resource records for the source DC. Ensure that the parent DNS zone contains delegation resource records (name server (NS) and host (A) resource records) for the correct DCs (the DCs that have been restored from backup) in the child zone.

- Make sure that the DC in the root domain is contacting the correct Key Distribution Center (KDC) in the non-root domain. To test this, at the command prompt, type the following command, and then press ENTER:

  ```
  nltest /dsgetdc:nonroot domain name /KDC /Force
  ```

**Resources to verify replication is working**

After you have restored or re-installed all DCs, you can verify that AD DS and SYSVOL are recovered and replicating correctly by using **repadmin /replsum**, which runs on any version of Windows Server.

**Tip**

You can also download and run the Active Directory Replication Status Tool (ADReplStatus), a free tool that monitors replication status of DCs and reports errors. ADReplStatus requires .NET Framework 4, which will be installed if it is not already present.

Check the DFS Replication log in Event Viewer for Event ID 4602 (or File Replication Service event ID 13516), which indicates SYSVOL has been initialized.

If the first recovered DC logs Event ID 4614 ("the domain controller is waiting to perform initial replication. The replicated folder will remain in the initial synchronization state until it has replicated with its partner") in the DFS Replication log, then Event ID 4602 does not appear and you need to perform the following manual steps to recover SYSVOL if it is replicated by DFSR:

1. When DFSR Event 4612 appears on the first restored DC perform a manual authoritative restore as described in 2218556: How to force an authoritative and non-authoritative synchronization for DFSR-replicated SYSVOL (like "D4/D2" for FRS) (http://support.microsoft.com/kb/2218556).

2. Set **SysvolReady Flag** to 1 manually, as described in 947022 The NETLOGON share is not present after you install Active Directory Domain Services on a new full or read-only Windows Server 2008-based domain controller.

You can also create a diagnostic report DFS Replication. For more information, see Create a Diagnostic Report for DFS Replication and DFS Step-by-Step Guide for Windows Server 2008. If the server is running Windows Server 2008 R2, you can use dfsrdiag.exe ReplicationState command line switch.

You can also run the Replications test using dcdiag.exe to check for replication errors. For more information, see Knowledge Base article 249256.

# Appendix B: Frequently Asked Questions

Updated: April 25, 2013

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

This appendix contains frequently asked questions (FAQs) regarding forest recovery:

- What can I do to speed up recovery?
- Can I automate the forest recovery process?

**What can I do to speed up recovery?**

Although speed of recovery is not the primary goal of this guide, you can achieve shorter recovery times by:

- Creating a detailed forest recovery plan, updating it on a regular basis, and practicing it in a simulated test environment of reasonable size at least once a year
- Using virtualized domain controller (DC) cloning

  Virtualized DC cloning expedites the process to get additional DCs running after one DC is restored from backup in each domain. The additional virtualized DCs can be cloned rather than waiting for potentially lengthy AD DS installations to be completed and for the completion of non-critical replication after installation.

  Forests where virtual DCs are hosted in a relatively small number of well-connected data centers potentially benefit most from cloning during recovery. However, any environment where multiple virtualized DCs for the same domain are co-located on the same hypervisor host should benefit.
- Deploying read-only domain controllers (RODCs)

  RODCs can provide business continuity during the recovery process because they do not have to be disconnected from the network as writable DCs do. RODCs do not perform outbound replication. Therefore, they do not present the same risk that writable DCs pose for replicating damaging data back into the recovered environment.

Other factors that affect the duration of the forest recovery process include the following:

- When you restore DCs from backups, it takes time to:
  - Locate the physical backup media, such as tapes.
  - Reinstall the operating system.
  - Restore data from backup media.

  You can reduce the time required to reinstall the operating system and restore data from backup by performing full server recovery instead of system state restore. Because full server recovery is binary-based, it completes much faster than system state restore.

  However, if the server contains data that is excluded from system state data that you do not want to restore, full server recovery might not be a viable alternative to system state restore. Consider the advantages of performing a full server recovery instead of a system state restore for your servers specifically, and prepare accordingly by performing the appropriate type of backup that you plan to restore later.
- When you rebuild DCs, it takes time to replicate data for network-based promotions.

You can decrease the time required for restoring DCs by performing the following steps:

- Reduce the time for retrieving backup media by:

    - Using the Active Directory Database Mounting Tool (Dsamain.exe) to identify the best backup to use for restore operations. For more information about using the Active Directory Database Mounting Tool, see the Active Directory Database Mounting Tool Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkId=132577).

    - Labeling the backup media clearly and storing the media in an organized fashion at a convenient, yet secure, location that allows fast retrieval.

    - Using the Volume Shadow Copy Service with a storage area network (SAN) to maintain backups from different points in time. For more information, see Windows Server 2003 Active Directory Fast Recovery with Volume Shadow Copy Service and Virtual Disk Service (http://go.microsoft.com/fwlink/?LinkId=70781).

- Force the removal of AD DS from the DCs instead of reinstalling the operating system. If the cause of the forest-wide failure has been identified to be purely within the scope of AD DS, you do not have to reinstall the operating system on the DCs.

  For more information about forcing the removal of AD DS from a DC that runs Windows Server 2008 or later, see Forcing the Removal of a Windows Server 2008 Domain Controller (http://go.microsoft.com/fwlink/?LinkId=132627). For more information about forcing the removal of AD DS from a DC that runs Windows Server 2003, see article 332199 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=70780).

- Use faster tape devices or disk backups to reduce the time that is required for restore operations.

You can also help accelerate AD DS installations by using the Install from Media (IFM) feature to rebuild DCs in each domain. IFM reduces the replication latency that is incurred when you rebuild DCs in each domain.

Businesses that have a more aggressive service-level agreement (SLA) might consider altering the forest recovery procedures to speed recovery.

**Can I automate the forest recovery process?**
Because of the complex and critical nature of the forest recovery process, there is currently no end-to-end automation of it. The forest recovery process is more a logistical and organizational challenge of restoring business continuity than a technical problem of process automation. Therefore, the individual who administers the environment should create a forest recovery plan that is specific to that environment and then automate sections of it that can be automated successfully.

You can perform most of the forest recovery steps by using command-line tools. Therefore, most of the steps are scriptable. For example, Ntdsutil.exe is one of the most frequently used tools in the forest recovery process.

Although scripts can speed recovery, you must thoroughly test these scripts before you apply them in a real environment. Also, you must update them according to changes in the Active Directory environment, such as the addition of a new domain or DC, or a new version of Active Directory.

# Appendix C: Recovering a Single Domain within a Multidomain Forest

Published: April 25, 2013

Updated: April 25, 2013

Applies To: Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

There can be times when it is necessary to recover only a single domain within a forest that has multiple domains, rather than a full forest recovery. This topic covers considerations for recovering a single domain and possible strategies for recovery.

A single domain recovery presents a unique challenge for rebuilding global catalog (GC) servers. For example, if the first domain controller (DC) for the domain is restored from a backup that was created one week earlier, then all other GCs in the forest will have more up-to-date data for that domain than the restored DC. To re-establish GC data consistency, there are a couple options:

- Unhost and then rehost all GCs in the forest, except those in the recovered domain, at the same time.
- Follow the forest recovery process to recover the domain, and then remove lingering objects from GCs in other domains.

The following sections provide general considerations for each option. The complete set of steps that need to be done for the recovery will vary for different Active Directory environments.

**Rehost all GCs**
Rehosting all GCs can be done using repadmin /unhost and repadmin /rehost commands (part of repadmin /experthelp). You would run the repadmin commands on every GC in each domain that is not recovered.

| ⚠**Warning** |
| --- |
| The password of the built-in Administrator account for all domains must be ready for use in case a problem prevents access to a GC for logon. |

This option can be advantageous for a small organization that has only a few domain controllers for each domain. All of the GCs could be rebuilt on a Friday night and, if necessary, complete replication for all read-only domain partitions before Monday morning. But if you need to recover a large domain that covers sites across the globe, rehosting the read-only domain partition on all GCs for other domains can significantly impact operations and potentially require down time.

**Remove lingering objects**
Similar to the forest recovery process, you restore one DC from backup in the domain that you need to recover, perform metadata cleanup of remaining DCs, and then re-install AD DS to build out the domain. On the GCs of all other domains in the forest, you remove the lingering objects for the read-only partition of the recovered domain.

The source for the lingering object cleanup must be a DC in the recovered domain. To be certain that the source DC does not have any lingering objects for any domain partitions, you can remove the global catalog if it was a GC.

Removing lingering objects is advantageous for larger organizations that cannot risk the down time associated with the other options.

For more information, see Use Repadmin to remove lingering objects.

# Appendix D: Forest Recovery with Windows Server 2003 Domain Controllers

Published: April 25, 2013

Updated: April 25, 2013

Applies To: Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

This topic includes forest recovery procedures for domain controllers (DCs) that run Windows Server 2003. The general process for forest recovery is no different with Windows Server 2003 DCs, but specific procedures can differ because of different tools. For example, Ntdsutil.exe can be used to backup and restore DCs that run Windows Server 2003 DCs, whereas Windows Server Backup or Wbadmin.exe is used for DCs that run Windows Server 2008 or later.

- Backing up the System State data
- Performing a nonauthoritative restore
- Install and configure the DNS Server service

**Backing up the System State data**

Use the following procedure to back up the System State data, along with any other data you have selected for the current backup operation, of a DC that runs Windows Server 2003. Windows Server 2003 includes the Ntbackup tool, which you can use to back up System State data.

Membership in **Administrators** or **Backup Operators**, or equivalent, is the minimum required to back up files and folders. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

If you are backing up the System State data to a tape, and the Backup program indicates that there is no unused media available, you might have to use Removable Storage. This adds your tape to the free media pool so that Backup can use it.

You can only back up the System State data on a local computer. You cannot back it up on a remote computer.

**To back up the System State data on a domain controller that runs Windows Server 2003**

1. Click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
2. On the **Welcome** page, click **Advanced Mode**.
3. On the **Backup** tab, select the check box for any drive, folder, or file that you want to back up.
4. Select the **System State** check box.
5. Click **Start Backup**.

**Performing a nonauthoritative restore**

Use the following procedure to perform a nonauthoritative restore of a DC that runs Windows Server 2003. By performing a nonauthoritative restore on Active Directory in Windows Server 2003, you automatically perform a nonauthoritative restore of SYSVOL. No additional steps are required.

| **Note** |
| --- |
| If you are also reinstalling the Windows Server 2003 operating system, you might or might not join the computer to the domain and you can give any name to the computer during setup of the operating system. Do not install Active Directory. After reinstalling the operating system, go directly to step 4. |

On Windows Server 2003 domain controllers where you have restored only system state data, you need to also reinstall any software applications that were running on DCs before recovery. Restoring AD DS on the first DC in the domain also restores the registry because they both are part of System State data. Keep this in mind if you had any applications running on these DCs and if they had any information stored in the registry.

To save time required to re-install software, determine if applications that need to be installed on the DCs are compatible with virtual DC cloning. Such applications can be installed on the source DC prior to cloning in order to save the time and effort required to install them on the cloned virtual DCs.

**To perform a nonauthoritative restore**

1. After you start the DC, press F8 to restart the computer in Directory Services Restore Mode (DSRM).
2. Select **Directory Services Restore Mode (Windows domain controllers only)**.
3. Select the operating system that you want to start in restore mode.
4. Log on as an administrator (you can only use a local computer account, no domain logon option is available).
5. At a command prompt, type **ntbackup**, and then press ENTER.
6. On the **Welcome** page, click **Advanced Mode**, and then select the **Restore and Manage Media** tab. (Do not select **Restore Wizard**.)
7. Select the appropriate backup file to restore from and ensure that the **System disk** and **System State** check boxes are selected.
8. Click **Start Restore**.
9. When the restore operation is complete, restart the computer.

Use the following procedure to perform an authoritative (also known as primary) restore of SYSVOL on a DC that runs Windows Server 2003. Perform this procedure only on the first Windows Server 2003 DC that is restored in the domain.

**To perform an authoritative restore of SYSVOL**

1. Perform steps 1 through 8 in the previous procedure.
2. In the **Confirm Restore** dialog box, click **Advanced**.
3. To perform an authoritative restore of SYSVOL, select the check box **When restoring replicated data sets, mark the restored data as the primary data for all replicas**.

   | **Note** |
   | --- |
   | Marking the restored data as the primary data in the Backup is equivalent to setting the **BurFlags** entry to D4 under the following registry subkey:<br><br>**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Cumulative Replica Sets***GUID* |

4. When the restore operation is complete, restart the computer.

**Install and configure the DNS Server service**

If the DC that you restored from backup is running Windows Server 2003, you can install DNS server without connecting the DC to any network.

**To install and configure the DNS Server service**

1. Open Windows Components Wizard. To open the wizard:
   - Click **Start**, click **Control Panel**, and then click **Add or Remove Programs**.
   - Click **Add/Remove Windows Components**.
2. In **Components**, select the **Networking Services** check box, and then click **Details**.
3. In **Subcomponents of Networking Services**, select the **Domain Name System (DNS)** check box, click **OK**, and then click **Next**.
4. If you are prompted, in **Copy files from**, type the full path of the distribution files, and then click **OK**.

After the installation, complete the following steps to configure the DNS server.

5. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **DNS**.

6. Create DNS zones for the same DNS domain names that were hosted on the DNS servers before the critical malfunction. For more information, see Add a Forward Lookup Zone (http://go.microsoft.com/fwlink/?LinkId=74574).

7. Configure the DNS data as it existed before the critical malfunction. For example:

   - Configure DNS zones to be stored in AD DS. For more information, see Change the Zone Type (http://go.microsoft.com/fwlink/?LinkId=74579).

   - Configure the DNS zone that is authoritative for domain controller locator (DC Locator) resource records to allow secure dynamic update. For more information, see Allow Only Secure Dynamic Updates (http://go.microsoft.com/fwlink/?LinkId=74580).

8. Ensure that the parent DNS zone contains delegation resource records (name server (NS) and glue host (A) resource records) for the child zone that is hosted on this DNS server. For more information, see Create a Zone Delegation (http://go.microsoft.com/fwlink/?LinkId=74562).

9. After you configure DNS, at the command prompt, type the following command, and then press ENTER:

   **net stop netlogon**

10. Type the following command, and then press ENTER:

    **net start netlogon**

| **Note** |
| --- |
| Net Logon will register the DC Locator resource records in DNS for this DC. If you are installing the DNS Server service on a server in the child domain, this DC will not be able to register its records immediately. This is because it is currently isolated as part of the recovery process, and its primary DNS server is the forest root DNS server. Configure this computer with the same IP address as it had before the disaster to avoid DC service lookup failures. |

# Additional Resources

Updated: April 25, 2013

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

This section contains additional resources related to forest recovery.

The following resources are useful for recovering domain controllers that run Windows Server 2012:

- Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100)
- Maintaining Business Continuity of Virtualized Environments with Hyper-V Replica: scenario overview. For more information about setting up Hyper-V Replica, see
- Deploy Hyper-V Replica
- Remote Server Administration Tools for Windows 8

The following resources are useful for recovering domain controllers that run Windows Server 2008 R2 or Windows Server 2008:

- Remote Server Administration Tools for Windows 7 with Service Pack 1 (SP1)
- Microsoft Remote Server Administration Tools for Windows Vista (http://go.microsoft.com/fwlink/?LinkID=115118)
- Active Directory Database Mounting Tool Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkId=132577)
- Ntdsutil (http://go.microsoft.com/fwlink/?LinkId=132629)
- Forcing the Removal of a Windows Server 2008 Domain Controller (http://go.microsoft.com/fwlink/?LinkID=132627)
- Installing Active Directory Domain Services from Media (http://go.microsoft.com/fwlink/?LinkId=132630)
- Performing an Unscheduled Backup of a Domain Controller (http://go.microsoft.com/fwlink/?LinkId=132632)
- Performing a Nonauthoritative Restore of AD DS (http://go.microsoft.com/fwlink/?LinkId=132637)

# Enabling Advanced Windows Server 2003 Active Directory Features

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Microsoft® Windows® Server 2003 Active Directory® directory service enables you to introduce advanced features into your environment by raising the domain or forest functional level. You can raise the functional level when all domain controllers in the domain or forest are running an appropriate version of Windows. Raising the functional level allows you to introduce new features but also limits the versions of Windows that can run on domain controllers in your environment.

**In This Chapter**

Overview of Enabling Advanced Active Directory Features

Preparing to Enable Functional Levels

Enabling Windows Server 2003 Active Directory Functional Levels

Additional Resources for Enabling Advanced Windows Server 2003 Active Directory Features

**Related Information**

- For more information about domain and forest functional levels, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).
- For more information about enabling functional levels in a new Microsoft® Windows® Server 2003 environment, see "Deploying the Windows Server 2003 Forest Root Domain" in this book.
- For more information about enabling functional levels after upgrading from Microsoft® Windows NT® 4.0, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.
- For more information about enabling functional levels after upgrading from Microsoft® Windows® 2000, see "Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" in this book.

# Overview of Enabling Advanced Active Directory Features

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Functional levels in Windows Server 2003 Active Directory enable you to implement advanced features — such as efficient group membership replication, deactivation and redefinition of attributes and classes in the schema, and domain rename — that require that domain controllers within a domain or forest be running the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; or Windows® Server 2003, Datacenter Edition operating systems. If you want to enable these advanced Windows Server 2003 Active Directory features in your organization, you must raise the domain and/or forest to the appropriate functional level.

Before you can identify and enable the functional level that best meets the needs of your organization, you must identify the Windows operating systems that you are currently running and that you plan to maintain in your environment after you deploy Windows Server 2003.

If you are currently running Windows NT 4.0 and you do not plan to deploy Windows 2000 in your environment, after you deploy the first Windows Server 2003–based domain controller, raise the forest functional level to Windows Server 2003 interim to take advantage of the advanced features available at that forest functional level.

If you are currently running both Windows 2000 and Windows NT 4.0 in your environment, after you deploy a Windows Server 2003–based domain controller, keep the forest functional level set to Windows 2000. This enables you to take advantage of all advanced features available at that forest functional level.

If you are currently running only Windows 2000 in your environment or you are planning to install any number of Windows 2000–based domain controllers in the future, after you deploy a Windows Server 2003–based domain controller, keep the forest functional level set to Windows 2000. This enables you to take advantage of all advanced features available at that forest functional level.

If you are deploying a new Windows Server 2003 environment and plan to run only Windows Server 2003–based domain controllers, after you deploy the first Windows Server 2003–based domain controller you can raise the forest functional level to Windows Server 2003 to take advantage of all available Windows Server 2003 Active Directory features.

**Note**

- For a list of the job aids that are available to assist you in enabling functional levels, see "Additional Resources for Enabling Advanced Windows Server 2003 Active Directory Features" later in this chapter.

# Process for Enabling Advanced Active Directory Features

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Enabling advanced Active Directory features involves identifying the operating systems that are running on the domain controllers in your environment and the functional level that best meets the needs of your organization based on your existing infrastructure, and raising the domain or forest functional level as appropriate. Figure 5.1 shows the process for enabling advanced Active Directory features.

**Figure 5.1 Enabling Advanced Active Directory Features**



# Functional Levels Background Information

Updated: July 9, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 Active Directory functional levels expand on the mixed and native modes introduced in the Windows 2000 operating system. In Windows 2000, a mixed mode domain supports domain controllers running either Windows 2000 or the Windows NT 4.0 operating system. Domains in native mode only support Windows 2000–based domain controllers. If all domain controllers in a mixed mode domain are upgraded to Windows 2000, the domain administrator can change the mode to native, making additional Windows 2000 features available.

In Windows Server 2003, the functional level of a domain or forest defines the set of advanced Windows Server 2003 Active Directory features that are available in that domain or forest. The functional level of a domain or forest also defines the set of Windows operating systems that can run on the domain controllers in that domain or forest.

**Note**

- The functional level of a domain or forest defines only the set of Windows operating systems that can run on domain controllers. It does not define the client operating systems that are supported in the forest.

When the first Windows Server 2003–based domain controller is deployed in a domain or forest, a set of default Active Directory features becomes available. Table 5.1 summarizes the Active Directory features that are available by default on any domain controller running Windows Server 2003.

**Table 5.1 Default Windows Server 2003 Active Directory Features**

| Feature | Functionality |
|---|---|
| Multiple selection of user objects | Allows you to modify common attributes of multiple user objects at one time. |
| Drag and drop functionality | Allows you to move Active Directory objects from container to container by dragging one or more objects to a location in the domain hierarchy. You can also add objects to group membership lists by dragging one or more objects (including other group objects) to the target group. |
| Efficient search capabilities | Search functionality is object-oriented and provides an efficient search that minimizes network traffic associated with browsing objects. |
| Saved queries | Allows you to save commonly used search parameters for reuse in Active Directory Users and Computers |
| Active Directory command-line tools | Allows you to run new directory service commands for administration scenarios. |
| InetOrgPerson class | The inetOrgPerson class has been added to the base schema as a security principal and can be used in the same manner as the user class. |
| Application directory partitions | Allows you to configure the replication scope for application-specific data among domain controllers. For example, you can control the replication scope of Domain Name System (DNS) zone data stored in Active Directory so that only specific domain controllers in the forest participate in DNS zone replication. |
| Ability to add additional domain controllers by using backup media | Reduces the time it takes to add an additional domain controller in an existing domain by using backup media. |

| | |
|---|---|
| Universal group membership caching | Prevents the need to locate a global catalog across a wide area network (WAN) when logging on by storing universal group membership information on an authenticating domain controller. |
| Secure Lightweight Directory Access Protocol (LDAP) traffic | Active Directory administrative tools sign and encrypt all LDAP traffic by default. Signing LDAP traffic guarantees that the packaged data comes from a known source and that it has not been tampered with. |
| Partial synchronization of the global catalog | Provides improved replication of the global catalog when schema changes add attributes to the global catalog partial attribute set. Only the new attributes are replicated, not the entire global catalog. |
| Active Directory quotas | Quotas can be specified in Active Directory to control the number of objects a user, group, or computer can own in a given directory partition. Members of the Domain Administrators and Enterprise Administrators groups are exempt from quotas. |

For more information about the default Active Directory features that are available on any Windows Server 2003 domain controller, see "New features for Active Directory" in Help and Support Center for Windows Server 2003.

When the first Windows Server 2003–based domain controller is deployed in a domain or forest, the domain or forest operates by default at the lowest functional level that is possible in that environment. This allows you to take advantage of the default Active Directory features while running versions of Windows earlier than Windows Server 2003.

When you raise the functional level of a domain or forest, a set of advanced features becomes available. For example, the Windows Server 2003 interim forest functional level supports more features than the Windows 2000 forest functional level, but fewer features than the Windows Server 2003 forest functional level supports. Windows Server 2003 is the highest functional level that is available for a domain or forest. The Windows Server 2003 functional level supports the most advanced Active Directory features; however, only Windows Server 2003 domain controllers can operate in that domain or forest.

If you raise the domain functional level to Windows Server 2003, you cannot introduce any domain controllers that are running versions of Windows earlier than Windows Server 2003 into that domain. This applies to the forest functional level as well.

Table 5.2 lists the Windows Server 2003 domain functional levels, the operating systems that they support, and the Windows Server 2003 features that are available at each domain functional level.

**Table 5.2 Windows Server 2003 Domain Functional Levels**

| **Windows Server 2003 Domain Functional Level** | **Supported Domain Controller Operating Systems** | **Advanced Features Available at Each Domain Functional Level** |
|---|---|---|
| Windows 2000 mixed | Windows NT 4.0<br>Windows 2000<br>Windows Server 2003 | All default Active Directory features, and:<br>• Universal Groups are enabled for distribution groups, but are disabled for security groups. |

| | | |
|---|---|---|
| Windows 2000 native | Windows 2000<br>Windows Server 2003 | All default Active Directory features, all features from the Windows 2000 mixed domain functional level, and:<br>• Universal Groups are enabled for both distribution and security groups.<br>• Group conversion is enabled, allowing conversion between security and distribution groups.<br>• Group nesting is available, allowing nesting of groups within other groups.<br>• Security identifier (SID) history is available, allowing the migration of security principals from one domain to another. |
| Windows Server 2003 interim | Windows NT 4.0<br>Windows Server 2003 | Same as Windows 2000 mixed. |
| Windows Server 2003 | Windows Server 2003 | All default Active Directory features, all features from the Windows 2000 native domain functional level, and:<br>• Supports new functionality of the netdom.exe tool to prepare domain controllers for rename. It is recommended that you rename a domain controller by using netdom.exe to ensure that all appropriate steps are taken.<br>• Enables updates to the logon timestamp attribute. The *lastLogonTimestamp* attribute is updated with the last logon time of the user or computer. This attribute is replicated within the domain.<br>• Provides the ability to set the *userPassword* attribute as the effective password on inetOrgPerson and user objects.<br>• Provides the ability to redirect the Users and Computers containers in order to define a new well-known location for user and computer accounts.<br>• Allows for authorization manager to store its authorization policies in Active Directory.<br>• Includes constrained delegation, which allows applications to take advantage of the secure delegation of user credentials by means of Kerberos authentication protocol. Delegation can be configured to be allowed only to specific destination services.<br>• Supports selective authentication, by which it is possible to specify the users and groups from a trusted forest who are allowed to authenticate to resource servers in a trusting forest. |

Table 5.3 lists the Windows Server 2003 forest functional levels, the operating systems that they support, and the Windows Server 2003 features that are available at each forest functional level.

**Table 5.3 Windows Server 2003 Forest Functional Levels**

| Windows Server 2003 Forest Functional Level | Supported Domain Controller Operating Systems | Advanced Features Available at Each Forest Functional Level |
|---|---|---|

| | | |
|---|---|---|
| Windows 2000 | Windows NT 4.0<br>Windows 2000<br>Windows Server 2003 | All default Active Directory features. |
| Windows Server 2003 interim | Windows NT 4.0<br>Windows Server 2003 | All default Active Directory features, and:<br>• Linked value replication.<br>• Improved KCC algorithms and scalability.<br>• The following attributes included in the global catalog:<br>  • Ms-DS-Trust-Forest-Trust-Info<br>  • Trust-Direction<br>  • Trust-Attributes<br>  • Trust-Type<br>  • Trust-Partner<br>  • Security-Identifier<br>  • Ms-DS-Entry-Time-To-Die<br>  • MSMQ-Secured-Source<br>  • MSMQ-Multicast-Address<br>  • Print-Memory<br>  • Print-Rate<br>  • Print-Rate-Unit<br>  • MS-DRM-Identity-Certificate |

| | | |
|---|---|---|
| Windows Server 2003 | Windows Server 2003 | All Active Directory features available at the Windows Server 2003 interim level, and:<br>• The ability to create instances of the dynamic auxiliary class called *dynamicObject* in a domain naming context.<br>• The ability to convert an *inetOrgPerson* object instance into a *User* object instance and vice versa.<br>• The ability to create instances of the new group types *basic* and *query based*, used by the role–based Authorization Manager.<br>• Deactivation and redefinition of attributes and classes in the schema. The following attributes can be reused: ldapDisplayName, schemaIdGuid, OID, and mapiID.<br>• Forest trust.<br>• Domain rename. |

**Guidelines for Raising Domain Functional Levels**
The following guidelines apply to raising the domain functional level:

- You must be a member of the Domain Admins group to raise the domain functional level.
- You can raise the domain functional level on the primary domain controller (PDC) emulator operations master only. The Active Directory administrative tools used to raise the domain functional level (Active Directory Domains and Trusts and Active Directory Users and Computers) automatically target the PDC emulator when you raise the domain functional level.
- You can raise the functional level of a domain only if all domain controllers in the domain are running the version or versions of Windows that the new functional level supports.
- You cannot lower the functional level of a domain after it has been raised.

**Guidelines for Raising Forest Functional Levels**
The following guidelines apply to raising the forest functional level:

- You must be a member of the Enterprise Admins group to raise the forest functional level.
- You can raise the forest functional level on the schema operations master only. The Active Directory Domains and Trusts console automatically targets the schema operations master when you raise the forest functional level.
- You can raise the functional level of a forest only if all domain controllers in the forest are running the version or versions of Windows that the new functional level supports.
- You can raise the forest to the Windows Server 2003 functional level only if all domains are at either the Windows 2000 native or Windows Server 2003 functional level.
- You cannot lower the functional level of a forest after it has been raised.

**Important**

- Raising the domain and forest functional levels are one-way operations that cannot be reversed. In the event that you need to revert to a lower functional level, you need to rebuild the domain or forest or restore it from a backup. For more

information about domain and forest recovery, see the Best Practices: Active Directory Forest Recovery link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

When you raise the forest functional level to Windows Server 2003, Active Directory automatically raises all domains that are operating at the Windows 2000 native domain functional level to the Windows Server 2003 domain functional level. However, if any domains in your environment are operating at the Windows 2000 mixed domain functional level, you cannot raise the forest functional level to Windows Server 2003.

For more information about raising functional levels, see "Raising domain and forest functional levels" in Help and Support Center for Windows Server 2003.

# Preparing to Enable Functional Levels

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can enable domain and forest functional levels, you need to evaluate your current environment and identify the functional level scenario that best meets the needs of your organization. For a worksheet to assist you in preparing to enable functional levels, see "Assess Your Current Environment" later in this chapter.

Figure 5.2 shows the process for preparing to enable functional levels.

**Figure 5.2 Preparing to Enable Functional Levels**



# Assess Your Current Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assess your current environment by identifying the domains in your forest, the domain controllers that are located in each domain, the operating system that each domain controller is running, and the date that you plan to upgrade the domain controller. If you plan to retire a domain controller, document the reasons for this decision.

Circumstances that might prevent you from upgrading an earlier version of the Windows operating system and enabling the Windows Server 2003 functional level include:

- Insufficient hardware
- A domain controller running an antivirus program that is incompatible with Windows Server 2003
- Use of a version-specific program that does not run on Windows Server 2003
- The need to perform a Service Pack upgrade

Documenting this information will help you identify the steps that are required for you to achieve a fully functional Windows Server 2003 environment.

For a worksheet to assist you in assessing your current environment, see "Domain Controller Assessment" (DSSPFL_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Domain Controller Assessment" on the Web at http://www.microsoft.com/reskit). Complete a separate worksheet for each domain, regardless of your forest structure.

Figure 5.3 shows an example of a completed worksheet for a domain assessment.

**Figure 5.3 Example of a Domain Controller Assessment Worksheet**

| Domain Controller Assessment | | | | |
|---|---|---|---|---|
| Prepared By | Forest Owner | | Date | 1/2/03 |
| Forest Name | Concorp.contoso.com | | | |
| Domain Name | Noam.concorp.contoso.com<br>☑ Mixed Mode ☐ Native Mode | | | |
| DC Name | Operating System/Service Pack | Deployment Blocker | Upgrade Date | Retire Date |
| SEA-NOAM-DC01 | Windows NT 4.0 SP6 | Insufficient hardware | | 3/10/03 |
| SEA-NOAM-DC02 | Windows 2000 Advanced Server | None | 3/3/03 | |
| SEA-NOAM-DC03 | Windows 2000 Advanced Server | None | 3/3/03 | |
| SEA-NOAM-DC04 | Windows NT 4.0 SP4 | Need Service Pack Upgrade to SP6 | 3/10/03 | |

# Identify Your Functional Level Scenario

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you assess your current environment, identify the functional level scenario — Windows NT 4.0 environment, Windows 2000 mixed-mode environment, Windows 2000 native-mode environment, or new Windows Server 2003 forest — that applies to your organization.

**Windows NT 4.0 environment** You have a pure Windows NT 4.0 environment consisting of one or more Windows NT 4.0 PDCs and backup domain controllers (BDCs). You want to upgrade directly to Windows Server 2003 and take advantage of all Windows Server 2003 forest- and domain-level features without deploying any Windows 2000 domain controllers in the environment.

**Windows 2000 mixed mode environment** You have a mixed mode Windows 2000 domain that includes both Windows 2000 and Windows NT 4.0–based domain controllers. You want to upgrade to Windows Server 2003 to take advantage of all Windows Server 2003 forest- and domain-level features.

**Windows 2000 native mode environment** You have a native mode Windows 2000 domain consisting of only Windows 2000–based domain controllers. You want to upgrade to Windows Server 2003 to take advantage of all Windows Server 2003 forest- and domain-level features.

**New Windows Server 2003 forest** You are creating a new Windows Server 2003 forest by installing Active Directory on a Windows Server 2003–based member server. You want to take advantage of all Windows Server 2003 forest- and domain-level features.

# Enabling Windows Server 2003 Active Directory Functional Levels

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Enabling advanced Windows Server 2003 Active Directory features in your environment involves installing Windows Server 2003 Active Directory, determining the functional level that is appropriate for your environment, and then raising domain and forest functional levels to meet your requirements. If you choose to raise your existing infrastructure to the Windows Server 2003 functional level, you can take advantage of all the Windows Server 2003 Active Directory features that are available.

You can determine the current domain functional level by viewing the properties of the domain object in either Active Directory Users and Computers or Active Directory Domains and Trusts. You can determine the current forest functional level by using Active Directory Domains and Trusts to view the properties of the **Active Directory Domains and Trusts** node.

To raise the forest functional level to Windows Server 2003, use Active Directory Domains and Trusts. To raise the domain functional level to Windows Server 2003 or Windows 2000 native, use Active Directory Domains and Trusts or Active Directory Users and Computers. For more information about how to view and raise domain and forest functional levels, see "Raise the domain functional level" and "Raise the forest functional level" in Help and Support Center for Windows Server 2003.

**Figure 5.4 Enabling Windows Server 2003 Active Directory Functional Levels**



# Enabling Windows Server 2003 Functional Levels in a Windows NT 4.0 Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If all of the domain controllers in your environment are running Windows NT 4.0, and you plan to upgrade them to Windows Server 2003 *without* ever upgrading to Windows 2000 or installing a new Windows 2000–based domain controller, maintain the Windows Server 2003 interim functional level in your domains and forest until you upgrade all Windows NT 4.0 domain controllers to Windows Server 2003.

**Important**

- If you choose to raise the forest and domain functional level to Windows Server 2003 interim, you cannot return to the Windows 2000 mixed domain functional level or the Windows 2000 forest functional level, and therefore you cannot add Windows 2000–based domain controllers to the forest.

For more information about deploying Windows Server 2003 in a Windows NT 4.0 environment, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.

If you intend to add one or more Windows 2000–based domain controllers instead of having only domain controllers running Windows Server 2003 in your environment, see "Enabling Windows Server 2003 Functional Levels in a Mixed Windows 2000 Environment" later in this chapter.

**Important**

- If you are running Windows NT 4.0 or Windows 2000 domain controllers in your environment, do not raise the functional level of your domain or forest to Windows Server 2003. You cannot operate at the Windows Server 2003 functional level until all of your domain controllers are running Windows Server 2003.

Windows 2000 Active Directory group replication limits the size of groups in a Windows 2000 forest. You must divide groups that include more than 5,000 members into smaller groups when you upgrade to Windows 2000. The Windows Server 2003 interim forest functional level is ideal if the groups in any domains in your existing Windows NT 4.0 environment include more than 5,000 members. When you are operating at the Windows Server 2003 interim functional level, you can take advantage of group membership replication improvements, which support large groups of more than 5,000 members.

When upgrading your Windows NT 4.0 environment to Windows Server 2003, you can choose to do one of the following:

- Upgrade to a regional domain in an existing Windows Server 2003 forest.
- Upgrade to a single domain forest.

Whether you decide to upgrade to a regional domain in an existing Windows Server 2003 forest or upgrade to a single domain forest, if you choose to raise the forest functional level to Windows Server 2003 interim, you must remain at the Windows Server 2003 interim functional level until you upgrade all other Windows NT 4.0–based domain controllers to Windows Server 2003 or retire them from service. The Windows Server 2003 interim functional level supports both Windows NT 4.0–based domain controllers and Windows Server 2003–based domain controllers.

**Upgrading to a Regional Domain in an Existing Windows Server 2003 Forest**

When you upgrade a Windows NT 4.0 domain to a regional domain in an existing Windows Server 2003 forest, it is recommended that you raise the forest functional level of the existing forest to Windows Server 2003 interim before upgrading the Windows NT 4.0 PDC to take advantage of the added features of the Windows Server 2003 interim functional level. After you raise the forest functional level of the existing forest to Windows Server 2003 interim, the domain functional level of the forest root domain and all subsequent regional domains is set by default to Windows Server 2003 interim.

When you upgrade a Windows NT 4.0 domain to a regional domain in an existing Windows Server 2003 forest, where the forest functional level is set to Windows 2000, functional levels are set in the new regional domain to the following by default, and they remain in effect until you raise them manually:

- Windows 2000 mixed domain functional level
- Windows 2000 forest functional level

You cannot use Active Directory administrative consoles to raise the forest functional level of the existing Windows Server 2003 forest root domain to Windows Server 2003 interim. Instead, use a Lightweight Directory Access Protocol (LDAP) application such as ADSI Edit or LDP in Windows Support Tools to edit the value of the *msDS-Behavior-Version* attribute.

**To raise the forest functional level of the existing forest to Windows Server 2003 interim by using ADSI Edit**

1. In ADSI Edit, expand the Configuration partition, and expand **CN=Configuration,DC=forestname,DC=domainname,DC=com**.
2. Right-click **CN=Partitions**, and then click **Properties**.
3. Select the **msDS-Behavior-Version** attribute.
4. Click **Edit**.
5. In the **Value** field, type **1** to raise the forest functional level to Windows Server 2003 interim.
6. Click OK.

After you raise the forest functional level to Windows Server 2003 interim forest, you cannot add Windows 2000–based domain controllers to the forest.

If you are deploying a new Windows Server 2003 forest root domain and are planning to upgrade a Windows NT 4.0 domain to a regional domain in this new environment, after you raise the forest functional level to Windows Server 2003 interim, upgrade the Windows NT 4.0 domain to Windows Server 2003. Select **Child domain in an existing domain tree** when prompted by the Active Directory Installation Wizard.

For more information about deploying a Windows Server 2003 forest root domain, see "Deploying the Windows Server 2003 Forest Root Domain" in this book.

**Upgrading to a Single Domain Forest**

When upgrading to a new Windows Server 2003 single domain forest by upgrading an existing Windows NT 4.0 PDC to Windows Server 2003, you are prompted to use the Active Directory Installation Wizard to install Active Directory. The wizard gives you the option of setting the forest functional level to Windows Server 2003 interim during the Active Directory installation process.

If you set the functional level during the Active Directory installation, both the domain and forest will be set at Windows Server 2003 interim after the installation process is complete and the computer is restarted.

**Important**

- If you do not set the functional level to Windows Server 2003 interim during the Active Directory installation process, functional levels are set by default to the following:
    - Windows 2000 forest functional level
    - Windows 2000 mixed domain functional level
- Use the preceding procedure to use ADSI Edit to manually raise the forest functional level to Windows Server 2003 interim after the Active Directory installation process is complete and the computer is restarted.

# Raise the Domain Functional Level to Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you upgrade all Windows NT 4.0–based domain controllers in a domain to Windows Server 2003, you can raise the functional level of each domain in the forest to Windows Server 2003. Before you raise the domain functional level, however, you must ensure that no Windows NT 4.0–based domain controllers remain in the domain.

**Warning**

- If Windows NT 4.0–based domain controllers are running in a domain when you raise the domain functional level to Windows Server 2003, they will no longer be able to communicate with the new Windows Server 2003 domain controllers and will not receive necessary updates.

Use the following LDAP query to identify any Windows NT 4.0 domain controllers remaining in the domain. Run the LDAP query against the Domain container in Active Directory Users and Computers. If you have not manually changed the value of the *operatingSystemVersion* attribute of the computer object, this query is conclusive for domain controllers running Windows NT 4.0. You must be a member of the Domain Admins group to run the following query.

**To identify Windows NT 4.0–based domain controllers in a domain**

1. From any Windows Server 2003–based domain controller, open Active Directory Users and Computers.
2. If the domain controller is not already connected to the appropriate domain, connect it to the domain as follows:
   a. Right-click the current domain object, and then click **Connect to domain**.
   b. In the **Domain** dialog box, type the DNS name of the domain that you want to connect to, or click **Browse** to select the domain from the domain tree, and then click **OK**.
3. Right-click the domain object, and then click **Find**.
4. In the **Find** dialog box, click **Custom Search**.
5. Click the domain for which you want to change the functional level.
6. Click the **Advanced** tab.
7. In the **Enter LDAP query** box, type the following, leaving no spaces between any characters (the query is not case-sensitive):

   ```
   (&(objectCategory=computer)(operatingSystemVersion=4*)
   (userAccountControl:1.2.840.113556.1.4.803:=8192))
   ```

8. Click **Find Now**. This produces a list of the computers in the domain that are running Windows NT 4.0 and functioning as domain controllers.

   A domain controller might appear in the list for any of the following reasons:

   - The domain controller is running Windows NT 4.0 and must be upgraded.
   - The domain controller has been upgraded to Windows Server 2003, but the change has not replicated to the target domain controller.
   - The domain controller is no longer in service, but its computer object has not been removed from the domain.

Before you can change the domain functional level to Windows Server 2003, you must physically locate any domain controller in the list, determine its current status, and either upgrade or remove the domain controller as appropriate.

For more information about LDAP queries, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

# Raise the Forest Functional Level to Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After all domains are operating at the Windows Server 2003 functional level, raise the forest functional level to Windows Server 2003. This enables you to take advantage of all Windows Server 2003 forest-level features.

If any domains in the forest are still operating at the Windows Server 2003 interim functional level, you will be unable to raise the forest functional level to Windows Server 2003. Ensure that all domains are operating at the Windows Server 2003 functional level before you raise the forest functional level.

# Enabling Windows Server 2003 Functional Levels in a Mixed Windows 2000 Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your Windows 2000 forest includes one or more domains that contain Windows NT 4.0–based domain controllers, those domains are in Windows 2000 mixed mode. Domains that include only Windows 2000–based domain controllers might be in Windows 2000 mixed mode or native mode. Functional levels in a mixed Windows 2000 forest are set by default when you deploy the first Windows Server 2003–based domain controller.

For more information about deploying Windows Server 2003 in a mixed Windows 2000 environment, see "Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" in this book.

You can introduce a Windows Server 2003–based domain controller in a mixed environment in one of two ways:

- By installing a new Windows Server 2003–based domain controller.
- By upgrading an existing Windows 2000 domain controller in the forest to Windows Server 2003.

Functional levels are set at the following levels by default, and remain at these levels until they are raised manually:

- Windows 2000 mixed or Windows 2000 native domain functional level, depending on whether the domain was in mixed mode or native mode prior to the upgrade.
- Windows 2000 forest functional level.

If the domain functional level is set to Windows 2000 mixed after the initial upgrade, the domain must remain at that level for as long as Windows NT 4.0–based domain controllers are in the domain. If you upgrade all Windows NT 4.0–based domain controllers to either Windows 2000 or Windows Server 2003 and decommission the Windows NT 4.0–based domain controllers that you do not intend to upgrade, you can raise the domain functional level to Windows 2000 native.

If the domain functional level is set to Windows 2000 native after the initial upgrade, the domain must remain at that level for as long as Windows 2000–based domain controllers are operating in the domain.

**Note**

- This also applies to Windows NT 4.0 environments in which you intend to deploy one or more Windows 2000 domain controllers in the future. After the initial upgrade, the domain must remain at a functional level of Windows 2000 mixed.

After you upgrade all Windows 2000–based domain controllers to Windows Server 2003, you can raise the functional levels of the domains in the forest to Windows Server 2003. Before you raise the domain functional level, you must verify that no Windows NT 4.0–based domain controllers remain in the domain. For more information about identifying Windows NT 4.0–based domain controllers in a domain, see "Enabling Windows Server 2003 Functional Levels in a Windows NT 4.0 Environment" earlier in this chapter.

If all domain controllers in the domain are running Windows Server 2003, you can raise the domain functional level from Windows 2000 mixed to Windows Server 2003 directly. Alternatively, you can raise the functional level step by step — from Windows 2000 mixed to Windows 2000 native and then to Windows Server 2003.

After you upgrade all domain controllers in the forest to Windows Server 2003 and raise all domains to the Windows 2000 native or Windows Server 2003 functional level, you can raise the forest functional level to Windows Server 2003. This automatically raises the functional level of any remaining domains that are operating at the Windows 2000 native functional level to Windows Server 2003.

# Enabling Windows Server 2003 Functional Levels in a Native Windows 2000 Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If the domains in your Windows 2000 forest include only Windows 2000 domain controllers and are in Windows 2000 native mode, deploy a Windows Server 2003–based domain controller to enable functional levels.

For more information about deploying Windows Server 2003 in a Windows 2000 environment, see "Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" in this book.

In an environment that contains only domain controllers running Windows 2000, you can introduce a Windows Server 2003–based domain controller in one of two ways:

- By installing a new Windows Server 2003–based domain controller.
- By upgrading an existing Windows 2000 domain controller in the forest to Windows Server 2003.

Functional levels are set by default to the following levels, and they remain at these levels until they are raised manually:

- Windows 2000 native domain functional level
- Windows 2000 forest functional level

**Note**

- If your Windows 2000 forest consists solely of Windows 2000–based domain controllers, but one or more of your domains are operating in mixed mode, see "Enabling Windows Server 2003 Functional Levels in a Mixed Windows 2000 Environment" earlier in this chapter.

To take advantage of the Windows Server 2003 domain-level features without waiting to complete the upgrade of your Windows 2000 forest to Windows Server 2003, raise only the domain functional level to Windows Server 2003. Before you raise the domain functional level, you must upgrade all Windows 2000–based domain controllers in the domain to Windows Server 2003.

After you upgrade all Windows 2000–based domain controllers in the forest to Windows Server 2003, make sure that the domain functional level of each domain is set to Windows 2000 native or higher. Then raise the forest functional level to Windows Server 2003. Raising the forest functional level to Windows Server 2003 automatically raises the functional level of all domains in the forest that are set to Windows 2000 native or higher to Windows Server 2003.

# Enabling Windows Server 2003 Functional Levels in a New Windows Server 2003 Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have installed the first domain controller in a new Windows Server 2003 forest, functional levels are set by default to the following levels, and remain at these levels until they are raised manually:

- Windows 2000 mixed domain functional level
- Windows 2000 forest functional level

Functional levels are set at these levels to allow you the option of adding Windows 2000 or Windows NT 4.0–based domain controllers to your new Windows Server 2003 forest.

After you create a forest root domain, the domain functional level for each additional domain that you add to the Windows Server 2003 forest is set to Windows 2000 mixed.

**Important**

- If the forest is operating at the Windows Server 2003 functional level, and you attempt to install Active Directory on a Windows 2000–based member server, the installation will fail. If you install Active Directory on a Windows Server 2003–based member server in order to create a new regional domain, the domain functional level is set to Windows Server 2003.

After you deploy the new Windows Server 2003 forest and the domain functional level is set in all domains, raise the domain functional level and then the forest functional level to Windows Server 2003. This enables you to take advantage of all Windows Server 2003 forest- and domain-level features. Thereafter, all new domains that you create are set at the Windows Server 2003 domain functional level.

# Additional Resources for Enabling Advanced Windows Server 2003 Active Directory Features

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Deploying the Windows Server 2003 Forest Root Domain" in this book.
- "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.
- "Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" in this book.
- The *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit) for more information about Active Directory functional levels.
- Article 322692, "HOW TO: Raise the domain functional level in Windows Server 2003," in the Microsoft Knowledge Base for more information about raising functional levels. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Tools**

- ADSI Edit

  The ADSI Edit tool (Adsiedit.exe) is a Microsoft Management Console snap-in that you can use to edit objects in the Active Directory database. For more information about Adsiedit.exe, in Help and Support Center for Windows Server 2003, click **Tools,** and then click Install Windows Support Tools.

- LDP

  LDP provides an interface to perform LDAP operations against Active Directory. For more information about LDP, in Help and Support Center for Windows Server 2003, click **Tools,** and then click Install Windows Support Tools.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "New features for Active Directory" in Help and Support Center for Windows Server 2003 for more information about the default Active Directory features that are available on any Windows Server 2003 domain controller.
- "Raising domain and forest functional levels" in Help and Support Center for Windows Server 2003 for more information about raising functional levels.

**Related Job Aids**

- "Domain Controller Assessment" (DSSPFL_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Controller Assessment" on the Web at http://microsoft.com/reskit).

# Deploying the Windows Server 2003 Forest Root Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying your forest root domain is the first step in deploying the Active Directory® directory service infrastructure in your organization.

**In This Chapter**


Overview of Deploying the Forest Root Domain

Reviewing the Active Directory Design

Configuring DNS for the Forest Root Domain

Creating the Forest Root Domain

Raising the Functional Level

Additional Resources for Deploying the Root Domain


**Related Information**

- For more information about designing the Active Directory logical structure and the Domain Name System (DNS) infrastructure needed to support Active Directory, see "Designing the Active Directory Logical Structure" in this book.
- For more information about designing the Active Directory site topology, see "Designing the Site Topology" in this book.
- For more information about domain controller capacity planning for Active Directory, see "Planning Domain Controller Capacity" in this book.
- For more information about designing and deploying a DNS infrastructure for name resolution for your network, see "Process for Deploying DNS" in *Deploying Network Services* in this kit.

# Overview of Deploying the Forest Root Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first domain that you create in your Active Directoryforest is automatically designated as the forest root domain. The forest root domain provides the foundation for your Active Directory forest infrastructure. You must create the forest root domain before you create regional domains or upgrade other Microsoft® Windows NT® 4.0 domains in order to join them to an existing forest. In addition, services that are running on forest root domain controllers, such as the Kerberos version 5 authentication protocol, must be highly available to ensure that users maintain access to resources throughout the forest.

Before you deploy your forest root domain, your design team must design your Active Directory logical structure and site topology and plan your hardware requirements for domain controllers that are running the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems. During the forest root domain deployment, you begin to implement the Active Directory design that your design team has provided, including the DNS infrastructure that Active Directory requires.

The forest owner is responsible for deploying the forest root domain. After the forest root domain deployment is complete, deploy the remainder of your Active Directory forest as specified by your Active Directory design. The tasks that you must perform to deploy the remainder of your Active Directory forest depend on whether your design specifies a single domain forest or a multiple domain forest.

- **Single domain forest.** If your Active Directory forest design requires only a single domain, then the forest root domain will also contain all your users, groups, and resources. To deploy this model, you can create an organizational unit (OU) structure after the forest root domain deployment is complete. Then you can restructure Windows NT account and resource domains into the forest root domain.

- **Multiple domain forest.** In a multiple domain design, the forest root domain can be a dedicated root used only for administration of the forest, or it can contain users, groups, and resources in addition to the forest administration accounts. Once the forest root domain is deployed, the forest owner will create one or more regional child domains to complete the Active Directory forest hierarchy. The regional domains can be created either by upgrading existing Windows NT 4.0 or Microsoft® Windows® 2000 domains or by deploying additional new domains.

For more information about upgrading Windows NT domains, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book. For more information about deploying additional regional domains, see "Deploying Windows Server 2003 Regional Domains" in this book. For more information about restructuring Windows NT domains, see "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.

# Process for Deploying the Forest Root Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying the forest root domain involves completing the tasks that are shown in Figure 6.1. Only organizations with an existing DNS infrastructure or organizations whose DNS services are provided by an Internet service provider will configure DNS for the forest root domain. Other organizations will skip the task, allowing the Active Directory Installation Wizard to automatically create an internal root zone, which acts as the authoritative root for the organization.

**Figure 6.1 Deploying the Forest Root Domain**



# Background Information for Deploying the Forest Root Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy the forest root domain, understand the importance of maintaining the availability of forest root domain controllers. Also become familiar with how you can save time during the deployment process by automating installations and by using the Active Directory Installation Wizard.

**Forest Root Domain Controller Availability**
The forest root domain has different characteristics than other domains in the forest. It is the center of Kerberos authentication referrals and the domain controllers in the forest root domain host the forest-level operations master roles, in addition to domain-level operations master roles.

Because it has a unique role in the forest, special considerations must be made to ensure forest root domain controller availability.

- Because forest root domain controllers are used for Kerberos authentication referrals, ensure that at least one forest root domain controller is always online.
- Changes to the schema require the schema master to be online.
- Changes to the domain infrastructure, including adding, removing, or renaming domains or application directory partitions, require the domain naming master to be online.

Maintain current backups of forest root domain controllers. Although corrupted data does not replicate out to other domains, maintaining current backups allows you to quickly restore a domain controller to its original state if data on the domain controller ever becomes corrupted.

For information about planning forest root domain controller placement to ensure availability, see "Designing the Site Topology" in this book.

**Automated Installations**
You can perform an unattended installation of Active Directory by supplying an answer file when you run the Active Directory Installation Wizard. The answer file options for installing Active Directory are in the DCInstall section of the Unattend.txt file. You can find instructions for creating an answer file for an Active Directory installation in the Deploy.cab file in the Support\Tools folder on the Windows Server 2003 product CD. Inside the Deploy.cab file, open Ref.chm to access the Unattend.txt file. In the left pane, expand **Unattend.txt**, and click **DCInstall**. Use Windows Explorer to extract this document. You can rename the Unattend.txt file and save it in any location. The Unattend.txt file also contains sample answer files.

You can also automate the installation of Windows Server 2003 by using Sysprep.exe or an unattended installation. For more information about automating installations of Windows Server 2003, see *Automating and Customizing Installations* in this kit.

**The Active Directory Installation Wizard**
Windows Server 2003 includes improvements to the Active Directory Installation Wizard. When you install the first domain controller in a domain, you can allow the wizard to automatically install and configure Active Directory–integrated DNS. Even if you need to manually configure some settings later, the wizard saves time and prevents errors during the initial configuration.

# Reviewing the Active Directory Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Begin your forest root domain deployment by reviewing Active Directory design information, as shown in Figure 6.2.

**Figure 6.2 Reviewing the Active Directory Design**



# Review the Active Directory Logical Structure Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review the Active Directory logical structure design that your design team completed, including the DNS infrastructure that is planned for supporting Active Directory. If your organization has an existing DNS infrastructure, review current network diagrams and DNS domain hierarchy diagrams. Also review the existing DNS zone configuration, replication, and resource records that are used for delegation and forwarding.

Document the information that you will need to install Windows Server 2003 and to configure DNS on each domain controller in the forest root domain. For a worksheet to assist you in documenting this information, see "Domain Controller Configuration" (DSSDFR_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Domain Controller DNS Configuration" on the Web at http://www.microsoft.com/reskit).

**Example: Reviewing the Active Directory Logical Structure Design for Trey Research**
The Active Directory logical structure design for Trey Research requires the deployment team to create a new dedicated forest root domain, trccorp.treyresearch.net. They will then create two regional domains:

- A new regional domain, west.trccorp.treyresearch.net.
- A regional domain, east.trccorp.treyresearch.net**,** created by upgrading an existing Windows NT 4.0 domain, and then restructuring several other existing Windows NT 4.0 domains into it.

The deployment team reviews the existing DNS infrastructure for the Trey Research business unit. Their existing DNS infrastructure provides name resolution for all internal resources, including:

- Any servers, such as Web or mail servers, residing in the perimeter network (also known as the DMZ) and are accessed by Internet users.
- Any computers, or other network devices, that reside in the private network and run a non-Windows operating system, such as UNIX or Macintosh operating systems.

**Note**

- Windows NT 4.0–based computers in the private network use Windows Internet Name Service (WINS) to provide name resolution.

Trey Research's registered DNS domain name is treyresearch.net. This DNS domain name:

- Provides DNS naming for computers that are accessed by Internet users.
- Represents the *external* DNS namespace for the business unit.
- Runs on the Berkeley Internet Name Domain (BIND) DNS servers (SEA-TREY-DNS-01 and SEA-TREY-DNS-02) that are placed in the perimeter network.

Figure 6.3 shows an example of a completed domain controller configuration worksheet, showing the TCP/IP client settings for the domain controllers planned for the Trey Research forest root domain. For each domain controller, the preferred DNS server is the local domain controller, and the alternate is the closest DNS server. Initially, however, Trey Research configures the first domain controller in the domain to use a DNS server in its parent DNS domain as its preferred DNS server. During installation of Active Directory on the first domain controller, the preferred DNS server is changed to the local domain controller. Then, after the second domain controller is online, Trey Research reconfigures the first domain controller to use the second domain controller as its alternate DNS server.

**Figure 6.3 Example of a Domain Controller Configuration Worksheet**

| Domain Controller Configuration | | | |
|---|---|---|---|
| Prepared By | Active Directory Deployment Team | Date | 01/01/03 |

| | |
|---|---|
| Forest Name | trccorp.treyresearch.net |
| Forest Owner | forest owner |
| Domain Name | trccorp.treyresearch.net |

| Domain Controller | IP Address | Mask | Preferred DNS Server | Alternate DNS Server |
|---|---|---|---|---|
| SEA-TRC-DC01 | 172.16.16.2 | 255.255.252.0 | IP - SEA-TREY-DNS-01 (temp)<br>172.16.16.3 (final) | (none)(temp)<br>172.16.16.3 (final) |
| SEA-TRC-DC02 | 172.16.16.3 | 255.255.252.0 | 172.16.16.3 | 172.16.16.2 |
| SEA-TRC-DC03 | 172.16.16.4 | 255.255.252.0 | 172.16.16.4 | 172.16.16.3 |
| NYC-TRC-DC01 | 172.16.28.8 | 255.255.252.0 | 172.16.28.8 | 172.16.16.2 |

# Review Site Topology Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review your site topology design information, including:

- Site and subnet design.
- Site link design.

The site topology design team can provide worksheets that document the site topology information that you will need to configure the site topology for the forest.

Figure 6.4 shows an example of a completed Associating Subnets with Sites worksheet for Trey Research. The worksheet lists the sites that will be created for Trey Research and the corresponding locations and subnets to be included in each site. The Trey Research site topology design specifies that the Phoenix subnets are associated with the Seattle site.

**Figure 6.4 Example of an Associating Subnets with Sites Worksheet**

| Associating Subnets with Sites | | | |
|---|---|---|---|
| Prepared By | Site Topology Owner | Date | 1/1/03 |
| Forest Name | Trccorp.treyresearch.net | | |

| Site Candidate | Location | Active Directory Subnets |
|---|---|---|
| Seattle | Seattle<br><br>Phoenix | 172.16.16.0/22<br>172.16.24.0/22<br><br>172.16.64.0/22 |
| Boston | Boston | 172.16.36.0/22<br>172.16.40.0/22 |
| New York | New York | 172.16.28.0/22<br>172.16.32.0/22 |
| Los Angeles | Los Angeles | 172.16.136.0/22<br>172.16.92.0/22 |
| Washington DC | Washington DC | 172.16.116.0/22<br>172.16.140.0/22 |

Figure 6.5 shows an example of a completed Sites and Associated Site Links worksheet for Trey Research, including the communication links between locations.

**Figure 6.5 Example of a Sites and Associated Site Links Worksheet**

| Sites and Associated Site Links | | | |
|---|---|---|---|
| Prepared By | Site Topology Owner | Date | 1/1/03 |
| Forest Name | Trccorp.treyresearch.net | | |

| Site Names | Linked Sites | Name of the Site Link | Link Type |
|---|---|---|---|
| Seattle | LosAngeles | LAX-SEA | ISDN (128.8 Kbps) |
| Seattle | NewYork | SEA-NYC | T1 (1.544 Mbps) |
| NewYork | WashingtonDC | NYC-WDC | T1 (1.544 Mbps) |
| NewYork | Boston | NYC-BOS | ISDN (128.8 Kbps) |

# Review Hardware Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Ensure that each computer that you plan to use as a domain controller meets the hardware requirements for running Windows Server 2003. For more information about assessing the hardware requirements of domain controllers in a Windows Server 2003 domain, see "Planning Domain Controller Capacity" in this book.

# Configuring DNS for the Forest Root Domain

Updated: October 19, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure Domain Name System (DNS) for the forest root domain, the DNS administrator of your organization delegates the zone that matches the name of the forest root domain to the DNS servers (domain controllers) that you plan to install in the forest root domain.

**Important**

When no DNS infrastructure exists, you do not have to configure the forest root domain. To complete the next step in the deployment process, see Creating the Forest Root Domain. The remainder of this document describes the process of configuring and delegating a zone in the existing DNS internal namespace.

If you plan to implement an Active Directory domain namespace that does not match the DNS namespace, first read Disjoint Namespace, and then follow the steps in Create a Disjoint Namespace. If a disjoint namespace already exists, and you want to use a contiguous namespace, see Transition from a Disjoint Namespace to a Contiguous Namespace.

The following illustration shows when you configure DNS for the forest root domain in the forest root domain deployment process.




In preparation for the forest root domain deployment, you create a delegation for the DNS servers that will run on the domain controllers in the forest root domain. To create the delegation, you add DNS name server (NS) and host (A) resource records to the parent DNS zone.

**Note**

The delegation that occurs references the first forest root domain controller, which does not currently exist. You install and configure the DNS service on the first forest root domain controller in a subsequent step.

**To delegate the DNS zone for the forest root domain**

1. Create a name server (NS) resource record in the parent zone.

   Use the left-most portion of the forest root domain name, and the full DNS name of the domain controller, as shown in the following syntax:

   forest_root_domain IN NS domain_controller_name

2. Create a host (A) resource record in the parent zone.

   Use the full DNS name of the domain controller, as shown in the following syntax:

   domain_controller_name IN A domain_controller_ip_address

   For example, the DNS administrator for Trey Research created the following DNS resource records in the parent zone, treyresearch.net:

- trccorp in the name server (NS) SEA-TRC-DC01.trccorp.treyresearch.net
- SEA-TRC-DC01.trccorp.treyresearch.net in the 172.16.16.2 domain

# Creating the Forest Root Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in creating the forest root domain is deploying the first forest root domain controller. The forest owner is responsible for deploying the forest root domain.

Figure 6.7 illustrates the process for creating the forest root domain.

**Figure 6.7 Creating the Forest Root Domain**



# Deploy the First Forest Root Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To deploy the first domain controller in the forest root domain, complete the following tasks:

- Install Windows Server 2003
- Install Active Directory
- Verify the Active Directory installation
- Configure the Windows Time Service
- Verify DNS server recursive name resolution

# Install Windows Server 2003 on the First Forest Root Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in deploying the first forest root domain controller is to install Windows Server 2003 on the computer that you want to make the domain controller.

**Note**

- Before installing Windows Server 2003, ensure that DNS was never previously installed on the computer. If DNS was previously installed, configuration of the DNS resolver and forwarders might fail.

  - Insert the Windows Server 2003 operating system CD in the CD-ROM drive of the domain controller and select the option to install the operating system, or use an automated installation method. If the Windows Server 2003 media is shared on the network, run the Winnt32.exe command. Use the NTFS file system to format the partitions. Enter the computer name, static IP address, and subnet mask as specified by your design.

  - In TCP/IP Properties, configure the DNS client settings by using the information documented in the "Domain Controller Configuration" worksheet. The DNS settings are temporary and will be changed after you install Active Directory.

  - Enable Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode) to enable administrators to log on remotely if necessary. To enable Remote Desktop for Administration, in Control Panel, double-click **System**, select the **Remote** tab, and then select **Allow users to connect remotely to this computer**.

  - Install Windows Support Tools, which are available in the \Support\Tools folder on the Windows Server 2003 operating system CD.

# Install Active Directory on the First Forest Root Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Install Active Directory by running the Active Directory Installation Wizard on the computer that you want to make the first forest root domain controller. The Active Directory Installation Wizard creates the Active Directory database and initializes the directory data in the database.

In addition, on the first domain controller in a domain, the wizard also:

- Prompts the administrator to verify the installation and configuration of the DNS Server service.
- Configures DNS recursive name resolution by forwarding, by adding the IP addresses of the existing entries for **Preferred DNS server** and **Alternate DNS server** to the list of DNS servers on the **Forwarders** tab of the **Properties** sheet in the DNS snap-in for the domain controller.

  **Notes**

  - If you want to set different forwarders, or do not want to enable forwarding, you can change this setting manually by using the DNS snap-in.
  - If your domain controller is multihomed, forwarding is not configured automatically.
  - For manual configuration instructions, see "Verify DNS Server Recursive Name Resolution on the First Forest Root Domain Controller" later in this chapter.
- Configures DNS recursive name resolution by root hints, by adding the root hints that are configured on the **Preferred DNS server**
- Configures the **Preferred DNS server** to point to the DNS server that is running locally on the domain controller, and configures the **Alternate DNS server** to point to the DNS server that is connected through the minimum number of network segments.
- Creates two application directory partitions that are used by DNS. The DomainDnsZones application directory partition holds domain-wide DNS data, and the ForestDnsZones application directory partition holds forest-wide DNS data.

**To install Active Directory on the first forest root domain controller**

1. Log on to the Windows Server 2003–based member server.
2. At the command line, type:

   **dcpromo**

   – Or –

   Open **Administrative Tools** and click **Configure Your Server Wizard**. Select **Domain Controller (Active Directory)** to configure your domain controller. After the Configure Your Server Wizard finishes, the Active Directory Installation Wizard begins.

   Use Table 6.1 to complete the Active Directory Installation Wizard. Table 6.1 includes the specific actions taken by Trey Research as they deploy their first forest root domain controller, SEA-TRC-DC01.

   **Table 6.1 Information to Install Active Directory on the First Forest Root Domain Controller**

| Wizard Page or Dialog Box | Action | Example |
|---|---|---|
| Domain Controller Type | Select **Domain controller for a new domain** | |
| Create New Domain | Select **Domain in a new forest** | |
| New Domain Name | Type the full DNS name of the domain. | trccorp.treyresearch.net |
| NetBIOS Domain Name | Confirm or type the NetBIOS name. | TRCCORP |
| Database and Log Folders | Type the folder locations specified by your design. | Database folder: C:\WINNT\NTDS<br><br>Log folder: D:\Logs |
| Shared System Volume | Confirm or type the location specified by your design. | C:\WINNT\SYSVOL |
| DNS Registration Diagnostics | For organizations with an existing DNS infrastructure, a message will indicate that the wizard cannot contact the DNS server with which this DC will be registered. This is because the pre-created delegation record points to the local computer and DNS has not been installed on the domain controller at this point.<br><br>Select **Install and configure the DNS server on this computer and set this computer to use this DNS server as its preferred DNS server**. | Before running the Active Directory Installation Wizard, the Trey Research deployment team set the Preferred DNS server to 172.16.24.4, which is the IP address of a DNS server in the parent zone, treyresearch.net. This address will be automatically moved to the list of forwarders, and Preferred DNS server will be set to the local host. |
| Permissions | Select the security level specified by your design:<br>• Permissions compatible with pre-Windows 2000 server operating systems<br>• Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems | Because Trey Research currently has server programs running on Windows NT 4.0 servers, they selected **Permissions compatible with pre-Windows 2000 server operating systems**. |
| Directory Service Restore Mode Administration Password | In the **Password** and **Confirm password** boxes, type any strong password. | |

For more information about installing and removing Active Directory, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

# Verify the Active Directory Installation on the First Forest Root Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To verify the Active Directory installation on the first forest root domain controller:

1. Review the Windows Server 2003 event log and resolve any errors.

2. At the command line, run Dcdiag.exe and Netdiag.exe and resolve any errors that are reported.

   For more information about tests you can perform by using Dcdiag and Netdiag, see the Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search under "Administration and Configuration Guides" and download the *Active Directory Operations Guide*.

3. Run Task Manager and verify that the processor and memory system resources are within acceptable limits.

4. Open the DNS snap-in, navigate to **Forward Lookup Zones**, and verify that the zones _msdcs.*forest_root_domain_name* and *forest_root_domain_name* were created. Expand the *forest_root_domain_name* node and verify that **DomainDnsZones** and **ForestDnsZones** were created.

# Configure the Windows Time Service

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When deploying the forest root domain, it is important to correctly configure the Windows Time Service to meet your organization's needs. The Windows Time Service provides time synchronization to peers and clients, ensuring consistent time throughout an enterprise.

By default, the first domain controller deployed holds the primary domain controller (PDC) emulator operations master role. Set the PDC emulator to synchronize from a valid Network Time Protocol (NTP) source. If no source is configured, the service will log a message to the event log, and use the local clock when providing time to clients.

Follow these best practices for configuring the time source on the forest-root PDC emulator, in this order of preference:

- Install a hardware clock, such as a radio or GPS device, as the source for the PDC. There are many consumer and enterprise devices that use the Network Time Protocol (NTP), allowing you to install the device on an internal network for usage with the PDC.

- Configure the Windows Time service to synchronize with an external time server. External time servers allow users to synchronize computer clocks by means of dial-up, network, and radio links.

  The Microsoft time server (time.windows.com) uses NIST, the National Institute of Standards and Technology, located in Boulder, Colorado, as its external time provider. NIST provides the Automated Computer Time Service (ACTS), which can set a computer clock with an uncertainty of less than 10 milliseconds. The U.S. Naval Observatory (USNO) Time Service Department in Washington D.C. is another source for accurate time synchronization in the United States. Many other sites exist throughout the world that can be used for time synchronization. To find them, search for "time synchronization" on the Internet.

| **Note** |
| --- |
| Because synchronization with an external time source is not authenticated, it is less secure. |

Repeat this operation if you transfer or seize the PDC emulator operations master role to another domain controller in the forest root domain.

**To configure the Windows Time Service on first forest root domain controller**

1. Log on to the domain controller.

2. Type the following command to display the time difference between the local computer and a target computer, and then press ENTER:

   **w32tm /stripchart /computer:***target***/samples:***1***/dataonly**

   Where *target* specifies the DNS name or IP address of the NTP server that you are comparing the local computer's time against, such as time.windows.com and *1* specifies the number of time samples that will be returned from the target computer. In this example, only one sample will be returned to test basic NTP communication.

3. Open UDP port 123 for outgoing traffic if needed.

4. Open UDP port 123 (or a different port you have selected) for incoming NTP traffic.

5. Type the following command to configure the PDC emulator and then press ENTER:

   **w32tm /config /manualpeerlist:***peers* **/syncfromflags:manual /reliable:yes /update**

   *Peers* specifies the list of DNS names and/or IP addresses of the NTP time source that the PDC emulator synchronizes

from. For example, you can specify time.windows.com. When specifying multiple peers, use a space as the delimiter and enclose them in quotation marks. For example, if you wanted to configure time synchronization with a server named Ntp in the contoso.com domain and a server named Time in the fabrikam.com domain, you would enter the following command: **w32tm /config /update /manualpeerlist:"Ntp.contoso.com Time.fabrikam.com time.windows.com" /syncfromflags:manual /update**.

**Note**

- When specifying a manual peer, do not use the DNS name or IP address of a computer that uses the forest root domain controller as its source for time, such as another domain controller in the forest. The time service will not operate correctly if there are cycles in the time source configuration.

For more information about configuring and deploying the Windows Time Service, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**See Also**
**Other Resources**
Videos about the Windows Time Service

# Verify DNS Server Recursive Name Resolution on the First Forest Root Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DNS server recursive name resolution is configured automatically during the Active Directory installation process, as described in "Install Active Directory on the Second Domain Controller" earlier in this chapter. If your design specifies a different configuration, you can use the DNS snap-in or Dnscmd.exe to modify these settings.

**To verify DNS server recursive name resolution on the first forest root domain controller**

- Use the DNS snap-in to verify DNS server recursive name resolution for the method used in your organization based on the information in Table 6.2.

**Table 6.2 Information to Verify DNS Server Recursive Name Resolution**

| Method | Configuration |
|---|---|
| Recursive name resolution by root hints | Root hints are the recommended method to use for recursive name resolution in a Windows Server 2003 environment.<br><br>No additional configuration is necessary. When the DNS server specified as the Preferred DNS server during the installation process is properly configured, the root hints are automatically configured. To verify the root hints by using the DNS snap-in:<br><br>In the console tree, right-click the domain controller name, and then click **Properties**.<br><br>In the **Properties** sheet for the domain controller, view the root hints on the **Root Hints** tab. |
| Recursive name resolution by forwarding | Only use Forwarders if that is what your organization's design specifies. Root hints are the recommended method to use for recursive name resolution in a Windows Server 2003 environment.<br><br>Forward unresolved queries to specified DNS servers. To verify forwarding by using the DNS snap-in:<br><br>In the console tree, right-click the domain controller name, and then click **Properties**.<br><br>On the **Forwarders** tab, in the selected domain's **Forwarders** list, verify that the IP addresses match those specified by your design. |
| No existing DNS infrastructure | No additional configuration is necessary.<br><br>In this environment, if you want to configure internal DNS servers to resolve queries for external names, then configure this DNS server to forward unresolved queries to an external server, such as one in your perimeter network, or one hosted by an Internet service provider. |

# Deploy the Second Domain Controller in the Same Site

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you deploy the first forest root domain controller, deploy the second forest root domain controller in the same site, according to the design provided by your design team. To deploy the second forest root domain controller, complete the following tasks:

- Install Windows Server 2003
- Install Active Directory
- Install DNS Server service
- Verify the Active Directory installation

# Install Windows Server 2003 on the Second Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in deploying the second forest root domain controller is to install Windows Server 2003 on the computer that you want to make the second domain controller.

**Note**

- Before installing Windows Server 2003, ensure that DNS was not previously installed.

  - Insert the Windows Server 2003 operating system CD in the CD-ROM drive of the domain controller and select the option to install the operating system, or use an automated installation method. If the Windows Server 2003 media is shared on the network, run the Winnt32.exe command. Use NTFS to format the partitions. Enter the computer name, static IP address, and subnet mask as specified by your design.

  - Configure the DNS client settings by using the information documented in the "Domain Controller Configuration" worksheet (DSSDFR_1.doc).

  - Enable Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode) to enable administrators to log on remotely if necessary. To enable Remote Desktop for Administration, in Control Panel, double-click **System**, select the **Remote** tab, and then select **Allow users to connect remotely to this computer**.

  - Install Windows Support Tools, which are available in the \Support\Tools folder on the Windows Server 2003 operating system CD.

# Install Active Directory on the Second Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Install Active Directory on the computer that you want to make the second forest root domain controller by running the Active Directory Installation Wizard.

The Active Directory Installation Wizard:

- Creates the Active Directory database.
- Initializes the directory data in the database.

On domain controllers other than the first domain controller in a domain, installation of DNS is not automatic.

To deploy an additional domain controller in an existing domain, you can either let replication copy domain information from an existing source domain controller over the network or you can use the install from media feature, new in Windows Server 2003. Install from media allows you to pre-populate Active Directory with System State data backed up from an existing domain controller. This backup can be present on local CD, DVD, or hard disk partition. Installing from media drastically reduces the time required to install directory information by reducing the amount of data that is replicated over the network. Installing from media is most beneficial in environments with very large domains or for installing new domain controllers that are connected by a slow network link. To use the install from media feature, you first create a backup of System State from the existing domain controller, and then restore it to the new domain controller by using the **Restore to: Alternate location** option.

**To install Active Directory on the second domain controller**

1. Log on to the Windows Server 2003–based member server.
2. If you want to copy domain information from restored backup files, at the command line, type:

   **dcpromo /adv**

   – or –

   If you want to copy domain information over the network, either type **dcpromo** (without the /adv switch) or open **Administrative Tools**, click **Configure Your Server Wizard**, andselect **Domain Controller (Active Directory)** to configure your domain controller. After the Configure Your Server Wizard finishes, the Active Directory Installation Wizard begins.
3. Use Table 6.3 to help you complete the Active Directory Installation Wizard. Table 6.3 also includes the specific actions taken by Trey Research as they deployed their second forest root domain controller, SEA-TRC-DC02.

   **Table 6.3 Information to Install Active Directory on the Second Forest Root Domain Controller**

| Wizard Page or Dialog Box | Action | Example |
|---|---|---|
| Domain Controller Type | Select **Additional domain controller for an existing domain** | |

| | | |
|---|---|---|
| Copying Domain Information | (This dialog box appears only when you started the Active Directory Installation Wizard by typing dcpromo with the /adv switch at the command line or used the Configure Your Server Wizard.)<br><br>Select either:<br><br>● Over the network from a domain controller<br><br>● From these restored backup files | Trey Research is copying from the first TRCCORP domain controller, SEA-TRC-DC01,which is in the same location as the new one, so they selected **Over the network** to copy the information in the shortest time. |
| Global Catalog | (This dialog box appears only when **From these restored backup files** was selected, if the domain controller that you backed up was a global catalog server.)<br><br>Specify whether this domain controller should be configured as a global catalog server. | |
| Network Credentials | In the **User name** box, type a user account that has sufficient rights to add a domain controller, typically a member of Domain Admins.<br><br>In the **Password** box, type the password of the user account. | |
| Additional Domain Controller | (This dialog box appears only when **Over the network** was selected.)<br><br>Confirm or type the full DNS name of the forest root domain. | trccorp.treyresearch.net |
| Database and Log Folders | Type the folder locations specified by your design. | Database folder: C:\WINNT\NTDS<br><br>Log folder: D:\Logs |
| Shared System Volume | Confirm or type the location specified by your design. | C:\WINNT\SYSVOL |
| Directory Service Restore Mode Administration Password | In the **Password and Confirm password** boxes, type any strong password. | |

# Install DNS Server on the Second Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After Active Directory installation has finished and the computer has restarted, install DNS on the second Windows Server 2003–based domain controller that is added to the domain.

**To install DNS on additional domain controllers by using the Windows Components Wizard**

1. Click **Start**, point to **Settings**, and click **Control Panel**.
2. Double-click **Add or Remove Programs**, and then click **Add/Remove Windows Components**.
3. In **Components**, select the **Networking Services** check box, and then click **Details**.
4. In **Subcomponents of Networking Services**, select the **Domain Name System (DNS)** checkbox, click **OK**, and then click **Next**.
5. If prompted, in **Copy files from**, type the full path to the distribution files and then click **OK**. The required files will be copied to your hard disk.

# Verify the Active Directory Installation on the Second Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the same steps as shown in the procedure for the first domain controller, but instead of verifying that **DomainDnsZones** and **ForestDnsZones** were created, use the repadmin /showreps command to verify that the **ForestDnsZones** and **DomainDnsZones** application directory partitions are replicated successfully. Use the DNS snap-in to verify that DNS server recursive name resolution is configured according to the method used by your organization.

# Reconfigure the DNS Service

Updated: October 19, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Reconfigure the DNS service to account for the addition of the second domain controller in the forest root domain. You can also use these procedures as you deploy additional domain controllers that are running the DNS service. To reconfigure the DNS service:

- Enable Aging and Scavenging for DNS
- Configure the DNS Client Settings of the First and Subsequent Domain Controllers
- Update the DNS Delegation

# Enable Aging and Scavenging for DNS

Updated: October 19, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Enable aging and scavenging on two Windows Server 2003–based domain controllers running the DNS Server service per domain, to allow automatic cleanup and removal of stale resource records (RRs), which can accumulate in zone data over time.

With dynamic update, RRs are automatically added to zones when computers start on the network. However, in some cases, they are not automatically removed when computers leave the network. For example, if a computer registers its own host (A) RR at startup and is later improperly disconnected from the network, its host (A) RR might not be deleted. If your network has mobile users and computers, this situation can occur frequently.

If left unmanaged, the presence of stale RRs in zone data might cause problems including:

- If a large number of stale RRs remain in server zones, they can eventually take up server disk space and cause unnecessarily long zone transfers.
- DNS servers loading zones with stale RRs might use outdated information to answer client queries, potentially causing the clients to experience name resolution problems on the network.
- The accumulation of stale RRs at the DNS server can degrade its performance and responsiveness.

**Caution**

- By default, the aging and scavenging mechanism for the DNS Server service is disabled. Enable aging and scavenging only after you understand all parameters. Otherwise, the server could be accidentally configured to delete resource records that should not be deleted. If a resource record is accidentally deleted, not only will users fail to resolve queries for that resource record, but any user can create the resource record and take ownership of it, even on zones configured for secure dynamic update.
- For more information about how to configure aging and scavenging, see "Understanding aging and scavenging" in Help and Support Center for Windows Server 2003.

To enable the aging and scavenging features, perform the following steps to configure the applicable server and its Active Directory–integrated zones:

- Enable aging and scavenging at the server. These settings determine the effect of zone-level properties for any Active Directory–integrated zones loaded at the server.
- Enable aging and scavenging for selected zones at the DNS server. When zone-specific properties are set for a selected zone, these settings apply only to the applicable zone and its resource records. Unless these zone-level properties are otherwise configured, they inherit their default settings from comparable settings maintained in the DNS server aging/scavenging properties.

**To set aging and scavenging properties for the DNS server**

1. Log on to the computer that is running the DNS Server service with an account that is a member of the local Administrators group.
2. In the DNS console tree, right-click the applicable DNS server, and then click **Set Aging/Scavenging for all zones**.
3. Select the **Scavenge stale resource records** check box.
4. Modify other aging and scavenging properties as needed.

**To set aging and scavenging properties for a zone**

1. Log on to the computer that is running the DNS Server service with an account that is a member of the local Administrators group.

2. In the DNS console tree, right-click the applicable zone, then click **Properties**.

3. On the **General** tab, click **Aging**, and then select the **Scavenge stale resource records** check box.

4. Modify other aging and scavenging properties as needed.

# Configure the Domain Name System (DNS) Client Settings of the First and Subsequent Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have deployed an additional domain controller, modify the DNS client settings on the first forest root domain controller. Because no other domain controllers were running when you deployed the first forest root domain controller, modify the DNS client settings on the first forest root domain controller to include the additional domain controller. As you deploy more domain controllers, you might also need to modify the Alternate DNS server setting specified on existing domain controllers to ensure that this setting points to the DNS server that is connected through the minimum number of network segments.

**To configure the DNS client settings on previously installed domain controllers**

1. Open **Network Connections**, double-click your **Local Area Connection**, click **Properties**, click **Internet Protocol (TCP/IP)** to highlight it, and then click **Properties**.

2. For the **Preferred DNS server**, type the IP address of the DNS server that is running locally on the domain controller (local host).

3. Determine whether a newly installed domain controller is now closer to this domain controller than the domain controller that you originally specified as the **Alternate DNS server**. If it is, for **Alternate DNS server**, type the IP address of the newly installed domain controller.

# Update the DNS Delegation

Updated: October 19, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you install the DNS Server service on new domain controllers, update the DNS delegation for the forest root domain on a DNS server that is authoritative for the parent zone.

**To update the DNS delegation records for the additional domain controller**

1. Create a name server (NS) resource record in the parent zone. Use the left-most portion of the forest root domain name, and the full DNS name of the additional domain controller.

   **forest_root_domain** IN NS **additional_domain_controller_name**

2. Create a host address (A) resource record in the parent zone. Use the full DNS name of the domain controller.

   **additional_domain_controller_name** IN A **additional_domain_controller_ip_address**

   For example, the DNS administrator for Trey Research created the following DNS resource records in the parent zone, treyresearch.net:

   - **trccorp** IN NS **SEA-TRC-DC02.trccorp.treyresearch.net**
   - **SEA-TRC-DC02.trccorp.treyresearch.net** IN A **172.16.16.3**

# Transition Namespaces

Updated: May 3, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you decide to modify your network name resolution infrastructure after you read Disjoint Namespace, you must perform several steps. Whether you want to create a disjoint namespace or transition from a disjoint namespace to a contiguous namespace, you must configure the following components appropriately:

- Primary Domain Name System (DNS) suffixes
- Service principal names (SPNs)
- Group Policy objects (GPOs)
- Permissions to Active Directory objects

The primary differences between implementing a disjoint namespace and transitioning to a contiguous namespace are as follows:

- When you transition to a contiguous namespace from a disjoint namespace, you must ensure that the Active Directory domain name and the DNS namespace are identical.
- When you create a disjoint namespace, you must create different Active Directory and DNS namespaces.

If you create a disjoint namespace accidentally, you can reverse the erroneous steps. The process for doing so is similar to transitioning a namespace, but not as complex.

For more information about working with disjoint namespaces successfully, see the following documents:

- Create a Disjoint Namespace
- Reverse an Accidentally Created Disjoint Namespace
- Transition from a Disjoint Namespace to a Contiguous Namespace

**See Also**
**Other Resources**
Disjoint Namespace

# Transition from a Disjoint Namespace to a Contiguous Namespace

Updated: October 19, 2007

Applies To: Windows Server 2003, Windows Server 2003 with SP2

If your network uses a disjoint namespace and you decide to change to a contiguous namespace, there are several tasks to complete:

- Configure member computers to use a primary Domain Name System (DNS) suffix that is the same as the Active Directory domain name.
- Remove any Group Policy objects (GPOs) that specify disjoint namespaces.
- Clear DNS suffixes from the client search lists that no longer apply.
- Remove no longer needed permissions for Active Directory.

The actual number of tasks that are required for the change depends on how much configuration was accomplished to make the disjoint namespace work. For example, if custom DNS suffix search lists were never created, there is obviously no need to remove them.

| **Note** |
| --- |
| If any of your servers have service principal names (SPNs) that were manually configured, you may have to modify them. For instructions on how to manually configure SPNs, see Setspn Overview. |

**Configure domain member computers to use the same DNS suffix as the Active Directory domain**

You can configure the DNS client's primary DNS suffix by using Group Policy, manually through **System Properties** on each computer, or through a registry modification. These options are described in the following sections.

**Configure primary DNS suffixes by using a GPO**

On Windows Server 2003 servers, use a GPO to configure the primary DNS suffix for each member computer. Set the DNS Client policy for **Primary DNS Suffix** to the desired primary DNS suffix to match DNS suffix of the Active Directory domain name. The **Primary DNS Suffix** setting is located in the following GPO path: **Computer Configuration\Administrative Templates\Network\DNS Client**.

| **Note** |
| --- |
| For more information about managing Group Policy, see Group Policy Object Editor Tools and Settings (http://go.microsoft.com/fwlink/?LinkId=102368) and Enterprise Management with the Group Policy Management Console (http://go.microsoft.com/fwlink/?LinkId=29909) |

**Configure the primary DNS suffix by using System Properties**

To change from a disjoint namespace that is created on a local computer running Microsoft Windows 2000, Windows XP, or Windows Vista, complete the following procedure.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To configure the primary DNS suffix by using System Properties**

1. Open the **System Properties** dialog box:
    a. To open the **System Properties** dialog box in Windows 2000 or Windows XP, click **Start**, click **Run**, type **sysdm.cpl**, and then press ENTER.

b. To open the **System Properties** dialog box in Windows Vista, click **Start**, in **Start Search** type **sysdm.cpl**, and then press ENTER.

2. On the **Computer Name** tab, click **Change**, and then click **More**.
3. Verify that the value in **Primary DNS suffix of this computer** is the same as the Active Directory domain name.
4. Verify that the **Change primary DNS suffix when domain membership changes** check box is selected, and then click **OK** twice.
5. Click **OK** to confirm that the computer must be restarted, and then click **Close**.
6. When you are prompted to restart your computer, click **Restart Now**.

   If you choose to restart your computer later, any changes that you make will take effect at that time.

**Change the primary DNS Suffix in the Registry**

If you use a script (or directly modify the registry) to configure the Primary DNS suffix, there are two values to modify: **Domain** and **SyncDomainWithMembership**. Both of these values are in the **Parameters** key under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip**. To create a contiguous namespace, enable **SyncDomainWithMembership** by setting its value to **1**, and then enter the primary DNS suffix that is identical to the Active Directory domain for the **Domain** value. If the **SyncDomainWithMembership** value does not exist, you can create it as a **REGDWORD**. For more information about registry scripts, see article 264584 in the Microsoft Knowledge base (http://go.microsoft.com/fwlink/?LinkId=102370). If you configure these values in the registry, be sure to restart the computers to have the settings take effect.

**Remove unnecessary entries in the DNS suffix search list**

When you change from a disjoint namespace to a contiguous namespace, there are likely to be several entries in the DNS suffix search list that clients no longer need. There may also be some different values that should be added to the DNS suffix search list. The following sections describe the three main methods for modifying the DNS suffix search list, which use:

- Group Policy
- Network Connections
- The registry

**Changing the DNS suffix search list with a GPO**

On Windows Server 2003 servers, remove any GPOs or settings that were used to create custom DNS suffix search lists that applied to the disjoint namespace. Remove the **DNS Suffix Search List** or configure it with the correct DNS suffixes for the contiguous namespace. The GPO path for the **DNS Suffix Search List** policy setting is Computer Configuration\Administrative Templates\Network\DNS Client.

**Changing the DNS suffix search list by using Network Connections**

To change the DNS suffix search list by using the graphical user interface (GUI) on Windows 2000, Windows XP, or Windows Vista, complete the following procedure.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To change the DNS suffix search list by using Network Connections**

1. Open Network Connections.

   a. To open the **Network Connections** dialog box in Windows 2000 or Windows XP, click **Start**, click **Run**, type **ncpa.cpl**, and then press ENTER.

   b. To open the **Network Connections** dialog box in Windows Vista, click **Start**, type **ncpa.cpl** in **Start Search**, and then press ENTER.

2. Right-click the icon that represents the computer's local network connection (**Local Area Connection**, by default), and then click **Properties**.

3. In **This connection uses the following items**, select **Internet Protocol (TCP/IP)**, and then click **Properties**.

4. Click **Advanced**.

5. On the **DNS** tab, click **Append these DNS suffixes (in order)**,and then click **Add**. Type the first DNS suffix that you want the client to use when it searches for single-label names (names that do not have DNS suffixes), and then click **Add**. Repeat this step for each successive DNS suffix that you want clients to use when they try to resolve a name.

   | **Note** |
   | --- |
   | As an alternative, you can select **Append primary and connection specific DNS suffixes** and, as an option, also select **Append parent suffixes of the primary DNS suffix**, assuming that there is no longer a need for a custom list that includes other domains. |

6. Click **OK** twice to confirm your changes.

7. Close the **Local Area Connection Properties** and the **Network Connections** dialog boxes.

**Changing the DNS suffix search list in the registry**

If you want to use a script (or directly modify the registry) to configure the DNS suffix search list, you can modify the **SearchList** value that is located in the **Parameters** key under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip**. For more information about registry scripts, see article 264584 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=102370) .

**Remove no-longer-needed DNS suffixes from Active Directory**

To remove permissions that may have been added to grant client computers permission to modify their primary DNS suffixes in Active Directory, complete the following procedure.

| **Note** |
| --- |
| The ADSI Edit tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270). |

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To remove no-longer-needed DNS suffixes from Active Directory**

1. On any domain controller in the Active Directory domain, click **Start**, click **Run**, type **adsiedit.msc**, and then press ENTER.

   | **Note** |
   | --- |
   | If you recently installed the Support Tools to the default location, you may have to type **C:\Program Files\Support Tools\ADSIEdit.msc**, and then press ENTER. If you did not install the Support Tools to the default location, use the file path to which the tools were installed. |

2. Expand the **Domain Naming Context (NC)** to expose the directory partition that relates to the Active Directory domain for which you want to allow a disjoint namespace.

3. Right-click the Active Directory domain partition, and then click **Properties**. For example, DC=fabrikam,DC=com is the domain partition for Fabrikam.com.

4. In **Attributes**, locate and click the **msDS-AllowedDNSSuffixes** attribute, and then click **Edit**.

5. If there are any DNS suffixes entered in **Values**, select each one separately, and then click **Remove**.

6. When all the unnecessary DNS suffixes have been removed, click **OK** on the two open dialog boxes to confirm your changes.

7. Close ADSI Edit.

**Remove no-longer-needed permissions for Windows 2000 computers**

If Active Directory was modified to allow member computers or domain controllers running Windows 2000 to update the **dNSHostName** and **servicePrincipleName** attributes of their Active Directory computer accounts, you can revert those changes by completing the following procedure.

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To remove no-longer-needed permissions for Windows 2000 computers**

1. Open Active Directory Users and Computers. To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then press ENTER.
2. Click **View**, and then ensure that the **Advanced Features** check box is selected. If it is not selected, click **Advanced Features**. If the domain to which you want to allow a disjoint namespace does not appear in the console, do the following:
    a. In the console tree, right-click **Active Directory Users and Computers**, and then click **Connect to Domain**.
    b. In the **Domain** box, type the name of the Active Directory domain to which you want to allow the disjoint namespace (or use the **Browse** button to locate it), and then click **OK**.
3. In the console tree, right-click the object that represents the domain to which you want to allow a disjoint namespace, and then click **Properties**.
4. Right-click the domain, and then click **Properties**.
5. On the **Security** tab, click **Advanced**.
6. On the **Permissions** tab, click **Add**.
7. In the **Enter the object name to select** box, type **self**, and then click **OK**.
8. Configure the **Apply onto** box for **Computer objects**.
9. At the bottom of the **Permissions** box, ensure that the **Allow** check boxes that correspond to the **Validated write to DNS host name** and **Validated write to service principal name** permissions are cleared, and then click **OK** on the three open dialog boxes to confirm your changes.
10. Close Active Directory Users and Computers.

# Reverse an Accidentally Created Disjoint Namespace

Updated: October 19, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An administrator can configure a computer to be in a disjoint namespace accidentally in the either of the following ways:

- The local administrator can clear the **Change primary DNS suffix** policy setting when domain membership changes.
- On servers running Windows Server 2003, a domain or enterprise administrator can add a Domain Name System (DNS) domain name to the **Primary DNS Suffix** policy setting on the DNS client. This policy setting overrides the **Change primary DNS suffix when domain membership changes** policy setting on DNS clients.

The following sections describe how to reverse a disjoint namespace that an administrator accidentally created.

**Reverse a locally configured disjoint namespace**
Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To reverse a disjoint namespace that an administrator created on a local computer running the Microsoft Windows 2000, Windows XP, or Windows Vista operating systems, perform the following steps:

**Reverse a locally configured disjoint namespace**

1. Open the **System Properties** dialog box.
   a. To open the **System Properties** dialog box in Windows 2000 or Windows XP, click **Start**, click **Run**, type **sysdm.cpl**, and then press ENTER.
   b. To open System Properties using Windows Vista, click **Start**, in **Start Search** type **sysdm.cpl**, and then press ENTER.
2. On the **Computer Name** tab, click **Change**, and then click **More**.
3. Verify that the value in **Primary DNS suffix of this computer** is the same as the Active Directory domain name.
4. Verify that the **Change primary DNS suffix when domain membership changes** check box is selected, and then click **OK** twice.
5. Click **OK** to confirm that the computer must be restarted, and then click **Close**.
6. When you are prompted to restart your computer, click **Restart Now**.

   If you choose to restart your computer later, any changes that you made will take effect at that time.

**Reverse a disjoint namespace that was created using Group Policy**
Domain controllers can use Group Policy to deploy a primary DNS suffix to computers that are members of the domain. The Group Policy object (GPO) path for the **Primary DNS Suffix** policy setting is Computer Configuration\Administrative Templates\System\DNS Client.

To reverse a disjoint namespace that an administrator created using Group Policy, remove or revise the GPO that implements the **Primary DNS Suffix** policy setting. To determine which policy settings are implemented, you can use the Group Policy Results tool (Gpresult.exe) from the DNS client. For more information about Gpresult, see article 321709 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=102289). You can use the Group Policy Management Console (GPMC) to determine which GPO implements the policy. For more information, see GPMC How To (http://go.microsoft.com/fwlink/?LinkID=102282).

**Note**

The **Primary DNS Suffix** policy setting does not apply to Windows 2000-based domain controllers. If you are running only Windows 2000-based domain controllers, you must manually modify the namespace of those computers. For more information, see Primary DNS Suffix (http://go.microsoft.com/fwlink/?LinkId=102283).

# Create a Disjoint Namespace

Updated: October 19, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you decide to create a disjoint namespace on your network, first read about the advantages and disadvantages of doing so that are discussed in Disjoint Namespace. To create the disjoint namespace, you perform the following tasks:

1. Create appropriate DNS forward lookup zones
2. Configure different primary DNS suffixes
3. Enable computers to register disjoint names
4. Update the DNS suffix search ist

Although you do not have to complete these tasks before you deploy Active Directory, you must complete them for the disjoint namespace to function properly or to receive support from Microsoft for problems with the disjoint namespace.

| **Note** |
| --- |
| If there are manually configured service principal names (SPNs) on any of your servers, they may require modification. For instructions on how to manually configure SPNs, see Setspn Overview. |

The following sections describe the tasks that you must complete to create a working disjoint namespace.

**Create appropriate DNS forward lookup zones**
On a DNS server for the domain, you must add the appropriate forward lookup zone for each primary DNS suffix. For information about adding a forward lookup zone to Windows 2000 Server, see article 308201 in the Microsoft Knowledge Base (http://support.microsoft.com/kb/308201). For information about adding a forward lookup zone using a Windows Server 2003–based computer, see article 323445 in the Microsoft Knowledge Base (http://support.microsoft.com/kb/323445).

**Configure different primary DNS suffixes**
You can configure the **Primary DNS Suffix** policy setting on the DNS client by using Group Policy, by manually changing the settings in the **System Properties** dialog box on each computer, or by modifying the registry on each computer. The following sections describe these options.

**Configure primary DNS suffixes by using a GPO**
On Windows Server 2003–based servers, you use a Group Policy object (GPO) to configure the primary DNS suffix for each member computer. You set the **Primary DNS Suffix** policy setting to the primary DNS suffix that you want to use on the member computers that will apply this policy setting. The GPO path for the **Primary DNS Suffix** policy setting is Computer Configuration\Administrative Templates\Network\DNS Client.

| **Note** |
| --- |
| For more information about managing Group Policy, see Group Policy Object Editor Tools and Settings (http://go.microsoft.com/fwlink/?LinkId=102368) and Enterprise Management with the Group Policy Management Console (http://go.microsoft.com/fwlink/?LinkId=29909) |

**Change the primary DNS suffix by using System Properties**
Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To manually set the primary DNS suffix on a local computer running Windows 2000, Windows XP, or Windows Vista, complete the following procedure.

**Change the primary DNS suffix by using System Properties**

1. Open the **System Properties** dialog box.

   a. To open the **System Properties** dialog box in Windows 2000 or Windows XP, click **Start**, click **Run**, type **sysdm.cpl**, and then press ENTER.

   b. To open the **System Properties** dialog box in Windows Vista, click **Start**, in **Start Search** type **sysdm.cpl**, and then press ENTER.

2. On the **Computer Name** tab, click **Change**, and then click **More**.
3. In **Primary DNS suffix of this computer**, set the primary DNS suffix to the domain name that you want.
4. Verify that the **Change primary DNS suffix when domain membership changes** check box is cleared, and then click **OK** twice.
5. Click **OK** to confirm that the computer must be restarted, and then click **Close**.
6. When you are prompted to restart your computer, click **Restart Now**.
7. If you choose to restart the computer later, any changes that you made will take effect at that time.

**Change the primary DNS suffix by using the registry**

If you plan to use a script (or to modify the registry manually) when you configure the primary DNS suffix, you must modify the **Domain** and **SyncDomainWithMembership** values. Both of these values are in the **Parameters** key under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip**. To allow a disjoint namespace to be created, disable the **SyncDomainWithMembership** value by setting its value to **0**, and then enter the name of the primary DNS suffix that you want to use for the **Domain** value. If the **SyncDomainWithMembership** value does not exist, you can create it as a REGDWORD data type. For more information about registry scripts, see article 264584 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=102370). If you configure these values using a registry script or Registry Editor, you must restart the computer for the settings to take effect.

**Enable computers to register disjoint names**

By default, Active Directory does not allow you to register another domain name as the DNS suffix for a domain. If you want to use a disjoint namespace, you must modify the security settings on the **msDS-AllowedDNSSuffixes** attribute. Otherwise, the default security settings prevent member computers from registering a different namespace.

**Allow additional primary DNS suffixes**

To add the additional primary DNS suffixes to the list of allowed suffixes in Active Directory, perform the following steps.

> **Note**
>
> ADSI Edit (AdsiEdit.msc) is a Windows Server tool that you can use to view and edit raw Active Directory directory service attributes through the Active Directory Service Interfaces (ADSI) protocol. ADSI Edit is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center. To download Windows Server 2003 Service Pack 2 32-bit Support Tools, see (http://go.microsoft.com/fwlink/?LinkId=100114). For information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To add the additional primary DNS suffixes to the list of allowed suffixes in Active Directory**

1. On any domain controller in the Active Directory domain, click **Start**, click **Run**, type **adsiedit.msc**, and then press ENTER.

   > **Note**
   >
   > If you recently installed Windows Support Tools to the default location, you may have to type **C:\Program Files\Support Tools\ADSIEdit.msc**, and then press ENTER. If you did not install Windows Support Tools to the default location, use the file path to which the tools were installed.

2. Expand the domain naming context (NC) to expose the directory partition that relates to the Active Directory domain for

which you want to allow a disjoint namespace.

3. Right-click the Active Directory domain partition, and then click **Properties**.

   For example, DC=fabrikam,DC=com is the domain partition for Fabrikam.com.

4. In **Attributes**, click the **msDS-AllowedDNSSuffixes** attribute, and then click **Edit**.

5. In **Value to add**, type the name of a primary DNS suffix that you want to allow.

   For example, type na.corp.fabrikam.com if you have computers that will use that primary DNS suffix, and then click **Add**. To configure multiple DNS suffixes, repeat this step for each one.

6. When you have entered all the additional DNS suffixes that you would like to allow, click **OK** twice to confirm your changes.

7. Close ADSI Edit.

If the domain has no Windows 2000–based computers, you should skip the following section and go directly to Update the DNS Suffix Search List.

**Allow Windows 2000–based computers to update attributes**

If the domain has member computers or domain controllers running Windows 2000, you must allow these computers to update the **dNSHostName** and **servicePrincipleName** attributes of their Active Directory computer accounts.

| **Note** |
| --- |
| These steps do not have to be performed if you have no Windows 2000 computers on your network. |

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To allow Windows 2000–based computers to update attributes**

1. Open Active Directory Users and Computers.

   To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then press ENTER.

2. Click **View**, and then select the **Advanced Features** check box, if it is not already selected.

3. If the domain to which you want to allow a disjoint namespace is not in the console tree, perform the following steps:

   a. In the console tree, right-click Active Directory Users and Computers, and then click **Connect to Domain**.

   b. In **Domain**, type the name of the Active Directory domain to which you want to allow the disjoint namespace, and then click **OK**. (You can also use the **Browse** button to locate the name of the domain.)

4. In the console tree, right-click the domain to which you want to allow a disjoint namespace, and then click **Properties**.

5. On the **Security** tab, click **Advanced**.

6. On the **Permissions** tab, click **Add**.

7. In **Enter the object name to select**, type **self**, and then click **OK**.

8. In the **Apply onto** box, select **Computer objects**.

9. At the bottom of the **Permissions** box, select the **Allow** check boxes that correspond to the **Validated write to DNS host name** and **Validated write to service principal name** permissions.

10. Click **OK** three times to confirm your changes and close the dialog boxes.

11. Close Active Directory Users and Computers.

**Update the DNS suffix search list**

When you transition to a disjoint namespace, consider creating customized DNS suffix search lists to ensure that clients can

locate services and other computers when they perform single-label name queries. Single-label name queries occur when a computer or user queries a name, but does not provide a DNS suffix. The DNS suffixes that are you add to the DNS suffix search list are used by the DNS client to create fully qualified domain names (FQDNs) when resolving single-label names. For example, if your computer sends a query for the name computer2, the DNS suffix search list can be used by your computer to make a name resolution query for computer2.contoso.com, computer2.na.corp.contoso.com, computer2.corp.contoso.com, or whatever is in the DNS suffix search list.

You can configure the DNS suffix search list by using Group Policy or by modifying each computer on the network through the user interface (UI) or registry. The following sections discuss the three primary methods for modifying the DNS suffix search list:

- Group Policy
- Network connections
- The registry

**Change the DNS suffix search list by using a GPO**

On Windows Server 2003–based servers, you use a GPO to configure the DNS suffix search list for each member computer. You set the DNS Client policy for **DNS Suffix Search List** to the DNS suffix search list that you want to use on the member computers that will apply this policy setting. The GPO path for the **DNS Suffix Search List** policy setting is Computer Configuration\Administrative Templates\Network\DNS Client.

| **Note** |
| --- |
| By default, the **DNS Suffix Search List** policy setting is not available on a server running Windows 2000 Server. However, you can add it by using the Windows XP administrative templates. For more information, see article 307900 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=102373). |

For more information about managing Group Policy, see Group Policy Object Editor Tools and Settings (http://go.microsoft.com/fwlink/?LinkId=102368) and Enterprise Management with the Group Policy Management Console (http://go.microsoft.com/fwlink/?LinkId=22814)

**Change the DNS suffix search list by using network connections**

To change the **DNS Suffix Search List** policy setting by using the UI in Windows 2000, Windows XP, or Windows Vista, complete the following procedure.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To change the DNS suffix search list by using network connections**

1. Open the **System Properties** dialog box.
   a. To open the **System Properties** dialog box in Windows 2000 or Windows XP, click **Start**, click **Run**, type **sysdm.cpl**, and then press ENTER.
   b. To open the **System Properties** dialog box in Windows Vista, click **Start**, in **Start Search** type **sysdm.cpl**, and then press ENTER.
2. Right-click the icon that represents the computer's local network connection, and then click **Properties**.

   **Local Area Connection** is the default connection.
3. In **This connection uses the following items**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Advanced**.
5. On the **DNS** tab, click **Append these DNS suffixes (in order)**
6. Click **Add**, and then type the first DNS suffix that you want the clients to use when they search for single-label names (those that do not have DNS suffixes).

7. Repeat the previous step for each DNS suffix that you want clients to try when resolving a name.

| **Note** |
| --- |
| Be sure to include the Active Directory domain name and all other domain names in your organization that clients might have to resolve in the appropriate order. |

8. Click **OK** twice to confirm your changes.

9. Close the **Local Area Connection Properties** and the **Network Connections** dialog boxes.

**Change the DNS suffix search list in the registry**

If you plan to use a script (or to modify the registry manually) when you configure the DNS suffix search list, you can modify the **SearchList** value that is located in the **Parameters** key under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip**. For more information about registry scripts, see article 264584 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=102370).

# Configure Site Topology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The site topology owner configures the site topology for the forest. Configuring the site topology for the forest involves the following tasks:

- Delegate Active Directory site administration
- Create Active Directory sites
- Create and assign Active Directory subnets
- Create Active Directory site links

# Delegate Active Directory Site Topology Administration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configuring the site topology for the forest starts when the forest owner delegates administration of the Active Directory sites and site topology to the site topology owner.

**To delegate Active Directory site topology administration in your environment**

1. Create a global group named **SiteAdmins** in the forest root domain.

2. Add administrative users to the **SiteAdmins** global group.

   **Note**

   - The user accounts that you add must reside in the forest root domain. If you want to add users from regional domains to this group, it must be a universal group, and the forest root domain and the regional domain must be operating at the Microsoft® Windows® 2000 native or Windows Server 2003 functional level.

3. In **Active Directory Sites and Services**, right-click the **Sites** node, and then click **Delegate Control**.

4. Complete the Delegation of Control Wizard to delegate Full Control of the **Sites** node to **SiteAdmins**.

# Create Active Directory Sites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Create Active Directory sites by using Active Directory Sites and Services.

**To create the Active Directory sites**

1. Review the site topology design information in the Associating Subnets with Sites worksheet provided by your design team.

2. Create the sites specified in the site topology design. For more information about how to create site objects, see "Create a site" in Help and Support Center for Windows Server 2003.

# Create and Assign Active Directory Subnets

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Create Active Directory subnets by using Active Directory Sites and Services.

**To create Active Directory subnets and associate them with sites**

1. Review the site topology design information in the "Associating Subnets with Sites" worksheet provided by your design team.

2. Create the Active Directory subnets specified in the worksheet and associate the Active Directory subnet with the appropriate site. For more information about how to create subnet objects and associate subnets to Active Directory sites, see "Create a subnet" and "Associate a subnet with a site" in Help and Support Center for Windows Server 2003.

# Create Active Directory Site Links

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Create Active Directory site links by using Active Directory Sites and Services.

**To create Active Directory site links**

1. Review the site topology design information in the "Sites and Associated Site Links" worksheet provided by your design team.

2. Create the Active Directory site links and configure the site link properties as specified in the site topology design. For more information about how to create site link objects, see "Create a site link" on the Microsoft Web site.

# Deploy Additional Domain Controllers in Other Sites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your design specifies deployment of additional forest root domain controllers in other sites, deploy them by using the procedures in "Deploy the Second Domain Controller in the Same Site" earlier in this chapter.

# Configure Operations Master Roles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure the forest-level and domain-level operations master roles for the forest root domain. By default, the first domain controller in the forest root domain is assigned all operations master roles.

- If your design specifies that all domain controllers in the forest root domain are global catalog servers, leave all five operations master roles on the first domain controller, and designate the second domain controller as the standby.

- If any domain controllers in the forest root domain will not be global catalog servers, move all operations master roles from the first domain controller to the second domain controller, and ensure that the second domain controller will never be a global catalog server. Designate the third domain controller as the standby, and never make it a global catalog server.

  **Note**

  - In a single domain forest, the database content of a domain controller and a global catalog server are the same. Therefore, to load balance client lookups across global catalog servers in a single domain forest, ensure that all domain controllers are global catalog servers.

For a procedure to help you transfer operations master roles, see "Transfer operations master roles" in Help and Support Center for Windows Server 2003.

If your Active Directory design specifies that you designate a standby operations master for the current operations master role holder, configure the current role holder and the standby as direct replication partners by manually creating a connection object between them. Designating a standby operations master can save some time if you must reassign any operations master roles to the standby operations master.

Of all the operations master roles, the PDC emulator role has the highest impact on the performance of the domain controller hosting that role. In domains with more than 10,000 users, it might be necessary to reduce the number of authentication requests performed by the PDC emulator to decrease its workload and allow it to perform other tasks. If CPU utilization is higher than 50 percent or disk queues remain higher than 2 for several hours or days, reduce the number of client authentication requests received by the PDC emulator.

**Note**

- Other factors that can increase the workload on the PDC emulator include pre-Active Directory clients or applications that have been written to contact the PDC emulator.

To reduce the number of client authentication requests that are processed by the PDC emulator, adjust its weight or its priority in the DNS environment. If you want to proportionately reduce the number of client authentication requests received by the PDC emulator, adjust its weight. If you want to ensure that the PDC emulator does not receive any client authentication requests, adjust its priority.

Active Directory assigns a default value of 100 for the weight. By creating a new registry entry for the weight and assigning it a decreased value of 50, you can proportionately reduce the number of client authentication requests sent to the PDC emulator. This ensures that the PDC emulator will authenticate half of the number of clients than it would if the weight value remained at 100.

Active Directory assigns a default value of zero for the priority. By creating a new registry entry for the priority and assigning it an increased value of 200, you can ensure that the PDC emulator will never receive client authentication requests unless it is the only accessible domain controller.

Repeat these procedures if the PDC emulator operations master role is transferred or seized to another domain controller in the forest root domain.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to

reinstall Windows. If you must edit the registry, back it up first and see the "Registry Reference" on the *Microsoft Windows Server 2003 Deployment Kit* companion CD or on the Microsoft Web site.

**To change the weight for DNS SRV records by using the registry**

1. In the **Run** dialog box, type **regedit**, and press ENTER.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters.
3. Click **Edit**, click **New**, and then click **DWORD value**.
4. For the new entry name, type **LdapSrvWeight** and press ENTER. (The value name is not case sensitive.)
5. Double-click the entry name you just typed to open the **Edit DWORD Value** dialog box.
6. Choose **Decimal** as the **Base** option.
7. Enter a value from 0 through 65535. The recommended value is 50.
8. Click **OK**.
9. Click **File**, and then click **Exit** to close the registry editor.

Adjusting the priority of the domain controller also reduces the number of client referrals. However, rather than reducing it proportionally to the other domain controllers, changing the priority causes DNS to stop referring all clients to this domain controller unless all domain controllers with a lower priority setting are unavailable.

**Note**

- A lower value entered for **LdapSrvPriority** indicates a higher priority. A domain controller with an **LdapSrvPriority** setting of 100 has a lower priority than a domain controller with a setting of 10. Therefore, clients attempt to use the domain controller with the setting of 10 first.

**To change the priority for DNS SRV records by using the registry**

1. In the **Run** dialog box, type **regedit**, and press ENTER.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters
3. Click **Edit**, click **New**, and then click **DWORD value**.
4. For the new entry name, type **LdapSrvPriority**, and press ENTER.
5. Double-click the entry name that you just typed to open the **Edit DWORD Value** dialog box.
6. Choose **Decimal** as the **Base** option.
7. Enter a value from 0 through 65535. The recommended value is 200.
8. Click **OK**.
9. Click **File**, and then click **Exit** to close the registry editor.

For more information about adjusting the weight or the priority of the PDC emulator, see the Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search under "Administration and Configuration Guides" and download the *Active Directory Operations Guide*.

# Raising the Functional Level

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you deploy the first Windows Server 2003–based domain controller in your forest root domain, the forest operates by default at the Windows 2000 forest functional level, and the domain operates by default at the Windows 2000 mixed functional level. If your organization has only Windows NT 4.0 domains and your forest design requires upgrading Windows NT 4.0 domains and joining them to this Active Directory forest, you might want to raise the forest functional level to Windows Server 2003 interim after you deploy the forest root domain and before you begin the process for upgrading your Windows NT 4.0 domain to Windows Server 2003 Active Directory. If you raise the forest functional level to Windows Server 2003 interim, you can take advantage of the advanced features available at that functional level.

Figure 6.8 shows raising the functional level as the last step in the forest root domain deployment process.

**Figure 6.8 Raising the Functional Level**




Although the Windows Server 2003 domain functional level provides a number of features and advantages, only enable this functional level when your environment is ready and all of your Windows NT 4.0–based backup domain controllers (BDCs) have been upgraded.

Although it is possible for a domain to include both Windows NT 4.0–based and Windows Server 2003–based domain controllers, the Windows Server 2003 domain functional level provides more features.

When you have determined that your environment is ready, use Active Directory Domains and Trusts to enable the Windows Server 2003 domain functional level.

After all domains are operating at the Windows Server 2003 functional level, raise the forest functional level to Windows Server 2003 to take advantage of all Windows Server 2003 forest-level features.

To determine when to raise the functional level, and for procedures to perform those tasks, see "Enabling Advanced Windows Server 2003 Active Directory Features" in this book. For more information about upgrading Windows NT domains to Windows Server 2003 Active Directory, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.

# Additional Resources for Deploying the Root Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing the Active Directory Logical Structure" in this book.
- "Designing the Site Topology" in this book.
- "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.
- "Deploying Windows Server 2003 Regional Domains" in this book.
- The Active Directory Branch Office Planning Guide link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for a complete guide to information involving Active Directory branch office implementations.
- The Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search under "Administration and Configuration Guides" and download the *Active Directory Operations Guide*.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Active Directory" in Help and Support Center for Windows Server 2003.
- "Install Windows Support Tools" under "Tools" in Help and Support Center for Windows Server 2003.
- "Configure Site Settings" in Help and Support Center for Windows Server 2003 for more information about creating site objects, subnet objects, and associating subnets with sites.
- "Transfer operations master roles" in Help and Support Center for Windows Server 2003 for a procedure to help you transfer operations master roles.
- "Understanding aging and scavenging" in Help and Support Center for Windows Server 2003 for more information about how to configure aging and scavenging of stale resource records.

**Related Job Aids**

- "Domain Controller Configuration" (DSSDFR_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Controller Configuration" on the Web at http://www.microsoft.com/reskit).
- "Sites and Associated Site Links" (DSSTOPO_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Sites and Associated Site Links" on the Web at http://www.microsoft.com/reskit).
- "Associating Subnets with Sites" (DSSTOPO_6.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Associating Subnets with Sites" on the Web at http://www.microsoft.com/reskit).

# Deploying Windows Server 2003 Regional Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying Microsoft® Windows® Server 2003 regional domains involves creating new geographically based child domains under the forest root domain. Deploying new Windows Server 2003 regional domains allows you to create a new domain structure or improve your existing domain structure to take advantage of available bandwidth and optimize replication traffic.

**In This Chapter**


Overview of Deploying Regional Domains

Reviewing the Regional Domain Design

Delegating the DNS Domain for the New Regional Domain

Deploying the First Domain Controller in a New Regional Domain

Deploying Additional Domain Controllers in a New Regional Domain

Reconfiguring the DNS Service

Configuring Operations Master Roles

Additional Resources for Deploying Windows Server 2003 Regional Domains


**Related Information**

- For more information about the Active Directory® directory service logical structure, see "Designing the Active Directory Logical Structure" in this book.
- For more information about Active Directory site topology, see "Designing the Site Topology" in this book.
- For more information about hardware requirements for domain controllers, see "Planning Domain Controller Capacity" in this book.
- For more information about deploying the forest root domain, see "Deploying the Windows Server 2003 Forest Root Domain" in this book.

# Overview of Deploying Regional Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you complete the deployment of the forest root domain, you are ready to deploy any new Windows Server 2003 regional domains specified by your design. Windows Server 2003 regional domains include domain controllers that are running the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; or Windows® Server 2003, Datacenter Edition operating systems.

The design team determines which regional domains to deploy. As an alternative or in addition to creating a new regional domain, your Active Directory design might require upgrading existing domains. For more information about performing an in-place upgrade of a Microsoft® Windows NT® 4.0 account domain to a regional domain, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book. For information about performing an in-place upgrade of a Microsoft® Windows® 2000 domain to a regional domain, see "Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" in this book.

**Note**

- You can simultaneously create multiple regional domains and perform in-place upgrades of multiple Windows NT 4.0 or Windows 2000 domains. You do not need to complete the deployment of one regional domain before proceeding to the next regional domain.

After you deploy a new regional domain, you can create an organizational unit (OU) structure and either add new users, groups, and resources to the domain, or use the domain as a target for restructuring existing Windows NT 4.0 or Windows 2000 domains.

# Process for Deploying Regional Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying regional domains involves completing the tasks that are shown in Figure 7.1.

**Figure 7.1 Deploying Regional Domains**



# Background Information for Deploying Regional Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy regional domains, become familiar with how you can save time during the deployment process by automating installations and using the Active Directory Installation Wizard.

**Automating Installations**

You can perform an unattended installation of Active Directory by supplying an answer file when you run the Active Directory Installation Wizard. For more information about using an answer file to install Active Directory, in Help and Support Center for Windows Server 2003, click **Tools**, and click Windows Support Tools. Search for Ref.chm in the Deploy.cab file for examples and instructions about creating an answer file.

You can automate the installation of Windows Server 2003 by using Sysprep.exe or an unattended installation. For more information about automating installations, see *Automating and Customizing Installations* in this kit.

**Using the Active Directory Installation Wizard**

Windows Server 2003 includes improvements to the Active Directory Installation Wizard. When you install the first domain controller in a domain, you can allow the wizard to automatically install and configure Active Directory–integrated DNS zones. Even if you choose to manually configure some settings later, the wizard saves time and prevents errors during the initial configuration.

# Reviewing the Regional Domain Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In preparation for deployment, gather and review all design information from the design team and ensure that the domain controller hardware you plan to use meets the requirements for Windows Server 2003.

Reviewing the regional domain design involves two tasks, as shown in Figure 7.2.

**Figure 7.2 Reviewing the Regional Domain Design**



# Collect Regional Domain Design Information

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin the deployment, collect the design information that documents domain controller assignments. Record the computer name, static IP address, subnet mask, and DNS client settings for each domain controller in the regional domain. For a blank worksheet to assist you in documenting this information, see "Domain Controller Configuration" (DSSDFR_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Domain Controller Configuration" on the Web at http://www.microsoft.com/reskit).

**Example: Reviewing the Regional Domain Design for Trey Research**
The regional domain design for Trey Research specifies deployment of a new regional domain, west.trccorp.treyresearch.net. This regional domain is a child domain of the dedicated Windows Server 2003 forest root domain, trccorp.treyresearch.net.

**Note**

- The other regional domain in the example,east.trccorp.treyresearch.net, is created by upgrading a Windows NT 4.0 account domain, then restructuring other Windows NT 4.0 resource domains. Those portions of the example are covered in "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory," and "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.

The West domain will have domain controllers in two locations: Seattle and Los Angeles. Figure 7.3 shows an example of a completed regional domain controller configuration worksheet for the West regional domain for Trey Research. For each domain controller, the preferred DNS server is the local domain controller, and the alternate is the closest DNS server. Initially, however, Trey Research configures the first domain controller in the domain to use a DNS server in its parent domain as its preferred DNS server. During installation of Active Directory on the first domain controller, the preferred DNS server is changed to the local domain controller. Then, after the second domain controller is online, Trey Research reconfigures the first domain controller to use the second domain controller as its alternate DNS server.

**Figure 7.3 Example of a Domain Controller Configuration Worksheet**

| Domain Controller Configuration | | | | |
|---|---|---|---|---|
| Prepared By | Active Directory Deployment Team | | Date | 01/01/03 |
| Regional Domain Name | west.trccorp.treyreserach.net | | | |
| Forest Owner | forest owner | | | |
| Parent Domain Name | trccorp.treyresearch.net | | | |
| Domain Controller | IP Address | Mask | Preferred DNS Server | Alternate DNS Server |
| SEA-WEST-DC01 | 172.16.16.20 | 255.255.252.0 | 172.16.16.2 (temp) 172.16.16.20 (final) | 172.16.16.3 (temp) 172.16.16.21 (final) |
| SEA-WEST-DC02 | 172.16.16.21 | 255.255.252.0 | 172.16.16.21 | 172.16.16.2 |
| LAX-WEST-DC01 | 172.16.136.2 | 255.255.252.0 | 172.16.136.2 | 172.16.136.3 |
| LAX-WEST-DC02 | 172.16.136.3 | 255.255.252.0 | 172.16.136.3 | 172.16.136.2 |

# Review Hardware Requirements for Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Ensure that each computer that you plan to use as a domain controller meets the hardware requirements for running Windows Server 2003. For more information about assessing the hardware requirements of domain controllers running Windows Server 2003, see "Planning Domain Controller Capacity" in this book.

# Delegating the DNS Domain for the New Regional Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory DNS owner is responsible for delegating the zone matching the name of the regional domain to the DNS servers (domain controllers) that you will be installing in the regional domain.

Figure 7.4 shows when the DNS domain is delegated during the regional domain deployment process.

**Figure 7.4 Delegating the DNS Domain for the New Regional Domain**




Before you create the new regional domain, delegate the DNS domain for the new Windows Server 2003 regional domain:

- In the forest root domain DNS zone.
- On any forest root domain controller.

**To delegate the DNS domain for the new regional domain**

1. Open the DNS snap-in from any domain controller in the forest root domain.
2. In the console tree, right-click the forest root domain zone, and then click **New Delegation**.
3. Complete the New Delegation Wizard by using the information documented in your Domain Controller Configuration worksheet and the instructions in Table 7.1. Accept the default settings when no information is supplied. Table 7.1 includes sample data entered by Trey Research for their first two regional domain controllers, SEA-WEST-DC01 and SEA-WEST-DC02.

   **Table 7.1 Delegating the DNS Domain for the Regional Domain**

| Wizard Page | Action | Example |
|---|---|---|
| Delegated Domain Name | In the **Delegated Domain** box, type the name of the regional domain. | **west** |
| Name Servers | Click **Add**. In the **New Resource Record** dialog box, in the **Server name** box, type the name of the first domain controller you plan to deploy. | **SEA-WEST-DC01. west.trccorp.treyresearch.net** |

| | |
|---|---|
| In the **New Resource Record** dialog box, in the **IP address** box, type the corresponding IP address of the domain controller, click **Add**, and then click **OK**. | **172.16.16.20** |
| Click **Add**, and in the **New Resource Record** dialog box, in the **Server name** box, type the name of another domain controller you plan to deploy in the regional domain. | **SEA-WEST-DC02. west.trccorp.treyresearch.net** |
| In the **New Resource Record** dialog box, in the **IP address** box, type the corresponding IP address of the other domain controller, click **Add**, and then click **OK**. | **172.16.16.21** |

# Deploying the First Domain Controller in a New Regional Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For each regional domain, the forest owner is responsible for deploying the domain controllers. Figure 7.5 shows the process for deploying the first domain controller in a new regional domain.

**Figure 7.5 Deploying the First Domain Controller in a New Regional Domain**



# Install Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in deploying a regional domain is to install Windows Server 2003 on the first computer that you want to make the domain controller.

**Note**

- Before installing Windows Server 2003, ensure that DNS was never previously installed on the computer. If DNS was previously installed, configuration of the DNS resolver and forwarders might fail.

  - Insert the Windows Server 2003 operating system CD in the CD-ROM drive of the domain controller and select the option to install the operating system, or use an automated installation method. If the Windows Server 2003 media is shared on the network, run the Winnt32.exe command. Use the NTFS file system to format the partitions. Enter the computer name, static IP address, and subnet mask as specified by your design.

  - Configure the DNS client settings by using the information documented in the Domain Controller Configuration worksheet. The DNS settings are temporary and will be changed after you install Active Directory.

  - Enable Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode) to enable administrators to log on remotely if necessary. To enable Remote Desktop for Administration, in Control Panel, double-click **System**, select the **Remote** tab, and then select **Allow users to connect remotely to this computer**.

  - Install Windows Support Tools, which are available in the \Support\Tools folder on the Windows Server 2003 operating system CD.

# Install Active Directory by Running the Active Directory Installation Wizard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Install Active Directory by running the Active Directory Installation Wizard on the computer that you want to make the first regional domain controller. The Active Directory Installation Wizard creates the Active Directory database and initializes the directory data in the database.

In addition, on the first domain controller in a domain, after you run the wizard and restart the computer, the wizard also:

- Prompts the administrator to verify the installation and configuration of the DNS Server service.
- Configures DNS recursive name resolution by forwarding, by adding the IP addresses of the existing entries for **Preferred DNS server** and **Alternate DNS server** to the list of DNS servers on the **Forwarders** tab of the **Properties** sheet in the DNS snap-in for the domain controller.
- Configures DNS recursive name resolution by root hints, by adding the root hints that are configured on the **Preferred DNS server**.

  **Notes**

  - If you want to set different forwarders, or do not want to enable forwarding, you can change this setting manually by using the DNS snap-in.
  - If your domain controller is multihomed, forwarding is not configured automatically.
  - For manual configuration instructions, see "Verify DNS Server Recursive Name Resolution" later in this chapter.
- Configures the **Preferred DNS server** to point to the DNS server that is running locally on the domain controller, and configures the **Alternate DNS server** to point to the DNS server that is connected through the minimum number of network segments.

**To install Active Directory on the first regional domain controller**

1. Log on to the Windows Server 2003–based member server.
2. Verify that you can contact the parent DNS zone. At the command line, type:

   **nslookup parent_domain**

   (using the fully qualified domain name of the parent DNS zone of the regional domain.)
3. At the command line, type:

   **dcpromo**

   - or -

   Open **Administrative Tools** and click **Configure Your Server Wizard**. Select **Domain Controller (Active Directory)** to configure your domain controller. When the Configure Your Server Wizard finishes, the Active Directory Installation Wizard begins.
4. Use Table 7.2 to complete the Active Directory Installation Wizard for the first domain controller in the regional domain. Table 7.2 also includes information entered by Trey Research as they deploy their first regional domain controller, SEA-

WEST-DC01.

**Table 7.2 Installing Active Directory on the First Domain Controller in a New Regional Domain**

| Wizard Page or Dialog Box | Action | Example |
|---|---|---|
| Domain Controller Type | Select **Domain controller for a new domain**. | |
| Create New Domain | Select **Create a new child domain in an existing domain tree**. | |
| Network Credentials | In the **User name** box, type a user account that has sufficient rights to add a domain, typically a member of Enterprise Admins.<br><br>In the **Password** box, type the password of the user account.<br><br>In the **Domain** box, type the name of the forest root domain. | Forest root domain: **trccorp.treyresearch.net** |
| Child Domain Installation | Click **Browse**.<br><br>In the **Browse for Domain** dialog box, click the Parent domain name, and then click **OK**.<br><br>In the **Child Domain** box, type the NetBIOS name of the regional domain. | Parent domain: **trccorp.treyresearch.net**<br><br>Child Domain: **west** |
| NetBIOS Domain Name | Confirm or type the NetBIOS name. | **WEST** |
| Database and Log Folders | Type the folder locations specified by your design. | Database folder: C:\WINNT\NTDS<br><br>Log folder: D:\Logs |
| Shared System Volume | Confirm or type the location specified by your design. | C:\WINNT\SYSVOL |
| DNS Registration Diagnostics | A message will indicate that the wizard cannot contact the DNS server with which this domain controller will be registered. This is because the pre-created delegation record points to the local computer and DNS has not been installed on the domain controller at this point.<br><br>Select **Install and configure the DNS server on this computer and set this computer to use this DNS server as its preferred DNS server**. | |
| Permissions | Select the security level specified by your design:<br><br>• Permissions compatible with pre-Windows 2000 server operating systems<br><br>• Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems | Because Trey Research currently has server programs running on Windows NT 4.0–based servers, they selected **Permissions compatible with pre-Windows 2000 server operating systems**. |

| | | |
|---|---|---|
| Directory Service Restore Mode Administration Password | In the **Password** and **Confirm password** boxes, type any strong password. | |

For more information about installing and removing Active Directory, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

# Verify the Active Directory Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you run the Active Directory Installation Wizard to install Active Directory, verify the Active Directory installation by performing the following tasks:

1. Review the Windows Server 2003 event log and resolve any errors.

2. At the command line, run Dcdiag.exe and review and resolve any errors that are reported.

3. Run Task Manager and verify that the processor and memory system resources are within acceptable limits.

4. Open the DNS snap-in, expand **Forward Lookup Zones**, and verify that the zones _msdcs.*forest_root_domain_name* and *regional_domain_name* are present.

# Verify DNS Server Recursive Name Resolution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DNS server recursive name resolution is configured automatically during the Active Directory installation process. If your design specifies a different configuration, you can use the DNS snap-in or Dnscmd.exe to modify these settings.

**To verify name resolution on the regional domain controller**

- Use the DNS snap-in to verify DNS server recursive name resolution for the method used in your organization, based on the information in Table 7.3.

**Table 7.3 Information to Verify DNS Server Recursive Name Resolution**

| Method | Configuration |
|---|---|
| Recursive name resolution by root hints | Root hints are the recommended method to use for recursive name resolution in a Windows Server 2003 environment.<br><br>No additional configuration is necessary.<br><br>When the DNS server specified as the Preferred DNS server during the installation process is properly configured, the root hints are automatically configured.<br><br>To verify the root hints by using the DNS snap-in:<br><br>In the console tree, right-click the domain controller name, and then click **Properties**. Click the **Root Hints** tab to view the root hints. |
| Recursive name resolution by forwarding | Only use Forwarders if that is what your organization's design specifies. Root hints are the recommended method to use for recursive name resolution in a Windows Server 2003 environment.<br><br>Forward unresolved queries to specified DNS servers.<br><br>To verify forwarding by using the DNS snap-in:<br><br>In the console tree, right-click the domain controller name, and then click **Properties**.<br><br>On the **Forwarders** tab, in the selected domain's **Forwarders** list, verify that the IP addresses match those specified by your design. |

# Deploying Additional Domain Controllers in a New Regional Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 7.6 shows the process for deploying additional domain controllers in a new regional domain.

**Figure 7.6 Deploying Additional Domain Controllers in a New Regional Domain**




**Install Windows Server 2003**

Install Windows Server 2003 on additional computers by completing the same steps as for installing Windows Server 2003 on the first domain controller. Note that you do not need to configure temporary DNS client settings on additional domain controllers; instead, enter the IP addresses of the DNS servers that are connected through the minimum number of network segments, as documented in the Domain Controller Configuration worksheet. For more information about installing Windows Server 2003, see "Install Windows Server 2003" earlier in this chapter.

**Install Active Directory**

Install Active Directory on the computer that you want to make an additional regional domain controller by running the Active Directory Installation Wizard. You must complete the Active Directory installation and verify that installation is correct on the first domain controller before beginning installation of Active Directory on the second and subsequent domain controllers in the domain. For more information about verifying the Active Directory installation on the first regional domain controller, see "Verify the Active Directory Installation" earlier in this chapter.

The procedure for installing Active Directory on additional domain controllers in a domain differs slightly from the procedure for installing Active Directory on the first domain controller. For additional domain controllers in an existing domain, you have the option of using the install from media feature, which is new in Windows Server 2003. Install from media allows you to pre-populate Active Directory with System State data backed up from an existing domain controller. This backup can be present on local CD, DVD, or hard disk partition. Installing from media drastically reduces the time required to install directory information by reducing the amount of data that is replicated over the network. Installing from media is most beneficial in large domains or for installing new domain controllers that are connected by a slow network link. To use the install from media feature, you first create a backup of System State from the existing domain controller, then restore it to the new domain controller by using the **Restore to: Alternate location** option.

**To install Active Directory on additional regional domain controllers**

1. Log on to the Windows Server 2003–based member server.

2. If you want to copy domain information from restored backup files, at the command line, type:

   **dcpromo /adv**

– or –

If you want to copy domain information over the network, either type **dcpromo** (without the /adv switch) or open **Administrative Tools**, click **Configure Your Server Wizard**, andselect **Domain Controller (Active Directory)**.

3. Use Table 7.4 to complete the Active Directory Installation Wizard for subsequent domain controllers in the regional domain. Table 7.4 includes information entered by Trey Research as they deploy their second regional domain controller, SEA-WEST-DC02.

**Table 7.4 Installing Active Directory on Additional Domain Controllers in a New Regional Domain**

| Wizard Page or Dialog Box | Action | Example |
|---|---|---|
| Domain Controller Type | Select **Additional domain controller in an existing domain** | |
| Copying Domain Information | (This dialog box appears only when you started dcpromo with the /adv switch or used the **Configure Your Server Wizard**.)<br>Select either:<br>• Over the network from a domain controller<br>• From these restored backup files | Trey Research is copying from the first WEST domain controller, SEA-WEST-DC01,which is in the same location, so they selected **Over the network** to copy the information in the shortest time. |
| Global Catalog | (This dialog box appears only when **From these restored backup files** was selected, if the domain controller you backed up was a global catalog server.)<br>Specify whether to configure this domain controller as a global catalog server. | |
| Network Credentials | In the **User name** box, type a user account that has sufficient rights to add a domain controller, typically a member of Domain Admins.<br>In the **Password** box, type the password of the user account. | |
| Additional Domain Controller | (This dialog box appears only when **Over the network** was selected.)<br>Confirm or type the full DNS name of the forest root domain. | **west.trccorp.treyresearch.net** |
| Database and Log Folders | Type the folder locations specified by your design. | Database folder: C:\WINNT\NTDS<br>Log folder: D:\Logs |
| Shared System Volume | Confirm or type the location specified by your design. | C:\WINNT\SYSVOL |

| | | |
|---|---|---|
| Directory Service Restore Mode Administration Password | In the **Password and Confirm password** boxes, type any strong password. | |

For more information about installing and removing Active Directory, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

**Install DNS Server Service**

After Active Directory installation has finished and the computer has restarted, install DNS on the additional Windows Server 2003–based domain controller that is added to the domain.

**To install DNS on additional domain controllers using the Windows Components Wizard**

1. Click **Start**, point to **Settings**, and click **Control Panel**.
2. Double-click **Add or Remove Programs**, and then click **Add/Remove Windows Components**.
3. In **Components**, select the **Networking Services** check box, and then click **Details**.
4. In **Subcomponents of Networking Services**, select the **Domain Name System (DNS)** checkbox, click **OK**, and then click **Next**.
5. If prompted, in **Copy files from**, type the full path to the distribution files and then click **OK**. The required files will be copied to your hard disk.

**Verify the Active Directory Installation**

Use the same steps as shown in the procedure for the first domain controller, but instead of verifying that **DomainDnsZones** and **ForestDnsZones** were created, use the repadmin /showreps command to verify that the **ForestDnsZones** and **DomainDnsZones** application partitions are replicated successfully. Use the DNS snap-in to verify that DNS server recursive name resolution is configured according to the method used by your organization.

For more information about verifying the Active Directory installation, see "Verify the Active Directory Installation" earlier in this chapter.

# Reconfiguring the DNS Service

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Reconfigure the DNS service to enable aging and scavenging of stale resource records and to account for the additional domain controllers in the regional domain. Figure 7.7 shows the process for reconfiguring the DNS service.

**Figure 7.7 Reconfiguring the DNS Service**



# Enable Aging and Scavenging for Domain Name System

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Enable aging and scavenging on two Windows Server 2003–based domain controllers running the DNS Server service per domain, to allow automatic cleanup and removal of stale resource records (RRs), which can accumulate in zone data over time.

With dynamic update, RRs are automatically added to zones when computers start on the network. However, in some cases, they are not automatically removed when computers leave the network. For example, if a computer registers its own host (A) RR at startup and is later improperly disconnected from the network, its host (A) RR might not be deleted. If your network has mobile users and computers, this situation can occur frequently.

If left unmanaged, the presence of stale RRs in zone data might cause problems including:

- If a large number of stale RRs remain in server zones, they can eventually take up server disk space and cause unnecessarily long zone transfers.
- DNS servers loading zones with stale RRs might use outdated information to answer client queries, potentially causing the clients to experience name resolution problems on the network.
- The accumulation of stale RRs at the DNS server can degrade its performance and responsiveness.

**Caution**

- By default, the aging and scavenging mechanism for the DNS Server service is disabled. Enable aging and scavenging only after you understand all parameters. Otherwise, the server could be accidentally configured to delete resource records that should not be deleted. If a resource record is accidentally deleted, not only will users fail to resolve queries for that resource record, but any user can create the resource record and take ownership of it, even on zones configured for secure dynamic update.
- For more information about how to configure aging and scavenging, see "Understanding aging and scavenging" in Help and Support Center for Windows Server 2003.

To enable the aging and scavenging features, perform the following steps to configure the applicable server and its Active Directory–integrated zones:

- Enable aging and scavenging at the server. These settings determine the affect of zone-level properties for any Active Directory–integrated zones loaded at the server.
- Enable aging and scavenging for selected zones at the DNS server. When zone-specific properties are set for a selected zone, these settings apply only to the applicable zone and its resource records. Unless these zone-level properties are otherwise configured, they inherit their default settings from comparable settings maintained in the DNS server aging/scavenging properties.

**To set aging and scavenging properties for the DNS server**

1. Log on to the computer that is running the DNS Server service with an account that is a member of the local Administrators group.
2. In the DNS console tree, right-click the applicable DNS server, and then click **Set Aging/Scavenging for all zones**.
3. Select the **Scavenge stale resource records** check box.
4. Modify other aging and scavenging properties as needed.

**To set aging and scavenging properties for a zone**

1. Log on to the computer that is running the DNS Server service with an account that is a member of the local Administrators group.

2. In the DNS console tree, right-click the applicable zone, then click **Properties**.

3. On the **General** tab, click **Aging**, and then **s**elect the **Scavenge stale resource records** check box.

4. Modify other aging and scavenging properties as needed.

# Configure the DNS Client Settings of the First and Subsequent Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have deployed an additional domain controller, modify the DNS client settings on the first domain controller. Because no other domain controllers were running when you deployed the first domain controller, modify the DNS client settings on the first domain controller to include the additional domain controller. As you deploy more domain controllers, you might also need to modify the **Alternate DNS server** setting specified on existing domain controllers to ensure that this setting points to the DNS server that is connected through the minimum number of network segments.

**To configure the DNS client settings on previously installed domain controllers**

1. Open **Network Connections**, double-click your **Local Area Connection**, click **Properties**, click **Internet Protocol (TCP/IP)** to highlight it, and then click **Properties**.

2. For the **Preferred DNS server**, type the IP address of the DNS server that is running locally on the domain controller (local host).

3. Determine whether a newly installed domain controller is now closer to this domain controller than the domain controller that you originally specified as the **Alternate DNS server**. If it is, for **Alternate DNS server**, type the IP address of the newly installed domain controller.

# Update the DNS Delegation for the Regional Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deployed the first regional domain controllers, the DNS administrator delegated the DNS domain for the new regional domain. During that process, the DNS administrator added DNS name server (NS) and address (A) resource records to the parent DNS zone for each planned domain controller in your design. If you deploy additional domain controllers that you have not previously created resource records for, update the DNS delegation.

You can update the DNS delegation by using the DNS snap-in on any forest root domain controller, as described in "Delegating the DNS Domain for the New Regional Domain" earlier in this chapter.

# Configuring Operations Master Roles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure operations master roles to ensure the availability and performance of those domain controllers. Figure 7.8 shows when configuring operations master roles occurs in the regional domain deployment process.

**Figure 7.8 Configuring Operations Master Roles**




It is generally recommended that you leave the operations master roles on the first domain controller in the regional domain. Monitor this domain controller closely and ensure that it is not a global catalog server. Deploy an additional domain controller in the same site to act as the standby operations master.

Host the PDC emulator role on a powerful and reliable domain controller to ensure that it is available and capable of handling the workload.

Of all the operations master roles, the PDC emulator role has the highest impact on the performance of the domain controller hosting that role. In domains with more than 10,000 users, it might be necessary to reduce the number of authentication requests performed by the PDC emulator to decrease its workload and allow it to perform other tasks. If CPU utilization is higher than 50 percent or disk queues remain higher than 2 for several hours or days, reduce the number of client authentication requests received by the PDC emulator.

**Note**

- Other factors that can increase the workload on the PDC emulator include pre-Active Directory clients or applications that have been written to contact the PDC emulator.

To reduce the number of client authentication requests that are processed by the PDC emulator, adjust its weight or its priority in the DNS environment. If you want to proportionately reduce the number of client authentication requests received by the PDC emulator, adjust its weight. If you want to ensure that the PDC emulator does not receive any client authentication requests, adjust its priority.

Active Directory assigns a default value of 100 for the weight. By creating a new registry entry for the weight and assigning it a decreased value of 50, you can proportionately reduce the number of client authentication requests sent to the PDC emulator. This ensures that the PDC emulator will authenticate half of the number of clients that it would if the weight value remained at 100.

Active Directory assigns a default value of zero for the priority. By creating a new registry entry for the priority and assigning it an increased value of 200, you can ensure that the PDC emulator will never receive client authentication requests unless it is the only accessible domain controller.

Repeat these procedures if the PDC emulator operations master role is transferred or seized to another domain controller in the forest root domain.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Microsoft Windows Server 2003 Deployment Kit* companion CD or on the Web at http://www.microsoft.com/reskit.

**To change the weight for DNS SRV records by using the registry**

1. In the **Run** dialog box, type **regedit**, and press ENTER.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters.
3. Click **Edit**, click **New**, and then click **DWORD value**.
4. For the new entry name, type **LdapSrvWeight** and press ENTER. (The value name is not case sensitive.)
5. Double-click the entry name you just typed to open the **Edit DWORD Value** dialog box.
6. Choose **Decimal** as the **Base** option.
7. Enter a value from 0 through 65535. The recommended value is 50.
8. Click **OK**.
9. Click **File**, and then click **Exit** to close the registry editor.

Adjusting the priority of the domain controller also reduces the number of client referrals. However, rather than reducing it proportionally to the other domain controllers, changing the priority causes DNS to stop referring all clients to this domain controller unless all domain controllers with a lower priority setting are unavailable.

**Note**

- A lower value entered for **LdapSrvPriority** indicates a higher priority. A domain controller with an **LdapSrvPriority** setting of 100 has a lower priority than a domain controller with a setting of 10. Therefore, clients attempt to use the domain controller with the setting of 10 first.

**To change the priority for DNS SRV records by using the registry**

1. In the **Run** dialog box, type **regedit**, and press ENTER.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters
3. Click **Edit**, click **New**, and then click **DWORD value**.
4. For the new entry name, type **LdapSrvPriority**, and press ENTER.
5. Double-click the entry name that you just typed to open the **Edit DWORD Value** dialog box.
6. Choose **Decimal** as the **Base** option.
7. Enter a value from 0 through 65535. The recommended value is 200.
8. Click **OK**.
9. Click **File**, and then click **Exit** to close the registry editor.

For more information about adjusting the weight or the priority of the PDC emulator, see the Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search under "Administration and Configuration Guides" and download the *Active Directory Operations Guide*.

# Additional Resources for Deploying Windows Server 2003 Regional Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing the Active Directory Logical Structure" in this book.
- "Designing the Site Topology" in this book.
- "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.
- The Active Directory Branch Office Planning Guide link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for a complete guide to information involving Active Directory branch office implementations.
- The Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search under "Administration and Configuration Guides" and download the *Active Directory Operations Guide*.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Active Directory" in Help and Support Center for Windows Server 2003.
- "Windows Support Tools" under "Tools" in Help and Support Center for Windows Server 2003.

**Related Job Aids**

- "Domain Controller Configuration" (DSSDFR_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Controller Configuration" on the Web at http://www.microsoft.com/reskit).

# Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Upgrading your domains from Microsoft® Windows NT® 4.0 to Windows® Server 2003 Active Directory® directory service enables your organization to improve the security and scalability of your network infrastructure while reducing administrative overhead. As an alternative to restructuring Windows NT 4.0 domains, the in-place upgrade is an efficient, time-saving process that minimizes the effect on the Windows NT 4.0 production environment.

**In This Chapter**


Overview of Upgrading Windows NT 4.0 Domains

Collecting Design Information

Complete Pre-Upgrade Tasks

Upgrading Domains from Windows NT 4.0 to Windows Server 2003 Active Directory

Complete Post-Upgrade Tasks

Additional Resources for Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory


**Related Information**

- For more information about restructuring domains when upgrading from Windows NT 4.0 to Windows Server 2003, see "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.
- For more information about the Active Directory logical structure, see "Designing the Active Directory Logical Structure" in this book.
- For more information about Windows Server 2003 Active Directory Functional Levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.
- For more information about Active Directory site topology, see "Designing the Site Topology" in this book.

# Overview of Upgrading Windows NT 4.0 Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Upgrading your Windows NT 4.0 domains to the Microsoft® Windows® Server 2003, Standard Edition and Windows® Server 2003, Enterprise Edition operating systems enables you to simplify and reduce network administration. Windows Server 2003 Active Directory integrates with other applications and services and allows you to delegate administrative responsibility at the appropriate level when you have multiple organizations existing in a single domain structure.

When you upgrade your Windows NT 4.0 domains to Windows Server 2003 Active Directory, you improve scalability because Active Directory domains can scale to meet the needs of your organization. You also gain new capabilities by using Group Policy, and you gain more flexibility for business units.

In addition, performing an in-place upgrade of Windows NT 4.0 domains to Windows Server 2003 Active Directory has no adverse effect on your Windows NT 4.0 production environment. There are fewer administrative complexities than with restructuring your environment, such as maintaining access to shared directories, files, and printers. Groups and group memberships are retained. You do not need to migrate local profiles, and you retain the existing passwords and profiles for domain users.

Before planning and implementing Windows NT 4.0 in-place upgrades, ensure that your organization has already:

- Designed the Active Directory logical structure of the forest and Domain Name System (DNS) for your Active Directory environment.
- Designed a site topology to efficiently locate domain controllers.
- Deployed a Windows Server 2003 forest root domain if that is the upgrade path that your organization has decided on. For more information about the paths for in-place upgrading a Windows NT 4.0 environment, see "Upgrading Domains from Windows NT 4.0 to Windows Server 2003 Active Directory" later in this chapter.

After completing the in-place upgrade process, you can perform an in-place upgrade for any remaining Windows NT domains or restructure them into your new Windows Server 2003 forest. For information about restructuring Windows NT 4.0 domains to a Windows Server 2003 forest, see "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.

**Note**

- For a list of the job aids that are available to assist you in upgrading your Windows NT 4.0 domains to Windows Server 2003 Active Directory, see "Additional Resources for Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" later in this chapter.

# Process for Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Upgrading your Windows NT 4.0 domains to Windows Server 2003 Active Directory involves first completing the necessary preparation tasks and then following the steps to complete the upgrade. Figure 8.1 shows the process for upgrading Windows NT 4.0 domains to Windows Server 2003 Active Directory.

**Figure 8.1 Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory**



# Background Information for Upgrading to Windows Server 2003 Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin the Windows NT 4.0 in-place domain upgrade, become familiar with some important issues that affect the upgrade process.

**PDC Offline Operations**
During the process of upgrading the operating system on the primary domain controller (PDC) from Windows NT 4.0 to Windows Server 2003 and installing Active Directory, client operations such as logon and resource access will continue to function because these services are provided by backup domain controllers. However, because the PDC will be offline during most phases of the upgrade process, typically between one and three hours, operations that require data to be written to the domain will not succeed. For example, users will not be able to change their passwords and administrators will not be able to create, delete, or unlock user accounts. Administrative tools, such as User Manager for Domains or Server Manager, can be used only in read-only mode on backup domain controllers in the domain. In addition, you will not be able to create new objects, such as users and groups, while the PDC is offline.

**Full Synchronization of the Local Security Authority Database**
After upgrading a Windows NT 4.0 PDC, or after transferring the PDC role to another domain controller, the LSA will perform a single full synchronization of all objects in the database. This synchronization causes events to be logged in Event Viewer; specifically, Event Viewer in Windows Server 2003 will log Event ID 5713 and Event Viewer in Windows NT 4.0 will log Event ID 5717. However, the LSA database contains relatively few objects and the full synchronization does not affect network performance.

Do not confuse the full synchronization of the LSA database with a backup domain controller (BDC) full synchronization. A BDC full synchronization typically happens when too many changes occur on a PDC before the PDC can replicate the changes to a BDC. The number of objects that are replicated during a BDC full synchronization and the amount of network traffic that is generated depends on the number of users, groups, and workstations in the domain.

**Domain Users and Client Workstation Operating Systems**
When Microsoft® Windows® 2000, Microsoft® Windows® XP, and Windows Server 2003 clients attempt to authenticate with a domain controller, they first retrieve a list of domain controllers from either DNS or WINS, and will then authenticate with the first domain controller that responds to their authentication request. The first domain controller to respond is usually a domain controller located closest to the client. The client and the domain controller will then negotiate which authentication protocol to use.

When Windows 2000, Windows XP, and Windows Server 2003 clients are members of a Windows NT 4.0 domain, they will only use the NTLM protocol to authenticate because that is the only authentication protocol supported by Windows NT 4.0. Windows 2000 and Windows Server 2003 domain controllers are capable of using either the NTLM or the more secure Kerberos authentication protocol.

When performing an in-place upgrade of a Windows NT 4.0 domain to Windows Server 2003, the first domain controller upgraded is the Windows NT 4.0 PDC. If clients in the domain running Windows 2000, Windows XP, and Windows Server 2003 select the new Active Directory domain controller for authentication, the negotiation of the authentication protocol will reveal that there are now domain controllers in the domain that support the Kerberos protocol. These clients will then upgrade their secure channel to exclusively use the Kerberos protocol for authentication requests and will no longer attempt to authenticate using the NTLM protocol, potentially causing the new Active Directory domain controller to become overloaded with authentication requests.

To prevent Windows Server 2003–based domain controllers from being overloaded with authentication requests, configure each Windows Server 2003–based domain controller to emulate a Windows NT 4.0–based domain controller during the upgrade process. Configuring a newly upgraded Windows Server 2003–based domain controller to emulate a Windows NT 4.0–based domain controller by using the NT4Emulator registry entry shields the new domain controller from getting too many authentication requests from Active Directory clients. Shielding the Active Directory domain controller takes place before the operating system is upgraded to Windows Server 2003 to prevent clients running Windows 2000, Windows XP, and Windows Server 2003 from ever establishing exclusive communications with a Windows Server 2003–based domain controller.

When upgrading additional Windows NT 4.0–based domain controllers after the PDC has been configured to emulate a

Windows NT 4.0–based domain controller, you must remember to configure the computer you are upgrading with the NeutralizeNT4Emulator registry entry. This is so that the additional domain controller will recognize the upgraded PDC that is emulating a Windows NT 4.0–based domain controller as an Active Directory domain controller. If the computer is not configured to neutralize emulation, you will not be able to install Active Directory because the additional domain controller will not be able to authenticate to an Active Directory domain controller.

For each site in which clients are running Windows 2000, Windows XP, and Windows Server 2003, ensure that you have enough Windows Server 2003–based domain controllers deployed in that site before removing Windows NT 4.0 emulation.

For more information about emulating Windows NT 4.0–based domain controllers, see "Configure Protection Against Domain Controller Overload on Additional Domain Controllers" later in this chapter.

For more information about domain controller placement, see "Designing the Site Topology" in this book. For more information about domain controller capacity planning and determining the number of domain controllers needed in each site to service Active Directory clients, see "Planning Domain Controller Capacity" in this book.

**Service Compatibility**

In Windows NT 4.0 and earlier server operating systems, services running in the context of the Local System account communicate with other services over the network by using *null sessions* (a session in which a user name or password is not provided). In Windows 2000 and later operating systems, services running in the context of the Local System account on the local computer use the local computer account to authenticate to other servers. By default, Active Directory does not accept null session queries.

Of all the services that run in the context of the Local System account, Remote Access Services (RAS) is the most prominent. You cannot use null sessions to access network resources by using NTLM authentication unless the remote computer allows access with null credentials.

In an Active Directory environment containing both Windows NT 4.0–based and Windows Server 2003–based domain controllers, a member server that is running Windows NT 4.0 and is configured as a RAS server cannot retrieve information from a Windows Server 2003–based domain controller. For example, if a caller tries to dial into your network and accesses a Windows NT 4.0 member server that is configured as a RAS server, the RAS server must query a domain controller first to verify whether the caller has permission to dial into the network. Therefore, RAS operates correctly only if the domain controller responding to the RAS authentication request is a Windows NT 4.0–based BDC or the Active Directory domain has been configured to allow resources to be accessed by using null credentials. By upgrading the operating system on Windows NT 4.0 member servers that are configured as RAS servers to Windows Server 2003, you ensure that RAS callers are successfully authenticated by a Windows Server 2003 Active Directory–based domain controller.

The recommended solution is to upgrade the RAS servers to Windows Server 2003. However, if this cannot be done, the alternatives are:

- While installing Active Directory on the upgraded Windows NT 4.0 PDC, on the Permissions page of the Active Directory Installation wizard, select **Permissions compatible with pre-Windows 2000 Server operating systems**.

  – Or –

  Add the Everyone group and the Anonymous Logon group to the Pre-Windows 2000 Compatible Access built-in group by using Active Directory Users and Computers or the command line.

**To add the Everyone group to the Pre-Windows 2000 Compatible Access Group by using the command line**

- At the command line, type:

```
net localgroup "Pre-Windows 2000 Compatible Access" Everyone /add
```

**To add the Anonymous Logon group to the Pre-Windows 2000 Compatible Access Group by using the command line**

- At the command line, type:

```
net localgroup "Pre-Windows 2000 Compatible Access" “Anonymous Logon” /add
```

**Note**

- After this update to the Pre-Windows 2000 Compatible Access group replicates, you must restart the Server service on all domain controllers.

Both of these methods combined allow null sessions to read information out of the directory. After you upgrade all RAS servers, and when you no longer need backward compatibility with operating systems earlier than Windows 2000, remove the Everyone group and the Anonymous Logon group from the Pre-Windows 2000 Compatible Access built-in group. For more information about removing the Everyone group and the Anonymous Logon group from the Pre-Windows 2000 Compatible Access group, see "Eliminate Anonymous Connections to Domain Controllers" later in this chapter.

**LAN Manager Replication Service and the File Replication Service**
In Windows NT 4.0, the LAN Manager Replication (LMRepl) service provides single master replication of logon scripts and other database information located in the NETLOGON share on a Windows NT 4.0–based domain controller that is designated as an export server to all other Windows NT 4.0–based domain controllers in the domain. LMRepl can be configured only on Windows NT 4.0–based domain controllers. In Windows 2000 and Windows Server 2003, logon scripts and profile information are stored in the NETLOGON shared folder (which contains policies and scripts for non-Active Directory clients) and the SYSVOL shared folder (which contains Group Policy files and scripts for Active Directory clients). The *File Replication Service (FRS)*, a multimaster replication engine that runs automatically on all Windows Server 2003–based domain controllers, replaces the LMRepl service and replicates the NETLOGON and SYSVOL shared folders between domain controllers in a Windows Server 2003 domain.

During the in-place domain upgrade process, your environment includes Windows NT 4.0–based BDCs operating with Windows Server 2003–based domain controllers. FRS and LMRepl are not backward compatible. Therefore, to provide support for the LMRepl service in the Active Directory environment, you need to create a bridge between LMRepl and FRS to replicated new files created in the NETLOGON folder on Windows Server 2003 domain controllers to the Windows NT 4.0 export server. The bridge is created by using the Lbridge.cmd script and the Robocopy.exe tool so that both services can operate autonomously. Do this by selecting one Windows Server 2003–based domain controller to copy the SYSVOL shared folder to the Windows NT 4.0 export directory of the Windows NT 4.0 export server. You can use a regularly scheduled script to copy the shared folder. For more information about creating this script, see "Synchronize File Replication Services" later in this chapter.

**Security Policy Considerations when Upgrading from Windows NT 4.0 to Windows Server 2003**
Server message block (SMB) packet signing and secure channel signing are security policies that are enabled by default on Windows Server 2003–based domain controllers. To allow clients running earlier versions of Windows to communicate with domain controllers running Windows Server 2003, you might need to temporarily disable these security policies during the upgrade process.

**SMB packet signing**
SMB packet signing is a security mechanism that protects the data integrity of SMB traffic between client computers and servers, and prevents man-in-the-middle attacks by providing a form of mutual authentication. This is done by placing a digital security signature into each SMB packet, which is then verified by the receiving party. Server-side SMB signing is required by default on Windows Server 2003–based domain controllers, which means that all clients are required to have SMB packet signing enabled.

Clients running Windows NT 4.0 with Service Pack 2 or earlier, and clients running Microsoft® Windows® 95 without the Directory Service Client Pack, do not support SMB packet signing. These clients will not be able to authenticate to a Windows Server 2003–based domain controller. To ensure successful authentication, upgrade these clients to a later version of the operating system or Service Pack. However, if you cannot upgrade your clients, you can allow them to be authenticated by configuring SMB packet signing on all Windows Server 2003–based domain controllers so that SMB packet signing is preferred but not required.

For more information about SMB packet signing, see "Microsoft network server: Digitally sign communications (always)" on the Microsoft Web site.

For more information about configuring SMB packet signing on Windows Server 2003–based domain controllers, see "Modify Security Policies" later in this chapter.

For more information about the Directory Services Client Pack, see article 323466, "Availability of the Directory Services Client Update for Windows 95 and Windows 98" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Secure channel signing and encryption**
When a computer becomes a member of a domain, a computer account is created. Each time the computer starts, it uses the computer account password to create a secure channel with a domain controller for its domain. This secure channel is used to ensure secure communications between a domain member and a domain controller for its domain. Secure channel signing is required by default on Windows Server 2003–based domain controllers, which means that all clients must enable secure

channel signing and encryption.

Clients running Windows NT 4.0 with Service Pack 3 or earlier installed do not support secure channel signing. These clients will not be able to establish communications with a Windows Server 2003–based domain controller. To ensure successful communication, upgrade these clients to a later version of the operating system or Service Pack. However, if you cannot upgrade your clients, you must disable secure channel signing on all Windows Server 2003–based domain controllers so that the traffic passing through the secure channel is not required to be signed or encrypted.

**Note**

- Unlike SMB packet signing, secure channel signing does not affect Windows 95 clients.

For more information about secure channel signing, see "Domain member: Digitally encrypt or sign secure channel data (always)" in Help and Support Center for Windows Server 2003.

For more information about configuring secure channel signing on Windows Server 2003–based domain controllers, see "Modify Security Policies" later in this chapter.

# Collecting Design Information

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In preparation for deployment, the forest owner in your organization is responsible for working with the deployment team to acquire the following information:

- Documentation of your current Windows NT 4.0 environment.
- Names of the Windows NT 4.0 domains that will be upgraded and the order in which to upgrade them.
- Supported operating system upgrade paths for your Windows NT 4.0–based domain controllers.

Information such as domain diagrams, network services, and trust relationships might have been documented as part of the design process, and collecting it will be a matter of querying the design team. However, information such as the existing network and hardware configuration of each domain controller might have to be collected or documented by the forest owner during the deployment phase of the project.

In addition, the forest owner is responsible for developing a test plan and for developing a recovery plan in the event that the deployment does not complete successfully.

Figure 8.2 shows the steps involved in collecting the design information that will be used to upgrade Windows NT 4.0 domains to Windows Server 2003 Active Directory.

**Figure 8.2 Collecting Design Information**



# Document the Existing Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before upgrading a Windows NT 4.0 domain to Windows Server 2003 Active Directory, document the existing Windows NT 4.0 domain structure.

Create a diagram that includes the following information:

- The names of all account and resource domains.
- The inbound and outbound trust relationships that each domain shares.

If documentation already exists for your domain, review the existing documentation for accuracy and clarity.

Figure 8.3 shows an example of the existing Windows NT 4.0 domain structure for a fictitious company, Trey Research.

**Figure 8.3 Example of a Windows NT 4.0 Domain Diagram**




In addition to documenting the existing domain structure, document the following:

- The domain controllers and the services that each provides in the domain.
- The existing hardware configuration on all domain controllers in the domain.
- The existing network configuration, including IP address and network adapter information for each domain controller.
- The current domain controller assignments and the role that you plan to assign to each domain controller after the in-place domain upgrade.

# Document Domain Controllers and Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Identify and document the domain controllers in the existing Windows NT 4.0 domain. Include in your documentation the role that each domain controller assumes in the domain and the services that each domain controller provides. Identify domain controllers that provide Remote Access Service and the LAN Manager Replication (LMRepl) service, because upgrading to Windows Server 2003 Active Directory affects these services.

For a worksheet to assist you in documenting domain controllers and services see "Windows NT 4.0 Domain Controllers and Services" (DSSUPNT_1.doc) or "Windows NT 4.0 Domain Controller Documentation" (DSSUPNT_5.xls) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Windows NT 4.0 Domain Controllers and Services" or "Windows NT 4.0 Domain Controller Documentation" on the Web at http://www.microsoft.com/reskit). "Windows NT 4.0 Domain Controller Documentation" is a master worksheet combining the information from all four individual worksheets in this section.

**Example: Documenting Windows NT 4.0 Domain Controllers and Services**
Trey Research has a Windows NT 4.0 account domain that includes nine domain controllers running Windows NT 4.0. Because the resource domains hold all of the application servers, the account domain does not include member servers. The PDC, SEA-EAST-DC01, is also a Windows Internet Name Service (WINS) server, as are two BDCs, BOS-EAST-DC01 and BOS-EAST-DC02.

Trey Research documented the domain controllers and services in their Windows NT 4.0 domain, as shown in Figure 8.4.

**Figure 8.4 Example of Windows NT 4.0 Domain Controllers and Services Worksheet**

| Windows NT 4.0 Domain Controllers and Services | | | |
|---|---|---|---|
| Prepared By | Forest Owner | Date | 1/2/03 |
| Domain Name | EAST | | |

| Domain Controller Name | Domain Controller Role | Services |
|---|---|---|
| SEA-EAST-DC01 | PDC | WINS, LMRepl export server |
| SEA-EAST-DC02 | BDC | LMRepl import server, RAS server |
| WDC-EAST-DC01 | BDC | LMRepl import server |
| WDC-EAST-DC02 | BDC | LMRepl import server |
| NYC-EAST-DC01 | BDC | LMRepl import server, RAS server |
| NYC-EAST-DC02 | BDC | LMRepl import server, RAS server |
| NYC-EAST-DC03 | BDC | LMRepl import server |
| BOS-EAST-DC01 | BDC | WINS, LMRepl import server |
| BOS-EAST-DC02 | BDC | WINS, LMRepl import server |

For more information about the effect of upgrading to Windows Server 2003 Active Directory on the RAS service and the LMRepl service, see "Background Information for Upgrading to Windows Server 2003 Active Directory" earlier in this chapter.

# Document the Existing Hardware Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review and document the existing hardware configuration of each domain controller that you plan to upgrade to Windows Server 2003. Use this information to identify the domain controllers in your environment that you can upgrade to Windows Server 2003 and the domain controllers that do not meet the hardware requirements for Windows Server 2003. Retain at least one domain controller that does not meet Windows Server 2003 hardware requirements to serve as a rollback server in the event that you must roll back your deployment.

If the PDC does not meet the hardware requirements, you can transfer the PDC role to a backup domain controller (BDC) that does meet the hardware requirements and upgrade it. If none of your Windows NT 4.0 domain controllers meet Windows Server 2003 hardware requirements, install a Windows NT 4.0 BDC on a computer that does meet the hardware requirements for a domain controller that is running Windows Server 2003 and transfer the PDC role to it.

You can also add a Windows Server 2003–based member server to a Windows NT 4.0 domain at any time before you upgrade to Windows Server 2003 Active Directory because Windows Server 2003–based member servers can operate within a Windows NT 4.0 environment. You can install Active Directory on the member server after you upgrade the PDC.

For more information about the hardware requirements of domain controllers in a Windows Server 2003 domain, see "Planning Domain Controller Capacity" in this book. To determine whether your hardware configuration is compatible with Windows Server 2003, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For a worksheet to assist you in documenting your existing domain controller hardware configuration, see "Windows NT 4.0 Hardware Configuration" (DSSUPNT_2.doc) or "Windows NT 4.0 Domain Controller Documentation" (DSSUPNT_5.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows NT 4.0 Hardware Configuration" or "Windows NT 4.0 Domain Controller Documentation" on the Web at http://www.microsft.com/reskit).

**Example: Documenting the Windows NT 4.0 Hardware Configuration**

Figure 8.5 shows an example of a Hardware Configuration worksheet for the Windows NT 4.0–based domain controllers in the EAST domain for Trey Research.

**Figure 8.5 Example of a Windows NT 4.0 Hardware Configuration Worksheet**

| Windows NT 4.0 Hardware Configuration | | | |
|---|---|---|---|
| Prepared By | Forest Owner | Date | 1/2/03 |
| Domain Name | EAST | | |

| DC Name | Operating System | Hardware Information | | | Upgradeable to Windows Server 2003 DC? |
|---|---|---|---|---|---|
| | | RAM | CPU | Disk Space | |
| SEA-EAST-DC01 | Windows NT 4.0 | 2 GB | 1 x 850 MHz | 20 GB | Yes |
| SEA-EAST-DC02 | Windows NT 4.0 | 2 GB | 1 x 850 MHz | 20 GB | Yes |
| WDC-EAST-DC01 | Windows NT 4.0 | 1 GB | 1 x 850 MHz | 20 GB | Yes |
| WDC-EAST-DC02 | Windows NT 4.0 | 256 MB | 1 x 850 MHz | 10 GB | No |
| NYC-EAST-DC01 | Windows NT 4.0 | 2 GB | 4 x 850 MHz | 40 GB | Yes |
| NYC-EAST-DC02 | Windows NT 4.0 | 2 GB | 4 x 850 MHz | 40 GB | Yes |
| NYC-EAST-DC03 | Windows NT 4.0 | 2 GB | 4 x 850 MHz | 40 GB | Yes |
| BOS-EAST-DC01 | Windows NT 4.0 | 1 GB | 1 x 850 MHz | 20 GB | Yes |
| BOS-EAST-DC02 | Windows NT 4.0 | 256 MB | 1 x 850 MHz | 8 GB | No |

Domain controllers BOS-EAST-DC02 and WDC-EAST-DC02 do not meet the minimum memory requirements for a Windows Server 2003–based domain controller. Therefore, Trey Research has determined that BOS-EAST-DC02 will be used as the Windows NT 4.0 rollback server if a problem occurs during the in-place upgrade process and WDC-EAST-DC02 will be assigned as a member server in the Windows Server 2003 forest. All other Windows NT 4.0–based domain controllers are capable of supporting Windows Server 2003 Active Directory.

# Document the Existing Network Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Document the existing network configuration for your Windows NT 4.0 domain. Some network adapter drivers that are included with earlier versions of the operating system are not distributed with Windows Server 2003. If you attempt to upgrade a Windows NT 4.0–based domain controller to Windows Server 2003 and a network adapter is installed for which a driver is not provided, your network information might be lost or detected incorrectly during the upgrade.

**Note**

- You can install device drivers that are not included on the Windows Server 2003 operating system CD from the vendor's Web site.

Create a network configuration table listing the type of network adapter that each domain controller uses. Also include the TCP/IP configuration information for each domain controller, including IP address, subnet mask, and default gateway. You can run the ipconfig command at the command line to determine IP address, subnet mask, and default gateway. For more information about the ipconfig command, type **ipconfig /?** at the command line.

To determine whether the network card is supported by Windows Server 2003, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For a worksheet to assist you in documenting your existing Windows NT 4.0 network configuration, see "Windows NT 4.0 Network Configuration" (DSSUPNT_3.doc) or "Windows NT 4.0 Domain Controller Documentation" (DSSUPNT_5.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows NT 4.0 Network Configuration" or "Windows NT 4.0 Domain Controller Documentation" on the Web at http://www.microsoft.com/reskit).

Figure 8.6 shows an example of a network configuration worksheet for the EAST domain for Trey Research.

**Figure 8.6 Example of a Windows NT 4.0 Network Configuration Worksheet**

# Windows NT 4.0 Network Configuration

| Prepared By | Forest Owner | Date | 1/2/03 |
|---|---|---|---|

| Domain Name | EAST |
|---|---|

| DC Name | Network Adapter | IP Address | Subnet Mask | Default Gateway |
|---|---|---|---|---|
| SEA-EAST-DC01 | NETGEAR FA310TX Fast Ethernet Adapter | 172.16.16.10 | 255.255.252.0 | 172.16.16.1 |
| SEA-EAST-DC02 | Intel PRO/100 VE Desktop Connection | 172.16.16.11 | 255.255.252.0 | 172.16.16.1 |
| WDC-EAST-DC01 | IBM Netfinity 10/100 Ethernet Adapter | 172.16.116.2 | 255.255.252.0 | 172.16.116.1 |
| WDC-EAST-DC02 | IBM Netfinity 10/100 Ethernet Adapter | 172.16.116.3 | 255.255.252.0 | 172.16.116.1 |
| NYC-EAST-DC01 | 3Com EtherLink III LAN PC Card (3C589) (Ethernet) | 172.16.28.2 | 255.255.252.0 | 172.16.28.1 |
| NYC-EAST-DC02 | IBM Netfinity 10/100 Ethernet Adapter | 172.16.28.3 | 255.255.252.0 | 172.16.28.1 |
| NYC-EAST-DC03 | IBM Netfinity 10/100 Ethernet Adapter | 172.16.28.4 | 255.255.255.0 | 172.16.28.1 |
| BOS-EAST-DC01 | Intel PRO/100 VE Desktop Connection | 172.16.36.2 | 255.255.252.0 | 172.16.36.1 |
| BOS-EAST-DC02 | Kingston EtherRx KNE111TX PCI Fast Ethernet Adapter | 172.16.36.3 | 255.255.252.0 | 172.16.36.1 |

# Document Domain Controller Role Assignments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As part of your in-place domain upgrade plan, assign the existing Windows NT 4.0–based domain controllers roles that they will assume in the Windows Server 2003 domain after the upgrade is complete. Assign one of the following three roles to Windows NT 4.0–based domain controllers in a Windows Server 2003 domain:

- **Windows Server 2003–based domain controller.** Assign the role of Windows Server 2003–based domain controller to any Windows NT 4.0 PDCs and other Windows NT 4.0–based domain controllers that meet the appropriate hardware and software requirements.

- **Rollback server.** Assign the role of rollback server in the Windows Server 2003 domain to a Windows NT 4.0 BDC that does not meet the Windows Server 2003 domain controller hardware requirements.

- **Windows Server 2003–based member server.** Assign the role of member server in the Windows Server 2003 domain to a Windows NT 4.0–based BDC that does not meet the Windows Server 2003 domain controller hardware requirements.

For more information about the software and hardware requirements for Windows Server 2003–based domain controllers, see "Determine Supported Operating System Upgrades" later in this chapter and "Document the Existing Hardware Configuration" earlier in this chapter.

Create a domain controller assignment table that outlines the roles that you plan to assign to your Windows NT 4.0–based domain controllers in the Windows Server 2003 domain. In this table, list the Windows NT 4.0–based domain controllers in your domain, indicate whether they meet the hardware requirements for Windows Server 2003, and list the role for each domain controller before and after you upgrade the domain, as shown in Figure 8.7.

For a worksheet to assist you in documenting Windows NT 4.0–based domain controller roles, see "Windows NT 4.0 Domain Controller Role Assignment" (DSSUPNT_4.doc) or "Windows NT 4.0 Domain Controller Documentation" (DSSUPNT_5.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows NT 4.0 Domain Controller Role Assignment" or "Windows NT 4.0 Domain Controller Documentation" on the Web at http://www.microsoft.com/reskit).

**Figure 8.7 Example of a Windows NT 4.0 Domain Controller Role Assignment Worksheet**

| Windows NT 4.0 Domain Controller Role Assignment | | | |
|---|---|---|---|
| Prepared By | Forest Owner | Date | 1/2/03 |
| Domain Name | EAST | | |

| DC Name | Meets hardware requirements? | Role before upgrade | Role after upgrade |
|---|---|---|---|
| SEA-EAST-DC01 | Yes | Windows NT 4.0 PDC | Windows Server 2003 domain controller |
| SEA-EAST-DC02 | Yes | Windows NT 4.0 BDC | Windows Server 2003 domain controller |
| WDC-EAST-DC01 | Yes | Windows NT 4.0 BDC | Windows Server 2003 domain controller |
| WDC-EAST-DC02 | No | Windows NT 4.0 BDC | Windows Server 2003 member server |
| NYC-EAST-DC01 | Yes | Windows NT 4.0 BDC | Windows Server 2003 domain controller |
| NYC-EAST-DC02 | Yes | Windows NT 4.0 BDC | Windows Server 2003 domain controller |
| NYC-EAST-DC03 | Yes | Windows NT 4.0 BDC | Windows Server 2003 domain controller |
| BOS-EAST-DC01 | Yes | Windows NT 4.0 BDC | Windows Server 2003 domain controller |
| BOS-EAST-DC02 | No | Windows NT 4.0 BDC | Windows NT 4.0 Rollback Server |

# Determine the Domain Upgrade Order

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin the in-place domain upgrade process, determine the order in which you plan to upgrade your Windows NT 4.0 domains. Because account domains generally contain more objects than resource domains, upgrade your account domains before upgrading your resource domains. This allows your organization to take advantage of Windows Server 2003 security and administration features early in the upgrade process.

The order in which you upgrade account domains in your organization can affect the efficiency of your in-place domain upgrade process. Use the following guidelines to determine the order in which to upgrade multiple account domains:

- Upgrade domains that will become targets for restructuring first. After upgrading these domains, you can restructure remaining domain objects into the restructuring target. Target domains must be set at the Windows 2000 native domain functional level before restructuring objects into them.

- Upgrade domains over which you have direct control and to which you have easy access. This allows convenient access to these domains in the event that you must roll back your deployment if the upgrade does not go as planned.

For more information about restructuring Windows NT domains, see "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.

# Determine Supported Operating System Upgrades

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Identify the Windows NT 4.0 platforms that are running in your environment and determine whether an operating system upgrade to Windows Server 2003 is supported, or whether you must perform a clean operating system installation.

Table 8.1 lists the Windows NT 4.0 platforms and indicates which platforms you can upgrade directly to each edition of Windows Server 2003. You do not need to reinstall applications on platforms that you can upgrade directly to Windows Server 2003.

**Table 8.1 Supported Upgrade Paths to Windows Server 2003**

| Platform | Upgrade to Windows Server 2003, Standard Edition | Upgrade to Windows Server 2003, Enterprise Edition | Upgrade to Windows Server 2003, Datacenter Edition |
|---|---|---|---|
| Windows NT 4.0 Server, Standard Edition | ● | ● | |
| Windows NT 4.0 Terminal Server | ● | ● | |
| Windows NT 4.0 Server, Enterprise Edition | | ● | |

**Important**

- All versions of Windows NT 4.0 must have Service Pack 5 or later installed before upgrading to Windows Server 2003.

If you have computers in your environment that are running operating systems that you cannot upgrade directly to a version of Windows Server 2003, such as Windows NT 3.51, you must do one of the following:

- If you need to retain applications that are located on those computers, upgrade the computers to run an operating system that you can upgrade to Windows Server 2003 after verifying that those applications will function on and are supported by Windows Server 2003.

- If you do not need to retain applications that are located on those computers, perform a clean installation of Windows Server 2003 on those computers.

# Develop a Test Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Develop a plan for testing your in-place domain upgrade procedures throughout the in-place domain upgrade process to ensure that they have completed successfully and to determine whether the process of upgrading Windows NT 4.0 domains to Windows Server 2003 Active Directory was successful.

Table 8.2 lists the Active Directory configurations that you must test and the tools that you can use to test each configuration. For more information about the options that are available for these tools, see "Active Directory support tools" in Windows Server 2003 Help on the Microsoft Web site. For more information about specific configuration and functionality tests that you can perform before and after the Active Directory installation, see the Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search under "Administration and Configuration Guides" and download the *Active Directory Operations Guide*.

**Table 8.2 Active Directory Configuration Test Components**

| Configuration | Tool | Purpose |
|---|---|---|
| Active Directory service | Dcdiag.exe | Tests for successful Active Directory connectivity and functionality. Confirms that the domain controller has passed the diagnostic tests (such as connectivity and replicated objects). Each test must return a "passed" result. |
| | Netdiag.exe | Diagnoses networking and connectivity problems by performing a series of tests to determine the state of your network client and whether it is functional. |
| Active Directory replication | Repadmin.exe /replsum | Returns all replication events taking place between the forest root domain and other Active Directory domain controllers. This must return a successful replication event with all inbound and outbound replication partners. |
| BDC replication status | Nltest.exe /bdc_query:*domainname* | Shows connection status for all the BDCs. This must show "status = success" for each domain controller within the domain. |

After you confirm that the Active Directory configuration is correct, you need to verify that Active Directory is functioning correctly. Table 8.3 lists the Active Directory functions that you need to test and the methods that you can use to perform the tests.

**Table 8.3 Active Directory Functionality Test Components**

| Function | Test | Method |
|---|---|---|
| Trust relationships | Verify the transitive trusts with the parent domain and the one-way trusts with Windows NT 4.0 domains. | Use the verify feature in Active Directory Domains and Trusts on the upgraded PDC to validate the trust relationships that are in place. |
| New user creation | Create a new user on the Windows Server 2003–based domain controller. | Log on with administrator credentials and use Active Directory Users and Computers to verify that the new user was created successfully. |
| New user object replication | After replication to BDCs takes place, determine whether new user is replicated to BDCs. | 1. Type **Net User** at a command prompt on a Windows NT 4.0–based domain controller, and then verify that the new user account exists.<br>2. Modify a property of an existing user and verify that the modified property replicates with the user. |

| | | |
|---|---|---|
| Successful logon request | Verify that users can log on successfully. | 1. Disconnect the Windows Server 2003–based domain controller to confirm that the Windows NT 4.0–based domain controller is validating the user logon request.<br><br>2. Verify that you can log on successfully by using the new user account credentials from each client computer.<br><br>3. Verify that all client operating systems in the upgraded domain and the domains that it trusts can log on successfully.<br><br>4. Repeat step number two over trust relationships where the trusting domain controller has a secure channel with the Windows NT 4.0–based and Windows Server 2003–based domain controllers in the trusted domain. |
| Successful resource access | Verify that the user can access important resources. | 1. Access e-mail resources.<br><br>2. Access roaming profiles.<br><br>3. Access printers.<br><br>4. Resource permissions belonging to the user and a group. |

# Develop a Recovery Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Create a recovery plan for use if the in-place domain upgrade process does not go as planned. Select a Windows NT 4.0 BDC to be used as a rollback server. Synchronize the BDC with the PDC and take the rollback server offline in the event that it must be promoted to a PDC to restore the domain to its original state. Although you are unlikely to need the offline domain controller, it is recommended that you take one offline as a precautionary step if the Security Accounts Manager (SAM) account database on all domain controllers becomes corrupt.

Include the following in your recovery plan:

- The steps needed for recovery. Be sure to provide clear instructions so that the deployment team can restore normal operations to the organization if necessary.

- The estimated time that can elapse before recovery must take place. When elements of the upgrade process test unsuccessfully, you might spend unanticipated amounts of time identifying and correcting errors. Establish clear guidelines for the time period after which the deployment team must restore operations for end users.

- Team review and sign-off. All members of the deployment team must sign off on the recovery plan. This signifies consensus about the recovery plan and reduces the chances that misunderstandings occur when the upgrade process does not go as planned.

**Restoring the Domain to its Original State**

If your in-place upgrade process fails, you can roll back a Windows Server 2003 Active Directory domain to its original state as a Windows NT 4.0 domain. There are two ways to roll back the deployment to its original state:

**Note**

- The first recovery method is preferred for restoring a domain to its original state. The second recovery method should only be used if the SAM database on all domain controllers becomes corrupt.

1. Remove (either by disconnecting the network cable or turning off) any Windows Server 2003–based domain controllers from the domain.

2. Promote a Windows NT 4.0 BDC to become the PDC.

3. Synchronize all Windows NT 4.0–based domain controllers.

4. Test Windows NT 4.0 server operations and domain validation.

5. Document the reasons for the unsuccessful domain upgrade and communicate them to your design team.

6. Restart the design phase for the in-place domain upgrade. Be sure to include steps to mitigate the factors that caused the first in-place domain upgrade to fail.

   – Or –

   If a failure occurs after performing the steps above, remove all Windows Server 2003–based domain controllers from the network and promote the Windows NT 4.0 BDC that has been designated as the rollback server to become the PDC.

7. Perform a full synchronization of all Windows NT 4.0 BDCs.

8. Test Windows NT 4.0 server operations and domain validation.

9. Document the reasons for the unsuccessful domain upgrade and communicate them to your design team.

10. Restart the design phase for the in-place domain upgrade. Be sure to include steps to mitigate the factors that caused the first in-place domain upgrade to fail.

**Important**

- You must take all Windows Server 2003–based domain controllers offline before you promote the rollback server to become the new PDC. If any Windows Server 2003–based domain controllers remain online in the domain, the promotion of the BDC to a PDC will not work.

# Complete Pre-Upgrade Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you create your plan for upgrading your Windows NT 4.0 domains to Windows Server 2003 Active Directory, you must complete the pre-upgrade tasks shown in Figure 8.8 before beginning the in-place upgrade process for your domain.

**Figure 8.8 Completing Pre-Upgrade Tasks**



# Relocate the LMRepl File Replication Service

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To maintain the replication of files in the NETLOGON shared folder from the Windows NT 4.0 export server to all other Windows NT 4.0 BDCs running the LMRepl replication engine during the in-place domain upgrade process, upgrade all servers that are hosting import directories before you upgrade the server that is hosting the export directory.

If the server hosting the export directory is the PDC, you can do one of the following:

- Promote a BDC that meets the Windows Server 2003 domain controller hardware requirements to become the new PDC and demote the existing PDC to serve as a BDC hosting the export server.

  – Or –

  Reconfigure the LMRepl export server on a BDC and remove it from the PDC.

To test the new configuration to ensure that LMRepl continues to work correctly, place an empty file on the export server and verify that the file is replicated to the import directories during replication. Next, delete the replicated file from the import directory, and then verify that the file is deleted during the next replication.

# Ensure Remote Access Service Compatibility

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To ensure remote access compatibility in a mixed Windows NT 4.0 and Windows Server 2003 environment, upgrade the operating system on all remote access servers in the domain to Windows Server 2003 before you begin the in-place domain upgrade process. If RAS is running on a domain controller, upgrade that domain controller early in the in-place domain upgrade process to minimize security risks.

# Enable the Windows NT 4.0 Environment Change Freeze

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you upgrade the PDC in your Windows NT 4.0 domain to Windows Server 2003 Active Directory, you must freeze the Windows NT 4.0 environment to ensure that no other domain changes occur until after the PDC is upgraded. Freeze the Windows NT 4.0 environment when:

- You have completed all of the updates to the Windows NT 4.0 domain and have replicated them to all domain controllers.
- You have synchronized a BDC and have taken it offline for recovery purposes.

When you freeze the Windows NT 4.0 environment, no additional domain changes can take place until you upgrade the Windows NT 4.0 PDC to Windows Server 2003. Communicate to all appropriate individuals that changes to the environment, such as password updates, will not be accepted after a specific date.

# Upgrading Domains from Windows NT 4.0 to Windows Server 2003 Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin the Windows NT 4.0 in-place upgrade process, determine the upgrade path that your Active Directory design specifies. The Active Directory design will specify one of two possible in-place upgrade paths:

- Upgrade to a regional domain in an existing forest.

  Before upgrading a Windows NT 4.0 domain and joining an existing forest as a regional domain, you must first deploy a Windows Server 2003 forest root domain. Complete the planning and design phases of your Active Directory deployment and then complete the process for deploying the forest root domain. After the forest root domain is deployed, complete the in-place domain upgrade process by following the steps outlined in "Upgrade to a Regional Domain in an Existing Forest" later in this chapter. For more information about deploying the Windows Server 2003 forest root domain, see "Deploying the Windows Server 2003 Forest Root Domain" in this book.

  **Note**

  - If your organization already has a Windows 2000 or Windows Server 2003 Active Directory infrastructure in place, complete the in-place upgrade process by upgrading to a regional domain in an existing forest.

  To help illustrate the process for upgrading to a regional domain in an existing forest, sample data for Trey Research, is provided within the context of the tasks that must be performed.

- Upgrade to a single domain forest.

  To create a new single domain forest, complete the in-place domain upgrade process by following the steps outlined in "Upgrade to a Single Domain Forest" later in this chapter. To help illustrate the process for upgrading to a single domain forest, sample data for a fictitious company, Fabrikam, Inc, is provided within the context of the tasks that must be performed.

For more information about designing an Active Directory logical structure and determining what forest design model best suits your organization, see "Designing the Active Directory Logical Structure" in this book.

Figure 8.9 shows the two paths available for upgrading domains from Windows NT 4.0 to Windows Server 2003 Active Directory and additional tasks that all organizations must perform regardless of which option is specified by the Active Directory design. The additional tasks, including modifying security policies, synchronizing file replication services, recreating trusts, using DNS registration to decrease the workload on the PDC emulator, and upgrading additional domain controllers, are performed after the PDC is upgraded.

**Figure 8.9 Upgrading Domains from Windows NT 4.0 to Windows Server 2003 Active Directory**




After the in-place domain upgrade is complete, you can upgrade additional Windows NT 4.0 domains in-place or restructure the remaining Windows NT 4.0 domains into your Windows Server 2003 Active Directory environment. For more information about restructuring Windows NT 4 domains, see "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.

# Upgrade to a Regional Domain in an Existing Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To complete the process for upgrading to a regional domain in an existing forest, perform the following tasks:

1. Back up all domain data.
2. Enable the Windows Server 2003 interim forest functional level in the existing forest.
3. Delegate the DNS zone in the forest root domain.
4. Configure protection against domain controller overload.
5. Upgrade the operating system of the Windows NT 4.0 PDC.
6. Install Active Directory.
7. Perform post-upgrade tests.

# Back Up Your Domain Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Back up your Windows NT 4.0 domain data before you begin the upgrade. This task varies according to the operations and procedures that already exist in your environment. At minimum, complete the following steps:

- Back up the PDC.
- Back up the BDC that you designated as the rollback server.
    - Test all backup media to ensure that the data can be restored successfully.

**Important**

- Store backup media in a secure offsite location designated by and accessible to the deployment team before you begin the upgrade process.

# Enable the Windows Server 2003 Interim Forest Functional Level

Updated: August 3, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If all domain controllers in the existing forest are running Windows Server 2003, the functional level of the forest is set at Windows 2000, and the functional level of the forest root domain is set at Windows 2000 mixed, you can raise the forest functional level to Windows Server 2003 interim. Raising the forest functional level to Windows Server 2003 interim is recommended in order to take advantage of the Windows Server 2003 Active Directory features available at that level. However, if you are considering adding Windows 2000–based domain controllers to your environment at any time, you can maintain the Windows 2000 forest functional level and still upgrade your Windows NT 4.0 domains.

Raise the forest functional level in the existing forest to Windows Server 2003 interim before upgrading the PDC and joining the existing forest during the Active Directory installation. By raising the forest functional level in the existing forest before you upgrade the PDC, any additional domains that you upgrade as regional domains will automatically join the Windows Server 2003 forest at the Windows Server 2003 interim domain functional level.

**Important**

- If you raise the forest and domain functional level to Windows Server 2003 interim, you cannot return to the Windows 2000 mixed domain functional level or to the Windows 2000 forest functional level. After you raise the functional level to Windows Server 2003 interim, the environment only supports Windows NT 4.0– and Windows Server 2003–based domain controllers. You can no longer add Windows 2000–based domain controllers into this environment.

You cannot use Active Directory administrative consoles to raise the forest functional level to Windows Server 2003 interim. Instead, use a Lightweight Directory Access Protocol (LDAP) application such as ADSI Edit or LDP in Windows Support Tools to edit the value of the *msDS-Behavior-Version* attribute. You must be a member of the Enterprise Admins group to raise the forest functional level, and you must do this on the domain controller that holds the schema master role.

**To raise the forest functional level to Windows Server 2003 interim by using ADSI Edit**

1. In ADSI Edit, expand the **Configuration** partition, expand **CN=Configuration,DC=*forestname*,DC=*domainname*,DC=com**

2. Right-click **CN=Partitions,** and then click **Properties**.

3. Select the **msDS-Behavior-Version** attribute, and then click **Edit**.

4. In the **Value** field, type **1** to raise the forest functional level to Windows Server 2003 interim, and then click **OK**.

For more information about raising functional levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.

# Delegate the DNS Zone for the New Regional Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory DNS owner in your organization is responsible for delegating the zone that matches the name of the regional domain to the DNS servers that are running on the domain controllers in the regional domain.

Before you create the new regional domain, delegate the DNS zone for the new Windows Server 2003 regional domain on any domain controller in the forest root domain DNS zone.

**To delegate the DNS zone for the new regional domain**

1. Open the DNS snap-in from any domain controller in the forest root domain.

2. In the console tree, right-click the forest root domainzone, and then click **New Delegation**.

3. Table 8.4 lists information to complete the New Delegation Wizard, as well as sample data for delegating the DNS domain for the first two regional domain controllers in the east.trccorp.treyresearch.net domain, SEA-EAST-DC01 and SEA-EAST-DC02. Accept the default settings when no information is supplied.

**Table 8.4 Delegating the DNS Domain for the New Regional Domain**

| Wizard Page | Action | Example |
|---|---|---|
| Delegated Domain Name | In the **Delegated Domain** box, type the name of the regional domain. | **East** |
| Name Servers | a. Click **Add**. In the **New Resource Record** dialog box, in the **Server** name box, type the name of the first domain controller you plan to deploy.<br><br>b. In the **New Resource Record** dialog box, in the **IP address** box, type the corresponding IP address of the domain controller, click **Add**, and then click **OK**.<br><br>c. Click **Add**, and in the **New Resource Record** dialog box, in the **Server name** box, type the name of another domain controller you plan to deploy in the regional domain.<br><br>d. In the **New Resource Record** dialog box, in the **IP address box**, type the corresponding IP address of the other domain controller, click **Add**, and then click **OK**. | **SEA-EAST-DC01.trccorp.treyresearch.net**<br><br>**172.16.16.10**<br><br>**SEA-EAST-DC02.trccorp.treyresearch.net**<br><br>**172.16.16.11** |

# Configure Protection From Domain Controller Overload

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before installing Windows Server 2003 on the Windows NT 4.0 PDC, shield the domain controller by configuring it to emulate a Windows NT 4.0–based domain controller. By shielding the domain controller, clients running Windows 2000, Windows XP, and Windows Server 2003 will not recognize it as an Active Directory domain controller. Clients will authenticate with the new Windows Server 2003–based domain controller as if it were a Windows NT 4.0–based domain controller. This step protects the domain controller from being overloaded with authentication requests from Active Directory clients.

Maintain the emulation setting until enough Windows Server 2003–based domain controllers are in each site to service all Active Directory clients.

**Note**

- After removing the NT4Emulator registry entry, Windows 2000, Windows XP, and Windows Server 2003 clients will not immediately begin to use the Kerberos authentication protocol. This will be delayed until each client resets its secure channel or is restarted.

If no Windows 2000, Windows XP, or Windows Server 2003 clients are running in a particular site, or if a Windows Server 2003–based domain controller has the capacity to support the number of clients that are present in the site, you do not need this configuration.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference.

**To configure emulation on a Windows NT 4.0–based domain controller before upgrade**

1. In the **Run** dialog box, type **regedit**, and then press ENTER.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters
3. Click **Edit**, click **New**, and then click **DWORD Value**.
4. For the new entry name, type **NT4Emulator**, and then press ENTER.
5. Double-click the entry name that you typed in the previous step.
6. In the **Edit DWORD Value** dialog box, type **1** in the **Value data:** box, and then click **OK**.
7. Click **Registry**, and then click **Exit** to close the registry editor.

Repeat this procedure as needed on each Windows NT 4.0–based domain controller that you plan to upgrade to Windows Server 2003.

After you protect the PDC from becoming overloaded, you must be sure to neutralize the emulation on any additional domain controllers you upgrade. Additional domain controllers in the same domain must be able to contact an Active Directory domain controller in their domain for the Active Directory installation to succeed.

On Windows NT 4.0 BDCs, setting the NT4Emulator registry entry before the operating system upgrade will protect the domain controller from overload. Setting the NeutralizeNT4Emulator registry entry immediately afterward allows the BDC to contact an Active Directory domain controller that has the NT4Emulator registry entry set and successfully install Active Directory. For more information about neutralizing Windows NT 4.0 emulation, see "Neutralize Windows NT 4.0 Domain Controller Emulation" later in this chapter.

After you upgrade all domain controllers, or you have enough Windows Server 2003–based domain controllers to authenticate the clients in your domain that are running Windows 2000, Windows XP, and Windows Server 2003, you can reverse this configuration by editing the registry again and removing the NT4Emulator registry entry.

# Upgrade the Operating System of the Windows NT 4.0 PDC

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before upgrading the operating system to Windows Server 2003, use the Winnt32.exe command-line tool to detect any upgrade problems you might have to resolve. This tool reports potential upgrade problems, such as inadequate hardware resources or compatibility problems.

**To determine potential upgrade problems**

- At the command line, connect to the I386 directory located at your installation source and type the following command:

```
winnt32 /checkupgradeonly
```

Resolve any reported problems before performing the upgrade.

To install the operating system on the computer, insert the Windows Server 2003 operating system CD in the CD-ROM drive of the domain controller and select the option to install the operating system, or use an automated installation method. If the Windows Server 2003 media is shared on the network, run the Winnt32.exe command from the installation source.

Complete the operating system installation by doing the following:

- Select **Upgrade** for the Installation type.
- Use the NTFS file system to convert the partitions. The installation of Active Directory will not succeed if you do not have at least one NTFS partition available on which to locate the SYSVOL shared folder.
- Verify that you are using a static IP address.
- Configure DNS client settings by using the IP address of the closest DNS server for the **Preferred DNS server** setting and either leave the **Alternate DNS server** setting blank or use the IP address of the closest DNS server. These DNS client settings are temporary and will be changed during the installation of Active Directory.
- Install Windows Support Tools, which are available in the \Support\Tools folder on the Windows Server 2003 operating system CD.

During the operating system upgrade the computer will restart three times. After you upgrade the operating system on a Windows NT 4.0 domain controller to Windows Server 2003, the computer is in an intermediate state, meaning that the computer is no longer a Windows NT 4.0–based domain controller, and it is not a Windows Server 2003–based member server or domain controller until Active Directory is installed. After the computer restarts for the last time, the Welcome to the Active Directory Installation Wizard appears.

# Installing Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory Installation Wizard creates the Active Directory database and moves objects from the Windows NT 4.0 Security Accounts Manager (SAM) to the Active Directory database.

**Note**

- When you are upgrading to a regional domain in an existing Active Directory forest, ensure that the domain naming master in the forest root domain is running Windows Server 2003 before installing Active Directory on the newly upgraded PDC. This ensures that application directory partitions are created on the first domain controller in the new regional domain.

In addition, on the first domain controller in a new regional domain in an existing forest, the wizard does the following:

- Prompts the administrator to verify the installation and configuration of the DNS Server service.
- Configures DNS recursive name resolution forwarding by adding the IP addresses of the existing entries for **Preferred DNS server** and **Alternate DNS server** to the list of DNS servers on the **Forwarders** tab of the **Properties** sheet for the domain controller.
- Configures DNS recursive name resolution by root hints, by adding the root hints that are configured on the **Preferred DNS server** to the list of DNS servers on the **Root Hints** tab of the **Properties** sheet for the domain controller.
- Configures the **Preferred DNS server** to point to the DNS server that is running locally on the domain controller, and configures the **Alternate DNS server** to point to the closest DNS server.
- Creates the DomainDnsZones application directory partition that is used by DNS to hold domain-wide DNS data.

Table 8.5 lists information to install Active Directory on an upgraded Windows NT 4.0 PDC and sample data for installing Active Directory on the first domain controller in a new regional domain in the trccorp.treyresearch.net forest, SEA-EAST-DC01.

**Table 8.5 Information to Install Active Directory on a Windows NT 4.0 PDC**

| Wizard Page or Dialog Box | Action | Example |
|---|---|---|
| Create New Domain | Select **ChildDomain in an existing domain tree** | |
| Network Credentials | Type the user name and password of an account with sufficient privileges to install Active Directory on this computer, and the fully qualified domain name of the parent domain. | |
| Child Domain Installation | Enter the full DNS name of the parent domain and the single label name of the new regional domain. | **trccorp.treyresearch.net**<br>**east** |
| Database and Log Folders | Type the folder locations specified by your design | The design for Trey Research specifies that the database folder remain in the default location: **C:\Winnt\Ntds**, and that the log folder is placed on a separate partition: **D:\Logs** |

| | | |
|---|---|---|
| Shared System Volume | Confirm or type the location specified by your design | **C:\Winnt\Sysvol** |
| DNS Registration Diagnostics | DNS Registration Diagnostics will indicate that it cannot find the name and address of the DNS server with which this domain controller will be registered. This is because the pre-created delegation record points to the local computer and DNS has not been installed on the domain controller at this point.<br><br>Select the option to **Install and configure the DNS server on this computer and set this computer to use this DNS server as its preferred DNS server.** | |
| Permissions | Select the security level specified by your design:<br><br>• Permissions compatible with pre-Windows 2000 server operating systems<br><br>• Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems | Because Trey Research currently has services running on Windows NT 4.0–based servers under the context of the Local System account, they selected **Permissions compatible with pre-Windows 2000 server operating systems**. |
| Directory Service Restore Mode Administration Password | In the **Password** and **Confirm password** boxes, type any strong password | |

Verify that all information on the Summary page is accurate, and then click **Finish.** After the Active Directory Installation Wizard finishes, you will be prompted to restart the computer. The installation will not be complete until the computer restarts.

For more information about installing and removing Active Directory, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

After you install Windows Server 2003 Active Directory, enable Remote Desktop for Administration, formerly known as Terminal Services in Remote Administration mode, to enable administrators to log on remotely if necessary. To enable Remote Desktop for Administration, in Control Panel, double-click **System**, select the **Remote** tab, and then select **Allow users to connect remotely to this computer**.

# Verify the DNS Server Recursive Name Resolution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DNS server recursive name resolution is configured automatically during the Active Directory installation process. If your design specifies a different configuration, you can use the DNS snap-in or Dnscmd.exe to modify these settings. Use the DNS snap-in to verify DNS server recursive name resolution based on the information in Table 8.6.

**Table 8.6 Information to Verify DNS Server Recursive Name Resolution**

| Method | Configuration |
|---|---|
| Recursive name resolution by root hints | No additional configuration is necessary. When the DNS server specified as the Preferred DNS server during the installation process is correctly configured, the root hints are automatically configured. To verify the root hints by using the DNS snap-in:<br>1. In the console tree, right-click the domain controller name, and then click **Properties**.<br>2. In the **Properties** sheet for the domain controller, view the root hints on the **Root Hints** tab.<br>Root hints are the recommended method to use for recursive name resolution in a Windows Server 2003 environment. |
| Recursive name resolution by forwarding | Forward unresolved queries to specified DNS servers. To verify forwarding by using the DNS snap-in:<br>1. In the console tree, right-click the domain controller name, and then click **Properties**.<br>2. On the **Forwarders** tab, in the selected domain's Forwarders list, verify that the IP addresses match those specified by your design.<br>Forwarders should be used only if that is what your organization's design specifies. Root hints are the recommended method to use for recursive name resolution in a Windows Server 2003 environment. |
| No existing DNS infrastructure | No additional configuration is necessary.<br>In this environment, if you want to configure internal DNS servers to resolve queries for external names, configure this DNS server to forward unresolved queries to an external server, such as one in your perimeter network, or one hosted by an Internet service provider. |

# Perform Post-Upgrade Tests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After the Active Directory Installation Wizard completes, verify the Active Directory installation. Review the Windows Server 2003 Event Viewer for any errors and use the DNS snap-in to verify that the DomainDnsZones was created under the DNS root zone.

You can also verify that Event ID 4500 has been logged in DNS Events to ensure that application directory partitions have been created.

Next, perform the tests defined in your test plan to test the Active Directory configuration and verify whether Active Directory is functioning correctly. For more information about developing a test plan, see "Develop a Test Plan" earlier in this chapter.

After performing post-upgrade tests and verifying that the upgrade of the Windows NT 4.0 PDC and the installation of Active Directory succeeded, complete the upgrade process by continuing to "Modify Security Policies" later in this chapter.

# Upgrade to a Single Domain Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To complete the process for upgrading to a single domain forest, perform the following tasks:

1. Back up all domain data.
2. If you have an existing DNS infrastructure, delegate the DNS zone for the new Windows Server 2003 domain.
3. Configure protection against domain controller overload.
4. Upgrade the operating system of the Windows NT 4.0 PDC.
5. Install Active Directory.
6. Configure the site topology.
7. Configure the Windows Time Service.
8. Enable aging and scavenging for DNS.
9. Verify DNS server recursive name resolution.
10. Perform post-upgrade tests.

**Important**

- When upgrading to a single domain forest, any individual who is a member of the Domain Admins group in the existing Windows NT 4.0 domain will become a member of the Domain Admins and Enterprise Admins groups. Before upgrading the first Windows NT 4.0 domain, remove users whom you do not want to have full access to the entire forest from both the Administrators and Domain Admins groups.

# Back Up the Domain Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Back up your Windows NT 4.0 domain data before you begin the upgrade. This task varies according to the operations and procedures that already exist in your environment. It is recommended that you complete the following steps:

- Back up the PDC.
- Back up the BDC that you designated as the rollback server.
    - Test all backup media to ensure that the data can be restored successfully.

**Important**

- Before you begin the upgrade process, store the backup media in a secure offsite location designated by and accessible to the upgrade team.

# Delegate the DNS Zone for the Windows Server 2003 Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization has an existing DNS infrastructure, review current network diagrams and DNS domain hierarchy diagrams. Also review the existing DNS zone configuration, replication, and resource records that are used for delegation and forwarding. To configure the DNS zone for the single domain forest, the DNS administrator of your existing DNS infrastructure delegates the zone matching the name of the new Windows Server 2003 domain to the DNS servers that are running on the domain controllers in the single domain forest.

**Important**

- When no DNS infrastructure exists, skip this step in the process for upgrading to a single domain forest and proceed to the next step, "Configure Protection Against Domain Controller Overload on Additional Domain Controllers" later in this chapter. The remainder of this step describes the process of configuring and delegating a zone in the existing DNS internal namespace.

In preparation for the deployment of the single domain forest, create a delegation for the DNS servers that will be running on the domain controllers in the Windows Server 2003 domain. Create the delegation by adding DNS name server (NS) and address (A) resource records to the parent DNS zone.

**Note**

- The delegation that occurs in this step references the first Windows Server 2003–based domain controller, which does not currently exist. The DNS service is installed and configured on the first Windows Server 2003–based domain controller in a later step.

**To delegate the DNS zone for the Windows Server 2003 domain**

1. Create a name server (NS) resource record in the parent zone. Use the full DNS name of the domain controller.

   ***forest_root_domain*** IN NS ***domain_controller_name***

2. Create a host address (A) resource record in the parent zone. Use the full DNS name of the domain controller.

   ***domain_controller_name*** IN A ***domain_controller_ip_address***

   For example, the DNS administrator for Fabrikam created the following DNS resource records in the parent zone, fabrikam.com:

   - **fabricorp** IN NS **SEA-FAB-DC01.fabricorp.fabrikam.com**
   - **SEA-FAB-DC01.fabricorp.fabrikam.com** IN A **172.16.16.2**

# Configure Protection Against Domain Controller Overload

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before installing Windows Server 2003 on the Windows NT 4.0 PDC, shield the domain controller by configuring it to emulate a Windows NT 4.0–based domain controller. By shielding the domain controller, clients running Windows 2000, Windows XP, and Windows Server 2003 will not recognize it as an Active Directory domain controller. Clients will authenticate with the new Windows Server 2003–based domain controller as if it were a Windows NT 4.0–based domain controller, protecting it from being overloaded with authentication requests from Active Directory clients.

Maintain the emulation setting until there are enough Windows Server 2003–based domain controllers in each site to service all Active Directory clients.

**Note**

- After removing the NT4Emulator registry entry, Windows 2000, Windows XP, and Windows Server 2003 clients will not immediately begin to use the Kerberos authentication protocol. This will be delayed until each client resets its secure channel or is restarted.

If no Windows 2000, Windows XP, or Windows Server 2003 clients are running in a particular site, or if the Windows Server 2003–based domain controller has the capacity to support the number of clients that are present in the site, this configuration is not needed.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back up first and see the Windows Server 2003 Resource Kit Registry Reference.

**To configure emulation on a Windows NT 4.0–based domain controller before upgrade**

1. In the **Run** dialog box, type **regedit**, and press ENTER.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters
3. Click **Edit**, click **New**, and then click **DWORD Value**.
4. For the new entry name, type **NT4Emulator**, and then press ENTER.
5. Double-click the name that you typed in the previous step.
6. In the **Edit DWORD Value** dialog box, type **1** in the **Value data:** box, and then click **OK**.
7. Click **Registry**, and then click **Exit** to close the registry editor.

Repeat this procedure as needed on each Windows NT 4.0–based domain controller that you plant to upgrade to a Windows Server 2003.

After you protect the PDC from becoming overloaded, you must neutralize the emulation on any additional domain controllers that you plan to upgrade. For the Active Directory installation to succeed, additional domain controllers in the same domain must be able to contact an Active Directory domain controller in their domain.

On Windows NT 4.0 BDCs, setting the NT4Emulator registry entry before the operating system upgrade will protect the domain controller from overload. Setting the NeutralizeNT4Emulator registry entry immediately afterward will allow the BDC to contact an Active Directory domain controller that has the NT4Emulator registry entry set and successfully install Active Directory. For more information about neutralizing Windows NT 4.0 emulation, see "Neutralize Windows NT 4.0 Domain Controller Emulation" later in this chapter.

After you upgrade all domain controllers, or you have enough Windows Server 2003–based domain controllers to authenticate the clients in your domain that are running Windows 2000, Windows XP, and Windows Server 2003, you can reverse this configuration by editing the registry again and removing the NT4Emulator registry entry.

# Upgrade the Operating System of the Windows NT 4.0 Primary Domain Controller (PDC)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before upgrading the operating system to Windows Server 2003, use the Winnt32.exe command-line tool to detect any upgrade problems. This tool reports potential upgrade problems, such as inadequate hardware resources or compatibility problems.

**To determine potential upgrade problems**

- At the command line, connect to the I386 directory located at your installation source and type the following command:

```
winnt32 /checkupgradeonly
```

Resolve reported problems before performing the upgrade.

To install the operating system on the computer, insert the Windows Server 2003 operating system CD in the CD-ROM drive of the domain controller and select the option to install the operating system, or use an automated installation method. If the Windows Server 2003 media is shared on the network, run the Winnt32.exe command.

Complete the operating system installation by doing the following:

1. Select **Upgrade** for the Installation type.

2. Use NTFS to convert the partitions. The installation of Active Directory will not succeed if you do not have at least one NTFS partition available on which to locate the SYSVOL shared folder.

3. Verify that you are using a static IP address.

4. Configure DNS client settings by using the IP address of the closest DNS server for the **Preferred DNS Server** settings. If you have more than one DNS server, add the IP address of the next closest DNS server to the **Alternate DNS server** setting. If there are no other DNS servers, leave the alternate setting blank. These DNS client settings are temporary and will be changed during the installation of Active Directory.

5. Install Windows Support Tools, which are available in the \Support\Tools folder on the Windows Server 2003 operating system CD.

During the operating system upgrade the computer will restart three times. After you upgrade the operating system on a Windows NT 4.0 domain controller to Windows Server 2003, the computer is in an intermediate state, meaning that the computer is no longer a Windows NT 4.0–based domain controller, and it is not a Windows Server 2003–based member server or domain controller until Active Directory is installed. After the computer restarts for the last time, the Welcome to the Active Directory Installation Wizard appears.

# Install Active Directory by Continuing the Active Directory Installation Wizard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Immediately proceed with the installation of Active Directory by continuing the Active Directory Installation Wizard. The Active Directory Installation Wizard creates the Active Directory database and moves objects from the Windows NT 4.0 SAM to the Active Directory database. In addition, on the first domain controller in a new domain, the wizard proceeds through the following tasks:

- Prompts the administrator to verify the installation and configuration of the DNS Server service.
- Configures DNS recursive name resolution forwarding by adding the IP addresses of the existing entries for **Preferred DNS server** and **Alternate DNS server** to the list of DNS servers on the **Forwarders** tab of the **Properties** sheet for the domain controller.
- Configures DNS recursive name resolution by root hints, by adding the root hints that are configured on the **Preferred DNS server** to the list of DNS servers on the **Root Hints** tab of the **Properties** sheet for the domain controller.
- Configures the **Preferred DNS server** to point to DNS server that is running locally on the domain controller, and configures the **Alternate DNS server** to point to the closest DNS server.
- Creates two application directory partitions that are used by DNS. The DomainDnsZones application directory partition holds domain-wide DNS data, and the ForestDnsZones application directory partition holds forest-wide DNS data.
- Prompts the administrator to select the forest functional level.

Table 8.7 lists information to install Active Directory on a Windows NT 4.0 PDC, and lists sample data for installing Active Directory on the first domain controller in the single domain forest for Fabrikam, SEA-FAB-DC01.

**Table 8.7 Information to Install Active Directory on a Windows NT 4.0 PDC**

| Wizard Page or Dialog Box | Action | Example |
|---|---|---|
| Create New Domain | Select **Domain in a new forest** | |
| New Domain Name | Type the full DNS name of the domain | **Fabricorp.fabrikam.com** |
| Forest Functional Level | Choose **Windows Server 2003 interim** | Because Fabrikam does not plan to add any Windows 2000–based domain controllers to their forest at any time they chose the **Windows Server 2003 interim** forest functional level. |
| Database and Log Folders | Type the folder locations specified by your design | The design for Trey Research specifies that the database folder remain in the default location: **C:\Winnt\Ntds**, and that the log folder is placed on a separate partition: **D:\Logs** |

| | | |
|---|---|---|
| Shared System Volume | Confirm or type the location specified by your design | **C:\Winnt\Sysvol** |
| DNS Registration Diagnostics | DNS Registration Diagnostics will indicate that it cannot find the name and address of the DNS server with which this domain controller will be registered. This is because the pre-created delegation record points to the local computer and DNS has not been installed on the domain controller at this point.<br><br>Select the option to **Install and configure the DNS server on this computer and set this computer to use this DNS server as its preferred DNS server.** | |
| Permissions | Select the security level specified by your design:<br><br>• Permissions compatible with pre-Windows 2000 server operating systems<br><br>• Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems | Because Fabrikam currently has services running on Windows NT 4.0–based servers under the context of the Local System account, they selected **Permissions compatible with pre-Windows 2000 server operating systems** |
| Directory Service Restore Mode Administration Password | In the **Password** and **Confirm password** boxes, type any strong password | |

Verify that all information on the Summary page is accurate, and then click **Finish**. After the Active Directory Installation Wizard finishes, you will be prompted to restart the computer. The installation will not be complete until the computer restarts.

For more information about installing and removing Active Directory, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

After you install Windows Server 2003 Active Directory, enable Remote Desktop for Administration, formerly known as Terminal Services in Remote Administration mode, to enable administrators to log on remotely if necessary. To enable Remote Desktop for Administration, in Control Panel, double-click **System**, select the **Remote** tab, and then select **Allow users to connect remotely to this computer**.

# Configure the Site Topology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you install Active Directory on the first domain controller in the forest, a site object named Default-First-Site-Name is created in the Sites container in Active Directory. The server object for the first domain controller is created in this site.

If no additional sites have been defined in Active Directory, the server object for all subsequent domain controllers will be added to the Default-First-Site-Name site object. However, if additional sites are defined in Active Directory and the IP address of the installation computer matches an existing subnet in a defined site, the domain controller is added to that site.

**Note**

- Domain controllers are only added to sites based on their IP address at the time of installation. After installation, if the IP address, subnet, or site information of a domain controller changes, an administrator must manually move the domain controller to the new site.

To simplify the placement of the domain controller into the appropriate site, configure your site topology before you install Active Directory on additional domain controllers. After all sites are created, a server object for each additional domain controller is created in the appropriate site according to its IP address.

Configure your Active Directory site topology as specified in your site topology design. For information about creating a site topology design, see "Designing the Site Topology" in this book. For more information about configuring your site topology, see "Configure Site Settings" and "Configure Replication Between Sites" in Help and Support Center for Windows Server 2003.

# Configure the Windows Time Service on the Forest Root Domain Controller

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When deploying a single domain forest, it is important to correctly configure the Windows Time Service on the forest root domain controller to meet your organization's needs. The Windows Time Service provides time synchronization to peers and clients, ensuring that there is consistent time throughout an enterprise.

By default, the first domain controller that is deployed holds the PDC emulator operations master role, and should be set to synchronize from a valid Network Time Protocol (NTP) source. If no source is configured, the service will log a message to the event log, and use the local clock when providing time to clients. Although internet NTP sources are valid for this configuration, it is recommended that a dedicated hardware device, such as a GPS, or Radio clock be employed in the interest of security.

It is recommended that you repeat this operation when the PDC emulator operations master role is transferred or seized in the forest root domain.

**To configure the Windows Time Service on first forest root domain controller**

1. Log on to the domain controller.

2. At the command line, type:

   ```
   W32tm /config /manualpeerlist:<peers> /syncfromflags:manual
   ```

   where <peers> is a space–delimited list of DNS and/or IP addresses. When specifying multiple peers, enclose the list in quotation marks.

3. Update the Windows Time Service configuration. At the command line, type:

   ```
   W32tm /config /update
   ```

   – or –

   ```
   Net stop w32time
   Net start w32time
   ```

### Note

- When specifying a manual peer, do not use the DNS name or IP address of a computer that uses the forest root domain controller as its source for time, such as another domain controller in the forest. The time service will not operate correctly if there are cycles in the time source configuration.

For more information about configuring and deploying the Windows Time Service, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

**See Also**
**Other Resources**
Videos about the Windows Time Service

# Enable Aging and Scavenging for Domain Name System (DNS)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a new single domain forest you will need to enable aging and scavenging on Windows Server 2003–based domain controllers running the DNS Server service to allow automatic cleanup and removal of stale resource records (RRs), which can accumulate in zone data over time.

With dynamic update, RRs are automatically added to zones when computers start on the network. However, in some cases, they are not automatically removed when computers leave the network. For example, if a computer registers its own host (A) RR at startup, and is later incorrectly disconnected from the network, its host (A) RR might not be deleted. If your network has mobile users and computers, this situation can occur frequently.

If left unmanaged, the presence of stale RRs in zone data might cause problems, including the following:

- If a large number of stale RRs remain in server zones, they can eventually take up server disk space and cause unnecessarily long zone transfers.
- DNS servers loading zones with stale RRs might use outdated information to answer client queries, potentially causing the clients to experience name resolution problems on the network.
- The accumulation of stale RRs at the DNS server can degrade its performance and responsiveness.

**Caution**

- By default, the aging and scavenging mechanism for the DNS Server service is disabled. Enable aging and scavenging only after you understand all parameters. Otherwise, the server could be accidentally configured to delete resource records that should not be deleted. If a resource record is accidentally deleted, not only will users fail to resolve queries for that resource record, but any user can create the resource record and take ownership of it, even on zones configured for secure dynamic update.
- For more information about how to configure aging and scavenging, see "Understanding aging and scavenging" in Help and Support Center for Windows Server 2003.

To enable the aging and scavenging features, and to configure the applicable server and its Active Directory–integrated zones, perform these tasks:

- Enable aging and scavenging at the server. These settings determine the effect of zone-level properties for any Active Directory–integrated zones loaded at the server.
- Enable aging and scavenging for selected zones at the DNS server. When zone-specific properties are set for a selected zone, these settings apply only to the applicable zone and its resource records. Unless these zone-level properties are otherwise configured, they inherit their defaults from comparable settings maintained in server aging/scavenging properties.

**To set aging and scavenging properties for the DNS server**

1. Log on to the computer that is running the DNS Server service with an account that is a member of the local Administrators group.
2. In the DNS console tree, right-click the applicable DNS server, and then click **Set Aging/Scavenging for all zones**.
3. Select the **Scavenge stale resource records** check box.
4. Modify other aging and scavenging properties as needed.

**To set aging and scavenging properties for a zone**

1. Log on to the computer that is running the DNS Server service with an account that is a member of the local Administrators group.

2. In the DNS console tree, right-click the applicable zone, and then click **Properties**.

3. On the **General** tab, click **Aging**, and then select the **Scavenge stale resource records** check box.

4. Modify other aging and scavenging properties as needed.

# Verify DNS Server Recursive Name Resolution When Upgrading to a Single Domain Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DNS server recursive name resolution is configured automatically during the Active Directory installation process. If your design specifies a different configuration, you can use the DNS snap-in or Dnscmd.exe to modify these settings. Use the DNS snap-in to verify DNS server recursive name resolution based on the information in Table 8.8.

**Table 8.8 Information to Verify DNS Server Recursive Name Resolution**

| Method | Configuration |
|---|---|
| Recursive name resolution by root hints | No additional configuration is necessary. When the DNS server specified as the Preferred DNS server during the installation process is correctly configured, the root hints are automatically configured. To verify the root hints by using the DNS snap-in:<br>1. In the console tree, right-click the domain controller name, and then click **Properties**.<br>2. In the **Properties** sheet for the domain controller, view the root hints on the **Root Hints** tab.<br>Root hints are the recommended method to use for recursive name resolution in a Windows Server 2003 environment. |
| Recursive name resolution by forwarding | Forward unresolved queries to specified DNS servers. To verify forwarding by using the DNS snap-in:<br>1. In the console tree, right-click the domain controller name, and then click **Properties**.<br>2. On the **Forwarders** tab, in the selected domain's Forwarders list, verify that the IP addresses match those specified by your design.<br>Forwarders should be used only if that is what your organization's design specifies. Root hints are the recommended method to use for recursive name resolution in a Windows Server 2003 environment. |
| No existing DNS infrastructure | No additional configuration is necessary.<br>In this environment, if you want to configure internal DNS servers to resolve queries for external names, then configure this DNS server to forward unresolved queries to an external server, such as one in your perimeter network, or one hosted by an Internet service provider. |

# Perform Post-Upgrade Tests After the Active Directory Installation Wizard Completes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After the Active Directory Installation Wizard completes, verify the Active Directory installation. Review the Windows Server 2003 event log for any errors and use the DNS snap-in to verify that the following two DNS zones were created under the DNS root zone:

- DomainDnsZones
- ForestDnsZones

You can also verify that Event ID 4500 has been logged in DNS Events to ensure that application directory partitions have been created.

Next, perform the tests defined in your test plan to test the Active Directory configuration and verify whether Active Directory is functioning correctly. For more information about developing a test plan, see "Develop a Test Plan" earlier in this chapter.

After performing post-upgrade tests and verifying that the upgrade of the Windows NT 4.0 PDC and the installation of Active Directory succeeded, complete the upgrade process by continuing to "Modify Security Policies" later in this chapter.

# Modify Security Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To ensure that clients running earlier versions of the Windows operating system can access domain resources in the new Windows Server 2003 domain, you might have to modify default security policies.

In order to increase security, Windows Server 2003–based domain controllers require by default that clients attempting to authenticate to them use SMB packet and secure channel signing. Clients running the Windows 95 operating system without the Directory Service Client Pack or Windows NT 4.0 with Service Pack 2 and earlier do not support SMB packet signing and will not be able to log on or access domain resources on the network. Clients running Windows NT 4.0 with Service Pack 3 and earlier do not support secure channel signing and will not be able to establish communications with a domain controller in their domain.

The most secure way to enable these clients to log on and access domain resources on the network is to apply either the appropriate service pack or the Directory Service Client Pack. If you cannot apply either of these, configure all Windows Server 2003–based domain controllers to not require SMB packet signing and secure channel signing. To do this, disable the following settings in the Default Domain Controllers Policy:

- **Microsoft network server: Digitally sign communications (always)**
- **Domain member: Digitally encrypt or sign secure channel data (always)**

**Important**

- If you modify these policies, the default security policies in your environment are weakened. However, this is necessary to ensure that some clients running earlier versions of Windows can access domain resources. After all the clients in your environment are running versions of Windows that support SMB packet and secure channel signing, you can re-enable these security policies to increase security. It is recommended that you upgrade your Windows clients as soon as possible.

**To make SMB packet and secure channel signing optional on Windows Server 2003–based domain controllers**

1. Open **Active Directory Users and Computers**, right-click the **Domain Controllers** container, and then click **Properties**.
2. Select the **Group Policy** tab, and then click **Edit**.
3. Under **Computer Configuration**, navigate to Windows Settings\Security Settings\Local Policies\Security Options.
4. In the details pane, double-click **Microsoft network server: Digitally sign communications (always)** and then click **Disabled** to prevent SMB packet signing from being required.
5. Click **OK**.
6. In the Details pane, double-click **Domain member: Digitally encrypt or sign secure channel data (always),** click **Disabled** to prevent secure channel signing from being required, and then click **OK**.
7. To apply the Group Policy change immediately, either restart the domain controller, or run the **gpupdate /force** command.

**Note**

- Modifying these settings in the Domain Controllers container will change the Default Domain Controllers Policy. Policy changes that are made here are replicated to all other domain controllers in the domain, requiring you to modify these policies only one time.

For more information about SMB packet signing and secure channel signing, see "Background Information for Upgrading to Windows Server 2003 Active Directory" earlier in this chapter.

For more information about security policies, see "Security Options" on the Microsoft Web site.

# Synchronize File Replication Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After upgrading the Windows NT 4.0 PDC, create a script to copy logon script and profile information from the NETLOGON shared folder on the new Windows Server 2003–based domain controller to the REPL$ share on the Windows NT 4.0 BDC that is providing export services to other Windows NT 4.0 BDCs in the domain.

**To create logon script and profile replication between Windows Server 2003–based and Windows NT 4.0–based domain controllers**

1. Create a user account on a Windows NT 4.0 BDC in the Windows Server 2003 domainby using User Manager for Domains and the information in Table 8.9.

   **Table 8.9 User Account Information for Logon Script and Profile Replication**

   | When Prompted For | Use |
   |---|---|
   | User name | LbridgeAcct |
   | Description | Account used by lbridge.cmd for replication. |
   | Password | *password* (where *password* is any password that meets the security requirements for your organization). |
   | Select **User Must Change Password At Next Logon** to clear the checkbox. | |

2. Use the information listed in Table 8.10 to ensure that the LbridgeAcct has the correct permissions on both the Windows Server 2003–based domain controller and on the Windows NT 4.0 BDC.

   **Table 8.10 Permissions for the LbridgeAcct User Account**

   | Folder | Permission |
   |---|---|
   | On the Windows NT 4.0 BDC, ensure that the LbridgeAcct is granted Full Control to the REPL$ share.<br><br>In Server Manager, select the computer configured as the export server, click **Computer**, and select **Shared Directories.** Select **REPL$**, and then click **Properties**. In the **Share Properties** dialog box, click **Permissions**, click **Add**, and then click **Show Users**. Select the LbridgeAcct. In the **Type of Access** drop-down list box, select **Full Control**. | Full Control |
   | On the Windows Server 2003–based domain controller in the new Windows Server 2003 domain, ensure that the LbridgeAcct is granted Read access to the NETLOGON shared folder.<br><br>Access the NETLOGON shared folder by typing *win_dc*\Netlogon (where *win_dc* is the name of the Windows Server 2003–based domain controller) in the Run dialog box. | Read |

3. Create a destination folder on the Windows Server 2003–based domain controller where you will install the **Lbridge.cmd** script and the **Robocopy.exe** tool.

4. Modify the path statement in Environment Variables to include the destination folder. Right-click **My Computer**, click **Properties**, click the **Advanced** tab, and then click **Environment Variables**. In the System Variables list, select **Path** and click **Edit**. Append the Variable value with the location of the destination folder (*;C:\destination folder*).

   The **Lbridge.cmd** script and **Robocopy.exe** tools are available on the *Windows Server 2003 Deployment Kit* companion CD.

5. On the Windows Server 2003–based domain controller, in Windows Explorer, right-click the Lbridge.cmd script, and then

click **Edit**. Edit as indicated in Table 8.11.

**Table 8.11 Modifications to the lbridge.cmd Script**

| Script Line | Change To |
|---|---|
| Set L-Destination=%1 | Set L-Destination=***winnt_dc*\REPL$** (where *winnt_dc* is name of the Windows NT 4.0 BDC hosting the LMRepl export server. |
| Call :Xcopy | @Rem Call :Xcopy |
| @Rem Call :Robocopy | Call :Robocopy |
| Echo Robocopy %L-Source%%L-Destination% /E /PURGE | Robocopy %L-Source%%L-Destination% /E /PURGE |

6. On the Windows Server 2003–based domain controller, open **Control Panel**, point to **Scheduled Tasks**, and then click **AddScheduled Task**.

7. Complete the Scheduled Task Wizard by using the information in Table 8.12. Accept the default settings when no information is supplied.

**Table 8.12 Scheduled Task Wizard Actions for Lbridge.cmd**

| Wizard Page | Action |
|---|---|
| Click the program you want Windows to run | Click **Browse**.<br><br>In the **Select Program to Schedule** dialog box, click lbridge.cmd. |
| Type a name for this task | Type **FRS - LMRepl Replication Bridge**. |
| Perform this task | Select **Daily**. |
| Start time | Enter the time and date that you want the replication to start. |
| Enter the user name | Type **LbridgeAcct**. |
| Enter the password | Type the password that you have chosen for LbridgeAcct. |
| Confirm the password | Confirm the password for LbridgeAcct. |
| Open advanced properties for this task when I click Finish | Select the check box. |

The **FRS - LMRepl Replication Bridge** dialog box opens.

8. In the **FRS - LMRepl Replication Bridge** dialog box, on the **Schedule** tab, click **Advanced**.

9. In the **Advanced Schedule Options** dialog box, select the **Repeat task** check box.

10. In the **Every** box, specify how often you want the script to run.

11. In the **Duration** box, specify how long you want the script to run, and then click **OK**.

12. In the **FRS - LMRepl Replication Bridge** dialog box, on the **Schedule** page, click **OK**.

To verify that the script ran successfully and replication is occurring from the Windows Server 2003–based domain controller to all Windows NT 4.0–based domain controllers, place a file called Test.txt into the

***Win_dc*\SYSVOL\sysvol*domainname*\scripts** folder. After replication is scheduled to take place, verify that the Test.txt file has replicated to the **\Winnt_dc\system32\REPL\Import\scripts** folder.

# Recreate Trusts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Trust relationships between Windows NT 4.0 domains use NetBIOS domain names. During the in-place upgrade of your Windows NT 4.0 environment, if some of your Windows NT 4.0 domains have trust relationships with other Windows NT 4.0 domains that are then upgraded into separate forests, those trust relationships between the domains in different forests remain, but continue to use the NetBIOS domain name. It is recommended that trust relationships between domains in different forests use the DNS name for the domain in order to gain better functionality in a Windows Server 2003 environment. To rename the trust relationship by using the DNS name for the domain, delete and recreate external trust relationships that exist between Windows NT 4.0 domains and Active Directory domains in different forests. Trusts that use NetBIOS names and exist between Windows NT 4.0 domains can be left in place.

# Use DNS Registration to Decrease the Workload on the PDC Emulator

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After upgrading the Windows NT 4.0 PDC, the domain controller hosts the PDC emulator operations master role. Of all the operations master roles, the PDC emulator role has the greatest effect on the domain controller that is hosting that role because the PDC emulator fulfills additional tasks in the domain, such as processing password changes, processing authentication requests if the password fails on the authenticating domain controller, and all write operations to the domain that are requested or performed by applications or by clients running Windows 2000, Windows XP, and Windows Server 2003.

In domains with more than 10,000 users, it might be necessary to reduce the number of authentication requests received by the PDC emulator and allow it to perform other tasks. If, after upgrading the Windows NT 4.0 PDC, CPU utilization is higher than 50 percent or if disk queues remain higher than two for several hours or days, reduce the number of client authentication requests that are received by the PDC emulator.

**Note**

- Other factors that can increase the workload on the PDC emulator include pre-Active Directory clients or applications that have been written to contact the PDC emulator.

To reduce the number of client authentication requests that are processed by the PDC emulator, adjust its weight or its priority in the DNS environment. If you want to proportionately reduce the number of client authentication requests received by the PDC emulator, adjust its weight. If you want to ensure that the PDC emulator does not receive any client authentication requests, adjust its priority.

Active Directory assigns a default value of 100 for the weight. By creating a new registry entry for the weight and assigning it a decreased value of 50, you can proportionately reduce the number of client authentication requests that are sent to the PDC. This ensures that the PDC will authenticate half of the number of clients that it would if the weight value remained at 100.

Active Directory assigns a default value of zero for the priority. By creating a new registry entry for the priority and assigning it an increased value of 200, you can ensure that the PDC will never receive client authentication requests unless it is the only accessible domain controller.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the "Registry Reference" on the *Windows Server 2003 Deployment Kit* companion CD or on the Microsoft Web site.

**To change the weight for DNS SRV records by using the registry**

1. In the **Run** dialog box, type **regedit**, and then press ENTER.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters
3. Click **Edit**, click **New**, and then click **DWORD value**.
4. For the new entry name, type **LdapSrvWeight**, and then press ENTER. (The value name is not case sensitive.)
5. Double-click the entry name you just typed.
6. In the **Edit DWORD Value** dialog box, select **Decimal** as the **Base** option.
7. Enter a value between 0 and 65535 (the recommended value is 50), and then click **OK**.
8. Click **File**, and then click **Exit** to close the registry editor.

Adjusting the priority of the domain controller also reduces the number of client referrals. However, rather than reducing it proportionally to the other domain controllers, changing the priority causes DNS to stop referring all clients to this domain controller unless all domain controllers with a lower priority setting are unavailable.

**Note**

- A lower value entered for **LdapSrvPriority** indicates a higher priority. A domain controller with an **LdapSrvPriority** setting of 100 has a lower priority than a domain controller with a setting of 10. Therefore, clients attempt to use the domain controller with the setting of 10 first.

**To change the priority for DNS SRV records by using the registry**

1. In the **Run** dialog box, type **regedit**, and then press ENTER.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters
3. Click **Edit**, click **New**, and then click **DWORD value**.
4. For the new entry name, type **LdapSrvPriority**, and then press ENTER.
5. Double-click the entry name that you just typed.
6. In the **Edit DWORD Value** dialog box, select **Decimal** as the **Base** option.
7. Enter a value between 0 and 65535 (the recommended value is 200), and then click **OK**.
8. Click **File**, and then click **Exit** to close the registry editor.

For more information about adjusting the weight or the priority of the PDC emulator, see the Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search under "Administration and Configuration Guides" and download the *Active Directory Operations Guide*.

# Upgrade Additional Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you upgrade the operating system and install Active Directory on the Windows NT 4.0 PDC, add another Windows Server 2003–based domain controller to the domain as soon as possible. This provides redundancy for any clients running in the environment.

You can add additional domain controllers to the Windows Server 2003 domain by upgrading Windows NT 4.0–based BDCs and installing Active Directory, or by adding Windows Server 2003–based member servers to the domain and installing Active Directory on the member servers.

To complete the process for upgrading additional domain controllers, perform the following tasks:

1. Configure protection against domain controller overload.
2. Neutralize Windows NT 4.0 domain controller emulation.
3. Upgrade the operating system of Windows NT 4.0 BDCs.
4. Install Active Directory.
5. Install DNS on additional domain controllers.
6. Reconfigure the DNS Service.
7. Add Windows NT 4.0 BDCs to the Windows Server 2003 domain if necessary.
8. Perform post-upgrade tests.

# Configure Protection Against Domain Controller Overload on Additional Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure an additional domain controller against overload, perform the same steps that were performed to configure protection on the Windows NT 4.0 PDC. Configure the domain controller to emulate a Windows NT 4.0–based domain controller before you upgrade the operating system and install Active Directory.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference.

**To configure emulation on a Windows NT 4.0–based domain controller before upgrade**

1. In the **Run** dialog box, type **regedit**, and then press ENTER.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters
3. Click **Edit**, click **New**, and then click **DWORD Value**.
4. For the new entry name, type **NT4Emulator**, and then press ENTER.
5. Double-click the entry name that you typed in the previous step.
6. In the **Edit DWORD Value** dialog box, type **1** in the **Value data:** box, and then click **OK**.

Keep the registry editor open to perform the next step in the upgrade process.

# Neutralize Windows NT 4.0 Domain Controller Emulation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

On all additional domain controllers, you must be sure to neutralize Windows NT 4.0 emulation before installing Active Directory. When deploying additional domain controllers, the computers must be able to contact an Active Directory domain controller during the installation of Active Directory. If the Active Directory domain controllers that you have already upgraded have been configured to protect against domain controller overload by setting the value of the NT4Emulator registry entry to one, the additional domain controllers will only recognize them as Windows NT 4.0–based domain controllers and the Active Directory installation will fail.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference.

**To neutralize Windows NT 4.0 emulation**

1. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters
2. Click **Edit**, click **New**, and then click **DWORD Value**.
3. For the new entry name, type **NeutralizeNT4Emulator**, and then press ENTER.
4. Double-click the entry name that you typed in the previous step.
5. In the **Edit DWORD Value** dialog box, type **1** in the **Value data:** box, and then click **OK**.
6. Click **Registry**, and then click **Exit** to close the registry editor.

You can also administer Windows Server 2003–based domain controllers that have been configured to emulate a Windows NT 4.0–based domain controller from a workstation running Windows 2000 or Windows XP. If you intend to use a management workstation that is running Windows 2000 or Windows XP, you must first neutralize the emulation mode on the management workstation so that Windows Server 2003–based domain controllers will respond to it.

# Upgrade Windows NT 4.0 BDCs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can upgrade any Windows NT 4.0 BDC to a Windows Server 2003–based domain controller as long as it meets the hardware requirements for a domain controller running Windows Server 2003. For more information about the hardware requirements for Windows Server 2003–based domain controllers, see "Planning Domain Controller Capacity" in this book.

Before upgrading the operating system to Windows Server 2003, use the Winnt32.exe command-line tool to detect any upgrade problems. This tool reports potential upgrade problems, such as inadequate hardware resources or compatibility problems.

**To determine potential upgrade problems**

- At the command line, connect to the I386 directory located at your installation source and type the following command:

```
winnt32 /checkupgradeonly
```

Resolve reported problems before performing the upgrade.

To install the operating system on the computer, insert the Windows Server 2003 operating system CD in the CD-ROM drive of the domain controller and select the option to install the operating system, or use an automated installation method. If the Windows Server 2003 media is shared on the network, run the Winnt32.exe command.

To complete the operating system installation, perform these tasks:

1. Select **Upgrade** for the Installation type.

2. Use NTFS to convert the partitions. The installation of Active Directory will not succeed if you do not have at least one NTFS partition available on which to locate the SYSVOL shared folder.

3. Verify that you are using a static IP address.

4. On the first additional domain controller upgraded, configure DNS client settings by using the IP address of the PDC for the **Preferred DNS server** setting and do not specify an IP address in the **Alternate DNS server** setting.

   On all remaining domain controllers that are upgraded, configure DNS client settings by using the IP address of the PDC for the **Preferred DNS server** setting and use the IP address of the second domain controller upgraded for the **Alternate DNS server** setting.

   These DNS client settings are temporary and will be changed during the installation of Active Directory.

5. Install Windows Support Tools, which are available in the \Support\Tools folder on the Windows Server 2003 operating system CD.

During the operating system upgrade the computer will restart three times. After the computer restarts for the last time, the **Welcome to the Active Directory Installation Wizard** appears.

# Install Active Directory on the Additional Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process for installing Active Directory on additional domain controllers is identical whether you upgraded to a regional domain controller in an existing domain or upgraded to a single domain forest. After upgrading the operating system on a Windows NT 4.0 additional domain controller to Windows Server 2003, the computer is in an intermediate state, meaning that the computer is no longer a Windows NT 4.0–based domain controller, nor is it a Windows Server 2003–based member server or domain controller.

The Active Directory Installation Wizard allows you to create an additional domain controller or a member server in the new domain. If you will be installing Active Directory by replicating the directory data over the network or from another media source, select the **Member Server** option in the Active Directory Installation Wizard. Selecting **Member Server** will configure the computer to be a Windows Server 2003–based member server, allowing you to install Active Directory at a later time.

**To install Active Directory on a Windows Server 2003–based member server**

- At the command Line, type **Dcpromo**.

  – or –

  Open **Administrative Tools**, and then click **Configure Your Server Wizard**. Select **Domain Controller (Active Directory)** to configure your domain controller. After the Configure Your Server Wizard finishes, the Active Directory Installation Wizard begins.

You can also install Active Directory by using the install from media feature, new in Windows Server 2003. Install from media allows you to pre-populate Active Directory with System State data backed up from an existing Windows Server 2003–based domain controller. This backup can be present on local CD, DVD, or hard disk partition. Installing from media drastically reduces the time required to install directory data by reducing the amount of data that is replicated over the network. Installing from media is most beneficial in environments with very large domains or for installing new domain controllers that are connected by a slow network link.

**To install Active Directory on a Windows Server 2003–based member server from media**

- Type **dcpromo /adv** in the **Run** dialog box.

  The wizard prompts you to choose a network share or a backup as the installation source. If you are installing from backup files, you must identify the location of the files. If the domain controller from which you restored the System State data was a global catalog server, you will have the option make this new domain controller a global catalog server. The wizard will then proceed with the installation.

For more information about installing and removing Active Directory, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

Table 8.13 lists information for installing Active Directory on additional domain controllers, as well as sample data for installing Active Directory on additional domain controllers in a regional domain in the existing Trey Research forest or in the Fabrikam single domain forest. Trey Research will install Active Directory immediately after upgrading the operating system. Fabrikam will use the **dcpromo /adv** command to install Active Directory on a member server by copying directory data over the network from a domain controller.

**Table 8.13 Installing Active Directory on Additional Domain Controllers**

| Wizard Page or Dialog Box | Action | Example |
|---|---|---|

| | | |
|---|---|---|
| Additional Domain Controller or Member Server | Select whether you want the computer to become a member server or an additional domain controller for the domain. | **Upgrading to a regional domain in an existing forest:**<br><br>Trey Research will select Additional domain controller to install Active Directory immediately.<br><br>**Upgrading to a single domain forest:**<br><br>Fabrikam will select Member Server. They will install Active Directory at a later time using the **dcpromo /adv** command. |
| Domain Controller Type | Select **Additional domain controller for an existing domain**. | **Upgrading to a regional domain in an existing forest:**<br><br>Trey Research will not see this wizard page.<br><br>**Upgrading to a single domain forest:**<br><br>When Fabrikam initiates the Active Directory Installation Wizard by using the **dcpromo /adv** command, this is the first wizard page that appears. |
| Copying Domain Information | Select either:<br><br>• Over the network from a domain controller<br><br>• From these restored backup files | **Upgrading to a regional domain in an existing forest:**<br><br>Trey Research will not see this wizard page because they chose to install Active Directory immediately following the operating system upgrade.<br><br>**Upgrading to a single domain forest:**<br><br>Fabrikam will copy domain information from the first domain controller that is deployed, SEA-FAB-DC01,which is in the same location as the new one. Therefore, they selected **Over the network from a domain controller** to copy the information in the shortest time. |
| Network Credentials | Type the user name and password of an account with sufficient privileges to install Active Directory on this computer, and the fully qualified domain name of the domain in which the computer will become an additional domain controller. | |
| Additional Domain Controller | Type the full DNS name of the forest root domain. | **Upgrading to a regional domain in an existing forest:**<br><br>Trey Research will not see this wizard page. It appears only if you are installing Active Directory over the network from a domain controller.<br><br>**Upgrading to a single domain forest:**<br><br>**Fabricorp.fabrikam.com** |
| Database and Log Folders | Type the folder locations specified by your design. | Database folder: **C:\Windows\NTDS**<br><br>Log folder: **D:\Logs** |
| Shared System Volume | Confirm or type the location specified by your design. | **C:\Windows\SYSVOL** |
| Directory Service Restore Mode Administration Password | In the **Password and Confirm password** boxes, type any strong password. | |

Verify that all information on the Summary page is accurate, and then click Finish. After the Active Directory Installation Wizard

finishes, you will be prompted to restart the computer. The installation will not be complete until the computer restarts.

**Note**

- All additional domain controllers added to a single domain forest should be configured as Global Catalog servers. For more information about global catalog server placement, see "Designing the Site Topology" in this book.

After you install Windows Server 2003 Active Directory, enable Remote Desktop for Administration, formerly known as Terminal Services in Remote Administration mode, to enable administrators to log on remotely if necessary. To enable Remote Desktop for Administration, in Control Panel, double-click **System**, select the **Remote** tab, and then select **Allow users to connect remotely to this computer**.

# Install DNS on Additional Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Install DNS on all Windows Server 2003–based domain controllers that are added to the domain.

**To install DNS on additional domain controllers using the Windows Components Wizard**

1. Click **Start**, point to **Settings**, and click **Control Panel**.
2. Double click **Add or Remove Programs**, and then click **Add/Remove Windows Components**.
3. In **Components**, select the **Networking Services** check box, and then click **Details**.
4. In **Subcomponents of Networking Services**, select the **Domain Name System (DNS)** checkbox, click **OK**, and then click **Next**.
5. If prompted, in **Copy files from**, type the full path to the distribution files and then click **OK**. The required files will be copied to your hard disk.

# Reconfigure the Domain Name System (DNS) Service

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After deploying additional domain controllers in either a new regional domain in an existing forest or in a single domain forest, do the following to reconfigure the DNS service:

- Configure the DNS client settings of the first and subsequent domain controllers

  After you have deployed an additional domain controller, modify the DNS client settings on the first domain controller. Because no other domain controllers were running when you deployed the first domain controller, modify the DNS client settings on the first domain controller to include the additional domain controller. As you deploy more domain controllers, you might also need to modify the Alternate DNS server setting specified on existing domain controllers to ensure that this setting points to the closest DNS server.

- Update the DNS delegation

  If you have delegated the DNS zone to an existing DNS server, update the DNS delegation for the domain after you install the DNS Server service on new domain controllers.

# Add Windows NT 4.0 BDCs to Windows Server 2003 Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have applications in your environment that can run only on a Windows NT 4.0–based domain controller and if all the Windows NT 4.0 BDCs have been upgraded to Windows Server 2003 or the existing Windows NT 4.0 BDC in your environment becomes unavailable, you might need to add an additional Windows NT 4.0 BDC to your environment. You can do this by installing a new Windows NT 4.0 BDC in the domain. Prior to installing the new Windows NT 4.0 BDC in the domain, you must first add the new computer account to the Windows Server 2003 domain.

**Note**

- You will not be able to install a new Windows NT 4.0–based BDC in your environment if you have SMB packet signing and secure channel signing enabled. If these security policies are enabled in your environment, modify them before installing a new Windows NT 4.0–based BDC. For information about modifying security policies, see "Modify Security Policies" earlier in this chapter.

**To add a Windows NT 4.0 BDC to a Windows Server 2003 domain**

1. In Active Directory Users and Computers, right-click the **Domain Controllers** folder.
2. Point to **New**, and then click **Computer**.
3. Type the computer name of the BDC.
4. Ensure that the checkboxes are selected for **Assign this computer account as a pre-Windows 2000 Computer** and **Assign this computer account as a backup domain controller**.
5. Install the BDC in to the domain.

# Perform Post-Upgrade Tests on Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After each additional domain controller is deployed, verify that the upgrade was successful. Use the same tests and tools that you used to verify that the upgrade of the Windows NT 4.0 PDC was successful. For more information about developing a test plan, see "Develop a Test Plan" earlier in this chapter.

Also verify that DNS recursive name resolution is configured according to your organization's DNS design. For more information about verifying recursive name resolution, see "Verify DNS Server Recursive Name Resolution on the First Forest Root Domain Controller" earlier in this chapter.

# Complete Post-Upgrade Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you upgrade all domain controllers in the domain to Windows Server 2003, complete the post-upgrade tasks. These tasks are the final step in the process for upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory, as shown in Figure 8.10.

**Figure 8.10 Completing Post-Upgrade Tasks**



# Eliminate Anonymous Connections to Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you upgrade all the servers in the domain hosting services that run as Local System and use Anonymous or null credentials when accessing a domain controller, such as Windows NT 4.0 RAS servers, remove the Everyone and Anonymous Logon groups from the Pre-Windows 2000 Compatible Access built-in group. This task increases the security of your domain by preventing anonymous connections to domain controllers.

**To remove groups from the Pre-Windows 2000 Compatible Access Group using the command line**

- At a command prompt, type:

```
net localgroup "Pre-Windows 2000 Compatible Access" GroupName /delete
```

When using the **net localgroup** command to add or delete any group or group member name that includes spaces, such as the Anonymous Logon group, you must enclose the group name in quotation marks.

# Raise Your Domain and Forest Functional Levels

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although the Windows Server 2003 domain functional level provides a number of features and advantages, enable this functional level only when all your Windows NT 4.0 BDCs have been upgraded and you are certain that your environment is ready.

**Important**

- If you raise the domain and forest functional levels to Windows Server 2003, this action cannot be reversed and you cannot add Windows NT 4.0–based or Windows 2000–based domain controllers to the environment. Any existing Windows NT 4.0 or Windows 2000–based domain controllers in the environment will no longer function. Before you raise functional levels to take advantage of advanced Windows Server 2003 features, ensure that you will never need to install domain controllers that run Windows NT 4.0 or Windows 2000 in your environment.

After you determine that your environment is ready, use Active Directory Domains and Trusts to enable the Windows Server 2003 domain functional level.

After all domains are operating at the Windows Server 2003 functional level, raise the forest functional level to Windows Server 2003 to take advantage of all Windows Server 2003 forest-level features.

For more information about enabling functional levels and the features available at the Windows Server 2003 domain and forest functional levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.

# Redirect the Users and Computers Containers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The default CN=Users and CN=Computers containers that are created when Active Directory is installed are not organizational units (OUs). Objects in the default containers are more difficult to manage because Group Policy cannot be applied directly to them. New user accounts, computer accounts, and security groups that are created by using earlier versions of user interface and command-line management tools, such as the net user and net computer commands, the net group command, the netdom add command where the /ou command is either not specified or supported, or Windows NT 4.0 tools such as User Manager for Domains, do not allow administrators to specify a target organizational unit and therefore create these objects in either the CN=Computers container or the CN=User container by default.

It is recommended that administrators who upgrade Windows NT 4.0–based and Windows 2000–based domain controllers to Windows Server 2003 redirect the well-known path for the CN=Users and CN=Computers to an OU that is specified by the administrator so that Group Policy can apply to containers that are hosting newly created objects.

**Important**

- The CN=Users and CN=Computers containers are computer-protected objects. You cannot (and must not) remove them for backward compatibility purposes. However, you can rename these objects.

In Windows Server 2003 Active Directory, when the domain functional level has been raised to Windows Server 2003, you can redirect the default CN=Users and CN=Computers containers to organizational units that you specify so that each can support Group Policy, making them easier to manage.

**To redirect the Users container**

1. In Active Directory Users and Computers, create an organizational unit container to which you will redirect users that were created with earlier versions of user interface and command-line management tools.
2. At the command line, change to the system32 directory by typing:

```
Cd %systemroot%\system32
```

3. At the %systemroot%\system32 directory, type:

```
Redirusr ou=
newuserou
,DC=
domainname
,dc=com
```

**To redirect the Computers container**

1. In Active Directory Users and Computers, create an organizational unit container to which you will redirect computer objects that were created with earlier versions of user interface and command-line management tools.
2. At the command line, change to the system32 directory by typing:

```
Cd %systemroot%\system32
```

3. At the %systemroot%\system32 directory, type:

```
Redircmp ou=
newcomputerou
```

```
,DC=
domainname
,dc=com
```

For more information about creating an organizational unit design, see "Designing the Active Directory Logical Structure" in this book.

# Complete the Upgrade

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Complete the following tasks to finalize the process:

- Review, update, and document the domain architecture to reflect any changes that you made during the in-place domain upgrade process.
- Review your operating procedures and administrative tasks to determine whether new Windows Server 2003 features, such as Group Policy objects or distributed administration, affect the operations environment. Be sure to document any changes that you identify.
- Remove the FRS script from the domain controller that you scheduled to provide the daily script export to an LMRepl server.
- After you ensure that your Windows Server 2003 Active Directory environment is operating successfully for a period of time, you can redeploy the rollback server that you reserved for the recovery process. If you do not need the Windows NT 4.0 BDC to achieve the required load balance among your domain controllers, maintain the rollback server for one week. Maintain the backup of the rollback server for a longer period of time to be safe.
- Some Windows NT 4.0 applications, such as Microsoft® Systems Management Server (SMS), can have an unpredictable effect on the domain when installed after the domain has been upgraded to Active Directory. Ensure that you are running SMS 2.0 and have installed Service Pack 4. For more information about SMS, see the SMS Downloads link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

After the above tasks have been completed successfully, you will have completed the in-place upgrade process.

# Additional Resources for Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book for more information about restructuring domains when upgrading from Windows NT 4.0 to Windows Server 2003.
- "Designing the Active Directory Logical Structure" in this book for more information about the Active Directory logical structure.
- "Designing the Site Topology" in this book for more information about Active Directory site topology.
- "Enabling Advanced Windows Server 2003 Active Directory Features" in this book for more information about enabling functional levels.
- "Deploying DNS" in *Deploying Network Services* for more information about deploying DNS.

**Related Tools**

- Adsiedit.exe

  The ADSI Edit tool (Adsiedit.exe) is a Microsoft Management Console snap-in that you can use to edit objects in the Active Directory database. For more information about Adsiedit.exe, in Help and Support Center for Windows Server 2003, click **Tools,** and then click Install Windows Support Tools.

- Ldp.exe

  Ldp.exe provides an interface to perform LDAP operations against Active Directory. For more information about Ldp.exe, in Help and Support Center for Windows Server 2003, click **Tools,** and then click Install Windows Support Tools.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Active Directory" in Help and Support Center for Windows Server 2003.
- "Install Windows Support Tools" under "Tools" in Help and Support Center for Windows Server 2003.
- "Configure Site Settings" in Help and Support Center for Windows Server 2003 for more information about creating site objects, subnet objects, and associating subnets with sites.
- "Understanding aging and scavenging" in Help and Support Center for Windows Server 2003 for more information about how to configure aging and scavenging of stale resource records.

**Related Job Aids**

- "Windows NT 4.0 Domain Controllers and Services" (DSSUPNT_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows NT 4.0 Domain Controllers and Services" on the Web at http://www.microsoft.com/reskit).

- "Windows NT 4.0 Hardware Configuration" (DSSUPNT_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows NT 4.0 Hardware Configuration" on the Web at http://www.microsoft.com/reskit).

- "Windows NT 4.0 Network Configuration" (DSSUPNT_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows NT 4.0 Network Configuration" on the Web at http://www.microsoft.com/reskit).

- "Windows NT 4.0 Domain Controller Role Assignment" (DSSUPNT_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows NT 4.0 Domain Controller Role Assignment" on the Web at http://www.microsoft.com/reskit).

- "Windows NT 4.0 Domain Controller Documentation" (DSSUPNT_5.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows NT 4.0 Domain Controller Documentation" on the Web at http://www.microsoft.com/reskit).

# Upgrading from Windows 2000 Domains to Windows Server 2003 Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Upgrading your network operating system from Microsoft® Windows® 2000 to Windows® Server 2003 requires minimal network configuration and typically has a low impact on user operations. The upgrade process is straightforward, efficient, and allows your organization to take advantage of the improved security that is offered by Windows Server 2003.

**In This Chapter**


Overview of Upgrading Your Windows 2000 Domains to Windows Server 2003 Domains

Planning to Upgrade Windows 2000 Domains to Windows Server 2003 Domains

Completing Pre-Upgrade Tasks

Upgrading Windows 2000 Domains to Windows Server 2003

Completing Post-Upgrade Tasks

Additional Resources for Upgrading from Windows 2000 Domains to Windows Server 2003 Domains


**Related Information**

- For more information about designing the Active Directory® directory service logical structure and the DNS infrastructure needed to support Active Directory, see "Designing the Active Directory Logical Structure" in this book.
- For more information about Active Directory functional levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.
- For more information about upgrading from Microsoft® Windows NT® version 4.0 to Windows Server 2003 Active Directory, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.
- For more information about deploying a DNS infrastructure for name resolution on your network, see "Deploying DNS" in *Deploying Network Services* of this kit.

# Overview of Upgrading Your Windows 2000 Domains to Windows Server 2003 Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By upgrading your network operating system from Microsoft® Windows® 2000 Server to the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; or Windows® Server 2003, Datacenter Edition operating system, you can maintain your current network and domain configuration while improving the security, scalability, and manageability of your network infrastructure.

Prior to upgrading your Windows 2000 domains, review your business objectives and decide how they relate to your existing Active Directory infrastructure. Although your objectives might not require other significant changes to your existing environment, the operating system upgrade is an opportune time to review your existing Active Directory design, including your Active Directory logical structure, site topology, and domain controller capacity. You might find opportunities for increased efficiencies and cost savings that you can incorporate into your upgrade process.

Additionally, ensure that you test your upgrade process in a lab and pilot program. For more information about lab testing and piloting, see "Planning an Active Directory Deployment Project" in this book.

When the domain upgrade process is complete, all domain controllers will be running Windows Server 2003, and the Active Directory domain and forest will be operating at the Windows Server 2003 functional level. At the Windows Server 2003 forest functional level, you can take advantage of all advanced Active Directory features. For more information about advanced Active Directory features related to Active Directory functional levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.

**Note**

- For a list of the job aids that are available to assist you in upgrading from Windows 2000 Server to Windows Server 2003, see "Additional Resources for Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" later in this chapter.

# Process for Upgrading Windows 2000 Domains to Windows Server 2003 Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Upgrading your Windows 2000 Active Directory environment to a Windows Server 2003 Active Directory environment involves completing several tasks. Figure 9.1 shows the tasks in the upgrade process.

**Figure 9.1 Upgrading Windows 2000 Domains to Windows Server 2003 Domains**



# Background Information for Upgrading Windows 2000 Domains to Windows Server 2003 Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin the process of upgrading your Windows 2000 Active Directory environment to Windows Server 2003 Active Directory, become familiar with some important issues that affect the upgrade process.

**Active Directory Preparation Tool**
To prepare your Windows 2000 forest and domains for upgrade to Windows Server 2003 Active Directory, or for the introduction of a new Windows Server 2003–based domain controller, you must use the Active Directory Preparation tool (Adprep.exe). Adprep.exe is located on the Windows Server 2003 and Windows Server 2003 Service Pack 1 (SP1) operating system CD.

Adprep.exe prepares the forest and domains for an Active Directory upgrade by performing a collection of operations prior to the installation of the first Windows Server 2003–based domain controller, including:

- Extending your current schema with new schema information that Adprep.exe provides, while preserving previous schema modifications in your environment.
- Resetting permissions on containers and objects throughout the directory for improved security and interoperability with new Windows Server 2003 domains.
- Copying administrative tools to manage new Windows Server 2003 domains to the local computer.

### Note

- Changes made by Adprep.exe do not affect the functioning of Windows NT 4.0–based or Windows 2000–based domain controllers.

For more information about using Adprep.exe to prepare your environment, see "Prepare Your Infrastructure for Upgrade" later in this chapter.

**Application Directory Partitions for DNS**
Application directory partitions provide storage for application-specific data that can be replicated to a specific set of domain controllers in the same forest. If you have at least one domain controller in your forest running Windows Server 2003 and the domain naming master is also running Windows Server 2003, you can take advantage of application directory partitions.

For example, you can use application directory partitions to store DNS data on Windows Server 2003–based domain controllers. DNS-specific application directory partitions are automatically created in the forest and in each domain when the DNS service is installed on new or upgraded Windows Server 2003–based domain controllers. If application directory partition creation fails during Active Directory installation, DNS attempts to create the partitions again every time that the service starts.

### Note

- The creation and deletion of application directory partitions, including the default DNS application directory partitions, requires that the domain naming master role holder reside on a Windows Server 2003–based domain controller.

The following DNS-specific application directory partitions are created during Active Directory installation:

- ForestDnsZones — A forest-wide application directory partition shared by all DNS servers in the same forest
- DomainDnsZones — Domain-wide application directory partitions for each DNS server in the same domain

**SRV resource records**
A Windows Server 2003–based domain controller's Net Logon service uses dynamic updates to register SRV resource records

in the DNS database, as described in "A DNS RR for specifying the location of services (DNS SRV)." For more information about this draft, see the Internet Engineering Task Force (IETF) web page. This SRV record is used to map the name of a service, such as the Lightweight Directory Access Protocol (LDAP) service, to the DNS computer name of a server that offers that service. In a Windows Server 2003 network, an LDAP resource record locates a domain controller. A workstation that is logging on to a Windows Server 2003 domain queries DNS for SRV records in the general form:

*_Service* **.**_ *Protocol* **.** *DnsDomainName*

where *Service* is the service requested, *Protocol* is the protocol requested, and *DnsDomainName* is the fully qualified DNS name of the Active Directory domain.

Active Directory servers offer the LDAP service over the TCP protocol; therefore, clients find an LDAP server by querying DNS for a record of the form:

**_ldap._tcp.** *DnsDomainName*

**Note**

- The service and protocol strings require an underscore (_) prefix to prevent potential collisions with existing names in the namespace.

This format is applicable for implementations of LDAP servers other than Windows Server 2003–based domain controllers and also possible implementations of LDAP directory services that employ Global Catalog servers other than servers running Windows Server 2003.

**_msdcs.domain_name subdomain**
This Microsoft-specific subdomain allows location of domain controllers that have Windows Server 2003–specific roles in the domain, as well as the location by globally unique identifier (GUID) when a domain has been renamed.

To facilitate location of Windows Server 2003–based domain controllers, the Net Logon service in addition to the standard *_Service._Protocol.DnsDomainName* format records , also registers SRV records that identify the well-known server-type pseudonyms "dc" (domain controller), "gc" (Global Catalog), "pdc" (primary domain controller), and "domains" (GUID) as prefixes in the _msdcs.domain_name subdomain. To accommodate locating domain controllers by server type or by GUID (abbreviated "dctype"), Windows Server 2003–based domain controllers register SRV records in the following form in the _msdcs.*domain_name* subdomain:

*_Service._Protocol.DcType.* **_msdcs** *.DnsDomainName*

**_msdcs.forest_root_domain subdomain**
The _msdcs.*forest_root_domain* subdomain stores forest-wide resource records that are of interest to clients and domain controllers from all parts of the forest. For example, all domain controllers in the forest register CNAME and LDAP, Kerberos, and GC SRV resource records in the msdcs.*forest_root_domain* subdomain. The CNAME resource records are used by the replication system to locate replication partners and the GC SRV resource records are used by clients to lookup global catalog servers.

For any two domain controllers to replicate with each other, including two domain controllers from the same domain, they must be able to look up forest-wide locator records. For a newly created domain controller to participate in replication, it must be able to register its forest-wide records in DNS, and other domain controllers must be able to look up these records. Therefore, the DNS servers that are authoritative for the _msdcs.*forest_root_domain* subdomain needs to be available for replication and global catalog lookups.

For this reason, it is recommended that you create a separate _msdcs.*forest_root_domain* zone and define its replication scope so that it is replicated to all DNS servers in the forest. For more information about creating a separate _msdcs.*forest_root_domain* zone, see KB article 817470. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Some organizations running Windows 2000 Active Directory already created an _msdcs.*forest_root_domain* to help clients locate domain controllers more efficiently. If an _msdcs.*forest_root_domain* already exists in your Windows 2000 environment, then it is recommended that you move the zone to the ForestDnsZones application directory partition after all domain controllers in the forest are running Windows Server 2003. In addition, for each domain in the forest, move the _msdcs.*domain_name* zone to the DomainDnsZones application directory partition for that domain.

Moving the Active Directory–integrated DNS zones into the domain and forest-wide application directory partitions provides the following benefits:

- Because the forest-wide application directory partition can replicate outside a specified domain, and because moving the _msdcs.*forest_root_domain* into the forest-wide application directory partition replicates it to all domain controllers in the forest that are running the DNS service, you do not have to use DNS zone transfer to replicate the zone file information to DNS servers outside the domain.

- Domain-wide replication can be targeted to minimize replication traffic because administrators can specify which of the domain controllers running the DNS service receive the DNS zone data.

- Forest-wide replication can be targeted to minimize replication traffic because DNS data is no longer replicated to the global catalog.

- DNS records located on global catalog servers in the forest are removed, minimizing the amount of information replicated with the global catalog.

For more information about using application directory partitions to store DNS data, see "Use DNS Application Directory Partitions" later in this chapter.

**Intrasite Replication Frequency**

Windows 2000 domain controllers that are upgraded to Windows Server 2003 maintain their default intrasite replication frequency of 300/30. This means that any changes made to Active Directory replicate to all other domain controllers in the same site five minutes (300 seconds) after a change is made, with a 30-second offset before notifying the next domain controller, until the forest functional level is raised to Windows Server 2003. When the forest functional level is raised to Windows Server 2003, the replication frequency of Active Directory is changed to the Windows Server 2003 default setting of 15/3. This means that changes will replicate to all domain controllers in the same site 15 seconds after a change is made, with a 3-second offset before notifying the next domain controller. If you modified the 300/30 default replication frequency setting in Windows 2000, the setting does not change to the 15/3 default setting in Windows Server 2003 after you complete the upgrade. However, a new installation of Windows Server 2003 will always use the 15/3 intrasite replication frequency setting.

**Important**

- Do not modify the default 300/30 intrasite replication frequency on Windows 2000 domain controllers. Instead, upgrade your Windows 2000 domain to Windows Server 2003 and raise the forest functional level to Windows Server 2003 to take advantage of the 15/3 intrasite replication frequency.

**New Groups and New Group Memberships Created After Upgrading the PDC**

After upgrading the Windows 2000–based domain controller holding the role of the PDC emulator in each domain in the forest to Windows Server 2003, several new well-known and built-in groups are created and some new group memberships are established. If you transfer the PDC emulator role to a Windows Server 2003–based domain controller instead of upgrading it, these groups will be created when the role is transferred. The new well-known and built-in groups are:

- Builtin\Remote Desktop Users

- Builtin\Network Configuration Operators

- Performance Monitor Users

- Performance Log Users

- Builtin\Incoming Forest Trust Builders

- Builtin\Performance Monitoring Users

- Builtin\Performance Logging Users

- Builtin\Windows Authorization Access Group

- Builtin\Terminal Service License Server

The newly established group memberships are:

- If the Everyone group is in the Pre-Windows 2000 Compatible Access group, the Anonymous Logon group and Authenticated Users group are also added to the Pre-Windows 2000 Compatible Access group.
- The Network Servers group is added to the Performance Monitoring alias.
- The Enterprise Domain Controllers group is added to the Windows Authorization Access group.

In addition, when upgrading the Windows 2000 domain controller that holds the role of the PDC emulator in the forest root domain, the following additional security principals are created:

- LocalService
- NetworkService
- NTLM Authentication
- Other Organziation
- Remote Interactive Logon
- SChannel Authentication
- This Organization

For more information about new well-known and built-in groups in Windows Server 2003, see "Default groups" in Help and Support Center for Windows Server 2003.

**Security Policy Considerations When Upgrading from Windows 2000 to Windows Server 2003**
Server message block (SMB) packet signing and secure channel signing are security policies enabled by default on Windows Server 2003–based domain controllers. To allow clients running earlier versions of Windows to communicate with domain controllers running Windows Server 2003, you might need to temporarily disable these security policies during the upgrade process.

**SMB packet signing**
SMB packet signing is a security mechanism that protects the data integrity of SMB traffic between client computers and servers, and prevents man-in-the-middle attacks by providing a form of mutual authentication. This is done by placing a digital security signature into each SMB packet, which is then verified by the receiving party. Server-side SMB signing is required by default on Windows Server 2003–based domain controllers, which means that all clients are required to have SMB packet signing enabled.

Clients running Windows NT 4.0 with Service Pack 2 (SP2) or earlier and clients running Microsoft® Windows® 95 without the Directory Service Client Pack do not support SMB packet signing. These clients will not be able to authenticate to a Windows Server 2003–based domain controller. To ensure successful authentication, upgrade these clients to a later version of the operating system or service pack. However, if you cannot upgrade your clients, you can allow them to be authenticated by configuring SMB packet signing on all Windows Server 2003–based domain controllers so that SMB packet signing is allowed but not required.

For more information about SMB packet signing, see "Microsoft network server: Digitally sign communications (always)" in Help and Support Center for Windows Server 2003.

For more information about configuring SMB packet signing on Windows Server 2003–based domain controllers, see "Modify Security Policies" later in this chapter.

**Secure channel signing and encryption**
When a computer becomes a member of a domain, a computer account is created. Each time the computer starts, it uses the computer account password to create a secure channel with a domain controller for its domain. This secure channel is used to ensure secure communications between a domain member and a domain controller for its domain. Secure channel signing is required by default on Windows Server 2003–based domain controllers, which means that all clients must enable secure channel signing and encryption.

Clients running Windows NT 4.0 with Service Pack 3 (SP3) or earlier installed do not support secure channel signing. These clients will not be able to establish communications with a Windows Server 2003–based domain controller. To ensure successful communication, upgrade these clients to a later version of the operating system or service pack. However, if you cannot upgrade your clients, you must disable secure channel signing on all Windows Server 2003–based domain controllers so that the traffic passing through the secure channel is not required to be signed or encrypted.

For more information about configuring secure channel signing on Windows Server 2003–based domain controllers, see "Modify Security Policies" later in this chapter.

For more information about secure channel signing, see "Domain member: Digitally encrypt or sign secure channel data (always)" in Help and Support Center for Windows Server 2003.

# Planning to Upgrade Windows 2000 Domains to Windows Server 2003 Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Planning to upgrade your Windows 2000 environment to Windows Server 2003 Active Directory involves completing the tasks and procedures that are shown in Figure 9.2.

**Figure 9.2 Planning to Upgrade Windows 2000 Domains to Windows Server 2003 Domains**



# Create a Pre-Upgrade Task Checklist

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can create a pre-upgrade task checklist to help organize the tasks necessary to prepare for a successful domain upgrade. In your checklist, include the tasks listed in the sample checklist in Figure 9.3, in addition to any additional tasks specific to your organization.

For a worksheet to assist you in creating your own pre-upgrade task checklist, see "Pre-Upgrade Task Checklist" (DSSUPWN_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Pre-Upgrade Task Checklist" on the Web at http://www.microsoft.com/reskit).

**Figure 9.3 Example of a Pre-Upgrade Task Checklist**

| Pre-Upgrade Task Checklist | | | |
|---|---|---|---|
| Prepared By | Forest Owner | Date | 1/2/03 |
| Forest Name | Concorp.contoso.com | | |

| Task | Method for Completion | Completed By | Date Completed |
|---|---|---|---|
| Gather job aids | Create new job aids or use job aids found on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit. | | |
| Assign appropriate credentials | Assign Schema Admins and Enterprise Admins credentials to the individual who will prepare the Active Directory forest for upgrade. Assign Domain or Enterprise Admins credentials to the individual who will prepare each Active Directory domain for upgrade. | | |
| Introduce a Windows Server 2003-based member server | Install a Windows Server 2003 based member server into the forest root domain. | | |
| Verify that your upgrade is supported | Review Table 9.2, "Supported Upgrade Paths to Windows Server 2003." | | |
| Assess all current domain controller hardware | Assess your current hardware and review requirements for domain controllers running Windows Server 2003 as specified in "Planning Domain Controller Capacity" in this book. | | |
| Document the name, IP addresses, locations, operations master roles held by each domain controller before and after the upgrade, and the order in which you will upgrade the operating system on each domain controller. | Complete the worksheet "Windows 2000 Domain Controller Documentation" (DSSUPWN_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD. | | |
| Develop a test plan | Create a matrix outlining the tests and tools to be used throughout the domain upgrade process. | | |
| Develop and approve recovery plan | Call team meeting to determine plan for recovery and approve completed recovery plan. | | |

| | | | |
|---|---|---|---|
| Determine required service pack levels | Use the **Repadmin /showattr** command to inventory the operating system and service pack revision level. | | |
| Backup domain data | Back up all data according to the backup plan specified for your organization. | | |
| Resolve upgrade and application compatibility issues | Use the **winnt32 /checkupgradeonly** command to find any upgrade issues. | | |
| Prepare the forest and each domain for Windows Server 2003 Active Directory | Apply hotfixes to resolve schema compatibility issues with Exchange 2000 and Services for UNIX 2.0.<br>Use ADPrep.exe in the forest and in each regional domain. | | |

# Assign Appropriate Credentials

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assign appropriate credentials to the users who are responsible for preparing the forest and domain for an Active Directory upgrade. The **adprep /forestprep** command requires a user account that is a member of the Schema Admins, Enterprise Admins, and Domain Admins groups. The **adprep /domainprep** command requires a user account that is a member of the Domain Admins group in the targeted domain.

Additionally, the security context can affect the ability of an administrator to complete the upgrade from Windows 2000 to Windows Server 2003. Members of the Builtin\Administrators group can upgrade the operating system and install software on a computer. The following groups are members of the Builtin\Administrators group by default:

- The Enterprise Admins group is a member of Builtin\Administrators in the forest root domain and in each regional domain in the forest.
- The Domain Admins group is a member of Builtin\Administrators in their domain.
- The Domain Admins group is a member of Builtin\Administrators on member servers in their domain.

Table 9.1 shows the credentials that are required to upgrade servers, depending on the domain membership of the servers.

**Table 9.1 Credentials Required to Upgrade Servers to Windows Server 2003**

| Credential | Domain Controller in Forest Root Domain | Member Server in Forest Root Domain | Domain Controller in Regional Domain | Member Server in Regional Domain |
|---|---|---|---|---|
| Enterprise Admins in forest root domain | ● | | ● | |
| Domain Admins in forest root domain | ● | ● | | |
| Builtin\Administrators in forest root domain | ● | | | |
| Domain Admins in regional domain | | | ● | ● |
| Builtin\Administrators in regional domain | | | ● | |

You also need to ensure that the administrator who is upgrading the domain controllers has the following rights:

- Backup files and directories (SE_BACKUP_NAME)
- Modify firmware environment values (SE_SYSTEM_ENVIRONMENT_NAME)
- Restore files and directories (SE_RESTORE_NAME)
- Shut down the system (SE_SHUTDOWN_NAME)

The setup program cannot run properly if these rights are not defined, or if they are disabled by a domain Group Policy setting on the computer.

**To verify if user rights assignments are disabled by a domain Group Policy setting**

1. In the **Run** dialog box, type **mmc**, and then click **OK**.
2. Click **File**, and then click **Add/Remove snap-in**.
3. In the Add/Remove snap-in dialog box, click **Add**.
4. In the Available Standalone snap-ins dialog box, select **Group Policy**, and then click **Add**.
5. At the Welcome to the Group Policy Wizard screen, verify that **Local Computer** appears in the **Group Policy Object** box, and then click **Finish**.
6. Close the Add/Remove snap-in dialog box and the Add Standalone snap-in dialog box.
7. In the Console Root, navigate to the Local Computer Policy\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment folder.
8. In the details pane, verify that the user who will perform the upgrade is a member in one of the groups that has the necessary rights assigned. The policies are named identically to the user rights listed above.

Assign the appropriate credentials in advance to allow both testing and deployment to proceed without unexpected security delays.

# Introduce a Windows Server 2003-Based Member Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin the domain upgrade process, introduce a Windows Server 2003–based member server into your environment. Installing Active Directory on a Windows Server 2003–based member server facilitates the domain upgrade process by allowing all existing services to run uninterrupted while you are upgrading to Windows Server 2003 Active Directory.

You can introduce the member server to any domain in the forest; however, if your forest root domain is a dedicated root, it is recommended that you introduce the member server into the forest root domain. Placing the member server into a dedicated root domain has the lowest impact on your environment because users generally do not log on to a dedicated forest root domain; therefore, user authentications are minimal.

After you prepare the Windows 2000 forest and domain by using the Active Directory Preparation tool, install Active Directory on the member server, creating an additional domain controller in the existing domain. The Windows Server 2003–based member server will become the first Windows Server 2003–based domain controller in the forest.

# Determine Supported Software Upgrades

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Identify the versions of Windows 2000 that are running in your environment, and then determine whether you can upgrade the operating system on your computers to Windows Server 2003, or whether you must perform a clean operating system installation.

Table 9.2 lists the Windows 2000 operating system platforms and indicates which platforms can be upgraded directly to each version of Windows Server 2003.

**Table 9.2 Supported Upgrade Paths to Windows Server 2003**

| Platform | Upgrade to Windows Server 2003, Standard Edition | Upgrade to Windows Server 2003, Enterprise Edition | Upgrade to Windows Server 2003, Datacenter Edition |
|---|---|---|---|
| Microsoft® Windows® 2000 Professional | | | |
| Windows 2000 Server | • | • | |
| Microsoft® Windows® 2000 Advanced Server | | • | |
| Microsoft® Windows® 2000 Datacenter Server | | | • |

# Assess Hardware Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review and document the existing hardware configuration of each domain controller that you plan to upgrade to Windows Server 2003. Use this information to identify the domain controllers in your environment that you can upgrade to Windows Server 2003 and the domain controllers that do not meet the hardware requirements necessary to run Windows Server 2003. You can retain domain controllers that do not meet the necessary hardware requirements to serve as rollback servers in the event that you must roll back your deployment. In most cases, a Windows 2000–based domain controller meets the requirements to be upgraded to Windows Server 2003, as long as it has adequate disk space.

At minimum, a domain controller requires available free disk space for the Active Directory database, Active Directory log files, SYSVOL, and the operating system. Use the following guidelines to determine how much disk space to allot for your Active Directory installation:

- On the drive that will contain the Active Directory database, NTDS.dit, provide 0.4 gigabytes (GB) of storage for each 1,000 users. For example, for a forest with two domains (domain A, domain B), with 10,000 and 5,000 users respectively, provide a minimum of 4 GB of disk space for each domain controller that hosts domain A and a minimum of 2 GB of disk space for each domain controller that hosts domain B. Available space must equal at least 10 percent of your existing database size, or at least 250 megabytes (MB), whichever is greater.

- On the drive containing the Active Directory log files, provide at least 500 MB of available space.

- On the drive containing the SYSVOL shared folder, provide at least 500 MB of available space.

- On the drive containing the Windows Server 2003 operating system files, to run setup, provide at least 1.25 GB to 2 GB of available space.

For more information about assessing the hardware requirements of domain controllers in a Windows Server 2003 domain, see "Planning Domain Controller Capacity" in this book.

# Determine Domain Controller Upgrade Order

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Determine the order in which you will upgrade your domain controllers before beginning the domain upgrade process. Record the name, IP address, the domain in which the domain controller will be located, and the operations master roles held by each domain controller before and after the upgrade. Finally, record the order in which you will upgrade the operating system on each domain controller.

The recommended order for upgrading domain controllers from Windows 2000 to Windows Server 2003 is:

- Install Active Directory on a Windows Server 2003–based member server in the forest root domain by using the Active Directory Installation Wizard. This creates the first Windows Server 2003–based domain controller.
- Upgrade the operating system on the Windows 2000–based domain controller holding the role of domain naming master. If you choose not to upgrade the domain controller, transfer the domain naming master role to a domain controller running Windows Server 2003.
- Upgrade the domain controller that holds the PDC emulator role in each domain, or transfer the roles to Windows Server 2003–based domain controllers.
- Continue upgrading all Windows 2000–based domain controllers to Windows Server 2003 until the domain upgrade is complete.

**Note**

- This order for upgrading or installing Windows Server 2003 domain controllers is a recommendation only. It is safe to upgrade the domain controllers holding the domain naming master and PDC emulator roles at any time in the upgrade process.

Use a domain controller documentation table to document information about each domain controller in the forest. For a worksheet to assist you in documenting your domain controller information, see "Windows 2000 Domain Controller Documentation" (DSSUPWN_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows 2000 Domain Controller Documentation" on the Web at http://www.microsoft.com/reskit).

Figure 9.4 shows an example of a completed domain controller documentation table for Contoso.

**Figure 9.4 Example of a Windows 2000 Domain Controller Documentation Table**

| Windows 2000 Domain Controller Documentation | | | |
|---|---|---|---|
| Prepared By | Forest Owner | Date | 1/2/03 |

| | |
|---|---|
| Forest Name | Concorp.contoso.com |
| Domain Name | Noam.concorp.contoso.com |

| DC Name | IP Address | Location | Role Before Upgrade | Role After Upgrade | Order |
|---|---|---|---|---|---|
| SEA-CON-DC01 | 172.16.4.2 | Forest Root Domain | Member Server | Domain Controller | 1 |
| SEA-CON-DC02 | 172.16.4.3 | Forest Root Domain | Schema Master, Domain Naming Master | Schema Master, Domain Naming Master | 2 |
| SEA-CON-DC03 | 172.16.4.4 | Forest Root Domain | Infrastructure Master, PDC, RID Master | Infrastructure Master, PDC, RID Master | 3 |
| SEA-NOAM-DC01 | 172.16.4.10 | North America Regional Domain | Domain Controller | Domain Controller | 5 |
| SEA-NOAM-DC02 | 172.16.4.11 | North America Regional Domain | Infrastructure Master, PDC, RID Master | Infrastructure Master, PDC, RID Master | 4 |
| SEA-EMEA-DC01 | 172.16.4.20 | Europe/ Middle East/Asia Regional Domain | Domain Controller | Domain Controller | 7 |
| SEA-EMEA-DC02 | 172.16.4.21 | Europe/ Middle East/Asia Regional Domain | Infrastructure Master, PDC, RID Master | Infrastructure Master, PDC, RID Master | 6 |

# Developing a Test Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to develop a plan for testing your domain upgrade procedures throughout the upgrade process. Before you begin, test your existing domain controllers to ensure that they are functioning properly, and continue to test your domain controllers throughout the process to verify that Active Directory replication is consistent and successful.

Many of the tools required to verify your domain upgrade procedures are located in the Support\Tools folder on the Windows Server 2003 operating system CD. Install the Windows Server 2003 support tools on a client computer running Microsoft® Windows® XP Professional or on a Windows Server 2003–based member server.

Table 9.3 lists the tools and log files to use in your test plan.

**Table 9.3 Tools and Logs Used to Test Domain Upgrade Procedures**

| Tool/Log File | Description | Location |
|---|---|---|
| Repadmin.exe | Checks replication consistency and monitors both inbound and outbound replication partners. Displays replication status of inbound replication partners and directory partitions. | Windows Server 2003 operating system CD in the Support\Tools folder. |
| Dcdiag.exe | Diagnoses the state of domain controllers in a forest or enterprise, tests for successful Active Directory connectivity and functionality, and returns the results as passed or failed. | Windows Server 2003 operating system CD in the Support\Tools folder. |
| Netdiag.exe | Diagnoses networking and connectivity problems by performing a series of tests to determine the state of your network client and whether it is functional. | Windows Server 2003 operating system CD in the Support\Tools folder. |
| Nltest.exe | Queries and checks the status of trusts and can forcibly shut down domain controllers. | Windows Server 2003 operating system CD in the Support\Tools folder. |
| Dnscmd.exe | Provides the properties of DNS servers, zones, and resource records. | Windows Server 2003 operating system CD in the Support\Tools folder. |
| Adprep Log | Provides a detailed progress report of the forest and domain preparation process. | %systemroot%\System32\Debug\Adprep folder. |
| Dcpromoui.log and Dcpromo.log | Provides a detailed progress report of the Active Directory installation. Includes information regarding replication and services in addition to applicable error messages. | %systemroot%\Debug folder. |
| Adsiedit.exe | A Microsoft Management Console (MMC) snap-in that acts as a low-level editor for Active Directory and allows you to view, add, delete, and move objects and attributes within the directory. | Windows Server 2003 operating system CD in the Support\Tools folder. |

For more information about Windows Support Tools, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Install Windows Support Tools.

Create a test matrix that meets your needs, based on the services that you require to support your environment. For a worksheet to assist you in documenting your test matrix, see "Windows 2000 Upgrade Test Matrix" (DSSUPWN_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows 2000 Upgrade Test Matrix" on the Web at http://www.microsoft.com/reskit).

Figure 9.5 shows an example of a completed upgrade test matrix.

**Figure 9.5 Example of a Windows 2000 Upgrade Test Matrix**

# Windows 2000 Upgrade Test Matrix

| Prepared By | Forest Owner | Date | 1/2/03 |
|---|---|---|---|

| | |
|---|---|
| Forest Name | Concorp.contoso.com |
| Domain Name | Noam.concorp.contoso.com |

| Test | Expected Results | Tool | Success/ Fail | Notes |
|---|---|---|---|---|
| Container Creation | After you prepare the forest and domain for upgrade, the CN=ForestUpdates container is present in the Configuration partition.<br>Expand CN=ForestUpdates, and verify that the CN=Windows2003Update container has been created.<br>Expand the DomainUpdates container in Active Directory Users and Computers, and verify that the Windows2003Update container was created. | ADSI Edit for ForestUpdates container<br>Active Directory Users and Computers for DomainUpdates container | | |
| Trusts | All trusts are active. | Nltest.exe and Dcdiag.exe | | |
| New User Creation | Create a user called "testl" to ensure that accounts can be created. | Active Directory Users and Computers | | |
| Replication | Successful replication. | Repadmin.exe | | |
| DNS | Successful DNS registration and configuration. | Dnscmd.exe | | |
| Active Directory | Successful Active Directory connectivity and functionality. | Dcdiag.exe | | |
| Successful logon request | Account is valid, and access is granted. | Login dialog box | | |

# Develop a Recovery Plan for Your Domain Upgrade Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Develop a recovery plan for use in the event that some portion of your domain upgrade process fails. A successful recovery plan includes:

- Step-by-step instructions, so that the upgrade team can restore normal operations to the organization.
- A sign-off process, ensuring that all team members review, agree upon, and sign off on the recovery plan.

For more information about developing a recovery plan, see the Active Directory Disaster Recovery link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Completing Pre-Upgrade Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Prior to upgrading your Windows 2000 environment to Windows Server 2003 Active Directory, you must complete several pre-upgrade tasks.

Figure 9.6 shows the process for completing pre-upgrade tasks.

**Figure 9.6 Completing Pre-Upgrade Tasks**



# Determine Service Pack Levels

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before preparing your infrastructure for upgrade and installing the Windows Server 2003 operating system, all Windows 2000–based domain controllers in the forest must be running Windows 2000 Service Pack 1 (SP1) with hotfix 265089, or Windows 2000 SP2 (or later).

Hotfix 265089 (included in Windows 2000 SP2 and later) is required to prevent potential domain controller corruption. Use the **repadmin /showattr** command to inventory the operating system and service pack revision level on all domain controllers in a particular domain.

**Note**

- When administering Windows 2000–based domain controllers from a computer running Windows XP Professional or Windows Server 2003, you might experience interoperability problems with the Windows Server 2003 administrative tools unless your Windows 2000–based domain controllers are running Windows 2000 SP3 or later. Some Windows Server 2003 Active Directory administrative tools sign and encrypt all LDAP traffic. Computers running Windows 2000 SP3 or later can interpret the signed and encrypted LDAP traffic.

**To determine domain controller operating system and service pack levels**

- For each domain in the forest, type the following command at the command line of a computer that has the Windows Server 2003 support tools installed:

  ```
  repadmin /showattr domain controller in target domain ncobj:domain: /filter:"(&
  (objectcategory=computer)(primaryGroupID=516))" /subtree
  /atts:operatingSystem,operatingSystemVersion,operatingSystemServicePack
  ```

  The following text is sample output from this command:

  ```
  DN: CN=NA-DC-01,OU=Domain Controllers,DC=company,DC=com
      1> operatingSystem: Windows Server 2003
      1> operatingSystemVersion: 5.2 (3663)
  DN: CN=NA-DC-02,OU=Domain Controllers,DC=company,DC=com
      1> operatingSystem: Windows 2000 Server
      1> operatingSystemVersion: 5.0 (2195)
      1> operatingSystemServicePack: Service Pack 3
  ```

  **Note**

  - The **repadmin /showattr** command does not show any hotfixes that might be installed on a domain controller.

Upgrade domain controllers to the appropriate service pack as needed. For more information about recommended hotfixes to use with Service Pack 2, see article 331161, "List of Fixes to Use on Windows 2000 Domain Controllers Before You Run the Adprep/Forestprep Command" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Back Up Domain Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Back up your Windows 2000 domain data before you begin the upgrade. This task varies according to the operations and procedures that already exist in your environment. At minimum, complete the following steps:

- To allow for fault tolerance, ensure successful replication between two domain controllers in each domain.
- Back up two domain controllers in each domain in the forest, including System State data.
    - Test all backup media to ensure that the data can be restored successfully.

**Important**

- Store backup media in a secure off-site location designated by, and accessible to, the upgrade team before you begin the upgrade process.

# Resolve Upgrade and Application Compatibility Problems

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before upgrading a server to Windows Server 2003, use the **Winnt32.exe** command-line tool with the **/checkupgradeonly** parameter to identify potential upgrade problems, such as inadequate hardware resources or compatibility problems.

Two application compatibility problems you might need to resolve include:

- Distributed File System (DFS) root shares are not supported if they are hosted on a file allocation table (FAT) partition.

  In Windows Server 2003, DFS root shares must be located on NTFS partitions with no files or directories under the DFS link.

  For more information about deploying DFS, see "Designing and Deploying File Servers" in *Planning Server Deployments* in this kit.

- Windows 2000–based computers running Remote Installation Services (RIS) might cause errors in a Windows Server 2003 Active Directory domain.

  When using Windows 2000 RIS server in your Windows Server 2003 Active Directory Domain, you might receive the following error when using the Client Installation Wizard (CIW):

  ```
  " Unable to create or Modify Computer account"
  Error: 00004E4F
  ```

  This error occurs because Windows Server 2003 creates machine account objects differently from Windows 2000. To prevent this error from occurring when creating machine accounts, configure the Windows 2000–based RIS servers in your environment to point to a domain controller running Windows 2000. This is done by adding the DefaultServer registry parameter to the Windows 2000 RIS servers.

  For more information about configuring optional registry parameters for the Boot Information Negotiation Layer (BINL) service, see article 235979, "Optional Registry Parameters for the BINL Service" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

  You must remove the Windows 2000 Administration Tools Pack before upgrading to Windows Server 2003. For more information about Windows 2000 administration tools and upgrade issues, see article 304718, "Administering Windows 2000–Based Computers Using Windows XP Professional–Based Clients," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**To identify potential upgrade and compatibility problems**

- At the command line, connect to the I386 directory located at your installation source and type the following command:

  ```
  winnt32 /checkupgradeonly
  ```

Resolve any reported problems prior to performing the upgrade.

# Prepare Your Infrastructure for Upgrade

Updated: February 22, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Preparing your infrastructure for upgrade involves resolving any Active Directory schema compatibility issues, such as issues that might occur with Microsoft® Exchange 2000 and Services for UNIX 2.0, and then running Adprep.exe to prepare the forest and domains for the upgrade. Before you upgrade the first Windows 2000–based domain controller to Windows Server 2003 Active Directory, you must use Adprep.exe to:

- Run **adprep /forestprep** once on the schema master to prepare the forest.
- Run **adprep /domainprep** once on the infrastructure master in each domain in which you plan to place a Windows Server 2003–based domain controller.

  For more information about the changes that adprep /forestprep and adprep /domainprep perform, see Running Adprep.exe (http://go.microsoft.com/fwlink/?LinkId=142597).
- If you are using the version of Adprep.exe that is included with Windows Server 2003 Service Pack 1 (SP1), run **adprep /domainprep /gpprep** on the infrastructure master in each domain in which you plan to place a Windows Server 2003–based domain controller. For more information about this new adprep.exe switch, see Enhancements to Adprep.exe in Windows Server 2003 Service Pack 1.

When you are upgrading the operating system on a Windows 2000–based domain controller to Windows Server 2003, Setup (Winnt32.exe) verifies that the forest and domain have been prepared. If you have not prepared the forest and the domain in which the domain controller will be a member, or if the changes have not fully replicated, Winnt32.exe fails, the upgrade terminates, and you are notified that you must run Adprep.exe /forestprep in the forest and Adprep.exe /domainprep in the target domain.

**Note**

- You can run Adprep.exe multiple times, but it performs actions only once. For example, Adprep.exe does not adjust access control lists (ACLs) each time you run the command.

You must prepare your infrastructure before using the Active Directory Installation Wizard to install Active Directory on a Windows Server 2003–based member server. The Active Directory installation fails if the wizard detects that the forest and domain have not been prepared.

**Caution**

- Adprep.exe is the only supported method of upgrading the Windows 2000 Active Directory schema to Windows Server 2003. Attempting to use any other script or tool for this purpose can cause problems with the schema and is not supported by Microsoft.

To prepare your Windows 2000 Active Directory forest and domain for the upgrade to Windows Server 2003 Active Directory, Adprep.exe performs the following tasks:

- Updates the Active Directory schema.

  **Note**

  - Changes that are made to the global catalog by Adprep.exe do not cause a full synchronization of the global catalog because the partial attribute set is not changed.
- Improves default security descriptors.
- Upgrades display specifiers.

- Adjusts access control lists on Active Directory objects and on files in the SYSVOL shared folder to allow domain controller access.

  In versions of Windows earlier than Windows Server 2003, including the Everyone security identifier (SID) on an ACL or group membership allows authenticated users, guest users, and anyone with an anonymous logon to gain access to many resources. Windows 2000–based domain controllers also use anonymous access to gain control of some Active Directory objects and files. In Windows Server 2003, the Everyone group no longer contains the anonymous users group, thus restricting domain controller access to particular objects. Adprep.exe adjusts the ACLs on these objects so that domain controllers can still access them.

  In the updated version of adprep.exe, this change is performed by the new **adprep /domainprep /gpprep** switch. For more information about this new adprep.exe switch, see Enhancements to Adprep.exe in Windows Server 2003 Service Pack 1.

- Creates new objects that are used by applications such as COM+ and Windows Management Instrumentation (WMI).

- Creates new containers in Active Directory that are used to verify that the preparation was successful.

You can run Adprep.exe only at the command line.

**Enhancements to Adprep.exe in Windows Server 2003 Service Pack 1**

In the original release version of Windows Server 2003, one of the operations that is performed by the **adprep /domainprep** command adds more restrictive security descriptors to all Group Policy objects (GPOs) in the SYSVOL shared resource. When you modify the permissions on all the GPOs in the SYSVOL tree, the NT File Replication Service (NTFRS) on the originating domain controller must send all the GPOs to all the other domain controllers in that domain, which causes an increase in volume of replication. If you have a large number of domain controllers or a large number of GPOs in the SYSVOL tree, you must assess the potential impact that this replication traffic will have on your network. If you determine that your network will be affected negatively, you can use the **adprep /domainprep /gpprep** switch separately from the **adprep /domainprep** switch to modify the permissions in the SYSVOL shared resource during nonpeak hours when network bandwidth is optimal.

The updated version of Adprep.exe supports the following commands and enhancements. These enhancements help administrators upgrade successfully to Windows Server 2003.

- **adprep /forestprep**

  The **adprep /forestprep** command performs the same operations as it does in the original version of Windows Server 2003. The syntax of this command is unchanged. Enhancements include better error message handling in configurations that prevent the **adprep /forestprep** command from running successfully.

- **adprep /domainprep**

  The **adprep /domainprep** command performs the same operations as it does in the original version of Windows Server 2003. However, the updated command does not modify permissions on GPOs unless the new **/gpprep** switch is used.

  After you run the updated **adprep /domainprep** command, you receive the following message:

```
The new cross domain planning functionality for Group Policy, RSOP Planning Mode,
requires file system and Active Directory permissions to be updated for existing
Group Policy Objects (GPOs). You can enable this functionality at any time by
running "adprep.exe /domainprep /gpprep" on the DC that holds the infrastructure
operations master role.
This operation will cause all GPOs located in the policies folder of the SYSVOL to
be replicated once between the domain controllers in this domain. Microsoft
recommends reading KB Q324392, particularly if you have a large number of Group
Policy Objects.
```

- **adprep /domainprep /gpprep**

The functionality of the **adprep domainprep /gpprep** command depends on the state of the domain. If the updated **adprep /domainprep** command has not been run, this command is the functional equivalent of the **adprep /domainprep** command in the original version of Windows Server 2003.

If the updated **adprep /domainprep** command has already been run, the **adprep /domainprep /gpprep** command adds only the inheritable access control entries (ACEs) on GPOs in the SYSVOL shared resource.

| **Note** |
| --- |
| The new cross-domain planning functionality for Group Policy, Resultant Set of Policy (RSOP) Planning Mode, requires file system and Active Directory permissions to be updated for existing GPOs. You can enable this functionality anytime by running **adprep.exe /domainprep /gpprep** on the domain controller that holds the infrastructure operations master role. For more information about RSoP functionality, see "RSoP Overview" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=48185. |

**Resolve Adprep.exe Compatibility Problems with Exchange 2000**

When you prepare the forest by using the Active Directory Preparation tool in a Windows 2000 forest containing the Exchange 2000 schema, the LDAP display names of the three Windows Server 2003 InetOrgPerson attributes *Secretary, labeledURI, and houseIdentifier* conflict with the non-RFC-compliant versions added by Exchange 2000. On the domain controller that receives the Windows Server 2003 schema updates, the *lDAPDisplayName* attributes for the Exchange 2000 definitions of these attributes are modified to prevent a conflict. When the changes are replicated in Active Directory, however, the additional domain controllers inadvertently detect the changes as a schema collision because duplicate names are present.

When Active Directory detects a duplicate name, it modifies the name of one of the objects by adding "Dup" and some unique characters to the beginning of the name. For example, the *Secretary, labeledURI,* and *houseIdentifier* name collisions appear similar to the following:

```
lDAPDisplayName: DUP-labeledURI-9591bbd3-d2a6-4669-afda-48af7c35507d
lDAPDisplayName: DUP-secretary-c5a1240d-70c0-455c-9906-a4070602f85f
lDAPDisplayName: DUP-houseIdentifier-e7c5d1bd-a422-4b9e-b4db-ecad2b6839cf
```

If you are already running Exchange 2000, you need to run the fixup script found in article 314649, "ADPREP Command Causes Mangled Attributes in Windows 2000 Forests That Contain Exchange 2000 Servers." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

If you have not yet deployed Exchange 2000 in your environment, you can avoid name collisions by preparing the Active Directory forest by using **adprep /forestprep** to create the initial definition of the *Secretary, labeledURI,* and *houseIdentifier* attributes before installing Exchange 2000. Specifically, you can avoid LDAP display name collision problems by doing one of the following:

- Run the Active Directory Preparation tool in a Windows 2000 forest before you install Exchange 2000.
- Add Exchange 2000 to an existing Windows Server 2003 forest.

For more information about schema collisions between Exchange 2000 and Windows Server 2003, see article 314649, "ADPREP Command Causes Mangled Attributes in Windows 2000 Forests That Contain Exchange 2000 Servers," and article 325379 "How to Upgrade Windows 2000 Domain Controllers to Windows Server 2003" in the Microsoft Knowledge Base. To find these articles, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Resolve Adprep.exe Compatibility Problems with Services for UNIX 2.0**

Adprep.exe prepares the forest or domain with the schema attribute CN=uid, which is compliant with RFC 2307 for use by the Server for Network Information System (NIS) component of Services for UNIX. However, in Services for UNIX 2.0, the Server for NIS component uses a different attribute schema: CN-uid,CN=msSFUName. This discrepancy can cause the upgrade to Windows Server 2003 to fail. To solve this problem, either upgrade to Windows Services for UNIX 3.0 or install the Q293783_sfu_2_x86_en.exe hotfix.

**To resolve Server for NIS compatibility issues with Windows Server 2003**

1. Run Q293783_sfu_2_x86_en.exe on the domain controller that holds schema master role.

2. Review the Hotfix.txt file that is included with the hotfix for installation specifics.

3. Verify end-to-end Active Directory replication of the schema throughout the forest.

For more information about Services for UNIX 2.0 application compatibility issues and the hotfix installation file, see article 293783, "Cannot Upgrade Windows 2000 Server to Windows Server 2003 with Windows Services for UNIX 2.0 Installed" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Prepare the Forest for the Upgrade**
Before preparing the forest for the upgrade, use your preferred monitoring tool to verify that replication is functioning properly. Changes that are made by Adprep.exe must replicate for the upgrade process to succeed; therefore, if domain controllers are not replicating properly, the upgrade preparation process will fail.

After verifying successful replication, use the **adprep /forestprep** command to prepare the forest for the upgrade.

**To prepare the Active Directory forest for the upgrade**

1. In the forest root domain, log on to the domain controller that holds the schema master role with an account that is a member of the Schema Admins security group.

2. Verify that the schema operations master has performed inbound replication of the schema directory partition by typing the following at a command prompt:

   - **repadmin /showreps**

3. A previous release of this documentation suggested that you take the schema master offline before running adprep /forestprep as a precautionary measure to protect the Active Directory schema from corruption. It has since been determined that this is not necessary and might cause a schema operations master to reject schema changes when it is restarted on a private network.

4. Insert the Windows Server 2003 operating system CD, or connect to the network installation shared folder, and then locate and open the I386 folder. On the schema operations master, type the following command at the command line:

   ```
   adprep /forestprep
   ```

   The following warning appears:

   ```
   ADPREP WARNING:
   Before running adprep, all Windows 2000 domain controllers in the forest should be
   upgraded to Windows 2000 Service Pack 1 (SP1) with QFE 265089, or to Windows 2000
   SP2 (or later).
   QFE 265089 (included in Windows 2000 SP2 and later) is required to prevent
   potential domain controller corruption.
   For more information about preparing your forest and domain see KB article Q331161
   at http://support.microsoft.com.
   [User Action]
   If ALL your existing Windows 2000 domain controllers meet this requirement, type C
   and then press ENTER to continue. Otherwise, type any key and press ENTER to quit.
   ```

   You can either continue with the preparation process or quit and install SP2. It is recommended that you install SP2 or later on all Windows 2000–based domain controllers before continuing.

After **adprep /forestprep** has finished, verify that all operations have completed successfully.

**To verify that the Active Directory Preparation tool has completed all operations successfully**

1. At the command line, type:

1. At the command line, type:

   ```
   adsiedit.msc
   ```

   **Important**

   - Adsiedit.exe is one of the Windows 2000 support tools, which is still installed on the computer at this point in the domain upgrade process. If you have removed the Windows 2000 support tools, you can reinstall them from the Support\Tools folder on the Windows 2000 operating system CD. For more information about Adsiedit.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Install Windows Support Tools.

2. Expand the **Configuration** container and verify that **CN=ForestUpdates** has been created.

3. Expand **CN=ForestUpdates** and verify that **CN=Windows2003Update** is present.

4. Examine the Event Log for any event messages that indicate that the domain controller is not functioning properly.

5. Verify that the changes that Adprep.exe made to the schema operations master are being replicated to all other the domain controllers in the forest.

   Successful replication is necessary when preparing an entire forest for Active Directory upgrade because you can prepare a domain controller by using the **adprep /domainprep** command only if it has received the changes made by the **adprep /forestprep** command. Attempting to upgrade a domain controller that has not received the changes generates an error message. Allow enough time for the changes to replicate to all domains in the forest.

   **Tip**

   - Adprep.exe creates a log file each time it runs that can help you troubleshoot errors. The log file documents each step of the forest preparation process. Each Adprep.exe log file is located in a subfolder in the %systemroot%\System32\Debug\Adprep folder. Each subfolder is stamped with the date and time when Adprep.exe was run.

**Prepare the Domain for Upgrade**

After you prepare the forest for the upgrade, you must also prepare each domain in which you plan to operate a Windows Server 2003–based domain controller.

**To prepare the Active Directory domain for upgrade**

1. Log on to the infrastructure master by using Domain Admins credentials.

2. Insert the Windows Server 2003 operating system CD, or connect to the network installation shared folder, and then locate and open the I386 folder. At the command line, type:

   ```
   adprep /domainprep
   ```

After **adprep /domainprep** has finished, verify that all operations have completed successfully.

If you are upgrading to Windows Server 2003 with SP1 or you have the updated version of adprep.exe, use the **adprep /domainprep /gpprep** switch at the command line to prepare the domain for upgrade. This command will cause a full synchronization of the SYSVOL folder. Therefore, if you are concerned about the impact of the large-scale replication, you can run **adprep /domainprep** to prepare the domain and then run **adprep /domainprep /gpprep** to synchronize the GPOs in the SYSVOL folder at a later time, when network bandwidth permits.

**To verify that the Active Directory Preparation tool has completed all operations successfully**

- Using Adsiedit.exe, expand the **Domain** container, and then go to **DC=*domainname*,DC=com**,**CN=System**,**CN=DomainUpdates**. Verify that **CN=Windows2003Update** is present.

–Or–

In **Active Directory Users and Computers**, from the **View** menu, select **Advanced Features**. Expand the **System** container, go to the **DomainUpdates** container, and then expand it. Verify that the **Windows2003Update** container is present.

If you receive an error message, do one of the following, based on the error message text:

- Run the **adprep /forestprep** command.
- Wait for replication to complete.
- Troubleshoot replication.

# Upgrading Windows 2000 Domains to Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can begin the actual domain upgrade process after the forest is prepared, all changes made by the **adprep /forestprep** command have replicated throughout the forest, and all applicable domains have been prepared with **adprep /forestprep**.

Figure 9.7 shows the process for upgrading Windows 2000–based domain controllers to Windows Server 2003 Active Directory.

**Figure 9.7 Upgrading Windows 2000 Domains to Windows Server 2003 Domains**




Windows Server 2003–based domain controllers can be introduced in your environment by either installing Active Directory on a Windows Server 2003–based member server by using the Active Directory Installation Wizard or by upgrading the operating system of an existing Windows 2000–based domain controller. Refer to your domain controller documentation table, and follow the upgrade order determined earlier in the planning process. For more information about the order in which to upgrade your domain controllers, see "Determine Domain Controller Upgrade Order" earlier in this chapter.

**Note**

- Before you attempt to upgrade a domain controller in another domain to Windows Server 2003 Active Directory, remember that you must first run the **adprep /domainprep** command on the infrastructure master role holder in that domain. Run **adprep /forestprep** only once in the forest root domain, and run **adprep /domainprep** once in each domain in the forest in which you plan to locate a Windows Server 2003-based domain controller.

# Install Active Directory on Windows Server 2003-Based Member Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Install Active Directory on a Windows Server 2003–based member server that is located in the forest root domain by using the Active Directory Installation Wizard. When you install Active Directory, the member server becomes a domain controller. You can install Active Directory on any Windows Server 2003–based member server that meets the domain controller hardware requirements. This is the recommended method for introducing Windows Server 2003 into your environment.

The Active Directory Installation Wizard:

- Allows you to create an additional domain controller in the existing domain.
- Configures the local server to host the directory service.
- Creates directory partitions and default domain security principals.
- Allows you to install or configure DNS.

To install Active Directory on a Windows Server 2003–based member server, start the Active Directory Installation Wizard by using one of the following methods:

**To install Active Directory on a Windows Server 2003–based member server**

- At the command line, type:

  ```
  dcpromo
  ```

  – or –

  Open **Administrative Tools**, and then click **Configure Your Server Wizard**. Select **Domain Controller (Active Directory)** to configure your domain controller. After the Configure Your Server Wizard finishes, the Active Directory Installation Wizard begins.

After the first Windows Server 2003–based domain controller has been deployed, you can install Active Directory on additional domain controllers by installing from media, a new installation feature of Windows Server 2003. Installing from media allows you to pre-populate Active Directory with System State data backed up from an existing Windows Server 2003–based domain controller. This backup can be present on local CD, DVD, or hard disk partition. Installing from media drastically reduces the time required to install directory information by reducing the amount of data that is replicated over the network. Installing from media is most beneficial in environments with very large domains or for installing new domain controllers that are connected by a slow network link.

**To install Active Directory on a Windows Server 2003–based member server from media**

- In the **Run** dialog box, **type dcpromo /adv**, and then click **OK**.

The wizard prompts you to choose a network share or a backup as the installation source. If you are installing from backup files, you must identify the location of the files. If the domain controller from which you restored the System State data was a global catalog, you will have the option to make this new domain controller a global catalog. The wizard will then proceed with the installation.

Table 9.4 lists information for installing Active Directory on a Windows Server 2003–based member server, in addition to sample data for installing Active Directory on an additional domain controller in the existing Contoso forest. Contoso will install Active Directory from a Windows Server 2003, Enterprise Edition CD by using the **dcpromo** command.

**Table 9.4 Installing Active Directory on Windows Server 2003–Based Member Servers**

| Wizard Page or Dialog Box | Action | Example |
|---|---|---|
| Domain Controller Type | Select **Additional domain controller for an existing domain**. | |
| Network Credentials | Type the user name and password of an account with sufficient privileges to install Active Directory on this computer, and the fully qualified domain name of the domain in which the computer will become an additional domain controller. | |
| Additional Domain Controller | Type the full DNS name of the forest root domain. | **Concorp.contoso.com** |
| Database and Log Folders | Type the folder locations specified by your design. | Database folder: C:\Windows\NTDS<br><br>Log folder: D:\Logs |
| Shared System Volume | Confirm or type the location specified by your design. | C:\Windows\SYSVOL |
| Directory Service Restore Mode Administration Password | In the **Password and Confirm password** boxes, type any strong password. | |

Verify that all information on the Summary page is accurate, and then click **Finish**. After Active Directory is installed, you will be prompted to restart the computer. The installation will not be complete until the computer restarts.

After you install Active Directory on the Windows 2003–based member server, allow sufficient time for replication to occur and other domain controllers to synchronize with the new domain controller.

For more information about installing and removing Active Directory, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

# Upgrade Existing Windows 2000-Based Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you upgrade the operating system on a Windows 2000–based domain controller to Windows Server 2003, the computer immediately assumes the role of domain controller after the final restart of the computer. It is not necessary to install Active Directory by using the Active Directory Installation Wizard. Upgrade the following existing Windows 2000–based domain controllers early in the upgrade process:

- The Windows 2000–based domain controller that holds the role of the domain naming master. This will ensure the creation of application directory partitions that will later be used for DNS. If you choose not to upgrade the domain naming master, you must transfer the role of the domain naming master to a Windows Server 2003 domain controller.

  When a domain controller running the DNS Server service restarts for the first time after the operating system has been upgraded to Windows Server 2003, it will try to subscribe to existing application directory partitions or create them if it does not detect them. If the domain naming master is not a Windows Server 2003–based domain controller, the creation of the application directory partitions will fail and errors will be generated.

- The Windows 2000–based domain controller that holds the PDC emulator role in the forest root domain. This will ensure that additional security principals are created for the forest. For more information about the additional security principals that are created after the PDC emulator in the forest root domain is upgraded, see "Background Information for Upgrading Windows 2000 Domains to Windows Server 2003 Domains" earlier in this chapter.

- All other Windows 2000–based domain controllers that hold the PDC emulator role. This will ensure that all new Windows Server 2003 groups and group memberships are created. If you choose not to upgrade the PDC emulator for each domain, you must transfer the PDC role to a Windows Server 2003–based domain controller.

To initiate the installation of the operating system on a domain controller, insert the Windows Server 2003 operating system CD on the domain controller, or, if the Windows Server 2003 media are shared over the network, run the Winnt32.exe command-line tool.

You can also perform an unattended installation of Windows Server 2003. Instructions for creating an answer file for an Active Directory installation are located in the Deploy.cab file in the Support\Tools folder on the Windows Server 2003 operating system CD. Inside the Deploy.cab file, open Ref.chm to access the Unattend.txt file. Expand **Unattend.txt** in the left pane, and then click **DCInstall**.

# Modify Default Security Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To ensure that clients running older versions of the Windows operating system will be able to access domain resources in the new Windows Server 2003 domain, you might have to modify default security policies.

**Important**

- Be aware that by modifying these policies you are weakening the default security policies in your environment. However, this is necessary to ensure that some clients running earlier versions of Windows will be able to access domain resources. After all clients in your environment are running versions of Windows that support SMB packet and secure channel signing, you can re-enable these security policies to increase security. It is recommended that you upgrade your Windows–based clients as soon as possible.

In order to increase security, Windows Server 2003–based domain controllers require by default that clients attempting to authenticate to them use SMB packet and secure channel signing. Clients running Windows 95 or Windows NT 4.0 with Service Pack 2 (SP2) and earlier without the Directory Service Client Pack do not support SMB packet signing and will not be able to log on or access domain resources on the network. Clients running Windows NT 4.0 with Service Pack 3 (SP3) and earlier do not support secure channel signing and will not be able to establish communications with a domain controller in their domain.

The most secure way to enable these clients to logon and access domain resources on the network is to apply the appropriate service pack or the Directory Service Client Pack. If you cannot apply the most recent service pack or the Directory Service Client Pack, configure all Windows Server 2003–based domain controllers to not require SMB packet signing or secure channel signing by disabling the following settings in the Default Domain Controllers Policy:

- **Microsoft network server: Digitally sign communications (always)**
- **Domain member: Digitally encrypt or sign secure channel data (always)**

Back up the Default Domain Controllers Policy Group Policy object before modifying it. Use the Group Policy Management Console (GPMC) to back up the Group Policy object so that it can be restored if necessary. The Group Policy Management Console (GPMC) is a tool that permits you to manage Group Policy for multiple domains and sites in one or more forests. GPMC is the recommended method for managing Group Policy; however this chapter does not assume that you are using GPMC for security policy management and deployment.

GPMC is not included with Windows Server 2003. To obtain GPMC, see the Group Policy Management Console (GPMC) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**To disable SMB packet and secure channel signing enforcement on Windows Server 2003–based domain controllers**

1. Open **Active Directory Users and Computers**, right-click the **Domain Controllers** container, and then click **Properties**.
2. Click the **Group Policy** tab, and then click **Edit**.
3. Under **Computer Configuration**, go to the Windows Settings\Security Settings\Local Policies\Security Options folder.
4. In the details pane, double-click **Microsoft network server: Digitally sign communications (always)**, and then click **Disabled** to prevent SMB packet signing from being required.
5. Click **OK**.
6. In the details pane, double-click **Domain member: Digitally encrypt or sign secure channel data (always)**,and then click **Disabled** to prevent secure channel signing from being required.
7. Click **OK**.

To apply the Group Policy change immediately, either restart the domain controller, or type **gpupdate /force** at a command line, and then press ENTER.

**Note**

- Modifying these settings in the Domain Controllers container will change the Default Domain Controllers Policy. Policy changes made here will be replicated to all other domain controllers in the domain, so you only need to modify these policies one time to affect the Default Domain Controllers Policy on all domain controllers.

For more information about SMB packet signing and secure channel signing, see "Background Information for Upgrading Windows 2000 Domains to Windows Server 2003 Domains" earlier in this chapter.

For more information about security policies, see "Security Setting Descriptions" on the Microsoft Web site.

For more information about managing and deploying security policies and the Group Policy Management Console (GPMC), see "Deploying Security Policy" in *Designing a Managed Environment* in this kit.

# Update Group Policy Permissions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy Modeling is a new feature of the GPMC that simulates the resultant set of policy for a particular configuration. The simulation is performed by a service that runs on Windows Server 2003–based domain controllers. To perform the simulation across domains, the service must have read access to all Group Policy objects (GPOs) in the forest.

In a Windows Server 2003 domain that has been upgraded from Windows 2000 or newly installed, the Enterprise Domain Controllers group is automatically given read access to all newly created GPOs. This ensures that the service can read all GPOs in the forest.

However, if the domain was upgraded from Windows 2000, the Enterprise Domain Controllers group will not have read access to any existing GPOs that were created prior to the upgrade. The Group Policy Management Console detects this when you click a GPO and notifies the user that Enterprise Domain Controllers do not have read access to all GPOs in this domain. To solve this problem, use the sample script that is provided with the Group Policy Management Console, GrantPermissionOnAllGPOs.wsf. This script will update the permissions on all GPOs in the domain. You must be a member of the Domain Admins group or have permissions to modify security on all GPOs in the domain to run this script.

**Note**

- To download the GPMC, see the Group Policy Management Console link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**To update permissions on all GPOs in a domain**

1. At the command line, change to the %programfiles%\Gpmc\Scripts folder.

2. Type the following:

   ```
   GrantPermissionOnAllGPOs.wsf "Enterprise Domain Controllers" /permission:read
   /domain:DNSDomainName /Replace
   ```

   Using the Replace switch removes existing permissions for the group or user before making the change. If a group or user is already granted a permission type higher than the new permission type, and you do not specify Replace, no change is made.

For more information about using GPMC for deploying Group Policy, see "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* in this kit.

For more information about Group Policy Management Console (GPMC) scripting, see the Platform SDK: Group Policy Management Console (http://go.microsoft.com/fwlink?linkid=17912).

# Perform Clean-up Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After upgrading to Windows Server 2003, perform the following clean-up operations:

- After the security descriptor propagator has finished building the single instance store, perform an offline defragmentation of the database on each upgraded domain controller. This reduces the size of Active Directory on the file system by up to 40 percent, reduces the memory footprint, and updates pages in the database to Windows Server 2003 format. For more information about performing an offline defragmentation of the Active Directory database, see the Active Directory link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search under "Administration and Configuration Guides," and then download the *Active Directory Operations Guide*. The procedure for performing an offline defragmentation of the database applies to both Windows 2000 and Windows Server 2003.

- Create a new System State backup for at least two domain controllers in your environment. For more information about backing up Active Directory, see the *Active Directory Operations Guide*. The procedure for backing up System State data applies to both Windows 2000 and Windows Server 2003. Be sure to label all backup tapes with the operating system version that the domain controller is running, including service packs and hotfixes.

# Completing Post-Upgrade Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After upgrading the operating system on all domain controllers in the forest to Windows Server 2003, complete the domain upgrade by raising the domain and forest functional levels to Windows Server 2003, and use newly created application directory partitions to store DNS information. You must then redirect the Users and Computers containers.

Figure 9.8 shows the tasks necessary to complete your upgrade to Windows Server 2003 Active Directory.

**Figure 9.8 Completing the Upgrade to Windows Server 2003 Active Directory**



# Raise the Functional Levels of Domains and Forests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To enable all Windows Server 2003 Active Directory features, raise the functional level of your forest to Windows Server 2003. This will automatically raise the functional level of all domains to Windows Server 2003.

**Important**

- Raising the domain and forest functional levels to Windows Server 2003 is a nonreversible task and prohibits the addition of Windows NT 4.0–based or Windows 2000–based domain controllers to the environment. Any existing Windows NT 4.0 or Windows 2000–based domain controllers in the environment will no longer function. Before raising functional levels to take advantage of advanced Windows Server 2003 features, ensure that you will never need to install domain controllers running Windows NT 4.0 or Windows 2000 in your environment.

After all domains are operating at the Windows Server 2003 functional level, raise the forest functional level to Windows Server 2003. To do this, right-click **Active Directory Domains and Trusts** in Active Directory Domains and Trusts, and select **Raise Forest Functional Level**. This allows you to take advantage of all Windows Server 2003 forest-level features.

**Note**

- You can only raise the functional level of the forest to Windows Server 2003 if all domains are set to the Windows 2000 native functional level or higher.

For more information about enabling functional levels after upgrading from a Windows 2000 environment and the features available at the Windows Server 2003 domain and forest functional levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.

# Use DNS Application Directory Partitions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use application directory partitions for Active Directory–integrated DNS zones to reduce replication traffic and the amount of data stored in the global catalog.

After completing the upgrade of all Windows 2000–based domain controllers in the forest to Windows Server 2003, move the Active Directory–integrated DNS data on all DNS servers from the domain partition into the newly created DNS application directory partitions. This is done by changing the replication scope of the DNS zones.

Move the DNS zones that you want to replicate to all DNS servers in the forest to the forest-wide DNS application directory partition, ForestDnsZones. For each domain in the forest, move the DNS zones that you want to replicate to all DNS servers in the domain to the domain-wide DNS application directory partition, DomainDnsZones.

**Important**

- Before you attempt to move DNS data to an application directory partition, make sure that the domain naming master is hosted on a Windows Server 2003–based domain controller.

If there is an existing _msdcs.*forest_root_domain* zone on your DNS server, move it to the ForestDnsZones application directory partition.

If the _msdcs.*forest_root_domain* zone is not present as a separate zone on your DNS server, you do not need to perform this procedure because the DNS data that is stored in the _msdcs.*forest_root_domain* is moved with the forest root domain zone to the domain-wide application directory partition, DomainDnsZones.

**Note**

- For more information about DNS and application directory partitions, see "Background Information for Upgrading Windows 2000 Domains to Windows Server 2003 Domains" earlier in this chapter.

**To change the replication scope of the domain-wide DNS zone by using a DNS application directory partition**

1. On a domain controller that hosts a DNS server in the particular domain, open the DNS snap-in, right-click the DNS zone that uses the fully qualified domain name of the Active Directory domain, and then click **Properties**.
2. Click the **Change** button next to **Replication: All domain controllers in the Active Directory domain**.
3. Click **To all DNS servers in the Active Directory domain *domainname***.

**To change the replication scope of the _msdcs.forest_root_domain DNS zone by using a DNS application directory partition**

1. On a domain controller that hosts a DNS server in the forest root domain, open the DNS snap-in, right-click the _msdcs.*forest_root_domain* DNS zone, and then click **Properties**.
2. Click the **Change** button next to **Replication: All domain controllers in the Active Directory domain**.
3. Click **To all DNS servers in the Active Directory forest *forestname***.

For more information about creating, enlisting in, and removing application directory partitions, see Help and Support Center for Windows Server 2003. For more information about creating a DNS design for Active Directory, see "Designing the Active Directory Logical Structure" in this book.

# Redirect Users and Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The default CN=Users and CN=Computers containers that are created when Active Directory is installed are not organizational units. Objects in the default containers are more difficult to manage because Group Policy cannot be applied directly to them. New user accounts, computer accounts, and security groups that are created by using earlier versions of user interface and command-line management tools, do not allow administrators to specify a target organizational unit and, therefore, to create these objects in either the CN=Computers container or the CN=User container by default. Examples of these earlier versions include the **net user** and **net computer** commands, the **net group** command, the **netdom add** command where the **/ou** parameter is either not specified or supported, or Windows NT 4.0 tools such as User Manager for Domains.

It is recommended that administrators who upgrade Windows NT 4.0–based and Windows 2000–based domain controllers to Windows Server 2003 redirect the well-known path for the CN=Users and CN=Computers containers to an organizational unit specified by the administrator so that Group Policy can apply to containers hosting newly created objects.

**Important**

- The CN=Users and CN=Computers containers are computer-protected objects. For backward compatibility reasons, you cannot (and must not) remove them. However, you can rename these objects.

In Windows Server 2003 Active Directory, when the domain functional level has been raised to Windows Server 2003, you can redirect the default CN=Users and CN=Computers containers to organizational units that you specify so that each can support Group Policy, making them easier to manage.

**To redirect the Users container**

1. In Active Directory Users and Computers, create an organizational unit container to which you will redirect user created with earlier versions of user interface and command-line management tools.
2. At the command line, change to the System32 folder by typing:

```
cd %systemroot%\system32
```

3. At the %systemroot%\System32 folder, type the following, where *newuserou* is the name of the new user OU and *domainname* is the name of the domain:

```
redirusr ou=
newuserou
,DC=
domainname
,dc=com
```

**To redirect the Computers container**

1. In Active Directory Users and Computers, create an organizational unit container to which you will redirect computer objects created with earlier versions of user interface and command-line management tools.
2. At the command line, change to the System32 folder by typing:

```
cd %systemroot%\system32
```

3. At the %systemroot%\System32 folder, type the following, where *newcomputerou* is the name of the new computer OU and *domainname* is the name of the domain:

```
redircmp ou=
newcomputerou
,DC=
domainname
,dc=com
```

For more information about creating an organizational unit design, see "Designing the Active Directory Logical Structure" in this book.

# Completing the Upgrade

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Complete the following tasks to finalize the process:

- Review, update, and document the domain architecture to reflect any changes that you made during the domain upgrade process.
- Verify that the NETLOGON and SYSVOL shared folders exist and that the File Replication Service (FRS) is functioning without error by checking the Event Viewer.
- Verify that Group Policy is being applied successfully by checking the application log in Event Viewer for Event ID 1704.
- Verify that all SRV, CNAME, and A resource records have been registered in DNS.
- Continuously monitor your domain controllers and Active Directory. Using a monitoring solution such as Microsoft® Operations Manager to monitor the distributed Active Directory service and the services that it relies on helps maintain consistent directory data and a consistent level of service throughout the forest.

After these tasks have been completed successfully, you will have completed the in-place upgrade process.

# Additional Resources for Upgrading from Windows 2000 Domains to Windows Server 2003 Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Enabling Advanced Windows Server 2003 Active Directory Features" in this book for more information about advanced Active Directory features and how they are related to functional levels.
- "Designing the Active Directory Logical Structure" in this book.
- "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.
- "Process for Deploying DNS" in *Deploying Network Services* of this kit for more information about deploying DNS to support name resolution on your network.
- The Active Directory Branch Office Planning Guide link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Tools**

- Adsiedit.exe

  The ADSIEdit.exe tool is an MMC snap-in that you can use to edit objects in the Active Directory database. For more information about ADSIEdit.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Install Windows Support Tools.

- Repadmin.exe

  The Repadmin.exe tool can be used to administer replication between domain controllers in Active Directory. For information about how to use the Repadmin.exe tool, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Install Windows Support Tools.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Active Directory" in Help and Support Center for Windows Server 2003.
- "Installing and Upgrading the Operating System" in Help and Support Center for Windows Server 2003 for more information about Active Directory preparation and the Active Directory Preparation tool (Adprep.exe).
- "Managing Core Network Services" in Help and Support Center for Windows Server 2003 for more information about application directory partitions.

**Related Job Aids**

- "Pre-Upgrade Task Checklist" (DSSUPWN_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Pre-Upgrade Task Checklist" on the Web at http://www.microsoft.com/reskit).
- "Windows 2000 Domain Controller Documentation" (DSSUPWN_2.doc) on the *Windows Server 2003 Deployment Kit*

companion CD (or see "Windows 2000 Domain Controller Documentation" on the Web at http://www.microsoft.com/reskit).

- "Windows 2000 Upgrade Test Matrix" (DSSUPWN_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows 2000 Upgrade Test Matrix" on the Web at http://www.microsoft.com/reskit).

# Restructuring Windows NT 4.0 Domains to an Active Directory Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Microsoft® Windows NT® 4.0 operating system domains to an Active Directory® directory service forest involves migrating the contents of a number of Windows NT 4.0 domains into a smaller number of Active Directory domains. By restructuring your environment to an Active Directory forest, you can take advantage of new and enhanced Active Directory features, while reducing the complexity of your organization and the associated administrative costs.

**In This Chapter**

Overview of Restructuring Windows NT 4.0 Domains to an Active Directory Forest

Planning to Restructure Windows NT 4.0 Domains to an Active Directory Forest

Preparing the Source and Target Domains for Restructuring

Restructuring Account Domains

Restructuring Resource Domains

Additional Resources for Restructuring Windows NT 4.0 Domains

**Related Information**

- For more information about designing the Active Directory logical structure, see "Designing the Active Directory Logical Structure" in this book.
- For more information about designing the Active Directory site topology, see "Designing the Site Topology" in this book.
- For more information about creating an Active Directory forest, see "Deploying the Windows Server 2003 Forest Root Domain" in this book.

# Overview of Restructuring Windows NT 4.0 Domains to an Active Directory Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are upgrading from Windows NT 4.0 domains to an Active Directory forest, you can consolidate your existing account and resource domains by restructuring them into a smaller number of larger Active Directory domains. You can restructure Windows NT 4.0 domains into Active Directory domains running the Microsoft® Windows® 2000 or Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; or Windows® Server 2003, Datacenter Edition operating systems.

Before you begin to restructure your Windows NT 4.0 domains, ensure that your organization has designed an Active Directory logical structure and a site topology. In a Windows 2000 environment, ensure that any domains that you plan to use as target domains are operating in Windows 2000 native mode. In a Windows Server 2003 environment, ensure that for any domains that you plan to use as target domains, you have raised the Active Directory domain functional level to Windows 2000 native or Windows Server 2003.

For more information about upgrading Windows NT 4.0 domains in-place, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.

After the restructuring process is complete, you can take advantage of the features available in Active Directory. By consolidating multiple Windows NT 4.0 account and resource domains into a single or a few Active Directory domains, you can create an Active Directory domain structure that is more efficient than your existing Windows NT 4.0 domain structure.

**Note**

- For a list of the job aids that are available to assist you in restructuring Windows NT 4.0 domains, see "Additional Resources for Restructuring Windows NT 4.0 Domains" later in this chapter.

# Process for Restructuring Windows NT 4.0 Domains to an Active Directory Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Windows NT 4.0 domains to an Active Directory forest involves completing the necessary planning and preparation tasks in the appropriate order, and then migrating Windows NT 4.0 accounts and resources to an Active Directory domain or domains.

Figure 10.1 shows the process for restructuring Windows NT 4.0 domains to an Active Directory forest.

**Figure 10.1 Restructuring Windows NT 4.0 Domains to an Active Directory Forest**



# Windows NT 4.0 Domain Migration Background Information

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Windows NT 4.0 domains involves migrating user, group, and computer objects from a Windows NT 4.0 account or resource source domain into an Active Directory target domain. The target domain can either be a newly created Active Directory domain, or a Windows NT 4.0 domain that has been upgraded in-place.

After the migration process is complete, clients can log on to the original Windows NT 4.0 domain or the new Active Directory domain for authentication. Clients can also access resources in either the source or the target domain, until the source domain is decommissioned.

To successfully consolidate Windows NT 4.0 account and resource domains into Active Directory domains, it is important to be familiar with the capabilities of both operating systems in the context of restructuring. You must also be familiar with the tools that support object migrations.

The migration process between forests is non-destructive. Because a Windows NT 4.0 domain cannot be part of an Active Directory forest, when objects are migrated from a Windows NT 4.0 domain to an Active Directory domain, they are copied from the source domain to the target domain. The source account primary security identifier (SID) can be maintained in the SID history of the object. You can use the Active Directory Migration Tool (ADMT) version 2 or a third-party restructuring tool to restructure domains.

For more information about domain migration, SIDs, and SID history, see the Domain Migration Cookbook link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Terms and Definitions for Domain Restructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following terms apply to the Windows NT 4.0 domain restructure process.

**Domain restructure** A migration process that involves changing the domain structure of a forest. A domain restructure can involve either consolidating or adding domains, and can take place between forests or within a forest.

**Domain consolidation** A restructuring process that involves eliminating Windows NT 4.0 or Active Directory domains by merging their content with that of other domains.

**Migration objects** Domain objects that you plan to migrate to a new domain by using a tool, such as ADMT. These objects can be user accounts, service accounts, groups, or computers.

**Source domain** The domain that contains the objects that you plan to migrate to a new domain.

**Target domain** The domain that receives the migration objects from the source domain. In a Windows 2000 environment, the target domain must be operating in Windows 2000 native mode. In a Windows Server 2003 environment, the target domain must be at the Windows 2000 native or Windows Server 2003 Active Directory domain functional level.

**Built-in groups** Default security groups with common sets of rights and permissions. These rights and permissions are assigned to any accounts or groups that you designate as members of these groups. Built-in account SIDs are identical in every domain, and therefore, built-in groups cannot be migration objects.

**Password export server (PES)** A domain controller in the source domain that is used to securely copy account passwords from the source domain to the target domain. When you are migrating accounts and passwords between domains, one source domain controller must be configured as the PES.

# Active Directory Migration Tool 2.0

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ADMT enables you to copy objects from Windows NT 4.0 domains to an Active Directory forest. ADMT includes wizards that automate migration tasks such as copying users, groups, and service accounts; moving computers; migrating trusts; and performing security translation.

When you use ADMT to restructure Windows NT 4.0 domains, ADMT copies the accounts that are migrated, so that when the accounts are created in the target domain, they continue to exist in the source domain. The primary security identifiers (SIDs) for the accounts can be migrated to the SID history in the target domain.

ADMT is available on the Windows Server 2003 operating system CD, in the Admigration.msi file in the \i386\admt directory.

You can perform ADMT tasks by using the ADMT console, by using a command-line procedure, or by using a script. When you are running ADMT from the command line, it is often more efficient to use an option file to specify command-line options. You can use the ADMT option file reference in Listing 10.1 to assist you in creating option files.

Listing 10.1 lists all common options that apply to several migration tasks. Each type of migration task has a section that lists options that are specific to that task. The section name corresponds to the task name used when executing the ADMT command-line tool. Items can be commented out with a semicolon. In Listing 10.1, the default values are commented out.

When running ADMT at the command line, you do not need to include an option in your command if you want to accept the default value. In this chapter, however, tables are provided for reference. The tables list the command-line equivalent of each option that is shown in the corresponding ADMT console procedure, including those where you are accepting the default value.

**Listing 10.1 ADMT Option File Reference**

```
 [Migration]
;TestMigration=No
;IntraForest=No
SourceDomain=" source_domain_name"
;SourceOu=" source_ou_path"
TargetDomain=" target_domain_name"
TargetOu=" target_ou_path"
;RenameOption=dont
;RenamePrefixOrSuffix=""
;PasswordOption=Complex
;PasswordServer=""
;PasswordFile=""
;ConflictOptions=Ignore
;ConflictPrefixOrSuffix=""
;UserPropertiesToExclude=""
;InetOrgPersonPropertiesToExclude=""
;GroupPropertiesToExclude=""
;ComputerPropertiesToExclude=""

[User]
;DisableOption=EnableTarget
;SourceExpiration=None
MigrateSIDs=Yes
;TranslateRoamingProfile=No
;UpdateUserRights=No
;MigrateGroups=No
;UpdatePreviouslyMigratedObjects=No
;FixGroupMembership=Yes
;MigrateServiceAccounts=No
;UpdateGroupRights=No

[Group]
MigrateSIDs=Yes
;UpdatePreviouslyMigratedObjects=No
;FixGroupMembership=Yes
;UpdateGroupRights=No
;MigrateMembers=No
;DisableOption=EnableTarget
;SourceExpiration=None
;TranslateRoamingProfile=No
```

```
;MigrateServiceAccounts=No

[Security]
TranslationOption=Add
;TranslateFilesAndFolders=No
;TranslateLocalGroups=No
;TranslatePrinters=No

;TranslateRegistry=No
;TranslateShares=No
;TranslateUserProfiles=No
;TranslateUserRights=No

;SidMappingFile=" SidMappingFile.txt"
```

For a file that contains the ADMT options, see "ADMT Option File Reference" (DSSRENT_3.txt) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "ADMT Option File Reference" on the Web at http://www.microsoft.com/reskit).

When running ADMT from the command line, it is often more efficient to use an include file to list the users, groups, or computers that will be migrated. For example, to migrate a small number of computers, you might type each computer name at the command line, using the /N option as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
"
```

*Computer_name1* and *computer_name2* are the names of computers in the source domain that you are migrating in this batch.

If you have a large number of users, groups, or computers to migrate, you can list them in an include file. For example, to create an include file for a batch of computers, create a plain text file and list the computer names, each name on a separate line. Then specify the include file name with the /F option, as follows:

```
ADMT COMPUTER /F "
includefile_name
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
"
```

ADMT supports use of ANSI or standard Unicode files for option files and include files.

For a complete list of all available ADMT options and syntax, run ADMT and select **Help** from the menu.

The sample scripts provided in this chapter reference the symbolic constants that are defined in the AdmtConstants.vbs file. Listing 10.2 shows the ADMT constants VBScript file.

**Listing 10.2 ADMT Scripting Constants**

```
Option Explicit
'-----------------------------------------------------------------------------
' ADMT Scripting Constants
'-----------------------------------------------------------------------------
' RenameOption constants
Const admtDoNotRename      = 0
Const admtRenameWithPrefix = 1
```

```
Const admtRenameWithSuffix = 2

' PasswordOption constants
Const admtPasswordFromName = 0
Const admtComplexPassword  = 1
Const admtCopyPassword     = 2

' ConflictOptions constants
Const admtIgnoreConflicting           = &H0000
Const admtReplaceConflicting          = &H0001
Const admtRenameConflictingWithPrefix = &H0002
Const admtRenameConflictingWithSuffix = &H0003
Const admtRemoveExistingUserRights    = &H0010
Const admtRemoveExistingMembers       = &H0020
Const admtMoveReplacedAccounts        = &H0040

' DisableOption constants
Const admtEnableTarget       = 0
Const admtDisableSource      = 1
Const admtDisableTarget      = 2
Const admtTargetSameAsSource = 4

' SourceExpiration constant
Const admtNoExpiration = -1

' Translation Option
Const admtTranslateReplace = 0
Const admtTranslateAdd     = 1
Const admtTranslateRemove  = 2

' Report Type
Const admtReportMigratedAccounts  = 0
Const admtReportMigratedComputers = 1
Const admtReportExpiredComputers  = 2
Const admtReportAccountReferences = 3
Const admtReportNameConflicts     = 4

' Option constants
Const admtNone     = 0
Const admtData     = 1
Const admtFile     = 2
Const admtDomain   = 3
Const admtRecurse           = &H0100
Const admtFlattenHierarchy  = &H0000
Const admtMaintainHierarchy = &H0200
```

For a file that contains the ADMT scripting constants, see "ADMT Scripting Constants" (AdmtConstants.vbs) on the *Windows Server 2003 Deployment Kit* companion CD (or see "ADMT Scripting Constants" on the Web at http://www.microsoft.com/reskit).

# Planning to Restructure Windows NT 4.0 Domains to an Active Directory Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Planning your domain restructure involves defining the accounts and resources to be migrated, creating a test plan, and creating a rollback plan for use in the event that the migration fails. You must also plan for user profile migration, establish administrative procedures to ensure that users can continue to log on to the network and access resources during the migration, and create an end-user communication plan.

Figure 10.2 shows the planning steps associated with restructuring Windows NT 4.0 domains to an Active Directory forest.

**Figure 10.2 Planning to Restructure Windows NT 4.0 Domains to an Active Directory Forest**



# Assign Object Locations and Roles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Large organizations are likely to be migrating thousands of user, group, and resource objects from numerous Windows NT 4.0 domains to Active Directory domains. It is important to keep track of these objects by assigning them roles and locations in the target domain. It is useful to document the assignments for all of the objects that you are migrating in object assignment worksheets. Create one worksheet for account objects, such as users, groups, and service accounts, and one worksheet for resource objects, such as workstations, profiles, and domain controllers. In your worksheets, list the source and target locations for all objects to be migrated. For information about identifying service accounts to be migrated, see "Service Account Identification" later in this chapter.

**Note**

- Built-in accounts (such as Administrators, Users, and Power Users) and well-known accounts (such as Domain Admins and Domain Users) cannot be ADMT migration objects.

For a worksheet to assist you in documenting account object assignments, see "User and Group Object Assignment Table — Windows NT 4.0 Source" (DSSRENT_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "User and Group Object Assignment Table — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

Figure 10.3 shows an example of a completed user and group object assignment worksheet.

**Figure 10.3 Example of a Completed User and Group Object Assignment Table Worksheet — Windows NT 4.0 Source**

| User and Group Object Assignment Table - Windows NT 4.0 Source | | | |
|---|---|---|---|
| Prepared By | Active Directory Deployment Team | Date | 01/01/03 |
| Source Domain Name | Boston | | |
| Target Domain Name | east.trccorp.treyresearch.net | | |
| Name | Type | Source Location | Target Location |
| Finance | Group | Boston domain | East\Boston Groups OU |
| Accounting | Group | Boston domain | East\Boston Groups OU |
| JBrown | User | Boston domain | East\Boston Users OU |
| MNguyen | User | Boston domain | East\Boston Users OU |
| !Scheduler | Service Account | Boston domain | East\Service Accounts OU |

Creating a resource object assignment worksheet also involves identifying the target organizational unit (OU) for each object and noting the physical location and role in the target domain. For a worksheet to assist you in documenting resource object assignments, see "Resource Object Assignment Table — Windows NT 4.0 Source" (DSSRENT_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Resource Object Assignment Table — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

Figure 10.4 shows an example of a completed resource object assignment worksheet.

**Figure 10.4 Example of a Completed Resource Object Assignment Table Worksheet — Windows NT 4.0 Source**

| Resource Assignment Table - Windows NT 4.0 Source | | | |
|---|---|---|---|
| Prepared By | Active Directory Deployment Team | Date | 01/01/03 |

| | |
|---|---|
| Source Domain Name | PROD-APPS |
| Target Domain Name | east.trccorp.treyresearch.net |

| Resource Name | Resource Type | Target OU | Physical Location | Role in Target Domain |
|---|---|---|---|---|
| FS01 | Member server (file server) | Member Servers | Boston | Member Server |
| UserWrk01 | Computer account | Workstations | Boston | Computer account |
| UserWrk02 | Computer account | Workstations | Boston | Computer account |
| DC01 | Domain Controller | Domain Controllers | New York | Domain controller |
| DC02 | Domain Controller | Member Servers | Washington DC | Member server |

# Developing a Migration Test Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to develop a test plan to assist you in systematically testing the results of the migration, and identifying and correcting any problems that might occur. Testing allows you to ensure that users who are migrated from the source to the target domain are able to log on, to access resources based on group membership or user credentials, and that they have retained their desktop settings. Testing also allows you to ensure that users are able to access the resources that you migrate.

Use the following process to test your account object and resource object migration:

1. Create a test user in the source domain that you include with your migrations.

2. Join that user to the appropriate global groups to enable resource access.

3. Log on to the source domain as the test user and verify that you can access resources as appropriate.

4. After you migrate the user account, translate the user profile, and migrate the workstation of the user, log on to the target domain as the test user and verify that the user has retained all necessary access and functionality. For example, you might test to verify that:

   - The user can log on successfully.

   - The user has access to all appropriate resources, such as file and print shares; access to services such as messaging; and access to line-of-business applications. It is especially important to test access to applications developed internally that access database servers.

   - The user profile was successfully translated, and the user retains desktop settings, desktop appearance, shortcuts, and access to the My Documents folder. Also, verify that applications appear in and start from the Start Menu.

5. After you migrate resources, log on as the test user in the target domain and verify that you can access resources as appropriate.

If any steps in the test process fail, identify the source of the problem and determine whether you can correct the problem before the object needs to be accessible in the target domain. If you cannot correct the problem before access to the object is required, roll back to your original configuration to ensure access to the user or resource object. For more information about creating a rollback plan, see "Creating a MigrationvRollback Plan" later in this chapter.

As part of your test plan, create a migration test matrix. Complete a test matrix for each step that you complete in the migration process. For example, if you migrate 10 batches of users, complete the test matrix 10 times, once for each batch that you migrate. If you migrate 10 member servers, complete the test matrix for each of the 10 servers.

For a worksheet to assist you in creating a test matrix, see "Migration Test Matrix" (DSSREER_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migration Test Matrix" on the Web at http://www.microsoft.com/reskit).

Figure 10.5 shows an example of a completed migration test matrix.

**Figure 10.5 Example of a Completed Migration Test Matrix Worksheet**

# Migration Test Matrix

| Prepared By | Active Directory Deployment Team | Date | 01/01/03 |
|---|---|---|---|

| | |
|---|---|
| Source Domain Name | Boston |
| Target Domain Name | east.trccorp.treyresearch.net |

| Test Name | Performed Before Migration | Result After Migration | Notes |
|---|---|---|---|
| Create test user for users in Sales department "batch1" | Yes | Success | |
| Join user to sales group | Yes | Success | |
| Log on to source domain and verify access to required resources | Yes | Success | |
| Log on to target domain | No | Success | |
| Verify access to file and print shares | No | Success | |
| Verify access to customer management application | No | Success | |
| Verify user desktop settings and access to My Documents folder | No | Success | |
| Verify that applications appear and start from the Start menu | No | Success | |

# Creating a MigrationvRollback Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Reduce the risk of disrupting end users in your organization by establishing a rollback plan. In general, it is possible to isolate and resolve any problems that occur during each phase of the migration. However, it is important to analyze potential risks and identify the levels of user impact and downtime that might necessitate rolling back the migration. You might be required to roll back your migration if any of the following occur:

- Users cannot log on to their accounts after migration.
- Users cannot access resources after migration.
- User migration is incomplete; for example, passwords did not migrate.
- User migration was successful, but user workstation migration or local profile translation failed.

If user impact or downtime reaches a level that you have defined as unacceptable in your organization, you can implement your rollback plan and continue to operate in your premigration environment. Because the source domain remains intact when you are restructuring Windows NT 4.0 domains, you can restore the original environment by completing a few steps.

To roll back the migration of Windows NT 4.0 account objects:

1. Enable the user accounts in the source domain (if you disabled the accounts during the migration process).
2. Notify the users to log off of the Active Directory target domain and log on to the source domain.
3. Verify that users are able to access resources.
4. Verify that the logon scripts and user profiles for users work as configured in the source domain.

The rollback process for Windows NT 4.0 resource objects is similar to that for account objects. To roll back the migration of Windows NT 4.0 resource objects:

1. Change the domain membership for the resource object to the source domain.
2. Restart the server or workstation.
3. Log on as a user and verify that you can access the resource.

**Note**

- Before you modify objects such as member servers or domain controllers during the migration, back up all data.

# Planning for User Profile Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

User profiles store user data and information about the user's desktop settings. User profiles can either be roaming or local. The migration process is slightly different for local and for roaming profiles.

Profiles are translated during the migration process. Only the primary SID is used to assign a profile to a user; SIDs in SID history are ignored. If you translate profiles in add mode, the user accounts in the target and the source domains have access to the profile. In this way, if you roll the user account back to the source domain, the user account in the source domain can use the profile. If you translate profiles in replace mode, you must retranslate the profile by using a SID mapping file to restore the environment to the premigration state.

**Important**

- In the event that you need to roll back to your original configuration, notify users that profile changes made in the target domain are not reflected in the source domain.

Some organizations might choose not to migrate user profiles. Other organizations might choose to replace users' workstations during the user account migration process, and use a tool such as the User State Migration Tool (USMT) to migrate user data and settings to the users' new computers. Table 10.1 summarizes the migration requirements for user profiles.

Table 10.1 Migration Requirements for User Profiles

| Type | Description | Migration Requirements |
|---|---|---|
| Roaming profiles | User profiles are stored centrally on servers. Profiles are available to the user, regardless of the workstation in use. | Select the profile migration option during the user account migration process. Then migrate local user profiles for a batch of users immediately after migrating those users. |
| Local profiles | User profiles are stored locally on the workstation. When a user logs on to another workstation, a unique local user profile is created. | Migrate as a separate step from the user account migration process. Migrate local user profiles for a batch of users immediately after migrating those users. |
| Profiles not managed | Same as local profiles. | No migration requirements. Users lose their existing profiles when their user accounts are migrated. |
| Hardware refresh | User state information is stored locally on the workstation. | Migrate as a separate step from the user account migration. Migrate the profiles to the user's new computer by means of a tool such as USMT. |

# Establishing Administrative Procedures

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must define how the objects that you are migrating are to be administered during the restructure process. Establishing administrative procedures for migration objects enables you to preserve the objects both in the source and the target domains, and therefore enables you to fall back to the premigration environment in the event that the restructure process is not successful.

Plan for the administration and management of the following types of account migration objects:

- User accounts, including SIDs
- Global group memberships

**Administering User Accounts**

During the migration process, user accounts exist in both the source and the target domains. Continue to administer the accounts in the source domain while the migration is taking place. Use the **replace conflicting accounts** option in ADMT to remigrate user accounts as often as necessary during the migration process. This ensures that changes made to the account in the source domain are propagated to the account in the target domain. This operation merges the existing account and the new account so that administration of the object can continue in the source domain for the duration of the migration process.

**Administering Global Groups**

Continue to administer the groups in the source domain during the migration process. Remigrate groups as often as necessary by using the **replace conflicting accounts** option in ADMT to ensure that changes made to group membership in the source domain are propagated to the group in the target domain.

# Creating an End-User Communication Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Develop a plan to inform all affected users about the upcoming account migration, to ensure that they understand their responsibilities, the impact of the migration, and who to contact for help and support.

Before you begin the user migration process, send a notice to all users who are scheduled to be migrated. Because you typically migrate users in batches of approximately one hundred users at a time, it is also helpful to send a final notice to the users in each batch two to three days before their batch is scheduled. If your organization maintains an intranet, publish the account migration schedule and the information contained in the user mail on an easily accessible Web page.

Include the following information in your end-user communication.

**General information**
Alert users to the fact that their user account is scheduled to be migrated to a new domain, as part of the upgrade to Windows 2000 or Windows Server 2003. Point users to a Web page or internal resource where they can find additional information and view a migration schedule.

Inform users of their new domain name. Be sure to let them know that their account passwords will not change. Let them know that the original domain account will be disabled immediately following the migration, and the disabled account will be deleted after a specified period of time.

**Impact**
Make sure that users understand that when their account is migrated, it is possible that they could be unable to access some resources, such as Web sites, shared folders, or resources that are not widely used by individuals in their group or division.

Provide information to users about whom to contact for assistance in regaining access to required resources.

**Logon status during migration**
Make sure that users understand that during the migration process, they will be unable to log on to the domain or access e-mail or other resources. Be sure to specify the period of time for which they will be unable to log on.

**Premigration steps**
Alert users to any steps that they need to complete before the migration process begins. For example, they must decrypt files encrypted by means of Encrypting File System (EFS). Failure to decrypt encrypted files will result in loss of access to encrypted files following the migration.

Users must also ensure that their computers are connected to the network when their account is scheduled to be migrated.

**Expected changes**
Describe other changes that users can expect to experience following the migration, such as changes in use of smart cards, secure mail, or instant messaging.

**Scheduling and support information**
Provide information about where users can go to find more information; for example, an internal Web site where you post information about the migration. Also, provide information about who to contact if a user has a conflict with the date scheduled for the migration.

# Preparing the Source and Target Domains for Restructuring

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you restructure your Windows NT 4.0 domains, ensure that your target domains are set to a mode or functional level that supports restructuring. In a Windows 2000 environment, the target domains must be operating in Windows 2000 native mode. In a Windows Server 2003 environment, the target domains must be at the Windows 2000 native or Windows Server 2003 Active Directory domain functional level.

Figure 10.6 shows the steps involved in preparing the source and target domains for restructuring.

**Figure 10.6 Preparing the Source and Target Domains for Restructuring**



# Installation of High Encryption Software

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are migrating passwords, both the domain controller on which ADMT is installed in the target domain and the PES in the source domain require 128-bit high encryption.

**Note**

- Installing high encryption in a Windows NT 4.0 domain does not affect the ability of domain controllers to replicate, and does not affect the user logon process or users' access to resources.

This encryption is standard on domain controllers running Windows Server 2003 or Windows 2000 Service Pack 3 (SP3) or later. If you plan to install ADMT on a domain controller that is running Windows 2000 SP2 or earlier, install the 128-bit high encryption pack from the Windows 2000 High Encryption Pack link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For the PES in the Windows NT 4.0 source domain, install the 128-bit high encryption pack from the Windows NT 4.0 Service Pack 6a — High Encryption link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Important**

- The high encryption pack exists for North America Windows NT 4.0 Service Packs. If you need high encryption on an international Windows NT 4.0 domain controller, you need to install Microsoft Internet Explorer 4.1 or later and the high encryption pack for this Internet Explorer version.

# Establishing Required Trusts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can migrate accounts and resources from a Windows NT 4.0 source domain to an Active Directory target domain, you must ensure that the appropriate external trusts are in place. Trust relationships between the domains enable ADMT to migrate objects and translate local user profiles from the source to the target domains. In addition, depending on how they are configured, trust relationships can enable users in the source domains to access resources in the target domains, and can enable users in the target domains to access resources in source domains that have not yet been migrated.

In order to migrate users and global groups, you must establish a one-way trust between the Windows NT 4.0 source account domain and the target domain, so that the source domain trusts the target domain.

To migrate resources or translate local profiles, you must do one of the following:

- Create a one-way trust between the source resource domain and the target domain, so that the source resource domain trusts the target domain. This enables migrated users in the target domain to access resources in source resource domains that have not yet been migrated.

- Create a two-way trust between the source resource domain and the target domain. A two-way trust also enables any users that have been migrated to the target domain to access resources in the source domain.

# Establishing Migration Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To migrate accounts and resources, you must establish a migration account or accounts and assign the appropriate credentials to those accounts. Because ADMT requires only a limited set of credentials, creating separate migration accounts enables you to simplify administration. If the migration tasks for your organization are distributed across more than one group, it is helpful to create a migration account for each group involved in performing the migration.

Table 10.2 lists the credentials that are required in the source and target domains for different migration objects.

**Table 10.2 Migration Account Credentials**

| Migration Object | Credentials Necessary In Source Domain | Credentials Necessary in Target Domain |
|---|---|---|
| User/Group with SID history | Local administrator or domain administrator | Delegated permission on the user OU or the group OU, extended permission to migrate SID history, and local administrator on the computer on which ADMT is installed. |
| Computer | Domain administrator or administrator in the source domain and on each computer | Delegated permission on the computer OU and local administrator on the computer on which ADMT is installed. |
| Profile | Local administrator or domain administrator | Delegated permission on the user OU and local administrator on the computer on which ADMT is installed. |

The procedures provide examples for creating groups or accounts to migrate accounts and resources. Procedures differ according to whether a one-way trust or a two-way trust exists. The procedure for creating migration groups when a one-way trust exists is more complex than the procedure for creating migration accounts when a two-way trust exists, because you must add the migration group to the local Administrators group on local workstations. The procedure for creating migration groups when a one-way trust exists involves the creation of separate groups for migrating accounts and resources; however, you can combine *acct_migrators* and *res_migrators* into one group, if you do not need to separate them to delegate different sets of permissions.

**To create an account migration group when a one-way trust exists in which the source domain trusts the target domain**

1. In the target domain, create a global group, *acct_migrators*.

2. In the target domain, add the *acct_migrators* group to the Domain Admins group, or delegate administration of OUs that are targets for account migration to this group.

3. If you are migrating SID history, and you did not place the *acct_migrators* group in the Domain Admins group, grant the *acct_migrators* group the extended permission **Migrate SID History** on the target domain object. To do this:

    a. Start Active Directory Users and Computers, right-click the domain object, and then click **Properties**.

    b. Click the **Security** tab, click **Add**, and select **acct_migrators**.

    c. In the **Permissions for acct_migrators** box, click **Allow** for the Migrate SID History permission.

4. In the source domain, add the *acct_migrators* group to the Administrators group.

5. On each computer on which you plan to translate local profiles, add the *acct_migrators* group to the local Administrators group.

**To create the resource migration group when a one-way trust exists in which the source domain trusts the target domain**

1. In the target domain, create a global group, *res_migrators*.

2. In the target domain, add the *res_migrators* group to the Domain Admins group, or delegate administration of OUs that are targets for resource migration to this group.

3. In the source domain, add the *res_migrators* group to the Administrators group.

4. On each computer that you plan to migrate or on which you plan to perform security translation, add the *res_migrators* group to the local Administrators group.

**To create a resource migration account when a two-way trust exists between the source and target domains**

1. In the source domain, create a user account, *res_migrator*.

2. In the source domain, add the *res_migratos* account to the Domain Admins group. (The Domain Admins group is a member of the local Administrators group on every computer in the domain by default; therefore, you do not need to add it to the local Administrators group on every computer.)

3. In the target domain, delegate permissions on OUs that are targets for resource migration to the *res_migrators* account.

# Configuring the Source and Target Domains to Migrate SID History

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order to migrate SID history, you must ensure that the following conditions are met:

- A local group used to audit SID history operations exists in the source domain.
- TCP/IP client support is enabled on the source domain primary domain controller (PDC).
- Audit policies are enabled in the source and target domains.

You can configure these items manually before beginning the migration, or you can allow ADMT to configure them automatically the first time that it runs. If you want to configure them manually, use the following procedures:

**To create a local group in the source domain to support auditing**

- On the source domain PDC, create a local group *source_domain*$$$, where domain is the name of your source domain; for example, in a domain called Boston, create a local group called boston$$$.

**To enable TCP/IP client support on the source domain PDC**

1. On the source domain PDC, navigate to the following registry subkey:

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa**

2. Modify the registry entry **TcpipClientSupport**, of data type REG_DWORD, by setting the value to 1.
3. Restart the computer.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or on the Web at http://www.microsoft.com/reskit.

**To enable auditing in the Active Directory target domain**

1. Log on as an administrator to any domain controller in the target domain.
2. Start Active Directory Users and Computers, expand the domain, and then double-click the **Domain Controllers** OU.
3. On the **Group Policy** tab, click **Default Domain Controllers Policy**, and then click **Edit**.
4. Expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
5. Double-click **Audit account management**, click **Define these policy settings**, and then click both **Success** and **Failure**.

**To enable auditing in the Windows NT 4.0 source domain**

1. Log on as an administrator to any domain controller in the source domain.

2. Open **User Manager for Domains**, click **Policies**, and then click **Audit**.

3. For **User and Group Management**, verify that **Audit These Events** is selected and that **Success** and **Failure** are both selected.

# Configure the Target Domain OU Structure for Administration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory design team designs the OU structure for the target domain. They also define the groups that are responsible for the administration of each OU and the memberships of each group. You can use that information and the following procedures to configure the target domain for administration.

**To configure the target domain OU structure for administration**

1. Log on as an administrator to any domain controller in the target domain.
2. Start Active Directory Users and Computers**,** and then create the OU structure specified by your design team.
3. Create administrative groups and assign users to these groups.
4. Delegate the administration of the OU structure to groups as defined by your design team.

# Installing ADMT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Install and run ADMT on a domain controller in the target domain. The target domain owner is responsible for the migration of all Windows NT 4.0 accounts and resources into the Active Directory domain.

**To install ADMT in the target domain**

1. Log on to a domain controller in the target domain by using your ADMT migration account.

2. Insert the Windows Server 2003 operating system CD in the drive, and then navigate to \i386\admt.

3. Run admigration.msi.

# Enabling Password Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ADMT version 2 supports interforest password migration. If you intend to migrate passwords when you restructure your Windows NT 4.0 domains, you must enable password migration.

When you migrate passwords by using ADMT, you must use a PES to host password migration support DLLs. The PES can be any domain controller in the source domain that supports 128-bit encryption. Use the following procedures to enable support for password migration.

In the target domain, ensure that the built-in group Pre-Windows 2000 Compatible Access contains the Everyone and Anonymous Logon groups. The groups will not be present if you selected **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems** when you installed Active Directory in the target domain.

If the groups are not present, ADMT logs an Access Denied error. You must then manually add them to enable support for password migration. To do this, type the following at the command line on a target domain controller:

```
        NET LOCALGROUP "Pre-Windows 2000 Compatible Access" Everyone /ADD
NET LOCALGROUP "Pre-Windows 2000 Compatible Access" "Anonymous Logon" /ADD
```

After this update to the Pre-Windows 2000 Compatible Access group replicates, you must restart the Server service on all domain controllers in the target domain.

**To create an encryption key**

1. Log on to the domain controller in the target domain on which you installed ADMT by using your ADMT migration account.
2. Place a blank, formatted diskette in the floppy drive.
3. Open a command window, navigate to the folder in which ADMT is installed, and then, at the command line, type:

   ```
   ADMT KEY
   Source_DomainFloppy_Drive [password] [*]
   ```

   The *Source_Domain* is the NetBIOS name, not the fully qualified domain name (FQDN). A password, which provides key encryption, is optional. To protect the shared key, type either the password or an asterisk on the command line. The asterisk causes you to be prompted for a password that is not displayed on the screen.

| **Note** |
| --- |
| For security reasons, providing a password is strongly recommended. |

**To enable password migration on the source domain**

1. On the PES, insert the encryption key diskette.
2. Insert the Windows Server 2003 operating system CD in the CD drive.
3. Navigate to the \i386\admt\pwdmig folder, and then run pwdmig.exe. If you set a password during the key generation process on the domain controller in the target domain, the **Key Password Required** dialog box appears.
4. Type the password, and then complete the setup process.
5. Navigate to the following registry subkey on the source domain PES:
   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\LSA**
6. Modify the registry entry **AllowPasswordExport**, of data type REG_DWORD, by setting the value to 1.

| **Caution** |
| --- |

The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the "Registry Reference" on the *Windows Server 2003 Deployment Kit* companion CD or on the Microsoft Web site.

# Initializing ADMT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Initialize ADMT by running a test migration of a global group with the option to migrate SID history selected. If you did not previously configure the source and target domains to migrate SID history, you will receive an error and a prompt for each item that has not yet been configured. When you accept each prompt, ADMT automatically completes the following tasks, which are required in order to enable SID history migration:

- Creates a local group, *source_domain*$$$, in the source domain, which is used to audit SID history operations. Do not add members to this group; if you do, SID history migration will fail.

- Enables TCP/IP client support on the source domain PDC by setting to 1 the value of the registry entry **TcpipClientSupport**, which is locatedin the following subkey:

  HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa

  Setting **TcpipClientSupport** to 1 enables remote procedure calls (RPCs) over the TCP transport while preserving the security of the system.

- Enables audit policies in the source and target domains.

For more information about configuring the source and target domains to migrate SID history, see "Configuring the Source and Target Domains to Migrate SID History" earlier in this chapter.

**To initialize ADMT by running a test migration of a global group**

1. In the Active Directory Migration Tool console, on the **Action** menu, click **Group Account Migration Wizard**.

2. Complete the Group Account Migration Wizard by using the information provided in Table 10.3. Accept default settings when no information is specified.

**Table 10.3 Migrating Global Groups by Using the Group Account Migration Wizard**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Select **Test the migration settings and migrate later?** |
| Domain Selection | In the **Sourcedomain** box, type the source domain name.<br><br>In the **Target domain** box, type the target domain name. |
| Group Selection | Click **Add.**<br><br>In the **Select Groups** dialog box, select a group. |
| Organizational Unit Selection | Click **Browse**.<br><br>In the **Browse for Container** dialog box, navigate to the **Groups** OU. |
| Group Options | Select the **Migrate Group SIDs to target domain** check box.<br><br>Click **Do not rename accounts**. |
| User Account | Type the name and password of the migration account that you created for migrating accounts.<br><br>In the **Domain** box, type the NetBIOS name of the target domain. |

| | |
|---|---|
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

3. When the wizard has finished running, click **View Log**, and then review the migration log for any errors.

4. Verify that the test migration configured ADMT properly by ensuring that:

   - A new local group *source_domain$$$* exists in the source domain. This account supports ADMT auditing of SIDhistory migration.

   - The registry entry **TcpipClientSupport** is created, and its value set to 1, in the following subkey on the source domain PDC:

     HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa

   - The audit policy for account management is enabled on the source and target domains.

# Service Account Identification

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Identify the member servers and domain controllers that run applications as services by using a service account. A service account is a user account created explicitly to provide a security context for these applications. The service account is a standard user account that is assigned the **Log on as a service** user right.

**Note**

- This step does not migrate the accounts but only identifies them for later migration. To identify service accounts, ADMT dispatches an agent to the server in the source domain that records the services that are running on that server.

ADMT includes these accounts in the list of service accounts identified as not running under Local System account. These accounts must be updated on the computer after the migration of the accounts.

When service accounts are migrated later by means of the User Account Migration Wizard, ADMT makes the proper changes on the member server. If the service account is assigned rights by means of a group membership, you must also run the Security Translation Wizard to update the user rights and group memberships of the service account. For more information about migrating service accounts, see "Transitioning Service Accounts" later in this chapter.

You can identify service accounts by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To identify service accounts by using the ADMT console**

1. On the domain controller on which ADMT is installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool console, and then select **Service Account Migration Wizard**.
3. Complete the Service Account Migration Wizard by using the information in Table 10.4.

   **Table 10.4 Using the Service Account Migration Wizard**

| Wizard Page | Action |
|---|---|
| Domain Selection | In the **Source domain** box, type or select the NetBIOS name of the source domain.<br>In the **Target domain** box, type or select the NetBIOS or DNS name of the target domain. |
| Update Information | Select **Yes, update the information**. |
| Service Account Selection | Click **Add**.<br>In the **Select Computers** list box, select the names of all servers that have service accounts.<br>Click **OK**, and then click **Next**. |
| Service Account Information | Select any user accounts that do not need to be marked as service accounts in the ADMT database, and click **Skip/Include** to mark them as **Skip**. |

4. When the wizard has finished running on all computers, on the Server List page, click **View Log**. The log contains information about why the wizard did not succeed on any computers. For example, invalid permissions cause the wizard to not succeed.

**To identify service accounts by using the ADMT console**

1. On the domain controller on which ADMT is installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool console, and then select **Service Account Migration Wizard**.
3. Complete the Service Account Migration Wizard by using the information in Table 10.4.

**Table 10.4 Using the Service Account Migration Wizard**

| Wizard Page | Action |
|---|---|
| Domain Selection | In the **Source domain** box, type or select the NetBIOS name of the source domain.<br>In the **Target domain** box, type or select the NetBIOS or DNS name of the target domain. |
| Update Information | Select **Yes, update the information**. |
| Service Account Selection | Click **Add**.<br>In the **Select Computers** list box, select the names of all servers that have service accounts.<br>Click **OK**, and then click **Next**. |
| Service Account Information | Select any user accounts that do not need to be marked as service accounts in the ADMT database, and click **Skip/Include** to mark them as **Skip**. |

4. When the wizard has finished running on all computers, on the Server List page, click **View Log**. The log contains information about why the wizard did not succeed on any computers. For example, invalid permissions cause the wizard to not succeed.

**To identify service accounts by using the ADMT command-line option**

1. On a domain controller in the target domain on which ADMT installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT SERVICE /N "
computer_name1
" "
computer_name2
" [
parameters
]
```

*Computer_name1* and*computer_name2* are the names of computers in the source domain that run service accounts. You can append parameters to the command as follows:

```
ADMT SERVICE /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
"
```

Alternatively, you can include parameters in an options file that is specified on the command line as follows:

```
ADMT SERVICE /N "
computer_name1
" "
```

```
computer_name2
" /O:"
option_file.txt
"
```

Table 10.5 lists the parameters required for the identification of service accounts, along with the command-line parameter and option file equivalents.

**Table 10.5 Parameters Required for Service Account Identification**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |

3. Review the results that are displayed on the screen for any errors.

**To identify service accounts by using a script**

- Prepare a script that incorporates ADMT commands and options for identifying service accounts by using the sample script shown in Listing 10.3.

**Listing 10.3 Identifying Service Accounts**

```
<Job id=" IdentifyingServiceAccounts" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit
   Dim objMigration
   Dim objServiceAccountEnumeration
   '
   'Create instance of ADMT migration objects.
   '
   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objServiceAccountEnumeration = _
        objMigration.CreateServiceAccountEnumeration
   '
   'Specify general migration options.
   '
   objMigration.SourceDomain = " source domain"
   '
   'Enumerate service accounts on specified computers.
   '
   objServiceAccountEnumeration.Enumerate admtData, _
  Array(" computer name1" ," computer name2" )
   Set objServiceAccountEnumeration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a script file to assist you in creating a script to identify service accounts, see "Identifying Service Accounts" (DSSRENT_5.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Identifying Service Accounts" on the Web at http://www.microsoft.com/reskit).

# Restructuring Account Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Windows NT 4.0 account domains involves migrating the user, group, and service accounts that exist in the Windows NT 4.0 domain into an Active Directory domain. If account domains include resources, you must also migrate those resources. The Active Directory target domain owner is responsible for the migration of all Windows NT 4.0 accounts. The migration process must be completed for every Windows NT 4.0 source account domain.

Figure 10.7 shows the process for restructuring Windows NT 4.0 account domains.

**Figure 10.7 Restructuring Account Domains**



# Transitioning Service Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Begin the process of migrating objects by transitioning the service accounts that you selected when you identified service accounts. For more information about identifying service accounts, see "Service Account Identification" earlier in this chapter.

To transition service accounts, use ADMT to complete the following tasks:

- Migrate the service accounts from the Windows NT 4.0 source account domain to the target domain.
- Modify the services on each server in the source domain to use the service account in the target domain in place of the service account in the Windows NT 4.0 source account domain.

**Note**

- This process only updates the **Log on as a service** right. If a service account has local rights based on its membership in a local group, then you must update these rights by running the Security Translation Wizard.

You can transition service accounts by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To transition service accounts by using the ADMT console**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool console, and select **User Account Migration Wizard**.
3. Complete the User Account Migration Wizard by using the information provided in Table 10.6.

**Table 10.6 Using the User Account Migration to Transition Service Accounts**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| User Selection | Click **Add.**<br>In the **Select Users** dialog box, select all of the service accounts that were identified by the Service Account Migration Wizard, and then click **Add**.<br>Click **OK**. |
| Organizational Unit Selection | Click **Browse**.<br>In the **Browse for Container** dialog box, navigate to the target domain and select the container for the service accounts, and then click **OK**. |
| Password Options | Click **Complex passwords**. |
| Account Transition Options | Click **Enable target accounts**.<br>Select the **Migrate user SIDs to target domains** check box. |

| | |
|---|---|
| User Account | Type the user name, password, and domain of a user account that has administrative credentials. |
| User Options | Select the **Update user rights** check box.<br>Make sure that **Do Not Rename Accounts** is selected.<br>Ensure that no other settings are selected, including the **Migrate associated user groups** option. |
| Naming Conflicts | Select **Ignore conflicting accounts**. |
| Service Account Information | Select **Migrate all service accounts and update SCM for items marked include**. If you are also migrating other user accounts that are not service accounts, this wizard page tells you that you have selected some accounts that are marked as service accounts in the ADMT database. By default, the accounts are marked as Include. To change the status of the account, select the account, and then click the **Skip/Include** button.<br>Click **Next** to migrate the accounts. |

4. When the wizard has finished running, click **View Log**, and then review the migration log for any errors.
5. Start Active Directory Users and Computers, navigate to the OU that you created for service accounts, and then verify that the service accounts exist in the target domain OU.
6. Confirm that each application for which the service account was relocated continues to function correctly.

**To transition service accounts by using the ADMT command-line option**

1. On a domain controller in the target domain on which ADMT installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT USER /N "
user_name1
" "
user _name2
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT USER /N "
user _name1
" "
user _name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /MSS:
YES
/UUR:YES
```

Alternatively, you can include parameters in an option file that is specified on the command line as follows:

```
ADMT USER /N "
user_name1
" "
user_name2
" /O:"
option_file.txt
"
```

Table 10.7 lists the parameters required for transitioning service accounts, along with the command-line parameter and option file equivalents.

**Table 10.7 Parameters Required for Transitioning Service Accounts**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| Disable accounts | /DOT:ENABLETARGET (default) | DisableOption=ENABLETARGET (default) |
| Migrate password | /PO:COMPLEX (default) | PasswordOption= COMPLEX |
| Migrate user SIDs = YES | /MSS:YES | MigrateSIDs=YES |
| Update user rights=YES | /UUR:YES | UpdateUserRights=YES |
| Do not rename accounts | /RO:DONT (default) | RenameOption=DONT (default) |
| Ignore conflicting accounts | /CO:IGNORE (default) | ConflictOptions=IGNORE (default) |

3. Review the results that are displayed on the screen for any errors.
4. Start Active Directory Users and Computers and navigate to the target service account OU. Verify that the service accounts exist in the target domain OU.
5. Confirm that each application for which the service account was relocated continues to function correctly.

**To transition service accounts by using a script**

- Prepare a script that incorporates ADMT commands and options for transitioning service accounts by using the sample script shown in Listing 10.4.

**Listing 10.4 Transitioning Service Accounts**

```
<Job id=" TransitioningServiceAccountsNTSource" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit
   Dim objMigration
   Dim objUserMigration
   '
```

```
   'Create instance of ADMT migration objects.
   '
   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objUserMigration = objMigration.CreateUserMigration
   '
   'Specify general migration options.
   '
   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"
   objMigration.ConflictOptions = admtIgnoreConflicting
   '
   'Specify user migration specific options.
   '
   objUserMigration.DisableOption = admtEnableTarget
   objUserMigration.MigrateSIDs = True
   objUserMigration.UpdateUserRights = True
   objUserMigration.MigrateServiceAccounts = True

   '
   'Migrate specified service accounts.
   '
   objUserMigration.Migrate admtData, _
  Array(" service account name1" ," service account name2" )
   Set objUserMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a script file to assist you in creating a script to transition service accounts, see "Transitioning Service Accounts — Windows NT 4.0 Source" (DSSRENT_6.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Transitioning Service Accounts —Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

# Migrating Global Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order to preserve the global group user memberships, you must migrate global groups before you migrate users.

Global group migration involves the following steps:

1. The administrator selects the global group objects in the source domain. ADMT reads the objects in the source domain.
2. A new global group object is created in the target domain. A new primary SID is created for the object in the target domain.
3. To preserve resource access if you are migrating SID history, ADMT adds the SID of the global group in the source domain to the SID history attribute of the new global group in the target domain.

Following the migration, audit events are logged in both the source and the target domain.

**Note**

- If the user account migration process takes place over an extended period of time, then you might need to remigrate global groups from the source to the target domain to propagate membership changes that are made in the source domain before the migration process is complete. For more information about remigrating global groups, see "Remigrating Global Groups" later in this chapter.

You can migrate global groups by using the Active Directory Migration Tool console, by using the ADMT command-line option, or by using a script. For a worksheet to assist you in completing your initial global group migration, see "User and Group Object Assignment Table — Windows NT 4.0 Source" (DSSRENT_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "User and Group Object Assignment Table — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

**To migrate global groups from a Windows NT 4.0 source domain by using the ADMT console**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool console, and then select **Group Account Migration Wizard**.
3. Complete the Group Account Migration Wizard by using the information provided in Table 10.8.

**Table 10.8 Using the Group Account Migration Wizard to Migrate Groups**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Select **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS name of the source domain or select the name from a list.<br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain.<br>If ADMT includes the names of the source and target domains, ensure that they are correct. |
| Group Selection | Click **Add.**<br>In the **Select Groups** dialog box, select all global groups that you want to migrate (except built-in groups), click **Add**, and then click **OK**. |

| | |
|---|---|
| Organizational Unit Selection | Type in the name of the OU or click **Browse**.<br>In the **Browse for Container** dialog box, find the container in the target domain you want to move the global groups into, and then click **OK**. |
| Group Options | Select **Migrate Group SIDs to target domain**.<br>Select **Do not rename accounts**.<br>Make sure that all other options are not selected. |
| User Account | Enter **User name**, **Password**, and **Domain** of an account that has administrative rights in the source domain. |
| Naming Conflicts | Select **Ignore conflicting accounts and don't migrate**. |

4. When the wizard has finished running, click **View Log**, and then review the migration log for any errors.
5. Start Active Directory Users and Computers, and then navigate to the target OU. Verify that the global groups exist in the target domain OU.

**To migrate global groups from a Windows NT 4.0 source domain by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on using the ADMT account migration account.
2. At the command line, type:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /MSS:YES
```

Alternatively, you can include parameters in an option file that is specified on the command line as follows:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /O: "
option_file.txt
"
```

Table 10.9 lists the parameters required for migrating global groups, along with the command-line parameter and option file equivalents.

**Table 10.9 Parameters Required for Global Group Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|

| | | |
|---|---|---|
| Source domain | /SD:"*source_domain*" | SourceDomain="*source_domain*" |
| Target domain | /TD:"*target_domain*" | TargetDomain="*target_domain*" |
| Target OU location | /TO:"*target_OU*" | TargetOU="*target_OU*" |
| Migrate GG SIDs | /MSS:YES | MigrateSIDs=YES |
| Do not rename accts | /RO:DONT (default) | RenameOption=DONT |
| Ignore conflicting accts and do not migrate them | /CO:IGNORE (default) | ConflictOptions=IGNORE |

3. Review the results that are displayed on the screen for any errors.
4. Start Active Directory Users and Computers, and then navigate to the target OU. Verify that the global groups exist in the target domain OU.

**To migrate global groups by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating global groups by using the sample script shown in Listing 10.5.

**Listing 10.5 Migrating Global Groups**

```
<Job id=" MigratingGlobalGroupsNTSource" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objGroupMigration

   '
   'Create instance of ADMT migration objects.
   '
   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objGroupMigration = objMigration.CreateGroupMigration

   '
   'Specify general migration options.
   '
   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"

   '
   'Specify group migration specific options.
   '
   objGroupMigration.MigrateSIDs = True

   '
   'Migrate specified group objects.
   '
   objGroupMigration.Migrate admtData, Array(" group name1" ," group name2" )

   Set objGroupMigration = Nothing
```

```
   Set objMigration = Nothing
</Script>
</Job>
```

For a script file to assist you in creating a script to migrate global groups, see "Migrating Global Groups — Windows NT 4.0 Source" (DSSRENT_7.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Global Groups —Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

# Migrating Users in Batches

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To begin migrating users, first migrate a small number of users as a pilot and test whether they are able to access resources after the migration. Then migrate the remaining users in the source account domain in small batches of up to 100 users per batch. This helps to keep the migration process manageable.

Complete the following steps to migrate user accounts to the target domain:

1. Create test accounts that you can use to verify the success of each batch.
2. Migrate a single batch of users.
3. Translate any local user profiles for user accounts that are included in the batch.
4. Migrate workstations for users that are included in the batch.
5. Remigrate all global groups to update any group membership changes.
6. Notify users in the batch to log on to the target domain.
7. Repeat the steps for the next batch of users, until you have migrated all user accounts.

Until you migrate all user and group accounts, continue to administer global group membership in the Windows NT 4.0 source account domain. To support a rollback strategy, manually synchronize any changes you make to users in the target domain with the existing user accounts in the Windows NT 4.0 source account domain.

Figure 10.8 shows the steps you must complete for each batch of user accounts.

**Figure 10.8 Migrating Users in Batches**



# Migrating User Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The ADMT user account migration process includes the following steps:

1. The administrator selects the source user objects to be migrated.
2. ADMT reads attributes from the source user objects.
3. ADMT creates a new user object in the target domain and a new primary SID for the new user account.
4. If you are migrating SID history, ADMT adds the original SID of the user account to the SID history attribute of the new user account.
5. ADMT migrates the password for the user account.

You cannot migrate every user property when you migrate user accounts. For example, Protected Storage (Pstore) contents for Windows NT 4.0 workstations, including Encrypting File System (EFS) private keys, are not migrated by ADMT when you migrate user accounts. To migrate Pstore contents, you must export and import keys during the migration process.

For clients that are running Windows 2000 or later operating systems, data that is protected by the Data Protection API (DPAPI) is also not migrated. DPAPI helps protect the following items:

- Web page credentials (for example, passwords)
- File share credentials
- Private keys associated with EFS, S/MIME, and other certificates
- Program data that is protected by using the CryptProtectData() function

For this reason, it is important to test user migrations by using your test migration account to identify any properties that did not migrate and update user configurations in the target domain accordingly.

Following the migration, audit events are logged in both the source and the target domains if you are migrating SID history.

Using ADMT to migrate user accounts preserves group memberships. Because global groups can contain only members from the domain in which the group is located, when users are migrated to a new domain, the user accounts in the target domain cannot be members of the global groups in the source domain.

As part of the migration process, ADMT identifies the global groups in the source domain that the user account belongs to, and then determines whether the global groups have been migrated. If ADMT identifies global groups in the target domain that the users belonged to in the source domain, the tool adds the users to the appropriate groups in the target domain.

Using ADMT to migrate user accounts also preserves user passwords. After the user accounts are migrated to and enabled in the target domain, the users can log on to the target domain by using their original passwords. If the user account migration is successful but the password migration process fails, then ADMT creates a new complex password for the user account in the target domain. By default, ADMT stores new complex passwords that are created in the default drive, in Program Files\Active Directory Migration Tool\Logs\Password.txt file.

If you have a Group Policy setting on the target domain that does not allow blank passwords (the **Default Domain Policy/Computer Configuration/Security Settings/Account Policies/Password Policy/Minimum password length** setting is set to any number other than zero), then password migration will fail for any user who has a blank password. ADMT generates a complex password for that user, and writes an error to the error log.

Establish a method for notifying users who have been assigned new passwords. For example, you can create a script to send an e-mail message to users to notify them of their new passwords.

**Important**

- Because only a hash of a user password exists in the source domain, the password filter cannot verify whether the password meets complexity or length requirements. The target domain controller used to set the password can verify the password history because it compares the hash of the password against previous hashes.

You can migrate user accounts by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate the current batch of users by using the ADMT console**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.

2. Open the Active Directory Migration Tool, and then select **User Account Migration Wizard**.

3. Complete the User Account Migration Wizard by using the information provided in Table 10.10.

   **Table 10.10 Using the User Account Migration Wizard to Migrate User Accounts**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the NetBIOS name of the source domain.<br><br>In the **Target domain** box, type or select the NetBIOS or DNS name of the target domain. |
| User Selection | Click **Add.**<br><br>In the **Select Users** dialog box, select the user accounts you want to migrate, and click **Add**.<br><br>Click **OK**. |
| Organizational Unit Selection | ADMT lists an OU here. Ensure that this is the correct target OU. If it is not correct, type in the correct OU or click **Browse**.<br><br>In the **Browse for Container** dialog box, navigate to the target domain and OU, and then click **OK**. |
| Password Options | Select **Migrate Passwords**. |
| Account Transition Options | Select **Enable target accounts**.<br><br>Select **Days Until Source Account Expires** and enter a number of days. A value of 7 is commonly used.<br><br>Select the **Migrate user SIDs to target domains** check box. |
| User Account | Type in the **User name**, **Password**, and **Domain** of a user account that has administrative credentials. |
| User Options | Select the **Translate roaming profiles** check box.<br><br>Select the **Update user rights** check box.<br><br>Clear the **Migrate associated user groups** check box.<br><br>Select the **Fix users' group memberships**.<br><br>Select the **Do not rename accounts** check box. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

4. When the wizard has finished running, click **View Log** and review the migration log for any errors.

5. Start Active Directory Users and Computers, and then verify that the user accounts exist in the appropriate OU in the

target domain.

**To migrate the current batch of users by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT installed, log on by using the ADMT account migration account.

2. To migrate a set of specified users "*user_name1*" "*user_name2*," at the command line, type:

   ```
   ADMT USER /N "user_name1" "user_name2" [parameters]
   ```

   You can append parameters to the command as follows:

   ```
   ADMT USER /N "user_name1" "user_name2" /SD:"source_domain" /TD:"target_domain"
   /TO:"target_OU" /MSS:YES TRP:YES /UUR:YES
   ```

   Alternatively, you can include parameters in an option file that is specified on the command line as follows:

   ```
   ADMT USER /N "user_name1" "user_name2" /O "option_file.txt"
   ```

   Table 10.11 lists the parameters required for migrating user accounts, along with the command-line parameter and option file equivalents.

   **Table 10.11 Parameters Required for User Migrations**

| Options | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:`"*source_domain*" | `SourceDomain=`"*source_ domain*" |
| Target domain | `/TD:`"*target_domain*" | `TargetDomain=`"*target_ domain*" |
| Target OU location | `/TO:`"*target_OU*" | `TargetOU=`"*target_OU*" |
| Migrate SIDs | `/MSS:YES` | `MigrateSIDs=YES` |
| Password options | `/PO:COPY`<br>`/PS:`"*server name*" | `PasswordOption=COPY`<br>`PasswordServer=`"*server name*" |
| Do not rename accts | `/RO:DONT` (default) | `RenameOption=DONT` |
| Ignore conflicting accts and do not migrate them | `/CO:IGNORE` (default) | `ConflictOptions=IGNORE` |
| Translate Roaming Profile | `/TRP:YES (default)` | `TranslateRoamingProfile=YES` |
| Update User Rights | `/UUR:YES(/UUR:NO is the default)` | `UpdateUserRights=NO` |

3. Review the results that are displayed on the screen for any errors.

4. Start Active Directory Users and Computers, and then navigate to the target OU. Verify that the users exist in the target domain OU.

**To migrate users by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating users by using the sample script shown in Listing 10.6.

  **Listing 10.6 Migrating User Accounts**

```
<Job id=" MigratingUserAccountsNTSource" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objUserMigration

   '
   'Create instance of ADMT migration objects.
   '
   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objUserMigration = objMigration.CreateUserMigration

   '
   'Specify general migration options.
   '
   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"
   objMigration.PasswordOption = admtCopyPassword
   objMigration.PasswordServer = " password export server name"
   objMigration.ConflictOptions = admtIgnoreConflicting

   '
   'Specify user migration specific options.
   '
   objUserMigration.MigrateSIDs = True
   objUserMigration.TranslateRoamingProfile = True
   objUserMigration.UpdateUserRights = True
   objUserMigration.FixGroupMembership = True
   objUserMigration.MigrateServiceAccounts = False

   '
   'Migrate specified user objects.
   '

   objUserMigration.Migrate admtData, Array(" user name1" ," user name2" )

   Set objUserMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in creating a script to migrate user accounts, see "Migrating User Accounts — Windows NT 4.0 Source" (DSSRENT_8.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating User Accounts — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

# Translating Local User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Local profiles contain the desktop state and user data for users in the source account domain. ADMT can migrate local profiles only for workstations running Windows NT 4.0, Windows 2000, and Microsoft® Windows® XP. Migrate local user profiles for a batch of users immediately after migrating the batch of users, and before the users log on to the target domain.

It is important to verify that user profile translation succeeded for each user before that user is allowed to log on. If the user profile translation fails for a user, that user must not log on to the target domain. Roll the user back manually, by disabling the user account in the target domain, and enabling the account in the source domain.

You can migrate roaming user profiles during the user account migration process. However, you still need to translate the local copy of the profile.

**Note**

- Migrate the local user profiles the night before you notify the users to log on by using their new accounts in the target domain. Migrating profiles the night before ensures that the new user profile reflects the most current user settings.

Before you migrate local user profiles, create a list of the workstations in the source domain. You can migrate local user profiles by running the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate local user profiles by using the ADMT console**

1. For each workstation in the source domain that is running Windows NT 4.0, verify that the ADMT account migration account was added as a member of the local Administrators group.
2. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT resource migration account.
3. Open the Active Directory Migration Tool, and then run the Security Translation Wizard.
4. Complete the Security Translation Wizard by using the information in Table 10.12.

**Table 10.12 Using the Security Translation Wizard to Migrate Local User Profiles**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Select **Migrate Now?** |
| Security Translation Options | Select **Previously migrated objects**. |
| Domain Selection | In the **Source domain** box, type or select the name of the domain that contains the user accounts.<br>In the **Target domain** box, type or select the name of the target domain. |
| Computer Selection | Click **Add** and add computers from the domain that contains the computer accounts. (Depending on your existing domain structure, this might be a resource domain.) |
| Translate Objects | Select **User Profiles**. |
| Security Translation Options | Select **Replace**. |

After migrating local user profiles, you must verify that the migration process was successful. Complete the following steps to

verify that the user profile migration was successful:

1. Review the status message for each computer for which you migrated the user profile.

2. For computers for which the status is not **Success**, review the message to identify the cause of the failure. For more details about the cause of the failure, click **View Dispatch Log**.

3. For each computer for which the agent completed successfully, select the computer name from the list, and click **Agent Detail**.

4. Click **View Log** and review the log for any errors that occurred during the local profile migration process.

5. For computers to which an agent failed to be dispatched, fix the cause of the problem, and then run the Retry Wizard to dispatch agents to those computers.

6. For any computers for which the agent failed to migrate the local profile, fix the cause of the problem, and then perform the local profile migration again on those computers.

You can use the following procedure to migrate local user profiles by using the ADMT command-line option

**To migrate local user profiles by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TOT:ADD /TUP:YES
```

Alternatively, you can include parameters in an option file that is specified on the command line as follows:

```
ADMT USER /N "
user_name1
" "
user_name2
" /O:"
option_file.txt
"
```

Table 10.13 lists the parameters required for migrating user accounts, along with the command-line parameter and option file equivalents.

**Table 10.13 Parameters Required for Local User Profile Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_ domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_ domain*`"` |
| Translate option | /TOT:REPLACE | TranslateOption=REPLACE |
| Modify local user profile security | /TUP:YES | TranslateUserProfiles=YES |

1. Review the results that are displayed on the screen for any errors.

2. Review the dispatch log to determine whether agents were successfully dispatched to every computer.

3. Review the migration log for errors related to local profiles on workstations.

4. For computers to which the agent failed to be dispatched, fix the cause of the problem, and then run the Retry Wizard to dispatch agents to those computers.

5. For any computers for which the agent failed to migrate the local profile, fix the cause of the problem, and then perform the local profile migration again on those computers.

Note that it is important to verify the success of the local profile migration before users attempt to log on to the target domain.

If a user logs on to the target domain by using their new target account and their local profile did not migrate successfully, then you must do the following:

1. On the computer of the user, delete the following registry subkey, where *SID* is the SID of the user account in the target domain:

   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\ProfileList*SID*

2. Remigrate the local user profile for that computer.

**Caution**

- The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the "Registry Reference" on the *Windows Server 2003 Deployment Kit* companion CD or on the Microsoft Web site.

**To migrate local user profiles by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating local profiles by using the sample script shown in Listing 10.7. To verify that the local user profile migration was successful, use the verification process listed in steps 3 through 7 of the command-line procedure.

**Listing 10.7 Migrating Local Profiles**

```
<Job id=" MigratingLocalProfilesNTSource" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
```

```
   Dim objSecurityTranslation

   '
   'Create instance of ADMT migration objects.
   '
   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objSecurityTranslation = objMigration.CreateSecurityTranslation

   '
   'Specify general migration options.
   '
   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " Computers"

   '
   'Specify security translation specific options.
   '
   objSecurityTranslation.TranslationOption = admtTranslateReplace
   objSecurityTranslation.TranslateUserProfiles = True

   '
   'Perform security translation on specified computer objects.
   '
   objSecurityTranslation.Translate admtData, _
  Array(" computer name1" ," computer name2" )

   Set objSecurityTranslation = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script to assist you in migrating local profiles, see "Migrating Local Profiles — Windows NT 4.0 Source" (DSSRENT_9.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Local Profiles — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

# Migrating User Workstations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For each batch of user accounts that you migrate, migrate a batch containing those users' workstations, so the user account and the workstation exist in the target domain when the user logs on. When you move workstations between domains, the Security Accounts Manager (SAM) database moves along with them. Accounts located in the local SAM database (such as local groups) that are used to enable access to resources always move with the computer, and therefore do not need to be migrated.

**Note**

- Restart workstations immediately after joining them to the target domain, by selecting a low number (such as 1) for the **RestartDelay** parameter. Resources that are not restarted following migration are in an indeterminate state.

You can migrate workstations by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate workstations by using the ADMT console**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT resource migration account.
2. Open the Active Directory Migration Tool and select **Computer Account Migration Wizard**.
3. Complete the Computer Account Migration Wizard by using the information provided in Table 10.14.

   **Table 10.14 Using the Computer Account Migration Wizard to Migrate Workstations**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Select **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the NetBIOS name of the source domain.<br>In the **Target domain** box, type or select the NetBIOS or DNS name of the target domain. |
| Computer Selection | Click **Add.**<br>In the **Select Computers** dialog box, select the names of all the workstations in the source domain that are being migrated in this batch, click **Add**, and then click **OK**. |
| Organizational Unit Selection | Click **Browse**.<br>In the **Browse for Container** dialog box, find the container in the target domain to which the computers are migrating, and then click **OK**. |
| Translate Objects | Ensure that no boxes are selected. |
| Computer Options | In the **Minutes before computer restart after wizard completion** box, type **1**.<br>Select **Do not rename computers**. |
| Naming Conflicts | Select **Ignore conflicting accounts and don't migrate.** |

4. When the wizard has finished running on all computers, click **View Log** and review the migration log for any errors. If errors are reported, run the Retry Wizard to dispatch agents to workstations for which the migration process failed.
5. Start Active Directory Users and Computers, then and verify that the workstations exist in the appropriate OU in the target domain.

**To migrate workstations by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT installed, log on by using the ADMT resource migration account.
2. At the command line, type:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /RDL:1
```

Alternatively, you can include parameters in an option file that is specified on the command line as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /O "
option_file.txt
"
```

Table 10.15 lists the parameters required for migrating workstations, along with the command-line parameter and option file equivalents.

**Table 10.15 Parameters Required for Workstation Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:`"*source_domain*" | `SourceDomain=`"*source_ domain*" |
| Target domain | `/TD:`"*target_domain*" | `TargetDomain=`"*target_ domain*" |
| Target OU location | `/TO:`"*target_OU*" | `TargetOU=`"*target_OU*" |
| • Restart computer one minute after wizard completes | /RDL:1 | RestartDelay=1 |

| | | |
|---|---|---|
| • Do not rename accounts | /RO:DONT (default) | RenameOption=DONT |
| • Ignore conflicting accounts and do not migrate them | /CO:IGNORE (default) | ConflictOptions=IGNORE |

3. Review the results that are displayed on the screen for any errors.
4. Start Active Directory Users and Computers, and then navigate to the target OU. Verify that the workstations exist in the target OU.

**To migrate workstations by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating workstations by using the sample script shown in Listing 10.8.

**Listing 10.8 Migrating Workstations or Member Servers**

```
<Job id=" MigratingWorkstationsMemberServersNTSource" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objComputerMigration

   '
   'Create instance of ADMT migration objects.
   '
   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objComputerMigration = objMigration.CreateComputerMigration

   '
   'Specify general migration options.
   '
   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " Computers"

   '
   'Specify computer migration specific options.
   '
   objComputerMigration.RestartDelay = 1

   '
   'Migrate computer objects on specified computer objects.
   '
   objComputerMigration.Migrate admtData, _
  Array(" computer name1" ," computer name2" )

   Set objComputerMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in migrating workstations, see "Migrating Workstations or Member Servers — Windows NT 4.0 Source" (DSSRENT_10.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations or Member Servers — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

# Remigrating Global Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A large user account migration might take place over an extended period of time. For this reason, you might need to remigrate global groups from the source to the target domain to reflect changes made to the groups in the source domain after the initial global group migration occurred. This procedure updates global group membership without overwriting previously migrated user accounts.

| **Note** |
| --- |
| Do not migrate global groups during peak work hours. The global group migration process consumes a large amount of network resources and resources on the domain controller that is running ADMT. |

You can migrate global groups by using the Active Directory Migration Tool console, by using the ADMT command-line option, or by using a script.

**To remigrate global groups from a Windows NT 4.0 source domain by using the ADMT console**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool console, and select **Group Account Migration Wizard**.
3. Complete the Group Account Migration Wizard by using the information provided in Table 10.16.

**Table 10.16 Using the Group Account Migration Wizard to Remigrate Groups**

| Wizard Page | Action |
| --- | --- |
| Test or Make Changes | Select **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS name of the source domain or select the name from a list.<br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain.<br>If ADMT includes the names of the source and target domains, ensure that they are correct. |
| Group Selection | Click **Add**.<br>In the **Select Groups** dialog box, select all global groups that you want to migrate (except built-in groups), click **Add**, and then click **OK**. |
| Organizational Unit Selection | Type in the name of the OU or click **Browse**.<br>In the **Browse for Container** dialog box, find the container in the target domain you want to move the global groups into, and then click **OK**. |
| Group Options | Select **Update user rights**.<br>Ensure that **Copy group members** is not selected.<br>Ensure that **Update previously migrated objects** is not selected.<br>Select **Fix membership of group**.<br>Select **Migrate Group SIDs to target domain**.<br>Select **Do not rename accounts**. |

| | |
|---|---|
| User Account | Enter **User name**, **Password**, and **Domain** of an account that has administrative rights in the source domain. |
| Naming Conflicts | Select **Replace conflicting accounts**.<br>Select **Remove existing user rights**.<br>Select **Remove existing members of groups being replaced**.<br>Ensure that **Move replaced accounts to specified target Organizational Unit** is not selected. |

4. When the wizard has finished running, click **View Log** and review the migration log for any errors.
5. Start Active Directory Users and Computers, and then navigate to the target OU. Verify that the global groups exist in the target domain OU.

**To remigrate global groups from a Windows NT 4.0 source domain by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on using the ADMT account migration account.
2. At the command line, type:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /MSS:YES /UGR:YES /MMS:NO
/CO:REPLACE+REMOVEUSERRIGHTS+REMOVEMEMBERS
```

Alternatively, you can include parameters in an option file that is specified on the command line as follows:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /O: "
option_file.txt
"
```

Table 10.17 lists the parameters required for remigrating global groups, along with the command-line parameter and option file equivalents.

**Table 10.17 Parameters Required for Global Group Remigrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:`"*source_domain*" | `SourceDomain=`"*source_ domain*" |
| Target domain | `/TD:`"*target_domain*" | `TargetDomain=`"*target_ domain*" |
| Target OU location | `/TO:`"*target_OU*" | `TargetOU=`"*target_OU*" |

| | | |
|---|---|---|
| Migrate GG SIDs | /MSS:YES | MigrateSIDs=YES |
| Do not rename accts | /RO:DONT (default) | RenameOption=DONT |
| Update user rights | /UGR:YES | UpdateGroupRights=YES |
| Copy group members | /MMS:NO | MigrateMembers=NO |
| Ignore conflicting accts and do not migrate them | /CO:REPLACE+REMOVEUSER RIGHTS+REMOVEMEMBERS | ConflictOptions=REPLACE+REMOVEUSERRIGHTS+REMOVE MEMBERS |

3. Review the results that are displayed on the screen for any errors.
4. Start Active Directory Users and Computers, and then navigate to the target OU. Verify that the global groups exist in the target domain OU.

**To remigrate global groups by using a script**

- Prepare a script that incorporates ADMT commands and options for remigrating global groups by using the sample script shown in Listing 10.9.

**Listing 10.9 Remigrating Global Groups**

```
<Job id=" RemigratingGlobalGroupsNTSource" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objGroupMigration

   '
   'Create instance of ADMT migration objects.
   '
   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objGroupMigration = objMigration.CreateGroupMigration

   '
   'Specify general migration options.
   '
   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"
   objMigration.ConflictOptions = admtReplaceConflicting + _
  admtRemoveExistingUserRights + admtRemoveExistingMembers

   '
   'Specify group migration specific options.
   '
   objGroupMigration.MigrateSIDs = True
   objGroupMigration.UpdateGroupRights = True
   objGroupMigration.UpdatePreviouslyMigratedObjects = False
   objGroupMigration.FixGroupMembership = True
   objGroupMigration.MigrateMembers = False

   '
   'Migrate specified group objects.
   '
```

```
   objGroupMigration.Migrate admtData, Array(" group name1" ," group name2" )

   Set objGroupMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a script file to assist you in creating a script to migrate global groups, see "Remigrating Global Groups — Windows NT 4.0 Source" (DSSRENT_11.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Remigrating Global Groups — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

# Completing the Account Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you migrate all batches of accounts from the source to the target domain, you can complete the account migration process. This involves the following steps:

- Transferring user and group administration to the target domain by doing the following:
    - Freezing user and group administration in the source domain.
    - Maintaining at least two operational domain controllers in the source domain until the account migration process is complete.
    - Backing up the two domain controllers in the source domain.
    - Beginning administration of user and group accounts in the target domain.
- Decommissioning the Windows NT 4.0 source account domain by doing the following:

  **Note**

    - Ensure that you retain a full system backup of the PDC for each account domain. This enables you to bring the account domain back online.
    - Removing all trust relationships involving the Windows NT 4.0 source account domain.
    - Repurposing any remaining account domain controllers in the source account domain that you did not migrate to the target domain.
    - Disabling all accounts that you created during the migration process, including those accounts to which you assigned administrative credentials.

When you decommission the source account domain, shared local groups and local groups in resource domains that you have not translated by using the Security Translation Wizard display group members as **account unknown** because member names from the source domain do not resolve. Those group memberships still exist, however, and this does not impact users.

**Note**

- Do not delete shared local groups and local groups that display **account unknown** because this disables the access facilitated by SID history. Run the Security Translation Wizard to remove these entries.

# Restructuring Resource Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Windows NT 4.0 resource domains into an Active Directory domain involves migrating resources, such as workstation and server computer accounts, and the associated shared local groups to the new environment. The target domain owner is responsible for the migration of all Windows NT 4.0 resources. The migration process must be completed for every Windows NT 4.0 source resource domain.

Figure 10.9 shows the process for restructuring Windows NT 4.0 resource domains.

**Figure 10.9 Restructuring Windows NT 4.0 Resource Domains**



# Migrating Workstations and Member Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Migrate remaining workstations that you did not migrate during the user account migration process, along with member servers, in small batches of up to 100 computers. Workstation account and member server migration is a straightforward process. Workstations and member servers have their own SAM account database. When you move workstations and member servers between domains, the database moves along with them. Accounts located in the local SAM database (such as local groups) that are used to enable access to resources always move with the computer, and therefore do not need to be migrated.

**Note**

- Restart member workstations and servers immediately after joining them to the target domain, by selecting a low number (such as 1) for the RestartDelay parameter. Resources that are not restarted following migration are in an indeterminate state.

You can migrate workstations and member servers by running the ADMT console, by using the ADMT command-line procedure, or by using a script.

**To migrate Windows NT 4.0 workstations and member servers by using the ADMT console**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT resource migration account.
2. Open the Active Directory Migration Tool and select **Computer Account Migration Wizard**.
3. Complete the Computer Account Migration Wizard by using the information provided in Table 10.18.

   **Table 10.18 Using the Computer Account Migration Wizard to Migrate Workstations and Member Servers**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Select **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the NetBIOS name of the source domain.<br>In the **Target domain** box, type or select the NetBIOS or DNS name of the target domain. |
| Computer Selection | Click **Add.**<br>In the **Select Computers** dialog box, select the names of all the workstations and member servers in the source domain, click **Add**, and then click **OK**. |
| Organizational Unit Selection | Click **Browse**.<br>In the **Browse for Container** dialog box, find the container in the target domain to which the computers are migrating, and then click **OK**. |
| Translate Objects | Ensure that no check boxes are selected. |
| Computer Options | In the **Minutes before computer restart after wizard completion** box, type **1**. |
| Naming Conflicts | Select **Ignore conflicting accounts and don't migrate.** |

4. When the wizard has finished running on all computers, click **View Log** and review the migration log for any errors. If errors are reported, run the Retry Wizard to dispatch agents to workstations and member servers for which the migration process failed.
5. Start Active Directory Users and Computers**,** and then verify that the workstations and member servers exist in the

appropriate OU in the target domain.

**To migrate workstations and member servers by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /RDL:1
```

Alternatively, you can include parameters in an option file that is specified on the command line as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /O "
option_file.txt
"
```

Table 10.19 lists the parameters required for migrating workstations and member servers, along with the command-line parameter and option file equivalents.

**Table 10.19 Parameters Required for Workstation and Member Server Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:`"*source_domain*" | `SourceDomain=`"*source_ domain*" |
| Target domain | `/TD:`"*target_domain*" | `TargetDomain=`"*target_ domain*" |
| Target OU location | `/TO:`"*target_OU*" | `TargetOU=`"*target_OU*" |

| | | |
|---|---|---|
| • Restart computer one minute after wizard completes | /RDL:1 | RestartDelay=1 |
| • Do not rename accounts | /RO:DONT (default) | RenameOption=DONT |
| • Ignore conflicting accounts and do not migrate them | /CO:IGNORE (default) | ConflictOptions=IGNORE |

3. Review the results that are displayed on the screen for any errors.
4. Start Active Directory Users and Computers, and then navigate to the target OUs. Verify that the workstations and member servers exist in the target OUs.

**To migrate workstations and member servers by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating workstations and member servers.

  For a sample script file to assist you in migrating workstations and member servers, see "Migrating User Workstations" earlier in this chapter and "Migrating Workstations or Member Servers — Windows NT 4.0 Source" (DSSRENT_10.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations or Member Servers — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

# Migrating Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In some domains, domain controllers are dedicated servers used for authentication only. In other domains, resources are located on domain controllers. Shared local groups provide resource permissions to all domain controllers that support a resource.

Migrating domain controllers is a more complicated process than migrating member servers for the following reasons:

- Windows NT 4.0 backup domain controllers (BDCs) cannot be returned to member server status or moved between domains.
- All domain controller access rights are based on shared local groups that must also be migrated.

For this reason, only migrate domain controllers running Windows NT 4.0 when shared resources are located on the domain controller. Decommission any domain controllers that do not have shared resources.

When deciding the order in which to migrate resource domains, consider the following:

- Migrate resource domains that use domain controllers as dedicated domain controllers first, then migrate resource domains that use domain controllers hosting resources.
- Migrate all the backup domain controllers in a domain first and then migrate the PDC.

**Migrating Shared Local Groups**

Migrate the shared local groups from the source domain to the target domain before you upgrade the Windows NT 4.0 domain controllers to Windows 2000 or Windows Server 2003 and move them to the target domain. Note that it is not necessary to change any access control lists (ACLs) as part of this process. The ACLs continue to reference the shared local groups in the source resource domain. Because the shared local groups were migrated to the target domain as domain local groups (and have the original SID in the SID history), users are able to access the resources as before. ADMT retains the membership of the local group during the migration.

**To migrate shared local groups by using the ADMT console**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT resource migration account.
2. Open the Active Directory Migration Tool and select **Group Account Migration Wizard**.
3. Complete the Group Account Migration Wizard by using the information provided in Table 10.20.

**Table 10.20 Using the Group Account Migration Wizard to Migrate Shared Local Groups**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Select **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS name of the source domain or select the name from a list.<br><br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain. |
| Group Selection | Click **Add.**<br><br>In the **Select Groups** dialog box, select all domain local groups that you need to migrate (except built-in groups), click **Add**, and then click **OK**. |

| | |
|---|---|
| Organizational Unit Selection | Type in the name of the OU or click **Browse**.<br><br>In the **Browse for Container** dialog box, find the container in the target domain to which the local groups are migrating, and then click **OK**. |
| Group Options | Select **Fix membership of group**.<br><br>Select **Migrate Group SIDs to target domain**.<br><br>Select **Do not rename accounts**.<br><br>Ensure that all other options are not selected. |
| User Account | Enter the **User name**, **Password**, and **Domain** of an account that has administrative rights in the source domain. |
| Naming Conflicts | Select **Ignore conflicting accounts and don't migrate**. |

4. When the wizard has finished running, click **View Dispatch Log**. Review the migration log for any errors.

5. Start Active Directory Users and Computers, and then navigate tothe target OU. Verify that the shared local groups exist in the target domain OU.

**Migrating Backup Domain Controllers**

Use the following steps to migrate Windows NT 4.0 BDCs from the source resource domain to the target domain.

**To migrate backup domain controllers**

1. Take the BDC offline, then use Server Manager to promote it to a PDC.

2. Upgrade the PDC to a Windows 2000– or Windows Server 2003–based domain controller.

3. Return the domain controller to stand-alone server status by running the Active Directory Installation Wizard. If you are removing Active Directory from the last domain controller in the domain, select **Domain controller is the last controller in the domain**.

4. Join the stand-alone server to the target domain.

# Completing the Resource Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you migrate all the resources from the source to the target domain, complete the restructuring process by doing the following:

- Transferring the administration of user and group accounts from the source to the target domain.
- Ensuring that at least two domain controllers are operational in the source resource domain until the resource migration process is complete.
- Backing up the two domain controllers in the source resource domain.

After completing these steps, you can translate security on the member servers in the target domain and decommission the source resource domain.

# Translating Security on Member Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Translate security on member servers to clean up ACLs. After users and groups are migrated to the target domain, ACLs still contain security IDs from the source domain. The SIDs from the source domain remain in the ACLs, and users can access resources by using SID history while the migration is in progress. However, after the migration is complete, the SIDs from the source domain must be translated to SIDs in the target domain. Use the Security Translation Wizard in ADMT to translate the source domain SIDs to the target domain SIDs that are contained in the ACLs.

**To translate security on member servers**

- Complete the Security Translation Wizard in ADMT by using the information provided in Table 10.21.

**Table 10.21 Using the ADMT Security Translation Wizard**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Select **Migrate Now?** |
| Security Translation Options | Select **Previously migrated objects**. |
| Domain Selection | In the **Source domain** box, type or select the name of the domain that contains the user accounts.<br><br>In the **Target domain** box, type or select the name of the target domain. |
| Translate Objects | Clear the **User Profiles** check box.<br><br>Select all other options. |
| Security Translation Options | Select **Add**. |

# Decommissioning the Source Resource Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you complete the resource migration, decommission the Windows NT 4.0 source resource domains.

**Note**

- Ensure that you retain a full system backup of the PDC for each resource domain. This enables you to bring the resource domain back online.

To decommission the Windows NT 4.0 source resource domain:

1. Remove all trust relationships between the Windows NT 4.0 source resource domain and the Active Directory target domain.

2. Repurpose any remaining domain controllers in the source resource domain that you did not migrate to the target domain.

3. Disable all accounts that you created during the migration process, including those accounts to which you assigned administrative permissions.

# Additional Resources for Restructuring Windows NT 4.0 Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing the Active Directory Logical Structure" in this book.
- "Designing the Site Topology" in this book
- "Deploying the Windows Server 2003 Forest Root Domain" in this book.
- "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.
- The Active Directory Branch Office Planning Guide link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for a complete guide to information involving Active Directory branch office implementations.

**Related Tools**

- ADMT v2

  The Active Directory Migration Tool (ADMT) version 2 is a tool that you can use to perform the most common migration tasks, such as migrating users, groups, and computers, mirroring trusts, and translating SIDs. It is implemented as an MMC snap-in, a command-line executable, and a COM object implemented in ADMTScript.dll.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- Windows Support Toolsunder Tools in Help and Support Center for Windows Server 2003.

**Related Job Aids**

- "User and Group Object Assignment Table — Windows NT 4.0 Source" (DSSRENT_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "User and Group Object Assignment Table — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).
- "Resource Object Assignment Table — Windows NT 4.0 Source" (DSSRENT_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Resource Object Assignment Table — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).
- "Migration Test Matrix" (DSSREER_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migration Test Matrix" on the Web at http://www.microsoft.com/reskit).
- "ADMT Option File Reference" (DSSRENT_3.txt) on the *Windows Server 2003 Deployment Kit* companion CD (or see "ADMT Option File Reference" on the Web at http://www.microsoft.com/reskit).
- "ADMT Scripting Constants" (AdmtConstants.vbs) on the *Windows Server 2003 Deployment Kit* companion CD (or see "ADMT Scripting Constants" on the Web at http://www.microsoft.com/reskit).

- "Identifying Service Accounts" (DSSRENT_5.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Identifying Service Accounts" on the Web at http://www.microsoft.com/reskit).

- "Transitioning Service Accounts — Windows NT 4.0 Source" (DSSRENT_6.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Transitioning Service Accounts — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

- "Migrating Global Groups — Windows NT 4.0 Source" (DSSRENT_7.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Global Groups — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

- "Migrating User Accounts — Windows NT 4.0 Source" (DSSRENT_8.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating User Accounts — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

- "Migrating Local Profiles — Windows NT 4.0 Source" (DSSRENT_9.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Local Profiles — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

- "Migrating Workstations or Member Servers —Windows NT 4.0 Source" (DSSRENT_10.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations or Member Servers — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

- "Remigrating Global Groups — Windows NT 4.0 Source" (DSSRENT_11.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Remigrating Global Groups — Windows NT 4.0 Source" on the Web at http://www.microsoft.com/reskit).

# Restructuring Active Directory Domains Between Forests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Active Directory® directory service domains between forests involves relocating objects from source domains in one forest to target domains in another forest. If you are migrating a pilot domain into your production environment, merging with another organization and consolidating the two IT infrastructures, or consolidating resource and account domains that you upgraded in-place from a Microsoft® Windows NT® 4.0 operating system environment, then restructuring Active Directory domains between forests enables you to reduce the complexity of your organization and the associated administrative costs. However, if your Active Directory environment consists of multiple forests, you might also move objects between forests on a regular basis.

**In This Chapter**

Overview of Restructuring Active Directory Domains Between Forests

Active Directory Merger, Acquisition, and Divestiture: Restructuring Limitations

Planning to Restructure Active Directory Domains Between Forests

Preparing the Source and Target Domains

Migrating Accounts

Migrating Resources

Completing the Migration

Additional Resources for Retructuring Active Directory Domains Between Forests

**Related Information**

- For more information about the logical structure of Active Directory and about designing a Domain Name System (DNS) for Active Directory, see "Designing the Active Directory Logical Structure" in this book.
- For more information about designing the Active Directory site topology, see "Designing the Site Topology" in this book.

# Overview of Restructuring Active Directory Domains Between Forests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Microsoft® Windows® 2000, Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems domains between forests enables you to reduce the number of domains in your organization and therefore the administrative complexity and associated overhead costs of your Active Directory environment. Restructuring domains involves copying accounts and resources from a Windows 2000 or Windows Server 2003 Active Directory source domain to a target domain in a different Active Directory forest. The target domain must be a Windows 2000 native mode or Windows Server 2003 functional level domain.

If you completed an in-place upgrade of more than one domain from Windows NT 4.0 to multiple Windows 2000 or Windows Server 2003 forests, you might need to restructure your domains between forests to consolidate the objects because the Windows NT 4.0 Security Accounts Manager (SAM) size limits that restrict the number of objects in a domain do not apply to an Active Directory environment. If your organization has recently merged with another organization or IT infrastructure, then restructuring enables you to consolidate accounts and resources between the two infrastructures.

**Note**

- For a list of the job aids that are available to assist in you restructuring Active Directory domains between forests, see "Additional Resources for Retructuring Active Directory Domains Between Forests" later in this chapter.

# Process for Restructuring Active Directory Domains Between Forests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Active Directory domains between forests involves planning and preparing for the domain restructure for your organization and successfully migrating accounts and resources to an Active Directory domain in another forest. Figure 11.1 shows the process for restructuring Active Directory domains between forests.

**Figure 11.1 Restructuring Active Directory Domains Between Forests**



# Background Information for Restructuring Active Directory Domains Between Forests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The migration of domain objects between forests is a nondestructive process. The source and target domain environments exist simultaneously during the migration; therefore, you have the option to roll back to the source environment if necessary. The Active Directory Migration Tool (ADMT) version 2 enables you to migrate accounts and resources between domains while preserving user and object permissions. During the interforest restructure process, users have continuous access to required resources, and users, groups, and resources can be moved independently of each other.

Before you begin to restructure Active Directory domains between forests, you must be familiar with the account and resource migration process, Windows Server 2003 functional levels, and ADMT.

#### Terms and Definitions

The following terms apply to the Active Directory domain restructure process.

**Migration** The process of moving or copying an object from a source domain to a target domain, while preserving or modifying characteristics of the object to make it accessible in the new domain.

**Domain Restructure** A migration process that involves changing the domain structure of a forest. A domain restructure can involve either consolidating or adding domains, and can take place between forests or within a forest.

**Migration objects** Domain objects that are moved from the source domain to the target domain during the migration process. Migration objects can be user accounts, service accounts, groups, or computers.

**Source domain** The domain from which objects are moved during a migration. When restructuring Active Directory domains between forests, the source domain is an Active Directory domain in a different forest from the target domain.

**Target domain** The domain to which objects are moved during a migration.

**Built-in accounts** Default security groups that have common sets of rights and permissions. You can use built-in accounts to grant permissions to any accounts or groups that you designate as members of these groups. Built-in account security identifiers (SIDs) are identical in every domain and therefore built-in accounts cannot be migration objects.

#### Account Migration Process

Restructuring accounts between Active Directory forests involves the copying of users, groups, and local profiles from the source domain to the target domain, while preserving the access rights and attributes of those objects.

When user accounts are migrated between Active Directory domains in different forests, the original account remains in place in the source domain and a new account is created in the target domain. Because the SID of a security principal (user or group) always contains an identifier for the domain in which the security principal is located, a new SID is created for the user in the target domain. Because ADMT can migrate the SID of the original security principal to the security principal in the target domain, you do not need to perform additional tasks to ensure resource access, unless you are using SID filtering between the forests.

If you are using Microsoft® Exchange Server version 5.5, use the ADMT Exchange Server Migration Wizard to translate security on the mailboxes for migrated users. If you are using Exchange 2000 servers, ADMT does not provide tools for mailbox migration. In this case, plan to migrate mailboxes first by using the Exchange 2000 mailbox migration tool and then migrate user accounts.

If you are using Group Policy to manage folder redirection or software distribution, you need to ensure that these policies continue to apply when you migrate user accounts to a new forest. Also, if you are using a Group Policy object (GPO) to grant or deny remote access in the source domain and not the target domain, then ADMT cannot determine which remote access to assign to the user.

If you are using Group Policy to manage folder redirection, then Offline Files does not work after the user account is migrated to a new forest. Offline Files stores the SID of the user as owner; the SID changes when the user account is migrated. To restore ownership of Offline Files, use the ADMT Security Translation Wizard to replace the permissions on the files and folders on the client computer containing the offline files cache.

To ensure that users continue to have access to Offline Files after you migrate user accounts to the target domain, you can do

the following:

1. Translate security on client computers to update the Offline Files.

2. If the SID history of the user account was not migrated to the target domain, translate security on the server that hosts redirected folders.

If you are using folder redirection, one of the following occurs:

- If the folder redirection path is different in the new environment, then users can access the folder if the SID history of the user account was migrated to the target domain. The folder redirection extension copies the files from the original location in the source domain to the new location in the target domain. SID history enables the user account to access the source folders.

- If the folder redirection path is the same in the new environment, then users cannot access the redirected folder because folder redirection will check ownership of the redirected folder and will fail. You must then translate security on the redirected folder on the server.

Because some software installation (.msi) packages require access to the original source for operations such as repair and remove, if you are using Group Policy to manage software installation, then you need to translate security on the software distribution point after you migrate users to ensure that software installation continues to function properly in the target domain.

**Resource Migration Process**
Active Directory domains include three types of resources:

- Workstation accounts

- Member server accounts

- Resources on member servers

The migration of workstations and member servers is a straightforward process. The local groups that you create to assign permissions to users are located in the local SAM database and are moved when you move the server. You do not have to reconfigure access control lists (ACLs) to enable users to access resources after the migration.

In Active Directory, domain controllers can be migrated between domains. To do this, you must remove Active Directory from the domain controller, migrate it as a member server to the target domain, and then reinstall Active Directory. If you have deployed any domain controllers in the target domain that are running Windows Server 2003, then you can only reinstall Active Directory on the migrated server if the server is running Windows Server 2003 or if the forest and domain are operating at the Windows 2000 functional level.

For more information about functional levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.

**Functional Levels**
*Functional levels* in Windows Server 2003 identify the level of functionality that the domain controllers running different Windows operating system versions of Active Directory can use, such as schema compatibility and replication functionality. The functional level of a domain or forest defines the set of Windows operating systems that can run on the domain controllers in that domain or forest. The functional level of a domain or forest also defines the additional Windows Server 2003 Active Directory features that are available in that domain or forest.

All target domains to which you migrate objects when restructuring Windows Server 2003 domains between forests must be operating at the Windows 2000 native mode or Windows Server 2003 functional level.

**Active Directory Migration Tool**
ADMT version 2 enables you to copy objects between Active Directory forests. ADMT includes wizards that automate migration tasks such as copying users, groups, service accounts, and computers; migrating trusts; and performing security translation.

When you use the ADMT tool to perform an interforest restructure, ADMT copies the accounts that are migrated, so that when the accounts are created in the target domain, they continue to exist in the source domain. The primary SIDs for the accounts

can be migrated to the SID history in the target domain.

ADMT is available on the Windows Server 2003 operating system CD, in the Admigration.msi file in the \i386\admt directory.

You can perform ADMT tasks by using the ADMT console, by using a command-line procedure, or by using a script. When running ADMT from the command line, it is often more efficient to use an option file to specify command-line options. You can use the ADMT option file reference in Listing 11.1 to assist you in creating option files.

Listing 11.1 lists all common options that apply to several migration tasks. Each type of migration task has a section listing options specific to that task. The section name corresponds to the task name when executing the ADMT command-line tool. Items can be commented out with a semicolon. In the following list, the default values are commented out.

When running ADMT at the command line, you do not need to include an option in your command if you want to accept the default value. In this chapter, however, tables are provided for reference. The tables list the command-line equivalent of each option that is shown in the corresponding ADMT console procedure, including those where you are accepting the default value.

**Listing 11.1 ADMT Option File Reference**

```
 [Migration]
;TestMigration=No
;IntraForest=No
SourceDomain=" source_domain_name"
;SourceOu=" source_ou_path"
TargetDomain=" target_domain_name"
TargetOu=" target_ou_path"
;RenameOption=Dont
;RenamePrefixOrSuffix=""
;PasswordOption=Complex
;PasswordServer=""
;PasswordFile=""
;ConflictOptions=Ignore
;ConflictPrefixOrSuffix=""
;UserPropertiesToExclude=""
;InetOrgPersonPropertiesToExclude=""
;GroupPropertiesToExclude=""
;ComputerPropertiesToExclude=""

[User]
;DisableOption=EnableTarget
;SourceExpiration=None
MigrateSIDs=Yes
;TranslateRoamingProfile=No
;UpdateUserRights=No
;MigrateGroups=No
;UpdatePreviouslyMigratedObjects=No
;FixGroupMembership=Yes
;MigrateServiceAccounts=No
;UpdateGroupRights=No

[Group]
MigrateSIDs=Yes
;UpdatePreviouslyMigratedObjects=No
;FixGroupMembership=Yes
;UpdateGroupRights=No
;MigrateMembers=No
;DisableOption=EnableTarget
;SourceExpiration=None
;TranslateRoamingProfile=No
;MigrateServiceAccounts=No

[Security]
TranslationOption=Add
;TranslateFilesAndFolders=No
;TranslateLocalGroups=No
;TranslatePrinters=No
;TranslateRegistry=No
;TranslateShares=No
;TranslateUserProfiles=No
;TranslateUserRights=No
;SidMappingFile=" SidMappingFile.txt"
```

For a file that contains the ADMT options, see "ADMT Option File Reference" (DSSRENT_3.txt) on the *Microsoft® Window® Server 2003 Deployment Kit* companion CD (or see "ADMT Option File Reference" on the Web at http://www.microsoft.com/reskit).

When running ADMT at the command line, it is often more efficient to use an include file to list the users, groups, or computers that will be migrated. For example, to migrate a small number of computers, you might type each computer name at the command line, using the /N option as follows:

**ADMT COMPUTER /N** " *computer_name1* " " *computer_name2* " **/SD:**" *source_domain* " **/TD:**" *target_domain* " **/TO:**" *target_OU* "

*Computer_name1* and *computer_name2* are the names of computers in the source domain that you are migrating in this batch.

If you have a large number of users, groups, or computers to migrate, you can list them in an include file. For example, to create an include file for a batch of computers, create a plain text file and list the computer names, each on a separate line. Then specify the include file name with the /F option, as follows:

**ADMT COMPUTER /F** " *includefile_name* " **/SD:"source_domain" /TD:"target_domain" /TO:"target_OU"**

ADMT supports the use of American National Standards Institute (ANSI) or standard Unicode files for option files and include files.

To specify names of users, groups, or computers, use one of the following conventions:

- The SAM account name. Note that to specify a computer name in this format, you must append a $ to the computer name. For example, to specify a computer with the name Workstation01, use Workstation01$.
- The Windows NT account name. This includes the domain name and the SAM account name. For example, the Windows NT account name of a computer named Workstation01 that is in the Asia domain is Asia\ Workstation01$.
- The relative distinguished name (also known as RDN). For example, cn= Workstation01.
- The canonical name. This can be specified as *DNS domain name/ou_path/object_name*, or *ou_path/object_name*; for example, Asia.trccorp.treyresearch.net/Computers/ Workstation01 or Computers/Workstation01.

For a complete list of all available ADMT options and syntax, run ADMT and select **Help** from the menu.

The sample scripts provided in this chapter reference the symbolic constants that are defined in the AdmtConstants.vbs file. Listing 11.2 shows the ADMT constants VBScript file.

**Listing 11.2 ADMT Scripting Constants**

```
Option Explicit

'-----------------------------------------------------------------------------
' ADMT Scripting Constants
'-----------------------------------------------------------------------------

' RenameOption constants

Const admtDoNotRename      = 0
Const admtRenameWithPrefix = 1
Const admtRenameWithSuffix = 2

' PasswordOption constants

Const admtPasswordFromName = 0
Const admtComplexPassword  = 1
Const admtCopyPassword     = 2

' ConflictOptions constants

Const admtIgnoreConflicting           = &H0000
Const admtReplaceConflicting          = &H0001
Const admtRenameConflictingWithPrefix = &H0002
Const admtRenameConflictingWithSuffix = &H0003
```

```
Const admtRemoveExistingUserRights    = &H0010
Const admtRemoveExistingMembers       = &H0020
Const admtMoveReplacedAccounts        = &H0040

' DisableOption constants

Const admtEnableTarget       = 0
Const admtDisableSource      = 1
Const admtDisableTarget      = 2
Const admtTargetSameAsSource = 4

' SourceExpiration constant

Const admtNoExpiration = -1

' Translation Option

Const admtTranslateReplace = 0
Const admtTranslateAdd     = 1
Const admtTranslateRemove  = 2

' Report Type

Const admtReportMigratedAccounts  = 0
Const admtReportMigratedComputers = 1
Const admtReportExpiredComputers  = 2
Const admtReportAccountReferences = 3
Const admtReportNameConflicts     = 4

' Option constants

Const admtNone     = 0
Const admtData     = 1
Const admtFile     = 2
Const admtDomain   = 3
Const admtRecurse           = &H0100
Const admtFlattenHierarchy  = &H0000
Const admtMaintainHierarchy = &H0200
```

For a script file that contains the ADMT scripting constants, see "ADMT Scripting Constants" (AdmtConstants.vbs) on the *Windows Server 2003 Deployment Kit* companion CD (or see "ADMT Scripting Constants" on the Web at http://www.microsoft.com/reskit).

# Active Directory Merger, Acquisition, and Divestiture: Restructuring Limitations

Updated: May 3, 2010

Applies To: Windows SBS 2003, Windows SBS 2008, Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 Foundation, Windows Server 2008 R2

When a company acquires another company, business unit, or product line, the purchasing company may also want to acquire corresponding IT assets from the seller. Specifically, the buyer may want to acquire some or all of the domain controllers that host the user accounts, computer accounts, or security groups that correspond to the business assets that are to be purchased. The only supported methods for the buyer to acquire the IT assets that are stored in the seller's Active Directory forest are as follows:

1. Acquire the only instance of the forest, including all domain controllers and directory data in the seller's entire forest.

2. Migrate the needed directory data from the seller's forest or domains to one or more of the buyer's domains. The target for such a migration may be an entirely new forest or one or more existing domains that are already deployed in the buyer's forest. We recommend that you migrate the directory data without security identifier (SID) history. If you migrate the directory data with SID history, information about the seller's forest will be retained in the new forest of the buyer. For more information about migrating directory data without SID history, see Migrating Accounts Without Using SID History (http://go.microsoft.com/fwlink/?LinkId=113694).

This support limitation exists because:

- Each domain in an Active Directory forest is assigned a unique identity during the creation of the forest. Copying domain controllers from an original domain to a cloned domain compromises the security of both the domains and the forest. Threats to the original domain and the cloned domain include the following

    - Sharing of passwords that can be used to gain access to resources

    - Insight regarding privileged user accounts and groups

    - Mapping of IP addresses to computer names

    - Additions, deletions, and modifications of directory information if domain controllers in a cloned domain ever establish network connectivity with domain controllers from the original domain

- Cloned domains share a common security identity; therefore, trust relationships cannot be established between them, even if one or both of the domains have been renamed.

# Planning to Restructure Active Directory Domains Between Forests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Completing the necessary planning tasks before you begin your migration enables you to ensure that users can continue to log on to the network and access resources during the migration. Planning includes defining the accounts and resources to be migrated, creating a test plan, and creating a rollback plan for use in the event that the migration fails. It also includes planning for user profile migration, establishing administrative procedures, and creating an end-user communication plan. To prepare for the restructuring process, the Active Directory deployment team must be sure to obtain the necessary design information from the Active Directory design team.

Figure 11.2 shows the steps involved in planning to restructure Windows Server 2003 domains between forests.

**Figure 11.2 Planning to Restructure Active Directory Domains Between Forests**



# Determining Your Account Migration Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ADMT enables you to use SID history to maintain resource permissions when you migrate accounts. However, if SID filtering is enabled between your source and target domains, and you do not trust the administrators in the source domain, then you cannot disable SID filtering and use SID history to enable access to resources in the source domain. In this case, you must use a different migration process.

You can choose one of the following three methods to migrate accounts between forests while maintaining user rights to access resources in the source domain:

- Migrate user accounts without using SID history for resource access, but translate security for all resources before the migration process to ensure resource access. For more information about migrating accounts without using SID history, see "Migrating Accounts Without Using SID History" later in this chapter.

- Migrate user accounts while using SID history for resource access. With this method, you remove SID filtering on the trusts between the domains to enable users to access resources in the source domain by means of their SID history credentials.

    - If you have a cross-forest trust in place, you remove SID filtering on the cross-forest trust. (You can also override the cross-forest trust by creating an external trust so that the domain that holds the resources trusts the target domain, and then removing SID filtering on the external trust.)

    - If you do not have a cross-forest trust in place, you establish external trusts between the source and target domains. You then need to remove SID filtering on the external trusts if the domain controller used to create the trust is running Windows Server 2003 or Windows 2000 Service Pack 4 (SP4) or later.

  For more information about this process, see "Migrating Accounts While Using SID History" later in this chapter.

- Migrate all users, groups, and resources to the target domain in one step. For more information about this process, see "Migrating Accounts While Using SID History" later in this chapter.

To determine which account migration process is best for your organization, you must first determine if you can disable SID filtering and migrate accounts while using SID history for resource access. You can safely do this if the administrators of the source domain fully trust the administrators of the target domain. You might choose to disable SID filtering if one of the following conditions applies:

- The administrators of the trusting domain are the administrators for the trusted domain.

- The administrators of the trusting domain trust the administrators of the trusted domain and are confident that they have secured the domain appropriately.

If you disable SID filtering, you remove the security boundary between forests, which otherwise provides data and service isolation between the forests. For example, an administrator in the target domain who has service administrator rights or an individual who has physical access to a domain controller can modify the SID history of an account to include the SID of a domain administrator in the source domain. When the user account for which the SID history has been modified logs on to the target domain, it presents valid domain administrator credentials for and can obtain access to resources in the source domain.

For this reason, if you do not trust the administrators in the target domain or do not believe that the domain controllers in the target domain are physically secure, enable SID filtering between your source and target domains, and migrate user accounts without using SID history for resource access.

Figure 11.3 shows the decision process involved in determining which migration process is appropriate for your organization.

**Figure 11.3 Determining Your User Account Migration Process**


Do you fully trust the source domain?
No
Can you migrate users and resources at the same time?
No
Yes
Yes
Do you have a cross-forest trust in place?
No
Establish trusts between source and target domain
Migrate resources at the same time that you migrate users
Yes
Remove SID filtering on cross-forest trust
Remove SID filtering on explicit trust if between Windows Server 2003 or Windows 2000 SP4 or later domain controllers
Proceed with migration without using SID history for resource access
Proceed with migration while using SID history for resource access

# Using SID History to Preserve Resource Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The best practice for granting access to resources is to use global groups to group users, and local groups to protect resources. Place global groups into a local group to grant the members of the global group access to the resource. A global group can only contain members from its own domain. When a user is migrated between domains, any global groups to which the user belongs must also be migrated in order for users to continue to have access to resources protected by discretionary access control lists (DACLs) referring to global groups. After migrating an account and maintaining the SID history of the source domain account, when a user logs on to the target domain, both the new SID and the original SID from the SID history attribute are added to the access token of the user and determine the local group memberships of the user. The SIDs of the groups in which the user is a member are then added to the access token, together with the SID history of those groups.

Resources within the source and target domains resolve their ACLs to SIDs and then check for matches between their ACLs and the access token when granting or denying access. If the SID or the SID history matches, access to the resource is granted or denied, according to the access specified in the ACL. If the resource is in the source domain and you have not run security translation, it uses the SID history of the user account to grant access.

You can also preserve the original SID for global groups and universal groups in the SID history of the global group or universal group in the target domain. Because local group memberships are based on SIDs, when you migrate the SID to the SID history of the global group or universal group in the target domain, the local group memberships of the global group or universal group are preserved automatically.

SID history is used for roaming user profile access, certification authority access, and software installation access, as well as resource access. If you are not using SID history for resource access, you still need to migrate SID history to facilitate access to those items.

# Using SID Filtering When Migrating User Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

SID filtering does not allow for the use of SIDs from outside the forest to enable access to any resource within the forest. You can enable the SID of a user in a different forest to access a resource within a forest that has SID filtering enabled by translating security on the resource to include the user SID in the permission list. Because SID filtering does not apply to authentication within a domain, it is also possible to allow access to resources by means of SID history if the resource and the account are in the same domain.

To allow users or groups to access a resource by using SID history, the forest in which the resource is located must trust the forest in which the account is located. SID filtering is applied by default when a forest trust is established between two forest root domains. Also, SID filtering is enabled by default when external trusts are established between domain controllers running Windows Server 2003 or Windows 2000 SP4 or later. This prevents potential security attacks by an administrator in a different forest.

For more information about SID history–based attacks and SID filtering, see Configuring SID Filtering Settings at http://go.microsoft.com/fwlink/?LinkId=73446.

# Assigning Object Locations and Roles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assign your migration objects locations and roles in the target domain. It is useful to create object assignment tables in which to document the assignments for all of the objects that you are migrating. Create one table for account objects, such as users, groups, and service accounts, and one table for resource objects, such as workstations, profiles, and domain controllers. In your tables, list the source and target locations for all objects to be migrated.

Before you create your account object assignment table, determine whether the domain organizational unit (OU) structures for the source and target domains are the same. If the OU structures are not the same, you must identify the source and target OU in your object assignment tables.

For a worksheet to assist you in creating an account object assignment table, see "User and Group Object Assignment Table" (DSSREER_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "User and Group Object Assignment Table" on the Web at http://www.microsoft.com/reskit).

Figure 11.4 shows an example of an object assignment table for users and groups.

**Figure 11.4 Example of a User and Group Object Assignment Table**

| User and Group Object Assignment Table | | | |
|---|---|---|---|
| Prepared By | Active Directory Deployment Team | Date | 01/01/03 |
| Source Forest Name | Trccorp.treyresearch.net | | |
| Target Forest Name | Concorp.contoso.com | | |
| Source Domain Name | Asia.trccorp.treyresearch.net | | |
| Target Domain Name | Emea.concorp.contoso.com | | |

| Name | Type | Source Location | Target Location |
|---|---|---|---|
| Finance | Group | Asia\Groups OU | EMEA\Groups OU |
| Accounting | Group | Asia\Groups OU | EMEA\Groups OU |
| JBrown | User | Asia\Users OU | EMEA\Users OU |
| MNguyen | User | Asia\Users OU | EMEA\Users OU |
| !Scheduler | Service Account | Asia\Users OU | EMEA\Service Accounts OU |

Creating a resource object assignment table also involves identifying the source and target OU for each object and noting the physical location and role in the target domain. For a worksheet to assist you in creating a resource object assignment table, see "Resource Object Assignment Table" (DSSREER_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Resource Object Assignment Table" on the Web at http://www.microsoft.com/reskit).

Figure 11.5 shows an example of a resource object assignment table.

**Figure 11.5 Example of a Resource Object Assignment Table**

| Resource Assignment Table | | | |
|---|---|---|---|
| Prepared By | Active Directory Deployment Team | Date | 01/01/03 |

| | |
|---|---|
| Source Forest Name | Trccorp.treyresearch.net |
| Target Forest Name | Concorp.contoso.com |
| Source Domain Name | Asia.trccorp.treyresearch.net |
| Target Domain Name | Emea.concorp.contoso.com |

| Name | Type | Source OU | Target OU | Physical Location | Role in Target Domain |
|---|---|---|---|---|---|
| FS01 | Member server (file server) | Computers | Member Servers | Hong Kong SAR | Member Server |
| UserWrk01 | Computer account | Default | Wrk | Hong Kong SAR | Computer account |
| UserWrk02 | Computer account | Wrk | Wrk | Boston | Computer account |
| DC01 | Domain Controller | Domain Controllers | Domain Controllers | Boston | Domain controller |
| DC02 | Domain Controller | Domain Controllers | Member Servers | Hong Kong SAR | Member server |

# Developing a Test Plan for Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to develop a test plan to assist you in systematically testing each object after it is migrated to the new environment, and identifying and correcting any problems that might occur. Testing to verify that your migration is successful allows you to ensure that users who are migrated from the source to the target domain are able to log on, to access resources based on group membership, and to access resources based on user credentials. Testing also allows you to ensure that users are able to access the resources that you migrate.

Use the following process to test your account object and resource object migration:

1. Create a test user in the source domain. Include this test user with your migrations.

2. Join that user to the appropriate global groups to enable resource access.

3. Log on to the source domain as the test user and verify that you can access resources as appropriate.

4. After you translate the user profile, migrate the workstation of the user, and migrate the user account, log on to the target domain as the test user and verify that the user has retained all necessary access and functionality. For example, you might test to verify that:

   - The user can log on successfully.

   - The user has access to all appropriate resources, such as file and print shares; access to services such as messaging; and access to line of business applications. It is especially important to test access to internally developed applications that access database servers.

   - The user profile was successfully translated, and the user retains desktop settings, desktop appearance, shortcuts, and access to the My Documents folder. Also, verify that applications appear in and start from the **Start** menu.

After you migrate resources, log on as the test user in the target domain and verify that you can access resources as appropriate.

If any steps in the test process fail, identify the source of the problem and determine whether you can correct the problem before the object needs to be accessible in the target domain. If you cannot correct the problem before access to the object is required, roll back to your original configuration to ensure access to the user or resource object. For more information about creating a rollback plan, see "Creating a Rollback Plan" later in this chapter.

As part of your test plan, create a migration test matrix. Complete a test matrix for each step that you complete in the migration process. For example, if you migrate 10 batches of users, complete the test matrix 10 times, once for each batch that you migrate. If you migrate 10 member servers, complete the test matrix for each of the 10 servers.

For a worksheet to assist you in creating a test matrix, see "Migration Test Matrix" (DSSREER_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migration Test Matrix" on the Web at http://www.microsoft.com/reskit).

Figure 11.6 shows an example of a completed migration test matrix.

**Figure 11.6 Example of a Completed Migration Test Matrix**

# Migration Test Matrix

| Prepared By | Active Directory Deployment Team | Date | 01/01/03 |
|---|---|---|---|

| | |
|---|---|
| Source Forest Name | Trccorp.treyresearch.net |
| Target Forest Name | Concorp.contoso.com |
| Source Domain Name | Asia.trccorp.treyresearch.net |
| Target Domain Name | Emea.concorp.contoso.com |

| Test Name | Performed Before Migration | Result After Migration | Notes |
|---|---|---|---|
| Create test user for users in Sales department "batch1" | Yes | Success | |
| Join user to sales group | Yes | Success | |
| Log on to source domain and verify access to required resources | Yes | Success | |
| Log on to target domain | No | Success | |
| Verify access to file and print shares | No | Success | |
| Verify access to customer management application | No | Success | |
| Verify user desktop settings and access to My Documents folder | No | Success | |
| Verify that applications appear and start from the Start menu | No | Success | |

# Creating a Rollback Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Reduce the risk of disrupting end users in your organization by establishing a rollback plan. In general, it is possible to isolate and resolve any problems that occur during each phase of the migration. However, it is important to analyze potential risks and identify the levels of user impact and downtime that might necessitate rolling back the migration. You might be required to roll back your migration if any of the following occur:

- Users cannot log on to their accounts after migration.
- Users cannot access resources after migration.
- User migration is incomplete; for example, passwords did not migrate.
- User migration was successful, but user workstation migration or local profile translation failed.

If user impact or downtime reaches a level that you have defined as unacceptable in your organization, you can implement your rollback plan and continue to operate in your premigration environment. Because the source domain remains intact during an interforest restructure, you can restore the original environment by completing a few key steps.

To roll back to the premigration environment after migrating account objects:

1. Enable the user accounts in the source domain (if you disabled the accounts during the migration process).
2. Notify the users to log off of the target domain.
3. Notify the users to log on to the source domain.
4. Verify that users are able to access resources.
5. Verify that the logon scripts and user profiles for users work as configured in the source domain.

The rollback process for resource objects is similar to that for account objects. To roll back to the premigration environment after migrating resource objects:

1. Change the domain membership for the server or workstation to the source domain.
2. Restart the server or workstation.
3. Log on as a user and verify that you can access the resource.

**Note**

- If you need to modify objects such as member servers or domain controllers in order to migrate them to the target domain, back up all the data before making the modifications and performing the migration.

# Planning for a User Profile Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

User profiles store user data and information about the desktop settings of the user. User profiles can either be roaming or local. The migration process is different for local and for roaming profiles.

Profile translation is one type of security translation, and profiles are translated during the migration process. If you perform security translation in add mode, the SIDs in the target and the source domains both have access to the profile. Therefore, if you need to roll back to the source environment, the SID in the source domain can use the profile. If you perform security translation in replace mode, you must retranslate the profile by using a SID mapping file (undoing the security translation) to roll back to the source environment.

**Important**

- In the event that you need to roll back to your original configuration, notify users that profile changes made in the target domain are not reflected in the source domain.

Some organizations might choose not to migrate user profiles. Other organizations might choose to replace users' workstations during the user account migration process, and use a tool such as the User State Migration Tool (USMT) to migrate user data and settings to the users' new computers. Table 11.1 summarizes the migration requirements for user profiles.

Table 11.1 Migration Requirements for User Profiles

| Type | Description | Migration Requirements |
|---|---|---|
| Roaming profiles | User profiles are stored centrally on servers. Profiles are available to the user, regardless of the workstation in use. | Migrate local user profiles for a batch of users before migrating those users. Then select the profile migration option during the user account migration process to migrate the roaming profile. |
| Local profiles | User profiles are stored locally on the workstation. When a user logs on to another workstation, a unique local user profile is created. | Migrate as a separate step from the user account migration process. Migrate local user profiles for a batch of users before migrating those users. |
| Profiles not managed | Same as local profiles. | No migration requirements. Users lose their existing profiles when their user accounts are migrated. |
| Hardware refresh | User state information is stored locally on the workstation. | Migrate as a separate step from the user account migration. Migrate the profiles to the user's new computer by means of a tool such as USMT. |

# Creating Administrative Procedures

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must define how the objects that you are migrating are to be administered during the interforest restructure process. Establishing administrative procedures for migration objects enables you to preserve the objects both in the source and the target domains, and therefore enables you to fall back to the premigration environment in the event that the restructure process is not successful.

Plan for the administration and management of the following types of account migration objects:

- User accounts, including SIDs.
- Global group membership.
- User profiles.

**Administering User Accounts**

During the migration process, user accounts exist in both the source and the target domains. Administer changes to user accounts in the domain in which the user object is active. Continue to administer changes to group memberships in the source domain while the migration is taking place. Use the **Replace conflicting accounts** option in ADMT to remigrate user accounts as often as necessary during the migration process. This ensures that changes made to the account in the source domain are propagated to the account in the target domain. This operation merges the existing account and the new account so that administration of the object can continue in the source domain for the duration of the migration process.

The **Replace conflicting accounts** option applies the following guidelines when an account is migrated:

- If you change an attribute in the target domain and it is not used in the source domain, then it is not overwritten with the NULL value from the source domain.
- If you change an attribute in the target domain and it is used in the source domain, then it is overwritten with the value from the source domain.
- If the user has group memberships, the memberships are merged together from the source memberships and the target memberships.

If this is not the desired behavior, you can configure ADMT to exclude attributes from being migrated so that attributes in the target domain are retained.

For example, if you migrate a user and after you complete the migration you set attributes on the new user object in the target domain, such as a telephone number or office number, and you remigrate the user by using the **Replace conflicting accounts** option in ADMT, the new information is retained in the target domain. If you changed the group memberships for the user in the source domain, the changes are propagated to the target domain when you perform the remigration.

Some attributes are excluded from the migration. These include:

- Attributes that are always excluded by the system
- Attributes that are in the system exclusion list
- Attributes that are configured by the administrator to be excluded

**Attributes that are always excluded by the system**

Some attributes are always excluded from the migration by ADMT and cannot be configured to be migrated. This protects system-owned attributes. These attributes include:

- Object globally unique identifier (GUID)

- SIDs (although SIDs can be added to the SID history of the object in the target domain)
- LegacyExchangeDN

**System attribute exclusion list**
The first time that you run an ADMT user migration, ADMT generates a system attribute exclusion list, which it stores in its database. The system attribute exclusion list contains two attributes by default: **mail** and **proxyAddresses**. ADMT also reads the schema in the target domain, and adds any attributes to the list that are not part of the base schema. Attributes in this list are excluded from migration operations even if the attribute is not specified in the attribute exclusion list. An administrator can change the system attribute exclusion list only by using a script. This protects attributes that are important in order for server-based applications, such as Exchange, to work. If the target domain schema is further extended after ADMT has generated the list, administrators must manually add the new attributes to the list, unless they are certain that copying the values of these attributes from the source domain will not interfere with server-based applications.

**Attribute exclusion list**
Administrators can define a list of attributes that are excluded from each migration. This is called an attribute exclusion list. By default, when using the ADMT console, state information for attributes that are configured to be excluded is stored in the UI and included in the exclusion list for the next migration. Scripting and command-line attributes do not have state information and therefore are not stored in the UI. These attributes must be added to the attribute exclusion list for each migration operation, either by means of the attribute name or by means of an option file.

**Administering Global Groups**
Continue to administer the groups in the source domain during the migration process. Remigrate groups as often as necessary by using the **Replace conflicting accounts** option in ADMT to ensure that changes made to group memberships in the source domain are propagated to the groups in the target domain.

# Creating the End-User Communication Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Develop a plan to inform all affected users about the upcoming account migration, to ensure that they understand their responsibilities, the impact of the migration, and who to contact for help and support.

Before you begin the user migration process, send a notice to all users who are scheduled to be migrated. Because you typically migrate users in batches of approximately one hundred users at a time, it is also helpful to send a final notice to the users in each batch two to three days before their batch is scheduled. If your organization maintains an intranet, publish the account migration schedule and the information contained in the user mail on an easily accessible Web page.

Include the following information in your end-user communication.

**General information**
Alert users to the fact that their user accounts are scheduled to be migrated to a new domain. Point users to a Web page or internal resource where they can find additional information, and view a migration schedule.

Inform users of their new domain name. Be sure to let them know that their account passwords will not change. Let them know that the original domain account will be disabled immediately following the migration, and the disabled account will be deleted after a specified period of time.

**Impact**
Make sure that users understand that when their account is migrated, it is possible that they could be unable to access some resources, such as Web sites, shared folders, or resources that are not widely used by individuals in their group or division.

Provide information to users about who to contact for assistance in regaining access to required resources.

**Logon status during migration**
Make sure that users understand that during the migration process, they will be unable to log on to the domain or access e-mail or other resources. Be sure to specify the period of time for which they will be unable to log on.

**Premigration steps**
Alert users to any steps that they need to complete before the migration process begins. For example, they must decrypt files encrypted by means of Encrypted File System (EFS). Failure to decrypt encrypted files will result in loss of access to encrypted files following the migration.

Users must also ensure that their computers are connected to the network when their account is scheduled to be migrated.

**Expected changes**
Describe other changes that users can expect to experience following the migration, such as changes in use of smart cards, secure mail, or instant messaging, if applicable.

**Scheduling and support information**
Provide information about where users can go to find more information; for example, an internal Web site where you post information about the migration. Also, provide information about who to contact if a user has a conflict with the date scheduled for the migration.

# Preparing the Source and Target Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin to migrate your accounts from the source to the target domains, you need to prepare the source and target domains for the migration. Figure 11.7 shows the tasks required to prepare your source and target domains for the interforest domain restructure process.

**Figure 11.7 Preparing the Source and Target Domains**



# Installing 128-Bit High Encryption Software

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The computer on which ADMT is installed requires 128-bit high encryption. This encryption is standard on computers running Windows Server 2003 and Windows 2000 Server SP3 or later. If you plan to install ADMT on a computer that is running Windows 2000 SP2 or earlier, you must install the 128-bit high encryption pack.

You can download the encryption pack from the Windows 2000 High Encryption Pack link on the Web Resources page at http://www.microsoft.com/windows/webresources.

# Establishing Required Trusts for Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can migrate accounts and resources from a source domain to a target domain in a different Active Directory forest, you must ensure that the appropriate trusts exist between the forests. Trust relationships between the forests that you are restructuring enables ADMT to migrate users and service accounts and translate local user profiles from the source to the target domains. In addition, trust relationships enable users in the source domains to access resources in the target domains, and users in the target domains to access resources in the source domains that have not yet been migrated.

To migrate users and global groups, you must establish a one-way trust between the source domain and the target domain, so that the source domain trusts the target domain.

To migrate resources or translate local profiles, you must do one of the following:

- Create a one-way trust between the source domain and the target domain.
- Create a two-way trust between source and target domains.

# Establishing Migration Accounts for Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To migrate accounts and resources between forests, you must establish a migration account or accounts and assign the appropriate credentials to those accounts. ADMT uses the migration accounts to migrate the objects that you identify. Because ADMT requires only a limited set of credentials, creating separate migration accounts enables you to simplify administration. If the migration tasks for your organization are distributed across more than one group, it is helpful to create a migration account for each group involved in performing the migration

To simplify administration, create a single account in the source domain and a single account in the target domain for all objects, with the required credentials to modify the objects, such as users, global groups, and local profiles, to be migrated by that account. For example, a migration account that you use to migrate user accounts along with SID history, global groups along with SID history, computers, and user profiles has local administrator or domain administrator credentials in the source domain, and delegated permission on the user, group, and computer OUs in the target domain, with the extended right to migrate SID history on the user OU. The user needs to be a local administrator on the computer in the target domain on which ADMT is installed. A migration account that you use to migrate workstations and domain controllers must have local administrator or source domain administrator credentials on the workstations or the account must have source domain administrator credentials on the domain controller, or both.

In the target domain, it is necessary to use an account that has delegated permissions on the computer OU and the user OU. You might want to use a separate account for the migration of workstations if this migration process is delegated to administrators that are in the same location as the workstations.

Table 11.2 lists the credentials that are required in the source and target domains for different migration objects.

**Table 11.2 Migration Account Credentials**

| Migration Object | Credentials Necessary in Source Domain | Credentials Necessary in Target Domain |
|---|---|---|
| User/Group without SID history | Delegated **Read all user information** permission on the user OU or group OU. | Delegated **Create user objects** permission on the user OU or group OU and local administrator on the computer on which ADMT is installed. |
| User/Group with SID history | Local administrator or domain administrator | Delegated permission on the user OU or the group OU, extended permission to migrate SID history, and local administrator on the computer on which ADMT is installed. |
| Computer | Domain administrator or administrator in the source domain and on each computer | Delegated permission on the computer OU and local administrator on the computer on which ADMT is installed. |
| Profile | Local administrator or domain administrator | Delegated permission on the user OU and local administrator on the computer on which ADMT is installed. |

The following procedures provide examples for creating groups or accounts to migrate accounts and resources. Procedures differ according to whether a one-way trust or a two-way trust exists The procedure for creating migration groups when a one-way trust exists is more complex than the procedure for creating migration accounts when a two-way trust exists, because you must add the migration group to the local Administrators group on local workstations. The sample procedure for creating migration groups when a one-way trust exists involves the creation of separate groups for migrating accounts and resources; however, you can combine *acct_migrators* and *res_migrators* into one group, if you do not need to separate them to delegate different sets of permissions.

**To create an account migration group when a one-way trust exists in which the source domain trusts the target domain**

1. In the target domain, create a global group, *acct_migrators*.

2. In the target domain, add the *acct_migrators* group to the Domain Admins group, or delegate administration of OUs that are targets for account migration to this group.

3. If you are migrating SID history, and you did not place the *acct_migrators* group in the Domain Admins group, grant the *acct_migrators* group the extended permission **Migrate SID History** on the target domain object. To do this:

    a. Start Active Directory Users and Computers, right-click the domain object, and then click **Properties**.

    b. Click the **Security** tab, click **Add**, and select ***acct_migrators***.

    c. In the **Permissions for acct_migrators** box, click **Allow** for the Migrate SID History permission.

4. In the source domain, add the *acct_migrators* group to the Administrators group.

5. On each computer on which you plan to translate local profiles, add the *acct_migrators* group to the local Administrators group.

**To create the resource migration group when a one-way trust exists in which the source domain trusts the target domain**

1. In the target domain, create a global group, *res_migrators*.

2. In the target domain, add the *res_migrators* group to the Domain Admins group, or delegate administration of OUs that are targets for resource migration to this group.

3. In the source domain, add the *res_migrators* group to the Administrators group.

4. On each computer that you plan to migrate or on which you plan to perform security translation, add the *res_migrators* group to the local Administrators group.

**To create a resource migration account when a two-way trust exists between the source and target domains**

1. In the source domain, create an account, *res_migrator*.

2. In the source domain, add the *res_migrator* account to the Domain Admins group. (The Domain Admins group is a member of the local Administrators group on every computer in the domain by default; therefore, you do not need to add it to the local Administrators group on every computer.)

3. In the target domain, delegate permissions on OUs that are targets for resource migration to the *res_migrator* account.

# Configuring the Source and Target Domains for SID History Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To migrate SID history, three conditions must be met:

- A local group used to audit SID history operations exists in the source domain.
- TCP/IP client support must be enabled on the source domain primary domain controller (PDC) emulator.
- Audit policies must be enabled.

You can configure these items manually before beginning the migration, or you can allow ADMT to configure them automatically the first time it runs. If you want to configure them manually, use the following procedures.

**To create a local group in the source domain to support auditing**

- In the source domain, create a local group *source_domain*$$$, where domain is the NetBIOS name of your source domain, for example, boston$$$. Do not add members to this group; if you do, SID history migration will fail.

  **Caution**

  - The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

**To enable TCP/IP client support on the source domain PDC emulator**

1. On the domain controller in the source domain that holds the PDC emulator role, use the registry editor to navigate to the following registry subkey:

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa**

2. Modify the registry entry **TcpipClientSupport**, of data type REG_DWORD, by setting the value to 1.
3. Restart the computer.

**To enable auditing on the Windows Server 2003 domains**

1. Log on as an administrator to any domain controller in the target domain.
2. Open **Active Directory Users and Computers**, expand the domain, and double-click the **Domain Controllers** OU.
3. On the **Group Policy** tab, click **Default Domain Controllers Policy**.
4. Expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
5. Double-click **Audit account management**, and then click both **Success** and **Failure**.
6. Repeat steps 1-5 on the source domain.

**Note**

- To assist you in using the event log to troubleshoot errors in the migration process, be sure to synchronize the time on all computers that are involved in the migration.

# Configuring the Target Domain OU Structure for Administration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory design team creates the OU structure for the target domain. They also define the groups responsible for the administration of each OU and the membership of each group. You can use that information to configure the target domain for administration. This involves the following steps:

1. Log on as an administrator to any domain controller in the target domain.
2. Open Active Directory Users and Computers and create the OU structure specified by your design team.
3. Create administrative groups and assign users.
4. Delegate the administration of the OU structure to groups as defined in your design.

# Installing ADMT in the Target Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Install and run ADMT on a domain controller in the target domain.

**To install ADMT in the target domain**

1. Log on to a computer in the target domain by using your ADMT migration account.
2. Insert the Windows 2000 or Windows Server 2003 operating system CD in the drive and navigate to \i386\admt.
3. Run admigration.msi.

# Enabling Migration of Passwords

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ADMT version 2 supports password migration. When you migrate passwords by using ADMT, you must use a password export server (PES) to host password migration support dynamic-link libraries (DLLs). The PES can be any domain controller in the source domain that supports 128-bit encryption.

Passwords are copied from the source domain to the target domain in hash form; therefore, it is not possible for a password filter to verify that the complexity or length of the passwords meet the requirements of the organization. The target domain controller used to set the password can, however, verify password history by comparing the hash of the password against previous hashes.

In the target domain, ensure that the built-in group Pre-Windows 2000 Compatible Access contains the Everyone and Anonymous Logon groups. The groups will not be present if you selected **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems** when you installed Active Directory in the target domain.

If the groups are not present, ADMT logs an Access Denied error. You must then manually add them to enable support for password migration. To do this, type the following at the command line on a target domain controller:

```
        NET LOCALGROUP "Pre-Windows 2000 Compatible Access" Everyone /ADD
NET LOCALGROUP "Pre-Windows 2000 Compatible Access" "Anonymous Logon" /ADD
```

After this update to the Pre-Windows 2000 Compatible Access group replicates, you must restart the Server service on all domain controllers in the target domain.

Use the following procedures to enable support for password migration.

**To generate an encryption key**

1. Log on to the domain controller in the target domain on which you installed ADMT by using your ADMT migration account.
2. Select the location to which to save the file. The location can be any local drive, including the floppy drive.
3. Open a command window, change to the drive on which ADMT is installed, and at the command line, type:

```
ADMT KEY
Source_Domain Folder
[
password
] [*]
```

The *Source_Domain* can be specified as either the DNS or NetBIOS name. A password, which provides key encryption, is optional. To protect the shared key, type either the password or an asterisk (*) at the command line. The asterisk prompts for a password that is not displayed on the screen.

**To generate an encryption key**

1. Log on to the domain controller in the target domain on which you installed ADMT by using your ADMT migration account.
2. Select the location to which to save the file. The location can be any local drive, including the floppy drive.
3. Open a command window, change to the drive on which ADMT is installed, and at the command line, type:

```
ADMT KEY
Source_Domain Folder
[
password
```

```
] [*]
```

The *Source_Domain* can be specified as either the DNS or NetBIOS name. A password, which provides key encryption, is optional. To protect the shared key, type either the password or an asterisk (*) at the command line. The asterisk prompts for a password that is not displayed on the screen.

| **Note** |
| --- |
| For security reasons, it is strongly recommended that you provide a password. |

**To enable password migration on the source domain**

1. On the PES in the source domain, insert the encryption key disk.
2. Insert the Windows Server 2003 operating system CD in the CD drive.
3. Navigate to the \i386\admt\pwdmig directory, and run pwdmig.exe. If you set a password during the key generation process on the domain controller in the target domain, the Key Password Required dialog box appears.
4. Enter the password and then complete the setup.
5. On the source domain PES, use the registry editor to navigate to the following registry subkey:

   ```
   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\LSA
   ```

6. Modify the registry entry **AllowPasswordExport**, of data type REG_DWORD, by setting the value to 1.

| **Caution** |
| --- |
| The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference. |

# Initializing ADMT by Running a Test Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Initialize ADMT by running a test migration of a global group with the option to migrate SID history selected. If you did not previously configure the source and target domains to migrate SID history, you will receive an error and a prompt for each item that has not yet been configured. When you accept each prompt, ADMT automatically completes the following tasks, which are required to enable SID history migration:

- Creates a local group, *source_domain*$$$, in the source domain, which is used to audit SID history operations. Do not add members to this group; if you do, SID history migration will fail.

- Enables TCP/IP client support on the source domain PDC emulator by setting to 1 the value of the registry entry **TcpipClientSupport**, which is locatedin the following subkey:

  HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa

  Setting **TcpipClientSupport** to 1 enables remote procedure calls (RPCs) over the TCP transport while preserving the security of the system.

- Enables audit policies in the source and target domains.

For more information about configuring the source and target domains to migrate SID history, see "Configuring the Source and Target Domains for SID History Migration" earlier in this chapter.

**To initialize ADMT by running a test migration of a global group**

1. In the Active Directory Migration Tool console, on the **Tools** menu, click **Group Account Migration Wizard**.

2. Complete the Group Account Migration Wizard by using the information provided in Table 11.3. Accept default settings when no information is specified.

**Table 11.3 Migrating Global Groups by Using the Group Account Migration Wizard**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Test the migration settings and migrate later?** |
| Domain Selection | In the **Sourcedomain** box, type the source domain name.<br><br>In the **Target domain** box, type the target domain name. |
| Group Selection | Click **Add**.<br><br>In the **Select Groups** dialog box, select a group. |
| Organizational Unit Selection | Click **Browse**.<br><br>In the **Browse for Container** dialog box, locate the **Groups** OU. |
| Group Options | Select the **Migrate Group SIDs to target domain** check box.<br><br>Click **Do not rename accounts**. |
| User Account | Type the name and password of the migration account that you created for migrating accounts.<br><br>In the **Domain** box, type the NetBIOS name of the target domain. |

| | |
|---|---|
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

3. When the wizard has finished running, click **View Log** and review the migration log for any errors.

4. Verify that the test migration configured ADMT properly by ensuring that:

   - A new local group *source_domain*$$$ exists in the source domain. This account supports ADMT internal auditing.

   - The registry entry **TcpipClientSupport** is created as data type REG_DWORD with a value of 1 in the following subkey on the source domain PDC emulator:

   HKEY_LOCAL_MACHINE\System\CCS\Control\Lsa

   - The audit policy for account management is enabled on the target domain.

# Identifying Service Accounts for Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Identify the member servers and domain controllers in the source domain that run applications in the context of a service account. A service account is a user account created as a means to provide a security context for applications. The service account is a standard user account that is granted permission to log on as a service. ADMT does not migrate services that run in the context of the Local System account because they are migrated when the computer is migrated; however, services that run in the context of a user account must be updated on the computer after you have completed the account migration process. ADMT also cannot migrate the Local Service or Network Service accounts because they are well-known accounts that always exist in Windows Server 2003.

The process of identifying, migrating, and updating services that run in the context of user accounts involves three steps. First, the administrator starts ADMT from the target Active Directory domain controller and runs the Service Account Migration Wizard. The Service Account Migration Wizard sends an agent to a computer that is specified and identifies all services on the computer that are running in the context of a user account. The wizard only identifies service accounts that are running in the context of a user account; it does not actually migrate the accounts. The next step, which can occur later in the migration process, is to migrate the accounts when other user accounts are migrated by means of the User Account Migration Wizard.

The Service Account Migration Wizard scans an administrator-defined list of servers for services that are configured to use a domain account to authenticate. The accounts are then flagged as service accounts in the ADMT database. The password is never migrated when a service account is migrated. Instead, ADMT uses a clear-text representation of the password to configure the services after the service account migration. An encrypted version of the password is then stored in the "password.txt" file in the ADMT installation folder.

An administrator of a workstation or server can install any service and configure the service to use any domain account. If the administrator cannot configure the service to authenticate with the correct password, then the service will be unable to start. After the service account is migrated, ADMT configures the service on the workstation or the server to use the new password and the service will now start under the user account.

It is important to include in the Service Account Migration Wizard only those servers that are managed by trusted administrators. Do not use the Service Account Migration Wizard to detect service accounts on computers that are not managed by trusted administrators, such as workstations.

Dispatch agents to all servers managed by trusted administrators in the domain to ensure that you do not overlook any service accounts. If you miss a service account that shares an account with a service that has already been migrated, it is not possible for ADMT to synchronize them. You must then manually change the password for the service account and then reset the service account password on each server that is running that service.

When the accounts that the Service Account Migration Wizard identifies in the ADMT database as running in the context of a user account are migrated to the target domain, ADMT grants each account the right to log on as a service. If the service account is assigned rights by means of its membership in a group, the Security Translation Wizard updates the account to assign those rights. For more information about running the Security Translation Wizard, see "Transitioning Service Accounts in Your Migration" later in this chapter.

**To identify service accounts by using the ADMT console**

1. On the domain controller on which ADMT is installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool console and then select **Service Account Migration Wizard**.
3. Complete the Service Account Migration Wizard by using the information in Table 11.4.

**Table 11.4 Using the Service Account Migration Wizard**

| Wizard Page | Action |
|---|---|
| Domain Selection | In the **Source domain** box, type or select the NetBIOS or DNS name of the source domain.<br>In the **Target domain** box, type or select the NetBIOS or DNS name of the target domain. |
| Update Information | Click **Yes, update the information**. |

| | |
|---|---|
| Service Account Selection | Click **Add**.<br>In the **Select Computers** list box, click the names of all servers that are capable of running services in the source domain and in all domains that trust the source domain.<br>Click **OK**, and then click **Next**. |
| Service Account Information | Select any user accounts that do not need to be marked as service accounts in the ADMT database, and then click **Skip/Include** to mark them as **Skip**. |

4. When the wizard has finished running on all computers, on the Server List page, click **View Log**, and review the migration log for any errors.

**To identify service accounts by using the ADMT command-line option**

1. On a domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT SERVICE /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
"
```

*Computer_name1* and*computer_name2* are the names of computers in the source domain that run service accounts.

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT SERVICE /N "
computer_name1
" "
computer_name2
" /O: "
option_file.txt
"
```

Table 11.5 lists the common parameters used for the identification of service accounts, along with the command-line parameter and option file equivalents.

**Table 11.5 Common Parameters Used for Service Account Identification**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |

3. Review the results that are displayed on the screen for any errors.

**To identify service accounts by using a script**

- Create a script that incorporates ADMT commands and options for identifying service accounts by using the sample

script shown in Listing 11.3.

**Listing 11.3 Identifying Service Accounts**

```
<Job id=" IdentifyingServiceAccounts" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objServiceAccountEnumeration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objServiceAccountEnumeration = _
  objMigration.CreateServiceAccountEnumeration

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"

   '
   'Enumerate service accounts on specified computers.
   '

   objServiceAccountEnumeration.Enumerate admtData, _
  Array(" computer name1" ," computer name2" )

   Set objServiceAccountEnumeration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a script file to assist you in creating a script to identify service accounts, see "Identifying Service Accounts" (DSSRENT_5.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Identifying Service Accounts" on the Web at http://www.microsoft.com/reskit).

# Migrating Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process of migrating account objects from a source domain to a target domain in another Active Directory forest involves first migrating service accounts, and then migrating global groups. After the groups are in place in the target domain, you can migrate users according to the process that you selected, either while using SID history for resource access, or without using SID history for resource access. When the account object migration process is complete, you can instruct users from the source domain to log on to the target domain. Figure 11.8 shows the process for migrating accounts between domains in different forests.

**Figure 11.8 Migrating Accounts**



# Transitioning Service Accounts in Your Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Begin the process of migrating objects by migrating service accounts. For information about identifying service accounts for migration, see "Identifying Service Accounts" earlier in this chapter. To transition service accounts, you must complete the following tasks:

- Migrate the service accounts from the source domain to the target domain.
- Modify the services on each server in the source domain to use the service account in the target domain in place of the service account in the source domain.

You can transition service accounts by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To transition service accounts by using the ADMT console**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool console, and then select **User Account Migration Wizard**.
3. Complete the User Account Migration Wizard by using the information in Table 11.6.

**Table 11.6 Using the User Account Migration Wizard to Transition Service Accounts**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| User Selection | Click **Add**.<br>In the **Select Users** dialog box, select all of the service accounts that were identified by the Service Account Migration Wizard, and then click **Add**.<br>Click **OK**. |
| Organizational Unit Selection | Click **Browse**.<br>In the **Browse for Container** dialog box, locate the source domain, select the container for the service accounts, and then click **OK**. |
| Password Options | Click **Complex passwords**. |
| Account Transition Options | Click **Enable target accounts**.<br>Click the **Migrate user SIDs to target domains** check box. |
| User Account | Type the user name, password, and domain of a user account that has administrative credentials. |

| | |
|---|---|
| User Options | Click the **Update user rights** check box.<br>Make sure that **Do Not Rename Accounts** is selected.<br>Ensure that no other settings are selected, including **Migrate associated user groups**. |
| Naming Conflicts | Click **Ignore conflicting accounts**. |
| Service Account Information | Click **Migrate all service accounts and update SCM for items marked include**. If you are also migrating other user accounts that are not service accounts, this wizard page tells you that you have selected some accounts that are marked as service accounts in the ADMT database. By default, the accounts are marked as Include. To change the status of the account, select the account, and then click **Skip/Include**.<br>Click **Next** to migrate the accounts. |

4. When the wizard has finished running, click **View Log** and review the migration log for any errors.
5. Open the Active Directory Users and Computers console, locate the OU that you created for service accounts, and verify that the service accounts exist in the target domain OU.

**To transition service accounts by using the ADMT command-line option**

1. On a domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT USER /N "
user_name1
" "
user_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /MSS:
YES
```

*Server name1* and *Server_name2* are the names of servers in the source domain that run service accounts. Alternatively, you can include parameters in an option file that is specified on the command line as follows:

```
ADMT USER /N "
server_name1
" "
server_name2
" /O:"
option_file.txt
"
```

Table 11.7 lists the common parameters used for transitioning service accounts, along with the command-line parameter and option file equivalents.

**Table 11.7 Common Parameters Used for Transitioning Service Accounts**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| Disable accounts | `/DOT:ENABLETARGET` *(default)* | `DisableOption=ENABLETARGET` *(default)* |
| Migrate password | `/PO:COMPLEX` (default) | `PasswordOption=COMPLEX` |
| Migrate user SIDs = YES | `/MSS:YES` | `MigrateSIDs=YES` |
| Update user rights=YES | `/UUR:YES` | `UpdateUserRights=YES` |
| Do not rename accounts | `/RO:DONT` (default) | `RenameOption=DONT` (default) |
| Ignore conflicting accounts | `/CO:IGNORE` (default) | `ConflictOptions=IGNORE` (default) |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers and locate the target service account OU. Verify that the service accounts exist in the target domain OU.

**To transition service accounts by using a script**

- Prepare a script that incorporates ADMT commands and options for transitioning service accounts by using the sample script shown in Listing 11.4.

**Listing 11.4 Transitioning Service Accounts Between Forests**

```
<Job id=" TransitioningServiceAccountsBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objUserMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objUserMigration = objMigration.CreateUserMigration

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " source container"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"
   objMigration.ConflictOptions = admtIgnoreConflicting

   '
   'Specify user migration specific options.
   '
```

```
   objUserMigration.MigrateSIDs = True
   objUserMigration.UpdateUserRights = True
   objUserMigration.MigrateServiceAccounts = True

   '
   'Migrate specified service accounts.
   '

   objUserMigration.Migrate admtData, _
  Array(" service account name1" ," service account name2" )

   Set objUserMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a script file to assist you in creating a script to transition service accounts, see "Transitioning Service Accounts Between Forests" (DSSREER_4.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Transitioning Service Accounts Between Forests" on the Web at http://www.microsoft.com/reskit).

# Migration of Global Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To preserve the global group user memberships, you must migrate global groups before you migrate users.

| **Note** |
| --- |
| Do not migrate global groups during peak work hours. The global group migration process can consume a large amount of network resources and resources on the domain controller that is running ADMT. |

Global group migration involves the following steps:

1. The administrator selects global group objects in the source domain.
2. A new global group object is created in the target domain. A new primary SID is created for the object in the target domain.
3. To preserve resource access, ADMT adds the SID of the global group in the source domain to the SID history attribute of the new global group in the target domain.

Following the migration, events are logged in both the source and the target domain.

| **Note** |
| --- |
| If the user account migration process takes place over an extended period of time, then you might need to remigrate global groups from the source to the target domain to propagate membership changes that are made in the source domain before the migration process is complete. For more information about remigrating global groups, see "Remigrating All Global Groups After All Batches Are Migrated" later in this chapter. |

You can migrate global groups by using the Active Directory Migration Tool console, by using the ADMT command-line option, or by using a script.

**To migrate global groups by using the ADMT console**

1. On the domain controller in the target domain on which ADMT installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool, and then select **Group Account Migration Wizard**.
3. Complete the Group Account Migration Wizard by using the information in Table 11.8.

**Table 11.8 Using the Group Account Migration Wizard to Migrate Groups**

| Wizard Page | Action |
| --- | --- |
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS or DNS name of the source domain or select the name from a list.<br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain.<br>If ADMT includes the names of the source and target domains, ensure that they are correct. |

| | |
|---|---|
| Group Selection | Click **Add**.<br>In the **Select Groups** dialog box, select all the global groups that you want to migrate (except built-in groups), click **Add**, and then click **OK**. |
| Organizational Unit Selection | Type the name of the OU, or click **Browse**.<br>In the **Browse for Container** dialog box, locate the container in the target domain you want to move the global groups into, and then click **OK**. |
| Group Options | Click **Migrate Group SIDs to target domain**.<br>Click **Do not rename accounts**<br>Make sure that all other options are not selected. |
| User Account | Type the user name, password, and domainof an account that has administrative rights in the source domain. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

4. When the wizard has finished running, click **View Log**, and review the migration log for any errors.
5. Open the Active Directory Users and Computers console and locate the target OU. Verify that the global groups exist in the target domain OU.

**To migrate global groups by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" [
parameters
]
```

Alternatively, you can include parameters in an option file that is specified on the command line as follows:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /O: "
option_file.txt
"
```

Table 11.9 lists the common parameters used for migrating global groups, along with the command-line parameter and

option file equivalents.

**Table 11.9 Common Parameters Used for Global Group Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Source OU location | `/SO:"`*source_OU*`"` | `SourceOU="`*source_OU*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| Migrate GG SIDs | `/MSS:YES` | `MigrateSIDs=YES` |
| Do not rename accts | `/RO:DONT` *(default)* | `RenameOption=DONT` |
| Ignore conflicting accts and do not migrate them | `/CO:IGNORE` *(default)* | `ConflictOptions=IGNORE` |

3. Review the results that are displayed on the screen for any errors.
4. Open the Active Directory Users and Computers console and locate the target OU. Verify that the global groups exist in the target domain OU.

**To migrate global groups by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating global groups by using the sample script shown in Listing 11.5

**Listing 11.5 Migrating Global Groups Between Forests**

```
<Job id=" MigratingGlobalGroupsBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objGroupMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objGroupMigration = objMigration.CreateGroupMigration

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
    objMigration.SourceOu = " source container"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"

   '
   'Specify group migration specific options.
   '
```

```
   objGroupMigration.MigrateSIDs = True

   '
   'Migrate specified group objects.
   '

   objGroupMigration.Migrate admtData, Array(" group name1" ," group name2" )

   Set objGroupMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a script file to assist you in creating a script to migrate global groups, see "Migrating Global Groups Between Forests" (DSSREER_5.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Global Groups Between Forests" on the Web at http://www.microsoft.com/reskit).

# Migrating Accounts While Using SID History

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To migrate accounts while using SID history, first migrate all user accounts, but do not enable them in the target domain, to prepopulate the target domain and allow migration of user profiles. After all user accounts are successfully migrated, begin migrating users in batches by migrating first the user profile, then the workstation, then the user account. Before migrating all user accounts, ensure that you have created test accounts that you can use to verify the success of each batch.

You cannot migrate every user property when you migrate user accounts. For example, Protected Storage (Pstore) contents for Windows NT 4.0 workstations, including Encrypting File System (EFS) private keys, are not migrated by ADMT when you migrate user accounts. To migrate Pstore contents, you must export and import keys during the migration process.

For clients that are running Windows 2000 or later operating systems, data that is protected by the Data Protection API (DPAPI) is also not migrated. DPAPI helps protect the following items:

- Web page credentials (for example, passwords)
- File share credentials
- Private keys associated with EFS, S/MIME, and other certificates
- Program data that is protected by using the CryptProtectData() function

For this reason, it is important to test user migrations by using your test migration account to identify any properties that did not migrate and update user configurations in the target domain accordingly.

Complete the following steps to migrate user accounts to the target domain:

1. Migrate all user accounts with the account enabled in the source domain, disabled in the target domain, with complex password selected, and with no attributes migrated.
2. Translate local user profiles for a batch of users.
3. Migrate workstations in batches that correspond to the user account batches.
4. Before migrating the batch of user accounts, verify that local profile and workstation migration succeeded for all users in the batch. Do not migrate any user account for which profile or workstation migration failed, because this will result in users overwriting their existing profiles when they log on to the target domain.
5. Remigrate user accounts in batches with the account set to expire in the source domain in seven days, the target account enabled, with password migration selected, and all attributes migrated.
6. After each batch, you might want to remigrate all global groups to update any group membership changes.
7. Notify users in the batch to log on to the target domain.
8. After all users are migrated, run a final global group migration to update any group membership changes.

Migrating user accounts in batches enables you to track the accounts that have been migrated and to test the success of each migration step. If the OU structure for the target domain is the same as the OU structure for the source domain, then migrate groups of users based on OU. If the OU structures are not the same, select an alternate way to group users based on the structure of your organization. For example, you might migrate users by business unit or by floor to enable you to consolidate help desk resources.

If you plan to retain your source domain OU structure, migrate the OU along with the users that they contain. For example, if your source domain is a Windows Server 2003 Active Directory environment that has a functional OU structure, and the target

domain does not have an OU structure, migrate OUs from the source domain.

If you created a new OU structure in the target domain, migrate batches of users without the OUs. For example, if your source environment was a Windows NT 4.0 domain that you upgraded to a Windows Server 2003 domain, the source domain might not have an existing OU structure; therefore, you can migrate users without migrating OUs.

For more information about designing an OU structure, see "Designing the Active Directory Logical Structure" in this book.

Until you migrate all user and group accounts, continue to administer global group membership in the source domain. To support a rollback strategy, manually synchronize any changes you make to users in the target domain with the existing user accounts in the source domain. For more information about administering users and groups during the interforest restructure process, see "Creating Administrative Procedures" earlier in this chapter.

If you are migrating OUs when you migrate user accounts, migrate the groups that belong to those OUs to the target domain OU during the user account migration process. When you migrate global groups by using the global group migration process, they are placed in the target OU in the target domain. If you migrate OUs from the source to the target domain, select the option to move the global groups to the target domain at the same time so that the groups are moved from the target OU that they were placed in during the initial global group migration to the OU in which they belong.

Using ADMT to migrate user accounts preserves group memberships. Because global groups can contain only members from the domain in which the group is located, when users are migrated to a new domain, the user accounts in the target domain cannot be members of the global groups in the source domain. As part of the migration process, ADMT identifies the global groups in the source domain that the user accounts belong to, and then determines whether the global groups have been migrated. If ADMT identifies global groups in the target domain that the migrated users belonged to in the source domain, the tool adds the users to the appropriate global groups in the target domain.

Using ADMT to migrate user accounts also preserves user passwords. After the user accounts are migrated to and enabled in the target domain, the users can log on to the target domain by using their original passwords. If the user account migration process is successful but the password migration process fails, ADMT creates a new complex password for the user account in the target domain. By default, ADMT stores new complex passwords in the C:\Program Files\Active Directory Migration Tool\Logs\Password.txt file.

If you have a Group Policy setting on the target domain that does not allow blank passwords (the **Default Domain Policy/Computer Configuration/Security Settings/Account Policies/Password Policy/Minimum password length** setting is set to any number other than zero), then password migration will fail for any user who has a blank password. ADMT generates a complex password for that user, and writes an error to the error log.

Establish a method for notifying users who have been assigned new passwords. For example, you can create a script to send an e-mail message to users to notify them of their new passwords.

Figure 11.9 shows the steps involved in migrating accounts if you are using SID history for resource access.

**Figure 11.9 Migrating Accounts While Using SID History**



# Migrating All User Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Begin the user account migration process by migrating all users. This enables you to translate local profiles and ensure that users continue to have the appropriate resource access following the migration.

**Notes**

- Built-in accounts (such as Administrators, Users, and Power Users cannot be ADMT migration objects. Because built-in account SIDs are identical in every domain, migrating these accounts to a target domain results in duplicate SIDs in a single domain. Every SID in a forest must be unique.
- Well-known accounts (such as Domain Admins and Domain Users) also cannot be ADMT migration objects

The ADMT user account migration process includes the following steps:

1. ADMT reads the attributes of the source user objects.
2. ADMT creates a new user object in the target domain and a new primary SID for the new user account.
3. ADMT adds the original SID of the user account to the SID history attribute of the new user account.
4. ADMT migrates the password for the user account.
5. If ADMT identifies global groups in the target domain that the migrated users belonged to in the source domain, the tool adds the users to the appropriate global groups in the target domain.

During the migration, audit events are logged in both the source and the target domains.

You can migrate user accounts by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate user accounts by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool, and then select **User Account Migration Wizard**.
3. Complete the User Account Migration Wizard by using the information in Table 11.10.

**Table 11.10 Using the User Account Migration Wizard to Migrate User Accounts**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?**. |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| User Selection | Click **Add**.<br>In the **Select Users** dialog box, click all the user accounts, and then click **Add**. By default, the wizard migrates the accounts to the Users container. Click **Do Not Migrate Passwords** (use complex passwords).<br>Click **OK**. |

| | |
|---|---|
| Organizational Unit Selection | ADMT lists an OU here. Ensure that this is the correct target OU. If it is not correct, type the correct OU or click **Browse**.<br>In the **Browse for Container** dialog box, locate the target domain and OU, and then click **OK**. |
| Password Options | Click **Do NotMigrate Passwords**.<br>Click **Complex Passwords**. |
| Account Transition Options | Click **Disable target accounts**.<br>Click **Enable source account**.<br>Click the **Migrate user SIDs to target domains** check box. |
| User Account | Type the user name, password, and domainof a user account that has administrative credentials. |
| User Options | Click the **Translate roaming profiles** check box.<br>Click the **Update user rights** check box.<br>Clear the **Migrate associated user groups** check box.<br>Click **Fix users' group memberships**.<br>Click the **Do not rename accounts** check box. |
| Object Property Exclusion | Clear the **Exclude specific object properties from migration** check box. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**.<br>Clear the **Remove existing user rights** check box.<br>Clear the **Move replaced accounts to specific target Organizational Unit** check box. |

4. When the wizard has finished running, click **View Log** and review the migration log for any errors.
5. Open Active Directory Users and Computers and verify that the user accounts exist in the appropriate OU in the target domain.

**To migrate user accounts by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT USER /N "user_name1" "user_name2"  /SD:"source_domain"
/TD:"target_domain"  /TO:"target_OU"  [parameters]
```

You can append parameters to the command as follows:

```
ADMT USER /N "user_name1" "user_name2"  /SD:"source_domain"
/TD:"target_domain"  /TO:"target_OU" /MSS:YES TRP:YES /UUR:YES
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT USER /N "user_name1" "user_name2" /O "option_file.txt"
```

Table 11.11 lists the common parameters used for migrating user accounts, along with the command-line parameter and

option file equivalents.

**Table 11.11 Common Parameters Used for User Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| Migrate SIDs | `/MSS:YES` | `MigrateSIDs=YES` |
| Do not rename accts | `/RO:DONT` *(default)* | `RenameOption=DONT` |
| Ignore conflicting accts and not migrate them | `/CO:IGNORE` *(default)* | `ConflictOptions=IGNORE` |
| Translate Roaming Profile | `/TRP:YES` *(default)* | `TranslateRoamingProfile=YES` |
| Update User Rights | `/UUR:NO` | `UpdateUserRights=NO` |
| Password Options | `/PO:COMPLEX` | `PasswordOption=COMPLEX` |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers and locate the target OU. Verify that the users exist in the target OU.

**To migrate user accounts by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating users by using the sample script shown in Listing 11.6.

**Listing 11.6 Migrating All User Accounts Between Forests**

```
<Job id=" MigratingAllUserAccountsBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objUserMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objUserMigration = objMigration.CreateUserMigration

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " source container"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"
   objMigration.PasswordOption = admtComplexPassword
   objMigration.ConflictOptions = admtIgnoreConflicting
```

```
    '
    'Specify user migration specific options.
    '

    objUserMigration.MigrateSIDs = True
    objUserMigration.TranslateRoamingProfile = True
    objUserMigration.UpdateUserRights = True
    objUserMigration.FixGroupMembership = True
    objUserMigration.MigrateServiceAccounts = False

    '
    'Migrate specified user objects.
    '

    objUserMigration.Migrate admtData, Array(" user name1" ," user name2" )

    Set objUserMigration = Nothing
    Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in creating a script to migrate user accounts, see "Migrating All User Accounts Between Forests" (DSSREER_6.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating All User Accounts Between Forests" on the Web at http://www.microsoft.com/reskit).

# Remigrating User Accounts and Workstations in Batches

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remigrating user accounts and workstations in batches helps you track the migration process. For each batch of users, first translate local user profiles, and then migrate workstations. Verify that the profile and workstation migration succeeded, and then migrate the user accounts. Remigrate global groups after each batch.

**Translating Local User Profiles**
ADMT only translates profiles for computers running Windows NT 4.0, Windows 2000, Microsoft® Windows® XP, or Windows Server 2003.

User profiles are stored locally on the workstation. When a user logs on to another workstation, he or she must create a new, unique local user profile. Translate the local user profiles immediately after migrating a batch of users.

Local profiles are translated in replace mode because if you perform the profile translation in add mode, certain aspects of software installation that use software deployment Group Policies might not work. Any application that is packaged with Windows Installer version 2.0 (which is included on workstations running Windows 2000 SP3 or later or Windows XP SP1 or later, as well as many common software packages) might not function after the profile is translated. For example, the application executable files might not be removed after the last user removed the application. When the ADMT Security Translation Wizard translates local profiles in replace mode, it reverts to add mode if a profile cannot be loaded because it is already in use by the user or a service using the user account. In some cases, this can cause some aspects of the software deployment Group Policies to fail to function after the profile is translated.

| **Note** |
| --- |
| Translate the local user profiles the night before you notify the users to log on by using their new accounts in the target domain. Migrating profiles the night before ensures that the new user profile reflects the most current user settings. |

You can translate local user profiles by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To translate local user profiles by using the ADMT console**

1. For each workstation in the source domain that is running Windows NT 4.0, Windows 2000, or Windows XP, add the ADMT resource migration account to the local Administrators group.
2. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
3. Open the Active Directory Migration Tool, and then select **Security Translation Wizard**.
4. Complete the Security Translation Wizard by using the information in Table 11.12.

   **Table 11.12 Using the Security Translation Wizard to Translate Local User Profiles**

   | Wizard Page | Action |
   | --- | --- |
   | Test or Make Changes | Click **Migrate Now?** |
   | Security Translation Options | Click **Previously migrated objects**. |
   | Domain Selection | In the **Source domain** box, type or select the name of the source account domain.<br><br>In the **Target domain** box, type or select the name of the target domain. |
   | Translate Objects | Click **User Profiles**. |
   | Security Translation Options | Click **Replace**. |

5. When the wizard has finished running on all computers, click **View Log** and review the dispatch log for any errors.

**To translate local user profiles by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" [
```

```
parameters
]
```

You can append parameters to the command as follows:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TOT:REPLACE /TUP:YES
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" /O "
option_file.txt
"
```

Table 11.13 lists the common parameters used for migrating user accounts, along with the command-line parameter and option file equivalents.

**Table 11.13 Common Parameters Used for Local User Profile Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Translate option = REPLACE | /TOT:REPLACE | TranslateOption=REPLACE |
| Modify local user profile security | /TUP:YES | TranslateUserProfiles=YES |

3. Review the results that are displayed in the dispatch log for any errors.

**To translate local user profiles by using a script**

- Prepare a script that incorporates ADMT commands and options for translating local user profiles by using the sample script shown in Listing 11.7.

**Listing 11.7 Translating Local User Profiles Between Forests**

```
<Job id=" TranslatingLocalProfilesBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objSecurityTranslation

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objSecurityTranslation = objMigration.CreateSecurityTranslation

   '
```

```
    'Specify general migration options.
    '

    objMigration.SourceDomain = " source domain"
    objMigration.TargetDomain = " target domain"
    objMigration.TargetOu = " Computers"

    '
    'Specify security translation specific options.
    '

    objSecurityTranslation.TranslationOption = admtTranslateReplace
    objSecurityTranslation.TranslateUserProfiles = True

    '
    'Perform security translation on specified computer objects.
    '

    objSecurityTranslation.Translate admtData, _
   Array(" computer name1" ," computer name2" )

    Set objSecurityTranslation = Nothing
    Set objMigration = Nothing
</Script>
</Job>
```

For a sample script to assist you in translating local user profiles, see "Translating Local User Profiles Between Forests" (DSSREER_7.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Local User Profiles Between Forests" on the Web at http://www.microsoft.com/reskit).

**Migrating Workstations in Batches**

After you migrate a batch of local user profiles, migrate the corresponding batch of user workstations. When you move workstations between domains, the SAM database moves along with them. Accounts located in the local SAM database (such as local groups) that are used to enable access to resources always move with the computer, and therefore do not need to be migrated.

**Note**

- Use a low value in the ADMT computer restart option, to restart workstations immediately after joining them to the target domain, or as soon as possible thereafter. Resources that are not restarted following migration are in an indeterminate state.

You can migrate workstations and member servers by using the AMDT console, by using the ADMT command-line option, or by using a script.

**To migrate workstations by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT resource migration account.
2. Open the Active Directory Migration Tool, and then select **Computer Account Migration Wizard**.
3. Complete the Computer Account Migration Wizard by using the information in Table 11.14.

**Table 11.14 Using the Computer Account Migration Wizard to Migrate Workstations**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br><br>In the **Target domain** box, type or select the name of the target domain. |
| Computer Selection | Click **Add**.<br><br>In the **Select Computers** dialog box, select the names of all the workstations and member servers in the source domain, click **Add**, and then click **OK**. |

| | |
|---|---|
| Organizational Unit Selection | Click **Browse**.<br><br>In the **Browse for Container** dialog box, locate the target domain **Computers** container or the appropriate OU, and then click **OK**. |
| Security Translation Options | Click the **Local groups** check box.<br><br>Click the **User rights** check box. |
| Translate Objects | Click **Add**. |
| Computer Options | In the **Minutes before computer restart after wizard completion** box, type **1**. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate.** |

4. When the wizard has finished running on all computers, click **View Log** and review the migration log and dispatch log for any errors.
5. If errors are reported, run the Retry Wizard to dispatch agents to workstations for which the migration process failed.
6. Open Active Directory Users and Computers and verify that the workstations exist in the appropriate OU in the target domain.

**To migrate workstations by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT installed, log on by using the ADMT resource migration account.
2. At the command line, type:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /RDL:1
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /O "
option_file.txt
"
```

Table 11.15 lists the common parameters used for workstation migration, along with the command-line parameter and option file equivalents.

**Table 11.15 Common Parameters Used for Workstation Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Source OU location | `/SO:"`*source_OU*`"` | `SourceOU="`*source_OU*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| • Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| • Restart Delay | `/RDL:1` | `RestartDelay=1` |
| • Do not rename accounts | `/RO:DONT` *(default)* | `RenameOption=DONT` |
| • Translate Option | `/TOT:ADD` | `TranslateOption=YES` |
| • Translate Option | `/TUR:YES` | `TranslateOption=ADD` |
| • Translate User Rights | `/TLG:YES` | `TranslateUserRights=YES` |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers and locate the target OU. Verify that the workstations and member servers exist in the target OU.

**To migrate workstations by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating workstations by using the sample script shown in Listing 11.8.

**Listing 11.8 Migrating Workstations Between Forests**

```
<Job id=" MigratingWorkstationsBwtweenForest" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objComputerMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objComputerMigration = objMigration.CreateComputerMigration

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " Computers"
   objMigration.TargetDomain = " target domain"
```

```
    objMigration.TargetOu = " Computers"

    '
    'Specify computer migration specific options.
    '

    objComputerMigration.RestartDelay = 1
    objComputerMigration.TranslationOption = admtTranslateAdd
    objComputerMigration.TranslateLocalGroups = True
    objComputerMigration.TranslateUserRights = True



    '
    'Migrate computer objects on specified computer objects.
    '

    objComputerMigration.Migrate admtData, _
   Array(" computer name1" ," computer name2" )

    Set objComputerMigration = Nothing
    Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in migrating workstations between forests, see "Migrating Workstations Between Forests" (DSSREER_9.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations Between Forests" on the Web at http://www.microsoft.com/reskit).

**Remigrating User Accounts in Batches**

After you have verified the success of local user profile and user workstation migration for the user batch, migrate the user accounts for that batch. You can migrate user accounts in batches by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate the current batch of user accounts by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool, and then select **User Account Migration Wizard**.
3. Complete the User Account Migration Wizard by using the information in Table 11.16.

**Table 11.16 Using the User Account Migration Wizard to Migrate User Accounts**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br><br>In the **Target domain** box, type or select the name of the target domain. |
| User Selection | Click **Add**.<br><br>In the **Select Users** dialog box, select the user accounts that you want to migrate, and then click **Add**. By default, the wizard migrates the accounts to the Users OU. If you need to change the OU, click **Location**, and then browse for the appropriate OU.<br><br>Click **Migrate Passwords**.<br><br>Click **OK**. |
| Organizational Unit Selection | Ensure that ADMT lists the correct target OU. If it is not correct, type the correct OU or click **Browse**.<br><br>In the **Browse for Container** dialog box, locate the target domain **Users** container or appropriate OU, and then click **OK**. |
| Password Options | Click **Migrate Passwords**.<br><br>Click **Password migration source DC.** |

| | |
|---|---|
| Account Transition Options | Click **Enable target accounts**.<br><br>Click **Days until source account expires,** and then type **7** in the text box.<br><br>Click the **Migrate user SIDs to target domains** check box. |
| User Account | Type the user name, password, and domainof a user account that has administrative credentials. |
| User Options | Click the **Translate roaming profiles** check box.<br><br>Click the **Update user rights** check box.<br><br>Clear the **Migrate associated user groups** check box<br><br>Click **Fix users' group memberships**.<br><br>Click **Do not rename accounts** check box. |
| Object Property Exclusion | Clear the **Exclude specific object properties from migration** check box. |
| Naming Conflicts | Click **Replace conflicting accounts and don't migrate**.<br><br>Clear the **Remove existing user rights** check box.<br><br>Clear the **Move replaced accounts to specific target Organizational Unit** check box. |

4. When the wizard has finished, click **View Log**, and review the migration log for any errors.
5. Open Active Directory Users and Computers and verify that the user accounts exist in the appropriate OU in the target domain.

**To migrate the current batch of users by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT USER /N "user_name1""user_name2" /SD:"source_domain"
/TD:"target_domain"  /TO:"target_OU"  [parameters]
```

You can append parameters to the command as follows:

```
ADMT USER /N "user_name1""user_name2" /SD:"source_domain"
/TD:"target_domain"  /TO:"target_OU"  /MSS:YES TRP:YES /UUR:YES
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT USER /N "user_name1""user_name2" /O "option_file.txt"
```

Table 11.17 lists the common parameters used for migrating user accounts, along with the command-line parameter and option file equivalents.

**Table 11.17 Common Parameters Used for User Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Source OU location | `/SO:"`*source_OU*`"` | `SourceOU="`*source_OU*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |

| | | |
|---|---|---|
| Migrate SIDs | /MSS:YES | MigrateSIDs=YES |
| Do not rename accts | /RO:DONT *(default)* | RenameOption=DONT |
| Ignore conflicting accts and not migrate them | /CO: REPLACE+REMOVEUSERRIGHTS+MOVEREPLACEDACCOUNTS | ConflictOptions=REPLACE+REMOVEUSERRIGHTS+MOVEREPLACEDACCOUNTS |
| Translate Roaming Profile | /TRP:YES *(default)* | TranslateRoamingProfile=YES |
| Update User Rights | /UUR:NO | UpdateUserRights=NO |
| Password Options | /PO:COPY /PS:<name of PES server> | PasswordOption=COPY<br>PasswordServer=<name of server> |
| Expire accounts | /SEP:7 | SourceExpiration=7 |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers and locate the target OU. Verify that the users exist in the target OU.

**To migrate the current batch of user accounts by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating users by using the sample script shown in Listing 11.9.

**Listing 11.9 Migrating User Accounts in Batches Between Forests**

```
<Job id=" MigratingUserAccountsInBatchesBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objUserMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objUserMigration = objMigration.CreateUserMigration

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " source container"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"
   objMigration.PasswordOption = admtCopyPassword
   objMigration.PasswordServer = " password export server name"
   objMigration.ConflictOptions = admtReplaceConflicting + _
  admtRemoveExistingUserRights + admtMoveReplacedAccounts
   '
   'Specify user migration specific options.
   '
   objUserMigration.SourceExpiration = 7
   objUserMigration.MigrateSIDs = True
```

```
    objUserMigration.TranslateRoamingProfile = True
    objUserMigration.UpdateUserRights = False
    objUserMigration.FixGroupMembership = True
    objUserMigration.MigrateServiceAccounts = False

    '
    'Migrate specified user objects.
    '

    objUserMigration.Migrate admtData, Array(" user name1" ," user name2" )

    Set objUserMigration = Nothing
    Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in creating a script to migrate user accounts, see "Migrating User Accounts in Batches Between Forests" (DSSREER_8.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating User Accounts in Batches Between Forests" on the Web at http://www.microsoft.com/reskit).

**Remigrating All Global Groups Following User Account Migration**

A large user account migration might take place over an extended period of time. For this reason, you might need to remigrate global groups from the source to the target domain after you migrate each batch of users, to reflect changes made to the membership of groups in the source domain after the initial global group migration occurred. For more information about and procedures for remigrating global groups, see "Remigrating All Global Groups After All Batches Are Migrated" later in this chapter.

# Remigrating All Global Groups After All Batches Are Migrated

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After all batches have been migrated, perform a final global group remigration to ensure that any late changes made to global group membership in the source domain are reflected in the target domain. You can remigrate global groups by using the Active Directory Migration Tool console, by using the ADMT command-line option, or by using a script.

**To remigrate global groups by using the ADMT console**

1. On the domain controller in the target domain on which ADMT installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool, and then select **Group Account Migration Wizard**.
3. Complete the Group Account Migration Wizard by using the information in Table 11.18.

   **Table 11.18 Using the Group Account Migration Wizard to Remigrate Global Groups**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS or DNS name of the source domain or select the name from a list.<br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain.<br>If ADMT includes the names of the source and target domains, ensure that they are correct. |
| Group Selection | Click **Add**.<br>In the **Select Groups** dialog box, select all global groups, click **Add**, and then click **OK**. |
| Organizational Unit Selection | Type the name of the OU, or click **Browse**.<br>In the **Browse for Container** dialog box, locate the container in the target domain you want to move the global groups into, and then click **OK**. |
| Group Options | Click **Update user rights**.<br>Ensure that **Copy group members** is not selected**.**<br>Ensure that **Update previously migrated objects** is not selected.<br>Click **Fix membership of group**.<br>Click **Migrate Group SIDs to target domain**.<br>Click **Do not rename accounts** |
| User Account | Type the user name, password, and domain of an account that has administrative rights in the source domain. |
| Object Property Exclusion | Clear the **Exclude specific object properties from migration** check box. |
| Naming Conflicts | Click **Replace conflicting accounts** (all other options are cleared). |

4. When the wizard has finished running, click **View Log**, and review the migration log for any errors.

5. Open Active Directory Users and Computers and locate the target OU. Verify that the global groups exist in the target domain OU.

**To remigrate global groups by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /MSS:YES
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /O: "
option_file.txt
"
```

Table 11.19 lists the common parameters used for migrating global groups, along with the command-line parameter and option file equivalents.

**Table 11.19 Common Parameters Used for Global Group Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Source OU location | `/SO:"`*source_OU*`"` | `SourceOU="`*source_OU*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |

| | | |
|---|---|---|
| Target OU location | /TO:"*target_OU*" | TargetOU="*target_OU*" |
| Migrate GG SIDs | /MSS:YES | MigrateSIDs=YES |
| Do not rename accts | /RO:DONT *(default)* | RenameOption=DONT |
| Ignore conflicting accts and do not migrate them | /CO:REPLACE | ConflictOptions=REPLACE |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers and locate the target OU. Verify that the global groups exist in the target domain OU.

**To remigrate global groups by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating global groups by using the sample script shown in Listing 11.10.

**Listing 11.10 Remigrating Global Groups Between Forests**

```
<Job id=" RemigratingGlobalGroupsBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objGroupMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objGroupMigration = objMigration.CreateGroupMigration

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " source container"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"
   objMigration.ConflictOptions = admtReplaceConflicting

   '
   'Specify group migration specific options.
   '

   objGroupMigration.MigrateSIDs = True

   '
   'Migrate specified group objects.
   '

   objGroupMigration.Migrate admtData, Array(" group name1" ," group name2" )

   Set objGroupMigration = Nothing
   Set objMigration = Nothing
```

```
</Script>
</Job>
```

For a script file to assist you in creating a script to migrate global groups, see "Remigrating Global Groups Between Forests" (DSSREER_10.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Remigrating Global Groups Between Forests" on the Web at http://www.microsoft.com/reskit).

# Migrating Accounts Without Using SID History

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are not using SID history for resource access because SID filter quarantining is in place between your forests, your migration process involves first migrating all user accounts, but not enabling them in the target domain, to prepopulate the target domain and allow migration of user profiles. Then you run security translation on all resources that the users access across forests. The next step is to migrate users in batches by migrating first the user profile, then the workstation, then the user account. Finally, you must remigrate the global groups to apply any changes made to the global groups in the source domain, and translate security in remove mode.

It is still important to migrate SID history although user accounts will not use SID history for resource access. This ensures that operations such as Offline Files continue to function within the forest. Migrating SID history does not present a security risk because SID filter quarantining is in place between the source and target forests. Before migrating all user accounts, ensure that you have created test accounts that you can use to verify the success of each batch.

Complete the following steps to migrate user accounts to the target domain:

1. Migrate all users, and use the **Fix users' group membership** option to have ADMT identify global groups in the target domain that the user belonged to in the source domain and add the user to the appropriate global group in the target domain. For this initial user migration, leave the user account enabled in source domain, and disabled in the target domain.

2. Translate security in add mode for files, shares, printers, local groups, and domain local groups.

3. Translate local user profiles for a batch of users.

4. Migrate workstations in batches that correspond to the user account batches.

5. Before migrating the batch of user accounts, verify that local profile and workstation migration succeeded for all users in the batch. Do not migrate any user account for which profile or workstation migration failed, because this will result in users overwriting their existing profiles when they log onto the target domain.

6. Remigrate user accounts in small batches with the accounts in the source domain set to expire in seven days, the target accounts enabled, password migration selected, and all attributes selected for migration.

7. After each batch, you might want to remigrate all global groups to update any group membership changes.

8. After all users are migrated, run a final global group migration to update any group membership changes

9. Translate security in remove mode for files, shared folders, printers, local groups, and domain local groups.

10. Notify users in the batch to log on to the target domain.

Until you migrate all user and group accounts, continue to administer global group membership in the source domain.

Figure 11.10 shows the steps involved in migrating accounts that are not using SID history for resources access.

**Figure 11.10 Migrating Accounts Without Using SID History**


Migrate all user accounts
Translate security in add mode
Remigrate user accounts and workstations in batches
Remigrate all global groups
Translate security in remove mode
Translate local user profiles
Migrate workstations in batches
Remigrate user accounts in batches
Remigrate all global groups
All batches migrated?
No
Yes
Done

# Migration of All User Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Begin the user account migration process by migrating all users. This enables you to translate security on all files, printers, shared folders, local groups, and domain local groups to ensure that users continue to have the appropriate resource access following the migration.

**Notes**

- Built-in accounts (such as Administrators, Users, and Power Users) cannot be ADMT migration objects. Because built-in account SIDs are identical in every domain, migrating these accounts to a target domain results in duplicate SIDs in a single domain. Every SID in a forest must be unique.

- Well-known accounts (such as Domain Admins and Domain Users) also cannot be ADMT migration objects.

The ADMT user account migration process includes the following steps:

1. ADMT reads the attributes of the source user objects.

2. ADMT creates a new user object in the target domain and a new primary SID for the new user account.

3. ADMT adds the original SID of the user account to the SID history attribute of the new user account.

4. ADMT migrates the password for the user account.

5. If ADMT identifies global groups in the target domain that the migrated users belonged to in the source domain, the tool adds the users to the appropriate global groups in the target domain.

During the migration, audit events are logged in both the source and the target domains.

You can migrate user accounts by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate user accounts by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool, and then select **User Account Migration Wizard**.
3. Complete the User Account Migration Wizard by using the information in Table 11.20.

**Table 11.20 Using the User Account Migration Wizard to Migrate User Accounts**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| User Selection | Click **Add**.<br>In the **Select Users** dialog box, select all user accounts, and then click **Add**. By default, the wizard migrates the accounts to the Users OU. Click **Do Not Migrate Passwords** (use complex passwords).<br>Click **OK**. |

| | |
|---|---|
| Organizational Unit Selection | Ensure that ADMT lists the correct target OU. If it is not correct, type the correct OU or click **Browse**.<br>In the **Browse for Container** dialog box, locate the target domain **Users** container or appropriate OU, and then click **OK**. |
| Password Options | Click **Do NotMigrate Passwords**.<br>Click **Complex Passwords**. |
| Account Transition Options | Click **Disable target accounts**.<br>Click **Enable source account**.<br>Click the **Migrate user SIDs to target domains** check box. |
| User Account | Type the user name, password, and domain of a user account that has administrative credentials. |
| User Options | Click the **Translate roaming profiles** check box.<br>Click the **Update user rights** check box.<br>Clear the **Migrate associated user groups** check box.<br>Click **Fix users' group membership**.<br>Select the **Do not rename accounts** check box. |
| Object Property Exclusion | Clear the **Exclude specific object properties from migration** check box. |

4. When the wizard finishes, click **View Log**, and review the migration log for any errors.
5. Open Active Directory Users and Computers and verify that the user accounts exist in the appropriate OU in the target domain.

**To migrate user accounts by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT USER /N "user_name1" "user_name2" /SD:"source_domain"
/TD:"target_domain" /TO:"target_OU"  [parameters]
```

You can append parameters to the command as follows:

```
ADMT USER /N "user_name1" "user_name2" /SD:"source_domain"
/TD:"target_domain" /TO:"target_OU" /MSS:YES TRP:YES /UUR:YES
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT USER /N "user_name1" "user_name2" /O "option_file.txt"
```

Table 11.21 lists the common parameters used for migrating user accounts, along with the command-line parameter and option file equivalents.

**Table 11.21 Common Parameters Used for User Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|

| | | |
|---|---|---|
| Source domain | /SD:"*source_domain*" | SourceDomain="*source_domain*" |
| Source OU location | /SO:"*source_OU*" | SourceOU="*source_OU*" |
| Target domain | /TD:"*target_domain*" | TargetDomain="*target_domain*" |
| Target OU location | /TO:"*target_OU*" | TargetOU="*target_OU*" |
| Migrate SIDs | /MSS:YES | MigrateSIDs=YES |
| Do not rename accts | /RO:DONT *(default)* | RenameOption=DONT |
| Ignore conflicting accts and not migrate them | /CO:IGNORE *(default)* | ConflictOptions=IGNORE |
| Translate Roaming Profile | /TRP:YES *(default)* | TranslateRoamingProfile=YES |
| Update User Rights | /UUR:YES | UpdateUserRights=YES |
| Password Options | /PO:COMPLEX | PasswordOption=COMPLEX |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers and locate the target OU. Verify that the users exist in the target OU.

**To migrate user accounts by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating users.

  For a sample script file to assist you in creating a script to migrate all user accounts, see "Migrating All User Accounts Between Forests" (DSSREER_6.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating All User Accounts Between Forests" on the Web at http://www.microsoft.com/reskit).

# Translating Security in Add Mode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Translate security on servers to add the SIDs of the user and group accounts in the target domain to the ACLs of the resources. After objects are migrated to the target domain, the objects contain the ACL entries from both the source and the target domains. Use the Security Translation Wizard in ADMT to add the target domain SIDs from the migrated objects. Run the Security Translation Wizard on all files, shares, printers, local groups, and at least one domain controller (to translate security on shared local groups).

You can translate security in add mode on objects by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To translate security in add mode on objects by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT account migration account.

2. Open the Active Directory Migration Tool, and then select **Security Translation Wizard**.

3. Complete the Security Translation Wizard by using the information in Table 11.22.

**Table 11.22 Using the ADMT Security Translation Wizard in Add Mode**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Security Translation Options | Click **Previously migrated objects**. |
| Domain Selection | In the **Source domain** box, type or select the name of the source account domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| Translate Objects | Clear the **User Profiles** check box.<br>Select all other check boxes. |
| Security Translation Options | Click **Add**. |

**To translate security in add mode on objects by using a script**

- Prepare a script that incorporates ADMT commands and options for translating security in add mode on objects by using the sample script shown in Listing 11.11.

**Listing 11.11 Translating Security in Add Mode on Objects Between Forests**

```
<Job id=" TranslatingSecurityInAddModeOnObjectsBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objSecurityTranslation

   '
   'Create instance of ADMT migration objects.
   '
```

```
   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objSecurityTranslation = objMigration.CreateSecurityTranslation

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " Computers"

   '
   'Specify security translation specific options.
   '

   objSecurityTranslation.TranslationOption = admtTranslateAdd
   objSecurityTranslation.TranslateFilesAndFolders = True
   objSecurityTranslation.TranslateLocalGroups = True
   objSecurityTranslation.TranslatePrinters = True
   objSecurityTranslation.TranslateRegistry = True
   objSecurityTranslation.TranslateShares = True
   objSecurityTranslation.TranslateUserProfiles = False
   objSecurityTranslation.TranslateUserRights = True

   '
   'Perform security translation on specified computer objects.
   '

   objSecurityTranslation.Translate admtData, _
  Array(" computer name1" ," computer name2" )

   Set objSecurityTranslation = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in creating a script to translate security in add mode on objects, see "Translating Security in Add Mode on Objects Between Forests" (DSSREER_11.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Security in Add Mode on Objects Between Forests" on the Web at http://www.microsoft.com/reskit).

# Remigration of User Accounts and Workstations in Batches

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remigrating user accounts and workstations in batches helps you track the migration process. For each batch of users, first translate local user profiles, and then migrate workstations. Verify that the profile and workstation migration succeeded, and then migrate the user accounts. Remigrate global groups after each batch.

**Translating Local User Profiles**
ADMT only translates profiles for computers running Windows NT 4.0, Windows 2000, Windows XP, or Windows Server 2003.

Local profiles are translated in replace mode because if you perform the profile translation in add mode, software installation by means of software deployment Group Policies might not work. Any application that is packaged with Windows Installer version 2.0 (which is used on workstations running Windows 2000 SP3 or later, Windows XP SP1 or later, or many common software packages) might not function after the profile is translated. When the ADMT Security Translation Wizard is translating local profiles in replace mode, it reverts to add mode if a profile is locked. This might result in a successful profile translation; however, application installations might not function after the profile is translated.

Before you start the local user profile translation, allow enough time for the workstations to restart after you move them to the target domain. Allow for the ADMT time delay factor (one minute) plus the time required for a restart cycle for your workstations.

| **Note** |
| --- |
| Translate the local user profiles the night before you notify the users to log on by using their new accounts in the target domain. Migrating profiles the night before ensures that the new user profile reflects the most current user settings. |

You can translate local user profiles by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To translate local user profiles by using the ADMT console**

1. For each workstation in the source domain that is running Windows NT 4.0, Windows 2000, or Windows XP, add the ADMT resource migration account to the local Administrators group.
2. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
3. Open the Active Directory Migration Tool, and then select **Security Translation Wizard**.
4. Complete the Security Translation Wizard by using the information in Table 11.23.

   **Table 11.23 Using the Security Translation Wizard to Translate Local User Profiles**

   | Wizard Page | Action |
   | --- | --- |
   | Test or Make Changes | Click **Migrate Now?** |
   | Security Translation Options | Click **Previously migrated objects**. |
   | Domain Selection | In the **Source domain** box, type or select the name of the source account domain.<br>In the **Target domain** box, type or select the name of the target domain. |
   | Translate Objects | Click **User Profiles**. |
   | Security Translation Options | Click **Replace**. |

5. When the wizard has finished running on all computers, click **View Log** and review the dispatch log for any errors.
6.

**To translate local user profiles by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
```

```
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TOT:REPLACE /TUP:YES
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" /O:"
option_file.txt
"
```

Table 11.24 lists the common parameters used for migrating user accounts, along with the command-line parameter and option file equivalents.

**Table 11.24 Common Parameters Used for Local User Profile Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Translate option = REPLACE | `/TOT:REPLACE` | `TranslateOption=REPLACE` |
| Modify local user profile security | `/TUP:YES` | `TranslateUserProfiles=YES` |

3. Review the results that are displayed in the dispatch log for any errors.

**To translate local user profiles by using a script**

- Prepare a script that incorporates ADMT commands and options for translating local user profiles.

  For a sample script to assist you in translating local user profiles, see "Translating Local User Profiles Between Forests" (DSSREER_7.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Local User Profiles Between Forests" on the Web at http://www.microsoft.com/reskit).

**To translate local user profiles by using a script**

**Migrating Workstations in Batches**

After you migrate a batch of local user profiles, migrate the corresponding batch of user workstations. When you move workstations between domains, the SAM database moves along with them. Accounts located in the local SAM database (such as local groups) that are used to enable access to resources always move with the computer, and therefore do not need to be migrated.

**Note**

- Use a low value in the ADMT computer restart option, to restart workstations immediately after joining them to the target domain, or as soon as possible thereafter. Resources that are not restarted following migration are in an indeterminate state.

You can migrate workstations by using the Active Directory Migration Tool console, by using the ADMT command-line option, or by using a script.

**To migrate workstations by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT resource migration account.
2. Open the Active Directory Migration Tool, and then select **Computer Account Migration Wizard**.
3. Complete the Computer Account Migration Wizard by using the information in Table 11.25.

**Table 11.25 Using the Computer Account Migration Wizard to Migrate Workstations**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br><br>In the **Target domain** box, type or select the name of the target domain. |
| Computer Selection | Click **Add**.<br><br>In the **Select Computers** dialog box, select the names of the workstations, click **Add**, and then click **OK**. |
| Organizational Unit Selection | Click **Browse**.<br><br>In the **Browse for Container** dialog box, locate the target domain **Computers** container, and then click **OK**. |
| Security Translation Options | Click the **Local groups** check box.<br><br>Click the **User rights** check box. |
| Translate Objects | Click **Add**. |
| Computer Options | In the **Minutes before computer restart after wizard completion** box, type **1**. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

4. When the wizard has finished running on all computers, click **View Log** and review the migration log and dispatch log for any errors.
5. If errors are reported, run the Retry Wizard to dispatch agents to workstations for which the migration process failed.
6. Open Active Directory Users and Computers and verify that the workstations exist in the appropriate OU in the target domain.

**To migrate workstations by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT installed, log on by using the ADMT resource migration account.
2. At the command line, type:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /SD:"
```

```
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /RDL:1
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /O "
option_file.txt
"
```

Table 11.26 lists the common parameters used for workstation migration, along with the command-line parameter and option file equivalents.

**Table 11.26 Common Parameters Used for Workstation Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Source OU location | `/SO:"`*source_OU*`"` | `SourceOU="`*source_OU*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| • Restart computer one minute after wizard completes | `/RDL:1` | `RestartDelay=1` |
| • Do not rename accounts | `/RO:DONT` *(default)* | `RenameOption=DONT` |
| • Ignore conflicting accounts and do not migrate them | `/CO:IGNORE` *(default)* | `ConflictOptions=IGNORE` |
| • Translate Option | `/TOT:ADD` | `TranslateOption=YES` |
| • Translate User Rights | `/TUR:YES` | `TranslateUserRights=YES` |
| • Translate Local Groups | `/TLG:YES` | `TranslateLocalGroups=YES` |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers and locate the target OU. Verify that the workstations exist in the target OU.

**To migrate workstations by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating workstations.

  For a sample script file to assist you in migrating workstations between forests, see "Migrating Workstations Between Forests" (DSSREER_9.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations Between Forests" on the Web at http://www.microsoft.com/reskit).

**Remigrating User Accounts in Batches**

After you have verified the success of local user profile and user workstation migration for the user batch, migrate the user accounts for that batch. You can migrate user accounts in batches by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To remigrate the current batch of user accounts by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool, and then select **User Account Migration Wizard**.
3. Complete the User Account Migration Wizard by using the information in Table 11.27.

**Table 11.27 Using the User Account Migration Wizard to Migrate User Accounts**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br><br>In the **Target domain** box, type or select the name of the target domain. |
| User Selection | Click **Add**.<br><br>In the **Select Users** dialog box, select the user accounts that you want to migrate, and then click **Add**. By default, the wizard migrates the accounts to the Users container. If you need to change the OU, click **Location**, and then find the OU that contains the accounts.<br><br>Click **Migrate Passwords**.<br><br>Click **OK**. |
| Organizational Unit Selection | Ensure that ADMT lists the correct target OU. If it is not correct, type the correct OU or click **Browse**.<br><br>In the **Browse for Container** dialog box, locate the target domain **Users** container or appropriate OU, and then click **OK**. |
| Password Options | Click **Migrate Passwords**.<br><br>Select **Password migration source DC** from the drop-down list box. |
| Account Transition Options | Click **Enable target accounts**.<br><br>Click **Days until source account expires,** and then type **7** in the text box.<br><br>Click the **Migrate user SIDs to target domains** check box. |
| User Account | Type the user name, password, and domain of a user account that has administrative credentials. |
| User Options | Click the **Translate roaming profiles** check box.<br><br>Click the **Update user rights** check box.<br><br>Click the **Migrate associated user groups** check box.<br><br>Click **Fix users' group memberships**.<br><br>Click **Do not rename accounts**. |
| Object Property Exclusion | Clear the **Exclude specific object properties from migration** check box. |

4. When the wizard finishes, click **View Log**, and review the migration log for any errors.
5. Open Active Directory Users and Computers and verify that the user accounts exist in the appropriate OU in the target domain.

**To remigrate the current batch of users by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.
2. At the command line, type:

```
ADMT USER /N "user_name1""user_name2" /SD:"source_domain"
```

```
/TD:"target_domain" /TO:"target_OU"  [parameters]
```

You can append parameters to the command as follows:

```
                ADMT USER /N "user_name1""user_name2" /SD:"source_domain"
/TD:"target_domain"  /TO:"target_OU" /MSS:YES TRP:YES /UUR:YES
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
                ADMT USER /N "user_name1""user_name2" /O "option_file.txt"
```

Table 11.28 lists the common parameters used for migrating user accounts, along with the command-line parameter and option file equivalents.

**Table 11.28 Common Parameters Used for User Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Source OU location | `/SO:"`*source_OU*`"` | `SourceOU="`*source_OU*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| Migrate SIDs | `/MSS:YES` | `MigrateSIDs=YES` |
| Do not rename accts | `/RO:DONT` *(default)* | `RenameOption=DONT` |
| Ignore conflicting accts and not migrate them | `/CO: REPLACE+REMOVEUSERRIGHTS+MOVEREPLACEDACCOUNTS` | `ConflictOptions=REPLACE+REMOVEUSERRIGHTS+MOVEREPLACEDACCOUNTS` |
| Translate Roaming Profile | `/TRP:YES` *(default)* | `TranslateRoamingProfile=YES` |
| Update User Rights | `/UUR:YES` | `UpdateUserRights=YES` |
| Password Options | `/PO:COPY /PS:`<name of PES server> | `PasswordOption=copy`<br>`PasswordServer=`<name of server> |
| Expire accounts | `/SEP:7` | `SourceExpiration=7` |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers and locate the target OU. Verify that the users exist in the target OU.

**To remigrate the current batch of user accounts by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating users by using the sample script shown in Listing 11.12.

**Listing 11.12 Remigrating User Accounts in Batches Between Forests**

```
 <Job id=" RemigratingUserAccountsInBatchesBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit
```

```
    Dim objMigration
    Dim objUserMigration

    '
    'Create instance of ADMT migration objects.
    '

    Set objMigration = CreateObject(" ADMT.Migration" )
    Set objUserMigration = objMigration.CreateUserMigration

    '
    'Specify general migration options.
    '

    objMigration.SourceDomain = " source domain"
    objMigration.SourceOu = " source container"
    objMigration.TargetDomain = " target domain"
    objMigration.TargetOu = " target container"
    objMigration.PasswordOption = admtCopyPassword
    objMigration.PasswordServer = " password export server name"
    objMigration.ConflictOptions = admtReplaceConflicting

    '
    'Specify user migration specific options.
    '
    objUserMigration.SourceExpiration = 7
    objUserMigration.MigrateSIDs = True
    objUserMigration.TranslateRoamingProfile = True
    objUserMigration.UpdateUserRights = True
    objUserMigration.FixGroupMembership = True
    objUserMigration.MigrateServiceAccounts = False

    '
    'Migrate specified user objects.
    '

    objUserMigration.Migrate admtData, Array(" user name1" ," user name2" )

    Set objUserMigration = Nothing
    Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in creating a script to migrate user accounts, see "Remigrating User Accounts in Batches Between Forests" (DSSREER_12.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Remigrating User Accounts in Batches Between Forests" on the Web at http://www.microsoft.com/reskit).

**Remigrating All Global Groups Following User Account Migration**

A large user account migration might take place over an extended period of time. For this reason, you might need to remigrate global groups from the source to the target domain after you migrate each batch of users, to reflect changes made to the membership of groups in the source domain after the initial global group migration occurred. For more information about and procedures for remigrating global groups, see "Remigrating All Global Groups After All Batches Are Migrated" later in this chapter.

# Remigration of All Global Groups After All Batches Are Migrated

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remigrating All Global Groups After All Batches Are Migrated**

After all batches have been migrated, perform a final global group remigration to ensure that any late changes made to global group membership in the source domain are reflected in the target domain. You can remigrate global groups by using the Active Directory Migration Tool console, by using the ADMT command-line option, or by using a script.

**To remigrate global groups by using the ADMT console**

1. On the domain controller in the target domain on which ADMT installed, log on by using the ADMT account migration account.
2. Open the Active Directory Migration Tool, and then select **Group Account Migration Wizard**.
3. Complete the Group Account Migration Wizard by using the information in Table 11.29.

**Table 11.29 Using the Group Account Migration Wizard to Remigrate Global Groups**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS or DNS name of the source domain or select the name from a list.<br><br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain.<br><br>If ADMT includes the names of the source and target domains, ensure that they are correct. |
| Group Selection | Click **Add**.<br><br>In the **Select Groups** dialog box, select all global groups, click **Add**, and then click **OK**. |
| Organizational Unit Selection | Type the name of the OU or click **Browse**.<br><br>In the **Browse for Container** dialog box, locate the container in the target domain you want to move the global groups into, and then click **OK**. |
| Group Options | Click **Update user rights**.<br><br>Ensure that **Copy group members** is not selected**.**<br><br>Ensure that **Update previously migrated objects** is not selected.<br><br>Click **Fix membership of group**.<br><br>Click **Migrate Group SIDs to target domain**.<br><br>Click **Do not rename accounts** |
| User Account | Type the user name, password, and domain of an account that has administrative rights in the source domain. |
| Object Property Exclusion | Clear **Exclude specific object properties from migration**. |
| Naming Conflicts | Click **Replace conflicting accounts**. Clear all other options. |

4. When the wizard has finished running, click **View Log**, and review the migration log for any errors.

5. Open Active Directory Users and Computers and locate the target OU. Verify that the global groups exist in the target domain OU.

**To remigrate global groups by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT is installed, log on by using the ADMT account migration account.

2. At the command line, type:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /MSS:YES
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /O:"
option_file.txt
"
```

Table 11.30 lists the common parameters used for migrating global groups, along with the command-line parameter and option file equivalents.

**Table 11.30 Common Parameters Used for Global Group Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| Migrate GG SIDs | `/MSS:YES` | `MigrateSIDs=YES` |

| | | |
|---|---|---|
| Do not rename accts | `/RO:DONT` *(default)* | `RenameOption=DONT` |
| Ignore conflicting accts and do not migrate them | `/CO:REPLACE` | `ConflictOptions=REPLACE` |

3. Review the results that are displayed on the screen for any errors.

4. Open Active Directory Users and Computers and locate the target OU. Verify that the global groups exist in the target domain OU.

**To remigrate global groups by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating global groups.

  For a script file to assist you in creating a script to remigrate global groups, see "Remigrating Global Groups Between Forests" (DSSREER_10.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Remigrating Global Groups Between Forests" on the Web at http://www.microsoft.com/reskit).

# Translating Security in Remove Mode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Translate security on objects to remove the SIDs of the accounts in the source domain from the ACLs of the migrated objects. Do this only after all of the source accounts are disabled. Run the Security Translation Wizard on all files, shared folders, printers, and local groups, and at least one domain controller (to translate security on shared local groups).

You can translate security in remove mode on objects by using the ADMT console or by using a script.

**To translate security in remove mode on objects by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT account migration account.

2. Open the Active Directory Migration Tool, and then select **Security Translation Wizard**.

3. Complete the Security Translation Wizard by using the information in Table 11.31.

**Table 11.31 Using the ADMT Security Translation Wizard in Remove Mode**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Security Translation Options | Click **Previously migrated objects**. |
| Domain Selection | In the **Source domain** box, type or select the name of the source account domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| Translate Objects | Clear the **User Profiles** check box.<br>Select all the other check boxes. |
| Security Translation Options | Click **Remove**. |

**To translate security in remove mode on objects by using a script**

- Prepare a script that incorporates ADMT commands and options for translating security in remove mode on objects by using the sample script shown in Listing 11.13.

**Listing 11.13 Translating Security in Remove Mode on Objects Between Forests**

```
<Job id=" TranslatingSecurityInRemoveModeOnObjectsBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objSecurityTranslation

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objSecurityTranslation = objMigration.CreateSecurityTranslation
```

```
    '
    'Specify general migration options.
    '

    objMigration.SourceDomain = " source domain"
    objMigration.TargetDomain = " target domain"
    objMigration.TargetOu = " Computers"

    '
    'Specify security translation specific options.
    '

    objSecurityTranslation.TranslationOption = admtTranslateRemove
    objSecurityTranslation.TranslateFilesAndFolders = True
    objSecurityTranslation.TranslateLocalGroups = True
    objSecurityTranslation.TranslatePrinters = True
    objSecurityTranslation.TranslateRegistry = True
    objSecurityTranslation.TranslateShares = True
    objSecurityTranslation.TranslateUserProfiles = False
    objSecurityTranslation.TranslateUserRights = True

    '
    'Perform security translation on specified computer objects.
    '

    objSecurityTranslation.Translate admtData, _
   Array(" computer name1" ," computer name2" )

    Set objSecurityTranslation = Nothing
    Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in creating a script to translate security in remove mode on objects, see "Translating Security in Remove Mode on Objects Between Forests" (DSSREER_13.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Security in Remove Mode on Objects Between Forests" on the Web at http://www.microsoft.com/reskit).

# Migrating Resources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process of migrating resources between Active Directory domains in different forests involves completing the migration of workstation accounts, member servers, and domain controllers. When you have successfully migrated all resource objects to the target domain, you can decommission the source domain.

Figure 11.11 shows the process for migrating resource objects between Active Directory domains in different forests.

**Figure 11.11 Migrating Resources Between Active Directory Domains in Different Forests**



# Migration of Workstations and Member Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Migrate remaining workstations that you did not migrate during the user account migration process, along with member servers, in small batches of up to 100 computers. Workstation account and member server migration is a straightforward process. Workstations and member servers have their own SAM account database. When you move workstations and member servers between domains, the database moves along with them. Accounts located in the local SAM database (such as local groups) that are used to enable access to resources always move with the computer, and therefore do not need to be migrated. Because the migration requires that workstations and member servers restart, it is important to schedule the migration for a time when the server is not servicing requests.

**Note**

- Use a low value in the ADMT computer restart option, to restart member servers immediately after joining them to the target domain, or as soon as possible thereafter. Resources that are not restarted following migration are in an indeterminate state.

You can migrate member servers by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate workstations and member servers by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT resource migration account.
2. Open the Active Directory Migration Tool, and then select **Computer Account Migration Wizard**.
3. Complete the Computer Account Migration Wizard by using the information in Table 11.32.

**Table 11.32 Using the Computer Account Migration Wizard to Migrate Workstations and Member Servers**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| Computer Selection | Click **Add**.<br>In the **Select Computers** dialog box, select the names of all of the workstations and member servers in the source domain that you want to migrate, click **Add**, and then click **OK**. |
| Organizational Unit Selection | Click **Browse**.<br>In the **Browsefor Container** dialog box, locate the target domain **Computers** container, or the appropriate OU, and then click **OK**. |
| Security Translation Options | Click the **Local groups** check box.<br>Click the **User rights** check box. |
| Translate Objects | Click **Add**. |
| Computer Options | In the **Minutes before computer restart after wizard completion** box, type **1**. |

| | |
|---|---|
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

4. When the wizard has finished running on all computers, click **View Log** and review the migration log and dispatch log for any errors.
5. If errors are reported, run the Retry Wizard to dispatch agents to workstations and member servers for which the migration process failed.
6. Open Active Directory Users and Computers and verify that the workstations and member servers exist in the appropriate OU in the target domain.

**To migrate workstations and member servers by using the ADMT command-line option**

1. On the domain controller in the target domain on which ADMT installed, log on by using the ADMT resource migration account.
2. At the command line, type:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
" /RDL:1
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /O "
option_file.txt
"
```

Table 11.33 lists the common parameters used for workstation and member server migration, along with the command-line parameter and option file equivalents.

**Table 11.33 Common Parameters Used for Workstation and Member Server Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| • Restart computer one minute after wizard completes | `/RDL:1` | `RestartDelay=1` |
| • Do not rename accounts | `/RO:DONT` *(default)* | `RenameOption=DONT` |
| • Ignore conflicting accounts and do not migrate them | `/CO:IGNORE` *(default)* | `ConflictOptions=IGNORE` |
| • Translate Option | `/TOT:ADD` | `TranslateOption=YES` |
| • Translate User Rights | `/TUR:YES` | `TranslateUserRights=YES` |
| • Translate Local Groups | `/TLG:YES` | `TranslateLocalGroups=YES` |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers and locate the target OU. Verify that the workstations and member servers exist in the target OU.

**To migrate workstations and member servers by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating workstations and member servers by using the sample script shown in Listing 11.14.

**Listing 11.14 Migrating Workstations or Member Servers Between Forests**

```
<Job id=" MigratingWorkstationsMemberServersBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objComputerMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
```

```
   Set objComputerMigration = objMigration.CreateComputerMigration

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " Computers"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " Computers"

   '
   'Specify computer migration specific options.
   '
   objComputerMigration.TranslationOption = admtTranslateAdd
   objComputerMigration.TranslateLocalGroups = True
   objComputerMigration.TranslateUserRights = True
   objComputerMigration.RestartDelay = 1

   '
   'Migrate computer objects on specified computer objects.
   '

   objComputerMigration.Migrate admtData, _
  Array(" computer name1" ," computer name2" )

   Set objComputerMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a script file to assist you in creating a script to migrate member servers, see "Migrating Workstations or Member Servers Between Forests" (DSSREER_14.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations or Member Servers Between Forests" on the Web at http://www.microsoft.com/reskit).

# Migrating Domain and Shared Local Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Shared local groups are local groups in Windows NT 4.0 and Active Directory domains that can be used in ACLs on domain controllers. When a domain is configured to operate either in Windows 2000 native mode or at the Windows Server 2003 domain functional level, shared local groups are automatically changed to domain local groups. These groups can then be used in ACLs on member servers and workstations. If either domain local groups or shared local groups are used in ACLs on either domain controllers or member servers, then they need to be migrated to the target domain before the server is migrated.

Note that it is not necessary to change any ACLs as part of the migration process. The ACLs continue to reference the domain local groups or shared local groups in the source domain. Because the domain local groups or shared local groups can be migrated to the target domain while using SID history, users maintain access to the resources. ADMT retains the membership of the local group during the migration.

You can migrate domain or shared local groups by using the ADMT console or by using a script.

**To migrate domain and shared local groups by using the ADMT console**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT resource migration account.

2. Open the Active Directory Migration Tool, and then select **Group Account Migration Wizard**.

3. Complete the Group Account Migration Wizard by using the information in Table 11.34.

   **Table 11.34 Using the Group Account Migration Wizard to Migrate Domain and Shared Local Groups**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS or DNS name of the source domain or select the name from a list.<br><br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain. |
| Group Selection | Click **Add**.<br><br>In the **Select Groups** dialog box, select all domain and shared local groups that you need to migrate (except built-in groups), click **Add**, and then click **OK**. |
| Organizational Unit Selection | Type the name of the OU or click **Browse**.<br><br>In the **Browse for Container** dialog box, locate the container in the target domain to which the local groups are migrating, and then click **OK**. |
| Group Options | Click **Migrate Group SIDs to target domain**.<br><br>Click **Do not rename accounts**<br><br>Ensure that all other options are not selected. |
| User Account | Type the user name, password, and domainof an account that has administrative rights in the source domain. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

4. When the wizard has finished running, click **View Log**. Review the migration log for any errors.

5. Open Active Directory Users and Computers, locate the target OU, and then verify that the shared local groups exist in

the target domain OU.

**To migrate domain and shared local groups by using a script**

- Prepare a script that incorporates ADMT commands and options for migrating domain and shared local groups by using the sample script shown in Listing 11.15.

**Listing 11.15 Migrating Domain and Shared Local Groups Between Forests**

```
<Job id=" MigratingDomainAndSharedLocalGroupsBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objGroupMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objGroupMigration = objMigration.CreateGroupMigration

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " source container"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"

   '
   'Specify group migration specific options.
   '

   objGroupMigration.MigrateSIDs = True

   '
   'Migrate specified group objects.
   '

   objGroupMigration.Migrate admtData, _
  Array(" local group name1" ," local group name2" )

   Set objGroupMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a script file to assist you in creating a script to migrate domain and shared local groups, see "Migrating Domain and Shared Local Groups Between Forests" (DSSREER_15.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Domain and Shared Local Groups Between Forests" on the Web at http://www.microsoft.com/reskit).

# Migration of Domain Controllers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Active Directory, domain controllers can be migrated between domains. To do this, you must remove Active Directory from the domain controller, migrate it as a member server to the target domain, and then reinstall Active Directory. Note that if the server is running Windows 2000, you cannot install Active Directory in the target domain if the target domain is already at the Windows Server 2003 functional level. In this case, you must upgrade the server to Windows Server 2003 before installing Active Directory.

# Completing the Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you migrate all accounts and resources from the source to the target domain, complete the restructuring process by doing the following:

- Transferring the administration of user and group accounts from the source to the target domain.
- Ensuring that at least two domain controllers continue to operate in the source domain until the resource migration process is complete.
- Backing up the two domain controllers in the source domain.

Figure 11.12 shows the process for completing the migration of Active Directory domains between forests.

**Figure 11.12 Completing the Migration of Active Directory Domains Between Forests**



# Translating Security on Your Member Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Translate security on member servers to clean up the ACLs of the resources. After objects are migrated to the target domain, resources contain the ACL entries of the source domain objects. If you are using SID history to provide access to resources during the migration, the SIDs from the source domain remain in the ACLs to enable users to access resources while the migration is in progress. After the migration is complete, the SIDs from the source domain are no longer needed. Use the Security Translation Wizard in ADMT to replace the source domain SIDs with the target domain SIDs.

You do not need to perform this procedure if you are not using SID history for resource access because you should have already run security translation in remove mode after the user migration.

**To translate security on member servers**

1. On the domain controller in the target domain on which you installed ADMT, log on by using the ADMT account migration account.

2. Open the Active Directory Migration Tool, and then select **Security Translation Wizard**.

3. Complete the Security Translation Wizard by using the information in Table 11.35.

**Table 11.35 Using the ADMT Security Translation Wizard**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Security Translation Options | Click **Previously migrated objects**. |
| Domain Selection | In the **Source domain** box, type or select the name of the source account domain.<br><br>In the **Target domain** box, type or select the name of the target domain. |
| Translate Objects | Clear the **User Profiles** check box.<br><br>Select all other options. |
| Security Translation Options | Click **Replace**. |

**To translate security on member servers by using a script**

- Prepare a script that incorporates ADMT commands and options for translating security on member servers by using the sample script shown in Listing 11.16.

**Listing 11.16 Translating Security on Member Servers Between Forests**

```
<Job id=" TranslatingSecurityOnMemberServersBetweenForests" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objSecurityTranslation

   '
   'Create instance of ADMT migration objects.
   '
```

```
   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objSecurityTranslation = objMigration.CreateSecurityTranslation

   '
   'Specify general migration options.
   '

   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " Computers"

   '
   'Specify security translation specific options.
   '

   objSecurityTranslation.TranslationOption = admtTranslateReplace
   objSecurityTranslation.TranslateFilesAndFolders = True
   objSecurityTranslation.TranslateLocalGroups = True
   objSecurityTranslation.TranslatePrinters = True
   objSecurityTranslation.TranslateRegistry = True
   objSecurityTranslation.TranslateShares = True
   objSecurityTranslation.TranslateUserProfiles = False
   objSecurityTranslation.TranslateUserRights = True

   '
   'Perform security translation on specified computer objects.
   '

   objSecurityTranslation.Translate admtData, _
  Array(" computer name1" ," computer name2" )

   Set objSecurityTranslation = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in creating a script to translate security on member servers, see "Translating Security on Member Servers Between Forests" (DSSREER_16.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Security on Member Servers Between Forests" on the Web at http://www.microsoft.com/reskit).

# Decommissioning the Source Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you complete the migration of the accounts and resources in your source domain, decommission the source domain. Ensure that you retain a full system state backup of a domains controller so that you can bring the domain back online at any time.

To decommission the source domain:

1. Remove all trust relationships between the source domain and the target domain.

2. Repurpose any remaining domain controllers in the source domain that you did not migrate to the target domain.

3. Disable all accounts that you created during the migration process, including those accounts to which you assigned administrative permissions.

**Note**

- When you decommission the source domain, shared local groups and local groups that you have not translated by using the Security Translation Wizard display group members as "account unknown" because member names from the source domain do not resolve. Those group memberships still exist, however, and this does not affect users. Do not delete "account unknown" entries because this disables the access facilitated by SID history. Run the Security Translation Wizard to remove these entries.

# Additional Resources for Retructuring Active Directory Domains Between Forests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing the Active Directory Logical Structure" in this book for more information about designing the Active Directory logical structure and designing DNS for Active Directory.
- "Designing the Site Topology" in this book for more information about the Active Directory site topology.
- "Deploying DNS" in *Deploying Network Services* for more information about deploying DNS.
- "Enabling Advanced Windows Server 2003 Active Directory Features" in this book for more information about functional levels.

**Related Tools**

- ADMT

  You can use ADMT version 2 to migrate accounts to restructure your Windows Server 2003 domains. ADMT includes wizards that you can use to migrate objects such as groups, users, computers, and service accounts. These wizards assist you in performing the appropriate migration for each object type.

**Related Job Aids**

- "User and Group Object Assignment Table" (DSSREER_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "User and Group Object Assignment Table" on the Web at http://www.microsoft.com/reskit).
- "Resource Object Assignment Table" (DSSREER_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Resource Object Assignment Table" on the Web at http://www.microsoft.com/reskit).
- "Migration Test Matrix" (DSSREER_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migration Test Matrix" on the Web at http://www.microsoft.com/reskit).
- "ADMT Option File Reference" (DSSRENT_3.txt) on the *Windows Server 2003 Deployment Kit* companion CD (or see "ADMT Option File Reference" on the Web at http://www.microsoft.com/reskit).
- "ADMT Scripting Constants" (AdmtConstants.vbs) on the *Windows Server 2003 Deployment Kit* companion CD (or see "ADMT Scripting Constants" on the Web at http://www.microsoft.com/reskit).
- "Identifying Service Accounts" (DSSRENT_5.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Identifying Service Accounts" on the Web at http://www.microsoft.com/reskit).
- "Transitioning Service Accounts Between Forests" (DSSREER_4.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Transitioning Service Accounts Between Forests" on the Web at http://www.microsoft.com/reskit).
- "Migrating Global Groups Between Forests" (DSSREER_5.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Global Groups Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Migrating All User Accounts Between Forests" (DSSREER_6.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating All User Accounts Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Translating Local User Profiles Between Forests" (DSSREER_7.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Local User Profiles Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Migrating User Accounts in Batches Between Forests" (DSSREER_8.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating User Accounts in Batches Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Migrating Workstations Between Forests" (DSSREER_9.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Remigrating Global Groups Between Forests" (DSSREER_10.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Remigrating Global Groups Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Translating Security in Add Mode on Objects Between Forests" (DSSREER_11.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Security in Add Mode on Objects Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Remigrating User Accounts in Batches Between Forests" (DSSREER_12.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Remigrating User Accounts in Batches Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Translating Security in Remove Mode on Objects Between Forests" (DSSREER_13.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Security in Remove Mode on Objects Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Migrating Workstations or Member Servers Between Forests" (DSSREER_14.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations or Member Servers Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Migrating Domain and Shared Local Groups Between Forests" (DSSREER_15.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Domain and Shared Local Groups Between Forests" on the Web at http://www.microsoft.com/reskit).

- "Translating Security on Member Servers Between Forests" (DSSREER_16.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Security on Member Servers Between Forests" on the Web at http://www.microsoft.com/reskit).

# Restructuring Active Directory Domains Within a Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Active Directory® directory service domains within a forest with the goal of reducing the number of domains allows you to decrease the administrative overhead for your organization by reducing replication traffic, reducing the amount of user and group administration that is required, and simplifying the administration of Group Policy. If users change positions or move between locations, you might also move objects between forests on a regular basis. The process for restructuring Active Directory domains within a forest differs from the process of restructuring Active Directory domains between forests and requires careful planning and testing.

**In This Chapter**

Overview of Restructuring Active Directory Domains Within a Forest

Preparing to Restructure Active Directory Domains Within a Forest

Migrating Domain Objects Between Active Directory Domains

Completing Post-Migration Tasks

Additional Resources for Restructuring Active Directory Domains Within a Forest

**Related Information**

- For more information about designing the Active Directory logical structure, see "Designing the Active Directory Logical Structure" in this book.
- For more information about designing the Active Directory site topology, see "Designing the Site Topology" in this book.
- For more information about creating an Active Directory forest, see "Deploying the Windows Server 2003 Forest Root Domain" in this book.

# Overview of Restructuring Active Directory Domains Within a Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The most efficient Active Directory design includes the smallest possible number of domains. By minimizing the number of domains running the Microsoft® Windows® Server 2003, Standard Edition; Microsoft® Windows® Server 2003, Enterprise Edition; and Microsoft® Windows® Server 2003, Datacenter Edition operating systems, you can reduce administrative costs and increase the efficiency of your organization.

You might need to restructure domains within your forest if, for example, your organization closes a regional office location, and the regional domain for that location is no longer needed. You might also restructure domains within your forest if you upgrade your network infrastructure, thereby increasing your bandwidth capacity and your replication capacity and enabling you to reduce the size of your Active Directory logical structure.

The process of restructuring Active Directory domains within a forest is similar to the process of migrating accounts between domains. When you migrate accounts and resources between domains, you migrate objects from the source domain to the target domain without decommissioning the source domain. When you restructure Active Directory domains, you eliminate the source domain from the forest after you complete the migration of all domain objects.

Before you begin the process for restructuring Active Directory domains within a forest, ensure that the source and target domains are operating at the Microsoft® Windows® 2000 native or Windows Server 2003 domain functional level. Restructuring source domains that are operating at the Windows 2000 mixed domain functional level, which can include domain controllers that are running Microsoft® Windows NT® 4.0, is not recommended.

After you complete the process for restructuring Active Directory domains within a forest, you can reduce overhead and simplify administration in your organization.

**Note**

- For a list of the job aids that are available to assist you in restructuring Active Directory domains within a forest, see "Additional Resources for Restructuring Active Directory Domains Within a Forest" later in this chapter.

# Process for Restructuring Active Directory Domains Within a Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Active Directory domains within a forest requires careful planning and preparation. To maintain user access to resources during the process of restructuring domains, you must perform the migration steps in a specific order. Figure 12.1 shows the process for restructuring Active Directory domains within a forest.

**Figure 12.1 Restructuring Active Directory Domains Within a Forest**



# Background Information for Restructuring Active Directory Domains Within a Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Active Directory domains within a forest involves migrating accounts and resources from the source domains to the target domains. Unlike the process for restructuring Active Directory domains between forests, when you restructure domains within a forest, the migrated objects no longer exist in the source domain. Additionally, migrating user accounts, resources, and groups requires special consideration when you restructure Active Directory domains within a forest because of the containment rules that apply to Active Directory groups. For these reasons, the challenge when you restructure Active Directory domains within a forest is to ensure that users have continuous access to resources during the migration process.

**Active Directory Migration Tool**
The Active Directory Migration Tool (ADMT) version 2 allows you to migrate objects within Active Directory forests. ADMT includes wizards that automate migration tasks such as migrating users, groups, service accounts, computers, and trusts, and performing security translation.

ADMT is available in the Admigration.msi file in the \i386\admt directory on the Windows Server 2003 operating system CD.

You can perform ADMT tasks by using the ADMT console, a command line, or a script. When running ADMT at the command line, it is often more efficient to use an option file to specify command-line options. You can use the ADMT option file reference shown in Listing 12.1 to assist you in creating option files. Examples of command-line syntax are provided for each task that you must perform to restructure the domains within the forest.

Listing 12.1 lists common options that apply to several migration tasks. Each type of migration task has a section that lists options that are specific to that task. The section name corresponds to the task name when you run ADMT at the command line. Items can be commented out with a semi-colon. In Listing 12.1, the default values are commented out.

When running ADMT at the command line, you do not need to include an option in your command if you want to accept the default value. In this chapter, however, tables are provided for reference. The tables list the command-line equivalent of each option that is shown in the corresponding ADMT console procedure, including those where you are accepting the default value.

**Listing 12.1 ADMT Option File Reference**

```
 [Migration]
;TestMigration=No
;IntraForest=No
SourceDomain=" source_domain_name"
;SourceOu=" source_ou_path"

TargetDomain=" target_domain_name"
TargetOu=" target_ou_path"
;RenameOption=Dont
;RenamePrefixOrSuffix=""
;PasswordOption=Complex
;PasswordServer=""
;PasswordFile=""
;ConflictOptions=Ignore
;ConflictPrefixOrSuffix=""
;UserPropertiesToExclude=""
;InetOrgPersonPropertiesToExclude=""
;GroupPropertiesToExclude=""
;ComputerPropertiesToExclude=""

[User]
;DisableOption=EnableTarget
;SourceExpiration=None
MigrateSIDs=Yes
;TranslateRoamingProfile=No
;UpdateUserRights=No
;MigrateGroups=No
;UpdatePreviouslyMigratedObjects=No
;FixGroupMembership=Yes
;MigrateServiceAccounts=No
```

```
;UpdateGroupRights=No

[Group]
MigrateSIDs=Yes
;UpdatePreviouslyMigratedObjects=No
;FixGroupMembership=Yes
;UpdateGroupRights=No
;MigrateMembers=No
;DisableOption=EnableTarget
;SourceExpiration=None
;TranslateRoamingProfile=No
;MigrateServiceAccounts=No

[Security]
TranslationOption=Add
;TranslateFilesAndFolders=No
;TranslateLocalGroups=No
;TranslatePrinters=No
;TranslateRegistry=No
;TranslateShares=No
;TranslateUserProfiles=No
;TranslateUserRights=No
;SidMappingFile=" SidMappingFile.txt"
```

For a file that includes common options that are available for ADMT, see "ADMT Option File Reference" (DSSRENT_3.txt) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "ADMT Option File Reference" on the Web at http://www.microsoft.com/reskit).

When running ADMT from the command line, it is often more efficient to use an include file to list the users, groups, or computers that will be migrated. For example, to migrate a small number of computers, you might type each computer name at the command line, using the /N option as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /IF:YES /SD:"
source_domain
" /TD:"
target_domain
" /TO:"
target_OU
"
```

*Computer_name1* and *computer_name2* are the names of computers in the source domain that you are migrating in this batch.

If you have a large number of users, groups, or computers to migrate, you can list them in an include file. For example, to create an include file for a batch of computers, create a plain text file and list the computer names, each name on a separate line. Then specify the include file name with the /F option, as follows:

```
ADMT COMPUTER /F "includefile_name" /IF:YES /SD:"source_domain"
/TD:"target_domain" /TO:"target_OU"
```

ADMT supports the use of American National Standards Institute (ANSI) or standard Unicode files for option files and include files.

To specify the names of users, groups, or computers, use one of the following conventions:

- The Security Accounts Manager (SAM) account name. Note that to specify a computer name in this format, you must append a $ to the computer name. For example, to specify a computer with the name Workstation01, use Workstation01$.

- The Windows NT 4.0 account name. This includes the domain name and the SAM account name. For example, the Windows NT 4.0 account name of a computer named Workstation01 that is in the Asia domain is Asia\ Workstation01$.

- The relative distinguished name (also known as RDN). For example, cn= Workstation01.

- The canonical name. This can be specified as *DNS domain name/ou_path/object_name*, or *ou_path/object_name*; for example, Asia.trccorp.treyresearch.net/Computers/ Workstation01 or Computers/Workstation01.

For a complete list of all available ADMT options and syntax, run ADMT, and then select **Help** on the menu.

The sample scripts provided in this chapter reference the symbolic constants that are defined in the AdmtConstants.vbs file. Listing 12.2 shows the ADMT constants VBScript file.

**Listing 12.2 ADMT Scripting Constants**

```
Option Explicit

'-----------------------------------------------------------------------------
' ADMT Scripting Constants
'-----------------------------------------------------------------------------

' RenameOption constants

Const admtDoNotRename      = 0
Const admtRenameWithPrefix = 1
Const admtRenameWithSuffix = 2

' PasswordOption constants

Const admtPasswordFromName = 0
Const admtComplexPassword  = 1
Const admtCopyPassword     = 2

' ConflictOptions constants

Const admtIgnoreConflicting           = &H0000
Const admtReplaceConflicting          = &H0001
Const admtRenameConflictingWithPrefix = &H0002
Const admtRenameConflictingWithSuffix = &H0003
Const admtRemoveExistingUserRights    = &H0010
Const admtRemoveExistingMembers       = &H0020
Const admtMoveReplacedAccounts        = &H0040

' DisableOption constants

Const admtEnableTarget       = 0
Const admtDisableSource      = 1
Const admtDisableTarget      = 2
Const admtTargetSameAsSource = 4

' SourceExpiration constant

Const admtNoExpiration = -1

' Translation Option

Const admtTranslateReplace = 0
Const admtTranslateAdd     = 1
Const admtTranslateRemove  = 2

' Report Type

Const admtReportMigratedAccounts  = 0
Const admtReportMigratedComputers = 1
Const admtReportExpiredComputers  = 2
Const admtReportAccountReferences = 3
Const admtReportNameConflicts     = 4

' Option constants

Const admtNone     = 0
```

```
Const admtData      = 1
Const admtFile      = 2
Const admtDomain    = 3
Const admtRecurse           = &H0100
Const admtFlattenHierarchy  = &H0000
Const admtMaintainHierarchy = &H0200
```

For a script file that includes the constants that are available for ADMT scripts, see "ADMT Scripting Constants" (AdmtConstants.vbs) on the *Windows Server 2003 Deployment Kit* companion CD (or see "ADMT Scripting Constants" on the Web at http://www.microsoft.com/reskit). The constants are also provided on the Windows Server 2003 operating system CD, in the TemplateScript.vbs file in the \i386\admt directory.

**Closed Sets and Open Sets**

When you restructure Active Directory domains within a forest, you must be concerned with two types of closed sets. The first type is users and groups, and the second type is resources and local groups.

The first type of closed set includes user accounts, all the global groups to which the users belong, and all the other members of the global groups. Because global groups are limited to members of the domain where the global group exists, if you migrate a user account to a new domain but do not migrate the global groups to which the user belongs, the user is no longer a valid member of those global groups and cannot access resources that are based on membership in those global groups. Therefore, when you are moving accounts between domains in a forest, it is necessary to move closed sets so that users retain access to their resources.

Built-in accounts (such as Administrators, Users, and Power Users) and well-known accounts (such as Domain Admins and Domain Users) cannot be ADMT migration objects. However, migrating these groups in closed sets is not a common problem because using them in access control lists (ACLs) or membership in domain local groups is not an effective way to assign resource permissions.

When you migrate users, ADMT makes the user a member of the domain users group in the target domain but does not maintain permissions for other built-in groups, such as Server Operators, Backup Operators, or well-known groups, such as Domain Admins. If you have used built-in or well-known groups to assign resource permissions, you must reassign those permissions to a new domain local group before you begin the migration. Reassigning permissions includes the following steps:

1. Create a new domain local group in the source domain.

2. Create a new global group in the source domain that contains users who need access to the resource.

3. Add the new global group to the domain local group.

4. Run security translation by using a SIDmapping file that maps the well-known group to the new domain local group (created in the first step) on all resources that assign permissions using well-known groups. For information about performing a security translation by using a SIDmapping file, see "Translate Security by Using a SID mapping File" later in this chapter.

In small domain environments that have few global groups, you might be able to identify closed sets of users and groups. If you can identify closed sets, you can migrate users and groups at the same time. In a large domain environment, a user can belong to a number of global groups; therefore, it is difficult to identify and migrate only closed sets of users and groups. For this reason, it is best to migrate groups before you migrate user accounts.

For example, User 1 belongs to global groups Global A and Global B and is a member of Domain 1. If an administrator moves User 1 and Global A to Domain 2, these accounts no longer exist in Domain 1; they exist only in Domain 2. Global B group remains in Domain 1. This creates an *open set*, or a set that includes users and groups in more than one domain. Because global groups can only contain members from the domain where the global group exists, the membership of User 1 in Global B is no longer valid, and User 1 can no longer access resources based on membership in Global B. Therefore, it is best to migrate both global groups before you migrate User 1.

If you are migrating an open set of objects in an environment where the functional level for both the source domain and the target domain is Windows 2000 native or higher, ADMT transforms the global group into a universal group so that it can contain users from other domains and retain the group membership. When the set becomes a closed set, ADMT changes the group back to a global group. The benefit of this process is that ADMT ensures that all closed set problems are resolved. However, replication of the global catalog is affected while the groups are universal groups because membership is copied to

the global catalog.

**Note**

- If the functional level of the source domain is Windows 2000 mixed, ADMT cannot transform the global group into a universal group because universal groups cannot exist at that functional level. Even if the target domain is in native mode, however, users in mixed mode domains would not get the SIDs of universal groups in their access tokens, if the groups are from outside the domain. Therefore, ADMT creates a copy of the global group in the target domain and adds all migrated users to the copy of that group. This group has a new security identifier (SID) and no SID history. This method does not preserve access to resources unless you run the ADMT Security Translation Wizard in Add mode to update permissions, which delays and complicates the migration process. For this reason, it is not recommended that you restructure domains that are operating at the Windows 2000 mixed functional level or the Windows Server 2003 interim domain functional level.

The second type of closed set is resources and local groups. In most cases, resources have permissions assigned to computer local groups or domain local groups. Because computer local groups are migrated when you migrate the computer, these groups are a natural closed set. However, domain local groups can be used on multiple computers to assign permissions. In this case, you can either migrate all the computers that use the domain local group at the same time the domain local group is migrated to the target domain, or you can manually change the domain local group to a universal group and then migrate the universal group. Changing the domain local group to a universal group is a manual process because ADMT does not automatically perform this task. Although this change can affect the size of your global catalog, over a limited time period, it is an effective way to migrate resources and domain local groups as a closed set.

#### SID History

SID history enables you to maintain user access to resources during the process of restructuring Active Directory domains. When you migrate an object to another domain, the object is assigned a new SID. Because you assign permissions to objects based on SIDs, when the SID changes, the user loses access to that resource until you can reassign permissions. When you use ADMT to migrate objects between domains, the SID history is automatically retained. In this way, the SID from the source domain remains as an attribute of the object after the object is migrated to the target domain.

For example, an organization that is restructuring its Active Directory domains moves universal and global groups from a source domain to the target domain before moving user accounts. Because this is a migration within a forest and the functional level of the source domain is Windows 2000 native, these groups cease to exist in the source domain and exist only in the target domain. Because the SID history of both users and groups is migrated, the users continue to have access to resources in the source domain based on their membership in a group that exists in the target domain.

#### Resource Assignment to Groups

The most effective way to assign permissions to resources is to assign users to global groups, place global groups within domain local groups, and then assign permissions to the domain local groups. Assigning permissions to resources in this way simplifies the migration process.

#### Terms and Definitions

The following terms apply to the process for restructuring Active Directory domains within a forest.

**Migration** The process of moving an object from a source domain to a target domain, while preserving or modifying characteristics of the object to make it accessible in the new domain.

**Domain Restructure** A migration process that involves changing the domain structure of a forest. A domain restructure can involve either consolidating or adding domains, and can take place between forests or within a forest.

**Domain Consolidation** A restructuring process that involves eliminating Windows NT 4.0 or Active Directory domains by merging their contents with that of other domains.

# Preparing to Restructure Active Directory Domains Within a Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Careful planning is an important part of restructuring Active Directory domains. Completing the necessary preparatory tasks before you restructure Active Directory domains enables you to reduce the effect on users of migrating objects from source to target domains. To prepare for the restructuring process, the Active Directory deployment team must obtain the necessary design information from the Active Directory design team.

Figure 12.2 shows the steps that are involved in preparing to restructure Active Directory domains within a forest.

**Figure 12.2 Preparing to Restructure Active Directory Domains Within a Forest**



# Design the New Active Directory Forest Structure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Evaluate the domain structure of your existing Active Directory forest, and then identify the domains that you want to restructure by consolidating them with other domains.

To design your new Active Directory forest structure, you must identify the source domain that you will migrate objects from and the target domains where you will place those objects. You must also evaluate the structure of the target domain.

**Identify the Target and Source Domains**
The source domain or domains are the domains that you want to migrate objects from and that you plan to decommission. When you restructure Active Directory domains, it is best to migrate the smallest possible number of objects. When you select source domains, identify the domains that have the fewest objects to migrate.

**Evaluate the OU Structure of the Target Domain**
Identify the organizational units (OUs) from the source domain that you need in the target domain, and then determine whether you need to create new OUs in the target domain.

You can migrate the OU structure when you migrate users, groups, or computers if you are using ADMT in command-line or scripting mode. The OUs are always copied between the domains and are not deleted in the source domain. To successfully migrate an OU, a source OU and a target OU must be specified in ADMT and the target OU must exist. All objects in the source OU and all sub-OUs are migrated. The specified source OU itself is not migrated. The migrated objects and sub-OUs are migrated to the target OU.

For more information about creating an OU structure, see "Designing the Active Directory Logical Structure" in this book.

# Assign Domain Object Roles and Locations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use a domain object assignment table to document object roles and locations. In the domain object assignment table, include the domain objects that exist in the source domain and their planned roles and locations in the target domain. For a worksheet to assist you in creating a domain object assignment table, see "Domain Object Assignment Table" (DSSRERA_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Object Assignment Table" on the Web at http://www.microsoft.com/reskit).

For example, Contoso Corporation plans to migrate their Africa domain objects to their EMEA domain. The design team creates a domain object assignment table that includes every group, user, and computer that needs to be migrated from the Africa domain to the EMEA domain. The completed table permits the organization to assess the amount of work and time that is required to complete the migration of the Africa domain.

Figure 12.3 shows the domain object assignment table that the Contoso design team created.

**Figure 12.3 Example of a Domain Object Assignment Table**

| Domain Object Assignment Table | | | |
|---|---|---|---|
| Prepared By | Design Team | Date | 01/01/03 |
| Forest Name | Concorp.contoso.com | | |
| Source Domain | Africa.concorp.contoso.com | | |
| Target Domain | Emea.concorp.contoso.com | | |

| Name | Type | Source Location | Target Location |
|---|---|---|---|
| Sales | Global Group | Africa\Groups OU | EMEA\Groups OU |
| Marketing | Global Group | Africa\Groups OU | EMEA\Groups OU |
| Engineering | Global Group | Africa\Groups OU | EMEA\Groups OU |
| AFR-AFRICA-DC01 | Domain Controller | Africa\Domain Controllers | EMEA\Domain Controllers |
| AFR-AFRICA-DC02 | Domain Controller | Africa\Domain Controllers | EMEA\Domain Controllers |
| SSmith | User | Africa\Users OU | EMEA\Users OU |
| EJohnson | User | Africa\Users OU | EMEA\Users OU |
| _SQL user | Service Account | Africa\Users OU | EMEA\Users OU |
| AFR-AFRICA-W2P1 | Computer Account | Africa\Wrk OU | EMEA\Wrk OU |
| AFR-AFRICA-W2P2 | Computer Account | Africa\Wrk OU | EMEA\Wrk OU |

# Plan for Group Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Unless you can identify closed sets when you are restructuring Active Directory domains within a forest, it is recommended that you migrate groups and users separately to ensure that users continue to have access to required resources.

Table 12.1 lists each type of group and where the group is physically located.

**Table 12.1 Location of Each Group Type**

| Group Type | Location |
|---|---|
| Global Group | Active Directory |
| Universal Group | Active Directory |
| Domain Local Group | Active Directory |
| Computer Local Group | Database of the local computer |

Each type of group is migrated differently based on the group's physical location and its rules for group membership. Global, universal, and domain local groups are migrated by using ADMT and can be transformed into universal groups for the duration of the migration if you are not migrating closed sets. You can update computer local group membership by running security translation.

Each group type has different rules for membership and serves a different purpose. This affects the order that the groups are migrated from the source to the target domains. Table 12.2 summarizes the groups and their membership rules.

**Table 12.2 Groups and Membership Rules**

| Type of Group | Rules and Membership |
|---|---|
| Universal groups | Universal groups can contain members from any domain in the forest and replicate group membership to the global catalog. For this reason, they can be used for administrative groups. When you restructure domains, migrate universal groups first. |
| Global groups | Global groups can include only members from the domain to which they belong. ADMT automatically changes the global group in the source domain to a universal group when it is migrated to the target domain if the functional level of both domains is Windows 2000 native or higher. ADMT automatically changes universal groups back to global groups after all members of the group are migrated to the target domain. |
| Domain local groups | Domain local groups can contain users from any domain. They are used to assign permissions to resources. When you restructure domains, you must migrate domain local groups when you migrate the resources to which they provide access, or you must change the group type to universal group. This minimizes the disruption in user access to resources. |

# Plan for Service Account Transitioning

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most services run within the context of the Local System account and because of this, they do not need any maintenance when they are migrated to a different domain. Some services, however, run in the context of a user account instead of the Local System account.

Service account transitioning refers to the process of identifying, migrating, and updating services that run in the context of user accounts. This process has three steps. First, start ADMT on a Windows Server 2003–based member server in the target domain, and then run the Service Account Migration Wizard. The Service Account Migration Wizard sends an agent to a computer that you specify in the wizard and identifies all services on the computer that are running in the context of a user account. The wizard only identifies service accounts that are running in the context of a user account and saves the information in the ADMT database; it does not actually migrate the accounts. The next step, which can occur later in the transition process, is to migrate the accounts by using the User Account Migration Wizard when you migrate the other user accounts.

For information about installing and initializing ADMT, see "Install ADMT" later in this chapter.

The Service Account Migration Wizard checks every service on a computer to identify services that run in the context of a user account. It is possible to create a security hole during the migration of service accounts if someone who is not a service administrator enters an account with administrative permissions in the source domain but uses an invalid password on their computer to start the service. The service will not start before the account migration, because the password is not correct, but it will work after migration because ADMT resets the password of the service account and configures all services that are using that service account with the new password.

To eliminate this possible security problem, it is important to include in the Service Account Migration Wizard only those servers that are managed by trusted administrators. Do not use the Service Account Migration Wizard to detect service accounts on computers that are not managed by trusted administrators, such as workstations.

If you do not identify and transition a trusted computer that therefore does not get its service account updated, you will need to manually set the new password created by ADMT. To do this, obtain the password from the Password.txt file, and then manually enter that account and password information for the service on the computer that did not get transitioned.

When the accounts that the Service Account Migration Wizard identifies in the ADMT database as running in the context of a user account are migrated to the target domain, ADMT grants each account the right to log on as a service.

**To run the Service Account Migration Wizard**

1. In ADMT, start the Service Account Migration Wizard.

2. Complete the wizard by using the information in Table 12.3.

**Table 12.3 Using the Service Account Migration Wizard**

| Wizard Page | Action |
|---|---|
| Domain Selection | In the **Source domain** box, type or select the NetBIOS or DNS name of the source domain.<br><br>In the **Target domain** box, type or select the NetBIOS or DNS name of the target domain. |
| Update Information | Click **Yes, update the information**. |
| Service Account Selection | Click **Add**.<br><br>In the **Select Computers** list box, select all servers that have service accounts.<br><br>Click **OK**,and then click **Next.** |
| Service Account Information | Select any user accounts that do not need to be marked as service accounts in the ADMT database, and then click **Skip/Include** to mark the accounts as **Skip**. |

The wizard connects to the selected computers, and then sends an agent to check every service on the remote computers. The Service Account Information page lists the services that are running in the context of a user account and the name of that user account. ADMT notes in its database that these user accounts need to be migrated as service accounts. If you do not want a user account to be migrated as a service account, select the account, and then click **Skip/Include** to change the status from **Include** to **Skip**.

You use **Update SCM** to update the Service Control Manager with the new information. Unless you have a failure in reaching a computer to update the service, the **Update SCM** button is not available. If you have a problem updating a service account after the account was identified and migrated, ensure that the computer that you are trying to reach is available, and then restart the Service Account Migration Wizard. In the wizard, click **Update SCM** to try to update the service. If you ran the Service Account Migration Wizard previously and the **Update SCM** button is not available, examine the ADMT log files to determine the cause of the problem. After you correct the problem and the agent can connect successfully, the **Update SCM** button becomes available.

# Plan for Test Migrations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can create a pilot for the migration process by migrating users and the global groups of which they are members to the target domain. However, unless you are migrating a closed set of users and groups in your pilot, the global groups will become universal groups and replicate their membership to the global catalog for the duration of the migration process. Use a test lab that simulates your production environment to test your migration process and to resolve any issues before you begin to restructure your Active Directory domains. Continue to test your migration process until you correct as many errors as possible.

Make sure that you test user access to resources after each step of the migration. If users do not have access to resources during the migration process, the migration is not successful.

For more information about setting up a pilot to test migration, see "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book.

# Create a Rollback Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you begin the migration process, you cannot roll back the changes that you make to the Active Directory domains in your forest. Because accounts are moved and not copied from one domain to another when you restructure domains, the changes are not reversible. If your plans change after you begin the migration process, the only way to return your source domain is to remigrate the accounts. Create a rollback plan in case you need to remigrate accounts after you have begun to restructure your domains. To create a rollback plan, select the method that you will use to remigrate accounts.

You can use the ADMT wizards to remigrate accounts from the target domain back to the source domain. In this case, the original target domain becomes the new source domain, and the original source domain becomes the new target domain. Follow the same steps in the wizards that you used earlier to migrate the accounts. If you remigrate the accounts, the objects that have been migrated to the target domain and then remigrated to the source domain will have new SIDs. However, they will have the original SID in their SID history so they will not be identical to the accounts before the migration, but they will have the same functionality.

Another option for remigrating accounts is to use the Undo Wizard in ADMT. The Undo Wizard uses fewer steps than the other ADMT wizards to remigrate accounts back to their original domain. However, the Undo Wizard only remigrates accounts that were originally migrated by using the User Account Migration Wizard, the Group Account Migration Wizard, and the Computer Account Migration Wizard. It does not undo changes that were made by the Service Account Migration Wizard or the Security Translation Wizard. The Undo Wizard only reverses the last operation that was performed.

If you want to reverse a service account migration, you must enumerate the services again, and then remigrate the service accounts by reversing the target and source domains.

If you use scripts to perform the original migration, using scripts to remigrate accounts is the fastest method to roll back the changes. Simply reverse the objects used for the source and target domains in the script to remigrate the objects.

**Note**

- If the functional level of the original source domain is Windows 2000 mixed, you cannot use a rollback method to undo the changes and migrate the accounts back to the source domain. A remigration requires that the original source domain becomes the target domain, and the functional level of the target domain must be Windows 2000 native or Windows Server 2003. For this reason, it is not recommended that you restructure domains that are operating at the Windows 2000 mixed functional level or the Windows Server 2003 interim domain functional level.

After you create your rollback plan, make sure to test it to identify and correct any problems before you begin to restructure your Active Directory domains.

# Create an End-User Communication Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Develop a plan to inform all affected users about the upcoming account migration, to ensure that they understand their responsibilities, the impact of the migration, and who to contact for help and support.

Before you begin the user migration process, send a notice to all users who are scheduled to be migrated. Because you typically migrate users in batches of approximately one hundred users at a time, it is also helpful to send a final notice to the users in each batch 2-3 days before their batch is scheduled. If your organization maintains an intranet, publish the account migration schedule and the information contained in the user mail on an easily accessible Web page.

Include the following information in your end-user communication.

**General information**
Alert users to the fact that their user accounts are scheduled to be migrated to a new domain. Point users to a Web page or internal resource where they can find additional information, and view a migration schedule.

Inform users of their new domain name. Be sure to let them know that their account passwords will not change. Let them know that the original domain account will be disabled immediately following the migration, and the disabled account will be deleted after a specified period of time. This is not needed if users use User Principal Names (UPNs) to log on.

**Impact**
Make sure that users understand that when their account is migrated, it is possible that they could be unable to access some resources, such as Web sites, shared folders, or resources that are not widely used by individuals in their group or division.

Provide information to users about who to contact for assistance in regaining access to required resources.

**Logon status during migration**
Make sure that users understand that during the migration process, they will be unable to log on to the domain or access e-mail or other resources. Be sure to specify the period of time for which they will be unable to log on.

**Premigration steps**
Alert users to any steps that they need to complete before the migration process begins. For example, they must decrypt files encrypted by means of Encrypted File System (EFS). Failure to decrypt encrypted files will result in loss of access to encrypted files following the migration.

**Expected changes**
Describe other changes that users can expect to experience following the migration, such as changes in use of smart cards, secure mail, or instant messaging if applicable.

**Scheduling and support information**
Provide information about where users can go to find more information; for example, an internal Web site where you post information about the migration. Also, provide information about who to contact if a user has a conflict with the date scheduled for the migration.

# Create Migration Account Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To migrate accounts and resources within a forest, you must establish an account migration group and a resource migration group with the appropriate credentials. You must then add the accounts that will be performing the ADMT migrations to the account migration and resource migration groups, as appropriate. Because ADMT requires only a limited set of permissions, creating separate migration groups allows you to simplify administration by creating the groups, assigning the appropriate permissions, and then adding the necessary administrators to those groups. If the migration tasks for your organization are distributed across more than one administrative group, create separate migration groups for each administrative group that performs the migration.

To simplify administration, create a single migration group in the source domain and a single migration group in the target domain for all objects. Assign the required permissions to modify objects such as users, global groups, and local profiles according to Table 12.4. The user who is running ADMT must be an administrator on the computer where ADMT is installed.

In the target domain, use a group with delegated control of the computer OU and the user OU. You might want to use a separate group for the migration of workstations if this migration process is delegated to administrators who are in the same location as the workstations.

Use the information in Table 12.4 to determine the credentials that are required for your migration.

**Table 12.4 Migration Account Group Credentials**

| Migration Object | Credentials Necessary in Source Domain | Credentials Necessary in Target Domain |
|---|---|---|
| User/Group | Local administrator or domain administrator, or delegated rights to delete the objects in the source OU | Delegated control of the user OU or the group OU and local administrator on the computer where ADMT is installed |
| Computer | Domain administrator or delegated rights to delete the objects in the source OU and member of Administrators group on each computer | Delegated control on the computer OU and local administrator on the computer where ADMT is installed |
| Profile (needed for Windows NT 4.0 computers only) | Local administrator or domain administrator | Delegated control on the computer OU and local administrator on the computer where ADMT is installed |

# Install ADMT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have not installed and initialized ADMT in your environment, you must install and run ADMT on a member server or domain controller in the target domain or the source domain.

**To install ADMT in the target domain or the source domain**

1. Log on to a member server or domain controller in the domain by using your ADMT migration account.
2. Insert the Windows Server 2003 CD into the CD-ROM drive, and then locate the \i386\admt folder.
3. Double-click **admigration.msi**.

Initialize ADMT by running a test migration of your global groups.

**To initialize ADMT by running a test migration of a global group**

1. In the ADMT console, select the **Group Account Migration Wizard**.
2. Complete the **Group Account Migration Wizard** by using the information provided in Table 12.5. Accept default settings if no information is specified.

   **Table 12.5 Migrating Global Groups by Using the Group Account Migration Wizard**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Test the migration settings and migrate later?** |
| Domain Selection | In the **Sourcedomain** box, type the source domain name.<br>In the **Target domain** box, type the target domain name. |
| Group Selection | Click **Add.**<br>In the **Select Groups** dialog box, select a group. |
| Organizational Unit Selection | Click **Browse**.<br>In the **Browse for Container** dialog box, locate the targetOU. |
| Group Options | Click **Do not rename accounts**. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

3. After the wizard runs, click **View Log**, and then review the migration log for any errors.

# Example: Preparing to Restructure Active Directory Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso Corporation upgraded its hardware to increase its network bandwidth and the amount of replication traffic that it can support. As a result, the company is consolidating the Africa domain into the EMEA domain.

The Africa domain is the source domain and the EMEA domain is the target domain for the migration. The organization needs to migrate a total of 1,800 users from the Africa domain to the EMEA domain. In addition to the user accounts, they must also migrate resources such as workstations, servers, and groups. Because Contoso Corporation is a large organization with many global groups, closed sets are difficult to identify, so the company decided to migrate global groups as universal groups. They can do this because the infrastructure of the corporation can handle the increased replication of the universal groups and because both the Africa and EMEA domains are operating at the Windows 2000 native functional level. The company created identical OU structures in the Africa and EMEA domains; therefore, they do not need to create a new OU structure or migrate OUs.

Contoso Corporation created a list of computers that run service accounts, so that it can use the Service Account Migration Wizard to identify services that run in the context of user accounts. The company is most concerned about a set of accounts that access a Microsoft® SQL Server™ database. Access to this database is an important part of their business.

The company decides to use ADMT as its migration tool and to use the wizards. The company installs ADMT and creates two account migration groups to use for the migration process. They assign high-level permissions to the first group, and then add the appropriate deployment team members to that group. The centralized deployment team will use this account to migrate users. They assign workstation and local resource permissions to the second group. The deployment team will use the second group to migrate resources at the remote locations.

# Migrating Domain Objects Between Active Directory Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restructuring Active Directory domains within a forest involves migrating domain objects in a specific order to ensure that users maintain access to resources. Figure 12.4 shows the process for migrating domain objects between Active Directory domains.

**Figure 12.4 Migrating Domain Objects Between Active Directory Domains**



# Migrate Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To protect your system against the problem of open sets when you restructure Active Directory domains within a forest, migrate groups before you migrate the user accounts that are members of those groups. If you only migrate groups when you migrate users, you might not migrate nested groups, thereby creating an open set. Migrate universal groups first, followed by global groups. Migrate domain local groups when you migrate the resources (domain controllers and member servers) on which they are used to assign permissions. Computer local groups migrate with the computer.

# Migrate Universal Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Migrate universal groups, without migrating users who are members of these groups at the same time, from the source domain to the target domain. Migrating universal groups without the users helps to protect against the problem of open sets. SID history allows group members to continue to have access to resources based on universal group membership. When you migrate universal groups to the target domain, they cease to exist in the source domain.

**Note**

- If you are migrating a small number of universal groups, you can migrate universal groups at the same time that you migrate global groups.

You can migrate universal groups by using the Group Account Migration Wizard or by using a script.

**To migrate universal groups without members**

- Complete the Group Account Migration Wizard by using the information provided in Table 12.6.

**Table 12.6 Using the Group Account Migration Wizard to Migrate Universal Groups**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS or DNS name of the source domain or select the name from a list.<br><br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain.<br><br>If ADMT includes the names of the source and target domains, ensure that they are correct. |
| Group Selection | Click **Add.**<br><br>In the **Select Groups** dialog box, select all universal groups that you want to migrate, click **Add**, and then click **OK**. |
| Organizational Unit Selection | Type the name of the OU, or click **Browse**.<br><br>In the **Browse for Container** dialog box, find the container in the target domain that you want to move the universal groups into, and then click **OK**. |
| Group Options | The **Migrate Group SIDs to target domain** and **Fix Group Membership** check boxes are checked and appear dimmed.<br><br>Click **Do not rename accounts**.<br><br>Ensure that no other options are selected. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

**To migrate universal groups by using a script**

- Use Listing 12.3 to prepare a script that incorporates ADMT commands and options for migrating groups within a forest.

**Listing 12.3 Migrating Groups Within a Forest**

```
<Job id=" MigratingGroupsWithinForest" >
```

```
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objGroupMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objGroupMigration = objMigration.CreateGroupMigration

   '
   'Specify general migration options.
   '

   objMigration.IntraForest = True
   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " source container"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"

   '
   'Migrate specified group objects.
   '

   objGroupMigration.Migrate admtData, Array(" group name1" ," group name2" )

   Set objGroupMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script to assist you in migrating groups, see "Migrating Groups Within a Forest" (DSSRERA_2.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Groups Within a Forest" on the Web at http://www.microsoft.com/reskit).

# Migrate Global Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Migrate global groups, without members, from the source domain to the target domain to protect against the problem of open sets. After global groups are migrated to the target domain, they cease to exist in the source domain if the source domain has a functional level of Windows 2000 native or higher.

Because global groups only contain members from their own domain, they cannot be migrated from one domain to another. ADMT changes global groups to universal groups when they are migrated. The universal group in the target domain retains the SID history of the global group in the source domain, which enables users to continue to access resources in the source domain after the global groups are migrated. ADMT changes the universal groups back to global groups after all members of the global group are migrated from the source domain to the target domain.

You do not need to include built-in and well-known global groups in your migration. Built-in and well-known groups already exist in the target domain. If a built-in or well-known global group is selected for migration, ADMT does not migrate it; instead, ADMT makes a note in the log and continues to migrate other global groups.

The procedure for using the Group Account Migration Wizard to migrate global groups is the same as that for migrating universal groups. For more information about the procedure for migrating global groups and universal groups, see "Migrate Universal Groups" earlier in this chapter.

After you complete the global group migration process, use Active Directory Users and Computers to verify that the global groups migrated successfully. Verify that the global groups no longer exist in the source domain and that the groups appear in the target domain in the OU that you specified during the migration process. The global groups are listed as universal groups in the target domain if they still have members in the source domain. To view a list of members of the universal group, right-click the group, click **Properties**, and then click the **Members** tab. The original members of the global group are listed. Note, however, that user accounts have not yet been migrated.

You can migrate global groups by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate global groups by using the ADMT console**

1. On the member server in the target domain where ADMT is installed, log on by using a user account that is a member of the ADMT account migration group.
2. Start ADMT, and select **Group Account Migration Wizard**.
3. Complete the Group Account Migration Wizard by using the information provided in Table 12.7.

**Table 12.7 Using the Group Account Migration Wizard to Migrate Groups**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS or DNS name of the source domain or select the name from a list.<br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain.<br>If ADMT includes the names of the source and target domains, ensure that they are correct. |
| Group Selection | Click **Add**.<br>In the **Select Groups** dialog box, select all global groups that you want to migrate (except built-in and well-known groups), click **Add**, and then click **OK**. |
| Organizational Unit Selection | Type the name of the OU, or click **Browse**.<br>In the **Browse for Container** dialog box, find the container in the target domain that you want to move the global groups into, and then click **OK**. |

| | |
|---|---|
| Group Options | Click **Migrate Group SIDs to target domain**.<br>Click **Do not rename accounts**.<br>Ensure that no other options are selected. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

4. After the wizard runs, click **View Log**, and review the migration log for any errors.
5. Open Active Directory Users and Computers, and then locate the target domain OU. Verify that the global groups exist in the target domain OU.

**To migrate global groups by using the ADMT command-line option**

1. On the member server in the target domain where ADMT is installed, log on by using a user account that is a member of the ADMT account migration group.
2. At a command line, type:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /IF:YES /SD:"source_domain" /TD:"target_domain" /TO:"target_OU"
[parameters]
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT GROUP /N "
group_name1
" "
group_name2
" /O: "
option_file.txt
"
```

Table 12.8 lists the parameters that are required for migrating global groups, the command-line parameters, and option file equivalents.

**Table 12.8 Parameters Required for Global Group Migrations**

| Parameters | Command-Line Syntax | Option File Syntax |
|---|---|---|
| Intra-Forest | `/IF:YES` | `IntraForest=YES` |
| Target domain | `/TD:`"*target_domain*" | `TargetDomain=`"*target_domain*" |
| Target OU location | `/TO:`"*target_OU*" | `TargetOU=`"*target_OU*" |
| Do not rename accts | `/RO:DONT` (default) | `RenameOption=DONT` |
| Ignore conflicting accts and do not migrate them | `/CO:IGNORE` (default) | `ConflictOptions=IGNORE` |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers, and then locate the target domain OU. Verify that the global groups exist in the target domain OU.

**To migrate global groups by using a script**

1. Use a script that incorporates ADMT commands and options for migrating universal groups. For more information about migrating universal groups, see "Migrate Universal Groups" earlier in this chapter.
2. After completing the global group migration by using a script, view the migration log. The migration.log file is stored in the folder where you installed ADMT, typically %Program Files%\Active Directory Migration Tool.

   For a sample script to assist you in migrating groups, see "Migrating Groups Within a Forest" (DSSRERA_2.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Groups Within a Forest" on the Web at http://www.microsoft.com/reskit).

| **Note** |
| --- |
| Because universal groups are replicated to the global catalog, converting global groups to universal groups can affect replication traffic. When the domain is operating at the Windows Server 2003 functional level, this impact is reduced because only changes to the universal group membership are replicated. However, if the domain is not operating at the Windows Server 2003 functional level, the entire group membership replicates every time universal group membership changes. |

# Migrate Service Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Migrate the service accounts that you identified earlier in the intraforest restructure process by using the Service Account Migration Wizard. This wizard marked the accounts as service accounts in the ADMT database. For more information about using ADMT to identify service accounts that are running in the context of a user account, see "Plan for Service Account Transitioning" earlier in this chapter.

You can migrate service accounts by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate service accounts by using the ADMT console**

- Complete the User Account Migration Wizard by using the information provided in Table 12.9.

**Table 12.9 Using the User Account Migration Wizard to Migrate Service Accounts**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| User Selection | Click **Add**.<br>In the **Select Users** dialog box, select all the service accounts that were identified by the Service Account Migration Wizard, and then click **Add**. By default, the wizard adds the service accounts to the Users container. If you need to change the OU, click the **Location** button, and then locate the OU that contains the service accounts.<br>Click **OK**. |
| Organizational Unit Selection | Click **Browse**.<br>In the **Browse for Container** dialog box, locate the source domain, select the container for the service accounts, and then click **OK**. |
| User Options | Select the **Update user rights** check box.<br>Ensure that the **Do Not Rename Accounts** check box is selected.<br>Ensure that no other settings are selected, including the **Migrate associated user groups** option. A warning box will appear to inform you that if the global groups to which the user accounts belong are not also migrated, users will lose access to resources. Select **OK** to continue with the migration. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |
| Service Account Information | Click **Migrate all service accounts and update SCM for items marked include**. The wizard will present you with a list of the service accounts that you are migrating (if you are migrating accounts that are not service accounts, they will be migrated but will not be listed). By default, the accounts are marked as **Include**. To change the status of the account, select the account, and then click the **Skip/Include** button.<br>Click **Next** to migrate the accounts. |

A **Migration Progress** dialog box updates you on the status of the migration. During this time, ADMT moves the accounts to the target domain, generates a new password for the accounts, assigns the accounts the right to log on as a service, and

provides this new information to the services that use the accounts. When the status is listed as **Completed** in the **Migration Progress** dialog box, you can continue with the rest of the intraforest migration. Before the migration of the service accounts is completed, users might experience interruptions when they use the services because the service still uses the account that has been migrated until the service is restarted. For any services that continually use credentials, such as search services, manually restart the services to ensure optimal results.

**To migrate service accounts by using the ADMT command-line option**

1. On a member server in the target domain where ADMT is installed, log on by using a user account that is a member of the ADMT account migration group.
2. At the command line, type:

```
ADMT USER /N "
server_name1
" "
server_name2
" /IF:YES /SD:"source_domain" /TD:"target_domain" /TO:"target_OU"
/MSA:YES
```

   *Server_name1* and *Server_name2* are the names of servers in the source domain that run service accounts. Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT USER /N "
server_name1
" "
server_name2
" /O: "
option_file.txt
"
```

   Table 12.10 lists the parameters that are required for migrating service accounts, the command-line parameters, and option file equivalents.

   **Table 12.10 Parameters Required for Migrating Service Accounts**

| Parameters | Command Line Syntax | Option File Syntax |
|---|---|---|
| Intra-Forest | `/IF:YES` | `IntraForest=YES` |
| Target domain | `/TD:`"*target_domain*" | `TargetDomain=`"*target_domain*" |
| Target OU location | `/TO:`"*target_OU*" | `TargetOU=`"*target_OU*" |
| Migrate Service Accounts | `/MSA:YES` | `MigrateServiceAccounts=YES` |
| Update user rights | `/UUR:YES` | `UpdateUserRights=YES` |
| Do not rename accounts | `/RO:DONT` *(default)* | `RenameOption=DONT` *(default)* |
| Ignore conflicting accounts | `/CO:IGNORE` *(default)* | `ConflictOptions=IGNORE` *(default)* |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers, and locate the target domain OU. Verify that the service accounts exist in the target domain OU.

**To migrate service accounts by using a script**

- Use Listing 12.4 to prepare a script that incorporates ADMT commands and options for migrating service accounts.

**Listing 12.4 Migrating Service Accounts Within a Forest**

```
<Job id=" MigratingServiceAccountsWithinForest" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objUserMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objUserMigration = objMigration.CreateUserMigration

   '
   'Specify general migration options.
   '

   objMigration.IntraForest = True
   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " source container"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"

   '
   'Specify user migration specific options.
   '

   objUserMigration.UpdateUserRights = True
   objUserMigration.MigrateServiceAccounts = True

   '
   'Migrate specified service accounts.
   '

   objUserMigration.Migrate admtData, _
  Array(" service account name1" ," service account name2" )

   Set objUserMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script to assist you in migrating service accounts, see "Migrating Service Accounts Within a Forest" (DSSRERA_3.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Service Accounts Within a Forest" on the Web at http://www.microsoft.com/reskit).

# Migrate User Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domains can include a large number of user accounts. To make the migration of user accounts manageable, use a technique called *phased transitioning*, by which you place your user accounts into smaller batches and migrate each of the smaller batches individually. You can group the users in any way that you prefer.

You cannot migrate every user property when you migrate user accounts. For example, Protected Storage (Pstore) contents for Windows NT 4.0 workstations, including Encrypting File System (EFS) private keys, are not migrated by ADMT when you migrate user accounts. To migrate Pstore contents, you must export and import keys during the migration process.

For clients that are running Windows 2000 or later operating systems, data that is protected by the Data Protection API (DPAPI) is also not migrated. DPAPI helps protect the following items:

- Web page credentials (for example, passwords)
- File share credentials
- Private keys associated with EFS, S/MIME, and other certificates
- Program data that is protected by using the CryptProtectData() function

For this reason, it is important to test user migrations by using your test migration account to identify any properties that did not migrate and update user configurations in the target domain accordingly.

If you are using Group Policy objects to manage software installation, you need to consider that some software installation (.msi) packages require access to the original source for certain operations, such as repair and uninstall. The administrator must reassign permissions to the software distribution point to provide access to any account.

To retain software distribution access, perform these tasks:

- Migrate users by using ADMT.
- Run the Security Translation Wizard on the software distribution point.

# Migrating OUs and Subtrees of OUs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you want to copy OUs and subtrees of OUs to your target domain, you can use either the command-line or scripting option and substitute the appropriate parameters. You must specify a source OU and a target OU, and the target OU must exist. All objects in the source OU and all sub-OUs are migrated, but the specified source OU itself is not migrated. The migrated objects and sub-OUs are migrated to the target OU.

If you are using the command-line option to migrate your accounts, groups, or computers, and you also want to migrate OUs, modify the command line to use the /D option. Instead of using the /N (/IncludeName) option, you must use the /D (/IncludeDomain) option with RECURSE and MAINTAIN as follows:

```
ADMT /D:RECURSE+MAINTAIN /O "option_file.txt"
```

If you are migrating accounts, groups, or computers by using the scripting option, and you also want to migrate OUs, modify your script to use the admtDomain option. Instead of using the admtData or admtFile option, you must use the admtDomain option with admtRecurse and admtMaintainHierarchy as follows:

```
objUserMigration.Migrate admtDomain + admtRecurse + admtMaintainHierarchy
```

# Migrate Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can migrate each batch of user accounts by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate user accounts by using the ADMT console**

- Complete the User Account Migration Wizard by using the information provided in Table 12.11.

  **Table 12.11 Using the User Account Migration Wizard to Move User Accounts**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| User Selection | Click **Add**.<br>In the **Select Users** dialog box, select the user accounts that you want to migrate in the current batch, and then click **Add**.<br>Click **OK**. |
| Organizational Unit Selection | Ensure that ADMT lists the correct target OU. If it is not correct, type the correct OU, or click **Browse**.<br>In the **Browse for Container** dialog box, locate the target domain and OU, and then click **OK**. |
| User Options | Click the **Translate roaming profiles** check box.<br>Click the **Update user rights** check box.<br>Clear the **Migrate associated user groups** check box. A warning box appears that states that if the global groups to which the user accounts belong are not also migrated, users will lose access to resources. Click **OK** to continue with the migration. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

After you click **Finish** in the User Account Migration Wizard, the **Migration Progress** dialog box appears. After the status changes to **Completed**, view the migration log to determine whether any errors occurred in the migration process. In the **Migration Progress** dialog box, click **Close**.

**To migrate the user accounts by using the ADMT command-line option**

1. On the member server in the target domain where ADMT is installed, log on by using a user account that is a member of the ADMT account migration group.
2. At the command line, type:

```
ADMT USER /N "
user_name1
" "
user_name2
" /IF:YES /SD:"source_domain" /TD:"target_domain" /TO:"target_OU"
[parameters]
```

You can append parameters to the command as follows:

```
ADMT USER /N "
user_name
1" "
user_name2
" /IF:YES /SD:"source_domain" /TD:"target_domain" /TO:"target_OU"
TRP:YES /UUR:YES
```

Alternatively, you can include parameters in an option file that is specified on the command line as follows:

```
ADMT USER /N "
user_name1
" "
user_name2
" /O "
option_file
.txt"
```

Table 12.12 lists the parameters that are required for migrating user accounts, the command-line parameters, and option file equivalents.

**Table 12.12 Parameters Required for Migrating User Accounts**

| Parameters | Command Line Syntax | Option File Syntax |
|---|---|---|
| Intra-Forest | `/IF:YES` | `IntraForest=YES` |
| Source domain | `/SD:"`*source_domain*`"` | `SourceDomain="`*source_domain*`"` |
| Source OU location | `/SO:"`*source_OU*`"` | `SourceOU="`*source_OU*`"` |
| Target domain | `/TD:"`*target_domain*`"` | `TargetDomain="`*target_domain*`"` |
| Target OU location | `/TO:"`*target_OU*`"` | `TargetOU="`*target_OU*`"` |
| Do not rename accts | `/RO:DONT` (default) | `RenameOption=DONT` |
| Ignore conflicting accts and not migrate them | `/CO:IGNORE` (default) | `ConflictOptions=IGNORE` |
| Translate Roaming Profile | `/TRP:YES` (default) | `TranslateRoamingProfile=YES` |
| Update User Rights | `/UUR:YES` | `UpdateUserRights=YES` |

3. Review the results that are displayed on the screen for any errors.
4. Open Active Directory Users and Computers, and then locate the target domain OU. Verify that the users exist in the target domain OU.

**To migrate user accounts by using a script**

- Use Listing 12.5 to prepare a script that incorporates ADMT commands and options for migrating user accounts within a forest.

**Listing 12.5 Migrating User Accounts Within a Forest**

```
<Job id=" MigratingUserAccountsWithinForest" >
```

```
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objUserMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objUserMigration = objMigration.CreateUserMigration

   '
   'Specify general migration options.
   '

   objMigration.IntraForest = True
   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " source container"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " target container"

   '
   'Specify user migration specific options.
   '

   objUserMigration.TranslateRoamingProfile = True
   objUserMigration.UpdateUserRights = True
   objUserMigration.FixGroupMembership = True
   objUserMigration.MigrateServiceAccounts = False

   '
   'Migrate specified user objects.
   '

   objUserMigration.Migrate admtData, Array(" user name1" ," user name2" )

   Set objUserMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script file to assist you in migrating user accounts, see "Migrating User Accounts Within a Forest" (DSSRERA_4.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating User Accounts Within a Forest" on the Web at http://www.microsoft.com/reskit).

# Translate Local User Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Translate local user profiles after you migrate the user accounts. You only need to translate local user profiles for Windows NT 4.0 computers. To minimize the disruption to users, translate local user profiles immediately after you migrate a batch of users. If your source domain includes only a small number of pre-Active Directory clients, migrate them as a group, and then translate their user profiles before you migrate the next batch of users.

**Note**

- Because roaming profiles are stored on a server, you do not need to translate them.

Local profiles are translated in replace mode because if you perform the profile translation in add mode, software installation by using software deployment group policies might not work. Any application that is packaged with Microsoft® Windows® Installer version 2.0 (which is used on workstations with Windows 2000 Service Pack 3 (SP3) or later, Microsoft® Windows® XP Service Pack 1 (SP1) or later, and many common software packages) might not function after the profile is translated. When the Security Translation Wizard is translating local profiles in replace mode, it reverts to add mode if a profile is locked. This could result in a successful profile translation, but application installations might not function after the profile is translated.

No action is required to translate a local profile on Windows 2000 or Windows XP clients between domains in the same forest because the globally unique identifier (GUID) of the user remains the same. The local profile can use the SID-to-GUID mapping that it preserves in the registry to reassign the profile of the user, and then reassociate it with the new SID.

For profile translations if a user is using offline files on a client running Windows XP SP1, the user loses access to the files in the offline folder. Although the SID of the user changes, the owner in the access control lists of the files and folders does not change. On Windows XP SP1 clients, users will not have access to content in offline folders unless they are owners of the files and folders. Therefore, to give the user access to the offline file folder, you must run the Security Translation Wizard on the profile folder.

If you are migrating the user account to a domain within the forest, and the path for the local profile is different, the user profile is modified, and a new profile folder is created on the server with the correct ACLs. The administrator must make sure that the user has access to the profile folder.

You can translate local user profiles by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To translate local user profiles by using the ADMT console**

- Complete the Security Translation Wizard by using the information provided in Table 12.13.

  **Table 12.13 Using the Security Translation Wizard to Translate Local User Profiles**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Security Translation Options | Click **Previously migrated objects**. |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| Translate Objects | Click **User Profiles**. |
| Security Translation Options | Click **Replace**. |

**To translate local user profiles by using the ADMT command-line option**

1. On the member server in the target domain where ADMT is installed, log on by using a user account that is a member of the ADMT account migration group.

2. At the command line, type:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" [
parameters
]
```

You can append parameters to the command as follows:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" /IF:YES /SD:"source_domain" /TD:"target_domain" /TOT:REPLACE
/TUP:YES
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT SECURITY /N "
computer_name1
" "
computer_name2
" /O:"
option_file.txt
"
```

Table 12.14 lists the parameters that are required for translating local user profiles, command-line parameters, and option file equivalents.

**Table 12.14 Parameters Required for Migrating Local User Profiles**

| Parameters | Command Line Syntax | Option File Syntax |
|---|---|---|
| IntraForest | /IF:YES | IntraForest=YES |
| Source domain | /SD:"*source_domain*" | SourceDomain="*source_domain*" |
| Target domain | /TD:"*target_domain*" | TargetDomain="*target_domain*" |
| Target domain | /TOT:REPLACE | TranslateOption=REPLACE |
| Modify local user profile security | /TUP:YES | TranslateUserProfiles=YES |

3. Review the results that are displayed in the migration log for any errors.

**To translate local user profiles by using a script**

- Use Listing 12.6 to prepare a script that incorporates ADMT commands and options for translating local user profiles.

**Listing 12.6 Translating Local User Profiles Within a Forest**

```
<Job id=" TranslatingLocalProfilesWithinForest" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
```

```
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objSecurityTranslation

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objSecurityTranslation = objMigration.CreateSecurityTranslation

   '
   'Specify general migration options.
   '

   objMigration.IntraForest = True
   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"


   '
   'Specify security translation specific options.
   '

   objSecurityTranslation.TranslationOption = admtTranslateReplace
   objSecurityTranslation.TranslateUserProfiles = True

   '
   'Perform security translation on specified computer objects.
   '

   objSecurityTranslation.Translate admtData, _
  Array(" computer name1" ," computer name2" )

   Set objSecurityTranslation = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script to assist you in migrating local user profiles, see "Translating Local Profiles Within a Forest" (DSSRERA_5.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Local Profiles Within a Forest" on the Web at http://www.microsoft.com/reskit).

# Migrate Workstations and Member Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Migrate workstations and member servers from the source domain to the target domain. When you migrate computers, the changes do not take effect until the computer is restarted. Restart the computers that you are migrating as soon as possible to complete the migration process.

Computer accounts are treated differently than user and group accounts during a migration between domains within an Active Directory forest. Computer accounts are copied instead of moved during an intraforest migration, and an account for the computer remains in the source domain after the computer has been successfully migrated to the target domain. It is recommended that you delete the computer account in the source domain seven days after the migration is complete.

You can migrate workstations and member servers by using the ADMT console, by using the ADMT command-line option, or by using a script.

**To migrate workstations and member servers by using the ADMT console**

- Complete the Computer Account Migration Wizard by using the information provided in Table 12.15.

**Table 12.15 Using the Computer Account Migration Wizard to Migrate Workstations and Member Servers**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| Computer Selection | Click **Add**.<br>In the **Select Computers** dialog box, select the names of the workstations and member servers in the source domain that you want to migrate, click **Add**, and then click **OK**. |
| Organizational Unit Selection | Click **Browse**.<br>In the **Browse for Container** dialog box, locate the OU in the target domain to which the computers are migrating, and then click **OK**. |
| Translate Objects | Click the **Local groups** check box**.**<br>Click the **User rights** check box. |
| Security Translation Options | Click **Add**. |
| Computer Options | In the **Minutes before computer restart after wizard completion** box, type **1**. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

**To migrate workstations and member servers by using the ADMT command-line option**

1. On the member server in the target domain where ADMT installed, log on by using a user account that is a member of the ADMT resource migration group.
2. At the command line, type:

```
ADMT COMPUTER /N "
computer_name1
"
"
computer_name2
" /IF:YES /SD:"source_domain" /TD:"target_domain" /TO:"target_OU"
[parameters]
```

You can append parameters to the command as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /IF:YES /SD:"source_domain" /TD:"target_domain" /TO:"target_OU"
/RDL:1
```

Alternatively, you can include parameters in an option file that is specified at the command line as follows:

```
ADMT COMPUTER /N "
computer_name1
" "
computer_name2
" /O "
option_file.txt
"
```

Table 12.16 lists the parameters that are required for workstation and member server migration, the command-line parameters, and option file equivalents.

**Table 12.16 Parameters Required for Workstation and Member Server Migrations**

| Parameters | Command Line Syntax | Option File Syntax |
|---|---|---|
| IntraForest | `/IF:YES` | `IntraForest=YES` |
| Source domain | `/SD:`"*source_domain*" | `SourceDomain=`"*source_domain*" |
| Target domain | `/TD:`"*target_domain*" | `TargetDomain=`"*target_domain*" |
| Target OU location | `/TO:`"*target_OU*" | `TargetOU=`"*target_OU*" |
| • Restart computer one minute after wizard completes | `/RDL:1` | `RestartDelay=1` |
| • Do not rename accounts | `/RO:DONT` (default) | `RenameOption=DONT` |
| • Ignore conflicting accounts and do not migrate them | `/CO:IGNORE` (default) | `ConflictOptions=IGNORE` |
| • Translate Option | `/TOT:ADD` | `TranslateOption=YES` |

| | | |
|---|---|---|
| • Translate User Rights | /TUR:YES | TranslateUserRights=YES |
| • Translate Local Groups | /TLG:YES | TranslateLocalGroups=YES |

3. Review the results that are displayed on the screen for any errors. The migration log lists computers, completion status, and the path to the DCTLog.txt file for each computer. If an error is reported for a computer, you will need to refer to the DCTLog.txt file on that computer to review any problems with local groups.
4. Open Active Directory Users and Computers, and then locate the target domain OU. Verify that the workstations and member servers exist in the target domain OU.

**To migrate workstations and member servers by using a script**

- Use Listing 12.7 to prepare a script that incorporates ADMT commands and options for migrating workstations and member servers within a forest.

**Listing 12.7 Migrating Workstations and Member Servers Within a Forest**

```
<Job id=" MigratingWorkstationsMemberServersWithinForest" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objComputerMigration

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objComputerMigration = objMigration.CreateComputerMigration

   '
   'Specify general migration options.
   '

   objMigration.IntraForest = True
   objMigration.SourceDomain = " source domain"
   objMigration.SourceOu = " Computers"
   objMigration.TargetDomain = " target domain"
   objMigration.TargetOu = " Computers"

   '
   'Specify computer migration specific options.
   '

   objComputerMigration.TranslationOption = admtTranslateAdd
   objComputerMigration.TranslateLocalGroups = True
   objComputerMigration.TranslateUserRights = True
   objComputerMigration.RestartDelay = 1

   '
   'Migrate computer objects on specified computer objects.
   '

   objComputerMigration.Migrate admtData, _
```

```
  Array(" computer name1" ," computer name2" )

   Set objComputerMigration = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script to assist you in migrating workstations and member servers, see "Migrating Workstations and Member Servers Within a Forest" (DSSRERA_6.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations and Member Servers Within a Forest" on the Web at http://www.microsoft.com/reskit).

# Migrate Domain Local Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Migrate the domain local groups that exist in Active Directory. You can migrate domain local groups by using the ADMT console or by using a script.

**To migrate domain local groups by using the ADMT console**

- Complete the Group Account Migration Wizard by using the information provided in Table 12.17.

**Table 12.17 Using the Group Account Migration Wizard to Migrate Domain Local Groups**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Domain Selection | In the **Source domain** box, type the NetBIOS or DNS name of the source domain, or select the name from a list.<br><br>In the **Target domain** box, type the NetBIOS or DNS name of the target domain. |
| Group Selection | Click **Add**.<br><br>In the **Select Groups** dialog box, select all domain local groups that you need to migrate (except built-in and well-known groups), click **Add**, and then click **OK**. |
| Organizational Unit Selection | Type the name of the OU, or click **Browse**.<br><br>In the **Browse for Container** dialog box, locate the OU in the target domain to which the domain local groups are migrating, and then click **OK**. |
| Group Options | The **Migrate Group SIDs to target domain** and **Fix Group Membership** check boxes are selected and appear dimmed.<br><br>Click **Do not rename accounts**.<br><br>Ensure that no other options are selected. |
| Naming Conflicts | Click **Ignore conflicting accounts and don't migrate**. |

**To migrate domain local groups by using a script**

- Use a script that incorporates ADMT commands and options for migrating universal groups. For more information about migrating universal groups, see "Migrate Universal Groups" earlier in this chapter.

  For a sample script to assist you in migrating groups, see "Migrating Groups Within a Forest" (DSSRERA_2.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Groups Within a Forest" on the Web at http://www.microsoft.com/reskit).

# Example: Restructuring Active Directory Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso Corporation wants to migrate the Africa domain objects to the EMEA domain. To minimize the impact on users and reduce the time that network traffic is affected, Contoso decided to complete the migration of domain objects in as little time as possible. During the migration, all global groups that are migrated become universal groups until each user who belongs to the group is migrated to the target domain. Because the forest includes domains running at the Windows 2000 functional level, only incremental membership changes to the universal groups are replicated to the global catalog.

Contoso administrators set up an extensive test lab to test the migration process before proceeding. They determined by testing the procedures that they could complete the migration process in two weeks. Global and universal groups will be migrated on Monday and Tuesday of the first week. Service accounts will be migrated and services updated on Wednesday. Thursday, Friday, and Saturday are reserved for migrating user accounts. Servers will be migrated on Sunday of the first week. The second week is reserved for migrating workstations. Domain local groups will be migrated at the end of the second week.

# Completing Post-Migration Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you complete all the migration tasks that are required to restructure your Active Directory domains within a forest, you must verify that the migration occurred as planned and complete a few post-migration tasks. Figure 12.5 shows the process for completing post-migration tasks.

**Figure 12.5 Process for Completing Post-Migration Tasks**



# Examine Migration Logs for Errors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ADMT keeps a detailed log of every action that you perform when you migrate resources between Active Directory domains. Errors that occur during the migration process are noted in the migration log, although they might not produce a warning message in ADMT. Examining the migration log after a migration is a good way to verify that all tasks were completed successfully. Because it is important to complete the steps of the migration in a specific order, it is best to check the migration log after each step so that you can discover any failures in time to fix them.

# Verify Group Types

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ADMT changes global groups to universal groups when you migrate them from the source domain to the target domain. The change occurs automatically because global groups can only contain members of their own domain. Therefore, they cannot continue to be global groups when they are migrated to another domain until the group members are migrated. ADMT changes the universal groups back to global groups when the last member of the group is migrated to the target domain. Because universal groups replicate their membership to the global catalog, it is important to verify that the universal groups correctly changed back to global groups.

Use Active Directory Users and Computers to verify that universal groups migrated successfully. If you manually changed domain local groups into universal groups, make sure that you switch them back to domain local groups when all resources have been migrated.

# Translate Security on Member Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Translate security on member servers to clean up the ACLs of the resources. After objects are migrated to the target domain, resources contain the ACL entries of the source domain objects. Although SID history provides access to resources during the migration, access control lists should be cleaned up after the migration to contain the new primary SID of the migrated groups. Use the Security Translation Wizard in ADMT to replace the source domain SIDs with the target domain SIDs.

**To translate security on member servers by using the ADMT console**

- Complete the Security Translation Wizard by using the information provided in Table 12.18.

**Table 12.18 Using the ADMT Security Translation Wizard**

| Wizard Page | Action |
|---|---|
| Test or Make Changes | Click **Migrate Now?** |
| Security Translation Options | Click **Previously migrated objects**. |
| Domain Selection | In the **Source domain** box, type or select the name of the source domain.<br>In the **Target domain** box, type or select the name of the target domain. |
| Translate Objects | Click **file and folders**, **shares**, **printers**, **user rights**, and **registry**. |
| Security Translation Options | Click **Replace**. |

**To translate security on member servers by using a script**

- Use Listing 12.8 to prepare a script that incorporates ADMT commands and options to translate security on member servers.

**Listing 12.8 Translating Security on Member Servers Within a Forest**

```
<Job id=" TranslatingSecurityOnMemberServersWithinForest" >
<Script language=" VBScript"  src=" AdmtConstants.vbs" />
<Script language=" VBScript" >
   Option Explicit

   Dim objMigration
   Dim objSecurityTranslation

   '
   'Create instance of ADMT migration objects.
   '

   Set objMigration = CreateObject(" ADMT.Migration" )
   Set objSecurityTranslation = objMigration.CreateSecurityTranslation

   '
   'Specify general migration options.
   '

   objMigration.IntraForest = True
   objMigration.SourceDomain = " source domain"
   objMigration.TargetDomain = " target domain"
```

```
   objMigration.TargetOu = " Computers"

   '
   'Specify security translation specific options.
   '

   objSecurityTranslation.TranslationOption = admtTranslateReplace
   objSecurityTranslation.TranslateFilesAndFolders = True
   objSecurityTranslation.TranslateLocalGroups = True
   objSecurityTranslation.TranslatePrinters = True
   objSecurityTranslation.TranslateRegistry = True
   objSecurityTranslation.TranslateShares = True
   objSecurityTranslation.TranslateUserProfiles = False
   objSecurityTranslation.TranslateUserRights = True

   '
   'Perform security translation on specified computer objects.
   '

   objSecurityTranslation.Translate admtData, _
  Array(" computer name1" ," computer name2" )

   Set objSecurityTranslation = Nothing
   Set objMigration = Nothing
</Script>
</Job>
```

For a sample script to assist you in translating security on member servers, see "Translating Security On Member Servers Within a Forest" (DSSRERA_7.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Security On Member Servers Within a Forest" on the Web at http://www.microsoft.com/reskit).

# Translate Security by Using a SID mapping File

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To translate security so that permissions that are granted to the source account or group are now granted to the target account or group, use a SID mapping file to associate the two accounts. The SID mapping file is a comma-separated-value (CSV) format file that lists pairs of accounts in either a Windows NT account name (domain\name) format or a SID format. The account on the left side of the file is the source account, and the account on the right side of the file is the target account. ADMT security translation translates security from the source account to the target account.

To use a SID mapping file, you must run ADMT from the command line. The parameter is **/SMF**. The full command looks similar to the following:

**ADMT SECURITY /N** " *computer_name* " **/SMF:**" *sid_mapping_file_path* "

A large SID mapping file can take hours to load on the computer that is running ADMT. If the computer that is running ADMT is also an Active Directory domain controller, this problem may affect the performance of that domain controller. For more information and a hotfix to help alleviate this problem, see article 842533 in Microsoft Help and Support at http://go.microsoft.com/fwlink/?LinkId=74264.

# Decommission the Source Domain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have migrated all the objects from the source domain to the target domain, including all the computers and member servers, only the domain controllers remain in the source domain. To decommission the source domain, run the Active Directory Installation Wizard to remove Active Directory from the domain controllers in the source domain. Migrate the domain controllers from the source domain to the target domain as member servers. If necessary, depending on the new role that is planned for the servers in the target domain, use the Active Directory Installation Wizard to install Active Directory on the member servers to return them to domain controller status in the target domain. Run security translation on domain controllers if resources reside on the computer that will be used as the new domain controller.

# Example: Completing Post-Migration Tasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The post-migration team of Contoso Corporation starts the post-migration tasks during the first week of the migration. The team examines the migration log after the first group of migrations is completed on the first day. They analyze the migration log and define the action that is required to migrate any accounts that run into errors so that the migration team can continue with the migration without interruption.

During the second week of the migration process, the deployment team verifies that global groups have returned from universal group to global group status after the migration of users has completed. After the member servers are migrated, the deployment team runs the Security Translation Wizard to remove the source domain SIDs from the ACLs of the member servers. Finally, the deployment team decommissions the Africa domain at the end of the second week by removing Active Directory from the domain controllers in the Africa domain. They then migrate the domain controllers to the EMEA domain as member servers.

# Additional Resources for Restructuring Active Directory Domains Within a Forest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Designing the Active Directory Logical Structure" in this book for more information about the Active Directory logical structure.
- "Designing the Site Topology" in this book for more information about Active Directory site topology.
- "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" in this book for more information about restructuring Windows NT 4.0 domains to an Active Directory forest.
- "Designing a Resource Authorization Strategy" in this book for more information about creating a security group strategy.

**Related Tools**

- ADMT version 2

  You can use ADMT version 2 to migrate accounts to restructure your Active Directory domains. ADMT includes wizards that you can use to migrate objects such as groups, users, computers, and service accounts. These wizards assist you in performing the appropriate migration for each object type.

**Related Job Aids**

- "Domain Object Assignment Table" (DSSRERA_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Domain Object Assignment Table " on the Web at http://www.microsoft.com/reskit).
- "ADMT Option File Reference" (DSSRENT_3.txt) on the *Windows Server 2003 Deployment Kit* companion CD (or see "ADMT Option File Reference" on the Web at http://www.microsoft.com/reskit).
- "ADMT Scripting Constants" (AdmtConstants.vbs) on the *Windows Server 2003 Deployment Kit* companion CD (or see "ADMT Constants" on the Web at http://www.microsoft.com/reskit).
- "Migrating Groups Within a Forest" (DSSRERA_2.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Groups Within a Forest" on the Web at http://www.microsoft.com/reskit).
- "Migrating Service Accounts Within a Forest" (DSSRERA_3.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Service Accounts Within a Forest" on the Web at http://www.microsoft.com/reskit).
- "Migrating User Accounts Within a Forest" (DSSRERA_4.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating User Accounts Within a Forest" on the Web at http://www.microsoft.com/reskit).
- "Translating Local Profiles Within a Forest" (DSSRERA_5.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Local Profiles Within a Forest" on the Web at http://www.microsoft.com/reskit).
- "Migrating Workstations and Member Servers Within a Forest" (DSSRERA_6.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Migrating Workstations and Member Servers Within a Forest" on the Web at http://www.microsoft.com/reskit).

- "Translating Security On Member Servers Within a Forest" (DSSRERA_7.wsf) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Translating Security On Member Servers Within a Forest" on the Web at http://www.microsoft.com/reskit).

# Planning a Secure Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To plan a secure environment, you need a clear and consistent strategy for addressing the many aspects of the Microsoft® Windows® Server 2003 operating system, including security-related issues and features. You can begin by identifying the user-related requirements that impact security, and the other aspects of the network that comprise a secure common infrastructure. You can use this chapter to locate the chapters in this book that can help you address those requirements.

**In This Chapter**

Overview of a Secure Windows Server 2003 Environment

Addressing User-Related Requirements

Establishing a Secure Shared IT Infrastructure

Additional Resources for Planning a Secure Environment

**Related Information**

- For more information about deployment project planning, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* in this kit.
- For more information about Active Directory project planning, see "Planning an Active Directory Deployment Project" in this book.
- For more information about Group Policy-related features, see "Planning a Managed Environment" in *Designing a Managed Environment* in this kit.

# Overview of a Secure Windows Server 2003 Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Secure IT environments are the result of careful planning from the time you decide to deploy a new feature or service to the moment when responsibility for those features and services is handed over to those responsible for the day-to-day operation of the network. To ensure a secure IT environment, security must be addressed in every possible area of network design and planning, including such diverse areas as the Active Directory® directory service, networking, and client configuration. Even then, the best security plans and designs in the world cannot protect an organization if security is not an essential part of their operating procedures.

Every organization has its own unique mix of clients, servers, and user requirements that make planning a comprehensive, secure environment a major challenge. Without a consistent approach to security, some areas of the network might benefit from extremely rigorous security while others are only minimally secured.

Fortunately, the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; and Windows® Server 2003, Web Edition operating systems, and the Microsoft® Windows® XP Professional operating system provide many features and capabilities that you can use to configure and maintain a secure network operating environment. In fact, there are security capabilities in nearly every area of Windows Server 2003 and Windows XP Professional. Many of these security features and capabilities have been added or enhanced since the introduction of the Microsoft® Windows® 2000 Professional and Windows® 2000 Server operating systems.

This Windows Server 2003 security planning process is based on a high-level view of the security configuration options and capabilities. This security planning process is based on two organizing principles:

- **Users need access to resources.** This access can be very basic, including only desktop logon and the availability of access control lists (ACLs) on resources. This access can also include optional services such as remote network logons, wireless network access, and access for external users, such as business partners or customers.

- **The network requires a secure shared IT infrastructure.** This infrastructure includes security boundaries, secure servers and services, secure networking, and an effective plan for delegating administration.

Together, these building blocks of network operating system security can provide the trust and integrity needed in today's complex operating environments. By analyzing the security requirements of your organization by using a security planning process, you can establish a high-level security framework for your Windows Server 2003 deployment.

**Important**

- This security planning process is not intended to replace a detailed assessment of existing security systems, gaps, and solutions. For more information about designing and deploying a secure Windows Server 2003 environment you need to complete the assessment, design, and planning steps that are described throughout the *Microsoft® Windows® Server 2003 Deployment Kit*. You might also have to assess the non-operating system software, such as messaging, database, and office suite programs, that are used in your environment.

# Addressing User-Related Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

User-related requirements are essential considerations in network design. There are security requirements associated with almost every user-related design decision that you make. For example, who are your users? Do they log on from a computer that is connected to a network? Do they connect to the network while traveling or from home? Are they customers or business partners? Most importantly, how do you verify that your users are who they say they are?

Many chapters in the *Microsoft® Windows® Server 2003 Deployment Kit* require you to assess user and network requirements before you design or plan your implementation of the technologies and capabilities they cover. You should compile and evaluate these requirements across functional lines so that decisions you make in one area do not conflict with your requirements in other areas.

The following items are key security-related user requirements that every organization must address:

- Keyboard logons
- Access to resources

Other user-related requirements are not as universally applicable, though most organizations also implement security-related requirements in the following areas:

- Remote network access
- Wireless network access
- Standard Client configurations
- Encrypting File System (EFS)
- Certificates for secure e-mail and smart card logons
- Extranet access

# Keyboard Logons

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The most fundamental user-related security requirement is authentication of users who log on to the network from computers that are physically connected to the network. Authentication technology in the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition operating systems allows you to implement a variety of authentication strategies. These strategies can be based on factors such as the complexity of your organization, the quality of a user's credentials, how users access the network, and the clients that they use to gain access.

To plan for the secure use of user and computer accounts based on Windows Server 2003 authentication technologies, use the following guidelines:

- If you have implemented Active Directory in your environment, you can use it to manage users and computers across domains and forests. Deploying Active Directory enables the use of the Kerberos V5 protocol for authentication. This provides the highest level of security and manageability.

- If you are using UNIX, you can use a UNIX Kerberos Key Distribution Center (KDC) to provide authentication services for a realm. Although it is potentially as secure as an Active Directory environment, configuring interoperability with UNIX Kerberos realms is more complex to configure and maintain.

- If you have not upgraded to an Active Directory environment, you can rely on the Security Accounts Manager (SAM) and NTLM to authenticate local users. This option is not as secure as using Active Directory and the Kerberos V5 protocol to authenticate users.

You can also meet differing security needs by varying the strength of credentials that you enforce:

- You can rely on passwords, and define password complexity requirements and password expiration requirements.

- For greater credential strength, you can issue computer- or smart card–based certificates to users and computers. Although these certificates are more difficult to configure and administer, they make it possible to use credentials that are much more complex and secure than any user-conceived password.

For more information about creating a basic authentication framework, see "Designing an Authentication Strategy" in this book. For more information about using certificates to enhance the authentication process, see "Designing a Public Key Infrastructure" in this book. For more information about using smart cards for authentication, see "Planning a Smart Card Deployment" in this book.

# Access to Resources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There is little reason to authenticate users and computers unless you also intend to allow them to access resources on the network. Few organizations allow all security principals, no matter how carefully authenticated, to access all resources on the network. Therefore, you need a strategy to manage access to resources.

The first part of this process is identifying your users by workgroup, job function, or a combination of workgroup and job function. You can then identify the different types of resources that users might access, such as departmental or job-specific data. This step allows you to plan for different levels of access. In Windows Server 2003 and Windows XP Professional, there are four basic strategies for controlling access to resources:

- **User/ACL method.** Security principal accounts are added directly to the ACL for a resource. The ACL is then configured to allow the appropriate permissions for each security principal. This option is appropriate only for managing basic access to a limited number of resources for a small number of users. Otherwise, the complexity and administrator time needed to set and maintain the appropriate permissions becomes burdensome.

- **Account group/ACL method.** Security groups, rather than individual security principal accounts, are added to the resource ACL, and the group is given a set of access permissions. The Account group/ACL method is scalable because security groups can be nested if your domain is configured for Windows 2000 or Windows Server 2003 native mode functional levels. If groups from multiple domains or from multiple forests require identical access permissions, they can be grouped together into one all-encompassing security group, and you can add this security group to the resource ACL. However, if different groups require different access permissions, this method requires more administrative effort from the resource owner. In this case, each group must be added to the ACL separately, and access permissions must be granted to each group.

- **Account group/resource group method.** Users with similar access requirements are grouped into account groups, and the account groups are added to a resource group that is granted specific resource access permissions. The Account group/resource group method is often most appropriate for large organizations with many shared resources because it is scalable and can be maintained in all environments and at all domain functional levels. This method can be used in any environment, but it is most useful when you cannot nest groups because, for example, your domains are at Windows 2000 mixed functional levels or you need to share resources with Microsoft® Windows NT® version 4.0 operating system domains. The Account group/resource group method is also preferable if your resources require a stable set of common permissions.

- **Role-based authorization method.** Users with similar roles are authorized to perform predefined sets of tasks based on scripts called authorization rules. This allows you to apply fine-grained control over the mapping between access control and the structure and tasks performed in your organization. To use role-based authorization, however, programs need to be written to take advantage of this capability. Additionally, your environment must be operating at the Windows Server 2003 domain functional level.

For more information about establishing and deploying an access management strategy, see "Designing a Resource Authorization Strategy" in this book. For more information about role-based authorization, see the Microsoft Platform SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Your authorization strategy is also a factor in other design and deployment processes, such as server deployments. For more information about deploying file and print servers, see "Designing and Deploying File Servers" and "Designing and Deploying Print Servers" in *Planning Server Deployments* in this kit.

# Remote Network Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A growing number of employees require remote network access. Many of these employees are traveling and might have different levels of connectivity from one day to the next. Other employees simply find it efficient to work at a customer location or at home, at least some of the time. Configuring each user's computer for remote network access or relying on users to configure their own computers for remote network access can lead to heavy support responsibilities.

To address this problem, you can use Connection Manager to customize a self-installing service profile for the Connection Manager client. The Connection Manager client allows your users to either connect to your network directly or to create a virtual private network (VPN) connection from a remote location. This solution reduces administration by providing a single connection client for all remote access users. Additionally, it allows you to configure precise client security settings, such as idle timeout time, maximum session time, and encryption strength.

**Note**

- To provide this service, you need to configure remote access to support a dial-up network, a virtual private network (VPN), or a combination of both. For more information about these options, see "Securing Network Traffic" later in this chapter.

For more information about deploying Connection Manager solutions, see "Deploying Remote Access Clients Using Connection Manager" in *Deploying Network Services* in this kit.

# Wireless Network Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Direct wire-based connections to the network are not always practical, even within the physical premises of an organization. With a wireless local area network (WLAN), users of portable computers can access their network connections from any facility in which the organization has wireless access points.

However, wireless networks have significant security implications:

- Any person with a compatible wireless network adapter can potentially associate with your wireless access points. To counter this security risk, wireless access points must require authentication and authorization of the wireless client before data can be sent to, and received from, the network that is attached to the wireless access points.

- In some situations, you might want to allow limited unauthenticated wireless access. For example, in a corporate environment, you can grant guest access to visitors, such as consultants, by redirecting them to a specific virtual LAN (VLAN) that provides only limited network access, such as access to the Internet.

- Wireless networking signals use radio waves to send and receive information. Anyone within a certain distance of a wireless access point can detect and receive all data sent to and from the wireless access point. To counter this threat, you can encrypt the data that is sent between the wireless clients and the wireless access points by using Wired Equivalent Privacy (WEP). For increased security, you can use certificate-based encryption, which requires both a public key infrastructure (PKI) and that a certificate be installed on each wireless client.

For more information about designing a wireless LAN, see "Deploying a Wireless LAN" in *Deploying Network Services* in this kit. For information about deploying a public key infrastructure to support a wireless network, see "Designing a Public Key Infrastructure" in this book.

# Standard Client Configurations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most IT organizations think of their users as customers of their services. However, users can unintentionally become adversaries of the IT department when they install their own software and then require support for this software, when they tinker with and damage approved software programs, when they inadvertently allow viruses onto their computers and, ultimately, the network, and when they fail to protect their data by performing proper maintenance tasks, such as periodic backups.

For these reasons, many organizations consider establishing and maintaining the Standard Client configurations as one of the most significant security and cost-related challenges that they face. Consider the following questions:

- How do you define and implement secure client and server configurations?
- How do you deploy, update, and repair user programs?
- How do you update server and client operating systems with required service packs and patches?
- How do you ensure that user data is adequately protected against potential hardware failures?
- How do you prevent unauthorized users from viewing user data, even if they gain physical access to the computer?

Windows Server 2003 provides a number of features that you can use to enhance security by implementing and maintaining Standard Client configurations. These include:

- **IntelliMirror user settings management.** This feature can be used to define and implement a wide variety of user security settings.
- **IntelliMirror software installation and maintenance.** This feature can be used to deploy, update, and repair user programs.
- **Software Update Services.** This feature can be used to update the operating system itself with required service packs and patches.
- **IntelliMirror user data management.** This feature can be used to protect user data by storing it on a network server.
- **Encrypting File System.** This feature can be used to prevent unauthorized users from viewing data, even when they have physical access to a computer.

**Defining Secure Configurations**

The way you configure your clients and servers has a critical impact on the protection that you provide for network resources and how vulnerable you are to attacks. Creating secure configurations can be daunting. There are so many security-related options that configuring and applying each individual setting to individual clients and servers is virtually impossible.

To address this issue, Windows 2000 and Windows Server 2003 provide a wide variety of Group Policy–based security policy settings that can be configured once and applied at the same time to hundreds or thousands of servers and clients. These servers and clients must be running Windows Server 2003, Windows 2000, or Windows XP Professional. You can use Group Policy to deploy and manage the following types of security policies:

- **Password policies.** These policies are used to enforce restrictions that relate to password age, length, and complexity.
- **Account lockout policies.** These policies define options for the number of unsuccessful logon attempts that are allowed.
- **Audit policies.** These policies are used to track account logon, object access, and management events.

- **User rights assignments.** These policies are security options that apply to user accounts.
- **Security options.** These policies enable you to set and enforce well-known security-related registry options without editing the registry.
- **Public key policies.** These policies allow you to enable or prevent certificate auto-enrollment and renewal.
- **Encrypting file system policies.** These policies allow you to establish a data recovery agent for files and folders that have been encrypted by using EFS.
- **Software restriction policies.** These policies are used to protect computers from code that is not trusted by identifying and specifying which programs are allowed to run.
- **Internet Protocol security (IPsec) policies.** These policiesare used to configure IPsec services for authenticating or encrypting network traffic. An IPsec policy consists of a set of security rules, and each security rule consists of an IP filter with an action.

Setting and applying security policies to large numbers of computers by using Group Policy is much easier than setting security policy on individual computers. To simplify the task of selecting security policy settings, Windows Server 2003 comes with a number of security policy templates that combine groups of security policy settings for different scenarios. The following server security policy templates are available:

- **Setup Security.inf.** This policy template provides the default security settings that are applied during installation of the operating system, including the file permissions for the root of the system drive.
- **DC Security.inf.** This policy template provides the default security settings for domain controllers.
- **Compatws.inf.** This policy template changes the default file and registry permissions that are granted to users. It also removes all members of the Power Users group.
- **Secure.inf.** This policy template definesenhanced security settings that are least likely to impact application compatibility and restrict anonymous users. This policy template also limits the use of LanManager and NTLM authentication.
- **Hisec.inf.** This policy template includes supersets of the secure templates that impose further restrictions on the levels of encryption and signing required for authentication and for the data that flows over Secure Channels and between SMB clients and servers. It also removes all members of the Power Users group.
- **Rootsec.inf.** This policy template provides system root security, and defines the new permissions that were introduced with Windows XP Professional for the root of the system drive.
- **Notssid.inf.** This policy template removes the Terminal Server user security identifier (SID).

**Note**

- Similar client-specific versions of the Setup, Compatible, Secure, and Highly Secure templates are available for computers that are running Windows XP Professional.

You can apply these templates without modification to servers and clients in your organization. You can also modify a few settings on these templates and then apply the templates, or you can create your own security policy templates that contain only the security policy settings that you want.

Many organizations prefer to use the most secure policy templates that still allow all of their required programs and usage scenarios to succeed. Whichever option you select, you need to test your security policy settings thoroughly to ensure that they meet your security requirements without unduly restricting the ability of your users to complete their tasks.

For more information about selecting and deploying security policy settings and security policy templates that are appropriate

for your organization, see "Deploying Security Policy" in *Designing a Managed Environment* in this kit.

**Managing Application Software**
Many organizations rely on tens or even hundreds of software programs to conduct their business. Software administrators have to ensure that users have continuous, reliable access to these programs, including deploying new programs when they become available, updating them on a timely basis, and removing them when they are no longer needed.

At the same time, most administrators want to prevent users from installing or running unknown or potentially dangerous software. At the most benign level, users might merely be installing games or other software that distracts them from their assigned tasks. However, even games or other "harmless" software can at times corrupt a system or require help desk support. At the most dangerous level, unauthorized programs might release a harmful virus or open a security hole in the network.

To make matters more complex, software management is not a solution that can be configured once for all types of programs and all types of users. Certain programs, such as antivirus and e-mail programs, are needed by every user. Other programs might be needed only by users in certain departments or by users who perform certain tasks. Some users, such as frequent travelers, may only be connected to the network intermittently at varying speeds, and they might want to control when large programs are downloaded.

IntelliMirror software management allows you to address these concerns by doing the following:

- Mandating the installation of certain software on every computer.
- Requiring other programs only on certain computers.
- Making other programs available for installation on an as-needed basis or when the user has an adequate connection and time to complete the installation.
- Automatically repairing or updating programs as needed.

Base your software management options on the requirements, locations, and connectivity that are available to your users. For more information about selecting and configuring appropriate software management options for your organization, see "Deploying a Managed Software Environment" in *Deploying a Managed Environment* in this kit.

**Managing Operating System Updates**
Hostile users are constantly looking for vulnerabilities in operating systems. As new vulnerabilities are exposed and solutions become available, it is up to network administrators to ensure that patches and service packs are installed on a timely basis.

Maintaining operating system software through the deployment of patches and service packs has become an increasingly challenging part of organizational security strategies. Until recently, organizations had to check the Windows Update Web site for operating system patches, and then download, test, and distribute patches manually.

A new Windows Server 2003 feature, Microsoft Software Update Services (SUS), helps you collect, approve, and distribute critical operating system patches to resolve known security vulnerabilities and stability issues. You can use these services on computers that are running Windows 2000, Windows XP Professional, and Windows Server 2003.

As an alternative, you can use Microsoft® Systems Management Server (SMS) with an SUS Feature Pack to deploy and manage software patches.

When you choose whether to use Windows Update, SUS, or SMS, consider the following:

- For home offices and small businesses, manual deployment of operating system patches using Windows Updates remains the recommended solution.
- For medium-size organizations (one or more locations with up to 500 users), SUS is the recommended method for distributing operating system patches and service packs.
- For large organizations, SMS is the recommended software-management solution. However, SUS can also work in these environments if the following circumstances exist:
    - You do not need to support clients that are running Microsoft Windows NT 4.0.
    - You need a deployment method only for patches that are not critical updates.

- The needs of your organization are met by SUS and do not require the more comprehensive reporting tools that are included in SMS.

**User Data Management**

Users typically do not back up their data, so if a user has a computer failure or works on a different computer, his or her data is no longer available. IntelliMirror data management allows you to protect and improve the availability of users' data resources by storing them on a network server, where they can be backed up regularly and made available when a user logs on to any computer on the network.

Typically, users in larger organizations have a wide range of data management needs in addition to the data management needs of the organization itself. Users also typically have user settings that must be available to them wherever they log on to the network. Consider not only the needs of the organization, but also the different work needs of different types of users before you design and deploy services that define the users' computing environment.

To develop a plan for managing users' data and settings, categorize users according to their job types and computer experience. Determine if users need to be able to work from any computer on the network and if they need to retain access to network-based files when disconnected from the network.

When planning a deployment, determine the following:

- The number of users to be accommodated. The allocation of resources, such as disk space and network bandwidth varies depending on the number of users.

- The primary programs and computing services in use in the organization, such as databases and e-mail programs. Databases are stored on network shares and do not move with users. However, user settings typically must be available to users wherever they log on to the network. E-mail programs might store users' e-mail folders on network shares or on the local computer. In either case, these folders must remain available.

- If mobile users switch from one computer to another or each user uses a single portable computer in different locations. A user who switches between computers might benefit from Roaming User Profiles more than a user who works exclusively on a single computer.

- The security requirements of different users, programs, and data files. Security involves both access controls and backup requirements.

For more information about planning and deploying IntelliMirror data management, see "Planning a Managed Environment" in *Deploying a Managed Environment* in this kit.

# Encrypting File System (EFS)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In many organizations, a high percentage of users travel with portable computers that they use in customer facilities, airports, hotels, and at home. This means that valuable data is often beyond the control of the organization. If the user's computer is stolen, a hostile user can bypass many aspects of the operating system's security and possibly gain access to sensitive data. The same problem can exist even with desktop computers that are shared among multiple users.

One solution to this problem is to encrypt sensitive files by means of EFS, which adds another layer of security. When files are encrypted, user data is protected even if an attacker has full access to the computer's data storage. Users can encrypt files if they have permission to modify the files. To use EFS, all users must have EFS certificates. These certificates can be self-signed or they can be issued and managed by using a PKI. Businesses that have a PKI and use it to deliver EFS certificates to users can significantly simplify the manageability and recovery of data that is encrypted with EFS if the key becomes lost.

If EFS is not appropriate in your environment, or if you have files that you do not want to be encrypted, you can partially or completely disable EFS in various ways. There are also a number of ways in which you can configure EFS to meet the specific needs of your organization.

For more information about deploying EFS, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Securing E-mail

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many organizations place a high priority on implementing more secure mail services that provide confidential communication, data integrity, and nonrepudiation. In today's increasingly interconnected network environments, intruders can monitor mail servers and network traffic to obtain proprietary or sensitive information. You also risk exposure of proprietary and confidential business information when you send mail over the Internet from within your organization. Additionally, on IP networks, anyone can impersonate mail senders by using readily available tools to counterfeit the originating IP address and mail headers. When you use standard Internet mail, you can never be absolutely certain who really sent a message or whether the contents of the message are valid.

Windows Server 2003 allows you to enhance e-mail security by issuing user certificates that prove the sender's identity, the mail's point of origin, and the authenticity of the message. These certificates can also make it possible to encrypt mail to protect the confidentiality of data. However, encryption also results in a greater likelihood that the recipient will not be able to read the document if he or she does not have the needed decryption key.

- To issue e-mail certificates, you either need to contract for this service with an independent certification authority (CA) or configure one or more certification authorities for your own organization by using Certificate Services that are available in Windows Server 2003.

For more information about configuring a network of CAs to issue and manage certificates for secure-e-mail and other programs, see "Designing a Public Key Infrastructure" in this book.

# Securing the Logon Process with Smart Cards

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A growing number of organizations are concerned that conventional password-based logons do not provide sufficient security. They base this concern on potential vulnerabilities posed by:

- Interactive logons by users
- Interactive logons by administrators or other personnel who have access to sensitive systems or data
- Logons by users who are accessing the network from a remote location and computer

Requiring the use of secure encrypted logon credentials that are stored on a smart card can enhance the security of any or all of these scenarios. However, requiring the use of smart cards and smart card readers can impose significant costs and requires substantial administrative support.

Additionally, requiring smart cards for interactive user or administrator logons can be particularly problematic for organizations that are spread across different geographic locations. They can also be problematic for organizations that do not have security personnel in each location to administer and support smart cards.

Another problem that occurs with the use of smart cards for remote access logons is that remote users often own computer hardware and software that does not conform to minimum corporate standards, and therefore might not support smart card use. This complicates your ability to administer and promote the use of smart cards for remote access logons.

Finally, consider that users might experience longer logon times when they use smart cards, especially over slow dial-up connections.

For more information about evaluating and planning for the use of smart cards in your organization, see "Planning a Smart Card Deployment" in this book.

# Setting Up Extranet Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A growing number of organizations find it valuable to share some restricted data with business partners, contractors, suppliers, and customers. You can accomplish this by combining the same networking technologies that you use on your public Internet Web sites and internal intranet with a tight security model to form an extranet. In fact, extranets have become a primary means of interaction between companies that contribute to each other's everyday business.

There are many aspects involved in setting up a secure extranet. For many organizations, the two key steps involve:

- Preventing unauthorized users from accessing data that is not intended for them.
- Encrypting all data that is transferred between the client and the Web server.

The authentication methods that you select vary based on the method's ability to protect the users' credentials (user account and password information). To ensure your users' credentials are protected, select the strongest authentication method possible.

The following authentication methods are used for Web site authentication:

- **Anonymous authentication.** This method gives anonymous users the right to log on locally. You can set up different anonymous accounts for different Web sites, virtual directories, directories, and files. This provides a great deal of flexibility and fine control by specifying where specific accounts can be used within the site.
- **Basic authentication.** This authentication method is supported by most Web servers and Web browsers. With Basic authentication, you can restrict access to files on a server that is running Internet Information Services (IIS) 6.0 by using NTFS file system security. Users must enter credentials, which allows you to track access based on the user ID. You can use this method to restrict access to some parts of the Web server when you cannot guarantee that the user's browser supports integrated Windows authentication, and you must authenticate through a proxy server.
- **Digest authentication.** This method does not send the password in clear text, unlike Basic authentication. Additionally, Digest authentication can work through proxy servers, unlike integrated Windows authentication.
- **Advanced Digest authentication.** This method is a newer version of Digest Authentication that was introduced with Windows XP Professional. In Advanced Digest authentication, the client uses a hash of the password, which is the same as the hash of the password that is stored on the server. Older Web browsers do not support Advanced Digest authentication.
- **Integrated Windows authentication.** This method uses either NTLM authentication or Kerberos authentication and the current credentials of the logged on user to authenticate to the Web server. This method is primarily used for intranet situations.
- **Certificate authentication.** This method can be used to allow both servers and clients to authenticate each other. Server certificates usually contain information about your company and the organization that issued the certificate. Client certificates usually contain identifying information about the user and the organization that issued the certificate.

To provide confidential data exchange, you encrypt the data that is transferred between the client and the Web server. Based on the type of client that might access your program, select one of the following methods for encrypting the data:

- **Secure Sockets Layer (SSL).** This method can be used when you are creating an Internet solution, the client computers are not a part of your organization, and you want to avoid installing any files on the client computers.
- **IPsec.** This method can be used when you are creating an intranet or an extranet partner solution, the client computers are a part of your organization or a partner organization, it is acceptable to install certificates on the client computer, and the client computers have high-speed persistent connections to your private network.

- **VPN.** This method can be used when you are creating an intranet or an extranet partner solution, the client computers are a part of your organization or a partner organization, it is acceptable to install certificates on the client computers, and the client computers are connecting over direct or dial-up Internet connections.

For more information about security and IIS, see "Securing Web Sites and Applications" in the *Internet Information Services (IIS) 6.0 Deployment Guide* in this kit (or see "Securing Web Sites and Applications" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48683). For more information about designing and deploying IPsec, see "Deploying IPSec" in *Deploying Network Services* in this kit. For more information about designing and deploying VPN, see "Deploying Dial-up and VPN Remote Access Servers" in *Deploying Network Services* in this kit. For more information about designing a public key infrastructure and certificates for use with external clients, see "Designing a Public Key Infrastructure" in this book.

# Establishing a Secure Shared IT Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Not all security-related features apply directly to users. Many basic network services and configuration decisions involve creating and defining explicit boundaries, securing network traffic, and securing your servers.

# Creating and Enhancing Security Boundaries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As networks become more complex, they can become more difficult to manage. Windows NT 4.0 introduced the concept of domains to enhance an administrator's ability to manage the users and computers in a network. The domain concept has been enhanced and expanded considerably in Windows 2000 and Windows Server 2003 to address security concerns of organizations, including the following concerns:

- Organizations that have acquired subsidiaries with significantly different administrative and security requirements.
- International organizations that want to divide up administrative and security responsibilities along national or regional boundaries.
- Rapidly growing organizations that want to unify security and administrative responsibilities across disparate business units.

Active Directory in Microsoft Windows Server 2003 enables organizations to simplify user and resource management; support directory-enabled programs; and create a scalable, secure, and manageable infrastructure. A well-designed Active Directory logical structure provides the following benefits:

- Centralized management of Windows networks that contain large numbers of objects
- A consolidated domain structure and reduced hardware and administration costs
- The supporting framework for Group Policy–based user, data, and software management
- The ability to delegate administrative control over resources
- Integration with services such as Microsoft® Exchange 2000 Server, a PKI, and domain-based distributed file system (DFS)

Achieving these results requires careful planning of the following elements:

- **Domains.** An administrative unit in a computer network that groups a number of capabilities for management convenience, including network-wide user identity, authentication, trust relationships, policy administration, and replication.
- **Forests.** One or more Active Directory domains that share a schema and global catalog. Each forest is a single instance of the directory and defines a security boundary.
- **Organizational units.** Active Directory containers where you can place users, groups, computers, and other organizational units. You can use organizational units to create containers within a domain to represent the hierarchical and logical structures within your organization.

The way that you plan your domains, forests, and organizational units plays a critical role in defining your network's security boundaries. The relationship might sometimes be based on administrative requirements; at other times, the relationship might be defined by operational requirements such as controlling replication. Additionally, if you have multiple forests, you need to plan the logical trust relationships between forests that allows pass-through authentication.

**Domains, Forests, and Organizational Units**

Active Directory is a distributed database that stores and manages information about network resources. The way that you organize Active Directory determines how well you can manage network resources and distribute administrative responsibilities.

Active Directory allows administrators to organize elements of a network (such as users, computers, and devices) into a

hierarchical, tree-like structure based on containers. The top-level Active Directory container is the forest. Forests include domains, and domains include organizational units. Administrative ownership and control in Active Directory containers are organized in the following ways:

- The default administrative owner of a forest is the Domain Admins group of the forest root domain. The Domain Admins group of the forest root controls the membership of the Enterprise Admins and Schema Admins groups. By default, the Enterprise Admins and Schema Admins groups have control over forest-wide configuration settings, which also makes them service administrators.

- The default administrator of a domain is the Domain Admins group of the domain. Because the Domain Admins control domain controllers, they are also service administrators. All non-root Domain Admins in a forest are peers, regardless of their domain's position in the naming hierarchy.

- Control over an organizational unit and the objects within it is determined by the ACLs on the organizational unit and on the objects in the organizational unit. Users and groups that have control over objects in organizational units are data administrators.

To facilitate the management of large numbers of objects, Active Directory supports administrative delegation at the container level. If administrative control is the priority for your organization, base your logical structure design on forests and organizational units. Forests and organizational units are used to control the delegation of authority throughout the directory. Many organizations consolidate divisions into a single forest to enhance their users' ability to collaborate and to reduce costs.

If you choose to organize Active Directory according to geographic location, you must apply a domain model to your logical design. Domains let you control where information is replicated and let you partition data so that it can be stored where it is used most frequently. A well-designed domain model prevents unnecessary replication and promotes more efficient use of available bandwidth between remote locations.

To determine the number of forests that your organization requires, identify the isolation requirements for each division of the organization that will be using the directory service. Consider the following elements:

- Generally, a single forest deployment isolates data from parties outside the organization. If your organization includes more than one IT group, the only way to achieve isolation in a single forest environment is to select one IT group to act as the administrators of the forest, and then make the other IT groups in the organization relinquish control of the directory.

- If divisions of your organization require that you isolate data from the rest of the organization, you must deploy multiple forests. For example, you might need multiple forests if legal or contractual obligations require that your organization guarantees the security of data for a particular project.

- If your organization includes multiple divisions with separate IT groups, each IT group might prefer to manage its own forest; however, your business needs might require resource sharing between divisions. You can deploy multiple forests, each of which is managed by an individual IT group, and then establish external trusts between the forests to facilitate collaboration. In this type of environment, be careful to avoid granting administrative access to users in other forests.

**Trusts**

If your organization includes more than one forest, you must enable the forests to allow authentication and resource sharing. You can do this by establishing trust relationships between some or all of the domains in the forests. The types of trust relationships that you can establish depend on the versions of the operating system that are running in each forest:

- **Authentication between Windows Server 2003 forests.** When all domains in two forests trust each other and must authenticate users, establish a forest trust between the forests. When only some of the domains in two Windows Server 2003 forests trust each other, establish one-way or two-way external trusts between the domains that require interforest authentication.

- **Authentication between Windows Server 2003 and Windows 2000 forests.** It is not possible to establish transitive forest trusts between Windows Server 2003 and Windows 2000 forests. To enable authentication with Windows 2000 forests, establish one-way or two-way external trusts between the domains that need to share resources.

- **Authentication between Windows Server 2003 and Windows NT 4.0 forests.** It is not possible to establish transitive forest trusts between Windows Server 2003 and Windows NT 4.0 domains. Establish one-way or two-way

external trusts between the domains that need to share resources.

For more information about designing Active Directory forests, see "Designing the Active Directory Logical Structure" in this book. For more information about implementing trusts between forests, see "Designing an Authentication Strategy" in this book.

# Public Key Infrastructure (PKI)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A growing number of organizations are choosing to use digital certificates for programs, including the following:

- Secure e-mail
- Software code signing
- Internet authentication
- IP security
- Smart card logon
- Encrypting File System user and recovery certificates
- Wireless (802.1x) authentication

Before you can enable these programs, you must contract with a third-party authentication service or use Windows Server 2003 to configure one or more CAs as a PKI to issue and manage your digital certificates.

An outside authentication service might be a good option if you plan to use your certificates outside of your own organization, but they can also be expensive and relatively inflexible.

If you plan to use your certificates only within your organization or with a limited number of business partners with whom you share resources, setting up and maintaining an internal PKI can be both less expensive and open to modification as your needs change.

A Windows Server 2003 PKI makes it possible for an organization to perform the following tasks:

- **Publish certificates**. The PKI administrator makes certificate templates available to clients (users, services, programs, and computers) and enables additional CAs to issue certificates.
- **Enroll clients**. To participate in a PKI, users, services, and computers must request and receive certificates from an issuing CA. Typically, enrollment is initiated when a requester provides unique identifying information and a newly generated public key. The CA administrator or enrollment agent uses this unique identifying information to authenticate the identity of the requester before issuing a certificate.
- **Use certificates**. Client certificates are validated while CAs and certificate revocation lists are available to verify or deny their authenticity. If they are validated, a PKI provides an easy way for users to use keys in conjunction with programs that perform public key cryptographic operations. This makes it possible to provide security for e-mail, e-commerce, and networks.
- **Renew or revoke certificates**. A well-designed PKI makes it easy for you to renew or revoke existing certificates and to manage the trust levels that are associated with certificates used by different clients or for different programs.

Because certificates are intended to enhance the quality of a user's or computer's identification, almost every decision you make regarding certificate services — from the number and placement of certification authorities to the specifications for the useful life of each certificate — can have a significant impact on the effectiveness of a PKI as a security solution. For more information about setting up and maintaining a PKI, see "Designing a Public Key Infrastructure" in this book.

# Securing Network Traffic

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Hostile users have become increasingly sophisticated in monitoring and compromising network traffic. The tools that are available to assist them in their activities have also become more powerful every year. At the same time, most organizations want to open up their networks to support new methods of doing business that are enabled by telecommunications. Supporting these business scenarios without compromising security requires the use of very critical networking technologies, including the following:

- Remote access servers
- IPsec
- Internet Authentication Service (IAS)
- Microsoft Internet Security and Acceleration (ISA) Server

A consistent and integrated plan for using these features along with other components of the networking infrastructure can enhance the security of your environment.

**Remote Access Servers**

A growing number of users require access to the network when they are away from the organization's physical premises. However, opening up your network to the outside world poses potential security risks.

To support users who require access to your network from remote locations without compromising security, you can deploy a dial-up network, a virtual private network (VPN), or a combination of both.

- A *dial-up network* enables remote users to dial in directly to a remote access server on your network.
- A *VPN* enables remote users who are connected to the Internet to establish a connection to a VPN server on your network through a VPN connection.

Dial-up connections are inherently more private than a solution that uses a public network such as the Internet. However, with dial-up networking, your organization faces a large initial investment and continuing expenses throughout the life cycle of the solution. These expenses include:

- **Hardware purchase and installation.** Dial-up networking requires an initial investment in modems or other communication hardware, server hardware, and phone line installation.
- **Monthly phone costs.** Each phone line that is used for remote access increases the cost of dial-up networking. If you use toll-free numbers or the callback feature to defray long distance charges for your users, these costs can be substantial. Most businesses can arrange a bulk rate for long distance, which is preferable to reimbursing users individually at their more expensive residential rates.
- **Ongoing support.** The number of remote access users and the complexity of your remote access design significantly affects the ongoing support costs for dial-up networking. Support costs include network support engineers, testing equipment, training, and help desk personnel to support and manage the deployment. These costs represent the largest portion of your organization's investment.

If you use a VPN for remote access, users connect to your corporate network over the Internet. VPNs use a combination of tunneling, authentication, and encryption technologies to create secure connections. VPNs reduce remote access expenses by using the existing Internet infrastructure. You can use a VPN to partially or entirely replace your centralized, in-house, dial-up remote access infrastructure and legacy services.

VPNs offer two primary benefits:

- **Reduced costs.** Using the Internet as a connection medium saves long-distance phone expenses and requires less hardware than a dial-up networking solution.

- **Sufficient security.** Authentication prevents unauthorized users from connecting to your network. Strong encryption methods make it extremely difficult for a hostile party to interpret the data that is sent across a VPN connection.

For more information about deploying dial-up remote access server and VPN solutions, see "Deploying Dial-up and VPN Remote Access Servers" in *Deploying Network Services* in this kit.

**IPsec**

Although business networks typically separate their internal networks from the Internet by using a limited number of security gateways, internal corporate networks have become so complex that it is difficult to know that all critical data is protected at all times from external and internal attackers. Using a firewall to block incoming traffic by port is not always practical because it is difficult to identify all ports that require protection. IPsec provides a much stronger access model based on enforcing trusted communication.

You can use the Windows Server 2003 implementation of IPsec to compensate for the limited protection for network traffic provided by programs, or as a network-layer foundation of a defense-in-depth strategy. While the IPsec features in Windows Server 2003 provide greater security, they are not always compatible with earlier versions of Windows.

The first step in deploying IPsec on your network is to determine which data and which computers you need to secure and at what level of security. Although adding IPsec security can protect your network and data, it can also produce tradeoffs in performance, and some computers might not support the security you require.

For more information about planning and implementing IPsec in your environment, including how to choose between IPsec and other solutions, see "Deploying IPSec" in *Deploying Network Services* in this kit.

**IAS**

Not all users can be directly authenticated by Active Directory. For example, organizations that outsource network access or form joint ventures with other organizations require the authentication of user accounts from outside their private network. In addition, organizations that provide outsourcing services, such as Internet service providers (ISPs), face the increasingly complex challenge of managing all types of network access from a single point of administration, regardless of the type of network access equipment that is used.

IAS allows you to provide authentication and authorization for the following types of network access:

- Dial-up

- VPN

- Wireless

- Switch access

IAS also allows you to apply different types of security policies to different groups based on their mode of access and quality of credentials.

For more information about planning an IAS deployment, see "Deploying IAS" in *Deploying Network Services* in this kit.

**ISA Server**

The Internet is a tremendous source of information for the employees of many organizations. News updates, background information about customers, product research, and other activities are just a few of the business activities that are enhanced by the Internet.

Allowing employees to access the Internet poses significant reliability and security risks:

- **Reliability.** Network bandwidth might be compromised if too many users access the Internet at the same time.

- **Security.** Allowing users to access resources outside the organization's network can create a potential breach in the network's perimeter defenses.

ISA Server allows you to address these concerns with two important capabilities:

- **A multilayer firewall with integrated intrusion detection**. ISA is a complete firewall product that provides packet filtering, stateful packet inspection and application-level awareness for many common protocols, such as Simple Mail Transfer Protocol (SMTP) and Domain Name System (DNS). You can create access policies based on user-level information, IP addresses, or Web content.

- **A caching server.** You can configure ISA Server to cache content that your organization uses frequently or that your Internet clients access.

ISA Server can be configured to perform one or more of these tasks. In addition, it can be configured for the following uses:

- To connect to remote sites

- To secure network perimeters

- To support extranets

For more information about these design and deployment options, see "Deploying ISA Server" in *Deploying Network Services* in this kit.

# Securing Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In general, servers are only as secure as the configuration options that you enable or disable. The security of the configuration options depends on who can access the servers. A secure server configuration for one organization might not be secure for another organization. Whether security is adequate depends on a variety of factors, such as where the server is on the network, what data is stored on the server, what role the server plays, and who administers it. One point is certain, however: securing servers requires painstaking attention to detail based on strict evaluations of what services and configuration options are needed and not needed, and testing to ensure that access control mechanisms perform without preventing the server from performing its legitimate tasks.

The *Windows Server 2003 Deployment Kit* includes detailed information about how to configure and administer servers securely for a variety of specific roles. In addition, you can enhance general server security by doing the following:

- Whenever possible, keep critical servers offline or unavailable to the Internet.
- Require administrators to use smart cards when they access your servers.
- Implement a strategy for auditing all activities on critical servers.
- Create standard server configurations and require their use whenever a server is deployed to fill a similar role, even in remote offices or in subsidiaries.

# Planning for Secure Administration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Just as organizations need to configure and manage access by users to files, shares, and other resources, you also need to delegate administrative tasks in Windows Server 2003 environments. Delegating these tasks allows you to distribute the administrative workload without granting sweeping privileges to each person. From a high-level perspective, your decisions regarding delegation of administration will be based on the following four factors:

- Isolation: The degree to which administrative roles need to be separated.
- Autonomy: The limits that you want to place on the ability of any administrator or group to act unchallenged.
- Service administration: The authority to manage key network services.
- Data administration: The authority to manage key data resources.

For example, you might choose to delegate administration to meet one or more of the following needs:

- **Organizational structure requirements**. Parts of an organization might participate in a shared infrastructure to save costs, but they might also want the ability to operate independently from the rest of the organization. For example, a research division within a large organization might want to maintain control over all of its own research data.
- **Operational requirements**. One part of an organization or a program that leverages Active Directory might place unique constraints on directory service configuration, availability, or security. Examples are found in:
    - Military organizations
    - Hosting scenarios
    - Organizations that maintain a directory that is available both internally and externally (publicly accessible by users on the Internet)
- **Legal requirements**. Some organizations have legal requirements that require them to operate in a specific manner, such as restricting access to certain information. These requirements commonly apply to the following types of organizations:
    - Financial institutions that need to maintain private client financial records
    - Defense contractors working on classified military projects
    - Government organizations that maintain top secret data

For information about delegation of administration, see "Designing the Active Directory Logical Structure" in this book. For information about how to use security groups to delegate administrative tasks, see "Designing a Resource Authorization Strategy" in this book. For information about deploying smart cards for administrative logons, see "Planning a Smart Card Deployment" in this book.

# Additional Resources for Planning a Secure Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* in this kit.
- "Planning an Active Directory Deployment Project" in this book.
- "Planning a Managed Environment" in *Designing a Managed Environment* in this kit for more information about Group Policy-related features, and planning and deploying IntelliMirror data management.
- "Designing a Public Key Infrastructure" in this book for more information about using certificates to enhance the authentication process.
- "Planning a Smart Card Deployment" in this book.
- "Designing and Deploying File Servers" in *Planning Server Deployments* in this kit.
- "Designing and Deploying Print Servers" in *Planning Server Deployments* in this kit.
- "Deploying Security Policy" in *Designing a Managed Environment* in this kit for more information about selecting and deploying security policy settings and security policy templates that are appropriate for your organization.
- "Deploying a Managed Software Environment" in *Deploying a Managed Environment* in this kit for more information about selecting and configuring appropriate software management options for your organization.
- "Deploying IPSec" in *Deploying Network Services* in this kit.
- "Deploying Dial-up and VPN Remote Access Servers" in *Deploying Network Services* in this kit.
- "Designing an Authentication Strategy" in this book for more information about implementing trusts between forests.
- "Deploying IAS" in *Deploying Network Services* in this kit.
- "Deploying ISA Server" in *Deploying Network Services* in this kit.
- "Designing the Active Directory Logical Structure" in this book for information about delegation of administration.
- "Designing a Resource Authorization Strategy" in this book for information about how to use security groups to delegate administrative tasks.
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about Encrypting File System.
- "Securing Web Sites and Applications" in *Deploying Internet Information Services (IIS) 6.0* in this kit (or see "Securing Web Sites and Applications" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48683) for more information about security and IIS
- The Microsoft Platform SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about role-based authorization.

# Designing an Authentication Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most organizations need to support seamless access to the network for multiple types of users, such as workers in offices, employees who are traveling, and perhaps even business partners and customers. At the same time, organizations need to protect network resources from potential intruders. A well-designed authentication strategy can help you achieve this complex balance between providing reliable access for users and strong network security for your organization.

**In This Chapter**


Overview of the Authentication Strategy Design Process

Creating a Foundation for Authentication

Securing the Authentication Process

Extending Your Authentication Framework

Enabling Supplemental Authentication Strategies

Educating Users About the Authentication Process

Additional Resources for Designing an Authentication Strategy


**Related Information**

- For more information about the Kerberos version 5 authentication protocol, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

- For more information about the Active Directory® directory service logical structure, see "Designing the Active Directory Logical Structure" in this book.

- For more information about upgrading from the Microsoft® Windows NT® version 4.0 operating system to the Microsoft® Windows® Server 2003 operating system, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.

# Overview of the Authentication Strategy Design Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

One of the most fundamental elements of an organization's security strategy is verifying the identity of clients and granting them appropriate access to system resources based on their identity. By creating an authentication strategy for your organization, you can prevent attackers and malicious users from accessing and tampering with sensitive information, consuming computing power or other system resources, and impersonating users in order to send misleading or incorrect information.

Authentication technology in the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems allows you to implement a variety of authentication strategies based on the complexity of your organization, the quality of a user's credentials, the means by which users access the network, and the clients they use to gain access. In addition, Windows Server 2003 authentication technology allows you to establish a foundation for more efficient management of users, computers, and services on the network.

**Note**

- For a list of the job aids that are available to assist you in designing an authentication strategy, see "Additional Resources for Designing an Authentication Strategy" later in this chapter.

# Process for Designing an Authentication Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing an authentication strategy involves evaluating your existing infrastructure and creating accounts, establishing a means to secure the authentication process, and establishing standards for network authentication and time synchronization. You might also need to extend your authentication model to allow authentication between forests or between other Kerberos realms, and to enable delegated authentication in order to facilitate user access to system resources.

Figure 14.1 shows the process for designing an authentication strategy.

**Figure 14.1 Designing an Authentication Strategy**



# Authentication Background Information

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 authentication technology includes a number of features that provide solutions for a wide variety of business needs.

**Central administration of accounts**
Administrators can create a single account for each user that allows the user to access the appropriate network resources. Users can log on at different desktops, workstations, or notebooks in the domain by using the same user name and a password or smart card.

**Single sign-on environment**
Users are required to enter a user name and password or smart card only when first logging on to a Windows Server 2003–based computer. The Windows Server 2003 operating system automatically authenticates the user to the local computer, to the Active Directory domain, and to any other application or resource server in the forest that requires authentication prior to access.

When users change passwords, the updates are made to the user accounts in Active Directory. The password changes apply automatically to all resources in the domain or forest.

**Computer accounts in Active Directory**
Computer accounts in Active Directory for all of the computers within a domain allow many of the Windows Server 2003 security features that are designed for users to be applied to computers as well.

Computer accounts in Active Directory also allow you to add application servers as member servers within your trusted domains and to demand authentication from the users and other services that access these resource servers.

**Service accounts in Active Directory**
The services running on resource servers are authenticated automatically if the servers are members of a domain that trusts the user's account domain. In Windows Server 2003, all of the domains in a forest automatically have two-way transitive trust. Windows Server 2003 also supports transitive trust relationships between forests. In this way, when organizations add application servers to their domains, only authenticated users and services can access them.

**Smart card support**
Windows Server 2003 supports optional smart card authentication. A smart card contains a processor chip that stores the user's private key and public key certificate. The user inserts the card into a smart card reader attached to the computer. The user then types in a personal identification number (PIN) when requested, to enable access to the keys stored on the smart card. Authentication proceeds when the correct PIN enables access to the private key and the certificate on the card, allowing the Active Directory authentication service to verify the user's identity. In this way, computers that store highly sensitive data can be secured from attack without the need to store them in locked rooms. At the same time, authorized users can access information stored on high-security computers.

**Certification for Microsoft Windows**
Windows Server 2003 authentication interoperates with third-party applications designed according to the Application Specification for Windows 2000. The Application Specification defines the technical requirements for applications to earn the Certified for Microsoft Windows logo.

Applications can carry the Certified for Microsoft Windows logo when they have passed compliance testing and have executed a logo license agreement with Microsoft. To pass compliance testing, a server application must operate within the appropriate security context, reducing the risk posed by successful attacks, and perform Kerberos-based mutual authentication for all client requests, ensuring that clients know that the servers with which they are communicating are the intended parties, and not attackers posing as the server.

**Auditing authentication**
Windows Server 2003 provides security audit information to track attempts to log on to servers and workstations. This gives organizations the ability to detect unauthorized attempts to access the system.

**Kerberos V5 authentication protocol**
When a client attempts to connect to a resource server, the Kerberos Key Distribution Center (KDC), running on a domain controller, provides the client with a ticket to verify the user's identity to the server, and a shared secret key. The ticket allows the server to validate the user immediately and can be used multiple times. The shared secret key is passed to the server in

encrypted form, allowing both computers to use the shared secret key to encrypt any network data they exchange.

The Microsoft implementation of the Kerberos authentication protocol is based on industry standard specifications defined by the Internet Engineering Task Force (IETF). The Kerberos V5 authentication protocol provides the following advantages:

- **Efficient authentication to servers.** Because authentication takes place quickly, users do not lose productive work time. Clients can obtain a ticket for a particular server one time and reuse the ticket for multiple network sessions.

- **Mutual authentication.** By means of the shared secret key, parties at both ends of a network connection can verify each other's identities. This is a change from NTLM, which allows only servers to verify the identities of their clients.

- **Delegated authentication.** A service can impersonate a client when connecting to a network service, such as a database. Delegated authentication is not available in NTLM.

- **Interoperability**. Kerberos authentication in Windows Server 2003 can interoperate with the implementation of Kerberos authentication in other operating systems.

# Tools for Deploying Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following tools are available to assist you in deploying authentication:

- **Active Directory Users and Computers.** A Microsoft Management Console (MMC) snap-in that allows you to create user and computer accounts in the Active Directory.

- **Group Policy.** An MMC snap-in that allows you to apply authentication Group Policy, including Kerberos, auditing, and NTLM authentication behavior.

- **Certificate Services.** An MMC snap-in used to establish the certification authorities (CAs) and issue the certificates used in public key authentication.

For more information about these tools, see the *Distributed Services Guide* of the *Windows Server 2003 Resource Kit* (or see the Distributed Services Guide on the Web at http://www.microsoft.com/reskit).

# Creating a Foundation for Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you deploy Windows Server 2003, you can create a foundation for secure authentication of users, computers, and services in your organization by creating accounts in Active Directory for all entities that require authenticated access to resources. Because a number of factors impact the authentication strategy that you deploy, you must evaluate the structure of your existing environment before you create the foundation for your authentication strategy.

For a worksheet to assist you in creating a foundation for authentication, see "Authentication Strategy Planning" (DSSAUT_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Authentication Strategy Planning" on the Web at http://www.microsoft.com/reskit).

Figure 14.2 shows the process for creating a foundation for authentication.

**Figure 14.2 Creating a Foundation for Authentication**



# Evaluating Your Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you establish an authentication strategy for your organization, you must become familiar with your current environment, including the structure of your organization; the users, computers, and services in your organization that require authentication; and the applications and services that are in use. Specifically, identify the following:

- The number of domain controllers in your organization. Ensure that you have enough domain controllers in your environment to accommodate your users' authentication requests. If the number of domain controllers is insufficient, a large volume of client requests can result in failed authentication attempts. If you determine that you have an insufficient number of domain controllers, deploy more domain controllers to meet the logon needs of your users.

- The type of network connectivity between site locations in your organization. Domain controllers must be well connected to users to ensure reliable access for authentication. Clients that do not have access to local domain controllers might be unable to access resources if the network connection is unavailable. If the connectivity between domain controllers in remote sites is insufficient, deploy more domain controllers in those sites or improve the connectivity between the sites.

- The number of CAs that are available in your organization and their locations. As with domain controllers, a sufficient number of CAs must be available to handle client requests and they must be well connected in order to provide timely responses. For information about creating a CA infrastructure, see "Designing a Public Key Infrastructure" in this book.

- The number of users, groups, and computers in your organization and where computers are located. This impacts the number of domain controllers and CAs that are required to ensure consistent authentication.

- The number and locations of users who access the network by means of RADIUS and RAS servers.

  **Note**

  - Windows Server 2003 provides for remote user authentication by means of RADIUS and RAS servers. For more information about using RADIUS servers, see "Deploying IAS" in *Deploying Network Services* of this kit. For more information about using RAS servers, see "Deploying Dial-up and VPN Remote Access Servers" in *Deploying Network Services*.

- Whether your organization includes clients running versions of Windows earlier than the Microsoft® Windows® 2000 operating system or other non-native operating systems, or applications that require authentication protocols other than the Kerberos V5 authentication protocol or require special configuration to interoperate with the Kerberos authentication protocol. The operating systems and applications in use in your environment impact the authentication protocols that you can enable by means of authentication policy. For example, versions of Windows earlier than Windows 2000 require NTLM authentication or anonymous access. If clients in your environment are running these operating systems, you must configure the **LAN Manager authentication level** policy to enable those clients to access resources in your system.

  **Note**

  - When you enable LAN Manager authentication, you cannot take advantage of all of the security benefits that are available in Windows Server 2003. Therefore, if you do not need to support versions of Windows earlier than Windows 2000, it is best to use the Kerberos authentication protocol.

- The number and location of smart card users in your organization, if applicable, and any security-sensitive tasks or users, such as administrators, that might require smart cards in the future. The number of current and planned future smart card users in your organization impacts the number of CAs that you require.

# Creating User Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

User accounts are required for authentication. Assign users the appropriate permissions to access resources by creating user accounts in Active Directory and adding the accounts to the appropriate groups. Adding accounts to security groups and applying access control settings to resources allows users to utilize their authenticated identity to access resources, and facilitates account management.

It is best to grant users and groups access to only those resources that are required for them to complete their job tasks. In this way, if any user account is compromised by a malicious user, he or she has limited access to resources, and therefore can cause only minimal damage.

For more information about user accounts and security groups, see "Designing a Resource Authorization Strategy" in this book.

**Note**

- Do not allow users to share accounts or passwords or to use weak passwords. Shared accounts and weak passwords compromise the security of your environment. For more information about creating password policies, see "Creating a Strong Password Policy" later in this chapter.

Creating user accounts involves creating a plan for user account management in your organization.

# Creating a User Account Management Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you deploy Windows Server 2003 and establish the appropriate user accounts in Active Directory, you need to create a plan for user account management. Creating a user account management plan involves determining which individuals in your organization have the right to create new user accounts, and establishing a plan for the disabling of and resetting of user accounts.

**Assign the User Account Creation Right**

Assigning the right to create new user accounts involves carefully balancing strong security and timely response to requests to create new accounts.

Because misuse of the user account creation right presents a security risk to your organization, assign this right to trusted administrators only. For many organizations, it is sufficient to limit the ability to create new user accounts to the members of the Domain Administrators group. In large organizations or in situations where administrators need to delegate tasks, you might need to assign the right to create new user accounts to another group, such as the IT staff or the Human Resources group.

Whoever you designate to create user accounts, a general guideline is to assign one individual the right to create new user accounts for every 100 employees. However, you might need to adjust this number based on the expected growth of your organization. For example, if your organization regularly adds new divisions, acquires companies, or expands into other markets, you need to plan for the creation of new user accounts by assigning the right to create new user accounts to the appropriate number of individuals to meet the requirements for your anticipated growth.

**Establish a Plan for the Disabling of User Accounts**

Because unused but active user accounts are a common target for security attacks, you must establish a clear, consistent policy for disabling user accounts.

You can choose one of the following solutions for disabling active unused user accounts in your organization:

- Include disabling user accounts as part of the employee departure procedure. Establish a policy by which user accounts are deleted from Active Directory when employees leave your organization.

- Create scripts that search for user accounts that have not been logged on to for a period of time or have not had their password changed, and delete the accounts that the script identifies. For example, you might decide to create a script that identifies accounts that have not been logged on to for six weeks, or that have not had their passwords changed for twice the password lifetime prescribed by domain Group Policy, and delete those accounts.

**Establish a Plan for Resetting User Accounts**

When a user forgets his or her password, the account must be reset before it can be used. An effective way to enable the resetting of user accounts in your organization is to grant help desk staff the right to reset passwords. Delegate the right to reset passwords to help desk staff so that members of the Domain Administrators group are not required to reset user account passwords.

# Creating a Computer Account Management Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows 2000 and Windows Server 2003 computers have accounts in Active Directory and are authenticated in a separate process that is transparent to the user.

You can use computer authentication to apply uniform security policies to groups of computers, such as computers contained in a domain, a site, or an organizational unit (OU) based on how the computers are grouped and which rights and policies are granted and applied to each group. For example, you can configure an OU for computers that are public kiosks on a retail floor and apply limited permissions to users. You can configure another OU for a computer stored in a locked office and allow users greater access to resources.

Evaluate the security needs for different types of computers in your organization. Determine which computers are more vulnerable to compromise and therefore require stronger security settings, and then apply policies to the domains, sites, and OUs as appropriate to your security needs. For more information about applying security policies, see "Deploying Security Policy" in *Designing a Managed Environment* in this kit.

**Managing Computer Accounts**

You also need to establish a plan for managing computer accounts, including:

- The creation of new accounts
- The deletion of old accounts
- Resetting of computer account passwords.

Because new computer accounts are created automatically whenever a computer is added to a domain, you need to decide who has the right to add computers to domains. You can delegate this responsibility to an individual or group in your organization by adding them to the **Add workstations to domain** Group Policy.

You can choose to manage new computer account creation in your organization in one of the following ways:

- **Allow authenticated users to create new computer accounts.** This approach might be desirable in organizations where users can be largely trusted. However, if you only want to trust a limited group of users, such as developers, for example, to create new computer accounts, you can control this by using the Security Configuration Manager to either assign or deny this right to users. By default, authenticated users are assigned the **Add workstations to domain** user right on the Group Policy object on domain controllers. This enables them to create up to 10 computer accounts in the domain by using the Network Identification Wizard. The wizard requests information about the computer name, the domain or workgroup that the computer is joining, and the domain users that are to be added to the local groups for local computer access, and uses this information and the credentials of the authenticated user to create a new account in Active Directory.

  **Note**

  - After a computer account is created, administrators must ensure that the account is a member of the appropriate groups, so that the appropriate Group Policies are applied.

- **IT staff joins each new computer to the domain individually during installation.** Although this approach can work for small organizations in which computer account creation occurs infrequently, it is impractical for large organizations with a high volume of new computer accounts.

- **IT staff uses scripts to create new accounts ahead of time, and assigns new computers to existing accounts during installation.** You can use an Active Directory Service Interfaces (ADSI) script to create computer accounts in advance of installing new computers. As new computers are brought online, their computer names must match the names that you have specified in the script. This approach works well for organizations in which many similar computers need to be added to a domain simultaneously, such as in a training lab or server farm. For more information about using scripts to create new computer accounts, see Windows Deployment and Resource Kits at

http://www.microsoft.com/reskit, or see the MSDN Scripting Clinic link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Note**

- It is more secure to create new computer accounts from the computer itself, rather than creating the accounts remotely or by using scripts. An attacker who gains access to some part of a domain can use existing scripts or remote account creation processes to create accounts to further compromise the system. Requiring that new accounts be created from the new computer protects against such attacks.

You can choose to delete computer accounts in your organization in one of the following ways:

- Include deleting users' computer accounts as part of the employee departure procedure. When employees leave your organization, establish a policy by which their computer accounts are deleted from Active Directory.

- Create scripts that search for computer accounts that havenot been logged on to for a period of time or have not had their password changed, and delete those computer accounts. For example, you might create a script that identifies accounts that have not been logged on to for six weeks or that have not had their passwords changed for twice the password lifetime as prescribed by domain Group Policy, and delete those accounts.

If a computer is unable to contact a domain controller to initiate a password change, the account might become unsynchronized with the domain and require a password reset. An effective way to enable the resetting of computer accounts in your organization is to assign help desk staff the right to reset passwords. Delegate the right to reset computer passwords to help desk staff so that members of the Domain Administrators group are not required to reset computer account passwords.

**Important**

- If you are migrating from Windows NT 4.0 domains, you must create a plan for the creation of new computer accounts. Computers running Windows NT 4.0 do not have computer accounts.

# Creating Service Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Like users, services have accounts and authenticate to the network operating system. This ensures that only authorized services are able to complete tasks, and protects against attackers who create unauthorized services to infiltrate network systems.

Most service accounts are created automatically when a service is installed. Similarly, applications that act as services, such as print spoolers or messaging services, create accounts automatically to complete their tasks. Therefore, in general, you do not need to create or modify service accounts. However, if service accounts are deleted accidentally, you must recreate them manually.

Creating service accounts is similar to creating user accounts. The only additional configuration step that is needed is to set the service principal name (SPN) for the account. This needs to be done to ensure mutual authentication. For example, in the case of a web server, a SPN of http/*hostname* might need to be set for the service account. The SPN can be set for the account by using the Setspn utility. For more information about Setspn, in Help and Support Center for Windows Server 2003, click Tools, and then click Install Windows Support Tools.

There are also built-in service accounts that use the computer account credentials by default for network authentication. These include the LocalSystem account, which was already present in Windows 2000. However, LocalSystem is a privileged account and should be used only when required.

Windows Server 2003 includes the following new security contexts to provide a means by which you can further secure network service accounts:

- **LocalService.** This context is intended for services that run with limited access on local computers and do not require network authentication. In this way, a compromised service can do limited damage to the local computer and no damage to network computers.

- **NetworkService.** This context is intended for services that need to complete tasks on the network, but require only restricted local capabilities.

# Securing Service Accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most services have specific functions, so it is best to grant them only those rights that are required for the services to perform those functions. In this way, if attackers compromise a service account, they have limited access and can do only a limited amount of damage. If a service account has rights that extend beyond its specific function, an attacker who compromises the account can do extensive damage.

To ensure maximum security, avoid running services on domain controllers. For example, do not make your domain controller a mail server, Web server, and file and print server. Adding multiple services on a critical link such as the domain controller is risky, because it increases the complexity of the system and therefore increases the potential for compromise. A problem with a print server that might otherwise only give an attacker the ability to create unauthorized print jobs can instead grant the attacker access to Active Directory, a critical data repository. The security benefits of using separate computers for services outweigh the initial investment in hardware equipment.

Also, you might need to reset service account passwords. Do not modify service accounts unless a problem occurs that interferes with the functioning of a service.

**To reset service account passwords**

1. In **Active Directory Users and Computers**, right-click the user's account.
2. Click **Reset Password**.
3. Enter and confirm the new password.

You must ensure that the service uses the newly selected password before the service can take advantage of the reset password. Ensure that the password that the service uses and the password that you reset the service account to have are the same.

# Applying Authentication Policies to Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can manage authentication in your organization by adding user, computer, and service accounts to groups and then applying authentication policies to those groups. For example, you can apply the following policies to groups, based on their function in the organization:

- **Log on locally**
- **Access this computer from the network**
- **Log on over network**
- **Reset accounts**
- **Create accounts**

If you want to make a computer less accessible to others, including both legitimate users and attackers, you can use policies in the following ways to restrict access for less trusted groups (such as Anonymous):

- Assign the **Deny access to this computer from the network** policy.
- Assign the **Deny logon locally** policy.
- Remove the **Remove computer from docking station** policy.

Other policies that you might assign or deny to users can also increase security or maximize flexibility, such as **Deny logon as batch job** or **Log on as service**. For more information about Group Policies that impact authentication, see "Deploying Security Policy" in *Designing a Managed Environment* of this kit.

# Example: Creating a Foundation for Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization that includes 2,100 users and 3,700 computers created an authentication strategy when they deployed Windows Server 2003 in their environment. Because computers in their environment are running versions of the Windows operating system earlier than Windows 2000, they need to support LAN Manager authentication. They decided to make members of the help desk staff and the Administrators group responsible for user account management, and delegated computer account management to the help desk staff.

The organization secured their service accounts by running only required services on domain controllers and restricting the number of individuals who are able to administer services. They assigned the **Log on locally**, **Access this computer from the network,** and **Log on over network** rights to Domain Admins and Domain Users, but not to Guest accounts, to protect the security of their system. They granted the **Reset accounts** and **Create account** policies to help desk staff to reduce the administrative burden on domain administrators.

Figure 14.3 shows the worksheet that the organization created to document their authentication strategy plan.

**Figure 14.3 Example of an Authentication Strategy Planning Worksheet**

| Authentication Strategy Planning | | | |
|---|---|---|---|
| Prepared By | | Date | |
| **Current Environment** | | | |
| | Number | | Number |
| Domain controllers | 5 | Certification authorities | 3 |
| Users | 2,100 | Computers | 3,700 |
| Smart card users | | | |
| Operating systems running: | | | |
| Windows 2000, Windows NT 4.0, Windows ME, Windows 98, Linux, Macintosh OS X | | | |
| **User Accounts** | | | |
| Users or groups to whom the following rights are assigned: | | | |
| Reset user passwords | Help desk, domain administrators | | |
| Reset account lockouts | Help desk, domain administrators | | |
| Create new user accounts | Help desk, domain administrators | | |
| Disabling user accounts plan: | | | |
| When a user leaves the organization, human resources is responsible for notifying the Help Desk team of the need to disable the user account. | | | |
| **Computer Accounts** | | | |
| Security policies that apply to computer groups: | | | |
| Log on locally; Access this computer from the network; Log on over network. | | | |
| Users or groups to whom the following rights are assigned: | | | |
| Create computer accounts | Help Desk | | |
| Delete computer accounts | Help Desk | | |
| Reset computer accounts | Help Desk | | |
| **Service Accounts** | | | |
| Services authorized to run in the organization | IIS, FTP, telnet, Windows Time Service, KDC, DHSP, DNS, WINS, proxy service | | |
| Services authorized to run on domain controllers | Due to security concerns, only those services that are installed when the operating system is installed are permitted to run on domain controllers. | | |
| Users authorized to reset service accounts | Members of the Domain Admins group. | | |

| Authentication Strategy Planning | | | |
|---|---|---|---|
| Prepared By | | Date | |

| Policies Assigned to Groups | |
|---|---|
| Groups to which the following policies are assigned: | |
| Log on locally | Domain Users, Domain Admins |
| Deny log on locally | Domain Users, Domain Admins |
| Access this computer from the network | Domain Users, Domain Admins |
| Deny access to this computer from the network | Guest |
| Log on over network | Domain Users, Domain Admins |
| Reset accounts | Help desk |
| Create accounts | Help desk |
| Remove computer from docking station | Domain Admins |
| Other security-related policies that apply to groups: | |
| | |

# Securing the Authentication Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to secure your authentication process to protect your system against various types of security threats, such as password-cracking tools, brute-force or dictionary attacks, abuse of system access rights, impersonation of authenticated users, and replay attacks. In addition, if you share resources on your network with other organizations, you must ensure that your authentication policies interoperate with the authentication policies that are in place on other systems.

For a worksheet to use in documenting authentication security policies, see "Authentication Security" (DSSAUT_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Authentication Security" on the Web at http://www.microsoft.com/reskit).

Figure 14.4 shows the process for securing authentication.

**Figure 14.4 Securing Authentication**



# Creating a Strong Password Policy

Updated: September 30, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Given enough encrypted data, time, and computing power, attackers can compromise almost any cryptographic system. You can prevent such attackers from succeeding by making the task of cracking the password as difficult as possible. Two key strategies to accomplish this are to require users to set complex passwords and to require users to change their passwords periodically, so that attackers do not have sufficient time to crack the complex encryption code.

**Complex Passwords**
You should set password policy to require complex passwords, which contain a combination of uppercase and lowercase letters, numbers, and symbols, and are typically a minimum of seven characters long or more for all accounts, including administrative accounts, such as local administrator, domain administrator, and enterprise administrator.

In this way, when users submit a new password, the password policy determines whether the password meets established requirements. You can set more complex password requirements; however, such password policies can increase costs to the organization if they obligate users to select passwords that are difficult to remember. Users might be forced to call the help desk if they forget their passwords, or they might write down their passwords, thus making them vulnerable to discovery. For this reason, when you establish password policies, you need to balance the need for strong security against the need to make the password policy easy for users to follow.

**Older Client Operating Systems**
Versions of the operating system earlier than Windows Server 2003 cannot handle passwords that contain more than 14 characters. For example:

- Attempts to log on to a Windows 2000–based computer running Terminal Services by using automatic logon settings configured in Client Connection Manager fail if your password is more than 14 characters long. Client Connection Manager has a 14-character limitation for passwords used for automatic logon. To work around this problem, you must manually enter a password to be used for the connection when prompted. You can prevent this by modifying the password used in Client Connection Manager and on your domain to be no more than 14 characters long.

- In the Microsoft Windows NT 3.51 operating system, Run.exe allows users to start utilities. When users start utilities, they can specify a user account and password to be used to start the application. When the password parameter is used, Run.exe stores the values in buffers limited to 14 characters. Passwords longer than 14 characters are truncated for storage and then passed to domain controllers in truncated form, causing authentication failures.

You can solve many of these problems by applying the latest service packs for operating systems. If your organization includes client computers running versions of the operating system earlier than Windows Server 2003 that do not support longer passwords, be sure to account for this when you set your password policies.

**Selecting Password Policy Options**
Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 provide security policies that ensure that all users select strong passwords. Creating a password policy involves setting the following options in the Default Domain Group Policy object. These policies are enforced on all user accounts in a domain.

**Enforce password history**
This policy setting determines the number of unique new passwords that have to be associated with a user account before an old password can be reused. This feature enables administrators to enhance security by preventing users from reusing their old password continually. The default value on domain controllers is 24, and the default value on stand-alone servers is 0. Most IT departments choose a value greater than 10.

**Maximum password age**
This policy setting determines the period of time (in days) that a password can be used before the system requires the user to change it. The best defense against impersonation is to require that users change their passwords regularly. This reduces the amount of time available for attackers to crack unknown passwords. The maximum password age can be set between 1 and 999 days. The default value of 42 days is generally appropriate; however, some IT departments shorten this to 30 days. To specify that passwords never expire, set this value to 0.

**Minimum password age**

This policy setting determines the number of days that must pass before a user can change his or her password. Defining a minimum password age prevents users from circumventing the password history policy by defining multiple passwords in rapid succession until they can use their old password again. The default value on domain controllers is 1, and the default value on stand-alone servers is 0. A value of a few days discourages rapid password recycling while still permitting users to change their own passwords if desired.

If you enable the **Enforce password history** policy setting, the minimum password age must be configured to a value greater than 0.

If **Maximum password age** is set between 1 and 999 days, **Minimum password age** must be less than **Maximum password age**. If **Maximum password age** is set to 0, **Minimum password age** can be any value between 0 and 998 days.

| **Note** |
| --- |
| Setting this parameter to a value higher than **Maximum password age** forces users to call the IT department to change their passwords, which increases costs to the organization. |

**Minimum password length**
The policy setting determines the minimum number of characters that a user's password must contain. You can set a value between 1 and 14 characters. To specify that no password is required, set the value to 0. The default value on domain controllers is 7, and the default value on stand-alone servers is 0.

**Passwords must meet complexity requirements**
This policy setting determines whether new passwords must meet complexity requirements. If this policy setting is enabled, passwords must meet the following minimum requirements:

- Passwords cannot contain the user's account name or parts of the user's full name that exceed two consecutive characters.
- Passwords must be at least six characters in length or the number of characters specified in the **Minimum password length** policy setting.
- Passwords must contain characters from at least three of the following four categories:
    - English uppercase alphabet characters (A–Z)
    - English lowercase alphabet characters (a–z)
    - Base 10 digits (0–9)
    - Non-alphanumeric characters (for example, !$#,%)

This policy setting is enabled by default on domain controllers and disabled by default on stand-alone servers.

**Store passwords using reversible encryption**
This policy setting determines whether the operating system stores passwords by using reversible encryption. It also provides support for applications that use protocols that require knowledge of user's password for authentication purposes.

| **Important** |
| --- |
| Storing passwords by using reversible encryption is the same as storing plaintext versions of the passwords. For this reason, this policy should not be enabled unless application requirements outweigh the need to protect password information. |

This policy setting must be enabled when using:

- Challenge Handshake Authentication Protocol (CHAP) authentication through remote access
- Internet Authentication Service (IAS)
- Digest authentication in Internet Information Services (IIS).

This policy setting is disabled by default.

# Establishing an Account Lockout Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to establish an account lockout policy at the same time that you establish a password security policy. Account lockout policies protect your environment against brute-force or dictionary attacks. Given enough tries, even complex passwords can be guessed. Account lockout policies reduce the number of guesses that an attacker can make.

It is best to establish an account lockout policy that is restrictive enough to prevent attacks, while still allowing for the occasional user error. An account lockout policy that is too strict might increase the number of support calls in your organization as users who type their passwords incorrectly are mistakenly locked out.

Creating an account lockout policy involves setting the following options in the Default Domain Group Policy object.

**Account lockout threshold**
The account lockout threshold limits the number of times that anyone can attempt to log on to a computer from a remote location. This prevents attackers from trying all possible passwords over the network. This setting is disabled by default in the Default Domain Group Policy object. You can turn it on by setting the value to a number within the accepted range of 1 through 999. Set the value high enough to ensure that occasional errors do not result in account lockout.

Note that this setting does not apply to attempts to log on at the console of a locked workstation or to attempts to unlock a screensaver. Locked workstations cannot be forced to run password-cracking programs.

**Account lockout duration**
The account lockout duration determines how long, in minutes, an account that has exceeded the account lockout threshold remains locked before it is automatically unlocked. Valid settings range from 0 through 99,999 minutes, or about 10 weeks. When the value is set to 0, an administrator must manually unlock the account. Because account lockout policies are designed to protect against brute-force attacks, setting even a low value for the account lockout duration reduces the number of possible attacks considerably. Note that setting a high value for the account lockout duration can increase help desk calls when legitimate users are mistakenly locked out, and aside from indicating that an attack was attempted, provides little additional protection.

By default, this policy is not defined, because it is only applicable when an account lockout threshold is specified.

**Reset account lockout counter after**
This setting determines the number of minutes that must elapse after a failed logon attempt before the counter is reset to 0 bad logon attempts. The range is 1 through 99,999 minutes. This value must be less than or equal to the account lockout duration.

**Enforce user logon restrictions**
When this option is enabled, the KDC validates every request for a session ticket by examining the user rights policy on the target computer. The user requesting the session ticket must be assigned the **Log on locally** policy (if the requested service is running on the same computer) or the **Access this computer from the network** policy (if the requested service is on a remote computer) to receive a session ticket. This option also serves as a means to ensure that the requesting account is still valid. Verification is optional because the extra step takes time and might slow network access to services, but if account rights have changed or user accounts have been disabled between the time when the initial ticket was issued and the time when a service ticket was requested, these changes do not take effect.

By default, the policy is enabled in the Default Domain Group Policy object. If the policy is disabled, this check is not performed. For greater security in an environment in which user accounts change frequently, enable this setting. For faster performance, particularly in a more stable user account environment, disable this setting.

# Assigning Logon Hours

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can assign logon hours as a means to ensure that employees are using computers only during specified hours. This setting applies both to *interactive logon*, in which a user unlocks a computer and has access to the local computer, and *network logon*, in which a user obtains credentials that allow him or her to access resources on the network.

Assigning logon hours is useful for organizations in which some users are less trustworthy than others or require supervision. For example, you might want to restrict logon hours when:

- Logon hours are a condition for security certification, such as in a government network.
- Your organization includes shift workers. In this case, allow shift workers to log on only during their scheduled hours.
- Your organization includes temporary employees.

The logon schedule is enforced by the Kerberos Group Policy setting **Enforce User Logon Restrictions**, which is enabled by default in Windows Server 2003. Whether users are forced to log off when their logon hours expire is determined by the **Automatically log off users** setting.

By default, all domain users can log on at any time. You can use the following procedure to limit the logon hours of an individual domain user.

**To restrict the logon hours of a domain user**

1. In **Active Directory Users and Computers**, right-click the user's account.
2. Click **Properties**, and click the **Account** tab.
3. Click **Logon Hours**. In the **Logon Hours** dialog box, indicate the hours and/or days of the week in which you are restricting the user from logging on.

When you have set the logon hours for an individual, you can copy that account to apply the same settings to a new user in the same department.

**To restrict the logon hours for multiple users in the same OU**

1. In **Active Directory Users and Computers**, select the user accounts, and then right-click any of the selected items.
2. Use the **Properties of Multiple Objects** dialog box to alter the properties for all of the selected users.

When you restrict logon hours, you might also want to force users to log off after a certain point. If you apply this policy, users cannot log on to a new computer, but they can stay logged on even during restricted logon hours. To force users to log off when logon hours expire for their account, apply the **Network security: Force logoff when logon hours expire** policy.

# Creating a Ticket Expiration Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to establish reasonable lifetimes for tickets in your organization. Ticket lifetimes must be short enough to prevent attackers from cracking the cryptography that protects the ticket's stored credentials. However, ticket lifetimes must also be long enough to be convenient for users and to ensure that requests for new tickets do not overload the network.

Creating a ticket expiration policy involves setting the following options in the Default Domain Group Policy object.

**Maximum lifetime for user ticket**
This setting indicates the amount of time for which a ticket is valid before it expires. Generally, it is best if the Maximum Lifetime for User Ticket setting reflects the average amount of time that users access their computers in one day. This is set to 10 hours in the Default Domain Group Policy object. At the end of the ticket lifetime, the user either obtains a new ticket or renews the existing ticket. This process is performed transparently by the computer, but each ticket request or renewal produces network traffic and domain controller loading.

A short maximum ticket lifetime provides greater security but also increases network traffic. A long maximum ticket lifetime decreases network traffic but does not provide the same level of security.

**Maximum lifetime for service ticket**
This setting usually matches the established user ticket lifetime. It might be shorter, however, if there is a need in your organization for secure authentication to services beyond what is required for user authentication. It might be longer if users require uninterrupted access to services for long periods of time. For example, you might need to extend the ticket lifetime if your users run jobs that have a duration that is longer than the duration of the user ticket lifetime. If you do not have any special requirements for service ticket lifetime, do not extend the lifetime of the ticket.

The maximum service ticket lifetime must be greater than 10 minutes and less than or equal to the Maximum Lifetime for User Ticket setting. By default, this value is set to 600 minutes (10 hours) in the Default Domain Group Policy object (GPO). Ongoing operations are not interrupted if the session ticket used to authenticate the connection expires before the operation is complete.

**Maximum lifetime for user ticket renewal**
This setting determines the period of time (in days) during which a user's ticket-granting ticket (TGT) can be renewed. By default, this is set to seven days in the Default Domain GPO. Shorter renewal times make it easier to require users to reauthenticate in the event that you suspect that there has been a security breach. An attacker with a renewable user ticket can continue to renew that ticket for as long as the policy allows. Shortening renewal times makes an attacker's task more difficult, but it also increases the authentication load on domain controllers.

# Establishing Network Authentication Standards

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 allows for interoperability with earlier versions of Windows and other operating systems. Interoperating with other operating systems, however, can negatively impact your network security. It is important, therefore, to establish standards for network authentication to minimize the effect that this interoperability has on your organization's security. You can do this by restricting LAN Manager authentication and by restricting anonymous access. You must also establish a plan for upgrading Windows NT 4.0 domain controllers to balance the Kerberos authentication load in Windows Server 2003 and Windows 2000 domains.

# Restricting LAN Manager Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Due to advances in cracking tools and hardware capabilities, LAN Manager authentication encryption is more vulnerable to attack than newer forms of encryption. For this reason, it is important to restrict the use of LAN Manager authentication whenever possible. Windows Server 2003 supports all versions of LAN Manager authentication, including LM, NTLM, and NTLM version 2 (NTLMv2), to allow for compatibility with clients that do not support newer authentication protocols.

If it is necessary in your organization to support LAN Manager authentication, you can increase security by enabling support of NTLMv2 whenever possible. Reducing or eliminating the use of LAN Manager authentication and NTLM version 1 (NTLMv1) removes password hash values from the network, and therefore increases network security. You can enable NTLMv2 support by doing the following:

- Upgrading to at least Service Pack 4 (SP4) on all Windows NT 4.0–based clients

- Installing the directory services client on all client computers that are running the Microsoft® Windows® 95 or Windows® 98 operating system You can install the directory services client from the Windows Server 2003 operating system CD.

- Tightening LAN Manager authentication policies. If all clients support NTLMv2, set Domain Group Policy for **LAN Manager Authentication Level** to **Send NTLMv2 response only\refuse LM & NTLM**. This policy is under Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\. If some clients exist that do not support NTLMv2, set the **LAN Manager Authentication Level** to **Send NTLM response only**. This reduces the amount of ciphertext available to attackers.

**Note**

- Clients that do not typically support NTLMv2 include Macintosh and Windows Services for UNIX.

# Restricting Anonymous Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Windows Server 2003, access that was available to Anonymous users in Windows NT 4.0 is available only to Everyone and Guest accounts. However, in some of the following situations you might still need to allow Anonymous access to portions of your network.

Some of the services running versions of Windows earlier than Windows 2000 use anonymous access to request user account information from domain controllers and to list network shares on file servers and workstations.

You also might need to allow Anonymous access when an administrator in the trusting domain of a one-way trust relationship across forests needs to list users and shares in the trusted domain of another forest.

In addition, the Windows NT Remote Access Service (RAS) uses anonymous logon to determine whether a user has permission to establish a RAS connection. Anonymous access to Active Directory is used to change passwords from earlier systems. This form of anonymous access is enabled by the Pre-Windows 2000 compatible access security group, which is a local group found only on Windows 2000 and Windows Server 2003 domain controllers. By default, this group has read access to user and group objects in Active Directory.

If you need to support networks containing a mix of Windows NT 4.0, Windows 2000, and Windows Server 2003 desktops and servers, you must take into account the new restrictions on anonymous access by doing the following:

- First determine which services and applications require anonymous access to network resources, and identify the servers to which anonymous access is needed.

- Then decide whether to add the Anonymous Logon identity to specific access control lists (ACLs), or to make security policy changes that relax the restrictions that Windows Server 2003 places on anonymous access.

You can regulate anonymous access by doing the following:

- Edit the ACLs of the resources, adding the Anonymous Logon identity to the list of authorized users. This approach is the most secure, but requires editing the ACLs of each resource, which might be difficult to manage or troubleshoot.

- Use the **Do not allow anonymous enumeration of SAM accounts and shares policy** Group Policy object, which can be found in Computer Configuration/Windows Settings/Security Settings/Local Policies/Security Options, to prevent attackers from using anonymous connections to obtain information about accounts and shares on a computer. Preventing Security Accounts Manager (SAM) account enumeration can help thwart attacks, but also prevents legitimate users in other domains from obtaining this information.

- Disable or do not configure this policy if your domain includes computers running versions of Windows earlier than Windows 2000 or if it has an outbound, one-way trust relationship with a domain in another forest. The browser service on computers running Windows NT 4.0 and earlier requires the ability to enumerate shares anonymously when it connects to backup browsers, master browsers, and domain master browsers to retrieve server lists and domain lists. Users on the trusting side of a one-way trust relationship need the ability to enumerate SAM accounts anonymously when they add domain accounts and groups on the trusted side of the relationship to security groups in the trusting domain.

  For more information about domain and forest trust relationships, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

- Use the **Let Everyone permissions apply to anonymous users** policy to extend anonymous access to match the Windows NT 4.0 model. If this policy is enabled, Anonymous users can access any resource that Everyone is allowed to access. Do not enable this policy unless there is a compelling business reason to compromise the security provided by requiring some form of authentication. If you do enable this policy, work to disable it by editing the ACLs of specific resources to allow anonymous access, as required in particular cases.

- If you need to permit clients running versions of Windows earlier than Windows 2000 to change their passwords, add the Everyone and Anonymous Logon groups to the Pre-Windows 2000 compatible access group, enabling anonymous access to the accounts. The membership of this group is determined by a user option during the installation of the first domain controller in the domain. You can change the group membership if necessary.

# Creating a Plan for Windows NT 4.0 Domain Controller Upgrade

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Workstations running the Microsoft® Windows® XP operating system that are joined to an existing Windows NT 4.0 domain use NTLM to authenticate to the Windows NT 4.0 domain controllers. When you add a Windows Server 2003– or Windows 2000–based domain controller to a domain, all clients running the Microsoft® Windows® XP, or Windows® 2000 Professional operating system and all servers running Windows Server 2003 or Windows 2000 automatically use Kerberos authentication when users log on interactively. Users at these computers therefore cannot log on by using the Windows NT backup domain controllers. This shifts the Windows Server 2003 and Windows 2000 user authentication load to the existing Windows Server 2003– or Windows 2000–based domain controllers. If the primary domain controller becomes overloaded or fails for any reason, users cannot log on to computers running Windows Server 2003 or Windows 2000.

For this reason, when you add Windows Server 2003– or Windows 2000–based domain controllers to a domain, you must continue to add Windows Server 2003– or Windows 2000–based domain controllers to keep pace with client demands. You can do this by installing new domain controllers or by upgrading existing Windows NT 4.0–based domain controllers.

Review your organization's plans for domain upgrade and consolidation and for the deployment of new Windows Server 2003–based workstations, and ensure that workstation upgrade does not proceed more rapidly than the domain controller upgrade. Add Windows Server 2003– or Windows 2000–based domain controllers to a domain only when you are certain that the domain controllers have the capacity to meet the authentication needs of all Windows Server 2003–based workstations in the domain.

For more information about upgrading Windows NT 4.0 domains to Windows Server 2003 Active Directory, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book.

# Setting Clock Synchronization Tolerance to Prevent Replay Attacks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **Maximum tolerance for clock synchronization** Group Policy to protect your organization against replay attacks, in which attackers replay authentic network exchanges that they capture off the wire to cause the server to allow them access to the system. If your clock synchronization tolerance setting is low, the server rejects replayed messages for which the allowable time skew has passed.

The **Maximum tolerance for computer clock synchronization** Group Policy is set to 5 minutes by default. In most cases, this provides an acceptable level of security. You can increase protection against replay attacks by shortening the maximum tolerance for clock synchronization. Tighter synchronization requirements, however, might result in increased authentication traffic.

Shortening the maximum tolerance reduces replay attacks because the Kerberos V5 authentication protocol uses authenticators based on time to establish user identities. A shorter tolerance makes a replay attack more difficult.

The **Maximum tolerance for computer clock synchronization** Group Policy can be found in the Default Domain Policy object under Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy.

In general these settings should be changed only if there is a strong reason to believe you might be vulnerable to this type of attack.

For more information about time synchronization in Windows Server 2003, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Example: Securing Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization developed an authentication strategy to strengthen the default security that Windows Server 2003 provides. This is to protect data from attackers whose activity has been noted in audit logs and because management has made increasing the security of the system a top priority. To meet these security demands, administrators created password policies to ensure strong passwords, applied account lockout policies to prevent brute-force attacks, assigned logon hours to prevent users from working during unsupervised times, and established a ticket expiration policy to enforce logon hours for several groups, excluding batch jobs.

The organization chose to eliminate LAN Manager authentication by setting the **Restrict LanMan Authentication** policy to **Not supported** to prevent the use of authentication methods that are vulnerable to attack. They chose to eliminate anonymous logon by enabling **Restrict Anonymous Access** to limit the number of resources that attackers can access by impersonating anonymous users. In order to enable these policies, the organization chose to retire their Windows NT 4.0–based domain controllers, and replaced them with new Windows Server 2003–based computers.

Administrators in the organization accepted the default clock synchronization tolerance of five minutes. This setting protects the system against replay attacks while keeping authentication traffic to a minimum.

Figure 14.5 shows the worksheet that the organization created to document their authentication security plan.

**Figure 14.5 Example of an Authentication Security Worksheet**

## Authentication Security

| | | | |
|---|---|---|---|
| Prepared By | | Date | |
| Domain | | | |

### Existing Operating Systems

| | | | | Upgrade? (Y/N) |
|---|---|---|---|---|
| OS | Windows Server 2003 | SP | | N |
| OS | Windows 2000 | SP | 2 | N |
| OS | Windows 2000 | SP | | N |
| OS | Windows NT 4.0 | SP | 6a | Y |

### Password Policy

| | | | |
|---|---|---|---|
| Maximum Password Age | 42 days | Enforce Password History | Yes |
| Minimum Password Age | 1 day | Minimum Password Length | 7 |
| Passwords Must Meet Complexity Requirements (Y/N) | | | Y |

### Account Lockout Policy

| | | | |
|---|---|---|---|
| Account Lockout Threshold | 3 | Account Lockout Duration | 1 minute |
| Reset Account Lockout Counter After | 1 minute | Enforce User Logon Restrictions | Enabled |

### Logon Hours Policy

| Groups to which logon hours are assigned | Hours for which groups are permitted to log on | Force logoff when logon hours expire? (Y/N) |
|---|---|---|
| Assembly | 7:00 A.M. - 7:00 P.M. | Y |
| Data Entry | 7:00 A.M. - 7:00 P.M. | Y |
| Reception | 7:00 A.M. - 7:00 P.M. | Y |
| Batch | 9:00 A.M. - 5:00 P.M. | N |
| | | |

### Ticket Expiration Policy

| | | | |
|---|---|---|---|
| User Ticket Maximum Lifetime | 10 hours | Service Ticket Maximum Lifetime | 10 hours |
| User Ticket Renewal Maximum Lifetime | | 7 days | |

### Network Authentication Standards

Groups to which the following policies apply:

| Policy | Setting(s) |
|---|---|
| Restrict LanMan Authentication | No groups. Support clients that require LanMan authentication. |
| Restrict Anonymous Access | Enabled |

## Authentication Security

| | | | |
|---|---|---|---|
| Prepared By | | Date | |
| Domain | | | |

### Windows NT 4.0 Domain Controller Upgrade Plan

To upgrade your Windows NT 4.0 domain controllers, install your three new computers as domain controllers, and then remove the remaining Windows NT 4.0 domain controllers.

### Clock Synchronization Tolerance

| | |
|---|---|
| Maximum tolerance for computer clock synchronization | 5 minutes |

# Extending Your Authentication Framework

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your environment includes more than one forest, or your organization needs to exchange information with other Kerberos clients and servers, you need to extend your authentication framework to accommodate those relationships and resources. You do this by establishing trust relationships and by creating Active Directory accounts as appropriate.

For a worksheet to assist you in documenting your plan for extending your authentication framework, see "Extended Authentication Framework" (DSSAUT_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Extended Authentication Framework" on the Web at http://www.microsoft.com/reskit).

Figure 14.6 shows the process for extending your authentication framework.

**Figure 14.6 Extending Your Authentication Framework**



# Establishing Interforest Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization includes more than one forest, you need to enable the forests to allow authentication and resource sharing. You can do this by establishing trust relationships between some or all of the domains in the forests. The types of trust relationships that you establish will depend on the versions of the operating system that are running in each forest.

For more information about establishing trust relationships, see "Understanding Trusts" in Help and Support Center for Windows Server 2003.

**Authentication between Windows Server 2003 forests**
When all domains in two forests trust each other and need to authenticate users, establish a forest trust between the forests. When only some of the domains in two Windows Server 2003 forests trust each other, establish one-way or two-way external trusts between the domains that require interforest authentication.

**Authentication between Windows Server 2003 and Windows 2000 forests**
It is not possible to establish transitive forest trusts between Windows Server 2003 and Windows 2000 forests. To enable authentication with Windows 2000 forests, establish one-way or two-way external trusts between the domains that need to share resources.

**Authentication between Windows Server 2003 and Windows NT 4.0 forests**
It is not possible to establish transitive forest trusts between Windows Server 2003 and Windows NT 4.0 domains. Establish one-way or two-way external trusts between the domains that need to share resources.

In each of these cases you should consider whether the selective authentication option needs to be enabled. Selective authentication should only be enabled when the trusted domain is located in the extranet or in a different corporation, and therefore only requires access to a very limited set of resources.

The Kerberos V5 authentication protocol does not work across forests in Windows 2000 or Windows NT environments. In these situations, Windows Server 2003 relies on the NTLM protocol for authentication across forests. A direct trust relationship between two domains in separate forests enables NTLM authentication; however, NTLM enables client authentication only, and not mutual authentication. Therefore, if you must authenticate across forest boundaries, you need to ensure that all computers running versions of Windows earlier than Windows 2000 have been upgraded to use NTLMv2.

# Enabling Interoperability with Kerberos Clients and Servers Running Other Operating Systems

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some organizations have clients running UNIX or other operating systems. To allow for the secure exchange of information with other clients, you can configure the clients to authenticate to Windows Server 2003 domain controllers in order to obtain the required credentials. Similarly, Windows clients can be configured to authenticate to other KDCs.

To enable interoperability with UNIX or other clients, you must:

- Establish a Realm trust, a particular type of external trust with a UNIX realm. This enables an authentication that is completed in one realm or domain to be trusted by another realm or domain.
- Create account mappings so that other clients have mapped accounts in Active Directory. This enables other clients to access resources that are secured in a Windows environment.

**Establishing Trusts with Kerberos Realms**

To enable authentication between Windows domains and UNIX realms or other clients, you must establish a one-way or two-way trust between the two so that tickets generated in one are recognized and accepted by resources in the other. For example, a one-way trust relationship in which a Kerberos realm trusts a Windows Server 2003 domain allows Windows Server 2003 users to log on to the Kerberos realm; in other words, the UNIX server accepts or trusts the authentication performed by the Windows Server 2003 KDC. Another trust can be created so that users logged on to the Kerberos realm can access resources in the Windows Server 2003 domain.

**Configuring Accounts for Kerberos Clients Running Other Operating Systems**

After you have established trusts between a UNIX or other realm and a Windows domain, you must coordinate accounts, enabling users to authenticate and access resources.

**User accounts**

You must configure clients to authenticate to the appropriate KDC. For example, you might configure a Linux desktop to authenticate to a Windows Server 2003–based KDC at logon. Most Kerberos clients allow for the specification of a KDC for authentication as part of the logon to the local computer. Windows Server 2003 provides the Kerberos KDC services as part of the domain controller, so the clients log on to the domain controller itself. The domain controller locates the KDC by means of service location records in the DNS. This frees the administrator from having to maintain explicit Kerberos configuration data for each client.

**Service accounts**

Windows Server 2003 supports the authentication of other Kerberos services in a Windows Server 2003 domain. If you require services to access resources across the domain or realm, you must create service accounts in Active Directory to represent those services. For example, you can make a UNIX-based telnet service accessible to Kerberos clients in a Windows domain by creating a service account in Active Directory for that service. In this case, the telnet service is part of the Windows domain, rather than the other Kerberos realm, as is the case with trust relationships established between Windows and other Kerberos realms.

**Account mappings**

When a Windows Server 2003 domain trusts a Kerberos realm, the principals in the Kerberos realm do not contain the group associations that are used for access control in the Windows Server 2003 environment. You can use account mapping in the Windows Server 2003 domain to provide authorization information for Kerberos principals from trusted realms. You can either map accounts one-to-one, by mapping each account in a realm to a corresponding account in the Windows Server 2003 domain, or you can use one-to-many mapping, by which multiple individual accounts in a realm are mapped to one account in the Windows Server 2003 domain. For more information about account mapping, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

To ensure seamless interoperability, you must keep the accounts in the Kerberos realm and the Windows Server 2003 domain synchronized. You can use ADSI and Lightweight Directory Access Protocol (LDAP) in Active Directory to synchronize accounts, or use metadirectory technology such as Zoomit Via.

# Deploying Smart Cards

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization requires a more secure form of authentication, you can use certificates and smart cards for user logon. Smart cards provide additional security because they require both a password and a physical smart card for a user to log on. The smart card contains a public key certificate, and to unlock the certificate, the user must supply a password, or PIN. It is much more difficult for an attacker to obtain both a physical smart card and the PIN number that is used to unlock it in order to gain access to network resources.

In order to decide whether smart card authentication is appropriate for your organizations, evaluate the potential benefits against the following consideration:

- **Costs.** Deploying smart cards entails initial equipment costs for the purchase of smart cards and smart card readers, as well as administrative costs for preparing and distributing smart cards.

- **Infrastructure.** A public key infrastructure (PKI) is required for smart card authentication. For more information about establishing a PKI, see "Designing a Public Key Infrastructure" in this book.

- **Ongoing administration.** Unlike card keys, smart cards cannot be replaced easily if lost or forgotten. This introduces the potential for lost productivity if a user forgets or loses his or her card.

Because of the potential costs and administrative burden, many organizations choose to deploy smart cards only to certain groups of users, such as administrators or users who have access to extremely sensitive data.

For more information about deploying smart cards, see "Planning a Smart Card Deployment" in this book.

# Example: Extending the Authentication Framework

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization that is based on an Active Directory logical structure that includes four forests, forest A, forest B, forest C, and forest D, must extend its authentication framework in order to facilitate resource access for its clients in all locations. The organization needs to share resources with a Windows NT 4.0 domain, Domain E, and Kerberos clients running Unix in Unix realm F. Figure 14.7 shows the logical structure of the organization.

**Figure 14.7 Organization Logical Structure**




To enable the sharing of resources, administrators establish the following trust relationships:

- A forest trust between forests A and B
- External trusts from domains in forests A and B to domains in forests C, D, and E
- External trusts between domains in forests C, D, and E
- Realm trusts between all domains and realm F

The organization chooses to deploy smart cards to domain administrators, as these accounts are more sensitive to security attacks.

Figure 14.8 shows the worksheet that the organization created to document their extended authentication framework.

**Figure 14.8 Example of an Extended Authentication Framework Worksheet**

| Extended Authentication Framework | | | |
|---|---|---|---|
| Prepared By | | Date | |
| Domain | | | |

| Interforest Authentication | |
|---|---|
| List the Resources in other forests to which access must be enabled: | |
| Other Windows Server 2003 forests | Forest A, Forest B |
| Windows 2000 forests | Forest C, Forest D |
| Windows NT 4.0 domains | Domain E |
| Non-Windows structures such as UNIX realms | Unix Realm F |

| Smart Cards |
|---|
| List users or groups that require smart cards for authentication: |
| |

# Enabling Supplemental Authentication Strategies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some organizations require more complex authentication solutions to meet their user access needs. For example, if your organization includes Web-based clients that use applications such as Microsoft® Internet Explorer to access data stored on back-end servers, you must enable complex authentication strategies, such as delegation, constrained delegation, and protocol transfer, to allow those clients to access the requested services in a secure manner.

For a worksheet to assist you in documenting your plan for enabling supplemental authentication strategies, see "Supplemental Authentication Strategies" (DSSAUT_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Supplemental Authentication Strategies" on the Web at http://www.microsoft.com/reskit).

Figure 14.9 shows the process for enabling supplemental authentication strategies.

**Figure 14.9 Enabling Supplemental Authentication Strategies**



# Enabling Delegated Authentication

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Delegated authentication occurs when a network service accepts a request from a user and assumes that user's identity in order to initiate a new connection to a second network service.

To enable delegated authentication, you must establish front-end or first-tier servers, such as Web servers, that are responsible for handling client requests, and back-end or n-tier servers, such as large databases, that are responsible for storing information. You can delegate the right to enable delegated authentication to users in your organization in order to reduce the administrative load on your administrators. To delegate this right, assign the **Enable computer and user accounts to be trusted for delegation** user right to the selected individuals. Users who are assigned the right to enable delegated authentication can assign the **Trusted for delegation** right to computer and service accounts that are used to serve users information that is stored on back-end servers and must be accessed securely. The user account that is requesting the resource must not be marked as sensitive; marking an account as sensitive explicitly denies the right to delegation.

By establishing a service or computer as trusted for delegation, you enable that service or computer to complete delegated authentication, receive a ticket for the user who is making the request, and then access information for that user. Delegated authentication prevents an attacker who gains control of a front-end server, such as a Web server, from also gaining access to data stored on a back-end server. By requiring that all data be accessed by means of credentials that are delegated to the server for use on the client's behalf, you ensure that the server cannot be compromised and then used to gain access to sensitive information about other servers.

Delegated authentication is useful for multitier applications that are designed to use single sign-on capabilities across multiple computers. For example, domain controllers are automatically trusted for delegation. If this property is disabled on a domain controller, the Message Queuing service cannot run. Also, if you enable the Encrypting File System on a file server, the server must be trusted for delegation in order to store encrypted files on behalf of users. Delegated authentication is also useful on applications where Internet Information Services (IIS) supports a Web interface to a database running on another computer, such as Outlook Web Access in Exchange, or Web Enrollment Support pages for an enterprise certification authority, if the pages are installed on a separate Web server.

It is recommended that you deny the right to participate in delegated authentication to the computer accounts in Active Directory for computers that are not physically secure, and to domain administrator accounts. Domain administrator accounts have access to sensitive resources and, if compromised, poses a higher risk to your organization.

When computers that are trusted for delegation are compromised by an attacker, the attacker can use them to access data stored on other servers by using the delegated credentials of an authenticated user. Ensure that only secure computers are trusted for delegation, and do not allow the delegation of powerful user accounts, such as administrator accounts. Also, consider applying constrained delegation to computers that are trusted for delegation, to limit the ways in which delegated credentials can be used. In this way, an attacker who has access to the computer has access to only limited services.

For more information about enabling constrained delegation, see "Enabling Constrained Delegation" later in this chapter.

**To restrict delegated authentication**

1. In **Active Directory Users and Computers,** right-click the computer or user account and select **Properties**.

2. On the **Account** tab, under **Account Options**, select the **Account is sensitive and cannot be delegated** check box, and click **OK**.

3. You can also restrict delegated authentication to prevent the delegation of sensitive user accounts by marking the account as not enabled for delegation. Restrict delegated authentication for accounts that are less secure or that are particularly powerful.

# Enabling Constrained Delegation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Constrained delegation allows administrators to specify particular services from which a computer that is trusted for delegation can request resources. By using constrained delegation, you can prevent attackers who compromise a server from accessing resources beyond the limited scope of that server's range.

Before you enable constrained delegation, isolate critical data that you must keep secure from data to which users require frequent access. For example, if your organization maintains an e-commerce Web site, you might choose to isolate customer credit cards numbers, internal accounting, or human resources information from order status information that customers access frequently.

**To enable constrained delegation**

1. In **Active Directory Users and Computers**,right-click the computer account and select **Properties**.
2. On the **Delegation** tab, click **Trust this computer for delegation to specified services only**.
3. Select **Use Kerberos only**,or select **Use any authentication protocol**.
4. Click **Add** and, in **Add Services**, click **Users and Computers**.
5. In **Add Services**, select the service or services that are trusted for delegation, and click **OK**.

You can further restrict the scope of delegation that is permitted, for example to disable delegation for highly sensitive accounts such as administrator accounts.

**To restrict delegation**

1. In **Active Directory Users and Computers,** right-click the user account and select **Properties**.
2. Select the **Sensitive** attribute check box for the user account, and click **OK**.

# Example: Supplementary Authentication Strategies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization chose to extend its authentication framework by enabling delegation and constrained delegation. To achieve this, they assigned the right to enable delegated authentication to specific user accounts, identified the computer accounts that are to be trusted for delegation, and established who is responsible for applying these policies. This allowed them to strengthen the security of their system by limiting the resources to which computers that are trusted for delegation have access. For example, they enabled the Web interface human resources database to access confidential data stored in databases in other servers, assigned the trusted for delegation right to workstations, and restricted delegation on the domain administrator user account.

Figure 14.10 shows the worksheet that the organization created to document their supplementary authentication strategies.

**Figure 14.10 Example of a Supplementary Authentication Strategies Worksheet**

| Supplementary Authentication Strategies | | | |
|---|---|---|---|
| Prepared By | | Date | |
| Delegated Authentication | | | |
| User accounts for which delegation is not permitted: | | | Individual responsible for marking the account as sensitive: |
| Domain Administrators | | | Enterprise Administrators |
| | | | |
| Computers in the organization that are trusted for delegation: | | | Individual responsible for enabling the computer for delegation: |
| Workstations | | | Domain Administrators |
| | | | |
| Constrained Delegation | | | |
| Services to which users in your organization require access: | | | Individual responsible for permitting delegation to the specified SPNs on the computer account: |
| HR database | SPN | HRdb | Domain Administrators |
| | SPN | | |

# Educating Users About the Authentication Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Ensuring that users understand their role in the authentication process is an important step in establishing an authentication strategy. By giving users clear guidelines and explaining the justification behind the guidelines, you reduce the chances that users will fail to comply with the procedures that you put into place.

Figure 14.11 shows the process for educating users.

**Figure 14.11 Educating Users**



# Increasing Awareness of Social Engineering Attacks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Users must be cautioned against sharing their passwords with others. All legitimate users have their own accounts, and any administrator who needs to complete a task for a user can do so by using his or her own account, without knowledge of the user's password. Tasks such as the resetting of a user's password or the unlocking of a user's account do not require the use of a password.

One effective and simple way that an attacker can compromise the security of a system is to call users, claiming to be from the help desk, and request their passwords. Because users feel compelled to help solve problems, they are not motivated to question the authenticity of the caller. Caution users to beware of such calls and assure them that it is appropriate to be skeptical of requests for their passwords. Establish a procedure by which users who receive calls of this type request the caller's name and number and call them back. In this way, they can ensure that the call is legitimate before they reveal sensitive information.

# Communicating Password Creation Guidelines

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Educate your users about how to create strong passwords and how to keep them secret to help to protect your system from compromise as a result of simple carelessness.

Users must understand that their passwords must meet your organization's complexity guidelines. Incorporating uppercase and lowercase letters, numbers, and symbols into a password makes the password much more difficult to crack. Suggest to users that they insert numbers and characters into common phrases to protect against dictionary attacks. For example, the phrase "iamhappy" is easy to remember, but is hard to guess if some characters are changed so that the phrase appears as "1AmH@ppy!".

Caution users against writing their passwords down and leaving them in an accessible place. Be sure that users understand the danger of leaving their passwords in places where they can be discovered by an attacker. Although the need to remember complex passwords and reset them frequently might cause some inconvenience, the benefits of protecting your organization's resources far outweigh the costs.

# Additional Resources for Designing an Authentication Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following resources contain additional information and tools related to the chapter.

**Related Information**

- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about the Kerberos V5 authentication protocol and trust relationships.
- "Designing the Active Directory Logical Structure" in this book for more information about the Active Directory logical structure.
- "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory" in this book for more information about domain migration.
- "Designing a Public Key Infrastructure" in this book for more information about designing a public key infrastructure.
- "Deploying IAS" in *Deploying Network Services* of this kit for more information about using RADIUS servers.
- "Deploying Dial-up and VPN Remote Access Servers" in *Deploying Network Services* for more information about using RAS servers.
- "Designing a Resource Authorization Strategy" in this book for more information about user accounts and security groups.
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about time synchronization in Windows Server 2003.
- "Deploying Smart Cards" in this book for more information about deploying smart cards.
- "Deploying Security Policy" in *Designing a Managed Environment* of this kit for more information about Group Policies that impact authentication.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Understanding Trusts" in Help and Support Center for Windows Server 2003 for more information about establishing trust relationships.

**Related Job Aids**

- "Authentication Strategy Planning" (DSSAUT_1.doc) "Authentication Strategy Planning" on the Web at http://www.microsoft.com/reskit).
- "Authentication Security" (DSSAUT_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Authentication Security" on the Web at http://www.microsoft.com/reskit).
- "Extended Authentication Framework" (DSSAUT_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Extended Authentication Framework" on the Web at http://www.microsoft.com/reskit).

- "Supplemental Authentication Strategies" (DSSAUT_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Supplemental Authentication Strategies" on the Web at http://www.microsoft.com/reskit).

# Designing a Resource Authorization Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Authenticated users need access to shared resources such as file shares and printers. An authorization strategy based on security groups enables you to effectively manage users' access to these resources. To create a resource authorization strategy, you need to decide what types of security groups to use and plan how to manage and maintain these groups.

**In This Chapter**

Overview of Designing a Resource Authorization Strategy

Establishing a Resource Authorization Method

Defining Policies for Security Group Management

Delegating Security Group Maintenance

Additional Resources for Resource Authorization Strategy

**Related Information**

- For more information about domain and forest functional levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in this book and the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

- For more information about authentication, see "Designing an Authentication Strategy" in this book and the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

- For more information about authorization and access control, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Overview of Designing a Resource Authorization Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

How you choose to give and manage appropriate access to individual users and groups of users is the key to your resource authorization strategy. Logging on does not automatically give users access to the resources they require. Users must be authorized to access specific resources, but only at the level of access they need. Moreover, many users have identical needs for access to a network resource. For example, all users in the accounting department might need access to a specific color printer, so you can easily manage access by putting every member of the accounting department into a security group that is authorized to access that printer.

Because security groups are so critical to efficiently controlling access, they form the main component of your authorization strategy. Consequently, you need to know what security group types are available and how they should be used. Before you design a resource authorization strategy, you also need to be familiar with trust relationships, domain and forest functional levels, and basic security principles. By applying this information appropriately for your organization, you can design a resource authorization strategy that is scalable, easy to maintain, and cost-effective.

# Process for Designing a Resource Authorization Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing a resource authorization strategy involves selecting the best resource authorization method or methods for your environment. You must also consider and implement policies for security group management that determine who is allowed to create security groups, how security groups are named, how they can be nested, and how they are administered.

Figure 15.1 shows the process for designing a resource authorization strategy.

**Figure 15.1 Designing a Resource Authorization Strategy**



# Background Information for Designing a Resource Authorization Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you design your resource authorization strategy, it is important that you understand security group types and scope, domain and forest functional levels, group nesting, and trust relationships between domains and forests.

**Security Groups**
Security groups are used to combine user accounts, computer accounts, and other groups of accounts into manageable units. Rights or permissions are then assigned to the security groups, rather than to individual accounts. Security groups are classified according to scope. In the Microsoft® Windows® Server 2003 operating system, these scopes include:

- Local
- Domain local
- Global
- Universal

It is important that you understand group scopes and how these groups can be used before you design your resource authorization strategy.

For more information about group scope and group usage, see the following topics in Help and Support Center for Windows Server 2003:

- Group types
- Group scope
- Understanding Local Users and Groups
- Where groups can be created
- Nesting groups
- Accessing resources across domains

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

**Domain and Forest Functional Levels**
The security group options that are available depend on the domain and forest functional levels at which your organization is operating. For example, if your domain is operating at the Microsoft® Windows® 2000 mixed functional level, then you can only add global groups to local groups. However, if your domain is operating at the Windows Server 2003 functional level, you can nest global groups within other global groups, giving you much greater flexibility in managing your security groups.

If your forest is operating at the Windows Server 2003 functional level, you can also enable trusts across forests. These two-way trusts enable you to authorize users to access resources across forests. You can also enable external trusts between single domains in different forests.

For more information about domain and forest functional levels, group nesting, domain trusts, and trusts across forests, see the following topics in Help and Support Center for Windows Server 2003:

- Domain and forest functionality

- Nesting groups
- Understanding Trusts

**Terms and Definitions**

It is important to be familiar with the following terms as you design your resource authorization strategy.

**Account Group** An account group is a security group whose members are user or computer accounts, all of which require the same permissions for a resource.

**Resource Group** A resource group is a security group that has been added to the access control list (ACL) of a resource and granted a specific set of access permissions.

**Principle of Least Access** The principle of least access states that users must have access to the software, data, and devices required to perform their daily duties, but must not have access to local or network resources that are not required for their job tasks. This minimizes potential security risks.

# Establishing a Resource Authorization Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can select one of three resource authorization methods to enable users to access shared resources in your organization. It is important to select the method that is the most appropriate for your environment and for the specific resource. Figure 15.2 shows the steps in this process.

**Figure 15.2 Establishing a Resource Authorization Method**



# Selecting a Resource Authorization Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Depending on the resource and the needs of your organization, you can choose to apply any or all of the following authorization methods to enable access to shared resources:

- User/ACL method
- Account Group/ACL (AG/ACL) method
- Account Group/Resource Group (AG/RG) method
- Role-based authorization

# User/ACL Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the user/ACL method, the resource owner adds security principal accounts directly to the resource ACL. For example, Bob creates a file share and adds Alice as an authorized user, giving her read-only permission to the share. Individual resource owners often use this method.

Adding individual user accounts to individual resource ACLs might work for small organizations with few users and minimal resources. However, this approach does not scale well and is impractical for larger organizations.

Adding a large number of users to a large number of ACLs creates the following problems:

- Access to resources is inconsistent from one user to the next. For example, one engineer might have access to a laser printer, a plotter, a backup device, and many file shares. Another engineer in the same group might need access to the same resources, but have access to only a subset of those resources.

- The IT staff must handle a large volume of requests for updated or corrected access. In the previous example, Alice, the engineer who does not have sufficient access, might send individual requests to the IT group as she discovers that access to various resources is denied to her.

- There is no way to know which resources a specific user has access to without extensive tracking, and, therefore, permissions cannot be easily revoked when a user changes positions or leaves the organization.

- In very large organizations, unmanaged security descriptors waste disk space.

In most environments, it is better to use the AG/ACL or AG/RG methods rather than the limited user/ACL method. The general guideline is to assign permissions to groups, rather than to individual accounts. However, if a resource is sensitive for security reasons, it might be advisable to assign permissions to individual accounts.. For instance, suppose that you want to share a database of employee data and that you cannot control or easily determine which users belong to various account groups. In this case, rather than adding an account group to the database ACL or to the database's resource group, you might choose to add specific users directly to the database ACL. This requires more maintenance, but security concerns might outweigh maintenance requirements.

# AG/ACL Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the Account Group/ACL (AG/ACL) method, security groups, rather than individual user accounts, are added to the resource ACL, and the group is given a set of access permissions. For example, Bob might grant the group Everyone access to his public file share. This method is more scalable and more easily maintained than the user/ACL method. Typically, domain users are added to various global groups when they join an organization. Members of these global groups usually have something in common. They might work in the same building or on the same team or they might have the same job function. Rather than adding individual accounts to a resource ACL, the resource owner simply gives the domain global group access permissions. When a user joins or leaves the organization and is thus added or deleted from domain security groups, she or he is granted or denied access to the resource without any modification of resource permissions.

The AG/ACL method is scalable in that security groups can be nested if your domain is configured for Windows 2000 native or Windows Server 2003 functional level. Groups from multiple domains, or from multiple forests (if forest trusts are enabled) can be grouped together into one encompassing security group, and this security group can be added to the resource ACL. This method can scale without modification to the resource ACL, assuming that all the users require the same permissions.

This method requires more administrative effort from the resource owner, however, if different groups require different access permissions. Some groups might require read-only access, while others might require read-write access or full access. In this case, each group must be added to the ACL separately, and access permissions must be granted to each group.

Even more administrative effort from the resource owner is required if your domains are configured for Windows 2000 mixed functional level. At this functional level, group nesting is limited. Global groups can join only local groups, and global groups cannot be nested within other global groups. In this situation, neither the user/ACL nor the AG/ACL methods are typically acceptable.

If your domains are all at either Windows 2000 native or Windows Server 2003 functional level and if group nesting is allowed by your group nesting policy, then you might use either the AG/ACL method or the AG/RG method. If you have many account groups that must be given different access permissions, then the AG/ACL method might be advisable, especially if account groups and the permissions they require change frequently. Using the AG/RG method requires that you create a different resource groups for each combination of special permissions.

# AG/RG Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the Account Group/Resource Group (AG/RG) method, users with similar access requirements are grouped into account groups, and the account groups are added as members to a resource group that has been granted specific resource access permissions.

The AG/RG method is often most appropriate for large organizations with many shared resources because it is scalable and maintainable in all environments, at all domain functional levels. For example, suppose Bob maintains a print server that controls 50 printers that are used by 5,000 engineers. Bob might use the AG/RG method in several ways, but for this example he creates a local group on the printer server called Color Printer-Mail Room and gives this resource group permission to print to this print server. The group of 50 engineers working on the second floor needs access to this printer, and they are all members of an account group called Second Floor Engineers. All Bob needs to do is add the Second Floor Engineers account group to the Color Printer-Mail Room resource group. All 50 engineers now have access to the printer. Bob can use the AG/RG method for all 50 printers.

In this example, however, Bob could have used the AG/ACL method just as easily. The scalability of the AG/RG method is more apparent when Bob needs to add 25 different account groups to the printer. If he uses the AG/ACL method, he has to add 25 groups to the printer ACL and set permissions for each one. If he uses the AG/RG method, he only needs to add the account groups to the resource group.

The AG/RG method is also appropriate if a group of related resources are being shared to multiple account groups. For example, Bob might have a group of printers that are all shared to the same account groups. Suppose all of the printers on the sixth floor are used by five separate accounting teams. Rather than creating a separate local group for each printer, Bob creates one domain local group called Sixth Floor Printers. He adds this same resource group to the ACL for each printer and assigns permission to print. Now he can add the five global groups containing the five accounting team members to Sixth Floor Printers, and all five account groups gain access to all the printers on the sixth floor. If a new accounting team is formed and needs access to the printer, an account group containing the accountants' user accounts is added to the Sixth Floor Printers domain local group with no need to modify the ACLs of any printers.

Applying the AG/RG method facilitates the process of revoking users' access to one set of resources and granting them access to other resources. As users join, leave, or move within the organization, simply adding the user to or deleting the user from an account group allows or removes access to resources without the need to modify resource ACLs. For example, if engineers move to new offices on the third floor of their building, the resource owners can update printer access by removing the Engineering Users group from the Sixth Floor Printers group, and adding it to the Third Floor Printers group. In this way, all engineers receive immediate access to all of the printers near their new offices.

This method can be used in any environment. However, it is especially useful if your domains are at Windows 2000 mixed functional levels or if you need to share resources with Microsoft® Windows NT® 4.0 operating system domains. Without the ability to nest groups, it is easier to create a resource group with the appropriate permissions and to add account groups to the resource group rather than assigning individual permissions directly to each account group.

The AG/RG method is also preferable if your resources require a stable set of common permissions. For example, a file share typically requires full permissions for very few people, read-write permission for more people, and read-only permission for most people. In this situation, you might create three resource groups, one for each of the three common access levels. Account groups can be added to each resource group as needed to grant the correct access level.

The AG/RG method is also a good choice if you need to share groups of related resources that have the same access permissions, as described earlier in this section.

# Role-based Authorization

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Role-based authorization uses scripts called authorization rules to enable users with similar roles to be authorized to perform predefined sets of tasks in specific applications. This allows fine-grained control over the mapping between access control and the structure and tasks performed in your organization. In order to take advantage of role-based authorization, however, applications must be written to take advantage of this capability. In addition, the organization's Active Directory® directory service must be at the Windows Server 2003 domain functional level.

For more information about role-based authorization, see "Authorization Manager and role-based administration overview" in Help and Support Center for Windows Server 2003; also see "Role-based Access Control" in the Microsoft® Platform Software Development Kit (SDK).

# Selecting Local Groups or Domain Local Groups as Resource Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In both the AG/ACL and the AG/RG methods, you can select as resource groups either local groups on the computer that controls the resource or domain local groups in the computer's domain. This choice is typically relevant for administrators who manage network resources. If Bob shares a printer, he might create a local group to act as the resource group rather than ask a domain administrator to create a domain local group and grant him permission to manage it. Network resource managers, however, can choose either type of group as a resource group.

One advantage of using domain local groups as resource groups is that domain local groups can be managed anywhere in the domain. Creating local groups requires the resource manager to access the specific computer where the local group is to be created. Another advantage of using domain local groups as resource groups is that domain local groups are visible in Active Directory. Assuming that the domain local group is named in a meaningful way, it is relatively easy to locate the resource group and modify it to add or remove account groups.

However, there are also significant disadvantages to using domain local groups as resource groups. For example, suppose Bob is managing a file server with 500 shares, each of which has three resource groups to define the three common access levels. If Bob uses domain local groups for the resource groups, 1,500 groups will be displayed in Active Directory Users and Groups in the Microsoft Management Console (MMC). Furthermore, suppose that Bob has 50 file servers. Even with a meaningful group naming policy, it is not going to be easy to find the right resource group among 75,000 group names. In this case, Bob is much more likely to choose local groups as resource groups.

Another disadvantage of using domain local groups as resource groups is the relative difficulty of group retirement. As resources are changed, moved, or retired, the resource groups associated with the resource also must be changed, moved, or retired. A resource manager is more likely to keep up with the management of local groups than of domain local groups that are managed by another administrator. The closer the resource groups are to the actual resource, the more likely they are to be maintained or retired properly.

Security access token size also can be an issue if users are members of too many groups. A user's access token is built primarily from the security identifiers (SIDs) of the groups to which the user belongs. The default maximum token size for the Microsoft® Windows XP Professional operating system and Windows Server 2003 is 12,000 bytes. This token size enables users to belong to approximately 120 groups. If a user belongs to more than 120 groups, the buffer allocated for the user's token is exceeded. The result is that the user is unable to log on to the network or is otherwise denied access to network resources.

The maximum token size can be modified if necessary. With token-size modification, users can successfully belong to hundreds of groups. If you choose to use domain local groups as resource groups, token size issues are more likely to arise than if you use local groups on the computer where the resource is shared.

**Note**

- For more information about token size and default group membership limits, see article Q327825, "New Resolution for Problems That Occur When Users Belong to Many Groups" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at www.microsoft.com/windows/reskits/webresources.

# Defining Policies for Security Group Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Policies for security group management are a significant part of your resource authorization strategy. For example, you need to establish a policy defining who can create security groups and when they should be created. You also need policies for security group naming, security group nesting, and security group retirement. Figure 15.3 shows the tasks for defining security group policies.

**Note**

- For the purposes of this chapter, "security group policies" refers to policies for security group management and is not related to Group Policy or Group Policy objects. For more information about Group Policy, see "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit, and see "Group Policy" in Help and Support Center for Windows Server 2003.

**Figure 15.3 Defining Policies for Security Group Management**



# Defining a Security Group Creation Policy

Updated: December 7, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to define which members of your organization are allowed to create new security groups, and you need to identify the process that they use to create new security groups.

**Delegating Security Group Creation Rights**
Windows Server 2003 does not place any constraints on the ability to delegate permission to create security groups. You can edit the ACL of any security group and give any user in the forest permission to update the group's membership. This enables you to simplify the administration of groups in your organization.

By default, members of the Domain Admins, Enterprise Admins, and Account Operators groups have the Create Group Objects and Delete Group Objects permission. If it is appropriate for Account Operators in your organization to be responsible for managing the creation of security groups, there is no need to delegate this ability to others.

If there is a need for individuals other than members of the Administrators, Domain Admins, Enterprise Admins, or Account Operators groups to create security groups, you can delegate the ability to create and to delete new security groups within an organizational unit (OU) to individuals or groups in your organization. You might need to create a separate security group, such as Group Admins, and add the appropriate user accounts to this group. If the Group Admins have authority only within a single domain, this is a domain local group.

For more information about delegating the right to create security groups by using the Active Directory Delegation Wizard, see "Delegating administration" in Help and Support Center for Windows Server 2003.

**Defining a Security Group Request Process**
You need to define a process by which requests for the creation of security groups are submitted and approved. Requiring users to submit requests for the creation of new security groups and approving those requests before the groups are created limits the unnecessary proliferation of security groups in your organization.

A request for a new group should include the following information about the group:

- Purpose and scope
- Proposed membership
- Relationship to other groups
- Expected lifetime
- Group owner

Be sure that your IT staff keeps these requests on file or in a database as a source of information for various security maintenance tasks. This information is especially useful in identifying potentially obsolete groups. For more information about identifying obsolete security groups, see "Defining a Security Group Retirement Policy" later in this chapter.

# Defining a Security Group Naming Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Establishing a forest-wide or enterprise-wide naming convention for security groups helps to ensure that secure access control in your organization is not compromised. Without a universal naming convention, the potential for user error when adding or removing members and selecting the correct group increases substantially.

For example, when adding or removing members to groups, users select from a list of all groups in the forest, but cannot view information about the groups, such as membership or scope. Groups with similar or cryptic names might be selected accidentally. The consequences of granting access to the wrong group can be serious, causing members to have access to restricted resources or to be denied access to resources that are necessary for job tasks.

When establishing a security group naming convention for your organization, ensure that it does the following:

- Provides for the inclusion of information about the group's scope, purpose, and owner in its name and its description. This helps to differentiate each group from similar groups.
- Conforms to a hierarchy of standard labels to be used in a fixed order, beginning with the most general labels and ending with the most specific labels. This allows group names to be sorted alphabetically into an organized list.

In addition, consider the following when creating group names and descriptions:

- Both the name and description of a security group can include up to 256 characters.
- The first 20 characters of the name are usually visible in a list of security groups without resizing columns and scrolling. When viewing the **Properties** dialog box of the group, about 50 characters of the group name are viewable. You might want to abbreviate the organizational labels used in the group name to ensure that the distinguishing portion of the group name can be viewed in these environments.

You can apply any naming strategy that works for your organization, as long as group names provide enough information to distinguish them from other groups. A common approach is to create a security group naming standard that organizes groups according to the business structure of your organization. In this way, group names are composed of labels that represent your company's organization, such as division, department, team, and task.

Without descriptive labels, it is possible to create confusing group names. Three divisions, for example, might have similar group names for their sales groups. If someone searches for the sales group, it is difficult to know if "Sales," "Sales Group," or "AllSales" is the correct group. Adding more organizationally descriptive labels can take time and planning, but user group searches and rights assignments are more accurate as a result.

These labels are placed in order from most general to most specific, with the final label reflecting the purpose or content of the group, thus distinguishing the group name from similar names. For example, the following three group names represent different sales teams within the Avionics Division Sales Department of a corporation:

- Avionics Sales GovtContracts
- Avionics Sales RetailSales
- Avionics Sales TeleSales

Some organizations are based on domain membership. The following are examples of typical group names, where "Con" represents a domain in the fictional company "Contoso."

- Con-Sales-Clerical-Users
- Con-Sales-Mgmt-Users
- Con-Support-HelpDesk-Users

- Con-Support-HelpDesk-Resources
- Con-Support-Shipping-Users

Some organizations are based on geographical location. The following group names might represent printer resources in the Boston and Seattle remote offices:

- BOS-Printers-Floor3-Laser
- BOS-Printers-Floor3-Color
- SEA-Printers-Reception-Laser
- SEA-Printers-Lab-LargeFormat

An organized system for naming groups makes it easy to locate the correct security group, and helps protect against duplicate naming. Users can search for related groups by entering the first few letters of the group name in the **Select Users, Contacts, Computers or Groups** dialog box. For example, they can filter the list for "Avionics Sales" groups.

The list of group names in the **Select Users, Contacts, Computers or Groups** dialog box also displays each group's description string. This added field can be used to include other attributes of the group, such as the group scope, the name of the person who maintains the group, and a brief statement of the group's purpose. The following are examples of group descriptions:

- *Avionics Support HelpDesk Resources:* Domain local resource group for help desk resource ACLs, maintained by Bjones x2344.
- *Avionics HumanRes AllUsers:* Global account group of all human resources user accounts, maintained by Ssmith x9488.

**Important**

- Windows Server 2003 does not provide any software support for enforcing a naming standard. Establish your naming policy and communicate it to all employees in your organization who have been delegated the right to create security groups.

# Defining a Security Group Nesting Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Security group nesting occurs when one security group is made a member of another security group, and the nested group inherits all of the privileges and permissions that are granted to the parent.

**Note**

- A domain must be configured at the Windows 2000 native or Windows Server 2003 functional level in order for global groups to be nested within global groups or for domain local groups to be nested within domain local groups.

Unrestrained group nesting can result in access token size problems because the token contains the SIDs for each group of which the user is a member, either directly or indirectly. The default group membership limitation is 120 groups. For more information about issues related to group membership and access token size, see "Selecting Local Groups or Domain Local Groups as Resource Groups" earlier in this chapter. Additionally, if group nesting is not constrained by a nesting policy, it becomes difficult to know exactly which permissions might be inherited by members of a security group that is nested within another, or several other, security groups.

If your enterprise is organized according to the AG/RG model, you need to allow for some degree of nesting without allowing nested groups to proliferate. Design and follow an explicit nesting strategy so that nesting relationships among security groups are predefined and well understood. This helps to guard against including users in too many groups.

If you choose to apply a security group nesting policy in your organization, consider drafting specific guidelines for that policy and communicating them to all employees who have permission to add members to a group.

# Creating a Hierarchy of Nested Security Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Nested groups allow you to provide company-wide or department-wide access to resources with minimum maintenance. Placing every team account group into a single company-wide resource group is not an effective solution because it requires the creation and maintenance of a large number of membership links. To use nested groups, administrators create a series of account groups that represent the managerial divisions of the company.

The top account group might be called "All Employees," and would be attached to a resource group that gives access to resources and shared directories.

The next level might contain account groups that represent major divisions of the company. Each group at this level is a member of All Employees, and is attached to a resource group giving access to shares and other resources appropriate to the division it represents.

Within a division, the next level of account groups might represent departments. Shared resources for the department might include project schedules, meeting schedules, vacation schedules, or any network information appropriate to the whole department. The department account groups are all members of the division account group.

Within a department, the management structure can be organized into security groups to any required level of specificity. These might be team account groups, and might represent leaf nodes in the organization's hierarchical tree.

With this group hierarchy in place, you can give a new employee instant access to the resources of the team, the department, the division, and the company as a whole by placing him or her in a team account group. This system supports the principle of least access because the new employee cannot view the resources of adjacent teams, other departments, or other divisions.

Figure 15.4 shows an example of the nested security group hierarchy for the Avionics Division group of an organization; the vertical arrows represent group membership.

**Figure 15.4 Nested Account Group Hierarchy**



# Modifying a Nested Security Group Hierarchy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The hierarchical structure of nested security groups gives the broadest access to users in the leaf nodes of the tree, the team account groups. You can modify the structure to allow department managers full access to all team resources, for example, or to allow the company's CEO access to all company resources. Figure 15.5 shows how the nested security group hierarchy illustrated in Figure 15.4 might be modified to allow all managers full access to all team resources; the arrows represent group membership.

**Figure 15.5 Modified Nested Security Group Hierarchy**




To support this type of access, you create another hierarchical series of account groups, with the membership links going in the opposite direction. For example, you might create a group called Avionics Sales Managers, and add the appropriate users to the group, then add this group to the membership of all Sales team account groups. In this way, the Avionics Sales Managers are able to view the files of all Avionics Sales teams.

You can apply this model all the way up the organizational chart, giving the CEO complete access to every file in the forest. This approach, however, is problematic for two reasons. First, high-ranking executives are priority targets for attackers and do not need access to thousands of files scattered throughout the forest. An intruder who intercepts the password of the CEO can cause unlimited damage. Second, in a large organization this strategy can fill user access tokens with hundreds of group SIDs, leading to significant logon delays, denial of access to resources, or the inability to log on. For more information about issues related to group membership and access token size, see "Selecting Local Groups or Domain Local Groups as Resource Groups" earlier in this chapter.

# Defining a Security Group Retirement Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When security groups become obsolete due to personnel changes, project completion, or corporate reorganizations, they need to be identified and retired (deleted) to minimize security risks. This task usually falls to the organization's IT department. It is important to create a policy for the retirement of security groups.

In this policy, include procedures for:

- Gathering information about security groups, especially information related to the expected lifetime of the group or the group's renewal interval. For more information about sources for this information, see "Defining a Security Group Creation Policy" earlier in this chapter.
- Identifying potentially obsolete security groups.
- Deleting obsolete security groups from Active Directory.

**Important**

- Deleting a security group from Active Directory is non-reversible because access to resources is based on the group's unique SID. If the wrong group is deleted, or if a group is deleted because it is thought to be obsolete when it is not, the IT department must create a new group in place of the deleted group and reintegrate it into the domain.

**Identifying Obsolete Security Groups**

Although account groups for very small teams might not change frequently, large account groups experience almost constant turnover in membership. If an account group's membership has not changed at all for some time, the group might be obsolete. The Active Directory Users and Computers snap-in allows you to search for potentially obsolete groups by entering a Lightweight Directory Access Protocol (LDAP) query for a list of security groups whose definitions, including membership, have not changed since a specified date. A query in the format (&(objectCategory=group)(whenChanged<=YYYYMMDDHHMMSS.0Z)) returns groups that have not been modified past the specified date. For example, to query for groups that have not been modified since midnight on December 31, 2002 in Pacific Standard Time, the query is (&(objectCategory=group)(whenChanged<=20021231240000.0-8)). The last portion of the query (0-8) signifies the Pacific Standard time zone, which is eight hours less than Greenwich Mean Time.

**Note**

- For more information about building LDAP queries and about generalized time, see the Windows Platform SDK link on the Web Resources page at www.microsoft.com/windows/reskits/webresources.

After you have identified groups whose membership has not changed for a period of time, you can verify the status of each group by querying the group owner, the group members, or the appropriate departmental administrators. Be sure to keep a record of your investigation for future reference.

Resource groups typically become obsolete when the resource is no longer needed, perhaps because projects requiring the resource have ended or because the hardware resource is moved or retired. Because resource groups are often configured as local groups on the computer that controls the resource, identifying obsolete resource groups is often a task for the computer owner or the resource manager. Finding obsolete local groups configured as resource groups is more challenging than finding domain-based account groups. Computers controlling resources are typically not domain controllers, so Active Directory is not available, and LDAP queries for unmodified local groups are not possible.

One approach is not to track local resource groups because local groups do not populate Active Directory. If your reason for identifying obsolete groups is to reduce the number of Active Directory objects, monitoring domain-based account and resource groups might be all that you require.

Another approach is to document and create a database of shared network resources. Typically, this database would include information about corporate resources rather than file shares or similar resources shared from a user's desktop, for instance. If the share type, location, owner, creation date, and any other relevant information are recorded, then periodic reviews can be

conducted to verify that resources and their associated resource groups are needed or obsolete.

**Deleting Obsolete Security Groups**
When you have identified an obsolete group, you can disable it for a trial period before deleting it from your domain. You can disable a security group by changing it into a distribution group for a set period of time, such as two months. Changing a security group into a distribution group does not alter its SID, scope, or membership; however, it effectively disables all access provided by the group because SIDs of distribution groups are not included in a user's access token. If you do not receive any notification that a user's access permissions have changed after the established wait period has expired, you can safely delete the group.

You must have the appropriate permissions to delete a security group. By default, Account Operators, Domain Admins, and Enterprise Admins have the ability to delete security groups. This permission can be granted to other individuals or administrative groups by editing the ACLs of OUs and containers.

### Caution

- Group deletion is immediate and permanent. Limit the number of persons to whom you delegate the ability to delete groups.

# Delegating Security Group Maintenance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In large organizations, it is not possible for IT staff to be responsible for performing routine membership maintenance on large numbers of account groups, or to administer the ACLs and resource groups for resources in all departments. After a security group has been created and named, you can delegate the management of the group to other members or your organization. This can relieve the burden on your IT department.

Figure 15.6 shows the tasks for delegating security group maintenance.

**Figure 15.6 Delegating Security Group Maintenance**



# Identifying Individuals to Maintain Security Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By delegating security group maintenance to the appropriate individuals, you can ensure that requests for changes in membership are evaluated by individuals who can judge the appropriateness of the request, have the authority to make the change, and are motivated to keep group membership and access permissions correct and up-to-date.

Ideally, an employee who is familiar with the personnel, such as a department administrative assistant, is responsible for managing the membership of an account group. Because security groups in Windows Server 2003 can also be used as Microsoft® Exchange 2000 mailing lists, it is both convenient and cost-effective to have an individual maintain the account group for both purposes. This requires delegating read and write permissions for the group to the selected individual.

It is recommended that resource owners set ACLs because they are best able to supervise access to the resource. For example, the file server operator who maintains a department's shared directories is the logical candidate to be responsible for adding resource groups to the ACLs of the shares. Similarly, the person who is responsible for departmental printers is the best choice to manage the membership of the applicable printer resource groups.

# Delegating Account Group Maintenance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To delegate permission to maintain account group memberships, you can apply the *account group/resource group/admin group model*, by which you extend the AG/RG model to accommodate a third kind of security group, the departmental admin group.

*Admin groups* are usually domain local groups that consist of members of the local business unit, such as managers or trusted clerical workers. A *departmental* admin group has permission to manage the membership of the department's security groups, and sometimes manages other department resources as well. The departmental admin group is added to the ACL of a business unit's OU, or to the ACLs of the individual user and resource groups. Although the departmental admin group has permission to manage group membership, it cannot create and delete groups.

Each department has its own departmental admin group with specific permission to change membership of the department's account groups. Because membership in the departmental admin group gives a user the ability to add any forest user to departmental account groups, the departmental administrators must be highly trusted employees.

Consequently, you also must carefully determine who controls the membership of the admin group itself. Allowing departmental administrators to self-administer their admin groups is a security risk, because such delegated authority can be misused. Relatively little effort is required to maintain the departmental admin groups, so it is more secure to have the central IT department administer them. In this way, the IT department can delegate account group maintenance to specific individuals in each department, and these assignments cannot be changed without IT participation.

# Delegating Resource Group Maintenance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the AG/RG model, related resources are collected into resource groups, although the resources themselves are not members of the resource groups. Instead, the ACLs of the individual resources are edited to grant the resource group appropriate access, and then account groups are added as members of the resource group, which gives the members of the account group access to the resources.

If you have adopted this model for your organization, you need to decide who sets the ACLs on the resources, and who controls the membership of the resource groups. In some cases, the departmental administrators (who maintain departmental account groups) have control over both the resources and resource groups in their departments.

In large organizations, however, the departmental admin group can become overwhelmed with requests for resource access or with other resource management duties. One solution to this problem is to enable resource owners to control access to their resources. The resource owner has full knowledge of the resource and is motivated both to make the resource available to legitimate users and to protect it from unauthorized access. Resource owners have control over access to their resources by default in Windows Server 2003.

For resource owners to have full control over access to their resources, they also must control the membership of the resource group. In this way, they control which account groups can become members of their resource groups.

In some cases, you must delegate resource group maintenance to resource owners. To do so, you must establish that the resource owner is responsible for editing the ACLs of the resource in order to grant appropriate access to specific resource groups. By default, the owner of a file or directory, a Server Operator, a Domain Admin, or an Enterprise Admin has permission to edit the ACLs of a resource. You must also grant resource owners control over the membership of the resource groups (but not the right to create or delete the resource groups) that must access their resource.

For more information about editing ACLs, see "Access Control" in Help and Support Center for Windows Server 2003.

# Additional Resources for Resource Authorization Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this chapter.

**Related Information**

- The *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit) for more information about Active Directory functional levels.
- The *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit) for more information about authorization, access control, domain and forest trusts, and logon and authentication.
- "Designing an Authentication Strategy" in this book.
- "Enabling Advanced Windows Server 2003 Active Directory Features" in this book.
- The TechNet link on the Web Resources page at http://windows.microsoft.com/windows/reskits/webresources.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Group types"
- "Group scope"
- "Understanding Local Users and Groups"
- "Where groups can be created"
- "Nesting groups"
- "Accessing resources across domains"
- "Domain and forest functionality"
- "Nesting groups"
- "Understanding Trusts"
- "Authorization Manager and role-based administration overview" in for more information about role-based authorization.
- "Group Policy" in Help and Support Center .
- "Delegating administration" in Help and Support Center for more information about delegating the right to create security groups by using the Active Directory Delegation Wizard.
- "Access Control" in Help and Support Center for more information about editing ACLs.

# Designing a Public Key Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| **Note** |
| --- |
| To download a copy of this document, see http://go.microsoft.com/fwlink/?LinkId=50360. |

Windows® Server 2003 enables a variety of secure applications and business scenarios based on the use of digital certificates. Before you can use digital certificates, however, you need to design a public key infrastructure (PKI), which involves planning configuration options for one or more certification authorities, preparing certificates to meet the needs of your organization, and creating a PKI management plan.

**In This Chapter**


Overview of the PKI Design Process

Defining Certificate Requirements

Designing Your CA Infrastructure

Extending Your CA Infrastructure

Defining Certificate Configuration Options

Creating a Certificate Management Plan

Deploying the PKI

Additional Resources for Designing a Public Key Infrastructure


**Related Information**

- For more information about Windows Server 2003 public key features, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection.
- For more information about using certificates in conjunction with Encrypting File System, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection.
- For more information about deploying smart cards, see "Planning a Smart Card Deployment" in this book.

# Overview of the PKI Design Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Organizations use a variety of technology solutions to enable essential business processes, such as online ordering, exchanges of contracts, and remote access. A public key infrastructure based on Microsoft Windows Server 2003 Certificate Services provides a means by which organizations can secure these critical internal and external processes.

Deploying a PKI allows you to perform tasks such as:

- Digitally signing files such as documents and applications.
- Securing e-mail from unintended viewers.
- Enabling secure connections between computers, even if they are connected over the public Internet or through a wireless network.
- Enhancing user authentication through the use of smart cards.

If your organization does not currently have a public key infrastructure, begin the process of designing a new public key infrastructure by identifying the certificate requirements for your organization. If your organization already uses a public key infrastructure based on Microsoft® Windows NT® version 4.0, Microsoft® Windows® 2000, or third-party certificate services, you can improve your PKI capabilities by taking advantage of new and enhanced features in Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition. When you have completed the PKI design process, you can deploy a public key infrastructure that provides solutions for all of your internal security requirements, as well as security requirements for business exchanges with external customers or business partners.

**Note**

- For a list of the job aids that are available to assist you with the PKI design process, see "Additional Resources for Designing a Public Key Infrastructure" later in this chapter.

# Process for Designing a PKI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing a PKI for your organization involves defining your certificate requirements, creating a design for your infrastructure, creating a certificate management plan, and deploying your PKI solution. Figure 16.1 shows the steps that are involved in designing a public key infrastructure.

**Figure 16.1 Designing a PKI**




**Note**

- The steps involved in designing a PKI are interdependent. For example, defining certificate server configurations, locations, and roles has a significant impact on how you address key certificate management issues. Your evolving certificate management standards in turn have significant implications for the certificate server roles, locations, and configurations that you develop.

# Basic PKI Concepts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Public key infrastructure is the term used to describe the laws, policies, procedures, standards, and software that regulate or control the operation of certificates and public and private keys. More specifically, a PKI is a system of digital certificates, certification authorities, and other registration authorities that verify and authenticate the validity of each party involved in an electronic transaction.

A PKI consists of the following basic components:

**Digital certificates** Electronic credentials, consisting of public keys, which are used to sign and encrypt data. Digital certificates provide the foundation of a PKI.

**One or more certification authorities (CAs)** Trusted entities or services that issue digital certificates. When multiple CAs are used, they are typically arranged in a carefully prescribed order and perform specialized tasks, such as issuing certificates to subordinate CAs or issuing certificates to users.

**Certificate policy and practice statements** The two documents that outline how the CA and its certificates are to be used, the degree of trust that can be placed in these certificates, legal liabilities if the trust is broken, and so on.

**Certificate repositories** A directory service or other location where certificates are stored and published. In a Windows Server 2003 domain environment, the Active Directory® directory service is the most likely publication point for certificates issued by Windows Server 2003–based CAs.

**Certificate revocation lists (CRL)** Lists of certificates that have been revoked before reaching the scheduled expiration date.

**Note**

- With Certificate Services in Windows Server 2003, Microsoft introduces a new type of certificate revocation list called a delta CRL, which allows you to publish information about recently revoked certificates more frequently without using the bandwidth required for publishing full CRLs.

**Certificate trust lists** These are signed lists, which are located on the client, of trusted CA certificates. Certificate trust means that a certificate is part of a certificate trust list (CTL) or that the CTL contains a trusted certificate from another CA that is part of the certificate's certificate chain. Windows Server 2003 domain administrators can use Group Policy objects (GPOs) to publish and maintain CTLs.

**Key archival and recovery** A feature that makes it possible to archive and recover the private key portion of a public-private key pair, in the event that a user loses his or her private keys, or an administrator needs to assume the role of a user for data access or data recovery. Private key recovery does not recover any data or messages; it merely enables the recovery process.

**Public key standards** Standards developed to describe the syntax for digital signing and encrypting of messages and to ensure that a user has an appropriate private key. To maximize interoperability with third-party applications that use public key technology, the Windows Server 2003 PKI is based on the standards recommended by the Public-Key Infrastructure (X.509) (PKIX) working group of the Internet Engineering Task Force (IETF). Other standards that the IETF has recommended also have a significant impact on public key infrastructure interoperability, including standards for Transport Layer Security (TLS), Secure/Multipurpose Internet Mail Extensions (S/MIME), and Internet Protocol security (IPSec).

# Windows Server 2003 PKI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use PKI-based applications on workstations and servers running Microsoft® Windows® XP Professional, Windows Server 2003, Windows® 2000, or Windows NT 4.0, as well as on workstations running Microsoft® Windows® 95 and Microsoft® Windows® 98. The ability to create and manage a PKI is available in Microsoft® Windows NT® 4.0 Server, Microsoft® Windows® 2000 Server, and Windows Server 2003. However, Windows Server 2003 provides more extensive support for a PKI.

In addition, a growing number of applications and system services that require the secure transfer of information also rely on the Windows Server 2003 PKI. Applications that are enabled for certificate-based security include Microsoft® Outlook®, Internet Explorer®, Internet Information Services, Microsoft® Exchange Server, Microsoft® Commerce Server 2000 and Commerce Server 2002, Outlook Express, and Microsoft® SQL Server™. A number of third-party applications also take advantage of the Windows Server 2003 PKI.

**How a Public Key Infrastructure Works**

A Windows Server 2003 PKI makes it possible for an organization to do the following:

- **Publish certificates.** The PKI administrator makes certificate templates available to clients (users, services, applications, and computers) and enables additional CAs to issue certificates.

- **Enroll clients.** To participate in a PKI, users, services, or computers must request and receive certificates from an issuing CA or a Registration Authority (RA). Typically, enrollment is initiated when a requester provides unique information and a newly generated public key. The CA administrator or enrollment agent uses the information provided to authenticate the identity of the requester before issuing a certificate.

- **Use certificates.** Clients use their certificates, which are validated or invalidated in a timely manner as long as CAs and certificate revocation lists are available to verify or deny their authenticity. If they are validated, a PKI provides an easy way for users to use keys in conjunction with applications that perform public key cryptographic operations, making it possible to provide security for e-mail, e-commerce, and networks.

- **Renew or revoke certificates.** A well-designed PKI makes it easy for you to renew or revoke existing certificates, and to manage the trust level associated with certificates used by different clients or for different applications.

The status of a public key certificate is determined by means of the chain building process. Chain building is the process of building a trust chain, or certification path, from the end certificate to a root CA that is trusted by the security principal. Figure 16.2 shows a certification path in a two-level CA hierarchy.

**Figure 16.2 Certification Path in a Two-Level CA Hierarchy**




In this example, the issuing CA issued the User certificate, and the root CA issued the certificate of the issuing CA. This is considered a trusted chain, because it terminates with a root CA certificate that has been designed and implemented to meet the highest degree of trust.

The chain building process validates the certification path by checking each certificate in the certification path from the end certificate to the certificate of the root CA. If the CryptoAPI discovers a problem with one of the certificates in the path, or if it cannot find a certificate, the certification path is either considered invalid or is given less weight than a fully validated certificate.

For more information about how PKIs function, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Defining Certificate Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use a Windows Server 2003 public key infrastructure to provide a wide range of strong, scalable, cryptography-based solutions for network and information security. The value of the information that you want to protect, as well as the costs involved with implementing a strong security system, impact the level of security that you choose for your organization.

Figure 16.3 shows the steps that are involved in determining your certificate requirements.

**Figure 16.3 Defining Certificate Requirements**



# Determining Secure Application Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin to design your public key infrastructure and configure certificate services, you need to define the security needs of your organization. For example, does your organization require electronic purchasing, secure e-mail, secure connections for roaming users, or digital signing of files? If so, you need to configure CAs to issue and manage certificates for each of these business solutions.

A Windows Server 2003 PKI can support the following security applications:

- Digital signatures
- Secure e-mail
- Software code signing
- Internet authentication
- IP security
- Smart card logon
- Encrypting file system user and recovery certificates
- 802.1x authentication

**Digital Signatures**

A digital signature is a means for originators of a message, file, or other digitally encoded information to bind their identities to the data. This can be extremely useful for important documents such as legal opinions and contracts. The process of digitally signing information involves transforming the information, together with some secret information held by the sender, into a tag called a *signature*. Digital signatures are used in public key environments to help secure electronic commerce transactions by providing verification that the individual sending the message is who he or she claims to be, and by confirming that the message received is identical to the message sent.

You can use digital signatures even when data is distributed in plaintext, such as with e-mail. In this case, while the sensitivity of the message itself does not warrant encryption, it can be important as a means to ensure that the data is in its original form and has not been sent by an impostor.

One way that your organization can capitalize on the use of digital signatures is by using CAPICOM. CAPICOM is an ActiveX control that provides a COM interface to Microsoft CryptoAPI. It exposes a select set of CryptoAPI functions to enable application developers to incorporate digital signing and encryption functionality into their applications. Because CAPICOM uses COM, application developers can access this functionality in a number of programming environments, such as Microsoft® Visual Basic®, Microsoft® Visual Basic® Scripting Edition, Active Server Pages, Microsoft® JScript®, C++, and others. CAPICOM is packaged as an ActiveX control, allowing Web developers to use it in Web-based applications as well.

You can use CAPICOM for:

- Digitally signing data with a smart card or software key.
- Verifying digitally signed data.
- Displaying certificate information.
- Inspecting certificate properties such as subject name or expiration date.
- Adding and removing certificates from the certificate stores.

- Encrypting and decrypting data with a password.
- Encrypting and decrypting data by means of public keys and certificates.

**Secure E-mail**

Standard Internet mail is sent as plaintext over open networks with no security. In the increasingly interconnected network environments of today, intruders can monitor mail servers and network traffic to obtain proprietary or sensitive information. You also risk exposure of proprietary and confidential business information when you send mail over the Internet from within your organization.

Another form of intrusion is impersonation. On IP networks, anyone can impersonate mail senders by using readily available tools to counterfeit the originating IP address and mail headers. When you use standard Internet mail, you can never be sure who really sent a message or whether the contents of the message are valid. Moreover, malicious attackers can use mail to cause harm to the recipient computers and networks (for example, by sending attachments that contain viruses).

For these reasons, many organizations have placed a high priority on implementing secure mail services that provide confidential communication, data integrity, and non-repudiation. A Windows Server 2003 public key infrastructure allows you to enhance e-mail security by using certificates to prove the identity of the sender, the point of origin of the mail, and the authenticity of the message. It also makes it possible to encrypt mail. To provide message authentication, data integrity, and non-repudiation, secure mail clients can sign messages with the private key of the sender before sending the messages. The recipients then use the public key of the sender to verify the message by checking the digital signature.

S/MIME clients that run on any platform or operating system can exchange secure mail because all cryptographic functions are performed on the clients, not on the servers.

**Software Code Signing**

A growing number of applications, ActiveX® controls, and Java applets are being downloaded and installed on computers with little or no user notification.

In response to this problem, Microsoft introduced Authenticode™ digital signature technology in 1996, and in 1997 added significant enhancements to this technology. Authenticode technology allows software publishers to digitally sign any form of active content, including multiple-file archives. These signatures can be used to verify both the publishers of the content and the content integrity at time of download. Many software vendors already sign their applications and you can use these signatures to manage the software applications used on your network.

Authenticode relies on a certification authority structure in which a small number of commercial CAs issue software-publishing certificates. If you want to expand the use of software-publishing certificates in your own organization, the Windows 2000 and Windows Server 2003 PKI allows you to issue your own Authenticode certificates to internal developers or contractors and allows any employee to verify the origin and integrity of downloaded applications.

**Internet Authentication**

The Internet has become a key element in the growth of electronic commerce. However, for many users, security considerations impact how much and what kind of information they are willing to share across the Internet. The major concerns are:

- **Confidentiality.** Data that is transferred between clients and servers needs to be encrypted to prevent its exposure over public Internet links.
- **Server authentication.** Clients need a way to verify the identity of the servers they are communicating with.
- **Client authentication.** Servers need a way to verify the identity of clients.

Client authentication of the server takes place when the client verifies the cryptographic signatures on the certificate of the server, and any intermediate CA certificates, to a root CA certificate located in the trusted root store on the client. Server authentication of the client is accomplished when the server verifies the cryptographic signatures on the certificate of the client, and any intermediate CA certificates, to a root CA installed in the trusted root store on the server. When the identity of the client is verified, the server can establish a security context to determine what resources the client is allowed or not allowed to use on the server.

**IP Security**

Windows 2000 and Windows Server 2003 incorporate Internet Protocol security (IPSec) to protect data moving across the

network. IPSec is a suite of protocols that allows encrypted and digitally signed communication between two computers or between a computer and a router over an insecure network. The encryption is applied at the IP network layer, which means that it is transparent to most applications that use specific protocols for network communication. IPSec provides end-to-end security, meaning that the IP packets are encrypted or signed by the sending entity, are unreadable en route, and can be decrypted only by the recipient entity. Due to a special algorithm for generating the same shared encryption key at both ends of the connection, the key does not need to be passed over the network.

You do not need to use public key technology to use IPSec; instead you can use the Kerberos version 5 authentication protocol or shared secret keys that are communicated securely by means of an out-of-band mechanism at the network end points for encryption. However, if you use public key technology in conjunction with IPSec, you can create a scalable distributed trust architecture in which IPSec devices can mutually authenticate each other and agree upon encryption keys without relying on prearranged shared secrets, either out-of-band or in-band. This, in turn, yields a higher level of security than IPSec without a PKI.

For more information about deploying IPSec solutions, see "Deploying IPSec" in *Deploying Network Services* of this kit.

**Smart Card Logon**

Smart card logon is integrated with the Kerberos version 5 authentication protocol implemented in Windows Server 2003. When smart card logon is enabled, the system recognizes a smart-card insertion event as an alternative to the standard Ctrl + Alt + Del secure attention sequence to initiate a logon. The user is then prompted for the smart card PIN code, which controls access to operations performed by using the private key stored on the smart card. In this system, the smart card also contains a copy of the certificate of the user (issued by an enterprise CA). This allows the user to roam within the domain.

Smart cards enhance the security of your organization by allowing you to store extremely strong credentials in an easy-to-use form. Requiring a physical smart card for authentication virtually eliminates the potential for spoofing the identities of your users across a network. In addition, you can also use smart card applications in conjunction with virtual private networks and certificate mapping, and in e-commerce. For many organizations, the potential to use smart cards for logon is one of the most compelling reasons for implementing a public key infrastructure.

For more information about deploying smart cards, see "Planning a Smart Card Deployment" in this book.

**Encrypting File System Use and Recovery**

The Windows Server 2003 Encrypting File System (EFS) allows users and services to encrypt their data to prevent others who authenticate to the system from viewing the information. However, EFS also provides for data recovery if another means is needed to access this data — for example, if the user who encrypted the data leaves the organization, or if the original encryption key is lost. To support this requirement, EFS allows recovery agents to configure public keys that are used to enable file recovery. The recovery key only makes available the randomly generated file encryption key, not a private key of the user. This ensures that no other private information is accidentally revealed to the recovery agent.

### Note

- You do not need to have a public key infrastructure to use Windows Server 2003 Encrypting File System. However, the use of public keys improves the manageability of EFS. In a Windows domain environment, it is recommended that EFS be used in conjunction with a PKI.

For more information about EFS, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**Wireless (802.1x) Authentication**

A growing number of organizations and facilities such as airports and hotels are implementing wireless network access. This creates the challenge of ensuring that:

- Only authenticated users can access the wireless network.
- Data transmitted across the wireless network cannot be intercepted.

Public key infrastructures, in conjunction with the IEEE 802.1x standard for port-based network access control, support both of these goals by providing centralized user identification, authentication, dynamic key management, and accounting to provide authenticated network access to 802.11 wireless networks and to wired Ethernet networks.

For more information about deploying a wireless network, see "Deploying a Wireless LAN" in this book.

# Determining Certificate Requirements for Users, Computers, and Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have identified the security technologies that you need to implement to meet the business needs of your organization, you need to identify the categories of users, computers, and services that will use these technologies and for which you need to provide certificate services. For example, certificate use might be based on job function, location, organizational structure, or a combination of these three, or all computers or users in the organization might use certain certificate applications.

For each of the groups that you have identified, you need to determine:

- **The types of certificates to be issued.** This is based on the security application requirements of your organization and the design of your PKI infrastructure.
- **The number of users, computers, and applications that need certificates.** This number can include as few as one or as many users, computers, or applications as are in an entire organization.
- **The physical location of the users, computers, and applications that need certificates.** Different certificate solutions might be required for users in remote offices or for users who travel frequently than are required for users in the headquarters office of an organization. Also, requirements can differ based on geography. For example, you might want to restrict users in one country/region from using their certificates to access data in an organizational business unit in another country/region.
- **The level of security that is required to support the users, computers, and applications that need certificates.** Users who work with sensitive information typically require higher levels of security than other members of the organization.
- **The number of certificates required for each user, computer, and application.** In some cases, one certificate can meet all requirements. Other times, you need multiple certificates to enable specific applications and meet specific security requirements.
- **The enrollment requirements for each certificate that you plan to issue.** For example, do users have to present one or more pieces of physical identification, such as a driver's license, or can they simply request a certificate electronically?

**Note**

- For a worksheet to assist you in identifying user certificate requirements, see "Summary of User Certificate Requirements" (DSSPKI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Summary of User Certificate Requirements" on the Web at http://www.microsoft.com/reskit).

# Documenting Certificate Policies and Practices

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designing a public key infrastructure involves configuring certificates and certification authorities, developing support procedures, and establishing a system of checks and balances for administrative authority. Only by effectively addressing both the technical and administrative issues related to your public key infrastructure can you ensure that your certificate services provide the level of security that your organization requires

It is helpful to record the decisions that you make as you design your PKI by creating certificate policy statements and certificate practice statements. These documents assist you in planning and in communicating with individuals and businesses outside your organization. For many organizations and certificate uses, certificate policy statements and certificate practice statements are considered legal documents or legal disclaimers.

In general, the IT department is responsible for setting and maintaining PKI policies and practices. However, because of the legal, financial, and tactical uses of PKIs, representatives from outside the IT department, such as human resources, finance, legal, and marketing, might also be involved in establishing certificate policies.

A *certificate policy* is a set of rules that indicates the applicability of a certificate to a particular group of clients or applications that have common security requirements. Certificate policy statements generally include the following types of information:

- How users are authenticated to the CA.
- Legal issues, such as liability, that might arise if the CA is either compromised or used for something other than its intended purpose.
- The intended purpose of the certificate.
- Private key management requirements, such as storage on smart cards or other hardware devices.
- Whether the private key can be exported or archived
- Requirements for users of the certificates, including what users must do in the event that their private keys are lost or compromised.
- Requirements for certificate enrollment and renewal.
- Minimum length for the public key and private key pairs.

**Important**

- You can implement many of the decisions that you document in your certificate policy statements by creating a CAPolicy.inf file and copying it to the system directory of the CA before the CA is installed or renewed. For more information about CAPolicy.inf file contents and configuration, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

A *certificate practice statement* is a statement of the practices that IT uses to manage the certificates that it issues. It describes how the certificate policy of the organization is interpreted in the context of the system architecture of the organization and its operating procedures. The IT department is responsible for preparing and maintaining the certificate practice statement.

A certificate practice statement usually includes the following types of information:

- Positive identification of the CA (including CA name, server name, and DNS address).
- The certificate policies that are implemented by the CA and the certificate types that are issued.

- The policies, procedures, and processes for issuing, renewing, and recovering certificates.
- Cryptographic algorithms, cryptographic service providers (CSPs), and key length used for the CA certificate.
- Physical, network, and procedural security for the CA.
- The certificate lifetime of each certificate issued by the CA.
- Policies for revoking certificates, including conditions for certificate revocation, such as employee termination and misuse of user rights.
- Policies for CRLs, including where to locate CRL distribution points and how often CRLs are published.
- A policy for renewing the certificate of the CA before its expiration.

It is best to create a certificate practice statement for each CA in your public key infrastructure. The certificate practice statement associated with a CA can incorporate multiple certificate policies. Also, to consolidate information, the certificate practice statement for a subordinate CA can reference common or general information in the certificate practice statement of a parent CA.

For an outline to assist you in creating a certificate practice statement, see "Certificate Practice Statement Outline" (DSSPKI_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Certificate Practice Statement Outline" on the Web at http://www.microsoft.com/reskit).

**Important**

- In some situations, such as when digital signatures are used on binding contracts, the certificate practice statement can also be considered a legal statement about the level of security that is provided and the safeguards that are being used to establish and maintain the security level.

# Example: Defining Certificate Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization decides to implement a public key infrastructure because a number of business units within the organization are using certificate services independently. The business units use similar infrastructures that include many of the same components — such as CAs and certificate templates — and have similar goals. Therefore, the organization develops a PKI with a central corporate root that also allows individual business units to implement certificate services for their specific needs.

The organization chooses to use certificate services for the following:

- E-mail
- Internet authentication
- Encrypting File System
- Software code signing
- Smart card logon

In addition, they identify the following requirements:

- All users throughout the organization are required to use certificates in order to secure e-mail traffic.
- Individual business units need to use Internet authentication to facilitate the sharing of data on their local networks with their joint venture partners.
- All users are able to use Encrypting File System.
- Developers and network administrators must use software code signing for the custom applications and scripts of the organization.
- Administrators are required to log on using a smart card before they can perform certain tasks, such as administering domain controllers.

The organization then divides these requirements into the following security classifications:

- Medium security, which includes the e-mail and EFS certificates.
- Internal high security, which includes the software code signing and smart card logon certificates, and serves the needs of network administrators and developers.
- External high security, which includes the Internet authentication certificates and meets the need of the organization to share information with joint venture partners.

Figure 16.4 shows an example of the User Certificate Requirements worksheet that the organization created to summarize these classifications.

**Figure 16.4 Example of a User Certificate Requirements Worksheet**

| Summary of User Certificate Requirements | | | |
|---|---|---|---|
| Prepared By | | Date | |

Instructions: Col.1 List the names of groups of users, computers, services, business units, or network roles requiring certificates. Col. 2 Enter the number of clients in the group that will be issued certificates. Col. 3 Indicate whether the users are internal employees or external customers, and where they are located. If appropriate, include the owner name. Col. 4 List one or more certificate-enabled applications that will be used by the users in the group in that location. Indicate the name of either the parent domain or the account OU. Col 5 Enter a value such as high, medium, or low.

| Group | Number | Location | Applications | Security Level |
|---|---|---|---|---|
| Users | 45,000 | Internal | E-mail, EFS | Medium |
| Administrators and Developers | 250 | Internal | Smart card logon, code signing | High |
| Computers | 45,000 | Internal | Wireless Authentication | High |
| Partners | 35 | External | Internet Authentication | High |
| | | | | |

For a worksheet to assist you in documenting your certificate requirements, see "User Certificate Requirements" (DSSPKI_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "User Certificate Requirements" on the Web at http://www.microsoft.com/reskit).

After they have planned the trust relationships for the internal CA infrastructure and extended external CA infrastructure, the organization can design its certificates and certificate management processes. Administrators must examine the security and user requirements to develop a secure certificate services solution. For more information about designing certificates and configuring CAs, see "Creating a Certificate Management Plan" later in this chapter.

# Designing Your CA Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To support the certificate-based applications of your organization, you must establish a framework of linked CAs that are responsible for issuing, validating, renewing, and revoking certificates as needed. The goal in establishing a CA infrastructure is to provide reliable service to users, manageability for administrators, and flexibility to meet both current and future needs, while maintaining an optimum level of security for the organization.

Figure 16.5 shows the steps involved in designing your CA infrastructure.

**Figure 16.5 Designing Your CA Infrastructure**



# Planning Core CA Options

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can establish a CA infrastructure that meets the security needs and certificate requirements for your organization, you need to make decisions about a number of core CA options that are available. Planning the CA infrastructure for your organization involves making decisions about the following:

- Location of the root certification authorities.
- Internal versus third-party CAs.
- Requirements for CA capacity, performance, and scalability.
- Your Active Directory structure.
- Your PKI management model.
- CA types and roles.
- Use of hardware cryptographic service providers.
- Number of CAs required.

# Designing Root CAs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A CA infrastructure consists of a hierarchy of CAs that trust one another and authenticate certificates belonging to one another. Within this infrastructure, a final authority, called a root CA, must be in place. The root CA certifies other certification authorities to publish and manage certificates within the organization. Before you establish a CA hierarchy, you must determine the following:

- Who designates the root certification authority in the organization. For example, determine whether this is the responsibility of central IT, divisional IT departments, or a third-party organization.
- Where the root certification authority is to be located.
- Who manages the root certification authority.
- Whether the role of the root CA is only to certify other certification authorities, or also to serve certificate requests from users.

After you have made these determinations, you can define the roles for any additional certification authorities, including who manages them and what trust relationships they have with other CAs. For more information about CA roles, see "Defining CA Roles in the Trust Hierarchy" later in this chapter.

# Selecting Internal CAs vs. Third-Party CAs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Depending on the functionality that you require, the capabilities of your IT infrastructure and IT administrators, and the costs that your organization can support, you might choose to base your certification authority infrastructure on internal CAs, third-party CAs, or a combination of internal and third-party CAs.

**Internal CAs**
If your organization conducts most of its business with partner organizations and wants to maintain control of how certificates are issued, internal CAs are the best choice. Internal CAs:

- Allow an organization to maintain direct control over its security policies.
- Allow an organization to align its certificate policy with its overall security policy.
- Can be integrated with the Active Directory infrastructure of the organization.
- Can be expanded to include additional functionality and users at relatively little extra cost.

The disadvantages associated with using internal CAs include:

- The organization must manage its own certificates.
- The deployment schedule for internal CAs might be longer than that for CAs available from third-party service providers.
- The organization must accept liability for problems with the PKI.

**External CAs**
If your organization conducts most of its business with external customers and clients and wants to outsource certificate issuing and management processes, you might choose to use third-party CAs. Third-party CAs:

- Allow customers a greater degree of confidence when conducting secure transactions with the organization.
- Allow the organization to take advantage of the expertise of a professional service provider.
- Allow the organization to use certificate-based security technology while developing an internally managed PKI.
- Allow the organization to take advantage of the provider's understanding of the technical, legal, and business issues associated with certificate use.

The disadvantages associated with use of third-party CAs include:

- They typically involve a high per-certificate cost.
- They might require the development of two different management standards, one for internally issued certificates and one for commercially issued certificates.
- They allow less flexibility in configuring and managing certificates.
- The organization must have access to the third-party CAs in order to access the CRLs.
- Autoenrollment is not possible.
- Third-party CAs allow only limited integration with the internal directories, applications, and infrastructure of the

organization.

You might need to use both internal and third-party CAs. For information about using a combination of internal and third-party CAs in your organization, see "Extending Your CA Infrastructure" later in this chapter.

# Evaluating CA Capacity, Performance, and Scalability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Organizations must agree upon a definition of acceptable CA performance. To determine the appropriate number of CAs and the best configuration for your CA infrastructure, you need to evaluate and address the factors in your organization that impact CA capacity, performance, and scalability. These include:

- The number of certificates that you need to issue and renew.
- The key lengths of the issuing CA certificates.
- The type of hardware that is used for your CAs.
- The number and configuration of the client computers that you need to support.
- The quality of your network connections.

A stand-alone Windows Server 2003 CA supports more than 35 million certificates per physical CA without any degradation of performance.

An individual departmental certification authority running on a server with a dual processor and 512 megabytes (MB) of RAM can issue more than 2 million standard-key-length certificates per day. Even with an unusually large CA key, a single stand-alone CA with the appropriate hardware is capable of issuing more than 750,000 user certificates per day.

Using a greater number of small CAs with strategically located CRL distribution points reduces the risk that your organization might be forced to revoke and reissue all its certificates if a large CA is compromised. However, using a greater number of CAs might increase your administrative overhead.

For many organizations, the primary limitations to CA performance are the amount of physical storage available and the quality of the clients' network connectivity to the CA. If too many clients attempt to access your CA over slow network connections, autoenrollment requests can be delayed.

Another significant factor is the number of roles that a CA server performs on the network. If a CA server is operating in more than one capacity in the network — for example, if it also functions as a domain controller — it can negatively impact the capacity and performance of the CA. It can also complicate the delegation of administration for the CA server. For this reason, unless your organization is extremely small, use your CAs only to issue certificates.

Some hardware components impact PKI capacity and performance more than others. When you are selecting the server hardware for your CAs, consider the following:

- **Number of CPUs.** Large CA key sizes require more CPU resources. The greater the number of CPUs, the better the performance of the CA. CPU power is the most critical resource for a Windows Server 2003 certification authority.

  **Note**

  - Because of the architecture of their databases, Windows Server 2003 certification authorities are CPU-intensive and use a substantial amount of the disk subsystem. However, other hardware resources can also impact the performance of a CA when the system is put under stress.

- **Disk performance.** In general, a high-performance disk subsystem allows for a faster rate of certificate enrollment. However, key length impacts disk performance. With a shorter CA key length, the CPU has fewer calculations to perform and, therefore, it can complete a large number of operations. With longer CA keys, the CPU needs more time to issue a certificate and this results in a smaller number of disk input/output (IO) operations per time interval.

- **Number of disks.** You can improve performance slightly by using separate physical disks for the database and log files. You can improve performance significantly by placing the database and log files on RAID or striped disk sets. In general, the drive that contains the certification authority database is used more than the drive hosting the log file.

**Note**

- Using separate logical disks does not provide any performance advantages.

- **Amount of memory.** The amount of memory that you use does not have a significant impact on CA performance, but must meet general system requirements

- **Hard disk capacity.** Certificate key length does not affect the size of an individual database record. Therefore, the size of the CA database increases linearly as more records are added. In addition, the higher the capacity of the hard disk, the greater the number of certificates that a CA can issue.

**Tip**

- Plan for your hard disk requirements to grow over time. In general, every certificate that you issue requires 17 kilobytes (KB) in the database and 15 KB in the log file.

The type of hardware that your clients use can also impact performance. When you are selecting or evaluating the capabilities of the hardware for your CA clients, consider the following:

- **Key length.** The greater the key length of a requested certificate, the greater the impact on the CPU of the server hosting the CA.

- **Network bandwidth.** Assuming that the CA is not serving in more than one capacity, a 100-megabit network connection is sufficient to prevent performance bottlenecks.

As you plan your CA infrastructure, you also need to ensure that your design is flexible enough to accommodate changes to your organization. For example, you need to be able to accommodate:

- Changes in the functionality that you require from your public key infrastructure.

- Growth or decline in demand for certificates.

- The addition or removal of locations that CAs need to serve.

- The effect of revocation. Revoking large numbers of certificates can take several minutes and increase the size of the database.

Using multiple CAs is an excellent way to ensure that your infrastructure can support enterprise scalability. The use of multiple CAs, even for organizations with minimal certificate requirements, provides the following advantages:

- **Greater reliability.** If you need to take an individual CA offline for maintenance or backup, another CA can service its requests.

- **Scalability.** Increases in demand, either from new users or from new applications, can be accommodated more easily.

- **Distributed administration.** Many organizations distribute security administration across a number of IT administrators to prevent one individual or team from controlling the entire security technology infrastructure of the organization.

- **Improved availability.** Users in remote offices can access a CA that is local to them rather than accessing a CA across slow Wide Area Network (WAN) links.

**Note**

- You can reorganize your CA infrastructure by adding or removing a CA and its associated users from a CA hierarchy. However, you cannot move a subset of users on a single CA to a new CA without forcing the users to re-enroll with the

new CA.

# Integrating the Active Directory Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your CA infrastructure is independent of the domain structure of your Windows environment. For example, one CA can service requests from multiple domains, or multiple CAs can serve a single domain. CA hierarchies with stand-alone CAs can even span multiple Active Directory forests.

If possible, take your PKI requirements into account when you design your Active Directory infrastructure. Active Directory and PKI technology impact each other in the following ways:

- **Enterprise CAs are bound to the forest.** As a result, enterprise CAs can only issue certificates to computers and users in the forest. In addition, you cannot change the name of the CA or the computer after it is deployed. Moreover, the computer cannot be removed from the domain or forest. Because much of the security of an organization is established at the forest level, the security of an enterprise CA is connected to the forest in which it is located. For this reason, each forest requires its own enterprise CAs.

  **Note**

  - If certificates from stand-alone CAs are published to Active Directory, these stand-alone CAs cannot be renamed or removed from the forest without their certificates becoming invalid. However, you can rename stand-alone CAs that belong to workgroups without impacting the status of their certificates.

- **Certificate storage affects the size of your directory.** If you store certificates in user objects, the size of the directory increases and replication time might increase. Because the userCertificate attribute contains data about all the user certificates, the addition of a certificate to that multivalued attribute causes Active Directory to replicate attribute data for all certificates.

- **Complications such as failure to recognize the user or the certificate can occur.** This happens if you do not apply a consistent naming structure for both your distinguished names (also known as DNs) and your user principal names (UPNs).

- **Enterprise CAs rely on the existence of an Active Directory schema.** If your schema is based on Windows 2000 Active Directory, you might need to extend it to support Windows Server 2003 Certificate Services functionality, such as version 2 certificate templates. For more information about version 2 templates, see "Selecting Certificate Templates" later in this chapter.

For certificates with a long life, the availability of the CA services themselves is much less important than the availability of the directory that holds the certificates and the certificate revocation lists. If you integrate your CAs with Active Directory, your certificates and CRLs are automatically published to the directory and replicated throughout the forest as part of the global catalog.

**Note**

- If you use Active Directory to publish and replicate information about CRLs throughout your organization, be sure to review Active Directory replication schedules and policies in order to ensure that this data is distributed in a timely manner.

Windows Server 2003 Certificate Services functions whether Active Directory in your organization is based on Windows 2000 or Windows Server 2003. It also functions if your organization is operating in mixed mode.

#### Configuring Public Key Group Policy

If you have an Active Directory environment, Group Policy allows you to link certificate services to groups of users or computers based on their domains or organizational unit membership. You must configure public key Group Policy in order to perform the following tasks:

- **Add trusted root certificates for groups of computers.** You can define the following:

- Which root CAs users can trust when verifying certificates.
- Whether users are allowed to trust additional CAs of their own choosing.
- The purposes for which certificates issued by each CA can be used.

Enterprise root CAs within your domain forest are automatically added to these policies.

- **Distribute certificate trust lists for computers and users.** For more information about certificate trust lists, see "Evaluating Factors That Affect Extended Trusts" later in this chapter.
- **Enable autoenrollment.** For more information about autoenrollment, see "Selecting a Certificate Enrollment and Renewal Method" later in this chapter.
- **Designate EFS recovery agent accounts.** You can define an EFS recovery policy within the scope of the policy object. If a recovery policy is defined, it is populated with the certificates of the recovery agents.

In many organizations, users and computers are already organized into domains and organizational units that are based on the organization structure, location, and job function. If your organization has not already created an Active Directory domain structure, the best way for you to take advantage of Public Key Group Policy is to define the groups of users and computers that will use your Certificate Services and communicate this information to the Active Directory and Group Policy administrators, so that they can address your public key requirements in their planning.

For more information about how to plan for the use of Group Policy, see "Designing a Resource Authorization Strategy" in this book.

# Defining PKI Management and Delegation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to define a PKI management model early in the process of designing your CA infrastructure. This PKI management model must complement your existing security management delegation plan and help you to meet Common Criteria requirements for role separation. To ensure that a single individual cannot compromise PKI services, it is best to distribute management roles across different individuals in your organization. This involves deciding which individuals are to perform each of the following tasks:

- Creating or modifying existing CAs
- Managing certificate templates
- Issuing cross certificates
- Issuing or revoking user certificates
- Configuring and viewing audit logs

You can use discretionary access control lists (DACLs) to manage CA permissions and delegate CA management tasks.

Windows Server 2003 includes the following CA management roles:

- **Service Manager.** Configures and manages Certificate Services for local users, assigns certificate managers, and renews CA certificates.
- **Certificate Manager.** Issues and revokes certificates.
- **Auditor.** Audits the actions of local administrators, service managers, and certificate managers.

The extent to which you separate roles depends on the level of security that you require for a particular service. Assign the fewest possible rights to users in order to achieve the greatest level of security. For example, you can adopt the following rules:

- No user can assume the roles of both CA Administrator and Certificate Manager.
- No user can assume the roles of both User Manager and Certificate Manager.

If you need stricter guidelines, you can include the following:

- No user can assume the roles of both Auditor and Certificate Manager.

To facilitate this delegation process, you need to understand how various PKI administrative roles align with Windows Server 2003 administrative roles. Table 16.1 lists the Windows Server 2003 administrative roles that correspond to each PKI administrative role.

**Table 16.1 PKI Administrative Roles and Their Corresponding Windows Server 2003 Administrative Roles**

| PKI Administrative Role | Description | Windows Server 2003 Administrative Role |
|---|---|---|
| PKI Administrator | Configures, maintains, and renews the CA. | User |
| Backup Operator | Performs system backup and recovery. | Backup Operator on the server on which the CA is running |

| | | |
|---|---|---|
| Audit Manager | Configures, views, and maintains audit logs. | Local Administrator on the server on which the CA is running |
| Key Recovery Manager | Requests retrieval of a private key stored by the service. | User |
| Certificate Manager | Approves certificate enrollment and revocation requests. | User |
| User Manager | Manages users and their associated information. | Account Operators (or person delegated to create user accounts in Active Directory) |
| Enrollee | Requests certificates form the CA | Authenticated Users |

Table 16.2 lists the actions that each PKI administrative role can perform.

**Table 16.2 Actions Performed By PKI Administrative Roles**

| **Action** | **Enrollee** | **CA Admin** | **Certificate Manager** | **Audit Manager** | **Backup Operator** | **Local Server Admin** |
|---|---|---|---|---|---|---|
| Install a CA | | | | | | • |
| Configure a CA | | • | | | | • |
| Policy and exit module configuration | | • | | | | |
| Stop/start service | | • | | | | • |
| Change configuration | | • | | | | |
| Assign user roles | | • | | | | |
| Establish user accounts | | • | | | | • |
| Maintain user accounts | | • | | | | • |
| Configure profiles | | • | | | | • |
| Renew CA keys | | | | | | • |
| Define key recovery agent(s) | | • | | | | |
| Define officer roles | | • | | | | |
| Enable role separation | | • | | | | |
| Issue/Approve certificates | | | • | | | |
| Deny certificates | | | • | | | |
| Revoke certificates | | | • | | | |
| Unrevoke certificates | | | • | | | |

| | | | | | | |
|---|---|---|---|---|---|---|
| Renew certificates | | | ● | | | |
| Enable, publish, or configure CRL schedule | | ● | | | | |
| Configure audit parameters | | | | ● | | ● |
| Audit logs | | | | ● | | ● |
| Back up system | | | | | ● | ● |
| Restore system | | | | | ● | ● |
| Read CA properties, CRL | ● | | | | | |
| Request certificate | ● | | | | | |
| Read CA database | | ● | ● | ● | ● | |
| Read CA configuration information | | ● | ● | ● | ● | |
| Read issued, Revoked, pending certificates | | ● | ● | ● | ● | |

**Note**

- As you delegate roles and responsibilities, be sure to keep track of the permissions that you configure on certificate directories. Distributing access to a PKI to a number of individuals creates greater security risks.

# Defining CA Types and Roles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To plan your CA infrastructure, you need to understand the different types of CAs available with Windows Server 2003 and the roles that they can play. Windows Server 2003 Certificate Services supports the following two types of CAs:

- Enterprise
- Stand-alone

Enterprise and stand-alone CAs can be configured as either Root CAs or Subordinate CAs. Subordinate CAs can further be configured as either Intermediate CAs (also referred to as a policy CA) or Issuing CAs.

Before you create your CA infrastructure, you need to determine the type or types of CAs that you plan to use, and define the specialized roles that you plan to have each CA assume.

**Enterprise vs. Stand-Alone CAs**

*Enterprise CAs* are integrated with Active Directory. They publish certificates and CRLs to Active Directory. Enterprise CAs use information stored in Active Directory, including user accounts and security groups, to approve or deny certificate requests. Enterprise CAs use certificate templates. When a certificate is issued, the enterprise CA uses information in the certificate template to generate a certificate with the appropriate attributes for that certificate type.

If you want to enable automated certificate approval and automatic user certificate enrollment, use enterprise CAs to issue certificates. These features are only available when the CA infrastructure is integrated with Active Directory. Additionally, only enterprise CAs can issue certificates that enable smart card logon, because this process requires that smart card certificates be mapped automatically to the user accounts in Active Directory.

*Stand-alone CAs* do not require Active Directory and do not use certificate templates. If you use stand-alone CAs, all information about the requested certificate type must be included in the certificate request. By default, all certificate requests submitted to stand-alone CAs are held in a pending queue until a CA administrator approves them. You can configure stand-alone CAs to issue certificates automatically upon request, but this is less secure and is usually not recommended, because the requests are not authenticated.

From a performance perspective, using stand-alone CAs with automatic issuance enables you to issue certificates at a faster rate than you can by using enterprise CAs. However, unless you are using autoissuance, using stand-alone CAs to issue large volumes of certificates usually comes at a high administrative cost because an administrator must manually review and then approve or deny each certificate request. For this reason, stand-alone CAs are best used with public key security applications on extranets and the Internet, when users do not have Windows 2000 or Windows Server 2003 accounts, and when the volume of certificates to be issued and managed is relatively low.

You must use stand-alone CAs to issue certificates when you are using a third-party directory service or when Active Directory is not available.

**Note**

- You can use both enterprise and stand-alone certification authorities in your organization.

Table 16.3 lists the options that each type of CA supports.

**Table 16.3 Options for Enterprise vs. Stand-Alone CAs**

| Option | Enterprise CA | Stand-alone CA |
|---|---|---|
| Publish certificates in Active Directory and use Active Directory to validate certificate requests. | • | |
| Take the CA offline. | | • |
| Configure the CA to issue certificates automatically. | • | |

| | | |
|---|---|---|
| Allow administrators to approve certificate requests manually. | | • |
| Use certificate templates. | • | |
| Authenticate requests to Active Directory. | • | |

**Root CAs**

A *root CA* is the CA that is at the top of a certification hierarchy and must be trusted unconditionally by clients in your organization. All certificate chains terminate at a root CA. Whether you use enterprise or stand-alone CAs, you need to designate a root CA.

Because there is no higher certifying authority in the certification hierarchy, the subject of the certificate issued by a root CA is also the issuer of the certificate. Likewise, because the certificate chain terminates when it reaches a self-signed CA, all self-signed CAs are root CAs. Windows Server 2003 only allows you to designate a self-signed CA as a root CA. The decision to designate a CA as a trusted root CA can be made at either the enterprise level or locally, by the individual IT administrator.

A root CA serves as the foundation upon which you base your certification authority trust model. It guarantees that the subject public key belongs to the subject identity information that is contained in the certificates it issues. Different CAs might also verify this relationship by using different standards; therefore it is important to understand the policies and procedures of the root certification authority before choosing to trust that authority to verify public keys.

The root CA is the most important CA in your hierarchy. If your root CA is compromised, every other CA and certificate in your hierarchy might have been compromised. You can maximize the security of the root CA by keeping it disconnected from the network and using subordinate CAs to issue certificates to other subordinate CAs or to end users.

For more information about using a third-party CA as the root CA, see "Extending Your CA Infrastructure" later in this chapter. For more information about disconnecting CAs from the network, see "Using Offline CAs" later in this chapter.

**Subordinate CAs**

CAs that are not root CAs are considered subordinate. The first subordinate CA in a hierarchy obtains its CA certificate from the root CA. This first subordinate CA can, in turn, use this key to issue certificates that verify the integrity of another subordinate CA. These higher subordinate CAs are referred to as intermediate CAs. An intermediate CA is subordinate to a root CA, but also serves as a higher certifying authority to one or more subordinate CAs.

An intermediate CA is often referred to as a policy CA because it is typically used to separate classes of certificates that can be distinguished by policy. For example, policy separation includes the level of assurance that a CA provides or the geographical location of the CA to distinguish different end-entity populations. A policy CA can be online or offline.

### Note

- Most organizations use one root CA and two policy CAs — one to support internal users, the second to support external users.

The next level in the CA hierarchy usually contains the issuing CA. The issuing CA issues certificates to users and computers and is almost always online. In many CA hierarchies, the lowest level of subordinate CAs is replaced by RAs, which can act as an intermediary for a CA by authenticating the identity of a user who is applying for a certificate, initiating revocation requests, and assisting in key recovery. Unlike a CA, however, an RA does not issue certificates or CRLs; it merely processes transactions on behalf of the CA.

# Using Offline CAs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Securing your CA hierarchy is a critical task. If an intruder can gain access to a CA, either physically or by means of the network, he or she might retrieve the private key of the CA and then impersonate the CA to gain access to valuable network resources. The compromise of even one CA key invalidates the security protection that it and any CAs below it in the hierarchy provide. For this reason, it is important to avoid connecting root CAs to the network.

To ensure the reliability of your CA infrastructure, specify that any non-issuing root and intermediate CAs must be offline. This minimizes the risk of the CA private keys becoming compromised. You can take a CA offline in any of the following ways:

- By installing a CA on a stand-alone Windows 2000 or Windows Server 2003 and configuring it as a stand-alone CA.
- By physically removing the computer from the network.
- By shutting down the CA service.
- By shutting down the computer.

**Caution**

- Shutting down a CA computer prevents auditing from taking place. Therefore, if a CA computer is compromised, a hardware failure does not generate an audit notification.

Make sure that you keep CAs in a secure area with limited access.

Installing an offline CA on a server that is a member of a domain can cause problems with a secure channel when you bring the CA back online after a long offline period. This is because the computer account password changes every 30 days. You can get around this by making offline CA computers members of a workgroup. Installing an offline CA as an enterprise CA can cause Active Directory to have problems updating when you disconnect the server from the network. Therefore, do not use an enterprise CA as a root CA.

Because they can operate offline, it is a good idea to use stand-alone CAs for root and intermediate CAs.

When a CA is supposed to be an offline CA, you can still publish its certificate and CRL in Active Directory. You must be sure to bring an offline CA online at regular intervals, based on your CRL publication schedule, to generate a new CRL for the CA. You must also bring the CA online to process certificate requests for subordinate CA certificates.

**Important**

- In general, the CRL and Authority Information Access (AIA) paths of an offline CA have to be modified before the first certificate is issued because the CRL and AIA paths, by default, point to the local http server and the local file system. Because the CA is offline and not accessible to other members of a network, the functionality of the CA must be separated from CRL and AIA distribution.

Because offline CAs process a small number of certificate requests at infrequent intervals, the administrative costs of maintaining offline CAs is low.

The client-side certificate validation process is not affected when a CA is offline because the client verifies the validity of the certificate by checking the certificate chain and the CRL. You cannot store both sources on the offline CA because clients need access to the CRL and AIA paths that are part of the certificate.

**Important**

- Taking a root CA offline does not reduce its importance, so be sure to use reliable hardware for offline root CAs. A hardware failure on an offline CA prevents you from publishing CRLs or issuing certificates to new subordinate CAs.

# Using Hardware CSPs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Hardware CSPs can support a wide range of cryptographic operations and technologies. Keys stored in tamper-resistant hardware crypto-devices are more secure than keys stored on local computer hard disks. Therefore, keys stored in hardware cryptographic devices can have key lifetimes that are longer than keys stored by software CSPs on hard disks.

**Note**

- Another advantage to using hardware CSPs is that the key material is kept outside the memory of the computer and within the hardware device. This makes it impossible to access the key of the CA by means of a memory dump.

If you determine that a hardware CSP is too costly, consider using smart cards for key storage. When you store cryptographic keys on a smart card, no one in your organization can issue or revoke certificates without the appropriate smart card together with the correct personal identification number (PIN).

If you choose to use hardware cryptographic service providers for CA private key storage, you must ensure that the hardware device is physically secured, or at least back up the operator cards or tokens. You might, for example, keep it in a highly secured area in the computer room of your company, or lock it in a safe.

# Determining Number of CAs Required

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have identified your application and user requirements, you can begin to estimate the number of CAs that you need to deploy. If your organization has limited certificate requirements, a small user base, and limited expansion goals, a single CA might be sufficient. By using a single CA, you can still meet a variety of needs by customizing and deploying certificate templates and using role separation. However, if availability or distributed functionality of Certificate Services is a priority, you must deploy multiple CAs. You also need multiple CAs if you want separate CAs to issue certificates for different purposes.

To determine the number of CAs required, answer the following questions:

- Do you require more than one CA? If you are only supporting a single application and location, and if 100 percent availability of the CA is not critical, you might be able to use a single CA. Otherwise, you probably require at least one root and multiple subordinate CAs.

- If you need more than one CA, how many root CAs do you require? Generally, it is recommended that you have only one root CA as a single point of trust. This is because significant cost and effort is required to protect a root CA from compromise. With multiple root CAs, root maintenance becomes much more difficult.

  However, organizations with a decentralized security administration model, such as corporations with multiple, largely independent business units and no strong central administrative body, might require more than one root CA. For more information about using more than one root CA, see "Extending Your CA Infrastructure" later in this chapter.

- How many intermediate or policy CAs do you need?

- How many issuing CAs or RAs do you need?

  The number of intermediate and issuing CAs that you deploy depends on the following factors:

  - **Usage.** Certificates can be issued for a number of purposes (for example, secure e-mail, network authentication, and so on). Each of these uses might involve different issuing policies. Using separate CAs provides a basis for administering each policy separately.

  - **Organizational or geographic divisions.** You must have different policies for issuing certificates, depending on the role of an entity or its physical location in the organization. You can create separate subordinate CAs to administer these policies.

  - **Distribution of the certificate load.** You can deploy multiple issuing CAs to distribute the certificate load to meet site, network, and server requirements. For example, if network links between sites are slow or discontinuous, you might need to place issuing CAs at each site to meet Certificate Services performance and usability requirements.

  - **The need for flexible configuration.** You can tailor the CA environment (key strength, physical protection, protection against network attacks, and so on) to provide a balance between security and usability. For example, you can renew keys and certificates more frequently for the intermediate and issuing CAs that are at high risk for compromise, without requiring a change to established root trust relationships. Also, when you use more than one subordinate CA, you can turn off a subsection of the CA hierarchy without affecting established root trust relationships or the rest of the hierarchy.

  - **The need for redundant services.** If one enterprise CA fails, redundancy makes it possible for another issuing CA to provide users with uninterrupted service.

Strive to have only as many CAs and RAs as you need to function efficiently. Deploying more CAs than you need creates an unnecessary management burden, and introduces additional areas of security vulnerability.

**Note**

- You cannot install more than one CA on a server.

# Selecting a Trust Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Server 2003 PKI is based on a hierarchical CA model that is composed of well-defined trust and CA naming standards. This type of CA trust model provides scalability, easy administration, and consistency with a growing number of third-party CA products.

In a hierarchical CA model, multiple CAs are organized into clearly defined parent-child relationships. Child CAs are certified by parent CA-issued certificates, which bind the public key of a CA to its identity.

With a hierarchical CA model, you minimize the number of root CAs that you need in order to verify certificates. At the same time, hierarchical CAs allow you great flexibility in the number of certificate-issuing subordinate CAs that you can use.

The basic types of CA trust hierarchies include:

- **Rooted trust model.** In a rooted trust model, a CA is either a root or a subordinate, and you can use offline root CAs for the highest level of security.
- **Network (or cross-certification) trust model.** In a network trust model, every CA is both a root and a subordinate.
- **Hybrid trust model.** Hybrid trust models combine elements of both the rooted and network trust models.

Your PKI trust hierarchy must be based on one of these three trust models.

Whether you choose to apply a rooted, network, or hybrid trust model to your CA infrastructure, you need to base your trust structure on the business requirements of your organization and on the way your organization delegates responsibility for IT administration. In this way, your trust model might be based on one or a combination of the following:

- Quality of identification
- Organizational structure
- User location

# Rooted Trust Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a rooted trust model, the root CA is the trust anchor and has a self-signed certificate. The root CA issues a certificate to all direct subordinate CAs, if needed, which, in turn issue certificates to their subordinate CAs. A subordinate CA is trusted cryptographically, based on the signature of its parent.

Figure 16.6 illustrates the rooted trust model.

**Figure 16.6 Rooted Trust Model**




Numerous products and services offered by major software vendors, including Microsoft, support rooted trust hierarchies. You can add a new CA to a rooted trust hierarchy by enrolling it to a CA anywhere in the trust hierarchy. If you create a new trust hierarchy, it only needs to trust the root CA of the new PKI in order to trust all the subordinate CAs in the new hierarchy.

A rooted trust model enables you to compartmentalize risks, management, and certificate processing. Rooted trust hierarchies are more scalable and easier to administer than other hierarchies because each CA serves a single role within the hierarchy and is not operationally dependent on other CAs.

Any CA in a rooted trust hierarchy is either a root or a subordinate but never both. Each CA is responsible for processing requests and issuing certificates signed by its own key; each CA is responsible for revoking certificates and publishing CRLs to accessible locations; and each CA can be managed separately by different personnel in different parts of an organization.

Because CAs in a rooted trust hierarchy can be online or offline, rooted trust hierarchies allow great flexibility in the ways in which you can deploy and manage a PKI. You can protect the private key of a CA by taking the CA offline. Because offline CAs are typically the root and/or policy CAs that only issue certificates to other CAs, taking the CA offline does not impact other parts of the hierarchy.

Because most protocols deliver a chain of certificates that terminates in a trusted root CA, rooted trust hierarchies provide a straightforward means by which CAs can determine whether a certificate can be trusted.

**Note**

- If the certificate of a root CA expires, all certificates that are issued by the root CA or by its subordinate CAs also expire. For more information about managing certificate lifetimes, see "Selecting Certificate Security Options" later in this chapter.

# Network Trust Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization has multiple, distributed IT departments, you might not be able to establish a single, trusted root. In this situation, you can implement a network trust model, in which all CAs are self-signed and trust relationships between CAs are based on cross-certificates. Cross-certificates are special certificates that are used to establish complete or qualified one-way trusts between otherwise unrelated CAs. For more information about the use of cross-certificates and how to manage cross-certified relationships, see "Selecting an Extended CA Infrastructure Configuration" later in this chapter.

A network trust model can be viewed as a hierarchy because a cross-certificate is essentially the same as a subordinate CA certificate in a rooted trust model. The cross-certifying CA is the issuer and the cross-certified CA is the subject.

Because a cross-certificate is a logical subordination of one CA to another CA, a network trust model is in effect a hierarchy, with the added property that a root CA is also a subordinate CA in the cross-certifying PKI.

Unlike the rooted trust model, in which a global directory such as Active Directory is not required, a global directory is essential in a network trust hierarchy. Without a global directory, cross-certificates need to be preinstalled on all clients of the PKI; otherwise there is no way to discover them.

Figure 16.7 shows an example of a network trust model.

**Figure 16.7 Network Trust Model**




The trusts in Figure 16.7 are bidirectional, which means that CA1 issued a cross-certificate of trust to CA2 and CA2 issued a cross-certificate of trust to CA1. It is also possible to rescind trust for a CA by revoking its cross-certificate.

Cross-certification does not need to be bidirectional, and a cross-certifying CA does not need the cooperation of the CA being certified. For example, CA1 can cross certify CA2, without CA2 cross certifying CA1. In such a case, clients of CA1 trust CA2 and CA3, while clients of CA2 and CA3 do not trust CA1. To do this, CA1 creates a cross-certificate without the knowledge of CA2, because all that CA1 needs is the public key certificate of CA2. This is known as *unilateral cross-certification*, where one CA cross-certifies another CA but not the reverse.

Bidirectional cross-certificates are usually preferred, although with this model you need to manage a greater number of cross-relationships as the number of cross-certificates increases.

Full trust between cross-certified CAs also means that the client trusts all certificates issued by the other CA, regardless of the purpose of the certificate. In a native Windows Server 2003 environment, however, you can filter by certificate types. You can also limit trust between CAs by means of qualified subordination, which can be implemented in the form of name constraints, policy constraints, policy mapping, and path constraints. For more information about these methods, see "Extending Your CA Infrastructure" later in this chapter.

Cross-certification enables you to create bridges between separate PKIs without either PKI being directly subordinate to the other. Because cross-certification is an indirect subordination of one PKI to another, the trust point does not change relative to either PKI. In fact, bidirectional cross-certification models the way in which companies form relationships; that is, each side participates in establishing the relationship. A network trust model, however, is much more difficult to maintain and troubleshoot than a rooted trust model.

**Note**

- Use a network trust model only in conjunction with name constraints. For more information about name constraints, see "Using Name Constraints" later in this chapter.

# Hybrid Trust Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some organizations might find a pure rooted trust model too restrictive, because no single CA can serve as the root for all other CAs. At the same time, a pure network model can become prohibitively complex if too many different CAs are involved. If you use a hybrid approach, however, you can cross-certify only certain CAs and thus use the benefits of both the rooted and network trust models.

# Trust Hierarchy Based on Quality of Identification

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A trust hierarchy based on quality of identification enables an organization to configure CAs to issue certificates to specific groups of users. This type of trust hierarchy is ideal for organizations in which different identification and authentication requirements are applied to different groups of users, computers, and activities.

For example, an organization requires employees to appear in person to provide identification such as a driver's license or passport to a security officer, who checks an employee database to ensure that the individual is authorized, before they can receive appropriate credentials. However, because computers cannot assert an identity, managers in the organization are responsible for ensuring that computer names are correct and that computers are authorized to have a certificate. Because the organization requires CAs for employee certificates and computer certificates and each requires a different form of identification, the organization chooses to create a trust hierarchy based on quality of identification.

In a trust hierarchy based on quality of identification, the CAs subordinate to the root CA are organized according to the quality of identification required for the certificate to be issued. The subordinate CAs use certificates signed by the root CA in order to issue certificates to users, computers, services, or another CA.

A typical CA hierarchy based on quality of identification includes two or three issuing CAs for each of the following:

- Employee certificates
- Computer certificates
- Contractor certificates, if applicable. These certificates might require the same identification that employee certificates require, but contain an issuer statement stating that the individual is not a full-time employee.

# Trust Hierarchy Based on Organizational Structure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although a two- or three- tier trust hierarchy based on the quality of identification is sufficient for most organizations, some organizations might need to deploy a three-tier CA trust hierarchy based on the administrative structure of the organization.

In a trust hierarchy based on organizational structure, issuing CAs are configured to support different organizational divisions, such as permanent employees and contractors. The issuing policy, for example, might be based on the organization of user accounts, so that stronger security measures are applied to independent contractors, temporary employees, or external business partners.

Figure 16.8 shows a rooted trust hierarchy based on organizational structure.

**Figure 16.8 Rooted Trust Hierarchy Based on Organizational Structure**




Design your trust hierarchy according to organizational structure if your certificate requirements vary according to organizational units; for example, all employees receive certain certificates, all partners receive a different set of certificates, and so on. Do not use this type of design if you can define too many different groups of requirements; in this case, a trust hierarchy based on certificate usage is more appropriate.

# Trust Hierarchy Based on Location

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some organizations might find it necessary to implement a three-tier trust hierarchy based on location. This configuration allows regional administrators to manage the certificate requirements for users in a defined area such as a continent, country/region, or locale. Figure 16.9 shows a CA trust hierarchy based on location.

**Figure 16.9 Trust Hierarchy Based on Location**




Depending on the nature of your business, you might need to issue certificates based on location to comply with legal requirements — for example, if you perform work for a government agency — or other local regulations.

# Defining CA Roles in the Trust Hierarchy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have designed the trust hierarchy for your organization, you must define the roles for your root, policy, and issuing CAs.

The root CA, for example, might be used to sign, certify, and/or revoke subordinate CAs. Intermediate or policy CAs might serve internal or external customers, or, in larger organizations, might serve more specialized functions or locations. Issuing CAs and RAs might be defined according to the clients that they serve or the certificates that they issue.

You might choose to select some or all of the following roles for your intermediate and issuing CAs:

- **Intermediate CA.** Certifies subordinate CAs to issue certificates.
- **Rudimentary CA.** Issues certificates for the most basic operations, such as user authentication without an identity check.

  **Note**

  - Stand-alone CAs are primarily used in intermediate and rudimentary roles.

- **Basic security CA.** Issues certificates, based on an Active Directory identity check, to users and computers that do not have special security requirements.
- **Medium security CA.** Issues certificates to users and computers that meet special security requirements and whose identities are validated in Active Directory.
- **High security CA**. Issues certificates to users or computers that meet especially high security requirements, and whose identities must be verified by means of the examination of physical credentials.

  **Note**

  - Enterprise CAs are primarily used for basic, medium, and high security roles.

Keep the following considerations in mind as you define CA roles:

- Use a three-tier hierarchy with policy CAs only if necessary.
- Third-party CAs can form all or part of a Windows Server 2003 CA trust hierarchy.
- Some third-party products might require other CA trust models that might not be interoperable with rooted CA hierarchies. Windows Server 2003 and most commercial CAs support rooted CA hierarchies.

# Establishing a CA Naming Convention

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you configure CAs in your organization, you must establish a CA naming convention. Names for CAs cannot be more than 64 characters in length. You can create a name using any Unicode character, but you might want to use the ANSI character set if interoperability is a concern. The CA name does not have to be identical to the name of the computer.

The name that you specify when you configure a server to be a CA becomes, in Active Directory, the common name of the CA, and is reflected in every certificate that the CA issues. For this reason, it is important that you do not use the fully qualified domain name (FQDN) for the common name of the CA. This way, malicious users who obtain a copy of a certificate cannot identify and use the fully qualified domain name of the CA to create a potential security vulnerability.

**Note**

- You cannot change the name of a server after Certificate Services has been installed without invalidating all the certificates issued by the CA. To change the server name after Certificate Services has been installed, you must uninstall the CA, change the name of the server, reinstall the CA, and reissue all the certificates issued by the CA. You do not have to reinstall a CA if you rename a domain; however, you will have to reconfigure the CA to support the name change.

# Selecting a CA Database Location

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you install a CA in your organization, you must specify a location for the database and log files of the CA. You must also indicate whether you want to store the configuration information for the CA. Storing the CA configuration information is helpful for backing up and, if necessary, restoring your CA.

You can choose to copy the naming information and the certificate for the CA to the file system (the configuration directory is automatically shared by means of a share named certconfig).

**Note**

- You can change the location of the database and log files manually at a later time. However, you cannot perform this task by using the user interface.

Windows Server 2003 uses the JET database engine for the CA database. As with any JET database, it is a good idea to place the database and its log files on different physical disk drives, in order to improve fault tolerance and performance. By default, all these files are located in the certlog subdirectory of the system directory.

**Tip**

- Use a separate RAID for both the database and log files for the highest level of fault tolerance between backup intervals.

The CA database consists of the files listed in Table 16.4.

**Table 16.4 CA Database Files**

| Database file | Purpose |
|---|---|
| <CA name>.edb | The CA store |
| edb.log | The transaction log file for the CA store |
| res1.log | Reservation log file to store transactions if disk space is exhausted |
| res2.log | Reservation log file to store transactions if disk space is exhausted |
| edb.chk | Database checkpoint file |

**Note**

- You can determine the location of the database files for a CA by typing **`certutil -databaselocations`** at a command prompt or by looking in the Certificate Services snap-in user interface.

# Example: Designing a CA Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After an organization defines its certificate requirements, it creates a linked hierarchy of certification authorities to enable it to distribute certificates as needed, and to validate or reject certificates as appropriate.

In creating this CA infrastructure, the organization takes the following elements into account:

- **The security administration model of the organization.** For example, security administration is managed centrally from the headquarters of the organization, but individual business units create and support their own security requirements as needed for individual projects and business relationships. Some units operate autonomously, but report back to corporate IT.

- **The Active Directory infrastructure of the organization.** Because the organization has a single-forest logical structure, the CA infrastructure design is simple. The existing single-forest structure allows them to set up CAs, based on geography and bandwidth, to serve clients in multiple domains. For example, one or more common CAs support clients in offices on opposite coasts.

- **Potential use of a third-party CA.** The organization is concerned about IT costs and also prefers to manage its own security infrastructure. It addresses both concerns by creating and administering its own CA infrastructure. When joint venture business partners deploy PKIs, it is possible to integrate the two CA infrastructures without having to rely on a third-party CA.

  For more information about using third-party CAs to extend the CA infrastructure, see "Extending Your CA Infrastructure" later in this chapter.

Although the organization deploys Active Directory, it places a stand-alone root CA in a workgroup, rather than in the domain, for increased security. Also, it keeps this root CA offline and in a secure location that can only be accessed by an administrator who is authenticated by means of a smart card.

Directly below the root CA, the organization adds three policy CAs. One CA signs all certificates that have been issued to meet the high security standards of the organization, including software code signing, smart card logon, and Internet authentication certificates. The second CA signs all certificates that have been issued to meet the medium security standards of the organization, such as e-mail and EFS certificates. The third signs certificates for the CAs that issue certificates to external partners. These are also offline.

Figure 16.10 shows the CA infrastructure for the organization.

**Figure 16.10 Example of a CA Infrastructure of an Organization**



Table 16.5 summarizes the configuration of these CAs.

**Table 16.5 CA Configuration**

| CA | Name | State | Role | Domain |
|---|---|---|---|---|
| Root CA | RtCA01 | Offline | Stand-alone | None |
| Internal medium security policy | PolCA01 | Offline | Stand-alone | None |
| Internal high security policy | PolCA02 | Offline | Stand-alone | None |
| External high security policy | PolCA03 | Offline | Stand-alone | None |
| Internal medium security issuing 1 | IsCA01 | Online | Member server | Corp |
| Internal medium security issuing 2 | CA06 | Online | Member server | Corp |
| Internal medium security issuing 3 | CA07 | Online | Member server | Corp |
| Internal high security issuing 2 | CA08 | Online | Member server | Corp |
| Internal high security issuing 3 | CA09 | Online | Member server | Corp |
| External high security issuing 1 | CA01 | Online | Member server | Corp |

# Extending Your CA Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use a rooted CA hierarchy to enable many PKI applications. However, you might find that your PKI needs are too complex to be met by a simple rooted hierarchy. For example, you might need to extend your CA infrastructure to accommodate joint ventures, mergers, geographic, or other business requirements.

Figure 16.11 shows the steps involved in extending your CA infrastructure.

**Figure 16.11 Extending Your CA Infrastructure**



# Evaluating Factors That Affect Extended Trusts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Extending and refining your trust hierarchy is a critical step in the process of creating a secure PKI, and it involves complex decisions. It is important to define appropriate and inappropriate uses for certificates in your organization before you extend your CA infrastructure. Without proper planning, you might grant business partners and users more trust than you intend.

If you want to link your established Windows Server 2003 trust hierarchy with an external PKI, a number of factors can impact interoperability. Before you extend your CA infrastructure, evaluate the following features and standards in both PKIs:

- Standards support
- Algorithm support
- CRL distribution points
- Authority information access (AIA)
- Authority key identifier (AKI)
- Certificate extensions
- Key length
- Extended key usage (EKU)
- Directory integration

Determine whether any other PKIs with which you need to establish trust support these features and standards, and how you can accommodate differences. Addressing these issues when you design your PKI can help you to ensure the extensibility and interoperability of your PKI environment.

**Standards Support**
A number of technical standards provide a basis for interoperability between Windows Server 2003 and other PKI applications. To promote third-party interoperability with the Windows Server 2003 API, Microsoft supports the following standards:

- *PKIX*. Defines interoperable PKI standards for the Internet.
- *X.509*. Describes the standard format of a certificate.
- *PKCS*. Provides a standard for public key message exchanges.
- *TLS*. Provides a secure and authenticated channel between hosts on the Internet above the transport layer.
- *S/MIME*. Serves as a standard for secure e-mail across the Internet.
- *Kerberos authentication protocol*. Provides a symmetric key framework for authentication in large networks.
- *PC/SC*. Serves as a standard for integrating smart cards and smart card readers.

Most PKI vendors have adopted many or all of these PKI standards. Different vendors, however, can implement the standards in different ways. While it might be possible to link external PKI implementations to yours, this might involve making some changes to your existing design. For this reason, it is strongly recommended that you evaluate the external PKI to determine whether it meets all your critical requirements.

**Algorithm Support**
It is important for a PKI to interoperate with many different hardware vendors and to provide a hardware abstraction layer so that applications do not have to know where keys are stored.

Windows Server 2003 uses CryptoAPI to abstract hardware-based key management from applications, and it uses the PC/SC standard instead of PKCS#11 to communicate with smart cards and readers. Many third-party CAs have their own cryptographic APIs and use PKCS#11 to interface to hardware tokens such as smart cards. Because Windows 2000 and Windows Server 2003 require hardware devices to support Plug and Play and power management features, and PC/SC includes support for these ease-of-use features, Windows Server 2003 does not support PKCS#11.

### Note

- The Windows Server 2003 PKI can use third-party CSPs, and can enroll users for certificates that have keys that were generated by third-party CSPs.

**CRL Distribution Points**
The CRL distribution point (CDP) extension in a certificate identifies how revocation information for the certificate can be obtained. If a CRL distribution point is not always available, certificate chain building can be delayed, causing inconvenience for the user. If a CRL is not available at the distribution point that has been specified in the certificate, CRL retrieval might even fail and the certificate will be considered invalid.

### Note

- Publish CDP URLs for all CAs so that users who need to know whether or not issued certificates have been revoked can find that information.

You need to compare any third-party CRL support with the Windows Server 2003 CRL support. For example, the third-party PKIs might not support the Windows Server 2003 CRL process, which includes the use of delta CRLs. Conversely, the Windows Server 2003 PKI might not support the methods of the third-party PKI for processing CRL information. Your extended PKI deployment plan needs to account for these differences.

In general, follow these guidelines when you configure the CDP extension:

- If available, use Active Directory to support internal corporate clients.
- Use an externally referenced and trusted Lightweight Directory Access Protocol (LDAP) directory to support business partners and customers.
- Consider using HTTP distribution points, especially for certificates that will be used externally.

**Authority Information Access**
The Authority Information Access (AIA) extension is a pointer to the most currently published CA certificate of a CA. The AIA extension helps clients find CA certificates dynamically during chain building. The Windows Server 2003 PKI uses this extension to assist in building trust chains to validate certificates.

It is recommended that you publish AIA URLs for all PKIs for which users might need to retrieve up-to-date CA certificates. Whether a CA is online or offline, and whether it is a root, intermediate, or issuing CA, using the AIA extension minimizes the likelihood that PKI clients will encounter unverified certificate chains or revocation data. Such encounters can result in unsuccessful VPN connections, failed smart card logons, or unverified e-mail signatures.

Some third-party PKIs do not provide the AIA extension. In this case, parent certificates need to be distributed to domain clients so that the certificates are available before the chain building process begins. Cross-certificates must also be available locally on domain clients, because there is no information in a certificate specifying where it can be found.

**Authority Key Identifier**
The Authority Key Identifier (AKI) extension provides a means to identify the public key of the CA that validates the signature on a CRL. This identification is based on either the subject key identifier (SKI) or the issuer name and serial number from the certificate issued by the CRL issuer. The AKI extension is useful in cases when a CRL issuer has more than one signing key.

An organization that expects its PKI certificates to be used by other Windows Server 2003 PKIs must populate the Authority Key Identifier extension with a unique key identifier and an issuer name and serial number. The Windows Server 2003 PKI attempts to construct certificate chains by using the issuer name and serial number in the AKI first, and then the subject key

identifier.

### Note

- By default, Windows Server 2003 does not automatically add issuer names and serial numbers to the AKI extension. This data must be added manually by means of Certutil.exe, although in most cases it is not necessary to do so.

**Certificate Extensions**
Not all certificate extensions are universally recognized. If a CA does not recognize a certificate extension in a request and it has been marked critical, it rejects the certificate. Unless you intend to limit the use of the certificate to a specific application that understands the critical extension, avoid putting critical extensions in certificates because it limits interoperability.

**Key Length**
When different PKIs support different minimum and maximum key lengths, an interoperability problem results. Be sure that your internal PKI and the external PKI support the necessary encryption keys.

**Extended Key Usage**
The Extended Key Usage (EKU) extension indicates the purposes for which the public key contained in the certificate can be used. The Windows Server 2003 PKI uses the EKU extension to indicate certificates that support special functions, such as IPSec and EFS file encryption backup. The EKU extensions of other organizations might be used for different purposes.

**Directory Integration**
Windows Server 2003 PKI certificates can be published to any directory or repository, although the default CA exit module only supports Active Directory. However, by default, a Windows Server 2003 PKI relies on Active Directory and the LDAP for authentication, including smart card logons and certificate autoenrollment, as well as for certificate management.

With Microsoft Certificate Services, certificates issued by a third-party CA can be associated with a Windows Server 2003 user account stored in Active Directory. This is possible because applications such as Internet Explorer and Internet Information Services (IIS) can be used to authenticate a user to an account stored in Active Directory, based on the UPN name information in a certificate. The account to which the certificate maps provides information about user access rights on the server. This is an extremely powerful feature for Web-based applications and third-party CAs because it combines strong authentication by means of public key technology with the native authorization model of Windows Server 2003. For example, to enable extranet and remote access scenarios without requiring the application and certificate to manage access rights, administrators can use certificates from partner companies and map them to accounts in Active Directory by means of one-to-one or many-to-one mappings.

For more information about using one-to-one and many-to-one mapping, see "Map certificates to user accounts" later in this chapter.

# Selecting an Extended CA Infrastructure Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use one of three configurations to create an extended CA trust infrastructure:

- **Third-party root CA.** Use a third-party CA as a root CA for a new extended CA hierarchy shared between two organizations.
- **New root CA.** Establish your own new root CA to combine separate CA hierarchies for two organizations.
- **Cross-certification and qualified subordination .** Keep the existing CA hierarchies separate, but use cross-certification and qualified subordination to implement as much or as little trust as needed between the two organizations.

There are advantages and disadvantages to each approach. If you need to extend your CA infrastructure to include third-party PKIs, you need to evaluate the requirements of your organization to determine the method that is most appropriate for you.

# Using Third-Party Root CA Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Building a new public key hierarchy from an existing third-party root CA is an appropriate solution if you want to cross-certify with multiple business partners simultaneously. The third-party root CA is used to build a new public key hierarchy designed specifically to serve the needs of multiple organizations. Figure 16.12 shows an example of an extended CA infrastructure built from an existing root CA. In this example, organization A and organization B maintain their existing PKIs and share a new PKI that serves the specific needs of their business relationship.

**Figure 16.12 Extended CA Infrastructure Built from an Existing Third-Party Root CA**




The advantage to this approach is that you can off-load responsibility for maintaining the new infrastructure to an organization that specializes in this type of service. The intermediate and issuing CAs that you create on your side of the shared infrastructure can be administered separately from your existing internal PKI. In this way, the functions of the external PKI cannot compromise the internal PKI, and the organizations that share the extended infrastructure do not have to share transitive trust between their existing PKIs.

The disadvantage to this approach is that it involves the creation of a new, separate infrastructure with its own administrative requirements. Although much of this administration is off-loaded to a third party, this approach involves considerable additional cost. The costs can multiply each time you add a new business relationship that requires a separate shared PKI infrastructure.

You need to plan for an extended PKI based on an existing root CA in the same way that you plan for your existing internal PKI, in that you must decide where to locate intermediate and issuing CAs, how to manage them, how to protect them, and so on. In this case, you must work with your business partner and the root CA provider to make these decisions.

Third-party CAs can form all or part of a CA trust hierarchy. To ensure that third-party CAs provide interoperability with your existing infrastructure, test all proposed interoperability scenarios in your lab.

**Note**

- Some PKIs require CA trust models that are not interoperable with rooted CA hierarchies. Windows 2000 and most commercial CAs support rooted CA hierarchies.

**Adding Trusted Root Certificates to Group Policy**

If a stand-alone CA is not a domain member buts runs as a member of a workgroup, the root CA certificate must be added manually to the domain Group Policy. In contrast, when you install a stand-alone root CA on a computer that is a domain member, the certificate of the CA is added to the Trusted Root Certification Authorities Group Policy for the domain.

You can also add certificates for other root CAs to Trusted Root Certification Authorities Group Policy manually. These root CAs then become trusted root CAs for the computers within the scope of the Group Policy. For example, if you want to use a third-party CA as a root CA in a certification hierarchy, you must add the certificate for the third-party CA to the Trusted Root Certification Authorities Group Policy.

# Using a New Root CA Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You or your partner organization can create a new root CA to establish an extended CA infrastructure that supports your business requirements. The structure of this extended CA infrastructure is similar to that of an extended infrastructure based on a third-party rood CA. With a new root CA configuration, however, you and your partner organization must create a security management infrastructure, and must take responsibility for administering and maintaining the extended PKI. If one organization assumes this responsibility, the other organization must trust that its partner will protect the security interests of both parties.

This option can be more cost-effective than using a third-party root CA. In addition, you can use Windows Update to distribute new root certificates, improving reliability and decreasing costs.

The planning considerations for a new root CA–based extended infrastructure are similar to those that apply to your existing internal PKI. You and your partner organization are responsible for creating administrative policies for the root CA and enforcing the integrity of the new root.

# Using a Cross-Certification Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With the cross-certification method for extending the CA infrastructure, neither party creates a separate PKI; instead, cross-certificates, accompanied by qualified subordination, enable communication between existing public key infrastructures of two organizations to the degree of trust that their business relationship dictates.

Cross-certification creates a shared trust between two CAs that do not share a common root CA. These CAs exchange cross-certificates that allow their organizations to communicate. In this way, the organizations do not have to create and manage additional root CAs. Cross-certification might be the best option if a common root CA for both PKIs does not exist.

Figure 16.13 shows an example of an extended CA infrastructure based on cross-certification between the root CA of organization 1 and a subordinate CA in organization 2.

**Figure 16.13 Extended CA Infrastructure Based on Cross-Certification**




The advantages to using cross-certification to extend the PKI include low cost and a high degree of flexibility, as you can cross-certify at any level in the hierarchy. For example, if a division of organization 2 wants to share information with all of organization 1, the division can cross-certify with the root CA of organization 1. This, however, creates a security risk, as it exposes resources in parts of the organization that are not part of the business relationship. On the other hand, if a division of organization 1 and a division of organization 2 want to share information, the two divisions can cross-certify CAs that are lower in the CA hierarchy. This option is more secure, as the other divisions of the organizations are not unnecessarily exposed.

Cross-certification requires greater administrative overhead than other methods for extending the CA infrastructure, and entails the risk that outsiders might unintentionally be given access to internal resources. If an organization becomes involved in many cross-certification relationships with different levels of trust and different applications, the management overhead can be significant.

# Limiting Unplanned Trusts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you extend your CA trust infrastructure beyond the boundaries of the PKI of your organization, you can inadvertently create unplanned trust relationships.

Unplanned trusts that can occur include:

- Allowing certificates to be used for unintended applications
- Allowing external certificates to be used for longer than intended
- Enabling trust with the extended business partners of a business partners

For example, if company A trusts company B by means of an unconstrained cross-certificate, and company B trusts company C, then company A unintentionally trusts company C. Equally serious problems can occur when company A and company B cross-certify, and company A does not realize that company B does not have the same level of control over the manner in which its certificates are issued and used.

To limit the creation of unplanned trust relationships and the potential security risks that they pose, you can use CA constraints to define limits on your cross-certificate relationships. Constraints in Windows Server 2003 can be based on:

- Use and path length (basic constraints)
- Names
- Issuance policy
- Application policy
- Policy mapping

Implement these constraints when you configure your CA and end user certificates. For more information about defining constraints, see "Using Qualified Subordination to Restrict Certificates" later in this chapter.

**Using Certificate Trust Lists to Limit Unplanned Trusts**

You can use certificate trust lists (CTLs) to limit unplanned trusts. CTLs are the primary means of limiting unplanned trust relationships in Windows 2000 environments.

A CTL is a predefined list of certificates that is signed by a trusted entity. The CTL includes either hashes of certificates or a list of the actual certificate names. In most cases, the CTL is a list of hashed certificate contexts. The CTL allows you to limit the purposes for which certificates issued by an external CA can be used, and the validity period of those certificates.

Windows Server 2003 certificate trust lists allow you to do the following:

- **Create trust certificates from specific CAs without requiring broader trust for the root CA.** For example, you can use certificate trust lists on an extranet to trust certificates issued by certain commercial CAs. If you map certificates that are issued to an account stored in Active Directory, you can grant appropriate permission to users who need access to restricted extranet resources. This is possible because they have certificates issued by the trusted commercial CAs.
- **Restrict the permitted use of certificates issued by trusted CAs.** For example, you can use a certificate trust list on an extranet to restrict the permitted use of certificates to applications such as secure mail.
- **Control the period of time in which third-party certificates and CAs are valid.** For example, the CA of a business partner can have a lifetime of five years and issue certificates with lifetimes of one year. However, you can create a certificate trust list with a lifetime of six months to limit the time that certificates issued by the CA of the business partner are trusted on your extranet.

You might use a CTL to allow users to trust certificates that are issued by a commercial CA and restrict the permitted uses for those certificates. You might also use CTLs to control trust on an extranet for certificates that are issued by CAs that are managed by your business partners.

**Note**

- After a CTL is defined, it must be applied to client computers by means of Group Policy.

# Example: Selecting an Extended CA Infrastructure Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Organizations frequently enter into joint ventures, which can involve the sharing of confidential information, such as engineering data that stored is on an internal network. To facilitate this type of data sharing, an organization can initiate a cross-certified relationship that allows some, but not all, employees of another organization to access data on its network.

One way to enable this cross-certified relationship is to create a subordinate CA to a high security Issuing CA. This subordinate CA is then used to facilitate the joint venture relationship. Although it is possible to cross-certify directly with a corporate high security CA, the advantage of using a separate CA specifically for the joint venture is that it allows you to restrict the capabilities of the people who work for the other partners in the joint venture. They cannot, for example, use their certificates for unintended purposes or to access portions of the network that are not relevant to the joint venture.

Figure 16.14 illustrates the position of the new CA in the CA infrastructure of one organization.

**Figure 16.14 Extended CA Infrastructure**




Creating the CA alone does not enable the new joint venture operations. To enable this sharing, before the CAs are created, administrators must configure the cross-certificates that qualify the trust relationship between the two organizations. These cross-certificates define where, in the first organization, holders of the certificates belonging to the second organization can and cannot go, and which applications they can and cannot use. For information about how to implement these namespace and application limits, see "Using Constraints and Policy Mapping" later in this chapter.

# Defining Certificate Configuration Options

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After your CA infrastructure is in place, you can begin to define the certificate configuration options that you need to meet the requirements of your users, as well as the security needs of your organization. Figure 16.15 shows the steps involved in defining certificate configuration options.

**Figure 16.15 Defining Certificate Configuration Options**



# Selecting Certificate Templates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The certificate services that you deploy and the security requirements that are specific to your organization impact the types of certificates that you issue. You can issue multiple types of certificates to meet a variety of security requirements.

The certificate templates available with an enterprise CA in Windows Server 2000 and Windows Server 2003 provide the default contents of all certificates that can be requested from a Windows enterprise CA. These certificate templates are stored in Active Directory and cannot be used with stand-alone CAs.

Certificate templates can serve a single purpose or multiple purposes. Single-purpose templates generate certificates that can be used for a single application. For example, the Smart Card Logon certificate template is designed for smart card logon only. Multipurpose templates generate certificates that can be used for a number of applications, such as Secure Sockets Layer (SSL), S/MIME, and EFS. For example, a user certificate can be used for both user authentication and EFS encryption.

Both Windows 2000 and Windows Server 2003 support single-purpose and multipurpose templates. However, Windows 2000 and Windows Server 2003 Standard Edition only support version 1 templates, which have read-only attributes that cannot be customized or extended. Windows Server 2003, Enterprise Edition supports version 2 templates, which allow you to create new certificate templates, clone an existing template, and replace templates that are already in use.

**Important**

- If you are already using version 1 templates, you can upgrade them to version 2 templates. However, the domain admins in your top level domain must have full access control permissions on the version 1 templates in order to complete this upgrade. Domain administrators do not need to have full access control over the templates after the upgrade has been completed.

Both version 1 and version 2 certificate templates include the following information:

- Intended user of the certificate.
- CA that issued the certificate.
- Serial number that uniquely identifies each certificate.
- Public key value for the user identified in the subject field.
- Validity period of the certificate.
- Extensions, if any, which apply to the certificate, includingadditional information that can define certificate purposes, restrictions, and management.
- Digital signature of the CA, which verifies the relationship of the certificate to the issuing CA.

  **Note**

  - You can also create your own certificate templates.

Before the certificates are issued, you need to determine the following critical information:

- Certificate key length
- Certificate validity period
- Optional certificate extensions

**Note**

- Certificate templates, in conjunction with the CA policy module, allow you to define certificate policy for CA certificates.

In addition, version 2 templates allow you to configure the following:

- Customized enrollment policies
- Policies related to validity periods
- Policies related to application usage
- Policies related to key usage
- Policies related to key archiving
- Certificate authorization
- Domain authentication
- Certificate administrators
- Signed enrollment agents
- Key creation
- Key and CSP types
- Certificate contents

**Important**

- You must upgrade the schema in an Active Directory forest to Windows Server 2003 in order to support version 2 templates. You do not need to upgrade all domain controllers to Windows Server 2003 to perform a schema upgrade.

Certificate templates can only be used when the server that is running Certificate Services is an enterprise CA. Enterprise CAs can issue a variety of certificate types based on the templates. You can configure each enterprise CA to issue only specific types of certificates. Table 16.6 lists the different types of version 1 certificate templates that are available, and the purposes for each.

**Table 16.6 Version 1 Certificate Templates**

| Certificate template name | Certificate purposes | Issued to |
|---|---|---|
| Administrator | Code signing, Microsoft trust list signing, EFS, secure e-mail, client authentication | Users |
| Authenticated Session | Client authentication | Users |
| Basic EFS | Encrypting File System | Users |
| CEP Encryption | Act as a registration authority | Users |
| Code Signing | Code signing | Users |

| | | |
|---|---|---|
| Domain Controller | Client authentication, server authentication | Computers |
| EFS Recovery Agent | File recovery | Users |
| Enrollment Agent (Computer) | Certificate request agent | Computers |
| Exchange Enrollment Agent (Offline Request) | Certificate request agent | Users |
| Exchange User Signature | Secure e-mail, client authentication | Users |
| Exchange User | Secure e-mail, client authentication | Users |
| IPSEC | IP Security | Computers |
| IPSEC (offline request) | IP Security | Computers |
| Root Certification Authority | Identify the root CA | Computers |
| Router | Client authentication | Computers/routers |
| Smartcard Logon | Client authentication | Users |
| Smartcard User | Client authentication, secure e-mail | Users |
| Subordinate CA | All | Computers |
| Trust List Signing | Microsoft trust list signing | Users |
| User | Authentication, secure e-mail, and EFS | Users |
| User Signature | Secure e-mail, client authentication | Users |
| WebServer | Server authentication | Computers |

Table 16.7 lists the version 2 certificate templates that are available in Windows Server 2003 Advanced Server and the purposes for each.

**Table 16.7 Version 2 Certificate Templates**

| Certificate template name | Certificate purposes | Issued to |
|---|---|---|
| CA Exchange | CA encryption | Computer |
| Cross certification authority | Qualified subordination | Computer |
| Directory E-mail Replication | Directory replication | Users |
| Domain Controller Authentication | Client authentication, server authentication | Users |
| Key Recovery Agent | Key recovery | Users |

**Note**

- When you select and modify templates, create function-based names for the templates, such as domainA_e-mail or

legal_signing. Function-based names help users to select the appropriate certificate for the task that they need to perform.

**Delegating Administration of Certificate Templates**

Although the majority of CA-related tasks are performed by administering the CA itself, certain tasks, including the administration of certificate templates, are controlled through Active Directory.

**To delegate the administration of certificate templates:**

- Right-click the **Certificate Templates** node in the **Certification Authority** snap-in and select **Manage**.
- Double click a certificate template.
- Under the **Security** tab, check the **Allow** boxes for the **Read** and **Write** permissions.

For more information about certificate templates, see *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Selecting Certificate Security Options

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The security requirements for your certificates are based on the following four critical factors:

- **Risk of attack.** The security of your network, the value of the network resources protected by the CA trust chain, and the cost of initiating an attack all impact the security requirements for your certificates.

- **The degree to which you trust your certificate users.** In general, the less you trust your users, the shorter the certificate lifetimes and CRL lifetimes, and the tighter the control over renewals. For example, you might trust temporary users less than you trust normal business users, and therefore you might set shorter lifetimes on the certificates of temporary users and require stricter controls for their renewal.

- **The amount of administrative effort that you are willing to devote to certificate renewal and CA renewal.** For example, to reduce the administrative effort required to renew CAs, you can specify long, safe lifetimes for your certification trust hierarchies.

- **The business value of the certificate.** For example, you place tighter restrictions on certificates used to validate critical data such as purchase contracts than you place on certificates for routine e-mail.

The standard settings for certificates issued by Microsoft Certificate Services meet most security needs. However, you might want to specify stronger security settings for certificates that are used by certain user groups. For example, you can specify longer private key lengths and shorter certificate lifetimes for certificates used to provide security for valuable information.

To ensure that your certificates meet the security requirements of your organization, you need to configure the following:

- The cryptographic algorithms and private key lengths for CAs and issued certificates.

- The amount of time for which the certificates and their keys are valid.

- Certificate renewal and permitted renewal frequency.

# Selecting Cryptographic Algorithms and Key Lengths

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows 2000 and Windows Server 2003 support several well-known cryptographic algorithms that are also supported by other PKI products. When you install a Windows Server 2003 CA, you can select CA key lengths from 384 bits to 16,384 bits, depending on the CSP that you select. In a typical deployment, user certificates have 1,024-bit keys and root CAs have 4,096-bit keys.

Key length is determined in part by the cryptographic algorithm that you select. Table 16.8 lists the algorithms that Windows Server 2003 supports and their minimum and maximum key lengths.

**Table 16.8 Windows Server 2003 Algorithms and Key Lengths**

| Algorithm | Minimum Key Length | Maximum Key Length |
|---|---|---|
| RSA | 384-bit | 16,384-bit |
| DSA | 512-bit | 1,024-bit |

In most cases, the default key length provides an acceptable balance between security and CA performance. However, to provide the maximum possible protection without degrading CA performance, select the largest keys that are practical for CAs. Keys that are at least 1,024 bits long are the best choice for CA certificates.

Keep in mind that generating large keys can place a high load on computer processors and might also increase the amount of time needed for signing operations to excessive levels.

**Note**

- Key length has a minimal impact on the size of a certificate. However, it can be a significant consideration for smart card deployments because of the space constraints of the card.

In general, public and private key lengths do not impact interoperability as long as both environments support a common range of key lengths. If one PKI supports large public keys and another does not, however, the two cannot exchange symmetric keys or sign and verify data.

**Note**

- While key exchange and digital signature operations performed between PKIs do not require the same public and private key lengths, symmetric key algorithms do. In addition, if different key lengths are used, both key lengths must be supported in both environments.

When PKIs do not support the same key lengths, some applications cannot decrypt data that other applications have encrypted. In addition, the PKIs might not be able to establish secure communications channels between applications if the applications cannot agree on symmetric key lengths, as required by protocols such as SSL and TLS.

If a PKI uses public/private keys based on an algorithm such as RSA, all PKI operations can be accomplished with only one key pair. However, single key pairs might not meet the security requirements of your organization or its choice of algorithm. For this reason, Windows Server 2003 supports both single key pairs and dual key pairs. A good PKI is flexible enough to allow as many or as few key pairs as are required by applications.

If one PKI operates according to the number of keys that applications use, it can impact interoperability with other PKIs. For example, an e-mail application could sign a message with a signature-only key and include the associated certificate in a message sent to a recipient without also sending an encryption certificate as part of the message. The recipient might then be unable to discover the encryption certificate of the sender to reply with an encrypted message back to the sender.

# Establishing Certificate and Key Lifetimes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A number of factors impact certificate lifetimes, such as the type of certificate, the security requirements of your organization, the standard practices in your industry, and government regulations. In general, longer keys support longer certificate lifetimes and key lifetimes.

When establishing certificate and key lifetimes, you must consider how vulnerable your keys are to compromise and what the potential consequences of compromise are. The following factors impact the lifetimes that you choose for certificates and keys:

- **The length of private keys for certificates.** Because longer keys are more difficult to break, they justify longer safe key lifetimes.
- **The security of the CAs and their private keys.** In general, the more secure the CA and its private key, the longer the safe certificate lifetime. CAs that are operated offline and stored in locked vaults or data centers are the most secure.
- **The strength of the technology used for cryptographic operations.** In general, stronger cryptographic technology supports longer key lifetimes. You can extend key lifetimes if you enhance private key storage by using smart cards and other hardware cryptographic service providers. Some cryptographic technologies provide stronger security, as well as support for stronger cryptographic algorithms. For example, you might use smart cards for user logon or FIPS 140-1 Crypto Cards for secure mail and secure Web browsers.
- **The vulnerability of the CA certification chain to attack.** In general, the more vulnerable your CA hierarchy is to attack, the longer the CA private keys and the shorter the key lifetimes required.
- **The users of your certificates.** Organizations typically trust their own employees more than they trust employees of other organizations. If you issue certificates to external users, you might want to shorten the lifetimes of those certificates.
- **The number of certificates that have been signed by a dedicated CA.** The more public the CA public key that is used to sign an issued certificate, the more vulnerable it becomes to attempts to break its key.

An expiration date is defined for each certificate at the time that it is issued. An enterprise CA issues certificates with lifetimes that are based on the certificate template for the requested certificate type.

Most certificate templates specify a lifetime of one year. However, the following version 1 certificate templates specify a lifetime of two years:

- CEP Encryption (offline request)
- Enrollment Agent
- Enrollment Agent (computer)
- Enrollment Agent (offline request)
- IPSec
- IPSec (offline request)
- Router (offline request)
- Web Server

The following certificate templates specify a lifetime of five years:

- Domain Controller
- Subordinate certification authority

The lifetimes of certificates issued by stand-alone CAs are determined by system registry settings for the CA.

The certificates for enterprise root CAs and enterprise stand-alone root CAs have a default lifetime of two years. However, you can specify a different lifetime for the CA during installation. The parent CA that approves the certificate request and issues the certificate determines the lifetime of a subordinate CA certificate.

# Creating a Certificate Renewal Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

CAs continue to issue and renew certificates until they reach the end of their established lifetimes. Certificates expire when the issuing CA reaches the end of its established lifetime, unless:

- They are renewed with a new key pair to extend their lifetime.
- They are revoked before the expiration date is reached.
- They are considered to have expired because an issuing CA is unavailable to verify their validity.

Certificate lifetimes impact the security of your PKI for the following reasons:

- Over a period of time, encryption keys become more vulnerable to attack. In general, the longer the period of time that a key pair is in use, the greater the risk that the key can be compromised. To mitigate this risk, you must establish maximum allowable key lifetimes and renew certificates with new key pairs before these limits are exceeded.
- When a CA certificate expires, all subordinate CAs that depend upon this CA for validation also expire.
- When a CA certificate is renewed, all certificates that have been issued by the CA can be renewed for a period based on the nesting guidelines described in the following section.

To reduce the risk of a private key becoming compromised, the private key and public key sets for certificates can be renewed each time the certificates are renewed, instead of when the keys reach their maximum lifetimes. You can renew CAs by assigning them a new key pair or by using the existing key pair. If you create a new key pair and the original certificate has not yet expired, it must have a new Subject Key Identifier (SKI) and a separate CRL. Renewing certificates with new key sets is not possible for some hardware-based CSPs, either because key storage limits prohibit this or because key generation takes a long time.

**Note**

- For more information about configuring certificate renewal, see "Selecting a Certificate Enrollment and Renewal Method" later in this chapter. For more information about the impact of certificate renewal on the use of certificate revocation lists, see "Establishing Certificate Revocation Policies" later in this chapter.

Certificate lifetimes affect the number of certificate renewal requests that are transmitted across your network. For users in remote offices who are connecting to the network across slow links, you might want to lengthen certificate lifetimes to reduce the number and frequency of these requests.

To create a certificate renewal strategy, determine the following:

- Which certificates, if any, are you allowed to renew?
- How often can a certificate be renewed before its key is retired?

In general, certificates with stronger keys that are used less frequently and that are less available to potential attackers can justify longer lifetimes and at least one renewal. Certificates with average key lengths and shorter lifetimes can be renewed more frequently — but not beyond the validity date for the certificate that authorizes the CA that issued the certificate. This is called nested validity or nested expiration.

**Nesting Certificate Lifetimes**

In addition to defining certificate lifetimes for your Windows Server 2003 CAs, you need to confirm that certificate lifetimes and renewals do not extend beyond the lifetimes of the CAs that are above them in the hierarchy.

By default, the certificate for the root CA has a longer lifetime than certificates for the other CAs in the hierarchy. This is

because a Windows Server 2003 CA cannot issue certificates with a lifetime that extends beyond the validity period of its own certificate. If the lifetime specified for a requested certificate type exceeds the expiration date of the certificate of the CA, the CA truncates the lifetime of the issued certificate to match the expiration date for its own certificate.

For example:

- If the end date of a Windows Server 2003 root CA certificate is January 2, 2012, no Windows Server 2003 child CA in the chain below the root can issue a certificate with a date that is past January 2, 2012.

- If a Windows Server 2003 intermediate CA has a certificate end date of January 2, 2008, no Windows Server 2003 child CA can issue certificates with an end date that is past January 2, 2008.

- If a Windows Server 2003 issuing CA has a certificate end date of January 2, 2004, no certificate that the CA issues can have an end date that is past January 2, 2004.

- If the end date of a Windows Server 2003 CA certificate is January 2, 2004, and it receives a request to issue a one-year certificate on August 1, 2002, the CA issues the one-year certificate with an end date of July 31, 2003. However, if the CA receives a request to issue a one-year certificate on August 1, 2003, the CA issues the certificate with an end date of January 2, 2004.

- A Windows Server 2003 CA with a certificate lifetime of five years with an end date of January 2, 2007, can issue one-year certificates until January 2, 2006, or two-year certificates until January 2, 2005. After January 2, 2005, the CA does not issue two-year certificates. It truncates the validity end date to January 2, 2007. Likewise, after January 2, 2006, the CA truncates the validity end date of both one-year and two-year certificates to January 2, 2007.

The more nesting you have in your certification hierarchy, the shorter the certificate lifetimes become. Configure your certificate life cycles in such a way as to avoid short certificate lifetimes and certificate renewal cycles. If you specify long lifetimes for CAs and later discover that the CAs are not secure, you can renew CAs in the certification hierarchy with shorter lifetimes to reduce the potential security risks.

**Note**

- For a worksheet to assist you in preparing your certificate life cycle plan, see "Windows Server 2003 Certificate Life Cycle Plan" (DSSPKI_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows Server 2003 Certificate Life Cycle Plan" on the Web at http://www.microsoft.com/reskit).

# Using Qualified Subordination to Restrict Certificates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many of the certificates that you issue can be used without any further customization. However, you might want to limit the scope of your certificates, whether they are intended to validate a subordinate CA, to cross-certify an external CA, or to enable an end user application. You can limit the scope of a certificate by:

- Defining the namespaces for which a subordinate CA will issue certificates.
- Specifying the acceptable uses of certificates issued by a qualified subordinate CA.
- Creating trust between separate certification hierarchies.

Qualified subordination restricts the certificates issued by the qualified subordinate CA, or by CAs that chain through the qualified subordinate CA, that are acceptable to your organization. You accomplish this by defining the following in the Policy.inf file:

**Note**

- The Policy.inf file is different from the CAPolicy.inf file. The Policy.inf file impacts qualified subordination, whereas the CAPolicy.inf file impacts the CA certificate.

    - **Basic constraints.** Define the certification path length required and allowed for policy identifiers and policy mapping.
    - **Name constraints.** Define the range of namespaces that are permitted or excluded by the qualified subordinate CA and its subordinates.
    - **Issuance policies.** Define the extent to which your organization trusts the identity presented in a certificate. These policies are identified in a certificate by object identifiers.
    - **Application policies.** Define the applications that can be used in conjunction with certain certificates.

In addition, if you are attempting to connect two different PKIs, whether within your organization or with a third-party, you need to use policy mapping to achieve equivalency between the policy constraints that you have defined and the policy constraints defined in the other PKI. The use of constraint extensions and policy mapping allows you to control certificate usage more effectively, and to administer your certificates more effectively.

Qualified subordination allows you to ensure that specific constraints are applied when a CA issues or an application uses a certificate. These constraints ensure that all certificates issued by the CA apply the policy restrictions that you have defined.

By definition, your root CA applies all policies. You can use intermediate CAs to issue certificates that enable different levels of security, such as High Security, Medium Security, and so on. The security policies that you define are identified by means of object identifiers. When certain object identifiers are applied to a CA certificate, all certificates below that CA in the hierarchy must also have a subset of those object identifiers. If you create a certificate chain with no valid policy, any certificates that are issued are considered invalid. However, if you create a certificate chain with no policy object identifiers at all, then the certificates that you issue are considered to match the "any policy" object identifier. Figure 16.16 shows how policy is applied to CAs.

**Figure 16.16 How Policy Is Applied to CAs**




The policies and constraints of each qualified subordinate CA are a subset of the policies and constraints of the parent CA.

# Using Basic Constraints

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Basic constraints allow an application to determine whether a certificate is a CA certificate, which can then be used by the certificate chain engine to build certification paths, or an end-certificate, which cannot.

You can also use basic constraints to limit the maximum number of CA certificates that can be included in a CA path. For example, setting a path length of zero in the basic constraints section of the CAPolicy.inf file allows only certificates issued by that specific CA to be included in the CA path. A path length of two allows only a total of three CA certificates in a certification path. In the latter case, any certification paths that include more than three CAs are discarded.

Use basic constraints if you do not want to trust additional CAs that are created lower in the CA hierarchy of your organization. You can also use basic constraints in cross-certified relationships if you trust your business partner and the certificates from all their existing CAs, but you do not want to trust certificates from any additional CAs that they authorize.

# Using Name Constraints

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Name constraints allow you to designate which namespaces are either permitted or excluded for certificates issued by a qualified subordinate CA. When the qualified subordinate CA receives a request, it compares the names present in the subject and the subject alternate name fields to the configured name constraints, to determine whether the namespace is permitted or excluded. As you design your PKI, you need to decide which individual clients and business units are able to enroll for and use certain certificates. For many organizations, the selected users, computers, and services are members of specific Active Directory domains and subdomains.

You can base name constraints on any of the following types of name formats:

- **X.500 Directory name.** Distinguished names identify users and resources on the network in Active Directory. This allows you to constrain a qualified subordinate CA to permit or exclude users in Active Directory by using the distinguished names of the users. Active Directory also uses distinguished names to create and reference groups of objects in the directory, such as users and computers. The distinguished names of these object groups can also be used as name constraints, allowing you to constrain a qualified subordinate CA to permit and exclude certificate issuance for entire groups in the directory.

- **DNS domain name.** You can apply the DNS namespaces that your network uses for name resolution as name constraints for a qualified subordinate CA. When the qualified subordinate CA receives a certificate request, it compares the DNS name associated with the computer requesting the certificate to its DNS name constraints and decides whether or not to issue a certificate. You can specify a DNS name constraint as a **DNS host name**, such as host1.example.microsoft.com, or as a **DNS namespace**, wherein all DNS host names are permitted or excluded, such as .example.microsoft.com.

- **E-mail and user principal name.** You can specify e-mail and UPN name constraints for an individual subject, such as person@example.contoso.com, or you can specify constraints for all subjects whose e-mail names or UPNs end in a specific name, such as @example.contoso.com. Typically, you need to specify e-mail or UPN name constraints for all subjects whose e-mail addresses and UPNs end in a specific name.

- **Universal Resource Identifier (URI).** URIs are used to identify resources on the Internet by means of identifiers such as URL, FTP, HTTP, telnet, mailto, news, and gopher. When validating the URI names in a certificate request, the qualified subordinate CA ignores the protocol element in the URI, such as http:// or ftp://, and uses the domain or host names only.

- **IP address.** IP address name constraints follow the formatting conventions specified in RFCs 791 (IPv4) and 1883 (IPv6). The IP addresses contained in the certificate requests made to a qualified subordinate CA are compared to the IP addresses in the name constraints of the qualified subordinate CA.

You can configure name constraints to result in the following outcomes:

- **Permitted.** The certificate request contains all names that are listed as permitted in the CA name constraints extension of the issuer.

- **Not permitted.** The certificate request contains a name that is not listed as permitted in the name constraints extension of the issuer.

- **Excluded.** The certificate request contains a name that is listed as excluded in the name constraints extension of the issuer.

A CA certificate can contain name constraints that are applied to all certificate requests made to the CA. Each request is compared to the list of permitted and excluded names to determine whether the name in the certificate is considered permitted, not permitted, excluded, or not defined. When you include name constraints in a CA certificate, the following rules are applied to the subject name and alternate subject name fields:

- **Excluded namespaces take precedence over permitted namespaces.** A qualified subordinate CA will not issue a certificate to a user within an excluded namespace even if the user is also within a permitted namespace. For example, a user might be within the permitted Active Directory namespace *.contoso.com* but also within the excluded DNS namespace *.uvw.contoso.com*. The excluded DNS namespace overrides the permitted Active Directory namespace and the certificate request of the user fails.

- **If the name constraints extension exists in a CA certificate, all name constraints must be present in the appropriate format.** Any name formats that are not included are considered to be wild cards that match all possibilities. For example, if the DNS name constraint is absent, the entry is treated as DNS="".

- **All name constraints are considered, even if they are not specified.** No precedence is applied to the listed name constraints. For this reason, name constraints that are not present are treated as wildcards. For example if you only restrict the DNS name space, the Name Constraints extension sets the remaining name constraints to allow all name spaces.

- **Name constraints are applied to the Subject Name extension and any existing Subject Alternate Name extensions.** For example, if a user can be identified by a DNS domain name and an alternate e-mail name, name constraints apply to both.

- **Name constraints apply to all names contained in a certificate request.** Each name in the subject or subject alternate name extensions must match at least one of the name constraints listed for that name type. A certificate request that includes a subject name or subject alternate name that does not match a listed name type is rejected.

- **Name constraints are not case sensitive.** For example, *.contoso.com* is treated the same as *CONTOSO.COM or ConToso.Com*.

**Important**

- Name constraint validation is performed on the CA, not on the client. However, you must have Windows XP and Windows Server 2003 clients in order to use name constraints.

# Using Issuance Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use issuance policies to define the extent to which your organization trusts the identity presented in a certificate. For example, you can set an issuance policy stipulating that you only trust certificates that were issued during a face-to-face meeting with a network administrator, such as when a smart card certificate is issued.

An object identifier must describe every issuance policy that you define. The inclusion of an issuance policy object identifier in an issued certificate indicates that the certificate was issued in a manner that meets the issuance requirements associated with the issuance policy object identifier.

**Note**

- Issuance policy is only available on Windows Server 2003 CAs. Windows 2000 does not provide issuance policy.

You can use a specific certificate template to define one or more issuance policy object identifiers that need to be included in any certificates issued. Windows Server 2003 includes four predefined issuance policies:

- **All Issuance (2.5.29.32.0).** The all issuance policy indicates that the issuance policy contains all other issuance policies. Typically, this object identifier is only assigned to CA certificates.

- **Low Assurance (1.3.6.1.4.1.311.21.8.x.y.z.1.400).** The low assurance object identifier is used to represent certificates that are issued with no additional security requirements.

  **Note**

  - The *x.y.z* portion of the object identifier is a randomly generated numeric sequence that is unique for each Windows Server 2003 forest.

- **Medium Assurance (1.3.6.1.4.1.311.21.8. *x.y.z* .1.401).** The medium assurance object identifier is used to represent certificates that have additional security requirements for issuance. For example, a smart card certificate that is issued in a face to face meeting with a smart card issuer might be considered a medium assurance certificate and contain the medium assurance object identifier.

- **High Assurance (1.3.6.1.4.1.311.21.8. *x.y.z* .1.402).** The high assurance object identifier is used to represent certificates that are issued with the highest security. For example, the issuance of a key recovery agent certificate might require additional background checks and a digital signature from a designated approver because a person holding this certificate can recover private key material from a Windows Server 2003, Enterprise Edition CA.

In addition, you can create your own object identifiers to represent custom issuance policies. For example, two organizations involved in a purchaser/seller relationship can define custom object identifiers to represent digital signature certificates for specific purchase amounts. In such a case, an object identifier can be defined for purchase between $100,000 and $500,000 and another object identifier can be defined for purchases greater than $500,000. Applications can then use these object identifiers to recognize whether a person had the appropriate signing authority for a specific volume purchase.

**Applying Policy Mapping**

In many cases, the administrators of two PKIs define their own policies and object identifiers. In some cases these policies are identical, but in most cases there are small differences between them. For example, one organization might stipulate that one physical form of identification is sufficient to grant a certificate request, while a second organization requires three forms of physical identification to grant a similar request. In these cases, you need to negotiate with the administrators of the other PKI to define terms of equivalence before the cross-certified relationship can be established. This is called policy mapping.

Policy mapping enables interoperability between two organizations that apply similar issuance and application policies, but have deployed different object identifiers. If the policy object identifier (for example, 1.2.3.4) of one company represents a specific function, and the policy object identifier of another company (for example, 11.22.33.44) represents the same function, they can be mapped, so that 11.22.33.44 and 1.2.3.4 are interchangeable.

The qualified subordinate CA that contains this mapping is called the *issuer CA* and the subordinate CA whose policies have been mapped is called the *subject CA*. In mapping some or all of the policies of the subject CA to the policies of the issuer, the issuer CA effectively subordinates the subject CA. The result of this mapping is that users and computers in the issuer CA trust hierarchy can use their own certification paths to validate certificates from users and computers in the subject CA trust hierarchy. The separate trust hierarchy can be within the same intranet or in separate PKI environments over an extranet.

**To apply policy mapping**

1. Identify the trust hierarchy with which you want to establish a trust relationship.

2. Establish equivalence in the assurance levels used by the two trust hierarchies involved in the trust relationship.

3. Obtain the issuance and application policy object identifiers used in both trust hierarchies involved in the trust relationship.

4. Map the issuance and application policy object identifiers in the separate trust hierarchies, and define their policy constraints in the CA certificate request for the qualified subordinate CA that you are installing in your trust hierarchy.

5. Install the qualified subordinate CA with the policies, policy mappings, and policy constraints in your trust hierarchy.

**Constraining Policy Mapping**

You can refine policy mapping by setting parameters for how the issuance policy defined in qualified subordination affects other CAs below the qualified subordinate CA. These parameters can help to limit unplanned trust relationships. The following two settings define this relationship:

- **Require explicit policy.** Specifies the number of certificates that can exist in the hierarchy below a certificate before an explicit policy must exist. For example, if the explicit policy is configured with a setting of three, the defined issuance policy must exist for three layers of the hierarchy. The CA on which the qualified subordination is defined is the first level.

- **Inhibit policy mapping.** Specifies the number of additional certificates that can appear in the path before policy mapping is no longer permitted. For example, an inhibit policy mapping value of three restricts the policy mapping to only three levels of CAs below the qualified subordinate CA.

# Using Application Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certificates provide important information that is not specific to an application. However, you might need to define which applications can be used in conjunction with certain certificates. Application policy allows you to ensure that certificates are only used with the applications that you specify.

An application can also be written to accept only certificates that contain specific application policies. When the application receives signed information from a user, the application reviews the certificate associated with the private key used to sign the information, and ensures that the application policy extension contains the object identifiers required by the application.

Application policies are similar to the Extend Key Usage (EKU) extension in a certificate, as both use one or more object identifiers to prescribe how the public key in a certificate must be used. Windows Server 2003 supports Extend Key Usage to support PKIs that use this extension, but application policies are used in place of EKU.

Application policy is Microsoft specific and is treated much like Extended Key Usage. If a certificate has an extension containing an application policy and also has an EKU extension, the EKU extension is ignored. If, however, a certificate has only an EKU extension, the EKU extension is treated like an application policy extension. If a certificate has an application policy extension and an EKU property, the effective policy for the certificate is the common policy between the EKU property object identifiers and the application policy object identifiers.

**Note**

- If you are issuing certificates that include both application policy and EKU extensions, ensure that the two extensions contain identical object identifiers.

# Using Constraints and Policy Mapping

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The method or methods that you use to limit unplanned trust depend in large part on the security challenges that you must address in your organization.

For example, use name constraints to limit unplanned domain-based trusts. If your organization has more than one domain — such as contoso.com and subdomain.contoso.com — you might only want cross-certificates to map to CAs in the contoso.com domain and not to CAs in the subdomain.contoso.com domain. On the other hand, if you have five different domains and want your cross-certificates to apply to three of them, path constraints provide a more flexible solution.

Use path constraints if transitive trusts create problems for your organization — for example, if you have an environment in which users can freely install subordinate CAs and you do not have strong security guidelines governing CA creation and management.

Policy constraints are the most useful of the constraint options, both internally and externally. Name and path constraints might not provide you with sufficient protection in certain cross-certified relationships if the security standards of your business partners are not as strong as yours. For example, organization A might have a policy that all user certificates must be approved manually by an administrator, while organization B approves certificate requests automatically as long as the user has an e-mail account. The security administrators for organization A must ensure that the certificates that users in organization B use to access resources meet the higher security standards that are necessary in organization A. They can accomplish this by using policy constraints.

Use policy mapping if the organization that you are cross-certifying with has policies that are similar to those of your organization. Policy mapping is less useful when the policies of your organization are stricter than the policies of the other organization, or vice versa. If this is the case, use policy constraints to restrict your trust relationship instead of using policy mapping.

# Example: Configuring Certificates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After an organization has defined its certificate requirements, internal PKI configuration, and external infrastructure, it needs to determine the certificate lifetimes, encryption key lengths, renewal policies, and other restrictions, if any, that apply to the use of each type of certificate. Figure 16.17 shows the certificate design decisions of one organization.

**Figure 16.17 Example of a Windows Server 2003 Certificate Lifecycle Plan Worksheet**

Windows Server 2003 Certificate Lifecycle Plan

| Prepared By | | Date | |
|---|---|---|---|

Instructions: Col 1. Start with the standalone root CA certificate. Then, in subsequent rows, list the certificates for the policy CA certificates it authorizes. Next, list the issuing CA certificates authorized by the policy CAs. Finally, list the end user certificates issued by the issuing CAs. Col. 4. Describe how often the certificate can be renewed and how often a new key should be used.

Notes: (1) The greater the security requirement, the larger the key should be. (2) The root CA has to have the longest life of all CA certificates. (3) Set the certificate and key renewal frequency to avoid invalidating subordinate CA certificates.

| Name and Purpose of Certificate | Key Length (bits) | Certificate Life | Private Key Renewal Strategy |
|---|---|---|---|
| Stand-alone root CA | 4,096 | 20 years | Renew at least every 10 years to ensure that policy CA certificates can be issued with lifetimes of 10 years. Renew by using a new key at least every 20 years. |
| Stand-alone policy CAs | 2,048 | 10 years | Renew at least every 5 years to ensure that child issuing CAs can be issued for a full 5 years. Renew by using a new key at least every 10 years. |
| Enterprise issuing CA s for medium security certificates | 2,048 | 5 years | Renew at least every 3 years to ensure that certificates can be issued for a full 2 years. Renew by using a new key at least every 5 years. |
| Enterprise issuing CA s for high security certificates | 2,048 | 5 years | Renew at least every four years to ensure that certificates can be issued for a full year. Renew by using a new key at least every 5 years. |
| Enterprise issuing CA for external certificates | 2,048 | 5 years | Renew at least every 4 years to ensure that certificates can be issued for a full year. Renew by using a new key at least every 5 years. |
| Secure mail and secure browser certificates | 1,024 | 1 year | Renew by using a new key at least every 2 years. |
| Smart card certificates | 1,024 | 1 year | Renew by using a new key at least every 2 years. |
| Administrator certificates | 1,024 | 1 year | Renew by using a new key at least every 2 years. |

| Name and Purpose of Certificate | Key Length (bits) | Certificate Life | Private Key Renewal Strategy |
|---|---|---|---|
| Secure Web server certificates | 1,024 | 2 years | Renew by using a new key at least every 2 years. |
| Business partner user certificates for an extranet | 1,024 | 6 months | Renew by using a new key at least every year. |

For a worksheet to assist you in documenting your certificate lifecycle plan, see "Windows Server 2003 Certificate Lifecycle Plan" (DSSPKI_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows Server 2003 Certificate Lifecycle Plan" on the Web at http://www.microsoft.com/reskit).

All certificates are issued by Windows Server 2003 CAs. The certificates for the people working for the business partners (for the extranet) can be issued by the Windows Server 2003 CA or by the CA of the business partner. CTLs allow the extranet domain to trust the certificates of the business partners. Where appropriate, stand-alone CAs provide flexible lifetimes for CAs. The renewal of certificates with new keys limits the amount of time that keys are in use and reduces the risk of key compromise.

This organization does not have unusual security requirements that require the use of one cryptographic algorithm over another. Therefore, they chose to accept the default cryptographic algorithms that have been established for each type of certificate and CA.

The certificates issued to the business partners of this organization are constrained by namespace, by path length, and to specific applications. In addition, the corporation uses policy mapping to specify the authentication procedures required of business partner users who are issued certificates to access the resources of the first organization.

# Creating a Certificate Management Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have configured certificates for your organization, you need to create a plan for managing certificates throughout their lifetimes.

Figure 16.18 shows the process for creating a certificate management plan.

**Figure 16.18 Creating a Certificate Management Plan**



# Selecting a Certificate Enrollment and Renewal Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To enable enrollment, you need to specify the enrollment and renewal processes for your certificates. Enrollment involves either configuring permissions to establish which security principals have Enroll permissions for specific templates (in the case of enterprise CAs) or appointing a certificate administrator who reviews each certificate request and issues or denies the request based on the information provided.

Microsoft Certificate Services supports the ability to process certificate requests manually, if administrative approval is required, or automatically, if no approval is necessary. The following enrollment and renewal methods are available:

- **Certificate autoenrollment and renewal.** Allows you to automatically issue certificates that enable PKI applications, such as smart card logon, EFS, SSL, and S/MIME, to users and computers within an Active Directory environment. Certificate autoenrollment is based on a combination of Group Policy settings and certificate templates, which allows you to enroll computers when they start up and to enroll users when they log on to their domain.

  **Note**

  - To use autoenrollment, you need a Windows Server 2003 domain controller, a Windows XP Professional client, and a Windows Server 2003 Advanced Server enterprise CA.

- **Certificate Request Wizard and Certificate Renewal Wizard.** Available from the Certificates console, you can use the Certificate Request Wizard to request a certificate from an active enterprise CA on behalf of a user, computer, or service.

  **Note**

  - This option can only be used for Windows 2000, Windows Server 2003, and Windows XP users, computers, and services.

- **Web Enrollment Support pages.** Contain Active Server Pages and ActiveX controls that provide a Web-based user interface to a CA. By default, the Web Enrollment Support pages are automatically installed on the computer on which the CA is installed, but you can also install the Web Enrollment Support pages on another Windows Server 2003 computer. You can also customize Web Enrollment Support pages. For example, you can limit user options or provide additional links to online user instructions and user support information.

  **Note**

  - You can use Web Enrollment Support pages on stand-alone CAs to issue most of the same types of certificates that enterprise CAs can issue, with the exception of certificates for smart card logon and for autoenrollment, which must be issued and renewed by an enterprise CA. The Web Enrollment Support pages that are installed on stand-alone CAs do not use certificate templates, so all information about the certificate, including information about the requester (and, if asking for a specific application, a correct object identifier), must be specified in the certificate request.

- **Smart card enrollment station.** Advanced version of the Web Enrollment Support pages that allows trusted administrators or security personnel to enroll for smart card certificates on the behalf of other users. For more information about using the smart card enrollment station, see "Planning a Smart Card Deployment" in this book.

To select the certificate enrollment and renewal processes that are appropriate for your organization, you need to consider the following:

- **The users, computers, and services for which you intend to provide services.** Determine whether they are internal or external to the organization. Identify the operating systems they are running and determine whether or not they are connected to Active Directory.

- **The operating system that your clients are using.** Clients running Windows Server 2003 and Windows XP can use the Certificate Request Wizard, autoenrollment, or the smart card enrollment station. Windows 2000 supports the Certificate Request Wizard but does not support smart card autoenrollment. Autoenrollment and the smart card enrollment station also require Active Directory. Most other clients can use their Web browsers to access Web-based enrollment and renewal services.

- **The policies that you establish in order to manage certificate distribution.** This includes both the procedural policies that you establish for your PKI, and the Group Policy settings that you use to implement those policies.

- **The type of CA that is issuing the certificates.** For example, you must have a Windows 2000 or Windows Server 2003 enterprise CA to use the smart card enrollment station. Only Windows Server 2003 CAs support smart card autoenrollment.

Selecting certificate enrollment and renewal processes involves making decisions about the following:

- Automatic versus manual requests

- Automatic versus manual approval

- An enrollment and renewal user interface

- CA certificate renewal

# Selecting Automatic vs. Manual Requests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Whether you choose to generate certificate requests automatically or manually depends on the types of certificates that you intend to use and the number and type of clients that you enroll. For example, if you want all users or computers to use a certain type of certificate, it is not practical for you to require that each certificate be requested individually. Although rolling out a new certificate to all users or computers at one time can generate a large amount of network activity, you can control that activity by deploying the certificate requests for each organizational unit one at a time.

On the other hand, you might want to have users or an administrator request certain high-security certificates, such as those used for digital signing or administrative tasks, only when needed. This can improve administrative control over these certificates — particularly if certificate use is not limited by a user or computer OU, or security group membership.

You can improve control over your certificates by using one of the following options to limit user certificate requests:

- **Restrict access to specific templates.** Configure the discretionary access control list (DACL) for each template so that only the required security principals have Enroll and Read permissions for particular templates.

- **Automate the deployment of computer certificates.** Configure Group Policy to automatically assign the necessary computer certificates by adding the certificate template to the Automatic Certificate Request Settings option in Group Policy.

  **Tip**

  - Autoenrollment is most useful for issuing and renewing computer and IPSec certificates.

# Selecting Automatic vs. Manual Approval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Users can request a certificate from a Windows Server 2003 CA either manually or automatically. This request is held until an administrator approves it, if manual approval is required, or until the verification process is completed. When the certificate request has been approved, the autoenrollment process installs the certificate automatically, or automatically renews the certificate on behalf of the user, based on the specifications in the certificate template.

Most of the time, you choose the same method for certificate approval that you choose for certificate requests — but not always. For example, if you have the appropriate Group Policy and DACL restrictions on your certificate templates, you might decide to approve automatically a certificate request that was generated manually. Conversely, in some cases, it is appropriate to manually approve certificate requests that are automatically generated.

**Note**

- You can use strong authentication to enhance the security associated with autoenrollment. With strong authentication, the certificate template uses a specify policy object identifier to require an additional signature on the certificate request. For example, you can set a policy that requires the use of a smart card to provide a stronger authentication method for autoenrollment requests, or you can require approval for automatic certificate requests, so that administrators must approve pending requests.

However, in general:

- For routine and high volume certificates, such as e-mail certificates, automatic approval is the best option for certificate approval as long as the certificate requester has already been authenticated with a valid set of domain credentials.

- When a high degree of administrative oversight is required, such as for software code signing certificates, consider processing certificate requests manually. By using the Certificate Request Wizard, you can evaluate every certificate request individually — or you can delegate this responsibility to another administrator.

# Selecting an Enrollment and Renewal User Interface

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The user interface that you select for certificate request and approval processing depends on whether you choose automatic or manual certificate request and approval methods. If you decide to use autoenrollment for both certificate requests and certificate approval, you must use a minimal user interface.

However, if all or part of the enrollment process is manual, you must decide between using the Web Enrollment Support pages or the Certificate Request Wizard. The Web Enrollment Support pages are the easier interface for users to use. Users can perform the following tasks from the Web Enrollment Support pages:

- Request and obtain a basic user certificate.
- Request and obtain other types of certificates by using advanced options.
- Request a certificate by using a certificate request file.
- Renew certificates by using a certificate renewal request file.
- Save a certificate request to a file.
- Save the issued certificate to a file.
- Check on pending certificate requests.
- Retrieve a CA certificate.
- Retrieve the latest certificate revocation list from a CA.
- Request smart card certificates on behalf of other users (for use by trusted administrators).

However, administrators might prefer to use the Certificate Request and Renewal Wizard. You can start the wizard from the Certificates snap-in. Because the wizard is linked to the Certificates snap-in, you can also create custom snap-ins that you can distribute to certification authority administrators to whom you have delegated specific roles.

Unless an organization uses firewalls between one part of the organization and another, you can use the Certificates snap-in or the Web interface interchangeably. If a firewall exists between the CA and the requesting client, you must request certificates by means of the Web Enrollment Support pages or ensure that port 135 and a dynamic port above 1024 is open for MMC DCOM communication.

Whether you choose to use the Web Enrollment Support Pages or the Certificate Request and Renewal Wizard, you might need to prepare documentation that describes how users can request a user certificate, what users can expect after they request the certificate (for example, automatic enrollment or a delay pending administrator approval), and how they can use the certificates after they receive them.

# Using CA Certificate Renewal

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a CA has issued and supports a large number of certificates, the quality of service and user satisfaction can decline. However, enrollment and renewal are relatively infrequent activities that can be anticipated and therefore planned for well in advance. Many organizations fail to plan for certificate renewal because this activity does not occur until some point well into the future.

When the certificate of a CA expires, the CA can no longer provide certificate services. To provide uninterrupted certificate services, use the Certificates console to renew the CA certificate before its expiration date. The interval that is required for CA renewal depends on the certificate lifetime of the CA.

After you renew a CA, the CA continues to issue certificates by using the new CA certificate, and the cycle starts over. Unexpired certificates that were issued by the pre-renewal CA continue to be trusted until they expire or are revoked.

You can use the standard enrollment and renewal methods that are available in Windows Server 2003 to renew your CAs and certificates. You can renew certificates with the same private key and public key set or with new private and public keys. However, if you have special needs, you can develop custom certificate enrollment and renewal applications for CAs.

**Caution**

- Do not renew certificates with the same private and public key sets if the renewal period exceeds the maximum safe key lifetime. The safe key lifetime is based on your choice of key lengths. Longer keys allow for longer safe key lifetimes.

The Windows Server 2003 Certificate Services Entry module supports industry-standard certificate requests by using remote procedure call (RPC) requests or HTTP requests. You can develop custom applications that make certificate requests to Certificate Services CAs.

# Map certificates to user accounts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have decided how you are going to distribute your certificates, you must decide how to get the certificates to the intended clients, whether they are computers, internal users, or external users. You need to use certificate mapping for many types of certificates, such as those used for smart card logons.

If you have Active Directory, you can map certificates to clients based on their domain or organizational unit membership.

You must decide how you define subject and issuer name information in certificates because this directly impacts applications that use PKIs. For example, if a certificate does not contain the e-mail address name as part of the Subject or Subject Alternative name, some older e-mail applications cannot accept the certificate for digitally signing or encrypting e-mail messages.

Certificate mapping allows you to provide a more secure method for user authentication. With certificate mapping, you link a specific certificate to the account of a user. A server application can then use public key technology to authenticate the user by means of this certificate.

When certificate mapping is enabled, users are authenticated in Active Directory on the authority of the mapped certificates, and are granted rights and permissions based on the authentication.

You can map certificates to user accounts in the following ways:

- One-to-one mapping. This creates an association between an individual certificate and a corresponding user account in Windows 2000 or Windows Server 2003.

- Many-to-one mapping. This creates an association for all certificates from a specific CA to a Windows 2000 or Windows Server 2003 user account.

You can also use certificate mapping to authenticate external users who do not have an account in Active Directory.

**Using One-to-One Mapping**
One-to-one mapping requires more administrative effort than many-to-one mapping. Use one-to-one mapping when you have a relatively small number of clients. If you decide to use one-to-one mapping for a large number of clients, develop custom Web enrollment pages by using Active Server Pages (ASP) technology to automate the mapping process.

**Using Many-to-One Mapping**
Many-to-one mapping is useful for authenticating large numbers of users who require access to a given resource on your network, such as an internal Web site. The CA that issues certificates to these users must be chained to a trusted root for your site, domain, organizational unit (OU), or forest. You can then set rules that map all certificates issued by that CA to a single user account in Windows Server 2003.

Mapping rules check the information that is contained in the certificates of users, such as user organization and the issuing CA, to determine whether the information matches certain criteria. When the information in a user certificate matches the criteria, the user is mapped to a specified user account. The permissions set on the user account apply to all users who hold certificates issued from the trusted CA.

You can use separate many-to-one certificate mappings for different groups that require access to resources on your network. You can configure user accounts that grant different sets of rights and permissions on the basis of client ownership of valid certificates that match the mapping rules. For example, you can map your employees to a user account that grants read access to the entire Web site. Then, you can map consultants and employees of business partners to other user accounts that allow access only to non-confidential information and selected proprietary information.

**Selecting IIS vs. Active Directory Mapping**
You can use either Internet Information Services (IIS) or Active Directory to create your mapping. When IIS does the mapping, the certificate is compared to a list of rules that IIS maintains in its database until it finds a rule that matches the account indicated. You can configure IIS mapping for each Web server. This type of mapping is useful if you need only a limited number of mappings or a different mapping on each Web server.

In Active Directory mapping, when the IIS server receives a certificate from the user, it passes it on to Active Directory, which maps it to a Windows 2000 or Windows Server 2003 user account. The IIS server then logs on the account.

You can create an Active Directory mapping in one of two ways. You can rely on UPN mapping, or, if UPN mapping is not possible, you can manually map a certificate to the account of a user.

Use Active Directory mapping when the account mappings are identical on all IIS servers. Active Directory mapping is easier to maintain than IIS mapping because you only have to create the mapping in one location.

For more information about creating a mapping, see "Mapping certificates to user accounts" in Help and Support Center for Windows Server 2003.

# Establishing Certificate Revocation Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

All certificates have specified lifetimes. However, in some situations, you might need to invalidate a certificate before it has reached the end of its lifetime. Creating policies for certificate revocation involves the following tasks:

- Defining the conditions that warrant the revocation of a certificate.
- Selecting a certificate revocation list publication location.
- Selecting the type or types of CRLs that you intend to use.
- Establishing schedules for the publication of CRLs.
- Establishing a cached CRL validity period.

# Defining Conditions for Certificate Revocation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Not all PKIs need to be supported by the publication of CRLs. For example, if your certificates provide only a low- or medium-level of security and are unlikely to be misused, or if they have short lifetimes, there might not be a need to create and distribute lists of revoked certificates. If, on the other hand, your certificates have a high perceived value and a lifetime that is long enough to cause potential misuse, you need to create and distribute certificate revocation lists on a regular basis.

Before you create certificate revocation schedules, define all the circumstances that justify the revocation of certificates in your organization. For example, you might choose to revoke certificates for any of the following reasons:

- An unauthorized user has gained access to the private key of the certificate.
- An unauthorized user has gained access to the CA. In this case, all the certificates that the CA has published must be revoked and reissued.
- Certificate criteria have changed; for example, an employee has moved to a different department.
- The certificate has been superseded. For example, you might decide to use a different encryption protocol or a longer key.
- The CA that issued the certificate is no longer operating.
- The status of the certificate is on hold. When a certificate has been revoked, it cannot be renewed. However, if the status of a certificate is questionable, you can revoke it and then rescind the revocation if necessary, or re-revoke it for another reason.

  **Note**

  - Use Certificate Hold sparingly because issues can develop involving the period that the certificate was on hold. For example, if a certificate was on hold for three hours but the hold is then removed, attempts to use the certificate during the hold period can yield unexpected results.
- Users misuse their security permissions or the private keys of users are compromised (a smart card is lost, for example).
- A computer is replaced or permanently removed from service, or the private key of the computer is compromised.

**Note**

- A root certificate cannot be revoked by means of a CRL because a root CA is self-signed. A CRL includes only certificates that are issued by a dedicated CA. The alternative is to revoke all the certificates issued by the root CA. The CA certificate becomes obsolete if there are no more valid certificates.

# Selecting a CRL Publication Location

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Selecting a location for CRL publication involves answering the following questions:

- Are the certificate revocation lists needed internally, externally, or both?

  CRLs have to be published where they can be accessed to validate or invalidate certificates. If the PKI is within the firewall of the organization and certificates are published to Active Directory, then LDAP can be used to publish CRLs. If the certificates are used outside the organization or if there is no directory service, http can be used to publish CRLs to a Web server because HTTP traffic can travel more reliably through a firewall than LDAP traffic.

- Do you require multiple CRL publication locations for fault tolerance or to support greater numbers of geographically diverse clients?

  If the answer is yes, choose the domain controllers and Web servers that provide you with greater coverage and improved response times. This way, if one CRL publication point becomes unavailable, the information is available from other publication points.

Figure 16.19 shows this decision process.

**Figure 16.19 Selecting a CRL Publication Location**



# Selecting a CRL Distribution Point

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because CRLs are valid only for a limited time, PKI clients need to retrieve a new CRL periodically. Windows Server 2003 PKI applications look in the CRL distribution point extension for a URL that points to a network location from which the CRL object can be retrieved. Because CRLs for enterprise CAs are stored in Active Directory, they can be accessed by means of LDAP. In comparison, because CRLs for stand-alone CAs are stored in a directory on the server, they can be accessed by means of HTTP, FTP, and so on as long as the CA is online. Therefore, you should set the CRL distribution point after the CA has been installed.

The system account writes the CRL to its distribution point, whether the CRL is published manually or is published according to an established schedule. Therefore you must ensure that the system accounts for CAs have permission to write to the CRL distribution point.

Because the CRL path is also included in every certificate, you must define the CRL location and its access path before deploying certificates. If an application performs revocation checking and a valid CRL is not available on the local computer, it rejects the certificate.

You can modify the CRL distribution point by using the Certification Authority MMC snap-in. In this way, you can change the location where the CRL is published to meet the needs of users in your organization. You must move the CRL distribution point from the CA configuration folder to a Web server to change the location of the CRL, and you must move each new CRL to the new distribution point, or else the chain will break when the previous CRL expires.

**Note**

- On root CAs, you must also modify the CRL distribution point in the CAPolicy.inf file so that the root CA certificate references the correct CDP and AIA paths, if specified.

If you are using certificates on the Internet, you must have at least one HTTPs-accessible location for all certificates that are not limited to internal use.

# Selecting a CRL Type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 includes two types of CRLs: base CRLs and delta CRLs. Base CRLs include a complete list of revoked certificates, and therefore they can grow quickly in organizations that issue a large number of certificates. Because updating base CRLs consumes a large amount of network bandwidth, you must weigh the benefits of publishing expired certificates against the costs in terms of time and network resources. Base CRLs must be published frequently to ensure that they remain current.

Delta CRLs enable you to simplify CRL management. A delta CRL contains information only about the certificates that have been revoked after the last base or delta CRL was published; therefore the data in a delta CRL is accurate throughout its lifetime. Because delta CRLs are smaller than base CRLs, they require less bandwidth to replicate across the network, and they can be published more frequently, thereby enhancing the security of your PKI. By combining base and delta CRLs, you can maximize the effectiveness of the CRLs in your organization and reduce the management effort required.

Whether to use delta CRLs in conjunction with your base CRLs depends on the following factors:

- **Compatibility.** Only Windows XP clients can recognize delta CRLs.
- **Volume.** If you revoke a large number of certificates, your delta CRLs can approach base CRLs in size. Therefore, it is not useful to use delta CRLs when large numbers of certificates are revoked between base CRL publication dates.
- **Online status.** It is best if delta CRLs are not used with offline CAs.

Figure 16.20 shows the process for selecting a CRL type.

**Figure 16.20 Selecting a CRL Type**



# Establishing a CRL Publication Schedule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

CAs publish CRLs listing the serial numbers of certificates that have been revoked according to an established publication schedule. The frequency of publication is based on the number of changes that take place in a given period of time, and how critical the need to maintain up-to-date revocation information is to your organization.

As you create your CRL policies, you need to specify publication schedules for CRLs. By default, enterprise CAs publish CRLs weekly to Active Directory, and stand-alone and enterprise CAs publish CRLs weekly to a directory on the CA server. For example, you can specify that certain CRLs are distributed to Web pages as well as to Active Directory, and that certain CRLs are published daily instead of weekly.

If you use delta CRLs, a typical publication schedule might look like this:

- Publish base CRLs every week.
- Publish delta CRLs every day.

The CRL schedule for the certificates of your issuing CA must account for potential network and server downtime. In addition, it must account for latency in Active Directory replication. For these reasons, make the CRL publication schedule longer than the maximum replication latency.

Make sure that your publication schedule is shorter than the validity period of the CRL. With a validity period of one week, your CRL will be up-to-date for most purposes. If you require an additional buffer to protect against interruptions in communications, you can publish the CRL every two days.

Your strategy for renewing CAs also impacts your CRL publication strategy. If you choose to reuse the existing key pair when you renew a certificate for a CA, the existing CRL covers certificates issued under both the old and new CA certificates. If you choose to renew certificates by using a new key pair for the CA, you need to issue one CRL for the certificates issued with the old key pair, and another CRL for certificates issued with the new key pair. Although both old and new certificates were issued by the same CA, the validity of the older certificates will be validated against one CRL, and the validity of the newer certificates will be validated against the other CRL.

**Note**

- CRLs are published for all CA keys. You cannot selectively publish a CRL for only one CA key pair.

# Setting the Cached CRL Validity Period

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To reduce the amount of network bandwidth needed to retrieve CRLs, the CRL that is specified in the CRL attribute of the certificate is cached on the client system using the certificate. You can control the schedule by which the client retrieves updated CRLs by setting the CRL lifetime.

CRL publication and client use of the most recent CRL are independent. The client does not retrieve a new CRL from its distribution point unless the lifetime of a matching cached CRL has expired. Therefore, when you set the CRL validity period, be sure to balance the intended and actual CRL lifetime.

The only way to force a client to retrieve the latest CRL from the CRL distribution point before the CRL cache on the client has expired is by clearing the CRL cache — a task that is difficult to perform in many networks.

# Planning for Data Recovery and Key Recovery

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If public key pairs and certificates are lost due to system failure, it can be time consuming and expensive to replace them and the data that they protect. For this reason, as part of your certificate management plan, you need to evaluate the potential consequences of loss of public keys and the data that they secure, and create a strategy for data and key recovery.

Data recovery is a process by which data is encrypted in such a way that more than one person can retrieve the data in plaintext form. Data recovery does not always imply that private key recovery has occurred; however, key recovery is one method for data recovery.

Use data recovery if you need to be able to recover data in your organization, but do not need to have access to individual private keys of users.

The advantages of data recovery include:

- It does not require a certification authority or PKI infrastructure.
- Data recovery policies can be managed centrally by means of Active Directory.
- Users do not have to manage certificates or private keys for data recovery.
- Decryption can be limited to the user alone.

The disadvantages of data recovery include:

- An administrative process must recover user data. Users cannot recover their own data.
- You cannot define the scope of what data can be recovered by a data recovery agent and what data cannot be recovered by a data recovery agent.
- Data recovery occurs manually on a file-by-file basis.
- Only data is recovered, and not the user keys. Therefore, after data recovery is completed, the user must re-enroll for new certificates.
- It is assumed that, when a key is lost, the certificate is compromised. Therefore, administrators must revoke old certificates.
- Stand-alone workstations or workstations in non-Active Directory environments cannot be centrally managed.

Key recovery allows a trusted agent to gain access to user private keys. For this reason, it is best to use key recovery only if your organization permits a person other than the original requester to have access to the private key of another user.

Use key recovery if organization policy permits the retrieval of user private keys and certificates. Key archiving and recovery implies that a person such as an administrator can gain access to the private keys of another user. Even when policies and procedures are in place to protect against unauthorized key recovery, issues with non-repudiation might still exist. If your organization does not permit a person other than the original requester to have access to the private keys of another user, do not implement key archival and recovery.

The advantages of key recovery include:

- Users do not have to re-enroll for certificates, change security settings, and so on.
- Existing certificates do not have to be revoked.
- Users do not have to recover any data or e-mail due to lost private keys.

- All data encrypted by means of a public key in a certificate can be recovered after a private key has been recovered.
- Windows Server 2003 does not accept signing keys for archival and recovery.

The disadvantages of key recovery include:

- User key recovery is a manual process that must be performed by administrators and users.
- Key recovery allows administrative access to the private keys of users.
- Non-repudiation assurance might not be available with key archival and recovery.
- Key recovery does not work with hardware security tokens such as smart cards.

**Note**

- Only a Windows Server 2003 Enterprise Edition CA can implement Windows Server 2003 key recovery.

Windows Server 2003 includes a new certificate template to support the key recovery agent role. A Windows Server 2003 CA can use only key recovery agent certificates that have been properly formatted and that have not expired. To enable key recovery, you need to complete the following tasks:

- Configure the key recovery agent template.
- Configure the CA to allow key archiving.
- Enroll and archive users.

Do not use either data recovery or key recovery if your organization wants to protect data from all parties except for the original user.

# Configuring the Key Recovery Agent Certificate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you configure a key recovery agent certificate, you must decide which users or groups can have Read and Enroll permissions on the key recovery agent certificate template. By default, only an Enterprise Administrator or a Domain Administrator can request a key recovery agent certificate. If you choose to change these defaults, you need to configure the new Read and Enroll permissions on the template itself.

You must configure an enterprise CA to issue key recovery agent certificates.

When you have configured permissions on the key recovery agent template and authorized an enterprise CA to issue key recovery agent certificates, a user with the appropriate permissions can request a key recovery agent certificate.

You must also select an encryption key length for the key recovery agent certificate. An encryption key of 2,048 bits satisfies most security needs. Keys that are 8,192 bits or larger can take the client CSP several hours to generate and can slow down public key operations on the CA when keys are archived.

You must mark the keys as exportable to enable the key recovery agent to export the private keys from the local store of the workstation to a floppy disk or other medium for safe storage. It is also best to protect the key recovery agent certificate private key with a strong password requirement. You can use a smart card as a key recovery agent.

The default key recovery agent certificate template requires manual approval of requests for key recovery agent certificates. It is best if a certificate manager manually approves all key recovery agent certificate requests. The certificate manager might choose to use fewer key recovery agents than the number of available key recovery agent certificates. In this way, no individual key recovery agent can decrypt all the private keys in the CA database. The CA chooses the key recovery agent certificate randomly as a means to ensure that the key recovery agent selection is not predictable.

Several cautions apply to key archiving. First, the default templates in Windows Server 2003 do not allow for key archiving. You must create new version 2 templates, which are available only in Windows Server 2003, Enterprise Edition, to support user enrollment with archiving.

Second, although you can configure the cryptographic service providers that are used for the private keys that are to be archived, you can only archive keys that are generated by means of a Rivest-Shamir-Adleman (RSA)-based CSP. The Digital Signature Standard (DSS) and Diffie-Hellman CSPs are not supported for key archiving.

# Establishing Key Recovery Agent Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Allowing someone other than the original user to recover keys presents a security risk. Although you trust your administrators, there are limits to how much any individual can be trusted with the ability to recover other the key pairs of other users. For example, your key recovery agent might leave the organization, taking a copy of the key. Therefore it is recommended that you monitor key recovery plans carefully.

Consider limiting the time that any one individual serves as the key recovery agent, or consider dividing the responsibility between several individuals and requiring that a smart card be used to perform key recovery tasks.

In addition, employ the following key recovery strategies:

- If you know that a key has been compromised, revoke it immediately.
- Do not recover keys or certificates that are used to secure high-value transactions or are associated with high-value certificates.
- Do not archive or recover private keys that are used for signing. This creates uncertainty in situations in which non-repudiation is the primary concern.

If possible, recover encryption keys only after the original certificates have been revoked. Issue a new key at the time of recovery. Revocation ensures that the user can still decrypt data with the old key but cannot encrypt new data.

# Educating Users About Planning for Data Recovery and Key Recovery

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although the process of certificate enrollment and use is nearly transparent to end users, it is important to educate users about certificates and their use. Specifically, be sure to provide end users with the following information:

- How to use the certificate enrollment user interface and certificate-enabled applications.
- The capabilities of certificates.
- The limitations of certificates.
- Inappropriate or ineffective uses for certificates.
- How to evaluate certificates received from others.
- The importance of retaining expired certificates.
- What to do in the event of an error message or if certificates fail to function as expected.

# Example: Creating a Certificate Management Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A PKI cannot be effective and secure unless an organization implements a management plan that includes strategies for enrolling and renewing certificates, mapping certificates to user accounts, revoking certificates and distributing CRLs, and using key recovery.

Many organizations base their certificate enrollment and renewal methods on the level of security associated with each type of certificate and the volume of certificate requests that they anticipate. For example, an organization makes the following decisions regarding certificate enrollment and renewal:

- Autoenrollment is the preferred enrollment method for e-mail and EFS certificate requests, which represent the majority of their certificate activity. Only clients who have already been authenticated by the network can request these certificates. The risks associated with the use of these certificates are relatively low.

- Manual approval is required for all certificates that are needed to perform network administration and software development.

- Manual approval is required for certificates that are issued to joint venture partners.

The basic user certificates of the organization (for e-mail and EFS) are distributed according to the domain membership of a user.

The distribution of high-security certificates is enforced with a one-to-one mapping. This is intended to further enforce the usage restrictions that have been placed on these certificates. Also, to improve the ability of the organization to define which file shares and other resources are available to their joint venture partners, a many-to-one mapping to a single account in Active Directory restricts their joint venture certificates.

Similarly, organizations are concerned about the timeliness of CRLs associated with their high-security certificates. Therefore, they decide to distribute CRLs for these CAs once a day, with delta CRLs published every two hours, or as needed. Because network bandwidth and replication can impact the distribution of CRLs and delta CRLs to their remote offices, they choose a less stringent publication schedule for their medium security CAs — new CRLs are published once a week, and delta CRLs are published at the close of every business day. Publishing at the end of the business day ensures that the updated information is replicated overnight and is available on the next business day.

# Deploying the PKI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After your public key design has been validated and refined by lab testing and pilot programs, you can deploy the PKI in your production environment. A disciplined approach to deploying a PKI is absolutely essential in order to establish the security of the applications that you are enabling. Figure 16.21 shows the PKI deployment process.

**Figure 16.21 Deploying the PKI**




**Important**

- Advanced planning is critical to the success of your PKI. Therefore, it is strongly recommended that you complete your PKI design before you deploy your PKI.

# Schedule Production Rollout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For large enterprise deployments, schedule the public key production rollout in stages. You can roll out different portions of the infrastructure at different times as necessary to support your security goals and business needs.

For example, you might begin by deploying EFS and IPSec because you do not have to establish a large CA hierarchy to take advantage of the security benefits of these features. You might place the next highest priority on secure mail and smart card authentication. You can schedule the rollout of the secure mail infrastructure before rolling out the smart card infrastructure, or you can choose to schedule the deployment of secure mail to one group or site and simultaneously roll out the smart card infrastructure to another group or site.

**Important**

- In many organizations, responsibility for the PKI deployment is transferred from a high-level design team to a more tactically focused implementation team at this time. If this is the case in your organization, be sure to provide the implementation team with all the documentation, including worksheets that you have created up to this point.

# Install Certification Authorities

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most organizations require a CA hierarchy to provide the required certificate services to meet their security and business needs. If you plan to use a CA hierarchy, you must install the root CA first and then install each subordinate CA in the hierarchy. For example, to create a three-level CA hierarchy and trust chain, install CAs on server computers in the following order:

1. Root CA
2. Intermediate CAs
3. Issuing CAs

For more information about installing CAs, see "Installing and configuring a certification authority" in Help and Support Center for Windows Server 2003.

# Install Offline Root CAs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In addition to the configuration options that you selected for your offline root CA, configure the following options when installing your offline root CA:

- Select **Certificate Services Web Enrollment Support**. Hosting the Web Enrollment service for an Offline Root CA on a separate system forces you to run both systems during the enrollment or renewal of subordinate CAs, which requires you to enable network connectivity between the two systems at that time.

- In the **Public and Private Key Pair** dialog box, leave the default **CSP** selection (Microsoft Strong Cryptographic Provider) and default **Hash** selection (SHA-1). Increase the **Key length** to meet your needs — for most root CAs, use the largest interoperable key length (4,096 bits).

  **Note**

  - For the purposes of Microsoft CAs, the Strong and Enhanced CSPs are considered equivalent. Both provide support for large key lengths (1024-bit keys or greater). Also, it is recommended that you use a hardware cryptographic service provider (CSP) or Hardware Security Module (HSM) to enhance the security of the signing keys of the certification authority.

- If you are installing a stand-alone CA as the root CA, the CA identification data must be entered manually. If you plan to publish the root CA certificate and CRL in your Active Directory environment, you have to enter the namespace of your Active Directory forest as the distinguished name suffix. In the **CA Identifying Information** dialog box, enter a customized distinguished name if you plan to publish your offline CA to a directory other than Active Directory. Use a customized distinguished name if you plan to use the offline CA as a trust anchor outside the enterprise.

  **Note**

  - The **Common name for this CA** fieldmust be filled in, but the customized field **distinguished name suffix** is optional. Your common name and distinguished name for the CA must reflect the organization and purpose of the CA to make the CAs easy for administrators and users to identify. The name of the CA must be unique within the organization, and possibly outside the organization as well. This information is filled in automatically if your CA is joined to an Active Directory–based domain.

- When you are asked to enter the Data Storage Locations, format the paths as local paths (such as C:\WINDOWS\System32\CertLog).

  **Note**

  - Although it is generally good practice to place the Certificate Database and Certificate Database log directories on a separate volume from the system partition, you do not need to do this for a root CA. The only data that is generated and must be stored concerns the certificates that correspond to a few subordinate CAs.

# Install Intermediate and Subordinate CAs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When installing Windows Server 2003 intermediate and issuing CAs, you can request the CA certificate from an online CA, or you can save the certificate request to a request file and make the certificate request offline. If you make an offline CA certificate request, the CA is not valid immediately upon installation. You must use the Certification Authority MMC snap-in to import the certificate of the CA and complete the installation process after the parent CA signs the CA certificate. You can also use the Certification Authority MMC snap-in to import subordinate CA certificates that are issued by third-party parent CAs.

# Publish the Offline CA Certificate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use secure procedures to publish the certificate and certificate revocation list (CRL) of the offline root CA. You only need to publish the certificate of the root CA one time. However, the CRL for the root CA must be published at regular intervals that correspond to the CRL publication interval value configured in the **Revoked Certificates Properties** of the root CA.

If the root CA is maintained in a secure location, such as a data center or vault, it is best if more than one administrator or trusted person publishes the offline CRL within that location, as prescribed in the certificate policy and certificate practice statements for your organization. After the CRL is published, you must transfer it manually from the data center or vault to a location where it can be distributed to your CRL distribution points.

Publish the offline CRL at least several days before the previously issued CRL is set to expire. This allows you to correct any hardware problems or publication failures in advance, ensuring that no interruption in service happens when your offline CRLs are published and replicated to all CDP locations.

After the offline root CA is installed, configure the various constraint and policy options for certificates that the offline CA issues. These extensions are necessary to ensure that the applications and clients that use the certificates in the hierarchy can perform revocation and chain building as needed.

# Apply CA Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you intend to implement your certificate practice statement, you need to create and format a issuer policy statement file, and place this file in the %windir% path of the root or subordinate CA before the CA is installed. This file, named CAPolicy.inf, serves two purposes:

- It provides basic information about the root CA, such as distribution points for the self-signed certificate, and the object identifier (also known as OID) information.
- It includes information for certificate renewal, such as the certificate lifetime of the self-signed certificate.

CAPolicy.inf is processed for root CA and subordinate CA installations and renewals. The CDP and AIA extensions in CAPolicy.inf are used for root CA installations and renewals only. Subordinate CA certificates inherit the CDP and AIA extensions of the issuing parent CA.

The CPS statement extension is applied when root CA certificates and subordinate CA certificates are requested. The CAPolicy.inf mechanism can only be used to include a CPS statement extension in a CA certificate and not an end-client certificate.

**Important**

- When CAPolicy.inf is used to install a CA, it must also be used for renewal; otherwise, the settings that have been defined might not be retained when the CA keys are renewed.

For more information about creating and using CAPolicy.inf files, see "Installing and configuring a certification authority" in Help and Support Center for Windows Server 2003.

# Configure CDP and AIA Extensions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After a root or subordinate CA is installed, you must configure the Authority Information Access (AIA) and CRL distribution point (CDP) extensions before the CA issues any certificates. The AIA extension specifies where to find up-to-date certificates for the CA. The CDP extension specifies where to find up-to-date CRLs that are signed by the CA. These extensions apply to all certificates that are issued by that CA.

Configuring these extensions ensures that this information is included in each certificate that the CA issues so that it is available to all clients. This ensures that PKI clients experience the least possible number of failures due to unverified certificate chains or certificate revocations that can result in unsuccessful VPN connections, failed smart card logons, or unverified e-mail signatures.

Follow these guidelines when configuring CDP extension URLs:

- Avoid publishing delta CRLs on offline root CAs. Because you do not revoke many certificates on an offline root CA, a delta CRL is probably not needed.

- Adjust the default LDAP:/// and HTTP:// URL locations on the **Extensions** tabof the certification authority **Properties** page according to your needs. Do not remove the local CDP location, however. The CA requires the local CDP location in order to publish the CRL to itself. The CA uses the local CRL to validate all certificates before they are issued to users. The local path does not show in the CDP extension of issued certificates.

- Enable the publication of delta CRLs, regardless of whether delta CRLs are going to be published, to allow for the potential use of delta CRLs in the future. Enable delta CRL publication by selecting the **Publish Delta CRLs to this location** check box.

- Publish both the LDAP and HTTP URLs for CDP locations to enable clients to retrieve CRL data with HTTP and LDAP. If required, publish a CRL on an HTTP Internet or extranet location so that users and applications outside the organization can perform certificate validation.

- Consider using Active Directory–based publication. An LDAP certificate revocation list URL distributed by means of Active Directory is replicated in a fault-tolerant, distributed, highly available manner. However, replication of CRL data among Active Directory domain controllers introduces some latency.

- For certificates that are to be validated by clients that use Active Directory, place the LDAP CDP location first in the list to optimize client revocation checking. Windows clients always retrieve the list of URLs in sequential order until a valid CRL is retrieved.

- Provide an additional HTTP CDP location or an alternative LDAP path to CRLs for clients that cannot use Active Directory or LDAP.

Follow these guidelines when publishing HTTP-based CRLs:

- If you are providing an HTTP CDP location, use round robin DNS or Web server virtual names to provide redundancy in the HTTP URL.

- Use HTTP CDP locations to provide accessible CRL locations for non-Windows brand operating system clients.

- Place HTTP CDP URLs second in the list of the URLs in the CDP extension if the CRL is distributed with Active Directory as well.

# Configure Certificate Templates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, Windows Server 2003 enterprise CAs are enabled upon installation to issue a variety of types of certificates. You can use the Certification Authority MMC snap-in to make the following modifications to this default configuration:

- Specify the certificate types that are to be issued by each CA.
- Delete any default certificate templates that you do not want the CA to issue from the certificate templates container.
- Add additional certificate templates that the CA can issue.

You can configure CAs to support one or multiple security functions by:

- Configuring root or intermediate CAs to issue subordinate certification authority certificates only.
- Configuring an issuing CA that supports secure Web communication services to issue authenticated session, computer, and Web server certificates only.
- Configuring an issuing CA that supports general business users to issue user certificates only, or configuring a CA that supports administrators to issue administrator certificates only.
- Configuring an issuing CA that supports smart card enrollment to issue smart card logon and smart card user certificates only.

The access control lists (ACLs) for each certificate template control the permissions needed to request certificate types. An enterprise CA grants certificate requests only for users, computers, or services that have the Enroll permission selected in the ACLs for that certificate template. The ACLs for certificate templates are preconfigured to enable various default user accounts and security groups to enroll for certificate types.

You can use the Certificate Templates MMC snap-in to modify the ACLs for each certificate template. For example, by default, only members of the Domain Administrators security group can request and obtain enrollment agent certificates. However, to specify that only certain members of your security department can request and obtain enrollment agent certificates, you can change the ACLs for the enrollment agent certificate template. You can remove domain admins from the ACL and add the appropriate user accounts or security groups.

For Windows Server 2003 stand-alone CAs, information about the certificate type must be included in the certificate request because stand-alone CAs do not use certificate templates. You can use stand-alone CAs with custom policy modules and custom certificate request applications to control the types of certificates that are issued.

# Configure Public Key Group Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the Group Policy MMC snap-in to configure the following optional categories of public key Group Policy for sites, domains, and organizational units:

- **EFS Recovery Agents.** By default, the local Administrator user account for the first domain controller installed in the domain is the EFS recovery account for that domain. You can specify alternate encrypted data recovery agents for EFS by importing into the policy the EFS recovery agent certificate for the appropriate alternate agent. You must first issue EFS recovery agent certificates to the user accounts on the local computers that you want to use as alternate recovery agents.

- **Automatic Certificate Enrollment.** You can specify automatic enrollment and renewal for computer certificates. When automatic enrollment is configured, the specified certificate types are issued to all computers within the scope of the public key Group Policy. Computer certificates issued by means of automatic enrollment are renewed from the issuing CA. Automatic enrollment does not function unless at least one enterprise CA is on line to process certificate requests.

In addition, you can use Group Policy to configure a number of CA trust options. Group Policy trust is configured and enforced within a single domain only. This allows different users in different domains to trust different root CAs.

When possible, manage trust of third-party root certification authorities by using Group Policy, and limit their scope by using qualified subordination. Third-party root CAs can be constrained by namespace and purpose to prevent unwanted trust and namespace violations within the organization.

Group Policy trust configuration is found in the computer policy for \Windows Settings\Security Settings\Public Key Policies\Trusted Root Certification Authorities. Users inherit policy from the computer policy. You can enable all computers and users to trust a root CA by adding the root CA certificate to Group Policy.

You can configure the following alternate trust options by selecting the **Trusted Root Certification Authorities** node in the Domain Security Policy MMC snap-in:

- **Enable or disable the ability for users to trust root CAs on a per-user basis.** Use this option to disable users from trusting a root CA outside the Enterprise root trust, Group Policy, the default computer store root CA list, and the list of root CA certificates provided by Windows Update.

- **Allow both Enterprise CA trust and third-party CA trust or only Enterprise root trust.** You can disable trust of third-party root CAs in the domain outside the enterprise root CA trust, including root certificates downloaded from the Windows Update. This disables user installation of root CA certificates.

**Note**

- Disabling third-party CAs can impact user access to applications such as SSL-secured Web sites.

# Configure CRL Publication

Updated: May 1, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you have updated the CDP extensions on the CA, you need to publish a new CRL so that all clients can access the new CRL data. For more information about configuring CRL publication, see "Manage Certificate Revocation" in Help and Support Center for Windows Server 2003.

**Modifying the Default Certificate Publication Period**
Certificates that are revoked prior to expiration remain in a published base CRL for one full base CRL period (defined by the CA) after they expire. Certificates that expire are no longer included in published CRLs after one additional base CRL expires.

Although applications do not check CRLs for certificates that have expired, you might in some cases want to maintain a public list of signing certificates that have been revoked. You can enable a registry setting on a CA to ensure that revoked certificates that have expired are not removed from the CRL.

**To modify the default CRL publication period for revoked and expired certificates on a CA**

- At the command prompt, type:

```
certutil -setreg ca\CRLFlags +CRLF_PUBLISH_EXPIRED_CERT_CRLS
```

| **Warning** |
| --- |
| Consider that large CRLs can cause performance issues for certificate clients and domain controllers. For more information on what could cause large CRLs, see TechNet Wiki article Large CRLs: What is added to a Certificate Revocation List (CRL)? |

**Ensuring Application Reliability**
Many applications rely on CRL availability and fail if the CRL is inaccessible or out-of-date. Follow these guidelines when publishing CRLs to ensure the reliability of your applications:

- Configure the CRL to be valid for a long enough period of time to allow for the recovery of the CA if there is a hardware or software failure.
- Set a reasonable CRL overlap period to protect against CRL publication or replication failures.
- Keep the private key of the CA and a copy of the CRL in a secure offline location so that you can sign and publish a valid CRL manually by using Certutil.exe when a catastrophic failure occurs.
- Use Active Directory to publish CRLs whenever possible. This maximizes availability and network performance. However, always consider the amount of time needed to replicate Active Directory data between domain controllers.
- Do not publish CRLs to Active Directory when the CRL publication period is shorter than the replication convergence time for the Active Directory forest.
- To prevent the use of a logon certificate, disable the account in Active Directory.

**Controlling CRL Size**
You can partition a base CRL to control its size. In this way, you can control the amount of data that is replicated to Active Directory and the size of the data object that clients download when they perform revocation checks on certificates. You partition base CRLs by renewing the CA key. This creates a partitioned CRL for all certificates that are issued after the key is renewed.

The CRL increases by about 29 bytes for every certificate that is revoked, depending on the reasons that you specify for the revocation. You might want to use a new key to renew the CA every time it reaches 100-125 kilobytes (KB) in size, to minimize download times. This strategy is based on the assumption that approximately 10 percent of the certificates that you issue are revoked before their natural expiration date. If your actual or planned revocation rate is higher or lower than this, adjust your

key renewal strategy as needed.

**Removing Expired CRLs**
By default, a CA maintains an expired CRL in the database and keeps it in the directory at the last known CDP publication point.

When the key for a CA expires, the CA continues to publish CRLs that are signed by the expired CA certificate. It is recommended that you maintain this CRL in the CA database to allow for long-term validation and auditing. You can, however, remove the CRL to clean out the database.

**To remove a CRL after a CA key expires**

- At the command prompt, type:

```
certutil -setreg ca\CRLFlags + CRLF_DELETE_EXPIRED_CRLS
```

# Delegate CA Administration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin to issue certificates, you need to delegate CA administrator roles. For information about defining CA administrator roles, see "Defining PKI Management and Delegation" earlier in this chapter.

You can assign CA Administrator and Certificate Manager permissions by using the Certification Authority MMC snap-in. Other roles, users, and groups are specified in the consoles used to perform particular tasks, such as Certificate Templates and Certificates. To change the roles of a user, you must change the security permissions, group membership, or rights of the user.

# Configure Certificate Enrollment and Renewal

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft Certificate Services supports a variety of enrollment and renewal methods, such as the Certificate Request Wizard or the Microsoft Certificate Services Web pages. However, if you deploy third-party certificate services or custom certificate enrollment and renewal applications, you must perform any configuration required for those services and applications.

# Issue Certificates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can issue certificates to users, computers, and services after the required certificate services are installed and configured. Keep the following considerations in mind when you start to issue certificates:

- Certificates are issued for computers within the scope of the Automatic Certificate Request settings of the Group Policy. Domain Administrators can also manually request certificates for local computers by using the Certificate Request Wizard or the Microsoft Certificate Services Web pages. Consider scheduling manual enrollment in stages to help distribute the administrative workload for computer enrollment.

- Smart card administrators can start issuing smart card certificates by using the Smart Card Enrollment Station available on the Microsoft Certificate Services Web pages. Consider scheduling smart card enrollment in stages to help distribute the administrative workload for smart card enrollment.

During the transition to smart cards, both smart card authentication and interactive logon with domain credentials should be enabled. Because this weakens network security, configure user account policy to require smart cards for interactive logon as soon as smart card users are trained and are using their cards.

Monitor the performance of certificate services closely as you start issuing certificates to ensure that CAs handle the certificate load. To correct excessive load conditions, consider adding more issuing CAs or scheduling certificate enrollment in smaller stages. Certificate renewal might also produce excessive load conditions. Adding more CAs and scheduling certificate enrollment in smaller stages can also help distribute peak renewal loads.

# Additional Resources for Designing a Public Key Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- The *Distributed Services Guide* of the *Windows Server 2003 Resource Kit* (or see the Distributed Services Guide on the Web at http://www.microsoft.com/reskit) for more information about public key features in Windows Server 2003 and using certificates in conjunction with Encrypting File System.
- "Planning a Smart Card Deployment" in this book for more information about deploying smart cards.
- The Software Development Kit (SDK) information in the MSDN Library link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about developing CAPICOM and other public key–enabled applications.
- The Common Criteria link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about the Common Criteria for Information Technology Security Evaluation.
- The National Institute of Standards and Technology (NIST) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about public standards that apply to public key infrastructures.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Mapping certificates to user accounts" in Help and Support Center for Windows Server 2003 for information about creating a mapping.
- "Manage Certificate Revocation" in Help and Support Center for Windows Server 2003 for information about configuring CRL publication.
- "Installing and configuring a certification authority" in Help and Support Center for Windows Server 2003 for more information about creating and using CAPolicy.inf files.

**Related Job Aids**

- "Summary of User Certificate Requirements" (DSSPKI_1.doc**)** on the *Windows Server 2003 Deployment Kit* companion CD (or see "Summary of User Certificate Requirements" on the Web at http://www.microsoft.com/reskit).
- "Certificate Practice Statement Outline" (DSSPKI_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Certificate Practice Statement Outline" on the Web at http://www.microsoft.com/reskit).
- "Windows Server 2003 Certificate Lifecycle Plan" (DSSPKI_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Windows Server 2003 Certificate Lifecycle Plan" on the Web at http://www.microsoft.com/reskit).

# Planning a Smart Card Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Smart card support in Microsoft® Windows® Server 2003 enables you to enhance the security of many critical functions, including client authentication, interactive logon, and document signing, in your organization. If you are using or planning to use public key certificates, deploy smart cards to increase security for your network and critical applications.

**In This Chapter**


Overview of Smart Card Deployment

Creating a Plan for Smart Card Use

Selecting Smart Card Hardware

Creating a Smart Card Deployment Plan

Planning for Ongoing Smart Card Support

Additional Resources for Smart Card Deployment


**Related Information**

- For more information about creating a public key infrastructure, see "Designing a Public Key Infrastructure" in this book.
- For more information about Windows Server 2003 Certificate Services and public key infrastructure features, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Overview of Smart Card Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most organizations use passwords to manage access to computer networks and resources. However, some users set weak passwords, write passwords down in insecure locations, or forget their passwords and require help desk assistance for password reset. For this reason, passwords alone might not provide the level of security and manageability that your organization requires.

Smart card support in Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems provides users with stronger credentials than even the most complex passwords. If you use, manage, and deploy smart cards properly, you can enhance the security of your organization and reduce your support costs.

Smart cards offer the following benefits:

- **Protection.** Smart cards provide tamper-resistant storage for private keys and other data. If a smart card is lost or stolen, it is difficult for anyone except the intended user to use the credentials that it stores.

- **Isolation.** Cryptographic operations are performed on the smart card itself rather than on the client or on a network server. This isolates security-sensitive data and processes from other parts of the system.

- **Portability.** Credentials and other private information stored on smart cards can easily be transported between computers at work, home, or other remote locations.

The number and variety of smart card–enabled applications is growing to meet the needs of organizations that want to rely on smart cards to enable secure authentication and to facilitate services.

Before you can deploy smart cards in your organization, you must have a public key infrastructure (PKI) in place. Next, you need to identify applications to enable for use with smart cards, and plan how to implement and support a smart card infrastructure before you can take advantage of the security benefits of smart cards.

**Note**

- For a list of the job aids that are available to assist you in deploying smart cards, see "Additional Resources for Smart Card Deployment" later in this chapter.

# Process for Planning a Smart Card Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Planning a smart card deployment involves making decisions about technical standards, hardware purchases, smart card management, and the logistics of smart card distribution.

Figure 16.1 shows the process for planning a smart card deployment.

**Figure 16.1 Planning a Smart Card Deployment**



# Smart Card Fundamentals

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 supports a variety of secure smart card applications and business scenarios. Before you begin to plan your smart card deployment, it is important to understand the basic components of smart card technology.

**Components of a Smart Card Infrastructure**

A number of hardware and software components are required in order to support a smart card infrastructure.

**Certificates** Digital data that securely bind a public key to the entity that holds the corresponding private key.

**Certification authorities** Trusted entities or services that issue digital certificates.

**Active Directory** The Windows Server 2003 directory service that serves as a repository for account information, primarily user credentials, security group memberships, and certificate templates. In addition, you can also use the Active Directory® directory service to store certificates, certificate revocation lists, and delta certificate revocation lists, and to publish root certification authorities (CAs) and cross-certificates.

**Smart cards** Hardware tokens containing integrated processors and memory chips that can be used to store certificates and private keys and to perform public key cryptography operations, such as authentication, digital signing, and key exchange.

**Smart card readers** Devices that connect a smart card to a computer. Smart card readers can also be used to write certificates to the smart card.

**Smart card software** The software provided by the smart card vendor to manage smart cards. In some cases, organizations might choose to create their own software tools if customized functionality is required.

# Creating a Plan for Smart Card Use

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before deploying smart cards in your organization, you must determine which processes, users, and groups of users require smart cards.

Figure 16.2 shows the process for creating a plan for smart card use in your organization.

**Figure 16.2 Creating a Plan for Smart Card Use**



# Identifying the Processes That Require Smart Cards

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A smart card deployment can help your organization meet numerous sensitive business requirements. You can use smart cards for any or all of the following processes:

- Interactive user logons, including remote access connections to the network
- Administrator logons
- Third-party authentication across the Internet
- Signing and encrypting e-mail

Evaluate additional equipment and administrative costs, procedures, and changes to user work patterns that each smart card–enabled process requires. Ensure that the benefits of deploying smart cards for each process outweigh the costs from hardware, administration, and potential user difficulties.

For a worksheet to assist you in documenting the processes in your organization that require smart cards, see "User and Group Smart Card Requirements" (DSSSMC_1.doc) on the Microsoft® *Windows® Server 2003 Deployment Kit* companion CD (or see "User and Group Smart Card Requirements" on the Web at http://www.microsoft.com/reskit).

# Interactive User Logons

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use smart cards for an interactive user logons if you want to enforce the use of secure encrypted logon credentials. If you require users to log on by using smart cards, you do not have to worry about the quality and security of user passwords.

Requiring smart cards for interactive user logons requires additional network administration for smart card distribution and support. This is problematic for organizations that are spread across different geographic locations and that do not have network or physical security personnel in each location to administer and support smart cards.

You can also use smart cards for remote access logons, and for Terminal Services and shared client logons.

**Remote Access Logons**
Local interactive logons require that users have both physical access to a computer that is a logical member of the organization and a network password. Remote users, however, can log on from any computer outside of the organization. If a malicious user obtains the password of a remote user, he or she can use it to access network resources from any computer. For this reason, conventional password-based remote access logons are more vulnerable to attack than local interactive logons.

You can secure the remote access process by requiring users to use smart cards when they connect to the corporate network by means of remote access logon. This solution prevents attackers from using the remote access dial-up or Internet connections to compromise the network, even if they have physical access to laptops or home computers.

One problem with requiring the use of smart cards for remote access logons is the fact that remote users often own computer hardware and software that does not conform to minimum corporate standards and, therefore, might not support smart card use. This complicates the process of administering and supporting smart cards for remote access logons. Also, users might experience longer logon times when they use smart cards, especially over slow dial-up connections.

**Terminal Services and Shared Clients**
If your organization is deploying Terminal Services, consider using smart cards for kiosk computers that are shared by multiple users. This can improve security in environments in which multiple users share a single computer terminal, relocate frequently, and do not use the conventional logoff procedure every time they move away from the terminal. This is often the case in hospitals, factories, or other businesses.

**Note**

- Smart card logons require Microsoft® Windows® XP or Windows Server 2003 Terminal Services clients, even on computers running Microsoft® Windows® 2000.

Providing smart card support for kiosks or Terminal Services clients that are in critical locations in your organization and are shared by several users is less costly than providing smart card support for interactive user logons, because you do not need to purchase and deploy a large number of smart card readers.

For more information about deploying Windows Server 2003 Terminal Services, see "Hosting Applications with Terminal Server" in *Planning Server Deployments* in this kit.

# Administrator Logons

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There is greater potential for harm to the network when administrator credentials, as opposed to user credentials, are misused. As a result, preventing unauthorized users from using administrative credentials to access their network is an important security priority for most organizations. Another vulnerability is introduced when you allow people to perform network administration tasks by using generic administrator accounts that are shared by multiple users; this limits the ability of the organization to track which user performs a specific action. Allowing administrators to log on by using administrative credentials when they are not performing administrative tasks also creates a significant security risk because attackers who compromise an administrator account can do a greater amount of damage to the system.

By requiring individuals to use smart cards to perform administrative tasks, you can significantly reduce the possibility that unauthorized users can gain administrative access to your network. You can use smart cards for administrator logons in the following two ways:

- By using smart cards for individual administrative operations.
- By using smart cards for an administrative shell.

In most cases, the best solution is to use a combination of these two strategies. For example, you can require that all administrators use smart cards to access data center servers. If the administrator is using a Windows 2000 or Windows XP client, he or she can use a smart card and administrative credentials to open a Terminal Services client session in order to log on to the data center servers.

**Important**

- It is not possible to utilize multiple credentials stored on a single smart card. Therefore, administrators who have more than one domain account require a smart card for each account.

**Using Smart Cards for Individual Administrative Operations**

When you use smart cards for individual administrative operations, administrators log on by using their standard user credentials, and then use administrative credentials when they need to perform specific administrative operations. For example, you might require an administrator to log on by using a smart card in order to install Active Directory on a member server. Administrative credentials apply only to the specific operation, which helps to protect the security of the system.

An administrator can also use smart cards to perform individual administrative operations on target computers running versions of the Windows operating system earlier than Windows XP or Windows Server 2003, as long as they use a smart card to log on to a computer running Windows XP or Windows Server 2003.

Not all administrative tools work with smart cards. Therefore, before you implement this solution, test it to ensure that you can perform the required administrative tasks and use the necessary administrative tools. If some of your required tools and tasks are incompatible with using smart cards, you must communicate to your administrators which tasks require smart cards and which must be completed by using administrative credentials.

**Using Smart Cards for an Administrative Shell**

When you use smart cards for an administrative shell, administrators log on by using user credentials. Then, when the administrator needs to perform administrative operations, he or she logs on by using a smart card and administrative credentials to open a Terminal Services client session. The administrator then performs the required administrative operations within the administrative shell.

This approach simplifies the process of performing multiple sequential administrative operations during a single session. However, the server that has Terminal Server enabled must be running Windows Server 2003. Although the Windows XP Terminal Services client can run on Windows 2000, the server-side support is only provided by Windows Server 2003.

# Authenticating Third Parties

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use smart cards for third-party authentication if you want to verify that queries, orders, or other communications originate from the appropriate individual or organization and that they conform to preestablished standards, such as purchase order limits. For example, banks that allow users to check their transaction histories or pay bills online, and distributors that accept purchase orders over the Internet can benefit from using smart cards for third-party authentication.

Deploying smart cards to third parties, however, requires careful administration. For example, you must ensure that attackers cannot obtain smart cards and guess the PIN to gain unauthorized access to the system. Also, if the customer services that are based on smart card authentication are an important part of your business, you need to ensure that the services are always available. If you do not administer your third-party smart card authentication process effectively, it can have a negative impact on your Internet business transactions.

# Signing and Encrypting E-mail

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use smart cards to enable digital signing and the encryption of electronic communications such as e-mails or contracts.

If you choose to deploy smart cards for digital signing, you need to determine the types of e-mail messages that require smart card–validated digital signatures. Use smart cards for the digital signing of e-mail messages where it is important to verify the identity of the sender and that the message has not been tampered with while in transit. Digitally signing routine e-mails creates unnecessary network traffic and can slow down ordinary communication between users. Note that when you use smart cards for the digital signing of sensitive documents, such as legal contracts or purchase orders, you must configure the certificate policies and extensions that control smart card certificate use.

Depending on the types of documents that you want users to sign digitally, you also need to make additional decisions about smart card–enabled digital signatures, such as whether assistants are allowed to sign documents on behalf of their superiors, whether send and read receipts are required, and how the receipts are to be stored.

For more information about certificates and certificate use, see "Designing a Public Key Infrastructure" in this book.

**Note**

- You must ensure that users know how to verify digital signatures. Unlike a hand-written signature, a digital signature is not embedded in a message or document, and might be overlooked.

# Defining Smart Card Service Level Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy smart cards, establish service level agreements to help your IT organization align smart card performance with the objectives of the organization in areas such as reliability, response times, and support procedures.

For example, you need to define smart card service level standards for:

- **The types of identification required to obtain a smart card.** You might choose to require a specific type of personal identification, such as a driver's license or other photo ID, in order for a user to obtain a smart card.

- **Unique service guarantees for special classes of employees, such as executives or roaming employees.** Define whether certain classes of employees are permitted to operate under support agreements that differ from those of other users.

- **Acceptable time needed for users to log on.** It is best to ensure that the different steps and time needed for smart card logon time are comparable to the steps and time needed for conventional password logons.

- **Acceptable logon times for remote access users.** Remote access logon times are more vulnerable to slowdowns than local network connections, especially if users have slow dial-up access connections. You might need to upgrade your remote access configuration in order ensure acceptable logon times for remote users.

- **Remote access exceptions.** The computer configurations of some users might not be compatible with smart cards, and remote users might lose or forget their smart cards. Identify the circumstances, if any, in which remote users are allowed to use remote access without using a smart card.

- **Number of unsuccessful PIN entries allowed.** Do not allow an unlimited number of attempts to enter a PIN. Allowing three or four attempts is generally adequate.

- **PIN reset requirements.** Decide whether users are allowed to reset their own PINs, or whether they need to provide personal identification to security or help desk personnel to have their PINs reset. If you decide that users need to provide positive identification, decide whether the user must present the identification in person, such as a photo ID, or demonstrate knowledge of a predefined secret, such as a mother's maiden name.

- **Service guarantees to users who cannot use their smart cards because of loss, damage, or blocking.** This includes:

  - Establishing when and how users can regain access to the network.

  - Determining whether to restrict these users' access to the network to certain areas, or to allow them access to any areas of the network that were previously accessible to them.

  Defining these limits helps you to establish user expectations and support procedures.

Document your service level standards. You will need to apply these standards in your smart card operations plan, test them in your lab and pilot deployments, communicate them to help desk personnel and to your users, and include them in your support and maintenance plan.

For a worksheet to assist you in documenting your service level agreement, see "Smart Card Service Level Agreement" (DSSSMC_2 .doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Smart Card Service Level Agreement" on the Web at http://www.microsoft.com/reskit).

**Important**

- Incorporate your smart card service level agreements in the Certificate Practice and Policy Statements for your public key infrastructure. For more information about creating Certificate Practice and Policy Statements, see "Designing a Public

Key Infrastructure" in this book.

# Selecting Smart Card Hardware

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Single smart cards and smart card readers are relatively inexpensive. However, when you deploy smart cards and smart card readers to hundreds or even thousands of users, equipment cost becomes an important consideration. You must evaluate smart card hardware in order to select the devices that best meet the needs of your organization at the best price.

Figure 16.3 shows the process for selecting smart card hardware.

**Figure 16.3 Selecting Smart Card Hardware**



# Creating a Smart Card Specification

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A wide variety of smart cards and smart card readers are available to choose from. Windows Server 2003 is designed to work with any cryptographic smart card that has an associated CryptoAPI cryptographic service provider. The physical characteristics of smart cards and readers are governed by published standards. Cards from any manufacturer that adheres to the ISO 7816 standard will likely be compatible with the reader you select. Be sure, however, to test smart cards and smart card readers to verify compatibility before deploying them in your production environment.

For more information about testing smart cards and smart card readers, see "Evaluating Smart Cards and Readers" later in this chapter.

**Note**

- For more information about ISO 7816, see the Smart Card Alliance link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Because smart cards both store *and* process data, it is important to create a specification for your smart cards. Creating a smart card specification involves making decisions about the following:

- Smart card hardware type
- Amount of memory required
- Intended useful smart card lifetime
- Intended smart card roles
- Smart card reader hardware
- Smart card management software

Table 16.1 lists some of the critical specifications that you need to define when you create your smart card specification.

**Table 16.1 Smart Card Hardware Specifications**

| Specification | Description |
|---|---|
| Memory | How much data you need to store on the smart card. |
| Life expectancy | The useful lifetime of the smart card. |
| Reuse | Whether or not the smart card can be configured for use by a second user, if the original user leaves the organization. |
| Type of card | The type of card that is most appropriate for your organization. You might choose one or more of the following:<br>• Credit card or token style<br>• Single purpose or dual purpose |
| Card dimensions | The size, length, and thickness of the card, depending on the type of card that you specify. |

| | |
|---|---|
| Number of cards | How many cards you need.<br>• If you have more users than computers, you need fewer readers than smart cards.<br>• If you use your smart cards on multiple systems, you need more readers than smart cards.<br>• If you specify more than one type of card, indicate the number of each type. |
| Type of smart card reader | The type of reader that is most appropriate for your organization. Options include:<br>• USB<br>• PCMCIA<br>• Serial |
| Number of smart card readers | How many readers you need.<br>• If you have more users than computers, you need fewer readers than smart cards.<br>• If you use your smart cards on multiple systems, you need more readers than smart cards.<br>• If administrators use one smart card for user logons and a second smart card for logging on with their administrative credentials, this will also impact the number of smart card readers that you require. |
| Performance requirements | The type of performance that you can expect. This includes:<br>• Minimum acceptable logon times for direct network logons.<br>• Minimum acceptable logon times for remote access logons.<br>• Ability to handle alternate credentials.<br>• Ability to restrict logons by using alternate credentials. |
| Smart card management tools | The types and quality of the tools provided by the hardware vendor to manage smart cards. |

For a worksheet to assist you in preparing a product specification, see "Smart Card Hardware Specification" (DSSSMC_3 .doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Smart Card Hardware Specification" on the Web at http://www.microsoft.com/reskit).

In the beginning of your deployment, you can meet user needs by using a single type of smart card with a single configuration option. However, as you expand your smart card infrastructure, you might need to deploy a variety of smart card types and configuration options.

**Smart Card Type**
Two types of smart cards are available for use with Windows Server 2003 and Windows XP: conventional credit card–shaped contact cards and smaller token-style cards that plug directly into the USB port of a computer.

**Note**

- Another type of smart card, called a contactless smart card, is not supported by Windows XP or Windows Server 2003.

**Credit card–shaped contact cards**
Credit card–shaped smart cards are available in three-volt and five-volt versions. They are the most common smart card solution, in part because they resemble the corporate card keys or badges that many organizations use.

### Note

- You can specify that your smart cards be screen-printed with your corporate logo and a picture of the user. If you plan to add graphics to smart cards, ask your vendor about the methods available for bulk printing and customizing cards.

If your organization uses card keys or badges, you can apply smart card chips to the existing card key or badge as a sticker or "skin." However, your card keys or badges need to fit into your smart card readers with a minimal amount of friction; therefore, be sure to include the physical thickness of the smart card in your specifications. This is an important factor to consider when you select a vendor to manufacture the stickers, as the material thickness for smart card chips can vary.

**Token-style smart cards**
Token-style smart cards are typically the size of a house key or automobile key. They plug directly into a USB port, providing a more compact solution than separate cards and readers. Token-style smart cards are ideal for laptop users who want to carry a minimum number of peripherals, or for workers who use a number of different computers. However, you cannot use token-style smart cards if your computers do not have USB connections, or if the USB connections are full or difficult to access.

**Memory**
Your smart card requires enough memory to store the certificate of the user, the smart card operating system, and additional applications. Smart cards run embedded operating systems, and in many cases, a form of file system in which data can be stored. To enable Windows smart card logon, you must be able to program the card to store a user's key pair, retrieve and store an associated public key certificate, and perform public and private key operations on behalf of the user.

To calculate the amount of memory that you need, determine the space requirements for:

- User certificates. A certificate typically requires about 1.5 kilobytes (KB). A smart card logon certificate with a 1,024-bit key typically requires 2.5 KB of space.
- The smart card operating system. The Windows for Smart Cards operating system requires about 15 KB.
- Applications required by the smart card vendor. A small application requires between 2 KB and 5 KB.
- Your custom applications.
- Future applications.

Figure 16.4 shows the additional space requirements of a typical 32 KB smart card. The smart card operating system requires about 15 KB, leaving 17 KB for the file system, which includes space for the card management software, the certificate, and any other custom applications.

**Figure 16.4 Memory Use on a 32 KB Smart Card**




It is possible to configure smart card file systems into public and private spaces. For example, you can define segregated areas for protected information, such as certificates, e-purses, and entire operating systems, and mark this data as Read Only to ensure the security of the smart card and restrict the amount of data that can be modified. In addition, some vendors provide cards with sub-states, such as Add Only, which is useful for organizations that want to restrict the ability of a user to revise an existing credential, and Update Only, which is useful for organizations that want to restrict ability of a user to add new

credentials to a card.

The data capacity available on smart cards is increasing as smart card technology improves. However, storage space on smart cards is expensive. Card vendors often restrict the amount of storage available to individual applications so that multiple applications or services can be stored on the card. Therefore, in your vendor specification, define all of your anticipated present and future card usage requirements and the memory requirements for each certificate and application that you require. If you plan to use your smart cards for multiple purposes, such as physical access to facilities and user logon, or to store additional data, you must increase your memory requirements. Also, when planning storage space on the chip, allocate space for applications that you are planning for future implementation.

**Note**

- Windows Server 2003 and Windows XP do not support the use of multiple certificates on a smart card.

#### Life Expectancy

You must define the length of time for which you will use a smart card before you replace or upgrade it. Contact your vendor for information about smart card life expectancy based on normal wear and tear.

In addition, you must take into account your current and future space requirements, including the anticipated need for additional applications and certificates with larger keys. Anticipate adding new applications, and potentially issuing new smart cards, over an 18-24 month card lifecycle. In the future, vendors are likely to introduce smart cards with more memory and other enhancements for a lower cost.

Also, determine whether you want your smart cards to be reusable in the event that users leave the organization. Reusing smart cards reduces the costs associated with issuing new ones. However, the cost associated with removing existing data and writing new data and applications is often equal to or more than the cost of preparing and issuing new smart cards.

#### Smart Card Roles

You can use smart cards for one of three roles. Determine how many smart cards you need to issue for each of the following roles:

- **Enrollment card.** Issue enrollment cards to individuals who enroll smart cards on behalf of other users. Enrollment cards have a special enrollment agent certificate. Issue the smallest possible number of enrollment cards that will enable you to enroll all required smart card users. This protects the security of your system.

- **User cards.** These are the standard cards that you issue to each user. Two types of user cards are available:

    - **Permanent.** Permanent user cards are cards that employees carry with them. They contain the cardholders' credentials, certificates, data, and applications. They might also have a photograph or a decal applied to the card. In a Windows Server 2003 environment, the permanent card points to a permanent certificate server.

    - **Temporary.** Temporary cards are a limited-use cards that are issued to guests, temporary employees, and users who have forgotten their permanent cards. They point to a temporary certificate server and can have a limited lifetime.

For a worksheet to assist you in documenting the roles for the smart cards that you issue, see "Smart Card Hardware Specification" (DSSSMC_3 .doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Smart Card Hardware Specification" on the Web at http://www.microsoft.com/reskit).

**Important**

- To ensure system security, issue master and enrollment cards to the smallest possible number of trusted employees. For more information about issuing enrollment agent cards, see "Establishing Enrollment Agents" later in this chapter.

#### Smart Card Readers

A variety of types of smart card readers are available. The majority of smart card readers connect to the computer through an RS-232 serial port, a Type II Personal Computer Memory Card International Association (PCMCIA) slot, or a universal serial bus (USB) port.

Although USB-compatible smart readers are the simplest type of reader to connect, the USB ports on some computers might already be used. For this reason, it is best to order a mix of card reader connector types based on the types of connections that

are available on your systems.

For a list of Windows-compatible smart card readers, see the Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Smart Card Management Tools**

You can perform most smart card–related tasks by using the Windows Server 2003 Certificate Services and software tools provided by the smart card vendor. However, it is important to assess the smart card tools that are available to determine whether they are sufficient to meet your needs. You might need to create additional tools for some smart card tasks.

For example, you might require tools to assist you in moving from a limited pilot phase to a full production deployment. Also, developers in your organization might need to create a direct interface between the smart card certificate and your building access systems. You might also choose to write a script that automatically enters critical data into a database when a smart card is created. This includes data such as smart card serial numbers, the names or e-mail names of the users who are assigned smart cards, the types of certificates that are issued to the users, when the certificates are issued, and when they expire.

For more information about creating scripts for Windows Server 2003, see the at http://www.microsoft.com/reskit, or see the TechNet Script Center link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Evaluating Smart Cards and Readers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to evaluate your prospective smart cards and readers throughout your smart card deployment process. Initially, obtain and evaluate a variety of smart cards and smart card readers to determine which vendors provide the best balance of specifications, performance, and price. As you deploy your smart card infrastructure, continue to evaluate your hardware to make sure that it performs as expected.

The smart cards and smart card readers that you deploy and the smart card production processes that you develop are likely be used many times every day. Therefore, you must ensure that your hardware is reliable. The service level agreements that you created when you defined your smart card requirements provide objective standards for measuring and documenting satisfactory performance.

To minimize user dissatisfaction and maximize manageability, be sure to test the following:

- **Installation and removal of the smart card software.** Make sure the smart cards work after you install the software. If the installation is faulty, use the Windows Event Viewer to access error messages that might explain the cause of the failure.

- **Fit of smart cards in readers.** Smart card dimensions, such as thickness, are governed by international standards. However, some organizations have found that, if the card-to-reader interface is too tight or abrasive, the cards deteriorate more rapidly.

- **Reader reliability.** To test reliability, create an environment that includes systems that have slower CPUs and less memory than computers in your organization. Test how well your smart card readers operate in this environment, as well as in other configurations. You can, for example, run a number of memory-intensive applications or use the smart cards and readers over slow connections to evaluate how each combination of smart cards and readers functions in these conditions. Your smart card service level agreements provide objective criteria for acceptable and unacceptable performance.

- **Card production.** Slow card production processes can impede your deployment. If your organization is unable to produce cards efficiently, use a third-party vendor to produce smart cards.

- **Ability to deploy multiple types of cards and readers.** If you are unable to efficiently deploy the types of cards, readers, and servers that you require, your service might be inconsistent and inefficient.

For a worksheet to assist you in documenting the results of your smart card reader evaluation, see "Smart Card Reader Evaluation" (DSSSMC_4 .doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Smart Card Reader Evaluation" on the Web at http://www.microsoft.com/reskit).

Figure 16.5 shows an example of a completed Smart Card Reader Evaluation worksheet.

**Figure 16.5 Example of a Smart Card Reader Evaluation Worksheet**

| Smart Card Reader Evaluation | | | |
|---|---|---|---|
| Prepared by | | Date | |

| Criteria | Manufacturer 1 | Manufacturer 2 | Manufacturer 3 |
|---|---|---|---|
| Cost | $15 | $22 | $18 |
| Usability | Good | Fair | Poor |
| Card Fit | Good | Good | Fair |
| Software Included in Windows Server 2003 | | | |
| PCMCIA Version | Yes | Yes | Yes |
| USB Version | Yes | Yes | Yes |
| Serial Version | No | Yes | Yes |
| Relationship | Good | No data | No data |
| Future Tests | Need to test more portable computers and typical home office systems | Need to test more portable computers and typical home office systems | Need to test more portable computers and typical home office systems |

# Creating a Smart Card Deployment Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying a smart card infrastructure is a time-consuming process because it involves deploying physical components (smart cards and smart card readers) and issuing digital certificates individually to every user who requires a smart card. Careful planning can significantly reduce the amount of time this process takes and enable you to enhance the security of your organization. Figure 16.6 shows the steps involved in creating a smart card deployment plan.

**Figure 16.6 Creating a Smart Card Deployment Plan**



# Establishing Certification Authorities

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to ensure that your public key infrastructure can support the issuance and verification of smart card certificates for the users and applications that you have identified. To ensure that your PKI can support a smart card infrastructure, you must do the following

- Configure your certification authorities (CAs) as enterprise CAs. Windows Server 2003 smart card certificates require enterprise CAs.

  **Important**

  - CAs that issue smart card certificates need to be trusted in the CA hierarchy and must be continuously online while the user is enrolled.

- Make sure that your issuing CAs are installed on servers that have enough storage and central processing power to support the smart card users in your organization.

For more information about planning your CA infrastructure, see "Designing a Public Key Infrastructure" in this book.

# Planning Smart Card Certificate Templates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use any of the following types of Windows Server 2003 certificate templates to enable smart card use in the Windows Server 2003 PKI:

- **Enrollment Agent**. Allows an authorized user to serve as a certificate request agent on behalf of other users.
- **Smart Card User**. Enables a user to log on and sign e-mail.
- **SmartCardLogon**. Enables a user to log on by using a smart card.

You can also create your own certificate templates to serve multiple purposes. For example, the smart card logon certificate template is designed for smart card logon only. If you intend to use your smart card infrastructure to support multiple applications, you can choose multipurpose templates instead. Multipurpose templates generate certificates that you can use for multiple applications, such as smart card logon and e-mail signing.

**Note**

- Windows 2000 only supports version 1 templates, which cannot be customized or extended. Use Windows Server 2003, Enterprise Edition, which supports version 2 templates, if you need to create new certificate templates, copy an existing template, or replace templates that are already in use.

As part of your planning for smart card certificate templates, you need to establish values for public keys, certificate lifetimes, and certificate renewal policies. These values are interrelated. For example, if you select a larger key value, you can implement a longer certificate lifetime. Or, you can use a small public key value if a certificate has a relatively short lifetime. Note, however, that the amount of memory that is available on the smart cards that you select also limits the size of the public keys that you can use.

**Important**

- Many organizations pre-enroll users for smart card certificates several weeks before they distribute smart cards to users. The certificate lifetime is determined by the date that you issue the certificate, not the date that you distribute the card to the user. Therefore, factor any distribution delays into your certificate lifetime and renewal strategy.

A Windows Server 2003 CA allows you to select a certificate public key length from 384 bits for minimal security to 16,384 bits for maximum security. For typical logon applications, a 1,024-bit key is adequate.

You can establish certificate lifetimes that are as long or as short as you need, and you can configure certificates to be nonrenewable, renewable a finite number of times, or renewable indefinitely.

To define public key values and certificate lifetimes and renewal policies, take into account:

- **The physical capacity of your smart cards**. Most of the smart cards that are available today have adequate space for all but the largest certificates.
- **How you define acceptable logon times**. Public key–based authentication often takes longer than authentication without certificates.

  **Note**

  - The smart card and smart card reader that you choose might also impact logon performance. Test different combinations until the terms specified in your service level agreements are satisfied.
- **The nature of the business relationship**. Smart card certificates issued to permanent employees usually warrant a longer lifetime and renewal cycle than certificates issued to short-term workers or to nonemployees.

- **The level of security that you want to enforce**. Highly sensitive operations warrant larger public key values and, typically, shorter certificate lifetimes.

For more information about planning public key and certificate renewal values, see "Designing a Public Key Infrastructure" in this book. For more information about how to configure certificate templates, see "Certificate Templates" in Help and Support Center for Windows Server 2003.

# Establishing Issuance Processes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must establish a plan for the issuance of the smart cards and for the writing of smart card certificates to the cards. This involves making decisions about the following:

- Smart card distribution requirements
- Certificate enrollment options
- Physical distribution of smart cards
- A user preparation plan

# Defining Smart Card Distribution Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Define the procedures for preparing and distributing smart cards and smart card certificates and replacing lost, stolen, or damaged smart cards, as well as contingencies such as when employees change jobs, names, or occupational status.

If you have an existing employee badge process, one solution for smart card distribution is to combine smart card preparation and distribution with badge preparation and distribution. Obtaining a badge typically requires a visit to a security office where the employees must prove their identity and then have their pictures taken. With smart cards, employees can have company-issued certificates attached to the badges that they use for building entry. In this case, the security office requests and installs the certificates on employees' badges, which also serve as their smart cards.

Requiring a person to appear in person and with physical credentials such as a driver's license is the most secure way to distribute smart cards, but this is not always possible. If your organization includes remote offices or traveling users, you need to establish a distribution strategy that accommodates the users' circumstances while minimizing the security risk. For example, you can have a receptionist or administrative assistant give the user a blank smart card and then have the user download the smart card certificate by using self-enrollment. The administrative assistant has physical access to the cards, but not to the PINs or the certificates necessary to activate the card.

You can also use registered mail or another delivery service that requires a signature upon receipt to distribute smart cards to individuals who do not have access to a security office. Otherwise, you can choose a designated individual to physically distribute smart cards. This is the least secure method of smart card distribution.

You must also plan for the physical distribution of replacement cards, especially for mobile or remote office users.

For more information about replacement card planning, see "Planning for Ongoing Smart Card Support" later in this chapter.

# Selecting Certificate Enrollment Options

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, only domain administrators can modify smart card certificate templates. Domain administrators can modify the access permissions on the certificate template to enable either of the following enrollment options:

- Enrollment agents, which allows one or more agents to initialize smart cards on behalf of users.
- Self-enrollment, which allows end users to initialize their own smart cards.

You need to select between the enrollment agent or self-enrollment options, based on the security requirements of your organization and your plan for smart card management and distribution. Using an enrollment agent provides the greatest level of security, but requires the highest level of IT support and is the most expensive. Self-enrollment provides the greatest amount of flexibility, and accommodates remote users, but is not as secure.

**Establishing Enrollment Agents**
If you decide to control smart card issuance from a central location, you need to authorize one or more individuals within the organization to be enrollment agents.

The enrollment agent needs to be issued an Enrollment Agent certificate, which makes it possible for the agent to enroll for certificates on behalf of users.

The advantages of using an enrollment agent include:

- A highly trusted individual processes all certificate and smart card requests.
- Domain administrators can delegate a potentially time-consuming task.
- It simplifies the smart card setup process for users.

The disadvantages include:

- It is difficult to ensure that enrollment agents are trustworthy. One way to enhance this trust is to require approval from several enrollment agents.
- Users in remote locations might not be able to obtain new or replacement smart cards when and where they need them.

Enrollment agents are typically members of the security, IT security, or help desk teams because these individuals have already been trusted with safeguarding valuable resources. In some organizations, such as banks that have many branches, help desk and security workers might not be conveniently located to perform this task. In this case you might need to designate a branch manager or other trusted employee to act as an enrollment agent.

The number of enrollment agents you need depends on:

- The number and proximity of locations in your organization, especially if enrollment agents will be authenticating users in person.
- The number of smart cards that need to be prepared by the enrollment agent.
- The number of other duties that enrollment agents need to perform.

Select the individuals to whom you issue Enrollment Agent certificates carefully. These individuals can enroll for smart card certificates on behalf of any domain user, including an administrator. If these individuals are not trustworthy, they can compromise the security of your organization. To ensure the security of your organization, allow only a limited number of your most trusted employees to serve as enrollment agents.

If you decide to use enrollment agents, prevent unauthorized users from becoming enrollment agents by placing strict controls on the CA used to issue Enrollment Agent certificates. Establish a subordinate CA that is only used to issue Enrollment Agent certificates. After you issue the initial Enrollment Agent certificates, you can either disable certificate issuance or take the CA offline until you require additional enrollment certificates.

**Note**

- For information about delegating enrollment agent authority to individuals who are not domain administrators, see "Prepare a smart card certificate enrollment station" in Help and Support Center for Windows Server 2003.

**Pre-Enrolling User Smart Cards**

If you decide to pre-enroll users for smart card certificates, make sure that the enrollment agent has the blank smart cards as well as the following information:

- The CA selected to issue the smart card certificates to the user, if there are multiple CAs in the organization. If there is only one CA in the organization with smart card certificate templates enabled, that CA is automatically selected.
- The cryptographic service provider that matches the brand of smart card that is to be issued to the user.
- The name of the user to be enrolled. This user must have Enroll permissions for the Smart Card certificate template. A domain administrator can set this either for the individual or for a group of users, such as Authenticated Users or Purchasing.
- The default PIN for the smart card, which is set by the card manufacturer.

  **Note**

  - You can create a script to force the user to change the PIN upon first use of the smart card.

Also, ensure that your enrollment agents review the certificates to verify that the information is correct before they distribute them to users.

**Using Self-Enrollment**

Although using enrollment agents requires more administrative time than allowing users to enroll themselves, the security benefits usually outweigh the overhead costs. However, using enrollment agents might not always be possible or necessary. For example, if it is unlikely that smart cards will be misused, or if the consequences of misuse are minimal, then you might use self-enrollment. In situations in which physical distribution is not possible, self-enrollment is the best alternative.

If you decide to use certificate self-enrollment, users can request a certificate from a Windows Server 2003 CA either manually or automatically. This request can be held pending administrator approval, if you decide that manual approval is required, or until the verification process is completed. Whichever option you choose, the certificate self-enrollment process installs the certificate automatically, or automatically renews the certificate on behalf of the user as soon as the certificate request is approved, based on the specifications in the certificate template.

# Educating Users About Establishing Issuance Processes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

User education is an important component of a smart card management plan. Ensure that users understand the purpose of the smart card deployment. Educate them about proper smart card handling and protection so that they can help the organization to meet its security goals. Emphasize that a smart card is a valuable resource that needs to be protected.

For example, be sure that users understand:

- The hardware and software that they need in order to use smart cards.
- How to install and use their smart cards and readers.
- How they can obtain their smart cards and smart card readers.
- What they need to do in order to configure their systems to use their smart cards.

  **Note**

  - Setup instructions are particularly important for traveling users, or for users who use remote access connections on their home computers.

- What to do if a smart card is lost or stolen.
- Who to call or contact for help and support.

In addition, provide the following guidelines to users:

- **Protect the external smart card chip**. If the chip becomes damaged (scratched, dented, and so forth) the reader might not be able to read the data on the chip.
- **Do not bend the card**. This can break critical internal components. It is risky, for example, to put a smart card in a back pocket, because the individual might sit on it and break its internal components.
- **Do not expose the smart card to temperature extremes**. Leaving a smart card on the dashboard of a car on a hot day can melt or warp the card and harm the chip. Cold temperatures can make the smart card brittle and easier to break.
- **Keep the smart card away from magnetic sources**. This includes credit cards and scanners at retail stores.

It is useful to have a printed version of this user training information available for distribution along with the smart card itself. If your organization also maintains an intranet, publish this information as an easy-to-locate Web page so that users can refer back to your instructions at a later date.

# Preparing a Smart Card Deployment Schedule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying smart cards to a large organization is a major undertaking, often involving thousands of employees. It is essential that you plan the deployment schedule so that your deployment is orderly, controlled, efficient, and secure. For example, you might stagger your deployment so that a defined number of users migrate each week. You can organize this by building or by organizational unit structure.

For a worksheet to assist you in documenting your smart card deployment schedule, see "Smart Card Deployment Schedule" (DSSSMC_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Smart Card Deployment Schedule" on the Web at http://www.microsoft.com/reskit).

Figure 16.7 shows an example of a completed Smart Card Deployment Schedule worksheet for an organization that deployed smart cards according to building.

**Figure 16.7 Example of a Smart Card Deployment Schedule Worksheet**

| Smart Card Deployment Schedule | | | | | |
|---|---|---|---|---|---|
| Prepared by | | | Date | | |
| Location | Deployment Dates | Number of Cards | Cards Prepared (Y/N) | Cards Distributed (Y/N) | Users Enrolled (Y/N) |
| Building 1 | 1/14 to 1/17 | 350 | Yes | Yes | Yes |
| Building 2 | 1/21 to 1/23 | 350 | Yes | No | No |
| Toronto | 1/28 to 1/30 | 425 | No | No | No |
| Edmonton | 2/4 to 2/6 | 300 | No | No | No |
| | | | | | |

In this example, smart card–enabled badges replace the existing building access badges. The plan is to have one member of the distribution team verify the credentials and status in the corporate database, while a second person deactivates the old physical access badge and activates the new smart card. The users receive the smart card, and a document that includes basic instructions, links for more information, an initial PIN, and a reminder that they need to change their PINs within eight hours.

For more information about preparing users for the smart card deployment, see "Educating Users About Establishing Issuance Processes" earlier in this chapter.

# Planning for Ongoing Smart Card Support

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to have all of the smart card management processes and policies in place before you distribute the smart cards to your users. Therefore, you need to create a plan for the ongoing use of and support for smart cards in your organization. When creating a smart card management plan, balance the level of security that you require against the need for ease of use and supportability. Figure 16.8 shows the steps for planning for ongoing smart card support.

**Figure 16.8 Planning for Ongoing Smart Card Support**



# Selecting Group Policy Settings to Manage Smart Card Use

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Several Group Policy settings are specific to smart card management. You can use these Group Policy settings to manage smart cards in your organization.

**Note**

- Other security policy settings, such as lockout policy or restricted logon times, can also impact smart card users if they use their cards for account logon.

**Smart card required for interactive logon**

When you set this policy on a user account, the user cannot log on to the account by using a password. They can only log on by using a smart card.

The advantage of using this policy setting is that it enforces strict security. However, if users are unable to log on by using conventional passwords, you must provide an alternate solution in the event that smart cards become unusable.

**Note**

- This policy setting applies to interactive and network logons only. It does not apply to remote access logons, which are managed by policy settings that are configured on the remote access server.

The **Smart card required for interactive logon** policy is not recommended for users who need to:

- Join a computer to a domain.
- Perform administrative tasks such as installing Active Directory on a member server.
- Configure a network connection for remote access.

If you choose not to use this security policy setting, users can revert to their standard network passwords if their smart cards are damaged or unavailable. However, this weakens security. In addition, users who use their passwords infrequently might forget them, and either write them down, or call the help desk for a password reset, increasing help desk costs to the organization.

**On smart card removal**

Users who walk away from computers that are running an active logon session create a security risk. To enforce the security of your system, it is best if users either log off or lock their computers when they leave. The **On smart card removal** policy allows you to force users to log off or lock their computers when they remove their smart cards.

**Note**

- If you select the forced logoff option, users need to make sure they have saved changes to documents and other files before they remove their smart cards. Otherwise, they lose any changes they have made.

Whether or not you set the **On smart card removal** policy depends on how your users interact with their computers. For example, this policy is a good choice if using computers in an open floor or kiosk environment. This policy might not be necessary when users have dedicated computers or exclusive use of multiple computers. You can use a password-protected screensaver or other means to lock the computers of these users.

**Note**

- The On smart card removal policy is a local computer policy that is administered on a per computer basis. Set the On smart card removal policy on a per user account basis, along with other domain security policy settings.

**Do not allow smart card device redirection**

Use the **Do not allow smart card device redirection** policy if you do not want to use smart cards in conjunction with Terminal Services sessions. Restrict this use of smart cards if you are concerned about the network resources required for Terminal Services sessions in your environment.

**Account lockout threshold**

Use the **Account lockout threshold** policy to disable accounts after a set number of failed logon attempts. An account that is locked out cannot be used until an administrator resets it, or until the account lockout duration expires. You can specify a value of between 1 and 999 failed logon attempts, or you can specify that the account is never locked out by setting the value to 0.

To thwart unauthorized attempts to use a smart card and PIN, establish account lockout thresholds to a low value, such as four or five attempts.

For more information about creating a strategy for unlocking smart cards, see "Defining Administrative and Support Processes" later in this chapter.

# Defining Administrative and Support Processes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many smart card–related problems cannot be addressed by means of Group Policy. In some cases, such as with account lockout policy, Group Policy settings can create support issues for smart card users. You must create a support plan to address potential contingencies related to smart card use. Use the following steps to create your plan:

- Identify the person in your support organization who is authorized to perform security-related smart card administrative tasks, such as resetting PINs, distributing temporary smart cards, or enabling temporary passwords.
- Define and document your plan for escalating and resolving smart card logon problems.
- Define and document any special policies and procedures that you want to establish for specific types of employees, such as senior executives, support personnel, or traveling users.
- Define and document the processes for handling name changes.
- Define and document processes for handling changes in employee status, such as when employees change from contract to full-time employment, or from full-time to part-time employment.
- Define and document the level of smart card support that users can expect to receive outside of regular business hours.

**Creating a Plan for Forgotten Smart Cards**

Before you deploy smart cards, you must establish a contingency plan in the event that users forget their smart cards. You can choose to do one of the following:

- Allow users to maintain passwords that they use for interactive logons when their smart cards are not available. This is the simplest, but least secure option.
- Issue temporary smart cards with certificates that have short lifetimes, such as one day. This solution is useful if your organization has someone on site who is authorized to distribute alternate smart card credentials. However, if an authorized agent is unavailable when the user needs assistance, the user cannot access the network.
- Enable help desk or security personnel to issue users limited-time passwords that are valid for the day. This solution is helpful when no one at the location is authorized to distribute temporary or replacement smart cards, or when you need to provide assistance to traveling users. If you use this option, you must identify how the identity of the user can be validated. Make sure your plan requires the help desk or security administrator to reset the limited-time password after the allotted time has expired.

**Tip**

- When you issue temporary smart cards or passwords to users, you can place those users in a security group that is only able to access a subset of the network resources of your organization. In this way, they can perform their job tasks without compromising the security of your resources.

**Managing PINs**

Using smart cards in combination with PINs simplifies the credential management process for your network. However, you need to create a plan for use in the event that a PIN is forgotten or the account lockout policy causes contingencies related to PIN management. Specifically, you need to define:

- **How locked smart cards can be unlocked**. Decide whether you want the help desk to unlock smart cards or enable users to unlock their own smart cards.
- **How to enable the help desk to unlock locked smart cards**. If you determine that the help desk is responsible for

unlocking smart cards, you need to use some form of secret information to verify that the user is who he or she claims to be.For example, you could have every user complete a set of predefined questions and answers, which are then stored in a secure database. Help desk personnel cannot view the answers to questions. Instead they type in user-provided answers to the questions. The database reports back to the help desk staff member whether the answers match the predefined answers in the database.

**Tip**

- Allow users only one attempt to answer the questions correctly, so that an impersonator does not have multiple opportunities to guess the correct responses to the questions.

- **How to enable users to unlock their own smart cards**. You can enable users to unlock their own smart cards by using a secure Web page linked to a database. The user answers a list of personal questions, the correct answers to which are stored in the secure database. When users need to reset their PINs, they must answer the questions correctly. If they do so, they are allowed to reset their PINs.

  This solution does not require help desk support. However, it exposes the PIN reset user interface to a larger audience, which is a potential security risk. In addition, it requires the user to have Internet access, which might not be possible if the user's smart card is locked.

# Additional Resources for Smart Card Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

## Defining Administrative and Support Processes

These resources contain additional information and job aids related to this chapter.

**Related Information**

- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).
- "Connecting Remote Sites" in *Deploying Network Services* in the *Windows Server 2003 Deployment Kit*.
- "Hosting Applications with Terminal Server" in *Planning Server Deployments* in this kit.
- The Microsoft Platform SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Certificate Templates" in Help and Support Center for Windows Server 2003 for information about how to configure certificate templates.
- "Prepare a smart card certificate enrollment station" in Help and Support Center for Windows Server 2003 for information about delegating enrollment agent authority to individuals who are not domain administrators.

**Related Job Aids**

- "User and Group Smart Card Requirements" (DSSSMC_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "User and Group Smart Card Requirements" on the Web at http://www.microsoft.com/reskit).
- "Smart Card Service Level Agreement" (DSSSMC_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Smart Card Service Level Agreements" on the Web at http://www.microsoft.com/reskit).
- "Smart Card Hardware Specification" (DSSSMC_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Smart Card Hardware Specification" on the Web at http://www.microsoft.com/reskit).
- "Smart Card Reader Evaluation" (DSSSMC_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Smart Card Reader Evaluation" on the Web at http://www.microsoft.com/reskit).
- "Smart Card Deployment Schedule" (DSSSMC_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Smart Card Deployment Schedule" on the Web at http://www.microsoft.com/reskit).

# Planning, Testing, and Piloting Deployment Projects

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Welcome to *Planning, Testing, and Piloting Deployment Projects* of the *Microsoft Windows Server 2003 Deployment Kit*.

This book provides guidelines for planning, testing, and piloting the deployment of Windows Server 2003 and Windows XP Professional operating systems in medium and large organizations. IT professionals can use this book to create a comprehensive project plan, which serves as a framework for designing and deploying complex technologies and operating system features.

A key component of this book is a roadmap that summarizes each of the chapters in the *Windows Server 2003 Deployment Kit*. The roadmap can be used to match business solutions with the design guidelines discussed in the kit. In addition, because testing and piloting are critical tasks in any deployment project, this book contains guidelines for designing and setting up a test lab, testing applications and resolving application compatibility issues, and designing and implementing a pilot project.

This book contains the following chapters:

Planning for Deployment

Designing a Test Environment

Planning and Testing for Application Deployment

Designing a Pilot Project

# Planning for Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To deploy the Microsoft® Windows® Server 2003 operating system, you need a detailed plan that describes how to move from the current state of your IT infrastructure to the new state. To develop a deployment plan, you need to choose a planning framework, establish clear goals and objectives, assess your current environment, create functional designs, and test your deployment in a pilot. The plan you develop and the documents associated with it provide the foundation for your deployment.

**In This Chapter**

Overview of the Deployment Planning Process

Choosing a Deployment Planning Framework

Defining Project Scope and Objectives

Assessing the Current Environment

Creating and Testing Functional Designs

Testing and Conducting the Pilot

Rolling Out the Deployment

Additional Resources for Planning for Deployment

**Related Information**

- For more information about setting up a test lab, see "Designing a Test Environment" in this book.
- For more information about testing to ensure that your applications are compatible with Windows Server 2003, see "Planning and Testing for Application Deployment" in this book.
- For more information about conducting a pilot project, see "Designing a Pilot Project" in this book.

# Overview of the Deployment Planning Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft Windows Server 2003 can provide your organization a reliable, secure, and scalable platform for applications and network services. Windows Server 2003 takes the best of the Microsoft® Windows® 2000 Server operating system technology and makes it easier to deploy, manage, and use. The form your organization's deployment will take is influenced by your organization's size, capabilities, and goals.

Begin the deployment process by forming a deployment team and uniting them behind a single vision. Define the roles on your deployment team, and choose the personnel to perform those roles. Choose a deployment planning methodology, or framework, to organize your team and the tasks required to deploy Windows Server 2003. After assessing your business needs and the way that Windows Server 2003 can help you meet those needs, your team can identify the objectives and scope of the deployment. Next, conduct an inventory of your organization's hardware, software, and network infrastructure to assess the changes you will need to make to deploy the new operating system. Then, using the information contained in this kit, develop functional designs for each component of the deployment. Complete the planning process by developing plans for testing, piloting, and deploying Windows Server 2003.

Many organizations that plan to deploy Windows Server 2003 also plan to deploy Microsoft® Windows® XP Professional operating system on client computers. Throughout this kit, when a task requires Windows XP Professional, information about deploying that operating system is included with the task.

# Choosing a Deployment Planning Framework

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying a large-scale IT project, such as Windows Server 2003, can be a daunting prospect. In fact, many IT deployment projects are not completed on time and within budget, and often the deployment does not include all of the features initially specified.

The failure of an IT deployment to meet its goals generally can be attributed to one or more of the following problems:

- **Disparity between goal and function.** Often, IT projects are viewed by an organization exclusively in terms of the functions and features they provide, not in terms of the business goals that those functions and features can help the organization achieve.

- **Lack of a common language between business and technical groups.** An IT group must understand the organization's business problems and needs in order to assess how a proposed solution can solve those problems and meet those needs. Establishing good communication between business and IT groups ensures that the deployment meets the organization's needs.

- **Inflexible processes.** In large, complex projects, changes are inevitable. Rigid rules and practices prevent a team from responding when circumstances change.

- **Complexity in the technical environment.** Increased reliance on servers and Web-distributed applications demands greater collaboration and cooperation within the IT department to assure interoperability and reliable operations. Within highly structured organizations, creative problem-solving is both challenging and vital.

- **Resource constraints.** Staffing a deployment team can be difficult because the demand for highly qualified IT personnel is always high. Often, budgets for new project development are constrained.

To avoid or mitigate these pitfalls, your organization can employ an IT deployment planning methodology. If your organization already uses a deployment methodology, you can use it to deploy Windows Server 2003. If you have not used a deployment planning methodology, you might want to consider using the Microsoft Solutions Framework (MSF), upon which the deployment of Windows Server 2003 has been based, as described in this kit.

MSF is an IT deployment planning methodology that was developed by Microsoft Consulting Services in conjunction with Microsoft business groups, Microsoft partners, and other industry sources. MSF provides the principles, models, and disciplines for managing the people, processes, and technology elements of an IT deployment, such as the deployment of Windows Server 2003.

Two business models form the foundation of MSF:

- The MSF team model, which ensures that all project goals are met by linking each team role, or area of expertise, with a major project responsibility

- The MSF process model, which organizes the process of creating and deploying an IT project by dividing it into distinct phases marked by milestones

# MSF Team Model Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Identifying the individuals with the right mix of technical and project management skills for deploying Windows Server 2003 is one of the first decisions that you will make. Even before staffing can begin, however, you need to identify the team roles, or areas of expertise, required for an IT project.

The MSF team model is based on six interdependent, multidisciplinary roles: product management, program management, development, testing, user experience, and release management.

**Product management** Product management is responsible for identifying the organization's business needs and the needs of the end users, and for making sure those needs are met.

**Program management** The program management team's goal is to deliver Windows Server 2003 within project constraints. The team is responsible for managing the project schedule and budget, and for reporting project status, managing risk factors, and managing the functional design of the deployment.

**Development** The development team builds the Windows Server 2003 implementation according to specification. The team's responsibilities include specifying the features of the Windows Server 2003 physical design, estimating the time and effort required to develop each feature, and preparing the product for deployment.

**Testing** The testing team ensures that Windows Server 2003 is released into the production environment only after all quality issues have been identified and resolved. The team's responsibilities include developing the testing strategy, designing and building the test lab, developing the test plan, and conducting tests.

**User experience** The user experience team ensures that Windows Server 2003 meets the users' needs. The team gathers, analyzes, and prioritizes user requirements and develops user training.

**Release management** The release management team is responsible for deploying Windows Server 2003. By designing and managing a pilot deployment of Windows Server 2003, the release management team ensures that Windows Server 2003 is sufficiently stable for deployment into the production environment.

The MSF team roles are flexible; they can be adapted to your organization's own processes and management philosophy. In a small organization or a limited deployment, one individual might play multiple roles. In larger organizations, a team might be required to perform all of the tasks assigned to each role.

For more information about the MSF team model, see the "MSF Team Model" white paper in the MSF Resource Library, which is available from the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# MSF Process Model Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The MSF process model divides the process of creating and deploying a solution, such as Windows Server 2003, into five distinct, timed phases, as illustrated in Figure 1.1.

**Figure 1.1 Microsoft Solutions Framework Phases**

Each of the five phases focuses on one major cluster of activities required to successfully deploy a solution, and each phase culminates in the completion of a major milestone:

- During the envisioning phase, the team identifies the high-level requirements and business goals of the project. Using this information, the team then develops the Vision/Scope document, which states the project goals and scope. This phase ends when the Vision/Scope document is approved.

- During the planning phase, the team identifies the Windows Server 2003 features that it will implement and decides how to implement them. In addition, the team develops a schedule. This phase ends when the master project plan is approved.

- During the developing phase, the appropriate teams develop and test the Windows Server 2003 code, infrastructure, and documentation. This phase ends when all of the components are complete.

- During the stabilizing phase, the appropriate teams test, stabilize, and prepare your organization's Windows Server 2003 implementation for release into your production environment. This phase ends when your organization's management approves Windows Server 2003 for release.

- During the deploying phase, the release team deploys Windows Server 2003 into your production environment.

For more information about the MSF process model, see the "MSF Process Model" white paper in the MSF Resource Library, which is available from the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Defining Project Scope and Objectives

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Defining your deployment objectives helps you to determine which Windows Server 2003 features you need to implement and how those features relate to your existing environment. Defining the project's scope helps you to identify exactly what the team can and cannot deliver within the project constraints.

To identify your organization's objectives in deploying Windows Server 2003, consider the following questions:

- Why is your organization deploying Windows Server 2003?
- What business benefits will your organization derive from deploying Windows Server 2003?
- What IT benefits will your organization derive from deploying Windows Server 2003?
- How does your organization's current IT environment prevent you from meeting your organization's business needs?
- Who will be affected by the move from your current operating system to Windows Server 2003?
- What must the deployment accomplish to be deemed successful?
- To define the scope of the deployment, ask yourself what can reasonably be accomplished by this project and what cannot, given the time frame, your budget, and other resources.

For a list of Windows Server 2003 features, see the Windows Server 2003 Features link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Assessing the Current Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before your team can begin planning the deployment of Windows Server 2003, you need to assess your organization's current IT environment. Doing so helps you understand your organization's structure, so that you can create a deployment plan that addresses real business needs.

Document the hardware and software used in your organization by conducting an inventory. Then, document your network and your IT environment by drawing diagrams and work process flowcharts that represent infrastructure relationships and processes, such as the relationship between your network and Microsoft® Active Directory® directory services architecture. After you have collected all relevant data about your environment, include these diagrams and flowcharts in your deployment plan.

# Conducting Hardware and Software Inventories

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy Windows Server 2003, you might need to upgrade the hardware and software used in your organization. Therefore, it is essential that you conduct comprehensive hardware and software inventories of all computers on your network before you begin planning your deployment.

You can use Microsoft® Systems Management Server (SMS) to conduct an inventory of the hardware and software on your network. For more information about SMS, see the Microsoft Systems Management Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

To obtain a hardware and software inventory from Microsoft® Windows® and Microsoft® Windows NT® operating system networks that use Windows Management Instrumentation (WMI), you can use scripts and a variety of third-party applications. For more information about scripting, see the at http://www.microsoft.com/reskit, or see the MSDN Scripting Clinic link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Conducting a Hardware Inventory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For each computer in your organization, document the following:

- Location of the computer
- Number of processors it has and the speed of each processor
- Amount of RAM
- Available disk space
- Type of monitor
- BIOS settings
- Configurations for peripheral devices, such as printers, scanners, and input devices
- Driver version numbers and other software and firmware information

In addition, document every device that is connected to your network, including network adapters. Also note departmental and geographic differences in the types of hardware used within your organization.

# Conducting a Software Inventory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Conduct a complete inventory of the applications that are used in your organization, including all custom (in-house) applications. Documenting the software installed on each computer will help you to identify the applications that will need to be reinstalled or upgraded after you have installed Windows Server 2003 or Windows XP Professional.

For each client and server computer in your organization, document the following:

- Applications that are installed
- Operating system that is installed
- Service packs that have been applied to the operating system
- Service packs that have been applied to applications
- Version numbers (or date and time-stamp data) of dynamic-link libraries (DLLs) associated with each application

Note whether each application is a standard desktop application, a line-of-business application, or a server-based application or service, as described in Table 1.1.

**Table 1.1 Application and Service Types**

| Application and Service Type | Description |
|---|---|
| Standard desktop | An application, such as Microsoft® Word or e-mail client software, that is required by a wide variety of users throughout the organization. |
| Line of business | An application that is more specialized and that is used only by specific departments, groups of users (such as developers), or geographic locations. Line-of-business applications can be developed internally or by third-party vendors. In your inventory, note where these applications are used within the organization. |
| Server–based application or service | These might include Microsoft® Internet Information Services (IIS), Microsoft® Exchange Server, database software, security services, or other services hosted on servers within your organization. |

Also note how important each application is to your organization by assigning it to one of the categories listed in Table 1.2.

**Table 1.2 Application and Service Categories**

| Category | Description |
|---|---|
| **Mission-critical** | An application or service that is essential to business operations. |
| **Tactical** | An application or service that is valuable to the business, but does not provide optimal benefit. |
| **Legacy** | An application or service that is still used by some groups or departments, but is nearing the end of its usefulness. Your plan should call for removing legacy applications before they become obsolete. |

| | |
|---|---|
| **Obsolete** | An application or service that provides no benefit to the business. Plan to remove obsolete applications as soon as possible because they can present interoperability problems that are not worth solving. |

For more information about conducting a software inventory, see "Planning and Testing for Application Deployment" in this book.

# Determining Hardware Requirements for Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you have a complete list of the hardware and supporting software used in your organization, compare the data you collected with the hardware requirements in Table 1.3 and Table 1.4 to determine whether your hardware is compatible with Windows Server 2003 and — where applicable — with Windows XP Professional. All of the client and server computers on which you plan to install these operating systems must be capable of supporting the installation.

**Important**

- When determining the minimum hardware requirements for your deployment project, account for all of the applications and services that computers in your organization use. Some applications might require faster processors or more RAM than Windows Server 2003 requires. When evaluating the published hardware requirements for your applications, account for the complete desktop and server configuration.

**Table 1.3 Windows Server 2003 Hardware Requirements**

| Windows Server 2003 Edition | Number of Processors | Processor Speed | RAM | Available Disk Space (for Setup) | Monitor |
|---|---|---|---|---|---|
| Web | 1–2 | 133 megahertz (MHz) minimum; 550 MHz recommended | 128 megabytes (MB) minimum; 256 MB recommended; 2 GB maximum | 1.5 gigabytes (GB) | Video Graphics Adapter (VGA) or higher; Super VGA (SVGA) (800 × 600) or higher recommended |
| Standard | 1–4 | 133 MHz minimum; 550 MHz recommended | 128 MB minimum; 256 MB recommended; 4 GB maximum | 1.5 GB | VGA or higher; SVGA (800 × 600) or higher recommended |
| Enterprise | 1–8 | 133 MHz minimum; 550 MHz recommended | 128 MB minimum; 256 MB recommended; 32 GB maximum | 1.5 GB | VGA or higher; SVGA (800 × 600) or higher recommended |
| Datacenter | 8–32 | 400 MHz minimum | 512 MB minimum; 64 GB maximum | 1.5 GB | VGA or higher; SVGA (800 × 600) or higher recommended |

**Note**

- Upgrade your organization's devices, remote access, and applications before deploying Windows Server 2003. For information about upgrading, see "Designing Unattended Installations" in *Automating and Customizing Installations* of this kit.

Table 1.4 lists the minimum and recommended hardware requirements for installing Windows XP Professional.

**Table 1.4 Windows XP Professional Hardware Requirements**

| Minimum Requirements | Recommended Requirements |
|---|---|
| Intel Pentium (or compatible) 233 MHz or higher processor | Intel Pentium II (or compatible) 300 MHz or higher processor |
| 64 MB of RAM | 128 MB (4 GB maximum) of RAM |

| | |
|---|---|
| 2 GB hard disk with 650 MB of free disk space (additional disk space required if installing over a network) | 2 GB of free disk space |
| VGA or higher video adapter | SVGA video adapter and Plug and Play monitor |
| Keyboard, mouse, or other pointing device | Keyboard, mouse, or other pointing device |
| CD-ROM or DVD-ROM drive (required for CD installations) | CD-ROM or DVD-ROM drive (12X or faster) |
| Network adapter (required for network installation) | Network adapter (required for network installation) |

# Assessing Your Network Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assess your network infrastructure by identifying existing network protocols, network bandwidth, and network hardware, such as routers, switches, and firewalls. Be sure to assess topology, network size, type of network, and traffic patterns, in addition to the logical organization of your network, name-resolution and address-resolution methods, naming conventions, and network services in use. You must also include the location of network sites and the available bandwidth between sites: This information can help you choose an installation method.

After you have identified the primary features of your network, you can start documenting your network infrastructure. You can use Windows Server 2003 diagnostic applications, such as Network Monitor, to gather information about your network. Or, you can use troubleshooting and configuration software offered by OEMs or software vendors. Then, create physical and logical network diagrams using a product such as Microsoft® Visio® drawing and diagramming software.

The physical and logical network diagrams should include the following information:

- Physical communication links, including cables and the paths of analog and digital lines.
- Server names, IP addresses, and domain membership.
- Location of printers, hubs, switches, routers, bridges, proxy servers, and other network devices.
- Wide area network (WAN) communication links, their speed, and available bandwidth between sites. If you have slow or heavily used connections, it is important to note them.
- Physical network infrastructure.
- Addressing infrastructure.
- Naming infrastructure.
- Authentication infrastructure.
- Security infrastructure.
- Intranet infrastructure.
- Management infrastructure.

The logical network diagram should also include the following information:

- Domain architecture
- Server roles
- Trust relationships and any policy restrictions that might affect your deployment

In addition, you need to indicate where your Windows environment is connected to other networking environments, such as UNIX or Novell NetWare. Identify the tools that you use to maintain this connectivity, and determine whether they are supported under Windows Server 2003.

Finally, be sure to address network security measures in your assessment of the network. Include information about how you manage client authentication, user and group access to resources, and Internet security. Also include your firewall and proxy configurations.

For detailed information about planning your network infrastructure, see the related chapters in *Deploying Network Services* of

this kit.

# Assessing Your IT Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Document the following key components of your IT environment.

**Business organization and geographical requirements** Conduct a site analysis of each physical location in your organization. Document the number of physical locations, the geographic distribution of locations, and the speed and reliability of network connections between locations. For each location, record the number of client and server computers, the function of the client and server computers, and the availability of support staff.

**Information architecture** Diagram your key business processes to illustrate how the information needed to make critical decisions will be available when needed. For example, are sales and marketing personnel able to confirm precise delivery dates for customer orders?

**Administrative model** By examining your existing administrative model, you identify the administrative tasks that IT personnel perform in all areas of your organization. This helps determine whether you need to change any aspect of the administrative operations design to accommodate the Windows Server 2003 features that you want to deploy. This also will highlight opportunities for improving administrative efficiency and security in your organization.

**IT standards** Over time, the network and application standards in many organizations become fragmented or obsolete. This is common in organizations that have merged with or acquired other companies. Disparate systems, built over time, designed by different people, and often geographically separated, are a potential risk to a successful deployment. Audit your existing systems to identify problem areas.

**Active Directory design** If your organization is planning to upgrade Active Directory directory service, be sure to document your directory infrastructure. For information about upgrading Active Directory to meet your organization's growth and business needs see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit.

**Security architecture** Security issues, such as who has access to your network and to your databases and how your organization is protected from the Internet, are vital considerations in all aspects of your Windows Server 2003 deployment plan. For information about how to secure your Windows Server 2003 environment, see "Designing the Active Directory Logical Structure" and the security chapters in Part 2 of *Designing and Deploying Directory and Security Services* of this kit.

# Creating and Testing Functional Designs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The next step in the deployment planning process is creating the actual design, or functional specification, for the Windows Server 2003 features that you intend to implement in your organization. The process of creating the functional specification and testing your implementation of Windows Server 2003 begins in the MSF planning phase and continues through the developing and stabilizing phases.

In the functional specification, describe the operating system features that you will implement, and how they will be configured and deployed. Many features of Windows Server 2003 are interrelated, particularly for implementations that include Active Directory. For this reason, create a dependencies matrix by using the information provided in the tables in this section, and include it in your functional specification.

Also describe the different types of computer users in your organization, the key tasks they perform, how these tasks are currently performed, and how performance can be improved in the new network environment. If yours is a large organization with multiple sites, or an international organization, include geographical concerns in your functional specification.

Creating a functional specification is an iterative process. As you develop your functional specification, test your proposed implementation in your test lab. Testing your functional specification helps you determine the best way to use Windows Server 2003 features to meet your project objectives. For information about creating a test lab to test your deployment of Windows Server 2003, see "Designing a Test Environment" in this book.

The following tables describe the contents of each chapter in this kit. Use these tables to find the information that you need to create your functional specification and plan your deployment.

Each table contains the following information:

- Chapter title, along with a brief description of the topics addressed in the chapter.
- Prerequisites for reading the chapter. Prerequisites may include one or all of the following:
    - **Tasks**. Occasionally you must perform certain tasks before you design or implement the feature or technology described in the chapter. For example, most Windows Server 2003 deployments will include Active Directory and DNS. Organizations that also will deploy Remote Installation Services (RIS) will need to take RIS into account when creating their Active Directory and DNS designs.
    - **Reading**. Your network designers should consult other chapters in the kit before creating their functional specifications.
    - **Technical background**. Your network designers should understand certain technological concepts and issues before deploying the technology described in the chapter.

Table 1.5 and Table 1.6 describe the chapters in *Deploying Network Services*, which is divided into two parts.

**Table 1.5 "Part 1: Core Network Infrastructure"**

| Chapter | Prerequisites |
|---|---|

| | |
|---|---|
| Designing a TCP/IP Network<br><br>Describes how to design a TCP/IP infrastructure for a new environment or for an existing environment. Explains the fundamental design decisions that you need to make about your networking infrastructure, addressing such issues as routing, addressing, multicasting, security, and introducing IP version 6 (IPv6) in your IP version 4 (IPv4) network. | **Technical Background**<br><br>General knowledge of TCP/IP.<br><br>**Tasks**<br><br>Might require changes to your Internet Protocol security (IPSec) design.<br><br>**Reading**<br><br>"Deploying DHCP"<br><br>"Deploying Domain Name System (DNS)"<br><br>"Deploying WINS"<br><br>"Deploying IPSec"<br><br>"Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* |
| Deploying DHCP<br><br>Describes DHCP server design considerations, and explains how to integrate DHCP with other services, define and activate scopes for your DHCP clients, and implement your DHCP solution. | **Technical Background**<br><br>Familiarity with your organization's DNS and WINS design. |
| Deploying Domain Name System (DNS)<br><br>Describes planning and designing your DNS infrastructure (DNS namespace, DNS server placement, DNS zones, and DNS client configuration). Also explains how to plan the level of integration with Active Directory that you require and how to identify security, scalability, and performance requirements, which is a concern for those who are integrating DNS with Active Directory. | **Technical Background**<br><br>Familiarity with your organization's DHCP and WINS design.<br><br>**Tasks**<br><br>Might require changes to your organization's cross-forest trusts in your Active Directory installation.<br><br>**Reading**<br><br>"Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* |
| Deploying WINS<br><br>Describes how to develop a server strategy, design a replication strategy, secure your WINS solution, integrate WINS with other services, and implement your WINS solution. | **Technical Background**<br><br>Familiarity with your organization's DNS and DHCP design.<br><br>**Reading**<br><br>"Deploying DHCP"<br><br>"Deploying Domain Name System (DNS)" |

| | |
|---|---|
| Deploying ISA Server<br><br>Describes how to identify the most appropriate deployment mode (cache, firewall, or integrated) for your organization and how to integrate Microsoft Internet Security and Acceleration (ISA) Server with other components in your network. Other topics include security considerations and implementing your deployment design. | **Tasks**<br><br>You must deploy your network infrastructure first.<br><br>Might require simultaneously designing your remote access server deployment.<br><br>**Reading**<br><br>"Deploying Dial-up and VPN Remote Access Servers"<br><br>"Connecting Remote Sites" |
| Deploying IPSec<br><br>Describes the design decisions that you must make and the planning tasks you must perform to use Internet Protocol security (IPSec) for end-to-end secure networks. Topics include determining where and how you must secure computers on your network, and deciding which policy provides the security measures you require. | **Technical Background**<br><br>Familiarity with your organization's public key infrastructure (PKI) design and IP design.<br><br>**Tasks**<br><br>Usually requires deployment of Group Policy.<br><br>Might also require deployment of Active Directory.<br><br>**Reading**<br><br>"Designing a TCP/IP Network" |

**Table 1.6 "Part 2: Network Access Infrastructure"**

| Chapter | Prerequisites |
|---|---|
| Deploying Internet Authentication Service (IAS)<br><br>Describes how to decide what role Internet Authentication Service (IAS) should play in your network infrastructure, how to configure connection request and remote access policies that allow users to access the network, and how to secure and optimize the deployment of Windows Server 2003. | **Reading**<br><br>"Deploying Dial-up and VPN Remote Access Servers"<br><br>"Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services*<br><br>"Deploying a Wireless LAN" |
| Deploying Dial-up and VPN Remote Access Servers<br><br>Describes dial-up and virtual private network (VPN) remote access and helps you to decide which solution will best serve your organization. Also describes the network infrastructure needed to support your remote access server design. | **Reading**<br><br>"Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services*<br><br>"Deploying Internet Authentication Service (IAS)"<br><br>"Deploying ISA Server" |

| | |
|---|---|
| Deploying Remote Access Clients Using Connection Manager<br><br>Provides information about remote access clients and describes how to provide phone book support, how to implement Connection Manager, and how to customize Connection Manager for your organization using the Connection Manager Administration Kit (CMAK). | **Technical Background**<br><br>Familiarity with your organization's security practices, such as smart card usage.<br><br>**Tasks**<br><br>You must first deploy an authentication service, such as IAS, to enable authentication, authorization, and auditing.<br><br>**Reading**<br><br>"Deploying Dial-up and VPN Remote Access Servers"<br><br>"Deploying Internet Authentication Service (IAS)" |
| Connecting Remote Sites<br><br>Describes how to choose the remote site connection type and its configuration options, which security features to use, and how to integrate the remote site connection into your existing network infrastructure. Also describes how to deploy your remote site connection after all the design decisions have been made. | **Reading**<br><br>"Deploying Dial-up and VPN Remote Access Servers"<br><br>"Deploying a Public Key Infrastructure" and "Designing the Site Topology" in *Designing and Deploying Directory and Security Services* |
| Deploying a Wireless LAN<br><br>Describes how to design and deploy a secure, manageable wireless local area network (WLAN) in an enterprise environment. | **Tasks**<br><br>You must deploy your network infrastructure first. In addition, to provide authorization and authentication for wireless users, automatic IP address assignment, and name resolution, your networking infrastructure should include the following components and services:<br><br>• Active Directory<br><br>• Remote Authentication Dial-In User Service (RADIUS) servers and proxies<br><br>• A certificate infrastructure (also known as a public key infrastructure, or PKI)<br><br>• DHCP services<br><br>• DNS services |

Table 1.7 and Table 1.8 describe the chapters in Designing and Deploying Directory and Security Services, which is divided into two parts.

**Table 1.7 "Part 1: Designing and Deploying Directory Services"**

| Chapter | Prerequisites |
|---|---|
| Planning an Active Directory Deployment Project<br><br>Provides the planning information necessary to help you determine your Active Directory deployment strategy, which will vary depending on your existing network configuration. Also provides a roadmap of the rest of "Part 1: Designing and Deploying Directory Services" to help you identify your next steps in the Active Directory deployment process. | **Technical Background**<br><br>Familiarity with all applications in your organization that use Active Directory, such as Microsoft Exchange Server. |

| | |
|---|---|
| Designing the Active Directory Logical Structure<br><br>Provides detailed information on designing an Active Directory logical structure for organizations that have no Active Directory structure and for those that want to change their current Active Directory deployment. | None |
| Designing the Site Topology<br><br>Provides information on designing a site topology for your environment, including information on how to sufficiently route query and replication traffic in ways that minimize both cost and administrative effort. Also explains how to determine which locations require domain controllers and where you need to create sites, site links, and site link bridges. | **Reading**<br><br>"Designing the Active Directory Logical Structure" |
| Planning Domain Controller Capacity<br><br>Provides guidelines for planning domain controller hardware needs and placement in a Windows Server 2003 environment, with an emphasis on monitoring needs. | **Technical Background**<br><br>General knowledge of your organization's logical structure and site topology.<br><br>**Reading**<br><br>"Designing the Active Directory Logical Structure"<br><br>"Designing the Site Topology" |
| Enabling Advanced Windows Server 2003 Active Directory Features<br><br>Provides information about Active Directory functional levels and the features that are enabled at each functional level. Also describes the process for raising functional levels in a Windows Server 2003 environment. | **Reading** (especially for new deployments):<br><br>"Designing the Active Directory Logical Structure"<br><br>"Designing the Site Topology"<br><br>"Planning Domain Controller Capacity" |
| Deploying the Windows Server 2003 Forest Root Domain<br><br>Describes the process for deploying a forest root domain in an environment that has no Active Directory structure. Steps include reviewing the Active Directory designs created in preceding chapters and configuring DNS for the root domain. | **Reading**<br><br>"Designing the Active Directory Logical Structure"<br><br>"Designing the Site Topology"<br><br>"Planning Domain Controller Capacity"<br><br>"Enabling Advanced Windows Server 2003 Active Directory Features" |
| Deploying Windows Server 2003 Regional Domains<br><br>Describes the process for deploying Windows Server 2003 regional domains. Also includes steps for reviewing your regional domain design and for delegating the DNS domain for the new regional domain. | **Reading**<br><br>"Designing the Active Directory Logical Structure"<br><br>"Designing the Site Topology"<br><br>"Planning Domain Controller Capacity"<br><br>"Enabling Advanced Windows Server 2003 Active Directory Features"<br><br>"Deploying the Windows Server 2003 Forest Root Domain" |

| | |
|---|---|
| Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory<br><br>Provides detailed steps for upgrading from Windows NT 4.0 to an existing forest or to a single domain forest. Also describes tasks that need to be performed prior to and after upgrading. | **Reading**<br><br>"Deploying the Windows Server 2003 Forest Root Domain"<br><br>"Deploying Windows Server 2003 Regional Domains"<br><br>If you plan to upgrade some Windows NT 4.0 domains and restructure others, you also should read "Restructuring Windows NT 4.0 Domains to an Active Directory Forest" |
| Upgrading from Windows 2000 Domains to Windows Server 2003 Domains<br><br>Provides detailed steps for upgrading from Windows 2000. Also describes tasks that need to be performed prior to and after upgrading. | **Reading**<br><br>"Enabling Advanced Windows Server 2003 Active Directory Features"<br><br>If you plan to change your current Active Directory structure, also read "Designing the Active Directory Logical Structure" and "Designing the Site Topology" |
| Restructuring Windows NT 4.0 Domains to an Active Directory Forest<br><br>Describes how to complete the necessary planning and preparation tasks, and how to migrate Windows NT 4.0 Server accounts and resources to an Active Directory domain or domains using the Active Directory Migration Tool (ADMT v2). | **Reading**<br><br>"Deploying the Windows Server 2003 Forest Root Domain"<br><br>"Deploying Windows Server 2003 Regional Domains"<br><br>If you plan to upgrade some domains and restructure others, also read "Upgrading Windows NT 4.0 Domains to Windows Server 2003 Active Directory." |
| Restructuring Active Directory Domains Between Forests<br><br>Describes how to complete the necessary planning and preparation tasks, and how to migrate accounts and resources between Active Directory domains in two different forests using ADMT v2. | **Reading** (if you plan to change your existing Active Directory structure):<br><br>"Designing the Active Directory Logical Structure"<br><br>"Designing the Site Topology"<br><br>"Enabling Advanced Windows Server 2003 Active Directory Features"<br><br>"Designing and Deploying File Servers" in *Planning Server Deployments* (for potential issues related to File Replication service [FRS]) |

| | |
|---|---|
| Restructuring Active Directory Domains Within a Forest<br><br>Describes how to complete the necessary planning and preparation tasks, and how to migrate accounts and resources between Active Directory domains located in the same forest by using ADMT v2. | **Reading** (if you plan to change your existing Active Directory structure):<br><br>"Designing the Active Directory Logical Structure"<br><br>"Designing the Site Topology"<br><br>"Enabling Advanced Windows Server 2003 Active Directory Features"<br><br>"Designing and Deploying File Servers" in *Planning Server Deployments* (for potential issues related to FRS) |

**Table 1.8 "Part 2: Deploying Distributed Security Services"**

| Chapter | Prerequisites |
|---|---|
| Planning a Secure Environment<br><br>Provides an overview of the security-related solutions discussed throughout this kit. Introduces the components of a secure Windows Server 2003 environment, explains how to securely address user needs, and describes how to establish a secure common infrastructure. | None |
| Designing an Authentication Strategy<br><br>Describes how to create an authentication strategy that includes account management, authentication protocol selection, and configuration options. The strategy you devise enhances the security of the authentication process. | **Technical Background**<br><br>Familiarity with your organization's implementation of Microsoft® Outlook® messaging and collaboration client and Microsoft Exchange Server.<br><br>**Reading**<br><br>"Planning a Smart Card Deployment" (if you plan to use smart cards in your organization)<br><br>"Designing the Active Directory Logical Structure" |
| Designing a Resource Authorization Strategy<br><br>Provides background information on security groups and domain and forest functional levels. Also describes how to develop a strategy for using security groups to effectively manage access to resources, and how to delegate security group maintenance. | **Task**<br><br>You must deploy Active Directory first.<br><br>**Reading**<br><br>"Designing the Active Directory Logical Structure"<br><br>"Restructuring Windows NT 4.0 Domains to an Active Directory Forest"<br><br>"Deploying Windows Server 2003 Regional Domains" |
| Designing a Public Key Infrastructure<br><br>Describes how to create an infrastructure of certification authorities, for the purpose of issuing and managing certificates for a variety of purposes. Steps include defining your certificate requirements, creating a design for your infrastructure, creating a certificate management plan, and deploying your PKI solution. | **Task**<br><br>You must deploy Active Directory first, if your organization uses Active Directory. You do not need Active Directory to deploy a PKI; however, the capabilities of the Windows Server 2003 PKI are greatly enhanced by Active Directory. |

| | |
|---|---|
| Planning a Smart Card Deployment<br><br>Describes how to develop specifications and management protocols that make it possible to secure various user activities through the use of smart cards. | **Reading**<br><br>"Designing a Public Key Infrastructure"<br><br>"Deploying Remote Access Clients Using Connection Manager" in *Deploying Network Services*<br><br>"Designing an Authentication Strategy" |

Table 1.9 describes the contents of *Planning Server Deployments*.

**Table 1.9 Planning Server Deployments**

| Chapter | Prerequisites |
|---|---|
| Planning for Storage<br><br>Provides an overview of Windows Server 2003 storage solutions and discusses how to determine application and organizational storage requirements, and how to plan for storage scalability, storage area network (SAN) deployment, fault tolerance, and disaster recovery. | None |
| Designing and Deploying File Servers<br><br>Provides an overview of new file server features in Windows Server 2003 and discusses how to identify file service goals, design Distributed File System (DFS) namespaces, plan file server availability by using server clusters or File Replication service (FRS), design a standard file server configuration, and secure and deploy file servers. | **Task**<br><br>You must deploy Active Directory first.<br><br>**Reading**<br><br>"Designing and Deploying Server Clusters"<br><br>"Designing a Resource Authorization Strategy" in *Designing and Deploying Directory and Security Services* |
| Designing and Deploying Print Servers<br><br>Describes the processes for inventorying your current print environment, identifying and resolving client print driver compatibility issues, creating a standard print server hardware design, designing your new print environment, increasing your print server availability, and deploying Windows Server 2003 print servers. | **Task**<br><br>You must deploy Active Directory first.<br><br>**Reading**<br><br>"Designing and Deploying Server Clusters"<br><br>"Designing a Resource Authorization Strategy" in *Designing and Deploying Directory and Security Services* |
| Hosting Applications with Terminal Server<br><br>Describes how to create an application hosting plan, how to plan network and hardware resources for Terminal Server, and how to design user and domain configurations for Terminal Server. | **Task**<br><br>You must design your Windows Server 2003 domain infrastructure first. |

| | |
|---|---|
| Planning for Remote Server Management<br><br>Helps you determine which servers to manage and whether to manage them remotely or locally. Explains how to identify the appropriate tools to use for management. Describes the required supporting hardware components and configurations. Also discusses network and security considerations and remediation, and whether and how to deploy Windows Server 2003 remotely. | **Technical Background**<br><br>Knowledge of your organization's server configurations, such as server location and role.<br><br>Knowledge of the availability requirements for the servers you plan to manage remotely and who the administrators are. |
| Planning for High Availability and Scalability<br><br>Introduces IT practices and hardware solutions that can increase the availability of your Windows Server 2003 servers. Also introduces the two clustering technologies available in Windows Server 2003: server clusters and Network Load Balancing (NLB). | None |
| Designing and Deploying Server Clusters<br><br>Describes how to plan and deploy server clusters. Includes discussions of application analysis and deployment strategies, failover strategies, network configuration, data storage considerations, and security considerations. Also explains how to start planning geographically dispersed clusters. | **Reading**<br><br>"Planning for High Availability and Scalability"<br><br>"Designing and Deploying File Servers"<br><br>"Designing and Deploying Print Servers" |
| Designing Network Load Balancing<br><br>Describes how to create designs that include NLB clusters. Topics include how to determine the number of clusters required, how to determine the number of cluster hosts required, the use of virtual clusters, securing the administration of the cluster, scaling applications and services by using NLB, and improving the availability of services by using NLB. | **Reading**<br><br>"Planning for High Availability and Scalability" |
| Deploying Network Load Balancing<br><br>Describes how to deploy the design created in the previous chapter without affecting existing services. Includes information about how to install a new NLB cluster, upgrade an existing NLB cluster, or migrate an existing NLB cluster. | **Reading**<br><br>"Planning for High Availability and Scalability"<br><br>"Designing Network Load Balancing" |

Table 1.10 describes the content of *Designing a Managed Environment*.

**Table 1.10 Designing a Managed Environment**

| Chapter | Prerequisites |
|---|---|
| Planning a Managed Environment<br><br>Describes how to define the scope of your project, assess your current environment, determine your requirements for change and configuration management technologies, and design managed solutions. | None |

| | |
|---|---|
| Designing a Group Policy Infrastructure<br><br>Describes how to design your organizational unit (OU) structure, define the scope of the application of Group Policy, test your design, and deploy it. | **Reading**<br><br>"Planning a Managed Environment"<br><br>"Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* |
| Staging Group Policy Deployments<br><br>Describes the process for staging Group Policy, which includes creating a staging environment that mimics the production environment, testing new Group Policy settings, and deploying those settings in the production environment. | **Reading**<br><br>"Designing a Group Policy Infrastructure" |
| Deploying Security Policy<br><br>Provides specific instructions for using Group Policy to secure your managed environment. Includes information about how to assess potential security risks, how to analyze the consequences of security breaches and how to develop countermeasures for them, and how to plan a security strategy with Group Policy. | **Reading**<br><br>"Designing a Group Policy Infrastructure"<br><br>"Staging Group Policy Deployments"<br><br>All of the chapters in "Part 2: Deploying Distributed Security Services" in *Designing and Deploying Directory and Security Services* (see<br><br>Table 1.8) |
| Deploying Software Update Services<br><br>Describes using Software Update Services (SUS) to manage security updates for your clients and servers. Topics include how to design the SUS server configuration, how to deploy the SUS server component, and how to deploy the Automated Updates feature. | None |
| Migrating User State<br><br>Describes how to migrate user state, including identifying the best methods for collecting, storing, and restoring user state data and deciding which user data and settings should be migrated. Other topics include preparing a migration plan that addresses storage and security requirements; potential registry, drive, and domain changes; and training users. | None |
| Implementing User State Management<br><br>Describes how to use IntelliMirror® management technologies to ensure that a user's data, applications, and settings remain constant throughout the user's environment. This process includes defining your objectives for managing user data and settings, determining which IntelliMirror technologies meet the needs of your users and your organization, and testing your designs in a controlled environment. | **Reading**<br><br>"Planning a Managed Environment"<br><br>"Designing a Group Policy Infrastructure"<br><br>"Staging Group Policy Deployments" |

| | |
|---|---|
| Deploying a Managed Software Environment<br><br>Describes the process of deploying a managed software environment. Topics include how to prepare and deploy applications, how to migrate applications to a managed environment, and how to update, upgrade, and remove applications. | **Reading**<br><br>"Planning a Managed Environment"<br><br>"Designing a Group Policy Infrastructure"<br><br>"Staging Group Policy Deployments" |
| Deploying a Simple Managed Environment<br><br>Provides an end-to-end deployment scenario based on the information presented in this book. The scenario covers deploying your network infrastructure, preparing your logical infrastructure, creating shared folders, publishing software, configuring user data and settings, using RIS to deploy Windows XP Professional, testing your deployment, and backing up a domain controller. | **Reading**<br><br>All of the chapters in this book |

Table 1.11 describes the contents of *Deploying Internet Information Services (IIS) 6.0*.

**Table 1.11 Deploying Internet Information Services (IIS) 6.0**

| Chapter | Prerequisites |
|---|---|
| Overview of Deploying IIS 6.0<br><br>Describes the high-level processes for deploying IIS Web servers in your organization. | None |
| Deploying ASP.NET Applications in IIS 6.0<br><br>Describes how to deploy ASP.NET applications and individual application files to a Web server. | **Task**<br><br>You must first deploy Windows Server 2003 with the default options, and enable IIS 6.0 with the default settings. |
| Securing Web Sites and Applications<br><br>Describes how to configure IIS 6.0 to provide optimal security for Web sites and applications hosted on the Web server. Topics include reducing the attack surface of the Web server, preventing unauthorized access, and configuring user authentication. | **Task**<br><br>You must first deploy Windows Server 2003 with the default options, and enable IIS 6.0 with the default settings.<br><br>**Reading**<br><br>The other chapters in this kit that address security concerns |
| Ensuring Application Availability<br><br>Information in this chapter describes how to achieve high application availability by setting availability goals, how to configure IIS to achieve those goals, and how to test your applications to ensure that they perform to meet the availability goals you have set. | **Task**<br><br>You must first deploy Windows Server 2003 with the default options, and enable IIS 6.0 with the default settings.<br><br>**Reading**<br><br>"Securing Web Sites and Applications" |
| Upgrading an IIS Server to IIS 6.0<br><br>Describes how to upgrade existing Web servers running IIS 4.0 on Windows NT 4.0 and IIS 5.0 on Windows 2000 to IIS 6.0 on Windows Server 2003. Topics include how to prepare for the upgrade, how to perform the upgrade, and how to further customize the configuration of IIS 6.0 after performing the upgrade. | None |

| | |
|---|---|
| Migrating IIS Web Sites to IIS 6.0<br><br>Describes how to migrate Web sites and applications from existing Web servers running Internet Information Services (IIS) 4.0 or IIS 5.0 to a freshly installed Web server running IIS 6.0 and Windows Server 2003 manually or by using the IIS Migration Tool. Explains how to prepare for and perform the migration and how to further customize the configuration of IIS 6.0 after migration is complete. | **Reading**<br><br>The documentation for the IIS Migration Tool, which is installed with the tool. To find the documentation, on the **Start** menu, point to **All Programs**, point to **IIS Resources**, and then click **IIS 6.0 Tools Documentation**. |
| Migrating Apache Web Sites to IIS 6.0<br><br>Describes how to migrate Web sites from existing Web servers running Apache on Linux to a freshly installed Web server running IIS 6.0 and Windows Server 2003 by using the Apache to IIS 6.0 Migration Tool. | **Reading**<br><br>The documentation for the Apache to IIS 6.0 Migration Tool, which is installed with the tool. To find the documentation, on the **Start** menu, point to **All Programs**, point to **IIS Resources**, and then click **IIS 6.0 Tools Documentation**. |

Table 1.12 describes the contents of Automating and Customizing Installations.

**Table 1.12 Automating and Customizing Installations**

| Chapter | Prerequisites |
|---|---|
| Choosing an Automated Installation Method<br><br>Provides guidelines for choosing which automated installation solution is best for a particular large-scale operating system rollout. Three automated installation solutions are evaluated: unattended installation, image-based installation with the System Preparation (Sysprep) tool, and Remote Installation Services (RIS) installation. | **Task**<br><br>You must first design your network, Active Directory, and servers (such as file servers, terminal servers, or print servers). |
| Designing Image-based Installations with Sysprep<br><br>Describes how to plan and design disk-imaging deployment solutions, including how to use Sysprep to perform large-scale operating system rollouts. Also explains how to assess your current environment, define disk image requirements, and design the overall deployment process. | **Task**<br><br>You must first design your network, Active Directory, and servers (such as file servers, terminal servers, or print servers). |
| Designing Unattended Installations<br><br>Describes how to plan and design unattended installation solutions for large-scale operating system rollouts. Explains how to perform unattended clean installations from CD or from a distribution folder, and how to perform unattended upgrade installations. Also describes how to assess your current environment, decide whether to upgrade or to perform clean installations, and design the overall deployment process. | **Task**<br><br>You must first design your network, Active Directory, and servers (such as file servers, terminal servers, or print servers). |
| Designing RIS Installations<br><br>Describes how to plan and design RIS deployment solutions for large-scale operating system rollouts. RIS solutions that are image-based (Riprep solutions) and solutions that are distribution-folder-based (Risetup solutions) are both discussed. | **Task**<br><br>You must first design your network, Active Directory, and servers (such as file servers, terminal servers, or print servers).<br><br>**Reading**<br><br>"Deploying DHCP" in *Deploying Network Services* |

| | |
|---|---|
| Migrating User State<br><br>Describes how to migrate user state, including identifying the best methods for collecting, storing, and restoring user state data and deciding which user data and settings to migrate. Other topics include preparing a migration plan that addresses storage and security requirements; identifying potential registry, drive, and domain changes; and educating users. | None |

# Testing and Conducting the Pilot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the early phases of the project, you test the components of your Windows Server 2003 implementation to validate your design, as described in the functional specification. Later, you will test the interoperability of components to ensure that all of the elements of the implementation work together, and you will ensure that all the applications used in your organization are compatible with Windows Server 2003.

After verifying your Windows Server 2003 implementation in your test lab, conduct at least one pilot deployment before deploying the operating system into your production environment. In a pilot, you deploy a limited number of computers in a controlled environment, evaluate the results, and fix problems. Deploy successive pilots until you determine that Windows Server 2003 is ready for full deployment.

In addition to testing your implementation of Windows Server 2003, conducting a pilot provides an opportunity to test your deployment plan and the deployment processes. It helps you to determine how much time is required to install Windows Server 2003, and the personnel and tools needed to do so. It also provides an opportunity to train support staff and to gauge user reaction to the new operating system, so that you can anticipate support needs.

After your pilot is stable, meet with management to assess the functionality of the new Windows Server 2003 infrastructure and to receive approval to continue with deployment.

For more information about testing your Windows Server 2003 deployment plan, see "Designing a Test Environment" in this book. For more information about conducting a pilot, see "Designing a Pilot Project" in this book. For more information about testing the compatibility of your applications with Windows XP Professional and Windows Server 2003, see "Planning and Testing for Application Deployment" in this book.

# Rolling Out the Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After testing your implementation in the lab and conducting a pilot program to refine your plan and further test your implementation, you are ready to deploy Windows Server 2003 throughout your enterprise. For some organizations, the initial pilot project is the first phase of the rollout. Others remove their pilot project installations and start their production rollout with clean installations. During the production rollout, testing and support activities remain important while iterative cycles of deployment, testing, validation, and support become the primary focus.

After Windows Server 2003 has been rolled out, responsibility for maintaining the new operating system infrastructure is formally turned over to the operations and support groups. At this point, you might want to develop a project closeout report, which includes final versions of all major project deliverables, sign-off by the team and management, and a summary of the next steps that need to be taken.

You might also decide to conduct a project review. Use the project review to objectively assess the strengths and weaknesses of your entire project and to analyze how you could improve future infrastructure deployments.

# Additional Resources for Planning for Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Related Information**

- "Designing a Test Environment" in this book for more information about testing your deployment.
- "Designing a Pilot Project" in this book for more information about piloting your deployment.
- "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit for information about upgrading Active Directory.
- "Planning a Secure Environment" and "Designing an Authentication Strategy" in *Designing and Deploying Directory and Security Services* of this kit for information about how to secure your Windows Server 2003 environment.
- "Designing Unattended Installations" in *Automating and Customizing Installations* of this kit for information about upgrading.
- The *Microsoft® Windows® Server 2003 Deployment Kit* for more information about the tasks and subtasks you must perform to deploy the various Windows Server 2003 technologies.
- "Planning Deployments" in *Microsoft® Windows® XP Professional Resource Kit Documentation* (or see "Planning Deployments" on the Web at http://www.microsoft.com/reskit) for information about planning Windows XP Professional deployments.
- *Software Project Survival Guide* by Steve McConnell, 1998, Redmond, Washington: Microsoft Press.
- The "MSF Team Model" and "MSF Process Model" white papers in the MSF Resource Library, which is available from the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Windows Server 2003 Features link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for a list of Windows Server 2003 features.
- The Microsoft Systems Management Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The at http://www.microsoft.com/reskit, or see the MSDN Scripting Clinic link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about scripting.

# Designing a Test Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy Microsoft® Windows® Server 2003 or Microsoft® Windows® XP Professional operating systems, even in a pilot, you need to test your proposed design in an environment that simulates and protects your production environment. It is critical to the success of your deployment to develop an environment specifically for testing. In this test environment, you can test hardware, operating systems, or applications designed to run together before introducing them into your production environment.

**In This Chapter**


Overview of Designing a Test Environment

Creating a Test Plan

Planning the Test Lab

Designing the Test Lab

Developing the Test Lab

Designing Test Cases

Conducting Tests

Using the Test Lab After Deployment

Additional Resources for Designing a Test Environment


**Related Information**

- For more information about planning application testing, see "Planning and Testing for Application Deployment" in this book.
- For more information about planning a pilot in your production environment, see "Designing a Pilot Project" in this book.

# Overview of Designing a Test Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A key factor in the success of your Windows Server 2003 or Windows XP Professional deployment project is thorough testing in a test environment that simulates your production environment as closely as possible. A *test environment* consists of a test lab or labs and includes test plans that detail what you will test and test cases that describe how you will test each component.

The test lab can be comprised of a single lab or of several labs, each of which supports testing without presenting risk to your production environment. In the test lab, members of the testing team can verify their deployment design assumptions, discover deployment problems, and improve their understanding of the Windows Server 2003 and Windows XP Professional technology. Such activities reduce the risk of errors occurring during deployment and minimize downtime in your production environment.

You might already have a lab that you use for testing, or maybe you hope to build a new test lab for your Windows Server 2003 deployment project. Regardless of your current situation, you need to think through and clearly define your goals for the test lab and its long-term purpose.

# Test Environment Design and Testing Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many organizations divide their testing teams into two subteams: the design team and the deployment team. The testing design team is responsible for the first three phases of the testing process: creating a test plan, planning the test lab, and designing the test lab. The design team collects information that is vital to the deployment process, identifies immediate and long-term testing needs, proposes a test lab design, or recommends improvements to the existing test lab.

The testing deployment team completes the process by implementing the design team's decisions and performing tests. The deployment team's responsibilities include developing the test lab, designing test cases, conducting tests, and using the lab after deployment.

Figure 2.1 illustrates the process of designing a test environment.

**Figure 2.1 Designing a Test Environment and Conducting Tests**



# Using the Test Lab Throughout the Project Life Cycle

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A well-designed test lab provides a controlled environment for the full range of testing required throughout the project life cycle — from experimenting with the technology through fine-tuning the rollout process. Figure 2.2 lists some of the tasks that can be performed in the test lab and indicates the Microsoft Solution Framework (MSF) phases during which each of these activities might occur. The time frames are estimations and might vary from deployment to deployment.

**Figure 2.2 Role of the Test Lab in the Project Life Cycle**



# Setting Up a Preliminary Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need a preliminary test lab early in the project's planning phase for learning about the product, proving concepts, testing various scenarios against your business model, and validating solutions. During this phase, your project team can use the preliminary test lab to establish your basic namespace design and high-level deployment plan, which you can then use as a baseline for further testing and development. By using the preliminary test lab as a baseline configuration and then adding functionality in stages, you can avoid the problems that occur when component designs are developed independently of each other.

If you do not already have a lab, it is important to start building one as soon as possible. If you plan to build a permanent test lab, you might need to obtain management approval and funding that is independent of your deployment project. If so, start the approval process as early as possible.

You can build a preliminary test lab with just two or three servers and client computers. If your organization has a lab, you can set up and configure your computers there. If you do not have a lab, an office that will accommodate this minimal configuration will do.

After you have selected a site for the preliminary lab, start assessing hardware requirements, reconfiguring existing test lab equipment, and purchasing or recycling hardware for the lab as needed. Then, as you refine the lab design, start putting together the pieces for your formal test lab. As you make decisions regarding hardware, software, and personnel requirements for testing, be sure to document them in your test plan. For more information about the test plan, see "Creating a Test Plan" later in this chapter.

# Creating a Test Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Early in the deployment planning phase, the testing team creates a test plan. The *test plan* defines the objectives and scope of the testing effort, and identifies the methodology that your team will use to conduct tests. It also identifies the hardware, software, and tools required for testing and the features and functions that will be tested. A well-rounded test plan notes any risk factors that jeopardize testing and includes a testing schedule.

If your testing team is divided into technology subteams, each subteam should develop a test plan for that team's specific technology area. For example, the networking team would write a test plan for testing networking features. All members of each subteam should review and approve its team's test plan before it is integrated into the general test plan.

Figure 2.3 illustrates the tasks you must perform to create the test plan.

**Figure 2.3 Creating a Test Plan**




You will use the test plan to design your test lab and to develop test cases. For information about test cases, see "Designing Test Cases" later in this chapter.

# Defining Testing Scope and Objectives

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the scope and objectives section of the test plan, the testing team describes specifically what you want your testing to accomplish. For example, the objective of one team was to migrate the Microsoft® Windows NT® Server 4.0 operating system environment to a Windows Server 2003 environment, component by component, keeping the access control lists (ACLs) and Exchange permissions intact. Another team's objective was to provide a means to measure network traffic and observe server performance during specific directory service tasks.

Also, you need to define the scope of your testing by identifying what you will test and what you will not. For example, you might limit your testing of client computer hardware to the minimum supported configurations or to the standard configurations.

# Defining Testing Methodology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Describe the general methodology that your team will use for testing. For example, your approach to testing schema changes might be to configure an isolated domain in the test lab where schema changes can be applied without affecting other lab tests. This section of the test plan might address the following:

- The domain architecture used for testing
- The tools and techniques used to conduct the tests or to measure results
- Automated techniques you plan to use during testing

For information about methodologies for testing application compatibility, see "Planning and Testing for Application Deployment" in this book.

# Identifying Required Resources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Identify the resources required for the test lab:

**Hardware** List the hardware that you need to conduct tests. For example, identify the standard configurations you plan to support for client computers. Include components such as video cards, modems, and external drives.

**Software** List the software that you need to test the compatibility of your applications with Windows Server 2003. For example, include Microsoft® Systems Management Server (SMS) or other server-based products that you need for testing. For more information about software required for application compatibility testing, see "Planning and Testing for Application Deployment" in this book.

**Databases** Include databases that you need to prepare for testing applications. Also include a description of the resources that you need to populate the databases, such as personnel and business data.

**Personnel** Identify the number of testers you need and the skill level required. Include consultants and other support personnel, as necessary.

**Training** Specify the Windows XP Professional or Windows Server 2003 training that your testers need to complete prior to testing their assigned applications or technologies.

**Tools** Include all tools or scripts that you need to automate testing and to track test results. For example, if you do not have a second test lab that you can use for testing wide area network (WAN) links, include link simulators. For more information about developing a test result tracking system, see "Developing an Incident-Tracking System" later in this chapter.

# Identifying the Features and Functions to Test

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

List all the features or aspects of features that need to be tested. This part of the test plan describes what to test, not how to test.

The following is an example of a feature test description:

- Test 1 — Trust retention

  Description: All trusts to and from a domain should be retained when the domain controllers are upgraded to Windows Server 2003. Use the Domain Tree Manager to view the trusts. If the trusts do not appear, then the test failed.

Note that the description does not include instructions on how to perform the test.

Include tests that verify or address:

- The functionality of each feature and service that you will implement.
- Interoperability with existing components and systems in the production environment, both during the phase-in period, when there is a mix of old functionality and new Windows Server 2003 functionality, and after the Windows Server 2003 environment has been rolled out.
- Hardware and driver compatibility for every type of client computer that will be running Windows XP Professional.
- Application compatibility for every application that will run on Windows XP Professional.
- Application compatibility for every server application that will run on Windows Server 2003.
- Optimization of configurations, such as those for standardized desktops on client computers.

For more information about planning for application compatibility testing, see "Planning and Testing for Application Deployment" in this book.

In addition, list:

- *Baselines* (a range of measurements derived from performance monitoring that represents acceptable performance under typical conditions) for performance monitoring.
- Baselines and stress tests for capacity planning.
- Procedures for deployment and post-deployment administration, such as procedures for upgrading a client computer and for backing out of a faulty rollout process.
- Required tools and utilities.

During the developing phase, the testing team will develop test cases that describe how to perform each test listed in the test plan. For more information about test cases, see "Designing Test Cases" later in this chapter.

# Identifying Risk Factors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Describe the risk factors that could prevent the successful completion of all required tests. For example, you might find that the test lab is behind schedule, or that required hardware or software is unavailable, or that testers are working on other projects or need additional training. After you have identified the risk factors, decide what you will do to avoid or mitigate each risk.

# Establishing a Testing Schedule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Draft a preliminary schedule that includes each test listed in the test plan. The schedule can help you coordinate test lab use among subteams. Assign a team member, ideally the test lab manager, if your team has one, to maintain and update the lab schedule. Having an up-to-date schedule is critical when unscheduled lab requests are submitted.

# Planning the Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A *test lab* is a network that is designed for testing and that is isolated from the corporate network. In the test lab, you can verify that components and features work together in an integrated environment that simulates your target production environment. The simulated environment should reflect both the phase-in period, when you have a mix of functionality, and the end of the project, when you have completely implemented Windows Server 2003 and Windows XP Professional.

Begin the lab planning process by evaluating test lab models, and then select the appropriate site for your lab.

Figure 2.4 illustrates the activities you undertake when planning the test lab.

**Figure 2.4 Planning the Test Lab**




**Note**

- If you already have a permanent test lab that you plan to use for testing your Windows Server 2003 design, you can to go directly to "Designing the Test Lab" later in this chapter.

# Evaluating Test Lab Models

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As your team considers its test lab needs, you will need to make several basic decisions. Can you upgrade an existing lab or will you need to develop a new one? Will an ad hoc lab do, or can your organization afford to build a permanent lab? Will your organization be better served by a centralized lab or by several labs that are distributed across your organization's various sites?

Your team might determine that one type of lab fits your goals best, but circumstances might dictate establishing another type of lab. For example, you might see advantages to having a permanent test lab that you can use for testing software patches and upgrades, but realize that your organization does not have the budget to build and maintain a permanent lab. Although you need to weigh the possible outcomes of using the opposing model, you might come up with creative ways to support your ideal solution.

Ask yourself:

- How will the decision affect test quality?
- How will the decision affect team training and design support?
- Will the test lab also benefit other existing projects?
- Can other projects combine efforts and budgets to share a test lab?
- Can you build the test lab in stages, starting with the most essential components and then adding to it as the budget allows?
- Will hardware vendors agree to a special arrangement?

  For example, can the vendors provide equipment on loan before it is purchased, or will they give you the equipment in return for using your organization's name for marketing purposes?

# Upgrading an Existing Test Lab vs. Building a New Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization already has a test lab, your team should decide if the lab can be upgraded to meet your current and future testing needs. You might find that to accommodate Windows Server 2003 deployment testing you simply need to perform minor hardware and software upgrades. For example, you might need to upgrade your servers so that they all have the same amount of memory and hard disk capacity and use the same processor type and speed as the servers you plan to deploy.

If upgrading is not feasible, you will need to seek approval for creating a new test lab. To justify the investment to your project sponsors, take a broad view of the associated costs. Testing in a suitably equipped lab leads to cleaner implementations and reduced support costs. By using the test lab to develop operational efficiencies, such as automated administrative tools and remote installation procedures, you can reduce your organization's total cost of ownership. When viewed over time, therefore, the costs of building and maintaining a test lab dedicated to Windows Server 2003 deployment are likely to be dramatically lower than the costs of fixing problems in production, redeploying poorly developed or poorly tested solutions, or managing the production environment with resource-intensive processes.

In organizations that build separate test labs for different projects, economies of scale often can be realized by consolidating labs. This way, several projects can share in the reduced cost of a single test lab. If your team decides to share a lab, however, do so with a project that has compatible scheduling and equipment requirements.

# Ad Hoc Test Lab vs. Permanent Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many organizations build an ad hoc test lab every time they need test facilities for a new project. Other organizations build a permanent test lab that is scalable for various projects and, after the original project is deployed, use it for change management.

**Ad Hoc Test Lab**
An *ad hoc test lab* is set up for a specific project only. When the project is over, the equipment is redeployed for a different use. For example, lab equipment might be used in the production environment, become part of inventory, or be returned to the vendor from whom it was leased.

Because all the lab equipment is eventually redeployed for other uses, the short-term costs of creating and using ad hoc test labs might be a fraction of those for creating and running a permanent test lab. This view of cost is shortsighted, however, because a new test lab must be built for every project that requires one.

There are additional disadvantages of using ad hoc labs:

- When a new test lab is required for each new project, time becomes a critical factor. Because teams need the test lab early in a project, issues such as the following arise:
    - Can you obtain the appropriate hardware and software licenses in time?
    - Will hardware or software substitutions lead to inadequate testing?
    - Can you locate the representative equipment from each vendor, including the exact models and versions, needed to adequately test the production mix of hardware and software?
    - Can you reserve the physical space needed for building the network configuration and performing the tests?
    - Will the time spent building and debugging the test lab reduce the time available for actual testing, possibly resulting in incomplete testing?
- When many teams are searching for hardware and software licenses, it becomes difficult to track who is using what and who authorizes purchases. The resulting lack of accountability can lead to excess expenditures and increased costs.

If you have an ad hoc test lab and have adequate space and funding, consider converting it to a permanent test lab. You might be able to reduce your team's cost for creating a permanent lab by seeking funding from other groups that would benefit from having a dedicated test lab.

**Permanent Test Lab**
The problems associated with creating and using an ad hoc test lab present compelling reasons for building a permanent, formalized lab. A permanent lab has the additional advantage of providing a place to test changes to your production environment after you have deployed Windows Server 2003 or Windows XP Professional, and, as such, it becomes a change management tool. After deployment, your organization could use the lab to test:

- Network upgrades.
- Service pack additions and software patches.
- New business applications that might pose compatibility problems.
- New desktop configurations.
- New hardware platforms.

- New administrative and support processes.
- New client computer management tools.

A fully equipped, permanent test lab that is used for change management:

- Helps keep the test lab adequately equipped.

  By maintaining a permanent lab and using it for change management, your organization might be able to justify the purchase of equipment you need for specific testing requirements. You are also more likely to maintain the mix of equipment that accurately reflects your production environment when you have a permanent lab. Maintaining the proper mix of equipment in the test lab allows thorough regression testing during the change management process.

- Reduces risk to your business.

  Solid testing in a lab that mimics the production environment leads to cleaner implementations.

- Saves time for the project.

  Setup and debugging time are minimized because a lab can be upgraded faster than a new lab can be assembled. The time saving is critical if your lab will be used for ongoing developmental prototyping or for general development.

- Helps establish consistent methodologies.

  With a permanent test lab and continuity in lab management, your team can establish testing processes and techniques that produce consistent results, which can be compared over time.

- Saves costs in the long run.

  When viewed across projects, the cost of a permanent test lab is likely to be more reasonable than that of an ad hoc test lab for which purchases are not tracked or financial accountability is diluted.

# Centralized Test Lab vs. Distributed Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A test lab can be centralized in one location or distributed across numerous physical, or even geographical, locations.

**Centralized Test Lab**
A *centralized test lab* is located in one site. Consolidating all lab resources allows you to develop a more complex and complete lab. Configuring and managing a centralized lab is also easier. However, if your organization is decentralized and you choose to establish a centralized test lab, testers from other areas or regions will need to travel to the lab to perform tests, which is not always practical or cost effective.

If you choose to create a centralized test lab, you might decide to develop several independent, labs within your site for different testing purposes. For example, you might build one test lab for network infrastructure and server testing and another test lab for client computer and application testing. Or, you might have a network infrastructure test lab with multiple locations connected by a WAN for testing the effects of using various link speeds.

If you are deploying Windows Server 2003 and Windows XP Professional at the same time, many factors will influence your decision about whether to develop separate test labs for each or a shared lab. Factors include:

- The complexity of your deployment (such as the need to make changes in your production environment and the new features you plan to implement).
- The size, location, and structure of your project teams.
- The size of your budget.
- The availability of physical space.
- The location of testers.
- The planned use of the test lab(s) after deployment.

**Distributed Test Lab**
A *distributed test lab* consists of several individual labs. Often, the labs are located near the testing team or application developers; because of this proximity to development and testing, more resources are available to address problems that arise during testing. Some organizations distribute their labs across their various business or functional sites, which allows testing to be performed on location. The disadvantage of the distributed lab model is that the cost of maintaining multiple test labs can be substantial.

The following case studies describe two organizations that have chosen to build distributed test labs and how they use them.

**Case study 1: Functional test lab sites**
A large, high-technology hardware manufacturer is organized along functional lines. Its regional offices are located at various geographical sites, selected for proximity to the suppliers and vendors who support each region's particular function. This manufacturer developed a distributed test lab that spans three of its major United States sites, ranging across Southwestern and Western states. Each lab is designed to test the functions and configurations used for that site's business. Each one is a permanent lab that is used for managing change in the production environment.

Eventually the organization plans to expand the test lab to encompass international locations, such as cities in Asia, the Middle East, Eastern Europe, and the United Kingdom. The organization will use these remote sites to design and test factors or solutions for the challenges of operating a global enterprise, such as:

- Connectivity in controlled countries/regions.
- Slow links.
- Intermittent links.

- Multiple languages.
- Multiple time zones.
- International currencies.
- Variations in computer and networking hardware.

**Case study 2: Contingency test lab sites**

A large financial lending organization needs to be prepared in case of a disaster. This organization wants each of its geographically separated sites to be ready to function in the role of a centralized IT department, if necessary. In this organization, each test lab is a permanent change management lab that is also used for disaster-recovery testing.

In the event of a disaster, production computers at the alternative sites would be used to perform IT department functions. To maintain a state of readiness, the organization performs tests in each lab to ensure that all the needed hardware and software components are available and function properly at all the sites. Processes that are tested include:

- Loading applications and databases.
- Setting configurations.
- Running applications.

# Selecting a Test Lab Location

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

our decisions about the type of test lab your organization needs and its location are likely to be interrelated. The location of a permanent test lab to be used by a wide variety of groups requires more consideration than the location of an ad hoc lab to be used by a few groups. And, if you plan to use the lab long term for change management, room for growth will be an important issue. As you weigh all available options, consider the following:

- What test lab facilities already exist? How adequate are they? How easily can you modify them to fit test requirements?
- Can you consolidate existing test labs?
- What is the scale and complexity of the implementation?
- Does the test lab need to connect to the production network or to other test labs? If it needs to connect to the production network, how will you regulate the connection and configure routers to protect the production network?

  For more information about connecting the test lab to the production network, see "Simulating the Proposed Server Environment" later in this chapter.

Additional issues to consider when selecting a test lab location or locations include how much space and how many sites you need and whether the sites are accessible, secure, and appropriate for housing valuable testing equipment.

**Space** The complexity of your organization's production environment influences the complexity of the test lab and, therefore, its space requirements. Factors in your current and proposed production environment that can determine the complexity — and, therefore, the space requirements — of your test lab include:

- The number and combination of functions and features that you plan to implement.
- Whether you plan to implement a domain that spans multiple sites.
- Whether you plan to implement a virtual private network (VPN).
- The amount of variability in your production environment.
- Whether you have — and plan to deploy — standard equipment, applications, and configurations in your production environment or whether you will use equipment from several vendors or several models, versions, and configurations.
- The level of complexity in your network configuration.
- Whether you have more than one type of topology in your production network. For example, do you plan to have interfaces between Windows Server 2003 and mainframe, Macintosh, UNIX, or Linux systems?

In addition to factors in your production environment, some test situations also can affect the complexity of your test lab. For example, you might want additional servers so that you can isolate certain types of testing. For more information about isolating tests, see "Designing Domains for Testing" later in this chapter.

Space requirements also influence the number of people that are able to participate in testing. Consider how many testers you need to accommodate at one time, and configure the test lab accordingly.

**Number of locations** Your test lab might need to consist of multiple interconnected locations so that you can test the effect of geographically separated network segments. For example, if you plan to implement Microsoft® Active Directory® directory service with multiple Active Directory sites, you might want to test replication over a WAN or Internet connection similar to your production environment. For more information about Active Directory sites and replication, see "Designing the Site Topology" in *Designing and Deploying Directory and Security Services* of this kit.

In other cases, you might want to maintain multiple independent test labs for different uses. For example, you might want a separate test lab for application testing or perhaps separate test labs for Windows Server 2003 and Windows XP Professional testing.

**Accessibility** The test lab should be accessible to all of the groups using it. If you implement a program in which people from outside the project team come to test their own applications, the test lab should have facilities, such as parking, to accommodate visitors.

**Security** Ensure that you can physically secure the test lab to protect your equipment from unauthorized use. You can restrict access to your test lab by installing a keypad that requires a Personal Identification Number (PIN) or a smart card reader.

**Environmental conditions** Any site you consider for your test lab should be clean and should be able to maintain temperature and humidity levels appropriate for your team and your valuable equipment. The lab site also should be able to accommodate your power, cabling, and network connectivity requirements.

# Designing the Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The lab planning process includes documenting the proposed test lab configuration. To design a lab that mimics your future production environment, you will need to simulate the proposed server and client environments as closely as possible. You also need to design domains for testing.

Figure 2.5 illustrates the test lab design process.

**Figure 2.5 Designing the Test Lab**



# Gathering Information About the Current and Proposed Environments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin designing your test lab, you need to gather detailed information about both your current and proposed environments.

For your current environment, you need:

- The network design (logical and physical).
- An inventory of existing hardware (servers, client computers, and portable computers).

You can use the hardware inventory that you conducted to assess your current IT environment, as described in "Planning for Deployment" in this book. Or, you can use SMS to collect this information.

For your proposed environment, you need to compile the following from the master project plan:

- The proposed Windows Server 2003 design.
- A list of features to evaluate and explore.
- A list of hardware required for use with Windows Server 2003. Though this list might evolve during testing, you need a fairly complete list to equip the lab initially.
- A list of administrative tools required (Windows Server 2003, third-party, and custom built).
- A list of the upgrades, such as service packs, drivers, and BIOSs, that you need to install to be ready for a Windows Server 2003 deployment.

You also will need a high-level understanding of the Windows Server 2003 features and functions that you intend to include in your proposed design.

# Documenting the Test Lab Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Design the test lab in a way that allows you to test as much of the proposed logical and physical production environment as possible, including computer hardware, network topology, WAN connections, domain architecture, services, databases, business applications, administrative tools, security model, application deployment methodology, and network server storage methods.

Design the client computer portion of the test lab so that you can test the same functions and features that you currently use and that you plan to use in your production environment. Include the same types of hardware, applications, and network configurations.

As you design the test lab, document it with both a text description and a diagram. Post the diagram in the lab to provide easy access to lab information and to keep lab users up-to-date on design changes.

For additional information about components that you might need to include in your test lab to fully test Windows Server 2003 and Windows XP Professional, see "Simulating the Proposed Server Environment" "Simulating the Proposed Client Computer Environment" and "Designing Domains for Testing" later in this chapter.

# Test Lab Description

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Include information about the following in your test lab description:

- The domain structure, including:
  - DNS structure
  - The forest and tree hierarchy
  - Group Policy objects (settings and where they apply)
  - The method you will use to populate user account data
  - The site layout
  - Trust relationships (transitive and explicit)

  For more information about domain structure considerations for testing Windows Server 2003, see "Simulating the Proposed Server Environment" and "Simulating the Proposed Client Computer Environment" later in this chapter.

- Domain controllers, including:
  - Primary domain controllers (PDCs) and backup domain controllers (BDCs), if you are migrating from Windows NT Server 4.0
  - Servers to be promoted to domain controllers, if you are migrating from any other operating system
  - Replication topology
  - Global catalog layout
- Member servers, including the services that will run on them. Examples of services include:
  - DHCP topology
  - WINS topology
- Client computers, including:
  - Computer make and model
  - Amount of memory
  - Processor type and speed
  - Hard disk capacity
  - Graphics cards (type, resolution, and color depth)
  - Network adapter (model and speed)

- Wireless network adapter

For more information about client computer configurations for testing Windows Server 2003, see "Simulating the Proposed Client Computer Environment" later in this chapter.

- Requirements for specific tests, including:

  - Side-by-side testing (tests performed on identical computers, one running the current operating system and one running the new operating system)
  - Mixed-mode testing (tests performed while migrating or transitioning from one domain to another)
  - Native-mode testing (tests performed while maintaining your existing domain structure)
  - Functional level testing (tests performed at levels of domain and forest functionality)
  - Dial-up and other remote testing
  - Interoperability testing (with UNIX, Linux, mainframes, and other operating systems)
  - Replication and Active Directory site testing
  - WAN link testing

For more information about Active Directory functional levels, see the *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit).

# Test Lab Diagrams

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The test lab diagram should show both the logical and the physical structure of the lab. Depending on the complexity of your test lab network, the logical and physical diagrams can be combined into one diagram or they can remain separate.

**Logical Diagram**

Document the following in the logical diagram:

- The domain hierarchy, including forests and trees. For more information about domain structure considerations for testing Windows Server 2003, see "Simulating the Proposed Server Environment" and "Designing Domains for Testing" later in this chapter.
- Domain names.
- Active Directory sites.
- Special service servers (domain controllers, global catalog servers, DNS servers, DHCP servers, and WINS servers). For each server, include the following information:
  - Computer name
  - IP address
  - Server function
- Transitive trusts.
- Explicit or one-way trusts.

Figure 2.6 is an example of a logical diagram of a test lab. This test lab has one tree, which consists of a root and three child domains. Double-headed arrows indicate transitive trusts between Windows Server 2003 domains. The Windows NT Server 4.0/Microsoft® Windows® 2000 operating system domain has explicit one-way trust relationships with the Windows Server 2003 tree. This test lab does not have Active Directory sites. The test lab includes domain controllers (noted in the figure as DCs), some of which are also DNS servers supporting the dynamic update protocol, DHCP servers, and one global catalog server.

**Figure 2.6 Example of a Test Lab Logical Diagram**




**Physical Diagram**

Document the following in the physical diagram:

- Network components, such as:
    - Routers and bridges
    - Hubs
    - Link simulators
    - Proxy servers
    - Sniffers and traffic generators
    - Analog and Integrated Services Digital Network (ISDN) lines
    - LAN (local area network), WAN, and Internet connections and speeds
- Servers (include domain name, computer name, IP address, and server function)
- Client computers (include computer name and, if you use static addressing, IP address)

Figure 2.7 is an example of a physical diagram for the same test lab illustrated in the logical diagram in Figure 2.6. This diagram shows the three subnets for the three child domains. Each subnet has both a Windows XP Professional client computer and another type of client computer. The test lab uses simulated frame relay links and has a UNIX server.

**Figure 2.7 Example of a Test Lab Physical Diagram**




After you have completed the lab description and the logical and physical designs, expand your lab design to include all the components that you will need to fully simulate the Windows Server 2003 and the Windows XP Professional environments, as described in the following sections.

# Simulating the Proposed Server Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section lists some components you might need to include in your lab to fully test Windows Server 2003. The recommendations presented here might not apply to all Windows Server 2003 implementations. Focus on the considerations that apply to your design.

**Server hardware and drivers**
Use the same type of hardware components and drivers that you use, or plan to use, on servers in the production environment. Be sure to obtain an updated BIOS that is compatible with Windows Server 2003.

**Services and configurations**
Use the same services and configurations that you will use in the actual deployment. For example, duplicate the Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and Windows Internet Names Service (WINS) configurations.

**User accounts**
If you are migrating from Windows NT Server 4.0, set up your domain controllers as replicas of your production domain controllers, using copies of the production user accounts. You can use the Clone Principal tool (Clonepr.dll) to copy production users to your test domain. For more information about the Clone Principal tool, see the Microsoft TechNet link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, and then search for Tools. Coordinate with your IT security division whenever you copy production data to test lab databases. For more information about security, see "Planning a Secure Environment" in *Designing and Deploying Directory and Security Services* of this kit.

**Domain structure**
If you are implementing Active Directory, simulate the domain hierarchy; for example, include a forest with multiple trees, a tree with parent and child domains, and transitive and one-way trust relationships, as appropriate. Reflect your IT administration structure (either centralized or decentralized). Include Active Directory sites as appropriate. For more information, see "Designing Domains for Testing" later in this chapter.

**Group Policy settings**
Duplicate the Group Policy settings from your production environment in the test lab. Properly testing Group Policy settings in the lab helps you identify whether implementing new features or technologies will require that you make changes to your Group Policy settings.

**Server strategy**
Include domain controllers, file and print servers, application servers, Web servers, database servers, and other utility servers that are, or will be, included in your production environment. If you plan to use SMS to deploy Windows Server 2003, include it in the test lab.

**Mixed environments**
To accommodate both the mixed environment that exists during a phased rollout and the Windows Server 2003 environment that will exist after the completion of the rollout, plan to include some native-mode and mixed-mode domains.

By simulating the interim state, you can identify functional problems that might occur during the phased implementation. The servers running operating systems other than Windows Server 2003 should mirror the services they perform in the current production environment. For more information, see "Designing Domains for Testing" later in this chapter.

**Client computer configuration**
Client computers in the test lab should conform to the standard configuration that your organization has established. They should have the same components and peripherals found in your production environment.

Use the same mix of client computers that is present in your production environment. If you plan to deploy Windows Server 2003 first and Windows XP Professional later, include in your test lab the client computer operating system that you will be using until Windows XP Professional is deployed.

If you plan to deploy Windows XP Professional first, test to see how the extended server functionality will be introduced into your environment as the infrastructure is deployed.

If you plan to have a phased rollout, include the same mix of client computers in your test lab that will be present in your production environment. For example, have client computers with Windows NT Workstation 4.0 and client computers with Windows XP Professional.

**Network topology and protocols**
Mirror the network topology and protocols that you use in your production environment as closely as possible. For example, if your production network uses both Ethernet and Token Ring, the test lab should include both.

**WAN connections**
If you have a WAN, the test lab should include routers for testing network latency or for testing domain controller and global catalog replication across the WAN link. If you have the facilities and budget, you might want to set up a secondary test lab at a remote location to test network latency across the WAN link. If you have a multinational organization, place the secondary test lab in a different region to test actual latency problems.

If you do not have a secondary test lab located where you can test the WAN link, you can connect two routers in the same lab with a cable and use a link simulator to test the link.

**Remote connectivity**
In the test lab, provide the same types of remote connectivity — such as Routing and Remote Access service and VPN — that will be included in your production environment so that you can test the Point-to-Point Tunneling Protocol (PPTP), Internet Protocol security (IPSec), Layer 2 Tunneling Protocol (L2TP), and Demand Dial Routing.

**Peripherals**
Include a representative sample of the types of peripherals used in your production environment. For example, include the same types of printers and scanners, along with their associated drivers.

**Interoperability**
If you plan to implement Windows Server 2003 so that it operates with networks or computers using another operating system, mimic the interoperability infrastructure. For example, include connections to mainframe hosts, UNIX systems, or other network operating systems. To keep your test lab configuration manageable, decide which interoperability scenarios are most important to your organization and focus on those.

**Administrative tools**
Include the tools (Windows Server 2003, third-party, or custom built) that you currently use or plan to use for server-based administrative tasks. Test the tools for compatibility and effectiveness in the new environment.

**Fault tolerance techniques**
Test any fault tolerance techniques you plan to use in your production environment. For example, if you plan to use clustering, include a server cluster in the test lab.

**Terminal Services**
If you plan to implement Terminal Services, install the appropriate mix of applications on the server so that you can assess the effect of running applications in a multiuser environment. You might need to modify the default operating environment for some applications to obtain the desired functionality. For more information about hosting applications using Terminal Services, see "Hosting Applications with Terminal Server" in *Planning Server Deployments* of this kit.

**Protecting the production network**
Be sure you isolate the test lab from your corporate network. If you need to provide a connection from the test lab to the corporate network, plan how to regulate and control the connection and be sure to devise a way to quickly terminate the connection.

You can design router configurations to protect the production network. For example, consider using a multihomed router with two network adapters to connect the test lab to the production network for specific, controlled uses. Configure the router so that the production network can access the test lab network, but the test lab network cannot access the production network. This approach protects the production environment from events that might occur in the test lab, but allows a user in production to access test lab resources. For example, you could use this approach to test logon scripts on a test lab server with a limited number of users before moving the scripts to a pilot in the production environment.

# Simulating the Proposed Client Computer Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section lists some components you might need to include in your lab to fully test Windows XP Professional. The recommendations presented here might not apply to all Windows XP Professional implementations. Focus on the considerations that apply to your lab design.

**Client computer hardware**
Include at least one client computer from each vendor and one of each model that will run Windows XP Professional in your production environment. If your organization uses laptops, docking stations, or port replicators, be sure to include components that fully represent the vendors and models. Be sure to obtain a BIOS update that is compatible with Windows XP Professional.

Develop a standard hardware configuration for Windows XP Professional as part of your deployment project. Lab testing can help you define and refine a standard configuration that is compatible with Windows XP Professional. Some of the components that need to be tested for compatibility include:

- Universal serial bus (USB) adapters.
- Compact disc (CD) and digital video disc (DVD) drives.
- Sound cards.
- Network adapters.
- Video adapters.
- Small computer system interface (SCSI) adapters.
- Mass storage controllers.
- Removable storage devices.
- Pointing devices (mouse or trackball).
- Keyboards.

To determine whether a device is compatible with Windows XP Professional, see the Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, and then look up the component. If a device is not included in the Windows Catalog list, contact the device vendor to find out if a driver is available.

Alternatively, you can use the Windows XP Professional CD to check hardware and software incompatibility. To obtain a report that indicates hardware and software incompatibilities and identifies device drivers that need to be updated, insert the Windows XP Professional CD in your drive and select Check System Compatibility. If updated drivers for your devices are not included with Windows XP Professional, contact the device vendors to obtain them.

On computers running Microsoft® Windows® 95 or Microsoft® Windows® 98 operating systems, you can check hardware and software incompatibility by using the log file called Upgrade.txt. Upgrade.txt is located in the Windows Installation folder. On systems running Windows NT Server, you can check hardware and software incompatibility by using the log file Winnt32.log. Winnt32.log also is located in the Windows Installation folder.

**Note**

- Windows XP Professional does not support 16-bit drivers. If your components use 16-bit drivers, you need to obtain 32-bit drivers for those components.

**Network connectivity**

Provide connectivity to the same types of networks that you use in the production environment, such as a LAN, a WAN, or the Internet.

If you plan to use Routing and Remote Access or a proxy network service in the production environment, include these types of connections in the test lab.

**Server-based services**
Configure servers for the services used in the production environment. For example, include services such as:

- DNS, WINS, and DHCP.
- Directory services, such as X.500 and Novell NetWare.
- File sharing.
- Network printing.
- Server-based line-of-business applications, both centralized and decentralized.
- IntelliMirror® management technologies.

Remember to account for administrative services or tools, such as:

- Remote operating system installation.
- Server-based application deployment.
- Tools for managing client computers (SMS, for example).

**Domain authentication**
If your organization uses, or plans to use, domain authentication, simulate your authentication configuration in the test lab. If you are migrating from Windows NT Server 4.0 to Windows Server 2003, plan for authentication in the mixed environment that will occur during the phase-in period.

**Network management services**
Include network services used in your environment, such as Simple Network Management Protocol (SNMP).

**Network protocols**
Use the protocols you plan to use in the production environment. Verify that the protocols you use on client computers work properly before connecting the client computers to the production network.

**Applications**
You need licenses for, and access to, the software for all applications, stand-alone or server-based, that is to be supported on Windows XP Professional client computers.

**Peripherals**
Include a representative sample of the types of peripherals, such as printers and scanners, used in the business environment.

**Server platform interoperability**
Simulate the server platforms to be accessed by Windows XP Professional client computers. You might need to establish connections with the following systems:

- Windows Server 2003
- Windows 2000 Server
- Windows NT Server 4.0
- Mainframes supporting 3270 emulation

- UNIX or Linux
- Other server operating systems

If you have a separate server test lab, consider connecting the client computer test lab to it, instead of installing servers in the client computer lab.

If you plan to deploy Windows XP Professional at the same time as Windows Server 2003, include every type of server that a client computer will be able to access during the deployment period.

**Desktop configurations**
As part of your Windows XP Professional deployment, your organization might decide to evaluate Standard Client configurations and Group Policy objects for managing them. Lab tests can provide information that will help you recommend specific configurations and Group Policy objects to management. If you decide to perform this type of evaluative testing, include side-by-side comparisons of different configurations and Group Policy object settings.

Plan to have enough computers of the same make and model to allow for the side-by-side evaluations. Evaluate client configurations based on performance, ease of use, stability, hardware and software compatibility, functionality, and security model. Evaluate Group Policy objects by verifying that they produce the desired result, particularly when more than one object applies to a configuration, and that the resulting logon time is acceptable.

**Performance**
Use the test lab to start evaluating the effect on your network traffic by testing for changes in baseline traffic patterns when there is limited user activity, for example, when conducting your pilot.

**Production network connectivity**
If you have separate test labs for clients and servers, you need to isolate them both from the corporate network. If you need to provide a connection from the test lab to the corporate network, plan how you will use routers to separate the two networks.

# Designing Domains for Testing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Design the test lab domain structure to provide a consistent setup and configuration so that testers can rely on a standardized infrastructure. For example, allocate a single domain for migration and mixed-mode testing. If you do this, the domain should always be in the mixed-mode state except for scheduled periods when it is rolled back to the prior state to test the migration process. This way, test lab users always know what to expect.

Because some tests alter the lab environment, they can inadvertently affect other tests; for example, server upgrade tests change the services running on the server, which can affect any networking protocols that you are using. Care must be taken to isolate, coordinate, and manage these types of tests. Address issues such as these in the test lab domain design.

**Important**

- If you destroy the child domains and intend to reuse the same names upon reinstallation, you must clean up the metadata. To avoid this problem, use the Active Directory Installation Wizard and remove Active Directory from the domain controller prior to destroying the child domains.

Some of the types of testing that might require separate domains are:

- DNS
- Native-mode
- Mixed-mode
- Functional level
- Migration process
- Business data replication

**Case study: Designing test domains**

A large manufacturing company created a root domain with four child domains. This domain structure allowed the project team to use a separate domain to test:

- Windows Server 2003 functionality, including printing, in a native-mode domain.
- VPN.
- Mixed-mode interoperability and migration process.
- Microsoft® Exchange Server integration with Windows Server 2003.

In addition, an isolated domain allowed the team to test DNS without affecting other testing.

Figure 2.8 illustrates the logical domain structure of the test lab.

**Figure 2.8 Example of a Test Lab Logical Domain Design**




After you have designed and documented your test lab, have your project subteams review the plan to ensure that all the necessary requirements have been addressed.

# Developing the Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you finish the test lab planning process and receive management approval, you can build the lab. To ensure smooth operation of the lab, assign a test lab manager and develop test lab guidelines and procedures.

Figure 2.9 illustrates the steps you take to develop the test lab.

**Figure 2.9 Developing the Test Lab**




If you plan to rebuild the test lab periodically as your testing focus changes, consider using tools or products such as SMS to manage test lab changes. Also consider using Remote Installation Services (RIS), Sysprep, or unattended installations to rapidly make changes to the client computer configurations in the test lab. Active Directory Service Interfaces (ADSI) and Windows Script Host can help you quickly create, delete, or change users, groups, and organizational units in your test lab environment.

For more information about using RIS to automate client computer installations, see "Designing RIS Installations" in *Automating and Customizing Installations* of this kit. For more information about using Sysprep to automate client computer installations, see "Designing Image-based Installations with Sysprep" in *Automating and Customizing Installations* of this kit. For more information about using unattended installations to automate client computer installations, see "Designing Unattended Installations" in *Automating and Customizing Installations* of this kit.

# Assigning a Test Lab Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your test lab will be permanent or will be used by many groups, you will need to assign someone to manage it. This is particularly important if the test lab is to be used by several groups for testing during the change management process.

Smaller test labs or labs used by a single team might not need a dedicated manager. For such a lab, consider the tradeoffs between — and the costs of — either hiring a dedicated lab manager or assigning the role to a project team member. Evaluate whether assuming lab management duties in addition to other project responsibilities might be too burdensome for a team member. Even if you decide not to assign a dedicated, full-time manager, select someone to be responsible for the test lab.

A test lab manager is responsible for:

- Establishing and enforcing test lab policies and procedures.
- Establishing a test lab budget for support costs.
- Scheduling any configuration changes to servers or clients that might be required by certain tests, and communicating these requirements to other test lab users.
- Developing a bug or problem escalation plan.
- Developing and managing an incident-tracking system.
- Developing and monitoring the change control process.

  The change control process defines who is allowed to make changes to the test lab environment. For more information about change control, see "Using the Test Lab After Deployment" later in this chapter.
- Setting up an inventory control system.
- Maintaining test lab documentation (such as test lab descriptions, diagrams, and guidelines and procedures).
- Managing network taps and server capacity and configurations.
- Managing hardware and software configurations and updates.
- Establishing physical security.

  The test lab manager takes measures to prevent unauthorized use of lab equipment and manages lab access with keys or electronic locks.
- Implementing a preventive maintenance program for equipment.
- Establishing an approval process for removing equipment (for lending it to another group, for example).
- Making periodic server backups.
- Ensuring the test lab is kept clean and orderly.

# Building the Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To build the test lab**

1. Acquire hardware and software, including administrative tools.

2. Install and configure routers, switches, wireless and optical devices. Label all network cables.

3. Install and configure all servers.

   If you are redeploying servers, you might need to upgrade them to accommodate Windows Server 2003. Use the same memory, disk capacity, and CPU speed that you plan to deploy. Be sure to check for viruses and to defragment the hard disks. Partition the hard disks in the same way you plan to partition them during deployment.

   Install the appropriate operating system, either Windows Server 2003 or the operating system from which you plan to upgrade.

   If you are upgrading your domain controllers, you should consider using a disk-imaging program to create images of the domain controllers before you upgrade them. Test the images and store them in a safe place. If necessary, you can use the disk images to restore the servers to their original state. By making reliable disk images, you avoid intruding on your production environment if your upgrade process changes or fails.

   If you are purchasing new equipment, burn in the components for two or three days to be sure that they are working properly.

4. Test all network connections.

   Testing the network before you install the servers makes it easier to locate and solve problems.

5. Install application software when you are ready to include it in testing. Install all server-based applications, such as SMS, and business applications that are included in your production environment.

   Build or load copies of the associated databases. Install tools for testing and administration. If you are planning to verify network traffic or test performance, you might want to include a hardware sniffer or a software sniffer.

   For more information about testing tools, see "Planning and Testing for Application Deployment" in this book.

6. Install and configure all client computers.

7. If you are implementing Terminal Services, install a representative set of applications so you can test concurrent users.

8. If you plan to create backups for restoring baseline configurations, set up the baseline configurations and make the backups.

   For example, if you plan to upgrade Windows NT Workstation 4.0 to Windows XP Professional instead of performing a clean installation, back up a Windows NT Workstation 4.0 client computer that is running your standard set of applications.

   The baseline configurations should include all the service packs that are supported in your environment. Be sure to test and document the restore process.

9. Test the individual components in the lab to isolate problems that are not related to Windows Server 2003 and Windows XP Professional deployment.

When testing starts, you will want to spend your time debugging deployment problems, not fixing problems with the test lab.

10. If you need to provide connectivity to the production network, configure and test routers to isolate the test lab from production.

**Important**

- As the team builds and rebuilds the test lab, be sure to document, in chronological order, every change made to servers and client computers. This documentation can help resolve problems that might arise later and help team members understand why a specific computer behaves as it does over time. It also can help you reverse recent changes to solve a short-term problem.

After your team has built the test lab, establish lab guidelines and procedures. Then your team can begin testing.

# Developing Test Lab Guidelines and Procedures

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Test lab users need to know when they can perform their tests, whether their testing will disrupt other tests, and what the current state of the lab is. For example, if a domain in the lab is used for testing the migration process, in addition to mixed-mode functionality, test lab users need to know if the computers are ready for an upgrade or if they already have been upgraded. Your lab manager or an assigned team member should establish a communication system to disseminate information about the availability and state of the test lab and should establish guidelines for lab use.

Guidelines should be easy to follow and to remember, and should address each of the following issues.

**Roles and responsibilities** Identify who is responsible for tasks such as scheduling test lab use and performing backups.

**Facilities and guidelines for special types of tests** Identify the domains and configurations team members should use for testing the migration process.

**Test lab change control procedures** Identify who is allowed to make configuration changes and define the approval process for change requests. For example, identify who can make schema changes and who should be notified when a change is made. Describe how changes to the lab should be made and identify who is responsible for documenting these changes.

**Server initialization procedures** Document the steps for installing, configuring, and populating domain controllers and member servers. Include DNS settings if you do not plan to use the DNS Server service in Windows Server 2003.

**Test lab restore procedures for testing the rollout** Document the steps for restoring domain controllers to their original state and for refreshing user account data. Document all server configurations. Test the refresh process before you begin migration testing.

**Client computer restore procedures** If you plan to rebuild client computers frequently to test various configurations, document the tools required to quickly restore the computers to a known state. For example, if you use RIS, note this in your guidelines.

As part of the lab guidelines and procedures, the lab manager or an appointed team member also should develop an escalation plan and an incident-tracking system.

# Developing an Escalation Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An *escalation plan* describes what testers need to do when problems arise during testing. The escalation plan should address the following issues:

- Where test failures and other test issues should be posted; for example, you might decide to post failures and problems in the incident-tracking system or on an intranet Web site.

  For information about the incident tracking system, see "Developing an Incident-Tracking System" later in this chapter.

- The steps to follow before posting an issue or problem; for example, you might want to indicate whether the problem should be reproduced and, if so, by whom.

- The information to include when posting an issue or problem:

  - Contact information (phone number, pager number, and e-mail address for the subteam lead and the external support person)

  - The status of the problem (new or ongoing)

  - The priority of the problem and the business objective that the problem jeopardizes

  - The sequence of events leading to the problem (include relevant information such as IP addresses and domain name)

  - The causes (known or suspected)

  - Troubleshooting information (traces, diagnostics)

- The procedure for notifying the development team of issues or problems.

- The notification hierarchy.

- The team members responsible for reviewing and resolving problems and issues.

# Developing an Incident-Tracking System

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An *incident-tracking system* is used for recording and reporting bugs and other testing problems, recording how they are resolved, and recording test results. A tracking system helps you monitor testing progress and evaluate the test success rate.

Your incident-tracking system should provide fields for recording:

- The tester's name and department.
- The date and time the test was performed.
- The Windows product name (Windows Server 2003 or Windows XP Professional).
- A full description of the test results.
- A description of any problem that was encountered, and its resolution.

Some organizations use a paper-based system to document test results, recording results on the page that describes the test case. With such a system, however, it is difficult to keep track of which tests have been performed and to create summary reports.

Alternatively, your organization can purchase a commercial product for tracking test cases. Or, your team can develop a database application specifically for this purpose. With either of these tools, you can create automated reports to monitor test results and progress.

Another alternative is to create an intranet site for posting test results. In this case, you can designate the level of detail displayed to a user depending on that user's credentials. For example, executives and test lab personnel could be permitted to view all test result information. Testers or technicians could be permitted to view only test results generated from within their group. By creating an intranet site, you can make test results available anywhere within your intranet, without requiring team members to install additional applications.

Regardless of the approach to documenting test results that you choose, it is important that test results be accessed easily.

# Designing Test Cases

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A *test case* is a detailed procedure that fully tests a feature or an aspect of a feature. Whereas the test plan describes what to test, a test case describes how to perform a particular test. You need to develop a test case for each test listed in the test plan. Figure 2.10 illustrates the point at which test case design occurs in the lab development and testing process.

**Figure 2.10 Designing Test Cases**




A test case includes:

- The purpose of the test.
- Special hardware requirements, such as a modem.
- Special software requirements, such as a tool.
- Specific setup or configuration requirements.
- A description of how to perform the test.
- The expected results or success criteria for the test.

Test cases should be written by a team member who understands the function or technology being tested, and each test case should be submitted for peer review.

Organizations take a variety of approaches to documenting test cases; these range from developing detailed, recipe-like steps to writing general descriptions. In detailed test cases, the steps describe exactly how to perform the test. In descriptive test cases, the tester decides at the time of the test how to perform the test and what data to use.

Most organizations prefer detailed test cases because determining pass or fail criteria is usually easier with this type of case. In addition, detailed test cases are reproducible and are easier to automate than descriptive test cases. This is particularly important if you plan to compare the results of tests over time, such as when you are optimizing configurations. Detailed test cases *are* more time-consuming to develop and maintain. On the other hand, test cases that are open to interpretation are not repeatable and can require debugging, consuming time that would be better spent on testing.

Table 2.1 provides an example of the first few steps of a detailed test case.

**Table 2.1 Sample Detailed Test Case**

| Step | Procedure | Success Criteria | Outcome |
|---|---|---|---|
| 1 | Log off the server, and return to the netlogon screen. | None. | |
| 2 | Click the domain list to open it. | The local server name does not appear in the list. | |
| 3 | Click the domain list to open it. | The root domain appears in the list. | |
| 4 | Log on to the server using an account with administrative credentials. | The account logs on to the server without errors. | |

When planning your tests, remember that it is not feasible to test everything. Instead of trying to test every combination, prioritize your testing so that you perform the most important tests — those that focus on areas that present the greatest risk or have the greatest probability of occurring — first. For example, you might choose to test the slowest client computer, the busiest server, or the least reliable network link. Then, if time allows, you can perform lower priority tests.

For more information about developing test cases, see "Planning and Testing for Application Deployment" in this book.

# Conducting Tests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After your team has designed test cases, it can begin testing. Figure 2.11 illustrates the point at which this step occurs in the overall lab development and testing cycle.

**Figure 2.11 Conducting Tests**




When conducting tests, the tester must perform each test as described in the test case, evaluate the test results, escalate problems that arise until they are resolved, and document the test results. Figure 2.12 illustrates the testing process.

**Figure 2.12 The Testing Process**




**Note**

- Before testing begins, the tester might need to modify the test lab setup to meet the requirements specified in the test case.

When conducting a test, the tester needs to follow the written test cases carefully. To accurately assess results or to reproduce a test to compare results over time, team members need to know exactly which steps the tester took and the exact sequence in which those steps were performed.

When a test is complete, the tester must evaluate the results against the criteria in the test case to determine if the test passed or failed. Not all tests fail because there is a problem with your system implementation; tests can fail because there is a

problem with the test itself, with the test lab setup, or with the proposed design. If one of these problems causes a test to fail, the tester should take the appropriate measures to remedy it. Table 2.2 shows some typical test problems and solutions.

**Table 2.2 Test Problems and Solutions**

| Problem | Potential Solution |
|---|---|
| Test case problem | Revise the test case, taking care to document all the changes made, and then rerun the test. |
| Test lab setup problem | Following the change control process for the lab, reconfigure the lab, and rerun the test. |
| Design problem | Follow the escalation procedure to notify the appropriate individuals about the problem. Prioritize outstanding problems and track them until they are resolved and the corresponding tests have been rerun. To prioritize problems, consider the potential impact and the probability that they will occur. |

Record all test results in the incident-tracking system, so you can monitor test progress.

**Important**

- Remember that the lab will change frequently as tests are run and new tests are begun. Make backups of baseline configurations so testers can quickly restore a computer to its prior state. Be sure to test the restore process. Document the backup files and store them in a safe, accessible place.

# Using the Test Lab After Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The importance of testing changes to your computing environment after you have deployed Windows Server 2003 and Windows XP Professional cannot be overemphasized. Because of the potential effect changes can have in the production environment, many organizations test every update and service pack until the anticipated results are achieved before they set up a pilot or a rollout. Figure 2.13 shows that the lab can be used to manage post-deployment changes.

**Figure 2.13 Using the Test Lab After Deployment**




If you plan to use your test lab to manage changes to your production environment, you will need to develop appropriate guidelines. Describe, in a general way, how to prepare the lab for testing a change and identify procedures for conducting tests and reporting results. For each change that you test, you will need to identify:

- Components required before the change can be implemented in the test lab.
- Steps required to implement the change.
- Documentation to be produced in the test lab.
- Action to take if a lab test fails.
- Action to take if a lab test succeeds.

Remember to periodically evaluate test lab components to determine and counteract the effects that accumulated changes have on them. For example, a computer that has undergone many changes in the testing process might not behave in the same way that a computer that has been newly installed with the same configuration would.

# Additional Resources for Designing a Test Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this chapter.

**Related Information**

- "Planning and Testing for Application Deployment" in this book for information about application compatibility testing.
- "Designing a Pilot Project" in this book for information about designing a pilot project.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about Dynamic Host Configuration Protocol (DHCP).
- "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of this kit for information about how to deploy DNS specifically for Active Directory.
- "Planning an Active Directory Deployment Project" in *Designing and Deploying Directory and Security Services* of this kit for information about how to plan and deploy Active Directory in your environment.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for more information about the Routing and Remote Access service.
- The *Active Directory Collection* of the *Windows Server 2003 Technical Reference* (or see the Active Directory Collection on the Web at http://www.microsoft.com/reskit) for more information about Active Directory functional levels.

# Planning and Testing for Application Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying applications to computers running the Microsoft® Windows® Server 2003 or Windows® XP Professional operating system requires that you first thoroughly analyze your applications and test them for compatibility with the operating system and your environment.

**In This Chapter**


Overview of Testing and Deploying Applications

Defining Application Deployment Project Scope and Objectives

Planning for Server Application Compatibility

Creating an Application Compatibility Test Plan

Preparing for Application Testing

Testing for and Resolving Compatibility Problems

Deploying and Distributing Applications and .Sdb Files

Additional Resources for Planning and Testing for Application Deployment


**Related Information**

- For more information about testing, see "Designing a Test Environment" in this book.
- For more information about piloting, see "Designing a Pilot Project" in this book.

# Overview of Testing and Deploying Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you move from your current version of the Microsoft® Windows® operating system to Windows XP Professional or Windows Server 2003, you need to test your applications to ensure that they are compatible with the new operating system. An organization can have as many as several thousand applications installed across distributed networks. Compatibility problems with one or many of these applications can cause costly work stoppages.

Although most applications developed for earlier versions of Windows probably will perform well on the new versions, some applications might behave differently because of new technologies within the new versions. The applications that you need to test to ensure compatibility include core applications that are part of your standard desktop configurations, such as office productivity suites; line-of-business applications, such as enterprise resource-planning suites; administrative tools, such as antivirus, compression, backup, and remote-control applications; and custom tools, such as logon scripts. Low-level applications — such as antivirus applications, kernel-mode drivers, and filter drivers — are especially likely to pose problems. You also need to ensure that your server applications are compatible.

This chapter is designed to help IT professionals in medium to large organizations test and deploy applications. It explains primarily how to test and deploy desktop applications that run on client computers. However, many of the concepts and techniques for application compatibility testing and deployment apply whether an application runs on Windows XP Professional or Windows Server 2003. For specific information about server application compatibility, see "Planning for Server Application Compatibility" later in this chapter.

This chapter does not describe how to test applications that will run on Microsoft® Windows® XP 64-Bit Edition Version 2003 or the 64-bit versions of Windows Server 2003 family. Testing 16-bit and 32-bit applications on these 64-bit operating systems is beyond the scope of typical application compatibility testing. This is because the 64-bit versions of these operating systems do not support 16-bit applications, and because you must use the *x*86 emulator, known as WOW64 (Windows32 on Windows64), to run 32-bit applications on the 64-bit versions of these operating systems. For more information about porting 32-bit applications to 64-bit operating systems, see the Porting 32-bit Applications link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Application Testing and Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Application compatibility testing begins with identifying and prioritizing the applications in use throughout your organization, which helps you determine the objectives and scope of the project. After establishing priorities and examining special considerations for server applications, you can develop the test plan. As you encounter compatibility problems during testing, you need to develop solutions, test them, and then package them for deployment. Figure 3.1 illustrates this process.

**Figure 3.1 Application Testing and Deployment Process**




**Important**

- To create distributable application compatibility solutions for Windows XP Professional and Windows Server 2003, you need to download and install the Microsoft Windows Application Compatibility Toolkit. To download the current version of the toolkit, use the Windows Application Compatibility link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. You can also find the toolkit on the Microsoft® Windows® Server 2003 Deployment Kit companion CD. The Windows Application Compatibility Toolkit cannot be installed on the 64-bit versions of the operating systems.

# Application Compatibility Testing Within the Project Life Cycle

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Application compatibility testing is an important part of your deployment project. Although application testing takes place during the developing phase of the project, related tasks occur throughout the project life cycle. Figure 3.2 shows the phase of the deployment project during which each of the application compatibility testing and deployment tasks might occur. The time frames are estimations and might vary from deployment to deployment.

**Figure 3.2 Application Compatibility Testing Within the Project Life Cycle**



# Application Compatibility Fundamentals

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To resolve application compatibility problems, perform the following tasks:

1. **Identify the applications that you need to test.** This includes creating an inventory of your applications (if you do not already have one) and determining the certification status, level of importance to your organization, and the priority of each application, to help determine the level of testing required.

2. **Identify application compatibility problems.** This includes testing, and possibly debugging, your applications.

3. **Resolve application compatibility problems.** This includes identifying and creating application compatibility solutions. It also can include modifying the source code and recompiling applications for which you have the source code.

4. **Deploy or distribute applications and solutions.** You can use various Windows XP Professional and Windows Server 2003 deployment and distribution tools.

To perform these tasks, you need to understand how application compatibility problems arise, what you can do to resolve the problems, and which tools you can use to plan your application compatibility testing project and to test, fix, and deploy applications.

# Common Compatibility Problems

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An application written specifically for a different version of Windows — especially for the Microsoft® Windows® 95, Windows® 98, or Windows® Millennium Edition (Windows Me) operating system — might manifest problems for several reasons. Most problems occur in the following areas.

**Setup and installation** Problems can occur when an application copies files and shortcuts to folders that exist on an earlier version of the Windows operating system, but not on Windows XP Professional or Windows Server 2003. Problems also can occur when a program:

- Writes entries directly to the registry without using the Windows Installer or the shell APIs.
- Checks for a specific version of the operating system.
- Does not support hard drives larger than 2 gigabytes (GB).
- Does not support long file names.
- Tries to access hardware directly rather than by calling the appropriate APIs.

**Kernel-mode drivers** Problems can occur when a driver tries to perform a task that is allowed in an earlier version of Windows but not in Windows XP Professional or Windows Server 2003. The most common types of programs that exhibit kernel-mode driver problems include antivirus programs, personal firewall programs, disk defragmenting programs, and CD burning programs.

**Permissions** Permissions problems can occur when an application tries to access areas of the file system or registry that are accessible by all users and applications in an earlier version of Windows but are no longer accessible in Windows XP Professional and Windows Server 2003. In Windows XP Professional and Windows Server 2003, user accounts can write only to the following locations:

- The registry key HKEY_CURRENT_USER, except its subkeys \Software\Policies and \Software\Microsoft\Windows\CurrentVersion\Policies
- The user's own user profile folders
- Shared Documents folders
- Any folder that the user creates from the root of the system hard disk

If a user tries to write to any other location, an error occurs.

**Heap management** Heap management problems can occur because Windows XP Professional and Windows Server 2003 immediately detect problems that frequently took some time to appear in earlier versions of Windows.

In addition, application compatibility problems can arise because of the following changes to existing features, and the addition of new features, in Windows XP Professional and Windows Server 2003:

- **Windows File Protection.** This feature prevents applications from replacing system files. Earlier versions of Windows allowed applications to replace shared system files during installation, which frequently caused problems ranging from application errors to an unstable operating system. Windows File Protection verifies that protected system files are the correct Microsoft version and restores the correct version when a file is replaced.
- **Enumeration of hardware devices.** Some hardware devices are no longer supported, which means they are not detected or enumerated during startup. This might cause problems for applications that expect a device to be present.
- **Enumeration of fonts.** The list of fonts has changed. Because registry keys have been added to support

internationalization, the operating system might not provide applications with the font choices that the application expects.

- **Windows Messaging Service.** Applications that expect the operating system to provide Windows Messaging Service (WMS) will not find it.

- **File input/output security.** Security has been enhanced for file input and output. Applications that use file filters, such as antivirus programs, might lose significant functionality.

- **Privileged instructions.** Use of privileged instructions is more restrictive. A *privileged instruction* is an operating system call that is executed in kernel mode rather than in user mode. An application might shut down unexpectedly if it uses privileged instructions such as **cli**, **sti**, **in**, or **out**. At other times, it might generate access violations that can be ignored, allowing the application to continue running without consequence.

# Application Compatibility Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows XP Professional and Windows Server 2003 resolve application compatibility problems by dynamically matching applications with solutions. The matching mechanism is controlled by a dynamic-link library (DLL) that starts and runs the matching mechanism whenever a user installs or runs an application. The matching mechanism is transparent and has a minimal effect on the performance of the operating system or installed applications.

To match applications with solutions, the matching mechanism relies on several databases. The databases contain a list of applications that have known problems and instructions for solving those problems. The databases are saved in the *systemroot*\AppPatch folder on Windows XP Professional and Windows Server 2003.

Three types of application compatibility databases are available: migration databases, prepackaged databases, and custom databases. The files containing the migration databases have an .inf extension; the files containing the prepackaged and custom databases have an .sdb extension. Throughout this chapter the term *.sdb file* is used to refer to prepackaged and custom application compatibility databases.

**Migration databases** These databases, which are provided by Microsoft, are used during an operating system upgrade to determine whether an application has compatibility problems, and to warn users about these compatibility problems if they prevent the application from running on the new operating system. Two migration databases are provided — one for checking applications that run on Windows 95, Windows 98, and Windows Me, and one for checking applications that run on the Microsoft® Windows NT® and Microsoft® Windows® 2000 operating systems. Migration databases have an .inf file name extension.

**Prepackaged databases** These databases also are provided by Microsoft, and are used when an application installs or runs, to determine whether the application has compatibility problems and to provide solutions to those problems. Prepackaged databases have the .sdb file name extension.

Prepackaged databases can contain any of the following solutions:

- **An application compatibility fix.** This is a segment of code that replaces a problematic function call with a function call that is compatible with Windows XP Professional or Windows Server 2003. This ensures that the operating system returns the result that the application expects.

- **An application compatibility mode.** This is a collection of application compatibility fixes that, together, emulate a specific operating system environment. For example, a Windows 95 application compatibility mode contains approximately 50 of the most common application compatibility fixes that applications designed to run on Windows 95 need in order to run on Windows XP Professional or Windows Server 2003.

- **An Application Help message.** This is a message that is displayed when a user installs or runs an application that is incompatible with the operating system, and an application compatibility fix or mode is not available or effective. An Application Help message can warn a user about an incompatibility but still let the user install or run the incompatible application, or it can block the user from installing or running the application.

**Custom databases** These databases are created by you, and can contain the same types of solutions as the prepackaged databases — application compatibility fixes, application compatibility modes, and Application Help messages. The difference is that these databases contain solutions for applications that are not addressed in the prepackaged databases. You can give a custom database any name, but it must have the .sdb file name extension. You also can save a custom database anywhere on a computer, but it is usually most convenient and efficient to save your custom databases in a subfolder in the *systemroot*\AppPatch folder (for example, *systemroot*\AppPatch\Custom).

In addition to the databases, the matching mechanism relies on several .dll files. These .dll files contain the code segments that run when an application requires an application compatibility fix or mode. These files begin with the letters "ac," and are installed by the operating system in the *systemroot*\AppPatch folder.

Table 3.1 describes all of the application compatibility databases that are prepackaged by Microsoft. The .inf files are stored on the Windows XP Professional and Windows Server 2003 operating system CDs. The .sdb files are stored in the *systemroot*\AppPatch folder.

**Table 3.1 Application Compatibility Databases**

| File | Description |
| --- | --- |
| MigDB.inf | Migration database that contains a list of Windows 95, Windows 98, and Windows Millennium Edition applications that are incompatible with Windows XP Professional or Windows Server 2003. |
| NTCompat.inf | Migration database that contains a list of applications that are designed to run on the Microsoft® Windows NT® Server 4.0 and Windows 2000 operating systems and are incompatible with Windows XP Professional or Windows Server 2003. |
| Apphelp.sdb | Prepackaged database that contains a list of third-party applications that have compatibility problems and associated Application Help messages. You can add third-party applications and custom Application Help messages to this database, but you cannot change or delete the existing list of names and Application Help messages. |
| Sysmain.sdb | Prepackaged database that contains a list of third-party applications that have compatibility problems and their associated application compatibility fixes and modes. You cannot change or delete the information in this database, but you can use the application compatibility fixes and modes that it contains to create custom databases. |
| Drvmain.sdb | Prepackaged database that contains a list of device drivers that have compatibility problems and their associated Application Help messages. |
| Msimain.sdb | Prepackaged database that contains a list of Windows Installer packages (.msi files) that have compatibility problems and their associated Application Help messages. |

# Application Compatibility Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft provides a variety of tools and features to help you to plan your application compatibility testing project and to test, fix, and deploy applications. These tools and features are provided for each of the four major phases in the overall application compatibility testing process: planning, testing, resolving, and deploying.

**Planning tools** These include Microsoft Application Compatibility Analyzer, Microsoft® Systems Management Server (SMS), Windows Catalog, and Windows Upgrade Advisor. Use Application Compatibility Analyzer, Windows Catalog, and Windows Upgrade Advisor to collect information about the applications in your organization and to identify applications that are known to have compatibility problems. Use SMS to create a software inventory.

**Testing tools** These include Microsoft Application Verifier and debugging tools. Use these tools to create a test environment for identifying hard-to-find application compatibility, stability, and security problems. Testing tools also include Windows Upgrade Advisor and the Windows Catalog, which identify applications that have already been tested and certified for use with Windows XP Professional and Windows Server 2003.

**Resolution tools** These include Microsoft Compatibility Administrator, the Program Compatibility Wizard, and the Compatibility property sheet. You can use Compatibility Administrator to apply compatibility modes and fixes to custom .sdb files, which you can distribute throughout your organization. The Program Compatibility Wizard and the Compatibility property sheet are ideal for resolving compatibility problems on a stand-alone computer (a computer that is not connected to a domain or a network).

**Deployment tools** These include automated installation tools, such as the System Preparation tool (Sysprep) and Remote Installation Services (RIS), and software management tools, such as Group Policy Software Installation and logon scripts. You can use these tools to deploy applications and .sdb files during an operating system rollout, or to distribute applications and .sdb files to computers that already have an operating system installed. Deployment tools also include the Application Compatibility Database Installer (Sdbinst.exe) and the Windows Installer program (Msiexec.exe), both of which are used in conjunction with the deployment and distribution tools to install applications and custom database (.sdb) files. You also can use the Windows Installer Software Development Kit to package applications, application updates, and .sdb files into Windows Installer packages (.msi files).

**Note**

- The Program Compatibility Wizard and the Compatibility property sheet are rarely used to address application compatibility issues in a large enterprise. In addition, you cannot use these two tools to resolve issues with kernel-mode drivers; doing so can result in system instability.

Table 3.2 provides a list of common tasks, the tools that you use to perform them, and where those tools can be found.

**Table 3.2 Tools for Performing Common Application Compatibility Tasks**

| To perform this task... | Use this tool or feature... | Available... |
|---|---|---|
| Create a software inventory that is adequate for general deployment planning tasks, including application compatibility testing. | SMS | For purchase from Microsoft. |
| Create a software inventory that is tailored for application compatibility testing. | Microsoft Application Compatibility Analyzer | In the Windows Application Compatibility Toolkit. |

| | | |
|---|---|---|
| Quickly check the compatibility of applications. | Windows Catalog<br>–or–<br>Windows Upgrade Advisor* | On the Microsoft Web site. Windows Upgrade Advisor is also included on the Windows XP Professional and Windows Server 2003 operating system CDs. |
| Test applications for hard-to-find stability, compatibility, and security problems. | Application Verifier | In the Windows Application Compatibility Toolkit. |
| Investigate and resolve complex application compatibility problems when you have access to the application's source code. | Microsoft Debugging Tools for Windows | On the Microsoft Web site. |
| Quickly resolve common application compatibility problems on a stand-alone computer. | Program Compatibility Wizard<br>–or–<br>Compatibility property sheet | In Windows XP Professional and Windows Server 2003. |
| Resolve application compatibility problems by creating a custom application compatibility database (.sdb file). | Compatibility Administrator | In the Windows Application Compatibility Toolkit. |
| Install and register custom application compatibility files and .sdb files on destination computers. | Application Compatibility Database Installer (Sdbinst.exe) | In Windows XP Professional and Windows Server 2003. |
| Create a Windows Installer package (.msi file). | Windows Installer Software Development Kit | On the Microsoft Web site. |
| Install a Windows Installer package (.msi file) on a destination computer. | Windows Installer tool (Msiexec.exe) | In Windows XP Professional and Windows Server 2003. |
| Deploy applications and .sdb files during an operating system rollout. | Windows System Preparation tool (Sysprep)<br>-or-<br>Unattended installation answer file (Unattend.txt)<br>-or-<br>Remote Installation Services (RIS) | On the Windows XP Professional and Windows Server 2003 operating system CDs. |

| | | |
|---|---|---|
| Distribute applications and .sdb files after an operating system rollout. | Group Policy Software Installation tool<br>-or-<br>Logon scripts<br>-or-<br>Software Update Services (SUS)<br>-or-<br>SMS | All tools except for SMS are included in Windows XP Professional and Windows Server 2003. SMS is available for purchase from Microsoft. |

*Windows Upgrade Advisor is used rarely in enterprise environments, although it is a viable method of checking the compatibility of an application.

For more information about planning tools, see "Generating an Application Inventory" later in this chapter. For more information about testing tools, see "Testing with Windows Application Verifier" later in this chapter. For more information about resolution tools, see "Applying and Testing Solutions with Compatibility Administrator" later in this chapter. For more information about deployment tools, see "Deploying and Distributing Applications and .Sdb Files" later in this chapter.

# Defining Application Deployment Project Scope and Objectives

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The fundamental objective of any application compatibility testing and deployment project is to successfully deploy applications that are compatible with the new operating system. To plan your application deployment project, you must first define the project scope and objectives. To do this, you need to conduct an inventory of the software currently in use in your organization, and then analyze this information to determine the constraints and risks of the compatibility testing project. You can use this information to begin creating the deployment project plan.

Figure 3.3 illustrates when this step occurs in the application compatibility testing process.

**Figure 3.3 Defining Application Deployment Project Scope and Objectives**



# Generating an Application Inventory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To determine the scope of the project, you need to conduct an inventory of the applications used in your organization. If you do not already have a complete software inventory, gather information about both client and server applications that are installed on your computers. Include limited-use applications for special projects within the organization and nonapproved software that will inevitably be present. By having a complete inventory of the applications used in your organization, you can ensure that an otherwise overlooked application is not incompatible.

Several tools are available for creating a software inventory. The Application Compatibility Analyzer, included in the Windows Application Compatibility Toolkit, gathers data that can make your application compatibility testing project easier and more efficient to implement. You also can use SMS or third-party tools to inventory your software. SMS is a robust systems management tool, which, although not designed specifically for application compatibility projects, provides adequate inventory information for application compatibility testing.

Regardless of which tool you use to generate your software inventory, your inventory should include the following information:

- The application name, vendor, and version number
- Service packs or hot fixes that have been applied or are required
- An evaluation of how critical the application is to the success of your business
- The target platform for the application and a description of the client/server components
- The development platform used to create the application
- Historical data, including any known problems or bugs that were identified in previous testing

Some of this information is gathered by the tool that you use to generate the software inventory. You will need to manually collect any additional information that is required. For information about the type of data that a tool collects, see the documentation for the tool.

**Note**

- In addition to generating an inventory of the applications currently in use in your organization, be sure to account for applications that you plan to introduce with the new operating system. Because these applications cannot be inventoried, you need to collect information about them manually.

# Using the Application Compatibility Analyzer to Generate an Application Inventory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Application Compatibility Analyzer helps you collect software-related information from the computers in your organization. The Application Compatibility Analyzer tool has two components:

- A console component, the Application Compatibility Analyzer, which you install on the administrator's computer.
- A command-line component, the Microsoft Collector, which you distribute to each computer from which you want to gather an application inventory. The Collector runs in the background on each computer to which it is distributed to collect data about every application on that computer. The Collector records the data in an XML-formatted log file for that computer.

You can compile the data gathered by the Collector into a Microsoft® SQL Server™ database or a Microsoft® Access database by using the Application Compatibility Analyzer console. Then, you can use the Application Compatibility Analyzer console to generate reports that describe the compatibility of the applications in your inventory. To do this, you use the Application Compatibility Analyzer console to compare your software inventory to a database of application compatibility information (known as the *application compatibility database*). The application compatibility database is maintained by Microsoft, and contains application compatibility information from the Windows Catalog, Microsoft internal testing, and outside parties. It is periodically updated by Microsoft. You can compare your inventory with information in the database directly over the Web, or you can download the compatibility information and add it to your database so you can use the compatibility information offline.

**Note**

- If you download the application compatibility database, you can modify the database to include the results of in-house application testing.

The Application Compatibility Analyzer is available for downloading from the Microsoft Web site. See the Windows Application and Customer Experience link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Installing and Using the Application Compatibility Analyzer**

You can install the Application Compatibility Analyzer on any computer that is running Windows 2000, Windows XP Professional, or Windows Server 2003. It requires Microsoft® Internet Explorer version 5.5 or later.

Install the Application Compatibility Analyzer on the administrator's computer on which you will review Application Compatibility Analyzer reports. Then, use Application Compatibility Analyzer to define the database that will store the information that Collector collects as either a SQL Server database or an Access database. If you are compiling Collector logs in an enterprise, a SQL Server database is recommended.

After installing the Application Compatibility Analyzer, distribute the Collector to your client and server computers. You can run the Collector on the Windows 95, Windows 98, Windows Me, Windows 2000® Professional, or Windows XP Professional operating system, and on servers running Windows NT Server 4.0, Windows 2000, or Windows Server 2003. The Collector need not run under an administrator account.

You can configure the Collector using command-line arguments or initialization (.ini) scripts. You can specify which drives, either network or local, and which paths the Collector should search, and whether the Collector should detect server functions. You also can configure the Collector to define the scope of the inventory, and you can specify where to save the log files (by default the log files are stored on the computer on which the Collector is run). For more information about customizing the Collector, see the Collector Help files.

**Important**

- Set up a secure share with write-only access on which to save your log files. The log files contain computer information that should be accessible only by administrators.

You can distribute the Collector manually on floppy disks or on CDs, over the network using a network distribution share or a hyperlink in e-mail, or through Group Policy. These distribution methods all require that end users actively install and run the Collector on their computers. Alternatively, you can automate Collector distribution by using logon scripts or SMS. For information about distributing the Collector, including sample logon scripts, see "Collection Scenarios" in Application Compatibility Analyzer Help.

**Note**

- If you plan to use logon scripts to run the Collector, use the **/CW** switch to configure the Collector to wait for five minutes before starting collection. This reduces the CPU load at startup.

When the Collector finishes collecting information from all of your computers, you can use the Application Compatibility Analyzer console to merge all of the Collector log files into the SQL or Access database that you defined earlier. For more information about merging Collector log files into a database, see "Merger" in Application Compatibility Analyzer Help.

**Example: Using the Application Compatibility Analyzer to Generate an Application Inventory**
The following example illustrates how the Application Compatibility Analyzer might be used in a network environment.

A midsized company is preparing to deploy Windows XP Professional to all desktop computers. The organization is divided among three physical locations connected by high-speed data connections. The organization uses the Active Directory® directory service within a single domain, and has created an organizational unit (OU) for each physical location (HQ, East, and West).

On Server1, the IT department creates a shared folder for storing all log files generated by the Collector, with the Universal Naming Convention (UNC) path \\Server1\Analyzer, and another shared folder for storing the Collector executable, with the path \\Server1\Collector. The IT department decides to use a logon script assigned through Active Directory to distribute and run the Collector. Because there are three physical locations represented by three separate OUs within the domain, the organization decides to gather data by OU. To accomplish this, the IT department adds the following lines to the logon script:

```
Copy \\server1\collector\collector.exe c:\
C:\collector.exe /O \\server1\analyzer /N /E HQ /CW
```

The first line of code copies the Collector executable program from the shared folder on Server1 to drive C on the client computer. The second line of the code instructs Collector to do the following: send log file output to the shared folder on Server1 that is specified for storing logs (with the **/O** switch); scan network drives that might be mapped (with the **/N** switch); mark the data with an "HQ" designation for the HQ OU (with the **/E** switch); and wait 5 minutes before starting the collection. The IT department can include similar lines for the logon scripts in the East and West OUs, with the exception that the **/E** switch would define the OU name in each case.

# Organizing and Analyzing Inventory Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After gathering your inventory data into a single database, you can analyze the applications currently installed in your organization. You can begin to determine which applications are compatible with the new operating system and which are not compatible or require special considerations, and you can make decisions about consolidating the application base in your organization.

**Checking Application Compatibility**

Using the Application Compatibility Analyzer, you can connect to the application compatibility database at Microsoft to determine whether the applications in your organization are compatible with Windows XP Professional or Windows Server 2003. When the Application Compatibility Analyzer tool generates a compatibility report, it automatically accesses the application compatibility database and updates your software inventory with compatibility information. If you use a different tool to generate an inventory, or if you want to check the compatibility of an application that you have not yet deployed, you can use the Windows Catalog instead.

You can use an application's compatibility rating to assess the level of testing that application requires. If your organization uses an application listed as compatible or as designed for Windows XP Professional or Windows Server 2003 without customization, you need only minimally test the product. However, if you customize applications with macros, templates, or custom tools, you need to thoroughly test those applications.

**Checking application compatibility using the Application Compatibility Analyzer**

When you use the Application Compatibility Analyzer tool to request compatibility information for an application from the application compatibility database, one of four levels of compatibility is returned:

- **Compatible.** The application has no known issues.
- **Compatible with Issues.** The application has one or more minor functionality issues.
- **Incompatible.** The application's general functionality fails.
- **Unknown.** The application's compatibility is unknown.

The compatibility levels returned by the application compatibility database are combined with the information in your inventory to create a local application compatibility database for your organization. The local compatibility database provides details about the operating systems and scenarios under which the applications were tested. It also contains other useful information, such as the number of computers on which each application is installed and which departments are using which applications. You can use this information to categorize your applications and prioritize them for testing.

For more information about checking application compatibility with the Application Compatibility Analyzer tool, see "Connecting to the Compatibility Database" in Application Compatibility Analyzer Help.

As you find out more about an application, either through your own testing or from the application vendor, you can update its compatibility information in your local compatibility database.

**Checking application compatibility using the Windows Catalog**

If you gathered your inventory using SMS or a third-party inventory tool, or if you have not yet deployed the application whose compatibility you want to check, you can use the Windows Catalog to check compatibility. In Windows XP Professional or Windows Server 2003, you can open the Windows Catalog by clicking **Start,** clicking **All Programs**, and then clicking **Windows Catalog**. Alternatively, you can use the Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

If an application is listed as Compatible in the Windows Catalog, either Microsoft or its manufacturer has tested it for compatibility with Windows XP Professional, and it meets specific compatibility criteria. The Windows Catalog also highlights products that are part of the Designed for Windows Logo Program for Windows XP Professional. Products with the Designed for Windows XP logo are specifically created to take advantage of the features of Windows XP Professional, which reduces the number of problems that you might otherwise have in using these applications.

**Gathering Additional Application Information**

Regardless of which tool you use to generate your inventory, you need to manually collect information that is organization-specific or group-specific for each application. This type of information cannot be collected by an inventory tool, but it is important for creating your test plan.

Talk with representatives from the groups that use each application to document:

- A description of the application's function.
- The name of the application vendor (if applicable).
- The version number of the application.
- The importance of the application to your organization.
- The current status of the application (in production, under development, no longer used).
- The development platform for the application (if the application was developed internally).
- Contact names and phone numbers for internal contacts and vendors.
- Other information that pertains to the organization or group's unique use of the application.
- Additional certification or testing that might be necessary to meet government regulations. For example, certain applications used in the development and production of pharmaceuticals must meet certification requirements set by the Food and Drug Administration (FDA). In some cases, if you change the operating system on which the application runs, you might also have to recertify the application.

You can use this information to help prioritize your testing, especially if you have limited resources and are unable to fully test each application.

**Consolidating Your Application Base**

Review the inventory to see if you can consolidate your organization's application base. Limiting the applications used in your organization can minimize your testing effort, decrease configuration variability during deployment, and increase the likelihood of a successful deployment.

Look for the following in your inventory:

- **Several versions of the same application.** Consider updating older applications to newer versions or moving all users to a full-featured version of a particular application. For example, if some of your users use Microsoft® Office 2000 and others use Microsoft® Office XP, you might decide to support only Office XP in the future.
- **Redundant applications.** If groups in your organization are using different applications to accomplish the same tasks, consider moving everyone to the same application.
- **Obsolete applications.** Review your inventory for applications that are rarely or never used in your organization, and consider retiring them.

# Defining Project Constraints and Risks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Identify the project constraints and risks in order to help your team define the parameters of the application deployment project and develop the project plan.

**Defining Project Constraints**

To prioritize and define the scope of the application deployment project, gather information about the constraints of your project. Constraints often include:

- **Resources.** Identify the equipment, software, staff, and space that are available for the project.
- **Time.** Identify the date by which the application deployment project must be completed, and how the application testing process fits into the larger deployment project.
- **Organizational issues.** If the project will not involve the entire organization, identify which groups in your organization will be affected by it. Additionally, determine if a particular group in the organization needs the new operating system sooner than others. If so, you might decide to perform a staged rollout.
- **Access to developers.** Identify applications that were developed in-house or especially for your organization. Access to the developers of these applications is critical during the testing and issue resolution phases of the project. Such access also can be an invaluable aid with retail applications.

**Identifying Risk Factors**

By identifying risk factors, you can identify the potential barriers to a successful deployment and assess the cost of failure for each potential blocking issue. With this information in hand, you can create contingency plans to help mitigate or avoid blocking issues. For application compatibility testing, it is important to think about the project risks early in the project rather than later, during the test planning phase.

The main risk for any application testing and deployment project is the business impact of application failure. For business-critical applications, a failure in deployment could have severe financial repercussions for the organization, whereas the failure to properly deploy an application that is used infrequently by a few users might have very little effect. Assess the financial impact that could be caused by a failure at any time during the deployment project.

For information about assessing constraints and risks, see the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Planning for Server Application Compatibility

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assessing the compatibility of server applications with Windows Server 2003 presents some unique challenges. Because server applications tend to span several servers and to do so in several tiers, generating an application inventory, as well as deploying application compatibility fixes for server applications, can be challenging.

Figure 3.4 illustrates when the compatibility of server applications is assessed in the application compatibility testing process.

**Figure 3.4 Planning for Server Application Compatibility**



# Inventory Considerations for Server Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Server applications are often tiered applications with different components running on different servers. This can make generating an inventory of these applications more difficult than doing so for client applications, because the various components of the application might not be associated in the inventory. Work with the owners of the servers in your network to manually associate the various components of the applications within your database.

The Application Compatibility Analyzer does not detect Internet Authentication Service (IAS) applications. If your organization has IAS applications, you must gather inventory data about them manually.

# Obtaining ISV Support for Server Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In large organizations, it is often important to have your Independent Software Vendor (ISV) support all server applications that will be running on the new operating system. Commonly, large organizations check with the ISV to see if there is an update or upgrade that solves application compatibility problems, rather than create an application compatibility fix themselves.

Alternatively, you can obtain a statement of support from an ISV. A statement of support means that the ISV has designed and tested the application to work well with the new operating system. Microsoft continues to work with ISVs to encourage them to provide statements of support for widely used server applications. You can find ISV statements of support for application compatibility in the Windows Server Catalog, or through the online compatibility database that you can access through the Application Compatibility Analyzer. If no statement of support is posted for your application, contact the ISV and request that they publish one through the Windows Server Catalog.

The Windows Server Catalog also contains the names of applications that are certified for Windows Server 2003. Applications that are certified for Windows Server 2003 must pass a rigorous third-party testing process, which provides the highest assurance that they will run on Windows Server 2003.

For more information about the Windows Server Catalog, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

In some cases, you might find that an ISV does not support the version of the software that you are using in the production environment, but does support a later version. If this is so, consider upgrading to the supported version.

**Note**

- An ISV can submit information about an application through the Designed for Windows Logo Program. For more information, see the Designed for Windows Logo Program link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Special Considerations for Server Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following sections provide information about server application compatibility.

**Active Directory Considerations**
Several changes in the implementation of Active Directory between Windows 2000 and Windows Server 2003 might affect your applications.

**Anonymous LDAP operations**
In earlier versions of Windows, Active Directory accepted anonymous client requests. With Windows Server 2003, only authenticated users can initiate a Lightweight Directory Access Protocol (LDAP) request against Windows Server 2003–based domain controllers. For more information about anonymous LDAP requests, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, and search for article number 326690 "Anonymous LDAP Operations to Active Directory Are Disabled on Windows Server 2003 Domain Controllers."

**Microsoft Exchange Server 2000 schema conflicts**
The Microsoft® Exchange Server 2000 schema defines several attributes differently from those defined in RFC 2798, "Definition of the inetOrgPerson LDAP Object Class." The attributes — **Secretary**, **labeledURI**, and **houseIdentifier** — also are defined by the base schema. This deviation can cause problems when you upgrade from Windows 2000 to Windows Server 2003. During the upgrade to Windows Server 2003, the **ldapDisplayNames** values for these attributes are renamed to permit the upgrade to succeed, causing duplicate names. This renaming does not cause problems with the functioning of Active Directory or the operating system, because the names are not used in this context. However, you might have an application that needs to access these values, which will be inaccessible.

This situation is not a problem if:

- You add the Windows 2000 InetOrgPerson Kit to a Windows 2000 forest before you run the Windows Server 2003 **adprep** command.
- You run the Windows Server 2003 **adprep** command in a Windows 2000 forest before you install Exchange Server 2000.
- You add Exchange Server 2000 to an existing Windows Server 2003 forest.

However, you may encounter problems if:

- You add the Exchange Server 2000 schema to a Windows 2000 forest before you add the InetOrgPerson class from the InetOrgPerson Kit.
- You add the Exchange Server 2000 version of the InetOrgPerson class to a Windows 2000 forest before you run the Windows Server 2003 **adprep** /**forestprep** command.
- A Windows 2000 domain controller that contains the Exchange Server 2000 definition of InetOrgPerson does not receive updates to Active Directory after you run InetOrgPerson-Fix.ldf from the Windows 2000 InetOrgPerson Kit.

For more information, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources and search for article number 314649 "Windows Server 2003 ADPREP Command Causes Mangled Attributes in Windows 2000 Forests That Contain Exchange 2000 Servers," and 328661 "XADM: Running Exchange 2000 Setup with /Forestprep Switch Produces Error 0XC1037AE6."

**Message Queuing**
Several features of previous versions of Microsoft® Message Queuing (MSMQ) have been removed from MSMQ version 3.0, which ships with Windows Server 2003. In addition, one feature has been changed. Consequently, you might need to modify applications that rely on features that are no longer relevant or useful, or that have been superseded by other features.

Feature changes include:

- The MSMQ Exchange connector is no longer supported in Windows XP Professional and Windows Server 2003. It continues to be available for Windows NT Server 4.0 and Windows 2000.

- The Internetwork Packet Exchange (IPX) protocol is not supported in MSMQ 3.0.

- The MSMQ 3.0-dependent client supports only the MSMQ version 2.0 level of functionality. This means that new MSMQ 3.0 features, such as distribution lists, are not available to MSMQ 3.0-dependent clients. The alternative is to deploy DCOM-based solutions instead.

If your applications use any of these features, verify that they work with MSMQ 3.0.

For more information about applications and MSMQ, see the "Programming Best Practices with Microsoft Message Queuing Services (MSMQ)" white paper, available through the Message Queuing link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Internet Information Services**
Windows Server 2003 includes Internet Information Services (IIS) version 6.0. Changes made in IIS 6.0 might affect your Web applications unless you address the following configuration issues during the installation of Windows Server 2003:

- IIS is not installed by default during a clean installation, except on the Microsoft® Windows® Server 2003, Web Edition, operating system. If you plan to run Web-based applications on other editions of the Windows Server 2003 operating system (Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003 Enterprise Edition; or Windows® Server 2003, Datacenter Edition), you must install IIS.

- Active Server Pages (ASP), ASP .NET, Server Extensions (installed and enabled) for the Microsoft® FrontPage® Server Extensions (installed and enabled), Web-based Distributed Authoring and Versioning (WebDAV), and all other Internet Server API (ISAPI) extensions and Common Gateway Interface (CGI) applications by default are not enabled when you enable the Application Server Role in **Add/Remove Windows Components**. You can install these by using **Add/Remove Windows Components**, and then enable them in the Web Service Extensions node in IIS Manager.

- If you perform a clean installation of IIS 6.0, IIS 6.0 runs in worker process isolation mode by default. To take advantage of the new security, availability, and scalability features in IIS 6.0, IIS 6.0 must be running in worker process isolation mode. However, for compatibility purposes, you can change the application isolation mode to IIS 5.0 isolation mode if the application must run in the Inetinfo.exe process or Dllhost.exe process, or when the application uses ISAPI filters that perform raw data reads.

For more information about IIS and compatibility problems for Web applications, see "Migrating IIS Web Sites to IIS 6.0" in *Deploying Internet Information Services (IIS) 6.0* of this kit (or see "Migrating IIS Web Sites to IIS 6.0" on the Web at http://www.microsoft.com/reskit).

**Terminal Server**
Because Terminal Server is available for multisession or multiuser use, and because display and keystroke information travels over the network, applications that have certain characteristics might perform poorly with Terminal Server. When planning to host applications with Terminal Server in your organization, establish a set of acceptable performance guidelines and determine through testing whether such applications run better on the user's local computer or over your intranet. For more information about application compatibility with Windows Server 2003 Terminal Server, see "Hosting Applications with Terminal Server" in *Planning Server Deployments* of this kit.

For more information about server application compatibility, see the Server Roles: The Application Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Enhanced Security Configuration in Internet Explorer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To prevent security risks, the default configuration of Internet Explorer has changed in Windows Server 2003. This new default configuration, known as Internet Explorer Enhanced Security Configuration, restricts the ability of a Web site to download files and to run scripts and ActiveX® components. In addition, the Internet Explorer Enhanced Security Configuration disables Internet Explorer's ability to detect whether a Web site is in the Local intranet zone, which restricts access to such things as: intranet Web sites, Web-based applications that run over the intranet, and files or folders on network shares. The Local intranet zone is one of the four Web content zones, along with Internet, Trusted sites, and Restricted sites.

Prior to the implementation of the Internet Explorer Enhanced Security Configuration, anything that was in the Local intranet zone was considered trustworthy. For example, if Internet Explorer determined that a Web site was in the Local intranet zone, and the Web site requested a user's credentials, Internet Explorer would automatically pass the credentials to the Web site without prompting the user. Under the Internet Explorer Enhanced Security Configuration, intranet Web sites and universal naming convention (UNC) paths must be explicitly specified in the Local intranet zone or the Trusted sites zone to be considered trustworthy.

Because of the Enhanced Security Configuration settings, some applications might experience problems interacting with Web sites, network resources, and ActiveX components and scripts. These problems can occur when an application is being installed and when an application is running. For example, the Enhanced Security Configuration can prevent an application from being installed properly when the application attempts to download required files from a site that is restricted. Similarly, the Enhanced Security Configuration can prevent an application from running properly when it attempts to use an Internet Explorer feature that has been disabled. In addition, an application might not be able to use a UNC path to access a shared folder across the network.

To overcome these limitations, you can:

- Add trusted sites and UNC paths to the Trusted sites zone or to the Local intranet zone.
- Modify Internet Explorer security settings for a specific content zone or all content zones.
- Disable Internet Explorer Enhanced Security Configuration settings.

If you know that an application requires access to an Internet site that is restricted, you also can request that the application vendor change the application setup program so that the required Web site is added to the Trusted sites zone during installation.

**Caution**

- Do not disable the Enhanced Security Configuration settings unless you have no other alternative. Disabling the Enhanced Security Configuration increases the risk of security attacks.

For more information about Enhanced Security Configuration settings, see "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003.

# Testing Applications for Compatibility with the Enhanced Security Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a user or an administrator opens Internet Explorer for the first time on Windows Server 2003, a default home page that contains information about the Internet Explorer Enhanced Security Configuration is displayed in the browser window. If the user is not a member of the Administrators group, the home page describes the Enhanced Security Configuration and explains how the user can access a Web site that is not trusted. If the user is a member of the Administrators group, the home page describes the Enhanced Security Configuration in detail and explains how the administrator can modify it. Because OEMs can set a different default home page as part of their system configuration, a warning explaining that the Enhanced Security Configuration is enabled is also displayed in a separate dialog box. Users can permanently dismiss this dialog box.

When users start Internet Explorer directly — for example, by clicking a shortcut to Internet Explorer or by clicking Internet Explorer on the **Start** menu — these messages provide clear guidance for what to do if a Web site does not display properly or if a network resource cannot be reached. However, if an application starts Internet Explorer, the user might not receive these warnings and, consequently, will not understand why the application is failing or how to resolve the problem.

There are four primary types of problems to look for when you test server applications under the Internet Explorer Enhanced Security Configuration settings.

**Sites and UNC paths are not listed in the trusted Web content zone** Look for cases where an application starts Internet Explorer so that the user can view a Web site, download files, or run a Web-based application. If the application cannot access these Web sites, applications, or resources, see if you can resolve the problem by adding the Web site or UNC path to the Trusted sites zone or to the Local intranet zone (if appropriate).

**Use of client-side redirection** Look for cases where an application's online content relies on client-side redirection (for example, an application tries to open a Web page, but then gets routed automatically to a different URL). Client-side redirection is disabled under the Enhanced Security Configuration. For client-side redirection to work properly, the target Web site needs to be added to either the Trusted sites zone or to the Local intranet zone (if appropriate).

**Dependency on ActiveX components** Look for cases where an application accesses a Web site that relies on Active X components to display properly, such as a Web site that requires Macromedia Flash or Adobe Acrobat. You need to download and install these types of Web-based applications and ActiveX components on a per-computer basis, as you would do with any other application software.

**Use of Internet Explorer features** Although they are not obvious, look for cases where an application relies on Internet Explorer features and components — for example, Install on Demand or multimedia features.

While most applications that rely on Internet Explorer features will work properly, in some cases an application, or some extensible component within an application, might attempt to access an Internet Web site that relies on a script or an ActiveX component. In these cases, if the Web site is not listed in the Trusted sites zone, the script or ActiveX component will not run properly, and the user will be prompted to add the Web site to the Trusted sites zone. You can eliminate the need for user intervention by identifying the scripts and ActiveX components that an application relies on and by deploying those scripts and ActiveX components to your user base, just as you would any application.

# Modifying Enhanced Security Configuration Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can modify Enhanced Security Configuration settings manually by using **Add or Remove Programs** and **Internet Options** in Control Panel. Usually, you modify settings manually on a single machine when you are testing applications and trying to identify and resolve problems. You also can automate the modification of Enhanced Security Configuration settings during an unattended installation or during an image-based installation with Sysprep by changing answer file settings. Usually, you automate configuration tasks when you are deploying an operating system and applications to many computers.

**Manually Modifying Enhanced Security Configuration Settings**

When Enhanced Security Configuration settings are enabled, it reduces the exposure of your server to security attacks from Web sites that are not listed in the Local intranet zone or in the Trusted sites zone. By adding intranet Web sites and UNC paths that you trust to the Local intranet zone, or by adding Internet Web sites to the Trusted sites zone, you can maintain the security of your server computers while allowing access to trusted Web sites.

**To add Web sites and UNC paths to the Local intranet or Trusted sites zones**

1. In Control Panel, double-click the icon for **Internet Options**.
2. On the **Security** tab, click **Local intranet** or **Trusted sites**, depending on which Web content zone you want to add a Web site or a UNC path to.
3. Click **Sites**.
4. In the **Add this Web site to the zone** text box, type the URL of the Web site or UNC path that you want to add, and then click **Add**.

**Important**

- Do not add Internet Web sites to the Local intranet zone. User credentials can be passed to Web sites in the Local intranet zone without notifying the user.

When the Enhanced Security Configuration is enabled, several of the extensibility and security features of Internet Explorer are adjusted to decrease security risks. As a consequence, some scripts, ActiveX components, and other applications might not run. By changing these settings, you can maintain the security of your server computers and still allow certain scripts, ActiveX components, and applications to run.

**Important**

- Changing extensibility and security features requires in-depth knowledge and understanding of the consequences and is not recommended.

**To modify Internet Explorer extensibility and security features**

1. In Control Panel, double-click the icon for **Internet Options**.
2. On the **Advanced** tab, configure extensibility and security features by selecting or clearing the check box next to the feature.

The Enhanced Security Configuration can be enabled or disabled based on group membership. You can disable the Enhanced Security Configuration for members of the Administrators group (it is enabled by default), and you can enable Enhanced Security Configuration for members of the Users group (it is disabled by default). Usually, you disable the Enhanced Security Configuration for members of the Administrators group only when you are testing applications or when you are configuring a computer prior to putting the computer into a production environment. You should always enable the Enhanced Security Configuration for members of the Administrators group before you put a computer into a production environment.

**To enable or disable Enhanced Security Configuration**

1. In Control Panel, double-click **Add or Remove Programs**.

2. Click **Add/Remove Windows Components**.

3. Click **Internet Explorer Enhanced Security Configuration**, and then click **Details**.

4. Select the group for which you want to enable or disable Enhanced Security Configuration, and then click **OK**.

**Automating the Modification of Enhanced Security Configuration Settings**

There are two ways to modify Enhanced Security Configuration settings during an unattended installation or during an image-based installation with Sysprep: you can add trusted Web sites to the Local intranet zone or to the Trusted sites zone, or you can disable or enable the Enhanced Security Configuration.

Usually, during an automated installation, you do not need to disable the Enhanced Security Configuration. However, there is one case in which you might need to do so. If you are performing image-based installations with Sysprep, and you set up your master installation by performing an unattended installation, you might need to disable the Enhanced Security Configuration so that you can download and install ActiveX components, device drivers, and applications on the master installation. You can then enable the Enhanced Security Configuration during the Factory mode phase of the image-based installation. For more information about automated installations, see "Deploying and Distributing Applications and .Sdb Files" later in this chapter.

**Adding Web sites to Web content zones during automated installations**

You can add Web sites to the Trusted sites zone and the Local intranet zone by configuring the [IEHardening] section of the unattended installation answer file (Unattend.txt) and the Factory mode answer file (Winbom.ini). The unattended installation answer file is used to automate an unattended installation; the Factory mode answer file is used to automate the Factory mode phase of an image-based installation that uses Sysprep. The syntax is the same for both answer files.

You can add Web sites to the Trusted sites zone by using the following entry in the [IEHardening] section of Unattend.txt or Winbom.ini:

```
TrustedSites = url_1 [,url_2]…
```

You can add Web sites to the Local Intranet zone by using the following entry in the [IEHardening] section of Unattend.txt or Winbom.ini:

```
LocalIntranetSites = url_1 [,url_2]…
```

**Enabling or disabling Enhanced Security Configuration during automated installations**

You can enable or disable the Enhanced Security Configuration by configuring the [Components] section in the unattended installation answer file (Unattend.txt) and the Factory mode answer file (Winbom.ini). The syntax is the same for both answer files.

You can enable or disable the Enhanced Security Configuration for members of the Administrators group by using the following entry in the [Components] section of Unattend.txt or Winbom.ini:

```
IEHardenAdmin = On | Off
```

You can enable or disable Enhanced Security Configuration for members of the Users group by using the following entry in the [Components] section of Unattend.txt or Winbom.ini:

```
IEHardenUser = On | Off
```

# Creating an Application Compatibility Test Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have gathered the preliminary information that you need, you can begin creating your application compatibility test plan. Your application compatibility test plan should address all of the issues that are typically addressed in a standard test plan — testing scope and objectives, testing hardware and software requirements, incident tracking procedures, testing scheduling, and so on. Most of these issues are discussed earlier in this chapter and in "Designing a Test Environment" in this book.

You can use this information to create the foundation of your application compatibility test plan. In addition, your application compatibility test plan should:

- Define your testing methodology.
- Identify which applications you need to test.

Figure 3.5 shows when these activities occur in the application compatibility testing process.

**Figure 3.5 Creating an Application Compatibility Test Plan**




For complete information about creating your overall test plan, see "Designing a Test Environment" in this book.

# Defining a Test Methodology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Two testing methodologies are useful for testing application compatibility: Windows Fundamentals Test Methodology and Windows Applications Exploratory Test Methodology. Use Windows Fundamentals Test Methodology when your testing team thoroughly understands how an application works and how the groups who use the application expect it to perform. Use Windows Applications Exploratory Test Methodology, a more ad hoc form of testing, when this is not the case.

For more information about Windows Fundamentals testing and Windows Applications Exploratory testing, see "Developing Test Cases" later in this chapter.

# Identifying the Applications to Test

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most medium-to-large organizations use so many applications that it is not possible to test them all thoroughly. Consequently, most organizations prioritize the applications they plan to test.

After you have consolidated your application base, so that only the applications that you plan to support with the new operating system are listed, you can prioritize your list based on factors such as whether applications are compatible with Windows Server 2003, how critical they are to your business operations, and the number of users who depend upon them. If the amount of time available for testing or your resources are limited, you can use this prioritized list to decide whether to perform a reduced set of tests on an application or set of applications or whether to not test certain applications at all.

**Important**

- If time and resources permit, full testing of all applications is recommended.

The ultimate goal of prioritizing inventoried applications is to identify the core group of applications that must function properly before you begin to roll out the new version of the Windows operating system. Whether you use existing priority categories — such as those you use for disaster recovery — or use the categories suggested here, the categories need to account for this core group of applications.

Use the information in the following sections and any other considerations unique to your particular situation to establish the testing priority for each application.

**Prioritizing Application Testing Based on Levels of Compatibility**

Using the compatibility information that you obtained through the Application Compatibility Analyzer or the Windows Catalog, you can begin to prioritize the level of application testing. Use the information in Table 3.3 to assist with this task.

**Table 3.3 Application Compatibility Priority Categories**

| Application's Compatibility | Priority Recommendation |
|---|---|
| **Compatible**. The application has been tested and is deemed compatible with Windows XP Professional or Windows Server 2003. | If the application is used without customization, you might assign it low priority. However, it is recommended that you test compatible applications with your organization's Active Directoryand Group Policy infrastructure, at a minimum, to ensure that there are no problems in your particular environment.<br><br>If you use the application with customizations such as templates or macros, at the very least test the customizations. |
| **Compatible with Issues**. The application typically is compatible, but might have problems when run in certain contexts. | Consult the Application Details section of the Application Compatibility Analyzer report to identify the situations under which the application has problems, and test the application in those situations. Also test compatible applications with your organization's Active Directory and Group Policy infrastructure to ensure that there are no problems in your particular environment. |
| **Incompatible**. The application is not compatible with Windows XP Professional or Windows Server 2003. | If the Application Details section of the Application Compatibility Analyzer report lists fixes for the application, fully test the fixes within your environment. |
| **Unknown**. Either the application or its compatibility with Windows XP Professional or Windows Server 2003 is unknown to Microsoft. | Fully test the application if you plan to use it in your production environment. |

**Prioritizing Application Testing Based on How Critical Applications Are to Your Business**

By using the information that you collected from the groups who use the applications, you can prioritize the testing of your applications based on how critical they are to your organization's business. Use the categories listed in Table 3.4 to prioritize

your applications for testing using this criterion.

**Table 3.4 Business Priority Categories**

| Category | Impact of Failure | Priority Recommendations |
|---|---|---|
| **Mission-critical**. Required in order to collect revenue or to fulfill legal obligations; must be first back online after a disaster occurs. | High | Thoroughly test the application in the proposed environment even if it is listed as compatible with the new operating system. Situations in your network or domain infrastructure might cause unforeseen compatibility problems with a mission-critical application, which must be fixed before you deploy the application and new operating system. |
| **Business-critical**. Required in order to run the business infrastructure; should be second back online after a disaster occurs. | Moderate | Thoroughly test the application in the proposed environment; however, if an application is listed as compatible, and it has not been customized, you might consider testing only the most important functionality if time is limited. |
| **Required**. Required in order to run the business, but can be offline longer. | Low | Thoroughly test the application in the proposed environment; however, if an application is listed as compatible, and it has not been customized, you might consider not testing it if time is limited. |
| **Other**. Not required in order to run the business. | Very low | Thorough testing is not necessary for the application. |

**Prioritizing Application Testing Based on Usage**

In prioritizing the applications that you plan to test, consider the impact of the failure of an application on the organization as well as the number of employees who use it.

The report generated by the Application Compatibility Analyzer can tell you how many desktops in the organization each application is installed on. You should test an application that is used by many employees, regardless of whether the application is listed as compatible with the operating system. An example of this type of application is antivirus software that runs on every desktop in the organization.

If only a few employees use an application, but its failure will affect a great number of employees, you should test the application. For example, only a few people in the accounting department use the payroll software. However, if that software does not work properly, all employees might receive their paychecks late.

# Preparing for Application Testing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before application testing can begin, you must prepare your test environment. It is assumed that your organization has a test lab and has established standards and best practices for testing. If this is not the case, see "Designing a Test Environment" in this book for information about designing and setting up your test lab.

To prepare your test environment, you need to set up your test lab for application compatibility testing, install tools for testing applications and for debugging them, gather all of the information and components needed to install and test each application, and develop test cases. Figure 3.6 shows when these tasks occur within the application compatibility testing process.

**Figure 3.6 Preparing for Application Testing**



# Setting Up a Test Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To determine whether your applications are compatible with Windows XP Professional or Windows Server 2003, you must test them in a lab environment on computers that represent the hardware and software configurations found in your organization. These hardware and software configurations include variables such as all of the operating systems that interoperate in your environment; a mix of clients; the mix of applications to be installed on a single computer; and all the supporting programs, hardware components, and required files. Therefore, the size and complexity of the test lab is determined by the complexity of the applications to be tested and by the network environment where you plan to deploy them.

In addition to duplicating hardware and software configurations, you need to simulate the way that you install and use the applications in your production environment. For example, one organization has a production environment that encompasses thousands of applications running on both Windows NT Server 4.0 and Windows 2000. This organization uses scripted application installations to control the application environment and to prevent unauthorized copying or installation of third-party products. In addition, it standardizes drive mappings for production applications. For example, drive F is reserved for authorized applications for which users have read-only access, drive G is reserved for authorized applications for which users have read and write access, and drive J is reserved for the users' customized and personal settings, which are retained across sessions.

The test lab includes separate network shares to simulate the drive mappings. Windows XP Professional is installed on all of the test lab computers, except for two baseline computers that have Windows NT Server 4.0 and Windows 2000 installed. Each computer is set up with a core set of applications that are common to all users — such as an office productivity suite, Internet browser, antivirus application, and administrative tools. This core set of applications is pretested for compatibility. Each additional application is then tested using the production application installation script in an environment where the core applications are installed.

# Determining a Lab Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Depending on your circumstances and environment, you might decide to test applications in the same lab that you use to test your Windows Server 2003 family deployment. However, if you have the budget and the space, it might be more efficient to have a separate lab for testing applications. For example, if you have many applications, a lab that can be used full-time only for testing applications can make it easier for you to meet your schedule.

The lab that you use for testing applications can be centralized or distributed. If your application owners and testers are widely dispersed, a distributed lab can make resources more readily available during testing and also can help end users become involved in testing. Because distributed labs can require a greater investment in equipment and can be more complex to manage, you might combine both approaches. For example, you might use a large, centralized lab to test critical line-of-business applications that require specialized servers and configurations, and use smaller, distributed labs to test more independent workstation applications. A distributed lab in this case might be as simple as a workstation on a tester's desk.

When deciding your lab strategy, consider the following factors:

- The number and complexity of applications to be tested
- The location of application testers
- The location of application owners
- Budget and space constraints
- Hardware constraints
- Transportation time and costs

For more information about centralized and distributed testing facilities, see "Designing a Test Environment" in this book.

# Configuring the Lab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the test lab, configure test stations that can be independently managed and scheduled. If enough computers are available, include at least two workstations in each test station to support side-by-side testing. Configure one with the operating system that is currently deployed and one with the new operating system. If you do not have enough computers for each test station, provide a separate group of computers running the current operating system that you can assign as needed.

If you plan to deploy Windows XP Professional or Windows Server 2003 by performing system upgrades rather than clean installations, you need to test how applications behave on the new operating system without reinstalling them. To test, include a workstation that is running your current version of the operating system in each test station, so that you can install each application before you upgrade to the new operating system.

Provide each test station with its own dedicated monitor, keyboard, and mouse. To isolate problems that might occur at a test station, configure each one so that it can be detached quickly from the rest of the test network, for example, by unplugging a cable or making a programmatic change. Provide access to peripheral devices such as printers and scanners, so that testers can test the full functionality of applications.

To facilitate testing, configure the workstations with the following features:

- **Enable auditing to monitor permission failures.** Because the default settings for Windows XP Professional and Windows Server 2003 are more secure than those in earlier versions of Windows, it is not uncommon for an application to fail because permissions are invalid, especially when the application is installed or when it accesses files. Use the audit log and the event log to identify permissions that need to be changed to allow an application to be installed and to run. For more information about the event log and audit log, see Event Viewer overview in Help and Support Center for Windows Server 2003.

- **Enable crash dump files in Dr. Watson.** *Crash dump files* are binary files that capture the contents of memory when a fatal error occurs. Crash dump files can be loaded into a debugger. For information about configuring Dr. Watson to enable crash dump files, see "Using Dr. Watson" in Help and Support Center for Windows XP Professional.

Mirror the corporate domain structure in the lab, and duplicate user rights for both the domain and the workstations. If you plan to use Group Policy settings in your new environment, implement those policies in the lab.

As a rule, try to duplicate resources such as databases and client/server applications in the lab. If you cannot duplicate some resources, such as mainframe databases, and therefore need to access them over the corporate network, plan how to regulate and control the connection, and devise a quick way to terminate it if that becomes necessary.

As a best practice, rebuild the workstations and, possibly, the servers before testing each application. This approach minimizes the possibility of introducing problems that are not related to application compatibility and helps testers identify problems. To ensure a consistent starting point and to reduce setup time, use scripts or computer images to rebuild the workstations. Your baseline image should include your organization's core applications so that as you test each application, you also test for possible negative interactions between applications.

For more information about designing and configuring a test lab, see "Designing a Test Environment" in this book.

# Installing Application Testing Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After setting up the test lab, gather and install application testing tools. Be sure to install any testing tools that you use to automate tests, such as scripts and batch files, as well as the tools in the Windows Application Compatibility Toolkit that are used to apply application compatibility fixes.

**Testing Tools**
Some testing tools gather or define test data, others run tests, and still others monitor what happens to files and the registry when tests run. The following types of tools automate various steps in the testing process:

- Test automation tools, such as record and playback tools, or test case generators
- Custom scripts, such as installation scripts
- Data extraction tools, such as data warehouse tools
- Test databases, such as those created for regression testing
- Comparison and monitoring tools, such as tools that track file and registry changes

**Windows Application Compatibility Toolkit**
The Windows Application Compatibility Toolkit, provided by Microsoft, contains Windows Application Verifier, a tool that helps uncover heap problems and other problems that are difficult to locate, and Compatibility Administrator, a tool that applies application compatibility fixes to an application when it runs in order to resolve compatibility problems. The application compatibility fixes help resolve compatibility problems when you do not have access to the source code or when you do not have sufficient time to debug, modify, and implement code changes. Because application compatibility fixes do not modify source code, you must apply them whenever you install an application on a workstation.

# Installing Debugging Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If the source code for an application is available, and you want to resolve the underlying cause of compatibility problems by modifying the source code, use a debugging tool to investigate the problem. You can use your own debugging tools, or you can use Debugging Tools for Windows.

Debugging Tools for Windows is a collection of kernel-mode and user-mode debuggers provided by Microsoft. Applications run in one of two modes: *kernel mode* or *user mode*. Operating system services and some low-level programs, such as device drivers and antivirus applications, run in kernel mode. Kernel-mode applications can access any part of the system. User-mode applications are limited to their allocated area in memory and have restricted access to system interfaces and system data. Most end-user applications run in user mode.

Choose a tool based on the type of application that you need to debug (a kernel-mode or user-mode application) and whether you will use a local or remote computer for debugging. Most of the applications that you need to test for compatibility with Windows XP Professional are user-mode applications; however, you might need to test kernel-mode applications if you develop device drivers or other kernel-mode applications in-house. If you have kernel-mode applications that have been developed by third-party vendors, contact those vendors for upgrades.

Table 3.5 lists each debugger included in Debugging Tools for Windows, telling the processing mode for which it is designed and whether it must be used from the local computer or a remote computer.

**Table 3.5 Debugging Tools for Windows**

| Tool | Processing Mode | Debugging Computer |
|---|---|---|
| Microsoft Console Debugger (CDB) | User mode | Local or remote |
| Microsoft NT Symbolic Debugger (NTSD) | User mode | Local or remote |
| Microsoft Windows Debugger (WinDbg) | User mode and kernel mode | Local or remote |
| Microsoft Kernel Debugger (KD) | Kernel mode | Remote |

**Important**

- Debugging information for all system applications, drivers, and DLLs resides in separate files known as *symbol files*. By separating debugging information from the application, symbol files allow applications to be smaller and faster, but still permit them to be debugged if the symbol files are installed. If you use a Microsoft debugger on a third-party application that has no symbol file, the debugger tries to extract symbolic information from the executable file and indicates that the characters are export symbols. These are not actual symbols, and they should not be used for debugging analysis.

In addition to the debuggers, Debugging Tools for Windows provides some other tools that can help you during debugging. The Logger and LogViewer, in particular, can help you debug application compatibility problems. Logger and LogViewer can help you analyze the API calls made by an application. If an application worked on an earlier version of Windows, differences in API implementations probably are causing the incompatibilities. One of the most effective ways to debug a third-party application is to examine the Win32 APIs that it calls.

**CDB and NTSD**

Microsoft Console Debugger (CDB) and the Microsoft NT Symbolic Debugger (NTSD) are console debuggers that allow you to analyze user-mode memory and constructs on target applications that are either console or graphical Windows applications. CDB and NTSD have the same functionality, except that NTSD opens a second command window when it is started. When you use these tools to debug an application that has failed, they can obtain a stack trace or display invalid parameters. You also can use them to observe the behavior of an application that is functioning properly.

By using CDB or NTSD, you can display and execute program code, set breakpoints, and examine and change values in memory. CDB and NTSD can disassemble binary code and display assembly instructions. They also enable you to analyze source code directly.

**Note**

- A copy of NTSD is installed in the \Windows*system32* directory when Windows XP Professional or Windows Server 2003 is installed, but the version included with Debugging Tools for Windows is the most recent version. The version of NTSD installed with Windows XP Professional or Windows Server 2003 can be used only for local debugging.

**WinDbg**

WinDbg is a graphical interface debugger that supports full source-level debugging for the Windows kernel, kernel-mode drivers, and system services, as well as user-mode applications and drivers. WinDbg is the most versatile debugger in Debugging Tools for Windows. It can monitor variables, CPU registers, and parameters passed in function calls. You also can use it to review a crash dump file from Dr. Watson, which captures the state of the system when a failure occurs. By using WinDbg, you can view source code, set breakpoints, and view variables (including C++ objects), stack traces, and memory.

**Logger and LogViewer**

Logger monitors the actions of a user-mode application and records all of its API calls. You can display the resulting information in a debugger, save it as a text file, or display it in an interactive format in the LogViewer.

You can start Logger by using the stand-alone Logger.exe program, or by starting NTSD or WinDbg, and using the LogExts.dll debugger extensions. Both methods produce the same type of log output. By starting Logger with NTSD or WinDbg, however, you have access to all the functionality of both a debugger and Logger.

After logging the API calls, you can use LogViewer to view the log file. You can view the list of all the API calls in the order in which they were logged, filter the list, export the list to a text file, and view the values of parameters sent to or received from each API call.

**Obtaining Debugging Tools for Windows**

Debugging Tools for Windows are available for downloading. To use the tools, you should also install the appropriate symbol files for the version of Windows that is running the application that you need to debug. Symbol files contain a variety of information that is not required to run a program, but that is useful for debugging. This information is contained in a symbol file instead of in the executable file for performance purposes. To download the most recent version of Debugging Tools for Windows and the appropriate symbol files, see the Debugging Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. For more information about debugging, including information about debugging tools, see *Using Debugging Tools for Windows*, which is installed when you install the Debugging Tools for Windows. To open *Using Debugging Tools for Windows* on Windows XP Professional, click **Start**, click **All Programs**, click **Debugging Tools for Windows**, and then click **Debugging Help**. To open *Using Debugging Tools for Windows* on Windows Server 2003, click **Start**, click **Programs**, click **Debugging Tools for Windows**, and then click **Debugging Help**.

# Gathering Application Information

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin testing, gather everything that you need for testing each application. Microsoft customers have found that the more effort they put into preparation and information gathering before beginning to test, the faster and more easily testing progresses. Having all related information readily available also makes it easier to resolve any problems that testing uncovers.

For each application, gather its software and documentation along with supporting components required to install and run the application, and find out whether the application is dependent on the network basic input/output system (NetBIOS).

**Gathering Software and Supporting Information**

Work with the application owners within your organization to locate each application and its documentation. Before installing each application in the test lab, gather the supporting information and components that are required to run it. For example, gather the following information:

- Required security clearances, such as user IDs, passwords, database permissions, group memberships, and so on.
- Required servers and IP addresses.
- Accurate installation instructions, installation media, and contact information for application experts.
- Service packs or patches that have been applied or are required.
- Software dependencies, such as the need for SQL Server or third-party components.
- Whether the application is client-based or server-based, and where each component resides — on the client or on the server.
- The version of Windows on which the application currently runs.
- Application configuration information, such as search paths and working directories.
- The directory path where the application files are installed.
- Registry settings that the application creates or modifies.
- The location where user data is stored.
- URLs for Web applications.
- The location of symbol files that were created when the application was compiled, for use with debuggers.

**Identifying NetBIOS Dependencies**

Determine whether any applications require NetBIOS. If any do, ensure that NetBIOS is included in the new Windows environment.

On computers running Windows 95, Windows 98, and Windows NT, NetBIOS is enabled by default. On computer running Windows 2000, Windows XP Professional, and Windows Server 2003, NetBIOS is disabled by default. If an application runs on a system without NetBIOS and calls functions that use NetBIOS, the functions return errors. In such a situation, you must either enable NetBIOS or change the application to use functions that do not depend on NetBIOS.

# Developing Test Cases

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before starting to test, develop and document *test cases*, which describe the tests to be performed and, if possible, the expected results. The number and complexity of test cases that you need depends in part on how extensively you need to test an application. The extent of testing required for an application depends on several factors, including the application's level of compatibility, how critical the application is to your business, and how often the application is used. In addition, the extent of testing required depends on whether the application is a third-party product, has been customized, or was developed in-house. For third-party applications, the extent of testing also depends on whether the application is known to be compatible with the new version of Windows that you are deploying.

Applications that require limited testing fall into the following categories:

- Applications that are listed as Compatible in the Windows Catalog
- Applications that are not required to run the business
- Applications that are required to run the business, but can be offline for long periods of time without impacting the business
- Applications that do not impact the organization severely if they fail

Usually, you do not need to develop complex or wide-ranging test cases for these types of applications.

Applications that require in-depth testing fall into one or more of the following categories:

- Applications that have been developed in-house
- Applications that have been customized for your organization
- Applications whose compatibility is listed as Unknown by the Application Compatibility Analyzer
- Applications that are listed as Compatible with Issues by the Application Compatibility Analyzer
- Applications that are not listed in the Windows Catalog

For these applications, you need to develop a wide range of test cases that encompass all the functionality required in your environment.

Applications that are listed as Compatible by the Application Compatibility Analyzer or in the Windows Catalog or that are listed as Designed for Windows XP in the Windows Catalog do not need to be tested as extensively as those in the preceding list. Even so, you should perform general testing to ensure that these applications perform as expected, given the way they are used with your hardware and software configuration. Develop test cases for the functions that are used most frequently and for any atypical functions or unique techniques that are used. For example, if your organization uses scripts to manage the desktop environment, test the applications with those scripts. If your organization hosts applications with Terminal Server, remember to include test cases to test the applications with Terminal Server.

The testing methodology that you choose for each application determines how you develop the test cases. Many testing methodologies exist, and you might already have one that you use. If you do not already have a methodology for developing test cases, use one of the following methodologies:

- Windows Fundamentals Test Methodology, which helps you develop test cases when you have an in-depth understanding of the application
- Windows Applications Exploratory Test Methodology, which helps you develop test cases as you learn about an application that is unfamiliar to you

# Using Windows Fundamentals Test Methodology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Fundamentals Test Methodology provides guidelines for developing test cases when you are familiar with an application, its end users, and its use in your environment. When you have this level of knowledge about an application, you can develop test cases based on the way your organization uses it.

When using the Windows Fundamentals test method, define the test cases ahead of time, before you start testing the application. Develop test cases for the following four areas: installation, operating system interaction, application functionality, and uninstallation. For a checklist to help you track your Windows Fundamentals tests, see the "Windows Fundamentals Testing Checklist" (PDP_APPA_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see the "Windows Fundamentals Testing Checklist" on the Web at http://www.microsoft.com/reskit).

**Developing Test Cases for Installation**

Even if you plan to upgrade from an earlier version of Windows to Windows XP Professional or Windows Server 2003, you need to test application installation. This phase of testing verifies that an application can be installed correctly — whether by a user, a power user, or an administrator account — and that the various installation methods — typical, network, full, and minimal — work correctly. If you use Windows Installer, you should also test the application's advertising and publishing functions.

Develop test cases to verify that the application:

- Installs successfully with the types of user account to be used in your environment (user, power user, or administrator).
- Installs successfully with the types of installation to be performed in your environment (minimal, typical, full, or custom).
- Installs successfully with the methods to be used in your environment (network installation or a local installation using **Add or Remove Programs**).
- Is listed correctly in **Add or Remove Programs**.
- Checks the Windows version correctly.
- Installs on a partition that has more than 2 GB of free disk space.
- Is listed on the **Start** menu, and has desktop shortcuts for all users (if applicable).

If you use Windows Installer, develop test cases to verify that the application:

- Advertises and publishes correctly.
- Uses system variables for the locations of user data settings.
- Respects Windows File Protection.
- Supports DLL redirection of shared components.
- Does not use outdated techniques such as placing .ini files in the system folder.

**Developing Test Cases for Operating System Interaction**

Tests of operating system interaction verify that an application works correctly with the fundamental features of the operating system. These tests use standard Windows operations, such as opening and closing files, manipulating windows, printing, and viewing Help and Support Center topics.

The following lists provide a starting point for identifying test cases that verify an application's interaction with the operating system. This list is not exhaustive; it merely provides an idea of the scope of features to test. The "Windows Compatibility

Checklist" white paper, which is included in the Windows Application Compatibility Toolkit, is oriented toward applications under development and might provide additional ideas. If an application uses peripherals such as scanners or digital cameras, you might also need test cases to verify the application's interaction with these devices.

**Basic application testing**

Develop test cases to verify that the application:

- Interacts with the other applications used in your environment.
- Starts successfully in a variety of ways (from the **Start** menu, from shortcuts, from the autorun file on the application CD, from the console, and from the Quick Launch bar).
- Starts successfully when you double-click a file that has the file name extension associated with the application, if applicable.
- Starts successfully when another instance of the application is running (if applicable; some applications do not allow multiple instances).
- Has associated shortcuts on its **Start** menu, such as for documentation or related programs, that work correctly, if applicable.
- Handles copy, paste, and edit functions correctly.
- Handles changes to the monitor resolution correctly (such as a change from 640 × 480 pixels to 1024 × 768 pixels).
- Works correctly with Windows display themes.
- Starts successfully when you open an application document by using the wheel on a wheel mouse, if applicable.
- Scrolls correctly when you use the wheel on a wheel mouse, if applicable.
- Opens Help from the **Help** menu and with the F1 key, if applicable.

**Testing file system interaction**

Develop test cases to verify that the application can successfully:

- Handle long file names and universal naming convention (UNC) paths.
- Save a file to, and open a file from, an NTFS file system folder that has restricted access permissions.
- Save a file to, and open a file from, a disk partition that has more than 2 GB of free disk space.
- Save a file that has a 255-character file name to a UNC path and to a mapped drive, and open the file from both locations.
- Save a file that has a 255-character file name to the My Documents folder, and open it from that folder.
- Save and open a file that has supported special characters — such as a plus sign, a comma, a semicolon, an equal sign, or brackets ([]) — in the file name.
- Save and open all document types supported by the application.
- Import and export files of all formats that are supported by the application, such as JPG, GIF, or BMP.
- Open documents on the NTFS file system, even when they were created on the FAT16 or the FAT32 file system.
- Open and save a document on both a drive formatted for the FAT16 file system and a drive formatted for the NTFS file

system.

- Open and save a document on both a drive formatted for the FAT32 file system and a drive formatted for the NTFS file system.
- Save and open a file on a distributed file system (DFS).
- Open documents that were copied from Windows NT Server 4.0–based or Windows 2000–based NTFS computers to Windows XP Professional–based NTFS computers.

**Testing printing**
Develop test cases to verify that the application can successfully:

- Print a document to a printer that has a long name.
- Print a document to a printer by using a UNC path.
- Print a document to a local printer and to a network printer.
- Print a document to a PostScript printer.
- Print to a file on a generic text printer.
- Print to a file on a color inkjet printer.
- Handle a print command when no printer drivers are installed.
- Print a document from an LPT port other than LPT1 or LPT2.
- Handle a print command when the printer is turned off while a document is printing.
- Print a document in both portrait and landscape orientations.

**Testing Advanced Configuration and Power Interface (ACPI) support**
Develop test cases to verify that the application:

- Handles ACPI notification and events correctly.
- Allows the monitor to turn off correctly when turn-off time is set to 1 minute.
- Allows the hard drive to turn off correctly when turn-off time is set to 2 minutes.
- Goes into standby and hibernation, and then resumes without loss of functionality.
- Wakes from standby and hibernation — whether triggered by normal time-out or by a low battery alarm — without losing data.

**Testing support for multiple monitors**
Develop test cases for configurations with multiple monitors, if appropriate, to verify that the application handles the following correctly:

- Monitor placement, left and right
- Monitor placement, top and bottom
- Different monitor resolutions on both monitors

- Different monitor color depths on both monitors
- Spanning across monitors, both from left to right and from top to bottom
- Mouse movement, back and forth from left to right and up and down from top to bottom, on both monitors
- Windows set to maximum size and minimum size on both monitors

**Developing Test Cases for Application Functionality**

Tests for application functionality validate that the application features used in your environment work as expected and that known input yields expected results. You need to test application functionality whether or not you have access to source code.

If you have developed or customized an application internally, you might have test scripts and data you can use to verify the application's compatibility with the new operating system. If you do not have test scripts, work with the people in the organization who are familiar with the application to develop test cases that exercise all essential functions.

**Developing Test Cases for Uninstallation**

Tests for removing an application verify that all components are removed from the computer without causing problems.

Develop test cases to verify the following:

- The application is removed from **Add or Remove Programs**.
- All shortcuts are removed from the **Start** menu.
- All files in the installation directory are removed.
- No system files or shared components are removed.
- All application-specific registry entries are removed.

# Using the Windows Applications Exploratory Test Methodology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Applications Exploratory Test Methodology defines a process for learning about an application so that you can identify the most important aspects to test when no one who is knowledgeable about the application is available for testing.

When you use the exploratory test methodology, you do not define the test cases in advance. Rather, you define the test cases and execute the tests as you explore and learn about the application.

**Exploratory Testing Overview**
The purpose of exploratory testing is to determine, in a limited amount of time, whether an application works well on the new version of Windows for the majority of users. Exploratory testing addresses both application functionality and stability. This method is not a comprehensive test method, and it is not a substitute for a more formal method if application experts are available.

Exploratory testing is an interactive process during which you concurrently explore the application, design the tests, and execute the tests. Although the method is freeform in many respects, the process includes specific tasks, objectives, and deliverables, making it a systematic process with auditable results.

This method is based on the *General Functionality and Stability Test Procedure* developed by James Bach. For more information about the exploratory test methodology, see the "Windows Applications Exploratory Test Procedure" white paper in the Application Compatibility Toolkit.

**Exploratory Testing Process**
The exploratory testing process consists of five tasks:

1. Identify the purpose of the application.
2. Identify the functions that the application supports.
3. Identify potential areas of application instability.
4. Document issues and questions that arise about the application during the exploration process.
5. Create a test outline, and use it to execute the tests.

As you identify the purpose, functions, and areas of potential instability, you create a test case outline and use it to test each function and to record problems. The results of the exploratory testing process are notes and documented application failures.

Because it is not possible to test all functions of an application in a limited amount of time, you can simplify the testing process by making risk-based decisions about how much attention to give each function. To help you make these decisions, the methodology categorizes functions as *primary* or *contributing*.

To determine which functions are primary and which are contributing, you must understand the purpose of the application. A *primary function* is associated with the purpose of the application and is essential to that purpose. Any function that contributes to the utility of the application, but is not essential to it, is a *contributing function*. Typically, primary functions warrant the most testing. Sometimes a group of contributing functions might be considered a single primary function, or a single primary function might be separated into primary and contributing subfunctions. For more information about how to classify a function, see the "Windows Applications Exploratory Test Procedure" white paper in the Application Compatibility Toolkit.

After identifying and categorizing the functions, test all of the primary functions that you can reasonably test in the time available. Then test a sample of significant contributing functions. Many contributing functions probably will be tested as you test the primary functions.

Also test areas where application instability is most likely to occur, especially in primary functions. Functions that are potentially unstable include:

- Functions that interoperate with other applications.
- Functions that handle events that are external to the application, for example, waking up a sleeping computer when a fax arrives.
- Functions that make intensive use of memory.
- Functions that interact extensively with the operating system.
- Functions of unusual complexity.
- Functions that change operating parameters, such as preference settings.
- Functions that modify the operating system configuration.
- Functions that intercept or recover from errors.
- Functions that replace basic operating system functions.
- Any function or set of functions that involves multiple simultaneous processes.
- Functions that manipulate multiple files simultaneously.
- Functions that open files over a network.

# Testing for and Resolving Compatibility Problems

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After developing your test cases, you can start to test applications and record the results. Perform your basic testing first, using either the Windows Applications Exploratory Test Methodology or the Windows Fundamentals Test Methodology. Then, if you have access to source code, use Application Verifier to test for subtle problems that are difficult to find. As you detect compatibility problems, use Compatibility Administrator to apply and test application compatibility fixes to resolve the problems. For some special cases in which application compatibility fixes do not resolve the problem, you might need to try special techniques. Figure 3.7 shows when these testing and problem resolution tasks occur during the application compatibility testing process.

**Figure 3.7 Testing for and Resolving Compatibility Problems**



# Testing Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you conduct tests, evaluate and record the results. Before you assume the problem is a Windows compatibility issue, always perform the following two steps:

- **Verify that the problem is not a test environment issue**. The problem might be the result of a hardware, software, or network configuration issue rather than an application compatibility problem.

- **Verify that the problem does not occur on the version of Windows you currently use**. Perform exactly the same test on the version of Windows on which the application currently runs to see if it works correctly there. If the same error occurs on the current platform, the problem is not a compatibility issue specific to Windows XP Professional or Windows Server 2003.

The testing process is an iterative process of testing, identifying problems, resolving problems, and retesting. The steps in the process vary depending on the nature of the problem encountered and the nature of the resolution. Problems that are resolved by applying an application compatibility fix, for example, are managed differently from problems resolved by modifying source code or obtaining vendor upgrades.

For problems that can be resolved by applying an application compatibility fix, the resolution phase can be built into the testing phase: As you encounter problems, you apply application compatibility fixes and retest until the application is ready to deploy.

If you modify source code, the iterative process is the more traditional cycle of testing, debugging, coding and recompiling, and retesting. If you cannot identify an application compatibility fix to resolve the problem, and you do not have access to source code, you need to contact the application vendor for an application upgrade.

Figure 3.8 illustrates the testing process.

**Figure 3.8 Testing Applications**



# Testing with Windows Application Verifier

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you have access to source code, Windows Application Verifier can help you uncover subtle stability, reliability, compatibility, and security problems. Windows Application Verifier is not an automated test tool. To use it, you must run an application with Application Verifier enabled and exercise all of your test cases.

Windows Application Verifier helps you to identify code that needs to be modified by running in either of two modes of operation: debugging or logging. In *debugging mode*, it detects problems in the application and transfers the results to a debugger. In *logging mode*, it writes all of the output that it generates to a text file for later review.

Figure 3.9 shows the main window of Application Verifier.

**Figure 3.9 Application Verifier Main Window**




To start Application Verifier on Windows XP Professional, click **Start**, click **All Programs**, click **Microsoft Windows Application Compatibility Toolkit**, and then click **Windows Application Verifier**. On Windows Server 2003, click **Start**, click **Programs**, click **Microsoft Windows Application Compatibility Toolkit**, and then click **Windows Application Verifier**.

To use Application Verifier, you must be logged on as an Administrator. Click **Add** to add the applications that you want to test in the left pane. To test an application, select it in the left pane, select the tests that you want to apply to it in the right pane, and then start the application in your normal way and run the test cases. After you add an application to the list, Application Verifier is enabled for that application, whether you start the application from this tool, from a shortcut, or from the command line. The tests that you select are applied whenever you run the application on that computer until you remove the application from the list in Application Verifier.

**Note**

- Application Verifier tests require the symbol files for the version of Windows and the application that you are testing.

The tests detect the most common application problems, such as incorrect version checking, bad registry usage, and hard-coded file paths. You can view a list of the detected problems in the Application Verifier log, which is located in the \Documents and Settings\All Users\Shared Documents\AppVerifierLogs folder with a file name corresponding to the name of the

executable file for the application. You also can export the log to a text file so that you can send it to others.

Some of the test options provided by this tool are designed to cause the application to stop responding when it encounters a problem. These problems are not written to the Application Verifier log. To determine the cause of the problem in such a case, run the application under a debugger with Application Verifier running. When using a debugger, you can set Application Verifier so that it breaks into the debugger for any event that is written to the Application Verifier log.

For more information about the Application Verifier tool, see Windows Application Compatibility Toolkit Help.

Table 3.6 describes the tests that you can perform with Application Verifier.

**Table 3.6 Application Verifier Tests**

| Test | Description |
|---|---|
| **PageHeap — Detect heap corruptions** | Performs regular checks of the heap and adds guard pages at the end of each allocation to catch possible heap overruns. |
| **Locks — Check lock usage** | Looks for common errors with locks, and displays the output in a separate debugger application. This test can cause access violations if an error is found. |
| **Handles — Detect invalid handle usage** | Checks for common problems with handles, and displays the output in a separate debugger application. This test can cause access violations if an error is found. |
| **Stacks — Check for adequate stack** | Disables stack growth, causing a stack overflow exception if the initial allocation is too small. |
| **LogStartAndStop — Logs start and stop** | Enters log information when the application starts or stops, making it easier for you to review test data. |
| **FilePaths — Checks system path usage** | Monitors the application's attempts to obtain file path information to determine if the application uses hard-coded paths or a nonstandard method of gathering the path information. This test can cause the application to stop responding if the application uses an improper method for determining file paths. |
| **HighVersionLie — Checks version handling** | In the past, many applications were written to run on a specific version of Windows. This test returns a very high version number to test how the application handles a version that does not yet exist. |
| **RegistryChecks — Checks registry usage** | Checks how the application uses the system registry to determine whether the application uses any inappropriate or dangerous calls. Logs any problems that it detects. |
| **WindowsFileProtection — Logs changes to Windows File Protection files** | Verifies that applications are not improperly replacing files that are protected by Windows File Protection. |
| **DXFileVersionInfo — Logs DirectX file checks** | Logs information about any attempt that the application makes to check the version of Microsoft® DirectX® application programming interface installed on the local computer. |
| **LogRegistryChanges — Logs registry changes** | Logs any changes that the application makes to the registry. |
| **LogFileChanges — Logs file system changes** | Logs information about any change that the application makes to the file system. |
| **ObsoleteAPICalls — Logs calls made to obsolete APIs** | Logs any attempt to use a *deprecated API*, which is an API that is still supported but is likely to be unsupported in the future. |

| | |
|---|---|
| **KernelModeDriveInstall — Logs installation of kernel-mode drivers** | Logs data about any kernel-mode driver that is installed for use by the application. |
| **SecurityChecks — Logs potential security problems** | Identifies many common security problems. |

# Applying and Testing Solutions with Compatibility Administrator

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Compatibility Administrator is an all-purpose tool that you can use to apply and test application compatibility fixes and compatibility modes, and then package them for deployment. Application compatibility fixes are prepackaged solutions provided by Microsoft to resolve known compatibility problems. Compatibility modes are collections of application compatibility fixes that resolve problems that commonly occur under certain circumstances. By using Compatibility Administrator, you can fine-tune the solutions to be applied to an application by mixing and matching application compatibility fixes and compatibility modes.

To simplify distribution, you can create custom databases that consolidate all of the solutions that are required in order to support multiple applications. For example, you can create one custom database for each standardized desktop in a managed environment and then deploy those databases to the computers running Windows XP Professional and Windows Server 2003 in your environment.

In addition, use Compatibility Administrator to:

- Create custom compatibility modes.
- Create custom Application Help messages.
- Search the computer's hard disk for applications whose compatibility problems are resolved with application compatibility fixes, compatibility modes, or Application Help messages.
- Query compatibility databases for solutions that meet specified criteria.
- Temporarily disable or enable application compatibility fixes.

Use this tool whether or not you have access to the source code for your applications. You do not need to log on under an Administrator account to run Compatibility Administrator, but you do need to log on under an Administrator account to install a custom database.

# Strategies for Using Compatibility Administrator

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you have determined that an application has compatibility problems, you can use the Compatibility Administrator to apply one or more application compatibility fixes to resolve the problems. The steps that follow are a suggested approach to isolating the appropriate resolution.

To resolve compatibility problems for an application by using the Compatibility Administrator:

1. Follow the suggestions contained in Application Help messages that appear for the application. For example, an Application Help message might indicate where you can obtain an update.

2. To quickly determine whether application compatibility fixes resolve the problems, try a prepackaged compatibility mode. Start with the compatibility mode that targets the environment in which the application currently runs. For example, if the application currently runs on Windows 98, use the Windows 98/Windows Me compatibility mode, or, if you are testing as a restricted user, try Limited User Access (LUA) mode. You might need to try several compatibility modes until you find one that works. Include the color, resolution, and visual themes compatibility modes, if necessary.

3. If a compatibility mode does not resolve all of the problems, try a combination of a compatibility mode and application compatibility fixes. When you apply a compatibility mode, you can add or remove any combination of application compatibility fixes. For information about which application compatibility fixes to apply for different types of problems, see "Determining the Application Compatibility Fix to Apply" later in this section.

If you resolve compatibility problems by applying a compatibility mode, you should then determine the precise mix of application compatibility fixes required for the application by iteratively deleting application compatibility fixes and retesting. Doing so can prevent introducing new problems with application compatibility fixes that are not necessary. However, you might need to balance the very low risk of introducing new problems against the time you have available for testing applications.

Follow these guidelines as you apply and test application compatibility fixes and modes:

- Do not use compatibility modes to bypass version warnings for the installation programs of kernel-mode applications such as antivirus, firewall, CD burning, disk management, or backup applications. Some of these applications are designed for a specific operating system and, therefore, they intentionally prohibit the use of later versions. Installing such applications can potentially cause system instability. Replace such applications with ones designed specifically for Windows XP Professional or Windows Server 2003.

- Before you start testing your applications, be sure that you have installed the most current critical updates and operating system updates. Critical updates and operating system updates can contain application compatibility fixes. You can use Windows Update or Software Update Services (SUS) to obtain critical updates and operating system updates. To access the Windows Update Web site, click the Home button in Help and Support Center for Windows Server 2003, and then, under **Pick a Task,** click **Keep your computer up-to-date with Windows Update**. For more information about SUS, see "Deploying and Distributing Applications and .Sdb Files" later in this chapter.

- If you find a combination of application compatibility fixes that you think will work for multiple applications, create a custom compatibility mode. If you need to include a custom compatibility mode in a different custom database, copy it. For information about how to create a custom compatibility mode and copy it to a different database, see "Creating Custom Compatibility Modes" later in this chapter.

- For simplicity in deployment, create as few custom databases as possible. Even if you create several custom databases during the testing process, you can consolidate them into a single database by using Compatibility Administrator. For information about how to copy to a database, see "Creating Custom Compatibility Modes" later in this chapter.

# Applying and Testing Application Compatibility Fixes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you encounter an application that has a compatibility problem, use Compatibility Administrator, shown in Figure 3.10, to select application compatibility fixes, either singly or in combination, and then test the application to see if the fixes resolve the problem. To find application compatibility fixes for a wide range of problems, use the "Compatibility Solutions Spreadsheet" (PDP_APP_1.xls) on the Windows Server 2003 Deployment Kit companion CD (or see the "Compatibility Solutions Spreadsheet" on the Web at http://www.microsoft.com/reskit) or, see Table 3.9 through Table 3.31.

**Figure 3.10 Compatibility Administrator**




To start Compatibility Administrator on Windows XP Professional, click **Start**, click **All Programs**, click **Microsoft Windows Application Compatibility Toolkit**, and then click **Compatibility Administrator Tool**. On Windows Server 2003, click **Start**, click **Programs**, click **Microsoft Windows Application Compatibility Toolkit**, and then click **Compatibility Administrator Tool**.

The left pane displays the contents of all the application compatibility databases on the computer. When you select an application that is listed in a database, the right pane displays detailed information. For example, if you select an application in a custom database in the left pane, the right pane displays the application file name, the compatibility mode or application compatibility fixes that apply to the application, and the application matching information.

**To select and test application compatibility fixes**

1. On the **File** menu, click **New** to create a new custom database.

2. On the **Database** menu, click **Create New,** and then click **Application Fix**.

3. Follow the steps in the Create New Application Fix Wizard to select the application to test and the compatibility mode or application compatibility fixes to apply, and then click **Test Run** to test the application from the wizard.

4. Continue selecting application compatibility fixes and retesting them until the problem is resolved.

5. Continue following the steps in the wizard to select the type of information that you want to use to match the application to the application compatibility fix.

6. Click **Finish** to save the application compatibility fixes in the database.

7. On the **File** menu, click **Save** to save the custom database. Provide a descriptive name for the database; the system appends an .sdb file name extension. Then save the database in the *systemroot*\AppPatch\Custom folder (or any other folder that you choose).

Alternatively, you can use this procedure to add more applications, with their associated application compatibility fixes, to an existing custom database. You also can copy application compatibility fixes from other application compatibility databases by using the **Copy** command on the shortcut menu, which is displayed by right-clicking the application name in the left pane of Compatibility Administrator.

This procedure saves the custom database but does not install it on the test computer. This means that the next time you run the application on the test computer without Compatibility Administrator running, the application compatibility fixes will not be applied to the application. To apply the application compatibility fixes to an application every time that you run it, you must install the custom database on the computer on which the application will run.

**To install a custom database on the test computer**

1. In Compatibility Administrator, select the custom database that you want to install.

2. On the **File** menu, click **Install**.

### Note

- You can run Compatibility Administrator in expert mode by typing **compatadmin.exe /x** at the command line. Expert mode provides more than 200 additional fixes, many of which are application-specific, that you can apply to a custom database. For more information about expert mode, see the Application Compatibility Toolkit version 3.0 or later.

For information about deploying and installing custom databases on Windows XP Professional or Windows Server 2003 computers, see "Deploying and Distributing Applications and .Sdb Files" later in this chapter.

# Applying and Testing Compatibility Modes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Sometimes you might want to get an application to work quickly by applying all of the application compatibility fixes that are commonly required by applications written for a specific computing environment, instead of testing for the specific application compatibility fixes needed by the application. In this case, you can apply a prepackaged compatibility mode. A *prepackaged compatibility mode* is a group of application compatibility fixes that resolve problems that commonly occur in certain environments. For example, one prepackaged compatibility mode includes application compatibility fixes that address problems typically found in applications that are run by users with limited user rights.

Table 3.7 lists the prepackaged compatibility modes that Compatibility Administrator provides and describes what they do.

**Table 3.7 Prepackaged Compatibility Modes**

| Prepackaged Compatibility Mode | Description |
|---|---|
| Windows 95 | Resolves the most common Windows 95 incompatibilities. |
| Windows 98 / Windows Me | Resolves the most common Windows 98 and Windows Millennium Edition incompatibilities. |
| Windows NT 4.0 SP 5 | Returns version credentials for Windows NT Server 4.0 with Service Pack 5. |
| Windows 2000 | Returns version credentials for Windows 2000. |
| 256 Color | Sets the screen to 256 colors; this is most useful for games and educational titles. |
| 640 x 480 Screen Resolution | Sets the screen resolution to 640 x 480 pixels; this is most useful for games and educational titles. |
| Disable Visual Themes | Disables Windows XP Professional or Windows Server 2003 visual themes. |
| International | Handles double-byte character sets. |
| Limited User Access | Redirects file-system and registry write requests to nonrestricted areas when the user has insufficient access permissions; useful for applications written for versions of Windows that have no security-related user accounts. |
| Profile Setup Support | Installs the application for all users of a computer rather than only for the current user. |

You can apply a prepackaged compatibility mode in two ways:

- Use Compatibility Administrator to add a prepackaged compatibility mode to a custom application compatibility database.
- Use a batch file or script to apply a prepackaged compatibility mode directly to an application.

**Applying Prepackaged Compatibility Modes by Using Compatibility Administrator**

You can use Compatibility Administrator to add any of the prepackaged compatibility modes to an application compatibility database. The procedure for doing this is identical to the procedure described in "Applying and Testing Application

Compatibility Fixes" earlier in this chapter. However, instead of adding application compatibility fixes to an application compatibility database, you add one of the prepackaged compatibility modes listed in Table 3.7.

**Applying Compatibility Modes in Batch Files and Scripts**

If you have a batch file or script that executes programs that have compatibility problems, you can apply prepackaged compatibility modes from the script or batch file instead of packaging them in a custom database and deploying the database to the computer where the programs run. When you apply compatibility modes in a batch file or script, you can ensure that they are recursively applied to processes initiated by any command in the batch file.

The syntax for applying compatibility modes in a script or batch file is:

```
set __COMPAT_LAYER=[!]CompatModeName1 [CompatModeName2]
```

The keyword **set** is followed by a space and then two underscores (__). *CompatModeName* represents the short name for the compatibility mode to be applied, as defined in Table 3.8. For example, the Windows 98 compatibility mode is "Win98." The parameter "!" specifies that a compatibility mode is not to be applied if other application compatibility fixes are already being applied to the application.

**Table 3.8 Scripting Short Names for Compatibility Modes**

| Prepackaged Compatibility Mode | Scripting Short Name |
|---|---|
| Windows 95 | Win95 |
| Windows 98/Windows Me | Win98 |
| Windows NT 4.0 SP 5 | NT4SP5 |
| Windows 2000 | Win2000 |
| 256 Color | 256Color |
| 640 x 480 Screen Resolution | 640x480 |
| Disable Visual Themes | DisableThemes |
| International | International |
| Limited User Access | LUA |
| Profile Setup Support | ProfilesSetup |

For example, to apply the Windows 98 and Disable Visual Themes compatibility modes, use this command:

```
set __COMPAT_LAYER=Win98 DisableThemes
```

You can disable the compatibility mode after applying it in this way by using the same command without specifying the compatibility mode. For example, you can cancel the compatibility mode if it is no longer needed at a certain point in the batch file. Any processes that start while the compatibility mode is in effect continue to run with the compatibility mode activated until they terminate. The syntax for canceling a compatibility mode in a script or batch file is:

```
set __COMPAT_LAYER=
```

To see a simple example of the **set** command in a batch file, see article Q286705, "HOW TO: Script Compatibility Layers in Windows XP," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Creating Custom Compatibility Modes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you identify the application compatibility fixes that are needed for your applications, you might find a group of them that are needed for several applications, or you might want to customize a prepackaged compatibility mode by adding or removing application compatibility fixes. In this case, you can create a custom compatibility mode and include it in multiple custom databases.

**To create a custom compatibility mode**

1. In the Compatibility Administrator, select the custom database to which you want to add the custom compatibility mode.
2. On the **Database** menu, click **Create New**, and then click **Compatibility Mode**.
3. Give the custom compatibility mode a descriptive name, and select the application compatibility fixes to include in it.
4. On the **File** menu, click **Save** to save the updates to the custom database.

Include the new custom compatibility mode in other custom databases by using the **Copy** and **Paste** commands on the shortcut menu, which is displayed by right-clicking the compatibility mode name.

# Creating Custom Application Help Messages

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a compatibility problem cannot be resolved with an application compatibility fix, an Application Help message is the last option. Application Help messages can prevent users from running an application that might cause system instability or that might affect other applications adversely. An *Application Help message* either blocks the application from running or displays an informative message before allowing the application to run. For example, an Application Help message might notify the user that the application requires an update and provide the URL to the vendor's Web site.

Microsoft, in conjunction with ISVs, has developed prepackaged Application Help messages for some third-party applications that have compatibility problems. The prepackaged Application Help messages are most commonly used to block the installation of low-level programs — such as antivirus programs and disk access utilities — that were not designed to run on Windows XP Professional or Windows Server 2003. The messages are displayed when a user attempts to install or run one of these applications.

You also can create custom Application Help messages. Because Application Help messages do not resolve compatibility issues, but instead only notify or stop the user, use this type of solution only when you can find no other solution. Whether or not you choose to use an Application Help message that blocks the application from running, provide an informative message that describes possible results or provides instructions about where to find an update.

**To create a custom Application Help message**

1. In the Compatibility Administrator, select the custom database to which you want to add the custom Application Help message.

2. On the **Database** menu, click **Create New**, and then click **Apphelp Message**.

3. Follow the steps in the Create a Custom AppHelp Message Wizard to select the application to test and the type of information to use to match the application to the Application Help message.

4. Follow the steps in the wizard to specify whether you want the application to run or not to run after the message is displayed, and to specify the message text.

5. Click **Test Run** to test the application with the new Application Help message.

6. After resolving any problems, on the **File** menu, click **Save** to save the updates to the custom database.

# Determining the Application Compatibility Fix to Apply

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The tables in this section include descriptions of some compatibility problems that might occur in applications written for earlier versions of Windows, along with the application compatibility fixes that resolve them. As you test applications and encounter compatibility problems, these tables can help you isolate the most appropriate application compatibility fix to apply. If you are still in the planning phase, skip to "Identifying Strategies for Resolving Special Problems" later in this chapter for additional information on resolving problems.

In some cases, you might need to look at multiple tables to find the application compatibility fix that resolves your compatibility problem. For example, you might want to review both color palette issues (Table 3.12) and display issues (Table 3.13) if elements in the application window are not displayed correctly.

Table 3.9 identifies the application compatibility fixes that you can apply for 16-bit applications that have compatibility problems.

**Table 3.9 Application Compatibility Fixes for 16-Bit Application Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| The installation program or application expects to find Microsoft Windows 95 or Microsoft Windows 98 running. | WOWCFEX_PLATFORMVERSIONLIE |
| The installation program or application expects to find Microsoft Windows NT running. | WOWCF2_HACKWINFLAGS |
| The installation program or application expects to find Service Pack 2 installed. | WOWCFEX_SP2CSDVERSIONLIE |
| The application window is hidden by a second window. | ForceSeparateVDM |
| The application becomes unresponsive. | WOWCF_NOWAITFORINPUTIDLE |
| The application behaves unpredictably. | WOWCF_USER_DDENOSYNC |
| A minimized application cannot be restored. | WOWCF2_FIXLUNATRAYRECT |
| Messages are displayed behind the application. | WOWCF2_SETFOREGROUND |
| A message indicates that a file has not been found. | WOWCF2_SYNCSYSFILE |
| A message indicates that the application requires 256 colors. | WOWCFEX_DISPMODE256 |
| A message indicates that the application requires 256 colors. | WOWCFEX_EATDEVMODEMSG |
| A general protection fault (GP fault) occurs. | WOWCFEX_FORCEINCDPMI |
| The application returns a divide-by-zero error. | WOWCF2_DELAYTIMEGETTIME |
| The Help file does not open. | WOWCFEX_USEWINHELP32 |

Table 3.10 identifies the application compatibility fixes that you can apply when the application cannot be removed successfully.

**Table 3.10 Application Compatibility Fixes for Application Removal Problems**

| Symptom | Application Compatibility Fix |
| --- | --- |
| The application returns the "<DirName> refers to a location that is unavailable" error message when an application is removed. | DelayShowGroup |
| During uninstallation, the application returns an error message stating that registry entries cannot be removed. | LUARedirectReg_Cleanup |
| Files are left behind after an application has been uninstalled. | LUARedirectFS_Cleanup |
| Files are left behind after an application has been uninstalled. | MoveWinInitRenameToReg |

Table 3.11 identifies the application compatibility fixes that you can apply when problems occur with the Clipboard when you perform copy and paste operations.

**Table 3.11 Application Compatibility Fixes for Clipboard Problems**

| Symptom | Application Compatibility Fix |
| --- | --- |
| The application does not perform copy or paste operations correctly. | EmptyClipboardtoSet |
| When bitmaps are copied to or from the Clipboard, the application returns the "Access violation" error message. | EmulateClipboardDIBFormat |

Table 3.12 identifies the application compatibility fixes that you can apply when the color palette is distorted or when components of the application window are displayed in the wrong color.

**Table 3.12 Application Compatibility Fixes for Color Palette Problems**

| Symptom | Application Compatibility Fix |
| --- | --- |
| The application window does not display as expected at startup and during execution. | IgnoreHungAppPaint |
| The application changes the system colors. | IgnoreSysColChanges |
| After the application switches tasks or modes, colors change. | PaletteRestore |
| A palette error occurs. | EmulateGetDeviceCaps |
| A palette error occurs when you quit the application, and it returns the "PM_free():Allocation 001B4498 failed integrity check! (InitPalette)" error message. | EmulateGetDeviceCaps |
| The application returns the "Unhandled Error: User Category 10 Fatal Error: Code 2 Palette Error" message. | EmulateDeleteObject |
| The palette is distorted. | ForceDefaultSystemPaletteEntries |
| After the application switches tasks or modes, images appear corrupted. | PaletteRestore |
| After the application switches tasks or modes, the background turns black. | PaletteRestore |
| After the application switches tasks or modes, buttons disappear. | PaletteRestore |

Table 3.13 identifies the application compatibility fixes that you can apply when the application has problems with the screen display, such as incorrect screen resolution and color quality settings.

**Table 3.13 Application Compatibility Fixes for Display Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| The application starts up in a 640 x 480 pixel window and does not switch to a full screen. | Force640x480 |
| The application starts up in a 640 x 480 pixel window and does not switch to a full screen. | Force640x480x16 |
| The application starts up in a 640 x 480 pixel window and does not switch to a full screen. | Force640x480x8 |
| Most of the application window runs off the desktop. | LazyReleaseDC |
| The application returns a message that indicates that the display is corrupted. | Force8BitColor |
| Graphics are corrupted in full-screen mode. | ForceDirectDrawEmulation |
| The application returns a message that indicates that 256 colors are required. | Force8BitColor |
| The application returns a message that indicates that 256-color mode is required. | ForceDisplayMode |
| Introductory videos do not play. | ForceDirectDrawEmulation |
| The taskbar is visible when the application is in full-screen mode. | HideTaskBar |
| When the application switches modes, the screen becomes blank. Pressing ALT+TAB restores the video display. | IgnoreNoModeChange |
| When the application closes, the system display mode changes. | ForceTemporaryModeChange |
| When the application closes, the display mode is not restored. | EmulateUSER |

Table 3.14 identifies the application compatibility fixes that you can apply when the application fails during execution.

**Table 3.14 Application Compatibility Fixes for Execution Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| Access violations occur regularly during application execution. | DisableBoostThread |
| The application shuts down unexpectedly during execution. | HeapIgnoreMoveable |
| The application shuts down unexpectedly during execution. | HeapLookasideFree |
| The application shuts down unexpectedly during execution. | HeapPadAllocation |
| The application shuts down unexpectedly during execution. | HeapValidateFrees |
| The application shuts down unexpectedly when it accesses configuration or user settings. | LUARedirectFS |
| The application periodically shuts down unexpectedly, typically at the same interval at which the screen saver is set to activate. | DisableScreenSaver |

| | |
|---|---|
| Access violations occur during normal functioning as the heap becomes corrupted, causing the application to become unresponsive. | EmulateHeap |
| An access violation occurs when the application switches contexts. | HeapForceGrowable |
| The application returns a "Runtime error 6" message when you install add-ons to the application. | IgnoreVBOverflow |
| The application returns a "Runtime error 6" message when you attempt to use Help. | IgnoreVBOverflow |
| Executable programs designed to run at Windows startup, such as application toolbars, do not run for all users. | LUARedirectReg |
| A default screensaver cannot be applied. | MoveIniToRegistry |

Table 3.15 identifies the application compatibility fixes that you can apply when access violations occur while an application is exiting.

**Table 3.15 Application Compatibility Fixes for Exit Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| Access violations occur when the application is closed. | EmulateHeap |
| Access violations occur when the application is closed. | IgnoreFreeLibrary |
| Access violations occur when the application is closed. | IgnoreLoadLibrary |
| Access violations occur when the application is closed. | IgnoreCRTExit |
| Access violations occur when the application is removed. | IgnoreCRTExit |
| Access violations occur when dialog boxes are canceled. | IgnoreCRTExit |

Table 3.16 identifies the application compatibility fixes that you can apply when the application has file access problems.

**Table 3.16 Application Compatibility Fixes for File Input/Output Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| During installation or operation, the application reports the incorrect amount of available disk space. | EmulateGetDiskFreeSpace |
| When a user with a DBCS (double-byte character set) name starts the application, the application returns an error message indicating that there are problems with the user's temp path. | RedirectDBCSTempPath |
| The application returns the "The system file does not exist" error message. | RedirectWindowsDirToSystem32 |
| When you attempt to open certain Audio Video Interleaved (.avi) files, the application returns the "Access violation" error message. | EmulateWriteFile |
| When you attempt to open a file, the application stops responding. | RemoveOverlappedFlagFromCreateFile |
| When you attempt to open a file, the application says that the file is read-only. | RemoveReadOnlyAttribute |
| When you attempt to open a folder, the application says that the folder is read-only. | RemoveReadOnlyAttribute |

| | |
|---|---|
| When you move files or projects, the application returns the "Access violation" error message. | ChangeAuthenticationLevel |
| The application returns the "File mapping not created" error message. | CorrectCreateEventName |
| The application returns the "File mapping not created" error message. | EmulateCreateFileMapping |
| When accessing a CD, the application returns either the "Unable to locate files" or the "Unable to install" error message. | EmulateCDFS |
| When accessing a CD, the application quits without generating an error message. | RemoveNoBufferingFlagFromCreateFile |
| The application returns the "Files do not exist" error message. | EmulateEnvironmentBlock |
| Files are not found or error messages or paths are truncated. | EmulateGetCommandLine |
| The Uninstall command fails to remove files and folders from program files. You need to use this fix with the CorrectFilePaths compatibility fix (in Table 3.26). | CorrectFilePathsUninstall |

Table 3.17 identifies the application compatibility fixes that you can apply when focus within an application is not maintained correctly.

**Table 3.17 Application Compatibility Fixes for Focus Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| When the application starts, the configuration window is minimized. | ForceKeepFocus |
| The application window becomes inactive. | ForceApplicationFocus |
| A message box becomes inactive. | ForceMessageBoxFocus |

Table 3.18 identifies the application compatibility fixes that you can apply when graphical elements in the application are not displayed correctly.

**Table 3.18 Application Compatibility Fixes for Graphics Control Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| When the application loads, the graphics become corrupted. | EmulateBitmapStride |
| When a minimal installation of the application is performed, graphics are missing and video clips do not play. | ForceCDStop |
| When the application starts, the screen becomes blank. | ForceSimpleWindow |
| When the application starts, the screen becomes blank. | EmulateDirectDrawSync |
| When the application starts, the screen becomes blank. | WinExecRaceConditionFix |
| Either the introductory video does not appear, or the application won’t play in hardware mode. | ForceSimpleWindow |
| Inactive windows do not display as expected. | AlwaysActiveMenus |

| | |
|---|---|
| Clicking on an active bitmap image causes the application to become unresponsive. Bitmap images do not load. | CorrectBitmapHeader |
| Instead of an icon, a shadowed block appears. | CorrectCreateIcon |
| When the application loads bitmaps, they do not display as expected. | CorrectCreateSurface |
| The characters of Asian fonts do not display correctly. | CorrectFarEastFont |
| Pressing ESC returns you to a blank window rather than to the application menu. | DelayWin95VersionLie |
| Menu or list box contents do not appear or they do not appear as expected. | DisableAnimation |
| When you enable the FilterKeys option, the graphics do not display as expected. | DisableFilterKeys |
| The application alternates between the currently applied themed view and the nonthemed view. | DisableThemes |
| When the application runs under new themes, the standard message boxes do not appear. | DisableThemes |
| The system buttons (Minimize and Close) do not work. | DisableThemes |
| The title bar of the application window hides the toolbars. | DisableThemes |
| Dialog boxes are truncated. | DisableThemes |
| Buttons overlap. | DisableThemes |
| When the application window is resized, it does not display as expected. | DisableThemes |
| Text does not appear in the menu or in message boxes. | EmulateDrawText |
| The application fails when it attempts to display the menu or message boxes. | EmulateDrawText |
| Bitmaps in the Autorun dialog boxes do not display as expected. | EmulateGetProfileString |
| The application scrolls text too rapidly. | EmulateGetStringType |
| Typed text or text that should be displayed on the menu buttons or on labels does not appear. | EmulateTextColor |
| The application does not display graphics. | EnlargeGetObjectBufferSize |
| A file cannot be inserted as an icon into another file. Attempting to do so produces an "Out of memory or system resources" error. | ExtractAssociatedIcon |
| An application’s attempts to access menus, specifically the File menu, produces an "Access violation" error. | FailObsoleteShellAPIs |
| The menu and toolbar buttons are the wrong color. | FakeThemeMetrics |
| Typed text or text that is displayed on the menu buttons or in labels is the wrong color. | EmulateTextColor |
| Menu commands are replaced with black boxes. | ForceDefaultSystemPaletteEntries |

| | |
|---|---|
| The application fails when it returns to the main menu. | HeapDelayLocalFree |
| The mouse pointer appears unexpectedly. | HideCursor |
| A button or other control is corrupted. | LoadComctl32Version5 |
| Resizing a window causes the icons and other visual elements to disappear. | NoGdiBatching |
| The application window flickers and is repeatedly refreshed. | NoShadow |
| The user interface is unresponsive. | RemoveInvalidW2KWindowStyles |
| After you quit the application, the system cursor is not restored. | RestoreSystemCursors |
| Images, and the text associated with them, are displayed backwards. | UnMirrorImageList |
| When you insert an image file, the application returns an "Access violation" error message. | IgnoreLoadLibrary |
| When you modify an image, an access violation occurs. | DisableThemes |

Table 3.19 identifies the application compatibility fixes that you can apply when the application fails to install.

**Table 3.19 Application Compatibility Fixes for Installation Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| When you attempt to install the application, the application returns the "ComponentMoveData" error message. | CorrectOpenFileExclusive |
| When you attempt to perform a default installation, text in the Destination Directory dialog box is unreadable. | EmulateGetProfileString |
| When you attempt to install the application, the application returns "The application or DLL <drive>:\<%windir%>\system32\ole2.dll is not a valid Windows image" error message. | CorrectVerInstallFile |
| When you attempt to install an application using an Autorun file, the installation fails. | EmulateCreateProcess |
| Attempts to install the application with InstallShield version 5.0 fail when user names with East Asian characters are used. | IgnoreOemToChar |
| The application appears to become unresponsive during setup because a dialog box containing an error message occupies space beyond the visible desktop. | KeepWindowOnMonitor |
| The application fails to install. | RedirectEXE |
| Setup fails unexpectedly. | EmulateLZHandles |
| On a dual processor machine, an application fails during installation or when you attempt to start it. | SingleProcAffinity |
| The application returns the "<x> not installed correctly, please install again" error message, where x is an application component. | SyncSystemAndSystem32 |

| | |
|---|---|
| The application cannot find a path during installation when a double-byte character set (DBCS) user name is used. | HandleDBCSUserName |
| The application cannot find a path during installation when a DBCS user name is used. | HandleDBCSUserName2 |
| Access violations occur during installation. | FileVersionInfoLie |
| During installation, when the application is searching for the presence of previous versions, an access violation occurs. | LimitFindFile |
| When the application is installed, an access violation occurs. | RemoveBroadcastPostMessage |
| When the application is installed, an access violation occurs. | IgnoreFreeLibrary |
| The application does not install from an Autorun file, causing an access violation. | HideDisplayModes |
| During installation, the application returns the "Cannot open files" error message. | FailOpenFile |
| After you run the setup program, the application returns the "Access violation" error message. | DuplicateHandleFix |
| The application returns the "Setup32: This product does not support Windows NT" error message. | DeRandomizeExeName |
| The application returns the "Unknown button state" error message. | DisableW2KOwnerDrawButtonStates |
| You are notified that a certain DLL is required for installation. | EmulateGetProfileString |
| The application returns the "Must be an administrator to install" error message. | ForceAdminAccess |
| The Configuration menu is blank. After the application's setup program finishes, InstallShield fails. | ForceWorkingDirectoryToEXEPath |
| If you do not register the application, the application does not start. | HeapLookasideFree |
| The application fails to start. | RedirectEXE |
| When the application is started, an access violation occurs. | RemoveBroadcastPostMessage |
| The system fails to restart after you confirm the application's restart request. | EnableRestarts |
| The application fails to register services. | RedirectEXE |
| The Readme.txt file cannot be opened. | HandleRegExpandSzRegistryKeys |

Table 3.20 identifies the application compatibility fixes that you can apply when entering data with the keyboard or using certain key sequences causes problems.

**Table 3.20 Application Compatibility Fixes for Keyboard Input Problems**

| Symptom | Application Compatibility Fix |
|---|---|

| | |
|---|---|
| When you attempt to type in logon information, the application returns the "Unexpected error" message. | AliasDXDC |
| The application does not correctly display the shortcut keys. | DisableKeyboardCues |
| Using the ALT+TAB key combination causes the application window to become blank. | IgnoreAltTab |
| The application becomes unresponsive when you use the ALT+TAB key combination. | IgnoreAltTab |
| The application stops responding when you use the SHIFT key, the CTRL key, or the ALT key 5 times or more in succession. | DisableStickyKeys |
| The keyboard locks up during normal operation. | EmulateUSER |
| The keyboard is sluggish or doesn't function at all, and the system becomes generally unresponsive. | LowerThreadPriority |

Table 3.21 identifies the application compatibility fixes that you can apply when the application has problems with sound, graphics, animation, or video.

**Table 3.21 Application Compatibility Fixes for Multimedia Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| Sounds do not play. | DelayWinMMCallback |
| No sound or limited sound plays. | CorrectACMStreamOpen |
| Sounds do not play because the application does not detect sound devices. | AddWritePermissionsToDeviceFiles |
| The application cannot use the sound hardware, so it returns the "Sound hardware already in use" error message. | AddWritePermissionsToDeviceFiles |
| Sounds do not play with Audio Compression Manager drivers. | CorrectACMArgs |
| The application returns the "Active Movie Control: An object or name was not found" error message. | CorrectActiveMoviePath |
| A message indicates that the sound device driver failed to load. | CorrectSoundDeviceId |
| The sound loops perpetually, and the application does not always respond to shutdown requests. | EmulatePlaySound |
| Audio Video Interleaved (.avi) files continue to run behind the menu screen. | ForceAVIWindow |
| The cursor freezes momentarily during background music changes. | IgnoreMCISTOP |
| The application attempts to send sound to the wrong device (usually to an audio card instead of to the desired USB devices). | WaveOutUsePreferredDevice |

Table 3.22 identifies the application compatibility fix that you can apply when network connections cannot be set up correctly.

**Table 3.22 Application Compatibility Fix for a Network Connection Problem**

| Symptom | Application Compatibility Fix |
|---|---|

| | |
|---|---|
| The application automatically selects Serial Connection when you attempt to select a network connection. | DirectPlayEnumOrder |

Table 3.23 identifies the application compatibility fix that you can apply when permission problems occur.

**Table 3.23 Application Compatibility Fix for a Permissions Problem**

| Symptom | Application Compatibility Fix |
|---|---|
| Insufficient user rights produce errors. | LUATrackFS |

Table 3.24 identifies the application compatibility fixes that you can apply when printing does not work correctly.

**Table 3.24 Application Compatibility Fixes for Printing Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| Access violations occur when the application starts or when it searches for printers. Consequently, the application returns error messages when you attempt to find, select, or print to a printer. | EmulatePrinter |
| The application shuts down unexpectedly when you attempt to print. | ForceAnsiWindowProc |

Table 3.25 identifies the application compatibility fixes that you can apply when invalid registry entries cause problems.

**Table 3.25 Application Compatibility Fixes for Registry Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| The application makes a call to a statically loaded .dll file that has not yet been loaded. | InjectDll |
| When you use a limited user account (LUA), the application returns the error message "Failed to update system registry. Please try using RegEdit." | LUARedirectFS |

Table 3.26 identifies the application compatibility fixes that you can apply when **Start** menu entries or shortcuts are not created correctly or do not work correctly.

**Table 3.26 Application Compatibility Fixes for Start Menu and Shortcut Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| Shortcuts are not created for all users. | CorrectFilePaths |
| Shortcuts are not created for all users. | LUARedirectReg |
| A desktop shortcut is not created. | ForceAnsiGetDisplayNameOf |
| Shortcuts are not removed for all users. | CorrectFilePathsUninstall |
| A valid and an invalid shortcut appear in the Start menu during application installation. | DelayShowGroup |
| The application does not appear on the Start menu. | ProfilesEnvStrings |
| The application does not appear in the Start menu. | ProfilesGetFolderPath |
| The application does not appear in the Start menu. | ProfilesRegQueryValueEx |

| | |
|---|---|
| Shortcuts do not appear on the desktop. | ProfilesEnvStrings |
| Shortcuts do not appear on the desktop. | ProfilesGetFolderPath |
| Shortcuts do not appear on the desktop. | ProfilesRegQueryValueEx |

Table 3.27 identifies the application compatibility fixes that you can apply when problems occur during application startup or during the startup of a process initiated by the application.

**Table 3.27 Application Compatibility Fixes for Startup Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| The setup program fails to install the application. | HandleWvsprintfExceptions |
| The application cannot load. | IgnoreDebugOutput |
| The application does not start after it is installed. | EmulateGetStdHandle |
| The application does not start, but displays error messages or indicates that an access violation has occurred. | StackSwap |
| When starting, the application causes an access violation. | HeapPadAllocation |
| When you attempt to start the application, it returns an error message indicating that the application is not installed. | FailGetStdHandle |
| When you attempt to start the application, it returns an error message indicating that the application must be installed. | EmulateGetStdHandle |
| The application fails to start, and returns an error message. | EmulateVerQueryValue |
| When the application is started directly after setup is complete, it returns the "No Disc Inserted" error message because it does not detect the CD drive. | AddWritePermissionsToDeviceFiles |
| An access violation occurs when the application starts. | EmulateGetDeviceCaps |
| When you start the application, it returns the "Computer memory full" error message. | GlobalMemoryStatusLie |
| When you attempt to start the application, it returns the "Component failed to be located" error message. | VirtualRegistry |
| The application stops responding immediately after starting. | HideDisplayModes |
| The application quits immediately after starting. | HeapClearAllocation |
| The application becomes unresponsive after starting. | AddWritePermissionsToDeviceFiles |
| The application becomes unresponsive or closes unexpectedly shortly after starting. | EmulateFindHandles |
| When the application starts, it returns an exception error. | IgnoreCoCreateInstance |
| When the application starts, it returns the ".dll not compatible with the application" error message. | LoadLibraryCWD |

| | |
|---|---|
| A component or file is not correctly registered or is missing, or the application returns an "Invalid" error message. | CorrectFilePathInSetDlgItemText |
| When the application starts, a window containing only a cursor opens. | IgnoreException |
| The application returns the "Component failed to initialize" error message. | DelayDllInit |
| When you start the application, it automatically minimizes, but sound effects continue to play. The application might or might not run when selected from the taskbar. | NoGhost |
| Runtime errors occur when you start the application or when you start specific components. | SetEnvironmentVariable |
| When the application expects to find a joystick, it quits unexpectedly. | EmulateJoystick |
| When you attempt to start executable programs, such as Scandisk, Disk Defragmenter, or Help, the application displays an error message. | EmulateMissingEXE |
| When the application runs on a high-end machine, its window is blank or it displays an "Access violation" error message. | EmulateSlowCPU |
| When you attempt to start two instances of the application, the application returns an error message. | EmulateToolHelp32 |
| Starting a Safedisk program causes access violations. | FailCloseProfileUserMapping |
| Logging on and starting services take a long time. | FixServiceStartupCircularDependency |
| The application returns a message asking that the correct CD be put in the drive or stating that there is no CD in the drive. | ForceCDStop |
| The application fails to run its startup movie (.avi file) and shuts down. | ForceCoInitialize |
| The application returns the "DirectX does not support the Windows version on the system" error message. | ForceDXSetupSuccess |
| The application does not run in hardware mode. | ForceSimpleWindow |
| The **Autorun** menu does not appear. | CorrectBitmapHeader |
| The **Autorun** dialog box does not get refreshed after the application is installed, so that the option to start the application is still disabled. | HandleAPIExceptions |
| The application becomes unresponsive immediately after starting. | IgnoreException |
| The application displays message boxes that give debugging or other extraneous information. | IgnoreMessageBox |
| The application fails to start or performs slowly. | IgnoreScheduler |
| The application does not start, and it returns the "Desktop is running in 0 bit color mode" error message. | LazyReleaseDC |
| The application does not start, or you get error messages when you attempt to start it. | MapMemoryB0000 |

| | |
|---|---|
| When the application is run by a user other than the user who installed it, it displays error messages such as "Cannot find <X>" (where x is a folder or file that does exist on the system) or "<X> components were not found or are corrupt, run setup." | PopulateDefaultHKCUSettings |
| When you start the application, it returns the "Not a valid Win32 application" or "Unable to run the command" error message. | RecopyExeFromCD |
| An access violation occurs when you attempt to start another program through the application by linking to Windows components (Internet Explorer Web Links or Windows Explorer folders, for example). | RemoveDDEFlagFromShellExecuteEx |
| When the application attempts to start Windows components (such as Microsoft Paint), it returns the "File or program cannot be found" error message. | SearchPathInAppPaths |
| The application returns the "Files cannot be found or do not exist" error message. | SetEnvironmentVariable |
| Entry-point errors occur for functions within earlier versions of Cryptsys and DirectX DLLs. | TerminateExe |
| The application returns the error message "WinG32.dll has been incorrectly installed in the system folder. Please move it to the System32 folder." | WinG32SysToSys32 |

Table 3.28 identifies the application compatibility fixes that you can apply when the telephone dialer does not work correctly.

**Table 3.28 Application Compatibility Fixes for Telephony Application Programming Interface (TAPI) Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| A runtime error occurs in Mfc40.dll when the call is disconnected. | IgnoreTAPIDisconnect |
| The phone dialer does not hang up. | IgnoreTAPIDisconnect |

Table 3.29 identifies the application compatibility fixes that you can apply when the application expects a version of Windows that is not available.

**Table 3.29 Application Compatibility Fixes for Versioning Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| The application requests that Windows 95 be installed. | Win95VersionLie |
| The application requests that Windows 98 be installed. | Win98VersionLie |
| The application requests that Windows NT 4.0 Service Pack 5 be installed. | WinNT4SP5VersionLie |
| The application requests that Windows 2000 be installed. | Win2000VersionLie |
| The application requests that Windows 2000 Service Pack 1 be installed. | Win2000SP1VersionLie |
| The application requests that a previous version of DirectX be installed. | Win2000VersionLie |
| The application requests that a previous version of DirectX be installed. | WinNT4SP5VersionLie |
| The application requests that a previous version of DirectX be installed. | Win2000SP1VersionLie |
| The application reports that it will not function on a Windows NT system. | Win95VersionLie |

| | |
|---|---|
| The application reports that it will not function on a Windows NT system. | Win98VersionLie |

Table 3.30 identifies the application compatibility fixes that you can apply when application windows do not work correctly.

**Table 3.30 Application Compatibility Problems for Window Management Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| An unexpected dialog box appears. | ForceShellLinkResolveNoUI |
| A dialog box does not come to the foreground; you must press ALT+TAB to see it. | GiveupForeground |

Table 3.31 includes miscellaneous application compatibility fixes that do not fit within any other category.

**Table 3.31 Application Compatibility Fixes for Miscellaneous Problems**

| Symptom | Application Compatibility Fix |
|---|---|
| A spawned process fails to operate properly, although the parent process functions normally. For example, Setup.exe works correctly, but the spawned process, Setup2.exe, fails. | PropagateProcessHistory |
| The application does not display system information properly when you use the System Info button in the **About** dialog box. | RemoveIpFromMsInfoCommandLine |

# Identifying Strategies for Resolving Special Problems

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some special problems cannot be resolved by applying application compatibility fixes or by debugging and modifying source code. If application compatibility fixes do not resolve an application problem, try one of the following solutions.

**Reinstall the application by using an Administrator account**
If the problematic application was written for Windows 95, Windows 98, or Windows Millennium Edition, use an Administrator account to install the application. These versions of Windows did not differentiate between Administrator and limited user accounts, so an application that runs on these versions might write to a location that is not allowed by Windows XP Professional.

**Use the current Microsoft Virtual Machine**
Use the most current version of the Microsoft Virtual Machine, which enables Java programs to run on Windows.

**Remove and reinstall the application after upgrading to Windows XP Professional**
If you upgraded the computer to Windows XP Professional, and the application is listed in the Windows Catalog as compatible with Windows XP Professional, try to remove and then reinstall the application after performing the upgrade. If removing and reinstalling the application on Windows XP Professional does not solve the problem, the application might not have been entirely removed. Contact the application vendor for specific instructions on how to remove the application manually. For example, it might be necessary to remove certain folders or registry entries in order to fully reinstall the application.

**Verify MDAC and DirectX versions**
If the application uses Microsoft Data Access Components (MDAC) or DirectX, ensure that you are using the correct versions. Windows XP Professional with Service Pack 1 and Windows Server 2003 support MDAC version 2.7 with Service Pack 1 and DirectX version 9.0. Some older applications attempt to install an earlier version of MDAC or DirectX, but Windows File Protection replaces the earlier version with the newer version. If the application requires the earlier version, errors can occur. In this case, you might need to create a redirection file, as described in the following section, to ensure that the required version is loaded when the application runs.

**Load earlier versions of .dll files**
Some older applications require earlier versions of .dll files than are supplied by Windows XP Professional or Windows Server 2003. If a newer version of a .dll file is loaded when an application runs, problems can occur. To ensure that the application uses the version of the .dll file that it requires, create a redirection file. The redirection file must be created in the same directory as the .dll files and the application.

A redirection file contains no data or information (the contents of the redirection file are ignored), but its name must have the format *ApplicationExecutable.local*. For example, if an application's name is Editor.exe, the redirection file must be named Editor.exe.local.

Even if you do not use a redirection file, always install the .dll files for an application in the same directory as the application. Doing so ensures that the application does not overwrite copies of the .dll files, which might be used by other applications. Doing so also ensures that the application loads the .dll files that it requires.

Placing a redirection file in the directory with an application ensures that all the .dll files in that directory are loaded when the application runs, instead of the .dll files supplied by the operating system. Without a redirection file, the .dll files supplied by the operating system are loaded when the application runs.

### Important

- You do not need to use a redirection file with Windows XP Professional with Service Pack 1 or higher, or with Windows Server 2003. You do need to use a redirection file with Windows XP Professional without Service Pack 1.

**Create application manifest files**
As you migrate applications to Windows XP Professional or Windows Server 2003, you should write an application manifest, a .man file, for each application. The *manifest* is a text file in XML format that lists the dynamic-link libraries that the application should be executed with, along with their version numbers. In the manifest, list the DLL versions with which the application was originally built and tested. When Windows XP Professional or Windows Server 2003 runs an .exe file for which a manifest exists, it loads only the DLL versions specified in the manifest. For more information about specific DLLs, see the DLL Help Database link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Change the short date format**

Some applications that were designed for Windows 95 do not handle short dates correctly when they run on Windows XP Professional. If an application does not accept short dates that are entered by a user, try changing the short date format in **Regional and Language Options**. For example, if the short date format is M/d/yyyy (the default), try changing the format to MM/dd/yyyy.

**To change the short date format on Windows Server 2003**

1. Open **Regional and Language Options** in Control Panel.
2. On the **Regional Options** tab, click **Customize**.
3. On the **Date** tab, next to **Short date format**, select a short date format.

**To change the short date format on Windows XP Professional**

1. Open Control Panel, and then click **Date, Time, Language, and Regional Options**.
2. Click **Regional and Language Options**.
3. On the **Regional Options** tab, click **Customize**.
4. On the **Date** tab, next to **Short date format**, select a short date format.

**Verify that the Kodak Image Edit control is not required**

Ensure that the application does not depend on the Kodak Image Edit control. This control has been removed from Windows XP Professional and Windows Server 2003.

**Verify the length of network printer names**

Some older applications fail when they attempt to print to a network printer with a name longer than 31 characters. To prevent this problem, create a local connection to the network printer for the application.

**To create a local printer connection**

1. Start the Add Printer Wizard.
2. On the Local or Network Printer page, select **Local printer attached to this computer**,and clear **Automatically detect and install my Plug and Play printer**.
3. On the Select a Printer Port page, select the port if it is listed. Otherwise, select **Create a new port,** and then for **Type of Port**, select **Local Port**.
4. Click **Next** to display the **Port Name** window, and enter the port name for the printer in the format *PrintServer**PrinterName*.
5. On the Install Printer Software page, select the printer manufacturer and the appropriate printer.
6. On the Use Existing Driver page, select **Keep existing driver**, if there is one.
7. On the Name Your Printer page, type a name using fewer than 31 characters and finish the wizard as appropriate for your environment.

# Deploying and Distributing Applications and .Sdb Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After adding application compatibility fixes and modes to your custom application compatibility databases (.sdb files), you must deploy or distribute the .sdb files to destination computers in your organization. Usually, you also must deploy or distribute the applications to which the fixes apply. In this context, *deploying* means you are installing .sdb files and applications during an operating system rollout. *Distributing* means that you are installing .sdb files and applications as part of your ongoing operations after rolling out an operating system.

Typically, an automated installation method, such as image-based installation or unattended installation, is used to deploy .sdb files and applications to destination computers during an operating system rollout. By contrast, a software management mechanism, such as Group Policy Software Installation, logon scripts, or Software Update Services (SUS), is typically used to distribute .sdb files and applications to destination computers that already have an operating system installed.

Your deployment team is usually responsible for deploying applications and .sdb files during an operating system rollout. Your operations team is usually responsible for distributing applications and .sdb files on destination computers that already have operating systems installed.

Figure 3.11 shows where the deployment and distribution tasks fit into the overall application compatibility planning, testing, and deployment process.

**Figure 3.11 Deploying and Distributing Applications and .Sdb Files**




This section describes how to use various tools to deploy and distribute applications and .sdb files. This section does not help you choose a deployment or distribution tool. That choice depends on several factors, including your deployment design, your IT standards and practices, and your operational procedures. It is assumed that you have already chosen a deployment or distribution tool based on these factors. For more information about choosing an automated deployment tool, see "Choosing an Automated Installation Method" in *Automating and Customizing Installations* of this kit. For more information about choosing an application distribution tool, see the "Application Deployment Using Microsoft Management Technologies" white paper, available through the Application Deployment link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Reviewing Application Deployment and Distribution Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft provides several tools to help you deploy or distribute applications and .sdb files. These tools include:

- Deployment tools, such as the System Preparation (Sysprep) tool and Remote Installation Services (RIS). You use these automated installation tools to deploy applications and .sdb files during an operating system rollout.

- Distribution tools, such as SUS, Group Policy Software Installation, logon scripts, and SMS. You use these tools to distribute applications and .sdb files to computers that already have an operating system installed.

In addition, Microsoft provides two tools that you use in conjunction with the deployment tools and the distribution tools to package and install applications and .sdb files. These tools are:

- Application Compatibility Database Installer (Sdbinst.exe), which you use to install and register .sdb files.

- Windows Installer, which consists of two tools: a software packaging tool that you use to package application installation files and .sdb files, and an installer tool that you use to install the packages.

This section provides you with an overview of these tools.

**Deployment Tools**

There are three primary methods of deploying applications and .sdb files during an operating system rollout: unattended installation, image-based installation, and RIS installation. Each method requires one or more tools.

**Unattended installation** Unattended installation is a method of automating the Windows Setup program so that no user input is required when you install an operating system or applications. To perform unattended installations, you need to create an answer file that contains answers to the questions that Windows Setup normally prompts for during an installation. By using an answer file to control Windows Setup, you can automate the entire installation process, including the installation of applications and .sdb files. Unattended installation is the most flexible method of deploying applications and .sdb files: It can be used to perform operating system upgrades or clean installations on computers with different hardware and software configurations.

**Image-based installation** Image-based installation is a method of copying, also known as *cloning*, preconfigured operating systems and software applications onto clients and servers. To perform image-based installations, you need the Sysprep tool, which can be found on the Windows XP Professional or Windows Server 2003 operating system CD, and you need a third-party disk-imaging tool. Image-based installation is the fastest method of installing operating systems and applications, but you cannot use it to perform operating system upgrades.

**RIS installation** RIS installation is a method of cleanly installing an operating system and applications across a network with minimal end-user interaction. RIS uses Pre-Boot eXecution Environment (PXE) technology to enable client computers without an operating system to start up and connect remotely over a TCP/IP network connection to a RIS server, which installs a supported operating system. RIS installations require substantial planning and design, but are an efficient deployment method if you need to install clean operating systems and preconfigured applications onto a large number of clients.

To obtain the unattended and image-based installation tools, open the Support\Tools folder on any Windows XP Professional or Windows Server 2003 operating system CD, and then open Deploy.cab. The Readme.txt file in Deploy.cab explains how to install and use the tools. The RIS installation tools are built into Windows Server 2003. For more information about unattended installations, image-based installations, and RIS installations, see "Designing Unattended Installations," "Designing Image-based Installations with Sysprep," and "Designing RIS Installations," respectively, in *Automating and Customizing Installations* of this kit.

**Distribution Tools**

The four most commonly used methods of distributing applications and .sdb files are the use of Group Policy Software Installation, logon scripts, the Software Update Services (SUS), and SMS.

**Group Policy Software Installation** You can use Group Policy Software Installation to centrally install and manage software applications throughout an organization. Using this management system, you can distribute applications and .sdb files to

groups of users and computers. It is the most common method of distributing applications and .sdb files to computers that already have an operating system installed. Active Directory is required in order to use Group Policy Software Installation.

**Logon scripts** Logon scripts run when a user logs on to a computer. Logon scripts typically are written as command (.cmd) files, batch (.bat) files, or Microsoft® Visual Basic® Scripting Edition (.vbs) files.

**Software Update Services (SUS)** You can use SUS to download updates for Windows XP Professional and Windows Server 2003. The updates can include updated versions of the prepackaged application compatibility databases and updated versions of the corresponding code segments that run when an application needs an application compatibility fix, application compatibility mode, or Application Help message. SUS provides dynamic notification of updates to computers, and it provides a simple, automatic solution for distributing updates to networked clients and servers.

**SMS** You can use SMS to distribute applications and .sdb files. SMS includes features such as software metering and software distribution-point population and maintenance. SMS does not require Active Directory.

**Note**

- You also can use a third-party software management program to distribute applications and .sdb files.

Group Policy Software Installation and SUS are built into Windows Server 2003. For more information about Group Policy, see "Group Policy" in Help and Support Center for Windows Server 2003 and "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit. For more information about SUS, see the "Software Update Services Deployment" white paper, available through the Software Update Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, and the Deploying Software Update Services White Paper on the Microsoft Web site.

SMS is a Microsoft product available for purchase. For more information about SMS, see the SMS Product Information link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Application Compatibility Database Installer**

The Application Compatibility Database Installer (Sdbinst.exe) installs and registers your .sdb files on destination computers. Sdbinst.exe is a command-line program that is installed by default into the *systemroot*\System32 folder on computers running Windows XP Professional or Windows Server 2003. Although you can use Compatibility Administrator to install and register .sdb files, using Sdbinst.exe is the only way to automate this task on a large number of computers.

To use Sdbinst.exe, you first copy your .sdb files to the destination computers, and then run Sdbinst.exe on each destination computer. Usually, you copy your .sdb files to a subfolder in the *systemroot*\AppPatch folder (for example, *systemroot*\AppPatch\Custom) on each of the destination computers. This keeps all of the .sdb files in one place, because the prepackaged databases are stored in the *systemroot*\AppPatch folder on Windows XP Professional and Windows Server 2003. If you prefer, you can copy your application compatibility database files to a different folder, but if you do this you must set permissions on the folder granting the Everyone group Read & Execute access and the Administrators group Full Control access. This ensures that your application fixes get implemented regardless of who runs an application. It also ensures that only those users with Administrator credentials can register or alter your application compatibility database files.

Sdbinst.exe requires the following syntax:

```
          sdbinst [-?] [-q] [-u] [-g] [-n] ApplicationCompatibilityDatabase.sdb |
{GUID} | "InternalApplicationName"
```

The parameters are defined in Table 3.32.

**Table 3.32 Command-Line Parameters for Sdbinst.exe**

| Parameter | Description |
|---|---|
| **-q** | Runs Sdbinst.exe in quiet mode, which suppresses message boxes during the database installation process. Use this parameter if you are using an automated installation method to install application compatibility databases. |
| **-u** | Uninstalls an application compatibility database. To identify the application compatibility database that is to be uninstalled, you must specify a file name, a globally unique identifier (GUID), or an internal name. |

| | |
|---|---|
| **-g** | Specifies the GUID of the application compatibility database that you want to uninstall. Use only with the **-u** parameter. You must specify a GUID if you use this parameter. |
| **-n** | Specifies the internal name of the application compatibility database to be uninstalled. This is the name assigned internally to the database when it was created in Compatibility Administrator. Use only with the **-u** parameter. You must specify the internal name of the database if you use this parameter. |
| *ApplicationCompatibilityDatabase* | Specifies the name of the .sdb file that you want to install or uninstall. |
| *GUID* | Specifies the globally unique identifier of the application compatibility database that you want to uninstall. Use only with the **-g** parameter. |
| *InternalApplicationName* | Specifies the name of the application compatibility database that you want to uninstall. This is the name assigned internally to the database when it was created in Compatibility Administrator. Use only with the **-n** parameter. |
| **-?** | Displays the Help documentation for Sdbinst.exe. |

**Windows Installer**

Windows Installer consists of several tools and features, two of which are useful for deploying and distributing applications and .sdb files: Windows Installer packages (.msi files), and the Windows Installer tool (Msiexec.exe). As a stand-alone tool, Msiexec.exe is not a solitary solution for deploying and distributing applications and .sdb files. You use Msiexec.exe in conjunction with your deployment or distribution tools. Likewise, you use .msi files in conjunction with your deployment or distribution tools.

**Windows Installer packages**

A Windows Installer package (.msi file) contains all of the information that the Windows Installer service requires to run an application's setup program and to install and configure the application. In addition, a Windows Installer package can contain any files that are required by the application's setup program, such as dynamic-link library (.dll) files, information (.inf) files, executable (.exe) files, or cabinet (.cab) files. Packaging your applications and your .sdb files into Windows Installer packages makes application deployment and distribution easier and more efficient, and prevents you from having to manage numerous installation files for each of your applications. Application developers usually create Windows Installer packages at the end of the software development process, but you can create a Windows Installer package after an application has been developed.

You can use either Microsoft® Visual Studio® .NET development system or a third-party software program to create Windows Installer packages. For more information about creating Windows Installer packages, see the Windows Installer Software Developer's Kit (SDK) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Windows Installer tool (Msiexec.exe)**

Msiexec.exe carries out the instructions that are contained in a Windows Installer package. Msiexec.exe performs all installation-related tasks, including copying files onto the hard disk, modifying the registry, and creating shortcuts on the desktop. Msiexec.exe is a command-line tool that is built into Windows XP Professional and Windows Server 2003.

For more information about Msiexec.exe command-line syntax and usage, see "Msiexec (command-line options)" in Help and Support Center for Windows Server 2003.

# Deploying Applications and .Sdb Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use any of the following automated installation methods to deploy applications and .sdb files during an operating system rollout: unattended installation, image-based installation with the Sysprep tool, and RIS installation. The method that you use depends on several factors, such as the operating system that you are installing, your organization's network and hardware configuration, and the applications that you are deploying. With a few exceptions, no single factor determines which installation method you should use.

This section does not provide guidance for choosing an automated installation method. It is assumed that you have evaluated all of the factors that affect your choice of an automated installation method already, and that you have determined which automated installation method best suits your needs. It is also assumed that you understand how to use the tools associated with each of these automated installation methods.

**Note**

- For more information about unattended installation and image-based installation with Sysprep, see the *Microsoft® Windows® Server 2003 Corporate Deployment Tools User's Guide* (Deploy.chm and Ref.chm). Deploy.chm and Ref.chm are included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD. Deploy.chm and Ref.chm must be stored in the same folder. You can open these files on the operating system CD, or open them on your computer by installing Deploy.cab in a folder on your computer.

# Deploying Applications and .Sdb Files During Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To deploy applications and .sdb files during an unattended installation, you need to configure the following items.

**Your distribution share** Add your applications and .sdb files to your distribution share so that they can be installed on your destination computers.

**Your Cmdlines.txt file** Add the appropriate commands to your Cmdlines.txt file so that Sdbinst.exe runs and installs your .sdb files after an operating system is installed on a destination computer, but before the destination computer restarts and a user logs on.

**The [GuiRunOnce] section of your answer file** Add the appropriate commands to the [GuiRunOnce] section of your answer file so that your applications are installed on each of your destination computers after the computer restarts and a user logs on.

Your unattended installation must adhere to the following guidelines:

- You cannot perform an unattended installation from a Windows XP Professional or Windows Server 2003 operating system CD. You must use a distribution share.
- You cannot name your answer file Winnt.sif. You must name it Unattend.txt or another name. You can use the name Winnt.sif only if you are starting your destination computers from a Windows XP Professional or Windows Server 2003 operating system CD, which you cannot do if you are deploying applications or .sdb files.
- You cannot install applications or .sdb files during an unattended upgrade. This is because you cannot use a distribution share during an unattended upgrade, and you need to use a distribution share to install applications and .sdb files.

#### Configuring Your Distribution Share

You need to add the following files to your distribution share to deploy applications and .sdb files during an unattended installation.

**.Sdb files** Add your .sdb files to your distribution share so that the .sdb files are copied to every destination computer. Although you can save .sdb files in any folder on a destination computer, the most commonly used location is a folder named *systemroot*\AppPatch\Custom. To copy your .sdb files to the *systemroot*\AppPatch\Custom folder on each of your destination computers, save your .sdb files in the $OEM$\$$\AppPatch\Custom folder of your distribution share. To save your .sdb files to another folder on your destination computers, save your .sdb files in a corresponding folder in the $OEM$ folder of your distribution share. For example, if you want to save your .sdb files in *systemdrive*\Myapps\Sdbfiles, save your .sdb files in the $OEM$\$1\Myapps\Sdbfiles folder of your distribution share.

**Installation files for applications** Add the installation files for your applications to your distribution share. This ensures that the installation files are copied to every destination computer, and that the files are available during your unattended installation. You can save installation files, such as .msi files, in any folder that you choose on a destination computer. For example, if you are installing an application called Myapp, you might save the installation files for Myapp in *systemdrive*\Apps\Myapp. To create this folder structure on every destination computer, add $OEM$\$1\Apps\Myapp to your distribution share.

For more information about structuring your distribution share, see "Designing Unattended Installations" in *Automating and Customizing Installations* of this kit.

#### Configuring Your Cmdlines.txt File

To install .sdb files during an unattended installation, you need to configure your Cmdlines.txt file. If you are not already using a Cmdlines.txt file to perform other installation and configuration tasks during an unattended installation, you need to create a Cmdlines.txt file.

**To create a Cmdlines.txt file**

1. Use Notepad, or any text editor, to create a text file named Cmdlines.txt.

2. Add the following section to the Cmdlines.txt file:

```
[Commands]
```

3. Save the Cmdlines.txt file in the $OEM$ folder of your distribution share.

4. Add the following section, entry, and value to your unattended installation answer file:

```
[Unattended]
OemPreinstall = Yes
```

**To configure a Cmdlines.txt file to install .sdb files**

- Add the **Sdbinst.exe** command, with the appropriate parameters, to the [Commands] section of your Cmdlines.txt file. The command and associated parameters must be enclosed in quotation marks.

Be sure to use the **/c** parameter with the Cmd.exe command. This terminates Cmd.exe when Sdbinst.exe has finished running. Also, be sure to use the **-q** parameter with the **Sdbinst.exe** command, which instructs Sdbinst.exe to run in quiet mode. For example, to install *systemroot*\AppPatch\Myapp.sdb, create the following entry in your Cmdlines.txt file, below the [Commands] section heading:

```
" cmd.exe /c sdbinst -q %systemroot%\AppPatch\Myapp.sdb"
```

**Configuring the [GuiRunOnce] Section in Your Answer File**

To install applications on a destination computer during an unattended installation, you need to add commands to the [GuiRunOnce] section in your answer file. The commands that you use depend on how you are installing the application. If you are installing an application that is packaged into an .msi file, use the **Msiexec** command to start the Windows Installer tool. If you are installing an application that is not packaged into an .msi file, use whatever command is necessary to run the application's setup program (usually, the command is **Setup**).

For example, to install an application that is packaged into an .msi file, add the following command to the [GuiRunOnce] section of your answer file:

```
" cmd.exe /c msiexec /qn /i Myapp.msi"
```

The **/c** parameter instructs the command console program to close when Msiexec.exe stops running. The **/i** parameter instructs the Windows Installer program to install the .msi package, and the **/qn** parameter instructs the Windows Installer program to run without displaying a user interface.

Likewise, to install an application that is not packaged into an .msi file and that has its own setup program named Setup.exe, add the following command to the [GuiRunOnce] section of your answer file:

```
" cmd.exe /c Setup.exe"
```

Regardless of whether you install applications with the Windows Installer program or with an application's setup program, you must adhere to the following guidelines in setting up an unattended installation.

**Installation programs cannot require a restart** You can automate installation tasks only if you can prevent the installation program from restarting the computer. When a computer restarts, all remaining commands in the [GuiRunOnce] section are lost. If the system restarts before running the commands listed in the [GuiRunOnce] section, the remaining commands will not run. Therefore, you need to suppress restarts. If you cannot suppress a restart within the installation program, you can try to repackage the application into a Windows Installer package by using a third-party application. For more information about Windows Installer packaging, see the Windows Installer Software Developer's Kit (SDK) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Installation programs cannot rely on the Windows Explorer shell** The Windows Explorer shell is not loaded when the operating system starts running commands, programs, scripts, or batch files that are listed in the [GuiRunOnce] section of an answer file. In some cases, the application vendor can give you an updated installer program that does not rely on the Windows Explorer shell. If the vendor cannot provide this, you might be able to repackage the application as a Windows Installer package (.msi file).

**Installation programs might need to be run from a batch file** Installation programs often start and stop several different

processes. In some cases when you are installing multiple applications, this can inadvertently cause the next command listed in the [GuiRunOnce] section to start before the previous command finishes running. When this occurs, the second installation program usually fails. To prevent this, you can run the installation programs from batch files by using the **Start** command with the **/wait** parameter. This forces each installation program to run until it is completed before the next command listed in [GuiRunOnce] runs. For more information about using the **start** command to run installation programs in batch files, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Command-line reference A-Z.

# Deploying Applications and .Sdb Files During Image-based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Three primary methods are available for deploying applications and .sdb files during an image-based installation:

- You can install and configure your applications and .sdb files on the master installation and then create a disk image of the master installation. The *master installation* is the reference computer that contains the operating system, software applications, and configuration settings that you plan to install on destination computers in your organization. When you copy the disk image to a destination computer, and then start the destination computer, the applications and .sdb files will be installed, configured, and ready for use by an end user on the destination computer.

- You can install applications and .sdb files by using a Sysprep.inf file to automate Mini-Setup. Mini-Setup is a subset of Windows Setup that gathers and configures user-specific information when a user starts a destination computer for the first time. This method uses a Cmdlines.txt file and the [GuiRunOnce] section of the Sysprep.inf file.

- You can use the Factory mode feature of Sysprep, and a Winbom.ini file, to attach or detach applications and run Sdbinst.exe on each destination computer. When you attach or detach applications, your application files and .sdb files exist on your disk image, but the applications and .sdb files are not installed. Then, when you copy the disk image to a destination computer and start the destination computer, you can use a Winbom.ini file to attach (install) or detach (remove) your application files and .sdb files.

The first and second methods are used most commonly in corporate environments in which a single standard software configuration is deployed. The Factory mode method is most often used in factory environments or corporate environments where several different software configurations are being deployed.

**Installing and Configuring Applications and .Sdb Files on a Disk Image**
To deploy installed and configured applications and .sdb files on a disk image:

1. Install and configure your applications on your master installation.

2. Copy your .sdb files to your master installation. Usually, .sdb files are saved in the *systemroot*\AppPatch\Custom folder, but you can save the files in any folder.

3. Run Sdbinst.exe on your master installation to install and register your .sdb files.

4. Run the Sysprep tool on your master installation.

5. Run your disk imaging program, and save the disk image in a shared folder or on some type of removable media.

6. Use your disk imaging tool to copy the disk image to your destination computers.

This is only a summary of the tasks that you need to perform to deploy installed and configured applications and .sdb files on a disk image. For more information about designing and performing an image-based installation, see "Designing Image-based Installations with Sysprep" in *Automating and Customizing Installations* of this kit.

**Installing and Configuring Applications and .Sdb Files During Mini-Setup**
To install and configure applications and .sdb files during Mini-Setup:

1. Copy your application files to the Sysprep folder on your master installation. Do not install the applications.

2. Copy your .sdb files to your master installation. Usually, .sdb files are saved in the *systemroot*\AppPatch\Custom folder, but you can save the files in any folder.

3. Create and configure your Cmdlines.txt file as described in "Deploying Applications and .Sdb Files During Unattended Installations" earlier in this chapter. Save the file in the Sysprep folder on your master installation.

4. Specify the path to your Cmdlines.txt file by configuring the InstallFilesPath entry in the [Unattended] section in your Sysprep.inf file.

5. Create and configure the [GuiRunOnce] section of your Sysprep.inf file in the same way you created and configured the [GuiRunOnce] section of your unattended installation answer file. For more information about configuring the [GuiRunOnce] section of Sysprep.inf, see "Deploying Applications and .Sdb Files During Unattended Installations" earlier in this chapter.

6. Configure any other sections and entries in your Sysprep.inf file, and save the file in the same location as Sysprep.exe and Setupcl.exe, which are usually in the *systemdrive*\Sysprep folder on your master installation.

7. Run the Sysprep tool on your master installation.

8. Run your disk imaging program, and save your disk image in a shared folder or on some type of removable media.

9. Use your disk imaging tool to copy the disk image to a destination computer.

This is only a summary of the tasks that you need to perform to deploy applications and .sdb files during Mini-Setup. For more information about designing and performing an image-based installation, see "Designing Image-based Installations with Sysprep" in *Automating and Customizing Installations* of this kit.

**Installing and Configuring Applications and .Sdb Files in Factory Mode**
To install and configure applications and .sdb files by using the Factory mode feature of Sysprep:

1. Copy your application files to your master installation. Do not install the applications.

2. Copy your .sdb files to your master installation. Usually, .sdb files are saved in the *systemroot*\AppPatch\Custom folder, but you can save the files in any folder.

3. Create a Winbom.ini file by using Notepad or any other text editor, and save the Winbom.ini file on a floppy disk, on a CD, or in the same folder as the Factory.exe program (usually, the *systemdrive*\Sysprep folder).

4. Configure the [OEMRunOnce] section of the Winbom.ini file. This section contains the list of commands to be run while the computer is running in Factory mode. For example, you can attach (install) or detach (remove) applications, and run Sdbinst.exe.

5. Configure the WinbomType entry in the [Factory] section of the Winbom.ini file. This entry specifies how the Winbom.ini file is processed; it is required.

6. Run the Sysprep tool, with the **/factory** parameter, on your master installation.

7. Run your disk imaging program, and save your disk image in a shared folder or on some type of removable media.

8. Use your disk imaging tool to copy the disk image to a destination computer.

9. Start the destination computer.

10. After the computer processes your Winbom.ini file and installs your applications and .sdb files, run Sysprep with the **/reseal** parameter.

This is only a summary of the tasks that you need to perform to deploy applications and .sdb files by using the Factory mode feature of Sysprep. Although Factory mode lets you install and configure applications quickly and efficiently, using it requires substantial planning and design. For more information about planning, designing, and implementing image-based installations with Factory mode, see "Designing Image-based Installations with Sysprep" in *Automating and Customizing Installations* of

this kit.

# Deploying Applications and .Sdb Files During RIS Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Two methods are available for using a Remote Installation Services (RIS) server to deploy applications and .sdb files: the Risetup method, which uses Remote Installation Services setup (Risetup.exe), and the Riprep method, which uses Remote Installation Preparation Wizard (Riprep.exe). With either method, you can deploy an operating system, applications, and .sdb files across your network to destination computers throughout your organization.

**Risetup** In the Risetup method, a distribution share is used to deploy an operating system, applications, and .sdb files. The distribution share is similar in structure to the distribution share that is used during an unattended installation. You can add applications and .sdb files to the distribution share, and use answer files to install the applications and .sdb files.

**Riprep** In the Riprep method, disk images are used to deploy an operating system, applications, and .sdb files. The disk images are similar to the disk images that are used during an image-based installation with Sysprep, but instead of using Sysprep to prepare the master installation, you use Riprep. Riprep converts the master installation into a remote installation image that you can install on multiple destination computers.

Using the Riprep tool is the most common way to use RIS to deploy an operating system, applications, and .sdb files. It is faster than the Risetup method, and it requires less configuration after an image is copied to a destination computer.

**Deploying Applications and .Sdb Files Using Riprep**

To deploy applications and .sdb files using Riprep:

1. Install and configure the operating system and applications on your master installation.

2. Copy your .sdb files to your master installation. Usually, .sdb files are saved in the *systemroot*\AppPatch\Custom folder, but you can save the files in any folder.

3. Run Sdbinst.exe on your master installation to install and register your .sdb files.

4. Test the master installation to verify that the applications function properly and that the configuration is correct.

5. Run Remote Installation Services setup (Riprep.exe) on your master installation.

   Riprep converts the master installation into a remote installation image that you can install on multiple destination computers. It also replicates the image to a RIS server, where the image is available for installation on remote-boot-enabled client computers.

   Riprep configures various operating system settings on the master installation to ensure that every copy of the master installation's disk image is unique when you install it on destination computers. For example, Riprep resets the security identifiers (SIDs). Riprep also configures the master installation so that every destination computer starts in Mini-Setup mode.

6. Use RIS to copy the disk image to your destination computers.

   Before a RIS server can deploy an operating system image to a client computer, the client computer must initiate a remote network boot. When a client computer initiates a remote network boot, it is assigned a RIS server, which downloads the disk image to the client computer.

   After the disk image is downloaded to the client computer and the client computer is restarted, Mini-Setup starts. Mini-Setup then uses Plug and Play to detect hardware differences between the master installation and destination computers. This enables you to use a single master installation to deploy an operating system, applications, and .sdb files to destination computers that have different hardware configurations.

This is only a summary of the tasks that are required in order to deploy installed and configured applications and .sdb files using RIS. For more information about designing and performing a RIS installation, see "Designing RIS Installations" in *Automating and Customizing Installations* of this kit.

# Distributing Applications and .Sdb Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To distribute applications and .sdb files to computers that already have an operating system installed, you can use any of the following tools: the Group Policy Software Installation tool, logon scripts, SUS, and SMS. The tool that you use depends on several factors, such as the operating systems that you are running, your Active Directory configuration, and the design of your managed environment. With a few exceptions, no single factor determines which distribution tool you should use. This section does not provide guidance for choosing a distribution tool. It is assumed that you have already evaluated all of the factors that affect your choice of an application distribution tool, and that you have determined which tool best suits your needs.

**Tip**

- You can also distribute applications and .sdb files by writing stand-alone scripts that can be run any time (as opposed to logon scripts), and then running the scripts remotely with Windows Script Host. For information about Windows Script Host, see the Windows Script Host link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Distributing Applications and .Sdb Files Using Group Policy Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can distribute applications and .sdb files to remote computers in managed environments by using Group Policy Software Installation to assign Windows Installer packages to computers. This is the most common method of distributing applications and .sdb files in managed environments. Windows Installer packages are deployed and managed within a Group Policy object, which is, in turn, associated with an Active Directory container — either a site, a domain, or an organizational unit. Therefore, you need either to create a new Group Policy object and link it to one of these Active Directory containers or to edit an existing Group Policy object and link it to a site, domain, or organizational unit.

**Important**

- To complete this procedure, you must log on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

**To deploy software to remote computers**

1. Open the Group Policy Object Editor.
2. Do one of the following:
   - To assign or publish software applications to computers, in the console tree, double-click **Computer Configuration**.
   - To assign or publish software applications to users, in the console tree, double-click **User Configuration**.
3. Double-click **Software Settings**, and then click **Software Installation**.
4. Right-click **Software Installation**, click **New**, and then click **Package**.
5. Click the Windows Installer package that you want to assign, and then click **Open**.
6. In **Deploy Software,** click **Assigned**.

For information about using Group Policy, see "Group Policy" in Help and Support Center for Windows Server 2003.

# Distributing Applications and .Sdb Files Using Logon Scripts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

*Logon scripts* are batch files, programs, or scripts that can be assigned to user accounts. Logon scripts run automatically when a user logs on. A logon script can be used to configure a user's working environment when the user logs on. It allows an administrator to influence a user's environment without managing all aspects of it. You also can use a logon script to install applications, application updates, and .sdb files.

A logon script can be assigned to a user account or to a security group in two ways: You can assign a logon script to run on a stand-alone computer by using Group Policy, or you can assign a logon script to run on a group of computers in a site, organizational unit, or domain by using Active Directory.

If you use Group Policy to assign logon scripts on a stand-alone computer, the scripts are saved locally in the *systemroot*\System32\GroupPolicy\Machine\Scripts\Logon folder on each stand-alone computer. If you use Active Directory to assign logon scripts to a group of computers, the scripts are saved in the *systemroot*\SYSVOL\sysvol*DomainName*\Scripts folder on your domain controllers. This folder is shared using the share name NETLOGON. Computers running Windows XP Professional or Windows Server 2003 look for logon scripts in this folder on the authenticating domain controller when a user logs on.

To assign a logon script to a user or group, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

**To assign a logon script to a local user account or security group**

1. In Control Panel, double-click **Administrative Tools**, and then double-click **Computer Management**.
2. In the console tree, double-click **Local Users and Groups**.
3. In the console tree, click **Users** or **Groups**.
4. Double-click the name of the user or group to whom you want to assign a logon script.
5. Click the **Profile** tab.
6. In the **Logon script** text box, enter the path and name of the logon script that you want to assign to the user or group, and then click **OK**.

**To assign a logon script to a user in Active Directory**

1. In Control Panel, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
2. In the console tree, click **Users**.
3. Double-click the name of the user to whom you want to assign the logon script.
4. Click the **Profile** tab.
5. In the **Logon script** text box, type the path and name of the logon script that you want to assign to the user, and then click **OK**.

For more information about configuring logon scripts, see "Logon Scripts" in Help and Support Center for Windows Server 2003.

# Distributing Patches and .Sdb Files Using SUS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

SUS replaces Corporate Windows Update. It provides a means for medium-sized enterprises to manage and distribute critical Windows patches. These patches resolve security vulnerabilities and stability problems with the Windows XP Professional and Windows Server 2003 operating systems. The patches also can include updates to the prepackaged databases that are installed in the *systemroot*\AppPatch folder. Although SUS is well suited to an Active Directory environment, Active Directory is not required in order to use SUS.

SUS provides dynamic notification and automatic distribution of critical updates to computers that are running Windows XP Professional or Windows Server 2003. In addition, only one of your computers needs to be connected to the Internet for SUS to function.

An SUS solution has three components: an SUS server, an Automated Updates client, and an Intranet-hosted Windows Update server.

**An SUS server** The server component of an SUS solution is installed inside the corporate firewall. With SUS installed, the internal server can synchronize with the Windows Update Web site whenever critical updates are available. The synchronization can be automatic, or the administrator can perform it manually. After the updates are downloaded to the server that is running SUS, the administrator can decide which updates to publish. All clients that are configured to use the server that is running SUS can then install the updates.

**An Automatic Updates client** The client component of a SUS solution is installed on client computers running Windows XP Professional or Windows Server 2003 so that they can connect to the internal server that is running SUS. The administrator can control which server each client computer connects to, and schedule when the client performs all installations of critical updates.

**An intranet-hosted Windows Update server** This server acts as the virtual Windows Update server for client computers. It contains the synchronization service and administrative tools for managing updates. It services requests for approved updates by the client computers connected to it using the Hypertext Transfer Protocol (HTTP). This server can also host critical updates downloaded from the synchronization service and refer client computers to those updates.

For more information about configuring and deploying SUS, see the "Software Update Services Deployment" white paper, available through the Software Update Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, and the Deploying Software Update Services White Paper on the Microsoft Web site.

# Distributing Applications and .Sdb Files Using SMS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

SMS includes a software distribution feature that can automate the distribution of applications and .sdb files. By using the software distribution feature, you can centrally manage and monitor the installation of applications and .sdb files on remote computers. SMS is a Microsoft product available for purchase. For more information about SMS, see the SMS Product Information link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Additional Resources for Planning and Testing for Application Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this chapter.

**Related Information**

- "Designing a Test Environment" in this book for more information about testing.
- "Designing a Pilot Project" in this book for more information about piloting.
- "Choosing an Automated Installation Method" in *Automating and Customizing Installations* of this kit for information about choosing an automated deployment tool.
- "Designing Image-based Installations with Sysprep" in *Automating and Customizing Installations* of this kit for information about designing and performing an image-based installation.
- "Designing RIS Installations" in *Automating and Customizing Installations* of this kit for more information about designing and performing a RIS installation.
- "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit for information about Group Policy Software Installation.
- "Hosting Applications with Terminal Server" in *Planning Server Deployments* of this kit for information about application compatibility with Windows Server 2003 Terminal Server.
- The Deploying Software Update Services White Paper on the Microsoft Web site for information about SUS.
- The Server Roles: The Application Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about server application compatibility.
- The Porting 32-bit Applications link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about porting 32-bit applications to 64-bit operating systems.
- The Windows Application Compatibility link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources to download the Application Compatibility Toolkit.
- The Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for the Windows Catalog.
- The Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about the Windows Server Catalog.
- The Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about assessing project constraints and risks.
- The Designed for Windows Logo Program link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- Article number 326690 "Anonymous LDAP Operations to Active Directory Are Disabled on Windows Server 2003 Domain Controllers," available through the Microsoft Knowledge Base link on the Web Resources page at

http://www.microsoft.com/windows/reskits/webresources for information about anonymous LDAP requests.

- Article number 314649 "Windows Server ADPREP Command Causes Mangled Attributes in Windows 2000 Forests That Contain Exchange 2000 Servers," and article number 328661 "XADM: Running Exchange 2000 Setup with /Forestprep Switch Produces Error 0XC1037AE6," available through the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- "Migrating IIS Web Sites to IIS 6.0" in *Deploying Internet Information Services (IIS) 6.0* of this kit (or see "Migrating IIS Web Sites to IIS 6.0" on the Web at http://www.microsoft.com/reskit).

- The Debugging Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources to download the most recent version of Debugging Tools for Windows and symbol files.

- The SMS Product Information link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about SMS.

- The Windows Installer Software Developer's Kit (SDK) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about creating Windows Installer packages.

- The Windows Script Host link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about Windows Script Host.

- The following white papers from the Application Compatibility Toolkit, available through the Windows Application Compatibility link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources:

    - "Common Application Compatibility Issues," for detailed information about common compatibility problems and examples of source code that you can use to avoid them.

    - The "Windows Compatibility Checklist," to help you identify test cases that verify an application's interaction with the operating system.

    - The "Windows Applications Exploratory Test Procedure," for information about the exploratory test methodology.

- The "Programming Best Practices with Microsoft Message Queuing Services (MSMQ)" white paper, available through the Message Queuing link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about applications and MSMQ.

- The "Software Update Services Deployment" white paper, available through the Software Update Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about SUS.

- The "Application Deployment Using Microsoft Management Technologies" white paper, available through the Application Deployment link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about choosing an application distribution tool.

**Related Tools**

- Application Compatibility Toolkit

  Use the Application Compatibility Toolkit to test, identify, and resolve application compatibility problems. To download the Application Compatibility Toolkit, see the Windows Application Compatibility link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Sysprep.exe, Setupcl.exe, and Factory.exe

  Use Sysprep.exe, Setupcl.exe, and Factory.exe to prepare a hard disk for disk imaging. To obtain Sysprep, open the **Support\Tools** folder on the Windows XP Professional or Windows Server 2003 operating system CD, and then open Deploy.cab. You can also find these tools by clicking the Windows Downloads link on the Web Resources page at

http://www.microsoft.com/windows/reskits/webresources.

- Debugging Tools for Windows

  Use the Debugging Tools for Windows to identify complex or hard-to-find application compatibility problems when you have the source code and symbol files for an application. To download the most recent version of Debugging Tools for Windows and the appropriate symbol files, see the Debugging Tools link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Job Aids**

- "Compatibility Solutions Spreadsheet" (PDP_APPA_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Compatibility Solutions Spreadsheet" on the Web at http://www.microsoft.com/reskit) for a spreadsheet to help you find application compatibility fixes for applications compatibility problems.

- "Windows Fundamentals Testing Checklist" (PDP_APPA_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see the "Windows Fundamentals Testing Checklist on the Web at http://www.microsoft.com/reskit) for a checklist to help you track your Windows Fundamentals tests.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Msiexec (command-line options)" in Help and Support Center for Windows Server 2003 for information about Msiexec.exe command-line syntax and usage.

- "Start" in Help and Support Center for Windows Server 2003 for more information about using the **Start** command to run installation programs in batch files. To find "Start" in Help and Support Center for Windows Server 2003, look under **Support Tasks**, click **Tools**, and then click **Start**.

- "Group Policy" in Help and Support Center for Windows Server 2003.

- "Logon Scripts" in Help and Support Center for Windows Server 2003 for information about configuring logon scripts.

- "Event Viewer overview" in Help and Support Center for Windows Server 2003 for information about the event log and audit log.

# Designing a Pilot Project

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Conducting the pilot is the last major step before deployment of the Microsoft® Windows® Server 2003 operating system. During the pilot, your release management team tests your design in a controlled environment in which users perform their normal business tasks using the new features. This demonstrates that the design works in the production environment as expected and that it meets the organization's business requirements.

**In This Chapter**


Overview of Designing a Pilot Project

Creating a Pilot Plan

Preparing for the Pilot

Deploying and Testing the Pilot

Evaluating the Pilot

Preparing for Production Deployment

Additional Resources for Designing a Pilot Project


**Related Information**

- For more information about Microsoft Solutions Framework, see the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- For more information about deployment project planning, see "Planning for Deployment" in this book.
- For more information about testing your implementation design prior to conducting the pilot, see "Designing a Test Environment" in this book.
- For more information about how to pilot specific technologies, such as Microsoft® Active Directory® directory service, see the related books in this kit.

# Overview of Designing a Pilot Project

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A *pilot* release is a deployment of your Windows Server 2003 or Microsoft® Windows® XP Professional operating system to a subset of the live production environment or user group. The primary purposes of a pilot are to demonstrate that your design works in the production environment as you expected and that it meets your organization's business requirements. Conducting a pilot reduces your organization's risk of encountering problems during full-scale deployment. To further minimize your risk during deployment, you might want to conduct multiple pilots consisting of separate pilots for different technologies or operating systems, or you might want to conduct a full-scale pilot in phases.

Depending on the context of the project, the pilot release can take the following forms:

- In a general business enterprise, a pilot can be conducted with a group of users on a set of servers in a datacenter.
- In a Web development enterprise, a pilot can be conducted by hosting site files on a staging server or servers or in folders that are live on the Internet, using a test Web address.
- Commercial software vendors, such as Microsoft, often conduct a pilot with a special group of early adopters.

The pilot participants provide feedback about how well the design and features are working. The release management team uses this feedback to resolve issues that arise or to create a contingency plan. Ultimately, the pilot leads to a decision to proceed with a full deployment or to slow down so you can resolve problems that could jeopardize your deployment.

The Microsoft Solutions Framework (MSF) training course offers in-depth instruction on the practical use of MSF for planning and conducting a deployment project. For more information about the MSF training course, see the Microsoft Training and Certification link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, and search for MSF.

# Process for Designing a Pilot Project

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The pilot process is iterative. You deploy Windows Server 2003 on a limited number of servers in a controlled environment, evaluate the results, fix problems, and deploy an updated release of your pilot, repeating this process until you are satisfied that you are ready for a full deployment.

If you are deploying more than one operating system, or if you are deploying into an environment that contains more than one operating system, you might want to conduct multiple pilots. This allows you to isolate problems more quickly, because you do not have to research problems on multiple platforms. For example, if you are deploying both the Windows XP Professional and Windows Server 2003 operating systems, you can conduct separate pilots for the client operating system and for the server operating system, and then conduct a final pilot in which the two operating systems are deployed at the same time.

Figure 4.1 illustrates the process of planning and conducting a pilot project.

**Figure 4.1 Designing a Pilot Project**



# Pilot Planning and Deployment Within the Project Cycle

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The pilot is conducted during the stabilizing phase of the MSF project cycle. However, planning for the pilot occurs during the envisioning and planning phases of the deployment project, and preparing for the pilot occurs during the developing phase. Figure 4.2 illustrates the tasks involved in planning for and conducting a pilot and shows the MSF phase during which each of these activities might occur. The time frames are estimations, and they might vary from deployment to deployment.

**Figure 4.2 Role of the Pilot in the Project Life Cycle**



# Creating a Pilot Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

During the planning phase of the deployment project, product management, program management, and release management teams collaborate to create the pilot plan. The pilot plan defines the scope and objectives of the pilot and identifies pilot participants and where the pilot will be conducted. It includes a schedule for deploying and conducting the pilot and plans for training and communicating with pilot participants, evaluating the pilot, identifying risks and contingencies, and other key activities that occur during a pilot deployment.

You have various options for how to conduct your pilot. For example, you might have multiple pilots, with a separate pilot for Windows XP Professional and one for Windows Server 2003, or you might introduce your implementation progressively with a staggered rollout, gradually adding more user groups to your pilot.

If you plan to conduct multiple pilots, develop a pilot plan for each. For example, if the subteam responsible for deploying a particular operating system plans to conduct its own pilot, that subteam should write a pilot plan. Even if you do not plan to conduct multiple pilots, each subteam can contribute to the overall pilot plan.

Figure 4.3 illustrates where creating a pilot plan occurs in the deployment pilot process.

**Figure 4.3 Creating a Pilot Plan**




When the pilot plan is ready for review, have project team members, necessary support personnel, and management representatives read and approve the plan. Be sure that the supervisors of everyone directly affected by the pilot have a chance to review the plan. For example, if the schedule allots time for a particular user group to participate in the pilot, have the supervisor of that group review the schedule.

# Defining the Pilot Scope and Objectives

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in planning your pilot is to define what you want to accomplish (the objectives, or goals) and what you plan to include and exclude (the scope). Be sure to align the pilot objectives and scope with those for the deployment project as a whole, as defined in the master project plan. Ensure that the pilot plan includes an opportunity for the team to evaluate features identified in the project plan as high priority, to ascertain that they successfully meet all of your business objectives.

If you plan to conduct multiple pilots, define the scope and objectives of each by using these guidelines.

# Defining Pilot Objectives

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Explicitly state the objectives of the pilot. Use the objectives to identify criteria for measuring the success of your pilot.

Many organizations have primary pilot objectives, such as:

- Ensure that the system works properly in the business environment.
- Ensure that the design meets the business requirements.
- Build user support for the Windows Server 2003 deployment project.

Many organizations have additional objectives, such as:

- Test the deployment process.
- Train the installation team.
- Create documentation for the full deployment.
- Train the support and Help desk teams.
- Train the administrative teams.
- Gather information for estimating future support requirements.
- Gather information for estimating actual hardware requirements.

  **Note**

  - Sometimes a pilot reveals that there is a greater amount of data in the production environment than anticipated and that storage requirements exceed those that were proposed in the test lab. As a result, hardware requirements might exceed those previously estimated.

- Meet the baseline requirements for functionality that were established in testing.
- Develop and test end-user training materials.

You might need to define objectives that are related to specific technologies, such as Active Directory, or to your network infrastructure. If you plan to conduct multiple pilots, define the objectives for each.

If possible, use the objectives to develop metrics for evaluating the pilot. For example, if your team's objective is to gather information for estimating future support requirements, establish baseline metrics for current system functionalities, such as system availability, performance, and reliability, and collect similar metrics for the system after the pilot is complete. These metrics can then be used to evaluate the success of the pilot and to estimate the level of support required.

# Defining Pilot Scope

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Define the scope of the pilot by clearly stating which services and features will be included and which will not. When you list the services and features you plan to include in the pilot, also state how you expect them to perform. Describe the areas of functionality that the pilot implementation affects, and note to what extent they are affected, and in which situations they are affected.

Avoid the mistake of making pilot testing too easy. Test as many situations that might arise in production as you possibly can.

If certain technologies have particular requirements, include those specific requirements in your pilot. For example, if you are going to deploy Windows Server 2003 on an international scale in which users do not all speak one language, be sure to address international language issues in the pilot.

Do not expect to test every feature or service during the pilot. Focus on processes that present the greatest risk and events that are most likely to occur. Prioritizing the features to be tested in the pilot is particularly important if your team plans to conduct multiple pilots. When multiple pilots are planned, start small and gradually increase the scope of successive pilots.

If certain aspects of your design cannot be covered by the pilot, describe them. For example, if your organization plans to upgrade your domain using your existing architecture without any restructuring, and then restructure the architecture later, you might choose to exclude the restructuring process from the first pilot.

Also specify the duration of the pilot, in terms of either time or of the criteria to be met.

Be sure to describe how you expect to proceed after the pilot is complete. If you plan to keep some functions in place and remove others for the full production rollout, identify the features that will be removed. For example, if you are redesigning your namespace, you might want the option of changing it after the pilot concludes. Indicate how to back out features that will be removed.

# Defining the Pilot Group

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Carefully select the pilot participants and the sites where the pilot will be conducted. First establish the selection criteria for participants, and then select candidates based on these criteria.

**Selecting the Pilot Group and the Pilot Sites**
The pilot group should include end users who are typical of those in your organization. This might include a wide range of users, some of whom are technically proficient and others who are less comfortable with technology. If you plan to conduct multiple pilots, the type of user you select might vary as the pilots progress. For an early pilot, however, good candidates have some of the following traits:

- Are able to derive tangible benefit from Windows Server 2003.
- Play a noncritical role in day-to-day operations.

  The group should be able to absorb some downtime or reduction in performance if problems occur.
- Represent the target environment.

  Choose groups or sites that do not have unique requirements or operating environments, because the pilot should predict how your design and rollout will work in your organization at large. Most pilot participants should have the same technical expertise as the average user in your organization.
- Perform a variety of activities using a variety of computer hardware.
- Are enthusiastic about the Windows Server 2003 project.
- Are comfortable with technology.

  This is not an essential attribute for pilot participants. However, if participants are comfortable with technology, they tend to be more patient with problems that occur during a pilot, and they are more likely to test the limits of the Windows Server 2003 operating system. This type of participant, however, might accept problems that should be reported. Encourage participants to report every problem they encounter.

You can identify candidates by conducting interviews, issuing and evaluating questionnaires, or requesting volunteers. Volunteers, in particular, can be very helpful, because their offer to participate indicates an interest in the Windows Server 2003 technology. You might also consider soliciting recommendations for users who might be good candidates for the pilot; for example, you might ask supervisors to identify power users of a particular technology that is going to be deployed.

The selection of a pilot site or sites often depends on the type and location of the pilot participants who have been selected and the number of support staff available to help them. Determine the number of pilot sites and the size of the pilot users group based on:

- The objectives of the pilot.
- The number of functions and features you are testing.
- The size of your support staff.
- Specific technology requirements that can be piloted only by using particular sites or user groups.

Choosing suitable pilot groups and sites helps to ensure a successful pilot project.

**Using the IT Group as a Pilot Group**

Many organizations use their IT group for the first pilot. Because this group is familiar with testing software and hardware changes, they are adaptable and can report issues quickly. After problems with the first pilot have been resolved, you can begin your full-scale pilot in the production environment.

**Selecting a User Liaison**

After you select your participants, choose one as your user liaison. Select someone who has good communication skills and a good relationship with both the pilot group and the project team.

Work with the user liaison as you plan the pilot. The liaison can provide you with information about the type of work the pilot group performs and can prepare the group for its role in the pilot.

# Creating Pilot Plan Documents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The pilot plan includes a number of documents that provide procedures for successfully rolling out the pilot in a business environment. The pilot plan should include:

- A training plan that describes what the pilot participants need to know before they begin the pilot and describes how you plan to train them.
- A support plan that describes how problems that arise for participants during the pilot will be resolved.
- A communication plan that explains how you will keep participants informed about what is happening in the project.
- An evaluation plan that describes how you plan to obtain feedback from participants.

  This can be used to assess design changes that might be required before the pilot is rolled out to the rest of the organization.
- A risk and contingency plan that describes potential risk factors and how you plan to assess and diffuse them.
- A backup and recovery plan that describes how to back up data and system configurations, how to test those backups, and how to roll back to previous configurations if problems arise that might affect business operations or participants' access to data.

  Having a backup and recovery plan is particularly important because the pilot is a rollout in a business environment.
- A schedule that includes milestones at which you plan to evaluate and make necessary changes to the pilot.

The pilot plans you create help to ensure the smooth deployment and completion of your pilot project by clarifying the tasks and roles of the release management team and the pilot participants.

# Creating a Pilot Training Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Well before you begin the pilot, your team needs to decide which features of the new system the participants need to be trained on, the form training will take, and when training will occur.

Training for pilot participants should be limited to the features that they must understand to do their work. Include a plan to provide training for the support and operations teams, which might need to be more detailed than the training presented to pilot participants.

How you train participants depends greatly on your facilities and your budget. Options include hiring an external trainer, holding brown-bag seminars, developing a "train-the-trainer" program, or using media technology to broadcast training. You might also consider using the test lab for training after testing is complete.

After you have determined what the training should cover and how it will be accomplished, estimate how long it will take and decide when to begin. Many organizations find that it works best to provide training just prior to pilot installation. Remember to include training in the pilot plan schedule.

# Creating a Pilot Support Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The support plan identifies who will provide support for pilot participants, the level of support required, and how users can report problems. Be sure to develop your support plan early, because doing so will help you identify the type of training the support staff requires.

**Defining Support Team Roles**
For each pilot, identify who will support participants: project team members such as developers and testers, Help desk personnel, or outside resources. Even if the project team will not be providing primary support for the pilot, decide what the role of the team will be. If the project team is divided into subteams for different technologies, such as Active Directory or network infrastructure, determine whether the subteams should play a role in supporting participants and, if so, what that role should be.

If Help desk personnel will provide support, or if one of your pilot objectives is to train Help desk personnel, determine how they will be trained. If external resources are required, ensure that they are aware of the schedule and have been properly trained.

Develop a contact list for support staff that includes the names, phone numbers, and e-mail addresses of all support staff members and other key contacts. For example, provide the contact name and information for a manufacturer who supplies key equipment.

**Service Levels for Support**
Determine the level of service that the support team can provide during the pilot. The level of service depends on the number of support personnel available and the number of problems they can support. For example, if support staff can install upgrades to 20 client computers, but can provide support to only five participants if they all have problems, plan to upgrade only five client computers, so the team can provide the appropriate support level.

When determining support requirements, some typical questions to ask include:

- Must critical problems be resolved within a specific number of hours? If so, what is that number?
- During which hours must support be available to the users?

The level of service your support team can provide dictates the size of your pilot and the timeline for its deployment. In addition, determining the level of service helps you establish a baseline for a level of service that can be used to create final Service Level Agreements (SLAs) for the organization.

**Problem Tracking and Resolution**
When problems arise during pilot testing, participants must have a way to report them to the team. The incident tracking system, problem escalation process, and issue resolution process your team develops should identify:

- Where participants post their problems.

  Can they report problems to an existing incident-tracking system, or does the team need to develop a new mechanism, such as a Web site, for participants to use to log their problems?

  If your testing team set up an incident-tracking system to record and report testing problems, you can use the same system to report and track problems that arise during the pilot. If your organization has not developed an incident-tracking system, set one up before the pilot begins. Be sure that the incident-tracking system is well documented so pilot participants can easily enter problems into the system.

- How problems will be reviewed, prioritized, and fixed.

  For example, support staff can review reported problems on your incident-tracking system daily, and a problem can be assigned to the support staff personnel responsible for the applicable technology or functional area. After a problem has been assigned, the support team member responsible for it can prioritize the problem and estimate the amount of time required to fix it. If this information is entered into the incident-tracking system, participants who report problems can

track the status of their problems.

- The escalation process the team will use to notify the appropriate team members.

  The incident-tracking process should have a procedure for escalating problems that can not easily be resolved to the appropriate team for resolution. Determine who is responsible for deciding when a problem must be escalated, and who should contact the team who can resolve the issue.

  During a pilot, problems must be escalated and resolved much more quickly than during the testing phase, because pilot participants are working in a business environment. How quickly problems will be resolved after the escalation process has begun should be spelled out and understood by both the pilot participants and the team. Set reasonable expectations for issue resolution so both support staff and pilot participants are comfortable with the time allowed for the escalation and resolution of problems.

- How change requests are submitted, approved, tested, and implemented.

  Your change management plan should address these issues. The way your team handles change management can vary depending upon when the change is requested: during testing, during the pilot, or when Windows Server 2003 has been rolled out to the business environment.

Note that the procedures for resolving a problem that arises during the pilot might vary significantly from those used during full production deployment. For example, a bug in the pilot might be discovered and fixed, and an updated release deployed to participants very quickly. Depending on the severity of a problem, when a bug is discovered in production, the team might be required to go through a lengthy change management and deployment process before the bug is resolved and an update to the system is made available in the production environment.

# Creating a Communication Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the communication plan, describe how you will communicate with participants both before the pilot begins, to prepare them for the pilot, and during the pilot, to exchange status reports. Use the communication plan to identify:

- The type of information you will communicate, to whom it will be communicated, by what means, and how often.

  If you provide too much information too often or to too many people, participants might start ignoring your communications. Remember that it is important to report successes, too.

- Mechanisms for communicating information about the pilot, such as Web sites, frequently asked questions (FAQ) pages, procedures, and status reports. As you determine how you will communicate with participants during the pilot, begin creating the required mechanisms.

  For example, set up e-mail distribution lists for the various groups that need to receive specific types of information. Note the types of information you plan send to those on each distribution list. For example, from time to time you might want to communicate information about the pilot to others in the organization to keep them informed about the progress of the pilot and to encourage their support for the change when it is their turn to upgrade.

- What you need participants to do to prepare for the pilot. In your instructions to pilot participants, include the ramifications of failing to comply with the instructions. For example, if you need participants to log off at night so you can update user logon settings, tell them that failure to log off might result in their inability to log on at all because their settings could not be updated.

**Tip**

- Use links in your e-mail correspondence to pilot participants to provide detailed information without specifying it in the e-mail. For example, if you have users of varied technical competency, provide a link to a Web site that contains additional technical information for those who need the detail, instead of providing it in the e-mail. This approach also allows you to update instructions and details on the Web site as the need arises without having to resend the information.

# Creating an Evaluation Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The evaluation plan describes the way you plan to obtain and assess feedback from the team and pilot participants. This feedback can be used to evaluate design changes that might be required before the deployment is rolled out to the rest of the organization. In addition, it helps your team decide whether the deployment project should proceed and, if so, how.

In the evaluation plan, define the process you plan to use to solicit feedback from participants and the team. Determine the types of questions the evaluation will ask. Decide on a review process, and determine who is responsible for making the decision about the next step to be taken.

# Creating a Risk and Contingency Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The risk and contingency plan describes the risk factors that could prevent the pilot from being deployed successfully. Risks might include the test lab being behind schedule, required hardware or software being unavailable, or participants working on other projects or needing additional training.

Properly addressing risk factors that could prevent a successful deployment in a pilot environment reduces the likelihood of encountering those same problems when you deploy into your production environment.

# Creating a Backup and Recovery Plan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The backup and recovery plan establishes guidelines and procedures to prevent problems that might cause data loss or interruptions to your organization's operations, and to allow recovery as quickly as possible if such events do occur.

Consider planning downtime or outages for the pilot in order to test rollback procedures, and, if applicable, disaster recovery and business continuity plans.

**Creating a Backup Plan**

The importance of a backup plan cannot be overstated. When you begin rolling out Windows Server 2003 in the business environment, problems might arise that even the most thorough testing could not reveal. By making regular and reliable backups, you ensure that the team can restore the system to its original state if your pilot rollout process changes or fails.

The backup plan should define procedures for:

- Backing up baseline configurations (the state of a computer before it is upgraded) so a computer can quickly be restored to its prior state.
- Backing up servers before they are upgraded.
- Backing up the most recent system and user data before you begin switching systems.
- Testing the restore process for each of the above processes using the backup files.

In addition, the backup plan should identify who is responsible for performing backups, and should include the schedule for all periodic backups and periodic testing of backups, as well as instructions for labeling and storing all backup files.

For more information about creating backup plans, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Creating a Recovery Plan**

The recovery plan describes the recovery and rollback process, which allows you to return your production system to whatever earlier state you require. Depending on the severity of a problem encountered in the pilot, you might need to return your production system to a baseline configuration or just roll it back to the state it was in at a particular point in time.

Include the following elements in the recovery plan:

**A list of scenarios** Analyze all of the systems involved in the rollout and identify the situations, or scenarios, under which problems are likely to occur. Determine which systems might be affected and the functional dependencies among them so you have a clear understanding of the larger impact that a single failure might have. Use these scenarios to create strategies that identify when and how to run backups and the types of recovery for which you need to plan.

**A definition of acceptable downtime** Define how much downtime your organization can accommodate. If your organization cannot afford for systems to go down during normal business hours, you might plan to roll out the pilot, or parts of it, at night or over a weekend. If systems must be operational at all times, you might plan to deploy servers and desktops on new computers and then quickly replace the old ones, instead of upgrading computers.

**A list of critical systems and processes** In the event that a failure does occur, you need to know which systems are the most critical and must be brought back online first. If resources such as bandwidth are limited, you need to know which systems have the highest priority and which should not take up network traffic. When evaluating how critical a system or process is, consider factors such as its effect on human health and safety, the legal liability it exposes, the risk to corporate confidentiality, and the cost of replacement.

**A recovery strategy** Your recovery strategy should define how you will recover data or systems in each of the scenarios you define in the recovery plan. This might include restoring data from backup tapes, switching over to redundant systems, rolling back to previous configurations, or other strategies. Include an additional procedure for recovering from severe data corruption in your directory service if that becomes necessary. By having your recovery strategy in place, you can quickly restore your production environment to the required state so that work can continue with minimal interruption.

**A rollback strategy** The rollback strategy defines how you plan to use backup and recovery procedures to return your pilot or production environment to the state it was in before changes were made. Specify the criteria that a problem should meet to warrant rolling the environment back to its previous state. For example, you might establish a system for classifying the severity of problems and describe which type of response is warranted by certain levels of severity. Also decide whether you need to have different rollback strategies for different types of problems. For example, you might develop one procedure for backing out the entire pilot if the problem is pervasive and another procedure for backing out specific components if the problem is isolated.

**The roles and responsibilities for team members** Make sure that every task in the plan is assigned to an appropriate team member, and that that person has the information needed to successfully perform required tasks. Consider including training in the plan.

Have the backup and recovery plan reviewed by the project team and by those responsible for potentially affected systems. After the plan has been approved, test it to ensure that the processes you put in place work as expected.

# Creating a Pilot Schedule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

One of the earliest activities in planning a pilot is to draft a schedule, which is usually included in the master project schedule. When drafting the pilot schedule, be sure to schedule time for:

- Planning the pilot.
- Training participants, support staff, and the installation team.
- Creating an inventory of pilot sites.
- Upgrading hardware.
- Developing the support and communication mechanisms you identified during planning.
- Deploying the pilot.
- Testing by pilot participants during the pilot.
- Evaluating the pilot.

To accurately schedule pilot deployment, you need to know:

- How many computers or servers need to be upgraded and the approximate time it will take to upgrade each one.

  After you have gathered this information, determine how many machines can be upgraded per day and the sequence in which they should be upgraded.

- Which hours of the day and which days of the week are best for upgrading servers and client computers. Questions to ask include:

  - Should computers be upgraded during off hours to avoid disrupting users?
  - Should client computers be upgraded during working hours so that the users can attend training during that time?

- The types of upgrades that are required.

  If you are upgrading servers, the pilot might require a longer period than if you are upgrading only client computers, because some server processes occur on a monthly, rather than a daily, schedule.

- Whether participants should receive training before Windows Server 2003 is upgraded or installed on their computers. If so, the training schedule is a dependency for your pilot deployment.

After you have drafted the pilot schedule, review it with management and all affected users. After you deploy the pilot, update the master schedule based on your installation experience so that it more accurately reflects the time that will be required for installation during the full deployment.

# Preparing for the Pilot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Preparation for the pilot deployment begins in the developing phase of the project. The release management team sends communications to the participants about their role in the pilot and prepares an inventory of the software and hardware that will be used with Windows Server 2003. Finally, the testing team prepares procedures and tests to aid in the rollout and the validation of the success of the rollout.

Figure 4.4 illustrates what is involved in preparing for the pilot and when this step occurs in the deployment pilot project.

**Figure 4.4 Preparing for the Pilot**



# Preparing Pilot Sites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Prepare the pilot sites in advance so that the installation team can begin to install upgrades when the pilot begins. If you have not already done so, inventory the computers and network equipment that are used at the pilot site. Then determine whether the following hardware or software upgrades are required:

- Client computer upgrades to meet your minimum supported hardware configuration (including memory, hard disk capacity, processor speed and type, and network adapters). Your minimum supported hardware configuration is based on the baseline configuration metrics established in your pilot objectives.
- Server upgrades for optimal hardware configurations.
- Network upgrades to meet design requirements.
- Client and server upgrades (hardware, applications, or drivers) to provide compatibility with Windows Server 2003.

For more information about compiling an inventory of your network equipment, see "Planning for Deployment" in this book.

At the very least, identify the required upgrades and acquire the required components. If possible, have the new components installed and tested before the pilot is scheduled to begin.

In addition, use the inventory of computers and system components to identify the following:

- All applications in use at the pilot site
- Any special security measures required at the site
- All special connectivity requirements for the site
- Minimum system requirements as determined by your baseline functionality metrics

This information allows the release management team to plan for the types of testing that might be necessary during the pilot, such as tests that push the limits of available network connectivity for a site.

Ensure that the testing team has tested all the hardware and software present at the pilot site to verify that they are compatible with Windows Server 2003. Also make certain that the installation team is aware of any special requirements and is equipped to handle them.

# Preparing Pilot Participants

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Establish communication with the pilot group as early in the pilot planning process as possible. Your initial contact should open the channel of communication and set participants' expectations. Then, as the start date for the pilot approaches, prepare participants for their role by keeping them apprised of pilot deployment plans and target dates and by providing necessary training.

**Establishing Early Communication**
Soon after you select the participants, meet with them to:

- Obtain their commitment to the pilot.
- Select a user liaison.
- Provide them with a timeline for the pilot.
- Clarify participants' responsibilities.
- Describe the type of testing they are to perform.

  Although pilot participants will continue to perform their daily business tasks, you might want to specify certain areas they should focus on. Provide test plans for any special tasks they are to perform.
- Discuss communication mechanisms that have been established and how they should be used by participants.
- Discuss the level of support you plan to provide.
- Clarify problem reporting and issue resolution procedures.
- Discuss rollback plans.

Pilot participants need to understand what the pilot entails. Ensure that they understand how the pilot might affect their work. Be sure to address any concerns they might have about the pilot or their role in the deployment process.

**Keeping Participants Informed**
As the pilot plans develop, the user liaison can keep the release management team informed of the participants' concerns and the participants informed about new developments.

As the training plan progresses, tell participants about the type of training they will receive and when they can expect it. When the support plan is finalized, explain to participants how and when they should request support. Also explain the mechanisms provided for reporting problems or issues.

As the pilot deployment date approaches, the release management team should provide participants with the following information:

- Scheduled dates for training and for upgrading computers
- Procedures they need to follow before their computers are upgraded
- Contact names and numbers for support

Regular and comprehensive communication with pilot participants helps to ensure that participants are committed to the success of the pilot project.

**Training the Pilot Participants**

Train pilot participants just before the pilot is scheduled to start. The duration and type of training you provide depends on the scope of the pilot, as determined in the training plan.

Training might be required for both support staff and pilot participants. If so, train the support staff before the pilot participants, so that the participants can have assistance as soon as the pilot begins if they need it.

# Testing the Rollout Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although the test lab is a good place to debug problems with the proposed deployment of Windows Server 2003, the pilot provides an opportunity to assess the accuracy and efficiency of deployment procedures in the actual production environment. If one of your pilot goals is to test the Windows Server 2003 rollout process, schedule time during the testing phase for the installation team to develop, document, and test the rollout procedures. Preparing the pilot site also can provide valuable information about the rollout process, such as the time required to upgrade each computer and the number of concurrent upgrades your support team can accommodate, and information about other upgrade tasks.

Procedures and resources required for deploying Windows Server 2003 on various types of computers should be carefully documented throughout the planning, testing, and developing phases of the project. These documents will form the core of the installation guidelines and should include:

- A list of the required tools and supplies.
- A list of automated installation scripts and where they are stored.

  Ensure that any scripts and tools for automating installations obtained from the testing team are appropriate for the computers in the pilot environment. For more information about automating installations, see "Choosing an Automated Installation Method" in *Automating and Customizing Installations* of this kit.
- A list of backups that installers must make before and during deployment.

  Include backups of user data on client computers.
- Procedures for performing both automated and manual computer installations and upgrades.

  If the automated installation method does not work properly, you can install the operating system manually. For more information about automating installations, see "Choosing an Automated Installation Method" in *Automating and Customizing Installations* of this kit.
- The acceptance tests that the installers must perform during pilot deployment and immediately after it to verify that the deployment works as expected. These tests are developed by the testing team during the developing phase. For more information, see "Designing a Test Environment" in this book.
- The operational procedures that installers and administrators must perform (for example, resetting permissions, changing passwords, and restoring user data).
- Procedures for backing out if the pilot fails. Review the backup and recovery plan for information about to how to back out system changes. For more information about creating the backup and recovery plan, see "Creating a Pilot Plan" earlier in this chapter.

Most of this documentation, including a list of automated installation scripts and operational procedures, can be obtained from the testing team. Use the documentation to create a checklist of deployment tasks that can be used to ensure that all necessary steps have been completed.

# Deploying and Testing the Pilot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you deploy your pilot, you begin testing your Windows Server 2003 implementation under live conditions. Start the pilot deployment process with a trial run of the pilot to identify problems with the deployment and the pilot plan. Then, when the full pilot begins, keep track of which deployment tasks have been completed so that the team can monitor the progress of the pilot.

As participants use the system, have the release management team track the progress of the pilot, and pinpoint areas of concern. Encourage all participants to use the incident-tracking system to report problems and to use the escalation plan when immediate problem resolution is not possible.

In MSF, the pilot is conducted during the stabilizing phase. Figure 4.5 illustrates the steps involved in deploying and testing a pilot.

**Figure 4.5 Deploying and Testing the Pilot**




**Conducting a Trial Run of the Pilot**

Before you deploy the pilot, perform a trial run of the process. A trial run is a rehearsal of all the elements of the deployment that allows your team to identify problems prior to deploying the pilot.

To perform a trial run, schedule a time during nonbusiness hours to perform the entire upgrade process, to test the new setup thoroughly, and then to back out all components of the implementation. Use the same deployment task checklist that you plan to use for deploying the pilot, as described in the next section.

**Deploying the Pilot**

As the installation team deploys the pilot, they need to:

- Validate all backups, clearly label them, and store them in a safe place.
- Record how long installation takes so the team can update the schedule.
- Identify and document inefficiencies, and use this information to refine the rollout process.
- Update the production deployment documents to reflect all changes that have been made during the pilot.
- Ensure that a system administrator who has full security administrative credentials, including rights to administer mail

and database server passwords, is available during the pilot deployment.

Instruct the installation team to check off each of the above steps in the deployment tasks checklist as it is completed.

**Resolving Issues During the Pilot Deployment**
During the pilot, the participants perform their regular, daily work and report any problems by using the incident-tracking system.

Problems that are reported are reviewed, prioritized, and fixed according to procedures outlined in the issue resolution plan. Problems can be resolved either by further development, by documenting resolutions and workarounds for the installation teams and support staff, or by incorporating the resolution or workaround as supplemental material in training courses.

Depending on your team's rollback strategy, resolving problems might require rolling back the pilot. Clear communication with participants is essential when a rollback is required.

If multiple pilots are scheduled, evaluate the results of the current pilot to determine if it meets the success criteria before beginning the next pilot.

**Monitoring the Pilot**
During the pilot, the release management team should continually monitor the pilot system, including the network used by the pilot group, looking for bottlenecks and areas that need to be fine-tuned. The more information you collect during the pilot, the more accurately you can evaluate its success and recommend how to proceed with the full deployment to the entire business environment.

Have the team check problem reports frequently and look for trends and for problems with traffic flow and application performance, such as:

- Degradation in performance, such as slower response time.
- Improvement in performance, such as quicker response time.
- Tasks that could be performed before, but not after, the upgrade.
- Tasks that can be accomplished only by using a workaround.

It is important that release team members visit the pilot site periodically. Talking with users frequently reveals issues or problems that might otherwise go unnoticed.

As the release management team receives feedback, they will need to assess problems that pose risk to the overall project. Specifically, they should look for situations that might result in:

- Scope changes
- Cost increases
- Interoperability problems
- Unanticipated downtime

Assessing the severity of these issues allows the release management team to make informed decisions about whether to proceed with the deployment.

# Evaluating the Pilot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When the pilot is complete, obtain feedback from a variety of sources, including participants, pilot release management and support teams, and other observers, to evaluate the pilot's success.

Figure 4.6 illustrates when pilot evaluation occurs in the pilot deployment process.

**Figure 4.6 Evaluating the Pilot**



# Obtaining Feedback

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Have your incident-tracking system in place when the pilot begins, and encourage pilot participants to use it to report problems. Participants often neglect to report problems, either because they think a problem is insignificant or because they find a way to work around it. To accurately assess your pilot, however, you need participants to report every problem.

**Tip**

- To encourage participants to offer feedback, consider developing an incentive program. For example, you might give away prizes or have managers give recognition to participants who make a special effort to contribute during the pilot.

Gather feedback from a variety of sources and through a variety of means, including:

- Web site feedback forms
- Sessions with business managers
- Problem reports
- Participant surveys
- IT project team observations

Try to obtain information about both the Windows Server 2003 design and the deployment process, including information about:

- Training
- The rollout process
- Support
- Communications
- Problems encountered
- Problems that were fixed without the need for incident-tracking
- Potential improvements

Review the parts of the pilot that were successful and the types of problems that were reported so that you can revise and improve your plan. Determine whether the pilot met the success criteria you defined before the pilot began. If you established metrics to measure your success, evaluate how the pilot measured up. Use the feedback to evaluate whether the delivered design meets the design specification, as well as the business requirements.

# Deciding the Next Step

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Once enough pilot data has been collected and participant feedback has been evaluated, the team must decide how to proceed. Depending on how well the pilot meets the success criteria, there are a number of strategies that your team can employ at this point in the pilot deployment:

- **Stagger the pilot forward.** If the pilot was deemed partially successful, deploy the pilot to the next pilot group or, if your team has planned to conduct multiple pilots, proceed with the next pilot.

- **Roll back the pilot.** When the pilot is not completely successful, it is often necessary to roll it back so that issues can then be resolved. Typically, a rollback is required when:

    - The production environment contains invalid or problematic data that was not discovered in testing and that caused the deployment to fail. Roll back the pilot to the configuration used before the pilot began, clean up the data, and redeploy the pilot so that the pilot group can continue with its work.

    - Production configurations and settings, such as Group Policy settings or security restrictions, cause problems with the deployment that could not be detected in the test lab. After rolling the pilot back to the configuration used before the pilot began, resolve the problems with the conflicting configurations or settings and redeploy the pilot.

- **Suspend the pilot.** If the pilot is not successful and issues cannot be resolved easily, suspend the entire pilot, halting all pilot testing until the issues have been resolved and the pilot can be redeployed. This requires rolling the pilot back to the configuration used before the pilot began.

- **Update the pilot and continue.** If the pilot is not successful, but the issues raised are easily fixed, issue the same pilot group an update to existing code.

- **Proceed to the production deployment phase.** If the pilot is deemed successful and ready for production, you can proceed with your plans for full deployment.

The pilot is not complete until the team ensures that the proposed solution is viable in the production environment and that every component of the solution is ready for deployment.

# Preparing for Production Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After your team has agreed that the pilot has been successfully completed and has obtained management approval for proceeding, the next step in your deployment project is to fully deploy the system to the entire organization. During this phase, the release team deploys the core technology and site components, stabilizes the deployment, transitions the management of the project to the operations and support teams, and obtains final management approval of the project.

Before the actual deployment begins, ensure that the production deployment plan is complete. The production deployment plan contains many of the same types of documents as the pilot plan. For a discussion of pilot plan documents, see "Creating a Pilot Plan" earlier in this chapter.

Keep in mind that your production deployment plan must include a strategy for deploying the Windows Server 2003 implementation across your entire organization, and, as such, must address such issues as how to manage specific requirements for particular sites, how to train all users, and how the deployment schedule will affect work in various departments. The production deployment plan is a living document that should continually be updated to reflect changes that are made when testing of the release components and processes reveals a deficiency.

Many of the tools that your team developed for the pilot deployment also can be used for the production deployment. These include the list of tasks that must be performed to deploy Windows Server 2003, the issue resolution process your team established, and the mechanism for monitoring your system.

MSF contains information about deployment that your team might find helpful for planning your production deployment. For more information about MSF, see the "MSF Process Model" white paper, available through the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Additional Resources for Designing a Pilot Project

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2



**Related Information**

- "Planning for Deployment" in this book for more information about deployment project planning.
- "Designing a Test Environment" in this book for more information about testing prior to the pilot.
- "Choosing an Automated Installation Method" in *Automating and Customizing Installations* of this kit for more information about automating installations.
- "Designing Unattended Installations" in *Automating and Customizing Installations* of this kit for more information about automating installations using unattended installations.
- "Designing Image-based Installations with Sysprep" in *Automating and Customizing Installations* of this kit for more information about automating installations using image-based installation methods.
- "Designing RIS Installations" in *Automating and Customizing Installations* of this kit for more information about automating installations using Remote Installation Services (RIS).
- The *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit) for more information about creating backup plans.
- The Microsoft Training and Certification link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about MSF training courses.
- The "MSF Process Model" white paper, available through the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about the Microsoft Solutions Framework.

# Planning Server Deployments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Welcome to *Planning Server Deployments* of the *Microsoft Windows Server 2003 Deployment Kit*.

This book provides comprehensive information about planning server storage and designing and deploying file servers, print servers, and terminal servers in medium and large organizations. You can also use the guidelines in this book to maximize the availability and scalability of your servers by planning for remote server management, designing and deploying server clusters, and designing and deploying Network Load Balancing clusters.

This book contains the following chapters:


Planning for Storage

Designing and Deploying File Servers

Designing and Deploying Print Servers

Hosting Applications with Terminal Server

Planning for Remote Server Management

Planning for High Availability and Scalability

Designing and Deploying Server Clusters

Designing Network Load Balancing

Deploying Network Load Balancing

# Planning for Storage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For many organizations today, data is their most important asset. Losing data can mean losing thousands or even millions of dollars in revenue. Organizations can help protect valuable data by ensuring adequate, available, and secure storage. To manage and secure data, you can use a number of technologies with the Microsoft® Windows® Server 2003 operating system, including direct-attached storage, storage area networks (SANs), network-attached storage, Windows server clusters, Volume Shadow Copy service, Distributed File System (DFS), and Automated System Recovery (ASR).

**In This Chapter**


Overview of Storage Planning

Determining Organizational and Application Requirements

Reviewing Storage Architectures

Planning for Storage Scalability

Planning for SAN Deployment

Planning to Protect User Data

Planning for Fault Tolerance

Planning for Disaster Recovery

Additional Resources for Planning for Storage


**Related Information**

- For information about file servers, Distributed File System (DFS), File Replication service (FRS), Windows server clusters, NTFS file system permissions, disk quotas, shadow copies, and redundant array of independent disks (RAID) levels, see "Designing and Deploying File Servers" in this book.

- For information about high availability planning, see "Planning for High Availability and Scalability" in this book.

- For information about backup and recovery, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Overview of Storage Planning

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Several factors have contributed to the growth of both online and offline organizational storage requirements. Among these factors are:

- Increasing amounts of data
- Increasing needs for high availability and fast recovery
- Declining costs of storage

These developments have caused many organizations to place increased priority on storage planning and design. Planning and designing a storage solution involves defining the specific requirements of your organization and carefully evaluating how those requirements can work with or benefit from operating system services and features. Windows Server 2003 includes support for critical storage technologies such as direct-attached storage, network-attached storage, and SANs, as well as features such as Automated System Recovery (ASR), shadow copies, and open file backup that should be considered as you plan to meet your organization's storage requirements. If your storage plan includes using additional storage management or backup tools, you will also need to plan for integration of the operating system's storage features with these tools.

You can use the information in this chapter to assess your organization's specific storage needs, and consider the storage technologies and operating system features that are available to meet those needs. You can then follow the processes defined here to develop a storage plan that meets your needs for scalability, security, availability, and recoverability. While this chapter does discuss how to plan for and work with various storage architectures and technologies, you should already have a working knowledge of each of these technologies.

# The Storage Planning Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A properly implemented storage infrastructure will meet your organization's needs if it is sufficiently scalable, protects valuable data, is fault tolerant, and is recoverable in the event of disaster. With this in mind, you need to address both your organizational needs and specific methods for achieving these requirements when planning for storage. Figure 1.1 shows the steps involved in planning for storage.

**Figure 1.1 Planning for Storage**



# New Storage Features in Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 provides many new features that aid in the planning and deployment of storage within your organization. Table 1.1 describes these features.

**Table 1.1 Windows Server 2003 Storage Improvements**

| Feature | Description |
|---|---|
| Automated System Recovery (ASR) | ASR allows you to initiate *bare metal restores* of a system. With bare metal restores, you can boot a failed system by using the installation CD, insert an ASR floppy disk (created during an ASR backup), and recover the system state and required operating system files on the boot and system volumes from a backup. ASR can be combined with Remote Installation Services (RIS) to automate the recovery of several systems on your network. |
| Volume Shadow Copy service | The Volume Shadow Copy service provides an infrastructure for creating a volume that is a point-in-time image of the original volume. This image is known in the storage industry as a volume snapshot. The shadow copy can be used for the purpose of backing up the files within the volume; the copy is identical to the files at the instant the shadow copy was taken.<br><br>Solutions built on the Volume Shadow Copy service can produce much higher quality snapshots than other technologies because of the ability to integrate with business applications and coordinate with storage hardware. As a result, high-fidelity backup recovery and data mining are possible without significantly affecting performance. |
| Shadow copies | You can give users access to previous versions of files by enabling shadow copies, which provide point-in-time copies of files stored on servers running Windows Server 2003. Because users can restore or roll back their files to a previous state, you can reduce the administrative burden of restoring previously backed up files for users who accidentally delete or overwrite them. |
| Open file backup | By integrating with the Volume Shadow Copy service, the Backup program (NTBackup.exe) available in Windows Server 2003 can now back up open files that are locked by a user or application. Previously, files had to be closed in order to be successfully backed up. |
| Distributed File System (DFS) | Servers running Microsoft® Windows® Server 2003, Enterprise Edition and Windows® Server 2003, Datacenter Edition support multiple DFS roots on a single server. Other new DFS features include the ability to configure DFS to choose an alternate target based on cost if no same-site targets are available. |
| File Replication service (FRS) | You can use the Distributed File System snap-in to configure the FRS replication topologies (ring, hub and spoke, and custom), and configure connection priorities. FRS suppresses excessive replication and provides better staging directory management. |
| SAN support | Windows Server 2003 provides many new features that make it easier to work with data and devices on a SAN. These features include improved support for booting a system from a storage device on a SAN, greater control of mounting and unmounting volumes on the SAN, including the ability to suppress the automatic mounting of volumes, and improved handling of Fibre Channel host bus adapters (HBAs) for easy SAN integration and interoperability. |
| Virtual Disk service (VDS) | VDS provides a standard set of application programming interfaces (APIs) to storage hardware and storage management programs that makes it easy for storage applications to work with storage hardware that includes a VDS provider. The new command-line tool DiskRaid (which is available in the *Microsoft® Windows® Server 2003 Deployment Kit* ) allows you to manage hardware devices that have VDS hardware providers without the need for other management applications. |

| | |
|---|---|
| Multipathing | Multipathing provides for high availability data access by allowing a host to have up to 32 paths to access an external storage device, which facilitates failover and load balancing. Multipathing is not a feature of the operating system, but is supported through the MPIO Driver Development Kit (DDK), which provides a means for storage vendors to create interoperable multipathing solutions. |
| Disk Defragmenter | Disk Defragmenter performs faster and is more efficient than in the Microsoft® Windows® 2000 operating system, also now supports online defragmentation of the master file table (MFT), and can defragment NTFS volumes with any cluster size. |

# Determining Organizational and Application Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Determining your organization's storage requirements not only involves considering your company's business needs, but also any requirements related to your line of business applications. For example, a client application might need to access data on file servers by using a single UNC path, but your storage plan calls for providing fault tolerance by storing the data on two servers. The requirements of this client application could be met by implementing DFS. The first step in planning for storage, consequently, involves determining both the needs of your organization and the needs of your applications. Collectively, they will define your storage requirements. This process is shown in Figure 1.2.

**Figure 1.2 Determining Organizational and Application Requirements**



# Determining Organizational Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Organizational storage requirements go beyond having sufficient capacity for the data stored on each server and workstation. You must also satisfy redundancy and replication requirements. For example, to provide fault tolerance, you might decide to replicate data between two different file servers at two sites. In this case, you will need to ensure that there are sufficient storage resources at each site to store the replicated data.

To determine your organizational requirements, it will be helpful to answer the following questions.

**What are your current priorities?** You should focus on prioritizing which systems to address in what order. For example, mission-critical servers are more important than individual workstations. Prioritizing systems will help to define where you should begin when defining storage requirements, and is also important when budget concerns limit the amount of storage that you can immediately upgrade. When prioritizing, focus on maximizing return on investment and distributing the benefit of improvements to the largest possible number of individuals in your organization.

**What is the expected growth of data?** Do not plan only for your current volume of data. Your plan should also include estimates of the size and expected growth rate of your organization's data.

**What are your budget constraints?** Budget is always a limiting factor. Understanding your budget constraints reduces wasted design efforts and helps focus on improvements that can actually be delivered. You might consider a solution that technically addresses all of your needs, but you will need to find a compromise or an alternative if the necessary funds are not available.

**Can you consolidate storage?** Consolidating storage reduces the cost of managing storage on multiple servers and increases the efficiency of storage allocation and backup tasks. However, consolidation also introduces risk, because more users depend on fewer servers; so an outage affects more users than in a more distributed installation. If you plan to consolidate storage, you should also investigate ways to increase storage availability.

**Will training be needed?** It might be necessary for your administrators to get additional training. For example, administrators might be comfortable with planning simple clusters, but they should not attempt state-of-the-art Fibre Channel solutions without specific training. Make plans to build necessary training into your project and budget if you plan to implement storage solutions with a high degree of complexity, or with complex management interfaces.

**Will your storage needs change?** Your storage needs are probably growing. It is easier to manage a flat growth rate than an explosive one, but in either case it is best to build flexibility into your storage plan at the outset. Implementing flexible Windows storage features now lets you adapt your plan later with little or no disruption to users. For example, DFS allows you to add servers transparently to your environment, and SAN-based storage arrays and network-attached storage appliances let you achieve higher storage utilization than direct-attached storage. Look for any upcoming projects or other business changes that might require additional storage.

**How many copies of data are maintained?** You need to plan adequate storage not only for primary copies of data, but also for any additional copies of the data that you need. Be sure to plan for the appropriate amount of backup media, and for sufficient online disk storage for shadow copies and replication.

**What are your performance requirements?** Your choice of storage solutions is normally influenced by performance considerations. For example, your local area network (LAN) and wide area network (WAN) network components will likely need to be considered when planning for replication or data recovery between two systems in the same site, or between two or more systems in remote sites. Other considerations regarding performance include the type of server disks used, such as Enhanced Integrated Drive Electronics (EIDE) or Small Computer System Interface (SCSI), storage interfaces, and the speed of backup storage devices, such as tape drives.

**What are your backup scheduling requirements?** You might have a backup window that must be met for nightly backups, or you might be able to run backups over their own network during the day. In addition to impacting the network, even backups made to direct-attached storage devices consume CPU cycles. Your backup plan should ensure that servers are available for clients and applications when they are needed most, and that backup activity does not impact network traffic during peak usage periods.

**How important is ease of backup and recovery?** You might have the resources to manually back up your data nightly, and, in case of a system disaster, you might be able to bring in additional staff on an overtime basis to recover the information. Alternatively, you might require an automated approach that relies on high-end enterprise level backup software. Your plan for backup and recovery of data must balance your need for automation against the increased costs of various automated solutions.

**How long can data recovery take?** Some critical systems must be continuously available. For example, an online banking enterprise might require that its system experience zero downtime. Other organizations might be able to tolerate a one-hour window of downtime. Your plan needs to take into account your organization's specific tolerance for downtime during data recovery.

**Will you use legacy equipment?** Budget considerations sometimes dictate that you not replace existing equipment as you build your new storage environment. This is especially true for expensive storage libraries. The presence of older servers, disk arrays, cabling, and other equipment might save your organization money, but it might also dictate compromises in your storage plan. Also, if your plan calls for keeping existing technology, you might need to budget for any additional hardware needed to allow you to use the technology in your new storage environment.

# Determining Application Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When determining storage requirements, start at the application and operating system level, noting the storage and disk requirements for each system. For example, you might decide that your Microsoft® Exchange Server requires 500 gigabytes (GB) of total storage, but you also need fault-tolerance. To configure a 500 GB RAID 0+1 disk array to provide fault-tolerance, you would need 1 terabyte (TB) of physical disk storage. You would also need to allocate drives for the Exchange transaction log files.

Taking this approach server-by-server and application-by-application, you can determine the number of physical disks that you require, as well as the amount of total storage you need. When you estimate your total storage requirements, do not forget to include backup requirements. For example, you might plan to back up your Exchange server to a magnetic storage array for very fast backup and recovery performance, and then copy the backup data to tape for offsite storage. To keep a single full backup onsite would require that you double your storage estimate for the Exchange server. You will also need to estimate the growth rate of each database so that your storage meets both your immediate and future needs for a period of time as defined by your organization.

When considering storage requirements for the operating system, you should note which servers require disk mirroring to protect their boot and system volumes, and note the number of required disks along with the amount of required storage for each.

Estimating application storage requirements can be difficult without in-depth knowledge of all the applications that run in an organization. If your organization has designated specialists for your applications, consult with them to arrive at a proper estimate. The documentation for each application can also serve as an excellent source for compiling storage planning information

# Reviewing Storage Architectures

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you plan a new storage configuration or reengineer your existing configuration, you will need to make decisions about the types of storage architecture to use. These architectures are direct-attached storage, network-attached storage, and storage area networks. The different storage architectures are not necessarily exclusive, but are actually complementary and often coexist on the same network. Because these architectures are well known, the discussion that follows includes only a brief overview of each architecture, and focuses on the planning considerations related to using each architecture with Windows Server 2003. The process for reviewing the storage architectures is shown in Figure 1.3.

**Figure 1.3 Reviewing Storage Architectures**




**Direct-Attached Storage**

Direct-attached storage refers to a storage device, such as a hard drive or tape drive, that is directly connected to a single computer. These connections are usually made by one of the following methods:

- Enhanced Integrated Drive Electronics (EIDE)
- Small Computer Systems Interface (SCSI)
- Fibre Channel

EIDE connects internal Advanced Technology Attachment (ATA) storage to a computer, SCSI provides a means to connect both internal and external storage to a computer, and Fibre Channel connects external storage to a computer. Fibre Channel is most often used with external storage in a SAN. Although Fibre Channel can be used for direct-attached storage, less expensive SCSI storage can offer similar performance, but works only over limited distances due to the physical limitations of the SCSI bus. When external direct-attached storage devices are located more than twelve meters away from a server, then Fibre Channel must be used.

Direct-attached storage retains its high popularity because of its low entry cost and ease of deployment. The simple learning curve associated with direct-attached storage technologies is also a factor many organizations consider. Direct-attached storage also makes it easy to logically and physically isolate data, because the data can only be directly accessed through a single server.

Although it is simple to deploy, there are other management considerations to take into account with direct-attached storage:

- Direct-attached storage can be more expensive to manage because you cannot redeploy unused capacity, which results in underutilization.

- Having storage distributed throughout the organization makes it difficult to get a consolidated view of storage across the organization.

- Disaster recovery scenarios are limited because a disaster will cause both server and storage outages.

- For data backup and recovery, you need to choose whether to attach local backup devices to each server, install dual network adapters in each server and back up the data over a separate LAN, or back up the server over the corporate LAN. Large organizations have found that placing stand-alone tape drives in individual servers can quickly become expensive and difficult to manage, especially when the number of servers in the organization grows into the hundreds. In this situation, it is often best to back up servers over a network to a storage library, which offers backup consolidation and eases management.

**Network-Attached Storage**

Network-attached storage solutions are ideal for both addressing immediate storage needs and storage consolidation projects. Their high capacity and ease of deployment allow you to implement network-attached storage with little planning. With network-attached storage, you can fulfill your storage requirements with a single purchase. This is because many network-attached storage devices are built to be *appliances*. Like home appliances, such as a refrigerator or microwave oven, you can literally purchase a network-attached storage appliance and plug it into your network. This is possible because network-attached storage appliances come prepackaged with network interfaces, storage interfaces (SCSI, Fibre Channel, or both), internal or external magnetic storage, and an operating system. Only the services required for file serving, security, and management are installed on the appliance, thus making it a high-performance file server. A network-attached storage implementation is shown in Figure 1.4.

**Figure 1.4 Network-Attached Storage Implementation**




Network-attached storage has risen in popularity primarily due to the quick-and-easy nature of its implementation: Many Windows Powered Network Attached Storage appliances can be installed and running within minutes, and easily integrate into any enterprise environment.

The operating systems on Windows Powered Network Attached Storage appliances are either Windows 2000 or Windows Server 2003. Windows Powered Network Attached Storage appliances offer the file system, security, reliability, and scalability features common to all Windows operating systems, and come preconfigured with network file system (NFS), File Transfer Protocol (FTP), AppleTalk, Hypertext Transfer Protocol (HTTP), Web Distributed Authoring and Versioning (WebDAV), and NetWare protocol support. Because Windows Powered Network Attached Storage appliances support other operating system environments such as UNIX and Linux, you can seamlessly integrate them into any existing computer infrastructure, whether it uses Windows or not. They also support standard Web protocols, which allows you to manage Windows Powered Network Attached Storage appliances from a Web browser.

Windows Powered Network Attached Storage appliances are well suited for use as both departmental file servers and enterprise file servers. Network-attached storage solutions built on Windows offer data protection and high availability features like clustering and snapshots. In addition, network-attached storage appliances offer a highly scalable solution that can grow with the business. For example, storage can be added on the fly to address growing business needs.

To summarize, the key benefits of Windows Powered Network Attached Storage are:

- **Ease of deployment.** Windows Powered Network Attached Storage appliances can be installed in a heterogeneous environment within minutes.

- **Active Directory integration.** Windows Powered Network Attached Storage appliances can be integrated with the Active Directory® directory service to take advantage of features such as Kerberos authentication, Internet Protocol security (IPSec), Encrypting File System (EFS), and Group Policy objects (GPOs).

- **ISV utility support.** You can install several independent software vendor (ISV) utilities on Windows Powered Network Attached Storage appliances, including quota, backup, antivirus, and replication utilities. This is not possible with many

other network-attached storage appliances.

- **Simple management.** Remote management is supported through Terminal Services sessions as well as through a Web interface. Administrators do not have to learn a new operating system to operate a network-attached storage appliance, because it uses Windows 2000 or Windows Server 2003.

- **Enhanced snapshot support.** Windows Powered Network Attached Storage appliances maximize the availability of data by supporting up to 250 snapshots, which provides an advantage compared to other network-attached storage solutions. *Snapshots* are similar to the shadow copy feature in Windows Server 2003.

- **DFS and FRS integration.** You can use DFS to build and manage a single, hierarchical view of multiple file servers and their shares. To provide redundancy and load balancing, you can store shares with identical data on multiple servers, and DFS will automatically redirect clients to alternate servers if a server becomes unavailable. To keep these shares synchronized, you can use FRS.

For more information about Windows Powered Network Attached Storage appliances, see the Windows Powered Network Attached Storage link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Storage Area Networks**

Unlike direct-attached or network-attached storage, SANs require careful planning before implementation. SANs differ from network-attached storage in that they store and access data at the block level, whereas network-attached storage appliances store data at the file level. File level access is preferable for users and applications needing to access a particular file, whereas block level data access is better for applications that need to quickly access data. SANs also differ from other storage architectures in that they are dedicated storage networks, and use their own network protocols and hardware components.

A simple SAN configuration is illustrated in Figure 1.5.

**Figure 1.5 SAN Implementation**




For more information about specific considerations for deploying SANs with Windows Server 2003 see "Planning for SAN Deployment" later in this chapter.

# Planning for Storage Scalability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A properly planned storage infrastructure should meet both your current and future needs. It is crucial to understand the storage design and implementation limitations of your existing systems, as well as methods to overcome those limits. The general process for storage scalability planning is shown in Figure 1.6.

**Figure 1.6 Planning for Storage Scalability**



# Reviewing Storage Limits

Updated: August 3, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2008 R2 with SP1

To accurately plan for storage scalability, you must take into account the physical and logical design limits of your servers. The following sections describe these limits.

**Reviewing the physical storage limits**

From a hardware perspective, your choice of storage bus will determine the physical limits of your storage scalability. For example, local storage with two Integrated Device Electronics (IDE) controllers will allow a system to have up to four IDE storage devices; with SCSI storage, the SCSI bus type will determine the maximum amount of allowable storage devices on each bus. If you need to extend a system's local storage beyond the limitations of IDE or SCSI, you can use Fibre Channel–based storage or use SCSI protocol over TCP/IP (iSCSI), a new IP-based storage networking standard that carries SCSI commands over IP networks. However, these newer, more flexible architectures are more expensive than traditional IDE or SCSI buses.

| **Important** |
|---|
| Unless otherwise noted, all storage limits discussed in the following sections apply to Windows Server® 2003, Windows Server® 2008, and Windows Server® 2008 R2 operating systems. |

**Reviewing the limits for basic and dynamic Disks**

In addition to reviewing the physical storage limits, you also need to keep in mind the volume types in Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 and their limits. All versions of these operating systems provide enhanced volume management capabilities. You can plan for both basic disks and volumes and dynamic disks and volumes.

*Basic disks* use the same disk structures as those used in Microsoft® Windows NT® version 4.0 and Windows 2000. When you use basic disks, you are limited to creating four primary partitions per physical disk, or three primary partitions and one extended partition that can contain multiple logical drives. Primary partitions and logical drives on basic disks are known as *basic volumes*. As shown in Table 1.2, basic volumes are limited to 2 TB. Even if you create multiple volumes on a single logical unit, the combined size of all those volumes cannot exceed 2 TB. If you want to use volumes larger than 2 TB, you must use dynamic spanned, striped, or RAID-5 volumes.

Dynamic disks, which offer features not available in basic disks, were introduced in Windows 2000. Dynamic disks contain *dynamic volumes*, including simple volumes, spanned volumes, striped volumes, mirrored volumes, and RAID-5 volumes. Dynamic disks offer greater flexibility for volume management because they use a hidden database, instead of the disk's partition table, to track information about dynamic volumes on the disk and about other dynamic disks in the server. This flexibility allows you to create spanned, striped (RAID-0), and RAID-5 volumes that exceed the 2-TB size limit of basic volumes. Simple and mirrored volumes cannot exceed 2 TB.

Table 1.2 lists the contrasts between basic and dynamic Master Boot Record (MBR) disks.

**Table 1.2 Comparison of Basic and Dynamic MBR Disks**

| Storage Capabilities | Basic Disks | Dynamic Disks |
|---|---|---|
| Volume size (maximum) | 2 TB | 2 TB for simple and mirrored volumes.<br><br>Up to 64 TB for spanned and striped volumes. (2 TB per disk with a maximum of 32 disks per volume.)<br><br>Up to 62 TB for RAID-5 volumes. (2 TB per disk with a maximum of 32 disks per volume and 2 TB used for parity.) |
| Supported RAID implementation | Hardware RAID only | Hardware or software RAID |
| Storage of boot and system volumes | Yes | Simple or mirrored volumes only |

| | | |
|---|---|---|
| Shared cluster storage in server clusters | Supported | Not supported* |

* For more information about using dynamic disks on shared cluster storage, see article 237853, in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=99316).

Use the following guidelines to choose between basic disks and dynamic disks.

Use basic disks if:

- You do not need to create volumes that exceed 2 TB.
- You are configuring shared cluster storage on a server cluster.
- The server runs other operating systems that cannot access dynamic disks. The only server operating systems that can access dynamic disks are Microsoft® Windows® 2000 Server and Windows Server 2003.

Use dynamic disks if:

- You want to create RAID-0 volumes or fault-tolerant volumes (RAID-1 or RAID-5) and the server does not contain hardware RAID.
- You want to combine logical units (LUNs) in a RAID array to create a volume larger than 2 TB.
- You want to extend a volume, but the underlying hardware cannot dynamically increase the size of LUNs, or the hardware has reached its maximum LUN size.

**Reviewing the limits for NTFS and the number of volumes per server**
When you configure storage in Windows operating systems, it is important to understand NTFS limits and volume limits for server storage.

Table 1.3 describes the NTFS and volume limits.

**Table 1.3 NTFS and Volume Limits**

| Description | Limit |
|---|---|
| Maximum size of an NTFS volume | $2^{32}$ clusters minus 1 cluster<br><br>Using a 64-kilobyte (KB) cluster (the maximum NTFS cluster size), the maximum size of an NTFS volume is 256 TB minus 64 KB.<br><br>Using a 4-KB cluster (the default NTFS cluster size), the maximum size of an NTFS volume is 16 TB minus 4 KB. |
| Maximum number of dynamic volumes per disk group | 1,000<br><br>A disk group is collection of dynamic disks. Windows Server 2003 supports one disk group per server. |
| Maximum volumes per server | Approximately 2,000 volumes<br><br>Up to 1,000 of these volumes can be dynamic volumes; the rest are basic volumes. Boot times increase as you increase the number of volumes. In addition, the use of multipathing can reduce the number of volumes per server. |

**Reviewing the number of supported disk drives**
When you select the number of disk drives in Windows operating systems, it is important to understand the number of drives supported in different types of file systems. The number of disk drives you can use for each server is limited only by the available memory for FAT16, FAT32, NTFS, and UDF files systems.

**Reviewing methods for adding storage space to volumes**
If you create volumes smaller than the maximum sizes listed earlier, and you later want to add storage space to those volumes, you can do so by increasing the size of the volume or by creating a mounted drive.

**Extending volumes**
When a defined volume begins to reach its allocated capacity, but there is still free unallocated space on a hard disk or hardware RAID array, you can use the DiskPart command-line tool (Diskpart.exe), available in the Windows Server 2003 operating system, to extend the volume size without any loss of data. When using DiskPart, consider the following guidelines:

- You cannot use DiskPart to extend system or boot volumes; DiskPart can only be used to extend data volumes. This is not applicable for Windows Server 2008.
- The volume to be extended must be formatted with the NTFS file system.
- For basic volumes, the unallocated disk space used for the extension must be the next contiguous space on the same disk. (The disk can be a single physical disk or a group of physical disks presented by a RAID adapter to the operating system as a single disk, often referred to as a LUN.)
- For dynamic volumes, the unallocated disk space can be any empty space on any dynamic disk on the system.
- If you upgraded from Windows 2000 to Windows Server 2003, you cannot extend a simple or spanned volume that was originally created as a basic volume and converted to a dynamic volume on Windows 2000. For more information about basic and dynamic volumes, see Storage Technologies Collection (http://go.microsoft.com/fwlink/?linkid=4681).

For more information about using DiskPart, see "DiskPart" in Help and Support Center for Windows Server 2003. For more information about choosing volume sizes based on Chkdsk and recovery times, see "Designing and Deploying File Servers" in this book.

**Using mounted drives**
Mounted drives are useful when you want to add more storage to an existing volume without having to extend the volume. A mounted drive is a local volume attached to an empty folder on an NTFS volume. Mounted drives are not subject to the 26-drive limit imposed by drive letters, so you can use mounted drives to access more than 26 drives on your computer. For more information about mounted drives, including information about creating mounted drives on server clusters, see "Using NTFS mounted drives" in Help and Support Center for Windows Server 2003.

# Planning for Storage Virtualization by Using DFS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Storage virtualization is a means to organize storage virtually instead of physically. With virtualized storage, the user sees all available storage, but is insulated from the actual physical structure and processes of storage, such as the configuration, location, and capacity of the individual devices that make up the storage system. At the network level, you can achieve virtualization by using DFS. DFS allows you to move data between servers, load balance data access among servers, and change or expand your storage configuration without impacting users.

Even if your network incorporates only a single file share, you should still build your shared storage infrastructure to use DFS. This way, as your storage needs expand, the infrastructure is already in place to support that growth. With the transparency that DFS provides, you can configure and reconfigure storage at any time with no downtime or interruption to users or applications.

The benefits of using DFS include:

**Unified namespace** A DFS namespace links together shared folders on different servers to create a hierarchical structure that behaves like a single high-capacity hard disk. Users can navigate the logical namespace without having to know the physical server names or shared folders that host the data.

**Location transparency** DFS simplifies migrating data from one file server to another. Because users do not need to know the name of each physical server or shared folder that contains the data, you can physically move files to another server without having to reconfigure applications and shortcuts, and without having to reeducate users about where they can find their data.

**Storage scalability** You can deploy additional or higher-performance file servers and present the storage on the new servers as new folders within an existing namespace.

**Increased availability of file server data** When multiple servers running Windows Server 2003 host a domain-based DFS root, clients are redirected to the next available root server if any of these servers fail, which provides fault-tolerant data access. To ensure the availability of stand-alone DFS namespaces, you can create the root on a clustered file server.

**Load sharing** DFS provides a degree of load sharing by mapping a single logical name to shared folders on multiple file servers. For example, suppose that \\Company\StockInfo is a heavily used shared folder. By using DFS, you can associate this location with multiple shared folders on different servers, even if the servers are located in different sites.

**Simplified maintenance** Using multiple copies of shared folders also allows administrators to transparently perform preventive maintenance or upgrades on servers. A server that hosts one target can be taken offline without affecting users, because DFS automatically routes requests to a target that is online.

For more information about DFS, see "Designing and Deploying File Servers" in this book.

# Planning for SAN Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organizational and application requirements can best be satisfied by using a SAN, you need to follow the steps in this section to adequately plan for your SAN implementation. Several new or enhanced storage features are included with Windows Server 2003 to better support storage devices located on a SAN. Planning for optimal SAN deployment in a Windows Server 2003 network includes careful consideration of these SAN-related features so that you can seamlessly integrate your SAN with your Windows network. You can then plan how to configure Windows for your specific SAN implementation, and how to integrate additional storage technologies from other vendors. The process of planning for SAN deployment in a Windows network is shown in Figure 1.7.

**Figure 1.7 Planning for SAN Deployment**



# Planning to Use Windows Storage Features with SANs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you plan your SAN implementation, consider which operating system features you can use to optimize the storage located on your SAN. Windows Server 2003 adds several services, components, and capabilities to better support storage devices located on a SAN. These include:

- Virtual Disk service (VDS)
- Volume Shadow Copy service
- Volume Automount
- Multipath Input/Output (MPIO)

**Virtual Disk Service**

Until recently, storage applications managed storage by directly interfacing with storage devices, based on the hardware vendor's device-specific applications. In Windows Server 2003, the Virtual Disk service (VDS) provides a standard set of APIs that storage software can use to communicate with VDS, which then sends the necessary queries or commands to the storage hardware. VDS simplifies storage management by providing a single management interface for multivendor storage devices, whether in direct-attached storage or on SANs. By using VDS, you can configure a mixed storage environment without needing to use the individual management tools supplied by each hardware vendor. To support VDS, hardware vendors must supply a VDS *provider* for the storage hardware. The provider translates the VDS APIs into instructions specific to the hardware.

VDS insulates applications from the complexities associated with storage implementations, making query and configuration operations common across all managed devices. The specific scope of VDS functionality depends on the capabilities of the underlying hardware.

Windows Server 2003 components that work with VDS as management applications include the Disk Management snap-in, the DiskPart command-line tool, and the DiskRaid command-line tool (Diskraid.exe), which is available in the *Windows Server 2003 Deployment Kit*. DiskRaid requires at least one VDS provider supplied by a storage vendor. For more information about DiskRaid, click **Tools** in Help and Support Center for Windows Server 2003, and then click Windows Resource Kit Tools Help.

Because of the advantages of using VDS, you should consider making VDS support a criterion when evaluating hardware or software for use in your storage infrastructure. Devices with VDS providers will be compatible with Windows Server 2003 and with other storage applications that are written for VDS. If your plan calls for storage devices that are not VDS compliant, check to ensure that the devices are compatible with Windows Server 2003 and are supported by your storage applications. For more information about device compatibility with Windows Server 2003, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources

**Volume Shadow Copy Service**

The Volume Shadow Copy service adds new features for direct-attached storage and external storage arrays, including arrays that are SAN-based. It provides an infrastructure that makes possible high-fidelity backup, recovery, and data mining without significantly affecting performance.

The Volume Shadow Copy service is not a user-accessible feature, but an operating system service that works by creating a volume that can be accessed by other applications. This volume appears as a read-only point-in-time copy of the original volume. By default, the shadow copy does not appear to the server as an additional volume, but other applications can mount it and read from it. Following the creation of a shadow copy, the original volume continues to be used by the system in the same way it was used before the shadow copy was created.

The Volume Shadow Copy service can be extended with providers supplied by storage vendors, so that backup and data management applications can leverage the advanced capabilities of different storage arrays without having to deal directly with the complexities of a particular array. This allows storage management applications to more easily support a number of different hardware devices.

The Volume Shadow Copy service also enables additional capabilities in Windows Server 2003, including the following:

- **Open-file backups.** Backup applications that work with Volume Shadow Copy service, such as the Backup program in Windows Server 2003, can back up open files. For more information about backing up open files, see "Planning to Back Up Open Files" later in this chapter.

- **Shadow copies.** After an administrator has configured and enabled shadow copies, users themselves can restore earlier versions of files from point-in-time copies stored on file servers that are running Windows Server 2003. For more information about shadow copies, see "Protecting User Data by Using Shadow Copies" later in this chapter.

**Volume Automount**

By default, volumes automatically mount when discovered in Microsoft® Windows® Server 2003, Standard Edition and Windows® Server 2003, Web Edition operating systems. In the Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition operating systems, automount is suppressed by default to prevent these servers from aggressively mounting all volumes attached to the SAN. Automount can be enabled or disabled by using the DiskPart command-line tool. When volumes are not automounted, you can mount them and assign drive letters to them by using DiskPart.

For more information about using DiskPart, see "DiskPart" in Help and Support Center for Windows Server 2003.

**Multipath Input/Output**

Multipathing increases availability by providing multiple paths from the host to the external storage device. The Multipath Input/Output (MPIO) Driver Development Kit (DDK) allows storage vendors to create interoperable multipathing solutions that work with Windows Server 2003. Up to 32 paths to each logical unit are supported. For devices that can support multiple active paths, a vendor can improve performance by implementing load balancing.

To take advantage of MPIO on your SAN, you must ensure that your storage uses multipath drivers that support Microsoft MPIO technology. For more information about compatibility with Windows Server 2003, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Configuring Windows Server 2003 for SANs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you create your plan for deploying servers running Windows Server 2003 that use SAN-based storage, you need to resolve several configuration issues. In particular, consider the following:

- Is the SAN topology compatible with Windows Server 2003?
- Can the SAN storage be configured to be shared between server cluster nodes?
- Can Windows Server 2003 boot from the storage connected to the SAN?
- How will zoning and LUN masking be configured?

The sections that follow assume that you understand the basic concepts of SANs, including Fibre Channel SAN topologies and hardware components.

# SAN Configurations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your SAN configuration directly affects the speed and reliability with which your systems can access storage devices on the SAN. Two critical considerations are your choice of a Fibre Channel SAN topology and the use of redundant components in the access path to storage devices.

**Supported Fibre Channel Topologies**
Windows Server 2003 supports all Fibre Channel SAN topologies (point-to-point, loop, and fabric). For attaching single storage devices to Windows Server 2003 systems using a Fibre Channel host bus adapter (HBA), the point-to-point topology is ideal. There are, however, some limitations with respect to the loop (FC-AL) topology. Loop topologies are inherently limited in operation and scalability, because they often require full loop reinitialization when devices are added or removed from the SAN. This can cause data to be inaccessible for periods ranging up to minutes. Also, loops can only scale to support up to 127 devices, and, like token ring networks, only allow a single device on the SAN to transmit data at one time, which might not make full use of the available bandwidth on the SAN.

Windows Server 2003 supports FC-AL topologies when a single server is connected to storage. However, the only multiserver configuration supported on a loop is a single two-node cluster. This limitation also applies to the derivatives of FC-AL topologies: Fibre Channel Private Loop SCSI Direct Attachment (FC-PLDA), Fibre Channel Fabric Loop Attachment (FC-FLA), and Fibre Channel Private Loop Attachment (FC-PLA). This means that your SAN design should incorporate a true fabric, or switch-based fabric topology, using a fabric SAN switch (hubs or loop switches should not be used) for any SAN implementation that contains more than two components (servers or storage devices).

**Important**

- Placing new devices on an FC-AL–based SAN causes Loop Initialization Primitives (LIPs), which appear to the operating system as storage bus resets. LIPs can result in data unavailability and even data loss, if a LIP interrupts a write operation that is only partially committed to the storage device. LIPs can also disrupt backup operations to tape devices and can cause your devices to be assigned new addresses. Therefore, you should use persistent binding. You can limit the propagation of LIPs by using only fabric-based SAN topologies and compatible devices.

For more information about supported Fibre Channel topologies, see article 317162, "Supported Fibre Channel Configurations." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Redundant Components**
Although storage data on a SAN offers many benefits, it also introduces additional components whose failure can result in loss of access to data. As on a LAN, the failure of an adapter, switch, or cable can interrupt or remove access to a storage device. The best means to provide for highly available storage is to implement a SAN that uses redundant storage (such as hardware RAID), redundant data access paths (including cables and switches) and server HBAs, and MPIO storage drivers.

# Server Cluster Integration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For server clusters of three or more nodes, each node must connect to shared storage over a fabric-based SAN. When planning an SAN storage solution that uses server clusters, ensure that your plan meets the following requirements:

- Each cluster on a SAN should be deployed in its own zone if possible. The mechanism the cluster uses to protect access to the disks can adversely affect other clusters or servers in the same zone. You can prevent this by using zoning to separate the SAN cluster traffic from other traffic.

- All HBAs in a single cluster must be the same type and have the same firmware, driver versions, and driver parameters. Many storage and switch vendors require that all HBAs on the same zone (and, in some cases, on the same fabric) share these features.

- All other storage drivers in the same cluster must have the same software versions.

- Only one node on a cluster should access the disks that are to be managed by the cluster when you first create the cluster, otherwise resource contention might cause the installation to fail.

  This is a requirement because, when the volume automount option is enabled, Windows Server 2003 mounts any volumes it can detect when the operating system boots (this is disabled by default on Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition).

  **Caution**

  - If multiple nodes that are not in the same cluster access the same disk, data corruption results.

- Tape devices must not be located in the same zone as cluster disk storage devices unless Storport miniport drivers are used. A tape device will interpret a bus reset as an instruction to rewind and possibly dismount the tape, which will disrupt any backup operation. If this occurs, the backup must be restarted.

For a white paper that provides more information about using SANs with server clusters, see the Windows Clustering: Storage Area Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresourceson the Web Resources page at http://www.microsoft.com/windows/reskits/webresources

# SAN Boot Considerations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your plan requires servers attached to a SAN to boot from disks located on the SAN, you must verify with the disk vendor that the disks can be used to boot a Windows server. When a server's system and boot volumes are stored on a SAN, the SAN vendor must ensure that the disks are accessible during system startup. Make sure that the vendor provides specific configuration information, including needed firmware revisions and setup instructions. You will also need to ensure that each server's HBA is configured according to the disk vendor's guidelines for booting from a SAN.

Additional requirements for successfully booting servers from a SAN are:

- Proper fiber channel topology is used.

  The SAN should be configured in a fabric topology, or the host can be directly attached to the Fibre Channel storage device. For more information about Fibre Channel topologies, see "SAN Configurations" earlier in this chapter.

- Multipathing and redundant SAN links are used.

  You can avoid single points of failure in your SAN design by implementing redundant paths to the critical storage on the SAN, redundant HBAs, and MPIO-compliant drivers.

  Finally, use of the new Storport port driver and appropriate miniports from your HBA vendor will allow more flexibility in configuring your SAN. With Storport miniports, it is possible to have the boot disk and cluster disks on the same Fibre Channel connection. For more information about this configuration, see the Windows Clustering: Storage Area Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- The boot disk LUN is dedicated to a single host.

  A server that boots from the SAN must have exclusive access to the disk that it is booting from. Exclusively associating a storage device with a server can be accomplished either through zoning or LUN masking, which is configured at the SAN switch, HBA, or storage subsystem level. Storage-based LUN masking is the best choice when you cannot use switch zoning to isolate systems. For more information about LUN masking and zoning see "LUN Masking vs. Zoning" later in this chapter.

For more information about booting from SANs, see article 305547, "Support for Booting from a Storage Area Network (SAN)." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# LUN Masking vs. Zoning

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The potential for any server on a SAN to mount and access any drive on the SAN can create several problems. The two most prominent ones are disk resource contention and data corruption. To deal with these problems, you can isolate and protect storage devices on a SAN by using zoning and LUN masking, which allow you to dedicate storage devices on the SAN to individual servers.

**Zoning**

Many devices and nodes can be attached to a SAN. When data is stored in a single *cloud*, or storage entity, it is important to control which hosts have access to specific devices. Zoning controls access from one node to another. Zoning lets you isolate a single server to a group of storage devices or a single storage device, or associate a grouping of multiple servers with one or more storage devices, as might be needed in a server cluster deployment.

Zoning is implemented at the hardware level (by using the capabilities of Fibre Channel switches) and can usually be done either on a port basis (hard zoning) or on a *World-Wide Name* (WWN) basis (soft zoning). WWNs are 64-bit identifiers for devices or ports. All devices with multiple ports have WWNs for each port, which provides more detailed management. Because of their length, WWNs are expressed in hexadecimal numbers, similarly to MAC addresses on network adapters. Zoning is configured on a per-target and initiator basis. Consequently, if you need to attach multiple nonclustered nodes to the same storage port, you must also use LUN masking.

**LUN masking**

LUN masking, performed at the storage controller level, allows you to define relationships between LUNs and individual servers. Storage controllers usually provide the means for creating LUN-level access controls that allow access to a given LUN by one or more hosts. By providing this access control at the storage controller, the controller itself enforces access policies to the devices. LUN masking provides more detailed security than zoning, because LUNs provide a means for sharing storage at the port level.

When properly implemented, LUN masking fully isolates servers and storage from events such as resets. This is critical for preventing the problems previously noted. It is important to thoroughly test your design and implementation of LUN masking, especially if you use LUN masking in server clusters.

# Integrating Vendor Features

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is likely that your SAN will include applications and services from various vendors. For optimal performance and reliability, you should ensure that all devices and applications on the SAN are fully compatible with Windows Server 2003. Use the following guidelines for all applications on your SAN:

- Device drivers for SAN components should conform to Plug and Play specifications. Failure to use Plug and Play drivers can unexpectedly prevent the operating system from recognizing a device and installing the correct device drivers, which can lead to data loss.

- Drivers for all devices on the SAN, such as HBA and multipathing drivers, should have a digital signature from Microsoft or the storage vendor. Device drivers that contain a digital signature assure you that the product was tested for compatibility with Windows, comes from the specified source, and has not been altered since testing. If the driver does not contain a digital signature from Microsoft or the storage vendor, contact the vendor directly and ask for a driver that meets the Designed for Microsoft Windows Server 2003 logo requirements.

  You can use the Windows Server Catalog to check for software and hardware compatibility with Windows Server 2003. To find the Windows Server Catalog, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Windows Server Catalog also provides information about additional compatibility requirements for software that qualifies for the Supported & Certified for Windows Server 2003 program and hardware that qualifies for the Designed for Windows Server 2003 program. Qualifying products are specifically created to take advantage of the new features in Windows Server 2003.

- Add VDS and Volume Shadow Copy service providers for storage devices that have them. This significantly improves the compatibility of the devices with Windows Server 2003 and gives applications from various vendors the ability to use the capabilities of the devices.

- Contact your storage vendor for a Microsoft MPIO-based multipath solution.

- Make sure that all device drivers are qualified by the storage vendor for use with the specific storage products from that vendor (disk, tape or software) that you will use in your SAN.

- If you are implementing Windows Clustering, it is extremely important that you select hardware configurations that have been approved by Microsoft for Windows server clusters. Microsoft supports only complete server clusters systems chosen from the Windows Server Catalog. For more information about approved hardware configurations, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about using SANs with server clusters, see "Designing and Deploying Server Clusters" in this book.

# Planning to Protect User Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 provides new features that allow you to automate protecting and recovering user data. Most notably, using shadow copies provides a high level of user data protection, and also frees administrators or help desk personnel from dealing with user requests to restore single files. You will also want to consider including folder redirection in your data protection plan. The process of planning to protect user data is shown in Figure 1.8.

**Figure 1.8 Planning to Protect User Data**



# Protecting User Data by Using Shadow Copies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Shadow copies increase the availability of data and reduce the administrative burden of restoring files. After an administrator has configured and enabled shadow copies, users themselves can restore earlier versions of files from point-in-time copies stored on general-purpose file servers that are running Windows Server 2003. Shadow copies are also available on server clusters running Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.

You can enable shadow copies by using Windows Explorer, the Disk Management snap-in, or Shared Folders extension to the Computer Management snap-in. You can also automate administrative tasks that involve making shadow copies by using the command-line tool Vssadmin.exe.

Shadow copies work by making a block-level copy of any changes that have occurred to files since the last shadow copy. Only the changes are copied, not the entire file. As a result, previous versions of files do not usually take up as much disk space as the current file, although the amount of disk space used for changes can vary depending on the application that changed the file. For example, some applications rewrite the entire file when a change is made, whereas other applications append changes to the existing file. If the application rewrites the entire file to disk, the shadow copy contains the entire file. Therefore, you should consider how the applications in your organization manage changes to files, as well as the frequency and number of updates to files, when you plan how much disk space to allocate for shadow copies.

Shadow copies are designed for volumes that store user data, and are best suited for protecting user data stored on servers from the following problems:

- Accidental deletion of files
- Accidental overwriting of files (for example, when a user forgets to use **Save As** to create a new version of a file)
- File corruption

Although shadow copies are made for an entire volume, users must use shared folders to access shadow copies. This is because shadow copies are designed to be accessed using the Common Internet File Sharing (CIFS) protocol. If you or your users want to access a previous version of a file that is not located in a shared folder, you must first share the folder. If you want to access shadow copies on a local server, you must use the *servername**sharename* path to access shadow copies.

The method of providing access to shadow copies from a given computer differs depending upon the operating system installed on the computer. Shadow copies can be accessed by computers running Windows Server 2003 and by computers running Microsoft® Windows® XP Professional on which you have installed the Previous Versions Client pack by running Twcli32.msi. This file is located in Windows Server 2003 in *windir*\system32\clients\twclient. You can install this file manually on clients or deploy the file by using the software distribution component of Group Policy. For more information about software distribution, see "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit.

To access shadow copies from previous versions of Windows, including Windows 2000 and Windows XP Professional, you can download and install the Shadow Copy Client. For more information and to download the Shadow Copy Client, see the Shadow Copy Client Download link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Note**

- The Previous Versions Client and the Shadow Copy Client provide the same functionality, but the Shadow Copy Client can be installed on multiple operating systems, such as Windows 2000 and Windows XP Professional, whereas the Previous Versions Client can only be installed on Windows XP Professional.

If you have not yet deployed these operating systems or client packs on your clients, you can deploy one or more computers from which users can restore previous versions of files. You can also distribute the client pack on a case-by-case basis to users who request that files be restored.

For more information about shadow copies, see the following sources:

- For more information about designing a shadow copy strategy, see "Designing and Deploying File Servers" in this book.

- For more information about using shadow copies on server clusters, see "Using Shadow Copies of Shared Folders in a server cluster" in Help and Support Center for Windows Server 2003.

- For more information about using Vssadmin.exe to create shadow copies, see "Vssadmin" in Help and Support Center for Windows Server 2003.

# Protecting User Data by Using Folder Redirection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If users store their data on their local workstations, you can protect it by using folder redirection. With folder redirection, user folders, such as the My Documents folder, can be redirected to a Windows Server 2003 file server. With all files located on a central server, you can configure shadow copies on the server so that users can recover earlier versions of their files, and you can back up their files in accordance with your organizational standards. You can easily implement folder redirection in your organization by using Group Policy.

Folder redirection offers the following advantages:

- Centralized user data for ease of backup
- Excellent integration and transparency with roaming user profiles
- Easily enforced disk quotas

To redirect special folders, in the Group Policy Object Editor snap-in to Microsoft Management Console (MMC), select the Group Policy object (GPO) that is linked to the site, domain, or organizational unit (OU) that contains the users whose folders you want to redirect. In the **User Configuration\Windows Settings\Folder Redirection** node of the Group Policy Object Editor, double-click **Folder Redirection**, right-click the special folder to redirect, and select **Properties**. Then you can configure redirection as either Basic or Advanced. With Basic redirection, the specified folders of all users are redirected to the same network share. With Advanced redirection, you can specify network share locations for redirection by user group.

If you have mobile users who travel with portable computers, working exclusively with folder redirection might not be feasible. To provide the benefits of folder redirection to mobile users while still making their data available to them when they travel, you can combine Offline Files with folder redirection. Offline Files allows users to download copies of their files from a server when they are connected to the network. If they modify files when they are not connected to the network, those files can be resynchronized with the server the next time the system is connected to the network. You can configure Offline Files options in a Group Policy object in the **Computer Configuration\Administrative Templates\Network\Offline Files** node of the Group Policy Object Editor.

For more information about redirecting folders, see "Implementing User State Management" in *Designing a Managed Environment* of this kit.

# Planning for Fault Tolerance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Organizations are increasingly finding that any downtime that results in mission-critical data being unavailable is unacceptable. Windows Server 2003 offers several solutions to increase the availability of data. Solutions that use RAID and clustering technologies are especially suited to providing fault-tolerant storage. The process for analyzing and selecting among these solutions is illustrated in Figure 1.9.

**Figure 1.9 Planning for Fault Tolerance**



# Achieving Fault Tolerance by Using RAID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

RAID is commonly implemented for both performance and fault tolerance. With RAID, you can choose to assemble disks to provide fault tolerance, performance, or both, depending on the RAID level that you configure. Table 1.4 summarizes commonly available RAID levels.

**Table 1.4 RAID Comparison**

| RAID Level | Description | Minimum Disks Required | Effective Capacity |
|---|---|---|---|
| 0 | Disk Striping. Two or more disks appear to the operating system as a single disk. Data is striped across each disk during read/write operations. Potentially increases disk access speeds 2X or better. Not fault tolerant. | 2 | S*N<br>*N* = # of disks in array<br>*S* = Size of smallest disk in array |
| 1 | Disk mirroring. Data is mirrored on two or more disks. Provides fault tolerance, but at a higher cost (space required is double the amount of data). Read performance is increased as well. | 2 | S<br>*S* = Size of smallest disk in array |
| 0+1 | Combines RAID 0 and RAID 1; offers the performance of RAID 0 and the protection of RAID 1. | 4 | S*N/M<br>*N* = # of disks in array<br>*S* = Size of smallest disk in array<br>*M* = Number of mirror sets |
| 5 | Disk striping with parity. Provides slower performance than RAID 0, but provides fault tolerance. A single disk can be lost without any data loss. Parity bits are distributed across all disks in the array. | 3 | S*(N-1)<br>*N* = # of disks in array<br>*S* = Size of smallest disk in array |

From a design perspective, your choice of a RAID solution should be dictated by the type of data being stored. Although RAID 0 offers the fastest read and write performance, it does not offer any fault tolerance, so that if a single disk in a RAID 0 array is lost, all data is lost and will need to be recovered from backup. This might be a good choice for high performance workstations, but might not be suited to mission-critical servers.

RAID 1 allows you to configure two or more disks to mirror each other. This configuration produces slow writes, but relatively quick reads, and provides a means to maintain high data availability on servers, because a single disk can be lost without any loss of data. When more than two disks make up the mirror, the RAID 1 array can lose multiple disks so long as a complete mirrored pair is not lost. When planning a RAID 1 solution, remember that the amount of physical disk space required is twice the space required to store the data.

RAID 0+1 combines the performance benefit of striping with the fault tolerance of mirroring. Compared to RAID 0, writes are

slower, but reads are equally fast. Compared to RAID 1, RAID 0+1 offers faster writes and reads but also requires additional storage to create the mirrored stripe sets. This configuration is often ideal for mission-critical database storage, because it offers both fast read access and fault tolerance.

RAID 5 provides fault tolerance: you can lose a single disk in an array with no loss of data. However, RAID 5 operates much more slowly than RAID 0 because a parity bit must be calculated for all write operations. RAID-5 volumes are well suited for reads and also work well in the following situations:

- In large query or database mining applications where reads occur much more often than writes. Performance degrades as the percentage of write operations increases. Database applications that read randomly work well with the built-in load balancing of a RAID 5 volume.

- Where a high degree of fault tolerance is required without the cost of the additional disk space needed for a RAID 1 volume. A RAID 5 volume is significantly more efficient than a mirrored volume when larger numbers of disks are used. The space required for storing the parity information is equivalent to 1/Number of disks, so a 10-disk array uses 1/10 of its capacity for parity information. The disk space that is used for parity decreases as the number of disks in the array increases.

**Choosing Between Hardware and Software RAID**

An additional consideration with RAID implementations is the choice between hardware-based and software-based RAID. With hardware RAID, a hardware RAID controller allows you to configure the RAID level of attached disks. With software RAID, the operating system manages the RAID configuration, along with data reads and writes.

Windows Server 2003 supports the following software RAID types:

- RAID 0: Up to 32 disks striped

- RAID 1: Two disks mirrored

- RAID 5: Up to 32 disks striped with parity

To configure software RAID, a disk must be configured as a dynamic disk. Although software RAID has lower performance than hardware RAID, software RAID is inexpensive and easy to configure because it has no special hardware requirements other than multiple disks. If cost is more important than performance, software RAID is appropriate. If you plan to use software RAID for write-heavy workloads, use RAID-1 instead of RAID-5. Using software-based RAID with Windows Server 2003 dynamic disks is a good choice for providing fault tolerance using SCSI or EIDE disks to non-mission-critical servers that require fault tolerance and can accommodate the added CPU load imposed by software RAID. This solution is ideal for small to medium organizations that need to add a level of fault tolerance while avoiding the cost of hardware RAID controllers.

For more information about software RAID and disk management, see Storage Technologies Collection (http://go.microsoft.com/fwlink/?linkid=4681).

**Using Dynamic Disks with Hardware RAID**

You can also use dynamic disks with hardware-based RAID solutions. Using dynamic disks with hardware RAID can be useful in the following situations:

- You want create a large volume by using software RAID, such as RAID-0, across hardware RAID LUNs.

- You want to extend a volume, but the underlying hardware cannot dynamically increase the size of LUNs.

- You want to extend a volume, but the hardware has reached its maximum LUN size.

Before converting hardware RAID disks to dynamic disks, review the following restrictions:

- You cannot use dynamic disks on shared cluster storage. However, you can use the DiskPart command-line tool to extend basic volumes on shared cluster storage. For more information, see "Extend a basic volume" in Help and Support Center for Windows Server 2003.

- If you create a software RAID-0 volume across multiple hardware arrays, you cannot later extend the RAID-0 volume to increase its size. If you anticipate needing to extend the volume, create a spanned volume instead.

**Preparing to Upgrade Servers That Contain Multidisk Fault-Tolerant Volumes**

Servers running Microsoft® Windows NT® Server version 4.0 can contain volume sets, mirror sets, stripe sets, and stripe sets with parity created by using the fault-tolerant driver Ftdisk.sys. To encourage administrators to begin using dynamic volumes, Windows 2000 offers limited support for Ftdisk volumes. Completing this transition, Windows Server 2003 does not support multidisk volumes. If you plan to upgrade a server that contains multidisk volumes, review the following issues:

- If you are upgrading from Windows NT Server version 4.0 to Windows Server 2003, back up and then delete all multidisk volumes before upgrading. This is necessary because Windows Server 2003 cannot access these volumes. Be sure to verify that your backup was successful before deleting the volumes. After you finish upgrading to Windows Server 2003, create new dynamic volumes, and then restore the data from your backup.

- If you are upgrading from Windows NT Server 4.0 to Windows Server 2003, and the paging file is located on a multidisk volume, you must use **System** in Control Panel to move the paging file to a primary partition or logical drive before beginning Setup. For more information about moving the paging file, see article 123747 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=99315).

- If you are upgrading from Windows 2000 Server to Windows Server 2003, you must use the Disk Management snap-in to convert all basic disks that contain multidisk volumes to dynamic disks before beginning Setup. If you do not do this, Setup does not continue. For information about converting basic disks to dynamic disks, see "Change a basic disk into a dynamic disk" in Help and Support Center for Windows Server 2003.

For more information about multidisk volumes, see Storage Technologies Collection (http://go.microsoft.com/fwlink/?linkid=4681).

# Achieving Fault Tolerance by Using Clustering

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many organizations require that critical data be continuously available. Cluster technology provides a means of configuring storage to help meet that goal. Simply put, a cluster is two or more computer systems that act and are managed as one. Clients access the cluster by using a single host name or IP address; their request is answered by one of the systems in the cluster.

The purpose of cluster technology is to eliminate single points of failure. When availability of data is your paramount consideration, clustering is ideal. Using a cluster avoids all of these single points of failure:

- Network card failure
- Processor failure
- Motherboard failure
- Power failure
- Cable failure
- Storage adapter failure

With a cluster, you can essentially eliminate nearly any hardware failure associated with using a single computer. If hardware associated with one system fails, the other system automatically takes over. Two types of clustering solutions that accomplish this are *server clusters* and *Network Load Balancing clusters*. Both types of clustering are available on Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition. In addition, Network Load Balancing clusters are available on Windows Server 2003, Web Edition and Windows Server 2003, Standard Edition.

**Server Clusters**

Server clusters are often implemented to offer high availability solutions to applications that need both read and write access to data, such as database, e-mail, and file servers. Server clusters can be configured with up to eight computers, or nodes, participating in the cluster. To share the same data source, server cluster nodes connect to external disk arrays by using either a SCSI or Fibre Channel connection. Fibre Channel is required for interconnecting clusters of three or more nodes to shared storage. For the 64-bit versions of Windows Server 2003, you must always use Fibre Channel hardware to connect the nodes to shared storage.

When planning to deploy server clusters in your storage solution, you must take into account the following considerations:

- The boot and system disks of each cluster node must not be located on the same storage bus as the shared storage devices, unless you use a Storport driver for your HBAs.
- Shared cluster disks cannot be configured as dynamic disks.
- Shared cluster disks must be formatted as basic disks with the NTFS file system.
- For the 64-bit versions of the Windows Server 2003 family, the shared cluster disks must be partitioned as master boot record (MBR) and not as GUID partition table (GPT) disks.
- You cannot use Remote Storage with shared cluster storage.
- You should not enable write caching on shared cluster disks unless they are logical units on an external RAID subsystem that has proper power protection (such as multiple power supplies, multiple feeds from the power grid, or adequate battery backup).
- Because cluster disks must be basic disks, you cannot use software RAID. For disk fault tolerance, you must use a hardware-based RAID solution.

For more information about server clusters, see "Designing and Deploying Server Clusters" in this book.

**Network Load Balancing Clusters**

Network Load Balancing clusters maintain their own local copy of data and are ideal for load balancing access to static data, such as Web pages. Up to 32 computers can participate in a Network Load Balancing cluster. Because they manage their own local data, Network Load Balancing clusters are much easier to plan and implement. By using the Network Load Balancing Manager, you can quickly configure all Network Load Balancing clusters in your enterprise from a single server.

Using Network Load Balancing clusters is the best choice for several data availability needs. For any server that has difficulty meeting the load demands of its clients, Network Load Balancing is an ideal solution, and is commonly used to provide fault tolerance and load balancing for:

- Web Servers
- FTP Servers
- Streaming Media Servers
- VPN Servers
- Terminal Servers

For each of these, Network Load Balancing is ideal, not only because it is easy to implement, but also because of how easily a Network Load Balancing cluster can scale as your company grows. Because of their simple scalability, your initial estimates of the number of servers you will require need not be perfect with Network Load Balancing clusters. As the load on a Network Load Balancing cluster grows, you can balance the increased load by simply adding additional hosts to the cluster.

Because each Network Load Balancing cluster host maintains its own local copy of storage, storage planning with Network Load Balancing clusters is not as complex as with server clusters. Many of the disk restrictions of server clusters do not apply to Network Load Balancing clusters. The general storage considerations for Network Load Balancing cluster planning are:

- Network Load Balancing cluster hosts can use any local storage space, including space on boot or system volumes.
- Local storage can consist of basic and dynamic disks.
- Hardware or software RAID can be used to add additional fault tolerance. If the Network Load Balancing cluster services a high level of traffic and you need disk fault tolerance, you must use hardware RAID.

For more information about Network Load Balancing clusters, see "Designing Network Load Balancing" and "Deploying Network Load Balancing" in this book.

# Planning for Disaster Recovery

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A number of technologies are available that provide fault tolerance in the event of a failure, such as fault tolerant disk configurations with hot swappable drives, server clusters, and uninterruptible power supplies, However, all of these high availability technologies cannot substitute for having a reliable backup of mission-critical data.

In a complete site disaster, for example, it is possible that both online and offline availability technologies are destroyed (for example, all cluster nodes and all disks in a RAID array). Following such a system failure or disaster at a particular site, you must be able to recover data and systems from backup. Recovering systems or sites from failure is a daunting task, unless you have thoroughly planned and prepared by implementing scheduled backups, providing for backups of open files, and configuring for Automated System Restore. You must also thoroughly test your ability to restore data. The tasks involved in planning for disaster recovery are shown in Figure 1.10.

**Figure 1.10 Planning for Disaster Recovery**



# Planning a Backup Schedule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The backup schedule is perhaps the most important consideration when planning for disaster recovery. It is critical to plan your backups around your restore requirements. For example, you run only incremental backups of a file server over an extended period of time. This approach delivers very fast backups that use little network bandwidth. However, to restore the server, you would need to restore all of the incremental backups, which could take days to complete. A backup plan presumes a restore plan: Your backup plan should always be based on your requirements for restoring data. You should carefully test and document all of the processes and time required to fully recover your servers from backup data. Use the results of this testing to establish the requirements for your backup schedules.

# Backup Decision Points

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To plan an effective backup schedule that meets your organization's data protection and recovery requirements, consider both the available backup types and your restore plan. When considering backup types, remember that your choice must accommodate restoring a system within the time specified in your restore plan. Available backup types can be categorized as fundamental and advanced.

**Fundamental Backup Types**

The fundamental backup types available are:

- **Full.** Full backups are the baseline for all other backups, and contain all the data on a system (or all the data in the folders or volumes that are defined to be backed up). Because full backups secure all server data, frequent full backups can provide you with greater guarantees as to the speed and success of restore operations. Remember that when you add additional backup types to the backup set, restore jobs are prolonged.

- **Incremental.** An incremental backup stores all files that have changed since the last backup, regardless of the backup type. The advantage of incremental backups is that they take the least time to complete. However, during a restore operation, each incremental backup in the backup set is applied, which could result in a lengthy restore job.

- **Differential.** Differential backups contain all data that has changed since the last full backup. The advantage of differential backups is that they can shorten restore time compared to a backup set that includes just full and incremental backups. However, if you perform too many differential backups within a single backup set, the size of the differential backup might grow to be as large as the baseline full backup.

- **Copy.** Copy backups are identical to full backups, with the exception that they do not mark files as backed up, and thus do not impact any other backup types. A full backup marks the beginning of a new backup cycle, while a copy backup does not. Copy backups are most frequently used to create offsite copies of backup data. In this way, you can maintain a local full backup for quick restore purposes, and keep an offsite copy backup for use in the event that your local backup copy is lost.

Your backup schedule should probably include some combination of all backup types. Regular full backups take the longest to complete, but offer the quickest restore performance. Because full backups copy all files defined in the backup job, they also consume the most backup media. Incremental backups consume the least backup media and take the shortest time to complete; however, they result in lengthier restores. Consequently, it is best to balance each of the backup types so as to satisfy your organizations restore needs. For example, you might elect to perform weekly full backups of a crucial file server, and supplement that with daily differential backups, and incremental backups every two hours. This approach gives you a high degree of data protection, while keeping restore time within reasonable boundaries.

**Advanced Backup Types**

Advanced backup types include:

- **Backup applications that leverage the Volume Shadow Copy service.** Backup applications that support the Volume Shadow Copy service can back up open files, and all the backed-up files are from a single point-in-time. For more information about backing up open files, see "Planning to Back Up Open Files" later in this chapter.

- **Server-free backups.** In a server-free backup, a *data mover* on a SAN copies data on a SAN disk to another storage device, such as a library on the SAN. A data mover is software that can reside on a SAN device, such as a switch or router, or on a server connected to the SAN. Because the data mover performs the copy operations, the server can maintain a high level of performance for clients. Server-free backup applications must have a way of creating a point-in-time snapshot, such as by leveraging the Volume Shadow Copy service or by using some other technology.

- **Hardware-based shadow copies.** External storage arrays that support hardware-based shadow copies (also called *snapshots*) can implement Volume Shadow Copy service providers to fully utilize advanced storage array features.

Your choice of how to integrate advanced backup types into your backup schedule will be dictated by your restore requirements. Some systems might require hourly shadow copies, while other, low priority systems might only warrant weekly full backups and daily incremental backups. Your best strategy is to align your backup plan for each system with the restore plan for that system. If the system requires quick recovery of user files, frequent shadow copies are ideal. If the data on the same system is required to be retained for seven years, then you will also need to perform frequent full backups of the system to media that can be stored in an offsite location.

For more information about backup and recovery planning, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Planning to Back Up Open Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Open files are files that are normally skipped during the backup process. This is usually because the files are locked by a service or application, such as an operating system, word processing program, or application database. In Windows 2000 and Windows NT 4.0, if the operating system could not back up a file that was locked by an application, the file was skipped, and thus not backed up. In Windows Server 2003, open files are backed up by using the Volume Shadow Copy service.

When a backup is initiated by an application that can use the Volume Shadow Copy service, such as the Backup program in Windows Server 2003, the Volume Shadow Copy service makes a shadow copy of the volume to be backed up. The shadow copy constitutes a read-only copy of the volume data that is read by the backup application during the backup job. Applications can continue to access the files on the volume itself, uninterrupted by the backup. After the backup is completed, the shadow copy of the volume is deleted, because it is no longer needed. The backed-up data is stored on the backup media.

By default, Windows Server 2003 temporarily consumes free disk space on a volume for the shadow copy. The amount of disk space consumed depends on the amount of data that changes on the volume during the backup.

In the event that a shadow copy is unsuccessful, for example, when there is not enough temporary disk space available on the volume, Backup continues without using shadow copy techniques and, as in previous versions of Windows, reads files from the original volume and does not back up any open files.

To take advantage of open-file backups, purchase a backup application that works with the Volume Shadow Copy service.

# Planning for Bare Metal Restores with ASR

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Automated System Recovery (ASR) is a new tool for use with Backup (NTBackup.exe) and other programs created by ISVs. It has a limited but critical purpose: to help you automatically restore your system after a system failure. Previously, restoring the complete system required that you first reinstall the operating system and then restore the data. With a bare metal restore, you can boot a system from the operating system CD, and then use an ASR floppy disk to recover the system directly from a backup.

ASR works with Windows Setup to rebuild the storage configuration of the physical disks and writes the critical operating system files to the boot and system partitions in order to allow the system to boot successfully. This process is referred to as a *bare metal restore*, because the system is restored to hardware that has no installed software. The process uses an ASR floppy disk that defines the state of the storage prior to the disaster and the process to be used for restoring the server. After an ASR restore completes, you can restore any needed user or application files.

# How ASR Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ASR represents a new approach to backup and recovery. Prior to ASR, after a large-scale failure, you needed to reinstall Windows, configure all physical storage to the original settings, and then perform a complete restore of the data. The process of rebuilding the operating system could be lengthy, and you needed to perform many of these tasks at the local computer. ASR significantly automates this process. In addition to automating the restore of a single system, ASR can be used with Remote Installation Services (RIS) to automate the system state recovery of several systems across the network.

To prepare for ASR recovery, you must run the Automated System Recovery Wizard, which is part of Backup. To access this wizard when you are running Backup in Advanced Mode, click **Tools** and select **ASR Wizard**. You can start Backup in Advanced Mode by clearing the **Always start in Wizard Mode** check box when Backup starts.

The wizard backs up the operating system boot volumes and system volumes, but does not back up other volumes, such as program or data volumes. To secure data on other volumes, you must back up those volumes separately by using Backup or another backup tool. You can, however, choose to back up **All information on this computer** when running Backup. This option creates a full backup of your entire system, including ASR data. This means that you can recover the entire system through the ASR process in the event of failure.

When an ASR restore is initiated, ASR first reads the disk configurations from the ASR floppy disk and restores all disk signatures and volumes on the disks from which the system boots. In the ASR process, these are known as *critical disks*, because they are required by the operating system. Noncritical disks — disks that might store user or application data — are not backed up as a part of a normal ASR backup, and are not included in an ASR restore. If these disks are not corrupted, their data will still be accessible after the ASR restore completes. If you want to secure data on noncritical disks from disk failure, you can do so by backing it up separately.

After the critical disks are recreated, ASR performs a simple installation of Windows Server 2003 and automatically starts a restore from backup using the backup media originally created by the ASR Wizard. During an ASR restore, any Plug and Play devices on the system are detected and installed.

Before performing an ASR restore, ensure that the target system to which the restore will be made meets the following requirements:

- The target system hardware (except for hard disks, video cards, and network adapters) is identical to that of the original system.
- There are enough disks to restore all the critical system disks.
- The number and storage capacity of the critical disks are at least as great as those of the corresponding original disks.

**Caution**

- Do not depend on ASR to back up and recover user data files stored on the boot and system volumes. In addition, because your system volume is formatted during the ASR recovery process, any user files or directories located on those volumes are lost.

You normally access the ASR state file (Asr.sif) through a local floppy disk drive. If the computer does not have a floppy disk drive, or you want to perform an ASR restore over a network or remotely, you can use a Remote Installation Services (RIS) server to fully automate the ASR process. RIS uses Pre-boot eXecution Environment (PXE) technology to enable client computers without an operating system to boot remotely to a RIS server that performs installation of a supported operating system over a TCP/IP network connection. Consequently, the remote installation client computer must have a PXE-enabled network adapter.

For more information about using RIS to perform remote installations, see "Designing RIS Installations" in *Automating and Customizing Installations* of this kit.

**Note**

- ASR behaves differently from the Emergency Repair Disk feature in Windows 2000 Server, which ASR replaces. Emergency Repair Disk replaces missing or corrupt system files without formatting drives or reconfiguring storage. ASR,

by contrast, always formats the boot volume and might format the system volume.

# Guidelines for Using ASR

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To successfully use ASR in your disaster recovery plan, you should include the following guidelines:

- Run ASR backup regularly, preferably by using automatic settings.
- Plan for making the required resources, including tape backup drives and removable and hard disks, available for ASR recovery.
- Perform any needed file system conversions before running your first ASR backup.
- Plan for conditions that might prevent a fully successful ASR restore.

  Under the following conditions, ASR might not be able to restore all disk configurations

  - If a critical volume is not accessible during an ASR restore, the restore will fail.
  - Noncritical disks that are a part of the ASR backup are not restored if they are not found during the ASR restore, but the balance of the restore will complete successfully. Disk types that might not appear to the restore process include IEEE 1394, USB, or Jaz disks.
- Plan to protect the critical files Asr.sif and Asrpnp.sif generated by Backup and copied to your ASR floppy disk.

  If the ASR floppy disk that contains these files is lost, you can recover the files from the *systemroot*\Repair folder on the host system. If these files are not accessible on the original host, you can recover them from the ASR backup media by using another system. By storing these files in three locations — the ASR floppy disk, the Repair folder, and on ASR backup media — you have three levels of protection against their loss.

For more information about ASR, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit) and see "Automated System Recovery (ASR) overview" in Help and Support Center for Windows Server 2003.

# Testing Restores

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The single most overlooked aspect of disaster recovery planning is testing restores. Receiving confirmation that a backup completed successfully does not guarantee that a backup can be restored. To prepare for recovery and to validate backup data, you should periodically test your backups of mission-critical servers. If a server cannot be taken offline for testing restores, you can instead restore its backup data to a test server. Practicing restore operations allows you to prepare for problems that you might encounter when recovering a complete system after a failure. These problems include the following:

- When restoring to a new server with different hardware — for example, if the original server used SCSI storage and the server where you restore the data uses IDE storage — you might need to filter files such as Boot.ini from the restore job. If the restored data replaces the Boot.ini file, the system might not be able to boot following the restore, and the parameters in the restored Boot.ini file would need to be modified.

- When restoring to a new server with different hardware, you might need different drivers, such as drivers for the network adapter or video adapter. Make sure that these drivers are available during the restore.

- The time required to restore data might exceed the time allotted in your recovery plan, which means you must change your backup schedule, your process, or your equipment to improve the speed of the restoration.

- Backup media might be corrupt. Your disaster recovery plan should specify preventive measures (such as redundant backups), and indicate the procedures to follow in the event of corrupt media (such as the location of any redundant backups, or whether out-of-date backups are to be used).

Testing restores can help you plan how to deal with such problems. For example, if your organization requires that a critical file server be returned to full operation within a four hour window, but your test restore took six hours to complete, you can plan to either adjust your backup schedule or else justify the purchase faster backup media drives or network components. Without thoroughly testing your restore procedures, you cannot conclusively document recovery procedures or recovery time, both of which are crucial to your organization's disaster recovery plan.

# Additional Resources for Planning for Storage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Related Information**

- "Designing and Deploying File Servers" in this book.
- "Planning for High Availability and Scalability" in this book.
- "Designing and Deploying Server Clusters" in this book.
- "Designing Network Load Balancing" in this book.
- "Deploying Network Load Balancing" in this book.
- The *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit) for more information about disk management and file systems.
- The Windows Powered Network Attached Storage link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Storage Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Windows Clustering: Storage Area Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Using NTFS mounted drives" in Help and Support Center for Windows Server 2003.
- "Using Shadow Copies of Shared Folders in a server cluster" in Help and Support Center for Windows Server 2003.
- "Extend a basic volume" in Help and Support Center for Windows Server 2003.
- "Change a basic disk into a dynamic disk" in Help and Support Center for Windows Server 2003.
- "Automated System Recovery (ASR) overview" in Help and Support Center for Windows Server 2003.

**Related Tools**

- Diskraid.exe

  Use Diskraid.exe to manage and configure RAID subsystems that have the necessary VDS hardware providers. For more information about Diskraid.exe, in Help and Support Center for Windows Server 2003 click **Tools**, and then click Command-line reference A-Z.

- Diskpart.exe

  Use Diskpart.exe to manage disks and volumes from the command line. For more information about Diskpart.exe, in Help and Support Center for Windows Server 2003 click **Tools**, and then click Command-line reference A-Z.

# Designing and Deploying File Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

File servers running the Microsoft® Windows® Server 2003 operating system are ideal for providing access to files for users in medium and large organizations. Windows Server 2003 offers a number of file server solutions, such as Distributed File System (DFS), File Replication service (FRS), Windows server clusters, NTFS permissions, disk quotas, and shadow copies, for enhancing the manageability, scalability, availability, and security of file servers.

**In This Chapter**


Overview of Designing and Deploying File Servers

Identifying File Services Goals

Designing DFS Namespaces

Planning File Server Availability

Designing a Standard File Server Configuration

Planning File Server Security

Deploying File Servers

Additional Resources for Designing and Deploying File Servers


**Related Information**

- For information about installing and managing software applications by using Group Policy, see "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit.
- For information about managing user desktops, settings, and data by storing data and settings on network servers, see "Implementing User State Management" in *Designing a Managed Environment*.
- For information about server clusters, see "Designing and Deploying Server Clusters" in this book.

# Overview of Designing and Deploying File Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

File servers provide users a way to access shared data within their department or organization. The simplest way to create a file server is to share a folder on a server. However, this solution does not provide file server manageability, scalability, availability, or security. To achieve these goals, you can deploy the following Windows Server 2003 solutions:

**Shadow copies** If users frequently require administrators to restore deleted or overwritten files from tape, shadow copies provide point-in-time copies of files in shared folders, allowing users to recover files that were accidentally deleted or overwritten.

**DFS** If users need to access files on multiple file servers without having to keep track of all the server names, you can use DFS to logically group physical shared folders located on different servers by transparently connecting them to one or more hierarchical namespaces. DFS also provides fault-tolerance and load-sharing capabilities.

**FRS and Windows server clusters** Windows Server 2003 provides two independent solutions, FRS and server clusters, to ensure that important business data is always available, even if a server fails or is taken offline for maintenance.

**Disk quotas** By using the disk quotas feature in Windows Server 2003, you can track files on a per-volume, per-user basis to monitor disk space use and to prevent file servers from filling to capacity without warning.

**NTFS permissions** To prevent unauthorized users from accessing folders, you can use NTFS file system permissions to specify the groups and users whose access you want to restrict or allow and then select the type of access.

You can use the following information to design your organization's file services. Or, if your organization already has file servers running Microsoft® Windows® NT 4.0 or Microsoft® Windows® 2000 operating systems, you can use the following information to improve the design of your existing file services to take advantage of the new and enhanced features in the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems.

# File Server Design and Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 2.1 outlines the general process of file server deployment. This process consists of a design phase and a deployment phase. A file services designer or design team performs the design process. The file services deployment team implements the design by deploying file servers running Windows Server 2003 and configuring the file services solutions described in this chapter. Because some features of DFS and FRS require the Active Directory® directory service, this process assumes that your organization already has an existing Active Directory infrastructure.

**Figure 2.1 Designing and Deploying a File Server**



# New and Enhanced File Server Features in Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following sections summarize the new and enhanced features in Windows Server 2003 that are described in this chapter.

**Shadow copies**
The shadow copy feature provides point-in-time copies of files on a volume, allowing users to view the contents of shared folders as they existed at points of time in the past. After you enable this feature, users can recover files that they accidentally delete or overwrite.

**DFS enhancements**
Windows Server 2003 includes the following DFS enhancements:

- You can create multiple DFS roots on servers running Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.
- You can configure DFS to choose an alternate target based on the lowest connection cost if no same-site targets are available.
- You can use the Distributed File System snap-in to choose replication topologies that complement your network infrastructure.
- You can move a root target or a link target from one Active Directory site to another, and DFS will update the root or link information for the new site within 25 hours.
- To reduce network traffic to the server acting as primary domain controller (PDC) emulator master, you can configure DFS to get namespace updates from the closest domain controller, instead of from the server acting as the PDC emulator master. This mode is known as *root scalability mode*.
- By using the Offline Files feature, you can make shared folders that correspond to DFS link targets available offline to clients running Microsoft® Windows® XP Professional and Windows Server 2003. (You cannot make DFS link targets available offline to clients running Windows 2000.)

**FRS enhancements**
Windows Server 2003 includes the following FRS enhancements:

- FRS detects and suppresses excessive replication.
- FRS manages the staging directory to prevent it from becoming full.
- FRS supports connection priorities, which allow you to control the sequencing of the initial synchronization that occurs when you add a new member to the replica set or when you perform a *nonauthoritative restore*, which is used to bring a failed replica member back into synchronization with its partners.

**Security enhancements**
When you share a folder, the default permission is the Read permission for the Everyone group. This default is different from the default share permission in Windows 2000, which was the Full Control permission for the Everyone group.

**Server cluster enhancements**
Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition include the following server cluster enhancements:

- You can store encrypted files by using Encrypting File System (EFS) on cluster storage.

- You can create multiple stand-alone DFS roots on a clustered file server, and the roots can reside on any of the nodes in the cluster.

- You can enable client-side caching to cache files from a cluster file share onto client computers. Clients can then access these files even when the client computer is disconnected from the network.

# Identifying File Services Goals

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some of the major steps for designing and deploying file servers might not apply to your organization. Therefore, the first step in the design process is to identify the goals that you want to achieve by deploying file servers running Windows Server 2003, as shown in Figure 2.2. These goals determine the design and deployment steps that are necessary for your organization.

**Figure 2.2 Identifying File Services Goals**




The following sections describe common goals for file services. Use the information in these sections to identify the goals for your organization and to find the relevant sections in this chapter or other sources of information to help you achieve those goals.

**Improving the way users access files on file servers**

If you want to improve how users access files on file servers, you might have the following goals:

- Providing an intuitive way for users to access multiple file servers throughout the organization.
- Making data on multiple file servers appear as though it were available on a single file server.
- Making data available in multiple sites so that users in each site use fast, inexpensive bandwidth to access the data.
- Reducing delays that occur when users access heavily used shared folders.
- Providing fault-tolerant access to shared folders.
- Consolidating file servers or migrating data without affecting how users locate data.

For more information about improving how users access files on file servers, see "Designing DFS Namespaces" later in this chapter.

**Managing applications and user data and settings**

If you are managing applications, user data, and settings, you might have the following goals:

- Enabling users to access files even when they are not connected to the network.
- Storing application files on file servers so that users can install the applications from the network to their local workstations.
- Using Group Policy–based software management to deploy, upgrade, update, and remove users' applications without

going to individual workstations.

- Allowing users to run applications from the file server.

For more information about managing applications and user data and settings, see "Implementing User State Management" and "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit. For more information about hosting applications in a central location, see "Hosting Applications with Terminal Server" in this book.

**Adding storage to file servers**
If you plan to add storage to file servers, you might have the following goals:

- Transparently adding more storage to a file server.
- Making data on multiple volumes or disks in a file server appear within a single volume or drive letter.
- Creating more than 26 volumes on a server without being limited by the 26-drive letter limit.

For more information about adding storage to file servers, see "Using NTFS mounted drives" in Help and Support Center for Windows Server 2003.

**Planning for file server availability and reliability**
If you are planning for file server availability and reliability, you might have the following goals:

- Choosing file server hardware for reliability and availability.
- Ensuring data availability if a file server fails or is taken offline for maintenance.
- Making data available in multiple sites to provide inexpensive access to users within each site.

For more information about planning for file server availability and reliability, see "Planning File Server Availability" later in this chapter.

**Choosing file server hardware and settings**
If you are choosing file server hardware and settings, you might have the following goals:

- Choosing compatible file server hardware that meets your performance and storage requirements.
- Increasing file server performance.
- Consolidating file servers to reduce management costs and increase storage allocation efficiency.
- Enabling users to access previous versions of files on the file server.
- Monitoring and controlling disk space use.

For more information about choosing file server hardware and settings, see "Designing a Standard File Server Configuration" later in this chapter.

**Planning for file server security**
If you are planning for file server security, you might have the following goals:

- Protecting file servers from viruses.
- Preventing unauthorized users from accessing data on file servers.
- Allowing users to store encrypted files on a file server.

For more information about planning for file server security, see "Planning File Server Security" later in this chapter.

# Designing DFS Namespaces

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Users might have difficulty finding information in shared folders that are located on numerous file servers. Because shared folders are usually associated with physical servers, the user must first determine which physical server is hosting the shared folder. For example, a user might need to access product information about a server named \\Building 4\Marketing2\Prod_Info and on a server named \\Corporate\Floor 4\Sales\Prod_Info.

You can use DFS to address this challenge by consolidating a large set of physical shared folders into one or more virtual namespaces. You do not need to modify the shared folders to add them to the namespace, and users can navigate the namespace without having to know the physical server names or shared folders hosting the data.

Figure 2.3 outlines the general process of designing one or more DFS namespaces. For an Excel spreadsheet to assist you in documenting your DFS namespace design decisions, see "DFS Configuration Worksheet" (Sdcfsv_1.xls) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**Figure 2.3 Designing DFS Namespaces**




For more information about DFS security, see "Planning DFS and FRS Security" later in this chapter. For in-depth technical and troubleshooting information about DFS, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Deciding Whether to Implement DFS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Organizations of any size, with any number of file servers, can benefit from implementing DFS. DFS is especially beneficial for organizations in which any of the following conditions exist:

- The organization plans to deploy additional file servers or consolidate existing file servers.
- The organization has data that is stored in multiple file servers.
- The organization wants to replace physical servers or shared folders without affecting how users access the data.
- The organization has data located on servers in multiple sites and wants clients to connect to the closest servers.
- Most users require access to multiple file servers.
- Users experience delays when accessing file servers during peak usage periods.
- Users require uninterrupted access to file servers.

Even if you are busy planning your organization's migration to Windows Server 2003, you can make plans to implement DFS without immediately designing your entire namespace. You do not need to deploy DFS all at one time; you can choose to add as much or as little of your organization's physical storage as you need to the DFS namespace, at a pace that works with your overall migration schedule.

When deciding whether to implement DFS, do the following:

1. Review DFS terminology.
2. Review the benefits of using DFS.
3. Evaluate clients and servers for compatibility.

The following sections describe each of these steps.

# Reviewing DFS Terminology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are not familiar with DFS, review the following terms and definitions to understand the important elements of a DFS configuration. For visual examples of these concepts, see Figure 2.4 through Figure 2.7 later in this section.

**DFS namespace** A virtual view of shared folders on different servers as provided by DFS. A DFS namespace consists of a root and many links and targets. The namespace starts with a root that maps to one or more root targets. Below the root are links that map to their own targets.

**DFS root** The starting point of the DFS namespace. The root is often used to refer to the namespace as a whole. A root maps to one or more root targets, each of which corresponds to a shared folder on a separate server. The DFS root must reside on an NTFS volume. A DFS root has one of the following formats: *servername**rootname* or *domainname**rootname*.

**Root target** A physical server that hosts a DFS namespace. A domain-based DFS root can have multiple root targets, whereas a stand-alone DFS root can only have one root target.

**Stand-alone DFS namespace** A DFS namespace whose configuration information is stored locally in the registry of the host server. The path to access the root or a link starts with the host server name. A stand-alone DFS root has only one root target. Stand-alone roots are not fault tolerant; when the root target is unavailable, the entire DFS namespace is inaccessible. You can make stand-alone DFS roots fault tolerant by creating them on clustered file servers.

**Domain-based DFS namespace** A DFS namespace that has configuration information stored in Active Directory. The path to access the root or a link starts with the host domain name. A domain-based DFS root can have multiple root targets, which offers fault tolerance and load sharing at the root level.

**DFS path** Any Universal Naming Convention (UNC) path that starts with a DFS root.

**Link** A component in a DFS path that lies below the root and maps to one or more link targets.

**Link target** The mapping destination of a link. A link target can be any UNC path. For example, a link target could be a shared folder or another DFS path.

Figure 2.4 illustrates the elements of a stand-alone DFS namespace in the Distributed File System snap-in. These elements include a stand-alone DFS root, a single root target, and multiple links.

**Figure 2.4 Elements of a Stand-Alone DFS Namespace**




Figure 2.5 illustrates the elements of a domain-based DFS namespace in the Distributed File System snap-in. Notice that the \\Reskit.com\Public root has two root targets on different servers.

**Figure 2.5 Elements of a Domain-based DFS Namespace**



Figure 2.6 illustrates multiple link targets for the Software link. Notice that the link targets exist on three different servers and that the administrator has disabled referrals to the link target on \\dfs-03. DFS will not refer clients to the link target on \\dfs-03 until the administrator enables referrals.

**Figure 2.6 Multiple Link Targets**




The roots and links displayed in the Distributed File System snap-in also appear on each root server's local storage as follows:

- When you create a DFS root, you specify a shared folder to use as the root folder. If you add multiple root targets to a domain-based DFS root, you specify a shared folder on each of those root targets. (The shared folder names should always match the root name.)

- When you add links to the root, DFS creates special folders under each root folder. These folders, called *link folders*, are actually reparse points, and they display the following error message if you try to access them on the local server:

  E:\Public\GroupData is not accessible. The network location cannot be reached.

  Users who access the link folders from across the network are redirected to the appropriate link target.

Figure 2.7 illustrates volume E:\ on the local storage of one of the root targets. The volume contains root and link folders for the \\Reskit.com\Public namespace.

**Figure 2.7 Root and Link Folders**



# Reviewing the Benefits of Using DFS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you evaluate DFS for your organization, it is helpful to understand the benefits that your organization can gain after designing and implementing a DFS namespace. The following list describes the benefits of using DFS:

**Unified namespace** A DFS namespace links together shared folders on different servers to create a hierarchical structure that behaves like a single high-capacity hard disk. Users can navigate the logical namespace without having to know the physical server names or shared folders hosting the data.

**Location transparency** DFS simplifies migrating data from one file server to another. Because users do not need to know the name of each physical server or shared folder that contains the data, you can physically move data to another server without having to reconfigure applications and shortcuts, and without having to re-educate users about where they can find their data.

**Storage scalability** You can deploy additional or higher-performance file servers and present the storage on the new servers as new folders within an existing namespace.

**Namespace scalability** Servers running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition can host multiple domain-based DFS roots and stand-alone DFS roots. This feature improves the scalability of DFS, enabling you to build many large namespaces without having to add file servers to host the roots.

**Increased availability of file server data** When multiple servers running Windows Server 2003 host a domain-based DFS root, clients are redirected to the next available root server if any of these servers fail, providing fault-tolerant data access. To ensure the availability of stand-alone DFS namespaces, you can create the root on a clustered file server.

**Alternate site selection based on cost** By default, if a target in the same site as the users fails, or if no same-site target exists, DFS refers clients to a random target. If you configure the optional site costing feature, DFS can use the site information in Active Directory to locate an alternate target that has the lowest-cost network connection as defined by the administrator in the Active Directory Sites and Services snap-in. After site costing is enabled, clients can access data on DFS targets over the optimum network connection.

**Load sharing** DFS provides a degree of load sharing by mapping a given logical name to shared folders on multiple file servers. For example, suppose that \\Company\StockInfo is a heavily used shared folder. By using DFS, you can associate this location with multiple shared folders on different servers, even if the servers are located in different sites.

**Intelligent client caching** When a user requests access to a target that is a part of a DFS namespace, a referral containing the target's information is cached on the client. The next time the client requires access to that portion of the namespace, the client uses the cached referral instead of obtaining a new referral, and connects directly to one of the target computers. For more information about client caching in DFS, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**Support for offline folders** If your clients are running Microsoft® Windows® XP or Windows Server 2003, you can make DFS link targets available offline by using the Offline Files feature. You can also use this feature to automatically cache programs so that users can run the programs locally instead of from the server. Using this feature for link targets that host applications can reduce network traffic and improve server scalability.

**Simplified maintenance** If a link has multiple link targets, administrators can perform preventive maintenance, repairs, or upgrades on servers by disabling referrals to specific link targets. While the referral to the link target is disabled, DFS automatically routes new requests to the remaining link targets that are online.

**Dynamic site discovery** DFS now supports dynamic site discovery. In Windows 2000, DFS maintained static site information. After the site information for a particular network resource was known, DFS used that information indefinitely, regardless of any changes in the site information of the resource. In Windows Server 2003, when you move a resource from one site to another, the information used by DFS converges to the new site information within 25 hours.

**Security integration** You do not need to configure additional security for DFS namespaces, because file and folder access is enforced by existing NTFS and share permissions on each link target. For example, a user navigating a DFS namespace is permitted to access only the files or folders for which he or she has appropriate NTFS or share permissions. If you use FRS to replicate content among multiple targets, FRS also replicates access control lists (ACLs) for each file and folder. For more information about DFS and FRS security, see "Planning DFS and FRS Security" later in this chapter.

# Evaluating Client and Server Compatibility

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you implement a DFS namespace, review the types of clients and servers in your organization to make certain that the servers can host targets and that the clients can access targets in the DFS namespace. For example, if you have UNIX clients, they cannot access the DFS namespace and must instead access the files by using the UNC path to the various file servers. Table 2.1 summarizes DFS interoperability.

**Table 2.1 DFS Interoperability**

| Platform | Act as DFS Clients? | Host DFS Roots? | Act as a Link Target? |
|---|---|---|---|
| Microsoft® Windows® Server 2003, Web Edition | Yes | Yes. Can host one stand-alone DFS root or one domain-based DFS root per server. | Yes |
| Windows Server 2003, Standard Edition | Yes | Yes. Can host one stand-alone DFS root or one domain-based DFS root per server. | Yes |
| Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition | Yes | Yes. Can host multiple stand-alone DFS roots and multiple domain-based DFS roots per server. | Yes |
| Windows XP | Yes | No | Yes |
| Windows Preinstallation Environment (WinPE) | Yes | No | No |
| Microsoft® Windows® 2000 Server family* | Yes | Yes, one stand-alone DFS root or domain-based DFS root per server. | Yes |
| Microsoft® Windows® 2000 Professional | Yes | No | Yes |
| Microsoft® Windows NT® Server 4.0 with Service Pack 6a | Yes | Yes, a single stand-alone DFS root per server. | Yes |
| Windows NT Workstation 4.0 with Service Pack 6a | Yes | No | Yes |
| Microsoft® Windows® Millennium Edition (Me) | Yes, client for stand-alone DFS included. Because Windows Me is designed specifically for home use, no domain-based DFS client is provided. | No | Yes |

| | | | |
|---|---|---|---|
| Microsoft® Windows® 98 | Yes, client for stand-alone DFS included; install the Active Directory client extension for Microsoft® Windows® 95 or Microsoft Windows 98 to access domain-based DFS namespaces. | No | Yes |

* Applies to general-purpose servers and Windows Powered Network Attached Storage solutions running Windows Server 2003.

**Note**

- The Active Directory client extension for Windows 95 or Windows 98 is available on the Windows 2000 operating system CD, or see the Active Directory Client Extensions link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

When evaluating client compatibility, review the following important considerations:

- Clients must be members of a domain before they can access a domain-based DFS namespace.

- Link targets can use other protocols, such as NetWare Core Protocol (NCP) for NetWare and Network Filesystem (NFS) for UNIX, but clients must have the appropriate redirector installed to access those link targets.

- In organizations that have a large number of domains, clients might have difficulty accessing link targets in other domains or forests. In addition, clients running Windows 98 might not be able to access any domain-based DFS namespace and might also have difficulty accessing links that point to other DFS namespaces. For more information about clients running Windows 98, see "Designing a DFS Namespace" later in this chapter.

# Choosing the DFS Namespace Type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When creating a DFS namespace, you create either a stand-alone DFS root or a domain-based DFS root. Table 2.2 describes the differences between domain-based DFS namespaces and stand-alone DFS namespaces.

**Table 2.2 How DFS Namespace Types Differ**

| Characteristic | Domain-based | Stand-Alone |
|---|---|---|
| Path to DFS namespace | *domainname**rootname*<br>*Netbiosdomainname**rootname*<br>*DNSdomainname**rootname* | *servername**rootname* |
| Group memberships required to create and administer namespaces | For DFS administrators who are not members of the Domain Admins group, it is recommended that you delegate permissions so that administrators can create new domain-based DFS namespaces. Administrators must also be members of the local Administrators group on each of the root targets to be able to add and delete links and add and remove the root targets. | DFS administrators must be members of the local Administrators group on the local server to create new stand-alone DFS roots and add or delete links. |
| Where DFS root information is stored | In Active Directory. DFS root information is replicated to all servers that host domain-based DFS roots. | In the registry of the root server. |
| DFS namespace size restrictions | Large domain-based DFS namespaces might cause significantly increased network traffic due to the size of the DFS Active Directory object. As a result, Microsoft recommends using fewer than 5,000 links in domain-based DFS namespaces. | The largest recommended namespace size for a stand-alone root is 50,000 links. |
| Supported methods to ensure DFS root availability | Create multiple DFS root targets in the same domain. | Create a stand-alone DFS root on a clustered file server. |
| Supported methods to ensure link target availability | Create multiple link targets and replicate files by using one of the following methods:<br>• Enabling FRS<br>• Copying files manually or by using scripts<br>• Using a third-party replication tool | Create multiple link targets and replicate files by using one of the following methods:<br>• Copying files manually or by using scripts<br>• Using a third-party replication tool |

**Note**

- For information about DFS namespace size restrictions, see "Reviewing DFS Size Recommendations" later in this chapter.

Use the following guidelines to choose a DFS namespace type.

Choose stand-alone DFS namespaces if:

- Your organization does not use Active Directory.
- You need to create a DFS namespace and are not part of the Domain Admins group, or company policy prevents you from delegating authority to manage a domain-based DFS namespace.
- You need to create a single namespace with more than 5,000 links. (If you can divide your links among two or more namespaces, domain-based DFS is an option.)
- You want to ensure the availability of the namespace by using a clustered file server.

Choose domain-based DFS namespaces if:

- You plan to use FRS to replicate data and you want to use the Distributed File System snap-in to configure and administer replication.
- You want to ensure the availability of the namespace by using multiple root targets.

As described in Table 2.2, you can increase the availability of roots and links in both types of DFS namespaces. For more information and specific guidelines about increasing the availability of roots and links, see "Increasing the Availability of DFS Namespaces" later in this chapter.

# Reviewing DFS Size Recommendations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you design your DFS namespace, use the guidelines in Table 2.3 to avoid potential performance problems that can arise when size recommendations are exceeded.

**Table 2.3 DFS Size Recommendations**

| Description | Recommendation* | Explanation |
|---|---|---|
| Path limit | Less than 260 characters | Win32 application programming interfaces (APIs) have a maximum path limit of 260 characters, so applications will fail when trying to access a namespace that goes beyond that limit. If the path length of the DFS namespace exceeds the Win32 API limit of 260 characters, users must map part of the namespace to a drive letter and access the longer namespace through the mapped drive letter. |
| Number of DFS roots per server running Windows Server 2003, Standard Edition | One | Windows Server 2003, Standard Edition is limited to one root per server. |
| Number of DFS roots per server running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition | Varies | There is no limit to the number of DFS roots you can create on a server running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. However, as you increase the number of roots per server, the Distributed File System service takes longer to initialize and uses more memory. |
| Number of root targets per domain-based DFS root | No fixed limit | If you do not enable root scalability mode, Microsoft recommends using 16 or fewer root targets to limit traffic to the server acting as the primary domain controller (PDC) emulator master. |
| Number of links per DFS namespace | 5,000 for domain-based DFS<br>50,000 links for stand-alone DFS | When the number of links exceeds the recommended limit, you might experience performance degradation when making changes to the DFS configuration. For stand-alone DFS, namespace initialization after server startup might also be delayed. |
| Size of each DFS Active Directory object (applies to domain-based DFS namespaces only) | 5 megabytes (MB) | The size of the DFS Active Directory object is determined by the number and path length of roots, links, comments, and targets in the namespace. Microsoft recommends using no more than 5,000 links in a domain-based namespace to prevent the DFS Active Directory object from exceeding 5 MB. Limiting the size of the Active Directory object is important because large domain-based DFS configurations can cause significantly increased network traffic originating from updates made to those roots, links, and targets. |

* The figures in this table are based on information gathered in a test environment. The numbers in an operational DFS configuration might exceed the numbers described here and still provide acceptable performance.

**Notes**

- You can check the size of an existing DFS namespace by using the following syntax in Dfsutil.exe:
- dfsutil /root:\\domainname\rootname /view (for domain-based DFS)
- dfsutil /root:\\servername\rootname /view (for stand-alone DFS)

- The command output displays the number of links and, for domain-based DFS namespaces, the size of the DFS Active Directory object (described as *blob size*).

If your organization plans to create large namespaces, there are a number of strategies you can implement to work within the size recommendations shown in Table 2.3.

**Keep comments to a minimum**
When you add a root target or link target in the Distributed File Systems snap-in, you can enter comments that describe the target. If you plan to create a large namespace, use minimal comments, if any, because they can increase the overall size of the namespace.

### Note

- Comments are visible only within the DFS administration tools, and they are not visible to users when they navigate the namespace.

**Create multiple namespaces**
If you need to create more than 5,000 links in a domain-based DFS namespace, you can create multiple DFS namespaces that meet the recommended sizes and then link them together. For more information about creating multiple namespaces, see "Designing a DFS Namespace" later in this chapter.

**Enable root scalability mode**
You enable root scalability mode by using the /RootScalability parameter in Dfsutil.exe, which you can install from the \Support\Tools folder on the Windows Server 2003 operating system CD. When root scalability mode is enabled, DFS root servers get updates from the closest domain controller instead of the server acting as the PDC emulator master. As a result, root scalability mode reduces network traffic to the PDC emulator master at the expense of faster updates to all root servers. (When you make changes to the namespace, the changes are still made on the PDC emulator master, but the root servers no longer poll the PDC emulator master hourly for those changes; instead, they poll the closest domain controller.) With this mode enabled, you can have as many root targets as you need, as long as the size of the DFS Active Directory object (for each root) is less than 5 MB. For more information about the 5-MB limit, see the entry describing the size of the DFS Active Directory object in Table 2.3 earlier in this chapter.

Do not use root scalability mode if any of the following conditions exist in your organization:

- Your namespace changes frequently, and users cannot tolerate having inconsistent views of the namespace.

- Domain controller replication is slow. This increases the amount of time it takes for the PDC emulator master to replicate DFS changes to other domain controllers, which, in turn, replicate changes to the root servers. Until this replication completes, the namespace will be inconsistent on all root servers.

### Note

- After you enable root scalability mode in a mixed domain, root servers running Windows Server 2003 can obtain updates from the closest domain controller; however, root servers running Windows 2000 Server still obtain updates from the PDC emulator master.

For information about installing Windows Support Tools, see "Install Windows Support Tools" in Help and Support Center for Windows Server 2003.

**Migrate root servers running Windows 2000 Server to Windows Server 2003**
Root servers running Windows Server 2003 do not add site information to the DFS Active Directory object. As a result, if all root servers run Windows Server 2003, DFS can store more root and link information to the DFS Active Directory object before reaching the recommended 5-MB limit. For more information about using a mix of root servers running Windows 2000 Server and Windows Server 2003, see "Designing a DFS Namespace" later in this chapter.

# Planning the Number of DFS Namespaces

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your next step is to plan the number of namespaces you want in your domain. For an Excel spreadsheet to assist you in documenting your namespace decisions, see "DFS Configuration Worksheet" (Sdcfsv_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

Medium organizations might require only a single namespace, while large organizations might need multiple DFS namespaces. You can determine the number of namespaces you require by reviewing the following factors.

**Scope of your domain**
If your domain has a broad scope — geographically, organizationally, or functionally — you should plan for multiple DFS namespaces so that administrators in the geographical, organizational, or functional departments can define their own namespaces. On the other hand, if the domain has a narrow scope geographically, organizationally, or functionally, you might want to define a single DFS namespace.

**Size of your DFS namespace**
If your DFS namespace exceeds the recommended number of links per namespace, as discussed earlier in "Reviewing DFS Size Recommendations" create multiple DFS namespaces, each of which does not exceed the recommended size. In this way, you can provide a single namespace to users by creating a single DFS namespace with links that point to other DFS namespaces. For more information about linking from one namespace to another, see "Designing a DFS Namespace" later in this chapter.

**Administrative boundaries**
How DFS namespaces are administered can also affect the number of DFS namespaces your organization requires. For example, your organization might have the following administrative boundaries:

- **Geographic.** Geographically diverse sites can each have an administrator who creates and manages the DFS namespace located in that site.
- **Departmental or group ownership.** Individual departments or groups can create and manage a DFS namespace that is used by members of that department or group.
- **Political.** Individual departments or groups can create and manage a DFS namespace that is used by members of that department or group.

If groups in your organization will create and manage their own DFS namespaces, you can build an extensive DFS namespace out of smaller, more focused DFS namespaces. One benefit of this method is that you can present specific DFS roots to some users as the true top of the hierarchy and also present a set of those DFS roots to other users as the only DFS links in a larger hierarchy. By using a hierarchy of DFS roots, you can scale the namespace as your organization grows and tailor the namespace for distributed management.

For more information about linking from one namespace to another, see "Designing a DFS Namespace" later in this chapter.

**DFS namespace depth**
Limit the depth of DFS namespaces to 260 characters. The 260-character limit includes the fully qualified domain name (FQDN) of the domain hosting the DFS root as well as the DFS root name. If you exceed this limit, applications will fail when trying to access the namespace. To work around this issue, users must map part of the namespace to a drive letter and then access the longer namespace through the mapped drive letter.

# Developing Root and Link Naming Standards

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you roll out DFS, you have the opportunity to implement consistent namespace designs. Developing naming standards first — and ensuring that you adhere to the naming standards during implementation — makes it easier to use and manage any DFS namespace, both from a user perspective and an administrative perspective. Even if you do not expect to implement DFS until a later phase of your Windows Server 2003 deployment, it is important to begin thinking about namespace design early in the planning process. For an Excel spreadsheet to assist you in documenting your DFS namespace design decisions, see "DFS Configuration Worksheet" (Sdcfsv_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**Creating DFS Root Names**

A DFS root name, significant primarily to users, is the point beyond the server name or domain name that is at the top of the hierarchy of the logical namespace. Standardized and meaningful names at this level are very important, especially if you have more than one DFS namespace in a domain, because the DFS root name is where users enter the namespace. The contents of a DFS namespace must be as clear as possible to the users so that they do not follow the wrong path, possibly across expensive WAN connections, and have to backtrack.

When creating the DFS root names for servers that will contain multiple roots, review the following restrictions:

- Each root requires its own shared folder.
- When you create a domain-based DFS root, the share name in the UNC path *servername**sharename* must be the same name as the DFS root name in *domainname**rootname*. For example, if you want to create a domain-based DFS root \\Reskit.com\Public on Server1, the UNC path to the shared folder must be \\Server1\Public.
- A root cannot be nested within another root. For example, if C:\Root is a shared folder that uses the share name Public, and you use this shared folder as a stand-alone DFS root (*servername*\Public) or domain-based DFS root (*domainname*\Public), you cannot create another root in the folder C:\Root\Software. Similarly, if you create a root by using the root folder C:\, you cannot create another root at C:\Root.
- On server clusters, do not create clustered DFS roots that have the same name as nonclustered DFS roots or shared folders.
- Shared folders on domain controllers must not have the same name as any domain-based DFS roots in the domain. If they do, clients who try to access the shared folder on the domain controller are redirected to the domain-based DFS root. For example, if Reskit.com has a domain controller named DC1 that contains a shared folder named Tools (\\DC1\Tools), do not create a domain-based DFS namespace using a root named Tools (\\Reskit.com\Tools). Otherwise, when users attempt to access \\DC1\Tools, they are redirected to \\Reskit.com\Tools.

**Creating DFS Link Names**

A DFS link is a component in a DFS namespace that lies below the root and maps to one or more link targets. Because DFS link names are exposed to users, it is important to develop standardized, meaningful names for DFS links. Another important design goal is to develop a DFS namespace that provides intuitive navigation within the hierarchy that the namespace represents. Keep in mind that comments entered in the Distributed File System snap-in are not visible to users. For this reason, the namespace must be as clear as possible at all levels.

Designing a clear naming scheme is even more important for a DFS namespace than for a physical namespace, because a user might jump to a shared folder on a different file server when he or she selects a link in the DFS namespace. As a result, a session has to be set up with that physical server (if one does not already exist), which might delay access. Therefore, you want to minimize the number of times that users traverse a wrong path. Clear and meaningful naming standards can help.

Try to keep links at the same level in the DFS namespace consistent in context. For example, you probably would not want to have links named New York, Seattle, and Milan mixed with other links named Sales, Marketing, and Consulting. To help you create a consistent namespace, DFS supports adding one or more folder names to the link name so that you can create a meaningful hierarchy of link names. In the previous example, you could create links such as the following:

- Branches\New York, Branches\Seattle, and Branches\Milan

- Departments\Sales, Departments\Marketing, and Departments\Consulting

When users browse this namespace, they will see a folder called Branches and another called Departments, which they can use to navigate to folders for branch offices (New York, Seattle, and Milan) and departments (Sales, Marketing, and Consulting).

**Note**

- If you want to encourage users to access the DFS namespace instead of going to individual servers, you can use a dollar sign ($) at the end of the shared folder name to hide it from casual browsers. The shared folder will still appear in the DFS namespace with the link name you specify. Doing so prevents users from accessing the shared folders by specifying individual server names. Instead, users must access the shared folders by using the namespace, which enables DFS to load share requests across multiple link targets and allows clients to be directed to another link target if the previously used target is unavailable.

# Designing a DFS Namespace

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you design one or more DFS namespaces for your organization, you need to make a number of decisions about the structure and capacity of the namespaces. For an Excel spreadsheet to assist you in documenting your DFS namespace decisions, see "DFS Configuration Worksheet" (Sdcfsv_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**Determining Who Can Manage the Namespace**

You can delegate administrative authority to individual users so that they can manage a DFS namespace. Table 2.4 describes the permissions and group memberships that you must delegate before users can manage a namespace on a member server. Administering DFS namespaces on a domain controller or configuring FRS replication requires membership in the Domain Admins group.

**Table 2.4 Permissions or Group Memberships Required to Administer DFS Namespaces**

| Task | Permissions or Group Membership Required |
|---|---|
| Creating or removing a domain-based DFS root on a member server. | One of the following:<br>• Membership in the Domain Admins group.<br>• Full Control permission on the DFS-Configuration container in Active Directory and membership in the local Administrators group on the root server. |
| Adding or removing a root target from an existing domain-based DFS root on a member server. | One of the following:<br>• Membership in the Domain Admins group.<br>• Full Control permission on the DFS-Configuration container in Active Directory and membership in the local Administrators group on the root server. |
| Creating or deleting a stand-alone DFS root on a member server. | Membership in the local Administrators group on the root server. |
| Adding a link to a domain-based DFS namespace or adding a target to an existing link on a member server. | Membership in the local Administrators group on each of the root target servers. |
| Removing a link from a domain-based DFS namespace or removing a target from an existing link on a member server. | Membership in the local Administrators group on each of the root target servers. |
| Changing root-related or link-related information, such as comments, referral status, and cache limits on a member server. | Membership in the local Administrators group on each of the root target servers. |
| Performing any of the tasks in this table on a domain controller. | Membership in the Domain Admins group. |
| Enabling replication on links in a domain-based DFS namespace. | • Membership in the Domain Admins group. |

You can also limit delegated authority to just one domain-based DFS namespace on a member server by granting a user or group Full Control permission on the DFS root object contained in the DFS-Configuration container. Doing so allows the administrator to add or remove root targets from a specific namespace. For more information about how to delegate

permission to manage a DFS namespace, see "Deploying DFS" later in this chapter.

**Selecting Servers to Host Roots**

Table 2.5 describes the guidelines for servers hosting stand-alone DFS roots and domain-based DFS roots.

**Table 2.5 Guidelines for Servers That Host DFS Roots**

| Server Hosting Stand-Alone DFS Roots | Server Hosting Domain-based DFS Roots |
|---|---|
| • Must contain an NTFS volume to host the root.<br>• Can be a member server or domain controller.<br>• Can be a general-purpose server or Windows Powered Network Attached Storage.<br>• Can be a clustered file server. | • Must contain an NTFS volume to host the root.<br>• Must be a member server or domain controller in the domain in which the DFS namespace is configured. (This requirement applies to every root target for a given domain-based DFS namespace.)<br>• Can be a general-purpose server or Windows Powered Network Attached Storage.<br>• Cannot be a clustered file server unless you host the domain-based DFS root on the local storage of a node in the server cluster. |

The following sections provide other factors to consider when selecting root servers.

**Restrictions for servers running Windows Server 2003, Standard Edition**

If you are using servers running Windows Server 2003, Standard Edition, you can create only one DFS root per server, which means that you need one server for each root you plan to host.

**Root servers that have the RestrictAnonymous registry value**

Before you create a DFS root on a server, verify that the RestrictAnonymous registry value is not set on the server. This registry value restricts anonymous access and causes DFS referral failures. For more information about this registry value, see "Planning DFS and FRS Security" later in this chapter.

**Root server performance**

When evaluating the hardware specifications of the servers that host roots, note that clients access the root server to get referrals, and then the clients cache the referrals locally. Therefore, root servers do not typically experience high CPU usage. However, as the size of the namespace grows, the DFS service uses more memory, so consider using more than the minimum recommended RAM for servers that host large DFS namespaces and for servers that host multiple DFS namespaces. For more information about choosing RAM and CPU speed for file servers, see "Determining RAM and CPU Specifications" later in this chapter.

**Important**

- Servers running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition can host multiple roots of any type (stand-alone or domain-based). However, as you increase the number of roots (namespaces) on a server, you also increase the number of namespaces that will be unavailable if the server fails. Therefore, devise a plan that allows you to increase the availability of the namespaces. For more information, see "Increasing the Availability of DFS Namespaces" later in this chapter.

**Hosting roots on domain controllers**

When deciding whether to host a DFS root on a domain controller, consider the following factors:

- Only members of the Domain Admins group can manage a DFS namespace hosted on a domain controller.
- If you plan to use a domain controller to host a DFS root, the server hardware must be sized to handle the additional load. As described earlier, root servers that host large or multiple namespaces require additional memory. For information about capacity planning for domain controllers, see "Planning Domain Controller Capacity" in *Designing and Deploying Directory and Security Services* of this kit.

**Using root servers running Windows 2000 Server and Windows Server 2003**

If you plan to host a domain-based DFS root on servers running a mix of Windows 2000 and Windows Server 2003, you need

to understand how DFS handles site information in each of these operating systems. These differences are important because Windows Server 2003 does not store site information in the DFS Active Directory object; instead, root servers running Windows Server 2003 obtain site information directly from Active Directory. If you have root servers running Windows 2000 Server, those servers try to obtain site information from the DFS Active Directory object, and unless you use Dfsutil.exe to manually update the site info in the DFS Active Directory object, the root servers running Windows 2000 Server might provide referrals that lead to a target outside of the client's site.

Table 2.6 describes how DFS root servers handle site information.

**Table 2.6 How Root Servers Handle Site Information**

| Site Difference | Windows 2000 Server | Windows Server 2003 |
|---|---|---|
| Where DFS stores and retrieves site information for root and link targets | DFS stores a copy of site information for root and link targets in the DFS object in Active Directory. | DFS uses IP addresses of root and link targets to obtain site information directly from Active Directory. By default, DFS does not store site information in the Active Directory object. |
| Characteristic of site information | Static | Dynamic |
| Method for updating site information after moving a link target to a different site | Remove the link target from the namespace, and then add it back. | By default, site information automatically updates every 25 hours. |
| How root servers use site information for referrals | Root servers running Windows 2000 Server use the link target's site information only if the link target was created by using Windows 2000 Server. If the link target was created by using Windows Server 2003, no site information is stored, which means that the referral could lead to a target outside of the client's site. | Root servers running Windows Server 2003 ignore any site information in the Active Directory object; instead, they use site information directly from Active Directory. |

Running Windows Server 2003 on every root server is recommended for a number of reasons:

- Site information is always up to date because DFS obtains site information directly from Active Directory instead of storing a copy of site information in the DFS Active Directory object.
- The DFS Active Directory object can hold additional root and link targets because it does not contain site information.

If you want referrals from root servers running Windows 2000 Server to be ordered according to site information, you can use the /UpdateWin2kStaticSiteTable parameter in Dfsutil.exe to update the static site information for all root and link targets in the DFS Active Directory object. If you plan to use this parameter, review the following issues:

- Using this parameter increases the size of the DFS Active Directory object, possibly making it exceed the 5-MB recommended size limit.
- You need to run this parameter each time you want to update site information.
- Root servers running Windows Server 2003 continue to get site information directly from Active Directory, and they ignore all site information in the DFS Active Directory object.

When all root servers are running Windows Server 2003, you can use the /PurgeWin2kStaticSiteTable parameter in Dfsutil.exe to remove site information from the DFS Active Directory object, providing you additional space for creating root and link targets.

**Important**

- If you are using a mix of root servers running Windows 2000 Server and Windows Server 2003, use the version of the Distributed File System snap-in available in either Windows Server 2003 or the Windows Server 2003 Administration Tools Pack to manage the namespace. Do not use the version of the Distributed File System snap-in available in Windows 2000 Server.

**Determine the Root and Link Referral Time to Live Values**

You can use the Distributed File System snap-in to specify the length of time that clients cache referrals to DFS root targets and link targets by adjusting the Time to Live value for each DFS root or link. When a client receives a referral to a target, the client will continue to access that target (either a particular root target or link target) until one of the following occurs: the client computer is restarted, the user clears the cache, or the Time to Live value for the root or link expires. The client will then obtain a new referral the next time it attempts to access the target. However, if the client continues to access the root or link within the Time To Live value, the Time to Live value is renewed each time and the client never requests a new referral. Clients that resume from hibernation do not request new referrals either.

Windows Server 2003 uses a default Time to Live value of 300 seconds (5 minutes) for DFS roots and 1800 seconds (30 minutes) for DFS links. The default Time to Live values work well in organizations where the namespace changes frequently and it is important that clients have up-to-date referrals. When using these values, make sure that your network has adequate bandwidth and your root servers have adequate resources to handle the traffic generated when clients contact the root servers for referrals.

Consider increasing the Time to Live value in the following situations:

- An established root or link has a single link target. In this case, you do not need to load balance requests among multiple servers, so it is appropriate to have a longer Time to Live value for that root or link.

- Your organization has clients in multiple Active Directory sites, but you do not have root servers in all sites. In this situation, you can increase the Time to Live value for the root to ensure that the client does not have to go out of its site for a root referral if more than 5 minutes have passed since the client last accessed the namespace. If your root information is static, you can set this value to be several hours.

For more information about Time to Live values, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**Determining the Contents of the Root Folder**

The root folder of a DFS namespace is a launch point into the namespace — that is, a placeholder for the namespace. Keep the root folder as uncluttered as possible. For example, you might place a single file in the root folder (a readme file) that describes the contents and purpose of the namespace. When users access the namespace, they will see the readme file and the links that point to one or more link targets. For more information about the root folder, see "Reviewing DFS Terminology" earlier in this chapter.

**Choosing Shared Folders to Add to the Namespace**

Add shared folders to your namespace only if they are well established and unlikely to be retired in the near future. If you have targets whose underlying physical name is dynamic, include them in the DFS namespace only if you can tolerate the added administrative overhead or develop automated scripts to update the DFS links.

To secure files and folders, use NTFS as the underlying file system for shared folders that you add to a DFS namespace. You can add shared folders residing on file allocation table (FAT) volumes in Windows-based servers and shared folders residing on servers running other operating systems, such as UNIX or Novell NetWare servers; however, you cannot use NTFS security features to secure these folders. In addition, if you plan to use FRS to replicate the contents of shared folders, you must use NTFS as the file system.

To help users find data in DFS namespaces more easily, you can publish the UNC path for DFS links as shared folders in Active Directory. For example, if an administrator of the Reskit.com domain wants to publish the DFS link where users can install applications, the administrator can specify \\Reskit.com\Public\Apps as the path to the published folder. The administrator can also specify key words for that shared folder, such as "applications" or "Office." Domain users in a forest can locate the applications by using the Active Directory search tool in My Network Places to query for the application folder by name or key words.

For information about publishing shared folders in Active Directory, see "Publish a shared folder" in Help and Support Center for Windows Server 2003.

**Reviewing Rules for Creating Links**

When designing the namespace, review the following rules regarding link creation.

**Creating new links under existing links**
You cannot create a new link under an existing link. To create a hierarchy of links, specify additional folder names within the link name. For example, instead of creating a link called Groups, create links called Groups\Development, Groups\Management, Groups\Administrative, and so forth.

**Creating links to shared folders in other domains or forests**
If a client can access a link target in another trusted domain or trusted forest by using the target's UNC path, the client can also access the link target by using its DFS path, but only if the list of domains fits into the client's cache, which is 4 KB by default (roughly 2000 characters). If the list of domains is too large to fit into the 4-KB cache, the following actions occur:

- Clients running Windows 98 cannot access any domain-based DFS namespaces. To notify you of this issue, DFS writes an entry with the ID 14537 in the system log in Event Viewer on the domain controller that enumerates the domains.

- Computers running Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003 automatically increase their cache size to accept the list of domains, up to a maximum of 56 KB.

If the list of domains exceeds 56 KB, DFS puts as many domains in the cache as it can until the cache reaches 56 KB. DFS then writes an entry with the ID 14536 in the system event log in Event Viewer to notify you of this issue. When populating the cache, DFS gives preference to local and explicitly trusted domains by filling the cache with their names first. Consequently, by creating explicit trust relationships with domains that host important DFS namespaces you can minimize the possibility that those domain names might be dropped from the list that is returned to the client.

**Important**

- To make sure that clients can access link targets in other trusted domains or trusted forests, you must use DNS names for all link targets and configure DFS to use fully qualified domain names in referrals. For more information, see article Q244380, "How to Configure Dfs to Use Fully Qualified Domain Names in Referrals." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Creating links to mounted drives**
A mounted drive is a local volume attached to an empty folder on an NTFS volume. If C:\Root is a DFS root folder on Server1 (\\Server1\Root), and you create a folder C:\Root\Link1, where Link1 is a mounted drive (for example, Link1 points to drive D), you cannot create a link named Link1 in that DFS namespace.

**Creating links to different namespaces**
Windows Server 2003 supports creating links that point to other DFS namespaces. Linking to other namespaces is common in organizations that want to combine the availability benefits of domain-based DFS namespaces with the scalability of stand-alone DFS namespaces. For example, if an organization needs to create 10,000 links but does not want to divide these between two domain-based DFS namespaces, the organization can take the following steps:

1. Create a stand-alone DFS namespace with 10,000 links.

2. Create a domain-based DFS root.

3. Under the domain-based DFS root, create a link that points to the stand-alone DFS namespace.

When linking to other namespaces, you must follow these guidelines to make certain that clients can be redirected properly if a target is unavailable:

- If you plan to specify a domain-based DFS namespace as a link target (either the root or a link within that namespace), you cannot specify alternate link targets. (Windows Server 2003 enforces this restriction.)

- If you plan to specify a stand-alone DFS namespace as a link target (either the root or a link within that namespace), you can specify alternate link targets that are either stand-alone DFS roots or links within the stand-alone DFS namespace. Do not specify domain-based DFS roots or shared folders as alternate targets.

**Important**

- The DFS tools do not prohibit you from specifying domain-based DFS roots or shared folders as alternate targets. Therefore, follow these guidelines carefully.

When linking to other namespaces, review the following restrictions:

- A DFS path can consist of no more than eight hops through other DFS namespaces.
- Clients running Windows 98 might not correctly access links pointing to other DFS namespaces. Windows 98–based clients can only access the following types of links to other namespaces:
    - A link in a stand-alone DFS namespace that points to a stand-alone DFS root or link.
    - A link in a domain-based DFS namespace that points to a stand-alone DFS root. (This works only if the client has the latest Active Directory client installed, as described in article Q323466, "Directory Services Client Update for Windows 95 and Windows 98." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For additional rules for specifying multiple link targets, see "Choosing an Availability Method for Data in Link Targets" later in this chapter.

# Increasing the Availability of DFS Namespaces

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you create your initial namespace design, you need to review that design to determine whether part or all of the namespace needs to be available at all times. You can ensure DFS namespace availability at both the root and link target level. This helps prevent the situation in which the target servers are up and running but users cannot access them by using the namespace.

To increase the availability of DFS namespaces, take the following steps:

1. Choose an availability method for DFS roots.
2. Choose an availability method for data in link targets.
3. Determine where to place multiple targets.
4. Choose a replication method.

The following sections describe each of these steps.

# Choosing an Availability Method for DFS Roots

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you plan to create DFS namespaces that must be highly available, such as those that provide access to business-critical data, the availability method you choose depends on the root type.

**Stand-alone DFS roots**
You ensure the availability of a stand-alone DFS root by creating it on the cluster storage of a clustered file server by using the Cluster Administrator snap-in. For more information about making stand-alone DFS roots highly available, see "Increasing Data Availability by Using Clustering" later in this chapter.

**Domain-based DFS roots**
You ensure the availability of domain-based DFS roots by creating multiple root targets on nonclustered file servers or on the local storage of the nodes of server clusters. (Domain-based DFS roots cannot be created on cluster storage.) All root targets must belong to the same domain. To create root targets, use the Distributed File System snap-in or the Dfsutil.exe command-line tool. (For information about choosing servers to host root targets, see "Designing a DFS Namespace" earlier in this chapter.)

To ensure the availability of domain-based DFS roots, you must have at least two domain controllers and two root targets within the domain that is hosting the root. If you have only one domain controller, and it becomes unavailable, the namespace is inaccessible. Similarly, if you have only a single root target, and the server hosting the root target is unavailable, the namespace is also unavailable.

### Note

- If you plan to use more than 16 root targets, or if you have root servers in remote sites that connect to the PDC emulator master across slow links, consider enabling root scalability mode. For more information about root scalability mode, see "Reviewing DFS Size Recommendations" earlier in this chapter.

After you determine which roots need to be highly available, document your decisions. For an Excel spreadsheet to assist you in documenting the high-availability requirements of DFS roots, see "DFS Configuration Worksheet" (Sdcfsv_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

# Choosing an Availability Method for Data in Link Targets

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you design your namespace, you need to identify link targets whose data must be highly available. There are two ways to increase the availability of data in link targets:

- Create a single link that points to a link target on a clustered file server.
- Create multiple link targets and replicate content among them.

You can create link targets that point to clustered file servers in both types of namespaces. However, if you want to replicate content among multiple link targets, the type of namespace determines your replication options.

**Using replication in stand-alone DFS namespaces**
In a stand-alone DFS namespace, you must replicate the files by copying them manually, using scripts, using Robocopy.exe, which is available in the *Windows Server 2003 Deployment Kit,* or by using other replication tools. The Distributed File System snap-in does not provide a user interface for configuring FRS replication in stand-alone DFS namespaces. To configure replication manually, consult the documentation supplied with your replication tools.

**Using replication in domain-based DFS namespaces**
The Distributed File System snap-in in Windows Server 2003 provides a user interface for creating the FRS topology and schedule on servers running Windows Server 2003. If you do not want to use FRS in a domain-based DFS namespace, you can replicate files by copying them manually or by using third-party replication tools.

For more information about replication, see "Choosing a Replication Method" later in this chapter.

**Note**

- The Distributed File System snap-in is also part of the Windows Server 2003 Administration Tools Pack; you can install this pack on computers running Windows XP with Service Pack 1 (SP1) or later and create FRS schedules and topologies on remote servers running Windows 2000. For more information, see article Q304718, "Administering Windows 2000-Based and Windows Server 2003-Based Computers Using Windows XP Professional-Based Clients." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

If you plan to use multiple link targets to ensure data availability, you need to configure link targets correctly. Each link can have targets that correspond to only one of the following options:

- One or more shared folders.
- One or more stand-alone DFS paths anywhere in the stand-alone DFS namespace, including the root.
- A single domain-based DFS path anywhere in the domain-based DFS namespace, including the root.

**Important**

- The DFS tools do not prohibit you from creating links that conflict with these guidelines. Therefore, follow these guidelines carefully.

After you determine which link targets need to be highly available, document your decisions. For an Excel spreadsheet to assist you in documenting which links require multiple link targets, see "DFS Configuration Worksheet" (Sdcfsv_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

# Placing Multiple Targets

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you evaluate where to place multiple targets, you should plan to place at least one target in the same Active Directory site where users access the data. Doing so enables clients to use fast, inexpensive bandwidth to access the target. Use multiple same-site targets to ensure namespace availability in each site and to avoid the need to use expensive bandwidth if one of the same-site targets fails or is taken offline. Using multiple same-site targets also provides load sharing among the targets in the site.

After you determine where to place multiple targets, you also need to consider where clients will be redirected if the primary target is unavailable. DFS supports three methods of target selection, which are described in the following sections. These methods apply to both stand-alone and domain-based DFS namespaces. After you determine the method of target selection for each link or root, document your decisions. For an Excel spreadsheet to assist you in documenting the target selection methods, see "DFS Configuration Worksheet" (Sdcfsv_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**Important**

- If you have a mix of root servers running Windows 2000 Server and Windows Server 2003, target selection is random if a root server running Windows 2000 Server provides a referral for a link target created in Windows Server 2003, regardless of which target method you configure. For more information about target selection, see "Designing a DFS Namespace" earlier in this chapter.

**Default target selection**

If the last (or only) target in an Active Directory site fails or is taken offline, DFS directs clients to another target in the same site, if a target is available. Clients are directed to a random target if no same-site targets are available. DFS does not consider bandwidth cost, connection speed, or the target server's processing load when choosing the random target.

**Restricted same-site target selection**

By using Dfsutil.exe /InSite parameter, you can limit client access to only those targets that are in the same site as the client. You enable this feature on a DFS root or on individual links in the namespace. If you enable this feature on the root, referrals for any link in the namespace return only targets that are in the same site as the client. If this functionality is disabled on the root, the individual settings on each link are used. When using this feature, plan to have at least one target (or two targets, for fault tolerance) in every site, and plan to monitor servers to make sure they are online and accessible. If no same-site targets exist, clients in that site are denied access to the data in the namespace.

**Note**

- The /InSite parameter takes effect after you stop and restart the Distributed File System service on each root server or when the service reads the DFS metadata, which happens every hour by default.

**Least expensive target selection**

If you create a stand-alone or domain-based DFS root on a server running Windows Server 2003, and the domain controller acting as the Intersite Topology Generator (ISTG) is also running Windows Server 2003, you can use the /SiteCosting parameter in Dfsutil.exe to enable DFS to choose an alternate target based on connection cost if no same-site targets are available. Windows Server 2003 uses the site and costing information in Active Directory to determine whether sites are linked by inexpensive, high-speed links or by expensive WAN links.

Site costing is not available in the following situations:

- When a stand-alone DFS namespace is hosted on a server that is not part of any domain.

- When the closest domain controller acting as the ISTG is running Windows 2000 Server.

  (This situation occurs when there are only domain controllers running Windows 2000 Server in the site of the DFS root server, or when there are no domain controllers in the site of the DFS root server and the closest site with at least one domain controller has only Windows 2000 Server domain controllers in that site. The closest site is defined in Active Directory.)

If you plan to enable site costing, review the following:

- You can enable site costing on a per-namespace basis.

- When the domain controller acting as the ISTG is running Windows 2000 Server, DFS uses the default target selection, described earlier in this section. If domain controllers running Windows 2000 Server and Windows Server 2003 exist in a site, the ISTG role is automatically given to the domain controller running Windows Server 2003.

For more information about defining sites in Active Directory, see "Designing the Site Topology" in *Designing and Deploying Directory and Security Services* of this kit.

# Choosing a Replication Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you plan to use multiple targets and you want to synchronize data in those targets, you need to choose a replication method. You have several methods for replicating data:

- A manual replication method (such as using the command-line tool Robocopy.exe, which is available in the *Windows Server 2003 Deployment Kit*)
- FRS
- A third-party replication tool

It is not mandatory to use FRS to keep targets synchronized. In fact, by default, FRS is not enabled for DFS targets in domain-based DFS namespaces. However, in general you do want to make sure that the underlying shared folders that correspond to DFS links and targets are synchronized to present the same data to users, regardless of the folder that they want to access.

The following sections describe when to use manual replication or FRS. For an Excel spreadsheet to assist you in documenting the replication method for each target, see "DFS Configuration Worksheet" (Sdcfsv_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit). For information about using third-party replication tools, consult the documentation provided by your software vendor.

**When to Use Manual Replication**

If the data in the shared folder is static, you can replicate the data by doing a one-time copy of the data to a target in the replica set. Even if the data in the shared folder is dynamic but changes infrequently, you might want to keep the targets synchronized by downloading the initial copies over the network and then manually updating them with changes. You must use manual replication if you plan to use a stand-alone DFS namespace or if one or more link targets for a particular link do not run Windows 2000 Server or Windows Server 2003.

**When to Use FRS**

FRS works by detecting changes to file and folders in a replica set and replicating those changes to other file servers in the replica set. When a change occurs, FRS replicates the entire file, not just the changed bytes. You can use FRS only if you are using a domain-based DFS namespace, and only servers running a Windows 2000 Server or Windows Server 2003 operating system can be part of a replica set. In addition, all replica sets must be created on NTFS volumes.

FRS offers a number of advantages over manually copying files, including:

**Continuous replication** FRS can provide continuous replication, subject to server and network loads. When a file or folder change occurs, and the file or folder is closed, FRS can begin replicating the changed file or folder to outbound partners (that is, the replica members that will receive the changed file or folder) within five seconds.

**Replication scheduling** You can schedule replication to occur at specified times and durations as needed by your organization. Scheduling replication to occur during evening hours, for example, can reduce the cost of transmitting data over expensive WAN links. Replicating data during off-hours also frees up network bandwidth for other uses.

**Compression** To save disk space, FRS compresses files in the staging directory by using NTFS compression. Files sent between replica members remain compressed when transmitted over the network.

**Authenticated RPC with encryption** To provide secure communications, FRS uses Kerberos authentication protocol for authenticated remote procedure call (RPC) to encrypt the data sent between members of a replica set.

**Fault-tolerant replication path** FRS does not rely on broadcast technology, and it can provide fault-tolerant distribution via multiple connection paths between members. If a given replica member is unavailable, the data will flow via a different route. FRS uses logic that prevents a file from being sent more than once to any given member.

**Conflict resolution** FRS can resolve file and folder conflicts to make data consistent among the replica members. If two identically named files on different servers are added to the replica set, FRS uses a "last writer wins" algorithm, which means that the most recent update to a file in a replica set becomes the version of the file or folder that replicates to the other members of the replica set. If two identically named folders on different servers are added to the replica tree, FRS identifies the conflict during replication and renames the folder that was most recently created. Both folders are replicated to all servers in

the replica set, and administrators can later merge the contents of two folders or take some other measure to reestablish the single folder.

**Replication integrity** FRS relies on the update sequence number (USN) journal to log records of files that have changed on a replica member. Files are replicated only after they have been modified and closed. As a result, FRS does not lose track of a changed file even if a replica member shuts down abruptly. After the replica member comes back online, FRS replicates changes that originated from other replica members, as well as replicating local changes that occurred before the shutdown. This replication takes place according to the replication schedule.

When you use FRS, the link targets might not always be completely synchronized. As a result, one client's view of a link in a DFS namespace can be different from another client's view of the same link. This inconsistency can happen when clients have been referred to different link targets in the namespace. Link targets do become consistent with time, but you might experience temporary inconsistencies due to replication latency when updates are occurring. For more information about using FRS, see "Choosing an Availability Strategy for Business-Critical Data" later in this chapter.

FRS is typically used to keep link targets synchronized. It is also possible to put files and folders directly in a domain-based DFS root target and then enable replication on the root so that the files and folders are replicated to all root targets. However, avoid enabling replication on domain-based DFS roots for the following reasons:

- Morphed folders can occur under the DFS root folder. Morphed folders occur when two or more identically named folders on different servers are added to the replica tree. FRS identifies the conflict during replication, and the receiving member protects the original copy of the folder and renames (*morphs*) the later inbound copy of the folder. The morphed folder names have a suffix of "_NTFRS_xxxxxxxx," where "xxxxxxxx" represents eight random hexadecimal digits.

  Morphed folders occur in replicated roots for the following reason: When you create a link in the namespace, DFS creates a link folder under each root folder on every root server. For example, if you add 1,000 links to a namespace, DFS creates a link folder under the DFS root folder for each of those 1,000 links on every root server. When you enable FRS replication on the root, FRS attempts to replicate its local copy of the folder structure to every root server. Because each root server has a local copy of the same folder structure as the incoming changes, FRS identifies the duplicate folder names and renames the folders that were most recently created. FRS then replicates all morphed folders to all root targets in the replica set.

  **Note**

  - If morphed folders do occur, you must use the /RemoveReparse:<DirectoryName> parameter in Dfsutil.exe to delete each morphed folder. For more information about morphed folders, see "Choosing an Availability Strategy for Business-Critical Data" later in this chapter.

- When adding a new root target to an FRS replicated root, you cannot replicate the contents of individual folders in the root based on business priority. Instead, the entire contents of the root are replicated to the new root target. On the other hand, if you enable replication only on individual links, you are creating multiple replica sets, which allows you to enable replication on the most important links first and then enable replication on the links in the namespace as desired.

- You cannot take individual root targets offline. For example, if you are adding a new root target, users who are referred to the new member might see incomplete data until replication is complete. On the other hand, if you enable replication on individual links, you can take a new link target offline while the initial replication takes place or whenever you want to restrict access to a particular link target.

# Example: An Organization Designs a DFS Namespace

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization with 35,000 users in one site designs a stand-alone DFS namespace to provide unified access to 1.4 terabytes (TB) of both business-critical and nonessential software. The team responsible for designing and deploying the namespace chooses a stand-alone DFS namespace for the following reasons:

- The team is not a member of the Domain Admins group, and corporate security policy prevents them from becoming members of this group.
- Corporate security policy also prohibits the team from creating a domain to host a domain-based DFS namespace.

When designing the namespace, the team needs to provide access to the following types of software:

- Business-critical software and operating systems
- Previous (archived) versions of software and operating systems
- Multimedia software that runs from the network
- Training courseware

To ensure the availability of the namespace, the team creates the stand-alone DFS root, \\software\public, on a two-node clustered file server. Each node has two Pentium III 1.26-gigahertz (GHz) CPUs and 256 MB of RAM. The root server does not provide any other services.

When evaluating users' availability requirements for the different types of software, the team determines that the business-critical software and operating systems must be available at all times. In addition, because this is the most frequently accessed software, the team wants good response times. To provide the desired availability and performance, the team uses four servers as link targets. These servers each have two Pentium III 733-MHz CPUs and 256 MB of RAM.

Although the multimedia software is not business critical, the team uses two servers as link targets for this software to improve server response times, because the client portion of the multimedia software accesses files from the server. The team does not use redundant servers for the remaining software, because it is not business critical, and users can tolerate temporary downtime of those servers.

Figure 2.8 describes the DFS namespace for this example.

**Figure 2.8 A Stand-Alone DFS Namespace Used for Software Installation**




When updating software on the \\archsvr, \\mmsoft, and \\trainsvr servers, the team adds new files directly to the servers. To update software on the business-critical \\software servers, the team copies the new files to a staging server, \\stagesvr. The team uses Robocopy scripts on \\stagesvr to copy the data to the \\software servers. By making changes only at the staging server, the team makes sure that no accidental changes are made on the \\software servers. The team also uses the staging server as the backup source, because it contains source files for a number of servers in addition to the \\software servers. The backup team backs up the other servers individually.

To prevent the staging server from becoming overloaded, the team does not make the staging server part of the namespace, so users do not directly access the server. The team also increases the hardware specifications of this server to meet the increased disk and CPU utilization required for copying files to the \\software servers and to support backups. The hardware configuration has four Pentium III Xeon 700-MHz processors and 1 gigabyte (GB) of RAM.

While deploying the namespace in the production environment, the team learned a number of valuable lessons that can help other organizations as they test and deploy their namespaces:

- When setting permissions on the clustered root server, follow the guidelines in "Planning Cluster Security" later in this chapter. Test the clustered root server to verify that permissions work correctly after a failover occurs.

- Be sure to change the Cluster service password as specified by your organization's password policy. Windows Server 2003 does not send notification that the password is about to expire. If the password expires, problems can occur at failover. For more information about Cluster service passwords, see "Change the Cluster service account password" in Help and Support Center for Windows Server 2003.

- Stagger new software announcements so that users do not overload the servers while trying to install new software. For example, send e-mail notification to a percentage of users each day.

# Planning File Server Availability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you plan for file server availability, begin by ensuring the uptime of the physical file server. Next, determine whether the file server contains business-critical data. If it does, you need to decide whether to use an availability strategy, such as replication or clustering, to ensure the availability of the data. If data on the file server is temporary or not business critical, continue to "Planning File Server Security" later in this chapter.

Figure 2.9 describes the planning process for file server availability.

**Figure 2.9 Planning File Server Availability**



# Planning for File Server Uptime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following guidelines provide basic steps to increase the uptime of file servers. For comprehensive coverage of these topics, see "Planning for High Availability and Scalability" in this book.

**Choosing Hardware for Reliability and Availability**
Implement a well-planned hardware strategy to increase file server availability while reducing support costs and failure recovery times. You can choose hardware for reliability and availability by following these guidelines:

- Choose hardware from the Windows Server Catalog for products in the Windows Server 2003 family.
- Establish hardware standards.
- Keep spares or stand-by systems for redundancy.
- Use fault-tolerant network components, server cooling, and power supplies.
- Use error checking and correcting (ECC) memory. ECC memory uses a checking scheme to guarantee that the failure of any one bit out of a byte of information is corrected.
- Use fault-tolerant storage components, such as redundant disk controllers, hot-swappable disks and hot spares, and disks configured as redundant arrays of independent disks (RAID). For more information about RAID, see "Planning the Layout and RAID Level of Volumes" later in this chapter.
- Use disk resource management tools, such as disk quotas, to ensure that users always have available disk space.
- Keep a log or database of changes made to the file server. The log should contain dated entries for hardware failures and replacements, service pack and software installations, and other significant changes.
- Deploy FRS-compliant antivirus software on link target computers prior to adding files to FRS replicated links or adding new members to the replica set.

**Maintaining the Physical Environment**
Maintain high standards for the environment in which the file servers must run. Neglecting the environment can negate all other efforts to maintain the availability of file servers. Take the following measures to maintain the physical environment:

- Maintain proper temperature and humidity.
- Protect servers from dust or contaminates.
- For power outages, provide a steady supply of power for the servers by using uninterruptible power supply (UPS) units or standby generators.
- Maintain server cables.
- Secure the server room.

**Planning for Backup and Preparing for Recovery**
Backups are essential for high-availability file servers, because the ultimate recovery method is to restore from backup. To plan for backup and recovery:

- Create a plan for backup.

- Monitor backups.
- Decide between local and network backups.
- Check the condition of backup media.
- Perform trial restorations on a regular basis. A trial restoration confirms that your files are properly backed up and can uncover hardware problems that do not show up with software verifications. Be sure to note the time it takes to either restore or re-replicate the data so that you will know how long it takes to bring a server back online after a failure.
- Perform regular backups of clustered file servers and test cluster failures and failover policies. For more information about backing up server clusters and testing policies, see "Backing up and restoring server clusters" and "Test Cluster Failures and Failover Policies" in Help and Support Center for Windows Server 2003.
- If you are using DFS and FRS, put a procedure in place for recovering failed members of an FRS replica set. For more information about troubleshooting FRS, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**Choosing Software for Reliability**

Installing incompatible or unreliable software reduces the overall availability of file servers. To choose software for reliability, follow these guidelines:

- Select software that is compatible with Windows Server 2003.
- Select signed drivers. Microsoft promotes driver signing for designated device classes as a mechanism to advance the quality of drivers, to provide a better user experience, and to reduce support costs for vendors and total cost of ownership for customers. For more information about driver signing, see the Driver Signing and File Protection link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- Select software that supports the high-availability features you require, such as server clusters and online backup.

# Identifying Business-Critical File Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you plan for file server uptime, identify file servers that must be highly available. These file servers typically contain business-critical data that is required for an organization's central purpose, such as software distribution points, e-mail or business databases, and internal or external Web sites. File servers need to be highly available for the following reasons as well:

- **The file server contains one or more stand-alone DFS roots.** If you create a stand-alone DFS root on a file server, and the server fails or is taken offline for maintenance, the entire DFS namespace is unavailable. Users can access the shared folders only if they know the name of the file servers where the shared folders are located. To make stand-alone DFS namespaces fault tolerant, you can create the roots on clustered file servers.

- **Your organization is consolidating file servers.** Consolidation leads to a greater dependency on fewer file servers, so you must ensure the availability of the remaining file servers.

- **Your organization has existing SLAs.** Service Level Agreements (SLAs) and Organization Level Agreements (OLAs) specify the required uptime for file servers, usually defined as the percentage of time that the file server is available for use. For example, your organization might require that file servers have 99.7-percent uptime regardless of the type of data they contain. To meet these agreements, some organizations deploy clustered file servers and assign experienced administrators to manage those file servers. In addition to deploying hardware, these administrators use defined and tested processes to fulfill these agreements.

You do not need to implement additional availability strategies on file servers that store temporary or non-business-critical files. For example, if the file server fails in some way, but users can tolerate the loss of the file server for the time it takes to repair the file server or restore the data from backup, you do not need to use an availability strategy such as replication or clustering. Instead, you can work to decrease the amount of time it takes to restore the file server.

# Choosing an Availability Strategy for Business-Critical Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If a file server contains business-critical data, you need to make certain that the data is highly available. Windows Server 2003 provides two primary strategies for increasing data availability: FRS and clustering.

**FRS** This strategy involves creating one or more domain-based DFS namespaces, using link targets that point to multiple file servers and using File Replication service (FRS) to synchronize the data in the link targets. This chapter describes the design and deployment process for FRS, although you can also synchronize data manually by using tools such as Robocopy or by using third-party replication tools.

**Clustering** A server cluster is a group of individual computer systems working together cooperatively to provide increased computing power and to provide continuous availability of business-critical applications or resources. This group of computers appears to network clients as if it were a single system, by virtue of a common cluster name. A cluster can be configured so that the workload is distributed among the group, and if one of the cluster members fails, another cluster member automatically assumes its duties.

Both of these strategies involve using multiple file servers to ensure data availability. If for some reason you cannot use multiple file servers, follow the guidelines in "Planning for File Server Uptime" earlier in this chapter to increase the availability of the physical server.

When evaluating these two strategies, you must keep in mind your organization's tolerance for inconsistent data. FRS can cause temporary data inconsistency as data is replicated across multiple servers. Clustered file servers maintain only one copy of the data; therefore, data inconsistency does not occur.

**Note**

- If your organization plans to implement geographically dispersed clusters for disaster tolerance, you need to understand your data consistency needs in different failure and recovery scenarios and work with the solution vendors to match your requirements. Different geographically dispersed cluster solutions provide different replication and redundancy strategies, ranging from synchronous mirroring across sites to asynchronous replication. For more information about geographically dispersed clusters, see "Designing and Deploying Server Clusters" in this book.

**Using FRS as an Availability Strategy**

You can use FRS to replicate data in domain-based DFS namespaces on file servers running a Windows 2000 Server or Windows Server 2003 operating system. When evaluating FRS, you must determine whether your organization can tolerate periods of inconsistent data that can occur within a replica set. Data inconsistency can occur at the file and folder level as follows:

- 07FRS uses a "last writer wins" algorithm for files. This algorithm is applied in two situations: when the same file is changed on two or more servers, and when two or more different files with the same name are added to the replica tree on different servers. The most recent update to a file in a replica set becomes the version of the file that replicates to the other members of the replica set, which might result in data loss if multiple masters have updated the file. In addition, FRS cannot enforce file-sharing restrictions or file locking between two users who are working on the same file on two different replica set members.

- FRS uses a "last writer wins" algorithm when a folder on two or more servers is changed, such as by changing folder attributes. However, FRS uses a "first writer wins" algorithm when two or more identically named folders on different servers are added to the replica tree. When this occurs, FRS identifies the conflict during replication, and the receiving member protects the original copy of the folder and renames (*morphs*) the later inbound copy of the folder. The morphed folder names have a suffix of "_NTFRS_xxxxxxxx," where "xxxxxxxx" represents eight random hexadecimal digits. The folders are replicated to all servers in the replica set, and administrators can later merge the contents of the folders or take some other measure to reestablish the single folder.

Temporary data inconsistency due to replication latency is more likely to occur in geographically diverse sites with infrequent replication across slow WAN links. If you want to use replication among servers in the same site, consistency is probably not an issue, because the replication can occur quickly after the file changes — assuming that only one user makes changes to the data. If two users make changes to the data, replication conflicts occur and one user loses those changes.

Replication works well in the following scenarios.

**When the data is read-only or changes infrequently**
Because changes occur infrequently, the data is usually consistent. In addition, FRS has less data to replicate, so network bandwidth is not heavily affected.

**When the sites are geographically dispersed and consistency is not an issue**
Geographically dispersed sites might have slower bandwidth connections, but if your organization does not require the data in those sites to always be consistent with each other, you can configure replication in those sites on a schedule that make sense for your organization. For example, if your organization has sites in Los Angeles and Zimbabwe, you can place one or more replicas of the data in servers in those sites and schedule replication to occur at night or during periods of low bandwidth use. Because in this scenario replication could take hours or days to update every member, the delay must be acceptable to your organization.

**When each file is changed by only one person from one location**
Replication conflicts rarely occur if only a single user changes a given file from a single location. Some common scenarios for single authorship are redirected My Documents folders and other home directories. Conversely, if users roam between sites, replication latency could cause the file to be temporarily inconsistent between sites.

**When replication takes place among a small number of servers in the same site**
Replication latency is reduced by frequently replicating data using high-speed connections. As a result, data tends to be more consistent.

Replication should not be used in the following scenarios.

**In organizations with no operations group or dedicated administrators**
Organizations that do not have the staff or the time to monitor FRS event logs on each replica member should not implement FRS. Organizations must also have well-defined procedures in place to prevent the accidental or unintentional deletion of data in the replica set, because deleting a file or folder from one replica member causes the file or folder (and its contents) to be deleted from all replica members. In addition, if a folder is moved out of the replica tree, FRS deletes the folder and its contents on the remaining replica members. To avoid having to restore the files or folders from backup, you can enable shadow copies on some of the replica members so that you can easily restore a file or folder that was accidentally deleted. For more information about shadow copies, see "Designing a Shadow Copy Strategy" later in this chapter. For more information about FRS logs, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**In organizations that do not update virus signatures or closely manage folder permissions**
A virus in FRS-replicated content can spread rapidly to replica members and to clients that access the replicated data. Viruses are especially damaging in environments where the Everyone group has share permissions or NTFS permissions to modify content. To prevent the spread of viruses, it is essential that replica members have FRS-compatible, up-to-date virus scanners installed on the servers and on clients that access replicated data. For more information about preventing the spread of viruses, see "Planning Virus Protection for File Servers" and "Planning DFS and FRS Security" later in this chapter.

**When the rate of change exceeds what FRS can replicate**
If you plan to schedule replication to occur during a specified replication window, verify that FRS can replicate all the changed files within the window. Replication throughput is determined by a number of factors:

- The number and size of changed files
- The speed of the disk subsystem
- The speed of the network
- Whether you have optimized the servers by placing the replica tree, the staging directory, and the FRS data on separate disks.

Each organization will have different FRS throughput rates, depending on these factors. In addition, if your data compresses extremely well, your file throughput will be higher. To determine the replication rate, perform testing in a lab environment that resembles your production environment.

If the amount of data changes exceeds what FRS can replicate in a given period of time, you need to change one of these factors, such as increasing the speed of the disk subsystem (number of disks, mechanical speed, or disk cache) or network. If no change is possible, FRS is not recommended for your organization.

**In organizations that always use clustered file servers**
Some organizations use clustered file servers regardless of whether the server contains business-critical data. Although storing FRS-replicated content on the cluster storage of a clustered file server might imply increased availability of the data, combining clustering and FRS is not recommended. Data might become inconsistent among the members of the replica set, thus defeating the purpose of clustering, which is to have highly available data that remains consistent because only one copy of the data exists. In addition, Windows Server 2003 does not support configuring FRS to replicate data on cluster storage.

**In organizations that use Remote Storage**
Remote Storage is a feature in Windows Server 2003 that automatically copies infrequently used files on local volumes to a library of magnetic tapes or magneto-optical disks. Organizations that use Remote Storage must not use FRS on the same volume. Specifically, do not perform any of the following tasks:

- Do not create a replica set on a volume that is managed by Remote Storage.
- Do not add a volume that contains folders that are part of an FRS replica set to Remote Storage.
- If you use Remote Storage for volumes that contain FRS replica sets, backup tapes might be damaged or destroyed if FRS recalls a large number of files from Remote Storage. The damage occurs because FRS does not recall files in media order. As a result, files are extracted randomly, and the process can take days to complete and might damage or destroy the tape in the process. Random extraction from magneto-optical disks can also be extremely time consuming.

### Caution

- Windows Server 2003 does not prevent you from using Remote Storage and FRS replica sets on the same volumes, so take extra precautions to avoid using these two features on the same volume.

**When locks by users or processes prevent updates to files and folders**
FRS does not replicate locked files or folders to other replica members, nor does FRS update a file on a replica member if the local file is open. If users or processes frequently leave files open for extended periods, consider using clustering instead of FRS.

**When the data to be replicated is on mounted drives**
If a mounted drive exists in a replica tree, FRS does not replicate the data in the mounted drive.

**When the data to be replicated is encrypted by using EFS**
FRS does not replicate files encrypted by using EFS, nor does FRS warn you that EFS-encrypted files are present in the replica set.

**When the FRS jet database, FRS logs, and staging directory are stored on volumes where NTFS disk quotas are enabled**
If you plan to store a replica set on a volume where disk quotas are enabled, you must move the staging directory, FRS jet database, and FRS logs to a volume where disk quotas are disabled. For more information, see "Planning the Staging Directory" later in this chapter.

**Using Clustering as an Availability Strategy**
If the data changes frequently and your organization requires consistent data that is highly available, use clustered file servers. Clustered file servers allow client access to file services during unplanned and planned outages. When one of the servers in the cluster is unavailable, cluster resources and applications move to other available cluster nodes. Server clusters do not guarantee nonstop operation, but they do provide sufficient availability for most business-critical applications, including file services. A cluster service can monitor applications and resources and automatically recognize and recover from many failure conditions. This ability provides flexibility in managing the workload within a cluster and improves overall system availability.

Server cluster benefits include the following:

**High availability** Ownership of resources, such as disks and IP addresses, is automatically transferred from a failed server to a surviving server. When a system or application in the cluster fails, the cluster software restarts the failed application on a surviving server, or it disperses the work from the failed node to the remaining nodes. As a result, users experience only a momentary pause in service.

**Manageability** You can use the Cluster Administrator snap-in to manage a cluster as a single system and to manage applications as if they were running on a single server. You can move applications to different servers within the cluster, and you can manually balance server workloads and free servers for planned maintenance. You can also monitor the status of the cluster, all nodes, and resources from anywhere on the network.

**Scalability** Server clusters can grow to meet increased demand. When the overall client load for a clustered file server exceeds the cluster's capabilities, you can add additional nodes.

Clustered file servers work well in the following scenarios.

**When multiple users access and change the files**
Because only one copy of the file exists, Windows Server 2003 can enforce file locking so that only one user can make changes at a time. As a result, data is always consistent.

**When large numbers of users access data in the same site**
Clustered file servers are useful for providing access to users in a single physical site. In this case, you do not need a replication method to provide data consistency among sites.

**When files change frequently and data consistency is a must**
Even with a large number of changes, data is always consistent and there is no need to replicate the changes to multiple servers.

**When you want to reduce the administrative overhead associated with creating many shared folders**
On clustered file servers, you can use the Share Subdirectories feature to automatically share any folders that are created within a folder that is configured as a File Share resource. This feature is useful if you need to create a large number of shared folders.

**When you want to ensure the availability of a stand-alone DFS root**
Creating stand-alone DFS roots on clustered file servers allows the namespaces to remain available, even if one of the nodes of the cluster fails.

**When you want to make encrypted files highly available**
Windows Server 2003 supports using the EFS clustered file servers. Using EFS in FRS replica sets is not supported. For more information about using EFS on clustered file servers, see "Planning Encrypted File Storage" later in this chapter.

Some issues to consider when using clustered file servers include the following.

**Dynamic disks are not available**
If you want to use dynamic disks, you must use them on nonclustered file servers or on the local storage devices of the cluster. If the node hosting the local storage devices fails, the data becomes unavailable until the node is repaired and brought back online.

If you need to extend the size of basic volumes used for shared cluster storage, you can do so by using DiskPart. For more information about extending basic volumes, see "Extend a basic volume" in Help and Support Center for Windows Server 2003.

**Clustered file servers must use complete cluster systems**
For your server clusters to be supported by Microsoft, you must choose complete cluster systems from the Windows Server Catalog for the Windows Server 2003 family. For more information about support for server clusters, see article Q309395, "The Microsoft Support Policy for Server Clusters and the Hardware." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Increasing Data Availability by Using FRS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To design an FRS replication strategy, you need to take the following steps:

1. Identify data to replicate.
2. Choose the replication topology.
3. Design the replication schedule.
4. Plan the staging directory.
5. Determine connection priorities.

The following sections describe each of these steps.

**Important**

- If you configure FRS to replicate files on file servers running Windows 2000, it is highly recommended that you install the Windows 2000 Service Pack 3 (SP3) and the post-SP3 release of Ntfrs.exe, or install later service packs that include this release. For more information about updating FRS, see article 811370, "Issues That Are Fixed in the Post-Service Pack 3 Release of Ntfrs.exe." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For more information about FRS security, see "Planning DFS and FRS Security" later in this chapter.

# Identifying Data to Replicate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you identify which DFS roots and links require multiple replicated targets, as described in "Increasing the Availability of DFS Namespaces" earlier in this chapter, you also need to consider the following issues.

**Size of the USN journal**
The default USN journal size in Windows Server 2003 is 512 MB. If your volume contains 400,000 files or fewer, no additional configuration is required. For every 100,000 additional files on a volume containing FRS-replicated content, increase the update sequence number (USN) journal by 128 MB. If files on the volume are changed or renamed frequently (regardless of whether they are part of the replica set), consider sizing the USN journal larger than these recommendations to prevent USN journal wraps, which can occur when large numbers of files change so quickly that the USN journal must discard the oldest changes (before FRS has a chance to detect the changes) to stay within the specified size limit. To recover from a journal wrap, you must perform a nonauthoritative restore on the server to synchronize its files with the files on the other replica members. For more information about USN journal wraps, see article Q292438, "Troubleshooting Journal_Wrap Errors on Sysvol and DFS Replica Sets." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Where to store the replica tree**
To provide the best consistency and minimize administrative intervention, carefully plan where to store the replica tree. Your choice will depend on the amount of storage on the server, the volume layout, the number of files on the volume that do not participate in a replica set, and the number of replica sets on the server.

- If you have a single replica tree, store it on its own volume if possible, or store it on a volume that also stores other user data. Adjust the USN journal size based on the total number of files on the volume.
- If you have more replica trees than volumes, store multiple replica trees on each volume and adjust the USN journal size based on the total number of files on the volume.
- If you can create a volume for each replica tree, store each replica tree on its own volume and adjust the USN journal accordingly.

The staging directory also requires careful placement. For more information, see "Planning the Staging Directory" later in this chapter.

**Replicated link targets to omit from the DFS namespace**
When you omit a replicated link target from the namespace, you disable referrals so that DFS does not direct clients to the link target. However, the link target still replicates with other members. Disabling DFS referrals for replica members is useful for a number of scenarios. For example, you can use such a member for any of the following purposes:

- As a backup server. For example, assume you have five servers in a replica set. You might make four of the servers part of the DFS namespace and use the fifth server as a backup source. In this scenario, DFS never refers clients to the fifth server.
- As a reference server. For example, you omit (from the namespace) the server where all changes originate, thus ensuring that you can always introduce new content without being blocked because of open files. The reference server becomes the "known good" server; all other servers should contain identical content.
- As a way to update application data. For example, you have two replica members, ServerA and ServerB, which contain application data that is frequently updated. You can disable referrals for ServerA, update the application data, and enable referrals to ServerA. You can then repeat the procedure on ServerB. In this scenario, one of the replica members is always available while the other is being updated.

To omit a replicated link target from the namespace, configure the link target within the DFS namespace and configure replication as desired by using the Distributed File System snap-in. Then use the **Enable or Disable Referral** command in the Distributed File System snap-in to disable referrals to the link targets that you want to omit from the namespace.

**Files or folders to exclude from replication**
You can use the Distributed File System snap-in to set filters that exclude subfolders (and their contents) or files from replication. You exclude subfolders by specifying their name, and you exclude files by using wild cards to specify file names and extensions. By default, no subfolders are excluded. The default file filters exclude the following files from replication:

- File names starting with a tilde (~) character
- Files with .bak or .tmp extensions

Filters act as exclusion filters only for new files and folders added to a replica set. They have no effect on existing files in the replica set. For example, if you change the existing file filter from "*.tmp, *.bak" to "*.old, *.bak," FRS does not go through the replica set and exclude all files that match *.old, nor does it go through the replica set and begin to replicate all files that match *.tmp. After the filter change, new files matching *.old that are added to the replica set are not replicated. New files matching *.tmp that are added to the replica set are replicated.

Any file in the replica set that was excluded from replication under the old file filter (such as Test.tmp, created when the old filter was in force) is automatically replicated under the new file filter only after the file is modified. Likewise, changes to any file that was not excluded from replication under the old filter (such as Test.old, created when the old filter was in force) continue to replicate under the new filter, until you explicitly delete the file.

These rules apply in the same manner to the directory exclusion filter. If a directory is excluded, all subdirectories and files under that directory are also excluded.

Regardless of the filters you set, FRS always excludes the following from replication:

- NTFS mounted drives
- Files encrypted by using EFS
- Any reparse points except those associated with DFS namespaces. If a file has a reparse point used for Hierarchical Storage Management (HSM) or Single Instance Store (SIS), FRS replicates the underlying file but not the reparse point.
- Files on which the temporary attribute has been set

For more information about configuring filters, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

**Number of servers in the replica set**
Will you replicate data among three servers or three hundred servers? As the number of servers increases, configuring the topology and schedule becomes more complex, and it takes longer to replicate data to all members of the replica set. If many members exist in the replica set, and if any member can originate new data, the amount of aggregate bandwidth consumed can be very high. Also note that any server with outbound connections is a server that can potentially originate content and thus requires careful monitoring.

**Number of replica sets per server**
Although there is no limit to the number of replica sets that can be hosted on a single server, it is recommended that you host no more than 150 replica sets on a single server to provide optimal replication performance. The optimal number of replica sets for servers in your organization depends on the CPU, memory, disk input/output (I/O) throughput, and the amount of data changed.

**Network bandwidth between sites**
Replicating data between sites that are connected with slow WAN links requires careful planning of the topology and schedule, because FRS does not provide bandwidth throttling. If the sites have a high-bandwidth connection, but business-critical databases and other applications use that connection as well, schedule replication so that it does not consume bandwidth that is needed for other uses.

**Amount of data in the replica tree**
If the replica tree includes a large amount of data, and a new member is served by a low-bandwidth link, prestage the replica tree on the new member. Prestaging a replica tree preserves security descriptors and object IDs and involves restoring the data to the new member without using the low-bandwidth link. For example, you can restore a backup of the replica tree on the new member, or you can place the new member in the same site as an existing member and prestage the replica tree by using a high-bandwidth local area network (LAN) connection. For more information about adding new replica members, see "Adding

a New Member to an Existing Replica Set" later in this chapter.

**How frequently the data changes**
If the data that is to be replicated changes frequently, estimate the total amount of data that will be generated by each replica member over a period of time. Data that changes frequently is more difficult to keep synchronized across multiple servers, especially if those servers are located across WAN links. If replication latency is a concern, consider clustering instead of replication.

**Whether you want to use single master replication**
FRS is a multimaster replication engine, which means that new content and changes to existing content can originate from any member. If you want to limit the servers that can originate data, you can set up single master replication in one of the following ways:

- Set NTFS permissions on targets to control who can update data in the replica tree.
- Configure the replication topology so that a target does not have any outbound connections.

### Important

- If you use this method, do not directly make changes on replica members that have no outbound connections. Changes made directly on a server that has no outbound connections will not replicate, and the server replica tree will be inconsistent and potentially unpredictable.

**Speed of recovery from hardware failures**
If a server contains a large amount of replicated, business-critical data, be sure to use redundant server components, such as hardware RAID, redundant disk controllers, and redundant network cards to reduce the likelihood that a hardware failure could cause the server to become unavailable. To further increase availability, use multiple servers in each site so that if a server in the site is unavailable due to hardware problems, clients in that site can access the remaining servers as you repair and restore the data on the failed server. Using multiple servers is also useful for remote sites where repair time is lengthy and in cases where the amount of replicated data is so large that it cannot be restored from a central site within the required restoration window.

**Expected growth of replicated data**
You need to know whether you plan to replicate larger and larger amounts of data over time so that you can make sure that your topology, schedule, and bandwidth can handle the additional data.

**Expected increase in the number of replica members**
If you plan to deploy a small number of servers at first and then deploy additional servers over time, you need to make sure that your topology and bandwidth can handle the new servers. For example, if you plan to add 100 more servers, and each of those servers can originate data, you must make sure that your network has the bandwidth to replicate this data. On the other hand, if you are deploying 100 new servers, but only one of those servers can originate data, bandwidth might not be a limiting factor, although you do need to make sure that your bandwidth can handle the initial synchronization of data among servers.

# Choosing a Replication Topology

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The replication topology describes the logical connections that FRS uses to replicate files among servers. To minimize the network bandwidth required for replication, identify where your bandwidth is highest and lowest on your network, and model your FRS replication topology after the physical topology of your network.

You use the Distributed File System snap-in to select a replication topology. After you add a second link target to a link, you are prompted to configure replication by using the Configure Replication Wizard, which allows you to perform the following tasks:

- Choose the initial master whose contents are replicated to the other targets. The initial master is only relevant during the creation of the replica set. After that, the server that acted as the initial master is treated no differently from any other server.
- Choose the location of the staging directory. By default, the staging directory is placed on a different volume than the content to be replicated.
- Choose the replication topology: ring, hub and spoke, full mesh, or custom. If you choose custom, you can add or delete connections to or from the replica set. For the other three standard topologies, you can only enable or disable connections between target servers.

The following describes the four topology types available in the Distributed File System snap-in. For an Excel spreadsheet to assist you in documenting your decision after you choose the topology, see "FRS Configuration Worksheet" (Sdcfsv_2.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "FRS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**Ring topology**

Figure 2.10 illustrates a ring topology. In a ring topology, files replicate from one server to another in a circular configuration, with each server connected to the servers on either side of it in the ring. Choose a ring topology if your physical network topology resembles a ring topology. For example, if each server is located in a different site and has existing connectivity with neighboring servers, you can choose the ring topology so that each server connects only to neighboring servers. Because the ring topology is bidirectional, each connection is fault tolerant. If a single connection or server fails, data can still replicate to all members in the opposite direction.

**Figure 2.10 Ring Topology**




If you plan to create a ring topology with more than seven members, consider adding shortcut connections between some of the members to reduce the number of hops required for data to replicate to all members.

**Hub and spoke topology**

Figure 2.11 illustrates hub and spoke topology. In a hub and spoke topology, you designate one server as the *hub*. Other servers, called *spokes*, connect to the hub. This topology is typically used for WANs that consist of faster network connections between major computing hubs and slower links connecting branch offices. In this topology, files replicate from the hub server to the spoke servers and vice versa, but files do not replicate directly between two spoke servers. When you choose this topology, you must choose which server will act as the hub. If you want to set up multiple hubs, use a custom topology. Using multiple hubs is recommended, because using only one hub means that the hub is a single point of failure.

**Figure 2.11 Hub and Spoke Topology**




**Full mesh topology**

Figure 2.12 illustrates a full mesh topology. In a full mesh topology, every server connects to every other server. A file created on one server replicates directly to all other servers. Because each member connects to every other member, the propagation of change orders for replicating files can impose a heavy burden on the network. To reduce unnecessary traffic, use a different topology or delete connections you do not actually need. A full mesh topology is not recommended for replica sets with five or more members.

**Figure 2.12 Full Mesh Topology**




**Custom topology**

With a custom topology, you create the connections between the servers yourself. One example of a custom topology is a redundant hub and spoke topology. In this configuration, a hub site might contain two file servers that are connected by a high-speed link. Each of these two hub servers might connect with four branch file servers in a hub and spoke arrangement. An example of this topology is shown in "Example: An Organization Designs a Replication Strategy" later in this chapter.

Figure 2.13 illustrates another type of hub and spoke topology known as a *multitier redundant hub and spoke* with an optional staging server used for introducing changes into the replica set.

**Figure 2.13 Multitier Redundant Hub and Spoke Topology**



# Designing a Replication Schedule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you use the Distributed File System snap-in to configure replication, FRS schedules replication to take place 24 hours a day, seven days a week. Whenever you copy a file to a target that participates in replication, or when an existing file in the replica set is changed, FRS replicates the entire file to the other targets using the connections specified in the replication topology. Continuous replication is advised only for environments that meet the following criteria:

- The replica members are all connected by high-bandwidth connections, such as hub servers in a redundant hub and spoke topology.
- The amount of data to be replicated is not so large that it interferes with the bandwidth required for other purposes during peak or off-peak hours.
- The application whose data is being replicated needs rapid distribution of file changes, such as antivirus signature files or login scripts.

If your environment does not meet these criteria, plan to create a replication schedule so that data is replicated at night or during other times as appropriate. A replication schedule involves enabling and disabling replication so that it occurs at specified times on specified days. For example, you might schedule replication to occur beginning at 12:00 midnight and ending at 2:00 A.M. every day.

When determining the replication schedule, consider the following issues:

**Amount of data to be replicated and the length of the replication window** Estimate the amount of data to be replicated so that you can choose a duration that is long enough to replicate all the data. In DFS namespaces, replication stops when the schedule window closes, and any remaining files are delayed until the next replication window opens. If you occasionally put a large amount of data into the replica set, FRS might take several replication periods to replicate all the data.

**Amount of latency your organization can tolerate** If you plan to replicate data that changes frequently, consider the amount of latency that you can tolerate in keeping your targets synchronized and the amount of bandwidth that is consumed during replication.

**Whether you want to stagger schedules** If you are using a redundant hub and spoke topology with several hubs, you can split the replication load among the hubs by staggering the schedule. An example of this schedule is described in "Example: An Organization Designs a Replication Strategy" later in this chapter.

For an Excel spreadsheet to assist you in documenting the replication schedule, see "FRS Configuration Worksheet" (Sdcfsv_2.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "FRS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

# Planning the Staging Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The staging directory acts as a buffer by retaining copies of updated files until replication partners successfully receive them and move them into the target directory. The following sections describe the benefits and design considerations for staging directories. For an Excel spreadsheet to assist you in documenting the staging directory settings and location, see "FRS Configuration Worksheet" (Sdcfsv_2.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "FRS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**Why FRS Uses a Staging Directory**

The staging directory acts as a queue for changes to be replicated to downstream partners. After the changes are made to a file and the file is closed, the file content is compressed, written to the staging directory, and replicated according to schedule. Any further use of that file does not prevent FRS from replicating the staging file to other members. In addition, if the file is replicated to multiple downstream partners or to members with slow data links, using a staging file ensures that the underlying file in the replica tree can still be accessed.

### Note

- A new or updated file remains in the staging directory of a replica member with downstream partners even after the file has replicated to all members. The file is retained for seven days in order to optimize future synchronizations with new downstream partners. After seven days, FRS deletes the file. Do not attempt to delete any files in the staging directory yourself.

**Staging Directory Size**

The size of the staging directory governs the maximum amount of disk space that FRS can use to hold those staging files and the maximum file size that FRS can replicate. The default size of the staging directory is approximately 660 MB, the minimum size is 10 MB, and the maximum size is 2 TB. The largest file that FRS can replicate is determined by the staging directory size on both the upstream partner (the replica member that sends out the changed file) and downstream partners (the replica member that receives the changed file) and whether the replicated file, to the extent that it can be compressed, can be accommodated by the current staging directory size. Therefore, the largest file that FRS can replicate is 2 TB, assuming that the staging directory size is set to the maximum on upstream and downstream partners.

**Managing Staging Directory Capacity**

When FRS tries to allocate space for a staging file and is not successful, either because there is not enough physical disk space or because the size of the files in the staging directory has reached 90 percent of the staging directory size, FRS starts to remove files from the staging directory. Staged files are removed (in the order of the longest time since the last access) until the size of the staging directory has dropped below 60 percent of the staging directory limit. Additionally, staging files for downstream partners that have been inaccessible for more than seven days are deleted. As a result, FRS does not stop replicating if the staging directory runs out of free space. This means that if a downstream partner goes offline for an extended period of time, the offline member does not cause the upstream partner's staging directory to fill with accumulated staging files.

### Note

- In Windows Server 2003, FRS detects and suppresses excessive replication that could be caused when you use Group Policy to overwrite file permissions or when you use antivirus software that overwrites security descriptors. In these examples, writes to files cause no net content changes. When FRS detects a significant increase in the number of changes made to a file, FRS logs an event with ID 13567 in the File Replication service event log. Despite the FRS response to excessive replication caused by applications, services, or security policies, you should investigate the source of the excessive replication and eliminate it.

The fact that FRS removes files from the staging directory does not mean that the underlying file is deleted or will not be replicated. The change order in the outbound log still exists, and the file will eventually be sent to downstream partners when they process the change order. However, before replication can take place, the upstream partner must recreate the file in the staging directory, which can affect performance. Recreating the staging file can also cause a replication delay if the file on the upstream partner is in use, preventing FRS from creating the staging file.

A staging directory can reach 90 percent of its capacity in the following situations:

- **More data is being replicated than the staging directory can hold.** If you plan to replicate more than 660 MB of data, or if you expect that the largest file to replicate will be 660 MB or larger, you must increase the size of the staging directory. Use Table 2.7 for sizing guidelines.

- **A slow outbound partner.** Construct replica connections that have comparable bandwidth for all outbound partners. It is also a good idea to balance that bandwidth for inbound and outbound connections. If you cannot balance the bandwidth, increase the size of the staging directory.

- **Replication is disabled.** If a file is changed or added to the replica set while replication is disabled, the staging directory must be large enough to hold data awaiting replication during the off-time or else older staged files are removed. Replication still occurs, but less optimally.

Table 2.7 describes a range of sizing guidelines for staging directories. Use the numbers in Table 1.7 as a baseline for testing and then adjust the numbers as necessary for your environment. Also, note that the numbers in Table 1.7 apply to each replica set on a server. If a server contains multiple replica sets, you must follow these guidelines for each staging directory.

**Table 2.7 Staging Directory Size Guidelines per Replica Set**

| Scenario | Minimum | Acceptable | Best Performance |
|---|---|---|---|
| Adding a new replica member | Whichever is larger:<br>• 660 MB<br>• Take the size of the 128 largest files in the replica tree, multiplied by the number of downstream partners, and then multiply that number by 1.2. | Whichever is larger:<br>• 660 MB<br>• Take the size of the 128 largest files in the replica tree, multiplied by the number of downstream partners, and then multiply that number by 1.2. | Whichever is larger:<br>• 660 MB<br>• Take the size of the 128 largest files in the replica tree, multiplied by the number of downstream partners, and then multiply that number by 1.2. |
| Increasing the staging directory space to account for backlog caused by replication schedules | No additional requirement | Add space equal to the maximum quantity of expected file changes in a seven-day period multiplied by 1.2. | Add space equal to the expected size of the replica tree, multiplied by 1.2. |
| Additional recommendations | No additional recommendations | Use dedicated disks for the staging directory. | Use dedicated, high-performance disks for the staging directory. |
| Suitability for large replica sets | Not recommended | Recommended | Recommended for organizations that require highest performance. |

**Note**

- The recommendations in Table 2.7 offer sufficient staging directory space for a worst-case scenario. In fact, the staging directory needs to be only as large as the largest 128 files, multiplied by the number of downstream partners that are concurrently performing a full synchronization against the replica member.

For more information about configuring the staging directory, see article Q329491, "Configuring Correct Staging Area Space for Replica Sets." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For information about changing the size of the staging directory, see "Deploying FRS" later in this chapter.

**Staging Directory Compression Considerations**

FRS replica members running Windows 2000 Server with Service Pack 3 (SP3) or Windows Server 2003 compress the files replicated among them. Compression reduces the size of files in the staging directory on the upstream partners, over the network between compression-enabled partners, and in the staging directory of downstream partners prior to files being moved into their final location.

To maintain backward compatibility, servers running Windows 2000 Server with SP3 or Windows Server 2003 generate uncompressed staging files when they send changed files to downstream partners running Windows 2000 Server with Service Pack 2 (SP2) or earlier (or to servers running Windows 2000 Server with SP3 or Windows Server 2003 that have compression explicitly disabled in the registry).

To accommodate environments that contain a combination of compression-enabled and compression-disabled partners, the server originating or forwarding changes must generate two sets of staging files for each modified file or folder: one compressed version for servers running Windows 2000 Server with SP3 or Windows Server 2003, and one uncompressed version for partners that have compression explicitly disabled or that are running Windows 2000 Server with SP2 or earlier. To mitigate the generation of two sets of staging files, you can take one of the following steps:

- Confirm that you have sufficient disk space and an appropriately sized staging directory on all replica members.
- Explicitly disable compression in the registry until all members of the replica set are running Windows 2000 Server with SP3 or Windows Server 2003. For information about disabling replication, see article Q288160, "Error Message: Error Compressing Data," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- Rapidly deploy Windows 2000 Server with SP3 or Windows Server 2003 on all members of the replica set, especially if you are replicating a large amount of content.

**Choosing the Staging Directory Location**

You can choose the location of the staging directory when you use the Distributed File System snap-in to configure a replica set. By default, all replica sets on the server use the same staging directory, and the staging directory space limit applies to all staging directories on the local server. When the staging directory is shared among multiple replica sets, any one replica set can cause the staging directory to reach 90% capacity, at which point FRS begins to delete the oldest staging files for any of the replica sets.

To simplify staging directory management and any future recovery procedures, it is recommended that you configure replica sets in one of the following ways:

- Give each replica set its own staging directory, and place the staging directories on the same volume. This solution provides good recoverability while minimizing cost.
- Give each replica set its own staging directory, and put each staging directory on its own volume. This solution provides better recoverability, but at a higher cost.

Do not store the staging directory on a volume where NTFS disk quotas are enabled.

**Improving FRS Performance**

To distribute disk traffic, store the staging directory, FRS logs, the FRS jet database (Ntfrs.jdb), and replica trees on separate disks. Using separate disks for the log files is especially important when you configure FRS logging by using a high severity level. For the best replication performance, distribute disk I/O by placing the FRS logs and the staging directory on separate disks from that of the operating system and the disks containing the replica tree. In a hub and spoke topology, using separate disks is especially recommended for the hub servers.

### Important

- If you store the FRS jet database on a disk that has its write cache enabled, FRS might not recover if power to the drive is interrupted and critical updates are lost. Windows Server 2003 warns you about this by adding an entry with Event 13512 in the File Replication service event log. To leave write caching enabled for improved performance and to ensure the integrity of the FRS jet database, use an uninterrupted power supply (UPS) or a disk controller and cache with a battery backup. Otherwise, disable write caching until you can install a backup power supply.

When moving the FRS logs and FRS jet database, be sure that NTFS disk quotas are disabled on the target volume. For more information about moving the FRS logs and FRS jet database, see article Q221093, "HOW TO: Relocate the NTFRS Jet Database and Log Files," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Determine Connection Priorities

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Connection priorities control the sequencing of the initial synchronization that occurs when you add a new member to a replica set or when you perform a nonauthoritative restore. By changing connection priorities, you can control which inbound partners go first based on resource considerations. (A server's inbound partners, also known as upstream partners, are those servers from which the new or recovered member can receive replicated content.) For example, you can specify that a new member first synchronize with a partner:

- Across a high-bandwidth network connection.
- That has low server activity.
- That has the most up-to-date files.

You configure connection priorities by using the three priority levels in the Distributed File System snap-in:

- **High.** All connections marked **High** must successfully synchronize before FRS attempts to synchronize medium-priority connections.
- **Medium.** At least one medium-priority connection must successfully complete an initial synchronization before FRS attempts to synchronize any low-priority connections. FRS attempts to synchronize all connections in the medium-priority level, but only one must be successful before FRS attempts to synchronize low-priority connections.
- **Low.** The default connection priority is low. For these connections, FRS attempts to synchronize one time, but any failure does not delay other synchronization attempts. If no medium- or high-priority connections exist, at least one low-priority connection must succeed before FRS considers the initial synchronization operation complete. If medium- or high-priority connections exist, none of the connections in the low-priority class needs to succeed for FRS to consider the initial synchronization operation complete.

**Important**

- Connection priorities are used for initial synchronizations only on custom topologies. If you use a ring, hub and spoke, or full mesh topology, connection priorities are used only on reinitialization during a nonauthoritative restore operation.

When evaluating your topology to determine connection priority, identify the following details for each replica member:

- All inbound connections
- The speed of those connections

Based on this data, determine which connection priority to assign to each inbound connection. The following guidelines will assist you in choosing the appropriate connection priority:

- Use medium priority or high priority only for high-bandwidth connections.
- If you plan to create a high-priority connection, note that if this connection fails, no other synchronization takes place for lower-priority partners until all high-priority connections have succeeded.
- Use low priority for low-bandwidth or unreliable connections.

**Note**

- Keep in mind that after the initial synchronization, connection priorities are ignored until you need to perform a nonauthoritative restore to recover a failed replica member.

For an Excel spreadsheet to assist you in documenting connection priorities, see "FRS Configuration Worksheet" (Sdcfsv_2.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "FRS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

Figure 2.14 illustrates an organization's planned topology. The letters A through C represent inbound connections. The hub servers are connected by a high-bandwidth network connection, and two of the hub servers connect with branch servers by using low-bandwidth network connections.

**Figure 2.14 Sample FRS Topology**




In this topology, the administrator assigns connection priorities as follows:

- "A" connections use medium priority because the hubs are connected by high-bandwidth connections and the hub servers are expected to be up to date.

- "B" and "C" connections use low priority because the branch servers connect to the hubs across a low-bandwidth network.

When the hub servers in this example are deployed, Hub1 will be the initial master. Because the "A" connections are marked as medium priority, the two other hub servers will not attempt replication with the branch servers until initial replication completes with at least one other hub server.

The administrator specifies connections "B" and "C" as low priority because the branch servers have low-bandwidth network connections. In the event of a nonauthoritative restore for any of the hubs, using low priority causes the repaired hub server to replicate first from other hubs before replicating with the branch servers.

**Note**

- To increase the availability of data, the administrator can place multiple servers in each branch. If the servers in the branch use high-bandwidth connections to each other, the administrator can set medium priority for those inbound connections. That way, if one of the branch servers fails and is later restored, it will first attempt to replicate with a local branch server before replicating from the hub server across a slow network connection.

# Example: An Organization Designs a Replication Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization uses Group Policy to distribute 10 GB of software to users in four sites. The software changes nightly, and it is important that clients have access to these updates the following day. To make the software files and updates highly available, the organization creates a domain-based DFS namespace for each branch and then uses the Distributed File System snap-in to configure replication. To make the updated software available even if a hub server is offline, the organization uses redundant hub servers that are connected by high-speed LAN links. The organization also uses staggered replication schedules to distribute the load between the hub servers. Figure 2.15 describes the FRS topology.

**Figure 2.15 Redundant Hub and Spoke Topology for Software Distribution to Branch Offices**




The organization uses a separate namespace for each branch so that if the local branch server is not available, the clients transparently access one of the hubs instead of another branch server. The organization enables the least-expensive target selection feature of DFS so that the clients always choose the local branch server when it is available.

In this scenario, the organization configures replication as follows:

- Each branch has its own replica set, and each replica set contains three members: the branch and the two hubs.
- Replication between the hubs is enabled at all times, so changes at one hub are immediately replicated to the other hub.
- Inbound connections between the hub servers are medium priority; all other inbound connections are low priority.
- Odd-numbered branches replicate with Hub1 from 6:00 P.M. until 7:30 P.M. and with Hub2 from 8:00 P.M. until 9:30 P.M.
- Even-numbered branches replicate with Hub2 from 6:00 P.M. until 7:30 P.M. and with Hub1 from 8:00 P.M. until 9:30 P.M.

As shown in Figure 2.15, all software updates originate at the hub servers and are replicated to the branch servers on a staggered schedule. If replication completes when the branch is connected to its primary hub, no replication takes place when the branch connects to its alternate hub, because the data will have already replicated between the hubs. Similarly, if one of the hub servers fails or is taken offline for maintenance, the branch servers replicate with the remaining hub, and when the offline hub is brought back online, the two hubs replicate immediately to synchronize their data.

Because of the large amount (10 GB) of software files, the organization prestages new branch servers by restoring the software files to the new branch servers while they are located at the staging site. The organization then transports the new branch servers to the branch sites for deployment. By prestaging the software files, the organization does not need to replicate the

initial 10 GB of files across the network. For more information about prestaging replica sets, see "Adding a New Member to an Existing Replica Set" later in this chapter.

# Increasing Data Availability by Using Clustering

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The information presented in this section assumes an understanding of basic cluster terminology. If you are not familiar with clustering, you can refer to a number of sources for further information:

- For general information about server clusters, including terminology, procedures, checklists, and best practices, see Help and Support Center for Windows Server 2003.
- For information about designing server clusters, see "Designing and Deploying Server Clusters" in this book.

You can use clustered file servers to make sure that users can access business-critical data, such as home directories and software distribution shares. You can also use clustered servers to consolidate multiple nonclustered file servers onto a clustered file server. Each server is then re-created on the cluster as a virtual server.

**Important**

- Because multiple network names might belong to the same server, File Share resources might appear under these names and might also be accessible through them. To avoid client connectivity interruptions, make sure that clients use the network name that is associated with the group in which the File Share resource belongs. For more information about using the proper network name, see article Q170762, "Cluster Shares Appear in Browse List Under Other Names" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Choosing Virtual Server Names**

If you are migrating nonclustered file servers to clustered file servers, and you do not currently use (or plan to use) DFS, you should use the IP address and network name of each migrated server to create an identical virtual server on the clustered file server. Doing so allows users to continue to access data by using the server names that they are familiar with. You also ensure that links and shortcuts continue to work as they did before the migration.

**Choosing the File Share Resource Type**

A clustered file server uses the File Share resource type to make data available to users. You have three options when configuring this resource type:

**Normal share**

Normal shares function similarly to shared folders created on nonclustered file servers, except that you use the Cluster Administrator snap-in to create them. When you create a normal share, you publish a single folder to the network under a single name. Normal shares do not have required dependencies; however, if the normal share is located on cluster storage, the normal shares should, at a minimum, have dependencies on the cluster storage device (the Physical Disk or other storage-class resource), with the preferential addition of dependencies on the network name and IP address.

You can create a maximum of 1,674 resources, including File Share resources, on a cluster. To provide good failover performance, do not approach this limit in production environments. In addition, if you need to create a large number of normal shares, it is better to use the Share Subdirectories option.

**Share subdirectories (or hide subdirectories)**

A share subdirectories share publishes several network names: one each for a folder and all of its immediate subfolders. For example, if you create a share subdirectories share called Users, any folder you add to the Users folder is automatically shared.

**Important**

- When sharing subdirectories, do not use the same name for subfolders under different folders (for example, \folder1\name1 and \folder2\name1). The first shared subdirectory to come online will remain online, but the second shared subdirectory will not be initialized and will not come online.

Share subdirectories offer a number of advantages:

- They do not require you to create a File Share resource for every folder. As a result, you reduce the potential time and

CPU load needed to detect failures.

- They are an efficient way to create large numbers of related file shares on a single file server. For example, you can use this method to create a home directory for each user who has files on the server, and you can hide the subdirectory shares to prevent users from seeing the subdirectory shares when they browse the network.

- They allow administrators who are not experienced with server clusters or using the Cluster Administrator snap-in to easily create folders that are automatically shared.

When using the share subdirectories feature, determine the number of file shares you plan to host on the cluster, because the number of file shares affects server cluster capacity planning and failover policies. Specifically, you need to determine the following:

- The number of nodes you plan to have in the server cluster

- The number of node failures you want the server cluster to withstand while still providing acceptable performance

- Whether the remaining nodes can handle the load of the failed nodes

The number of nodes is important because a single node can support a limited number of file shares, which varies according to the amount of RAM in the server and which is described in "Reviewing File Server Limits" later in this chapter. If you want the cluster to be able to survive the loss of all but one node, make sure that the cluster hosts no more than the maximum number of file shares that can be hosted by a single node. This is especially important for two-node clusters, where the failure of one node leaves the single remaining node to pick up all the file shares.

In a four-node or eight-node cluster, you have other options that might be more appropriate, depending on the failure scenarios that you want to protect against. For example, if you want a four-node cluster to survive one node failing at any point, you can configure the file shares so that if one node fails, its file shares are spread across the remaining three nodes. In this scenario, each node can be loaded to 66 percent of the maximum number of file shares and still be within the maximum limit of a single node in the event of a single failure. In this case, the cluster can host three times the number of file shares that a single server can host. If you want to survive two nodes failing, a four-node cluster can hold twice as many files shares (because if two nodes fail, the remaining two nodes pick up the load from the two failed servers) and so on.

For more information about server cluster capacity planning and failover policies, see "Designing and Deploying Server Clusters" in this book.

**Note**

- For more information about using the Share Subdirectories option to create home directories, see article Q256926, "Implementing Home Folders on a Server Cluster," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**DFS root**

You can use the File Share resource type to create a resource that manages a stand-alone DFS root on a cluster. The DFS root File Share resource has required dependencies on a network name, which can be either the cluster name or any other network name for a virtual server. However, it is recommended that you do not use the cluster name (in the Cluster group) for any resources.

The Cluster service manages a resource group as a single unit of failover. Therefore, if you want four DFS roots to fail over independently, create them in different groups. Each group has a network name or the virtual server name with which the root is associated.

When creating a DFS root in clustered file servers, review the following issues:

- The name that you specify for one DFS root cannot overlap with the names for the other DFS roots in the same cluster. For example, if you name one DFS share C:\Dfsroots\Root1, you cannot create other roots using names that are derived from the first DFS root, such as C:\Dfsroots or C:\Dfsroots\Root1\Root2, and so forth.

- On server clusters, do not create clustered DFS roots that have the same name as nonclustered DFS roots or shared folders.

- Clustered file servers running Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition support multiple stand-alone DFS roots. These DFS roots can exist in multiple resource groups, and each group can be hosted on a different node.

- Clustered file servers running Microsoft® Windows® 2000 Advanced Server or Windows® 2000 Datacenter Server support one DFS root per server. Mixed-version clusters running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition on some nodes and Windows 2000 on others only support one DFS root per cluster as well. For more information about mixed-version clusters and rolling upgrades, see "Deploying Clustered File Servers" later in this chapter.

  **Caution**

  - Do not make DFS configuration changes (for example, creating new roots, adding new links, new link targets, and so on) while operating a mixed-version cluster, because all DFS changes are lost upon failover.

For more information about creating DFS roots on a clustered file server, see article Q301588, "HOW TO: Use DFS on a Server Cluster to Maintain a Single Namespace," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Using Shadow Copies on Server Clusters**

Windows Server 2003 supports creating shadow copies on cluster storage. Shadow copies are point-in-time copies of files that are stored on file servers running Windows Server 2003. By enabling shadow copies, you can reduce the administrative burden of restoring previously backed-up files for users who accidentally delete or overwrite important files.

If you plan to enable shadow copies on a server cluster, review the following issues:

- Cluster-managed shadow copies can be created only on cluster storage with a Physical Disk resource. In a cluster without cluster storage, shadow copies can be created and managed only locally.

- The recurring scheduled task that generates volume shadow copies must run on the same node that currently owns the storage volume (where the shadow copies are stored).

- The cluster resource that manages the scheduled task must be able to fail over with the Physical Disk resource that manages the storage volume.

- If you place the source volume (where the user files are stored) and the storage volume (where the shadow copies are stored) on separate disks, those disks must be in the same resource group.

- When you configure shadow copies using the procedure described in "Enable Shadow Copies for Shared Folders in a cluster" in Help and Support Center for Windows Server 2003, the resources and dependencies are set up automatically. Do not attempt to modify or delete these resources and dependencies directly.

To ensure availability, plan to enable shadow copies before you deploy the server cluster. Do not enable shadow copies in a deployed cluster. When you enable shadow copies, the shadow copy volumes (as well as all resources that directly or indirectly depend on the disk) go offline for a brief period while the resource dependencies are created. The shadow copy volumes are not accessible to applications and users while they are offline.

**Note**

- If you enable shadow copies before clustering your servers, you must disable, and then re-enable, shadow copies. If you cluster a disk containing a previously created shadow copy, the shadow copy might stop functioning after that disk's Physical Disk resource fails over.

For additional guidelines and information about shadow copies, see "Designing a Shadow Copy Strategy" later in this chapter. For in-depth information about using shadow copies on server clusters, see "Using Shadow Copies of Shared Folders in a server cluster" in Help and Support Center for Windows Server 2003.

# Example: An Organization Designs a Clustered File Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A large organization uses centralized file servers to store home directories for 3,000 users. The organization has an existing SLA that specifies its file servers must be 99.99 percent available, which means that the file servers can be offline no more than 53 minutes a year. To meet this level of availability, the organization implements clustered file servers in its data center, using a storage area network (SAN) for storage.

After evaluating storage and performance requirements, the organization determines that it needs one file server per 1,000 users, for a total of three file servers. The organization could implement a three-node cluster, but it decides to implement a four-node cluster instead. It will use three nodes to provide access to user directories and the fourth node to perform the following functions:

- Host a stand-alone DFS namespace to provide a unified view of the user directories on the remaining three nodes.
- Perform backups of the other three nodes in the evenings.
- Take over the resources of another node if the node fails or is taken offline for maintenance.

Figure 2.16 illustrates how the organization designs its four-node cluster.

**Figure 2.16 Four-Node Cluster**




To create the user directories, the administrator creates a File Share resource on three nodes and specifies the Share Subdirectories option. The administrator places each File Share resource in a virtual server (a resource group containing IP address and Network Name resources) and sets up dependencies between the File Share resource and the Network Name resource. The administrator then uses a script to quickly create the user directories, which are automatically shared after creation, and to set appropriate permissions for each user directory. The administrator creates another File Share resource on the fourth node and specifies the DFS root option and then uses the Distributed File System snap-in to create a link to each user directory. Users can access their shares by using the DFS path \\Home\Folders*Username*. Logon scripts also map user shares to each user's U: drive.

**Note**

- If the organization needs to provide access to user data when the complete site fails, due either to a total loss of power as a result of a natural disaster or another event, the organization can implement a geographically dispersed cluster.

For more information about designing server clusters, including capacity planning, failover policies, and geographically dispersed clusters, see "Designing and Deploying Server Clusters" in this book.

# Designing a Standard File Server Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To maximize the reliability, availability, and performance of file servers, design one or more standard hardware and software configurations for file servers in your organization. Using a standard configuration for file servers provides a number of benefits:

- You experience less complexity managing and maintaining the hardware. Instead of learning to perform hardware tasks on multiple types of file servers, you need only to learn the task once.
- You reduce the amount of testing that must be done when updating drivers or applications on the file server. Instead of testing the fixes on multiple hardware configurations, you can test the updates on one file server and then deploy the updates to the other file servers.
- You can keep fewer spares on-site. For example, if you use the same type of hard disks in all file servers, you need fewer spare disks, reducing the cost and complexity of providing spares.
- You reduce the number of disk images or answer files you need to maintain, which is useful if you are using automated installation technologies, such as unattended installation or Remote Installation Services (RIS), to deploy your file servers.

**Important**

- If you have read "Planning File Server Availability" earlier in this chapter and you have decided to implement clustered file servers, review the chapter "Designing and Deploying Server Clusters" in this book for important guidelines on choosing and configuring cluster hardware and storage.

Figure 2.17 outlines the process for designing a standard file server configuration.

**Figure 2.17 Designing a Standard File Server Configuration**



# Evaluating Existing Hardware

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you evaluate your existing hardware, you must determine if the hardware meets the recommended hardware requirements for Windows Server 2003 and verify that the hardware is listed in the Windows Server Catalog for Windows Server 2003. If your hardware does not meet these requirements, you might need to acquire new hardware or upgrade your existing hardware.

Your organization might also have existing standards in place that specify the standard configuration for file servers. If so, make sure that your file servers meet those requirements as well.

**Reviewing the Recommended Hardware Requirements and the Windows Server Catalog**
Microsoft recommends that file servers running Windows Server 2003 should meet the hardware requirements listed in Table 2.8.

**Table 2.8 Recommended Hardware Requirements for Windows Server 2003**

| Operating System | CPU | RAM |
|---|---|---|
| Windows Server 2003, Standard Edition | 550 MHz | 256 MB |
| Windows Server 2003, Enterprise Edition | 550 MHz | 256 MB |
| Windows Server 2003, Datacenter Edition | 550 MHz* | 512 MB minimum |

*For servers on which you are upgrading from Windows 2000 Datacenter Server, the minimum processor speed is 400 MHz.

The section "Determining RAM and CPU Specifications" later in this chapter provides specific guidelines for how increasing CPU and RAM can improve file server performance. If the server provides additional services or acts as a domain controller, the server must meet higher requirements than those listed here. For more information about hardware requirements for domain controllers, see "Planning Domain Controller Capacity" in *Designing and Deploying Directory and Security Services* of this kit.

In addition to verifying that your hardware meets the recommended requirements for file servers, verify that all hardware devices on the file server are listed in the Windows Server Catalog. Only Designed for Windows Server 2003 hardware is included in the Windows Server Catalog, which means that the hardware is designed to take advantage of new features in Windows Server 2003.

For more information about hardware requirements for server clusters, see "Designing and Deploying Server Clusters" in this book.

Before you install Windows Server 2003, you can run a hardware and software compatibility check by running the Windows Upgrade Advisor from the Windows Server 2003 operating system CD. The compatibility check does not require you to actually begin an installation. To run the check, insert the Windows Server 2003 operating system CD in the CD-ROM drive and, when a display appears, follow the prompts for checking system compatibility. You will be offered the option to download the latest Setup files (by using Dynamic Update) when you run the check. If you have Internet connectivity, it is recommended that you allow the download.

Another way to run the Windows Upgrade Advisor is to insert the Windows Server 2003 operating system CD in the CD-ROM drive, open a command prompt, and type:

**d:\I386\winnt32 /checkupgradeonly**

In this example, **d** represents the CD-ROM drive.

For more information about hardware and software compatibility, see the following resources:

- To download the Windows Upgrade Advisor, see the Windows Upgrade Advisor link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- For more information about hardware compatibility, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- For more information about application compatibility, see "Planning and Testing for Application Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

**Verifying That Mass Storage Controllers Are Supported**

If your file servers use a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for disks, confirm that it is compatible with products in the Windows Server 2003 family.

If your controller is compatible with products in the Windows Server 2003, but you are aware that the manufacturer has supplied a separate driver file for use with your operating system, obtain the file (on a floppy disk) before you begin Setup. During the early part of Setup, a line at the bottom of the screen prompts you to press F6. Further prompts guide you in supplying the driver file to Setup so that it can gain access to the mass storage controller.

If you are not sure whether you must obtain a separate driver file from the manufacturer of your mass storage controller, you can try running Setup. If the controller is not supported by the driver files on the Windows Server 2003 operating system CD, and it therefore requires a driver file supplied by the hardware manufacturer, Setup stops and displays a message saying that no disk devices can be found, or it displays an incomplete list of controllers. After you obtain the necessary driver file, restart Setup, and press F6 when prompted.

**Verifying That Your Hardware Meets Your Requirements**

If your organization has a standard configuration for file servers, verify that your existing file servers meet these requirements. If they do, proceed to "Planning File Server Security" later in this chapter. If your organization plans to update its standard file server configuration, continue to "Choosing the File Server Type."

**Identifying Servers to Use as Windows Server 2003 File Servers**

Identify the file servers on which you plan to run Windows Server 2003, as well as the servers that you plan to decommission as a result of consolidation or failure to meet hardware requirements.

# Choosing the File Server Type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Two types of servers run Windows Server 2003: general-purpose servers and Windows Powered Network Attached Storage.

**General-purpose servers** A general-purpose server can function in a number of roles, including domain controller, print server, WINS server, DHCP server, application server, or an e-mail or database server. General-purpose servers provide administrators with the flexibility to configure server software and hardware as they choose. Administrators perform the necessary configuration and optimization that the server requires to perform its role.

**Windows Powered Network Attached Storage** Windows Powered Network Attached Storage solutions are dedicated file servers running Windows 2000 Server or Windows Server 2003. With built-in support for UNIX and Apple file protocols, Windows Powered Network Attached Storage solutions offer seamless integration in a heterogeneous environment.

Table 2.9 describes the features and capabilities of each file server type.

**Table 2.9 Features and Capabilities of File Server Types**

| Description | General-Purpose Server | Windows Powered Network Attached Storage |
|---|---|---|
| Comes preinstalled with Windows 2000 or Windows Server 2003. | ●* | ● |
| Offers the file system, security, reliability, and scalability features of Windows 2000 or Windows Server 2003. | ● | ● |
| Can provide services for purposes other than file sharing, such as domain controller or print, WINS, or DHCP server. | ● | |
| Can run applications, such as Microsoft® SQL Server 2000 or Exchange 2000 Server. | ● | |
| Comes preconfigured with NFS, File Transfer Protocol (FTP), AppleTalk, HTTP, WebDAV, and NetWare protocol support. | | ● |
| Can be deployed in as little as 15 minutes. | | ● |

* Some general-purpose servers come preinstalled with the appropriate operating system; however, other general-purpose servers require you to configure the volumes and install the operating system.

Some hardware vendors offer Windows Powered Network Attached Storage solutions in a cluster configuration. If your organization requires highly available servers, consult your hardware vendor for more information.

# Reviewing File Server Limits

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you plan your file server configuration, keep in mind file system, storage, and other limits related to file servers. Table 2.10 describes these limits.

**Table 2.10 File System, Storage, and File Server Limits for Windows Server 2003**

| Description | Limit |
|---|---|
| Maximum size of a basic volume | 2 TB |
| Maximum size of a dynamic volume | 2 TB for simple and mirrored (RAID-1) volumes.<br>Up to 64 TB for spanned and striped (RAID-0) volumes. (2 TB per disk with a maximum of 32 disks per volume.)<br>Up to 62 TB for RAID-5 volumes. (2 TB per disk with a maximum of 32 disks per volume and 2 TB used for parity.) |
| Maximum number of dynamic volumes per disk group | 1,000<br>A disk group is collection of dynamic disks. Windows Server 2003 supports one disk group per server. |
| Maximum size of an NTFS volume | $2^{32}$ clusters minus 1 cluster<br>Using a 64-kilobyte (KB) cluster (the maximum NTFS cluster size), the maximum size of an NTFS volume is 256 TB minus 64 KB.<br>Using a 4-KB cluster (the default NTFS cluster size), the maximum size of an NTFS volume is 16 TB minus 4 KB. |
| Maximum file size on an NTFS volume | 16 TB ($2^{44}$ bytes) minus 64 KB |
| Maximum number of files on an NTFS volume | 4,294,967,295 ($2^{32}$ minus 1 file)<br>There is no limit to the number of files that can be stored in a folder. For recommendations on limiting the number of files stored on a volume, see "Determining Maximum Volume Size" later in this chapter. |
| Maximum number of clusters on an NTFS volume | 4,294,967,296 ($2^{32}$) |
| Maximum volumes per server | Approximately 2,000 volumes.<br>Up to 1,000 of these volumes can be dynamic volumes; the rest are basic volumes. Boot times increase as you increase the number of volumes. In addition, you must use mounted drives to access volumes when all drive letters on a server have been used. For more information about mounted drives, see "Using NTFS mounted drives" in Help and Support Center for Windows Server 2003. |

| | |
|---|---|
| Maximum number of shared folders on a server | Varies. The number of shares on a server affects server boot time. On a server with typical hardware and thousands of shares, boot time can be delayed by minutes. Exact delays depend on server hardware.<br><br>Shared folder information is stored in the system hive of the registry. For systems with less than 800 MB of RAM, the System hive can be as large as one-quarter of the physical memory. For systems with more than 800 MB of RAM, the maximum size of the System hive is 200 MB. If the system hive exceeds this limit, the server cannot mount the registry at startup and Windows Server 2003 cannot start. |

For information about optimizing NTFS performance, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Determining RAM and CPU Specifications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As hardware continues to improve, hardware vendors frequently change the RAM and CPU configurations in the servers they offer. Use the information about performance improvements and scaling factors in the sections that follow to choose the hardware that best meets your organization's needs.

When designing RAM and CPU specifications for server clusters, also consider the failover policy that you plan to use and the maximum number of File Share resources and share subdirectories that will be hosted on any one node after a failure. Each node must have the RAM and CPU resources required to host the resources of one or more failed nodes, depending on your failover policy. For more information about failover policies and server cluster capacity planning, see "Designing and Deploying Server Clusters" in this book.

**Note**

- Many of the figures presented in this section are derived from NetBench statistics for file server throughput. NetBench is a portable Ziff Davis Media benchmark program that measures how well a file server handles file I/O requests from 32-bit Windows clients. NetBench provides an overall I/O throughput score and average response time for servers, along with individual scores for clients. You can use these scores to measure, analyze, and predict how well your server can handle file requests from clients.

**Reviewing Windows Server 2003 CPU Specifications**

Table 2.11 describes the recommended CPU speed and number of processors supported by Windows Server 2003.

**Table 2.11 CPU Requirements for Windows Server 2003**

| Specification | Windows Server 2003, Standard Edition | Windows Server 2003, Enterprise Edition | Windows Server 2003, Datacenter Edition |
|---|---|---|---|
| Minimum recommended CPU speed | 550 MHz | 550 MHz | 550 MHz |
| Number of CPUs supported | 1–4 | 1–8 | 8–32 |

**Reviewing Operating System Performance Improvements**

Even if you plan to use existing hardware to run Windows Server 2003, you can benefit from performance enhancements available in Windows Server 2003, as well as client and server protocol improvements available when using clients running Windows XP Professional.

Table 2.12 describes performance improvements that can be gained by migrating to new operating systems on identical hardware.

**Table 2.12 Operating System Performance Improvements on the Same Hardware**

| Current Server and Client Operating Systems | New Server and Client Operating Systems | Improvement Factor |
|---|---|---|
| Windows NT Server 4.0 with Microsoft® Windows NT® Workstation 4.0 clients | Windows 2000 Server with Windows 2000 Professional clients | Up to 1.25X |
| Windows 2000 Server with Windows 2000 Professional clients | Windows Server 2003 with Windows XP Professional clients | Up to 2.2X |
| Windows NT Server 4.0 with Windows NT Workstation 4.0 clients | Windows Server 2003 with Windows XP Professional clients | Up to 2.75X |

These figures are based on the following assumptions:

- The server is uniprocessor (UP), 2P, 4P, or 8P.

- For each comparison, the server hardware is the same.
- No memory, disk, or network bottlenecks prevent the processor from performing at full capacity.

**Reviewing Performance Improvements Gained by Upgrading Processors**

Table 2.13 describes performance improvements that can be gained by upgrading server processors. Upgrading processors improves processing power, memory bandwidth, I/O bandwidth, and the system bus. These figures are based on actual processor improvements, not operating system improvements. If you plan to use processors that are faster than those listed here, performance will be greater than the following figures show.

**Table 2.13 Performance Improvements Gained by Upgrading Processors**

| Old Processor | New Processor (Server Class) | Improvement Factor |
|---|---|---|
| 200 MHz Intel Pentium Pro | 400 MHz Intel Pentium II Xeon | 2X |
| 400 MHz Intel Pentium II Xeon | 900 MHz Intel Pentium III Xeon | 2X |
| 200 MHz Intel Pentium Pro | 900 MHz Intel Pentium III Xeon | 4X |

**Reviewing Performance Improvements Gained by Adding Processors**

To increase performance, consider using more than one processor in your file servers. One advantage of using multiple processors is the ability to handle more concurrent clients, resulting in higher scaling factors at high client loads. Table 2.14 describes the NetBench throughput improvements gained by adding processors on file servers running Windows Server 2003.

**Table 2.14 Performance Improvements Gained by Adding Processors**

| Original Number of Processors | After Upgrade | Scaling Factor* |
|---|---|---|
| 1 | 2 | 1.4X to 1.6X |
| 2 | 4 | 1.3X to 1.4X |
| 1 | 4 | 1.8X to 2.3X |
| 4 | 8 | 1.3X to 1.4X |
| 1 | 8 | 2.4X to 3.2X |

* The scaling factors are based on a range of client loads.

**Determining the Client Load Based on Processor Utilization**

NetBench stresses the system by applying a heavy load on the file server. Microsoft used the most intensive CPU operations — file opens and file creates (subsequently described as *opens/creates*) — to translate a NetBench client load to a more realistic client load.

To determine the client load, Microsoft observed approximately 820 opens/creates per second at peak NetBench throughput (100-percent CPU utilization) on a Xeon 900-MHz server with a single processor (UP). Table 2.15 describes active client loads for light, medium, and heavy user loads at 70-percent CPU utilization. The client loads are defined as follows:

- Light: one open/create every 10 seconds
- Medium: two opens/creates every 10 seconds
- Heavy: three opens/creates every 10 seconds

Assuming that a light user load causes one open/create every 10 seconds, a UP Xeon 900-MHz server can handle 5,700 users at 70-percent CPU utilization and 8,200 users at 100-percent CPU utilization. (These figures are derived by dividing the opens/creates per second at peak NetBench throughput by the opens/creates per second for light users.)

The figures in Table 2.15 for the UP Xeon 900-MHz server are based on the following assumptions:

- No memory, disk, or network bottlenecks prevent the processor from performing at 100-percent capacity.
- The clients are running Windows XP.

The figures for 2P, 4P, and 8P Xeon 900-MHz servers were calculated by using the scaling factors described in Table 2.14. For example, on a 4P Xeon 900-MHz server, under a heavy client load of three opens/creates every 10 seconds, the server can handle 3,400 to 4,400 active users. This figure is derived by taking the 1,900 heavy-load users supported on a UP Xeon 900-MHz server and multiplying that figure by the UP-to-4P scaling factors of 1.8X to 2.3X provided in Table 2.14.

**Table 2.15 Number of Active Users Supported Based on a NetBench-Type Workload**

| Processor | Heavy Load | Medium Load | Light Load |
|---|---|---|---|
| UP Xeon, 900 MHz | 1,900 | 2,800 | 5,700 |
| 2P Xeon, 900 MHz | 2,600 to 3,000 | 3,900 to 4,500 | 8,000 to 9,100 |
| 4P Xeon, 900 MHz | 3,400 to 4,400 | 5,000 to 6,400 | 10,300 to 13,100 |
| 8P Xeon, 900 MHz | 4,600 to 6,100 | 6,700 to 9,000 | 13,700 to 18,200 |

**Determining RAM Specifications**

Using adequate RAM in file servers ensures that Windows Server 2003 can temporarily cache (store) files in memory, reducing the need to retrieve files from disk. Table 2.16 describes the minimum recommended RAM and maximum RAM for Windows Server 2003.

**Table 2.16 Minimum and Maximum RAM for Windows Server 2003**

| RAM Specification | Windows Server 2003, Standard Edition | Windows Server 2003, Enterprise Edition | Windows Server 2003, Datacenter Edition |
|---|---|---|---|
| Minimum recommended RAM | 256 MB | 256 MB | 512 MB minimum |
| Maximum RAM | 4 GB | 32 GB | 64 GB |

To determine the amount of RAM required to support the file server workload, review the number of remote file handles that can be efficiently supported by a file server running Windows Server 2003. Next, review how additional RAM affects the total size of files, or *file set size*, that can be held in memory at any time.

**Remote concurrent file handles**

A file server running Windows Server 2003 with 1 GB of RAM can efficiently support approximately 100,000 remote concurrent file handles, regardless of the size of the files. If your users are likely to have more than 100,000 files open at a time, plan to split this load across two or more servers.

**File set size**

On a file server with 1 GB of RAM, Windows Server 2003 can hold approximately 500 MB of file content and NTFS metadata in memory. (The amount of memory used for NTFS metadata depends on the depth of the directory hierarchy and query distribution, among other factors.) Windows Server 2003 uses the rest of the RAM for providing nonpaged pool and other operating system functions. For each additional gigabyte of RAM that you add, Windows Server 2003 can use the entire RAM capacity for storing file content in memory. For example, a file server with 3 GB of RAM can support approximately 2.5 GB of file content in memory. When the file set size exceeds the amount of memory, files are paged to disk. This paging can result in disk bottlenecks, though using a fast disk subsystem can alleviate this problem.

When determining how much RAM you plan to install in file servers, consider the following guidelines:

- When users typically access the same files, the file set is known as "hot," because the files are frequently stored in memory. For hot file sets, invest in more RAM to accommodate the entire hot file set. Typically, hot file sets are less than 1 percent of the file set, although this figure can vary.
- When users access random files, the file set is known as "cold." For cold file sets, invest in faster disks, because users

typically open files that are not already in memory, and the response time is limited by disk latency. For example, if you can cut disk latency in half by using a faster disk subsystem (including number of disks, mechanical speed, and disk cache), compare the cost of doing so to the amount of RAM it would take to achieve similar performance. Using a faster disk subsystem might be less expensive.

# Determining Storage Specifications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to make sure that file servers have adequate, fault-tolerant storage. When determining storage specifications, perform the following tasks:

1. Estimate the amount of data to be stored.
2. Plan the amount of storage for each server configuration.
3. Determine the maximum volume size.
4. Plan the layout and RAID level of volumes.

The following sections describe each of these steps.

For more information about storage requirements for server clusters, see "Designing and Deploying Server Clusters" in this book.

# Estimating the Amount of Data to Be Stored

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you estimate the amount of data to be stored on a new file server, include the following information:

- The amount of data currently stored on any file servers that will be consolidated onto the new file server.
- If the file server will be a replica member, the amount of replicated data that will be stored on the new file server.
- The amount of data that you will need to store on the file server in the future.

A general guideline is to plan for faster growth than you experienced in the past. Investigate whether your organization plans to hire a large number of people and whether any groups in your organization are planning large projects that require extra storage.

You must also take into account the amount of space used by operating system files, applications, RAID redundancy, log files, and other factors. Table 2.17 describes some factors that affect file server capacity.

**Table 2.17 Factors That Affect File Server Capacity**

| Factor | Storage Capacity Required |
|---|---|
| Operating system files | At least 1.5 GB. To allow space for optional components, future service packs, and other items, allow an additional 3 GB to 5 GB for the operating system volume. |
| Paging file | 1.5 times the amount of RAM by default. |
| Memory dump | Depending on the memory dump file option that you choose, the amount of disk space required can be as large as the amount of physical memory plus 1 MB. |
| Applications | Varies according to the application, which can include antivirus, backup, and disk quota software; database applications; and optional components, such as Recovery Console, Windows Services for UNIX, and Windows Services for NetWare. |
| Log files | Varies according to the application that creates the log file. Some applications allow you to configure a maximum log file size. You must make sure that you have adequate free space to store the log files. |
| RAID solution | Varies. For more information about the RAID solutions, see "Planning the Layout and RAID Level of Volumes" later in this chapter. |
| Shadow copies | Ten percent of the volume by default, although increasing this size is recommended. For more information about shadow copies, see "Designing a Shadow Copy Strategy" later in this chapter. |
| Snapshots available on Windows Powered Network Attached Storage solutions | Some Windows Powered Network Attached Storage solutions support snapshots that are similar to the shadow copies feature available in Windows Server 2003. The amount of space that these snapshots consume varies according to the frequency of the snapshots and the amount of changed data. The default setting is 20 percent of the volume. |

# Planning Storage for Each Server Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The amount of disk space that is supported by a file server can be in the terabytes if the server is connected to high-capacity, direct-attached storage arrays or SAN-connected external storage arrays. However, your goal should not be to attach as much storage as possible to a file server. Instead, plan the amount of storage on a file server by examining your current backup solution and answering the following questions:

- How much data can your backup solution contain? Your backup solution should have the capacity to back up the entire file server.
- How much time do you have to complete your backup? You must be able to perform a full backup within a period of time that is acceptable to your organization and users. For organizations that use hardware-based snapshots and backup software that can back up open files, the backup window might not be an issue.
- How quickly do you need to restore data? In other words, do you need to be able to restore data and get the file server up and running within eight hours? Twenty-four hours? Verify that you can restore data within the required time.

If your organization provides file services to groups who have different storage requirements, you can design file server configurations with different amounts of storage. For example, a low-end file server might have 100 GB of storage, but a high-end file server might have 400 GB of storage. Because the high-end server will typically support more users, make sure that the RAM and CPU configuration is more robust than for the low-end server.

**Important**

- If you choose file server hardware that supports future storage expansion, answer the previous questions taking the additional storage into account. In other words, if you add more storage to an existing file server but cannot back up or restore the file server in the required time, you must deploy additional file servers or acquire a backup solution that can handle the extra storage.

Backup solutions are becoming faster and more sophisticated. If your current backup solution is not meeting your organization's needs, investigate new solutions, such as those that provide the following types of benefits:

- Many backup programs, including the Backup program that is available in Windows Server 2003, can back up files that are open, allowing administrators to back up a server without having to disconnect users or shut down processes that might have files open.
- Some backup programs offer software-based snapshots that take point-in-time images of a volume, creating virtual replicas without physically copying the data. Servers running Windows Server 2003 provide a similar feature, known as *shadow copies*. For more information about shadow copies, see "Designing a Shadow Copy Strategy" later in this chapter.
- Hardware vendors offer hardware-based snapshots (also called *split mirrors*, *snapshot mirrors*, or *clones*) that you can use to create mirror images of volumes to be used for online backup, application development, and testing purposes.
- Backup solutions designed for SAN-connected external storage arrays can perform backups without sending the data across the network. (These are known as *LAN-free backups*.) SAN backup solutions also offer serverless backups that move data directly from disk to tape.

# Determining Maximum Volume Size

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Volume sizes often vary, but it is important to set a maximum volume size. NTFS supports volumes up to 256 TB minus 64 KB ($2^{32}$ clusters minus 1 cluster), but it is recommended that you take the following factors into account when you determine the size of volumes in your file servers.

**Backup and restoration times**
Sizing your volumes to be the same size or smaller than your backup solution's capacity is a good guideline but not an absolute rule. For example, some organizations might back up files on a per-folder basis rather than on a per-volume basis. In this case, the organization would make sure that the size of each folder is smaller than the backup solution. This method, however, requires more backups — one for every folder. In the event that you must restore the data, this method also requires more restorations to restore all the folders on the volume, increasing the overall complexity and length of the process. Because emergency restorations are often done in a hurry and under pressure, let restoration time and simplicity be your guide for sizing volumes.

**Chkdsk times**
Although file system errors are rare on NTFS volumes, you need to consider the time required to run Chkdsk to repair any errors that occur. Chkdsk times are determined by the number of files on the volume and by the number of files in the largest folder. Chkdsk performance has improved significantly since Windows NT 4.0 and Windows 2000. As a result, downtime due to Chkdsk should be minimal. However, it is important to avoid having volumes that contain so many files that Chkdsk requires longer to complete than the amount of downtime your users can tolerate.

To determine how long Chkdsk takes to complete, run Chkdsk on a test server with a similar number of files as those stored on a typical file server volume. Based on these results, you can limit the number of files to be stored on a volume so that Chkdsk can complete within the acceptable time limit.

**Important**

- Windows Server 2003 provides some Chkdsk parameters that shorten the time required to run Chkdsk. However, if the volume contains file system errors, you must consider the volume at risk until you are able to run a full Chkdsk. For more information about running Chkdsk, see "Chkdsk" in Help and Support Center for Windows Server 2003.

If you need to store a large number of files on a file server and are concerned about Chkdsk times, you can distribute the files on separate volumes on the file server. If, however, you need those files to appear as though they are on a single volume, you can use mounted drives. When you create a mounted drive, you mount one local volume at any empty folder on another local NTFS volume. For example, you can create a folder called Data on volume E and then mount volume F to that folder. When you open the Data folder, the contents of volume F are displayed.

Mounted drives are also useful when you want to add more storage to an existing volume without having to extend the volume. For more information about mounted drives, including information on creating mounted drives on clustered file servers, see "Using NTFS mounted drives" in Help and Support Center for Windows Server 2003.

# Planning the Layout and RAID Level of Volumes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you design a standard hardware configuration, determine what type of data you plan to store on the file server, the RAID type and level that is appropriate for each data type, and how you plan to divide the disk space into volumes.

**Evaluating Data Types**

Table 2.18 describes the types of data that are typically stored on file servers and the disk I/O characteristics of the data. Knowing the I/O characteristics can help you choose the RAID level that offers the best performance for that particular type of I/O. If you have other types of data stored on the file server, be sure to estimate their I/O characteristics as well.

**Table 2.18 Data Types and Their Characteristics**

| Data Type | Description | Disk I/O Characteristics |
|---|---|---|
| Operating system | The operating system, drivers, and any applications installed on the server, such as antivirus, backup, or disk quota software. | Server startup consists mostly of large reads, because the data required for startup is optimized on the disk. After startup completes, the disk I/O is mostly small reads. |
| Paging space | Space used by the paging file. | If the server memory is sized correctly, the disk I/O consists of small reads and writes. If the server experiences heavy reads and writes, the server needs more memory. |
| User and shared data | Documents, spreadsheets, graphics, and other data created by users and stored on the file server. | For user and shared data, disk I/O must be measured on a workload-by-workload basis. |
| Application files | Software that users install or run from the network. | For application files, disk I/O must be measured on a workload-by-workload basis. |
| Log files | Files used by database, communications, or transaction applications to store a history of operations (for example, the FRS log file). | Small or large writes, depending on the application. |

**Choosing Between Hardware and Software RAID**

Many computer manufacturers provide server-class computers that support hardware-based RAID. With hardware RAID, you use the configuration utility that is provided by the hardware manufacturer to group one or more physical disks into what appears to the operating system as a single disk, sometimes called a *virtual disk* or *logical unit (LUN)*. When you create the virtual disk, you also select a RAID level. Hardware RAID typically provides a choice of RAID-0 (striped), RAID-1 (mirrored), RAID-5 (striped with parity), and in some servers, RAID-0+1 (mirrored stripe). After creating these virtual disks, you use the Disk Management snap-in or the command-line tool DiskPart.exe in Windows Server 2003 to create one or more volumes on each virtual disk.

If your server hardware does not have built-in hardware RAID support, you can use the software RAID provided in Windows Server 2003 to create RAID-0 volumes, RAID-1 volumes, and RAID-5 volumes on dynamic disks. Although software RAID has lower performance than hardware RAID, software RAID is inexpensive and easy to configure because it has no special hardware requirements other than multiple disks. If cost is more important than performance, software RAID is appropriate. If you plan to use software RAID for write-heavy workloads, use RAID-1, not RAID-5.

**Note**

- Before you can use the software RAID in Windows Server 2003, you must convert basic disks to dynamic disks. In addition, dynamic disks are not supported on cluster storage. For more information about basic disks, dynamic disks, software RAID, and disk management, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Choosing the RAID Level**

Choosing the RAID level is a trade-off among the following factors:

- Cost
- Performance
- Availability
- Reliability

You can determine the best RAID level for your file servers by evaluating the read and write loads of the various data types and then deciding how much you are willing to spend to get the performance, reliability, and availability that your organization requires. Table 2.19 describes four common RAID levels; their relative costs, performance, and availability; and their recommended uses. The performance descriptions in Table 2.19 compare RAID performance to the performance of *just a bunch of disks (JBOD)* — a term used to describe an array of disks that is not configured using RAID.

**Table 2.19 Comparing RAID Levels**

| **Factors** | **RAID-0** | **RAID-1** | **RAID-5** | **RAID-0+1** |
|---|---|---|---|---|
| Minimum number of disks | 2 | 2 | 3 | 4 |
| Usable storage capacity | 100 percent | 50 percent | (N-1)/N disks, where N is the number of disks | 50 percent |
| Fault tolerance | None. Losing a single disk causes all data on the volume to be lost. | Can lose multiple disks as long as a mirrored pair is not lost. | Can tolerate the loss of one disk. | Can lose multiple disks as long as a mirrored pair is not lost. |
| Read performance | Generally improved by increasing concurrency. | Up to twice as good as JBOD (assuming twice the number of disks). | Generally improved by increasing concurrency. | Improved by increasing concurrency and by having multiple sources for each request. |
| Write performance | Generally improved by increasing concurrency. | Between 20 percent and 40 percent worse than JBOD for most workloads. | Worse unless using full-stripe writes (large requests). Can be as low as 25 percent of JBOD. | Can be worse or better depending on request size. |
| Best use | • Temporary data only | • Operating system<br>• Log files | • Operating system<br>• User and shared data[1]<br>• Application files[2] | • Operating system<br>• User and shared data<br>• Application files<br>• Log files |

[1] Place user and shared data on RAID-5 volumes only if cost is the overriding factor.

[2] Place application files on RAID-5 volumes only if cost is the overriding factor.

When determining the number of disks to compose a virtual disk, keep in mind the following factors:

- Performance increases as you add disks.
- Mean time between failure (MTBF) decreases as you add disks to RAID-5 or RAID-0 arrays.
- RAID-0+1 almost always offers better performance than RAID-1 when you use more than two disks.

- Usable storage capacity increases as you add disks, but so does cost.

**Using Dynamic Disks with Hardware RAID**

Using dynamic disks with hardware RAID can be useful in the following situations:

- You want to extend a volume, but the underlying hardware cannot dynamically increase the size of LUNs.
- You want to extend a volume, but the hardware has reached its maximum LUN size.

Before converting hardware RAID disks to dynamic disks, review the following restrictions:

- You cannot use dynamic disks on cluster storage. However, you can use the command-line tool DiskPart.exe to extend basic volumes on cluster storage. For more information about extending basic volumes, see "Extend a basic volume" in Help and Support Center for Windows Server 2003.
- If you plan to use hardware snapshots of dynamic disks, you cannot import the snapshot again on the same server, although you can move the snapshot to other servers.
- If you create a software RAID-0 volume across multiple hardware arrays, you cannot extend the RAID-0 later to increase its size. If you anticipate needing to extend the volume, create a spanned volume instead.

For more information about dynamic disks, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit). For best practices on using dynamic disks, see article Q329707, "Best Practices for Using Dynamic Disks on Windows 2000-Based Computers," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Choosing the Cluster Size**

A *cluster* (or *allocation unit*) is the smallest amount of disk space that can be allocated to hold a file. NTFS uses 4-KB clusters for volumes over 2 GB. Using a 4-KB cluster size allows you to use NTFS compression on the volume. However, in some cases you might want to use a larger cluster size to optimize NTFS performance.

To determine the cluster size, evaluate the types of files to be stored on the volume so that you can determine whether to use the default 4-KB cluster size. Some important questions to answer include the following:

- Are the files typically the same size?
- Are the files smaller than the default cluster size?
- Do the files remain the same size or grow larger?

If the files are typically smaller than the default cluster size (for example, 4 KB) and do not increase, use the default cluster size to reduce wasted disk space. However, smaller clusters can increase fragmentation, especially when files grow to fill more than one cluster. Therefore, plan to adjust the cluster size accordingly when you format the volume. If the files you store tend to be large or increase in size, and you do not want to use NTFS compression, use 16-KB or 32-KB clusters to optimize performance.

**Important**

- If you plan to defragment volumes on which shadow copies are enabled, it is recommended that you use a cluster size of 16 KB or larger. If you do not, the number of changes caused by defragmentation can cause shadow copies to be deleted faster than expected. Note, however, that NTFS compression is supported only if the cluster size is 4 KB or smaller. For more information about shadow copies, see "Designing a Shadow Copy Strategy" later in this chapter.

**Determining the Volume Layout**

When you determine the volume layout, you evaluate the type of data to be stored and the number of volumes that you want to create. For better manageability, use separate volumes for each data type where possible, though the operating system and paging file are often placed on a single volume. For example, use one volume for the operating system and paging file and one or more volumes for shared user data, applications, and log files.

If performance is important, place different data types in separate volumes on different virtual disks. Using separate virtual disks is especially important for any data types that create heavy write loads, such as log files. For example, dedicate a single set of disks (composing a virtual disk) to handling the disk I/O created by the updates to the log file. Placing the paging file on a separate virtual disk can provide some minor improvements in performance, but not typically enough to make it worth the extra cost.

To gain some performance benefits while minimizing cost, it is often useful to combine different data types in one or more volumes on the same virtual disks. A common method is to store the operating system and paging file on one virtual disk and the user data, applications, and log files in one or more volumes on the remaining virtual disk.

**Note**

- On servers running Windows Server 2003, you can increase the size of basic volumes and dynamic volumes. If you plan to use basic volumes (primary partitions or logical drives), you can extend them on the same disk only, and the basic volume must be followed by contiguous unallocated space. If your hardware supports adding more physical disks to a virtual disk, make sure that the basic volume is the last volume on the disk. For more information about extending volumes and disk management, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Determining the Number of File Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following section describes reasons why you might want to deploy additional file servers and how you can determine the number of file servers that you require.

**Consolidating Older File Servers**
Many organizations today are consolidating older file servers throughout the organization into fewer larger, more powerful file servers. Consolidation reduces the cost of managing multiple file servers and increases the efficiency of storage allocation and backup tasks.

If you plan to consolidate file servers running Windows NT 4.0 or Windows 2000 onto new file servers running Windows Server 2003, determine the number of file servers running Windows Server 2003 that is necessary to match or exceed the performance of your existing file servers. You can derive these figures by using the scaling factors provided in "Determining RAM and CPU Specifications" earlier in this chapter.

The following scenarios describe two consolidation situations based on the scaling factors presented in "Determining RAM and CPU Specifications" earlier in this chapter.

**Scenario 1: Consolidating file servers running Windows NT Server 4.0**
An organization plans to migrate its file services to file servers running Windows Server 2003. It currently has 100 Windows NT 4.0 file servers with the following hardware:

- UP Pentium Pro, 200 MHz
- 256 MB RAM
- Fast Ethernet network adapters

The organization plans to use the following hardware to consolidate file servers:

- 4P 900-MHz Intel® Pentium® III Xeon™
- 3 GB RAM
- Gigabit network adapters

To determine the number of file servers the organization needs when migrating to Windows Server 2003, use the scaling factors described in Table 2.20.

**Table 2.20 Consolidating File Servers Running Windows NT Server 4.0**

| Area of Improvement | Details | Scaling Factor |
|---|---|---|
| Operating system | As shown in Table 2.12, migrating from Windows NT 4.0 to Windows Server 2003 provides a 2.75X performance improvement. | 2.75X |
| Hardware upgrade | As shown in Table 2.13, migrating from a Pentium Pro, 200 MHz, to a Xeon, 900 MHz, provides a 4X performance improvement. | 4X |
| Processor scaling | As shown in Table 2.14, migrating from UP (one processor) servers to 4P servers provides a 1.8X performance improvement. | 1.8X |
| Total improvements | Multiply the scaling factors (2.75 x 4 x 1.8) | 19.8X |

Based on these figures, the organization can migrate 100 file servers running Windows NT 4.0 to five 4P Xeon 900-MHz

servers running Windows Server 2003. (To derive this number, divide 100 by 19.8.) The organization must also be sure that the five Xeon servers have adequate storage space to store the consolidated data, as well as future data, and that 3 GB of memory can handle the total workload of the 100 servers.

**Scenario 2: Consolidating file servers running Windows 2000 Server**
An organization plans to migrate its file services to file servers running Windows Server 2003. It currently has five Windows 2000 file servers with the following hardware:

- 4P 400-MHz Intel® Pentium® II Xeon™
- 1 GB RAM
- Gigabit network adapters

The organization plans to use the following hardware to consolidate file servers:

- 8P 900-MHz Intel® Pentium® III Xeon™
- 4 GB RAM
- Gigabit network adapters

To determine the number of file servers the organization needs when migrating to Windows Server 2003, use the scaling factors described in Table 2.21.

**Table 2.21 Consolidating File Servers Running Windows 2000 Server**

| Area of Improvement | Details | Scaling Factor |
|---|---|---|
| Operating system | As shown in Table 2.12, migrating from Windows 2000 Server to Windows Server 2003 provides a 2.2X performance improvement. | 2.2X |
| Hardware upgrade | As shown in Table 2.13, migrating from a Xeon, 400 MHz, to a Xeon, 900 MHz, provides a 2X performance improvement. | 2X |
| Processor scaling | As shown in Table 2.14, migrating from 4P servers to 8P servers provides a 1.3X performance improvement. | 1.3X |
| Total improvements | Multiply the scaling factors (2.2 x 2 x 1.3) | 5.7X |

Based on these figures, the organization can use a single 8P Xeon 900-MHz file server, assuming it has enough memory, disk, and network I/O available. To ensure the availability of the server, the organization should consider using a server cluster instead of a single stand-alone file server. For more information about increasing the availability of file servers, see "Planning File Server Availability" earlier in this chapter.

**Providing Additional Storage**
If your organization uses direct-attached storage for file servers, deploy as many file servers as it takes to meet your storage needs. If each of your file servers provides 500 GB of storage, and you need at least 2 TB of storage, you need a minimum of five file servers. You also need to provide the new file servers with enough processing power and memory to serve your clients. Use the factors described in "Determining RAM and CPU Specifications" earlier in this chapter.

**Providing More Processing Power**
If users experience delays when accessing file servers, there are a number of ways you can improve file server performance:

- Enable the Offline Files feature so that clients cache files locally instead of going to the file server whenever they need a file. You can also use this feature to enable automatic caching for programs. When you enable program caching on a shared folder, and a user runs a program from the shared folder, Windows Server 2003 copies the application .exe and .dll files as they are used and runs the files from the client's local cache, which reduces server load and network traffic. For simple programs that use a single .exe file, this feature enables the program to run correctly when the server is down or the client is not connected to the network. For programs that consist of multiple files (.exe and .dll), all files must be

cached on the local computer so that the program can run while offline. The user can ensure that all files are cached by choosing the **Make Available Offline** option from the **File** menu in Windows Explorer. For more information about the Offline Files feature, see "Offline settings for shared resources" and "Make a file or folder available offline" in Help and Support Center for Windows Server 2003.

- Place copies of heavily used shared folders with primarily read-only files on multiple file servers, and use DFS to provide load sharing. If you also plan to use the Offline Files feature, note that you can use this feature in conjunction with DFS only for clients running Windows XP Professional. For more information about Offline Files, see "Offline Files overview" in Help and Support Center for Windows Server 2003.

- Upgrade the processor, RAM, or disk subsystem in your existing file servers. Use the factors described in "Determining RAM and CPU Specifications" earlier in this chapter to determine the performance improvements you can gain by upgrading existing hardware.

**Providing Same-Site Data Access for Users**

If you have remote sites with expensive WAN links, you can deploy one or more file servers in those sites to provide users with fast, inexpensive LAN access to the file servers. For each site, take into account the storage and processing power needed by users in the site. You also need a backup and recovery procedure in place at each site in case one of the file servers fails.

# Determining the File Server Hardware Life Cycle

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The hardware life cycle is the amount of time that hardware is used before it is replaced with new hardware. Typical hardware life cycles are 12, 24, or 36 months. If your organization frequently adds new file servers to increase capacity or add new functionality, consider using a relatively short replacement cycle for file servers. This recommendation is based on the high rate of change in computers, disk subsystems, and other components. Maintaining a shorter cycle prevents the problem of servers that were purchased at the beginning of the cycle becoming obsolete before the end of the cycle. However, a shorter life cycle affects the manageability of file servers, because you must replace hardware, install the operating system and applications, and migrate data more frequently.

To transparently migrate files from decommissioned servers to new servers without affecting users, implement one or more DFS namespaces. For more information about DFS, see "Designing DFS Namespaces" earlier in this chapter.

For more information about deploying new servers and migrating data, see "Deploying File Servers" later in this chapter.

# Designing a Shadow Copy Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can give users access to previous versions of files by enabling shadow copies, which provide point-in-time copies of files stored on file servers running Windows Server 2003. By enabling shadow copies, you can reduce the administrative burden of restoring previously backed up files for users who accidentally delete or overwrite important files. Shadow copies work for both open and closed files; therefore, shadow copies can be taken even when files are in use.

Shadow copies work by making a block-level copy of any changes that have occurred to files since the last shadow copy. Only the changes are copied, not the entire file. As a result, previous versions of files do not usually take up as much disk space as the current file, although the amount of disk space used for changes can vary depending on the application that changed the file. For example, some applications rewrite the entire file when a change is made, whereas other applications append changes to the existing file. If the entire file is rewritten to disk, the shadow copy contains the entire file. Therefore, consider the type of applications in your organization, as well as the frequency and number of updates, when you determine how much disk space to allocate for shadow copies.

**Important**

- Shadow copies do not eliminate the need to perform regular backups, nor do shadow copies protect you from media failure. In addition, shadow copies are not permanent. As new shadow copies are taken, old shadow copies are purged when the size of all shadow copies reaches a configurable maximum or when the number of shadow copies reaches 64, whichever is sooner. As a result, shadow copies might not be present for as long as users expect them to be. Be sure to consider user needs and expectations when you configure shadow copies.

Shadow copies are designed for volumes that store user data, such as home directories and My Documents folders that are redirected by using Group Policy, or other shared folders where users store data. Shadow copies work with compressed or encrypted files, and they retain whatever permissions were set on the files when the shadow copies were taken. For example, if a user is denied permission to read a file, that user would not be able to restore a previous version of the file, nor would the user be able to read the file after it has been restored.

Although shadow copies are taken for an entire volume, users must use shared folders to access shadow copies. Administrators on the local server must also specify the *servername**sharename* path to access shadow copies. If you or your users want to access a previous version of a file that does not reside in a shared folder, you must first share the folder.

**Collecting Information to Design a Shadow Copy Strategy**

When you design a shadow copy strategy, review the following topics. For an Excel spreadsheet to assist you in documenting your shadow copy design decisions, see "Shadow Copy and Disk Quota Configuration Worksheet" (Sdcfsv_7.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**Shadow copy support in client operating systems**

Shadow copies can be accessed by computers running Windows Server 2003 and by computers running Windows XP Professional on which you have installed the Previous Versions Client pack (Twcli32.msi). This file is located in the Windows Server 2003 operating system in *windir*\system32\clients\twclient. You can install this file manually on clients or deploy the file by using the software distribution component of Group Policy. For more information about software distribution, see "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit.

To access shadow copies from previous versions of Windows, including Windows 2000 and Windows XP Professional, you can download and install the Shadow Copy Client. For more information and to download the Shadow Copy Client, see the Shadow Copy Client Download link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Note**

- The Previous Versions Client and the Shadow Copy Client provide the same functionality, but the Shadow Copy Client can be installed on multiple operating systems, such as Windows 2000 and Windows XP Professional, whereas the Previous Versions Client can only be installed on Windows XP Professional.

If you have not yet deployed these operating systems or client packs on your clients, you can deploy a single computer (or as many as necessary) from which users can restore previous versions of files. You can also distribute the client pack on a case-

by-case basis to users who request that files be restored.

**Shadow copy support in server operating systems**
Shadow copies are available only on file servers running Windows Server 2003.

**Shadow copy support on server clusters**
There are a number of important considerations for managing shadow copies on cluster storage. For more information about managing shadow copies on server clusters, see "Using Shadow Copies of Shared Folders in a server cluster" in Help and Support Center for Windows Server 2003.

**File system requirements**
Shadow copies are available only on NTFS volumes.

**Recommended scenarios for using shadow copies**
Shadow copies work best when the server stores user files such as documents, spreadsheets, and graphics files. Do not use shadow copies to provide access to previous versions of application or e-mail databases.

**Amount of volume space to allocate to shadow copies**
When you enable shadow copies on a volume, you can specify the maximum amount of volume space to be used for the shadow copies. The default limit is 10 percent of the source volume (the volume being copied). Increase the limit for volumes where users frequently change files. Also, setting the limit too small causes the oldest shadow copies to be deleted frequently, which defeats the purpose of shadow copies and which will likely frustrate users. In fact, if the amount of changes is greater than the amount of space allocated to storing shadow copies, no shadow copy is created. Therefore, carefully consider the amount of disk space that you want to set aside for shadow copies, while keeping in mind user expectations for how many versions they want to be available. Your users might expect only a single shadow copy to be available, or they might expect three days' or three weeks' worth of shadow copies. The more shadow copies the users expect, the more storage you need to allocate for storing them.

Setting the limit too small can adversely affect other programs, such as the Backup program in Windows Server 2003 and other backup programs that support the Volume Shadow Copy service. During the backup process, these programs create temporary shadow copies for consistent backups. The temporary shadow copies count towards the volume space limit you specify for shadow copies. If the available volume space remaining in the limit is too small, the Volume Shadow Copy service can delete existing shadow copies to free up space for the temporary shadow copy. If there are no existing shadow copies, and the volume space limit cannot accommodate the temporary shadow copy, the backup might fail. For more information about the Volume Shadow Copy service, see the Storage Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

### Important

- Regardless of the volume space that you allocate for shadow copies, you can have a maximum of 64 shadow copies for any volume. When the 65th shadow copy is taken, the oldest shadow copy is purged.

**Frequency at which Windows Server 2003 creates shadow copies**
By default, Windows Server 2003 creates shadow copies at 7:00 A.M. and at 12:00 noon Monday through Friday. However, you can change the schedule to better accommodate users. Keep in mind that the more shadow copies you create, the more disk space the shadow copies can consume, especially if files change frequently. When you determine the schedule, avoid scheduling shadow copies to occur more than once per hour.

**Storing shadow copies on separate disks**
You can dedicate a volume on separate disks for storing the shadow copies of another volume on the same file server. For example, if user files are stored on H:\, you might use another volume, such as S:\, to store the shadow copies. Using a separate volume on separate disks provides better performance, and it is recommended for heavily used file servers. If you plan to use a separate volume for the storage area (where the shadow copies are stored), be sure to change the maximum size to **No Limit** to reflect the space available on the storage area volume instead of the source volume (where the user files are stored).

### Important

- If you plan to store the shadow copies on the same volume as the user files, note that a burst of disk I/O can cause all shadow copies to be deleted. If you cannot tolerate the sudden deletion of shadow copies, use a volume that will not be shadow copied, preferably on separate disks, for storing shadow copies.

**File servers containing mounted drives**
A mounted drive is a local volume attached to an empty folder (called a *mount point*) on an NTFS volume. If you enable shadow copies on a volume that contains mounted drives, the mounted drives are not included when shadow copies are taken.

In addition, if you share a mounted drive and enable shadow copies on it, users cannot access the shadow copies if they traverse from the host volume (where the mount point is stored) to the mounted drive.

For example, assume you have the folder E:\Data\Users, and the Users folder is a mount point for F:\. You enable shadow copies on both E:\ and F:\, you share E:\Data as \\Server1\Data, and you share E:\Data\Users as \\Server1\Users. In this example, users can access previous versions of \\Server1\Data and \\Server1\Users but not \\Server1\Data\Users.

**Defragmenting the volume where shadow copied files are stored**
If you plan to defragment volumes on which shadow copies are enabled, it is recommended that you use a cluster (or allocation unit) size of 16 KB or larger. If you do not, the number of changes caused by the defragmentation process can cause shadow copies to be deleted faster than expected. Note, however, that NTFS compression is supported only if the cluster size is 4 KB or smaller.

### Note

- To check the cluster size of a volume, use the **fsutil fsinfo ntfsinfo** command. If the volume contains data and you want to change the cluster size, you must back up the data on the volume, reformat it using the new cluster size, and then restore the data.

**Converting basic disks to dynamic disks**
If the volumes that contain the original files and the shadow copy storage area are on separate basic disks and you want to convert both of the disks to dynamic disks, you must follow these directions:

- **If the shadow copy storage area is not on the boot volume.** First dismount and take offline the volume that contains the original files. To do this, use Mountvol.exe with the /p parameter. Next, convert the disk that contains the storage area volume to a dynamic disk. After the conversion, you have 20 minutes to mount the volume that contains the original files and bring it online by using Mountvol.exe or the Disk Management snap-in; if you do not, you will lose the existing shadow copies. After you bring the volume that contains the original files back online, you can convert that disk to a dynamic disk.

- **If the shadow copy storage area is on the boot volume.** You can convert the disk that contains the storage area volume to a dynamic disk without having to dismount the volume that contains the original files. To complete the conversion, you must restart the computer twice. Next, convert the disk that contains the original files to a dynamic disk.

### Important

- If the disk that contains the original files is converted to a dynamic disk first, the shadow copies are deleted when you convert the disk that contains the storage area volume to a dynamic disk.

**Using DFS to provide access to volumes that contain shadow copies**
You can create DFS link targets on volumes that contain shadow copies, and users can retrieve previous versions of files from the DFS link target, just as they can from a regular shared folder. However, if you use multiple link targets for a DFS link, each of those link targets can reside on a different volume with its own shadow copies. As a result, the previous versions of files can vary, depending on which link target the user last accesses to change the file.

**Backing up shadow copies**
Windows Server 2003 does not support backing up shadow copies. When you back up a volume, shadow copies are not backed up.

**Using scripting to create shadow copies**
To create shadow copies by using scripts, use Vssadmin.exe to create the shadow copies and, optionally, use Schtasks.exe to schedule the creation of shadow copies. For more information about using these command-line tools, see "Vssadmin" and "Schtasks" in Help and Support Center for Windows Server 2003.

For more information about shadow copies, see "Shadow Copies for Shared Folders" in Help and Support Center for Windows Server 2003.

**Example: An Organization Designs a Shadow Copy Strategy**
A large organization allows users to redirect their My Documents folders to a central file server. When a user accidentally deletes or overwrites a file and requests that the file be restored, the process requires between 1 day and 3 days and up to three escalations before the backup/restore group can restore a file. The organization determines that it costs approximately $300 for support and escalation costs plus lost productivity while the restoration takes place.

To decrease the cost associated with restoring files and to increase user satisfaction, the organization enables shadow copies on its file servers. During the pilot program, the organization determines that the default settings for the schedule and storage volume allow two weeks of previous versions. When a user contacts the support group to have a file restored, the support group provides a copy of the Previous Versions Client pack to the user. It takes the user approximately five minutes to install the software and another five minutes to restore the file. As a result, a user can restore a file in 10 minutes instead of 1 day to 3 days, and user satisfaction is greatly increased.

# Designing a Disk Quota Strategy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To prevent file servers from filling to capacity without warning, use disk quotas to track and control disk space usage on file servers. Windows Server 2003 provides disk quota functionality that tracks quotas on a per-user, per-volume basis. After you enable the warning level and limit, they apply to all users who own files stored on the volume. Any user who creates a new file is automatically assigned the current warning level and limit.

Windows Server 2003 quotas are well suited for volumes that store user data, such as redirected My Documents folders or other per-user folders. For example, an organization uses Group Policy to redirect its employees' My Documents and My Pictures folders to a file server volume that stores only user data, not group data. After evaluating the storage needs of users in the organization, the administrator chooses to set storage limits by enforcing a 300-MB disk quota for each user. The administrator also sets a 270-MB warning level. If a user needs additional storage, the user must get approval from his or her manager before the administrator can increase the quota for the individual user.

Windows Server 2003 disk quotas are not designed for project-oriented scenarios in which a shared folder is given a quota that is shared among a specified set of users. A wide range of Windows Server 2003–compatible, third-party products provide group quotas, per-folder quotas, and threshold event features that this scenario requires.

To design a quota strategy for volumes that store user data, review the following topics. For an Excel spreadsheet to assist you in documenting your disk quota design decisions, see "Shadow Copy and Disk Quota Configuration Worksheet" (Sdcfsv_7.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**File system requirements**
Disk quotas are available only on NTFS volumes.

**Tracking quotas vs. enforcing quotas**
Tracking quotas is useful if you want to track use of disk space without denying users access to a volume. Enforcing quotas denies disk space to users who exceed their limit. Users can determine how much disk space they have available on the volume by viewing the properties of a mapped shared folder in Windows Explorer.

**Determining how much disk space to provide to each user**
You can change quotas on a per-user basis if some users require additional disk space. For example, the current limit might be 200 MB, but you can specify which users have a 500-MB limit. You cannot assign quotas on a per-group basis.

When setting quotas, make sure that you have enough disk space to accommodate existing users and that you have enough extra free space to accommodate future growth. Each file that is stored on the volume can use up to 64 KB of NTFS metadata that is not charged to a user's quota. To avoid running out of disk space, you must leave sufficient disk space to accommodate this metadata.

**Using disk quotas on server clusters**
Use only domain-level accounts, because local accounts cannot be resolved when the disk fails over. For more information about configuring disk quotas on shared cluster disks, see article Q278365, "HOW TO: Configure Disk Quotas for a Shared Disk in a Cluster" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Logging quota events**
Windows Server 2003 provides quota reporting by logging events in the event log when users exceed their warning level or quota limit. You can use Windows Management Instrumentation (WMI) for customized reporting by using the WMI classes Win32_DiskQuota, Win32_QuotaSetting, and Win32_VolumeQuotaSetting. For more information about WMI, see the Microsoft Windows Management Instrumentation (WMI) SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Using disk quotas on DFS link targets**
DFS provides a simple way to create custom quotas on different volumes that appear to users as different folders on the same volume. If you are using DFS to provide multiple link targets, note that each of these link targets can reside on a volume with different quota settings. To keep the user experience consistent, use the same quota settings for each volume.

**Using disk quotas on shadow-copied volumes**
You can enable shadow copies and disk quotas on the same volume, and shadow copies do not count against a user's disk

quota.

**Note**

- For more information about file systems and Windows Server 2003 disk quotas, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Planning File Server Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Planning for file server security is vital for protecting your organization's sensitive or business-critical files. When planning for file server security, you need to protect not only the data but the physical server as well. In addition, if you plan to implement availability strategies, such as DFS or clustering, you need to take additional steps to secure the resources associated with these features.

Figure 2.18 describes the process for planning file server security.

**Figure 2.18 Planning File Server Security**



# Ensuring the Physical Security of Each File Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For file servers that must maintain high availability, restrict physical access to only designated individuals. In addition, consider to what extent you need to restrict physical access to network hardware. The details of how you implement physical security depend on your physical facilities and your organization's structure and policies. You should also implement methods to restrict access to backup media and any instruction sheets that you create, such as the instructions that go in the recovery manual for a file server. Allowing unauthorized people to study documentation or configuration manuals means that they can quickly cause harm to the system if they are able to obtain access.

Even if the physical server is in a secure room, the file server might still be accessible through remote administration tools. Therefore, implement methods for restricting access to remote administration of file servers, and ensure that remote administration tools do not weaken your organization's security model. For example, remote administration tools do not always use strong authentication protocols, such as Kerberos V5, to authenticate users across the network. You might be able to implement weaker protocols, such as NTLM, depending on the remote management tool you use and the operating system that is running on the host you are administering. In addition, certain remote administration tools might transmit unencrypted data (plaintext) across the network. This makes your data vulnerable to network sniffers.

For more information about security and remote administration, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Planning Baseline Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows 2000 Server Baseline Security Checklist provides instructions for configuring a baseline level of security on servers running Windows Server 2003. The checklist contains tasks such as using NTFS, using strong passwords, disabling unnecessary services, disabling unnecessary accounts, and more. For more information about the baseline security checklist, see the Windows 2000 Server Baseline Security Checklist link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Run the Microsoft Baseline Security Analyzer (Mbsa.exe for the graphical user interface version; Mbsacli.exe for the command-line version). For more information about the Microsoft Baseline Security Analyzer, see the MBSA link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

To further enhance security, review the Microsoft Windows Update Web site regularly for patches that fix vulnerabilities and provide security enhancements. For more information about Windows Update, see the Windows Update link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Planning Virus Protection for File Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To protect file servers from viruses, plan to take the following precautions:

- Use Windows Server 2003–compatible antivirus software, and regularly update virus signature files.
- Back up files regularly so that damage is minimized if a virus attack does occur.
- With clustered file servers, you must use an antivirus program that is cluster aware. If it is not, failover might not occur correctly. If your organization does not have cluster-aware antivirus software, you can install antivirus software on a nonclustered server and use that server to periodically scan the drives that are mapped to the clustered file server. For more information about running antivirus software on server clusters, see article Q250355, "Antivirus Software May Cause Problems with Cluster Services," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- For FRS-replicated content, you must use antivirus programs that are FRS compatible and that do not change the security descriptor of files. For more information about FRS compatible antivirus programs, see article Q815263, "Antivirus, Backup and Disk Optimization Programs That Are Compatible with the File Replication Service." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Planning Access to Shared Folders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you plan access to shared folders, determine the type of permissions to use, who needs access to the folders, and the level of access that users require. You can also disable administrative shares and hide shared folders.

**Determining the Type of Permissions to Use**
Permissions define the type of access granted to a user or group for a file or folder. Windows Server 2003 offers two types of permissions:

- NTFS permissions restrict local and remote access to files and folders on NTFS volumes. When you create a new folder, it inherits permissions from its parent folder. When you create a file in a folder, the file inherits permissions from the parent folder.

- Share permissionsrestrict remote access to shared folders, but share permissions do not restrict access to users who log on to the server locally. Share permissions are available on both FAT and NTFS volumes.

To simplify administering and troubleshooting permissions, use NTFS permissions to control user and group access to file system resources.

Although NTFS is recommended as the primary method for securing folders, you must keep in mind that default share permissions are assigned when you share a folder, and the default share permissions have changed for Windows Server 2003. Windows 2000 and Windows XP grant the Everyone group the Full Control share permission, but Windows Server 2003 grants the Everyone group the Read share permission. This change increases the security of shared folders and helps prevent the spread of viruses.

Because the more restrictive permissions always apply when you use a combination of share and NTFS permissions, you might need to change the default share permissions if you want users to be able to add or change files in the folder. If you do not change the default share permissions, users will have the Read share permission even if you grant users NTFS permissions such as Write or Modify.

If you use a clustered file server, you must create share permissions by using the Cluster Administrator snap-in, not Windows Explorer. In addition, if you plan to use the Share Subdirectories option, you must use NTFS permissions to secure the subdirectories. For more information about these options, see "Planning Cluster Security" later in this chapter.

**Determining Who Needs Access to the Folders**
To increase security and prevent users from browsing through shared folders that are not relevant to their jobs, assign permissions only to groups that require access to the shared folders.

To reduce administrative overhead when assigning permissions, do the following:

- Assign permissions to groups rather than to users.

- Place users in global groups or universal groups, nest these groups within domain local groups, and then assign the domain local groups permissions to the folder.

You do not need to deny permissions for specific groups. When permission to perform an operation is not explicitly granted, it is implicitly denied. For example, if you allow the Marketing group, and only the Marketing group, permission to access a shared folder, users who are not members of the Marketing group are implicitly denied access. The operating system does not allow users who are not members of the Marketing group to access the folder.

Deny access to folders only in the following scenarios:

- You want to exclude a subset of a group (for example, an individual user) that has permissions.

- You want to exclude one or more special permissions when you have already granted Full Control to a user or group.

**Note**

- If you plan to redirect your users' My Documents folders, note that each user is granted exclusive access to his or her My Documents folder on the file server. If you need to access a user's My Documents folder, you have two choices: take ownership of the folder or follow the instructions provided in article Q288991, "Enabling the Administrator to Have Access to Redirected Folders" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

#### Determining the Level of Access That Users Require

Assign the most restrictive permissions that still allow users to perform required tasks. The following descriptions explain the permissions that are associated with folders on NTFS volumes.

**Write** Users can copy or paste new files and subfolders in the folder and change folder attributes. However, users cannot open or browse the folder unless you grant the Read permission. Assigning Write permission is useful for folders where users can file confidential reports, such as timesheets, that only the manager or shared folder administrator can read.

**Read** Users can see the names of files and subfolders in a folder and view folder attributes, ownership, and permissions. Users can open and view files, but they cannot change files or add new files. Assign the Read permission if users need only to read information in a folder and they do not need to delete, create, or change files.

**List Folder Contents** Users can see the names of files and subfolders in the folder. However, users cannot open files to view their contents.

**Read & Execute** Users have the same rights as those assigned through the Read permission, as well as the ability to traverse folders. Traverse folders rights allow a user to reach files and folders located in subdirectories, even if the user does not have permission to access portions of the directory path.

**Modify** Users can delete the folder and perform the actions permitted by the Write and Read & Execute permissions. Because Modify gives users the ability to delete the folder, use Modify permission only for administrators or for the group or department owner of the folder.

**Full Control** Users can change permissions, take ownership, delete subfolders and files, and perform the actions granted by all other permissions. Because Full Control gives users the ability to delete the folder, use Full Control permission only for administrators or for the group or department owner of the folder.

For more information about permissions and file servers, see "Permissions on a file server" in Help and Support Center for Windows Server 2003.

#### Determining Whether to Disable Administrative Shares

Windows Server 2003 creates shared folders, known as *administrative shares*, by default when you start a server or when you stop and then start the Server service. These folders are shared for administrative purposes, and they allow users and applications with the appropriate administrative rights to gain access to the system remotely. For example, some backup software applications use these shares to remotely connect to systems to back up data.

Administrative shares have default share permissions that restrict access to members of only a few security groups. Each share name is appended with a dollar sign ($), which hides the share from users who browse the server. One type of administrative share is the root folder of every volume (C$, D$, and so on).

You can disable these administrative shares temporarily or permanently. For more information about disabling administrative shares and an overview of remote administration, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

#### Determining Whether to Hide Shared Folders

You can hide a shared folder by appending a dollar sign ($) to the shared folder name. Hiding shared folders is useful when you want to make a shared folder available over the network while keeping it hidden from people browsing on the network.

Hiding shared folders does not necessarily make them more secure, because anyone who knows the name of the server and the shared folder can connect to it. Therefore, you must set the necessary NTFS permissions on the shared folder so that access is granted only to the appropriate groups.

# Planning Encrypted File Storage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows 2000, Windows XP, and Windows Server 2003 support storing files that are encrypted using EFS. However, remote decryption is a potential security risk, because files are decrypted before transmission on the local server, and they are transmitted unencrypted over the network in plaintext. Therefore, before you allow encrypted files to be stored on file servers, decide whether the risk associated with transmitting unencrypted files over the network is acceptable.

You can greatly reduce or eliminate this risk by enabling Internet Protocol security (IPSec) policies, which encrypts data that is transmitted between servers, or by using Web Distributed Authoring and Versioning (WebDAV) folders. WebDAV folders have many advantages compared to shared folders, so you should use them whenever possible for remote storage of encrypted files. WebDAV folders require less administrative effort and provide greater security than shared folders. WebDAV folders can also securely store and deliver files that are encrypted with EFS over the Internet by means of Hypertext Transfer Protocol (HTTP).

Before users can encrypt files that reside on a remote file server, you must designate the file server as trusted for delegation. Doing so allows all users with files on that server to encrypt their files. For more information about enabling encryption on a file server, see "Enable a remote server for file encryption" in Help and Support Center for Windows Server 2003. Note that when encrypting files on a WebDAV server, the server does not need to be trusted for delegation.

**Important**

- To enable EFS on a clustered file server, you must perform a number of steps to configure the environment correctly. For more information about enabling EFS on server clusters, see "Create a cluster-managed encrypted file share" in Help and Support Center for Windows Server 2003.

If you allow users to store encrypted files on file servers, review the following issues:

- Users can encrypt files on remote NTFS volumes only when both the user's computer and the file server are members of the same Windows Server 2003 forest. (This restriction does not apply to WebDAV folders.)
- Users must have Write or Modify permissions to encrypt or decrypt a file.
- Users cannot encrypt files that are compressed. If users encrypt a compressed file or folder, the file or folder is uncompressed.

For more information about EFS and using WebDAV folders to store encrypted files, see "Encrypting and decrypting data" and "Internet Information Services (IIS) 6.0 overview" in Help and Support Center for Windows Server 2003. For more information about configuring IPSec, see the *Networking Guide* of the *Windows Server 2003 Resource Kit* (or see the Networking Guide on the Web at http://www.microsoft.com/reskit).

# Planning DFS and FRS Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When planning to secure DFS namespaces and content replicated by FRS, follow these guidelines:

- Use NTFS permissions to secure DFS targets. If you are using FRS to replicate DFS link target information, any permission changes you make on one member of the replica set replicate to other members. If you are not using FRS for automatic replication, you must set the permissions on targets and manually propagate any changes that occur.

- When setting NTFS permissions, always use the path of the physical folder (*servername**sharename*) instead of navigating through the DFS namespace to set permissions. This is especially important when you have multiple link targets for a given link. Setting permissions on a folder by using its DFS path can cause the folder to inherit permissions from its parent folder in the namespace. In addition, if there are multiple link targets, only one of them gets its permissions updated when you use the DFS path.

- If you plan to use share permissions, note that FRS does not replicate share permissions; therefore, you must plan to implement identical share permissions for each shared folder in a replica set. If you do not, users might have inconsistent access to shared folders across the network.

- To prevent the spread of viruses in read-only FRS-replicated content, give the appropriate groups the NTFS Read & Execute permission, create a group for administrators who update content, and assign that group the NTFS Modify permission. Do not grant permissions to the Everyone group. For additional recommendations, see "Permissions on a file server" in Help and Support Center for Windows Server 2003.

- For FRS-replicated content, you must use antivirus programs that are FRS compatible and that do not change the security descriptor of files. For more information about FRS compatible antivirus programs, see article Q815263, "Antivirus, Backup and Disk Optimization Programs That Are Compatible with the File Replication Service." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- You must have permissions on the DFS configuration object in Active Directory to add and delete roots to a domain-based DFS namespace.

- You can create DFS link targets that point to shared folders containing data that is encrypted by using EFS. However, you cannot use FRS to replicate those files among multiple link targets.

- Do not enable the RestrictAnonymous registry value on DFS root servers. Doing so restricts anonymous access and causes DFS referral failures. This registry value is also part of the HiSecWeb security template, which is designed to help secure Internet Information Services (IIS) at the operating system level. For more information about the RestrictAnonymous registry value, see article Q246261, "How to Use the RestrictAnonymous Registry Value in Windows 2000." For more information about the HiSecWeb template, see article Q316347, "IIS 5: HiSecWeb Potential Risks and the IIS Lockdown Tool," in the Microsoft Knowledge Base. To find these articles, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Planning Cluster Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When planning to secure clustered file servers, follow these guidelines:

- After you create a folder by using Windows Explorer, verify that the Cluster service account has the Read permission on the folder so that you can share the folder properly by using Cluster Administrator. (Do not share the folder by using Windows Explorer.)

- Use Cluster Administrator to set share permissions. If you change file share permissions using Windows Explorer or My Computer, instead of using **Permissions** on the **Parameters** tab in Cluster Administrator, the permissions are lost when the resource is taken offline.

**Note**

- When you set file share permissions by using Cluster Administrator, the default permissions give the Everyone group the Read permission. When you set file share permissions by using Cluster.exe, the Everyone group has the Full Control permission.

    - To secure File Share resources on the local server, use Windows Explorer to assign NTFS permissions on the physical folder, because share permissions apply only when users connect to the clustered file server across the network.

    - Do not assign NTFS permissions to local groups on clustered file servers. These permissions will have no meaning when the clustered disk resource is moved to another server. Therefore, always assign permissions to a domain local group.

    - By default, access to cluster file shares is disabled to anonymous users. To allow anonymous access to specific file shares, you can either enable Kerberos V5 authentication on the Network Name resource that is associated with the file share or you can change the local security policy setting. For more information about configuring these Kerberos properties, see "Enable Kerberos authentication for virtual servers" in Help and Support Center for Windows Server 2003.

    - When you create File Share resources by using the Share Subdirectories option, the subdirectories inherit the permissions of the parent. If you are using the subdirectories as user folders, and you want to allow only the user and administrator to access the folder, set NTFS permissions on each subfolder.

    - To enable EFS on a clustered file server, you must perform some steps to configure the environment correctly. These steps are described in "Create a cluster-managed encrypted file share" in Help and Support Center for Windows Server 2003.

For more information about server cluster security, see "Best practices for securing server clusters" in Help and Support Center for Windows Server 2003.

# Deploying File Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying file servers involves installing Windows Server 2003 on the physical servers, adding the servers to the network, and then implementing the design decisions you made throughout this chapter. Organizations that plan to deploy clustered file servers must take the additional step of configuring the clusters; all other organizations can begin by choosing the installation method. Figure 2.19 describes the file server deployment process.

**Figure 2.19 File Server Deployment Process**



# Installing Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you install Windows Server 2003, you need to decide whether to perform clean installations or upgrades. You also need to decide if you want to perform an automated installation.

**Choosing Between Clean Installations and Upgrades**
If your servers are new and do not contain operating systems, you must perform clean installations. If you have existing servers, you can perform clean installations or upgrades. Consider the following points when choosing between a clean installation and an upgrade.

Choosing a new installation:

- If you reformat your hard disk and then perform a new installation, the efficiency of your disk might improve (compared to not reformatting it). Reformatting also gives you the opportunity to modify the size or number of disk partitions to make them match your requirements more closely.

- If you want to practice careful configuration management (for example, for a server where high availability is important), you might want to perform a new installation on that server instead of an upgrade. This is especially true on servers on which you have upgraded the operating system several times in the past.

Choosing an upgrade:

- With an upgrade, configuration is simpler, and you retain your existing users, settings, groups, rights, and permissions.

- With an upgrade, you do not need to reinstall files and applications. As with any major changes to the hard disk, however, it is recommended that you back up the disk before beginning an upgrade.

**Preparing for a Clean Installation on Existing File Servers**
Review the following issues before performing a clean installation on existing file servers:

- If you are performing a clean installation on a file server that contains a domain-based DFS root, remove the server as a DFS root before you begin the installation.

- If you are performing a clean installation on a member of an FRS replica set, use the Distributed File System snap-in (available in Windows Server 2003 or in the Windows Server 2003 Administration Tools Pack) to remove all inbound and outbound FRS connections to the server before beginning the clean installation.

**Upgrading Existing File Servers**
If you plan to upgrade file servers to Windows Server 2003, review the following issues:

- **Service pack requirements for servers running Windows NT 4.0.** Before you can upgrade a server that is running Windows NT 4.0 to Windows Server 2003, you must first install Windows NT 4.0 Service Pack 5 or a later version.

- **Upgrading servers that contain multidisk volumes.** If a server contains multidisk volumes, review the following issues:

    - If you are upgrading from Windows NT Server 4.0 to Windows Server 2003, verify that your backup software and hardware are compatible with both Windows NT Server 4.0 and Windows Server 2003. Next, back up and then delete all multidisk volumes (volume sets, mirror sets, stripe sets, and stripe sets with parity) before you upgrade, because Windows Server 2003 cannot access these volumes. Be sure to verify that your backup was successful before deleting the volumes. After you finish upgrading to Windows Server 2003, create new dynamic volumes, and then restore the data.

    - If you are upgrading from Windows NT Server 4.0 to Windows Server 2003, and the paging file resides on a

multidisk volume, you must use System in Control Panel to move the paging file to a primary partition or logical drive before beginning Setup.

- If you are upgrading from Windows 2000 Server to Windows Server 2003, you must use Disk Management to convert all basic disks that contain multidisk volumes to dynamic disks before beginning Setup. If you do not do this, Setup does not continue.

- **Upgrading servers that contain multiple DFS roots.** If you are running a prerelease version of Windows Server 2003, Standard Edition, and the server hosts multiple DFS roots, only one of those roots will be available after the upgrade. For more information about this and other DFS issues during upgrades, see "Deploying DFS" later in this chapter.

- **Upgrading servers that are part of FRS replica sets.** To keep the data in a replica set available during the upgrade, stagger the upgrades of replica members.

**Choosing Automated Installations**

Many large organizations use automated installations to install Windows Server 2003. Bulk installations are faster, easier, less expensive, more consistent, and more efficient than having users or IT professionals sit at servers and enter data manually. Windows Server 2003 provides the following tools and methods that you can use to design and deploy very simple or very sophisticated automated installations:

- Remote Installation Services (RIS)

- System Preparation Tool (Sysprep)

- Unattended Installation

These automated installation methods address a variety of specific issues, and each method has inherent advantages and disadvantages, depending on the environment in which you use these methods. To determine the best methods to use in the context of your specific environment, see "Choosing an Automated Installation Method" in *Automating and Customizing Installations* of this kit.

# Deploying Clustered File Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The cluster installation method you choose depends on whether you have existing clustered file servers. If you do have existing clustered file servers, you have several options when upgrading server clusters:

- Perform a new installation of Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, and configure the cluster at the same time.
- Upgrade a cluster that is running Windows NT Server 4.0, Enterprise Edition.
- Upgrade a cluster that is running Windows 2000, possibly through a rolling upgrade.

There are two major advantages to a rolling upgrade. First, there is minimal interruption of service to clients. (However, server response time might decrease during the phases in which fewer nodes handle the work of the entire cluster.) Second, you do not have to recreate your cluster configuration. The configuration remains intact during the upgrade process.

**Caution**

- Do not make DFS configuration changes (for example, adding new links, new link targets, and so on) while operating a mixed-version cluster, because all DFS changes are lost upon failover.

You cannot perform a rolling upgrade directly from Windows NT Server 4.0, Enterprise Edition to Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. Instead you have two options:

- You can maintain cluster availability by performing an upgrade to Windows 2000 first (as specified in the Windows 2000 documentation), and then upgrade to Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.
- You can perform a nonrolling upgrade directly from Windows NT 4.0 to Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. Note that a nonrolling upgrade does not allow you to maintain cluster availability.

For comprehensive information about choosing the cluster installation method and performing the actual installation, see "Designing and Deploying Server Clusters" in this book and see "Installing and upgrading on cluster nodes" in Help and Support Center for Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.

Regardless of the installation type, evaluate your cluster hardware for compatibility with Windows Server 2003. For more information about choosing cluster hardware, see "Designing and Deploying Server Clusters" in this book.

After you have installed Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition and configured the server cluster, use the following sections for information about creating mounted drives and File Share resources on server clusters.

**Creating Mounted Drives on Server Clusters**

If you assign drive letters to the cluster storage devices, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, limit the total number of such devices to 23. Mounted drives are not subject to this limit; you can use mounted drives to access more than 23 cluster storage devices in your server cluster.

When using NTFS mounted drives with server clusters, follow these recommendations:

- Make sure that you create unique mounted drives so that they do not conflict with existing local drives on any node in the cluster.
- Do not create mounted drives between disks on the cluster storage device (cluster disks) and local disks.
- Do not create mounted drives from the cluster disk that contains the quorum resource (the quorum disk). You can,

however, create a mounted drive from the quorum disk to a clustered disk.

- Mounted drives from one cluster disk to another must be in the same cluster resource group, and they must be dependent on the root disk.

- Use Event Viewer to check the system log for any Cluster service errors or warnings indicating mount point failures. These errors are listed as ClusSvc in the **Source** column and as Physical Disk Resource in the **Category** column.

For more information about creating mounted drives, see "Create a mounted drive" in Help and Support Center for Windows Server 2003.

**Creating File Share Resources**

After you choose the cluster installation method and create the server cluster, migrate existing data on existing file servers, if necessary, and then create one or more File Share resources by using the Cluster Administrator snap-in or by creating scripts. For more information about creating a cluster-managed file share, see "Create a cluster-managed file share" in Help and Support Center for Windows Server 2003.

If you prefer to use scripts to create File Share resources, see article Q284838, "How to Create a Server Cluster File Share with Cluster.exe" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. For more information, see "Managing a Server Cluster from the Command Line" in Help and Support Center for Windows Server 2003.

After you create the File Share resources, you can enable shadow copies on cluster storage. For more information about enabling shadow copies on server clusters, see "Enable Shadow Copies for Shared Folders in a cluster" in Help and Support Center for Windows Server 2003.

# Migrating File Server Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you plan to consolidate file servers or replace outdated servers with new servers, migrate existing applications and data to the new servers. To ensure the success of the migration, create a detailed and well-tested migration plan that does the following:

- Minimizes the amount of time that data is unavailable to users
- Minimizes the impact on network bandwidth
- Minimizes the cost of the migration
- Ensures that no data loss occurs during the migration

The following sections describe the tasks for creating a migration plan

# Identifying Data to Migrate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Large organizations rarely have time to take all production data offline before migrating it. Instead, they typically move data gradually, migrating different classes of data in stages. A migration plan should identify which data to move, where to move it, and the order in which it should be moved. File server data can usually be classified in the following classes:

- Business-critical data
- Application data
- Personal data (such as home directories)
- Profiles or other configuration data
- Departmental or group data
- Other data types, such as special projects or temporary data

**Important**

- Be sure to review any legal or copyright issues that might arise when you migrate data. For example, if users are storing MP3 files on the file server, you might violate copyrights by copying those files. If you are copying applications, make sure that doing so does not violate any licensing agreements.

When creating your migration plan, determine whether you want to consolidate different classes of data at different rates, onto different hardware, onto servers with different SLAs, or to different locations. Also, consider the importance of the data when planning your migration. For example, to minimize the impact of any problems during the migration, you might choose to move nonessential data first, followed by business-critical data.

In addition to identifying which data to move, you can also identify data that you do not need to move, such as duplicate, obsolete, or non-business-related files. Eliminating these files from your migration plan decreases the number of files that you need to migrate and increases the amount of available storage on the target server after the migration.

# Identifying Migration Risks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The next step is to assess the risks associated with data migration. The best way to identify these risks is to set up a test lab with clients and servers, similar to those in your production environment, and conduct trial migrations. In addition to testing your actual migration method, test any applications that are installed on client computers to determine if the applications are affected by the migration. For example, applications that depend on components stored on a particular file server might not work correctly if the file server or share is renamed after the migration is complete. If you identify problems, update your migration plan so that you prevent or mitigate those problems.

There are a number of solutions you can implement either before or during the migration to prevent problems from occurring. For example, you can ensure that shortcuts and links on client computers work correctly after the migration by doing one of the following:

- **Implementing DFS.** If you have implemented DFS before the migration, the migration is transparent to users; you just need to update the link target location.

- **Migrating to clustered file servers.** In this case, the migration is transparent if you create virtual servers that use the same name as the previously used file servers. You can also transparently migrate any stand-alone DFS namespaces to clustered file servers. Using the previous server names, create your virtual servers, and then migrate the namespaces by using Dfsutil.exe.

- **Using a third-party migration tool.** Use a tool that can identify and fix broken links, such as OLE links or application-related links in the registry.

No migration plan is complete until you develop and test "back-out" procedures to follow if the migration fails at any stage in the process. Verify that you can restore data to its previous state and location in the event that you need to roll back the migration.

The following sections describe additional risks to consider during the migration.

**NTFS permissions**

When evaluating migration methods, determine whether the methods support migrating NTFS permissions. If you use domain local groups to assign permissions to files and folders, members of those groups will have the same access to the files and folders after they are moved to the new server. However, if you assign permissions to computer local groups, members of those groups cannot access the files after the migration. Either reconfigure permissions to use domain local groups before the migration, or plan to create new computer local groups on the target server and give those groups the appropriate permissions to the files after the migration is complete.

**NTFS compression**

If you use compression on the source servers, determine if you still require compression on the target server. Hardware on the target server is typically more powerful, with greater storage capacity. Therefore, compression might not be necessary, although you do need to account for the additional space required by the uncompressed data. Also, when you copy compressed files from one server to another, the compression attribute is lost unless you enable compression on the target volume.

**Files encrypted by using EFS**

Migrating encrypted files has a number of risks. Only administrators who are EFS recovery agents can copy or move files that are encrypted by other users. When these administrators copy or move encrypted files to a remote shared folder, the files are decrypted locally, transmitted in plaintext, and then re-encrypted on the target server only if the remote computer is trusted for delegation and the target volume uses NTFS. When the files are encrypted again on the target server, they have new file encryption keys that are encrypted by using the administrator's public key, if it is available, or by using a new public key, which EFS generates if the profile is unavailable. As a result, the users who originally encrypt the files are no longer able to access them. Therefore, do not copy or move encrypted files from one server to another; use a backup and restore method instead.

**Important**

- When EFS recovery agents copy encrypted files to a target server that is not trusted for delegation, or to a target volume that uses FAT, the files become plaintext on the target server.

For more information about EFS recovery agents, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services* of this kit.

**Choosing a Migration Method**
Because many migration methods require downtime to move data, migrating data can be challenging in the following situations:

- The data has high uptime requirements as specified by SLAs in your organization, and the migration method might require more downtime than the SLA allows.
- Users expect to be able to access that data throughout their workday or at all times.

When planning your migration, consider the advantages and disadvantages of each migration method that is available to you, such as those described in the following sections. Your organization might use other methods not described here.

**Backup and restore data**
This method involves making a backup of the source server and restoring the data on a target server. Depending on your backup method and hardware, this method might involve moving tapes from one backup device to another (for direct-attached backup hardware) or restoring data from a tape library.

Advantages:

- Large organizations routinely back up and restore data, so you can use existing backups to begin the migration.
- Backing up and restoring data does not impact network bandwidth when it is performed on the local servers.
- You can use this method to migrate encrypted files.
- Disadvantages:
- Both the source server and the target server must support the backup device and its related software.
- If users are accessing files during the backup, you must make arrangements to migrate files that change while the backup occurs.
- Backup programs do not migrate shared folder information to the destination server. As a result, folders that are shared on the source server are not shared when you restore them on the destination server. You must share the destination folders manually or by using scripts, and then use Permcopy.exe, which is available on the *Windows Server 2003 Deployment Kit* companion CD, to migrate any share permissions. For more information about Permcopy.exe, click **Tools** in Help and Support Center for Windows Server 2003, and then click "Windows Resource Kit Tools Help"

**Use third-party migration tools, hardware, or services provided by hardware vendors**
A number of third-party migration tools can automate the entire migration process. These migration tools typically offer features such as bandwidth throttling, scheduling, incremental file migration, and monitoring. Many hardware vendors also provide hardware solutions and services that can assist you in migrating data.

**Copy the data**
This method involves using Robocopy.exe to copy data from one server to another and then using Permcopy.exe to migrate share permissions. Use the **/SEC** parameter in Robocopy.exe to copy NTFS permissions from the source folder to the destination folder. After you finish moving the data, share the folders on the target server (either manually or by using scripts), and then use Permcopy.exe to copy share permissions from the source server to the target server.

Permcopy.exe and Robocopy.exe are available on the *Windows Server 2003 Deployment Kit* companion CD. For more information about Permcopy.exe and Robocopy.exe, click **Tools** in Help and Support Center for Windows Server 2003, and then click "Windows Resource Kit Tools Help."

Advantages:

- These tools are available as part of the *Windows Server 2003 Deployment Kit*.

- This method is simple if you are migrating a small amount of data.
- Robocopy.exe gives you a number of options for migrating data, including the following:
    - Using file names, wildcard characters, paths, or file attributes to include or exclude source files as candidates for migrating.
    - Controlling the number of times that Robocopy.exe retries an operation after encountering a recoverable network error.
    - Scheduling copy jobs to run automatically.
- Disadvantages:
- Files are copied across the network, which can impact network bandwidth.
- This method should not be used to migrate encrypted files.
- If users are accessing files as the files are copied, you must make arrangements to migrate files that change while the copy occurs.
- This method might take a long time to complete if you are migrating large amounts of data.
- This method does not migrate shared folder information to the destination server. As a result, folders that are shared on the source server are not shared when you copy or restore them on the destination server. You must share the destination folders manually or by using scripts, and then use Permcopy.exe to migrate any share permissions.

**Completing the Migration**

After you migrate data, verify that users can access the data on the new servers by completing the following tasks:

- Verify that NTFS and share permissions are migrated correctly.
- If a logon script maps drives to the old file servers, update the scripts so that they point to the new file servers.
- If the migrated folders are DFS link targets, update the link information so that it points to the new file server.
- If the migrated folders are redirected My Documents folders, use Group Policy to specify the new file server.
- If you migrate data to a clustered file server, create the necessary File Share resources.

# Deploying DFS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you upgrade a server containing a DFS root, ensure that the namespace will be available after the upgrade. After Windows Server 2003 is installed, either through a clean installation or upgrade, you can create new DFS namespaces or migrate existing DFS namespaces to servers running Windows Server 2003. If you create a new namespace, you can use the information provided by the design team in the "DFS Configuration Worksheet" (Sdcfsv_1.xls).

**Upgrading Servers That Contain DFS Namespaces**
If a file server contains existing DFS roots, they are converted as follows:

- When you upgrade a server running Windows NT 4.0 to Windows Server 2003, any DFS roots are converted to Windows Server 2003 stand-alone DFS roots.

- When you upgrade a server running Windows 2000 to Windows Server 2003, stand-alone DFS roots and domain-based DFS roots are converted to Windows Server 2003 stand-alone DFS roots and domain-based DFS roots, respectively.

In the following scenarios, DFS roots are not available after upgrade:

- If you upgrade a server that hosts a root on a FAT volume, the namespace is unavailable until you convert the volume to NTFS.

**Upgrading Prerelease Versions of Windows Server 2003, Standard Edition**
If you are running a prerelease version of Windows Server 2003, Standard Edition, and that server hosts multiple DFS roots, only one of those roots will be available after the upgrade. This is because Windows Server 2003, Standard Edition cannot host multiple namespace roots until you apply a hotfix (as described below).

To work around this issue, upgrade to Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, or use the following procedure:

1. Before upgrading to Windows Server 2003, Standard Edition, open a command prompt window and use the **Dfsutil.exe** command with the **/Export** parameter to export all but one of the namespaces to text files.

2. Use the **Dfsutil.exe** command with the **/RemFtRoot** parameter to remove any extra domain-based roots from the server, or the **/RemStdRoot** parameter to remove a standalone namespace root.

| **Note** |
| --- |
| If you cannot remove a namespace root using the **/RemFtRoot** or **/RemStdRoot** parameters, use the **Dfsutil.exe/UnmapFtRoot** or **/Clean** parameters to force the removal of the roots. |

1. At the command prompt, type **netstop dfs & net start dfs**

   The DFS service will be briefly unavailable when the service is stopped and restarted, so perform this procedure during a period of low namespace usage.

2. Upgrade the server to Windows Server 2003, Standard Edition, and then install the hotfix described in Microsoft Knowledge Base article 903651 on the Microsoft Web site at http://support.microsoft.com/default.aspx?scid=kb;en-us;903651. This hotfix enables a server running Windows Server 2003, Standard Edition to host multiple domain-based roots.

3. Use **Dfsutil.exe** with the **/Import** parameter to import each namespace you deleted.

**Delegating the Administration of DFS Namespaces**
Use the procedures below to allow members of the local Administrators group on each root server to create and manage domain-based DFS namespaces. For more information about delegating permission to manage a DFS namespace, see "Designing a DFS Namespace" earlier in this chapter.

The following procedure grants the selected user the ability to create new DFS namespaces as well as administer existing ones.

**To delegate a user to administer DFS**

1. In the Active Directory Users and Computers snap-in, on the **View** menu, click **Advanced Features**.
2. In the console tree, double-click the **System** folder to expand it.
3. Click the **DFS-Configuration** folder.

   Any existing root objects appear in the details pane.
4. Right-click **DFS-Configuration**, and then click **Properties**.
5. On the **Security** tab, click **Add**.
6. Type the name of the user to whom you want to delegate administrative rights, and then click **OK**.
7. Select the user Full Control permission, and then click **OK**.

Use the following procedure to allow a user to have DFS administrative permissions only within a single DFS namespace.

**To grant a user permission to administer only a single DFS namespace**

1. In the Active Directory Users and Computers snap-in, on the **View** menu, click **Advanced Features**.
2. In the console tree, double-click the **System** folder to expand it.
3. Click the **DFS-Configuration** folder.

   Any existing root objects appear in the details pane.
4. Right-click the root object that you want to allow the user to administer, and then click **Properties**.
5. On the **Security** tab, click **Add**.
6. Type the name of the user, and then click **OK**.
7. Verify that the user is granted the Full Control permission, and then click **OK**.

**Creating New DFS Namespaces on Stand-Alone Servers**
Use the Distributed File System snap-in or the command-line tools Dfscmd.exe and Dfsutil.exe to create the root and link targets on stand-alone (nonclustered) file servers. The following topics in Help and Support Center for Windows Server 2003 provide information about using these tools:

- "Checklist: Creating a distributed file system"
- "Dfscmd"

For more information about installing Dfsutil.exe, see Windows Support Tools.

If you plan to enable FRS on link targets in the DFS namespace, see "Deploying FRS" later in this chapter.

**Creating New DFS Namespaces on Clustered Servers**
Use the Cluster Administrator snap-in to create a stand-alone DFS namespace on a clustered file server. For more information about creating a DFS root on server clusters, see "Create a cluster-managed file share" in Help and Support for Windows Server 2003.

**Migrating or Consolidating Existing Namespaces on New Servers**
If you have namespaces on existing file servers that you want to consolidate onto a file server that is running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, or if you want to move a namespace from one server to another, you can use the Dfsutil.exe support tool to export the namespace from the source server to the destination server.

In the following example, an administrator wants to migrate the following namespaces on different servers to a single server running Windows Server 2003, Enterprise Edition.

- \\NT4SVR\Marketing (a stand-alone DFS root on a server running Windows NT Server 4.0)
- \\W2KSVR\Public (a stand-alone DFS root on a server running Windows 2000 Server)

First, the administrator creates the following stand-alone DFS roots on the server running Windows Server 2003, Enterprise Edition:

- \\NETSVR\Marketing
- \\NETSVR\Public

Next, the administrator installs Windows Support Tools from the Windows Server 2003 operating system CD and then uses the Dfsutil.exe tool to run the following commands:

```
          Dfsutil /Root:\\NT4SVR\Marketing /export:Nt4.txt
Dfsutil /Root:\\W2KSVR\Public /export:w2k.txt
```

Finally, the administrator runs the following commands to import the namespaces onto the server running Windows Server 2003, Enterprise Edition:

```
          Dfsutil /Root:\\NETSVR\Marketing /import:Nt4.txt /set
Dfsutil /Root:\\NETSVR\Public /import:w2k.txt /set
```

**Using Dfsutil.exe to Customize the Namespace**
Use the Dfsutil.exe parameters listed in Table 2.22 to customize the namespace. For more information about using Dfsutil.exe, see Windows Support Tools.

**Table 2.22 Dfsutil.exe Parameters Used to Customize the Namespace**

| Customization | Dfsutil.exe Parameter |
|---|---|
| Enable root scalability mode | /RootScalability |
| Add site information to root servers running Windows 2000 Server | /UpdateWin2kStaticSiteTable |
| Remove site information from root servers running Windows 2000 Server | /PurgeWin2kStaticSiteTable |
| Enable restricted same-site target selection | /InSite |
| Enable closest or least expensive target selection | /SiteCosting |

# Deploying FRS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have deployed a domain-based DFS namespace, you are ready to deploy FRS. This section describes the following FRS deployment scenarios:

- Deploying a new replica set
- Adding a new replica member to an existing replica set

This section also provides an overview of the tasks required to deploy each scenario. For detailed procedures for each scenario, see the worksheets specified in the relevant sections. The worksheets also provide information about using Sonar.exe to monitor the status of your FRS deployment. Sonar.exe is available on *Windows Server 2003 Deployment Kit* companion CD. For more information about Sonar.exe, click **Tools** in Help and Support Center for Windows Server 2003, and then click Windows Resource Kit Tools Help.

# Deploying a New Replica Set

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are deploying a new replica set, use the "FRS Configuration Worksheet" (Sdcfsv_2.xls) that was completed by the file services design team to begin the deployment. If the design team did not use this job aid, you will need the following information:

- The names of the link targets to be replicated. These link targets form the replica set.
- The FRS topology (ring, full mesh, hub and spoke, or custom) to use for each replica set. If the topology is custom, you also need the inbound and outbound connections for each replica member.
- The FRS replication schedule for each connection.
- The location and size of the staging directory.
- The connection priorities for each inbound connection.

Use the following information to choose the appropriate FRS deployment scenario:

- See "Deploying an Empty Replica Set" later in this section if you have two or more existing link targets on which you want to enable replication and the link targets do not contain data.
- See "Deploying a Replica Set with Existing Content" later in this section if you have two or more existing link targets that contain data and the link targets are currently in use in your organization.
- See "Adding a New Member to an Existing Replica Set" later in this chapter if you have an existing replica set and you want to add a new member.

**Deploying an Empty Replica Set**

This scenario requires that you have deployed your domain-based DFS namespace but you do not have any files and folders in the link targets that you plan to replicate. This scenario is the most efficient way to deploy FRS, because replica members perform their join operation on an empty tree. When you add files to the tree, FRS builds a single staging file for each file in the tree and uses those files to source all direct outbound partners.

**Important**

- When you deploy a new replica set using these scenarios, it is recommended that you deploy two replica members connected by high-bandwidth links and then prestage additional members, especially if they are located in remote sites with slow-bandwidth connections. For more information about adding members to a replica set, see "Adding a New Member to an Existing Replica Set" later in this chapter.

Before you begin, choose a procedure for either a standard topology (ring, hub and spoke, or full mesh) or a custom topology.

**Deploying a standard FRS topology**

To deploy a ring, hub and spoke, or full mesh topology, perform the following tasks:

1. Disable referrals to the link targets to be replicated.
2. Configure the USN journal.
3. Configure replication for a standard topology, and choose the staging directory location.
4. Adjust the size of the staging directory.

5. Verify that replication is working.
6. Set the initial replication schedule.
7. Configure filters to exclude files or folders from replication (optional).
8. Copy the data into the replica tree, and verify that the data is consistent.
9. Optimize the replication schedule (optional).
10. Enable referrals to the replicated link targets.
11. Configure connection priorities (optional).

For a Word document that provides detailed instructions on completing each of these tasks, "Deploying a Standard FRS Topology for an Empty Replica Set" (Sdcfsv_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Deploying a Standard FRS Topology for an Empty Replica Set" on the Web at http://www.microsoft.com/reskit). After you complete these tasks, you can notify users that the DFS namespace is available.

**Deploying a custom FRS topology**
To deploy a custom FRS topology, perform the following tasks:

1. Disable referrals to the link targets to be replicated.
2. Configure the USN journal.
3. Configure replication for a custom topology, choose the staging directory location, and configure connection priorities.
4. Adjust the size of the staging directory.
5. Verify that replication is working.
6. Set the initial replication schedule (optional).
7. Configure filters to exclude files or folders from replication (optional).
8. Copy the data into the replica tree, and verify that the data is consistent.
9. Enable referrals to the replicated link targets.
10. Optimize the replication schedule (optional).

For a Word document that provides detailed instructions on completing each of these tasks, see "Deploying a Custom FRS Topology for an Empty Replica Set" (Sdcfsv_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Deploying a Custom FRS Topology for an Empty Replica Set" on the Web at http://www.microsoft.com/reskit). After you complete these tasks, you can notify users that the DFS namespace is available.

**Deploying a Replica Set with Existing Content**
In this scenario, you have deployed a DFS namespace with existing content in one or more link targets, but you are not using FRS to keep those link targets synchronized. Users know about and use the namespace, and they might be accessing files from the link targets. This scenario is more complex than deploying FRS in an empty directory tree for a number of reasons:

- Users or applications might have locks on files or folders, which can prevent FRS from originating outgoing changes and receiving incoming changes.
- If files exist on the initial master, FRS builds a unique staging file for each file in the tree for each direct outbound partner. In other words, each partner gets its own set of files, which is more disk I/O and disk space intensive.

- The files in the replica tree of the initial master become authoritative, which means that after replication completes, the replica tree on each member is identical to the replica tree on the initial master. For servers that are not the initial master, FRS moves any files that existed in the replica tree to a folder named "NtFrs_PreExisting___See EventLog."

For these reasons, it is recommended that you create a new link and link targets (essentially setting up an empty replica set), configure replication as desired, and copy the existing data into the new replica tree. After replication is complete, change the name of the link that is associated with the replica set so that users can access the new replica set by using the existing link name.

**Important**

- This procedure requires enough disk space on each replica member to store two copies of the files — one copy is the existing data, and the second copy is the data in the newly created replica tree.

In the following procedures, "existing link" refers to the link that already contains data in its link targets, and "new link" is the link that you create with replicated link targets.

**To create a new link and link targets**

1. If users can modify files in the link targets, disable referrals to the link targets for the existing link, and then close all open files.

2. Create a new link, and then follow the procedure in one of the following documents to set up an empty replica set:

   - "Deploying a Standard FRS Topology for an Empty Replica Set" (Sdcfsv_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Deploying a Standard FRS Topology for an Empty Replica Set" on the Web at http://www.microsoft.com/reskit).

   - "Deploying a Custom FRS Topology for an Empty Replica Set" (Sdcfsv_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Deploying a Custom FRS Topology for an Empty Replica Set" on the Web at http://www.microsoft.com/reskit).

   Be sure to disable referrals to the new link targets so that users do not access them.

3. Copy the data to be replicated into the replica tree.

   This step is complete when the files are replicated successfully to all replica members.

4. In the Distributed File System snap-in, under the existing link, add the same link targets that you added under the new link. Be sure to clear the **Add this target to the replication set** check box.

5. If you have not done so already, disable referrals, and then close all open files so that users are no longer accessing any files in the existing link targets.

6. Under the existing link, remove the link targets that are not replicated. After you complete this step, the existing link and the new link have identical link targets.

Figure 2.20 shows how two example links, Existing Link and New Link, appear in the Distributed File System snap-in at the end of this procedure.

**Figure 2.20 Two Links in the Distributed File System Snap-in**




Next, use the following procedure to change the link name that is associated with the new replica set to match the existing link name.

**To match the new link name with the existing link name**

1. In the Active Directory Users and Computers snap-in, click the **View** menu, and then click **Advanced Features**.

2. In the console tree, under the domain where the DFS root is located, navigate to **System**, **File Replication Service**, **DFS Volumes**.

   When you expand **DFS Volumes**, you should see existing DFS replica sets, including the one that you created in the previous procedure. The replica set is identified as **rootname|newlinkname**. Identify the correct replica set before you continue.

3. To change the name of the link that is associated with the replica set, right-click the replica set, and then click **Rename**.

4. After **rootname|**, type the name of the existing link over the name of the new link.

Figure 2.21 shows how the **rootname|newlink** name appears at Step 2 and after Step 4.

**Figure 2.21 Changing the Link Name**




To complete the deployment, delete the new link.

**To delete the new link**

1. In the Distributed File System snap-in, locate the new link.

   The new link should no longer have replication enabled. It might be necessary to close and reopen the snap-in for the updates to appear.

2. Right-click the new link, and then click **Delete Link**.

3. Enable referrals to the link targets for the existing link.

# Adding a New Member to an Existing Replica Set

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you add a new member to an existing replica set, determine how to source the new member. You have two choices:

- **Replicate the files to the new member across the network.** When you replicate files across the network to the new member, you add the server as a link target within an existing replica set. When you add the new member to the topology, you can create a single inbound connection so that you source from a specific server, such as a server with fast connectivity, or you can add a computer with redundant inbound connections where it sources according to connection priority and schedule.

- **Prestage the files on the new member.** When you prestage files, you use a backup program to back up the replica tree and restore it on the new member. Using a backup program is required to preserve the file object IDs and security descriptors. (Command-line tools such as Copy.exe, Xcopy.exe, and Robocopy.exe do not preserve these items, and they cannot be used to prestage files.) After you restore the files, FRS replicates to the new member any files that changed since the backup was created, thus minimizing the number of files that are replicated across the network.

The method you choose typically depends on the size of the replica set and the available network bandwidth. You should replicate files over the network only if one of the following conditions exists:

- You have a high-bandwidth network, and you can afford to use part of that bandwidth for the initial replication.

- You have a low-bandwidth network, and you can afford to use that bandwidth for initial replication, and you also have the time to wait for the initial replication to complete.

If neither of these conditions exists, you should prestage the new member.

The following sections describe three possible scenarios for adding a new replica member:

- Prestaging a new replica member to avoid replicating files across the network.

- Adding a new replica member that contains no content. In this scenario, files are replicated across the network.

- Adding a new replica member that contains a copy of the content. In this scenario, the new member has an existing copy of the files, but these files are not prestaged. The files from the initial master are replicated across the network.

**Prestaging a New Replica Member**

In this scenario, you have an existing replica set, and you want to add a new replica member by prestaging it. After you prestage the new member and enable replication, FRS moves the prestaged files in the replica tree on the new member to the NtFrs_PreExisting____See_EventLog folder, but FRS then moves back to the replica tree any files that have the same object ID and content as the files in the master replica set. FRS replicates across the network only the files that were added or changed since you prestaged the files.

Before you can prestage a new replica member, one of the following events must occur:

- At least seven days have passed since you enabled replication between the first two members.

- You clear the outbound change log by modifying the registry on all direct upstream partners.

To prestage a new replica member, perform the following tasks:

1. Configure the USN journal.

2. Prestage the new replica member.

3. Create a new link target, add it to the replica set, and then specify the staging directory location.
4. Disable referrals to the new link target.
5. Configure connection objects, the initial schedule, and connection priorities (custom topologies only).
6. Adjust the size of the staging directory.
7. Verify that replication is working on the new member.
8. Enable referrals to the link target.
9. Optimize the replication schedule.
10. Configure connection priorities (standard topologies only).

For a Word document to assist you in completing each of these tasks, see "Prestaging a New Replica Member" (Sdcfsv_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Pre-staging a New Replica Member" on the Web at http://www.microsoft.com/reskit). After you complete these tasks, you can notify users that the DFS namespace is available.

**Adding a New Replica Member That Contains No Content**
In this scenario, the new replica member does not contain a copy of the replica tree. Data will replicate across the network according to the replication schedule.

To add a new replica member that contains no content, perform the following tasks:

1. Configure the USN journal.
2. Create a new link target, add it to the replica set, and then specify the staging directory location.
3. Disable referrals to the new link target.
4. Configure connection objects, the initial schedule, and connection priorities (custom topologies only).
5. Adjust the size of the staging directory.
6. Verify that replication is working on the new member.
7. Enable referrals to the link target.
8. Optimize the replication schedule (optional).
9. Configure connection priorities (standard topologies only).

For a Word document to assist you in completing each of these tasks, see "Adding a New Replica Member that Contains No Content" (Sdcfsv_6.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Adding a New Replica Member that Contains No Content" on the Web at http://www.microsoft.com/reskit). After you complete these tasks, you can notify users that the DFS namespace is available.

**Adding a New Replica Member That Contains a Copy of the Content**
In this scenario, the new replica member already contains a copy of the replica tree. You might have copied the data to the server manually (using Robocopy, for example), but you have not prestaged the new member. (Remember that prestaging preserves the file object IDs and security descriptors.) Because the data is not prestaged, the initial master replicates the entire replica tree to the new member, even if the new member already contains identical files. On the new replica member, FRS moves any files that existed in the replica tree to a folder named "NtFrs_PreExisting____See_EventLog." After replication completes, you can either delete this folder or move any unique files back into the replica tree.

To add a new replica member that contains a copy of the content, perform the following tasks:

1. Follow the procedures in "Adding a New Replica Member that Contains No Content" (Sdcfsv_6.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Adding a New Replica Member that Contains No Content" on the Web at http://www.microsoft.com/reskit).

2. After replication completes, locate the folder "NtFrs_PreExisting____See_EventLog" on the new replica member.

3. Do one of the following:

   - If the folder contains unique files that are not part of the replica set, copy those files back into the replica tree.

   - If you do not want to save any of the files, delete the folder.

# Configuring File Server Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many procedures for configuring file server settings are provided in Help and Support Center for Windows Server 2003. Table 2.23 lists subjects and file server topics that can be useful for configuring file server settings. For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

**Table 2.23 Help and Support Center Topics Related to File Server Settings**

| Subject | Topic Title |
|---|---|
| Sharing folders | "Share a folder or drive" |
| Enabling disk quotas | "Checklist: Setting up disk quotas" |
| Enabling shadow copies | "Checklist: Deploying Shadow Copies for Shared Folders" |
| Setting NTFS permissions | "Set, view, change, or remove permissions on files and folders" |
| Enabling encryption on file servers | "Enable a remote server for file encryption" |
| Remotely managing file servers | "Using Web Interface for Remote Administration" |
| Managing disks and volumes | "Disk Management overview" and "DiskPart" |
| Enabling encryption on clustered file servers | "Create a cluster-managed encrypted file share" |
| Enabling client-side caching on clustered file servers | "Set client-side caching for a File Share resource" |
| Enabling shadow copies on cluster storage | "Enable Shadow Copies for Shared Folders in a cluster" |

# Additional Resources for Designing and Deploying File Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Related Information**

- "Hosting Applications with Terminal Server" in this book.
- "Planning for High Availability and Scalability" in this book.
- "Designing and Deploying Server Clusters" in this book.
- "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit.
- "Implementing User State Management" in *Designing a Managed Environment* of this kit.
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about the Distributed File System and the File Replication service.
- The *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit) for more information about disk management and file systems.
- The *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit) for information about interoperability with NetWare and interoperability with UNIX.
- "Deployment, Monitoring and Troubleshooting of the Windows 2000 File Replication Service Using the SONAR, TOPCHK, CONSTAT and IOLOGSUM Tools" (Troubleshooting_frs.doc), included with the version of Sonar.exe that is available from the Free Tool Downloads link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The File Services Community Center link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about how Windows server technologies deliver reliable file services that are essential to enterprise computing infrastructures.
- The Windows Powered Network Attached Storage link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Storage Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The File Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Microsoft® Windows SharePoint® Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for information about sharing information within organizations and over the Internet.

**Related Job Aids**

- "DFS Configuration Worksheet" (Sdcfsv_1.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).
- "FRS Configuration Worksheet" (Sdcfsv_2.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "FRS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

- "Deploying a Standard FRS Topology for an Empty Replica Set" (Sdcfsv_3.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Deploying a Standard FRS Topology for an Empty Replica Set" on the Web at http://www.microsoft.com/reskit).

- "Deploying a Custom FRS Topology for an Empty Replica Set" (Sdcfsv_4.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Deploying a Custom FRS Topology for an Empty Replica Set" on the Web at http://www.microsoft.com/reskit).

- "Prestaging a New Replica Member" (Sdcfsv_5.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Prestaging a New Replica Member" on the Web at http://www.microsoft.com/reskit).

- "Adding a New Replica Member That Contains No Content" (Sdcfsv_6.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Adding a New Replica Member That Contains No Content" on the Web at http://www.microsoft.com/reskit).

- "Shadow Copy and Disk Quota Configuration Worksheet" (Sdcfsv_7.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "DFS Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Disk Management overview" and "DiskPart" in Help and Support Center for Windows Server 2003 for information about managing disks and volumes.

- "Extend a basic volume" in Help and Support Center for Windows Server 2003.

- "Using NTFS mounted drives" and "Create a mounted drive" in Help and Support Center for Windows Server 2003.

- "Shadow Copies for Shared Folders" in Help and Support Center for Windows Server 2003.

- "Vssadmin" and "Schtasks" in Help and Support Center for Windows Server 2003 for information about scheduling shadow copies from the command line.

- "Checklist: Setting up disk quotas" in Help and Support Center for Windows Server 2003.

- "Share a folder or drive" in Help and Support Center for Windows Server 2003.

- "Publish a shared folder" in Help and Support Center for Windows Server 2003 for information about publishing shared folders in Active Directory.

- "Offline settings for shared resources," and "Make a file or folder available offline" in Help and Support Center for Windows Server 2003.

- "Permissions on a file server" in Help and Support Center for Windows Server 2003.

- "Set, view, change, or remove permissions on files and folders" in Help and Support Center for Windows Server 2003.

- "Enable a remote server for file encryption" in Help and Support Center for Windows Server 2003.

- "Encrypting and decrypting data" and "Internet Information Services (IIS) 6.0 overview" in Help and Support Center for Windows Server 2003 for information about EFS and using WebDAV folders to store encrypted files.

- "Install Windows Support Tools" in Help and Support Center for Windows Server 2003 for information about installing Windows Support Tools, including Dfsutil.exe.

- "Using Web Interface for Remote Administration" in Help and Support Center for Windows Server 2003.
- "Chkdsk" in Help and Support Center for Windows Server 2003 for information about running Chkdsk.exe.
- "Checklist: Creating a distributed file system" and "Dfscmd" in Help and Support Center for Windows Server 2003 for information about creating a DFS namespace.
- "Installing and Upgrading on Cluster Nodes" in Help and Support Center for Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition for comprehensive information about choosing the cluster installation method and performing the actual installation.
- "Managing a Server Cluster from the Command Line" in Help and Support Center for Windows Server 2003.
- "Create a cluster-managed file share" in Help and Support for Windows Server 2003.
- "Create a cluster-managed encrypted file share" in Help and Support Center for Windows Server 2003 for information about enabling EFS on server clusters.
- "Set client-side caching for a File Share resource" in Help and Support Center for Windows Server 2003 for information about configuring client-side caching on a server cluster.
- "Using Shadow Copies of Shared Folders in a server cluster" and "Enable Shadow Copies for Shared Folders in a cluster" in Help and Support Center for Windows Server 2003 for in-depth information about using shadow copies on server clusters.
- "Change the Cluster service account password" in Help and Support Center for Windows Server 2003.
- "Enable Kerberos authentication for virtual servers" in Help and Support Center for Windows Server 2003 for information about allowing anonymous access to specific file shares on a server cluster.
- "Backing up and restoring server clusters" and "Test Cluster Failures and Failover Policies" in Help and Support Center for Windows Server 2003.
- "Best practices for securing server clusters" in Help and Support Center for Windows Server 2003.

**Related Tools**

- Dfscmd.exe

  Use Dfscmd.exe to create links and add link targets. For more information about Dfscmd.exe, see Command-line reference A-Z.

- Dfsutil.exe

  Use Dfsutil.exe to configure and troubleshoot DFS namespaces. For more information about Dfsutil.exe, see Windows Support Tools.

- DiskPart.exe

  Use DiskPart.exe to manage disks and volumes from the command line. For more information about DiskPart.exe, see Command-line reference A-Z.

- Permcopy.exe

  Use Permcopy.exe to copy shared folder permissions from one shared folder to another. For more information about

Permcopy.exe, see Windows Resource Kit Tools Help.

- Robocopy.exe

  Use Robocopy.exe to maintain identical copies of a folder tree in multiple locations. For more information about Robocopy.exe, see Windows Resource Kit Tools Help.

- Sonar.exe

  Use Sonar.exe to monitor key statistics about FRS members in a replica set. For more information about Sonar.exe, see Windows Resource Kit Tools Help.

- Windows Server 2003 Administration Tools Pack

  Use the Windows Server 2003 Administration Tools Pack to remotely manage servers running Windows Server 2003 from computers running Windows XP. For more information about installing this tools pack, see article Q304718, "Administer Windows 2000–Based Computers Using Windows XP Professional Clients," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Designing and Deploying Print Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Adding or upgrading print servers, consolidating servers, or creating print server clusters can increase the reliability, availability, and manageability of your printing environment. Network and print administrators in medium-to-large organizations can deploy and manage print servers more efficiently by properly planning and deploying a network printing solution.

**In This Chapter**

Overview of the Print Server Deployment Process

Inventorying the Current Printing Environment

Designing the New Printing Environment

Deploying Print Servers

Additional Resources for Designing and Deploying Print Servers

**Related Information**

- For conceptual information about network printing and Windows clustering, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).
- For more information about deploying Windows server clusters, see "Designing and Deploying Server Clusters" in this book.
- For information about deploying file servers, see "Overview of Designing and Deploying File Servers" in this book.

# Overview of the Print Server Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying print servers running the Microsoft® Windows® Server 2003, Standard Edition or Windows® Server 2003, Enterprise Edition operating systems lets your organization share printing resources across the network, allowing client computers to send print jobs to printers attached locally to a print server and to printers accessible across the Internet. Windows Server 2003 provides a number of solutions that can benefit your organization.

**Cross-platform printing support** Windows Server 2003 supports printing from Microsoft® Windows® 95, Windows® 98, Windows® Millennium Edition, Windows NT® version 4.0, Windows® 2000, Windows® XP Professional, or UNIX or Linux operating systems.

**Increased availability** If your business relies on high availability of network printers, you can cluster print servers or install a standby server to increase print server availability.

**Driver versioning** To improve stability, Windows Server 2003 restricts the installation of older printer drivers, which can cause the system to become unresponsive. You can override these restrictions if your network has clients that require older printer drivers.

**Consolidation tools** If you have too many print servers to manage efficiently in your current printing environment, you can use the Print Migrator tool to help automate printer consolidation on Windows Server 2003. Instead of recreating all of the printer objects manually, you can back up existing printers and restore them to your new print servers.

**Rich printer status reporting** If you use the standard port monitor, you can receive status reports when the printer runs out of paper, runs low on toner, or is stopped by a paper jam.

**Easy location of nearby printers** If you implement the Active Directory® directory service, your users can easily find printers that are published in Active Directory based on attributes that you designate.

**Improved security** New local Group Policy settings can control which individuals have access to the print queue over the network.

**Centralized printer management** Windows Server 2003 lets you remotely manage and configure printers from any computer running Windows Server 2003.

If your current network environment has a mix of clients and print servers running different operating systems, you can improve the interoperability, availability, and efficiency of your print servers by properly designing your new printing environment. After the design team creates a design, and the deployment team implements that design in a test lab, you are ready to deploy your new Windows Server 2003 print servers in a production environment.

Effectively deploying print servers that run Windows Server 2003 is a three-phase process, as illustrated in Figure 3.1.

**Figure 3.1 Designing and Deploying Print Servers**




A print server design team completes the first two phases: examining your current environment and designing a new printing environment. During these phases, the design team collects information that is vital to the deployment process, determines your immediate and long-term printing needs, recommends improvements to the printing environment, and creates an appropriate printing strategy.

A print server deployment team completes the final phase of the deployment process by implementing the design team's

decisions. The deployment team's responsibilities include adding printing capabilities to existing or newly created clustered servers, installing Windows Server 2003 on the print servers, creating the printers, and ensuring that clients receive the required printer connections.

**Note**

- For a list of the job aids that are available to assist you in deploying print servers, see "Additional Resources for Designing and Deploying Print Servers" later in this chapter.

# Inventorying the Current Printing Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first phase in moving to a Windows Server 2003–based printing environment is gathering information about your current printing environment. Document the port types, printer names, share names, IP address, driver models and driver versions of each printer that you plan to replace or migrate. This information helps you plan details of the deployment process and gives you a way to rebuild printers if a server fails.

Figure 3.2 illustrates the place of this phase in the design and deployment process.

**Figure 3.2 Inventorying the Current Printing Environment**




A number of tools can help you inventory your current printing environment. These tools can also be useful if you decide to migrate your existing print servers to new print servers that are running Windows Server 2003. The tools you can use depend on your current operating system:

- Windows Server 2003 includes scripts written in the Microsoft® Visual Basic® Scripting Language that you can use to inventory your current printing environment.
- If you are in a Windows NT 4.0 domain, you must use PrnAdmin to inventory your current printing environment.

**Note**

- For a worksheet to assist you in inventorying your printing environment, see "Print Server Inventory Worksheet" (Sdcpsv_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Print Server Inventory Worksheet" on the Web at http://www.microsoft.com/reskit).

#### Visual Basic Scripts

Table 3.1 lists information that you might require while inventorying your printing environment and the scripts you can use to obtain that information. These scripts are located in the Windows\System32 directory.

**Table 3.1 Scripts That Provide Print Server Information**

| Information Needed | Script |
|---|---|
| Print server name | Prnmngr.vbs |
| Print queue name | Prnmngr.vbs |
| Printer driver name | Prnmngr.vbs |
| Print port type | Prnport.vbs |
| Print port configuration | Prnport.vbs |
| Print device IP address | Prnport.vbs |

For a list of parameters for each script listed in Table 6.1, perform the following procedure.

**To list the parameters for the Visual Basic scripts listed in Table 6.1**

1. Open a **Command Prompt**, and then move to the Windows/System32 directory.

2. At the command prompt, type:

   **cscript *scriptname*.vbs**

**PrnAdmin**
If you are migrating from Windows NT 4.0 to Windows Server 2003, you can use Prnadmin.dll to inventory your current environment. PrnAdmin enumerates ports, drivers, printers, or forms on local or remote printers. The Visual Basic script for the PrnAdmin utility is named Prnmgr.vbs.

You can find the PrnAdmin utility on the *Windows Server 2003 Deployment Kit* companion CD. To install PrnAdmin, type the following command:

```
regsvr32 "C:\Program Files\Resource Kit\PrnAdmin.dll"
```

### Important

- The preceding command syntax assumes that you installed the contents of the *Windows Server 2003 Resource Kit* companion CD in the default location, C:\Program Files\Resource Kit. If you install the contents in a different location, you must use that path in the syntax.

**PrnAdmin syntax**
The command-line syntax for PrnAdmin is:

```
cscript prnmgr.vbs [-adx[c]] [l] [-cserver] [-bprinter] [-mdriver model]
[-pdriver path] [-r port] [-f file]
```

**PrnAdmin example**
The following command lists all printers on the print server \\adminserver:

```
cscript prnmgr.vbs -l -c \\adminserver
```

# Designing the New Printing Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After inventorying your current printing environment, your design team can start designing your new printing environment. Figure 3.3 illustrates the component tasks in this phase and its position in the design and deployment process.

**Figure 3.3 Designing the New Printing Environment**




**Note**

- For a worksheet to assist you in designing your printing environment, see "Print Server Design Worksheet" (Sdcpsv_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Print Server Design Worksheet" on the Web at http://www.microsoft.com/reskit).

# Increasing the Efficiency of the Printing Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When designing your new printing environment, examine your existing environment to determine if you need to make more efficient use of your resources. Examine the efficiency of both printing performance and printing management.

**Printing Performance**

You might need to make adjustments to servers that are not responding at the levels that you require. If users are waiting an unacceptable amount of time for jobs to print, you can improve the performance of your print servers by replacing older print devices with newer, faster devices.

Another option is to create a printer pool in which two or more of the same model printers connect to one print server and behave as a single printer. You can make the printer pool appear to users as a single printer by using multiple ports on the print server to connect the printers to one logical printer (the printer's software interface on the print server). When a user prints a document to the printer pool, the document prints to the first available printer in the pool. Users do not have to know which printer is available. The logical printer checks for an available port and then sends documents to ports in the order in which the documents are added to the print queue. A printer pool also eases administration, because you can manage multiple printers from the same logical printer on the print server.

If you have remote offices that experience bandwidth limitations while connected to a central print server, consider placing print servers in those offices. However, increasing performance in this manner has a drawback: If the print server fails, someone on-site must troubleshoot or restart the server. In addition, if your print server also hosts other services, restarting the server causes those other services to go offline as well.

**Printing Management**

You might also have more print servers than you need. If you have more print servers than you can manage efficiently, consider consolidating some of your print servers. For more information about consolidating print servers, see "Designing for Efficient Use of Hardware in Printing" later in this chapter.

# Optimizing Print Server Availability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By evaluating your availability requirements, you can determine if you need to increase print server availability and minimize downtime if a server stops responding. Certain users or groups in your organization might need highly available network printing. An accounting group, for example, might print time-sensitive documents such as payroll checks. In such a case, you need to examine high-availability solutions for that group.

If you determine that you have users who need highly available printing, you should consider either creating a print server cluster or putting a standby system in place in the event that your print server becomes unavailable. Both of these methods have trade-offs in cost and manageability. The following sections describe these trade-offs and can help you decide which method is more appropriate for your organization.

**Creating Clustered Print Servers to Increase Availability**

A *cluster* is a group of individual computer systems working together to provide increased computing power and to ensure continuous availability of mission-critical applications or services. The computers in this group use a common cluster name, appear to network clients as a single system, and distribute the cluster's workload among the group.

Clustered print servers provide failover and failback capabilities, which enhance the availability of print servers by allowing individual print servers, known as *nodes*, to fill in for each other if one of them fails (although there might be a short period when the printers are unavailable while the resources fail over to the other nodes). With clustered print servers, you also can perform maintenance on one node, without interrupting printing operations, by moving the cluster resources to another node.

You can create clusters only on servers that run Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. However, after you set up a cluster on these computers, you can administer the cluster, including installing and configuring printers, from any computer on the network that is running Windows 2000, Windows XP Professional with Service Pack 1 and the Administration Tools Pack, or Windows Server 2003. For information about installing the Administration Tools Pack, see "Windows Server 2003 Administration Tools Pack Overview" in Help and Support Center for Windows Server 2003.

Before deciding whether to use clustered print servers to increase printer availability, consider the following requirements:

- Clustered print servers have additional costs and hardware requirements. A clustered print server requires two or more physical servers running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. Because of the specialized hardware and the software requirements, the overall cost of implementing clustered print servers can be greater than the cost of providing a standby print server.

- Clustered print servers require administrator expertise. Implementing and maintaining clustered print servers requires experience or specialized training. Less experienced administrators might cause significant problems and delays, and mistakes can lower the availability of the print servers.

- For more information about deploying Windows server clusters, see "Designing and Deploying Server Clusters" in this book.

**Creating Standby Servers to Increase Availability**

Using a standby server to increase print server availability is a simple, low-cost alternative to clustering. The best method for setting up a standby server is to make sure that the hardware and software, including drivers, are exactly the same as the standard print server configuration that the design team has created. If you have a standby print server based on the standard print server configuration, you can replace any print server on your network that is based on the same configuration. In the event that a print server stops responding, rename the standby server to the same name as the server that you are replacing, and then connect the standby server to your network. For more information about standard print server configurations, see "Designing for Efficient Use of Hardware in Printing" later in this chapter.

Consider the following issues before creating standby print servers:

- Standby print servers carry additional costs. Providing a standby server requires two servers for each of the print servers that you want to be highly available. Ideally, these two servers have identical hardware and software configurations, including patches and printer drivers. The minimum requirements for a standby server are the same as the requirements for any server running Windows Server 2003.

- Standby print servers require additional maintenance. If you use standby servers, any changes that you make to the primary print server need to be made to the standby server also. If a server stops responding, you can rename the standby server and remove the primary print server from the network by using a Remote Desktop Connection.

- This method of increasing availability can result in some downtime, which is limited to the time that it takes to rename the standby print server, attach it to the network, and remove the primary print server.

- Users must resubmit any print jobs that failed to print. Print jobs that have been spooled but have not yet printed when the server stops responding must be resubmitted to the standby server after it replaces the primary print server.

# Determining Hardware and Driver Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before deploying a new Windows Server 2003 print environment, examine your existing print server hardware. If any of your hardware does not meet the recommended hardware requirements for Windows Server 2003 and does not appear in Windows Server Catalog, you must acquire new hardware or upgrade your existing hardware. For more information about hardware requirements, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources

You must also consider which printer driver versions you will need. Depending on the mix of client operating systems in your environment, you might need to install different printer driver versions on different client computers in order to print to the same printer.

# Designing for Efficient Use of Hardware in Printing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When determining your print server hardware requirements, consider both the sizing and placement of servers to maximize availability of printing resources, and determine whether you want to consolidate any existing print servers. After you make these design choices, consider creating a standard print server configuration to simplify the management and maintenance of your hardware.

**Determining Print Server Sizing**

When planning the printing strategy for your organization, consider the following factors as you estimate the number and types of printers that you need, both now and in the future:

- Determine how to divide and allocate printing resources. High-volume printers generally have more features, but a breakdown affects more users.
- Consider the printer features that you need, such as color printing, duplex printing, multibin mailboxes, envelope feeders, and stapling. Determine which users need these features and the physical locations of these users.
- Estimate any additional hardware requirements. Managing a large number of printers or many large documents requires a large amount of memory and disk space, and, potentially, a more powerful computer. If you need to store a copy of everything you print, you must provide additional disk space.
- Consider dedicating a server exclusively to printing. If you are considering using a server for both file and print sharing, remember that printing operations require a large amount of data input/output (I/O), and that the operating system gives priority to file operations over printing operations.
- Match printing volume as closely as possible to a printer's duty cycle (the number of pages the printer was designed to print per month). Maintenance problems are generally reduced if printing volume is less than or equal to the printer's duty cycle.
- Consider printing-speed requirements. In addition to the processing speed of the printer itself, network considerations also affect printing speed. Typically, printers that are attached directly to the network with network adapters offer faster throughput than printers attached with parallel buses. However, print throughput rates also depend on network traffic, the type of network adapter, and the protocol used.
- To increase the throughput of a print server, relocate the print spool folder. For best results, move the spool folder to a dedicated disk drive that does not share any files, including paging files, with the operating system.

The appropriate hardware for an organization varies widely depending on business requirements. Many businesses have relatively simple printing needs, while others have more complex printing needs. For example, when the graphics department of an advertising agency determines that it needs a high-speed color printer, the administrator locates the printer in the same area of the building as the graphics department. To avoid conflicts with other I/O operations on the server, the administrator also designates the D: drive on the print server as the print spool folder. For more information about print server sizing, see the Print Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources

**Determining Print Server Placement**

The placement of printers and print servers can affect both the efficiency of users and the congestion on your network. To minimize printing traffic in your network environment, locate printers close to the people who use them. Also, if some users print sensitive documents, consider locating their printers in an area of the building that only authorized personnel can access.

**Determining Whether to Consolidate Printers**

If you determine that you have more print servers than you can manage efficiently, consolidate some of them onto one of your more powerful print servers. Central print server management offers an easier way to administer print servers, and consolidation reduces the cost incurred by maintaining more print servers than you need. At the same time, consolidating servers increases the value of recovery planning and procedures.

As you consider whether or not to consolidate print servers, weigh the following advantages and disadvantages.

**Advantages of consolidating printers**

- You have fewer print servers to maintain.
- You can reassign the servers that no longer serve as print servers.

**Disadvantages of consolidating printers**

- If a server stops responding, your users have fewer servers to depend on.
- If the print server is no longer near users on the network, your users might experience delays in processing print jobs.

Using Print Migrator version 3.0 helps you automate the process of consolidating your print servers. You merely back up your existing servers and restore them to the new Windows Server 2003 print server. For more information about Print Migrator 3.0, see "Choosing a Print Server Migration Method" later in this chapter.

**Creating a Standard Print Server Configuration**

Creating a standard print server configuration is highly recommended. This involves defining a specific hardware and software combination that works well for your organization and operates with minimal maintenance. The specific configuration varies from business to business. The needs of your organization might necessitate more than one standard print server configuration.

Designing a standard configuration for print servers provides two important benefits:

- It reduces the complexity of managing and maintaining hardware. You need to learn hardware management and maintenance tasks for only one type of print server instead of for multiple types.
- It reduces the amount of testing that you must perform when updating drivers on the print server. Instead of testing the updates on multiple hardware and software configurations, you can test the updates on one print server and then deploy the updates to the other print servers.

If your organization has selected Cluster service to provide the highest level of availability, create a standard configuration for clustered print servers as well. For more information about clustering, see "Designing and Deploying Server Clusters" in this book.

As you design your configuration, consider creating a logical and standard naming convention for your printers, which can make it easier for users to locate printers across a network if you decide to publish your printers in Active Directory.

# Designing for Printer Driver Compatibility

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Designing for Printer Driver Compatibility**

When designing your new printing environment, you must understand how different types of drivers work on different versions of the Windows operating system. Print servers running Windows Server 2003 use two types of drivers:

- Version 2 printer drivers operate in *kernel mode* (inside the operating system kernel memory space) and are designed for Windows NT 4.0.
- Version 3 printer drivers operate in *user mode* (outside the operating system kernel memory space) and are designed for Windows 2000, Windows XP, and Windows Server 2003.

Windows Server 2003 supports both Version 2 and Version 3 printer drivers. To ensure the stability of your print servers, use Version 3 drivers whenever possible. If some computers in your organization run Windows NT 4.0 and other computers run Windows 2000, Windows XP Professional, or Windows Server 2003, you must ensure that you have the proper mix of drivers to support all of these clients, and you might have to modify the Windows Server 2003 print server to include additional Version 2 drivers that work with Windows NT 4.0 clients.

**Note**

- Poorly written device drivers have been identified as a major cause of system instability. To ensure that your drivers are reliable, use drivers approved under the Windows Hardware Logo Program on your print servers.

Having the correct drivers installed on your server is important if you plan to use the Point and Print feature. By using Point and Print, a user running Windows NT 4.0, Windows 2000, Windows XP Professional, or Windows Server 2003 can create a connection to a remote printer without using disks or other installation media. All necessary files and configuration information automatically download from the print server to the client. Point and Print eases your administrative overhead by eliminating the need to visit each workstation to install the appropriate device drivers.

Windows Server 2003 also provides Point and Print support for Windows 95, Windows 98, and Windows Millennium Edition clients on your network. However, Point and Print does not support automatic driver upgrades for these clients.

Your new print servers must include printer drivers that support all the clients on your network. Maintain a list of printers that are using Version 2 drivers, so that you can replace these drivers with Version 3 drivers if you upgrade the client base to Windows 2000 or Windows XP Professional. Additionally, if you have PostScript printers, you can use PostScript drivers to increase interoperability in a mixed Windows NT 4.0, Windows 2000, and Windows XP Professional client environment.

Table 3.2 notes the compatibility of Version 2, Version 3, and PostScript drivers with different versions of the Windows operating systems.

**Table 3.2 Printer Driver Compatibility with Windows Operating Systems**

| Driver type | Version 2 | Version 3 | PostScript |
|---|---|---|---|
| Windows NT 4.0 | ● | Supported as additional drivers* | ● |
| Windows 2000 | ● | ● | ● |
| Windows XP Professional | ● | ● | ● |
| Windows Server 2003 | Supported as additional drivers | ● | ● |

* Must be installed remotely from Windows 2000 and Windows XP Professional clients to allow those clients to print to a Windows NT 4.0 print server.

**Version 2 and Version 3 Drivers**

Version 3 drivers are recommended over Version 2 drivers because they support advanced printing features and improved performance. If a server stops responding as a result of a Version 3 driver problem, you only need to stop and restart the print

spooler instead of restarting the server. Additionally, Version 3 drivers can reduce printer downtime, because Windows Server 2003 provides automatic recovery of the spooler service by default. If a Version 2 driver stops responding, you must reboot the server, taking all other services offline while the server restarts.

Often the existing Version 2 drivers on your print server function properly with Windows Server 2003. Although a default local Group Policy setting in Windows Server 2003 prohibits the installation of new Version 2 drivers, this policy setting does not affect existing Version 2 drivers when you perform an upgrade unless a driver is known to be problematic. To identify drivers that will not carry through the upgrade, use the command-line tool Fixprnsv.exe prior to upgrading.

For more information about Fixprnsv.exe, see "Identifying Interoperability Issues with Print Server Upgrades" later in this chapter. For more information about disallowing installation of printers using kernel-mode drivers, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**PostScript Drivers**

Microsoft recommends that you consider using PostScript minidrivers (.ppd files) to increase compatibility in mixed environments with Windows NT 4.0, Windows 2000, and Windows Server 2003 systems. Computers running Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003 can all run the same PSCRIPT5 core code, regardless of whether these computers are running Version 2 or Version 3 drivers. This offers a higher level of interoperability among clients running different versions of Windows.

**Choosing Drivers for Point and Print Connectivity**

If you are using Point and Print to connect clients to printers, driver compatibility is particularly important. Clients that use the same driver architecture as the print server can simply download the primary driver used on the server and therefore do not require the installation of any additional drivers on the print server for full support. For example, a Windows Server 2003 print server using a Version 3 driver for a shared printer does not require the installation of any additional drivers to provide full Point and Print support for Windows 2000 and Windows XP Professional clients, because those clients also support Version 3 drivers.

If Windows 95, Windows 98, Windows Millennium Edition, and Windows NT 4.0 clients connect to the print server, you must install the appropriate driver as an additional driver on the Windows Server 2003 print server.

For additional information about network printing, including Point and Print driver installation, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Point and Print for clients running Windows NT 4.0 and earlier versions of Windows**

Setting up printer drivers on a print server running Windows Server 2003 to support Point and Print is generally the same for clients running Windows NT 4.0 and earlier versions of Windows as it is for clients running Windows 2000 or Windows XP. However, you may encounter some specific differences.

You can select the printer driver versions to install on the print server on the **Drivers** tab of the **Server Properties** property sheet, or, when you install additional drivers using the Add Printer Driver Wizard, in the **Additional Drivers** dialog box shown in Figure 3.4.

**Figure 3.4 Additional Drivers Dialog Box**




Before you load any additional drivers on x86-based systems, one of three printer driver versions is selected by default,

represented by the three following options:

- **Windows 2000, Windows XP, or Windows Server 2003.** If this option is selected before you load any additional drivers, it indicates that the server is running a Version 3 driver written specifically for the Windows 2000, Windows XP, or Windows Server 2003 platform. You cannot send that driver to a Windows NT 4.0 client for Point and Print functionality. You need to add a Version 2 driver to support Point and Print on Windows NT 4.0 client computers.

- **Windows NT 4.0.** This option indicates that a Windows NT 4.0 Version 2 driver is running on a server that is running Windows 2000 or Windows Server 2003. No additional drivers are required to support Point and Print on Windows NT 4.0–based client computers.

- **Windows 95, 98, or Windows Millennium Edition.** If this option is selected, users can connect to and download the existing driver from the server. However, those clients will not receive automatic updates when the driver on the server is updated.

**Point and Print for servers running Unidrv Version 3 drivers**

When a Version 3 driver written specifically for the Windows Server 2003 or Windows 2000 platforms is installed on the server, you must install a Version 2 driver with the same model name on Windows NT 4.0 clients in order to support Point and Print. Changes in the internal data structures between the Version 2 and Version 3 Unidrv.dll can cause problems in sharing or preserving settings between the clients and the server. These problems primarily manifest themselves in two ways:

- New or advanced settings on the Print Server might not be visible from the property sheet of a client printer.

- The print server might not correctly interpret the client printer properties, causing the properties to be lost. As a result, the Point and Print connection allows the client to print, but might not allow modification of certain device settings or preferences if they do not match those on the server.

### Note

- Incompatible data structures can cause certain printer settings — such as **Print Text as Graphics** or **Print Optimization** — to be lost between computers running Windows NT 4.0 and computers running Windows 2000. Additionally, settings such as paper tray assignments and media types might not be preserved between driver models.

# Choosing the Print Server Installation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Print server installation includes installing or upgrading to the Windows Server 2003 operating system and creating or migrating printers. The type of installation that you choose depends on your existing printing environment. If you have new servers with no operating system installed, you must perform a clean installation of the Windows Server 2003 operating system. If you have existing servers, you can perform a clean installation or, if the server is running Windows 2000 or Windows NT 4.0, you can upgrade.

Other considerations might also play a part in your decision. For example, some organizations have software standards that prohibit upgrading the operating system on servers. These organizations choose to perform only clean installations. With a clean installation, you can install the latest printer drivers from the independent hardware vendor (IHV), thus eliminating older, potentially problematic drivers.

Consider the following points when choosing between a clean installation and an upgrade.

**Choosing a clean installation**

- In a clean installation, you must create all new printers or migrate existing printers.
- If you want to practice good print queue management (for example, by complying with the company's print queue naming conventions), you might want to perform a clean installation instead of upgrading a server.

**Choosing to upgrade**

- Configuration is simpler when you upgrade. You retain your existing print queues, drivers, and ports, minimizing the impact on users.
- You might encounter interoperability issues with your existing printer drivers.

# Choosing a Print Server Migration Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you decide that a clean installation of Windows Server 2003 is best for your organization, you must also decide if moving some of your existing printers from other print servers makes sense. Migrating a print server is a method of moving existing printers onto your new server. If you intend to add new print servers to your existing environment, consider migrating some of your existing print servers to these new, typically more powerful servers to take full advantage of your improved hardware resources.

Three approaches are available for migrating existing print server configurations:

- Manual migration
- Automated migration
- Partially automated migration

**Manual Print Server Migration**
A manual print server migration requires that you document the current printing environment on your existing Windows print servers (including ports, printer names, share names, and driver models and versions). Then, using the Add Printer Wizard, recreate each printer on your new Windows Server 2003 print servers. The advantage of this method is that you can create all of the queues by using the most recent corporate standard and the newest available driver versions. The disadvantage is that this method can be very time consuming for large print environments — and it is more error-prone than the automated method.

**Automated Print Server Migration with Print Migrator**
Microsoft has developed a printer migration utility called Print Migrator, which can back up and restore print server configurations with minimal user intervention. Print Migrator handles the migration of print queues, drivers, and printer ports. You can also use Print Migrator to consolidate print servers.

Print Migrator does not change the driver versions when moving printers to the Windows Server 2003. For example, Print Migrator copies any Version 2 drivers used on a Windows NT 4.0 print server and recreates them on the new Windows Server 2003 print server. This process provides a much higher level of interoperability with any Windows NT 4.0 clients that might still be using print services from that server. To download Print Migrator, see the Print Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Note**

- Print Migrator ships with documentation about its operation and use. Review the material carefully for more details about this tool's capabilities and limitations.

The newest version of Print Migrator, Print Migrator 3.0, provides the following functionalities:

- Version support across operating systems
- Cluster support
- Port conversion from Line Printer Remote (LPR)-to-standard port monitor (standard TCP/IP port monitor)

**Version support across operating systems**
Differences in the way that Version 3 and Version 2 drivers work present interoperability issues when you migrate printers across versions of the Windows operating system. Cross-version support in Print Migrator 3.0 ensures successful migration from Windows NT 4.0 to Windows 2000 or Windows Server 2003.

**Cluster support**
Print Migrator 3.0 also supports migrating your print cluster from a server running Windows NT 4.0 or Windows 2000 to a print server running Windows Server 2003.

**LPR to standard TCP/IP port conversion**

If your print servers currently have LPR ports, consider converting the ports to standard TCP/IP ports. The standard TCP/IP port is the preferred printer port in Windows Server 2003. You can use other printer ports with Windows Server 2003, but they offer limited functionality.

**Partially Automated Print Server Migration**

Administrators can migrate printers by using a combination of manual and automated methods. Scripts supported by Windows 2000, Windows XP, and Windows Server 2003 can perform most of the necessary tasks. For more information about using scripts to partially automate your migration, see "Managing printing from the command line" in Help and Support Center for Windows Server 2003.

# Identifying Interoperability Issues with Print Server Upgrades

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Upgrading print servers from previous versions of Windows can present interoperability issues because of changes in the printer driver model beginning with Windows 2000 and continuing with Windows Server 2003. For more information about driver compatibility, see "Designing for Printer Driver Compatibility" earlier in this chapter.

To handle driver-related issues more easily, use the command-line utility Fixprnsv.exe, a powerful tool for working with printer driver issues on printer servers that are running Windows NT 4.0, Windows 2000, or Windows Server 2003. The Fixprnsv.exe identifies incompatible drivers before you upgrade, finds all known bad drivers on your system, and identifies Version 2 drivers that work properly with Windows Server 2003.

The Windows Server 2003 installation CD-ROM includes Fixprnsv.exe. The executable file is located in Windows\Printers directory. To view a list of commands used with this utility, use the command **fixprnsv.exe /?** from that directory.

For more information about specific compatibility issues involved in upgrading from Windows NT or Windows 2000, see "Deploying Print Servers with Upgrade Installations" later in this chapter.

# Choosing an Installation Method for Clustered Print Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Several options are available for upgrading a clustered print server. You can use any of the following methods:

- Upgrade a cluster that is running Windows NT Server 4.0, Enterprise Edition with Service Pack 4 or later.
- Upgrade a cluster that is running Microsoft Windows 2000 Advanced Server or Microsoft Windows 2000 Datacenter Server.
- Perform a clean installation of Windows Server 2003, Enterprise Edition and at the same time upgrade an existing stand-alone server to a server cluster.

**Upgrading a Windows NT 4.0 Server Cluster**

You cannot perform a rolling upgrade (upgrading the nodes sequentially) directly from Windows NT Server 4.0, Enterprise Edition to Windows Server 2003, Enterprise Edition. Instead, you have two options:

- Perform a nonrolling upgrade directly from Windows NT 4.0, Enterprise Edition with Service Pack 4 or later to Windows Server 2003, Enterprise Edition by taking the cluster resources offline and upgrading each node. With a nonrolling upgrade, you cannot maintain cluster availability during the upgrade.
- To maintain cluster availability, perform a rolling upgrade to Windows 2000 and then perform a second rolling upgrade to Windows Server 2003.

  To perform the rolling upgrade to Windows 2000, first move all of the cluster resources from the target node to other nodes. Install Windows 2000 Advanced Server or Windows 2000 Datacenter Server on the target node, and move the cluster resources back to the newly upgraded cluster node. Then install Windows 2000 Advanced Server or Windows 2000 Datacenter Server on the remaining nodes in the cluster.

  Complete the rolling upgrade by following the same process to install Windows Server 2003 on your cluster.

**Upgrading a Windows 2000 Server Cluster**

The easiest method for upgrading a Windows 2000 Server cluster is to perform a rolling upgrade. A rolling upgrade has two major advantages:

- Clients experience minimal interruption of service. (However, server response time might increase during the phases in which fewer nodes handle the work of the entire cluster.)
- You do not have to recreate your cluster configuration, because the configuration remains intact during the upgrade process.

**Upgrading an Existing Stand-Alone Server to a Cluster**

To upgrade an existing stand-alone print server to a print server cluster, use Print Migrator 3.0 to back up and restore the print server configurations to the newly created server cluster. You must create the cluster first, before adding the print spooler resource.

Additional hardware requirements exist for using the Windows Cluster service. For a list of supported cluster hardware, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources and search for "cluster."

**Note**

- Only Standard TCP/IP ports and LPR ports are supported on a server cluster.

For more information about clustering, see "Designing and Deploying Server Clusters" in this book.

# Determining the Print Server Network Connection Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After installing Windows Server 2003, choose one of the following three methods to connect your print servers to the network:

- Whenever possible, use the standard TCP/IP port, also know as the *standard port*. The standard port is the preferred printer port in Windows Server 2003.
- For older devices that do not function correctly with the standard port, use the Line Printer Remote (LPR) port.
- To connect to printers over the Internet or an intranet, use the Internet Printing Protocol (IPP).

When creating printer ports, Windows Server 2003 installs the associated port monitor for that port type.

**Connecting Through the Standard TCP/IP Port**
The standard port is designed for Windows Server 2003–based print servers that use TCP/IP to communicate with printing devices, including network-ready printers, network adapters such as Hewlett-Packard JetDirect, and external print servers such as Intel NetPort. The standard port uses the standard TCP/IP port monitor (TCPMon.dll), also known as the *standard port monitor*. The standard port monitor can support many printers on one server. It is faster and easier to configure than the LPR Port Monitor (Lprmon), and it provides rich status information and error condition reporting that eases administration of your print servers.

The standard port monitor uses Simple Network Management Protocol (SNMP) and the standard printer message information base as defined in RFC 1759, "Printer MIB." As a result, the standard port monitor provides much more detailed status information than does LPR when, for example, the printer generates an error indicating a paper jam or low toner, or a print job is not responding and must be purged.

Although Windows Server 2003 requires TCP/IP on the print server in order to communicate with the printer, clients do not need TCP/IP. Because clients communicate with the print server rather than with the printer, they can do so using any common network protocol.

For example, a Novell NetWare client might use the Internetwork Packet Exchange (IPX) protocol to send documents to a Windows Server 2003 print server. The server then passes the document to the printer using the TCP/IP protocol. Status and error information is relayed from the printer to the print server using TCP/IP, and from the print server to the client using IPX.

**Connecting Through the LPR Port**
The LPR port uses the LPR monitor **Lprmon**, which is provided with Print Services for UNIX. The LPR port acts as the client for printing according to guidelines defined in RFC 1179, "Line Printer Daemon Protocol." The LPR port spools to a target host enabled for Line Printer Daemon (LPD). The target is usually a computer running one of the following:

- UNIX
- Multiple Virtual Storage (MVS)
- Virtual Address Extension (VAX)/Virtual Memory System (VMS)

However, the target alternatively can be one of the following:

- A printer with a network adapter that supports LPD
- An external print server
- A computer running Windows XP Professional, Windows Server 2003, Windows 2000, or Windows NT
- A print server running Print Services for UNIX or another version of LPD.

Before you can create an LPR port, the client must have TCP/IP and Print Services for UNIX installed. The target server must be running a Berkeley Software Distribution (BSD)–compatible LPD utility that is installed according to RFC 1179. If the server is running Windows Server 2003, Windows 2000, or Windows NT 4.0, make sure that Print Services for UNIX is started.

**Using the Internet Printing Protocol (IPP)**

You can make Windows Server 2003 printing features accessible over the Internet. To install a printer over the Internet, use the printer's URL as the name of the printer. Within an intranet, you can alternatively use the printer's share name (for example, http://*servername/sharename*).

If you are connecting to a print server on the same intranet, it is recommended that users' intranet security settings for Microsoft® Internet Explorer be set to medium-low or lower. Having this setting throughout an intranet creates a true Point and Print connection, which offers many advantages over an IPP connection. You can view, connect to, and manage printers from any browser. On clients running Windows 95, Windows 98, or Windows Millennium Edition, you must install the Web Client in order to manage printers by using a browser. For clients running Windows NT 4.0, you can manage printers from a browser, but cannot connect to them for printing.

To configure a computer to print to a printer by using a URL address, perform the following procedure.

**To configure a computer to print to a printer by using a URL address**

1. Install Microsoft Internet Information Services (IIS) on a computer that is running Windows XP Professional

   or

   Install Microsoft Peer Web Services (PWS) on a computer running Windows 2000 Professional.

2. Install IIS on a print server that is running Windows Server 2003.

   The low-level protocol used for job submission is Internet Printing Protocol (IPP) version 1.0, which uses the Hypertext Transfer Protocol (HTTP) protocol as a carrier.

3. Ensure that all clients use Internet Explorer version 4.0 or later to connect to a printer.

For information about Internet Printing, see "Internet Printing" in Help and Support Center for Windows Server 2003.

# Planning for Print Server Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Planning for print server security is vital in order to protect your organization's resources. As with any production server, you need to protect the physical print server and safeguard access to data stored on the server. Consequently, your security plan must address three areas:

- Physical location
- Group Policy settings
- Printer permissions

**Ensuring the Physical Security of Each Print Server**

Locate your print servers in a physically secure location that only designated individuals can access. Allowing unauthorized access to your print servers risks harm to the system. In addition, consider to what extent you also need to restrict physical access to network hardware. The details of implementing these security measures depend on your physical facilities as well as your organization's structure and policies.

**Securing the Print Environment**

Windows Server 2003 adds new Group Policy settings that affect how clients connect to print servers on the network. Two of these policy settings are particularly useful for security.

**Allow print spooler to accept client connections** This Group Policy setting, which is configured on the server, determines how clients access the print server over the network. If an individual with administrative credentials creates shared printers for use by managed clients, the spooler automatically allows connections upon creation of the first shared printer. If a virtual spooler resource is created on a clustered server, the spooler likewise automatically allows connections. If no shared printers or virtual spooler resources already exist, you might need to enable this policy setting by using the Computer Management snap-in from a remote computer. To administer print services on a server running Windows Server 2003, log on to the server locally, or log on remotely through a Remote Desktop session.

**Point and Print restrictions** This Group Policy setting, which is configured on client computers, determines the print servers to which the client can connect. To provide a higher level of security for managed workstations, this policy setting controls a client computer's ability to connect to and install a printer driver from specified print servers. By default, managed clients can use Point and Print only with servers that are within their forest. An administrator can use this policy to add additional servers to the list of trusted print servers. Alternatively, administrators can disable this policy to enable managed clients to connect to any accessible print server and install a printer driver from it.

**Using Printer Permissions to Control Access to Shared Printers**

Even if the physical server is in a secure room, the print server might still be accessible through remote administration tools. Therefore, you need to implement methods for restricting access to remote administration of print servers. You can restrict access to a print server by setting printer permissions.

You might also want to restrict access to a particular network printer. For example, the graphics department in your organization might use a high-speed color printer regularly. The cost of supplies for a color printer is substantially higher than for a black-and-white printer, so you decide to deny access to anyone outside the graphics department. The best way to restrict access is to establish a new security group. You establish a group named Graphics Department, and grant access only to members of this group.

For a procedure describing how to restrict access to printers by using printer permissions, see "Restricting Access to Printers" later in this chapter.

# Testing the Print Server Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Testing is a key part of designing your print environment. Before you deploy your print servers in a production environment, make sure that your users can print successfully in a test environment. In most cases, thorough testing of your print servers during the design phase minimizes any impact on your users' ability to print.

It is recommended that you conduct testing in two stages. During the first stage, test your Windows Server 2003 deployment in a laboratory or controlled environment. After verifying that your design is successful in the lab, conduct a pilot deployment in a controlled real-world environment in which users perform their normal business tasks using the new printing features. This two-stage approach reduces your risk of encountering problems during full-scale deployment.

Before conducting the tests, make sure that the computers used in your pilot test run the same operating systems that you have in your production environment, and that the appropriate drivers are installed.

The following are some typical scenarios that you might want to test:

- If you list your printers in Active Directory, make sure that all the printers appear as choices when users try to add a network printer.

- Make sure that all users participating in the pilot test can print after connecting to each of the printers on the new servers.

- Have several users send various print jobs to the print server to measure how well the server holds up under a typical load.

- If you implement default printer location strings by using Group Policy objects, verify that these Group Policy settings are working as expected for the various groups to which they are assigned.

  For more information about Printer Location Strings, see the *Server Management Guide* of the *Windows Server 2003 Resource Kit* (or see the Server Management Guide on the Web at http://www.microsoft.com/reskit).

- If you use the standard port monitor with SNMP enabled, let the printer run out of paper to ensure that you receive an "out of paper" message when you send the next print job. For more information about standard port monitor, see "Determining the Print Server Network Connection Method" earlier in this chapter.

- If you set up a print server cluster, turn off one of the servers and see if the resources fail over to the other nodes. For more information about initiating a resource failure, see "Initiate a resource failure in Help and Support Center for Windows Server 2003.

# Deploying Print Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After your design team completes its work, the print server deployment team can implement the design for your new printing environment. Deploying print servers involves installing Windows Server 2003 on the servers, adding the servers to the network, and then adding the printers. Organizations that plan to deploy clustered print servers must take the additional step of configuring the clusters before creating the print spooler resource.

Figure 3.5 shows the tasks involved in deploying print servers.

**Figure 3.5 Deploying Print Servers**



# Installing Windows Server 2003 on Print Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The details of the print server deployment process vary depending on whether your design specifies a clean installation or an upgrade installation. If you are upgrading existing stand-alone or clustered servers, the upgrade process also differs depending on the version of the Windows operating system installed on your existing print servers.

# Deploying Print Servers with Clean Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you deploy print servers with clean installations of Windows Server 2003, you can decide whether to create new printers or migrate existing printers to the new servers. If you create new printers, you can enforce the standard printer naming convention that the design team developed during the design phase. If you have existing printers, it is a good idea to migrate them by using Print Migrator 3.0.

**Setting Up a New Print Server**

The first step in setting up a new print server is to install and configure Windows Server 2003. Identify each operating system that the printer server must support to determine any special selections that you must make while setting up your server. For example, if you support computers running Windows NT 4.0, you must install additional drivers during the upgrade process. After installing Windows Server 2003, you can begin adding printers.

**Migrating Printers by Using Print Migrator 3.0**

By using Print Migrator 3.0, you can back up your current print server configuration and restore the settings on a new print server, eliminating the need to manually recreate print queues and printer ports, install drivers, and assign IP addresses. Having to manually perform these tasks is a significant obstacle for most organizations. Using Print Migrator greatly reduces the manual intervention required when migrating your print servers from previous versions of Windows.

For example, an administrator obtains three new high-end servers to use as print servers. Rather than manually configure each server, the administrator uses Print Migrator 3.0 to back up three of the existing print servers and then restores the settings on the new servers. For more information about Print Migrator 3.0, see "Choosing a Print Server Migration Method" earlier in this chapter.

To ensure that your printer migration goes smoothly, read the documentation provided with Print Migrator 3.0.

# Deploying Print Servers with Upgrade Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploying an upgraded installation of Windows Server 2003 can be a simple process depending on the printer drivers installed on your servers. One of the biggest challenges of upgrading is resolving any printer driver problems. Before upgrading your servers, use the command-line utility Fixprnsv.exe, provided with Windows Server 2003, to help you identify any printer driver problems. If you are upgrading from Windows NT 4.0 or Windows 2000, consider the issues presented in the following sections.

**Important**

- As with any major software installation, it is recommended that you back up the hard disk before beginning an upgrade.

**Using Fixprnsv.exe to Resolve Driver Issues**

To manage driver-related issues, use the command-line utility Fixprnsv.exe. Fixprnsv.exe automatically replaces incompatible printer drivers or those with known problems. It locates existing printer drivers that can replace unsuitable drivers. In many cases, IHVs provide new printer drivers for this purpose. If replacement drivers are available, Fixprnsv.exe replaces problem drivers with Microsoft-provided drivers. If Fixprnsv.exe does not find a suitable replacement driver, it displays a message advising you to check the printer manufacturer's Web site for a newer version of the driver. Fixprnsv.exe installs drivers only for printers that are already configured on the print server.

The Fixprnsv.exe utility is located on the Windows Server 2003 installation CD-ROM in the \Windows\Printers\ directory. For a list of commands for use with Fixprnsv.exe, use the following procedure.

**To list the commands for Fixprnsv.exe**

1. Insert the Windows Server 2003 installation CD-ROM in the CD-ROM drive.

2. At the command prompt, change to the drive that contains the CD-ROM and type:

   **fixprnsv.exe /?**

**Upgrading from Windows 2000**

If you are upgrading a Windows 2000 print server that does not have Version 2 device drivers installed, your upgrade can be seamless. If your Windows 2000 print server has Version 2 drivers, you might encounter some of the issues involved in upgrading from Windows NT 4.0, which are discussed in the following section.

**Upgrading from Windows NT 4.0**

When upgrading from Windows NT 4.0 to Windows Server 2003, one or more of the following situations might occur:

- Windows NT 4.0 drivers that shipped with the operating system are upgraded to the new version of these drivers in Windows 2000, Windows XP, or Windows Server 2003. The Windows NT 4.0 driver remains as an additional driver.

- If the name of an IHV driver matches the name of a driver that ships with Windows 2000, Windows XP, or Windows Server 2003, the installation might upgrade the driver to the a Version 3 inbox driver (a driver that ships with Windows Server 2003). This upgrade occurs with no user intervention.

- Unstable drivers are blocked and are not carried through the upgrade. Printers that use these drivers are removed during the upgrade process. If the printer is connected directly to the computer and has a recognized Plug and Play ID, Windows searches for a suitable driver. If a suitable driver is found, the printer is installed during the upgrade as a new printer.

- Drivers that are not blocked and do not have name matches are carried through the upgrade unchanged.

**Windows NT 4.0 inbox drivers**

All Windows NT 4.0 Printer Control Language (PCL) drivers and Raster Device Drivers (RASDD) are upgraded to the latest Unidrv drivers as part of the upgrade process. The latest Unidrv drivers include UNIDRV5 for Windows 2000, and UNIDRV5.1

for Windows XP or Windows Server 2003. PostScript drivers that were shipped with Windows NT 4.0 are automatically upgraded to PostScript 5.0 in Windows 2000 or PostScript 5.2 in Windows XP or Windows Server 2003.

**Note**

- Drivers that you have installed as additional drivers for Point and Print on clients running Windows 95, Windows 98, or Windows Millennium Edition are not preserved during an upgrade from Windows NT 4.0 to Windows 2000 or Windows Server 2003. After you configure the print server, you must reinstall these additional drivers.

**IHV drivers with matching names**

IHV Windows NT 4.0 drivers are upgraded to Windows 2000, Windows XP, or Windows Server 2003 drivers with no user intervention if the driver name matches the existing Windows NT 4.0 driver name or if a newer version of the driver is available. In this situation, the IHV driver is treated exactly the same as the Microsoft driver.

Because of interaction problems between Windows 2000 drivers and Windows NT 4.0 drivers in Point and Print environments, many IHVs recommend that you reinstall their Windows NT 4.0 driver following the upgrade. This applies to users who plan to use Point and Print between Windows NT and Windows 2000, Windows XP, or Windows Server 2003.

**Blocked IHV drivers**

Microsoft designed the upgrade path from Windows NT 4.0 so that the new operating system replaces an old driver if a newer driver is available. Windows preserves a driver during an upgrade if the installation utility does not find a newer version and does not identify a driver as causing problems in Windows 2000, Windows XP, or Windows Server 2003. A driver is blocked for one of two reasons:

- Microsoft determines, through testing, that the driver causes substantial instability of the operating system.
- The IHV requests that the driver be blocked based on the IHV's own testing and available updates.

The system file Printupg.inf contains a list of known bad drivers. Drivers in this list might have an alternative inbox driver. If an alternative driver does not exist, the driver is not upgraded during the operating system installation. Instead, you need to install a newer version of the driver from the IHV. The information found in Printupg.inf can be helpful in identifying whether or not a new version from an IHV is blocked. The upgrade report also contains information about all drivers slated for removal during the upgrade.

In addition, the Fixprnsv.exe tool automatically replaces known bad drivers if inbox substitutes are available. If a compatible driver cannot be located, Fixprnsv.exe reports that fact and refers the user to the IHV's Web site. By running Fixprnsv.exe prior to upgrading, you can identify printers that are slated for deletion during the upgrade ahead of time. Otherwise, you might discover after the upgrade that some printers are unexpectedly missing.

**Non-blocked IHV drivers**

Certain Windows NT 4.0 drivers are not blocked and do not match the driver name of an inbox driver. These drivers proceed through the upgrade process without being altered or replaced. Use Fixprnsv.exe to identify drivers that must be replaced. For drivers that are not blocked (either because they are not in the Printupg.inf file or because they have a later date than the date of similar drivers listed in Printupg.inf), Fixprnsv.exe takes no action.

# Adding Printers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After installing Windows Server 2003 on the print server, you must add the printers to the server. Depending on the decisions that the design team made, adding printers can involve the following steps:

- Adding spooler resources to server clusters
- Installing printer ports
- Adding additional printer drivers
- Publishing printers in Active Directory
- Connecting clients to printers
- Restricting access to printers
- Setting up Internet printing

# Adding Spooler Resources to Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order to add printers to server clusters, you first need to create and configure print spooler resources. After creating the spooler resource, you can add printers, publish them in Active Directory, and add any additional drivers that your users require. You must also create and specify the spooler resource in order to use Print Migrator 3.0 to restore printers from another print server.

**Creating a Spooler Resource**
To create a spooler resource for a cluster, you must be the administrator of the cluster as well as of each node within the cluster. You must also have the administration software installed on your computer. Cluster Administrator is the graphical application that is supplied with the Cluster service to manage clusters. Alternatively, you can use Cluster.exe, a command-line tool, or develop custom administration tools by using the Cluster service command interfaces. Cluster Administrator is included in the Administrative Tools Pack (Adminpak.msi). The Cluster.exe tool and Adminpak.msi files are located in the Windows\System32 directory of a Windows 2003 Server.

The procedure for creating a print spooler resource is like the procedure for creating any resource in a cluster service. For more information about creating a cluster-managed printer, see "Create a cluster-managed printer in Help and Support Center for Windows Server 2003.

After you create the print spooler resource, migrate existing printers from existing print servers, if necessary, by using Print Migrator 3.0. For more information about Print Migrator, see "Choosing a Print Server Migration Method" earlier in this chapter.

**Configuring the Resource**
Configure the spooler resource by selecting the print spooler resource type in the Cluster Administrator window or by using Cluster.exe. Each cluster group can have only one print spooler resource, because resources are organized into interdependent groups that must fail over together.

After you create and configure the print spooler resource, you might need to install any third-party port monitors and print processors that are cluster-compatible on each node of the cluster. To do this, take each spooler resource offline, and then bring it online again so that the new resource is visible to the cluster. Then add printers to the clustered spooler. For more information about printing and the Cluster service, see "Print Spooler resource type" in Help and Support Center.

**Adding a Printer to a Cluster**
After creating a group and resources, you can add printers to the cluster. Each node must have connectivity to the remote print device. A printer that is locally connected to a node cannot be part of a cluster configuration, because the printer connects directly to the node and does not fail over if that node goes down.

Adding a printer to a cluster is very similar to adding the printer to any other computer, with the following exceptions:

- You never start from the local **Printers and Faxes** folder. The cluster always appears remote, even if you are working on the active cluster node. Instead type the virtual server name (for example, \\Virtual_Server_Prn) in the **Run** dialog box, and then click the remote **Printers and Faxes** folder that is displayed.

- If the Add Printer Wizard does not appear when you open the remote **Printers and Faxes** folder, you cannot continue. One of three things might be wrong:

    - The associated print spooler resource is not online or is not configured

    - You are not logged on as the administrator.

    - The spooler service is not started on the local computer. (This is unlikely.)

When a cluster group containing a print spooler resource fails over to another node, the document that is currently being sent to the printer is restarted on the other node after the failure. When you move a print spooler resource or take it offline, Cluster service waits until all documents are spooled or the configured wait time elapses. The cluster discards any documents that are submitted while the spooler resource is unavailable or the cluster is offline, and users must resubmit those print jobs.

For more information about clustering, see "Designing and Deploying Server Clusters" in this book.

# Installing Printer Ports

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your users print over a network, you must create printer ports to enable connections between print servers and printers. You can install printer ports for stand-alone servers and clusters from the Printers and Faxes folder on the print server. The printer port can be one of two port types:

- Standard TCP/IP port
- LPR port

If you are using clustered print servers, you need to install the printers on each node and ensure that each node has the appropriate protocols, port monitors (not ports), and print processors installed. You can do this remotely, but you must address each node by its node name, not the cluster name. Ensure that the settings on all nodes are identical, because a specification made on one node does not carry over to the other nodes.

**Standard TCP/IP Port**

The standard port is the preferred printer port in Windows Server 2003. The standard port uses the standard TCP/IP port monitor (standard port monitor) and is designed for Windows Server 2003–based print servers that communicate with printers by using TCP/IP. For more information about the advantages of using the standard port and the prerequisites for installing a standard port, see "Determining the Print Server Network Connection Method" earlier in this chapter.

**To install a standard port**

1. In Control Panel, open the **Printers and Faxes** folder, right-click the printer that you want to configure, and click **Properties**.
2. On the **Ports** tab, click **Add Port**.
3. In the **Printer Ports** dialog box, select **Standard TCP/IP port**, and click **New Port** to launch the **Add Standard TCP/IP Printer Port Wizard**.
4. Complete the **Add Standard TCP/IP Printer Port Wizard by** using the information provided in Table 3.3.

   **Table 3.3 Using the Add Standard TCP/IP Printer Port Wizard**

   | Wizard Page | Action |
   |---|---|
   | Printer Name or IP Address | To identify the printer that will be connected to the port, type its name in the **Printer Name** or its IP address in the **IP Address** box. |
   | Port Name | In the **Port Name** box, type a port name, which can be any character string, or use the default name that the wizard supplies. |
   | Additional Port Information Required | To configure a standard port, click **Standard**, and then select one of the listed devices. If you do not know the details of the port, try using the **Generic Network Card**.<br><br>To create a custom configuration, click **Custom**, and then configure the port by using the **Configure Standard TCP/IP Port Monitor** screen that appears. |

   If the wizard cannot determine the appropriate protocol for the port, it prompts you for the information. Follow the vendor's instructions for selecting either the RAW or LPR option.

   If you are not prompted for more information, continue to step 5.

5. Review the port information, and click **Finish**.

The new port is listed on the **Ports** tab of the **Properties** property sheet.

**Note**

- With Windows Server 2003, administrators can remotely configure and manage ports from any server running Windows Server 2003. This feature applies to local ports, the standard TCP/IP port, and LPR ports. You must configure AppleTalk ports locally on the server.

You can reconfigure the standard port monitor by adjust the settings in the property sheet for the print server.

**Caution**

- The Configure Port dialog box does not validate the settings created in the following procedure. If they are incorrect, the port no longer works.

**To reconfigure the standard TCP/IP port**

1. In Control Panel, open the **Printers and Faxes** folder, right-click the appropriate printer, and click **Properties**.
2. On the **Ports** tab, click the **Configure Port** button.
3. In the **Configure Standard TCP/IP Port Monitor** dialog box, click either the RAW or LPR protocol.
4. To configure the protocol, take one of the following actions:
   - For **RAW Settings**, type the port number that the printer vendor specified (usually 9100).
   - For **LPR Settings**, type the LPR queue name that the printer vendor specified.
5. If the printer supports SNMP and RFC 1759, select the **SNMP Status Enabled** check box.
6. If **SNMP Status Enabled** is selected, you can change both the SNMP community name and the host device index:
   - The community name is usually "Public," but you can enter another community name if you want to limit access to the printer.
   - The device index is used mainly for multiport devices that support several printers; each port on a multiport device has a different device index, specified by the device vendor.

**LPR Port**

If your network includes clients that are running UNIX, install LPD, which acts as the client for printing according to RFC 1179 guidelines. For more information about the LPR port, see "Determining the Print Server Network Connection Method" earlier in this chapter.

**To install an LPR port**

1. In Control Panel, open the **Printers and Faxes** folder.
2. Under Printer Tasks, click **Add a printer** to open the **Add Printer Wizard**. Then click **Next**.
3. Complete the **Add Printer Wizard** byusing the information provided in Table 3.4.

**Table 3.4 Steps for Completing the Add Printer Wizard**

| Wizard Page | Action |
|---|---|
| Local or Network Printer | Click **Local printer attached to this computer**, clear the **Automatically detect and install my Plug and Play printer** check box, and then click **Next**. |
| Select a Printer Port | Click **Create a new port**, and then select **LPR Port**. |
| Add LPR compatible printer | If **LPR Port** is not available, click **Cancel** to stop the Wizard. To add the LPR Port, install the optional networking component, **Print Services for UNIX.** |
| Name or address of server providing LPD | Type the Domain Name System (DNS) name or Internet Protocol (IP) address of the host for the printer that you are adding. The host might be the direct-connect TCP/IP printing device or the UNIX computer to which the printing device is connected. The DNS name can be the name specified for the host in the Hosts file. |
| Name of printer or print queue on that server | Type the name of the printer as it is identified by the host, which is either the direct-connect printer itself or the UNIX computer. |

**Configuring LPR to print text files to a PostScript printer in UNIX**

In some cases, sending an ASCII text file to a PostScript printer on a UNIX-based computer can cause PostScript code to be output on the printer. This can occur because the PostScript printer is processing the text as RAW data without interpreting the PostScript code.

LPR sends a processing instruction in each print job in the form of a control command indicating the data format. Table 3.5 shows the LPR commands for specific data types. Under **Lprmon**, the default is **l** (RAW data); under Lpr.exe, the default is **f** (text).

**Table 3.5 Control Commands for LPR Printing**

| LPR Control Command | Data Format Transmitted |
|---|---|
| f, p | Text data type |
| L | RAW data type |
| O | RAW data type formatted for a PostScript printer |

Using the default value of **l** can cause the PostScript code to be printed on a UNIX-based system. The **l** value sets the RAW data type, which causes the text file to be printed ignoring the PostScript instructions. To correct this, set **0** (RAW data type formatted for a PostScript printer) as the default value. This might not be necessary on some UNIX systems, which include software that scans arriving documents for PostScript code with the **l** value. If the software detects the **l** value, the document goes directly to the printer; otherwise, the software adds PostScript code. You can change the default control command that LPR sends by editing the **PrintSwitch** entry in the registry.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD, or at http://www.microsoft.com/reskit.

**To change the default control command for a printer**

1. In the registry editor, navigate to:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Print \Monitors\LPRPort\Ports\Port-name\IP Address or Host name: Printer-name

2. Double-click **PrintSwitch**.

3. In the **Data type** box, type REG_SZ, and then in the **Value Data** box, type the control command you want to specify. (See Table 3.5 for a list of control commands and the value data types that they represent.)

**Respooling**

*Respooling* is a method of spooling a document twice to enable the print monitor to locate the appropriate LPD print server. LPR must include an accurate byte count in the control file, but it cannot get the byte count from the local print provider. After Lprmon receives a document from the local print provider, it spools the document a second time to a temporary file in the default spool directory, finds the size of that file, and then sends the size to the LPD print server.

**Status reporting**

LPR is a one-way protocol and does not return a detailed error status report. If a problem occurs, the message is always ERROR. For more information about LPR or LPR errors, see the Microsoft Knowledge Base link on the Web Resources page at www.Microsoft.com/windows/reskits/webresources. Search the Knowledge Base using the keywords "LPR" or "LPR Errors".

**Pooling printers**

Printer pooling is especially useful in high-volume printing environments. Pooled printers appear to clients as a single printer, but printing throughput is increased because the load is distributed among the printers in the pool.

Before you set up a printer pool, consider the following issues:

- Two or more printers are required; Windows Server 2003 does not limit the number of printers in a pool.
- The printers in the pool must be of the same model, and they must use the same printer driver.
- Printer ports can be of the same type or mixed (such as parallel, serial, and network).
- If you want to ensure that documents are first sent to the faster printers, add the faster printers to the pool first and the slower printers last. Print jobs are routed in the order in which you create the ports.
- Because users do not know which printer prints their documents, it is a good idea to locate all of the pooled printers in the same physical location. Otherwise, users might not be able to find their printed documents.

**To create a printer pool**

1. In Control Panel, open the **Printers and Faxes** folder, right-click the appropriate printer, and click **Properties**.
2. On the **Ports** tab, select the **Enable printer pooling** check box.
3. In the list of ports, select the check boxes for the ports connected to the printers that you want to pool.
4. Repeat steps 2 and 3 for each additional printer to be included in the printer pool.

# Adding Additional Printer Drivers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By adding additional drivers to your print server, you can support clients that are running various versions of the Windows operating systems. It is preferable to use Version 3 drivers rather than Version 2 drivers as your primary drivers, because Version 2 drivers can cause the system to become unstable or stop responding. However, you might need to use Version 2 drivers as additional drivers if you support clients running Windows NT 4.0. If your server has known bad drivers when it is upgraded, those drivers are deleted during the upgrade process. Any additional drivers that are added must have the same name as the primary driver unless they are Windows 95, Windows 98, or Windows Millennium Edition drivers.

**Note**

- With Windows Server 2003, you no longer need to install printer drivers on each node of a cluster. The drivers automatically propagate to the other nodes of the cluster after they are installed on the virtual server.

# Publishing Printers in Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Publishing your shared printers in Active Directory can make locating printers across a network more efficient for your users. In Windows Server 2003, the print subsystem is tightly integrated with Active Directory, making it possible to search across a domain for printers at different locations. By using the standard printer object that Windows Server 2003 provides, you enable users to search for printer-based attributes such as printing capabilities (including PostScript, color, and paper size) and printer locations (allowing users to find printers located near them).

If you plan to publish printers to Active Directory, follow the naming standard that the print services design team created when you fill out the location fields on the printer properties pages across your network. This enables users to enter a standard string to search for printers by location. If you use subnets to define the sites within your organization, Active Directory can find nearby printers – a process that has formerly been difficult for both administrators and users.

For example, if you are in Los Angeles and want to find all the Los Angeles printers in your deployment, search for a network printer by using the Add Printer Wizard and type **US/LAX** in the **Location** dialog box. If **US/LAX** matches the printer location syntax in Active Directory, your search might return the following results:

- US/LAX/1/101
- US/LAX/2/103

These results indicate that two printers are available in Los Angeles, located in buildings 1 and 2, in rooms 101 and 103.

Printer publishing is controlled by the **List in the Directory** check box on the **Sharing** tab of the **Properties** sheet for each printer. Printers that are added by using the Add Printer Wizard are published by default, and the wizard does not let you change this setting. If you do not want to publish a printer in Active Directory, after adding the printer, open the printer property sheet, and clear the **List in the directory** check box on the **Sharing** tab.

The printer is placed in the **Computer** object in Active Directory on the print server. After publishing the printer in Active Directory, you can move or rename the object by using the Users and Computers snap-in in Active Directory.

**To publish a printer in Active Directory**

1. In Control Panel, open the **Printers and Faxes** folder.
2. Right-click the icon for the printer that you want to publish in Active Directory, and click **Properties**.
3. On the **Sharing** tab, select the **List in the directory** check box.

# Connecting Clients to Printers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Generally, clients can establish a connection to a shared network printer hosted on a Windows Server 2003 print server in one of three ways:

- **Entering the Universal Naming Convention (UNC) path.** In the **Run** dialog box, type the UNC path for the printer (for example, \\PrintServer\Printer), and then click **OK**.

- **Using the Add Printer Wizard.** From the **Printer and Faxes** folder, select **Add a Printer**. When the Add Printer Wizard opens, click **Next**, select the network printer option, and either type or browse to the path for the shared printer.

- **Using drag-and-drop.** In the **Run** dialog box, open the shared printers folder on the remote print server. Then drag the desired printer icon into the **Printers and Faxes** folder on the local workstation.

# Restricting Access to Printers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your design calls for restricting access to certain printers, you can do so by using printer permissions. To do this, it is recommended that you create a user group and then limit access to a printer to members of the group.

**To create a local group**

1. Right-click **My Computer**, and click **Manage**.
2. In the console tree, double-click **Local Users and Groups**.
3. To add the group, right-click **Groups**, and click **New Group**.
4. In the **Group name** box, type the name of the new local group.
5. To add users to the group, click **Add**, and enter the user names in the **Enter the object names to select** box.

**To restrict access to a specific printer**

1. In Control Panel**,** open the **Printers and Faxes** folder.
2. Right-click the icon for the printer to which you want to restrict access, and click **Properties**.
3. To remove extraneous group members, on the **Security** tab, remove all entries in the **Groups or user names** list box except Administrator and Creator Owner.
4. To grant access to the printer, click **Add**, and then enter the names of the groups and users that you want to grant access to this printer.

# Enabling Internet Printing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using Internet printing, you can manage print resources from your Web browser. To be able to print over the Internet, clients within the same Local Area Network (LAN) must use a remote procedure call (RPC) to connect to the print server. For more information about prerequisites for Internet printing, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Installing a Printer from a Web Page**

To install a printer for Internet printing, you can either view a Web page to find a printer that is identified by a URL, or — if the client is running Windows Server 2003, Windows XP Professional, Windows 2000, Windows 95, Windows 98, or Windows Millennium Edition — connect to a printer share through a Web page.

Open a web browser, and type: http*://servername*/printers. Under **Printer Actions,** click **Connect**. The installation page displays available options based on your permissions. Windows Server 2003 downloads the printer software to the client, and the printer is displayed in the **Printers and Faxes** folder on the client.

The installation route depends on whether the client and the print server are on the same intranet and are both running Windows XP Professional, Windows Server 2003, Windows 2000, or Windows NT 4.0. If they are, the client and print server communicate by means of an RPC, and the installed printer continues to use an RPC to link the client and the server even if HTTP is not specified in the address.

The installation uses HTTP instead of RPC in the following instances:

- The client and server are not on the same intranet.
- The client is not running Windows Server 2003, Windows XP Professional, Windows 2000, or Windows NT 4.0.
- The printer contains an internal network adapter, supports Internet Printing Protocol 1.0, and is not connected to a server.

With HTTP, the print server generates a .cab file containing the required .inf and installation files and sends the .cab file to the client. On the client computer, the .cab file starts the Add Printer Wizard to complete the installation. A progress report is displayed in HTML while the wizard is working.

**Important**

- Installation is not automatic for Web-based printers with internal network adapters. You must start the Add Printer Wizard, enter the printer's URL instead of a UNC path, and manually enter information that the wizard requires. You can use this method to install any URL-identified printer by means of HTTP.

**Security for Internet Printing**

Print server security is provided by IIS, which runs on the print server. IIS allows basic authentication, which all browsers support. The administrator must select basic authentication to enable the print server to support all browsers and all Internet clients. IIS and PWS allow the use of Integrated Windows authentication and Kerberos authentication, both of which are supported by Internet Explorer.

The authentication method for Internet printing in IIS or PWS is set in the print server's property sheet on the **Directory Security** tab.

By default, print jobs are sent over HTTP as RAW data. If it is important to keep this data secure, use either a Virtual Private Network (VPN) or Secure Socket Layer (SSL) connection.

**To select an authentication method**

1. In the console tree of the IIS console, expand the node for the server, expand the Web Sites node, expand the Default Web Site node, and then expand the Printers node.
2. Click the icon at the Printers node.

This node represents a virtual directory that is used to set all security for Internet printing. A list of Application Server Pages (ASP) appears in the details pane.

3. In the console tree, right-click the printer, click **Properties**, and then click the **Directory Security** tab.

4. Choose one of the following Directory Security options by clicking the respective **Edit** button:

   - **Authentication and access control**
   - **IP address and domain name restrictions**
   - **Secure communications**

Typically, administrators select **Enable anonymous access**,which allows a client to access each server resource by impersonating the Anonymous account IUSR_*computername*. No user action is required. If a user attempts to connect to another domain or proxy server that does not allow anonymous access, a dialog box prompts for the user name and password.

**To choose anonymous access authentication**

1. On the **Directory Security** tab of the **Printers Properties** page, click the **Edit** button for **Authentication and access control**.

2. Select the **Enable Anonymous access** check box.

3. Clear the **Windows Integrated authentication** check box.

**Note**

- **Integrated Windows authentication** is checked by default and takes precedence over other types of authentication. To ensure that users are authenticated anonymously, clear all check boxes except **Enable anonymous access**.

Integrated Windows authentication is more secure, because it does not send the password. During Integrated Windows authentication, IIS applies either challenge and response encryption technology, or Kerberos encryption technology, depending on the capability of the client. For more information about IIS security, see the *Internet Information Services (IIS) 6.0 Resource Guide* of the *Windows Server 2003 Resource Kit* (or see the Internet Information Services (IIS) 6.0 Resource Guide on the Web at http://www.microsoft.com/reskit).

# Additional Resources for Designing and Deploying Print Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Related Information**

- The *Server Management Guide* of the *Windows Server 2003 Resource Kit* (or see the Server Management Guide on the Web at http://www.microsoft.com/reskit).
- The at http://www.microsoft.com/reskit, or see the MSDN Scripting Clinic link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources, for more information about scripting.
- RFC 1759: *Printer MIB* for more information about Printer Message Information Base.
- RFC 1179: *Line Printer Daemon Protocol* for more information about using an LPR port that uses the Line Printer Daemon Protocol.

**Related Help topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Initiate a resource failure in Help and Support Center for Windows Server 2003 for more information about initiating a resource failure.
- "Create a cluster-managed printer in Help and Support Center for Windows Server 2003 for more information about creating a cluster-managed printer.
- "Print Spooler resource type" in Help and Support Center or more information about printing and the Cluster service.
- "Windows Server 2003 Administration Tools Pack Overview in Help and Support Center for Windows Server 2003 for more information about the Administration Tools Pack.
- "Managing printing from the command line" in Help and Support Center for Windows Server 2003 for more information about managing printers from the command line.
- "Internet Printing" in Help and Support Center for Windows Server 2003 for more information about the Internet Printing.

**Related Job Aids**

- "Print Server Inventory Worksheet" (Sdcpsv_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Print Server Inventory Worksheet" on the Web at http://www.microsoft.com/reskit).
- "Print Server Design Worksheet" (Sdcpsv_2.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Print Server Design Worksheet" on the Web at http://www.microsoft.com/reskit).

**Related Tools**

- Print Migrator 3.0 (Printmig.exe)

  Use Printmig.exe to help automate printer consolidation on Windows Server 2003. Instead of recreating all of the printer objects manually, you can back up existing printers and restore them to your new print servers. To download Print Migrator, see the Print Services link on the Web Resources page at

http://www.microsoft.com/windows/reskits/webresources.

- Fixprnsv.exe

  Use Fixprnsv.exe to help resolve printer driver compatibility issues between Windows NT 4.0, Windows 2000 and Windows Server 2003. Fixprnsv.exe is located on the Windows Server 2003 installation CD-ROM. The executable file is located in Windows\Printers directory. To view a list of commands used with this utility, use the command **fixprnsv.exe /?** from that directory.

- PrnAdmin

  Use PrnAdmin to help migrate from Windows NT 4.0 to Windows Server 2003. You can use Prnadmin.dll to inventory your current environment. PrnAdmin enumerates ports, drivers, printers, or forms on local or remote printers. The Visual Basic script for the PrnAdmin utility is named Prnmgr.vbs. PrnAdmin is located on the *Windows Server 2003 Deployment Kit* companion CD. To install PrnAdmin, type the following command:

  **regsvr32 "C:\Program Files\Resource Kit\PrnAdmin.dll"**

# Hosting Applications with Terminal Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With Terminal Server you can host applications in a central location, which saves management and deployment time and increases security because you install the applications once on a server, rather than on every desktop.

**In This Chapter**


Overview of the Terminal Server Deployment Process

Planning for Application Hosting

Planning Network and Hardware Resources

Designing Domain and User Configuration

Designing Terminal Server Installation and Configuration

Designing the Remote Desktop Connection Installation and Configuration

Additional Resources for Hosting Applications with Terminal Server


**Related Information**

- For information about using Terminal Services and other remote management tools for remote administration, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

- For more information about Remote Desktop, see "Configuring Remote Desktop" in *Microsoft® Windows® XP Professional Resource Kit Documentation* (or see "Configuring Remote Desktop" on the Web at http://www.microsoft.com/reskit).

# Overview of the Terminal Server Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Hosting applications from a central location with Terminal Server enables centralized management of those applications. You also gain added security because data is stored on the server, rather than on local computers. The following examples illustrate the kinds of business solutions that Terminal Server can provide.

**Bandwidth-constrained locations** In areas where high bandwidth is not available or cost-effective, deploying applications on a terminal server can improve performance for remote users because large amounts of data are not being transmitted over the connection.

**Applications in development or transition** You can simplify administration of crucial line-of-business applications that are in development or require frequent updating by deploying these applications on a terminal server. You can ensure that all users have access to the latest version by updating the application on the server rather than on individual client computers, reducing the total number of updates required.

**Use of third-party platforms** You can reduce desktop hardware costs for users who have a primary line-of-business application that is available only on a desktop running an operating system other than a Microsoft® Windows® operating system, but who still require access to Windows-based applications. Users can connect to a terminal server to run the Windows-based applications.

**Use of thin clients** Data-entry workers can access their primary application from a Windows-based terminal (or thin client) rather than a personal computer, reducing total cost of ownership. Using thin clients also minimizes work disruption for task workers. If a terminal stops running, you can replace it quickly, with minimal setup.

For more details about the solutions that Terminal Server provides, see "Identifying the Role of Terminal Server in Your Organization" later in this chapter.

Before you plan your deployment of Terminal Server, you need to design your domain infrastructure. After you complete the tasks outlined in this chapter, you can conduct a pilot test of your Terminal Server solution, adjust your design if necessary, and deploy the completed solution to the production environment. For more information about pilot testing, see Planning for Deployment in *Planning. Testing, and Piloting Deployment Projects* of this kit. For more information about implementing your Terminal Server solution, see "Terminal server role: Configuring a terminal server" in the Configure Your Server Wizard and "Checklists: Setting up Terminal Server" in Help and Support Center for the Microsoft® Windows® Server 2003 operating system.

Terminal Server is available with the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems.

# Process for Planning and Deploying Terminal Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 4.1 illustrates the high-level process for planning and deploying Terminal Server for hosting applications. Although the primary tasks remain the same for anyone planning to deploy Terminal Server, the steps differ depending on the solution you want to provide. Those differences are described in context throughout the chapter.

**Figure 4.1 Planning and Deploying Terminal Server**




For a job aid to assist you in recording your design decisions, see "Terminal Server Planning and Design Worksheet" (SDCTS_1.xls) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Terminal Server Planning and Design Worksheet" on the Web at http://www.microsoft.com/reskit). For a job aid to assist you in recording your Terminal Server Group Policy configuration decisions, see "Group Policy Configuration Worksheet" (SDCTS_2.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Group Policy Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

# Basic Terminal Server Concepts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Terminal Server component of Windows Server 2003, which was called Terminal Services in Application Server Mode in the Microsoft® Windows® 2000 operating system, can deliver the Windows desktop, in addition to Windows-based applications, from a centralized server to virtually any desktop computing device, including those that cannot run Windows.

Terminal Server uses the Remote Desktop Protocol (RDP) to communicate between client and server. After you deploy an application on a terminal server, clients can connect over a remote access connection, local area network (LAN), wide area network (WAN), or the Internet. The client computers can run Windows (including the Microsoft® Windows® CE operating system) or run on other operating systems such as the Apple Macintosh or even UNIX (using a third-party add-on). When a user accesses an application on a terminal server that is running Windows Server 2003, all of the work of the application takes place on the server and only the keyboard, mouse, and display information are transmitted over the network to the user desktop. Each user sees only their individual session, which is managed transparently by the server operating system and is independent of any other client session.

Windows Server 2003 Terminal Server offers the following new features and improvements:

- For improved load-balancing performance, the new Session Directory feature automatically reconnects users to an existing session within a load-balanced server farm, rather than just being directed to the least loaded server when they reconnect.

- For improved logon security, you can use a smart card to log on to a terminal server running Windows Server 2003 and run applications from client computers with a smart card subsystem (including computers running the Microsoft® Windows® XP operating system, Windows 2000, and the Windows CE operating system). This provides an additional level of physical security to your network environment.

- For a better user experience, Terminal Services supports a wider variety of data redirection types (including file system, serial port, printer, audio, and time zone) and supports connections in up to 24-bit color.

- For improved management, there is an expanded set of Group Policy settings for Terminal Server and a full Windows Management Instrumentation (WMI) provider allowing for scripted configuration of Terminal Server settings.

- For improved license management, you can install Terminal Services Enterprise Licensing on either a domain controller or a member server. Terminal Services Licensing wizards have been improved to reduce the complexity of activating a terminal server and assigning it licenses. There is also a new security group that allows only specific terminal servers to request licenses from a license server in a domain.

For more information about improvements in Windows Server 2003 Terminal Server, see the "What's New in Terminal Server" white paper at the Terminal Services: Community Center link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources and Using Terminal Server to host applications centrally in Help and Support Center for Windows Server 2003.

# Planning for Application Hosting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As shown in Figure 4.2, the first step in planning your deployment of Windows Server 2003 Terminal Server is to clearly identify the scope of your deployment. This step provides you with a high-level specification on which to base the design of your network and domain infrastructure.

**Figure 4.2 Planning for Application Hosting with Terminal Server**



# Identifying the Role of Terminal Server in Your Organization

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The way in which you plan to use Terminal Server has an effect on how you deploy it. The following sections outline the various ways in which you might use Terminal Server in your organization. Use the information in these sections to clarify how you are going to use Terminal Server in your organization.

**Hosting Line-of-Business Applications**

If your organization or certain groups within your organization use specialized applications to do their work, it might be beneficial to host the applications with Terminal Server. For example, you might decide to use Terminal Server in the following situations:

- **Custom applications.** If your line-of-business application is developed in-house or especially for your organization, and it tends to require frequent updating or repair, deploying the application once on a terminal server can simplify administration of the application. This is especially useful if your environment is geographically dispersed or if you are deploying Terminal Server to centrally serve your organization's branch offices.

- **Large central data pool.** Applications that rely on access to a single data source often run better on a terminal server because large amounts of data do not have to travel across the network to users. Instead, the data processing takes place on the server. Only the keystrokes and display information travel across the network, so you can use lower bandwidth connections. This is especially useful if the users of the data pool are located remotely, for example in a branch office with a slow connection to the database server.

- **Task workers.** In environments where you want workers to access only the application that they need to perform their jobs, you can centralize the administration of these users by using Terminal Server.

- **IT Admin tools.** System administrators can perform tasks that require Enterprise Admin or Domain Admin permissions on a terminal server that has the necessary tools installed on it. They can also run their desktop applications on their local computer without these permissions.

- **Upgrading operating systems.** If your organization uses a line-of-business application that is not optimized for your desktop operating system, you can host the application with Terminal Server rather than change operating systems.

**Hosting the Desktop**

You can use Terminal Server to host users' entire desktop environments, so that when users log on, they see their usual desktop environments or desktop environments especially designed for their remote use. In this situation, users can open and close the applications they choose in the same way that they access applications from the Windows desktop on the local computer. You can host the desktop in the situations discussed in the following sections.

**Remote users**

Hosting the desktop with Terminal Server can provide higher levels of consistency and performance for users in remote locations because large amounts of application data are not being transmitted over the connection. For example, you might use Terminal Server in the following situations:

- **Bandwidth-constrained locations.** In areas where high bandwidth is not available or cost-effective, deploying applications on a terminal server can improve performance for users who are connecting to the network from a remote location over slow dial-up connections.

- **Mobile users .** For users who are traveling and tend to access the corporate network over connections of varying bandwidth, Terminal Server can provide a more consistent experience.

**Client heterogeneity**

If your organization is in the process of converting all users to the same platform or upgrading desktop hardware, you can use Terminal Server to quickly deliver the most up-to-date version of the operating system and applications to the user while enabling you to spread the desktop platform or hardware conversion over a longer period of time. You can also deliver a highly controlled standard desktop to users by using Terminal Server, as illustrated in more detail in the following list:

- **Mixed-platform environment.** If you have users who require applications based on operating systems other than Windows to perform their jobs, but your organization is transitioning to or requires a Windows-based desktop, you can host the desktop with Terminal Server. This requires the use of third-party software in conjunction with Terminal Server.

- **Upgrading hardware.** If your organization is planning to upgrade to Windows XP on the desktop, but not all of the desktop hardware is compatible, you can use Terminal Server to host the desktops of users who have older hardware while you are in the process of upgrading the hardware. All users can have the same desktop environment and run the latest versions of the applications designed for Windows XP regardless of their desktop hardware.

- **Highly controlled environments.** In situations where you want to deliver a standardized and controlled environment to users, you can host the desktop with Terminal Server to centralize management.

**Hardware Considerations**

You can reduce hardware costs by hosting applications with Terminal Server and using thin client devices or older hardware on your users' desktops.

- **Use of thin clients.** Windows Powered thin clients (sometimes called Windows-based terminals) offer an alternative to personal computers and traditional "green screen" terminals by enabling easy remote access to productivity and line-of-business applications that are hosted on Windows-based terminal servers.

- **Extending the life of older hardware.** Rather than replacing older hardware that is no longer capable of running new Windows-based applications, you can use that older hardware much like a thin client to access the desktop and applications on the server rather than on the local computer.

**Using the Remote Desktop Web Connection**

The Remote Desktop Web Connection is an ActiveX control that provides virtually the same functionality as the executable version of Remote Desktop Connection, but it delivers this functionality over the Web even if the executable version is not installed on the client computer. When hosted in a Web page, the ActiveX Client Control allows a user to log on to a terminal server through a TCP/IP Internet or intranet connection and view a Windows desktop inside Internet Explorer.

The Remote Desktop Web Connection provides an easy way to offer Terminal Server through a URL. Consider using the Remote Desktop Web Connection in the following situations:

- **Roaming users.** Users who are away from their computers can use Remote Desktop Web Connection to gain secure access to their primary workstations from any computer running Windows and Internet Explorer, provided they can reach the target computers.

- **Delivery of extranet applications**. You can use Remote Desktop Web Connection to allow business partners or customers access to internal applications over the Internet. Users who gain access in this manner do not need to reconfigure their computers, and they do not gain access to your internal network.

- **Deployment transition.** You can deploy the Remote Desktop Web Connection quickly and use it while you are deploying your full Remote Desktop Connection infrastructure.

For more information about the Remote Desktop Web Connection, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit)..

# Choosing Applications to Host

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before deploying applications with Terminal Server, you must determine which applications are ideal candidates for hosting in your organization, and determine if those applications will work well in a Terminal Server environment. You might also want to host the entire desktop through Terminal Server.

# Identifying Ideal Candidates for Hosting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When choosing applications to host with Terminal Server, consider the compatibility of those applications with Terminal Server. Many desktop applications load and run correctly with Terminal Server. Applications that have been certified for Windows are generally compatible with Terminal Server. For more information about the Certified for Windows program, see the Certified for Windows link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

With any Windows-based application the definitive source of compatibility information should always be the application's vendor or developer. For more information about application compatibility with Terminal Services, see "Program Compatibility" in Help and Support Center for Windows Server 2003. There are also application compatibility notes for some applications available on the VeriTest Catalog of Certified Windows Server Applications link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Application Characteristics That Affect Performance with Terminal Server**
Because Terminal Server is available for multisession (multiuser) use, and because the display and keystroke information travels over the network, applications that have certain characteristics might perform poorly with Terminal Server. When planning for your organization, establish a set of acceptable performance guidelines and determine through testing whether such applications run better on the user's local computer. For example:

- Multimedia applications or applications that have very large graphical output do not run well with Terminal Server. Because large amounts of display information travel over the network, the performance of these types of applications is often unacceptable, especially over low-bandwidth connections. Many computer-aided design (CAD) and streaming media applications fall into this category.

- Running 16-bit programs with Terminal Server can reduce the number of users a processor supports and increases the memory required for each user. Windows Server 2003 translates 16-bit programs in enhanced mode through a process called Windows on Windows (WOW). This causes 16-bit programs to consume additional system resources. A similar situation exists with running 32-bit applications on the 64-bit versions of Windows Server 2003.

  **Note**

  - You can prevent 16-bit applications from running on Windows Server 2003 by enabling the **Prevent access to 16-bit applications** Group Policy setting. This setting is located in Administrative Templates/Windows Components/Application Compatibility. You can set this setting for both the computer and the user.

  There are third-party add-ons that you can use to help with the performance of 16-bit applications and other application compatibility issues. For more information about these add-ons, see the AppSense and Tame Software links on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Applications that are based on the Microsoft® MS-DOS® operating system use keyboard polling that can consume all the available resources on a CPU.

- Applications with features that run continuously in the background (for example automatic spelling and grammar checking in Microsoft® Word) can affect performance with Terminal Server. In some cases, you can restrict user or group access to certain application types, disable unnecessary features that require the most resources, or install applications on separate servers to minimize the performance effects.

- Complex automation or macros can be CPU-intensive and are not recommended. If you use or plan to use macros with the applications you plan to host, test your plans thoroughly to ensure adequate performance.

- If you plan to use Terminal Server to host custom applications or applications developed in-house, you need to consider the specification for applications for the Certified for Windows program. For more information about the Certified for Windows program, see the Certified for Windows link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. For information about guidelines for developing applications for use with Terminal Server, see the Optimizing Applications for Windows 2000 Terminal Services link on the Web

Resources page at http://www.microsoft.com/windows/reskits/webresources.

- If you plan to host older applications with Windows Server 2003 Terminal Server, for example applications that you have hosted with the Microsoft® Windows® NT 4.0 Terminal Server Edition operating system, be aware that you might need to run Terminal Server in Relaxed Security mode, which allows access to system files and the registry on the server, in order for these applications to function properly. For more information about security modes, see "Designing the Terminal Server Configuration" later in this chapter.

After you choose the applications you want to host, analyze these applications for potential problems in a multiuser environment and thoroughly test them, including end-user testing and pilot testing, to ensure that they work in your environment. For more information about piloting and testing, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

**Compatibility Scripts**

You might have to modify or provide support for custom applications or applications that are not written to accommodate multiuser access. Such applications might not use per user data storage, the user-specific portions of the registry, or appropriate permissions. You can run compatibility scripts for these applications after you install the application. These scripts are available:

- **With the Windows Server 2003 operating system**. These scripts are located in the *systemroot*\Application Compatibility Scripts\Install folder.

- **From your application vendor**. These scripts might be available on the application CD, on the Web, or by special order from the company.

- **By developing them yourself**. Use the scripts in the *systemroot*\Application Compatibility Scripts\Install\Template folder as a template.

Information about running these scripts is available in the "Designing Application Installation" section later in the chapter. For more information about the Windows Application Compatibility Toolkit, see "Testing Application Compatibility" in *Automating and Customizing Installations* of this kit.

# Hosting Full Desktops with Terminal Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you determine that hosting with Terminal Server is an ideal way to manage an application, consider running just the application (not the entire desktop) on the terminal server. This can save significant resources on the terminal server and can allow many more users to log on to the server simultaneously.

If, however, you want your users to access their full desktop from the terminal server as outlined in "Identifying the Role of Terminal Server in Your Organization" earlier in this chapter, be sure to fully test for performance and server capacity using the full load of applications to which your users will have access. For more information about server capacity, see "Terminal Server Capacity Planning" later in this chapter.

**Choosing the Desktop Theme**

If you host the full desktop with Terminal Server, the desktop environment is like the Windows desktop. The default desktop theme for Windows Server 2003, however, is Windows Classic. You can use the Windows XP default theme, however this theme affects performance for the user because it is more graphic intensive than Windows Classic. Test the responsiveness for the end user and perform real-user testing to ensure that the performance is satisfactory when using this theme. For more information about testing, see "Terminal Server Capacity Planning" later in this chapter.

You can also choose a specific theme for your end users or enforce the use of the Windows Classic theme through Group Policy. For more information, see "Configuring User Group Policy Settings" later in this chapter.

# Choosing the Licensing Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To use Terminal Server in your organization, you are required to have a Windows Server 2003 license for every terminal server that you deploy in your organization as well as Terminal Server Client Access Licenses (CALs) for devices that access the terminal servers. For terminal servers that are running Windows Server 2003, there are two types of Terminal Server CALs:

- Per Device
- Per User

Which CAL you choose depends on how you plan to use Terminal Server. By default, Terminal Server is configured in Per Device mode, but it can be switched to Per User mode using the Terminal Services Connection Configuration (TSCC) tool or by using Windows Management Instrumentation (WMI). You can serve both license types from the same license server. For more information about how to set your licensing mode, see "Designing the Terminal Server Configuration" later in this chapter.

**Note**

- Windows 2000 Internet Connector Licensing has been replaced by Terminal Server External Connector Licensing in Windows Server 2003. Improvements include licensing qualification extended to business partners in addition to customers, authenticated access, and unlimited concurrent users per server. For information about External Connector Licensing, see "External Connector License Overview" at http://go.microsoft.com/fwlink/?LinkID=100438. Windows Server 2003 retains support for Windows 2000 Internet Connector Licensing.

A Terminal Server license server on your network manages the Terminal Services CALs. A license server stores all Terminal Server CAL tokens that have been installed for a terminal server and tracks the license tokens that have been issued to clients. For more information about setting up a license server, see "Planning the License Server" later in this chapter.

For more information about Terminal Server licensing, see Terminal Server Licensing overview in Help and Support Center for Windows Server 2003 and "Microsoft Windows Server 2003 Terminal Server Licensing" at http://go.microsoft.com/fwlink/?LinkID=26220. For general information about licensing, see "Microsoft Volume Licensing" at http://go.microsoft.com/fwlink/?LinkId=100481.

**Note**

- Each service or application that users access from the terminal server must be licensed appropriately. Typically each device requires application licenses and CALs associated with it, even if the application or service is accessed indirectly through the terminal server. For more information, check the product documentation, End User License Agreement (EULA), or any other document that specifies product usage rights.

**Per Device Licensing Mode**

A Per Device CAL provides each client computer the right to access a terminal server that is running Windows Server 2003. The Per Device CAL is stored locally and presented to the terminal server each time the client computer connects to the server.

Per Device licensing is a good choice for:

- Hosting a user's primary desktop for devices the customer owns or controls.
- Thin clients or computers that connect to a terminal server for a large percentage of the working day.
- Hosting line-of-business applications that are used for the bulk of your users' work.

This type of licensing is a poor choice if you do not control the device accessing the server, for example computers in an Internet café, or if you have a business partner who connects to your terminal server from outside your network.

**Per User Licensing Mode**

In Per User licensing mode you must have one license for every user. With Per User licensing, one user can access a terminal

server from an unlimited number of devices and only needs one CAL rather than a CAL for each device.

**Note**

- At the release of Windows Server 2003, Per User licensing is not enforced. However the terminal server must be able to discover a license server after the 120-day grace period expires. Otherwise, clients are denied access to the terminal server. For more information about the 120-day grace period, see "Planning the License Server" later in this chapter. For more information about Per User licensing, see "Microsoft Windows Server 2003 Terminal Server Licensing" at http://www.go.microsoft.com/fwlink/?LinkID=26220. For information about plans for enforcing per user licensing, see "Microsoft Volume Licensing" at http://go.microsoft.com/fwlink/?LinkId=100481.

Per User licensing is a good choice in the following situations:

- Providing access for roaming users.
- Providing access for users who use more than one computer, for example, a portable and a desktop computer.
- Providing ease of management for organizations that track access to the network by user, rather than by computer.

In general, if your organization has more computers than users, Per User licensing might be a cost-effective way to deploy Terminal Server because you only pay for the user to access Terminal Server, rather than paying for every device from which the user accesses Terminal Server. Check the end-user license agreement for the applications that you plan to host to determine if they support per user licensing.

# Planning Network and Hardware Resources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you select the applications to host with Terminal Server, assess your network infrastructure, hardware, and security components to be sure your network is ready to deploy and run Terminal Server. Figure 4.3 shows the steps in this process.

**Figure 4.3 Planning Network and Hardware Resources**



# Planning Server Resources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Careful planning of the servers and server infrastructure on which you deploy Terminal Server determines the number of users who can access a terminal server at one time and the quality of their experience. You can also load balance your terminal servers for fault tolerance.

# Terminal Server Capacity Planning

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Server Capacity Planning**
Size your terminal servers with sufficient CPU, memory, and disk resources to handle the client demand. For adequate performance, a terminal server requires a minimum of 128 megabytes (MB) RAM, plus additional RAM for each user running applications on the server, depending on the type of user. For more information about server capacity requirements for Terminal Server, see the Terminal Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Important**

- Guidelines outlined on this Web site are for specific server configurations, and provide a baseline for estimating server capacity. Terminal Server capacity can vary depending on factors such as type of user, server and network configuration, and the applications you are hosting. Test your design thoroughly and monitor the server load after deployment to be sure your servers have adequate capacity to provide users with an acceptable experience.

A multiprocessor configuration can maximize CPU availability. In general, processor and memory requirements scale linearly. You can support nearly double the number of users on a multiprocessor-capable Pentium system by doubling the number of processors and doubling the amount of memory. For this reason, purchasing a system that supports multiple processors, even if you initially purchase only one processor, allows you to add capacity easily as your requirements grow.

**Estimating User Demand**
The amount of RAM and CPU that Terminal Server users consume depends on the application features that they use, how often they use the application, and how much work they accomplish in any unit of time. If you understand the work patterns of your users, you can use the Terminal Server capacity planning tools on the *Windows Server 2003 Deployment Kit* companion CD — Roboserver (Robosrv.exe) and Roboclient (Robocli.exe) — to more accurately emulate the activity level of your users. For more information about Roboserver and Roboclient, click **Tools** in Help and Support Center, and then click Windows Resource Kit Tools Help. How to use the capacity planning tools to emulate the activity of your users is discussed later in this section.

**Application Considerations**
When determining server sizing, examine the types of applications your server runs. Check system requirements for each application carefully and consider that RAM and CPU requirements increase according to the number of sessions expected to run simultaneously. Also, a terminal server shares executable resources among individual users, just as Windows shares executable resources among individual programs. As a result, the memory requirements for additional users running the same program are typically less than the requirements for the first user who loads the application. Although you cannot make precise estimates based on these factors, they give you a basis for projecting program performance.

**Testing Server Capacity**
Test the servers on which you plan to host applications in order to ensure adequate server capacity. Using a server that you estimate to be of the correct capacity, construct a test environment with the appropriate number of sessions running the application for a specified period of time and at a level of activity that simulates actual use. Because adequate performance can be subjective, you need to also perform real-user testing. Ask users to identify perceived performance issues, and make sure that you establish baseline task completion times based on user feedback, and use a stopwatch to accurately measure the effect of server capacity on performance throughout testing. A good indicator of unacceptable performance is when the time to complete the tasks in the simulation takes 10 percent longer than the baseline.

**Using the Capacity Planning Tools**
To test the capacity of your terminal server to handle the estimated user demand, the *Windows Server 2003 Deployment Kit* provides Terminal Server capacity planning tools — Roboserver (Robosrv.exe) and Roboclient (Robocli.exe) — which include application scripting support. You can use these tools, which are available on the *Windows Server 2003 Deployment Kit* companion CD, to easily place and manage simulated loads on a server. This information can help you determine whether your environment can handle the load that users place on it. Used together, these tools make capacity planning easier and more automated. For more information about Roboserver and Roboclient, click **Tools** in Help and Support Center, and then click Windows Resource Kit Tools Help.

To use these tools, you will need to equip your test lab with the following hardware and software:

- A domain controller

- A server running Terminal Server
- Any additional servers needed to emulate the intended use of Terminal Server in the production environment, such as:
    - A server running Microsoft® SQL Server™ (if you are going to host an application with a SQL Server back end)
    - A server running Microsoft® Exchange (if you intend to host e-mail)
- Client computers

**Note**

- The scripts included with the tools are designed for use with an Exchange server, and emulate typical usage patterns for knowledge workers and data entry workers. You can use these scripts as a template to create custom scripts for your needs. You should alter these scripts to emulate the users and infrastructure in your own environment.

The following sections summarize the hardware requirements for capacity planning tests.

For more information about using the capacity planning tools, click **Tools** in Help and Support Center for Windows Server 2003, and then click Windows Resource Kit Tools Help.

**Note**

- Using these tools can also help you identify compatibility issues with the applications you plan to host with Terminal Server.

**Domain controller** Roboserver runs on the domain controller. The domain controller should run Windows Server 2003, Enterprise Edition, or Microsoft® Windows® 2000 Advanced Server and be configured as the domain controller of a private domain. It should also be configured as a DNS server with both a Forward Lookup Zone and a Reverse Lookup Zone.

**Terminal server** The server that runs Terminal Server should have the same hardware and be configured the same as the server planned for the production environment.

**Client computers** Roboclient runs on the client computers. You need enough client computers to represent the number of users you plan to access the terminal server. You can emulate several users by using a single client computer. It is recommended that you limit the number to five, but if resources are limited, a low-end client computer can emulate as many as 20 users. In general, limit the number of users emulated on each client computer so as to keep the CPU usage of the computer below 20 percent.

**Note**

- Emulating several users by using a single computer in the test lab alters network traffic differently than in the production environment. Network usage is higher in the production environment.

You can also use System Monitor, a Windows Server 2003 support tool, to monitor the performance on the terminal servers in your test lab. To start System Monitor, at the command line type **perfmon**. For more information about System Monitor, see Perfmon in Help and Support Center for Windows Server 2003.

**Note**

- Measure disk performance from a remote computer because having the perfmon log on the disk that is being measured can affect the performance results.

Microsoft Operations Manager (MOM) is also a very useful tool for performance monitoring and scalability planning. For more information, see the Microsoft Operations Manager link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Real-User Capacity Testing**

Because a terminal server running at or above capacity affects performance, it is important that you also test server capacity with a sampling of real users. Also, the Terminal Server capacity planning tools simulate user activity but do not reflect actual human behavior, such as breaks and peaks and lulls in activity, and adequate performance for users is subjective. You can use your pilot test to further determine whether the capacity of your terminal servers is adequate. For more information about pilot testing, see "Planning for Deployment" in *Planning. Testing, and Piloting Deployment Projects* of this kit.

**Note**

- It is highly recommended that you conduct a pilot test for your design before implementing it in the production environment. For Terminal Server, pilot testing is crucial for determining correct server size and satisfactory application performance levels in the production environment.

**Data Storage Considerations**

With Terminal Server, there are some storage access regions that you should consider distributing across several hard drives or redundant array of independent disks (RAID) array, especially for implementations serving large number of users:

- System binaries
- Application binaries
- User profiles and data
- Page files

For Terminal Server, it is best to put your page files on a dedicated hard drive. In addition, because of the multiuser nature of Terminal Server, be sure that your page file is twice the size of your RAM. For example, if your Terminal Server has 2 gigabytes (GB) of RAM, you should have a 4 GB paging file. You should also have two paging files if necessary. For information about how to change the default paging file size, see "Change the size of the virtual memory paging file" in Help and Support Center for Windows Server 2003.

For improved performance of storage devices with a backup power supply, you can use the following procedure.

**To enable advanced disk performance for storage devices**

1. From the **Start** menu, click **All Programs**, click **Administrative Tools**, and then click **Computer Management**.
2. In **Computer Management**, under **Storage**, click **Disk Management**.
3. Right-click the disk for which you want to enable advanced performance and click **Properties**.
4. On the **Policies** tab, check the **Enable advanced performance** checkbox.

# Load Balancing Terminal Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Using a load-balancing solution with Terminal Server distributes sessions across multiple servers for improved performance. Terminal Services Session Directory, which is available with Windows Server 2003, Enterprise Edition, works with your load-balancing solution. Session Directory is a database that tracks user sessions that are running on load-balanced terminal servers. It provides information when a user reconnects (after disconnecting intentionally or because of a network failure) to ensure that the user reconnects to the same session rather than starting a new session. Session Directory, which can support several thousand sessions, is also cluster-aware, so that you can support users who have concurrent sessions on different terminal server farms without confusion

**Note**

- Load-balancing solutions that do not allow direct client network access to servers in the load-balanced farm need to support redirection tokens to work with Session Directory.

To implement load balancing and Session Directory with Terminal Server, you need to have the following in place:

- A load-balancing solution.
- Two or more terminal servers logically grouped into a Terminal Server farm.
- A Session Directory server.

**Setting Up a Terminal Server Load-Balanced Farm**

In general, set up your terminal server load-balanced farm and your load-balancing solution as you would in any other situation. Consult the documentation for your load-balancing solution.

There are a few things to consider for Terminal Server:

- Split network traffic between two network adapters — one used for Terminal Server, and the other for access to other network resources and infrastructure — placed on different subnets. By allowing RDP traffic only over the Terminal Services adapter, you can have more consistent traffic analysis and better security. More information about this is available later in this section.
- For easier administration, place load-balanced terminal servers into an organizational unit (OU) and apply Group Policy settings to that OU. If you have more than one Terminal Server load-balanced farm in your organization, place each in its own OU inside your Terminal Server OU so that you can apply Terminal Server–specific settings to the overall OU, but manage the separate farms individually.
- Configure your home directories and other user data storage in such a way that your users can easily access their data in the event of their being on a different server the next time they log on. For more information about how to do this, see "Planning Per-User Requirements" later in this chapter.
- Consider placing your terminal server farm and your clients on the same network backbone with your user profile servers and at least one domain controller. This allows for faster logon times.

**Network Separation**

Network separation is the separation of RDP traffic from other network traffic protocols. By dedicating one network adapter to Terminal Server traffic and the RDP, and one network adapter to application traffic, you can realize better overall performance on the Terminal Server load-balanced farms. This arrangement reduces the risk of having a network adapter bottleneck on the servers. Also, you can provide additional security by assigning different IP addresses and subnets to application traffic and to the RDP traffic. You can then constrain these separate routes as necessary by using routers, or switches and firewalls. This also tends to provide better auditing of users and traffic occurring on the network.

You can use IP packet filtering to achieve network separation. For more information about IP packet filtering, see "Designing a

TCP/IP Network" in *Deploying Network Services* of this kit. You can specify the network adapter on which you want to place the RDP traffic on the **Network Adapter** tab of the TSCC. For more information about Terminal Server configuration tools, see "Designing the Terminal Server Configuration" later in this chapter.

#### Selecting a Host Server for Session Directory

The Session Directory server can be any server on the network that is running the Terminal Server Session Directory service. It is best if the Session Directory server is a highly available network server that is not running Terminal Server. However, you can place the Session Directory on a member of the cluster, if necessary. The Session Directory requires very little CPU, memory, and hard drive resources, so you can use a low-end member server to host the Session Directory service.

**Important**

- Running the Session Directory service on a domain controller is not recommended. On a member server, the Session Directory Computers group is a local group, but on a domain controller this group is a domain local group and is available on all domain controllers. The Session Directory Computers group is discussed in the following section.

One Session Directory host server can service multiple load-balanced clusters, and it is cluster-aware so it can handle users who have sessions running on different clusters seamlessly. You can also cluster the Session Directory itself for improved reliability.

#### Configuring Session Directory

After you choose a server on which to host Session Directory, you must start and configure the Session Directory service and configure the servers it will serve. You must configure the Session Directory host server to accept connections from authorized computers and you must configure the load-balanced servers to use Session Directory.

#### Host server configuration

When the Session Directory service starts, by default it creates the Session Directory Computers group (if one does not already exist). The group is empty and you need to add to this group the load-balanced terminal servers that will use this Session Directory server.

#### Session directory settings

If you are load balancing several terminal servers and using Session Directory, you can configure the servers to use Session Directory through TSCC (under Server Settings), Group Policy, or WMI.

Set the following settings to configure your load-balanced Terminal Server farm to use the session directory server:

- **Terminal Server IP Address Redirection.** When the client computer cannot connect directly to the terminal server, you can use the **Terminal Server IP Address Redirection** Group Policy setting to mask the IP address of the destination server in a load-balanced farm.

  **Notes**

  - You can only configure this setting through Group Policy or WMI.
  - This policy setting is enabled by default. Disable this setting only if both of the following are true:

Your load-balancing solution does not allow direct connectivity from the client computer to the terminal server (for example, if your load-balancing solution is also a router).

Your load-balancing solution supports the use of Session Directory routing tokens.

- **Join Session Directory.** Enable this setting and apply it to your load-balanced Terminal Server OU to allow your server farm to use Session Directory. When you enable this setting, you must set the Session Directory Server and Session Directory Cluster Name settings.
- **Session Directory Server.** Enable this setting and enter the Domain Name System (DNS) name, IP address, or fully qualified domain name of the Session Directory server.
- **Session Directory Cluster Name.** Enable this setting and enter the DNS name for the load-balanced farm.
- **Network adapter and IP address session directory should redirect users to.** If your load-balanced farm is

configured so that network traffic is separated between two or more network adapters, use the drop-down list in TSCC to choose the network adapter on the load-balanced server to which client computers should be directed.

**Note**

- You can only configure this setting through TSCC or WMI.

**Implementing a Load-Balancing Solution with Session Directory**

After you install and configure the terminal servers, you can set up your load-balancing solution and the Session Directory service if you are planning to use this service. To set up load balancing, consult the documentation for your solution. To use Group Policy to configure and manage your load-balanced Terminal Server farm centrally, create an OU and add each computer in the farm to the OU.

To set up Session Directory, you need to set up both the Session Directory host server and the servers that will be using Session Directory.

**To set up the Session Directory host computer**

1. Right-click **My Computer** and click **Manage.**

2. Navigate to Services and Applications/Services, double-click **Terminal Services Session Directory**, and set the Startup type to Automatic.

   This service is off by default, and it is set to Manual. Starting this service and setting it to Automatic ensures that the service starts when the host server is turned on.

3. Add the computer group containing the load-balanced Terminal Server farm to the Session Directory Computers group that is created when the Session Directory service starts.

   **Note**

   - By default, the Session Directory Computers group is empty. You can also create this group prior to starting the Session Directory service.

4. Configure the servers that will use the Session Directory service according to your plans from the previous section.

# Running Terminal Server with Other Applications and Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Choosing the servers on which you deploy your applications depends on the applications you plan to host with Terminal Server and the applications and services you plan to run on the server with it. For example, you might want to conserve resources by hosting Terminal Server on a computer that also serves other purposes, or an application you plan to host might require another back-end service to run.

**Sharing Terminal Servers with Other Applications and Services**

To conserve resources, you can deploy Terminal Server on a server with other applications and services on the network. However, because Terminal Server is optimized for the desktop experience, this can affect the performance and scalability of both Terminal Server and the other application or service. You should, for example, avoid deploying server-based services, such as SQL Server and Exchange Server, with Terminal Server. Likewise, deploying Terminal Server on a server along with file, print, or other services can affect the performance of these services. This can cause problems especially if non-Terminal Server users access these services. Also, some services can cause conflicts with certain applications that you are hosting with Terminal Server if the client and server applications use different versions of DLLs. Test your setup if you plan to deploy Terminal Server on a server with other applications and services.

**Hosting Data on a Separate Server or Drive**

If you are planning to use Terminal Server to host a server-side, data-intensive application (for example a custom line-of-business application that has a SQL Server back end), you can improve performance by hosting the database and Terminal Server on separate servers. Placing the terminal server and the database on the same high-speed link or even on the same subnet increases the responsiveness of the application. You can optimize disk access performance by configuring multiple SCSI channels and distributing your operating system and application traffic across different physical drives. You can also dedicate a network adapter to operating system and application traffic, while the other network adapter handles RDP traffic.

# Planning the License Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A license server is required when Terminal Server is deployed for hosting applications and desktops. A license server stores all Terminal Server client license tokens that have been installed for a terminal server or group of terminal servers. It tracks the license tokens that have been issued by the Microsoft License Clearinghouse (the service that Microsoft maintains to activate license servers and to issue licenses to the license servers that request them). The license server manages client licenses through the terminal server. There is no direct communication between the client and the license server. For more information about Terminal Services licensing and how it works, see Terminal Server Licensing overview in Help and Support Center for Windows Server 2003 and the "Microsoft Windows Server 2003 Terminal Server Licensing" white paper at the Terminal Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

If you have a terminal server that is running Windows Server 2003, you must also have a license server that is running Windows Server 2003. You can use a license server that is running Windows Server 2003 to serve license tokens onto the terminal server that is running Windows Server 2003, or you can upgrade an existing license server that is running Windows 2000 to Windows Server 2003. You can set the **Prevent License Upgrade** Group Policy setting to prevent the license server from handing out Windows Server 2003 CALs to terminal servers that are running Windows 2000.

**Note**

- For Windows Server 2003 you are no longer required to host Terminal Server licensing on a domain controller, but if you do not, you must configure the Terminal Server to discover the license server. For more information, see "Configuring the Preferred License Server" later in this section.

For information about how to transfer licenses from one license server to another, see Repeat the installation of a client access license in Help and Support Center for Windows Server 2003.

A terminal server must be able to connect to an activated license server before license tokens can be issued. To help you with your deployment planning and to give you time to install a license server and add license tokens, there is a 120-day grace period after the first client connects to a terminal server that is running Windows before you must have a license server activated. After this period, if the terminal server does not find an activated license server, clients are denied access. This grace period also ends when the first license token is issued to a client.

You can manage your license servers and Terminal Services license tokens using the Terminal Server Licensing administrative tool. For more information, see Terminal Services administrative tools in Help and Support Center for Windows Server 2003.

When planning your license server requirements, consider your organization's domain structure in order to choose a license server type and to select a server that is a good candidate for hosting Terminal Server Licensing. If, for example, you are deploying terminal servers for different business units or cost centers within your organization, consider deploying a license server for each Terminal Server deployment rather than a license server to service an entire domain or site. This enables each business unit to manage its own licenses and related costs. You must, however, configure your terminal servers to discover the correct license server in this situation. There is information about configuring a preferred license server later in this section.

**Choosing the License Server Type**

There are two types of license servers for Windows Server 2003, a domain license server and an enterprise license server. Decide which of the two types of license servers you require.

**Domain license server**

A domain license server is appropriate for managing Terminal Server users within a domain. Use a domain license server if you have only one domain, or if you want to manage licenses in each of your domains separately. If you have workgroups or Windows NT 4.0 domains, a domain license server is the only type that you can install, and any terminal server within the Windows NT 4.0 domain or workgroup can discover the license server. Terminal servers in a Windows 2000 or Windows Server 2003 domain can discover domain license servers only if they are in the same domain as the license server.

**Enterprise license server**

An enterprise license server can serve terminal servers in any domain within a site, but the domain must be a Windows Server 2003 or Windows 2000 Active Directory domain. Both the license server and Terminal Server must be within the same site.

This type of license server is appropriate if you manage users across multiple domains within a site, or if you are planning to restructure your domains. Enterprise license servers can be installed only by using Add or Remove Programs, not during

Windows Server 2003 Setup. For more information about installing the license server, see "Designing Server Setting Configurations" later in this chapter.

#### Choosing a License Server Host Server

Servers running Windows Terminal Services can work only with license servers that are running Windows. In Windows Server 2003 domains, you can deploy the enterprise license server on either a domain controller or a member server. You are no longer required to install the license server on a domain controller, as was required in Windows 2000. For information about upgrading your license servers, see "Upgrading to Windows Server 2003 Terminal Server" later in this chapter.

### Important

- If you plan to deploy a domain license server, and you choose not to deploy your license server on a domain controller, you must directly edit the registry to configure Terminal Server to locate the license server. You must do this even if the license server and the terminal server are the same computer. For more information about this process, see "Configuring the Preferred License Server" later in this section.

Take the following situations into consideration when choosing where to host Terminal Server Licensing:

- If you are deploying a license server to serve only one terminal server, consider installing the license server on the same server as Terminal Server. Terminal Server licensing causes negligible network traffic, but placement on the same server as Terminal Server eliminates the need for the server to go over the network to obtain licenses, and Terminal Server users experience no downtime if problems occur on the network.

- If you are deploying a license server for a terminal server that is hosting applications for remote clients or for users who access applications over the Internet, place the license server so that it can be accessed easily by the terminal server because there is no direct communication between the client and the license server.

Terminal Services Licensing takes about a megabyte (MB) of disk space to install, and about 10 MB of additional disk space per 1,000 licenses. If you plan to activate the license server and install subsequent license key packs automatically through the Terminal Services Licensing tool, the license server must communicate with the Microsoft License Clearinghouse using Secure Sockets Layer (SSL). SSL is a secure connection; however, as an added security precaution, avoid placing your license server on a server that hosts sensitive data that you would not otherwise want exposed to the Internet. For more information about methods to activate the license server, see Activating a Terminal Server license server in Help and Support Center for Windows Server 2003. For more information about methods for installing subsequent Client Access Licenses, see Install Client Access Licenses in Help and Support Center for Windows Server 2003.

#### Setting Up Fault Tolerant License Servers

One activated license server can serve many terminal servers and Terminal Server users simultaneously. In large implementations of Terminal Services, you might want to deploy at least two licensing servers for fault tolerance. Terminal Server Licensing is not cluster-aware. However, you can improve fault tolerance for the license server by deploying two license servers per domain.

### Note

- You can host Per Device and Per User licensing on the same server.

Choose two host servers for licensing and install all licenses on one license server; do not install any licenses on the backup server. The full server functions as the primary license server and can service all license requests. If the primary license server becomes unavailable, the backup license server can issue 90-day temporary license tokens to clients that need them. When the primary license server is back online, it replaces the temporary license tokens the next time the clients connect.

#### Configuring the Preferred License Server

If you choose not to deploy your license server on a domain controller, you must configure Terminal Services to locate the license server by directly editing the registry. For more information about the Terminal Services license server discovery process for Windows Server 2003, see article 279561, "How to Override the License Server Discovery Process in Windows Server 2003 Terminal Services." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

### Caution

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing

settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD, or at http://www.microsoft.com/reskit.

**To select a specific license server**

1. On each terminal server, in the **Run** dialog box, type **regedit**, and then click **OK**.

2. Locate the following path in the registry and select it:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\TermService\Parameters

3. On the **Edit** menu, point to **New**, and then click **Key** to add a subkey named LicenseServers, if one does not already exist.

4. Select the LicenseServers subkey, and on the **Edit** menu, point to **New**, and then click **Key** and type in the NetBIOS name, fully qualified domain name, or IP address of the appropriate license server to add the new subkey under LicenseServers.

5. Repeat step 4 for each license server you want the terminal server to discover.

The order of the license servers listed in the registry subkey does not guarantee the order in which the license servers are discovered. If you have more than one implementation of Terminal Server in your organization, you should also configure the license servers themselves to discover only the servers intended to service that implementation by using this procedure.

When Terminal Server discovers a license server, it continues to use that license server until it becomes unavailable. If the licence server is out of licenses, it will forward the request to a license server that has the requested type of license.

**Note**

- If your license servers are not on domain controllers, they will only discover other license servers if their registries are configured for them to do so.

If you are deploying more than one license server to serve separate Terminal Server implementations and if you do not need to keep the licenses allocated to each implementation separate, consider adding the names of all of the license servers in your site or domain to the registry subkey LicenseServers for fault tolerance. However, if you need to keep the licenses allocated to each implementation separate, for example if they are managed by a different administrator or are in different cost centers, you should add to the registry subkey only the name of the license server or servers specific to each terminal server. In this situation, consider deploying two license servers if fault tolerance is a concern.

# Planning Network Connectivity and Bandwidth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Terminal Server uses RDP to communicate between client and server. RDP works only across a TCP/IP connection, such as a local area network (LAN), wide area network (WAN), dial-up, Integrated Services Digital Network (ISDN), digital subscriber line (DSL), or virtual private network (VPN) connection. You can still use other protocols, such as Internetwork Packet Exchange (IPX) or NetBIOS Extended User Interface (NetBEUI), as the transport protocol for non–Terminal Server traffic, such as network file or printer sharing, or between the client portion of a client-server application and its server.

**Note**

- The Remote Desktop client requires TCP/IP to connect to the server, but after Terminal Server users connect to a terminal server, they can use IPX to gain access to Novell servers if necessary.

**Bandwidth Considerations**

Terminal Server works very well over low-bandwidth connections and uses whatever IP connection you provide. However, you can optimize both application and overall network performance by making sure the type of connection is appropriate to the work that is done. For example, a single user can connect over a low-bandwidth modem line and realize good performance, but it is not appropriate to share a 28.8-kilobit line among an active office of 100 people. Also, consider whether the security that an IP connection provides is appropriate to the data that you plan to transmit. For more information about bandwidth requirements, see the Terminal Services link and the Remote Desktop Protocol (RDP) Features and Performance link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Note**

- Printing, sound, drive redirection, and user file transfer requirements can increase the bandwidth requirement and might cause performance to drop below a level that is considered acceptable performance for users.

**Access to Terminal Server over a VPN**

Users can gain access to Terminal Server over a VPN by using Layer Two Tunneling Protocol (L2TP) or Point-to-Point Tunneling Protocol (PPTP). By using encryption, either one of these tunneling options provides secure access to a private network for users operating over a public medium.

**Access to Terminal Server Over a Wide Area Network**

Terminal Services can provide remote users with access to applications that would otherwise be unusable because of poor performance across dial-up or slow WAN connections.

Give additional consideration to applications with offline capabilities and synchronization if users will be accessing Terminal Server over a WAN. No offline synchronization of the local application is performed over the terminal server connection.

# Planning Network Security Components

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use firewalls and smart cards with Terminal Server to increase security. Use the information in the following sections to help you plan for the use of these security components in your environment. You also need to carefully choose the file system you plan to use with Terminal Server.

**Using Firewalls with Terminal Server**

If your organization uses a firewall for security and if you need clients to connect from the other side of the firewall, keep port 3389 open for RDP connections between the client and server and implement filter rules to ensure this traffic can reach only the terminal servers. This is not necessary if you use PPTP, L2TP, or other VPN technologies to tunnel through the firewall, because the port will be available through the tunnel, without being explicitly closed by the firewall. If you are using the Remote Desktop Web Connection, you must also check to be sure that port 80 is open.

You can change the RDP port if necessary, but you must apply the modification to both the server and clients. However changing the RDP port, for example to a well known and already open port such as 80, makes separation, identification, and audit of RDP traffic much more difficult. For more information about changing the RDP port, see article 187623, "How to Change Terminal Server's Listening Port" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For best results, use a firewall that uses user-based authentication. This is especially important for secure access over a WAN, because a firewall that grants access based on an IP address (rather than based on user credentials) allows users through if the IP address of the server running Terminal Services has been granted access. For more information, see "Planning Network Connectivity and Bandwidth" earlier in this chapter.

Determine if the firewall used in your organization is a packet-level or application-level firewall. Packet-level firewalls are easier to configure for new protocols. If your organization uses an application-level firewall, check to see if the vendor has defined a filter for the RDP.

**Using NTFS with Terminal Server**

Because of the multiuser nature of Terminal Services, it is strongly recommended that you use the Windows Server 2003 version of NTFS as the only file system on the server, rather than file allocation table (FAT). FAT does not offer any user and directory security, whereas with NTFS you can limit subdirectories to certain users or groups of users. This is important in a multiuser system such as Terminal Services. Without the security that NTFS provides, any user has access to every directory and file on the terminal server. For more information about file systems, see "Designing and Deploying File Servers" in this book.

**Using Smart Cards with Terminal Server**

With Windows Server 2003, you can now enable users to log on to a remote session in an Active Directory domain using a smart card. Smart cards allow you to require strong credentials from users in a manageable way, providing a more secure environment.

To use smart cards with Windows Server 2003 Terminal Server, you must have Active Directory® directory service deployed in your organization and your client computers must be running a Microsoft client operating system with built-in smart card support, such as Windows XP or Windows 2000, or most devices running Microsoft® Windows® CE .NET. You must also install smart card readers on the client computers.

Otherwise, deploying smart cards for use with Windows Server 2003 Terminal Server is the same as deploying smart cards that will not be used with Terminal Server. For more information see "Planning a Smart Card Deployment" in the *Designing and Deploying Directory and Security Services* book of this kit.

# Choosing Client Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Client computers, or terminals, connect to a terminal server that is running Windows Server 2003 using the Remote Desktop Connection tool, which you can install on a disk or in firmware. The Remote Desktop Connection tool is installed by default when you install Windows Server 2003, Windows XP (you need to install Windows XP Service Pack 1 to benefit from features such as Auto-Reconnect), and most versions of Windows CE. On earlier versions of Windows and on the Pocket PC, you have to manually install Remote Desktop.

You can also install this tool manually from the Windows Server 2003 operating system CD or download it from the Web for use on a computer running the Microsoft® Windows® 95, Microsoft® Windows® 98, Microsoft® Windows® Millennium Edition, Windows NT 4.0, or Microsoft® Windows® 2000 Professional operating systems. Remote Desktop Connection has very low physical RAM requirements and generally works on any device that meets the minimum requirements for the operating system on which it runs. For information, see the Terminal Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

A Remote Desktop Connection Client for Mac OS X is also available from Microsoft. For more information, see the Mactopia Downloads: Remote Desktop Connection Client link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For other platforms such as UNIX, you need to obtain a third-party tool to connect to a terminal server. There are a number of third-party RDP clients available. Many of these are reverse engineered and are not supported or endorsed by Microsoft. These clients might cause performance and compatibility problems, and lack of functionality. Check with the software vendor to ensure that your RDP software is a licensed client. For more information about third-party Terminal Server solutions, see the Citrix and New Moon Systems links on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Your choice of client platform depends on the current installed base and individual user need. Consider the following issues when choosing a client computer or platform from which to connect to your terminal server:

- You can use Windows-based terminals in the following situations:
    - Task workers can access their primary application from a Windows-based terminal (or thin client) rather than a personal computer, reducing hardware costs and work disruption. If a terminal stops responding, you can replace it quickly, with minimal setup.
    - In locations where a personal computer does not fare well — such as in exceptionally dirty, hot, or cold environments — consider deploying Windows-based thin clients, which typically have no moving parts and have a sealed-case design.
    - Windows-based thin clients can be used to replace green-screen terminals.
- You can use client computers running earlier versions of Windows to provide access to applications that require a more recent version of Windows by upgrading your terminal server and hosting the application.
- To save on hardware costs, you can use older computers that are not capable of running newer versions of Windows to access the latest application versions through Terminal Server.
- The Remote Desktop Web Connection requires a browser that supports ActiveX (available with Internet Explorer).
- Non-Windows-based computers such as Macintosh or UNIX computers need additional client software to connect to Terminal Server.

# Planning Use of Peripheral Hardware

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use desktop peripherals such as bar code readers, scanners, or card swipes with Windows Server 2003 Terminal Server if your line-of-business application requires them. You can use peripheral hardware attached to serial and parallel ports of the client computer. You can also use printers attached to the client, as well as network printers, with Terminal Server. Peripherals that connect to the local computer through a specialized third-party hardware card are not recognized by the RDP-based Remote Desktop client, but they might be available using third-party software. For more information, see Peripherals that affect performance in Help and Support Center for Windows Server 2003.

**Port Redirection**

Terminal Server supports port redirection in Windows Server 2003. Using port redirection, applications running within the Terminal Server session have access to ports on the client, enabling you to use devices like bar code readers or scanners with Terminal Server. Test your devices for use with Terminal Server before deploying your solution in the production environment.

By default Windows Server 2003 Terminal Server enables LPT and COM port redirection. If you are not planning to enable these types of peripherals, you might want to disable them so that the client computer is not vulnerable to security threats that could access your computer through those ports. You can disable LPT and COM port redirection by using Terminal Server Group Policy or TSCC. For more information, see "Designing the Terminal Server Configuration" later in this chapter.

By default, Terminal Server does not allow redirection of FireWire (IEEE 1394) ports with Windows XP and Windows 2000 clients. However, you can enable FireWire port redirection on these clients by enabling all ports to be redirected by using the following procedure to modify the registry on the client computer.

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.

**To filter ports for redirection**

1. On the client computer, in the **Run** dialog box, type **regedit**, and then click **OK**.
2. Locate the following subkey in the registry and select it:

   HKEY_CURRENT_USER\Software\Microsoft\Terminal Server Client\Default\AddIns\RDPDR
3. On the **Edit** menu, point to **New**, and then click **DWORD Value**.
4. Type **FilterQueueType**, and then press **ENTER**.
5. With **FilterQueueType** selected, on the **Edit** menu, click **Modify**.
6. Type **FFFFFFFF**, and then click **OK**.

For more information about port redirection, see article 302361, "Multifunction Printers That Use DOT4 Ports Are Not Redirected By Using Remote Desktop." in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Printer Redirection**

When the user logs on to the terminal server, the server detects the client's local printer and installs the appropriate printer driver on the remote computer. If multiple printers are connected to the client computer, Terminal Services defaults to routing all print jobs to the client computer's default printer. Only printers whose drivers are available on the Windows client computer appear as available in a Remote Desktop session for local redirected printers (server-side printers are always available). If the driver for your printer is not included with the client operating system, you must manually install it on the server.

# Designing Domain and User Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Developing your design also involves planning the location of Terminal Server within your proposed Windows Server 2003 domain infrastructure, as shown in Figure 4.4. After you have determined how Terminal Server integrates into your domain model, you can plan the user and security settings that can be managed through Active Directory and Group Policy.

**Figure 4.4 Planning Domain and Security Configuration**



# Integrating Terminal Server into Your Domain Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Terminal Server need not be in an Active Directory domain to function, but without a domain architecture, users must have separate accounts on every computer running Terminal Server. This limits manageability and makes it more difficult to administer groups of users.

**Integrating with Existing Windows NT 4.0 Domain Structure**
If your organization does not currently use Active Directory, you can use an existing Windows NT 4.0 domain, which allows you to take advantage of the new features available in Windows Server 2003 Terminal Server without affecting the production environment. However, limitations apply, such as the existing Security Accounts Manager (SAM) 40,000-objects-per-domain limitation of the Windows NT 4.0 domain model. Administrators have the option of adding Terminal Server–specific attributes to users' accounts. This adds a small amount of information, typically 1 kilobyte (KB) or less, to a user's entry in the domain SAM database.

**Integrating with the Windows Server 2003 and Windows 2000 Active Directory Infrastructure**
This option takes full advantage of Active Directory, giving you the option of applying Group Policy settings to control the Terminal Server environment. Just as you are likely to manage your portable computers or domain controllers in a manner different from your desktop computers, you also manage your terminal servers and Terminal Server users differently. When you define your Active Directory structure, it is recommended that you place your terminal servers in a separate Terminal Services OU. Reserve this OU for Terminal Services computers. Do not include other users or non-Terminal Services computer objects. In addition, if you are deploying load-balanced server farms for Terminal Services, place each farm in a separate OU within the Terminal Services OU. It is also recommended that you place your Terminal Server users in a separate Terminal Server users OU.

In an Active Directory environment, avoid configuring Terminal Server as a domain controller for the following reasons:

- Any user rights policies you apply to such a server apply to all domain controllers in the domain. For example, to use Terminal Services, users must be authorized to log on locally to the server. If the server running Terminal Services is a domain controller, users can log on locally to all domain controllers in the Terminal Services domain, presenting a serious security risk.

- Domain controller functions place a heavy load on system resources and would thus have an effect on the user's Terminal Server experience.

- By default, enabling Terminal Services sets the server process-scheduling priority to favor interactive applications. The system does not assign top priority to critical domain-level processes such as user count replication, logon requests, logon script replication, and authentication requests.

# Planning Per-User Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because of the multiuser nature of Terminal Server, it is important to plan per-user settings and data storage carefully for ease of management and an optimal user experience. How you plan to use Terminal Server in your organization affects your choice of user profile types, and your plan for using these profiles.

# Using User Profiles in Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A profile describes the Windows Server 2003 configuration for a specific user, including the user's environment and preference settings. Profiles typically contain such user-specific information as installed applications, desktop icons, and color options. To plan for user profiles in a Terminal Server environment, choose the solution that is best for your environment, and then plan for the storage of the profiles. For more information about user profiles, see User profiles overview in Help and Support Center for Windows Server 2003. For information about general planning for user profiles, see "Implementing User State Management" in *Designing a Managed Environment* of this kit.

Unless you plan carefully for the use of user profiles, they tend to grow in size. This is a problem in a Terminal Server environment because user profiles are stored on the terminal server by default. If you have many users accessing the terminal server, the user profile files soon consume a large amount of space on the server hard drive. You should store user data and profiles on a separate drive from the system installation hard drive.

There are three different types of profiles you can use with Terminal Server:

- Terminal Server–specific profile
- Windows Server 2003 mandatory roaming profile
- Windows Server 2003 local profile

When a user logs on to a server running Terminal Server, the server first searches for the Terminal Server–specific profile. If Terminal Server cannot locate this profile, it attempts to load the user Windows Server 2003 roaming profile or Windows Server 2003 local profile.

It is recommended that you plan to use either Terminal Server–specific or roaming user profiles for your Terminal Server users, rather than local profiles, in order to better manage the size of the profiles and optimize the user experience. Terminal Server–specific profiles are recommended in most cases. Consider the following situations when choosing which type of user profile to use with Terminal Server:

- If you are planning to keep the environment for your Terminal Server users standardized and under tight control, you can use mandatory roaming user profiles to restrict access to certain applications. You can also use mandatory roaming user profiles to assign users profiles that cannot be changed.
- If you assign roaming user profiles to users who tend to access the terminal server from different computers (for example IT administrators, users who access the application from a kiosk, or users who work in certain task-worker environments), those users can retain their settings regardless of where they log on.
- If you are using Terminal Server to deliver a consistent desktop to client computers of varying platforms or configurations, you cannot use roaming user profiles unless you can group the different configurations and platforms into different OUs.
- If you are using Terminal Server in a load-balanced farm, you should plan to use roaming user profiles.

**Using Terminal Server–Specific Profiles**

Use Terminal Server–specific profiles to present a session to the user that is different from the user's desktop or to create user profiles optimized to the Terminal Services environment. The following are some of the situations where using Terminal Server–specific profiles might be advantageous:

- To provide users who are accessing Terminal Server with an environment that is different from the environment on their local computers.
- To provide a different look and feel for different users on the same terminal server, for example, if you have task workers and a manager on the same server.

- To better manage the size of user profiles for Terminal Services users who do not have controlled user environments that have been set through assigned or mandatory user profiles. You can use Group Policy to manage the profiles on the server that stores your Terminal Server profiles.

You can configure Terminal Services–specific profile settings for each user by using the following procedure.

**To configure Terminal Services–specific profile settings**

1. Open Active Directory Users and Computers.
2. Right-click the user for which you want to set profile settings, and then click Properties.
3. Click the Terminal Services profile tab.

You can configure the following Terminal Services–specific profile settings:

- **Terminal Services User Profile path**. You can choose a place to store users' Terminal Services profiles other than the default location.

  **Note**

  - You can also set this through Group Policy under Computer Configuration\Administrative Templates\Windows Components\Terminal Services. For more information, see "Designing the Terminal Server Configuration" later in this chapter.

- **Terminal Services home folder**. You can specify a path to a home folder for use with Terminal Server sessions. This directory can be either a local folder or a network share.

For information about setting Terminal Server profiles, see Terminal Services Profile in Help and Support Center for Windows Server 2003.

**Using Roaming Mandatory User Profiles**

Roaming user profiles allow users to move between different computers and maintain the same environment and preference settings. A roaming mandatory user profile is a preconfigured user profile that you assign to users. Because users cannot change a roaming mandatory profile, using this type of profile ensures that these user profiles remain at a manageable size. Additionally, you can assign one mandatory profile to all users who require identical desktop configurations. This allows you to change the desktop environments for all those users by changing only one profile.

Take the following issues into consideration when planning to use roaming mandatory user profiles with Terminal Server:

- When planning for the use of profiles for a large number of Terminal Server users, consider using Terminal Server profiles rather than roaming user profiles.
- If you are combining Folder Redirection and roaming user profiles, it is recommended you not use quotas on the profile.
- If your users roam between computers that are running Windows XP Professional, Windows XP 64-Bit Edition, Windows Server 2003, and Windows 2000, you can use the **Prevent Roaming Profile changes from being propagated to the server** Group Policy setting to be sure that each client computer receives only the profile that applies to the particular platform that the user is logged on to. For more information, see "Group Policy in multiplatform networks" in Help and Support Center for Windows Server 2003. To find this topic, click **Index** in Help and Support Center, type the keywords "Group Policy," and then select the topic "multiplatform networks."

The roaming profile information is stored on the local hard drive of the terminal server. It is recommended that this information be deleted after the user logs off. You can do this by enabling the **Delete cached copies of roaming profiles** Group Policy setting (in System/User Profiles under User Configuration in the Group Policy Object Editor) and applying the setting to your Terminal Server OU.

**Important**

- In order to use roaming profiles on a group of Terminal Services computers, the Terminal Services computers must be identical in application and operating system configuration, such as the location of the systemroot folder and the installation location of all applications. Otherwise, group different configurations into different OUs and administer the roaming profiles separately.

For information about how to set or change a user's roaming profile path, see Change a user's Terminal Services profile path in Help and Support Center for Windows Server 2003.

**Planning for User Profile Storage and Management**

Unless you manage user profiles correctly, they can become very large and can cause problems for your Terminal Server users. In order to keep the size of your user profiles for Terminal Server under control, use the **Limit profile size** Group Policy setting or use mandatory profiles. You can find the **Limit profile size** Group Policy setting under User Configuration\Administrative Templates\System\User Profiles.

The profile path copies all user profiles to drive C of the terminal server by default. Depending on the number of users accessing your terminal server, this could greatly deplete the free space on this disk. Choose a location on a file or print server that has enough space to store the profiles and that is readily available to Terminal Server users, and then create a Windows Server 2003 share that users can access with read/write permissions. Do not store Terminal Server profiles and users' primary desktop profiles in the same location. You should store profiles in a different location from user home directories. For information, see "Change a user's Terminal Services profile path" in Help and Support Center for Windows Server 2003.

**Increasing Time-out Values for Profiles with Terminal Server**

**Caution**

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Windows Server 2003 Deployment Kit* companion CD, or at http://www.microsoft.com/reskit.

In a Terminal Server environment, because many users tend to access the terminal server and the profile server at the same time, the server can develop bottlenecks or the network itself can become saturated. This can cause problems with user profiles primarily because a time-out can occur during profile unloading or write back. As a result, changes to the profile are not saved. By increasing the time-out values when you set up Terminal Server, you can reduce the incidence of profile-related issues. You can increase profile time-out values by using the following procedure.

**To increase profile time-out values**

1. In the Group Policy Object Editor, navigate to the **Maximum retries to unload and update user profile** policy, which is located in Computer Configuration/Administrative Templates/System/User Profiles.

2. Enable this setting and set it to 120.

3. In the **Run** dialog box, type **regedit**, and then click **OK**.

4. Locate the following subkey in the registry and select it:

   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Terminal Server

5. On the **Edit** menu, click **Add,** and then click **DWORD Value.**

6. Add a registry entry named LogoffTimeout with the following settings:

   - Base: Decimal

   - Value: 120 (4 minutes, time-out expressed in 2-second units)

     **Note**

Do not set this value lower than 3 minutes or higher than 15 minutes.

For more information, see article 299386, "Logoff Process May Not Be Completed Because Time-Out Is Too Slow" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Using Home Directories with Terminal Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important that you plan for use of home directories in a Terminal Server environment because most applications must install user-specific information or copy configuration files for each user. By default, Windows Server 2003 defines a home directory for each user. For Terminal Server users, the default user's home directory is his or her user profile directory on the terminal server, for example \Wtsrv\Profiles*Username*. This directory contains user settings. Terminal Services writes user-specific application files, such as .ini files, and by default refers any application seeking the Windows system directory to the user's home directory.

Users typically save their personal files to their home directory, in the My Documents folder. This can be a problem if roaming profiles are used and the home directory is located within the user's profile directory. Windows Server 2003 copies everything in the user's profile directory to the profile cache on the local computer each time the user logs on. This can take considerable time and resources, particularly if the roaming profile is stored across the network or over a slow link. You can use Group Policy to redirect the My Documents folder to a central non–Terminal Server computer. For more information, see "Change a user's Terminal Services profile path" in Help and Support Center for Windows Server 2003.

It is recommended that you use Terminal Services–specific home directories. Choose a location on a file or print server for the home directories. Share the file and give Change permissions to Everyone, then change the home directory path for Terminal Server users. By default, users have full access to their individual home directory. Administrators can copy files into the directory, but not read or delete files there.

You can specify a location for redirecting the home directory for Terminal Services users with the **TS User Home Directory** Group Policy setting. With this setting you can specify the location of the home directory on the network or on the local computer, the root path, and the network drive letter if the root path is located on the network.

**Note**

- To facilitate the use of application compatibility scripts, use the same virtual drive letter for all user home directory redirection points. The first time you run an application compatibility script on a server, the server prompts you to set the drive letter that references the root of the user home directory. This drive letter is used for all subsequent application compatibility scripts. For terminal server farms, it is essential that the same drive letter be used on all the servers within the farm.

# Using Folder Redirection with Terminal Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Folder Redirection allows users and administrators to redirect the path of a folder to a new location. The new location can be a folder on the local computer or a directory on a network share. It is recommended that you redirect users' My Documents folder to a private server share by using Folder Redirection. Users can then access their My Documents folders from either a Windows XP Professional client computer or Terminal Server session as if they were accessing their local drives. This is especially useful for roaming users who access the terminal server from different computers at different times.

For general information about Folder Redirection, see "Folder Redirection" in Help and Support Center for Windows Server 2003. To find this topic, click **Index** in Help and Support Center, type the keywords "Folder Redirection," and then select the topic "overview."

When setting Folder Redirection, make sure that the user has write access to the folder in which you create the new folder for that user. If the user does not have sufficient access to the shared folder, the folder redirection fails and defaults to a dedicated folder created on the server or desktop to which the user logs on.

# Configuring User Group Policy Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Under User Configuration in the Group Policy Object Editor, you can set several Group Policy settings that are particularly useful for Terminal Server. Use these settings to control the user experience and prevent access to areas of the terminal server. For more information about each of the settings listed here, see the Group Policy Explain text associated with each setting. For a job aid to assist you in recording your Terminal Server Group Policy configuration decisions, see "Group Policy Configuration Worksheet" (SDCTS_2.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Group Policy Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

See the following resources for more specific information about using Group Policy:

- For general information about Group Policy, see the Management Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- For more information about designing a Group Policy infrastructure, see "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit.
- For information about using Group Policy to lock down a Terminal Server session, see article 278295, "How to Lock Down a Windows 2000 Terminal Server Session" in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- For more information about applying Group Policy to Terminal Server, see article 260370, "How to Apply Group Policy Objects to Terminal Services Servers." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Note**

- Because these settings apply to the user, and not the computer, they affect the user environment regardless of which computer the user accesses. When applying settings that you want to apply only when users have a session on the terminal server (as opposed to their own desktop computer), use computer settings that apply to the terminal server. For more information, see "Designing Terminal Server Installation and Configuration" later in this chapter.

**Configuring the User Display**

A graphic-intensive display can affect performance for users of Terminal Server. To ensure the best possible performance, you can control what users can put on their desktops by configuring the Group Policy settings located under User Configuration/Administrative Templates/Control Panel/Display.

**Configuring desktop items**

Many organizations permit users to choose their own desktop wallpaper or screen savers. However, in a Terminal Server environment, these graphics can have an effect on performance. Use the following Group Policy settings to control users' ability to change wallpaper and screen savers.

**Screen savers** You can use several Group Policy settings to affect the user's screen saver. You can disable screen savers altogether by disabling the **Screen Saver** policy. You can also specify the screen saver by enabling this policy and also by enabling and specifying the screen saver executable name in the **Screen Saver executable name** policy. For more information about these Group Policy settings, see the **Explain** tab located on the property sheet for each policy.

**Wallpaper** By enabling the **Prevent changing wallpaper** setting you can disable all the options in the **Desktop** tab of Display in Control Panel. This includes changing the wallpaper and changing the appearance of the desktop icons. By not allowing these changes, you can ensure that users do not choose desktop display items that might affect the performance of the server.

**Configuring the desktop theme**

If you are hosting the full desktop with Terminal Server, by default the desktop environment resembles a Windows Classic desktop. By default Windows Server 2003 does not have themes enabled. You can enable themes by starting the Themes service and configuring it to start automatically. For more information about starting the Themes service, see "Configure how a service is started" in Help and Support Center for Windows Server 2003.

After you have configured the Themes service to start automatically, you can enforce a specific desktop theme or the Windows XP theme for your Terminal Server users by using the following procedure. For more information about choosing to use desktop themes with Terminal Server, see "Hosting Full Desktops with Terminal Server" earlier in this chapter.

**To load a specific theme for the desktop**

1. In the Group Policy Object Editor, navigate to User Configuration/Administrative Templates/Control Panel/Display/Desktop Themes.

2. Open the **Load a specific visual style file or force Windows Classic** setting.

3. Take one of the following actions, depending on what you are trying to achieve:

    - To force Windows Classic, enable this setting.

    - To load the Windows XP theme, enable the setting and type **%windir%\resources\Themes\Luna\Luna.msstyles** in the **Path to Visual Style** dialog box. For information about using the Windows XP theme with Terminal Server, see "Choosing Applications to Host" earlier in this chapter.

    - To load another theme or a custom theme, type the path to that theme in the dialog box.

**Restricting Access to Drives on a Terminal Server**

You can use Group Policy settings to hide and restrict access to drives on the terminal server. By enabling these settings you can ensure that users do not inadvertently access data stored on other drives, or delete or damage program or other critical system files on the C drive. The following settings are located in the Group Policy Object Editor under User Configuration/Administrative Templates/Windows Components/Windows Explorer:

- **Hide these specified drives in My Computer.** You can remove the icons for specified drives from a user's My Computer folder by enabling this setting and using the drop-down list to select the drives you would like to hide. However, this setting does not restrict access to these drives.

- **Prevent access to drives from My Computer.** Enable this setting to prevent users from accessing the chosen combination of drives. Use this setting to lock down the terminal server for users accessing it for their primary desktop.

**Configuring Start Menu and Taskbar Items**

You can use Group Policy settings to remove and to restrict access to items from the **Start** menu for Terminal Server users. The following settings are located in User Configuration/Administrative Templates/Start Menu and Taskbar:

- Enabling the **Remove Run menu from Start Menu** setting removes this menu from the **Start** menu. It also removes the **New Task** command from Task Manager and blocks the user from accessing Universal Naming Convention (UNC) paths, local drives, and local folders from the Internet Explorer address bar. While these are not the only methods for running applications, enabling this setting makes it difficult for users to access resources on the server or network.

- Enabling the **Remove Logoff on the Start Menu** setting prevents users from logging off the server from the Start menu. Enabling this setting does not prevent users from logging off using CTRL+ALT+DEL.

- Enabling the **Remove and Prevent access to the Shut Down command** prevents administrators from accidently shutting down the terminal server.

- Enabling the **Remove links and access to Windows Update** setting prevents users from attempting to download updates to Windows on to the server.

- Enabling the **Remove Favorites menu from Start Menu** setting reduces confusion for users who do not have access to the Internet.

# Planning Terminal Server User Rights and Logon

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the Remote Desktop Users group to manage user rights on a terminal server. There are also special logon configurations and Internet Explorer security configurations you can set for Terminal Server.

**The Remote Desktop Users Group**

Before users can create a remote connection with Remote Desktop, they must have the appropriate permissions. By default, members of the Administrator group can connect remotely to the server. Members of the Remote Desktop Users built-in local group also have remote logon permissions. This built-in group gives administrators control over the resources that Terminal Server users can access. Access to Terminal Server is distributed with a default set of user rights that you can change for extra security. To provide users or groups with the appropriate rights, use the Terminal Services Configuration snap-in Permissions tab to add these groups or users and to modify permissions. For more information about managing permissions, see Managing Permissions on Connections in Help and Support Center for Windows Server 2003.

**Important**

- By default, the Remote Desktop Users local group for Windows Server 2003 Terminal Server is empty.

Populate the Remote Desktop Users group with your Terminal Services users group by using the Computer Management tool. For more information about populating the Remote Desktop Users group, see "Add users to the Remote Desktop Users group" in Help and Support Center for Windows Server 2003.

To control who can add members to the Remote Desktop Users group, add this group to **Restricted Groups** by using the following procedure:

**To add the Remote Desktop Users group to Restricted Groups**

1. In the Security Templates Microsoft Management Console (MMC) snap-in, create a new template, or use an existing one.
2. In the navigation pane, right-click **Restricted Groups** in the template, click **Add Group**, and then type **Remote Desktop Users**.
3. In the details pane, double-click **Remote Desktop Users**, click **Add Members**, and select the users who you want to add to this group.

**Planning for Automatic Logon**

With Terminal Server, you can allow users to connect without entering a user name and password. You can do this on a per-user basis through the Remote Desktop Connection tool or on a per-server basis through TSCC or through Group Policy.

When you enable this, anyone with a Remote Desktop client can log on to the server. Use this connection method only in conjunction with starting users directly into a line-of-business application, especially if the application itself requires a password for access. You can enable this setting on a per-user basis through the Remote Desktop Connection tool or Group Policy or on a per-server basis through TSCC or through Group Policy. A per-server automatic logon policy is appropriate when a server is dedicated to a particular task-based application. If a server hosts more than one application, assign automatic logon on a per-user basis.

**Editing User-Specific Logon Information**

When users log on to the system, Terminal Services executes a batch file called UsrLogon.cmd in the system32 directory to make any modifications to the end-user environment and to ensure that users can run their applications correctly. If Terminal Server modifications are necessary to the user environment, you can make them by editing this file. Be aware that editing this file can affect the logon compatibility scripts that were written for applications. For more information about compatibility scripts, see "Identifying Ideal Candidates for Hosting" earlier in this chapter.

In your logon scripts, consider checking for the presence of the environment variables *clientname* or *sessionname*. These environment variables are Terminal Server–specific, and they only appear in a user environment when the user is logged on to a terminal server. You might choose to make changes to the user environment, for example omitting the execution of antivirus software, if the script determines that the environment is running on Terminal Server.

**Internet Explorer Enhanced Security Configuration**

Windows Server 2003 is installed with the Internet Explorer Enhanced Security Configuration enabled. This configuration decreases the exposure of your server to attacks that can occur through Web content and application scripts. As a result, some Web sites might not display or perform as expected. For a better user experience with Terminal Server, remove the enhanced security configuration from members of the Users group. Because these users have fewer privileges on the server, they present a lower level of risk if they are victims of an attack. This configuration allows users to browse Internet and intranet sites much as if they were using a stand-alone desktop computer.

For more information about Internet Explorer Enhanced Security Configuration settings, see "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003.

You can also configure Internet Explorer Enhanced Security Configuration through Unattended Setup. For more information, see "Enabling Terminal Server Using an Automated Installation Method" later in this chapter.

# Designing Terminal Server Installation and Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By carefully planning how you install and configure the various components of your Terminal Server solution, you can save considerable time in the installation and make management of your servers easier. Figure 4.5 describes the steps in this process. The components needed for a complete Terminal Server solution include the terminal servers themselves, the license server, and the applications that you plan to host, as well as a load-balancing solution and a Session Directory server if you plan to use them.

**Figure 4.5 Designing Terminal Server Installation and Configuration**



# Designing the Installation Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Selecting the best way to deploy your terminal servers depends on the scope of your use of Terminal Server. You can enable Terminal Server on a server that has Windows Server 2003 already installed by using Add or Remove Programs in Control Panel or the Configure Your Server Wizard. You can also deploy Windows Server 2003 with Terminal Server enabled by using an automated installation method. These options are discussed in further detail in the following sections.

# Enabling Terminal Server Manually

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are deploying just one terminal server, or if you are upgrading to Windows Server 2003, it might be easiest to enable Terminal Server on a server that already has the operating system installed and configure the server individually. You can accomplish this through the Configure Your Server Wizard or through Add or Remove Programs in Control Panel. For information about how to use the Configure Your Server Wizard, see "Install Terminal Server" in Help and Support Center for Windows Server 2003. To enable Terminal Server through Add or Remove Programs, open Add or Remove Programs from Control Panel, or in the **Run** dialog box, type **control appwiz.cpl**, and click Add/Remove Windows Components.

**Note**

- You can also install Terminal Server Licensing through Add or Remove Programs.

**Internet Explorer Enhanced Security Configuration Considerations**

If you choose to install Terminal Server by using Add or Remove Windows Components in Add or Remove Programs, a warning appears if the installation will proceed with Internet Explorer Enhanced Security Configuration enabled, and allows you to choose whether you want to proceed with this configuration. It is recommended that you choose No and return to Add or Remove Windows Components and disable this configuration to allow users to browse the Internet.

If you install Terminal Server through the Configure Your Server Wizard, no warning appears. When the Manage Your Server window appears after you have added the Terminal Server role, a message appears that indicates the status of the Internet Explorer Enhanced Security Configuration, and gives instructions about how to change it.

For more information about Internet Explorer Enhanced Security Configuration settings, see "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003.

# Enabling Terminal Server Using an Automated Installation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

One way to ensure that your servers have a consistent configuration is by using an automated installation method to install Windows Server 2003 with Terminal Server enabled. The following automated installation methods are available in Windows Server 2003:

- Unattended Setup
- Sysprep
- Remote Installation Services (RIS)

If you are deploying Terminal Server with several configurations, you can enable Terminal Server on all of the servers by using an automated installation method, and then configure the servers individually after the installation.

You cannot configure settings such as Group Policy settings within the automated installation itself. However you can set up the installation to run scripts that can configure these settings and perform a wide variety of other configurations after the setup is complete. For more information about the options for automated installations, see the design chapters for the various automated installation methods in *Automating and Customizing Installations* of this kit and the "Microsoft Windows Corporate Deployment Tools User's Guide" in the \Support\Tools folder of the Windows Server 2003 operating system CD.

**Tip**

- If you deploy Terminal Server for use in a load-balanced farm, an automated installation method is an easy way to ensure an identical configuration on all the servers in the farm.

**Unattended Setup**

Unattended Setup uses an answer file to automate the setup of Windows Server 2003. It is the only way to automate an upgrade to Windows Server 2003. Unattended Setup is a good choice for installing Terminal Server if, for example, you are planning to deploy terminal servers with a few applications for different purposes. For more information about designing an Unattended Setup, see "Designing Unattended Installations" in *Automating and Customizing Installations* of this kit.

For Terminal Server, you need to make entries in a few of sections of the answer file (Unattend.txt) in order to complete the automated installation:

**[Components]**

The [Components] section contains entries for installing the components of Windows XP or Windows Server 2003 operating systems. These components use the values **On** or **Off** to install or not install the component.

**Note**

- The entries in this section can also be set in the Sysprep Factory mode, **sysprep -factory**, which is used for the preinstallation of Windows in an OEM manufacturing environment.

For Terminal Server, you can install the following components:

- **LicenseServer**. Set this to **On** if you are setting up the server as a license server for Terminal Server.
- **TerminalServer**.Set this to **On** if you are setting up the server as a terminal server.

  **Important**

  - When you install Terminal Server through Unattended Setup, it is installed by default with Internet Explorer Enhanced Security Configuration enabled for administrators, and disabled for users. To change Internet Explorer

Enhanced Security Configuration during Unattended Setup, you can configure IEHArdenUser and IEHArdenAdmin, also found in the [Components] section of the answer file. For more information about Internet Explorer Enhanced Security Configuration settings, see "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003.

- **TSWebClient**. Set this to **On** to install the ActiveX control and sample pages for hosting Remote Desktop client connections over the Web.

- **IEHardenUser.** This is set to **Off** by default, which removes the Internet Explorer Enhanced Security Configuration from members of the Restricted Users and Guests groups. For more information about Internet Explorer Enhanced Security Configuration settings, see "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003.

- **IEHardenAdmin**. This is set to **On** by default, which applies the Internet Explorer Enhanced Security Configuration to members of the Administrators and Power Users groups. For more information about Internet Explorer Enhanced Security Configuration settings, see "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003.

**[TerminalServices]**
If you set the TerminalServer entry to **On** in the answer file, you also need to set the following entries in the [TerminalServices] section:

**Note**

- You can also set these settings through the TSCC. If these settings are changed in TSCC after the installation, the new settings take precedence.

- **AllowConnections**. When set to **1** this entry enables Terminal Server by allowing connections. The default for this setting is **0**. You can also set this through Group Policy.

- **LicensingMode**. This setting configures how Terminal Services manages CALs — **PerDevice** or **PerUser**. The default for this setting is **PerDevice**. For more information about choosing a licensing mode, see "Choosing the Licensing Model" earlier in this chapter.

- **PermissionsSetting**. This setting allows you to choose the security mode for Terminal Server users during terminal server sessions. This entry maps to the Permission Compatibility setting in TSCC. The default setting for this entry is **0**, which maps to Full Security in TSCC.

  **Important**

  - Only set PermissionSetting to **1**, which maps to Relaxed Security in TSCC, if you are running older applications and your testing has determined that your applications will not run properly with the setting at **0**. For more information, see "Designing the Terminal Server Configuration" later in this chapter.

**Sysprep**
Sysprep is an image-based method that you can use for clean installations of Windows Server 2003 with Terminal Server enabled. Using Sysprep is the quickest way to install Terminal Server along with many applications. Sysprep is a good choice for installing Terminal Server if, for example, you are deploying a terminal server farm.

If you plan to use a single image for many server configurations, you can use Sysprep in factory mode. For Terminal Server, you can configure the entries in the [Components] section of the Winbom.ini file for use with Sysprep in factory mode. These settings are the same as for the [Components] section of the Unattend.txt file for Unattended Setup discussed earlier in this section. For more information about using Sysprep for automated installations, see "Designing Image-based Installations with Sysprep" in *Automating and Customizing Installations* of this kit.

**RIS**
RIS requires the deployment of a network infrastructure specifically for its use and requires that you install Active Directory in your organization. Plan to use RIS for the automated installation of Terminal Server only if you are planning to use it throughout your organization for installing Windows Server 2003. If you do plan to use RIS for installing Terminal Server, you

need to be sure to image a terminal server that does not already have licenses. You need to install licenses on the servers after the installation. For more information about using RIS for automated installations, see "Designing RIS Installations" in *Automating and Customizing Installations* of this kit.

# Upgrading to Windows Server 2003 Terminal Server

Updated: March 10, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The best way to install Terminal Server is by performing a clean installation. However, if you already have a Windows 2000 Terminal Services in Application Mode or Windows NT 4.0 Terminal Server Edition infrastructure in place, you might want to perform an upgrade. There are a number of other reasons why you might perform an upgrade, for example, if you are transitioning gradually to Windows Server 2003 or if you want to retain the ability to use your older software and device drivers. If you are upgrading your terminal servers from Windows 2000 to Windows Server 2003, you must take into consideration a number of changes and new requirements with Windows Server 2003.

**Note**

- When performing an operating system upgrade on a terminal server, the security template that is applied does not reset the access control lists (ACLs). This is in contrast to a non-Terminal Server upgrade, which resets the ACLs. If you are moving from Windows NT 4.0 Terminal Server Edition or Windows 2000 in Relaxed Security mode to Windows Server 2003 Terminal Server, consider performing a clean installation.

**Upgrading Terminal Server Licensing**

With Windows Server 2003, you must use a license server that is running Windows Server 2003. You can still issue licenses to Windows 2000 Terminal Server with a license server that is running Windows Server 2003, if you plan to gradually upgrade your terminal servers. Also, with Windows Server 2003 you no longer need to host your license server on a domain controller. For more information about choosing a host server for Terminal Server Licensing, see "Planning the License Server" earlier in this chapter.

Table 4.1 summarizes the licensing issues to be aware of when upgrading from Windows 2000 Terminal Services to Windows Server 2003 Terminal Server.

**Table 4.1 Licensing Issues for Upgrading Terminal Server**

| If you are... | You need to... |
|---|---|
| Gradually upgrading your terminal servers to Windows Server 2003 from Windows 2000. | Upgrade your license server to Windows Server 2003 at the same time that you upgrade the first terminal server. You can also ensure that the license server issues the appropriate CAL for Windows 2000 or Windows Server 2003 through Group Policy. For more information, see "Designing the Terminal Server Configuration" later in this chapter. |
| Performing an in-place upgrade of the domain controller that hosts Terminal Server Licensing. | Consider keeping your licenses on the same server whether or not you plan to host your license server on a domain controller. You can demote the domain controller either before or after you upgrade the server to Windows Server 2003 by using the Active Directory Installation Wizard (dcpromo.exe).<br><br>If you upgrade your license server, you must reactivate the license server.<br><br>For information about upgrading a domain controller, see "Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" in *Designing and Deploying Directory and Security Services* of this kit. |

| | |
|---|---|
| Performing a clean installation of the server that hosts Terminal Server Licensing. | Deploy Terminal Server Licensing on a server that runs Windows Server 2003, migrate the client access licenses (TS CALs) to the new license server, and then deactivate the previous license server by doing the following:<br><br>1. Determine the server on which you want to install Terminal Server Licensing and whether you want it to serve as an Enterprise license server or as a Domain license server. For more information, see Terminal Server license server roles (http://go.microsoft.com/fwlink/?LinkId=62981).<br><br>2. Install and then activate Terminal Server Licensing on the new Windows Server 2003 license server. For step-by-step instructions, see Install Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkId=62988) and Activate a Terminal Server License Server (http://go.microsoft.com/fwlink/?LinkId=62991).<br><br>3. Install the same number and type of TS CALs on the new license server that you had on the Windows 2000 license server. To do this, call the Microsoft Clearinghouse. When you call the Clearinghouse, make sure that you have your License Purchase Agreement information (documentation of your purchased TS CALs) readily available to facilitate the re-issuance of the CALs. For step-by-step instructions on how to locate the correct telephone number for the Clearinghouse, see Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkId=62994).<br><br>4. Make sure that your terminal server can discover the new license server. For step-by-step instructions, see "verify whether the terminal server can discover the license server" in The terminal server cannot locate the license server (http://go.microsoft.com/fwlink/?LinkId=62995). Clients that received CALs from the previous license server will continue to operate as normal until their CALs expire. When the previously issued CALs expire, new licenses will be requested for these clients from the new license server. For information about license issuance and expiration, see Recovering client access licenses on a Terminal Server license server (http://go.microsoft.com/fwlink/?LinkId=62996).<br><br>**Important**<br>To be discoverable in a domain, a license server must be a domain controller. If you did not install Terminal Server Licensing on a domain controller, then, on the terminal servers, configure the new license server as a preferred license server, so that it can be discovered in the domain. For step-by-step instructions, see "Set a preferred license server" in The terminal server cannot locate the license server (http://go.microsoft.com/fwlink/?LinkId=62995). For detailed information about the license server discovery process, see "Licensing Service Discovery" in Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkId=26220).<br><br>5. Deactivate the previous license server. For step-by-step instructions, see Deactivate a Terminal Server license server (http://go.microsoft.com/fwlink/?LinkId=62998). |

| | |
|---|---|
| Redeploying Terminal Server Licensing to a new server. | Deploy Terminal Server Licensing on a server that runs Windows Server 2003, migrate the client access licenses (TS CALs) to the new license server, and then deactivate the previous license server by doing the following:<br><br>1. Determine the server on which you want to install Terminal Server Licensing and whether you want it to serve as an Enterprise license server or as a Domain license server. For more information, see Terminal Server license server roles (http://go.microsoft.com/fwlink/?LinkId=62981).<br><br>2. Install and then activate Terminal Server Licensing on the new Windows Server 2003 license server. For step-by-step instructions, see Install Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkId=62988) and Activate a Terminal Server License Server (http://go.microsoft.com/fwlink/?LinkId=62991).<br><br>3. Install the same number and type of TS CALs on the new license server that you had on the Windows 2000 license server. To do this, call the Microsoft Clearinghouse. When you call the Clearinghouse, make sure that you have your License Purchase Agreement information (documentation of your purchased TS CALs) readily available to facilitate the re-issuance of the CALs. For step-by-step instructions on how to locate the correct telephone number for the Clearinghouse, see Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkId=62994).<br><br>4. Make sure that your terminal server can discover the new license server. For step-by-step instructions, see "verify whether the terminal server can discover the license server" in The terminal server cannot locate the license server (http://go.microsoft.com/fwlink/?LinkId=62995). Clients that received CALs from the previous license server will continue to operate as normal until their CALs expire. When the previously issued CALs expire, new licenses will be requested for these clients from the new license server. For information about license issuance and expiration, see Recovering client access licenses on a Terminal Server license server (http://go.microsoft.com/fwlink/?LinkId=62996).<br><br>**Important**<br>To be discoverable in a domain, a license server must be a domain controller. If you did not install Terminal Server Licensing on a domain controller, then, on the terminal servers, configure the new license server as a preferred license server, so that it can be discovered in the domain. For step-by-step instructions, see "Set a preferred license server" in The terminal server cannot locate the license server (http://go.microsoft.com/fwlink/?LinkId=62995). For detailed information about the license server discovery process, see "Licensing Service Discovery" in Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkId=26220).<br><br>5. Deactivate the previous license server. For step-by-step instructions, see Deactivate a Terminal Server license server (http://go.microsoft.com/fwlink/?LinkId=62998). |

**Group Policy Upgrade Issues**

Windows Server 2003 offers more Terminal Server–specific Group Policy settings than did earlier versions of the Windows operating system. Your initial Group Policy settings vary depending on whether you perform an upgrade or clean installation. When you upgrade your terminal servers to Windows Server 2003, the new Group Policy settings are left Not Configured rather than set to their default. When you perform a clean installation, the default settings are used. Carefully review the Group Policy settings for Terminal Server to choose the best settings for your environment. For more information, see "Designing the Terminal Server Configuration" later in this chapter.

Also new with Windows Server 2003 is the Remote Desktop Users group. Users are required to be members of this built-in local group in order to gain access to the terminal server. This group is empty by default. For more information, see "Planning Terminal Server User Rights and Logon" earlier in this chapter.

**Permission Compatibility**

If you are upgrading from Windows 2000 Terminal Services or Windows NT 4.0 Terminal Server Edition, and are planning to host the same applications with Windows Server 2003, be aware that, in order for them to work, some older applications require access to system resources, such as the registry, for members of the Users group on a terminal server. If your older applications require this type of access, you can control the Permission Compatibility setting in TSCC on the **Server**

**Configuration** tab for application compatibility. This setting is set to Full Security by default, which restricts access to system resources.

**Important**

- Do not set this setting to Relaxed Security unless you have performed thorough testing and have determined that your applications will not work properly otherwise. Relaxed Security mode gives users access to system components on the server, such as the ability to modify the system32 directory where operating system files are stored, access to the Program Files directories, and read/write access to registry settings in HKEY_LOCAL_MACHINE.

For more information about configuring this setting, see "Designing the Terminal Server Configuration" later in this chapter.

**Automated Installation**

If you are planning to use an automated installation method to upgrade your terminal servers, you can do this only by using Unattended Setup. For more information about using Unattended Setup with Terminal Server, see "Designing Application Installation" later in this chapter. For more information about installing Windows Server 2003 by using an automated installation method, see *Automating and Customizing Installations* of this kit.

**Internet Explorer Enhanced Security Configuration**

When you upgrade from Windows 2000 Terminal Services, Internet Explorer Enhanced Security Configuration is enabled for administrators and disabled for users by default. With this configuration, users are able to browse and download files from the Internet without restriction, but administrators cannot do so. For more information about Internet Explorer Enhanced Security Configuration settings, see "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003.

# Designing Application Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are several ways in which to install the applications you plan to host on the terminal server. The size of your deployment and whether you have Active Directory installed are factors in deciding which method to use.

**Installing Applications Manually**
With manual installations, in order for Terminal Services to replicate the necessary registry entries or .ini files for each user, the user must install the application by using **Add or Remove Programs** in Control Panel. You can also install applications from the command line by using the **change user /install** command, but using **Add or Remove Programs** is preferable.

**Installing Applications Using Group Policy and Windows Installer**
A recommended way to distribute applications to a server in a Terminal Services farm is to use Group Policy. Active Directory is required for this. By separating the terminal servers into their own OU, you can create a Group Policy object that is linked to only that OU, and you can then assign Windows Installer (.msi) packages to it. Assigned applications are installed on the server when the server is turned on. For more information, see "Assigned and published programs" in Help and Support Center for Windows Server 2003. For more information about deploying software using Group Policy, see "Deploying and upgrading software" in Help and Support Center for Windows Server 2003 and "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit.

**Note**

- Terminal Server cannot accept published programs because publishing occurs on a per-user basis. Additionally, you must assign programs on a per-computer basis, rather than on a per-user basis.

Some applications require a transform file (.mst) when when you install them by using Windows Installer. *Transform files* are modifications to .msi packages that you create to instruct Windows Installer to install the application and all the needed components locally on the terminal server. Test your application installation to determine if a transform file is needed. For more information about .mst files, see "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit.

To create and install applications using an .mst file, see the Office 2000 Resource Kit at the MSDN Library link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. Search for the topics under "Installing Office in a Windows Terminal Server Environment" in the Office 2000 Resource Kit on this Web site.

**Note**

- Microsoft® Office XP installs and runs on a terminal server correctly without requiring a transform file.

**Running Compatibility Scripts**
If it is necessary to run application compatibility scripts for any of the applications that you are hosting with Terminal Server, you must run them after you install the application, but before you restart the server. For more information about compatibility scripts, see "Identifying Ideal Candidates for Hosting" earlier in this chapter and the Optimizing Applications for Windows 2000 Terminal Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

When you run the first compatibility script, you are asked to select a drive letter to be used by this script and all future compatibility scripts. This drive letter must be the same as the users' home directory. For more information about using home directories with Terminal Server, see "Using Home Directories with Terminal Server" earlier in this chapter.

# Supporting Multilingual and International Users

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 with Terminal Server enabled offers multilingual support. Using this capability, Terminal Services can simultaneously serve users in any language that is installed on the server. Terminal servers that are running Windows Server 2003 can take advantage of the multilingual user interface (MUI). This interface allows you to install multiple languages on the system and configure them on a per user basis. You must load MUI to take advantage of this functionality, which simplifies the deployment of Terminal Services within a multinational organization. It is also recommended that you install the keyboard drivers for any language-specific keyboards that you use in your organization as part of the Terminal Services installation. You can support your users around the world who can understand and work in English by hosting all users on the International English version of Windows Server 2003. This ensures that the organization is not in violation of United States export laws regulating strong encryption. However, using the International English Version might be inadequate if the organization must provide support for different languages. By default, Terminal Server installs all the available keyboard layouts, including support for Asian non-Input Method Editor keyboards.

While loading the MUI, install all of the languages that you expect users to need. The installation copies the appropriate language DLLs and Help files. By using Regional and Language Options in Control Panel, users can select their default language and keyboard settings. Regional and Language Options also includes such settings as date and currency formatting. Because this setting is stored in the user profile, individual users can adjust the settings to match their locality. If a user has a roaming profile that specifies a different language than the language that is loaded in the user profile, then the system uses the default user language from the profile.

**Important**

- Roaming users should not use Folder Redirection with MUI. If they do, they might create multiple language versions of My Documents and other per-user folders on the computer. If the user interface (UI) language files that are needed to support the roaming user's default UI language have not been installed on the computer, the localized names of the new folders might not display correctly.

For more information about MUI, see the MultiLanguage Version link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Designing the Terminal Server Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure Terminal Services by using TSCC Microsoft Management Console (MMC) snap-in, Group Policy, or WMI. Some server configurations settings for Terminal Services, for example the Sessions settings, are also configurable for the user. Because the server settings can override the user settings, it is recommended that you set these settings on the server whenever possible.

# Choosing a Terminal Server Configuration Tool

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Choose a configuration tool based on your level of permission, the tasks you want to accomplish, and the level at which you want to apply settings. If you are a domain administrator configuring global settings on a number of terminal servers, it is recommended that you use Group Policy. WMI is a good choice for automating the configuring of global settings on a number of terminal servers if you are not a domain administrator and if you are familiar with scripting. However, you must be a local administrator on every computer you want to configure using WMI. You can also invoke WMI to configure settings automatically after you install Windows Server 2003 with an automated installation method. For more information about configuring automated installation to invoke scripts, see the design chapter in *Automating and Customizing Installations* that corresponds to the automated installation method you have chosen to use. The TSCC snap-in is an easy-to-use tool for configuring a single terminal server, and is useful for configuring unique settings for a particular server. Table 4.2 summarizes these tools, and the following sections describe each tool in further detail.

**Table 4.2 Benefit and Restriction Comparison for Terminal Server Configuration Tools**

| Tool | Benefits | Restrictions |
|---|---|---|
| Group Policy | Can centrally configure terminal servers and Terminal Server users by applying policies to OUs.<br>Always overrides configurations set by using other tools. | Administrator must be a domain administrator to apply Group Policy settings to OUs.<br>Must have Active Directory in place. |
| WMI Terminal Server provider | Can configure many terminal servers or Terminal Server users using scripts. | Administrator must know how to write scripts. |
| Terminal Server Connection Configuration snap-in | Can configure unique per-server settings.<br>Some configurations only available in TSCC snap-in. | Can be overwritten by Group Policy settings.<br>Can be applied only to a single terminal server and its users.<br>Cannot be used to configure a remote server. |

Your choice of tools might also depend on your server environment or number of connections. Take the following information into consideration when making your choice:

- In an operating system environment where only Windows Server 2003 is run, you can use Terminal Server Group Policy settings to configure all settings that apply across an OU. You can also configure individual Windows Server 2003 operating systems using Group Policy on the local Group Policy object.

- In server environments where several different versions of the Windows operating system are run, you might need to use a combination of tools. For example, you might configure the Windows Server 2003 operating systems with Group Policy, while using TSCC to configure servers that are running earlier versions of Windows.

- If you have two or more connections on the same computer, and you want to configure each connection differently, you cannot use Group Policy. Instead, use the TSCC tool, which allows you to configure Terminal Services settings on a per-connection basis.

**Using Group Policy**

You can use Group Policy to configure Terminal Server connection settings, set user policies, configure terminal server clusters, and manage Terminal Server sessions. You can apply Group Policy settings for users of a computer through the Remote Desktop Users group, for individual computers through local Group Policy, or for groups of computers through a Terminal Server OU. To set local policies for users of a particular computer, you must be an administrator for that computer. To set policies for an OU in a domain, you must be an administrator for that domain. Settings that are specific to Terminal Server are located under Computer Configuration/Administrative Templates/Windows Components/Terminal Services. For more information, see "Configuring Terminal Services with Group PolicyConfiguring Terminal Services with Group Policy" in Help

and Support Center for Windows Server 2003.

**Note**

- There are many new Group Policy settings with Windows Server 2003. If you are upgrading your terminal servers from Windows 2000, the new Group Policy settings will be set to Disabled by default. For new Windows Server 2003 deployments, the defaults are stated in the "Designing Terminal Server Connection Configurations" section later in this chapter.

**Using the Terminal Services WMI provider**

The Terminal Services WMI provider allows administrators to create custom scripts for configuring, managing, and querying terminal servers. It contains properties and methods that can perform the same tasks as Terminal Services configuration tools and command-line tools, but remotely and through scripted applications. A *provider* is an architectural element of WMI that extends the WMI schema of classes to allow WMI to work with new types of objects. The Terminal Services provider defines classes for querying and configuring Terminal Services. The Terminal Services WMI provider is defined in *systemroot*\System32\Wbem\tscfgwmi.mof. For more information about WMI, see "Configuring Terminal Services with WMIConfiguring Terminal Services with WMI" in Help and Support Center for Windows Server 2003.

**Note**

- Configuration settings applied with the Terminal Services WMI provider operate in the same order of precedence as they do if applied with the corresponding configuration tool. In general, Group Policy settings always override settings applied with WMI. For more information about using WMI, see the TechNet Script Center link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Using the TSCC snap-in**

With the TSCC snap-in you can configure the RDP connection parameters and connection permissions for the terminal server. You can apply settings only to a single terminal server and its users, however, and you cannot use TSCC to configure a remote server. In addition, there are several settings that you can set only by using TSCC. For more information, see "Configuring Terminal Services with TSCC" in Help and Support Center for Windows Server 2003.

# Designing Terminal Server Connection Configurations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following guidelines to design the configuration of the connections to your terminal servers. If you are using TSCC to configure a single server, you can use the Terminal Services Connection Wizard to configure these settings when you create a new connection. For more information about configuring Terminal Server with TSCC when you create a new connection, see "Make a new connectionMake a new connection" in Help and Support Center for Windows Server 2003. You can also use Group Policy or WMI to configure the connections to many terminal servers. For a job aid to assist you in recording your Terminal Server Group Policy configuration decisions, see "Group Policy Configuration Worksheet" (SDCTS_2.xls) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Group Policy Configuration Worksheet" on the Web at http://www.microsoft.com/reskit).

**Data Encryption**

You can assign data transfer encryption levels between the Remote Desktop client and Terminal Server by using either Group Policy or TSCC. The default RDP encryption level is **Client Compatible**. You can choose one of the following possible choices for encryption:

- **FIPS Compliant.** Encrypts traffic between client and server to meet the Federal Information Processing Standard 140-1 (FIPS 140-1). Use this level when Terminal Services connections require the highest degree of encryption, such as those required by the U.S. federal government.

- **Client Compatible.** With Client Compatible encryption, traffic between the client and the server is encrypted using the RC4 algorithm and the strongest key the client supports (40-bit, 56-bit, or 128 bit). The server negotiates with the client to determine the key strength on connection, however the server does not accept non-encrypted client connections.

- **High.** Traffic in both directions is encrypted using the RC4 algorithm and a 128-bit key only. If a client does not support 128-bit encryption, it is not permitted to connect.

- **Low.** Traffic from the client to the server only is encrypted, at the strongest key that the client supports. This can improve performance on the client because the client does not have to decrypt the screen update data coming from the server. The client still encrypts the keystroke and mouse data that it sends to the server. This also allows you to use products to improve performance over a WAN, for example to use between a branch and a home office. Use this setting only if you are planning to use these products. A malicious user can monitor documents and data coming from the server over the link if this setting is used.

On the **General** tab of TSCC or on the Data Encryption property page of the Terminal Services Connection Wizard, the **Use standard Windows authentication** check box is cleared by default. If you select this check box, lower security authentication mechanisms are permitted.

**Logon Settings**

By default, users are allowed access to the terminal server using the information that they provided to log on to their remote desktop client. For a more secure system, you can require users to provide logon credentials to access the terminal server.

You can allow access to the server based on the credentials provided here by clicking the **Always use the following logon information** radio button on the **Logon Settings** tab of TSCC or by enabling the **Always prompt client for password upon connection** Group Policy setting located in the Encryption and Security folder under Terminal Services in Computer Configuration. Use one of these options only if you are hosting an application that requires a password for access. For more information, see "Planning Terminal Server User Rights and Logon" earlier in this chapter.

**Remote Procedure Call (RPC) Security Policy**

You can enable this Group Policy setting to allow Terminal Server to accept only authenticated and encrypted requests. It is recommended that you enable this setting for increased security.

**Sessions**

You can set Terminal Server time-out and reconnection settings on the server to help manage the number of sessions held by a terminal server at any one time, to help manage unreliable connections for remote users, or to reduce the impact on the CPU of many users logging on to the server at the same time. You can set most of these settings by using TSCC or Group Policy. Exceptions are noted.

### Note

- You can also set these settings per user through Group Policy User Configuration and through Active Directory Users and Computers user properties. However, the Group Policy computer settings listed here take precedence over user settings.

**End a disconnected session**
You can set a limit on the time that a disconnected session continues to exist on the server. There can be many reasons why a session becomes disconnected, for example if a user's computer fails or if the user places the session into a disconnected state in order to access the same session from another location. The programs and processes that the user had running before the disconnection continue to run during a disconnected session. Because a user can have a disconnected session running on the server without realizing it, it is best to set a limit on how long disconnected sessions continue to run on the server. However, by setting this as high as possible you can achieve better server performance by reducing the additional CPU usage needed when a user reconnects to the server.

**Active session limit**
You can set a limit on how long a user can maintain an active session with the server. If you choose **Never**, the server allows an active session to continue forever.

**Idle session limit**
You can set a limit on how long an idle session remains open. An idle session occurs if there has been no mouse or keyboard activity for a certain period of time. This can indicate that the user has stepped away from the computer, presenting someone else with the opportunity to use his or her session. When a session has been idle for more time than you have specified, the user is notified and given two minutes to place the session back into an active state. If the two minutes elapse, the server disconnects or ends the session, depending on the settings you choose. Just as with choosing a timeframe for ending a disconnected session, it is important to understand users' work patterns and needs when choosing a limit for idle sessions.

### Note

- In a load-balanced farm, an idle session cannot be moved to a different server within the farm.

**Session reconnection**
By default, users can reconnect to a disconnected session from a computer other that the one from which they originally connected to the session. However, if a user connects from a Citrix ICA client, you can restrict the user to reconnecting only from the computer that originally connected to the session. This setting does not apply to Windows clients.

**Session limit behavior**
If you choose to enforce session limits for your users, you can ensure that users do not choose to have their sessions end when the connection to the server is broken for whatever reason. This way, you can be sure that all users can pick up where they left off in the event of a server failure, thereby reducing the loss of work and helpdesk incidents. If you choose to have sessions disconnect, you can use TSCC or Group Policy to specify the amount of time the session remains in a disconnected state.

**Shut down options**
You can use the following two Group Policy settings, found in the Terminal Services folder under Computer Configuration/Administrative Templates/Windows Components, to remove items from the **Start** menu and **Shut Down** dialog box so that users cannot use certain methods to disconnect or log off from their session:

- Enable the **Remove Windows Security item from the Start menu** policy to prevent users from unintentionally logging off of Terminal Server.

- Enable the **Remove Disconnect option from Shut Down dialog** policy to prevent users from using this method to disconnect from Terminal Server.

**Automatic reconnection**
You can allow a session to automatically reconnect to Terminal Server if the network connection is lost.

**User Environment**
Use the following settings to control the users' desktop environment.

**Launch application on connection**
If you are hosting a single application with Terminal Server, for example a line-of-business application or an application for

task workers who use the server for only one thing, you can have that application start automatically when the user logs on. This eliminates the possibility of the user running unauthorized applications on the server or accessing other parts of the server or the network through the server. For information about how to start an application on connection, see "Specify a program to start on session connectionSpecify a program to start on session connection" in Help and Support Center for Windows Server 2003. For more information about automatic logon, see "Planning Terminal Server User Rights and Logon" earlier in this chapter. You can configure this setting by using Group Policy (which you can apply to both computers and users), TSCC, and for users through the Remote Desktop Connection tool.

**Desktop wallpaper**

By default, sessions connecting to Windows Server 2003 Terminal Server do not display desktop wallpaper. For sessions connecting to servers running previous versions of Windows server operating systems and clients running Windows XP Professional or earlier, you can accomplish this by using Group Policy. For more information, see "Enforce removal of Remote Desktop wallpaperEnforce removal of Remote Desktop wallpaper" in Help and Support Center for Windows Server 2003.

### Note

- In Windows 2000 you can disable desktop wallpaper on the client from the **Environment** tab of the Terminal Services Connection Configuration tool.

**Remote Control**

You can use Remote Control to control or troubleshoot a user's session from a remote location. You can configure this setting by using Group Policy settings (which you can apply to both computers and users) or TSCC. If you choose to enable this in your organization, you can configure this setting in the following ways:

- Full Control with user's permission
- Full Control without user's permission
- View Session with user's permission
- View Session without user's permission

It is recommended that you configure this setting so that the user's permission is required to allow another person to access their computer through Remote Control. Keep in mind that even if you require a user to give permission, the user can choose to allow anyone who has acquired the correct permissions to access his or her computer. Also, some countries and regions have laws that do not allow the **View Session with user's permission** setting. If your organization has offices in several countries and regions, check the local laws before configuring Remote Control in this way. For more information about configuring sessions for Remote Control, see "Configure remote control settingsConfigure remote control settings" in Help and Support Center for Windows Server 2003.

**Client Settings**

You can use the settings discussed in this section to control certain aspects of the user experience and allow users to perform certain operations through their desktop computer rather than through the server.

**Client/Server data redirection**

You can use data redirection with Terminal Server to enable users to access and use resources from their desktop computers rather than from the terminal server. The most notable of these resources is printing, but you can also enable drive, audio, smart card, and clipboard redirection to the client computer. In general, restricting user's options to only those required to do their jobs can minimize the likelihood of introducing a security issue. You can configure most of these settings through Group Policy (which you can apply to computers only) or TSCC. Exceptions are noted.

### Note

- You can also configure redirection for disk drives, printer, and serial ports by using the Remote Desktop Connection tool on the client. For more information, see "Configuring Remote Desktop Connection" later in this chapter.

Table 4.3 summarizes the data redirection settings.

**Table 4.3 Data Redirection Settings**

| Data redirection type | Characteristics |
|---|---|
| Time zone | Only configurable by using Group Policy.<br><br>By default the session time zone is the same as the time zone of the terminal server. This can be an issue if you are using Terminal Server for remote or mobile users, especially when your line-of-business applications (for example, financial applications) have time dependencies. |
| Clipboard | By default, you can copy and paste between the terminal server and the Remote Desktop client. Disable this ability if you have sensitive data on the terminal server that should not be shared outside of the application in which it is used. |
| Smart card | Only configurable by using Group Policy.<br><br>By default the ability to log on to the Remote Desktop client is allowed. For more information about using smart card with Terminal Server, see the "Planning Network Security Components" earlier in this chapter. |
| Audio | By default, users cannot play audio on the Remote Desktop client. The sound plays on the server rather than the client computer. If you enable this setting, users can specify on the Remote Desktop Connection tool whether to play audio at their computer or the server, or to not have the sound play at all. |
| Serial port | By default, Terminal Server allows users to redirect data to peripherals attached to the serial (COM) port. Disable this capability unless there is a requirement for it. This prevents users from printing or copying sensitive data stored on the terminal server, and reduces vulnerabilities to security threats that could access your computer through these ports. |
| Client printer | By default, users can redirect print jobs to a printer attached to their client computer. Unless users need to print to a local printer, disable this capability so that users cannot print or copy sensitive data stored on the terminal server. |
| Parallel port | By default, Terminal Server allows users to redirect data to peripherals attached to the parallel (LPT) port. Disable this capability unless there is a requirement for it. This prevents users from printing or copying sensitive data stored on the terminal server, and reduces vulnerabilities to security threats that could access your computer through these ports. |
| Drive | By default, Terminal Server allows users to redirect data to the drives on the client computer. Unless there is a requirement for this, you should disable this capability so that users cannot copy sensitive data stored on the terminal server onto their local computer. |
| Default printer | By default, Terminal Server designates the client default printer as the default printer in a session. Unless users need to be able to print to a local printer, disable this capability so that users cannot print or copy sensitive data stored on the terminal server. |

**Color depth**
You can reduce or increase the maximum color depth depending on your bandwidth and fidelity requirements (greater color depth requires more bandwidth and resources on the terminal server).

**Number of connections**
By default, an unlimited number of sessions are permitted on the terminal server. Restricting the number of sessions improves the performance of Terminal Server because fewer sessions are demanding system resources.

You can configure this setting in the Terminal Services folder of the Group Policy Object Editor. In TSCC, you can configure the number of settings on the **Network Adapter** tab. You can also select the network adapter you want to use for the RDP connection traffic.

**Permissions**
You can use TSCC to change your permissions lists by adding and removing users and groups. You can also customize the permissions for users or groups on a per-connection basis. It is recommended that you give permissions to as few users and groups as is necessary, and to give those users and groups the lowest level permissions necessary for them to do their jobs. For more information about how to set permissions, see the topics under "Managing Permissions on ConnectionsManaging Permissions on Connections" in Help and Support Center for Windows Server 2003.

# Designing Server Setting Configurations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following guidelines to design the configuration of your terminal servers.

**Keep-Alive Connections**
Configure this Group Policy setting to Enabled to ensure that the session state remains consistent with the client state. Use this setting only if you are having problems with users who cannot reconnect.

**Temporary Folders**
These settings can be set in the Server Settings folder of TSCC or in the Temporary Folders node of the Terminal Services administrative template in the Group Policy Object Editor. In TSCC, both of these settings are configured Yes by default, which means that temporary folders are used per session and deleted upon logging off. It is recommended that you keep the default settings unless there is a compelling business reason to change them.

**Delete or retain temporary folders when exiting** It is recommended that you configure your servers so that the data stored in the temporary folder for user sessions is deleted when users log off. This provides added security by eliminating a point of access for user data. It also provides a way to manage the load on the server, because temporary folders tend to grow quickly in size in a multiuser environment. If you are using Group Policy to set this setting (set to Disable to delete temporary folders), you must also configure the server to use per-session temporary folders.

**Use separate temporary folders for each session** This setting keeps each session's temporary folders in a separate folder, which enables you to configure the server to delete temporary folders when a single user logs off without affecting other users' sessions.

**Active Desktop**
You can restrict users from using Active Desktop by using TSCC or Group Policy under User Configuration/Administrative Templates/Desktop/Active Desktop. Active Desktop allows users to choose JPEG and HTML wallpaper, both of which can affect performance for the user because of the amount of data that needs to transfer from the server to the desktop. The default setting is Enable. It is recommended that you disable Active Desktop.

**Session Directory Settings**
If you are load balancing several terminal servers and using Session Directory, you can use TSCC or Group Policy to configure the servers to use Session Directory. Session Directory is a database that tracks user sessions that are running on load-balanced terminal servers. For more information about using Session Directory (including configuring Session Directory), see "Load Balancing Terminal Servers" earlier in this chapter. For more information about Session Directory, see the "Session Directory and Load Balancing Using Windows Server 2003 Terminal Server" white paper at the Terminal Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Licensing**
Use the following settings to configure the license server for Terminal Server. Most of these settings can be configured only through Group Policy. Exceptions are noted.

**Licensing Mode** You can set the licensing mode to Per Device or Per User through the TSCC Server Settings folder. For more information about licensing modes for Terminal Server, see "Choosing the Licensing Model" earlier in this chapter.

**License Server Security Group** Enabling this Group Policy setting and applying it to your license server creates a local group called Terminal Services Computers. The license server issues licenses only to the terminal servers in this group. You must add both the terminal servers for which you need to provide licenses and any license servers that might need to acquire licenses to this group for each license server. For ease of management in a large Terminal Server deployment, create an Active Directory global group named Terminal Server Licensing and add all of your terminal servers and all of your license servers to this group. Then add this group to the Terminal Services Computers group of each license server. Whenever you deploy a new terminal server or license server, if you add it to the Terminal Server Licensing group, it automatically appears in the Terminal Services Computers group for each license server.

**Prevent License Upgrade** In an environment with both Windows Server 2003 and Windows 2000 Terminal Server, enable this Group Policy setting to prevent the license server from handing out Windows Server 2003 CALs to terminal servers that are running Windows 2000.

**Limit users to one remote session**
By default, the TSCC setting that restricts users to one session is set to Yes. This ensures that a user who disconnects can

reconnect to the same session. This also conserves server resources by keeping the number of sessions on your server to a known number. You can enforce this setting domain-wide by using Group Policy. Limit users to one session unless you have a valid business reason for allowing more than one.

**Permission Compatibility**
This TSCC setting is set to Full Security by default. It restricts access to system resources, such as the registry, to members of the Users group on a terminal server.

### Important

- Although some older applications require such access, do not set this setting to Relaxed Security unless you have determined through testing that your applications do not work properly in Full Security mode.

**Home Directory**
Use this Group Policy setting to set the path for storing your user home directories. For more information about planning for storage of home directories, see "Using Home Directories with Terminal Server" earlier in this chapter.

**Roaming Profiles**
Use this Group Policy setting to set the path for storing your roaming user profiles. For more information about planning for storage of roaming user profiles, see "Using User Profiles in Windows Server 2003" earlier in this chapter.

# Using Software Restriction Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Software restriction policies, which are new with Windows XP and Windows Server 2003, enable you to identify software running on computers in your domain and to control whether a user can run them. Restricting certain types of applications can, for example, protect your organization against viruses. As a way to lock down the user environment on a terminal server, you can set up software restriction policies that allow users to run only specific applications on the server.

Software restriction policies are located in the Group Policy Object Editor under Windows Settings/Security Settings. Windows Installer operates with applications permitted by these Software Restriction Policies. For more information, see "Software Restriction Policies" in Help and Support Center for Windows Server 2003.

You can use software restriction policies with Terminal Server by using path rules, as shown in Table 4.4. These rules allow groups of users, when separated into different OUs, to access only the applications or application components that you want the groups of users to access on the server. For example, a company has a terminal server with a line-of-business application and a few productivity applications for the use of the accounts payable department. The company has decided that account managers need access to all of the available applications for that department, but the data-entry workers in that department need access only to the line-of-business application. The company sets the default rule to Disallowed and configures the software restriction policies as outlined in the following table.

**Table 4.4 Example Software Restriction Policy Configuration**

| Path Rule | Security Level |
|---|---|
| **Terminal Server OU** | |
| %windir% | Unrestricted |
| %windir%\regedit.exe | Disallowed |
| %windir%\system32\cmd.exe | Disallowed |
| %windir%\system32\command.com | Disallowed |
| %windir%\system32\dllcache | Disallowed |
| %windir%\system32\gpresult.exe | Disallowed |
| %windir%\system32\gpupdate.exe | Disallowed |
| %ProgramFiles%\Windows NT\Accessories | Unrestricted |
| **Data Entry OU** | |
| %ProgramFiles%\Accounts Payable Software | Unrestricted |
| **Account Managers OU** | |
| %ProgramFiles%\Microsoft Office\Office | Unrestricted |
| %ProgramFiles%\Internet Explorer | Unrestricted |

# Configuring Terminal Server for Differing Time Zones

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default Terminal Server keeps track of time according to the time zone in which it has been configured, rather than on a per-user basis. This can be a problem when a user connects to a terminal server outside of the time zone in which the user is located because the local computer uses the time configured on the terminal server rather than the local time. If you are hosting time-sensitive applications on your terminal server or if you have processes in place that depend on the user's current local time, such as financial systems and calendaring, you might want to enable the Allow Time Zone Redirection Group Policy setting. This policy is located in Windows Components/Terminal Services. With this setting enabled, Terminal Services uses the server base time on the terminal server and the client time zone information to calculate the time on the session.

# Designing the Remote Desktop Connection Installation and Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Remote Desktop Connection is the client-side tool that enables the connection to the terminal server. You need to install the Remote Desktop Connection on clients running Windows 2000 or earlier versions of the Windows operating system. On Windows XP clients, the Remote Desktop Connection is already installed. However you can preconfigure and distribute a connection to ensure the optimal settings for your users. For more information about the Remote Desktop Connection, see "Configuring Remote Desktop" in the *Windows XP Professional Resource Kit Documentation* (or see "Configuring Remote Desktop" on the Web at http://www.microsoft.com/reskit). Figure 4.6 shows the process for designing the Remote Desktop Connection installation and configuration.

**Figure 4.6 Designing the Remote Desktop Connection Installation and Configuration**



# Designing the Remote Desktop Connection Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you currently have Windows XP deployed to the clients, the Remote Desktop Connection is already built in. No further deployment is required unless you want to configure the connection.

The Remote Desktop Connection Windows Installer Setup package is a tool for deploying the Remote Desktop Connection and associated files for specially configured connections and on computers running Windows NT 4.0, Windows 2000 Server, Windows 2000 Professional, Windows 98, and Windows 95. You can also place the Client Windows Installer Setup package on a share for distribution through Group Policy to workstations running Windows Server 2003.

# Configuring Remote Desktop Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using the Remote Desktop Connection tool, you can set from the client many of the same settings that you can set on the server using TSCC, Group Policy, or WMI. Because server settings override client settings, it is recommended that you set these settings on the server for ease of management. However, there are some situations where you might need to set these settings using the Remote Desktop Connection tool. An example is if you are providing access to the same terminal server for users with slightly different user roles or purposes. These settings are summarized in Table 4.5. For detailed explanations and recommendations for these settings, see "Designing the Terminal Server Configuration" earlier in this chapter.

**Table 4.5 Remote Desktop Connection Tool Settings**

| Tab | Description |
|---|---|
| General | Enter or change logon and connection settings. |
| Display | Change remote Desktop size (resolution) and color depth, and choose whether to display the connection bar in full-screen mode. |
| Local Resources | Control sound, the behavior of keyboard combinations, and connections to local devices (drives, printers, and serial ports). |
| Programs | Start a program upon connection to the terminal server. |
| Experience | Configure the connection speed and adjust the desktop environment for optimal performance |

With the Remote Desktop Connection tool you can also create and save a preconfigured Remote Desktop Connection for distribution.

**Preconfiguring a Remote Desktop Connection**

Remote Desktop Connection is installed by default on all versions of Windows Server 2003, most versions of Windows CE, and all versions of Windows XP operating systems. On earlier versions of Windows and on the Pocket PC you must manually install Remote Desktop Connection in order to connect to Terminal Server.

**Note**

- Windows 2000 includes the 32-bit Terminal Services Client for connecting to Terminal Services. It is recommended that you install the latest version of the Remote Desktop Connection to get the latest performance improvements.

You can preconfigure a Remote Desktop Connection for distribution to your client computers through Group Policy as an .msi package along with associated files.

**To create a preconfigured Remote Desktop Connection for distribution**

1. Open Remote Desktop Connection.

2. Expand the dialog box by clicking the **Options** button and configure the connection.

3. On the **General** tab of the Remote Desktop Connection tool, click the **Save As** button, navigate to the folder in which you want to save the configured connection, rename the file, and click **Save**.

4. Create an .msi package and distribute it to your Terminal Server users OU. For information about how to create an .msi package, see the Windows Installer documentation link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Note**

- Users need to enter their own user name and password into the connection.

**Configuring the Desktop Experience**

You can configure the connection speed and fine-tune the desktop environment for optimal performance by using the settings located on the **Experience** tab of the Remote Desktop Connection tool. The settings available on this tab are summarized in Table 4.6. Most of these settings relate to the graphical nature of the desktop. Because the processing for graphics happens on the server and must be transmitted to the desktop over the network connection, you can increase performance by allowing only the minimal settings that are necessary for your users. When you choose a connection speed, the check boxes are automatically selected to indicate the recommended settings for the richest visual experience possible at that speed.

**Table 4.6 Desktop Experience Settings**

| Setting | Description |
|---|---|
| Desktop background | Allows user to choose a custom background or wallpaper for their desktop. Because these can be very graphic-intensive, disable this setting unless it is necessary for your users. |
| Show contents of window while dragging | Redraws the contents of a window rather than showing just the frame of the window when you move the window across the screen. Because this redrawing takes place on the server and needs to be transmitted over the network connection to display on the user's desktop, disable this setting unless it is necessary. |
| Menu and window animation | Enables cascading menus and scroll transitions for menus and tool tips. Because this rendering takes place on the server and needs to be transmitted over the network connection to display on the user's desktop, disable this setting unless it is necessary. |
| Themes | Allows user to choose a background plus a set of sounds, icons, and other elements for their desktop. Because many themes are graphic-intensive, disable this setting unless it is necessary for your users. |
| Bitmap caching | Stores frequently used images on the local computer to speed up the connection. By enabling this setting, these images do not have to be transmitted across the network connection each time they appear on the desktop. This setting is enabled by default. |

**Locking Down the User Session**

You can configure Group Policy settings and apply them to your Terminal Server Users group to control your users sessions and user access to certain features on the server. For more information about controlling user sessions, see article 278295, "How to Lock Down a Windows 2000 Terminal Server Session." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. You can also lock down your users' sessions by using the Security Templates MMC snap-in or the settings under Computer Configuration/Windows Settings/Security Settings in the Group Policy Object Editor.

# Additional Resources for Hosting Applications with Terminal Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- The *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit) for information about using Terminal Services for remote administration.
- "Configuring Remote Desktop" in the *Windows XP Professional Resource Kit Documentation* (or see "Configuring Remote Desktop" on the Web at http://www.microsoft.com/reskit).
- The Terminal Services link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Terminal Services Community Center link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- The Terminal Services Scaling link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Related Tools**

- Roboserver and Roboclient

  Using these tools you can easily place and manage simulated loads on a terminal server. This information can help you determine whether your environment can handle the load that users place on it. Used together, these tools make capacity planning easier and more automated. For more information about Roboserver and Roboclient, click **Tools** in Help and Support Center, and then click Windows Resource Kit Tools Help.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

Topics under "Terminal Services overview" in Help and Support Center for Windows Server 2003.

# Planning for Remote Server Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you deploy servers running the Microsoft® Windows® Server 2003 operating system, determine the extent to which system administrators need to manage them remotely. If you need to perform only typical administrative tasks over the network, conventional tools suffice. For remotely administered servers — those having no keyboard, video display, or mouse — or for any servers having high availability requirements, you might need to use Emergency Management Services, possibly in conjunction with other hardware and software components, to remotely bring servers back into service even when they are not accessible by means of the standard network connection.

**In This Chapter**


Overview of Remote Management Planning

Evaluating Remote Management Needs

Choosing Software and Hardware Tools

Designing the Hardware Configuration

Configuring Your Infrastructure for Remote Management

Planning for Remote Management Deployment

Additional Resources for Remote Server Management


**Related Information**

- For more information about Emergency Management Services for either remotely administered or standard servers, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

- For an overview of remote administration and information about associated remote management tools, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Overview of Remote Management Planning

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although local administration is the most secure way to manage the computers in a networked environment, cost, staffing, and system availability requirements typically preclude this approach in medium to large organizations. Remote management can solve many of these problems by increasing productivity, decreasing time to resolution, decreasing staffing requirements, and allowing more flexibility in server placement. Remote management can also resolve physical accessibility issues for servers that are remotely located or rack mounted — whether they are remotely administered or use keyboard-video-mouse (KVM) switches.

Windows Server 2003 provides a variety of tools and technologies to help you remotely administer the servers in your network. Conventional tools help you perform common tasks on computers that are functioning and available over the network. These tools are sometimes referred to as *in-band* tools because they function through the standard network connection. Other tools allow you to connect to a computer that is not responding to the standard network connection for some reason, such as when a Stop error occurs or the network adapter fails. These tools are sometimes referred to as *out-of-band* tools because they use a connection that does not depend on network drivers.

Emergency Management Services is a new feature in Windows Server 2003 that supports out-of-band connections. Emergency Management Services is included with Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; and Windows® Server 2003, Web Edition.

Before you begin managing servers remotely using in-band or out-of-band tools, you need to develop a remote management plan. A remote management plan ensures that you choose the appropriate remote management tools and management configurations for your organization and that you address all the impacts they might have on your infrastructure.

To develop your plan, you first need to know the server configuration in your organization: where the servers are located and what roles they perform. You also need to know the availability requirements for the servers you plan to manage remotely and who the administrators are.

Follow the planning process presented here to develop a plan for remotely managing the servers in your environment. This planning process addresses tools for Windows Server 2003 remote installation, in-band management tasks — such as configuring, monitoring, and troubleshooting — and out-of-band management tasks — such as restarting the computer and viewing Stop error messages. It does not address workstation management or software distribution tools for applications and updates.

# Remote Management Planning Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To produce an effective remote management plan, you need to determine which servers to manage and when to manage them remotely or locally, the tools to use, the required supporting hardware components and configurations, network and security considerations and remediation, and whether and how to deploy Windows Server 2003 remotely. These factors are included in the process presented here. Figure 5.1 illustrates the steps in the remote management planning process.

**Figure 5.1 Planning Remote Management**



# Remote Management Concepts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Conventional remote management involves establishing an in-band connection to a server to manage it across the network. As a class, in-band management tools are robust, versatile, and secure. Whenever a server is functioning and available across the network, in-band management tools are the tools of choice. In the past, this type of management was the only type of remote management available for servers running Windows operating systems.

By using Windows Server 2003, you can also manage servers by using out-of-band connections that make it possible to perform management tasks when the server is not responding to the standard network connection. When you lose the capability to manage the server with the in-band connection — such as when the firmware is initializing, a Stop error occurs, or the server stops responding — you need to use an out-of-band connection to manage it remotely.

A primary purpose of the out-of-band connection is to provide a means for you to return the server to a functioning state so that you can continue to manage it with your conventional in-band tools. Out-of-band management also makes it possible for you to configure your servers for remote administration, which means the local keyboard and monitor are optional. When combined with appropriate out-of-band hardware components, Windows Server 2003–based servers need to be managed locally only for hardware installation or replacement.

**Note**

- For a document to assist you in setting up servers for remote administration, see "Headless Server Quick Start" (SDCEMS_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "Headless Server Quick Start" on the Web at http://www.microsoft.com/reskit).

**In-Band Management**

Conventional remote management tools, such as Telnet, use in-band connections for communicating with the managed server. In-band connections rely on operating system network drivers for establishing connections between computers, so a server must be initialized and operational to be managed with these tools.

The most common in-band connection hardware device is a network adapter, such as an Ethernet adapter, analog modem, or Integrated Services Digital Network (ISDN) modem. The most typical method of in-band remote management is connecting through the network directly to the server, but you can also use Windows Server 2003 remote access and virtual private network (VPN) connections to manage servers through an in-band connection.

Typical tasks performed by using in-band management tools include routine configuration, monitoring, troubleshooting, and maintenance.

**Out-of-Band Management**

When a server is not in a functional state and cannot be accessed by using the standard network connection, you need to use an out-of-band connection to manage it remotely. Out-of-band connections do not rely on operating system network drivers for establishing connections between computers. The following are typical situations when you might need to manage a computer by using an out-of-band connection:

- The server is powered down.
- The BIOS is conducting the POST.
- A Stop error occurs.
- The server is too low on resources to respond adequately.
- The network adapter malfunctions or fails.
- The Windows loader or Recovery Console is running.
- The server is not fully initialized.

Common out-of-band connection hardware devices include serial ports, analog modems, and ISDN modems. In addition, out-of-band hardware components known as service processors can in some cases provide out-of-band connections over the network. A *service processor* is a microprocessor that functions independently of the CPU(s) in a computer and provides additional server management functionality for any operating state, whether or not the operating system is functioning.

Emergency Management Services, as well as features built into some hardware components and firmware, support out-of-band connections. Hardware components that support out-of-band connections include service processors, terminal concentrators, intelligent uninterruptible power supplies (UPSs), and intelligent power switches. A *terminal concentrator* is a hardware device that allows you to monitor multiple servers simultaneously by connecting to their out-of-band serial ports through a single network connection. An *intelligent UPS* or *intelligent power switch* is one that provides some remote functionality, such as powering up or down or resetting a computer. These components are described in detail in "Choosing Software and Hardware Tools" later in this chapter.

*Console redirection* is a key out-of-band feature that sends keyboard input and character-based output destined for the local display device to the out-of-band port so that you can view the information on a remote computer. Console redirection does not preclude locally attached monitors and keyboards: a computer with console redirection can still process input from a local keyboard and display output to a local monitor.

Emergency Management Services, service processors, and some system firmware provide console redirection. The extensible firmware interface (EFI) on Itanium-based computers typically provides console redirection. The BIOS on x86-based computers might or might not provide console redirection; contact your computer manufacturer to find out.

The specific component that controls console redirection changes as the operating state changes. For example, a service processor or firmware provides console redirection during power up and during the Power On Self Test (POST), and Emergency Management Services provides console redirection as soon as the Windows loader (Ntldr) starts. You can transfer control of console redirection to another component, such as a service processor, by entering specific escape sequences from the management computer.

Additional components that support out-of-band connections are described later in this chapter.

**Emergency Management Services**

Emergency Management Services is a new feature in Windows Server 2003 that provides out-of-band connections through a serial port or, in the case of some service processors, an alternate network connection. With Emergency Management Services, you can perform administrative tasks remotely using an out-of-band connection. When you combine Emergency Management Services with the appropriate out-of-band hardware, you can perform all administrative tasks remotely, except for installing or replacing hardware.

Emergency Management Services features are fully or partially available when the Windows Server 2003 operating system is loading, is running, or is in distress — such as when it is slow or not responding or when a Stop error occurs.

The three Emergency Management Services features that support out-of-band management are console redirection, Special Administration Console (SAC), and !Special Administration Console (!SAC). The Windows Server 2003 loader or kernel must be at least partially functioning for these features to be available.

**Console redirection**

Emergency Management Services can redirect keyboard input and character-based output when any of the following Windows Server 2003 components are running:

- The operating system loader on x86-based multiple-boot computers
- Recovery Console
- The Windows Server 2003 kernel
- The command prompt (cmd.exe)
- Text-mode Setup, during Windows Server 2003 installation
- Remote Installation Services (RIS) Setup

For more information about how Emergency Management Services console redirection works with these components and when it is not available, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Special Administration Console**
Special Administration Console (SAC) is the primary Emergency Management Services command-line environment. It provides a variety of commands for monitoring server status and troubleshooting problems during emergencies. SAC is available whenever the Windows Server 2003 kernel is running in normal mode, safe mode, and during the graphical user interface (GUI)–mode phase of Setup.

### Caution

- Using SAC during an operating system upgrade or installation might cause the upgrade or installation to fail or become unstable.

Using SAC, you can establish multiple user sessions, called *channels*, and switch between them. Although SAC is separate from the Windows Server 2003 command prompt (cmd.exe), you can establish a command prompt channel from SAC and then switch between SAC and the command prompt channels by using SAC commands or escape sequences. During the GUI-mode phase of Setup, you can also establish a channel for viewing setup logs so that you can troubleshoot unresponsive or failed operating system installations. You can access only one channel at a time, which means that multiuser access to SAC is not available. To create a SAC channel, you must use a local Administrator account; therefore, all the commands that you run in a channel run as local administrator.

For more information about SAC capabilities and how to use them, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit). For more information about SAC commands, see "Special Administration Console (SAC) and SAC commands" in Help and Support Center for Windows Server 2003.

**!Special Administration Console**
!Special Administration Console (!SAC) is a scaled-down version of SAC that is designed to recover an unresponsive system — for example, due to a Stop error. You cannot directly invoke !SAC; it becomes available automatically when a server experiences a system failure or fault. !SAC is a last-resort tool that lets you view Stop messages, obtain computer identification information, and restart the server.

**Remotely Administered Servers**
Remotely administered servers (also known as "headless servers") are computers that can operate without a keyboard, mouse, and local monitor. If the system firmware supports it, the video adapter and keyboard controller are also optional. Emergency Management Services, in combination with the appropriate hardware and system firmware console redirection, makes it possible for you to configure servers running Windows Server 2003 for remote administration. Redirecting firmware allows the system to pass the POST without a video adapter. When a server runs Emergency Management Services, you can manage it with both in-band and out-of-band tools, making locally attached input and output devices unnecessary.

Using remotely administered servers does not preclude having locally attached input devices. Depending on your hardware configuration, you might be able to attach and remove local devices, such as Universal Serial Bus (USB) keyboards and mouse devices, as needed.

Computer manufacturers offer systems preconfigured for remote administration that have no keyboard controller, video adapter, or mouse. For more information about purchasing a computer that you can administer remotely or configuring your existing computers for remote administration, contact your computer manufacturer.

# Evaluating Remote Management Needs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To implement an efficient, cost effective, secure remote management plan, you need to determine which servers to manage remotely by considering the tradeoffs — including cost, convenience, and system availability — for various ways of managing the servers in your environment. Before you decide whether or what to manage remotely, you need to evaluate remote and local management, and in-band and out-of-band management.

With Windows Server 2003, you can choose to remotely perform server management tasks that previously could be done only locally. By using Emergency Management Services with the appropriate hardware and software tools, you can remotely perform almost everything except hardware installation. When considering this type of solution, you need to weigh the level of support provided by the solution against the possible need to use legacy hardware. You might need to make a compromise between the level of support you want and the cost of additional hardware. You might also want to evaluate the savings you can obtain by configuring some of your servers for remote administration.

The first step in developing a remote management plan is to determine which servers in your environment need to be managed remotely. Figure 5.2 illustrates the place of this step in the planning process.

**Figure 5.2 Evaluating Your Remote Management Needs**



# Assessing Your Server Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assessing your server environment includes determining how many and what kind of servers you have, where they are located, and who administers them. The goal is to decide which servers should be managed remotely and which should be managed locally. For servers you want to manage remotely, you need to decide whether you want to perform only in-band tasks — those performed when the operating system is up and running — or both in-band and out-of-band tasks — those performed when the operating system is not running or is not responding appropriately. For example, you might want to configure some servers so that you can remotely install the operating system and perform all management tasks remotely, regardless of the state of the system, while for other servers you might only want to perform routine maintenance tasks remotely.

**Identifying Computers for Remote Management**

When determining which servers to manage remotely and which locally, consider the following issues:

- **Number of servers.** Does your organization have many servers or just a few? How manageable are they in relation to the number of technicians available for local support? If your organization has few servers and technicians are readily available, you might not need to manage remotely. On the other hand, if you have many servers and limited support resources, you probably want to manage at least some of them remotely.

- **Location of servers.** Are the servers in your organization consolidated in a single or a few locations where they are easily accessed by technicians, or are they widely dispersed? If the computers are distributed throughout your organization, what distances are involved? For example, are they located internationally, in different buildings in the same city, or down the hall? The more servers you have and the farther apart their locations, the greater the savings you are likely to achieve by remote management.

- **Server roles supported.** Do you have different management requirements for different types of servers, such as domain controllers, print servers, file servers, Web servers, application servers, and database servers? You might decide to manage different types of servers differently for reasons such as security or criticality to your business. For example, you might want to manage application servers remotely for mission-critical applications, but domain controllers locally for security reasons.

- **Availability requirements.** What percentage of the time must a server be up and running? For example, if you have servers with an availability rating of 99.99 percent, you can't afford a downtime of anything more than minutes. How many users typically access a server? For example, if a domain controller has a problem, many users might be affected, while if a print server in a small department fails, only a few users are affected. What is the impact to your business if a server is not available for a period of time? Remotely managing servers that have high availability requirements can significantly save downtime while waiting for someone to perform local management.

- **Reliability history.** How reliable is the server? How frequently does it typically encounter problems? The lower the reliability, the more benefit you can derive from managing it remotely.

- **Number and expertise of on-site technicians.** How many technicians do you have for local management and where are they based? For example, are there administrators on-site in branch offices? How experienced are they? If you have few technicians in certain areas, or they are inexperienced, you can probably benefit from remote management.

- **Security issues.** Do you have servers you prefer to manage locally for security reasons? For example, do you have public servers that you prefer to manage locally rather than by establishing a VPN connection? Or do have specific tasks that you prefer to perform locally and others you prefer to perform remotely?

**Identifying Tasks for Remote Management**

After you know which servers you want to manage remotely, the next step is to assess which kinds of management tasks to perform remotely and which to perform locally. For example, do you want to perform tasks remotely only when a server can be accessed across the network, or also when it cannot be accessed across the network?

If you have enough experienced technicians available or have highly reliable servers, in-band remote management support

might be sufficient. On the other hand, if some of your servers have high availability requirements or a history of poor reliability, or if you have insufficient on-site technical support, you might need out-of-band remote management support.

If you have many servers mounted in racks and tethered to KVM switches, you might decide to configure them for remote administration. You can then perform all administrative tasks remotely. Because remotely administered servers do not have monitors, keyboards, or mouse devices and might not require video adapters, they can present significant savings in space, energy consumption, and hardware costs.

Another consideration is whether your current servers are designed for out-of-band management. If you plan to use the hardware you already own and it is not designed for out-of-band support, your options might be more limited. For more information about deciding whether to use existing or new hardware, see "Planning for Remote Management Deployment" later in this chapter.

If you decide you need out-of-band remote management support, the next step is to determine how much support you need. As you determine the degree of out-of-band support you want, you need to weigh the costs of any additional hardware components — and possibly the cost of upgrading existing computers — against the savings in system uptime and reduced on-site support.

By using Emergency Management Services with nothing but a serial port, you can access nearly all Windows Server 2003 operating system states. By combining Emergency Management Services with appropriate out-of-band hardware, however, you can perform nearly every administrative task remotely — even before the operating system is fully initialized or after it stops responding — and you can effectively manage many servers from a single management computer.

To help determine the degree of out-of-band remote support you need, decide which of the following tasks you need to perform remotely:

- Start up a server
- Change BIOS settings
- View POST results
- Select which operating system to start
- View Stop error messages
- Reset a server
- Access the server when it is inaccessible through the network
- Install the Windows Server 2003 operating system

Different out-of-band tasks might require different types of hardware components. For more information about which components are needed for which tasks, see "Choosing Out-of-Band Management Tools" later in this chapter.

**Identifying Responsibilities for Remote Management**

As you plan for remote server management, you also need to decide whether you want different administrators to be responsible for different groups of tasks or whether you want all administrators to be authorized to perform the entire range of tasks. Depending on the size of your organization, the number and skill level of available administrators, and security requirements, you might decide to make some administrators responsible for only certain administrative tasks. To decide how to assign remote management tasks, consider the following issues:

- **Security issues.** Do you want to divide up tasks for security reasons? For example, you might decide that certain tools — based on the power of the tool or its level of built-in security — are to be used only by some administrators or that specific servers are to be accessible only to some administrators.
- **Skill level of administrators.** Do you have highly skilled administrators in all locations or only in some locations? For locations with less skilled administrators, you might decide to use remote management rather than telephone support.

# Examples: Determining Remote Management Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As you evaluate your needs for remote management, you develop a list of requirements. These requirements determine which tools and out-of-band configurations are most appropriate for your environment. The following descriptions illustrate three different types of remote management requirements and the level of remote management support appropriate to each. The tools and configurations appropriate to these levels of remote management are described later in this chapter.

**Minimal Remote Management**
You need minimal support for remote management if you have requirements such as the following:

- The site has few servers and those servers are located close together.
- The servers have low availability requirements.
- Experienced on-site technicians are typically available.
- Some routine management tasks are performed centrally.

An example business situation with minimal requirements for remote management might be a branch office that has a single, or just a few, servers running Windows Server 2003. The servers are primarily file and print servers and have low availability requirements. The branch has some on-site support available.

In this situation, on-site technicians perform most administrative tasks locally, but centrally located administrators perform some routine in-band tasks remotely over the network, such as changing a static IP address or other configuration settings,

**Moderate Remote Management**
You need moderate support for remote management if you have requirements such as the following:

- The site has many servers.
- The servers have availability requirements that range from low to high.
- Experienced on-site technicians are typically available to perform out-of-band tasks for servers with low to medium availability requirements.
- Many routine management tasks can be performed efficiently by using in-band remote management tools.

A situation with moderate requirements for remote management might be a data center that has many computers, some of which are legacy systems that do not support Emergency Management Services.

In this situation, the organization plans to supplement existing systems with out-of-band hardware components in order to perform some out-of-band management tasks remotely, even though the computers are not built for Emergency Management Services.

Administrators use in-band remote management tools for the systems in the data center. They use out-of-band management for some systems that have high availability requirements, but not for others that do not have such requirements. For example, the print servers are very reliable, so the organization decided it was not cost effective to support out-of-band management for these servers.

**Maximum Remote Management**
You need maximum support for remote management if you have requirements such as the following:

- The servers have high availability requirements.
- You need to perform a number of out-of-band tasks remotely.

- The servers are remotely administered.

A situation with maximum requirements for remote management might exist when a group of computers has very high availability requirements. These computers support remote administration and have built-in service processors. They are mounted in racks in a highly secured room.

Administrators go into the highly secured room only when they add or replace hardware. All other administrative tasks are performed remotely, by using an in-band connection when possible and by using an out-of-band connection when a server cannot be accessed through the standard network. The Windows Server 2003 operating system is installed on these computers remotely by using RIS.

# Choosing Software and Hardware Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you determine how much you want to manage remotely, the next step is to select the tools and supporting components you need to accomplish your remote management tasks. Figure 5.3 illustrates the place of this step in the process.

**Figure 5.3 Choosing Software Tools and Hardware Components**




As you select your tools, think about the tasks you want to perform remotely when you have network access — by using in-band connections — and those you want to perform remotely when you do not have network access — by using out-of-band connections. As you select tools, evaluate their potential impact on your environment and build any needed environmental adjustments into your remote management plan.

**Selecting In-Band or Out-of-Band Tools**

In-band management is always the method of choice for managing servers when you can access them through their standard connections. If a server is functional enough to respond through the standard connection, conventional in-band management tools can provide a much broader range of functionality — and possibly greater security — than you might achieve with out-of-band management.

**Note**

- While the security of in-band management is highly dependent on the individual management tool, the security of out-of-band management is highly dependent on your out-of-band component configuration. For example, in a configuration that uses a remote serial connection, the security of the out-of-band management is dependent on the security built into the modem. For information about the security implications for different out-of-band configurations, see "Designing the Hardware Configuration" later in this chapter.

Keep in mind that out-of-band management is a last resort when you cannot access the server in any other way. The goal of out-of-band management is always to bring a server back into service so that you can manage it with in-band tools.

Table 5.1 shows whether to use in-band or out-of-band tools for various types of tasks during various operating states. After you know whether to use an in-band tool or an out-of-band tool, you can select the most appropriate specific tool or component, as described later in this chapter, for the tasks you want to perform remotely.

**Table 5.1 Choosing In-Band or Out-of-Band Tools**

| During This Operating State | For This Type of Task | Use This Type of Tool |
|---|---|---|

| | | |
|---|---|---|
| System powering on or off, or resetting | Power up, power down, reset | Out-of-band and in-band with Remote Desktop for Administration |
| Firmware initializing | Configure firmware, troubleshoot, restart | Out-of-band with supporting firmware |
| Operating system loading | Choose operating system to start, troubleshoot | Out-of-band, including Emergency Management Services |
| Text mode setup | Monitor, troubleshoot | Out-of-band, including Emergency Management Services |
| GUI mode setup | Monitor, troubleshoot | Out-of-band, including Emergency Management Services |
| Operating system fully functional | Monitor, troubleshoot, modify configuration settings | In-band |
| Operating system not responding on network | Troubleshoot, restart | Out-of-band, including Emergency Management Services |
| Stop message occurred | Troubleshoot, restart | Out-of-band, including Emergency Management Services |
| System extremely slow responding on network | Troubleshoot, restart | In-band and out-of-band, including Emergency Management Services |

**Evaluating Tools for Environmental Impact**

As you evaluate the software and hardware tools to use, consider the impact they might have on your network environment. For example, some tools present more security risks than others, and some increase network traffic more than others. Considerations such as these might influence your selection of one tool over another, or they might identify additional changes you need to make to your environment to mitigate the impact. The documentation provided with a remote management tool might contain information indicating its potential impact on your environment and any configuration changes needed for its use. For more information about configuring your environment for remote management, see "Configuring Your Infrastructure for Remote Management" later in this chapter.

As you develop your remote management plan, include the following lists:

**Tasks you plan to perform remotely** The more comprehensive you make this list, the easier it will be to identify all the tools you need. A task can be broad in scope (for example, manage DHCP servers), or it can be narrow in scope (for example, change the static IP address on a server). This list should include not only in-band tasks but also any required out-of-band tasks, such as remotely installing the operating system or powering up the computer.

**Tools for performing the tasks** Typically, you can use several different remote management tools to perform the same task. Include in your list all the tools that apply to each remote task you want to perform. In some cases, you do not need to use a specific tool to perform a remote administration task; rather, you simply need to change a configuration setting. If a task does not require a specific tool, note this in your remote administration plan. If your environment includes a mix of operating systems, you might need to look for tools that provide interoperability for some tasks. Make sure this list also includes any out-of-band tools or components you plan to obtain or install.

**Network impacts to be addressed** Remote administration can have any of several impacts on your network: it can increase network traffic, decrease server performance, or create security vulnerabilities. You might need to reconfigure network, system, or security settings to mitigate or eliminate these impacts. Include in this list each potential impact and the specific steps you plan to take to address it.

# Choosing In-Band Management Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 supports a wide variety of in-band remote management tools that you can use to manage servers. Use in-band tools when your Windows Server 2003–based server is functioning and accessible through your standard network connection.

Tools for remotely managing servers are available from many sources. Some of the tools are specific to a task, while others support a range of tasks. Some provide a command-line environment, while others provide a graphical user interface (GUI) environment. Some tools work best for managing a single computer at a time, while others support sessions with multiple computers.

In addition to the many tools built in to Windows Server 2003, management tools are available from the following sources:

- Windows Support Tools, located in the Support\Tools folder on the Windows Server 2003 operating system CD, provide command-line tools for specific management tasks in a variety of areas such as performance, security, and deployment. For more information about Windows Support Tools, in Help and Support Center for Windows Server 2003, click **Tools**, and then click "Windows Support Tools."

- Resource Kit Tools provide a variety of command-line tools for specific tasks. For more information about Resource Kit Tools, click **Tools** on Help and Support Center for Windows Server 2003, and then click Windows Resource Kit Tools Help.

- Third-party tools available from independent software vendors (ISVs) provide a wide variety of specific or general remote management capabilities.

This section describes major characteristics of some common remote management tools for servers. Some of these tools can also be used for performing management tasks through out-of-band connections. For in-depth information about the technical considerations and potential impacts of specific remote management tools, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit). For information about software distribution tools, see "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit.

Table 5.2 summarizes the characteristics of common remote management tools.

**Table 5.2 Common Remote Management Tools**

| Tool | Key Characteristics |
|---|---|
| Telnet | Command line; efficient and versatile; provides interoperability in mixed environments; in general, not secure |
| Windows Management Instrumentation Command-line (WMIC) | Customized applications and command-line scripts for remote management |
| Windows Script Host (WSH) | Customized scripts for remote management |
| Microsoft Management Console (MMC) | Multiple sessions; variety of snap-ins for various administrative tasks |
| Remote Desktop for Administration | GUI; multiple sessions; high resource usage |
| Group Policy | Efficient way to manage a variety of settings for groups of servers |

**Telnet**

Telnet is a global, versatile tool that has minimal system resource and network bandwidth requirements and that provides interoperability with other operating systems. With Windows Server 2003 Telnet Server, any client that supports the Telnet

protocol can connect to Windows–based systems. For example, a UNIX Telnet client can connect to a Windows–based server.

By using Telnet, you can establish a command console session on a remote computer and use it to run command-line programs and shell commands, interacting with the remote server as though you were logged on locally. Telnet can establish any number of connections and supports interactive scripts.

The Windows Server 2003 32-bit version of Telnet does not support secure logon, while the 64-bit version provides secure logon by using NTLM authentication. Some versions of Telnet provided with terminal concentrators also support secure logon. Telnet does not support encryption.

By using Telnet, you can perform out-of-band management tasks by establishing a network connection to a terminal concentrator that is connected to servers through their serial ports. For more information about terminal concentrators, see "Using Terminal Concentrators for Remote Management" later in this chapter.

**Windows Management Instrumentation**
Windows Management Instrumentation (WMI) is an infrastructure that enables you to access and modify standards-based information about objects — such as computers, applications, and network components — in your enterprise environment. Using WMI, you can create powerful administration applications to monitor and respond to specific events in your environment. For example, you can create applications to check CPU usage on your Windows Server 2003–based servers and notify you when it exceeds a specified level. Although WMI is a powerful tool for building customized applications, it does require a certain amount of developer time and expertise.

Windows Management Instrumentation Command-line (WMIC) provides a simplified interface to WMI. By using WMIC, you can access WMI-based information using the command line or scripts. You can use WMIC from any computer where WMIC is enabled to manage any remote computer. WMIC does not have to be available on the remote computer.

For technical information about developing applications using WMI and using WMIC, see the WMI SDK link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Windows Script Host**
Windows Script Host (WSH) is a language-independent scripting infrastructure that allows you to write scripts for local or remote management tasks. You can use WSH to write scripts that include WMI, Active Directory® directory service, and other application programming interface (API) calls. WSH typically is used for noninteractive scripts, such as logon and computer automation scripts.

**Microsoft Management Console**
Microsoft Management Console (MMC) is a framework for hosting tools, also known as snap-ins, that you can use to manage servers locally or remotely. With MMC, you can create consoles that include the tools you use most often.

Each MMC snap-in has unique advantages and disadvantages that make it suitable in some cases and unsuitable in others. For example, some are suitable for slow network connections and some transmit encrypted data. Before you use a snap-in to perform a remote management task, make sure that it is the best remote management tool for the task. For more information about using MMC snap-ins for remote management, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Remote Desktop for Administration**
Remote Desktop for Administration is an MMC snap-in that you can use to establish a remote console session on one or more servers and switch between sessions. By using Remote Desktop for Administration, you can log on to a remote server and use the server's desktop to perform administrative tasks, just as if you were logged on locally. Remote Desktop for Administration supports Kerberos authentication and built-in encryption.

Remote Desktop for Administration is a versatile remote management tool because it supports both GUI and command-line interfaces. Because Remote Desktop for Administration uses Remote Desktop Protocol (RDP), it efficiently transmits the user interface from the server to the client and keyboard sequences and mouse clicks from the client to the server. Nevertheless, this tool requires more memory and network bandwidth resources than many other tools.

### Important

- Remote Desktop for Administration is affected by the Internet Explorer Enhanced Security Configuration, which places your server and Microsoft Internet Explorer in a configuration that decreases the exposure of your server to attacks that can occur through Web content and application scripts. As a result, some Web sites might not display or perform as expected. For more information, see "Setting up Remote Desktop Web Connection" and "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003.

**Windows Server 2003 Administration Tools Pack**

The Windows Server 2003 Administration Tools Pack includes several of the most common tools for remotely managing servers from a Microsoft® Windows® XP Professional–based computer with Service Pack 1. Many of the tools are MMC snap-ins. The tools pack is included on the 32-bit version of the Windows Server 2003 operating system CD, and the Windows Installer package — Adminpak.msi — is placed in C:\windows\system32\adminpak during the operating system installation. For 64-bit versions, use Remote Desktop for Administration instead.

For more detailed information about the tools pack, see "Windows Server 2003 Administration Tools Pack Overview" in Help and Support Center for Windows Server 2003. Help topics for the Administration Tools Pack are installed when you install the tools.

### Note

- You cannot install the Windows Server 2003 Administration Tools Pack on a server that is running a member of the Microsoft® Windows® 2000 Server or Windows Server 2003 family operating systems. The administrative tools already exist on all servers running these systems. You can install the tools pack only on a computer that is running Windows XP Professional with Service Pack 1.

**Group Policy**

In an Active Directory environment, you can use Group Policy to control such things as permissions, application availability, and security for member servers, domain controllers, and any other server running Windows Server 2003 within the scope of management. You can use Group Policy to manage registry-based policy by using Administrative Templates and to assign scripts, such as for startup and shutdown. For more information about Group Policy, see the following:

- "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit.

- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit).

# Choosing Out-of-Band Management Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you cannot manage a server by using conventional in-band tools for some reason, an out-of-band connection is the only way to remotely manage it. If you configure your servers for remote administration, and you cannot access them by using your in-band tools, an out-of-band connection is the only way to manage them.

To configure a server for out-of-band management, you need to consider software, firmware, and hardware. Emergency Management Services, which is included with Windows Server 2003, is the principal out-of-band component. With only Emergency Management Services and a serial port, you can manage most Windows Server 2003 operating states. When you combine Emergency Management Services with supporting firmware and hardware components, you can also perform tasks ranging from powering up computers to recovering unresponsive systems — everything, in fact, except for replacing and installing hardware.

The additional components you choose depend on which tasks you want to perform remotely, how much you are willing to spend for extra features, and how many servers you have to manage. The following tools and components work with Emergency Management Services to support out-of-band remote management:

- Firmware — BIOS for x86-based computers or EFI for Itanium-based computers — that provides console redirection
- Serial ports and modems
- Terminal concentrators
- Service processors
- Intelligent UPSs or intelligent power switches

By selecting the optional components that best meet your remote management requirements, you can capitalize on the full range of out-of-band management capabilities supported by Emergency Management Services.

For more information about selecting hardware components, see "Best practices for selecting and configuring hardware" in Help and Support Center for Windows Server 2003.

Table 5.3 shows which out-of-band tools support various operating states and remote management tasks.

**Table 5.3 Components Required for Out-of-Band Situations**

| Operating State or Task | Type of Tool |
|---|---|
| Windows Server 2003 is starting | Emergency Management Services |
| Server fails to fully initialize | Emergency Management Services |
| Administrator needs to run Recovery Console | Emergency Management Services |
| Server is not functioning due to stop message | Emergency Management Services |
| System is low on resources, resulting in slow or no response to requests | Emergency Management Services |
| Network stack has malfunctioned or failed | Emergency Management Services |
| System is not responding on the network | Emergency Management Services |

| | |
|---|---|
| System is not responding on the network or to Emergency Management Services | Service processor |
| System is powered down | Wake-on-LAN network adapter*, intelligent UPS, intelligent power switch, or service processor |
| BIOS is conducting POST | Redirecting firmware or service processor |
| Change firmware configuration settings | Redirecting firmware or service processor |
| Operating system installation by using RIS | Emergency Management Services (see "Selecting the Installation Method" later in this chapter |

* Wake-on-LAN is a combined hardware and software technology that allows you to remotely turn on Advanced Configuration and Power Interface (ACPI)–compliant computers. Several vendors provide Wake-on-LAN remote management solutions. Some vendors support this functionality across a local area network (LAN) or wide area network (WAN), while some support it over the Internet. For more information about Wake-on-LAN technology, use a Web search engine and search using the keyword "Wake-on-LAN."

Some trade-offs you might experience with out-of-band components include:

- Limited maximum throughput.
- No GUI support.
- Optionally, additional hardware requirements.

# Emergency Management Services in Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To manage a server from a remote computer when the server is not available on the network, you must enable Emergency Management Services. Emergency Management Services is a Windows Server 2003 service that runs on the managed server. This service is not enabled by default when you install the Windows Server 2003 operating system, but you can enable it during installation or at any later time.

Emergency Management Services features are available when the Windows Server 2003 loader or kernel is at least partially running. You can access all Emergency Management Services output by using terminal emulator software that supports VT100, VT100+, or VT-UTF8 protocols on the management computer, although VT-UTF8 is the preferred protocol. For more information about terminal emulator software and the supported protocols, see "Management Software for Out-of-Band Connections" later in this chapter. For more information about enabling, configuring, and using Emergency Management Services, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

When Emergency Management Services is enabled:

- Console redirection automatically sends output to the out-of-band port for any supported operating state, as indicated in Table 5.4.
- You can use SAC to issue supported commands or switch to the command shell (cmd.exe) whenever the kernel is running.
- You can view logs during the GUI-mode phase of Setup.
- !SAC automatically becomes available whenever a system failure occurs.

Table 5.4 shows when you can use Emergency Management Services features for remote management, with or without special out-of-band hardware.

**Table 5.4 Using Emergency Management Services Features**

| Task | Feature |
|---|---|
| Selecting operating system during system load | Console redirection |
| Running Recovery Console | Console redirection |
| Viewing text mode setup messages | Console redirection |
| Viewing GUI mode setup messages | SAC, including setup logs |
| Viewing RIS loading messages | Console redirection |
| Viewing Stop error messages | Console redirection |
| Monitoring and managing with out-of-band connections | SAC |
| Performing last-resort system recovery | !SAC |

**Emergency Management Services Console Redirection**

Emergency Management Services console redirection redirects the output from supported Windows Server 2003 functions to the out-of-band port. When Emergency Management Services is enabled, you can perform remote management through the out-of-band port, as shown in Table 5.5.

**Table 5.5 Emergency Management Services Console Redirection**

| Managed Operating State | Example Tasks |
|---|---|
| Windows Server 2003 Loader | • Select the operating system to load on x86-based multiple-boot systems.<br>• Verify the load of Windows Server 2003 components before in-band tools become available. |
| Kernel at least partially functioning | • Perform SAC commands, such as changing the priority of a process.<br>• Perform !SAC commands, such as viewing Stop messages when a system problem occurs. |
| Recovery Console running | • Troubleshoot startup problems. |
| Text-mode Setup | • View Windows Server 2003 Setup progress.<br>• Respond to text-mode Setup prompts. |
| GUI-mode Setup | • Perform SAC commands and monitor setup logs. |
| RIS-based Setup | • Respond to the **F12** prompt to initiate RIS-based Setup. |

**Note**

- You must have firmware redirection to view server information before the Windows Server 2003 operating system starts.

**Special Administration Console**

When Emergency Management Services is enabled, SAC is always available through the specified out-of-band port, as long as the Windows Server 2003 kernel is running. You can use SAC at any time to carry out out-of-band management commands during the following system operating states:

- Normal system operation
- Windows Server 2003 components initialization
- Safe mode
- GUI-mode during Windows Server 2003 Setup

The SAC prompt appears when you connect to a server that is running Emergency Management Services. The SAC command-line environment supports a specific set of commands. For information about SAC commands, see "Special Administration Console (SAC) and SAC commands" in Help and Support Center for Windows Server 2003.

Using SAC, you can perform management tasks such as the following:

- Gathering server information, such as computer name and IP address.
- Changing a server's TCP/IP networking information to resolve issues caused by incorrect parameters or a duplicate IP address.
- Obtaining a list of processes and threads running on the computer to determine if they are causing a system

performance problem, if you cannot perform this task by using in-band tools.

- Raising or lowering the priority of a process, or ending a process that is consuming excessive server processor resources or other system resources to eliminate performance issues.
- Restarting or shutting down a server as part of unplanned maintenance task, when the in-band mechanism fails.
- Setting the system time and date, for example, for Kerberos authentication.
- Starting a command shell and running text-based tools, and switching between the command prompt and SAC.
- Viewing setup logs during GUI mode setup and switching between the setup logs and SAC.

**!Special Administration Console**

When Emergency Management Services is enabled and a system failure occurs, !SAC — an abbreviated form of SAC — automatically replaces SAC as the command-line environment. For information about !SAC commands, see " !Special Administration Console (!SAC) and !SAC commands" in Help and Support Center for Windows Server 2003.

**Important**

- !SAC is not available if the debugger is running or the system is set to restart automatically when Stop errors occur.

Using !SAC, you can perform tasks such as the following:

- View redirected Stop messages.
- Display computer identification information.
- View an abbreviated log of loaded drivers and some kernel events.
- Restart the computer.

# Using serial ports for out-of-band management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To perform out-of-band management, you need to establish a secure connection through a serial port, phone line, or an additional network connection. The serial port, also known as a COM port, is the most common out-of-band interface. It is the default out-of-band device for Emergency Management Services. You can provide remote access to an out-of-band serial port by using modems or terminal concentrators, as described later in this section.

To use the serial port as an out-of-band device with Emergency Management Services, it must meet the following requirements:

- The serial port must be a standard 16450 or 16550 Universal Asynchronous/Receive Transmit (UART) device. Windows Server 2003 tests the device for compliance before using it with Emergency Management Services.
- The serial port interface must be provided by hardware, not by a Windows driver.
- If the system firmware is compatible with Emergency Management Services, the firmware and the serial port must be configured to use the same serial port settings.
- A kernel debugger cannot share the same COM port. To avoid this problem, disable kernel debugging on servers with Emergency Management Services enabled. For more information about using kernel debuggers with Emergency Management Services, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).
- The serial port must be the only out-of-band management port. Emergency Management Services does not support one out-of-band port for outbound communication and a second port for inbound communication.

Direct serial connections provide no logical security and therefore must be secured physically. For more information about security considerations, see "Providing Security for Remote Management" later in this chapter.

**Note**

- When you use Emergency Management Services with a serial port, use null modem cables that support the carrier detect signal.

For Emergency Management Services technical details and information about configuring serial port settings, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Using Modems for Remote Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A modem can provide remote access to a single server through its out-of-band serial port. Modems can be useful when you have just a few servers to manage remotely, such as a branch office with one or two servers. For an example of this type of implementation, see "Designing the Hardware Configuration" later in this chapter.

Take the following considerations into account when selecting a modem to use for out-of-band remote management:

- Security for modem connections is determined by the modem. The preferred security mechanism for modems is dial-back.

- The modem must be configurable and must not rely on initialization. Emergency Management Services does not initialize the modem, so you must configure the modem to answer or dial back automatically and pass all serial data through unchanged. For information about configuring modems for callback, see "Configure client callback options" in Help and Support Center for Windows Server 2003.

# Using Terminal Concentrators for Remote Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Terminal concentrators provide remote access to multiple servers through their out-of-band serial ports. The servers connect to the serial ports on the terminal concentrator with null modem cables. The remote management computer establishes a network connection to the terminal concentrator by using its network port. Typically, you use Telnet or a Web interface to remotely perform management tasks on the servers connected to the terminal concentrator.

Terminal concentrators facilitate remote management of servers in the following ways:

- You can manage servers using a serial connection without being within the distance of a serial cable length.
- You can monitor and manage multiple servers simultaneously from a single management computer.
- Several administrators can simultaneously view information for different servers.

Setup, configuration, and features for terminal concentrators vary by manufacturer. When choosing a terminal concentrator, assess the following features:

- Number of available serial ports.
- Built-in security features, such as use of passwords and encryption.
- Power switch capabilities.
- Number of available Ethernet ports, if important in your environment.

Some terminal concentrators support Secure Shell (SSH), which is a secure command-line alternative to Telnet. SSH is a protocol for establishing secure connections over networks. It provides logical security for the in-band connection from the management computer by supporting strong authentication and encryption and protecting against a variety of network-level attacks. Because SSH is independent of the operating system, it provides interoperability in environments with mixed operating systems. Several vendors provide Windows implementations of SSH clients and servers. Use a Web search engine and search using the keyword "SSH" to find a variety of SSH vendors, as well as frequently asked questions (FAQs) and other documentation.

You need to provide physical security for the serial connections from the servers to the terminal concentrator. Because the security features for terminal concentrators are not standardized, you might need to provide your own logical security for the in-band connection. If your terminal concentrator does not support authentication and encryption, consider using one of the following techniques to secure the connection:

- Use a secondary private management network that you can access with direct-dial remote access or with a VPN connection.
- Use a router to secure the network traffic.
- Use SSH, if the terminal concentrator supports it, instead of Telnet to provide authentication and encryption.

# Console Redirection by Firmware

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Console redirection provided by system firmware (either BIOS for x86-based computers or EFI for Itanium-based computers) provides out-of-band access to server information before the Windows Server 2003 operating system starts. Firmware console redirection works together with Emergency Management Services console redirection to provide out-of-band support for any operating state.

If your firmware does not provide console redirection — and you do not have a service processor that provides console redirection, as described later in this chapter — you cannot remotely manage servers during the time between system restart and the initial loading of the Windows Server 2003 operating system.

If you configure your servers for remote administration, you need to have either firmware console redirection or a service processor with console redirection so that you can access the servers when the Windows Server 2003 operating system is not functioning.

**Note**

- Firmware console redirection typically redirects only during text mode, not during GUI mode.

By using firmware console redirection, you can perform the following out-of-band tasks from a remote computer:

- View server status before the operating system starts up. For example, you can view POST status or disk-related error messages. Firmware console redirection typically allows the POST to complete without a local keyboard, mouse, or monitor.
- View and make modifications to firmware settings, such as disabling a peripheral device or changing boot sequence, with the built-in firmware configuration program.
- View master boot record (MBR) errors.
- Start a RIS-based setup by responding to the **F12** network boot prompt. This support is required only if the **F12** prompt is presented by the firmware.
- Boot the computer from the CD drive by responding to the **Press Any Key to Boot from CD** prompt.

When assessing firmware console redirection for use in conjunction with Emergency Management Services, verify that the firmware meets the following criteria:

- Shares the serial port with Emergency Management Services and releases control to Emergency Management Services after the Windows operating system starts.
- Supports VT-UTF8, VT100+, or, at minimum, VT100 terminal emulator conventions. For more information, see "Management Software for Out-of-Band Connections" later in this chapter.
- Preferably, supports the Serial Port Console Redirection (SPCR) table.

An SPCR table is provided with some ACPI-compliant system firmware for specifying how the out-of-band management port is used. For example, if the out-of-band port is a serial port, the SPCR table contains information such as serial port number, baud rate, terminal type, and other settings used for out-of-band communication. When this table exists, Emergency Management Services uses the information in it to ensure a consistent transition between firmware console redirection and Emergency Management Services console redirection. The SPCR table is recommended for use with both 32-bit and 64-bit versions of Windows Server 2003. For more information about the SPCR table and Emergency Management Services, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Using service processors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Service processors provide robust remote management support that is independent of the operating system. Because Emergency Management Services is available only when the loader or kernel is at least partially running, you might need such support when the system experiences severe problems that cause it to stop responding completely. Table 5.3, included earlier in this chapter, shows some of the operating states supported by service processors but not by Emergency Management Services. Consider a service processor if you need a high degree of reliability and availability for your servers or you decide to configure your servers for remote administration.

Typically, service processors are integrated into the system motherboard or into an add-in PCI adapter. Servers that have on-board service processors might offer higher out-of-band throughput by using higher-speed serial or Ethernet connections. Service processors operate independently from the main processor, use their own custom firmware, and sometimes include their own power supply. When you connect to a server through an out-of-band connection, you can communicate directly with the service processor.

Service processor features, client interfaces, and management tools vary by manufacturer. If you plan to use the service processor with Emergency Management Services, it is recommended that the service processor support these functions:

- Console redirection
- Remote power on and power off
- Remote reset
- Access to Emergency Management Services at all times

To be compatible with Emergency Management Services, make sure that the service processor also meets the following requirements:

- If the service process uses the serial port as its interface, it must share the serial port with Emergency Management Services and must release control to Emergency Management Services after the operating system has started.
- The UART interface must be described in the SPCR table, or in the EFI console device path for the 64-bit versions of Windows Server 2003.
- It supports VT-UTF8, VT100+, or, at minimum, VT100 terminal emulator conventions. For more information, see "Management Software for Out-of-Band Connections" later in this chapter.

Manufacturers offer a wide range of additional features. Evaluate the features and tools provided to ensure that they meet your needs. Additional features you might consider include:

- Access to Emergency Management Services through hardware interfaces other than serial, such as modem or RJ-45 Ethernet. The type of connection determines the additional components you need and the security requirements for out-of-band access to your server. For more information about configuring components for a service processor, contact the manufacturer.
- Console redirection of GUI screens.
- Any of a variety of management and troubleshooting tools.
- Client interfaces that range from simple Telnet consoles to complex Web browsers.

# Intelligent UPSs and Intelligent Power Switches

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An uninterruptible power supply (UPS) provides backup power to a server in the event of a power failure. Some UPS units, known as intelligent UPSs, and intelligent power switches can provide limited remote management capabilities such as powering up, powering down, and resetting a server.

If an intelligent UPS or intelligent power switch is integrated with terminal concentrator functionality, it can provide pass-through serial connection between the management computer and the server running Windows Server 2003. In this case, the port on the management computer has a serial connection to an external serial port on the UPS or power switch, which in turn has a serial connection to the server.

The situation just described can provide a more economical solution than an internal service processor: You can access the intelligent UPS, firmware console redirection, and Emergency Management Services through the same communication channel if all these components use the same terminal conventions, such as VT-UTF8. When the components share the same terminal conventions, each component can consistently interpret escape sequences passed to it. For an example of this type of configuration, see "Designing the Hardware Configuration" later in this chapter.

If the intelligent UPS or intelligent power switch shares the same management channel with Emergency Management Services, the UPS or power switch must passively monitor the serial data stream and respond only when it detects VT-UTF8, VT100+, or VT100 escape sequences that apply to it. For more information about using an intelligent UPS or intelligent power switch with Emergency Management Services, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

If you plan to use an intelligent UPS or intelligent power switch with Emergency Management Services, the server running Windows Server 2003 must be configured to start automatically when power is applied.

**Tip**

- The firmware for your server might provide a configuration option for automatically starting the computer when power is applied.

# Management Software for Out-of-Band Connections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, you use terminal emulation software on the management computer to connect to and communicate with a server through an out-of-band connection. The two most common methods are the following:

- Use Telnet — or a secure alternative such as SSH — to connect to a terminal concentrator through an in-band connection, which then connects to the server through an out-of-band connection.
- Use HyperTerminal to connect directly to the server.

If you use a service processor, it might require specific software to work with it and to interact with Emergency Management Services. For example, manufacturers might provide a Web browser or custom software.

Make sure that the terminal emulation software you use supports serial port and terminal definition settings that are compatible with Emergency Management Services, as well as with your service processor or system firmware. If possible, use terminal emulation software that supports the VT-UTF8 protocol because VT-UTF8 support for Unicode provides for multilingual versions of Windows. If English is the only language you need to support, the VT100+ terminal definition is sufficient. At minimum, you can use the VT100 definition, but this terminal definition requires that you manually enter escape sequences for function keys and so forth. For more information about terminal definitions and what they support, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit). For more information about the VT-UTF8 terminal definition, see the Emergency Management Services Design link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

If your hardware and firmware use the same terminal definition settings that Emergency Management Services uses, you can always use the same escape sequences for managing computers, regardless of what is controlling the port — hardware, firmware, or Emergency Management Services. If you use different terminal types, the escape sequences you need to use vary depending on what is controlling the serial port, making it difficult to determine the appropriate sequence to send.

# Examples: Selecting Remote Management Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you define your remote management requirements, you can select the tools and components you need to perform remote management tasks. The following illustrations describe the tools for the business situations and remote management levels described in "Examples: Determining Remote Management Requirements" earlier in this chapter.

**Minimal Remote Management**
Because the branch office in this business situation has low availability requirements and has on-site support, this situation requires no special out-of-band support. Emergency Management Services together with the basic serial port can provide out-of-band management support for the few times when administrators need to perform remote out-of-band tasks. Although this organization decided to use existing hardware, the OEM can provide a BIOS upgrade that supports console redirection. Remote Desktop for Administration provides remote in-band support in this situation.

For this branch office, administrators can perform tasks such as the following remotely:

- All in-band tasks, such as monitoring, configuring, and troubleshooting when the operating system is responding over the network.
- Monitoring operating system startup.
- Troubleshooting when the operating system is low on resources, such as when CPU usage is excessively high.
- Running Recovery Console.
- Monitoring firmware initialization and configuring firmware settings.

Administrators must perform the following types of tasks locally:

- Powering up and powering down.
- Resetting.
- Viewing POST status.
- Troubleshooting when the operating system is unresponsive.

For an illustration of out-of-band hardware configuration for this example, see "Configuring for Direct Serial Connection" later in this chapter.

**Moderate Remote Management**
Because the data center in this business situation has many high-availability servers, the organization does not want to incur the costs of upgrading to new computers to obtain out-of-band management support. The BIOS for some of the existing servers can be upgraded to provide console redirection; for other servers, such an upgrade is not available. Emergency Management Services together with a serial port provides out-of-band support when the Windows Server 2003 operating system is functioning.

This data center is configured with terminal concentrators to consolidate access to many servers by using a single in-band network connection. These terminal concentrators have SSH built in to provide authenticated and encrypted remote console sessions. They also have a built-in intelligent UPS for power backup and remote power functionality.

In this data center, administrators can perform tasks such as the following remotely:

- All in-band tasks, such as monitoring, configuring, and troubleshooting when the operating system is responding over the network.
- Monitoring operating system startup.

- Troubleshooting when the operating system is not responding over the network, such as when CPU usage is excessively high.
- Running Recovery Console.

Administrators must perform the following tasks locally:

- For servers without firmware console redirection, viewing POST results.
- For servers without firmware console redirection, monitoring firmware initialization and configuring firmware settings.
- Troubleshooting when the operating system is not responding to Emergency Management Services or the network.

For an example configuration for this example, see "Configuring for Terminal Concentrator Connections" later in this chapter.

**Maximum Remote Management**
In this business situation, the remotely administered servers support Emergency Management Services and include firmware console redirection and built-in service processors. The servers are configured with a terminal concentrator that has built-in intelligent UPS functionality and built-in SSH for secure communications.

Administrators perform all tasks — except hardware replacement — remotely. The Windows Server 2003 operating system is also installed remotely by using RIS.

For possible out-of-band hardware configurations that are appropriate for this business situation, see "Configuring for Intelligent UPS and Terminal Concentrator" and "Configuring for Network Service Processor" later in this chapter.

# Designing the Hardware Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because in-band management occurs through the standard network connection to a responding operating system, in-band management tools are typically software applications and tools. Out-of-band management, on the other hand, might require hardware components to access a non-responding server or to consolidate access to multiple servers. After you have determined which tasks you want to perform remotely and which tools and components to use, the next step is to determine how to configure the components, as Figure 5.4 illustrates.

**Figure 5.4 Designing the Hardware Configuration**




This section describes and illustrates ways to configure various out-of-band components. These configurations are only sample solutions. Depending on your hardware, you might require a different solution. For more detailed information about setting up out-of-band connections in an Emergency Management Services installation, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

Although you can use Emergency Management Services with only a serial port or modem to manage most Windows Server 2003 operating states, this minimal configuration is useful only if you are managing just one or two servers. If you are managing many servers, this minimal configuration might be of limited practical application.

When you combine Emergency Management Services with additional components, such as a terminal concentrator, an intelligent UPS or intelligent power switch integrated with terminal concentrator functionality, or a network-enabled service processor, you can use a single network connection from your management computer to manage multiple servers through their out-of-band ports. This configuration can be effective when you need to manage many servers in a single location.

The following configurations are described in order from least to most robust.

# Configuring for Direct Serial Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 5.5 illustrates a direct serial connection. Using a null modem cable, you can connect the management computer to the server that has Emergency Management Services enabled. This configuration is the simplest out-of-band configuration.

**Figure 5.5 Direct Serial Connection**




The management computer runs terminal emulation software that preferably supports VT-UTF8 terminal conventions.

Because null modem connections have no built-in security, this configuration depends on physical security. For more information about planning for security, see "Providing Security for Remote Management" later in this chapter.

This type of configuration has the following advantages:

- No need to buy additional hardware
- Easiest to set up and configure
- Good for situations with very few servers to manage

Disadvantages of this type of configuration include:

- Management computer must be in close proximity for physical security
- Least robust configuration
- Logical security is not supported
- Ineffective in situations with many servers to manage

A variant of this configuration might be a server with a service processor that provides the serial connection. If a service processor and firmware console redirection are included in the configuration, you can manage all operating states through the out-of-band connection. If Emergency Management Services is the sole provider of console redirection, you can perform out-of-band management only when the Windows Server 2003 loader or kernel is at least partially running.

With this configuration, you can be ready to quickly connect a laptop to a remotely administered server to perform management tasks.

# Configuring for Remote Serial Connection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 5.6 illustrates a remote serial connection using a modem. The modem is connected to the server that is running Emergency Management Services, and the remote management computer connects over the phone line.

**Figure 5.6 Serial Connection with Modem**




The management computer runs terminal emulation software that preferably supports VT-UTF8 terminal conventions. The modem functionality might be provided by an external modem or by a service processor with built-in modem functionality.

This configuration depends on physical security between the server and the modem and modem-based security for the phone connection. For security, modems that support dial-back are preferred. If you use a service processor, it might provide additional security options.

This type of configuration has the following advantages:

- No need to buy additional hardware
- Easy to set up and configure
- Good for situations with very few servers to manage
- Management computer can be in a remote location

Disadvantages of this type of configuration include:

- Security is not robust
- Ineffective in situations with many servers to manage

This configuration can be used for the minimal remote management example described in "Examples: Selecting Remote Management Tools" earlier in this chapter.

# Configuring for Terminal Concentrator Connections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 5.7 illustrates a terminal concentrator with direct serial connections to multiple servers running Emergency Management Services. The optional modem connects to the terminal concentrator. The remote management computer can connect to the terminal concentrator over the network for an in-band connection or over the phone line for an out-of-band connection when the network is not available.

**Figure 5.7 Remote Connections with Terminal Concentrator**




The servers running Emergency Management Services connect to the terminal concentrator with serial connections using the VT-UTF8 terminal definition. The serial connections might be provided by null modem cables or by service processors. The number of serial ports provided by the terminal concentrator determines the number of servers that can be supported in this configuration. If the terminal concentrator is integrated with intelligent power switch functionality, you can use this configuration to power servers off and on.

The remote management computer connects to the terminal concentrator over the standard network by using an in-band tool, such as Telnet or SSH, that supports the VT-UTF8 terminal definition.

This configuration requires physical security between the servers and the terminal concentrator. In addition, it is important to provide logical security for the terminal concentrator. Because access to a terminal concentrator provides access to multiple servers, it is important to ensure that only authorized persons can gain access, as described in "Providing Security for Remote Management" later in this chapter, thus protecting the connected servers from denial-of-service attacks. The terminal emulation software supported by the terminal concentrator helps determine the logical security between the management computer and terminal concentrator. For example, some terminal concentrators have SSH built in to provide authentication and encryption. If the terminal emulation software does not support authentication and encryption and the network is accessible to outsiders, an unauthorized person can sniff the network to obtain credentials for your servers. Another solution is to use a private, secondary network for remote management.

The advantages of this type of configuration are:

- Supports remote management for multiple servers
- Combined with firmware console redirection, provides broad functionality
- Supports logical security
- Can support the use of legacy computers

The primary disadvantage of this configuration is that it requires additional hardware.

This configuration is an excellent alternative to service processors. When combined with firmware redirection, this configuration provides functionality comparable to service processors without the additional costs. It provides a way to obtain state-of-the-art technology with legacy computers. For a document to assist you in setting up remotely administered servers in

this type of configuration, see "Headless Server Quick Start" (SDCEMS_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Headless Server Quick Start" on the Web at http://www.microsoft.com/reskit).

This configuration can be used for the moderate remote management example described in "Examples: Selecting Remote Management Tools" earlier in this chapter.

# Configuring for Intelligent UPS and Terminal Concentrator

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 5.8 illustrates a remote management computer with a network connection to an intelligent UPS that also functions as a terminal concentrator.

**Figure 5.8 Intelligent UPS and Terminal Concentrator**




The servers running Emergency Management Services have two connections to the intelligent UPS/terminal concentrator: an AC power connection and a serial connection that supports the VT-UTF8 convention. The serial connection can be provided by a null modem cable or by a service processor. The number of ports provided by the UPS determines the number of servers that can be supported in this configuration. The remote management computer connects to the UPS over the standard network by using an in-band tool, such as Telnet.

In this configuration, terminal concentrator and intelligent UPS functionality are combined, allowing you to manage multiple servers, as well as remotely control power to them.

This configuration requires physical security between the servers and the intelligent UPS. The network connection requires logical security, as described for the terminal concentrator configuration.

This configuration can provide an economical alternative to internal service processors for legacy computers that do not support service processors.

Advantages of this configuration include:

- Provides much of the functionality of a service processor
- Does not require a new component for each server (as service processors do)
- Supports remote management of multiple servers
- Supports logical security

The primary disadvantage of this configuration is that it requires additional hardware.

This configuration can be used for the moderate remote management example described in "Examples: Selecting Remote Management Tools" earlier in this chapter.

# Configuring for Network Service Processor

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 5.9 illustrates a remote management computer connected to servers with network-enabled service processors.

**Figure 5.9 Network Service Processor**




The servers in this configuration each have an internal network-enabled service processor. The service processors might be integrated into the motherboard or on an add-in PCI card.

The remote management computer connects to the servers through the network using an in-band tool such as Telnet or a proprietary tool provided by the service processor manufacturer.

The security in this configuration depends on the service processors. Using the Secure Sockets Layer (SSL) protocol is the preferred method. Using a secondary management network is another option for providing security.

Advantages of this configuration include:

- Does not require physical security for connections to servers
- Supports remote management for multiple servers

The primary disadvantage of this configuration is the cost of service processors, which are needed for each managed server.

This configuration can be used for the maximum remote management example described in "Examples: Selecting Remote Management Tools" earlier in this chapter.

# Configuring Your Infrastructure for Remote Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remote administration, whether in-band or out-of-band, might require changes to your environment in areas such as network connectivity, memory and processor resources, and security. After you select the tools you plan to use for remote management, assess the impact they have on your environment, and decide how to configure hardware components, the next step is to determine the changes you need to make to your infrastructure to mitigate any potential negative impacts. The place of this step in the process is illustrated in Figure 5.10.

**Figure 5.10 Configuring Your Infrastructure for Remote Management**




For more information about ways that remote administration might affect your environment, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Ensuring Network Connectivity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remote administration requires a stable, unrestricted connection between the management computer and the server being managed. In some cases, you also might need a high-speed connection. The tools you select and the location of the management computer in relation to the server being managed determine the network connectivity issues you need to address. Issues you need to address include available network bandwidth, the type of connection you plan to use, firewalls that might prevent access, and routers that might prevent access due to IP packet filtering.

**Network bandwidth**
Because remote administration tools can increase network traffic, you need to assess your current traffic levels and the capacity of your network to handle additional traffic that might result from the tools you decide to use. Typically, remote administration tools increase traffic only by a small amount, but some tools can increase traffic significantly. Remote Desktop for Administration, for example, is resource intensive and can create a significant amount of traffic.

**Type of connection**
If you select resource intensive remote management tools, you also need to assess the type of connection you plan to use with the tool. Typically, high-speed connections available on a LAN or WAN can, depending on your current traffic level, accommodate resource intensive tools. Dial-up, digital subscriber line (DSL), or broadband digital cable connections across a VPN, however, are relatively slow compared to LAN and WAN connections and might result in poor performance for some remote management tools.

If you plan to perform remote management tasks across a VPN, you need to configure the networking settings for the connection. For information about specifying these settings, see "VPN remote access for employees" in Help and Support Center for Windows Server 2003.

**Firewalls**
Because firewalls restrict network traffic, they can prevent you from remotely managing computers from outside the firewall. If you plan to use a management computer outside the firewall to manage servers inside the firewall, ensure that you can connect to them remotely through the firewall and that the firewall supports the network protocols and ports used by your tools. You might need to reconfigure firewall settings to permit remote management through the firewall.

For more information about the impact of firewall settings on remote management, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**IP packet filtering**
If you use routers that filter IP packets based on source and destination IP addresses and TCP ports, ensure that you can remotely access the servers you need to manage remotely. You might need to reconfigure IP packet filtering settings on your routers to permit remote management from your management computer.

For more details about configuring your network for remote management, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Ensuring Memory and Processor Resources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remote management software tools that run on a management computer use memory, disk, and processor resources. Remote management services that run on a managed server use memory and processor resources. Although you typically do not need to add memory or processor resources for remote management, you might want to ensure that the tools you choose do not degrade server performance, especially on critical servers such as domain controllers.

Remote management tools that support graphical user interfaces use more processor resources than command-line tools do. If you are concerned about the effect of a specific remote management tool on server performance, test the tasks you plan to perform with the tool to ensure that the current memory and processor resources can accommodate remote management. The task you perform with a tool affects processor activity more than the tool itself.

Remote management tools use memory on both the management computer and the server. If the computers meet minimum hardware requirements, remote management tools typically do not cause memory usage performance problems. If you optimize memory by terminating all nonessential services and processes, such as on a Web server, you might want to verify that your remote management tools do not have an impact on performance. For more information about assessing memory usage on remote management computers and servers, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Providing Security for Remote Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remote administration introduces new security considerations into your environment. When you manage servers remotely, sensitive information that normally is not transmitted across a network is sent over your network. For example, server identifying information, configuration information, and other sensitive management information such as user names and passwords can be transmitted. You need to ensure that your remote management tools and tasks do not expose this sensitive data to someone sniffing or eavesdropping on your network. In addition, when you use serial ports for out-of-band management, the null modem connections between the servers and the management computer or other out-of-band hardware component provide no logical security against unauthorized access.

When planning security solutions for remote management, you need to protect against intentional acts as well as accidents. For in-band remote management, you need to consider solutions such as authentication and encryption. If you plan to use dial-up, DSL, or broadband digital cable connections across a VPN, you also need to plan your firewall configuration. For out-of-band remote management, you need to plan physical security solutions to protect the inherently insecure serial connections. Finally, you need to determine a strategy for user rights and shared folder permissions so that only authorized administrators can perform authorized management tasks.

As you plan your remote management security strategy, you need to make sure that:

- The server allows administrative commands only from an authenticated computer.
- The server accepts administrative commands only from an authenticated administrator.
- Confidential information — including administrative commands and configuration settings — cannot be intercepted, read, or changed by intruders.
- Log files are viewed by using a secure method.

A secondary network built specifically for remote management can increase security, performance, and availability. You can control access to such a management network by using a secure router.

For more detailed information about assessing security risks inherent in remote management and an overview about how to mitigate or eliminate these security vulnerabilities, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Authentication**

When you perform remote administration, you need to log on to the remote computer you want to manage. Remote management tools use several different authentication protocols — some stronger than others — to ensure that only authorized users can access computers remotely. For example, some tools use the Kerberos version 5 authentication protocol and others use the NTLM authentication protocol. Kerberos authentication is more secure than NTLM authentication.

You can mitigate the vulnerabilities of less secure authentication protocols by configuring one or more Group Policy settings. Configure these policy settings for maximum protection if either of the following is true:

- You are administering remote computers in an environment that forces NTLM authentication.
- You are administering remote computers with remote management tools that use NTLM authentication.

For information about environments that force NTLM authentication and the description and location of Group Policy settings you can use with NTLM, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Encryption**

Some remote management tools encrypt data — including passwords — before transmitting it across the network, while others do not. Unencrypted data makes your network vulnerable to eavesdropping and sniffing.

If you decide to use a remote management tool that does not encrypt or otherwise secure data, you can mitigate the security

issue by using Internet Protocol security (IPSec) to encrypt the communication between the management computer and the server. When you use IPSec, IP packets can pass securely through routers or other computers that do not support IPSec. You administer IPSec by using policies, which you can configure for the specific security requirements of individual computers, domains, organizational units, sites, or your entire enterprise. If you plan to support dial-up remote management, consider using IPSec across a VPN connection. For more information about VPN and about using IPSec with VPN, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

For information about IPSec policies, see "Internet Protocol Security (IPSec)" in Help and Support Center for Windows Server 2003. For more information about IPSec in general, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

For detailed information about using encryption for remote management, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Physical Security**

Although corporate servers must always be situated in secure locations, out-of-band management introduces another physical security issue: the serial connections between servers and out-of-band management components, such as a remote management computer or a terminal concentrator, need to be protected physically because null modem connections provide no logical security. Some ways to provide physical security include:

- Keeping server rooms locked with secured access, such as keys, smart cards, or passwords.
- Using terminal concentrators or intelligent UPSs to consolidate access to servers and keeping these out-of-band hardware components in the same secured room with the servers.
- Keeping cable lengths short to prevent the possibility of extending them outside the secured room.

**Rights and Permissions**

After you know which servers you plan to manage remotely and which administrators are responsible for specific administrative tasks, you need to set up security groups and assign administrators membership in order to grant them access to remote resources. As you define your security groups, set up administrative tasks with the minimum necessary administrative credentials. By using this technique, you can avoid assigning users a higher security level than they need to perform the tasks for which they are responsible. For recommendations about assigning permissions and user rights, see "Best practices for permissions and user rights" in Help and Support Center for Windows Server 2003.

Two types of security considerations are important in remote administration: user rights and shared folder permissions.

**User rights** User rights control the tasks you can perform on a computer, such as setting up user accounts or installing hardware. Depending on the security model and the group structure you use, you might have to configure user rights on each server and management computer, or you might be able to configure them on the domain controller.

**Shared folder permissions** Shared folder permissions control which users or groups can gain access to the contents of a shared folder remotely over the network, as well as which actions users or groups can perform on the contents of those folders. You can configure shared folder permissions on the server and enable users to gain access to the folders remotely over the network. For example, you can assign Read or Full Control.

You need to configure user rights and shared folder permissions if administrators need to do the following:

- Access the administrative shares on a remote computer.
- Log on to computers remotely by using terminal emulation or command console programs.
- Access files or folders on a remote computer.

You can centrally control remote management by using Group Policy settings related to remote management. Group Policy settings for computer configuration include security settings that restrict how a user can access files, folders, and computers, as well as administrative template settings that change the behavior and appearance of remote management tools and technologies, such as Terminal Services.

**Important**

- Terminal Services is affected by the Internet Explorer Enhanced Security Configuration, which places your server and Microsoft Internet Explorer in a configuration that decreases the exposure of your server to attacks that can occur through Web content and application scripts. As a result, some Web sites might not display or perform as expected. For more information, see "Before Installing Terminal Server" and "Internet Explorer Enhanced Security Configuration" in Help and Support Center for Windows Server 2003.

For information about groups, user rights, permissions, and authorization and access control, see the *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit). For information about configuring user rights and permissions for remote management, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Secondary Management Network**

In addition to authentication, encryption, and user rights, you can add an extra layer of network security by placing your remote management system on a separate network segment and control access by using a secure router, as shown in Figure 5.11. You can use this configuration to control exactly which users and computers are allowed access to the management system.

**Figure 5.11 Secondary Management Network**




In this configuration, the servers are connected to the terminal concentrator with null modem cables, and all these components are located in a secure room. The management computer can access the servers by connecting to the terminal concentrator through the secure router. The management computer can use an in-band connection or a remote access connection through a remote access server.

# Planning for Remote Management Deployment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you create your remote management strategy and determine your tools and configurations, the remaining step is to plan for deployment, as illustrated in Figure 5.12.

**Figure 5.12 Planning for Remote Management Deployment**




If you decide to use Emergency Management Services for out-of-band management, evaluate your existing server hardware to determine whether it supports the out-of-band functionality you need. After you decide whether to use your existing hardware or to purchase new servers or components, you can decide how you want to install the Windows Server 2003 operating system on your remote servers. Then you need to test your plan in a lab environment and identify the tasks that you need to perform to implement your plan.

# Choosing New or Existing Hardware

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have older server hardware, it might not support some of the enhanced remote management capabilities provided by newer hardware. For many computers, however, you can purchase hardware components that provide enhanced out-of-band features. If you plan to purchase new servers, you can purchase ones that have been designed for compatibility with Emergency Management Services.

As you compare your existing server hardware with your remote management needs, choose one of the following approaches:

- Use existing servers only
- Use existing servers and purchase new peripherals
- Purchase new servers

These three choices are not mutually exclusive; you can decide the most cost-effective solution for each situation on a case-by-case basis. If you want to eventually replace existing servers with new ones but do not want to incur the entire cost at once, you can also plan to prestage new servers, and gradually deploy them over time.

# Using Existing Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You might be able to use your existing servers, even if they include legacy hardware, but they might provide more limited out-of-band options. If you use existing servers, out-of-band management functionality is limited to Emergency Management Services functionality, which means that you can perform out-of-band management tasks remotely only when the operating system is functioning. You might decide on this option if your servers can be easily and quickly accessed for local management when the operating system is not responding.

If your server was recently manufactured, you might be able to upgrade the system firmware. Firmware that supports console redirection lets you manage remotely before the operating system starts. Upgraded firmware can also make it possible to configure your servers for remote administration. Without firmware that supports console redirection, it is not possible to start your server without a keyboard controller and video adapter. Check with your original equipment manufacturer (OEM) to see if upgraded firmware is available.

If you decide to use existing servers, advantages include the following:

- You save on the cost of purchasing new servers.
- You can gain experience with the technology before investing in new hardware.
- You might be able to upgrade the operating system instead of performing a new installation, and you might not need to reinstall all the applications.

For more information about considerations for upgrading computers, see "Choosing an Automated Installation Method" in *Automating and Customizing Installations* in this kit.

Disadvantages of this approach include the following:

- You have only limited out-of-band functionality: with only Emergency Management Services, you can manage remotely only when the loader or kernel is at least partially running.
- You might experience compatibility issues with older hardware. Check the Windows Server Catalog to verify that all installed devices are listed. To find the Windows Server Catalog, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Using Existing Servers with New Peripherals

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You might decide to use your existing servers but purchase out-of-band peripherals to provide enhanced out-of-band functionality. For example, you can purchase an intelligent UPS to provide remote power options or a terminal concentrator to provide access to multiple servers from a network connection. Devices such as a PCI adapter with a service processor are also available.

Advantages of this approach include the following:

- You save on the cost of new hardware because new peripherals cost less than new servers.
- You might be able to upgrade to Windows Server 2003 instead of performing a new installation.
- You can gain experience with the new technology as you phase in new servers and components.
- You can obtain increased out-of-band functionality.

Disadvantages of this approach include the following:

- You incur the cost of new hardware components.
- If your OEM does not provide out-of-band support, your out-of-band functionality is limited.

If your OEM does not provide out-of-band-related hardware upgrades, or does not support use of such devices, then your options are limited to using existing servers or purchasing new ones. If you need out-of-band management support when the operating system is not functioning, you need to purchase new hardware. Newer systems that are compliant with Emergency Management Services typically have some or all of the out-of-band management features described in this chapter.

# Using New Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you plan to purchase new servers as part of an organization-wide Windows Server 2003 deployment and you plan to implement remote out-of-band management, look for features that enhance Emergency Management Services when you purchase the new computers. Newer computers that are compliant with Emergency Management Services typically include firmware that provides console redirection and might include service processors.

Advantages of purchasing new servers include the following:

- You can simplify the hardware configuration if the servers have components such as service processors or intelligent power switches built in.
- You are much less likely to encounter component incompatibilities because the Emergency Management Systems–compliant computers packaged by OEMs are sold as a single entity and tested together for compatibility.
- Servers that are compliant with Emergency Management Services might have features such as console redirection, a service processor, and an intelligent power switch built in, providing immediate access to enhanced out-of-band features.
- Servers that are compliant with Emergency Management Services typically are compatible with in-band Windows Server 2003 remote management technologies, including RIS.
- Remotely administered server configurations are an option.

The primary disadvantage of purchasing new servers is the cost of the new hardware.

# Prestaging Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you prestage servers you gradually replace existing servers in your organization with new ones preconfigured and tested in your lab. With this approach, you can use existing servers for your initial Windows Server 2003 deployment. Then, as servers reach approximately 90 percent usage, a technician can go onsite to replace the existing servers with new ones that are built to support out-of-band management.

Advantages of this approach include the following:

- You can spread out the cost of new hardware over time.
- You can gain experience with the technology before investing in new hardware and refine your buying criteria as you phase out old equipment and replace it with new servers.
- You gradually get increased out-of-band functionality as servers are replaced.
- You might be able to upgrade to Windows Server 2003 during the initial deployment, instead of performing new installations and reinstalling applications.

The primary disadvantage of prestaging is that you have the possibility of interim compatibility issues.

# Selecting the Installation Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your decision to use existing servers or to purchase new ones might influence the method you choose to install Windows Server 2003. If you decide to use existing servers, you might be able to upgrade the operating system.

There are many issues in addition to upgrade or new installation to consider when you choose an installation method. For a complete discussion of the considerations when deciding on a method, see "Choosing an Automated Installation Method" in *Automating and Customizing Installations* of this kit. This section covers only the installation issues involved with Emergency Management Services and remotely administered servers.

# Selecting RIS for Installation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

On 32-bit and 64-bit systems, you can use Remote Installation Services (RIS) to start a network-based Windows Server 2003 installation on a computer without an operating system. To install the Windows Server 2003 operating system by using RIS, you must have a network adapter that is Pre-Boot eXecution Environment (PXE) capable. A PXE-capable network adapter lets you boot the server from the network.

Special versions of Startrom.com support Emergency Management Services console redirection and let you perform an unattended installation on a server from a remote management computer. These versions of Startrom.com let you perform a remote unattended installation even if the firmware on the server does not provide console redirection. Without the Emergency Management Services version of Startrom.com, you cannot respond to the **Press F12 for network boot** prompt unless the server has firmware that provides console redirection. If you have firmware console redirection, however, it is preferable to rely on it instead of Emergency Management Services console redirection and to use the standard Startrom.com file, because the standard version supports a wider variety of baud rates. The special versions of Startrom.com support only a baud rate of 9600.

**Note**

- If you use one of the special versions of Startrom.com, you must also configure the management computer to use a baud rate of 9600 baud in order to read the redirected output.

For more information about using RIS with Emergency Management Services, see the *Server Management Guide* of the *Windows Server 2003 Technical Reference* on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=4681). For more information about RIS in general, see "Designing RIS Installations" in the *Automating and Customizing Installations* topic.

# Selecting Network-Based Unattended Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you want to fully automate Setup and enable Emergency Management Services, you can use a network-based unattended installation. If your firmware supports the Serial Port Console Redirection (SPCR) table, Setup automatically detects the out-of-band port parameters and enables Emergency Management Services. If your firmware does not support the SPCR table or you want to fully automate Setup, you can configure Setup by using the Unattend.txt file. You can also use the Unattend.txt file to disable Emergency Management Services if your firmware supports the SPCR table, but you do not want to enable Emergency Management Services.

When Emergency Management Services is enabled, you can use Emergency Management Services console redirection to remotely monitor text-mode setup. Although you cannot use either firmware or Emergency Management Services console redirection to redirect GUI-mode setup graphics such as windows and cursors, you can avoid pauses for user input by specifying parameters in your Unattend.txt file. For remotely administered computers, you need to use an Unattend.txt file or run Setup by using a Remote Desktop for Administration connection.

If your Unattend.txt file does not contain the information necessary for GUI-mode setup to proceed — such as your accepting the End User License Agreement (EULA), entering a 25-digit product key, and providing an administrator password — you will be prompted through the serial port to provide that input. If the server has a local monitor, a dialog will show that the system is waiting for input through the serial port. You can optionally use the local keyboard to cancel this. If more information is missing from your answer file, a default answer is created so that you are not prompted again. For example, if the computername parameter is missing, the random computer name, "computername=*" is automatically generated. You can then change this information after GUI-mode setup is complete.

A sample Unattend.txt file is on the Windows Server 2003 operating system CD. You can use default settings or create customized installations by modifying or adding parameters. For more details about Emergency Management Server parameters for the Unattend.txt file, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit). For more information about network-based unattended installations, see "Designing Unattended Installations" in the *Automating and Customizing Installations* book of this kit.

# Selecting Unattended CD-Based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your system supports firmware console redirection, you can enable Emergency Management Services as part of a CD-based installation. If your firmware supports the SPCR table, Setup automatically detects the out-of-band port parameters and enables Emergency Management Services. You can then use Windows Emergency Management Services console redirection to remotely monitor text-mode setup, but you cannot use console redirection to monitor GUI-mode setup. To fully automate a CD-based installation that includes customized parameters, you need to configure an unattended setup by using the Winnt.sif file.

If your firmware supports the SPCR table, console redirection is enabled, and you do not plan to use unattended files during setup, you can use Express Setup. When you boot a remotely administered server from CD and do not use the Winnt.sif file, Setup automatically provides default parameters for the computer configuration so that the GUI-mode phase can complete without displaying any prompts. After Setup is complete, you can use SAC to set parameters such as IP address or view the computer name, if the IP address was assigned using DHCP. Then you can use a Remote Desktop for Administration connection to complete your custom configuration.

To begin a CD-based Windows Server 2003 setup, you must enable firmware console redirection to remotely view and respond to the **Press any key to boot from CD-ROM** prompt that appears when starting the computer by using the Windows Server 2003 operating system CD.

**Tip**

- CD-based setups are most useful for local installations during testing.

For more details about using CD-based setups and configuring the Winnt.sif file for Emergency Management Services, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Selecting Unattended Image-Based Installations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can also use an Unattend.txt file to fully automate Setup and enable Emergency Management Services as part of an image-based installation on remotely administered servers. You can distribute the CD image you create to a remote server, power on that server, and proceed through an unattended GUI-mode setup just as you would proceed through an unattended, network-based installation. You can also use Express Setup if your firmware supports the SPCR table, console redirection is enabled, and you do not plan to use unattended files during setup.

# Enable Emergency Management Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Emergency Management Services functionality is built into Windows Server 2003 and is installed, but not enabled, when the operating system is installed. You can enable Emergency Management Services at any of the following times:

- During a new Windows Server 2003 installation, including RIS-based setups.
- During a Windows Server 2003 upgrade (x86-based systems only).
- After you complete a Windows Server 2003 installation.

For information about how to enable Emergency Management Services during or after installation, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Preparing for Deployment of Your Remote Management Solution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have followed the planning steps in this chapter to identify your organization's remote management requirements and to design a solution that meets those requirements, you need to test your solution in a lab environment that emulates your production environment as closely as possible. Because out-of-band management and support for remotely administered servers is new to the Windows environment, it is particularly important to test your remotely administered servers and out-of-band configurations. Then develop a list of tasks you need to perform before you can implement your plan. Include tasks for preparing your network, hardware and software configurations, and security.

**Network connectivity**
Identify the steps you need to take to prepare your network infrastructure to ensure connectivity and stable access. Include tasks to accomplish the following:

- Configure your network for increased bandwidth, or establish a secondary network for remote management, if necessary.
- Configure Dial-up Networking settings, if you plan to support dial-up remote management.
- Configure firewall settings, if you plan to remotely manage across a firewall.
- Configure IP packet filter settings, if you plan to remotely manage across a router that filters packets.

**Server resources**
Identify the steps for assessing computer resources and for upgrading them, if necessary, such as:

- Upgrade server and management computer memory and processor, if necessary, to support remote management tools.
- Upgrade or purchase new computers that provide enhanced out-of-band support, if necessary.

**Security**
Identify the steps for ensuring both physical and logical security, for both in-band and out-of-band management, such as:

- Provide secured access to servers.
- Provide physical security for out-of-band serial connections.
- Provide for authentication and encryption protocols.
- Configure Group Policy settings to mitigate vulnerabilities of less secure authentication protocols.
- Configure IPSec policies to provide security if you plan to use tools without encryption or dial-up connections through a VPN.
- Establish a secondary management network or configure secured routers, if necessary.
- Set up security groups and assign membership for administrators.
- Set up shared folders as necessary, and assign permissions.
- Configure Group Policy settings to restrict the types of administrative tasks remote users can perform.

**Configure out-of-band settings**
Identify the steps for configuring computers for out-of-band management:

- Enable firmware console redirection.
- Verify and, if necessary, configure hardware resource settings for serial ports.
- Configure SPCR table settings.
- Configure the service processor, if necessary.
- Select consistent terminal definition settings for firmware, Emergency Management Services, and client terminal software.

For more information about specific settings for Emergency Management Services, see the *Server Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Connect out-of-band hardware components**
Identify the steps for configuring the out-of-band infrastructure:

- Set up serial port null modem connections.
- Connect terminal concentrators, modems, service processors, and intelligent UPSs.

For more information about setting up null modem connections in an Emergency Management Services installation, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

**Install and enable Emergency Management Services**
Identify the steps for installing Windows Server 2003 and enabling Emergency Management Services:

- Build and test files and directories for unattended network, unattended CD, or RIS installations of Windows Server 2003.
- Enable Emergency Management Services during or after Windows Server 2003 installation.

For more information about enabling Emergency Management Services, see the *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit).

# Additional Resources for Remote Server Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information and tools related to this chapter.

**Related Information**

- The *Storage Technologies Collection* of the *Windows Server 2003 Technical Reference* (or see the Storage Technologies Collection on the Web at http://www.microsoft.com/reskit) for more information about remote management, remote management tools, and configuring and using Emergency Management Services.
- "Deploying a Managed Software Environment" in *Designing a Managed Environment* of this kit for information about software distribution tools.
- "Designing a Group Policy Infrastructure" in *Designing a Managed Environment* of this kit for more information about Group Policy.
- The *Windows Security Collection* of the *Windows Server 2003 Technical Reference* (or see the Windows Security Collection on the Web at http://www.microsoft.com/reskit) for more information about Group Policy, and also for more information about groups, user rights, permissions, and authorization and access control.
- The *Internetworking Collection* of the *Windows Server 2003 Technical Reference* (or see the Internetworking Collection on the Web at http://www.microsoft.com/reskit) for more information about VPN and about using IPSec with VPN.
- "Choosing an Automated Installation Method" in *Automating and Customizing Installations* in this kit for more information about choosing an installation method for the Windows Server 2003 operating system.
- "Designing RIS Installations" in *Automating and Customizing Installations* of this kit for more information about Remote Installation Services (RIS).
- "Designing Unattended Installations" in *Automating and Customizing Installations* of this kit for more information about network-based unattended installations.
- The WMI Software Development Kit (SDK) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about developing applications using WMI.

**Related Tools**

- Windows Server 2003 Administration Tools Pack

  This set of tools includes several of the most common tools for remotely managing servers from a Windows XP Professional–based computer. Install the tools pack by running C:\windows\system32\adminpak\adminpak.msi.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Special Administration Console (SAC) and SAC commands" in Help and Support Center for Windows Server 2003 for information about SAC commands.
- " !Special Administration Console (!SAC) and !SAC commands" in Help and Support Center for Windows Server 2003 for information about !SAC commands.

- "Windows Server 2003 Administration Tools Pack Overview" in Help and Support Center for Windows Server 2003 for more information about the Administration Tools Pack. This topic is installed when you install the tools.

- "Configure client callback options" in Help and Support Center for Windows Server 2003 for information about configuring modems for callback.

- "VPN remote access for employees" in Help and Support Center for Windows Server 2003 for information about network settings for VPN connections.

- "Best practices for selecting and configuring hardware" in Help and Support Center for Windows Server 2003 for information about selecting hardware components to use with Emergency Management Services.

- "Internet Protocol Security (IPSec)" in Help and Support Center for Windows Server 2003 for information about IPSec policies.

- "Best practices for permissions and user rights" in Help and Support Center for Windows Server 2003 for information about assigning permissions and user rights for remote management.

**Related Job Aids**

For a document to assist you in setting up remotely administered servers with a terminal concentrator, see "Headless Server Quick Start" (SDCEMS_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Headless Server Quick Start" on the Web at http://www.microsoft.com/reskit).

# Planning for High Availability and Scalability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Business-critical applications such as corporate databases and e-mail often need to reside on systems and network structures that are designed for high availability. The same is true for retail Web sites and other Web-based businesses. Knowing about high availability concepts and practices can help you to maximize the availability (extremely low downtime) and scalability (the ability to grow as demand increases) of your server systems. Using sound IT practices and fault-tolerant hardware solutions in your deployment can increase availability and scalability. Additionally, the Microsoft® Windows® Server 2003 operating system offers two clustering technologies — server clusters and Network Load Balancing — that provide both the reliability and scalability that most enterprises need.

**In This Chapter**

Overview of Planning for High Availability and Scalability

Defining Availability and Scalability Goals

Using IT Procedures to Increase Availability and Scalability

Evaluating Windows Clustering

Using Clusters to Increase Availability and Scalability

Additional Resources for Planning Server Deployments

**Related Information**

- For more information about planning for server clusters, see "Designing and Deploying Server Clusters" in this book.
- For more information about Network Load Balancing, see "Designing Network Load Balancing" in this book.
- For more information about fault-tolerant storage solutions, see "Planning for Storage" in this book.
- For information about deploying Emergency Management Services for managing remote servers, see "Planning for Remote Server Management" in this book.

# Overview of Planning for High Availability and Scalability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A highly available system reliably provides an acceptable level of service with minimal downtime. Downtime penalizes businesses, which can experience reduced productivity, lost sales, and lost faith from clients, partners, and customers.

By implementing recommended IT practices, you can increase the availability of key services, applications, and servers. These practices also help you minimize both planned downtime, such as for maintenance or service pack installations, and unplanned downtime, such as downtime caused by a server failure.

Additionally, the clustering technologies available in Windows Server 2003 can provide added levels of availability and scalability. A cluster is a group of independent computers that work together to provide a common set of services. If a cluster node (a computer in a cluster) fails, other nodes in the cluster assume the functions of the failed node. Clusters enhance the availabilityof critical applications and services. Clusters also increase the scalabilityof your deployment by allowing you to scale your solution over time to support additional clients or increased client demands. In most deployments, server clustersprovide availability, and Network Load Balancing clustersincrease scalability.

If you need the highest levels of availability (and if your budget can accommodate the expense), use both recommended IT practices and clustering for a truly comprehensive high availability deployment that is also scalable. If you determine that your organization needs to increase availability and scalability, this chapter can help you create an effective plan for implementing a high availability and scalability solution that meets your organization's current and future needs.

**Note**

- Server clusters are supported by the Microsoft® Windows® Server 2003, Enterprise Edition, and Windows® Server 2003, Datacenter Edition, operating systems. Network Load Balancing is supported by all Windows Server 2003 operating systems, including Windows® Server 2003, Standard Edition, and Windows® Server 2003, Web Edition.

# High Availability and Scalability Planning Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 6.1 illustrates the high availability design process and defines the steps you can take to ensure a deployment that meets your requirements for high availability.

**Figure 6.1 Planning for High Availability and Scalability**



# Basic High Availability and Scalability Concepts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To begin designing a Windows Server 2003 deployment for maximum availability and scalability, familiarize yourself with these fundamental high availability and scalability concepts.

**Basic High Availability Concepts**
An integral part of building large-scale, mission-critical systems that your business and your users can rely on is to ensure that no single point of failure can render a server or network unavailable. There are several types of failures you must plan against to keep your system highly available.

**Storage failures**
There are many ways to protect against failures of individual storage devices using techniques such as Redundant Array of Independent Disks (RAID). Storage vendors provide hardware solutions that support many different types of hardware redundancy for storage devices, allowing devices as well as individual components in the storage controller itself to be exchanged without losing access to the data. Software solutions also provide similar capabilities.

**Network failures**
There are many components to a computer network, and there are many typical network topologies that provide highly available connectivity. All types of networks need to be considered, including client access networks and management networks. In storage area networks (SANs), failures might include the storage fabrics that link the computers to the storage units. For more information about SANs, see "Planning for Storage" and "Designing and Deploying Server Clusters" in this book.

**Computer failures**
Many enterprise-level server platforms provide redundancy inside the computer itself, such as through redundant power supplies and fans. Vendors also allow components such as peripheral component interconnect (PCI) cards and memory to be swapped in and out without removing the computer from service. In cases where a computer fails or needs to be taken out of service for routine maintenance or upgrades, clustering provides redundancy to enable applications or services to continue. This redundancy happens automatically in clustering, either through failover of the application (transferring client requests from one computer to another) or by having multiple instances of the same application available for client requests.

**Site failures**
In extreme cases, a complete site can fail due to a total power loss, a natural disaster, or other unusual occurrences. More and more businesses are recognizing the value of deploying mission-critical solutions across multiple geographically dispersed sites. For disaster tolerance, a data center's hardware, applications, and data can be duplicated at one or more geographically remote sites. If one site fails, the other sites continue offering service until the failed site is repaired. Sites can be *active-active*, where all sites carry some of the load, or *active-passive*, where one or more sites are on standby.

You can prevent most of these failures by using the following methods:

- **Proven IT practices.** IT strategies can help your organization avoid some or all of the above failures. IT practices take on added importance when a clustering solution is not applicable to, or even possible in, your particular deployment. Whether or not you choose to deploy a clustering solution, all Windows deployments should at a minimum follow the guidelines and reference the resources listed in this chapter for fault-tolerant hardware solutions.

- **Clustering.** This chapter introduces Windows Server 2003 clustering technologies and provides an overview of how they work. Different kinds of clusters can be used together to provide a true end-to-end high availability and scalability solution. For more information about end-to-end solutions, see "Using Clusters to Increase Availability and Scalability" later in this chapter. For complete information about designing clusters, see "Designing and Deploying Server Clusters" and "Designing Network Load Balancing" in this book.

**Basic Scalability Concepts**
In general deployments, scalability is the measure of how well a service or application can grow to meet increasing performance demands. When applied to clustering, scalability is the ability to incrementally add systems to an existing cluster when the overall load of the cluster exceeds the cluster's capabilities.

**Scaling up**
*Scaling up* involves increasing system resources (such as processors, memory, disks, and network adapters) to your existing

hardware or replacing existing hardware with greater system resources. Scaling up is appropriate when you want to improve client response time, such as on a Network Load Balancing cluster. If the required number of users are not properly supported, adding random access memory (RAM) or central processing units (CPUs) to the server is one way to meet the demand.

Windows Server 2003 supports single or multiple CPUs that conform to the symmetric multiprocessing (SMP) standard. Using SMP, the operating system can run threads on any available processor, which makes it possible for applications to use multiple processors when additional processing power is required to increase the capability of a system. For more information about scaling up, see "Using Clusters to Increase Availability and Scalability" later in this chapter.

**Scaling out**

*Scaling out* involves adding servers to meet demand. In a server cluster, this means adding nodes to the cluster. Scaling out is also appropriate when you want to improve client response time with your servers, and when you have the hardware budget to purchase additional servers as needed. For more information about scaling out, see "Using Clusters to Increase Availability and Scalability" later in this chapter.

**Testing and Pilot Deployments**

Before you deploy any new solution, whether it is a fault-tolerant hardware or networking component, a software monitoring tool, or a Windows clustering solution, you should thoroughly test the solution before deploying it. After testing in an isolated lab, test the solution in a pilot deployment in which only a few users are affected, and make any necessary adjustments to the design. After you are satisfied with the pilot deployment, perform a full-scale deployment in your production environment. Depending on the number of users you have, you might want to perform your full-scale deployment in stages. After each stage, verify that your system can accommodate the increased processing load from the additional users before deploying the next group of users.

Clusters in particular need to be tested and then deployed in a pilot environment before you deploy them in a production environment. For more information about testing clusters, see "Designing and Deploying Server Clusters" and "Deploying Network Load Balancing" in this book.

For complete information about setting up a test environment for your Windows Server 2003 deployment and implementing a pilot deployment see "Designing a Test Environment" and "Designing a Pilot Project" in *Planning, Testing, and Piloting Deployment Projects* in this kit.

# Defining Availability and Scalability Goals

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Defining availability and scalability goals is the first step toward ensuring that your efforts are focused on the elements of the system that matter the most to your organization. Figure 6.2 shows how to quantify and determine your availability goals.

**Figure 6.2 Defining Availability and Scalability Goals**




Availability goals allow you to accomplish the following tasks:

- Design, operate, and evaluate your systems in relationship to a consistent set of priorities, and place new requests or problems in context.
- Keep efforts focused where they are needed. Without clear goals, efforts can become uncoordinated, or resources can be spread so thin that none of the organization's most important needs are met.
- Limit costs. You can direct expenditures toward the areas where they make the most difference.
- Recognize when tradeoffs must be made, and make them in appropriate ways.
- Clarify areas where one set of goals might conflict with another, and avoid making plans that require a system to handle two conflicting goals simultaneously.
- Provide a way for operators and support staff to prioritize unexpected problems when they arise by referring to the availability goals for that component or service.

# Quantifying Availability and Scalability for Your Organization

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your goal in quantifying availability is to compare the costs of your current IT environment — including the actual costs of outages — and the cost of implementing high availability solutions. These solutions include training costs for your staff as well as facilities costs, such as costs for new hardware. After you have calculated the costs, IT managers can use these numbers to make business decisions, not just technical decisions, about your high availability solution. For information about monitoring tools that can help you measure the availability of your services and systems, see "Implementing Software Monitoring and Error-Detection Tools" later in this chapter.

Scalability is more difficult to quantify because it is based on future needs and therefore requires a certain amount of estimation and prediction. Remember, though, that scalability is tied to availability because if your system cannot grow to meet increased demand, certain services will become less available to your users.

#### Determining Availability Requirements

Availability can be expressed numerically as the percentage of the time that a service is available for use. The exact level of availability must be determined in the context of the service and the organization that uses the service. Table 6.1 displays common availability levels that many organizations try to achieve. The following formula is used to calculate these levels:

**Percentage of availability = (total elapsed time – sum of downtime)/total elapsed time**

**Table 6.1 Availability Measurements and Yearly Downtime**

| Availability | Yearly Downtime |
|---|---|
| 99.999% | 5 minutes |
| 99.99% | 53 minutes |
| 99.9% | 8 hours, 45 minutes |

Availability requirements can vary depending on the server role. Your users can probably continue to work if a print server is down, for example, but if a server hosting a mission-critical database fails, your business might feel the effects immediately.

#### Determining Reliability Requirements

Reliability is related to availability, and it is generally measured by computing the time between failures. Mean time between failures (MTBF) is calculated by using the following equation:

**MTBF = (total elapsed time – sum of downtime)/number of failures**

A related measurement is mean time to repair (MTTR), which is the average amount of time that it takes to bring an IT service or component back to full functionality after a failure.

A system is more reliable if it is *fault tolerant*. Fault tolerance is the ability of a system to continue functioning when part of the system fails. This is achieved by designing the system with a high degree of hardware redundancy. If any single component fails, the redundant component takes its place with no appreciable downtime. For more information about fault-tolerant components, see "Planning and Designing Fault-Tolerant Hardware Solutions" later in this chapter.

#### Determining Scalability Requirements

You need to consider scalability now to provide your organization a certain amount of flexibility in the future. If you believe your hardware budget will be sufficient, you can plan to purchase hardware at regular intervals to add to your existing deployment. The amount of hardware you purchase depends on the exact increase in demand. If you have budget limitations, purchase servers that you can scale up later by adding RAM or CPUs to meet a rise in users or client requests.

Looking at past growth can help you determine how demand on your IT system might grow. However, because business technology is becoming increasingly complex, and reliance on that technology is growing every year, you must consider other factors as well. If you anticipate growth, realize that some aspects of your deployment may grow at different rates. You might need many more Web servers than print servers, for example, over a certain period of time. For some types of servers, it might be sufficient to add CPU power when network traffic increases, while in other cases, such as with a Network Load Balancing cluster, the most practical scaling solution might be to add more servers.

Recreate your Windows deployment as accurately as possible in a test environment and, either manually or through a

simulation program, put as much workload as possible on different areas of your deployment. Observing your system under such circumstances can help you formulate scaling priorities and anticipate where you might need to scale first.

After your system is deployed, software-monitoring tools can alert you when certain components of your system are near or at capacity. Use these tools to monitor performance levels and system capacity so that you know when a scaling solution is needed. For more information about monitoring performance levels, see "Implementing Software Monitoring and Error-Detection Tools" later in this chapter.

# Analyzing Risk

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When planning a highly available Windows Server 2003 environment, consider all available alternatives and measure the risk of failure for each alternative. Begin with your current organization and then implement design changes that increase reliability to varying degrees. Evaluate the costs of each alternative against its risk factors and the impact of downtime to your organization.

Often, achieving a certain level of availability can be relatively inexpensive, but to go from 98 percent availability to 99 percent, for example, or from 99.9 percent availability to 99.99 percent, can be very costly. This is because bringing your organization to the next level of availabity might entail a combination of new or costly hardware solutions, additional staff, and support staff for non-peak hours. As you determine how important it is to maintain productivity in your IT environment, consider whether those added days, hours, and minutes of availability are worth the price.

Every operations center needs a risk management plan. When assessing risks in your proposed Windows Server 2003 deployment, remember that network and server failures can cause considerable loss to businesses. After you evaluate risks versus costs, and after you design and deploy your system, your IT staff needs sound guidelines and plans of action in case a failure in the system does occur.

For more information about designing a risk management plan for your Windows Server 2003 deployment, see the Microsoft Operations Framework (MOF) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. The MOF Web page provides complete information about creating and maintaining a flexible risk management plan, with emphasis on change management and physical environment management, in addition to staffing and team role recommendations.

# Developing Availability and Scalability Goals

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Begin establishing goals by reviewing information that is readily available within your organization. Existing Service Level Agreements, for example, define the availability goals for specific IT services or systems. Gather information from those individuals and groups who are most directly affected, such as the users or departments that depend on the services and the people who make decisions about IT staffing.

The following questions provide a starting point for developing a list of availability goals. These goals, and the factors that influence them, vary from organization to organization. By identifying the goals appropriate to your situation, you can clarify your priorities as you work to increase system availability and reliability.

**Organization's Central Purposes**
These fundamental questions will help you prioritize the applications and services that are most important to your organization and the extent to which you rely on your IT infrastructure for certain key tasks.

- What are the organization's central purposes?
- What must the organization accomplish to survive and flourish?

**Details on Record That Help Define Availability Requirements**
The questions in this section can help you quantify your availability needs, which is the first step in addressing those needs.

- If your organization has attempted to evaluate the need for high availability in the past, do you have existing documents that already outline availability goals?
- Do you have current or previous Service Level Agreements, Operating Level Agreements, or similar agreements that define service levels?
- Have you defined acceptable and unacceptable service levels?
- Do you have data about the cost of outages or the effect of service delays or outages (for example, information about the cost of an outage at 9 A.M. versus the cost of an outage at 9 P.M.)?
- Do you have any data from groups that practice incident management, problem management, availability management, or similar disciplines?

**Users of IT Services**
It is important to define the needs of your users to provide them with the availability they need to do their work. There is often a tradeoff between providing high availability and paying the cost of hardware, training, and support. Categorizing your users can make these kinds of business decisions easier.

- Who are the end users? What groups or categories do they fall into? What expertise levels do they have?
- How important is each user group or category to the organization's central goals?

**Requirements and Requests of End Users**
These questions help pinpoint the needs of your users. You can more easily customize your high availability solutions and anticipate scalability issues if you know exactly what your users need.

- Among the tasks that users commonly perform, which are the most important to the organization's central purposes?
- When end users try to accomplish the most important tasks, what do they expect to see on their screens (or access through some other device)? Described another way, what data (or other resources) do users need to access, and what applications or services do they need when working with that data?

- For the users and tasks most important to the organization, what defines a satisfactory level of service?

**Requirements for User Accounts, Networks, or Similar Types of Infrastructure**

It is important to know about supporting services — even services that you do not control — when evaluating availability needs and defining availability goals. Your system will be only as fault tolerant as the systems that support it.

- What types of network infrastructure and directory services are required so that users can accomplish the tasks that you have identified as requirements for end users? In other words, what types of behind-the-scenes services do users require?

- For these behind-the-scenes services, what defines the difference between satisfactory and unsatisfactory results for the organization?

**Time Requirements and Variations**

Keeping support staff on-site to maintain a system can be expensive. Costs can be minimized if support personnel are on-site only during critical periods. Similarly, knowing when workload is highest can help you anticipate when availability is most important, and possibly when a failure is likely to occur.

- Are services needed on a 24-hours-a-day, 7-days-a-week basis, or on some other schedule (such as 9 A.M. to 5 P.M. on weekdays)?

- What are the normal variations in load over time?

- What increments of downtime are significant (for example, five seconds, five minutes, an hour) during peak and nonpeak hours?

# Using IT Procedures to Increase Availability and Scalability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following sections introduce best practices for optimizing your Windows Server 2003 deployment for high availability and scalability. A well-planned deployment strategy can increase system availability and scalability while reducing the support costs and failure recovery times of a system. Figure 6.3 displays the process for deploying your servers and network infrastructure in a fault-tolerant manner that also provides manageability.

**Figure 6.3 Using IT Procedures to Increase Availability and Scalability**




To aid in this planning, Microsoft recommends the Microsoft Operations Framework (MOF). MOF is a flexible, open-ended set of guidelines and concepts that you can adapt to your specific operations needs. Adopting MOF practices provides greater organization and contributes to regular communication between your IT department, your end users, and other departments in your company that might be affected. For complete information about MOF and how to implement it in your organization, see the Microsoft Operations Framework (MOF) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Planning and Designing Hardware for High Availability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

*Fault tolerance* means there can be no single point of failure that can cause system failures. The following sections contain highlights of IT practices for ensuring reliable hardware performance and a fault-tolerant IT infrastructure. Included is a summary of ways to safeguard your servers for optimal performance. For detailed information about IT fault-tolerant hardware strategies and highly available system designs, see the Microsoft Solutions Framework link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Planning and Designing Fault-Tolerant Hardware Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An effective hardware strategy can improve the availability of a system. These strategies can range from adopting commonsense practices to using expensive fault-tolerant equipment.

**Using Standardized Hardware**
To ensure full compatibility with Windows operating systems, choose hardware from the Windows Server Catalog only. For more information, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

When selecting your hardware from the Windows Server Catalog, adopt one standard for hardware and standardize it as much as possible. To do this, pick one type of computer and use the same kinds of components, such as network cards, disk controllers, and graphics cards, on all your computers. Use this computer type for all applications, even if it is more than you need for some applications. The only parameters that you should modify are the amount of memory, number of CPUs, and the hard disk configurations.

Standardizing hardware has the following advantages:

- Having only one platform reduces the amount of testing needed.
- When testing driver updates or application-software updates, only one test is needed before deploying to all your computers.
- With only one system type, fewer spare parts are required.
- Because only one type of system must be supported, support personnel require less training.

For help choosing standardized hardware for your file and print servers, see "Designing and Deploying File Servers" and "Designing and Deploying Print Servers" in this book.

**Using Spares and Standby Servers**
This chapter discusses clustering as a means of providing high availability for your applications and services to your end users. However, there are two clustering alternatives that provide flexibility or redundancy in your hardware design: spares and standby systems.

**Spares**
Keep spare parts on-site, and include spares in any hardware budget. One of the advantages of using a standard configuration is the reduced number of spares that must be kept on-site. If all of the hard drives are of the same type and manufacturer, for example, you can keep fewer drives in stock as spares. This reduces the cost and complexity associated with providing spares.

The number of spares that you need to keep on hand varies according to the configuration and failure conditions that users and operations personnel can tolerate. Another concern is availability of replacement parts. Some parts, such as memory and CPU, are easy to find years later. Other parts, like hard drives, are often difficult to locate after only a few years. For parts that may be hard to find, and where exact matches must be used, plan to buy spares when you buy the equipment. Consider using service companies or contracts with a vendor to delegate the responsibility, or consider keeping one or two of each of the critical components in a central location.

**Standby Systems**
Consider the possibility of maintaining an entire standby system, possibly even a hot standby to which data is replicated automatically. For file servers, for example, the Windows Server 2003 Distributed File System (DFS) allows you to logically group folders located on different servers by transparently connecting them to one or more hierarchical namespaces. When DFS is combined with File Replication service (FRS), clients can access data even if one of the file servers goes down, because the other servers have identical content. DFS and FRS are discussed in detail in "Designing and Deploying File Servers" in this book.

If the costs of downtime are very high and clustering is not a viable option, you can use standby systems to decrease recovery times. Using standby systems can also be important if failure of the computer can result in high costs, such as lost profits from server downtime or penalties from a Service Level Agreement violation.

A standby system can quickly replace a failed system or, in some cases, act as a source of spare parts. Also, if a system has a catastrophic failure that does not involve the hard drives, it might be possible to move the drives from the failed system to a working system (possibly in combination with using backup media) to restore operations relatively quickly. This scenario does not happen often, but it does happen, in particular with CPU or motherboard component failures. (Note that this transfer of data after a failure is performed automatically in a server cluster.)

One advantage to using standby equipment to recover from an outage is that the failed unit is available for careful after-the-fact diagnosis to determine the cause of the failure. Getting to the root cause of the failure is extremely important in preventing repeated failures.

Standby equipment should be certified and running on a 24-hours-a-day, 7-days-a-week basis, just like the production equipment. If you do not keep the standby equipment operational, you cannot be sure it will be available when you need it.

**Using Fault-Tolerant Components**

Using fault-tolerant technology improves both availability and performance. The following sections describe some basic fault-tolerant considerations in two key areas of your deployment: storage and network components. In both cases you should also consult hardware vendors for details specific to each product, especially if you are considering deploying server clusters. For more information about storage options and strategies for server clusters, see "Designing and Deploying Server Clusters" in this book.

**Storage Strategies**

When planning how to store your data, consider the following points:

- The type and quantity of information that must be stored. For example, will a particular computer be used to store a large database needing frequent reads and writes?

- The cost of the equipment. It does not make sense to spend more money on the storage system than you expect to recover in saved time and data if a failure occurs.

- Specific needs for protecting data or making data constantly available. Do you need to prevent data loss, or do you need to make data constantly available? Or are both necessary? For preventing data loss, a RAID arrangement is recommended. For high availability of an application or service, consider multiple disk controllers, a RAID array, or a Windows clustering solution. (Clustering is discussed later in this chapter.)

- A good backup and recovery plan is essential. Downtime is inevitable, but a sound and proven backup and recovery plan can minimize the time it takes to restore services to your users. For more information, see "Backing up and recovering data" in Help and Support Center for Windows Server 2003.

- Physical memory copying, or memory *mirroring*, provides fault tolerance through memory replication. Memory-mirroring techniques include having two sets of RAM in one computer, each a mirror of the other, or mirroring the entire system state, which includes RAM, CPU, adapter, and bus states. Memory mirroring must be developed and implemented in conjunction with the original equipment manufacturer (OEM).

For more information about storage strategies, see "Planning for Storage" in this book.

**Network Components**

The network adapter is a potential single point of failure. Fortunately, the network adapter is, on average, very reliable. However, other components outside the computer can fail, causing the same effect that you would experience with the loss of the network adapter. These include the network cable to the computer; the switch or hub; the router; and the Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and Windows Internet Name Service (WINS) systems. Any one of these components can fail and cause the failure of one or more servers and, potentially, all the servers.

You can contend with such failures through redundancy in your network design. Many components lend themselves to backup or load-sharing strategies. The following list describes redundancy strategies for the network hardware (hub or switch, network adapter, and wiring), the routers, and DNS or WINS.

**Note**

- The following is a list of general planning considerations for all networks. For more information about network considerations for specific cluster types, see "Designing Network Load Balancing" and "Designing and Deploying Server Clusters" in this book.

**Network hardware** Although hubs, switches, network adapters, and wiring are very reliable, if a service must be guaranteed, it is still important to use redundancy for these components. Consult with the vendors who provide your network hardware and support for recommendations on how to build redundancy into your network. For more information about building redundancy into your network, see "Designing a TCP/IP Network" in *Deploying Network Services* of this kit.

**Routers** Routers do not frequently fail, but when they do, entire computer centers can go down. Having redundant routing capability in the computer center is critical. Your router vendor is a recommended source of information about how to protect against router failures.

**DHCP** For the servers on which you must maintain the highest degree of availability, use fixed IP addresses on servers and do not use DHCP. This prevents an outage due to the failure of the DHCP server. This can improve address resolution by DNS servers that do not handle the dynamic address assignment provided by DHCP. For more information about DHCP, see "Deploying DHCP" in *Deploying Network Services* of this kit.

**DNS and WINS** DNS and WINS infrastructure components are easy to replicate. Both were designed to support replication of their name tables and other information. Make sure that when you use multiple DNS and WINS servers, you place them on different network segments. For information about WINS and DNS servers and replication, see "Deploying DNS" and "Deploying WINS" in *Deploying Network Services* of this kit.

For information about replication options for WINS servers, see "Configuring WINS replication" in Help and Support Center for Windows Server 2003. For information about replicating DNS zones, see "DNS zone replication in Active Directory" in Help and Support Center for Windows Server 2003.

For more information about network infrastructure, see the Microsoft Systems Architecture link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Safeguarding the Physical Environment of the Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An important practice for high availability of servers is to maintain high standards for the environment in which the servers must run. The following list contains information to consider if you want to increase the longevity and reliability of your hardware:

- **Temperature and humidity.** Install mission-critical servers in a room set aside for that purpose, where you can carefully control temperature and humidity. Computers perform best at approximately 70 degrees Fahrenheit. In an office, temperature is not normally an issue, but be aware of the effect of a long holiday weekend in the summer with the air conditioning turned off.

- **Dust or contaminants.** Protect servers and other equipment from dust and contaminants where possible, and check for dust periodically. Dust and other contaminants can cause components to short-circuit or overheat, which can cause intermittent failures. Whenever the case of a server is opened for any reason, perform a quick check to determine whether the unit needs cleaning. If so, check all the other units in the area.

- **Power supplies.** Planning for power outages, like any disaster-recovery planning, is best done long before you anticipate outages, and it involves identifying the resources that are most critical to the operation of the company. When possible, provide power from at least two different circuits to the computer room and divide redundant power supplies between the power sources. Ideally, the circuits should originate from two different sources external to the building. Be aware of the maximum amount of power a location can provide. It is possible that a location could have so many servers that there is not sufficient power for any additional servers you might want to install. Consider a backup power supply for use in the event of a power failure in your computer center. It may be necessary to continue providing computer service to other buildings in the area or to areas geographically remote from the computer center. Short outages can be dealt with through uninterruptible power supply (UPS) units. Longer duration outages can be handled using standby generators. Include network equipment, such as routers, when reviewing equipment that requires backup power during an outage.

- **Maintenance of cables.** Prevent physical damage to cables in the computer room by making sure cables are neat and orderly, either with a cable management system or tie wraps. Cables should never be loose in a cabinet, where they can be disconnected by mistake. Make sure all cables are securely attached at both ends where possible, and make sure pull-out, rack-mounted equipment has enough slack in the cables, and that the cables do not bind and are not pinched or scraped. Set up good pathways for redundant sets of cables. If you use multiple sources of power or network communications, try to route the cables into the cabinets from different points. If one cable is severed, the other can continue to function. Do not plug dual power supplies into the same power strip. If possible, use separate power outlets or UPS units (ideally, connected to separate circuits) to avoid a single point of failure.

- **Security of the computer room.** For servers that must maintain high availability, restrict physical access for all but designated individuals. In addition, consider the extent to which you need to restrict physical access to network hardware. The details of how you implement this depend on your physical facilities and your organization's structure and policies. When reviewing the security in place for the computer room, also review your methods for restricting access to remote administration of servers. Make sure that only designated individuals have remote access to your configuration information and your administration tools.

# Implementing Software Monitoring and Error-Detection Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Constant vigilance of your network and applications is essential for high availability. Software-monitoring tools and techniques allow you to determine the health of your system and identify potential trouble spots before an error occurs.

This section assumes that you have selected software that supports the high availability features you require. Not all software supports features such as redundancy or clustering. For an application that requires 99 percent uptime, this might not matter. An application that requires 99.9 percent or greater availability must support such features. Monitoring tools can reveal performance trends and other indications that a potential loss of service is eminent before an error actually occurs. If an error does occur, monitoring tools can provide analytic data that administrators can use to prevent the problem from happening again.

Before deploying software, check the application's hardware requirements and consult the documentation or software vendor to be sure the application supports online backup. When your monitoring tools detect a problem or error in an application, online backup allows you to fix the problem with no disruption of service.

# Planning Long-Term Monitoring for Trend Analysis

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Whenever possible, establish monitoring over the long term so you can carry out trend analysis and capacity planning. The advantages of long-term monitoring include:

- Increasing your ability to predict when expansion of the system is necessary.
- Helping with troubleshooting for problems such as memory leaks or abnormal disk consumption.
- Assisting in identifying strategies for load balancing.

Capacity planning plays an important part in the long-term success of a high availability system. A good capacity plan can limit avoidable failures by anticipating system usage, scalability, and growth requirements. Do not limit capacity planning to items such as disks. Ensure that your plans encompass all parts of a system that could become bottlenecks.

Potential bottlenecks for a server include the CPU, memory, disks, and networking components such as cables, routers, switches, and telecommunications carriers.

**CPU** As a general guideline, provide enough processing power to keep CPU utilization below 75 percent.

**Memory** It is important to understand physical memory and its relation to caching, virtual memory, and disk paging. The goal is to provide enough memory to avoid excessive paging, because paging can never be as efficient as directly transferring information to or from physical memory.

**Disks** A general observation for disks is that utilization above 85 percent tends to cause longer seek times, adding to overall response time. As a general guideline, size your disks, and distribute the load between them, to keep utilization below 85 percent.

**Networking, including cabling, routers, switches, and telecommunications carriers** Create a logical map that includes each type of network component and its location along with servers and user access points. This type of diagram provides an understanding of how information travels back and forth between a given user and a given server. Along with using the diagram, gather and organize as much information as possible about the capacities of your network components and about the network loads generated in your enterprise.

**Domain controller traffic** When monitoring network performance, be sure to consider domain controllertraffic and other services inherent in your networking infrastructure. The number of domain controllers that you deploy and their location in your topology can affect performance. For more information about domain controllers, see "Planning Domain Controller Capacity" in *Designing and Deploying Directory and Security Services* of this kit.

# Choosing Monitoring Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you deploy your software, establish routine and automated monitoring and error detection for your operating system and applications. If you can detect application and system errors immediately after they occur, you have a better chance of responding before a system shutdown. Monitoring can also alert you if scaling is necessary somewhere in your organization. For example, if one or more servers are operating at capacity some or all of the time, you can decide if you need to add more servers or increase the CPUs of existing servers.

You can use the following tools and methods to monitor your programs. For more information about monitoring, consult "Monitoring and status tools" in Help and Support Center for Windows Server 2003.

When choosing software to run in a server cluster or Network Load Balancing cluster, see "Designing Network Load Balancing" and "Designing and Deploying Server Clusters" in this book. For additional information, see "Choosing applications to run on a server cluster" and "Determining which applications to use with Network Load Balancing" in Help and Support Center for Windows Server 2003.

**Windows Management Instrumentation**

Windows Management Instrumentation (WMI) helps you manage your network and applications as they become larger and more complex. With WMI, you can monitor, track, and control system events that are related to software applications, hardware components, and networks. WMI includes a uniform scripting application programming interface (API), which defines all managed objects under a common object framework that is based on the Common Information Model (CIM). Scripts use the WMI API to access information from different sources. WMI can submit queries that filter requests for very specific information, and it can subscribe to WMI events based on your particular interests, rather than being limited to events predefined by the original developers. For more information, see Windows Management Instrumentation overview in Help and Support Center for Windows Server 2003.

**Microsoft Operations Manager 2000**

The Microsoft® Operations Manager 2000 management packs have a full set of features to help administrators monitor and manage both the events and performance of their IT systems based on the Microsoft® Windows® 2000 or Windows Server 2003 operating systems. Microsoft Operations Manager 2000 Application Management Pack improves the availability of Windows-based networks and server applications. Microsoft Operations Manager 2000 and Windows Server 2003 are sold separately. For information about Microsoft Operations Manager 2000 and the Application Management Pack, see the Microsoft Operations Manager (MOM) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Simple Network Management Protocol**

Simple Network Management Protocol (SNMP) allows you to capture configuration and status information about systems in your network and have the information sent to a designated computer for event monitoring. For more information, see "SNMP" in Help and Support Center for Windows Server 2003.

**Event logs**

When you diagnose a system problem, event logs are often the best place to start. By using the event logs in Event Viewer, you can gather important information about hardware, software, and system problems. Windows Server 2003 records this information in the system log, the application log, and the security log. In addition, some system components such as the Cluster service and FRS also record events in a log. For more information about event logs, see "Checking event logs" in Help and Support Center for Windows Server 2003.

**Service Control Manager**

Service Control Manager (SCM), a tool introduced with the release of the Microsoft® Windows NT® version 4.0 operating system, maintains a database of installed services in the registry. SCM can provide high availability because you can configure it to autorestart services after they have failed. For more information about SCM, see the topic "Service Control Manager" in the Windows Platform SDK documentation.

**Performance Logs and Alerts**

Performance Logs and Alerts collects performance data automatically from local or remote computers. You can collect a variety of information about key resources such as CPU, memory, disk space, and the resources needed by the application. When planning your performance logging, determine the information you need and collect it at regular intervals. Be aware, however, that performance sampling consumes CPU and memory resources, and that excessively large performance logs are hard to store and hard to extract useful information from. For more information about automatically collecting performance

data, see "Performance Logs and Alerts overview" in Help and Support Center for Windows Server 2003.

**Shutdown Event Tracker**

You can document the reasons for shutdowns and save the information in a standard format by using Shutdown Event Tracker. You can use codes to categorize the major and minor reasons for each shutdown and record a comment for the shutdown. For more information, see "Shutdown Event Tracker overview" in Help and Support Center for Windows Server 2003.

# Planning for Unreliable Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An unreliable application is an application, usually proprietary to an organization, that your business cannot do without, but that does not meet high standards for reliability. There are two basic approaches you can take if you must work with such an application:

- Remove the unreliable applications from the servers that are most critical to your enterprise. If an application is known to be unreliable, take steps to isolate it, and do not run the application on a mission-critical server.

- Provide sufficient monitoring, and use automatic restarting options where appropriate. Sufficient monitoring requires taking snapshots of important system performance measurements at regular intervals. You can set up automatic restarting of an application or service by using the Services snap-in. For more information about services, see "Services overview" in Help and Support Center for Windows Server 2003.

# Avoiding Operational Pitfalls

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following sections describe operational practices that can limit the availability of applications and computers, in both clustered and nonclustered environments.

**Supporting multiple versions of the operating system, service packs, and out-of-date applications**
Support of a highly available system becomes much more difficult when multiple combinations of different versions of software (and hardware) are used together in one system or in systems that interact on the network. Older software, protocols, and drivers (and the associated hardware) become impractical when they do not support new technologies.

Set aside resources and time for planning, testing, and installing new operating systems, applications, and (where appropriate) hardware. When planning software upgrades, work with users to identify the features they require. Provide training to ease users through software transitions. In your budget for software and support, provide funds for upgrading applications and operating systems in the future.

**Installing incompatible hardware**
Maintain and follow a hardware standard for new systems, spare parts, and replacement parts.

**Failing to plan for future capacity requirements**
Capacity planning is critical to the success of highly available systems. Study and monitor your system during peak loads to understand how much extra capacity currently exists in the system.

**Performing outdated procedures**
Make sure you remove any outdated procedures from operation and support schedules when a root system problem is fixed. For example, when software is replaced or upgraded, certain procedures might become unnecessary or no longer be valid. Pay special attention to procedures that may have become routine. Be sure that all procedures are necessary and not simply temporary fixes for issues for which the root cause has not been found.

**Failing to monitor the system**
If you do not use adequate monitoring, you might not have the ability to catch problems before they become critical and cause system failures. Without monitoring, an application or server failure may be the only notification you receive of a problem.

**Failing to determine the nature of the problem before reacting**
If the operations staff is not trained and directed to analyze problems carefully before reacting, your personnel can spend large amounts of time responding inappropriately to a problem. They also might not use monitoring tools effectively in the crucial time between the first signs of a problem and an actual failure.

**Treating symptoms instead of root cause**
Symptom treatment is an effective strategy for restoring service when an unexpected failure occurs or when performing short-term preventative maintenance. However, symptom treatments that are added to standard operating procedures can become unmanageable. Support personnel can be overwhelmed with symptom treatment and might not be able to react properly to new failures.

**Stopping and restarting to end error conditions**
Stopping and restarting a computer may be necessary at times. However, if this process temporarily fixes a problem but leaves the root cause untouched, it can create more problems than it solves.

# Evaluating Windows Clustering

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 provides two clustering technologies: server clusters and Network Load Balancing. The clustering technologies in Windows Server 2003 provide different availability and scalability solutions, and they can be implemented throughout your organization.

You can also utilize an additional type of clustering, Component Load Balancing (CLB) clusters, in your Windows Server 2003 deployment. CLB clusters, however, are a feature of Microsoft® Application Center 2000, and they are not discussed in detail here.

**Note**

- You can use CLB clusters on Windows Server 2003, provided that you use Application Center 2000 Service Pack 2 or later.

For complete information about CLB clusters, see your Microsoft Application Center 2000 documentation. Figure 6.4 illustrates the high-level process for evaluating clustering technologies.

**Figure 6.4 Evaluating Clustering**



# Evaluating the Benefits of Clustering

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A cluster is two or more computers working together to provide higher availability, reliability, and scalability than can be obtained by using a single system. When failure occurs in a cluster, resources are redirected and the workload is redistributed. Microsoft cluster technologies guard against three specific types of failure:

- Application and service failures, which affect application software and essential services.
- System and hardware failures, which affect hardware components such as CPUs, drives, memory, network adapters, and power supplies.
- Site failures in multisite organizations, which can be caused by natural disasters, power outages, or connectivity outages.

**Benefits of Clustering**

If one server in a cluster stops working, a process called *failover* automatically shifts the workload of the failed server to another server in the cluster. Failover ensures continuous availability of applications and data.

This ability to handle failure allows clusters to meet two requirements that are typical in most data center environments:

- **High availability.** The ability to provide end users with access to a service for a high percentage of time while reducing unscheduled outages.
- **High reliability.** The ability to reduce the frequency of system failure.

Additionally, Network Load Balancing clusters address the need for high scalability, which is the ability to add resources and computers to improve performance.

**Limitations of Clustering**

Server clusters are designed to keep applications available, rather than keeping data available. To protect against viruses, corruption, and other threats to data, organizations need solid data protection and recovery plans. Cluster technology cannot protect against failures caused by viruses, software corruption, or human error.

The Cluster service, the service behind server clusters, depends on compatible applications and services to operate properly. The software must respond appropriately when a failure occurs. Administrators must be able to configure where an application stores its data on the server cluster. Also, clients that are accessing a clustered application or service must be able to reconnect to the cluster virtual server after a failure has occurred and a new cluster node has taken over the application.

Only services and applications that use TCP/IP for client-server communication are supported on Network Load Balancing clusters and server clusters.

You cannot use Windows Server 2003 File Replication service (FRS) on shared cluster storage. You also cannot create domain-based Distributed File System (DFS) roots on shared cluster storage. Finally, without the proper management tools, you also cannot use dynamic disks on shared cluster storage. For more information about using dynamic disks on shared cluster storage, see article 237853, "Dynamic Disk Configuration Unavailable for Server Cluster Resources" In the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. DFS and FRS are discussed in detail in "Designing and Deploying File Servers" in this book.

**Clustering vs. Fault-Tolerant Hardware Components**

Both clustering and fault-tolerant hardware protect your system from failures of components such as the CPU, memory, fan, or PCI bus. Fault-tolerant hardware is discussed in the section "Planning and Designing Fault-Tolerant Hardware Solutions" earlier in this chapter. Clustering and fault tolerance can be used together in a complete end-to-end solution, but be aware that the two technologies provide high availability in different ways.

Clustering can protect your system against an application or operating system failure, but a fault-tolerant standby server (or a server that uses *hot-swappable* hardware, which allows a device to be added while the server is running) cannot. It is also possible to upgrade an application or operating system or to install a service pack or hotfix without taking a cluster offline.

Upgrades on standby servers, however, are only possible by taking that hardware offline.

# Evaluating Cluster Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you decide to deploy clustering with Windows Server 2003, choosing a cluster technology depends greatly on the application or service that you want to host on the cluster. Server clusters and Network Load Balancing both provide failover support for IP-based applications and services that require high scalability and availability. Each type of cluster, however, is intended for different kinds of services. Your choice of cluster technologies depends primarily on whether you run *stateful* or *stateless* applications:

- **Server clusters are designed for stateful applications.** Stateful applications have long-running in-memory state, or they have large, frequently updated data states. A database such as Microsoft® SQL Server™ 2000 is an example of a stateful application.

- **Network Load Balancing is intended for stateless applications.** Stateless applications do not have long-running in-memory state. A stateless application treats each client request as an independent operation, and therefore it can load-balance each request independently. Stateless applications often have read-only data or data that changes infrequently. Web front-end servers, virtual private networks (VPNs), and File Transfer Protocol (FTP) servers typically use Network Load Balancing.

**Server Clusters**

Server clusters provide high availability for more complex, stateful applications and services by allowing the failover of resources. Server clusters also maintain client connections to applications and services. If your application is stateful, with frequent changes to data, a server cluster is a more appropriate solution. Server clusters run on Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

Server clusters, which use the Cluster service, maintain data integrity and provide failover support and high availability for mission-critical applications and services on your back-end servers, including databases, messaging systems, and file and print services. Organizations can use server clusters to make applications and data available on multiple servers that are linked together in a server cluster configuration. Back-end applications and services, such as messaging applications like Microsoft Exchange or database applications like Microsoft SQL Server, are ideal candidates for server clusters.

In server clusters, nodes share access to data. Nodes can be either active or passive, and the configuration of each node depends on the operating mode (active or passive) and how you configure failover in the cluster. A server that is designated to handle failover must be sized to handle the workload of the failed node in addition to its own workload.

In Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, server clusters can contain up to eight nodes. Each node is attached to one or more cluster storage devices, which allow different servers to share the same data. Because nodes in a server cluster share access to data, the type and method of storage in the server cluster is very important.

For information about server clusters, including resource groups, failover and failover policies, storage configurations, cluster-aware applications, and optimizing your network for server clusters, see "Planning your server cluster" in Help and Support Center for Windows Server 2003 and "Designing and Deploying Server Clusters" in this book.

**Network Load Balancing**

If your application is stateless or can otherwise be cloned with no decline in performance, consider deploying Network Load Balancing. Network Load Balancing provides failover support for IP-based applications and services that require high scalability and availability. Network Load Balancing can run on all editions of Windows Server 2003.

Network Load Balancing addresses bottlenecks caused by front-end services, providing continuous availability for IP-based applications and services that require high scalability. Network Load Balancing clusters are used to provide scalability for Web services and other front-end servers, such as VPN servers and firewalls. Organizations can build groups of clustered computers to support load balancing of TCP and User Datagram Protocol (UDP) traffic requests.

Network Load Balancing clusters are groups of identical, typically cloned computers that, through their numbers, enhance the availability of Web servers, Microsoft® Internet Security and Acceleration (ISA) servers (for proxy and firewall servers), and other applications that receive TCP and UDP traffic. Because Network Load Balancing cluster nodes are usually identical clones of each other and can therefore operate independently, all nodes in a Network Load Balancing cluster are active.

You can scale out Network Load Balancing clusters by adding additional servers. As demand on the cluster increases, you can

scale out Network Load Balancing clusters to as many as 32 servers if necessary. Each node runs a copy of the IP-based application or service that is being load balanced and stores all the data necessary for the application or service to run on local drives.

In clusters, stateless applications are typically cloned, so that multiple instances of the same code are executed on the same dataset. Figure 6.5 shows a cloned application (called "App") deployed in a cluster. Each instance of the cloned application is self-contained, so that a client can make a request to any instance and will always receive the same result.

**Figure 6.5 Cloned Application**




Changes made to one instance of a cloned stateless application can be replicated to the other instances, because the dataset of stateless applications is relatively static. Because stateful applications such as Microsoft Exchange or Microsoft SQL Server are updated with new data frequently, they cannot be easily cloned and they are not good candidates for hosting on Network Load Balancing clusters.

For more information about Network Load Balancing, see "Designing Network Load Balancing" and "Deploying Network Load Balancing" in this book.

**Component Load Balancing**

CLB clusters address the unique scalability and availability needs of middle-tier (business) applications that use the COM+ programming model. Organizations can load balance COM+ components over more than one node to dramatically enhance the availability and scalability of software applications. CLB clusters, however, are a feature of Microsoft Application Center 2000. For information about CLB clusters, see your Microsoft Application Center 2000 documentation.

# Using Clusters to Increase Availability and Scalability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Different types of clusters provide different benefits. Network Load Balancing clusters are designed to provide scalability because you can add nodes as your workload increases. Server clusters increase availability of stateful applications, and they can also allow you to consolidate servers and save on hardware costs. Figure 6.6 shows the steps for planning cluster deployment.

**Figure 6.6 Planning Cluster Deployment**




**Scaling by Adding Servers**

Scaling by adding servers to a Network Load Balancing cluster is also known as *scaling out*. In Windows Server 2003, Network Load Balancing clusters can contain up to 32 nodes. Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, support server clusters containing up to eight nodes. For this reason, it is recommended that you use Network Load Balancing rather than server clusters to increase scalability. Additional servers allow you to meet increased demand on your Network Load Balancing clusters; however, not every organization has the budget to readily add hardware.

Table 6.2 summarizes the number of servers a cluster can contain, by cluster type and operating system.

**Table 6.2 Maximum Number of Nodes in a Cluster**

| Operating System | Network Load Balancing | Component Load Balancing* | Server Cluster |
|---|---|---|---|
| Microsoft® Windows® 2000 Advanced Server | 32 | 12 | 2 |
| Microsoft® Windows® 2000 Datacenter Server | 32 | 12 | 4 |
| Windows Server 2003, Standard Edition | 32 | 12 | N/A |
| Windows Server 2003, Enterprise Edition | 32 | 12 | 8 |
| Windows Server 2003, Datacenter Edition | 32 | 12 | 8 |

* Component Load Balancing is not included with the Windows Server 2003 operating system and runs only on Windows Application Center 2000. You can use CLB clusters with Windows Server 2003, provided that you use Windows Application Center 2000 Service Pack 2 or later. For complete information about Component Load Balancing, see your Windows Application Center 2000 documentation.

Network Load Balancing Clusters also run on Windows Server 2003, Web Edition; the maximum number of nodes is 32.

**Scaling by Adding CPUs and RAM**

Your options for adding CPUs and RAM, also known as *scaling up*, depend on the server operating system. This is a less expensive scaling option than adding nodes, but you are limited by the capacity of the operating system.

Table 6.3 illustrates the CPU and RAM capacity in Windows 2000 and Windows Server 2003.

**Table 6.3 Maximum Number of Processors and RAM**

| Operating System | Number of Processors | Maximum RAM |
|---|---|---|
| Windows 2000 Advanced Server | 8 | 8 GB |
| Windows 2000 Datacenter Server | 32 | 64 GB |
| Windows Server 2003, Enterprise Edition | 8 | 32 GB |
| Windows Server 2003, Datacenter Edition | 32 | 64 GB |

**Availability and Server Consolidation with Server Clusters**

While Network Load Balancing can meet your scaling needs, server clusters provide high availability for hosted applications by reducing the interruption of service due to unplanned outages and hardware failure.

Increasing the number of nodes in a server cluster proportionally increases the availability of the services and applications running on that server cluster. This is because spreading the workload over a greater number of nodes allows each node to run at a lower capacity. When failover occurs in a server cluster with many nodes, there are more servers available to accept the workload of the failed node. Similarly, when a server running at low capacity takes on additional processing after a failover, the failover is less likely to result in a decline in performance.

The hardware required for server cluster nodes is expensive, however. In some cases, it might be possible for you to save money on hardware by consolidating the number of servers or nodes you have deployed in your server clusters. You can still provide high availability of your mission-critical applications after server consolidation, but be aware that combining clusters to reduce the total number of nodes also reduces availability. This is a tradeoff that your organization should evaluate.

Figure 6.7 shows two clusters before consolidation, where two separate two-node clusters, each with one active and one passive node, provide service for a group of clients.

**Figure 6.7 Two Server Clusters Before Consolidation**




Two clusters, each dedicated to a separate application, generally have higher availability than one cluster hosting both applications. This is depends on, among other factors, the available capacity of each server, other programs that may be running on the servers, and the hosted applications themselves. However, if a potential loss in availability is acceptable to your organization, you can consolidate servers.

Figure 6.8 shows what the clusters in Figure 6.7 would look like if they were consolidated into a single three-node cluster. There is a potential loss in availability because in the event of a failure, both active clusters will fail over to the same passive node, whereas before consolidation, each active node had a dedicated passive node. In this example, if both active nodes were to fail at the same time, the single passive node might not have the capacity to take on the workloads of both nodes simultaneously. Your organization must consider such factors as the likelihood of multiple failures in the server cluster, the importance of keeping the applications in this server cluster available, and if potential loss in services is worth the money saved by server consolidation.

**Figure 6.8 Consolidated Server Cluster**


A
A
P
Cluster

# Evaluating Combining Cluster Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use several Windows clustering solutions together, depending on the services and applications that you need to make highly available. Each cluster type performs a different role in a multicluster deployment. The most common scenario that combines all three clusters is an e-commerce Web site, where the front-end Web servers are configured in Network Load Balancing clusters that receive client requests, the middle-tier applications (such as the Microsoft® BizTalk® Server) use CLB clusters running on Windows Application Center 2000, and the back-end SQL database servers operate in a server cluster. Figure 6.9 illustrates a typical deployment, utilizing all three Windows clustering technologies.

**Figure 6.9 E-Commerce Clustering Scenario**



# Evaluating Clustering Across Multiple Sites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many organizations build disaster recovery and increased availability into their infrastructure by using multiple physical sites. Most designs involve a primary site and one or more remote sites, where a remote site mirrors the primary site. The level at which components are mirrored between sites depends on the Service Level Agreement and the business requirements.

**Geographically Dispersed Clusters**

Geographically dispersed server clusters are possible in Windows Server 2003 by using virtual local area networks (VLANs)to connect SANs over long distances. Majority node set, a new type of quorum resource in Windows Server 2003, makes it easier for cluster servers to be geographically separated while maintaining consistency in the event of a node failure.

A geographically dispersed cluster is a combination of pieces supplied by hardware OEMs and software vendors. The configurations required can be complex, and the clusters must use only components supported by Microsoft. Geographically dispersed clusters should be deployed only in conjunction with vendors who provide qualified configurations.

For more information about geographically dispersed clusters and SANs, see "Designing and Deploying Server Clusters" in this book. For complete information about geographically dispersed clusters, see the Geographically Dispersed Clusters link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**MultiSite Network Load Balancing Clusters**

Like geographically dispersed server clusters, Network Load Balancing clusters that span more than one physical site (sometimes known as *stretch clusters*) are intended to protect data in the event of a disaster. Additionally, multisite configurations can be used to provide local access to clients that are spread nationally or globally. For example, a Network Load Balancing cluster split between a data center on the East Coast and a data center on the West Coast can provide services to users in either location.

Unlike geographically dispersed server clusters, in most cases a single Network Load Balancing cluster is not deployed across multiple sites. Instead, multiple, independent Network Load Balancing clusters are deployed, typically with one cluster at each location. The sites are then combined into a single service using features such as DNS round robinto load balance client requests across the sites. Figure 6.10 illustrates a simple version of this scenario.

**Figure 6.10 Multisite Network Load Balancing Configuration**




In Figure 6.10, the client uses DNS to resolve the service name to an IP address. Depending on the IP address that the client receives, the request is sent to Site 1 or Site 2. In this type of configuration, the sites are essentially clones of each other, and the sites provide the same set of services. A Web site, for example, would be identical at either site, and clients receive the same response to a given query regardless of the site they are directed to. Not only are client requests load-balanced between the sites, but in the event that one site fails, the other site can continue responding to client requests. It is worth noting that when one site goes down, there is potential for a longer response time to individual requests, depending on geographical factors and network demand.

# Additional Resources for Planning Server Deployments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Related Information**

- "Designing and Deploying Server Clusters" in this book for more information about planning and deploying server clusters.
- "Designing Network Load Balancing" in this book for more information about designing Network Load Balancing clusters.
- "Deploying Network Load Balancing" in this book for information about deploying Network Load Balancing clusters.
- "Designing and Deploying File Servers" in this book for information about clustered file servers.
- "Designing and Deploying Print Servers" in this book for information about clustered print servers.
- "Deploying DHCP" in *Deploying Network Services* of this kit for information about deploying clustered DHCP servers.
- "Deploying WINS" in *Deploying Network Services* of this kit for information about deploying clustered WINS servers.

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Configuring WINS replication" in Help and Support Center for Windows Server 2003.
- "DNS zone replication in Active Directory" in Help and Support Center for Windows Server 2003.
- "Planning Your Server Cluster" in Help and Support Center for Windows Server 2003.
- "Choosing applications to run on a server cluster" in Help and Support Center for Windows Server 2003.
- "Determining which applications to use with Network Load Balancing" in Help and Support Center for Windows Server 2003.
- "WMI Overview" in Help and Support Center for Windows Server 2003 for information about Windows Management Instrumentation, which is used to manage networks and networked applications.
- "Services" in Help and Support Center for Windows Server 2003 for information about the Services snap-in.
- "Backing Up and Recovering Data" in Help and Support Center for Windows Server 2003.
- "Monitoring and Status Tools" in Help and Support Center for Windows Server 2003.
- "Performance Logs and Alerts overview" in Help and Support Center for Windows Server 2003.
- "SNMP" in Help and Support Center for Windows Server 2003.
- "Checking event logs" in Help and Support Center for Windows Server 2003.
- "Shutdown Event Tracker overview" in Help and Support Center for Windows Server 2003.

# Designing and Deploying Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Server clusters ensure that applications continue to run in the event of planned service downtime due to maintenance or during unplanned downtime due to hardware or network failure. In order to design server clusters, organizations need to understand the applications that they intend to host on a server cluster, namely the best way to deploy and configure the application in a clustering environment, and they also need to know the server cluster's storage and capacity requirements.

**In This Chapter**


Overview of the Server Clusters Design Process

Analyzing High-Availability Requirements

Determining How to Deploy Applications

Designing the Support Structure for Server Clusters

Protecting Data from Failure or Disaster

Securing Server Clusters

Deploying Server Clusters

Additional Resources for Server Clusters


**Related Information**

- For more information about planning for high availability see "Planning for High Availability and Scalability" in this book.
- For more information about Network Load Balancing clusters, see "Designing Network Load Balancing" in this book.
- For more information about clustered print servers, see "Designing and Deploying Print Servers" in this book.
- For more information about clustered file servers, see "Designing and Deploying File Servers" in this book.
- For more information about designing storage systems, see "Planning for Storage" in this book.

# Overview of the Server Clusters Design Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Mission-critical applications, such as corporate databases and e-mail, must reside on systems that are designed for high availability and scalability. Deploying server clusters with the Microsoft® Windows® Server 2003, Enterprise Edition or the Windows® Server 2003, Datacenter Edition operating system minimizes the amount of planned and unplanned server downtime. Server clusters can benefit your organization if:

- Your users depend on regular access to mission-critical data and applications to do their jobs.
- Your organization has established a limit on the amount of planned or unplanned service downtime that you can sustain.
- The cost of the additional hardware that server clusters require is less than the cost of having mission-critical data and applications offline during a failure.

Windows Server 2003 provides two different clustering technologies: server clusters (discussed in this chapter) and Network Load Balancing. An overview of these cluster solutions is provided in "Planning for High Availability and Scalability" in this book. For more information about Network Load Balancing, see "Designing Network Load Balancing" and "Deploying Network Load Balancing" in this book. Note that this chapter focuses on using server clusters to make applications highly available. For information about using server clusters with resources such as file and print services, see "Designing and Deploying Print Servers" and "Designing and Deploying File Servers" in this book.

A third clustering technology, Component Load Balancing (CLB), is available in Microsoft® Application Center 2000 but not in Windows Server 2003. CLB clusters enable COM+ applications to be distributed across multiple servers. For more information about CLB, see the documentation for Microsoft Application Center 2000.

If your business is considering server clusters to provide increased availability for one or more applications, and you plan to use Windows Server 2003 — either Enterprise Edition or Datacenter Edition — this chapter will help you evaluate, design, and implement a server cluster compatible with those applications. If you are already running server clusters on Windows NT 4.0 or Windows 2000, this chapter helps you prepare to upgrade those clusters to Windows Server 2003.

# Server Cluster Design and Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You begin the server cluster design and deployment process by defining your high-availability needs. After you determine the applications or services you want to host on a server cluster, you need to understand the clustering requirements of those applications or services. Next, design your server cluster support network, making sure that you protect your data from failure, disaster, or security risks. After you evaluate and account for all relevant hardware, software, network, and security factors, you are ready to deploy a server cluster. Figure 7.1 illustrates the server cluster design process.

**Figure 7.1 Designing and Deploying Server Clusters**



# Server Cluster Fundamentals

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following sections provide an overview of key clustering concepts and summarize new clustering features that have been added to Windows Server 2003.

A *cluster* is a group of individual computer systems working together cooperatively to provide increased computing power and to ensure continuous availability of mission-critical applications or services. From the client viewpoint, an application that runs on a server cluster is no different than an application that runs on any other server, except that availability is higher.

**Clustering Terms**

The following terms are fundamental clustering concepts that are common to almost every clustering decision. For more information about the following terms and for an overview of server clusters, see "Server Cluster Components" and "Cluster Objects" in Help and Support Center for Windows Server 2003.

**Node** A computer system that is a member of a server cluster. Windows Server 2003 supports up to eight nodes in a server cluster.

**Resource** A physical or logical entity that is capable of being managed by a cluster, brought online, taken offline, and moved between nodes. A resource can be owned only by a single node at any point in time.

**Resource groups** A collection of one or more resources that are managed and monitored as a single unit. Resource groups can be started and stopped independently of other groups (when a resource group is stopped, all resources within the group are stopped). In a server cluster, resource groups are indivisible units that are hosted on one node at any point in time. During failover, resource groups are transferred from one node to another.

**Virtual server** A collection of services that appear to clients as a physical Windows-based server but are not associated with a specific server. A virtual server is typically a resource group that contains all of the resources needed to run a particular application and can be failed over like any other resource group. All virtual servers must include a Network Name resource and an IP Address resource.

**Failover** The process of taking resource groups offline on one node and bringing them back online on another node. When a resource group goes offline, all resources belonging to that group go offline. The offline and online transitions occur in a predefined order. Resources that are dependent on other resources are taken offline before and brought online after the resources upon which they depend.

**Failback** The process of moving resources, either individually or in a group, back to their original node after a failed node rejoins a cluster and comes back online.

**Quorum resource** The quorum-capable resource selected to maintain the configuration data necessary for recovery of the cluster. This data contains details of all of the changes that have been applied to the cluster database. The quorum resource is generally accessible to other cluster resources so that any cluster node has access to the most recent database changes. By default there is only one quorum resource per cluster.

**Cluster Hardware Requirements**

The following are general hardware requirements for building Windows Server 2003 clusters. More detailed information is provided throughout this chapter where applicable. In order for server clusters to be supported by Microsoft, all hardware must be selected from a list of qualified clustering solutions. See the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. For more information about support for server clusters, see article Q309395, "The Microsoft Support Policy for Server Clusters and the Hardware." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

- Use a minimum of two computers running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition (computer hardware must also be listed in the Windows Server Catalog). You cannot mix x86-based and Itanium architecture–based computers in the same server cluster.

- Meet minimum storage requirements. Minimum storage requirements depend on a number of factors, such as whether or not your server cluster uses storage area network (SAN) technology, or the type of quorum resource used in the server cluster. For more information about storage solutions, see "Designing the Support Structure for Server Clusters" later in this chapter. Note that all storage solutions must also be listed in the Windows Server Catalog.

- Use a minimum of two network adapters for each node. In recommended configurations, one network adapter connects the node to the other nodes in the cluster for communication and configuration purposes (private network). The second adapter connects the cluster to both an external network (public network) and the private network.

**Important**

- Server clusters do not protect against hardware failures. Microsoft recommends a redundant array of independent disks (RAID) solution for all cluster disks to eliminate disk drives as a potential single point of failure.

**New in Windows Server 2003**

Windows Server 2003 introduces a number of new clustering features that are highlighted below.

**More nodes in a cluster** Windows Server 2003 supports up to eight nodes per cluster.

**Support for 64-bit versions of Windows Server 2003 operating system** The 64-bit versions of Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition support the Cluster service. Note that you cannot use GUID partition table (GPT) disks for shared cluster storage. A GPT disk is an Itanium–based disk partition style in the 64-bit versions of Windows Server 2003. For 64-bit versions of Windows Server 2003, you must partition cluster disks on a shared bus as master boot record (MBR) disks and not as GPT disks.

**Simplified cluster configuration and setup** Server cluster system files are installed by default with Windows Server 2003. New features include the ability to create new server clusters — or add nodes to an existing cluster — remotely. Another new feature, the New Server Cluster Wizard, analyzes hardware and software configurations to identify potential problems before installation.

**Security enhancements** You can reset the Cluster service account password without stopping the Cluster service, allowing you to maintain corporate password policies without compromising availability. In addition, server clusters now support the Kerberos version 5 authentication protocol. For more information, see "Securing Server Clusters" later in this chapter.

**Scripting** An application can be made server cluster-aware through scripting (both VBScript and Jscript are supported), rather than through resource dynamic-link library (DLL) files. Unlike DLLs, scripting does not require knowledge of the C or C++ programming languages, which means scripts are easier for developers and administrators to create and implement. Scripts are also easier to customize for your applications.

**Majority node set clusters** In every cluster, the quorum resource maintains all configuration data necessary for the recovery of the cluster. In majority node set clusters, the quorum data is stored on each node, allowing for, among other things, geographically dispersed clusters. For more information about quorums, see "Designing the Support Structure for Server Clusters" later in this chapter. For more information about majority node set clusters, see "Protecting Data from Failure or Disaster" later in this chapter.

# Analyzing High-Availability Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After determining that you have mission-critical applications or services that need to be highly available, analyze your business needs to determine if clustering is the best solution. You must then evaluate your applications for compatibility with clustering. Figure 7.2 illustrates the process for determining if server clusters are appropriate for your needs.

**Figure 7.2 Analyzing High-Availability Requirements**



# Analyzing High-Availability Business Needs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This chapter assumes that you have already decided to make one or more applications or services highly available, based on your organization's business needs. You need high availability if users require regular access to data or if your organization cannot tolerate downtime from an application or service for more than a minimal amount of time. Other common business motivations for high availability include contractual obligations, concern about losing sales during the time it takes to rebuild a server, and the potential for lost employee productivity when data is inaccessible.

If your business or organization requires high availability, answer the following questions before you begin designing a server cluster:

- How many users depend on a particular application?
- Are your users located at more than one physical site?
- What is the minimum acceptable downtime for a particular application?
- How many failures does your cluster need to tolerate? One node failure? Two nodes or more?
- How much hardware (particularly cluster nodes) can be dedicated to the server cluster?

Additionally, if you have more than one application that needs to be made highly available, do some applications take priority over other applications that will be running on the server cluster? If the answer is yes, consider prioritizing the applications. For example, if more than one node fails in your server cluster, and the remaining hardware is not sufficient to take on every failed node's workload, which applications and services must continue working?

Understanding the business requirements that must be handled by your server cluster is important, particularly later in the server cluster design process when you design failover policies. For more information about these requirements, see "Determining How to Deploy Applications" later in this chapter.

If a server cluster is the most appropriate high-availability solution for your business or organization, the next step is to make sure your application can run on a server cluster.

# Determining Application Compatibility for Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Not all applications can be adapted to run on a server cluster. If your application meets all the following conditions, it will likely work in a server cluster:

- Your application uses a protocol based on Internet Protocol (IP).
- Your application can specify where application data is stored.
- Your application attempts to reconnect or recover from network failures.

IP-based protocols include Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Distributed Component Object Model (DCOM), named pipes, and remote procedure call (RPC) over TCP/IP. NetBEUI and Internetwork Packet Exchange (IPX) protocols do not work on a server cluster.

**Tip**

- For information about pretesting to determine if your application is cluster ready, or for information about whether your application supports failover, whether clients can survive failure and subsequent restart of your application, and if clients can survive a failure without a restart of your application, see "Choosing applications to run on a server cluster" in Help and Support Center for Windows Server 2003.

Additionally, some applications do not react to cluster events. An application must be *cluster-aware* (must communicate with the Cluster application programming interface) to receive status and notification information from the server cluster. The Cluster application programming interface (API) is a collection of functions that are implemented by the cluster software and used to manage the cluster, cluster objects, and the cluster database.

Most older applications are not cluster-aware, but they can still run in a server cluster and fail over. These applications respond to only the most basic methods of failure detection and shutdown, and they are not capable of the more advanced initialization and cleanup tasks performed by applications that can react to cluster events. For this reason, these applications might not be as highly available in a server cluster as are applications that can interact with the cluster.

For more information about applications that can run on a cluster, see the Software Development Kit (SDK) link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources and search for the help topic "Server Cluster Application Types."

Before designing your server cluster, contact the software vendor or consult the documentation to see if your application can run on a cluster. This information can affect how you install your application in a server cluster, and might also determine the best way to upgrade an application once it is installed in a server cluster.

# Determining How to Deploy Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some of your deployment decisions are dictated by the application you plan to deploy in the server cluster. Figure 7.3 displays the process for analyzing your application and determining the best way to deploy the application in a server cluster.

**Figure 7.3 Determining How to Deploy Applications**



# Evaluating the Application

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need a thorough understanding of your application's installation needs before deploying the application in a server cluster. For every application that you intend to install in a server cluster, check the documentation or contact the software vendor about installation issues in general and clustering considerations in particular.

The information you need before deploying an application in a server cluster includes:

- **Setup information.** How the application setup works, such as where the application is installed and where supporting data for the application is installed. When installing Microsoft® Exchange 2000 Server on a server cluster, for example, the log files and database in each storage group must be on the cluster's shared disk. This allows the log files and the storage group databases to fail over to another node if the Exchange virtual server goes offline.

- **Upgrade information.** When it is time to upgrade the application, can the application be upgraded online, or will the server cluster need to be taken offline?

- **Application behavior during failover.** Anticipating application behavior during failover will also help to optimize your deployment. In Microsoft® SQL Server™, Microsoft recommends making all transactions as small as possible because the startup process for SQL virtual servers includes going through the transaction log for each database and rolling transactions back or forward. Therefore, larger transactions, especially with a large volume of transactions, could result in slower failover time.

# Choosing a Method for Deploying an Application on a Server Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The best method for deploying an application on a server cluster depends greatly on what type of application it is. Applications running on a cluster can be characterized in two ways: single-instance and multiple-instance. The way your application behaves in a server cluster depends on which type (single-instance or multiple-instance) it is. This means that your deployment method can also vary by application type.

# Deploying Single-Instance Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With single-instance applications, only one application instance is running on the cluster at any time, and the application typically has data sets that cannot be partitioned across multiple nodes. The Dynamic Host Configuration Protocol (DHCP) service is an example of a single-instance application. The set of leased IP addresses that the application provides is small, but would be complicated to replicate across a cluster. Therefore the DHCP instance exists as a single instance, and high availability is provided by failing over the instance to another node in the cluster.

Figure 7.4 illustrates a single-instance application, named "App," in a four-node cluster. Only one node is currently hosting the application. If that node fails, the application can fail over to one of the other three nodes (the failover options are represented by dotted lines). All four nodes are connected to the application data source, but only the node currently hosting the application can access the data source.

**Figure 7.4 Single-Instance Application**



# Deploying Multiple-Instance Applications

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With a multiple-instance application, either multiple instances of the same code, or different pieces of code that cooperate to provide a single service, can run on a cluster. In both cases, the client or end user sees only one partition of an application. Applications can be either *cloned* or *partitioned* to create multiple instances in a cluster.

**Cloned Applications**

In a cloned deployment, multiple instances of the same code are executed on the same data set. Cloned applications are more commonly deployed in Network Load Balancing clusters than in server clusters, because cloning is ideal for *stateless* applications, which are applications with data that is rarely updated or that is read-only. Server clusters are designed for applications such as databases and messaging systems that have frequently updated data. For more information about Network Load Balancing clusters, see "Designing Network Load Balancing" in this book.

**Partitioned Applications**

Server clusters are designed for *stateful* applications, which are applications with frequently updated data sets and long-running in-memory states. Most stateful applications can be partitioned. A large database, for example, can be partitioned so that names beginning with A through L are in one database partition, and those beginning with M through Z are in another database partition. In other applications, functionality — rather than data — is partitioned. After they are partitioned, the different functions are deployed across a set of nodes in a cluster.

In order to represent a single application instance to clients or end users, a partitioned application requires an application-dependent decomposition, routing, and aggregation mechanism that allows a single client request to be distributed across the set of partitions. Then the results from the partitions need to be combined into a single response back to the client. This mechanism depends on the type of application and type of data returned, and is therefore application-specific. The query engine in SQL Server performs this splitting and aggregating. In Exchange 2000 Server clusters, the Exchange front-end server uses the Active Directory® directory service to perform this task.

Figure 7.5 shows an example of a partitioned application in which a database of customer names has been partitioned on a four-node cluster. The solid lines represent each node's connection to its partition of the database; the dotted lines represent the application's access to a given partition as needed.

**Figure 7.5 Data Partitioning for an Application on a Cluster**




When a client issues a query such as "return all records" across multiple instances, the request is split across the instances and — after results are collected — aggregated back into a single result, which is passed back to the client.

A deployment such as the one in Figure 7.5 allows for scalability, because more servers can be added to house additional partitions, but it does not provide high availability. If one of the nodes fails, that part of the data set is no longer accessible.

As mentioned above, an application can also be partitioned by functionality. Figure 7.6 shows a multiple-instance application deployment, where functionality, not data, is shared among the nodes.

**Figure 7.6 Functional Partitions**




When a client request is introduced to the server cluster in Figure 7.6, the request is directed to the appropriate node, depending on the nature of the request. A billing inquiry, for example, is handled by the node designated for that function. The billing node is the only node with access to the billing database.

The deployment of an application across a server cluster depends on the style of application (single-instance versus multiple-instance) and the number of applications that you plan to deploy in the same cluster.

The following sections describe the most common application deployment options.

**One single-instance application** One node in the cluster executes the application, servicing user requests. The remaining nodes are on standby, ready to host the application in the event of failover. This type of deployment is well suited for a two-node active-passive cluster in which one node runs the applications and the second node is on standby. This deployment is not typically used in clusters with three or more nodes, however, because of hardware costs. In a cluster of *N* nodes, only 1/*N* of total processing power is used. Although availability of the single application is high, even in a two-node cluster hardware costs are higher than with other deployment options because of the cost of purchasing standby servers. Some multiple-instance applications, such as SQL Server, can be deployed as a single instance. In the case of SQL Server, when a node in the cluster running SQL Server fails, the entire SQL Server workload is transferred to another, identically configured node.

**Several single-instance applications** Each application is independent from the others, and different applications can fail over to different nodes in the cluster. For more information about defining failover targets, see "Creating Failover Policies" later in this chapter. This type of deployment provides high availability for applications, but careful capacity planning is required to ensure that each node has sufficient processing capacity, memory, and bandwidth to accommodate additional loads in the event of failover.

**One multiple-instance application** The cluster supports only one application, with pieces of the application hosted on different nodes. When there is a failover with a partitioned application, a node can host multiple partitions until failback. In SQL Server and other multiple-instance applications, it is important to make sure that each node has the capacity to take on additional partitions in the event of failover.

**Several multiple-instance applications** Each node hosts one instance of each application. In such a deployment, both capacity planning and defining failover targets become increasingly important and complex.

# Planning for Server Capacity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before deploying your application on a high-availability server cluster, you need to answer some key business questions:

- How many servers are available for your cluster and what is the capacity of each server?
- In the case of a failover, what kind of performance is acceptable until failback returns the cluster to full processing capacity?
- Given that it is possible for more than one node in a cluster to fail, how many failures do you want your cluster to be able to withstand before the application becomes unavailable?

The workload of a server both before and after failover is closely tied to your server resource planning. The way your application adapts to the server cluster depends on how much of the server resources it consumes. The consumption of server resources affects the number of servers you need to ensure high availability for the application, in addition to how many servers are required to pick up the workload for a given node in the event of a failover.

Consider a simple example: a two-node cluster, with Node A and Node B. If Node A fails over, Node B has to continue processing its own jobs in addition to processing the jobs that fail over from Node A. This means that prior to failover, both nodes can operate at a maximum of 50 percent capacity. After failover, Node B operates at 100 percent capacity. When any server operates at maximum capacity, performance is compromised. However, even reduced performance is preferable to completely losing Node A's workload, which is what would happen in a nonclustered environment. In this example, an organization would have to determine if this level of availability is sufficient for its needs, or if an expenditure on more hardware is warranted.

If you add a third node to the above example and all three nodes operate at 50 percent, the two remaining nodes can divide the workload of a failed node. If Node A fails over, Node B and Node C assume Node A's work and operate at 75 percent capacity, which provides better application performance than a single server operating at 100 percent capacity. However, if two nodes in this cluster were to fail over, the remaining node would not have sufficient capacity to assume the processing for the entire cluster.

Windows Server 2003 supports up to eight nodes in a cluster. With additional nodes, you have more options for distributing the workload in a cluster.

**Cluster Server Capacity Examples**

The following illustrations of a simple four-node cluster demonstrate how server load can be balanced across nodes. These examples also demonstrate the relationship between hardware resources, fault tolerance, and availability.

**Important**

- The examples in this section are theoretical and intended for discussion purposes only. Be aware that in an actual server cluster deployment, applications cannot always be distributed as evenly as they are in the following examples.

In Figure 7.7, all four nodes are operational, each running at 25 percent capacity.

**Figure 7.7 Cluster Under Normal Operating Conditions**




Figure 7.8 shows the cluster after one of the nodes fails. The remaining three nodes each receive a share of the failed node's work.

**Figure 7.8 Cluster After a Single Node Failure**




After two nodes fail, shown in Figure 7.9, the remaining two nodes are each operating at approximately 50 percent capacity.

**Figure 7.9 Cluster After Two Failures**




Finally, when three nodes fail, as shown in Figure 7.10, the remaining node has enough capacity to take over the clusters' entire workload (though application performance might be degraded somewhat with the server operating at 100 percent capacity).

**Figure 7.10 Cluster After Three Failures**




**Deployment Considerations**

When determining the appropriate size of your server cluster (the number of nodes and their capacity) and the failover policies for one or more applications, consider the following points:

- The cluster configuration in Figure 7.7 through Figure 7.10 is able to tolerate three failures with no loss of availability for clients and users. The cost of this availability is the price of four servers of sufficient capacity. For example, if the above nodes were running at 50 percent or 75 percent capacity, the same amount of work could be done with fewer nodes (or the same number of nodes could do more work). However, the cluster would then be unable to accommodate three failures; the remaining nodes would already be at full capacity by the time the third node failed. If you don't need to anticipate three failures, you could use fewer nodes running at a higher capacity than 25 percent.

- In Windows Server 2003, Cluster Administrator allows you to designate applications to fail over to specific nodes. If your application cannot afford to lose any performance when it fails over to another node, configure failover distribution so that the application fails over to a node with plenty of room (running at low capacity). Again, this requires a hardware investment in order to have a server standing by for the purpose of failover.

- If some decline in performance is acceptable for your application, specifically during the time it takes to get the failed node restarted, you can fail over the application to a node that is operating at a higher capacity. This means you can use fewer nodes, but the chances are higher that performance on the new node will decrease after failover.

- If some applications are a higher priority than others to your organization, you can use Cluster Administrator to fail over the high-priority applications before the others. This is useful in deployments where multiple failures cannot be tolerated due to capacity limitations. In this case, if both a high-priority application and a lower-priority application fail over and need to find room elsewhere in the cluster, the high-priority application always takes precedence.

- In planning server capacity (and in all other aspects of cluster deployment planning), it is important that you consider scalability and future demands on your cluster. It is very likely that over time your applications will grow, or you will need to deploy more applications on a cluster, or both. When calculating how many servers you will need and how much capacity each server needs to have, a general IT principal is to increase your estimates by 30 percent, and then add capacity for anticipated future growth. The extra 30 percent provides a safety buffer against calculation error or unforeseen load additions.

# Creating Failover Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To provide high availability, single-instance applications — and the individual partitions of multiple-instance applications — fail over from one cluster node to another. In a two-node cluster, failover involves transferring the resources of one node to the remaining node. As the number of nodes in a cluster increases, different failover policies are possible.

When you create failover policies, you must perform each of the following tasks:

1. Make key failover policy decisions.
2. Evaluate N+I failover policies.
3. Evaluate random versus customized failover control.

# Making Key Decisions in Failover Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to make the following decisions when planning and configuring your failover policies:

- How many nodes, such as standby servers, do you deploy for failover?
- What resources do you fail over to which nodes?
- Do you partition applications and services across cluster nodes, or do you distribute the load differently?
- How do you consolidate resources into resource groups? A group is the indivisible unit (a collection of resources) that fails over from one node to another.

**Note**

- A failover policy often involves targeting a specific node for a given application. Such targeting can be configured through Cluster Administrator. For more information about targeting failover nodes, see "Determining failover and move policies for groups" in Help and Support Center for Windows Server 2003.

Many of the key decisions in failover policies involve weighing the benefits of implementing the best possible solution against the cost of purchasing and maintaining hardware. For example, do your applications and services need the highest level of availability, or is your organization willing to settle for a slightly lower level of availability at a greatly reduced cost? If you decide to purchase standby servers that are dedicated to receiving specific applications in the event of failover, be sure your budget can accommodate the additional cost.

# Evaluating N+I Failover Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The most common modeling used for failover policies is *N+I,* where *N* nodes are hosting applications and *I* (idle) nodes are standing by as spares. The simplest method for providing failover is to match each node with its own idle failover node: in a three-node cluster that would be 3+3. This design achieves the highest level of availability, but requires an investment in hardware that remains idle much of the time.

A failure on a cluster might result in an impact on users that can range from slow performance to application unavailability. When deploying an N+I solution, extensive testing must be done to ensure that, in the case of a massive failure, one node can run all applications and services with little or no user impact.

N+I configurations are scalable and good for handling multiple failures. An example of an N+I configuration is shown in Figure 7.11.

**Figure 7.11 N+I Configuration**




The configuration shown in Figure 7.11 is a good configuration if your application must continue running without any decrease in performance after failover. This configuration guarantees that the application will fail over to a node with 100 percent available resources and not to a node already in use by another application or another partition (assuming there has not already been another failover in the same server cluster).

Exchange 2000 Server is a good example of an application suited for an N+I server cluster deployment. In a four-node cluster, Microsoft recommends three active nodes and one node on standby (3+1). Because Exchange can have large memory and processing requirements, it is better to have a resource group in an Exchange cluster fail over to a node with 100 percent available resources. The potential drawback is that only one node failure is supported in this scenario.

**Standby Server**

A variation of N+I, a standby server configuration has idle nodes prepared to take on the work of one or more active nodes. Standby servers are effective for maintaining a consistent performance level for your users even after failover. A three-node version of a standby server configuration is illustrated in Figure 7.12.

**Figure 7.12 Standby Server Configuration**




Standby servers are good for larger applications that require a lot of resources such as processing power, memory, and input/output (I/O) bandwidth. In the event of a failure on one node, all applications are still hosted on separate servers. In Figure 7.12, for example, App1 and App2 are large applications. Combining them on the same node might require more processing or capacity than a single node could provide. With a standby server configuration, if Node A were to fail, App1 would move to Node B and the applications would remain on separate nodes. It is easy to plan server capacity for this failover policy; the size of each node is based on the application that it will be running, and the capacity of the spare server is sized to the maximum of the other nodes.

The drawback of this failover policy is that there is a single point of failure. In Figure 7.12, for example, if Node A failed, and App1 moved to Node B, no failover could take place if Node B were to also fail. This type of configuration also does not support multiple failures well. For example, if both Node A and Node C were to fail, Node B would not have sufficient resources to run both applications. This example shows a three-node cluster, however. In larger clusters (up to eight nodes) you can configure a variety of standby combinations to eliminate a single point of failure.

**Failover Pairs**

Failover pairs are another common N+I variation. Failover pairs are often used in clusters with four or more nodes, and they are ideal for ensuring performance consistency after failover. Failover policies can be defined so that each application can fail over between two nodes, such as in the four-node cluster illustrated in Figure 7.13. Here, Node A is paired with Node B, and Node C is paired with Node D.

**Figure 7.13 Failover Pair Server Configuration**




Failover pairs are also good for larger, resource-intensive applications such as databases. For capacity planning, you must ensure that each node in the pair can accommodate the application in question. Compared to a standby server policy, failover pairs have the added benefit of being part of a larger cluster. If the failover node needs to be brought down for maintenance, for example, a different failover node for the application can be assigned dynamically.

The disadvantage of failover pairs is that in simpler configurations, such as the one illustrated in Figure 7.13, only 50 percent of the cluster resources are ever in use. Also, depending on the number of nodes and the specific failover policies configured on the nodes, multiple failures in this kind of cluster might require administrators to manually failback resources when nodes are restored.

# Evaluating Random vs. Customized Failover Control

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Depending on the size and complexity of your server cluster, and the resources at your disposal, you might want to enact precise control over failover policies. Alternatively, you can let Windows Server 2003 determine the failover location for an application randomly.

**Random Failover Policies**
In clusters that run several small application instances, where the capacity of any node is large enough to support several instances at a time, random failover policies work well. Random failover policies are not recommended for large clusters (four or more nodes) because, given the great number of combinations of nodes, resources, and failures, there is no way to prevent some nodes from becoming overloaded while other nodes are underused.

When a resource group has an empty preferred owners list, and the node hosting that group fails, Windows clustering randomizes the failover target for the group. Assuming there are sufficient resources and application instances in a cluster, random failover in a small cluster can actually load balance applications across cluster nodes. Random failover is also well able to handle multiple failures.

The disadvantage of random failover policies is that you might find it difficult to plan capacity because statistically there is no guarantee that applications will be equally divided across nodes. Similarly, the effect on the performance of an application after failover is also difficult to predict. Random failover policies are not recommended in clusters with large, resource-intensive applications.

**Customized Failover Policies**
Targeting specific applications to specific nodes gives the administrator full control over what happens when a failure occurs. It also makes capacity planning easy, because you know which application is going to run on which server.

*Group affinity* is an attribute of a dependency between groups that determines the groups' relative location in the server cluster. For resources and resource groups that have no requirements for specific hardware or configurations, it might not matter where the resource groups are hosted on the cluster. There are circumstances, however, where the location of one group can affect the location of other groups. Groups can be set up to have strong or weak *affinity*. With strong affinity, they will be placed together on the same node, if possible. They can also be set up to have strong or weak *anti-affinity*. With strong anti-affinity, groups will be kept apart if at all possible. Exchange virtual servers are an example of groups that use anti-affinity, because running more than one instance of Exchange on the same node is not recommended.

If you have multiple applications, defining good policies for each can be difficult. In the event of multiple failures, it is difficult to account for all the possible combinations of application and server failures. For more information about configuring group affinity, see article Q296799, "How to Configure Windows Clustering Groups for Hot Spare Support." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Designing the Support Structure for Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A server cluster depends on network connectivity for communication with clients and end users, and for intra-cluster configuration information between nodes. The other fundamental requirement for a server cluster is storage ¯ the storage of application information that is shared and failed over between nodes, and the storage of cluster configuration data. Note that your storage design, in addition to providing sufficient storage capacity, must be able to deliver information to the cluster in a timely and efficient manner.

Before you design the support structure for your server clusters, assemble the information that you collected during the steps outlined earlier in this chapter, including:

- Usage statistics, such as the number of users and clients who will be using applications running on your clusters.
- The number of clusters you plan to run on your network, and the number of nodes in each cluster.
- The amount of storage your applications need.
- The scalability needs in your organization, particularly projections for anticipated growth in storage needs and network traffic in the near future.

Figure 7.14 shows the high-level process for analyzing your network and infrastructure needs.

**Figure 7.14 Designing the Support Structure for Server Clusters**




**Important**

- All storage and networking solutions for a server cluster must be supported by Microsoft. For more information about qualified storage and networking solutions, see the Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Choosing a Cluster Storage Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Currently, the most common storage I/O technologies used with Windows Server 2003 clusters are Parallel SCSI and Fibre Channel.

Beginning with the Windows Server 2003 release, SCSI interconnects are supported only on two-node clusters running the 32-bit version of Windows Server 2003, Enterprise Edition. SCSI is not supported on Windows Server 2003, Datacenter Edition, or any 64-bit version of the Windows Server 2003 family.

SCSI storage and two types of Fibre Channel — arbitrated loops and switched network — are currently qualified storage configurations in the Windows Catalog. Both SCSI and Fibre Channel are discussed later in this section.

**Note**

- This section provides an overview of storage technology and implementation recommendations. The information here is intended to provide deployment guidelines and help administrators make informed deployment decisions. For procedural information about installing or configuring storage in your server cluster, see "Storage configuration options" in Help and Support Center for Windows Server 2003.

In addition to configuring adequate storage for applications, you also need a dedicated disk to use as the quorum device. This disk must have a minimum capacity of 50 megabytes (MB). For optimal NTFS file system performance, it is recommended that the quorum disk have at least 500 MB. For more information about the quorum device, see "Choosing the Cluster Model" later in this chapter. It is recommended that you do not use the quorum disk for storing applications or anything other than quorum data.

# Using RAID in Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is strongly recommended that you use a RAID solution for all disks running the Windows Server 2003 operating system. A RAID configuration prevents disk failure from being a single point of failure in your server cluster.

It is possible to use hardware-based RAID with Windows Server 2003. By using a controller provided by the hardware vendor or RAID solution provider, you can configure physical disks into fault-tolerant sets or volumes. The sets can be made visible to clusters either as whole disks or as smaller partitions.

Cluster nodes are unaware of the physical implementation of disks and treat each volume as a physical disk. For an ideal configuration, it is recommended that you associate the logical volume size as closely as possible with the actual physical disk size. This practice avoids taxing the physical disk with too many I/O operations from multiple logical volumes.

If your RAID controller supports dynamic logical unit (LUN) expansion, cluster disks can be extended without rebooting, providing excellent scalability for your organization's needs. The command-line tool DiskPart.exe lets administrators apply the physical extension of the disks to the logical partitions with no disruption of service. For more information about DiskPart.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Command-line reference A-Z.

# Using Storage Interconnects with Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Make sure that all interconnects used on a server cluster are supported by Microsoft by checking the Windows Catalog. This also applies to any software that is used to provide fault tolerance or load balancing for adapters or interconnects. For more information about qualified interconnects and software, see the Windows Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

Consider establishing redundant access routes to stored data. Providing multiple paths to your storage is another way to provide high availability. With versions of Windows operating systems earlier than Windows Server 2003, vendors and storage providers implemented two or more storage fabrics that were configured for load balancing and fault tolerance. The specific solution was designed by the vendor and required specially designed configurations and drivers. With the release of Windows Server 2003, Microsoft developed and supplied vendors with a multipath driver, which they can use in place of customized drivers.

# Using SCSI Storage with Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

SCSI is supported only on two-node clusters running the 32-bit version of Windows Server 2003, Enterprise Edition. Server clusters using SCSI are not supported on the 64-bit version of Windows Server 2003. General guidelines for deploying SCSI in your server cluster are listed below:

- All devices on the SCSI bus, including disks, must have unique SCSI IDs (for example, by default, most SCSI adapters have an ID of 7).
- SCSI hubs are not supported in server clusters.
- The SCSI bus must be terminated. Use physical terminating devices, and do not use controller-based or device-based termination.

# Using Fibre Channel Storage with Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are two supported methods of Fibre Channel-based storage in a Windows Server 2003 server cluster: arbitrated loops and switched fabric.

**Important**

- When evaluating both types of Fibre Channel implementation, read the vendor's documentation and be sure you understand the specific features and restrictions of each.

Although the term *Fibre Channel* implies the use of fiber-optic technology, copper coaxial cable is also allowed for interconnects.

**Arbitrated Loops (FC-AL)**

A Fibre Channel arbitrated loop (FC-AL) is a set of nodes and devices connected into a single loop. FC-AL provides a cost-effective way to connect up to 126 devices into a single network. As with SCSI, a maximum of two nodes is supported in an FC-AL server cluster configured with a hub. An FC-AL is illustrated in Figure 7.15.

**Figure 7.15 FC-AL Connection**




FC-ALs provide a solution for two nodes and a small number of devices in relatively static configurations. All devices on the loop share the media, and any packet traveling from one device to another must pass through all intermediate devices.

If your high-availability needs can be met with a two-node server cluster, an FC-AL deployment has several advantages:

- The cost is relatively low.
- Loops can be expanded to add storage (although nodes cannot be added).
- Loops are easy for Fibre Channel vendors to develop.

The disadvantage is that loops can be difficult to deploy in an organization. Because every device on the loop shares the media, overall bandwidth in the cluster is lowered. Some organizations might also be unduly restricted by the 126-device limit.

**Switched Fabric (FC-SW)**

For any cluster larger than two nodes, a switched Fibre Channel fabric (FC-SW) is the only supported storage technology. In an FC-SW, devices are connected in a many-to-many topology using Fibre Channel switches (illustrated in Figure 7.16).

**Figure 7.16 FC-SW Connection**



When a node or device communicates with another node or device in an FC-SW, the source and target set up a point-to-point connection (similar to a virtual circuit) and communicate directly with each other. The fabric itself routes data from the source to the target. In an FC-SW, the media are not shared, any device can communicate with any other device, and communication occurs at full bus speed. This is a fully scalable enterprise solution and, as such, is highly recommended for deployment with server clusters.

FC-SW is the primary technology employed in SANs. Other advantages of FC-SW include ease of deployment, the ability to support millions of devices, and switches that provide fault isolation and rerouting. Also, there are no shared media as there are in FC-AL, allowing for faster communication. However, be aware that FC-SWs can be difficult for vendors to develop, and the switches can be expensive. Vendors also have to account for interoperability issues between components from different vendors or manufacturers.

# Using SANs with Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For any large-scale cluster deployment, it is recommended that you use a SAN for data storage. Smaller SCSI and stand-alone Fibre Channel storage devices work with server clusters, but SANs provide superior fault tolerance.

A SAN is a set of interconnected devices (such as disks and tapes) and servers that are connected to a common communication and data transfer infrastructure (FC-SW, in the case of Windows Server 2003 clusters). A SAN allows multiple server access to a pool of storage in which any server can potentially access any storage unit.

The information in this section provides an overview of using SAN technology with your Windows Server 2003 clusters. For more information about deploying server clusters on SANs, see the Windows Clustering: Storage Area Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Note**

- Vendors that provide SAN fabric components and software management tools have a wide range of tools for setting up, configuring, monitoring, and managing the SAN fabric. Contact your SAN vendor for details about your particular SAN solution.

The following sections provide an overview of SAN concepts that directly affect a server cluster deployment.

**HBAs**

Host bus adapters (HBAs) are the interface cards that connect a cluster node to a SAN, similar to the way that a network adapter connects a server to a typical Ethernet network. HBAs, however, are more difficult to configure than network adapters (unless the HBAs are preconfigured by the SAN vendor).

**Zoning and LUN Masking**

Zoning and LUN masking are fundamental to SAN deployments, particularly as they relate to a Windows Server 2003 cluster deployment. Both methods provide isolation and protection of server cluster data within a SAN, and in most deployments one or the other is sufficient. Work with your hardware vendor to determine whether zoning or LUN masking is more appropriate for your organization.

**Zoning**

Many devices and nodes can be attached to a SAN. With data stored in a single *cloud*, or storage entity, it is important to control which hosts have access to specific devices. Zoning allows administrators to partition devices in logical volumes and thereby reserve the devices in a volume for a server cluster. That means that all interactions between cluster nodes and devices in the logical storage volumes are isolated within the boundaries of the zone; other noncluster members of the SAN are not affected by cluster activity.

Figure 7.17 is a logical depiction of two SAN zones (Zone A and Zone B), each containing a storage controller (S1and S2, respectively).

**Figure 7.17 Zoning**



In this implementation, Node A and Node B can access data from the storage controller S1, but Node C cannot. Node C can access data from storage controller S2.

Zoning needs to be implemented at the hardware level (with the controller or switch) and not through software. The primary reason is that zoning is also a security mechanism for a SAN-based cluster, because unauthorized servers cannot access devices inside the zone (access control is implemented by the switches in the fabric, so a host adapter cannot gain access to a device for which it has not been configured). With software zoning, the cluster would be left unsecured if the software component failed.

In addition to providing cluster security, zoning also limits the traffic flow within a given SAN environment. Traffic between ports is routed only to segments of the fabric that are in the same zone.

**LUN masking**

A LUN is a logical disk defined within a SAN. Server clusters see LUNs and think they are physical disks. LUN masking, performed at the controller level, allows you to define relationships between LUNs and cluster nodes. Storage controllers usually provide the means for creating LUN-level access controls that allow access to a given LUN to one or more hosts. By providing this access control at the storage controller, the controller itself can enforce access policies to the devices.

LUN masking provides a more detailed level of security than zoning, because LUNs provide a means for zoning at the port level. For example, many SAN switches allow overlapping zones, which enables a storage controller to reside in multiple zones. Multiple clusters in multiple zones can share the data on those controllers. Figure 7.18 illustrates such a scenario.

**Figure 7.18 Storage Controller in Multiple Zones**




LUNs used by Cluster A can be masked, or hidden, from Cluster B so that only authorized users can access data on a shared storage controller.

**Requirements for Deploying SANs with Windows Server 2003 Clusters**

The following list highlights the deployment requirements to meet when using a SAN storage solution with your server cluster. For a white paper to provide more complete information about using SANs with server clusters, see the Windows Clustering: Storage Area Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

### Note

- A utility that might be helpful in deploying SANs with server clusters is Mountvol.exe. Mountvol.exe is included with the Windows Server 2003 operating system, and provides a means of creating volume mount points and linking volumes without requiring a drive letter. For more information about Mountvol.exe, see "Mountvol" in Help and Support Center for Windows Server 2003.

  - Each cluster on a SAN must be deployed in its own zone. The mechanism the cluster uses to protect access to the disks can have an adverse effect on other clusters that are in the same zone. By using zoning to separate the cluster traffic from other cluster or noncluster traffic, there is no chance of interference.

  - All HBAs in a single cluster must be the same type and have the same firmware version. Many storage and switch vendors require that all HBAs in the same zone — and, in some cases, the same fabric — share these characteristics.

  - All storage device drivers and HBA device drivers in a cluster must have the same software version.

  - Never allow multiple nodes access to the same storage devices unless they are in the same cluster.

**Caution**

If multiple nodes in different clusters can access a given disk, data corruption will result.

- Never put tape devices into the same zone as cluster disk storage devices. A tape device could misinterpret a bus rest and rewind at inappropriate times, such as during a large backup.

**Guidelines for Deploying SANs with Windows Server 2003 Clusters**

In addition to the SAN requirements discussed in the previous section, the following practices are highly recommended for server cluster deployment:

- In a highly available storage fabric, you need to deploy clustered servers with multiple HBAs. In these cases, always load the multipath driver software. If the I/O subsystem sees two HBAs, it assumes they are different buses and enumerates all the devices as though they were different devices on each bus. The host, meanwhile, is seeing multiple paths to the same disks. Failure to load the multipath driver will disable the second device, because the operating system sees what it thinks are two independent disks with the same signature.

- Hardware snapshots can go either to a server outside the server cluster or to a backup *clone* disk. When the original disk fails, you can replace it in the server cluster with its clone. When using cloned disks, it is very important that only one disk with a given disk signature be exposed to the server cluster at one time. Many controllers provide snapshots at the controller level that can be exposed to the cluster as a completely separate LUN. Cluster performance is degraded when multiple devices have the same signature. If a hardware snapshot is exposed back to the node with the original disk online, the I/O subsystem attempts to rewrite the signature. However, if the snapshot is exposed to another node in the cluster, the Cluster service does not recognize it as a different disk and the result could be data corruption. Although this is not specifically a SAN issue, the controllers that provide this functionality are typically deployed in a SAN environment.

  **Caution**

  - Exposing a hardware snapshot to a node on the cluster might result in data corruption.

# Choosing the Cluster Model

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The term *cluster model* refers to the manner in which the quorum resource is used in the server cluster. Server clusters require a quorum resource, which contains all of the configuration data necessary for recovery of the cluster. The cluster database, which resides in the Windows Server 2003 registry on each cluster node, contains information about all physical and logical elements in a cluster, including cluster objects, their properties, and configuration data. When a cluster node fails and then comes back online, the other nodes update the failed node's copy of the cluster database. It is the quorum resource that allows the Cluster service to keep every active node's database up to date.

The quorum resource, like any other resource, can be owned by only one node at a time. A node can form a cluster only if it can gain control of the quorum resource. Similarly, a node can join a cluster (or remain in an existing cluster) only if it can communicate with the node that controls the quorum resource.

Servers negotiate for ownership of the quorum resource in order to avoid split-brain scenarios, which occur when nodes lose communication with one another and multiple partitions of the cluster converge with a well-defined set of members, each partition believing that it is the one and only instance of the cluster. In such a case, each partition would continue operating in the belief that it had control of any cluster-wide shared resources or state, ultimately leading to data corruption.

The concept of a quorum ensures that only one partition of a cluster survives a split-brain scenario. When the cluster is partitioned, the quorum resource is used as an arbiter. The partition that owns the quorum resource is allowed to continue. The other partitions of the cluster are said to have *lost quorum*, and Cluster service — along with any resources hosted on the nodes that are not part of the partition that has quorum — is terminated.

The Cluster service stores cluster configuration data in a quorum log file. This file is usually located on a shared disk that all nodes in the cluster can access, and it acts as the definitive version of the cluster configuration. It holds cluster configuration information such as which servers are part of the cluster, what resources are installed in the cluster, and the state of those resources (for example, online or offline). Because the purpose of a cluster is to have multiple physical servers acting as a single virtual server, it is critical that each of the physical servers has a consistent view of how the cluster is configured.

A quorum resource can be any resource that:

- Provides a means for arbitration leading to membership and cluster state decisions.
- Provides physical storage to store configuration information.
- Uses NTFS.

The quorum resource you choose is dictated by the server cluster model you deploy (number of servers, number of sites, and so on). Do not build your server cluster around a particular type of quorum resource; instead, design the cluster that best supports your applications, and then choose one of the cluster models described in the following sections.

Three types of cluster models are available for deployment with server clusters:

- Local quorum
- Single quorum device
- Majority node set

For more information about the cluster database and the quorum resource, see "Quorum resource" in Help and Support Center for Windows Server 2003.

**Local Quorum Cluster**

This cluster model is for clusters that consist of only one node. This model is also referred to as a local quorum. It is typically used for:

- Deploying dynamic file shares on a single cluster node, to ease home directory deployment and administration.

- Testing.
- Development.

**Single Quorum Device Cluster**
The most widely used cluster model is the single quorum device cluster (also referred to as the standard quorum model), which is illustrated in Figure 7.19. The definitive copy of the cluster configuration data is on a single cluster storage device connected to all nodes.

**Figure 7.19 Single Quorum Device Cluster**




**Majority Node Set Cluster**
With Windows Server 2003, Microsoft introduces the majority node set cluster. The majority node set quorum model is intended for sophisticated, end-to-end clustering solutions. Each node maintains its own copy of the cluster configuration data. The quorum resource ensures that the cluster configuration data is kept consistent across the nodes. For this reason, majority node set quorums are typically found in geographically dispersed clusters. Another advantage of majority node set quorums is that a quorum disk can be taken offline for maintenance and the cluster will continue to operate.

The major difference between majority node set clusters and single quorum device clusters is that single quorum device clusters can survive with just one node, but majority node set clusters need to have a majority of the cluster nodes survive a failure for the server cluster to continue operating. A majority node set cluster is illustrated in Figure 7.20.

**Figure 7.20 Majority Node Set Cluster**




A majority node set cluster is a good solution for controlled, targeted scenarios as part of a cluster solution offered by an independent software vendor (ISV) or independent hardware vendor (IHV). By abstracting storage from the Cluster service, majority node set clusters provide ISVs with much greater flexibility for designing sophisticated cluster scenarios.

There are strict requirements to adhere to when you deploy majority node set clusters. The deployment of majority node set clusters is appropriate in some specialized configurations that need tightly consistent cluster features and do not have shared disks, for example:

- Clusters that host applications that can fail over, but that use other, application-specific methods to keep data consistent between nodes. Database log shipping and file replication for relatively static data are examples of this kind of application.
- Clusters that host applications that have no persistent data, but that need to cooperate in a tightly coupled manner to provide consistent volatile state.
- Multisite clusters, which span more than one physical site.

For more information about majority node sets and geographically dispersed clusters, see "Protecting Data from Failure or Disaster" later in this chapter.

# Designing the Cluster Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A network performs one of the following roles in a cluster:

- A private network carries internal cluster communication.
- A public network provides client systems with access to cluster application services.
- A mixed network (public and private) carries internal cluster communication and connects client systems to cluster application services.

Figure 7.21 shows a typical four-node server cluster deployment with a public (intranet) and private network.

**Figure 7.21 Four-Node Cluster Service Cluster with Private and Public Networks**




A network might have no role in a cluster; this is also known as disabling a network for cluster use. In the case of single quorum device clusters, the Cluster service will not use the network for internal traffic, nor will it bring IP Address resources online. Other cluster-related traffic, such as traffic to and from a domain controller for authentication, might or might not use this network.

Use these guidelines for designing the network segments of the server clusters:

- You can provide fault tolerance in your network by grouping network adapters on multiple ports to a single physical segment. This practice is known as *teaming network adapters* and is described in "Eliminating Single Points of Failure in Your Network Design" later in this chapter.
- You can use any network adapter that complies with Network Driver Interface Specification (NDIS).
- Use only network adapters that comply with specifications for the Peripheral Component Interconnect (PCI) local bus. Do not use network adapters compatible with an Industrial Standard Architecture (ISA) or Extended Industrial Standard Architecture (EISA) bus.
- For each server cluster network, use only identical network adapters on all nodes (that is, the same make, model, and firmware version). This is not a requirement, but it is strongly recommended. Consider the loads on each type of adapter and plan accordingly; for example, the private network adapters can require 10 MB to 100 MB of throughput, although the public network adapters can require 1 gigabyte (GB) of throughput.

# Eliminating Single Points of Failure in Your Network Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must incorporate failsafe measures within the cluster network design. To achieve this end, adhere to the following rules:

- Configure at least two of the cluster networks for internal cluster communication to eliminate a single point of failure. A server cluster whose nodes are connected by only one network is not supported.
- Design each cluster network so that it fails independently of other cluster networks. This requirement implies that the components that make up any two networks must be physically independent. For example, the use of a multiport network adapter to attach a node to two cluster networks does not meet this requirement if the ports share circuitry.
- If using multihoming nodes, make sure the adapters reside on separate subnets.
- Do not connect multiple adapters on one node to the same network; doing so does not provide fault tolerance or load balancing.

  **Note**

  - If a cluster node has multiple adapters attached to one network, the Cluster service will recognize only one network adapter per node per subnet.

The best way to guard against network failure in your server cluster is by teaming network adapters. By grouping network adapters on multiple ports to a single physical segment, you can provide fault tolerance to your cluster network. If a port fails — whether the failure occurs on the adapter, cable, switch port, or switch — another port takes over automatically. To the operating system and other devices on the network, this failover is transparent.

An important consideration, however, is that networks dedicated to internal server cluster communication cannot be teamed. Teaming network adapters is supported only on networks that are not dedicated to internal cluster traffic. Teaming network adapters on all cluster networks concurrently is not supported. There are alternative methods for achieving redundancy on networks that are dedicated to internal cluster traffic. These alternatives include adding a second private network, which avoids the cost of adding network adapters (however, this second network must be for intra-cluster network traffic only).

Teaming network adapters on other cluster networks is acceptable; however, if communication problems occur on a teamed network, it is recommended that you disable teaming.

# Determining Domain Controller Access for Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For Windows Server 2003 clusters to function properly, any node that forms part of the cluster must validate the Cluster service account in the local domain. To do this, each node must be able to establish a secure channel with a domain controller. If a domain controller cannot validate the account, the Cluster service does not start. This is also true for other clustered applications that need account validation, such as SQL Server and Microsoft® Exchange 2000. There are three ways to provide the necessary access, presented here in order of preference:

- Configure cluster nodes as member servers within a Windows domain and give them fast, reliable connectivity to a local domain controller.
- If the connectivity between cluster nodes and domain controllers is slow or unreliable, locate a domain controller within the cluster.
- Configure the cluster nodes as domain controllers, so that the Cluster service account can always be validated.

The security rights applicable to a domain controller administrator are often inappropriate or nonfunctional for a cluster administrator. A domain controller administrator can apply global policy settings to a domain controller that might conflict with the cluster role. Domain controller administrative roles span the network; this model generally does not suit the security model for clusters.

It is strongly recommended that you do not configure the cluster nodes as domain controllers, if at all possible. But if you must configure the cluster nodes as domain controllers, follow these guidelines:

- If one node in a two-node cluster is a domain controller, both nodes must be domain controllers. If you have more than two nodes, you must configure at least two nodes as domain controllers. This gives you failover assurance for the domain controller services.
- A domain controller that is idle can use between 130 MB and 140 MB of Random Access Memory (RAM), which includes running the Cluster service. If these domain controllers have to replicate with other domain controllers within the domain and across domains, replication traffic can saturate bandwidth and degrade overall performance.
- If the cluster nodes are the only domain controllers, they must be Domain Name System (DNS) servers as well. You must address the problem of not registering the private interface in DNS, especially if the interface is connected by a crossover cable (two-node cluster only). The DNS servers must support dynamic updates. You must configure each domain controller and cluster node with a primary — and at least one secondary — DNS server. Each DNS server needs to point to itself for primary DNS resolution and to the other DNS servers for secondary resolution.
- If the cluster nodes are the only domain controllers, they must also be global catalog servers.
- The first domain controller in the forest takes on all single-master operation roles. You can redistribute these roles to each node; however, if a node fails over, the single-master operation roles that the node has taken on are no longer available.
- If a domain controller is so busy that the Cluster service is unable to gain access to the quorum as needed, the Cluster service might interpret this as a resource failure and cause the resource group to fail over to the other node.
- Clustering other applications, such as SQL Server or Exchange 2000, in a scenario where the nodes are also domain controllers can result in poor performance and low availability due to resource constraints. Be sure to test this configuration thoroughly in a lab environment before you deploy it.
- You must promote a cluster node to a domain controller by using the Active Directory Installation Wizard before you create a cluster on that node or add it to an existing cluster.

- Practice extreme caution when demoting a domain controller that is also a cluster node. When a node is demoted from a domain controller, security settings are changed. For example, certain domain accounts and groups revert back to the default built-in accounts and groups originally on the sever. This means the security ID (SID) of the Domain Admins group changes to that of the local Administrators group. As a result, the Administrators entry in the security descriptor for the Cluster service is no longer valid because its SID still matches the Domain Admins group, and not the SID for the local Administrators group.

# Protecting Data from Failure or Disaster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must take steps to protect and ensure the integrity of your data. If your organization spans more than one physical site, for example, you might want to cluster resources across sites in what is known as a geographically dispersed cluster. Geographically dispersed clusters maintain high availability even when communication with another site is disrupted or lost. And all server clusters, whether in a single site or spread across two sites, need a plan for backing up application data regularly, and for recovering data in the event of a failure or disaster.

Figure 7.22 displays the process for ensuring the data integrity for your server cluster.

**Figure 7.22 Protecting Data from Failure or Disaster**



# Designing a Geographically Dispersed Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Geographically dispersed server clusters ensure that a complete outage at one site does not cause a loss of access to the application being hosted. All nodes hosting an application must exist within the same cluster. Therefore, to provide fault tolerance, a single cluster spans multiple sites.

**Important**

- Windows Server 2003 supports two-site geographically dispersed clusters. However, Microsoft does not provide a software mechanism for replicating application data from one site to another. Instead, Microsoft works with hardware and software vendors to provide a complete solution. For more information about qualified clustering solutions, see the Windows Server Catalog link or the Geographically Dispersed Clusters link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

A geographically dispersed cluster requires multiple storage arrays, at least one in each site, to ensure that in the event of failure at one site, the remaining site will have local copies of the data. In addition, the nodes of a geographic cluster are connected to storage in such a way that when there is a failure at one site or a failure in communication between sites, the functioning nodes can still connect to storage at their own site. For example, in the simple two-site cluster configuration shown in Figure 7.23, the nodes in Site A are directly connected to the storage in Site A, so they can access data with no dependencies on Site B.

**Figure 7.23 Geographically Dispersed Cluster**




In Figure 7.23, Nodes A and B are connected to one array in Storage Controller A, while Node C and Node D are connected to another array in Storage Controller B. These storage arrays present a single view of the disks spanning both arrays. In other words, Disks A and B are combined into a single logical device (by using mirroring, either at the controller level or the host level). Logically, the arrays appear to be a single device that can fail over between Nodes A, B, C, and D.

**Note**

- For a job aid to assist you in gathering your requirements for a multisite server cluster and organizing the necessary, preparatory information, see "Planning Checklist for Geographically Dispersed Clusters" (Sdcclu_1.doc) on the Microsoft® *Windows® Server 2003 Deployment Kit* companion CD (or see "Planning Checklist for Geographically Dispersed Clusters" on the Web at http://www.microsoft.com/reskit).

# Design Requirements for Geographically Dispersed Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although there are many requirements in setting up a geographically dispersed cluster, at the most fundamental level your design has to meet two requirements:

- Both sites must have independent copies of the same data. The goal of a geographically dispersed cluster is that if one site is lost, such as from a natural disaster, the applications can continue running at the other site. For read-only data, the challenge is relatively simple: the data can be cloned, and one instance can run at each site. However, in a typical stateful application, in which data is updated frequently, you must consider how changes made to the data at one site are replicated to the other site.

- When there is a failure at one site, the application must restart at the other site. Even if the application data is replicated to all sites, you need to know how the application is restarted at an alternate site when the site that was running the application fails.

**Important**

- The hardware and software configuration of a geographically dispersed cluster must be certified and listed in the Windows Server Catalog. Also, be sure to involve your hardware manufacturer in your design decisions. Frequently, third-party software and drivers are required for geographically dispersed clusters to function correctly. Microsoft Product Support Services might not be aware of how these components interact with Windows Clustering. For more information about hardware and software configurations, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Network Requirements for Geographically Dispersed Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following network requirements pertain to all server clusters, but are particularly important with regard to geographically dispersed clusters:

- All nodes must be on the same subnet. The nodes in the cluster can be on different physical networks, but the private and public network connections between cluster nodes must appear as a single, nonrouted Local Area Network (LAN). You can do this by creating virtual LANs (VLANs).

- Each VLAN must fail independently of all other cluster networks. The reason for this rule is the same as for regular LANs — to avoid a single point of failure.

- The roundtrip communication latency between any pair of nodes cannot exceed 500 milliseconds. If communication latency exceeds this limit, the Cluster service assumes the node has failed and will potentially fail over resources.

# Data Replication in Geographically Dispersed Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Data can be replicated between sites with different techniques, and at different levels in the clustering infrastructure. At the block level (known as *disk-device-level replication* or *disk-device-level mirroring*), replication is performed by the storage controllers or by mirroring the software. At the file-system level (replication of file system changes), the host software performs the replication. Finally, at the application level, the applications themselves can replicate data. An example of application-level replication is Microsoft SQL Server log shipping.

The method of data replication that is used depends on the requirements of the application and the business needs of the organization that owns the cluster.

**Synchronous vs. Asynchronous Replication**

When planning geographically dispersed clusters, you need to understand your data consistency needs in different failure and recovery scenarios and work with the solution vendors to meet your requirements. Different geographically dispersed cluster solutions provide different replication and redundancy strategies. Determine the support requirements of your applications with regard to replication -- in geographically dispersed server clusters, the type of data replication is just as important as the level at which it occurs. There are two types of data replication: synchronous and asynchronous.

*Synchronous replication* is when an application performs an operation on one node at one site, and then that operation is not completed until the change has been made on the other sites. Using synchronous, block-level replication as an example, if an application at Site A writes a block of data to a disk mirrored to Site B, the I/O operation will not be completed until the change has been made to both the disk on Site A and the disk on Site B. Because of this potential latency, synchronous replication can slow or otherwise detract from application performance for your users.

*Asynchronous replication* is when a change is made to the data on Site A and that change eventually makes it to Site B. In asynchronous replication, if an application at Site A writes a block of data to a disk mirrored to Site B, then the I/O operation is complete as soon as the change is made to the disk at Site A. The replication software transfers the change to Site B (in the background) and eventually makes that change to Site B. With asynchronous replication, the data at Site B can be out of date with respect to Site A at any point in time.

Different vendors implement asynchronous replication in different ways. Some preserve the order of operations and others do not. This is very important, because if a solution preserves ordering, then the disk at Site B might be out of date, but it will always represent a state that existed at Site A at some point in the past. This means Site B is *crash consistent*: the data at Site B represents the data that would exist at Site A if Site A had crashed at that point in time. Conversely, if a solution does not preserve ordering, the I/O operations might be applied at Site B in an arbitrary order. In this case, the data set at Site B might never have existed at Site A. Many applications can recover from crash-consistent states; very few can recover from out-of-order I/O operation sequences.

**Caution**

- Geographically dispersed server clusters must never use asynchronous replication unless the order of I/O operations is preserved. If this order is not preserved, the data that is replicated to the second site can appear corrupt to the application and be totally unusable.

Mirroring and replication solutions are implemented differently, depending on the vendor, the business needs of the organization, and the logistics of the cluster. In general terms, however, for every disk there is a master copy and one or more secondary copies. The master is modified, and then the changes are propagated to the secondary copies.

In the case of disk-device-level replication, the secondary disk might not be visible to the applications. If it is visible, it will be a read-only copy of the device. When there is a failover, a cluster resource typically designates one of the secondary disks to be the new primary, and the old primary becomes a secondary. In other words, most of the mirroring solutions are master-secondary, one-way mirror sets. This is usually on a per-disk basis, so some disks might have the master at one site and others might have the master at another site.

# Majority Node Set Quorum in Geographically Dispersed Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although geographically dispersed clusters can use a standard quorum, presenting the quorum as a single logical shared drive between sites can create design issues. Majority node set clusters solve these issues by allowing the quorum to be stored on the local hard disk of each node.

Applications in a geographically dispersed cluster are typically set up to fail over in the same manner as those in a single-site cluster; however, the total failover solution is inherently more complex. The Cluster service provides health monitoring and failure detection of the applications, the nodes, and the communications links, but there are cases where it cannot differentiate between various failure modes.

Consider two identical sites, each having the same number of nodes and running the same software. If a complete failure of all communication (both network and storage fabric) occurs between the sites, neither site can continue without human intervention because neither site has sufficient information about the other site's continuity.

As discussed in "Choosing the Cluster Model" earlier in this chapter, the server cluster architecture requires that a single-quorum resource in the cluster be used as the tiebreaker to avoid split-brain scenarios. Although split-brain scenarios can happen in single-site clusters, they are much more likely to occur in geographically dispersed clusters.

If communication between two sites in a geographically dispersed server cluster were to fail, none of the cluster nodes in either site could determine which of the following is true:

- Communication between sites failed and the other site is still alive.
- The other site is dead and no longer available to run applications.

This problem is solved when one of the partitions takes control of the quorum resource. However, with majority node set clusters, the process is different than in a traditional (single-site, single-quorum) server cluster. Traditional clusters can continue as long as one of the nodes owns the quorum disk, and traditional server clusters can therefore continue even if only one node out of the configured set of nodes is running (if that node owns the quorum disk). By contrast, a server cluster running with a majority node set quorum resource will start up (or continue running) only if a majority of the nodes configured for the cluster are running, and if all nodes in that majority can communicate with each other.

When designing your server cluster with a majority node set quorum, consider this difference, because it can affect how majority node set clusters behave when the server cluster is partitioned and during split-brain scenarios. This "more than half" quorum requirement makes the number of nodes in your server cluster very important. For more information about majority node set quorum cluster models, see "Model 3: Majority node set server cluster configuration" in Help and Support Center for Windows Server 2003.

# Single Quorum in Geographically Dispersed Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is possible to use a single quorum disk resource in a geographically dispersed cluster, provided the quorum is replicated across sites, such as through disk mirroring. When a single quorum is used across multiple sites, the quorum data must always use synchronous replication, although application data can be replicated either synchronously or asynchronously (assuming there are no application-specific restrictions). This is because the quorum could become corrupt during asynchronous replication, rendering different quorum data on each site, and your server cluster would then fail.

# Backing Up and Restoring Cluster Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order to maintain high availability on your server clusters, it is important to back up server cluster data on a regular basis. The Backup or Restore Wizard in Windows Server 2003 (Ntbackup.exe) has been enhanced to enable backups and restores of the local cluster database, in addition to the capacity to restore configurations locally and to all nodes in a cluster.

**Note**

- In addition to the Backup or Restore Wizard, Ntbackup.exe is also available as a command-line tool. For more information about Ntbackup.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Command-line reference A-Z.

Although there are some scenarios in which you can restore data without having a backup data set available, your options are limited. Providing true high availability to your users is possible only by backing up your data on a regular basis. For information about what you can do to restore your server cluster without a backup, see the Server Cluster (MSCS) Backup and Recovery Best Practices link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

A new backup and restore feature in Windows Server 2003 is Automated System Recovery (ASR), which helps you recover a system that will not start. ASR is a two-part recovery option consisting of ASR backup and ASR restore, both of which are accessed through the Advanced mode of the Backup or Restore Wizard. Note that ASR backs up only the partition used by the operating system. You must back up other partitions by using the standard Backup or Restore Wizard.

You can use ASR to recover a cluster node, the quorum disk, and the signatures on any lost shared disks. You can use ASR to restore a Windows Server 2003 installation when:

- A node cannot start as the result of damaged or missing Windows Server 2003 system files. ASR verifies and replaces damaged critical system files.
- There is a hardware failure, such as disk failure. ASR restores the installation to the state of your most recent ASR backup.
- The Cluster service will not start on the local node because the Cluster database is corrupt or missing.
- The disk signature has changed on one of the shared disks as a result of a disk replacement or other disk-related issue, causing a shared disk to fail to come online.

For more information about ASR, see "Automated System Recovery (ASR) overview" in Help and Support Center for Windows Server 2003.

For specific information about backing up cluster disk signatures and partition layouts, the cluster quorum, and data on cluster nodes, see "Backing up and restoring server clusters" in Help and Support Center for Windows Server 2003. Details include how to restore the cluster database on a local node, how to restore the contents of a cluster quorum disk for all nodes in a cluster, and how to repair a damaged cluster node by using ASR.

Help and Support Center for Windows Server 2003 also includes:

- The specific permissions that a cluster administrator needs in order to perform each backup and restore procedure. For example, to back up the data on cluster nodes or to restore a damaged cluster node using Automated System Recovery, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. For the specific access control permissions you need to perform each backup and restore operation, see "Backing up and restoring server clusters" in Help and Support Center for Windows Server 2003.
- Ten failover and restore scenarios that describe failure from cluster disk data loss, cluster quorum corruption, cluster disk corruption or failure, and complete cluster failure. Scenarios include symptoms that are common indications of a server cluster failure, and methods for recovering from such failures. For more information see "Backing up and restoring server clusters" in Help and Support Center for Windows Server 2003.

**Important**

- For information about backing up applications that run on your server cluster, consult the documentation for the application.

# Securing Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Security considerations must be present in all aspects of your cluster design. After your clustering infrastructure is in place, you need to consider a number of recommended policies and configurations to ensure that your server clusters are secure. Figure 7.24 shows an overview of security deployment considerations.

**Figure 7.24 Securing Your Server Clusters**




Before beginning your security design, see "Best practices for securing server clusters" in Help and Support Center for Windows Server 2003.

**Note**

- Do not weaken the default security settings for clustering objects in the operating system. This includes registry keys, files, and devices.

# Ensuring the Physical Security of Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For servers that must maintain high availability, it is recommended that you restrict physical access to a small number of essential, designated individuals. Locate server clusters and their storage in physically secure locations.

When you review your policies for keeping your computer room secure, be sure to include the following:

- Restrict access to the instruction sheets that you create, such as the directions for server recovery.
- Review your administration practices to ensure that only trusted, authorized personnel have access to restricted data, such as logs and backup and restore media.

# Ensuring a Secure Networking Infrastructure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Standard security devices — such as firewalls, network probes, and management tools to detect irregular traffic — must be in place before you put sensitive data on server clusters (or any other servers).

**Note**

- Although you can use IP security (IPSec) with server clusters, IPSec was not designed for use in failover situations. You can, however, use IPSec if your business need for secure connectivity outweighs your need for continuous client access during failover. For more information about IPSec and clusters, see article Q306677, "IPSec Is Not Designed for Failover," in the Microsoft Knowledge Base. To find this article, click the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

In order to create a secure environment for your server clusters, begin with the following basic network precautions:

- Secure your network with firewalls and management tools that can detect irregular network traffic.
- Restrict physical access to network hardware (routers, hubs, and switches) to protect them from unauthorized individuals.
- Make sure administrative permissions are adhered to, and that logs and other resources are protected by access control lists (ACLs).
- Ensure that network services such as the Active Directory, DNS, and DHCP, and Windows Internet Name Service (WINS) are protected from compromise. Any compromise of these infrastructure services can lead to a compromise of the Cluster service itself.

# Protecting Server Clusters from Network Flooding

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Cluster service uses User Data Protocol (UDP) port 3343 for intracluster communication. This communication includes heartbeat traffic, which detects node failure, and cluster control operations. Some of this communication (such as heartbeats) is time sensitive. If there is a significant load on the ports due to network flood attacks, the result can be false node-failure detection, which can cause unnecessary failover operations, leading to application downtime. To ensure that node-failure detection cannot be compromised in this way, secure the private network or networks between cluster nodes so that only trusted nodes can send data.

# Protecting Server Clusters from Unauthorized Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Cluster service provides remote management capabilities through the Cluster API so that you can manage the cluster from a management workstation. If you are administrating clusters remotely, be sure to use only trusted computers. The Cluster API uses the NTLM authentication protocol to authenticate the administrator on the cluster. This allows the cluster to verify that the administrator has sufficient administrative credentials to manage the cluster; however, it does not provide mutual authentication. In other words, the client has no guarantee that the node or cluster it is connected to is the real cluster.

If an unauthorized computer appears on the network with the same IP address or network name assigned to the cluster (assuming DNS information was compromised), the unauthorized computer can masquerade as the cluster. If that unauthorized computer also appears to implement the Cluster API, the unauthorized computer can intercept any management commands sent to the cluster. In many cases, this might not represent a threat because the administrator simply receives a false positive or negative that an action was performed. However, there are some sensitive operations (such as changing the cluster configuration) during which an unauthorized recipient could potentially collect cluster configuration data that might help to extend the attack surface.

This type of masquerade attack cannot occur until a number of conditions are met:

- The unauthorized computer must be visible on the same subnet as the target cluster and respond to traffic to the cluster IP address (this can be the IP address of the physical computers or the IP address of virtual servers hosted by the computer).
- In a typical environment, the unauthorized computer can take over the IP address only if the cluster node or virtual server is not running.
- The unauthorized computer must appear to implement the Cluster API.
- The administrator or an operational procedure must change the configuration that includes the sensitive data (for example, the administrator must change the Cluster service account password).

To protect against such attacks, the subnets on which the server clusters reside must be protected in the following ways:

- The subnet used by the cluster nodes must not extend beyond the set of nodes that can be physically secured or trusted. You must ensure that unauthorized or potentially compromised computers cannot attach to the subnet that contains the clusters. Because this network typically contains domain controllers, DNS servers, WINS servers, DHCP servers, and other network infrastructure, you must take steps to ensure that these infrastructure servers are also secure.
- Make sure that only nodes in the cluster are visible on a private network (multiple clusters can use the same private network). Make sure that no other network infrastructure servers or other application servers exist on the private subnet. This can be achieved by physically constraining the network itself (for example, ensuring that the private network is a LAN that is only connected to the cluster nodes), or by isolating the private networks by using VLAN-capable switches.

# Protecting the Cluster Disks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Clustering supports only NTFS on cluster disks. This ensures that file protection can be used to safeguard data on the cluster disks. Because the cluster disks can fail over between nodes, you must use only domain user accounts (or Local System, Network Service, or Local Service) to protect files. Local user accounts on one computer have no meaning on other computers in the cluster.

Cluster disks are periodically checked for health. The Cluster service account must have Write access to the top-level directory of all cluster disks. If the Cluster service account does not have Write access, the disk might be mistakenly declared as failed.

**Evaluating Upgrade Risks**

On the Microsoft® Windows® 2000 Server and Microsoft® Windows NT® version 4.0 operating systems, the default security attributes of the cluster log files in the *windir*\Cluster directory and the default security attributes of the quorum directory allow any authenticated user to read the contents. The security of these directories has been tightened in Windows Server 2003 to stop nonadministrator access altogether, ensuring that unauthorized users cannot gain information about the cluster configuration. However, these security attributes are not modified when you upgrade to Windows Server 2003, so on an upgraded Windows Server 2003 cluster node, all authenticated users might have Read access to these directories. Be sure to manually set the permissions to conform to the minimum access requirements.

### Note

- Be aware that if you install any service packs in the future, the service packs might reset permissions that have been configured manually.

**Protecting the Quorum Disk**

Never store other application data on the quorum disk. The quorum disk should contain only quorum data.

The quorum disk health determines the health of the entire cluster. If the quorum disk fails, the Cluster service becomes unavailable on all cluster nodes. The Cluster service checks the health of the quorum disk and negotiates for exclusive access to the physical drive by using standard I/O operations. These operations are queued to the device along with any other I/O operations to that device. If extremely heavy traffic delays the Cluster service I/O operations, the Cluster service declares the quorum disk as failed and forces a regroup event to bring the quorum back online somewhere else in the cluster.

To protect against malicious applications that could fill up the quorum disk, or flood the quorum disk with I/O operations, restrict access to the quorum disk to the local Administrators group on the local computer and the Cluster service account. If the quorum disk fills up, the Cluster service might be unable to log required data. In this case, the Cluster service fails, potentially on all cluster nodes.

**Protecting the Server Cluster Data Disks**

As it does with the quorum disk, the Cluster service periodically checks the health of the cluster data disks. If malicious applications flood the cluster data disks with I/O operations, the Cluster service health check can fail, thereby causing the disk (and any applications that depend on the disk) to fail over to another cluster node, which can result in a denial-of-service attack. To avoid this possibility, restrict access to the cluster data disks to only those applications that store data on the specific disks.

# Creating Secure Administrative Policies for Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To provide adequate and manageable security for sever clusters, you need to implement appropriate administrative practices and policies. The most important secure administrative policy is to permit only trusted, specially designated workstations to access the Cluster service. Any compromise on computers used to administer the cluster can compromise the cluster. For example, if a workstation from which administrative tools are run can be accessed by unauthorized users, these users can run unreliable or malicious code on the cluster without the knowledge of the cluster administrator. For this reason, it is important to audit both administrative access to the server cluster and changes to the Cluster service account.

The following sections describe best practices for securely administrating a server cluster.

**Note**

- To change the password of the Cluster service account without taking any nodes offline, it is essential that each node in the cluster have the same Cluster service account.

**Verifying Permissions for the Cluster Service Account**

The nodes in a server cluster use authenticated communication mechanisms to ensure that only valid members of the cluster can participate in intra-cluster protocols. Authentication of the communication is based on the Cluster service account. The Cluster service creates and maintains files, devices, registry keys, and other objects in the operating system. The default security setting of these objects ensures that unauthorized users cannot impact the cluster configuration or the applications running on the cluster. Making these security settings less restrictive can lead to the cluster being compromised and application data being corrupted.

For the server cluster to function properly, the Cluster service account must have certain permissions associated with it. During cluster installation, some of these rights are granted directly to the Cluster service account while others are inherited when the Cluster service account is made a member of the local Administrators group. The Cluster service account must have rights that allow it to perform the following actions:

- Act as part of the operating system.
- Back up files and directories.
- Adjust memory quotas for a process.
- Increase scheduling priority.
- Log on as a service.
- Restore files and directories.
- Debug programs.
- Manage auditing and security logs.
- Impersonate a client after authentication.

In Windows Server 2003, the Cluster service can publish virtual servers as computer objects in Active Directory. To ensure correct operation, the Cluster service account needs appropriate permissions to manipulate these objects in the Active Directory Computers container.

**Note**

- Although the Network Name resource publishes a computer object in Active Directory, that computer object must not be used for administrative tasks such as applying Group Policy settings. The only roles for the virtual server computer object are to support Kerberos authentication, and for cluster-aware services that can use Active Directory (such as Message

Queuing) to publish service provider information.

**Administering Clusters**

Cluster administrators can grant permissions to groups or individuals to manage the cluster. There is no fine level of control: Either a user has credentials to administer the cluster, or the user does not. Because of the high degree of impact administrators can have on your system's security, granting administrative credentials to a user must be done with careful consideration.

The security descriptor for the Cluster service contains the accounts that are authorized to administrate the cluster. By default, a cluster node's local Administrators group is added to the Cluster service security descriptor. The service accounts LocalSystem and NetworkService are also added to the security descriptor. The local Administrators group and the two accounts cannot be removed from the security descriptor. Be aware that adding a domain user or global group to the local Administrators group gives cluster administrator permissions to that group or account. If a node is evicted from a cluster, or if the last node is removed, the Cluster service account is not removed from the local Administrators group. When you remove a node from the cluster, you must manually remove the Cluster service account from the local Administrators group.

Cluster administrators can manage all aspects of the cluster configuration, including:

- Taking resources offline and bringing resources online.
- Adding and removing nodes from the cluster.
- Adding and removing resources from the cluster.

In addition, cluster administrators are able to shut down the Cluster service on nodes, provided they are also member of the local Administrators group. Apart from the local Administrators group on a node, all other members of the cluster security descriptor must be either domain user accounts, built-in local accounts such as System or NetworkService, or global groups. This ensures that the account is an identical, well-defined, and authorized account on all nodes in the cluster.

The Cluster service account does not need to be a member of the Domain Admins group, because the Cluster service account does not need domain administrator permissions. As a general security guideline, give all accounts the minimal possible permissions.

If you are deploying multiple clusters in a single domain, you can make administration easier by using the same Cluster service account on all nodes. However, you must balance ease of management against the potential security risks associated with using a single account for many clusters. If the account is compromised, the scope of the impact might exceed the benefits you gain by increasing the ease of management. With Windows Server 2003, the Cluster service account password on multiple clusters can be changed at the same time, as long as every node in the cluster is using the same Cluster service account.

Cluster administrators and the Cluster service need to use different accounts to administer the cluster. This allows more specific auditing and allows policy settings (such as password expiration) to be applied to the Cluster service account and the accounts used to administer the cluster.

If you plan to deploy multiple clusters that have different Cluster service accounts, create a global group or universal group that implements all the policy settings described earlier in this section. Then place each Cluster service account into the group. This eases management of the Cluster service accounts by providing a single container for all Cluster service accounts, and a single point of management for changing account policy settings. For example, you could put all cluster nodes and the Cluster service accounts into a single organizational unit (OU) in Active Directory.

Your OU model depends on your Active Directory implementation. If your Windows Server 2003 clusters reside in a Windows NT 4.0 domain, OUs and universal groups are not available.

Follow these guidelines for ease of use and best security:

- If there are multiple clusters in a single domain, use the same Cluster service account on all nodes to make administration easier.
- If you have password expiration policy settings on your Cluster service account, do not use this account for other services.
- Do not use the Cluster service account for SQL Server or Exchange 2000 if you have password expiration policy settings. When multiple services use the same account, coordinating the password change across the Cluster service and other

services is complex and can cause the entire cluster or service to become unavailable during a password rotation. It is recommended that you use a dedicated account, which can be maintained independently, for each service.

- With Windows Server 2003, the Cluster service account password can be changed online without taking down the cluster, but only if the Cluster service account is not used by other services.

**Caution**

- Be sure to change the Cluster service account password before it expires. The cluster will stop functioning when the password expires, because intra-cluster communication can no longer be successfully authenticated.

# Applying Kerberos Authentication in a Clustered Environment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Carefully plan your use of Kerberos authentication in a server cluster. By default, Kerberos authentication support for the network name resource is turned off. When Kerberos authentication support is enabled, the Cluster service account must be able to create a virtual computer object in Active Directory.

By default, all users have **Add workstations to the domain** permissions,which allows the creation of computer objects in Active Directory.

**Note**

- By default, authenticated users can join up to ten machine accounts to the domain. This limit does not apply to members of the Administrators or Domain Admins groups, and to those users who have delegated permissions on containers in Active Directory to create and delete computer accounts.

If the **Add workstations to the domain** permission has been removed from the Cluster service account, before enabling Kerberos authentication, a member of the Domain Administrators group must perform one of the following actions:

- Grant the Cluster service account the **Create Computer Objects** permission on the Computers organizational unit in Active Directory.

- Create the computer object manually in Active Directory before enabling Kerberos authentication. If the object is manually created, the Cluster service account must have **Write all properties** access permission to allow it to manipulate the computer object.

**Caution**

- Although the Network Name resource supports the changing of its **Name** property, many services, such as Message Queuing and Microsoft® SQL Server™ 2000, do not support changing the Network Name resource **Name** property. Do not change this property unless you fully understand the implications of doing so. In some cases, changing the Network Name resource **Name** property can lead to loss of data or service failure.

Before you allow clients to access a server cluster, or before failing over the resources to another node, ensure that the domain controllers have replicated newly created computer objects associated with Network Name resources. Until replication is complete, clients might fail to authenticate, or they could authenticate with the default NTLM authentication protocol and not with Kerberos authentication. One way to avoid replication issues with your clients is to not notify clients that the service is available until you are sure replication is complete. The amount of time it takes your directory to replicate information depends on many factors, such as the topology and the amount of network traffic. If you are not sure of the time necessary to allow for replication, replication can be forced with tools such as Ntdsutil.exe. For more information about Ntdsutil.exe, see "Using Ntdsutil" in Help and Support Center for Windows Server 2003.

After you enable Kerberos authentication on a server cluster, do not disable Kerberos authentication on a virtual server without knowing the effects the disabling will have on other services that use that virtual server. Microsoft® Message Queuing (MSMQ), for example, relies on the presence of a virtual computer object and ceases to function if its dependent Network Name resource has Kerberos authentication support disabled.

Use extreme care when you remove Kerberos authentication from a Network Name resource. When you disable Kerberos authentication from a Network Name resource, the computer object is disabled, leaving the system administrator with an explicit decision to delete the computer object. If the object remains in Active Directory, the Network Name resource does not go online. If the computer object is deleted, properties attached by applications that can use Active Directory are also deleted, and the applications might no longer function correctly.

For more information about enabling Kerberos authentication on a virtual server, see "Enable Kerberos authentication for virtual servers" in Help and Support Center for Windows Server 2003, or see articles Q302389, "Description of the Properties of the Cluster Network Name Resource in Windows Server 2003," and Q307532, "How to Troubleshoot the Cluster Service Account When It Modifies Computer Objects," in the Microsoft Knowledge Base. To find these articles, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Deploying Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After your application has been evaluated for server cluster deployment, and your hardware and network are in place, you can deploy your server cluster on Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. Figure 7.25 illustrates the process for installing or upgrading a server cluster.

**Figure 7.25 Deploying a Server Cluster**




**Important**

- Server clusters have very specific hardware requirements. It is extremely important that you select hardware configurations approved by Microsoft for Windows server clusters. For more information about approved hardware configurations, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Installing a New Server Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section is designed as an overview of the new server cluster installation process and is not intended as an installation guide. For detailed information about creating server clusters, see "Checklists: Creating Server Clusters" in Help and Support Center for Windows Server 2003.

Before installing Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition on your cluster nodes, consult "Planning for Deployments" in *Planning, Testing, and Piloting Deployment Projects* in this kit. Information in that chapter includes methods for taking inventory of your current IT environment and how to create a functional specification that you can use to create a clear and thorough cluster deployment plan. In addition, gather the following materials and have them available for reference during installation:

- A list of all services and applications to be deployed on server clusters.
- A plan that defines which applications are to be installed on which nodes.
- Failover policies for each service or application, including resource group planning.
- A selected quorum model.
- A physical and logical security plan for the cluster.
- Specifications for capacity requirements.
- Documentation for your storage system.
- The Windows Server Catalog — approved device drivers for network hardware and storage systems. For more information about support for server clusters, see article Q309395, "The Microsoft Support Policy for Server Clusters and the Hardware Compatibility List." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- Documentation supplied with all cluster hardware and all applications or services that will be deployed on the server cluster.
- An IP addressing scheme for the cluster networks, both private and public.
- A selected cluster name, its length limited by NetBIOS parameters.

Before installing, it is recommended that you complete the checklists in "Installing and upgrading on cluster nodes" in Help and Support Center for Windows Server 2003, Datacenter Edition and Windows Server 2003, Enterprise Edition.

Completing these checklists and reviewing the recommended information in Help and Support Center for Windows Server 2003, Datacenter Edition and Windows Server 2003, Enterprise Edition ensures that your hardware and network components are compatible and supported by Microsoft.

Additionally, be sure to read "Best practices for installing and upgrading cluster nodes" and the conceptual topics under "Installing Windows Server 2003, Enterprise Edition on cluster nodes" in Help and Support Center for Windows Server 2003, Enterprise Edition.

# Upgrading a Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are upgrading an existing server cluster from the Microsoft® Windows NT Server 4.0, Enterprise Edition or Windows® 2000 Advanced Server operating system, you have a choice of the following:

- Take the servers offline, upgrade the operating system, and rebuild the server cluster.
- Perform a rolling upgrade.

A *rolling upgrade* is the process of upgrading cluster nodes by turns while the other nodes continue to provide service. The server cluster downtime is reduced to a few minutes, which is the time needed to move resources from one node to another. Note that you cannot perform a rolling upgrade directly from Windows NT 4.0 to Windows Server 2003, but you can upgrade to Windows 2000 Advanced Server and then upgrade to Windows Server 2003, Enterprise Edition. Rolling upgrades are described in "Rolling Upgrades" later in this chapter.

**Important**

- Consult with the software vendors, or your software documentation, to be certain that all applications and services currently running on your server clusters will be able to function on Windows Server 2003.

# Upgrading by Rebuilding a Server Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process of upgrading a server cluster by taking the nodes offline, upgrading the operating system, and then rebuilding the cluster requires the same preparation as described in "Installing a New Server Cluster" earlier in this chapter.

**Checklists for Upgrading a Server Cluster**
Before upgrading a Windows NT Server 4.0, Enterprise Edition server cluster or a Windows 2000 Advanced Server cluster, it is recommended that you complete the checklist "Preparation for upgrading a cluster" in Help and Support Center for Windows Server 2003, Datacenter Edition and Windows Server 2003, Enterprise Edition.

- Additionally, be sure to read the relevant conceptual topics under "Upgrading to Windows Server 2003, Enterprise Edition on cluster nodes" in Help and Support Center for Windows Server 2003, Enterprise Edition and "Upgrading to Windows Server 2003, Datacenter Edition on cluster nodes" in Help and Support Center for Windows Server 2003, Enterprise Edition. Also available is information specific to Windows NT 4.0 server clusters running the Microsoft Distributed Transaction Coordinator (DTC) and MSMQ applications, and Windows NT 4.0 server clusters running Internet Information Services (IIS).

- For information about upgrading clusters from Windows 2000 Server, see "Upgrade a cluster from Windows 2000" in Help and Support Center for Windows Server 2003, Enterprise Edition. In Help and Support Center for Windows Server 2003, Datacenter Edition, search for this title to view individual topics about how to upgrade clusters containing IIS and clusters containing Message Queueing.

**Upgrading from Windows NT Server 4.0, Enterprise Edition**
There are extra considerations when upgrading directly to Windows Server 2003, Enterprise Edition from Windows NT Server 4.0, Enterprise Edition. You must be running Windows NT Server 4.0, Enterprise Server, Service Pack 5 or later, and you have to make sure that your server cluster hardware is compatible with Windows Server 2003 (and has enough processing power and memory). Check the Windows Server Catalog to be certain Microsoft will support your Windows NT 4.0 Server cluster after you upgrade to the new operating system.

It is also important to check with your software and service vendors, or the supporting documentation, to be certain that your services and applications still function after you upgrade.

For information about upgrading and reinstating your cluster nodes, see "Upgrade a cluster from Windows NT 4.0" in Help and Support Center for Windows Server 2003, Enterprise Edition. Note that by upgrading directly to Windows Server 2003, there is no way for you to provide service to your users until the upgrade of the entire cluster is complete. For more information about supported hardware for server clusters, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

# Rolling Upgrades

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Rolling upgrades allow you to upgrade your server clusters while continuing to provide service to your users. Rolling upgrades can be performed for the following scenarios only:

- From Windows 2000 Advanced Server to Windows Server 2003, Enterprise Edition.
- From the Microsoft® Windows® 2000 Datacenter Server operating system to Windows Server 2003, Datacenter Edition.

You cannot perform a rolling upgrade directly from Windows NT 4.0 to Windows Server 2003. To upgrade from Windows NT 4.0, you can perform a rolling upgrade to Windows 2000 and then perform another rolling upgrade to Windows Server 2003. Consult your Windows 2000 documentation for the best way to upgrade your server cluster to Windows 2000 Advanced Server.

As with a standard (offline) upgrade, before performing a rolling upgrade you have to check the Windows Server Catalog to make sure your cluster hardware is supported, and then you have to confirm that your clustered applications and services will be able to run in a Windows Server 2003 cluster. For more information about cluster hardware support, see the Windows Server Catalog link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

For detailed procedures for performing a rolling upgrade, see "Perform a rolling upgrade from Windows 2000" in Help and Support Center for Windows Server 2003. Also available is specific information for upgrading Windows 2000 clusters running IIS and MSMQ, and for upgrading a print spooler resource.

During the process of a rolling upgrade, some of your nodes will be upgraded to Windows Server 2003 while others will still be running Windows 2000, temporarily resulting in a mixed-version cluster. Although mixed version clusters are supported by Windows Server 2003, Datacenter Edition, this is not a desirable configuration. You should upgrade all nodes to the same operating system as soon as possible.

The operation of a mixed-version cluster is complicated if a resource type that you add to the cluster is supported in one version of the operating system but not in the other. For example, the Cluster service in Windows Server 2003, Datacenter Edition supports the Generic Script resource type. However, older versions of the Cluster service do not support it. A mixed-version cluster can run a Generic Script resource on a node running Windows Server 2003, Datacenter Edition, but not on a node running Windows 2000. The Cluster service transparently sets the possible owners of new resource types to prevent these resources from failing over to a Windows 2000 node of a mixed-version cluster. Consequently, when you view the possible owners of a new resource type, a Windows 2000 node will not be in the list, and you will not be able to add this node to the list. If you create such a resource during the mixed-version phase of a rolling upgrade, the resource groups containing those resources will not fail over to a Windows 2000 node.

For information about which services support rolling upgrades, see "Resource behavior during rolling upgrades" in Help and Support Center for Windows Server 2003. Also, before you perform a rolling upgrade, read Relnotes.htm in the \Docs folder of the Windows Server 2003, Enterprise Edition operating system CD, in addition to the product documentation that comes with each application or resource.

If you have a service that is not fully supported during rolling upgrade, you can still perform a last node rolling upgrade. For more information, see "Perform a last node rolling upgrade from Windows 2000" in Help and Support Center for Windows Server 2003. This process entails moving all the applications and resources on your server cluster that do not support rolling upgrades to a single Windows 2000 node. You then upgrade the remaining nodes in the cluster, transfer the applications and resources to the newly upgraded nodes, and upgrade the final node to Windows Server 2003.

# Testing and Deploying the Cluster Configuration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You must always test a server cluster configuration prior to live deployment. Ensure that your failover policies work and that the receiving nodes have the required capacity for all failover resources. Perform a benchmark on each server upon initial installation, and monitor the servers to ensure they are staying within operational guidelines. Collect performance data at peak and off-peak periods.

For information about ensuring that your server cluster is ready for a test environment, see "Test Cluster Failures and Failover Policies" in Help and Support Center for Windows Server 2003.

After verifying that your servers are running within acceptable tolerances and that failover is properly achieved, pilot test your server clusters in a production environment. A small pilot deployment in a production environment lets you test the effectiveness of your clusters in an environment that simulates the conditions of a full deployment. This allows you to make any necessary adjustments to your cluster design before performing the full deployment.

Complete information about testing and piloting Windows Server 2003 deployments is available in "Designing a Test Environment" and "Designing a Pilot Project" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

# Additional Resources for Server Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Related Information**

- "Designing Network Load Balancing" and "Designing Network Load Balancing" in this book for more information about designing and deploying Network Load Balancing clusters.
- "Deploying DHCP" in *Deploying Network Services* of this kit for information about deploying clustered DHCP servers.
- "Deploying WINS" in *Deploying Network Services* of this kit for information about deploying clustered WINS servers.
- The SQL Server 2000 Failover Clustering link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about deploying clustering solutions with SQL Server 2000.
- The Deploying Microsoft Exchange 2000 Server Service Pack 2 Clusters link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about using Exchange 2000 Server clusters running on Windows 2000 Server.
- Article 814607, "Microsoft Support for Server Clusters with 3rd Party System Components." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources .
- Article 309395, "The Microsoft Support Policy for Server Clusters and the Hardware Compatibility List." To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources .

**Related Tools**

- DiskPart.exe

  You can use DiskPart.exe to manage objects (disks, partitions, or volumes) by using scripts or direct input from a command prompt. For more information about DiskPart.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Command-line reference A-Z.

- Mountvol.exe

  You can use this command-line utility to create, delete, or list a volume mount point. Mountvol.exe is a way to link volumes without requiring a drive letter. For more information, see "Mountvol" in Help and Support Center for Windows Server 2003.

- Ntbackup.exe

  You can use this backup tool to perform backup operations at a command prompt or from a batch file using the ntbackup command followed by various parameters. For more information about Ntbackup.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Command-line reference A-Z.

- Clusdiag.exe

  You can use the Cluster Diagnostics and Verification Tool (ClusDiag.exe) to run diagnostic tests. These tests will help you verify the functionality of a cluster and troubleshoot cluster failures using the generated log files. For more information about Clusdiag.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Command-line reference A-Z.

**Related Job Aids**

- "Planning Checklist for Geographically Dispersed Clusters" (Sdcclu_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "Planning Checklist for Geographically Dispersed Clusters" on the Web at http://www.microsoft.com/reskit).

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- "Server Cluster Components" and "Cluster Objects" in Help and Support Center for Windows Server 2003 for an overview of server clusters.
- "Choosing applications to run on a server cluster" in Help and Support Center for Windows Server 2003.
- "Determining failover and move policies for groups" in Help and Support Center for Windows Server 2003.
- "Storage configuration options" in Help and Support Center for Windows Server 2003.
- "Mountvol" in Help and Support Center for Windows Server 2003.
- "Quorum resource" in Help and Support Center for Windows Server 2003.
- "Model 3: Majority node set server cluster configuration" in Help and Support Center for Windows Server 2003.
- "Automated System Recovery (ASR) overview" in Help and Support Center for Windows Server 2003.
- "Backing up and restoring server clusters" in Help and Support Center for Windows Server 2003.
- "Best practices for securing server clusters" in Help and Support Center for Windows Server 2003.
- "Enable Kerberos authentication for virtual servers" in Help and Support Center for Windows Server 2003.
- "Using Ntdsutil" in Help and Support Center for Windows Server 2003.
- "Best practices for installing and upgrading cluster nodes" and the conceptual topics under "Installing Windows Server 2003, Enterprise Edition on cluster nodes" in Help and Support Center for Windows Server 2003, Enterprise Edition
- "Installing Windows Server 2003, Enterprise Edition on cluster nodes" in Help and Support Center for Windows Server 2003, Enterprise Edition
- "Upgrading to Windows Server 2003, Enterprise Edition on cluster nodes" in Help and Support Center for Windows Server 2003, Enterprise Edition.
- "Checklists: Creating Server Clusters" in Help and Support Center for Windows Server 2003.
- "Upgrade a cluster from Windows 2000" in Help and Support Center for Windows Server 2003, Enterprise Edition. In Help and Support Center for Windows Server 2003, Datacenter Edition, search for this title to view individual topics telling how to upgrade clusters containing IIS and clusters containing Message Queuing.
- "Upgrade a cluster from Windows NT 4.0" in Help and Support Center for Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.
- "Perform a rolling upgrade from Windows 2000" in Help and Support Center for Windows Server 2003, Enterprise

Edition and Windows Server 2003, Datacenter Edition.

- "Resource behavior during rolling upgrades" in Help and Support Center for Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.

- "Perform a last node rolling upgrade from Windows 2000" in Help and Support Center for Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.

- "Test Cluster Failures and Failover Policies" in Help and Support Center for Windows Server 2003.

# Designing Network Load Balancing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many of your deployments will include mission-critical applications and services. The servers that host your applications and services must be able to support projected increases in the number of users, and they must ensure that users can access mission-critical applications. To fulfill these requirements, your solution must be highly available and scalable. Network Load Balancing (NLB) improves availability and scalability in your solutions by distributing application load across multiple servers.

**In This Chapter**

Overview of the NLB Design Process

Identifying Applications That Benefit from NLB

Determining the Core Specifications in the NLB Design

Securing NLB Solutions

Scaling NLB Solutions

Ensuring Availability in NLB Solutions

Additional Resources for Designing Network Load Balancing

**Related Information**

- For information about deploying Network Load Balancing, see "Deploying Network Load Balancing" in this book.
- For information about server clusters, see "Designing and Deploying Server Clusters" in this book.

# Overview of the NLB Design Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Improving availability and scalability in your solution depends on the applications and services in your organization. A computer running the Microsoft® Windows® Server 2003 operating system can provide a high level of reliability and scalable performance. However, a Network Load Balancing cluster can achieve the higher levels of availability and performance required by mission-critical servers.

A Network Load Balancing cluster comprises multiple servers running any version of the Windows Server 2003 family, including the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; and Windows® Server 2003, Web Edition operating systems. The servers are combined to provide greater scalability and availability than is possible with an individual server. Network Load Balancing distributes client requests across the servers to improve scalability. If a server fails, client requests are redistributed to the remaining servers to improve availability. Network Load Balancing can improve scalability and availability for applications and services that communicate with clients that use Transmission Control Protocol (TCP) or User Datagram Protocol (UDP).

The Network Load Balancing design process assumes that you are creating new clusters or that you are redesigning existing Windows Load Balancing Service (WLBS) or Network Load Balancing clusters. Upon completion of the Network Load Balancing design process, you will have a solution that meets or exceeds your scalability and availability requirements.

**Note**

- For a list of additional information to assist you in designing Network Load Balancing clusters, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

# NLB Design Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Creating a Network Load Balancing design involves more than documenting Windows Server 2003 and Network Load Balancing settings on the individual application servers. You must identify the applications and services that can benefit from Network Load Balancing, determine the core set of specifications in your design, ensure that your solution is secure, and provide for scalability and availability. The process for creating your Network Load Balancing design is shown in Figure 8.1.

**Figure 8.1 Designing a Network Load Balancing Solution**




As you create your Network Load Balancing design, document your decisions and use that information to deploy your Network Load Balancing solution.

**Note**

- For a Word document to assist you in documenting your Network Load Balancing design decisions, see "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Microsoft® Windows® Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

# NLB Fundamentals

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To create a successful Network Load Balancing design and to ensure that Network Load Balancing is correct for your solution, you need to know the fundamentals of how Network Load Balancing provides improved scalability and availability, and how Network Load Balancing compares with other strategies for providing scalability and availability.

**How NLB Provides Improved Scalability and Availability**

Network Load Balancing improves scalability and availability by distributing client traffic across the servers that you include in the Network Load Balancing cluster. Each *cluster host* (a server running on a cluster) runs an instance of the applications supported by your cluster. Network Load Balancing transparently distributes client requests among the cluster hosts. Clients access your cluster by using one or more virtual IP addresses. From the perspective of the client, the cluster appears to be a single server that answers the client request.

As the scalability and availability requirements of your solution change, you can add or remove servers from the cluster as necessary. Network Load Balancing automatically distributes client traffic to take advantage of any servers that you add to the cluster. In addition, when you remove a server from the cluster, Network Load Balancing redistributes the client traffic among the remaining servers in the cluster.

As an example, assume that your organization has a Web application farm running Microsoft® Internet Information Services (IIS) version 6.0 that hosts your organization's Internet presence. As seen in Figure 8.2, Network Load Balancing allows your individual Web application servers to service client requests from the Internet by distributing them across the cluster. On each of the servers, you install IIS 6.0 and Network Load Balancing. By combining the individual Web application servers into a Network Load Balancing cluster, you can load balance the requests to improve client response times and to provide improved fault tolerance in the event that one of the Web application servers fails.

**Figure 8.2 Network Load Balancing Cluster in a Web Farm**




Network Load Balancing automatically detects and recovers when the entire server fails or is manually disconnected from the network. However, Network Load Balancing is unaware of the applications and services running on the cluster, and it does not detect failed applications or services. To provide awareness of application or service failures, you need to add management software, such as Microsoft® Operations Manager (MOM), Microsoft® Application Center 2000, a third-part party application, or software developed by your organization.

When your design requires fault tolerance for servers that support your Network Load Balancing cluster, such as servers running Microsoft® SQL Server™ 2000, include Microsoft server clusters. For example, you can improve the availability of the network database (SQLCLSTR-01 in Figure 8.2) by creating a two-node server cluster. For more information about server

clusters, see "Designing and Deploying Server Clusters" in this book.

Network Load Balancing runs as an intermediate network driver in the Windows Server 2003 network architecture. Network Load Balancing is logically situated beneath higher-level application protocols, such as Hypertext Transfer Protocol (HTTP) and File Transfer Protocol (FTP), and above the network adapter drivers. Figure 8.3 illustrates the relationship of Network Load Balancing in the Windows Server 2003 network architecture.

**Figure 8.3 Network Load Balancing in the Windows Server 2003 Network Architecture**




To maximize throughput and to provide high availability in your solution, Network Load Balancing uses a distributed software architecture. A copy of the Network Load Balancing driver runs on each host in the cluster. The Network Load Balancing drivers allow all hosts in the cluster to concurrently receive incoming network traffic for the cluster.

On each host in the cluster, the driver acts as an intermediary between the network adapter driver and the TCP/IP stack. This allows a subset of the incoming network traffic to be received by the host. Network Load Balancing uses this filtering mechanism to distribute incoming client requests among the servers in the cluster.

Network Load Balancing architecture maximizes throughput by using a common media access control (MAC) address to deliver incoming network traffic to all hosts in the cluster. As a result, there is no need to route incoming packets to the individual hosts in the cluster. Because filtering unwanted network traffic is faster than routing packets (which involves receiving, examining, rewriting, and resending), Network Load Balancing delivers higher network throughput than dispatcher-based software load balancing solutions. Also, as you add hosts to your Network Load Balancing cluster, the scalability grows proportionally, and any dependence on a particular host diminishes.

Because Network Load Balancing load balances client traffic across multiple servers, it provides higher availability in your solution. One or more cluster hosts can fail, but the cluster continues to service client requests as long as any cluster hosts are running.

**NLB and Round Robin DNS**

Round robin Domain Name System (DNS) is a software method for distributing workload among multiple servers, but does not prevent clients from detecting server outages. If one of the servers fails, round robin DNS continues sending client requests to the server until a network administrator detects the failure and removes the server from the DNS address list. This results in service disruption for clients.

In contrast, Network Load Balancing automatically detects servers that have been disconnected from the cluster and redistributes client requests to the remaining servers. Unlike round robin DNS, this prevents clients from sending requests to the failed servers.

For more information about methods for improving availability and scalability, see "Planning for High Availability and Scalability" in this book.

# Identifying Applications That Benefit from NLB

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in creating your Network Load Balancing design, as illustrated in Figure 8.4, is to determine if any of your applications can benefit from the improved scalability and availability provided by Network Load Balancing. Your applications include not only applications in the traditional sense, such as Web applications, but also Windows Server 2003 network services, such as virtual private network (VPN) remote access.

**Figure 8.4 Identifying Applications That Benefit from Network Load Balancing**




Beyond using TCP/IP, an application that works well with Network Load Balancing has one or more of the following characteristics:

- Avoids instructing the client to open a subsequent connection that must be serviced by the same cluster host.

  For example, a Web application sends a response instructing a client to open an additional Secure Sockets Layer (SSL) session, and the application expects the SSL session to be established with the same cluster host that originated the response.

- Avoids retaining application state on a specific server in the cluster, such as session state saved by Active Server Pages (ASP).

  For example, some Web applications use a common database, or service, to maintain session state. A cookie is sent to the client on the initial request. On subsequent requests, the client sends the cookie to the cluster. Any cluster host in the cluster is capable of restoring the application session state by using the cookie sent by the client to retrieve the session state from the session database, or service. Applications that use a common database, or service, in conjunction with client-side cookies to maintain session state are more likely to work with Network Load Balancing than those that retain application state on a specific host in the cluster.

Applications and services that are stateless are the most likely to run without modifications on Network Load Balancing. Applications and services are said to be *stateless* when each client request is considered to be a new request that is unrelated to any previous request. Anytime a client request is considered to be a continuation of a previous request, the applications and services maintain *session state*.

Applications that maintain session state are said to be *stateful*. Stateful applications require *affinity* with the cluster host, when the session state is maintained locally.

You can set affinity between the client computer and a specific server within your cluster by using the port rules in Network Load Balancing. For more information about port rules, see "Identifying Applications or Services That Require Custom Port Rules" later in this chapter.

Regardless of which category your application belongs to, always test the application with Network Load Balancing in your lab environment during the proof-of-concept stage in your design. Ensure that your application is compatible with Network Load

Balancing before continuing with the design process.

Some of the common applications and services that run on Network Load Balancing include:

- Web applications running on IIS 6.0

  One of the most common of the solutions that use Network Load Balancing is an IIS 6.0 Web farm. A typical challenge in supporting Web applications occurs when an application must maintain a persistent connection to a specific cluster host. For example, if a Web application uses Hypertext Transfer Protocol Secure (HTTPS), the application should contact the same cluster hosts within the cluster, for efficiency. Connecting to a different cluster host requires establishing a new SSL session, which creates excess network traffic and overhead on the client and server. Network Load Balancing maintains affinity and reduces the possibility that a new SSL session needs to be established.

- VPN remote access running on Routing and Remote Access

  Another solution that uses Network Load Balancing involves using the Routing and Remote Access service in Windows Server 2003 to provide VPN remote connectivity. In the VPN solution, you combine multiple remote access servers running Windows Server 2003 and Routing and Remote Access to create a VPN remote access server farm.

- Web content caching and firewall running on Microsoft® Internet Security and Acceleration (ISA) Server 2000

  You can also use Network Load Balancing in solutions that include ISA Server to provide network security, network isolation, network address translation, or Web content caching. In ISA Server solutions, the design and deployment are integral parts of the ISA Server design and deployment process.

  For more information about creating ISA Server designs and deploying ISA Server in your organization, see "Deploying ISA Server" in *Deploying Network Services* of this kit and see the documentation that accompanies ISA Server.

- Application hosted on Terminal Services

  When you run applications on Terminal Services, the Terminal Services clients can be load balanced across a number of computers running Terminal Services. Network Load Balancing is combined with the Session Directory service in Terminal Server to provide improved scalability and availability for Terminal Services.

  For more information about creating Terminal Services designs to host applications, see "Hosting Applications with Terminal Server" in this book, or see the Session Directory and Load Balancing Using Terminal Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

  **Note**

  - Network Load Balancing is an inappropriate solution if you are using Terminal Services to provide remote administration for specific servers, because Network Load Balancing can distribute your administration traffic to any one of the cluster hosts within the cluster. When you want to use Terminal Services to remotely administer cluster hosts, connect to the specific cluster host through the dedicated IP address on the cluster adapter or through a separate management adapter. For more information about the adapters in a cluster host, see "Selecting the Number of Network Adapters in Each Cluster Host" later in this chapter.

- Custom applications

  Network Load Balancing might be an appropriate method of improving scalability and availability for applications that your organization or third-party organizations have developed. Custom applications must adhere to the same criteria listed earlier in this section.

  The considerations for custom applications are similar to those for Web applications running on IIS 6.0. For example, any application behavior that maintains local session state across multiple TCP connections is less likely to benefit from Network Load Balancing. For more information about determining if an application requires cluster host affinity, see "Identifying Applications or Services That Require Custom Port Rules" later in this chapter.

  In addition to being compatible with Network Load Balancing, custom applications must be compatible with Windows

Server 2003. For help in determining if your application is compatible with Windows Server 2003, use the Windows Application Compatibility Toolkit on the *Windows Server 2003 Deployment Kit* companion CD. For more information about the Windows Application Compatibility Toolkit, in Help and Support Center for Windows Server 2003, click **Tools**, and then click Windows Resource Kit Tools Help.

# Determining the Core Specifications in the NLB Design

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you identify applications that can benefit from Network Load Balancing, you are ready to design the core specifications for your Network Load Balancing design. These core specifications form the foundation on which you create your cluster. They include the design process steps that are required for all Network Load Balancing solutions. Figure 8.5 illustrates the current step in the process for creating your Network Load Balancing design. The steps that occur later in the design process depend on the design decisions that you make about these essential aspects of your design.

**Figure 8.5 Determining the Core Specifications in the NLB Design**



# Combining Applications on the Same Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you create your Network Load Balancing design, one of the first steps is to determine if you can combine the applications and services in your solution on the same cluster. One of the primary concerns when combining applications on the same cluster is determining if the applications are compatible with each other. Table 8.1 lists the categories of common applications and services that run on Network Load Balancing, and it describes how you can combine them on the same cluster.

**Table 8.1 Compatibility of Applications and Services on a Cluster**

| Application or Service | Combined on the Same Cluster |
|---|---|
| IIS 6.0 Web applications | • Can be combined on the same cluster; however, might require customized port rules.[1] |
| Terminal Services | • Can run any combination of applications as long as the applications are compatible with Terminal Services.<br>• Avoid combining Terminal Services (when Terminal Services is hosting applications) with other application platforms or services, such as IIS 6.0 or VPN remote access.[2] |
| VPN remote access | • Can be combined with ISA Server to combine remote access server and firewall features in the same server.<br>• Otherwise, avoid combining with other application platforms and services, such as IIS 6.0. |
| ISA Server | • Can be combined with Routing and Remote Access to combine remote access server and firewall features.<br>• Can be combined with IIS 6.0 to provide default Web site redirection to a Web site that is local to the server running ISA Server and IIS 6.0. |
| Custom applications | • Can be combined on the same cluster; however, might require customized port rules.[1] |

[1] For more information, see "Identifying Applications or Services That Require Custom Port Rules" later in this chapter.

[2] If you are not hosting applications with Terminal Services, you can combine Terminal Services with IIS 6.0 Web applications, VPN remote access servers, ISA Server, and your custom applications to provide remote administration of those servers.

As mentioned in Table 8.1, avoid running some applications and services, such as Terminal Services and VPN remote access, on the same cluster. The reasons for not running these applications and services on the same cluster include system resource constraints, security constraints, and ease of management.

In some instances, the cluster hosts might not have sufficient system resources to run the combined applications. For example, if you use Terminal Services to host applications, Terminal Services consumes a significant amount of processor and memory resources. Hosting applications with Terminal Services on the same cluster as IIS 6.0 might cause significant delays in running Web applications. For more information about scaling applications and services on Network Load Balancing clusters, see "Scaling NLB Solutions" later in this chapter.

Also, security considerations might require running the applications and services on separate clusters. For example, you might want to avoid running an FTP site that allows anonymous access on a cluster that also runs a secured e-commerce Web application.

Finally, combining applications on the same cluster can make the applications unmanageable and difficult to administer. The complexity of combining applications can correspondingly increase the complexity of managing and operating the cluster. For more information about creating clusters that are easy to manage and operate, see "Ensuring Ease of Cluster Management and Operations" later in this chapter.

# Example: Combining Applications on the Same Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A fictitious organization, Contoso, is in the process of restructuring its existing network infrastructure and application platforms. Contoso wants to reduce the number of Network Load Balancing clusters required to support its new network infrastructure and application platforms. The company plans to run the following applications and services on Network Load Balancing clusters:

- A VPN remote access server farm, based on Routing and Remote Access
- Three e-commerce Web applications, based on IIS 6.0

  One of the applications is based on static Hypertext Markup Language (HTML) content with Common Gateway Interface (CGI) scripts, another application is an ASP application, and the third application is an ASP.NET application.
- A customer support FTP site, based on IIS 6.0

  The customer support FTP site supports the downloading of information and documents to customers and the uploading of customer documents, files, and other information to the FTP site. The customer support FTP site allows anonymous access for users who want to download and upload files.

For each of these applications and services, Contoso must determine if the application can be combined with other applications and services or if the application requires its own cluster. Table 8.2 lists the clusters required to support the applications and services and the reasons for including those clusters in the design.

**Table 8.2 Clusters Required by Contoso**

| Cluster Required | Applications Running on the Cluster | Reason for Inclusion |
|---|---|---|
| NLBClusterA | VPN remote access server farm | VPN remote access servers should not be combined with other applications platforms or services. |
| NLBClusterB | E-commerce Web applications | The three -commerce Web applications are compatible with one another. |
| NLBClusterC | Customer support FTP site | The customer support FTP site supports anonymous user access and cannot be combined with the e-commerce Web applications. |

# Specifying Cluster and Cluster Host Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For each Network Load Balancing cluster, there are settings that are common to all cluster hosts and there are other settings that are unique to each cluster host. The cluster parameters define the cluster and establish cluster-wide configurations, such as the virtual IP address assigned to the cluster. The cluster host parameters define the cluster host role and identity within the cluster, such as the cluster host's priority within the cluster or the dedicated IP address that is unique to the cluster host.

Specify the cluster and cluster host parameters by completing the following steps:

1. Specify the settings that are common to all cluster hosts within the same cluster.

2. Specify the settings that are unique to each cluster host.

**Note**

- For a Word document to assist you in documenting cluster and cluster host parameters, see "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

# Specifying the Cluster Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Cluster parameters are the configuration settings that define the cluster and the configuration settings common to all cluster hosts within the cluster. The cluster parameters must be unique within your organization's network.

Specify the cluster parameters by completing the following steps:

1. Specify the cluster IP address.
2. Specify the cluster fully qualified domain name (FQDN).
3. Specify the cluster operation mode.
4. Specify the remote control settings.

**Tip**

- When you configure cluster parameters by using Network Load Balancing Manager, enter the cluster parameters once during the creation of the cluster. As cluster hosts are added to the cluster, Network Load Balancing Manager automatically configures the cluster parameters for the new cluster hosts. When you configure them by using other methods, configure the cluster parameters identically for all cluster hosts within the same cluster. For more information about configuring the cluster parameters, see "Network Load Balancing parameters" in Help and Support Center for Windows Server 2003.

**Specifying the Cluster IP Address**

The cluster IP address is the virtual IP address that is assigned to the cluster. Client requests are sent to the cluster IP address. The cluster IP address has a corresponding subnet mask that is part of the cluster IP address specifications. The criteria for specifying the cluster IP address and subnet mask are the same as for all other IP addresses. When you configure the cluster manually, you must configure the cluster IP address and subnet mask identically on all cluster hosts within the cluster.

**Tip**

- The cluster IP address must appear in the list of IP addresses in the TCP/IP properties. Network Load Balancing Manager automatically configures the TCP/IP properties so that the cluster IP address is in the list. When you configure the TCP/IP properties by other methods, you must ensure that the cluster IP address is in the list of IP addresses in the TCP/IP properties.

**Specifying the Cluster FQDN**

For each cluster, you must designate the FQDN to be assigned to the cluster. The FQDN is registered in DNS later in the deployment process, when the cluster is deployed and ready to service client requests. The cluster FQDN refers to the cluster as a whole.

**Tip**

- You can enter the cluster FQDN in the full Internet name setting. The full Internet name is not automatically registered in DNS or used by other Windows components. As a result, treat the full Internet name setting as a comment that allows administrators to easily identify the cluster's FQDN. If you elect to leave the full Internet name setting blank, the operation of any Windows component, including Network Load Balancing, is not affected.

In addition to the cluster FQDN, designate a FQDN for each application and service running on the cluster. Ultimately these FQDNs become DNS entries as you deploy the corresponding application or service.

**Specify the Cluster Operation Mode**

The cluster operation mode determines the method, *unicast* or *multicast*, that is used to propagate incoming client requests to all the cluster hosts. For more information about determining which method to select for the cluster, see "Selecting the Unicast

or Multicast Method of Distributing Incoming Requests" later in this chapter.

**Specify the Remote Control Settings**

The remote control settings provide the ability to perform certain remote administration tasks, such as starting and stopping a cluster host, from the command line utility Nlb.exe. Because of security-related concerns, avoid enabling the remote control settings. Nlb.exe can only perform limited administrative tasks. Network Load Balancing Manager can perform all administrative tasks, and is the preferred method for administering clusters. Remotely administering a cluster by using Network Load Balancing Manager is not affected by the remote control settings. For more information about specifying the remote control settings, see "Securing NLB Solutions" later in this chapter.

# Specifying the Cluster Host Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The cluster host parameters are the configuration settings that define each cluster host, including the configuration settings that are unique to each cluster host within the cluster. The cluster host parameters must be unique within a cluster and within your organization's network.

Specify the cluster host parameters by completing the following steps:

1. Specify the cluster host priority.
2. Specify the dedicated IP address configuration.
3. Specify the initial host state.

For more information about configuring the cluster host parameters, see "Network Load Balancing parameters" in Help and Support Center for Windows Server 2003.

**Tip**

- Network Load Balancing Manager prevents common cluster host configuration errors, such as duplicate cluster host priorities or duplicate dedicated IP addresses. When you use other methods to configure the cluster host parameters, you must ensure that the cluster host parameters are configured appropriately for all cluster hosts within the cluster.

**Specifying the Cluster Host Priority**

For each cluster host, you must specify a cluster host priority that is unique within the cluster. During cluster convergence, the cluster host with the lowest numeric value for the cluster host priority triggers the end of convergence. For example, if three cluster hosts have the priorities of 3, 16, and 22, 3 is the cluster host with the highest priority and it will trigger the end of convergence.

Specify the cluster host priority for your cluster by using any unique identifier between 1 and 32. Any sequence can be used as long as the cluster host priorities are unique.

**Note**

- If you specify the same cluster host priority for two cluster hosts, the last cluster host that starts fails to join the cluster. An error message describing the problem is written to the Windows system event log. The existing cluster hosts continue to operate as before.

**Specifying the Dedicated IP Address Configuration**

The dedicated IP address is an IP address that is assigned to each cluster host for network traffic that is not associated with the cluster, such as Telnet access to a specific host within a cluster. This IP address is used to individually address each host in the cluster; therefore, it should be unique for each host. Enter this parameter in standard Internet dotted notation (for example, w.x.y.z).

Traffic that is sent to the dedicated IP address is not load balanced by Network Load Balancing. Network Load Balancing ensures that all traffic to the dedicated IP address is unaffected by the Network Load Balancing current configuration, including:

- When a host is running as part of the cluster.
- When Network Load Balancing is disabled as a result of parameter errors in the registry.

**Tip**

- The dedicated IP address must be the first IP address in the list of IP addresses in the TCP/IP properties. Network Load Balancing Manager automatically configures the TCP/IP properties so that the dedicated IP address is first. When you use other methods to configure the TCP/IP properties, you must ensure that the dedicated IP address is the first IP address

listed in the TCP/IP properties.

**Specifying the Initial Host State**

The *initial host state* specifies whether Network Load Balancing starts and whether the cluster host joins the cluster when the operating system starts. You need to determine the correct initial host state for the applications and services running on the cluster.

Network Load Balancing starts very early in the system start sequence. As a result, a cluster host can join the cluster before the applications and services running on the cluster host are ready to handle traffic. In this situation, clients might be directed to the cluster host and experience outages.

In some instances, management software, such as MOM or Applications Center 2000, is responsible for starting Network Load Balancing. The management software monitors the applications and services running on the cluster host and starts Network Load Balancing when the applications are fully operational.

In other instances, you might decide to start Network Load Balancing manually to ensure that applications and services are running before Network Load Balancing starts.

Table 8.3 lists the possible settings for the initial host state and the reasons for selecting the specific initial host state.

**Table 8.3 Selecting the Appropriate Initial Host State**

| Initial Host State | Reasons for Selecting the Initial Host State |
|---|---|
| Started | • The applications and services running on the cluster start before Network Load Balancing.<br>• The length of time between the start of Network Load Balancing and the start of applications and services running on the cluster is negligible. |
| Stopped | • Management software, such as MOM or Application Center 2000, is responsible for starting Network Load Balancing automatically.<br>• The applications and services running on the cluster start after Network Load Balancing, and you want to start the cluster host manually. |
| Suspended | • You have performed maintenance on the cluster host, and you want to prevent the cluster from responding to clients after a restart of the cluster host. |

# Example: Specifying Cluster and Cluster Host Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A fictitious organization, Contoso, is designing a VPN remote access solution, based on Routing and Remote Access and Network Load Balancing. The VPN remote access solution provides remote access to Contoso's private network by establishing Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) VPN tunnels through the Internet. The VPN remote access server farm contains five servers that have identical system resources.

Table 8.4 lists the cluster parameter design decisions for the VPN remote access server farm and the reasons for making those decisions.

**Table 8.4 Cluster Parameter Settings for the VPN Remote Access Server Farm**

| Decision | Reason for the Decision |
|---|---|
| The cluster IP address and subnet mask is assigned an IP address and subnet mask that is accessible from the Internet. | Remote access users need to access the VPN server farm running on the cluster from the Internet. |
| The cluster is assigned the FQDN of vpn.contoso.com. | The FQDN assigned to the cluster is the name used by remote access users when accessing the VPN server farm running on the cluster. |
| The cluster operation mode is set to unicast. | The network infrastructure supports the unicast cluster operation mode. For more information about the decisions regarding setting this mode, see "Selecting the Unicast or Multicast Method of Distributing Incoming Requests" later in this chapter. |
| The remote control settings are not enabled. | The remote control settings are not required, because Network Load Balancing Manager will be used to administer the cluster. |

Table 8.5 lists the cluster host parameter design decisions for the VPN remote access server farm.

**Table 8.5 Cluster Host Parameter Settings for the VPN Remote Access Server Farm**

| Cluster Host Name | Host Priority | Dedicated IP Address | Initial Host State |
|---|---|---|---|
| NLBClusterA-01 | 1 | Not specified | Started |
| NLBClusterA-02 | 2 | Not specified | Started |
| NLBClusterA-03 | 3 | Not specified | Started |
| NLBClusterA-04 | 4 | Not specified | Started |
| NLBClusterA-05 | 5 | Not specified | Started |

A unique cluster host priority is assigned to each cluster host. The cluster host initial host state is set to Started to ensure that Network Load Balancing starts automatically and that the cluster host automatically joins the cluster.

# Controlling the Distribution of Client Traffic Within the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

One of the intended purposes of Network Load Balancing is to distribute incoming client traffic within the cluster. You can control the distribution of client traffic within the cluster by using Network Load Balancing port rules. Port rules are criteria-based policies that allow you to direct client requests to specific cluster hosts, based on TCP and UDP port numbers.

A default port rule is created during the installation of Network Load Balancing. In many instances, the default port rule is sufficient for some of the applications and services that use Network Load Balancing. When the default port rule is insufficient, you can create custom port rules. For more information about the default port rule, see "Identifying the Behavior of the Default Port Rule" later in this chapter.

The default port rule is sufficient for the following applications and services:

- VPN remote access farms with Routing and Remote Access
- Load-balancing application hosting with Terminal Services

**Tip**

- The ISA Server setup process defines the port rules that are necessary for ISA Server. No custom port rules are necessary for ISA Server.

If the default port rule is sufficient for your solution, and creating custom port rules is unnecessary, see "Specifying Cluster Network Connectivity" later in this chapter.

For IIS 6.0 Web farms or for custom applications, custom port rules might be required.

In some instances, the applications and services might require that the same client traffic be handled differently for the same or different applications and services. The virtual IP address assigned to the cluster can handle client traffic only in one way. However, you can specify a *virtual cluster* for each of the applications, allowing each application to have its own load-balancing behavior. Virtual clusters are a logical construct within the cluster, and they require no additional hardware.

For example, two Web applications might require different load-balancing behavior for HTTP (TCP port 80). You can create a virtual cluster for each Web application that allows different load-balancing behavior for the HTTP client traffic.

You can create a virtual cluster by specifying a virtual IP address in a Network Load Balancing port rule. The virtual IP address that is assigned in the port rule is associated with the application that requires the different load-balancing behavior.

Figure 8.6 illustrates the relationship between a Network Load Balancing cluster and the virtual clusters specified for the cluster. Each of the applications — Web applications A, B, and C — requires different load-balancing behavior. A virtual IP address is assigned to each virtual cluster and associated with each application. A DNS entry associates the virtual IP address with a URL for the corresponding application.

**Figure 8.6 Relationship Between an NLB Cluster and Virtual Clusters**




For more information about including Network Load Balancing port rules in your design, see "Identifying Applications or Services That Require Custom Port Rules" later in this chapter.

Control the distribution of client traffic within a cluster by completing the following steps:

1. Identify the behavior of the default port rule.

2. Identify applications or services that require custom port rules.

3. Specify the client traffic to be affected by the custom port rule.

4. Specify the affinity and load-balancing behavior of the custom port rule.

**Important**

- The port rules applied to each cluster host must be identical, with the exception of the load weight (in the multiple hosts filter mode) and the handling priority (in the single hosts filter mode). If there is a discrepancy between port rules on existing cluster hosts, the cluster will not converge.

# Identifying the Behavior of the Default Port Rule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When Network Load Balancing is installed on a cluster host, a default port rule is created. Table 8.6 lists the configuration of the default port rule that is created during Network Load Balancing installation.

**Table 8.6 Specifications for the Network Load Balancing Default Port Rule**

| Default Port Rule Setting | Set to This Value |
|---|---|
| Cluster IP address | All |
| Port range: From | 0 |
| Port range: To | 65535 |
| Protocols | Both |
| Filtering Mode | Multiple Hosts with Single affinity and Equal load weight |

Specify that the default port rule be deleted unless the default port rule is:

- Appropriate for the applications and services installed on the cluster.
- Modified for the applications and services installed on the cluster.

# Identifying Applications or Services That Require Custom Port Rules

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As previously mentioned, the default port rule is sufficient for some applications and services. However, many applications running on IIS 6.0 or custom applications that are developed by your organization might require customized port rules.

These applications and services can require customized port rules to influence how load is directed to hosts in the cluster. In addition, the applications and services might require a virtual cluster when the same client traffic must be handled differently for the same or different applications and services. With virtual clusters, you can use different port rules for different Web sites or applications hosted on the cluster, provided each Web site or application has a different virtual IP address.

**Identifying Applications and Services That Need Persistent Sessions**
The applications and services that run on Network Load Balancing include stateful applications (those that maintain session state) and stateless applications. Maintaining session state means that the application or service collects information when first connecting to a cluster host and then retains the information for subsequent requests. During a user session, the same server must handle all the requests from the user in order to access that information. Applications and services that are stateless maintain no user or communication information for subsequent connections.

With a single server, maintaining session state presents no difficulty, because the user always connects to the same server. However, when client requests are load balanced within a cluster, without some type of persistence, the client might not be directed to the same cluster host for a series of client requests.

In Network Load Balancing, you maintain session state with a *port rule affinity* between the client and a specific cluster host. Port rule affinity directs all client requests from the same IP address to the same cluster host. You can use port rules to specify the port rule affinity between clients and cluster hosts. For more information about specifying port rule affinity between clients and cluster hosts, see "Specifying the Affinity and Load-Balancing Behavior of the Custom Port Rule" later in this chapter.

Port rule affinity might be required by the following:

- An application or service running on the cluster
- Session-oriented protocols that are used by the application or service running on the cluster if the session lifetime extends across multiple TCP connections

**Identifying applications that require affinity**
Applications and services require persistent sessions when the application or service retains information established in a client request that is used in subsequent client requests. In some instances, the application or service is aware that load balancing is occurring, and it uses an application-specific method, such as client-side cookies, for maintaining session state. In other instances, the application or service is unaware that load balancing is occurring, and it requires Network Load Balancing to maintain affinity between the client and specific cluster hosts.

An example of an application that maintains session state is a Web application that requires a user to log on to buy products through a shopping cart application. After the application authenticates the user, the user's information, including shipping information, billing information, and items in the shopping cart, is retained for subsequent requests. The session state is maintained until the user completes or cancels the purchase.

Table 8.7 lists common Web application types and their requirements for affinity provided by cluster port rules.

**Table 8.7 Web Applications and Their Requirements for Port Rule Affinity**

| Web Application | Description | Requires Port Rule Affinity |
|---|---|---|

| | | |
|---|---|---|
| Static, HTML-based applications | Application maintains no user information. | |
| Static, HTML-based applications with CGI | CGI portion of the application might retain user information. | • |
| Web application that uses client-side cookies | Application sends a cookie to the client when the application session is initiated. On subsequent requests, the client sends the cookie along as part of the request. The instance of the application running on individual cluster hosts is capable of retrieving application session state by using the cookie in the request. | |
| ASP.NET applications with session state persistence | ASP.NET applications support a method for maintaining session state on a centralized session state server or on a server running SQL Server 2000. Because the session state is managed centrally, any cluster host can recover session state information, and affinity is not required. For more information about ASP.NET application session state, see "Deploying ASP.NET Applications in IIS 6.0" in *Deploying Internet Information Services (IIS) 6.0*, or see "Deploying ASP .NET Applications in IIS 6.0" at the Microsoft Web site. | |
| ASP applications and ASP.NET applications without session state | These applications create session information when the user first starts the application. The session information is managed by the instance of the ASP or ASP.NET dynamic-link libraries (DLLs) running on the cluster host. On subsequent requests, the session state is maintained for the application on that specific cluster host. No session state information is shared among cluster hosts. | • |

For applications that are written by your organization, consult with the application developers to determine if, or how, any session information is retained. When your applications run on Terminal Services, port rule affinity is required.

**Identifying session-oriented protocols that require affinity**

Even when the application or service is stateless, the protocol required by the application or service might be session-oriented. For example, an application might be based on static HTML, but it might use SSL to provide security.

However, other session-oriented protocols, such as SSL, can operate correctly without affinity, but they incur a significant decrease in performance. When an SSL session is established with a cluster host, an SSL session ID is exchanged between the cluster host and the client. If the client makes a subsequent request to the same cluster host, the same session ID can be used.

If the client makes a subsequent request to a different cluster host, a new session ID must be obtained. The overhead for obtaining a new SSL session ID is five times more than for reusing a session ID. By using port rule affinity, SSL session IDs are reused whenever possible.

**Compensating for Differences in System Resources**

During the IT life cycle of your cluster, new cluster hosts might be added to the cluster. These new cluster hosts are typically new computers with improved system resources, such as processor speed, number of processors, or memory. Because of the improved system resources, the new cluster hosts are capable of servicing more clients.

The default port rule, created during the installation of Network Load Balancing, evenly distributes network requests among cluster hosts, regardless of the available system resources on individual cluster hosts. Cluster hosts with inadequate system resources can provide slower response times to clients than other cluster hosts with adequate system resources. You can compensate for differences in available system resources on individual cluster hosts by directing a higher percentage of client requests to the cluster hosts with adequate system resources.

The method of compensating for the differences in system resources is based on the filter mode in the port rule. Only the Multiple Hosts port filter mode supports compensating for differences in cluster host resources. For more information about port rule filter modes and, specifically, the Multiple Hosts port filter mode, see "Specifying the Affinity and Load-Balancing Behavior of the Custom Port Rule" later in this chapter.

**Deciding to Include Virtual Clusters**

In some instances, you might want a cluster to appear logically as multiple Network Load Balancing clusters. In Network Load Balancing, you can define virtual clusters that allow you to apply a unique set of port rules to each virtual cluster.

Include virtual clusters in the following situations:

- An application or service running on a cluster must exhibit different affinity behavior.

  For example, you might have an FTP site hosted on multiple servers running IIS 6.0. You want FTP downloads to be load balanced across all cluster hosts. However, you want FTP uploads to be sent only to one cluster host. You can create a virtual cluster that supports FTP downloads with no affinity and another virtual cluster that directs FTP uploads to a specific cluster host.

- Multiple applications or services running on the same cluster require different affinity behaviors for the same client traffic type (same TCP/UDP port number).

  For example, you might have two Web applications hosted on the same cluster running IIS 6.0. Both applications use HTTP (TCP port 80), but one application maintains session state while the other application does not. You can create a virtual cluster that supports the application that maintains session state with affinity and another virtual cluster that supports the stateless application without affinity.

- The maintenance and operations of an application or service must be independent of other applications and services running on the cluster.

  For example, you might have two applications hosted on the same cluster. For ease of maintenance, upgrade, or other operations tasks, you can create a virtual cluster for the application. You can stop client communication with the virtual cluster and then perform application maintenance without affecting other applications running on the cluster.

  **Tip**

  - You can stop client communications with a virtual cluster by using the **drain** parameter from Nlb.exe or Network Load Balancing Manager, both of which are a part of Windows Server 2003.

Virtual clusters are specialized port rules that require a virtual IP address, which is assigned to the virtual cluster. You define a virtual cluster by specifying one or more port rules that share the same virtual IP address. As with regular port rules, the port rules that define the virtual cluster must be the same for all cluster hosts in the cluster. For more information about port rule specifications, see "Specifying Client Traffic To Be Affected by the Custom Port Rule" later in this chapter.

# Specifying Client Traffic To Be Affected by the Custom Port Rule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Port rules can be divided into two parts. The first part of a port rule identifies the client traffic to be affected by the port rule. You identify the client traffic affected by the port rule by specifying the cluster IP address and the TCP (or UDP) port range. Network Load Balancing examines all client requests and determines if a port rule applies to the client request.

The second part of a port rule determines the affinity and load-balancing characteristics of the cluster. For more information about the affinity and load-balancing characteristics of the cluster, see "Specifying the Affinity and Load-Balancing Behavior of the Custom Port Rule" later in this chapter.

Specify the client traffic to be affected by the port rule by completing the following steps:

1. Designate the value for the cluster IP address.
2. Specify the port range for the port rule as a starting port number (From) and an ending port number (To).
3. Specify the protocol for the port range that you specified.

**Note**

- For a Word document to assist you in documenting your port rule settings, see "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

**Designating the Value for the Cluster IP Address**

If you are specifying a:

- Global port rule for a cluster, specify All.
- Port rule for a virtual cluster, specify the virtual IP address for the virtual cluster.

For more information about deciding when to include virtual clusters in your design, see "Identifying Applications or Services That Require Custom Port Rules" earlier in this chapter.

**Specifying the Port Range**

Specify the port range for the port rule:

- Based on the TCP and UDP ports used by the applications and services running on the Network Load Balancing cluster.
- As a starting port number (From) and an ending port number (To).

For example, if your Network Load Balancing cluster supports Web applications that use only TCP port 80, you specify a starting port number of 80 and an ending port number of 80.

**Specifying the Protocol for the Port Range**

If the applications and services are using:

- TCP, specify TCP.
- UDP, specify UDP.
- Both TCP and UDP, specify Both.

# Specifying the Affinity and Load-Balancing Behavior of the Custom Port Rule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The second part of a port rule affects the affinity and load-balancing behavior of the port rule. The affinity and load-balancing behavior of a port rule are specified by the filter mode of the port rule. Only one filter mode can be selected for each port rule. Table 8.8 describes the filter modes that are available.

**Table 8.8 Selecting the Filter Mode for the Port Rule**

| Filter Mode | Description |
|---|---|
| Multiple Hosts | Permits all cluster hosts to actively respond to client requests. This is the most common filter mode, because it allows the affinity and load-balancing characteristics to be customized. |
| Single Host | Allows only one cluster host in the cluster to actively respond to client requests. Could be useful when providing backup servers. For example, in an FTP site where users upload files, single host mode allows one host to receive FTP uploads. If the host fails, the host with the next highest priority takes over for the failed host. |
| Disable | Prevents the cluster from responding to a specific type of client traffic. |

**Note**

- For a Word document to assist you in documenting your port rule settings, see "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

#### Specifying Multiple Hosts Filter Mode Settings

The Multiple Hosts filter mode provides load balancing across multiple cluster hosts. One of the complexities introduced by load balancing is persistent relationships (affinity) between clients and cluster hosts. For more information about affinity, see "Identifying Applications or Services That Require Custom Port Rules" earlier in this chapter.

Another complexity introduced by load balancing is compensation for differences in cluster host system resources on the cluster over time. Based on differences in available system resources or applications, specific cluster hosts might not be capable of managing the same number of client requests. For more information about compensating for differences in cluster host system resources, see "Identifying Applications or Services That Require Custom Port Rules" earlier in this chapter.

Specify the settings for the Multiple Hosts filter mode by completing the following steps:

1. Preserve persistent application sessions by specifying the affinity between a client and a specific cluster host.

2. Compensate for differences in system resources by specifying the load weight settings.

#### Selecting the affinity between a client and a specific cluster host

You can override the default behavior by selecting other port rule affinity options. Table 8.9 lists the port rule affinity options and the reasons for selecting them.

**Table 8.9 Including Port Rule Affinity Options in Your Design**

| Option | Reasons for Selecting This Option |
|---|---|

| | |
|---|---|
| None | • You want to ensure even load balancing among cluster hosts.<br>• Client traffic is stateless (for example, HTTP traffic). |
| Single | • You want to ensure that requests from a specific client (IP address) are sent to the same cluster host.<br>• Client state is maintained across TCP connections (for example, HTTPS traffic). |
| Class C | • Client requests from a Class C IP address range (instead of a single IP address) are sent to the same cluster host.<br>• Clients use multiple proxy servers to access the cluster, and they appear to have multiple IP addresses within the same Class C IP address range.<br>• Client state is maintained across TCP connections (for example, HTTPS traffic). |

Single affinity is the most common selection when applications require that information about the user state be maintained across TCP connections. Examples of these types of applications include applications that use SSL or applications that retain user information, such as e-commerce shopping cart applications.

Applications that use SSL with Single affinity are efficient because the SSL session IDs are reused. Negotiating a new SSL session ID requires five times the amount of overhead as reusing a SSL session ID. Although negotiating the SSL session ID is transparent to the client, the cumulative increase in overhead could degrade the performance of the cluster.

Applications that retain user information can resolve the affinity requirement by using Network Load Balancing affinity or by using a common session state database or server. Applications that have a common session state database or server are combined with cookie-based affinity to allow any cluster host to restore the appropriate session state. If an application has a common session state database or server, you can select a port rule affinity of **None** if SSL is not part of the solution.

If an application retains user information by requiring that client transactions are completed on the same cluster host, select a port rule affinity of Single. For a discussion of when an application requires persistent connections to specific cluster hosts, along with an explicit discussion regarding SSL, see "Identifying Applications or Services That Require Custom Port Rules" earlier in this chapter.

Class C is used when Internet clients connect through proxy servers with different IP addresses within the same Class C IP address range. Some Internet service providers, such as America Online (AOL) in the United States, have proxy servers with different Class C IP addresses. Using different Class C IP addresses on the proxy servers can break the affinity between the client and the cluster host. In situations like this, other methods might be required to preserve cluster host affinity. For example, Application Center 2000 supports cookie-based affinity, which solves the problem of proxy servers with different Class C IP addresses. Alternatively, the application can be redesigned to maintain application session state in a common database or session state server, such as that provided by IIS 6.0 and ASP.NET. When possible, allow the application to maintain session state in this manner because the application will be more robust and scalable.

**Specifying the load weight settings for each cluster host**

The default port rule that is created during the installation of Network Load Balancing uses equal load weights. To override the default behavior, you can specify a custom load weight to be handled by each cluster host, as shown in Table 8.10.

**Table 8.10 Selecting the Method for Distributing Client Requests**

| Method | Description |
|---|---|
| Equal | Evenly distributes client requests across all cluster hosts when the available system resources are the same. |
| Load weight | Distributes client requests based on the available system capacity because of differences in the:<br>• Hardware configuration of each cluster host.<br>• Applications and services running on each cluster host. |

Specify the amount of client requests to be handled by each cluster host by completing one of the following steps:

- Specify Equal to ensure that client requests are evenly distributed across all cluster hosts, unless the cluster hosts do not have the same available system capacity.

- Otherwise, specify the load weight value for the cluster host.

  The load weight value can range from 0 (zero) to 100. You can prevent a cluster host from handling any client requests by specifying a load weight of 0.

  The percentage of client requests that are handled by each cluster host is computed by dividing the local load weight by the sum of all load weights across the cluster. Table 8.11 shows an example of the relationship between load weight and the percentage of client requests that are handled by each cluster host.

  **Table 8.11 Using Load Weights and Client Request Percentages**

| Cluster Host | Load Weight | Percentage of Client Requests |
|---|---|---|
| ClusterHost-A | 50 | 40% |
| ClusterHost-B | 50 | 40% |
| ClusterHost-C | 25 | 20% |

**Specifying Single Host Filter Mode Settings**

As with the cluster host priority setting for Network Load Balancing, which is discussed in "Specifying the Cluster Host Parameters" earlier in this chapter, you must specify the handling priority value for each cluster host that has the port rule.

Specify the handling priority for Single Host filter mode by completing the following steps:

1. Specify a value of 1 as the priority for the cluster host that will always handle the client traffic designated in the port rule criteria.

2. Increase the value assigned to the previous cluster host by 1, assign that value to the next cluster host, and repeat the process until you have specified the handling priority for all cluster hosts.

### Important

- If you specify the same handling priority for two cluster hosts, the last cluster host that starts will fail to join the cluster. An error message describing the problem will be written to the Windows system event log. The existing cluster hosts will continue to operate as before if cluster convergence completed previously. Otherwise, convergence must complete before traffic is handled by the cluster.

**Specifying Disable Filter Mode Settings**

As previously mentioned, the Disabled filter mode means that the client traffic corresponding to the port rule is blocked. Unlike the other filter modes, after you select this mode, no additional settings or values are required.

# Example: Controlling the Distribution of Client Traffic Within the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization implements the following solutions, which include Network Load Balancing, to reduce outages and improve performance:

- VPN remote access, based on Routing and Remote Access
- E-commerce Web applications, based on IIS 6.0
- A customer support FTP site, based on IIS 6.0

**VPN remote access solution with Routing and Remote Access**
The VPN remote access solution, based on Routing and Remote Access, provides remote access to the organization's private network by establishing PPTP and L2TP VPN tunnels through the Internet.

The VPN remote access server farm contains five servers that have identical system resources. Network Load Balancing is enabled on the network adapters connected to the Internet. A cluster IP address has been configured for the cluster, and the appropriate DNS entries have been determined but not yet created.

Because the default port rule provides the appropriate affinity and load balancing, no custom port rules are required for the cluster.

**E-commerce Web application solution with IIS 6.0**
The e-commerce Web application solution, based on IIS 6.0, includes the following applications:

- Web-based, e-commerce application built with static HTML
- Web-based, e-commerce application built with ASP

Both applications use:

- HTTPS (TCP port 443)
- HTTP (TCP Port 80)

The application that uses ASP maintains user session information after the user is authenticated.

The two e-commerce Web application solutions are combined on the same IIS 6.0 Web farm and, subsequently, the same Network Load Balancing cluster. The IIS 6.0 Web farm contains four computers that have identical system resources. Network Load Balancing is enabled on the network adapters connected to the Internet. A cluster IP address is configured for the cluster, and the appropriate DNS entries have been determined but not yet created.

To facilitate the combination of the two e-commerce Web applications into a single Web farm and cluster, each e-commerce application is assigned to a virtual cluster. By assigning each e-commerce application to a virtual cluster, client access can be stopped individually to allow operations tasks, such as upgrades, to be performed without disrupting client access to the other e-commerce application.

Table 8.12 lists the cluster, virtual cluster, and cluster hosts in the organization's e-commerce solution.

**Table 8.12 Network Load Balancing Clusters and Cluster Hosts in the E-Commerce Solution**

| Cluster Name | Solution | Type | Cluster Host |
|---|---|---|---|

| | | | |
|---|---|---|---|
| NLBCluster-B | Web-based, e-commerce applications | Cluster | • NLBClusterB-01<br>• NLBClusterB-02<br>• NLBClusterB-03<br>• NLBClusterB-04 |
| VirCluster-A | IIS 6.0 and HTML | Virtual cluster | Same as NLBCluster-B |
| VirCluster-B | IIS 6.0 and ASP | Virtual cluster | Same as NLBCluster-B |

Table 8.13 lists the port rules that meet the requirements of the e-commerce Web application solution that includes IIS 6.0 and Network Load Balancing.

**Table 8.13 Port Rules for IIS 6.0 E-Commerce Solution on NLBCluster-B**

| Cluster IP Address | Start | End | Protocol | Filtering Mode | Load Weight | Affinity |
|---|---|---|---|---|---|---|
| VirtualIP-A | 80 | 80 | TCP | Multiple | Equal | None |
| VirtualIP-A | 443 | 443 | TCP | Multiple | Equal | Single |
| VirtualIP-B | 80 | 80 | TCP | Multiple | Equal | Single |
| VirtualIP-B | 443 | 443 | TCP | Multiple | Equal | Single |

Because the port rules, listed in Table 8.13, have a specified load weight of Equal, the same port rules are used for all the cluster hosts. The virtual clusters, VirtualIP-A and VirtualIP-B, are dedicated to the respective e-commerce applications. Because both applications use HTTP (TCP port 80) and HTTPS (TCP port 443), port rules must be specified for each protocol in each virtual cluster.

**Customer support FTP site solution with IIS 6.0**

The customer support FTP site solution:

- Is based on IIS 6.0.
- Provides secured and unsecured access to files.
- Allows users to upload files to specified areas on the FTP site.
- Requires file uploads that are centralized on one FTP server to avoid users uploading duplicate files.
- Uses TCP port 20 for FTP.
- Uses TCP port 21 for FTP.

The customer support FTP site runs on an IIS 6.0 farm and, subsequently, a Network Load Balancing cluster. The IIS 6.0 farm contains three computers that have identical system resources. Network Load Balancing is enabled on the network adapters connected to the Internet. A cluster IP address has been configured for the cluster, and the appropriate DNS entries have been determined but not yet created.

To support the number of simultaneous users who are performing FTP downloads, all FTP download requests must be load balanced across the entire IIS 6.0 farm. However, to ensure that users upload files to only one location, FTP uploads must be directed to only one server in the IIS 6.0 farm.

To facilitate the differences in cluster host affinity between FTP uploads and downloads, each direction of FTP transfer is assigned to a different virtual cluster. By assigning FTP uploads to a virtual cluster and FTP downloads to another virtual cluster, the organization ensures that FTP downloads can be load balanced across all cluster hosts, while FTP uploads are sent to only one cluster host in the cluster.

Table 8.14 lists the cluster, virtual cluster, and cluster hosts selected in the organization's customer support FTP site solution.

**Table 8.14 Clusters and Cluster Hosts in the Customer Support FTP Site Solution**

| Cluster Name | Solution | Type | Cluster Host |
|---|---|---|---|
| NLBCluster-C | FTP site | Cluster | • NLBClusterC-01<br>• NLBClusterC-02<br>• NLBClusterC-03 |
| VirCluster-C | FTP download | Virtual cluster | All cluster hosts can be used for download. |
| VirCluster-D | FTP upload | Virtual cluster | NLBClusterC-03 is to be used for upload. |

Table 8.15 lists the port rules that meet the requirements of the organization's customer support FTP site solution, which includes IIS 6.0 and Network Load Balancing.

**Table 8.15 Port Rules for an IIS 6.0 FTP Site Solution on NLBCluster-C**

| Cluster IP Address | Start | End | Protocol | Filtering Mode | Load Weight | Affinity | Handling Priority |
|---|---|---|---|---|---|---|---|
| VirtualIP-C | 20 | 21 | TCP | Multiple Hosts | Equal | Single | NA |
| VirtualIP-D | 20 | 21 | TCP | Single Host | | | • NLBClusterC-03 = 1<br>• NLBClusterC-01 = 2<br>• NLBClusterC-02 = 3 |

The port rules that define VirCluster-C are identical for all the cluster hosts in NLBCluster-C. The port rules that define VirCluster-D are unique for each cluster host in NLBCluster-C, because the handling priority for each cluster host is unique. NLBClusterC-03 is the cluster host that is designated for FTP uploads, and it is assigned a handling priority of 1 to ensure that all file uploads are sent to NLBClusterC-03. The handling priority for NLBClusterC-01 and for NLBClusterC-02 must be unique and of a lower priority than for NLBClusterC-03.

# Specifying Cluster Network Connectivity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The scalability and availability improvements provided by your cluster depend on the network connectivity that is provided to the cluster. An improperly designed network infrastructure can cause client response-time problems and application outages. Specify the connections between the cluster and the client computers, the other servers within your organization's network, and the operation consoles to achieve the connectivity goals of your network.

The goals of network connectivity to the cluster are to provide:

- High-capacity network connectivity to ensure adequate client response time.
- Redundant routing and switching infrastructure.
- Restricted management of the cluster and cluster resources.
- Valid IP configuration for each network interface in each cluster host.

Specify the cluster network connectivity by completing the following steps:

1. Select the number of network adapters in each cluster host.
2. Select the unicast method or the multicast method of distributing incoming client requests to cluster hosts.
3. Include support for teaming network adapters.
4. Determine the network infrastructure requirements.

After you have specified the cluster network connectivity, document your decisions on your organization's network diagram (for example, on a Microsoft® Visio® drawing).

**Note**

- For a Word document to assist you in documenting your decisions, see "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

# Selecting the Number of Network Adapters in Each Cluster Host

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

At a minimum, you must connect each cluster host in your cluster to a network segment that has connectivity to the client computers. In most solutions, you need to connect each cluster host to other servers and to management and operations consoles in your organization.

With the appropriate cluster network connectivity, your solution will be properly secured, highly available, highly scalable, and easy to manage. Any design deficiencies in the cluster network connectivity portion of the design can compromise the security, availability, scalability, and manageability of the cluster.

Select the number of network adapters in each cluster host by completing the following steps:

1. Identify a network interface, referred to as the *cluster adapter*, which provides connectivity to the client computers.

   Include the appropriate IP configuration (IP address, subnet mask, and so forth) for the cluster adapter, so that the cluster host is on the same physical subnet or virtual subnet.

   All hosts of the same cluster must be on the same physical or virtual LAN (VLAN). In instances where the cluster hosts are not connected to the same physical subnet, ensure that the cluster hosts are connected to a virtual subnet. If all hosts are not connected to the same physical subnet, make sure that your routing and switch infrastructure supports virtual subnets.

2. Specify a network interface, referred to as the *management adapter*, which will provide connectivity to other servers within your organization and to management and operations consoles within your organization.

   Include the appropriate IP configuration (IP address, subnet mask, and so forth) for the management adapter.

**Important**

- To prevent unauthorized altering of the Network Load Balancing configuration, enable Network Load Balancing administration only on the management adapter. This is done by restricting traffic through your firewalls or routers.

In instances where only the cluster adapter is included and your applications require peer-to-peer communications between cluster hosts (beyond the cluster heartbeat traffic), see the discussion on unicast mode and multicast mode in "Selecting the Unicast or Multicast Method of Distributing Incoming Requests" later in this chapter.

Figure 8.7 illustrates cluster network connectivity that includes connectivity to clients, other servers in an organization, and management and operations consoles.

**Figure 8.7 Cluster Network Connectivity**


Internet
Firewall-01
Firewall-02
Segment that will connect cluster adapters to the clients
Router-01
Router-02
Switch-01 belongs to VLAN-01
Switch-02 belongs to VLAN-01
VPN-01
VPN-02
VPN-03
VPN-04
VPN-05
VPN-06
VPN-07
VPN-08
NLBCluster-01
DC-01
DC-02
IAS-01
IAS-02
Segment that connects management adapters to organization's resources and operations consoles
Firewall-03

# Selecting the Unicast or Multicast Method of Distributing Incoming Requests

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

All cluster hosts in a cluster receive all incoming client requests that are destined for the virtual IP address that is assigned to the cluster. The Network Load Balancing load-balancing algorithm, which runs on each cluster host, is responsible for determining which cluster host processes and responds to the client request.

You can distribute incoming client requests to cluster hosts by using unicast or multicast methods. Both methods send the incoming client requests to all hosts by sending the request to the cluster's MAC address.

When you use the unicast method, all cluster hosts share an identical unicast MAC address. Network Load Balancing overwrites the original MAC address of the cluster adapter with the unicast MAC address that is assigned to all the cluster hosts.

When you use the multicast method, each cluster host retains the original MAC address of the adapter. In addition to the original MAC address of the adapter, the adapter is assigned a multicast MAC address, which is shared by all cluster hosts. The incoming client requests are sent to all cluster hosts by using the multicast MAC address.

Select the unicast method for distributing client requests, unless only one network adapter is installed in each cluster host and the cluster hosts must communicate with each other. Because Network Load Balancing modifies the MAC address of all cluster hosts to be identical, cluster hosts cannot communicate directly with one another when using unicast. When peer-to-peer communication is required between cluster hosts, include an additional network adapter or select multicast mode. When the unicast method is inappropriate, select the multicast method.

For more information about the interaction between the method of distributing incoming requests and layer 2 switches, see article Q193602, "Configuration Options for WLBS Hosts Connected to a Layer 2 Switches," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at www.microsoft.com/windows/reskits/webresources.

**Selecting the Unicast Method**

In the unicast method:

- The cluster adapters for all cluster hosts are assigned the same unicast MAC address.
- The outgoing MAC address for each packet is modified, based on the cluster host's priority setting, to prevent upstream switches from discovering that all cluster hosts have the same MAC address.

  The modification of the outgoing MAC address is appropriate for switches. When a hub is used to connect the cluster hosts, disable the modification of the outgoing MAC address. On Windows Server 2003, you can disable modification of outgoing addresses by setting the value of the registry entry **MaskSourceMAC**, of data type REG_DWORD, to 0x0. **MaskSourceMAC** is located in HKLM\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*Adapter-GUID* (where *Adapter-GUID* is the long GUID assigned to the network adapter in the server).

  **Caution**

  - Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference on the *Microsoft Windows Server 2003 Deployment Kit* companion CD or at http://www.microsoft.com/reskit.
- The unicast MAC address is derived from the cluster's IP address to ensure uniqueness outside the cluster hosts.
- Communication between cluster hosts, other than Network Load Balancing–related traffic (such as heartbeat), is only available when you install an additional adapter, because the cluster hosts all have the same MAC address.

Although the unicast method works in all routing situations, it has the following disadvantages:

- A second network adapter is required to provide peer-to-peer communication between cluster hosts.
- If the cluster is connected to a switch, incoming packets are sent to all the ports on the switch, which can cause switch flooding.

**Selecting the Multicast Method**

In the multicast method:

- The cluster adapter for each cluster host retains the original hardware unicast MAC address (as specified by the hardware manufacture of the network adapter).
- The cluster adapters for all cluster hosts are assigned a multicast MAC address.
- The multicast MAC is derived from the cluster's IP address.
- Communication between cluster hosts is not affected, because each cluster host retains a unique MAC address.

By using the multicast method with Internet Group Membership Protocol (IGMP), you can limit switch flooding, if the switch supports IGMP *snooping*. IGMP snooping allows the switch to examine the contents of multicast packets and associate a port with a multicast address. Without IGMP snooping, switches might require additional configuration to tell the switch which ports to use for the multicast traffic. Otherwise, switch flooding occurs, as with the unicast method.

The multicast method has the following disadvantages:

- Upstream routers might require a static Address Resolution Protocol (ARP) entry. This is because routers might not accept an ARP response that resolves unicast IP addresses to multicast MAC addresses.
- Without IGMP, switches might require additional configuration to tell the switch which ports to use for the multicast traffic.
- Upstream routers might not support mapping a unicast IP address (the cluster IP address) with a multicast MAC address. In these situations, you must upgrade or replace the router. Otherwise, the multicast method is unusable.

# Including Support for Teaming Network Adapters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many hardware vendors support *teaming network adapters* to increase network bandwidth capacity and to provide an additional level of network redundancy. Teaming network adapters are multiple network adapters in the same computer that logically act as a single, virtual network adapter. The virtual network adapter can provide load balancing of traffic between the physical network adapters or automatic failover in the event that one of the network adapters fails.

Teaming network adapter drivers and Network Load Balancing both try to manipulate the MAC addresses of the network adapters. As a result, teaming network adapters might work only in limited configurations. As a general rule, avoid including teaming network adapters in your Network Load Balancing solutions.

**Note**

- Because each implementation of teaming network adapters is vendor-specific, the support for teaming network adapters is provided by the hardware vendor, and it is beyond the scope of Microsoft Product Support Services. For more information about teaming network adapters in your Network Load Balancing solution, consult the teaming network adapter hardware recommendations.

# Determining the Network Infrastructure Requirements

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The network infrastructure affects your Network Load Balancing solution more than any other component. No matter how much Network Load Balancing enables you to scale out your solution, an inadequate routing and switching infrastructure can create a number of problems.

Even if you optimize the Network Load Balancing cluster, an inadequate routing and switching infrastructure can restrict available network throughput and prevent clients from achieving any improvement in response times. Additionally, an inadequate routing and switching infrastructure can allow a single failure in a router, switch, or network path to disrupt communications with the cluster.

Determine the network infrastructure requirements for your cluster by completing the following steps:

1. Determine the IP subnet requirements for the cluster.

2. Determine how the cluster handles inbound and outbound traffic.

3. Determine when to include switches or hubs to connect cluster hosts to one another.

For more information about how the network infrastructure affects availability, see "Ensuring Availability in NLB Solutions" later in this chapter. For more information about how the network infrastructure affects scalability, see "Scaling NLB Solutions" later in this chapter.

**Determining Cluster IP Subnet Requirements**

Network Load Balancing requires that all cluster hosts be on the same IP subnet or virtual IP subnet (VLAN). This is because all cluster hosts share the cluster's virtual IP address. The routing infrastructure sends client requests to all cluster hosts by using a virtualized unicast MAC address or a multicast MAC address. For more information about unicast and multicast addressing, see "Selecting the Unicast or Multicast Method of Distributing Incoming Requests" earlier in this chapter.

**Determining How the Cluster Handles Inbound and Outbound Traffic**

Network Load Balancing cluster traffic is handled differently for inbound and outbound traffic. The differences between inbound client requests and outbound responses affect the network infrastructure and the use of switches and hubs.

Inbound cluster traffic is sent to all cluster hosts by using broadcasts or multicast traffic, which is required so that all cluster hosts can receive the inbound traffic. Switches learn where to send packets by watching the source packets that are sent from the computers directly connected to the switch. After the switch learns a computer's location, the switch sends subsequent packets to the same switch port.

From an inbound-traffic perspective, the behavior of the switch is similar to a hub, because inbound cluster traffic is sent to all switch ports. However, from an outbound-traffic perspective, the switch provides isolation by preventing other cluster hosts from seeing a response to a client from the cluster host that services the client request.

For more information about the behavior of Network Load Balancing traffic, see the Network Load Balancing Technical Overview link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

**Determining When to Include Switches or Hubs for Interconnecting Cluster Hosts**

For most networks, switches are the preferred technology for connecting network devices to one another. Although many existing network infrastructures have hubs, switches are typically used in new deployments.

Although Network Load Balancing works in most configurations of switches and hubs, some configurations allow for optimal performance, and they provide ease of maintenance and operations. The recommended configuration is to use switches for interconnecting cluster hosts. Configurations that include hubs are also supported, but they require more complex network infrastructure and configuration.

The supported configurations for switches and hubs include:

- Cluster hosts connected to a switch that is dedicated to the cluster.

- Cluster hosts connected to a switch that is shared with other devices.

- Cluster hosts connected to a hub that is connected to a switch.

**Cluster hosts connected to a switch that is dedicated to the cluster**
In this configuration, the cluster hosts connect to a switch that is dedicated to the cluster. Because the inbound cluster traffic is sent to all ports on the switch, the primary advantage of a switch for inbound traffic — segregating traffic to a limited number of switch ports — is lost.

However, for outbound traffic, only the cluster host responding to the client request is aware of the traffic, because the switch isolates the outbound traffic to the port connected to the responding cluster host. This reduces the congestion of traffic for all the cluster hosts that are connected to the switch.

**Cluster hosts connected to a switch that is shared with other devices**
In this configuration, the cluster hosts are connected to a switch that is shared with other devices. The switch sends inbound cluster traffic to the other devices as well, creating unnecessary traffic for the other devices.

You can isolate the inbound cluster traffic to only the cluster hosts by establishing a VLAN comprising the ports that connect to the cluster hosts. After you establish the VLAN, inbound cluster traffic is sent only to the cluster hosts and not to the other devices that are attached to the same switch. Using a VLAN to segregate inbound client traffic works for the unicast or multicast method of distributing incoming requests. You can also use the multicast method with IGMP to limit inbound cluster traffic only to the cluster hosts.

Because outbound cluster traffic is sent directly to the client that originated the request, only the cluster host responding to the request sees the outbound cluster traffic. All other cluster hosts, and the other devices connected to the same switch, are unaware of the outbound cluster traffic.

**Cluster hosts connected to a hub that is connected to a switch**
From the perspective of inbound cluster traffic, a switch provides no additional benefit over a hub, except for being able to define a VLAN that isolates inbound traffic from other devices attached to the switch. Because all cluster hosts receive the inbound cluster traffic, the switch sends all inbound traffic to all ports that are attached to the switch (or ports that are assigned to the same VLAN).

From the perspective of outbound cluster traffic, a switch provides a greater advantage, because it reduces the network contention for outbound cluster traffic. With a switch, only the cluster host originating the outbound cluster traffic is aware of the traffic, because outbound traffic is switched. Using a hub causes network contention for all cluster hosts. With a hub, all cluster hosts receive, and subsequently determine if they need to process, the outbound cluster traffic.

To eliminate the network contention caused by using a hub, install an additional network adapter in each cluster host for the purpose of responding to client traffic. Inbound requests are received through the cluster adapters in each cluster host, while the outbound responses are sent to the clients through the additional network adapter.

Because of the increased complexity of configuration (adding an additional network adapter to each cluster host) and network infrastructure (adding an additional switch to connect all the additional network adapters to the clients), this is not a recommended configuration. Because a switch is required for the additional network adapters, it is easier to eliminate the additional network adapters and connect the cluster adapters to a switch.

If you are considering adding network adapters to improve cluster performance, consider increasing the data rate of the cluster network adapters and corresponding network infrastructure, instead of adding network adapters. For example, to increase available network bandwidth, specify 100 megabits per second (Mbps) network adapters instead of 10 Mbps network adapters (along with the appropriate upgrades to intermediary switches and routers). For more information about increasing the available network bandwidth to the cluster, see "Increasing Available Network Bandwidth to the Cluster" later in this chapter.

# Example: Specifying Cluster Network Connectivity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization has e-commerce Web applications that are accessed by users on the Internet. The organization's design includes Network Load Balancing to eliminate any application outages and improve performance. The e-commerce Web applications, running on IIS 6.0 and Windows Server 2003, will reside in the organization's perimeter network, which is located between the Internet and the organization's private network.

Figure 8.8 illustrates the existing firewalls, Internet connectivity, and connectivity to the organization's private network.

**Figure 8.8 Existing Network Diagram Before the E-Commerce Web Application Solution**




The e-commerce Web application requires support for 1,500 simultaneous users and an acceptable data transfer rate of 10 kilobits per second (Kbps) for each user, for a total aggregate data rate of 15 Mbps.

The organization has conducted lab testing on the devices in the organization's design, including routers, switches, network segments, and servers, and it has determined their capacity as listed in Table 8.16.

**Table 8.16 Device Capacity Results from Lab Testing**

| Device | Results Verified in the Lab |
|---|---|
| IIS server | • Supports 200 simultaneous users.<br>• Provides total aggregate data rates up to 50 Mbps. |
| Router | • Supports virtual IP subnets (VLANs).<br>• Provides total aggregate data rate of 20 Mbps.<br>• Requires manual ARP registration of a unicast IP address with a multicast MAC address. |
| Switch | • Supports virtual IP subnets (VLANs).<br>• Provides total aggregate data rate of 60 Mbps. |

Figure 8.9 illustrates the organization's network infrastructure after including the e-commerce Web solution.

**Figure 8.9 E-Commerce Web Solution with Network Load Balancing**




Table 8.17 lists the design decisions in specifying cluster network connectivity and the reasons for making the decisions.

**Table 8.17 Cluster Network Connectivity Design Decisions and Their Justification**

| Decision | Reason for the Decision |
|---|---|
| Include eight IIS servers (cluster hosts). | Each IIS server supports a maximum of 200 users, and eight servers are required to achieve 1,500 simultaneous users. |
| Specify unicast mode for distributing incoming client requests. | Network infrastructure supports unicast mode. |
| Specify VLAN-01 between Switch-01 and Switch-02. | All hosts on the cluster must belong to the same IP subnet. |
| Place cluster hosts on Switch-01 and Switch-02. | Bandwidth of a single switch is insufficient to handle cluster host traffic. |

# Ensuring Ease of Cluster Management and Operations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After your Network Load Balancing cluster is deployed, the operations staff in your organization takes primary responsibility for the day-to-day operations of the cluster. In addition to normal administration tasks required to maintain the operating system, such as applying service packs and upgrading the operating system, over the IT life cycle of your cluster, the operations team also maintains and upgrades applications on the cluster.

Follow the guidelines presented in Table 8.18 to help create Network Load Balancing clusters that are easier to manage and operate.

**Table 8.18 Guidelines for Designing NLB Clusters for Easy Management**

| Guideline | Explanation |
|---|---|
| Include sufficient cluster hosts to support maintenance and failover needs. | Cluster management is easier if enough cluster hosts exist to support client requests when one of the hosts is offline due to maintenance, upgrade, or failure. During lab testing, determine the number of cluster hosts required to support the required client traffic, and then add at least one additional cluster host. |
| Create a network infrastructure design that can accommodate a change in the number of cluster hosts. | The network infrastructure design must be flexible enough to allow the removal and addition of cluster hosts. When creating the supporting network infrastructure design, include additional ports on the switches or hubs connecting cluster hosts to support the addition of cluster hosts, as required.<br><br>By adding additional network infrastructure support, cluster hosts can be added easily by the operations team without redesigning or recabling the existing network infrastructure. |

During the IT life cycle of the cluster, upgrades are performed on the cluster by the operations team. The operations team typically performs these as *rolling upgrades*, by upgrading individual cluster hosts, one at a time, until the entire cluster is upgraded. For more information about performing rolling upgrades, see "Deploying Network Load Balancing" in this book.

Scaling out improves cluster performance by adding cluster hosts to the cluster and distributing client traffic across more cluster hosts. For more information about scaling your cluster, see "Scaling NLB Solutions" later in this chapter.

# Example: Ensuring Ease of Cluster Management and Operations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization has five mission-critical applications that must be run by remote users. These remote users connect to the organization's private network through a VPN tunnel. The remote users run the applications on servers running Terminal Services.

Because of the number of simultaneous remote users, three servers running Terminal Services are required. The number of servers was determined during the lab-testing phase of the solution. To provide load balancing and fault tolerance, the servers are part of a Network Load Balancing cluster.

The number of remote users that simultaneously access the Terminal Services farm is expected to increase over the lifetime of the farm. Also, the applications running on the farm need to be upgraded periodically, and this requires that the servers be restarted.

Table 8.19 lists the design decisions for ensuring the cluster is easier to maintain and operate and the reasons for making the decisions.

**Table 8.19 Cluster Maintenance and Operations Decisions and Their Justification**

| Decision | Reason for the Decision |
|---|---|
| Include an additional cluster host. | Ensures that the upgrades to individual servers in the Terminal Services farm can be performed without affecting client response times. |
| Provide unused ports on Switch-01. | Provides expansion for the addition of cluster hosts (servers running Terminal Services). |

Figure 8.10 illustrates the network infrastructure after the optimization of the design to ensure ease of cluster management and operations.

**Figure 8.10 Infrastructure Optimized for Ease of Cluster Management and Operations**



# Securing NLB Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The fact that your solution requires Network Load Balancing indicates that the applications and services are critical to the operations of your organization. The stability and integrity of your solution depend on how secure you make the resources that are used by the applications services. Besides protecting confidential data, the security that you include in your design ensures the stability of your applications and services. Use the features of Windows Server 2003 to protect the resources and applications running on cluster hosts. Plan security for your network load balancing solution immediately after determining the core specifications for your Network Load Balancing design, as illustrated in Figure 8.11.

**Figure 8.11 Securing Network Load Balancing Solutions**




**Note**

- For a Word document to assist you in documenting your security settings, see "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

One of the most important methods of protecting resources and applications on Network Load Balancing clusters is to ensure that only authorized users can remotely manage your clusters. To remotely manage clusters, users must be members of the local Administrators group on every cluster host in the cluster.

**Note**

- Users who are members of the local Administrators group can administer Network Load Balancing locally, regardless of other restrictions.

Ensure that only authorized users can remotely manage clusters by following these practices:

- Specify that remote management by using Nlb.exe be disabled on all cluster hosts.

  By default, remote management through Nlb.exe is disabled on clusters. Remote administration by using Nlb.exe is not recommended, because this method presents many security risks, including the possibility of data tampering, denial-of-service attacks, and information disclosure.

  If you choose to enable remote control for Nlb.exe, restrict access by specifying a strong remote-control password. Also, include a firewall to protect the Network Load Balancing UDP ports (the ports that receive remote control commands) by allowing access only to authorized client computers.

- Specify that Network Load Balancing Manager be used for all administration.

- Specify that administration be performed only from a secure, trusted computer that is within your organization's private

network or that is connected through a VPN remote access connection.

In addition to the security provided by restricting the administration of Network Load Balancing, you can further protect the applications, services, and confidential data on the cluster by using other Windows Server 2003 security features. These security features are specific to the applications and services running on the cluster. For more information about the security features provided by the applications and services running on the cluster, see the security content for the documents listed in "Additional Resources for Deploying Network Load Balancing" later in this chapter.

# Example: Securing NLB Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization has e-commerce Web applications that are accessed by Internet users. The design includes Network Load Balancing to eliminate any application outages and improve performance. The e-commerce Web applications, running on IIS 6.0 and Windows Server 2003, resides in the organization's perimeter network, which is located between the Internet and the organization's private network.

Figure 8.12 illustrates the e-commerce Web farm network design after the placement of the following components:

- Network infrastructure, including firewalls, routers, switches, and network segments
- IIS 6.0 servers into the perimeter network
- Active Directory® directory service domain controllers
- A computer running Windows Server 2003 and SQL Server 2000 on a server cluster
- A computer running Windows Server 2003 and file services on a server cluster

**Figure 8.12 IIS 6.0 E-Commerce Web Farm Solution**




Table 8.20 lists the design decisions that were made to protect the applications and resources on the cluster and the reasons for making those decisions.

**Table 8.20 Cluster Security Design Decisions**

| Decision | Reason for the Decision |
|---|---|
| IIS 6.0 servers are placed behind Firewall-01 and Firewall-02. | Prevents unwanted client traffic by filtering unnecessary client traffic. |
| Isolated network for communication is placed between the cluster and other servers in organization. | Prevents users outside the organization from viewing communications between the cluster and other servers in the organization. |

| | |
|---|---|
| Remote management is disabled. | Prevents unauthorized users from managing the cluster. |

# Scaling NLB Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

One of the primary reasons for including Network Load Balancing in your network design solution is to provide the ability to scale out your network by adding additional servers. Any other methods of improving client response times are not directly related to Network Load Balancing, but rather to the network infrastructure and system hardware associated with the cluster. Figure 8.13 illustrates the steps involved in incorporating scaling options in your Network Load Balancing solution.

**Figure 8.13 Scaling Network Load Balancing Solutions**




Also, you can ensure the scalability of applications and services by using methods that are specific to the applications and services running on the cluster. For more information about improving the scalability of services running on Network Load Balancing, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

**Note**

- Document your design on your organization's network diagram. For a Word document to assist you in documenting your Network Load Balancing design decisions, see "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

# Scaling Up the Solution by Adding System Resources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

During the lab testing and pilot testing of your Network Load Balancing design, you might discover that the cluster hosts are not responding within the requirements for your solution. Also, as you anticipate the increase in growth that your design must support, the cluster hosts need to continue to respond in accordance with your solution requirements. When a cluster becomes unable to respond the client requests within a specified time, you can improve the performance by *scaling up* your solution.

Take the following actions to scale up your solution:

- Increase system resources (such as processors, memory, disks, and network adapters) for the existing cluster host.
- Replace the existing cluster host with another system that has greater resources.

Scaling up is appropriate in the following situations:

- You want to reduce the total cost of ownership (TCO) by including fewer, more powerful cluster hosts in your design.
- The client response time is insufficient for a cluster host that is supporting the required number of users.

Scale up your solution by adding system resources by completing the following steps:

1. Identify the applications or services running on the Network Load Balancing cluster.
2. Identify the system resources that are consumed by the applications or services running on the cluster.

   Because Network Load Balancing requires only minimal system resources, the applications and services running on the cluster determine the system resources that are consumed first. For more information about the system resource requirements for services running on Network Load Balancing, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.
3. Verify the system resource estimates in your deployment lab.
4. Modify the system resource requirements based on the data from your testing in the deployment lab.
5. Document the system requirements in your design.

**Note**

- You can reduce the system resource requirements for your cluster hosts and avoid scaling up by eliminating any unnecessary applications or services. For more information about the required applications or services in your solution, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

# Scaling Out the Solution by Adding Cluster Hosts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Another method that you can use to improve client response time is *scaling out*. You scale out your solution when you add cluster hosts to existing clusters.

Scaling out your solution is appropriate in the following situations:

- You want to reduce TCO by including more, less-expensive cluster hosts in your design.
- You cannot further upgrade the system resources of the existing cluster hosts.

**Tip**

- To further improve client response time, consider scaling up the cluster host system resources as you scale out.

Scale out your solution by adding cluster hosts to the cluster and completing the following steps:

1. Identify the applications or services that require scaling out.
2. Identify any port rules that are associated with the application or services.
3. Include the additional cluster hosts in the cluster.
4. When appropriate, divide the cluster hosts into separate clusters and use round robin DNS to distribute client traffic across the clusters.

# Scaling Out Within a Single Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Determine the number of cluster hosts that are required in your design by estimating the number of clients each host can support and then testing your estimates in your lab. The maximum number of clients that each cluster host can support depends on the applications and services running on the host. For more information about capacity planning for the applications and services running on Network Load Balancing clusters, see "Additional Resources for Designing Network Load Balancing" later in this chapter.

The maximum number of cluster hosts in a cluster depends on a number of factors, such as the network infrastructure or protocols used by the applications and services running on the cluster, in addition to any Network Load Balancing limitations.

The only method of determining the maximum number of hosts in a cluster is by testing your configuration in a lab. When your lab tests indicate that you are unable to support the number of simultaneous clients by scaling out within a single cluster, scale out by using multiple clusters.

Network infrastructure devices, such as routers or switches, can also limit the network throughput to the cluster and ultimately affect how much you can scale out a cluster. To overcome this limitation, you can increase the network infrastructure capacity by changing the number and placement of routers and switches. However, the changes to the network infrastructure capacity might prevent you from keeping the cluster hosts on the same subnet or VLAN. Because all the cluster hosts must be on the same subnet or VLAN, scaling out with a single cluster is inadequate and you must scale out by using multiple clusters on separate subnets or VLANs.

# Scaling Out by Using Multiple Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are unable to scale out your applications on a single cluster, you can scale out by using multiple clusters. Then distribute the client traffic across the multiple clusters by using round robin DNS.

For example, an organization is getting ready to deploy a new high-volume Web site for its Internet presence. To support the required number of client requests, 72 cluster hosts are necessary. In lab testing, the organization determines that the network infrastructure can support only 9 clusters hosts per VLAN. As a result, eight clusters are required to scale out the organization's Web farm. Table 8.21 lists the clusters and the IP addresses that are assigned to each cluster.

**Table 8.21 Clusters and Their IP Addresses in Round-Robin DNS**

| Network Load Balancing Cluster | Assigned IP Address |
|---|---|
| IISNLB-01 | 10.0.1.100 |
| IISNLB-02 | 10.0.2.100 |
| IISNLB-03 | 10.0.3.100 |
| IISNLB-04 | 10.0.4.100 |
| IISNLB-05 | 10.0.5.100 |
| IISNLB-06 | 10.0.6.100 |
| IISNLB-07 | 10.0.7.100 |
| IISNLB-08 | 10.0.8.100 |

Figure 8.14 illustrates an IIS 6.0 Web farm in which client traffic is distributed across multiple Network Load Balancing clusters by using round robin DNS.

**Figure 8.14 Using Round Robin DNS to Distribute Client Traffic Across Clusters**




The DNS name of the IIS 6.0 application Web farm is www.contoso.com. The DNS zone for contoso.com contains multiple

entries for www.contoso.com, as required by round robin DNS. Each DNS entry corresponds to a Network Load Balancing cluster in the IIS 6.0 application Web farm. Round robin DNS distributes client traffic to the clusters, and then Network Load Balancing distributes client traffic within the cluster.

One of the assumptions in combining Network Load Balancing and round robin DNS is that an entire cluster is never taken offline. When an entire cluster is offline, DNS still directs client requests to the offline cluster and some clients will experience service outages, as they do when round robin DNS entries point to an individual server that is offline. To prevent clients from experiencing service outages, take only individual cluster hosts offline for maintenance or for other reasons.

For more information about round robin DNS, see "Deploying DNS" in *Deploying Network Services* of this kit.

# Increasing Available Network Bandwidth to the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Other steps in improving client response time directly affect the cluster hosts themselves. However, even with a highly optimized cluster, inadequate network bandwidth between the clients and the cluster can increase client response times to unacceptable levels.

Increase the available bandwidth between clients and the cluster by completing the following steps:

1. Identify the intermediary network segments, routers, and switches between the clients and the clusters.

2. Calculate the aggregate bandwidth requirements for the maximum number of simultaneous users.

3. Determine the intermediary network segments, routers, and switches between the clients and the clusters that are unable to support the aggregate bandwidth requirements.

4. Modify your design to increase the available bandwidth, based on the information in Table 8.22.

**Table 8.22 Increasing Bandwidth Based on Network Infrastructure Limitations**

| Available Bandwidth Limited By | Increase Bandwidth by One or More of These Methods |
|---|---|
| Network segment available capacity | • Increase the data rate of the network segment. (Upgrade from Integrated Services Digital Network (ISDN) to digital subscriber line (DSL) or T1.)<br>• Include additional network segments, and distribute the network traffic across the network segments. |
| Switch capacity | • Increase the data rate of the switch and the devices that are connected to the switch. (Upgrade from 10BaseT to 100BaseT or 1000BaseT.)<br>• Include additional switches, and distribute the computers that are connected to the original switch across the switches. |
| Router capacity | • Increase the data rate of the router and the devices that are connected to the router. (Upgrade from 10BaseT to 100BaseT or 1000BaseT.)<br>• Include additional routers, and create redundant routes to load balance the network traffic between the routers. |

# Example: Scaling NLB Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization provides VPN remote access to the organization's users through the Internet. The design includes Network Load Balancing to eliminate any application outages and to improve performance. The VPN remote access servers, which run Routing and Remote Access and Windows Server 2003, reside in the organization's perimeter network, which is located between the Internet and the organization's private network.

Pilot testing of the VPN remote access solution indicates the need for an increase in remote access client performance. Figure 8.15 illustrates the Routing and Remote Access VPN design to be tested.

**Figure 8.15 VPN Solution Pilot Test Environment**




Table 8.23 lists the results of the pilot test for each portion of the design that is illustrated in Figure 8.15.

**Table 8.23 Results of VPN Pilot Test**

| Design Portion Tested | Results |
|---|---|
| Network infrastructure | • Router-01 was saturated with client traffic.<br>• Switch-01 was saturated with client traffic.<br>• The network segment between Switch-01 and Router-01 was saturated.<br>• The connection to the Internet was saturated. |
| Cluster hosts hardware | • Individual cluster hosts lacked sufficient memory to handle the required number of simultaneous users.<br>• The total number of cluster hosts was insufficient to handle total number of simultaneous users. |

After the pilot test, the organization modifies the VPN remote access design. Figure 8.16 illustrates the modified version of the VPN design.

**Figure 8.16 Revised VPN Remote Access Design**




The organization made the following design decisions to improve scalability for the cluster:

- Add Firewall-02 and Router-02 to provide additional network bandwidth to the Internet and to solve the saturation problem on Router-01.
- Add Switch-02 to provide additional network bandwidth to the Internet and to solve the saturation problem on Switch-01.
- Add memory for all cluster hosts to ensure sufficient system resources for the hosts.
- Add VPN-07 and VPN-08 to support the maximum number of simultaneous users.

# Ensuring Availability in NLB Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By including multiple cluster hosts that provide the same applications and services, Network Load Balancing inherently provides fault tolerance. However, to provide a complete high-availability solution, your design must include more than just Network Load Balancing. The network infrastructure and system hardware that are associated with the cluster also affect the availability of the applications and services running on the cluster. In addition, include application-level monitoring, such as the monitoring provided by MOM or Application Center 2000, to ensure that applications are operating correctly. Ensuring availability is the final task in designing a network load balancing solution, as shown in Figure 8.17.

**Figure 8.17 Ensuring Availability in Network Load Balancing Solutions**




Include the following items to ensure high availability for clients that access applications and services on the cluster:

- Cluster hosts with fault-tolerant hardware
- Signed device drivers and software only
- Fault-tolerant network infrastructure

Also, you can improve the availability of applications and services by using methods that are specific to the applications and services running on the cluster. For more information about improving the availability of services running on Network Load Balancing, see "Additional Resources for Designing Network Load Balancing" later in this chapter.

**Note**

- After you design the specifications for ensuring availability, document your decisions. For a Word document to assist you in recording your decisions, see "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

#### Including Fault-Tolerant Hardware on Cluster Hosts

The cluster host hardware that you specify in your design can affect the uptime of the applications and services in your solution. Including system hardware with a longer mean time between failure (MTBF) can ensure that you experience fewer cluster host failures. In addition, including cluster hosts with fault-tolerant hardware can prevent unnecessary outages in your cluster.

For more information about including fault-tolerant hardware in your design, see "Planning for High Availability and Scalability" in this book.

#### Including Signed Device Drivers and Software Only

Another method of improving application and services uptime is to include only signed device drivers and software on the cluster hosts. Drivers and software that are signed have been certified by Microsoft, your organization, or third-party

companies that your organization trusts. Because unstable drivers and software can affect cluster uptime, including only signed device drivers and software helps ensure the stability of the cluster.

You can specify Group Policy settings in Active Directory to centrally configure the cluster hosts for the appropriate driver signing settings. When you are unable to specify driver signing by using Active Directory, specify the Local Security policies for each cluster host.

For more information about signed device drivers and software, see "Driver signing for Windows" in Help and Support Center for Windows Server 2003. For more information about specifying Group Policy settings, see "Designing a Group Policy Infrastructure in *Designing a Managed Environment* of this kit.

**Including a Fault-Tolerant Network Infrastructure**

Even if you perform all the previous steps to ensure fault tolerance for improving application and services uptime, your solution is not complete. Even with a highly optimized cluster, failures in the network infrastructure between the clients and the cluster can reduce uptime for applications and services.

To include a fault-tolerant network infrastructure between the clients and the cluster, complete the following steps:

1. Identify the intermediary network segments, routers, and switches between the clients and the cluster.

2. Determine if any of the intermediary network segments, routers, and switches between the clients and the cluster are potential points of failure that can cause application and services outages.

3. Modify your design to provide a fault-tolerant network infrastructure, based on the information in Table 8.24.

**Table 8.24 Providing Network Infrastructure Fault Tolerance Based on Limitations**

| Potential Failure Points | Include Any of These Fault-Tolerance Solutions |
|---|---|
| Network connection failure | Redundant network connections to provide fault tolerance in the event that a network connection fails. For example, if you are connected to the Internet by a single T1 connection, a failure of the T1 connection would prevent clients from accessing the cluster. Specify a redundant T1 connection to help prevent this type of failure. |
| Switch failure | Redundant switches to provide fault tolerance in the event that a switch fails. |
| Router failure | Redundant routers and redundant routes to provide fault tolerance in the event that a router fails. |

# Example: Ensuring Availability in NLB Solutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization provides VPN remote access to the organization's users through the Internet. The organization's design includes Network Load Balancing to eliminate any application outages and improve performance. The VPN remote access servers, running Routing and Remote Access and Windows Server 2003, reside in the organization's perimeter network, which is located between the Internet and the organization's private network.

The design includes ISA Server, which protects the VPN remote access servers in the perimeter network. The ISA Server servers are in a cluster (ISANLB-01) that provides load balancing and fault tolerance.

During the pilot testing of the Web content caching solution, with ISA Server, the deployment team experiences a number of outages that affect the entire solution. Figure 8.18 illustrates the Web content caching design, incorporating ISA Server, that is tested.

**Figure 8.18 VPN Remote Access Test Environment**




Table 8.25 lists the results of the pilot test for each portion of the design that is illustrated in Figure 8.18.

**Table 8.25 Results of VPN Pilot Test**

| Design Portion Tested | Results |
|---|---|
| Network infrastructure | • A failure of Router-01 resulted in total outage of VPN services.<br>• A failure of Switch-01 resulted in total outage of VPN services.<br>• A total outage of VPN services occurred for failure of the network segment between Router-01 and the Internet, or between Switch-01 and Router-01. |

| | |
|---|---|
| Cluster host hardware | • A failure of a disk drive in a cluster host resulted in a total cluster host failure.<br>• Network adapters in the cluster hosts have unsigned device drivers. |

After the pilot test, the VPN remote access design is modified. Figure 8.19 illustrates the modified version of the VPN design.

**Figure 8.19 Revised VPN Remote Access Design**




Table 8.26 lists the design decisions that the organization makes to improve the uptime for the VPN remote access solution and the reasons for making those decisions.

**Table 8.26 Improving VPN Solution Uptime Design Decisions and Their Justification**

| Decision | Reason for the Decision |
|---|---|
| Added Router-02 and additional Internet connection. | Provides a redundant route path to the Internet in the event that Router-01 or the corresponding Internet connection fails. |
| Added Switch-02. | Provides redundant paths in the event that Switch-01 fails. |
| RAID disk controllers were used in each cluster host. | Provide disk fault tolerance to help prevent disk failures and cluster host failures. |
| Group Policy was established to allow cluster hosts to load signed device drivers. | Provides trusted software to help ensure a stable environment, and prevents cluster host failure. |

# Additional Resources for Designing Network Load Balancing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Related Information**

- "Deploying Dial-up and VPN Remote Access Servers" in *Deploying Network Services* in this kit.
- "Deploying ISA Server" in *Deploying Network Services* in this kit.
- "Hosting Applications with Terminal Server" in this book.
- "Planning for High Availability and Scalability" in this book.
- "Deploying DNS" in *Deploying Network Services* in this kit.
- "Ensuring Application Availability" in *Deploying Internet Information Services (IIS) 6.0* in this kit, or see "Ensuring Application Availability (IIS 6.0)" at the Microsoft Web site.
- The Network Load Balancing Technical Overview link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources for more information about Network Load Balancing.
- The Session Directory and Load Balancing Using Terminal Server link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.
- Article 193602, "Configuration Options for WLBS Hosts Connected to a Layer 2 Switches," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at www.microsoft.com/windows/reskits/webresources.

**Related Job Aids**

- "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Network Load Balancing parameters" in Help and Support Center for Windows Server 2003.
- "Driver signing for Windows" in Help and Support Center for Windows Server 2003.

# Deploying Network Load Balancing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After completing the design for the applications and services in your Network Load Balancing cluster, you are ready to deploy the cluster running the Microsoft® Windows® Server 2003 operating system in your pilot and production network environments. A successful deployment ensures that your Network Load Balancing cluster meets or exceeds the specifications in the design. In addition, you must ensure that the deployment of your Network Load Balancing cluster does not disrupt the operation of any existing applications or services.

**In This Chapter**


Overview of the NLB Deployment Process

Selecting the Automated Deployment Method

Implementing a New Cluster

Upgrading Existing Clusters

Migrating Existing Clusters

Additional Resources for Deploying Network Load Balancing


**Related Information**

- For more information about the Network Load Balancing design process, see "Designing Network Load Balancing" in *Planning Server Deployments* of this kit.
- For more information about designing server clusters, see "Designing and Deploying Server Clusters" in *Planning Server Deployments*.

# Overview of the NLB Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A Network Load Balancing cluster comprises multiple servers running any version of the Windows Server 2003 family of operating systems, including Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; and Windows® Server 2003, Web Edition.

Clustering allows you to combine application servers to provide a level of scaling and availability that is not possible with an individual server. Network Load Balancing distributes incoming client requests among the servers in the cluster to more evenly balance the workload of each server and prevent overload on any one server. To client computers, the Network Load Balancing cluster appears as a single server that is highly scalable and fault tolerant.

The Network Load Balancing deployment process assumes that your design team has completed the design of the Network Load Balancing solution for your organization and has performed limited testing in a lab. After the design team tests the design in the lab, your deployment team implements the Network Load Balancing solution first in a pilot environment and then in your production environment.

Upon completing the deployment process presented here, your Network Load Balancing solution (the Network Load Balancing cluster and the applications and services running on the cluster) will be in place. For more information about the procedures for deploying Network Load Balancing on individual servers, see "New ways to do familiar tasks" in Help and Support Center for Windows Server 2003, and then click "Network Load Balancing Manager".

**Note**

- As you implement your Network Load Balancing design, use the information for each cluster host recorded by your design team in the "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

# NLB Deployment Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deploy your Network Load Balancing solution by implementing the decisions made by the design team. The process for deploying Network Load Balancing clusters is shown in Figure 9.1.

**Figure 9.1 Deploying Network Load Balancing**




During the Network Load Balancing design process, the design team created and documented the design of the network configuration by using drawing and diagramming software (such as Microsoft® Visio®), and the "NLB Cluster Host Worksheet" job aid that your design team completed for each cluster host. (This worksheet is associated with "Designing Network Load Balancing" in *Planning Server Deployments* of this kit.) The Network Load Balancing design documents describe the number of Network Load Balancing clusters, the placement of the Network Load Balancing clusters in your network environment, the number of Network Load Balancing cluster hosts, and other specifications.

Before deploying Network Load Balancing in your production environment, test your deployment solution in a lab and in a pilot deployment. For more information about testing your design, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

After you test your design, you can use different options for deploying each Network Load Balancing cluster, based on the requirements of your organization. Table 9.1 lists deployment requirements and the corresponding options for deploying Network Load Balancing clusters. For each cluster in the Network Load Balancing design, select one of these options for deploying Network Load Balancing clusters.

**Table 9.1 Deployment Requirements and Corresponding Options for Deploying NLB Clusters**

| Deployment Requirements | Implement | Upgrade | Migrate |
|---|---|---|---|
| Fresh installation of Windows Server 2003 is required | • | | • |
| Clusters are to be deployed on new cluster hardware | • | | • |
| Two or more existing clusters to be consolidated into a single cluster | | | • |
| A corresponding cluster currently exists in the network environment | | • | • |
| Existing hardware is capable of running Windows Server 2003 | | • | |

| | | | |
|---|---|---|---|
| Existing Windows registry settings to be retained | | ● | |

When your Network Load Balancing solution includes multiple Network Load Balancing clusters, you might need to use more than one method for deploying Network Load Balancing clusters. For each Network Load Balancing cluster in your solution, read the corresponding section in this chapter for the deployment method you select:

- Implementing New Clusters
- Upgrading Existing Clusters
- Migrating Existing Clusters

# Selecting the Automated Deployment Method

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Begin the deployment of your Network Load Balancing solution by determining how to automate your Network Load Balancing deployment. Automate the deployment of your Network Load Balancing solution to ensure the consistency of installation, to reduce the time required to deploy your solution, and to assist in restoring failed Network Load Balancing cluster hosts. Manually deploy your Network Load Balancing solution only if creating and testing the automation files and scripts would take more time than manually configuring the cluster. Figure 9.2 shows the process for determining the best method to automate the deployment of your Network Load Balancing solution.

**Figure 9.2 Selecting the Automated Deployment Method**




You can automate the deployment of Windows Server 2003 and Network Load Balancing cluster hosts by using one of the following methods:

- Unattended installation
- Sysprep
    - Remote Installation Services (RIS)

Table 9.2 compares the characteristics of the various automated deployment methods.

**Table 9.2 Comparing Automated Deployment Methods**

| Deployment Characteristics | Unattended Installation | Sysprep | RIS |
|---|---|---|---|
| Uses images to deploy installation | | ● | ● |
| Uses scripts to customize installation | ● | ● | |
| Supports post installation scripts to install applications | ● | ● | ● |
| Deployed by using local drives on the target server | ● | ● | |

| | | | |
|---|---|---|---|
| Initiated by remotely administered servers | | | ● |
| Deployed from a network share | ● | ● | |

Depending on the requirements of each cluster, more than one method might be required to deploy all your Network Load Balancing clusters. For more information about these methods, see "Designing Unattended Installations" "Designing Image-based Installations with Sysprep" and "Designing RIS Installations" in *Automating and Customizing Installations* of this kit.

The unattended installation and Sysprep methods require that you create script files to automate the deployment. The Network Load Balancing–specific installation and configuration information is required in two sections of the unattended and Sysprep script files: the [MS_WLBS parameters] and [MS_TCPIP parameters] sections.

For more information about the process for configuring the script files, see the *Microsoft® Windows® Server 2003 Corporate Deployment Tools User's Guide* (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD. For more information about the parameters used by the script files, see the Windows Preinstallation Environment link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. There are no special considerations for automating Network Load Balancing deployment by using the RIS method.

For Network Load Balancing, you need to perform specific tasks when creating the unattended installation and Sysprep script files:

- Review the content in the Microsoft Windows Preinstallation Reference that relates to the [*MS_WLBS parameters*] section.

  WLBS stands for "Windows NT Load Balancing Service," the name of the load balancing service used in Microsoft® Windows NT® Server version 4.0. For reasons of backward compatibility, WLBS continues to be used in certain instances.

- Ensure that the IP addresses for the cluster and all virtual clusters are entered in the IPAddress parameter under the [*MS_TCPIP parameters*] section.

  Typically, Network Load Balancing Manager automatically adds the cluster and virtual cluster IP addresses to the list of IP addresses. Both unattended installation and Sysprep require that you add the addresses to the IPAddress parameter under the [*MS_TCPIP parameters*] section of the script.

# Implementing a New Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Many of the mission-critical applications deployed within your organization will be new applications, which require you to deploy new Network Load Balancing clusters. The process for implementing a new cluster involves more than installing Windows Server 2003 and Network Load Balancing on the individual application servers. Implementing a new cluster might require additional network infrastructure, network services, file services, database services, and security services. Figure 9.3 shows the steps that you must complete before and after the implementation of the new cluster.

**Figure 9.3 Implementing a New Cluster**



# Preparing to Implement the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your new cluster is dependent upon the network infrastructure and other network services in your total solution. Ensure that these network infrastructure and other network services are deployed prior to implementing your cluster.

Prepare for the implementation of the new cluster by using the information documented in the "NLB Cluster Host Worksheet" and other documentation (such as Visio drawings of the network environment) that your design team completed for a specific cluster host during the design process. Coordinate with the operations team during this step in the process to review the changes that will occur in your organization's network environment.

To prepare for the implementation of the new cluster, complete the following tasks:

1. Implement the network infrastructure required by the cluster and by the applications and services running on the cluster.
2. Implement any networking services required by the applications and services running on the cluster.
3. Select the method for automating any additional Network Load Balancing configuration.

**Implementing the Network Infrastructure**

Before you implement the cluster, you must implement the network infrastructure that connects the cluster to client computers, to other servers within your organization, and to management consoles. The network infrastructure components include:

- Network cables
- Hubs
- Switches
- Routers
- Firewalls

**Note**

- Any references to switches in this chapter refer to Layer 2 switches.

When implementing the network infrastructure, make sure to have specifications about your current network environment available. Specifically, your hardware and software inventory and a map of network topology can be helpful. For more information about creating those documents, see "Planning for Deployment in *Planning, Testing, and Piloting Deployment Projects* of this kit.

**Implementing Any Required Networking Services**

Network Load Balancing is independent of the other Windows Server 2003 network services. As a result, you do not need to implement networking services for Network Load Balancing.

However, the applications and services running on the cluster can be dependent on other Windows Server 2003 networking services. For more information about requirements for implementing networking services used by the services and applications running on the cluster, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

**Automating Additional Configurations**

In many instances, it is possible that Windows Server 2003, Network Load Balancing, or the applications and services running on a particular cluster might require additional configuration after the installation is complete. You can use any combination of the following methods for automating these additional configurations:

- Microsoft Visual Basic Scripting Edition (VBScript)

- Windows Management Instrumentation (WMI)
- Active Directory Service Interfaces (ADSI)

# Example: Preparing to Implement the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization is deploying a new Web application that will be accessed by a large volume of Internet users. Because of scaling and availability considerations, the organization will deploy the new, high-volume Web application on a Network Load Balancing cluster. Figure 9.4 illustrates the organization's network environment prior to the implementation of the new cluster.

**Figure 9.4 Network Environment Before Implementing New Cluster**




In addition to deciding that the new Web farm will run on a Network Load Balancing cluster, the design team also made the following configuration decisions:

- No single router, switch, or Internet Information Services (IIS) Web farm server failure will prevent users from running the Web application.
- Web application will store data in a clustered SQL server running Microsoft® SQL Server™ 2000 on a server cluster.
- Web application executables Active Server Pages (ASP), Hypertext Markup Language (HTML) pages, and other executable code will be stored on a file server running on a server cluster.
- Accounts used for authenticating Internet users will be stored in Active Directory® directory service.

As the first step in the organization's deployment of the Web application, the IIS 6.0 Web farm, and Network Load Balancing, the organization must restructure the network infrastructure to support the new Web farm and cluster. Figure 9.5 illustrates the organization's network environment after preparing for the implementation.

**Figure 9.5 Network Environment After Preparing to Implement a New Cluster**




Table 9.3 lists the deployment steps that were performed prior to the implementation of the new cluster and the reasons for performing those steps.

**Table 9.3 Deployment Steps Prior to Implementation of the New Cluster**

| Deployment Step | Reason |
|---|---|
| Add Firewall-02 and Firewall-03 | Provide redundancy and load balancing. |
| Add Switch-01 and Switch-02 | Provide redundancy and load balancing. |
| Add network segments on Switch-01 and Switch-02 | Connect the IIS 6.0 Web farm to the network. |
| Configure Switch-01 and Switch-02 to belong to the same VLAN | Provide load balancing of client requests by using Network Load Balancing. |
| Add SQLCLUSTR-01 | Provide database support for the Web application on a Microsoft server cluster. |
| Add FILECLUSTR-01 | Provide secured storage for the Web application executables and content on a Microsoft server cluster. |
| Add DC-01 and DC-01 | Provide storage and management of user accounts used in authenticating Internet users. |

# Implementing the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To implement the new cluster, complete the following tasks:

1. Install and configure the hardware and Windows Server 2003 for each cluster host.
2. Install and configure the first Network Load Balancing cluster host.
3. Install and configure additional Network Load Balancing cluster hosts.

**Note**

- When your Network Load Balancing solution includes multiple Network Load Balancing clusters with round robin DNS, complete these tasks for each Network Load Balancing cluster in the solution. For more information about combining multiple Network Load Balancing clusters with round robin DNS, see "Scaling NLB Solutions" in "Designing Network Load Balancing" in this book.

# Installing and Configuring the Hardware and Windows Server 2003

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in performing the implementation of the new cluster is to install and configure the hardware and Windows Server 2003 for each cluster host. Install and configure all cluster host hardware at the same time to ensure that you eliminate any configuration errors prior to installing and configuring the Network Load Balancing cluster.

To install and configure Windows Server 2003 on the cluster host hardware, you must be logged on as a user account that is a member of the local administrators group on all cluster hosts. Install and configure the cluster host by using the information documented in the "NLB Cluster Host Worksheet" that your design team completed for that host during the design process.

To install and configure the hardware and Windows Server 2003 on each cluster host in the new cluster, complete the following tasks:

1. Install the cluster host hardware in accordance with the manufacturer's recommendations.

2. Connect the cluster host hardware to the network infrastructure.

3. Install Windows Server 2003 with the default options and specifications from the worksheet for the cluster host.

4. Install any additional services (such as IIS 6.0 or Routing and Remote Access) by using the design specifications for the service.

   For detailed instructions on installing additional services, see the resources related to the corresponding service in "Additional Resources for Deploying Network Load Balancing" later in this chapter.

5. Configure the TCP/IP property settings and verify connectivity for the cluster adapters.

   **Caution**

   - Configure the dedicated IP address at this time. The cluster IP address and any virtual IP addresses for port rules are added later in the deployment process through Network Load Balancing Manager.

6. If a separate management network is used, configure the TCP/IP property settings and verify connectivity for the management adapter.

   Although not required, it is recommended that you use a separate management network adapter to provide a communication path that is isolated both from the cluster adapter and from the clients. For more information about the benefits of including a management network adapter in your design, see "Selecting the Number of Network Adapters in Each Cluster Host" in "Designing Network Load Balancing" in this book.

7. Configure each server to be a member server in a domain created specifically for managing the cluster and other related servers.

   Although not required, creating a domain for management of the cluster provides a centralized method of controlling security to the cluster. Management of clusters installed in a workgroup is more difficult and time-consuming. When the cluster resides in a perimeter network, create a separate forest for the exclusive purpose of managing servers (including cluster hosts) in the perimeter network.

# Installing and Configuring the First Cluster Host

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have installed and configured the hardware for the cluster, you are ready to install and configure the first cluster host. The first cluster host acts as a master copy when you use an image-based deployment method (such as RIS, Sysprep, or a third-party product) to deploy the remaining cluster hosts.

An image-based deployment is faster and ensures consistency when implementing the remaining cluster hosts, by reducing or eliminating manual configuration. In addition, the same image-based deployment method can be reused after the deployment to restore failed cluster hosts.

Depending on the type of image-based deployment, different methods are used to customize the cluster host after the image has been restored. For example, Sysprep can allow you to interactively customize the image or use a configuration file — sysprep.inf — to customize the image. For more information about customizing the restored image, see "Designing Unattended Installations," "Designing Image-based Installations with Sysprep," and "Designing RIS Installations" in *Automating and Customizing Installations* of this kit.

Perform the following task on the first cluster host by using the "NLB Cluster Host Worksheet" that your design team completed for the first cluster host:

1. If you did not use the automated installation process to create the new cluster, start Network Load Balancing Manager and create a new cluster.

   **Tip**

   - You can start Network Load Balancing Manager by running Nlbmgr.exe.

2. Install the applications and services on the first cluster host.

   Examples of Windows Server 2003 services to be installed at this time include IIS or Terminal Services. For more information about installing Windows Server 2003 services, see the chapters that discuss those services in the *Microsoft® Windows® Server 2003 Deployment Kit*.

   Examples of applications to be installed at this time include, Web applications or Windows applications that run on Terminal Services. For more information about installing the applications running on your cluster, see the documentation that accompanies your application.

3. Enable monitoring and health checking on the first cluster host.

   A Microsoft Operations Manager (MOM) Management Pack exists for Network Load Balancing. When your organization uses MOM to monitor and manage the servers within your organization, include the MOM Management Pack for Network Load Balancing on the cluster hosts.

   For location of additional information about monitoring and health checking the applications and services running on the cluster, review the resources in "Additional Resources for Deploying Network Load Balancing" later in this chapter.

4. Verify that the first Network Load Balancing cluster host responds to client queries by directing requests to the cluster IP address.

   Test the first cluster host by specifying the cluster IP address or a virtual cluster IP address in the client software that is used to access the application or service running on the cluster. For example, a client accessing an IIS application would put the cluster IP address or virtual cluster IP address in the Web browser address line.

**Important**

- Create an entry in DNS for the cluster only after you have completed the deployment of the entire cluster. Prematurely

publishing the applications and services in DNS might result in overwhelming the cluster hosts before all cluster hosts are installed

# Installing and Configuring Additional Cluster Hosts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have installed and configured the first cluster host, you are ready to install and configure the remaining cluster hosts in the cluster. The first cluster host acts as a master copy when you use an image-based deployment method (such as RIS, Sysprep, or a third-party product) to deploy the remaining cluster hosts.

Perform the following tasks on the remaining cluster hosts by using the "NLB Cluster Host Worksheet" that your design team completed for each cluster host:

1. Create an image of the first cluster host that has just been deployed (discussed in the previous section) as required by one of the following image-based automated installation methods:

   - Sysprep

     For more information about creating Sysprep images, see "Designing Image-based Installations with Sysprep" in *Automating and Customizing Installations* of this kit.

   - RIS

     For more information about creating RIS images, see "Designing RIS Installations" in *Automating and Customizing Installations*.

   - Third-party products

     For more information about creating images with third-party products, see the documentation provided with the third-party image deployment software.

2. Restore the image of the first cluster host (created in step 1) to one of the remaining cluster host, following the directions provided in the documentation for the image-base installation method you used.

3. Configure any computer specific information (such as computer name and IP address) on the newly deployed cluster host.

4. Enable monitoring and health checking for the additional cluster host.

   Use the same methods as described for the first cluster host.

5. Verify that the additional cluster host responds to client requests.

   Use the same methods as described for the first cluster host.

6. Complete steps 2 through 5 for each remaining cluster host in the Network Load Balancing cluster.

7. Ensure that the cluster is load balancing requests across all cluster hosts (based on the port rules of the cluster).

The time required to create and test the images used in an image-based deployment can be prohibitive. It might take you less time to install and configure the remaining cluster hosts in the same way that you installed and configured the first Network Load Balancing cluster host. For example, you could deploy a cluster that consists of three cluster hosts. If you decide to deploy the cluster hosts using a method other than image-based deployment, you must ensure that you can restore a failed cluster host.

# Example: Implementing the New Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The organization mentioned in "Example: Preparing to Implement the Cluster" earlier in this chapter is now ready to implement the new IIS 6.0 Web farm that uses Network Load Balancing for load balancing and fault tolerance. The network infrastructure and additional networking services have been deployed in preparation for the implementation.

In this step, the organization installed and configured the first cluster host as a model for the remaining cluster hosts. Then the organization deployed the remaining cluster hosts by using an image-based deployment method. Figure 9.6 illustrates the network environment after the implementation of the new IIS 6.0 Web farm and Network Load Balancing.

**Figure 9.6 Network Environment After Installing the New Cluster**




Table 9.4 lists the deployment steps that were performed to implement the new cluster and the reasons for performing those steps.

**Table 9.4 Deployment Steps for Implementing the New NLB Cluster**

| Deployment Step | Reason |
|---|---|
| Add IIS-01, IIS-02, IIS-03, IIS-04, IIS-05, IIS-06, IIS-07, and IIS-08 server hardware. | Server hardware needs to be connected to network infrastructure in preparation for Network Load Balancing deployment. |
| Install Windows Server 2003 and Network Load Balancing on IIS-01 by using unattended installation. | Unattended setup is chosen because of the limited number of hosts to be deployed. |
| Create an image of IIS-01 to use as a model for RIS deployment. | RIS allows the servers to be reimaged in the event of a server failure. |
| Deploy the image on IIS-02, IIS-03, IIS-04, IIS-05, IIS-06, IIS-07, and IIS-08. | Image deployment ensures a consistent configuration on all servers in the Network Load Balancing cluster. |
| Verify the Web farm responds to client requests. | Verification ensures that the Web farm is properly configured and that Network Load Balancing is load balancing. |

# Verifying the Cluster and Enabling Client Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The final step in implementing your new Network Load Balancing cluster is to ensure that you have properly implemented and configured your cluster. In "Installing and Configuring the First Cluster Host" earlier in this chapter, you verified your implementation of individual servers. Now you must ensure that the entire cluster is secure and is properly monitored and health checked.

To verify the cluster and enable client access, complete the following tasks:

1. Verify the cluster host restoration process.
2. Verify that identified security threats are mitigated.
3. Perform monitoring and health checking on the complete cluster.
4. Verify proper operation of applications and services running on the cluster.
5. Enable client access to the cluster.

# Verifying the Cluster Host Restoration Process

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before placing the cluster into a pilot or production environment, you need to verify cluster host restoration to ensure that you can properly restore a cluster host that has failed.

To verify the cluster host restoration process, complete the following tasks:

1. Remove a cluster host from the cluster by performing a drainstop on the cluster host.

   A drainstop prevents a cluster host from handling new client requests. While draining, a cluster host continues to complete any outstanding requests and remains in the cluster until all active requests are completed. Then the cluster host stops all cluster operation.

2. Remove all disk volumes and disk partitions on the cluster host.

3. Restore the cluster host based on the installation method selected earlier in the deployment process.

4. Restart the cluster host.

5. Verify that the System event log of each cluster host contains no errors and that the restored cluster host responds properly to client requests.

# Verifying That Identified Security Threats are Mitigated

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to verify that all the identified security threats are properly handled in your new Network Load Balancing solution.

To verify that the identified security threats are mitigated, complete the following tasks:

1. Connect a client computer to the network such that the clients access the cluster by using the same route path that a typical client computer would use to connect to the cluster.

   For example, when clients connect to the cluster through a series of firewalls and routers to connect to the cluster over the Internet, ensure the client computer used for testing connects to the cluster through the same firewalls and routers.

2. Log on to the client computer with the user rights identified in your security threats.

3. For each identified security threat, reproduce the steps that result in the security compromise of the cluster.

4. Document the results and report the findings to the design team.

5. With the assistance of the design team, resolve any outstanding security issues.

   **Important**

   - Resolve all security threats before proceeding further in the deployment process.

# Monitoring and Health Checking the Complete Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your next step in completing the implementation of your new cluster is to enable monitoring and health checking on the entire cluster. In "Installing and Configuring the First Cluster Host" and "Installing and Configuring Additional Cluster Hosts" earlier in this chapter, you enabled monitoring and health checking on individual cluster hosts. However, in this step you are ensuring that the cluster is monitored as a complete unit. Enable monitoring and health checking on the cluster before allowing users to access the cluster in a pilot or production environment.

As clients begin to access the applications and services in your cluster, continue to provide monitoring and health checking as described in "Installing and Configuring the First Cluster Host" earlier in this chapter. Verify that the cluster performs as expected with live client traffic.

After the deployment process is complete, ensure that your operations staff continues the monitoring and health checking process in their long-term operations processes as part of your ongoing operations.

# Verifying Proper Operations of Applications and Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before placing the complete cluster into a pilot or production environment, you need to verify that applications and services are running correctly on the cluster.

To verify proper operations of applications and services on the new cluster, complete the following tasks:

1. Temporarily connect a client computer to the same switch used by the cluster.
2. From the client computer, verify that the applications respond to client requests as expected.
3. Disconnect the client computer from the switch.

# Enable Client Access to the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your last step in the implementation of your new cluster is to allow clients access to the applications and services running on the cluster. Be sure that you successfully complete all previous steps in the process before enabling users to access the cluster in a pilot or production environment.

Enable client access to the applications and the services in the cluster by creating DNS entries. Users will access your applications and services by using user-friendly names or Uniform Resource Locaters (URLs), such as http://www.microsoft.com, which correspond to the individual applications or services on the Network Load Balancing cluster. The DNS entries allow the translation of the user-friendly name to at least one IP address. When round robin DNS is used for load balancing between clusters, you must create a DNS entry for each cluster.

Table 9.5 lists the criteria for determining the number of DNS entries required for your new cluster.

**Table 9.5 Criteria for Determining the DNS Entries for Your Cluster**

| Solution Includes One of the Following | Required DNS Entries |
|---|---|
| Only one Network Load Balancing cluster. | A DNS entry for the cluster and a DNS entry for each virtual cluster. |
| More than one Network Load Balancing cluster with client traffic distributed across Network Load Balancing clusters by using round robin DNS. | A round robin DNS entry for each cluster and a round robin DNS entry for each virtual cluster. |

# Example: Verifying the Cluster and Enabling Client Access

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The organization mentioned in the examples earlier in this chapter is now ready to complete the implementation of the new IIS 6.0 Web farm. The Web farm servers are implemented and basic connectivity is provided to the Web farm. The organization has verified that Network Load Balancing is distributing client traffic evenly and that all cluster hosts are servicing client requests.

In this step, the organization verifies that the Web farm and the cluster function as a whole. Although operation of each cluster host was verified during the implementation process, now the organization must ensure that the new Web farm meets or exceeds the design specifications established by your design team before enabling client access to the applications.

Figure 9.7 illustrates the network environment after the implementation of the new IIS 6.0 Web farm and Network Load Balancing.

**Figure 9.7 Network Environment After Implementing the New Cluster**




Table 9.6 lists the deployment steps that were performed to verify the new cluster and enable client access.

**Table 9.6 Verifying the New Cluster and Enabling Client Access**

| Deployment Step | Reason |
|---|---|
| IIS-03 taken off line. | Automatic failover to other Web servers must be proven. |
| IIS-02 restored. | Restoration process for Web servers must be proven. |
| Client attached to the Internet and security attacks performed. | Mitigation of security threats must be proven. |
| Monitoring and health checking enabled on IIS-01, IIS-02, IIS-03, IIS-04, IIS-05, IIS-06, IIS-07 and IIS-08. | • Proper operation of the IIS 6.0 Web farm must be verified during and after the implementation process to ensure that load balancing is occurring and that system resources are adequate for client requests. |

| | |
|---|---|
| A DNS entry created for the cluster IP address of NLBCluster-01. | Clients must have a DNS entry to access the Web applications running on the cluster. |

# Upgrading Existing Clusters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization has applications and services that currently run on the Windows NT 4.0 or Microsoft® Windows® 2000 operating systems, you can upgrade your existing WLBS or Network Load Balancing clusters to take advantage of the improved security and performance of Windows Server 2003 and Network Load Balancing. Figure 9.8 shows the process for upgrading an existing cluster.

**Figure 9.8 Upgrading an Existing Cluster**




The Network Load Balancing upgrade process assumes the following conditions:

- Your applications are running on an existing IIS 4.0 (on Windows NT 4.0) or IIS 5.0 (on Windows 2000).
- You are upgrading the operating system and services running on your existing application servers.
- The system resources of the computers in your existing application farm are sufficient, or can be upgraded, to support Windows Server 2003 and IIS 6.0.

You can upgrade a cluster by taking the entire cluster offline and upgrading all the hosts, or you can leave the cluster on line and perform a *rolling upgrade*. A rolling upgrade entails taking individual cluster hosts offline one at a time, upgrading each host, and returning the host to the cluster. You continue upgrading individual cluster hosts until the entire cluster is upgraded. A rolling upgrade allows the cluster to continue running during the upgrade.

The decision to use rolling upgrades is based on the applications and services running on your existing cluster. If the applications and services support rolling upgrades, then perform a rolling upgrade on the existing cluster. Otherwise, perform the upgrade process recommended for the applications and services running on the cluster. For more information on the upgrade process for the applications and services running on your cluster, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

Upgrade the cluster by using the information documented in the "NLB Cluster Host Worksheet" that your design team completed for each cluster host during the design process.

# Preparing to Upgrade the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Network Load Balancing runs independently of other networking services provided by Windows Server 2003. On the other hand, the applications and services running on the cluster can be dependent upon the network infrastructure and other network services in your existing environment. Prior to upgrading the existing cluster, deploy any network infrastructure components or networking services that are required by the applications and services running on the cluster.

Prepare to upgrade the cluster by performing the following tasks:

1. Verify that applications and services running on the cluster are compatible with Windows Server 2003.
2. Upgrade the network infrastructure as required by the applications and services running on the cluster.
3. Upgrade any networking services as required by the applications and services running on the cluster.

**Verifying Applications and Services Are Compatible with Windows Server 2003**
Before you upgrade the existing cluster, ensure that the applications and services running on the cluster are compatible with Windows Server 2003. For help in determining if your application is compatible with Windows Server 2003, use the Windows Application Compatibility Toolkit. To download the toolkit, see the Windows Application Compatibility link on the Web Resources page at http://www.microsoft.com/reskits/webresources.

**Upgrading Necessary Network Infrastructure**
Before you upgrade the existing cluster, ensure that the final configuration of the cluster can be supported by the network infrastructure that connects the cluster to client computers, to other servers within your organization, and to management consoles. Perform only the network infrastructure upgrades required by the applications and services running on the cluster. Avoid performing upgrades to the network infrastructure for other reasons at the same time that you are upgrading the cluster. This minimizes the number of changes to the environment and reduces the likelihood of problems occurring during the upgrade process.

The network infrastructure to upgrade includes the following components:

- Network cabling
- Hubs
- Switches
- Routers
- Firewalls

**Note**

- When performing this step, make sure to have specifications about your current network environment available for use. Specifically, your hardware and software inventory, and a map of network topology can be helpful. For more information about creating those documents, see "Planning for Deployment" in *Planning, Testing, and Piloting Deployment Projects* of this kit.

**Upgrading Any Required Networking Services**
Network Load Balancing is independent of the other Windows Server 2003 network services. As a result, no networking services upgrades are required for Network Load Balancing.

However, the applications and services running on the cluster can be dependent on other Windows Server 2003 networking services. For more information about requirements that the services and applications running on the cluster might have for upgrading networking services, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

# Example: Preparing to Upgrade the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An organization is upgrading their existing virtual private network (VPN) remote access solution. The existing solution has a VPN remote access server farm that supports Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP). Currently the VPN remote access server farm is running WLBS on Windows NT 4.0. Figure 9.9 illustrates the organization's network environment prior to preparing for the upgrade of the existing WLBS cluster that hosts the VPN remote access server farm.

**Figure 9.9 Network Environment Before Preparing to Upgrading the Cluster**




As the first step in upgrading the existing VPN remote access server farm, the organization installs any networking services required by the VPN remote access server farm. No networking infrastructure or networking services upgrades are required for the upgrade.

In the future, the organization plans to deploy Internet Authentication Service (IAS) servers to provide centralized management of remote access policies. However, the deployment of IAS is scheduled to occur after the upgrade to Windows Server 2003 to prevent any unnecessary complications during the upgrade process.

# Upgrading the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The upgrade of your existing cluster can be done one cluster host at a time (a rolling upgrade) or by taking the cluster offline to upgrade all the hosts at the same time (a nonrolling upgrade). It is recommended that you perform a rolling upgrade. Otherwise, perform the upgrade process recommended for the applications and services running on the cluster. For more information about the upgrade process for the applications and services running on your cluster, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

Determine if you can perform a rolling upgrade in your lab prior to your pilot or production deployments. If you can perform rolling upgrades, ensure that you perform the post-upgrade processes described in this section for each individual cluster host.

To perform a rolling upgrade on existing WLBS or Network Load Balancing clusters, complete the following tasks:

1. Prevent clients from accessing the cluster host to be upgraded by performing a drainstop on the cluster host.
2. Monitor client activity on the cluster host until all client activity ceases.
3. Upgrade the cluster host.
4. Verify that the applications and services are running correctly on the upgraded cluster host.
5. Add the cluster host back to the cluster.
6. Upgrade the remaining cluster hosts by performing steps 1 though 5 for each cluster host.

# Preventing Clients From Accessing the Cluster Host

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you upgrade a cluster host, you must ensure that no clients have active sessions running on the cluster host. During the upgrade process, the cluster host is still connected to the network infrastructure.

However, to prevent new clients from starting sessions or applications on the cluster host, perform a drainstop on the cluster host to be upgraded. Performing a drainstop prevents new clients from accessing the cluster while allowing existing clients to continue until they have completed their current operations.

In addition, configure the **Default state** of the **Initial host state** to **Stopped**. Configuring the cluster host in this manner ensures the cluster host cannot rejoin the cluster during the upgrade process. Because the upgrade process requires the cluster host to restart, you need to verify that the upgrade completed successfully before adding the cluster host back to the cluster.

**Tip**

- You can configure the Initial host state through Network Load Balancing Manager or the property settings of the cluster network adapter.

For more information about performing a drainstop on the source cluster and changing the **Default state** of cluster host, see "Create and Manage Network Load Balancing Clusters" in Help and Support Center for Windows Server 2003.

# Monitoring Client Activity on the Cluster Host

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you run the drainstop on the cluster host, monitor client activity on the cluster host to determine when clients are no longer using the cluster host. The method for determining when clients are no longer using the cluster host is specific to the applications and services running on the cluster host.

For example, for a VPN remote access solution, monitor for active VPN connections to the cluster host. When there are no active VPN connections, the cluster host is ready to be upgraded.

For more information about monitoring the applications and services running on the cluster for client activity, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

# Upgrading the Cluster Host

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you ensure that no clients are accessing the cluster host, upgrade the cluster host to Windows Server 2003. Ensure that during the upgrade process you apply the latest service packs and hotfixes.

After the upgrade is complete, check the upgrade logs (Setuperr.log in the *windir* folder of the system volume) to identify any problems that occurred during the upgrade process. Many applications and services running on the cluster have log files that identify problems that occur during the upgrade process. For example, IIS 6.0 creates a separate log that documents the upgrade process for IIS components. Make certain that you review any upgrade logs for the applications and services running on the cluster.

For more information about these upgrade logs, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

# Verifying the Applications and Services are Running Correctly

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After the upgrade of the cluster host is complete, you need to verify that the applications and services are running correctly on the cluster host. You need to do this before starting the cluster service again in order to allow the cluster host to rejoin the cluster. The process presented here is specific to Network Load Balancing. This process might apply to the applications and services running on the cluster, however, the applications and services running on the cluster might have a different verification process. For more information about the verification process for specific applications and services running on the cluster, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

To verify that the applications and services are running correctly after you upgrade the cluster host, complete the following tasks:

1. Temporarily connect a client computer to the same switch used by the cluster.

2. From the client computer, verify that the applications respond to client requests as they did prior to the upgrade.

3. Verify that the identified security threats are mitigated.

   The only action you need to take to mitigate Network Load Balancing–specific security threats is to ensure that unauthorized clients cannot remotely administer the cluster. Unless the network infrastructure, including firewalls or routers, changed significantly since the start of the upgrade, the Network Load Balancing security threats should still be mitigated.

   However, you must mitigate security threats that are unique to the applications and services running on the cluster. For more information about mitigating security threats for specific applications and services running on the cluster, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

4. Enable monitoring and health checking, if they are not already enabled.

   A Microsoft Operations Manager (MOM) Management Pack exists for Network Load Balancing. When your organization uses MOM to monitor and manage the servers within your organization, include the MOM Management Pack for Network Load Balancing on the cluster hosts.

   For more information about monitoring and health checking the applications and services running on the cluster, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

5. Disconnect the client computer from the switch.

# Adding the Cluster Host Back to the Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you complete the upgrade and verify that the applications and services are running correctly, add the cluster host back to the cluster. Because the cluster host is still connected by cable to the existing network infrastructure, adding the cluster host back to the cluster requires you to start Network Load Balancing on the cluster host.

To add the cluster host back to the cluster, use Network Load Balancing Manager to start the Network Load Balancing service on the cluster host. In addition, configure the **Default state** of the **Initial host state** to match the setting that existed prior to the upgrade process. Configuring the cluster host in this manner ensures that the clients accessing the applications and services running on the cluster do not encounter problems when you restart the cluster host and return it to the cluster.

The reason to configure the cluster host with the same settings, instead of selecting **Started**, is that Network Load Balancing loads very early in the operating system boot process. The cluster host can join the cluster before other services are running. This means the cluster host might receive load before other services, such as IIS, are started. Allowing Network Load Balancing to start before other services could result in the denial of service to users until the service starts.

You can use management software, such as MOM, to monitor the services running on the host and to start Network Load Balancing when the appropriate services are running. Using management software to start Network Load Balancing is recommended to prevent users from experiencing outages in service.

For more information on performing a drainstop on the source cluster and changing the **Initial host state** of cluster host, see "Create and Manage Network Load Balancing Clusters" in Help and Support Center for Windows Server 2003.

# Example: Upgrading the Existing NLB Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The organization mentioned in the examples earlier in this chapter is now ready to upgrade the existing VPN server farms. In this step, the organization performs a rolling upgrade on each cluster host in the VPN remote access server farm.

Each of the VPN remote access servers were upgraded in turn until the entire VPN remote access server farm was upgraded. Figure 9.10 illustrates the organization's network environment after the upgrade of the cluster.

**Figure 9.10 Network Environment After Upgrading the Cluster**




To upgrade the VPN remote access server farm, the following tasks were performed:

1. Prevented client access on the VPN-01 cluster host by performing a drainstop on VPN-01.
2. Monitored client activity on VPN-01 until all client activity ceased.
3. Upgraded VPN-01.
4. Verified that the applications and services are running correctly on VPN-01.
5. Added VPN-01 back to the cluster.
6. Performed steps 1 through 5 on the remaining VPN remote access servers.

# Migrating Existing Clusters

Updated: December 21, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Your organization might have applications and services that currently run on Windows NT 4.0 or Windows 2000 clusters. If you want these applications and services to take advantage of the improved security and performance of Windows Server 2003, and if you also want to run these applications on newly installed Windows Server 2003 servers, you can migrate the applications and servers from existing clusters to newly installed Network Load Balancing clusters.

| **Note** |
| --- |
| For information about upgrading an existing Network Load Balancing Cluster from Windows Server 2003 to Windows Server 2008 or Windows Server 2008 R2, see Upgrading an Existing Network Load Balancing Cluster. |

The migration process is appropriate in the following situations:

- You want to migrate one WLBS or Network Load Balancing cluster to a target Network Load Balancing cluster.
- You want to consolidate multiple WLBS or Network Load Balancing clusters onto a target Network Load Balancing cluster.

Migrating WLBS clusters from Windows NT 4.0 or Network Load Balancing clusters from Windows 2000 requires only minimal effort for Network Load Balancing. The majority of the migration effort involves the applications and services running on the cluster. The process for migrating an existing cluster to a target Network Load Balancing cluster running Windows Server 2003 is shown in Figure 9.11.

**Figure 9.11 Migrating Existing Clusters**




Migrate your WLBS or Network Load Balancing cluster by using the information documented in the "NLB Cluster Host Worksheet" that your design team completed for each cluster host during the design process.

**Network Load Balancing Migration Assumptions**

The Network Load Balancing migration process to Windows Server 2003 assumes the following conditions:

- Your applications and services are running on an existing cluster running WLBS (on Windows NT 4.0) or Network Load Balancing (on Windows 2000).
- You are installing a new Network Load Balancing cluster as the target cluster.

- Your applications and services are compatible with Windows Server 2003.

In addition to compatibility with Network Load Balancing, the applications and services must be compatible with Windows Server 2003. For help in determining if your application is compatible with Windows Server 2003, use the Windows Application Compatibility Toolkit on the *Windows Server 2003 Deployment Kit* companion CD.

For information about the migration process for the application or service running on your cluster, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

**Network Load Balancing Migrations and Consolidations**

In addition to migrating an existing cluster to the target cluster, you can use the Network Load Balancing migration process to consolidate multiple existing clusters to the target cluster. The consolidation process is identical to the migration process, except that, for each migration, you are migrating applications and services that formerly ran on multiple clusters onto a single target cluster.

As a part of consolidation, the port rules and cluster virtual IP addresses from all the source clusters must be created on the target cluster. For each cluster that you migrate, you must create port rules on the target cluster that handle the client traffic in the same manner as the traffic was handled on the source cluster. For more information about the migration of port rules, see "Migrating to the Target Cluster" later in this chapter.

# Implementing the Target Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in the migration process is to implement the target Network Load Balancing cluster. If the target Network Load Balancing cluster already exists, proceed to the next step in the process, described in "Migrating to the Target Cluster" later in this chapter

The process for implementing the target is identical to the process for installing a new Network Load Balancing cluster. Complete the process described in "Implementing a New Cluster" earlier in this chapter. Ensure that implementing the target cluster does not disrupt applications running on the source cluster. For example, avoid modifying the network infrastructure that actively connects to the source cluster and avoid modifying DNS entries that direct clients to the source cluster.

# Example: Implementing the Target Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso Pharmaceuticals is consolidating existing IIS Web farms onto a new IIS 6.0 Web farm. There are four independent IIS Web farms that need to be migrated. Two of the Web farms are running IIS 4.0 and WLBS on Windows NT 4.0. The other two Web farms are running IIS 5.0 and Network Load Balancing on Windows 2000.

During the migration, Contoso plans to consolidate the separate Web farms into one Web farm that supports all of their Web applications. Figure 9.12 illustrates Contoso's network environment prior to the migration of the existing WLBS and Network Load Balancing clusters that host the IIS Web farms.

**Figure 9.12 Network Environment Before Migrating the IIS Web Farms**




As the first step in Contoso's migration of the Web application server farm, Contoso restructured the network infrastructure to support the new IIS and Network Load Balancing deployment. Additional networking services were deployed in preparation for the implementation.

Then Contoso installed and configured the first server in the Web farm as a model for the remaining servers. Contoso deployed the remaining servers by using an image-based deployment method. Figure 9.13 illustrates Contoso's network environment after the implementation of the target IIS 6.0 Web farm (Network Load Balancing cluster).

**Figure 9.13 Network Environment After Installing the Target Cluster**




During the implementation of the new target cluster, the existing clusters were not affected. Clients continued to access the applications and services running on the existing cluster during the implementation of the target cluster.

# Migrating to the Target Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As discussed in "Implementing the Target Cluster" earlier in this chapter, the only part of Network Load Balancing to be migrated is the port rules. All other aspects of the migration of the cluster are specific to the applications and services running on the cluster. For more information about migrating applications and services running on the cluster, see "Additional Resources for Deploying Network Load Balancing" later in this chapter.

Because the number of port rules on a cluster is small, the existing port rules can be easily recreated on the target cluster. If you are consolidating multiple clusters to a target cluster, the creation of port rules is slightly different from the creation of port rules if you are migrating a single cluster to a target cluster. Primarily, when consolidating, you need to ensure that consolidated port rules do not conflict with one another.

Migrate or consolidate the port rules on the source cluster to the target cluster by completing the following steps:

1. Document the port rules on the source cluster.
2. Identify any differences between port rules that exist on the target cluster and the port rules on the source cluster.
3. Create the port rule on the target cluster, as required, with new or different virtual IP addresses.
4. Migrate the applications and services running on the source cluster.
5. Enable client access to the applications on the target cluster.

# Documenting Port Rules on the Source Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Document the port rules on the source cluster by completing the following steps:

1. View the port rules in Network Load Balancing Manager or by running Nlb.exe at a command prompt.

   To view port rules in Network Load Balancing Manager:

   a. Start **Network Load Balancing Manager**.

   b. On the **Host** menu, click **Properties**.

   c. On the **Port Rules** tab, view the port rules in the **Defined port rules** box.

   To view port rules by running Nlb.exe, in the command prompt, type **Nlb display *cluster_IP_address*** (where *cluster_IP_address* is the cluster IP address of the source cluster), and then press **Enter**.

   For more information about viewing port rules in Network Load Balancing Manager or by using Nlb.exe, see "Create and Manage Network Load Balancing Clusters" in Help and Support Center for Windows Server 2003.

2. Record the port rules on the job aid "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

# Identifying Differences Between Port Rules

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When migrating a source cluster to a target cluster, create port rules that are identical to the source cluster. When consolidating two or more clusters to a target cluster, create port rules that are identical to the source cluster unless port rules exist on the target cluster that specify the same TCP/UDP port number that is specified by a port rule on the source cluster.

During consolidation, two clusters might have port rules that relate to the same TCP/UDP port number, but have different affinity or load weight. In these situations you must differentiate between the different affinity or load weight behaviors by associating each of the source port rules with a different virtual cluster.

For example, you are migrating two clusters, Cluster A and Cluster B, to the same target cluster. Cluster A has a port rule for HTTP (TCP port 80) traffic that specifies no affinity, while Cluster B has a port rule for HTTP traffic that specifies Single affinity. The applications on these servers expect the affinity behavior. To ensure these applications behave as they did on the source servers, specify a port rule, with unique virtual IP address, for each of the affinity behaviors.

# Creating Port Rules on the Target Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Create the port rules on the target cluster by using Network Load Balancing Manager. Network Load Balancing Manager is recommended because it ensures that the port rules remain consistent across all the cluster hosts. Other methods require you to manually ensure that the port rules are consistent.

To create the port rules on the target cluster by using Network Load Balancing Manager:

1. Start **Network Load Balancing Manager**.
2. If **Network Load Balancing Manager** does not already list the cluster, connect to the cluster.
3. Right-click the cluster, and then click **Cluster Properties**.
4. Click the **Port Rules** tab.
5. Click **Add**.
6. Using the port rules documented in previous steps, specify values for the following:
   - **Cluster IP address**, which is the virtual IP address that you want this rule to apply to. Enter a specific virtual IP address to create a virtual cluster, or check **All** to apply the rule to all virtual IP addresses.
   - **Port range**
   - **Protocols**
   - **Filtering mode** and as appropriate, values for **Affinity**, and **Load weight**, and then click **OK**.
7. Repeat steps 5 and 6 until all port rules are created.

For more information about creating port rules in Network Load Balancing Manager, see "Create and Manage Network Load Balancing Clusters" in Help and Support Center for Windows Server 2003.

# Migrating the Applications and Services Running on the Source Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For information about the specific migration steps for the applications and services running on the source cluster, see "Additional Resources for Deploying Network Load Balancing" later in this chapter. After you migrate the applications and services, ensure that the applications and services are running correctly before enabling client access to the applications and services on the target cluster.

# Enabling Client Access to the Applications on the Target Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you have migrated the cluster port rules, applications, and services from the source cluster to the target cluster, you are ready to enable client access to the applications on the target cluster. Enable client access to the target cluster while maintaining the DNS entries to the source cluster. Over a period of time, remove the DNS entries that point to the source cluster.

Enable client access to the applications on the target cluster by completing the following steps:

1. Create the appropriate DNS entries for the applications running on the target cluster.

   For more information about how to create DNS entries for your applications see "Managing resource records" in Help and Support Center for Windows Server 2003.

2. Monitor client traffic to determine if clients are accessing the target cluster.

3. Over a period of time, such as a few hours or a day, confirm clients accessing target cluster are experiencing normal response times and application responses.

4. Remove the DNS entries pointing to the source cluster.

   For more information about how to remove DNS entries for your applications see "Managing resource records" in Help and Support Center for Windows Server 2003.

5. Prevent new clients from accessing the source cluster by performing a drainstop on the source cluster.

   For more information about how to perform a drainstop, see "Create and Manage Network Load Balancing Clusters" in Help and Support Center for Windows Server 2003.

6. Monitor client traffic to the source cluster to determine when clients are no longer accessing the source cluster.

7. When clients no longer access the source cluster, decommission the source cluster hardware.

# Example: Migrating to the Target Cluster

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The fictitious organization Contoso, mentioned in examples earlier in this chapter, is now ready to migrate their existing applications running on the four IIS Web farms to the target IIS 6.0 Web farm (and subsequently the target Network Load Balancing cluster). Contoso's target Web farm has been installed and their applications have been verified to work with Windows Server 2003 and Network Load Balancing.

Figure 9.14 illustrates Contoso's network environment after the migration of the existing applications to the target Web farm and Network Load Balancing cluster.

**Figure 9.14 Network Environment After Migrating Applications and Services**




Table 9.7 lists the applications, the Web farms that run the applications, and the DNS entries that correspond to the Web applications.

**Table 9.7 Applications, Web Farms, and DNS Entries**

| Application | Runs on | DNS Entries |
|---|---|---|
| Organization's Internet presence | • IISNLB-01<br>• IISNLB-02 | • www.contoso.com 10.0.0.111<br>• www.contoso.com 10.0.0.121 |
| Customer support site | • IISNLB-03 | • support.contoso.com 10.0.0.131 |
| E-commerce site | • IISNLB-03 | • sales.contoso.com 10.0.0.132 |
| Nonprofit third-party organization's Internet presence | • IISNLB-04 | • www.treyresearch.net 10.0.0.141 |

In this step, Contoso creates the port rules on the target cluster based on the port rules from each of the existing clusters

(IISNLB-01, IISNLB-02, IISNLB-03, and IISNLB-04). Table 9.8, Table 9.9, Table 9.10, and Table 9.11 lists the port rules for the existing clusters.

**Table 9.8 Port Rules for IIS Applications on IISNLB-01**

| Cluster IP Address | Start | End | Protocol | Filtering Mode | Load Weight | Affinity |
|---|---|---|---|---|---|---|
| All | 80 | 80 | TCP | Multiple | Equal | None |
| All | 443 | 443 | Both | Multiple | Equal | Single |

**Table 9.9 Port Rules for IIS Applications on IISNLB-02**

| Cluster IP Address | Start | End | Protocol | Filtering Mode | Load Weight | Affinity |
|---|---|---|---|---|---|---|
| All | 80 | 80 | TCP | Multiple | Equal | None |
| All | 443 | 443 | Both | Multiple | Equal | Single |

**Table 9.10 Port Rules for IIS Applications on IISNLB-03**

| Cluster IP Address | Start | End | Protocol | Filtering Mode | Load Weight | Affinity |
|---|---|---|---|---|---|---|
| All | 80 | 80 | TCP | Multiple | Equal | Single |
| All | 443 | 443 | Both | Multiple | Equal | Single |

**Table 9.11 Port Rules for IIS Applications on IISNLB-04**

| Cluster IP Address | Start | End | Protocol | Filtering Mode | Load Weight | Affinity |
|---|---|---|---|---|---|---|
| All | 80 | 80 | TCP | Multiple | Equal | None |

The port rules are recorded for each of the source clusters. Because each of the clusters has different port rules, virtual clusters are created to facilitate the consolidation of the applications onto the target cluster IISNLB-05. Table 9.12 lists the port rules that were created to support the applications on IISNLB-05.

**Table 9.12 Port Rules Created for IIS Applications on IISNLB-05**

| Cluster IP Address | Start | End | Protocol | Filtering Mode | Load Weight | Affinity | Created to Support |
|---|---|---|---|---|---|---|---|
| All | 80 | 80 | TCP | Multiple | Equal | None | All source clusters |
| All | 443 | 443 | Both | Multiple | Equal | Single | All source clusters |
| 10.0.0.153 | 80 | 80 | TCP | Multiple | Equal | Single | IISNLB-03 |

The port rules on IISNLB-01 and IISNLB-02 are combined into a pair of port rules on IISNLB-05 because IISNLB-01 and IISNLB-02 support the same application. The two clusters, IISNLB-01 and IISNLB-02, were created to provide scaling in the original design. Because the application — www.contoso.com — is being consolidated onto one Web farm on a single cluster, the port rules can be consolidated as well.

The virtual clusters on IISNLB-03 are created on IISNLB-05 to provide separate virtual IP addresses for support.contoso.com and sales.contoso.com.

A virtual cluster, 10.0.0.154, was created to support the application on IISNLB-04, www.treyresearch.net.

After the applications are migrated to the target cluster, IISNLB-05, and are verified for proper operation, DNS entries are created to direct clients to the applications running on the target cluster. Table 9.13 lists the applications, the Web farms that run the applications, and the DNS entries that correspond to the Web applications that include the target cluster, IISNLB-05.

**Table 9.13 Applications, Web Farms, and DNS Entries**

| Application | Runs on | DNS Entries |
|---|---|---|
| Organization's Internet presence | • IISNLB-01<br>• IISNLB-02<br>• IISNLB-05 | • www.contoso.com 10.0.0.111<br>• www.contoso.com 10.0.0.121<br>• www.contoso.com 10.0.0.151 |
| Customer support site | • IISNLB-03<br>• IISNLB-05 | • support.contoso.com 10.0.0.131<br>• support.contoso.com 10.0.0.152 |
| E-commerce site | • IISNLB-03<br>• IISNLB-05 | • sales.contoso.com 10.0.0.132<br>• sales.contoso.com 10.0.0.153 |
| Nonprofit third-party organization's Internet presence | • IISNLB-04<br>• IISNLB-05 | • www.treyresearch.net 10.0.0.141<br>• www.treyresearch.net 10.0.0.154 |

After a few days of operation, clients are experiencing no difficulty with IISNLB-05, so Contoso modifies the DNS entries to remove the original Web servers (IISNLB-01, IISNLB-02, IISNLB-03, and IISNLB-04). Table 9.14 reflects the DNS entries after the removal of the original Web servers.

**Table 9.14 Applications, Web Farms, and DNS Entries**

| Application | Runs on | DNS Entries |
|---|---|---|
| Organization's Internet presence | • IISNLB-05 | • www.contoso.com 10.0.0.151 |
| Customer support site | • IISNLB-05 | • support.contoso.com 10.0.0.152 |
| E-commerce site | • IISNLB-05 | • sales.contoso.com 10.0.0.153 |
| Nonprofit third-party organization's Internet presence | • IISNLB-05 | • www.treyresearch.net 10.0.0.154 |

To finalize the completion of the migration, Contoso performs the following tasks:

1. Contoso monitors the source clusters for client traffic.

2. After all client traffic ceases on the source clusters, Contoso decommissions the computers in the source clusters.

# Additional Resources for Deploying Network Load Balancing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Example: Migrating to the Target Cluster

**Related Information**

- "Deploying Dial-up and VPN Remote Access Servers in *Deploying Network Services* of this kit.
- "Deploying ISA Server in *Deploying Network Services*.
- "Deploying DNS" in *Deploying Network Services* for more information about deploying Network Load Balancing with round robin DNS.
- "Upgrading an IIS Server to IIS 6.0" in *Deploying Internet Information Services (IIS) 6.0*, or see "Upgrading an IIS Server to IIS 6.0" at the Microsoft Web site.
- "Migrating IIS Web Sites to IIS 6.0" in *Deploying Internet Information Services (IIS) 6.0*, or see "Migrating IIS Web Sites to IIS 6.0" at the Microsoft Web site.

**Related Job Aids**

- "NLB Cluster Host Worksheet" (Sdcnlb_1.doc) on the *Windows Server 2003 Deployment Kit* companion CD (or see "NLB Cluster Host Worksheet" on the Web at http://www.microsoft.com/reskit).

**Related Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Network Load Balancing Manager" under "New ways to do familiar tasks" in Help and Support Center for Windows Server 2003 for more information about the procedures for deploying Network Load Balancing on individual servers.
- "Create and Manage Network Load Balancing Clusters" in Help and Support Center for Windows Server 2003.
- "Managing resource records" in Help and Support Center for Windows Server 2003.

# Windows Server 2003 Performance Counters Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This searchable online file provides a reference to the performance counters used by System Monitor and Performance Logs and Alerts, hosted by the Performance console in Windows Server 2003. These counters collect and display information about system performance when you configure a graph, log, or alert.

This reference contains the following topics:

| Topic | Description |
|---|---|
| Overview of Performance Data Collection | Defines performance objects, counters, and instances. |
| Counters by Feature | Describes the performance counters that are associated with optional services or features. |
| Counters by Object | Describes performance counters, grouped by object. |
| Bottleneck-Detection Counters | Describes performance counters that you can use to identify bottlenecks in your system. |

# Overview of Performance Data Collection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Server 2003 family defines the performance data it collects in terms of *objects*, *counters*, and *instances*. A *performance object* is any resource, application, or service that can be measured. Using System Monitor and Performance Logs and Alerts, you can select performance objects, counters, and instances to collect and present data about the performance of system components or installed software.

Each object has *performance counters* that are used to measure various aspects of performance, such as transfer rates for disks or, for processors, the amount of processor time consumed. The object may also have an *instance*, which is a unique copy of a particular object type; not all object types support multiple instances.

This Help file describes objects, counters, and instances using the following syntax:

```
\\Computer_name\Object(ParentInstance/ObjectInstance#InstanceIndex)\Counter
```

- The \\Computer element is optional; it specifies the name of the computer. The default is the local computer.
- The \Object element specifies the performance object that contains one or more performance counters. If this object has a variable list of instances, then an instance string is included.
- The (ParentInstance/ObjectInstance#InstanceIndex) element is only required if the object supports multiple instances. The format of the string depends on the object type. If the object has single instances, the format is the instance name enclosed in parentheses. For example:

  ```
  (Explorer)
  ```

  For objects with multiple instances, the parent instance name comes before the object instance, and is separated by a forward slash ( / ) character. For example:

  ```
  (Explorer/0)
  ```

  If the object has multiple instances that have the same name string, they can be indexed sequentially by specifying the instance index (a number starting with "1") preceded by a pound sign. Instance indexes are zero-based, so all instances have an implicit zero index. For example:

  ```
  (Explorer/0#1)
  ```

- The \Counter element is required; it specifies the performance counter.

An instance called *_Total instance* is available on most objects and represents the sum of the values for all instances of the object for a specific counter.

In addition to the counter descriptions provided in this Help file, to display counter descriptions you can click **Explain** in the **Add Counters** dialog box in System Monitor or Performance Logs and Alerts. You can also click **Explain** in the **Add Objects** dialog box in Performance Counter Logs to display performance object descriptions.

# Counter Types

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Each counter is assigned a counter type. The counter type determines how the counter data is calculated, averaged, and displayed. The Performance console supports more than 35 counter types, although many of the available counter types are not implemented in the counters installed with the Windows Server 2003 family. All supported counter types are described, including the following information.

**Counter type name** The name of the counter type. The System Monitor counter descriptions in this Help file include the name of the counter type. You can click the counter name in the description to see a definition of the counter.

**Description** A brief description of the counter type, including a description of the formula by which counters of that type are calculated and displayed.

**Generic type** The general category that represents the display format counter. Generic types include:

- Average. These counters measure a value over time and display the average of the last two measurements.
- Difference. These counters subtract the last measurement from the previous one and display the difference, if it is positive; if negative, they display a zero.
- Instantaneous. These counters display the most recent measurement.
- Percentage. These counters display calculated values as a percentage.
- Rate. Similar to an average counter, these counters sample an increasing count of events over time. The count is divided by the change in time to display a rate of activity.

**Formula** Describes how to convert the raw data and other components, such as performance frequency, to arrive at the formatted counter value for display in the Performance console. The variables N and D appear in the formula to describe the two distinct data components, such as raw data and performance frequency, used to calculate the displayed counter values.

For each counter type, the notation of the formula can include the following variables:

| Variable | Description |
|---|---|
| B | A base-value variable. Some counter types require the use of a counter base that supplies a stored value for calculating their formatted value. That base value is the denominator in the formula for the counter type. When describing counter types with multiple samplings, B is shown with numbered subscripts, such as $B_0$ for the oldest sampling, $B_1$ for the next most recent sampling, and $B_n$ for the latest sampling. |
| N | The numerator, typically the first variable. For the purposes of the calculation formulas in this section, N is used to describe the raw counter data. |
| D | The denominator, typically the second variable, as needed. |

The following table describes the types of counter bases and the counter types that they support.

| Base Type | Description | Corresponding Counter Types |
|---|---|---|
| PERF_AVERAGE_BASE | Stores the denominator for calculating a counter to present "time per operation" or "count per operation." | PERF_AVERAGE_BULK<br>PERF_AVERAGE_TIMER |

| | | |
|---|---|---|
| PERF_COUNTER_MULTI_BASE | Indicates the number of items sampled. It is used as the denominator in the calculations to get an average among the items sampled when taking timings of multiple, but similar, items. | PERF_100NSEC_MULTI_TIMER<br>PERF_100NSEC_MULTI_TIMER_INV<br>PERF_COUNTER_MULTI_TIMER<br>PERF_COUNTER_MULTI_TIMER_INV |
| PERF_RAW_BASE | Stores the denominator of a counter that presents a general arithmetic fraction. | PERF_RAW_FRACTION |
| PERF_SAMPLE_BASE | Stores the number of sampling interrupts taken and is used as a denominator in the sampling fraction. The sampling fraction is the number of samples that were 1 (or TRUE) for a sample interrupt. | PERF_SAMPLE_FRACTION |

**Average** The mathematical formula used to calculate averages of the formatted counter data.

**Example** An example of the specified counter type. Click the counter name in the example to see a description of that counter.

# PERF_100NSEC_MULTI_TIMER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the active time of one or more components as a percentage of the total time of the sample interval. It measures time in 100 nanosecond units.<br>This counter type is a multitimer. Multitimers are designed to monitor more than one instance of a component, such as a processor or disk. |
| Generic type | Percentage. |
| Formula | $((N_1 - N_0) / (D_1 - D_0)) \times 100 / B$, where the denominator (D) represents the total elapsed time of the sample interval, the numerator (N) represents the portions of the sample interval during which the monitored components were active, and the variable B represents the base count for the monitored components. |
| Average | $((N_x - N_0) / (D_x - D_0)) \times 100 / B$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_100NSEC_MULTI_TIMER_INV

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the active time of one or more components as a percentage of the total time of the sample interval. Counters of this type measure time in 100 nanosecond units.<br><br>This counter type is an inverse multitimer. Multitimers are designed to monitor more than one instance of a component, such as a processor or disk. Inverse counters measure the time that a component is *not* active and derive its active time from the measurement of inactive time. |
| Generic type | Percentage. |
| Formula | $(B - ((N_1 - N_0) / (D_1 - D_0))) \times 100$, where the denominator (D) represents the total elapsed time of the sample interval, the numerator (N) represents the time during the interval when monitored components were inactive, and the variable B represents the number of components being monitored. |
| Average | $(B - ((N_x - N_0) / (D_x - D_0))) \times 100$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

**See Also**
**Concepts**
PERF_100NSEC_MULTI_TIMER

# PERF_100NSEC_TIMER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the active time of a component as a percentage of the total elapsed time of the sample interval. It measures time in units of 100 nanoseconds. Counters of this type are designed to measure the activity of one component at a time. |
| Generic type | Percentage. |
| Formula | $(N_1 - N_0) / (D_1 - D_0) \times 100$, where the denominator (D) represents the total elapsed time of the sample interval, and the numerator (N) represents the portions of the sample interval during which the monitored components were active. |
| Average | $(N_x - N_0) / (D_x - D_0) \times 100$ |
| Example | Processor\ % User Time |

# PERF_100NSEC_TIMER_INV

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the average percentage of active time observed during the sample interval. This is an inverse counter. Inverse counters are calculated by monitoring the percentage of time that the service was inactive and then subtracting that value from 100 percent. |
| Generic type | Percentage. |
| Formula | $(1- ((N_1 - N_0) / (D_1 - D_0))) \times 100$, where the denominator (D) represents the total elapsed time of the sample interval, and the numerator (N) represents the time during the interval when the monitored components were inactive. |
| Average | $(1- ((N_x - N_0) / (D_x - D_0))) \times 100$ |
| Example | Processor\ % Processor Time |

# PERF_AVERAGE_BASE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter is used as the base data (denominator) in the computation of time or count averages for the PERF_AVERAGE_BULK and PERF_AVERAGE_TIMER counter types. This counter type collects the last observed value only. |
| Generic type | Instantaneous |
| Formula | None. This counter uses raw data in factional calculations without displaying an output. |
| Average | SUM (N) / x |
| Example | There are no counters of this type installed with Windows Server 2003 family. However, counters of this type can be used by the Windows Server 2003 family to calculate the Average Disk Transfer counters. |

# PERF_AVERAGE_BULK

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows how many items are processed, on average, during an operation. Counters of this type display a ratio of the items processed (such as bytes sent) to the number of operations completed. The ratio is calculated by comparing the number of items processed during the last interval to the number of operations completed during the last interval. |
| Generic type | Average. |
| Formula | $(N_1 - N_0) / (D_1 - D_0)$, where the numerator (N) represents the number of items processed during the last sample interval, and the denominator (D) represents the number of operations completed during the last two sample intervals. |
| Average | $(N_x - N_0) / (D_x - D_0)$ |
| Example | PhysicalDisk\ Avg. Disk Bytes/Transfer |

# PERF_AVERAGE_TIMER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type measures the average time it takes to complete a process or operation. Counters of this type display a ratio of the total elapsed time of the sample interval to the number of processes or operations completed during that time. This counter type measures time in ticks of the system clock. The variable F represents the number of ticks per second. The value of F is factored into the equation so that the result is displayed in seconds. |
| Generic type | Average. |
| Formula | $((N_1 - N_0) / F) / (D_1 - D_0)$, where the numerator (N) represents the number of ticks counted during the last sample interval, the variable F represents the frequency of the ticks, and the denominator (D) represents the number of operations completed during the last sample interval. |
| Average | $((N_x - N_0) / F) / (D_x - D_0)$ |
| Example | PhysicalDisk\ Avg. Disk sec/Transfer |

# PERF_COUNTER_100NS_QUEUELEN_TYPE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type measures the queue-length space-time product using a 100-nanosecond time base. |
| Generic type | Rate. |
| Formula | $(TB(X_1-X_0)) / (Y_1-Y_0)$, where TB represents the performance frequency time base, the denominator (Y) represents the performance time measurement, and the numerator (X) represents counter data. |
| Average | $(TB(X_1-X_0)) / (Y_1-Y_0)$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_BULK_COUNT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the average number of operations completed during each second of the sample interval. Counters of this type measure time in ticks of the system clock. The variable F represents the number of ticks per second. The value of F is factored into the equation so that the result is displayed in seconds. This counter type is the same as the PERF_COUNTER_COUNTER type, but it uses larger fields to accommodate larger values. |
| Generic type | Difference. |
| Formula | $(N_1 - N_0) / ( (D_1 - D_0) / F$, where the numerator (N) represents the number of operations performed during the last sample interval, the denominator (D) represent the number of ticks elapsed during the last sample interval, and the variable F is the frequency of the ticks. |
| Average | $(N_x - N_0) / ((D_x - D_0)/ F)$ |
| Example | System\File Read Bytes/sec |

# PERF_COUNTER_COUNTER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the average number of operations completed during each second of the sample interval. Counters of this type measure time in ticks of the system clock. The F variable represents the number of ticks per second. The value of F is factored into the equation so that the result can be displayed in seconds. |
| Generic type | Difference. |
| Formula | $(N_1 - N_0) / ( (D_1 - D_0) / F)$, where the numerator (N) represents the number of operations performed during the last sample interval, the denominator (D) represents the number of ticks elapsed during the last sample interval, and F is the frequency of the ticks. |
| Average | $(N_x - N_0) / ((D_x - D_0) / F)$ |
| Example | System\File Read Operations/sec |

# PERF_COUNTER_DELTA

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the change in the measured attribute between the two most recent sample intervals. |
| Generic type | Difference. |
| Formula | $N_1 - N_0$ |
| Average | $(N_x - N_0) / x$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_LARGE_DELTA

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the change in the measured attribute between the two most recent sample intervals. It is the same as the PERF_COUNTER_DELTA counter type, except that it uses larger fields to accommodate larger values. |
| Generic type | Difference. |
| Formula | $N_1 - N_0$ |
| Average | $(N_x - N_0) / x$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_LARGE_QUEUELEN_TYPE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type monitors the average length of a queue to a resource over time. Counters of this type display the difference between the queue lengths observed during the last two sample intervals, divided by the duration of the interval. This counter type is the same as the PERF_COUNTER_QUEUELEN_TYPE counter type, except that it uses larger fields to accommodate larger values. |
| Generic type | Average. |
| Formula | $(N_1 - N_0) / (D_1 - D_0)$, where the numerator (N) represents the number of items in a queue, and the denominator (D) represents the time elapsed during the sample interval. |
| Average | $(N_x - N_0) / (D_x - D_0)$ |
| Example | PhysicalDisk\ Avg. Disk Queue Length |

# PERF_COUNTER_LARGE_RAWCOUNT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the last observed value only, not an average. It is the same as the PERF_COUNTER_RAWCOUNT counter type, except that it uses larger fields to accommodate larger values. |
| Generic type | Instantaneous. |
| Formula | None. Shows raw data as collected. |
| Average | SUM (N) / x |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_LARGE_RAWCOUNT_HEX

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the last observed value, in hexadecimal format. It is the same as the PERF_COUNTER_RAWCOUNT_HEX counter type, except that it uses larger fields to accommodate larger values. |
| Generic type | Instantaneous. |
| Formula | None. Shows raw data as collected. |
| Average | SUM (N) / x |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_MULTI_TIMER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type is a multitimer. Multitimers collect data from more than one instance of a component, such as a processor or disk. Counters of this type display the active time of one or more components as a percentage of the total time of the sample interval. Because the numerator records the active time of components operating simultaneously, the resulting percentage can exceed 100 percent.<br><br>This counter type differs from PERF_100NSEC_MULTI_TIMER in that it measures time in units of ticks of the system performance timer, rather than in 100 nanosecond units. |
| Generic type | Percentage. |
| Formula | $((N_1 - N_0) / (D_1 - D_0)) \times 100$, where the denominator (D) represents the total elapsed time of the sample interval, and the numerator (N) represents the portions of the sample interval during which the monitored components were active. They are multiplied by 100 to display the value as a percentage. |
| Average | $(N_x - N_0) / (D_x - D_0) \times 100$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_MULTI_TIMER_INV

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the active time of one or more components as a percentage of the total time of the sample interval.<br><br>This counter type is an inverse multitimer. Multitimers monitor more than one instance of a component, such as a processor or disk. Inverse counters measure the time that a component is *not* active, and derive the active time from that measurement.<br><br>This counter differs from PERF_100NSEC_MULTI_TIMER_INV in that it measures time in units of ticks of the system performance timer, rather than in 100 nanosecond units. |
| Generic type | Percentage. |
| Formula | (B- ((N$_1$ - N$_0$) / (D$_1$ - D$_0$))) x 100, where the denominator (D) represents the total elapsed time of the sample interval, the numerator (N) represents the time during the interval when monitored components were inactive, and the variable B represents the number of components being monitored. |
| Average | (B- ((N$_x$ - N$_0$) / (D$_x$ - D$_0$))) x 100 |
| Example | There are no counters of this type installed with Windows Server 2003 family. |

# PERF_COUNTER_OBJ_QUEUELEN_TYPE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type measures the queue-length space-time product using an object-specific time base. |
| Generic type | Percentage. |
| Formula | $(X_1-X_0) / (Y_1-Y_0)$, where the denominator (Y) represents the performance time measurement, and the numerator (X) represents counter data. |
| Average | $(X_1-X_0) / (Y_1-Y_0)$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_QUEUELEN_TYPE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type is designed to monitor the average length of a queue to a resource over time. It shows the difference between the queue lengths observed during the last two sample intervals divided by the duration of the interval. |
| Generic type | Average. |
| Formula | $(N_1 - N_0) / (D_1 - D_0)$, where the numerator (N) represents the number of items in the queue, and the denominator (D) represents the time elapsed during the last sample interval. |
| Average | $(N_x - N_0) / (D_x - D_0)$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_RAWCOUNT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the last observed value only. It does not display an average. |
| Generic type | Instantaneous. |
| Formula | None. Shows raw data as collected. |
| Average | SUM (N) / x |
| Example | Memory\Available Bytes |

# PERF_COUNTER_RAWCOUNT_HEX

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the most recently observed value, in hexadecimal format. It does not display an average. |
| Generic type | Instantaneous. |
| Formula | None. Shows raw data as collected. |
| Average | SUM (N) / x |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_TEXT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows a variable-length text string. It does not display calculated values. |
| Generic type | Text. |
| Formula | None. |
| Average | None. |
| Example | There are no counters of this type installed with Windows Server 2003 family. |

# PERF_COUNTER_TIMER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the average time that a component was active as a percentage of the total sample time. |
| Generic type | Percentage |
| Formula | $(N_1 - N_0) / (D_1 - D_0)$, where the denominator (D) represents the total elapsed time of the sample interval, and the numerator (N) represents the portions of the sample interval during which the monitored components were active. |
| Average | $(N_x - N_0) / (D_x - D_0)$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_COUNTER_TIMER_INV

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This is an inverse counter type. Inverse counters measure the time that a component is *not* active, and derive the active time from that measurement. Counters of this type display the average percentage of active time observed during sample interval. The value of these counters is calculated by monitoring the percentage of time that the service was inactive and then subtracting that value from 100 percent.<br><br>This counter type is the same as the PERF_100NSEC_TIMER_INV counter type, except that it measures time in units of ticks of the system performance timer, rather than in 100 nanosecond units. |
| Generic type | Percentage. |
| Formula | $(1- ((N_1 - N_0) / (D_1 - D_0))) \times 100$, where the denominator (D) represents the total elapsed time of the sample interval, and the numerator (N) represents the time during the interval when the monitored components were inactive. |
| Average | $(1- ((N_x - N_0) / (D_x - D_0))) \times 100$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_ELAPSED_TIME

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the total time between when the component or process started and the time when this value is calculated. The variable F represents the number of time units that elapse in one second. The value of F is factored into the equation so that the result is displayed in seconds. |
| Generic type | Difference. |
| Formula | $(D_0 - N_0) / F$, where the denominator (D) represents the current time, the numerator (N) represents the time the object was started, and the variable F represents the number of time units that elapse in one second. |
| Average | $(D_x - N_0) / F$ |
| Example | System\System Up Time |

# PERF_LARGE_RAW_BASE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type collects the last observed value. It is the same as the PERF_RAW_BASE counter type except that it uses larger fields to accomaodate larger values. |
| Generic type | Instantaneous |
| Formula | None. Uses raw data in factional calculations without displaying an output. |
| Average | SUM (N) / x |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_OBJ_TIME_TIMER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This 64-bit counter type is a timer that displays in object-specific units. |
| Generic type | Percentage. |
| Formula | $(X_1-X_0) / (Y_1-Y_0)$, where the denominator (Y) represents the performance time measurement, and the numerator (X) represents counter data. |
| Average | $(X_1-X_0) / (Y_1-Y_0)$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_PRECISION_100NS_TIMER

Updated: September 12, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows a value that consists of two counter values: the count of the elapsed time of the event being monitored, and the "clock" time from a private timer in the same units. It measures time in 100 nanosecond units. This counter type differs from other counter timers in that the clock tick value accompanies the counter value eliminating any possible difference due to latency from the function call. Precision counter types are used when standard system timers are not precise enough for accurate readings. |
| Generic type | Percentage. |
| Formula | $N1 - N_0 / D_1 - D_0$, where the numerator (N) represents the counter value, and the denominator (D) is the value of the private timer. The private timer has the same frequency as the 100 nanosecond timer. |
| Average | $N1 - N_0 / D_1 - D_0$ |
| Example | PhysicalDisk\% Disk Time |

# PERF_PRECISION_TIMER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows a value that consists of two counter values: the count of the elapsed time of the event being monitored, and the frequency from the system performance timer. This counter type differs from other counter timers in that the clock tick value accompanies the counter value, eliminating any possible difference due to latency from the function call. Precision counter types are used when standard system timers are not precise enough for accurate readings. |
| Generic type | Percentage. |
| Formula | $N_x - N_0 / D_1 - D_0$, where the numerator (N) represents the counter value, and the denominator (D) represents the value of the system performance timer. |
| Average | $N_x - N_0 / D_1 - D_0$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_PRECISION_OBJECT_TIMER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows a value that consists of two counter values: the count of the elapsed time of the event being monitored, and the frequency specified in the PerfFreq field of the object header. This counter type differs from other counter timers in that the clock tick value accompanies the counter value so as to eliminate any possible difference due to latency from the function call. Precision counter types are used when standard system timers are not precise enough for accurate readings. |
| Generic type | Percentage. |
| Formula | $N_x$ - $N_0$ / $D_1$ - $D_0$, where the numerator (N) represents the counter value and the denominator (D) represents the value of the private timer. |
| Average | $N_x$ - $N_0$ / $D_1$ - $D_0$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_RAW_BASE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type collects the last observed value only. |
| Generic type | Instantaneous |
| Formula | None. Uses raw data in factional calculations without displaying an output. |
| Average | SUM (N) / x |
| Example | There are no counters of this type installed with the Windows Server 2003 family. However, this type of counter is used to calculate the Disk Queue Length counters. |

# PERF_RAW_FRACTION

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the ratio of a subset to its set as a percentage. For example, it compares the number of bytes in use on a disk to the total number of bytes on the disk. Counters of this type display the current percentage only, not an average over time. |
| Generic type | Instantaneous, Percentage. |
| Formula | ($N_0$ / $D_0$), where the denominator (D) represents a measured attribute, and the numerator (N) represents one component of that attribute. |
| Average | SUM (N / D) /x |
| Example | Paging File\% Usage Peak |

# PERF_SAMPLE_COUNTER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the average number of operations completed in one second. It measures time in units of ticks of the system performance timer. The variable F represents the number of ticks that occur in one second. The value of F is factored into the equation so that the result is displayed in seconds. |
| Generic type | Average. |
| Formula | $(N_1 - N_0) / ((D_1 - D_0) / F)$, where the numerator (N) represents the number of operations completed, the denominator (D) represents elapsed time in units of ticks of the system performance timer, and the variable F represents the number of ticks that elapse in one second. |
| Average | $(N_x - N_0) / ((D_x - D_0) / F)$ |
| Example | There are no counters of this type installed with the Windows Server 2003 family. |

# PERF_SAMPLE_FRACTION

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| | |
|---|---|
| Description | This counter type shows the average ratio of hits to all operations during the last two sample intervals. |
| Generic type | Percentage; Hits % |
| Formula | $(N_1 - N_0) / (D_1 - D_0)$, where the numerator (N) represents the number of successful operations during the last sample interval, and the denominator (D) represents the change in the number of all operations (of the type measured) completed during the sample interval. |
| Average | $(N_x - N_0) / (D_x - D_0)$ (See Note below) |
| Example | Cache\Pin Read Hits % |

**Note**

- The Hit % counters are best displayed in System Monitor in chart view. Hits often appear in short bursts that are not visible in reports. Also, the average displayed for Hit % on the status bar in chart view might not match the average displayed in report view because they are calculated differently. In chart view, the Hit % is an average of all changes in the counter during the test interval; in report view, it is the average of the difference between the first and last counts during the test interval.

# Counters by Feature

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some Windows Server 2003 family features or services install one or more performance objects to measure the activity of the feature or service. The following is a list of features or services and their corresponding performance objects:

- Active Directory
    - NTDS Object
- File Replication Service
    - FileReplicaConn Object
    - FileReplicaSet Object
- Indexing Service
    - Indexing Service Object
    - Indexing Service Filter Object
    - Http Indexing Service Object
- Internet Information Services (IIS)
    - Active Server Pages Object
    - FTP Service Object
    - Web Service Object
    - Internet Information Services Global Object
- Message Queuing
    - MSMQ Session Object
    - MSMQ Queue Object
    - MSMQ Service Object
- Quality of Service (QoS) Admission Control
    - ACS/RSVP Service Object
    - ACS/RSVP Interfaces Object
    - ACS Policy Object
- Remote Access Service (RAS)
    - RAS Port Object

    - RAS Total Object
- Terminal Services
    - Terminal Services Session Object
- Windows System Resource Manager (WSRM)
    - WSRM Policy Object
    - WSRM Process Object
    - WSRM Process Matching Criteria Object

# Counters by Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section includes all of the performance objects installed with the Windows Server 2003 family, grouped alphabetically. If objects are installed on your computer, they appear in the performance tools user interfaces (UI). You typically install an object at the same time that you install its associated service or feature.

# ACS Policy Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Quality of Service (QoS) ACS Policy performance object provides data about the operation of Microsoft Identity Local Policy Module. Network administrators should monitor the counters provided by this object when the QoS Admission Control Service (ACS) or Resource Reservation Protocol (RSVP) is enabled on a subnet. The QoS ACS Policy object is available only if the QoS ACS or RSVP has won the Designated Subnet Bandwidth Manager (DSBM) election and is running the Policy-based Admission Control. The Subnet Bandwidth Manager is a QoS signaling technique supported by the Windows Server 2003 family that is used to manage bandwidth on shared media such as 802.1 local area network (LAN) topologies.

There is one object for each subnet being managed by this QoS ACS host.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bad User Identity Policy Elements | Shows the number of user Identity Policy Elements that could not be identified by Microsoft Identity Local Policy Module. | PERF_COUNTER_RAWCOUNT |
| DS Cache Size | Shows the number of entries in the directory service cache. | PERF_COUNTER_RAWCOUNT |
| Flow Count PATH Rejections | Shows the number of PATH messages rejected due to flow count policy. | PERF_COUNTER_RAWCOUNT |
| Flow Count RESV Rejections | Shows the number of RESV messages rejected due to flow count policy. | PERF_COUNTER_RAWCOUNT |
| Flow Duration PATH Rejections | Shows the number of PATH messages rejected due to flow duration policy. | PERF_COUNTER_RAWCOUNT |
| Flow Duration RESV Rejections | Shows the number of RESV messages rejected due to flow duration policy. | PERF_COUNTER_RAWCOUNT |
| Flow Rate PATH Rejections | Shows the number of PATH messages rejected due to flow rate policy. | PERF_COUNTER_RAWCOUNT |
| Flow Rate RESV Rejections | Shows the number of RESV messages rejected due to flow rate policy. | PERF_COUNTER_RAWCOUNT |
| Other Policy PATH Rejections | Shows the number of PATH messages rejected due to other policies on the subnet. | PERF_COUNTER_RAWCOUNT |
| Other Policy RESV Rejections | Shows the number of RESV messages rejected due to other policies on the subnet. | PERF_COUNTER_RAWCOUNT |
| PATH Message Changes Admitted | Shows the number of changes to PATH messages admitted. | PERF_COUNTER_RAWCOUNT |
| PATH Messages Admitted | Shows the number of PATH messages admitted. | PERF_COUNTER_RAWCOUNT |
| Peak Rate PATH Rejections | Shows the number of PATH messages rejected because their peak rate exceeded peak rate policy. | PERF_COUNTER_RAWCOUNT |
| Peak Rate RESV Rejections | Shows the number of RESV messages rejected because their peak rate exceeded peak rate policy. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| RESV Message Changes Admitted | Shows the number of changes to RESV messages admitted. | PERF_COUNTER_RAWCOUNT |
| RESV Messages Admitted | Shows the number of RESV messages admitted. | PERF_COUNTER_RAWCOUNT |
| Total Flow Rate PATH Rejections | Shows the number of PATH messages rejected because their total peak rate exceeded total flow rate policy. | PERF_COUNTER_RAWCOUNT |
| Total Flow Rate RESV Rejections | Shows the number of RESV messages rejected because their total peak rate exceeded the total flow rate policy. | PERF_COUNTER_RAWCOUNT |
| Total Peak Rate PATH Rejections | Shows the number of PATH messages rejected because their total peak rate exceeded the total peak rate policy. | PERF_COUNTER_RAWCOUNT |
| Total Peak Rate RESV Rejections | Shows the number of RESV messages rejected because their total peak rate exceeded the total peak rate policy. | PERF_COUNTER_RAWCOUNT |
| User ID Change PATH Rejections | Shows the number of PATH messages rejected because the user ID changed. | PERF_COUNTER_RAWCOUNT |
| User ID Change RESV Rejections | Shows the number of RESV messages rejected because the user ID changed. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Counters by Feature

# Active Server Pages Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Server Pages performance object consists of counters that monitor errors, requests, sessions, and other data generated by Active Server Pages (ASP pages) applications.

| Counter Name | Description | Counter Type |
|---|---|---|
| Debugging Requests | Shows the number of ASP pages sent to the script debugger. | PERF_COUNTER_RAWCOUNT |
| Errors During Script Runtime | Shows the number of requests that fail due to run-time errors. | PERF_COUNTER_RAWCOUNT |
| Errors From ASP Preprocessor | Shows the number of requests that fail due to preprocessor errors. | PERF_COUNTER_RAWCOUNT |
| Errors From Script Compilers | Shows the total number of requests that fail due to script compilation errors. | PERF_COUNTER_RAWCOUNT |
| Errors/sec | Shows the number of errors that occur per second. | PERF_COUNTER_COUNTER |
| Request Bytes In Total | Shows the total size of all requests. | PERF_COUNTER_RAWCOUNT |
| Request Bytes Out Total | Shows the total size of all responses sent to clients. This does not include the standard Hypertext Transfer Protocol (HTTP) response headers. | PERF_COUNTER_RAWCOUNT |
| Request Execution Time | Shows the time, in milliseconds, that it took to execute the most recent request. | PERF_COUNTER_RAWCOUNT |
| Request Wait Time | Shows the time, in milliseconds, that the most recent request was waiting in the queue. | PERF_COUNTER_RAWCOUNT |
| Requests Disconnected | Shows the number of requests that were disconnected due to a communications failure. | PERF_COUNTER_RAWCOUNT |
| Requests Executing | Shows the number of requests that are currently being executed. | PERF_COUNTER_RAWCOUNT |
| Requests Failed Total | Shows the total number of requests that fail due to errors, authorization failure, or rejections. | PERF_COUNTER_RAWCOUNT |
| Requests Not Authorized | Shows the number of requests that fail due to insufficient access rights. | PERF_COUNTER_RAWCOUNT |
| Requests Not Found | Shows the number of requests for files that were not found. | PERF_COUNTER_RAWCOUNT |
| Requests Queued | Shows the number of requests waiting for service from the queue. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Requests Rejected | Shows the total number of requests that were not executed due to insufficient resources for processing the requests. This may be due to configuration problems or a memory limitation. | PERF_COUNTER_RAWCOUNT |
| Requests Succeeded | Shows the number of requests that were successfully executed. | PERF_COUNTER_RAWCOUNT |
| Requests Timed Out | Shows the number of requests that timed out. | PERF_COUNTER_RAWCOUNT |
| Requests Total | Shows the total number of requests since the service was started. | PERF_COUNTER_RAWCOUNT |
| Requests/sec | Shows the number of requests that were executed per second. This value is important for monitoring server performance. | PERF_COUNTER_COUNTER |
| Script Engines Cached | Shows the number of script engines that are in the cache. | PERF_COUNTER_RAWCOUNT |
| Session Duration | Shows the time, in milliseconds, that the most recent session lasted. | PERF_COUNTER_RAWCOUNT |
| Sessions Current | Shows the number of sessions that are being serviced. | PERF_COUNTER_RAWCOUNT |
| Sessions Timed Out | Shows the number of sessions that have timed out. | PERF_COUNTER_RAWCOUNT |
| Sessions Total | Shows the total number of sessions that have been opened since the service was started. | PERF_COUNTER_RAWCOUNT |
| Templates Cached | Shows the number of templates that are currently in the ASP page file cache. | PERF_COUNTER_RAWCOUNT |
| Template Cache Hit Rate | Shows the percentage of requests found in the template cache over the sample interval. | PERF_COUNTER_RAWCOUNT |
| Template Notifications | Shows the number of templates invalidated in the cache due to change notification. | PERF_COUNTER_RAWCOUNT |
| Transactions Aborted | Shows the number of transactions that were aborted (rolled back). | PERF_COUNTER_RAWCOUNT |
| Transactions Committed | Shows the number of transactions that were committed. | PERF_COUNTER_RAWCOUNT |
| Transactions Pending | Shows the number of transactions that have not yet been committed. | PERF_COUNTER_RAWCOUNT |
| Transactions Total | Shows the total number of transactions since the service was started. This counter reports the sum of the values of Transactions Aborted, Transactions Committed, and Transactions Pending. | PERF_COUNTER_RAWCOUNT |
| Transactions/sec | Shows the number of transactions started per second. | PERF_COUNTER_COUNTER |

**See Also**

**Concepts**

Counters by Feature

# Apple Talk Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The AppleTalk performance object consists of counters that monitor traffic on the AppleTalk network.

| Counter Name | Description | Counter Type |
|---|---|---|
| AARP Packets/sec | Shows the rate, in incidents per second, at which AppleTalk Address Resolution Protocol (AARP) packets were received on this port. | PERF_COUNTER_COUNTER |
| ATP ALO Response/sec | Shows the rate, in incidents per second, at which AppleTalk Transaction Protocol (ATP) At-Least-Once transaction responses were received on this port. | PERF_COUNTER_COUNTER |
| ATP Packets/sec | Shows the rate, in incidents per second, at which ATP packets were received on this port. | PERF_COUNTER_COUNTER |
| ATP Recvd Release/sec | Shows the rate, in incidents per second, at which ATP transaction release packets were received on this port. | PERF_COUNTER_COUNTER |
| ATP Response Timeouts | Shows the number of ATP release timers that have expired on this port since the computer was last started. | PERF_COUNTER_RAWCOUNT |
| ATP Retries Local | Shows the number of ATP requests that were retransmitted on this port since the computer was last started. | PERF_COUNTER_RAWCOUNT |
| ATP Retries Remote | Shows the number of ATP requests that were retransmitted to this port since the computer was last started. | PERF_COUNTER_RAWCOUNT |
| ATP XO Response/sec | Shows the rate, in incidents per second, at which ATP Exactly-Once transaction responses were received on this port. | PERF_COUNTER_COUNTER |
| Average Time/AARP Packet | Shows the average time, in milliseconds, that was needed to process an AARP packet on this port. | PERF_AVERAGE_BULK |
| Average Time/ATP Packet | Shows the average time, in milliseconds, that was needed to process an ATP packet on this port. | PERF_AVERAGE_BULK |
| Average Time/DDP Packet | Shows the average time, in milliseconds, that was needed to process a Data Delivery Protocol (DDP) packet on this port. | PERF_AVERAGE_BULK |
| Average Time/NBP Packet | Shows the average time, in milliseconds, that was needed to process a Name Binding Protocol (NBP) packet on this port. | PERF_AVERAGE_BULK |
| Average Time/RTMP Packet | Shows the average time, in milliseconds, that was needed to process a Routing Table Maintenance Protocol (RTMP) packet on this port. | PERF_AVERAGE_BULK |

| | | |
|---|---|---|
| Average Time/ZIP Packet | Shows the average time, in milliseconds, that was needed to process a Zone Information Protocol (ZIP) packet on this port. | PERF_AVERAGE_BULK |
| Bytes In/sec | Shows the rate, in incidents per second, at which bytes were received on this port. | PERF_COUNTER_BULK_COUNT |
| Bytes Out/sec | Shows the rate, in incidents per second, at which bytes were sent on this port. | PERF_COUNTER_BULK_COUNT |
| Current Nonpaged Pool | Shows the current amount of nonpaged memory resources that were used by AppleTalk. | PERF_COUNTER_RAWCOUNT |
| DDP Packets/sec | Shows the rate, in incidents per second, at which DDP packets were received on this port. | PERF_COUNTER_COUNTER |
| NBP Packets/sec | Shows the rate, in incidents per second, at which NBP packets were received on this port. | PERF_COUNTER_COUNTER |
| Packets Dropped | Shows the number of packets that were dropped due to resource limitations on this port since the computer was last started. | PERF_COUNTER_RAWCOUNT |
| Packets In/sec | Shows the rate, in incidents per second, at which packets were received on this port. | PERF_COUNTER_COUNTER |
| Packets Out/sec | Shows the rate, in incidents per second, at which packets were sent on this port. | PERF_COUNTER_COUNTER |
| Packets Routed In/sec | Shows the rate, in incidents per second, at which packets were routed on this port. | PERF_COUNTER_COUNTER |
| Packets Routed Out/sec | Shows the rate, in incidents per second, at which packets were routed out on this port. | PERF_COUNTER_COUNTER |
| RTMP Packets/sec | Shows the rate, in incidents per second, at which RTMP packets were received on this port. | PERF_COUNTER_COUNTER |
| ZIP Packets/sec | Shows the rate, in incidents per second, at which ZIP packets were received on this port. | PERF_COUNTER_COUNTER |

# Battery Status Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Battery Status performance object reports the battery status on laptops and on other computers that use batteries, based on Windows Management Instrumentation (WMI). Battery Status is only displayed in Performance Console when batteries are available in your computer and the WMI Driver is installed.

| **Counter Name** | **Description** | **Counter Type** |
|---|---|---|
| Charge Rate | Shows the rate of battery charge in milliwatts (mW). If the battery is not charging, or is discharging, then this value will be 0. | WMI-Based Driver |
| Discharge Rate | Shows the rate of battery discharge in milliwatts (mW). If the battery is not discharging, or is charging, then this value will be 0. | WMI-Based Driver |
| Remaining Battery Capacity | Shows the amount of charge remaining in the battery in milliwatts per hour (mWh). | WMI-Based Driver |
| Battery Tag | Shows the number of times that the battery has been changed. | WMI-Based Driver |
| Voltage | Shows the voltage of the battery in milliwatts (mW). | WMI-Based Driver |

# Browser Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Browser performance object consists of counters that measure the rates of announcements, enumerations, and other browser transmissions.

| Counter Name | Description | Counter Type |
|---|---|---|
| Announcements Domain/sec | Shows the rate at which a domain has announced itself to the network. | PERF_COUNTER_BULK_COUNT |
| Announcements Server/sec | Shows the rate at which the servers in this domain have announced themselves to the server. | PERF_COUNTER_BULK_COUNT |
| Announcements Total/sec | Shows the sum of the values of Announcements Domain/sec and Announcements Server/sec. | PERF_COUNTER_BULK_COUNT |
| Duplicate Master Announcements | Shows the number of times that the master browser has detected another master browser on the same domain. | PERF_COUNTER_RAWCOUNT |
| Election Packets/sec | Shows the rate at which browser election packets have been received by the workstation. | PERF_COUNTER_COUNTER |
| Enumerations Domain/sec | Shows the rate at which domain browse requests have been processed by this workstation. | PERF_COUNTER_COUNTER |
| Enumerations Other/sec | Shows the rate of browse requests processed by this workstation that are not domain or server browse requests. | PERF_COUNTER_COUNTER |
| Enumerations Server/sec | Shows the rate at which server browse requests have been processed by this workstation. | PERF_COUNTER_COUNTER |
| Enumerations Total/sec | Shows the rate at which browse requests have been processed by this workstation. This value is the sum of the values of Enumerations Domain/sec, Enumerations Other/sec, and Enumerations Server/sec. | PERF_COUNTER_COUNTER |
| Illegal Datagrams/sec | Shows the rate at which incorrectly formatted datagrams have been received by the workstation. | PERF_COUNTER_BULK_COUNT |
| Mailslot Allocations Failed | Shows the number of times that the datagram receiver has failed to allocate a buffer to hold a user mailslot write. | PERF_COUNTER_RAWCOUNT |
| Mailslot Opens Failed/sec | Shows the rate at which this workstation received mailslot messages to be delivered to mailslots that are not present on this workstation. | PERF_COUNTER_COUNTER |
| Mailslot Receives Failed | Shows the number of mailslot messages that could not be received due to transport failures. | PERF_COUNTER_RAWCOUNT |
| Mailslot Writes Failed | Shows the total number of mailslot messages that have been successfully received, but could not be written to the mailslot. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Mailslot Writes/sec | Shows the rate at which mailslot messages have been successfully received. | PERF_COUNTER_COUNTER |
| Missed Mailslot Datagrams | Shows the number of mailslot datagrams that have been discarded due to configuration or allocation limits. | PERF_COUNTER_RAWCOUNT |
| Missed Server Announcements | Shows the number of server announcements that have been missed due to configuration or allocation limits. | PERF_COUNTER_RAWCOUNT |
| Missed Server List Requests | Shows the number of times a request to retrieve a list of browser servers was received by this workstation, but could not be processed. | PERF_COUNTER_RAWCOUNT |
| Server Announce Allocations Failed/sec | Shows the rate at which server (or domain) announcements have failed due to lack of memory. | PERF_COUNTER_COUNTER |
| Server List Requests/sec | Shows the rate at which requests to retrieve a list of browser servers have been processed by this workstation. | PERF_COUNTER_COUNTER |

# Cache Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Cache performance object consists of counters that monitor the file system cache, an area of physical memory that stores recently used data as long as possible to permit access to the data without having to read from the disk. Because applications typically use the cache, the cache is monitored as an indicator of application input/output (I/O) operations. When memory is plentiful, the cache can grow, but when memory is scarce, the cache can become too small to be effective.

| **Counter Name** | **Description** | **Counter Type** |
|---|---|---|
| Async Copy Reads/sec | Shows the rate at which read operations from pages of the file system cache involve a memory copy of the data from the cache to the application's buffer. The application regains control immediately, even if the disk must be accessed to retrieve the page. | PERF_COUNTER_COUNTER |
| Async Data Maps/sec | Shows the rate at which an application that uses a file system (such as the NTFS file system) to map a page of a file into the file system cache does not wait for the page to be retrieved, if the page is not in main memory. | PERF_COUNTER_COUNTER |
| Async Fast Reads/sec | Shows the rate at which read operations from the file system cache bypass the installed file system and retrieve the data directly from the cache. Normally, file I/O requests invoke the file system to retrieve data from a file, but this path permits data to be retrieved from the cache directly (without file system involvement) if the data is in the cache. Even if the data is not in the cache, one invocation of the file system is avoided. If the data is not in the cache, the request does not wait until the data has been retrieved from the disk, but gets control immediately. | PERF_COUNTER_COUNTER |
| Async MDL Reads/sec | Shows the rate at which read operations from the file system cache use a memory descriptor list (MDL) to access the pages. The MDL contains the physical address of each page in the transfer, thus permitting direct memory access (DMA) of the pages. If the accessed pages are not in main memory, the calling application program does not wait for the pages to be retrieved from disk through a page fault. | PERF_COUNTER_COUNTER |
| Async Pin Reads/sec | Shows the rate at which data is read into the file system cache before writing the data back to disk. Pages read in this fashion are "pinned" in memory at the completion of the read. *Pinned* pages are those that are read into the file system cache before the system writes data back to the disk at the completion of the read. The file system regains control immediately, even if the disk must be accessed to retrieve the page. While pinned, a page's physical address is not altered. | PERF_COUNTER_COUNTER |
| Copy Read Hits % | Shows the percentage of cache copy read requests that did not require a disk read to access the page in the cache. A *copy read* is a file read operation that is satisfied by a memory copy from a page in the cache to the application's buffer. The LAN redirector uses this method for retrieving information from the cache, as does the LAN server for small transfers. This method is also used by the disk file systems. | PERF_SAMPLE_FRACTION |
| Copy Reads/sec | Shows the rate at which read operations from pages of the file system cache involve a copy read. | PERF_COUNTER_COUNTER |
| Data Flush Pages/sec | Shows the rate at which the file system cache has flushed its contents to disk in response to a request to flush, or to satisfy a write-through file write request. More than one page can be transferred on each flush operation. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Data Flushes/sec | Shows the rate at which the file system cache has flushed its contents to disk in response to a request to flush, or to satisfy a write-through file write request. | PERF_COUNTER_COUNTER |
| Data Map Hits % | Shows the percentage of data maps in the file system cache that can be resolved without having to retrieve a page from the disk, because the page is already in physical memory. | PERF_SAMPLE_FRACTION |
| Data Map Pins/sec | Shows the rate at which data maps in the file system cache resulted in pinning a page in main memory. | PERF_SAMPLE_FRACTION |
| Data Maps/sec | Shows the rate at which a file system, such as NTFS, maps a page of a file into the file system cache to read the page. | PERF_COUNTER_COUNTER |
| Fast Read Not Possibles/sec | Shows the rate at which an application programming interface (API) function call attempts to bypass the file system to get to data in the file system cache that required invoking the file system. | PERF_COUNTER_COUNTER |
| Fast Read Resource Misses/sec | Shows the rate at which cache misses occur because there are not enough resources to satisfy the request. | PERF_COUNTER_COUNTER |
| Fast Reads/sec | Shows the rate at which read operations from the file system cache bypass the installed file system and retrieve the data directly from the cache. | PERF_COUNTER_COUNTER |
| Lazy Write Flushes/sec | Shows the rate at which the Lazy Writer thread writes to disk. | PERF_COUNTER_COUNTER |
| Lazy Write Pages/sec | Shows the rate at which the Lazy Writer thread has written to disk. | PERF_COUNTER_COUNTER |
| MDL Read Hits % | Shows the percentage of MDL read requests to the file system cache that did not have to access the disk to provide memory access to the page or pages in the cache. | PERF_SAMPLE_FRACTION |
| MDL Reads/sec | Shows the rate at which read operations from the file system cache use an MDL to access the data. The LAN server uses this method for large transfers out of the server. | PERF_COUNTER_COUNTER |
| Pin Read Hits % | Shows the percentage of pin read requests that did not have to access the disk to provide access to the page in the file system cache. | PERF_SAMPLE_FRACTION |
| Pin Reads/sec | Shows the rate at which data is read into the file system cache before writing the data back to disk. Pages read in this fashion are pinned in memory at the completion of the read. | PERF_COUNTER_COUNTER |
| Read Aheads/sec | Shows the rate at which read operations from the file system cache detect sequential access to a file. *Read aheads* permit data to be transferred in larger blocks than those being requested by the application, reducing the overhead per access. | PERF_COUNTER_COUNTER |
| Sync Copy Reads/sec | Shows the rate at which read operations from pages of the file system cache involve a memory copy of the data from the cache to the application's buffer. The file system does not regain control until the copy operation is complete, even if the disk must be accessed to retrieve the page. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Sync Data Maps/sec | Shows the rate at which a file system (such as NTFS) maps a page of a file into the file system cache to read the page, and — if the page is not in main memory — waits for it to be retrieved. | PERF_COUNTER_COUNTER |
| Sync Fast Reads/sec | Shows the rate at which read operations from the file system cache bypass the installed file system and retrieve the data directly from the cache. | PERF_COUNTER_COUNTER |
| Sync MDL Reads/sec | Shows the rate at which read operations from the file system cache use an MDL to access the pages. | PERF_COUNTER_COUNTER |
| Sync Pin Reads/sec | Shows the rate at which data is read into the file system cache before it is written back to disk. Pages read in this fashion are pinned in memory at the completion of the read. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Finding Cache Bottlenecks
Find memory bottlenecks
Windows Server 2003 System Performance Counters

# Client Service for NetWare Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Client Service for NetWare performance object consists of counters that measure packet transmission rates, logon attempts, and connections.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Received/sec | Shows the rate at which bytes are coming in to the redirector from the network. This counter includes all application data in addition to network protocol information (such as packet headers). | PERF_COUNTER_BULK_COUNT |
| Bytes Total/sec | Shows the rate at which the redirector is processing data bytes. This counter includes all application and file data in addition to network protocol information (such as packet headers). | PERF_COUNTER_BULK_COUNT |
| Bytes Transmitted/sec | Shows the rate at which bytes are leaving the redirector to the network. This counter includes all application data in addition to network protocol information (such as packet headers). | PERF_COUNTER_BULK_COUNT |
| Connect NetWare 2.x | Shows the number of connections to NetWare 2.*x* servers. | PERF_COUNTER_RAWCOUNT |
| Connect NetWare 3.x | Shows the number of connections to NetWare 3.*x* servers. | PERF_COUNTER_RAWCOUNT |
| Connect NetWare 4.x | Shows the number of connections to NetWare 4.*x* servers. | PERF_COUNTER_RAWCOUNT |
| File Data Operations/sec | Shows the rate at which the redirector is processing data operations. In general, one operation includes many bytes, although this may not always be the case because of operation overhead. You can determine the efficiency of this path by dividing the value of Bytes/sec by the value of File Data Operations/sec to determine the average number of bytes transferred per operation. | PERF_COUNTER_COUNTER |
| File Read Operations/sec | Shows the rate at which applications are asking the redirector for data. Each call to a file system (or similar application programming interface (API) call) counts as one operation. | PERF_COUNTER_COUNTER |
| File Write Operations/sec | Shows the rate at which applications are sending data to the redirector. Each call to a file system (or similar API call) counts as one operation. | PERF_COUNTER_COUNTER |
| Packet Burst IO/sec | Shows the sum of the values of Packet Burst Read NCP Count/sec and Packet Burst Write NCP Count/sec. | PERF_COUNTER_COUNTER |
| Packet Burst Read NCP Count/sec | Shows the rate at which NetWare Core Protocol (NCP) requests packet burst reads. The windowing Packet Burst Protocol is built on top of Internetwork Packet Exchange (IPX). It can speed the transfer of multiple-packet NCP file reads and writes. | PERF_COUNTER_COUNTER |
| Packet Burst Read Timeouts/sec | Shows the rate at which the Client Service for NetWare was required to retransmit a packet burst read request because the NetWare server took too long to respond. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Packet Burst Write NCP Count/sec | Shows the rate at which NCP requests packet burst writes. | PERF_COUNTER_COUNTER |
| Packet Burst Write Timeouts/sec | Shows the rate at which the Client Service for NetWare was required to retransmit a burst write request because the NetWare server took too long to respond. | PERF_COUNTER_COUNTER |
| Packets Received/sec | Shows the rate at which the redirector receives packets, also called server message blocks (SMBs). Network transmissions are divided into packets. The average number of bytes received in a packet can be obtained by dividing the value of Bytes Received/sec by the value of Packets Received/sec. Some packets received might not contain incoming data (for example, an acknowledgment to a write operation made by the redirector would count as an incoming packet). | PERF_COUNTER_BULK_COUNT |
| Packets Transmitted/sec | Shows the rate at which the redirector sends packets. The average number of bytes transmitted in a packet can be obtained by dividing the value of Bytes Transmitted/sec by the value of Packets Transmitted/sec. | PERF_COUNTER_BULK_COUNT |
| Packets/sec | Shows the rate at which the redirector processes data packets. In general, one packet includes many bytes, although this may not always be the case because each packet has protocol overhead. The average number of bytes transferred in a packet can be obtained by dividing the value of Bytes/sec by the value of Packets/sec. You can also divide this counter by the value of Operations/sec to determine the average number of packets per operation, another measure of efficiency. | PERF_COUNTER_BULK_COUNT |
| Read Operations Random/sec | Shows the rate at which, on a file-by-file basis, nonsequential read operations are made. If a read is made using a particular file handle, and then is followed by another read that is not the next contiguous byte, this counter is incremented by one. | PERF_COUNTER_COUNTER |
| Read Packets/sec | Shows the rate at which read packets are being placed on the network. Each time a single packet is sent with a request to read data remotely, this counter is incremented by one. | PERF_COUNTER_COUNTER |
| Server Disconnects | Shows the number of times that a server has disconnected the redirector. | PERF_COUNTER_RAWCOUNT |
| Server Reconnects | Shows the number of times that the redirector reconnected to a server to complete a new active request. The redirector may be disconnected by the server if it remains inactive for too long. Even if all of your remote files are closed, the redirector keeps your connections intact on your local computer for about 10 minutes. Such inactive connections are called dormant connections. Reconnecting may take a substantial period of time. | PERF_COUNTER_RAWCOUNT |
| Server Sessions | Shows the total number of security objects managed by the redirector. Logging on to a server followed by a network access to the same server establishes one connection, but two sessions. Each session corresponds to one security object. | PERF_COUNTER_RAWCOUNT |
| Write Operations Random/sec | Shows the rate at which, on a file-by-file basis, nonsequential write operations are made. If a write is made using a particular file handle, and then is followed by another write that is not the next contiguous byte, this counter is incremented by one. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Write Packets/sec | Shows the rate at which write operations are sent to the network. Each time a single packet is sent with a request to write remote data, this counter is incremented by one. | PERF_COUNTER_COUNTER |

# Database Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Database performance object consists of counters that describe the activity of the Extensible Storage Engine (ESE), a transacted database management system that stores all Microsoft Active Directory objects. The objects in the ESE database can access the Database performance object, which provides counters to measure the activity of the ESE. The Database object is not installed by default; use Lodctr.exe in the Microsoft Platform Software Development Kit (SDK) to install it. For more information about the Platform SDK, see the Software Development Kit (SDK) information in the MSDN Library link on the Web Resources page.

These counters provide performance statistics only for processes that use the ESE. This section does not describe counters that are for developer use only.

| Counter Name | Description | Counter Type |
|---|---|---|
| Cache % Hit | Shows the percentage of database file page requests that were fulfilled by the database cache manager without causing a file operation. If this percentage is too low, the database cache size may be too small. | PERF_SAMPLE_FRACTION |
| Cache Page Fault Stalls/sec | Shows the number of page faults that cannot be serviced because there are no pages available for allocation from the database cache. If this counter is nonzero most of the time, the clean threshold may be too low. | PERF_COUNTER_COUNTER |
| Cache Page Faults/sec | Shows the number of database file page requests that require the database cache manager to allocate a new page from the database cache. If this rate is too high, the database cache size may be too small. | PERF_COUNTER_COUNTER |
| Cache Size | Shows the amount of system memory used by the database cache manager to hold commonly used information from the database file(s), in order to reduce the number of file operations the system must perform. If the database cache size seems to be too small for optimal performance and there is very little available memory on the system (see Memory\Available Bytes), adding more memory to the system may increase performance. If there is a lot of available memory on the system and the database cache size does not grow beyond a certain point, the database cache size may be too small. Increasing this limit may increase performance. | PERF_COUNTER_RAWCOUNT |
| File Bytes Read/sec | Shows the rate at which bytes are read from the database file(s) into the database cache. If this rate is too high, the database cache size may be too small. | PERF_COUNTER_COUNTER |
| File Bytes Written/sec | Shows the rate at which bytes are written to the database file(s) from the database cache. If this rate is too high, the database cache size may be too small. | PERF_COUNTER_COUNTER |
| File Operations Pending | Shows the number of reads and writes issued by the database cache manager to the database file(s) that are currently being processed by the operating system. A large number of pending file operations can increase system throughput but will increase response time for individual operations. Typically, a larger number means that file operations to the database file(s) may be a bottleneck. | PERF_COUNTER_RAWCOUNT |
| File Operations/sec | Shows the number of reads and writes issued by the database cache manager to the database file(s). If this rate is too high, the database cache size may be too small. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Log Record Stalls/sec | Shows the number of log records that cannot be added to the log buffers because the buffers are full. If this counter is nonzero most of the time, the log buffer size may be a bottleneck. | PERF_COUNTER_COUNTER |
| Log Threads Waiting | Shows the number of threads waiting for their data to be written to the log during a database update. If this number is too high, the log may be a bottleneck. | PERF_COUNTER_RAWCOUNT |
| Log Writes/sec | Shows the number of times that the contents of the log buffers are written to the log. If this number approaches the maximum write rate for the media holding the log, the log may be a bottleneck. | PERF_COUNTER_COUNTER |
| Table Open Cache % Hit | Shows the percentage of database tables opened by using cached schema information. If this percentage is too low, the table cache size may be too small. | PERF_SAMPLE_FRACTION |
| Table Open Cache Hits/sec | Shows the number of database tables opened by using cached schema information. If this rate is too low, the table cache size may be too small. | PERF_COUNTER_COUNTER |
| Table Open Cache Misses/sec | Shows the number of database tables opened without using cached schema information. If this rate is too high, the table cache size may be too small. | PERF_COUNTER_COUNTER |
| Table Opens/sec | Shows the number of database tables opened. | PERF_COUNTER_COUNTER |

# DHCP Server Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The DHCP Server performance object consists of counters that monitor DHCP service activity. The DHCP service enables the server to perform as a DHCP server and configure DHCP-enabled client computers on your network, as described in the DHCP standard. This object is available when the service is running.

| Counter Name | Description | Counter Type |
|---|---|---|
| Acks/sec | Shows the rate at which DHCP acknowledgements have been sent by the DHCP server. A sudden or abnormal increase in this number indicates that a large number of clients are being renewed by the DHCP server. This may indicate that scope lease times are too short. | PERF_COUNTER_COUNTER |
| Active Queue Length | Shows the number of packets in the processing queue of the DHCP server. This number equals the number of unprocessed messages received by the server. A large number may indicate heavy server traffic. | PERF_COUNTER_RAWCOUNT |
| Conflict Check Queue Length | Shows the number of packets waiting in the DHCP server queue due to conflict detection. This queue holds messages that cannot be responded to while the DHCP server performs address conflict detection. A large value may indicate that the DHCP server property **Conflict detection attempts** has been set too high, or that there is heavy lease traffic on the server. | PERF_COUNTER_RAWCOUNT |
| Declines/sec | Shows the rate at which DHCP declines have been received by the DHCP server. A high value indicates that several clients have found their addresses to be in conflict, possibly indicating network trouble. You may want to enable conflict detection on the DHCP server. Conflict detection should only be used temporarily on the server. | PERF_COUNTER_COUNTER |
| Discovers/sec | Shows the rate at which DHCP discovers have been received by the DHCP server. A sudden or abnormal increase indicates that a large number of clients are probably attempting to initialize and obtain an IP address lease from the server, such as when a number of client computers are started at one time. | PERF_COUNTER_COUNTER |
| Duplicates Dropped/sec | Shows the rate at which the DHCP server received duplicate packets. A large number indicates that clients are probably timing out too fast or that the server is not responding very fast. | PERF_COUNTER_COUNTER |
| Informs/sec | Shows the rate at which DHCP informs have been received by the DHCP server. DHCP inform messages are used when the DHCP server queries for the directory service for the enterprise root, and when dynamic update updates are being done on behalf of clients by the server. | PERF_COUNTER_COUNTER |
| Milliseconds per Packet (Avg). | Shows the average time per packet taken by the DHCP server to send a response. This number can vary depending on the server hardware and its input/output (I/O) subsystem. A sudden or unreasonable increase may indicate a problem, either with the I/O subsystem slowing or because of intrinsic processing overhead on the server computer. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Nacks/sec | Shows the rate at which DHCP negative acknowledgments (Nacks) have been sent by the DHCP server. A very high value may indicate potential network problems, such as misconfiguration of clients or the server. One possible cause of server problems is a deactivated scope. For clients, a very high value could be caused by computers moving between subnets, such as portable computers or other mobile devices. | PERF_COUNTER_COUNTER |
| Offers/sec | Shows that rate at which DHCP offers have been sent out by the DHCP server. A sudden or abnormal increase in this number indicates heavy traffic on the server. | PERF_COUNTER_COUNTER |
| Packets Expired/sec | Shows the rate at which packets expire in the DHCP server message queue. Packets expire because they are in the server's internal message queue for too long. A large number indicates that the server is either taking too long to process some packets while other packets are queued and becoming stale, or traffic on the network is too high for the DHCP server to handle. This could suggest a disk or memory bottleneck. | PERF_COUNTER_COUNTER |
| Packets Received/sec | Shows the rate at which packets are received by the DHCP server. A large number indicates heavy DHCP-related message traffic to the server. | PERF_COUNTER_COUNTER |
| Releases/sec | Shows the rate at which DHCP releases have been received by the DHCP server. This number is only incremented when a client manually releases its address, such as when the **ipconfig /release** command is used at the client computer. Because clients rarely release their addresses, this counter should not be high for most networks and configurations. | PERF_COUNTER_COUNTER |
| Requests/sec | Shows the rate at which DHCP requests have been received by the DHCP server. A sudden or abnormal increase in this number indicates that a large number of clients are probably trying to renew their leases with the DHCP server. This may indicate that scope lease times are too short. | PERF_COUNTER_COUNTER |

# Distributed Transaction Coordinator Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The performance object provides statistics about Microsoft Distributed Transaction Coordinator (MS DTC), part of Component Services (formerly known as Microsoft Transaction Server) that coordinates external (two-phase) transactions. Among other things, MS DTC is used to coordinate two-phase transactions by Message Queuing.

| Counter Name | Description | Counter Type |
|---|---|---|
| Aborted Transactions | Shows the number of transactions that were aborted. | PERF_COUNTER_RAWCOUNT |
| Aborted Transactions/sec | Shows the rate, in incidents per second, at which transactions were aborted. | PERF_COUNTER_RAWCOUNT |
| Active Transactions | Shows the number of transactions that were active during the sample period. | PERF_COUNTER_RAWCOUNT |
| Active Transactions Maximum | Shows the maximum number of transactions that were ever concurrently active. | PERF_COUNTER_RAWCOUNT |
| Committed Transactions | Shows the number of transactions that were committed. | PERF_COUNTER_RAWCOUNT |
| Committed Transactions/sec | Shows the rate, in incidents per second, at which transactions were committed. | PERF_COUNTER_RAWCOUNT |
| Force Aborted Transactions | Shows the number of transactions that were aborted by the system administrator. | PERF_COUNTER_RAWCOUNT |
| Force Committed Transactions | Shows the number of transactions that were committed by the system administrator. | PERF_COUNTER_RAWCOUNT |
| In Doubt Transactions | Shows the number of transactions that were in doubt; that is, neither committed nor aborted. | PERF_COUNTER_RAWCOUNT |
| Response Time – Average | Shows the average time between when transactions were begun and when they were committed. | PERF_COUNTER_RAWCOUNT |
| Response Time – Maximum | Shows the maximum time between when transactions were begun and when they were committed. | PERF_COUNTER_RAWCOUNT |
| Response Time – Minimum | Shows the minimum time between when transactions were begun and when they were committed. | PERF_COUNTER_RAWCOUNT |
| Transactions/sec | Shows the rate, in incidents per second, at which transactions were performed. | PERF_COUNTER_RAWCOUNT |

# DNS Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Domain Name System (DNS) is an Internet and TCP/IP standard name service. The DNS performance object consists of counters for monitoring various areas of performance of the DNS service, which enables client computers on your network to register and resolve DNS domain names. These names are used to find and access resources offered by other computers on your network or other networks, such as the Internet. The data reported by the DNS object includes the number of overall queries and responses processed by a DNS server, the number of DNS queries and responses processed using either the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP), the number of queries and responses that used recursion to look up and fully resolve DNS names on behalf of requesting clients, and the number of queries and responses made to servers using the WINS lookup integration features of the DNS service.

| Counter Name | Description | Counter Type |
|---|---|---|
| AXFR Request Received | Shows the total number of full zone transfer requests that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| AXFR Request Sent | Shows the total number of full zone transfer requests that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| AXFR Response Received | Shows the total number of full zone transfer responses that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| AXFR Success Received | Shows the total number of successful full zone transfers that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| AXFR Success Sent | Shows the total number of successful full zone transfers that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Caching Memory | Shows the total amount of caching memory that was used by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Database Node Memory | Shows the total amount of database node memory that was used by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Dynamic Update NoOperation | Shows the total number of NoOperation/empty dynamic update requests that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Dynamic Update NoOperation/sec | Shows the rate, in incidents per second, at which NoOperation/empty dynamic update requests were received by the DNS server. | PERF_COUNTER_COUNTER |
| Dynamic Update Queued | Shows the total number of dynamic updates that were queued by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Dynamic Update Received | Shows the total number of dynamic update requests that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Dynamic Update Received/sec | Shows the rate, in incidents per second, at which dynamic update requests were received by the DNS server. | PERF_COUNTER_COUNTER |
| Dynamic Update Rejected | Shows the total number of dynamic updates that were rejected by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Dynamic Update TimeOuts | Shows the total number of dynamic update timeouts of the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Dynamic Update Written to Database | Shows the total number of dynamic updates that were written to the database by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Dynamic Update Written to Database/sec | Shows the rate, in incidents per second, at which dynamic updates were written to the database by the DNS server. | PERF_COUNTER_COUNTER |
| IXFR Request Received | Shows the total number of incremental zone transfer requests that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| IXFR Request Sent | Shows the total number of incremental zone transfer requests that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| IXFR Response Received | Shows the total number of incremental zone transfer responses that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| IXFR Success Received | Shows the total number of successful incremental zone transfers that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| IXFR Success Sent | Shows the total number of successful incremental zone transfers that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| IXFR TCP Success Received | Shows the total number of successful TCP incremental zone transfers that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| IXFR UDP Success Received | Shows the total number of successful UDP incremental zone transfers that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Nbstat Memory | Shows the total amount of Nbstat memory that was used by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Notify Received | Shows the total number of notifies that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Notify Sent | Shows the total number of notifies that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Record Flow Memory | Shows the total record flow memory that was used by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Recursive Queries | Shows the total number of recursive queries that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Recursive Queries/sec | Shows the rate, in incidents per second, at which recursive queries were received by the DNS server. | PERF_COUNTER_COUNTER |
| Recursive Query Failure | Shows the total number of failures that occurred while processing recursive queries, since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Recursive Query Failure/sec | Shows the rate, in incidents per second, of failures while processing recursive queries. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Recursive Send TimeOuts | Shows the total number of recursive queries sending timeouts since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Recursive TimeOut/sec | Shows the rate, in incidents per second, of recursive queries sending timeouts. | PERF_COUNTER_COUNTER |
| Secure Update Failure | Shows the total number of secure update failures of the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Secure Update Received | Shows the total number of secure update requests that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Secure Update Received/sec | Shows the rate, in incidents per second, at which secure update requests were received by the DNS server. | PERF_COUNTER_COUNTER |
| TCP Message Memory | Shows the total amount of TCP message memory that was used by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| TCP Query Received | Shows the total number of TCP queries that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| TCP Query Received/sec | Shows the rate, in incidents per second, at which TCP queries were received by the DNS server. | PERF_COUNTER_COUNTER |
| TCP Response Sent | Shows the total number of TCP responses that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| TCP Response Sent/sec | Shows the rate, in incidents per second, at which TCP responses are were sent by the DNS server. | PERF_COUNTER_COUNTER |
| Total Query Received | Shows the total number of queries that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Query Received/sec | Shows the rate, in incidents per second, at which queries were received by the DNS server. | PERF_COUNTER_COUNTER |
| Total Response Sent | Shows the total number of responses that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Response Sent/sec | Shows the rate, in incidents per second, at which responses were sent by the DNS server. | PERF_COUNTER_COUNTER |
| UDP Message Memory | Shows the total amount of UDP message memory that was used by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| UDP Query Received | Shows the total number of UDP queries that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| UDP Query Received/sec | Shows the rate, in incidents per second, at which UDP queries were received by the DNS server. | PERF_COUNTER_COUNTER |
| UDP Response Sent | Shows the total number of UDP responses that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| UDP Response Sent/sec | Shows the rate, in incidents per second, at which UDP responses were sent by the DNS server. | PERF_COUNTER_COUNTER |
| WINS Lookup Received | Shows the total number of WINS lookup requests that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| WINS Lookup Received/sec | Shows the rate, in incidents per second, at which WINS lookup requests were received by the DNS server. | PERF_COUNTER_COUNTER |
| WINS Response Sent | Shows the total number of WINS lookup responses that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| WINS Response Sent/sec | Shows the rate, in incidents per second, at which WINS lookup responses were sent by the server. | PERF_COUNTER_COUNTER |
| WINS Reverse Lookup Received | Shows the total number of WINS reverse lookup requests that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| WINS Reverse Lookup Received/sec | Shows the rate, in incidents per second, at which WINS reverse lookup requests were received by the DNS server. | PERF_COUNTER_RAWCOUNT |
| WINS Reverse Response Sent | Shows the total number of WINS reverse lookup responses that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| WINS Reverse Response Sent/sec | Shows the rate, in incidents per second, at which WINS reverse lookup responses were sent by the server. | PERF_COUNTER_COUNTER |
| Zone Transfer Failure | Shows the total number of failed zone transfers that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Zone Transfer Request Received | Shows the total number of zone transfer requests that were received by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Zone Transfer SOA Request Sent | Shows the total number of zone transfer start of authority (SOA) requests that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Zone Transfer Success | Shows the total number of successful zone transfers that were sent by the DNS server since the service was last started. | PERF_COUNTER_RAWCOUNT |

# Fax Service Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Shared Fax Service performance object tracks fax activity while the Fax service is running. It is recommended that you use the counters in conjunction with event log data about fax events to assess service performance. The following counters can indicate system problems that prevent the fax from being sent or received, such as installation of the wrong modem driver.

| Counter Name | Description | Counter Type |
|---|---|---|
| Inbound Bytes | Shows the number of inbound bytes received by Fax Service. | PERF_COUNTER_RAWCOUNT |
| Inbound Failed Receptions | Shows the number of faxes that were not received, such as a call that was answered without a fax being received. | PERF_COUNTER_RAWCOUNT |
| Inbound Faxes | Shows the number of inbound faxes received by Fax Service. | PERF_COUNTER_RAWCOUNT |
| Inbound Minutes Receiving | Shows the amount of time, in minutes, spent receiving inbound messages. | PERF_COUNTER_RAWCOUNT |
| Inbound Pages | Shows the number of inbound pages received by Fax Service. | PERF_COUNTER_RAWCOUNT |
| Outbound Bytes | Shows the number of outbound bytes sent by Fax Service. | PERF_COUNTER_RAWCOUNT |
| Outbound Failed Connections | Shows the number of faxes that were not sent because there was no dial tone, no answer, busy, or another unknown error. | PERF_COUNTER_RAWCOUNT |
| Outbound Failed Transmissions | Shows the number of faxes that could not be sent. | PERF_COUNTER_RAWCOUNT |
| Outbound Faxes | Shows the number of outbound faxes sent by Fax Service. | PERF_COUNTER_RAWCOUNT |
| Outbound Minutes Sending | Shows the amount of time, in minutes, spent sending outbound messages. | PERF_COUNTER_RAWCOUNT |
| Outbound Pages | Shows the number of outbound pages sent by Fax Service. | PERF_COUNTER_RAWCOUNT |
| Total Bytes | Shows the total number of bytes sent and received by Fax Service. | PERF_COUNTER_RAWCOUNT |
| Total Faxes | Shows the total number of faxes sent and received by Fax Service. | PERF_COUNTER_RAWCOUNT |
| Total Minutes Sending/Receiving | Shows the total amount of time, in minutes, spent sending and receiving. | PERF_COUNTER_RAWCOUNT |
| Total Pages | Shows the total number of pages sent and received by Fax Service. | PERF_COUNTER_RAWCOUNT |

# FileReplicaConn Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The FileReplicaConn performance object displays performance statistics of the Replicaconn object that defines replica connections to Distributed File System (DFS) roots.

| Counter Name | Description | Counter Type |
|---|---|---|
| Authentications | Shows the number of authentications that were performed. | PERF_COUNTER_COUNTER |
| Authentications in Error | Shows the number of authentications that were performed in error. | PERF_COUNTER_COUNTER |
| Bindings | Shows the number of bindings that were completed. | PERF_COUNTER_COUNTER |
| Bindings in Error | Shows the number of bindings that were completed in error. | PERF_COUNTER_COUNTER |
| Communication Timeouts | Shows the number of communications that were timed out. | PERF_COUNTER_COUNTER |
| Fetch Blocks Received | Shows the number of fetch blocks that were received. | PERF_COUNTER_COUNTER |
| Fetch Blocks Received in Bytes | Shows the number of fetch blocks, in bytes, that were received. | PERF_COUNTER_BULK_COUNT |
| Fetch Blocks Sent | Shows the number of fetch blocks that were sent. | PERF_COUNTER_COUNTER |
| Fetch Blocks Sent in Bytes | Shows the number of fetch blocks, in bytes, that were sent. | PERF_COUNTER_BULK_COUNT |
| Fetch Requests Received | Shows the number of fetch requests that were received. | PERF_COUNTER_COUNTER |
| Fetch Requests Sent | Shows the number of fetch requests that were sent. | PERF_COUNTER_COUNTER |
| Inbound Change Orders Dampened | Shows the number of inbound change orders that were dampened. | PERF_COUNTER_COUNTER |
| Join Notifications Received | Shows the number of join notifications that were received. | PERF_COUNTER_COUNTER |
| Join Notifications Sent | Shows the number of join notifications that were sent. | PERF_COUNTER_COUNTER |
| Joins | Shows the number of joins. | PERF_COUNTER_COUNTER |
| Local Change Orders Sent | Shows the number of local change orders that were sent. | PERF_COUNTER_COUNTER |
| Local Change Orders Sent At Join | Shows the number of local change orders that were sent at join. | PERF_COUNTER_COUNTER |
| Outbound Change Orders Dampened | Shows the number of outbound change orders that were dampened. | PERF_COUNTER_COUNTER |
| Packets Sent | Shows the number of packets that were sent. | PERF_COUNTER_COUNTER |
| Packets Sent in Bytes | Shows the number of packets, in bytes, that were sent. | PERF_COUNTER_BULK_COUNT |

| | | |
|---|---|---|
| Packets Sent in Error | Shows the number of packets that were sent in error. | PERF_COUNTER_COUNTER |
| Remote Change Orders Received | Shows the number of remote change orders that were received. | PERF_COUNTER_COUNTER |
| Remote Change Orders Sent | Shows the number of remote change orders that were sent. | PERF_COUNTER_COUNTER |
| Unjoins | Shows the number of unjoins. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
FileReplicaSet Object
Counters by Feature

# FileReplicaSet Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The FileReplicaSet performance object displays performance statistics of the Replicaset object that defines a *replica set* (one or more *replicas*, shared volumes or directories that store duplicates of the contents of an original share). The object reports statistics for the computer that constitutes the original share. For example, in a bidirectional ring topology of three computers A, B, and C, where computer A and computer C both replicate from computer B, the FileReplicaSet counters would provide data for computer B.

| Counter Name | Description | Counter Type |
|---|---|---|
| Authentications | Shows the number of authentications that were performed. | PERF_COUNTER_COUNTER |
| Authentications in Error | Shows the number of authentications that were performed in error. | PERF_COUNTER_COUNTER |
| Bindings | Shows the number of bindings that were completed. | PERF_COUNTER_COUNTER |
| Bindings in Error | Shows the number of bindings that were completed in error. | PERF_COUNTER_COUNTER |
| Bytes of Files Installed | Shows the number of bytes of files that were installed. | PERF_COUNTER_BULK_COUNT |
| Bytes of Staging Fetched | Shows the number of bytes of staging that were fetched. Staging is the temporary storage of files prior to replication. | PERF_COUNTER_BULK_COUNT |
| Bytes of Staging Generated | Shows the number of bytes of staging that were generated. | PERF_COUNTER_BULK_COUNT |
| Bytes of Staging Regenerated | Shows the number of bytes of staging that were regenerated. | PERF_COUNTER_BULK_COUNT |
| Change Orders Aborted | Shows the number of change orders that were aborted. | PERF_COUNTER_COUNTER |
| Change Orders Evaporated | Shows the number of change orders that were evaporated. | PERF_COUNTER_COUNTER |
| Change Orders Issued | Shows the number of change orders that were issued. | PERF_COUNTER_COUNTER |
| Change Orders Morphed | Shows the number of change orders that were morphed. | PERF_COUNTER_COUNTER |
| Change Orders Propagated | Shows the number of change orders that were propagated. | PERF_COUNTER_COUNTER |
| Change Orders Received | Shows the number of change orders that were received. This is an important counter to monitor. In an idle state this counter should be zero. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Change Orders Retired | Shows the number of change orders that were retired. | PERF_COUNTER_COUNTER |
| Change Orders Retried | Shows the number of change orders that were retried. | PERF_COUNTER_COUNTER |
| Change Orders Retried at Fetch | Shows the number of change orders that were retried at fetch. | PERF_COUNTER_COUNTER |
| Change Orders Retried at Generate | Shows the number of change orders that were retried at generate. | PERF_COUNTER_COUNTER |
| Change Orders Retried at Install | Shows the number of change orders that were retried at install. | PERF_COUNTER_COUNTER |
| Change Orders Retried at Rename | Shows the number of change orders that were retried at rename. | PERF_COUNTER_COUNTER |
| Change Orders Sent | Shows the number of change orders that were sent. This is an important counter to monitor. In an idle state this counter should be zero. | PERF_COUNTER_COUNTER |
| Communication Timeouts | Shows the number of communications that were timed out. | PERF_COUNTER_COUNTER |
| DS Bindings | Shows the number of directory service bindings. | PERF_COUNTER_COUNTER |
| DS Bindings in Error | Shows the number of directory service bindings that were in error. | PERF_COUNTER_COUNTER |
| DS Objects | Shows the number of directory service objects. | PERF_COUNTER_COUNTER |
| DS Objects in Error | Shows the number of directory service objects that were in error. | PERF_COUNTER_COUNTER |
| DS Polls | Shows the number of directory service polls. | PERF_COUNTER_COUNTER |
| DS Polls with Changes | Shows the number of directory service polls with changes. | PERF_COUNTER_COUNTER |
| DS Polls without Changes | Shows the number of directory service polls without changes. | PERF_COUNTER_COUNTER |
| DS Searches | Shows the number of directory service searches. | PERF_COUNTER_COUNTER |
| DS Searches in Error | Shows the number of directory service searches that were in error. | PERF_COUNTER_COUNTER |
| Fetch Blocks Received | Shows the number of fetch blocks that were received. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Fetch Blocks Received in Bytes | Shows the number of fetch blocks, in bytes, that were received. | PERF_COUNTER_BULK_COUNT |
| Fetch Blocks Sent | Shows the number of fetch blocks that were sent. | PERF_COUNTER_COUNTER |
| Fetch Blocks Sent in Bytes | Shows the number of fetch blocks, in bytes, that were sent. | PERF_COUNTER_BULK_COUNT |
| Fetch Requests Received | Shows the number of fetch requests that were received. | PERF_COUNTER_COUNTER |
| Fetch Requests Sent | Shows the number of fetch requests that were sent. | PERF_COUNTER_COUNTER |
| Files Installed | Shows the number of files that were installed. This is an important counter to monitor. In an idle state this counter should be zero. | PERF_COUNTER_COUNTER |
| Files Installed with Error | Shows the number of files that were installed with error. | PERF_COUNTER_COUNTER |
| Inbound Change Orders Dampened | Shows the number of inbound change orders that were dampened. | PERF_COUNTER_COUNTER |
| Join Notifications Received | Shows the number of join notifications that were received. | PERF_COUNTER_COUNTER |
| Join Notifications Sent | Shows the number of join notifications that were sent. | PERF_COUNTER_COUNTER |
| Joins | Shows the number of joins. | PERF_COUNTER_COUNTER |
| KB of Staging Space Free | Shows the number of kilobytes (KB) of staging space that were free. This is an important counter to monitor. In an idle state, this counter should be 660 KB. | PERF_COUNTER_LARGE_RAWCOUNT |
| KB of Staging Space In Use | Shows the number of kilobytes (KB) of staging space that were in use. This is an important counter to monitor. If all available staging space is in use, then replication stops. | PERF_COUNTER_LARGE_RAWCOUNT |
| Local Change Orders Aborted | Shows the number of local change orders that were aborted. | PERF_COUNTER_COUNTER |
| Local Change Orders Issued | Shows the number of local change orders that were issued. | PERF_COUNTER_COUNTER |
| Local Change Orders Morphed | Shows the number of local change orders that were morphed. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Local Change Orders Propagated | Shows the number of local change orders that were propagated. | PERF_COUNTER_COUNTER |
| Local Change Orders Retired | Shows the number of local change orders that were retired. | PERF_COUNTER_COUNTER |
| Local Change Orders Retried | Shows the number of local change orders that were retried. | PERF_COUNTER_COUNTER |
| Local Change Orders Retried at Fetch | Shows the number of local change orders that were retried at fetch. | PERF_COUNTER_COUNTER |
| Local Change Orders Retried at Generate | Shows the number of local change orders that were retried at generate. | PERF_COUNTER_COUNTER |
| Local Change Orders Retried at Install | Shows the number of local change orders that were retried at install. | PERF_COUNTER_COUNTER |
| Local Change Orders Retried at Rename | Shows the number of local change orders that were retried at rename. | PERF_COUNTER_COUNTER |
| Local Change Orders Sent | Shows the number of local change orders that were sent. | PERF_COUNTER_COUNTER |
| Local Change Orders Sent At Join | Shows the number of local change orders that were sent at join. | PERF_COUNTER_COUNTER |
| Outbound Change Orders Dampened | Shows the number of outbound change orders that were dampened. | PERF_COUNTER_COUNTER |
| Packets Received | Shows the number of packets (change notifications, file data, or other command packets) that were received. This is an important counter to monitor. In an idle state, there should be no packets received unless a computer is having trouble joining with other computers in the replica set. | PERF_COUNTER_COUNTER |
| Packets Received in Bytes | Shows the number of packets, in bytes, that were received. | PERF_COUNTER_BULK_COUNT |
| Packets Received in Error | Shows the number of packets received that were in error. | PERF_COUNTER_COUNTER |
| Packets Sent | Shows the number of packets that were sent. This is an important counter to monitor. | PERF_COUNTER_COUNTER |
| Packets Sent in Bytes | Shows the number of packets, in bytes, that were sent. | PERF_COUNTER_BULK_COUNT |

| | | |
|---|---|---|
| Packets Sent in Error | Shows the number of packets that were sent in error. | PERF_COUNTER_COUNTER |
| Remote Change Orders Aborted | Shows the number of remote change orders that were aborted. | PERF_COUNTER_COUNTER |
| Remote Change Orders Issued | Shows the number of remote change orders that were issued. | PERF_COUNTER_COUNTER |
| Remote Change Orders Morphed | Shows the number of remote change orders that were morphed. | PERF_COUNTER_COUNTER |
| Remote Change Orders Propagated | Shows the number of remote change orders that were propagated. | PERF_COUNTER_COUNTER |
| Remote Change Orders Received | Shows the number of remote change orders that were received. | PERF_COUNTER_COUNTER |
| Remote Change Orders Retired | Shows the number of remote change orders that were retired. | PERF_COUNTER_COUNTER |
| Remote Change Orders Retried | Shows the number of remote change orders that were retried. | PERF_COUNTER_COUNTER |
| Remote Change Orders Retried at Fetch | Shows the number of remote change orders that were retried at fetch. | PERF_COUNTER_COUNTER |
| Remote Change Orders Retried at Generate | Shows the number of remote change orders that were retried at generate. | PERF_COUNTER_COUNTER |
| Remote Change Orders Retried at Install | Shows the number of remote change orders that were retried at install. | PERF_COUNTER_COUNTER |
| Remote Change Orders Retried at Rename | Shows the number of remote change orders that were retried at rename. | PERF_COUNTER_COUNTER |
| Remote Change Orders Sent | Shows the number of remote change orders that were sent. | PERF_COUNTER_COUNTER |
| Replica Sets Created | Shows the number of replica sets that were created. | PERF_COUNTER_COUNTER |
| Replica Sets Deleted | Shows the number of replica sets that were deleted. | PERF_COUNTER_COUNTER |
| Replica Sets Removed | Shows the number of replica sets that were removed. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Replica Sets Started | Shows the number of replica sets that were started. | PERF_COUNTER_COUNTER |
| Staging Files Fetched | Shows the number of staging files that were fetched. | PERF_COUNTER_COUNTER |
| Staging Files Generated | Shows the number of staging files that were generated. | PERF_COUNTER_COUNTER |
| Staging Files Generated with Error | Shows the number of staging files that were generated with error. | PERF_COUNTER_COUNTER |
| Staging Files Regenerated | Shows the number of staging files that were regenerated. | PERF_COUNTER_COUNTER |
| Threads exited | Shows the number of threads that were exited. | PERF_COUNTER_COUNTER |
| Threads started | Shows the number of threads that were started. | PERF_COUNTER_COUNTER |
| Unjoins | Shows the number of unjoins. | PERF_COUNTER_COUNTER |
| USN Reads | Shows the number of update sequence number (USN) reads. | PERF_COUNTER_COUNTER |
| USN Records Accepted | Shows the number of USN records that were accepted. This is an important counter to monitor. Replication is triggered by entries to the NTFS USN journal. A high value on this counter, such as one every five seconds, indicates heavy replication traffic and may result in replication latency. | PERF_COUNTER_COUNTER |
| USN Records Examined | Shows the number of USN records that were examined. | PERF_COUNTER_COUNTER |
| USN Records Rejected | Shows the number of USN records that were rejected. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
FileReplicaConn Object
Counters by Feature

# FTP Service Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The FTP Service performance object consists of counters that show data about the Internet Information Services (IIS) FTP service.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Received/sec | Shows the rate at which data bytes are received by the FTP service at the application layer. This value does not include protocol headers or control bytes associated with the data received. | PERF_COUNTER_BULK_COUNT |
| Bytes Sent/sec | Shows the rate at which data bytes are sent by the FTP service at the application layer. This value does not include protocol headers or control bytes associated with the data sent. | PERF_COUNTER_BULK_COUNT |
| Bytes Total/sec | Shows the rate at which data bytes are sent and received by the FTP service at the application layer. This value does not include protocol headers or control bytes associated with the data transmitted. Bytes Total/sec is the sum of the values of Bytes Sent/sec and Bytes Received/sec. | PERF_COUNTER_BULK_COUNT |
| Current Anonymous Users | Shows the number of anonymous connections to the FTP service. *Anonymous connections* are connections whose requests either did not contain a user name and password, or whose user name and password were ignored because authentication is not permitted on the server. If anonymous connections are not permitted on the server, the value of Current Anonymous Users is zero. Current Anonymous Users shows successful connections only. If a client request for an anonymous connection is rejected and the client responds with valid authenticating data, the connection is counted as nonanonymous. There might be a lag between the counter value and the exact number of connections: new and closed connections might not be reflected immediately in the counter value. | PERF_COUNTER_RAWCOUNT |
| Current Connections | Shows the number of connections to the FTP service. There might be a lag between the counter value and the exact number of connections: new and closed connections might not be reflected immediately in the counter value. | PERF_COUNTER_RAWCOUNT |
| Current NonAnonymous Users | Shows the number of nonanonymous connections to the FTP service. *Nonanonymous connections* are connections whose requests contain a valid user name and password, or other authentication that is required by a custom authentication scheme. If authentication is not enabled on the server, and none of the applications that run on the server request or require authentication, the value of Current NonAnonymous Users is zero. Current NonAnonymous Users shows successful connections only. If a client request for an anonymous connection is rejected and the client responds with valid authenticating data, the connection is counted as nonanonymous. There might be a lag between the counter value and the exact number of connections: new and closed connections might not be reflected immediately in the counter value. | PERF_COUNTER_RAWCOUNT |
| Maximum Anonymous Users | Shows the largest number of anonymous connections to the FTP service maintained simultaneously since the FTP service was started. There might be a lag between the counter value and the exact number of connections: new and closed connections might not be reflected immediately in the counter value. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Maximum Connections | Shows the largest number of connections to the FTP service simultaneously maintained since the FTP service started. There might be a lag between the counter value and the exact number of connections: new and closed connections might not be reflected immediately in the counter value. | PERF_COUNTER_RAWCOUNT |
| Maximum NonAnonymous Users | Shows the largest number of nonanonymous connections to the FTP service maintained simultaneously since the FTP service was started. Maximum NonAnonymous Users shows successful connections only. If a client request for an anonymous connection is rejected and the client responds with valid authenticating data, the connection is counted as nonanonymous. There might be a lag between the counter value and the exact number of connections: new and closed connections might not be reflected immediately in the counter value. | PERF_COUNTER_RAWCOUNT |
| Total Anonymous Users | Shows the number of anonymous connections to the FTP service that have been established since the FTP service was started. Total Anonymous Users shows successful connections only. If a client request for an anonymous connection is rejected and the client responds with valid authenticating data, the connection is counted as nonanonymous. There might be a lag between the counter value and the exact number of connections: new and closed connections might not be reflected immediately in the counter value. | PERF_COUNTER_RAWCOUNT |
| Total Connection Attempts | Shows the total number of connections to the FTP service that have been established since the service was started. The value of this counter either remains the same or increases over time. This value does not include failed attempts to connect to the FTP server. To monitor all connection attempts, use the Connection counters on the TCP performance object. To monitor the number of connections to the FTP server that are currently active, use the FTP Service\Current Connections counter. | PERF_COUNTER_RAWCOUNT |
| Total Files Received | Shows the total number of files received by the FTP service since the service was started. | PERF_COUNTER_RAWCOUNT |
| Total Files Sent | Shows the total number of files sent by the FTP service since the service was started. | PERF_COUNTER_RAWCOUNT |
| Total Files Transferred | Shows the total number of files transmitted by the FTP service since the service was started. Total Files Transferred is the sum of the values of Total Files Sent and Total Files Received. | PERF_COUNTER_RAWCOUNT |
| Total Logon Attempts | Shows the total number of times a user logged on successfully to the FTP service since the service was started. This counter is incremented when a connection satisfies the criteria for logging on to FTP, such as providing user authentication requirements. The FTP service supports clear-text logon authentication. The value of Total Logon Attempts does not include failed attempts to log on to the FTP service. However, the value of Total Logon Attempts is a subset of the value of Total Connection Attempts. Failed logon attempts (situations where clients were able to connect but not to log on) can be calculated by subtracting the value of Total Logon Attempts from the value of Total Connection Attempts. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Total NonAnonymous Users | Shows the number of nonanonymous connections to the FTP service that have been established since the FTP service was started. Total NonAnonymous Users shows successful connections only. If a client request for an anonymous connection is rejected and the client responds with valid authenticating data, the connection is counted as nonanonymous. There might be a lag between the counter value and the exact number of connections: new and closed connections might not be reflected immediately in the counter value. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Counters by Feature

# Http Indexing Service Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Http Indexing Service object reports statistics about queries run by the Indexing Service through Internet Information Services (IIS). Unlike other Indexing Service counters, which can report values even if the Indexing Service is not being used in conjunction with a Web server, these counters only report values when Indexing Service queries are executed through the Web server.

| Counter Name | Description | Counter Type |
|---|---|---|
| % Cache Hits | Shows the percentage of queries that were found in the cache, out of the total number of times the cache was queried since the Indexing Service was lasted started. | PERF_RAW_FRACTION |
| % Cache Misses | Shows the percentage of times that queries were not found in the cache, out of the total number of times the cache was queried since the Indexing Service was lasted started. | PERF_RAW_FRACTION |
| Active Queries | Shows the current number of running queries since the Indexing Service was lasted started. | PERF_COUNTER_RAWCOUNT |
| Cache Items | Shows the number of completed queries that are stored in the cache since the Web Service was last started and the first query was posted. | PERF_COUNTER_RAWCOUNT |
| Current Requests Queued | Shows the current number of query requests that are pending. | PERF_COUNTER_RAWCOUNT |
| Queries per Minute | Shows the rate, in incidents per minute, of queries by the server. This counter can be used to determine server usage. | PERF_COUNTER_RAWCOUNT |
| Total Queries | Shows the total number of queries run since the service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Requests Rejected | Shows the total number of query requests that were rejected because the server was too busy. If this value begins to increase steadily, a bottleneck is developing. To alleviate this condition, reconfigure the server by adding more RAM, or use less resource-intensive queries. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Counters by Feature

# IAS Accounting Clients Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IAS Accounting Clients performance object is installed by the Internet Authentication Service (IAS). IAS uses the Remote Authentication Dial-In User Service (RADIUS) protocol to perform remote authentication. IAS performance objects report activity for servers or clients, including user authentication, authorization, and accounting (usage).

When selecting counters for this object, note that each client is a separate instance.

| Counter Name | Description | Counter Type |
|---|---|---|
| Accounting-Requests | Shows the number of packets received from this client on the accounting port. | PERF_COUNTER_RAWCOUNT |
| Accounting-Requests/sec | Shows the rate, in seconds, at which packets were received from this client on the accounting port. | PERF_COUNTER_COUNTER |
| Accounting-Responses | Shows the number of RADIUS Accounting-Response packets sent to this client. | PERF_COUNTER_RAWCOUNT |
| Accounting-Responses/sec | Shows the rate in seconds at which duplicate RADIUS Accounting-Request packets were received from this client. | PERF_COUNTER_COUNTER |
| Bad Authenticators | Shows the number of packets that contained invalid Signature attributes. | PERF_COUNTER_RAWCOUNT |
| Bad Authenticators/sec | Shows the rate, in seconds, at which packets that contained invalid Signature attributes were received. | PERF_COUNTER_COUNTER |
| Dropped Packets | Shows the number of incoming packets silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_RAWCOUNT |
| Dropped Packets/sec | Shows the rate in seconds at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_COUNTER |
| Duplicate Accounting-Requests | Shows the number of duplicate RADIUS Accounting-Request packets received from this client. | PERF_COUNTER_RAWCOUNT |
| Duplicate Accounting-Requests/sec | Shows the rate, in seconds, at which duplicate RADIUS Accounting-Request packets were received from this client. | PERF_COUNTER_COUNTER |
| Malformed Packets | Shows the number of malformed packets received. Bad authenticators or unknown types are not included. | PERF_COUNTER_RAWCOUNT |
| Malformed Packets/sec | Shows the rate, in seconds, at which malformed packets were received. | PERF_COUNTER_COUNTER |
| No Record | Shows the number of RADIUS Accounting-Request packets that were received and responded to but not recorded. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| No Record/sec | Shows the rate, in seconds, at which RADIUS Accounting-Request packets were received and responded to but not recorded. | PERF_COUNTER_COUNTER |
| Packets Received | Shows the number of packets received. | PERF_COUNTER_RAWCOUNT |
| Packets Received/sec | Shows the rate, in seconds, at which packets were received. | PERF_COUNTER_COUNTER |
| Packets Sent | Shows the number of packets sent. | PERF_COUNTER_RAWCOUNT |
| Packets Sent/sec | Shows the rate, in seconds, at which packets were sent. | PERF_COUNTER_COUNTER |
| Unknown Type | Shows the number of packets of unknown type that were received. An unknown type is a RADIUS packet that contains an unrecognized entry in the **Code** field, used to identify the packet type. Consequently IAS does not support packets of unknown type. | PERF_COUNTER_RAWCOUNT |
| Unknown Type/sec | Shows the rate, in seconds, at which packets of unknown type were received. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Internet Authentication Service

# IAS Accounting Proxy Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IAS Accounting Proxy performance object is installed by Internet Authentication Service (IAS). IAS uses the Remote Authentication Dial-In User Service (RADIUS) protocol. IAS performance objects report activity for servers or clients, including user authentication, authorization, and accounting (use).

| Counter Name | Description | Counter Type |
|---|---|---|
| Accounting-Requests | Shows the number of packets received from this client on the accounting port. | PERF_COUNTER_RAWCOUNT |
| Accounting-Requests/sec | Shows the rate, in incidents per second, at which packets were received from this client on the accounting port. | PERF_COUNTER_COUNTER |
| Accounting-Responses | Shows the number of RADIUS Accounting-Response packets sent to this client. | PERF_COUNTER_RAWCOUNT |
| Accounting-Responses/sec | Shows the rate, in incidents per second, at which RADIUS Accounting-Response packets were sent to this client. | PERF_COUNTER_COUNTER |
| Bad Authenticators | Shows the number of packets that contained invalid Message-Authenticator attributes. | PERF_COUNTER_RAWCOUNT |
| Bad Authenticators/sec | Shows the rate, in incidents per second, at which packets that contained invalid Message-Authenticator attributes were received. | PERF_COUNTER_COUNTER |
| Dropped Packets | Shows the number of incoming packets that were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_RAWCOUNT |
| Dropped Packets/sec | Shows the rate, in incidents per second, at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_COUNTER |
| Invalid Addresses | Shows the number of packets received from unknown addresses. | PERF_COUNTER_RAWCOUNT |
| Invalid Addresses/sec | Shows the rate, in incidents per second, at which packets were received from unknown addresses. | PERF_COUNTER_COUNTER |
| Malformed Packets | Shows the number of malformed packets received. Bad authenticators or unknown types are not included. | PERF_COUNTER_RAWCOUNT |
| Malformed Packets/sec | Shows the rate, in incidents per second, at which malformed packets were received. Bad authenticators or unknown types are not included. | PERF_COUNTER_COUNTER |
| Packets Received | Shows the number of packets received. | PERF_COUNTER_RAWCOUNT |
| Packets Received/sec | Shows the rate, in incidents per second, at which packets were received. | PERF_COUNTER_COUNTER |
| Pending Requests | Shows the number of requests destined for this server that have not yet timed out or received a response. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Request Timeouts | Shows the number of requests to this server that have timed out. | PERF_COUNTER_RAWCOUNT |
| Request Timeouts/sec | Shows the rate, in incidents per second, at which requests to this server have timed out. | PERF_COUNTER_COUNTER |
| Retransmissions | Shows the number of requests that were retransmitted to this server. | PERF_COUNTER_RAWCOUNT |
| Retransmissions/sec | Shows the rate, in incidents per second, at which requests were retransmitted to this server. | PERF_COUNTER_COUNTER |
| Unknown Type | Shows the number of packets of unknown type that were received. | PERF_COUNTER_RAWCOUNT |
| Unknown Type/sec | Shows the rate, in incidents per second, at which packets of unknown type were received. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Internet Authentication Service

# IAS Accounting Server Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IAS Accounting Server performance object is installed by Internet Authentication Service (IAS). IAS uses the Remote Authentication Dial-In User Service (RADIUS) protocol to perform remote authentication. IAS performance objects report activity for servers or clients, including user authentication, authorization, and accounting (usage).

When selecting counters for this object, note that each client is a separate instance.

| Counter Name | Description | Counter Type |
|---|---|---|
| Accounting-Requests | Shows the number of packets received from this client on the accounting port. | PERF_COUNTER_RAWCOUNT |
| Accounting-Requests/sec | Shows the rate, in incidents per second, at which packets were received from this client on the accounting port. | PERF_COUNTER_COUNTER |
| Accounting-Responses | Shows the number of RADIUS Accounting-Response packets sent to this client. | PERF_COUNTER_RAWCOUNT |
| Accounting-Responses/sec | Shows the rate, in incidents per second, at which RADIUS Accounting-Response packets were sent to this client. | PERF_COUNTER_COUNTER |
| Bad Authenticators | Shows the number of packets that contained invalid Signature attributes. | PERF_COUNTER_RAWCOUNT |
| Bad Authenticators/sec | Shows the rate, in incidents per second, at which packets that contained invalid Signature attributes were received. | PERF_COUNTER_COUNTER |
| Dropped Packets | Shows the number of incoming packets that were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_RAWCOUNT |
| Dropped Packets/sec | Shows the rate, in incidents per second, at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_COUNTER |
| Duplicate Accounting-Requests | Shows the number of duplicate RADIUS Accounting-Request packets that were received from this client. | PERF_COUNTER_RAWCOUNT |
| Duplicate Accounting-Requests/sec | Shows the rate, in incidents per second, at which duplicate RADIUS Accounting-Request packets were received from this client. | PERF_COUNTER_COUNTER |
| Invalid Requests | Shows the number of packets received from unknown addresses. | PERF_COUNTER_RAWCOUNT |
| Invalid Requests/sec | Shows the rate, in incidents per second, at which packets were received from unknown addresses. | PERF_COUNTER_COUNTER |
| Malformed Packets | Shows the number of malformed packets received. Bad authenticators or unknown types are not included. | PERF_COUNTER_RAWCOUNT |
| Malformed Packets/sec | Shows the rate, in incidents per second, at which malformed packets were received. Bad authenticators or unknown types are not included. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| No Record | Shows the number of RADIUS Accounting-Request packets that were received and responded to but not recorded. | PERF_COUNTER_RAWCOUNT |
| No Record/sec | Shows the rate, in incidents per second, at which RADIUS Accounting-Request packets were received and responded to but not recorded. | PERF_COUNTER_COUNTER |
| Packets Received | Shows the number of packets received. | PERF_COUNTER_RAWCOUNT |
| Packets Received/sec | Shows the rate, in incidents per second, at which packets were received. | PERF_COUNTER_COUNTER |
| Packets Sent | Shows the number of packets sent. | PERF_COUNTER_RAWCOUNT |
| Packets Sent/sec | Shows the rate, in incidents per second, at which packets were sent. | PERF_COUNTER_COUNTER |
| Server Reset Time | Shows the time elapsed, to the hundredths of a second, since the server configuration was reset. | PERF_COUNTER_RAWCOUNT |
| Server Up Time | Shows the time elapsed, to the hundredths of a second, since the server process was started. | PERF_COUNTER_RAWCOUNT |
| Unknown Type | Shows the number of packets of unknown type that were received. | PERF_COUNTER_RAWCOUNT |
| Unknown Type/sec | Shows the rate, in incidents per second, at which packets of unknown type were received. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Internet Authentication Service

# IAS Authentication Clients Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IAS Authentication Clients performance object is installed by Internet Authentication Service (IAS). IAS uses the Remote Authentication Dial-In User Service (RADIUS) protocol to perform remote authentication. IAS performance objects report activity for servers or clients, including user authentication, authorization, and accounting (usage).

When selecting counters for this object, note that each client is a separate instance.

| Counter Name | Description | Counter Type |
|---|---|---|
| Access-Accepts | Shows the number of RADIUS Access-Accept packets sent to this client. | PERF_COUNTER_RAWCOUNT |
| Access-Accepts/sec | Shows the rate, in incidents per second, at which RADIUS Access-Accept packets were sent to this client. | PERF_COUNTER_COUNTER |
| Access-Challenges | Shows the number of RADIUS Access-Challenge packets sent to this client. | PERF_COUNTER_RAWCOUNT |
| Access-Challenges/sec | Shows the rate, in incidents per second, at which RADIUS Access-Challenge packets were sent to this client. | PERF_COUNTER_COUNTER |
| Access-Rejects | Shows the number of RADIUS Access-Reject packets sent to this client. | PERF_COUNTER_RAWCOUNT |
| Access-Rejects/sec | Shows the rate, in incidents per second, at which RADIUS Access-Reject packets were sent to this client. | PERF_COUNTER_COUNTER |
| Access-Requests | Shows the number of packets received on the authentication port from this client. | PERF_COUNTER_RAWCOUNT |
| Access-Requests/sec | Shows the rate, in incidents per second, at which packets were received on the authentication port from this client. | PERF_COUNTER_COUNTER |
| Bad Authenticators | Shows the number of packets that contained invalid Message-Authenticator attributes. | PERF_COUNTER_RAWCOUNT |
| Bad Authenticators/sec | Shows the rate, in incidents per second, of packets that contained invalid Message-Authenticator attributes. | PERF_COUNTER_COUNTER |
| Dropped Packets | Shows the number of incoming packets that were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.. | PERF_COUNTER_RAWCOUNT |
| Dropped Packets/sec | Shows the rate, in incidents per second, at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_COUNTER |
| Duplicate Access-Requests | Shows the number of duplicate RADIUS Access-Request packets received from this client. | PERF_COUNTER_RAWCOUNT |
| Duplicate Access-Requests/sec | Shows the rate, in incidents per second, at which duplicate RADIUS Access-Request packets were received from this client. | PERF_COUNTER_COUNTER |
| Malformed Packets | Shows the number of malformed packets received. Bad authenticators or unknown types are not included. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Malformed Packets/sec | Shows the rate, in incidents per second, at which malformed packets were received. Bad authenticators or unknown types are not included. | PERF_COUNTER_COUNTER |
| Packets Received | Shows the number of packets received. | PERF_COUNTER_RAWCOUNT |
| Packets Received/sec | Shows the rate, in incidents per second, at which packets were received. | PERF_COUNTER_COUNTER |
| Packets Sent | Shows the number of packets sent. | PERF_COUNTER_RAWCOUNT |
| Packets Sent/sec | Shows the rate, in incidents per second, at which packets were sent. | PERF_COUNTER_COUNTER |
| Unknown Type | Shows the number of packets of unknown type that were received. | PERF_COUNTER_RAWCOUNT |
| Unknown Type/sec | Shows the rate, in incidents per second, at which packets of unknown type were received. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Internet Authentication Service

# IAS Authentication Proxy Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IAS Authentication Proxy performance object is installed by Internet Authentication Service (IAS). ISA uses the Remote Authentication Dial-In User Service (RADIUS) protocol. IAS performance objects report activity for servers or clients, including user authentication, authorization, and accounting (use).

| Counter Name | Description | Counter Type |
|---|---|---|
| Access-Accepts | Shows the number of RADIUS Access-Accept packets sent to this server. | PERF_COUNTER_RAWCOUNT |
| Access-Accepts/sec | Shows the rate, in incidents per second, at which RADIUS Access-Accept packets were sent to this server. | PERF_COUNTER_COUNTER |
| Access-Challenges | Shows the number of Access-Challenge packets sent to this server. | PERF_COUNTER_RAWCOUNT |
| Access-Challenges/sec | Shows the rate, in incidents per second, at which Access-Challenge messages were processed. | PERF_COUNTER_COUNTER |
| Access-Rejects | Shows the number of RADIUS Access-Rejects packets sent to this server. | PERF_COUNTER_RAWCOUNT |
| Access-Rejects/sec | Shows the rate, in incidents per second, at which Access-Reject messages were processed. | PERF_COUNTER_COUNTER |
| Access-Requests | Shows the number of packets that were received on the authentication port from this server. | PERF_COUNTER_RAWCOUNT |
| Access-Requests/sec | Shows the rate, in incidents per second, at which packets were received on the authentication port from this server. | PERF_COUNTER_COUNTER |
| Bad Authenticators | Shows the number of packets that contained invalid Message-Authenticator attributes. | PERF_COUNTER_RAWCOUNT |
| Bad Authenticators/sec | Shows the rate, in incidents per second, at which packets that contained invalid Message-Authenticator attributes were received. | PERF_COUNTER_COUNTER |
| Dropped Packets | Shows the number of incoming packets that were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_RAWCOUNT |
| Dropped Packets/sec | Shows the rate, in incidents per second, at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_COUNTER |
| Invalid Addresses | Shows the number of packets received from unknown addresses. | PERF_COUNTER_RAWCOUNT |
| Invalid Addresses/sec | Shows the rate, in incidents per second, at which packets were received from unknown addresses. | PERF_COUNTER_COUNTER |
| Malformed Packets | Shows the number of malformed packets received. Bad authenticators or unknown types are not included. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Malformed Packets/sec | Shows the rate, in incidents per second, at which malformed packets were received. Bad authenticators or unknown types are not included. | PERF_COUNTER_COUNTER |
| Packets Received | Shows the number of packets received. | PERF_COUNTER_RAWCOUNT |
| Packets Received/sec | Shows the rate, in incidents per second, at which packets were received. | PERF_COUNTER_COUNTER |
| Pending Requests | Shows the number of requests destined for this server that have not yet timed out or received a response. | PERF_COUNTER_RAWCOUNT |
| Request Timeouts | Shows the number of requests to this server that have timed out. | PERF_COUNTER_RAWCOUNT |
| Request Timeouts/sec | Shows the rate, in incidents per second, at which requests to this server timed out. | PERF_COUNTER_COUNTER |
| Retransmissions | Shows the number of requests that were retransmitted to this server. | PERF_COUNTER_RAWCOUNT |
| Retransmissions/sec | Shows the rate, in incidents per second, at which requests to this server were retransmitted. | PERF_COUNTER_COUNTER |
| Unknown Type | Shows the number of packets of unknown type that were received. | PERF_COUNTER_RAWCOUNT |
| Unknown Type/sec | Shows the rate, in incidents per second, at which packets of unknown type were received. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Internet Authentication Service

# IAS Authentication Server Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IAS Authentication Server performance object is installed by Internet Authentication Service (IAS). IAS uses the Remote Authentication Dial-In User Service (RADIUS) protocol to perform remote authentication. IAS performance objects report activity for servers or clients, including user authentication, authorization, and accounting (usage).

When selecting counters for this object, note that each client is a separate instance.

| Counter Name | Description | Counter Type |
|---|---|---|
| Access-Accepts | Shows the number of RADIUS Access-Accept packets sent to this server. | PERF_COUNTER_RAWCOUNT |
| Access-Accepts/sec | Shows the rate, in incidents per second, at which RADIUS Access-Accept packets were sent to this server. | PERF_COUNTER_COUNTER |
| Access-Challenges | Shows the number of Access-Challenge packets sent to this server. | PERF_COUNTER_RAWCOUNT |
| Access-Challenges/sec | Shows the rate, in incidents per second, at which Access-Challenge messages are being processed. | PERF_COUNTER_COUNTER |
| Access-Rejects | Shows the number of RADIUS Access-Reject packets sent to this server. | PERF_COUNTER_RAWCOUNT |
| Access-Rejects/sec | Shows the rate, in incidents per second, at which Access-Reject messages are being processed. | PERF_COUNTER_COUNTER |
| Access-Requests | Shows the number of packets that were received on the authentication port from this server. | PERF_COUNTER_RAWCOUNT |
| Access-Requests/sec. | Shows the rate, in incidents per second, at which packets were received on an authentication port from this server. | PERF_COUNTER_COUNTER |
| Bad Authenticators | Shows the total number of packets that contained invalid Message-Authenticator attributes. | PERF_COUNTER_RAWCOUNT |
| Bad Authenticators/sec | Shows the rate in, incidents per second, at which packets that contained invalid Message-Authenticator attributes were received. | PERF_COUNTER_COUNTER |
| Dropped Packets | Shows the number of incoming packets dropped for reasons other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_RAWCOUNT |
| Dropped Packets/sec | Shows the rate, in incidents per second, at which packets were dropped for reasons other than being malformed, bad authenticators, or of an unknown type. | PERF_COUNTER_COUNTER |
| Duplicate Access-Requests | Shows the number of duplicate RADIUS Access-Request packets received from this server. | PERF_COUNTER_RAWCOUNT |
| Duplicate Access-Requests/sec | Shows the rate, in incidents per second, at which duplicate RADIUS Access-Request packets were received from this server. | PERF_COUNTER_COUNTER |
| Invalid Requests | Shows the number of packets received from unknown addresses. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Invalid Requests/sec | Shows the rate, in incidents per second, at which packets were received from unknown addresses. | PERF_COUNTER_COUNTER |
| Malformed Packets | Shows the number of malformed packets received. Bad authenticators or unknown types are not included. | PERF_COUNTER_RAWCOUNT |
| Malformed Packets/sec | Shows the rate, in incidents per second, at which malformed packets were received. Bad authenticators or unknown types are not included. | PERF_COUNTER_COUNTER |
| Packets Received | Shows the number of packets received. | PERF_COUNTER_RAWCOUNT |
| Packets Received/sec | Shows the rate, in incidents per second, at which packets were received. | PERF_COUNTER_COUNTER |
| Packets Sent | Shows the number of packets sent. | PERF_COUNTER_RAWCOUNT |
| Packets Sent/sec | Shows the rate, in incidents per second, at which packets were sent. | PERF_COUNTER_COUNTER |
| Server Reset Time | Shows the time elapsed, to the hundredths of a second, since the server configuration was reset. | PERF_COUNTER_RAWCOUNT |
| Server Up Time | Shows the time elapsed, to the hundredths of a second, since the server process was started. | PERF_COUNTER_RAWCOUNT |
| Unknown Type | Shows the number of packets received that were of an unknown type. | PERF_COUNTER_RAWCOUNT |
| Unknown Type/sec | Shows the rate, in incidents per second, at which packets of an unknown type were received. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Internet Authentication Service

# IAS Remote Accounting Servers Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS Remote Accounting Servers Object**

The IAS Remote Accounting Servers performance object is installed by Internet Authentication Service (IAS). IAS uses the Remote Authentication Dial-In User Service (RADIUS) protocol. IAS performance objects report activity for servers or clients, including user authentication, authorization, and accounting (use).

| Counter Name | Description | Counter Type |
|---|---|---|
| Accounting-Requests | Shows the number of packets received from this client on the accounting port. | PERF_COUNTER_RAWCOUNT |
| Accounting-Requests/sec | Shows the rate, in seconds, at which packets were received from this client on the accounting port. | PERF_COUNTER_COUNTER |
| Accounting-Responses | Shows the number of RADIUS Accounting-Response packets sent to this client. | PERF_COUNTER_RAWCOUNT |
| Accounting-Responses/sec | Shows the rate, in seconds, at which RADIUS Accounting-Response packets were sent to this client. | PERF_COUNTER_COUNTER |
| Bad Authenticators | Shows the number of packets that contained invalid Message-Authenticator attributes. | PERF_COUNTER_RAWCOUNT |
| Bad Authenticators/sec | Shows the rate, in seconds, at which packets that contained invalid Message-Authenticator attributes were received. | PERF_COUNTER_COUNTER |
| Dropped Packets | Shows the number of incoming packets that were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_RAWCOUNT |
| Dropped Packets/sec | Shows the rate, in seconds, at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_COUNTER |
| Duplicate Accounting-Requests | Shows the number of duplicate RADIUS Accounting-Request packets that were received from this client. | PERF_COUNTER_RAWCOUNT |
| Duplicate Accounting-Requests/sec | Shows the rate, in seconds, at which duplicate RADIUS Accounting-Request packets were received from this client. | PERF_COUNTER_COUNTER |
| Invalid Requests | Shows the number of packets received from unknown addresses. | PERF_COUNTER_RAWCOUNT |
| Invalid Requests/sec | Shows the rate, in seconds, at which packets were received from unknown addresses. | PERF_COUNTER_COUNTER |
| Malformed Packets | Shows the number of malformed packets received. Bad authenticators or unknown types are not included. | PERF_COUNTER_RAWCOUNT |
| Malformed Packets/sec | Shows the rate, in seconds, at which malformed packets were received. Bad authenticators or unknown types are not included. | PERF_COUNTER_COUNTER |
| No Record | Shows the number of RADIUS Accounting-Request packets that were received and responded to but not recorded. | PERF_COUNTER_RAWCOUNT |
| No Record/sec | Shows the rate in seconds at which RADIUS Accounting-Request packets were received and responded to but not recorded. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Packets Received | Shows the number of packets received. | PERF_COUNTER_RAWCOUNT |
| Packets Received/sec | Shows the rate, in seconds, at which packets were received. | PERF_COUNTER_COUNTER |
| Packets Sent | Shows the number of packets sent. | PERF_COUNTER_RAWCOUNT |
| Packets Sent/sec | Shows the rate, in seconds, at which packets were sent. | PERF_COUNTER_COUNTER |
| Server Reset Time | Shows the time elapsed, in hundredths of a second, since the server configuration was reset. | PERF_COUNTER_RAWCOUNT |
| Server Up Time | Shows the time elapsed, in hundredths of a second, since the server process was started. | PERF_COUNTER_COUNTER |
| Unknown Type | Shows the number of packets of unknown type that were received. | PERF_COUNTER_RAWCOUNT |
| Unknown Type/sec | Shows the rate, in seconds, at which packets of unknown type were received. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
IAS Accounting Server Object
Internet Authentication Service

# IAS Remote Authentication Servers Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IAS Remote Authentication Servers performance object is installed by Internet Authentication Service (IAS). IAS uses the Remote Authentication Dial-In User Service (RADIUS) protocol to perform remote administration. IAS performance objects report activity for servers or clients, including user authentication, authorization, and accounting (usage).

| Counter Name | Description | Counter Type |
|---|---|---|
| Access-Accepts | Shows the number of RADIUS Access-Accept packets sent to this server. | PERF_COUNTER_RAWCOUNT |
| Access-Accepts/sec | Shows the rate, in incidents per second, at which RADIUS Access-Accept packets were sent to this server. | PERF_COUNTER_COUNTER |
| Access-Challenges | Shows the number of Access-Challenge packets sent to this server. | PERF_COUNTER_RAWCOUNT |
| Access-Challenges/sec | Shows the rate, in incidents per second, at which Access-Challenge messages were processed. | PERF_COUNTER_COUNTER |
| Access-Rejects | Shows the number of RADIUS Access-Rejects packets sent to this server. | PERF_COUNTER_RAWCOUNT |
| Access-Rejects/sec | Shows the rate, in incidents per second, at which Access-Reject messages were processed. | PERF_COUNTER_COUNTER |
| Access-Requests | Shows the number of packets that were received on the authentication port from this server. | PERF_COUNTER_RAWCOUNT |
| Access-Requests/sec | Shows the rate, in incidents per second, at which packets were received on the authentication port from this server. | PERF_COUNTER_COUNTER |
| Bad Authenticators | Shows the total number of packets that contained invalid Message-Authenticator attributes. | PERF_COUNTER_RAWCOUNT |
| Bad Authenticators/sec | Shows the rate, in incidents per second, at which packets that contained invalid Message-Authenticator attributes were received. | PERF_COUNTER_COUNTER |
| Dropped Packets | Shows the number of incoming packets that were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_RAWCOUNT |
| Dropped Packets/sec | Shows the rate, in incidents per second, at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types. | PERF_COUNTER_COUNTER |
| Last Round-Trip Time | Shows the time interval, to the hundredths of a second, between the most recent request and the response that matched it from this server. | PERF_COUNTER_RAWCOUNT |
| Malformed Packets | Shows the number of malformed packets received. Bad authenticators or unknown types are not included. | PERF_COUNTER_RAWCOUNT |
| Malformed Packets/sec | Shows the rate, in incidents per second, at which malformed packets were received. Bad authenticators or unknown types are not included. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Packets Received | Shows the number of packets received. | PERF_COUNTER_RAWCOUNT |
| Packets Received/sec | Shows the rate, in incidents per second, at which packets were received. | PERF_COUNTER_COUNTER |
| Pending Requests | Shows the number of requests destined for this server that have not yet timed out or received a response. | PERF_COUNTER_RAWCOUNT |
| Port Number | Shows UDP port the proxy is using to send requests to this server. | PERF_COUNTER_RAWCOUNT |
| Request Timeouts | Shows the number of requests to this server that timed out. | PERF_COUNTER_RAWCOUNT |
| Request Timeouts/sec | Shows the rate, in incidents per second, at which requests to this server timed out. | PERF_COUNTER_COUNTER |
| Retransmissions | Shows the number of requests retransmitted to this server. | PERF_COUNTER_RAWCOUNT |
| Retransmissions/sec | Shows the rate, in incidents per second, at which requests to this server were retransmitted. | PERF_COUNTER_COUNTER |
| Unknown Type | Shows the number of packets of unknown type that were received. | PERF_COUNTER_RAWCOUNT |
| Unknown Type/sec | Shows the rate, in incidents per second, at which packets of unknown type were received. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
IAS Authentication Server Object
Internet Authentication Service

# ICMP Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The ICMP performance object consists of counters that measure the rates at which Internet Control Message Protocol (ICMP) messages are sent and received by using ICMP. It also includes counters that monitor ICMP errors.

| Counter Name | Description | Counter Type |
|---|---|---|
| Messages Outbound Errors | Shows the number of ICMP messages that were not sent due to problems within ICMP, such as lack of buffers. This value does not include errors that were discovered outside the ICMP layer, such as those recording the failure of IP to route the resultant datagram. In some implementations, none of the error types are included in the value of this counter. | PERF_COUNTER_RAWCOUNT |
| Messages Received Errors | Shows the number of ICMP messages that were received but had errors, such as bad ICMP checksums, bad length, etc. | PERF_COUNTER_RAWCOUNT |
| Messages Received/sec | Shows the rate, in incidents per second, at which ICMP messages were received. The rate includes messages received in error. | PERF_COUNTER_COUNTER |
| Messages Sent/sec | Shows the rate, in incidents per second, at which the server attempted to send messages. The rate includes messages sent in error. | PERF_COUNTER_COUNTER |
| Messages/sec | Shows the total rate, in incidents per second, at which ICMP messages were sent and received by the entity. The rate includes messages received or sent in error. | PERF_COUNTER_COUNTER |
| Received Address Mask | Shows the number of ICMP Address Mask Request messages that were received. | PERF_COUNTER_RAWCOUNT |
| Received Address Mask Reply | Shows the number of ICMP Address Mask Reply messages that were received. | PERF_COUNTER_RAWCOUNT |
| Received Dest. Unreachable | Shows the number of ICMP Destination Unreachable messages that were received. | PERF_COUNTER_RAWCOUNT |
| Received Echo Reply/sec | Shows the rate, in incidents per second, at which ICMP Echo Reply messages were received. | PERF_COUNTER_COUNTER |
| Received Echo/sec | Shows the rate, in incidents per second, at which ICMP Echo messages were received. | PERF_COUNTER_COUNTER |
| Received Parameter Problem | Shows the number of ICMP Parameter Problem messages that were received. | PERF_COUNTER_RAWCOUNT |
| Received Redirect/sec | Shows the rate, in incidents per second, at which ICMP Redirect messages were received. | PERF_COUNTER_COUNTER |
| Received Source Quench | Shows the number of ICMP Source Quench messages that were received. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Received Time Exceeded | Shows the number of ICMP Time Exceeded messages that were received. | PERF_COUNTER_RAWCOUNT |
| Received Timestamp Reply/sec | Shows the rate, in incidents per second, at which ICMP Timestamp Reply messages were received. | PERF_COUNTER_COUNTER |
| Received Timestamp/sec | Shows the rate, in incidents per second, at which ICMP Timestamp Request messages were received. | PERF_COUNTER_COUNTER |
| Sent Address Mask | Shows the number of ICMP Address Mask Request messages that were sent. | PERF_COUNTER_RAWCOUNT |
| Sent Address Mask Reply | Shows the number of ICMP Address Mask Reply messages that were sent. | PERF_COUNTER_RAWCOUNT |
| Sent Destination Unreachable | Shows the number of ICMP Destination Unreachable messages that were sent. | PERF_COUNTER_RAWCOUNT |
| Sent Echo Reply/sec | Shows the rate, in incidents per second, at which ICMP Echo Reply messages were sent. | PERF_COUNTER_COUNTER |
| Sent Echo/sec | Shows the rate, in incidents per second, at which ICMP Echo messages were sent. | PERF_COUNTER_COUNTER |
| Sent Parameter Problem | Shows the number of ICMP Parameter Problem messages that were sent. | PERF_COUNTER_RAWCOUNT |
| Sent Redirect/sec | Shows the rate, in incidents per second, at which ICMP Redirect messages were sent. | PERF_COUNTER_COUNTER |
| Sent Source Quench | Shows the number of ICMP Source Quench messages that were sent. | PERF_COUNTER_RAWCOUNT |
| Sent Time Exceeded | Shows the number of ICMP Time Exceeded messages that were sent. | PERF_COUNTER_RAWCOUNT |
| Sent Timestamp Reply/sec | Shows the rate, in incidents per second, at which ICMP Timestamp Reply messages were sent. | PERF_COUNTER_COUNTER |
| Sent Timestamp/sec | Shows the rate, in incidents per second, at which ICMP Timestamp Request messages were sent. | PERF_COUNTER_COUNTER |

# Image Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Image performance object consists of counters that monitor virtual address usage of images executed by processes on the computer. This object contains "costly" counters for which data collection is typically either processor-intensive or memory-intensive. Costly counters cannot be selected with System Monitor but can be accessed by using Showperf.exe, a tool included in the Windows Resource Kit Tools Help for Windows Server 2003.

| Counter Name | Description | Counter Type |
|---|---|---|
| Exec Read Only | Shows memory that can be executed as well as read. Image Space is the virtual address space used by the selected image with this protection. | PERF_COUNTER_RAWCOUNT |
| Exec Read/Write | Shows memory that can be executed by programs, in addition to being read and written. | PERF_COUNTER_RAWCOUNT |
| Exec Write Copy | Shows memory that can be executed by programs, in addition to being read and written. This type of protection is used when memory needs to be shared between processes. If the sharing processes only read the memory, they all use the same memory. If a sharing process wants to have write access to this shared memory, a copy of this memory is made for that process. | PERF_COUNTER_RAWCOUNT |
| Executable | Shows memory that can be executed by programs, but cannot be read or written. This type of protection is not supported by all processor types. | PERF_COUNTER_RAWCOUNT |
| No Access | Prevents a process from writing or reading these pages, and generates an access violation if either is attempted. | PERF_COUNTER_RAWCOUNT |
| Read Only | Prevents the contents of these pages from being modified. Any attempt to write or modify these pages generates an access violation | PERF_COUNTER_RAWCOUNT |
| Read/Write | Allows a process to read, modify and write to these pages. | PERF_COUNTER_RAWCOUNT |
| Write Copy | Used when memory is shared for reading but not for writing. If the sharing processes only read this memory, they all use the same memory. If a sharing process wants read/write access to this shared memory, a copy of this memory is made for that process. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Process Address Space Object
Job Object Object
Job Object Details
Thread Details Object

# Indexing Service Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Indexing Service performance object reports statistics about the creation of indexes, and about the merging and querying of indexes by the Indexing Service.

| Counter Name | Description | Counter Type |
|---|---|---|
| # Documents Indexed | Shows the number of documents indexed since the current indexing session started. This does not include the documents indexed in prior sessions of the Indexing Service. | PERF_COUNTER_RAWCOUNT |
| Deferred for Indexing | Shows the number of documents that need to be indexed but cannot because they were in use. | PERF_COUNTER_RAWCOUNT |
| Files to be Indexed | Shows the minimum number of documents that were known to need indexing. | PERF_COUNTER_RAWCOUNT |
| Index Size (MB) | Shows the size, in megabytes (MB), of the content index. | PERF_COUNTER_RAWCOUNT |
| Merge Progress | Shows the percentage of the current merge that has been completed. | PERF_COUNTER_RAWCOUNT |
| Running Queries | Shows the number of queries that were running. | PERF_COUNTER_RAWCOUNT |
| Saved Indexes | Shows the number of indexes that were saved. | PERF_COUNTER_RAWCOUNT |
| Total # Documents | Shows the total number of documents known to the index. | PERF_COUNTER_RAWCOUNT |
| Total # of Queries | Shows the total number of queries made since the index was mounted. | PERF_COUNTER_RAWCOUNT |
| Unique Keys | Shows the number of unique words in the index. | PERF_COUNTER_RAWCOUNT |
| Word Lists | Shows the number of word lists in the index. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Counters by Feature

# Indexing Service Filter Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Indexing Service Filter object monitors the filtering activity of the Indexing Service.

| Counter Name | Description | Counter Type |
|---|---|---|
| Binding Time (msec) | Shows the average time, in milliseconds (msec), for the Indexing Service to bind to a filter file. | PERF_COUNTER_RAWCOUNT |
| Indexing Speed (MB/hr) | Shows the speed, in megabytes per hour (MB/hr), at which documents were indexed. This does not include generating abstracts. | PERF_COUNTER_RAWCOUNT |
| Total Indexing Speed (MB/hr) | Shows the speed, in megabytes per hour (MB/hr), at which documents were indexed. This includes generating abstracts. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Counters by Feature

# Internet Information Services Global Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Internet Information Services Global performance object consists of counters that apply jointly to the Hypertext Transport Protocol (HTTP) service and File Transfer Protocol (FTP) service.

| Counter Name | Description | Counter Type |
|---|---|---|
| Active Flushed Entries | Shows the number of cached file handles that will be closed when all current transfers are completed. | PERF_COUNTER_RAWCOUNT |
| BLOB Cache Flushes | Shows the number of flushes of the binary large object (BLOB) cache since the service was started. | PERF_COUNTER_RAWCOUNT |
| BLOB Cache Hits | Shows the total number of successful lookups in the BLOB cache. | PERF_COUNTER_RAWCOUNT |
| BLOB Cache Hits % | Shows the ratio of BLOB cache hits to total cache requests. | PERF_RAW_FRACTION |
| BLOB Cache Misses | Shows the total number of unsuccessful lookups in the BLOB cache. | PERF_COUNTER_RAWCOUNT |
| Current BLOBs Cached | Shows the number of BLOB information blocks currently in the cache for HTTP and FTP services. | PERF_COUNTER_RAWCOUNT |
| Current Blocked Async I/O Requests | Shows the number of current requests temporarily blocked due to bandwidth throttling settings. Blocked requests are held in a buffer and then unblocked if more bandwidth becomes available (unless a timeout limit is reached). | PERF_COUNTER_RAWCOUNT |
| Current File Handles Cached | Shows the number of open file handles currently in the cache for HTTP and FTP services. | PERF_COUNTER_RAWCOUNT |
| Current URIs Cached | Shows Uniform Resource Identifier (URI) information blocks currently in the cache for HTTP and FTP services. | PERF_COUNTER_RAWCOUNT |
| File Handle Cache Flushes | Shows the number of file handle cache flushes since the service was started. | PERF_COUNTER_RAWCOUNT |
| File Handle Cache Hits | Shows the total number of successful lookups in the file handle cache. | PERF_COUNTER_RAWCOUNT |
| File Handle Cache Hits % | Shows the ratio of file handle cache hits to total cache requests. | PERF_RAW_FRACTION |
| File Handle Cache Misses | Shows the total number of unsuccessful lookups in the file handle cache. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Measured Async I/O Bandwidth Usage | Shows the amount of measured bandwidth of asynchronous input/output (I/O) averaged over a minute. This is a measure of the total amount of user traffic on the server. | PERF_COUNTER_RAWCOUNT |
| Total Allowed Async I/O Requests | Shows the total number of requests allowed by bandwidth throttling settings since the service was started. When you throttle bandwidth, the number of allowed user requests will be restricted. | PERF_COUNTER_RAWCOUNT |
| Total BLOBs Cached | Shows the total number of BLOB information blocks added to the cache for HTTP and FTP services | PERF_COUNTER_RAWCOUNT |
| Total Blocked Async I/O Requests | Shows the total number of requests temporarily blocked due to bandwidth throttling settings since the service was started. Blocked requests are held in a buffer and then unblocked if more bandwidth becomes available (unless a timeout limit is reached). | PERF_COUNTER_RAWCOUNT |
| Total File Handles Cached | Shows the total number of file handles added to the cache for HTTP and FTP services | PERF_COUNTER_RAWCOUNT |
| Total Flushed BLOBs | Shows the number of BLOB information blocks that have been removed from the cache since the service was started. | PERF_COUNTER_RAWCOUNT |
| Total Flushed Files | Shows the number of file handles that have been removed from the cache since the service was started. | PERF_COUNTER_RAWCOUNT |
| Total Flushed URIs | Shows the number of URI information blocks that have been removed from the cache since the service was started. | PERF_COUNTER_RAWCOUNT |
| Total Rejected Async I/O Requests | Shows the total requests rejected due to bandwidth throttling settings since the service was started. Unlike a blocked request, a rejected request is not held in a buffer. | PERF_COUNTER_RAWCOUNT |
| Total URIs Cached | Shows the total number of URI information blocks added to the cache for HTTP and FTP services. | PERF_COUNTER_RAWCOUNT |
| URI Cache Flushes | Shows the total number of URI cache flushes since the service was started. | PERF_COUNTER_RAWCOUNT |
| URI Cache Hits | Shows the total number of successful lookups in the URI cache. | PERF_COUNTER_RAWCOUNT |
| URI Cache Hits % | Shows the ratio of URI cache hits to total cache requests. | PERF_RAW_FRACTION |
| URI Cache Misses | Shows the total number of unsuccessful lookups in the URI cache. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Counters by Feature

# IP Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The IP performance object includes counters that describe the rates at which Internet Protocol (IP) datagrams are transmitted by a computer using the IP protocol. It also describes various IP error counts.

| Counter Name | Description | Counter Type |
|---|---|---|
| Datagrams Forwarded/sec | Shows the rate, in incidents per second, at which attempts were made to find routes to forward input datagrams to their final destination, because the local server was not the final IP destination. In servers that do not act as IP gateways, this rate includes only packets that were source-routed via this entity where the source-route option processing was successful. | PERF_COUNTER_COUNTER |
| Datagrams Outbound No Route | Shows the number of IP datagrams that were discarded because no route could be found to transmit them to their destination. This counter includes any packets counted in Datagrams Forwarded/sec that meet this "no route" criterion. | PERF_COUNTER_RAWCOUNT |
| Datagrams Outbound Discarded | Shows the number of output IP datagrams that were discarded even though no problems were encountered to prevent their transmission to their destination (for example, lack of buffer space). This counter includes datagrams counted in Datagrams Forwarded/sec that meet this criterion. | PERF_COUNTER_RAWCOUNT |
| Datagrams Received Address Errors | Shows the number of input IP datagrams that were discarded because the IP address in their IP header destination field was not valid for the computer. This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For entities that are not IP gateways and do not forward datagrams, this counter includes datagrams that were discarded because the destination was not a local address. | PERF_COUNTER_RAWCOUNT |
| Datagrams Received Delivered/sec | Shows the rate, in incidents per second, at which input IP datagrams were successfully delivered to IP user-protocols, including Internet Control Message Protocol (ICMP). | PERF_COUNTER_COUNTER |
| Datagrams Received Discarded | Shows the number of input IP datagrams that were discarded even though problems prevented their continued processing (for example, lack of buffer space). This counter does not include any datagrams discarded while awaiting reassembly. | PERF_COUNTER_RAWCOUNT |
| Datagrams Received Header Errors | Shows the number of input IP datagrams that were discarded due to errors in the IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered while processing their IP options, etc. | PERF_COUNTER_RAWCOUNT |
| Datagrams Received Unknown Protocol | Shows the number of locally-addressed datagrams that were successfully received but were discarded because of an unknown or unsupported protocol. | PERF_COUNTER_RAWCOUNT |
| Datagrams Received/sec | Shows the rate, in incidents per second, at which IP datagrams were received from the interfaces, including those in error. Datagrams Received/sec is a subset of Datagrams/sec. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Datagrams Sent/sec | Shows the rate, in incidents per second, at which IP datagrams were supplied for transmission by local IP user-protocols (including ICMP). This counter does not include datagrams counted in Datagrams Forwarded/sec. Datagrams Sent/sec is a subset of Datagrams/sec. | PERF_COUNTER_COUNTER |
| Datagrams/sec | Shows the rate, in incidents per second, at which IP datagrams were received from or sent to the interfaces, including those in error. Forwarded datagrams are not included in this rate. | PERF_COUNTER_COUNTER |
| Fragment Reassembly Failures | Shows the number of failures detected by the IP reassembly algorithm, such as time outs, errors, etc. This is not necessarily a count of discarded IP fragments since some algorithms (notably RFC 815) lose track of the number of fragments by combining them as they are received. | PERF_COUNTER_RAWCOUNT |
| Fragmentation Failures | Shows the number of IP datagrams that were discarded because they needed to be fragmented but could not be (for example, because the "Don't Fragment" flag was set). | PERF_COUNTER_RAWCOUNT |
| Fragmented Datagrams/sec | Shows the rate, in incidents per second, at which datagrams were successfully fragmented. | PERF_COUNTER_COUNTER |
| Fragments Created/sec | Shows the rate, in incidents per second, at which IP datagram fragments were generated as a result of fragmentation. | PERF_COUNTER_COUNTER |
| Fragments Reassembled/sec | Shows the rate, in incidents per second, at which IP fragments were successfully reassembled. | PERF_COUNTER_COUNTER |
| Fragments Received/sec | Shows the rate, in incidents per second, at which IP fragments that need to be reassembled at this entity were received. | PERF_COUNTER_COUNTER |

# Job Object Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Job Object performance object reports the accounting and processor usage data collected by each active job object. A job object provides a means of managing groups of processes with respect to their processor usage and other factors.

| Counter Name | Description | Counter Type |
|---|---|---|
| Current % Kernel Mode Time | Shows the percentage of time that the processes in the job object spent executing code in kernel or privileged mode, during the sample interval. | PERF_100NSEC_TIMER |
| Current % Processor Time | Shows the percentage of time that the processes in the job object spent executing code during the sample interval. | PERF_100NSEC_TIMER |
| Current % User Mode Time | Shows the percentage of time that the processes in the job object spent executing code in user mode during the sample interval. | PERF_100NSEC_TIMER |
| Pages/sec | Shows the page fault rate, in incidents per second, of all the processes in the job object during the sample interval. | PERF_COUNTER_COUNTER |
| Process Count - Active | Shows the number of processes that were associated with the job object during the sample interval. | PERF_COUNTER_RAWCOUNT |
| Process Count - Terminated | Shows the number of processes that were terminated because of a limit violation during the sample interval. | PERF_COUNTER_RAWCOUNT |
| Process Count - Total | Shows the total number of processes, both active and terminated, that were associated with the job object during the sample interval. | PERF_COUNTER_RAWCOUNT |
| This Period mSec - Kernel Mode | Shows the amount of user mode processor time, in milliseconds, that was used by all the processes in the job object during the sample interval. This includes processes that have terminated, or that were no longer associated with the job object because of a time limit on the job. | PERF_COUNTER_LARGE_RAWCOUNT |
| This Period mSec - Processor | Shows the amount of processor time, in milliseconds, that was used by all the processes in the job object during the sample interval. This includes processes that have terminated, or that were no longer associated with the job object because of a time limit on the job. | PERF_COUNTER_LARGE_RAWCOUNT |
| This Period mSec - User Mode | Shows the amount of user mode processor time, in milliseconds, that was used by all the processes in the job object during the sample interval. This includes processes that have terminated, or that were no longer associated with the job object because of a time limit on the job. | PERF_COUNTER_LARGE_RAWCOUNT |

| | | |
|---|---|---|
| Total mSec - Kernel Mode | Shows the total amount of kernel mode processor time, in milliseconds, that was used by all the processes in the job object since it was created. This includes processes that have terminated, or that are no longer associated with the job object. | PERF_COUNTER_LARGE_RAWCOUNT |
| Total mSec - Processor | Shows the total amount of processor time, in milliseconds, that was used by all the processes in the job object since it was created. This includes processes that have terminated, or that are no longer associated with the job object. | PERF_COUNTER_LARGE_RAWCOUNT |
| Total mSec - User Mode | Shows the total amount of user mode processor time, in milliseconds, that was used by all the processes in the job object since it was created. This includes processes that have terminated, or that are no longer associated with the job object. | PERF_COUNTER_LARGE_RAWCOUNT |

**See Also**
**Concepts**
Job Object Details

# Job Object Details

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Job Object Details performance object counters show detailed performance information about the active processes that make up a job object.

| Counter Name | Description | Counter Type |
|---|---|---|
| % Privileged Time | Shows the percentage of elapsed time that the process threads have spent executing code in privileged mode. When a Windows 2000 system service is called, the service will often run in privileged mode to gain access to system-private data. Such data is protected from access by threads executing in user mode. Calls to the system can be explicit or implicit, such as page faults or interrupts. Unlike some early operating systems, Windows 2000 uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. Some work done by Windows 2000 on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process. | PERF_100NSEC_TIMER |
| % Processor Time | Shows the percentage of elapsed time that all of the process threads used the processor to execute instructions. An instruction is the basic unit of execution in a computer, a thread is the object that executes instructions, and a process is the object created when a program is run. Code executed to handle some hardware interrupts and trap conditions are included in this count. | PERF_100NSEC_TIMER |
| % User Time | Shows the percentage of elapsed time that the process threads spent executing code in user mode. Applications, environment subsystems, and integral subsystems execute in user mode. Code executing in user mode cannot damage the integrity of the Windows 2000 executive, kernel, and device drivers. Unlike some early operating systems, Windows 2000 uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. Some work done by Windows 2000 on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process. | PERF_100NSEC_TIMER |
| Creating Process ID | Shows the Process ID of the process that created the process. The creating process may have terminated, so this value may no longer identify a running process. | PERF_COUNTER_LARGE_RAWCOUNT |
| Elapsed Time | Shows the time, in seconds, that this process has been running. | PERF_ELAPSED_TIME |
| Handle Count | Shows the total number of handles currently open by this process. This is the sum of the handles currently open by each thread in this process. | PERF_COUNTER_RAWCOUNT |
| ID Process | Shows the unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process. | PERF_COUNTER_LARGE_RAWCOUNT |
| IO Data Bytes/sec | Shows the rate that the process is reading and writing bytes in I/O operations. It counts all I/O activity generated by the process to include file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |

| | | |
|---|---|---|
| IO Data Operations/sec | Shows the rate at which the process is issuing read and write I/O operations. This counter counts all I/O activity generated by the process to include file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Other Bytes/sec | Shows the rate at which the process is issuing bytes to I/O operations that don't involve data such as control operations. It counts all I/O activity generated by the process to include file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Other Operations/sec | Shows the rate at which the process is issuing I/O operations that are neither read nor write operations (for example, a control function). It counts all I/O activity generated by the process to include file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Read Bytes/sec | Shows the rate at which the process is reading bytes from I/O operations. It counts all I/O activity generated by the process to include file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Read Operations/sec | Shows the rate at which the process is issuing read I/O operations. It counts all I/O activity generated by the process to include file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Write Bytes/sec | Shows the rate at which the process is writing bytes to I/O operations. It counts all I/O activity generated by the process to include file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Write Operations/sec | Shows the rate at which the process is issuing write I/O operations. It counts all I/O activity generated by the process to include file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| Page Faults/sec | Shows the rate of page faults by the threads executing in this process. A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This will not cause the page to be fetched from disk if it is on the standby list and already in main memory, or if it is in use by another process that is sharing the page. | PERF_COUNTER_COUNTER |
| Page File Bytes | Shows the current number of bytes this process has used in the paging file(s). Paging files are used to store pages of memory used by the process that are not contained in other files. Paging files are shared by all processes; lack of space in paging files can prevent other processes from allocating memory. | PERF_COUNTER_LARGE_RAWCOUNT |
| Page File Bytes Peak | Shows the maximum number of bytes this process has used in the paging file(s). | PERF_COUNTER_LARGE_RAWCOUNT |
| Pool Nonpaged Bytes | Shows the size, in bytes, in the nonpaged pool, an area of system memory (physical memory used by the operating system) for objects that cannot be written to disk, but must remain in physical memory as long as they are allocated. Memory\Pool Nonpaged Bytes is calculated differently than Process\Pool Nonpaged Bytes, and might not equal Process(_Total)\Pool Nonpaged Bytes This counter displays the last observed value only; it is not an average. | PERF_COUNTER_RAWCOUNT |
| Pool Paged Bytes | Shows the size, in bytes, in the paged pool. Memory\Pool Paged Bytes is calculated differently than Process\Pool Paged Bytes, so it might not equal Process(_Total)\Pool Paged Bytes. This counter displays the last observed value only; it is not an average. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Priority Base | Shows the current base priority of this process. Threads within a process can raise and lower their own base priority relative to the base priority of the priority. | PERF_COUNTER_RAWCOUNT |
| Private Bytes | Shows the current size, in bytes, that this process has allocated that cannot be shared with other processes. | PERF_COUNTER_LARGE_RAWCOUNT |
| Thread Count | Shows the number of threads currently active in this process. An instruction is the basic unit of execution in a processor, and a thread is the object that executes instructions. Every running process has at least one thread. | PERF_COUNTER_RAWCOUNT |
| Virtual Bytes | Shows the current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and by using too much, the process can limit its ability to load libraries. | PERF_COUNTER_LARGE_RAWCOUNT |
| Virtual Bytes Peak | Shows the maximum number of bytes of virtual address space the process has used at any one time. | PERF_COUNTER_LARGE_RAWCOUNT |
| Working Set | Shows the current size, in bytes, in the Working Set of this process. The Working Set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the Working Set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from Working Sets. If they are needed, they will be soft-faulted back into the Working Set before leaving main memory. | PERF_COUNTER_RAWCOUNT |
| Working Set Peak | Shows the maximum size, in bytes, in the Working Set of this process at any point in time. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Job Object Object

# LogicalDisk Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The LogicalDisk performance object consists of counters that monitor logical partitions of hard or fixed disk drives. System Monitor identifies logical disks by their drive letter, such as "C."

| Counter Name | Description | Counter Type |
|---|---|---|
| % Disk Time | Shows the percentage of time that the selected disk drive was busy servicing read or write requests. | PERF_PRECISION_100NS_TIMER |
| % Disk Read Time | Shows the percentage of time that the selected disk drive was busy servicing read requests. | PERF_PRECISION_100NS_TIMER |
| % Disk Write Time | Shows the percentage of time that the selected disk drive was busy servicing write requests. | PERF_PRECISION_100NS_TIMER |
| % Free Space | Shows the percentage of the total usable space on the selected logical disk drive that was free. | PERF_RAW_FRACTION |
| Avg. Disk Bytes/Read | Shows the average number of bytes transferred from the disk during read operations. | PERF_AVERAGE_BULK |
| Avg. Disk Bytes/Transfer | Shows the average number of bytes transferred to or from the disk during write or read operations. | PERF_AVERAGE_BULK |
| Avg. Disk Bytes/Write | Shows the average number of bytes transferred to the disk during write operations. | PERF_AVERAGE_BULK |
| Avg. Disk Queue Length | Shows the average number of both read and write requests that were queued for the selected disk during the sample interval. | PERF_COUNTER_LARGE_QUEUELEN_TYPE |
| Avg. Disk Read Queue Length | Shows the average number of read requests that were queued for the selected disk during the sample interval. | PERF_COUNTER_LARGE_QUEUELEN_TYPE |
| Avg. Disk sec/Read | Shows the average time, in seconds, of a read operation from the disk. | PERF_AVERAGE_TIMER |
| Avg. Disk sec/Transfer | Shows the time, in seconds, of the average disk transfer. | PERF_AVERAGE_TIMER |
| Avg. Disk sec/Write | Shows the average time, in seconds, of a write operation to the disk. | PERF_AVERAGE_TIMER |
| Avg. Disk Write Queue Length | Shows the average number of write requests that were queued for the selected disk during the sample interval. | PERF_COUNTER_LARGE_QUEUELEN_TYPE |

| | | |
|---|---|---|
| Current Disk Queue Length | Shows the number of requests outstanding on the disk at the time that the performance data is collected. It also includes requests in service at the time of the collection. This is a snapshot, not an average over a time interval. Multispindle disk devices can have multiple requests that are active at one time, but other concurrent requests are awaiting service. This counter might reflect a transitory high or low queue length, but if there is a sustained load on the disk drive, it is likely that this will be consistently high. Requests experience delays that are proportional to the length of this queue, minus the number of spindles on the disks. For good performance, this difference should average less than two. | PERF_COUNTER_RAWCOUNT |
| Disk Bytes/sec | Shows the rate, in incidents per second, at which bytes were transferred to or from the disk during write or read operations. | PERF_COUNTER_BULK_COUNT |
| Disk Read Bytes/sec | Shows the rate, in incidents per second, at which bytes were transferred from the disk during read operations. | PERF_COUNTER_BULK_COUNT |
| Disk Reads/sec | Shows the rate, in incidents per second, at which read operations were performed on the disk. | PERF_COUNTER_COUNTER |
| Disk Transfers/sec | Shows the rate, in incidents per second, at which read and write operations were performed on the disk. | PERF_COUNTER_COUNTER |
| Disk Write Bytes/sec | Shows the rate, in incidents per second, at which bytes were transferred to the disk during write operations. | PERF_COUNTER_BULK_COUNT |
| Disk Writes/sec | Shows the rate, in incidents per second, at which write operations were performed on the disk. | PERF_COUNTER_COUNTER |
| Free Megabytes | Shows the unallocated space, in megabytes, on the disk drive. One megabyte is equal to 1,048,576 bytes. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Finding Cache Bottlenecks
Find disk bottlenecks

# MacFile Server Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The MacFile Server performance object consists of counters that provide information about a system running File Server for Macintosh.

| Counter Name | Description | Counter Type |
|---|---|---|
| Current Files Open | Shows the number of internal files currently open in the MacFile server. This count does not include files opened on behalf of Macintosh clients. | PERF_COUNTER_RAWCOUNT |
| Current Paged Memory | Shows the current amount of paged memory resources used by the MacFile server. | PERF_COUNTER_RAWCOUNT |
| Current Nonpaged Memory | Shows the current amount of nonpaged memory resources used by the MacFile Server. | PERF_COUNTER_RAWCOUNT |
| Current Sessions | Shows the number of sessions currently connected to the MacFile Server. This counter reports current server activity. | PERF_COUNTER_RAWCOUNT |
| Failed Logons | Shows the number of failed attempts to log on to the MacFile server. These values can indicate that password-guessing programs are being used to undermine the security on the server. | PERF_COUNTER_RAWCOUNT |
| Max Nonpaged Memory | Shows the maximum amount of nonpaged memory resources used by the MacFile Server. | PERF_COUNTER_RAWCOUNT |
| Max Paged Memory | Shows the maximum amount of paged memory resources used by the MacFile Server. | PERF_COUNTER_RAWCOUNT |
| Maximum Files Open | Shows the maximum number of internal files open at one time in the MacFile Server. This count does not include files opened on behalf of Macintosh clients. | PERF_COUNTER_RAWCOUNT |
| Maximum Sessions | Shows the maximum number of sessions connected at one time to the MacFile Server. This counter indicates the level of usage of the server. | PERF_COUNTER_RAWCOUNT |
| Data Read/sec | Shows the rate, in incidents per second, at which bytes were read from the disk. | PERF_COUNTER_BULK_COUNT |
| Data Written/sec | Shows the rate, in incidents per second, at which bytes were written to the disk. | PERF_COUNTER_BULK_COUNT |
| Data Received/sec | Shows the rate, in incidents per second, at which bytes were received from the network. This counter indicates how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Data Transmitted/sec | Shows the rate, in incidents per second, at which bytes were sent on the network. This counter indicates how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Current Queue Length | Shows the number of outstanding work items waiting to be processed. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Maximum Queue Length | Shows the maximum number of outstanding work items waiting at one time. | PERF_COUNTER_RAWCOUNT |
| Current Threads | Shows the current number of threads used by the MacFile Server. This counter indicates how busy the server is. | PERF_COUNTER_RAWCOUNT |
| Maximum Threads | Shows the maximum number of threads used by the MacFile Server. This counter indicates the peak usage level of the server. | PERF_COUNTER_RAWCOUNT |

# Memory Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Memory performance object consists of counters that describe the behavior of physical and virtual memory on the computer. Physical memory is the amount of RAM on the computer. Virtual memory consists of space in physical memory and on disk. Many of the memory counters monitor paging, which is the movement of pages of code and data between disk and physical memory. Excessive paging is a symptom of a memory shortage and can cause delays that interfere with all system processes.

| Counter Name | Description | Counter Type |
|---|---|---|
| % Committed Bytes In Use | Shows the ratio of Committed Bytes to the Commit Limit. Committed memory is physical memory in use, for which space has been reserved in the paging file(s) so that it can be written to disk. The commit limit is determined by the size of the paging file. If the paging file is enlarged, the commit limit increases, and the ratio is reduced. | PERF_RAW_FRACTION |
| Available Bytes | Shows the amount of physical memory, in bytes, immediately available for allocation to a process or for system use. It is equal to the sum of memory assigned to the standby (cached), free, and zero page lists. | PERF_COUNTER_RAWCOUNT |
| Available KBytes | Shows the amount of physical memory, in Kilobytes, immediately available for allocation to a process or for system use. It is equal to the sum of memory assigned to the standby (cached), free, and zero page lists. | PERF_COUNTER_RAWCOUNT |
| Available MBytes | Shows the amount of physical memory, in Megabytes, immediately available for allocation to a process or for system use. It is equal to the sum of memory assigned to the standby (cached), free, and zero page lists. | PERF_COUNTER_RAWCOUNT |
| Cache Bytes | Shows the sum of the values of System Cache Resident Bytes, System Driver Resident Bytes, System Code Resident Bytes, and Pool Paged Resident Bytes. | PERF_COUNTER_RAWCOUNT |
| Cache Bytes Peak | Shows the maximum number of bytes used by the file system cache since the system was last started. This might be larger than the current size of the cache. | PERF_COUNTER_RAWCOUNT |
| Cache Faults/sec | Shows the rate, in incidents per second, at which faults occured when a page that was sought in the file system cache was not found and was be retrieved either from elsewhere in memory (a soft fault) or from disk (a hard fault). This counter shows the total number of faults, without regard for the number of pages faulted in each operation. | PERF_COUNTER_COUNTER |
| Commit Limit | Shows the amount of virtual memory, in bytes, that can be committed without having to extend the paging file(s). Committed memory is physical memory that has space reserved on the disk paging file(s). There can be one or more paging files on each physical drive. If the paging file(s) are expanded, this limit increases accordingly. | PERF_COUNTER_RAWCOUNT |
| Committed Bytes | Shows the amount of committed virtual memory, in bytes. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Demand Zero Faults/sec | Shows the average rate, in incidents per second, at which page faults required a zeroed page to satisfy the fault. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval. Zeroed pages (pages emptied of previously stored data and filled with zeroes) prevent processes from seeing data stored by earlier processes that used the same memory space. This counter displays the number of faults, without regard to the number of pages retrieved to satisfy the fault. | PERF_COUNTER_COUNTER |
| Free System Page Table Entries | Shows the number of page table entries not in use by the system. | PERF_COUNTER_RAWCOUNT |
| Page Faults/sec | Shows the average number of pages faulted per second, which is equal to the number of page fault operations because only one page is faulted in each fault operation. This counter includes both hard faults (those that require disk access) and soft faults (where the faulted page is found elsewhere in physical memory). Most processors can handle large numbers of soft faults without significant consequence. However, hard faults can cause delaysbecause they require disk access.. | PERF_COUNTER_COUNTER |
| Page Reads/sec | Shows the rate, in incidents per second, at which the disk was read to resolve hard page faults. This counter shows numbers of read operations, without regard to the number of pages retrieved in each operation. Hard page faults occur when a process references a page in virtual memory that must be retrieved from disk because it is not in its working set or elsewhere in physical memory. This counter is a primary indicator for the kinds of faults that cause system-wide delays. It includes read operations to satisfy faults in the file system cache (usually requested by applications) and in noncached mapped memory files. Compare the value of Page Reads/sec to the value of Pages Input/sec to find an average of how many pages were read during each read operation. | PERF_COUNTER_COUNTER |
| Page Writes/sec | Shows the rate, in incidents per second, at which pages were written to disk to free up space in physical memory. Pages are written to disk only if they are changed while in physical memory, so they are likely to hold data, not code. This counter shows write operations, without regard to the number of pages written in each operation. | PERF_COUNTER_COUNTER |
| Pages Input/sec | Shows the rate, in incidents per second, at which pages were read from disk to resolve hard page faults. Hard page faults occur when a process refers to a page in virtual memory that must be retrieved from disk because it is not in its working set or elsewhere in physical memory. When a page is faulted, the system tries to read multiple contiguous pages into memory to maximize the benefit of the read operation. Compare Pages Input/sec to Page Reads/sec to find the average number of pages read into memory during each read operation | PERF_COUNTER_COUNTER |
| Pages Output/sec | Shows the rate, in incidents per second, at which pages were written to disk to free up space in physical memory. A high rate of pages output might indicate a memory shortage. The Windows Server 2003 family writes more pages back to disk to free up space when physical memory is in short supply. This counter shows numbers of pages, and can be compared to other counts of pages without conversion. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Pages/sec | Shows the rate, in incidents per second, at which pages were read from or written to disk to resolve hard page faults. This counter is a primary indicator for the kinds of faults that cause system-wide delays. It is the sum of Pages Input/sec and Pages Output/sec. It is counted in numbers of pages, so it can be directly compared to other counts of pages such as Page Faults/sec. It includes pages retrieved to satisfy faults in the file system cache (usually requested by applications) and noncached mapped memory files. | PERF_COUNTER_COUNTER |
| Pool Nonpaged Allocs | Shows the number of calls to allocate space in the nonpaged pool. This counter is measured in numbers of calls to allocate space, regardless of the amount of space allocated in each call. | PERF_COUNTER_RAWCOUNT |
| Pool Nonpaged Bytes | Shows the size, in bytes, of the nonpaged pool. Pool Nonpaged Bytes is calculated differently than Process\Pool Nonpaged Bytes, so it might not equal Process(_Total )\Pool Nonpaged Bytes. | PERF_COUNTER_RAWCOUNT |
| Pool Paged Allocs | Shows the number of calls to allocate space in the paged pool. This counter is measured in numbers of calls to allocate space, regardless of the amount of space allocated in each call. | PERF_COUNTER_RAWCOUNT |
| Pool Paged Bytes | Shows the size, in bytes, of the paged pool. Pool Paged Bytes is calculated differently than Process\Pool Paged Bytes, so it might not equal Process(_Total )\Pool Paged Bytes. | PERF_COUNTER_RAWCOUNT |
| Pool Paged Resident Bytes | Shows the size, in bytes, of the paged pool. Space used by the paged and nonpaged pools is taken from physical memory, so a pool that is too large denies memory space to processes. | PERF_COUNTER_RAWCOUNT |
| System Cache Resident Bytes | Shows the size, in bytes, of pageable operating system code in the file system cache. This value includes only current physical pages and does not include any virtual memory pages that are not currently resident. It does not equal the System Cache value shown in Task Manager. As a result, this value may be smaller than the actual amount of virtual memory in use by the file system cache. This value is a component of System Code Resident Bytes that represents all pageable operating system code that is currently in physical memory. | PERF_COUNTER_RAWCOUNT |
| System Code Resident Bytes | Shows the size, in bytes, of operating system code currently in physical memory that can be written to disk when not in use. This value is a component of System Code Total Bytes, which also includes operating system code on disk. System Code Resident Bytes (and System Code Total Bytes) does not include code that must remain in physical memory. | PERF_COUNTER_RAWCOUNT |
| System Code Total Bytes | Shows the size, in bytes, of pageable operating system code currently in virtual memory. It is a measure of the amount of physical memory being used by the operating system that can be written to disk when not in use. This value is calculated by adding the bytes in Ntoskrnl.exe, Hal.dll, the boot drivers, and file systems loaded by Ntldr/osloader. This counter does not include code that must remain in physical memory. | PERF_COUNTER_RAWCOUNT |
| System Driver Resident Bytes | Shows the size, in bytes, of pageable physical memory being used by device drivers. The counter is the working set (physical memory area) of the drivers. This value is a component of System Driver Total Bytes, which also includes driver memory that has been written to disk. Neither System Driver Resident Bytes nor System Driver Total Bytes includes memory that cannot be written to disk. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| System Driver Total Bytes | Shows the size, in bytes, of pageable virtual memory currently being used by device drivers. Pageable memory can be written to disk when it is not being used. It includes physical memory (System Driver Resident Bytes) and code and data written to disk. This counter is a component of System Code Total Bytes. | PERF_COUNTER_RAWCOUNT |
| Transition Faults/sec | Shows the rate, in incidents per second, at which page faults were resolved by recovering pages without additional disk activity, including pages that were being used by another process sharing the page, or that were on the modified page list or the standby list, or that were being written to disk at the time of the page fault. This counter is also equal to the number of pages faulted because only one page is faulted in each operation. | PERF_COUNTER_COUNTER |
| Write Copies/sec | Shows the rate, in incidents per second, at which page faults were caused by attempts to write that were satisfied by copying the page from elsewhere in physical memory. This is an economical way of sharing data since pages are only copied when they are written to; otherwise, the page is shared. This counter shows the number of copies, without regard to the number of pages copied in each operation. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Find memory bottlenecks
Finding Cache Bottlenecks

# MSMQ Queue Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The MSMQ Queue performance object monitors message statistics for a specified queue. There are instances for each queue on the computer. In addition, if a Message Queuing–based application has a private queue open on another computer, an instance for that queue is also available. The computer queue instance represents the computer's source journal and dead letter queues.

| Counter Name | Description | Counter Type |
|---|---|---|
| MSMQ Queue | Shows the Message Queuing queues in this computer. | PERF_COUNTER_RAWCOUNT |
| Bytes in Journal Queue | Shows the total number of bytes that reside in the journal queue. For the computer queue instance, this represents the computer journal queue. | PERF_COUNTER_RAWCOUNT |
| Bytes in Queue | Shows the total number of bytes that currently reside in the queue. For the computer queue instance, this represents the dead letter queue. | PERF_COUNTER_RAWCOUNT |
| Messages in Journal Queue | Shows the total number of messages that currently reside in the journal queue. For the computer queue instance, this represents the computer journal queue. | PERF_COUNTER_RAWCOUNT |
| Messages in Queue | Shows the total number of messages that currently reside in the queue. For the computer queue instance, this represents the dead letter queue. For queues residing on the local computer, this counter shows how many messages are waiting to be read. For queues residing on a remote computer, the counter shows the number of messages waiting to be sent. This counter is useful in determining if message handling is stalled. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
MSMQ Service Object
MSMQ Session Object

# MSMQ Service Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The MSMQ Service performance object monitors session and message statistics for a specified computer that is running Message Queuing.

| Counter Name | Description | Counter Type |
|---|---|---|
| Incoming Messages/sec | Shows the rate, in incidents per second, at which Message Queuing messages were received by the Message Queuing service. This counter is useful for determining if the local Message Queuing service has received messages. | PERF_COUNTER_COUNTER |
| IP Sessions | Shows the number of open IP sessions. | PERF_COUNTER_RAWCOUNT |
| MSMQ Incoming Messages | Shows the number of Message Queuing messages that were received by the Message Queuing service. | PERF_COUNTER_RAWCOUNT |
| MSMQ Outgoing Messages | Shows the number of Message Queuing messages that were sent by the Message Queuing service. | PERF_COUNTER_RAWCOUNT |
| Outgoing Messages/sec | Shows the rate, in incidents per second, at which Message Queuing messages were sent by the Message Queuing service. This counter is useful for determining if the local Message Queuing service has sent messages. | PERF_COUNTER_COUNTER |
| Sessions | Shows the total number of open network sessions. A session is opened whenever a message is sent or received, and the session is closed automatically after two to five minutes of inactivity. The Windows Server 2003 family closes a session after two minutes; Windows XP Professional and most other Windows operating systems close a session after five minutes. This counter indicates if messaging activity has occurred recently. | PERF_COUNTER_RAWCOUNT |
| Total Bytes in All Queues | Shows the total number of bytes in all active queues under the Message Queuing service. | PERF_COUNTER_RAWCOUNT |
| Total Messages in All Queues | Shows the total number of messages in all active queues under the Message Queuing service. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
MSMQ Queue Object
MSMQ Session Object

# MSMQ Session Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The MSMQ Session performance object monitors statistics about active sessions between computers running Message Queuing. You can choose to have an instance for each session.

| Counter Name | Description | Counter Type |
|---|---|---|
| Incoming Bytes | Shows the total number of bytes that were received through the selected session. | PERF_COUNTER_RAWCOUNT |
| Incoming Bytes/sec | Shows the rate, in incidents per second, at which bytes were received through the selected session. | PERF_COUNTER_RAWCOUNT |
| Incoming Messages | Shows the total number of messages that were received through the selected session. | PERF_COUNTER_RAWCOUNT |
| Incoming Messages/sec | Shows the rate, in incidents per second, at which Message Queuing messages were received through the selected session. | PERF_COUNTER_RAWCOUNT |
| Outgoing Bytes | Shows the total number of bytes that were sent through the selected session. | PERF_COUNTER_RAWCOUNT |
| Outgoing Bytes/sec | Shows the rate, in incidents per second, at which Message Queuing messages were sent through the selected session. | PERF_COUNTER_RAWCOUNT |
| Outgoing Messages | Shows the total number of Message Queuing messages that were sent through the selected session. | PERF_COUNTER_RAWCOUNT |
| Outgoing Messages/sec | Shows the rate, in incidents per second, at which Message Queuing messages were sent through the selected session. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
MSMQ Queue Object
MSMQ Service Object

# NBT Connection Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The NBT Connection performance object consists of counters that measure the rates at which bytes are sent and received over the NetBIOS over a TCP/IP (NetBT or NBT) connection between the local computer and a remote computer. The NetBT or NBT protocol provides NetBIOS support for the TCP/IP protocol. The NBT connection is identified by the name of the remote computer.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Received/sec | Shows the rate at which bytes are received by the local computer over an NBT connection to a remote computer. All bytes received by the local computer over the connection are counted. | PERF_COUNTER_BULK_COUNT |
| Bytes Sent/sec | Shows the rate at which bytes are sent by the local computer over an NBT connection to a remote computer. All bytes sent by the local computer over the connection are counted. | PERF_COUNTER_BULK_COUNT |
| Bytes Total/sec | Shows the rate at which bytes are transmitted by the local computer over an NBT connection to a remote computer. All bytes transmitted by the local computer over the connection are counted. | PERF_COUNTER_BULK_COUNT |

# DUPTILE1_Network Interface Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Network Interface performance object consists of counters that measure the rates at which bytes and packets are sent and received over a TCP/IP connection. It includes counters that monitor connection errors.

The Network Interface counters display data about the network adapters on the server computer. The first instance of the Network Interface object (Instance 1) that you see in System Monitor represents the *loopback*. The *loopback* is a local path through the protocol driver and the network adapter. All other instances represent installed network adapters.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Received/sec | Shows the rate, in incidents per second, at which bytes were received over each network adapter. The counted bytes include framing characters. Bytes Received/sec is a subset of Bytes Total/sec. | PERF_COUNTER_COUNTER |
| Bytes Sent/sec | Shows the rate, in incidents per second, at which bytes were sent over each network adapter. The counted bytes include framing characters. Bytes Sent/sec is a subset of Bytes Total/sec. | PERF_COUNTER_COUNTER |
| Bytes Total/sec | Shows the rate, in incidents per second, at which bytes were sent and received on the network interface, including framing characters. Bytes Total/sec is the sum of the values of Bytes Received/sec and Bytes Sent/sec. | PERF_COUNTER_COUNTER |
| Current Bandwidth | Shows an estimate of the current bandwidth of the network interface in bits per second (bps). For interfaces that do not vary in bandwidth, or for those where no accurate estimation can be made, this value is the nominal bandwidth. | PERF_COUNTER_RAWCOUNT |
| Output Queue Length | Shows the length, in number of packets, of the output packet queue. If this is longer than two packets, it indicates that there are delays, and if possible the bottleneck should be found and eliminated. Since the requests are queued by Network Driver Interface Specification (NDIS) in this implementation, this value is always 0. | PERF_COUNTER_RAWCOUNT |
| Packets Outbound Discarded | Shows the number of outbound packets to be discarded, even though no errors were detected to prevent transmission. One possible reason for discarding such a packet could be to free up buffer space. | PERF_COUNTER_RAWCOUNT |
| Packets Outbound Errors | Shows the number of outbound packets that could not be transmitted because of errors. | PERF_COUNTER_RAWCOUNT |
| Packets Received Discarded | Shows the number of inbound packets that were discarded, even though no errors were detected to prevent their being delivered to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space. | PERF_COUNTER_RAWCOUNT |
| Packets Received Errors | Shows the number of inbound packets that contained errors that prevented them from being delivered to a higher-layer protocol. | PERF_COUNTER_RAWCOUNT |
| Packets Received Non-Unicast/sec | Shows the rate, in incidents per second, at which non-unicast (subnet broadcast or subnet multicast) packets were delivered to a higher-layer protocol. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Packets Received Unicast/sec | Shows the rate, in incidents per second, at which subnet-unicast packets were delivered to a higher-layer protocol. | PERF_COUNTER_COUNTER |
| Packets Received Unknown | Shows the number of packets received through the interface that were discarded because of an unknown or unsupported protocol. | PERF_COUNTER_RAWCOUNT |
| Packets Received/sec | Shows the rate, in incidents per second, at which packets were received on the network interface. | PERF_COUNTER_COUNTER |
| Packets Sent/sec | Shows the rate, in incidents per second, at which packets were sent on the network interface. | PERF_COUNTER_COUNTER |
| Packets Sent Non-Unicast/sec | Shows the rate, in incidents per second, at which packets were requested to be transmitted to non-unicast (subnet broadcast or subnet multicast) addresses by higher-level protocols. This counter includes packets that were discarded or not sent. | PERF_COUNTER_COUNTER |
| Packets Sent Unicast/sec | Shows the rate, in incidents per second, at which packets were requested to be transmitted to subnet-unicast addresses by higher-level protocols. This counter includes the packets that were discarded or not sent. | PERF_COUNTER_COUNTER |

# NNTP Commands Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The NNTP Commands performance object includes counters for all Network News Transport Protocol (NNTP) commands processed by the NNTP service, which is installed as an optional component of Internet Information Services.

NNTP counters are started for each virtual server instance and are reset to zero when the virtual server is started. A single NNTP service can have many virtual servers (also called sites).

| Counter Name | Description | Counter Type |
|---|---|---|
| Article Commands | Shows the number of ARTICLE commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Article Commands/sec | Shows the rate, in incidents per second, at which ARTICLE commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Check Commands | Shows the number of CHECK commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Check Commands/sec | Shows the rate, in incidents per second, at which CHECK commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Group Commands | Shows the number of GROUP commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Group Commands/sec | Shows the rate, in incidents per second, at which GROUP commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Help Commands | Shows the number of HELP commands that were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Help Commands/sec | Shows the rate, in incidents per second, at which HELP commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| IHave Commands | Shows the number of IHAVE commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| IHave Commands/sec | Shows the rate, in incidents per second, at which IHAVE commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Last Commands | Shows the number of LAST commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Last Commands/sec | Shows the rate, in incidents per second, at which LAST commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| List Commands | Shows the number of LIST commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| List Commands/sec | Shows the rate, in incidents per second, at which LIST commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Logon Attempts | Shows the number of logon attempts that have been made to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Logon Attempts/sec | Shows the rate, in incidents per second, at which logon attempts have been made to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Logon Failures | Shows the number of logons that failed on the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Logon Failures/sec | Shows the rate, in incidents per second, at which logons failed on the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Mode Commands | Shows the number of MODE commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Mode Commands/sec | Shows the rate, in incidents per second, at which MODE commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Newgroups Commands | Shows the number of NEWGROUPS commands that were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Newgroups Commands/sec | Shows the rate, in incidents per second, at which NEWSGROUPS commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Newnews Commands | Shows the number of NEWNEWS commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Newnews Commands/sec | Shows the rate, in incidents per second, at which NEWNEWS commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Next Commands | Shows the number of NEXT commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Next Commands/sec | Shows the rate, in incidents per second, at which NEXT commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Post Commands | Shows the number of POST commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Post Commands/sec | Shows the rate, in incidents per second, at which POST commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Quit Commands | Shows the number of QUIT commands that were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Quit Commands/sec | Shows the rate, in incidents per second, at which QUIT commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Search Commands | Shows the number of SEARCH commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Search Commands/sec | Shows the rate, in incidents per second, at which SEARCH commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Stat Commands | Shows the number of STAT commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Stat Commands/sec | Shows the rate, in incidents per second, at which STAT commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Takethis Commands | Shows the number of TAKETHIS commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Takethis Commands/sec | Shows the rate, in incidents per second, at which TAKETHIS commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| XHdr Commands | Shows the number of XHDR commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| XHdr Commands/sec | Shows the rate, in incidents per second, at which XHDR commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| XOver Commands | Shows the number of XOVER commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| XOver Commands/sec | Shows the rate, in incidents per second, at which XOVER commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| XPat Commands | Shows the number of XPAT commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| XPat Commands/sec | Shows the rate, in incidents per second, at which XPAT commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| XReplic Commands | Shows the number of XREPLIC commands received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| XReplic Commands/sec | Shows the rate, in incidents per second, at which XREPLIC commands were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
NNTP Server Object

# NNTP Server Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The NNTP Server performance object monitors posting, authentication, and connection activity on a Network News Transport Protocol (NNTP) Server. The NNTP service is an optional component of Internet Information Services (IIS).

NNTP counters are started for each virtual server instance and are reset to zero when the virtual server is started. A single NNTP service can have many virtual servers (also called sites).

In addition to monitoring these counters, it is recommended that you also monitor the system counters Physical Disk Object, Physical Disk Object, Physical Disk Object, and Processor Object.

| Counter Name | Description | Counter Type |
|---|---|---|
| Article Map Entries | Shows the entries inserted into the article mapping table of the NNTP Server. | PERF_COUNTER_RAWCOUNT |
| Article Map Entries/sec | Shows the rate, in incidents per second, at which entries were inserted into the article mapping table of the NNTP Server. | PERF_COUNTER_COUNTER |
| Articles Deleted | Shows the number of articles deleted on the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Articles Deleted/sec | Shows the rate, in incidents per second, at which articles were deleted from the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Articles Posted | Shows the number of articles posted (committed) to the NNTP Server since the NNTP service was last started. This can be used to determine if the server is accepting posts. If this value is smaller than that of Articles Received, it indicates that the server is rejecting articles and you need to investigate the cause. | PERF_COUNTER_RAWCOUNT |
| Articles Posted/sec | Shows the rate, in incidents per second, at which articles were posted to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Articles Received | Shows the total number of files received by the NNTP Server since the NNTP service was last started. Files received can include rejected articles. If this value is greater than that of Articles Posted, this indicates that the server is rejecting articles and you need to investigate the cause. | PERF_COUNTER_RAWCOUNT |
| Articles Received/sec | Shows the rate, in incidents per second, at which files were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Articles Sent | Shows the total number of files sent by the NNTP Server to other servers since the NNTP service was last started. Files sent can include rejected articles. | PERF_COUNTER_RAWCOUNT |
| Articles Sent/sec | Shows the rate, in incidents per second, at which files were sent by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Articles Total | Shows the sum of NNTP Server\Articles Sent and NNTP Server\Articles Received. This is the total number of files transferred by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Bytes Received/sec | Shows the rate, in incidents per second, at which that data bytes were received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_BULK_COUNT |
| Bytes Sent/sec | Shows the rate, in incidents per second, at which data bytes were sent by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_BULK_COUNT |
| Bytes Total/sec | Shows the sum of NNTP Server\Bytes Sent/sec and NNTP Server\Bytes Received/sec. This is the total rate, in incidents per second, at which bytes were transferred by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_BULK_COUNT |
| Control Messages Failed | Shows the total number of control messages failed or not applied by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Control Messages Received | Shows the total number of control messages received by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Current Anonymous Users | Shows the number of anonymous users who are currently connected to the NNTP Server. | PERF_COUNTER_RAWCOUNT |
| Current Connections | Shows the current number of connections to the NNTP Server. | PERF_COUNTER_RAWCOUNT |
| Current NonAnonymous Users | Shows the number of nonanonymous users currently connected to the NNTP Server. | PERF_COUNTER_RAWCOUNT |
| Current Outbound Connections | Shows the number of current outbound connections being made by the NNTP Server. | PERF_COUNTER_RAWCOUNT |
| Failed Outbound Logons | Shows the number of failed outbound logons made by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| History Map Entries | Shows the number of entries inserted into the history mapping table of the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| History Map Entries/sec | Shows the rate, in incidents per second, at which entries were inserted into the history mapping table of the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |
| Maximum Anonymous Users | Shows the maximum number of anonymous users that were simultaneously connected to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Maximum Connections | Shows the maximum number of simultaneous connections to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Maximum NonAnonymous Users | Shows the maximum number of nonanonymous users that were simultaneously connected to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Moderated Postings Failed | Shows the total number of moderated postings that the NNTP Server failed to send to a Simple Mail Transfer Protocol (SMTP) Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Moderated Postings Sent | Shows the total number of moderated postings that the NNTP Server attempted to send to an SMTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Sessions Flow Controlled | Shows the number of client sessions that were currently in a flow controlled state in the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Anonymous Users | Shows the total number of anonymous users that have connected to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Connections | Shows the number of connections that have been made to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Total NonAnonymous Users | Shows the total number of nonanonymous users that have connected to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Outbound Connections | Shows the number of outbound connections that have been made by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Outbound Connections Failed | Shows the number of unsuccessful outbound connections that have been made by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Passive Feeds | Shows the number of passive feeds that were accepted by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Pull Feeds | Shows the number of pull feeds that were made by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Total Push Feeds | Shows the number of push feeds that were made by the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Total SSL Connections | Shows the number of Secure Sockets Layer (SSL) connections that have been made to the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Xover Entries | Shows the number of XOVER entries in the XOVER table of the NNTP Server since the NNTP service was last started. | PERF_COUNTER_RAWCOUNT |
| Xover Entries/sec | Shows the rate, in incidents per second, at which XOVER entries were inserted in the XOVER table of the NNTP Server since the NNTP service was last started. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
NNTP Commands Object

# NTDS Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The NTDS object consists of counters that provide statistics about the activity of the Active Directory directory service. The object provides information about the following:

- The Name Service Provider Interface (NSPI), used to facilitate communication between Active Directory and Exchange Directory Service (XDS)
- The Local Security Authority (LSA), a protected subsystem that maintains security for the local computer
- The Security Accounts Manager (SAM) interface, which provides compatibility between Windows Server 2003, Windows 2000, and Windows NT 4.0 domains
- The Lightweight Directory Access Protocol (LDAP) interface, which provides the application programming interface (API) for LDAP clients and exposes the Active Directory Services Interface (ADSI) so additional applications may be written that can talk to Active Directory
- The Directory System Agent (DSA), the Active Directory process that runs on each domain controller and manages all the directory service functions
- The Knowledge Consistency Checker (KCC), an Active Directory component that is responsible for generating the replication topology between domain controllers

| Counter Name | Description | Counter Type |
|---|---|---|
| AB ANR/sec | Shows the rate at which Address Book clients perform Ambiguous Name Resolutions (ANR) operations. | PERF_COUNTER_COUNTER |
| AB Browses/sec | Shows the rate at which Address Book clients perform browse operations. | PERF_COUNTER_COUNTER |
| AB Client Sessions | Shows the number of connected Address Book client sessions. | PERF_COUNTER_RAWCOUNT |
| AB Matches/sec | Shows the rate at which Address Book clients perform find operations. | PERF_COUNTER_COUNTER |
| AB Property Reads/sec | Shows the rate at which Address Book clients perform property read operations. | PERF_COUNTER_COUNTER |
| AB Proxy Lookups/sec | Shows the rate at which proxy clients perform search operations. | PERF_COUNTER_COUNTER |
| AB Searches/sec | Shows the rate at which Address Book clients perform key search operations. | PERF_COUNTER_COUNTER |
| DRA Highest USN Committed (High part) | Shows the high-order 32 bits of the highest update sequence number (USN) committed on the DSA. | PERF_COUNTER_RAWCOUNT |
| DRA Highest USN Committed (Low part) | Shows the low-order 32 bits of the highest USN committed on the DSA. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| DRA Highest USN Issued (High part) | Shows the high-order 32 bits of the highest USN issued on the DSA. | PERF_COUNTER_RAWCOUNT |
| DRA Highest USN Issued (Low part) | Shows the low-order 32 bits of the highest USN issued on the DSA. | PERF_COUNTER_RAWCOUNT |
| DRA Inbound Bytes Compressed (Between Sites, After Compression)/sec | Shows the compressed size of inbound compressed replication data (size after compression, from DSAs in other sites). | PERF_COUNTER_COUNTER |
| DRA Inbound Bytes Compressed (Between Sites, Before Compression)/sec | Shows the original size of inbound compressed replication data (size before compression, from DSAs in other sites). | PERF_COUNTER_COUNTER |
| DRA Inbound Bytes Not Compressed (Within Site)/sec | Shows the number of incoming bytes replicated that were not compressed at the source (that is, from DSAs in the same site). | PERF_COUNTER_COUNTER |
| DRA Inbound Bytes Total/sec | Shows the total number of bytes replicated in. This counter is the sum of the number of uncompressed bytes (never compressed) and the number of compressed bytes (after compression). | PERF_COUNTER_COUNTER |
| DRA Inbound Full Sync Objects Remaining | Shows the number of objects remaining until the full synchronization is completed (when set). | PERF_COUNTER_RAWCOUNT |
| DRA Inbound Object Updates Remaining in Packet | Shows the number of object updates received in the current directory replication update packet that have not yet been applied to the local server. | PERF_COUNTER_RAWCOUNT |
| DRA Inbound Objects Applied/sec | Shows the rate at which replication updates received from replication partners are applied by the local directory service. This counter excludes changes that are received but not applied (because, for example, the change has already been made). This indicates how much replication update activity is occurring on the server as a result of changes generated on other servers. | PERF_COUNTER_COUNTER |
| DRA Inbound Objects Filtered/sec | Shows the number of objects received from inbound replication partners that contained no updates that needed to be applied. | PERF_COUNTER_COUNTER |
| DRA Inbound Objects/sec | Shows the number of objects received from neighbors through inbound replication. A *neighbor* is a domain controller from which the local domain controller replicates locally. | PERF_COUNTER_COUNTER |
| DRA Inbound Properties Applied/sec | Shows the number of properties that are updated due to the incoming property's winning the reconciliation logic that determines the final value to be replicated. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| DRA Inbound Properties Filtered/sec | Shows the number of property changes that are received during the replication that we have already seen. | PERF_COUNTER_COUNTER |
| DRA Inbound Properties Total/sec | Shows the total number of object properties received from inbound replication partners. | PERF_COUNTER_COUNTER |
| DRA Inbound Values (DNs only)/sec | Shows the number of object property values received from inbound replication partners that are distinguished names (DNs) that reference other objects. DN values, such as group or distribution list memberships, are generally more expensive to apply than other kinds of values. | PERF_COUNTER_COUNTER |
| DRA Inbound Values Total/sec | Shows the total number of object property values received from inbound replication partners. Each inbound object has one or more properties, and each property has zero or more values. Zero values indicate property removal. | PERF_COUNTER_COUNTER |
| DRA Outbound Bytes Compressed (Between Sites, After Compression)/sec | Shows the compressed size of outbound compressed replication data (size after compression, from DSAs in other sites). | PERF_COUNTER_COUNTER |
| DRA Outbound Bytes Compressed (Between Sites, Before Compression)/sec | Shows the original size of outbound compressed replication data (size before compression, from DSAs in other sites). | PERF_COUNTER_COUNTER |
| DRA Outbound Bytes Not Compressed (Within Site)/sec | Shows the number of bytes replicated out that were not compressed (that is, from DSAs in the same site). | PERF_COUNTER_COUNTER |
| DRA Outbound Bytes Total/sec | Shows the total number of bytes replicated out. This counter is the sum of the number of uncompressed bytes (never compressed) and the number of compressed bytes (after compression). | PERF_COUNTER_COUNTER |
| DRA Outbound Objects Filtered/sec | Shows the number of objects that were determined by outbound replication to have no updates that the outbound partner did not already have. | PERF_COUNTER_COUNTER |
| DRA Outbound Objects/sec | Shows the number of objects replicated out. | PERF_COUNTER_COUNTER |
| DRA Outbound Properties/sec | Shows the number of properties replicated out. | PERF_COUNTER_COUNTER |
| DRA Outbound Values (DNs only)/sec | Shows the number of object property values containing DNs sent to outbound replication partners. DN values, such as group or distribution list memberships, are generally more expensive to read than other kinds of values. | PERF_COUNTER_COUNTER |
| DRA Outbound Values Total/sec | Shows the number of object property values sent to outbound replication partners. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| DRA Pending Replication Synchronizations | Shows the number of directory synchronizations that are queued for this server but not yet processed. | PERF_COUNTER_RAWCOUNT |
| DRA Reads | Shows the percentage of directory reads coming from the directory replication agent (DRA). | PERF_RAW_FRACTION |
| DRA Searches | Shows the percentage of directory searches coming from the DRA. | PERF_RAW_FRACTION |
| DRA Sync Failures on Schema Mismatch | Shows the number of synchronization requests made to neighbors that failed because their schema are not synchronized. | PERF_COUNTER_RAWCOUNT |
| DRA Sync Requests Made | Shows the number of synchronization requests made to neighbors. | PERF_COUNTER_RAWCOUNT |
| DRA Sync Requests Successful | Shows the number of synchronization requests made to neighbors that were successfully returned. | PERF_COUNTER_RAWCOUNT |
| DRA Writes | Shows the percentage of directory writes coming from the DRA. | PERF_RAW_FRACTION |
| DS Client Binds/sec | Shows the number of Ntdsapi.dll binds per second serviced by this domain controller. | PERF_COUNTER_COUNTER |
| DS Client Name Translations/sec | Shows the number of Ntdsapi.dll name translations per second serviced by this domain controller. | PERF_COUNTER_COUNTER |
| DS Directory Reads/sec | Shows the number of directory reads per second. | PERF_COUNTER_COUNTER |
| DS Directory Searches/sec | Shows the number of directory searches per second. | PERF_COUNTER_COUNTER |
| DS Directory Writes/sec | Shows the number of directory writes per second. | PERF_COUNTER_COUNTER |
| DS Monitor List Size | Shows the number of requests to be notified when objects are updated that are currently registered with this DSA. | PERF_COUNTER_RAWCOUNT |
| DS Name Cache hit rate | Shows the percentage of directory object name component lookups that are satisfied out of the DSA's name cache. | PERF_RAW_FRACTION |
| DS Notify Queue Size | Shows the number of pending update notifications that have been queued but not yet transmitted to clients. | PERF_COUNTER_RAWCOUNT |
| DS Other Reads | Shows the percentage of directory reads that do not come from SAM, DRA, LDAP, LSA, XDS, KCC, or NSPI. | PERF_RAW_FRACTION |
| DS Other Searches | Shows the percentage of directory searches that do not come from SAM, DRA, LDAP, LSA, XDS, KCC, or NSPI. | PERF_RAW_FRACTION |
| DS Other Writes | Shows the percentage of directory writes that do not come from SAM, DRA, LDAP, LSA, XDS, KCC, or NSPI. | PERF_RAW_FRACTION |

| | | |
|---|---|---|
| DS Search sub-operations/sec | Shows the number of search suboperations per second. One search operation is made up of many suboperations. Each suboperation roughly corresponds to an object that the search causes the directory service to consider. | PERF_COUNTER_COUNTER |
| DS Security Descriptor Propagations Events | Shows the number of security descriptor propagation events that are queued but not yet processed. | PERF_COUNTER_RAWCOUNT |
| DS Security Descriptor Propagator Average Exclusion Time | Shows the average length of time that the security descriptor propagator spends waiting for exclusive access to database elements during a security descriptor propagation suboperation. | PERF_COUNTER_RAWCOUNT |
| DS Security Descriptor Propagator Runtime Queue | Shows the number of objects that remain to be examined while the current directory service security descriptor propagator event is being processed. | PERF_COUNTER_RAWCOUNT |
| DS Security Descriptor sub-operations/sec | Shows the number of security descriptor propagation suboperations per second. One security descriptor propagation operation is made up of many suboperations. | PERF_COUNTER_COUNTER |
| DS Server Binds/sec | Shows the number of domain controller–to–domain controller binds per second that are serviced by this domain controller. | PERF_COUNTER_COUNTER |
| DS Server Name Translations/sec | Shows the number of domain controller–to–domain controller name translations per second that are serviced by this domain controller. | PERF_COUNTER_COUNTER |
| DS Threads in Use | Shows the current number of threads in use by the directory service (which is different from the number of threads in the directory service process). This is the number of threads currently servicing client API calls; it can be used to indicate whether additional processors should be used. | PERF_COUNTER_RAWCOUNT |
| KCC Reads | Shows the percentage of directory reads that come from KCC. | PERF_RAW_FRACTION |
| KCC Searches | Shows the percentage of directory searches that come from KCC. | PERF_RAW_FRACTION |
| KCC Writes | Shows the percentage of directory writes that come from KCC. | PERF_RAW_FRACTION |
| KDC AS Requests | Shows the number of Authentication Server (AS) requests serviced by the Kerberos Key Distribution Center (KDC) per second. AS requests are used by the client to obtain a ticket-granting ticket. | PERF_COUNTER_COUNTER |
| KDC TGS Requests | Shows the number of Ticket Granting Server (TGS) requests serviced by the KDC per second. TGS requests are used by the client to obtain a ticket to a resource. | PERF_COUNTER_COUNTER |
| Kerberos Authentications | Shows the number of times per second that clients use a ticket to this domain controller to authenticate to this domain controller. | PERF_COUNTER_COUNTER |
| LDAP Active Threads | Shows the current number of threads in use by the LDAP subsystem of the local directory service. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| LDAP Bind Time | Shows the time, in milliseconds, taken for the last successful LDAP bind. | PERF_COUNTER_RAWCOUNT |
| LDAP Client Sessions | Shows the number of currently connected LDAP client sessions. | PERF_COUNTER_RAWCOUNT |
| LDAP Searches | Shows the percentage of directory searches coming from LDAP. | PERF_RAW_FRACTION |
| LDAP Searches/sec | Shows the rate at which LDAP clients perform search operations. | PERF_COUNTER_COUNTER |
| LDAP Successful Binds | Shows the percentage of LDAP bind attempts that are successful. | PERF_RAW_FRACTION |
| LDAP Successful Binds/sec | Shows the number of LDAP binds per second. | PERF_COUNTER_COUNTER |
| LDAP UDP operations/sec | Shows the number of User Datagram Protocol (UDP) operations that the LDAP server is processing per second. | PERF_COUNTER_COUNTER |
| LDAP Writes | Shows the percentage of directory writes coming from LDAP. | PERF_RAW_FRACTION |
| LDAP Writes/sec | Shows the rate at which LDAP clients perform write operations. | PERF_COUNTER_COUNTER |
| LSA Reads | Shows the percentage of directory reads that come from LSA. | PERF_RAW_FRACTION |
| LSA Searches | Shows the percentage of directory searches that come from LSA. | PERF_RAW_FRACTION |
| LSA Writes | Shows the percentage of directory writes that come from LSA. | PERF_RAW_FRACTION |
| NSPI Reads | Shows the percentage of directory reads that come from NSPI. | PERF_RAW_FRACTION |
| NSPI Searches | Shows the percentage of directory searches that come from NSPI. | PERF_RAW_FRACTION |
| NSPI Writes | Shows the percentage of directory writes that come from NSPI. | PERF_RAW_FRACTION |
| NTLM Authentications | Shows the number of NTLM authentications per second serviced by this domain controller. | PERF_COUNTER_COUNTER |
| SAM Account Group Membership Evaluations/sec | Shows the number of SAM Account Group Membership Evaluations per second. | PERF_COUNTER_COUNTER |
| SAM Create Machine Attempts/sec | Shows the number of SAM create machine attempts per second. | PERF_COUNTER_COUNTER |
| SAM Create User Attempts/sec | Shows the number of SAM create-user attempts per second. | PERF_COUNTER_COUNTER |
| SAM Enumerations/sec | Shows the number of SAM enumerations per second. | PERF_COUNTER_COUNTER |
| SAM GC Evaluations/sec | Shows the number of SAM global catalog evaluations per second. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| SAM Membership Changes/sec | Shows the number of SAM membership changes. | PERF_COUNTER_COUNTER |
| SAM Non-Transitive Membership Evaluations/sec | Shows the number of SAM nontransitive membership evaluations per second. | PERF_COUNTER_COUNTER |
| SAM Password Changes/sec | Shows the number of SAM password changes per second. | PERF_COUNTER_COUNTER |
| SAM Query Displays/sec | Shows the number of SAM query displays per second. | PERF_COUNTER_COUNTER |
| SAM Reads | Shows the percentage of directory reads that come from SAM. | PERF_RAW_FRACTION |
| SAM Resource Group Membership Evaluations/sec | Shows the number of SAM resource group membership evaluations per second. | PERF_COUNTER_COUNTER |
| SAM Searches | Shows the percentage of directory searches that come from SAM. | PERF_RAW_FRACTION |
| SAM Successful Create Machines/sec | Shows the number of machines that were successfully created per second. | PERF_COUNTER_COUNTER |
| SAM Successful Create Users/sec | Shows the number of users that were successfully created per second. | PERF_COUNTER_COUNTER |
| SAM Transitive Membership Evaluations/sec | Shows the number of SAM transitive membership evaluations per second. | PERF_COUNTER_COUNTER |
| SAM Universal Group Membership Evaluations/sec | Shows the number of SAM universal group membership evaluations per second. | PERF_COUNTER_COUNTER |
| SAM Writes | Shows the percentage of directory writes that come from SAM. | PERF_RAW_FRACTION |
| XDS Client Sessions | Shows the number of connected XDS client sessions. This indicates the number of connections from other Windows Server 2003 services and the Exchange Administrator program. | PERF_COUNTER_RAWCOUNT |
| XDS Reads | Shows the percentage of directory reads that come from XDS. | PERF_RAW_FRACTION |
| XDS Searches | Shows the percentage of directory searches that come from XDS. | PERF_RAW_FRACTION |
| XDS Writes | Shows the percentage of directory writes that come from XDS. | PERF_RAW_FRACTION |

# NWLink IPX Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The NWLink IPX performance object consists of counters that measure data transmission to and from computers that use the Internetwork Packet Exchange (IPX) protocol. This object reports valid data only for counters that monitor datagrams and packets; for frame-related counters, the object tends to report zero values. In addition, connection-related data is not likely to be useful, because IPX is a connectionless protocol.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Total/sec | Shows the total rate at which bytes are sent to or received from the network by the protocol, but shows only the bytes in frames that carry data. This is the sum of the values of NWLink IPX\Frame Bytes/sec and NWLink IPX\Datagram Bytes/sec. | PERF_COUNTER_BULK_COUNT |
| Connection Session Timeouts | Shows the number of connections that are dropped because a session times out. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Connections Canceled | Shows the number of connections that are canceled. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Connections No Retries | Shows the number of connections that are successfully made on the first try. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Connections Open | Shows the number of connections currently open for this protocol. This counter shows the current count only and is not cumulative. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Connections With Retries | Shows the number of connections that are made after attempting to reconnect. A retry occurs when the first connection attempt fails. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Datagram Bytes Received/sec | Shows the rate at which datagram bytes are received by the computer. | PERF_COUNTER_BULK_COUNT |
| Datagram Bytes Sent/sec | Shows the rate at which datagram bytes are sent from the computer. | PERF_COUNTER_BULK_COUNT |
| Datagram Bytes/sec | Shows the rate at which datagram bytes are processed by the computer. This counter is the sum of the values of NWLink IPX\Datagram Bytes Received/sec and NWLink IPX\Datagram Bytes Sent/sec. | PERF_COUNTER_BULK_COUNT |
| Datagrams Received/sec | Shows the rate at which datagrams are received by the computer. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Datagrams Sent/sec | Shows the rate at which datagrams are sent from the computer. | PERF_COUNTER_COUNTER |
| Datagrams/sec | Shows the rate at which datagrams are processed by the computer. This counter shows the sum of the values of NWLink IPX\Datagrams Received/sec and NWLink IPX\Datagrams Sent/sec. | PERF_COUNTER_COUNTER |
| Disconnects Local | Shows the number of session disconnections that are initiated by the local computer. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Disconnects Remote | Shows the number of session disconnections that are initiated by the remote computer. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Expirations Ack | Shows the number of T2 timer expirations. | PERF_COUNTER_RAWCOUNT |
| Expirations Response | Shows the number of T1 timer expirations. | PERF_COUNTER_RAWCOUNT |
| Failures Adapter | Shows the number of times a connection is dropped because an adapter fails. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Failures Link | Shows the number of connections that are dropped due to a link failure. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Failures No Listen | Shows the number of connections that are rejected because the remote computer is not listening for connection requests. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Failures Not Found | Shows the number of connection attempts that fail because the remote computer cannot be found. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Failures Resource Local | Shows the number of connections that fail because of resource problems or shortages on the local computer. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Failures Resource Remote | Shows the number of connections that fail because of resource problems or shortages on the remote computer. This number is cumulative and shows a running total. Note that connection-related data is not likely to be useful, because IPX is a connectionless protocol. | PERF_COUNTER_RAWCOUNT |
| Frame Bytes Received/sec | Shows the rate at which data bytes are received by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Rejected/sec | Shows the rate at which data bytes are rejected by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |

| | | |
|---|---|---|
| Frame Bytes Re-Sent/sec | Shows the rate at which data bytes are re-sent by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Sent/sec | Shows the rate at which data bytes are sent by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes/sec | Shows the rate at which data bytes are processed by the computer. This counter is the sum of Frame Bytes Sent/sec and Frame Bytes Received/sec. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frames Received/sec | Shows the rate at which data frames are received by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames Rejected/sec | Shows the rate at which data frames are rejected by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames Re-Sent/sec | Shows the rate at which data frames are re-sent by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames Sent/sec | Shows the rate at which data frames are sent by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames/sec | Shows the rate at which data frames are processed by the computer. This counter is the sum of the values of NWLink IPX\Frames Received/sec and NWLink IPX\Frames Sent/sec. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Packets Received/sec | Shows the rate at which packets are received by the computer. This counter shows all packets processed (control as well as data packets). | PERF_COUNTER_COUNTER |
| Packets Sent/sec | Shows the rate at which packets are sent by the computer. This counter shows all packets sent by the computer (control as well as data packets). | PERF_COUNTER_COUNTER |
| Packets/sec | Shows the rate at which packets are processed by the computer. This count is the sum of the values of NWLink IPX\Packets Sent/sec and NWLink IPX\Packets Received/sec. This counter includes all packets processed (control as well as data packets). | PERF_COUNTER_COUNTER |
| Piggyback Ack Queued/sec | Shows the rate at which "piggybacked" acknowledgments, acknowledgements that are added by the receiver onto the next outgoing data frame, are queued. | PERF_COUNTER_COUNTER |
| Piggyback Ack Timeouts | Shows the number of times that a "piggybacked" acknowledgment, an acknowledgement that is added by the receiver onto the next outgoing data frame, cannot be sent because there is no outgoing frame. If no outgoing frame is sent within the time-out period, an acknowledgment frame is sent, and this counter is incremented. | PERF_COUNTER_RAWCOUNT |
| Window Send Average | Shows the average number of data bytes that are sent before waiting for an acknowledgment from the remote computer. | PERF_COUNTER_RAWCOUNT |
| Window Send Maximum | Shows the maximum number of data bytes that are sent before waiting for an acknowledgment from the remote computer. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
NWLink NetBIOS Object
NWLink SPX Object

# NWLink NetBIOS Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The NWLink NetBIOS performance object consists of counters that monitor Internetwork Packet Exchange (IPX) transport rates and connections. This object reports valid data only for counters that monitor datagrams and packets; for frame-related counters, the object tends to report zero values.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Total/sec | Shows the total rate of bytes sent to or received from the network by the protocol, but shows only the bytes in frames that carry data. This is the sum of the values of NWLink NetBIOS\Frame Bytes/sec and NWLink NetBIOS\Datagram Bytes/sec. | PERF_COUNTER_BULK_COUNT |
| Connection Session Timeouts | Shows the number of connections that are dropped because a session times out. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Connections Canceled | Shows the number of connections that are canceled. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Connections No Retries | Shows the number of connections that are successfully made on the first try. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Connections Open | Shows the number of connections currently open for this protocol. This counter shows the current count only and does not accumulate over time. | PERF_COUNTER_RAWCOUNT |
| Connections With Retries | Shows the number of connections that are made after attempting to reconnect. A retry occurs when the first connection attempt fails. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Datagram Bytes Received/sec | Shows the rate at which datagram bytes are received by the computer. | PERF_COUNTER_BULK_COUNT |
| Datagram Bytes Sent/sec | Shows the rate at which datagram bytes are sent from the computer. | PERF_COUNTER_BULK_COUNT |
| Datagram Bytes/sec | Shows the rate at which datagram bytes are processed by the computer. This counter is the sum of the values of NWLink NetBIOS\Datagram Bytes Received/sec and NWLink NetBIOS\Datagram Bytes Sent/sec. | PERF_COUNTER_BULK_COUNT |
| Datagrams Received/sec | Shows the rate at which datagrams are received by the computer. | PERF_COUNTER_COUNTER |
| Datagrams Sent/sec | Shows the rate at which datagrams are sent from the computer. | PERF_COUNTER_COUNTER |
| Datagrams/sec | Shows the rate at which datagrams are processed by the computer. This counter shows the sum of the values of NWLink NetBIOS\Datagrams Received/sec and NWLink NetBIOS\Datagrams Sent/sec. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Disconnects Local | Shows the number of session disconnections that are initiated by the local computer. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Disconnects Remote | Shows the number of session disconnections that are initiated by the remote computer. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Expirations Ack | Shows the number of T2 timer expirations. | PERF_COUNTER_RAWCOUNT |
| Expirations Response | Shows the number of T1 timer expirations. | PERF_COUNTER_RAWCOUNT |
| Failures Adapter | Shows the number of connections that are dropped because an adapter fails. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures Link | Shows the number of connections that are dropped due to a link failure. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures No Listen | Shows the number of connections that are rejected because the remote computer is not listening for connection requests. | PERF_COUNTER_RAWCOUNT |
| Failures Not Found | Shows the number of connection attempts that fail because the remote computer cannot be found. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures Resource Local | Shows the number of connections that fail because of resource problems or shortages on the local computer. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures Resource Remote | Shows the number of connections that fail because of resource problems or shortages on the remote computer. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Frame Bytes Received/sec | Shows the rate at which data bytes are received by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Rejected/sec | Shows the rate at which data bytes are rejected by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Re-Sent/sec | Shows the rate at which data bytes are re-sent by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Sent/sec | Shows the rate at which data bytes are sent by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes/sec | Shows the rate at which data bytes are processed by the computer. This counter is the sum of the values of NWLink NetBIOS\Frame Bytes Received/sec and NWLink NetBIOS\Frame Bytes Sent/sec. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frames Received/sec | Shows the rate at which data frames are received by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames Rejected/sec | Shows the rate at which data frames are rejected by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Frames Re-Sent/sec | Shows the rate at which data frames are re-sent by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames Sent/sec | Shows the rate at which data frames are sent by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames/sec | Shows the rate at which data frames are processed by the computer. This counter is the sum of the values of NWLink NetBIOS\Frames Received/sec and NWLink NetBIOS\Frames Sent/sec. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Packets Received/sec | Shows the rate at which packets are received by the computer. This counter shows all packets processed (control as well as data packets). | PERF_COUNTER_COUNTER |
| Packets Sent/sec | Shows the rate at which packets are sent by the computer. This counter shows all packets sent by the computer (control as well as data packets). | PERF_COUNTER_COUNTER |
| Packets/sec | Shows the rate at which packets are processed by the computer. This counter is the sum of the values of NWLink NetBIOS\Packets Received/sec and NWLink NetBIOS\Packets Sent/sec. This counter includes all packets processed (control as well as data packets). | PERF_COUNTER_COUNTER |
| Piggyback Ack Queued/sec | Shows the rate at which "piggybacked" acknowledgments, acknowledgements that are added by the receiver onto the next outgoing data frame, are queued. | PERF_COUNTER_COUNTER |
| Piggyback Ack Timeouts | Shows the number of times that a "piggybacked" acknowledgment, an acknowledgement that is added by the receiver onto the next outgoing data frame, cannot be sent because there is no outgoing frame. If no outgoing frame is sent within the time-out period, an acknowledgment frame is sent, and this counter is incremented. | PERF_COUNTER_RAWCOUNT |
| Window Send Average | Shows the average number of data bytes that are sent before waiting for an acknowledgment from the remote computer. | PERF_COUNTER_RAWCOUNT |
| Window Send Maximum | Shows the maximum number of data bytes that are sent before waiting for an acknowledgment from the remote computer. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
NWLink IPX Object
NWLink SPX Object

# NWLink SPX Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The NWLink SPX performance object consists of counters that measure data transmission and session connections for computers that use the Sequenced Packet Exchange (SPX) protocol. This object reports valid data only for counters that monitor datagrams and packets; for frame-related counters, the object tends to report zero values.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Total/sec | Shows the total rate of bytes sent to or received from the network by the protocol, but shows only the bytes in frames that carry data. This is the sum of the values of Frame Bytes/sec and Datagram Bytes/sec. | PERF_COUNTER_BULK_COUNT |
| Connection Session Timeouts | Shows the number of connections that are dropped because a session times out. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Connections Canceled | Shows the number of connections that are canceled. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Connections No Retries | Shows the number of connections that are successfully made on the first try. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Connections Open | Shows the number of connections currently open for this protocol. This counter shows the current count only and is not cumulative. | PERF_COUNTER_RAWCOUNT |
| Connections With Retries | Shows the number of connections that are made after attempting to reconnect. A retry occurs when the first connection attempt fails. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Datagram Bytes Received/sec | Shows the rate at which datagram bytes are received by the computer. | PERF_COUNTER_BULK_COUNT |
| Datagram Bytes Sent/sec | Shows the rate at which datagram bytes are sent from the computer. | PERF_COUNTER_BULK_COUNT |
| Datagram Bytes/sec | Shows the rate at which datagram bytes are processed by the computer. This counter is the sum of the values of Datagram Bytes Received/sec and Datagram Bytes Sent/sec. | PERF_COUNTER_BULK_COUNT |
| Datagrams Received/sec | Shows the rate at which datagrams are received by the computer. | PERF_COUNTER_COUNTER |
| Datagrams Sent/sec | Shows the rate at which datagrams are sent from the computer. | PERF_COUNTER_COUNTER |
| Datagrams/sec | Shows the rate at which datagrams are processed by the computer. This counter is the sum of the values of Datagrams Received/sec and Datagrams Sent/sec. | PERF_COUNTER_COUNTER |
| Disconnects Local | Shows the number of session disconnections that are initiated by the local computer. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Disconnects Remote | Shows the number of session disconnections that are initiated by the remote computer. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Expirations Ack | Shows the number of T2 timer expirations | PERF_COUNTER_RAWCOUNT |
| Expirations Response | Shows the number of T1 timer expirations. | PERF_COUNTER_RAWCOUNT |
| Failures Adapter | Shows the number of connections that are dropped because an adapter fails. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures Link | Shows the number of connections that are dropped due to a link failure. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures No Listen | Shows the number of connections that are rejected because the remote computer is not listening for connection requests. | PERF_COUNTER_RAWCOUNT |
| Failures Not Found | Shows the number of connection attempts that fail because the remote computer cannot be found. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures Resource Local | Shows the number of connections that fail because of resource problems or shortages on the local computer. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures Resource Remote | Shows the number of connections that fail because of resource problems or shortages on the remote computer. This number is cumulative and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Frame Bytes Received/sec | Shows the rate at which data bytes are received by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Rejected/sec | Shows the rate at which data bytes are rejected by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Re-Sent/sec | Shows the rate at which data bytes are re-sent by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Sent/sec | Shows the rate at which data bytes are sent by the computer. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes/sec | Shows the rate at which data bytes are processed by the computer. This counter is the sum of the values of Frame Bytes Received/sec and Frame Bytes Sent/sec. This counter shows only the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frames Received/sec | Shows the rate at which data frames are received by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames Rejected/sec | Shows the rate at which data frames are rejected by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames Re-Sent/sec | Shows the rate at which data frames are re-sent by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Frames Sent/sec | Shows the rate at which data frames are sent by the computer. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames/sec | Shows the rate at which data frames are processed by the computer. This counter is the sum of the values of Frames Received/sec and Frames Sent/sec. This counter shows only the frames that carry data. | PERF_COUNTER_COUNTER |
| Packets Received/sec | Shows the rate at which packets are received by the computer. This counter shows all packets processed (control as well as data packets). | PERF_COUNTER_COUNTER |
| Packets Sent/sec | Shows the rate at which packets are sent by the computer. This counter shows all packets sent by the computer (control as well as data packets). | PERF_COUNTER_COUNTER |
| Packets/sec | Shows the rate at which packets are processed by the computer. This counter is the sum of the values of Packets Sent/Sec and Packets Received/Sec. This counter includes all packets processed (control as well as data packets). NWLink SPX\Packets/sec is visible in IPX and NetBIOS. | PERF_COUNTER_COUNTER |
| Piggyback Ack Queued/sec | Shows the rate at which "piggybacked" acknowledgments, acknowledgements that are added by the receiver onto the next outgoing data frame, are queued. | PERF_COUNTER_COUNTER |
| Piggyback Ack Timeouts | Shows the number of times that a "piggybacked" acknowledgment, an acknowledgement that is added by the receiver onto the next outgoing data frame, cannot be sent because there is no outgoing frame. If no outgoing frame is sent within the time-out period, an acknowledgment frame is sent, and this counter is incremented. | PERF_COUNTER_RAWCOUNT |
| Window Send Average | Shows the running average number of data bytes that are sent before waiting for an acknowledgment from the remote computer. | PERF_COUNTER_RAWCOUNT |
| Window Send Maximum | Shows the maximum number of data bytes that are sent before waiting for an acknowledgment from the remote computer. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
NWLink NetBIOS Object
NWLink IPX Object

# Objects Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Objects performance object consists of counters that monitor logical objects in the system, such as processes, threads, mutexes, and semaphores. This information can be used to detect the unnecessary consumption of computer resources. Each object requires memory to store basic information about the object.

| Counter Name | Description | Counter Type |
|---|---|---|
| Events | Shows the number of events in the computer at the time of data collection. An event occurs when two or more threads try to synchronize execution. This is an instantaneous count, not an average over time. | PERF_COUNTER_RAWCOUNT |
| Mutexes | Shows the number of mutexes in the computer at the time of data collection. Mutexes are executive dispatcher objects that ensure threads are synchronized. Mutexes are used by threads to assure that only one thread is executing a particular section of code. This is an instantaneous count, not an average over time. | PERF_COUNTER_RAWCOUNT |
| Processes | Shows the number of processes in the computer at the time of data collection. Each process represents the running of a program. This is an instantaneous count, not an average over time. | PERF_COUNTER_RAWCOUNT |
| Sections | Shows the number of sections in the computer at the time of data collection. A section is a portion of virtual memory created by a process for storing data. A process can share sections with other processes. This is an instantaneous count, not an average over time. | PERF_COUNTER_RAWCOUNT |
| Semaphores | Shows the number of semaphores in the computer at the time of data collection. Threads use semaphores to obtain exclusive access to data structures that they share with other threads. This is an instantaneous count, not an average over time. | PERF_COUNTER_RAWCOUNT |
| Threads | Shows the number of threads in the computer at the time of data collection. A thread is the basic executable entity that can execute instructions in a processor. This is an instantaneous count, not an average over time. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Thread Object
Process Object
Image Object
Process Address Space Object

# Paging File Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Paging File performance object consists of counters that monitor the paging file(s) on the computer. The paging file is a reserved space on disk that backs up committed physical memory on the computer.

| Counter Name | Description | Counter Type |
|---|---|---|
| Usage Peak | Shows the amount of the page file instance in use. | PERF_RAW_FRACTION |
| % Usage Peak | Shows the greatest percentage of the paging file that was in use during the sample interval. | PERF_100NSEC_TIMER |

**See Also**
**Concepts**
Process Object

# PBServer Monitor Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The PBServer Monitor performance object monitors activity on the phone book server. Phone Book Service is a comprehensive phone book management system, and is a component of Connection Point Services. Phone Book Service provides a central location for uploading and distributing access numbers for both public and private networks. A few of the counters, such as Error Hits, can be used to detect errors.

These counter objects do not show up in System Monitor until the service has been started. The service starts when a phone book is requested.

| Counter Name | Description | Counter Type |
|---|---|---|
| Delta Upgrade Hits | Shows the total number of updated (delta) phone book files that have been successfully downloaded to the Connection Manager client. | PERF_COUNTER_RAWCOUNT |
| Delta Upgrade Hits/sec | Shows the rate, in incidents per second, at which updated (delta) phone book files were successfully downloaded to the Connection Manager client. | PERF_COUNTER_COUNTER |
| Error Hits | Shows the total number of requests that could not be satisfied by the Phone Book Service. This could be due to a problem with the client or due to corruption in the service database. For example, the client is looking for the wrong phone book. To investigate the cause, review the Web Service log for Windows Server 2003 family verbose logging. Connection Manager requests appear in the log as Get requests to the Pbserver.dll. Also check for the PB= value in the log; this will be name of phone book. Research this name; try to find the name in the database (Pbserver.mdb). It may be necessary to repost the phone book using administrative tools. | PERF_COUNTER_RAWCOUNT |
| Error Hits/sec | Shows the rate, in incidents per second, at which requests could not be satisfied by the Phone Book Service. | PERF_COUNTER_COUNTER |
| Full Upgrade Hits | Shows the total number of full phone book files that have been successfully downloaded. | PERF_COUNTER_RAWCOUNT |
| Full Upgrade Hits/sec | Shows the rate, in incidents per second, at which full phone book files were successfully downloaded. | PERF_COUNTER_COUNTER |
| No Upgrade Hits | Shows the total number of requests for phone book files that did not require an update to the phone book file on the user's computer. | PERF_COUNTER_RAWCOUNT |
| No Upgrade Hits/sec | Shows the rate, in incidents per second, at which requests for phone book files did not require an update to the phone book file on the user's computer. | PERF_COUNTER_COUNTER |
| Total Hits | Shows the total number of requests for phone book files since Web services were last started. If this value seems low in comparison with the amount of user traffic you expect, this could indicate a problem. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Total Hits/sec | Shows the rate, in incidents per second, of all requests for phone book files. This counter is useful for determining the load on the Phone Book Service. | PERF_COUNTER_COUNTER |

# Physical Disk Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The PhysicalDisk performance object consists of counters that monitor hard or fixed disk drives. Disks are used to store file, program, and paging data. They are read to retrieve these items, and are written to record changes to them. The values of physical disk counters are sums of the values of the logical disks (or partitions) into which they are divided.

| Counter Name | Description | Counter Type |
|---|---|---|
| % Disk Read Time | Shows the percentage of time that the selected disk drive was busy servicing read requests. | PERF_PRECISION_100NS_TIMER |
| % Disk Time | Shows the percentage of elapsed time that the selected disk drive was busy servicing read or write requests. | PERF_PRECISION_100NS_TIMER |
| % Disk Write Time | Shows the percentage of elapsed time that the selected disk drive was busy servicing write requests. | PERF_PRECISION_100NS_TIMER |
| % Idle Time | Shows the percentage of elapsed time during the sample interval that the selected disk drive was idle. | PERF_PRECISION_100NS_TIMER |
| Avg. Disk Bytes/Read | Shows the average number of bytes that were transferred from the disk during read operations. | PERF_AVERAGE_BULK |
| Avg. Disk Bytes/Transfer | Shows the average number of bytes that were transferred to or from the disk during write or read operations. | PERF_AVERAGE_BULK |
| Avg. Disk Bytes/Write | Shows the average number of bytes that were transferred to the disk during write operations. | PERF_AVERAGE_BULK |
| Avg. Disk Queue Length | Shows the average number of both read and write requests that were queued for the selected disk during the sample interval. | PERF_COUNTER_LARGE_QUEUELEN_TYPE |
| Avg. Disk Read Queue Length | Shows the average number of read requests that were queued for the selected disk during the sample interval. | PERF_COUNTER_LARGE_QUEUELEN_TYPE |
| Avg. Disk sec/Read | Shows the average time, in seconds, of a read of data from the disk. | PERF_AVERAGE_TIMER |
| Avg. Disk sec/Transfer | Shows the average time, in seconds, of a disk transfer. | PERF_AVERAGE_TIMER |
| Avg. Disk sec/Write | Shows the average time, in seconds, of a write of data to the disk. | PERF_AVERAGE_TIMER |
| Avg. Disk Write Queue Length | Shows the average number of write requests that were queued for the selected disk during the sample interval. | PERF_COUNTER_LARGE_QUEUELEN_TYPE |

| | | |
|---|---|---|
| Current Disk Queue Length | Shows the number of requests that were outstanding on the disk at the time that the performance data was collected. This is a snapshot, not an average over the time interval. It includes requests in service at the time of the collection. Multispindle disk devices can have multiple requests active at one time, but other concurrent requests are awaiting service. This counter might reflect a transitory high or low queue length, but if this counter is consistently high, then it is likely that there is a sustained load on the disk drive. Requests experience delays proportional to the length of this queue, minus the number of spindles on the disks. This difference should average less than two. | PERF_COUNTER_RAWCOUNT |
| Disk Bytes/sec | Shows the rate, in incidents per second, at which bytes were transferred to or from the disk during write or read operations. | PERF_COUNTER_BULK_COUNT |
| Disk Read Bytes/sec | Shows the rate, in incidents per second, at which bytes were transferred from the disk during read operations. | PERF_COUNTER_BULK_COUNT |
| Disk Reads/sec | Shows the rate, in incidents per second, at which read operations were performed on the disk. | PERF_COUNTER_COUNTER |
| Disk Transfers/sec | Shows the rate, in incidents per second, at which read and write operations were performed on the disk. | PERF_COUNTER_COUNTER |
| Disk Write Bytes/sec | Shows the rate, in incidents per second, at which bytes were transferred to the disk during write operations. | PERF_COUNTER_BULK_COUNT |
| Disk Writes/sec | Shows the rate, in incidents per second, at which write operations were performed on the disk. | PERF_COUNTER_COUNTER |
| Split IO/sec | Shows the rate, in incidents per second, at which input/output (I/O) requests to the disk were split into multiple requests. A split I/O might result from requesting data in a size that is too large to fit into a single I/O, or from a fragmented disk subsystem. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Find disk bottlenecks

# Print Queue Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Print Queue performance object is provided by Windows Server 2003 family printing features to track activity in the server print queue. The Print Queue counters are reset when either the print server or the spooler service is restarted.

| Counter Name | Description | Counter Type |
|---|---|---|
| Add Network Printer Calls | Shows the total number of calls from other print servers to add shared network printers to this print server since it was last started. | PERF_COUNTER_RAWCOUNT |
| Bytes Printed/sec | Shows the rate, in incidents per second, at which bytes were printed in a print queue. This is a rough indication of how busy a printer is. This counter can be used for bottleneck detection. | PERF_COUNTER_BULK_COUNT |
| Enumerate Network Printer Calls | Shows the total number of calls from browse clients to this print server to request network browse lists since the print server was last started. | PERF_COUNTER_RAWCOUNT |
| Job Errors | Shows the total number of job errors in a print queue since the print server was last started. Job errors can indicate network or connectivity problems with the print device. | PERF_COUNTER_RAWCOUNT |
| Jobs | Shows the current number of jobs in a print queue. This counter can be used to identify excessive use. | PERF_COUNTER_RAWCOUNT |
| Jobs Spooling | Shows the current number of spooling jobs in a print queue. | PERF_COUNTER_RAWCOUNT |
| Max Jobs Spooling | Shows the maximum number of spooling jobs in a print queue since the print server was last started. | PERF_COUNTER_RAWCOUNT |
| Max References | Shows the peak number of references (open handles) to this printer. | PERF_COUNTER_RAWCOUNT |
| Not Ready Errors | Shows the total number of printer not ready errors in a print queue since the print server was last started. | PERF_COUNTER_RAWCOUNT |
| Out of Paper Errors | Shows the total number of out of paper errors in a print queue since the print server was last started. | PERF_COUNTER_RAWCOUNT |
| References | Shows the current number of references (open handles) to a print queue. A reference can be a user or a program that is connecting to a printer and opening a print queue. | PERF_COUNTER_RAWCOUNT |
| Total Jobs Printed | Shows the total number of jobs printed on a print queue since the print server was last started. | PERF_COUNTER_RAWCOUNT |
| Total Pages Printed | Shows the total number of pages printed through Graphics Device Interface (GDI) on a print queue since the print server was last started. | PERF_COUNTER_RAWCOUNT |

# Process Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Process performance object consists of counters that monitor running application program and system processes. All the threads in a process share the same address space and have access to the same data.

| Counter Name | Description | Counter Type |
|---|---|---|
| % Privileged Time | Shows the percentage of non-idle processor time spent executing code in privileged mode. Privileged mode is a processing mode designed for operating system components and hardware-manipulating drivers. It allows direct access to hardware and memory. The alternative, user mode, is a restricted processing mode designed for applications, environmental subsystems, and integral subsystems. The operating system switches application threads to privileged mode to access operating system services.<br><br>% Privileged Time includes time servicing interrupts and DPCs. A high rate of privileged time might be attributed to a large number of interrupts generated by a failing device. When a Windows system service is called, the service often runs in privileged mode to gain access to system-private data. Such data is protected from access by threads executing in user mode. Calls to the system can be explicit or implicit, such as page faults or interrupts.<br><br>Unlike some earlier Windows operating systems, Windows 2000 and the Windows Server 2003 family use process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Windows 2000 and Windows Server 2003 on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process. | PERF_100NSEC_TIMER |
| % Processor Time | Shows the percentage of time that the processor spent executing a non-idle thread. It is calculated by measuring the duration that the idle thread is active during the sample interval, and subtracting that time from 100 %. (Each processor has an idle thread that consumes cycles when no other threads are ready to run.) This counter is the primary indicator of processor activity, and displays the average percentage of busy time observed during the sample interval. Code executed to handle some hardware interrupts and trap conditions are included in this count. | PERF_100NSEC_TIMER |
| % User Time | Shows the percentage of time that the processor spent executing code in user mode. Applications, environment subsystems, and integral subsystems execute in user mode. Code executing in user mode cannot damage the integrity of the Windows Executive, kernel, and/or device drivers. Unlike some earlier Windows operating systems, Windows 2000 and Windows Server 2003 use process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Windows 2000 and Windows Server 2003 on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process. | PERF_100NSEC_TIMER |

| | | |
|---|---|---|
| Page Faults/sec | Shows the rate, in incidents per second, at which page faults were handled by the processor. A page fault occurs when a process requires code or data that is not in its working set. This counter includes both hard faults (those that require disk access) and soft faults (those where the faulted page is found elsewhere in physical memory). Most processors can handle large numbers of soft faults without consequence. However, hard faults can cause significant delays. | PERF_COUNTER_COUNTER |
| Page File Bytes | Shows the the current amount of virtual memory, in bytes, that a process has reserved for use in the paging file(s). Paging files are used to store pages of memory used by the process. Paging files are shared by all processes, and the lack of space in paging files can prevent other processes from allocating memory. If there is no paging file, this counter reflects the current amount of virtual memory that the process has reserved for use in physical memory. | PERF_COUNTER_RAWCOUNT |
| Page File Bytes Peak | Shows the maximum amount of virtual memory, in bytes, that a process has reserved for use in the paging file(s). Paging files are used to store pages of memory used by the process. Paging files are shared by all processes, and the lack of space in paging files can prevent other processes from allocating memory. If there is no paging file, this counter reflects the maximum amount of virtual memory that the process has reserved for use in physical memory. | PERF_COUNTER_RAWCOUNT |
| Priority Base | Shows the current base priority of this process. Threads within a process can raise and lower their own base priority relative to the process's base priority. | PERF_COUNTER_RAWCOUNT |
| Private Bytes | Shows the size, in bytes, that this process has allocated that cannot be shared with other processes. | PERF_COUNTER_RAWCOUNT |
| Working Set | Shows the size, in bytes, in the working set of this process. The working set is the set of memory pages that were touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the working set of a process, even if they are not in use. When free memory falls below a threshold, pages are trimmed from working sets. If the pages are needed, they will be soft-faulted back into the working set before leaving main memory. | PERF_COUNTER_RAWCOUNT |
| Creating Process ID | Shows the identifier of the process that created the current process. Note that the creating process might have terminated since the current process was created, and so the current Creating Process ID counter may no longer identify a running process. | PERF_COUNTER_RAWCOUNT |
| Elapsed Time | Shows the time, in seconds, that this process has been running. | PERF_ELAPSED_TIME |
| Handle Count | Shows the total number of handles currently open by this process. This number is the equal to the sum of the handles currently open by each thread in this process. | PERF_COUNTER_RAWCOUNT |
| ID Process | Shows the unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process. | PERF_COUNTER_RAWCOUNT |
| IO Data Bytes/sec | Shows the rate, in incidents per second, at which the process was reading and writing bytes in I/O operations. It counts all I/O activity generated by the process including file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Data Operations/sec | Shows the rate, in incidents per second, at which the process was issuing read and write I/O operations. It counts all I/O activity generated by the process including file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |

| | | |
|---|---|---|
| IO Other Bytes/sec | Shows the rate, in incidents per second, at which the process was issuing bytes to I/O operations that do not involve data such as control operations. It counts all I/O activity generated by the process including file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Other Operations/sec | Shows the rate, in incidents per second, at which the process was issuing I/O operations that were neither read nor write operations (for example, a control function). It counts all I/O activity generated by the process including file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Read Bytes/sec | Shows the rate, in incidents per second, at which the process was reading bytes from I/O operations. It counts all I/O activity generated by the process including file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Read Operations/sec | Shows the rate, in incidents per second, at which the process was issuing read I/O operations. It counts all I/O activity generated by the process including file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Write Bytes/sec | Shows the rate, in incidents per second, at which the process was writing bytes to I/O operations. It counts all I/O activity generated by the process including file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| IO Write Operations/sec | Shows the rate, in incidents per second, at which the process was issuing write I/O operations. It counts all I/O activity generated by the process including file, network, and device I/Os. | PERF_COUNTER_BULK_COUNT |
| Pool Nonpaged Bytes | Shows the number of bytes in the nonpaged pool, an area of system memory (physical memory used by the operating system) for objects that cannot be written to disk. Nonpaged pool pages cannot be paged out to the paging file, but instead remain in main memory as long as they are allocated. | PERF_COUNTER_RAWCOUNT |
| Pool Paged Bytes | Shows the number of bytes in the paged pool, an area of system memory (physical memory used by the operating system) for objects that can be written to disk. Paged pool pages can be paged out to the paging file when not accessed by the system for sustained periods of time. Memory\Pool Paged Bytes and Process\Pool Paged Bytes are calculated differently than this counter, so the counters might not display the same amounts. | PERF_COUNTER_RAWCOUNT |
| Thread Count | Shows the number of threads that were active in this process. A thread is the object that executes instructions, which are the basic unit of execution in a processor. Every running process has at least one thread. | PERF_COUNTER_RAWCOUNT |
| Virtual Bytes | Shows the size, in bytes, of the virtual address space that the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and by using too much, the process can limit its ability to load libraries. | PERF_COUNTER_RAWCOUNT |
| Virtual Bytes Peak | Shows the maximum size, in bytes, of virtual address space that the process has used at any one time. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. However, virtual space is finite, and the process might limit its ability to load libraries by using too much. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Working Set Peak | Shows the maximum size, in bytes, in the working set of this process. The working set is the set of memory pages that were touched recently by the threads in the process. If free memory in the computer is above a certain threshold, pages are left in the working set of a process, even if they are not in use. When free memory falls below a certain threshold, pages are trimmed from working sets. If the pages are needed, they will be soft-faulted back into the working set before leaving main memory. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Find processor bottlenecks
Find memory bottlenecks
Finding Bottlenecks on Multiprocessor Computers

# Process Address Space Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Process Address Space performance object consists of counters that monitor memory allocation and use for a selected process.

This object contains "costly" counters for which data collection is typically either processor-intensive or memory-intensive. Costly counters cannot be selected with System Monitor but can be accessed by using Showperf.exe and other tools included in the Windows Resource Kit Tools Help for Windows Server 2003.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Free | Shows the total unused virtual address space of this process. | PERF_COUNTER_RAWCOUNT |
| Bytes Image Free | Shows the amount of virtual address space that is not in use and is not reserved by images that run within this process. | PERF_COUNTER_RAWCOUNT |
| Bytes Image Reserved | Shows the sum of all virtual memory reserved by images that run within this process. | PERF_COUNTER_RAWCOUNT |
| Bytes Reserved | Shows the total amount of virtual memory reserved for future use by this process. | PERF_COUNTER_RAWCOUNT |
| ID Process | Shows the unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process. | PERF_COUNTER_RAWCOUNT |
| Image Space Exec Read Only | Shows the virtual address space in use by the images that are being executed by the process. This counter is the sum of all the address space with this protection allocated by images that run within the selected process. *Execute read-only* memory can be executed by programs, in addition to being read. | PERF_COUNTER_RAWCOUNT |
| Image Space Exec Read/Write | Shows the virtual address space in use by the images that are being executed by the process. This counter is the sum of all the address space with this protection allocated by images that run within the selected process. *Execute read/write* memory can be executed by programs, in addition to being read and written and modified. | PERF_COUNTER_RAWCOUNT |
| Image Space Exec Write Copy | Shows the virtual address space in use by the images that are being executed by the process. This counter is the sum of all the address space with this protection allocated by images that run within the selected process. *Execute write copy* memory can be executed by programs, in addition to being read and written. This type of protection is used when memory needs to be shared between processes. If the sharing processes only read the memory, they all use the same memory. If a sharing process wants to have write access to this shared memory, a copy of this memory is made for that process. | PERF_COUNTER_RAWCOUNT |
| Image Space Executable | Shows the virtual address space in use by the images that are being executed by the process. This counter is the sum of all the address space with this protection allocated by images that run within the selected process. *Executable* memory can be executed by programs, but cannot be read or written. This type of protection is not supported by all processor types. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Image Space No Access | Shows the virtual address space in use by the images that are being executed by the process. This counter is the sum of all the address space with this protection allocated by images that run within the selected process. *No access* protection prevents a process from writing to or reading from these pages, and an access violation is generated if either is attempted. | PERF_COUNTER_RAWCOUNT |
| Image Space Read Only | Shows the virtual address space in use by the images that are being executed by the process. This counter is the sum of all the address space with this protection allocated by images that run within the selected process. *Read-only* protection prevents the contents of these pages from being modified. Any attempt to write or modify these pages generates an access violation. | PERF_COUNTER_RAWCOUNT |
| Image Space Read/Write | Shows the virtual address space in use by the images that are being executed by the process. This counter is the sum of all the address space with this protection allocated by images that run within the selected process. *Read/write* protection allows a process to read, modify, and write to these pages. | PERF_COUNTER_RAWCOUNT |
| Image Space Write Copy | Shows the virtual address space in use by the images that are being executed by the process. This counter is the sum of all the address space with this protection allocated by images that run within the selected process. *Write copy* protection is used when memory is shared for reading but not for writing. If the sharing processes only read the memory, they all use the same memory. If a sharing process wants to have read/write access to this shared memory, a copy of this memory is made for that process. | PERF_COUNTER_RAWCOUNT |
| Mapped Space Exec Read Only | Shows virtual memory that has been mapped to a specific virtual address (or range of virtual addresses) in the virtual address space of the process. | PERF_COUNTER_RAWCOUNT |
| Mapped Space Exec Read/Write | Shows virtual memory that has been mapped to a specific virtual address (or range of virtual addresses) in the virtual address space of the process. | PERF_COUNTER_RAWCOUNT |
| Mapped Space Exec Write Copy | Shows virtual memory that has been mapped to a specific virtual address (or range of virtual addresses) in the virtual address space of the process. | PERF_COUNTER_RAWCOUNT |
| Mapped Space Executable | Shows virtual memory that has been mapped to a specific virtual address (or range of virtual addresses) in the virtual address space of the process. | PERF_COUNTER_RAWCOUNT |
| Mapped Space No Access | Shows virtual memory that has been mapped to a specific virtual address (or range of virtual addresses) in the virtual address space of the process. | PERF_COUNTER_RAWCOUNT |
| Mapped Space Read Only | Shows virtual memory that has been mapped to a specific virtual address (or range of virtual addresses) in the virtual address space of the process. | PERF_COUNTER_RAWCOUNT |
| Mapped Space Read/Write | Shows virtual memory that has been mapped to a specific virtual address (or range of virtual addresses) in the virtual address space of the process. | PERF_COUNTER_RAWCOUNT |
| Mapped Space Write Copy | Shows virtual memory that has been mapped to a specific virtual address (or range of virtual addresses) in the virtual address space of the process. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Reserved Space Exec Read Only | Shows virtual memory that has been reserved for future use by a process, but has not been mapped or committed. | PERF_COUNTER_RAWCOUNT |
| Reserved Space Exec Read/Write | Shows virtual memory that has been reserved for future use by a process, but has not been mapped or committed. | PERF_COUNTER_RAWCOUNT |
| Reserved Space Exec Write Copy | Shows virtual memory that has been reserved for future use by a process, but has not been mapped or committed. | PERF_COUNTER_RAWCOUNT |
| Reserved Space Executable | Shows virtual memory that has been reserved for future use by a process, but has not been mapped or committed. | PERF_COUNTER_RAWCOUNT |
| Reserved Space No Access | Shows virtual memory that has been reserved for future use by a process, but has not been mapped or committed. | PERF_COUNTER_RAWCOUNT |
| Reserved Space Read Only | Shows virtual memory that has been reserved for future use by a process, but has not been mapped or committed. | PERF_COUNTER_RAWCOUNT |
| Reserved Space Read/Write | Shows virtual memory that has been reserved for future use by a process, but has not been mapped or committed. | PERF_COUNTER_RAWCOUNT |
| Reserved Space Write Copy | Shows virtual memory that has been reserved for future use by a process, but has not been mapped or committed. | PERF_COUNTER_RAWCOUNT |
| Unassigned Space Exec Read Only | Shows mapped and committed virtual memory in use by the process that is not attributable to any particular image being executed by the process. | PERF_COUNTER_RAWCOUNT |
| Unassigned Space Exec Read/Write | Shows mapped and committed virtual memory in use by the process that is not attributable to any particular image being executed by the process. | PERF_COUNTER_RAWCOUNT |
| Unassigned Space Exec Write Copy | Shows mapped and committed virtual memory in use by the process that is not attributable to any particular image being executed by the process. | PERF_COUNTER_RAWCOUNT |
| Unassigned Space Executable | Shows mapped and committed virtual memory in use by the process that is not attributable to any particular image being executed by the process. | PERF_COUNTER_RAWCOUNT |
| Unassigned Space No Access | Shows mapped and committed virtual memory in use by the process that is not attributable to any particular image being executed by the process. | PERF_COUNTER_RAWCOUNT |
| Unassigned Space Read Only | Shows mapped and committed virtual memory in use by the process that is not attributable to any particular image being executed by the process. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Unassigned Space Read/Write | Shows mapped and committed virtual memory in use by the process that is not attributable to any particular image being executed by the process. | PERF_COUNTER_RAWCOUNT |
| Unassigned Space Write Copy | Shows mapped and committed virtual memory in use by the process that is not attributable to any particular image being executed by the process. | PERF_COUNTER_RAWCOUNT |

**See Also**

**Concepts**

Image Object
Thread Object
Thread Details Object
Process Object

# Processor Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Processor performance object consists of counters that measure aspects of processor activity. The processor is the part of the computer that performs arithmetic and logical computations, initiates operations on peripherals, and runs the threads of processes. A computer can have multiple processors. The Processor object represents each processor as an instance of the object.

| Counter Name | Description | Counter Type |
|---|---|---|
| % C1 Time | Shows the percentage of time that the processor spent in the C1 low-power idle state. % C1 Time is a subset of the total processor idle time. C1 low-power idle state enables the processor to maintain its entire context and quickly return to the running state. Not all systems support the C1 state. | PERF_100NSEC_TIMER |
| % C2 Time | Shows the percentage of time that the processor spent in the C2 low-power idle state. % C2 Time is a subset of the total processor idle time. C2 low-power idle state enables the processor to maintain the context of the system caches. The C2 power state is a lower power and higher exit latency state than C1. Not all systems support the C2 state. | PERF_100NSEC_TIMER |
| % C3 Time | Shows the percentage of time that the processor spent in the C3 low-power idle state. % C3 Time is a subset of the total processor idle time. When the processor is in the C3 low-power idle state it is unable to maintain the coherency of its caches. The C3 power state is a lower power and higher exit latency state than C2. Not all systems support the C3 state. | PERF_100NSEC_TIMER |
| % DPC Time | Shows the percentage of time that the processor spent receiving and servicing deferred procedure calls (DPCs) during the sample interval. DPCs are interrupts that run at a lower priority than standard interrupts. % DPC Time is a component of % Privileged Time because DPCs are executed in privileged mode. They are counted separately and are not a component of the interrupt counters. This counter displays the average busy time as a percentage of the sample time. | PERF_100NSEC_TIMER |
| % Idle Time | Shows the time that the processor was idle during the sample interval. | PERF_100NSEC_TIMER |
| % Interrupt Time | Shows the percentage of time that the processor spent receiving and servicing hardware interrupts during the sample interval. This value is an indirect indicator of the activity of devices that generate interrupts, such as the system clock, the mouse, disk drivers, data communication lines, network interface cards and other peripheral devices. These devices normally interrupt the processor when they have completed a task or require attention. Normal thread execution is suspended during interrupts. Most system clocks interrupt the processor every 10 milliseconds, creating a background of interrupt activity. This counter displays the average busy time as a percentage of the sample time. | PERF_100NSEC_TIMER |

| | | |
|---|---|---|
| % Privileged Time | Shows the percentage of elapsed time that this thread spent executing code in privileged mode. When a Windows Server 2003 family operating system service is called, the service often runs in privileged mode in order to gain access to system-private data. Such data is protected from access by threads executing in user mode. Calls to the system can be explicit or implicit, such as page faults and interrupts. Unlike some early operating systems, Windows Server 2003 uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Windows Server 2003 on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process. | PERF_100NSEC_TIMER |
| % Processor Time | Shows the percentage of elapsed time that this thread used the processor to execute instructions. An instruction is the basic unit of execution in a processor, and a thread is the object that executes instructions. Code executed to handle some hardware interrupts and trap conditions is included in this count. | PERF_100NSEC_TIMER_INV |
| % User Time | Shows the percentage of elapsed time that this thread spent executing code in user mode. Applications, environment subsystems, and integral subsystems execute in user mode. Code executing in user mode cannot damage the integrity of the Windows Server 2003 Executive, Kernel, and device drivers. Unlike some early operating systems, Windows Server 2003 uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Windows Server 2003 on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process. | PERF_100NSEC_TIMER |
| C1 Transitions/sec | Shows the rate, in incidents per second, at which the CPU entered the C1 low-power idle state. The CPU enters the C1 state when it is sufficiently idle, and exits this state on any interrupt. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval. | PERF_COUNTER_BULK_COUNT |
| C2 Transitions/sec | Shows the rate, in incidents per second, at which the CPU entered the C2 low-power idle state. The CPU enters the C2 state when it is sufficiently idle, and exits this state on any interrupt. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval. | PERF_COUNTER_BULK_COUNT |
| C3 Transitions/sec | Shows the rate, in incidents per second, at which the CPU entered the C3 low-power idle state. The CPU enters the C3 state when it is sufficiently idle, and exits this state on any interrupt. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval. | PERF_COUNTER_BULK_COUNT |
| DPC Rate | Shows the rate at which deferred procedure calls (DPCs) were added to the processor's DPC queue between the timer ticks of the processor clock. DPCs are interrupts that run at a lower priority than standard interrupts. Each processor has its own DPC queue. This counter measures the rate at which DPCs were added to the queue, not the number of DPCs in the queue. This counter displays the last observed value only; it is not an average. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| DPCs Queued/sec | Shows the average rate, in incidents per second, at which deferred procedure calls (DPCs) were added to the processor's DPC queue. DPCs are interrupts that run at a lower priority than standard interrupts. Each processor has its own DPC queue. This counter measures the rate at which DPCs were added to the queue, not the number of DPCs in the queue. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval. | PERF_COUNTER_COUNTER |
| Interrupts/sec | Shows the rate, in incidents per second, at which the processor received and serviced hardware interrupts. It does not include deferred procedure calls (DPCs), which are counted separately. This value is an indirect indicator of the activity of devices that generate interrupts, such as the system clock, the mouse, disk drivers, data communication lines, network interface cards, and other peripheral devices. These devices normally interrupt the processor when they have completed a task or require attention. Normal thread execution is suspended during interrupts. Most system clocks interrupt the processor every 10 milliseconds, creating a background of interrupt activity. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Finding Bottlenecks on Multiprocessor Computers
Find processor bottlenecks

# PSched Flow Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The PSched Flow Object shows statistics from the packet scheduler, a feature of Counters by Feature

| Counter Name | Description | Counter Type |
|---|---|---|
| Packets dropped | Shows the number of packets dropped by the packet scheduler. | PERF_COUNTER_RAWCOUNT |
| Packets scheduled | Shows the number of packets that were scheduled in some way (rather than just being directly sent to the underlying miniport). | PERF_COUNTER_RAWCOUNT |
| Packets transmitted | Shows the number of packets from this flow that have been sent. | PERF_COUNTER_RAWCOUNT |
| Average packets in shaper | Shows the average number of packets in the scheduler over the last sampling period. | PERF_COUNTER_RAWCOUNT |
| Max packets in shaper | Shows the maximum number of packets that have ever simultaneously been in the scheduler. | PERF_COUNTER_RAWCOUNT |
| Average packets in sequencer | Shows the average number of packets in the sequencer over the last sampling period. | PERF_COUNTER_RAWCOUNT |
| Max packets in sequencer | Shows the maximum number of packets that have ever simultaneously been in the sequencer. | PERF_COUNTER_RAWCOUNT |
| Bytes scheduled | Shows the number of bytes that were scheduled in some way (rather than just being directly sent to the underlying miniport). | PERF_COUNTER_RAWCOUNT |
| Bytes transmitted | Shows the number of bytes from this flow that have been sent. | PERF_COUNTER_RAWCOUNT |
| Bytes transmitted/sec | Shows the number of bytes per second from this flow that have been sent. | PERF_COUNTER_COUNTER |
| Bytes scheduled/sec | Shows the number of bytes per second from this flow that have been scheduled. | PERF_COUNTER_COUNTER |
| Packets transmitted/sec | Shows the number of packets per second from this flow that have been sent. | PERF_COUNTER_COUNTER |
| Packets scheduled/sec | Shows the number of packets per second from this flow that have been scheduled. | PERF_COUNTER_COUNTER |
| Packets dropped/sec | Shows the number of packets that were dropped by the packet scheduler from this flow per second. | PERF_COUNTER_COUNTER |
| Nonconforming packets scheduled | Shows the number of packets that have entered the packet scheduler at a rate which exceeded the flow parameters. | PERF_COUNTER_RAWCOUNT |
| Nonconforming packets scheduled/sec | Shows the rate at which nonconforming packets have entered the packet scheduler. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Nonconforming packets transmitted | Shows the number of packets that have been sent by the packet scheduler at a rate which exceeded the flow parameters. | PERF_COUNTER_RAWCOUNT |
| Nonconforming packets transmitted/sec | Shows the rate at which nonconforming packets have been sent by the packet scheduler. | PERF_COUNTER_COUNTER |
| Maximum Packets in netcard | Shows the maximum number of packets that have been queued in the network adapter by this flow. | PERF_COUNTER_RAWCOUNT |
| Average Packets in netcard | Shows the average number of packets that have been queued in the network adapter by this flow. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
PSched Pipe Object

# PSched Pipe Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The PSched Pipe performance object reports statistics from the packet scheduler pipe. On a local area network (LAN) adapter, the pipe represents the network interface card, and on a wide area network (WAN) adapter, the pipe represents each WAN line.

| Counter Name | Description | Counter Type |
|---|---|---|
| Average Packets in Netcard | Shows the average number of packets in the pipe over the last sampling period. | PERF_COUNTER_RAWCOUNT |
| Average Packets in Sequencer | Shows the average number of packets in the sequencer over the last sampling period. | PERF_COUNTER_RAWCOUNT |
| Average Packets in Shaper | Shows the average number of packets in the scheduler shaper over the last sampling period. | PERF_COUNTER_RAWCOUNT |
| Flow Mods Rejected | Shows the number of times that a flow modification has been rejected. | PERF_COUNTER_RAWCOUNT |
| Flows Closed | Shows the number of flows that have been closed. | PERF_COUNTER_RAWCOUNT |
| Flows Modified | Shows the number of times that a flow has been modified succesfully. | PERF_COUNTER_RAWCOUNT |
| Flows Opened | Shows the number of flows that were opened on this pipe (some of which may now be closed). | PERF_COUNTER_RAWCOUNT |
| Flows Rejected | Shows the number of flow creations that were rejected. | PERF_COUNTER_RAWCOUNT |
| Max Packets in Netcard | Shows the maximum number of packets that were ever simultaneously in the network card. | PERF_COUNTER_RAWCOUNT |
| Max Packets in Sequencer | Shows the maximum number of packets that were ever simultaneously in the sequencer. | PERF_COUNTER_RAWCOUNT |
| Max Packets in Shaper | Shows the maximum number of packets that were ever simultaneously in the scheduler. | PERF_COUNTER_RAWCOUNT |
| Max Simultaneous Flows | Shows the maximum number of flows that were simultaneously open on this pipe. | PERF_COUNTER_RAWCOUNT |
| Nonconforming Packets Scheduled | Shows the number of packets that entered the packet scheduler at a rate that exceeded the packet's flow parameters. | PERF_COUNTER_RAWCOUNT |
| Nonconforming Packets Scheduled/sec | Shows the rate, in incidents per second, at which nonconforming packets entered the packet scheduler. | PERF_COUNTER_COUNTER |
| Nonconforming Packets Transmitted | Shows the number of packets that were sent by the packet scheduler at a rate that exceeded the packet's flow parameters. | PERF_COUNTER_RAWCOUNT |
| Nonconforming Packets Transmitted/sec | Shows the rate, in incidents per second, at which nonconforming packets were sent by the packet scheduler. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Out of Packets | Shows the number of times that the packet scheduler has been unable to allocate memory to send or receive a packet. | PERF_COUNTER_RAWCOUNT |

PSched Flow Object

# RAS Port Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The RAS Port performance object consists of counters that monitor individual remote access service (RAS) ports of the RAS device on the system.

| Counter Name | Description | Counter Type |
|---|---|---|
| Alignment Errors | Shows the total number of alignment errors on this connection. Alignment errors occur when a received byte is different from the expected byte. | PERF_COUNTER_RAWCOUNT |
| Buffer Overrun Errors | Shows the total number of buffer overrun errors on this connection. Buffer overrun errors occur when the software cannot handle the rate at which data is received. | PERF_COUNTER_RAWCOUNT |
| Bytes Received | Shows the total number of bytes that were received on this connection. | PERF_COUNTER_RAWCOUNT |
| Bytes Received/sec | Shows the rate, in incidents per second, at which bytes were received. | PERF_COUNTER_COUNTER |
| Bytes Transmitted | Shows the total number of bytes that were transmitted on this connection. | PERF_COUNTER_RAWCOUNT |
| Bytes Transmitted/sec | Shows the rate, in incidents per second, at which bytes were transmitted. | PERF_COUNTER_COUNTER |
| CRC Errors | Shows the total number of cyclic redundancy check (CRC) errors on this connection. CRC errors occur when the frame received contains erroneous data. | PERF_COUNTER_RAWCOUNT |
| Frames Received | Shows the total number of data frames received on this connection. | PERF_COUNTER_RAWCOUNT |
| Frames Received/sec | Shows the rate, in incidents per second, at which frames were received. | PERF_COUNTER_COUNTER |
| Frames Transmitted | Shows the total number of data frames transmitted on this connection. | PERF_COUNTER_RAWCOUNT |
| Frames Transmitted/sec | Shows the rate, in incidents per second, at which frames were transmitted. | PERF_COUNTER_COUNTER |
| Percent Compression In | Shows the compression ratio for bytes that were received. | PERF_COUNTER_RAWCOUNT |
| Percent Compression Out | Shows the compression ratio for bytes that were transmitted. | PERF_COUNTER_RAWCOUNT |
| Serial Overrun Errors | Shows the total number of serial overrun errors on this connection. Serial overrun errors occur when the hardware cannot handle the rate at which data is received. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Timeout Errors | Shows the total number of timeout errors on this connection. Timeout errors occur when an expected packet is not received in time. | PERF_COUNTER_RAWCOUNT |
| Total Errors | Shows the total number of CRC, timeout, serial overrun, alignment, and buffer overrun errors on this connection. | PERF_COUNTER_RAWCOUNT |
| Total Errors/sec | Shows the total rate, in incidents per second, of CRC, timeout, serial overrun, alignment, and buffer overrun errors. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
RAS Total Object

# RAS Total Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The RAS Total performance object consists of counters that combine values for all ports of the remote access service (RAS) device on the computer.

| Counter Name | Description | Counter Type |
|---|---|---|
| Alignment Errors | Shows the total number of alignment errors for this connection. Alignment errors occur when a received byte is different from the expected byte. | PERF_COUNTER_RAWCOUNT |
| Buffer Overrun Errors | Shows the total number of buffer overrun errors for this connection. Buffer overrun errors occur when the software cannot handle the rate at which data is received. | PERF_COUNTER_RAWCOUNT |
| Bytes Received | Shows the total number of bytes that were received for this connection. | PERF_COUNTER_RAWCOUNT |
| Bytes Received/sec | Shows the rate, in incidents per second, at which bytes were received. | PERF_COUNTER_COUNTER |
| Bytes Transmitted | Shows the total number of bytes that were transmitted for this connection. | PERF_COUNTER_RAWCOUNT |
| Bytes Transmitted/sec | Shows the rate, in incidents per second, at which bytes were transmitted. | PERF_COUNTER_COUNTER |
| CRC Errors | Shows the total number of cyclic redundancy check (CRC) errors for this connection. CRC errors occur when the frame that is received contains erroneous data. | PERF_COUNTER_RAWCOUNT |
| Frames Received | Shows the total number of data frames that were received for this connection. | PERF_COUNTER_RAWCOUNT |
| Frames Received/sec | Shows the rate, in incidents per second, at which frames were received. | PERF_COUNTER_COUNTER |
| Frames Transmitted | Shows the total number of data frames that were transmitted for this connection. | PERF_COUNTER_RAWCOUNT |
| Frames Transmitted/sec | Shows the rate, in incidents per second, at which frames were transmitted. | PERF_COUNTER_COUNTER |
| Percent Compression In | Shows the compression ratio for bytes that were received. | PERF_COUNTER_RAWCOUNT |
| Percent Compression Out | Shows the compression ratio for bytes that were transmitted. | PERF_COUNTER_RAWCOUNT |
| Serial Overrun Errors | Shows the total number of serial overrun errors for this connection. Serial overrun errors occur when the hardware cannot handle the rate at which data is received. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Timeout Errors | Shows the total number of timeout errors for this connection. Timeout errors occur when an expected packet is not received within the time expected. | PERF_COUNTER_RAWCOUNT |
| Total Connections | Shows the total number of remote access connections. | PERF_COUNTER_RAWCOUNT |
| Total Errors | Shows the total number of CRC, timeout, serial overrun, alignment, and buffer overrun errors for this connection. | PERF_COUNTER_RAWCOUNT |
| Total Errors/sec | Shows the total rate, in incidents per second, of CRC, timeout, serial overrun, alignment, and buffer overrun errors. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
RAS Port Object

# Redirector Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Redirector performance object consists of counters that monitor network connections originating at the local computer.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Received/sec | Shows the rate, in incidents per second, of bytes that were received by the redirector from the network. It includes all application data and network protocol information, such as packet headers. | PERF_COUNTER_BULK_COUNT |
| Bytes Total/sec | Shows the rate, in incidents per second, at which the redirector processed data bytes. This includes all application and file data as well as protocol information, such as packet headers. | PERF_COUNTER_BULK_COUNT |
| Bytes Transmitted/sec | Shows the rate, in incidents per second, at which bytes left the redirector to the network. It includes all application data and network protocol information, such as packet headers. | PERF_COUNTER_BULK_COUNT |
| Connects Core | Shows the number of connections to servers running the original MS-Net SMB protocol, including MS-Net itself, Xenix, and VAXs. | PERF_COUNTER_RAWCOUNT |
| Connects LAN Manager 2.0 | Shows the number of connections to LAN Manager 2.0 servers, including LMX servers. | PERF_COUNTER_RAWCOUNT |
| Connects LAN Manager 2.1 | Shows the number of connections to LAN Manager 2.1 servers, including LMX servers. | PERF_COUNTER_RAWCOUNT |
| Connects Windows NT | Shows the total number of connections to computers that are running Windows NT, Windows 2000, Windows XP, and the Windows Server 2003 family. | PERF_COUNTER_RAWCOUNT |
| Current Commands | Shows the number of requests to the redirector that were queued for service. If this number is much larger than the number of network adapter cards installed in the computer, then the network(s) and/or the server(s) being accessed are seriously bottlenecked. | PERF_COUNTER_RAWCOUNT |
| File Data Operations/sec | Shows the rate, in incidents per second, at which the redirector processed data operations. To determine the efficiency of this path, calculate the average number of bytes transferred per operation by dividing the value of Bytes/sec by the value of this counter | PERF_COUNTER_COUNTER |
| File Read Operations/sec | Shows the rate, in incidents per second, at which applications asked the redirector for data. Each call to a file system, or similar application program interface (API) call, is counted as one operation. | PERF_COUNTER_COUNTER |
| File Write Operations/sec | Shows the rate, in incidents per second, at which applications sent data to the redirector. Each call to a file system, or similar API call, counts as one operation. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Network Errors/sec | Shows the rate, in incidents per second, at which serious unexpected errors occurred. Such errors generally indicate that the redirector and one or more servers are having serious communication difficulties. For example, a Server Manager Block protocol error generates a network error. An entry is written in the system event log, and can provide details. | PERF_COUNTER_COUNTER |
| Packets Received/sec | Shows the rate, in incidents per second, at which the redirector received packets (also called server message blocks or SMBs). Network transmissions are divided into packets. To calculate the average number of bytes received in a packet, divide the value of Bytes Received/sec by the value of this counter. Some packets received might not contain incoming data (for example, an acknowledgment of a write operation made by the redirector is counted as an incoming packet). | PERF_COUNTER_BULK_COUNT |
| Packets Transmitted/sec | Shows the rate, in incidents per second, at which the redirector sent packets (also called server message blocks or SMBs). Network transmissions are divided into packets. To calculate the average number of bytes transmitted in a packet, divide the value of Bytes Transmitted/sec by the value of this counter. | PERF_COUNTER_BULK_COUNT |
| Packets/sec | Shows the rate, in incidents per second, at which the redirector processed data packets. To determine the efficiency of this path, calculate the average number of bytes transferred per packet by dividing the value of Bytes Transmitted/sec by the value of this counter. For another measure of efficiency, calculate the average number of packets transferred per operation by dividing the value of this counter by the value of Operations/sec. | PERF_COUNTER_BULK_COUNT |
| Read Bytes Cache/sec | Shows the rate, in incidents per second, at which applications used the redirector to access the file system. Some data requests are satisfied by retrieving the data from the cache. Requests that miss the cache cause a page fault. | PERF_COUNTER_BULK_COUNT |
| Read Bytes Network/sec | Shows the rate, in incidents per second, at which applications read data across the network. This occurs when data sought in the file system cache is not found there and must be retrieved from the network. Dividing this value by the value of Bytes Received/sec indicates the proportion of application data that is traveling across the network. | PERF_COUNTER_BULK_COUNT |
| Read Bytes Non-Paging/sec | Shows the number of bytes, in incidents per second, read by the redirector in response to normal file requests by an application when they were redirected to come from another computer. In addition to file requests, this counter includes other methods of reading across the network, such as Named Pipes and Transactions. This counter counts only application data, not network protocol information. | PERF_COUNTER_BULK_COUNT |
| Read Bytes Paging/sec | Shows the rate, in incidents per second, at which the redirector attempted to read bytes in response to page faults. Page faults are caused by loading modules (such as programs and libraries), by a miss in the file system cache, or by files directly mapped into the address space of applications. | PERF_COUNTER_BULK_COUNT |
| Read Operations Random/sec | Shows the rate, in incidents per second, at which, on a file-by-file basis, non-sequential read operations are made. If a read operation is made using a particular file handle, and then is followed by another read operation that is not the next contiguous byte, this counter is incremented by one. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Read Packets Small/sec | Shows the rate, in incidents per second, at which read operations of less than one-fourth of the server's negotiated buffer size were made by applications. A high rate of small packet read operations indicates the possibility that buffers are being used inefficiently. This counter is incremented once for each small read operation. It does not count packets. | PERF_COUNTER_COUNTER |
| Read Packets/sec | Shows the rate, in incidents per second, at which read packets were placed on the network. Each time a single packet is sent with a request to read data remotely, this counter is incremented by one. | PERF_COUNTER_COUNTER |
| Reads Denied/sec | Shows the rate, in incidents per second, at which the server was unable to accommodate requests for raw read operations. When a read operation is much larger than the server's negotiated buffer size, the redirector requests a raw read operation. If granted, a raw read operation permits the transfer of the data without much protocol overhead on each packet. However, to accomplish a raw read operation, the server must lock out other requests; if the server is too busy, the request is denied. | PERF_COUNTER_COUNTER |
| Reads Large/sec | Shows the rate, in incidents per second, at which read operations that were more than double the server's negotiated buffer size were made by applications. Too many large read operations could place a strain on server resources. This counter is incremented once for each large read operation. It does not count packets. | PERF_COUNTER_COUNTER |
| Server Disconnects | Shows the number of times that a server has disconnected the redirector. | PERF_COUNTER_RAWCOUNT |
| Server Reconnects | Shows the number of times that the redirector had to reconnect to a server to complete a new active request. Users can be disconnected by the server if they remain inactive for too long. Even if all remote files are closed, the redirector keeps connections intact locally for about ten minutes. Such inactive connections are called dormant connections. Reconnecting is expensive in time. | PERF_COUNTER_RAWCOUNT |
| Server Sessions | Shows the total number of security objects that are managed by the redirector. For example, logging on to a server followed by a network access to the same server establishes one connection, but two sessions. | PERF_COUNTER_RAWCOUNT |
| Server Sessions Hung | Shows the number of active sessions that were timed out and were unable to proceed due to a lack of response from the remote server. | PERF_COUNTER_RAWCOUNT |
| Write Bytes Cache/sec | Shows the rate, in incidents per second, at which applications on a computer used the redirector to write to the file system cache. The data might not leave the computer immediately; it can be retained in the cache for further modification before being written to the network, reducing network traffic. Every byte that is written to the cache is counted here. | PERF_COUNTER_BULK_COUNT |
| Write Bytes Network/sec | Shows the rate, in incidents per second, at which applications wrote data across the network. This occurs when the file system cache is bypassed, such as for named pipes or transactions, or else when the cache writes the bytes to disk to make room for other data. To calculate what proportion of application data is being written to the network, divide the value of this counter by the value of Bytes Transmitted/sec. | PERF_COUNTER_BULK_COUNT |

| | | |
|---|---|---|
| Write Bytes Non-Paging/sec | Shows the rate, in incidents per second, at which bytes were written by the redirector. This is in response to normal file outputs by an application, when the bytes were redirected to another computer. In addition to file requests, this counter includes other methods of writing across the network such as Named Pipes and Transactions. This counter only counts application data, and does not count network protocol information. | PERF_COUNTER_BULK_COUNT |
| Write Bytes Paging/sec | Shows the rate, in incidents per second, at which the redirector attempted to write bytes that were changed in the pages being used by applications. The program data changed by modules (such as programs and libraries) that were loaded over the network are paged out when no longer needed. Other output pages come from the cache. | PERF_COUNTER_BULK_COUNT |
| Write Operations Random/sec | Shows the rate, in incidents per second, at which, on a file-by-file basis, non-sequential write operations were made. If a write operation is made using a particular file handle and then is followed by another write operation that is not the next contiguous byte, this counter is incremented by one. | PERF_COUNTER_COUNTER |
| Write Packets Small/sec | Shows the rate, in incidents per second, at which write operations were made by applications that are less than one-fourth of the server's negotiated buffer size. Too many write operations can indicate a waste of buffers on the server. This counter shows write operations, not packets, and is incremented once for each write operation. | PERF_COUNTER_COUNTER |
| Write Packets/sec | Shows the rate, in incidents per second, at which write operations were sent to the network. Each time a single packet is sent with a request to write remote data, this counter is incremented by one. | PERF_COUNTER_COUNTER |
| Writes Denied/sec | Shows the rate, in incidents per second, at which the server was unable to accommodate requests for raw write operations. When a write operation is much larger than the server's negotiated buffer size, the redirector requests a raw write operation. If granted, a raw write operation permits the transfer of the data without much protocol overhead on each packet. However, to accomplish a raw write operation, the server must lock out other requests; if the server is too busy, the request is denied. | PERF_COUNTER_COUNTER |
| Writes Large/sec | Shows the rate, in incidents per second, at which write operations that were more than double the server's negotiated buffer size were made by applications. Too many large write operations could place a strain on server resources. This counter is incremented once for each large write operation. It does not count packets. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Server Object
Server Work Queues Object

# Server Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Server performance object consists of counters that measure communication between the local computer and the network.

| Counter Name | Description | Counter Type |
|---|---|---|
| Blocking Requests Rejected | Shows the number of times that the server rejected blocking server message block (SMB) requests due to insufficient count of free work items. This counter indicates whether the MaxWorkItem or MinFreeWorkItems server registry parameters might need tuning. | PERF_COUNTER_COUNTER |
| Bytes Received/sec | Shows the rate, in incidents per second, at which the server received bytes from the network. This counter indicates how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Bytes Total/sec | Shows the rate, in incidents per second, at which the server transmitted bytes through the network. This value provides an overall indication of how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Bytes Transmitted/sec | Shows the rate, in incidents per second, at which the server sent bytes on the network. This counter indicates how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Context Blocks Queued/sec | Shows the rate, in incidents per second, at which work context blocks were placed on the server's file system process queue to await server action. | PERF_COUNTER_COUNTER |
| Errors Access Permissions | Shows the number of times that attempts to open files on behalf of clients have failed with the message STATUS_ACCESS_DENIED. If this counter is high, it can indicate that someone is attempting to access random files without proper authorization. | PERF_COUNTER_RAWCOUNT |
| Errors Granted Access | Shows the number of times that attempts to access files successfully opened were denied. This counter can indicate attempts to access files without proper authorization. | PERF_COUNTER_RAWCOUNT |
| Errors Logon | Shows the number of failed logon attempts to the server. If this counter is high, it can indicate that password guessing programs are attempting to compromise the server's security. | PERF_COUNTER_RAWCOUNT |
| Errors System | Shows the number of times that an internal server error was detected. Errors can reflect problems with logon, security, memory allocation, disk operations, transport driver interface operations, communication (such as receipt of unimplemented or unrecognized SMBs), or I/O Request Packet stack size for the server.<br><br>Many of these errors are also written to the System log and Security log in Event Viewer. The server can recover from most the errors displayed by this counter, but they are unexpected and should be reported to Microsoft Product Support Services. | PERF_COUNTER_RAWCOUNT |
| File Directory Searches | Shows the number of searches for files currently active in the server. This counter indicates current server activity. | PERF_COUNTER_RAWCOUNT |
| Files Open | Shows the number of files currently opened on the server. This counter indicates current server activity. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Files Opened Total | Shows the number of successful attempts to open a file, performed by the server on behalf of clients. This counter is useful in determining the amounts of file I/O and overhead for path-based operations, and for determining the effectiveness of open locks. | PERF_COUNTER_RAWCOUNT |
| Logon Total | Shows all interactive logon attempts, network logon attempts, service logon attempts, successful logon attempts, and failed logon attempts since the computer was last started. | PERF_COUNTER_RAWCOUNT |
| Logon/sec | Shows the rate, in incidents per second, of all interactive logon attempts, network logon attempts, service logon attempts, successful logon attempts, and failed logon attempts since the computer was last started. | PERF_COUNTER_COUNTER |
| Pool Nonpaged Bytes | Shows the size, in bytes, of nonpageable computer memory that the server was using. This value is useful for determining the values of the MaxNonpagedMemoryUsage entry in the Windows Server 2003 family registry. | PERF_COUNTER_RAWCOUNT |
| Pool Nonpaged Failures | Shows the number of times that allocations from the nonpaged pool have failed. If this counter is high, it indicates that the computer's physical memory is too small. | PERF_COUNTER_COUNTER |
| Pool Nonpaged Peak | Shows the maximum size, in bytes, of the nonpaged pool that the server had in use at any one point. This counter indicates how much physical memory the computer should have. | PERF_COUNTER_RAWCOUNT |
| Pool Paged Bytes | Shows the size, in bytes, of pageable computer memory that the server was using. This counter can help to determine good values for the MaxPagedMemoryUsage registry entry. | PERF_COUNTER_RAWCOUNT |
| Pool Paged Failures | Shows the number of times that allocations from the paged pool have failed. If this counter is high, it indicates that the computer's physical memory or page file are too small. | PERF_COUNTER_RAWCOUNT |
| Pool Paged Peak | Shows the maximum size, in bytes, of the paged pool that the server has allocated. This counter indicates the proper sizes of the page file(s) and physical memory. | PERF_COUNTER_RAWCOUNT |
| Server Sessions | Shows the number of sessions currently active in the server. This counter indicates current server activity. | PERF_COUNTER_RAWCOUNT |
| Sessions Errored Out | Shows the number of sessions that have been closed due to unexpected error conditions. This counter indicates how frequently network problems caused dropped sessions on the server. This counter reports auto-disconnects as well as errored-out sessions. For a more accurate value for errored-out sessions, subtract the value of Sessions Timed Out from Sessions Errored Out. | PERF_COUNTER_RAWCOUNT |
| Sessions Forced Off | Shows the number of sessions that have been forced to log off. This counter can indicate how many sessions were forced to log off due to logon time constraints. | PERF_COUNTER_RAWCOUNT |
| Sessions Logged Off | Shows the number of sessions that have terminated normally. This counter is useful in interpreting the statistics from Sessions Timed Out and Sessions Errored Out. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Sessions Timed Out | Shows the number of sessions that have been closed because idle time exceeded the AutoDisconnect parameter for the server. This counter shows whether the AutoDisconnect setting is helping to conserve resources. | PERF_COUNTER_RAWCOUNT |
| Work Item Shortages | Shows the number of times that STATUS_DATA_NOT_ACCEPTED was returned at receive indication time. This occurs when a work item is not available or cannot be allocated to service the incoming request. This counter indicates whether the InitWorkItems or MaxWorkItems registry entries might need to be adjusted. | PERF_COUNTER_COUNTER |

# Server Work Queues Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Server Work Queues performance object consists of counters that monitor the length of queues and objects in the queues.

| Counter Name | Description | Counter Type |
|---|---|---|
| Active Threads | Shows the number of threads working on a request from the server client for this CPU. The system keeps this number as low as possible to minimize unnecessary context switching. This is an instantaneous count for the CPU, not an average over time. | PERF_COUNTER_RAWCOUNT |
| Available Threads | Shows the number of server threads on this processor that are not currently working on requests from a client. The server dynamically adjusts the number of threads to maximize server performance. | PERF_COUNTER_RAWCOUNT |
| Available Work Items | Shows the instantaneous number of available work items for this CPU. A work item is typically a request from a client to the server; the server maintains a pool of available work items per CPU to speed processing. A sustained near-zero value indicates the need to increase the MinFreeWorkItems registry value for the server service. This value is always zero in the Blocking Queue instance. | PERF_COUNTER_RAWCOUNT |
| Borrowed Work Items | Shows the number of free work items borrowed from another CPU because a CPU ran out of available work items in its pool. An increasing value of this running counter might indicate the need to increase the MaxWorkItems or MinFreeWorkItems registry values for the server service. This value is always zero in the Blocking Queue instance. | PERF_COUNTER_RAWCOUNT |
| Bytes Received/sec | Shows the rate, in incidents per second, at which the server received bytes from the network clients on this CPU. This value is a measure of how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Bytes Sent/sec | Shows the rate, in incidents per second, at which the server sent bytes to the network clients on this CPU. This value is a measure of how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Bytes Transferred/sec | Shows the rate, in incidents per second, at which the server exchanged bytes with the network clients on this CPU. This value is a measure of how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Context Blocks Queued/sec | Shows the rate, in incidents per second, at which work context blocks have to be placed on the server's file system process queue to await server action. | PERF_COUNTER_COUNTER |
| Current Clients | Shows the current count of the clients being processed by this CPU. The server actively balances the client load across all of the CPUs in the system. If a single work queue can easily and quickly handle all the requests from clients, the clients might not be spread across multiple queues until a higher load is encountered. This value is always zero in the Blocking Queue instance. | PERF_COUNTER_RAWCOUNT |
| Queue Length | Shows the current length of the server work queue for this CPU. A sustained queue length greater than four might indicate processor congestion. This is an instantaneous count, not an average over time. | PERF_COUNTER_RAWCOUNT |
| Read Bytes/sec | Shows the rate, in incidents per second, at which the server read data from files for the clients on this CPU. This value is a measure of how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Read Operations/sec | Shows the rate, in incidents per second, at which the server performed file read operations for the clients on this CPU. This value is a measure of how busy the server is. This value is always zero in the Blocking Queue instance. | PERF_COUNTER_BULK_COUNT |
| Total Bytes/sec | Shows the rate, in incidents per second, at which the server read and wrote data to and from the files for the clients on this CPU. This value is a measure of how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Total Operations/sec | Shows the rate, in incidents per second, at which the server performed file read and write operations for the clients on this CPU. This value is a measure of how busy the server is. This value is always zero in the Blocking Queue instance. | PERF_COUNTER_BULK_COUNT |

| | | |
|---|---|---|
| Work Item Shortages | Shows the number of times that a work item was not available or could not be allocated to service the incoming request. A work item is the location where the server stores a server message block (SMB). The amount of work items available fluctuates between a minimum and maximum value configured, based on how the server is configured and the amount of memory on the computer. If work item shortages are occurring, it might be caused by an overloaded server. If the Work Item Shortages counter value is increasing, consider changing the value of the registry entry HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\Maxworkitems. Allowing this value to be achieved consistently initiates flow control, which hurts performance. This value is always zero in the Blocking Queue instance. | PERF_COUNTER_RAWCOUNT |
| Write Bytes/sec | Shows the rate, in incidents per second, at which the server wrote data to files for the clients on this CPU. This value is a measure of how busy the server is. | PERF_COUNTER_BULK_COUNT |
| Write Operations/sec | Shows the rate, in incidents per second, at which the server performed file write operations for the clients on this CPU. This value is a measure of how busy the server is. This value is always zero in the Blocking Queue instance. | PERF_COUNTER_BULK_COUNT |

# SMTP Server Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The SMTP Server performance object consists of counters that describe the activity of the Exchange NTFS store driver, which is responsible for storing queued messages that are being processed by the Simple Mail Transfer Protocol (SMTP) service.

| Counter Name | Description | Counter Type |
|---|---|---|
| % Recipients Local | Shows the percentage of recipients that were delivered locally. | PERF_RAW_FRACTION |
| % Recipients Remote | Shows the percentage of recipients that were delivered remotely. | PERF_RAW_FRACTION |
| Avg Recipients/msg Received | Shows the average number of recipients per inbound message received. | PERF_RAW_FRACTION |
| Avg Recipients/msg Sent | Shows the average number of recipients per outbound message sent. | PERF_RAW_FRACTION |
| Avg Retries/msg Delivered | Shows the average number of retries per local delivery. | PERF_RAW_FRACTION |
| Avg Retries/msg Sent | Shows the average number of retries per outbound message sent. | PERF_RAW_FRACTION |
| Badmailed Messages (Bad Pickup File) | Shows the number of malformed pickup messages sent to the badmail directory, which is a directory used to save messages that cannot be handled. These are usually non-delivery reports (NDRs) generated in response to NDRs. | PERF_COUNTER_RAWCOUNT |
| Badmailed Messages (General Failure) | Shows the number of messages sent to badmail for reasons not associated with a specific counter. | PERF_COUNTER_RAWCOUNT |
| Badmailed Messages (Hop Count Exceeded) | Shows the number of messages sent to badmail because they exceeded the maximum hop count. | PERF_COUNTER_RAWCOUNT |
| Badmailed Messages (NDR of DSN) | Shows the number of Delivery Status Notifications (DSN) sent to badmail because a message could not be delivered. If a message is undeliverable, it is returned to the sender with an NDR. | PERF_COUNTER_RAWCOUNT |
| Badmailed Messages (No Recipients) | Shows the number of messages sent to badmail because they had no recipients. | PERF_COUNTER_RAWCOUNT |
| Badmailed Messages (Triggered via Event) | Shows the number of messages sent to badmail at the request of a server event sink. The server events mechanism is used by SMTP to notify event sinks when specific events occur. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Bytes Received Total | Shows the total number of bytes received by the server. | PERF_COUNTER_LARGE_RAWCOUNT |
| Bytes Received/sec | Shows the rate, in incidents per second, at which bytes were received by the server. | PERF_COUNTER_BULK_COUNT |
| Bytes Sent Total | Shows the total number of bytes sent by the server. | PERF_COUNTER_LARGE_RAWCOUNT |
| Bytes Sent/sec | Shows the rate, in incidents per second, at which bytes were sent by the server. | PERF_COUNTER_BULK_COUNT |
| Bytes Total | Shows the total number of bytes sent and received by the server. | PERF_COUNTER_LARGE_RAWCOUNT |
| Bytes Total/sec | Shows the rate, in incidents per second, at which bytes were sent and received by the server. | PERF_COUNTER_BULK_COUNT |
| Cat: Address Lookup Completions | Shows the number of address lookup completions processed by the server. | PERF_COUNTER_RAWCOUNT |
| Cat: Address Lookup Completions/sec | Shows the rate, in incidents per second, at which address lookup completions were processed by the server. | PERF_COUNTER_COUNTER |
| Cat: Address Lookups | Shows the number of Lightweight Directory Access Protocol (LDAP) Directory lookups for individual addresses by the server. | PERF_COUNTER_RAWCOUNT |
| Cat: Address Lookups Not Found | Shows the number of address lookups that did not find any LDAP Directory object. | PERF_COUNTER_RAWCOUNT |
| Cat: Address Lookups/sec | Shows the rate, in incidents per second, at which address lookups were dispatched to the LDAP Directory. | PERF_COUNTER_COUNTER |
| Cat: Categorizations Completed | Shows the total number of messages to the categorizer that have completed categorization. Categorizing means looking up sender and recipient properties in the directory service, including doing distribution list expansion. | PERF_COUNTER_RAWCOUNT |
| Cat: Categorizations Completed Successfully | Shows the number of categorizations that were completed without any errors. | PERF_COUNTER_RAWCOUNT |
| Cat: Categorizations Completed/sec | Shows the rate, in incidents per second, at which messages submitted to categorizer were completed. | PERF_COUNTER_COUNTER |
| Cat: Categorizations Failed (DS Connection Failure) | Shows the number of categorizations that failed due to an LDAP Directory connection failure. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Cat: Categorizations Failed (DS Logon Failure) | Shows the number of categorizations that failed due to an LDAP Directory logon failure. | PERF_COUNTER_RAWCOUNT |
| Cat: Categorizations Failed (Non-Retryable Error) | Shows the number of categorizations that failed with a hard error (one which doesn't allow attempts to retry the operation). | PERF_COUNTER_RAWCOUNT |
| Cat: Categorizations Failed (Out Of Memory) | Shows the number of categorizations that failed due to lack of available memory. | PERF_COUNTER_RAWCOUNT |
| Cat: Categorizations Failed (Retryable Error) | Shows the number of categorizations that failed with a error that allowed an attempt to retry the operation. | PERF_COUNTER_RAWCOUNT |
| Cat: Categorizations Failed (Sink Retryable Error) | Shows the number of categorizations that failed with a generic error that allowed an attempt to retry the operation. | PERF_COUNTER_RAWCOUNT |
| Cat: Categorizations In Progress | Shows the number of categorizations in progress. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Bind Failures | Shows the total number of LDAP bind failures. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Binds | Shows the total number of successful LDAP binds, also known as bind authentications, that were performed. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Connection Failures | Shows the total number of connection failures encountered for LDAP servers. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Connections | Shows the total number of LDAP connections opened since the computer was last started. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Connections Currently Open | Shows the number of LDAP connections currently open. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP General Completion Failures | Shows the number of LDAP completions with a generic failure. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Paged Search Completion Failures | Shows the number of LDAP paged searches that completed with a failure. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Paged Search Failures | Shows the number of failures to dispatch an asynchronous paged LDAP search. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Cat: LDAP Paged Searches | Shows the number of LDAP paged searches that were successfully dispatched. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Paged Searches Completed | Shows the number of paged LDAP completions processed. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Search Completion Failures | Shows the number of LDAP searches that completed with a failure. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Search Failures | Shows the number of failures to dispatch an asynchronous LDAP search. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Searches | Shows the number of LDAP searches that were successfully dispatched. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Searches Abandoned | Shows the umber of LDAP searches that were abandoned. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Searches Completed | Shows the number of LDAP search completions that were processed. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Searches Completed/sec | Shows the rate, in incidents per second, at which LDAP search completions were processed. | PERF_COUNTER_COUNTER |
| Cat: LDAP Searches Pending Completion | Shows the number of LDAP searches that were pending asynchronous completion. | PERF_COUNTER_RAWCOUNT |
| Cat: LDAP Searches/sec | Shows the rate, in incidents per second, at which number LDAP searches were successfully dispatched. | PERF_COUNTER_COUNTER |
| Cat: Mailmsg Duplicate Collisions | Shows the number of times that a duplicate recipient address was detected by either mailmsg or the categorizer. | PERF_COUNTER_RAWCOUNT |
| Cat: Messages Aborted | Shows the number of messages that were marked to be aborted by the categorizer. | PERF_COUNTER_RAWCOUNT |
| Cat: Messages Bifurcated | Shows the number of new messages that were created by the categorizer (bifurcation). | PERF_COUNTER_RAWCOUNT |
| Cat: Messages Categorized | Shows the number of messages that the categorizer has submitted to queuing. | PERF_COUNTER_RAWCOUNT |
| Cat: Messages Submitted | Shows the total number of messages that were submitted to the categorizer. | PERF_COUNTER_RAWCOUNT |
| Cat: Messages Submitted/sec | Shows the rate, in incidents per second, at which messages were submitted to the categorizer. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Cat: Recipients After Categorization | Shows the number of mail message recipients that were submitted from the categorizer to queuing. | PERF_COUNTER_RAWCOUNT |
| Cat: Recipients Before Categorization | Shows the number of mail message recipients that were submitted to the categorizer. | PERF_COUNTER_RAWCOUNT |
| Cat: Recipients In Categorization | Shows the number of recipients that the categorizer was currently processing. | PERF_COUNTER_RAWCOUNT |
| Cat: Recipients NDRd (Ambiguous Address) | Shows the number of recipients with addresses that match multiple LDAP Directory objects. | PERF_COUNTER_RAWCOUNT |
| Cat: Recipients NDRd (Forwarding Loop) | Shows the number of recipients that received an NDR from the categorizer due to forwarding loop detection. | PERF_COUNTER_RAWCOUNT |
| Cat: Recipients NDRd (Illegal Address) | Shows the number of recipients that received an NDR from the categorizer due to illegal addresses. | PERF_COUNTER_RAWCOUNT |
| Cat: Recipients NDRd (Sink Recip Errors) | Shows the number of recipients that received an NDR from the categorizer due to generic recipient failure. | PERF_COUNTER_RAWCOUNT |
| Cat: Recipients NDRd (Unresolved) | Shows the number of unresolved recipients (local addresses not found). | PERF_COUNTER_RAWCOUNT |
| Cat: Recipients NDRd by Categorizer | Shows the number of recipients set to receive an NDR from the categorizer. | PERF_COUNTER_RAWCOUNT |
| Cat: Senders Unresolved | Shows the number of senders that were not found in the LDAP Directory. | PERF_COUNTER_RAWCOUNT |
| Cat: Senders With Ambiguous Addresses | Shows the number of senders with addresses that match multiple LDAP Directory objects. | PERF_COUNTER_RAWCOUNT |
| Categorizer Queue Length | Shows the number of messages in the categorizer queue waiting to be categorized. | PERF_COUNTER_RAWCOUNT |
| Connection Errors/sec | Shows the rate, in incidents per second, at which connection errors occurred. | PERF_COUNTER_COUNTER |
| Current Messages In Local Delivery | Shows the number of messages that were currently being processed by a server event sink for local delivery. | PERF_COUNTER_RAWCOUNT |
| Directory Drops Total | Shows the total number of messages placed in a drop directory. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Directory Drops/sec | Shows the rate, in incidents per second, at which messages were placed in a drop directory. | PERF_COUNTER_COUNTER |
| DNS Queries Total | Shows the total number of DNS lookups. | PERF_COUNTER_RAWCOUNT |
| DNS Queries/sec | Shows the rate, in incidents per second, of DNS lookups. | PERF_COUNTER_COUNTER |
| ETRN Messages Total | Shows the total number of extended turn (ETRN) messages received by the server. | PERF_COUNTER_RAWCOUNT |
| ETRN Messages/sec | Shows the rate, in incidents per second, of ETRN messages. | PERF_COUNTER_COUNTER |
| Inbound Connections Current | Shows the total number of connections that were currently inbound. | PERF_COUNTER_RAWCOUNT |
| Inbound Connections Total | Shows the total number of inbound connections received. | PERF_COUNTER_RAWCOUNT |
| Local Queue Length | Shows the number of messages in the local queue. | PERF_COUNTER_RAWCOUNT |
| Local Retry Queue Length | Shows the number of messages in the local retry queue. | PERF_COUNTER_RAWCOUNT |
| Message Bytes Received Total | Shows the total number of bytes received in messages. | PERF_COUNTER_LARGE_RAWCOUNT |
| Message Bytes Received/sec | Shows the rate, in incidents per second, at which bytes were received in messages. | PERF_COUNTER_BULK_COUNT |
| Message Bytes Sent Total | Shows the total number of bytes sent in messages. | PERF_COUNTER_LARGE_RAWCOUNT |
| Message Bytes Sent/sec | Shows the rate, in incidents per second, at which bytes were sent in messages. | PERF_COUNTER_BULK_COUNT |
| Message Bytes Total | Shows the total number of bytes sent and received in messages. | PERF_COUNTER_LARGE_RAWCOUNT |
| Message Bytes Total/sec | Shows the rate, in incidents per second, at which bytes were sent and received in messages. | PERF_COUNTER_BULK_COUNT |
| Message Delivery Retries | Shows the total number of local deliveries that were retried. | PERF_COUNTER_RAWCOUNT |
| Message Received/sec | Shows the rate, in incidents per second, at which inbound messages were being received. | PERF_COUNTER_COUNTER |
| Message Send Retries | Shows the total number of outbound message sends that were retried. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Messages Currently Undeliverable | Shows the number of messages that were reported as currently undeliverable by routing. | PERF_COUNTER_RAWCOUNT |
| Messages Delivered Total | Shows the total number of messages that were delivered to local mailboxes. | PERF_COUNTER_RAWCOUNT |
| Messages Delivered/sec | Shows the rate, in incidents per second, at which messages were delivered to local mailboxes. | PERF_COUNTER_COUNTER |
| Messages Pending Routing | Shows the number of messages that were categorized but not routed. | PERF_COUNTER_RAWCOUNT |
| Messages Received Total | Shows the total number of inbound messages that were accepted. | PERF_COUNTER_RAWCOUNT |
| Messages Refused for Address Objects | Shows the total number of messages that were refused due to no address objects. | PERF_COUNTER_RAWCOUNT |
| Messages Refused for Mail Objects | Shows the total number of messages that were refused due to no mail objects. | PERF_COUNTER_RAWCOUNT |
| Messages Refused for Size | Shows the total number of messages that were rejected because they were too big. | PERF_COUNTER_RAWCOUNT |
| Messages Sent Total | Shows the total number of outbound messages that were sent. | PERF_COUNTER_RAWCOUNT |
| Messages Sent/sec | Shows the rate, in incidents per second, at which outbound messages were sent. | PERF_COUNTER_COUNTER |
| NDRs Generated | Shows the number of NDRs that were generated. | PERF_COUNTER_RAWCOUNT |
| Number of MailFiles Open | Shows the number of handles to open mail files. | PERF_COUNTER_RAWCOUNT |
| Number of QueueFiles Open | Shows the number of handles to open queue files. | PERF_COUNTER_RAWCOUNT |
| Outbound Connections Current | Shows the number of connections that were currently outbound. | PERF_COUNTER_RAWCOUNT |
| Outbound Connections Refused | Shows the number of outbound connection attempts that were refused by remote sites. | PERF_COUNTER_RAWCOUNT |
| Outbound Connections Total | Shows the total number of outbound connections that were attempted. | PERF_COUNTER_RAWCOUNT |
| Pickup Directory Messages Retrieved Total | Shows the total number of messages that were retrieved from the mail pick-up directory. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Pickup Directory Messages Retrieved/sec | Shows the rate, in incidents per second, at which messages were retrieved from the mail pick-up directory. | PERF_COUNTER_COUNTER |
| Remote Queue Length | Shows the number of messages that were in the remote queue. | PERF_COUNTER_RAWCOUNT |
| Remote Retry Queue Length | Shows the number of messages that were in the retry queue for remote delivery. | PERF_COUNTER_RAWCOUNT |
| Routing Table Lookups Total | Shows the total number of routing table lookups. | PERF_COUNTER_RAWCOUNT |
| Routing Table Lookups/sec | Shows the rate, in incidents per second, of routing table lookups. | PERF_COUNTER_COUNTER |
| Total Connection Errors | Shows the total number of connection errors. | PERF_COUNTER_RAWCOUNT |
| Total DSN Failures | Shows the total number of failed DSN generation attempts. | PERF_COUNTER_RAWCOUNT |
| Total messages submitted | Shows the total number of messages submitted to queuing for delivery. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
SMTP NTFS Store Driver Object

# SMTP NTFS Store Driver Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The SMTP NTFS Store Driver performance object consists of counters that describe the activity of the Exchange NTFS store driver, which is responsible for message transport to and from the message store. The Simple Mail Transport Protocol (SMTP) is the protocol in the TCP/IP suite that directs the exchange of electronic mail.

| Counter Name | Description | Counter Type |
|---|---|---|
| Messages Allocated | Shows the total number of messages allocated since the computer was last started. | PERF_COUNTER_RAWCOUNT |
| Messages Deleted | Shows the total number of messages deleted since the computer was last started. | PERF_COUNTER_RAWCOUNT |
| Messages Enumerated | Shows the total number of messages enumerated during startup. | PERF_COUNTER_RAWCOUNT |
| Messages in the Queue Directory | Shows the current number of messages in the queue directory. | PERF_COUNTER_RAWCOUNT |
| Open Message Bodies | Shows the current number of open message bodies. | PERF_COUNTER_RAWCOUNT |
| Open Message Streams | Shows the current number of open message streams. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
SMTP Server Object

# System Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The System performance object consists of counters that apply to more than one component of the computer. The data collected by the system counters is derived from activity in the processor, memory, or disk subsystems.

| Counter Name | Description | Counter Type |
|---|---|---|
| % Registry Quota In Use | Shows the percentage of the Total Registry Quota Allowed that was being used by the system. | PERF_RAW_FRACTION |
| Alignment Fixups/sec | Shows the rate, in incidents per second, at which alignment faults were fixed by the system. | PERF_COUNTER_COUNTER |
| Context Switches/sec | Shows the combined rate, in incidents per second, at which all processors on the computer were switched from one thread to another. It is the sum of the values of Thread\Context Switches/sec for each thread running on all processors on the computer, and is measured in numbers of switches. Context switches occur when a running thread voluntarily relinquishes the processor, or is preempted by a higher priority, ready thread. | PERF_COUNTER_COUNTER |
| Exception Dispatches/sec | Shows the rate, in incidents per second, at which exceptions were dispatched by the system. | PERF_COUNTER_COUNTER |
| File Control Bytes/sec | Shows the overall rate, in incidents per second, at which bytes were transferred for all file system operations that were neither read nor write operations, such as file system control requests and requests for information about device characteristics or status. | PERF_COUNTER_BULK_COUNT |
| File Control Operations/sec | Shows the combined rate, in incidents per second, of file system operations that were neither read nor write operations, such as file system control requests and requests for information about device characteristics or status. This is the inverse of File Data Operations/sec. | PERF_COUNTER_COUNTER |
| File Data Operations/sec | Shows the combined rate, in incidents per second, of read and write operations on disks, serial, or parallel devices. This is the inverse of File Control Operations/sec. | PERF_COUNTER_COUNTER |
| File Read Bytes/sec | Shows the overall rate, in incidents per second, at which bytes were read to satisfy file system read requests to all devices on the computer, including read operations from the file system cache. | PERF_COUNTER_BULK_COUNT |
| File Read Operations/sec | Shows the combined rate, in incidents per second, of file system read requests to all devices on the computer, including requests to read from the file system cache. | PERF_COUNTER_COUNTER |
| File Write Bytes/sec | Shows the overall rate, in incidents per second, at which bytes were written to satisfy file system write requests to all devices on the computer, including write operations to the file system cache. | PERF_COUNTER_BULK_COUNT |
| File Write Operations/sec | Shows the combined rate, in incidents per second, of file system write requests to all devices on the computer, including requests to write to data in the file system cache. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Floating Emulations/sec | Shows the rate, in incidents per second, of floating emulations performed by the system. | PERF_COUNTER_COUNTER |
| Processes | Shows the number of processes in the computer at the time of data collection. This is an instantaneous count, not an average over the time interval. Each process represents a program that is running. | PERF_COUNTER_RAWCOUNT |
| Processor Queue Length | Shows the number of threads in the processor queue. Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time, even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of greater than two threads generally indicates processor congestion. | PERF_COUNTER_RAWCOUNT |
| System Calls/sec | Shows the combined rate, in incidents per second, of calls to operating system service routines by all processes running on the computer. These routines perform all of the basic scheduling and synchronization of activities on the computer, and provide access to nongraphic devices, memory management, and name space management. | PERF_COUNTER_COUNTER |
| System Up Time | Shows the total time, in seconds, that the computer has been operational since it was last started. | PERF_ELAPSED_TIME |
| Threads | Shows the number of threads in the computer at the time of data collection. This is an instantaneous count, not an average over the time interval. A thread is the basic executable entity that can execute instructions in a processor. | PERF_COUNTER_COUNTER |

# Telephony Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Telephony performance object provides counters for monitoring telephone-communication activity on a computer that is running a Windows Server 2003 family operating system.

| Counter Name | Description | Counter Type |
|---|---|---|
| Active Lines | Shows the number of telephone or integrated services digital network (ISDN) lines that were serviced by the computer and that were in use by applications during the sample period. | PERF_COUNTER_RAWCOUNT |
| Active Telephones | Shows the number of telephone devices (telephones or speaker phones) that were in use by applications during the sample period. | PERF_COUNTER_RAWCOUNT |
| Client Apps | Shows the number of applications that that were using telephony services. | PERF_COUNTER_RAWCOUNT |
| Current Incoming Calls | Shows the number of incoming calls that that were being serviced by the computer during the sample period. | PERF_COUNTER_RAWCOUNT |
| Current Outgoing Calls | Shows the number of outgoing calls that that were being serviced by the computer during the sample period. | PERF_COUNTER_RAWCOUNT |
| Incoming Calls/sec | Shows the rate, in incidents per second, at which incoming calls were answered by this computer. | PERF_COUNTER_COUNTER |
| Lines | Shows the number of telephone lines that were serviced by the computer. | PERF_COUNTER_RAWCOUNT |
| Outgoing Calls/sec | Shows the rate, in incidents per second, at which outgoing calls were made by the computer. | PERF_COUNTER_COUNTER |
| Telephone Devices | Shows the number of telephone devices (telephones or speaker phones) that were serviced by the computer. | PERF_COUNTER_RAWCOUNT |

# TCP Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The TCP performance object consists of counters that measure the rates at which Transmission Control Protocol (TCP) segments are sent and received using TCP. It includes counters that monitor the number of TCP connections in each TCP connection state.

| Counter Name | Description | Counter Type |
|---|---|---|
| Connection Failures | Shows the number of times that TCP connections have made a direct transition to the CLOSED state from the SYN-SENT or SYN-RCVD state, plus the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state, since the server was last started. | PERF_COUNTER_RAWCOUNT |
| Connections Active | Shows the number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state, since the server was last started. | PERF_COUNTER_RAWCOUNT |
| Connections Established | Shows the number of TCP connections for which the state was either ESTABLISHED or CLOSE-WAIT, since the server was last started. | PERF_COUNTER_RAWCOUNT |
| Connections Passive | Shows the number of times that TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state, since the server was last started. | PERF_COUNTER_RAWCOUNT |
| Connections Reset | Shows the number of times that TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED or CLOSE-WAIT state, since the server was last started. | PERF_COUNTER_RAWCOUNT |
| Segments Received/sec | Shows the rate, in incidents per second, at which segments were received, including those received in error. This count includes segments received on currently established connections. Segments Received/sec is a subset of Segments/sec. | PERF_COUNTER_COUNTER |
| Segments Retransmitted/sec | Shows the rate, in incidents per second, at which segments containing one or more previously transmitted bytes were retransmitted. | PERF_COUNTER_COUNTER |
| Segments Sent/sec | Shows the rate, in incidents per second, at which segments were sent. This value includes those on current connections, but excludes those containing only retransmitted bytes. Segments Sent/sec is a subset of Segments/sec. | PERF_COUNTER_COUNTER |
| Segments/sec | Shows the rate, in incidents per second, at which TCP segments were sent or received using the TCP protocol. Segments/sec is the sum of the values of Segments Received/sec and Segments Sent/sec. | PERF_COUNTER_COUNTER |

# Terminal Services Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Terminal Services object provides per-session resource statistics such as packet and frame transmission data, errors, and cache activity for Terminal Services.

In addition, the Terminal Services object reports session-specific activity using the following Process object counters:

| Counter Name | Description | Counter Type |
|---|---|---|
| Input Async Frame Error | Shows the number of input async framing errors. These can be caused by a noisy transmission line. Using a smaller packet size may help in some cases. | PERF_COUNTER_RAWCOUNT |
| Input Bytes | Shows the number of bytes input on this session that includes all protocol overhead. | PERF_COUNTER_COUNTER |

# Terminal Services Session Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Terminal Services Session performance object provides per-session resource statistics, such as packet and frame transmission data, errors, and cache activity for Terminal Services.

In addition, the Terminal Services Session performance object reports session-specific activity using the following Process object counters:

| | |
|---|---|
| Process Object | Process Object |
| Process Object | Process Object |
| Process Object | Process Object |
| Process Object | Process Object |
| Process Object | Process Object |
| Process Object | Process Object |
| Process Object | Process Object |
| Process Object | |

| Counter Name | Description | Counter Type |
|---|---|---|
| Input Async Frame Error | Shows the number of input async framing errors. These can be caused by a noisy transmission line. Using a smaller packet size can help in some cases. | PERF_COUNTER_RAWCOUNT |
| Input Async Overflow | Shows the number of input async overflow errors. These can be caused by a lack of buffer space available on the host. | PERF_COUNTER_RAWCOUNT |
| Input Async Overrun | Shows the number of input async overrun errors. These can be caused by the baud rate being faster than the computer can handle, or by a serial line that does not use a 16550 universal asynchronous receiver-transmitter (UART). Overruns can also occur if too many high speed serial lines are active at one time for the processor's power. | PERF_COUNTER_RAWCOUNT |
| Input Async Parity Error | Shows the number of input async parity errors. These can be caused by a noisy transmission line. | PERF_COUNTER_RAWCOUNT |
| Input Bytes | Shows the number of bytes input on this session, which includes all protocol overhead. | PERF_COUNTER_COUNTER |
| Input Compress Flushes | Shows the number of input compression dictionary flushes. When the data cannot be compressed, the compression dictionary is flushed so that newer data has a better chance of being compressed. One possible cause of data not compressing is transferring compressed files using client drive mapping. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Input Compressed Bytes | Shows the number of bytes that were input after compression. This number compared with the Total Bytes input is the compression ratio. | PERF_COUNTER_COUNTER |
| Input Compression Ratio | Shows the compression ratio of the server input data stream. This is Input Compressed Bytes compared with the Total Bytes input. | PERF_COUNTER_RAWCOUNT |
| Input Errors | Shows the number of input errors of all types. Some examples of input errors are lost acknowledgments (ACKs) or badly formed packets. | PERF_COUNTER_RAWCOUNT |
| Input Frames | Shows the number of frames input on this session. | PERF_COUNTER_COUNTER |
| Input Timeouts | Shows the total number of timeouts on the communication line, as seen from the client side of the connection. These are typically the result of a noisy line. On some high-latency networks, this could be the result of the protocol timeout being too short. Increasing the protocol timeout on these types of lines will improve performance by reducing the need for retransmissions. | PERF_COUNTER_RAWCOUNT |
| Input Transport Errors | Shows the number of Terminal Services transport-level errors on input. | PERF_COUNTER_RAWCOUNT |
| Input WaitForOutBuf | Shows the number of times that a "wait for an available send buffer" was done by the protocols on the client side of the connection. | PERF_COUNTER_RAWCOUNT |
| Input WdBytes | Shows the number of bytes input on this session after all protocol overhead has been removed. | PERF_COUNTER_COUNTER |
| Input WdFrames | Shows the number of frames input after any additional protocol added frames have been removed. | PERF_COUNTER_COUNTER |
| Output Async Frame Error | Shows the number of output async framing errors. These could be caused by a hardware or line problem. | PERF_COUNTER_RAWCOUNT |
| Output Async Overflow | Shows the number of output async overflow errors. | PERF_COUNTER_RAWCOUNT |
| Output Async Overrun | Shows the number of output async overrun errors. | PERF_COUNTER_RAWCOUNT |
| Output Async Parity Error | Shows the number of output async parity errors. These can be caused by a hardware or line problem. | PERF_COUNTER_RAWCOUNT |
| Output Bytes | Shows the number of bytes output on this session, which includes all protocol overhead. | PERF_COUNTER_COUNTER |
| Output Compress Flushes | Shows the number of output compression dictionary flushes. When the data cannot be compressed, the compression dictionary is flushed so that newer data has a better chance of being compressed. One possible cause of data not compressing is transferring compressed files using client drive mapping. | PERF_COUNTER_RAWCOUNT |
| Output Compressed Bytes | Shows the number of bytes output after compression. This number compared with the Total Bytes output is the compression ratio. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Output Compression Ratio | Shows the compression ratio of the server output data stream. Output Compressed Bytes compared with the Total Bytes output is the compression ratio. | PERF_COUNTER_RAWCOUNT |
| Output Errors | Shows the number of output errors of all types. Some examples of output errors are lost ACKs or badly formed packets. | PERF_COUNTER_RAWCOUNT |
| Output Frames | Shows the number of frames that were output on this session. | PERF_COUNTER_COUNTER |
| Output Timeouts | Shows the total number of timeouts on the communication line, from the host side of the connection. These are typically the result of a noisy line. | PERF_COUNTER_RAWCOUNT |
| Output Transport Errors | Shows the number of Terminal Services transport-level errors on output. | PERF_COUNTER_RAWCOUNT |
| Output WaitForOutBuf | Shows the number of times that a "wait for an available send buffer" was done by the protocol on the server side of the connection. | PERF_COUNTER_COUNTER |
| Output WdBytes | Shows the number of bytes output on this session after all protocol overhead has been removed. | PERF_COUNTER_COUNTER |
| Output WdFrames | Shows the number of frames output before any additional protocol frames have been added. | PERF_COUNTER_COUNTER |
| Protocol Bitmap Cache Hit Ratio | Shows the hit ratio in the protocol bitmap cache. A higher hit ratio means better performance, since data transmissions are reduced. Low hit ratios are due to the screen updating with new information that is either not re-used, or is flushed out of the client cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Bitmap Cache Hits | Shows the number of hits in the protocol bitmap cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Bitmap Cache Reads | Shows the number of references to the protocol bitmap cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Brush Cache Hit Ratio | Shows the hit ratio in the protocol brush cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Brush Cache Hits | Shows the number of hits in the protocol brush cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Brush Cache Reads | Shows the number of references to the protocol brush cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Glyph Cache Hit Ratio | Shows the hit ratio in the protocol glyph cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Glyph Cache Hits | Shows the number of hits in the protocol glyph cache. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Protocol Glyph Cache Reads | Shows the number of references to the protocol glyph cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Save Screen Bitmap Cache Hit Ratio | Shows the hit ratio in the protocol save screen bitmap cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Save Screen Bitmap Cache Hits | Shows the number of hits in the protocol save screen bitmap cache. | PERF_COUNTER_RAWCOUNT |
| Protocol Save Screen Bitmap Cache Reads | Shows the number of references to the protocol save screen bitmap cache. | PERF_COUNTER_RAWCOUNT |
| Total Async Frame Error | Shows the total number of async framing errors. | PERF_COUNTER_RAWCOUNT |
| Total Async Overflow | Shows the total number of async overflow errors. | PERF_COUNTER_RAWCOUNT |
| Total Async Overrun | Shows the total number of async overrun errors. | PERF_COUNTER_RAWCOUNT |
| Total Async Parity Error | Shows the total number of async parity errors. | PERF_COUNTER_RAWCOUNT |
| Total Bytes | Shows the total number of bytes on this session, which includes all protocol overhead. | PERF_COUNTER_COUNTER |
| Total Compress Flushes | Shows the total number of compression dictionary flushes. | PERF_COUNTER_RAWCOUNT |
| Total Compressed Bytes | Shows the total number of bytes after compression. This number compared with the Total Bytes output is the compression ratio. | PERF_COUNTER_COUNTER |
| Total Compression Ratio | Shows the total compression ratio of the server data stream. This is the Total Compressed Bytes compared with the Total Bytes output. | PERF_COUNTER_RAWCOUNT |
| Total Errors | Shows the total number of errors of all types. | PERF_COUNTER_RAWCOUNT |
| Total Frames | Shows the total number of frames on this session. | PERF_COUNTER_COUNTER |
| Total Protocol Cache Hit Ratio | Shows the overall hit ratio for all protocol caches. | PERF_COUNTER_RAWCOUNT |
| Total Protocol Cache Hits | Shows the total hits in all protocol caches. The protocol caches Windows objects that are likely to be re-used, to avoid having to resend them on the transmission line. Examples of objects are Windows icons and brushes. Hits in the cache represent objects that did not need to be resent. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Total Protocol Cache Reads | Shows the total references to all protocol caches. | PERF_COUNTER_RAWCOUNT |
| Total Timeouts | Shows the total number of timeouts on the communication line, from both the host and client sides of the connection. | PERF_COUNTER_RAWCOUNT |
| Total Transport Errors | Shows the total number of Terminal Services transport-level errors. | PERF_COUNTER_RAWCOUNT |
| Total WaitForOutBuf | Shows the number of times that a "wait for an available send buffer" was done by the protocols on both the server and client sides of the connection. | PERF_COUNTER_COUNTER |
| Total WdBytes | Shows the total number of bytes on this session after all protocol overhead has been removed. | PERF_COUNTER_COUNTER |
| Total WdFrames | Shows the total number of frames input and output before any additional protocol frames have been added. | PERF_COUNTER_COUNTER |

Terminal Services Object

# Thread Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Thread performance object consists of counters that measure aspects of thread behavior. A thread is the basic object that executes instructions on a processor. All running processes have at least one thread.

| Counter Name | Description | Counter Type |
|---|---|---|
| % Privileged Time | Shows the percentage of non-idle processor time spent executing code in privileged mode. Privileged mode is a processing mode designed for operating system components and hardware-manipulating drivers. It allows direct access to hardware and memory. The alternative, user mode, is a restricted processing mode designed for applications, environmental subsystems, and integral subsystems. The operating system switches application threads to privileged mode to access operating system services.<br><br>% Privileged Time includes time servicing interrupts and DPCs. A high rate of privileged time might be attributed to a large number of interrupts generated by a failing device. When a Windows system service is called, the service often runs in privileged mode to gain access to system-private data. Such data is protected from access by threads executing in user mode. Calls to the system can be explicit or implicit, such as page faults or interrupts.<br><br>Unlike some earlier Windows operating systems, Windows 2000 and the Windows Server 2003 family use process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Windows 2000 and Windows Server 2003 on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process. | PERF_100NSEC_TIMER |
| % Processor Time | Shows the percentage of time that the processor spent executing a non-idle thread. It is calculated by measuring the duration that the idle thread is active during the sample interval, and subtracting that time from 100 %. (Each processor has an idle thread that consumes cycles when no other threads are ready to run.) This counter is the primary indicator of processor activity, and displays the average percentage of busy time observed during the sample interval. Code executed to handle some hardware interrupts and trap conditions are included in this count. | PERF_100NSEC_TIMER |
| % User Time | Shows the percentage of time that the processor spent executing code in user mode. Applications, environment subsystems, and integral subsystems execute in user mode. Code executing in user mode cannot damage the integrity of the Windows Executive, kernel, and/or device drivers. Unlike some earlier Windows operating systems, Windows 2000 and Windows Server 2003 use process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Windows 2000 and Windows Server 2003 on behalf of your application might appear in other subsystem processes in addition to the privileged time in your process. | PERF_100NSEC_TIMER |
| Context Switches/sec | Shows the combined rate, in incidents per second, at which all processors on the computer were switched from one thread to another Context switches occur when a running thread voluntarily relinquishes the processor, or is preempted by a higher priority, ready thread. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Elapsed Time | Shows the time, in seconds, that this process has been running. | PERF_ELAPSED_TIME |
| ID Process | Shows the unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process. | PERF_COUNTER_RAWCOUNT |
| ID Thread | Shows the unique identifier of this thread. ID Thread numbers are reused, so they only identify a thread for the lifetime of that thread. | PERF_COUNTER_RAWCOUNT |
| Priority Base | Shows the current base priority of this process. Threads within a process can raise and lower their own base priority relative to the process's base priority. | PERF_COUNTER_RAWCOUNT |
| Priority Current | Shows the current dynamic priority of this thread. The system can raise the dynamic priority of the thread above the base priority if the thread is handling user input, or lower it toward the base priority if the thread becomes compute bound. | PERF_COUNTER_RAWCOUNT |
| Start Address | Shows the starting virtual address for this thread. | PERF_COUNTER_RAWCOUNT |
| Thread State | Shows the current state of the thread. Valid values are: 0 Initialized 1 Ready 2 Running 3 Standby 4 Terminated 5 Waiting 6 Transition 7 Unknown | PERF_COUNTER_RAWCOUNT |
| Thread Wait Reason | Shows the reason that a thread was in the Wait state. This counter is only applicable when the thread is in the Wait state, which is indicated by a Thread State value of 5. Values for this counter are as follows: 0 Waiting for a component of the Windows NT Executive 1 Waiting for a page to be freed 2 Waiting for a page to be mapped or copied 3 Waiting for space to be allocated in the paged or nonpaged pool 4 Waiting for an Execution Delay to be resolved 5 Suspended 6 Waiting for a user request 7 Waiting for a component of the Windows NT Executive 8 Waiting for a page to be freed 9 Waiting for a page to be mapped or copied 10 Waiting for space to be allocated in the paged or nonpaged pool 11 Waiting for an Execution Delay to be resolved 12 Suspended 13 Waiting for a user request 14 Waiting for an event pair high 15 Waiting for an event pair low 16 Waiting for an LPC Receive notice 17 Waiting for an LPC Reply notice 18 Waiting for virtual memory to be allocated 19 Waiting for a page to be written to disk 20+ (Reserved for future use) | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Thread Details Object
Finding Bottlenecks on Multiprocessor Computers
Find processor bottlenecks

# Thread Details Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Thread Details performance object consists of counters that measure aspects of thread behavior that are difficult or time-consuming to collect. These counters are distinguished from those in the Thread performance object by their high overhead.

This object contains *costly* counters, for which data collection is typically either processor-intensive or memory-intensive. Costly counters cannot be selected with System Monitor but can be accessed using another parsing tool.

| Counter Name | Description | Counter Type |
|---|---|---|
| User PC | Shows the current user program counter for this thread. This is also known as an instruction pointer in CPU registers. | PERF_COUNTER_LARGE_RAWCOUNT_HEX |

**See Also**
**Concepts**
Thread Object
Process Address Space Object
Image Object

# UDP Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The UDP performance object consists of counters that measure the rates at which User Data Protocol (UDP) datagrams are sent and received using the UDP. It includes counters that monitor UDP errors.

| Counter Name | Description | Counter Type |
|---|---|---|
| Datagrams No Port/sec | Shows the rate, in incidents per second, at which UDP datagrams were received for which there was no application at the destination port. | PERF_COUNTER_COUNTER |
| Datagrams Received Errors | Shows the number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port. | PERF_COUNTER_RAWCOUNT |
| Datagrams Received/sec | Shows the rate, in incidents per second, at which UDP datagrams were delivered to UDP users. | PERF_COUNTER_COUNTER |
| Datagrams Sent/sec | Shows the rate, in incidents per second, at which UDP datagrams were sent from the computer. | PERF_COUNTER_COUNTER |
| Datagrams/sec | Shows the rate, in incidents per second, at which UDP datagrams were transferred by the computer. | PERF_COUNTER_COUNTER |

# Web Service Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Web Service performance object counters (installed with Internet Information Services) monitor file transfer rates, bandwidth usage, connection rates, errors, and numbers and types of users.

You can view performance data for all instances of this object (using the _Total instance) or for specific instances, such as the Default Web Site or Administration Web Site instances.

| Counter Name | Description | Counter Type |
|---|---|---|
| Anonymous Users/sec | Shows the rate, in incidents per second, at which users made anonymous connections using the Web service. | PERF_COUNTER_COUNTER |
| Bytes Received/sec | Shows the rate, in incidents per second, at which data bytes were received by the Web service. | PERF_COUNTER_BULK_COUNT |
| Bytes Sent/sec | Shows the rate, in incidents per second, at which data bytes were sent by the Web service. | PERF_COUNTER_BULK_COUNT |
| Bytes Total/sec | Shows the sum of Web Service\Bytes Sent/sec and Web Service\Bytes Received/sec. This is the total rate, in incidents per second, at which bytes were transferred by the Web service. | PERF_COUNTER_BULK_COUNT |
| CGI Requests/sec | Shows the rate, in incidents per second, at which Common Gateway Interface (CGI) requests were simultaneously being processed by the Web service. | PERF_COUNTER_COUNTER |
| Connection Attempts/sec | Shows the rate, in incidents per second, at which connections using the Web service were being attempted. | PERF_COUNTER_COUNTER |
| Copy Requests/sec | Shows the rate, in incidents per second, at which Hypertext Transfer Protocol (HTTP) requests using the COPY method were made. COPY requests are used for copying files and directories. | PERF_COUNTER_COUNTER |
| Current Anonymous Users | Shows the number of users who had an anonymous connection using the Web service. | PERF_COUNTER_RAWCOUNT |
| Current Blocked Bandwidth Bytes | Shows the current total of bytes that were temporarily blocked due to bandwidth throttling settings. | PERF_COUNTER_RAWCOUNT |
| Current CAL count for authenticated users | Shows the current count of licenses used simultaneously by the Web service for authenticated connections. | PERF_COUNTER_RAWCOUNT |
| Current CAL count for SSL connections | Shows the current count of licenses used simultaneously by the Web service for Secure Socket Layer (SSL) connections. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Current CGI Requests | Shows the current number of CGI requests that were simultaneously being processed by the Web service. | PERF_COUNTER_RAWCOUNT |
| Current Connections | Shows the current number of connections established with the Web service. | PERF_COUNTER_RAWCOUNT |
| Current ISAPI Extension Requests | Shows the current number of Internet Server Application Programming Interface (ISAPI) Extension requests that were simultaneously being processed by the Web service. | PERF_COUNTER_RAWCOUNT |
| Current Non-Anonymous Users | Shows the number of users who currently have a non-anonymous connection using the Web service. | PERF_COUNTER_RAWCOUNT |
| Delete Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the DELETE method were made. Delete requests are generally used for file removals. | PERF_COUNTER_COUNTER |
| Files Received/sec | Shows the rate, in incidents per second, at which files were received by the Web service. | PERF_COUNTER_COUNTER |
| Files Sent/sec | Shows the rate, in incidents per second, at which files were sent by the Web service. | PERF_COUNTER_COUNTER |
| Files/sec | Shows the rate, in incidents per second, at which files were transferred, that is, sent and received by the Web service. | PERF_COUNTER_COUNTER |
| Get Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the GET method were made. GET requests are generally used for basic file retrievals or image maps, though they can be used with forms. | PERF_COUNTER_COUNTER |
| Head Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the HEAD method are made. HEAD requests generally indicate a client is querying the state of a document they already have to see if it needs to be refreshed. | PERF_COUNTER_COUNTER |
| ISAPI Extension Requests/sec | Shows the rate, in incidents per second, at which ISAPI Extension Requests were simultaneously being processed by the Web service. | PERF_COUNTER_COUNTER |
| Lock Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the LOCK method were made. LOCK requests are used to lock a file for one user so only that user can modify the file. | PERF_COUNTER_COUNTER |
| Locked Errors/sec | Shows the rate, in incidents per second, of errors due to requests that could not be satisfied by the server because the requested document was locked. These are generally reported as an HTTP 423 error code to the client. | PERF_COUNTER_COUNTER |
| Logon Attempts/sec | Shows the rate, in incidents per second, at which logons using the Web service were being attempted. | PERF_COUNTER_COUNTER |
| Maximum Anonymous Users | Shows the maximum number of users who established concurrent anonymous connections using the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Maximum CAL Count for Authenticated Users | Shows the maximum count of licenses used simultaneously by the Web service for authenticated connections. | PERF_COUNTER_RAWCOUNT |
| Maximum CAL Count for SSL Connections | Shows the maximum count of licenses used simultaneously by the Web service for SSL connections. | PERF_COUNTER_RAWCOUNT |
| Maximum CGI Requests | Shows the maximum number of CGI requests simultaneously processed by the Web service. | PERF_COUNTER_RAWCOUNT |
| Maximum Connections | Shows the maximum number of simultaneous connections established with the Web service. | PERF_COUNTER_RAWCOUNT |
| Maximum ISAPI Extension Requests | Shows the maximum number of extension requests simultaneously processed by the Web service. | PERF_COUNTER_RAWCOUNT |
| Maximum Non-Anonymous Users | Shows the maximum number of users who established concurrent non-anonymous connections using the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Measured Async I/O Bandwidth Usage | Shows the measured bandwidth of asynchronous I/O, averaged over one minute. | PERF_COUNTER_RAWCOUNT |
| Mkcol Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the MKCOL method were made. MKCOL requests are used to create directories on the server. | PERF_COUNTER_COUNTER |
| Move Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the MOVE method were made. MOVE requests are used for moving files and directories. | PERF_COUNTER_COUNTER |
| Non-Anonymous Users/sec | Shows the rate, in incidents per second, at which users were making non-anonymous connections using the Web service. | PERF_COUNTER_COUNTER |
| Not Found Errors/sec | Shows the rate, in incidents per second, of errors due to requests that could not be satisfied by the server because the requested document could not be found. These are generally reported as an HTTP 404 error code to the client. | PERF_COUNTER_COUNTER |
| Options Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the OPTIONS method were made. | PERF_COUNTER_COUNTER |
| Other Request Methods/sec | Shows the rate, in incidents per second, at which HTTP requests were made that do not use the OPTIONS, GET, HEAD, POST, PUT, DELETE, TRACE, MOVE, COPY, MKCOL, PROPFIND, PROPPATCH, MS-SEARCH, LOCK, or UNLOCK methods. These may include LINK or other methods supported by gateway applications. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Post Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the POST method were made. POST requests are generally used for forms or gateway requests. | PERF_COUNTER_COUNTER |
| Propfind Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the PROPFIND method were made. PROPFIND requests retrieve property values on files and directories. | PERF_COUNTER_COUNTER |
| Proppatch Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the PROPPATCH method were made. PROPPATCH requests set property values on files and directories. | PERF_COUNTER_COUNTER |
| Put Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the PUT method were made. The PUT method requests that the document sent to the server be stored under the supplied Internet address. | PERF_COUNTER_COUNTER |
| Search Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the MS-SEARCH method were made. SEARCH requests are used to query the server to find resources that match a set of conditions provided by the client. | PERF_COUNTER_COUNTER |
| Service Uptime | Shows the uptime for the Web service since it was last start started. | PERF_COUNTER_RAWCOUNT |
| Total Anonymous Users | Shows the total number of users who established an anonymous connection with the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Blocked Bandwidth Bytes | Shows the total number of bytes that were temporarily blocked due to bandwidth throttling settings (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Bytes Received | Shows the number of data bytes that were received by the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Bytes Sent | Shows the number of data bytes that were sent by the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Bytes Transferred | Shows the total bytes that were either received or sent by the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total CGI Requests | Shows the total number of CGI requests. CGI requests are custom gateway executable files (.exe) that the administrator can install to add forms processing or other dynamic data sources. CGI requests initiate a process on the server that can be a large drain on server resources. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total Connection Attempts (all instances) | Shows the number of connections that were attempted using the Web service (counted since service startup). This counter is for all instances listed. | PERF_COUNTER_RAWCOUNT |
| Total Copy Requests | Shows the number of HTTP requests using the COPY method (counted since service startup). COPY requests are used for copying files and directories. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Total Count of Failed CAL Requests for Authenticated Users | Shows the number of HTTP requests that failed due to a license being unavailable for an authenticated user. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total count of failed CAL requests for SSL connections | Shows the total count of HTTP requests that failed due to a license being unavailable for SSL connections. | PERF_COUNTER_RAWCOUNT |
| Total Delete Requests | Shows the number of HTTP requests using the DELETE method (counted since service startup). DELETE requests are generally used for file removals. | PERF_COUNTER_RAWCOUNT |
| Total Files Received | Shows the total number of files received by the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Files Sent | Shows the total number of files sent by the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Files Transferred | Shows the sum of Files Sent and Files Received. This is the total number of files transferred by the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Get Requests | Shows the total number of HTTP requests using the GET method (counted since service startup). GET requests are generally used for basic file retrievals or image maps, though they can be used with forms. | PERF_COUNTER_RAWCOUNT |
| Total Head Requests | Shows the total number of HTTP requests using the HEAD method (counted since service startup). HEAD requests generally indicate that a client is querying the state of a document they already have to see if it needs to be refreshed. | PERF_COUNTER_RAWCOUNT |
| Total ISAPI Extension Requests | Shows the total number of ISAPI Extension Requests. ISAPI Extension Requests are custom gateway dynamic-link libraries (DLLs) that the administrator can install to add forms processing or other dynamic data sources. Unlike CGI requests, ISAPI requests are simple calls to a DLL routine, thus they are better suited to high performance gateway applications. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total Lock Requests | Shows the total number of HTTP requests using the LOCK method (counted since service startup). LOCK requests are used to lock a file for one user so that only that user can modify the file. | PERF_COUNTER_RAWCOUNT |
| Total Locked Errors | Shows the total number of requests that could not be satisfied by the server because the requested was locked. These are generally reported as an HTTP 423 error code to the client. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total Logon Attempts | Shows the total number of logons that have been attempted using the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Method Requests | Shows the total number of all HTTP requests (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Method Requests/sec | Shows the rate, in incidents per second, at which all HTTP requests were made. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Total Mkcol Requests | Shows the total number of HTTP requests using the MKCOL method (counted since service startup). MKCOL requests are used to create directories on the server. | PERF_COUNTER_RAWCOUNT |
| Total Move Requests | Shows the total number of HTTP requests using the MOVE method (counted since service startup). MOVE requests are used for moving files and directories. | PERF_COUNTER_RAWCOUNT |
| Total Non-Anonymous Users | Shows the total number of users who established a non-anonymous connection with the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Not Found Errors | Shows the total number of requests that could not be satisfied by the server because the requested document could not be found. These are generally reported as an HTTP 404 error code to the client. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total Options Requests | Shows the total number of HTTP requests using the OPTIONS method (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Other Request Methods | Shows the total number of HTTP requests that are not OPTIONS, GET, HEAD, POST, PUT, DELETE, TRACE, MOVE, COPY, MKCOL, PROPFIND, PROPPATCH, MS-SEARCH, LOCK, or UNLOCK methods (counted since service startup). These may include LINK or other methods supported by gateway applications. | PERF_COUNTER_RAWCOUNT |
| Total Post Requests | Shows the total number of HTTP requests using the POST method (counted since service startup). POST requests are generally used for forms or gateway requests. | PERF_COUNTER_RAWCOUNT |
| Total Propfind Requests | Shows the total number of HTTP requests using the PROPFIND method (counted since service startup). PROPFIND requests retrieve property values on files and directories. | PERF_COUNTER_RAWCOUNT |
| Total Proppatch Requests | Shows the total number of HTTP requests using the PROPPATCH method (counted since service startup). PROPPATCH requests set property values on files and directories. | PERF_COUNTER_RAWCOUNT |
| Total Put Requests | Shows the total number of HTTP requests using the PUT method (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Search Requests | Shows the total number of HTTP requests using the MS-SEARCH method (counted since service startup). SEARCH requests are used to query the server to find resources that match a set of conditions provided by the client. | PERF_COUNTER_RAWCOUNT |
| Total Trace Requests | Shows the total number of HTTP requests using the TRACE method (counted since service startup). TRACE requests allow the client to see what is being received at the end of the request chain and to use the information for diagnostic purposes. | PERF_COUNTER_RAWCOUNT |
| Total Unlock Requests | Shows the total number of HTTP requests using the UNLOCK method (counted since service startup). UNLOCK requests are used to remove locks from files. | PERF_COUNTER_RAWCOUNT |
| Trace Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the TRACE method were made. TRACE requests allow the client to see what is being received at the end of the request chain and to use the information for diagnostic purposes. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Unlock Requests/sec | Shows the rate, in incidents per second, at which HTTP requests using the UNLOCK method were made. UNLOCK requests are used to remove locks from files. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Web Service Cache Object

# Web Service Cache Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Web Service Cache object counters installed with Internet Information Services monitor file transfer rates, bandwidth usage, and connection rates.

You can view performance data for all instances of this object (using the _Total instance) or for specific instances, such as the Default Web Site or Administration Web Site instances.

In addition to these counters, counters supplied by the Internet Information Services Global performance object combine data for the Web service with the performance data they report for IIS as a whole. For more information, see the Internet Information Services Global Object.

| Counter Name | Description | Counter Type |
|---|---|---|
| Anonymous Users/sec | Shows the rate at which users are making anonymous connections using the Web service. | PERF_COUNTER_COUNTER |
| Bytes Received/sec | Showsthe rate that data bytes are received by the Web service. | PERF_COUNTER_BULK_COUNT |
| Bytes Sent/sec | Shows the rate that data bytes are sent by the Web service. | PERF_COUNTER_BULK_COUNT |
| Bytes Total/sec | Shows the sum of Web Service\Bytes Sent/sec and Web Service\Bytes Received/sec. This is the total rate of bytes transferred by the Web service. | PERF_COUNTER_BULK_COUNT |
| CGI Requests/sec | Shows the rate at which Common Gateway Interface (CGI) requests are simultaneously being processed by the Web service. | PERF_COUNTER_COUNTER |
| Connection Attempts/sec | Shows the rate at which connections using the Web service are being attempted. | PERF_COUNTER_COUNTER |
| Copy Requests/sec | Shows the rate at which HTTP requests using the COPY method are made. Copy requests are used for copying files and directories | PERF_COUNTER_COUNTER |
| Current Anonymous Users | Shows the number of users who currently have an anonymous connection using the Web service. | PERF_COUNTER_RAWCOUNT |
| Current Blocked Async I/O Requests | Shows current requests that are temporarily blocked due to bandwidth throttling settings. | PERF_COUNTER_RAWCOUNT |
| Current CAL count for authenticated users | Shows the current count of licenses used simultaneously by the Web service for authenticated connections. | PERF_COUNTER_RAWCOUNT |
| Current CAL count for SSL connections | Shows the current count of licenses used simultaneously by the Web service for SSL connections. | PERF_COUNTER_RAWCOUNT |
| Current CGI Requests | Shows the current number of CGI requests that are simultaneously being processed by the Web service. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Current Connections | Shows the current number of connections established with the Web service. | PERF_COUNTER_RAWCOUNT |
| Current ISAPI Extension Requests | Shows the current number of Extension requests that are simultaneously being processed by the Web service. | PERF_COUNTER_RAWCOUNT |
| Current NonAnonymous Users | Shows the number of users who currently have a non-anonymous connection using the Web service. | PERF_COUNTER_RAWCOUNT |
| Delete Requests/sec | Shows the rate at which HTTP requests using the DELETE method are made. Delete requests are generally used for file removals. | PERF_COUNTER_COUNTER |
| Files Received/sec | Shows the rate at which files are received by the Web service. | PERF_COUNTER_COUNTER |
| Files Sent/sec | Shows the rate at which files are sent by the Web service. | PERF_COUNTER_COUNTER |
| Files/sec | Shows the rate at which files are transferred, that is, sent and received by the Web service. | PERF_COUNTER_COUNTER |
| Get Requests/sec | Shows the rate at which HTTP requests using the GET method are made. Get requests are generally used for basic file retrievals or image maps, though they can be used with forms. | PERF_COUNTER_COUNTER |
| Head Requests/sec | Shows the rate at which HTTP requests using the HEAD method are made. Head requests generally indicate a client is querying the state of a document they already have to see if it needs to be refreshed. | PERF_COUNTER_COUNTER |
| ISAPI Extension Requests/sec | Shows the rate at which ISAPI Extension requests are simultaneously being processed by the Web service. | PERF_COUNTER_COUNTER |
| Lock Requests/sec | Shows the rate at which HTTP requests using the LOCK method are made. Lock requests are used to lock a file for one user so that only that user can modify the file. | PERF_COUNTER_COUNTER |
| Locked Errors/sec | Shows the rate of errors due to requests that couldn't be satisfied by the server because the requested document was locked. These are generally reported as an HTTP 423 error code to the client. | PERF_COUNTER_COUNTER |
| Logon Attempts/sec | Shows the rate at which logons using the Web service are being attempted. | PERF_COUNTER_COUNTER |
| Maximum Anonymous Users | Shows the maximum number of users who established concurrent anonymous connections using the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Maximum CAL count for authenticated users | Shows the maximum count of licenses used simultaneously by the Web service for authenticated connections. | PERF_COUNTER_RAWCOUNT |
| Maximum CAL count for SSL connections | Shows the maximum count of licenses used simultaneously by the Web service for SSL connections. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Maximum CGI Requests | Shows the maximum number of CGI requests simultaneously processed by the Web service. | PERF_COUNTER_RAWCOUNT |
| Maximum Connections | Shows the maximum number of simultaneous connections established with the Web service. | PERF_COUNTER_RAWCOUNT |
| Maximum ISAPI Extension Requests | Shows the maximum number of Extension requests simultaneously processed by the Web service. | PERF_COUNTER_RAWCOUNT |
| Maximum NonAnonymous Users | Shows the maximum number of users who established concurrent non-anonymous connections using the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Measured Async I/O Bandwidth Usage | Shows the measured bandwidth of asynchronous I/O averaged over a minute. | PERF_COUNTER_RAWCOUNT |
| Mkcol Requests/sec | Shows the rate at which HTTP requests using the MKCOL method are made. MKCOL requests are used to create directories on the server. | PERF_COUNTER_COUNTER |
| Move Requests/sec | Shows the rate at which HTTP requests using the MOVE method are made. MOVE requests are used for moving files and directories. | PERF_COUNTER_COUNTER |
| NonAnonymous Users/sec | Shows the rate at which users are making nonanonymous connections using the Web service. | PERF_COUNTER_COUNTER |
| Not Found Errors/sec | Shows the rate of errors due to requests that couldn't be satisfied by the server because the requested document could not be found. These are generally reported as an HTTP 404 error code to the client. | PERF_COUNTER_COUNTER |
| Options Requests/sec | Shows the rate at which HTTP requests using the OPTIONS method are made. | PERF_COUNTER_COUNTER |
| Other Request Methods/sec | Shows the rate at which HTTP requests are made that do not use the OPTIONS, GET, HEAD, POST, PUT, DELETE, TRACE, MOVE, COPY, MKCOL, PROPFIND, PROPPATCH, MS-SEARCH, LOCK or UNLOCK methods. These may include LINK or other methods supported by gateway applications. | PERF_COUNTER_COUNTER |
| Post Requests/sec | Shows the rate at which HTTP requests using the POST method are made. POST requests are generally used for forms or gateway requests. | PERF_COUNTER_COUNTER |
| Propfind Requests/sec | Shows the rate at which HTTP requests using the PROPFIND method are made. PROPFIND requests retrieve property values on files and directories. | PERF_COUNTER_COUNTER |
| Proppatch Requests/sec | Shows the rate at which HTTP requests using the PROPPATCH method are made. PROPPATCH requests set property values on files and directories. | PERF_COUNTER_COUNTER |
| Put Requests/sec | Shows the rate at which HTTP requests using the PUT method are made. The PUT method requests that the document sent to the server be stored under the supplied Internet address. | PERF_COUNTER_COUNTER |
| Search Requests/sec | Shows the rate at which HTTP requests using the MS-SEARCH method are made. Search requests are used to query the server to find resources that match a set of conditions provided by the client. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Total Allowed Async I/O Requests | Shows the total number of requests allowed by bandwidth throttling settings (counted since service startup). When bandwidth throttling is defined, the number of allowed user requests is restricted. | PERF_COUNTER_RAWCOUNT |
| Total Anonymous Users | Shows the total number of users who established an anonymous connection with the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Blocked Async I/O Requests | Shows the number of total requests temporarily blocked due to bandwidth throttling settings (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total CGI Requests | Shows the total number of CGI requests. CGI requests are custom gateway executables (.exe) the administrator can install to add forms processing or other dynamic data sources. CGI requests initiate a process on the server which can be a large drain on server resources. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total Connection Attempts (all instances) | Shows the number of connections that have been attempted using the Web service (counted since service startup). This counter is for all instances listed. | PERF_COUNTER_RAWCOUNT |
| Total Copy Requests | Shows the number of HTTP requests using the COPY method (counted since service startup). COPY requests are used for copying files and directories | PERF_COUNTER_RAWCOUNT |
| Total count of failed CAL requests for authenticated users | Shows the number of HTTP requests that failed due to a license being unavailable for an authenticated user. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total count of failed CAL requests for SSL connections | Shows the total count of HTTP requests that failed due to a license being unavailable for SSL connections. | PERF_COUNTER_RAWCOUNT |
| Total Delete Requests | Shows the number of HTTP requests using the DELETE method (counted since service startup). Delete requests are generally used for file removals. | PERF_COUNTER_RAWCOUNT |
| Total Files Received | Shows the total number of files received by the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Files Sent | Shows the total number of files sent by the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Files Transferred | Shows the sum of Files Sent and Files Received. This is the total number of files transferred by the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Get Requests | Shows the number of HTTP requests using the GET method (counted since service startup). GET requests are generally used for basic file retrievals or image maps, though they can be used with forms. | PERF_COUNTER_RAWCOUNT |
| Total Head Requests | Shows the number of HTTP requests using the HEAD method (counted since service startup). Head requests generally indicate a client is querying the state of a document they already have to see if it needs to be refreshed. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Total ISAPI Extension Requests | Shows the total number of ISAPI Extension Requests. ISAPI Extension Requests are custom gateway Dynamic Link Libraries (.dll) the administrator can install to add forms processing or other dynamic data sources. Unlike CGI requests, ISAPI requests are simple calls to a DLL library routine, thus they are better suited to high performance gateway applications. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total Lock Requests | Shows the number of HTTP requests using the LOCK method (counted since service startup). LOCK requests are used to lock a file for one user so that only that user can modify the file. | PERF_COUNTER_RAWCOUNT |
| Total Locked Errors | Shows the number of requests that couldn't be satisfied by the server because the requested was locked. These are generally reported as an HTTP 423 error code to the client. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total Logon Attempts | Shows the number of logons that have been attempted using the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Method Requests | Shows the number of all HTTP requests (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Method Requests/sec | Shows the rate at which all HTTP requests are made. | PERF_COUNTER_COUNTER |
| Total Mkcol Requests | Shows the number of HTTP requests using the MKCOL method (counted since service startup). MKCOL requests are used to create directories on the server. | PERF_COUNTER_RAWCOUNT |
| Total Move Requests | Shows the number of HTTP requests using the MOVE method (counted since service startup). MOVE requests are used for moving files and directories. | PERF_COUNTER_RAWCOUNT |
| Total NonAnonymous Users | Shows the total number of users who established a non-anonymous connection with the Web service (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Not Found Errors | Shows the number of requests that couldn't be satisfied by the server because the requested document could not be found. These are generally reported as an HTTP 404 error code to the client. The count is the total since service startup. | PERF_COUNTER_RAWCOUNT |
| Total Options Requests | Shows the number of HTTP requests using the OPTIONS method (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Other Request Methods | Shows the number of HTTP requests that are not OPTIONS, GET, HEAD, POST, PUT, DELETE, TRACE, MOVE, COPY, MKCOL, PROPFIND, PROPPATCH, MS-SEARCH, LOCK or UNLOCK methods (counted since service startup). These may include LINK or other methods supported by gateway applications. | PERF_COUNTER_RAWCOUNT |
| Total Post Requests | Shows the number of HTTP requests using the POST method (counted since service startup). POST requests are generally used for forms or gateway requests. | PERF_COUNTER_RAWCOUNT |
| Total Propfind Requests | Shows the number of HTTP requests using the PROPFIND method (counted since service startup). PROPFIND requests retrieve property values on files and directories. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Total Proppatch Requests | Shows the number of HTTP requests using the PROPPATCH method (counted since service startup). PROPPATCH requests set property values on files and directories. | PERF_COUNTER_RAWCOUNT |
| Total Put Requests | Shows the number of HTTP requests using the PUT method (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Rejected Async I/O Requests | Shows the total number of requests rejected due to bandwidth throttling settings (counted since service startup). | PERF_COUNTER_RAWCOUNT |
| Total Search Requests | Shows the number of HTTP requests using the MS-SEARCH method (counted since service startup). Search requests are used to query the server to find resources that match a set of conditions provided by the client. | PERF_COUNTER_RAWCOUNT |
| Total Trace Requests | Shows the number of HTTP requests using the TRACE method (counted since service startup). TRACE requests allow the client to see what is being received at the end of the request chain and use the information for diagnostic purposes. | PERF_COUNTER_RAWCOUNT |
| Total Unlock Requests | Shows the number of HTTP requests using the UNLOCK method (counted since service startup). UNLOCK requests are used to remove locks from files. | PERF_COUNTER_RAWCOUNT |
| Trace Requests/sec | Shows the rate at which HTTP requests using the TRACE method are made. TRACE requests allow the client to see what is being received at the end of the request chain and use the information for diagnostic purposes. | PERF_COUNTER_COUNTER |
| Unlock Requests/sec | Shows the rate at which HTTP requests using the UNLOCK method are made. UNLOCK requests are used to remove locks from files. | PERF_COUNTER_COUNTER |

# Windows Media Station Service Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Media Station Service performance object consists of counters that provide statistics about the Windows Media Station service, which provides multicasting, distribution, and storage functions for Advanced Systems Format (ASF) streams.

| Counter Name | Description | Counter Type |
|---|---|---|
| Stations | Shows the number of station objects that currently exist on the server. | PERF_COUNTER_RAWCOUNT |
| Controllers | Shows the number of controllers that are currently connected to the server. | PERF_COUNTER_RAWCOUNT |
| Streams | Shows the number of stream objects that currently exist on the server. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Windows Media Unicast Service Object

# Windows Media Unicast Service Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Media Unicast Service performance object consists of counters that provide statistics about the Windows Media Unicast service, which provides unicasting functions for Advanced Systems Format (ASF) streams.

| Counter Name | Description | Counter Type |
|---|---|---|
| Active Live Unicast Streams | Shows the number of live unicast streams that are being streamed. | PERF_COUNTER_RAWCOUNT |
| Active Streams | Shows the number of active files, or streams, currently being sent to the instances of Windows Media Player. This counter provides an overall indicator of system activity. A client is still connected once it has stopped playing a file, because the Windows Media Unicast service maintains the bandwidth allocation to that client for 60 seconds after the client stops playing (this is called the Client Inactivity Time-out period). After the Client Inactivity Time-out period has elapsed, the server actively checks the client's status by periodically sending a message to each client. The server keeps a client connection open as long as the client responds to a message within three minutes (this is called the Client Terminal Time-out). | PERF_COUNTER_RAWCOUNT |
| Active TCP Streams | Shows the number of active Transmission Control Protocol (TCP) streams that are currently being sent to the instances of Windows Media Player. | PERF_COUNTER_RAWCOUNT |
| Active UDP Streams | Shows the number of active User Datagram Protocol (UDP) streams currently being sent to the instances of Windows Media Player. | PERF_COUNTER_RAWCOUNT |
| Aggregate Read Rate | Shows the sum, in bytes per second, of the read rates for all files being sent to the instances of Windows Media Player. This counter indicates the speed at which the server is reading data from the disk. The graph usually appears jagged because the block read time is shorter than the block sends. | PERF_COUNTER_BULK_COUNT |
| Aggregate Send Rate | Shows the sum, in bytes per second, of the send rates for all files being sent to the instances of Windows Media Player. This counter indicates the speed at which the server is sending data to the players. The graph must appear relatively smooth; a jagged graph can indicate that streams are battling for bandwidth. | PERF_COUNTER_BULK_COUNT |
| Allocated Bandwidth | Shows the amount of bandwidth that the server allocates, based on the number of instances of Windows Media Player that are currently connected. Generally, clients connected to Windows Media Services that are not receiving streams are still allocated bandwidth. If a Windows Media Player has had bandwidth allocated to it, but does not send a command during the Client Inactivity Time-out period, its bandwidth is deallocated. The Client Inactivity Time-out period is 60 seconds. | PERF_COUNTER_RAWCOUNT |
| Authentication Requests | Shows the number of times an authentication plug-in is called to authenticate a client. Depending on the protocol in use — either HTTP or Microsoft Media Server (MMS) — a single client may be authenticated once or several times as it accesses each piece of unicast content. This is a general indicator of how often the server uses authentication plug-ins. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Authentications Denied | Shows the number of times that clients are denied access to unicast content by authentication plug-ins. Clients are denied access by an authentication plug-in when their IDs cannot be verified, no user accounts exist, or the authentication plug-in is not working. | PERF_COUNTER_COUNTER |
| Authorization Requests | Shows the number of times that an authorization plug-in is called to authorize a client. Each client action—such as opening or playing unicast content—is separately authorized, so the total number of authorization requests is greater than the number of connected clients. This is a general indicator of how often the server uses authorization plug-ins. | PERF_COUNTER_COUNTER |
| Authorizations Refused | Shows the number of times that clients are denied access to unicast content by authorization plug-ins. Plug-ins deny access to clients based on their function (for example, a billing plug-in may deny client access to unicast content if the user account contains insufficient funds). | PERF_COUNTER_COUNTER |
| Connected Clients | Shows the number of client programs that are currently connected. This indicates all instances of Windows Media Player that are connected to the system. This counter includes clients that are playing streamed content, clients in paused and connected states, and clients that the server has not yet detected as being disconnected. The latter can occur if a client has lost power, or has been disconnected from the network and has not been able to notify the server that it has been disconnected. | PERF_COUNTER_RAWCOUNT |
| Connection Rate | Shows the rate at which clients are connecting to the server. This is used to determine the number of instantaneous connections being serviced by the server. The default is 10 connections per second. This counter is useful for correlating client connections with the utilization of system resources. | PERF_COUNTER_COUNTER |
| HTTP Streams | Shows the number of active HTTP streams that are currently being sent to instances of Windows Media Player. This counter reports the sum of the values of HTTP Streams Reading Body and HTTP Streams Reading Header. | PERF_COUNTER_RAWCOUNT |
| HTTP Streams Reading Body | Shows the number of connected clients that have read the file header and are actively streaming unicast content via HTTP. | PERF_COUNTER_RAWCOUNT |
| HTTP Streams Reading Header | Shows the number of clients that are connected to the server via HTTP and are reading a file header, but are not yet streaming the unicast content. | PERF_COUNTER_RAWCOUNT |
| Late Reads | Shows the number of late read completions per second. A *late read* is a disk read operation that takes significantly longer than expected to complete. Typically, this counter is nonzero only when the server is under a very heavy load. If this counter is often nonzero, the system might not be properly configured (for example, you might need additional hard disk drives or a lower setting for the maximum number of clients), or other applications might be competing with Windows Media Services for use of the disk. | PERF_COUNTER_COUNTER |
| Pending Connections | Shows the number of clients that are attempting to connect to the server but are not yet connected. This number may be high if the server is running near maximum capacity and cannot process a large number of connection requests in a timely manner. | PERF_COUNTER_RAWCOUNT |
| Plugin Errors | Shows the number of times that the plug-ins failed to perform their function (for example, a plug-in could not reach its user database, or the plug-in crashed). | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Plugin Events | Shows the number of events reported to plug-ins. This is an indicator of how often the server uses external plug-ins. | PERF_COUNTER_COUNTER |
| Scheduling Rate | Shows the rate, in seconds, at which requests are being made for tasks to be scheduled. Scheduled tasks include reading data from the disk and writing data to the network. The scheduling rate is roughly proportional to the rate at which packets are being sent by the server; it does not represent a one-to-one correspondence. | PERF_COUNTER_COUNTER |
| Stream Errors | Shows the number of errors that occur per second. Stream errors represent the number of stream data packets that are discarded by the server. They are introduced by the server when it cannot meet the demand for data and must throw packets away to avoid running behind schedule indefinitely. Stream errors show up most often after late reads occur. They indicate that the system is not configured properly, or other applications are competing for system resources. Heavy network traffic can also contribute to stream errors. | PERF_COUNTER_COUNTER |
| Stream Terminations | Shows the rate at which streams are terminated due to errors. Stream terminations indicate that the system is not configured properly, or the file being streamed is corrupted. The server unilaterally terminates a stream under certain circumstances — most commonly, when the server is running significantly behind schedule. If the server must throw away several packets in a row from a given stream to remain on schedule, it terminates the stream. Streams can also be terminated if invalid data is encountered during file streaming. | PERF_COUNTER_COUNTER |
| UDP Resend Requests | Shows the number of times clients ask the Windows Media Services server to re-send data packets that were not received. This value may be high when the server cannot reliably send packets via UDP. This is a good indicator of server or network overload. | PERF_COUNTER_COUNTER |
| UDP Resends Sent | Shows the number of UDP re-send requests that are processed by the Windows Media Services server. If the server is under a heavy load, not all UDP re-send requests are processed. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Windows Media Station Service Object

# WINS Server Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The WINS Server performance object reports the activity of a computer that is running Windows Internet Name Service (WINS). This object's counters are cleared each time you start and stop the service. WINS provides a dynamic database service that can register and resolve NetBIOS names to the IP addresses that are used on your network. The Windows Server 2003 family provides WINS, which enables the server computer to act as a NetBIOS name server and to register and resolve names for WINS-enabled client computers on your network, as described in the NetBIOS over TCP/IP standards.

| Counter Name | Description | Counter Type |
|---|---|---|
| Unique Registrations/sec | Shows the rate, in incidents per second, at which unique registrations were received by the server. | PERF_COUNTER_COUNTER |
| Group Registrations/sec | Shows the rate, in incidents per second, at which group registrations were received by the server. | PERF_COUNTER_COUNTER |
| Total Number of Registrations/sec | Shows the total rate, in incidents per second, at which registrations were received by the server. This rate is the sum of the unique and group registration rates. | PERF_COUNTER_COUNTER |
| Unique Renewals/sec | Shows the rate, in incidents per second, at which unique renewals were received by the server. | PERF_COUNTER_COUNTER |
| Group Renewals/sec | Shows the rate, in incidents per second, at which group renewals were received by the server. | PERF_COUNTER_COUNTER |
| Total Number of Renewals/sec | Shows the total rate, in incidents per second, at which renewals were received by the server. This rate is the sum of the unique and group renewal rates. | PERF_COUNTER_COUNTER |
| Releases/sec | Shows the rate, in incidents per second, at which releases were received by the server. | PERF_COUNTER_COUNTER |
| Queries/sec | Shows the rate, in incidents per second, at which queries were received by the server. | PERF_COUNTER_COUNTER |
| Unique Conflicts/sec | Shows the rate, in incidents per second, at which unique registrations or renewals that were received by the server resulted in conflicts with records in the database. | PERF_COUNTER_COUNTER |
| Group Conflicts/sec | Shows the rate, in incidents per second, at which group registrations that were received by the server resulted in conflicts with records in the database. | PERF_COUNTER_COUNTER |
| Total Number of Conflicts/sec | Shows the total rate, in incidents per second, at which conflicts were received by the server. This rate is the sum of the unique and group conflict rates. | PERF_COUNTER_COUNTER |
| Successful Releases/sec | Shows the rate, in incidents per second, of successful releases. | PERF_COUNTER_COUNTER |
| Failed Releases/sec | Shows the rate, in incidents per second, of failed releases. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Successful Queries/sec | Shows the rate, in incidents per second, of successful queries. | PERF_COUNTER_COUNTER |
| Failed Queries/sec | Shows the rate, in incidents per second, of failed queries. | PERF_COUNTER_COUNTER |

# WMI Objects Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The WMI Objects performance object consists of counters that describe the activity of Windows Management Instrumentation (WMI) objects. WMI has a reverse adapter that makes WMI performance data visible in System Monitor. WMI performance data only appears when the reverse adapter for WMI is installed.

| **Counter Name** | **Description** | **Counter Type** |
|---|---|---|
| HiPerfClasses | Shows the number of whole classes that the reverse adapter obtained from WMI. | PERF_COUNTER_LARGE_RAWCOUNT |
| HiPerfValidity | Shows the validity of shared memory held in the performance structure. A value of 1 indicates that data is valid and a value of 0 indicates that data is invalid. | PERF_COUNTER_RAWCOUNT |

# WSRM Policy Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The WSRM Policy performance object, installed with Windows System Resource Manager (WSRM), consists of a counter that provides the percentage of CPU cycles of the entire system that are used by the current resource-allocation policy, since the time the policy was set to manage the behavior of WSRM. (The resource-allocation policy is an ordered list of resource allocations that specifies the management behavior of WSRM at a given time.) Each WSRM Policy object will have one instance whose name is the same as the current active policy.

| Counter Name | Description | Counter Type |
|---|---|---|
| System Managed CPU % | The percentage of CPU cycles managed by WSRM that are used by the combined process-matching criteria. A process-matching criterion is the mechanism in WSRM that matches running processes to a resource-allocation policy. | PERF_SAMPLE_FRACTION |

**See Also**
**Concepts**
Counters by Feature

# WSRM Process Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The WSRM Process performance object, installed with Windows System Resource Manager (WSRM), consists of a counter that provides the percentage of total managed CPU cycles used by a particular process, since the time that the resource-allocation policy was set to manage the behavior of WSRM. (The resource-allocation policy is an ordered list of resource allocations that specifies the management behavior of WSRM at a given time.) One instance of this object exists for every managed process. Note that Total Managed CPU cycles is the sum of the CPU cycles for all the managed processes.

| Counter Name | Description | Counter Type |
|---|---|---|
| Actual Managed CPU % | The percentage of total managed CPU cycles that have been used by this process-matching criterion. A process-matching criterion is the mechanism in WSRM that matches running processes to a resource-allocation policy. | PERF_RAW_FRACTION |

**See Also**
**Concepts**
Counters by Feature

# WSRM Process Matching Criteria Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The WSRM process-matching criteria performance object, installed with Windows System Resource Manager (WSRM), consists of counters that monitor the CPU usage and memory usage of the processes matched by the process-matching criteria. The criteria are included in the managing resource-allocation policy. Each object will have as many instances as the number of process-matching criteria within the current active policy.

| Counter Name | Description | Counter Type |
|---|---|---|
| Actual Managed CPU % | The percentage of total managed CPU cycles that have been used by this process-matching criterion. A process-matching criterion is the mechanism in WSRM that matches running processes to a resource-allocation policy. | PERF_RAW_FRACTION |
| Managed CPU % | The percentage of managed CPU cycles used by this process-matching criterion between the last process samplings. | PERF_SAMPLE_FRACTION |
| Page File (MB) | The total size of page files used by this process. | PERF_COUNTER_RAWCOUNT |
| Process Count | The number of processes associated with this process-matching criterion. | PERF_COUNTER_RAWCOUNT |
| Target Managed CPU % | The percentage of total managed CPU cycles allocated to this process-matching criterion. | PERF_RAW_FRACTION |
| Working Set (MB) | The total size of the set of memory pages used by the threads of this process. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Counters by Feature

# ACS/RSVP Interfaces Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

ACS/RSVP Interfaces performance object counters report statistics about Resource Reservation Protocol (RSVP) activity, flow, bandwidth, RSVP messages sent, and so on for a specific interface or IP address. The ACS/RSVP Interfaces object is available if QoS Admission Control Service (QoS ACS) or RSVP is running and there is one object for each IP address on the system.

| Counter Name | Description | Counter Type |
|---|---|---|
| Admitted Bandwidth | Shows the total admitted bandwidth on this interface, in bits per second. | PERF_COUNTER_RAWCOUNT |
| Blocked RESVs | Shows the number of blockaded reservations that failed due to failed reservations that have not been removed from this node. For more information about blockade states, see RFC 2205. | PERF_COUNTER_RAWCOUNT |
| General Failures | Shows the total number of non-admission control errors that have occurred on this interface. For more information about other traffic control failures, see RFC 2205. | PERF_COUNTER_RAWCOUNT |
| Maximum Admitted Bandwidth | Shows the highest amount of allocated bandwidth admitted on the interface during the period in which QoS ACS or RSVP has been running. | PERF_COUNTER_RAWCOUNT |
| Number of Active Flows | Shows the number of reserved flows that were installed and maintained by RSVP on this interface. | PERF_COUNTER_RAWCOUNT |
| Number of Incoming Messages Dropped | Shows the number of incoming or received messages that were dropped due to syntax errors, protocol violations, and other errors. | PERF_COUNTER_RAWCOUNT |
| Number of Outgoing Messages Dropped | Shows the number of outgoing messages that were dropped due to errors while sending. | PERF_COUNTER_RAWCOUNT |
| PATH ERR Messages Received | Shows the total number of PATH ERR messages that were received on this interface. This number indicates the number of times that a PATH has been rejected by a remote host. | PERF_COUNTER_RAWCOUNT |
| PATH ERR Messages Sent | Shows the total number of PATH ERR rejection messages that were sent by the QoS ACS or RSVP service. The PATH error messages indicate an error in response to the PATH message. | PERF_COUNTER_RAWCOUNT |
| PATH Messages Received | Shows the total number of PATH messages that have been received on this interface. | PERF_COUNTER_RAWCOUNT |
| PATH Messages Sent | Shows the total number of PATH messages that were sent by the QoS ACS or RSVP service on this the interface. | PERF_COUNTER_RAWCOUNT |
| PATH State Block Timeouts | Shows the total number of times that the RSVP PATH state timed out. This happens when no PATH messages were received within the designated timeout period. This error message can be caused by a route change. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| PATH TEAR Messages Received | Shows the total number of PATH TEAR messages that have been received on the interface. A PATH TEAR message removes the PATH state along the route and in this host. | PERF_COUNTER_RAWCOUNT |
| PATH TEAR Messages Sent | Shows the total number of PATH TEAR messages that were sent on the interface to tear down an RSVP PATH state in the network. | PERF_COUNTER_RAWCOUNT |
| Policy Control Failures | Shows the total number of RSVP requests that were not admitted on this interface due to a conflict with policy settings. | PERF_COUNTER_RAWCOUNT |
| Receive Messages Errors - Big Messages | Shows the total number of incoming RSVP messages that were dropped because they were bigger than the maximum configured RSVP message size. | PERF_COUNTER_RAWCOUNT |
| Receive Messages Errors - No Memory | Shows the total number of incoming RSVP messages that were dropped due to insufficient memory. | PERF_COUNTER_RAWCOUNT |
| Resource Control Failures | Shows the number of reservations that failed due to the lack of resources. | PERF_COUNTER_RAWCOUNT |
| RESV CONFIRM Messages Received | Shows the total number of RESV CONFIRM messages that have been received on this interface. | PERF_COUNTER_RAWCOUNT |
| RESV CONFIRM Messages Sent | Shows the total number of RESV CONFIRM messages that were sent on this interface in response to RESV messages that were sent by remote hosts. | PERF_COUNTER_RAWCOUNT |
| RESV ERR Messages Received | Shows the total number of RESV ERR messages that were received on this interface by the QoS ACS or RSVP service, rejecting incoming RESV messages. An RESV ERR message indicates an error in response to the RESV. | PERF_COUNTER_RAWCOUNT |
| RESV ERR Messages Sent | Shows the total number of RESV errors that were sent on this interface. This number indicates how many reservations have been rejected by the network. | PERF_COUNTER_RAWCOUNT |
| RESV Messages Received | Shows the total number of RESV messages that have been received on this interface. | PERF_COUNTER_RAWCOUNT |
| RESV Messages Sent | Shows the total number of RESV messages that were sent by the ACS/RSVP service on this the interface. | PERF_COUNTER_RAWCOUNT |
| RESV State Block Timeouts | Shows the total number of times that the RSVP RESV state timed out. This happens when no RESV messages are received within the designated timeout period. This error message can be caused by a route change. | PERF_COUNTER_RAWCOUNT |
| RESV TEAR Messages Received | Shows the total number of RESV TEAR messages that have been received on the interface. An RESV TEAR message removes the reservation along the route and in this host. | PERF_COUNTER_RAWCOUNT |
| RESV TEAR Messages Sent | Shows the total number of RESV TEAR messages that were sent on the interface to tear down an RESV state in the network. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Send Messages Errors - Big Messages | Shows the total number of outgoing RSVP messages that were dropped because they were bigger than the maximum configured RSVP message size. | PERF_COUNTER_RAWCOUNT |
| Send Messages Errors - No Memory | Shows the total number of outgoing RSVP messages that were dropped due to insufficient memory. | PERF_COUNTER_RAWCOUNT |
| Signaling Bytes Received | Shows the total, in bytes, of the RSVP signaling traffic that has been received by the ACS/RSVP service on this interface. | PERF_COUNTER_RAWCOUNT |
| Signaling Bytes Sent | Shows the total, in bytes, of RSVP signaling traffic that was sent by the ACS/RSVP service on this interface. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Counters by Feature

# ACS/RSVP Service Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The ACS/RSVP Service performance object reports statistics about the Resource Reservation Protocol (RSVP) service and about local applications that use Quality of Service (QoS). The ACS/RSVP Service object is available even if the RSVP service is not running.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes in QoS Notifications | Shows a running total of the number of bytes that were delivered in QoS notifications to applications that use QoS. | PERF_COUNTER_RAWCOUNT |
| Failed QoS Requests | Shows the number of QoS requests that were generated by applications that use QoS and that were rejected by the ACS/RSVP service. Failed QoS requests can be caused by invalid QoS requests. | PERF_COUNTER_RAWCOUNT |
| Failed QoS Sends | Shows the number of QoS notifications from the ACS/RSVP service that could not be sent to the QoS applications. Failed QoS sends can be caused by terminated applications. | PERF_COUNTER_RAWCOUNT |
| Network Interfaces | Shows the number of local network interfaces (IP addresses) that are visible to, and used by, the RSVP service. | PERF_COUNTER_RAWCOUNT |
| Network Sockets | Shows the total number of raw sockets that were opened for the purpose of RSVP signaling. | PERF_COUNTER_RAWCOUNT |
| QoS notifications | Shows the number of QoS notifications that were delivered to applications that use QoS by the QoS RSVP service. | PERF_COUNTER_RAWCOUNT |
| QoS Sockets | Shows the number of sockets that use QoS. | PERF_COUNTER_RAWCOUNT |
| QoS-Enabled Receivers | Shows the number of RESV messages sent for receivers that use QoS. | PERF_COUNTER_RAWCOUNT |
| QoS-Enabled Senders | Shows the number of PATH messages sent for senders that use QoS. | PERF_COUNTER_RAWCOUNT |
| RSVP Sessions | Shows the current number of active RSVP sessions on the ACS/RSVP service. | PERF_COUNTER_RAWCOUNT |
| Timers | Shows the number of timer events that are scheduled to take place. This counter shows the activity level of the ACS/RSVP service. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
Windows Server 2003 System Performance Counters
Counters by Feature

# Gateway Service for NetWare Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Gateway Service for NetWare performance object consists of counters that measure packet transmission rates and activity of the Gateway for NetWare service.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Received/sec | Shows the rate of bytes coming in to the Redirector from the network. It includes all application data in addition to network protocol information (such as packet headers.) | PERF_COUNTER_BULK_COUNT |
| Bytes Total/sec | Shows the rate the Redirector is processing data bytes. It includes all application and file data in addition to protocol information. | PERF_COUNTER_BULK_COUNT |
| Bytes Transmitted/sec | Shows the rate at which bytes are leaving the Redirector to the network. It includes all application in addition to network protocol information. | PERF_COUNTER_BULK_COUNT |
| Connect NetWare 2.x | Shows the number of connections to NetWare 2.x servers. | PERF_COUNTER_RAWCOUNT |
| Connect NetWare 3.x | Shows the number of connections to NetWare 3.x servers. | PERF_COUNTER_RAWCOUNT |
| Connect NetWare 4.x | Shows the number of connections to NetWare 4.x servers. | PERF_COUNTER_RAWCOUNT |
| File Data Operations/sec | Shows the rate at which the Redirector is processing data operations. One operation includes many bytes but can be reduced because each packet has protocol overhead. You can determine the efficiency of this path by dividing the value of Bytes/sec by this counter to determine the average number of bytes transferred per operation. | PERF_COUNTER_COUNTER |
| File Read Operations/sec | Shows the rate at which applications are asking the Redirector for data. Each call to a file system or similar API call counts as one operation. | PERF_COUNTER_COUNTER |
| File Write Operations/sec | Shows the rate at which applications are sending data to the Redirector. Each call to a file system or similar API call counts as one operation. | PERF_COUNTER_COUNTER |
| Packet Burst IO/sec | Shows the sum of the values of Gateway Services for NetWare\Packet Burst Read NCPs/sec and Gateway Services for NetWare\Packet Burst Write NCPs/sec. | PERF_COUNTER_COUNTER |
| Packet Burst Read NCP Count/sec | Shows the rate of NetWare Core Protocol requests for Packet Burst Read. Packet Burst is a windowing protocol that improves performance. | PERF_COUNTER_COUNTER |
| Packet Burst Read Timeouts/sec | Shows the rate the NetWare Service needs to retransmit a Burst Read Request because the NetWare server took too long to respond. | PERF_COUNTER_COUNTER |
| Packet Burst Write NCP Count/sec | Shows the rate of NetWare Core Protocol requests for Packet Burst Write. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Packet Burst Write Timeouts/sec | Shows the rate at which the NetWare Service needs to retransmit a Burst Write Request because the NetWare server took too long to respond. | PERF_COUNTER_COUNTER |
| Packets Received/sec | Shows the rate at which the Redirector is receiving packets (SMBs). Network transmissions are divided into packets. The average number of bytes received in a packet can be obtained by dividing the value of Gateway Services for NetWare\Bytes Received/sec by this counter. Some packets received might not contain incoming data, for example an acknowledgment to a write made by the Redirector would count as an incoming packet. | PERF_COUNTER_BULK_COUNT |
| Packets Transmitted/sec | Shows the rate at which the Redirector is sending packets. Network transmissions are divided into packets. The average number of bytes transmitted in a packet can be obtained by dividing the value of Bytes Transmitted/sec by this counter. | PERF_COUNTER_BULK_COUNT |
| Packets/sec | Shows the rate the Redirector is processing data packets. One packet typically includes many bytes but can be reduced because each packet has protocol overhead. You can determine the efficiency of this path by dividing the value of Gateway Services for NetWare\Bytes Total/sec by this counter to determine the average number of bytes transferred/packet. You can also divide this counter by the value of Operations/sec to determine the average number of packets per operation. | PERF_COUNTER_BULK_COUNT |
| Read Operations Random/sec | Shows the rate at which, on a file-by-file basis, nonsequential reads are made. If a read is made using a particular file handle, and then is followed by another read that is not the contiguous next byte, this counter is incremented by one. | PERF_COUNTER_COUNTER |
| Read Packets/sec | Shows the rate at which read packets are being placed on the network. Each time a single packet is sent with a request to read data remotely, this counter is incremented by one. | PERF_COUNTER_COUNTER |
| Server Disconnects | Shows the number of times that a server has disconnected your Redirector. See also Server Reconnects. | PERF_COUNTER_RAWCOUNT |
| Server Reconnects | Shows the number of times the Redirector has had to reconnect to a server to complete a new active request. You can be disconnected by the Server if you remain inactive for too long. Locally even if all your remote files are closed, the Redirector will keep your connections intact for ten minutes. Such inactive connections are called Dormant Connections. Reconnecting is expensive in time. | PERF_COUNTER_RAWCOUNT |
| Server Sessions | Shows the total number of security objects that the Redirector has managed. For example, a logon to a server followed by a network access to the same server will establish one connection, but two sessions. | PERF_COUNTER_RAWCOUNT |
| Write Operations Random/sec | Shows the rate at which, on a file-by-file basis, nonsequential writes are made. | PERF_COUNTER_COUNTER |
| Write Packets/sec | Shows the rate at which writes are being sent to the network. Each time a single packet is sent with a request to write remote data, this counter is incremented by one. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Internet Information Services Global Object

# NetBEUI Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The NetBEUI performance object consists of counters that measure data transmission for network activity which conforms to the NetBIOS End User Interface (NetBEUI) standard.

| Counter Name | Description | Counter Type |
|---|---|---|
| Bytes Total/sec | Shows the total rate of bytes sent to or received from the network by the protocol, but only counts the bytes in frames (packets) which carry data. This is the sum of the values of Frame Bytes/sec and Datagram Bytes/sec. | PERF_COUNTER_BULK_COUNT |
| Connection Session Timeouts | Shows the number of connections that were dropped due to a session time out. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| | Shows the number of connections that were canceled. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Connections No Retries | Shows the total count of connections that were successfully made on the first try. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Connections Open | Shows the number of connections currently open for this protocol. This counter shows the current count only and does not accumulate over time. | PERF_COUNTER_RAWCOUNT |
| Connections With Retries | Shows the total count of connections that were made after retrying the attempt. A retry occurs when the first connection attempt failed. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Datagram Bytes Received/sec | Shows the rate at which datagram bytes are received by the computer. A datagram is a connectionless packet whose delivery to a remote computer is not guaranteed. | PERF_COUNTER_BULK_COUNT |
| Datagram Bytes Sent/sec | Shows the rate at which datagram bytes are sent from the computer. | PERF_COUNTER_BULK_COUNT |
| Datagram Bytes/sec | Shows the rate at which datagram bytes are processed by the computer. This counter measures datagram bytes that are both received and sent. | PERF_COUNTER_BULK_COUNT |
| Datagrams Received/sec | Shows the rate at which datagrams are received by the computer. | PERF_COUNTER_COUNTER |
| Datagrams Sent/sec | Shows the rate at which datagrams are sent from the computer. | PERF_COUNTER_COUNTER |
| Datagrams/sec | Shows the rate at which datagrams are processed by the computer. This counter measures datagrams that are both received and sent. | PERF_COUNTER_COUNTER |
| Disconnects Local | Shows the number of session disconnections that were initiated by the local computer. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |

| | | |
|---|---|---|
| Disconnects Remote | Shows the number of session disconnections that were initiated by the remote computer. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Expirations Ack | Shows the count of T2 timer expirations. | PERF_COUNTER_RAWCOUNT |
| Expirations Response | Shows the count of T1 timer expirations. | PERF_COUNTER_RAWCOUNT |
| Failures Adapter | Shows the number of connections that were dropped due to an adapter failure. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures Link | Shows the number of connections that were dropped due to a link failure. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures No Listen | Shows the number of connections that were rejected because the remote computer was not listening for connection requests. | PERF_COUNTER_RAWCOUNT |
| Failures Not Found | Shows the number of connection attempts that failed because the remote computer could not be found. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures Resource Local | Shows the number of connections that failed because of resource problems or shortages on the local computer. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Failures Resource Remote | Shows the number of connections that failed because of resource problems or shortages on the remote computer. This counter is an accumulator and shows a running total. | PERF_COUNTER_RAWCOUNT |
| Frame Bytes Received/sec | Shows the rate at which data bytes are received by the computer. This counter only shows the frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Rejected/sec | Shows the rate at which data bytes are rejected. This counter only shows the bytes in data frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Re-Sent/sec | Shows the rate at which data bytes are re-sent by the computer. This counter only shows the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes Sent/sec | Shows the rate at which data bytes are sent by the computer. This counter only shows the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frame Bytes/sec | Shows the rate at which data bytes are processed by the computer. This counter is the sum of the values of Frame Bytes Received/sec and Frame Bytes Sent/sec. This counter only shows the bytes in frames that carry data. | PERF_COUNTER_BULK_COUNT |
| Frames Received/sec | Shows the rate at which data frames are received by the computer. This counter only shows the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames Rejected/sec | Shows the rate at which data frames are rejected. This counter only shows the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames Re-Sent/sec | Shows the rate at which data frames are re-sent by the computer. This counter only shows the frames that carry data. | PERF_COUNTER_COUNTER |

| | | |
|---|---|---|
| Frames Sent/sec | Shows the rate at which data frames are sent by the computer. This counter only shows the frames that carry data. | PERF_COUNTER_COUNTER |
| Frames/sec | Shows the rate at which data frames are processed by the computer. This counter is the sum of the values of Frames Received/sec and Frames Sent/sec and only shows the frames that carry data. | PERF_COUNTER_COUNTER |
| Packets Received/sec | Shows the rate at which packets are received by the computer. This counter shows all packets processed (control as well as data packets). | PERF_COUNTER_COUNTER |
| Packets Sent/sec | Shows the rate at which packets are sent by the computer. This counter shows all packets sent by the computer (control as well as data packets). | PERF_COUNTER_COUNTER |
| Packets/sec | Shows the rate at which packets are processed by the computer. This count is the sum of the values of Packets Sent/sec and Packets Received/sec and includes all packets processed (control as well as data packets). | PERF_COUNTER_COUNTER |
| Piggyback Ack Queued/sec | Shows the rate at which piggyback acknowledgments are queued. Piggyback acknowledgments are acknowledgments of received packets that are to be included in the next outgoing packet to the remote computer. | PERF_COUNTER_COUNTER |
| Piggyback Ack Timeouts | Shows the number of times that a piggyback acknowledgment could not be sent because there was no outgoing packet to the remote computer on which to piggyback. If no outgoing packet is sent within the time-out period, then an acknowledgment packet is sent and this counter is incremented. | PERF_COUNTER_RAWCOUNT |
| Window Send Average | Shows the average number of data bytes that were sent before waiting for an acknowledgment from the remote computer. | PERF_COUNTER_RAWCOUNT |
| Window Send Maximum | Shows the maximum number of data bytes that are sent before waiting for an acknowledgment from the remote computer. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
NetBEUI Resource Object

# NetBEUI Resource Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The NetBEUI Resource performance object consists of counters that track the use of buffers by the NetBEUI protocol.

| Counter Name | Description | Counter Type |
|---|---|---|
| Times Exhausted | Shows the number of times that all resources (buffers) were in use. The number in parentheses following the resource name is used to identify the resource in Event Log messages. | PERF_COUNTER_RAWCOUNT |
| Used Average | Shows the current number of resources (buffers) in use. The number in parentheses following the resource name is used to identify the resource in Event Log messages. | PERF_COUNTER_RAWCOUNT |
| Used Maximum | Shows the maximum number of NetBEUI resources (buffers) in use at any point in time. This value is useful in sizing the maximum resources provided. The number in parentheses following the resource name is used to identify the resource in Event Log messages. | PERF_COUNTER_RAWCOUNT |

**See Also**
**Concepts**
NetBEUI Object

# Network Segment Object

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Network Segment performance object provides network statistics for the local network segment; it is installed by the Network Monitor driver protocol on Windows 2000 Professional and by Network Monitor Tools on Windows 2000 Server.

| Counter Name | Description | Counter Type |
|---|---|---|
| % Broadcast Frames | Shows the percentage of network bandwidth that is made up of broadcast traffic on this network segment. | PERF_AVERAGE_BULK |
| % Multicast Frames | Shows the percentage of network bandwidth that is made up of multicast traffic on this network segment. | PERF_AVERAGE_BULK |
| % Network Utilization | %Shows the percentage of network bandwidth in use on this network segment. | PERF_COUNTER_MULTI_TIMER |
| Broadcast Frames Received/sec | Shows the number of broadcast frames received per second on this network segment. | PERF_COUNTER_COUNTER |
| Multicast Frames Received/sec | Shows the number of multicast frames received per second on this network segment. | PERF_COUNTER_COUNTER |
| Total Bytes Received/sec | Shows the number of bytes received per second on this network segment. | PERF_COUNTER_COUNTER |
| Total Frames Received/sec | Shows the total number of frames received per second on this network segment. | PERF_COUNTER_COUNTER |

**See Also**
**Concepts**
Network Monitor

# Windows Server 2003 System Performance Counters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following counters are standard in the Windows Server 2003 operating system, are installed by Setup by default. Note that some may only be visible if you install or start the associated service.

- ACS/RSVP Service Object
- Browser Object
- Cache Object
- Distributed Transaction Coordinator Object
- Http Indexing Service Object
- IAS Accounting Clients Object
- IAS Accounting Server Object
- IAS Authentication Clients Object
- IAS Authentication Server Object
- ICMP Object
- Indexing Service Object
- Indexing Service Filter Object
- IP Object
- Job Object Object
- Job Object Details
- LogicalDisk Object
- Memory Object
- NBT Connection Object
- DUPTILE1_Network Interface Object
- Objects Object
- Paging File Object
- Physical Disk Object
- Print Queue Object
- Process Object

- Processor Object
- Redirector Object
- Server Object
- Server Work Queues Object
- System Object
- TCP Object
- Telephony Object
- Thread Object
- UDP Object

In addition, some counters are considered "costly," meaning that it takes a long time to collect data from them. These counters do not appear in System Monitor and can only be viewed using Showperf.exe or other performance-related tools in the Windows Server 2003 Technical Reference. The following are examples of costly counters:

Image Object

Process Address Space Object

Thread Details Object

For information about bottlenecks in some of these subsystem components, see the following topics:

Find memory bottlenecks

Find processor bottlenecks

Finding Cache Bottlenecks

Find disk bottlenecks

When monitoring multiprocessor systems, see the following topic:

Finding Bottlenecks on Multiprocessor Computers

# Bottleneck-Detection Counters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some counters are useful for identifying bottlenecks, performance problems that affect the operating system or installed applications. A bottleneck exists when a single component is adversely affecting the performance of the whole system. This section lists counters to watch when trying to diagnose or prevent performance problems.

- Find memory bottlenecks
- Find processor bottlenecks
- Find disk bottlenecks
- Finding Cache Bottlenecks
- Finding Bottlenecks on Multiprocessor Computers
- Counters for Finding Network Bottlenecks

# Find memory bottlenecks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following counters to monitor the performance of physical and virtual memory on a computer that is running one of the Windows Server 2003 family operating systems:

- Memory Object
- Memory Object
- Memory Object
- Memory Object
- Memory Object
- Memory Object
- Memory Object
- Process Object
- Process Object
- Process Object
- Process Object

# Find processor bottlenecks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following counters to monitor the performance of one or more processors on a computer that is running one of the Windows Server 2003 family operating systems:

- Processor Object
- System Object
- Process Object
- Process Object
- Process Object
- Process Object
- Thread Object
- Thread Object
- Thread Object
- Thread Object
- Thread Object
- Thread Object
- Thread Object

For details about performance monitoring for multiprocessor computers, see Finding Bottlenecks on Multiprocessor Computers.

# Find disk bottlenecks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following counters to monitor the performance of physical disks and logical partitions of physical disks on a computer that is running one of the Windows Server 2003 family operating systems:

- Physical Disk Object
- Physical Disk Object
- Physical Disk Object
- Physical Disk Object
- Physical Disk Object
- Physical Disk Object
- Physical Disk Object
- Physical Disk Object
- Physical Disk Object
- Physical Disk Object
- LogicalDisk Object
- LogicalDisk Object
- LogicalDisk Object
- LogicalDisk Object
- LogicalDisk Object
- LogicalDisk Object
- LogicalDisk Object
- LogicalDisk Object
- LogicalDisk Object
- LogicalDisk Object
- LogicalDisk Object

# Finding Bottlenecks on Multiprocessor Computers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following counters to monitor the performance of multiprocessor computers that are running the Windows Server 2003 family:

- System Object
- Processor Object
- Process Object
- Thread Object

# Finding Cache Bottlenecks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following counters to monitor the performance of the file system cache on a computer that is running one of the Windows Server 2003 family systems.

- Memory Object
- Memory Object
- Memory Object
- Memory Object
- Memory Object
- Cache Object
- Cache Object
- Cache Object
- Cache Object
- Cache Object
- Cache Object
- Physical Disk Object
- Physical Disk Object

# Counters for Finding Network Bottlenecks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following counters to identify bottlenecked network resources:

- DUPTILE1_Network Interface Object
- DUPTILE1_Network Interface Object
- DUPTILE1_Network Interface Object
- *Protocol_layer_object*\Segments Received/sec
- *Protocol_layer_object*\Segments Sent/sec
- *Protocol_layer_object*\Frames Sent/sec
- *Protocol_layer_object*\Frames Received/sec

  *Protocol_layer_object* can be TCP Object, UDP Object, NetBEUI Object, NWLink IPX Object, NWLink NetBIOS Object, NWLink SPX Object, or other protocol layer performance objects. For NWLink performance objects, frame-related counters report only zeroes. Use datagram-based counters for these objects.

- Server Object
- Server Object
- Server Object

# Windows Server 2003 Resource Kit Registry Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Welcome to the Resource Kit Registry Reference for Windows Server 2003.

The Registry Reference is a searchable technical reference that provides detailed descriptions of selected Windows Server 2003 registry content. It describes the subtrees, keys, subkeys, and entries that advanced users want to know about, including many entries that cannot be changed by using Windows tools or programming interfaces. It also contains essential background information about the Windows registry, as well as methods of backing up and restoring the registry, and tips on how to edit the registry safely.

**Caution**

- Directly editing the registry bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you are certain that you must edit the registry, first back it up by using one of the methods described in Backing Up and Restoring the Registry. After you make changes, be sure to test them before returning your computer to normal operation.

For more information, see one of the following:

- HKEY_CURRENT_CONFIG
- HKEY_CLASSES_ROOT
- HKEY_CURRENT_USER
- HKEY_LOCAL_MACHINE
- SYSTEM
- About the Registry Reference

# Glossary of Registry Terms

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**

**Active Window Tracking**
A window is activated (and optionally brought to the foreground) when the user moves the mouse over it.

**asynchronous transfer mode (ATM)**
A high-speed, connection-oriented, virtual circuit-based packet switching protocol used to transport many different types of network traffic. ATM packages data in 53-byte, fixed-length cells that can be switched quickly between logical connections on a network.

**authentication package**
Dynamic-link libraries (DLLs) loaded by the Local Security Authority (LSA) and used for local logon and network authentication. MSV1_0, the default authentication package for Windows 2000, Windows XP, and the Windows Server 2003 family, identifies users by their user names and validates them by their passwords.

**B**

**black hole router (TCP/IP)**
A router that does not return ICMP Destination Unreachable messages when it needs to fragment a TCP packet on which the Don't-Fragment bit is set. TCP must receive Destination Unreachable messages to perform Path MTU Discovery.

**browse list**
Any list of items that can be browsed, such as a list of servers on a network, or a list of printers displayed in the Add Printer wizard.

**C**

**code page**
A means of providing support for character sets and keyboard layouts for different countries or regions. A code page is a table that relates the binary character codes used by a program to keys on the keyboard or to characters on the display.

**context-sensitive**
Menus, buttons, and help text are those that are associated with a particular point (or context) in a window or dialog box. In Windows, when you right-click a context-senstive point, a menu appears displaying options for that context, or a "What's this" button appears which, when pressed, displays information about the context you clicked. In many dialog boxes, you can display context-sensitive help, that is, information about the context point, by clicking the question mark button on the caption bar and then clicking the context point.

**control set**
A complete set of the configuration data required to start devices and system services. The system always maintains at least two control sets in the registry. It identifies one as CurrentControlSet and the other as the LastKnownGood control set. The current control set is the one that was most recently used to start the system. The system stores its control set designations in the Select subkey.

**cryptographic key cache**
A reserved memory space where Message Queuing saves recently used encrypted symmetric keys so that the keys can be used again. Reusing keys makes encrypting and decrypting messages much faster. The cryptographic key cache has two sections, one for keys found in incoming messages and another for keys that are generated for outgoing messages. The size of each section can be adjusted by using the CryptSendKeyCacheSize and CryptReceiveKeyCacheSize entries.

**current control set**
The control set that was used to start the system this time. Only a successfully started system has a current control set. At startup, the system selects from among the numbered control sets stored in the registry. The number of the control set selected to be the current one appears in the **Current** entry in the HKLM\SYSTEM\Select subkey. The current control set is stored in the HKLM\SYSTEM\CurrentControlSet subkey.

See also: control set

**D**

**deadlock**
A situation in which a thread will not relinquish its exclusive access to a critical section.

**E**

**enumerator**
The means by which Windows learns about a device. Enumerator is a concept is used in Windows Plug and Play. It is a property of a device.

**F**

**failover time**
The amount of time it takes a resource, either individually or in a group, to complete the failover process.

**file name completion**
A quick search feature of the Windows command processor, Cmd.exe.

- When you type a path and file name and then press the file name completion character, the command processor searches for all files whose path and file name match your entry.

- When you type a path and then press the directory name completion character, the command processor searches for all files whose path (but not necessarily the file name) matches your entry.

The command processor displays one of the matching file names each time you type a completion character. To go backward through the list of file names, press SHIFT + *completion character*.

**file system cache**
An area of physical memory that holds frequently used pages. It allows applications and services to locate pages rapidly and reduces disk activity.

**foreground process / background process**
A process moves to the foreground when a user interacts with it; otherwise, it is a background process. This distinction lets the operating system optimize the performance of processes that are most important to the user. The system often assigns higher a processor time priority to foreground processes. As a result, the threads of foreground processes can preempt (interrupt) a running thread of a lower priority process. The threads of foreground processes often receive more processor time each time they run.

**G**

**global catalog**
A directory database that applications and clients can query to locate any object in a forest. The global catalog is hosted on one or more domain controllers in the forest. It contains a partial replica of every domain directory partition in the forest. These partial replicas include replicas of every object in the forest, as follows: the attributes most frequently used in search operations and the attributes required to locate a full replica of the object.

In Microsoft Provisioning System, the Exchange server maintains a list of global catalogs, and it maintains a load balance across global catalogs.

**globally unique identifier (GUID)**
A 16-byte value generated from the unique identifier on a device, the current date and time, and a sequence number. A GUID is used to identify a particular device or component.

**H**

**hardware profile**
Data that describes the configuration and characteristics of specific computer equipment. This information can be used to configure computers for using peripheral devices.

**hot keys**
A Windows feature that allows quick activation of specified accessibility features through a combination of keys pressed in unison.

**hot tracking**
When the cursor moves over an item, highlighting it but not selecting it.

**hubbed mode**
A mode in which the Address Resolution Protocol/multicast address resolution service (ARP/MARS) provides asynchronous transfer mode (ATM) addresses to requesting clients in the form of a multicast server (MCS) list value. In this mode, ARP/MARS acts as a multicast server, providing active forwarding of all multicast and broadcast traffic destined for IP addresses contained within the ranges specified in the list.

See also: asynchronous transfer mode (ATM); multicast address resolution service (MARS); nonhubbed mode

**I**

**ILMI**
See definition for: integrated local management interface (ILMI)

**Input Method Editor (IME)**
A program used to enter the thousands of different characters in written Asian languages with a standard 101-key keyboard. An IME consists of both an engine that converts keystrokes into phonetic and ideographic characters and a dictionary of commonly used ideographic words. As the user enters keystrokes, the IME engine attempts to identify which character or characters that the keystrokes should be converted into.

**integrated local management interface (ILMI)**
A limited set of Simple Network Management Protocol (SNMP) functions included in the asynchronous transfer mode (ATM) specification for the ATM user network interface (UNI).

See also: asynchronous transfer mode (ATM); Simple Network Management Protocol (SNMP)

**interface (TCP/IP)**
In TCP/IP, a set of communications functions that permit core elements of the TCP/IP protocol suite to communicate with TCP/IP services. The most commonly known interfaces, the Transport Driver Interface (TDI) and the Network Device Interface (NDIS) are public, and their specifications are available from Microsoft. In addition, there are a number of higher level interfaces available to user-mode applications. The most commonly used are Windows Sockets, RPC, and NetBIOS.

**IPX/SPX**

**iteration**
See definition for: iterative query

**iterative query**
A query made to a DNS server for the best answer the server can provide without seeking further help from other DNS servers. Also called a *nonrecursive query*.

**J**
There are no glossary terms that begin with this letter.

**K**

**key**
In Registry Editor, a folder that appears in the left pane of the Registry Editor window. A key can contain subkeys and entries. For example, Environment is a key of **HKEY_CURRENT_USER**.

In IP security (IPSec), a value used in combination with an algorithm to encrypt or decrypt data. Key settings for IPSec are configurable to provide greater security.

See also: registry

**L**

**Language ID**
A 16-bit value which consists of a primary and a secondary language ID. Bits 0-9 make up the primary language ID and bits 10-15 comprise the sub-language ID.

**LastKnownGood control set**
The control set used most recently to start the system successfully. When the system identifies the LastKnownGood control set, such as ControlSet001, it stores the number of the control set in the **LastKnownGood** entry in the HKLM\SYSTEM\Select subkey.

See also: control set

**Layer Two Tunneling Protocol (L2TP)**
An industry-standard Internet tunneling protocol that provides encapsulation for sending Point-to-Point Protocol (PPP) frames across packet-oriented media. For IP networks, L2TP traffic is sent as User Datagram Protocol (UDP) messages. In Microsoft operating systems, L2TP is used in conjunction with Internet Protocol security (IPSec) as a virtual private network (VPN) technology to provide remote access or router-to-router VPN connections. L2TP is described in RFC 2661.

**Locale**
A group of user preference information, such as date and time formatting, represented as a list of values related to the user's language and sublanguage.

**Locale ID**
A 32-bit value which consists of the language ID in the low word (bits 0-15) and the sorting ID (bits 16-19) and a reserved value (bits 20-31) in the high word.

See also: Language ID

**M**

**multicast address resolution service (MARS)**
A service for resolving multicast IP addresses to the asynchronous transfer mode (ATM) addresses of the clients that have joined that multicast group. MARS can work in conjunction with the multicast server (MCS) and clients to distribute multicast data through point-to-multipoint connections.

See also: asynchronous transfer mode (ATM)

**N**

**NDIS**
See definition for: Network Driver Interface Specification (NDIS)

**Network Driver Interface Specification (NDIS)**
A Microsoft/3Com specification establishing a common shared interface for Microsoft operating systems to support protocol-independent transport of multiple network transport protocols (such as TCP/IP, NetBEUI, IPX/SPX, and AppleTalk). NDIS allows more than one transport protocol to be bound and to operate simultaneously over a single network adapter.

**nonhubbed mode**
A mode in which asynchronous transfer mode (ATM) Address Resolution Protocol/multicast address resolution service (ARP/MARS) does not forward multicast and broadcast traffic for multicast group clients. In this mode, the service returns a dynamic listing of ATM hosts currently registered for the multicast group address to requesting clients. Clients then use this list to initiate and establish their own point-to-multipoint virtual connections with each of the members in the multicast list.

See also: asynchronous transfer mode (ATM); hubbed mode; multicast address resolution service (MARS)

**nonpaged pool**
Operating system memory that is never paged to disk. Paging is the moving of infrequently used parts of a program's working memory from RAM to another storage medium, usually the hard disk. In Task Manager, the amount of memory used by a process, in kilobytes.

See also: Task Manager

**O**
There are no glossary terms that begin with this letter.

**P**

**paged pool**
The system-allocated virtual memory that has been charged to a process and that can be paged. Paging is the moving of infrequently used parts of a program's working memory from random access memory (RAM) to another storage medium, usually the hard disk.

In Task Manager, the amount of system-allocated virtual memory, in kilobytes, used by a process.

See also: Task Manager

**pipe**

A section of shared memory used for interprocess communication. A program can write to a pipe and a program in another process can read the entry and reply. Processes that create pipes are commonly known as pipe servers and processes that read from the pipe are known as pipe clients.

**pod**
A group of UNIX computers that are designated to receive password synchronization updates from a Windows-based server. When a password changes, the change is sent to one UNIX host in each pod. The UNIX computers that comprise the pod are listed in the value of Hosts.

**pool leaks**
Application errors that occur when an application allocates, but does not free, pool space.

**published printer**
A printer that appears in Active Directory. The list of published printers is accessible by using the Add Printer Wizard.

**Q**
There are no glossary terms that begin with this letter.

**R**

**raster fonts**
Fonts that are stored as bitmaps. Raster fonts are designed with a specific size and resolution for a specific printer and cannot be scaled or rotated. If a printer does not support raster fonts, it will not print them. The five raster fonts are Courier, MS Sans Serif, MS Serif, Small, and Symbol. Also called *bit-mapped fonts*.

**recursive resolution**
One of the two process types (iterative and recursive) for DNS name resolution. In this process, a resolver (a DNS client) will request that a DNS server provide a complete answer to a query that does not include pointers to other DNS servers. When a client makes a query and requests that the server use recursive resolution to answer, it effectively shifts the workload of resolving the query from the client to the DNS server. If the DNS server supports and uses recursive resolution, it contacts other DNS servers as necessary (using iterative queries on behalf of the client) until it obtains a definitive answer to the query. This type of resolution allows the client resolver to be small and simple.

See also: iteration; iterative query

**referenced entry (TCP/IP ARP cache table)**
An ARP cache table entry is considered to be referenced when an outbound IP packet is sent to the address described by the entry.

**registry**
A database repository for information about a computer's configuration. The registry contains information that Windows continually references during operation, such as:

- Profiles for each user
- The programs installed on the computer and the types of documents that each can create
- Property settings for folders and program icons
- What hardware exists on the system
- Which ports are being used

The registry is organized hierarchically as a tree, and it is made up of keys and their subkeys, hives, and entries.

See also: key

**S**

**scan code**
A code number transmitted to an IBM or compatible computer when a key is pressed or released. Each key on the keyboard has a unique scan code.

**service group**

A group of similar services that are loaded together at startup. Most services that appear in the HKLM\SYSTEM\CurrentControlSet\Services subkey are part of a service group. The system loads one service group at a time. Services that are not in a group are loaded after all service groups are loaded.

**Simple Network Management Protocol (SNMP)**
A network protocol used to manage TCP/IP networks. In Windows, the SNMP service is used to provide status information about a host on a TCP/IP network.

**site**
One or more well-connected (highly reliable and fast) TCP/IP subnets. A site allows administrators to configure Active Directory access and replication topology to take advantage of the physical network.

**SNMP**
See definition for: Simple Network Management Protocol (SNMP)

**sorting ID**
A code which dictates the correct order for sorting the characters of a language. The sorting ID is used in conjunction with a locale ID to distinguish between default sorts and alternate sorts for locales which have more than one sorting order.

**subtree**
Any node within a tree, along with any selection of connected descendant nodes.

The highest level of the registry (for example, **HKEY_LOCAL_MACHINE**).

See also: key; registry

**SYN flooding**
A malicious attempt to impair a server that occurs when a computer sends the server numerous Transmission Control Protocol (TCP) connection requests that contain fictitious IP addresses. As a result, the resources of the server are occupied by futile attempts to send acknowledgments to the fictitious IP addresses. The term *SYN* refers to the *synchronize sequence number* message that is used to initialize a TCP connection. Also known as a *SYN attack*.

See also: SYN flooding attack protection

**SYN flooding attack protection**
A Transmission Control Protocol (TCP) feature that detects symptoms of SYN flooding and responds by reducing the resources (memory and time) a server spends on connection requests that it cannot acknowledge.

See also: SYN flooding

**T**

**Task Manager**
A tool that provides information about programs and processes running on the computer. Using Task Manager, you can end or run programs, end processes, and display a dynamic overview of your computer's performance.

**twip**
A unit of measure used in typesetting and desktop publishing, equal to one-twentieth of a printer's point (1/1440 inch). The twip notation begins with a dash (-), which, in this context, indicates a twip unit, not a negative value.

**U**

**UNC**
See definition for: Universal Naming Convention (UNC)

**UNC provider**
Interprets names in Universal Naming Convention format.

**UNI**
See definition for: user-to-network interface

**Universal Naming Convention (UNC)**
A convention for naming files and other resources beginning with two backslashes (\), indicating that the resource exists on a network computer. UNC names conform to the *servername**sharename* syntax, where *servername* is the server's name and *sharename* is the name of the shared resource. The UNC name of a directory or file can also include the directory path after the share name, by using the following syntax: *servername**sharename**directory**filename*.

**user-to-network interface**
The interface and signaling protocol used to establish connections between ATM users or end stations and an ATM switch or network. The ATM Call Manager component of Windows ATM Services support UNI.

**V**

**virtual machine**
A program that provides an independent operating system environment within another operating system. A virtual machine permits the user to run programs that are native to a different operating system.

**W**
There are no glossary terms that begin with this letter.

**X**
There are no glossary terms that begin with this letter.

**Y**
There are no glossary terms that begin with this letter.

**Z**
There are no glossary terms that begin with this letter.

# About the Registry Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Registry Reference describes registry content that is useful to administrators and support personnel. It is not a complete listing of all registry contents. The reference documents many registry entries that you can only change by editing the registry, and also documents many entries that should be changed by using standard user interfaces. The reference also describes entries that you can view, but that you should not change for any reason.

**For more information, see the following topics:**

- How to Use the Registry Reference
- Registry Reference Document Conventions

# How to Use the Registry Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Registry Reference describes selected elements of the registry that is installed by default with Windows Server 2003, and to a limited extent, Windows XP Professional. In addition to the default content, this reference also describes subkeys and entries that you can add to your registry, and ones that the system adds when you install optional services (such as hardware drivers, software drivers, or network services). This reference is not a complete listing of all registry contents. It does not document subkeys or entries for applications (such as Microsoft Word), nor does it describe the registry editor itself, Regedit.exe.

**How to Find Registry Elements**

The Registry Reference is designed to mimic the structure of Regedit.exe, the registry editor included in Windows Server 2003. Registry elements appear in the Registry Reference in their registry path, just as they do in Regedit.exe.

To find the description of a registry element, use the **Search** tab or click the books and pages in the table of contents in the left pane. A description of each element appears in the right pane. For example, to find a description of the DNS service subkey, in the left pane click **HKEY_LOCAL_MACHINE**, click **SYSTEM**, click **CurrentControlSet**, click **Services**, and click **DNS**, just as you do in Regedit.exe. Or, on the **Search** tab, type **DNS** in the Search box, select **Search titles only**, and press ENTER.

**Generic Subkeys**

This reference includes generic subkeys that do not appear in the registry, but that represent a group of subkeys with common characteristics. You can identify these generic subkeys by the brackets surrounding their names. For example, the entries that are common to most service subkeys appear in the <Service-name> subkey. If you are looking for a common entry and do not find it in the subkey in which it appears in your registry, look for a generic subkey.

**Limitations of this Reference**

If you cannot find a registry element, it might be because it is not included in this reference. The Registry Reference does not describe all elements that might appear in the registry. Some registry elements are specific to your registry. Others are part of the Windows internal code and are not meaningful to users. Many entries just store values you have set in Windows. The Registry Reference primarily documents entries that advanced users might need to change but that cannot be changed by using the Windows interface or the administrative tools provided with Windows.

Although we make every effort to be timely and accurate, the registry is part of the Windows internal program code and is therefore subject to change at any time. Possible discrepancies include changes in the locations of the entries, the default values, the range of possible values, the use of the data, and the content of entries. For updates and corrections, see the Windows Deployment and Resource Kits Web Site link below.

# Registry Reference Document Conventions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Document Conventions**
Each registry element documented in the Registry Reference has a page that includes following details: the name of the registry element, its registry path, and a description of the element. Optionally, the page might also include a change method, activation method, notes, tips, cautions, or cross-references to related entries.

**Element name and Registry path**
Element names are not unique within the registry, but are unique within their path. The *registry path* shows the location of the element in the registry. You can use the path to navigate to the element.

**Data Type, Range, and Default Value**
For registry entries only, the page includes a table with the data type, range, and default value.

| Data type | Range | Default |
|---|---|---|
| Defines the length and type of data that can be stored in the entry. See Data Types in the Registry. | Demonstrates the syntax or format of the data in the entry and provides the range of valid values. For example, **0 \| 4** means that "0" and "4" are the only values permitted. Words that describe a range, but that are not part of the range, appear in italics. For example, *IP address*. Italics are also used to show the data units. For example, 1–60 *seconds*. | The default is the value of an entry that has not been changed. If the entry does not necessarily appear in the registry, the default shows the value that describes how the system behaves without the registry entry. For example, if the system behaves as though the value is 0 if the entry is not in the registry, then the default value of an entry is 0. |

**Description**
For registry folders (DfOfSubtree, DfOfKey, and DfOfSubkey), the description explains what kinds of data are stored in the folder. It also explains which programs or services create or use the data in the folder.

For DfOfEntry, the description explains how the program or service interprets the value of the entry. If the entry has a range of values, then the description includes a table describing the effect of each value on the program or service.

**Change Method**
The change method describes the graphical user interface (GUI) or command-line method you should use to modify the element. Use these recommended methods whenever possible. If there is no change method described, then the element can be changed only by editing the registry directly.

**Activation Method**
The activation method describes additional steps that you must take to make changes to the element effective. Many changes are effective immediately if you make them by using the standard interfaces, rather than by editing the registry. However, to make some changes effective, you must restart a service, restart Windows, or log off and log on again.

**Note**
Notes provide additional information about the element. Typically, notes explain when the system adds an entry to the registry, qualifies the use of an entry, or describes the order of precedence of related entries.

**Tip**
Tips provide suggestions for using or learning more about the element.

**Caution**
Cautions describe the harm that might result if you add, change, or delete a registry element. Please take the cautions seriously. They are intended to protect your system.

**Related Entries**
Related entries are entries that configure the same component, or configure the same features on a different component. Related entries help you find the entry that meets your needs.

# About the Windows Registry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows operating systems cannot function without a registry. Nearly every action the computer performs requires the operating system to examine data stored in the registry to determine details like what resource to use, where to write a result, or how to display activity for the user. A registry with erroneous data can cause the computer to perform poorly, incorrectly, or not at all. Erroneous data in the registry can be caused by something as simple as a typing error made while editing the registry, or as serious as a physical disk error. Because of the central role played by the registry, it is essential that you back it up regularly and before making any modifications.

**For more information, see the following topics:**

- Overview of the Windows Registry
- Editing the Registry
- Data Types in the Registry
- Backing Up and Restoring the Registry

# Overview of the Windows Registry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The registry is a hierarchical database that stores the values of variables in Windows and the applications and services that run on Windows. The operating system and other programs also use the registry to store data about users and about the current configuration of the system and its components. Because the registry is available whenever the system is running, programs that start and stop can keep persistent data in the registry.

**Structure of the Registry**
Much of the registry consists of nested containers known as subtrees, keys, and subkeys, which are like folders. The actual data is stored in the registry entries, the lowest level element in the registry. The entries are like files. The series of nested containers make up a path to each entry. Because entry names are unique only within a path, entries are referenced by their full registry path and name.

The following figure shows the registry of a typical computer, as viewed with the registry editor Regedit.exe. The figure shows the terminology that this reference uses to refer to different parts of the registry.

**Subtree**
Subtrees are the *root*, or primary divisions, of the registry.

The Windows Server 2003 registry is divided into five subtrees. Subtrees themselves do not contain configuration data; they contain the keys, subkeys, and entries in which the data is stored.

Click the name of a subtree to see its description:

- HKEY_CLASSES_ROOT
- HKEY_CURRENT_USER
- HKEY_LOCAL_MACHINE
- HKEY_USERS
- HKEY_CURRENT_CONFIG

**Key**
Keys are the next level down from subtree, which contain at least one subkey (for example, Hardware Key). Some subtrees do not have keys.

**Subkey**
Subkeys are the next level down from keys or subtrees (if the subtree has no keys). Subkeys store entries and other subkeys.

**Entry**
Entries are the lowest-level element in the registry. They appear in the right pane of the registry editor window. Each entry consists of the entry name, its Data Types in the Registry (which defines the length and format of data that the entry can store), and a field known as the *value* of the registry entry. Data is stored in the value. Entries are referenced by their registry path and name.

Entries store the actual configuration data for the operating system and the programs that run on the system. As such, they are different from subtrees, keys, and subkeys, which are merely containers.

**Hive Files**
The permanent parts of the registry are stored as a set of files called the hive files. You can find a list of locations for these files in the hivelist subkey in HKLM\SYSTEM\CurrentControlSet\Control. These files are saved in systemroot\System32\Config and updated with each login. They consist of the following files, which store four of the five keys in HKEY_LOCAL_MACHINE and one key in HKEY_USERS:

- **SAM** Contains information stored in the key HKLM\SAM about the Security Accounts Manager (SAM) service.

- **SECURITY** Contains the security information stored in the key HKLM\SECURITY.
- **SOFTWARE** Contains information stored in the key HKLM\SOFTWARE about the computer's software configuration.
- **SYSTEM** Contains information stored in the HKLM\SYSTEM about the computer's system configuration.
- **DEFAULT** Contains the default system information that is stored in the key HKEY_USERS\.DEFAULT.

HKEY_LOCAL_MACHINE\HARDWARE is not stored as a file, because it is recreated each time the system starts.

**Variability of the Registry**
No two registries are exactly alike. Because the registry stores data about the hardware and software that is installed on and available to your computer, the values of many registry entries are specific to the computer and its current configuration. In addition, some of the registry content is created anew each time you start the operating system or log on. As a result, many changes to the registry are lost when you restart the computer, or when you log off and log on again.

The location of registry entries can change. Entries for programs sometimes move when the program is updated. Some programs move entire sections of the registry to a new location when a service or feature is enabled.

Because of the variability of the registry, it is not prudent to refer directly to a registry entry in any program or script. If you do, the program or script might not operate correctly with future versions of the operating system. To retrieve, add, delete, or change a registry entry, use the standard Win32 application programming interfaces (APIs). The APIs are updated with each version of the system to find the data you need.

**How Registry Data is Used**
Programs store many types of data in registry entries. Typically, programs use standard Win32 APIs to retrieve data from the registry. The program specifies an entry by its path and name, and the API returns the value of the entry. Programs also use standard APIs to add and delete registry content, and to change the values of registry entries.

After it retrieves and reads the data in the value of an entry, each program interprets the data and implements its result independently, depending on how the program is written. For example, a value of 1 might indicate to a program that an optional feature is enabled or disabled, or it might force the program to wait for one millisecond or one day before taking action. A file location stored in a registry entry might tell the program where to find a file, or it might instruct the program to move the file to the specified location.

Windows Server 2003 operating system components and programs use the registry in the following ways:

- **Setup** Both the Windows Server 2003 Setup program and other setup programs (for programs or hardware) add configuration data to the registry. For example, new information is added when you install a new SCSI adapter or change the settings for your display. Setup also reads information from the registry to determine if the prerequisite components are installed.
- **Recognizer** Each time you start the computer, the recognizer places hardware configuration data in the registry. This data includes a list of hardware detected in your system. Hardware detection is done by the hardware recognizer (Ntdetect.com) and the Windows Server 2003 kernel (Ntoskrnl.exe) programs.
- **Kernel** During system startup, the kernel extracts information from the registry, such as which device drivers to load and their load order. The kernel also passes information about itself (such as its version number) to the registry.
- **Device drivers** Device drivers send and receive load parameters and configuration data from the registry. A device driver must report the system resources that it uses, such as hardware interrupts and DMA channels, so that the system can add this information to the registry. Programs and device drivers can access this registry information to provide users with smart installation and configuration programs.

Because programs are so variable, it is very difficult to predict how a specific program will interpret the registry data. It is important to remember that the registry belongs to programs, not to users. Therefore, unless you are very familiar with the program or the registry entry, do not attempt to modify the registry.

**Users and the Registry**
Most users never need to view or edit the registry. The administrative tools and Windows interface enable users to safely

change their preferences and the services and features of the operating system. However, in rare instances, the only way to change an operating system variable is by editing the registry. For more information, see Editing the Registry.

**For More Information About the Registry**

The registry topics in Help and Support Center for Windows Server 2003 provide essential background information, including procedures for viewing and editing the registry. Also see the Windows Deployment and Resource Kits Web site link below.

# Editing the Registry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most users never need to edit the registry. You can configure most system services by using the programs provided with Windows (such as Computer Management, Control Panel, and Group Policy), or by using the administrative tools that install with optional services (such as WINS Manager and Internet Service Manager).

**Using the Registry Editor**
If you need to view the registry or to change a value that can be changed only in the registry directly, use Regedit.exe, the registry editor installed with Windows Server 2003. An alternative registry editor, Regedt32.exe, which is included in Windows 2000 and earlier, is not included in Windows Server 2003, although many of its functions appear in the Windows Server 2003 version of Regedit.exe.

**Caution**

- Do not experiment with the registry or change any settings if you do not understand them completely. The vast majority of registry settings are maintained automatically, and many can only be maintained correctly by the operating system itself. Changing these settings can produce unexpected, contradictory, or uninterpretable instructions in operating system and application programs. Editing registry settings can also degrade the performance of your system, obscure the location of essential files, or invalidate your Windows licensing agreement. Before making any changes, be sure to back up the registry by using one of the methods described in Backing Up and Restoring the Registry. After you make changes, be sure to test them before returning your computer to normal operation.

To start Regedit.exe, click **Start**, click **Run**, and then type **Regedit**.

Regedit.exe does not display the registry directly. Instead, it constructs a virtual database from several sources of data. These sources include some of the files in %systemroot%\System32\config. Regedit.exe does not display some portions of the registry that are particularly sensitive and that should never be edited, such as a subtree related to performance counters. For more information about Regedit.exe, see Registry Editor in Help and Support Center for Windows Server 2003, or see Help in Regedit.exe.

**Making Registry Changes Effective**
When you use an administrative tool to change the configuration of a system feature or service, the change usually takes effect either immediately or soon.

However, if you make the same change by editing the registry, for the change to take effect you might need to log off and log on again, restart the service, or restart Windows.

In general, if you change the value of any entry in the CurrentControlSet subkey, you must restart the computer for the changes to take effect. Also, if you use the registry editor Regedit.exe to change values for most entries in HKEY_CURRENT_USER, you must log off and log on again for the changes to take effect.

When a change to the value of an entry is not immediate, the description of the entry contains an **Activation Method**, which describes what you need to do to make the change effective.

# Data Types in the Registry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Data types define what kind of data a registry entry can store. The Windows Server 2003 registry uses the following data types:

**REG_SZ**
A fixed-length text string. Boolean ("True" or "False") values and other short text values usually have this data type. For an example, see the Pattern Entry entry.

**REG_EXPAND_SZ**
A variable-length text string. REG_EXPAND_SZ data can include variables that are resolved when an application or service uses the data. For example, the value of the File entry includes the variable %systemroot%. When the Event Log service references this entry, the variable is replaced by the name of the directory containing the Windows Server 2003 system files.

**REG_MULTI_SZ**
Multiple text strings formatted as an array of null-terminated strings, and terminated by two null characters. Entries that contain lists or multiple data in a readable form are usually this data type. The values in a REG_MULTI_SZ entry can be separated by spaces, commas, or other marks. For an example, see the value of the Machine entry, which is a list of paths accessible by all remote users.

**REG_BINARY**
Raw binary data. Most hardware component information is stored as binary data. It can be displayed in an easy-to-read format by using Windows Server 2003 Diagnostics. For an example, see theCustomColors entry.

**REG_DWORD**
A 32-bit (4-byte) number. Boolean ("True" or "False") values and many entries for device drivers and services use this data type. REG_DWORD data can be displayed and entered in hexadecimal or decimal format in the registry editor Regedit.exe. For an example, see the ActivityLogFlag entry.

**REG_DWORD_BIG_ENDIAN**
Same as REG_DWORD. A 32-bit number in which the most significant byte is displayed as the leftmost (or high-order) byte. This is the most common format for storing numbers in computers that are running Windows Server 2003.

**REG_DWORD_LITTLE_ENDIAN**
A 32-bit number in which the most significant byte is displayed as the rightmost (or low-order) byte. This is opposite of the order in which bytes are stored in the REG_DWORD and REG_DWORD_BIG_ENDIAN data types.

**REG_LINK**
Indicates a symbolic link between system or application data and a registry value. You can use Unicode characters in a REG_LINK entry.

**REG_FULL_RESOURCE_DESCRIPTOR**
A series of nested arrays designed to store a resource list for a hardware component or driver.

# Backing Up and Restoring the Registry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to back up your registry regularly, and it is also important to back up your current, working registry before you edit it. The greatest challenge to using a backup to restore a severely damaged registry is that Windows requires the registry to operate. The system constantly reads and writes to the registry, so without a functional registry to run Windows, you cannot easily start the registry restore tools. Saving a backup of the registry on removable media, a compatible disk partition, or a secure network share, can make it much easier to restore a computer to operation if an error prevents you from accessing files on your hard disk.

**Security Considerations**

Maintain the highest possible security for the backup copy of your registry, especially if you are backing up the registry for a server. The registry contains sensitive security and user information that an attacker can use to circumvent your network security or to damage your system. Use the same physical security that you use on your servers to make sure that no one can access this information. If you are saving a backup to of the registry a network share, make sure that the share is secure. Also, make sure that the backup cannot be infected by a virus, so that you can restore your system if the entire network becomes infected.

**Improve Security by Using Runas**

Running your computer as a member of the Administrators group makes the system more vulnerable to security threats. It is a best practice to log on by using an account in the Users group and then to provide administrative credentials only to use administrative tools that require them. When backing up or editing the registry, consider using Runas, a command-line tool that allows you to issue commands as if they were issued by an account that is different from that of your current logon session.

For Runas syntax and examples, see Runas in Help and Support Center for Windows Server 2003.

**Test Your Backup Methods**

Before relying on any method for backing up or restoring the registry, be sure to test the method thoroughly on a non-critical system. It is important to verify that your restoration method will work with your particular system configuration. Practicing backing up and restoring the registry will allow you to identify any problems or limitations of your backup methods before you need to repair a non-working registry. For example, network policy settings might prevent you from using registry backup and restore tools on the affected system.

**When to Back up Your Registry**

**Back up your registry regularly.** It is a best practice to back up your system on a regular schedule. Make sure that you back up the entire registry as a part of this regular system backup.

**Back up your registry before you edit it.** Using the registry editors bypasses standard safeguards, and it is possible for you to make changes to the registry that degrade performance or cause system failure. Before you edit the registry, save a copy of each of the subkeys in which you plan to make changes. Backing up only a portion of the registry is typically faster than backing up the entire registry, and uses less storage space. However, backing up the entire registry can provide more complete recovery in the event of system failure. If you chose to create a partial backup, it should be in addition to the complete backup that you make on a regular schedule.

**Back up your registry before you install or remove hardware or software.** To avoid losing configuration information, back up the registry immediately before you intend to make hardware or software changes. This way, you can restore the system to its previous configuration if the changes conflict or degrade performance. Back up the registry again after you have made the changes and determined that they are successful.

**How to Back Up Your Registry**

Before restoring the registry from a backup of the working registry, be sure to create a second backup of the current non-working registry. In the event that you find that the settings of the registry files that you are about to restore are less functional than the current non-working registry settings, you might want to restore from the second backup.

The following are descriptions of several methods for backing up and restoring your registry. Choose the method or methods that best meet your backup requirements.

**Backup Utility**

*Backup* is a utility included with Windows Server 2003 that you can use to back up and restore the entire registry. If you already use Backup to regularly back up your hard disk, you can easily include the registry in these backups.

Backup can create backup files on many kinds of media, including local and networked disk resources. Backup can create backups of files that are open and in use, in addition to closed files, so it is your best choice for creating a complete backup copy without disturbing active users or processes.

To back up the registry by using Backup, when you select the drives, files, and folders that you want to include in a backup set, also select **System State**. System State is a collection of system-specific data maintained by the operating system, which includes the registry. When you choose to back up or restore System State data, all of the data that is relevant to your system configuration is backed up or restored. Because it must be backed up and restored as a unit, the file size of this backup can be hundreds of megabytes. To restore the registry by using Backup, restore System State.

For more information about using the Backup utility and System State to backup and restore your registry, see Backup overview and System State data in Help and Support Center for Windows Server 2003.

**Ntbackup.exe**

*Ntbackup.exe* is a command-line tool that allows you to create backup files of your registry. Ntbackup.exe cannot be used to restore your registry, but Ntbackup.exe files can be used with Backup to perform the restoration.

For more information about using Ntbackup.exe to backup your registry, see Ntbackup in Help and Support Center for Windows Server 2003.

**Manually Copying Files**

You can back up the registry by manually copying the registry files from their storage location in the file system to another location. However, to make use of these files in the event of registry failure, you need to be able to start the computer by using an operating system that can read both the file system of the backup media and of the partition in which the non-working operating system is installed. Because you cannot copy files while they are open and in use, and the files that constitute the registry are always open and in use while the operating system is running, you cannot simply copy them as you would a typical file. To copy these files, you must shutdown the operating system and use another installed operating system or utility that can access and read the file system.

To back up the registry by manually copying files, copy all files in Systemroot\System32\Config to removable media, a network share, or a compatible partition. Be sure to note the actual path to these files, because Systemroot resolves differently when your computer is running a different operating system.

To restore the registry by manually copying files, copy the backed up registry files from their backup location to their original location (the location indicated by Systemroot\System32\Config when the associated operating system was running). This overwrites all registry settings with those in the backed up registry files. Restart the computer by using the repaired operating system.

**Regedit.exe and Reg.exe**

The registry editor Regedit.exe and the command-line tool Reg.exe allow you to back up and restore individual subkeys, keys, or whole subtrees by saving parts of your registry as files on your computer, removable media, or network shares.

Reg.exe can only be used to modify local registries, but Regedit.exe can be used to modify both local and remote registries. If a problem prevents users from logging in but does not stop the system from starting up, you might be able to connect to that system remotely from another computer by using Regedit.exe.

For more information about using Regedit.exe to backup and restore your registry, see Registry Editor overview in Help and Support Center for Windows Server 2003, or see Help in Regedit.exe. For more information about using Reg.exe to backup and restore your registry, see Reg in Help and Support Center for Windows Server 2003.

**RegBack.exe, RegRest.exe, and Regedt32.exe**

The functionality of the Windows 2000 Resource Kit Tools RegBack.exe and RegRest.exe is now integrated into Reg.exe. The functionality of the Windows 2000 registry editor Regedt32.exe is now integrated into Regedit.exe.

**Other Recovery Methods**

If your system cannot be restored from an available backup, consider the following recovery methods.

**Last Known Good Configuration**

Do not rely on Last Known Good Configuration as your only method of restoring the registry. It is a limited restoration that only restores information in HKLM\SYSTEM\CurrentControlSet. For this reason, Last Known Good Configuration is a convenient way to restore your hardware configuration, but is only effective when your problem is the result of hardware-related registry settings.

**Disaster Recovery**

For more information about recovering a system that does not start, including startup options, Recovery Console, and

Automated System Recovery (ASR), see "Disaster Recovery" in Help and Support Center for Windows Server 2003.

# HKEY_CURRENT_CONFIG

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCC

**Description**
The HKEY_CURRENT_CONFIG (HKCC) subtree stores configuration data for the current hardware profile.

This subtree does not contain any data. It just stores a pointer to the content of the Hardware Profiles\Current subkey in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\CurrentControlSet\Hardware Profiles subkey. The content of the Current subkey appears in HKEY_CURRENT_CONFIG, and it can be viewed and changed in either location. This subtree provides easier access to the data.

### Note

- For information about the content of HKEY_CURRENT_CONFIG, see the description of the CurrentControlSet\Hardware Profiles subkey and its Hardware Profiles\Current subkey.

**See Also**
**Concepts**
Hardware Profiles\Current
CurrentControlSet\Hardware Profiles

# HKEY_CLASSES_ROOT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCR

**Description**
The HKEY_CLASSES_ROOT subtree contains two types of data:

1. Data that associates file types with programs. The file type subkeys in HKEY_CLASSES_ROOT have the same name as the file name extension for the file type, such as .exe. File type associations are stored in the registry, but you should use Windows Explorer to change them. In Windows Explorer, from the **Tools** menu, click **Folder Options**, and then click the **File Types** tab.

2. Configuration data for COM objects, Visual Basic programs, or other automation. The configuration subkeys use either the program IDs (such as for COM, Visual Basic, automation, and scripting) or parent keys for other classes of information (such as for CLSID, Interface, TypeLib, AppId, and so on).

The content of HKEY_CLASSES_ROOT comes from two sources: HKEY_LOCAL_MACHINE\SOFTWARE\Classes and HKEY_CURRENT_USER\SOFTWARE\Classes. If a subkey or entry appears in either location, then it also appears in HKEY_CLASSES_ROOT. If the values of entries in the two Classes subkeys conflict, then only the value in HKEY_CURRENT_USER\SOFTWARE\Classes appears in HKEY_CLASSES_ROOT.

In Windows NT 4.0 and earlier, HKEY_CLASSES_ROOT displayed the data only in HKEY_LOCAL_MACHINE\SOFTWARE\Classes. The current, merged configuration lets the system independently register program classes for each user. This feature is known as *per-user class registration*.

Per-user class registration has several benefits:

- Multiple users can customize a single computer. For example, one user can install an application that uses a particular file type, and another user can install a different application that uses the same file type without affecting the first user.

- Roaming class registration. Because class registrations are stored in the user profile, the class registrations are loaded on each computer through the application management extension after the profile is loaded. When the user logs on to a different computer, the registration travels with the user through the directory and is written to a local portion of the user profile.

  Note: Roaming class registration is only available when using Intellimirror Software Management.

- Increased registry security. Previously, users could change or delete registry keys in HKLM\SOFTWARE\Classes, thus breaking an application or changing its behavior. However, with class information now written to the user profile, a system administrator can limit access to HKLM\SOFTWARE\Classes.

**Notes**

- This is a change in functionality from Windows NT 4.0, where HKEY_CLASSES_ROOT was an alias of HKEY_LOCAL_MACHINE\SOFTWARE\Classes.

- If an entry exists in both HKEY_LOCAL_MACHINE\SOFTWARE\Classes and HKEY_CURRENT_USER\SOFTWARE\Classes, then the value of the entry in HKEY_CURRENT_USER\SOFTWARE\Classes takes precedence.

- If you create a new key or subkey in HKEY_CLASSES_ROOT, then it will also be created in HKEY_LOCAL_MACHINE\SOFTWARE\Classes. If you create a new entry in this same key or subkey, then that entry will also appear in that key or subkey in HKEY_LOCAL_MACHINE\SOFTWARE\Classes.

- If you create a new entry in an existing subkey, where that entry will appear depends on whether the subkey already exists in HKEY_LOCAL_MACHINE\SOFTWARE\Classes or HKEY_CURRENT_USER\SOFTWARE\Classes:

- If the subkey exists in both HKEY_CURRENT_USER\SOFTWARE\Classes and HKEY_LOCAL_MACHINE\SOFTWARE\Classes, then the entry will appear in that subkey in HKEY_CURRENT_USER\SOFTWARE\Classes only.

- If the subkey already exists only in HKEY_CURRENT_USER\SOFTWARE\Classes, then the entry will appear in that subkey in HKEY_CURRENT_USER\SOFTWARE\Classes.

- If the subkey already exists only in HKEY_LOCAL_MACHINE\SOFTWARE\Classes, then the entry will appear in that subkey in HKEY_LOCAL_MACHINE\SOFTWARE\Classes.

# HKEY_CURRENT_USER

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU

**Description**
The HKEY_CURRENT_USER subtree contains the user profile for the user who is currently logged on to the computer. The user profile includes environment variables, personal program groups, desktop settings, network connections, printers, and application preferences. The data in the user profile is similar to the data stored in the Win.ini file in Windows 3.*x*.

The HKEY_CURRENT_USER subtree does not contain any data. It just stores a pointer to the content of the HKEY_USERS*Security ID (SID) of current user* subkey. Therefore, the content of that subkey also appears in HKEY_CURRENT_USER, and it can be viewed and changed in either location. The HKEY_CURRENT_USER subtree just provides easier access to the data.

A new HKEY_CURRENT_USER subtree is created each time a user logs on. The data for the subtree comes from the profile of the current user. If no profile is available, then the subtree is built from the user profile settings established for a default user, which are stored in *System drive*\Documents and Settings\Default User (WINNT)\Ntuser.dat.

# Console

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Console**
HKCU

**Description**
The Console key stores data for the console subsystem, which hosts all character-mode applications, such as Command Processor. In addition, the Console key can contain subkeys for custom command windows.

**Change Method**
To change the console window settings, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Notes

- The values of the entries in this key reflect the default window settings.

- The entries in this key, or the ability to change them, might be affected by Group Policy. For detailed information about particular Group Policy settings, see the *Windows Server 2003 Technical Reference* Group Policy Reference. For general information about Group Policy, see Windows Server 2003 Help or Windows XP Professional Help.

# ColorTable#

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *RGB value* | (Depends on color.) |

**Description**
Defines the colors in the color well on the **Colors** tab of the console **Properties** or **Defaults** dialog box. Each of the 16 ColorTable*x* entries corresponds to one of the registry entries **ColorTable00–ColorTable15.**

**Change Method**
To change the available colors, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Colors** tab, and then select a color well. Use the RGB boxes in the **Selected Color Values** group box to change that color.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.
- The ColorTable# entry is a variable representing the entries **ColorTable00–ColorTable15** in the HKCU\Console key. It does not actually appear in the registry. This variable entry displays information that is common to the entries in the Console key.

# CursorSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Percentage* | 0x19 (25%) |

**Description**
Specifies the percentage of a character cell that the cursor occupies. This setting affects the default window.

| Value | Meaning |
|---|---|
| 0x19 | Small (25%) |
| 0x32 | Medium (50%) |
| 0x64 | Large (100%) |

**Change Method**
To change the size of the cursor, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Options** tab, and then use the **Cursor Size** group box to change the cursor size.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

# FaceName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Font name* | (None) |

**Description**
Specifies the name of an alternate command window font. If there is no font name in the value of this entry, the system usesraster fonts.

**Change Method**
To change the alternate font, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Font** tab, and then select a font from the **Fonts** list box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

# FontFamily

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Number* | 0 |

**Description**
Specifies the font family of the font used in the command window.

| Value | Meaning |
|---|---|
| 0 | Raster fonts |
| 48 | TrueType fonts |

**Change Method**
To change the font type, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Font** tab, and then select a font from the **Fonts** single selection list box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

# FontSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Height by width* (in pixels) | 0x0 |

**Description**
Specifies the size of the font used in the command window.

The value of this entry is an 8-byte hexadecimal value. The first four digits specify the height of characters in pixels. The last four digits specify the width of characters in pixels.

For example, if the value of this entry is 0xC0005, characters in the command window are 12 (0xC) pixels high and 5 pixels wide. A value of 0x0 indicates the default character size, which is 8 pixels high and 12 pixels wide.

**Change Method**
To change the font size, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Font** tab, and then select a size from the **Size** single selection list box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- A value of 0x0 is not a null value, but reflects the default window settings.

# FontWeight

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies the weight or thickness of the command window font.

| Value | Meaning |
|---|---|
| 0 | Normal weight |
| 1 | Bold text |

**Change Method**
To change the font weight, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Font** tab and either select or clear the **Bold fonts** check box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

# FullScreen

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the command window opens in full-screen mode.

| Value | Meaning |
|---|---|
| 0 | Window mode |
| 1 | Full-screen mode |

**Change Method**
To change the screen mode, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Options** tab, and then select either **Window** or **Full Screen** in the **Display Options** group box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.
- This entry is supported only by computers with CPUs that are of x86 architecture.

# HistoryBufferSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Number of commands* | 50 |

**Description**
Specifies the number of commands that can be stored in each command history buffer.

**Change Method**
To change the size of the history buffer, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Options** tab, and then change the **Buffer Size** in the **Command History** group box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

**See Also**
**Concepts**
NumberOfHistoryBuffers

# InsertMode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies how the system responds when the user types over existing characters.

| Value | Meaning |
|---|---|
| 0 | Overtype mode. Typed characters replace highlighted characters. |
| 1 | Insert mode. Typed characters are inserted before the highlighted characters. |

**Change Method**
To change the typing mode, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Options** tab, and then select **Insert Mode** in the **Edit Options** group box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

# LoadConIme

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the IME proxy process for the Windows Server 2003 console (Conime.exe) loads automatically when you log on to Windows Server 2003.

| Value | Meaning |
|---|---|
| 0 | Do not load Conime.exe automatically. |
| 1 | Load Conime.exe automatically when Windows Server 2003 starts. |

**Activation Method**
To make changes to this entry effective, you must log off of Windows Server 2003, and then log on again.

### Note

- This entry is used only when the locale of the computer is set to 932 (Japanese), 936 (Chinese), 949 (Korean Unified Hangul), or 950 (Chinese Big5 Extended).

# NumberOfHistoryBuffers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Number of buffers* | 4 |

**Description**
Specifies the number of history buffers allocated to store commands.

**Change Method**
To change the number of history buffers, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Options** tab, and then change the **Number of Buffers** in the **Command History** group box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

**See Also**
**Concepts**
HistoryBufferSize

# PopupColors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Background color number, text number* | 0xF5 (Bright White, Purple) |

**Description**

Specifies the colors used in popup windows, by color number.

The value of this entry is a hexadecimal value. The first significant number (seventh digit) represents the background color. The second significant number (eighth digit) represents the text color.

| Hexadecimal value | Color |
|---|---|
| 0 | Black |
| 1 | Blue |
| 2 | Green |
| 3 | Aqua |
| 4 | Red |
| 5 | Purple |
| 6 | Yellow |
| 7 | White |
| 8 | Gray |
| 9 | Light Blue |
| A | Light Green |
| B | Light Aqua |
| C | Light Red |
| D | Light Purple |
| E | Light Yellow |
| F | Bright White |

**Change Method**

To change the colors of the popup windows, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Colors** tab, and then use the color well control options to change the colors.

**Activation Method**

Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

**Note**

- The value of this entry reflects the default window settings.

# QuickEdit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether QuickEdit mode is enabled. In QuickEdit mode, users can cut and paste text by using the mouse.

| Value | Meaning |
|---|---|
| 0 | QuickEdit mode is disabled. Users must use commands on the **Edit** menu to cut and paste. |
| 1 | QuickEdit mode is enabled. Users can use either the mouse or commands on the **Edit** menu to cut and paste. |

**Change Method**
To change the editing options, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Options** tab, and then select **Quick Edit Mode** in the **Edit Options** group box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

# ScreenBufferSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Lines by columns* (in pixels) Maximum: 0x7FFF7FFF | 0x190050 (25 lines by 80 columns) |

**Description**
Specifies the size of the screen buffer (the screen that is retained in memory). If the size of the screen displayed on the monitor is smaller than the screen buffer, you can scroll to see the entire screen. You cannot display a console screen that is larger than the screen buffer.

The value of this entry is an 4-byte hexadecimal value. The first 2 bytes (high word) represent the number of lines that appear on the screen. The last 2 bytes (low word) represent the number of columns on the screen.

For example, the default value, 0x190050, indicates that the screen buffer is 25 (0x19) lines high and 80 (0x50) columns wide.

**Change Method**
To change the screen buffer size, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Layout** tab, and then use the **Screen Buffer Size** group box to set the width and height.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

**See Also**
**Concepts**
WindowSize

# ScreenColors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–F (*Color number*) | 0x7 |

**Description**
Specifies the colors of console windows.

The value of this entry is an 4-byte hexadecimal value. The first 2 bytes (high word) represent the background color. The last 2 bytes (low word) represent the text color.

| Hexadecimal value | Color |
|---|---|
| 0 | Black |
| 1 | Blue |
| 2 | Green |
| 3 | Aqua |
| 4 | Red |
| 5 | Purple |
| 6 | Yellow |
| 7 | White |
| 8 | Gray |
| 9 | Light Blue |
| A | Light Green |
| B | Light Aqua |
| C | Light Red |
| D | Light Purple |
| E | Light Yellow |
| F | Bright White |

**Change Method**
To change the screen colors, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Colors** tab, and then use the color well options to change the values.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

**Note**

- The value of this entry reflects the default window settings.
- For Command Processor windows, the value of DefaultColor Entry takes precedence over the value of **ScreenColors.**

# WindowPosition

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Position* Maximum: 0x7FFF7FFF | |

**Description**
Specifies the position of the command window on the user's screen.

The value of this entry is an 4-byte hexadecimal value. The first 2 bytes (high word) represent the position of the window on the X (horizontal) axis. The last 2 bytes (low word) represent the position of the window on the Y (vertical) axis.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the window position, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Layout** tab, and then use the **Window Position** group box.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Notes

- Setting the value of this entry greater than the maximum will result in serious Command-console display problems.
- If this entry does not appear in the registry, the system selects a position for the window. This is equivalent to the **Let system position window** option on the **Layout** tab in the **Console Properties** box.

# WindowSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Size (number of characters)* Maximum: 0x7FFF7FFF | 0x190050 |

**Description**
Specifies the size of the command window.

The value of this entry is an 4-byte hexadecimal value. The first 2 bytes (high word) represent the number of lines that appear on the screen. The last 2 bytes (low word) represent the number of lines of columns on the screen.

For example, the default value, 0x190050, indicates that the screen buffer is 25 (0x19) lines high and 80 (0x50) columns wide.

If the screen buffer (the size of the screen retained in memory) is larger than the window size, you can scroll to see the entire screen. You cannot make the window size larger than the screen buffer size.

**Change Method**
To change the window size, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only. Click the **Layout** tab, and then use the **Window Size** box to set the width and height of the window.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The value of this entry reflects the default window settings.

**See Also**
**Concepts**
ScreenBufferSize

# console-window-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console

**Description**
Represents a particular console window.

When you save the configuration of a character-based window, such as when you create a shortcut to a screen, the system adds a subkey to the Console key, in which it saves the configuration of the window. These subkeys are named for the command windows. When you open the window, the system uses the values of the entries in this subkey rather than the values of the entries in the root of the Console key.

The entries in the *console-window-name* subkey are the same as the entries in the Console key.

**Change Method**
To change the console window settings, click the system menu icon in the title bar, and then select either **Defaults** to change the settings for all instances of that window, or **Properties** to change the current window only.

**Activation Method**
Changes made in the **Properties** dialog box are effective immediately. Changes made either in the **Defaults** dialog box or in the registry are effective the next time you open a console window.

### Note

- The *console-window-name* subkey does not appear unless the current user has changed the properties of the command window.

- The *console-window-name* subkey is a variable representing the subkeys in the Console key. It does not actually appear in the registry. This variable subkey displays information that is common to the subkeys in the Console key.

# element-name Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKCU\Console*console-window-name*

**Description**
The entries in the *console-window-name* subkey are used instead of those in the Console key for that console window. The meanings of these entries are the same as those of the entries in the Console key.

### Notes

- The *element-name* entry does not appear unless the current user has changed the properties of the command window.
- *element-name* is a variable representing the entries in the *console-window-name* subkey. It does not actually appear in the registry. This variable entry displays information that is common to the entries in the *console-window-name* subkey.

# Control Panel Key

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Control Panel**
HKCU

**Description**
The Control Panel key and its subkeys contain entries that specify the accessibility, regional, and desktop appearance settings.

**Change Method**
To change the values of these entries, use Control Panel.

### Note

- The entries in this key can be overriden by entries of the same name in the Microsoft\Control Panel and HKCU\Software\Policies\Microsoft\Windows\Control Panel subkeys.

# Accessibility Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessibility**
HKCU\Control Panel

**Description**
The Accessibility subkey contains subkeys and entries that specify the accessibility settings.

**Change Method**
To change the values of these entries, either double-click **Accessibility Options** in **Control Panel** or use the Accessibility hot keys.

# Accessibility-option

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessibility-option**
HKCU\Control Panel\Accessibility

**Description**
Represents an accessibility option.

**Change Method**
To change the values of these entries, either double-click **Accessibility Options** in **Control Panel** or use the Accessibility hot keys.

### Note

- The entry *Accessibility_option* is a placeholder representing the entries in the HKCU\Control Panel\Accessibility subkey. "Accessibility_option" does not actually appear in the registry. This type of entry stores information that is common to the entries in the Accessibility subkey.

# Appearance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Appearance**
HKCU\Control Panel

**Description**
The Appearance subkey stores information about the appearance of the Windows elements.

### Note

- The entries in this subkey can be superseded by Group Policy settings included in Windows Server 2003.

# Current Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Current**
HKCU\Control Panel\Appearance

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Scheme name* | Windows Standard |

**Description**
Stores the name of the color scheme currently in use on the computer.

**Change Method**
To change the color scheme, in **Control Panel** double-click **Display**. Click the **Appearance** tab, and then select a scheme from the drop-down list.

### Note

- This entry appears in the registry only when the current scheme is not "Windows Standard."

# CustomColors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CustomColors**
HKCU\Control Panel\Appearance

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *RGB value* | 0xFFFFFF |

**Description**
Stores information about the custom colors that the user creates.

**Change Method**
To change the color scheme, in **Control Panel** double-click **Display**, click the **Appearance** tab, click the **Color** drop-down list, and then click **Other**. Select a color, and then click **Add to Custom Colors**.

# Schemes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Schemes**
HKCU\Control Panel\Appearance

**Description**
The Schemes subkey stores information about the schemes used on the computer.

# Scheme-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scheme-name**
HKCU\Control Panel\Appearance\Schemes

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *RGB values* | (None) |

**Description**
Specifies the appearance of each element of the window scheme.

**Change Method**
To change the scheme, in **Control Panel** double-click **Display**, click the **Appearance** tab, and then select a scheme from the **Scheme** list. To change any of the individual elements of the scheme, select the element from the **Item** list.

### Note

- The *Scheme-name* subkey is a placeholder that represents the appearance schemes subkeys in the subkey HKCU\Control Panel\Appearance\Schemes*Scheme-name*. "Scheme-name" does not actually appear in the registry. This type of subkey stores information that is common to screen savers, and it is stored in the Schemes subkey.

# Colors Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Colors**
HKCU\Control Panel

**Description**
The Colors subkey contains entries that determine the color of each element of the Windows Server 2003 desktop. The values are shown as a series of three numbers, which refer to the RGB (red-green-blue) value of the color.

**Change Method**
To change the colors of the elements of the Windows Server 2003 desktop, in **Control Panel** double-click **Display**, then click the **Appearance** tab. Either change the complete color scheme, or change the colors of individual items.

**Activation Method**
Changes made using the **Display Properties** dialog box are effective immediately; changes made using the registry are not effective until the computer is restarted.

# Colors\element-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**element-name**
HKCU\Control Panel\Colors

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *RGB values* | *RGB value* |

**Description**
Represents a color and stores entries that apply only to that color.

**Change Method**
To change the colors of the elements of the Windows Server 2003 desktop, in **Control Panel** double-click **Display**. In the **Display Properties** dialog box, click the **Appearance** tab.

**Activation Method**
Changes made using the **Display Properties** box are effective immediately; changes made using the registry are not effective until the computer has been restarted.

### Notes

- The *element-name* entry is a placeholder representing the entries in the HKCU\Control Panel\Colors subkey. "Element-name" does not actually appear in the registry. This type of entry stores information that is common to the entries in the Colors subkey.

- The entries in the Colors subkey contain values that specify the color of each element of the Windows Server 2003 desktop. The values are shown as a series of three numbers, which refer to the RGB (red-green-blue) value of the color.

# Control Panel\Current

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Current**
HKCU\Control Panel

**Description**
The Current subkey appears in the registry, but it is not used in Windows Server 2003.

# Color Schemes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Color Schemes**
HKCU\Control Panel\Current

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Color scheme* | Windows Standard |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# Custom Colors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Custom Colors**
HKCU\Control Panel

**Description**
The Custom Colors subkey defines the custom colors in the color palette. The entries in the Custom Colors subkey are named **ColorA** through **ColorP**, and they all have a default value of 0xFFFFFF.

**Change Method**
To change the custom colors, in **Control Panel** double-click **Display**, click the **Appearance** tab, click the **Colors** drop-down list, and then click **Other**.

# Color Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Color**
HKCU\Control Panel\Custom Colors

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *RGB value* | 0xFFFFFF |

**Description**
Represents a custom color.

**Change Method**
To change the custom colors, in **Control Panel** double-click **Display**, click the **Appearance** tab, click the **Colors** list, and then click **Other**.

### Note

- The *Color* entry is a placeholder that represents the entries in the HKCU\Control Panel\Custom Colors subkey. "Color" is not the name of an actual entry in this subkey. This type of entry stores information that is common to the entries in the Custom Colors subkey.

# Control Panel\Desktop Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Desktop**
HKCU\Control Panel

**Description**
The Desktop subkey contains entries that control the appearance of the screen background and the position of windows and icons on the screen.

# ActiveWndTrkTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ActiveWndTrkTimeout**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Milliseconds* | 0 |

**Description**
Specifies time (in milliseconds) that the mouse needs to remain over a window before the window is activated.

# AutoEndTasks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AutoEndTasks**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 0 |

**Description**
Specifies whether user processes end automatically when the user either logs off or shuts down Windows Server 2003.

| Value | Meaning |
|---|---|
| 0 | Processes do not end automatically. The system waits until the process ends, and, if the process takes more time than the value of the HungAppTimeout entry, then the **End Task** dialog box appears, stating that the application cannot respond to the End Task request. |
| 1 | Processes end automatically. |

**See Also**
**Concepts**
HungAppTimeout
WaitToKillAppTimeout

# CoolSwitch

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CoolSwitch**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 1 |

**Description**
This entry is not used by and has no effect on Windows Server 2003.

# CoolSwitchColumns

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CoolSwitchColumns**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of columns in decimal* | 7 |

**Description**
Specifies how many columns of icons can appear in the **Task Switcher** (ALT+TAB) box. If there are more tasks than columns, **Task Switcher** displays multiple rows.

**See Also**
**Concepts**
CoolSwitchRows

# CoolSwitchRows

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CoolSwitchRows**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of rows in decimal* | 3 |

**Description**
Specifies how many rows of icons can appear in the **Task Switcher** (ALT+TAB) box. If there are more tasks than can be displayed in the number of rows specified by this entry and the number of columns specified by the value of **CoolSwitchColumns**, CoolSwitchColumns scrolls the displayed icons.

**See Also**
**Concepts**
CoolSwitchColumns

# CursorBlinkRate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CursorBlinkRate**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Milliseconds in decimal* | 530 |

**Description**
Specifies how much time elapses between each blink of the selection cursor.

**Change Method**
To change the time between cursor blinks, in **Control Panel** double-click **Keyboard**, click the **Speed** tab, and then use the slider in the **Cursor blink rate** box.

# DragFullWindows

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DragFullWindows**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 1 |

**Description**
Specifies what appears on the screen while a user drags a window.

| Value | Meaning |
|---|---|
| 0 | Only the outline of the window moves. |
| 1 | The entire window and its contents move. |

**Change Method**
To change what appears on the screen while a user drags a window, in **Control Panel** double-click **Display**, click the **Effects** tab, and then click the **Show window contents while dragging** check box.

### Note

- When the system is first installed, the value of this entry is 2 until Win32k.sys is loaded during the very first login session and queries the video adapter for its capabilities. Depending on the answer, the value of this entry is then set to 0 or 1.

# DragHeight

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DragHeight**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of pixels in decimal* | 4 |

**Description**
Specifies the height of the rectangle used to detect the start of a drag operation.

**See Also**
**Concepts**
DragWidth

# DragWidth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DragWidth**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of pixels in decimal* | 4 |

**Description**
Specifies the width of the rectangle used to detect the start of a drag operation.

**See Also**
**Concepts**
DragHeight

# FontSmoothing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FontSmoothing**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 2 | 2 |

**Description**
Determines whether the edges of large fonts are painted in different shades of gray to make the font look smoother.

| Value | Meaning |
|---|---|
| 0 | Font smoothing is disabled. Edges are not painted different shades of gray. |
| 2 | Font smoothing is enabled. Edges are painted different shades of gray. |

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Display**, click the **Effects** tab, and then click the **Smooth edges of screen fonts** check box.

# ForegroundFlashCount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ForegroundFlashCount**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Number of flashes in decimal* | 3 |

**Description**
Specifies the number of times the taskbar button flashes to notify the user that the system has activated a background window. If the time elapsed since the last user input exceeds the value of the ForegroundLockTimeout entry, the window will automatically be brought to the foreground.

**See Also**
**Concepts**
ForegroundLockTimeout

# ForegroundLockTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ForegroundLockTimeout**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Milliseconds* | 200000 |

**Description**
Specifies the time, following user input, during which the system keeps applications from moving into the foreground automatically (without selection by the user).

**See Also**
**Concepts**
ForegroundFlashCount

# GridGranularity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**GridGranularity**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0–49 (*8-pixel units*) | 0 |

**Description**
Windows Server 2003 does not use this entry, however, the entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

**Caution**

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# HungAppTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**HungAppTimeout**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Milliseconds in decimal* | 5000 |

**Description**
Specifies how long the system waits for user processes to end after the user clicks the **End Task** command button in Task Manager or after the user has selected to restart or shutdown the system. If this threshold is exceeded, the **End Task** dialog box appears, stating that the process did not respond.

**See Also**
**Concepts**
AutoEndTasks
WaitToKillAppTimeout

# MenuShowDelay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MenuShowDelay**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Milliseconds in decimal* | 400 |

**Description**
Specifies the interval from the time the cursor is pointed at a menu until the menu items are displayed.

# PaintDesktopVersion

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PaintDesktopVersion**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the version and build number of Windows is visible on the desktop.

| Value | Meaning |
|---|---|
| 0 | Disabled. The version and build number of Windows does not appear on the desktop. |
| 1 | Enabled. The version and build number of Windows appears on the desktop. |

# Pattern Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Pattern**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | (None) \| b1 b2 b3 b4 b5 b6 b7 b8 (*in decimal*) | (None) |

**Description**
Specifies a pattern for the screen background.

The value of this entry (b1–b8) defines a two-color pattern bitmap that is 8 pixels wide and 8 pixels high. Each b-number value specifies the pattern for one row of the bitmap. The b-number values are 8-bit binary values, which are entered in decimal. The bits in each b-number value represent a row of 8 pixels. If a bit is set to 0, the pixel it defines is set to the background color. If a bit is set to 1, the pixel it defines is set to the foreground color.

For example, if you set the b1 value to the decimal value 128, the top row of pixels in the bitmap appears as the binary equivalent, 10000000. The system interprets this to mean that the first pixel has the foreground color and the seven remaining pixels have the background color.

The background color is determined by the value of the **Background** entry, and the foreground color is determined by the value of the **WindowText** entry, both of which appear in the Colors Subkey subkey.

**See Also**
**Concepts**
Colors Subkey

# ScreenSaveActive Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ScreenSaveActive**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 0 |

**Description**
Specifies whether a screen saver is selected. Windows Server 2003 has a default value of 0 for this entry (Windows XP has 1).

| Value | Meaning |
|---|---|
| 0 | No screen saver is selected. |
| 1 | A screen saver is selected. |

**Change Method**
To select a screen saver, in **Control Panel** double-click **Display**, then click the **Screen Saver** tab. You can then either select a screen saver from the list or select (None).

### Note

- This entry can be superseded by a Group Policy setting included in Windows Server 2003. While the **No screen saver** Group Policy setting is enabled, the system ignores this entry. The configuration of the **No screen saver** Group Policy setting is stored in the **ScreenSaveActive** entry, which is in the Policies subkey.

**See Also**
**Concepts**
SCRNSAVE.EXE
ScreenSaveTimeOut Entry

# ScreenSaveTimeOut Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ScreenSaveTimeOut**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Seconds in decimal* | 600 (10 minutes) |

**Description**
Specifies how long the system must remain idle before the screen saver starts.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Display**, click the **Screen Saver** tab, and then change the value in the **Wait** box.

**See Also**
**Concepts**
ScreenSaveActive Entry

# SCRNSAVE.EXE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SCRNSAVE.EXE**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *File name* | (None) |

**Description**
Specifies the name of the screen saver executable file.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Display**, click the **Screen Saver** tab, and then select a screen saver from the **Screen saver** list.

You can also change this entry with the application programming interface (API) function **InstallScreenSaver**, which is exported by Desk.cpl. You can access this function with the command line **rundll32.exe desk.cpl,InstallScreenSaver *%1***, where *%1* is the name of a screen saver executable file.

**Activation Method**
Changes made to this entry become effective the next time the system starts a screen saver. Changes made to this entry while a screen saver is running have no effect on the running screen saver.

### Notes

- Windows operating systems add this entry to the registry when you use **Display** in Control Panel to select a screen saver. If you disable all screen savers by choosing **(None)** from the Screen Saver list, then the operating system deletes this entry from the registry and calls the **SystemParametersInfo** function with *uiAction* equal to SPI_SETSCREENSAVEACTIVE and *uiParam* equal to FALSE.

- On a computer that is running an operating system in the Windows NT, Windows 2000, Windows XP, or Windows Server 2003 family, this entry can be superseded by a Group Policy setting. While the **Screen saver executable name** Group Policy setting is enabled, the system ignores this entry. The configuration of the **Screen saver executable name** Group Policy setting is stored in the **SCRNSAVE.EXE** entry, which is in the Policies subkey.

# TileWallpaper

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TileWallpaper**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 0 |

**Description**
Specifies how the desktop wallpaper is positioned on the screen.

| Value | Meaning |
|---|---|
| 0 | Wallpaper is centered on the screen. |
| 1 | Wallpaper is tiled across the screen. |

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Display**, click the **Desktop** tab, and then select **Center** or **Tile** from the **Position** box.

### Note

- In **Display Properties**, you can also select **Stretch**, in addition to **Center** and **Tile**. For both centering and stretching, the value of this entry is zero, and stretching is configured by the entry **WallpaperStyle**.

# WaitToKillAppTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WaitToKillAppTimeout**
HKCU\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Milliseconds in decimal* | 20000 (20 seconds) |

**Description**
Specifies how long the system waits for user processes to end after the user attempts to log off or to shut down Windows Server 2003. When the time specified in this entry expires, the **End Task** dialog box appears, stating that the process did not respond. If the value of AutoEndTasks is 1, the system ends the process automatically.

**See Also**
**Concepts**
AutoEndTasks
HungAppTimeout

# WindowMetrics

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WindowMetrics**
HKCU\Control Panel\Desktop

**Description**
The WindowMetics subkey contains the current dimensions of window display elements. The units of some of the entries are in pixels, others are in X-Y coordinates, and some are in twips.

# BorderWidth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BorderWidth**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 to -750 (twips) | -15 (15 twips) |

**Description**
Specifies the current setting for the width of the borders around all the windows that have borders that can be sized.

**Change Method**
To change the value of this entry, in **Control Panel**, double-click **Display**, and then click the **Appearance** tab. Click the **Advanced** button. Select **Active Window Border** from the **Item** list, and then set the value of **Size**. The numbers in the **Size** box correspond to pixels, where one pixel equals fifteen twips.

**Activation Method**
Changes made in Control Panel are effective immediately. Changes made by using the registry editor Regedit.exe are not effective until you log off and then log on again.

### Notes

- Values entered in twips (minus values) are rounded to the nearest pixel value. For instance, -17 converts to 1 pixel, and -28 converts to 2 pixels.

- The registry also accepts values in pixels (0 to 50); however, any values entered in Control Panel will be converted to twips in the registry.

# CaptionFont

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CaptionFont**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *a data structure (in binary) that describes a font* | *data structure that describes the font Tahoma* |

**Description**
Stores a LOGFONT data structure that describes a font. The system uses the font described by this entry as the caption font.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Display**, and then click the **Appearance** tab. Next, click the **Advanced** button. Select **Active Title Bar** from the **Item** list, and then make selections from the **Font** and **Size** lists.

**See Also**
**Concepts**
CaptionHeight

# CaptionHeight

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CaptionHeight**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of twips* | -270 (270 twips) |

**Description**
Specifies the height of caption buttons.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Display**, and then click the **Appearance** tab. Next, click the **Advanced** button. Select **Active Title Bar** from the **Item** list, and then make selections from the **Font** and **Size** lists.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
CaptionFont

# IconFont

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IconFont**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *Font name (in binary)* | MS Sans Serif |

**Description**
Specifies the font used for text associated with icons.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# IconSpacing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IconSpacing**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of* | -75 (75 twips) |

**Description**
Specifies the width of the grid cell used for the large view of an icon.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Display**, and then click the **Appearance** tab. Next, click the **Advanced** button. Select **Icon Spacing (Horizontal)** from the **Item** list, and then make selections from the **Font** and **Size** lists.

**See Also**
**Concepts**
IconTitleWrap

# IconTitleWrap

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IconTitleWrap**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 1 |

**Description**
Specifies whether the title text for an icon wraps or truncates if it is too long to fit on one line.

| Value | Meaning |
|---|---|
| 0 | Title text is truncated. |
| 1 | Title text wraps to the next line. |

**See Also**
**Concepts**
IconSpacing

# MenuFont

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MenuFont**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *a data structure (in binary) that describes a font* | *data structure that describes the font Tahoma* |

**Description**
Stores a LOGFONT data structure that describes a font. The system uses the font described by this entry for menu bars.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Display**, and then click the **Appearance** tab. Next, click the **Advanced** button. Select **Menu** from the **Item** list, and then make selections from the **Font** and **Size** lists.

**See Also**
**Concepts**
MenuWidth
MenuHeight

# MenuWidth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MenuWidth**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of twips* | -270 (270 twips) |

**Description**
Specifies the width of menu bar buttons.

**Change Method**
You can change this entry only by editing the registry.

**See Also**
**Concepts**
MenuFont

# MessageFont

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MessageFont**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *a data structure (in binary) that describes a font* | *data structure that describes the font Tahoma* |

**Description**
Stores a FONTLOG data structure that describes a font. The system uses the font described by this entry in message boxes.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Display**, and then click the **Appearance** tab. Click the **Advanced** button. Select **Message Box** from the **Item** list, and then make selections from the **Font** and **Size** lists.

**See Also**
**Concepts**
CaptionFont

# MinAnimate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MinAnimate**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 0 |

**Description**
Specifies whether the operating system animates a window while minimizing, maximizing, and restoring the size of the window.

| Value | Meaning |
|---|---|
| 0 | Disabled. The operating system does not animate a window when resizing it. |
| 1 | Enabled. The operating system animates a window when resizing it. |

**Change Method**
To change the value of this entry, right-click **My Computer** and click **Properties** on the pop-up menu. Then, click the **Advanced** tab, and under **Performance**, click the **Settings** button. On the **Visual Effects** tab, clear or check **Animate windows when minimizing and maximizing**.

**Activation Method**
When you change this entry in **System Properties**, to make the change effective, click either **OK** or **Apply**.

**See Also**
**Concepts**
MinHorzGap
MinArrange

# MinArrange

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MinArrange**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0–11 | 0 |

**Description**
Specifies the arrangement of minimized windows on the desktop. The value of this entry specifies the starting point and the direction in which the windows are arranged from that point.

| Value | Meaning |
|---|---|
| 0 | Bottom left, moving right |
| 1 | Bottom right, moving left |
| 2 | Top left, moving right |
| 3 | Top right, moving left |
| 4 | Bottom left, moving up |
| 5 | Bottom right, moving up |
| 6 | Top left, moving down |
| 7 | Top right, moving down |
| 8 | Bottom left, hidden |
| 9 | Bottom right, hidden |
| 10 | Top left, hidden |
| 11 | Top right, hidden |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- Minimized windows are visible only when Windows Explorer is inactive.

**See Also**
**Concepts**
MinHorzGap
MinAnimate

# MinHorzGap

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MinHorzGap**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of pixels* | 0 |

**Description**
Specifies the size, in pixels, of the horizontal gap between two minimized windows.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
MinArrange
MinAnimate

# MinWidth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MinWidth**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of pixels* | -154 (154 pixels) |

**Description**
Specifies the width, in pixels, of a minimized window.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
MinHorzGap
MinArrange
MinAnimate

# Shell Icon BPP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shell Icon BPP**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 4 \| 8 \| 16 \| 24 \| 32 | 4 |

**Description**
Specifies the color resolution of the icons on the Windows Server 2003 desktop.

The value of this entry specifies the number of bits used to describe each pixel.

| Value (number of bits) | Meaning |
|---|---|
| 4 | 16 colors |
| 8 | 256 colors |
| 16 | 65,536 colors |
| 24 | 16,777,216 colors |
| 32 | 16,777,216 colors |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To use the maximum number of colors as determined by your display hardware and drivers, in **Control Panel** double-click **Display**, click the **Effects** tab, and then click the **Show icons using all possible colors** check box .

# ShellIconSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ShellIconSize**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 16–48 (*pixels, in decimal*) | 32 (*pixels*) |

**Description**
Specifies the size of icons that the Windows Explorer interface uses.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Size

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Size**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of pixels* | (Depends on video driver.) |

**Description**
Specifies the size of the large icon that the Windows Explorer interface uses in the shell.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Caution

- The value of this entry is determined by the video driver on the computer. Do not change the value of this entry.

# SmCaptionFont

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SmCaptionFont**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *a data structure (in binary) that describes a font* | *data structure that describes the font Tahoma* |

**Description**
Stores a FONTLOG data structure that describes a font. The system uses the font described by this entry for small captions.

**See Also**
**Concepts**
SmCaptionHeight
SmCaptionWidth

# SmCaptionHeight

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SmCaptionHeight**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of twips* | -180 (180 twips) |

**Description**
Specifies the height of small captions.

**See Also**
**Concepts**
SmCaptionFont
SmCaptionWidth

# SmCaptionWidth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SmCaptionWidth**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of twips* | -180 (180 twips) |

**Description**
Specifies the width, in twips, of small caption buttons.

**See Also**
**Concepts**
SmCaptionHeight

# StatusFont

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**StatusFont**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *Font name* (in binary) | MS Shell Dlg |

**Description**
Stores information about the font used in status bars and tooltips.

# MenuHeight

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MenuHeight**
HKCU\Control Panel\Desktop\WindowMetrics

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of* | -270 (270 twips) |

**Description**
Specifies the height of menu bars.

**See Also**
**Concepts**
MenuWidth
MenuFont

# International Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**International**
HKCU\Control Panel

**Description**
The International subkey stores options that are selected in Regional and Language Options in Control Panel. These entries determine how the system and programs display dates, times, currency, and other locale-specific notation. The data for each user is stored separately.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# AddHijriDate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AddHijriDate**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *A valid adjustment value.* | (None) |

**Description**
Specifies adjustments to the date within the Hijri calendar.

The *Hijri calendar* is a lunar calendar used in the Islamic world. The months begin and end according to a decree that is based on the observation of the new moon. It is difficult to accurately predict the beginnings and endings of the months, and there are regional variations, so an adjustment of between zero and two days is made to the calendar.

The **AddHijriDate** entry stores this adjustment value as follows:

| Value | Meaning |
|---|---|
| AddHijriDate-2 | Subtract two days. |
| AddHijriDate | Subtract one day. |
| (blank) | No adjustment is necessary. (If the date has not been adjusted, this entry does not appear in the registry. If the date has been adjusted and then restored to its original value, this entry appears in the registry with no value.) |
| AddHijriDate+1 | Add one day. |
| AddHijriDate+2 | Add two days. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe or by using the **Change Method** described below.

**Change Method**
To change the value of this entry, or to add it to the registry, first install the files for complex script and right-to-left languages. In **Control Panel**, click **Regional and Language Options**, then click the **Languages** tab. Under **Supplemental language support**, select the checkbox to **Install files for complex script and right-to-left languages (including Thai)**. Click **OK**.

To change the value of this entry, or to add it to the registry, from **Control Panel**, click on **Regional and Language Options**. In the **Standards and Formats** box, select a locale that uses the Hijri calendar. Click on the **Customize** button, then click on the **Date** tab. In the **Adjust Hijri date to:** drop-down list, select the appropriate value. Click **OK**, then click **OK** again.

**Activation Method**
Changes to this entry made through **Control Panel** take effect immediately.

# iCalendarType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iCalendarType**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0–12 (*Calendar-ID*) | 1 |

**Description**
Specifies the type of calendar currently being used on the system, as represented by its calendar ID.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Date** tab, and then select a calendar from the **Calendar type** list. Click **OK**. Note that the **Calendar type** list only appears if the computer is configured for a locale that has more than one calendar. (You can change the locale on the **Regional Options** tab of **Regional and Language Options**.)

**Activation Method**
Changes you make in **Regional and Language Options** are effective immediately.

# iCountry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iCountry**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Country/region code* | 1 |

**Description**
Stores the country/region code. This number is the same as the country or region's international telephone code, except for Canada, which is 2.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, and then click the **Regional Options** tab. Next, select a country/region from the click the list under **Standards and formats** (not the list under **Location** on the same page). Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# iCurrDigits

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iCurrDigits**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0–9 (*in decimal*) | 2 |

**Description**
Specifies the number of digits displayed in the fractional portion of currency.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Currency** tab, and then select from the **No. of digits after decimal** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

**See Also**
**Concepts**
iDigits

# iCurrency

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iCurrency**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 \| 2 \| 3 | 0 |

**Description**
Specifies the format for displaying positive numbers as currency.

| Value | Meaning |
|---|---|
| 0 | $1.1 |
| 1 | 1.1$ |
| 2 | $ 1.1 |
| 3 | 1.1 $ |

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Currency** tab, and then select a format from the **Positive currency format** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

**See Also**
**Concepts**
sCurrency

# iDate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iDate**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 \| 2 | 0 |

**Description**
Specifies the display order of day, month, and year in short dates.

| Value | Meaning |
|---|---|
| 0 | Month, day, and year. |
| 1 | Day, month, and year. |
| 2 | Year, month, and day. |

For example, if the value of this entry is 0, then February 8, 1999, will appear as 02/08/99 (assuming that registry entry sShortDate is set to MM/dd/yy).

**See Also**
**Concepts**
iDate
sLongDate
sDate

# iDigits

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iDigits**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0–9 (*in decimal*) | 2 |

**Description**
Specifies the number of digits displayed after the decimal separator in numbers. For example, the default (2) displays digits to the hundredth of a whole number.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. Click the **Numbers** tab, and then select from the **No. of digits after decimal** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

**See Also**
**Concepts**
iCurrDigits

# iFirstDayOfWeek

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iFirstDayOfWeek**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0–6 | 6 |

**Description**
Specifies which day is considered first in a week.

| Value | Meaning |
|---|---|
| 0 | Monday |
| 1 | Tuesday |
| 2 | Wednesday |
| 3 | Thursday |
| 4 | Friday |
| 5 | Saturday |
| 6 | Sunday |

# iFirstWeekOfYear

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iFirstWeekOfYear**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 \| 2 | 0 |

**Description**
Specifies which week of the year is considered first.

| Value | Meaning |
|---|---|
| 0 | Week containing January 1 is the first week of the year. |
| 1 | First full week following January 1 is the first week of the year. |
| 2 | First week containing at least four days is the first week of the year. |

# iLZero

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iLZero**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 1 |

**Description**
Specifies whether a leading zero is used in decimal fractions.

| Value | Meaning | Example |
|---|---|---|
| 0 | Uses no leading zero. | .1 |
| 1 | Uses a leading zero. | 0.1 |

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. Click the **Numbers** tab, and then select from the **Display leading zeroes** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

**See Also**
**Concepts**
sDecimal

# iMeasure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iMeasure**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 1 |

**Description**
Specifies which measurement system is used.

| Value | Meaning |
|---|---|
| 0 | Metric |
| 1 | U.S. |

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. Click the **Numbers** tab, and then select from the **Measurement system** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# iNegCurr

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iNegCurr**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0–15 (*in decimal*) | 0 |

**Description**
Specifies the format of negative currency amounts.

| Value | Meaning | Example |
|---|---|---|
| 0 | ($1.1) | ($100.00) |
| 1 | -$1.1 | -$100.00 |
| 2 | $-1.1 | $-100.00 |
| 3 | $1.1- | $100.00- |
| 4 | (1.1$) | (100.00$) |
| 5 | -1.1$ | -100.00$ |
| 6 | 1.1-$ | 100.00-$ |
| 7 | 1.1$- | 100.00$- |
| 8 | -1.1 $ | -100.00 $ (space before $) |
| 9 | -$ 1.1 | -$ 100.00 (space after $) |
| 10 | 1.1 $- | 100.00 $- (space before $) |
| 11 | $ 1.1- | $ 100.00- (space after $) |
| 12 | $ -1.1 | $ -100.00 (space after $) |
| 13 | 1.1- $ | 100.00- $ (space before $) |
| 14 | ($ 1.1) | ($ 100.00) (space after $) |
| 15 | (1.1 $) | (100.00 $) (space before $) |

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. Click the **Currency** tab, and then select from the **Negative currency format** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# iNegNumber

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iNegNumber**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0–4 | 1 |

**Description**
Specifies the format of negative numbers.

| Value | Meaning |
|---|---|
| 0 | (1.1) |
| 1 | -1.1 |
| 2 | - 1.1 |
| 3 | 1.1- |
| 4 | 1.1 - |

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. Click the **Date** tab, and then select from the **Negative number format** list. Click **OK**.

**Activation Method**
Changes you make in **Regional and Language Options** are effective immediately.

# iTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iTime**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 1 |

**Description**
Specifies whether time is displayed with a 12-hour or 24-hour clock format.

| Value | Meaning |
|---|---|
| 0 | 12-hour clock |
| 1 | 24-hour clock |

**Change Method**
The value of this entry is set when you select a time format. To select a time format, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. Click the **Time** tab, and then select a format from the **Time format** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# iTimePrefix

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iTimePrefix**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 0 |

**Description**
Specifies whether the time marker (that is, AM or PM) appears before or after the time.

| Value | Meaning | Example |
|---|---|---|
| 0 | Suffix | 9:15 AM |
| 1 | Prefix | AM 9:15 |

# iTLZero

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**iTLZero**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 0 |

**Description**
Specifies whether leading zeros are inserted before single-digit hours in time displays.

| Value | Meaning | Example |
|---|---|---|
| 0 | A leading zero is not inserted before a single-digit hour. | 6:04 PM |
| 1 | A leading zero is inserted before a single-digit hour. | 06:04 PM |

**Change Method**
The value of this entry is set when you select a time format. To select a time format, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. Click the **Time** tab, and then select a format from the **Time format** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# International\Locale

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Locale**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | 00000409 (*U.S. English*) |

**Description**
Specifies the locale that the current user has selected. The locale identifies the geographic variation of the spoken language that the user has chosen. The chosen locale determines the default values for all other regional preferences, such as how dates and times are formatted and the symbols that are used in currency. It also determines which options appear in the drop-down lists on the tabs within the **Regional and Language Options** box.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **General** tab, and then select a locale from the **Your locale (location)** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# NumShape

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NumShape**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 \| 2 | 1 |

**Description**
Specifies the shape of digits in Unicode languages. By default, the digits 0–9 are displayed with the appropriate digit shape for the default locale (presuming that the appropriate font is available). The value of this entry determines whether the digits 0–9 are forced to display as European digits.

| Value | Meaning |
|---|---|
| 0 | The shape of the digits depends on the context. |
| 1 | Digits are diplayed according to the default locale. |
| 2 | Digit shapes are determined by the sNativeDigits entry. |

When the value of this entry is 0, European digits follow Latin text, Arabic-Indic digits follow Arabic text, and Thai digits follow Thai text. When there is no preceding text, the system takes both the locale and the displayed reading order into consideration, as shown in the following table.

| Locale | Reading order | Digits |
|---|---|---|
| Arabic | right-to-left | Arabic-Indic |
| Thai | left-to-right | Thai |
| All others | any | Locale default |

**See Also**
**Concepts**
sNativeDigits

# s1159

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**s1159**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *String* | AM |

**Description**
Specifies the text or symbols used in 12-hour time formats to indicate that the time displayed is before noon.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Time** tab, and then select a symbol from the **AM symbol** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# s2359

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**s2359**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *String* | PM |

**Description**
Specifies the text or symbols used to indicate that the time displayed is at or after noon in 12-hour time formats.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Time** tab, and then select a symbol from the **PM symbol** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# sCountry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sCountry**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Name of country or region* | United States |

**Description**
Specifies the fully localized name of the country or region. This name is based on the localized version of the product.

### Caution

- Do not change the value of this entry. This value must comply with ISO standards and correspond to ISO values for functionality and language selection.

# sCurrency

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sCurrency**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol or character* | $ |

**Description**
Specifies the currency symbol used to display monetary values.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Currency** tab, and then select a symbol from the **Currency symbol** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

### Note

- Regional Options in Control Panel restricts this entry from accepting the numeric digits 0–9. Using the registry editor Regedit.exe to set the value of this entry to a numeric digit could have one of two effects:
    - If an application calls the GetCurrencyFormat API, then the system will use the default symbol for the locale.
    - If an application does not use the GetCurrencyFormat API and the application does not check the format itself, then the numeric value will be used. This could cause errors if the value is used in calculations.

**See Also**
**Concepts**
iCurrency

# sDate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sDate**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol or character* | / |

**Description**
Specifies the symbol separating numbers when the date is displayed in short date (numeric) format, such as 6/5/69.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Date** tab, and then select a symbol from the **Date separator** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

**Note**

- Regional Options in Control Panel links the **sDate** and sShortDate values. This ensures that the date separator and the date format are synchronized. For example:

| sDate | sShortDate | Meaning |
|---|---|---|
| / | m/d/yy | 6/5/69 |

- Setting **sTime** and **sTimeFormat** independently in the registry can result in conflicting values. To repair these settings, use Regional Options in Control Panel.

- Setting only one value in the registry might not change the date as it appears in applications, because an application might read either **sDate** or **sShortDate.**

**See Also**
**Concepts**
sShortDate
iDate
sLongDate

# sDecimal

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sDecimal**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol or character* | . (period) |

**Description**
Specifies the punctuation used to separate the fractional part of a decimal number from the whole-number part.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Currency** tab, and then select a symbol from the **Decimal symbol** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

### Note

- Regional Options in Control Panel restricts this entry from accepting the numeric digits 0–9. Using the registry editor Regedit.exe to set this value to a numeric digit could have one of two effects:
    - If an application calls the GetNumberFormat API, then the system will use the default symbol for the locale.
    - If an application does not use the GetNumberFormat API and the application does not check the format itself, then the numeric value will be used. This could cause errors if the value is used in calculations.

**See Also**
**Concepts**
iLZero

# sGrouping

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sGrouping**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *number of digits*[;*number of digits*...] | 3;0 |

**Description**
Describes how digits in a whole number are grouped.

Each numeric value in this entry determines how many digits will be grouped together. The first number in this value determines the number of digits immediately preceding the decimal point, and each subsequent value determines the number of digits between thousands separators. If the last value is 0, then the preceding value is repeated.

| Value | Example |
|---|---|
| 3;2;0 | 12,34,56,789 |
| 3;0 | 123,456,789 |
| 0;0 | 123456789 |

The symbol used to separate the groups is determined by the value ofsThousand.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize** button. Next, click the **Numbers** tab, and then select a grouping from the **Digit grouping** list. Click **OK**.

**Activation Method**
Changes you make in **Regional and Language Options** are effective immediately.

### Note

- **Regional and Language Options** in Control Panel limits this value to the three options in the table, but you can enter up to five groups of digits, such as 1;2;3;4;5, in the value of this entry.

**See Also**
**Concepts**
sThousand

# sList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sList**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol or character* | , (comma) |

**Description**
Specifies the symbol used to separate items in a list.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Numbers** tab, and then select a grouping from the **List Separator** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

# sLongDate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sLongDate**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Date format* | dddd, MMMM dd, yyyy |

**Description**
Specifies the format, including the abbreviations and separators, used to display dates in long date (text) format.

| Value | Example |
|---|---|
| dddd, MMMM dd, yyyy | Saturday, October 12, 1986 |
| MMMM dd, yyyy | October 12, 1986 |
| dddd, dd MMMM, yyyy | Saturday, 12 October, 1986 |
| dd MMMM, yyyy | 12 October, 1986 |

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Date** tab, and then select a format from the **Long date format** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

### Note

- You can also create your own date format if you use the correct values. (For more information about these values, see *Day, Month, Year, and Time Format*.) Click in the **Short date format** box, rather than clicking on the arrow to the right of the box, and replace or insert the necessary characters.

**See Also**
**Concepts**
iDate
sDate
sShortDate

# sMonDecimalSep

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sMonDecimalSep**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol* | "." (period) |

**Description**
Specifies the symbol used to separate the decimal values in currency.

For example, in the value **$1.20**, a period separates the decimal values.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Currency** tab, and then select a symbol from the **Decimal symbol** list. Click **OK**.

**Activation Method**
Changes you make in **Regional and Language Options** are effective immediately.

### Note

- **Regional and Language Options** in Control Panel restricts this entry from accepting the numeric digits 0–9. Editing the registry directly to set this entry to a numeric digit could have one of two effects:
    - If an application calls the GetCurrencyFormat API, the system will use the default symbol for the locale.
    - If an application does not use the GetCurrencyFormat API and the application does not check the format itself, the numeric value will be used. This could cause errors if the value is used in calculations.

# sMonGrouping

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sMonGrouping**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *number of digits*[;*number of digits*...] | 3;0 |

**Description**
Specifies how digits are grouped to the left of the decimal in currency.

Each numeric value in this entry determines how many digits are grouped together. The first number in this value determines the number of digits immediately preceeding the decimal point, and each subsequent value determines the number of digits between thousands separators. If the last value is 0, then the preceding value is repeated.

| Value | Example |
|---|---|
| 3;2;0 | 12,34,56,789 |
| 3;0 | 123,456,789 |
| 0;0 | 123456789 |

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click **Customize**. Next, click the **Currency** tab, and then select a grouping from the **Digit grouping** list.

**Activation Method**
Changes you make in **Regional and Language Options** are effective immediately.

### Note

- **Regional and Language Options** in Control Panel limits this value to the three options displayed in the table below. However, you can enter up to five groups of digits in this value in the registry. For example, you can enter 1;2;3;4;5.

**See Also**
**Concepts**
sMonThousandSep

# sMonThousandSep

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sMonThousandSep**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol* | "," (comma) |

**Description**
Specifies the symbol used to group digits in currency. Digits groupings are specified by the value of sMonGrouping.

For example, if **sMonGrouping** specifies grouping by thousands and the value of this entry stores a comma, then one thousand dollars is displayed numerically as **$1,000**.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Currency** tab, and then select a symbol from the **Digit grouping symbol** list.

**Activation Method**
Changes you make in **Regional and Language Options** are effective immediately.

### Note

- **Regional and Language Options** in Control Panel restricts this entry from accepting the numeric digits 0–9. Editing the registry directly to set this value to a numeric digit could have one of two effects:
    - If an application calls the GetCurrencyFormat API, the system will use the default symbol for the locale.
    - If an application does not use the GetCurrencyFormat API and the application does not check the format itself, the numeric value will be used. This could cause errors if the value is used in calculations.

**See Also**
**Concepts**
sMonGrouping

# sNativeDigits

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sNativeDigits**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number set* | 0123456789 |

**Description**
Describes the ten characters that are the native equivalent to the ASCII 0–9, as determined by the locale.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Numbers** tab, and then select substitution scheme from the **Standard digits** list. Click **OK**.

Note that **Standard digits** only appears when, on the **Languages** tab of **Regional and Language Options**, **Install files for complex script and right-to-left languages** is selected.

**Activation Method**
Changes you make in **Regional and Language Options** are effective immediately.

sGrouping, sThousand

**See Also**
**Concepts**
sGrouping
sThousand

# sNegativeSign

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sNegativeSign**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol* | "-" (hyphen) |

**Description**
Specifies the symbol that is used to display negative numbers.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. On the **Numbers** tab, select a symbol from the **Negative sign symbol** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

### Note

- **Regional and Language Options** in Control Panel restricts this entry from accepting the numeric digits 0–9. Using the registry editor Regedit.exe to set this value to a numeric digit could have one of two effects:

    - If an application calls the GetNumberFormat or GetCurrencyFormat API, then the system will use the default symbol for the locale.

    - If an application does not use the GetNumberFormat or GetCurrencyFormat API and the application does not check the format itself, then the numeric value will be used. This could cause errors if the value is used in calculations.

# sPositiveSign

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sPositiveSign**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol* | (None) |

**Description**
Specifies the symbol used to display positive numbers.

### Note

- Currently, no locale uses a symbol for positive numbers.
- **Regional and Language Options** in Control Panel restricts this entry from accepting the numeric digits 0–9. Using the registry editor Regedit.exe to set this value to a numeric digit could have one of two effects:
    - If an application calls the GetNumberFormat or GetCurrencyFormat API, then the system will use the default symbol for the locale.
    - If an application does not use the GetNumberFormat or GetCurrencyFormat API and the application does not check the format itself, then the numeric value will be used. This could cause errors if the value is used in calculations.

# sShortDate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sShortDate**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Time format* | M/d/yy |

**Description**
Specifies the format used to display dates in short date (numeric) format, for example 6/5/69.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. Click the **Date** tab, and then select a format from the **Short date format** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

### Notes

- You can also create your own date format if you use the correct values. (For more information about these values, see *Day, Month, Year, and Time Format*.) Click in the **Short date format** box, rather than clicking on the arrow to the right of the box, and replace or insert the necessary characters.

- **Regional and Language Options** in Control Panel links the sDate and **sShortDate** entries. This ensures the date separator and the date format are synchronized. For example,

| sDate | sShortDate | Meaning |
|---|---|---|
| / | m/d/yy | 6/5/69 |

- Setting **sDate** and **sShortDate** independently in the registry can result in conflicting values. To repair these settings, use Regional Options in Control Panel.

- Setting only one value in the registry might not change the date as it appears in applications, because an application can read either **sDate** or **sShortDate.**

**See Also**
**Concepts**
iDate
sDate
sLongDate

# sThousand

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sThousand**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol or character* | , (comma) |

**Description**
Specifies the symbol used to group the digits in a whole number.

For example, if the value of this entry is a comma, and the value of sGrouping is 3;0, then the system displays the number 3000 as 3,000.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**. On the **Regional Options** tab, click the **Customize** button. On the **Numbers** tab, select a symbol from the **Digit grouping symbol** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

### Note

- **Regional and Language Options** in Control Panel restricts this entry from accepting the numeric digits 0–9. Using the registry editor Regedit.exe to set this value to a numeric digit could have one of two effects:
    - If an application calls the GetNumberFormat API, then the system will use the default symbol for the locale.
    - If an application does not use the GetNumberFormat API and the application does not check the format itself, then the numeric value will be used. This could cause errors if the value is used in calculations.

**See Also**
**Concepts**
sGrouping

# sTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sTime**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Symbol or character* | : (colon) |

**Description**
Specifies the character used to separate the hours, minutes, and seconds when displaying time. For example, if the value is a colon, noon appears as 12:00:00.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Time** tab, and then select a character from the **Time separator** list. Click **OK**.

**Activation Method**
Changes made in **Regional and Language Options** are effective immediately.

### Note

- **Regional and Language Options** in Control Panel links the **sTime** and sTimeFormat entries. This ensures that the time separator and the time format are synchronized. For example,

| sTime | sTimeFormat | Meaning |
|---|---|---|
| : | h:mm:ss tt | 8:20:00 AM |

- Setting **sTime** and **sTimeFormat** independently in the registry can result in conflicting values. To repair these settings, use Regional Options in Control Panel.

- Setting only one value in the registry might not change the time as it appears in applications, because an application can read either **sTime** or **sTimeFormat**.

**See Also**
**Concepts**
sTimeFormat

# sTimeFormat

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**sTimeFormat**
HKCU\Control Panel\International

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Time format* | h:mm:ss tt |

**Description**
Descibes the format in which time is displayed.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. Next, click the **Time** tab, and then select a format from the **Time format** list. Click **OK**.

**Activation Method**
Changes you make in **Regional and Language Options** are effective immediately.

**Note**

- **Regional and Language Options** in Control Panel links the sTime and **sTimeFormat** entries. This ensures that the time separator and time format are synchronized. For example:

| sTime | sTimeFormat | Meaning |
|---|---|---|
| : | h:mm:ss tt | 8:20:00 AM |

- Setting **sTime** and **sTimeFormat** independently in the registry can result in conflicting values. To repair these settings, use **Regional and Language Options** in Control Panel.

- Setting only one value in the registry might not change the time as it appears in applications, because an application can read either **sTime** or **sTimeFormat.**

**See Also**
**Concepts**
sTime

# Patterns

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Patterns**
HKCU\Control Panel

**Description**
The Patterns subkey contains entries that specify the colors of desktop patterns.

**Change Method**
To change the values of the entries in this subkey, in **Control Panel** double-click **Display**, click the **Background** tab, click the **Pattern** button, and then select a pattern from the list. You can edit the pattern by clicking the **Edit Pattern** button.

# Patterns\Pattern

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**>Pattern**
HKCU\Control Panel\Patterns

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | (None) \| b1 b2 b3 b4 b5 b6 b7 b8 (*in decimal*) | 170 85 170 85 170 85 170 85 |

**Description**
Represents a pattern.

**Change Method**
To change the values of the entries in this subkey, in **Control Panel** double-click **Display**, click the **Background** tab, click the **Pattern** button, and then select a pattern from the list. You can edit the pattern by clicking the **Edit Pattern** button.

### Notes

- The *Pattern* entry is a variable representing the entries in the HKCU\Control Panel\Patterns subkey. It does not actually appear in the registry. This variable entry is used to display information that is common to the entries in the Patterns subkey.

- The values b1–b8 define a two-color pattern bitmap that is 8 pixels wide and 8 pixels high. Each b-number specifies the pattern for one row of the bitmap. The b-numbers are 8-bit binary values, which are entered in decimal. The bits in each b-number value represent a row of 8 pixels. If a bit is set to 0, then the pixel it defines is set to the background color. If a bit is set to 1, then the pixel it defines is set to the foreground color.

- For example, if you set the b1 value to the decimal value 128, then the top row of pixels in the bitmap appears as the binary equivalent, 10000000. The system interprets this to mean that the first pixel has the foreground color and the seven remaining pixels have the background color.

- The background color is determined by the value of **Background** and the foreground color is determined by the value of **WindowText**, both of which appear in the Colors Subkey subkey.

- Each entry in the Patterns subkey is named for and describes a pattern, and each has the REG_SZ data type.

**See Also**
**Concepts**
Colors Subkey

# Screen Saver.Screensaver-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Screen Saver.Screensaver-name**
HKCU\Control Panel

**Description**
This subkey represents a screen saver, and it stores entries that apply only to that screen saver.

The Control Panel subkey can contain subkeys representing each screen saver installed on Windows Server 2003. The Screen Saver.*Screensaver-name* subkey is a variable representing these subkeys. Each Screen Saver.*Screensaver-name* subkey stores data that is specific to a screen saver.

**Change Method**
To change the screen saver, in **Control Panel** double-click **Display**, click the **Screen Saver** tab, and then select a screen saver from the **Screen Saver** list.

### Note

- The Screen Saver.*Screensaver-name* subkey is a placeholder representing the screen saver subkeys in HKCU\Control Panel subkey. "Screen Saver.Screensaver-name" does not actually appear in the registry. This is a type of subkey that stores information that is common to screen savers and that is stored in the Control Panel subkey.

# Screen Saver.Screensaver-name\element-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**element-name**
HKCU\Control Panel\Screen Saver.*Screensaver-name*

**Description**
Each *element-name* entry represents an element of the screen saver that the user can alter.

**Change Method**
To change the screen saver elements, in **Control Panel** double-click **Display**, click the **Screen Saver** tab, and then click **Settings**.

### Note

- The *element-name* entry is a placeholder that represents the entries in the subkey HKCU\Control Panel\Screen Saver.*Screensaver-name*. "Screen Saver.Screensaver-name" does not actually appear in the registry. This type of entry displays information that is common to the entries in the Screen Saver.*Screensaver-name* subkey.

# Environment Key

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Environment**
HKCU

**Description**
The Environment key stores the values of environment variables that users can set. Each entry in this key associates an environment variable with the path that is substituted for the variable. The user's profile defines the default values for these entries.

**Change Method**
To change the values of entries in this key, in **Control Panel** double-click **System**, click the **Advanced** tab, and then click **Environment Variables**. These entries appear in the **User variables for *user-name*** section.

**Activation Method**
To make changes to the entries in this key effective, you must log off and then log on again. Changes to these variables take effect when a non-Windows-based application is run or when the command prompt is used.

### Caution

- The data type of entries in this key must be REG_EXPAND_SZ. If a user or application changes the data type, the system will not replace the variable with its intended value.

**See Also**
**Concepts**
Environment Subkey

# TEMP Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TEMP**
HKCU\Environment

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\Temp |

**Description**
Specifies the path represented by the TEMP variable.

**Change Method**
To change the value of this entry, in **Control Panel** use **System**, click the **Advanced** tab, and then click **Environment Variables**. In the **User variables for *user-name*** dialog box, select **TEMP**, click **Edit**, and then enter the new data in the **Variable Value** dialog box.

**Activation Method**
To make changes to the entries in this subkey effective, you must log off and then log on again. Changes to these variables take effect either when a non-Windows-based application is run or when the command prompt is used.

### Caution

- The data type of entries in this subkey must be REG_EXPAND_SZ. If a user or application changes the data type, the system will not replace the variable with its intended value.

# TMP Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TMP**
HKCU\Environment

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\Temp |

**Description**
Specifies the path represented by the TMP variable.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **System**, click the **Advanced** tab, and then click **Environment Variables**. In the **User variables for *user-name*** dialog box, select **TMP**, click **Edit**, and then enter the new data in the **Variable Value** dialog box.

**Activation Method**
To make changes to the entries in this subkey effective, you must log off and then log on again. Changes to these variables take effect when a non-Windows-based application is run or when the command prompt is used.

### Caution

- The data type of entries in this subkey must be REG_EXPAND_SZ. If a user or application changes the data type, the system will not replace the variable with its intended value.

# Network Key

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network**
HKCU

**Description**
The Network key stores information about mapped network drives. Each subkey in Network represents a drive that is connected to the computer and that will be reconnected when the user logs on again. The subkeys are named for the drive letter to which the drive is mapped. The system adds a *Drive letter* subkey to the registry when a drive is mapped and **Reconnect at Logon** is selected. The entries in the *Drive letter* subkeys store data that is used to reconnect the drive.

**Change Method**
To map to a network drive, in **Control Panel**, select **Network and Dial-up Connections**. From the **Tools** menu, select **Map Network Drive**. In the **Folder** box, type the path to the drive, and then click **Finish**.

# Drive letter

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Drive letter**
HKCU\Network

**Description**
The *Drive letter* subkey represents a mapped network drive that the system recognizes.

**Change Method**
To map to a network drive, in Control Panel, from the **Tools** menu select **Map Network Drive**. In the **Folder** box, type the path to the drive, and then click **Finish**.

### Notes

- The *Drive letter* subkey is a variable representing the subkeys in the HKCU\Network key. It does not actually appear in the registry. This variable entry displays information that is common to the subkeys in the Network subkey.

- The subkeys are named for the drive letter to which the drive is mapped. The system adds a *Drive letter* subkey to the registry when a drive is mapped and **Reconnect at Logon** is selected. The entries in the *Drive letter* subkeys store data that is used to reconnect the drive.

# ConnectionType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ConnectionType**
HKCU\Network*Drive letter*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 \| 2 | 1 |

**Description**
Specifies how the drive is connected to the local computer.

| Value | Meaning |
|---|---|
| 1 | Drive redirection |
| 2 | Print redirection |

# ProviderName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ProviderName**
HKCU\Network*Drive letter*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Network name* | Microsoft Windows Network |

**Description**
Specifies the network provider used to map this particular network drive.

# ProviderType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ProviderType**
HKCU\Network*Drive letter*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *provider-number* | *provider-number* |

**Description**
Specifies a constant that identifies the provider used to map this network drive. The value for the Microsoft LanMan provider is 0x20000. Other values are assigned by Microsoft to third-party providers.

# SOFTWARE Subtree

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Software**
HKCU

**Description**
This subtree contains the user profile for the user who is currently logged on to the computer. The user profile includes environment variables, personal program groups, desktop settings, network connections, printers, and application preferences.

This subtree does not contain any data. It only stores a pointer to the content of the HKEY_USERS*Security ID (SID)* of the current user's subkey. Therefore, the content of that subkey also appears in this subtree, and it can be viewed and changed in either location. This subtree provides easier access to the data.

A new HKEY_CURRENT_USER subtree is created each time a user logs on. The data for the subtree comes from the profile of the current user. If no profile is available, then the subtree is built from the user profile settings established for a default user, which are stored in System drive\Documents and Settings\Default User (WINNT)\Ntuser.dat.

**See Also**
**Concepts**
HKLM\SOFTWARE

# Microsoft Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft**
HKCU\Software

**Description**
The Microsoft subkey stores configuration data for Microsoft programs running on the computer. Each subkey in the Microsoft subkey represents a different program. The entries in the program subkeys are specific to the current user.

The Microsoft subkey in HKEY_CURRENT_USER stores entries that apply only to the current user of the computer. The Software\Microsoft subkey in HKEY_LOCAL_MACHINE stores entries that apply to all users of the computer.

**See Also**
**Concepts**
SOFTWARE\Microsoft

# Command Processor Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Command Processor**
HKCU\Software\Microsoft

**Description**
The Command Processor subkey stores user preferences for the command processor (Cmd.exe).

**Change Method**
You can change the values of the entries in this subkey by using Command Processor (Cmd.exe), but only for the duration of that Cmd.exe's execution. For all other changes, you must edit the registry directly.

**Activation Method**
If you make changes to the value of the entries in this subkey by editing the registry directly, you must restart Cmd.exe to make the changes effective.

## Notes

- Entries in the Command Processor subkey can also appear in HKLM. If both appear, in most cases the values in HKCU take precedence over the values in HKLM. However, the HKLM **Autorun** entry executes before the HKCU **Autorun** entry. Command-line settings entered in Cmd.exe take precedence over registry settings.

- For more information about Command Processor, see Windows Server 2003 Help or Windows XP Professional Help.

**See Also**
**Concepts**
Microsoft\Command Processor
Console

# AutoRun Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AutoRun**
HKCU\Software\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *command* | (None) |

**Description**
Contains commands that run each time you start Command Processor (Cmd.exe).

**Change Method**
You can change the value of this entry by using Command Processor (Cmd.exe) and typing **CMD /D**, but only for the duration of that instance of Cmd.exe. For all other changes, you must edit the registry directly.

**Activation Method**
To make changes to this entry effective, you must restart Cmd.exe.

### Notes

- **AutoRun** can appear in HKLM and HKCU. If both appear in the registry, the value of the entry in HKLM executes first, then the value in HKCU executes.
- The command-line parameters take precedence over the registry settings.

### Tip

- For more information about the AutoRun feature, at the command prompt, type **CMD /?**.

**See Also**
**Concepts**
Command Processor\AutoRun

# CompletionChar Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CompletionChar**
HKCU\Software\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0x1F \| 0x20 | 0x9 (the **TAB** key) |

**Description**
Specifies a character that activates the filename completion feature in Command Processor (Cmd.exe).

To specify a character, enter the hexadecimal value of an ASCII control character other than 0x20 (space). For example, 0x4 represents Ctrl+D and 0x6 represents Ctrl+F.

| Value | Meaning |
|---|---|
| 0x0 | Disables completion. |
| 0x1–0x1F | Corresponding ASCII control character. |
| 0xD | Disables completion. |
| 0x20 | **SPACEBAR**. Disables completion. |

**Change Method**
You can change the value of this entry by using Command Processor (Cmd.exe), but only for the duration of that Cmd.exe's execution. For all other changes, you must edit the registry directly.

To enable or disable file name completion and directory name completion for one use of the command processor, at the command prompt, type **CMD /F:ON** (to enable) or **CMD /F:OFF** (to disable). This command sets the completion character to CTRL+F (0x6).

**Activation Method**
If you make changes to the value of this entry by editing the registry directly, you must restart Cmd.exe to make the changes effective.

### Notes

- **CompletionChar** can appear in HKLM and HKCU. If both appear in the registry, the value of the entry in HKCU takes precedence over the value in HKLM. The command-line switches take precedence over the registry settings while using Command Processor (Cmd.exe).
- Assigning the values 0x0, 0xD, or 0x20 to this entry disables file name completion rather than assigning a completion character.
- If you use the console Paste function to enter text into a CMD window, you might trigger completion if the characters you paste contain completion characters.

### Tip

- For more information about the file name completion feature, at the command prompt, type: **CMD /?**
- The datatype of this entry can be either REG_DWORD or REG_SZ.

**See Also**

**Concepts**

PathCompletionChar Entry

Command Processor\CompletionChar

# DefaultColor Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DefaultColor**
HKCU\Software\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFE | 0x0 |

**Description**
Specifies a default background color and text color for the Command Processor (Cmd.exe) window. The value consists of two hexadecimal digits. The first digit specifies the background color; the second digit specifies the text color. The hexadecimal values correspond to colors in the following table.

| Hexadecimal value | Color |
|---|---|
| 0 | Black |
| 1 | Blue |
| 2 | Green |
| 3 | Aqua |
| 4 | Red |
| 5 | Purple |
| 6 | Yellow |
| 7 | White |
| 8 | Gray |
| 9 | Light Blue |
| A | Light Green |
| B | Light Aqua |
| C | Light Red |
| D | Light Purple |
| E | Light Yellow |
| F | Bright White |

This entry is used only when its value is valid and not 0x0, and when the user has not specified the window color by using the COLOR or /T commands. Values are not valid if the background color and the text color are identical, such as 0x55. If the value of this entry is not valid or is 0x0, and the user has not entered a command to change the color, then Cmd.exe uses the colors specified for the console window in which Cmd.exe runs.

**Activation Method**
If you make changes to the value of this entry by editing the registry directly, you must restart Cmd.exe to make the changes effective.

## Notes

- **DefaultColor** can appear in HKLM and HKCU. If both appear in the registry, then the value of the entry in HKCU takes precedence over the value in HKLM.

- The datatype of this entry can be either REG_DWORD or REG_SZ.

**See Also**
**Concepts**
Command Processor\DefaultColor
ScreenColors

# DelayedExpansion Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DelayedExpansion**
HKCU\Software\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the command processor, Cmd.exe, supports delayed expansion of environment variables. If the value of this entry is 1, then the Cmd.exe interprets an exclamation point (!) as an environment variable that is expanded only when used.

By default, Cmd.exe expands variables when it reads a line of text. However, you can direct Cmd.exe to expand a variable only when it is used. If you do this, you can include variables in compound conditional statements and prevent them from being set before they are tested.

| Value | Meaning |
|---|---|
| 0 | Delayed expansion is not supported. Cmd.exe expands environment variables when it reads them, and it does not interpret exclamation points as delayed variables. |
| 1 | Delayed expansion is supported. The system expands environment variables surrounded by exclamation points, such as !VAR!, only when they are used. |

**Change Method**
You can change the value of this entry by using Cmd.exe and typing **cmd /v:on**, but only for the duration of that Cmd.exe's execution. For all other changes, you must use the registry editor Regedit.exe.

**Activation Method**
If you make changes to the value of this entry by using Regedit.exe, then you must restart Cmd.exe to make the changes effective.

## Notes

- **DelayedExpansion** can appear in HKLM and in HKCU. If both appear in the registry, the value of the entry in HKCU takes precedence over the value in HKLM.
- The datatype of this entry can be either REG_DWORD or REG_SZ.
- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe. Even when this entry appears in the registry, the setting of the command-line parameter **/v** takes precedence over it.

## Tip

- For more details, run Cmd.exe and type **help set** or **cmd /?**, and read the entry on the **/v** command.

**See Also**
**Concepts**
Command Processor\DelayedExpansion

# EnableExtensions Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**EnableExtensions**
HKCU\Software\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the extensions to the command processor (Cmd.exe), such as chdir and setlocal, are enabled. Because the extensions might be incompatible with a command-script language, Windows Server 2003 lets you disable the extensions. This entry affects all extensions. You cannot use this entry to selectively enable or disable extensions.

| Value | Meaning |
|---|---|
| 0 | Disabled |
| 1 | Enabled |

**Change Method**
You can change the value of this entry by using Command Processor (Cmd.exe) and typing **cmd /e:off**, but only for the duration of that Cmd.exe's execution. For all other changes, you must edit the registry directly.

**Activation Method**
If you make changes to the value of this entry by editing the registry directly, you must restart Cmd.exe to make the changes effective.

### Notes

- The datatype of this entry can be either REG_DWORD or REG_SZ.
- **EnableExtensions** can appear in HKLM and in HKCU. If both appear in the registry, the value of the entry in HKCU takes precedence over the value in HKLM.

**See Also**
**Concepts**
Command ProcessorEnableExtensions

# PathCompletionChar Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PathCompletionChar**
HKCU\Software\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0x1F \| 0x20 | 0x9 (the **TAB** key) |

**Description**
Specifies a character that activates the directory name completion feature in Command Processor (Cmd.exe).

To specify a character, enter the hexadecimal value of an ASCII control character other than 0x20 (SPACEBAR). For example, 0x4 represents Ctrl+D, and 0x6 represents Ctrl+F.

| Value | Meaning |
|---|---|
| 0x0 | Disables completion. |
| 0x1–0x1F | Corresponding ASCII control character. |
| 0xD | Disables completion. |
| 0x9 | **TAB** key. |
| 0x20 | **SPACEBAR**. Disables completion. |

**Change Method**
You can change the value of this entry by using Cmd.exe, but only for the duration of that Cmd.exe's execution. For all other changes, you must use the registry editor Regedit.exe.

**Activation Method**
If you make changes to the value of this entry by using Regedit.exe, you must restart Cmd.exe to make the changes effective.

### Notes

- **PathCompletionChar** can appear in HKLM and HKCU. If both appear in the registry, then the value of the entry in HKCU takes precedence over the value in HKLM.
- Assigning the values of 0x0, 0xD, or 0x20 to this entry disables file name completion rather than assigning a completion character.
- The command-line perameters take precedence over the registry settings while using Cmd.exe.
- The datatype of this entry can be either REG_DWORD or REG_SZ.
- If you use the console Paste function to enter text into a CMD window, then you might trigger completion if the characters you paste contain completion characters.

### Tip

- To enable or disable directory name completion for one use of Cmd.exe, at the command prompt, type **CMD /F:ON** (to enable) or **CMD /F:OFF** (to disable). This command sets the completion character to CTRL+D (0x4).
- For more information about the path completion feature, at the command prompt, type **CMD /?**.

**See Also**

**Concepts**

CompletionChar Entry

Command Processor\PathCompletionChar

# Driver Signing Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Driver Signing**
HKCU\Software\Microsoft

**Description**
The Driver Signing subkey in HKCU stores entries for the driver signing feature for Windows Server 2003. The entries in this subkey apply only to the current user of the computer.

Microsoft digitally signs driver files to affirm that the driver software has been tested with Windows and has not been altered since it was tested.

### Note

- The entries in this subkey apply only to the current user of the computer. The entries in the Microsoft\Driver Signing subkey in HKLM apply to all users of the computer.
- The system adds this subkey to the registry when you change the default value of one of its entries.

### Tip

- For more information about driver signing, see Windows Server 2003 Help or Windows XP Professional Help.

**See Also**
**Concepts**
Microsoft\Driver Signing

# Policy Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policy**
HKCU\Software\Microsoft\Driver Signing

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 1 |

**Description**
Specifies how the system responds when a user tries to install device-driver files that are not digitally signed.

| Value | Meaning |
|---|---|
| 0 | Ignore. The system proceeds with the installation even if it includes unsigned files. |
| 1 | Warn. The system notifies the user that files are not digitally signed, and it lets the user decide whether to stop the installation or to proceed with it and permit unsigned files to be installed. |
| 2 | Block. The system does not install unsigned files. As a result, the installation stops, and none of the files in the driver package are installed. |

This entry applies only to the current user of the computer. However, if a stricter setting is configured for all users of the computer (in the Driver Signing\Policy entry in HKLM) or if a user Group Policy setting establishes a stricter configuration (in the entry **BehaviorOnFailedVerify** in HKLM), then this value is ignored for the current user.

This entry stores the option selected in the **File signature verification** section of the **Driver Signing Options** dialog box when the **Apply setting as system default** box is not selected. If the box is selected, the selection applies to all users of the computer, and it is stored in the Driver Signing\Policy entry in HKLM.

**Change Method**
To change the value of this entry, in **Control Panel** double-click **System**, click the **Hardware** tab, click **Driver Signing**, and then select an option in **File signature verification**. (Do not select the **Apply setting as a system default** check box.)

### Note

- The system adds this entry to the registry when you change the default value by using System in Control Panel.
- This entry applies to all users of the computer. The **Policy** entry also can appear in the HKCU\SOFTWARE\Microsoft\Driver Signing for a particular user. Also, you can use the **Code signing for device drivers** Group Policy setting to add the entry**BehaviorOnFailedVerify** to the registry.

### Tip

- For more information about driver signing and code signing, see Windows Server 2003 Help and Support Center or Windows XP Professional Help and Support Center.

**See Also**
**Concepts**
Driver Signing\Policy

# Microsoft\Fax Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Fax**
HKCU\Software\Microsoft\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# fxsclnt

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**fxsclnt**
HKCU\Software\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The fxsclnt subkey stores configuration data of the current user for the Fax Console.

# FaxConsoleView

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FaxConsoleView**
HKCU\Software\Microsoft\Fax\fxsclnt

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The *FaxConsoleView* subkey is a type of subkey that stores configuration information for the way faxes are displayed in the Fax Console of the Fax service.

"*FaxConsoleView*" is a placeholder for actual subkeys with one of the following four names: InboxView, IncomingView, OutboxView, and SentItemsView.

# Archive

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Archive**
HKCU\Software\Microsoft\Fax\fxsclnt

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The Archive subkey stores user configuration data for the Fax Console.

**Change Method**
You can change this entry only by editing the registry.

# SortAscending

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SortAscending**
HKCU\Software\Microsoft\Fax\fxsclnt*FaxConsoleView*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies whether the sorting within the view is ascending or descending. *FaxConsoleView* is one of four views named InboxView, IncomingView, OutboxView or SentItemsView.

**Change Method**
This entry is maintained automatically by Fax Console.

### Note

- *FaxConsoleView* is a placeholder for the types of subkeys in the fxsclnt subkey; "FaxConsoleView" does not actually appear in the registry.

# SortColumn

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SortColumn**
HKCU\Software\Microsoft\Fax\fxsclnt*FaxConsoleView*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies the column by which Fax Console sorts. *FaxConsoleView* is one of four views named InboxView, IncomingView, OutboxView or SentItemsView.

**Change Method**
This entry is maintained automatically by Fax Console.

### Note

- *FaxConsoleView* is a placeholder for the types of subkeys in the fxsclnt subkey; "FaxConsoleView" does not actually appear in the registry.

# Columns

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Columns**
HKCU\Software\Microsoft\Fax\fxsclnt*FaxConsoleView*

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The Columns subkey stores current user configuration data for column order and appearance within the Fax Console folders.

# Order Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Order**
HKCU\Software\Microsoft\Fax\fxsclnt*FaxConsoleView*\Columns

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies the position of the column <Column ID> in Fax Console. *FaxConsoleView* is one of four views named InboxView, IncomingView, OutboxView or SentItemsView.

**Change Method**
This entry is maintained automatically by Fax Console.

### Note

- *FaxConsoleView* is a placeholder for the types of subkeys in the fxsclnt subkey; "FaxConsoleView" does not actually appear in the registry.

# MessagesPerCall

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MessagesPerCall**
HKCU\Software\Microsoft\Fax\fxsclnt\Archive

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies a number of fax messages within a single request during the Inbox or SentItems folder refresh. This entry can be used to optimize the refresh time of big archives.

**Change Method**
You can change this entry only by editing the registry.

# Confirm

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Confirm**
HKCU\Software\Microsoft\Fax\fxsclnt

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

# ItemDeletion

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ItemDeletion**
HKCU\Software\Microsoft\Fax\fxsclnt\Confirm

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
The value of this entry determines whether or not deletion of a fax message or cover page should be confirmed.

| Value | Meaning |
|---|---|
| 0 | The deletion of fax message or cover page should not be confirmed. |
| 1 | The deletion of fax message or cover page should be confirmed |

**Change Method**
You can change this entry only by editing the registry.

# MainFrame

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MainFrame**
HKCU\Software\Microsoft\Fax\fxsclnt

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The MainFrame sub key stores current user configuration data for the Fax Console main window size and position. This data is used to place the Fax Console window at the same place it was opened last time.

# Maximized

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maximized**
HKCU\Software\Microsoft\Fax\fxsclnt\MainFrame

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies whether or not Fax Console window was maximized.

| Value | Meaning |
|---|---|
| 0 | Fax Console window was not maximized. |
| 1 | Fax Console window was maximized. |

**Change Method**
This entry is maintained automatically by Fax Console.

# NormalPosBottom

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NormalPosBottom**
HKCU\Software\Microsoft\Fax\fxsclnt\MainFrame

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores the last bottom Y (vertical) coordinate of the Fax Console window.

**Change Method**
This entry is maintained automatically by Fax Console.

# NormalPosLeft

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NormalPosLeft**
HKCU\Software\Microsoft\Fax\fxsclnt\MainFrame

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores the last left X (horizontal) coordinate of the Fax Console window.

**Change Method**
This entry is maintained automatically by Fax Console.

# NormalPosRight

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NormalPosRight**
HKCU\Software\Microsoft\Fax\fxsclnt\MainFrame

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores the last right X (horizontal) coordinate of the Fax Console window.

**Change Method**
This entry is maintained automatically by Fax Console.

# NormalPosTop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NormalPosTop**
HKCU\Software\Microsoft\Fax\fxsclnt\MainFrame

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores the last upper Y (vertical) coordinate of the Fax Console window.

**Change Method**
This entry is maintained automatically by Fax Console.

# SplitterPos

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SplitterPos**
HKCU\Software\Microsoft\Fax\fxsclnt\MainFrame

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores the last X (horizontal) coordinate of the Fax Console splitter bar. The splitter bar is the vertical line that seperates the left tree view from the right list view.

**Change Method**
This entry is maintained automatically by Fax Console.

# Fax\Setup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setup**
HKCU\Software\Microsoft\Fax\

**Description**
The Setup subkey stores configuration information about the current user for the Fax Service.

# CfgWzdrUserInfo

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CfgWzdrUserInfo**
HKCU\Software\Microsoft\Fax\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies whether the Fax service shows the Sender Information page during invocation of the Fax Configuration Wizard.

**Change Method**
The value of this entry is set by the Fax Configuration Wizard upon first complete.

# WereCpesConverted

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WereCpesConverted**
HKCU\Software\Microsoft\Fax\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies whether or not Windows 95 CPE cover pages have been converted to the new COV format.

# UserInfo

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**UserInfo**
HKCU\Software\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The UserInfo subkey stores configuration information about the sender of a fax for the Fax Service. This information can be sent in a cover page to the recipient of a fax.

# FaxNumber

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FaxNumber**
HKCU\Software\Microsoft\Fax\UserInfo

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores the telephone number of the local Fax device. The Fax Service prints this number in the banner of pages that it transmits when the registry entry Branding is equal to 1.

**Change Method**
To change the value of this entry, click on the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Console**. In **Fax Console**, on the **Tools** menu, select **Sender Information**. In **Sender Information**, type the telephone number for the fax device in the **Fax number** box.

**Activation Method**
Changes made to this entry though **Fax Console** or through the Fax API are effective immediately.

**See Also**
**Concepts**
Branding

# File Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File Manager**
HKCU\Software\Microsoft

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this subkey or its entries from the registry, and do not change the values in the entries. These changes can cause serious, unexpected results.

# FxsCover

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FxsCover**
HKCU\Software\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The FxsCover subkey stores information for the Cover Page Editor, a component of the Fax Service.

# Ntbackup Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ntbackup**
HKCU\Software\Microsoft

**Description**
The Ntbackup subkey contains entries that store configuration data for Windows Backup (Ntbackup.exe). Windows Backup enables you to back up files to tape, and to restore files from tape. You can use Windows Backup to back up the files in which the registry is stored.

### Note

- When you upgrade from a previous version of Windows on which you used Windows Backup, delete the subkeys and entries in the Ntbackup subkey before running the upgraded version of Windows Backup. This ensures that registry entries are appropriate to the most recent version of Windows Backup.

- To run Windows Backup, from the **Start** menu, point to **Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

### Caution

- When you configure Windows Backup, it adds its default subkeys and entries to the registry. Some of these entries are no longer used by Windows Backup. The other entries are maintained by the system, and you should never change any of them by editing the registry. Altering these values can prevent Windows Backup from successfully backing up and/or restoring data, and it can cause your system to behave unpredictably.

# Backup Engine

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backup Engine**
HKCU\Software\Microsoft\Ntbackup

**Description**
The Backup Engine subkey stores user preference information for the Backup utility.

### Caution

- Do not change the values of the entries in this subkey. These entries are maintained by the system, and modifying them might cause Backup to improperly save or restore data, resulting in data loss.

# Use fast file restore

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use fast file restore**
HKCU\Software\Microsoft\Ntbackup\Backup Engine

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 1 |

**Description**
Specifies whether the fast file restore feature is enabled.

| Value | Meaning |
|---|---|
| 0 | Disables fast file restore. |
| 1 | Enables fast file restore. |

**Tip**

- Fast file restore is desirable for most tasks. However, you should disable fast file restore if you encounter problems restoring a file that resides entirely on the last tape of a multiple tape restore operation. Change the value of this entry to 0 by using the registry editor Regedit.exe. Rerun Backup, then, by using Regedit.exe, change the value of this entry back to 1.

# User Interface

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User Interface**
HKCU\Software\Microsoft\Ntbackup

**Description**
The User Interface subkey stores user preference information for Backup.

**Caution**

- Do not change the values of the entries in this subkey. These entries are maintained by the system, and modifying them might cause Backup to improperly save or restore data, resulting in data loss.

# RAS Phonebook

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RAS Phonebook**
HKCU\Software\Microsoft

**Description**
The RAS Phonebook subkey is created and maintained by the Routing and Remote Access service. It stores entries that configure dial-up features for each user, such as the options on the **Dial-up Preferences** dialog box in Network and Dial-up Connections.

**Change Method**
To change the values of entries in this subkey, use Network and Dial-up Connections. Do not edit the registry.

**Tip**

- To open the **Dial-up Preferences** dialog box, start Network and Dial-up Connections, click the **Advanced** menu, and then select **Dial-up Preferences**.

# Microsoft\Windows

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows**
HKCU\Software\Microsoft

**Description**
The Windows subkey stores configuration data for the programs running on the Win32 subsystem. The subkeys of the Windows subkey represent versions of the Win32 subsystem, such as CurrentVersion, which refers to the version currently running on the computer. Each version subkey contains subkeys representing programs that run on that version of the subsystem.

The Windows subkey in HKEY_CURRENT_USER stores entries that apply only to the current user of the computer. The SOFTWARE\Microsoft\Windows subkey in HKEY_LOCAL_MACHINE stores entries that apply to all users of the computer.

**See Also**
**Concepts**
Microsoft\Windows Subkey

# CurrentVersion Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CurrentVersion**
HKCU\Software\Microsoft\Windows

**Description**
The CurrentVersion subkey stores configuration data for programs running on the Win32 subsystem. The subkeys of CurrentVersion represent programs that run on the current version of the Win32 subsystem and that store user-specific configuration data for those programs.

# CurrentVersion\Explorer Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Explorer**
HKCU\Software\Microsoft\Windows\CurrentVersion

**Description**
The Explorer subkey stores configuration data for Windows Explorer and the Windows interface, including the **Start** menu, the desktop, common folders, and Control Panel.

# Shutdown Setting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shutdown Setting**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1 \| 0x2 \| 0x4 \| 0x8 \| 0x10 \| 0x20 \| 0x40 \| 0x80 | 0x2 |

**Description**
Specifies how the system interprets a command to shut down the system.

This entry stores the setting selected most recently from the list on the **Shut Down Windows** dialog box. This setting is displayed as the default the next time the dialog box appears.

| Value | Meaning |
|---|---|
| 0x1 | Log off. Ends the session, leaving the computer running on full power. |
| 0x2 | Shut down. Ends the session and shuts down Windows Server 2003 so that the power can be turned off safely. |
| 0x4 | Restart. Ends the session, shuts down Windows Server 2003, and starts Windows Server 2003 again. |
| 0x8 | Restart in MS-DOS mode. (Not available in Windows Server 2003.) |
| 0x10 | Stand by. Maintains the session, keeping the computer running on low power with data still in memory. (This setting requires power management hardware support.) |
| 0x20 | Stand by (with wakeup events disabled). Maintains the session, keeping the computer running on low power with data still in memory. (This setting requires power management hardware support.) |
| 0x40 | Hibernate. Saves the session to disk so that the power can be turned off safely. The session is restored the next time Windows Server 2003 is started. |
| 0x80 | Disconnect. Disconnects the session. The session can be reconnected the next time the user logs on. (Disconnect is available only in Terminal Server sessions.) |

**Change Method**
The value of this entry is changed whenever the user selects an option in the **Shut Down Windows** window and clicks **OK**.

### Note

- The list on the **Shut Down Windows** dialog box displays all of the shut-down options that the computer supports. If the value of this entry is changed to represent an unsupported option, then the system uses the default value of 0x2.

# Advanced Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Advanced**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer

**Description**
The Advanced subkey stores the settings of the options in the **Taskbar and Start Menu Properties** dialog box.

**Change Method**
To change the value of entries in this subkey, from the **Start** menu, point to **Settings**, click **Taskbar & Start Menu**, then make changes in the **Taskbar and Start Menu Properties** dialog box. Do not edit this entry directly in the registry.

# Intellimenus Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IntelliMenus**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\Advanced

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | YES \| No | YES |

**Description**
Enables or disables personalized menus.

This entry stores the setting of the **Use Personalized Menus** check box in the **Taskbar and Start Menu Properties** dialog box.

Windows Server 2003 personalizes a long menu by moving recently used items to the top of the menu and by hiding items that have not been used recently. Users can display the hidden items by clicking an arrow to extend the menu.

| Value | Meaning |
|---|---|
| YES | Enables personalized menus. |
| No | Disables personalized menus. |

**Change Method**
To change the value of this entry, from the **Start** menu, point to **Settings**, and then click **Taskbar & Start Menu**. On the **General** tab, select or clear the **Use Personalized Menus** check box.

### Note

- This entry can be superseded by a Group Policy setting included in Windows Server 2003. While the **Disable Personalized Menus** policy setting is in effect, the system removes the **Use Personalized Menus** check box from the **Start** menu and ignores this entry. The configuration of the **Disable Personalized Menus** policy setting is stored in the value of **Intellimenus** in the Policies subkey. If users cannot find the **Use Personalized Menus** option, then check the configuration of the **Disable Personalized Menus** policy setting.

- Personalized menus require user tracking. If the **Disable user tracking** policy setting is in effect, then Group Policy sets the value of **NoInstrumentation** to 1. As a result, the system disables user tracking and personalized menus, and ignores the value of this entry.

# Shell Folders Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shell Folders**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Note

- Paths to the Windows Explorer folders are stored in the User Shell Folders Subkey subkey.

### Caution

- Do not delete this subkey or its entries from the registry, and do not change the values in the entries. These changes can cause serious, unexpected results.

# User Shell Folders Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User Shell Folders**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer

**Description**
The User Shell Folders subkey stores the paths to Windows Explorer folders for the current user of the computer.

### Note

- The entries in this subkey can also appear in the subkey Shell Folders. Both the subkeys Shell Folders and User Shell Folders can appear in both the keys HKEY_LOCAL_MACHINE and HKEY_CURRENT_USER. The entries that appear in user User Shell Folders take precedence over those in Shell Folders. The entries that appear in HKEY_CURRENT_USER take precedence over those in HKEY_LOCAL_MACHINE.

**See Also**
**Concepts**
Shell Folders Subkey
Explorer\Shell Folders
Explorer\User Shell Folders

# User Shell Folders\Desktop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Desktop**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\Desktop |

**Description**
Stores the path to the folder containing the shortcuts and Internet addresses that appear on the user's desktop.

You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's profile folder. The default is *Systemroot*\Documents and Settings*username*.

**See Also**
**Concepts**
Common Desktop

# Favorites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Favorites**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\Favorites |

**Description**
Stores the path to the folder containing the Internet addresses that appear in the **Favorites** menu of the user's Internet browser.

You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's profile folder. The default is *Systemroot*\Documents and Settings*username*.

**See Also**
**Concepts**
Common Favorites

# NetHood

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetHood**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\NetHood |

**Description**
Stores the path to the folder containing shortcuts to servers that the user has added to My Network Places.

You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's profile folder. The default is *Systemroot*\Documents and Settings*username*.

# User Shell Folders\Personal Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Personal**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\My Documents |

**Description**
Stores the path to the user's Personal folder.

The Personal folder is the default location for saving application files. You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's Profiles folder. The default is *Systemroot*\Documents and Settings*username*.

**See Also**
**Concepts**
User Shell Folders\Personal

# PrintHood

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PrintHood**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\PrintHood |

**Description**
Stores the path to the folder containing shortcuts to printers that the user has added to the computer.

You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's Profiles folder. The default is *Systemroot*\Documents and Settings*username*.

# User Shell Folders\Programs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Programs**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\Start Menu\Programs |

**Description**
Stores the path to the folder that contains the items that appear in the **Programs** menu on the **Start** menu.

You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's Profiles folder. The default is *Systemroot*\Documents and Settings*username*.

**See Also**
**Concepts**
Common Programs

# SendTo

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SendTo**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\SendTo |

**Description**
Stores the path to the folder containing links to drives, communication programs, and applications.

You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's Profiles folder. The default is *Systemroot*\Documents and Settings*username*.

# Start Menu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Start Menu**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\Start Menu |

**Description**
Stores the path to the folder containing the items that appear in the **Start** menu.

You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's Profiles folder. The default is *Systemroot*\Documents and Settings*username*.

**See Also**
**Concepts**
Common Start Menu

# Startup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Startup**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\Start Menu\Programs\Startup |

**Description**
Stores the path to the folder containing shortcuts to items that appear on the **Programs** menu on the user's **Start** menu.

Windows Server 2003 automatically starts these items when the user logs on. You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's Profiles folder. The default is *Systemroot*\Documents and Settings*username*.

**See Also**
**Concepts**
Common Startup

# Recent

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recent**
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %USERPROFILE%\Recent |

**Description**
Stores the path to the folder containing the names of the documents that the user opened most recently.

These items appear in the **Documents** menu on the user's **Start** menu. You can enter any valid Universal Naming Convention (UNC) name or mapped drive in the value of this entry.

### Note

- The %USERPROFILE% variable represents the path to the user's Profiles folder. The default is *Systemroot*\Documents and Settings*username*.

# CurrentVersion\Policies Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policies**
HKCU\Software\Microsoft\Windows\CurrentVersion

**Description**
The Policies subkey contains entries that store Group Policy settings.

Group Policy adds, deletes, and changes the values of entries in this subkey when you enable or disable a policy setting. When the policy settings affecting a particular user or computer change (such as when the user moves to a new user group) Group Policy deletes the entries for the old group and adds the entries for the new group.

Group Policy manages registry entries in four Policies subkeys:

- Policies Subkey
- CurrentVersion\Policies
- Software\Policies
- CurrentVersion\Policies Subkey

The Policies subkeys in HKCU store policy-related entries that are configured separately for each user. The system adds these entries when you enable (and, for some policies, when you disable) a policy setting in the User Configuration folder in Group Policy.

The Policies subkeys in HKLM store policy-related entries that apply to all users of the computer. The system adds these entries when you enable (and, for some policies, when you disable) a policy setting in the Computer Configuration folder in Group Policy.

Many policy-related entries can appear in both HKCU and HKLM. Typically, when both entries exist, the entries in HKLM take precedence over the entries in HKCU.

This subkey stores policy-related entries for Windows NT 4.0 and earlier. The SOFTWARE\Policies subkey was added for Windows 2000, and it remains in Windows Server 2003. Please note that the different versions of Windows support different policy settings.

**Change Method**
To add, remove, or change the value of entries in this subkey, use Group Policy.

### Note

- Many of the entries in the Policies subkeys have counterparts in other registry subkeys. These counterparts, known as *preferences*, have the same effect as entries in the Policies subkey. However, when an entry appears in the Policies subkey, the preference is almost always ignored. For details, see the individual entry descriptions.

### Caution

- Do not use the registry editor to change the values in this subkey. If you do, they will be overwritten the next time your policy settings are refreshed.

# Policies\ActiveDesktop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ActiveDesktop**
HKCU\Software\Microsoft\Windows\CurrentVersion\Policies

**Description**
The ActiveDesktop subkey stores the entries created when you configure a Group Policy setting for Active Desktop in Windows Server 2003. Group Policy creates and maintains these entries, and Active Desktop reads and interprets them.

# AdminComponent Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AdminComponent**
HKCU\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

**Description**
The AdminComponent subkey stores entries that Group Policy settings add for Active Desktop. These entries specify items that an administrator has added to or deleted from Active Desktop by configuring a policy setting.

# Policies\Explorer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Explorer**
HKCU\Software\Microsoft\Windows\CurrentVersion\Policies

**Description**
The Explorer subkey contains entries that store the configurations of Group Policy settings for Windows Explorer and for the Windows interface. These include policy settings for the **Start** menu, taskbar, desktop, and Control Panel. Group Policy creates and maintains the entries in this subkey, and Windows Explorer reads and interprets them.

This subkey stores policy-related entries that are configured separately for each user. The Software\Microsoft\Windows\CurrentVersion\Policies\Explorer subkey in HKEY_LOCAL_MACHINE stores policy-related entries that apply to all users of the computer.

### Note

- Many of the entries in the Policies subkeys have counterparts in other registry subkeys. These counterparts, known as *preferences*, have the same effect as entries in the Policies subkey. However, when an entry appears in the Policies subkey, the preference is almost always ignored. For details, see the individual entry descriptions.

**See Also**
**Concepts**
Policies\Explorer Subkey

# Explorer\NoFileAssociate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NoFileAssociate**
HKCU\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 for servers running Terminal Services, 0 for all others |

**Description**
Disables the buttons on the **File Types** tab. As a result, users can view file type associations, but they cannot add, delete, or change the associations.

The **File Types** tab shows which programs open when a user double-clicks a type of file. The file types are represented by the file name extension.

| Value | Meaning |
|---|---|
| 0 | The user can add, delete, and change file type associations. |
| 1 | The user cannot add, delete, or change file type associations. |

By disabling this tab, you might prevent the current user from creating an association that starts a risky or destructive program. Also, the current user cannot create less-productive associations, such as associating the .doc file type with Microsoft Wordpad instead of Microsoft Word.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Notes

- Although this entry is located in the Policies\Explorer subkey, it is not associated with a Group Policy included in Windows Server 2003. You can create a policy by adding this entry to an Administrative Templates (.adm) file. For more information about creating custom .adm files, see Windows Server 2003 Help and Support Center or Windows XP Professional Help and Support Center.

- This entry applies only to the current user. To prevent any user of the computer from adding, deleting, or changing file type associations, use the NoFileAssociate Entry entry in the HKLM subtree. If the settings of the **NoFileAssociate** entries conflict, the setting in HKCU takes precedence over the setting in HKLM.

### Tip

- To see the **File Types** tab, in My Computer, from the **Tools** menu, click **Folder Options**, and then click the **File Types** tab.

### Caution

- This entry specifies only whether buttons on the **File Types** tab in **Folder Options** are available. Any other method of changing file type associations is unaffected by this entry.

**See Also**
**Concepts**
NoFileAssociate Entry

# Windows Help

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Help**
HKCU\Software\Microsoft

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this subkey or its entries from the registry, and do not change the values in the entries. These changes can cause serious, unexpected results.

# Microsoft\Windows NT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows NT**
HKCU\Software\Microsoft

**Description**
The Windows NT subkey stores configuration data for components of Windows Server 2003. The subkeys of the Windows NT subkey represent versions of the operating system, such as CurrentVersion, which refers to the version currently running on the computer. Each version subkey contains subkeys representing components that run on that version of the system.

The Windows NT subkey in HKEY_CURRENT_USER stores entries that apply only to the current user of the computer. The SOFTWARE\Microsoft\Windows NT subkey in HKEY_LOCAL_MACHINE stores entries that apply to all users of the computer.

**See Also**
**Concepts**
Microsoft\Windows NT Subkey

# Windows NT\CurrentVersion

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CurrentVersion**
HKCU\Software\Microsoft\Windows NT

**Description**
The CurrentVersion subkey stores configuration data for programs that are designed to run on the most current version of Windows Server 2003 installed on the computer. The subkeys of CurrentVersion correspond to programs that run on the current version of Windows Server 2003. Those subkeys store user-specific configuration data for those programs.

# Devices Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devices**
HKCU\Software\Microsoft\Windows NT\CurrentVersion

**Description**
The Devices subkey lists the local printers and network-accessible printers to which the client computer has a connection.

Each of the entries in this subkey has the following format:

| Entry name | Data type | Value |
|---|---|---|
| *printer-name* | REG_SZ | *driver-name*,*port*: |

*Printer-name* is either the name of the local printer or *****server-name******printer-name*.

*Driver-name* for Windows clients is **winspool**.

When the printer is a local printer, *port* is the name of the port to which it is connected; for example, **LPT1**. Otherwise, *port* is **Ne***xx*, where *xx* is a numeric value, starting with 00. This latter format is used for printers that are access through the network, for local printers with ports that begin with ****, when the port name has a space in it, and when the name of the port has a length greater than five.

# Extensions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Extensions**
HKCU\Software\Microsoft\Windows NT\CurrentVersion

**Description**
The Extensions subkey stores entries that associate file name extensions with the programs that run them at the command line. These entries are used only when a file name extension is not already associated with a program.

Each of the entries in this subkey has the following format:

| Entry name | Data type | Value |
|---|---|---|
| *File name extension* | REG_SZ | *Command line to run* |

In the following example, the entry indicates that the default action for all files with the extension .txt is to open the file with Microsoft Notepad:

| Entry name | Data type | Value |
|---|---|---|
| txt | REG_SZ | Notepad.exe^.txt |

In the following example, the entry indicates that the default action for all files with the extension .txt is to open the file with notepad and print it to the default printer. At run time, the caret symbol (^) is replaced with the name of the selected file in the command line:

| Entry name | Data type | Value |
|---|---|---|
| txt | REG_SZ | Notepad.exe^.txt -openandprint |

**Change Method**
To associate a file type with a program, use Windows Explorer. From the **Tools** menu, click **Folder Options**, and then click the **File Types** tab.

### Note

- This subkey exists to provide backward compatibility with older versions of Windows. In Windows Server 2003, file associations are provided primarily by the information stored in HKEY_CLASSES_ROOT.

### Tip

- For more information about creating a file association, see the Microsoft Platform Software Development Kit (SDK).

**See Also**
**Concepts**
HKEY_CLASSES_ROOT

# Network Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network**
HKCU\Software\Microsoft\Windows NT\CurrentVersion

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support administration tools designed for Windows NT 4.0 or earlier.

# Server Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server Manager**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Network

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 and earlier. Specifically, this subkey stores configuration data for Server Manager.

# User Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User Manager**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Network

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 and earlier. Specifically, this subkey stores configuration data for User Manager.

### Note

- In Windows Server 2003 and Windows XP Professional, use the Users and Groups snap-in. For more information about new ways to do familiar tasks, see Windows Server 2003 Help and Support Center or Windows XP Professional Help and Support Center.

# User Manager for Domains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User Manager for Domains**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Network

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 and earlier. Specifically, this subkey stores configuration data for User Manager for Domains.

### Notes

- In Windows Server 2003, use Active Directory Users and Computers. For more information about new ways to do familiar tasks, see Windows Server 2003 Help or Windows XP Professional Help.

- If you are administering an Active Directory domain, use Active Directory Users and Computers. If you are remotely administering a Windows NT 4.0 domain, you will need to use User Manager for Domains from a computer running Windows NT 4.0 or Windows 2000; User Manager for Domains is not supported on Windows Server 2003.

# Program Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Program Manager**
HKCU\Software\Microsoft\Windows NT\CurrentVersion

**Description**
Program Manager is the primary desktop interface for Windows NT 3.51 and earlier. In Windows Server 2003, the default interface is Windows Explorer, but you can still use Program Manager (Progman.exe).

# Program Manager\Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restrictions**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager

**Description**
The Restrictions subkey remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this subkey.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# NoClose Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NoClose**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Restrictions

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# NoFileMenu Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NoFileMenu**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Restrictions

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# NoRun Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NoRun**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Restrictions

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# NoSaveSettings Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NoSaveSettings**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Restrictions

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# Program Manager\Restrictions Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restrictions**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Restrictions

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
The Restrictions subkey remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this subkey.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# ShowCommonGroups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ShowCommonGroups**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Restrictions

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# EditLevel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**EditLevel**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Restrictions

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 \| 3 \| 4 | 0 |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# Program Manager\Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Settings**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager

**Description**
The Settings subkey remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this subkey.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# AutoArrange

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AutoArrange**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Settings

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# display.drv

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**display.drv**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Settings

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *File name* | vga.drv |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# MinOnRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MinOnRun**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Settings

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# SaveSettings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SaveSettings**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager\Settings

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this entry.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# UNICODE Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**UNICODE Groups**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Program Manager

**Description**
The UNICODE Groups subkey remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this subkey.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# TaskManager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TaskManager**
HKCU\Software\Microsoft\Windows NT\CurrentVersion

**Description**
The TaskManager subkey stores preferences for Task Manager.

# TrueType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TrueType**
HKCU\Software\Microsoft\Windows NT\CurrentVersion

**Description**
The TrueType subkey stores data that determines the use and display of TrueType fonts in Windows-based applications.

# TTEnable

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TTEnable**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\TrueType

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | 1 |

**Description**
Specifies whether **Add Fonts** in the Fonts folder displays TrueType fonts. When the value of this entry is 1, or if this value is absent from the registry, **Add Fonts** displays all font files. When the value of this entry is not 1, **Add Fonts** displays only fonts with the file name extensions .FON and .FO_.

**Change Method**
You should not change the value of this entry.

### Note

- Changing the value of this entry to a value other than 1 will not disable TrueType fonts. TrueType fonts are always enabled.

# TTonly

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TTonly**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\TrueType

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 0 |

**Description**
This entry is obsolete. It once determined whether fonts other than TrueType fonts can be used in Windows-based programs, and it can still be changed from the Fonts folder. However, its value is unused.

| Value | Meaning |
|---|---|
| 0 | All installed fonts can be used. |
| 1 | Only TrueType fonts can be used. |

**Change Method**
This entry is obsolete, though it can still be found in the Fonts folder. In **Control Panel** double-click **Fonts**. On the **Tools** menu, click **Folder Options**, and then click the **TrueType Fonts** tab. Select the **Show only TrueType Fonts in the programs on my computer** check box.

### Note

- For more information on Fonts, click the **Help** menu in the Fonts folder.

# CurrentVersion\Windows

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows**
HKCU\Software\Microsoft\Windows NT\CurrentVersion

**Description**
The Windows subkey contains entries that store user preferences.

**See Also**
**Concepts**
Control\Windows Subkey
CurrentVersion\Windows Subkey

# Device Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Device**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Output-device-name*,winspool,*Port-connection* | *Output-device-name*,winspool,*Port-connection* |

**Description**
Specifies the default printer. An explicit port and driver must be assigned to the printer device. The *device-driver* is the file name (without the extension) of the device-driver file.

**Change Method**
To change the default printer, from the **Start** menu, point to **Settings**, and click **Printers**. Right-click the printer, and then click **Set as Default Printer**.

# Documents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Documents**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *File name extensions* | (None) |

**Description**
Defines files that Windows Server 2003 treats as documents. Enter the file name extensions without the preceding periods, and separate the extensions from each other with a space.

Use this entry only to define document file name extensions that are not associated with any application and that are not listed in the Extensions subkey.

**See Also**
**Concepts**
Windows\Programs

# fPrintError

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**fPrintError**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether error messages generated during debugging are printed. Error messages usually indicate serious problems in the Window Manager code.

| Value | Meaning |
|---|---|
| 0 | Disabled. Error messages do not print. |
| 1 | Enabled. Error messages print. |

**Note**

- This entry appears only when running a debug build of the operating system.

**See Also**
**Concepts**
fPrintVerbose
fPromptOnError
fPromptOnVerbose
fPromptOnWarning

# fPrintFileLine

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**fPrintFileLine**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether to print the file and line number that generated an error message during debugging.

| Value | Meaning |
|---|---|
| 0 | Disabled. File and line number do not print. |
| 1 | Enabled. File and line number do print. |

**Note**

- This entry appears only when running a debug build of the operating system.

# fPrintVerbose

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**fPrintVerbose**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies the level of detail printed when debugging generates warnings and error messages. Verbose data is intended only for Microsoft internal development purposes.

| Value | Meaning |
|---|---|
| 0 | Standard debugging data is displayed. |
| 1 | All data is displayed. |

**Note**

- This entry appears only when running a debug build of the operating system.

**See Also**
**Concepts**
fPrintError
fPrintWarning

# fPrintWarning

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**fPrintWarning**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether warnings generated during debugging are printed. Warnings usually indicate errors in an application's interaction with the Window Manager.

| Value | Meaning |
|---|---|
| 0 | Disabled. Warnings do not print. |
| 1 | Enabled. Warnings print. |

**Note**

- This entry appears only when running a debug build of the operating system.

**See Also**
**Concepts**
fPrintError
fPrintVerbose
fPromptOnError
fPromptOnVerbose
fPromptOnWarning

# fPromptOnError

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**fPromptOnError**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the debugger prompts the user when debugging generates error messages. Errors usually indicate serious problems in Window Manager code.

| Value | Meaning |
|---|---|
| 0 | Do not prompt. |
| 1 | Prompt. |

**Note**

- This entry appears only when running a debug build of the operating system.

**See Also**
**Concepts**
fPrintError
fPrintVerbose
fPromptOnError
fPromptOnVerbose
fPromptOnWarning

# fPromptOnVerbose

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**fPromptOnVerbose**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies the level of detail of user prompts when warnings and errors are generated during debugging. Verbose data is intended for Microsoft internal development only.

| Value | Meaning |
|---|---|
| 0 | Standard data is displayed. |
| 1 | All data is displayed. |

**Note**

- This entry appears only when running a debug build of the operating system.

**See Also**
**Concepts**
fPrintError
fPrintVerbose
fPromptOnError
fPromptOnVerbose
fPromptOnWarning

# fPromptOnWarning

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**fPromptOnWarning**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the debugger prompts the user when debugging generates warnings. Warnings usually indicate errors in an application's interaction with the Window Manager.

| Value | Meaning |
|---|---|
| 0 | Do not prompt. |
| 1 | Prompt. |

**Note**

- This entry appears only when running a debug build of the operating system.

# load

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**load**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | [*Path*]*File name(s)* | (This entry has no default value.) |

**Description**
Specifies the programs that run when Windows Server 2003 starts. The programs are minimized (they appear as taskbar icons) while they are running.

This entry stores a list of file names of programs or files associated with an application. Each file name must be separated from the others by a space. If the file is not located in the *Systemroot* directory, then you must specify the path to the file.

**Change Method**
To load a program at startup, add the application to *Systemroot*\Documents and Settings\All Users\Start Menu\Programs\Startup. Right-click the application icon, click **Properties**, click the **Shortcut** tab, then in the **Run** list, select **Minimized**. This will not actually change the value of the **load** registry entry, but it will emulate the behavior.

**See Also**
**Concepts**
run Entry

# NetMessage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetMessage**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | yes \| no | no |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# NullPort

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NullPort**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *String* \| None | (None) |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

**Caution**

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# Windows\Programs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Programs**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *File name extensions* (no periods) | com exe bat pif cmd |

**Description**
Specifes the file types that Windows Server 2003 considers to be applications. Each file name extension must be separated from the others by a space.

**See Also**
**Concepts**
Documents

# run Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**run**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | [*Path*]*File name(s)* | (This entry has no default value.) |

**Description**
Specifies the applications that the user sets Windows Server 2003 to run at startup.

This entry stores a list of file names of applications or documents associated with an application. Each file name must be separated from the others by a space. If the file is not located in the *Systemroot* directory, you must specify the path to the file.

# Winlogon Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Winlogon**
HKCU\Software\Microsoft\Windows NT\CurrentVersion

**Description**
The Winlogon subkey stores configuration data for the Winlogon component of Windows Server 2003. Winlogon authenticates an interactive user and enables the user to log on. Winlogon supports desktops, the **Ctrl+Alt+Del** secure attention sequence, screen saver control, network provider support, authentication, and password updates.

### Note

- Many of the entries in this subkey can be configured or overridden by the Security Configuration Editor in Group Policy.

# ReportDC

Updated: November 10, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ReportDC**
HKU\.DEFAULT\Software\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Stores the user setting for the domain controller message. (The computer setting is stored in the value of ReportControllerMissing.)

This entry determines, in part, whether the system displays the following message when it cannot contact the domain controller that stores a user's roaming user profile:

```
A domain controller for your domain could not be contacted. You have been logged on using cached account
information. Changes made to your profile since you last logged on might not be available.
```

This entry stores the setting of the **Do not display this message again** check box on the domain controller message dialog box. When the checkbox is selected, the system sets the value of this entry to 0.

| Value | Meaning |
|---|---|
| 0 | Do not display the domain controller message. (The check box is selected.) |
| 1 (or not in the registry) | Display the domain controller message. (The check box is cleared.) |

Two registry entries manage the display of the domain controller message: **ReportDC** and ReportControllerMissing. The message is displayed only when both entries are set to their display setting, that is, when **ReportDC** is either set to 1 or does not appear in the registry, and when **ReportDomainController** appears in the registry with a value of TRUE. By default, the message is not displayed.

**Change Method**
The value of this entry is changed when you check or clear **Do not display this message again** on the domain controller message box.

### Note

- While the domain controller message is displayed, other elements of the user interface are disabled, but all processes continue to run normally. If the **OK** button accompanying the message is not clicked after two minutes, the message box disappears.

**See Also**
**Concepts**
NoPopUpsOnBoot
ReportControllerMissing

# Winlogon\RunLogonScriptSync

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RunLogonScriptSync**
HKCU\Software\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the system waits for the logon script to finish running before it starts Windows Explorer and creates the desktop.

| Value | Meaning |
|---|---|
| 0 | The logon script and Windows Explorer can run simultaneously. |
| 1 | Windows Explorer does not start until the logon script has finished running. |

- For detailed information about particular Group Policy settings, see the *Windows Server 2003 Technical Reference*.
- For general information about Group Policy, see Windows Server 2003 Help or Windows XP Professional Help.

**Change Method**
To change the value of this entry, use Group Policy. This entry corresponds to the **Run logon scripts synchronously** Group Policy setting (User Configuration\Administrative Templates\System\Logon/Logoff).

### Notes

- This entry also appears in the CurrentVersion\Winlogon subkey in HKEY_LOCAL_MACHINE (see Winlogon\RunLogonScriptSync Subkey). The value in HKEY_LOCAL_MACHINE applies to all users. The value in HKEY_CURRENT_USER is configured separately for each user.
- This entry can be superseded by Group Policy settings included in Windows Server 2003. While the **Run logon scripts synchronously** Group Policy setting is enabled (in the Computer Configuration or User Configuration policy folders), the system ignores this entry. The configuration of the **Run logon scripts synchronously** policy setting is stored in the values of **RunLogonScriptSync** in HKEY_LOCAL_MACHINE and of **RunLogonScriptSync** in HKEY_CURRENT_USER.

**See Also**
**Concepts**
Winlogon\RunLogonScriptSync Subkey

# Microsoft\Windows Script

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Script**
HKCU\Software\Microsoft\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Windows Script\Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Settings**
HKCU\Software\Microsoft\Windows Script\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Settings\Enabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabled**
HKCU\Software\Microsoft\Windows Script\Settings

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether Windows Script Host (WSH) is disabled for the current user. This entry disables all scripts that rely on WSH, including VBScript and Jscript. If WSH is disabled and the user attempts to run a WSH script, then a message alerts the user that WSH is disabled. The message appears even if the user attempts to run the script from a bat ch file or a designated script host.

| Value | Meaning |
|---|---|
| 0 | WSH is disabled for the current user. |
| 1 | WSH is enabled for the current user. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
Enabled Entry

# Software\Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policies**
HKCU\Software

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Policies\Microsoft

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft**
HKCU\Software\Policies

**Description**
The Microsoft subkey contains entries that store Group Policy settings for Microsoft products. Group Policy creates and maintains the entries in this subkey.

This subkey stores entries that are configured separately for each user. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the User Configuration folder in Group Policy.

The Policies\Microsoft Subkey subkey in HKEY_LOCAL_MACHINE stores entries for policy settings that apply to all users of the computer. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the Computer Configuration folder in Group Policy.

**Change Method**
To add, remove, or change the entries in this subkey, use Group Policy.

**See Also**
**Concepts**
Policies\Microsoft Subkey

# Microsoft\Control Panel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Control Panel**
HKCU\Software\Policies\Microsoft

**Description**
The Control Panel subkey contains entries that store Group Policy settings for Control Panel programs. Group Policy creates and maintains the entries in this subkey, and the Control Panel programs read and interpret them.

**Change Method**
To add, remove, or change the entries in this subkey, use Group Policy.

### Note

- The entries in this subkey take precedence over entries of the same name in the HKCU\Control Panel Key and HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Control Panel subkeys.

# Control Panel\International

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**International**
HKCU\Software\Policies\Microsoft\Control Panel

**Description**
The International subkey contains entries that store Group Policy settings for calendar options in Regional Options in Control Panel. Group Policy reads and maintains the entries in this subkey, and Regional Options reads and interprets them.

There is also a Regional Options setting for the desktop in Group Policy, which is stored in the values of the entries in the Control Panel\Desktop subkey.

**Change Method**
To change the value of this entry, use Group Policy.

### Note

- The entries in this subkey, which all relate to Group Policy setting, take precedence over the entries in the International Subkey subkey in the HKCU\Control Panel subkey, which do not relate to Group Policy settings.

**See Also**
**Concepts**
International Subkey

# Calendars

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Calendars**
HKCU\Software\Policies\Microsoft\Control Panel\International

**Description**
The Calendars subkey contains entries that store Group Policy settings for calendar options in Regional Options in Control Panel. Group Policy reads and maintains the entries in this subkey, and Regional Options reads and interprets them.

There is also a Regional Options setting for the desktop in Group Policy, which is stored in the values of the entries in the Control Panel\Desktop subkey.

**Change Method**
To change the value of this entry, use Group Policy.

### Note

- The entries in this subkey, which all relate to Group Policy settings, take precedence over the entries in HKCU\Control Panel\International\Calendars subkey, which do not relate to Group Policy settings.

# Microsoft\MMC Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MMC**
HKCU\Software\Policies\Microsoft

**Description**
The MMC subkey contains entries that store Group Policy settings for the Microsoft Management Console (MMC). The entries in this subkey are written and maintained by Group Policy, and the MMC program reads and interprets them.

# Class-ID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Class-ID**
HKCU\Software\Policies\Microsoft\MMC

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Policies\Microsoft\Windows Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows**
HKCU\Software\Policies\Microsoft

**Description**
The Windows subkey contains entries that store Group Policy settings for components of Windows Server 2003. Group Policy creates and maintains the subkeys and entries in the Windows subkey, and those subkeys and entries are read and interpreted by the component that implements them.

This subkey stores entries that are configured separately for each user. There is also a Windows Subkey subkey in the HKEY_LOCAL_MACHINE subtree, which stores entries that apply to all users of the computer.

**See Also**
**Concepts**
Windows Subkey

# Installer Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installer**
HKCU\Software\Policies\Microsoft\Windows

**Description**
The Installer subkey stores the entries created when you configure a Group Policy setting for Windows Installer. Group Policy creates and maintains these entries, and Windows Installer reads and interprets them.

This subkey stores policy-related entries that are configured separately for each user. The system adds these entries to the registry when you enable (and, for some policies, when you disable) a policy in the User Configuration folder in Group Policy.

There is also an Windows\Installer Subkey subkey in the Policies subkey in the HKEY_LOCAL_MACHINE subtree, which stores entries for policies that apply to all users of the computer. The system adds these entries to the registry when you enable (and, for some policies, when you disable) a policy in the Computer Configuration folder in Group Policy.

**Change Method**
To add, remove, or change the value of entries in this subkey, use Group Policy.

**See Also**
**Concepts**
Windows\Installer Subkey

# UNICODE Program Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**UNICODE Program Groups**
HKCU

**Description**
The UNICODE Program Groups key remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this key.

### Note

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# HKEY_LOCAL_MACHINE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM

**Description**
The HKEY_LOCAL_MACHINE subtree contains information about the local computer system, including hardware and operating system data, such as bus type, system memory, device drivers, and startup control parameters.

# Hardware Key

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM

**Description**
The HARDWARE key stores data about the computer hardware that the system detects. This key is recreated each time the system starts. It includes information about devices, and about the software drivers and other resources associated with devices.

# HARDWARE\DESCRIPTION

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\HARDWARE

**Description**
The DESCRIPTION subkey is a repository for data collected by the Windows Server 2003 hardware recognizer. It includes information about the BIOS, processors, and bus architecture. The DESCRIPTION subkey is recreated each time the system starts.

# System Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\HARDWARE\DESCRIPTION

**Description**
The System subkey stores data collected by the Windows Server 2003 hardware recognizer. It includes information about the BIOS, processors, and bus architecture. The System subkey is recreated each time the system starts.

# DEVICEMAP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\HARDWARE

**Description**
The DEVICEMAP subkey contains a set of subkeys representing the base classes of all devices on the computer. A driver that controls a device stores the base class name in an entry in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services subkey.

# SAM Key

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM

**Description**
The SAM key contains information used by the Security Accounts Manager. Servers that are running Windows Server 2003 and are acting as domain controllers use this key to store user and group information pertaining to the directory service restore mode. Domain controllers maintain domain user and group information in Active Directory to manage user security. However, Windows NT 4.0 and earlier, as well as Windows Server 2003s that are not part of a Windows Server 2003 domain, use SAM. All the data in this key is in binary form.

The HKLM\SECURITY\SAM subkey stores a duplicate of the information in this key.

The information in the SAM key is not viewable by users of Windows XP Professional, and should not be edited directly in Windows Server 2003.

**Change Method**
To change the values of the entries in this key, use Active Directory.

### Caution

- Do not alter or add entries in the SAM key. Doing so could prevent users from logging on to the domain or their computers, and it could require you to restore your entire system.

# Security Key

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM

**Description**
The SECURITY key contains security information used by the system and network.

The information in the SECURITY key is not viewable by users of Windows XP Professional, and should not be edited directly in Windows Server 2003. All the data in this key is in binary form.

**Change Method**
To change configurations that are specified by information in this key, use the tools designed for configuring Group Policy, Security Policy, and Active Directory, and other tools designed for managing computers and users.

### Caution

- Do not modify or add entries in the SECURITY key. Doing so could prevent users from logging on to the domain or their computers, and it could require you to restore your entire system.

# HKLM\SOFTWARE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM

**Description**
The SOFTWARE key stores program variables that apply to all users of the computer. The system stores some of its variables in the subkeys in SOFTWARE. In addition, the SOFTWARE subkey contains subkeys representing the vendors, including Microsoft, who produce the software running on the system.

**See Also**
**Concepts**
SOFTWARE Subtree

# SOFTWARE\Microsoft

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE

**Description**
The Microsoft subkey stores configuration data for Microsoft programs running on the computer. The subkeys in the Microsoft subkey represent each of the programs, and they store settings that affect only that program.

The Microsoft subkey in HKEY_LOCAL_MACHINE stores entries that apply to all users of the computer. The Microsoft subkey in HKEY_CURRENT_USER stores entries that apply only to the current user of the computer.

**See Also**
**Concepts**
Microsoft Subkey

# Microsoft\Command Processor

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

**Description**
The Command Processor subkey stores user preferences for the command processor (Cmd.exe).

**Change Method**
You can change the values of the entries in this subkey by using Command Processor (Cmd.exe), but only for the duration of that Cmd.exe's execution. For all other changes, you must edit the registry directly.

**Activation Method**
If you make changes to the value of the entries in this subkey by editing the registry directly, you must restart Cmd.exe to make the changes effective.

## Notes

- Entries in the Command Processor subkey can also appear in HKLM. If both appear, in most cases the values in HKCU take precedence over the values in HKLM. However, the HKLM **Autorun** entry executes before the HKCU **Autorun** entry. Command-line settings entered in Cmd.exe take precedence over registry settings.

- For more information about Command Processor, see Windows Server 2003 Help or Windows XP Professional Help.

**See Also**
**Concepts**
Command Processor Subkey

# Command Processor\AutoRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *command* | (None) |

**Description**
Contains commands that are run each time you start Command Processor (Cmd.exe).

**Change Method**
You can change the value of this entry by using Cmd.exe and typing **CMD /D**, but only for the duration of that Cmd.exe's execution. For all other changes, you must edit the registry directly.

**Activation Method**
To make the value of this entry effective, you must restart Cmd.exe.

### Note

- **AutoRun** can appear in HKLM and HKCU. If both appear in the registry, the system first runs the value of the entry in HKLM, and then it runs the value in HKCU.

### Tip

- For more information about the AutoRun feature, at the command prompt, type **CMD /?**.

**See Also**
**Concepts**
AutoRun Entry

# Command Processor\CompletionChar

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CompletionChar**
HKLM\SOFTWARE\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0x1F \| 0x20–0xFF | 0x40 (*Completion disabled*) |

**Description**
Specifies a character that activates the file-name completion feature in the command processor (Cmd.exe).

To specify a character, enter the hexadecimal value of an ASCII control character. For example, 0x4 represents CTRL+D and 0x6 represents CTRL+F. To disable the file-name completion feature, set the value of this entry to 0x0 or to a value greater than or equal to 0x20.

| Value | Meaning |
|---|---|
| 0x0 | Disables completion. |
| 0x1–0x1F | Corresponding ASCII control character activates completion. |
| 0x20 or greater | Disables completion. |

**Change Method**
To change the completion character in Windows Server 2003, use the registry editor Regedit.exe. You can also change the value of this entry by using Cmd.exe, but only for the duration of that Cmd.exe's execution.

**Activation Method**
If you make changes to the value of this entry by editing the registry directly, you must restart Cmd.exe to make the changes effective.

### Notes

- **CompletionChar** can appear in HKLM and HKCU. If both appear in the registry, then the value of the entry in HKCU takes precedence over the value in HKLM. The command-line parameters take precedence over the registry settings while using Cmd.exe.
- Assigning the values 0x0, 0xD, or 0x20 to this entry disables file name completion rather than assigning a completion character.
- If you use the console Paste function to enter text into a CMD window, you might trigger completion if the characters you paste contain completion characters.

### Tip

- For more information about the file name completion feature, at the command prompt, type: **CMD /?**
- The datatype of this entry can be either REG_DWORD or REG_SZ.

### Tip

- To enable or disable file name completion and directory name completion for one use of Cmd.exe, at the command prompt, type **CMD /F:ON** (to enable) or **CMD /F:OFF** (to disable). This command sets the completion character to CTRL+F (0x6).

- For more information about the file-name completion feature, at a command prompt, type **CMD /?**

**See Also**
**Concepts**
PathCompletionChar Entry
CompletionChar Entry

# Command Processor\DefaultColor

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFE | 0x0 |

**Description**
Specifies a default background color and text color for the window in Command Processor (Cmd.exe). The value consists of two hexadecimal digits. The first digit specifies the background color; the second digit specifies the text color. The hexadecimal values correspond to the following colors:

| Hexadecimal value | Color |
|---|---|
| 0 | Black |
| 1 | Blue |
| 2 | Green |
| 3 | Aqua |
| 4 | Red |
| 5 | Purple |
| 6 | Yellow |
| 7 | White |
| 8 | Gray |
| 9 | Light Blue |
| A | Light Green |
| B | Light Aqua |
| C | Light Red |
| D | Light Purple |
| E | Light Yellow |
| F | Bright White |

This entry is used only when its value is valid and not 0x0, and when the user has not specified the window color by using the COLOR or /T commands. Values are not valid if the background color and the text color are identical, such as 0x55. If the value of this entry is not valid or is 0x0, and the user has not entered a command to change the color, then the command processor uses the colors specified for the console window in which the command processor runs.

**Activation Method**
If you make changes to the value of this entry by editing the registry directly, then you must restart Cmd.exe to make the changes effective.

### Notes

- **DefaultColor** can appear in HKLM and HKCU. If both appear in the registry, the value of the entry in HKCU takes

precedence over the value in HKLM.

- The datatype of this entry can be either REG_DWORD or REG_SZ.

**See Also**
**Concepts**
DefaultColor Entry
ScreenColors

# Command Processor\DelayedExpansion

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the command processor supports delayed expansion of environment variables. If the value of this entry is 1, then the command processor interprets an exclamation point (**!**) as an environment variable expands only when used.

By default, the command processor expands variables when it reads a line of text. But you can direct the command processor to expand a variable only when it is used. As a result, you can include variables in compound conditional statements and prevent them from being set before they are tested.

| Value | Meaning |
|---|---|
| 0 | Delayed expansion is not supported. The command processor expands environment variables when it reads them. It does not support exclamation points as delayed variables. |
| 1 | Delayed expansion is supported. The system expands environment variables surrounded by exclamation points (such as !VAR!) only when they are used. |

**Change Method**
You can change the value of this entry by using Command Processor (Cmd.exe) and typing **cmd /v:on**, but only for the duration of that Cmd.exe's execution. For all other changes, you must use the registry editor Regedit.exe.

**Activation Method**
If you make changes to the value of this entry by using Regedit.exe, then you must restart Cmd.exe to make the changes effective.

## Notes

- The **DelayedExpansion** entry can appear in HKEY_LOCAL_MACHINE and in HKEY_CURRENT_USER. If both appear in the registry, then the value of the entry in HKEY_CURRENT_USER takes precedence over the value in HKEY_LOCAL_MACHINE.
- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe. Even when this entry appears in the registry, the setting of the command line parameter **/V** takes precedence over this value.
- The datatype of this entry can be either REG_DWORD or REG_SZ.

## Tip

- For more details, start the Command Processor (Cmd.exe), type either **help set** or **cmd /?**, and read the entry on the **/V** command.

**See Also**
**Concepts**
DelayedExpansion Entry

# Command ProcessorEnableExtensions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the extensions to Command Processor (Cmd.exe), such as chdir and setlocal, are enabled. Because the extensions might be incompatible with a command-script language, Windows Server 2003 lets you disable the extensions. This entry affects all extensions. You cannot use it to enable or disable extensions selectively. You can, however, disable extensions for a particular invocation of Cmd.exe by using the **/E:OFF** switch.

| Value | Meaning |
|---|---|
| 0 | Disabled |
| 1 | Enabled |

**Change Method**
You can change the value of this entry by using Cmd.exe and typing **cmd /e:off**, but only for the duration of that Cmd.exe's execution. For all other changes, you must edit the registry directly.

**Activation Method**
If you make changes to the value of this entry by editing the registry directly, you must restart Cmd.exe to make the changes effective.

### Notes

- **EnableExtensions** can appear in HKLM and HKCU. If both appear in the registry, then the value of the entry in HKCU takes precedence over the value in HKLM.

- The datatype of this entry can be either REG_DWORD or REG_SZ.

**See Also**
**Concepts**
EnableExtensions Entry

# Command Processor\PathCompletionChar

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Command Processor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0x1F \| 0x20–0xFF | 0x40 (*Completion disabled*) |

**Description**
Specifies a character that activates the directory name completion in the command processor (Cmd.exe).

To specify a character, enter the hexadecimal value of an ASCII control character. For example, 0x4 represents CTRL+D and 0x6 represents CTRL+F. To disable the directory name completion, set the value of this entry to 0x0 or to a value greater than or equal to 0x20.

| Value | Meaning |
|---|---|
| 0x0 | Disables completion. |
| 0x1–0x1F | Corresponding ASCII control character activates completion. |
| 0x20 or greater | Disables completion. |

**Change Method**
To change the completion character in Windows Server 2003, use the registry editor Regedit.exe. You can also change the value of this entry by using Cmd.exe, but only for the duration of that Cmd.exe's execution.

**Activation Method**
If you make changes to the value of this entry by editing the registry directly, you must restart Cmd.exe to make the changes effective.

### Note

- **PathCompletionChar** can appear in HKLM and HKCU. If both appear in the registry, then the value of the entry in HKCU takes precedence over the value in HKLM.

### Tip

- To enable or disable directory name completion for one use of Cmd.exe, at the command prompt, type **CMD /F:ON** (to enable) or **CMD /F:OFF** (to disable). This command sets the completion character to CTRL+D (0x4).

- For more information about the path completion feature, at the command prompt, type **CMD /?**

**See Also**
**Concepts**
CompletionChar Entry
PathCompletionChar Entry

# Microsoft\Driver Signing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

**Description**
The Driver Signing subkey in HKLM stores entries for the driver signing feature. The entries in this subkey apply to all users of the computer.

Microsoft digitally signs driver files to affirm that the driver software has been tested with Windows and has not been altered since it was tested.

### Note

- The entries in this subkey apply to all users of the computer. The entries in the Driver Signing Subkey subkey in HKCU apply only to the current user.

### Tip

- For more information about driver signing, see Windows Server 2003 Help or Windows XP Professional Help.

**See Also**
**Concepts**
Driver Signing Subkey

# Driver Signing\Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Driver Signing

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | 00 \| 01 \| 02 | 01 |

**Description**
Specifies how the system responds when a user tries to install device driver files that are not digitally signed.

| Value | Meaning |
|---|---|
| 00 | Ignore. The installation proceeds, even if it includes unsigned files. |
| 01 | Warn. The system notifies the user that files are not digitally signed, and it prompts the user to decide whether to stop the installation or to proceed with it and permit unsigned files to be installed. |
| 02 | Block. The system will not install unsigned files. As a result, the installation stops, and none of the files in the driver package are installed. |

This entry applies to all users of the computer. However, if a stricter setting is configured for the current user (in Policy Entry in HKCU) or a user Group Policy setting establishes a stricter setting (in **BehaviorOnFailedVerify**), then this value is ignored for the current user.

This entry stores the option selected in the **File signature verification** section of the **Driver Signing Options** dialog box when the **Apply setting as system default** box is selected. If the box is not selected, then the selection applies only to the current user, and it is stored in the value of Policy Entry in HKCU\SOFTWARE\Microsoft\Driver Signing.

**Change Method**
To change the value of this entry, use System in Control Panel. Click the **Hardware** tab, click **Driver Signing**, select an option in the **File signature verification** section, and then click the **Apply setting as a system default** box.

### Note

- This entry applies to all users of the computer. The **Policy** entry also can appear in HKCU for a particular user.

**See Also**
**Concepts**
Policy Entry

# DrWatson

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

**Description**
The DrWatson subkey stores configuration settings for Dr. Watson for Windows (drwtsn32.exe), the application error debugger. These entries specify how Dr. Watson for Windows notifies the user of an application error, and how its log file and crash dump file are configured.

Dr. Watson for Windows installs with Windows Server 2003. The registry adds these entries in this subkey the first time Dr. Watson for Windows starts on the system.

**Change Method**
To change the values of the entries in this subkey, use Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

# AppendToLogFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether Dr. Watson for Windows overwrites the existing log file with new data or appends new data to the end of the existing log file. By default, the log file is named Drwtsn32.log and is stored in the system root directory.

| Value | Meaning |
|---|---|
| 0 | Do not append. Dr. Watson for Windows overwrites the existing log file each time it records data about an application error. |
| 1 | Append. Dr. Watson for Windows appends new data to the end of the existing log file. |

**Change Method**
To change the value of this entry, use the **Append to Existing Log File** option in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

**Tip**

- For a detailed description of the log file of Dr. Watson for Windows, start Drwtsn32.exe, click **Help**, and then double-click **Log File Description**.

# CrashDumpFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *File path & name* | %windrive%\Documents and Settings\All Users\Application Data\Microsoft\Dr Watson\user.dmp |

**Description**
Specifies the location of the crash dump file for Dr. Watson for Windows. The crash dump file is a binary file that is created when a program stops abnormally. The crash dump file can be interpreted by the Windows Debugger.

All users must have permission to read and write to the crash dump file. By default, the crash dump file is named user.dmp and resides in the root of the system directory.

This entry is used only when the value of CreateCrashDump is 1.

**Change Method**
To change the value of this entry, enter a file name in the **Crash Dump** text box in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

**See Also**
**Concepts**
CreateCrashDump

# CreateCrashDump

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether Dr. Watson for Windows creates a binary crash dump file when a program stops abnormally. The crash dump file can be interpreted by the Windows Debugger.

If the value of this entry is 1, then the location of the crash dump file must appear in the value of the CrashDumpFile entry. By default, the crash dump file is named user.dmp and is in the system root directory.

**Change Method**
To change the value of this entry, use the **Create Crash Dump File** option in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

**See Also**
**Concepts**
CrashDumpFile

# DumpAllThreads

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies which application threads are included in the log file that Dr. Watson for Windows creates when an application error occurs. By default, the log file is named Drwtsn32.log and is in the system root directory.

| Value | Meaning |
|---|---|
| 0 | Error thread only. The log file includes data only for the thread that caused the application error. |
| 1 | All threads. The log file includes data for all threads that run in the process in which the application error occurred. |

**Change Method**
To change the value of this entry, use the **Dump All Thread Contexts** option in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

**Tip**

- For a detailed description of the log file of Dr. Watson for Windows, start Drwtsn32.exe, click **Help**, and then double-click **Log File Description**.

# DumpSymbols

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether Dr. Watson for Windows includes the symbol table for each application module in its log file. The symbol table consists of the name and memory address of each symbol. By default, because it can be very large, the symbol table is not included in the log file.

| Value | Meaning |
|---|---|
| 0 | Omit symbol table |
| 1 | Include symbol table |

**Change Method**
To change the value of this entry, use the **Dump Symbol Table** option in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

**Tip**

- For a detailed description of the log file of Dr. Watson for Windows, start Drwtsn32.exe, click **Help**, and then double-click **Log File Description**.

# Instructions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x1F4 (*0–500*) | 0xA (*10*) |

**Description**
Specifies how many instructions preceding and following the faulty instruction are included in the instruction disassembly portion of Drwtsn32.log, the log file for Dr. Watson for Windows.

The instruction disassembly includes:

- The function that was being executed when the error occurred;
- The memory address, raw machine instruction, and decoded machine instruction for each adjacent instruction; and
- A complete analysis of the faulty instruction.

**Change Method**
To change the value of this entry, enter a number in the **Number of Instructions** text box in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

**Tip**

- To see an example of the instruction disassembly portion of the Dr. Watson for Windows log file, start Drwtsn32.exe, click **Help**, double-click **Log File Description**, and then scroll through the file.
- For a detailed description of the Dr. Watson for Windows log file, start Drwtsn32.exe, click **Help**, and then double-click **Log File Description**.

# LogFilePath Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Path name* | %windrive%\Documents and Settings\All Users\Application Data\Microsoft\Dr Watson |

**Description**
Specifies the location of Drwtsn32.log, the log file for Dr. Watson for Windows.

The log file contains diagnostic data about application errors that occur. All users must have permission to read and write to the log file.

**Change Method**
To change the value of this entry, enter a path in the **Log File Path** box in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

**Tip**

- For a detailed description of the log file of Dr. Watson for Windows, start Drwtsn32.exe, click **Help**, and then double-click **Log File Description**.

# MaximumCrashes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0xA |

**Description**
Specifies how many errors Dr. Watson for Windows maintains at any given time in its application error viewer and in the Event Viewer application log. When the number of recorded errors reaches the value of this entry, Dr. Watson operates normally and adds new errors to its log and crash dump file, but it does not add any new errors to its application log viewer or to the Event Viewer application log until it is reset.

The user is not notified that the number of crashes has reached the maximum.

To reset Dr. Watson for Windows, run Drwtsn32.exe. You can either click **Clear** or increase the value in the **Number of Errors to Save** box. When you click **Clear**, you delete all errors from the application error viewer in Dr. Watson for Windows, and you delete all events from the Event Viewer application log, even those not generated by Dr. Watson for Windows. However, records of the application errors remain in the Drwtsn32.log and in the crash dump file, user.dmp, if one has been selected.

**Change Method**
To change the value of this entry, enter a number in the **Number of Errors to Save** box in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

**Tip**

- The number of application errors Dr. Watson for Windows has analyzed since it was last reset is stored in the value of the NumberOfCrashes entry.

# NumberOfCrashes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–*Value of the* MaximumCrashes*entry* | (None) |

**Description**
Counts the number of application errors Dr. Watson for Windows is maintaining in its application error viewer and in the Event Viewer application log. When this value reaches the value of MaximumCrashes**,** Dr. Watson for Windows operates normally and adds new errors to its log and crash dump file, but it does not add any new errors to its application log viewer or to the Event Viewer application log until it is reset.

The user is not notified that the number of crashes has reached the maximum.

To reset Dr. Watson for Windows, run Drwtsn32.exe. You can either click **Clear** or increase the value in the **Number of Errors to Save** box. When you click **Clear**, you delete all errors from the application error viewer in Dr. Watson for Windows, and you delete all events from the Event Viewer application log, even those not generated by Dr. Watson for Windows. However, records of the application errors remain in the Drwtsn32.log and in the crash dump file, user.dmp, if one has been selected.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

### Caution

- Do not change the value of this entry. Dr. Watson for Windows uses this entry to maintain an accurate count of application errors.

# SoundNotification

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether Dr. Watson for Windows generates a sound when it detects an application error. If the value of this entry is 1, then you can enter the name of a wave file (.wav) in the value of the WaveFile entry. Otherwise, Dr. Watson generates two standard beeps.

This entry is used only when the computer is equipped with a sound card.

| Value | Meaning |
|---|---|
| 0 | No sound |
| 1 | Beeps or plays a specified wave file when an application error is detected |

**Change Method**
To change the value of this entry, use the **Sound Notification** option in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

# VisualNotification

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether Dr. Watson for Windows displays a message box when it detects an application error.

While the message box is displayed, no new messages or windows can be displayed, but background processes are not affected. If the buttons on the message box are not clicked within five minutes, the message box is removed.

**Change Method**
To change the value of this entry, use the **Visual Notification** option in Dr. Watson for Windows. To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

# WaveFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\DrWatson

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *File name* | (None) |

**Description**
Specifies the file containing the sound that Dr. Watson for Windows plays when it detects an application error.

This entry is used only when the computer has a sound card and the value of the SoundNotification entry is 1.

**Change Method**
To change the value of this entry, use the **Wave File** option in Dr. Watson for Windows. (If the computer does not have a sound card, this option is disabled.) To start Dr. Watson for Windows, from the **Start** menu, click **Run**, and then type **drwtsn32**.

**Activation Method**
This entry appears in the registry when you run Dr. Watson for Windows or when it runs automatically.

# Fax Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\

**Description**
The Fax sub key stores configuration data for the Fax Service. This sub key and its contents appear only when you install the Fax Service.

**Change Method**
You must restart the Fax Service to make changes to entries in this subkey effective.

# LastUniqueLineId

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
| --- | --- | --- |
| REG_DWORD | | |

**Description**
Identifies fax devices. This entry is used and maintained automatically by the Fax Service.

...

### Caution

- Do not modify this entry.

# CfgWzdrDevice

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the Fax Configuration wizard has configured the Fax service. This entry is used by the Fax Configuration Wizard.

| Value | Meaning |
|---|---|
| 0 | The Fax service has not yet been configured by the Fax Configuration Wizard. |
| 1 | The Fax service has been configured by the Fax Configuration Wizard. |

**Change Method**
This entry is set by the Fax Configuration Wizard when it configures the Fax service.

# DisableRoutingExtensionConfiguration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the Fax Configuration Wizard displays the **Routing Options** page during configuration of the Fax service.

| Value | Meaning |
|---|---|
| 0 | The **Routing Options** page of the Fax Configuration Wizard is displayed. |
| 1 | The **Routing Options** page of the Fax Configuration Wizard is not displayed. |

The **Routing Options** page allows the user to specify the printer to which the Fax service sends faxes for printing and the folder to which it saves additional copies of incoming faxes. (Saved copies of incoming faxes are also available from the Inbox archive of the Fax console.)

**Change Method**
You can change this entry only by editing the registry.

**Tip**

- If you have a custom method for routing configuration, consider hiding the standard method so that users cannot select it.

# Inbound Profile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
| --- | --- | --- |
| REG_SZ | | |

**Description**
Obsolete. This entry is no longer used by the Fax service, but it might continue to appear in some registries.

# ManualAnswerDevice

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 0 |

**Description**
Specifies the ID of the fax device that is set to manual answer mode. Only one device at a time can be set to manual answer mode. If the value of this entry is zero, then no devices are configured for manual answer mode.

If a fax device is configured for manual answer mode, then Fax Monitor displays an alert when the telephone line of the fax device is ringing, and gives the user the option of answering the call with Fax Monitor to receive a fax transmission.

**Change Method**
To configure a fax device for manual answer mode, click the **Start** menu, point to **Settings**, and then select **Printers and Faxes**. Right-click the name of the fax device, and then select **Properties**. Click the **Devices** tab, and then click **Properties** under the name of the fax-capable hardware (usually a modem). In the new window that opens (its title is the same as the name of the fax-capable hardware), on the **Receive** tab, click **Manual** under **Answer Mode**.

You can also change this entry by using the Fax application programming interface (API).

**Activation Method**
Changes made to this entry through **Printers and Faxes**, or through the Fax API, are effective immediately.

# MaxLineCloseTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 300 |

**Description**
Specifies the time, in seconds, that the Fax Service waits between attempts to turn on a fax device that is turned off.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must restart the Fax Service for changes to this entry to take effect.

# NextJobNumber

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 1 |

**Description**

Specifies the next job number for the Fax Service. This entry is maintained automatically by the service.

**Caution**

Do not modify this entry.

# RecipientsLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 0 |

**Description**
Specifies the maximum number of recipients of a single fax broadcast. If the value of this entry is zero, or if this entry is absent from the registry, the Fax Service imposes no limit.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must restart the Fax Service for changes to this entry to take effect.

# QueueDirectory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | %ALLUSERSPROFILE%\Application Data\Microsoft\Windows NT\MSFax\Queue |

**Description**
Specifies the path of the folder in which the Fax Service stores queued faxes. This entry does not exist in the registry by default. You must add this entry if you want to change the default location of the fax queue.

**Change Method**
You can add or change this entry only by editing the registry.

**Activation Method**
You must restart the Fax Service for changes to this entry to take effect.

**Tip**

- You only need this entry if you want to change the default location of the fax queue.

# QueueState

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | bit mask | 0 |

**Description**
Specifies incoming and outgoing actions to prohibit on the fax queue of the Fax Service. This entry is a bit mask of the values shown in the following table.

| Value | Meaning |
|---|---|
| 1 | Fax Service does not receive new incoming faxes. |
| 2 | Fax Service does not accept new outgoing fax jobs. |
| 4 | Fax Service does not send outgoing faxes. |

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager from the **Computer Management** console.

To change the value of this entry by using the Computer Management console, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **General** tab. This entry corresponds to the three options on the **General** tab that you can clear or check:

- **Disable submission of new outgoing faxes**
- **Disable transmission of outgoing faxes**
- **Disable reception of new faxes**

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

# Branding

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether or not the pages of outgoing fax transmissions have a banner. The banner includes the date and time of transmission, the transmitting subscriber identification (TSID), the current page number, and the telephone number of the recipient's fax device. The banner is prepended to all pages of outgoing fax transmissions.

| Value | Meaning |
|---|---|
| 0 | The system does not add a banner to outgoing fax transmissions. |
| 1 | The system adds a banner to all pages of outgoing fax transmissions. |

**Change Method**
To change the value of this entry, click on the **Start** menu, point to **Settings**, and then select **Printers and Faxes**. Right-click the name of the fax device, and then select **Properties**. Click the **Devices** tab, and then click **Properties** under the name of the fax-capable hardware (usually a modem). In the new window that opens (its title is the same as the name of the fax-capable hardware), select or clear the **Include banner** check box.

This entry can also be changed by using the Fax application programming interface (API).

**Activation Method**
Changes made to this entry though **Printers and Faxes** or through the Fax API are effective immediately.

**See Also**
**Concepts**
{F10A5326-0261-4715-B367-2970427BBD99}\TSID
FaxNumber

# ActivityLogging

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

The ActivityLogging subkey stores settings for activity logging by the Fax service.

# DBFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\ActivityLogging

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | C:\Documents and Settings\All Users\Application Data\Microsoft\Windows NT\MSFax\ActivityLog |

**Description**
Specifies the location in the file system where the Fax Service creates and maintains files that log incoming and outgoing fax activity. The Fax Service only logs activity if the values in the registry entries LogIncoming and LogOutgoingspecify that the service keeps logs.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, click the **Activity Logging** tab, and then type a file path and name into the **Activity log folder location** box. Click **OK**.

**Activation Method**
Changes made through **Fax Service Manager** or **Computer Management** are effective immediately.

### Note

- You can configure the Fax Service to create and maintain logs of incoming and outgoing fax jobs. Incoming activity is logged to a single file (InboxLOG.txt) that contains status, call time, file name, and other fax-specific information. Outgoing activity is logged to a single file (OutboxLOG.txt) that contains fax-specific information that, in addition to the type of information found in the incoming activity log, includes job IDs, submission time, and fax header contents.

**See Also**
**Concepts**
LogIncoming
LogOutgoing

# LogAgeLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\ActivityLogging

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 0 |

**Description**
Specifies the time, in months, that the Fax Service retains activity logs. If the value of this entry is zero, the Fax Service uses a default age limit of one month.

The value of this entry only affects retention of entries in activity logs if the value of the registry entry LogLimitCriteria specifies age as the criteria. In addition, the values of the registry entries LogIncoming and LogOutgoing must specify that the Fax Service keep activity logs.

When the oldest entry in an activity log reaches the age specified by this entry, the Fax Service takes the action specified by the registry entry LogLimitReachedAction.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax service for changes to this entry to take effect.

### Note

- You can configure the Fax Service to create and maintain logs of incoming and outgoing fax jobs.
    - **Incoming activity** is logged to a single file that contains status, call time, file name, and other fax-specific information.
    - **Outgoing activity** is logged to a single file that contains fax-specific information that, in addition to the type of information found in the incoming activity log, includes job IDs, submission time, and fax header contents.

**See Also**
**Concepts**
LogLimitCriteria
LogLimitReachedAction
LogSizeLimit
LogIncoming
LogOutgoing

# LogIncoming

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\ActivityLogging

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 1 |

**Description**
Specifies whether the Fax Service logs incoming fax activity. The log is kept in a single file and contains status, call time, file name, and other fax-specific information for each log entry.

| Value | Meaning |
|---|---|
| 0 | The Fax Service does not log incoming fax activity. |
| 1 | The Fax Service logs incoming fax activity. |

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager from the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, click the **Activity Logging** tab, and then clear or check **Log incoming fax activity**.

**Activation Method**
Changes made through **Fax Service Manager** or **Computer Management** are effective immediately.

### Note

- Outgoing fax activity is logged to a single file that contains fax-specific information. In addition to the type of information found in the incoming activity log, the file includes information the Fax Service used to create the fax, like job IDs, submission time, and banner contents.

**See Also**
**Concepts**
LogOutgoing

# LogLimitCriteria

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\ActivityLogging

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 0 |

**Description**
Specifies the criteria that the Fax Service uses to limit its activity logs.

| Value | Meaning |
|---|---|
| 0 | The Fax Service does not limit its activity logs. This is the default configuration. |
| 1 | The Fax Service limits its activity logs by file size. |
| 2 | The Fax Service limits its activity logs by age. |

The size and age by which the Fax service limits its activity logs are specified, respectively, by the registry entries LogSizeLimit and LogAgeLimit. When an activity log reaches a limit, the Fax Service performs an action on the log as specified by the registry entry LogLimitReachedAction.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax Service to make changes to this entry effective.

**See Also**
**Concepts**
LogSizeLimit
LogAgeLimit
LogLimitReachedAction
LogIncoming
LogOutgoing

# LogLimitReachedAction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\ActivityLogging

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies the action that the Fax Service performs on its activity logs when a log reaches a limit.

| Value | Meaning |
|---|---|
| 0 | The Fax Service copies the activity log to a new file. Then the service erases the original log file, creates a replacement file, and continues logging activity. |
| 1 | The Fax Service deletes the log file. Then the service creates a replacement file and continues logging activity. |

The registry entry LogLimitCriteria specifies the type of limit that the Fax Service imposes on its activity logs. The registry entries LogSizeLimit and LogAgeLimit set the limits on size and age, respectively, for activity logs kept by the Fax Service.

**Change Method**
You can change this value only by editing the registry.

**Activation Method**
You must stop and restart the Fax Service to make changes to this registry entry effective.

### Notes

- Incoming fax activity is logged to a single file named InboxLOG.txt.
- Outgoing activity is logged a single file named OutboxLOG.txt.
- When the Fax Service is configured to copy a log that has reached its limit to a new file, it names the new file for the current date by using the format *4-digit-year-2-digit-month-2-digit-day*.txt. For examle, on April 02, 2002, a newly copied log file would be named 2002-04-02.txt.

**See Also**
**Concepts**
LogLimitCriteria
LogSizeLimit
LogAgeLimit
LogIncoming
LogOutgoing

# LogOutgoing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\ActivityLogging

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 1 |

**Description**
Specifies whether the Fax Service logs outgoing fax activity. The log is kept in a single file and contains job IDs, submission time, banner contents, status, call time, file name, and other fax-specific information for each log entry.

| Value | Meaning |
|---|---|
| 0 | The Fax Service does not log incoming fax activity. |
| 1 | The Fax Service logs incoming fax activity. |

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, click the **Activity Logging** tab, and then clear or check **Log outgoing fax activity**.

**Activation Method**
Changes made through **Fax Service Manager** or **Computer Management** are effective immediately.

### Note

- Incoming fax activity is logged to a single file that contains status, call time, file name, and other fax-specific information for each log entry. Because incoming faxes are not scheduled or created by the local Fax service, the incoming activity log does not contain job IDs, submission time, or other information used to create the fax.

**See Also**
**Concepts**
LogIncoming

# LogSizeLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\ActivityLogging

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 0 |

**Description**
Specifies the maximum size, in megabytes, of activity logs kept by the Fax Service. If the value of this entry is zero, the Fax Service uses a default size limit of 50MB. When an activity log reaches the maximum size specified by this entry, the Fax Service takes the action specified by the registry entry LogLimitReachedAction.

The value of this entry only affects the size of activity logs if the value of the registry entry LogLimitCriteria specifies size as the criteria. In addition, the values of the registry entries LogIncoming and LogOutgoing must specify that the Fax service keep activity logs.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax Service for changes to this entry to take effect.

### Note

- You can configure the Fax Service to create and maintain logs of incoming and outgoing fax jobs.
    - Incoming activity is logged to a single file that contains status, call time, file name, and other fax-specific information.
    - Outgoing activity is logged to a single file that contains fax-specific information that, in addition to the type of information found in the incoming activity log, includes job IDs, submission time, and fax header contents.

**See Also**
**Concepts**
LogLimitCriteria
LogLimitReachedAction
LogAgeLimit
LogIncoming
LogOutgoing

# Device Providers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\

**Description**
The Device Providers subkey contains the list of registered fax service providers (FSPs) for the Fax Service.

# {2172FD8F-11F6-11D3-90BF-006094EB630B}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The {2172FD8F-11F6-11D3-90BF-006094EB630B} subkey stores configuration information about the Microsoft Fax Service Provider (FSP) for the Fax Service.

### Notes

- The Microsoft Fax Service Provider (FSP) is identified by the globally unique identifier 2172FD8F-11F6-11D3-90BF-006094EB630B. The configuration data that the Fax Service uses for the Microsoft FSP is identified by the globally unique identifier F10A5326-0261-4715-B367-2970427BBD99.

- A globally unique identifier (GUID) is a 128-bit value that uniquely identifies objects such as OLE servers, interfaces, manager entry-point vectors, and client objects. GUIDs are used in cross-process communication, such as remote procedure calling (RPC) and OLE.

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol specified by the T.30 specification of the International Telecommunications Union for standard fax modems. Third parties can also supply fax service providers to support other fax-capable devices.

# APIVersion

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**

This entry is no longer used by the Fax Service, but it might still appear in some registries.

# Capabilities

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**

This entry is no longer used by the Fax Service, but it might still appear in some registries.

# FriendlyName Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the display name of the Microsoft Fax Service Provider (FSP). This entry is set when the Microsoft FSP is registered with the Fax Service. This name appears in Fax Service Manager in the list of providers.

**Change Method**
You can change this registry entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax service to make changes to this registry entry effective.

### Note

- Fax Service Manager is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

# ImageName Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | | %systemroot%\system32\fxst30.dll |

**Description**

Specifies the path and file name of the dynamic-link library (DLL) that provides the Fax Service. The value of this entry is set in the registry when a device is registered as a fax service provider by the Fax Service.

# LogFileNumberIncoming

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies the number for the log for the next incoming fax transmission. The Microsoft Fax Service Provider uses this number when naming the log file.

**Change Method**
Do not change this entry.

### Note

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol for standard modems, as described by the T.30 specification of the International Telecommunications Union. Third parties can also supply fax service providers to support other fax-capable devices.

### Caution

- This entry is maintained automatically by the Microsoft Fax Service Provider. Do not modify it.

**See Also**
**Concepts**
LogFileNumberOutgoing
LoggingFolder

# LogFileNumberOutgoing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies the number for the log for the next outgoing fax transmission. The Microsoft Fax Service Provider uses this number when naming the log file.

### Note

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol for standard modems, as described by the T.30 specification of the International Telecommunications Union. Third parties can also supply fax service providers to support other fax-capable devices.

### Caution

- This entry is maintained automatically by the Microsoft Fax Service Provider. Do not modify it.

**See Also**
**Concepts**
LogFileNumberIncoming
LoggingFolder

# LoggingEnabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 2 |

**Description**
Specifies whether or not the Microsoft Fax Service Provider writes logs of fax jobs. Logs are written to the location in the file system that is specified by the registry entry LoggingFolder. They are named for the fax jobs for which they provide transmission and reception detail.

A log consists of the conversation between the Microsoft Fax Service Provider and the fax device.

| Value | Meaning |
|---|---|
| 0 | The Fax Service does not keep a log. |
| 1 | The Fax Service logs all jobs. |
| 2 | The Fax Service logs all failed jobs. |

**Note**

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol specified by the T.30 specification of the International Telecommunications Union for standard fax modems. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
LoggingFolder

# LoggingFolder

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *path and file name* | %windir%\system32\logfiles\fax |

**Description**
Specifies the location in the file system where the Microsoft Fax Service Provider writes logs generated during transmission and reception of fax jobs. Which logs are written is specified by the registry entry LoggingEnabled, which can also specify that no logs are written.

In the folder specified by this entry, the Microsoft Fax Service Provider creates two subfolders, Outgoing and Incoming. Log files in the Outgoing folder are named FSPS*number*.log, where *number* is the number of the corresponding fax job. Similarly, log files in the Incoming folder are named FSPR*number*.log. Log files consist of the conversation between the Microsoft Fax Service Provider and the fax device that sent or received the fax job.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax Service to make changes effective.

### Note

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol for standard modems, as described by the T.30 specification of the International Telecommunications Union. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
LoggingEnabled

# MaxConsecErrorLinesPerPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 110 |

**Description**
Specifies the maximum number of consecutive lines with transmission errors in a received fax that the Microsoft Fax Service Provider allows before requesting that the sender restart transmission of the page.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax Service to make changes to this entry effective.

### Note

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol for standard modems, as described by the T.30 specification of the International Telecommunications Union. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
MaxErrorLinesPerPage

# MaxErrorLinesPerPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 110 |

**Description**
Specifies the maximum number of lines with transmission errors that the Microsoft Fax Service Provider allows in a received fax before requesting that the sender restart transmission of the page.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax Service to make changes to this entry effective.

### Note

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol specified by the T.30 specification of the International Telecommunications Union for standard fax modems. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
MaxConsecErrorLinesPerPage

# MaxLogFileCount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 10 on Windows XP; 100 on Windows Server 2003 |

**Description**
Specifies the number of log files that the Microsoft Fax Service Provider keeps for both types of fax transmissions: incoming and outgoing. Each log is for a single fax transmission and consists of a record of the conversation between the Microsoft Fax Service Provider and the fax device that sent or received the fax. The Microsoft Fax Service Provider stores logs in the location in the file system specified by the registry entry LoggingFolder

.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax Service to make changes effective.

### Note

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol for standard modems, as described by the T.30 specification of the International Telecommunications Union. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
LoggingFolder

# MaxLogFileSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 102400 bytes (100KB) |

**Description**
Specifies the maximum size, in bytes, of the log files created by the Microsoft Fax Service Provider. If the value of this entry is zero, the Microsoft Fax Service Provider does not limit the size of log files.

Each log is for a single fax transmission and consists of a record of the conversation between the Microsoft Fax Service Provider and the fax device that sent or received the fax. The Microsoft Fax Service Provider stores logs in the location in the file system specified by The registry entry LoggingFolder

.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax Service to make changes effective.

### Note

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supples a fax service provider that supports the fax protocol for standard modems, as described by the T.30 specification of the International Telecommunications Union. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
LoggingFolder

# RetriesBeforeDCN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 3 |

**Description**
Specifies the maximum number of times, before terminating the connection, that the Microsoft Fax Service Provider retransmits a page that a receiving fax device has failed to receive. The receiving fax device communicates with the sending fax device to request that pages be resent.

Before disconnecting, the Microsoft Fax Service Provider tries decreasing the transmission speed when the number of retransmission attempts exceeds the value of registry entry RetriesBeforeDropSpeed.

**Change Method**
You can change this entry only by editing the registry.

### Note

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol for standard modems, as described by the T.30 specification of the International Telecommunications Union. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
RetriesBeforeDropSpeed

# RetriesBeforeDropSpeed

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Device Providers\{2172FD8F-11F6-11D3-90BF-006094EB630B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 1 |

**Description**
Specifies the number of times, before trying a slower transmission speed, that the Microsoft Fax Service Provider retransmits a page that a receiving fax device has failed to receive. The receiving fax device communicates with the sending fax device to request that pages be resent.

**Change Method**
You can change this entry only by editing the registry.

### Note

- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax service uses to interact with one or more fax devices. Microsoft supples a fax service provider that supports the fax protocol for standard modems, as described by the T.30 specification of the International Telecommunications Union. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
RetriesBeforeDCN

# Fax\Devices

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\

**Description**
The Devices subkey stores configuration information about fax devices for the Fax Service. Subkeys in this subkey are named after the unique ID numbers assigned to the fax devices about which they store information.

# device-ID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

The *device-ID* subkeys store configuration information about fax devices for the Fax Service.

The name "*device-ID*" is a placeholder; subkeys of the type described by *device-ID* are named after the unique ID numbers assigned to the fax devices about which they store information.

# Permanent Lineid

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores a unique ID number for the fax device specified by *device-ID* in the registry path.

**Change Method**
This number is used internally by the Fax Service and should not be changed.

# device-ID\{F10A5326-0261-4715-B367-2970427BBD99}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The {F10A5326-0261-4715-B367-2970427BBD99} subkey stores configuration data for the Microsoft Fax Service Provider for the device specified by *device-ID* in the registry path.

### Notes

- The Microsoft Fax Service Provider (FSP) is identified by the globally unique identifier 2172FD8F-11F6-11D3-90BF-006094EB630B. The configuration data that the Fax Service uses for this FSP is identified by the globally unique identifier F10A5326-0261-4715-B367-2970427BBD99.

- A globally unique identifier (GUID) is a 128-bit value that uniquely identifies objects such as OLE servers, interfaces, manager entry-point vectors, and client objects. GUIDs are used in cross-process communication, such as remote procedure calling (RPC) and OLE.

# {F10A5326-0261-4715-B367-2970427BBD99}\CSID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*\{F10A5326-0261-4715-B367-2970427BBD99}

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the Called Subscriber ID (CSID) that is assigned to the fax device specified by *device-ID* in the registry path.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, expand **Services and Applications**, then expand **Fax**, then expand **Devices and Providers**, then expand **Devices**, and then click the name of the fax device for which you want to set the CSID. From the **Action** menu, click **Properties**, and then click the **General** tab. Next, check **Receive faxes** if it is not already checked, and then type the CSID in the **Called Subscriber ID (CSID)** box. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

### Note

- When receiving a fax transmission, the Fax Service sends the CSID to the sender to identify the receiver. A CSID usually consists of the telephone number of the receiving fax device and a name. The CSID is limited to 20 characters.

**See Also**
**Concepts**
{F10A5326-0261-4715-B367-2970427BBD99}\TSID

# {F10A5326-0261-4715-B367-2970427BBD99}\DESCRIPTION

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*\{F10A5326-0261-4715-B367-2970427BBD99}

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores a text description of the fax device specified by *device-ID* in the registry path. This descriptive text appears in the fax device properties sheet on the **General** tab.

**Change Method**
You can view and change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To view or change the value of this entry by using Computer Management, open Computer Management. In the console tree, expand **Services and Applications**, then expand **Fax**, then expand **Devices and Providers**, then expand **Devices**, and then click the name of the fax device for which you want to view or set the description. From the **Action** menu, click **Properties**, and then click the **General** tab. You can view the description and change it by typing in the **Description** box. When you are finished, click **OK** to save your changes, or click **Cancel** to lose them.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

# {F10A5326-0261-4715-B367-2970427BBD99}\Flags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*\{F10A5326-0261-4715-B367-2970427BBD99}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 2 |

**Description**
Specifies whether the fax device, that is specified by *device-ID* in the registry path, is enabled to automatically answer calls and receive faxes and whether the device is enabled to send faxes. This entry stores the values in the following table as a bit mask.

This entry also specifies whether the device is a virtual device and not available through the Telephony Application Programming Interface (TAPI).

| Value | Meaning |
|---|---|
| 1 | The fax device is enabled to automatically answer calls and receive faxes. |
| 2 | The fax device is enabled to send faxes. |
| 4 | The fax device is a virtual device and is not available to TAPI. |

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, expand **Services and Applications**, then expand **Fax**, then expand **Devices and Providers**, then expand **Devices**, and then click the name of the fax device for which you want to set the CSID. From the **Action** menu, click **Properties**, and then click the **General** tab. Next, check **Receive faxes** if it is not already checked, and then select either **Automatic Answer** or **Manual Answer**. Click **OK**.

Only the Fax Configuration Wizard sets the value that specifies a virtual device. You cannot do this by using the user interface.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

**See Also**
**Concepts**
{F10A5326-0261-4715-B367-2970427BBD99}\Rings

# {F10A5326-0261-4715-B367-2970427BBD99}\GUID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*\{F10A5326-0261-4715-B367-2970427BBD99}

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores the globally unique identifier (GUID) that identifies the fax service provider associated with the fax device specified by *device-ID* in the registry path.

A globally unique identifier (GUID) is a 128-bit value that uniquely identifies objects such as OLE servers, interfaces, manager entry-point vectors, and client objects. GUIDs are used in cross-process communication, such as remote procedure calling (RPC) and OLE.

**Change Method**
This entry is configured automatically when a fax device is installed in the system and should not be changed.

### Notes

The Microsoft Fax Service Provider (FSP) is identified by the GUID 2172FD8F-11F6-11D3-90BF-006094EB630B. The configuration data that the Fax service uses for this FSP is identified by the globally unique identifier F10A5326-0261-4715-B367-2970427BBD99.

A fax service provider is a dynamic-link library (DLL) that provide an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol specified by the T.30 specification of the International Telecommunications Union (ITU) for standard fax modems. Third parties can also supply fax service providers to support other fax-capable devices.

# LastDetected

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*\{F10A5326-0261-4715-B367-2970427BBD99}

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | | |

**Description**
Stores the last time that the device specified by *device-ID* in the registry path was detected and operational.

**Change Method**
Do not change this entry.

### Caution

- This entry is used internally by the Microsoft Fax Service Provider. Do not change it.

# {F10A5326-0261-4715-B367-2970427BBD99}\Rings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*\{F10A5326-0261-4715-B367-2970427BBD99}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0-99 | 2 |

**Description**
Specifies the number times that an incoming call is allowed to ring before the fax device specifed by *device-ID* in the registry path, answers. However, the fax device only answers calls if it is configured to do so by the registry entry {F10A5326-0261-4715-B367-2970427BBD99}\Flags.

**Change Method**
You can view and change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, expand **Services and Applications**, then expand **Fax**, then expand **Devices and Providers**, then expand **Devices**, and then click the name of the fax device for which you want to set the CSID. From the **Action** menu, click **Properties**, and then click the **General** tab. Next, check **Receive faxes** if it is not already checked, and select **Automatic Answer** if it is not already selected. Change the number in the **Rings before answering** box by selecting the number and typing a new number, or by clicking the arrow buttons next to the number. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

**See Also**
**Concepts**
{F10A5326-0261-4715-B367-2970427BBD99}\Flags

# TAPI Permanent Lineid

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*\{F10A5326-0261-4715-B367-2970427BBD99}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores a unique ID number that is used internally by the Telephony Application Programming Interface (TAPI) to identify the fax device specified by *device-ID* in the registry path.

**Change Method**
This entry is maintained automatically by the Fax service and should not be changed.

### Caution

- Do not delete this entry from the registry or change its value. Doing so can cause the Fax Service to fail.

# {F10A5326-0261-4715-B367-2970427BBD99}\TSID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices*device-ID*\{F10A5326-0261-4715-B367-2970427BBD99}

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the Transmitting Subscriber ID (TSID) that is assigned to the fax device specified by *device-ID* in the registry path. When sending a fax, the Fax Service sends the TSID to the receiver to identify the sender. A TSID usually consists of the telephone number of the sending fax device and a name. The TSID is limited to 20 characters.

The Fax Service sends the TSID, depending on the value of registry entry **UseDeviceTsid**. The receiving fax device might use the TSID in a variety of ways; fax machines often display the TSID when receiving a fax and also might print it on the top of each received page. The Fax Service can also include the TSID in a banner, which it prepends to each transmitted page, when the registry entry Branding is set to 1.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003.

You can view and change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, expand **Services and Applications**, then expand **Fax**, then expand **Devices and Providers**, then expand **Devices**, and then click the name of the fax device for which you want to set the CSID. From the **Action** menu, click **Properties**, and then click the **General** tab. Next, check **Send faxes** if it is not already checked, and then type the CSID in the **Transmitting Subscriber ID (TSID)** box. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

**See Also**
**Concepts**
Original Setup Data\TSID
{F10A5326-0261-4715-B367-2970427BBD99}\CSID

# UnassociatedExtensionData

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices

**Description**
The UnassociatedExtensionData subkey stores information for the Fax Service about third-party extensions to the Fax Service.

# Devices Cache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The Devices Cache subkey stores data about fax devices that have been registered with the Fax Service but are currently removed from the system. If a removed device is later restored to the system, then the Fax Service automatically uses the configuration stored in this subkey.

The subkeys in the Devices Cache subkey have unique names that identify specific devices.

**See Also**
**Concepts**
TAPI Permanent Line ID

# TAPI Permanent Line ID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices Cache\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The *TAPI Permanent Line ID* subkey stores data about a specific fax device that was registered with the Fax Service but is currently removed from the system. If this device is later restored to the system, the Fax Service automatically uses the configuration stored in this subkey.

### Note

- The *TAPI Permanent Line ID* subkey is a placeholder in the Devices Cache subkey. It represents the subkeys, which have unique names that identify specific devices. "*TAPI Permanent Line ID*" does not actually appear in the registry.

# ManualAnswer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices Cache*TAPI Permanent Line ID*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | |

**Description**
Specifies whether the fax device specified by *TAPI Permanent Line ID* in the registry path was configured for manual answering when the Microsoft Fax Service Provider copied the configuration settings of the device to the Devices Cache subkey. The Devices Cache subkey stores data about fax devices that have been registered with the Fax service but are currently removed from the system. If a removed device is later restored to the system, the Fax service automatically uses the configuration stored in this subkey.

### Note

- A fax service provider is a dynamic-link library (DLL) that provide an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol specified by the T.30 specification of the International Telecommunications Union for standard fax modems. Third parties can also supply fax service providers to support other fax-capable devices.

### Caution

- This entry is maintained automatically by the Microsoft Fax Service Provider. Do not change it.

# TAPI Permanent Line ID\{F10A5326-0261-4715-B367-2970427BBD99}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices Cache*TAPI Permanent Line ID*

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The {F10A5326-0261-4715-B367-2970427BBD99} subkey in HKLM\SOFTWARE\Microsoft\Fax\Devices Cache*TAPI Permanent Line ID* stores a copy of the contents of the subkey device-ID\{F10A5326-0261-4715-B367-2970427BBD99} in HKLM\SOFTWARE\Microsoft\Fax\Devices*Device ID*.

### Notes

- The Microsoft Fax Service Provider (FSP) is identified by the globally unique identifier 2172FD8F-11F6-11D3-90BF-006094EB630B. The configuration data that the Fax Service uses for the Microsoft FSP is identified by the globally unique identifier F10A5326-0261-4715-B367-2970427BBD99.

- A globally unique identifier (GUID) is a 128-bit value that uniquely identifies objects such as OLE servers, interfaces, manager entry-point vectors, and client objects. GUIDs are used in cross-process communication, such as remote procedure calling (RPC) and OLE.

- The Telephone Application Programming Interface (TAPI) is part of the Microsoft telephony architecture. In the registry path of this registry entry, *TAPI Permanent Line ID* is an ID number assigned to a modem device by the telephony subsystem.

**See Also**
**Concepts**
device-ID\{F10A5326-0261-4715-B367-2970427BBD99}
Devices Cache

# TAPI Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Devices Cache*TAPI Permanent Line ID*

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The TAPI Data subkey stores a copy of the contents of the *device-ID* subkey in HKLM\SOFTWARE\Microsoft\fax\TAPIDevices.

**See Also**
**Concepts**
TAPIDevices\DeviceID
Devices Cache

# Inbox

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The Inbox subkey stores management information for incoming faxes stored by the Fax Service, including folder location, maximum size, and expiration times.

# AgeLimit Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Inbox

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 90 |

**Description**
Specifies the time, in days, that the Fax Service allows a received fax to remain in the inbox folder. (The location of the inbox folder is specified by the registry entry Folder Entry.) If a received fax is in the inbox folder longer than this time, the Fax Service automatically deletes it. If the value of this entry is zero, there is no time limit, and the Fax Service does not delete faxes from the inbox.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003

You can view and change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, expand **Services and Applications**, then expand **Fax**, then expand **Devices and Providers**, then expand **Devices**, and then click the name of the fax device for which you want to set the CSID. From the **Action** menu, click **Properties**, and then click the **General** tab. Next, check **Receive faxes** if it is not already checked, and then type the CSID in the **Called Subscriber ID (CSID)** box. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

**See Also**
**Concepts**
SentItems\AgeLimit
Folder Entry

# Folder Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Inbox

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | C:\Documents and Settings\All Users\Application Data\Microsoft\Windows NT\MSFax\Inbox |

**Description**
Specifies the location of the inbox folder of the Fax Service. The Fax Service stores a received fax in the inbox folder when the fax has been successfully received and routed.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Inbox** tab. Under **Archive folder**, check **Archive all incoming faxes to this folder** if it is not already checked. Next, type the path for the folder in the **Archive all incoming faxes to this folder** box, or click **Browse...** and browse to the folder. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

### Note

- Before the Fax Service puts an incoming fax in the inbox folder, the Fax Service stores it in the incoming folder. The incoming folder is a repository for all incoming faxes from devices local and remote. The Fax Service only moves a fax from the incoming folder to the inbox folder when it has been received and routed in accordance with the policy specified by the fax administrator. If a fax reception is terminated while in progress, but a portion of the fax is successfully received, the fax is assigned the status of "partially received," and is moved to the inbox.

**See Also**
**Concepts**
SentItems\Folder

# Inbox\HighWatermark

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Inbox

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 120 |

**Description**
Specifies a size threshhold, in megabytes, for the inbox folder of the Fax Service. When the combined size of the contents of the inbox folder reach the size specified by this entry, the Fax Service enters a warning event into the Windows event log if the registry entry Inbox\SizeQuotaWarn specifies that the service do so. The inbox folder is the folder in which the Fax Service stores faxes that have been received and routed successfully.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Inbox** tab. Under **Archive rules**, check or clear **Generate warning in the Event Log,** if it is not already checked. Next, change the number in the **Size quota high watermark** box by selecting the number and typing a new number, or by clicking the arrow buttons next to the number. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

## Notes

- This entry specifies the higher of two threshholds for the inbox folder. The other threshhold is specified by the registry entry Inbox\LowWatermark.
- Checking or clearing **Generate warning in the Event Log** changes the value of the registry entry Inbox\SizeQuotaWarn.
- Before the Fax Service puts an incoming fax in the inbox folder, the Fax Service stores it in the incoming folder. The incoming folder is a repository for all incoming faxes from devices local and remote. The Fax Service only moves a fax from the incoming folder to the inbox folder when it has been received and routed in accordance with the policy specified by the fax administrator. If a fax reception is terminated while in progress, but a portion of the fax is successfully received, the fax is assigned the status of "partially received," and is moved to the inbox.

**See Also**
**Concepts**
Inbox\SizeQuotaWarn
Inbox\LowWatermark

# Inbox\LowWatermark

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Inbox

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 115 |

**Description**
Specifies a size threshhold, in megabytes, for the inbox folder of the Fax Service. When the combined size of the contents of the inbox folder reach the size specified by this entry, the Fax Service enters a warning event into the Windows event log if the registry entry Inbox\SizeQuotaWarn specifies that the service do so. The inbox folder is the folder in which the Fax Service stores faxes that have been received and routed successfully.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Inbox** tab. Under **Archive rules**, check or clear **Generate warning in the Event Log**, if it is not already checked. Next, change the number in the **Size quota low watermark** box by selecting the number and typing a new number, or by clicking the arrow buttons next to the number. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

## Notes

- This entry specifies the lower of two threshholds for the inbox folder. The other threshhold is specified by the registry entry Inbox\HighWatermark.

- Checking or clearing **Generate warning in the Event Log** changes the value of the registry entry Inbox\SizeQuotaWarn.

- Before the Fax Service puts an incoming fax in the inbox folder, the Fax Service stores it in the incoming folder. The incoming folder is a repository for all incoming faxes from devices local and remote. The Fax Service only moves a fax from the incoming folder to the inbox folder when it has been received and routed in accordance with the policy specified by the fax administrator. If a fax reception is terminated while in progress, but a portion of the fax is successfully received, the fax is assigned the status of "partially received," and is moved to the inbox.

**See Also**
**Concepts**
Inbox\SizeQuotaWarn
Inbox\HighWatermark

# Inbox\SizeQuotaWarn

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Inbox

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the Fax Service enters a warning event into the Windows event log when the contents of the Fax Service inbox folder reach a combined size equal to or greater than one of two size threshholds (specified by Inbox\LowWatermark and Inbox\HighWatermark). The inbox folder is the folder in which the Fax Service stores faxes that have been received and routed successfully.

| Value | Meaning |
|---|---|
| 0 | The Fax Service does not create warning events in the Windows event log regarding the size of the inbox. |
| 1 | The Fax Service creates a warning event in the Windows event log when the contents of the inbox reach a combined size equal to or greater than Inbox\LowWatermark or Inbox\HighWatermark). |

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Inbox** tab. Under **Archive rules**, check or clear **Generate warning in the Event Log**. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

### Note

- Before the Fax Service puts an incoming fax in the inbox folder, the Fax Service stores it in the incoming folder. The incoming folder is a repository for all incoming faxes from devices local and remote. The Fax Service only moves a fax from the incoming folder to the inbox folder when it has been received and routed in accordance with the policy specified by the fax administrator. If a fax reception is terminated while in progress, but a portion of the fax is successfully received, the fax is assigned the status of "partially received" and is moved to the inbox.

**See Also**
**Concepts**
Inbox\LowWatermark
Inbox\HighWatermark

# Inbox\Use

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Inbox

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether or not the Fax Service archives incoming faxes.

| Value | Meaning |
|---|---|
| 0 | The Fax Service does not archive incoming faxes. |
| 1 | The Fax Service archives incoming faxes. |

Archived incoming faxes are stored in the location specified by the registry entry Folder Entry (in the same subkey as this entry).

**Change Method**
You can change the value of this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, click the **Inbox** tab, and then clear or check **Archive all incoming faxes to this folder**. Click **OK**.

**Activation Method**
Changes that you make by using **Fax Service Manager** or **Computer Management** are effective immediately.

**See Also**
**Concepts**
Folder Entry
SentItems\Use

# Logging Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

The Logging subkey stores the configuration for logs that the Fax Service keeps.

# event-type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Logging

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The *event-type* subkey stores configuration data for event logs that are kept by the Fax Service. The four types of *event-type* subkeys listed in the following table categorize the type of event that the Fax Service logs.

| Value | Meaning |
|---|---|
| 1 | Initialization and termination of fax transmissions |
| 2 | Outgoing fax transmissions |
| 3 | Incoming fax transmissions |
| 4 | General fax-related activity |

**Note**

- The *event-type* subkey is a placeholder in the Logging Subkey subkey. It represents the subkeys, which have a specific number as their name; "*event-type*" does not actually appear in the registry.

**See Also**
**Concepts**
ActivityLogging

# Level

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Logging*event-type*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 - 3 | 2 |

**Description**
Specifies a level for entries that the Fax Service makes in the Windows event log. The Fax Service logs the greatest number of events when this entry is equal to 3, and the lowest number when this entry is equal to 1. When this entry is equal to zero, no events are logged.

The Fax Service categorizes the entries that it makes in the Windows event log into four types. This entry specifies the logging level for the category that is specified by the number *event-type* in the registry path.

**Change Method**
This entry and its value are created when the Microsoft Fax Service Provider is configured. You can change the value later by using Fax Service Manager, which is part of the Fax service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Event Reports** tab. For each type of event—**General**, **Incoming**, **Outgoing**, or **Initialization or shutdown**—for which you want to change the logging level, move the slider to the left to reduce the logging level, or to the right to increase the logging level. Click **OK**.

**Activation Method**
Changes that you make by using **Fax Service Manager** or **Computer Management** are effective immediately.

### Notes

- For more information about the event categories that the Fax Service uses, see the event-type subkey.
- This entry pertains to entries that the Fax Service creates about itself in the Windows event log. The Fax Service also maintains activity logs that record information about the faxes that it has sent and received. Activity logs are configured by the registry entries in the ActivityLogging subkey.

**See Also**
**Concepts**
event-type

# Outbound Routing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The Outbound Routing subkey stores the configuration for the routing of outbound faxes by the Fax Service.

# Outbound Routing\Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Outbound Routing

**Description**
The *Groups* subkey stores information for the Fax Service about routing for outgoing faxes. This subkey contains other subkeys that are named for groups of fax-capable devices that are registered with the Fax Service.

**See Also**
**Concepts**
group-name

# group-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Outbound Routing\Groups

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

The *group-name* subkeys have names that match the names of routing groups for outgoing faxes in the Fax Service. Routing groups are groups of fax-capable devices as defined by an administrator, and together with routing rules, an administrator uses them to control the routing of outgoing faxes.

# group-name\Devices

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Outbound Routing\Groups*group-name*

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | | |

**Description**
Stores a list of the devices in the outbound routing group of the Fax Service. Each group of four bytes in the binary value of this entry can identify a single device, and does so when a device is present in the system.

**Change Method**
This entry is automatically maintained by the Fax Service, and you should not change it by editing the registry. To manage devices that are used by the Fax Service, make changes by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003 (but not on Windows XP) To start Fax Service Manager, from the **Start** menu, point to **Accessories**, point to **Communications**, point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

# Rules

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Outbound Routing\

| Data type | Range | Default value |
|---|---|---|
| NULL | | |

**Description**
The Rules subkey stores routing rules for the Fax Service. Each subkey in the Rules subkey stores the configuration for a single routing rule for outgoing faxes. The names of these subkeys have the format *country code*:*area code*.

**See Also**
**Concepts**
CountryCode:AreaCode

# CountryCode:AreaCode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Outbound Routing\Rules

| Data type | Range | Default value |
|---|---|---|
| NULL | | |

**Description**
The *CountryCode:AreaCode* subkey stores the configuration of a routing rule for the Fax Service. Each subkey of this type stores the configuration for a single routing rule for outgoing faxes.

In the names of these types of subkeys, *CountryCode* is the telephone dialling code for a country or region and *AreaCode* is the telephone dialing code for a city or area within the specified country or region.

**Change Method**
You can add routing rules—and, thus, subkeys of this type—by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To add a routing rule, in Fax Service Manager, expand **Outgoing Routing** and click **Rules**. From the **Action** menu, point to **New**, and then click **Rule...**. In **Add New Rule** dialog box, type a number in the **Country/region code** box, or click **Select...** to select a code from a list. Next, under **Area code**, select **Specific** and type an area code in the **Specific** box, or select **All areas**. Under **Target device**, select **Device** and then, in the **Device** list, click the name of a device; or select **Routing group** and then, in the **Routing group** list, click the name of a routing group. Click **OK**.

To delete a routing rule, in Fax Service Manager, expand **Outgoing Routing** and click **Rules**. In the results pane, click the rule that you want to delete, and on the **Action** menu, click **Delete**.

**Activation Method**
Changes that you make by using **Fax Service Manager** or **Computer Management** are effective immediately.

# CountryCode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Outbound Routing\Rules*CountryCode:AreaCode*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 0 |

**Description**
Specifies the area code to which the Fax Service applies the routing rule that is named in the registry path by *CountryCode:AreaCode*. If the telephone number for an outgoing fax matches the area code and country code of a routing rule, the Fax Service sends the fax by using a fax device that the routing rule specifies.

If this entry is equal to zero, the routing rule applies to faxes with telephone numbers for any country or region. The Fax Service automatically creates such a routing rule when the service is installed, and routing rules added subsequently by an administrator have precedence when the Fax Service routes an outgoing fax.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change this entry, in Fax Service Manager, expand **Outgoing Routing** and click **Rules**. In the results pane, click the rule that you want to delete, and on the **Action** menu, click **Properties**. Under **Dialed number**, type a country or region code into the **Country/region code** box or click **Select** to select the code from a list. If you select the code from the list, click **OK** in the **Select Country or Region** dialog box. Click **OK** in the properties sheet for the routing rule.

**Activation Method**
Changes that you make by using **Fax Service Manager** or **Computer Management** are effective immediately.

**See Also**
**Concepts**
CountryCode:AreaCode
UseGroup

# DeviceID Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Outbound Routing\Rules*CountryCode:AreaCode*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies the fax device by which the Fax Service sends a fax for which the routing rule, named in the registry path by *CountryCode:AreaCode*, applies. If the telephone number for an outgoing fax matches the area code and country code of a routing rule, the Fax Service sends the fax according to the matching routing rule.

Routing rules can specify sending a fax by either a specific device or a group of devices. This registry entry only appears in the registry when the associated routing rule is configured to route faxes to a single fax device. When a routing rule specifies a routing group, this entry is replaced by the registry entry GroupName.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change this entry, in Fax Service Manager, expand **Outgoing Routing** and click **Rules**. In the results pane, click the rule that you want to delete, and on the **Action** menu, click **Properties**. Under **Target device**, select **Device** to route faxes to a single device or select **Routing group** to route faxes to a group of devices. To specify the device, in the **Device** list, click the name of a device. To specify a group, in the **Routing group** list, click the name of a routing group. Click **OK** in the properties sheet for the routing rule.

**Activation Method**
Changes that you make by using **Fax Service Manager** or **Computer Management** are effective immediately.

**Tip**

- You must create a routing group before you can specify it in a routing rule.

**See Also**
**Concepts**
UseGroup
GroupName
CountryCode:AreaCode

# GroupName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Outbound Routing\Rules*CountryCode:AreaCode*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | <All devices> |

**Description**
Specifies the routing group by which the Fax Service sends a fax for which the routing rule, named in the registry path by *CountryCode:AreaCode*, applies. If the telephone number for an outgoing fax matches the area code and country code of a routing rule, the Fax Service sends the fax according to the matching routing rule.

Routing rules can specify sending a fax by either a specific device or a group of devices. This registry entry only appears in the registry when the associated routing rule is configured to route faxes to a routing group. When a routing rule specifies a single fax device, this entry is replaced by the registry entry GroupName.

When this entry is "<All devices>", the default value, the Fax Service sends a fax that matches the associated rule by using any fax device installed on the system. If the routing rule specifies a particular routing group, this entry contains the name of the routing group as it appears in Fax Service Manager.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change this entry, in Fax Service Manager, expand **Outgoing Routing** and click **Rules**. In the results pane, click the rule that you want to delete, and on the **Action** menu, click **Properties**. Under **Target device**, select **Device** to route faxes to a single device or select **Routing group** to route faxes to a group of devices. To specify the device, in the **Device** list, click the name of a device. To specify a group, in the **Routing group** list, click the name of a routing group. Click **OK** in the properties sheet for the routing rule.

**Activation Method**
Changes that you make by using **Fax Service Manager** or **Computer Management** are effective immediately.

**Tip**

- You must create a routing group before you can specify it in a routing rule.

**See Also**
**Concepts**
UseGroup
DeviceID Entry
CountryCode:AreaCode

# UseGroup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Outbound Routing\Rules*CountryCode:AreaCode*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 1 |

**Description**
Specifies how the Fax Service selects a fax device when it sends a fax for which the routing rule, named in the registry path by *CountryCode:AreaCode*, applies. A routing rule can specify either a single specific fax device, or it can specify a routing group, a group of fax devices, from which the Fax Service selects a device by which to transmit the outgoing fax.

| Value | Meaning |
|---|---|
| 0 | The routing rule specifies a single fax device. |
| 1 | The routing rule specifies a routing group. |

When an administrator creates a routing rule, the default selection method for devices is by routing group.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change this entry, in Fax Service Manager, expand **Outgoing Routing** and click **Rules**. In the results pane, click the rule that you want to delete, and on the **Action** menu, click **Properties**. Under **Target device**, select **Routing group**.

To specify a routing group (the name of which is stored in the registry entry GroupName, in the **Routing group** list, click the name of a routing group.

To finish, click **OK**.

**Activation Method**
Changes that you make by using **Fax Service Manager** or **Computer Management** are effective immediately.

### Notes

- If the telephone number for an outgoing fax matches the area code and country code of a routing rule, the Fax Service sends the fax according to the matching routing rule.
- The default routing group, called "<All devices>", is a group that the Fax Service creates automatically that includes all fax devices installed on the system.

**See Also**
**Concepts**
GroupName
DeviceID Entry
CountryCode:AreaCode

# Receipts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

The Receipts subkey stores the configuration information for receipts that the Fax Service generates for delivered faxes.

**Note**

- You can configure the Fax service to create delivery receipts either as alerts on the computer that runs the Fax Service in your enterprise, or as e-mail that it sends through an SMTP server.

# From Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Receipts

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the e-mail address from which the Fax Service sends messages to senders of faxes to report success or failure in the delivery of a fax.

The Fax Service uses Simple Mail Transfer Protocol (SMTP) to send the messages. The Fax Service uses the string stored by this entry as the "From" address when it connects to the server specified by the registry entry Server Entry, and does so on the Internet Protocol (IP) port specified by the registry entry Port.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry, in Fax Service Manager, click **Fax**. From the **Action** menu, click **Properties**, and then click the **Receipts** tab. Under **SMTP e-mail**, check **Enable SMTP e-mail receipts delivery**, if it is not already checked. Next, in the **From e-mail address** box, type the e-mail address from which you want the Fax service to send receipts. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

### Note

- The delivery notification feature of the Fax Service provides users with information about the success or failure of their sent faxes. If delivery notification is enabled by an administrator, a user sending a fax with the Send Fax Wizard can choose to receive a delivery receipt for the sent fax, in accordance with enabled delivery notification methods. Having a delivery receipt means that a user does not need to track the progress of an outgoing fax from Fax Console.

**See Also**
**Concepts**
Server Entry
Port

# Password

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Receipts

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | | |

**Description**
Specifies the password that the Fax Service uses to authenticate itself to the Simple Mail Transfer Protocol (SMTP) service when it sends notification of delivery by e-mail.

Authentication also requires a user name, specified by the registry entry User. The type of authentication is specified by the registry entry SMTPAuth.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service onWindows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry in Fax Service Manager, click **Fax**. From the **Action** menu, click **Properties**, and then click the **Receipts** tab. Under **SMTP e-mail**, check **Enable SMTP e-mail receipts delivery**, if it is not already checked. Next, click **Authentication**. In the **Authentication** window, select **Basic authentication** or **Integrated Windows Authentication**. Then click **Credentials**. Enter the user name and password when you are prompted to do so, and then click **OK**.

When you are done configuring authentication, click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

## Notes

- The delivery notification feature of the Fax Service provides users with information about the success or failure of their sent faxes. If delivery notification is enabled by an administrator, a user sending a fax with the Send Fax Wizard can choose to receive a delivery receipt for the sent fax, in accordance with the enabled delivery notification methods. Having a delivery receipt means that a user does not need to track the progress of an outgoing fax from Fax Console.

- Previously known as Windows NT Challenge/Response (NT/CR) or NT LAN Manager (NTLM), Integrated Windows Authentication is a secure authentication method that does not transmit user names or passwords. Instead, it relies on a cryptographic exchange with the server. Integrated Windows Authentication can use either the Kerberos V5 authentication protocol or its own challenge/response protocol.

- **Basic authentication** sends a user name and password over the network to the server as unencrypted plain text.

**See Also**
**Concepts**
User
SMTPAuth
Server Entry

# Port

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Receipts

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 25 |

**Description**
Specifies the number of the port to which the Fax Service connects when it sends a message to report success or failure in the delivery of a fax.

The Fax Service sends the message by e-mail by using Simple Mail Transfer Protocol (SMTP). The Fax Service connects to the Internet Protocol (IP) port specified by this entry on the server specified by the registry entry Server Entry.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Standard Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry, in Fax Service Manager, click **Fax**. From the **Action** menu, click **Properties**, and then click the **Receipts** tab. Under **SMTP e-mail**, check **Enable SMTP e-mail receipts delivery** if it is not already checked. Next, delete the number in the **Port** box and type the correct number for the SMTP port for the SMTP service running on the mail server. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

## Notes

- The delivery notification feature of the Fax Service provides users with information about the success or failure of their sent faxes. If delivery notification is enabled by an administrator, a user sending a fax with Send Fax wizard can choose to receive a delivery receipt for the fax being sent, in accordance with the delivery notification methods enabled. Having a delivery receipt means that a user does not need to track the progress of an outgoing fax from Fax Console.

- Port 25 for SMTP services on the Internet is one of the "well-known" port number assignments made by the Internet Assigned Numbers Authority.

**See Also**
**Concepts**
Server Entry
From Entry

# Server Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Receipts

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the name of the server to which the Fax Service connects when it sends a message to report success or failure in the delivery of a fax.

The Fax Service sends the message by e-mail, using Simple Mail Transfer Protocol (SMTP). The Fax Service connects to the server specified by this entry and does so on the Internet Protocol (IP) port specified by the registry entry Port.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry, in Fax Service Manager, click **Fax**. From the **Action** menu, click **Properties**, and then click the **Receipts** tab. Under **SMTP e-mail**, check **Enable SMTP e-mail receipts delivery**, if it is not already checked. Next, in the **Server address** box, type the name of the server to which you want the Fax Service to connect. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

### Note

- The delivery notification feature of the Fax Service provides users with information about the success or failure of their sent faxes. If delivery notification is enabled by an administrator, a user sending a fax with Send Fax wizard can choose to receive a delivery receipt for the fax being sent, in accordance with the delivery notification methods enabled. Having a delivery receipt means that a user does not need to track the progress of an outgoing fax from Fax Console.

**See Also**
**Concepts**
Port
From Entry

# SMTPAuth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Receipts

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 0 |

**Description**
Specifies the type of authentication that the Fax Service uses when it connects to the SMTP service on a server to send fax delivery notification. To send the message in e-mail, the Fax Service transmits the message by connecting to a server (specified by the registry entry Server Entry) that has SMTP service.

| Value | Meaning |
|---|---|
| 0 | **Anonymous access**. The Fax service does not authenticate itself, because the SMTP service allows anonymous connections. |
| 1 | **Basic authentication**.The Fax Service authenticates itself by sending a user name and password in unencrypted plain text over the network. |
| 2 | **Integrated Windows Authentication**.The Fax Service authenticates itself by using Integrated Windows Authentication. |

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry, in Fax Service Manager, click **Fax**. From the **Action** menu, click **Properties**, and then click the **Receipts** tab. Under **SMTP e-mail**, check **Enable SMTP e-mail receipts delivery** if it is not already checked. Next, click **Authentication**. In the **Authentication** window, select **Anonymous access**, **Basic authentication**, or **Integrated Windows Authentication**.

If you select **Basic authentication** or **Integrated Windows Authentication**, you must also supply a user name and password for the Fax Service to use for authentication. To do so, click **Credentials...**, and enter a user name and password when prompted. Click **OK**.

When you are done configuring authentication, click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

### Notes

- The delivery notification feature of the Fax Service provides users with information about the success or failure of their sent faxes. If delivery notification is enabled by an administrator, a user sending a fax with Send Fax wizard can choose to receive a delivery receipt for the fax being sent, in accordance with the delivery notification methods enabled. Having a delivery receipt means that a user does not need to track the progress of an outgoing fax from Fax Console.

- Previously known as Windows NT Challenge/Response (NT/CR) or NTLM, Integrated Windows Authentication is a secure authentication method that does not transmit user names or passwords. Instead, it relies on a cryptographic exchange with the server. Integrated Windows Authentication can use either the Kerberos V5 authentication protocol or its own challenge/response protocol.

**See Also**

**Concepts**

User

Password

Server Entry

# Receipts\Type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Receipts

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 4 |

**Description**
Specifies the types of receipts that the Fax Service delivers to the senders of faxes. The value of this entry is a bit-wise combination of the values shown in the following table.

| Value | Meaning |
|---|---|
| 1 | The Fax Service notifies the sender in e-mail by using Simple Mail Transport Protocol (SMTP). |
| 2 | This value is not used. |
| 4 | The Fax Service notifies the sender by using **net send** to create a message window on the sender's computer. |

An administrator can configure the Fax Service to use either, both, or none of these notification methods.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry, in Fax Service Manager, click **Fax**. From the **Action** menu, click **Properties**, and then click the **Receipts** tab. To configure the Fax Service to create pop-up messages in the sender's login session, under **Message Box**, check **Enable message boxes as receipts**. Then, clear the check box to prevent the Fax Service from sending such messages. To configure the Fax Service to send delivery receipts by e-mail, under **SMTP e-mail**, check **Enable SMTP e-mail receipts delivery**, and then clear the check box to prevent the Fax Service from sending such e-mail.

When you finish setting the receipts configuration, click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

### Notes

- Note that, if you check **Enable SMTP e-mail receipts delivery**, you must also type an e-mail address in the **From e-mail address** box, the name of a server in the **Server address** box, and a port number in the **Port** box.

- The delivery notification feature of the Fax Service provides users with information about the success or failure of their sent faxes. If delivery notification is enabled by an administrator, a user sending a fax with Send Fax wizard can choose to receive a delivery receipt for the fax being sent, in accordance with the delivery notification methods enabled. Having a delivery receipt means that a user does not need to track the progress of an outgoing fax from Fax Console.

**See Also**
**Concepts**
From Entry
Server Entry
Port

# UseForMsRoute

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Receipts

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 0 |

**Description**
Specifies whether the e-mail configuration that the Fax Service uses for sending delivery receipts is also used as the e-mail configuration for routing incoming faxes. The Fax Service applies this method of routing incoming faxes, seen in Fax Service Manager as **Route through e-mail**, according to routing rules established by the administrator.

To route an incoming fax through e-mail, the Fax Service connects to the server specified by the registry entry Server Entry on the port specified by the registry entry Port, and sends the fax in an e-mail message from the address specified in the registry entry From Entry.

| Value | Meaning |
|---|---|
| 0 | The Fax Service does not use the e-mail configuration for delivery receipts as the configuration for the **Route through e-mail** routing method. |
| 1 | The Fax Service routes incoming faxes, according to routing rules established by the administrator, through e-mail, by using the e-mail configuration for sending delivery receipts. |

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry, in Fax Service Manager, click **Fax**. From the **Action** menu, click **Properties**, and then click the **Receipts** tab. Under **SMTP e-mail**, check or clear **Use these SMTP settings for the Route Through E-Mail incoming routing method**. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

## Notes

- If you check **Use these SMTP settings for the Route Through E-Mail incoming routing method**, you must also type an e-mail address in the **From e-mail address** box, the name of a server in the **Server address** box, and a port number in the **Port** box.

- The delivery notification feature of the Fax Service provides users with information about the success or failure of their sent faxes. If delivery notification is enabled by an administrator, a user sending a fax with Send Fax wizard can choose to receive a delivery receipt for the fax being sent, in accordance with the enabled delivery notification methods. Having a delivery receipt means that a user does not need to track the progress of an outgoing fax from Fax Console.

**See Also**
**Concepts**
From Entry
Server Entry
Port

# User

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Receipts

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the user name that the Fax Service uses to authenticate itself to the Simple Mail Transfer Protocol (SMTP) service when it sends notification of delivery by e-mail.

Authentication also requires a password, specified by the registry entry Password. The type of authentication is specified by the registry entry SMTPAuth.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry, in Fax Service Manager, click **Fax**. From the **Action** menu, click **Properties**, and then click the **Receipts** tab. Under **SMTP e-mail**, check **Enable SMTP e-mail receipts delivery** if it is not already checked. Next, click **Authentication**. In the **Authentication** window, select **Basic authentication** or **Integrated Windows Authentication**. Then click **Credentials**. Enter the user name and password when you are prompted to do so, and then click **OK**.

When you are done configuring authentication, click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

## Notes

- The delivery notification feature of the Fax Service provides users with information about the success or failure of their sent faxes. If delivery notification is enabled by an administrator, a user sending a fax with Send Fax wizard can choose to receive a delivery receipt for the fax being sent, in accordance with the delivery notification methods enabled. Having a delivery receipt means that a user does not need to track the progress of an outgoing fax from Fax Console.

- Previously known as Windows NT Challenge/Response (NT/CR) or NT LAN Manager (NTLM), Integrated Windows Authentication is a secure authentication method that does not transmit user names or passwords. Instead, it relies on a cryptographic exchange with the server. Integrated Windows Authentication can use either the Kerberos V5 authentication protocol or its own challenge/response protocol.

- **Basic authentication** sends a user name and password over the network to the server as unencrypted plain text.

**See Also**
**Concepts**
Password
SMTPAuth
Server Entry

# Routing Extensions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

The Routing Extensions subkey stores configuration information about routing extensions for the Fax Service.

# ExtensionName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The *ExtensionName* subkeys store configuration information for routing extensions to the Fax Service. A routing extension provides the Fax Service with routing functions for incoming faxes.

The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system, and its configuration is stored in a subkey with the same name. (The name "ExtensionName" does not actually appear in the registry; *ExtensionName* is a placeholder for a type of registry subkey.) Other routing extensions might also be installed on a system.

# ExtensionName\FriendlyName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | Microsoft Routing Extension |

**Description**
Specifies the name of the routing extension named in the registry path by *ExtensionName*. Fax Service Manager displays the string stored by this entry in **Incoming Routing** to identify for each routing method the extension that provides the method.

### Notes

- A routing extension provides the Fax Service with routing functions for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

# ExtensionName\ImageName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | | %SystemRoot%\system32\faxroute.dll |

**Description**
Specifies the file name and path of the dynamic-link library (DLL) that provides the routing extension named in the registry path by *ExtensionName*. When the routing extension is the Microsoft Routing Extension, this entry has the default value **%SystemRoot%\system32\faxroute.dll**. For other routing extensions, this entry will have a different value.

**Change Method**
This entry is configured automatically when the routing extension is registered with the Fax Service. The extension is registered automatically during installation of the Fax component on the operating system. Do not change this entry.

### Notes

- A routing extension provides the Fax Service with routing functions for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

# Routing Methods

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*

**Description**
The Routing Methods subkey stores configuration information for the routing methods available to the Fax Service. The routing extension that is identified in the registry path by *ExtensionName* provides the routing methods.

The Microsoft Routing Extension provides the Fax Service with three routing methods. These are identified by the names Email, Folder, and Printing, which appear by name as subkeys in this subkey.

### Notes

- A routing extension provides the Fax Service with routing functions for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

**See Also**
**Concepts**
MethodName
ExtensionName

# MethodName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*\Routing Methods

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The *MethodName* subkey stores configuration information for the routing method for incoming faxes named by *MethodName*. ("MethodName" does not actually appear in the registry; *MethodName* is a placeholder for a type of subkey.) The routing method is provided by the routing extension that is identified in the registry path by *ExtensionName*.

The Microsoft Routing Extension provides the Fax Service with three routing methods. These are identified by the names Email, Folder, and Printing, which appear in the registry path in the place of *MethodName*.

### Notes

- A routing extension provides the Fax Service with routing functions for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

**See Also**
**Concepts**
ExtensionName

# MethodName\FriendlyName Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*\Routing Methods*MethodName*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies a longer name for the routing method named in the registry path by MethodName. Fax Service Manager displays the string stored by this entry in **Incoming Routing** to identify the routing methods provided by the routing extension.

### Notes

- A routing extension provides the Fax Service with routing methods for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

**See Also**
**Concepts**
MethodName

# Function Name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*\Routing Methods*MethodName*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | FaxRouteEmail \| FaxRouteStore \| FaxRoutePrint |

**Description**
Specifies the name of the function exported by the routing extension that is identified in the registry path by Routing Extensions. The function provides the routing method identified in the registry path by MethodName.

When *ExtensionName* is Microsoft Routing Extension, this entry has a default value of FaxRouteEmail, FaxRouteStore, or FaxRoutePrint, which correspond to *MethodName* being Email, Folder, or Printing.

Instead of the name of the function, Fax Service Manager displays the string stored in the registry entry MethodName\FriendlyName Entry to identify the routing method in the user interface.

### Notes

- A routing extension provides the Fax Service with routing methods for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

### Caution

- Do not change this entry. Changing this entry will cause the corresponding routing method to fail.

**See Also**
**Concepts**
Routing Extensions
MethodName
MethodName\FriendlyName Entry

# MethodName\Guid

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*\Routing Methods*MethodName*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | {9d3d0c32-9af2-11d0-abf7-00c04fd91a4e} |

**Description**
Specifies the globally unique identifier (GUID) for the fax routing method identified in the registry path by MethodName. The routing method is provided to the Fax Service by the routing extension that is identified in the registry path by Routing Extensions.

### Notes

- A routing extension provides the Fax Service with routing methods for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

### Caution

- Do not change this entry. Changing this entry will cause the corresponding routing method to fail.

**See Also**
**Concepts**
Routing Extensions
MethodName

# Priority

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*\Routing Methods*MethodName*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 3 |

**Description**
Specifies the priority for the global routing method identified in the registry path by MethodName. Fax Service Manager applies global routing methods to incoming faxes in order of priority. An administrator can change the priority order in Fax Service Manager.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry, in Fax Service Manager, double-click **Fax** if **Incoming Routing** is not visible. Next, expand **Incoming Routing**, and then click **Global Methods**. In the results pane, click the name of the routing method that you want to change, and then, on the **Action** menu, click **Move Up** or **Move Down**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

### Notes

- A routing extension provides the Fax Service with routing methods for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

**See Also**
**Concepts**
Routing Extensions
MethodName

# Email

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*\Routing Methods

**Description**
The Email subkey stores configuration information for the routing method **Route through e-mail**, which is provided to the Fax Service by the Microsoft Routing Extension.

The Email subkey is a type of subkey described by the subkey *MethodName*. For more information about this type of subkey, see MethodName.

### Notes

- A routing extension provides the Fax Service with routing methods for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

**See Also**
**Concepts**
MethodName
Routing Methods\Folder
Printing Subkey
ExtensionName

# Routing Methods\Folder

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*\Routing Methods

**Description**
The Folder subkey stores configuration information for the routing method **Store in a folder**, which is provided to the Fax Service by the Microsoft Routing Extension.

The Folder subkey is a type of subkey described by the subkey *MethodName*. For more information about this type of subkey, see MethodName.

### Notes

- A routing extension provides the Fax Service with routing methods for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

**See Also**
**Concepts**
MethodName
Email
Printing Subkey
ExtensionName

# Printing Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions*ExtensionName*\Routing Methods

**Description**
The Printing subkey stores configuration information for the routing method **Print**, which is provided to the Fax Service by the Microsoft Routing Extension.

The Folder subkey is a type of subkey described by the subkey *MethodName*. For more information about this type of subkey, see MethodName.

### Notes

- A routing extension provides the Fax Service with routing methods for incoming faxes.
- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

**See Also**
**Concepts**
MethodName
Email
Routing Methods\Folder
ExtensionName

# Microsoft Routing Extension

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Routing Extensions\

**Description**
The Microsoft Routing Extension subkey stores configuration information for the Microsoft Routing Extension, which is included with the Fax Service. A routing extension provides the Fax Service with routing functions for incoming faxes.

The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system. For a description of the contents of these types of subkeys, see entries in the ExtensionName subkey.

**See Also**
**Concepts**
ExtensionName

# Fax\Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

**Description**
The Security subkey stores security information for the Fax Service.

# Descriptor

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Security

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | | |

**Description**
Stores the security descriptor for the Fax Service. This entry is configured automatically when the Fax Service registers a routing extension.

### Notes

- A security descriptor is a structure and associated data that contains the security information for a securable object. A security descriptor identifies the owner and primary group of an object. It can also contain a discretionary access-control list (DACL) that controls access to the object, and a system access-control list (SACL) that controls the logging of attempts to access the object.

- A routing extension provides the Fax Service with routing methods for incoming faxes.

- The Microsoft Routing Extension is installed automatically when you add the Fax Service component to the operating system. Other routing extensions might also be installed on a system.

### Caution

- Do not change this entry. Changing this entry will cause the Fax Service to fail.

# SentItems

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

The SentItems subkey stores the configuration for the folder in which the Fax Service stores outgoing faxes.

# SentItems\AgeLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\SentItems

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 60 |

**Description**
Specifies the time, in days, that the Fax Service allows a sent fax to remain in the Sent Items folder. (The location of the Sent Items folder is specified by the registry entry SentItems\Folder.) If a resent fax is in the Sent Items folder longer than this time, the Fax Service automatically deletes it. If the value of this entry is zero, there is no time limit, and the Fax Service does not delete faxes from the Sent Items folder.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Sent Items** tab. Under **Archive rules**, check **Automatically delete faxes older than** if it is not already checked. Next, change the number in the **days** box by selecting the number and typing a new number or by clicking the arrow buttons next to the number. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

**See Also**
**Concepts**
SentItems\AgeLimit
Folder Entry

# SentItems\Folder

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\SentItems

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | C:\Documents and Settings\All Users\Application Data\Microsoft\Windows NT\MSFax\SentItems |

**Description**
Specifies the location of the Sent Items folder of the Fax Service. The Fax Service stores a sent fax in the Sent Items folder when the fax has been sent successfully.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Sent Items** tab. Under **Archive rules**, check or clear **Generate warning in the Event Log**, if it is not already checked. Next, change the number in the **Size quota low watermark** box by selecting the number and typing a new number or by clicking the arrow buttons next to the number. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

**See Also**
**Concepts**
Folder Entry

# HighWatermark Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\SentItems

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 100 |

**Description**
Specifies a size threshhold, in megabytes, for the Sent Items folder of the Fax Service. When the combined size of the contents of the Sent Items folder reach the size specified by this entry, the Fax Service enters a warning event into the Windows event log if the registry entry SizeQuotaWarn Entry specifies that the service do so. The Sent Items folder is the folder in which the Fax Service stores faxes that have been sent successfully.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Sent Items** tab. Under **Archive rules**, check **Generate warning in the Event Log** if it is not already checked. Next, change the number in the **Size quota high watermark** box by selecting the number and typing a new number, or by clicking the arrow buttons next to the number. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

## Notes

- This entry specifies the higher of two threshholds for the Sent Items folder. The other threshhold is specified by the registry entry LowWatermark Entry.
- Checking or clearing **Generate warning in the Event Log** changes the value of the registry entry SizeQuotaWarn Entry.

**See Also**
**Concepts**
SizeQuotaWarn Entry
LowWatermark Entry

# LowWatermark Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\SentItems

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 95 |

**Description**
Specifies a size threshhold, in megabytes, for the Sent Items folder of the Fax Service. When the combined size of the contents of the Sent Items folder reach the size specified by this entry, the Fax Service enters a warning event into the Windows event log if the registry entry SizeQuotaWarn Entry specifies that the service do so. The Sent Items folder is the folder in which the Fax Service stores faxes that have been sent successfully.

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Sent Items** tab. Under **Archive rules**, check **Generate warning in the Event Log**, if it is not already checked. Next, change the number in the **Size quota low watermark** box by selecting the number and typing a new number, or by clicking the arrow buttons next to the number. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

## Notes

- This entry specifies the lower of two threshholds for the inbox folder. The other threshhold is specified by the registry entry HighWatermark Entry.
- Checking or clearing **Generate warning in the Event Log** changes the value of the registry entry SizeQuotaWarn Entry.

**See Also**
**Concepts**
SizeQuotaWarn Entry
HighWatermark Entry

# SizeQuotaWarn Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\SentItems

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the Fax Service enters a warning event into the Windows event log when the contents of the Fax Service Sent Items folder reach a combined size equal to or greater than one of two size threshholds (specified by LowWatermark Entry and HighWatermark Entry). The Sent Items folder is the folder in which the Fax Service stores faxes that have been sent successfully.

| Value | Meaning |
|---|---|
| 0 | The Fax Service does not create warning events in the Windows event log regarding the size of the Sent Items folder. |
| 1 | The Fax Service creates a warning event in the Windows event log when the contents of the Sent Items folder reach a combined size equal to or greater than Inbox\LowWatermark or Inbox\HighWatermark). |

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, and then click the **Sent Items** tab. Under **Archive rules**, check or clear **Generate warning in the Event Log**. Click **OK**.

**Activation Method**
Changes made to this entry through **Computer Management** or **Fax Service Manager** are effective immediately.

**See Also**
**Concepts**
LowWatermark Entry
HighWatermark Entry

# SentItems\Use

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\SentItems

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the Fax Service archives outgoing faxes.

| Value | Meaning |
|---|---|
| 0 | The Fax Service does not archive outgoing faxes. |
| 1 | The Fax Service archives outgoing faxes. |

Archived outgoing faxes are stored in the location specified by the registry entry Folder Entry (in the same subkey as this entry).

**Change Method**
You can change this entry by using Fax Service Manager, which is part of the Fax Service on Windows Server 2003, but not on Windows XP. To start Fax Service Manager, click the **Start** menu, point to **Accessories**, then point to **Communications**, then point to **Fax**, and then select **Fax Service Manager**. You can also access Fax Service Manager in the **Computer Management** console.

To change the value of this entry by using Computer Management, open Computer Management. In the console tree, click **Services and Applications**, and then click **Fax**. From the **Action** menu, click **Properties**, click the **Sent Items** tab, and check **Archive all outgoing faxes to this folder**, if it is not already checked. Next, type a valid folder path in the **Archive all incoming faxes to this folder** box, or click **Browse** to browse the file system for the folder. Click **OK**.

**Activation Method**
Changes that you make by using **Fax Service Manager** or **Computer Management** are effective immediately.

**See Also**
**Concepts**
Inbox\Use
Folder Entry

# Fax\Setup Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

**Description**
The Setup subkey stores information about the Fax Service, such as the version number of the service and which components are installed on the system.

# Installed

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 1 |

**Description**
Specifies whether the Fax Service component is installed on the operating system.

| Value | Meaning |
|---|---|
| 0 | The Fax Service is not installed. |
| 1 | The Fax Service is installed. |

This entry is set automatically when the Fax Service component is installed or removed from the operating system.

**Change Method**
Do not change this entry.

### Caution

- This entry is maintained automatically by the Fax Service setup program. Do not change it.

# InstalledComponents Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Indicates, to the Fax Service, which Fax Service modules are installed on the operating system. If a bit in the value of this entry is 1, the Fax Service behaves as if the corresponding module, as shown in the following table, is absent.

| Value | Meaning |
|---|---|
| 1 | Fax Service (Fxssvc.exe) |
| 2 | Fax Console (Fxsclnt.exe) |
| 4 | Fax Service Manager (Fxsadmin.dll) |
| 8 | Send wizard (Fxssend.exe) |
| 16 | Configuration wizard (Fxscfgwz.dll) |
| 32 | Cover page editor (Fxscover.exe) |
| 64 | Client Help (Fxsclnt.hlp) |
| 128 | Client context help (Fxsclnt.chm) |
| 256 | Fax Service Manager Help (Fxsadmin.hlp) |
| 512 | Fax Service Manager context help (Fxsadmin.chm) |
| 1024 | Cover page editor help (Fxscover.chm) |
| 2048 | Fax Monitor (Fxsst.dll) |
| 4096 | Fax printer driver (Fxsui.dll) |

This entry does not exist by default, and its absence indicates to the Fax Service that all modules are installed.

**Activation Method**
You must stop and restart the Fax Service to make changes to this entry effective.

### Note

- This entry was invented for testing purposes. There is no reason for an administrator to add this entry to the registry.

### Caution

- If you create this entry and set any of the bits of its value equal to 1, you might prevent the Fax Service from working successfully.

# W2K_Inbox

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores the location in the file system of the Inbox folder of Personal Fax for Windows. This entry is created automatically by the Fax Service setup program when the Fax component is installed as an upgrade.

### Note

- When you upgrade from an operating system running Personal Fax for Windows (for example, Windows 2000) to Windows Server 2003, the Windows Server 2003 Fax service replaces Personal Fax for Windows as your fax application, and keeps your previous fax configuration settings.

# W2K_SentItems

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores the location in the file system of the Sent Items folder of Personal Fax for Windows. This entry is created automatically by the Fax Service setup program when the Fax component is installed as an upgrade.

### Note

- When you upgrade from an operating system running Personal Fax for Windows (for example, Windows 2000) to Windows Server 2003, the Windows Server 2003 Fax Service replaces Personal Fax for Windows as your fax application, and keeps your previous fax configuration settings.

# ProductBuild

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Identifies the build number of the Fax Service that is currently installed on the operating system.

**Change Method**
Do not change this entry.

### Caution

- This entry is set automatically when the Fax Service component is installed on the operating system, and you should not change it.

# ServerCoverPageDir

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | Microsoft\Windows NT\MSFax\Common Coverpages |

**Description**
Specifies the location in the file system, relative to %ALLUSERSPROFILE%\Application Data, where the Fax Service stores cover page templates that are available to all users. This entry is configured automatically when the Fax Service component is installed on the operating system.

### Note

- Users can send faxes with one of the cover pages stored in the location specified by this entry, or they can use their own cover pages if the administrator enables personal cover pages in Fax Service Manager.

# Original Setup Data

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup

**Description**
The Original Setup Data subkey stores configuration information for the Fax service to use to configure newly discovered fax devices.

# Original Setup Data\CSID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup\Original Setup Data

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | Fax |

**Description**
Specifies the default Called Subscriber ID (CSID) for new fax devices that the Fax service registers.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax service to make changes to this entry effective.

### Note

- When receiving a fax transmission, the Fax service sends the CSID to the sender to identify the receiver. A CSID usually consists of the telephone number of the receiving fax device and a name. The CSID is limited to 20 characters.

**See Also**
**Concepts**
{F10A5326-0261-4715-B367-2970427BBD99}\CSID
Original Setup Data\TSID

# Original Setup Data\Flags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup\Original Setup Data

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 2 |

**Description**
Specifies whether a fax device, by default, is enabled to automatically answer calls and receive faxes, and whether the device is enabled to send faxes. The Fax Service configures this default for each newly discovered fax device when it registers the fax device.

This entry stores its configuration information as bit-masked values. The following table explains the possible values.

| Value | Meaning |
|---|---|
| 1 | The fax device is enabled to automatically answer calls and receive faxes. |
| 2 | The fax device is enabled to send faxes. |
| 3 | The fax device is enabled both to send and to receive faxes. |

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax Service to make changes to this entry effective.

**Tip**

- If you are deploying many systems on which you want a different default number of rings, you can change the value of this entry during unattended installation of the operating system. For more information about unattended installations, see "Planning and Designing Unattended Installations" in the *Windows Resource Kit Deployment Planning Guide*.

**See Also**
**Concepts**
{F10A5326-0261-4715-B367-2970427BBD99}\Flags
Original Setup Data\Rings

# Original Setup Data\Rings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup\Original Setup Data

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 2 |

**Description**
Specifies the default number of times that an incoming call is allowed to ring before the fax device answers. The Fax service configures this default for each newly discovered fax device when it registers the fax device.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax service to make changes to this entry effective.

### Note

- A fax device only answers calls if it is configured to do so by the device-specific registry entry {F10A5326-0261-4715-B367-2970427BBD99}\Flags.

**See Also**
**Concepts**
{F10A5326-0261-4715-B367-2970427BBD99}\Rings

# Original Setup Data\TSID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup\Original Setup Data

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | Fax |

**Description**
Specifies the default transmitting subscriber identification (TSID) for every new fax device installed on the system.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Fax service to make changes to this entry effective.

### Note

- The Fax service sends the TSID when making fax transmissions, depending on the value of registry entry **UseDeviceTsid**. The receiving fax device might use the TSID in a variety of ways; fax machines often display the TSID when receiving a fax and also might print it on the top of each received page. The local Fax service can also include the TSID in a banner, which is prepended to each transmitted page, when the value of registry entry Branding is set to 1.

**See Also**
**Concepts**
{F10A5326-0261-4715-B367-2970427BBD99}\TSID
Original Setup Data\CSID

# ServiceStartup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

The ServiceStartup subkey stores information for the Fax Service about the startup state of the service.

# RPCReady

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\Setup\ServiceStartup

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Specifies to the Fax Service that the Fax Service is ready for remote procedure call (RPC) connections.

**Caution**

- This registry entry is maintained automatically by the Fax Service. Do not modify it.

# TAPIDevices

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The TAPIDevices subkey stores information for the Fax Service about Telephony Application Programming Interface (TAPI) devices installed on the system.

# TAPIDevices\DeviceID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The *DeviceId* subkey stores configuration information for the Fax Service about the Telephony Application Programming Interface (TAPI) device specified by the name of this type of subkey. The Fax Service automatically maintains a *DeviceId* subkey for each TAPI device that it registers.

**Caution**

- Do not attempt to remove a device by deleting registry subkeys like this subkey. If you need to remove a fax device or prevent the Fax Service from using a device, then use Fax Service Manager to delete the device or create routing rules that exclude it from use. For more information about adding and removing fax devices and about routing rules, see Fax Service Manager Help.

# {Guid}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | | |

**Description**
Stores configuration information for the Fax service about a routing extension for a single fax device. The routing extension iis identified by the globally unique identifier (GUID) in the name of this entry, and the fax device is identified in the registry path by *DeviceId*. For example, the name of the folder for the Store In Folder routing method or the name of the printer for the Print routing method.

The values of this type of entry, which are specified by the GUID in their name and *DeviceID* in their path, are encrypted.

**Change Method**
These types of registry entries, which are specified by the GUID in their name and *DeviceID* in their path, are created and maintained automatically by the Fax service. Do not change them.

# Modem Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The Modem subkey stores configuration information for the modem specified by *DeviceId* in the registry path. The Fax Service creates the entries in this subkey during the first fax transmission that it makes with the specified device.

The Fax Service creates most of the entries in this subkey by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

**Note**

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

# AnswerCommand

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies commands that the Fax service sends to the specified modem to configure it for adaptive answer and to answer a call. The modem is specified by *DeviceId* in the registry path.

### Notes

- The Fax Service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

# EnableV17Recv

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | "1" |

**Description**
Specifies the command that the Fax Service sends to the modem to disable use of the communication protocol V.17 during reception of a fax.

| Value | Meaning |
|---|---|
| "0" | The Fax Service cannot use V.17 for reception of faxes. |
| "1" | The Fax Service can use V.17 for reception of faxes. |

Note that the numeric values shown in the preceding table are stored as strings by this entry.

**Notes**

- V.17 is a protocol established by the International Telecommunications Union (ITU) for fax communications. The Fax Service also supports ITU fax protocols V.29 and V.27.

- The Fax Service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

# HighestSendSpeed

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | "0" |

**Description**
Specifies, in bits per second, the highest over-the-wire speed that the Fax Service uses for sending a fax with the modem specified by *DeviceId* in the registry path. A value of "0" for this entry indicates the highest available speed.

### Notes

- This entry stores numeric values as string values.
- The T.30 protocol engine uses this setting when creating the DIS frame.
- The Fax service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.
- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

# ModemId

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the string for which the Microsoft Fax Service Provider (FSP) searches in the response from a modem. The FSP attempts to elicit the response from a modem by sending it the command specified in the ModemIdCmd entry. The modem is identified in the registry path by *DeviceId*. If the string is absent, the FSP reinstalls the modem in the operating system.

This entry is not always present in the registry. When it is absent, the FSP does not query the modem for identification.

### Notes

- This entry can by used to check the version of the firmware of the modem.
- The Fax Service creates most of the other entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.
- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.
- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol specified by the T.30 specification of the International Telecommunications Union (ITU) for standard fax modems. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
ModemIdCmd

# ModemIdCmd

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the command that the Microsoft Fax Service Provider (FSP) sends to a modem to request identification from the modem. The modem is identified in the registry path by *DeviceId*. If the FSP searches the response from the modem for the string specified by the ModemId entry, and if the string is absent, the FSP reinstalls the modem in the operating system.

This entry is not always present in the registry. When it is absent, the FSP does not query the modem for identification.

### Notes

- This entry can by used to check the version of the firmware of the modem.
- The Fax Service creates most of the other entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.
- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.
- A fax service provider is a dynamic-link library (DLL) that provides an application programming interface (API) that the Fax Service uses to interact with one or more fax devices. Microsoft supplies a fax service provider that supports the fax protocol specified by the T.30 specification of the International Telecommunications Union (ITU) for standard fax modems. Third parties can also supply fax service providers to support other fax-capable devices.

**See Also**
**Concepts**
ModemId

# ModemResponseDataConnect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Identifies for the Fax Service the response that the modem sends to indicate that it is connected to a remote modem on a data call. The modem is identified in the registry path by *DeviceId*.

### Notes

- The Fax Service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

**See Also**
**Concepts**
ModemResponseDataDetect

# ModemResponseDataDetect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Identifies the modem response that indicates that the incoming call is a data call. The Fax Service uses this identifier with the modem specified in the registry path by *DeviceId*.

### Notes

- The Fax Service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

# ModemResponseFaxConnect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Identifies for the Fax Service the response that modem identified in the registry path by *DeviceId* sends to indicate that it is connected to a remote modem on a fax call.

### Notes

- The Fax Service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

**See Also**
**Concepts**
ModemResponseFaxDetect

# ModemResponseFaxDetect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Identifies the modem response that indicates to that the incoming call is a fax call. The Fax Services uses this identifier with the modem specified in the registry path by *DeviceId*.

### Notes

- The Fax Service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

# ModemSendSpeeds

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores the sending capabilities of the modem specified by *DeviceId* in the registry path.

| Value | Meaning |
|---|---|
| 0 | V.27 (2400 bps) |
| 1 | V.29 (7200/9600 bps) |
| 2 | V.27 (2400/4800 bps) |
| 4 | V.33 (12000/14400 bps) |
| 8 | V.17 (7200/9600/12000/14400 bps) |

The Microsoft Fax Service Provider (FSP) determines the value of this entry when it detects and installs the modem to the operating system. The FSP sets the fifth bit (8) to 0 if the registry entry **EnableV17Send** specifies that the modem should not use the V.17 protocol to transmit faxes.

# PreAnswerCommand

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the modem command string that the Fax Service sends to the modem just before sending the answer command. By default, this entry stores an empty string.

This entry is created automatically by the Fax Service when it sends a fax for the first time using the modem identified in the registry path by *DeviceId*.

### Notes

- The Fax Service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

# PreDialCommand

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Specifies the modem command string that the Fax Service sends to the modem just before sending the dial command. By default, this entry stores an empty string.

Before sending the modem command specified by this entry, if there is a command specified, then the Fax Service sends the setup command, specified by the registry entry **SetupCommand**, and the modem class command, specified by the registry entry **FixModemClass**.

This entry is created automatically by the Fax Service when it sends a fax for the first time using the modem identified in the registry path by *DeviceId*.

### Notes

- The Fax Service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

# ResetCommand

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Fax\TAPIDevices*DeviceId*\Modem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | AT&FS0=0E0V1Q0 |

**Description**
Specifies the modem command string that the Fax Service sends to the modem before commencing each transmission and before answering each incoming call, unless the modem is in auto-answer mode. This reset command instructs the modem to load its factory default settings, disable auto-answer, turn off command echo, use words (rather than numbers) for results messages, and turn on results messaging. If the value of this entry is null, then the Fax Service does not send a reset command.

This entry is created automatically by the Fax Service when it sends a fax for the first time using the modem identified in the registry path by *DeviceId*.

### Notes

- The Fax Service creates entries in the subkey Fax\TAPIDevices*DeviceId*\Modem by copying the values of identically named entries in the part of the registry that stores settings for Unimodem 5. Unimodem 5 acquires some of these settings from the setup information (.inf) file for the modem when Windows installs the modem driver.

- Unimodem 5 is a Telephony Service Provider that provides an abstraction for modem devices such that applications can operate transparently across a wide variety of types and makes of modems. In addition, modem installation is greatly simplified since Unimodem 5 is aware of the operating characteristics of these modems, and their drivers are already available in the operating system.

- If the reset command fails, Unimodem 5 sends the command string **AT&FS0=0E0**. If this replacement command fails, Unimodem 5 sends the command string **AT&FE0**. If all of these reset commands fail, modem initialization fails.

# Internet Explorer Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Media

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Internet Explorer

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Autoplay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Internet Explorer\media

**Description**
Stores configuration data for the policy setting **Configure Media Explorer Bar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# MSMQ

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

**Description**
The MSMQ subkey stores configuration data for Message Queuing (also known as MSMQ). This subkey and its contents appear only if you have installed the Message Queuing service.

Message Queuing is a store-and-forward service that enables applications to communicate across heterogeneous networks and operating systems, even when the networks and applications are running at different times. Message Queuing stores messages in queues and forwards them when the destination networks and applications come online. The entries in this subkey apply to all users of the computer.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ

**Description**
The Parameters subkey stores configuration options for Message Queuing (also known as MSMQ).

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# AckTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*milliseconds*) | (Message Queuing calculates a value for this entry dynamically.) |

**Description**

Specifies how long Message Queuing (also known as MSMQ) waits for internal message acknowledgment of a network session before deciding that the session should be closed.

By default, Message Queuing calculates an optimal acknowledgment timeout based on the quality of the communication line. However, you can override the result of the calculation by adding this entry to the registry and setting it to a different value.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Adding this entry to the registry prevents the system from using an optimal value calculated for each line. If you set the value of this entry too low, then sessions can be disconnected before any messages are delivered.

# CertInfoCacheSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*certificates*) | 0x35 (53) |

**Description**
Specifies the maximum number of certificates stored in the certificate cache, which stores information about certificates found in incoming messages. This cache helps Message Queuing (also known as MSMQ) authenticate incoming messages more quickly.

If the cache is full (that is, if the number of certificates equals the value of the **CertInfoCacheSize** entry) when the next certificate arrives, then Message Queuing deletes half of the certificates in the cache. The oldest certificates are deleted first.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Entries in the certificate cache are not subject to an expiration time. Certificates are deleted from the cache only when the cache is full.

# MsmqIpPingPort

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*port number*) | 0xDC7 (3527) |

**Description**
Specifies the User Datagram Protocol (UDP) port that Message Queuing (also known as MSMQ) uses to ping servers. Message Queuing uses an internal ping algorithm to quickly determine the availability and configuration of a server.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- This entry is used only when the value of UsePing is 1.

### Caution

- This value must be the same on all computers in the enterprise. Using different UPD ports can significantly increase the time it takes for Message Queuing to establish network connections.

# MsmqIpPort

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*port number*) | 0x709 (1801) |

**Description**

Specifies the TCP port that Message Queuing (also known as MSMQ) uses to connect queue managers.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Caution

- Do not change the value of this entry. Changing this value can prevent Message Queuing computers from exchanging messages.

# PerThreadDSServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Enables applications to work with a different directory service server for each thread.

The default value allows all threads to work with the same directory service server, and if one thread changes the directory service, then it affects the others. In Windows Server 2003, some operations might require access to a specific directory service server. Enabling this mode guarantees that the directory service activity of the other threads will not change the server. This is relevant only in Message Queuing (also known as MSMQ) mixed mode, when some of the servers in the site are running MSMQ 1.0 and some are running Message Queuing.

| Value | Meaning |
|---|---|
| 0 | All threads access the same directory service server. |
| 1 | All threads can access a different directory service server. |

# QmPbKeyCacheSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *keys* | 0x35 (53 keys) |

**Description**
Specifies the maximum number of keys that can be stored in the queue manager (QM) public encryption key cache. This cache stores public encryption keys recently used in replication and notification messages.

If the cache is full (that is, if the number of keys is equal to the value of this entry) when the next key arrives, then Message Queuing (also known as MSMQ) deletes half of the keys in the cache. The oldest keys are deleted first.

All replication and notification messages are digitally signed by the Message Queuing Information Store (MQIS) that generated the message. The Message Queuing computers that receive the message validate its signature by using the public key of the sending MQIS. Storing recently used keys in a cache reduces the time it takes to verify the signature.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- Entries in the public encryption key cache are not subject to an expiration time. These entries are deleted only when the cache is full.

# QMThreadNo

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0x10 *threads* | Windows Server 2003: *(Number of processors x 5) + 3* Windows XP Professional: *Number of processors x 3* |

**Description**
Specifies the number of threads created in the Message Queuing (also known as MSMQ) process (Mqsvc.exe) to handle messages.

Increasing the value of this entry can improve the performance of Message Queuing, but too many threads can overload the processor.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# RemoteQMMachine

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Computer name* | *Computer name* |

**Description**
Specifies the name of the supporting server computer for the local Message Queuing (also known as MSMQ) dependent client. The supporting server routes messages to the client, transmits queries and responses, and updates the client's view of the enterprise.

**Change Method**
To change the value of this entry, in Control Panel, double-click **MS Message Queue**, and then click the **MSMQ Client** tab.

### Note

- This entry is used only when it appears in the registry of a Message Queuing client that is dependent on a parent server, and it appears there by default.

# RMFlushInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *milliseconds* | 0x1B7740 (1,800,000 millseconds = 30 minutes) |

**Description**
Specifies how often Message Queuing (also known as MSMQ) takes a checkpoint of Resource Manager transactions. The checkpoint saves the current transactions and exactly-once-delivery state of Message Queuing.

The value of this entry is used when no checkpoints have been taken in the Resource Manager as a result of other events, such as when its log file is more than half full. It guarantees that a checkpoint is performed no less frequently than the time specified.

If this value is too large, it might prolong the time it takes to recover files to their last checkpoint. However, if the value is too small, frequent checkpoints can degrade Message Queuing performance.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Do not change the value of this entry. An inappropriate value can interfere with logging, transactions, and the exactly-once delivery mechanism.

# RpcCancelTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *minutes* | 0x5 |

**Description**
Specifies how long a Message Queuing (also known as MSMQ) client waits after submitting a request before it notifies the user that the Message Queuing Information Store (MQIS) server has not responded. This entry appears only in the registries of Message Queuing client computers.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- This entry is used only when it appears in the registry of a Message Queuing client computer.

# CryptReceiveKeyCacheSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*keys*) | 0x7F (*127*) |

**Description**
Specifies the maximum number of symmetric encryption keys found in incoming messages that can be stored in the cryptographic key cache on the receiving server.

If the cache is full (that is, if the number of incoming keys is equal to the value of the **CryptReceiveKeyCacheSize** entry) when the next certificate arrives, then Message Queuing (also known as MSMQ) deletes half of the certificates in the cache. The oldest certificates are deleted first.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Incoming keys in the cryptographic key cache are not subject to an expiration time. Incoming keys are deleted from the cache only when the cache is full.

**See Also**
**Concepts**
CryptSendKeyCacheSize

# CryptSendKeyCacheSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*keys*) | 0x35 (*53*) |

**Description**
Specifies the maximum number of outgoing symmetric encryption keys that can be stored in the cryptographic key cache on the sending server.

If the cache is full (that is, if the number of outgoing keys is equal to the value of the **CryptReceiveKeyCacheSize** entry) when the next key arrives, then Message Queuing (also known as MSMQ) deletes half of the keys in the cache. The oldest keys are deleted first.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Incoming keys in the cryptographic key cache are not subject to expiration. Incoming keys are deleted from the cache only when the cache is full.

**See Also**
**Concepts**
CryptReceiveKeyCacheSize

# DeferredInit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Defers the contact between Message Queuing (also known as MSMQ) and the MQIS server that would otherwise occur automatically at startup.

| Value | Meaning |
|---|---|
| 0 | Automatic. Message Queuing contacts the MQIS server automatically when the service starts. |
| 1 | Deferred. Message Queuing does not contact the MQIS server automatically when the service starts. Instead, it contacts the MQIS when the first Message Queuing program requires access to the MQIS. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- The deferred setting can cause a brief delay while Message Queuing connects to the MQIS server. This delay might cause the first program that requires the MQIS server to appear momentarily unresponsive.

# DSAdsRefreshIPSitesIntervalSecs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*seconds*) | 0xE10 (*3,600 seconds = 1 hour*) |

**Description**

Specifies the minimum interval between successive updates of the IP subnets data for the site. Message Queuing (also known as MSMQ) servers use information from these updates to find the IP addresses of computer sites.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# DSAdsRefreshNT4SitesIntervalSecs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*seconds*) | 0x5460 (*21,600 seconds = 6 hours*) |

**Description**
Specifies the minimum interval between successive updates of the list of primary site controllers running Windows NT 4.0. Message Queuing (also known as MSMQ) uses information from these updates to determine whether it needs to support clients using previous versions of the operating system.

### Note

- This entry is used only when Windows Server 2003 or Windows 2000 Message Queuing servers are supporting Windows NT 4.0 clients.

# DSCliSearchAdsForServersIntervalSecs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*seconds*) | 0x708 (*1,800 seconds = 30 minutes*) |

**Description**
Specifies the minimum interval between successive searches through Active Directory for Message Queuing (also known as MSMQ) servers. Message Queuing searches the Active Directory when it cannot connect to known Message Queuing servers.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Caution

- Do not change the value of this entry. Changing it can prevent Message Queuing from locating known servers in the Active Directory.

# DSSecurityCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *Security descriptor* | (None. The security descriptor of the computer is copied from the MQIS.) |

**Description**

Stores the security descriptor of the local Message Queuing (also known as MSMQ) computer. The security descriptor of the computer determines, among other things, which users or groups of users have permission to create queues on the computer. The value of this registry entry is used if an offline user attempts to create a private queue when a queue request is processed.

The security descriptor is saved in the registry to allow private queues to be created while the computer is not connected to the network.

**Caution**

- Do not change the value of this entry. If you do, then the online operation of Message Queuing might be inconsistent with the offline operation.

# IdleAckDelay Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*milliseconds*) | 500 |

**Description**

Sets the maximum idle delay time between receiving a message and sending an acknowledgment to the source.

Idle time is the time during which no messages are sent on the session. It ensures that computers creating a session for a limited time only will get all acknowledgments possible during the session. This prevents unnecessary resending due to lost acknowledgments.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# This Topic Is No Longer Available

We're sorry—the topic that you requested is no longer available. Use the search box below to find related information.

# LastPrivateQueueId

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | 0x0–0xFFFFFFFF | 0xA (*10*) |

**Description**
Records the serial number of the last private queue created. Message Queuing (also known as MSMQ) updates this entry any time a private queue is created.

### Caution

- Do not change the value of this entry. Doing so could cause Message Queuing to generate the same identifier for a private queue and damage both the queue and the message service.

# LogDataCreated

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether log files are missing.

When Message Queuing (also known as MSMQ) runs for the first time, it creates local log files and sets the value of this entry to 1. Later, Message Queuing uses this entry to determine whether log files are missing. If they are, then Message Queuing does not start, and an appropriate event is generated.

| Value | Meaning |
|---|---|
| 0 | Local log files are not created on the computer. |
| 1 | Local log files are created on the computer. |

# LogMgrBuffers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*buffers*) | 0x190 (*400*) |

**Description**

Specifies how many buffers the Log Manager allocates for each new log file.

This entry is used to optimize Log Manager performance.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Do not change the value of this entry. An inappropriate value can interfere with logging, transactions, and the exactly-once delivery mechanism.

# LogMgrChkptInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*milliseconds*) | 0x2710 (*10 seconds*) |

**Description**
Specifies how often the Log Manager performs an internal checkpoint on the log. This entry is used to optimize Log Manager performance.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Do not change the value of this entry. An inappropriate value can interfere with logging, transactions, and the exactly-once delivery mechanism.

# LogMgrFileSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*bytes*) | 0x600000 (*6 MB*) |

**Description**

Specifies the size of the Log Manager file. The size of this file affects the maintenance of the Log Manager file and the performance of Message Queuing (also known as MSMQ). Message Queuing performs checkpoints, in which it saves the current log entries, when the log file occupies more than half of the space allocated for it. If the value of this entry is too small, then Message Queuing takes checkpoints frequently, diminishing overall performance. If the file is too large, then it consumes unnecessary disk space.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Do not change the value of this entry. An inappropriate value can interfere with logging and with the operation of Message Queuing.

# LogMgrFlushInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *milliseconds* | 0x32 (50) |

**Description**

Specifies how often the Log Manager flushes the log file.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Caution

- Do not change the value of this entry. An inappropriate value can interfere with logging, transactions, and the exactly-once delivery mechanism.

# LogMgrProbeInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *milliseconds* | 0x64 (100) |

**Description**
Specifies how often the Log Manager considers writing an empty record to the log file buffer. Empty records are written to fill the log file buffers so that the buffer contents are written to disk.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Do not change the value of this entry. An inappropriate value can interfere with logging, transactions, and the exactly-once delivery mechanism.

# LogMgrTimerInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *milliseconds* | 0xA (10) |

**Description**

Specifies how often the Log Manager considers performing a flush or internal checkpoint (*pulse interval*).

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Do not change the value of this entry. An inappropriate value can interfere with logging, transactions, and the exactly-once delivery mechanism.

# StoreAckTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *milliseconds* | (Not in registry) |

**Description**
Specifies how long Message Queuing (also known as MSMQ) waits for internal message acknowledgement of persistent messages on a network session before deciding that the session should be closed.

By default, Message Queuing calculates an optimal acknowledgment timeout based on the quality of the communication line. However, you can add this entry to the registry and then set it to a different value to override the result of the calculation.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Adding this entry to the registry prevents the system from using an optimal value calculated for each line. If you set the value of this entry too low, then sessions can be disconnected before any messages are delivered.

# StoreMqLogPath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Path and directory name* | *Installation path*\MSMQ\Storage |

**Description**
Specifies the directory where Message Queuing (also known as MSMQ) stores the transaction log file, also known as the QM log. Message Queuing stores the current state of active transactions and incoming transacted messages in the transaction log file.

**Change Method**
To view or change the value of this entry, in Control Panel, double-click **Message Queuing**. The directory where the transaction log and all checkpoint files are stored appears in the **Transaction Logger Folder** box.

### Note

- Message Queuing stores all checkpoint files in the same directory as transaction log file (QM Log). If you use **Message Queuing** in Control Panel to change the value in the **Transaction Logger Folder** box, then Message Queuing moves all checkpoint files and the transaction log to the new location.

# SysQueuePriority

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x7FFF (32,767) |

**Description**

Specifies the base priority of messages destined for system queues.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Caution

- Do not reduce the value of this entry. Permitting lower-priority messages to be sent to queues might delay messages with higher priority, including system messages.

# TCPNoDelay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Disables nagling. Nagling is a TCP feature that combines several small packets into a single, larger packet for more efficient transmission.

By default, Message Queuing (also known as MSMQ) enables nagling on the TCP sockets it uses. This improves overall performance, but it might briefly delay transmission of smaller packets. If the delay is undesirable or unacceptable, then you can add this entry to the registry to disable nagling.

| Value | Meaning |
|---|---|
| 0 | Nagling is enabled. |
| 1 | Nagling is disabled. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# TransactionMode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | DefaultCommit \| (blank) | DefaultCommit |

**Description**
Specifies whether Message Queuing (also known as MSMQ) uses the default-abort mode or the default-commit mode to process transactions.

When the value of this entry is blank, the service uses the default-abort mode. When the value of this entry contains DefaultCommit, the service uses the default-commit mode.

# UsePing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether Message Queuing (also known as MSMQ) uses an internal ping algorithm to determine the availability and configuration of a server. Using a ping, the service can find this information quickly.

| Value | Meaning |
|---|---|
| 0 | Message Queuing does not use the ping algorithm. |
| 1 | Message Queuing does use the ping algorithm. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
MsmqIpPingPort

# UserCacheSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *entries* | 0xFD (253 entries) |

**Description**

Specifies the maximum number of entries that can be stored in the user account information cache. This cache stores information about users who recently sent messages to the computer.

When Message Queuing (also known as MSMQ) receives a message, it uses user account information to verify that the sending user has permission to send messages to the particular queue. Storing user account information in the cache reduces the time it takes to verify user permissions.

If the cache is full (that is, if the number of entries is equal to the value of **UserCacheSize**) when the next entry arrives, then Message Queuing deletes half of the entries in the cache. The oldest entries are deleted first.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# WriteMsgTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *seconds* | 10 |

**Description**
Determines how long write-request messages are active. When the value of this entry expires, the write-request message is deleted.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- This entry is used only when it appears in the registry of a Message Queuing MQIS server.

# XactDeadLetterAlways

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether messages that are timed-out or undeliverable are placed in the XactDeadLetter queue.

In MSMQ 1.0, transactional messages messages that were timed-out or undeliverable were placed in the XactDeadLetter queue by default. In Windows Server 2003, the messages are placed in the XactDeadLetter queue only if the MSMQ_DEADLETTER property is selected.

Setting the value of this entry to 1 forces the local Message Queuing service to revert to MSMQ 1.0 behavior. Use this entry when the change in default behavior might break the existing application.

| Value | Meaning |
|---|---|
| 0 | Messages are placed in the XactDead Letter queue only if the message property (MSMQ_DEADLETTER) was set on the message. |
| 1 | Messages are placed in the XactDeadLetter queue. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# XactDelayReceiveNack

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFE *seconds* | (None) |

**Description**

Specifies an additional delay for timing out non-received transactional messages sent from this computer.

Message Queuing (also known as MSMQ) will wait an additional amount of time, equal to the value of the **XactDelayReceiveNack** entry, after the Time-To-Be-Received (TTBR) has expired and before placing the message in the XactDeadLetterQueue. The purpose of this delay is to allow positive receive acknowledgments (sent shortly before the TTBR expires) to reach the sender. If this value is not explicitly set in the registry, then the value of **XactDelayReceiveNack** will be automatically set to the smallest of the Time-To-Be-Received and Time-To-Reach-Queue properties of the message.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# MaxMessageSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *bytes* | 0x400000 (4 MB) |

**Description**
Specifies the maximum size of the message file. Because messages must fit in the file, the value of this entry also limits the size of messages that Message Queuing (also known as MSMQ) sends and receives.

When Message Queuing starts, it restores all persistent messages in message files. If it finds a message file larger than the value of this entry, then it might not operate correctly. Also, if it encounters messages larger than this value, then it blocks all further transports from the sending computer. To restore communication, you must set the value of this entry greater than or equal to the size of the message that exceeded this value, and then restart Message Queuing.

If you increase this value, you might be able to transmit larger messages between applications on the same computer. However, because this entry affects both incoming and outgoing messages, in order to send a larger message to other computers running Message Queuing, you must increase this value on all participating systems. If any system has a lower value, it might block transports from computers sending larger messages.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
Message Queuing checks this entry only at startup. If you change the value of this entry, you must restart Message Queuing to make the change effective.

### Note

- This entry affects Message Queuing only. It does not change characteristics of the network or network programs that might limit the size of messages.

### Caution

- Do not change the value of this entry. Increasing this value might prevent you from sending messages larger than 4 MB. Decreasing this value might block transports between Message Queuing computers.

# MaxUnackedPacket

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF*Messages* | 0x40 (64 Messages) |

**Description**

Specifies the maximum size of the message queue to each network session. The value of this entry limits the number and size of messages that can be stored between internal acknowledgments from the receiving server. If the queue fills before an acknowledgment is received, additional messages are discarded.

If the queue is too small, messages might be lost. But because messages in the queue are released again when a session fails, a large queue can flood the system with duplicate messages.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# SeqIDAtLastRestore

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | (None) |

**Description**

Records the value of the last assigned sequence ID at the time of the last restore in the registry.

This entry is used to distinguish between restored messages and those messages sent after the restore. This entry is necessary to guarantee the exactly-once-delivery of transactional messages when the backup/restore utility is used. Do not edit this value.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# SeqMaxAckDelay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFE *seconds* | 0xA (10 seconds) |

**Description**

Defines the length of delay in sending acknowledgements for transacted messages.

Message Queuing (also known as MSMQ) sends an order acknowledgment for the delivered, transacted messages. The value of this entry defines the delay in sending these messages. The receiving computer never sends an order acknowledgment more frequently than once per the number of seconds specified in the value of this entry.

This entry provides for batching of messages, which saves unnecessary traffic. Too large a value might cause resending of the messages because the sending computer might think the packets did not arrive. Too small a value might cause more order-acknowledgment packets to be sent.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# SeqResend10Time

Updated: November 6, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *seconds* | 0x5460 (21,600 seconds = 6 hours) |

**Description**

Determines how often outgoing, transacted messages are resent because they are unacknowledged.

The value of this entry specifies the interval between resends for the tenth and subsequent resend attempts. If the sender still does not receive an acknowledgement from the intended recipient, the message is resent at the rate specified by this value until it times out.

The interval between earlier resends is determined by the values of the following entries: SeqResend13TimeSeqResend46Time, SeqResend79Time.

| **Note** |
|---|
| Windows does not add this entry to the registry. You can add it by editing the registry or by using a program that edits the registry. |

| **Caution** |
|---|
| Do not change the value of this entry. An inappropriate value can interfere with message delivery. |

# SeqResend13Time

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *seconds* | 0x1E (30) |

**Description**

Specifies how often outgoing, transacted messages are resent because they are unacknowledged.

The value of this entry specifies the interval between resends for the first three times the message is resent. If, after the third resend, the sender still does not receive an acknowledgement from the intended recipient, then the message is resent at intervals determined by the values of the following entries: SeqResend46Time, SeqResend79Time, **SeqResend10Time**.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Do not change the value of this entry. An inappropriate value can interfere with message delivery.

# SeqResend46Time

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *seconds* | 0x12C (300 seconds = 5 minutes) |

**Description**

Specifies how often outgoing, transacted messages are resent because they are unacknowledged.

The value of this entry specifies the interval between resends for the fourth, fifth, and sixth times the message is resent. If, after the sixth resend, the sender still does not receive an acknowledgment from the intended recipient, then the interval between acknowledgments is determined by the values of the following entries: SeqResend79Time, **SeqResend10Time**. The interval between the first three resends is determined by the value of SeqResend13Time.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Do not change the value of this entry. An inappropriate value can interfere with message delivery.

# SeqResend79Time

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *seconds* | 0x708 (1800 seconds = 30 minutes) |

**Description**

Specifies how often outgoing, transacted messages are resent because they are unacknowledged.

The value of this entry specifies the interval between resends for the seventh, eighth, and ninth times the message is resent. If, after the ninth resend, the sender still does not receive an acknowledgment from the intended recipient, then the interval between acknowledgments is determined by the value of **SeqResend10Time**.

The interval between earlier resends is determined by the values of the following entries: SeqResend13Time, SeqResend46Time.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Do not change the value of this entry. An inappropriate value can interfere with message delivery.

# CryptKeyExpirationTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*milliseconds*) | 0x927C0 (*600,000 milliseconds = 10 minutes*) |

**Description**
Specifies how long outgoing symmetric keys remain in the cryptographic key cache. When the time specified by this value expires, the key is deleted from the cache of the source server, and symmetric keys used for encrypting messages are renewed.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
CryptSendKeyCacheSize

# CertificationAuthorities

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

**Description**
The CertificationAuthorities subkey contains X509-V3 certificates from external certification authorities (CAs). These subkeys are created by Microsoft Crypto API 2.0. Message Queuing (also known as MSMQ) uses the certificates to generate digital signatures that uniquely identify the message sender. The certificate and the digital signature are sent with each authenticated message.

**Change Method**
To enable or disable certificates, in Control Panel, double-click **MS Message Queue**, click the **Security** tab, and then click the **Certification Authorities** button.

### Note

- The entries in the subkeys in CertificationAuthorities store actual certification authority (CA) certificates, not just pointers to entries in another location.

### Caution

- Do not change the value of entries in this subkey by editing the registry. Changing these values can disable the security features of Message Queuing.

# ClusteredQMs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

**Description**
The ClusteredQMs subkey stores entire registry sections of Message Queuing (also knowns as MSMQ) services running as cluster resources.

For every clustered Message Queuing resource, there is a subkey in the ClusteredQMs subkey that holds all the entries of this Message Queuing resource.

**Note**

- Only edit the entries in this subkey that are on the node currently hosting the resource, and only when the resource is online.

# MachineCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

**Description**
The MachineCache subkey of the MSMQ\Parameters subkey stores configuration data that associates the local Message Queuing Information Store (MQIS) client computer with other computers in the enterprise managed by Message Queuing (also known as MSMQ).

# CurrentMQISServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Computer name* | *Computer name* |

**Description**
Stores the computer name of the MQIS server currently selected by Message Queuing (also known as MSMQ) on this computer. The MQIS client uses the server specified for MQIS queries.

### Note

- This entry is used only when it appears in the registry of a Message Queuing MQIS client, and is updated each time the system starts.

**See Also**
**Concepts**
StaticMQISServer

# EnterpriseId

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *Identifier* | (None) |

**Description**
Stores a globally unique identifier (GUID) that uniquely identifies the enterprise in which the computer is installed. The value of this entry must be the same in all Message Queuing (also known as MSMQ) computers in the enterprise.

**Caution**

- Do not change the value of this entry. If you change it, Message Queuing responds as though you removed the computer from the enterprise, and the computer will not be able to access any MQIS server or to establish any network connections.

# ForcedDSServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Computer Names* | (None) |

**Description**
Stores the MQIS server name, which Message Queuing (also known as MSMQ) uses when its querying of the directory service for the MQIS servers list either fails or returns no results.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- **ForcedDSServer** is the name of the Message Queuing directory service server used for the join domain. The value of this entry will not override a server name if one is found by the Message Queuing service.

# MachineJournalQuota

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFE *bytes* \| 0xFFFFFFFF | 0xFFFFFFFF |

**Description**

Limits the amount of data that is stored in the journal file on the computer.

| Value | Meaning |
|---|---|
| 0x0–0xFFFFFFFE *bytes* | Journal file size limit |
| 0xFFFFFFFF | No limit on the size of the journal file |

**Change Method**

To change the value of this entry, use Message Queuing Explorer.

# MachineQuota

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFE *kilobytes* \| 0xFFFFFFFF | 0xFFFFFFFF |

**Description**
Limits the total amount of data that is stored in message files on the computer. The sum of bytes in all message files on the computer cannot exceed the value of this entry.

| Value | Meaning |
|---|---|
| 0x0–0xFFFFFFFE kilobytes | Total message file size limit. |
| 0xFFFFFFFF | There is no limit on the message file size. |

**Change Method**
To change the value of this entry, use Message Queuing Explorer.

# MasterId

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Globally unique identifier (GUID)* | (None) |

**Description**
Stores a globally unique identifier (GUID) that uniquely identifies the computer to other MQIS servers in the enterprise.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- This entry is used only when it appears in the registry of a Message Queuing MQIS server. It appears by default only on Windows 4.0 computers.

### Caution

- Do not change the value of this entry. If you change it, other MQIS servers will not recognize the computer.

# QMId

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *Globally unique identifier (GUID)* | (None) |

**Description**
Stores the globally unique identifier (GUID) that uniquely identifies the computer in the enterprise. The GUID is generated when Message Queuing (also known as MSMQ) is installed.

### Caution

- Do not change the value of this entry. Message Queuing stores this GUID when Message Queuing is set up. Changing the value of this entry might cause Message Queuing to misidentify the computer.

# StaticMQISServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Computer name* | (None) |

**Description**

Specifies the server that Message Queuing (also known as MSMQ) uses for this client.

By default, this entry does not appear in the registry, and Message Queuing selects a current server randomly from the list of available backup site controller (BSC) and primary site controller (PSC) servers at the site. By adding this entry to the registry, you direct Message Queuing to use only the server specified by this value. The value of this entry replaces the values of the MQISServer and CurrentMQISServer entries, so that this server is the only one known to the client.

A static server is recommended when one server is clearly preferable to others, such as specifying a local server when other servers are remote and rely on slow network connections.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- This entry is used only when it appears in the registry of a Message Queuing MQIS client.

# MQISServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Computer names* | (None) |

**Description**
Stores a list of the Message Queuing (also known as MSMQ) servers in the current site.

For computers with Message Queuing belonging to a Windows 2000 site, this list contains all Windows 2000 domain controllers running Message Queuing, as well as all Windows 2000 servers running Message Queuing with Routing. For computers with Message Queuing belonging to Windows NT 4.0 MSMQ Enterprise, this list contains the primary site controller (PSC) and all backup site controllers (BSCs) at the local site. Message Queuing compiles the list at intervals specified by the value of the **DSListRefresh** entry, and it stores the list in the value of this entry. The Message Queuing client selects a server from this list.

### Note

- This entry is used only when it appears in the registry of a Message Queuing client.

**See Also**
**Concepts**
StaticMQISServer

# Migration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

**Description**
The Migration subkey stores entries that determine how Message Queuing (also known as MSMQ) data is propagated to clients running previous versions of the operating system. Entries in the Migration subkey are used only when there are MSMQ 1.0 servers running Windows NT 4.0 and earlier. They are not used when all servers are running Windows Server 2003 or Windows 2000.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# HelloInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Migration

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *seconds* | 0x4B0 (1,200 seconds = 20 minutes) |

**Description**
Specifies how often a Message Queuing (also known as MSMQ) server sends *hello* messages to each client.

*Hello* messages are Layer Two Tunneling Protocol (L2TP) control messages sent at regular intervals to prevent an active, but silent, connection from being mistaken for an idle line. They are also known as *keep-alive* messages.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- Entries in the Migration subkey are used only when there are MSMQ 1.0 servers that are running Windows NT 4.0 and earlier. They are not used when all servers are running Windows Server 2003 or Windows 2000.

**Tip**

- The length of the replication interval is a multiple of the length of the hello-message interval. The multiplier that determines the length of the replication interval is stored in the value of TimesHelloForReplicationInterval.

**See Also**
**Concepts**
TimesHelloForReplicationInterval

# PurgeBuffer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Migration

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x400 (1,024) |

**Description**
Specifies the number of recently deleted objects that cannot yet be permanently erased.

When objects are deleted, they remain in Active Directory until they are erased permanently. Message Queuing (also known as MSMQ) creates a buffer zone of recently deleted objects that cannot be erased. If an object is deleted accidentally, then you can retrieve it if it is still in the buffer zone. The value of this entry determines the number of recently deleted objects in the buffer zone.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Entries in the Migration subkey are used only when there are MSMQ 1.0 servers running Windows NT 4.0 and earlier. They are not used when all servers are running Windows Server 2003 or Windows 2000.

# ReplMsgTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Migration

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *seconds* | 0x4B0 (1,200 seconds = 20 minutes) |

**Description**

Specifies how long replication messages can remain in Active Directory before they are deleted.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- Entries in the Migration subkey are used only when there are MSMQ 1.0 servers running Windows NT 4.0 and earlier. They are not used when all servers are running Windows Server 2003 or Windows 2000.

# ReplThreads

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Migration

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x4 |

**Description**
Specifies the number of threads that Message Queuing (also known as MSMQ) allocates for answering replication and sync messages from servers running Windows NT 4.0.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Entries in the Migration subkey are used only when there are MSMQ 1.0 servers running Windows NT 4.0 and earlier. They are not used when all servers are running Windows Server 2003 or Windows 2000.

# TimesHelloForReplicationInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Migration

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0xFFFFFFFF |

**Description**
Specifies how often Message Queuing (also known as MSMQ) sends replication messages.

The interval between replication messages is based on the length of the interval between *hello* messages, as determined by the value of HelloInterval. The replication interval is expressed as the number of hello-message intervals. The value of this entry specifies the number of *hello* message intervals that must elapse between replication messages.

For example, if the value of **HelloInterval** is 60 seconds, and the value of **TimesHelloForReplicationInterval** is 60, then the interval between replication messages is one hour (60 × 60 seconds).

*Hello* messages are Layer Two Tunneling Protocol (L2TP) control messages sent at regular intervals to prevent an active, but silent, connection from being mistaken for an idle line. They are also known as *keep-alive* messages.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Entries in the Migration subkey are used only when there are MSMQ 1.0 servers running Windows NT 4.0 and earlier. They are not used when all servers are running Windows Server 2003 or Windows 2000.

**See Also**
**Concepts**
HelloInterval

# OCMsetup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

**Description**
The OCMsetup subkey stores information about the installation and configuration of Message Queuing (also known as MSMQ) when it is set up using the Optional Components Manager setup.

**Caution**

- These values are stored by Message Queuing for internal use. Do not add or change any entries or values in this subkey.

# Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\OCMsetup

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Path and directory name* | (None) |

**Description**
Stores the location of the directory where Message Queuing (also known as MSMQ) is installed.

**Caution**

- Do not change the value of this entry. Changing the value can prevent the Setup program from successfully installing, upgrading, or uninstalling Message Queuing.

# InstalledComponents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\OCMsetup

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | (None) |

**Description**

Records the Message Queuing (also known as MSMQ) components that were installed on this computer.

**Caution**

- Do not change the value of this entry. Changing the value can prevent the Setup program from successfully installing, upgrading, or uninstalling Message Queuing.

# ShortcutCount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\OCMsetup

| Data type | Range | Default value |
|---|---|---|
| | 0x0–0xFFFFFFFF | (None) |

**Description**

Records the number of shortcuts installed by Optional Components Manager (OCM) Setup.

**Caution**

- Do not change the value of this entry. Changing the value can prevent the Setup program from successfully installing, upgrading, or uninstalling Message Queuing (also known as MSMQ).

# ShortcutDirectory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\OCMsetup

| Data type | Range | Default value |
|---|---|---|
| | *Path and directory name* | (None) |

**Description**

Records the location of the directory where Message Queuing (also known as MSMQ) shortcuts are installed

**Caution**

- Do not change the value of this entry. Changing the value can prevent the Setup program from successfully installing, upgrading, or uninstalling Message Queuing.

# OnHold Queues

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

**Description**
This subkey contains a list of outgoing queues that are currently paused.

# Parameters\Security

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

**Description**
The Security subkey contains subkeys that contain information about the computer account and about whether server authentication is used (secured communication between client and server).

### Note

- The entries in the subkeys in the HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\CertificationAuthorities subkeys store actual certification authority (CA) certificates, not just pointers to entries in another location.

### Caution

- Do not change the values of entries in this subkey by editing the registry. Changing these values can disable the security features of Message Queuing.

# ServerCertDigest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Security

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | (None) |

**Description**
Specifies which server certificate is used if the certificate store contains several certificates. Message Queuing (also known as MSMQ) will use only the certificate whose hash (digest) matches this value.

The value of this entry is set when the user registers a server certificate through Control Panel. The **ServerCertDigest** entry is used only on a Message Queuing server that runs on a Windows Server 2003 domain controller.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# ServerCertStore

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Security

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *string* | (None) |

**Description**

Specifies the name of the machine certificate store where the server certificate for the local server is stored.

This entry is used only on a Message Queuing (also known as MSMQ) server that runs on a Windows Server 2003 domain controller. The server certificate is used for the Secured Communication feature, where Message Queuing clients can verify that they query an authentic Message Queuing server.

The value of this entry is set when the user registers a server certificate through Control Panel. The value is read by Message Queuing the first time a client wants to authenticate the server.

# ServersCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

**Description**
The ServersCache subkey of the MSMQ\Parameters subkey stores information pertaining to Message Queuing (also known as MSMQ) about each site in the enterprise. ServersCache contains a separate entry for each site in the enterprise. The entry is the name of the site, and its value consists of a list of all the MQIS servers in the site. These values are updated at an interval specified by the value of the **DSListRefresh** entry.

Each of the entries in the ServersCache subkey represents a site and stores the computer names and supported protocols of all Message Queuing servers at the site. These entries have the following format:

| Entry name | Data type | Value |
|---|---|---|
| Site name | REG_SZ | 01 \| 10 \| 11 *NumberComputer name*[,*NumberComputer name*...] |

Each computer name follows a number (01, 10, or 11), which indicates the communication protocols that each server supports. If the site includes multiple servers, the entries must be separated by commas.

- A value of 1 in the lower (right) digit indicates that the server supports the IPX protocol.
- A value of 1 in the higher (left) digit indicates that the server supports the IP protocol.

For example, the following entry indicates that, at the Central site, three computers are running Message Queuing: \\ServerOne supports both IP and IPX, \\ServerTwo supports only IP, and \\ServerThree supports only IPX:

| Entry name | Data type | Value |
|---|---|---|
| **Central** | REG_SZ | 11ServerOne,10ServerTwo,01ServerThree |

# Setup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters

**Description**
The Setup subkey stores entries that contain information about the operating system installed, the name of the computer, name of the local domain, and the name of the local organizational unit (OU).

**Note**

- The entries in the subkeys in CertificationAuthorities store actual certification authority (CA) certificates, not just pointers to entries in another location.

# JoinStatus

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0x4 | (None) |

**Description**
Implements transactional semantics for the join process, so that if a computer stops abnormally, the join process continues after Windows Server 2003 restarts.

Message Queuing (also known as MSMQ) adds this entry to the registry when Message Queuing detects that a computer joined a domain.

| Value | Meaning |
|---|---|
| 1 | MSMQ_JOIN_STATUS_START_JOINING |
| 2 | MSMQ_JOIN_STATUS_JOINED_SUCCESSFULLY |
| 3 | MSMQ_JOIN_STATUS_FAIL_TO_JOIN |
| 4 | MSMQ_JOIN_STATUS_UNKNOWN |

**Note**

- This entry appears by default only on computers installed in workgroup mode. For all other computers, it must be added by using the registry editor Regedit.exe.

# MachineDN

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *String* | (None) |

**Description**
Describes the local computer in the Active Directory.

The value of this entry is the distinguished name of the computer object and is used when the computer moves between domains to check if the Message Queuing (also known as MSMQ) configuration object is still in its old place. This entry is updated on every restart.

# MachineDomain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *String* | (None) |

**Description**

Enables the user to check if the computer has moved between domains.

The value of the **MachineDomain** entry is the NetBIOS name of the domain that owns the computer object. This entry is not erased when you move your computer from a domain to a workgroup.

# MSSQLServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Client Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSSQLServer

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ConnectTo

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\MSSQLServer\Client

| Data type | Range | Default value |
|---|---|---|
| NULL | | |

**Description**
The ConnectTo subkey stores connection information for Windows-based clients that connect to instances of the Microsoft SQL Server database engine. Registry entries in this subkey are of data type REG_SZ, and their values specify the Net-Library to use for connecting to the instance, as well as any network-specific parameters.

The connection information stored in this subkey is read by the OLE DB Provider for SQL Server, the SQL Server ODBC driver, or the SQL Server DB-Library dynamic-link library (DLL).

**Change Method**
You should not modify entries in this subkey by using the registry editor. Use the SQL Server Client Network Utility to configure server name and connection information. See SQL Server Client Network Utility Help for further instructions.

### Notes

- The ConnectTo subkey is used by the OLE DB Provider for SQL Server, the SQL Server ODBC Driver, and the SQL Server DB-Library DLL. Entries in this subkey identify which Net-Library these data access application programming interface (API) components call.

- Net-Libraries are SQL Server components that shield the data access API components from the details of using different Interprocess Communication (IPC) methods to communicate with an instancess of the SQL Server database engine.

### Caution

- Improperly updating the ConnectTo subkey might prevent applications from successfully connecting to an instance of the SQL Server database engine.

# BlockExeAttachments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Outlook Express

**Description**
Stores configuration data for the policy setting **Configure Outlook Express**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# Outlook Express_DepKit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Microsoft\PCHealth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ErrorReporting Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\PCHealth\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ErrorReporting\DW Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\PCHealth\ErrorReporting\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# PswdSync

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

**Description**
The PswdSync subkey stores configuration data for the Password Synchronization Service. The Password Synchronization Service lets users use a single password for computers that are running Windows Server 2003 and UNIX. It ensures that any password changes made in Windows Server 2003 are communicated to the UNIX system, in encrypted text. It also ensures that all passwords conform to the security standards implemented on the system by notification packages, such as Passfilt.dll.

The PswdSync subkey contains two subkeys, Domains Subkey and NotifyChange. The Domains subkey contains *UNIX-system-name* subkeys, each of which represents a UNIX system and stores data that applies to only to that system. The NotifyChange subkey is used internally by the Password Synchronization Service to detect and implement configuration changes while the service is running. Do not delete or change it.

The PswdSync subkey can also contain entries. The entries in PswdSync apply to all UNIX systems. These entries can also appear in the PswdSync\Domains*UNIX-system-name* subkeys, where they apply only to the UNIX systems. If the same entry appears in the PswdSync subkey and in a *UNIX-system-name* subkey, then the value in the *UNIX-system-name* subkey takes precedence over the value in PswdSync when configuring that UNIX system.

**Change Method**
To change the value of the entries in this subkey, use the Password Synchronization wizard in Services for UNIX Administrator. Do not edit the registry.

### Note

- The Pswdsync subkey appears in the registry when you configure the Password Synchronization Service by using the Password Synchronization wizard in Services for UNIX Administrator.

# Domains Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\PswdSync

**Description**
The Domains subkey stores configuration data for the UNIX systems that are associated with the Windows Server 2003 Password Synchronization Service. The Password Synchronization Service lets users use a single password for computers that are running Windows Server 2003 and UNIX.

The Domains subkey contains UNIX-system-name, each of which represents a UNIX system and stores data that applies only to that UNIX system.

**Change Method**
To change the value of the entries in this subkey, use the Password Synchronization wizard in Services for UNIX Administrator. Do not edit the registry.

### Note

- This subkey appears in the registry when you use the Password Synchronization wizard in Services for UNIX Administrator to configure a UNIX system.

# UNIX-system-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\PswdSync\Domains

**Description**
The *UNIX-system-name* subkey is named for a particular UNIX system associated with this Windows Server 2003, and it contains entries that apply only to that UNIX system. There is a *UNIX-system-name* subkey for each UNIX system associated with this Windows Server 2003.

The entries listed in *UNIX-system-name* can appear in any *UNIX-system-name* subkey. Some of the entries can also appear in the PswdSync subkey. However, when they appear in PswdSync, their values apply to all UNIX systems. When the same entry appears in PswdSync and in a *UNIX-system-name subkey*, the value in the *UNIX-system-name* subkey takes precedence over the value in the PswdSync subkey when configuring the named UNIX system.

**Change Method**
To change the value of entries in this subkey, use the Password Synchronization wizard in Services for UNIX Administrator. Do not edit the registry.

### Notes

- This subkey appears in the registry when you use the Password Synchronization wizard to configure a UNIX system.
- The *UNIX-system-name* subkey is a variable representing the subkeys in the Domains subkeys. It does not actually appear in the registry. This variable subkey displays the entries that can appear in any of the subkeys within the Domains Subkey subkey. The entries that actually appear in the registry depend upon the configuration of the UNIX system.

# LoginAccount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\PswdSync\Domains*UNIX-system-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Account-name* | root |

**Description**
Stores the account name used when using Rlogin to log on to a remote UNIX computer to synchronize passwords. Rlogin sends passwords in clear text.

This entry appears in the registry only when the system is running Services for Unix 1.0 and the **Unsecure using Rlogin** password propagation method is selected—that is, when the value of **PropagateMethod** is 0. Services for Unix 2.0 and Services for Unix 3.0 do not use this entry.

**Change Method**
To change the value of this entry, use the Password Synchronization Administrator. Do not edit the registry.

### Note

- This entry does not appear in the registry unless you install Services for Unix 1.0.

# SystemMonitor

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DisplaySingleLogSampleValue

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\SystemMonitor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | (None) |

**Description**
Specifies that System Monitor display only the average value of each group of data points. When System Monitor needs to display more than 100 data points, it groups the values into 100-point data buffers. System Monitor then calculates the minimum, average, and maximum value of each group and displays them as a vertical line. If this entry is present in the registry and its value is 1, System Monitor instead displays only the average value for each group of data points.

System Monitor ignores this entry when its data source is real-time data; it uses this entry when the data source is an SQL database or one or more log files.

**Change Method**
This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must stop and restart System Monitor for changes to take effect.

# DisplayThousandsSeparator

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\SystemMonitor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | (None) |

**Description**
Specifies whether System Monitor groups the digits of the values that it displays. For example, 50000000 with digits grouping is displayed as 50,000,000. If this entry exists in the registry and its value is 1, then System Monitor uses digit grouping. If this entry is absent from the registry or its value is zero, then System Monitor does not use digit grouping.

| Value | Meaning |
|---|---|
| 0 | System Montor does not group digits. |
| 1 | System Monitor uses digit grouping. |

**Change Method**
This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must stop and restart System Monitor for changes to take effect.

# Microsoft\Windows Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

**Description**
The Windows subkey stores configuration data for the programs running on the Win32 subsystem. The subkeys of the Windows subkey represent versions of the Win32 subsystem, such as CurrentVersion. Each version subkey contains subkeys representing programs that run on that version of the subsystem.

The Windows subkey in HKEY_LOCAL_MACHINE stores entries that apply to all users of the computer. There is also a SOFTWARE\Microsoft\Windows subkey in HKEY_CURRENT_USER that stores entries that apply only to the current user of the computer.

**See Also**
**Concepts**
Microsoft\Windows

# Windows\CurrentVersion

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows

**Description**
The CurrentVersion subkey stores configuration data for programs running on the Win32 subsystem. The subkeys of CurrentVersion represent programs that run on the current version of the Win32 subsystem, and they store configuration data for those programs that apply to all users of the computer.

# CurrentVersion\Explorer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion

**Description**
The Explorer subkey stores configuration data for Windows Explorer and for the Windows interface.

# Explorer\Shell Folders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Note

- Paths to the Windows Explorer folders are stored in the Explorer\User Shell Folders subkey. For more information about shell folders, see the Software Development Kit (SDK) information in the MSDN Library link on the Web Resources page.

### Caution

- Do not delete this subkey or its entries from the registry, and do not change the values in the entries. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
Explorer\User Shell Folders

# Explorer\User Shell Folders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer

**Description**
The User Shell Folders subkey of the HKEY_LOCAL_MACHINE subtree stores the paths to Windows Explorer folders for the current user of the computer.

Windows uses the value of this entry internally. To retrieve this value, use the SHGetFolderPath() application programming interface (API).

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
User Shell Folders Subkey

# Common Desktop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %ALLUSERPROFILE%\Desktop |

**Description**
Windows uses the value of this entry internally. To retrieve this value, use the SHGetFolderPath() application programming interface (API).

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
User Shell Folders\Desktop

# Common Documents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %ALLUSERPROFILE%\Documents |

**Description**
Windows uses the value of this entry internally. To retrieve this value, use the SHGetFolderPath() application programming interface (API).

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
Common Documents

# Common Favorites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %ALLUSERPROFILE%\Favorites |

**Description**
Windows Server 2003 uses the value of this entry internally. You should not modify it. To retrieve this value, use the SHGetFolderPath() application programming interface (API).

# Common Programs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %ALLUSERPROFILE%\Start Menu\Programs |

**Description**
Windows uses the value of this entry internally. To retrieve this value, use the SHGetFolderPath() application programming interface (API).

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
User Shell Folders\Programs

# Common Start Menu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %ALLUSERPROFILE%\Start Menu |

**Description**
Windows uses the value of this entry internally. To retrieve this value, use the SHGetFolderPath() application programming interface (API).

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# User Shell Folders\Personal

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Path* | *Systemroot*\My Documents |

**Description**
Windows uses the value of this entry internally. To retrieve this value, use the SHGetFolderPath() application programming interface (API).

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
User Shell Folders\Personal Entry

# Common Startup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %ALLUSERPROFILE%\Start Menu\Programs\Startup |

**Description**
Windows uses the value of this entry internally. To retrieve this value, use the SHGetFolderPath() application programming interface (API).

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
Startup

# CurrentVersion\Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ActiveDesktop Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# NoAddingComponents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit adding items**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# NoChangingWallPaper

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prevent changing wallpaper**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# NoClosingComponents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit closing items**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# NoComponents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Disable all items**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# NoDeletingComponents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit deleting items**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# NoDispAppearancePage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Appearance and Themes tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# NoDispScrSavPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Screen Saver tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# NoDispSettingsPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Settings tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# NoEditingComponents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit editing items**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# NoHTMLWallPaper

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Allow only bitmapped wallpaper**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# Add

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop\AdminComponent

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Add/Delete items**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# Delete

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop\AdminComponent

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Add/Delete items**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# ActiveDesktop\AdminComponent

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\ActiveDesktop

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# comdlg32

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Placesbar

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\comdlg32

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# NoBackButton

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\comdlg32\Placesbar

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide the common dialog back button**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer\Common Open File Dialog.

# NoFileMru

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\comdlg32\Placesbar

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide the dropdown list of recent files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer\Common Open File Dialog.

# NoPlacesBar

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\comdlg32\Placesbar

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide the common dialog places bar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer\Common Open File Dialog.

# Place0

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\comdlg32\Placesbar

**Description**
Stores configuration data for the policy setting **Items displayed in Places Bar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer\Common Open File Dialog.

# Place1

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\comdlg32\Placesbar

**Description**
Stores configuration data for the policy setting **Items displayed in Places Bar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer\Common Open File Dialog.

# Place2

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\comdlg32\Placesbar

**Description**
Stores configuration data for the policy setting **Items displayed in Places Bar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer\Common Open File Dialog.

# Place3

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\comdlg32\Placesbar

**Description**
Stores configuration data for the policy setting **Items displayed in Places Bar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer\Common Open File Dialog.

# Place4

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\comdlg32\Placesbar

**Description**
Stores configuration data for the policy setting **Items displayed in Places Bar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer\Common Open File Dialog.

# Btn_Back

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Copy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Cut

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Discussions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Edit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Encoding

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Favorites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Folders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Forward

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Fullscreen

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Home

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_MailNews

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Media

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Paste

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Print

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Refresh

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Search

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Size

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Stop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# Btn_Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# ClassicShell

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Turn on Classic Shell**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# ClearRecentDocsOnExit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Clear history of recently opened documents on exit**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# DisableLocalMachineRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not process the legacy run list**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Logon.

# DisableLocalMachineRunOnce

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not process the run once list**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Logon.

# DisableLocalUserRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Do not process the legacy run list**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Logon.

# DisableLocalUserRunOnce

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Do not process the run once list**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Logon.

# DisablePersonalDirChange

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit user from changing My Documents path**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# DisallowCpl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide specified Control Panel applets**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel.

# DisallowRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Don't run specified Windows applications**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# EnforceShellExtensionSecurity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Allow only per user or approved shell extensions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# ForceActiveDesktopOn

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enable Active Desktop**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# ForceStartMenuLogOff

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Add Logoff to the Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# GreyMSIAds

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Gray unavailable Windows Installer programs Start Menu shortcuts**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# HideClock

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove Clock from the system notification area**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# Explorer\Intellimenus

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Turn off personalized menus**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# LinkResolveIgnoreLinkInfo

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not track Shell shortcuts during roaming**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# LockTaskbar

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Lock the Taskbar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# MemCheckBoxInRunDlg

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Add Run in Separate Memory Space check box to Run dialog box**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoActiveDesktop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Disable Active Desktop**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# NoActiveDesktopChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit changes**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# NoAutoTrayNotify

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Turn off notification area cleanup**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoBandCustomize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Disable customizing browser toolbars**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# NoChangeAnimation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove UI to change menu animation setting**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoChangeKeyboardNavigationIndicators

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove UI to change keyboard navigation indicator setting**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoChangeStartMenu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Drag-and-drop context menus on the Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# Explorer\NoClose

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove and prevent access to the Shut Down command**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoCloseDragDropBands

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prevent adding, dragging, dropping and closing the Taskbar's toolbars**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoCommonGroups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove common program groups from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoComputersNearMe

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **No Computers Near Me in My Network Places**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoControlPanel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit access to the Control Panel**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel.

# NoDesktop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide and disable all items on the desktop**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoDesktopCleanupWizard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove the Desktop Cleanup Wizard**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoDisconnect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Disconnect option from Shut Down dialog**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# NoDrives

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide these specified drives in My Computer**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoDriveTypeAutoRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Turn off Autoplay**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# NoEncryptOnMove

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not automatically encrypt files moved to encrypted folders**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# NoFavoritesMenu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Favorites menu from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# SpecifyDefaultButtons

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# StartMenuLogOff

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Logoff on the Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# Explorer\NoFileMenu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove File menu from Windows Explorer**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoFind

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Search menu from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoFolderOptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Removes the Folder Options menu item from the Tools menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoHardwareTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Hardware tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoInstrumentation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Turn off user tracking**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoInternetIcon

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Internet Explorer icon on desktop**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoLogoff

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Logoff**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Ctrl+Alt+Del Options.

# NoManageMyComputerVerb

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hides the Manage item on the Windows Explorer context menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoMovingBands

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit adjusting desktop toolbars**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoNetConnectDisconnect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Map Network Drive and Disconnect Network Drive**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoNetHood

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide My Network Places icon on desktop**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoNetworkConnections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Network Connections from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoNTSecurity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Windows Security item from Start menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# NoPropertiesMyComputer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove Properties from the My Computer context menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoPropertiesMyDocuments

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove Properties from the My Documents context menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoPropertiesRecycleBin

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove Properties from the Recycle Bin context menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoRecentDocsHistory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not keep history of recently opened documents**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoRecentDocsMenu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Documents menu from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoRecentDocsNetHood

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not add shares of recently opened documents to My Network Places**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoResolveSearch

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not use the search-based method when resolving shell shortcuts**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoResolveTrack

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not use the tracking-based method when resolving shell shortcuts**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# Explorer\NoRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Run menu from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# Explorer\NoSaveSettings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Don't save settings at exit**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# NoSecurityTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove Security tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoSetFolders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove programs on Settings menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoSetTaskbar

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prevent changes to Taskbar and Start Menu Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoShellSearchButton

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Search button from Windows Explorer**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoSimpleStartMenu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Force classic Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoSMBalloonTip

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove Balloon Tips on Start Menu items**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoSMConfigurePrograms

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove Set Program Access and Defaults from Start menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoSMHelp

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Help menu from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoSMMyDocs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove My Documents icon from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoSMMyPictures

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove My Pictures icon from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoStartMenuEjectPC

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove the Undock PC button from the Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoStartMenuMFUprogramsList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove frequent programs list from the Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoStartMenuMorePrograms

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove All Programs list from the Start menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoStartMenuNetworkPlaces

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove My Network Places icon from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoStartMenuPinnedList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove pinned programs list from the Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoStartMenuSubFolders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove user's folders from the Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoTaskGrouping

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Prevent grouping of taskbar items**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoThemesTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove Theme option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display\Desktop Themes.

# NoToolbarCustomize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Disable customizing browser toolbar buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# NoToolbarsOnTaskbar

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Do not display any custom toolbars in the taskbar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoTrayContextMenu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove access to the context menus for the taskbar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoTrayItemsDisplay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Hide the notification area**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoUserNameInStartMenu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove user name from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# NoViewContextMenu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Windows Explorer's default context menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoViewOnDrive

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prevent access to drives from My Computer**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoWelcomeScreen

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Don't display the Getting Started welcome screen at logon**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# NoWindowsUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove links and access to Windows Update**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# RestrictCpl Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Show only specified Control Panel applets**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel.

# Policies\Explorer Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies

**Description**
The Explorer subkey contains entries that store Group Policy settings for Windows Explorer and for the Windows interface. These include policies for the Start menu, taskbar, desktop, and Control Panel. Group Policy creates and maintains the entries in this subkey, and Windows Explorer reads and interprets them.

This subkey in HKEY_LOCAL_MACHINE stores entries that apply to all users of this computer. The Software\Microsoft\Windows\CurrentVersion\Policies\Explorer subkey in HKEY_CURRENT_USER stores policy-related entries that are configured separately for each user.

**See Also**
**Concepts**
Policies\Explorer

# Btn_History

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Configure Toolbar Buttons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Toolbars.

# NoDFSTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove DFS tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoStartMenuMyMusic

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
Stores configuration data for the policy setting **Remove My Music icon from Start Menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Start Menu and Taskbar.

# RestrictRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Run only allowed Windows applications**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# NoFileAssociate Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 (Terminal Services systems) 0 (All other Windows Server 2003 systems) |

**Description**
Disables the buttons on the **File Types** tab. As a result, users can view file type associations, but they cannot add, delete, or change them.

The **File Types** tab shows which programs open when a type of file is double-clicked. The file types are represented by their file name extension.

| Value | Meaning |
|---|---|
| 0 | Users can add, delete, and change file type associations. |
| 1 | Users cannot add, delete, or change file type associations. |

By disabling this tab, you might prevent malicious users from creating an association that starts a risky or destructive program. Also, users cannot use the buttons on the **File Types** tab to create less productive associations, such as associating the .doc file type with Wordpad instead of Microsoft Word. This is a particular risk on multi-user systems, such as Terminal Services clients.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Notes

- Although this entry is located in the Policies\Explorer subkey, it is not associated with a Group Policy setting. You can create a policy by adding this entry to an Administrative Templates file. For more information on creating custom .adm files, see Windows Server 2003 Help or Windows XP Professional Help.

- This entry applies to all users of the computer. To set this entry for particular users, use Explorer\NoFileAssociate in the HKCU subtree. If the settings of these entries conflict, the setting in HKCU takes precedence over the setting in HKLM.

### Caution

- This entry specifies only whether buttons on the **File Types** tab in **Folder Options** are available. Any other method of changing file type associations is unaffected by this entry.

**See Also**
**Concepts**
Explorer\NoFileAssociate

# ForceClassicControlPanel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer\RestrictCpl

**Description**
Stores configuration data for the policy setting **Force classic Control Panel Style**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel.

# Explorer\RestrictCpl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# n Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer\Run

**Description**
Stores configuration data for the policy setting **Run these programs at user logon**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Logon.

# Run Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Explorer

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ConfirmFileDelete

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

**Description**
Stores configuration data for the policy setting **Display confirmation dialog when deleting files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# MaxRecentDocs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Maximum number of recent documents**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoCDBurning

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

**Description**
Stores configuration data for the policy setting **Remove CD Burning features**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoEntireNetwork

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **No Entire Network in My Network Places**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoOnlinePrintsWizard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

**Description**
Stores configuration data for the policy setting **Remove Order Prints from Picture Tasks**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoPublishingWizard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

**Description**
Stores configuration data for the policy setting **Remove Publish to Web from File and Folder Tasks**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoRecycleFiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

**Description**
Stores configuration data for the policy setting **Do not move deleted files to the Recycle Bin**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoRunasInstallPrompt

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not request alternate credentials**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoThumbnailCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

**Description**
Stores configuration data for the policy setting **Turn off caching of thumbnail pictures**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoWebServices

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

**Description**
Stores configuration data for the policy setting **Prevent Internet download for Web Publishing and Online Ordering wizards**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# NoWinKeys

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

**Description**
Stores configuration data for the policy setting **Turn off Windows+X hotkeys**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# PromptRunasInstallNetPath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Request credentials for network installations**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# RecycleBinSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Network

**Description**
Stores configuration data for the policy setting **Maximum allowed Recycle Bin size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Explorer.

# Policies\Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# NonEnum

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {450D8FBA-AD25-11D0-98A8-0800361B1103}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\NonEnum

**Description**
Stores configuration data for the policy setting **Remove My Documents icon on the desktop**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# {645FF040-5081-101B-9F08-00AA002F954E}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\NonEnum

**Description**
Stores configuration data for the policy setting **Remove Recycle Bin icon from desktop**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop.

# ConnectHomeDirToRoot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Connect home directory to root of the share**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# DisableBkGndGroupPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Turn off background refresh of Group Policy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# DisableChangePassword

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Change Password**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Ctrl+Alt+Del Options.

# DisableLockWorkstation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Lock Computer**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Ctrl+Alt+Del Options.

# DisableRegistryTools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prevent access to registry editing tools**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# DisableTaskMgr

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Task Manager**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Ctrl+Alt+Del Options.

# EnableProfileQuota

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Limit profile size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# HideLegacyLogonScripts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Run legacy logon scripts hidden**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Scripts.

# HideLogoffScripts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Run logoff scripts visible**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Scripts.

# HideLogonScripts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Run logon scripts visible**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Scripts.

# HideShutdownScripts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Run shutdown scripts visible**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Scripts.

# HideStartupScripts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Run startup scripts visible**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Scripts.

# IncludeRegInProQuota

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Limit profile size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# LogonType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

**Description**
Stores configuration data for the policy setting **Always use classic logon**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Logon.

# MaxGPOScriptWait

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Maximum wait time for Group Policy scripts**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Scripts.

# MaxProfileSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Limit profile size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# NoColorChoice

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

**Description**
Stores configuration data for the policy setting **Prohibit Theme color selection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display\Desktop Themes.

# NoDispBackgroundPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Desktop tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# NoDispCPL

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Display in Control Panel**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# NoSizeChoice

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

**Description**
Stores configuration data for the policy setting **Prohibit selection of font size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display\Desktop Themes.

# NoVisualStyleChoice

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

**Description**
Stores configuration data for the policy setting **Prevent selection of windows and buttons styles**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display\Desktop Themes.

# ProfileQuotaMessage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Limit profile size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# RunLogonScriptSync Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Run logon scripts synchronously**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Scripts.

# RunStartupScriptSync

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Run startup scripts asynchronously**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Scripts.

# SetVisualStyle

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

**Description**
Stores configuration data for the policy setting **Load a specific visual style file or force Windows Classic**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display\Desktop Themes.

# System\Shell

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Custom user interface**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# VerboseStatus

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Verbose vs normal status messages**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# Wallpaper

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Active Desktop Wallpaper**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# WallpaperStyle

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Active Desktop Wallpaper**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Desktop.

# WarnUser

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Limit profile size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# WarnUserTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Limit profile size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# Policies\System

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies

**Description**
The System subkey stores the entries created when you configure a Group Policy setting that affects a basic component of Windows Server 2003. Group Policy creates and maintains the entries in this subkey, and the component program reads and interprets them.

This subkey stores entries that apply to all users of this computer. The Software\Microsoft\Windows\CurrentVersion\Policies\System subkey in HKEY_CURRENT_USER stores policy setting related entries that are configured separately for each user.

**See Also**
**Concepts**
Microsoft\System Subkey
Windows\System Subkey

# DisableStatusMessages

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Boot / Shutdown / Logon / Logoff status messages**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# Uninstall

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DefaultCategory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Specify default category for Add New Programs**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoAddFromCDorFloppy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide the Add a program from CD-ROM or floppy disk option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoAddFromInternet

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide the Add programs from Microsoft option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoAddFromNetwork

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide the Add programs from your network option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoAddPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Add New Programs page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoAddRemovePrograms

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Add or Remove Programs**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoChooseProgramsPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

**Description**
Stores configuration data for the policy setting **Hide the Set Program Access and Defaults page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoRemovePage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Change or Remove Programs page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoServices

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Go directly to Components Wizard**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoSupportInfo

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove Support Information**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# NoWindowsSetupPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Uninstall

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Add/Remove Windows Components page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Add or Remove Programs.

# DisableWindowsUpdateAccess

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\WindowsUpdate

**Description**
Stores configuration data for the policy setting **Remove access to use all Windows Update features**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# WindowsUpdate Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# CurrentVersion\Run

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion

**Description**
The Run subkey of HKEY_LOCAL_MACHINE\Software stores the names of programs that Windows Server 2003 runs for all users of the computer each time the system starts up.

Each entry in the subkey represents a program. The value of the entry is the name of the executable file for the program. The entries have the following format:

| Entry name | Data type | Value |
|---|---|---|
| *Program Name* | REG_SZ | *File name* |

# RunOnce

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion

**Description**

The RunOnce subkey stores the names of programs that will be run the next time that a user logs on. These programs are run only one time, and then the registry entries that specified them are deleted from the subkey.

However, if the account of the next user to log on does not have permissions to delete entries from the RunOnce subkey, then the specified programs will not be run. The operating system will not run the programs until a user with adequate permissions logs on.

Each entry in the subkey represents a program. The value of the entry is the name of the executable file for the program. These entries have the following format:

| Entry name | Data type | Value |
|---|---|---|
| *Program Name* | REG_SZ | *File name* |

# Telephony Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Telephony\Server

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Telephony

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# HighSecurity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Telephony\Server

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFF | (None) |

**Description**
Specifies whether High Security mode is enabled on a Telephony Server. In High Security mode, the Telephony Server will accept RPC connections only from remote clients connecting with RPC_C_AUTHN_LEVEL_PKT_PRIVACY. Client connections that do not use RPC_C_AUTHN_LEVEL_PKT_PRIVACY will be rejected. Enabling High Security mode enhances the security of the Telephony Server component.

| Value | Meaning |
|---|---|
| 0x0 | High Security mode is disabled. |
| 0x1–0xFFFFFFFF (or not present in the registry) | High Security mode is enabled. |

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
You must stop and restart the Telephony service on the Telephony Server to make changes to this entry effective.

### Note

- Starting with Windows Server 2003 and Windows XP Professional Service Pack 2, the telephony clients connect to a remote Telephony Server with RPC_C_AUTHN_LEVEL_PKT_PRIVACY. Clients using earlier Windows operating systems do not use RPC_C_AUTHN_LEVEL_PKT_PRIVACY and will fail to connect to a Windows Server 2003 Telephony Server in High Security Mode, which is enabled by default. Disabling High Security mode on the Telephony Server ensures that clients running earlier versions of Windows can connect to the server.

### Caution

Setting the value of this entry to 0x0 disables High Security mode on the Telephony Server, which compromises the security of the Telephony Server by allowing connections from clients that do not use RPC_C_AUTHN_LEVEL_PKT_PRIVACY (the RPC parameters for these connections are not encrypted).

# Microsoft\Windows NT Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

**Description**
The Windows NT subkey stores configuration data for components of Windows Server 2003. The subkeys of the Windows NT subkey represent versions of the operating system, such as CurrentVersion, which refers to the version currently running on the computer. Each version subkey contains subkeys representing components that run on that version of the system.

This subkey stores entries that apply to all users of the computer. The Windows NT subkey in HKEY_CURRENT_USER stores entries that apply only to the current user of the computer.

**See Also**
**Concepts**
Microsoft\Windows NT

# CurrentVersion Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Version number* | 5.2 |

**Description**

Stores the version number of the Windows NT operating system that is currently running on the system.

# AeDebug

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
The AeDebug subkey stores user preferences for Dr. Watson and other debugging tools.

**Change Method**
Use the **Dr. Watson for Windows** dialog box to change the values of entries in this subkey. (Run drwtsn32 from the command prompt or use the Run command on the Start menu.) To disable Dr. Watson for Windows, delete the AeDebug subkey. To enable Dr. Watson for Windows, type **drwtsn -i** at the command prompt or from the Run command on the Start menu.

# Auto

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\AeDebug

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 1 |

**Description**
Specifies whether the debugger specified in the value of the Debugger entry starts automatically when an application error occurs.

| Value | Meaning |
|---|---|
| 0 | The system displays a message box notifying the user when an application error occurs. |
| 1 | The debugger starts automatically. |

If the **Debugger** entry contains the name of a valid debugger (such as Drwtsn32, WinDbg, or NTSD), then the user can either terminate the application by clicking **OK** or can run the debugger by clicking **Cancel**. If no debugger is specified, then the message box has no **Cancel** button.

# Debugger

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\AeDebug

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Name of debugger and switch settings.* | drwtsn32 -p %ld -e %ld -g |

**Description**

Specifies the debugger that activates when an application error occurs. If the value of the Auto entry is 1, then this debugger starts automatically. If the value of **Auto** is 0, then the user must start the debugger.

Dr. Watson for Windows (drwtsn32) is the default debugger, but another debugger (such as WinDbg or NTSD) can be installed as the automatic debugger. The values following the debugger name are specific to the debugger. The default switches for Dr. Watson are set as follows:

| Value | Meaning |
|---|---|
| -p | Stores the Process ID for the running process at the location specified by **%ld.** |
| -e | Stores the Event ID for the running process at the location specified by **%ld.** |
| -g | Tells the debugger to go (start). |

**Change Method**

You can change the value of this entry only by editing the registry. However, you can restore default settings by typing **drwtsn32 -i** at the command prompt.

### Note

- For more information on Dr. Watson for Windows, see Windows Server 2003 Help or Windows XP Professional Help.

# UserDebuggerHotKey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\AeDebug

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFF | 0x0 |

**Description**

Specifies the key that, when pressed, establishes a breakpoint in code being debugged.

The debugger interrupts code processing at the breakpoint so the programmer can examine a suspected problem.

The key specified in this value only sets a breakpoint. It does not invoke the debugger (the debugger must be running before the key is pressed), and it does not switch the debugger to single-step mode.

The value of this entry is a keyboard scan code. The default value (0x0) represents the F12 key on a 101-key keyboard or the "-" (hyphen, VK_SUBTRACT) key on an 82-key keyboard.

The following table lists the scan codes and virtual key codes for a standard keyboard.

| Scan code | Virtual key code |
|---|---|
| 0x0 | VK_SUBTRACT (F12) |
| 0x1 | VK_LBUTTON |
| 0x2 | VK_RBUTTON |
| 0x3 | VK_CANCEL |
| 0x4 | VK_MBUTTON |
| 0x8 | VK_BACK |
| 0x9 | VK_TAB |
| 0xC | VK_CLEAR |
| 0xD | VK_RETURN |
| 0x10 | VK_SHIFT |
| 0x11 | VK_CONTROL |
| 0x12 | VK_MENU |
| 0x13 | VK_PAUSE |
| 0x14 | VK_CAPITAL |
| 0x15 | VK_KANA, VK_HANGEUL, VK_HANGUL |
| 0x17 | VK_JUNJA |
| 0x18 | VK_FINAL |
| 0x19 | VK_HANJA, VK_KANJI |

| | |
|---|---|
| 0x1B | VK_ESCAPE |
| 0x1C | VK_CONVERT |
| 0x1D | VK_NONCONVERT |
| 0x1E | VK_ACCEPT |
| 0x1F | VK_MODECHANGE |
| 0x20 | VK_SPACE |
| 0x21 | VK_PRIOR |
| 0x22 | VK_NEXT |
| 0x23 | VK_END |
| 0x24 | VK_HOME |
| 0x25 | VK_LEFT |
| 0x26 | VK_UP |
| 0x27 | VK_RIGHT |
| 0x28 | VK_DOWN |
| 0x29 | VK_SELECT |
| 0x2A | VK_PRINT |
| 0x2B | VK_EXECUTE |
| 0x2C | VK_SNAPSHOT |
| 0x2D | VK_INSERT |
| 0x2E | VK_DELETE |
| 0x2F | VK_HELP |
| 0x30–0x39 | VK_0–VK_9 (ASCII 0–9) |
| 0x41–0x5A | VK_A–VK_Z (ASCII A–Z) |

# DNS Server_DepKit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Zones

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\DNS Server

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Zone-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\DNS Server\Zones

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
The *Zone-name* subkeys represent each of the authoritative zones on the server. They are named for the zone and store settings that apply only to the zone.

The entries in *Zone-name* can appear in any subkey in the HKLM\SYSTEM\CurrentControlSet\Services\DNS\Zones subkey.

**Activation Method**
The Domain Name System (DNS) reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console. If you change entries by editing the registry, the changes are not effective until you restart the DNS server, but changes made by using the DNS console are effective immediately.

### Notes

- The subkey *Zone-name* is a placeholder representing the zone subkeys in the subkey HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\DNS Server\Zones. "Zone-name" does not actually appear in the registry. Also, the entries documented in this reference for the *Zone-name* subkeys can appear in any registry, but the actual entries in your registry depend on the configuration of the zone.

- Note that in Windows 2000, Zone subkeys are located in HKLM\SYSTEM\CurrentControlSet\Services\DNS\Zones. Only zone-specific entries are relocated for Windows Server 2003; server-specific entries remain in the same location (HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters).

# LocalMasterServers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\DNS Server\Zones*Zone-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Contains the space-separated list of master servers to be used by this DNS server for the zone specified by this registry subkey. When the DNS server loads a stub zone, it checks the value of this entry for that zone. If this entry is empty or doesn't exist in the registry, then the server uses the master servers specified in Active Directory.

This entry is created when **Use local list of Master servers** is selected in the DNS snap-in of the Microsoft Management Console (MMC). When created, this entry is populated with the IP addresses of the master servers. When **Use local list of Master servers** is deselected, this entry is deleted from the registry.

**Change Method**
This entry is created and modified through the DNS snap-in of the Microsoft Management Console.

**Activation Method**
When this entry is modified through the DNS snap-in, changes take effect immediately.

### Note

- A stub zone is a type of DNS zone that contains the set of resource records necessary to enable the DNS server that is hosting the zone to send DNS queries for the zone directly to the DNS servers authoritative for the zone. Without the resource records of the stub zone, if the required resource records are not already in its cache, then the DNS server must perform an interative query for the records, starting with the name servers that are authoritative for the root zone.

# ScavengeServers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\DNS Server\Zones*Zone-name*

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *IP address[, IP address] ... (dotted decimal form)* | (None) |

**Description**
Limits the Domain Name System (DNS) servers that can scavenge the resource records for this zone to the servers listed in this entry.

When *scavenging*, the DNS server inspects the timestamps of resource records for this zone in the DNS database, and it deletes records that are out of date.

To enable scavenging, the value of the server's **Aging** entry must be set to 1.

By default, this entry does not appear in the registry and, therefore, all primary servers that are authoritative for this zone can scavenge its resource records. However, if this entry appears in the registry, only the servers specified in the value of this entry can scavenge the resource records for the zone.

**Change Method**
To change the value of this entry, use the *DNS Server Troubleshooting Tool* (Dnscmd.exe) included with Windows Server 2003. For more information, see Windows Server 2003 Tools Help.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

### Notes

- By default, scavenging is disabled on the DNS server and on all zones. To enable scavenging for any zone, you must enable it on the DNS server (as specified in **ScavengingInterval**) and on the zone (as specified by this entry). If scavenging is disabled on the DNS server, all entries that configure scavenging on any zone are ignored.

- If this entry appears in the registry, but its value is blank, no servers are authorized to scavenge the records of this zone. As a result, scavenging is effectively disabled on the zone.

- DNS adds this entry to the registry when you configure scavenging using the *DNS Server Troubleshooting Tool* (Dnscmd.exe) in the \Support\Tools\Support folder of the Windows Server 2003 operating system CD. You should not edit the registry to add, delete, or change this value.

# MasterServers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\DNS Server\Zones*Zone-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *IP address[, IP address] ... (dotted decimal form)* | (None) |

**Description**
Specifies the Domain Name System (DNS) master servers from which the DNS secondary server pulls zone transfers. This entry is required for secondary zones.

**Change Method**
To change the value of this entry, use the DNS console. Right-click the name of a secondary zone, click **Properties**, and then click the **General** tab. This value corresponds to the **Masters IP addresses** boxes. You should not change this value by editing the registry directly.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS snap-in. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

### Note

- This entry is effective only the when it appears in the registry of a secondary Domain Name System (DNS) zone. Otherwise, it is ignored.

# Font Drivers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
The Font Drivers subkey stores configuration data for external font drivers.

# driver-name Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Font Drivers

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *file name* | *file name* |

**Description**
Stores a list of external font drivers installed on the system.

The *driver-name* entry is a variable representing the entries in the Font Drivers subkey. It does not actually appear in the registry. This variable entry displays information that is common to the entries in the Font Drivers subkey.

The Font Drivers subkey contains an entry for each font driver installed in Windows Server 2003. Each entry has the following format:

| Entry name | Data type | Value |
|---|---|---|
| *Driver name* | REG_SZ | *File name* |

For example, the following entry indicates that the driver for Adobe Type Manager is in the file atmfd.dll. The driver files are stored in *Systemroot*\WINNT\system32:

| Entry name | Data type | Value |
|---|---|---|
| Adobe Type Manager | REG_SZ | atmfd.dll |

**Change Method**
To change the value of this entry, use the program supplied by the font vendor for adding and removing drivers.

### Note

- The bitmap, vector, and TrueType drivers are built in to Windows Server 2003 and do not appear in the value of this entry. The Adobe Type Manager driver for PostScript fonts, which is included with Windows Server 2003, does appear in this entry.

# FontDPI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
The FontDPI subkey stores information about the size of the system fonts.

# LogPixels

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\FontDPI

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x13–0x1E0 *pixels per inch* | 96 |

**Description**
Specifies the size of fonts on the desktop and in menus. This entry also determines the bitmap resolution of icons and toolbar buttons.

**Change Method**
To change the value of this entry, double-click **Display** in Control Panel. Click the **Settings** tab, click the **Advanced** button, and on the **General** tab, select a font size from the list.

**Activation Method**
You must restart the computer for changes to this entry to take effect.

### Note

- The actual range of this entry depends on the installed monitor.

# Fonts Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
The Fonts subkey stores the names of the files that support each font installed in Windows Server 2003. This information is used by applications running on Windows Server 2003.

Entries in this subkey are of data type REG_SZ and are named after the font stored in the file name specified by the entry. For example, the entry named **Arial (TrueType)**, if it exists in the registry, stores the string "ARIAL.TTF", the name of the file that has the Arial font. Entries can also store a combination of path and file name.

**See Also**
**Concepts**
Type 1 Fonts

# FontSubstitutes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
Stores the names of alternative fonts that can be used when a requested font or a font that appears in a document either is not available or is unsupported.

Entries in this subkey are of data type REG_SZ and are named after the fonts for which substitutes are defined. The value of an entry is the name of the font to be used in substitution for the font name matched by the name of the entry. For example, the entry **Helvetica**, if it exists in the registry, could store the string "Arial", which would configure Arial as a substitute for Helvetica.

# Font-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\FontSubstitutes

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *font name* | *substitute font* |

**Description**

Associates an alternative font with the actual font that a program requests or a document uses.

The substitute font can be used as the screen font, as the printer font, or as both. These entries have the following format:

| Entry name | Data type | Value |
|---|---|---|
| *Font name* | REG_SZ | *Substitute font name*[,*character set number*] |

For example, the following entry indicates that Arial can be substituted for Helvetica, if Helvetica is not available:

| Entry name | Data type | Value |
|---|---|---|
| Helvetica | REG_SZ | Arial |

The optional character set number must be one of the numbers in the following table.

| Character Set Number | Character Set Name |
|---|---|
| 0 | ANSI_CHARSET |
| 1 | DEFAULT_CHARSET |
| 2 | SYMBOL_CHARSET |
| 128 | SHIFTJIS_CHARSET |
| 129 | HANGEUL_CHARSET |
| 134 | GB2313_CHARSET |
| 136 | CHINESEBIG5_CHARSET |
| 161 | GREEK_CHARSET |
| 162 | TURKISH_CHARSET |
| 177 | HEBREW_CHARSET |
| 178 | ARABIC _CHARSET |
| 186 | BALTIC_CHARSET |
| 204 | RUSSIAN_CHARSET |
| 238 | EE_CHARSET |
| 255 | OEM_CHARSET |

**Note**

- The *Font-name* entry is a variable representing the entries in the FontSubstitutes subkey. It displays information that is

common to the entries in the FontSubstitutes subkey. It does not actually appear in the registry.

# GRE_Initialize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
The GRE_Initialize subkey holds screen raster fonts, which command windows use.

The fonts are stored in the *Systemroot*\WINNT\system32\dllcache folder, and they are hidden to prevent their accidental deletion. All the entries in this subkey have the following format:

| Entry name | Data type | Value |
|---|---|---|
| *Font name* | REG_SZ | *File name* |

**Change Method**
These fonts are used by the system. Do not change the values of these entries.

### Note

- For more information on raster fonts in Windows Server 2003, see the *Windows Server 2003 Technical Reference*.

# CurrentVersion\Network Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support administration tools designed for Windows NT 4.0 or earlier.

# Shared Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Network

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 or earlier.

# SMAddOns

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Network

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 or earlier.

# UMAddOns

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Network

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 or earlier.

# World Full Access Shared Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Network

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 or earlier.

# PDH

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
The PDH subkey stores configuration data for the Performance Data Helper dynamic-link library (Pdh.dll). PDH is a library of performance data collection and processing functions used by performance monitoring programs, such as System Monitor.

This subkey does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# CollectTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\PDH

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 30–3600 | None. |

**Description**
Specifies the amount of time, in seconds, that the system waits before timing out and returning an error when accessing performance counter information via the RegQueryValueEX interface of the Performance Data Helper library (Pdh.dll). When this entry is absent from the registry, the system waits 300 seconds.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
Changes to this entry are effective the next time that Pdh.dll is loaded. It might be necessary to stop and restart any program that uses Pdh.dll.

**Tip**

Consider increasing the value of this entry on heavily loaded systems. When the processing resources of a computer are under heavy use, a PDH call that uses the registry interface as a data source might take a long time to complete, and performance tools that use this interface can seem unresponsive. Increasing the value of this entry increases the amount of time that PDH waits before declaring the call unsuccessful.

# DefaultNullDataSource

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\PDH

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 \| 4 | 1 |

**Description**

Specifies the source from which the Performance Data Helper (PDH) functions retrieve performance counter data.

In all versions of Windows, the PDH functions use the registry to retrieve data collected by internal counters. In Windows 2000 and Windows Server 2003, PDH functions can also retrieve this data by using the Windows Management Instrumentation (WMI) interface.

| Value | Meaning |
|---|---|
| 1 | Registry. PDH functions retrieve performance counter data through registry interfaces. |
| 4 | WMI. PDH functions use WMI interfaces instead of registry interfaces to retrieve performance counter data. |

If the value of this entry is set to any value other than 4, the system behaves as if this entry is 1.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**

For changes to the entry to be effective, PDH must be reloaded.

# Perflib

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
The Perflib subkey stores configuration data for the Windows Performance Library, which collects and organizes data for performance tools, such as System Monitor.

In addition to entries, the Perflib subkey contains a Language-code subkey for each spoken language you configure for Windows Server 2003. The *Language-code* subkey stores performance counter names and their descriptions in the specified language. The *Language-code* subkey is named for the language code for that language. For example, the counters and descriptions for the English language are stored in a subkey named 009, the language code for English (United States).

**Activation Method**
To make changes to entries in this subkey effective, restart Windows.

**Tip**

- If you cannot display particular performance counters, check the Application Log in Event Viewer for events recorded by Perflib. When Perflib finds a disabling error in a performance counter dynamic-link library (DLL), it records an event in the log. Then it sets the value of Disable Performance Counters to 1 to disable the counters.

# Configuration Flags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x1F | 0x10 |

**Description**
Regulates the run-time test feature of the performance counter dynamic-link library (DLL).

The Performance Library (Perflib) tests performance counter DLLs whenever you select a performance counter in System Monitor or in a similar performance tool. It performs these reliability tests to identify and eliminate defective counter DLLs. A defective counter DLL can harm your Windows Server 2003 performance tools, preventing them from displaying any performance data and, in some cases, causing errors that shut down the system and start Dr. Watson for Windows.

Perflib can perform tests to detect the following errors:

- Slow procedure calls (non-trusted counter DLLs only).
- Data buffer alignment errors.
- Buffer integrity errors.
- Function errors.
- Program exceptions (critical).

The value of this entry determines which tests Perflib runs on trusted and non-trusted performance counter DLLs. It also determines how Perflib responds when a counter DLL fails a test. Perflib always records an error in the Application Log in Event Viewer (Source = Perflib). However, it can also disable the counter DLL so that the counters are not displayed in any performance tool until you manually re-enable the DLL.

By default, Perflib uses these test guidelines:

- Time the procedure calls of non-trusted counter DLLs. If a procedure call is slow, then record a message in the Application Log.
- Test trusted and non-trusted counter DLLs for data buffer alignment errors. If a counter DLL fails the test, then record a message in the Application Log.
- Thoroughly test the integrity of buffers that non-trusted counter DLLs return. If a counter fails the test, then disable the counter.
- Do not test the integrity of buffers that trusted counter DLLs return, unless the value of the ExtCounterTestLevel entry specifies otherwise. If a counter DLL fails a test that the **ExtCounterTestLevel** entry requires, then disable the counter.
- Detect errors in the function calls of trusted and non-trusted counter DLLs. If an error is detected, then record a message in the Application Log.
- Detect program exceptions in the function calls of trusted and non-trusted counter DLLs. If a DLL generates an exception, then disable the counters.

The value of this entry is a bitmask that you can use to change the default behavior. By default, only bit 0x10 (10000) is set. To

select a behavior, either set the bit representing that behavior to 1, or sum the values of all desired bits. For example, to prevent Perflib from testing for data buffer alignment (1 binary) and from disabling counters (10 binary), set the value of this entry to 0x3 (11 binary).

| Bit value | Test setting |
|---|---|
| 0x0 | Use default test behavior. |
| 0x1 (1 binary) | Do not test for data buffer alignment errors. (These are common errors that are not usually critical.) Perform all other tests. |
| 0x2 (10 binary) | Do not disable counters. Perform all tests, but if a counter generates a function error or a program exception, record a message in the Application Log. |
| 0x4 (100 binary) | Do not time procedure calls or perform buffer integrity tests. When set, the ExtCounterTestLevel entry and the entries that define a slow procedure call (see Notes below) are ignored. |
| 0x8 (1000 binary) | Disable counters if their procedure calls are slow. (Non-trusted counter DLLs only.) |
| 0x10 (10000 binary) | Collect debugging data. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make changes to this entry effective, either restart the Remote Registry Service or restart Windows.

### Notes

- Bit 0x2 (10) takes precedence over bit 0x8 (1000). If bit 0x2 is set, then Perflib does not disable performance counters.
- Perflib disables a counter by setting the value of the Disable Performance Counters entry for that service to 1. After you have repaired or replaced the counter DLL, you must re-enable it manually by changing the value of Disable Performance Counters to 1.
- Perflib does not time the procedure calls of trusted counters. To determine whether the procedure call of a non-trusted counter is abnormally slow, Perflib uses the standards established by the values of the OpenProcedureWaitTime, Open Timeout, and Collect Timeout entries.

### Tip

- If particular performance counters do not appear in your performance counter display tool, check the Application Log in Event Viewer for events recorded by Perflib.

### Caution

- Run-time tests can perceptibly delay the appearance of performance data, but they protect the system's monitoring capability.

**See Also**
**Concepts**
EventLogLevel Entry
Disable Performance Counters
ExtCounterTestLevel

# Disable Performance Counters Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Enables and disables the registry-based performance counters on the system.

If the value of this entry is 1, then Perflib does not retrieve performance counter data from the registry. As a result, System Monitor and other tools that use the data cannot display it. Instead, when a program requests performance data from the registry, the system returns an error to the program explaining that the Performance Library (Perflib) service is disabled.

| Value | Meaning |
|---|---|
| 0 | Enables performance counters. Programs can retrieve performance data from the registry. |
| 1 | Disables performance counters. Programs cannot retrieve performance data from the registry. |

**Activation Method**
To make changes to this entry effective, restart Windows.

### Note

- This entry affects all performance counters on the system. To enable or disable the performance counters for a particular service, use Exctrlst.exe to add Disable Performance Counters to the Performance Subkey subkey for the service.

### Tip

- If you cannot display particular performance counters, check the Application Log in Event Viewer for events recorded by Perflib. When Perflib finds a disabling error in a performance counter dynamic-link library (DLL), it records an event in the log. Then it sets the value of Disable Performance Counters to 1 to disable the counters.

### Caution

- If you change the value of this entry to 1, the **Performance object** and **Performance counter** lists in System Monitor will be empty. You will not be able to use System Monitor or any tool that uses Performance Library data to monitor your system.

**See Also**
**Concepts**
Disable Performance Counters

# EventLogLevel Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 \| 3 | 1 |

**Description**
Specifies which events the Performance Library records in the Application Log in Event Viewer.

| Value | Meaning |
|---|---|
| 0 | No recording. |
| 1 | Errors only. |
| 2 | Errors and warnings. |
| 3 | Errors, warnings, information, and success/failure conditions. |

Consider increasing the value of this entry if you are encountering problems starting or running a performance tool. If you find Perflib errors in the Application Log, increase this value again to record more detail on future errors.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make changes to this entry effective, either restart the Remote Registry Service or restart Windows.

### Note

- If a remote computer is using System Monitor to monitor the local computer, you must restart Windows on the local computer and stop and restart System Monitor on the remote computer before changes to this entry take effect.

### Tip

- If you cannot display particular performance counters, check the Application Log in Event Viewer for events recorded by Perflib. When Perflib finds a disabling error in a performance counter dynamic-link library (DLL), it records an event in the log. Then it sets the value of the Disable Performance Counters entry to 1 to disable the counters.

**See Also**
**Concepts**
ExtCounterTestLevel

# ExtCounterTestLevel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 \| 2 \| 3 \| 4 | 4 |

**Description**

Specifies the extent to which Perflib tests the data buffers of a trusted performance counter dynamic-link library (DLL).

If a counter DLL fails the buffer integrity test specified by the value of this entry, then, by default, Perflib records an error in the Application Log (Source = Perflib) and disables the counters. Perflib does not load the counters again until you have manually re-enabled them.

| Value | Meaning |
|---|---|
| 1 | High. The Performance Library thoroughly tests the pointers and buffers of returned trusted counter DLLs. To protect the user's buffer from damage, Perflib sends the DLL function a copy of the buffer, instead of the actual buffer. |
| 2 | Medium. The Performance Library checks pointers and buffer length, but it does not test pointer references or buffer contents. To protect the user's buffer from damage, Perflib sends the DLL function a copy of the buffer, instead of the actual buffer. |
| 3 | Low. The Performance Library does not test pointers or buffers. To protect the user's buffer from damage, Perflib sends the DLL function a copy of the buffer, instead of the actual buffer. |
| 4 | None. The Performance Library does not test pointers or buffers, and it sends the user's buffer, not a copy. This value minimizes System Monitor overhead, especially when monitoring applications that require large performance data buffers. |

Perflib tests the integrity of data buffers to identify and eliminate defective counter DLLs. One defective counter DLL can harm your Windows Server 2003 performance tools, preventing them from displaying any performance data and, in some cases, causing errors that shut down the system and start Dr. Watson for Windows.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**

To make changes to this entry effective, either restart the Remote Registry Service or restart Windows.

### Notes

- The data buffer integrity test is one of several Perflib run-time reliability tests. To configure the Perflib test feature, use the Configuration Flags entry.
- If the 0x4 bit (100) is set in the value of **Configuration Flags**, then Perflib does not perform data buffer integrity tests on any counters, including trusted counters.
- To prevent Perflib from disabling counters when they fail a test, set the 0x2 (10) bit in the value of the Configuration Flags entry. This bit prevents Perflib from disabling any counters that fail any tests. It is not specific to trusted counters or to data buffer integrity tests.
- Perflib disables a counter by setting the value of Disable Performance Counters for that service to 1. After you have repaired or replaced the counter DLL, you must re-enable it manually by changing the value of **Disable Performance Counters** to 1.
- This entry has a different effect on Windows 2000 and Windows Server 2003 than it has on Windows NT 4.0 and earlier.

In Windows NT 4.0 and earlier, it determines the testing level for extensible counters. For more information, see Regentry.hlp in the *Windows NT Resource Kit 4.0*.

**Tip**

- If a performance counter does not appear in your performance monitoring tool, check the Application Log in Event Viewer for events recorded by Perflib.

**See Also**
**Concepts**
Configuration Flags
EventLogLevel Entry
Library Validation Code
Disable Performance Counters

# OpenProcedureWaitTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xEE6B2800 *milliseconds* | 0x2710 (10 seconds) |

**Description**
Defines a slow Open procedure call for the performance counter dynamic-link libraries (DLLs) on the system.

When an Open procedure call takes longer to complete than the time specified in the value of this entry, then by default, Perflib records an error in the Application Log. However, if the value of Configuration Flags is 0x0, then Perflib disables the counter DLL and does not load the counters again until you have manually reenabled them.

Perflib times the procedure calls of non-trusted counters to identify counters that can delay or lock performance tools during initialization. While the Open procedure call proceeds, you cannot use the performance tools, and remote users cannot log on to the local computer.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make changes to this entry effective, restart Windows.

### Notes

- Perflib disables counters by setting the value of Disable Performance Counters for that service to 1. After you have repaired or replaced the counters, you must re-enable them manually by changing the value of **Disable Performance Counters** to 1.
- Perflib does not time the procedure calls of trusted performance counter DLLs This entry applies only to non-trusted DLLs.
- This entry applies to all counter DLLs on this system. To adjust the Open procedure time for the counters of a particular service, use the Open Timeout entry.
- The value of this entry is also the default value of the Collect Timeout entry.
- This entry has a different effect on Windows 2000 and Windows Server 2003 than it has on Windows NT 4.0 and earlier. In Windows NT 4.0 and earlier, it determines the Open procedure wait time only for extensible counters. For more information, see Regentry.hlp in the *Windows NT Resource Kit 4.0*.

### Tip

- If you suspect problems with procedure calls, check the Application Log in Event Viewer. Perflib records an error in the log whenever a procedure call exceeds its limit.

**See Also**
**Concepts**
Open Timeout
Configuration Flags
Collect Timeout

# TotalInstanceName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *String* | _Total |

**Description**
Specifies an alternate name for the **_Total** instance in performance tools, such as System Monitor.

The **_Total** instance represents the sum of values for all other instances of a performance object. For example, the **_Total** instance of Logical Disk: % Disk Time represents the sum of active time for all logical partitions on the computer. You can use this entry to change the name of the **_Total** instance as it appears in your performance tools.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make changes to this entry effective, restart Windows.

**Tip**

- If particular performance counters do not appear in your performance monitoring tool, check the Application Log in Event Viewer for events recorded by Perflib. When Perflib detects disabling errors in a counter dynamic-link library (DLL), it records an event in the log. Then it sets the value of Disable Performance Counters to 1 to disable the counters.

# Language-code

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib

**Description**
The *Language-code* subkey stores performance counter names and their descriptions in each spoken language you configure for Windows Server 2003. It is named for the language code associated with the language. For example, the counters and definitions for the English language are stored in a subkey named 009, the language code for English (United States).

The *Language-code* subkey is a variable representing the subkeys in the HKCU\Software\Microsoft\Perflib subkey. It does not actually appear in the registry. This variable subkey displays the entries that can appear in any of the language code subkeys in the Perflib subkey. The actual *Language-code* subkeys in the registry depend on the spoken languages chosen for the system.

**Activation Method**
To make changes to entries in this subkey effective, either restart the Remote Registry Service or restart Windows.

### Tip

- If you cannot display particular performance counters, check the Application Log in Event Viewer for events recorded by Perflib. When Perflib finds a disabling error in a performance counter dynamic-link library (DLL), it records an event in the log. Then it sets the value of Disable Performance Counters to 1 to disable the counters.

### Caution

- This subkey is maintained by the Performance Library. If you change the values, the performance tools on your system might not operate properly.

# Counter

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib*Language-code*

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Counter names and indices* | (None) |

**Description**
Stores the name and index of each performance counter installed on the system. The value of this entry comes from the .ini files for the counter.

**Activation Method**
To make changes to this entry effective, restart Windows.

**Tip**

- If you cannot display particular performance counters, check the Application Log in Event Viewer for events recorded by Perflib. When Perflib finds a disabling error in a performance counter dynamic-link library (DLL), it records an event in the log. Then it sets the value of Disable Performance Counters to 1 to disable the counters.

**Caution**

- Do not change the value of this entry. This value is maintained by the Performance Library. If you change it, the performance tools on your system will not operate properly.

# Help Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib*Language-code*

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Descriptions of counters* | (None) |

**Description**
Describes the performance counters installed on the system. The descriptions are designed for performance tools, such as System Monitor.

**Activation Method**
To make changes to this entry effective, either restart the Remote Registry Service or restart Windows.

**Tip**

- To view the description in System Monitor, in the **Add to Chart** dialog box, click **Explain**.
- For information about how to add descriptive text to the counters you develop, see the *Platform Software Development Kit* (SDK).

**Caution**

- Do not change the value of this entry. This value is maintained by the Performance Library. If you change it, the performance tools on your system will not operate properly.

# Print Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Printers Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print

**Description**
The Printers subkey contains information about the local printers.

### Note

- The printers listed in this subkey can be shared, or they can be accessible only to the host computer.
- The entries in this subkey are linked to the Print\Printers entries that appear in the SYSTEM key. Changes made in this subkey will be reflected in the SYSTEM key.

**See Also**
**Concepts**
server-name\Printers

# DefaultSpoolDirectory Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *path* | *Systemroot*\System32\spool\PRINTERS |

**Description**
Specifies the folder holding spooled files that the system is printing on a local printer.

This entry is backed up in the Printers subkey in the SYSTEM key. Any changes made to this entry in the SOFTWARE key will be reflected in the SYSTEM key.

**Change Method**
To change the value of this entry, from the **Print Server Properties** dialog box, click the **Advanced** tab, then type the new value in the **Spool folder** text box.

### Notes

- Before changing the spool folder, you should wait for all documents to finish printing.
- Do not set the value of this entry to be the root directory. If you do, then spooling reverts to the %Systemroot%\system32\spool\printer directory.

**See Also**
**Concepts**
Printers\DefaultSpoolDirectory

# printer-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers

**Description**
The *printer-name* subkey stores information about printers connected to the print server.

The subkeys and entries that appear in this subkey are backed up in the Control\Print subkey in HKEY_LOCAL_MACHINE\SYSTEM. Any changes made in the SOFTWARE\Printers subkey will be reflected in the SYSTEM\Printers subkey.

**Note**

- This subkey appears only on servers that host local printers. It does not appear on computers that connect only to network printers.

- The *printer-name* subkey is a placeholder representing the types of subkeys in the Printers subkey; "printer-name" does not actually appear in the registry.

# SpoolDirectory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers*printer-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *path* | (None) |

**Description**
Stores the location of a printer-specific location for spool files. The value of this entry takes precedence over the global default stored in the DefaultSpoolDirectory Entry entry.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Do not set the value of this entry to be the root directory. If you do, then spooling reverts to the %Systemroot%\system32\spool\printer directory.

# PrinterDriverData

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers*printer-name*

**Description**
The PrinterDriverData subkey contains configuration data for the specified printer.

**Change Method**
To change the value of entries in this subkey, use the SetPrinterData/GetPrinterData or the SetPrinterDataEx/GetPrinterDataEx APIs.

# Transparency

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers*printer-name*\PrinterDriverData

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | (blank) \| 1 | (None) |

**Description**
Specifies the background mode when the value of the Winprint_TextNoCRTranslation entry is set to 1. If the background mode is set to transparent, then the driver will not print the background of the document, regardless of color.

| Value | Meaning |
|---|---|
| (blank) | Background is opaque. |
| 1 | Background is transparent. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the SetPrinterData/GetPrinterData or the SetPrinterDataEx/GetPrinterDataEx APIs.

**See Also**
**Concepts**
Winprint_TextNoCRTranslation

# Winprint_TextNoCRTranslation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers*printer-name*\PrinterDriverData

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Enables or disables carriage return translation.

By default, Windows Server 2003 translates a carriage return to a carriage return plus a line feed.

| Value | Meaning |
|---|---|
| 0 | Carriage returns translate to a carriage return plus a line feed. |
| 1 | Carriage return translation is disabled. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the SetPrinterData/GetPrinterData or the SetPrinterDataEx/GetPrinterDataEx APIs.

**See Also**
**Concepts**
Transparency

# Winprint_TextNoTranslation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers*printer-name*\PrinterDriverData

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Enables or disables carriage return and line feed translation.

By default, Windows Server 2003 translates a carriage return or a line feed to a carriage return plus a line feed.

| Value | Meaning |
|---|---|
| 0 | Carriage return or line feed translate to a carriage return plus a line feed. |
| 1 | Carriage return and line feed translation is disabled. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the SetPrinterData/GetPrinterData or the SetPrinterDataEx/GetPrinterDataEx APIs.

# Providers Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print

**Description**
The Providers subkey contains information about the network printers. The printers listed in this subkey are not hosted locally, but the local computer is connected to them.

**See Also**
**Concepts**
Print\Providers
server-name\Providers

# LanMan Print Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Servers Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# server-name Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers

**Description**
The *server-name* subkey stores information about printers connected to the specified print server.

### Note

- The *server-name* subkey is a placeholder representing the types of subkeys in the Servers subkey; "server-name" does not actually appear in the registry.

# server-name\Printers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*

**Description**
The Printers subkey contains information about the individual network printers to which the local computer is connected.

**See Also**
**Concepts**
Printers Subkey

# Printers\DefaultSpoolDirectory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *path* | *Systemroot*\System32\spool\PRINTERS |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
DefaultSpoolDirectory Entry

# server-name\Providers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*

**Description**
The Providers subkey stores configuration data used by the specified server.

**See Also**
**Concepts**
Print\Providers
Providers Subkey

# Providers\EventLog

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*\Providers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x1F (bitmask) | Windows XP Professional: 0x1B (11011) Windows Server 2003: 0x1F (11111) |

**Description**
Specifies the level of detail of events that the print spooler logs for the specified server.

The value of this entry is a bitmask. To direct the print spooler to log entries of the type specified, set the bit representing that type to 1, or sum the hexadecimal values of all desired types. For example, the default value for Windows Server 2003, 11111, indicates that the print spooler logs all events. The default value for Windows XP Professional, 11011, indicates that the print spooler logs all events except information.

| Value | Event type |
|---|---|
| 0x0 (0 binary) | Success |
| 0x1 (1 binary) | Error |
| 0x2 (10 binary) | Warning |
| 0x4 (100 binary) | Information |
| 0x8 (1000 binary) | Audit success |
| 0x10 (10000 binary) | Audit failure |

**Change Method**
To change the value of this entry, from the **Start** menu, point to **Settings**, and then click **Printers**. From the **File** menu, click **Server Properties**, click the **Advanced** tab, and then select or clear the check boxes representing each event type.

### Note

- This entry only appears in the registry after a connection has been made to the remote machine.

**See Also**
**Concepts**
EventLog Entry

# RestartJobOnPoolEnabled Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*\Providers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether a print job on the specified network printer can be restarted after an error.

If the value of this entry is set to 1, then a **Cancel/Retry** dialog box appears on the user's system for the duration specified in the value of the Providers\RestartJobOnPoolError entry.

| Value | Meaning |
|---|---|
| 0 | Print jobs cannot be restarted after an error. |
| 1 | Print jobs can be restarted after an error. |

### Notes

- This entry applies only to pooled printers.
- The print job cannot be restarted if it has been completely sent to the printer and deleted from the print queue.

**See Also**
**Concepts**
Providers\RestartJobOnPoolEnabled
Providers\RestartJobOnPoolError

# Providers\RestartJobOnPoolError

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*\Providers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *number of seconds* | 0x258 (600 seconds = 10 minutes) |

**Description**
Specifies how long the print monitor for the specified network printer waits after an error occurs before starting a print job on another port in the print pool.

### Notes

- This entry applies only to pooled printers.
- Both the print monitor and the printer must support SNMP.
- The print job cannot be restarted if it has been completely sent to the printer and deleted from the print queue.

**See Also**
**Concepts**
RestartJobOnPoolError Entry

# Type 1 Installer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
Contains subkeys that store information about the Type 1 fonts installed on the system.

# Type 1 Fonts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Type 1 Installer

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | | |

**Description**
The Type 1 Fonts subkey contains the names of the Type 1 fonts that are installed on the system.

The fonts are stored in the *Systemroot*\WINNT\Fonts folder, and the corresponding entries store the following string values:

- *String 0* – Specifies with "T" for true and "F" for false whether *String 2* exists.
- *String 1* – Specifies the name of the Printer Font Metrics (.pfm) file.
- *String 2* – Specifies the name of the Printer Font Binary (.pfb) file.

For example, if the font Arial is installed as a Type 1 font, the corresponding value, when viewed in the Registry Editor (regedit.exe), could appear as follows: **T arial.pfm arial.pfb**.

**Change Method**
This entry is maintained automatically by the installation and removal of fonts in the Fonts folder, located in Control Panel.

### Notes

- The operating system uses Printer Font Metrics files to display a font on the screen.
- Printer Font Binary files contain mathematical descriptions for each character. Applications and printers use this information to print text in the described font.

**See Also**
**Concepts**
Fonts Subkey

# CurrentVersion\Windows Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
Stores configuration data used by 16-bit Windows 3.*x* applications running on Windows Server 2003 and earlier.

**See Also**
**Concepts**
CurrentVersion\Windows Control\Windows Subkey

# AppInit_DLLs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *DLL file names* | (This entry has no default value.) |

**Description**

Specifies which dynamic-link libraries (DLLs) installed on the computer are loaded by Windows 3.*x*–based applications.

# DeviceNotSelectedTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of seconds* (in decimal) | 15 |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows 3.51 or earlier.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# TransmissionRetryTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Number of seconds* (in decimal) | Windows NT 3.1: 45 Windows NT 3.51 and later: 90 |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

**Caution**

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# CurrentVersion\Winlogon

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
The Winlogon subkey stores configuration data for the Winlogon service. Winlogon provides support for interactive logon to Windows Server 2003.

This subkey stores data that applies to all users of the computer. The Winlogon Subkey subkey in HKEY_CURRENT_USER stores data that applies to a specific user.

**See Also**
**Concepts**
Winlogon Subkey

# AutoRestartShell

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**

Specifies whether the Windows user interface (typically, Explorer.exe) restarts automatically if it stops unexpectedly.

| Value | Meaning |
|---|---|
| 0 | If the Windows user interface or one of its components stops unexpectedly, the user must restart the interface. |
| 1 | If the Windows user interface or one of its components stops unexpectedly, the interface restarts automatically. |

**Note**

- The user can restart the interface either by logging off and logging on again, or by pressing Control-Alt-Del, selecting **Task Manager**, then from the **File** menu, selecting **Run** and typing "Explorer.exe".

# CachePrimaryDomain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | (None) |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

**Caution**

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# DCacheMinInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 120–86,400 *seconds* | 120 (2 minutes) |

**Description**

Specifies how often the list of domains in the **Unlock Workstation** dialog box is updated. The list of domains is updated before it is displayed to the user if its age exceeds the value of this entry.

A memory cache on the local computer stores the domain list. When the value of this entry is set to the default of two minutes, and a user unlocks a workstation or server, then the system retrieves domain data and updates the list. The short update interval ensures that the domain list is current. However, refreshing the list can cause a noticeable delay.

You can use this entry to increase the time between domain list updates, eliminating some of the delay. However, if the domain list changes between updates, users still will not see the most current list of domains.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# DCacheUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *Domain names in binary form* | (None) |

**Description**
Stores domain names in binary form for internal Winlogon programming code to use.

**Caution**

- Winlogon stores the value of this entry for its own use. Do not delete this entry or change its value. If you do, Windows Server 2003 might not operate properly.

# DebugServerCommand

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | (There is no range of values for this entry.) | (None) |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# DefaultDomainName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Domain-name* | NEWDOMAIN |

**Description**
Stores the name of the domain to which the user most recently logged on successfully. The value of this entry appears in the **Log On to Windows** dialog box the next time the dialog box is displayed.

**See Also**
**Concepts**
DefaultUserName

# DefaultUserName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *User-name* | Administrator |

**Description**
Stores the last user name entered in the **Log On to Windows** dialog box.

The value of this entry appears in the **Log On to Windows** dialog box the next time the dialog box is displayed. It is also used in automated logons and unattended setup.

### Note

- This entry is required if you have configured Windows Server 2003 to log on automatically by setting the value of the entry **AutoAdminLogon** to 1.

### Tip

- If you delete the value of this entry, then the system restores it the next time a user logs on to Windows Server 2003. To specify that the user-name field in the **Log On to Windows** dialog box be blank, set the value of the DontDisplayLastUserName entry to 1.

**See Also**
**Concepts**
DefaultDomainName
DontDisplayLastUserName

# DontDisplayLastUserName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ or REG_DWORD | 0 \| 1 (decimal only) | 0 |

**Description**
Specifies whether a user name appears in the **Log On to Windows** dialog box.

By default, Windows Server 2003 displays the user name of the last user who logged on successfully (as stored in the value of DefaultUserName) in the **Log On to Windows** dialog box. If the value of this entry is 1, then the **User name** box in the **Log On to Windows** dialog box is blank.

| Value | Meaning |
|---|---|
| 0 | The name of the last user who logged on successfully appears in the **Log On to Windows** dialog box. This setting is designed to make logging on faster and easier. |
| 1 | The **User name** field in the **Log On to Windows** dialog box is blank. This setting is designed to enhance the security of the system by not displaying a valid user name. |

**Note**

- A Group Policy setting in Windows Server 2003 can override the value of this entry. When the **Do not display last user name in logon screen** Group Policy setting (Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options) is configured, this entry is ignored.

**See Also**
**Concepts**
DefaultUserName

# KeepRasConnection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 0 |

**Description**
Specifies whether Windows Server 2003 closes all open dial-up (Routing and Remote Access) connections when a user logs off.

| Value | Meaning |
|---|---|
| 0 | RAS connections are closed. |
| 1 | RAS connections remain open. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# LogonPrompt

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *String* | (None) |

**Description**

Specifies a message for the **Log On to Windows** dialog box. By default, there is no logon prompt message.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Winlogon\ProfileDlgTimeOut

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *seconds* | 0x1E (30 seconds) |

**Description**
Defines the starting time for the User Environment Countdown Timer. This timer determines how long the system waits for a user response before it uses a default value. The value of this entry is used when either of the following occurs:

- The system detects a slow link between the client and the server that stores the client's profile.
- The system cannot access or update a server-based profile when the user logs on or off.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Notes

- The value of this entry can be overridden by a Group Policy setting. If your system does not respond as expected, the **Timeout for dialog boxes** Group Policy setting might be in effect. See the **ProfileDlgTimeout** entry in the **Policies** subkey.
- If you change this value by editing the registry, enter the value in decimal numbers.

**See Also**
**Concepts**
Winlogon\SlowLinkTimeOut

# RasForce

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 0 |

**Description**

Specifies whether the **Log on using dial-up connection** check box is selected by default when the system starts. This check box appears on the **Log On to Windows** dialog box that is displayed when the user logs on to Windows Server 2003.

The system uses the value of this entry only when you have created a dial-up connection by using Network and Dial-up Connections and the computer is a member of a domain.

| Value | Meaning |
|---|---|
| 0 | Cleared. The **Log on using dial-up connection** check box is cleared when displayed. Users can click it to select it. |
| 1 | Selected. The **Log on using dial-up connection** check box is selected when displayed. Users can click it to clear it. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# ReportBootOk

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | 0 \| 1 | 1 |

**Description**
Specifies whether Winlogon declares startup to be successful. If you change the value of this entry to 0, then you must use another program (such as Bootok.exe or Bootvrfy.exe) to declare the startup successful.

| Value | Meaning |
|---|---|
| 0 | Winlogon does not declare startup to be successful. |
| 1 | Winlogon declares startup to be successful after at least one user has logged on successfully. |

**Note**

- If this entry is deleted from the registry, Winlogon does not declare the startup to be successful.

**See Also**
**Concepts**
BootVerificationProgram
BootVerification

# ReportControllerMissing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | TRUE or any string | (None) |

**Description**
Stores the computer setting for the domain controller message. (The user setting is stored in the value of the ReportDC entry.)

This entry determines, in part, whether the system displays the following message when it cannot contact the domain controller that stores a user's roaming user profile:

```
A domain controller for your domain could not be contacted. You have been logged on using cached account
information. Changes made to your profile since you last logged on might not be available.
```

When the value of this entry is equal to **TRUE**, the system displays the message. When the value is any other string, and when this value is absent from the registry, the system does not display the message.

Two entries manage the display of the domain controller message: ReportDC and **ReportControllerMissing.** The system displays this message only when both entries are set to the display setting—that is, when **ReportDC** is either set to 1 or does not appear in the registry, and when **ReportDomainController** appears in the registry with a value of TRUE. By default, the message does not display.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
This entry does not exist in the registry by default, and you must add it to the registry and set its value equal to **TRUE** to cause the system to display the message when the system cannot contact the domain controller that stores the user's roaming profile.

## Notes

- While the domain controller message is displayed, other elements of the user interface are disabled, but all processes continue to run normally. If the **OK** button accompanying the domain controller message is not clicked after two minutes, then the message box disappears.

- The value of this entry must be in all capital letters. Otherwise, the system ignores the value and behaves as though the value is FALSE.

**See Also**
**Concepts**
ReportDC
NoPopUpsOnBoot

# Winlogon\RunLogonScriptSync Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the system waits for the logon script to finish running before it starts Windows Explorer and creates the desktop.

| Value | Meaning |
|---|---|
| 0 | The logon script and Windows Explorer can run simultaneously. |
| 1 | Windows Explorer does not start until the logon script has finished running. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To establish a policy setting related to this entry, use Group Policy. The value of this entry can be overridden by using the **Run logon scripts synchronously** Group Policy setting (User or Computer Configuration\Administrative Templates\System\Logon).

### Notes

- This entry applies to all users of the computer. The Winlogon\RunLogonScriptSync entry in HKEY_CURRENT_USER is configured separately for each user.

- Group Policy settings included in Windows Server 2003 can supersede this entry. While the **Run logon scripts synchronously** Group Policy setting is enabled (in the Computer Configuration or User Configuration policy folders), the system ignores this entry. The **Run logon scripts synchronously** policy setting is stored in the values of the **RunLogonScriptSync** entry in HKEY_LOCAL_MACHINE and the **RunLogonScriptSync** entry in HKEY_CURRENT_USER.

**See Also**
**Concepts**
Winlogon\RunLogonScriptSync

# ScreenSaverGracePeriod

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x20C49B (2,147,483) *seconds* (approximately 24 days) | 0x5 |

**Description**

Specifies when password protection of a screen saver becomes effective. This entry specifies the delay between the appearance of a password-protected screen saver and the enforcement of the password requirement.

Password protection of a screen saver is not effective immediately. By default, a brief period elapses within which the user can use the mouse or the keyboard to stop the screen saver without entering the password. This delay is designed to minimize the disruption that results when the screen saver starts while the user is working.

You can add this entry to the registry to adjust the length of the delay. To make password protection effective immediately, set the value of this entry to 0x0.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# SfcBugcheck

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 0 |

**Description**

Windows Server 2003 does not use this entry. It might appear in the registry when you upgrade from a pre-release version of Windows Server 2003.

# Winlogon\SFCDllCacheDir

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Fully qualified path* | *Systemroot*\System32\Dllcache |

**Description**
Specifies an alternate location for the Windows File Protection cache.

### Note

- This entry can be superseded by a Group Policy setting. While the policy setting **Specify Windows File Protection cache location** is enabled (which corresponds to registry entry SFCDllCacheDir Entry), the system ignores this entry. If the Windows File Protection cache does not appear in the expected location, check to see whether a policy setting is in effect.

**See Also**
**Concepts**
SFCDllCacheDir Entry

# Winlogon\SfcScan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 0 |

**Description**
Specifies when Windows File Protection scans protected files for changes. By default, it scans protected files only during setup.

| Value | Meaning |
|---|---|
| 0 | Windows File Protection scans files only during setup. |
| 1 | Windows File Protection scans files at setup and each time you start Windows Server 2003. This setting delays each startup. |
| 2 | Windows File Protection scans files the next time you start the system. |

**Change Method**
To change the value of this entry, use System File Checker (SFC.exe), a tool included in Windows Server 2003.

### Note

- This entry affects file scanning only. It does not affect the standard background file change detection that Windows File Protection provides.

- This entry can be superceded by a Group Policy setting included in Windows Server 2003. While the **Set Windows File Protection scanning** policy setting is enabled, the system ignores this entry. The configuration of the **Set Windows File Protection scanning** policy setting is stored in the value of **Sfcscan** in the Policies subkey. If Windows File Protection or System File Checker does not operate as expected, check to see whether a policy setting is in effect.

- Windows File Protection only scans when the current user is an administrator of the computer.

# SfcShowProgress Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Hides the file scan progress window. This window provides status information that might be helpful to advanced users, but it might confuse novices.

| Value | Meaning |
|---|---|
| 0 | The file scan progress window does not appear during file scanning. |
| 1 | The file scan progress window appears. |

**Notes**

- This entry can be superseded by a Group Policy setting included in Windows Server 2003. While the policy setting **Hide the file scan progress window** is enabled (which corresponds to registry entry Windows File Protection\SfcShowProgress), the system ignores this entry. If the scan progress window does not appear, check to see whether a policy setting is in effect.

- Windows File Protection only scans when the current user is an administrator of the computer.

**See Also**
**Concepts**
Windows File Protection\SfcShowProgress

# Shell Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Name of executable file* | Explorer.exe |

**Description**
Specifies the programs that provide the user interface to the operating system.

By default, Winlogon starts the programs specified in the value of **Userinit**, including Userinit.exe. Userinit.exe starts the user interface program. However, if Winlogon cannot start the programs specified in the value of **Userinit**, then Winlogon directly runs the programs specified in the value of this entry.

**See Also**
**Concepts**
Winlogon\System Subkey

# ShowLogonOptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**

Specifies whether logon options are displayed in the **Log on to Windows** dialog box when the dialog box opens.

The **Log on to Windows** dialog box has an **Options** button that alternately hides and displays the **Domain** box and the **Log on using dial-up connection** options.

| Value | Meaning |
|---|---|
| 0 | The options are hidden when the dialog box opens, but users can click the **Options** button to display them. |
| 1 | The options are displayed when the dialog box opens, but users can click the **Options** button to hide them. |

# Winlogon\SlowLinkProfileDefault

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Directs the system to wait for the remote copy of the roaming user profile to load, even if the profile is loading slowly. Also, the system waits for the remote copy when the user is notified about a slow connection but does not respond in the time allowed.

| Value | Meaning |
|---|---|
| 0 | Use the local copy. When a remote profile is slow to load, the system loads the local copy of the roaming user profile. The local copy is also used when the user is consulted (as set in the entry **SlowLinkUIEnabled**) but does not respond in the time allowed (as set in the Winlogon\ProfileDlgTimeOut entry). |
| 1 | The policy is enabled. The system waits for the remote copy of the user profile to load. |

Waiting for the remote profile is appropriate when a user moves between computers frequently and the local copy of the profile is not always current. Using the local copy is desirable when logging on quickly is a priority.

**Note**

- This entry can be superseded by a Group Policy setting included in Windows Server 2003. When the **Wait for remote user profile** policy setting is enabled, the system ignores this entry. The setting of the **Wait for remote user profile** policy setting is stored in the value of **SlowLinkProfileDefault** in the Policies subkey. If the slow link feature does not operate as expected, check to see if a policy setting is in effect.

# Winlogon\SlowLinkTimeOut

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–120,000 *milliseconds* | 2000 (2 seconds) |

**Description**
Defines which connections Winlogon considers to be slow. If the server storing a user's profile does not respond before the time specified in the value of this entry, then Winlogon considers the link to the server to be slow, and it offers the user the option of using a profile stored on the local computer instead of waiting for the server.

This entry is used only when the value of **SlowLinkDetectEnabled** is 1.

You can increase the value of **SlowLinkTimeOut** to favor the server-based profile. This is particularly useful for clients using addresses assigned by Dynamic Host Configuration Protocol (DHCP), or for computers accessing server-based profiles over slow wide area network (WAN) connections, such as dial-up connections.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use Group Policy. This entry corresponds to the **Slow network connection timeout** Group Policy setting (Computer Settings\Software Policies\Windows NT User Profiles).

**See Also**
**Concepts**
Winlogon\ProfileDlgTimeOut

# Winlogon\System Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Names of executable files* | (None) |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# Welcome Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *String* | (None) |

**Description**
Specifies the text that appears in the caption bar beside the title of the **Log On to Windows**, **Windows Security**, **Computer Locked**, and **Unlock Computer** dialog boxes.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Tip**

- The system adds only one space between the dialog box title and your message. To further separate the title from your message, begin your message with leading spaces.

# WOW

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion

**Description**
Contains subkeys and entries that control startup parameters affecting MS-DOS-based applications and applications created for 16-bit Windows 3.*x*.

**Caution**

- Do not change any of the subkeys in this subkey.

# boot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\WOW

**Description**
The boot subkey contains information about drivers and fonts that are loaded for each instance of the Windows on Windows (WOW) subsystem.

**Caution**

- Do not change any of the entries in this subkey.

# Microsoft\Windows Script Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Windows Script \Settings Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Microsoft\Windows Script

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Enabled Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Microsoft\Windows Script\Settings

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether Windows Script Host (WSH) is disabled for all users of the machine. This entry disables all scripts that rely on WSH, including VBScript and Jscript. If WSH is disabled and the user attempts to run a WSH script, then a message alerts the user that WSH is disabled. The message appears even if the user attempts to run the script from a batch file or a designated script host.

| Value | Meaning |
|---|---|
| 0 | WSH is disabled for all users of the machine. |
| 1 | WSH is enabled for all users of the machine. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
Settings\Enabled

# Policies Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Policies\Microsoft Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# CallSecurity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Set Call Security options**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# ConfigFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Enable Automatic Configuration**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# IntranetSupportURL

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Set the intranet support Web page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# MaxFileSendSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Limit the size of sent files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# MaximumBandwidth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Limit the bandwidth of Audio and Video**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Audio & Video.

# NoAddingDirectoryServers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent adding Directory servers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# NoAdvancedCalling

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Disable the Advanced Calling button**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Options Page.

# NoAllowControl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent Control**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Application Sharing.

# NoAppSharing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Disable application Sharing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Application Sharing.

# NoAudio

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Disable Audio**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Audio & Video.

# NoAudioPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Hide the Audio page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Options Page.

# NoAutoAcceptCalls

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent automatic acceptance of Calls**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# NoChangeDirectSound

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent changing DirectSound Audio setting**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Audio & Video.

# NoChangingCallMode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent changing Call placement method**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# NoChat

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Disable Chat**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# NoDirectoryServices

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Disable Directory services**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# NoFullDuplex

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Disable full duplex Audio**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Audio & Video.

# NoGeneralPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Hide the General page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Options Page.

# NoNewWhiteBoard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Disable Whiteboard**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# NoOldWhiteBoard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Disable NetMeeting 2.x Whiteboard**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# NoReceivingFiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent receiving files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# NoReceivingVideo

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent receiving Video**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Audio & Video.

# NoSecurityPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Hide the Security page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Options Page.

# NoSendingFiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent sending files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# NoSendingVideo

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent sending Video**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Audio & Video.

# NoSharing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent Sharing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Application Sharing.

# NoSharingDesktop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent Desktop Sharing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Application Sharing.

# NoSharingDosWindows

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent Sharing Command Prompts**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Application Sharing.

# NoSharingExplorer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent Sharing Explorer windows**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Application Sharing.

# NoTrueColorSharing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent Application Sharing in true color**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Application Sharing.

# NoVideoPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Hide the Video page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting\Options Page.

# NoWebDirectory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Prevent viewing Web directory**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# PersistAutoAcceptCalls

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Allow persisting automatic acceptance of Calls**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# Use AutoConfig

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Conferencing

**Description**
Stores configuration data for the policy setting **Enable Automatic Configuration**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\NetMeeting.

# Conferencing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

**Description**
The Conferencing subkey is one of two subkeys that store entries for Group Policy settings for NetMeeting. Group Policy creates and maintains these entries, and the NetMeeting program reads and interprets them.

This subkey stores entries that apply to all users of the computer. The system adds these entries when you configure a NetMeeting policy setting in the Computer Configuration folder in Group Policy.

The Conferencing subkey in HKEY_CURRENT_USER stores policy setting related entries that are configured separately for each user. The system adds these entries when you configure a NetMeeting policy setting in the User Configuration folder in Group Policy.

In Group Policy, the NetMeeting ploicy settings appear in the following folders:

- Computer Configuration\Administrative Templates\Windows Components\NetMeeting.
- User Configuration\Administrative Templates\Windows Components\NetMeeting.

The NetMeeting Group Policy settings appear only when the Conf.adm administrative template file is added to Group Policy. Conf.adm is included in Windows Server 2003 (at *Systemroot*\inf\conf.adm), but it is not loaded automatically. For instructions, see Help for Group Policy.

# Control Panel Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Control Panel\Desktop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Control Panel

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# MultiUILanguageID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Restrict selection of Windows menus and dialogs language**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Regional and Language Options.

# Microsoft\Internet Explorer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

**Description**
The Internet Explorer subkey is one of several subkeys that store entries for Group Policy settings for Internet Explorer. Group Policy creates and maintains these entries, and the Internet Explorer program reads and interprets them.

The Internet Explorer Group Policy settings appear when the Inetres.adm file is added to Group Policy. By default, this file is added automatically.

In Group Policy, the Internet Explorer settings appear in the following folders:

- Computer Configuration\Administrative Templates\Windows Components\Internet Explorer
- User Configuration\Administrative Templates\Windows Components\Internet Explorer

# DisableImportExportFavorites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer

**Description**
Stores configuration data for the policy setting **Disable importing and exporting of favorites**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Advanced Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Control Panel

**Description**
Stores configuration data for the policy setting **Disable changing Advanced page settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# AdvancedTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Control Panel

**Description**
Stores configuration data for the policy setting **Disable the Advanced page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel.

# ConnectionsTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Control Panel

**Description**
Stores configuration data for the policy setting **Disable the Connections page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel.

# GeneralTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Control Panel

**Description**
Stores configuration data for the policy setting **Disable the General page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel.

# HomePage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Control Panel

**Description**
Stores configuration data for the policy setting **Disable changing home page settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# PrivacyTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Control Panel

**Description**
Stores configuration data for the policy setting **Disable the Privacy page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel.

# ProgramsTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Control Panel

**Description**
Stores configuration data for the policy setting **Disable the Programs page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel.

# SecurityTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Control Panel

**Description**
Stores configuration data for the policy setting **Disable the Security page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel.

# ContentTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Control Panel

**Description**
Stores configuration data for the policy setting **Disable the Content page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel.

# Control Panel Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Infodelivery

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer

**Description**
The Infodelivery subkey is one of several subkeys that store entries for Group Policy settings for Internet Explorer. Group Policy creates and maintains these entries, and the Internet Explorer program reads and interprets them.

The Internet Explorer Group Policy settings appear when the Inetres.adm file is added to Group Policy. By default, this file is added automatically.

In Group Policy, the Internet Explorer settings appear in the following folders:

- Computer Configuration\Administrative Templates\Windows Components\Internet Explorer
- User Configuration\Administrative Templates\Windows Components\Internet Explorer

# Restrictions Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer\Infodelivery

**Description**
The Restrictions subkey stores the entries created when you configure Group Policy settings that restrict the installation, updating, or presentation of Internet Explorer. Group Policy creates and maintains these entries, and the Internet Explorer program reads and interprets them.

# MaxSubscriptionCount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Subscription Limits**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# MaxSubscriptionSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Subscription Limits**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# MaxWebcrawlLevels

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Subscription Limits**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# MinUpdateInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Subscription Limits**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoAddingChannels

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable adding channels**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoAddingSubscriptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable adding schedules for offline pages**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoBrowserClose

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **File menu: Disable closing the browser and Explorer windows**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoBrowserOptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Tools menu: Disable Internet Options... menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoBrowserSaveWebComplete

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **File menu: Disable Save As Web Page Complete**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoChannelLogging

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable offline page hit logging**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoChannelUI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable channel user interface completely**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoEditingScheduleGroups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable editing and creating of schedule groups**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoEditingSubscriptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable editing schedules for offline pages**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoFavorites

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Hide Favorites menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoJITSetup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable Automatic Install of Internet Explorer components**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# NoRemovingChannels

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable removing channels**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoRemovingSubscriptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable removing schedules for offline pages**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoScheduledUpdates

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable all scheduled offline pages**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoSearchCustomization

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Search: Disable Search Customization**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# NoSubscriptionContent

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Disable downloading of site subscription content**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# NoTheaterMode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **View menu: Disable Full Screen menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoViewSource

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **View menu: Disable Source menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# UpdateExcludeBegin

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Subscription Limits**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# UpdateExcludeEnd

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Infodelivery\Restrictions

**Description**
Stores configuration data for the policy setting **Subscription Limits**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Offline Pages.

# Persistence

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# 0

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer\Persistence

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DocumentLimit Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\0

**Description**
Stores configuration data for the policy setting **File size limits for Local Machine zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# 0\DomainLimit Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\0

**Description**
Stores configuration data for the policy setting **File size limits for Local Machine zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# 1\DocumentLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\1

**Description**
Stores configuration data for the policy setting **File size limits for Intranet zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# DomainLimit Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\1

**Description**
Stores configuration data for the policy setting **File size limits for Intranet zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# 1

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer\Persistence

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# 2

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer\Persistence

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# 2\DomainLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\2

**Description**
Stores configuration data for the policy setting **File size limits for Trusted Sites zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# 2\DocumentLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\2

**Description**
Stores configuration data for the policy setting **File size limits for Trusted Sites zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# 3\DocumentLimit Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\3

**Description**
Stores configuration data for the policy setting **File size limits for Internet zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# 3\DomainLimit Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\3

**Description**
Stores configuration data for the policy setting **File size limits for Internet zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# 3

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer\Persistence

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# 4

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer\Persistence

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# 4\DocumentLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\4

**Description**
Stores configuration data for the policy setting **File size limits for Restricted Sites zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# 4\DomainLimit Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Persistence\4

**Description**
Stores configuration data for the policy setting **File size limits for Restricted Sites zone**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Persistence Behavior.

# No_LaunchMediaBar

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Configure Media Explorer Bar**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# NoBrowserContextMenu

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Disable Context menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoBrowserSaveAs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **File menu: Disable Save As... menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoExternalBranding

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Disable external branding of Internet Explorer**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# NoFileNew

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **File menu: Disable New menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoFileOpen

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **File menu: Disable Open menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoFindFiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Search: Disable Find Files via F3 within the browser**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# NoHelpItemNetscapeHelp

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Help menu: Remove 'For Netscape Users' menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoHelpItemSendFeedback

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Help menu: Remove 'Send Feedback' menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoHelpItemTipOfTheDay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Help menu: Remove 'Tip of the Day' menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoHelpItemTutorial

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Help menu: Remove 'Tour' menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoOpeninNewWnd

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Disable Open in New Window menu option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# NoSelectDownloadDir

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Internet Explorer\Restrictions

**Description**
Stores configuration data for the policy setting **Disable Save this program to disk option**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Browser menus.

# Internet Explorer\Restrictions Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Internet Explorer

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Messenger_DepKit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Messenger\Client

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Messenger

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Client\PreventAutoRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Messenger\Client

**Description**
Stores configuration data for the policy setting **Do not automatically start Windows Messenger initially**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Messenger.

# PreventAutoRun Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Messenger\Client

**Description**
Stores configuration data for the policy setting **Do not automatically start Windows Messenger initially**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Messenger.

# Client\PreventRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Messenger\Client

**Description**
Stores configuration data for the policy setting **Do not allow Windows Messenger to be run**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Messenger.

# PreventRun Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Messenger\Client

**Description**
Stores configuration data for the policy setting **Do not allow Windows Messenger to be run**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Messenger.

# RestrictAuthorMode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC

**Description**
Stores configuration data for the policy setting **Restrict the user from entering author mode**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console.

# RestrictToPermittedSnapins

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Restrict users to the explicitly permitted list of snap-ins**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console.

# Microsoft\MMC

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {011BE22D-E453-11D1-945A-00C04FB984F9}\Restrict_Run

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{011BE22D-E453-11D1-945A-00C04FB984F9}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Security Configuration and Analysis**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {011BE22D-E453-11D1-945A-00C04FB984F9}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {03f1f940-a0f2-11d0-bb77-00aa00a1eab7}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# SMTP Protocol Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{03f1f940-a0f2-11d0-bb77-00aa00a1eab7}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **SMTP Protocol**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# System Properties Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{0F3621F1-23C6-11D1-AD97-00AA00B88E5A}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **System Properties**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {0F3621F1-23C6-11D1-AD97-00AA00B88E5A}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {0F6B957D-509E-11D1-A7CC-0000F87571E3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Administrative Templates (Computers) Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{0F6B957D-509E-11D1-A7CC-0000F87571E3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Administrative Templates (Computers)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# Administrative Templates (Users) Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{0F6B957E-509E-11D1-A7CC-0000F87571E3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Administrative Templates (Users)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# {0F6B957E-509E-11D1-A7CC-0000F87571E3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {18BA7139-D98B-43c2-94DA-2604E34E175D}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# .Net Framework Configuration Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{18BA7139-D98B-43c2-94DA-2604E34E175D}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **.Net Framework Configuration**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Routing and Remote Access Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{1AA7F839-C7F5-11D0-A376-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Routing and Remote Access**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {1AA7F839-C7F5-11D0-A376-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {1AA7F83C-C7F5-11D0-A376-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# AppleTalk Routing Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{1AA7F83C-C7F5-11D0-A376-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **AppleTalk Routing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# ADSI Edit Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{1C5DACFA-16BA-11D2-81D0-0000F87A7AA3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **ADSI Edit**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {1C5DACFA-16BA-11D2-81D0-0000F87A7AA3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {1F5EEC01-1214-4D94-80C5-4BDCD2014DDD}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Authorization Manager Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{1F5EEC01-1214-4D94-80C5-4BDCD2014DDD}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Authorization Manager**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# Wireless Monitor Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{23DC5869-BD9F-46fd-AADD-1F869BA64FC3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Wireless Monitor**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {23DC5869-BD9F-46fd-AADD-1F869BA64FC3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {243E20B0-48ED-11D2-97DA-00A024D77700}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Removable Storage Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{243E20B0-48ED-11D2-97DA-00A024D77700}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Removable Storage**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# Wireless Network (IEEE 802.11) Policies Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{2DA6AA7F-8C88-4194-A558-0D36E7FD3E64}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Wireless Network (IEEE 802.11) Policies**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# {2DA6AA7F-8C88-4194-A558-0D36E7FD3E64}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {2E19B602-48EB-11d2-83CA-00104BCA42CF}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# IAS Logging Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{2E19B602-48EB-11d2-83CA-00104BCA42CF}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IAS Logging**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {3060E8CE-7020-11D2-842D-00C04FA372D4}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Remote Installation Services Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{3060E8CE-7020-11D2-842D-00C04FA372D4}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remote Installation Services**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# {34AB8E82-C27E-11D1-A6C0-00C04FB94F17}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Public Key Policies Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{34AB8E82-C27E-11D1-A6C0-00C04FB94F17}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Public Key Policies**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# Event Viewer Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{394C052E-B830-11D0-9A86-00C04FD8DBF7}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Event Viewer**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {394C052E-B830-11D0-9A86-00C04FD8DBF7}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {3CB6973D-3E6F-11D0-95DB-00A024D77700}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Restrict_Run Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{3CB6973D-3E6F-11D0-95DB-00A024D77700}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Removable Storage Management**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Remote Desktops Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{3D5D035E-7721-4B83-A645-6C07A3D403B7}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remote Desktops**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {3D5D035E-7721-4B83-A645-6C07A3D403B7}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {3F276EB4-70EE-11D1-8A0F-00C04FB93753}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Certification Authority Policy Settings Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{3F276EB4-70EE-11D1-8A0F-00C04FB93753}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Certification Authority Policy Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# Scripts (Startup/Shutdown) Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{40B6664F-4972-11D1-A7CA-0000F87571E3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Scripts (Startup/Shutdown)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# {40B6664F-4972-11D1-A7CA-0000F87571E3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {40B66650-4972-11D1-A7CA-0000F87571E3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Scripts (Logon/Logoff) Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{40B66650-4972-11D1-A7CA-0000F87571E3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Scripts (Logon/Logoff)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# Disk Defragmenter Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{43668E21-2636-11D1-A1CE-0080C88593A5}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Disk Defragmenter**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {43668E21-2636-11D1-A1CE-0080C88593A5}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {45ac8c63-23e2-11d1-a696-00c04fd58bc3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# System Information Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{45ac8c63-23e2-11d1-a696-00c04fd58bc3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **System Information**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Certificates Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{53D6AB1D-2488-11D1-A28C-00C04FB94F17}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Certificates**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {53D6AB1D-2488-11D1-A28C-00C04FB94F17}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {57C596D0-9370-40C0-BA0D-AB491B63255D}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# IP Security Monitor Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{57C596D0-9370-40C0-BA0D-AB491B63255D}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IP Security Monitor**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Shared Folders Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{58221C65-EA27-11CF-ADCF-00AA00A80033}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Shared Folders**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {58221C65-EA27-11CF-ADCF-00AA00A80033}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {58221C66-EA27-11CF-ADCF-00AA00A80033}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Services Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{58221C66-EA27-11CF-ADCF-00AA00A80033}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Services**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Computer Management Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{58221C67-EA27-11CF-ADCF-00AA00A80033}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Computer Management**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {58221C67-EA27-11CF-ADCF-00AA00A80033}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {58221C69-EA27-11CF-ADCF-00AA00A80033}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Shared Folders Extension Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{58221C69-EA27-11CF-ADCF-00AA00A80033}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Shared Folders Ext**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# Remote Access Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{5880CD5C-8EC0-11d1-9570-0060B0576642}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remote Access**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {5880CD5C-8EC0-11d1-9570-0060B0576642}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {5ADF5BF6-E452-11D1-945A-00C04FB984F9}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Security Templates Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{5ADF5BF6-E452-11D1-945A-00C04FB984F9}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Security Templates**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# WMI Control Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{5C659257-E236-11D2-8899-00104B2AFB46}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **WMI Control**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {5C659257-E236-11D2-8899-00104B2AFB46}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {5D6179C8-17EC-11D1-9AA9-00C04FD8FE93}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Local Users and Groups Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{5D6179C8-17EC-11D1-9AA9-00C04FD8FE93}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Local Users and Groups**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Distributed File System Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{677A2D94-28D9-11D1-A95B-008048918FB1}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Distributed File System**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {677A2D94-28D9-11D1-A95B-008048918FB1}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {6DC3804B-7212-458D-ADB0-9A07E2AE1FA2}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Resultant Set of Policy snap-in Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{6DC3804B-7212-458D-ADB0-9A07E2AE1FA2}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Resultant Set of Policy snap-in**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy.

# Logical and Mapped Drives Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{6E8E0081-19CD-11D1-AD91-00AA00B8E05A}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Logical and Mapped Drives**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {6E8E0081-19CD-11D1-AD91-00AA00B8E05A}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {74246bfc-4c96-11d0-abef-0020af6b0b7a}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Device Manager Extension Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{74246bfc-4c96-11d0-abef-0020af6b0b7a}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Device Manager**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# Performance Logs and Alerts Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{7478EF61-8C46-11d1-8D99-00A0C913CAD4}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Performance Logs and Alerts**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {7478EF61-8C46-11d1-8D99-00A0C913CAD4}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {753EDB4D-2E1B-11D1-9064-00A0C90AB504}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# FAX Service Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{753EDB4D-2E1B-11D1-9064-00A0C90AB504}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **FAX Service**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# SNMP Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{7AF60DD3-4979-11D1-8A6C-00C04FC33566}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **SNMP**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {7AF60DD3-4979-11D1-8A6C-00C04FC33566}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {7E45546F-6D52-4D10-B702-9C2E67232E62}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Software Installation (Computers) Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{7E45546F-6D52-4D10-B702-9C2E67232E62}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Software Installation (Computers)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Resultant Set of Policy snap-in extensions.

# Security Settings Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{803E14A0-B4FB-11D0-A0D0-00A0C90F574B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Security Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# {803E14A0-B4FB-11D0-A0D0-00A0C90F574B}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {88E729D6-BDC1-11D1-BD2A-00C04FB9603F}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Folder Redirection Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{88E729D6-BDC1-11D1-BD2A-00C04FB9603F}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Folder Redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# Disk Management Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{8EAD3A12-B2C1-11d0-83AA-00A0C92C9D5D}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Disk Management**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {8EAD3A12-B2C1-11d0-83AA-00A0C92C9D5D}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {8F8F8DC0-5713-11D1-9551-0060B0576642}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Internet Authentication Service (IAS) Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{8F8F8DC0-5713-11D1-9551-0060B0576642}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Internet Authentication Service (IAS)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Group Policy Object Editor Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{8FC0B734-A0E1-11D1-A7D3-0000F87571E3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy Object Editor**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy.

# {8FC0B734-A0E1-11D1-A7D3-0000F87571E3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {90087284-d6d6-11d0-8353-00a0c90640bf}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Device Manager Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{90087284-d6d6-11d0-8353-00a0c90640bf}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Device Manager**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# IPX Routing Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{90810500-38F1-11D1-9345-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IPX Routing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {90810500-38F1-11D1-9345-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {90810502-38F1-11D1-9345-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# IPX RIP Routing Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{90810502-38F1-11D1-9345-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IPX RIP Routing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# IPX SAP Routing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{90810504-38F1-11D1-9345-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IPX SAP Routing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {90810504-38F1-11D1-9345-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {942A8E4F-A261-11D1-A760-00C04FB9603F}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Software Installation (Computers) Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{942A8E4F-A261-11D1-A760-00C04FB9603F}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Software Installation (Computers)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# Indexing Service Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{95AD72F0-44CE-11D0-AE29-00AA004B9986}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Indexing Service**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {95AD72F0-44CE-11D0-AE29-00AA004B9986}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {9EC88934-C774-11d1-87F4-00C04FC2C17B}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DCOM Configuration Extension

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{9EC88934-C774-11d1-87F4-00C04FC2C17B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **DCOM Configuration Extension**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {A841B6C2-7577-11D0-BB1F-00A0C922E79C}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Internet Information Services Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{A841B6C2-7577-11D0-BB1F-00A0C922E79C}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Internet Information Services**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {A994E107-6854-4F3D-917C-E6F01670F6D3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Certificate Templates Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{A994E107-6854-4F3D-917C-E6F01670F6D3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Certificate Templates**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Send Console Message Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{B1AFF7D0-0C49-11D1-BB12-00C04FC9A3A3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Send Console Message**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {B1AFF7D0-0C49-11D1-BB12-00C04FC9A3A3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {B52C1E50-1DD2-11D1-BC43-00C04FC31FD3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# RAS Dialin - User Node Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{B52C1E50-1DD2-11D1-BC43-00C04FC31FD3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **RAS Dialin - User Node**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# Administrative Templates (Computers) Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{B6F9C8AE-EF3A-41C8-A911-37370C331DD4}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Administrative Templates (Computers)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Resultant Set of Policy snap-in extensions.

# {B6F9C8AE-EF3A-41C8-A911-37370C331DD4}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {B6F9C8AF-EF3A-41C8-A911-37370C331DD4}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Administrative Templates (Users) Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{B6F9C8AF-EF3A-41C8-A911-37370C331DD4}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Administrative Templates (Users)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Resultant Set of Policy snap-in extensions.

# Extended View (Web View) Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{B708457E-DB61-4C55-A92F-0D4B5E9B1224}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Extended View (Web View)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {B708457E-DB61-4C55-A92F-0D4B5E9B1224}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {B91B6008-32D2-11D2-9888-00A0C925F917}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Terminal Services Configuration Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{B91B6008-32D2-11D2-9888-00A0C925F917}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Terminal Services Configuration**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Software Installation (Users) Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{BACF5C8A-A3C7-11D1-A760-00C04FB9603F}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Software Installation (Users)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# {BACF5C8A-A3C7-11D1-A760-00C04FB9603F}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {BD95BA60-2E26-AAD1-AD99-00AA00B8E05A}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Service Dependencies Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{BD95BA60-2E26-AAD1-AD99-00AA00B8E05A}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Service Dependencies**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# IP Routing Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{C2FE4500-D6C2-11D0-A37B-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IP Routing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {C2FE4500-D6C2-11D0-A37B-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {C2FE4502-D6C2-11D0-A37B-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DHCP Relay Management Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{C2FE4502-D6C2-11D0-A37B-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **DHCP Relay Management**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# RIP Routing Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{C2FE4504-D6C2-11D0-A37B-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **RIP Routing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {C2FE4504-D6C2-11D0-A37B-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {C2FE4506-D6C2-11D0-A37B-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# OSPF Routing Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{C2FE4506-D6C2-11D0-A37B-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **OSPF Routing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# IGMP Routing Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{C2FE4508-D6C2-11D0-A37B-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IGMP Routing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# {C2FE4508-D6C2-11D0-A37B-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {C2FE450B-D6C2-11D0-A37B-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Connection Sharing (NAT) Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{C2FE450B-D6C2-11D0-A37B-00C04FC9DA04}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Connection Sharing (NAT)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# ActiveX Control Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{C96401CF-0E17-11D3-885B-00C04F72C717}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **ActiveX Control**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {C96401CF-0E17-11D3-885B-00C04F72C717}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {C96401D1-0E17-11D3-885B-00C04F72C717}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Link to Web Address Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{C96401D1-0E17-11D3-885B-00C04F72C717}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Link to Web Address**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Component Services Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{C9BC92DF-5B9A-11D1-8F00-00C04FC2C17B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Component Services**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {C9BC92DF-5B9A-11D1-8F00-00C04FC2C17B}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {D70A2BEA-A63E-11D1-A7D4-0000F87571E3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Group Policy for Active Directory Tools Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{D70A2BEA-A63E-11D1-A7D4-0000F87571E3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy tab for Active Directory Tools**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy.

# Active Directory Sites and Services Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{D967F824-9968-11D0-B936-00C04FD8D5B0}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Active Directory Sites and Services**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {D967F824-9968-11D0-B936-00C04FD8D5B0}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {DAB1A262-4FD7-11D1-842C-00C04FB6C218}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Routing Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{DAB1A262-4FD7-11D1-842C-00C04FB6C218}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Routing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Extension snap-ins.

# Certification Authority Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{de751566-4cc6-11d1-8ca0-00c04fc297eb}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Certification Authority**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {de751566-4cc6-11d1-8ca0-00c04fc297eb}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {DEA8AFA0-CC85-11d0-9CE2-0080C7221EBD}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# IP Security Policy Management Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{DEA8AFA0-CC85-11d0-9CE2-0080C7221EBD}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IP Security Policy Management**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Group Policy Management Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{E12BBB5D-D59D-4E61-947A-301D25AE8C23}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy Management**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy.

# {E12BBB5D-D59D-4E61-947A-301D25AE8C23}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {E26D02A0-4C1F-11D1-9AA1-00C04FC3357A}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Telephony Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{E26D02A0-4C1F-11D1-9AA1-00C04FC3357A}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Telephony**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Active Directory Users and Computers Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{E355E538-1C2E-11D0-8C37-00C04FD8FE93}

| Data type | Range | Default value |
| --- | --- | --- |
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Active Directory Users and Computers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {E355E538-1C2E-11D0-8C37-00C04FD8FE93}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {EBC53A38-A23F-11D0-B09B-00C04FD8DCA6}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Active Directory Domains and Trusts Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{EBC53A38-A23F-11D0-B09B-00C04FD8DCA6}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Active Directory Domains and Trusts**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# Internet Explorer Maintenance Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{FC715823-C5FB-11D1-9EEF-00A0C90347FF}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Internet Explorer Maintenance**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions.

# {FC715823-C5FB-11D1-9EEF-00A0C90347FF}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {FD57D297-4FD9-11D1-854E-00C04FC31FD3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# QoS Admission Control Run Restriction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{FD57D297-4FD9-11D1-854E-00C04FC31FD3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **QoS Admission Control**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# FrontPage Server Extensions Run Restrictions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\{FF5903A8-78D6-11D1-92F6-006097B01056}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **FrontPage Server Extensions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins.

# {FF5903A8-78D6-11D1-92F6-006097B01056}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\MMC\

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Netlogon Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Netlogon\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Netlogon

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# AllowExclusiveScriptsShareAccess

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Netlogon share compatibility**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# AllowExclusiveSysvolShareAccess

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Sysvol share compatibility**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# AllowSingleLabelDnsDomain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Location of the DCs hosting a domain with single label DNS name**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# AutoSiteCoverage Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Automated Site Coverage by the DC Locator DNS SRV Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# Parameters\AvoidPdcOnWan Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Contact PDC on logon failure**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# BackgroundRetryInitialPeriod

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Initial DC Discovery Retry Setting for Background Callers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# BackgroundRetryMaximumPeriod

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Maximum DC Discovery Retry Interval Setting for Background Callers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# BackgroundRetryQuitTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Final DC Discovery Retry Setting for Background Callers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# BackgroundSuccessfulRefreshPeriod

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Positive Periodic DC Cache Refresh for Background Callers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# DnsAvoidRegisterRecords

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **DC Locator DNS records not registered by the DCs**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# DnsRefreshInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Refresh Interval of the DC Locator DNS Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# DnsTtl Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **TTL Set in the DC Locator DNS Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# ExpectedDialupDelay Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Expected dial-up delay on logon**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# GcSiteCoverage Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Sites Covered by the GC Locator DNS SRV Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# LdapSrvPriority

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Priority Set in the DC Locator DNS SRV Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# LdapSrvWeight

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Weight Set in the DC Locator DNS SRV Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# MaximumLogFileSize Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Maximum Log File Size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# NdncSiteCoverage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Sites Covered by the Application Directory Partition Locator DNS SRV Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# NegativeCachePeriod

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Negative DC Discovery Cache Setting**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# NonBackgroundSuccessfulRefreshPeriod

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Positive Periodic DC Cache Refresh for Non-Background Callers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# ScavengeInterval Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Scavenge Interval**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# SiteCoverage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Sites Covered by the DC Locator DNS SRV Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# SiteName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Site Name**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# UseDynamicDns

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Dynamic Registration of the DC Locator DNS Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon\DC Locator DNS Records.

# dbFlag

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Netlogon\Parameters

**Description**
Stores configuration data for the policy setting **Log File Debug Output Level**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Net Logon.

# PCHealth Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ForceQueueMode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting

**Description**
Stores configuration data for the policy setting **Report Errors**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting.

# IncludeKernelFaults

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting

**Description**
Stores configuration data for the policy setting **Report operating system errors**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting\Advanced Error Reporting settings.

# IncludeMicrosoftApps

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting

**Description**
Stores configuration data for the policy setting **Default application reporting settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting\Advanced Error Reporting settings.

# IncludeShutdownErrs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting

**Description**
Stores configuration data for the policy setting **Report unplanned shutdown events**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting\Advanced Error Reporting settings.

# IncludeWindowsApps

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting

**Description**
Stores configuration data for the policy setting **Default application reporting settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting\Advanced Error Reporting settings.

# ShowUI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting

**Description**
Stores configuration data for the policy setting **Display Error Notification**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting.

# PCHealth\ErrorReporting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\PCHealth

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# AllOrNone

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting

**Description**
Stores configuration data for the policy setting **Default application reporting settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting\Advanced Error Reporting settings.

# DoReport

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting

**Description**
Stores configuration data for the policy setting **Report Errors**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting.

# ErrorReporting\DW

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\PCHealth\ErrorReporting

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DWAllowHeadless

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting\DW

**Description**
Stores configuration data for the policy setting **Display Error Notification**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting.

# DWFileTreeRoot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting\DW

**Description**
Stores configuration data for the policy setting **Report Errors**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting.

# DWNoExternalURL

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting\DW

**Description**
Stores configuration data for the policy setting **Report Errors**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting.

# DWNoFileCollection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting\DW

**Description**
Stores configuration data for the policy setting **Report Errors**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting.

# DWNoSecondLevelCollection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting\DW

**Description**
Stores configuration data for the policy setting **Report Errors**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting.

# DWReporteeName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\ErrorReporting\DW

**Description**
Stores configuration data for the policy setting **Report Errors**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Error Reporting.

# Headlines

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\PCHealth\HelpSvc

**Description**
Stores configuration data for the policy setting **Do not allow Did you know content to appear**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Help and Support Center.

# HelpSvc

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\PCHealth

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Microsoft\System Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DFSClient

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\System

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DfsDcNameDelay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\System\DFSClient

**Description**
Stores configuration data for the policy setting **Sets how often a DFS Client discovers DC's**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network.

# NV PrimaryDnsSuffix

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\System\DNSClient

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Primary DNS Suffix**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\DNS Client.

# DNSClient Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\System

**Description**
The DNSClient subkey stores entries created when you configure a Group Policy setting for the Domain Name System (DNS) client. Group Policy creates and maintains the entries in this subkey, and DNS reads and interprets them.

# W32Time Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# W32Time\Config Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\W32Time

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# AnnounceFlags Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# EventLogFlags Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# Config\FrequencyCorrectRate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# HoldPeriod Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# LargePhaseOffset Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# Config\LocalClockDispersion

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# MaxAllowedPhaseOffset

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# MaxNegPhaseCorrection Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# MaxPollInterval Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# MinPollInterval Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# PhaseCorrectRate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# PollAdjustFactor Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# SpikeWatchPeriod Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# UpdateInterval Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# MaxPosPhaseCorrection Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\Config

**Description**
Stores configuration data for the policy setting **Global Configuration Settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service.

# CrossSiteSyncFlags Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32time\Parameters

**Description**
Stores configuration data for the policy setting **Configure Windows NTP Client**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service\Time Providers.

# Parameters\EventLogFlags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32time\Parameters

**Description**
Stores configuration data for the policy setting **Configure Windows NTP Client**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service\Time Providers.

# NtpServer Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32time\Parameters

**Description**
Stores configuration data for the policy setting **Configure Windows NTP Client**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service\Time Providers.

# ResolvePeerBackoffMaxTimes Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32time\Parameters

**Description**
Stores configuration data for the policy setting **Configure Windows NTP Client**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service\Time Providers.

# ResolvePeerBackoffMinutes Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32time\Parameters

**Description**
Stores configuration data for the policy setting **Configure Windows NTP Client**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service\Time Providers.

# SpecialPollInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32time\Parameters

**Description**
Stores configuration data for the policy setting **Configure Windows NTP Client**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service\Time Providers.

# Parameters\Type

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32time\Parameters

**Description**
Stores configuration data for the policy setting **Configure Windows NTP Client**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service\Time Providers.

# W32Time\Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\W32Time

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# TimeProviders Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\W32Time

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# TimeProviders\NtpClient Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\W32Time\TimeProviders

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# NtpClient\Enabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enable Windows NTP Client**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service\Time Providers.

# NtpServer\Enabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\W32Time\TimeProviders\NtpServer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enable Windows NTP Server**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows Time Service\Time Providers.

# TimeProviders\NtpServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\W32Time\TimeProviders

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Windows Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# App Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

**Description**
The App Management subkey contains entries storing the Group Policy settings that determine how programs run. Group Policy creates and maintains the subkeys and entries in the Windows subkey, and the component that implements the subkeys and entries reads and interprets them.

### Note

- Policy setting entries in the App Management subkey override default settings in the SOFTWARE\Microsoft\Windows\CurrentVersion\App Management subkeys.

**See Also**
**Concepts**
App Management

# App Management\COMClassStore

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\App Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Download missing COM components**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# App Management\COMClassStore Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\App Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Download missing COM components**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# DisableEngine

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\AppCompat

**Description**
Stores configuration data for the policy setting **Turn Off Application Compatibility Engine**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Application Compatibility.

# DisablePropPage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\AppCompat

**Description**
Stores configuration data for the policy setting **Remove Program Compatibility Property Page**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Application Compatibility.

# DisableWizard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\AppCompat

**Description**
Stores configuration data for the policy setting **Turn Off Program Compatibility Wizard**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Application Compatibility.

# LogAppHelpEvents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\AppCompat

**Description**
Stores configuration data for the policy setting **Turn On Application Help Log Events**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Application Compatibility.

# AppCompat\VDMDisallowed

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\AppCompat

**Description**
Stores configuration data for the policy setting **Prevent access to 16-bit applications**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Application Compatibility.

# AppCompat\VDMDisallowed Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\AppCompat

**Description**
Stores configuration data for the policy setting **Prevent access to 16-bit applications**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Application Compatibility.

# AppCompat

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# BITS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# JobInactivityTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\BITS

**Description**
Stores configuration data for the policy setting **Background Intelligent Transfer Service (BITS) inactive job timeout**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network.

# Windows\Control Panel Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Desktop Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\Control Panel

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Desktop\ScreenSaveActive

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Screen Saver**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# ScreenSaverIsSecure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Password protect the screen saver**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# Desktop\ScreenSaveTimeOut

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Screen Saver timeout**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# Desktop\SCRNSAVE.EXE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Control Panel\Desktop

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Screen Saver executable name**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Display.

# Windows\CurrentVersion Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

**Description**

The CurrentVersion subkey contains subkeys and entries that store Group Policy settings for components of Windows. The subkeys and entries in the CurrentVersion subkey are written and maintained by Group Policy, and they are read and interpreted by the component that implements them.

This subkey stores entries for policy settings that apply to all users of this computer. There is also a Software\Policies\Microsoft\Windows\CurrentVersion subkey in HKEY_CURRENT_USER, which stores policy setting related entries that are configured separately for each user.

# Identities

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\CurrentVersion\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Locked Down

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Identities

**Description**
Stores configuration data for the policy setting **Identity Manager: Prevent users from using Identities**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# Accessibility Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing accessibility settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Internet Settings Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\CurrentVersion

**Description**
The Internet Settings subkey is one of several subkeys that store entries for Group Policy settings for Internet Explorer. Group Policy creates and maintains these entries, and the Internet Explorer program reads and interprets them.

The Internet Explorer Group Policy settings appear when the Inetres.adm file is added to Group Policy. By default, this file is added automatically.

In Group Policy, the Internet Explorer settings appear in the following folders:

- Computer Configuration\Administrative Templates\Windows Components\Internet Explorer
- User Configuration\Administrative Templates\Windows Components\Internet Explorer

**Note**

- The Internet Explorer folder in Group Policy is created by the Inetres.adm administrative templates file. This file is included in Windows Server 2003, and it is loaded automatically when you start Group Policy.

# Autoconfig

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing Automatic Configuration settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Cache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing Temporary Internet files settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# CalendarContact

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing Calendar and Contact settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# Certificates Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing certificate settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Check_If_Default

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing default browser check**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# Internet Settings\Colors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing color settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Connwiz Admin Lock

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable Internet Connection wizard**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# DialupAutodetect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Use Automatic Detection for dial-up connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# DisplayScriptDownloadFailureUI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Display error message on proxy script download failure**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# EnableAutoProxyResultCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable caching of Auto-Proxy scripts**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Internet Settings\Fonts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing font settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# FormSuggest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable AutoComplete for forms**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# FormSuggest Passwords

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Do not allow AutoComplete to save passwords**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# History

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing history settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Languages

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing language settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# links

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing link color settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Messaging

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing Messaging settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing Profile Assistant settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# Proxy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing proxy settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# ProxySettingsPerUser

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Make proxy settings per-machine (rather than per-user)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# Ratings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing ratings settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# ResetWebSettings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable the Reset Web Settings feature**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# Security_HKLM_only

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Security Zones: Use only machine settings** .

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# Security_options_edit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Security Zones: Do not allow users to change policies**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# Security_zones_map_edit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Security Zones: Do not allow users to add/delete sites**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer.

# Connection Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

**Description**
Stores configuration data for the policy setting **Disable changing connection settings**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Explorer.

# {DED22F57-FEE2-11D0-953B-00C04FD9152D}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Carpoint**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# AllowedControls\{05589FA1-C356-11CE-BF01-00AA0055595A}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Media Player**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {05589FA1-C356-11CE-BF01-00AA0055595A} Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Media Player**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {22D6F312-B0F6-11D0-94AB-0080C74C7E95} Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Media Player**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# AllowedControls\{22D6F312-B0F6-11D0-94AB-0080C74C7E95}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Media Player**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {26F24A93-1DA2-11D0-A334-00AA004A5FC5}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **NetShow File Transfer Control**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {275E2FE0-7486-11D0-89D6-00A0C90C9B67}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Menu Controls**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {2D360201-FFF5-11D1-8D03-00A0C959BC0A}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **DHTML Edit Control**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {2FF18E10-DE11-11D1-8161-00A0C90DD90C}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **MSNBC**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {52ADE293-85E8-11D2-BB22-00104B0EA281}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Investor**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {7823A620-9DD9-11CF-A662-00AA00C066D2}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Menu Controls**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {9276B91A-E780-11d2-8A8D-00C04FA31D93}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Investor**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {AE24FDAE-03C6-11D1-8B76-0080C744F389}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Microsoft Scriptlet Component**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {BD1F006E-174F-11D2-95C0-00C04F9A8CFA}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Microsoft Survey Control**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {D27CDB6E-AE6D-11CF-96B8-444553540000}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Shockwave Flash**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {D45FD31B-5C6E-11D1-9EC1-00C04FD7081F}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Microsoft Agent**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {D6526FE0-E651-11CF-99CB-00C04FD64497}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Microsoft Chat**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# {F5131C24-E56D-11CF-B78A-444553540000}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\CurrentVersion\Internet Settings\AllowedControls

**Description**
Stores configuration data for the policy setting **Menu Controls**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in \Windows Components\Internet Explorer\Administrator Approved Controls.

# AllowedControls

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\CurrentVersion\Internet Settings

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# EnableFilter

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Directory UI

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enable filter in Find dialog box**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Directory.

# HideDirectoryFolder

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Directory UI

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Active Directory folder**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Directory.

# QueryLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Directory UI

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Maximum size of Active Directory searches**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Desktop\Active Directory.

# Directory UI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DriverSearching

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DontSearchCD

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\DriverSearching

**Description**
Stores configuration data for the policy setting **Configure driver search locations**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# DontSearchFloppies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\DriverSearching

**Description**
Stores configuration data for the policy setting **Configure driver search locations**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# DontSearchWindowsUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\DriverSearching

**Description**
Stores configuration data for the policy setting **Configure driver search locations**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# OnlyUseLocalAdminFiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

**Description**
Stores configuration data for the policy setting **Always use local ADM files for Group Policy Object Editor**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {0ACDD40C-75AC-47ab-BAA0-BF6DE7E7FE63}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# NoBackgroundPolicy Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{0ACDD40C-75AC-47ab-BAA0-BF6DE7E7FE63}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Wireless policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {0ACDD40C-75AC-47ab-BAA0-BF6DE7E7FE63}\NoGPOListChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{0ACDD40C-75AC-47ab-BAA0-BF6DE7E7FE63}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Wireless policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# NoSlowLink Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{0ACDD40C-75AC-47ab-BAA0-BF6DE7E7FE63}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Wireless policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {25537BA6-77A8-11D2-9B6C-0000F8080861}\NoGPOListChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{25537BA6-77A8-11D2-9B6C-0000F8080861}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Folder Redirection policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {25537BA6-77A8-11D2-9B6C-0000F8080861}\NoSlowLink

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{25537BA6-77A8-11D2-9B6C-0000F8080861}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Folder Redirection policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {25537BA6-77A8-11D2-9B6C-0000F8080861}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {35378EAC-683F-11D2-A89A-00C04FBBCFA2}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {35378EAC-683F-11D2-A89A-00C04FBBCFA2}\NoBackgroundPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{35378EAC-683F-11D2-A89A-00C04FBBCFA2}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Registry policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# NoGPOListChanges Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{35378EAC-683F-11D2-A89A-00C04FBBCFA2}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Registry policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {3610eda5-77ef-11d2-8dc5-00c04fa31a66}\NoBackgroundPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{3610eda5-77ef-11d2-8dc5-00c04fa31a66}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Disk Quota policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {3610eda5-77ef-11d2-8dc5-00c04fa31a66}\NoGPOListChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{3610eda5-77ef-11d2-8dc5-00c04fa31a66}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Disk Quota policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {3610eda5-77ef-11d2-8dc5-00c04fa31a66}\NoSlowLink

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{3610eda5-77ef-11d2-8dc5-00c04fa31a66}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Disk Quota policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {3610eda5-77ef-11d2-8dc5-00c04fa31a66}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {42B5FAAE-6536-11d2-AE5A-0000F87571E3}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {42B5FAAE-6536-11d2-AE5A-0000F87571E3}\NoBackgroundPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{42B5FAAE-6536-11d2-AE5A-0000F87571E3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Scripts policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {42B5FAAE-6536-11d2-AE5A-0000F87571E3}\NoGPOListChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{42B5FAAE-6536-11d2-AE5A-0000F87571E3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Scripts policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {42B5FAAE-6536-11d2-AE5A-0000F87571E3}\NoSlowLink

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{42B5FAAE-6536-11d2-AE5A-0000F87571E3}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Scripts policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {827D319E-6EAC-11D2-A4EA-00C04F79F83A}\NoGPOListChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{827D319E-6EAC-11D2-A4EA-00C04F79F83A}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Security policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {827D319E-6EAC-11D2-A4EA-00C04F79F83A}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {827D319E-6EAC-11D2-A4EA-00C04F79F83A}\NoBackgroundPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{827D319E-6EAC-11D2-A4EA-00C04F79F83A}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Security policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {A2E30F80-D7DE-11d2-BBDE-00C04F86AE3B}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {A2E30F80-D7DE-11d2-BBDE-00C04F86AE3B}\NoBackgroundPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{A2E30F80-D7DE-11d2-BBDE-00C04F86AE3B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Internet Explorer Maintenance policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {A2E30F80-D7DE-11d2-BBDE-00C04F86AE3B}\NoGPOListChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{A2E30F80-D7DE-11d2-BBDE-00C04F86AE3B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Internet Explorer Maintenance policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {A2E30F80-D7DE-11d2-BBDE-00C04F86AE3B}\NoSlowLink

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{A2E30F80-D7DE-11d2-BBDE-00C04F86AE3B}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Internet Explorer Maintenance policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {B1BE8D72-6EAC-11D2-A4EA-00C04F79F83A}\NoBackgroundPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{B1BE8D72-6EAC-11D2-A4EA-00C04F79F83A}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **EFS recovery policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {B1BE8D72-6EAC-11D2-A4EA-00C04F79F83A}\NoGPOListChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{B1BE8D72-6EAC-11D2-A4EA-00C04F79F83A}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **EFS recovery policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {B1BE8D72-6EAC-11D2-A4EA-00C04F79F83A}\NoSlowLink

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{B1BE8D72-6EAC-11D2-A4EA-00C04F79F83A}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **EFS recovery policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {B1BE8D72-6EAC-11D2-A4EA-00C04F79F83A}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {c6dc5466-785a-11d2-84d0-00c04fb169f7}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# {c6dc5466-785a-11d2-84d0-00c04fb169f7}\NoGPOListChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{c6dc5466-785a-11d2-84d0-00c04fb169f7}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Software Installation policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {c6dc5466-785a-11d2-84d0-00c04fb169f7}\NoSlowLink

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{c6dc5466-785a-11d2-84d0-00c04fb169f7}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Software Installation policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {e437bc1c-aa7d-11d2-a382-00c04f991e27}\NoBackgroundPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{e437bc1c-aa7d-11d2-a382-00c04f991e27}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IP Security policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {e437bc1c-aa7d-11d2-a382-00c04f991e27}\NoGPOListChanges

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{e437bc1c-aa7d-11d2-a382-00c04f991e27}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IP Security policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {e437bc1c-aa7d-11d2-a382-00c04f991e27}\NoSlowLink

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy\{e437bc1c-aa7d-11d2-a382-00c04f991e27}

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **IP Security policy processing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# {e437bc1c-aa7d-11d2-a382-00c04f991e27}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DCOption

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy Editor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy domain controller selection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# DisableAutoADMUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy Editor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Turn off automatic update of ADM files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# GPODisplayName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy Editor

**Description**
Stores configuration data for the policy setting **Default name for new Group Policy objects**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# NewGPOLinksDisabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy Editor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Create new Group Policy object links disabled by default**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# ShowPoliciesOnly

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Group Policy Editor

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enforce Show Policies Only**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# Group Policy Editor

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Windows\Installer Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

**Description**
The Installer subkey stores the entries created when you configure a Group Policy setting for Windows Installer in Windows Server 2003. Group Policy creates and maintains these entries, and Windows Installer reads and interprets them.

This subkey stores entries that apply to all users of the computer. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the Computer Configuration folder in Group Policy.

The Installer Subkey subkey in HKEY_CURRENT_USER stores entries that are configured separately for each user. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the User Configuration folder in Group Policy.

**Change Method**
To add, remove, or change the value of entries in this subkey, use Group Policy.

**See Also**
**Concepts**
Installer Subkey

# AllowLockdownBrowse

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enable user to browse for source while elevated**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# AllowLockdownMedia

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enable user to use media source while elevated**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# AllowLockdownPatch

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enable user to patch elevated products**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# AlwaysInstallElevated

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Always install with elevated privileges**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# DisableBrowse

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove browse dialog box for new source**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# DisableMedia

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prevent removable media source for any install**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# DisableMSI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Disable Windows Installer**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# DisablePatch

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit patching**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# DisableRollback

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit rollback**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# DisableUserInstalls

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

**Description**
Stores configuration data for the policy setting **Prohibit User Installs**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# EnableAdminTSRemote

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Allow admin to install from Terminal Services session**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# EnableUserControl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enable user control over installs**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# Installer\Logging

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Logging**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# SafeForScripting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Disable IE security prompt for Windows Installer scripts**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# SearchOrder

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Search order**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# TransformsSecure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Installer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Cache transforms in secure location on workstation**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Installer.

# SyncAtLogoff

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Synchronize all offline files before logging off**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# SyncAtLogon

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Synchronize all offline files when logging on**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# SyncAtSuspend

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache

**Description**
Stores configuration data for the policy setting **Synchronize offline files before suspend**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# DefCacheSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Default cache size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# NetCache\Enabled Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Allow or Disallow use of the Offline Files feature**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# GoOfflineAction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Action on server disconnect**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# NetCache\NoConfigCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit user configuration of Offline Files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# NoConfigCache Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit user configuration of Offline Files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# NetCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

**Description**
The NetCache subkey contains entries that store Group Policy settings for the Offline Files feature. Group Policy creates and maintains these entries, and the Offline Files program reads and interprets them. In Group Policy, these policy settings appear in the Computer Configuration\Administrative Templates\Network\Offline Files folder.

This subkey stores entries related to policy settings that apply to all users of this computer. The NetCache subkey in HKEY_CURRENT_USER stores policy setting related entries that are configured separately for each user.

Also, many of the policy setting related entries in this subkey take precedence over user settings while the policy setting is in effect. The affected user settings are stored in HKLM\Software\Microsoft\Windows\CurrentVersion\NetCache.

# AlwaysPinSubFolders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Subfolders always available offline**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# EncryptCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

**Description**
Stores configuration data for the policy setting **Encrypt the Offline Files cache**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# EventLoggingLevel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Event logging level**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# NoReminders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Turn off reminder balloons**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# PurgeAtLogoff

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **At logoff, delete local copy of userÆs offline files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# PurgeOnlyAutoCacheAtLogoff

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

**Description**
Stores configuration data for the policy setting **At logoff, delete local copy of userÆs offline files**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# ReminderBalloonTimeoutSeconds Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Reminder balloon lifetime**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# AssignedOfflineFolders\ReminderBalloonTimeoutSeconds

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Reminder balloon lifetime**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# ReminderFreqMinutes Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Reminder balloon frequency**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# AssignedOfflineFolders\ReminderFreqMinutes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Reminder balloon frequency**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# AssignedOfflineFolders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\NetCache

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# InitialBalloonTimeoutSeconds

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\AssignedOfflineFolders

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Initial reminder balloon lifetime**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# CustomGoOfflineActions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\NetCache

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ExcludeExtensions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\CustomGoOfflineActions

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Files not cached**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# NoCacheViewer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\CustomGoOfflineActions

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prevent use of Offline Files folder**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# NoMakeAvailableOffline

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\CustomGoOfflineActions

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Remove 'Make Available Offline'**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# DisableFRAdminPin

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\NoMakeAvailableOfflineList

**Description**
Stores configuration data for the policy setting **Do not automatically make redirected folders available offline**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# SlowLinkSpeed

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\NetCache\NoMakeAvailableOfflineList

**Description**
Stores configuration data for the policy setting **Configure Slow link speed**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Offline Files.

# NoMakeAvailableOfflineList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\NetCache

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# NC_AddRemoveComponents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit adding and removing components for a LAN or remote access connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_AdvancedSettings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit access to the Advanced Settings item on the Advanced menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_AllowAdvancedTCPIPConfig

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit TCP/IP advanced configuration**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_AllowNetBridge_NLA

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

**Description**
Stores configuration data for the policy setting **Prohibit installation and configuration of Network Bridge on your DNS domain network**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_ChangeBindState

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit Enabling/Disabling components of a LAN connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_DeleteAllUserConnection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Ability to delete all user remote access connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_DeleteConnection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit deletion of remote access connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_DialupPrefs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit access to the Remote Access Preferences item on the Advanced menu**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_EnableAdminProhibits

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

**Description**
Stores configuration data for the policy setting **Enable Windows 2000 Network Connections settings for Administrators**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_LanChangeProperties

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit access to properties of components of a LAN connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_LanConnect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Ability to Enable/Disable a LAN connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_LanProperties

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit access to properties of a LAN connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_NewConnectionWizard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit access to the New Connection Wizard**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_PersonalFirewallConfig

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

**Description**
Stores configuration data for the policy setting **Prohibit use of Internet Connection Firewall on your DNS domain network**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_RasAllUserProperties

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Ability to change properties of an all user remote access connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_RasChangeProperties

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit access to properties of components of a remote access connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_RasConnect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit connecting and disconnecting a remote access connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_RasMyProperties

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit changing properties of a private remote access connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_RenameAllUserRasConnection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

**Description**
Stores configuration data for the policy setting **Ability to rename all user remote access connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_RenameConnection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Ability to rename LAN connections or remote access connections available to all users**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_RenameLanConnection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

**Description**
Stores configuration data for the policy setting **Ability to rename LAN connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_RenameMyRasConnection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit renaming private remote access connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_ShowSharedAccessUI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit use of Internet Connection Sharing on your DNS domain network**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# NC_Statistics

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit viewing of status for an active connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# Windows\Network Connections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

**Description**
The Network Connections subkey contains entries that store Group Policy settings for Network and Dial-up Connections. Group Policy creates and maintains the entries in this subkey, and the Network and Dial-up Connections program reads and interprets them.

This subkey stores entries related to policy settings that apply to all users of this computer. The Software\Policies\Microsoft\Windows\Network Connections subkey in HKEY_CURRENT_USER stores policy setting related entries that are configured separately for each user.

# 8021XCARootHash

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Network Connections\8021X

**Description**
Stores configuration data for the policy setting **IEEE 802.1x Certificate Authority for Machine Authentication**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\Network Connections.

# 8021X

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\Network Connections

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# MaxOutstandingSends

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Limit outstanding packets**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler.

# TimerResolution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Set timer resolution**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler.

# NonBestEffortLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Limit reservable bandwidth**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler.

# Psched

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

**Description**
The Psched subkey stores entries created when you configure a Group Policy setting for the Quality of Service (QoS) Packet Scheduler.

Each entry in this subkey stores a policy setting used to configure the Packet Scheduler. The Packet Scheduler reads and interprets these entries when it implements the policy setting.

# DiffservByteMappingConforming\ServiceTypeBestEffort

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingConforming

**Description**
Stores configuration data for the policy setting **Best effort service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of conforming packets.

# DiffservByteMappingConforming\ServiceTypeControlledLoad

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingConforming

**Description**
Stores configuration data for the policy setting **Controlled load service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of conforming packets.

# DiffservByteMappingConforming\ServiceTypeGuaranteed

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingConforming

**Description**
Stores configuration data for the policy setting **Guaranteed service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of conforming packets.

# DiffservByteMappingConforming\ServiceTypeNetworkControl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingConforming

**Description**
Stores configuration data for the policy setting **Network control service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of conforming packets.

# DiffservByteMappingConformingServiceTypeQualitative

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingConforming

**Description**
Stores configuration data for the policy setting **Qualitative service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of conforming packets.

# DiffServByteMappingConforming

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\Psched

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

**See Also**
**Concepts**
DiffServByteMappingNonConforming

# DiffServByteMappingNonConforming

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\Psched

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

**See Also**
**Concepts**
DiffServByteMappingConforming

# ServiceTypeBestEffort Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingNonConforming

**Description**
Stores configuration data for the policy setting **Best effort service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of non-conforming packets.

# DiffservByteMappingNonConforming\ServiceTypeControlledLoad

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingNonConforming

**Description**
Stores configuration data for the policy setting **Controlled load service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of non-conforming packets.

# DiffservByteMappingNonConforming\ServiceTypeGuaranteed

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingNonConforming

**Description**
Stores configuration data for the policy setting **Guaranteed service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of non-conforming packets.

# DiffservByteMappingNonConforming\ServiceTypeNetworkControl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingNonConforming

**Description**
Stores configuration data for the policy setting **Network control service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of non-conforming packets.

# DiffservByteMappingNonConforming\ServiceTypeQualitative

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingNonConforming

**Description**
Stores configuration data for the policy setting **Qualitative service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\DSCP value of non-conforming packets.

# UserPriorityMapping\ServiceTypeBestEffort

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\UserPriorityMapping

**Description**
Stores configuration data for the policy setting **Best effort service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\Layer-2 priority value.

# UserPriorityMapping\ServiceTypeControlledLoad

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\UserPriorityMapping

**Description**
Stores configuration data for the policy setting **Controlled load service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\Layer-2 priority value.

# UserPriorityMapping\ServiceTypeGuaranteed

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\UserPriorityMapping

**Description**
Stores configuration data for the policy setting **Guaranteed service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\Layer-2 priority value.

# UserPriorityMapping\ServiceTypeNetworkControl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\UserPriorityMapping

**Description**
Stores configuration data for the policy setting **Network control service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\Layer-2 priority value.

# ServiceTypeNonConforming

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\UserPriorityMapping

**Description**
Stores configuration data for the policy setting **Non-conforming packets**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\Layer-2 priority value.

# UserPriorityMapping\ServiceTypeQualitative

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Psched\UserPriorityMapping

**Description**
Stores configuration data for the policy setting **Qualitative service type**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\QoS Packet Scheduler\Layer-2 priority value.

# UserPriorityMapping

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\Psched

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# AddAdminGroupToRUP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Add the Administrators security group to roaming user profiles**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# AllowX-ForestPolicy-and-RUP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Allow Cross-Forest User Policy and Roaming User Profiles**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# CompatibleRUPSecurity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Do not check for user ownership of Roaming Profile Folders**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# DeleteRoamingCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Delete cached copies of roaming profiles**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# DenyRsopToInteractiveUser

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Disallow Interactive Users from generating Resultant Set of Policy data**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# DenyUsersFromMachGP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Remove users ability to invoke machine policy refresh**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# DisableCMD

Updated: January 6, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 0 |

**Description**

Stores configuration data for the policy setting **Prevent access to the command prompt**.

| Value | Meaning |
|---|---|
| 0 | Allows users access to the command prompt. |
| 1 | Prevents users from running the interactive command prompt, Cmd.exe. This setting also restricts batch files and logon scripts (.cmd and .bat) from running on the computer.<br>**Caution**<br>If you enable this setting and the user tries to open a command window, the system displays a message explaining that a setting prevents the action. |
| 2 | Prevents users from running the interactive command prompt, Cmd.exe, but allows batch files and logon scripts to run.<br>**Tip**<br>You should use this setting if you run logon, logoff, startup, or shutdown batch file scripts, or if users run Remote Desktop Services (Terminal Services). |

**Change Method**

To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# System\DisableInHelp

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Restrict these programs from being launched from Help**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# System\DisableInHelp Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Restrict these programs from being launched from Help**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# DLT_AllowDomainMode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Allow Distributed Link Tracking clients to use domain resources**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# ExcludeProfileDirs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Exclude directories in roaming profile**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# GroupPolicyMinTransferRate Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy slow link detection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# System\GroupPolicyMinTransferRate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy slow link detection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# GroupPolicyRefreshTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy refresh interval for computers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# GroupPolicyRefreshTimeDC

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy refresh interval for domain controllers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# GroupPolicyRefreshTimeOffset

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy refresh interval for computers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# GroupPolicyRefreshTimeOffsetDC

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Group Policy refresh interval for domain controllers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# HelpQualifiedRootDir

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Restrict potentially unsafe HTML Help functions to specified folders**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# LocalProfile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Only allow local user profiles**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# System\ProfileDlgTimeOut

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Timeout for dialog boxes**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# ProfileErrorAction

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Log users off when roaming profile fails**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# ProfileUnloadTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Maximum retries to unload and update user profile**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# ReadOnlyProfile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Prevent Roaming Profile changes from propagating to the server**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# RSoPLogging

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

**Description**
Stores configuration data for the policy setting **Turn off Resultant Set of Policy logging**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# SlowLinkDetectEnabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Do not detect slow network connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# System\SlowLinkTimeOut

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Slow network connection timeout for user profiles**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# SlowLinkUIEnabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prompt user when slow link is detected**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# UserPolicyMode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **User Group Policy loopback processing mode**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Group Policy.

# UserProfileMinTransferRate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Slow network connection timeout for user profiles**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# System\SlowLinkProfileDefault

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Wait for remote user profile**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\User Profiles.

# Windows\System Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# PromptPasswordOnResume

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\System\Power

**Description**
Stores configuration data for the policy setting **Prompt for password on resume from hibernate / suspend**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Power Management.

# Power

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\System

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Allow Browse

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Task Scheduler5.0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit Browse**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Task Scheduler.

# Task Scheduler5.0\Disable Advanced

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Task Scheduler5.0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Advanced Properties Checkbox in Add Scheduled Task Wizard**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Task Scheduler.

# Disable Advanced Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Task Scheduler5.0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Advanced Properties Checkbox in Add Scheduled Task Wizard**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Task Scheduler.

# DragAndDrop

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Task Scheduler5.0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit Drag-and-Drop**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Task Scheduler.

# Execution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Task Scheduler5.0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prevent Task Run or End**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Task Scheduler.

# Property Pages

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Task Scheduler5.0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide Property Pages**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Task Scheduler.

# Task Creation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Task Scheduler5.0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit New Task Creation**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Task Scheduler.

# Task Deletion

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\Task Scheduler5.0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prohibit Task Deletion**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Task Scheduler.

# Task Scheduler5.0

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

**Description**
The Task Scheduler5.0 subkey stores the entries created when you configure a Group Policy setting for the Scheduled Tasks feature of Windows Server 2003. Group Policy creates and maintains these entries, and Scheduled Tasks reads and interprets them.

This subkey stores entries that apply to all users of the computer. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the Computer Configuration folder in Group Policy.

The Task Scheduler5.0 subkey in HKEY_CURRENT_USER stores entries that are configured separately for each user. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the User Configuration folder in Group Policy.

# Windows\WindowsUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# AU

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows\WindowsUpdate

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# AUOptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate\AU

**Description**
Stores configuration data for the policy setting **Configure Automatic Updates**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# NoAutoRebootWithLoggedOnUsers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate\AU

**Description**
Stores configuration data for the policy setting **No auto-restart for scheduled Automatic Updates installations**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# NoAutoUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate\AU

**Description**
Stores configuration data for the policy setting **Configure Automatic Updates**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# RescheduleWaitTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate\AU

**Description**
Stores configuration data for the policy setting **Reschedule Automatic Updates scheduled installations**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# ScheduledInstallDay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate\AU

**Description**
Stores configuration data for the policy setting **Configure Automatic Updates**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# ScheduledInstallTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate\AU

**Description**
Stores configuration data for the policy setting **Configure Automatic Updates**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# UseWUServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate\AU

**Description**
Stores configuration data for the policy setting **Specify intranet Microsoft update service location**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# WUServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate\AU

**Description**
Stores configuration data for the policy setting **Specify intranet Microsoft update service location**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# WUStatusServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate\AU

**Description**
Stores configuration data for the policy setting **Specify intranet Microsoft update service location**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Update.

# Policies\Microsoft\Windows NT Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

**Description**
The Windows NT subkey contains entries that store Group Policy settings for Windows Server 2003. Group Policy creates and maintains these entries, and the system components read and interpret them.

This subkey stores entries that apply to all users of the computer. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the Computer Configuration folder in Group Policy.

The Policies\Microsoft\Windows NT subkey in HKEY_CURRENT_USER stores policy setting-related entries that are configured separately for each user. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the User Configuration folder in Group Policy.

**Change Method**
To add, delete, or change an entry in this subkey, use Group Policy.

# Windows NT\CurrentVersion Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# MYS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\CurrentVersion

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DisableShowAtLogon

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\CurrentVersion\MYS

**Description**
Stores configuration data for the policy setting **Do not display Manage Your Server page at logon**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# SyncForegroundPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\CurrentVersion\Winlogon

**Description**
Stores configuration data for the policy setting **Always wait for the network at computer startup and logon**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Logon.

# CurrentVersion\Winlogon Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\CurrentVersion

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ApplyToRemovableMedia

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DiskQuota

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Apply policy to removable media**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Disk Quotas.

# Enable

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DiskQuota

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enable disk quotas**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Disk Quotas.

# Enforce

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DiskQuota

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Enforce disk quota limit**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Disk Quotas.

# LimitUnits

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DiskQuota

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Default quota limit and warning level**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Disk Quotas.

# LogEventOverLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DiskQuota

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Log event when quota limit exceeded**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Disk Quotas.

# LogEventOverThreshold

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DiskQuota

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Log event when quota warning level exceeded**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Disk Quotas.

# Threshold

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DiskQuota

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Default quota limit and warning level**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Disk Quotas.

# ThresholdUnits

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DiskQuota

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Default quota limit and warning level**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Disk Quotas.

# Limit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DiskQuota

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Default quota limit and warning level**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Disk Quotas.

# RegisterAdapterName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DNSClient

**Description**
Stores configuration data for the policy setting **Register DNS records with connection-specific DNS suffix**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\DNS Client.

# RegisterReverseLookup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DNSClient

**Description**
Stores configuration data for the policy setting **Register PTR Records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\DNS Client.

# RegistrationEnabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DNSClient

**Description**
Stores configuration data for the policy setting **Dynamic Update**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\DNS Client.

# RegistrationRefreshInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DNSClient

**Description**
Stores configuration data for the policy setting **Registration Refresh Interval**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\DNS Client.

# RegistrationTtl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DNSClient

**Description**
Stores configuration data for the policy setting **TTL Set in the A and PTR records**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\DNS Client.

# SearchList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DNSClient

**Description**
Stores configuration data for the policy setting **DNS Suffix Search List**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\DNS Client.

# UpdateSecurityLevel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DNSClient

**Description**
Stores configuration data for the policy setting **Update Security Level**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\DNS Client.

# UseDomainNameDevolution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\DNSClient

**Description**
Stores configuration data for the policy setting **Primary DNS Suffix Devolution**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\DNS Client.

# Windows NT\DNSClient

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Windows NT\Driver Signing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# BehaviorOnFailedVerify

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Driver Signing

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Code signing for device drivers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# PreventIISInstall

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\IIS

**Description**
Stores configuration data for the policy setting **Prevent IIS installation**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Internet Information Services.

# IIS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Windows NT\Printers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

**Description**
The Printers subkey contains entries that store Group Policy settings for printers. Group Policy creates and maintains these entries, and the printing functions in Windows Server 2003 read and interpret them.

This subkey stores entries that apply to all users of the computer. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the Computer Configuration folder in Group Policy.

The Policies\Microsoft\Windows NT\Printers subkey in HKEY_CURRENT_USER stores policy setting related entries that are configured separately for each user. The system adds these entries when you enable (and, for some policy settings, when you disable) a policy setting in the User Configuration folder in Group Policy.

# DisableWebPrinting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Web-based printing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# Immortal

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Allow pruning of published printers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# KMPrintersAreBlocked

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

**Description**
Stores configuration data for the policy setting **Disallow installation of printers using kernel-mode drivers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# PhysicalLocation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Computer location**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# PhysicalLocationSupport

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Pre-populate printer search location text**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# PruneDownlevel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Prune printers that are not automatically republished**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# PruningInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Directory pruning interval**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# PruningPriority

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Directory pruning priority**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# PruningRetries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Directory pruning retry**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# PruningRetryLog

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

**Description**
Stores configuration data for the policy setting **Log directory pruning retry events**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# PublishPrinters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Allow printers to be published**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# RegisterSpoolerRemoteRpcEndPoint

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

**Description**
Stores configuration data for the policy setting **Allow Print Spooler to accept client connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# ServerThread

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Printer browsing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# SupportLink

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Custom support URL in the Printers folder's left pane**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# VerifyPublishedState

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Check published state**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# PointAndPrint

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# InForest

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\PointAndPrint

**Description**
Stores configuration data for the policy setting **Point and Print Restrictions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Printers.

# NoAddPrinter

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\PointAndPrint

**Description**
Stores configuration data for the policy setting **Prevent addition of printers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Printers.

# NoDeletePrinter

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\PointAndPrint

**Description**
Stores configuration data for the policy setting **Prevent deletion of printers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Printers.

# Restricted

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\PointAndPrint

**Description**
Stores configuration data for the policy setting **Point and Print Restrictions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Printers.

# ServerList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\PointAndPrint

**Description**
Stores configuration data for the policy setting **Point and Print Restrictions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Printers.

# TrustedServers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\PointAndPrint

**Description**
Stores configuration data for the policy setting **Point and Print Restrictions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Printers.

# Auto Publishing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\Wizard

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Automatically publish new printers in Active Directory**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Printers.

# Default Search Scope

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\Wizard

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Default Active Directory path when searching for printers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Printers.

# Downlevel Browse

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\Wizard

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Browse the network to find printers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Printers.

# Printers Page URL

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Printers\Wizard

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Browse a common web site to find printers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Control Panel\Printers.

# Wizard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Printers

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ShutdownReasonOn

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Reliability

**Description**
Stores configuration data for the policy setting **Display Shutdown Event Tracker**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# ShutdownReasonUI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Reliability

**Description**
Stores configuration data for the policy setting **Display Shutdown Event Tracker**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# SnapShot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Reliability

**Description**
Stores configuration data for the policy setting **Activate Shutdown Event Tracker System State Data feature**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# TimeStampEnabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Reliability

**Description**
Stores configuration data for the policy setting **Enable Persistent Time Stamp**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# TimeStampInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Reliability

**Description**
Stores configuration data for the policy setting **Enable Persistent Time Stamp**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# Reliability

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Rpc

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ExtErrorInfoExceptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Rpc

**Description**
Stores configuration data for the policy setting **Propagation of extended error information**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Procedure Call.

# StateInformation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Rpc

**Description**
Stores configuration data for the policy setting **RPC Troubleshooting State Information**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Procedure Call.

# ExtErrorInformation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Rpc

**Description**
Stores configuration data for the policy setting **Propagation of extended error information**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Procedure Call.

# IgnoreDelegationFailure Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Rpc

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Rpc\IgnoreDelegationFailure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Rpc\IgnoreDelegationFailure

**Description**
Stores configuration data for the policy setting **Ignore Delegation Failure**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Procedure Call.

# MinimumConnectionTimeout Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Rpc\MinimumConnectionTimeout

**Description**
Stores configuration data for the policy setting **Minimum Idle Connection Timeout for RPC/HTTP connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Procedure Call.

# MinimumConnectionTimeout Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Rpc

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ServicePackSourcePath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Setup

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Specify Windows Service Pack installation file location**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# SourcePath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Setup

**Description**
Stores configuration data for the policy setting **Specify Windows installation file location**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System.

# Setup Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# SharedFolders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# PublishDfsRoots

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\SharedFolders

**Description**
Stores configuration data for the policy setting **Allow DFS roots to be published**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Shared Folders.

# PublishSharedFolders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\SharedFolders

**Description**
Stores configuration data for the policy setting **Allow shared folders to be published**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Shared Folders.

# ColorDepth

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Limit maximum color depth**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# DeleteTempDirsOnExit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not delete temp folder upon exit**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Temporary folders.

# fAllowFullControl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Solicited Remote Assistance**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Assistance.

# fAllowToGetHelp

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Solicited Remote Assistance**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Assistance.

# fAllowUnsolicited

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Offer Remote Assistance**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Assistance.

# fAllowUnsolicitedFullControl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Offer Remote Assistance**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Assistance.

# fDenyTSConnections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Allow users to connect remotely using Terminal Services**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# fDisableAutoReconnect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Automatic reconnection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# fDisableCam

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Allow audio redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection.

# fDisableCcm

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not allow COM port redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection.

# fDisableCdm

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not allow drive redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection.

# fDisableClip

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not allow clipboard redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection.

# fDisableCpm

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not allow client printer redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection.

# fDisableForcibleLogoff

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Deny log off of an administrator logged in to the console session**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# fDisableLPT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not allow LPT port redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection.

# fEnableSmartCard

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not allow smart card device redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection.

# fEnableTimeZoneRedirection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Allow Time Zone Redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection.

# fEncryptRPCTraffic

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Secure Server (Require Security)**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Encryption and Security\RPC Security Policy.

# fForceClientLptDef

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not set default client printer to be default printer in a session**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Client/Server data redirection.

# fInheritInitialProgram

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Start a program on connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# fNoRemoteDesktopWallpaper

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Enforce Removal of Remote Desktop Wallpaper**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# fPreventLicenseUpgrade

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Prevent License Upgrade**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Licensing.

# fPromptForPassword

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Always prompt client for password upon connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Encryption and Security.

# fReconnectSame

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Allow reconnection from original client only**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Sessions.

# fResetBroken

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Terminate session when time limits are reached**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Sessions.

# fSecureLicensing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **License Server Security Group**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Licensing.

# fSingleSessionPerUser

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Restrict Terminal Services users to a single remote session**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# fUseMailto

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Solicited Remote Assistance**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Assistance.

# fWritableTSCCPermTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not allow local administrators to customize permissions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# InitialProgram

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Start a program on connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# KeepAliveEnable

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Keep-Alive Connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# Terminal Services\KeepAliveInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Keep-Alive Connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# MaxConnectionTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Sets a time limit for active Terminal Services sessions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Sessions.

# MaxDisconnectionTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Set time limit for disconnected sessions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Sessions.

# MaxInstanceCount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Limit number of connections**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# MaxTicketExpiry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Solicited Remote Assistance**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Assistance.

# MaxTicketExpiryUnits

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Solicited Remote Assistance**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Remote Assistance.

# MinEncryptionLevel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Set client connection encryption level**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Encryption and Security.

# PerSessionTempDir

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Do not use temp folders per session**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Temporary folders.

# SessionDirectoryActive

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Join Session Directory**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Session Directory.

# SessionDirectoryClusterName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Session Directory Cluster Name**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Session Directory.

# SessionDirectoryExposeServerIP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Terminal Server IP Address Redirection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Session Directory.

# SessionDirectoryLocation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Session Directory Server**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Session Directory.

# Terminal Services\Shadow

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Sets rules for remote control of Terminal Services user sessions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# WFHomeDir

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **TS User Home Directory**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# WFHomeDirDrive

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **TS User Home Directory**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# WFHomeDirUNC

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **TS User Home Directory**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# WFProfilePath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Set path for TS Roaming Profiles**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# WorkDirectory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Start a program on connection**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services.

# MaxIdleTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT\Terminal Services

**Description**
Stores configuration data for the policy setting **Sets a time limit for active but idle Terminal Services sessions**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Terminal Services\Sessions.

# Terminal Services Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Windows File Protection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\Windows NT

**Description**
The Windows File Protection subkey contains entries that store Group Policy settings for the Windows File Protection feature of Windows Server 2003. Group Policy creates and maintains the entries in this subkey, and the Windows File Protection program reads and interprets them.

# SFCDllCacheDir Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Windows File Protection

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores configuration data for the policy setting **Specify Windows File Protection cache location**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows File Protection.

# SfcQuota

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Windows File Protection

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Limit Windows File Protection cache size**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows File Protection.

# Windows File Protection\SfcScan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Windows File Protection

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Set Windows File Protection scanning**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows File Protection.

# Windows File Protection\SfcShowProgress

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\Windows NT\Windows File Protection

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**
Stores configuration data for the policy setting **Hide the file scan progress window**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\System\Windows File Protection.

# QuickLaunchShortcut

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Prevent Quick Launch toolbar shortcut creation**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player.

# SetAndLockSkin

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Set and Lock Skin**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\User Interface.

# DefaultSkin

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Set and Lock Skin**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\User Interface.

# DesktopShortcut

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Prevent Desktop shortcut creation**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player.

# DisableAutoUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Prevent Automatic Updates**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player.

# DoNotShowAnchor

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Do Not Show Anchor**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\User Interface.

# DontUseFrameInterpolation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Prevent Video Smoothing**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player.

# EnableScreenSaver

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Allow Screen Saver**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Playback.

# GroupPrivacyAcceptance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Disable Setup First Use Configuration Dialogs**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player.

# HideNetworkTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Hide Network Tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# HidePrivacyTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Hide Privacy Tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\User Interface.

# HideSecurityTab

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Hide Security Tab**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\User Interface.

# PreventCDDVDMetadataRetrieval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Prevent CD and DVD media information retrieval**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player.

# PreventCodecDownload

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Prevent Codec Download**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Playback.

# PreventMusicFileMetadataRetrieval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer

**Description**
Stores configuration data for the policy setting **Prevent music file media information retrieval**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player.

# WindowsMediaPlayer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Protocols

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\WindowsMediaPlayer

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# HTTP Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols

**Description**
Stores configuration data for the policy setting **Streaming Media Protocols**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# TCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols

**Description**
Stores configuration data for the policy setting **Streaming Media Protocols**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# UDP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols

**Description**
Stores configuration data for the policy setting **Streaming Media Protocols**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# UDPPorts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols

**Description**
Stores configuration data for the policy setting **Streaming Media Protocols**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# WindowsMediaStreamingProtocols

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols

**Description**
Stores configuration data for the policy setting **Streaming Media Protocols**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# Multicast

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols

**Description**
Stores configuration data for the policy setting **Streaming Media Protocols**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# HTTP\AutodetectProxy Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\HTTP

**Description**
Stores configuration data for the policy setting **Configure HTTP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# BypassProxyLocal Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\HTTP

**Description**
Stores configuration data for the policy setting **Configure HTTP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# BypassProxyLocalAddress Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\HTTP

**Description**
Stores configuration data for the policy setting **Configure HTTP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# ProxyAddress Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\HTTP

**Description**
Stores configuration data for the policy setting **Configure HTTP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# HTTP\ProxyPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\HTTP

**Description**
Stores configuration data for the policy setting **Configure HTTP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# ProxyPort Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\HTTP

**Description**
Stores configuration data for the policy setting **Configure HTTP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# ProxyType Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\HTTP

**Description**
Stores configuration data for the policy setting **Configure HTTP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# UseBrowserProxy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\HTTP

**Description**
Stores configuration data for the policy setting **Configure HTTP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# UseProxy Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\HTTP

**Description**
Stores configuration data for the policy setting **Configure HTTP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# Protocols\HTTP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\WindowsMediaPlayer\Protocols

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# MMS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\WindowsMediaPlayer\Protocols

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# MMS\AutodetectProxy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\MMS

**Description**
Stores configuration data for the policy setting **Configure MMS Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# MMS\BypassProxyLocal

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\MMS

**Description**
Stores configuration data for the policy setting **Configure MMS Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# MMS\BypassProxyLocalAddress

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\MMS

**Description**
Stores configuration data for the policy setting **Configure MMS Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# MMS\ProxyAddress

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\MMS

**Description**
Stores configuration data for the policy setting **Configure MMS Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# MMS\ProxyPolicy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\MMS

**Description**
Stores configuration data for the policy setting **Configure MMS Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# MMS\ProxyPort

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\MMS

**Description**
Stores configuration data for the policy setting **Configure MMS Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# MMS\ProxyType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\MMS

**Description**
Stores configuration data for the policy setting **Configure MMS Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# AutodetectProxy Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\RTSP

**Description**
Stores configuration data for the policy setting **Configure RTSP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# ProxyPolicy Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\RTSP

**Description**
Stores configuration data for the policy setting **Configure RTSP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# RTSP\ProxyType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\RTSP

**Description**
Stores configuration data for the policy setting **Configure RTSP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# RTSP\UseProxy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\Microsoft\WindowsMediaPlayer\Protocols\RTSP

**Description**
Stores configuration data for the policy setting **Configure RTSP Proxy**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Windows Components\Windows Media Player\Networking.

# RTSP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\Microsoft\WindowsMediaPlayer\Protocols

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# SNMP Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# SNMP\Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\SNMP

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# PermittedManagers Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\SNMP\Parameters

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# PermittedManagers\n

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\SNMP\Parameters\PermittedManagers

**Description**
Stores configuration data for the policy setting **Permitted Managers**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\SNMP.

# TrapConfiguration Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\SNMP\Parameters

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# public

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\SNMP\Parameters\TrapConfiguration

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# public\n

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\SNMP\Parameters\TrapConfiguration\public

**Description**
Stores configuration data for the policy setting **Traps for public community**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\SNMP.

# ValidCommunities\n Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\Software\Policies\SNMP\Parameters\ValidCommunities

**Description**
Stores configuration data for the policy setting **Communities**.

**Change Method**
To change the value of this entry, use the Group Policy Object Editor (Gpedit.msc). The corresponding policy is located in Administrative Templates\Network\SNMP.

# ValidCommunities Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE\Policies\SNMP\Parameters

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Program Groups Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SOFTWARE

**Description**
The Program Groups subkey remains in the registry. However, Windows Server 2003 does not include Program Manager, the component that used this subkey.

In Windows NT 4.0 and earlier, the Program Groups subkey stored the contents of program groups the user defined in Program Manager. In Windows 2000 and Windows Server 2003, these Program Groups appear as shortcuts in subfolders of the **Start** menu.

**Note**

- If a user or program attempts to start Program Manager, Windows Server 2003 starts Windows Explorer instead.

# SYSTEM

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM

**Description**
The subkeys in the HKEY_LOCAL_MACHINE\SYSTEM key store entries for the current control set and for other control sets that can be used, or have been used, to run Windows Server 2003. Windows Server 2003 maintains at least two control sets to ensure that the system always starts.

# ControlSet00n

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM

**Description**
The *ControlSet00n* subkeys represent a control set for the system.

The numbered ControlSet00n subkeys, such as ControlSet001 and ControlSet002, contain control sets that can be used to start and run Windows Server 2003. Most systems have two numbered control sets (an original and a backup copy) of a control set that has been used to start the system successfully, but the system can maintain as many as four control sets.
Windows Server 2003 maintains backup copies so that you can undo configuration changes that might otherwise prevent you from starting the system.

Windows Server 2003 identifies the control set that was last used successfully, the current control set, and the default control set. These designations are stored in the Select subkey.

Windows Server 2003 automatically maintains the control sets and selects one of them when the system starts. It bases this selection on settings in the Select subkey and on the success or failure of system startup. Users can specify the LastKnownGood control set by pressing the F8 key on the **OS Choices Menu** during system startup to display the **Advanced Startup Options** menu.

All the numbered control set subkeys, including the one that appears in the CurrentControlSet subkey, can have the same subkeys and entries. In this reference, these subkeys are entries listed in the CurrentControlSet subkey.

**Note**

- The *ControlSet00n* subkey is a variable representing the numbered control sets in the HKLM\SYSTEM key. It does not actually appear in the registry.

**See Also**
**Concepts**
Select
CurrentControlSet

# CurrentControlSet

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM

**Description**
The CurrentControlSet subkey does not actually contain a control set. The subkeys and entries that appear in the CurrentControlSet subkey are the contents of one of the numbered control sets, such as ControlSet001. The control set that appears in CurrentControlSet is selected by the system at startup and is identified by the value of the **Current** entry in the Select subkey.

For example, if the value of the **Current** entry is 1, then the contents of ControlSet001 appear in CurrentControlSet.

# Control

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet

**Description**
The Control subkey and its subkeys contain entries that control system startup, such as the subsystems to load and the size and location of paging files.

**Activation Method**
You must restart the system before any changes in the Control subkey take effect.

# RegistrySizeLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**

This entry is not used by Windows Server 2003. In older versions of Windows Server, it specified the maximum amount of memory and disk space that the registry could consume.

# BackupRestore

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The BackupRestore subkey contains subkeys that determine which files the system does not back up and which registry subkeys and entries the system does not restore.

**Caution**

- Do not change the values of these entries. Altering these values can prevent Windows Backup from successfully backing up and/or restoring data.

# FilesNotToBackup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore

**Description**
The FilesNotToBackup subkey contains entries that specify the files that Backup omits.

Each of the entries in this subkey has the following format:

| Entry name | Data type | Value |
|---|---|---|
| *component* | REG_MULTI_SZ | *[drive][path]\filename* **[/s]** |

- *Drive* specifies the drive and is optional. For example, c:
- *Path* specifies the path and is optional. It cannot contain wildcard characters.
- *Filename* specifies the file and is required. It can contain wildcard characters.
- **/s** specifies that the file affects all subdirectories of the specified path. Environment variables can be substituted for all or part of the above.

| Entry | Default Value |
|---|---|
| Client Side Cache | %Systemroot%\csc* /s |
| ComPlus | %Systemroot%\registration*.clb |
| Internet Explorer | %UserProfile%\index.dat /s |
| Memory Page File | \Pagefile.sys |
| MS Distributed Transaction Coordinator | %Systemroot%\WINNT\System32\DTCLog\MSDTC.LOG |
| Power Options | \hiberfil.sys |
| Single Instance Storage | \SIS Common Store*.* /s |
| Temporary Files | %TEMP%* /s |

**Change Method**
To change the value of these entries, use Backup. From the **Tools** menu, click **Options**, then click the **Exclude Files** tab. Use the **Add new**, **Edit**, and **Remove** buttons to select those files not to be backed up.

**Caution**

- Do not change the values of the entries in this subkey. Altering these values can prevent Backup from successfully backing up or restoring data.

# KeysNotToRestore

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore

**Description**
The KeysNotToRestore subkey specifies the names of those registry subkeys and entries that Backup will not restore.

Each of the entries has the following format:

| Entry name | Data type | Value |
|---|---|---|
| *component* | REG_MULTI_SZ | *registry path* [*registry path*] |

The paths are relative from HKEY_LOCAL_MACHINE\SYSTEM and can take one of three forms:

- A path ending with "\" indicates that neither the subkey nor any of its subkeys or entries will be restored from the backup. Instead, the current values will be preserved.
- A path ending with "*" indicates that neither the subkey nor any of its subkeys and entries will be restored from the backup, but new entries will be merged from the backup.
- A path ending with an entry name indicates that the entry will not be restored from the backup.

| Entry | Default Value |
|---|---|
| Active Directory Restore | CurrentControlSet\Services\NTDS\Restore In Progress\<br>CurrentControlSet\Services\NTDS\Parameters\New Database GUID |
| Fault Tolerance | Disk\ |
| Installed Services | CurrentControlSet\Services* |
| LDM Boot Information | CurrentControlSet\Control\Services\dmio\boot info\ |
| Mount Manager | MountedDevices\ |
| Pending Rename Operations | CurrentControlset\Control\Session Manager\PendingFileRenameOperations |
| Plug & Play | CurrentControlSet\Session Manager\PendingFileRenameOperations |
| Session Manager | CurrentControlSet\Control\SessionManager\AllowProtectedRename |
| Windows Setup | Setup\SystemPartition |

**Caution**

- Do not change the values of the entries in this subkey. Altering these values can prevent Backup from successfully backing up or restoring data.

# BootVerificationProgram

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The BootVerificationProgram subkey stores data about custom startup verification programs. Startup verification programs specify the criteria for a successful startup and notify the system that the startup is successful—two functions otherwise performed by Winlogon. You can write a custom startup verification program and use this subkey to direct the Windows Server 2003 service controller to run it.

Some custom startup verification programs notify the system that startup was not successful, causing the system to restart using the LastKnownGood control set. Other programs determine whether their criteria were satisfied, and if so, they run Bootok.exe, a program that notifies the system that startup was successful. Bootok.exe is included in Windows Server 2003.

### Notes

- To run a custom startup verification program, you must disable the standard startup verification functions in Winlogon by setting the value of the ReportBootOk entry to 0.
- You cannot run a custom startup verification program if you are also running the Bootvrfy startup verification program.

### Caution

- Do not add or change the value of the entries in the BootVerificationProgram subkey unless you need to run a custom verification program.

**See Also**
**Concepts**
ReportBootOk

# BootVerificationProgram\ImagePath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\BootVerificationProgram

| Data type | Range | Default value |
|---|---|---|
| REG_SZ \| REG_EXPAND_SZ | Filename | (None) |

**Description**
Specifies the full path and file name of a custom startup verification program. At startup, the Windows Server 2003 service controller runs the program listed here.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- To run a custom startup verification program, you must disable the standard startup verification functions in Winlogon by setting ReportBootOk to 0.

**See Also**
**Concepts**
ReportBootOk

# Class

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The Class subkey stores configuration data for classes of hardware devices.

Each subkey in the Class subkey represents a device class and is named by the globally unique identifier (GUID) for the class. It appears as a 16-digit alphanumeric string within curly braces. These {*Class-GUID*} subkeys store configuration data that applies to the entire class of devices, such as the class of all modems.

Within the {*Class-GUID*} subkeys are device-specific subkeys, which represent particular devices that are instances of the class.

Because the Class subkey is hardware specific, its contents vary widely. The Class subkey on your computer might look quite different from the attributes described in this reference.

**Tip**

- To determine the type of hardware device represented by a {*Class-GUID*} subkey, use the registry editor Regedit.exe. Click the name of the {*Class-GUID*} subkey (it appears as a 16-byte alphanumeric string within brackets). The device type appears in the value of the first, unnamed entry in the subkey.

- To display hidden devices, non–Plug and Play devices, and devices not attached to the computer (commonly known as "ghosted" or "phantom" devices), type `set DEVMGR_SHOW_NONPRESENT_DEVICES=1` at the command line. Then, use Device Manager to remove or reconfigure these devices. Do not edit the registry. For more information about viewing hidden devices, see *Microsoft Windows XP Professional Resource Kit Documentation*.

**Caution**

- Do not change any entries in the Class subkey or any subkey within it. These subkeys are configured by the system to comply with Plug and Play standards. Changing any data can prevent your computer devices from starting or operating properly.

# (Network adapters) {4D36E972-E325-11CE-BFC1-08002BE10318}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# WanEndpoints

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class\(Network adapters) {4D36E972-E325-11CE-BFC1-08002BE10318}\WAN Miniport (L2TP) *Device-number (000n)*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x5 |

**Description**
Stores the number of wide area network (WAN) endpoints currently configured for use.

This entry stores the setting in the **Maximum ports** box in Routing and Remote Access.

**Change Method**
To change the value of this entry, use Routing and Remote Access. Right-click **Ports**, click **Properties**, click **WAN Miniport (L2TP)**, and then click **Configure**. This entry stores the value in the **Maximum ports** box.

### Note

- The value of this entry must be less than or equal to the value of the entry **MaxWanEndpoints**.

### Caution

- Do not change the value of this entry by editing the registry. If you do, the system cannot adjust corresponding values, and Layer Two Tunneling Protocol (L2TP) might not support the number of endpoints you request.

# {Class-GUID}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class

**Description**
The {*Class-GUID*} subkey represents the component associated with the class-GUID.

Each {*Class-GUID*} subkey in the Class subkey represents a device class and is named by the GUID for the class. It appears as a 16-digit alphanumeric string within curly braces. These {*Class-GUID*} subkeys store configuration data that applies to the entire class of devices, such as the class of all modems.

Within the {*Class-GUID*} subkeys are 4-digit numbered subkeys, which represent particular devices that are instances of the class. These subkeys store data that apply only to that device.

Because the Class subkey is hardware specific, its contents vary widely. The Class subkey on your computer might look quite different from the attributes described in this reference.

**Change Method**
It is strongly recommended that you do not make changes to entries in the Class subkey or any subkey within it.

### Note

- The {*Class-GUID*} subkey is a variable representing the subkeys in the HKLM\SYSTEM\CurrentControlSet\Control\Class subkey. It does not actually appear in the registry. This variable subkey displays the entries and subkeys that are common to many of the subkeys in the Class subkey. The actual Class subkeys in your registry depend on the hardware installed on the computer or available to the system.

### Tip

- To determine the type of hardware device represented by a {*Class-GUID*} subkey, use the registry editor Regedit.exe. Click the name of the class-GUID subkey (it appears as a 16-byte alphanumeric string within brackets). The device type appears in the value of the first, unnamed entry in the subkey.

### Caution

- Do not change any entries in the Class subkey or any subkey within it. These subkeys are configured by the system to comply with Plug and Play standards. Changing any data can prevent your computer devices from starting or operating properly.

# Device-number (000n) Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class*{Class-GUID}*

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Linkage Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class*{Class-GUID}**Device-number (000n)*

**Description**
The Linkage subkey stores configuration data for binding network components. It exists only for net class (adapters) devices that have a class globally unique identifier (GUID) of {4D36E972-E325-11CE-BFC1-08002bE10318}. The bindings associate devices and services with the transport protocols they use.

A binding consists of three entries: **UpperBind**, **Export**, and **RootDevice**. Each of these entries can include more than one component. The entries list the components in a precise order, so the first component in one entry corresponds with the first component in the other entries, and each subsequent component corresponds with subsequent components in the other entries.

In Windows NT 4.0 and earlier, all Linkage subkeys were stored in subkeys of the Services subkey (HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services*Service-name*). In Windows Server 2003, the Linkage subkeys for services are still in the Services subkey, but Linkage subkeys for devices are stored in HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Class*{Class-GUID}*, in the numeric subkey representing the particular device.

### Notes

- A valid binding requires a value for three entries: **UpperBind**, **Export**, and **RootDevice**.
- This information is provided only for reference. Use only published application programming interfaces. Do not include the Linkage subkey in Windows programs. This subkey might not be included in future versions of Windows.

### Tip

- To determine the type of hardware device represented by a class-GUID subkey, use the registry editor Regedit.exe. Click the name of the class-GUID subkey (it appears as a 16-byte alphanumeric string within brackets). The device type appears in the value of the first, unnamed entry in the subkey.
- Many digits in the GUID are unique to your system. However, the three digits preceding the first hyphen are specific to the device class and are the same on all systems. For example, the three digits preceding the first hyphen in the DVD/CD-ROM drives class are always 965.

### Caution

- Do not change any entries in the Class subkey or any subkey within it. The system configures these subkeys to comply with Plug and Play standards. Changing any data can prevent your computer devices from starting or operating properly.
- This value is maintained by the Network Configuration subsystem. Do not specify this value in an .inf file.

# Ndi Subkeys

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class*{Class-GUID}\Device-number (000n)*

**Description**
The Ndi subkeys store data for network driver installers (NDI). They appear only in the subkey for the Class GUID {4D36E972-E325-11CE-BFC1-08002bE10318}. They store data that the Network Class Installer uses to install and configure the component.

The entries in the Ndi subkeys vary with the requirements of the component. The entries described in this subkey are common to many components.

### Note

- Windows Server 2003 does not include NetRules subkeys. In Windows NT 4.0 and earlier, the NetRules subkeys contained entries that specified bind rules for network software components. Some bind specifications are stored in the Ndi subkeys that contain configuration data for network driver installers. However, these are primary upper bind and lower bind settings only.

**See Also**
**Concepts**
Linkage Subkey

# {4D36E972-E325-11CE-BFC1-08002BE10318}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class

**Description**
The {4D36E972-E325-11CE-BFC1-08002BE10318} subkey represents the class of network adapter devices that the system supports.

This subkey contains several 4-digit numbered subkeys, which represent particular network adapters. These subkeys store data that applies only to that device.

**Caution**

- Do not change any entries in the Class subkey or any subkey within it. The system configures these subkeys to comply with Plug and Play standards. Changing any data can prevent your computer devices from starting or operating properly.

# {4D36E972-E325-11CE-BFC1-08002bE10318}\Device-number (000n)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E972-E325-11CE-BFC1-08002bE10318}

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Ndi Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E972-E325-11CE-BFC1-08002bE10318}*Device-number (000n)*

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Device-number (000n)\Interfaces

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E972-E325-11CE-BFC1-08002bE10318}*Device-number (000n)*\Ndi

**Description**
The Interfaces subkey specifies the upper and lower binding interfaces for the component. It has two values:

- **Upper Range**. A REG-SZ value which is a comma-delimited string of interface names to which the compoonent can bind at its top edge.

- **Lower Range**. A REG-SZ value which is a comma-delimited string of interface names to which the component can bind at its lower edge.

# UpperRange

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E972-E325-11CE-BFC1-08002bE10318}*Device-number (000n)*\Ndi\Interfaces

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *List of interfaces* | (None) |

**Description**
Specifies the interfaces to which the driver can bind at the upper edge. An upper interface is a higher-level (farther from hardware) component with which the driver communicates.

### Note

- Windows Server 2003 does not include NetRules subkeys. In Windows NT 4.0 and earlier, the NetRules subkeys contained entries that specified bind rules for network software components. Some bind specifications are stored in the Ndi subkeys that contain configuration data for network driver installers. However, these are primary upper bind and lower bind settings only.

**See Also**
**Concepts**
Linkage Subkey

# CrashControl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CrashControl**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The Crash Control subkey stores entries that specify what happens when the system locks, fails, or terminates abnormally.

**Change Method**
To change the values of entries in this subkey, use System in Control Panel. Right-click **My Computer**, click **Properties**, click the **Advanced** tab, and then click **Startup and Recovery**.

# AutoReboot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AutoReboot**
HKLM\SYSTEM\CurrentControlSet\Control\CrashControl

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the system restarts automatically after it fails or locks.

This entry stores the setting of the **Automatically reboot** check box in System in Control Panel.

| Value | Meaning |
|---|---|
| 0 | Do not restart automatically. |
| 1 | Restart automatically. |

**Change Method**
To change the value of this entry, use the System in Control Panel. Click the **Advanced** tab, then click the **Startup and Recovery** button. When the **Automatically reboot** box is selected, the value of this entry is 1.

**Activation Method**
To make changes to this value effective, you must restart Windows Server 2003.

### Note

- To change the value of this entry, you must be an administrator of the computer.

# CrashDumpEnabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CrashDumpEnabled**
HKLM\SYSTEM\CurrentControlSet\Control\CrashControl

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 \| 3 | 3 |

**Description**
Specifies whether Windows Server 2003 writes the contents of the system memory to a log file when the system stops unexpectedly. The contents of this log file can be useful in determining why the system stopped.

This entry stores the setting of the **Write debugging information** box in System in Control Panel. The name and location of the file appear in the DumpFile entry.

| Value | Meaning |
|---|---|
| 0 | Debugging information is not written to a file. |
| 1 | Complete crash dump is written to a file. |
| 2 | Kernel memory dump is written to a file. |
| 3 | Small memory dump is written to a file. |

**Change Method**
To change the value of this entry, use the System in Control Panel. Click the **Advanced** tab, then click **Startup and Recovery**. Select one of the options from the **Write Debugging Information** box.

**Activation Method**
You must restart Windows to make changes to this entry effective.

### Note

- To change the value of this entry, you must be an administrator of the computer.

**See Also**
**Concepts**
DumpFile
Overwrite

# DumpFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DumpFile**
HKLM\SYSTEM\CurrentControlSet\Control\CrashControl

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %SystemRoot%\MEMORY.DMP |

**Description**
Specifies the location of the recovery log file. If the system stops abnormally, and the value of the **CrashDumpEnabled** entry is either 1 or 2, then Windows Server 2003 copies the contents of system memory to the recovery log file.

If you change the value of this entry, then Windows Server 2003 does not move any existing recovery log files. Instead, the next time the system stops abnormally, it creates a recovery log file in the new location.

This entry stores the value of the **Dump File** text box in **Write debugging information** in System in Control Panel.

**Change Method**
To change the value of this entry, use System in Control Panel. Click the **Advanced** tab, then click **Startup and Recovery**. Under **Write debugging information**, in the **Dump File** text box, enter the location of the log file.

**Activation Method**
You must restart Windows to make changes to this entry effective.

### Note

- To change the value of this entry, you must be an administrator of the computer.

### Tip

- For best performance, the recovery log file should be on the system drive. This prevents the crashdump engine from copying the file from the system drive, where it is created, to another drive.

**See Also**
**Concepts**
CrashDumpEnabled
Overwrite

# LogEvent

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**LogEvent**
HKLM\SYSTEM\CurrentControlSet\Control\CrashControl

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the system writes an error to the System Log in Event Viewer when Windows Server 2003 terminates abnormally.

This entry stores the setting of the **Write an event to the system log** check box in System in Control Panel.

| Value | Meaning |
|---|---|
| 0 | No events are logged when the system terminates abnormally. |
| 1 | The system writes an error message to the System Log when the system terminates abnormally. |

**Change Method**
To change the value of this entry, use System in Control Panel. Click the **Advanced** tab, then click **Startup and Recovery**. When the **Write an event to the system log** box is selected, the value of this entry is 1.

**Activation Method**
You must restart Windows to make changes to this entry effective.

### Note

- To change the value of this entry, you must be an administrator of the computer.

# MinidumpDir

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MinidumpDir**
HKLM\SYSTEM\CurrentControlSet\Control\CrashControl

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %SystemRoot%\Minidump |

**Description**
Specifies the location of the minidump recovery log file. If the system stops abnormally, and the value of the **CrashDumpEnabled** entry is 1, 2, or 3, then Windows Server 2003 stores the minidump recovery log file at the location specified in the value of this entry.

**Activation Method**
You must restart Windows to make changes to this entry effective.

### Note

- To change the value of this entry, you must be an administrator of the computer.

# NMICrashDump

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NMICrashDump**
HKLM\SYSTEM\CurrentControlSet\Control\CrashControl

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the system performs standard software error processing when a hardware error causes a nonmaskable interrupt (NMI).

This entry is used only on computers with x86 architecture, such as Intel (™) processor platforms. Other systems always perform standard software error processing when hardware interrupts are detected.

By default, an NMI interrupt stops the system and displays a hardware-malfunction message. If the value of this entry is 1, then the system also performs standard software error processing. The software error processing displays a Stop (*blue screen*) message containing information about the state of the operating system and of other programs when the NMI interrupt occurred.

If the value of the CrashDumpEnabled entry is also 1, then the standard error processing routine writes the contents of system memory to a recovery log file.

| Value | Meaning |
|---|---|
| 0 | NMI interrupts do not trigger standard software error processing. Only a hardware-malfunction message is displayed. |
| 1 | NMI interrupts display a hardware-malfunction message and also trigger standard software error processing, including a Stop (*blue screen*) message. |

This entry is designed to suppress the Stop message when a hardware error causes the system to stop. The text display associated with the Stop message might fill the screen and make it difficult to find the hardware malfunction message. Also, some processors cannot process the high volume of error messages associated with an NMI interrupt and a Stop message without error. For more information, see the Microsoft Knowledge Base link on the Web Resources page. Search the Knowledge Base for Article Q174830.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make changes to this value effective, you must restart Windows.

### Note

- To change the value of this entry, you must be an administrator of the computer.

**See Also**
**Concepts**
CrashDumpEnabled

# Overwrite

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Overwrite**
HKLM\SYSTEM\CurrentControlSet\Control\CrashControl

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the system writes over the old recovery file or creates a new recovery file each time the system stops abnormally.

When the system stops abnormally and the value of CrashDumpEnabled is either 1 or 2, Windows Server 2003 copies the contents of system memory to the recovery log file. The data in this file can help determine the cause of the failure. Note that minidumps are never overwritten, so this entry affects only Full and Kernel dumps.

This entry stores the setting of the **Overwrite any existing file** check box in System in Control Panel.

| Value | Meaning |
|---|---|
| 0 | The system creates a new recovery file for each abnormal stop. This value preserves a record of previous abnormal stops to aid your investigation into the cause of the stops. |
| 1 | The system overwrites the existing recovery file. This value prevents the system from accumulating unused recovery files. |

**Change Method**
To change the value of this entry, use System in Control Panel. Click the **Advanced** tab, then click the **Startup and Recovery** button. When the **Overwrite any existing file** check box is selected, the value of this entry is 1.

**Activation Method**
To make changes to this value effective, you must restart Windows Server 2003.

### Note

- To change the value of this entry, you must be an administrator of the computer.

**See Also**
**Concepts**
CrashDumpEnabled
DumpFile

# SendAlert

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SendAlert**
HKLM\SYSTEM\CurrentControlSet\Control\CrashControl

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the system alerts administrators of the computer when the system stops unexpectedly.

This entry stores the value of the **Send an administrative alert** check box in System in Control Panel.

| Value | Meaning |
|---|---|
| 0 | The system does not alert administrators when the system stops. |
| 1 | The system alerts administrators when the system stops. |

**Change Method**
To change the value of this entry, use System in Control Panel. Click the **Advanced** tab, then click the **Startup and Recovery** button. When the **Send an administrative alert** check box is selected, the value of this entry is 1.

**Activation Method**
To make changes to this value effective, you must restart Windows Server 2003.

### Note

- To change the value of this entry, you must be an administrator of the computer.

# CriticalDeviceDatabase

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CriticalDeviceDatabase**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The CriticalDeviceDatabase subkey stores the critical device database. This database stores configuration data for new devices that must be installed and started before the Windows Server 2003 components that normally install devices are started.

# FileSystem

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FileSystem**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The FileSystem subkey stores entries for file systems that Windows supports.

# NtfsAllowExtendedCharacterIn8dot3Name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NtfsAllowExtendedCharacterIn8dot3Name**
HKLM\SYSTEM\CurrentControlSet\Control\FileSystem

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the characters from the extended character set, including diacritic characters, can be used in short file names using the 8.3 naming convention on NTFS volumes.

| Value | Meaning |
|---|---|
| 0 | On NTFS volumes, file names using the 8.3 naming convention are limited to the standard ASCII character set (minus any reserved values). |
| 1 | On NTFS volumes, file names using the 8.3 naming convention may use extended characters. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# NtfsDisable8dot3NameCreation

Updated: November 6, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

**NtfsDisable8dot3NameCreation**
HKLM\SYSTEM\CurrentControlSet\Control\FileSystem

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 \| 3 | 2 |

**Description**
Specifies whether NTFS generates a short name in the 8.3 naming convention for long file names and for file names that contain characters from the extended character set. If the value of this entry is 0, then files can have two names: the name that the user specifies and the short name that NTFS generates. If the name that the user specifies conforms to the 8.3 naming convention, then NTFS does not generate a short name.

Changing this value does not change the file, but it does change the way that NTFS displays and manages the file. Also, files are named according to whatever rule is specified by this entry at the time of their creation; changing this entry does not alter the names of existing files.

| Value | Meaning |
|---|---|
| 0 | NTFS creates short file names. This setting enables applications that cannot process long file names and computers that use differentcode pages to find the files. |
| 1 | NTFS does not create short file names. Although this setting increases file performance, applications that cannot process long file names, and computers that use different code pages, might not be able to find the files. |
| 2 | NTFS sets the 8.3 naming convention creation on a per volume basis. |
| 3 | NTFS disables 8dot3 name creation on all volumes except the system volume. |

**Modifying the 8.3 naming convention settings on your system**
To disable/enable 8.3 naming convention, users can issue the set command as follows:

```
C:\Windows\system32>fsutil 8dot3name set /?
usage : set [0 through 3] | [<Volume Path> 1 | 0]
```

When a volume is specified the operation updates the on disk flag.

| Note |
|---|
| This operation is only meaningful if the registry value is set to 2. |

**Sample commands**
The following command disables 8dot3 name creation on all volumes:

```
fsutil 8dot3name set 1
```

The following command disables 8dot3 name creation on the C drive:

```
fsutil 8dot3name set C: 1
```

**Activation Method**
If you are using Windows Server 2003 or Windows Server 2003 R2, you must restart Windows to make changes to this entry

effective. If you are using Windows Server 2008 or Windows Server 2008 R2, then you do not have to restart Windows to make the changes to this entry effective.

# NtfsDisableLastAccessUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NtfsDisableLastAccessUpdate**
HKLM\SYSTEM\CurrentControlSet\Control\FileSystem

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether NTFS updates the last-accessed timestamp of a file when that file is opened.

Because updating the last-accessed timestamp requires writing data to the disk, an activity that accesses many files might be faster if this type of update is disabled. However, some applications may require that files have an accurate last-accessed timestamp.

| Value | Meaning |
|---|---|
| 0 | NTFS updates the last-accessed timestamp of a file whenever that file is opened. |
| 1 | NTFS does not update the last-access timestamp of a file when that file is opened. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# NtfsEncryptionService

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NtfsEncryptionService**
HKLM\SYSTEM\CurrentControlSet\Control\FileSystem

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Efs | Efs |

**Description**
Windows Server 2003 does not use this entry. It was used in Windows 2000 to specify which encryption service the NTFS file system used. Encrypting File System (EFS) is the encryption service built into Windows Server 2003.

# Win95TruncatedExtensions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Win95TruncatedExtensions**
HKLM\SYSTEM\CurrentControlSet\Control\FileSystem

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether NTFS and FAT generate file names for new files with the 8.3 naming convention.

Setting the value of this entry to 1 does not change any existing file name extensions, nor does it change the way these extensions are displayed or managed by Find, File Manager, or Windows Explorer. However, it causes NTFS and FAT to generate short names for new files, and to truncate the third character of file name extensions.

The following table shows how "filename.html" would be truncated:

| Value | Short name generated |
|---|---|
| 0 | filena~1.ht~ |
| 1 | filena~1.htm |

The following table provides an example of how this entry changes the file system's response to commands entered in a command prompt window.

| Value | Command: del *.htm | Command: del *.html |
|---|---|---|
| 0 | Deletes only files with extension .htm. | Deletes only files with extension .html. |
| 1 | Deletes files with extensions beginning with .htm, including *.htm and *.html but not *.ht~. | Deletes only files with extension .html. |

**Activation Method**
You must restart the computer for changes to this entry to take effect.

# GraphicsDrivers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**GraphicsDrivers**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DCI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DCI**
HKLM\SYSTEM\CurrentControlSet\Control\GraphicsDrivers

**Description**
Stores configuration data for graphics drivers, including those implementing DirectDraw for Windows Server 2003.

# Timeout Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Timeout**
HKLM\SYSTEM\CurrentControlSet\Control\GraphicsDrivers\DCI

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Seconds* | 7 |

**Description**
Specifies how long a DirectDraw application can keep a graphics-device frame-buffer locked while a dynamic-mode or security-desktop change is pending.

When the time specified in the value of this entry expires, the system requires the DirectDraw application to release the frame buffer. Requiring the release prevents the DirectDraw application from locking and thus not releasing its resources. However, requiring the release might cause an access violation error.

| Value | Meaning |
|---|---|
| 0 | DirectDraw acceleration is disabled and direct frame-buffer access is not permitted. However, DirectDraw continues to run in emulation mode. |
| *Seconds* | The number of seconds the system waits before requiring a DirectDraw application to release a frame buffer. |

**Activation Method**
Changes made to this entry take effect when you change video modes with **Display Properties** and when you restart the computer.

# GroupOrderList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**GroupOrderList**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
This subkey stores the order in which services in a service group are loaded when the system starts.

Services in a service group are assigned a tag. A *tag* is a numeric value that is unique within a service group. When the system loads a service group, it loads the services in the order in which their tags appear in the **GroupOrderList** entry for that service group.

Each entry in the GroupOrderList subkey represents a service group. The value of the entry contains the number of tags, followed by a series of tags in a specified order. All entries have the same format:

| | | |
|---|---|---|
| *Service group name* | REG_BINARY | *Number of tags Tag [Tag ...]* |

For example,

| | | |
|---|---|---|
| PointerPort | REG_BINARY | 03000000020000000100000003000000 |

The value in this example determines the order in which the three services in the PointerPort service group are loaded. The first significant digit (3) indicates that the value of the entry contains three tags. The remaining significant digits (2, 1, and 3) indicate that the service assigned tag 2 will load first in that group, followed by the service assigned tag 1, and then the service assigned tag 3.

**Tip**

- To find the services in any service group, use the registry editor Regedit.exe to open HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services. On the **Edit** menu, click **Find**, type the name of the service group, and select only the **Data** check box. Press F3 to find other services in the service group.

- To find the service group and tag number of any service, open the subkey representing the service in \CurrentControlSet\Services. Click **Group** to find its service group, and click **Tag** to find its tag number.

# hivelist

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**hivelist**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The hivelist subkey contains the locations of files that store the permanent keys of the registry. Permanent keys are not recreated each time the system starts. By default, the files for all permanent keys except those in HKEY_CURRENT_USER are in *Systemroot*\System32\Config. The HKEY_CURRENT_USER files are stored in the subdirectories of *System-drive*\Documents and Settings.

# \REGISTRY\MACHINE\SYSTEM

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**\REGISTRY\MACHINE\SYSTEM**
HKLM\SYSTEM\CurrentControlSet\Control\hivelist

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Path | \Device\Harddisk 0\Partition1\WINNT\System32\Config\System |

**Description**
Specifies the location of the files that store the HKEY_LOCAL_MACHINE\SYSTEM key.

# \REGISTRY\MACHINE\HARDWARE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**\REGISTRY\MACHINE\HARDWARE**
HKLM\SYSTEM\CurrentControlSet\Control\hivelist

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Path | (None) |

**Description**
This entry has no value because the HKEY_LOCAL_MACHINE\Hardware key is recreated each time the system starts. The contents of this entry are not saved in a file.

# \REGISTRY\MACHINE\SAM

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**\REGISTRY\MACHINE\SAM**
HKLM\SYSTEM\CurrentControlSet\Control\hivelist

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Path | \Device\Harddisk 0\Partition1\WINNT\System32\Config\SAM |

**Description**
Specifies the location of files that store the HKEY_LOCAL_MACHINE\SAM key.

# \REGISTRY\MACHINE\SECURITY

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**\REGISTRY\MACHINE\SECURITY**
HKLM\SYSTEM\CurrentControlSet\Control\hivelist

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Path | \Device\Harddisk 0\Partition1\WINNT\System32\Config\SECURITY |

**Description**
Specifies the location of the files that store the HKEY_LOCAL_MACHINE\Security key.

# \REGISTRY\MACHINE\SOFTWARE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**\REGISTRY\MACHINE\SOFTWARE**
HKLM\SYSTEM\CurrentControlSet\Control\hivelist

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Path | \Device\Harddisk 0\Partition1\WINNT\System32\Config\Software |

**Description**
Specifies the location of the files that store the HKEY_LOCAL_MACHINE\SOFTWARE key.

# \REGISTRY\USER\.DEFAULT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**\REGISTRY\USER\.DEFAULT**
HKLM\SYSTEM\CurrentControlSet\Control\hivelist

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Path | Device\Harddisk 0\Partition1\WINNT\System32\Config\Default |

**Description**
Specifies the location of the files that store the HKEY_USERS\.DEFAULT key.

# \REGISTRY\USER\Security ID number

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**\REGISTRY\USER\ Security ID number**
HKLM\SYSTEM\CurrentControlSet\Control\hivelist

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Path | \Device\Harddisk #\Documents and Settings*Username*\NTUSER.DAT |

**Description**
Specifies the location of the files that store the current user profile.

### Note

- The \REGISTRY\USER*Security ID number* entry is a variable representing the entries in the hivelist subkey. It does not actually appear in the registry. This variable entry displays information that is common to the entries in the hivelist subkey.

# IDConfigDB

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IDConfigDB**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The IDConfigDB subkey stores configuration data that identifies the current hardware configuration of the system.

# CurrentConfig

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CurrentConfig**
HKLM\SYSTEM\CurrentControlSet\Control\IDConfigDB

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 1 |

**Description**
Identifies the hardware profile currently in use on the computer.

The value of this entry corresponds to the number of one of the *Profile-number* subkeys in the CurrentControlSet\Hardware Profiles subkey of HKEY_LOCAL_MACHINE. (HKLM\System\CurrentControlSet\Hardware Profiles). It indicates the *Profile-number* subkey that stores the values that make up the current profile. The content of the indicated subkey also appears in the Hardware Profiles\Current subkey in Hardware Profiles.

This value also corresponds to the number of one of the *Profile-number* subkeys in the Hardware Profiles subkey of IDConfigDB (HKLM\System\CurrentControlSet\Control\IDConfigDB\Hardware Profiles). This subkey stores information about the hardware profiles configured for use on the computer, including FriendlyName, which stores the name of the profile used in System in Control Panel.

# Hardware Profiles Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware Profiles**
HKLM\SYSTEM\CurrentControlSet\Control\IDConfigDB

**Description**
The Hardware Profiles subkey stores information about the hardware profiles configured on the system.

A *hardware profile* consists of specified changes to the original hardware profile configured in the HKLM\SOFTWARE and SYSTEM keys of HKEY_LOCAL_MACHINE. These changes are stored in the CurrentControlSet\Hardware Profiles subkey in HKEY_LOCAL_MACHINE.

The Hardware Profiles subkey in the IDConfigDB subkey stores information about the profile, but not the profile itself.

You can use System in Control Panel to create a hardware profile. Also, the system creates new hardware profiles automatically when it encounters configurations that require new profiles, such as when a portable computer is undocked or is moved from one docking station to another.

**Tip**

- To create a hardware profile, double-click System in Control Panel, click the **Hardware** tab, and then click the **Hardware Profiles** button.

# Profile-number

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Profile-number**
HKLM\SYSTEM\CurrentControlSet\Control\IDConfigDB\Hardware Profiles

**Description**
A *Profile-number* subkey stores information about a particular hardware profile on the system.

The data stored in this subkey is specific to the hardware profile.

### Notes

- The *Profile-number* subkey is a variable representing the numbered subkeys in the IDConfigDB\Hardware Profiles subkey. It does not actually appear in the registry. This subkey describes features common to all numbered Hardware Profiles subkeys in the registry.

- This subkey stores information about the hardware profile, but it does not store the hardware profile. The changes that comprise the actual hardware profile are stored in a *Profile-number* subkey in HKLM\System\CurrentControlSet\Hardware Profiles.

### Tip

- To determine which numbered subkey in Hardware Profiles represents the current hardware profile, see the CurrentConfig entry. Its value corresponds to the number of the subkey that contains the current hardware profile.

**See Also**
**Concepts**
CurrentConfig

# Profile-number\FriendlyName Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FriendlyName**
HKLM\SYSTEM\CurrentControlSet\Control\IDConfigDB\Hardware Profiles*Profile-number*

**Description**
Stores the commonly used name for the hardware profile. This name is used to identify the hardware profile in System in Control Panel and in Device Manager.

**Tip**

- To determine which numbered subkey in Hardware Profiles represents the current hardware profile, see the CurrentConfig entry in HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\IDConfigDB. Its value corresponds to the number of the subkey that contains the current hardware profile.

# Lsa

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Lsa**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The Lsa subkey stores configuration data for the Local Security Authority (LSA).

**Tip**

- For more information about Windows Server 2003 security, see the *Windows Server 2003 Technical Reference*.

# Bounds

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Bounds**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | 0x0–0x00 00 0A 00 00 00 0A 00 | 0x00 30 00 00 00 20 00 00 (Upper bound = 0x3000 (12,288); lower bound = 0x2000 (8,192)) |

**Description**
Specifies thresholds for managing the length of the kernel-mode Local Security Authority (LSA) audit queue. The audit queue stores kernel-mode events destined for the Security Log in Event Viewer.

The value of this entry is an 8-byte binary field. The value of the first four bytes specifies the maximum number of items that can be held in the audit queue (the upper bound). When the number of audits exceeds this value, LSA discards all new audits until the number of audits remaining in the queue reaches the lower bound, as specified by the value of the last four bytes.

**Tip**

- The system does not notify you when the queue is nearing, has reached, or has exceeded its upper bound. To prevent the system from running when it cannot report all security events, set the value of CrashOnAuditFail to 1.

**See Also**
**Concepts**
CrashOnAuditFail

# CrashOnAuditFail

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CrashOnAuditFail**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 0 |

**Description**
Directs the system to halt when it cannot record new events in the Security Log in Event Viewer. This feature prevents unauthorized activities from occurring when they cannot be recorded in the Security Log.

The system also uses this entry to indicate that this feature has been triggered (a value of 2). When the value of this entry is 2, only members of the Administrators group can log on to the computer. This restricted state lets an Administrator log on to resolve the problem and to reset the value of this entry to 1.

| Value | Meaning |
|---|---|
| 0 | The feature is off. The system does not halt, even when it cannot record events in the Security Log. |
| 1 | The feature is on. The system halts when it cannot record an event in the Security Log. |
| 2 | The feature is on and has been triggered. The system halted because it could not record an auditable event in the Security Log. Only members of the Administrators group can log on. |

Typically, the system cannot record security events because the Security Log in Event Viewer is full or because the internal queue to the log has reached the maximum established by the value of Bounds.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
You must restart the computer before changes to this entry take effect.

**See Also**
**Concepts**
Bounds

# everyoneincludesanonymous

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**everyoneincludesanonymous**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the local *Everyone group* of users includes anonymous users. Specifically, this entry determines whether the Local Security Authority (LSA) includes the security identifier (SID) of the Everyone group in the security token of an anonymous user.

The *Everyone group* is a generic group of users that includes all users.

| Value | Meaning |
|---|---|
| 0 | The local Everyone group does not include anonymous users. |
| 1 | The local Everyone group includes anonymous users. |

**Change Method**
To change the value of this entry, use the **Let Everyone permissions apply to anonymous users** local security setting. In **Control Panel**, click **Performance and Maintenance**, then click **Administrative Tools**, and then click **Local Security Policy**. Double-click **Local Policies**, and then click **Security Options**. Right click **Network Access: Let Everyone permissions apply to anonymous users**, and then click **Properties**. If the local security setting is **Enabled**, then the value of this entry is 0. If the local security setting is **Disabled**, then the value of this entry is 1.

### Notes

For more information about the *LSA*, see the *Security Subsystem Architecture* section of the *Active Directory Data Storage* chapter in the *Distributed Systems Guide*.

### Tip

You can mimic the LSA behavior of Windows 2000 by setting the value of this entry to 1.

# NoLMHash

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NoLMHash**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether Security Accounts Manager (SAM) stores the LAN Manager (LM) hash of the user's password. The LM hash of the user's password is necessary to authenticate downlevel clients that cannot use NTLM or NTLMv2 authentication.

| Value | Meaning |
|---|---|
| 0 | SAM stores the message digest of the user's password. |
| 1 | SAM does not store the message digest of the user's password. |

# Notification Packages

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Notification Packages**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *File name*[ *file name*...] | FPNWCLNT RASSFM KDCSVC scecli |

**Description**
Specifies the dynamic-link libraries (DLLs) that are loaded or called when passwords are set or changed. To specify more than one file, list the file names one above the other by pressing ENTER between each file name.

This entry is often used to load libraries for services such as File and Print Services for NetWare.

You can use this entry to implement custom DLLs that specify alternate criteria for valid passwords, such as complex password rules. Notification packages that implement password verification are recommended for all primary and backup domain controllers. You can also write a custom password filter to fulfill special security requirements.

**Tip**

- To audit the loading of notification packages, from the **Start** menu, point to **Programs**, point to **Administrative Tools**, and then click **Local Security Policy**. Expand **Local Policies**, and then double-click **Audit Policy**. Double-click **Audit System Events**, and then use the **Success** and **Failure** check boxes.

# RestrictAnonymous

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RestrictAnonymous**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 0 |

**Description**
Restricts anonymous users from displaying lists of users and from viewing security permissions.

| Value | Meaning |
|---|---|
| 0 | Disabled. Anonymous users are not restricted. |
| 1 | Enabled. Users who log on anonymously (also known as *null session connections*) cannot display lists of domain user names, nor share names. Also, these users cannot view security permissions, and they cannot use all of the features of Windows Explorer, Local Users and Groups, and other programs that enumerate users or shares. |
| 2 | Anonymous users have no access without explicit anonymous permissions. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Do not set the value of this entry to 2 in mixed-mode environments. Only consider setting it to 2 in environments running only Windows Server 2003, and only after verifying that appropriate service levels and program function are maintained.

### Caution

- Pre-defined "High Secure" security templates set the value of this entry to 2. Use caution when using these templates.

# Kerberos

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Kerberos**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa

**Description**
The Kerberos subkey stores configuration data for the Kerberos V5 authentication protocol in Windows Server 2003. The Kerberos V5 protocol is the primary security protocol for authentication within a domain. It verifies the identity of the user and the system.

# KerberosDomains

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domains**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos

**Description**
The Domains subkey stores information about non-Windows Kerberos realms.

**Change Method**
To change the value of the entries in this subkey, use Kerberos Setup (Ksetup.exe), a tool included in Windows Support Tools. Do not use the registry editor Regedit.exe.

# Realm-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Realm-name**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Domains

**Description**
The *Realm-name* subkey stores configuration data for a non-Windows-based Kerberos realm. Non-Windows Kerberos servers use realms instead of Windows domains.

Kerberos Setup (Ksetup.exe), a tool included with the Windows Server 2003 Support Tools, adds this subkey to the registry when you use it to create a Kerberos realm. The Kerberos Security Provider in Windows Server 2003 uses the data in the subkey to locate a non-Windows Kerberos Key Distribution Center and to interoperate with the Kerberos realm.

This subkey bears the name you assign to the realm when you use the **/SetRealm** command in Ksetup.exe.

**Change Method**
To add, delete, or change the configuration of a Kerberos realm, use Ksetup.exe.

### Note

- The *Realm-name* subkey does not actually appear in the registry. It is a placeholder used to describe the elements common to realm-specific subkeys in the Domains subkey.

### Caution

If you use a domain's realm name in this configuration, the workstation treats the domain simply as a Kerberos realm.

# KdcNames

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**KdcNames**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Domains*Realm-name*

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Space-separated list of the DNS names of KDC servers* | (None. You must specify at least one KDC server when you create a realm.) |

**Description**
Stores the Domain Name System (DNS) names of computers that are running the Key Distribution Center (KDC) service for an Kerberos realm.

This entry stores the values that you enter when you use the **/AddKdc** command in Kerberos Setup (Ksetup.exe), a tool included in Windows Server 2003 Support Tools.

Computers that are running Windows Server 2003 can use a non-Windows Kerberos server to administer authentication, instead of using a Windows Server 2003 domain. These systems participate in an Kerberos realm instead of a Windows domain, and the DNS names for the KDC servers are stored in this entry in the registry, instead of on a domain controller.

**Change Method**
To change the value of this entry, use Ksetup.exe. Do not edit the registry.

### Note

- If this entry is absent from the *Realm-name* subkey, then the computer uses a DNS query (for SRV records) to locate the KDC for the realm.

# KpasswdNames

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**KpasswdNames**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Domains*Realm-name*

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *List of computer names* | (None. You must specify at least one password server when you create a realm.) |

**Description**
Lists the computer names of Kerberos Change Password Protocol Servers for the Kerberos realm. These servers communicate on port 464.

This entry stores the values that you enter when you use the **/AddKpasswd** command in Kerberos Setup (Ksetup.exe), a tool included in Windows Server 2003 Support Tools.

Computers running Windows Server 2003 can use a non-Windows Kerberos server to administer authentication, instead of using a Windows Server 2003 domain. These systems participate in a Kerberos realm instead of a Windows domain, and the names of the Key Distribution Center (KDC) password servers are stored in this entry in the registry, instead of on a domain controller.

**Change Method**
To change the value of this entry, use Ksetup.exe. Do not edit the registry.

### Caution

- If this entry is absent from the *Realm-name* subkey, then the computer attempts to locate the Kerberos Change Password Protocol Servers by using a DNS query for SRV records for the realm. If no servers can be located, then users are not be able to change their passwords in the realm.

# RealmFlags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RealmFlags**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Domains*Realm-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–7 | 0 |

**Description**
Specifies additional features of a non-Windows-based Kerberos realm.

Computers that are running Windows Server 2003 can use a non-Windows Kerberos server to administer authentication, instead of using a Windows Server 2003 domain. These systems participate in a Kerberos realm instead of a Windows domain. This entry establishes the features of the realm.

To enable a feature, set the bit representing the feature to 1, or sum the decimal or hexadecimal values representing each feature. For example, to enable all features, set the value of this entry to 7 (1 + 2 + 4).

| Value | Flag | Meaning |
|---|---|---|
| 0 | None | No additional features enabled. |
| 1 | SendAddress | Include IP numbers within tickets. Useful for solving some compatibility issues. |
| 2 | TcpSupported | Realm supports TCP rather than just UDP. |
| 4 | Delegate | Everyone in this realm is trusted for delegation. |
| 8 | NcSupported | Realm supports name canonicalization. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Kerberos\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Parameters**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos

**Description**
The Parameters subkey stores configuration options for the Kerberos V5 authentication protocol in Windows Server 2003.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# UserList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**UserList**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos

**Description**
The UserList subkey stores entries that associate a Kerberos security principal to a local Windows Server 2003 user account.

Computers that are running Windows Server 2003 can use a non-Windows Kerberos Key Distribution Center to administer authentication, instead of using a Windows Server 2003 domain. For ease of use, the system lets you map a Kerberos security principal, such as the name of a principal or a realm, to a local Windows user account.

This subkey stores mappings that you enter when you use the **/MapUser** command in Kerberos Setup (Ksetup.exe), a tool included in *Windows Server 2003 Support Tools*. Ksetup.exe adds the entries to the registry.

All entries in this subkey have the following format, where * indicates all users:

| Entry name | Data type | Value |
|---|---|---|
| *Kerberos-name* or * | REG_SZ | *Local-name* or * |

For example, if you enter the following mapping command in Ksetup.exe:

```
ksetup /mapuser user@domain.reskit.com  Guest
```

then Ksetup.exe adds the following entry to the UserList subkey:

| Entry name | Data type | Value |
|---|---|---|
| user@domain.reskit.com | REG_SZ | *Guest* |

**Change Method**
To change the value of the entries in this subkey, use Ksetup.exe. Do not edit the registry.

### Caution

- Mapping a Kerberos principal to a local account grants access to your computer to anyone who knows the password to the Kerberos principal. Use of a wildcard entry should be used cautiously, especially if it can reference a local account with administrative rights.

# MSV1_0

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MSV1_0**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa

**Description**
The MSV1_0 subkey stores information about the computer's security settings.

# NtlmMinClientSec

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NtlmMinClientSec**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa\MSV1_0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x10 \| 0x20 \| 0x80000 \| 0x20000000 | 0x0 |

**Description**
Specifies the minimum required security setting of client-side network connections for applications using the NTLM security support provider (SSP).

| Value | Meaning |
|---|---|
| 0x0 | None. No security is used for authentication or session security. |
| 0x10 | Message integrity. If the value of either this entry or the NtlmMinServerSec entry is 0x10, then the connection will fail unless message integrity is negotiated. |
| 0x20 | Message confidentiality. If the value of either this entry or the **NtlmMinServerSec** entry is 0x20, then the connection will fail unless message confidentiality is negotiated. |
| 0x80000 | NTLMv2 session security. If the value of either this entry or the **NtlmMinServerSec** entry is 0x80000, then the connection will fail unless NTLMv2 session security is negotiated. |
| 0x20000000 | 128-bit encryption. If the value of either this entry or the **NtlmMinServerSec** entry is 0x20000000, then the connection will fail unless 128-bit encryption is negotiated. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must restart Windows to make changes to this entry effective.

### Note

- These settings do not guarantee that an application uses message integrity or confidentiality, even if they are negotiated. For more information, see *Article Q147706* in the Microsoft Knowledge Base, or search by using the keywords *LM authentication*.

**See Also**
**Concepts**
NtlmMinServerSec

# NtlmMinServerSec

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NtlmMinServerSec**
HKLM\SYSTEM\CurrentControlSet\Control\Lsa\MSV1_0

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x10 \| 0x20 \| 0x80000 \| 0x20000000 | 0x0 |

**Description**
Specifies the minimum required security setting of server-side network connections for applications using the NTLM security support provider (SSP).

| Value | Meaning |
|---|---|
| 0x0 | None. No security is used for authentication or session security. |
| 0x10 | Message integrity. If the value of either this entry or the NtlmMinClientSec entry is 0x10, then the connection will fail unless message integrity is negotiated. |
| 0x20 | Message confidentiality. If the value of either this entry or the **NtlmMinClientSec** entry is 0x20, then the connection will fail unless message confidentiality is negotiated. |
| 0x80000 | NTLMv2 session security. If the value of either this entry or the **NtlmMinClientSec** entry is 0x80000, then the connection will fail unless NTLMv2 session security is negotiated. |
| 0x20000000 | 128-bit encryption. If the value of either this entry or the **NtlmMinClientSec** entry is 0x20000000, then the connection will fail unless 128-bit encryption is negotiated. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must restart Windows to make changes to this entry effective.

### Note

- These settings do not guarantee that an application uses message integrity or confidentiality, even if they are negotiated. For more information, see *Article Q147706* in the Microsoft Knowledge Base, or search by using the keywords *LM authentication*.

**See Also**
**Concepts**
NtlmMinClientSec

# Control\Network Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The Network subkey stores configuration data for networking components and connections.

# DUPTILE1_NetworkProvider

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetworkProvider**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The NetworkProvider subkey contains information about the available network providers.

**See Also**
**Concepts**
service-name\networkprovider

# NetworkProvider\Order

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Order**
HKLM\SYSTEM\CurrentControlSet\Control\NetworkProvider

**Description**
The Order subkey contains a list of the network providers that use the Microsoft network-independent APIs.

**See Also**
**Concepts**
service-name\networkprovider

# ProviderOrder

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ProviderOrder**
HKLM\SYSTEM\CurrentControlSet\Control\NetworkProvider\Order

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Provider name*[, *Provider name*, ...] | LanmanWorkstation (when a single network is installed) |

**Description**
Contains a list of the available network providers in the order in which they are accessed.

To appear in this list, a network provider must be represented by a subkey in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services. It must be a member of the NetworkProvider service group, and it must have a service-name\networkprovider subkey in its Services subkey. The NetworkProvider subkey stores information for the Multiple Provider Router. The service group for each service is stored in the value of the Group entry.

**Change Method**
To change the value of this entry, double-click Network in Control Panel. From the **Advanced** menu, choose **Advanced Settings** and click the **Provider Order** tab. Click on a provider and use the up and down arrows to change the order.

**Tip**

- To find all the services configured to run in the NetworkProvider service group, SC, a tool included in Windows Server 2003, and type:

```
sc queryex group= networkprovider state= all
```

**See Also**
**Concepts**
Group
service-name\networkprovider

# Nls

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Nls**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The Nls subkey stores data for the National Language Support (NLS) interface. This subkey contains information about supported locales and code pages.

The CodePage and Language subkeys of the Nls subkey associate code page and locale identification codes with the translation files required to use them.

**See Also**
**Concepts**
CodePage
Language

# CodePage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CodePage**
HKLM\SYSTEM\CurrentControlSet\Control\NIs

**Description**
Contains data that associates code page IDs with code page translation files installed on the system.

Each entry in the CodePage subkey represents one supported code page. If the code page is installed on the system, then the value of the entry is the name of the code page file. If the code page is supported but not installed, then the value of the entry is blank.

# code-page

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**code-page**
HKLM\SYSTEM\CurrentControlSet\Control\Nls\CodePage

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *NLS filename* | (*Depends on locale and language group installed.*) |

**Description**
Represents a code page.

Each entry in the CodePage subkey represents one supported code page. If the code page is installed on the system, then the value of the entry is the name of the code page file. If the code page is supported but not installed, then the value of the entry is blank.

Entries in the CodePage subkey have the following format:

| Entry name | Data type | Value |
|---|---|---|
| # *CodePageID* | REG_SZ | *Filename* \| (blank) |

For example, the following entry indicates that code page 10000 is installed on the system and is associated with code page file C_10000.nls:

| Entry name | Data type | Value |
|---|---|---|
| 10000 | REG_SZ | c_10000.nls |

As another example, the following entry indicates that code page 10006 is supported but is not installed on the system:

| Entry name | Data type | Value |
|---|---|---|
| 10006 | REG_SZ | |

**Note**

- The *code-page* entry is a variable representing the entries in the CodePage subkey. It does not actually appear in the registry. This variable entry displays information that is common to the entries in the CodePage subkey.

# ACP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ACP**
HKLM\SYSTEM\CurrentControlSet\Control\NIs\CodePage

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *code page* | 1252 (*U.S. English*) |

**Description**
Indicates the default code page for non-Unicode Windows characters. These code pages are supported for compatibility with non-Unicode programs. For U.S. English, the value of this entry indicates that code page 1252 and its file, C_1252.nls, are used for non-Unicode Windows characters.

| Value | Character set |
|---|---|
| 874 | Thai |
| 932 | ShiftJIS—Japan |
| 936 | GBK—Chinese (China, Singapore) |
| 949 | Korean Unified Hangul |
| 950 | Big5 Extended—Chinese (Taiwan, Hong Kong SAR) |
| 1250 | Central European |
| 1251 | Cyrillic |
| 1252 | U.S. (ANSI) |
| 1253 | Greek |
| 1254 | Turkish |
| 1255 | Hebrew |
| 1256 | Arabic |
| 1257 | Baltic |
| 1258 | Vietnamese |

**Caution**

- Changing the value of this entry could prevent the computer from properly displaying characters in non-Unicode programs, and could also result in data loss.

# MACCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MACCP**
HKLM\SYSTEM\CurrentControlSet\Control\NIs\CodePage

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *code page* | 10000 (*U.S. English*) |

**Description**
Indicates the default code page for characters in non-Unicode Macintosh programs. For U.S. English, the value of this entry indicates that code page 10000 and its file, C_10000.nls, are used for non-Unicode Macintosh programs.

| Value | Character set |
|---|---|
| 10000 | Roman |
| 10001 | Japanese |
| 10002 | Chinese (Big5) |
| 10003 | Korean |
| 10004 | Arabic |
| 10005 | Hebrew |
| 10006 | Greek I |
| 10007 | Cyrillic |
| 10010 | Romanian |
| 10017 | Ukrainian |
| 10029 | Latin 2 |
| 10079 | Icelandic |
| 10081 | Turkish |
| 10082 | Croatian |

**Caution**

- Changing the value of this entry could prevent the computer from properly displaying characters in non-Unicode programs, and could also result in data loss.

# OEMCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**OEMCP**
HKLM\SYSTEM\CurrentControlSet\Control\NIs\CodePage

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *code page* | 437 (*U.S. English*) |

**Description**
Indicates the default code page for OEM (original equipment manufacturers) characters, used primarily in DOS programs. For U.S. English, the value of this entry indicates that code page 437 and its file, C_437.nls, are used for OEM characters.

| Value | Character set |
|---|---|
| 437 | United States |
| 708 | Arabic (ASMO 708) |
| 720 | Arabic (Transparent ASMO) |
| 737 | Greek (formerly 437G) |
| 775 | Baltic |
| 850 | Multilingual (Latin I) |
| 852 | Slavic (Latin II) |
| 855 | Cyrillic |
| 857 | Turkish |
| 860 | Portuguese |
| 861 | Icelandic |
| 862 | Hebrew |
| 863 | Canadian-French |
| 864 | Arabic |
| 865 | Nordic |
| 866 | Russian |
| 869 | Modern Greek |
| 874 | Thai |
| 932 | ShiftJIS—Japan |
| 936 | GBK—Chinese (China, Singapore) |
| 949 | Korean Unified Hangul |
| 950 | Big5 Extended—Chinese (Taiwan, Hong Kong SAR) |

| | |
|---|---|
| 1258 | Windows Vietnamese |

**Caution**

- Changing the value of this entry could prevent the computer from properly displaying characters in non-Unicode programs, and could also result in data loss.

# Language

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Language**
HKLM\SYSTEM\CurrentControlSet\Control\NIs

**Description**
Contains data that associates locale IDs with internal supporting files installed on the system.

Each entry in the Language subkey represents one locale. If the locale is installed on the system, then the value of the entry is the name of its internal supporting file. If the locale is supported but not installed, then there is no value for the entry.

# language-ID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**language-ID**
HKLM\SYSTEM\CurrentControlSet\Control\Nls\Language

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *.nls filename* | (*Depends on locale and language group installed.*) |

**Description**
Stores the filename of the *casing table* for a particular locale.

A *casing table* specifies the use of upper and lower case letters for a particular language. For example, the character "a" uppercases to the character "A".

Each entry in the Language subkey represents one locale supported by Windows Server 2003. If the language group for that particular locale is installed on the system, then the casing table is also installed and the value of this entry is the name of the file that contains the casing table. The filename extension is .nls. If the language group is supported but not installed, then there is no value for the entry.

Each entry has the following format:

| Entry name | Data type | Value |
|---|---|---|
| # *LanguageID* | REG_SZ | *Filename.nls* \| (no value) |

For example, the following entry indicates that the language group for the locale 0402 is installed on the system and, therefore, the casing file L_intl.nls is also installed for locale 0402:

| Entry name | Data type | Value |
|---|---|---|
| 0402 | REG_SZ | l_intl.nls |

As another example, the following entry indicates that the language group for the locale 0403 is supported by Windows Server 2003, but is not installed on the computer, and no casing file is therefore supported for this locale:

| Entry name | Data type | Value |
|---|---|---|
| 0403 | REG_SZ | |

**Note**

- The *language-ID* entry is a variable representing the entries in the Language subkey. It does not actually appear in the registry. This variable entry displays information that is common to the entries in the Language subkey.

# Language Groups

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Language Groups**
HKLM\SYSTEM\CurrentControlSet\Control\Nls

**Description**
The Language Groups subkey contains information about the supported and installed language groups.

Each of the entries in this subkey represents a language group that is supported by Windows Server 2003.

| ***group-number*** | **Language Group** |
|---|---|
| 1 | Western Europe |
| 2 | Central Europe |
| 3 | Baltic |
| 4 | Greek |
| 5 | Cyrillic |
| 6 | Turkish |
| 7 | Japanese |
| 8 | Korean |
| 9 | Traditional Chinese |
| a | Simplified Chinese |
| b | Thai |
| c | Hebrew |
| d | Arabic |
| e | Vietnamese |
| f | Indic |
| 10 | Georgian |
| 11 | Armenian |

# group-number

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**group-number**
HKLM\SYSTEM\CurrentControlSet\Control\Nls\Language Groups

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | (Blank) \| 1 | (None) |

**Description**
Indicates which language groups are installed.

| Value | Meaning |
|---|---|
| (Blank) | The language group is supported but not installed. |
| 1 | The language group is installed. |

### Note

- This entry is a placeholder representing the type of entries in the Language Groups subkey; "group-number" does not actually appear in the registry.

# Nls\Locale

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Locale**
HKLM\SYSTEM\CurrentControlSet\Control\Nls

**Description**
The entries in this subkey list the locales that are installed on the system.

# Alternate Sorts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Alternate Sorts**
HKLM\SYSTEM\CurrentControlSet\Control\Nls\Locale

**Description**
The Alternate Sorts subkey contains data about the alternate sorts available for the installed locale. These sorting codes determine which sorting order is used for those locales that have more than one sorting order.

**Change Method**
For some locales, Windows provides multiple sorting methods. When a computer is configured for a locale with multiple sorting methods, you can change the sorting method from **Regional and Language Options**.

In **Control Panel** double-click **Regional and Language Options**, click the **Regional Options** tab, and then click the **Customize...** button. If the currently configured locale has multiple sorting methods, then Windows displays a **Sorting** tab. Click the **Sorting** tab, and then select a method from the **Select the sorting method to use with this language** list. Click **OK**.

# Control\Print

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Print**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The Print subkey stores printer information and settings that apply to all users of the computer.

### Note

- For more information about printing with Windows Server 2003, see *Administering Microsoft Windows XP Professional*.

# BrowsePrintWorkstations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BrowsePrintWorkstations**
HKLM\SYSTEM\CurrentControlSet\Control\Print

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0xFFFFFFFF | 2 |

**Description**
Used to calculate the number of clients in whose browse list the printer appears.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make the value of this entry effective, you must restart the Print Spooler service.

# EnableBroadcastSpoolerStatus

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**EnableBroadcastSpoolerStatus**
HKLM\SYSTEM\CurrentControlSet\Control\Print

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the system sends spooler and job status messages when the job count changes.

| Value | Meaning |
|---|---|
| 0 | Status messages are not sent. |
| 1 | Status messages are sent. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- A value of 1 can degrade the performance of your system by requiring programs that don't use spooler information to read and discard it.

# MaxRpcSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MaxRpcSize**
HKLM\SYSTEM\CurrentControlSet\Control\Print

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (in bytes) | 1048576 (1 MB) |

**Description**
Specifies the maximum size, in bytes, of incoming data blocks from a remote procedure call (RPC).

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make changes to the value of this entry effective, you must restart the Print Spooler service.

### Caution

- This entry can help prevent malicious denial-of-service attacks. If the data block of a remote procedure call (RPC) is larger than the value of this entry, the RPC run-time library rejects the call and sends an error message to the client. However, if you set the value of this entry to 0xFFFFFFFF, the limit on the size of incoming data blocks is removed, and your system is more vulnerable to attack.

# NetPrinterDecayPeriod

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetPrinterDecayPeriod**
HKLM\SYSTEM\CurrentControlSet\Control\Print

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF milliseconds | 0x36EE80 (3,600,000 milliseconds = 1 hour) |

**Description**
Specifies how long a network printer remains in the browse list.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# ServerThreadTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ServerThreadTimeout**
HKLM\SYSTEM\CurrentControlSet\Control\Print

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *milliseconds* | 0x927C0 (600,000 milliseconds = 10 minutes) |

**Description**
Specifies the amount of time that the server thread is idle before it notifies other print servers of its printers.

If the value of this entry is greater than the value of NetPrinterDecayPeriod, then the printer does not appear in the printer list.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make the value of this entry effective, you must restart the Print Spooler service.

**See Also**
**Concepts**
NetPrinterDecayPeriod

# SpoolerPriority

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SpoolerPriority**
HKLM\SYSTEM\CurrentControlSet\Control\Print

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x20 \| 0x40 \| 0x80 | 0x20 |

**Description**
Specifies the priority class of the print spooler process.

The priority of the print spooler process establishes a range of priority for print spooler threads.

| Value | Meaning |
|---|---|
| 0x20 | Normal priority |
| 0x40 | Idle |
| 0x80 | High priority |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Windows environment-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows environment-name**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Environments

**Description**
The Windows *environment-name* subkey represents a platform environment, and stores entries that apply only to that environment.

### Notes

- The Windows *environment-name* subkey is a placeholder representing the types of subkeys in the Environments subkey; "environment-name" does not actually appear in the registry.

- The Windows *environment-name* subkeys contain subkeys that store configuration data for printing on different platforms on which different versions of Windows run. Each subkey representing a platform has two subkeys, Drivers and Print Processors, which store configuration data for the print drivers and processors installed on Windows.

# Drivers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Drivers**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Environments\Windows *environment-name*

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Version-number

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Version-number**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Environments\Windows *environment-name*\Drivers

**Description**
The Version-*number* subkey contains information about the drivers for a print environment. There is a different number for each version of the printer driver. Version-2 drivers are for Windows NT 4.0; Version-3 drivers are for Windows 2000 and Windows Server 2003.

### Note

- The Version-*number* subkey is a placeholder representing the subkeys in the Drivers subkey, which have a specific number in their name. "Version-*number*" does not actually appear in the registry.

# driver-name Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**driver-name**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Environments\Windows *environment-name*\Drivers\Version-*number*

**Description**
The *driver-name* subkey in the Version-*number* subkey contains information about necessary files and settings for the specified printer driver.

### Note

- The *driver-name* subkey is a placeholder representing the types of entries in the Version-*number* subkey; "driver-name" does not actually appear in the registry.

# Version

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Version**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Environments\Windows *environment-name*\Drivers\Version-*number**driver-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Version number* | (Depends on printer driver.) |

**Description**
Contains the version number of the installed print driver.

### Note

- The print spooler adds this entry to the registry when a printer driver is installed into the print subsystem.

# Monitors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitors**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Description**
The Monitors subkey stores entries for print monitors installed on Windows Server 2003. The Monitors subkey contains a subkey for each port monitor that stores configuration data for the port monitor installed on the system.

# Monitor

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Monitors

**Description**
The *Monitor* subkey contains information about the status and drivers of the monitors used by the computer.

### Note

- The *Monitor* subkey is a placeholder representing the types of subkeys in the Monitors subkey; "Monitor" does not actually appear in the registry.

# Ports Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ports**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Monitors*Monitor*

**Description**
The Ports subkey contains information about the current availability of the ports.

# StatusUpdateEnabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**StatusUpdateEnabled**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Monitors*Monitor*\Ports

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the port monitor is able to update the spooler with the port status.

| Value | Meaning |
|---|---|
| 0 | Disabled. The monitor is not able to update the spooler with the port status. |
| 1 | Enabled. The monitor is able to update the spooler with the port status. |

This entry does not appear in every *Monitors* subkey. It appears only in the Standard TCP/IP Port subkey.

**Tip**

- For pooled printers, set the value of this entry to 1, and StatusUpdateInterval to a minumum value, to allow the spooler to send print jobs to an alternate printer when one printer in the pool is offline.

**See Also**
**Concepts**
Providers\RestartJobOnPoolEnabled
RestartJobOnPoolError Entry
StatusUpdateInterval

# StatusUpdateInterval

Updated: September 2, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**StatusUpdateInterval**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Monitors*Monitor*\Ports

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *number of minutes* | 0xA (10 minutes) |

**Description**
Specifies how often the port monitor updates the spooler with port status.

This entry does not appear in every *Monitors* subkey. It appears only in the Standard TCP/IP Port subkey.

**Tip**

- For pooled printers, setting the value of StatusUpdateEnabled to 1 and **StatusUpdateInterval** to a minumum value will allow the spooler to send print jobs to an alternate printer when one printer in the pool is offline.

**See Also**
**Concepts**
Providers\RestartJobOnPoolEnabled
RestartJobOnPoolError Entry
StatusUpdateEnabled

# Print\Providers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providers**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Description**
The Providers subkey contains information about the network printers. The printers listed in this subkey are not hosted locally, but they are printers to which the local computer is connected.

### Note

- Providers entries for LanMan Print Services are stored in the Providers Subkey subkey in the SOFTWARE key.

**See Also**
**Concepts**
Providers Subkey
server-name\Providers

# EventLog Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**EventLog**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x1F (bitmask) | Windows XP Professional: 0x1B (11011) Windows Server 2003: 0x1F (11111) |

**Description**
Specifies the level of detail of events that the print spooler logs for local print jobs.

The value of this entry is a bitmask. To direct the print spooler to log entries of the type specified, either set the bit representing that type to 1 or sum the hexadecimal values of all desired types. For example, the default value for Windows Server 2003 (11111) indicates that the print spooler logs all events. The default value for Windows XP Professional (11011) indicates that the print spooler logs all events except information.

| Value | Event type |
|---|---|
| 0x0 (0 binary) | Success |
| 0x1 (1 binary) | Error |
| 0x2 (10 binary) | Warning |
| 0x4 (100 binary) | Information |
| 0x8 (1000 binary) | Audit success |
| 0x10 (10000 binary) | Audit failure |

**Change Method**
To change the value of this entry, from the **Start** menu, point to **Settings**, and click **Printers**. From the **File** menu, click **Server Properties**, click the **Advanced** tab, and then select or clear the check boxes representing each event type.

**See Also**
**Concepts**
Providers\EventLog

# Providers\RestartJobOnPoolEnabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RestartJobOnPoolEnabled**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether a print job to a local printer can be restarted after an error.

If the value of this entry is set to 1, then a **Cancel/Retry** dialog box will appear on the user's system for the duration of the value of the RestartJobOnPoolError Entry entry.

| Value | Meaning |
|---|---|
| 0 | Print jobs cannot be restarted on error. |
| 1 | Print jobs can be restarted on error. |

### Notes

- This entry applies only to pooled printers.
- The print job cannot be restarted if it is completely sent to the printer and deleted from the print queue.

**See Also**
**Concepts**
RestartJobOnPoolEnabled Entry
RestartJobOnPoolError Entry

# RestartJobOnPoolError Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RestartJobOnPoolError**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *number of seconds* | 0x258 (600 seconds = 10 minutes) |

**Description**
Specifies how long the print monitor waits after an error occurs before starting a print job on another port in the print pool.

### Notes

- This entry applies only to pooled printers.
- Both the print monitor and the printer must support SNMP.
- The print job cannot be restarted if it has been completely sent to the printer and deleted from the print queue.

**See Also**
**Concepts**
Providers\RestartJobOnPoolError
StatusUpdateEnabled
StatusUpdateInterval

# NetWare or Compatible Network

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetWare or Compatible Network**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers

**Description**
Stores options for printing when using print resources on NetWare servers.

This subkey is used when you connect a Windows XP Professional system to a server running NetWare server software.

### Note

- This subkey is added to the registry when you install Client Service for NetWare.

**See Also**
**Concepts**
PrintOption

# DisplayName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DisplayName**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers\NetWare or Compatible Network

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *provider name* | NetWare or Compatible Network |

**Description**
Specifies the localized name of the print service provider displayed in the Print Manager and Printers main window.

### Note

- This entry only appears by default in the registry of a computer that has Client Services for NetWare installed.

# Name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Name**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers\NetWare or Compatible Network

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *file name* | nwprovau.dll |

**Description**
Specifies the dynamic-link library (DLL) for the NetWare or Compatible Print Service provider.

This DLL is stored in *Systemroot*\WINNT\system32.

### Note

- This entry only appears by default in the registry of a computer that has Client Services for NetWare installed.

# NetWare or Compatible Network\Timeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Timeout**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers\NetWare or Compatible Network

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x2710–0xFFFFFFFE *milliseconds* \| 0xFFFFFFFF | (None) |

**Description**
Specifies how often Windows Server 2003 updates the print queue of a NetWare server. The value of this entry specifies the time interval between automatic refresh updates.

| Value | Meaning |
|---|---|
| 0x2710–0xFFFFFFFE *milliseconds* | Specifies the time between updates of the NetWare print queue. |
| 0xFFFFFFFF | Disables automatic updates. Windows Server 2003 sends new print jobs to the NetWare server, but it does not update the queue list. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# ProductOptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ProductOptions**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
Stores information about the variations of Windows Server 2003 running on the computer.

**Caution**

- Do not delete this subkey or change its contents. This subkey stores protected operating system data. Editing this data violates your Windows Server 2003 licensing agreement and disqualifies you from receiving Microsoft Product Support Services. In addition, if a program tries to change the data, the system stops abnormally and displays a Stop error screen.

# ProductSuite

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ProductSuite**
HKLM\SYSTEM\CurrentControlSet\Control\ProductOptions

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Platform suite* or blank | (None) |

**Description**
Stores data that identifies a special operating system variation running on the computer. This entry lets programs identify and customize their features to the platform.

**Tip**

- To view this data, programs can use the GetVersionEx and VerifyVersionInfo application programming interfaces (APIs).

**Caution**

- Do not delete this entry or change its value. This entry stores protected operating system data. Editing this data violates your Windows Server 2003 licensing agreement and disqualifies you from receiving Microsoft Product Support Services service. If programs try to change the data, the system stops abnormally and displays a Stop error screen.

# ProductType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ProductType**
HKLM\SYSTEM\CurrentControlSet\Control\ProductOptions

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Platform identifier* | Windows XP Professional: WinNT Windows Server 2003: ServerNT or LanmanNT |

**Description**
Stores data that identifies the operating system platform running on the computer. This data lets programs identify and customize their features to the platform.

In Windows Server 2003, the following entries identify the platform options.

| Value | Meaning |
|---|---|
| WinNT | Windows XP Professional |
| ServerNT | Windows Server 2003 |
| LanmanNT | Windows Server 2003 |

**Tip**

- To view this data, programs can use the GetVersionEx and VerifyVersionInfo application programming interfaces.

**Caution**

- Do not delete this entry or change its value. This entry stores protected operating system data. Editing this data violates your Windows Server 2003 licensing agreement and disqualifies you from receiving Microsoft Product Support Services service. If a program tries to change the data, the system stops abnormally and displays a blue screen.

# SafeBoot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SafeBoot**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The entries in the SafeBoot subkey contain information about the safe mode settings of the computer. In safe mode, Windows Server 2003 uses default settings: VGA monitor, no network, Microsoft mouse driver, and the minimum device drivers required to start Windows.

For example, if your computer does not start after you install new software, you might be able to start it in safe mode with minimal services and then either change your computer settings or remove the newly installed software.

**Change Method**
To start Windows Server 2003 in safe mode, restart the computer and press F8 when you see the message **Please select the operating system to start**.

### Caution

- Changing the values of the entries in this subkey can prevent Windows Server 2003 from starting.

# AlternateShell

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AlternateShell**
HKLM\SYSTEM\CurrentControlSet\Control\SafeBoot

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Program name* | Cmd.exe |

**Description**
Lists the name of the alternative environment used when the **Safe mode with Command Prompt** option is selected.

**Change Method**
To start Windows Server 2003 in safe mode, restart the computer, and then press F8 when you see the message **Please select the operating system to start**.

### Note

- For information about starting the computer in safe mode, see Windows Server 2003 Help or Windows XP Professional Help.

### Caution

- Changing the value of this entry can prevent Windows Server 2003 from starting.

# Minimal

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Minimal**
HKLM\SYSTEM\CurrentControlSet\Control\SafeBoot

**Description**
The Minimal subkey stores the names of basic files and drivers (mouse, keyboard, mass storage, base video, and default system services) used when the **Safe Mode** startup option is selected.

**Change Method**
To start Windows Server 2003 in safe mode, restart the computer, and then press F8 when you see the message **Please select the operating system to start**.

### Note

- For information about starting the computer in safe mode, see Windows Server 2003 Help or Windows XP Professional Help.

### Caution

- Changing the value of this entry can prevent Windows Server 2003 from starting.

# file-name or driver-name Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**file-name or driver-name**
HKLM\SYSTEM\CurrentControlSet\Control\SafeBoot\Minimal

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *File name or driver name* | *File name or driver name* |

**Description**
Store the names of basic files and drivers (mouse, keyboard, mass storage, base video, and default system services) used when the **Safe Mode** startup option is selected.

**Change Method**
To start Windows Server 2003 in safe mode, restart the computer, and then press F8 when you see the message **Please select the operating system to start**.

### Note

- For information about starting the computer in safe mode, see Windows Server 2003 Help or Windows XP Professional Help.
- The ***file_name or driver_name*** subkey is a variable representing the subkeys in the HKLM\SYSTEM\CurrentControlSet\Control\SafeBoot\Minimal subkey. It does not actually appear in the registry. This variable entry displays information that is common to the entries in the Minimal subkey.

### Caution

- Changing the value of this entry can prevent Windows Server 2003 from starting.

# Safeboot\Network Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network**
HKLM\SYSTEM\CurrentControlSet\Control\SafeBoot

**Description**
The Network subkey stores the names of basic files, drivers, and network connections (mouse, keyboard, mass storage, base video, and default system services) used when the **Safe Mode with Networking** startup option is selected.

**Change Method**
To start Windows Server 2003 in safe mode, restart the computer, and then press F8 when you see the message **Please select the operating system to start**.

### Note

- For information about starting the computer in safe mode, see Windows Server 2003 Help or Windows XP Professional Help.

### Caution

- Changing the value of this entry can prevent Windows Server 2003 from starting.

# Network\file-name or driver-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**file-name or driver-name**
HKLM\SYSTEM\CurrentControlSet\Control\SafeBoot\Network

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *File name or driver name* | *File name or driver name* |

**Description**
Store the names of basic files and drivers (mouse, keyboard, mass storage, base video, and default system services) used when the **Safe Mode with Networking** startup option is selected.

**Change Method**
To start Windows Server 2003 in safe mode, restart the computer, and then press F8 when you see the message **Please select the operating system to start**.

### Note

- For information about starting the computer in safe mode, see Windows Server 2003 Help or Windows XP Professional Help.
- The *file_name or driver_name* subkey is a variable representing the subkeys in the HKLM\SYSTEM\CurrentControlSet\Control\SafeBoot\Network subkey. It does not actually appear in the registry. This variable entry displays information that is common to the entries in the Network subkey.

### Caution

- Changing the value of this entry can prevent Windows Server 2003 from starting.

# SecurePipeServers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SecurePipeServers**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The SecurePipeServers subkey contains the winreg subkey, which controls remote access to the registry. If these subkeys do not appear in the registry, then the permissions set in each registry subkey determine the remote access to the registry.

**Change Method**
To change the permissions set for the SecurePipeServers subkey, use the registry editor Regedit.exe. On the **Edit** menu, click **Permissions**.

### Note

- Remote registry security and the functions of the SecurePipeServers and winreg registry subkeys are included in Windows Server 2003 and Microsoft Windows XP Professional. However, they are only enabled by default on Windows Server 2003. If you add these subkeys to the registry of a computer running Microsoft Windows XP Professional, it will operate as described.

# winreg

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**winreg**
HKLM\SYSTEM\CurrentControlSet\Control\SecurePipeServers

**Description**
The permissions set for the winreg subkey determine who can remotely connect to a registry. If this subkey does not exist, then all users can remotely connect to the registry.

To remotely connect to a registry, a user must have at least Read Access to the winreg subkey on the target computer.

After the user is connected, the type of access that the user has to various registry keys and subkeys is determined by the permissions that have been granted to that user. Administrators have Full Control permission to winreg, and Backup Operators have read-only access to winreg.

**Change Method**
To view the permissions set for the winreg subkey, use the registry editor Regedit.exe. On the **Edit** menu, click **Permissions**.

### Note

- By default, the winreg subkey exists on both clients and servers. However, it is enabled by default only on Windows Server 2003. If you add this subkey to the registry of a computer that is running Windows XP Professional, then it will operate as described here.

# AllowedPaths

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AllowedPaths**
HKLM\SYSTEM\CurrentControlSet\Control\SecurePipeServers\winreg

**Description**
Contains a list of the registry paths that all users can access remotely, even if they do not otherwise have permission to access the winreg subkey. This subkey permits users to perform some system functions that require remote access to specific registry keys, such as checking the status of a printer, without having access to the winreg subkey.

**Change Method**
To view the permissions set for this subkey, use the registry editor Regedit.exe. On the **Edit** menu, click **Permissions**.

### Note

- The paths specified in this subkey can be accessed anonymously.

# Machine

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Machine**
HKLM\SYSTEM\CurrentControlSet\Control\SecurePipeServers\winreg\AllowedPaths

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Registry paths* | SYSTEM\CurrentControlSet\Control\ProductOptions<br>SYSTEM\CurrentControlSet\Control\Print\Printers SYSTEM\CurrentControlSet\Services\Eventlog<br>SOFTWARE\Microsoft\Windows NT\CurrentVersion |

**Description**
Lists registry paths in the HKEY_LOCAL_MACHINE subtree that all remote users can access. Remote users can access the paths listed in this subkey, and all subkeys of those paths, even if they do not otherwise have permission to access to the registry remotely.

### Note

- The paths specified can be accessed anonymously.

# ServiceGroupOrder

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ServiceGroupOrder**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The ServiceGroupOrder subkey contains a list of all service groups in the order in which the groups are loaded.

**Tip**

- To determine the service group to which a service belongs, find the value of its Group entry.
- To find all services in a service group, use the registry editor Regedit.exe. On the **Edit** menu, click **Find**, type the name of the service group, and then select only the **Data** check box.
- To specify the name of the group for a particular service, use SC, a tool included in Windows Server 2003.

**See Also**
**Concepts**
Group
GroupOrderList

# Session Manager

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Session Manager**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The Session Manager subkey contains data used by many different components of the Session Manager.

# LicensedProcessors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**LicensedProcessors**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 4 |

**Description**
Indicates how many processors the system is licensed to support.

This entry is protected by the Windows licensing agreement. If you change or delete the value or any properties of this entry, the system stops abnormally and displays a STOP 0x0000009A (SYSTEM_LICENSE_VIOLATION) blue screen. The STOP code is followed by the hexadecimal parameter, 0x00000005, indicating that off-line changes were made to the number of licensed processors.

**Caution**

- Deleting or changing this value violates your Windows licensing agreement and disqualifies you from receiving Microsoft Product Support Services services.

# RegisteredProcessors

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RegisteredProcessors**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 2 \| 4 | Windows Server 2003: 4 Windows XP Professional: 2 |

**Description**
This entry is for internal use only. Do not delete or change it.

# Environment Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Environment**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager

**Description**
The Environment subkey stores the values of static environment variables that the system creates. These variables are used by Winlogon and Windows Explorer in Windows NT 4.0 and later, and by Program Manager in Windows NT 3.51 and earlier.

This subkey does not store dynamic environment variables that the system creates. Also, it does not store environment variables that the user creates. User environment variables are stored in the Environment Key subkey.

**Change Method**
To change the values of entries in this subkey, in **Control Panel** double-click **System**, click the **Advanced** tab and then click **Environment Variables**. These entries appear in the **System variables** section.

### Caution

- The data type of entries in this subkey must be REG_EXPAND_SZ. If a user or application changes the data type, the system will not replace the variable with its intended value.

- Use extreme care when changing these the values of entries. If the operating system cannot find the files specified for a subsystem, you will not be able to run applications designed for other environments.

# ComSpec

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ComSpec**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Environment

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | Filename | %SystemRoot%\system32\cmd.exe |

**Description**
Stores the path and file name of the Windows Server 2003 command interpreter, Cmd.exe. The command interpreter is analogous to MS-DOS Command.com.

**Change Method**
To change the value of this entry, use System in Control Panel. Click the **Advanced** tab and then click **Environment Variables**. This entry appears in the **System variables** section.

### Caution

- The data type of this entry must be REG_EXPAND_SZ. If a user or application changes the data type, then the system will not replace the variable with its intended value.

# Path Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Path**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Environment

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Filename* | %SystemRoot%\system32; %SystemRoot% |

**Description**
Stores the path to system files.

**Change Method**
To change the value of this entry, double-click System in Control Panel. Click the **Advanced** tab, and then click **Environment Variables**. This entry appears in the **System variables** box.

### Caution

- The data type of this entry must be REG_EXPAND_SZ. If a user or an application changes the data type, then the system will not replace the variable with its intended value.

# PROCESSOR_ARCHITECTURE

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PROCESSOR_ARCHITECTURE**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Environment

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Processor type* | x86 |

**Description**
Stores the type of processors installed on the system.

For Windows Server 2003, the system always sets the value of this entry to x86. However, this entry remains in the system to anticipate future needs and to maintain compatibility with earlier systems.

# PROCESSOR_LEVEL

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PROCESSOR_LEVEL**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Environment

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Model number* | (None) |

**Description**
Stores the model number of the processor installed on the computer. Programs can use this entry to adjust their configuration for different systems.

# Environment\TEMP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TEMP**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Environment

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %SystemRoot%\Temp |

**Description**
Stores the path to the TEMP temporary directory for the system.

**Change Method**
To change the value of this entry, use System in Control Panel. Click the **Advanced** tab and then click **Environment Variables**. This entry appears in the **System variables** section.

### Caution

- The data type of this entry must be REG_EXPAND_SZ. If a user or application changes the data type, then the system will not replace the variable with its intended value.

# Environment\TMP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TMP**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Environment

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %SystemRoot%\Temp |

**Description**
Stores the path to the TMP temporary directory for the system.

**Change Method**
To change the value of this entry, use System in Control Panel. Click the **Advanced** tab and then click **Environment Variables**. This entry appears in the **System variables** section.

### Caution

- The data type of this entry must be REG_EXPAND_SZ. If a user or application changes the data type, then the system will not replace the variable with its intended value.

# windir

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**windir**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Environment

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Pathname* | %SystemRoot% |

**Description**
Stores the path to the executable file for Win16 on Win32 (WOW).

### Caution

- The data type of this entry must be REG_EXPAND_SZ. If a user or application changes the data type, then the system will not replace the variable with its intended value.

# Memory Management

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Memory Management**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager

**Description**
The Memory Management subkey stores configuration data for paging files.

**Change Method**
To create a new paging file or to change the size of a paging file, double-click **System** in Control Panel, click the **Advanced** tab, click the **Performance Options** button, and then click the **Change button** in the Virtual Memory box.

### Note

- For more information on memory management and paging files, see Windows Server 2003 Help or Windows XP Professional Help.

### Tip

- To create a new paging file or to change the size of a paging file, double-click **System** in Control Panel, click the **Advanced** tab, click the **Performance Options** button, and then click the **Change** button in the Virtual Memory box.

### Caution

- Do not change the size of the paging file by editing the registry. To create a new paging file or to change the size of a paging file, use Control Panel.

# ClearPageFileAtShutdown

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ClearPageFileAtShutdown**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether inactive pages in the paging file are filled with zeros when the system stops. This is a Windows Server 2003 security feature that prevents the pages from being read by another process. The system cannot fill active pages with zeros because some are being used by the system or by other active processes.

| Value | Meaning |
|---|---|
| 0 | Inactive pages are not filled with zeros. |
| 1 | Inactive pages are filled with zeros. |

**Caution**

- Changing this value can degrade the performance of your computer.

# DisablePagingExecutive

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DisablePagingExecutive**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether kernel-mode drivers and kernel-mode system code can be paged to disk when not in use.

| Value | Meaning |
|---|---|
| 0 | Drivers and system code can be paged to disk as needed. |
| 1 | Drivers and system code must remain in physical memory. |

**Activation Method**
To make changes to this entry effective, you must restart the system.

**Tip**

- Setting this value to 1 is useful when debugging drivers, because all of the code and data is always memory resident. It also improves performance on machines with a lot of memory, because it prevents disk reads to get code or data.

# LargeSystemCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**LargeSystemCache**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | Windows XP Professional: 0 Windows Server 2003: 1 |

**Description**
Specifies whether the system maintains a standard size or a large size file system cache, and influences how often the system writes changed pages to disk.

Increasing the size of the file system cache generally improves server performance, but it reduces the physical memory space available to applications and services. Similarly, writing system data less frequently minimizes use of the disk subsystem, but the changed pages occupy memory that might otherwise be used by applications.

| Value | Meaning |
|---|---|
| 0 | Establishes a standard size file-system cache of approximately 8 MB. The system allows changed pages to remain in physical memory until the number of available pages drops to approximately 1,000. This setting is recommended for servers running applications that do their own memory caching, such as Microsoft SQL Server, and for applications that perform best with ample memory, such as Internet Information Services (IIS). |
| 1 | Establishes a large system cache working set that can expand to physical memory, minus 4 MB, if needed. The system allows changed pages to remain in physical memory until the number of available pages drops to approximately 250. This setting is recommended for most computers running Windows Server 2003 on large networks. |

This entry and the **Size** entry (which is in the HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters subkey) store the setting of the **Optimization** section of the **Server Optimization** tab in Network and Dial-up Connections.

| Option setting | Large System Cache value | Size value |
|---|---|---|
| Minimize memory used | 0 | 1 |
| Balance | 0 | 2 |
| Maximize data throughput for file sharing | 1 | 3 |
| Maximize data throughput for network applications | 0 | 3 |

**Change Method**
To change the value of this entry on Windows Server 2003, use the **Server Optimization** tab in Network and Dial-up Connections. Right-click My Network Places, click **Properties**, right-click **Local Area Connection**, click **Properties**, click **File and Printer Sharing for Microsoft Networks**, and then click the **Properties** button. To set the value of this entry to 0, select the **Maximize data throughput for network applications** option. To set the value to 1, select **Maximize data throughput for file sharing**.

**Note**

- Because setting the value of this entry to 1 is not recommended for computers running Windows XP Professional, you cannot use Network and Dial-up Connections in Windows XP Professional to change the value of this entry.

- The Windows Server 2003 default value (1) is designed for file servers. Because this value can degrade service performance, it is not recommended for application servers. If you are running an application server, change the value of this entry to 0 by selecting the **Maximize data throughput for network applications** option in Network and Dial-up Connections.

- The system sets the value of this entry to 1 when you install Windows Server 2003. But many applications, such as SQL Server and Microsoft Exchange, change the value of this entry to 0.

# NonPagedPoolQuota

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NonPagedPoolQuota**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1–128 (megabytes) | 0 |

**Description**
Establishes the maximum space in the nonpaged pool that can be allocated by any process. Any attempt by a process to allocate more space in the nonpaged pool will fail.

| Value | Meaning |
|---|---|
| 0 | The system calculates an optimal value for this entry based on the amount of physical memory on the computer, and it adjusts this value if the amount of memory changes. This value is optimal for most systems. |
| 1–128 | Overrides the value calculated by the system, and prevents the system from adjusting this value dynamically. This setting can degrade performance. |

**Caution**

- Do not change the value of this entry. The system adjusts this value dynamically. Changing the value prevents the system from using its algorithms to calculate an optimal value for the system.

**See Also**
**Concepts**
NonPagedPoolSize
PagedPoolQuota

# NonPagedPoolSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NonPagedPoolSize**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| *Number of bytes* | 0 |

**Description**
Specifies the size of the nonpaged pool.

| Value | Meaning |
|---|---|
| 0 | The system calculates an optimal value for this entry based on the amount of physical memory on the computer, and it adjusts this value if the amount of memory changes. This value is optimal for most systems. |
| *Number of bytes* | Creates a nonpaged pool of the specified size. This takes precedence over the value calculated by the system, and it prevents the system from adjusting this value dynamically. This setting can degrade performance. |

**Activation Method**
You must restart the computer before a change to this value takes effect.

### Caution

- Changing the value of this entry can degrade the performance of your computer.

**See Also**
**Concepts**
NonPagedPoolQuota
NonPagedPoolSize

# PagedPoolQuota

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PagedPoolQuota**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1–128 (megabytes) | 0 |

**Description**
Establishes the maximum space in the paged pool that can be allocated by any process. Any attempt by a process to allocate more space in the paged pool will fail.

| Value | Meaning |
|---|---|
| 0 | The system calculates an optimal value for this entry based on the amount of physical memory on the computer, and adjusts this value if the amount of memory changes. This value is optimal for most systems. |
| 1–128 | Overrides the value calculated by the system and prevents the system from adjusting this value dynamically. This setting can degrade performance. |

**Activation Method**
You must restart the computer before a change to this value takes effect.

### Caution

- Do not change the value of this entry. The system adjusts this value dynamically based on internal algorithms. Changing this value prevents the system from calculating an optimal value. This is likely to degrade the performance of your system.

**See Also**
**Concepts**
NonPagedPoolQuota
NonPagedPoolSize

# PagedPoolSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PagedPoolSize**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0x20000000 *bytes* \| 0xFFFFFFFF | 0x0 |

**Description**
Specifies the size of the paged pool.

| Value | Meaning |
|---|---|
| 0x0 | The system calculates an optimal value for this entry based on the amount of physical memory in the computer, and it adjusts this value if the amount of memory changes. The system algorithm typically sets the size of the paged pool to approximately twice that of the nonpaged pool size. This value is optimal for most systems. |
| 0x1–0x20000000 *bytes* | Creates a paged pool of the specified size. This takes precedence over the value that the system calculates, and it prevents the system from adjusting this value dynamically. Limiting the size of the paged pool to 192 MB (or smaller) lets the system expand the file system (or system pages) virtual address space up to 960 MB. This setting is intended for file servers and other systems that require an expanded file system address space (meaning slightly faster access) at the expense of being able to actually cache less data. This only makes sense if you know the files your server frequently accesses already fit easily into the cache. |
| 0xFFFFFFFF | The system calculates the maximum paged pool allowed for the system in this configuration. For 32-bit systems, this is typically less than 500 MB. This setting is typically used for servers that are attempting to cache a very large number of frequently used small files, some number of very large size files, or both. In these cases, the file cache that relies on paged pool to manage its caching is able to cache more files (and for longer periods of time) if more paged pool is available. |

**Activation Method**
You must restart the computer before a change to this value takes effect.

### Note

- If the paged pool on a system is 192 MB or smaller (including 0, which means auto-size), then the system can expand the file system virtual address space or the system pages virtual address space up to 960 MB. For more information, see the SystemPages entry.

**See Also**
**Concepts**
NonPagedPoolSize
PagedPoolQuota
RegistrySizeLimit

# PagingFiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PagingFiles**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Name Minimum(MB) Maximum(MB)* | (*Varies from system to system*) |

**Description**
Specifies the name of the paging file and the minimum and maximum size range (in megabytes) for the paging file.

**Change Method**
To change the size range of the paging file or to create new paging files, double-click **System** in Control Panel, click the **Advanced** tab, click the **Performance Options** button, and then click the **Change** button in the **Virtual Memory** box.

### Note

- For more information about memory management and paging files, see Windows Server 2003 Help or Windows XP Professional Help.

# SystemPages

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SystemPages**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFE *entries* \| 0xFFFFFFFF | 0x0 |

**Description**
Defines the number of system page table entries that are reserved for mapping I/O buffers and other information into the system address space. Each system page table entry maps one page.

| Value | Meaning |
|---|---|
| 0x0 | The system calculates an optimal number of page table entries based on the platform type and the amount of memory available to the system. The system adjusts this value if the amount of memory changes. |
| 0x1–0xFFFFFFFE | Specifies the number of reserved page table entries. This value takes precedence over the system's calculated value, and it prevents the system from adjusting the value. |
| 0xFFFFFFFF | Maximize system pages. The system creates the largest number of page table entries possible within physical memory. The system monitors and adjusts this value dynamically when the configuration changes. |

**Caution**

- Do not change the value of this entry. Changing it prevents the system from calculating an optimal value for your system and adjusting the value when your system changes.

**See Also**
**Concepts**
NonPagedPoolSize

# SecondLevelDataCache

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SecondLevelDataCache**
HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1–*Cache size (in bytes)* | 0 (*256 KB*) |

**Description**
Specifies the number of bytes in the secondary processor cache, also known as the L2 cache. The value of this entry is used only when the Hardware Abstraction Layer (HAL) cannot detect the L2 cache.

| Value | Meaning |
|---|---|
| 0 | Sets the L2 cache size to 256 KB. |
| *Number of bytes* | Sets the L2 cache size to the value specified. |

# VirtualDeviceDrivers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VirtualDeviceDrivers**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The VirtualDeviceDrivers subkey stores configuration data for virtual device drivers.

# VDD

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VDD**
HKLM\SYSTEM\CurrentControlSet\Control\VirtualDeviceDrivers

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Filenames* | (None) |

**Description**
Specifies the Win32 dynamic-link libraries (DLLs) for virtual device drivers. The system loads the virtual device drivers in this list when a Windows Server 2003 virtual DOS machine is started. By default, this entry has no value because the Windows Server 2003 virtual device drivers are built into the system and do not need to be loaded separately.

# Control\Windows Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows**
HKLM\SYSTEM\CurrentControlSet\Control

**Description**
The Windows subkey contains configuration data for the Win32 subsystem in Windows Server 2003.

**See Also**
**Concepts**
Microsoft\Windows
Microsoft\Windows Subkey

# NoInteractiveServices

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NoInteractiveServices**
HKLM\SYSTEM\CurrentControlSet\Control\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Prevents background services from interacting with users. This entry was designed to improve performance of unattended computers.

This entry cannot be used to make a service interact with a user unless the service is written to be interactive. Also, it will not prevent a service from displaying error messages.

| Value | Meaning |
|---|---|
| 0 | Interactive background services (Type 0x110) can communicate with users. |
| 1 | Background services cannot interact with users. This value takes precedence over the Control Panel setting for interactive services (as described below) and the SERVICE_INTERACTIVE_PROCESS option in the CreateService API. |

### Note

- To determine whether a service is run in interactive mode, from the **Start** menu, point to **Programs**, point to **Administrative Tools**, and then double-click **Computer Management**. Double-click **System Tools**, double-click **Services**, right-click a service to see its properties, and then select the **Log On** tab. The **Allow Service to Interact with Desktop** check box is selected if the service is interactive, and cleared if it is not. Selecting the check box will not make a service interact with the user unless it is written to do so.

# NoPopUpsOnBoot

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NoPopUpsOnBoot**
HKLM\SYSTEM\CurrentControlSet\Control\Windows

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether pop-up error messages that appear when the system starts are suppressed. You can add **NoPopUpsOnBoot** to the registry to prevent Windows Server 2003 from displaying messages when no one is there to read them, such as computers that are configured to log on to Windows Server 2003 automatically.

| Value | Meaning |
|---|---|
| 0 | Pop-up error messages are displayed. |
| 1 | Pop-up error messages are suppressed. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- This entry does not suppress the message notifying the user that the domain controller cannot be contacted. To suppress that message, use ReportDC.

**See Also**
**Concepts**
ReportDC
ReportControllerMissing

# Enum Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet

**Description**
The Enum subkey contains a database representing all devices installed on the computer and recognized by the system. The database stores configuration data for hardware devices independent of the drivers they use. This subkey is critical to the configuration of Plug and Play devices in Windows.

**Change Method**
To change device settings, use Device Manager.

### Note

- Users, including administrators, cannot change the contents of the Enum subkey. These restrictions protect the integrity of the operating system and the devices on the computer.

### Caution

- Do not attempt to change the permissions on the Enum subkey. If you do, you might damage the sequence of inherited permissions in the registry. If you change the content of the Enum subkey, you can damage the Plug and Play functions of Windows Server 2003 and prevent devices from operating. If you must gain access to the Enum subkey, contact Microsoft Product Support Services.

**See Also**
**Concepts**
service-name\Enum

# Root

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# *PNP03<#>

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Device Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP03<#>\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# KeyboardDataQueueSize Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP03<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1 - 0xFFFFFFFF | 0x64 (100) |

**Description**
Specifies the number of keyboard events buffered by the keyboard driver. It is also used in calculating the size of the keyboard driver's internal buffer in nonpaged memory pool. To determine the number of bytes to allocate for the buffer, the system multiplies the size of the KEYBOARD_INPUT_DATA structure (defined in NTDDKBD.H) by the value of **KeyboardDataQueueSize**.

Consider increasing this value if the driver does not initialize or operate correctly and the System Log in Event Viewer contains the following message from the source, Inport:

```
"The ring buffer that stores incoming keyboard data has overflowed (buffer size is configurable via the Registry)."Text goes here.
```

### Note

- This entry is created by copying the value of Parameters\KeyboardDataQueueSize from the i8042prt subkey. Thereafter, the system uses only the entries in this location.

**See Also**
**Concepts**
Parameters\KeyboardDataQueueSize

# Device Parameters\PollStatusIterations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP03<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1 - 0xFFFFFFFF | 0x1 |

**Description**
Specifies the maximum number of times the system verifies interrupts on the i8042 controller status register. If the interrupt cannot be verified in the number of trials specified in **PollStatusIterations**, then the interrupt is ignored.

This entry is designed to minimize the effect of interrupts raised before the Output Buffer Full bit is set in the i8042 status register. This condition occurs on some computers, such as the Olivetti MIPS computer.

If the system seems to ignore keyboard interrupts, then increase the value of this entry. The system might mistakenly ignore an interrupt if the Output Buffer Full bit is not set when the interrupt is raised. However, if you increase the value of **PollStatusIterations**, then the Interrupt Service Routine will take longer to execute because there is a one microsecond delay following each check for Output Buffer Full.

### Note

- This entry is created by copying the value of PollStatusIterations Entry from the i8042prt subkey. Thereafter, the system uses only the entries in this location.

### Tip

- To determine whether the driver is processing keyboard interrupts correctly, press the NUM LOCK key. If the **Num Lock** light on the keyboard turns on or off, then this indicates that the i8042prt driver handled the keyboard interrupt correctly.

**See Also**
**Concepts**
PollStatusIterations Entry

# *PNP0F<#>

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# *PNP0F<#>\Device Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# EnableWheelDetection Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 1 |

**Description**
Specifies whether the driver attempts to detect and enable the wheel on the mouse device. Some devices are equipped with a mouse wheel to provide rapid scrolling and other control features when supported by an application.

| Value | Meaning |
|---|---|
| 0 | Disable mouse wheel detection. The driver does not attempt to determine whether the mouse has a wheel or to enable it. Even if the mouse has a wheel, the wheel does not function in programs. |
| 1 | Detect and enable the mouse wheel. If the driver detects a wheel on the mouse, then it enables the wheel. |
| 2 | Enable without detecting. The driver attempts to enable a mouse wheel without determining whether the mouse has a wheel. |

### Note

- This entry is created by copying the value of Parameters\EnableWheelDetection from the i8042prt subkey. Thereafter, the system uses only the entries in this location.
- For mouse devices that do not fully support all Intellimouse and PS/2 functions, including dynamic enabling, set the value of this entry to 2.

### Tip

- To view or change the value of this entry for a particular PS/2 mouse, use Device Manager. In the tree control, click **Device** and then, in the results pane, click **Mouse**, double-click the device name, and click the **Advanced Settings** tab. This value appears in the **Wheel Detection** box.

### Caution

- In rare instances, when the system tries to detect and enable the mouse wheel, the system stops responding. To recover, restart Windows. If the problem persists, restart by using the LastKnownGood control set.

**See Also**
**Concepts**
Parameters\EnableWheelDetection
WheelDetectIDs

# MouseDataQueueSize Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1 - 0xFFFFFFFF | 0x64 (100) |

**Description**
Specifies the number of mouse events buffered by the mouse driver. It is also used in calculating the size of the mouse driver's internal buffer in nonpaged memory pool. To determine the number of bytes to allocate for the buffer, the system multiplies the size of the MOUSE_INPUT_DATA structure (defined in NTDDMOU.H) by the value of **MouseDataQueueSize**.

Consider increasing this value if the driver does not initialize or operate correctly and the System Log in Event Viewer contains the following message from the source, Inport :

```
The ring buffer that stores incoming mouse data has overflowed (buffer size is configurable via the
Registry).
```

### Note

- This entry is created by copying the value of Parameters\MouseDataQueueSize from the i8042prt subkey. Thereafter, the system uses only the entries in this location.

### Tip

- To view or change the value of this entry for a particular PS/2 mouse, use Device Manager. In the tree control, click **Device** and then, in the results pane, click **Mouse**, double-click the device name, and click the **Advanced Settings** tab. This value appears in the **Input Buffer Length** box.

**See Also**
**Concepts**
Parameters\MouseDataQueueSize

# MouseInitializePolled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies how the mouse is initialized. The default value, whereby the driver initializes the mouse in response to an interrupt, is quicker. Polling can increase the time it takes to start Windows by 5 to 15 seconds. Do not change the default unless the system does not respond to the mouse or keyboard when the **Logon Information** (CTRL+ALT+DEL) screen is displayed.

| Value | Meaning |
|---|---|
| 0 | In response to an interrupt. A hardware interrupt signals the PS/2 mouse driver to initialize the mouse. |
| 1 | By polling the hardware. The driver attempts to initialize the mouse. If it is unsuccessful, it tries at regular intervals until it succeeds. |

**Tip**

- To view or change the value of this entry for a particular PS/2 mouse, use Device Manager. In the tree control, click **Device** and then, in the results pane, click **Mouse**, double-click the device name, and click the **Advanced Settings** tab. This value is set by using the **Fast Initialization** check box. If the box is checked, the mouse is initialized in response to an interrupt; if it is cleared, it is initialized by polling.

# MouseResolution Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Mouse resolution specifier* | 0x3 |

**Description**
Specifies the resolution of the mouse port. The number of counts per millimeter of resolution is determined by raising 2 to the power of the value of **MouseResolution**.

For example, if the value of MouseResolution is 3, then the mouse resolution is 2**3, or 8, counts per millimeter.

### Note

- This entry is created by copying the value of Parameters\MouseResolution from the i8042prt subkey. Thereafter, the system uses only the entries in this location.

**See Also**
**Concepts**
Parameters\MouseResolution

# MouseSynchIn100ns Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Time in 100-nanosecond units* | 20000000 (2 seconds) |

**Description**
Specifies the length of time after which the next mouse interrupt is assumed to indicate the start of a new mouse packet (partial packets are discarded). This enables the driver to synchronize its internal notion of the mouse packet state with the hardware state when a mouse interrupt is lost.

Consider modifying this value if the system appears to be generating extra mouse events, such as when objects are selected although the mouse button has not been pressed.

### Note

- This entry is created by copying the value of Parameters\MouseSynchIn100ns from the i8042prt subkey. Thereafter, the system uses only the entries in this location.

**See Also**
**Concepts**
Parameters\MouseSynchIn100ns

# Device Parameters\NumberOfButtons

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1 - 0xFFFFFFFF | 0x2 |

**Description**
Specifies the number of buttons on the mouse port mouse at startup. If the number of buttons detected at startup is incorrect, then you can override it by changing the value of this entry.

### Note

- This entry is created by copying the value of NumberOfButtons Entry from the i8042prt subkey. Thereafter, the system uses only the entries in this location.

**See Also**
**Concepts**
NumberOfButtons Entry

# Device Parameters\SampleRate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 20 \| 40 \| 60 \| 80 \| 100 \| 200 | 100 Hz |

**Description**
Specifies how often, in Hertz, the PS/2 driver measures the characteristics and activities of the PS/2 mouse. The driver uses the information gathered by sampling to optimize the operation of the mouse device.

### Note

- This entry is created by copying the value of SampleRate Entry from the i8042prt subkey. Thereafter, the system uses only the entries in this location.

### Tip

- To view or change the value of this entry for a particular PS/2 mouse, use Device Manager. In the tree control, click **Device** and then, in the results pane, click **Mouse**, double-click the device name, and click the **Advanced Settings** tab. This value appears in the **Sample Rate** box.

**See Also**
**Concepts**
SampleRate Entry

# WheelDetectIDs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *ASCII identification string* | (None) |

**Description**
A list of ASCII identification strings of wheel mice enabled on the system. By storing the identities of mouse wheel devices that have already been enabled, the system avoids redetecting and reenabling the device. This entry is used only when the value of EnableWheelDetection Entry is 1.

This entry is used to associate wheel mouse devices newly detected on the system with those already registered. When the system acquires an identification string from a new wheel mouse, it converts the string from scan codes into ASCII and compares it to the values stored in this entry. If the new string matches one of the values in this entry, then the mouse wheel is enabled without further delay. If no match is found, then the system attempts to enable the mouse wheel.

### Note

- Unlike some entries in this subkey, this entry remains in the Parameters subkey and is read from this location even after a PS/2 mouse is installed.

**See Also**
**Concepts**
Parameters\EnableWheelDetection
EnableWheelDetection Entry

# WheelDetectionTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Root*PNP0F<#>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 - 0xFFFFFFFF *microseconds* | 5DC (1.5 seconds) |

**Description**
Specifies how long the mouse driver waits between attempts to detect and enable the mouse wheel. This entry is used when the mouse driver is attempting to enable the mouse wheel. It is used only when the value of EnableWheelDetection Entry is 1 or 2.

By default, this entry does not have a value. However, the default interval is 1.5 seconds in the absence of a value for this entry.

### Note

- To override the default value, you must add this entry to the registry.

**See Also**
**Concepts**
EnableWheelDetection Entry
Parameters\EnableWheelDetection

# Serenum

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# <Device-name>

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Serenum\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# <Device-name>\Device Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Serenum\<*Device-name*>\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Device Parameters\NumberOfButtons Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Serenum\<*Device-name*>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1 - 0xFFFFFFFF | 0x2 |

**Description**
Stores the number of buttons on the serial mouse that are detected when the system starts. If the number of buttons detected at startup is incorrect, then you can override it by changing the value of this entry.

If you have upgraded from Windows NT 4.0 or earlier, then the system copies the value of the Parameters\NumberOfButtons Subkey entry in the HKLM\SYSTEM\CurrentControlSet\Services\Sermouse\Parameters subkey to the device-specific Enum subkeys of all new serial mouse devices that are enumerated by the system.

### Note

- The precise path to this entry varies with the device.

### Tip

- To access this value programmatically, use the SetupDi API or Configuration Manager API.

**See Also**
**Concepts**
Parameters\NumberOfButtons Subkey

# Device Parameters\SampleRate Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Serenum\<*Device-name*>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Mouse sampling rate in Hz* | 0x28 (40 Hz) |

**Description**
The system adds this entry to the registry, but does not use it. The Microsoft Serial Mouse Driver does not sample mouse activity.

### Note

- The precise path to this entry varies with the device.

**See Also**
**Concepts**
Parameters\SampleRate Subkey

# WaitEventMask

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Enum\Serenum\<*Device-name*>\Device Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1 - 0xFFFFFFFE \| 0xFFFFFFFF | 0x0 |

**Description**

Specifies whether the serial mouse driver detects when the serial mouse has been replaced on a serial port by another device.

The serial mouse alternates its use of serial ports with other serial devices, such as modems. By monitoring the state of modem lines, the driver can infer when the serial mouse has been removed dynamically from the port in favor of another device.

To monitor the modem lines, the serial mouse uses state data gathered by the serial port driver at initialization. From this data, it builds a bitmask representing the state of modem lines. A bit is set to 1 to indicate that a CTS (clear-to-send) or DSR (data-set-ready) signal is detected and that the line should be monitored for changes. It compares the bitmask to new data to detect changes in modem state and then infers when the serial mouse has been removed.

| Value | Meaning |
|---|---|
| 0x0 | Enables device removal detection. The serial driver creates a bitmask that represents all modem lines and monitors them for changes that indicate the removal of a device. |
| 0x1 - 0xFFFFFFFE | A bitmask representing the modem lines. If you enter a bitmask, the serial mouse driver uses the value you enter instead of building its own bitmask. |
| 0xFFFFFFFF | Disables serial device removal detection. The serial driver does not monitor modem lines. Instead, users must inform the system when modems can be removed. |

**Note**

- The precise path to this entry varies with the device.

**Caution**

- Do not change the value of this entry. If you enter a bitmask in this value, the serial mouse driver no longer creates a bitmask dynamically; Instead, it uses the mask you provide.

# CurrentControlSet\Hardware Profiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet

**Description**
The Hardware Profiles subkey stores the hardware profiles on your system.

A *hardware profile* consists of specified changes to the original hardware profile configured in the HKLM\SOFTWARE and SYSTEM keys of HKEY_LOCAL_MACHINE. To allow hardware profiles to be easily applied and removed, the original value in the SOFTWARE or SYSTEM key is not changed. Instead, the change is stored in a numbered subkey of the Hardware Profiles subkey.

You can use System in Control Panel to create a hardware profile. Also, the system creates new hardware profiles automatically when it encounters configurations that require one, such as when a portable computer is undocked or moved from one docking station to another.

Each hardware profile is stored in a *Profile-number* subkey of the Hardware Profiles subkey. The content of the numbered subkeys looks like a miniature version of the Software and System subkeys. It contains only those subkeys in the path to the entry that the hardware profile changes. When a hardware profile is applied, the values in the numbered subkeys override the corresponding values in Software and System.

The current hardware profile is stored in one of the numbered subkeys in the Hardware Profiles subkey, but it also appears in the Hardware Profiles\Current subkey of Hardware Profiles and in the HKEY_CURRENT_CONFIG subtree of the registry. Because the Current subkey and the HKEY_CURRENT_CONFIG subtree merely store pointers to the content of the numbered Hardware Profiles subkey currently in use, you can view and edit the data in any of these three locations.

### Note

- For information about the entries in a subkey of Hardware Profiles, see the description of that entry in the SOFTWARE or SYSTEM keys. For example, to find a description of the contents of Hardware Profiles*Profile-number*\Software\Fonts, see the description of HKLM\Software\Fonts.

### Tip

- To create hardware profiles, use System in Control Panel. Click the **Hardware** tab, and then click **Hardware Profiles**.

**See Also**
**Concepts**
Hardware Profiles\Current
HKEY_CURRENT_CONFIG

# Hardware Profiles\Current

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Hardware Profiles

**Description**
The Current subkey in the Hardware Profiles subkey displays configuration data for the hardware profile currently in use on the system.

The Current subkey does not actually contain any data. Instead, it contains a pointer to the *Profile-number* subkey in Hardware Profiles, which contains the current hardware profile. As such, the contents of the selected subkey appear in the Current subkey.

Also, the contents of the Current subkey appear in the HKEY_CURRENT_CONFIG subtree of the registry. Because the Current subkey and HKEY_CURRENT_CONFIG merely store pointers to the content of the numbered Hardware Profiles subkey currently in use, you can view and edit the data in any of the three locations.

The system always designates one profile as the current one.

- If the system does not have any hardware profiles that are appropriate to the current hardware configuration, then the system creates one automatically, and that profile becomes the current hardware profile.
- If the system includes only one hardware profile that is appropriate to the current hardware configuration, then that profile is automatically chosen to be the current hardware profile.
- If the system includes more than one hardware profile that is appropriate to the current hardware configuration, then when the system starts, it prompts the user to select from among the appropriate hardware profiles. The chosen profile becomes the current profile.

The current hardware profile remains current until the system stops or the hardware configuration changes, such as when an ACPI-capable (Advanced Configuration and Power Interface) portable computer docks or undocks.

**Tip**

- To determine which numbered subkey in Hardware Profiles represents the current hardware profile, see the CurrentConfig entry in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\IDConfigDB. The value of **CurrentConfig** corresponds to the number of the subkey that contains the current hardware profile.

**See Also**
**Concepts**
CurrentControlSet\Hardware Profiles
HKEY_CURRENT_CONFIG
CurrentConfig

# Services Subkeys

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet

**Description**
The Services subkeys contain entries for standard and optional Windows Server 2003 services, such as device drivers, file system drivers, and Win32 service drivers. Although the values of the entries differ for each service, most Services subkeys have the same subkeys and entries.

Each Services subkey bears the name of the service that uses it. Often, this is also the name of the file from which the service is loaded.

Some of the services and devices represented by subkeys in the Services subkey are installed on the computer, but some of the subkeys represent services that are not installed or not enabled. To determine which services are installed on the computer, click **Services** in Computer Management. To determine which devices are installed on the computer, use Device Manager.

**Note**

- Adding a subkey to the Services subkey does not install a service. Windows Server 2003 does not recognize a service unless it is properly installed using standard installation methods, such as the CreateService API or the SC CREATE command.

# service-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
Represents a service for Windows Server 2003.

Entries and subkeys common to Services subkeys appear in the *service-name* subkey. Entries and subkeys that are specific to a service are described in the subkey bearing the service name.

**Note**

- The *service-name* subkey is a variable representing the subkeys in the HKLM\SYSTEM\CurrentControlSet\Services subkey. It does not actually appear in the registry. This variable subkey displays the entries and subkeys that are common to many of the Services subkeys. The actual Services subkeys in the registry depend on the services installed on the computer or available to the system.

# ErrorControl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–3 | 1 |

**Description**
Specifies how to proceed if the driver for the service or device fails to load or initialize properly. By default, the value of this entry for all Windows Server 2003 services and drivers is 1.

| Value | Meaning |
|---|---|
| 0 | Ignore. Startup proceeds with no warning messages. |
| 1 | Normal. Startup proceeds, but a warning message appears. |
| 2 | Severe. The system declares that the startup has failed, and it restarts by using the LastKnownGood control set. If startup is already using the LastKnownGood control set, then it continues. |
| 3 | Critical. The system declares that the startup has failed, and it restarts by using the LastKnownGood control set. If startup is already using the LastKnownGood control set, then the system stops the startup procedures and runs a debugging program. |

**Change Method**
To change the value of this entry, you can use the Services snap-in. You can also use the ChangeServiceConfig API or use SC, a tool included in Windows Server 2003.

### Note

- This information is presented for reference only. You cannot change the load configuration of a service by editing this entry.

# Group

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Service group name* | (None) |

**Description**
Specifies the name of the service group to which the service belongs. If this entry does not appear in the registry, then the service does not belong to a service group. As a result, it will be loaded after all services in service groups are loaded.

**Change Method**
To change the value of this entry, you can use the Services snap-in. You can also use the ChangeServiceConfig API or use SC, a tool included in Windows Server 2003.

### Note

- This information is presented for reference only. You cannot change the service group of a service by editing this entry.

# service-name\ImagePath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path and file name* | *Systemroot*\System32\Drivers*Driver-name*.sys (Drivers) *Systemroot*\System32*Service-name*.exe (Services) |

**Description**

Specifies the location of the executable file for the driver or service. This entry is not used for network adapters.

**Caution**

Do not change the value of this entry.

# ObjectName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Object name* | Drivers: *Systemroot*\System32\Drivers*Driver-name*.sys Services: *Systemroot*\System32*Service-name*.exe |

**Description**
Specifies the name of the Windows Server 2003 driver object that I/O Manager uses to load the device driver, if the service represented by the subkey is either a kernel-mode driver or a file system driver.

If the service represented by the subkey is a Win32 service, then this entry specifies the account name that the service uses to log on to Windows Server 2003.

**Change Method**
To change the value of this entry, you can use the Services snap-in. You can also use the ChangeServiceConfig API or use SC, a tool included in Windows Server 2003.

### Note

- This information is presented for reference only. You cannot edit this entry to change the driver object or account name that Windows Server 2003 uses to start a service.

### Tip

- To determine what type of service the subkey represents, see the Type Entry entry.

# Start Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–4 | (None) |

**Description**
Specifies how the service is loaded or started. If the service is a Win32 service, then the value of this entry must be 2, 3, or 4. This entry is not used for network adapters.

| Value | Meaning |
|---|---|
| 0 | Boot (loaded by kernel loader). Components of the driver stack for the boot (startup) volume must be loaded by the kernel loader. |
| 1 | System (loaded by I/O subsystem). Specifies that the driver is loaded at kernel initialization. |
| 2 | Automatic (loaded by Service Control Manager). Specifies that the service is loaded or started automatically. |
| 3 | Manual. Specifies that the service does not start until the user starts it manually, such as by using Device Manager. |
| 4 | Disabled. Specifies that the service should not be started. |

**Change Method**
To change the value of this entry, you can use the Services snap-in. You can also use the ChangeServiceConfig API or use SC, a tool included in Windows Server 2003.

### Note

- This information is presented for reference only. You cannot change the starting rules for a service by editing this entry.

### Tip

- To determine what type of service is represented by the subkey, see Type Entry.

# Type Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 \| 2 \| 4 \| 8 \| 16 \| 32 \| 272 \| 288 | (None) |

**Description**
Identifies the type of service represented by the subkey.

| Value | Meaning |
|---|---|
| 1or 0x1 | A kernel-mode device driver. |
| 2 or 0x2 | A file system driver. |
| 4 or 0x4 | A set of arguments for an adapter. |
| 8 or 0x8 | A file system driver service, such as a file system recognizer. |
| 16 or 0x10 | A Win32 program that runs in a process by itself. This type of Win32 service can be started by the service controller. |
| 32 or 0x20 | A Win32 program that shares a process. This type of Win32 service can be started by the service controller. |
| 272 or 0x110 | A Win32 program that runs in a process by itself (like Type16) and that can interact with users. |
| 288 or 0x120 | A Win32 program that shares a process and that can interact with users. |

**Change Method**
To change the value of this entry, you can use the Services snap-in. You can also use the ChangeServiceConfig API or use SC, a tool included in Windows Server 2003.

### Note

- This information is presented for reference only. You cannot edit this entry to change the view of a service by Windows Server 2003.

**See Also**
**Concepts**
Group
ServiceGroupOrder
GroupOrderList

# DependOnGroup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Service group name* [*service group name*] | (None) |

**Description**
Specifies the service group(s) that must be loaded before this service is loaded. This entry ensures that prerequisite services are loaded successfully.

**Change Method**
To change the prerequisite services for a service, use the ChangeServiceConfig API or use SC, a tool included in Windows Server 2003.

### Note

- This information is presented for reference only. You cannot change the prerequisite services for a service by editing this entry.

# DependOnService

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Service group name* [*service name*] | (None) |

**Description**
Specifies the service(s) that must be loaded before this service is loaded. This entry ensures that prerequisite services are loaded successfully.

**Change Method**
To change the value of this entry, you can use the Services snap-in. You can also use the ChangeServiceConfig API or use SC, a tool included in Windows Server 2003.

### Note

- This information is presented for reference only. You cannot change the prerequisite services for a service by editing this entry.

# service-name\Enum

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

**Description**
The Enum subkey in each Services subkey stores information about the hardware associated with that service. The values in the Enum subkey associate the service with the instance of the device it controls.

The Enum subkey appears in subkeys that are in the Services subkey and that store values for drivers and other services controlling or interacting with devices.

**Note**

- Values of entries in this subkey are rewritten each time the system starts. Do not change the values of entries in this subkey.

**See Also**
**Concepts**
Enum Subkey

# service-name\Linkage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

**Description**
The Linkage subkey stores configuration data for binding network components. The bindings associate network services with protocols and with devices that support them.

The Linkage subkey appears in the Services subkeys for network services and protocols.

Hardware devices, such as network adapters, and network services, such as File and Print Services for Microsoft, also have Linkage subkeys. These Linkage subkeys store data about bindings from the perspective of the device. They associate devices and services with the transport protocols they use. In Windows Server 2003, the Linkage Subkey are stored in HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Class*{Class-GUID}*, in the numeric subkey representing the particular device.

Windows Server 2003 creates the bindings from data stored in the Linkage subkeys in both locations.

A binding consists of three entries, Bind, Export, and Route Entry. Each of these entries can include more than one component. The components are listed in the entries in a precise order, so that the first component in one entry corresponds with the first component in the other entries, and each subsequent component corresponds with subsequent components in the other entries.

### Note

- Windows Server 2003 does not include NetRules subkeys. In Windows NT 4.0 and earlier, the NetRules subkeys contained entries that specified bind rules for network software components. Some bind specifications are stored in the Ndi subkeys that contain configuration data for network driver installers. However, these are primary upper bind and lower bind settings only.

### Caution

- Do not change the values of these entries. Bindings are created by the system.

**See Also**
**Concepts**
Linkage Subkey
Ndi Subkeys

# Bind

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*\Linkage

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *ObjectName ObjectName* ... | (None) |

**Description**
Specifies the services, protocols, and devices to which this service is associated in a binding. The system builds the bind from the values in this entry and in the **Bind** entries of the associated protocols and devices.

More than one component of the same type can be listed in this value. For example, if the computer has two networks, then the value of **Bind** for a network adapter type contains two object names.

### Note

- A valid binding requires a value for three entries: **Bind**, Export, and Route Entry.

### Caution

- Bindings are created by the system. Do not change the values of these entries.

**See Also**
**Concepts**
Linkage Subkey

# Export

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*\Linkage

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *ObjectName ObjectName* ... | (None) |

**Description**
Specifies objects that the service creates to enable other services to access it. Objects named in the value of this entry must be installed on the system and available to the service or device.

### Note

- A valid binding requires a value for three entries: Bind, **Export**, and Route Entry.

### Caution

- Bindings are created by the system. Do not change the values of these entries.

**See Also**
**Concepts**
Linkage Subkey

# OtherDependencies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*\Linkage

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *ServiceName ServiceName* ... | (None) |

**Description**
Specifies services that the network component requires, in addition to those specified in the binding. This entry and its value are retained when bindings are reconfigured, even when the other entries in the Linkage subkey are deleted and recreated.

For example, in the NetBT\Linkage subkey, the value of the **OtherDependencies** entry is **Tcpip**. This value indicates that NetBIOS over TCP/IP (NetBT) requires TCP/IP, regardless of its current binding configuration.

### Notes

- A valid binding requires a value for three entries: Bind, Export, and Route Entry.
- Windows Server 2003 adds this entry to the subkey only if the network component requires it. This entry does not appear in all Linkage subkeys.

### Caution

- Bindings are created by the system. Do not change the values of these entries.

**See Also**
**Concepts**
Linkage Subkey

# Route Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*\Linkage

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *SubkeyName SubkeyName* ... | (None) |

**Description**
Specifies the Linkage subkeys from which data was collected to create the binding. Each name in the value of this entry indicates a path through the binding protocol.

The Linkage subkeys that appear in the value of this entry can represent services or devices.

- Linkage subkeys for services bear the name of the service (such as Tcpip) and appear in HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services.

- Linkage subkeys for devices are named by the globally unique identifier (GUID) of the device class and appear in HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Class.

### Note

- A valid binding requires a value for three entries: Bind, Export, and Route Entry.

### Caution

- Bindings are created by the system. Do not change the values of these entries.

**See Also**
**Concepts**
Linkage Subkey

# service-name\networkprovider

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

**Description**
The networkprovider subkey stores the name of the device and the provider path. It can appear in the Services subkeys for a network service. If a network service has a networkprovider subkey, the value of the Group entry in the *service-name* subkey must be **NetworkProvider**.

**See Also**
**Concepts**
DUPTILE1_NetworkProvider

# service-name\Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

**Description**
The Parameters subkeys contain entries specific to each service, such as a registry entry to enable an optional component of the service. Many entries in the Parameters subkeys are not added to the registry when you install the service. You must add them by editing the registry or by using a snap-in or administrative program that adds the entries to the registry.

# Performance Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

**Description**
The Performance subkey contains entries for the performance counter dynamic-link library (DLL) for the service. The performance library uses these values to collect data for the performance counters that monitor the service. These entries are maintained by the system and should not be changed.

**Activation Method**
To make changes to entries in this subkey effective, restart the Remote Registry Service or restart Windows Server 2003.

### Tip

- If particular performance counters do not appear in your performance monitoring tool, check the Application Log in Event Viewer for events recorded by Perflib. When Perflib detects disabling errors in a counter DLL, it records an event in the log and sets the value of Disable Performance Counters to 1 to disable the counters.

# Collect Timeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*\Performance

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xEE6B2800 *milliseconds* | (The value of OpenProcedureWaitTime.) |

**Description**
Defines a slow Collect procedure call for these counters.

This entry tells the Performance Library (Perflib) how long the Collect procedure call for these counters is expected to run. If the time specified in the value of this entry expires before the call completes, then Perfib considers the call to be abnormally slow, and it records an error in the Application Log in Event Viewer. Also, if the value of the Configuration Flags entry requires it, then Perflib disables the counters and does not display them again until you manually re-enable them.

Perflib times the procedure calls of non-trusted counters to identify counters that can delay or lock performance tools during initialization. While the Collect procedure call proceeds, you cannot use the performance tools, and remote users cannot log on to the local computer.

**Activation Method**
To make changes to this entry effective, restart the Remote Registry Service or restart Windows.

### Notes

- Timing procedure calls is one of several tests Perflib performs to maintain the integrity of the counters. To configure the Perflib test feature, use the Configuration Flags entry.
- Perflib does not time the procedure calls of trusted performance counter dynamic-link libraries (DLLs). This entry applies only to non-trusted DLLs.
- Perflib disables counters because of slow procedure calls only when the 0x8 (1000) bit is set and the 0x2 (10) bit is not set in the Configuration Flags entry.
- Perflib does not time procedure calls when the 0x4 (100) bit is set in the Configuration Flags entry.
- Perflib disables counters by setting the value of the Disable Performance Counters entry for that service to 1. After you have repaired or replaced the counter DLLs, you must re-enable them manually by changing the value of **Disable Performance Counters** to 0.
- Some performance counters DLLs add this entry to the registry. If this entry does not appear, you can add it to the registry to override the default value.

### Tip

- If you suspect problems with procedure calls, check the Application Log in Event Viewer. Perflib records an error in the log whenever a procedure call exceeds its limit.

**See Also**
**Concepts**
OpenProcedureWaitTime
Open Timeout

# Disable Performance Counters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*Service-name*\Performance

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Enables and disables registry-based performance counters for this service.

If the value of this entry is 1, then the Performance Library (Perflib) does not retrieve performance data about these counters from the registry. As a result, System Monitor and other tools that use the data cannot display it. Instead, the tools display a value of 0 or 100 percent, depending on the counter.

Perflib uses this entry to disable counters that cause serious errors. The disabling prevents those counters from jeopardizing the display of all performance counters on the system. You can use this entry to re-enable the counters for a service after the counter dynamic-link library (DLL) has been repaired or replaced.

| Value | Meaning |
|---|---|
| 1 | Disables performance counters for this service. |
| 2 | The Win32 version is disabled on computers running a Win64 version of the operating system. |
| 4 | The Win64 version of the service is disabled. |

**Change Method**
To change the value of this entry, use *Extensible Performance Counter List* (Exctrlst.exe) available with *Windows Server 2003 Resource Kit Tools*.

**Activation Method**
To make changes to this entry effective, restart Windows Server 2003.

### Note

- This entry affects only the performance counters for a particular service. To enable or disable all registry-based performance counters on the system, add the Disable Performance Counters Entry entry to the Perflib subkey.
- To prevent Perflib from disabling counters, set the 0x2 (10) bit in the value of the Configuration Flags entry.

### Tip

- If particular performance counters do not appear in your performance monitoring tool, then check the Application Log in Event Viewer for events recorded by Perflib. When Perflib detects disabling errors in a counter DLL, it records an event in the log and sets the value of Disable Performance Counters to 1 to disable the counters.

### Caution

- If you set the value of this entry to 1, then the performance objects and counters for this service do not appear in any performance tool. You will not be able to use System Monitor, nor any tool that uses Performance Library data, to monitor these counters.
- Do not use the registry editor Regedit.exe or Exctrlst.exe to enable counters that Perflib has disabled, unless the counter DLL has been repaired or replaced. If you do, then Perflib will detect the errors and disable the counters again. In addition, the troublesome counters can prevent you from monitoring any performance counters on your system.

**See Also**

**Concepts**

EventLogLevel Entry
Disable Performance Counters
ExtCounterTestLevel
OpenProcedureWaitTime
Open Timeout
Collect Timeout

# Library Validation Code

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*Service-name*\Performance

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *Unique code signature* | (None) |

**Description**
Identifies a trusted performance counter dynamic-link library (DLL).

This entry stores data identifying a particular version of a performance counter DLL as one that has been certified to be free of disabling errors. When the data for a DLL matches the value of this entry, the Performance Library (Perflib) considers the DLL to be trusted.

This entry is added to the registry by the script that installs the DLL in Windows Server 2003.

**Activation Method**
To make changes to this entry effective, restart the Remote Registry Service or restart Windows.

### Note

- By default, trusted counter DLLs are exempt from Perflib's extensive run-time testing, but you can use ExtCounterTestLevel to require some tests. **ExtCounterTestLevel** applies to all trusted counters on the system, not just those of a particular service.

### Tip

- If particular performance counters do not appear in your performance monitoring tool, check the Application Log in Event Viewer for events recorded by Perflib. When Perflib detects disabling errors in a counter DLL, it records an event in the log and sets the value of Disable Performance Counters to 1 to disable the counters.

### Caution

- Do not change the value of this entry by editing the registry. If you do, the performance counters for this service will not operate properly.

**See Also**
**Concepts**
ExtCounterTestLevel
Configuration Flags
Disable Performance Counters Entry
Disable Performance Counters

# Open Timeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*\Performance

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xEE6B2800 *milliseconds* | The value of OpenProcedureWaitTime |

**Description**
Defines a slow Open procedure call for these counters.

This entry tells the Performance Library (Perflib) how long the Open procedure call for these counters is expected to run. If the time specified in the value of this entry expires before the call completes, then Perfib considers the call to be abnormally slow, and it records an error in the Application Log in Event Viewer. Also, if the Configuration Flags entry requires it, then Perflib disables the counters and does not display the counters again until you manually re-enable them.

Perflib times the procedure calls of non-trusted counters to identify counters that can delay or lock performance tools during initialization. While the Collect procedure call proceeds, you cannot use the performance tools, and remote users cannot log on to the local computer.

**Activation Method**
To make changes to this entry effective, restart the Remote Registry Service or restart Windows.

### Note

- Timing procedure calls is one of several tests Perflib performs to maintain the integrity of the counters. To configure the Perflib test feature, use the Configuration Flags entry.
- Perflib does not time the procedure calls of trusted performance counter dynamic-link libraries (DLLs). This entry applies only to non-trusted DLLs.
- Perflib disables counters because of slow procedure calls only when, the 0x8 (1000) bit is set and the 0x2 (10) bit is not set in the value of Configuration Flags.
- Perflib does not time procedure calls when the 0x4 (100) bit is set in the value of Configuration Flags.
- Perflib disables counter by setting the value of Disable Performance Counters for that service to 1. After you have repaired or replaced the counter DLLs, you must re-enable them manually by changing the value of **Disable Performance Counters** to 0.
- Some performance counters DLLs add the **Collect Timeout** entry to the registry. If it does not appear, you can add this entry to the registry to override the default value.
- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Tip

- If you suspect problems with procedure calls, check the Application Log in Event Viewer. Perflib records an error in the log whenever a procedure call exceeds its limit.

**See Also**
**Concepts**
OpenProcedureWaitTime
Collect Timeout
Configuration Flags

# service-name\Security Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*

**Description**
The Security subkey stores information about the permissions set for a driver or a Win32 service.

# Security\Security Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services*service-name*\Security

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *Security descriptor* | (None) |

**Description**
Stores the security descriptor for the service. A security descriptor specifies the security attributes of the service, including the security identifier (SID) and access control lists associated with the service.

Applications view and edit this entry by using the QueryServiceObjectSecurity() and SetServiceObjectSecurity() Win32 APIs.

**Change Method**
To change the value of this entry, you can use the Services snap-in. You can also use the ChangeServiceConfig API or use SC, a tool included in Windows Server 2003.

### Note

- Not all services have this entry. In its absence, the Security Configuration Manager uses a default security descriptor.

# AtmUni

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The AtmUni subkey stores configuration data for the Asynchronous Transfer Mode (ATM) user-to-network interface (UNI) Call Manager. The ATM UNI Call Manager is a Windows Server 2003 component that creates switched and permanent virtual circuits for ATM signaling and communication.

The AtmUni\Parameters subkey stores entries that apply to all ATM network adapters. Entries that configure a particular adapter are stored in the subkey named by the globally unique identifier (GUID) of the adapter in the AtmUni\Parameters\Adapters subkey.

The AtmUni subkey and its subkeys appear in the registry only after you install an ATM network adapter. The system detects the adapter and installs the Windows ATM Services and their registry elements.

By default, the AtmUni subkey contains only the standard entries associated with a Windows service. However, Windows ATM Services adds entries to these subkeys when you use Network and Dial-up Connections to change the default properties of an ATM connection, the ATM Call Manager, or an ATM network adapter. You can also add the entries to configure your ATM adapters.

# AtmUni\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\AtmUni

**Description**
The AtmUni subkey stores configuration data for the asynchronous transfer mode (ATM) user-to-network interface (UNI) Call Manager. The ATM UNI Call Manager is a Windows Server 2003 service that creates switched and permanent virtual circuits for ATM signaling and communication.

The Parameters subkey stores entries that apply to all ATM network adapters.

In addition to these entries, the Parameters subkey contains an Adapters subkey, which contains a subkey for each ATM network adapter installed on the system. The subkeys within AtmUni\Parameters\Adapters are named by the globally unique identifier (GUID) of a particular network adapter, and they store entries that apply only to that adapter.

# LogFileName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\AtmUni\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Fully-qualified path to a text file (.txt)* | (None) |

**Description**
Specifies the location of the ATMUNI event log file. If you do not add this entry to the registry, then the AtmUni Call Manager does not create a log.

By adding this entry to the registry, you direct the AtmUni Call Manager to create an event log file at the specified location. To specify the types of events to record in the log, use the LogFlags entry. By default, no events are recorded in the log.

The AtmUni Call Manager creates a new log file at the specified location each time the service starts, and it records events in that file until the service stops. There is no limit on the log file size. To prevent a particular log file from being overwritten when the service starts again, you must rename the file or change the value of this entry.

This entry does not exist in the regitry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- The ATMUNI log file records events generated by activity on all asynchronous transfer mode (ATM) network adapters. You cannot configure a log file for a particular ATM adapter.

**See Also**
**Concepts**
LogFlags

# LogFlags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\AtmUni\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xF *bitmask value* | 0x0 |

**Description**
Specifies which types of events are recorded in the ATMUNI log file. The ATMUNI log file is created only when the value of this entry is greater than 0x0 and you have added a LogFileName entry to the AtmUni\Parameters subkey.

The value of this entry is a bitmask. To add an event type to the log, set the bit representing that event type to 1. To remove an event type from the log, set the bit representing that event type to 0. You can also determine the necessary value for this entry by summing the hexadecimal values of the event types that you want to appear in the log.

For example, to indicate that all event types should appear in the log, enter 0xF (1 + 2 + 4 + 8).

| Value | Meaning |
|---|---|
| 0x0 (0) | Do not write any events to the log. No log is created. |
| 0x1 (01) | Failed calls. |
| 0x2 (10) | Successful calls. |
| 0x4 (100) | Abnormal call terminations. |
| 0x8 (1000) | Normal call terminations. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- The ATMUNI log records events generated by the activity of all asynchronous transfer mode (ATM) network adapters. You cannot configure a log file for a particular ATM adapter.

**See Also**
**Concepts**
LogFileName

# Parameters\Adapters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\AtmUni\Parameters

**Description**
The Adapters subkey stores asynchronous transfer mode (ATM) configuration data that applies only to a particular adapter. (Entries that apply to all adapters appear in the AtmUni\Parameters subkey.) This subkey enables you to configure individual adapters independently.

The Adapters subkey contains one or more subkeys. Each of these subkeys represents a particular ATM network adapter, and each is named by the globally unique identifier (GUID) of that adapter. Each of the *Adapter-GUID* subkeys stores entries that apply only to the adapter represented by the subkey.

Some entries can appear in the AtmUni\Parameters subkey and in an *Adapter-GUID* subkey. If the entry appears in both locations, then the value in the *Adapter-GUID* subkey takes precedence over the value in the AtmUni\Parameters subkey when configuring that particular adapter.

# Adapter-GUID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\AtmUni\Parameters\Adapters

**Description**
Each *Adapter-GUID* subkey in the Adapters subkey represents a particular asynchronous transfer mode (ATM) network adapter, and each is named by the globally unique identifier (GUID) of the adapter. The actual *Adapter-GUID* subkeys in your registry depend on the ATM adapters that are installed on the computer or that are available to the system.

Some entries can appear in the AtmUni\Parameters subkey and in an *Adapter-GUID* subkey. If the entry appears in both locations, then the value in the *Adapter-GUID* subkey takes precedence over the value in the AtmUni\Parameters subkey when configuring that particular adapter.

### Note

- The *Adapter-GUID* subkey is a variable representing the subkeys in the HKLM\System\CurrentControlSet\Services\AtmUni\Parameters\Adapters subkey. It does not actually appear in the registry. This variable subkey displays the entries and subkeys that are common to many of the subkeys in Adapters.

# Adapter-GUID\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\AtmUni\Parameters\Adapters*Adapter-GUID*

**Description**
The Parameters subkey stores configuration data that applies only to the asynchronous transfer mode (ATM) network adapter represented by the subkey.

Some entries can appear in the AtmUni\Parameters subkey and in an *Adapter-GUID* subkey. If the entry appears in both locations, then the value in the *Adapter-GUID* subkey takes precedence over the value in the AtmUni\Parameters subkey when configuring that particular adapter.

# IlmiRegisterAddresses

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\AtmUni\Parameters\Adapters*Adapter-GUID*\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**

Specifies whether Integrated local management interface (ILMI) is enabled on the adapter.

By default, ILMI address registration is enabled on all ATM network adapters. To disable ILMI on a particular adapter, set the value of this entry to 0 in the *Adapter-GUID* subkey representing the target adapter.

| Value | Meaning |
|---|---|
| 0 | ILMI is disabled. |
| 1 | ILMI is enabled. |

# AuthSrv

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The AuthSrv subkey stores configuration data shared by the Internet Authentication Service (IAS) and Routing and Remote Access.

# AuthSrv\Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\AuthSrv

**Description**
The Parameters subkey stores configuration data shared by the Internet Authentication Service (IAS) and Routing and Remote Access.

# BootVerification

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The BootVerification subkey stores entries for Bootvrfy.exe, a program included in Windows Server 2003 that notifies the system that startup was successful. Bootvrfy.exe can be run on a local or remote computer.

By default, Winlogon determines whether a startup is successful, and it notifies the system of the success or failure. However, you can use Bootvrfy.exe to declare the logon successful instead. Typically, users place Bootvrfy.exe in a service group so that it runs after mission-critical applications have been loaded successfully. This ensures that startup is declared successful only after these applications have been loaded.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Notes

- To run a custom startup verification program, you must disable the standard startup verification functions in Winlogon by setting the value of the ReportBootOk entry to 0.

- You cannot run Bootvrfy.exe if you are also running a custom startup verification program. Custom boot verification programs are listed in the value of the BootVerificationProgram\ImagePath entry in the BootVerificationProgram subkey.

### Tip

- To run Bootvrfy.exe, use Regedit.exe to add the BootVerification subkey to the registry. Then, add the following entries to the subkey with the values specified:

| Entry name | Data type | Value |
|---|---|---|
| ErrorControl | REG_DWORD | 1 |
| service-name\ImagePath | REG_EXPAND_SZ | Bootvrfy.exe |
| ObjectName | REG_SZ | LocalSystem |
| Start Entry | REG_DWORD | 3 |
| Type Entry | REG_DWORD | 2 |

**See Also**
**Concepts**
BootVerificationProgram
BootVerificationProgram\ImagePath
ReportBootOk

# Cdrom

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The Cdrom subkey stores configuration data for CD-ROM drivers and devices.

**Activation Method**
You can activate changes to the values of entries in this subkey by restarting the computer.

You can also activate changes without restarting the computer. From Device Manager, in Computer Management, first disable all CD-ROM drives and then enable all CD-ROM drives.

**Caution**

Invalid configuration of entries in this subkey can prevent the computer from starting.

# Cdrom\AutoRun

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Cdrom

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the CD-ROM driver sends a Media Change Notification (MCN) message to clients, including the Windows interface, when it detects a change in media. The MCN message triggers media-related features, such as autoplay when a CD-ROM is inserted into the CD-ROM drive.

If the MCN message is disabled, then the media features that use it do not operate.

| Value | Meaning |
|---|---|
| 0 | Does not send an MCN message. |
| 1 | Sends an MCN message. |

**Activation Method**
You can activate changes to the value of this entry by restarting the computer. Or, in Computer Management, in Device Manager, disable all CD-ROM drives and then enable all CD-ROM drives.

### Note

- If the value of this entry is 0, then Autoplay does not operate, regardless of the value of the **NoDriveAutoRun** entry (in HKLM and HKCU) and the **NoDriveTypeAutoRun** entry (in HKLM and HKCU).

### Tip

- To disable Autoplay on particular drives, use **NoDriveAutoRun** (in HKLM or HKCU) or **NoDriveTypeAutoRun** (in HKLM or HKCU). Do not disable the MCN message unless you have no alternative.

### Caution

Disabling MCN may cause other system components, such as NT Media Services, to not work properly with the CD-ROMs in the system. Consider disabling only Autoplay, as described in Tips.

**See Also**
**Concepts**
AutoRunAlwaysDisable

# AutoRunAlwaysDisable

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Cdrom

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Vendor or product identfiers* | (List of identified products and product groups.) |

**Description**
Suppresses the Media Change Notification (MCN) message for particular CD-ROM drives.

By default, the system sends an MCN message when a CD-ROM is inserted in a CD-ROM drive. The MCN message triggers media features, such as Autoplay. When the MCN message is suppressed, these media features do not operate.

This entry identifies devices that do not support media features. Typically, this list consists of CD-ROM changers that do not support Autoplay because they must attempt to physically load the media to verify that the media is present.

The entries in this list do not need to be complete device identifiers. The system suppresses the MCN message for all devices that begin with the characters of an entry in the list. This method makes it easier to disable media features for all products of a certain vendor or product line.

### Notes

- If the value of this entry is not blank, then it takes precedence over a value of 1 in Cdrom\AutoRun for the particular devices identified.
- If a device is identified in this list, then Autoplay does not operate on the device, regardless of the value of the **NoDriveAutoRun** entry (in HKLM and HKCU) and the **NoDriveTypeAutoRun** entry (in HKLM and HKCU).
- The identifiers listed in this value are those reported to the system by the device. These identifiers might differ from those listed on the drive housing.
- The default value for this entry consists of products and product groups that have been identified through testing as being unable to support Autoplay. This list might change as testing progresses and products are updated.

### Tip

- To disable Autoplay on particular drives, use the **NoDriveAutoRun** entry (HKLM or HKCU) or the **NoDriveTypeAutoRun** entry (HKLM or HKCU). Do not disable the MCN message unless you have no alternative.

### Caution

- If you add or remove items from the list stored in this entry, then your CD-ROM drives might not operate properly.

**See Also**
**Concepts**
Cdrom\AutoRun

# DHCP Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Dhcp\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Dhcp

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Options

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Dhcp\Parameters

**Description**
The Options subkey stores configuration data for options that the client can request from the Dynamic Host Configuration Protocol (DHCP) service by using DHCP Manager.

The Options subkey does not contain entries. Instead, it contains numbered subkeys, each of which represents a DHCP option. The entries in each Option# subkey apply only to the option represented by the subkey.

**Caution**

- Do not directly modify the values of entries in the subkeys of this subkey by, for example, using the Registry Editor (regedit.exe). Values in this subkey should only by modified by using the standard interface.

# Option#

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Dhcp\Parameters\Options

**Description**
The Parameters subkey contains *Option#* subkeys, which store configuration data for options that the client can request from the Dynamic Host Configuration Protocol (DHCP) service by using DHCP Manager.

Most of the entries described in this reference as being stored in the *Option#* subkey can appear in any of the subkeys in the Options subkey.

**Note**

- "*Option#*" subkey is a placeholder representing the subkeys in the subkey DHCP\Parameters\Options "Option#" does not actually appear in the registry.

**Caution**

- Do not edit the value of these types of entries by using the Registry Editor (regedit.exe). Values in this subkey should only by modified by using the standard interface.

# KeyType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Dhcp\Parameters\Options*Option#*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–7 | 7 |

**Description**
Specifies the data type of the registry entry that stores configuration data for the option that the client can request from the Dynamic Host Configuration Protocol (DHCP) service. The data type defines the length and format of data that the entry can store.

| Value | Meaning |
|---|---|
| 1 | REG_SZ |
| 2 | REG_EXPAND_SZ |
| 3 | REG_BINARY |
| 4 | REG_DWORD or REG_DWORD_LITTLE_ENDIAN |
| 5 | REG_DWORD_BIG_ENDIAN |
| 6 | REG_LINK |
| 7 | REG_MULTI_SZ |

**Note**

- The name of the entry that stores configuration data appears in the value of the RegLocation entry.

**Tip**

- For more information about data types, see Data Types in the Registry.

**See Also**
**Concepts**
RegLocation

# RegLocation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Dhcp\Parameters\Options*Option#*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Registry path* | (Varies with the option.) |

**Description**
Specifies the path to the registry entry that stores configuration data for the option that the client can request from the Dynamic Host Configuration Protocol (DHCP) service. This data is obtained from the Dynamic Host Configuration Protocol (DHCP) service.

The path to the entry can include a variable (*?*) that represents the Services subkey for the network adapter. For example, if the value of **RegLocation** is System\CurrentControlSet\Services\?\Parameters\Tcpip\DhcpSubnetMaskOpt, it indicates that data for the option is stored in the value of the **DhcpSubnetMaskOpt** entry in HKLM\System\CurrentControlSet\Control\Services\AdapterName#\Parameters\Tcpip.

**Tip**

- The data type of the registry entry that stores configuration data appears in the value of the KeyType entry.

**See Also**
**Concepts**
KeyType

# DHCPServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The DHCPServer subkey stores configuration data for the Dynamic Host Configuration Protocol (DHCP) service. A DHCP server is a server that is running a Windows Server 2003 operating system, Microsoft TCP/IP, and the DHCP-compatible server software.

DHCP servers maintain a database of TCP/IP configuration information and provide that information to clients upon request.

# DHCPServer\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DHCPServer

**Description**
The Parameters subkey stores configuration options for the Dynamic Host Configuration Protocol (DHCP) service database on this server. DHCP servers maintain a database of TCP/IP configuration information and provide that information to clients upon request.

# ActivityLogFlag

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether Dynamic Host Configuration Protocol (DHCP) activity is audited and recorded in a log.

| Value | Meaning |
|---|---|
| 0 | DHCP activity is not audited or logged. |
| 1 | DHCP activity is audited and logged. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**

To change the value of this entry, use the DHCP Manager snap-in. Right-click the name of the server, click **Properties**, and then click the **General** tab. This entry stores the setting of the **Enable Audit Log** check box.

# APIProtocolSupport

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1 \| 0x2 \| 0x4 \| 0x5 \| 0x7 | 0x5 (5) |

**Description**

Specifies the supported protocols for the Dynamic Host Configuration Protocol (DHCP) server, and ensures that different computers running different protocols can access the DHCP server.

To enable support for a protocol, set the corresponding bit to 1, or add the hexadecimal value to the value of APIProtocolSupport. For example, to enable support for all protocols, set the value of this entry to 0x7 (111 binary).

| Binary value | Hexadecimal value | Meaning |
|---|---|---|
| 1 | 0x1 | Supports TCP/IP protocol |
| 10 | 0x2 | Supports named pipes protocols |
| 100 | 0x4 | Supports local procedure call (LPC) protocols |

# BackupDatabasePath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | *Systemroot*\System32\dhcp\backup |

**Description**
Specifies the directory where the backup copy of the Dynamic Host Configuration Protocol (DHCP) client information database is stored. If you change the value of this entry, then DHCP saves new backup copies in the new location, but it does not move existing backup copies of the database.

**Change Method**
To change the value of this entry, use the DHCPmgmt.msc snap-in or the netsh dhcp server commands.

**Activation Method**
To make changes to the value of this entry effective, you must restart DHCP server using the DHCPmgmt.msc snap-in, netsh, or services.msc

### Tip

- Store the backup file on a hard drive other than the one on which the database resides. This lets you recover the database even when the system drive fails. However, because DHCP Manager cannot access a network drive for database backup and recovery, you cannot store the backup file on a network drive.

**See Also**
**Concepts**
DatabasePath
BackupInterval

# BackupInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–71582 *minutes* | 60 |

**Description**
Specifies how often the Dynamic Host Configuration Protocol (DHCP) client information database is backed up.

### Note

- Because DHCP stores this value internally in milliseconds, this entry is limited to values less than 71583. When converted to milliseconds, higher values are too large to store in a DWORD field.

**See Also**
**Concepts**
BackupDatabasePath

# DatabaseCleanupInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *minutes* | 0x5A0 |

**Description**

Specifies how often the Dynamic Host Configuration Protocol (DHCP) Service deletes expired (*doomed*) records from the DHCP client information database. By default, records expire after four hours. Deleting expired records releases space for new IP addresses.

# DatabaseName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Filename* | dhcp.mdb |

**Description**
Specifies the name of the file that contains the Dynamic Host Configuration Protocol (DHCP) client information database.

**Change Method**
To change the value of this entry, use DHCP Manager.

**See Also**
**Concepts**
DatabasePath

# DatabasePath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | *Systemroot*\System32\dhcp |

**Description**
Specifies the directory where the Dynamic Host Configuration Protocol (DHCP) client information database files and audit log files are stored. If you change the value of this entry, DHCP stores new database files in this location, but it does not move existing database files to the new location.

By default, database files and audit log files are stored in the same directory. The location of log files is determined by the value of the DhcpLogFilePath entry.

**Change Method**
To change the value of this entry, use the DHCP Manager snap-in. Right-click the name of the server, click **Properties**, click the **Advanced** tab, and enter a new path in the **Database path** field.

**See Also**
**Concepts**
DatabaseName
DhcpLogFilePath

# DetectConflictRetries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–5 *retries* | 0 |

**Description**
Specifies how many times the Dynamic Host Configuration Protocol (DHCP) service issues Internet Control Message Protocol (ICMP) Ping messages if the messages are not successful. When the number of unsuccessful messages reaches the value of this entry, the DHCP service offers a new IP address.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the DHCP Manager snap-in. Right-click the name of the server, click **Properties**, and then click the **Advanced** tab. This entry stores the value in the **Conflict detection attempts** box.

### Note

- The maximum valid value for this entry is 5. If you enter a value greater than 5, then the DHCP service interprets it as 5.

# DhcpLogDiskSpaceCheckInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *number of audit log entries* | 0x32 (50) |

**Description**
Specifies how often Dynamic Host Configuration Protocol (DHCP) verifies that disk space and file size are adequate for its audit log. The value of this entry represents the number of audit log records entered between each verification check.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
DhcpLogMinSpaceOnDisk

# DhcpLogFilePath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Path name* | Same as DatabasePath |

**Description**
Specifies the directory where Dynamic Host Configuration Protocol (DHCP) audit logs are stored. If you change the value of this entry, then DHCP moves the audit log files to the new location.

By default, the audit logs are stored in the same directory as the DHCP client information database. The location of the database is specified in the value of the DatabasePath entry.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the DHCP Manager snap-in. Right-click the name of the server, click **Properties**, click the **Advanced** tab, and enter a new path in the **Log file path** box.

**See Also**
**Concepts**
DatabasePath

# DhcpLogFilesMaxSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *MB* | 0x7 (7 MB) |

**Description**
Specifies the maximum combined size of one week's worth of Dynamic Host Configuration Protocol (DHCP) audit logs. If the audit logs exceed the size specified by this value, then DHCP stops writing to the audit log until sufficient space becomes available.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# DhcpLogMinSpaceOnDisk

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *MB* | 0x14 (20 MB) |

**Description**
Specifies the minimum amount of free disk space required for audit logging.

The Dynamic Host Configuration Protocol (DHCP) periodically verifies that it has sufficient disk space to proceed with audit logging (the interval is specified in the value of the DhcpLogDiskSpaceCheckInterval entry). This entry specifies the minimum value that DHCP requires. If the amount of free disk space is less than the value of this entry, then the DHCP service does not write to the audit log. Logging is suspended until sufficient disk space is available.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
DhcpLogDiskSpaceCheckInterval

# IgnoreBroadcastFlag

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the Dynamic Host Configuration Protocol (DHCP) server complies with or ignores the value of the broadcast bit (broadcast flag) in a client DHCP Discover message.

By default, Windows Server 2003 DHCP servers ignore the broadcast bit on the client requests. Instead, they send all DHCP responses as IP broadcasts to the limited broadcast address (255.255.255.255). However, a DHCP server can be configured to send responses as unicast responses if the client request specifies that a broadcast response is not required.

On Microsoft DHCP clients, the broadcast bit is not set, that is, it is set to request initial Dhcpoffer and Dhcpack packets in unicast messages.

| Value | Meaning |
|---|---|
| 0 | Comply with broadcast flag in client request. |
| 1 | Ignore broadcast flag in client requests, and broadcast all DHCP packets |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must restart the system before a change to this value takes effect.

**Caution**

- Setting the value of this entry to 0 reduces broadcast traffic. However, clients cannot receive unicast Dhcpoffer and Dhcpack packets unless the Microsoft DHCP server and DHCP client are on a homogeneous Ethernet network or the same subnet of a Token Ring network. Do not change the default broadcast behavior of a Microsoft DHCP server if the server and client are on different Token Ring subnets or are separated by a router or bridge that does media access control (MAC) level address translation, for example, Ethernet to or from Token Ring.

# MaxActiveThreads

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–20 | 1 |

**Description**
Used to calculate the maximum number of Dynamic Host Configuration Protocol (DHCP) message processing threads that can be active at any given time. The DHCP service will activate no more threads than the value of this entry, plus one thread for each processor on the server computer.

The value of this entry must be no more than the value of the MaxProcessingThreads entry. If you enter a higher value, then the system uses the value of **MaxProcessingThreads**.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
MaxProcessingThreads

# MaxProcessingThreads

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x11 |

**Description**
Used to calculate the maximum number of message processing threads that the Dynamic Host Configuration Protocol (DHCP) service can create. The DHCP service creates no more threads than the value of this entry, plus one thread for each processor on the server computer. Although you can set the value of this entry higher than 20, DHCP creates no more than 20 threads.

This value must be greater than or equal to the value of the MaxActiveThreads entry.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
MaxActiveThreads

# DNS Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The DNS subkey stores configuration information for the Domain Name System (DNS) server.

DNS associates IP addresses and other data identifying a computer with the computer's display name, for example host.reskit.com.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console. If you change entries by editing the registry, the changes are not effective until you restart the DNS server.

**Tip**

- This subkey stores data for the DNS server. To configure DNS dynamic update on the DNS client, see the following entries in the Tcpip subkey: **DefaultRegistrationTTLEnableAdapterDomainNameRegistrationDisableDynamicUpdateDisableReplaceAddressesInConflicts**DisableReverseAddressRegistrations**UpdateSecurityLevel**

# DNS\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS

**Description**
The Parameters subkey stores Domain Name System (DNS) configuration data that applies to all zones in the network. Settings that apply to a specific zone are stored in a subkey representing that zone in the Zones subkey.

# AddressAnswerLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 5–28 *records* | 0 |

**Description**

Specifies the maximum number of A (host IP address) resource records that the Domain Name System (DNS) server can insert in the answer section of a response to an A record query (a query for an IP address). The value of this entry also influences the setting of the truncation bit. If the value of this entry is from 5 through 28, then the truncation bit is not set on the response, even when the packet space is exceeded.

The value of **AddressAnswerLimit** is imposed upon A record queries only; it does not affect other types of queries.

| Value | Meaning |
|---|---|
| 0 | There is no limit on the number of A records in a response. If the response, including all A records, does not fit in a UDP DNS packet, then the truncation bit is set. |
| 5 - 28 | The maximum permissible number of A records in a response. The truncation bit is not set, regardless of the size of the response. |

**Change Method**

To change the value of this entry, use Dnscmd.exe, a tool included in Windows Support Tools. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**

DNS reads its registry entries only when it starts. You can modify entries while the DNS server is running by using the DNS console. If you modify entries by editing the registry, the changes are not effective until you restart the DNS server.

### Note

- This entry was designed to resolve a problem in DNS for Windows 95 (without service packs) in which DNS fails when it receives more than 28 A records in a response to an A query. This occurs infrequently and only when numerous servers are supporting a Web site. However, because the truncation bit is not set, limiting A records has the added benefit of preventing remote DNS servers from using TCP to retry these queries. DNS servers often use TCP to retry queries when they receive a response with the truncation bit set.

- DNS does not add this entry to the registry. You can add it by using Dnscmd.exe, or by using the registry editor Regedit.exe.

# AutoCacheUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the DNS server attempts to update its cache file, Cache.dns, at startup. The cache file contains a list of NS records and A records for root servers. The cache file is updated with information that it obtains by querying root DNS servers for root NS type records .

| Value | Meaning |
|---|---|
| 0 | The DNS server does not update the cache file automatically. An administrator must update the cache file by using the DNS console or by editing the Cache.dns file manually. If you edit Cache.dns, you must restart the DNS server to make changes effective. |
| 1 | The DNS server attempts to update the cache file automatically with information that it obtains by querying root DNS servers. It queries the root servers the first time it does a recursive query. |

**Change Method**
To change the value of this entry, use Dnscmd.exe, a tool included in *Windows Server 2003 Support Tools*. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

### Notes

- This entry does not affect NS and A records that are added or changed by using the DNS snap-in or the New Server Wizard. Records added or changed in this manner are always written to the cache file.
- Some Internet root servers do not maintain an accurate set of root server data. As a result, even when updated, the cache file might not be complete.

### Tip

- To force an immediate update, in the DNS snap-in, right-click the name of the DNS server, and then click **Update Server Data Files**.

# AutoConfigFileZones

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 \| 3 | 1 |

**Description**
Specifies whether the Domain Name System (DNS) server updates the resource records of standard primary zones when the fully qualified domain name (that is, the name or the primary DNS suffix) of the local computer changes.

This entry establishes criteria for zone updates. If a zone meets the criteria, and the fully qualified domain name of the local computer changes, then the DNS server updates the host name on resource records that include the local computer name, such as SOA (Start of Authority), NS (name server), and A (address) records. Otherwise, the records are not updated when the computer name changes.

| Value | Meaning |
|---|---|
| 0 | No zones. The DNS server does not update resource records when the local computer name changes. |
| 1 | Dynamic update zones. The DNS server updates only the resource records of standard primary zones that permit dynamic updates. |
| 2 | No dynamic update zones. The DNS server updates only the resource records of standard primary zones that do not permit dynamic updates. |
| 3 | All zones. The DNS server updates the resource records of all standard primary zones. |

**Change Method**
To change the value of this entry, use Dnscmd.exe, a tool included in *Windows Server 2003 Support Tools*. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. You can modify entries while the DNS server is running by using the DNS console. If you modify entries by editing the registry, then the changes are not effective until you restart the DNS server.

**Tip**

- To find other registry entries related to DNS dynamic update, type "dynamic update" on the Search tab in this file.

**See Also**
**Concepts**
PreviousLocalHostname

# BindSecondaries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Enables the Domain Name System (DNS) server to communicate with non-Microsoft DNS servers that use an earlier, slower version of the DNS BIND service.

| Value | Meaning |
|---|---|
| 0 | The Microsoft DNS server places as many resource records as possible in each message to achieve maximum compression and speed of transfer. Microsoft DNS servers and non-Microsoft DNS servers that use BIND version 4.9.4 or higher can receive transfers in this format. |
| 1 | The Microsoft DNS server sends zone transfers to non-Microsoft DNS secondaries with one resource record in each message. Non-Microsoft DNS servers using BIND versions prior to 4.9.4 can receive the transfer only in this format. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the DNS console. Right-click the server name, click **Properties**, and then click the **Advanced** tab. This entry corresponds to the **Bind Secondaries** option in the Advanced server options list.

**Activation Method**
DNS reads its registry entries only when it starts. You can modify entries while the DNS server is running by using the DNS console. If you modify entries by editing the registry, then the changes are not effective until you restart the DNS server.

## Notes

- Transfers between Microsoft DNS servers always use the faster, high-compression method, regardless of the value of this entry.

- The default value of this entry is appropriate for most DNS servers. Change the value of this entry to 0 only if you have NEW BIND servers (or non-BIND,non-Windows servers) that are secondaries to a Microsoft DNS server and if transfer performance is a high priority.

# DatabaseDirectory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Directory path* | *Systemroot*\System32\Dns |

**Description**
Specifies the directory that contains the Domain Name System (DNS) database.

By default, the DNS database is located in *Systemroot*\System32\Dns, but you can add this entry to the registry to change the default location. If you do, then the DNS server loads from and writes new zone records to the new location. If you already have a DNS database in *Systemroot*\System32\Dns, then you must manually move it to the new location. DNS will not move or maintain the original database.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use Dnscmd.exe, a tool included in Windows Support Tools. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. To make changes to this entry effective, you must restart DNS.

# DisableAutoReverseZones

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Enables and disables automatic reverse zones, an optimizing feature of the Domain Name System (DNS). When this feature is enabled, DNS automatically creates three reverse lookup zones.

Reverse lookup zones enable the DNS server to be authoritative, that is, to know the answer in advance and to respond immediately to the most common name queries, thus eliminating unnecessary recursive queries. In accordance with pertinent Requests for Comments (RFCs), by default, the DNS server is authoritative for three reverse lookup zones:

- 0.in-addr.arpa (0.0.0.0)
- 127.in-addr.arpa (127.0.0.1 - loopback)
- 255.in-addr.arpa (255. 255. 255. 255 - broadcast)

| Value | Meaning |
|---|---|
| 0 | The automatic reverse zones feature is enabled. The DNS server automatically creates three reverse lookup zones. |
| 1 | The automatic reverse zones feature is disabled. The DNS server does not automatically create reverse lookup zones. |

**Change Method**
If Microsoft Product Support Services instructs you to change the value of this entry, use Dnscmd.exe, a tool included in the Windows Support Tools, and then restart the DNS server.

**Activation Method**
Because the DNS server creates zones only when it starts, you must restart the DNS server to make changes to this entry effective, even if you use Dnscmd.exe to make the change.

### Note

- DNS does not add this entry to the registry. You can add it by using the registry editor Regedit.exe, or by using a program that stores changes in the registry, such as Dnscmd.exe.

### Caution

- Do not change the value of this entry unless instructed to by Microsoft Product Support Services. Using automatic reverse zones improves the performance of the DNS server. The default value is optimal for almost all DNS server configurations.

# DoNotRoundRobinTypes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Contains a list of types of DNS resource records to exclude from round-robin name resolution. If the value of the **RoundRobin** entry is 1 (which configures the DNS service to use round-robin name resolution), then the types of resource records listed in this registry entry are excluded from round-robin name resolution.

Valid values for this registry entry are any standard DNS record type, including SOA, NS, A, PTR, CNAME, MX, and SRV. In addition, three types of resource records that are not part of standard DNS are also supported: WINS, WINS-R, and ATMA. For more information about resource records, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

### Notes

- To exlude all types of DNS records from round-robin name resolution, set the value of **RoundRobin** to zero.
- The default configuration of Windows Server 2003 DNS is for round-robin name resolution. This entry adds flexibility to Windows Server 2003 DNS, by allowing you to choose partial round-robin name resolution without requiring you to turn it off completely.

**See Also**
**Concepts**
LocalNetPriority

# Parameters\EventLogLevel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 \| 4 | 4 |

**Description**
Specifies which events Domain Name System (DNS) records in the DNS server log in Event Viewer.

| Value | Meaning |
|---|---|
| 0 | None. |
| 1 | Errors only. |
| 2 | Errors and warnings. |
| 4 | All events: errors, warnings, and informational events. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use Dnscmd.exe, a tool included in the Windows Support Tools. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using Dnscmd.exe. If you change entries by using the registry editor Regedit.exe, the changes are not effective until you restart the DNS server.

**See Also**
**Concepts**
Services\EventLog

# ForwardDelegations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the Domain Name System (DNS) server forwards queries about delegated subzones to servers outside of its authoritative zone. This entry is used only if the delegated subzone is within the DNS server's authoritative zone, even though the primary server for the subzone is in another zone.

By default, when the DNS server receives a query for a name outside its authoritative zone, it forwards the query to a similarly-named server in another zone. However, when it receives a query for a delegated subzone, and the subzone is within the DNS server's zone (even when the primary server for the site is in another zone), it sends the query directly to the subzone and does not forward it. When the value of this entry is 1, the DNS server forwards queries for delegated subzones to other servers, just as it does for names in other zones.

Consider adding this entry and setting the value to 1 if the delegation is at a remote site that a remote server can reach more easily.

| Value | Meaning |
|---|---|
| 0 | The DNS server sends queries about delegated subzones that are within its authoritative zone directly to the subzone. It does not forward them. |
| 1 | The DNS server forwards all name queries about delegated subzones to servers in other zones. |

**Change Method**
To change the value of this entry, use Dnscmd.exe, a tool included in the Windows Support Tools. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using Dnscmd.exe. If you change entries by using the registry editor Regedit.exe, then the changes are not effective until you restart the DNS server.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe, or by using a program that edits the registry, such as Dnscmd.exe.
- This entry is used only when forwarding is enabled. If the value of the **Forwarders** entry does not contain at least one valid IP address, then this entry is ignored.

**See Also**
**Concepts**
IsSlave
NoRecursion

# IsSlave

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies how the Domain Name System (DNS) server responds when it does not receive a response to a query that it has forwarded.

| Value | Meaning |
|---|---|
| 0 | If the forwarder servers do not respond, the DNS server issues standard iterative queries to try to resolve the remote name. |
| 1 | If the forwarder servers do not respond, the DNS server terminates the search and sends a SERVER_FAILURE response to the query. |

**Change Method**
To change the value of this entry, select or clear the **Do not use recursion** check box on the **Forwarders** tab in the DNS console. You can also use Dnscmd.exe, a tool included with the Windows Support Tools. Your changes are effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console or Dnscmd.exe. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

### Note

- If you start DNS by using a standard BIND file, the value of Slave in the BIND file takes precedence over the value of this registry entry. This entry might be deleted, or the value in the BIND file might replace the value of this entry.
- This entry is used only when forwarding is enabled, that is, when the value of the **Forwarders** entry includes at least one valid IP address.
- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Tip

- In general, a single forwarder is more efficient than multiple forwarders, because query results are concentrated in one forwarder's cache. To indicate that a forwarder should be queried more than once if an initial query is unsuccessful, list the IP address of the forwarder more than once on the **Forwarders** tab.

### Caution

- Do not change the value of this entry by editing the registry. If the **Forwarders** entry contains valid IP addresses and this entry is deleted, then the DNS server might not start or might not operate properly.

# ListenAddresses

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | *IP address[, IP address] ... (dotted decimal form)* | (None) |

**Description**
Lists the IP addresses bound to the Domain Name System (DNS) server. If this entry does not appear in the registry or its value contains no addresses, the DNS server attempts to bind to all IP addresses on the computer.

**Change Method**
To change the value of this entry, use the **Interfaces** dialog in the DNS console or Dnscmd.exe, a tool included with the Windows Support Tools. Your changes are effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

### Note

- This feature does not usually improve the performance of DNS servers that are running Windows Server 2003.
- The default value is optimal for most servers. However, you might need to exclude certain interfaces, especially if you are running Windows NT 4.0 prior to Service Pack 2, which cannot bind to more than 15 IP addresses. Windows Server 2003 is not limited. Also, in Windows NT 4.0 and earlier, the DNS server does not detect and bind to more than 35 interfaces properly.
- The DNS server must bind to individual IP addresses because many resolvers require that the DNS response comes from the IP address to which the query was sent. The responses may not appear to come from the same IP address for multi-IP servers, depending on TCP/IP stack's routing table.
- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# LocalNetPriority

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies the order in which the Domain Name System (DNS) server returns A (address) records when it has multiple A records for the same name.

| Value | Meaning |
|---|---|
| 0 | The DNS server returns the A records in the order in which they were added to the database. It does not sort the records. If the value of the **RoundRobin** entry is 1, then the server rotates among the records it returns. |
| 1 | The DNS server returns records in order of their similarity to the IP address of the querying client. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the DNS console. Right-click the server name, click **Properties**, and then click the **Advanced** tab. This entry corresponds to the **Enable Netmask Ordering** option. You can also use Dnscmd.exe, a tool included in Windows Support Tools.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console or Dnscmd.exe. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

### Note

- The order in which A records are returned depends on the value of this entry and of the **RoundRobin** entry, as follows:
- If both entries are set to 1 (or neither entry appears in the registry), then the DNS server rotates among the A records it returns in local net priority order, that is, in order of their similarity to the IP address of the querying client.
- If the value of **RoundRobin** is 0 and the value of **LocalNetPriority** is 1, then the DNS server returns the records in local net priority order. It does not rotate among available addresses.
- If the value of **RoundRobin** is 1 and the value of **LocalNetPriority** is 0, then the DNS server rotates among the available records in the order in which the records were added to the database.
- If the values of **RoundRobin** and **LocalNetPriority** are both 0, then the DNS server returns the records in the order in which they were added to the database. It does not attempt to sort them or to rotate among them.

# LogLevel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1 \| 0x10 \| 0x20 \| 0xFE\| 0x100 \| 0x200 \| 0x1000 \| 0x2000 \| 0x4000 \| 0x8000 \| 0xFFFF \| 0x10000 \| 0x20000 \| 0x1000000 \| 0x80000000 | 0x0 |

**Description**

Specifies which events the Domain Name System (DNS) server records in its transaction log, Dns.log.

This entry is a 32-bit binary value. To direct DNS to log a type of event, the DNS console sets the value of this entry to that event type, as listed below. To direct DNS to log more than one type of event, the DNS console sets the value to the sum of the values representing each event type you want to include in the log.

To get a complete log, set the value to 0x8100F331.

| Value | Meaning |
|---|---|
| 0x0 | None. The DNS Service does not create Dns.log. |
| 0x1 | Queries. |
| 0x10 | Notifications. |
| 0x20 | Updates. |
| 0xFE | Non-query transactions. |
| 0x100 | Questions. |
| 0x200 | Answers. |
| 0x1000 | Send packets. |
| 0x2000 | Receive packets. |
| 0x4000 | UDP (User Datagram Protocol) packets. |
| 0x8000 | TCP packets. |
| 0xFFFF | All packets. |
| 0x10000 | Active Directory write transaction. |
| 0x20000 | Active Directory update transaction. |
| 0x1000000 | Full packets. |
| 0x80000000 | Write-through transactions. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**

To change the value of this entry, use the DNS console. Right-click the server name, click **Properties**, and then click the **Logging** tab. This bitmask corresponds to the check boxes in the **Debug logging options** section.

**Activation Method**

DNS reads its registry entries only when it starts. You can modify entries while the DNS server is running by using the DNS console. If you change entries by editing the registry, the changes are not effective until you restart the DNS server.

**See Also**
**Concepts**
Parameters\EventLogLevel

# MaxCacheTtl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFF *seconds* | 0x15180 (86,400 seconds = 1 day) |

**Description**
Specifies how long the Domain Name System (DNS) server can save a record of a recursive name query.

If the value of this entry is 0x0, then the DNS server does not save any records.

The DNS server saves the records of recursive name queries in a memory cache so that it can respond quickly to new queries for the same name. Records are deleted from the cache periodically to keep the cache content current. The interval during which the records remain in the cache usually is determined by the value of the Time to Live (TTL) field in the record. This entry establishes the maximum time that records can remain in the cache. The DNS server deletes records from the cache when the value of this entry expires, even if the value of the TTL field in the record is greater.

**Change Method**
To change the value of this entry, use Dnscmd.exe, a tool included in Windows Support Tools. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. If you change the value of this entry by editing the registry, then the changes are not effective until you restart the DNS server.

### Notes

- This entry does not affect Windows Internet Name Service (WINS) data saved in the DNS memory cache. WINS data is saved until the Cache Timeout Value on the WINS record expires. To view or change the Cache Timeout Value on the WINS record, use the DNS snap-in. Right-click a zone name, click **Properties**, click the **WINS** tab, and then click the **Advanced** button.

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# NoRecursion

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the Domain Name System (DNS) server performs recursive resolution in accordance with the RecursionDesired bit in the DNS name query packet header. Recursion occurs only when the value of this entry is 0 and the RecursionDesired bit is set.

| Value | Meaning |
|---|---|
| 0 | The DNS server acts in accordance with the RecursionDesired bit set in the DNS query packet header. |
| 1 | The DNS server does not perform recursive resolution regardless of the setting of the RecursionDesired bit in the DNS query packet header. |

You might consider changing the value of this entry if clients are limited to the names on a server, such as within an intranet, or if the server cannot resolve external names.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the DNS console. Right-click the server name, click **Properties**, and then click the **Advanced** tab. This entry corresponds to the **Disable Recursion** option. Or use Dnscmd.exe, a tool included in Windows Support Tools. Using either method, the changes are effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console or Dnscmd.exe. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

### Note

- If you start DNS by using a standard BIND file, then the value of the RecursionDesired bit in the BIND file takes precedence over the value of this registry entry. The entry might be deleted or its value replaced by the value in the BIND file.

# PreviousLocalHostname

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Fully qualified domain name* | (Fully qualified domain name of the local computer) |

**Description**
Stores the fully qualified domain name of the local computer. When the computer name or the primary Domain Name System (DNS) suffix changes, the DNS server uses the previous name to update resource records on primary, Active Directory–integrated zones.

The DNS server copies the name of the local computer (stored in the **Hostname** entry) and the primary DNS suffix (stored in the **PrimaryDnsSuffix** entry) to the value of this entry. If the computer name or primary DNS suffix changes, then the DNS server searches SOA (Start of Authority), NS (name server), and A (address) resource records for the value of this entry, and it replaces each instance with the new computer name. Then, it copies the new value of **Hostname** and **PrimaryDnsSuffix** to the value of this entry for future changes.

### Note

- The DNS server updates resource records according to the standards set for the zone and stored in the AutoConfigFileZones entry.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
AutoConfigFileZones

# PublishAutonet

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the Domain Name System (DNS) server builds A (address) records for temporary IP addresses produced by Automatic Private IP Addressing (APIPA). This entry applies only when the DNS server is building A records for the computer on which it is running.

APIPA is designed for computers on single-subnet networks that do not include a Dynamic Host Configuration Protocol (DHCP) server. APIPA automatically assigns an IP address from its reserved range, 169.254.0.0 through 169.254.255.255. These APIPA IP addresses, commonly known as *autonet* addresses, are used only until the system finds a DHCP server. APIPA addresses cannot be used on the Internet.

| Value | Meaning |
|---|---|
| 0 | Exclude APIPA addresses. |
| 1 | Include APIPA addresses. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use Dnscmd.exe, a tool included in Windows Support Tools. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. If you change entries by using the registry, then the changes are not effective until you restart the DNS server.

# RecursionRetry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF *seconds* | 0x3 |

**Description**

Specifies how often Domain Name System (DNS) repeats recursive client queries when it does not receive a response from a remote server. If the DNS server does not receive a response before the time specified by the value of this entry expires, then the DNS server repeats the query to the same server or to other DNS servers.

The default value is appropriate for most servers. However, if this value is less than the time that a remote server requires to respond over a slow link, then increase this value so that it is slightly longer than the observed response time.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**

To change the value of this entry, use Dnscmd.exe, a tool included in Windows Support Tools.

**Activation Method**

Changes made by using Dnscmd.exe are effective immediately; you do not have to restart the DNS server. If you change the value of this entry by editing the registry, then the changes are not effective until you restart the DNS server.

# RpcProtocol

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x7 \| 0xFFFFFFFF | 0xFFFFFFFF |

**Description**
Specifies the protocols that administrative remote procedure calls (RPCs) use. Although these flags are not specific to Domain Name System (DNS), the DNS server establishes endpoints to allow connections using these protocols.

The value of this entry is a bitmap. You can set multiple bits by adding the bits together and setting the value of this entry to that sum:

| Bit | Meaning |
|---|---|
| 0x0 | No protocols (disables RPC for DNS) |
| 0x1 (001 binary) | TCP/IP |
| 0x2 (010 binary) | Named pipes |
| 0x4 (100 binary) | LPC |
| 0xFFFFFFFF | All protocols |

**Change Method**
To change the value of this entry, use the DNS console.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

# SecureResponses

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether the Domain Name System (DNS) server tries to eliminate illegitimate records by filtering the records that it saves in its memory cache.

The DNS server saves the records of recursive name queries in a memory cache so that it can respond quickly to new queries for the same name. By default, it saves all records. However, if the value of this entry is 1, then DNS saves only those query records for names that are in the same subtree as the server that provided them. For example, the DNS server would save an NS (name server) record for ns.reskit.com from the reskit.com server, but it would not save the NS record for ns.contoso.com from the reskit.com server. This filtering is designed to minimize the effect of malicious attacks on an Internet server, but it might generate additional network traffic.

| Value | Meaning |
|---|---|
| 0 | The DNS server saves all name query records in its memory cache. It does not attempt to filter out illegitimate records. |
| 1 | The DNS server saves only those records of names that are in the same subtree as the name in the original query. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**

To change the value of this entry, use the DNS console. Right-click the name of a DNS server, and then click the **Advanced** tab. This entry stores the setting of the **Secure cache against pollution** check box.

**Activation Method**

DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

# StrictFileParsing

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies how the Domain Name System (DNS) server responds when it receives zone files whose records contain errors that violate standards specified by Requests for Comments (RFCs). These include records for names outside of the zone, and CNAME (canonical name) records that contain other records.

| Value | Meaning |
|---|---|
| 0 | When the DNS server encounters an erroneous record, it writes an error to the DNS log in Event Viewer and continues to load. |
| 1 | When the DNS server encounters erroneous records, it writes an error to the DNS log in Event Viewer and stops without loading. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the DNS console. Right-click the server name, click **Properties**, and then click the **Advanced** tab. This entry corresponds to the **Fail On Load If Bad Zone Data** option. Or, use Dnscmd.exe, a tool included in Windows Support Tools.

**Activation Method**
Changes made by using either the DNS console or Dnscmd.exe are effective immediately; you do not have to restart the DNS server. If you change the value of this entry by editing the registry, then the changes are not effective until you restart the DNS server.

### Note

- The default behavior of DNS changed in Windows NT 4.0 with Service Pack 4. In Windows NT 4.0 (previous to Service Pack 4) and earlier, the DNS server does not start if it encounters incorrect zone records. Check the System Log in Event Viewer for errors.

# UpdateOptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x80000000 | 0x30F |

**Description**

Prohibits Domain Name System (DNS) dynamic update of certain types of records.

This entry is a bitmask. To disable DNS dynamic update on a record type, set the bit for that record type to 1, or sum the hexadecimal values of the record types.

For non-secure zones:

| Value | Meaning |
|---|---|
| 0x1 | SOA (Start of Authority) records. |
| 0x2 | NS (name server) records. |
| 0x4 | Delegation NS records. |
| 0x8 | A (address) records in the DNS server's own host record. |

For secure zones:

| Value | Meaning |
|---|---|
| 0x100 | SOA records. |
| 0x200 | NS records. |
| 0x400 | Delegation NS records. |
| 0x800 | A (address) records in the DNS server's own host record. |

Other possible values:

| Value | Meaning |
|---|---|
| 0x0 | Allows DNS dynamic update on all record types. |
| 0x30F | Prevents all of the update policy options for unsecure zones, and allows only SOA and zone root NS updates for secure zones. |
| 0x1000000 | Prevents the relay to a DNS server's peers of the server's A record updates. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**

To change the value of this entry, use Dnscmd.exe, a tool included in Windows Support Tools. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**

DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

- To find other registry entries related to DNS dynamic update, type "dynamic update" on the Search tab in this file.

# WriteAuthorityNs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies when the Domain Name System (DNS) server writes NS (name server) records in the Authority section of a response. This entry prevents the DNS server from writing unnecessary NS records in the Authority section, and it ensures that the DNS server complies with relevant Requests for Comments (RFCs).

| Value | Meaning |
|---|---|
| 0 | The DNS server writes NS records in the Authority section of referrals only. This setting complies with RFC 1034, Domain names—concepts and facilities, and with RFC 2181, Clarifications to the DNS Specification. |
| 1 | The DNS server writes NS records in the Authority section of all successful authoritative responses. These NS records are neither required nor useful. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**

To change the value of this entry, use Dnscmd.exe, a tool included in Windows Support Tools.

**Activation Method**

Changes made by using Dnscmd.exe are effective immediately; you do not have to restart the DNS server. If you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

### Note

- The default value is optimal for most DNS servers. Providing NS records in the Authority section consumes processor time and network bandwidth, and it is not recommended unless a network application or service requires it.

# XfrConnectTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\DNS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *seconds* | 1E (30 seconds) |

**Description**
Specifies how long the Domain Name System (DNS) server waits for the secondary server to connect to a primary server. If the connection is not established when the value of this entry expires, then the DNS server drops the connection.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use Dnscmd.exe, a tool included in *Windows Server 2003 Support Tools*. The change is effective immediately; you do not have to restart the DNS server.

**Activation Method**
DNS reads its registry entries only when it starts. You can change entries while the DNS server is running by using the DNS console. However, if you change entries by editing the registry, then the changes are not effective until you restart the DNS server.

# Services\EventLog

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The EventLog subkey stores configuration data that Event Viewer uses to manage the event logs.

This subkey contains the standard entries that appear in all Services subkeys. It also contains subkeys, called *Log-name*, representing each of the logs.

Windows Server 2003 creates subkeys for the three standard logs that appear in Event Viewer: Application, Security, and System. Windows services, such as the Domain Name System (DNS), also create subkeys for their own logs. Programmers can create additional log-specific subkeys for custom logs that they create.

# Log-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog

**Description**
The *Log-name* subkeys store configuration data for logs that are kept by Event Viewer. Each *Log-name* subkey is named for the log it configures. Logs include those created by Windows and its services, as well as those created by administrators. This reference documents entries and subkeys that are common to all of the logs.

**Note**

- "Log-name" does not actually appear in the registry. "*Log-name*" is a placeholder for the subkeys of this type in the Services\EventLog subkey.

# DisplayNameFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name*

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *File name* | *Systemroot*\system32\els.dll |

**Description**
Identifies the file that stores the localized name of the event log.

The name stored in the specified file appears as the log name in Event Viewer.

If this entry does not appear in the registry for an event log, then Event Viewer displays the name of the *Log-name* subkey as the log name.

**Caution**

- Do not change the value of this entry. If you do, Event Viewer cannot locate the file containing the log name.

**See Also**
**Concepts**
DisplayNameID

# DisplayNameID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | (None) |

**Description**
Stores the message identification number of the log name string.

This number indicates the message in which the localized display name appears. The message is stored in the file specified by the value of the DisplayNameFile entry.

For example, if the value of **DisplayNameID** is 0x5, then the localized display name for the log appears in Message ID 5.

If this entry does not appear in the registry for an event log, then Event Viewer displays the name of the *Log-name* subkey as the log name.

### Caution

- Do not change the value of this entry. If you do, Event Viewer cannot locate the correct display name for an event log.

**See Also**
**Concepts**
DisplayNameFile

# File

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name*

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path and file name* | %SystemRoot%\System32\Config\AppEvent.evt |

**Description**
Specifies the fully qualified path to the file where each event log is stored. This entry enables Event Viewer and other log processing programs to find the log files.

The value of this entry is displayed in **Log name** in the Properties dialog box for a log in Event Viewer.

**Note**

- If you must move the event log files, do not use variables other than %Systemroot% in the file paths. Event Viewer cannot resolve variables other than %Systemroot% on remote computers.

# MaxSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x10000–0xFFFF0000 (*bytes, in 64 KB increments*) (64 KB–4 GB) | 0x80000 (*512 KB*) |

**Description**
Specifies the maximum size of the log file.

This entry corresponds to the **Maximum log size** option on the Properties page for a log in Event Viewer.

**Change Method**
To change the value of this entry, use Event Viewer.

### Note

- You cannot resolve a full log by increasing the value of this entry. You must clear the log file manually by clicking **Clear all Events** on the Properties page for any log in Event Viewer.

# PrimaryModule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Subkey name* | (None) |

**Description**
Stores the name of the subkey in Services\Eventlog*Log-name* that contains default values for the entries in the Services\Eventlog*Log-name**Source-name* subkey.

These default values are used to display events in Event Viewer when the source program has not specified values in a *Source-name* subkey, or when files specified as the value of an entry cannot be opened or do not contain the required data.

For example, in the Services\Eventlog\Security subkey, the default value of **PrimaryModule** is Security. This value indicates that the Services\Eventlog\Security\Security subkey contains default values for entries that appear in the Services\Eventlog\Security*SourceName* subkey. These values are used to display events in the security log when the source program does not specify valid values for the entries.

**Change Method**
You can change this entry only by editing the registry.

**See Also**
**Concepts**
Log-name
Source-name

# Retention

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFE \| 0xFFFFFFFF (*seconds*) | 0x93A80 (*604,800 seconds = 7 days*) |

**Description**

Specifies the time, in seconds, that records of events are protected from being overwritten. When the age of an event reaches or exceeds the value of this entry, it can be overwritten.

| Value | Meaning |
|---|---|
| 0x0 | Event Viewer overwrites events only when the log reaches the maximum size, as defined by the value of the MaxSize entry. This value corresponds to the **Overwrite events as needed** option on the **Properties** dialog box for a log in Event Viewer. |
| 0x1–0xFFFFFFFE | Specifies how long records are protected. |
| 0xFFFFFFFF | Records are never overwritten. When the log file reaches its maximum size, as defined by the MaxSize entry, you must clear the log manually by clicking **Clear all Events**. Otherwise, the file discards new events. The log file must be cleared before any change to its size can take effect. This value corresponds to the **Do not overwrite events** option in Event Viewer. You can also activate a security feature that prevents a machine from being used if it cannot generate audit events. To activate this feature, set the value of the CrashOnAuditFail entry to a value other than 0. |

**Change Method**

To change the value of this entry, use Event Viewer. This entry corresponds to the options in the **Event log wrapping** section of the Properties page for any log.

### Note

- Event Viewer requires that you enter retention values in days, not seconds, and it enforces a range of one day (86,400 seconds) to 365 days (31,536,000 seconds). If you change the value of this entry by editing the registry, the value you enter is effective. However, if you change any values in the **Event Log Settings** Properties dialog box for any log in Event Viewer, then Event Viewer converts the value of this entry to the nearest day and replaces the value in the registry.

# Sources

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name*

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Names of programs* | (None) |

**Description**
Specifies the applications, services, or groups of applications that write events to this log. Each program listed in the value of this entry is represented by a subkey in the *Log-Name* subkey.

# Source-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name*

**Description**
The *Source-name* subkeys store configuration information for Event Viewer about the log sources specifed by *Source-name* in the registry path. The configuration information tells Event Viewer where to find the events for that source and how to display them. Event Viewer displays those events as part of the log specified by Log-name in the registry path.

A log source is a program, service, or operating system component that writes to an event log. For each log source, Event Viewer maintains a *Source-name* subkey that is named for the log source. "Source-name" does not actually appear in the registry; *Source-name* is a placeholder for this type of subkey.

For example, the Dr. Watson for Windows service records its events in the Application Log in Event Viewer. Therefore, the source, DrWatson, has a subkey in the EventLog\Application subkey.

### Note

- If a log source for which there is no *Source-name* subkey writes an event to an event log, then Event Viewer uses the configuration information in the subkey specified by the registry entry PrimaryModule to display the event.

### Caution

- Do not change the values of these entries by editing the registry. Log sources should use the standard Win32 application programming interfaces (APIs) to register a source program and to write events to the event log.

**See Also**
**Concepts**
Log-name

# CategoryCount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name**Source-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | (None) |

**Description**
Specifies the number of event categories defined by the source program.

This entry is required if the source uses categories and registers a category message file (as specified in the value of the CategoryMessageFile entry).

### Notes

If the **CategoryMessageFile** entry appears in the registry for a source, then a **CategoryCount** entry is also required for that source.

**See Also**
**Concepts**
CategoryMessageFile

# CategoryMessageFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name**Source-name*

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | Path and file name [; Path and file name] | (None) |

**Description**
Stores the location of the file containing categories for the events generated by the source program. Typically, event categories are included in the dynamic-link library (DLL) or executable file for the source program.

The data in the category message file associates events with categories. The category classifies the event, usually by identifying the component or function that generated the event. Each source program can define its own categories and associate them with events. The category of an event is displayed in the **Category** column in Event Viewer.

### Notes

- The same file can be used as the value of the EventMessageFile, **CategoryMessageFile**, and ParameterMessageFile entries.
- If **CategoryMessageFile** appears in the registry, the CategoryCount entry is required.

**See Also**
**Concepts**
CategoryCount
EventMessageFile
ParameterMessageFile

# EventMessageFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name**Source-name*

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | Path and file name [; Path and file name] | (None) |

**Description**
Stores the location of the file that contains event identifiers for the events generated by the source program, and text describing each event. Typically, event identifiers are included in the dynamic-link library (DLL) or executable file for the source program.

A program can enter more than one file name in this value. Because Event Viewer opens the files in the order in which they are listed, you should list files in the order in which event identifiers are most likely to be found.

Also, when displaying the event log of a remote computer, Event Viewer first tries to locate the event identifiers in the first event message file listed for the source on the remote computer. If Event Viewer cannot open the event message file or cannot find the event identifier, it searches in the first event message file listed for the source on the local computer before searching other event message files on the remote or local computers.

### Note

- The same file can be used as the value of the **EventMessageFile**, CategoryMessageFile, and ParameterMessageFile entries.

**See Also**
**Concepts**
CategoryMessageFile
ParameterMessageFile

# ParameterMessageFile

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name**Source-name*

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | Path and file name [; Path and file name] | (None) |

**Description**
Stores the location of the file that contains values for variables in the event description. This file is optional.

The descriptive text in the event message file can contain variables. The variables indicated by %n are resolved by substituting data from the event record. The variables indicated by %%n are resolved by substituting text from a parameter message file (stored in the ParameterMessageFile entry). The Event Log service locates the correct values and replaces the variables before displaying the text in the Description field in Event Viewer.

**Change Method**
You can change this entry only by editing the registry.

### Note

- The same file can be used as the value of the EventMessageFile, CategoryMessageFile, and **ParameterMessageFile** entries.

# TypesSupported

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\EventLog*Log-name**Source-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x1F (*bitmask*) | (None) |

**Description**

Indicates the types of events generated by the source program.

The value of this entry stores the settings of the **Types** check boxes in Event Viewer. (To view the **Types** check boxes, from the **View** menu in Event Viewer, click **Filter**.)

Each bit in the mask represents an event type. To indicate that the source program generates events of a certain type, set the bit for that type to 1. To indicate that the source program supports more than one type of event, set multiple bits, or sum the hexadecimal values for each event type.

| Bitmask value | Meaning |
|---|---|
| 0 | Success |
| 1 | Error |
| 10 binary (2) | Warning |
| 100 binary (4) | Information |
| 1000 binary (8) | Success audit |
| 10000 binary (16) | Failure audit |

**Caution**

- Do not change the value of this entry by editing the registry. Changing this value does not change the types of events generated by a source program. Source programs that add events of a new type should use the standard Win32 application programming interfaces (APIs) to change this value.

# i8042prt

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The i8042prt subkey stores data for the i8042prt driver. The i8042prt driver controls the keyboard and PS/2-compatible mouse for the Intel 8042 controller.

### Note

- For Windows Server 2003, many of the entries in the i8042prt\Parameters subkey are used only when a device is installed. When the system recognizes a new device, it copies these entries from Parameters to a device-specific subkey in the Enum Subkey subkey. Thereafter, it uses only the entries in Enum. For more information, see the entries in the Parameters subkey.

# i8042prt\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt

**Description**
The Parameters subkey stores data for the i8042prt driver. The i8042prt driver controls the keyboard and PS/2-compatible mouse for the Intel 8042 controller.

For Windows Server 2003, many of the entries in the Parameters subkey are used only when a device is installed. When the system recognizes a new device, it copies these entries from Parameters to a device-specific subkey in the Enum Subkey. Thereafter, it uses only the entries in Enum. For more information, see the entries in the Parameters subkey.

**Caution**

- Do not edit the Enum or Parameters subkeys. To change the value of entries in Enum or Parameters, use Device Manager.

# Parameters\EnableWheelDetection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 1 |

**Description**
Specifies whether the driver attempts to detect and enable the wheel on the mouse device. Some devices are equipped with a mouse wheel to provide rapid scrolling and other control features if supported by an application.

| Value | Meaning |
|---|---|
| 0 | Disable mouse wheel detection. The driver does not attempt to determine whether the mouse has a wheel or to enable it. Even if the mouse has a wheel, the wheel does not function in programs. |
| 1 | Detect and enable the mouse wheel. If the driver detects a wheel on the mouse, then it enables the wheel. |
| 2 | Enable without detecting. The driver attempts to enable a mouse wheel without determining whether the mouse has a wheel. |

**Note**

- This entry is used only when a device is installed. When the system recognizes a new device, it copies this entry to a device-specific subkey in the Enum Subkey subkey. Thereafter, it uses only the entries in Enum.
- For mouse devices that do not fully support all Intellimouse and PS/2 functions, including dynamic enabling, set the value of this entry to 2.

**Tip**

- To view or change the value of this entry for a particular PS/2 mouse, use Device Manager. In the tree control, click **Device** and then, in the results pane, click **Mouse**, double-click the device name, and click the **Advanced Settings** tab. This value appears in the **Wheel Detection** box.

**Caution**

- In rare instances, when the system tries to detect and enable the mouse wheel, the system stops responding. To recover, restart Windows Server 2003. If the problem persists, restart by using the LastKnownGood control set.

# Parameters\KeyboardDataQueueSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x64 (*100*) |

**Description**

Specifies the number of keyboard events that the keyboard driver buffers. This entry is also used in calculating the size of the keyboard driver's internal buffer in nonpaged memory pool. To determine the number of bytes to allocate for the buffer, the system multiplies the size of the KEYBOARD_INPUT_DATA structure (defined in NTDDKBD.H) by the value of **KeyboardDataQueueSize**.

Consider increasing this value if the driver does not initialize or operate correctly and the System Log in Event Viewer contains the following message from the source, i8042prt:

```
The ring buffer that stores incoming keyboard data has overflowed (buffer size is configurable via the registry).
```

### Note

- This entry is used only when a device is installed. When the system recognizes a new device, it copies this entry to a device-specific subkey in the Enum Subkey subkey. Thereafter, it uses only the entries in Enum.

# KeyboardDeviceBaseName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Base keyboard port device name* | KeyboardPort |

**Description**

This entry is added to the registry, but it is not used. In Windows NT 4.0 and earlier, the keyboard class driver uses this entry to locate keyboard port device objects. In Windows Server 2003, the keyboard class driver locates non–Plug and Play keyboards by using entries in the DEVICEMAP subkey (HKEY_LOCAL_MACHINE\Hardware\Devicemap\KeyboardClass).

# Parameters\MouseDataQueueSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x64 (*100*) |

**Description**

Specifies the number of mouse events that the mouse driver buffers. This entry is also used in calculating the size of the mouse driver's internal buffer in nonpaged memory pool. To determine the number of bytes to allocate for the buffer, the system multiplies the size of the MOUSE_INPUT_DATA structure (defined in NTDDMOU.H) by the value of **MouseDataQueueSize**.

You might consider increasing this value if the driver does not initialize or operate correctly and the System Log in Event Viewer contains the following message from the source, Inport:

```
The ring buffer that stores incoming mouse data has overflowed (buffer size is configurable via the
Registry).
```

### Note

- This entry is used only when a device is installed. When the system recognizes a new device, it copies this entry to a device-specific subkey in the Enum Subkey subkey. Thereafter, it uses only the entries in Enum.

### Tip

- To view or change the value of this entry for a particular PS/2 mouse, use Device Manager. In the tree control, click **Device** and then, in the results pane, click **Mouse**, double-click the device name, and click the **Advanced Settings** tab. This value appears in the **Input Buffer Length** box.

# MouseResendStallTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*microseconds*) | 3E8 (*1,000 microseconds = 1 millisecond*) |

**Description**
Specifies how long the mouse driver waits for an acknowledgement (ACK) of a reset if a RESEND message is returned without an ACK. This entry is used when the mouse driver interrupt service routine includes a reset.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Unlike some entries in this subkey, this entry remains in the Parameters subkey and is read from this location even after a keyboard is installed.

# Parameters\MouseResolution

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Mouse resolution specifier* | 0x3 |

**Description**
Specifies the resolution of the mouse port. The number of counts per millimeter of resolution is determined by raising 2 to the power of the value of **MouseResolution**.

For example, if the value of MouseResolution is 3, then the mouse resolution is 2**3, or 8 counts per millimeter.

### Note

- This entry is used only when a device is installed. When the system recognizes a new device, it copies this entry to a device-specific subkey in the Enum Subkey subkey. Thereafter, it uses only the entries in Enum.

# Parameters\MouseSynchIn100ns

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Time in 100-nanosecond units* | 20000000 (2 seconds) |

**Description**
Specifies the length of time after which the next mouse interrupt is assumed to indicate the start of a new mouse packet (partial packets are discarded). This enables the driver to synchronize its internal record of the mouse packet state with the hardware state when a mouse interrupt is lost.

Consider modifying this value if the system appears to be generating extra mouse events, such as when objects are selected although the mouse button has not been pressed.

**Note**

- This entry is used only when a device is installed. When the system recognizes a new device, it copies this entry to a device-specific subkey in the Enum Subkey subkey. Thereafter, it uses only the entries in Enum.

# NumberOfButtons Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x2 |

**Description**
Specifies the number of buttons on the mouse-port mouse at startup. If the number of buttons detected at startup is incorrect, then you can override it by changing the value of this entry.

### Note

- This entry is used only when a device is installed. When the system recognizes a new device, it copies this entry to a device-specific subkey in the Enum Subkey subkey. Thereafter, it uses only the entries in Enum.

# OverrideKeyboardSubtype

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x0 |

**Description**
Specifies the OEM-dependent keyboard subtype. You can add this entry to the registry to correct an error in the keyboard subtype detected at startup.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- Unlike some entries in this subkey, this entry remains in the Parameters subkey and is read from this location even after a keyboard is installed.

# OverrideKeyboardType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | (None) |

**Description**

Specifies the keyboard type. You can add this entry to the registry to correct an error in the keyboard type detected at startup.

The following table shows some commonly used values for this entry. This is not a complete list.

| Value | Meaning |
|---|---|
| 0x2 | Enhanced 101-key keyboard or compatible |
| 0x4 | Enhanced 102-key keyboard or compatible |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- Unlike some entries in this subkey, this entry remains in the Parameters subkey and is read from this location even after a keyboard is installed.

# PointerDeviceBaseName Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Base mouse port device name | PointerPort |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 and earlier.

In Windows NT 4.0 and earlier, the mouse port driver used this entry to locate pointer port device objects. In Windows Server 2003, the mouse port driver locates non–Plug and Play pointing devices by using entries in the DEVICEMAP subkey (HKEY_LOCAL_MACHINE\Hardware\Devicemap\PointerClass).

**Caution**

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

**See Also**
**Concepts**
DEVICEMAP

# PollingIterations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x400–0xFFFFFFFF (*attempts*) | 0x2EE0 (*12,000 attempts*) |

**Description**
Specifies the maximum number of times Windows Server 2003 polls the hardware. If the number of trials specified in this entry is exceeded, then Windows Server 2003 stops polling. This entry is used on keyboards other than old-style AT keyboards.

Consider increasing this value if the driver does not initialize or operate correctly and the System Log in Event Viewer contains the following message from the source, i8042prt:

```
The operation on <driver name> timed out (time out is configurable via the registry).
```

### Note

- Unlike some entries in this subkey, this entry remains in the Parameters subkey and is read from this location even after a keyboard is installed.

**See Also**
**Concepts**
PollingIterationsMaximum

# PollingIterationsMaximum

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x400–0xFFFFFFFF (*attempts*) | 0x2EE0 (*12,000 attempts*) |

**Description**
Specifies the maximum number of times Windows Server 2003 polls the hardware on older-style AT keyboards. If the number of trials specified in this entry is exceeded, then Windows Server 2003 stops polling.

You might consider increasing this value if the driver does not initialize or operate correctly and the System Log in Event Viewer contains the following message from the source, i8042prt:

```
The operation on <driver name> timed out (time out is configurable via the Registry).
```

### Note

- Unlike some entries in this subkey, this entry remains in the Parameters subkey and is read from this location even after a keyboard is installed.

**See Also**
**Concepts**
PollingIterations

# PollStatusIterations Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x1 |

**Description**
Specifies the maximum number of times the system verifies interrupts on the i8042 controller status register. If the interrupt cannot be verified in the number of trials specified in the value of this entry, then the interrupt is ignored.

This entry is designed to minimize the effect of interrupts raised before the Output Buffer Full bit is set in the i8042 status register. This condition occurs on some computers, such as the Olivetti MIPS computer.

If the system seems to ignore keyboard interrupts, then increase the value of this entry. The system might mistakenly ignore an interrupt if the Output Buffer Full bit is not set when the interrupt is raised. However, if you increase the value of this entry, then the Interrupt Service Routine will take longer to execute, because there is a one-microsecond delay following each check of the Output Buffer Full bit.

### Note

- This entry is used only when a device is installed. When the system recognizes a new device, it copies this entry to a device-specific subkey in the Enum Subkey subkey. Thereafter, it uses only the entries in Enum.

### Tip

- To determine whether the driver is processing keyboard interrupts correctly, press the NUM LOCK key. If the **Num Lock** light on the keyboard turns on or off, this indicates that the i8042prt driver handled the keyboard interrupt correctly.

# ResendIterations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x3 |

**Description**
Specifies the maximum number of times a hardware operation is attempted. If the number of trials exceeds the value of this entry, then Windows Server 2003 considers the operation to have failed.

You might consider increasing this value if the driver does not initialize or operate correctly and the System Log in Event Viewer contains the following message from the source, i8042prt:

```
Exceeded the allowable number of retries (configurable via the Registry) on device <device-name>.
```

### Note

- Unlike some entries in this subkey, this entry remains in the Parameters subkey and is read from this location even after a keyboard is installed.

# SampleRate Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 20 \| 40 \| 60 \| 80 \| 100 \| 200 | 100 Hz |

**Description**
Specifies how often, in Hertz, the PS/2 driver measures the characteristics and activities of the PS/2 mouse. The driver uses the information gathered through sampling to optimize the operation of the mouse device.

### Note

- This entry is used only when a device is installed. When the system recognizes a new device, it copies this entry to a device-specific subkey in the Enum Subkey subkey. Thereafter, it uses only the entries in Enum.

### Tip

- To view or change the value of this entry for a particular PS/2 mouse, use Device Manager. In the tree control, click **Device** and then, in the results pane, click **Mouse**, double-click the device name, and click the **Advanced Settings** tab. This value appears in the **Sample Rate** box.

# IAS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
Stores configuration data for the Internet Authentication Service (IAS) in Windows Server 2003.

IAS performs centralized authentication, authorization, auditing, and accounting of users who connect to a network by using the virtual private network (VPN) and dial-up functions.

# IAS\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\IAS

**Description**
The Parameters subkey stores configuration data for the Internet Authentication Service (IAS) in Windows Server 2003.

IAS performs centralized authentication, authorization, auditing, and accounting of users who connect to a network by using the virtual private network (VPN) and dial-up functions.

# Allow SNMP Set

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\IAS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether Internet Authentication Service (IAS) accepts incoming Simple Network Management Protocol (SNMP) *Set* messages. This entry applies only when SNMP is being used to monitor IAS.

*Set* messages change or update values in an SNMP management information database (MIB). IAS accepts these messages only when an agent has permission to write to the MIB.

Set messages are prohibited by default, but you can add this entry to the registry to permit the to accept Set messages.

| Value | Meaning |
|---|---|
| 0 | IAS rejects Set messages. |
| 1 | IAS accepts Set messages. |

# Ping User-Name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\IAS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *User name* | (None) |

**Description**
Specifies a user name (or user name pattern with variables) that the Internet Authentication Service (IAS) recognizes as a fictitious name. As a result, IAS rejects authentication requests, and it responds to all accounting requests from the network access server (NAS). Also, IAS does not record transactions involving this user in any log files.

Proxy servers implementing the Remote Authentication Dial-In User Service (RADIUS) protocol and network access servers periodically send authentication and accounting requests (known as *ping requests*) to verify that the server is responsive. These ping requests include fictional user names.

This entry helps IAS recognize the fictitious name or naming scheme used in ping requests. This method is likely to improve IAS performance and to make the event log easier to interpret.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Tip**

- To indicate more than one user name, enter a name pattern, such as a Domain Name System (DNS) name including wildcard characters. For more information about pattern-matching, see Help and Support Center for Windows Server 2003.

# LanmanServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The LanmanServer subkey stores configuration data for the LAN Manager 2.*x* Server service.

# LanmanServer\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer

**Description**
The Parameters subkey stores configuration data that is specific to the LAN Manager 2.*x* Server service.

# AlertSchedule

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFF (*1–65,535 minutes*) | 0x5 |

**Description**

Specifies how often the Server service checks alert conditions and sends alert messages. This entry also defines the time interval used in calculating the general error rate (**ErrorThreshold**) and the rate of network errors (**NetworkErrorThreshold**).

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# ConnectionNoSessionsTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF (*minutes*) | 0x2 |

**Description**

Specifies how long the Server service maintains an unused virtual circuit. If the time specified in value of this entry expires, and the client has not established a session, then the Server service closes the virtual circuit.

Decreasing this value can conserve server resources, but it might impair performance by requiring more virtual circuits to be re-established.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# DiskSpaceThreshold

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–99 (*percent*) | 10 |

**Description**
Specifies the percentage of disk space that must remain free for use. If the percentage of free space falls below the value of this entry, then the Server service records the following event in the System Log in Event Viewer:

```
"The <disk-letter> disk is at or near capacity. You might need to delete some files." (Srv Event ID 2013)
```

If the value of this entry is zero, then the system does not check the percentage of free space on the disk; it also does not check for free space in megabytes as specified by LowDiskSpaceMinimum.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
LowDiskSpaceMinimum

# EnableWFW311DirectIpx

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**

Specifies whether older, direct-hosted Internetwork Packet Exchange (IPX) clients can connect to this server.

Clients running earlier versions of Windows for Workgroups provide inadequate support for named pipes when running over direct-hosted IPX, causing named-pipe applications to stop responding. You can prevent these clients from connecting to the server by setting the value of this entry to 0.

| Value | Meaning |
|---|---|
| 0 | Direct-hosted IPX clients cannot connect to the server. |
| 1 | Direct-hosted IPX clients can connect to the server. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# hidden

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_BINARY | 0 \| 1 | 0 |

**Description**
Specifies whether the server's computer name is displayed to other computers in the domain.

If the value of this entry is 0, or if this entry is absent from the registry, then users can see the server's computer name in the following displays:

- In the Network Connection Wizard.
- In the **Browse for Network Connections** dialog box (In the **Map Network Drive** dialog box, click **Browse**, or in My Network Places, click **Add Network Place**).
- In My Network Places, by clicking **Entire Network** or **Computers Near Me**.

To omit this computer's name from the display, add this entry to the registry and change its value to 1.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# IRPStackSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 11–50 (*stack locations*) | 15 |

**Description**

Specifies how many stack locations the Server service establishes for I/O Request Packets (IRPs).

It might be necessary to increase this number for certain transports, or if you have many file system drivers installed on the system. Each stack uses 36 bytes of memory for each receive buffer (also known as a *work item*).

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- Some older applications may incorrectly set the value of this entry to a value that is outside (below) the valid range. By deleting the invalid value, you can restore this entry to its default value. You can also set the value yourself.

# LowDiskSpaceMinimum

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF MB | 400 MB |

**Description**

Specifies the amount of disk space that must remain free for use. If the amount of free space falls below the value of this entry and below the percentage specified by the value of the DiskSpaceThreshold entry, then the Server service records the following event in the System Log in Event Viewer:

```
"The <disk-letter> disk is at or near capacity. You might need to delete some files." (Srv Event ID 2013)
```

If the value of this entry is 0, then the system does not check whether free disk space has decreased below a specific number of megabytes. However, unlike DiskSpaceThreshold (which when equal to zero prevents both checks of free space), the value of this entry has no effect on the check specified by DiskSpaceThreshold.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# RestrictNullSessAccess

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the Server service limits access to the system by clients that are logged on to the system account without user name and password authentication.

| Value | Meaning |
|---|---|
| 0 | Null session access is not restricted. Unauthenticated users access all shared resources. |
| 1 | Null session access is restricted. Unauthenticated users can access only the server pipes listed in the value of the **NullSessionPipes** entry and the shared directories listed in the value of the **NullSessionShares** entry. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# SessConns

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–2,048 (*connections*) | 2048 |

**Description**
Specifies the maximum number of connections to network shares (tree connects) that are permitted for a single network connection. If the client requests more connections, then the Server service returns an error message.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# SessOpens

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–16,384 (*files*) | 16,384 |

**Description**
Specifies the maximum number of open files permitted for each connection. It is possible to have multiple sessions on a single connection.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# SessUsers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–2,048 (*users*) | 2,048 |

**Description**
Specifies the maximum number of users that can be logged on to each virtual circuit.

**Note**

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# SharingViolationDelay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–1,000 (*milliseconds*) | 200 |

**Description**
Specifies how often the Server service repeats a file operation when the initial request resulted in a sharing violation error. The value of this entry specifies the minimum time that must elapse between repeated attempts to open, rename, or delete a file.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
SharingViolationRetries

# SharingViolationRetries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–1,000 (*attempts*) | 5 |

**Description**
Specifies the maximum number of times the Server service repeats an attempted file operation when the request results in a sharing violation error.

The Server service repeats the operation, at a rate determined by the value of the SharingViolationDelay entry, until either the operation is successful or the value of this entry is reached. If a client requests more attempts than the value of this entry permits, then the Server service returns an error. This entry applies to open, rename, and delete file operations.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
SharingViolationDelay

# SizReqBuf

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 512–65,536 (*bytes*) | |

**Description**

Specifies the size of the request buffers that the Server service uses.

Small buffers use less memory, but large buffers can improve performance.

### Notes

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.
- For computers running Windows Server 2003 and with 512 MB or more of physical memory, the default size of the request buffers is 16,644 bytes; for servers with less physical memory, the default size is 4,356 bytes. If this entry is present in the registry, its value overrides the default value.

# srvcomment

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Text string* | (None) |

**Description**
Specifies the text that appears in the **Comment** field next to the name of this computer in My Network Places.

To view the comment text, in My Network Places, double-click **Entire Network** or **Computers Near Me**. From the **View** menu, either click **Details** or double click the name of any computer. Or, in My Network Places, double-click **Add Network Place,** and view the comment text in the **Browse for Network Resource** dialog box.

**Change Method**
To enter a computer description, right click **My Computer**, select **Properties**, and in the **Computer Name** tab, type the text you want to use in the **Computer description** text box, and then click **Apply**.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# ThreadPriority

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 \| 15 | 1 |

**Description**

Specifies the priority of all Server service threads in relation to the base priority of the Server service process.

| Value | Meaning |
|---|---|
| 0 | Run at the same priority as the process. |
| 1 | Run at one priority level higher than the process. This is equivalent to a process running in the foreground, where it interacts with users. |
| 2 | Run at two priority levels higher than the process. |
| 15 | Run at real-time priority. Threads running in real time can exclude all other threads, including essential system threads and threads that process mouse and keyboard commands. |

**Note**

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Users

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFE \| 0xFFFFFFFF (*users*) | 0xFFFFFFFF |

**Description**
Specifies whether a limit exists for the number of users that can be logged on to the Server service simultaneously, and if so, specifies what the limit is.

| Value | Meaning |
|---|---|
| 0x1–0xFFFFFFFE | Specifies the maximum number of users who can be logged on simultaneously. |
| 0xFFFFFFFF | There is no limit to the number of users who can be logged on simultaneously. |

**Change Method**
You can change this entry only by editing the registry.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# LanmanWorkstation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The LanmanWorkstation subkey stores configuration data for the Workstation service. The Workstation service provides network connections and communications on computers running Windows XP Professional and Windows Server 2003.

# LanmanWorkstation\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation

**Description**
The Parameters subkey stores configuration data for the Workstation service. The Workstation service provides network connections and communications on computers running Windows XP Professional and Windows Server 2003.

# BufNamedPipes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**

Specifies whether the redirector caches character-mode named pipes. If named pipes are not stored in the cache, then they are flushed to the server immediately, and read-ahead is disabled on these named pipes.

| Value | Meaning |
|---|---|
| 0 | Character-mode named pipes are not cached. |
| 1 | Permits the Workstation service to cache character-mode named pipes. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# KeepConn

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–65,535 (*seconds*) | 600 (*10 minutes*) |

**Description**

Specifies the maximum amount of time that an idle connection can remain open. If the idle time for a connection reaches the value of this entry, then the connection is closed.

Increase the value of this entry if your application closes and opens Universal Naming Convention (UNC) files on a server less frequently than every 10 minutes. This decreases the number of reconnections to a server.

# LockIncrement

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*milliseconds*) | 0xA (*10*) |

**Description**
Specifies how long an OS/2–based application can wait when a request for a lock fails. If the lock request cannot be granted before the time specified in the value of this entry expires, then the redirector rejects the lock request. This entry is used either if OS/2–based applications running on servers (except servers using LAN Manager version 2.0 and later) request that a lock operation wait indefinitely or if the lock cannot be granted immediately.

**Note**

- Do not change this value unless you are running an OS/2-based application that requests lock operations that might fail. This entry does not affect Windows 32-bit applications.

# LockMaximum

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*milliseconds*) | 0x1F4 (*500*) |

**Description**
Used to configure the lock backoff package. This entry prevents an application from consuming server time by issuing nonblocking requests when no data is available for the application.

# LockQuota

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*bytes*) | 0x1800 (*6 KB*) |

**Description**

Specifies the maximum amount of data that can be read from a file when the value of the UseLockReadUnlock entry is 1.

Consider increasing the value of this entry if your application performs a significant number of lock-and-read operations. These are operations in which an application locks data and then immediately reads the data.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Increasing the value of this entry to more than 2 MB and using an application that locks megabtyes of data can cause the system to deplete its paged pool.

# MaxCmds

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 50–65535 (*network control blocks*) | 50 |

**Description**

Specifies the maximum number of network control blocks that the redirector can reserve. The value of this entry coincides with the number of execution threads that can be outstanding simultaneously.

Increase this value to improve network throughput, especially if you are running applications that perform more than 15 operations simultaneously. However, because this entry also limits the number of outstanding execution threads, your network performance might not improve. Each additional execution thread uses a margin of 1 KB of nonpaged pool when the network is at capacity. However, these resources are not consumed until the user references data in the network control block.

# MaxCollectionCount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFF (*0–65,535 bytes*) | 0x10 (*16*) |

**Description**
Specifies the amount of data that must be present in the character-mode buffer of a named pipe to trigger a write operation. If the amount of data in the buffer meets or exceeds this value, then it is written immediately. Otherwise, it is retained in the buffer until either more data is added or the value of the **CollectionTime** entry expires.

Increasing the value of this entry can improve the performance of named-pipe applications, but it does not affect applications that do their own buffering, such as SQL Server applications.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
SizCharBuf

# SizCharBuf

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 64–4,096 (*bytes*) | 512 |

**Description**

Specifies the size of the character buffers for a named pipe. The redirector maintains a read-ahead buffer and a write-behind buffer for each pipe. This entry establishes the size of both buffers.

The buffer size determines the amount of data the redirector reads and writes.

When reading, the redirector attempts to read enough data to fill the read buffer. If the data it is reading is smaller than the buffer, then the redirector reads ahead until the buffer is full. If the data is larger than the buffer, then the redirector bypasses the character-mode buffer and reads the data directly into the user buffer.

When writing, the redirector collects data in the character-mode buffer until it meets or exceeds the size specified by the value of the MaxCollectionCount entry, or until the time specified by the value of the **CollectionTime** entry expires.

Increasing the value of this entry can improve the performance of named-pipe applications, but it does not affect applications that do their own buffering, such as SQL Server applications.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# UseLockReadUnlock

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the lock-and-read and write-and-unlock performance enhancements are enabled. These features improve performance when an application locks data and then immediately reads the data, or writes data and then immediately unlocks it.

| Value | Meaning |
|---|---|
| 0 | Lock-and-read and write-and-unlock features are not enabled. |
| 1 | Lock-and-read and write-and-unlock features are enabled. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Setting this value to 1 usually improves performance significantly. However, it degrades the performance of database applications that lock a range of data and do not allow data within that range to be read.

**See Also**
**Concepts**
LockQuota

# UseOpportunisticLocking

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the opportunistic-locking (oplock) performance enhancement is enabled. If enabled, the redirector requests an opportunistic lock on any file opened in "Deny None" mode. As a result, the server performs automatic read-ahead and write-behind caching on behalf of the redirector.

| Value | Meaning |
|---|---|
| 0 | Disables opportunistic locking |
| 1 | Enables opportunistic locking |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- Opportunistic locks improve performance significantly. Do not change the value of this entry unless you are testing the system.

# UseUnlockBehind

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the unlock-behind optimization feature is enabled. If enabled, the redirector unlocks data immediately in response to an unlock request. It does not wait for confirmation from the server that the unlock operation is complete.

| Value | Meaning |
|---|---|
| 0 | Disables the unlock-behind optimization feature |
| 1 | Enables the unlock-behind optimization feature |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- Unlock-behind optimization improves performance significantly. Do not change the value of this entry unless you are testing the system.

# Mouclass

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
Stores configuration data for the Mouseclass device driver. The Mouseclass device driver handles hardware-independent operations on a specific class of devices—in this case, the mouse and other pointing devices. The port drivers manage the hardware-specific operations.

# Mouclass\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Mouclass

**Description**
The Parameters subkey in the Mouclass subkey stores configuration data for the Mouseclass device driver. The Mouseclass device driver handles hardware-independent operations on a specific class of devices—in this case, the mouse and other pointing devices. The port drivers manage the hardware-specific operations.

# ConnectMultiplePorts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Mouclass\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether there is one or more than one port device object for each class device object. This entry is used primarily by device drivers.

| Value | Meaning |
|---|---|
| 0 | 1:1. Indicates that the class device object created by the driver is connected to one port device object. |
| 1 | 1:Many. Indicates that the class device object created by the driver is connected to multiple port device objects. As a result, the events generated by all pointing devices on the system are fed to the Windows subsystem in a single input stream. |

The mouse class driver can route all information from all the declared mouse ports in the system to a single class device object. By doing this, the class driver can work with older raw input threads that cannot open multiple, independent mouse devices. Windows Server 2003 reads all mouse devices independently, so this value is set to 0 by default.

**See Also**
**Concepts**
MaximumPortsServiced

# MaximumPortsServiced

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Mouclass\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x3 |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# Parameters\PointerDeviceBaseName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Mouclass\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Base port device name* | PointerClass |

**Description**
Specifies the base name for the device objects created by the mouse class device driver.

### Note

- The device driver stores information about the device in a subkey of DEVICEMAP (HKEY_LOCAL_MACHINE\HARDWARE\DEVICEMAP) that has the same name as the value of this entry. The pointer class driver uses the information in the DEVICEMAP subkey to locate all non–Plug and Play pointer port device objects. (The Plug and Play subsystem passes all instances of Plug and Play port device objects to the class driver's AddDevice routine, which negates the need for the class driver to search for devices.)

# MRxSmb

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The MRxSmb subkey stores configuration data for the Browser Server service. The Browser Server service provides the list of participating browser servers on the network.

# Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\MRxSmb

**Description**
The Parameters subkey stores configuration data for the Browser Server service. The Browser Server service provides the list of participating browser servers on the network.

# MailslotDatagramThreshold

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\MRxSmb\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFE \| 0xFFFFFFFF | 0xFFFFFFFF |

**Description**
Establishes the maximum number of mailslot requests that can be missed in one hour without generating an event. If more requests are missed, then the browser service writes an event to the System Log in Event Viewer.

| Value | Meaning |
|---|---|
| 0x0–0xFFFFFFFE | The number of requests that triggers an Event Log message. |
| 0xFFFFFFFF | No events are generated when requests are missed. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe. In Windows NT 3.1 and earlier, add this entry to the DGRcvr\Parameters subkey.

# NDIS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
Stores configuration data for the Network Driver Interface Specification (NDIS) interface in Windows Server 2003, Ndis.sys.

Ndis.sys supports communication between protocol drivers, network interface drivers, and network adapters. It acts as an intermediary between protocol drivers and network adapters, so that they can remain independent. This allows a network adapter to support multiple network protocols, such as TCP/IP and Internetwork Packet Exchange (IPX). It also allows multiple network interface drivers and multiple network adapters to co-exist within a system.

# NDIS/Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NDIS

**Description**
The Parameters subkey stores configuration data for the Network Driver Interface Specification (NDIS) 5.0 interface in Windows Server 2003, Ndis.sys.

Ndis.sys supports communication between protocol drivers, network interface drivers, and network adapters. Ndis.sys acts as an intermediary between protocol drivers and network adapters, so that they can remain independent. This allows a network adapter to support multiple network protocols, such as TCP/IP and IPX. It also allows multiple network interface drivers and multiple network adapters to co-exist within a system.

# ProcessorAffinityMask

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NDIS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFE \| 0xFFFFFFFF | 0xFFFFFFFF (*-1*) |

**Description**

Specifies which processors can be associated with a network adapter and can service the delayed procedure calls (DPCs) generated by that network adapter. This entry is designed for multiprocessor computers, specifically those with more than one network adapter.

Each of the 32 bits in the value of this entry represents one processor. The high bit represents processor 32, and the low bit represents processor 0.

If a bit is set to 1, then the processor represented by the bit can be associated with a network adapter and service all of the DPCs it generates.

If a bit is set to 0, then the processor represented by the bit cannot be associated with any particular network adapter. When a network adapter is not associated with a processor, its DPCs are serviced by the same processor that serviced the interrupt from which the DPC evolved.

| Value | Meaning |
|---|---|
| 0x0 \| 0xFFFFFFFF | If all of the bits are set to either 0 or 1, then none of the processors are associated with network adapters. DPCs are serviced by the same processor that serviced the interrupt. This setting is useful for platforms that distribute interrupts among all processors. These values are treated the same to skip the affinity association and thus avoid DPC redirecting. |
| 0x1–0xFFFFFFFE | If some of the bits are set to 0 and others are set to 1, then the system skips processors associated with bits set to 0, and it assigns network adapters to processors whose bits are set to 1. |

**Note**

- Windows Server 2003 sets this entry to the optimum value for your system. Do not modify the value of this entry.

# NdisWan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The NdisWan subkey stores configuration data for NDISWAN drivers supporting remote access service. This subkey appears in the registry only when remote access is installed on the computer.

### Note

- NDISWAN is an intermediate-level Network Device Interface Specification (NDIS) driver that lies between an 802.3 miniport and a WAN miniport, converting NDIS packets to WAN packets. NDISWAN supports Point-to-Point Protocol (PPP) protocol/link framing, compression, and encryption.

# NdisWan\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NdisWan

**Description**
The Parameters subkey stores configuration data for NDISWAN drivers supporting remote access service. The NdisWan subkey appears in the registry only when remote access service is installed on the computer.

NDISWAN is an intermediate-level Network Device Interface Specification (NDIS) driver that lies between an 802.3 miniport and a WAN miniport, converting NDIS packets to WAN packets. NDISWAN supports Point-to-Point Protocol (PPP) protocol/link framing, compression, and encryption.

# IGMPIdle

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NdisWan\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether NdisWan drivers consider Internet Group Management Protocol (IGMP) packets to be valid communications traffic.

Many servers send IGMP packets periodically to prevent a silent, but active, line from being mistaken for an idle line. These packets can be misinterpreted as communication. They can also prevent a truly idle line from being disconnected.

| Value | Meaning |
|---|---|
| 0 | IGMP packets are considered to be valid communications traffic, which keeps a line active. |
| 1 | IGMP packets are filtered out and not counted as valid communications traffic. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# NetBT

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# NetBT\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# BcastNameQueryCount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 - 0xFFFF | 3 |

**Description**
Specifies the number of times that NetBIOS over TCP/IP (NetBT) broadcasts a query for a unique name without receiving a response.

**Activation Method**
If this interface is configured by Dynamic Host Configuration Protocol (DHCP), then a change in parameters will take effect if you enter the command **ipconfig /renew** at the command line. Otherwise, you must restart the computer for a change to take effect.

### Note

- This entry is installed with its default value during installation of the TCP/IP components.

**See Also**
**Concepts**
BcastQueryTimeout
NameSrvQueryCount

# BcastQueryTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 100 to 0xFFFFFFFF | 750 (0x2EE) |

**Description**
Specifies the interval, in milliseconds, between successive broadcast queries for a single name. This entry is installed with its default value during installation of the TCP/IP components.

**Activation Method**
If this interface is configured by DHCP, then a change in parameters will take effect if you enter the command **ipconfig /renew** at the command line. Otherwise, you must restart the computer for a change to take effect.

**See Also**
**Concepts**
BcastNameQueryCount
NameSrvQueryTimeout

# BroadcastAddress

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 - 0xFFFFFFFF | |

**Description**
Specifies whether NetBIOS over TCP/IP (NetBT) uses a specific address for all name-related broadcasts. This entry does not exist by default, and you must add it to the registry to change the default broadcast address. Also, this entry sets the broadcast address to which NetBT sends name-related packets for all the subnets to which NetBT is bound.

**Activation Method**
If this interface is configured by DHCP, then enter the command **ipconfig /renew** at the command line for a change in parameters to take effect. Otherwise, you must restart the computer for a change to take effect.

### Note

- By default, NetBT uses the ones broadcast address appropriate for each net. For example, for a network of 11.101.0.0 with a subnet mask of 255.255.0.0, the subnet broadcast address is 11.101.255.255. If the network uses the zeroes broadcast address (which is configured by setting **UseZeroBroadcast**), then create this entry and set it to zero. In this example, the appropriate subnet broadcast address is 11.101.0.0, and the value of this entry is set to 0x0b650000.

# CacheTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 60000 to 0xFFFFFFFF | 600000 milliseconds (10 minutes) |

**Description**
Specifies the interval, in milliseconds, that NetBIOS names are cached by NetBIOS over TCP/IP (NetBT).

**Activation Method**
If this interface is configured by DHCP, then enter the command **ipconfig /renew** at the command line for a change in parameters to take effect. Otherwise, you must restart the computer for a change to take effect.

### Note

- This entry is installed with its default value during installation of the TCP/IP components.

# ConnectOnRequestedInterfaceOnly

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether NetBIOS connections to this computer are restricted to a particular adapter. When the redirector on a multihomed computer calls another computer name, it places calls on all NetBIOS over TCP/IP (NetBT) *transports* (protocol-adapter combinations) to which it is bound. Each transport independently attempts to reach the target name. Setting the value of this entry to 1 limits each transport to connecting to other computers that are reachable via its own adapter, preventing crossover traffic.

This entry is ignored if **CachePerAdapterEnabled** is absent from the registry or if it is present and its value is zero.

**Activation Method**
You must restart the computer for changes to take effect.

### Note

- This entry does not exist by default, and you must add it to the registry and set the value to 1 if you want to restrict NetBIOS connections to a specific interface. You can add this entry to the registry and set its value by using the registry editor Regedit.exe.

**See Also**
**Concepts**
RandomAdapter

# DhcpNodeType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 8 |

**Description**
This parameter specifies the NetBIOS over TCP/IP (NetBT) node type. The node type specified by the entry NodeType, if valid, overrides the value of this entry. If neither this entry nor **NodeType** exist in the registry, and there are no WINS servers configured on the network, then NetBT uses B-node name resolution. If there is at least one WINS server available, NetBT uses H-node name resolution.

| Value | Meaning |
|---|---|
| 1 | B-node. NetBT uses IP broadcast messages to register and resolve IP addresses from NetBIOS names. |
| 2 | P-node. NetBT uses point-to-point communication with a NetBIOS name server (in Windows Server 2003–based networks, this is the WINS server) to register and resolve IP addresses from NetBIOS names. |
| 4 | M-node. NetBT uses a mix of B-node and P-node to register and resolve NetBIOS names. M-node first uses broadcast resolution (B-node), then if necessary, uses a server query (P-node). |
| 8 | H-node. NetBT uses a hybrid of B-node and P-node. An H-node client always directly queries a name server first and then uses broadcasts only if the direct query fails. |

**Note**

- For more information on configuring WINS, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

**Caution**

This entry is created and maintained by the DHCP client service, if enabled, and you should not modify it. It is included here only for reference.

**See Also**
**Concepts**
NodeType

# DhcpScopeID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**
Stores the scope identifier that this computer received from the Dynamic Host Configuration Protocol (DHCP) server. The scope identifier, a dot-separated string (for example, microsoft.com) is appended to the NetBIOS over TCP/IP (NetBT) name, and should not be confused with DHCP scope (which specifies a range of addresses available for assignment by the DHCP server).

### Caution

This entry is created and maintained by the DHCP client service, if enabled, and you should not modify it. It is included here only for reference.

# EnableLMHOSTS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether NetBIOS over TCP/IP (NetBT) searches the Lmhosts file for names that it has not been able to resolve with name queries to a WINS server and broadcast name queries to the network. Though use of the Lmhosts file is enabled by default, there is no default Lmhosts file. When the value of this entry is 1 and there is no Lmhosts file, no action takes place.

| Value | Meaning |
|---|---|
| 0 | NetBT does not search the Lmhosts file. |
| 1 | NetBT searches the Lmhosts file, if it exists, for names that cannot be resolved by WINS or broadcast name queries. |

**Change Method**
To change the value of this entry, use **Network Connections** in Control Panel. Right-click **Local Area Connection** (or the name of another connection, if appropriate), click **Properties**, click **Internet Protocol (TCP/IP)** , and then click the **Properties** button. On the **Internet Protocol (TCP/IP) Properties** page, click the **Advanced** button, and click the **WINS** tab. This entry is associated with **Enable LMHOSTS lookup**, which you can check to enable and clear to disable.

**Activation Method**
Changes made to this entry through the **Properties** page for the interface take effect immediately.

### Note

- The path to the Lmhosts file is specified in the **DatabasePath** entry.

# InitialRefreshT.O.

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 960000 to 0xFFFFFFF | 960000 (16 minutes) |

**Description**
Specifies the initial refresh timeout, in milliseconds, used by NetBIOS over TCP/IP (NetBT) during name registration. NetBT tries to contact the WINS servers at 1/8th of this time interval when it is first registering names. When it receives a successful registration response, it uses the refresh interval specified in the response for subsequent registration attempts.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
If this interface is configured by DHCP, then a change in parameters will take effect if you enter the command **ipconfig /renew** at the command line. Otherwise, you must restart the computer for a change to take effect.

# LmhostsTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1000 to 0xFFFFFFFF | 6000 milliseconds (6 seconds) |

**Description**

Specifies the timeout interval, in milliseconds, for NetBIOS over TCP/IP (NetBT) name queries to Lmhosts and DNS. The units used by the timer are equal in length to the value of this entry, so the actual timeout can be as much as twice the value of this entry.

**Activation Method**

If this interface is configured by DHCP, then enter the command **ipconfig /renew** at the command line for a change in parameters to take effect. Otherwise, you must restart the computer for a change to take effect.

# MaxDgramBuffering

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 to 0xFFFFFFFF | 0x20000 (128 Kb) |

**Description**
Specifies the maximum amount of memory that NetBIOS over TCP/IP (NetBT) dynamically allocates for all outstanding sent datagrams. Once this limit is reached, attempting to send additional datagrams would fail due to insufficient resources.

**Activation Method**
If this interface is configured by DHCP, then enter the command **ipconfig /renew** at the command line for a change in parameters to take effect. Otherwise, you must restart the computer for a change to take effect.

# NameServerPort

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 to 65535 | 137 |

**Description**
Specifies the IP port number to which NetBIOS over TCP/IP (NetBT) sends packets related to name service, such as name queries and registrations. The Windows Internet Name Service (WINS), which provides name service for NetBIOS names, listens on port 137, the default for this entry. NetBIOS name servers from other vendors may listen on different ports. This entry is installed with its default value during installation of the TCP/IP components.

**Activation Method**
If this interface is configured by DHCP, then a change in parameters will take effect if you enter the command **ipconfig /renew** at the command line. Otherwise, you must restart the computer for a change to take effect.

# NameSrvQueryCount

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 to 0xFFFF | 3 |

**Description**
Specifies the number of times that NetBIOS over TCP/IP (NetBT) sends a query to a Windows Internet Name Service (WINS) server for a given name without receiving a response. This entry is installed with its default value during installation of the TCP/IP components.

**Activation Method**
If this interface is configured by Dynamic Host Configuration Protocol (DHCP), then a change in parameters will take effect if you enter the command **ipconfig /renew** at the command line. Otherwise, you must restart the computer for a change to take effect.

**See Also**
**Concepts**
BcastNameQueryCount
NameSrvQueryTimeout

# NameSrvQueryTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 100 - 0xFFFFFFFF | 1500 (1.5 seconds) |

**Description**
Specifies the interval, in milliseconds, between successive name queries to a Windows Internet Name Service (WINS) server for a given name. This entry is installed with its default value during installation of the TCP/IP components.

**Activation Method**
If this interface is configured by Dynamic Host Configuration Protocol (DHCP), then a change in parameters will take effect if you enter the command **ipconfig /renew** at the command line. Otherwise, you must restart the computer for a change to take effect.

**See Also**
**Concepts**
NameSrvQueryCount
BcastQueryTimeout

# NodeType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 \| 2 \| 4 \| 8 | 1 or 8, depending on how WINS is configured |

**Description**
Specifies what methods NetBIOS over TCP/IP (NetBT) uses to register and resolve names. Resolution via Lmhosts and/or DNS, if enabled, use these methods.

| Value | Meaning |
|---|---|
| 1 | B-node. NetBT uses IP broadcast messages to register and resolve IP addresses from NetBIOS names. |
| 2 | P-node. NetBT uses point-to-point communication with a NetBIOS name server (in Windows Server 2003–based networks, this is the WINS server) to register and resolve IP addresses from NetBIOS names. |
| 4 | M-node. NetBT uses a mix of B-node and P-node to register and resolve NetBIOS names. M-node first uses broadcast resolution (B-node), then if necessary, uses a server query (P-node). |
| 8 | H-node. NetBT uses a hybrid of B-node and P-node. An H-node client always directly queries a name server first and then uses broadcasts only if the direct query fails. |

This entry does not exist by default, and if you create it, its setting overrides the configuration specified by DhcpNodeType. If neither this entry nor **DhcpNodeType** exist in the registry, and there are no WINS servers configured on the network, then NetBT uses B-node name resolution. If there is at least one WINS server available, then NetBT uses H-node name resolution.

**Activation Method**
If this interface is configured by DHCP, then enter the command **ipconfig /renew** at the command line for a change in parameters to take effect. Otherwise, you must restart the computer for a change to take effect.

### Note

- For more information on configuring WINS, see the *Networking Collection* of the *Windows Server 2003 Technical Reference* (or see the Networking Collection on the Web at http://www.microsoft.com/reskit).

**See Also**
**Concepts**
DhcpNodeType

# NoNameReleaseOnDemand

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the computer releases its NetBIOS name when it receives a name-release request. If the value of this entry is 0 (the default), then the computer releases its NetBIOS name upon request; if the value of this entry is 1, then the computer does not release its NetBIOS name. Configuring the computer to refuse release of its NetBIOS name can protect against malicious requests for name release.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must restart the computer for changes to take effect.

# RandomAdapter

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies how NetBIOS over TCP/IP (NetBT) chooses the IP address to put in a name query. If the value of this entry is 1, then NetBT randomly chooses an IP address from all of its bound interfaces. If this entry is 0 or is absent from the registry, then NetBT uses the address of the interface on which the query arrived.

This entry is applicable only to multihomed computers. This feature is used by servers with multiple interfaces on the same network for load balancing. This entry does not exist in the registry by default, and you must add it if you want to configure NetBT to randomly choose an IP address from its available interfaces.

**Activation Method**
If this interface is configured by Dynamic Host Configuration Protocol (DHCP), then a change in parameters will take effect if you enter the command **ipconfig /renew** at the command line. Otherwise, you must restart the computer for a change to take effect.

**See Also**
**Concepts**
ConnectOnRequestedInterfaceOnly

# RefreshOpCode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 8 \| 9 | 8 |

**Description**
Specifies the operation code that NetBIOS over TCP/IP (NetBT) uses in name refresh packets.

The Microsoft implementation of NetBT uses a default of 8; however, other implementations, such as Ungermann-Bass, use a default of 9. Interoperation of clients and servers depends on them using the same operation code. If the default operation code (8) is incompatible with other systems on your network, then add this entry and set the value to 9.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
If this interface is configured by DHCP, then a change in parameters will take effect if you enter the command **ipconfig /renew** at the command line. Otherwise, you must restart the computer for a change to take effect.

# SessionKeepAlive

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 60000 to 0xFFFFFFFF | 3600000 milliseconds (1 hour) |

**Description**
Specifies the interval, in milliseconds, between keep-alive transmissions on a session. NetBIOS over TCP/IP (NetBT) sends keep-alive transmissions to verify that idle connections are still active. This prevents NetBT from inadvertently disconnecting active lines. Setting the value of this entry to 0xFFFFFFF disables keep-alive transmissions. This entry is installed with its default value during installation of the TCP/IP components.

**Activation Method**
If this interface is configured by DHCP, then a change in parameters will take effect if you enter the command **ipconfig /renew** at the command line. Otherwise, you must restart the computer for a change to take effect.

**See Also**
**Concepts**
Parameters\KeepAliveInterval

# TryAllIPAddrs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether, as a WINS client, this computer attempts to connect to all addresses it receives in response to a name query, or whether it attempts to connect to only the first address and stops if that fails. The default action is for the computer to try all addresses.

| Value | Meaning |
|---|---|
| 0 | Computer tries to connect only to the first address, and does not try another address in the event of failure. |
| 1 | Computer tries to connect to each address in the response to a name query, in order, from first to last. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must restart the computer for changes to take effect.

### Note

- When a WINS server answers a name query, it returns a list of IP addresses that are sorted into a preference order according to whether any of them are on the same subnet as an interface belonging to the client that sent the name query.

**See Also**
**Concepts**
TryAllNameServers
WinsDownTimeout

# TryAllNameServers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the client continues to query additional WINS servers from its list of WINS servers when a request to one of them for a NetBIOS session fails.

| Value | Meaning |
|---|---|
| 0 | Computer stops trying to query a WINS server after its first attempt fails. |
| 1 | Computer attempts to query all WINS servers in its list. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must restart the computer for changes to take effect.

**See Also**
**Concepts**
TryAllIPAddrs
WinsDownTimeout

# UseDnsOnlyForNameResolutions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether the computer should use only DNS for name resolution; that is, the computer does not attempt to resolve computer names through NetBIOS. The computer still registers and refreshes its NetBIOS names, and NetBIOS sessions are still allowed.

| Value | Meaning |
|---|---|
| 0 | Computer uses all available methods for name resolution. |
| 1 | Computer uses only DNS for name resolution. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**

You must restart the computer for changes to take effect.

**Tip**

To completely disable NetBIOS on an interface, see the regsitry entryNetbiosOptions.

# WinsDownTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1000 – 0xFFFFFFFF | 15,000 (15 seconds) |

**Description**
Specifies the amount of time that NetBIOS over TCP/IP (NetBT) waits before trying to use WINS again after it fails to contact any WINS server. This timeout is primarily for allowing computers that are temporarily disconnected from the network, such as portable devices, to proceed through their startup process without waiting for each name registration or query attempt to time out and fail individually.

**Activation Method**
You must restart the computer for changes to take effect.

### Note

- This entry does not exist in the registry by default. To configure a timeout of a length other than 15 seconds, add this entry by using the registry editor Regedit.exe.

### Tip

For more information about WINS name resolution, see "WINS Technical Reference" in the *Networking Collection* of the *Windows Server 2003 Technical Reference*.

**See Also**
**Concepts**
TryAllNameServers
TryAllIPAddrs

# Interfaces

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# interface-name Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\Interfaces

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# DHCPNameServerList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\Interfaces*interface-name*

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | | |

**Description**
Stores a list of IP addresses of Windows Internet Name Service (WINS) servers, as provided by the Dynamic Host Configuration Protocol (DHCP) service. The IP addresses are stored as dotted decimals (for example, 192.168.0.1) and are separated by spaces. If the entry NameServerList exists and contains one or more valid IP addresses, then this entry is ignored.

**Change Method**
Do not edit this entry. It is created and maintained by the DHCP client service for its own use.

### Note

- This entry replaces the entries **DHCPNameServer** and **DHCPNameServerBackup**, which were used in Windows NT 4.0.

**See Also**
**Concepts**
NameServerList

# DHCPNetbiosOptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\Interfaces*interface-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 \| 2 | 1 |

**Description**
Specifies whether NetBIOS is enabled on the interface named in the registry path, as configured by Dynamic Host Configuration Protocol (DHCP). If the entry NetbiosOptions exists in the registry, then this entry is ignored.

| Value | Meaning |
|---|---|
| 1 | NetBIOS is enabled on the interface named in the registry path of this entry. |
| 2 | NetBIOS is disabled on the interface named in the registry path of this entry. |

**Change Method**
This entry is created and maintained by the DHCP client service, if enabled, and you should not modify it. It is included here only for reference.

**See Also**
**Concepts**
NetbiosOptions

# NameServerList

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\Interfaces*interface-name*

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | | |

**Description**
Stores a list of IP addresses of Windows Internet Name Service (WINS) servers. The addresses are specified with dotted decimals (for example, 192.168.0.1), are separated by spaces, and are ordered according to preference.

**Change Method**
To change the value of this entry, use **Network Connections** in Control Panel. Right-click **Local Area Connection** (or the name of another connection, if appropriate), click **Properties**, click **Internet Protocol (TCP/IP)**, and then click the **Properties** button. On the **Internet Protocol (TCP/IP) Properties** page, click the **Advanced** button, then click the **WINS** tab. This entry is associated with the addresses shown under **WINS addresses, in order of use**. To modify the list, use the **Add**, **Edit**, or **Remove** buttons. Select an address and click the up or down arrow button to the right of the text box to move the address up or down in the list.

**Activation Method**
Changes made to this entry through the **Properties** page for the interface take effect immediately.

## Notes

- If this entry contains one or more valid addresses, then this entry overrides any and all addresses specified in the entry DHCPNameServerList.
- This entry replaces the registry entries **NameServer** and **NameServerBackup** that were used in Windows NT 4.0.

**See Also**
**Concepts**
DHCPNameServerList

# NetbiosOptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\Interfaces*interface-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 \| 2 | 1 |

**Description**
Specifies whether NetBIOS over TCP/IP (NetBT) is enabled on the interface named in the registry path.

| Value | Meaning |
|---|---|
| 1 | NetBT is enabled on the interface named in the registry path of this entry. |
| 2 | NetBT is disabled on the interface named in the registry path of this entry. |

**Change Method**
To change the value of this entry, use **Network Connections** in Control Panel. Right-click **Local Area Connection** (or the name of another connection, if appropriate), click **Properties**, click **Internet Protocol (TCP/IP)**, and then click the **Properties** button. On the **Internet Protocol (TCP/IP) Properties** page, click the **Advanced** button, and click the **WINS** tab. Select one of the following options by clicking on it: **Default**, **Enable NetBIOS over TCP/IP**, or **Disable NetBIOS over TCP/IP**.

The default option specifies that this computer use the NetBIOS setting specified by the Dynamic Host Configuration Protocol (DHCP) server. If this computer uses a static IP address rather than an address provided by a DHCP server, or if the DHCP server does not provide a NetBIOS setting, then NetBT is enabled.

**Activation Method**
Changes made to this entry through the **Properties** page for the interface take effect immediately.

### Note

- This entry overrides the state specified by the entry DHCPNetbiosOptions if they are both in the same registry path (that is, if they are configuring the same network interface).

**See Also**
**Concepts**
DHCPNetbiosOptions

# Services\Netlogon

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The Netlogon subkey stores information for the Net Logon service.

The Net Logon service verifies NTLM logon requests, and it registers, authenticates, and locates domain controllers. Also, to maintain backward compatibility, Net Logon manages replication of the user account database to back up domain controllers running Windows NT 4.0 and earlier.

**Note**

- For Windows Server 2003, Net Logon manages replication only when replication involves servers running Windows NT 4.0 and earlier. Net Logon for Windows Server 2003 does not manage replication between two servers running Windows Server 2003.

# Netlogon\Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon

**Description**
The Parameters subkey stores configuration data specific to the Net Logon service.

The Net Logon service verifies logon requests, and it registers, authenticates, and locates domain controllers. Also, to maintain backward compatibility, Net Logon manages replication of the user account database to backup domain controllers running Windows NT 4.0 and earlier.

**Note**

- For Windows Server 2003, Net Logon manages replication only when replication involves servers running Windows NT 4.0 and earlier. Net Logon for Windows Server 2003 does not manage replication between two servers running Windows Server 2003.

# Parameters\AutoSiteCoverage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the system can add sites to the coverage area of this domain controller.

Domain controllers *cover*, that is, provide services to, the site in which they reside and to other sites listed in the value of the entry**SiteCoverage**. In addition, when the value of **AutoSiteCoverage** is 1, the system can add sites that do not have domain controllers to this domain controller's coverage area.

| Value | Meaning |
|---|---|
| 0 | The system cannot add sites to the coverage area of this domain controller. |
| 1 | The system can add sites to the coverage area of this domain controller. |

The sites added to the domain controller's coverage are stored in memory, and a new list is assembled each time the Net Logon service starts or when Netlogon is notified of the site object changes. While Net Logon runs, it updates this list at an interval specified by the value of the entry **DnsRefreshInterval**.

To determine this extended site coverage, the Net Logon service uses information from Active Directory. It applies the criteria listed below to determine which domain controller can provide the best service to a non-covered site. By default, all domain controllers are candidates for a non-covered site. However, if you use this entry to exclude domain controllers from the selection pool, the system might not be able to select the best domain controller for the site.

The criteria are applied in the order shown, and each subsequent criterion is applied only if the candidate domain controllers are equal on all previous criteria.

- Connection cost: The connection to the domain controller is less costly than connections to other domain controllers in the domain.
- Most domain controllers: The site in which the domain controller resides has more domain controllers than other sites.
- Alphabetical order (tie breaker): The name of the site in which the domain controller resides is earlier in the alphabet than the site names of other domain controllers.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- This entry is used only when it appears in the registry of a domain controller.

**Tip**

- The list of additional sites computed for each domain controller is recorded in Netlogon.log when the value of the entry**DBFlag** entry is at least 0x20. To direct Net Logon to record the process it uses to compute site coverage, set the value of the entry **DBFlag** to 0x2800FFFF.

**See Also**
**Concepts**
Parameters\GcSiteCoverage

# AvoidPdcOnWan Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Prevents the backup domain controller (BDC) from sending new password information to a remote primary domain controller (PDC). Also, if a client password fails to authenticate on the BDC, the BDC does not attempt to authenticate that password on the remote PDC. Instead of Net Logon, the PDC and BDC use Active Directory replication to update password information.

This entry does not affect password sharing between BDCs and PDCs residing on the same site. If the BDC and PDC are on the same site, then the BDC sends password information regardless of the value of this entry.

Setting this value to 1 can reduce wide area network (WAN) traffic between domain controllers at remote locations, but PDCs might not always have the most current password data. As a result, legitimate users might not be authenticated.

| Value | Meaning |
|---|---|
| 0 | BDC sends password information to PDC in different site. |
| 1 | BDC does not send password information to PDC in different site. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# ChangeLogSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x10000–0x400000 (*64 KB–4 MB*) | 0x10000 (*64 KB*) |

**Description**
Specifies the maximum size of the Net Logon transaction log, Netlogon.chg. This limit applies to space allocated for the transaction log in physical memory and on disk in *Systemroot*\Netlogon.chg. The typical size of an entry is 32 bytes, so a 64 KB log holds approximately 2,000 entries.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Notes

- The value of this entry should be the same on all backup domain controllers. This consistency ensures that if a backup domain controller is promoted to a primary domain controller, it has the same value in this entry as the other backup domain controllers.

- Increasing the value of this entry to 4 MB does not degrade system performance. The 4 MB setting ensures that only those segments of the domain's database that changed are replicated.

- For Windows Server 2003, the Net Logon transaction log records only those replication-related changes involving servers running Windows NT 4.0 and earlier. Net Logon for Windows Server 2003 does not manage replication between servers running Windows Server 2003.

# CloseSiteTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x3C–0x418937 (*seconds*) (*1 minute–49.7 days*)49.7 days) | 0x384 (*15 minutes*) |

**Description**
Specifies how often the Net Logon service attempts to find a nearby domain controller. The value of this entry specifies the minimum time that must elapse between attempts. The value is used when Net Logon is using the secure channel to communicate with a domain controller (which it does, for example, when processing an interactive or network logon) in a site that is not nearby. In such a case, if the time elapsed since the last attempt to discover a close domain controller is greater than the time specified by this entry, Net Logon reattempts discovery of a nearby domain controller before sending the request to the remote domain controller.

This value does not apply when Net Logon is finding a domain controller to authenticate an interactive logon that is using pass-through authentication on a secure channel. These transactions do not need time between attempts.

This value helps separate attempts to locate domain controllers when a secure channel is used for other communications, such as pass-through authentication of a network logon.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Caution**

- Use caution when changing the value of this entry. If the value is too high, then the requesting client might be delayed significantly. If the value is too low, then repeated attempts to find the domain controller can create excessive network traffic.

# Parameters\DnsTtl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0x7FFFFFFF (*seconds*) | 0x258 (*600 seconds = 10 minutes*) |

**Description**
Specifies how long a memory cache in the client computer saves Domain Name System (DNS) records that the Net Logon service registers.

| Value | Meaning |
|---|---|
| 0x0 | DNS records registered by Net Logon are not cached. |
| 0x1–0x7FFFFFFF seconds | Specifies the maximum time that DNS records registered by Net Logon are saved in the client cache. |

This value applies to all records, regardless of how frequently or how recently the records might be referenced. This entry is relevant only on domain controllers.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Tip**

- Increasing the value of this entry might improve performance. However, if this value is too large, then the DNS records might not be refreshed on the client often enough.

# DuplicateEventlogTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x418937 (*seconds*) (*Up to 49.7 days*) | 0x3840 (*4 hours*) |

**Description**

Specifies the minimum time that must elapse between the recording of identical events in order for the Net Logon service to write both to the System Log in Event Viewer. If the time between recording is less than the value of this entry, Net Logon regards the second event as a continuation of the first event, and it does not write the second event to the System Log.

This entry minimizes the number of repeated entries in the System Log. To direct Net Logon to record all events, set the value of this entry to 0.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Parameters\ExpectedDialupDelay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–600 (*seconds*) | 0 |

**Description**
Specifies the additional time that the Net Logon service waits for a response from a domain controller. If the domain controller does not respond before the total time allowed elapses, then Net Logon assumes that the domain controller is not available, and it proceeds accordingly.

This entry is designed to allow extra time for a response if the domain controller is connected by a slow dialup link.

This entry extends time limits for the following conditions:

- When locating a domain controller in a trusted domain, such as when authenticating a user's logon credentials.

  Net Logon allows 15 seconds plus the value of this entry for the domain controller to respond. The time begins when the first of three datagrams is sent to the domain controller.

  ```
  Timeout = 15 seconds + ExpectedDialupDelay seconds
  ```

- When an application programming interface (API) call is issued over a secure channel to a domain controller that has already been located.

  Net Logon allows 45 seconds plus the value of this entry for the domain controller to respond.

  ```
  Timeout = 45 seconds + ExpectedDialupDelay seconds
  ```

To give remote domain controllers more time to respond, you can add this entry to the registry and increase its value.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Parameters\GcSiteCoverage

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | *Site names* | (None) |

**Description**

Specifies a list of sites in which this domain controller registers itself as a global catalog resource. These sites are in addition to those that Net Logon assumes to be closely related to the site based on their physical proximity, network topology, or bandwidth. As such, this list does not include the site in which the domain controller resides.

You can add this entry to the registry to add sites to the list. This list lets you associate sites that might be logically connected, but physically disparate.

The domain controller considers itself to be closely related to the sites in this list. Therefore, it replicates to the site and, when pinged by a client in the site, indicates that it is close by turning on the *closest site* bit.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- This entry is used only when it appears in the registry of a domain controller.

# Parameters\MaximumLogFileSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*bytes*) | 0x1312D00 (*20,000,000 bytes*) |

**Description**
Specifies the maximum size of the Net Logon debugging file, Netlogon.log. When the file reaches this size, any existing Netlogon.bak file is deleted, the current Netlogon.log file is renamed to Netlogon.bak, and a new Netlogon.log file is created. Logging is not interrupted.

The Netlogon.log file is created when the value of the **DBFlag** entry is greater than 0.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# MaximumMailslotMessages

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF (*messages*) | 0x1F4 (*500 messages*) |

**Description**

Specifies the maximum number of mailslot messages queued to the Net Logon service.

The Net Logon service is designed to process incoming mailslot messages immediately, but messages can accumulate when the server is processing many requests. Each mailslot message consumes about 1,500 bytes in the nonpaged memory pool until it is processed.

Reducing the value of this entry limits the amount of the nonpaged memory pool that Net Logon consumes. However, if this value is set too low, then Net Logon might reject important incoming mailslot messages.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# MaximumPasswordAge

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–1,000,000 (*days*) | 30 |

**Description**

Specifies how often the system changes the computer account password of the local computer. This entry is used only when the system is configured to change the computer password automatically at regular intervals, that is, when the value of the **DisablePasswordChange** entry is 0.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Pulse

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 60–172,800 (*seconds, up to 48 hours*) | 300 (*5 minutes*) |

**Description**
Specifies how often a primary domain controller (PDC) notifies backup domain controllers (BDCs) that changes have been made to the Security Accounts Manager (SAM) and Local Security Authority (LSA) databases. When the time specified by the value of this entry expires, the PDC sends notices only to the BDCs whose databases are not current.

The value of this entry does not affect database changes that trigger immediate updates. These include:

- Changing an account lockout policy.
- Changing a domain password policy.
- Changing a computer account password.
- Locking out an account.
- Changing the password for a trust account.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- In Windows Server 2003, the Net Logon service manages replication only if the replication involves servers running Windows NT 4.0 or earlier. Net Logon for Windows Server 2003 does not manage replication between two servers running Windows Server 2003.

**See Also**
**Concepts**
PulseConcurrency
PulseMaximum
PulseTimeout1
PulseTimeout2
Randomize

# PulseConcurrency

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–500 (*pulses*) | 10 |

**Description**
Defines the maximum number of simultaneous pulses that the primary domain controller (PDC) sends to backup domain controllers (BDCs).

This entry limits the workload that sending and responding to pulses places on the PDC. The PDC must be able to support the number of concurrent replication remote procedure calls stored in the value of this entry.

Increasing the value of this entry increases the load on the PDC. Decreasing the value increases the time that a domain with a large number of BDCs requires to replicate a change in the Security Accounts Manager (SAM) or Local Security Authority (LSA) databases to all of the backup controllers.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- In Windows Server 2003, the Net Logon service manages replication only if the replication involves servers running Windows NT 4.0 or earlier. Net Logon for Windows Server 2003 does not manage replication between two servers running Windows Server 2003.

**See Also**
**Concepts**
Pulse
PulseMaximum
PulseTimeout1
PulseTimeout2
Randomize

# PulseMaximum

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 60–172,800 (*seconds, up to 48 hours*) | 7200 (*2 hours*) |

**Description**
Defines the maximum time interval between pulses. If the time specified in the value of this entry expires, then the primary domain controller (PDC) sends at least one pulse, even if all backup domain controller (BDC) databases are current.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- In Windows Server 2003, the Net Logon service manages replication only if the replication involves servers running Windows NT 4.0 or earlier. Net Logon for Windows Server 2003 does not manage replication between two servers running Windows Server 2003.

**See Also**
**Concepts**
Pulse
PulseConcurrency
PulseTimeout1
PulseTimeout2
Randomize

# PulseTimeout1

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–120 (*seconds*) | 10 |

**Description**
Defines how long the primary domain controller (PDC) waits for a response from a backup domain controller (BDC) before it considers the BDC to be unresponsive. An unresponsive BDC is not counted against the limit specified in the PulseConcurrency entry, allowing the PDC to send a pulse to another BDC in the domain.

If this value is too high, then a domain with a large number of unresponsive BDCs takes a long time to complete a partial replication. But if this number is too low, then a slow BDC might mistakenly be treated as unresponsive. When the BDC finally does respond, it will partially replicate from the PDC, increasing the load on the PDC.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- In Windows Server 2003, the Net Logon service manages replication only if the replication involves servers running Windows NT 4.0 or earlier. Net Logon for Windows Server 2003 does not manage replication between two servers running Windows Server 2003.

**See Also**
**Concepts**
Pulse
PulseConcurrency
PulseMaximum
PulseTimeout2
Randomize

# PulseTimeout2

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 60–3,600 (*seconds*) | 300 (*5 minutes*) |

**Description**
Defines how long a primary domain controller (PDC) waits for a backup domain controller (BDC) to complete each step in the replication process. If the value of this entry is exceeded, then the PDC regards the BDC as unresponsive and can send a pulse to another BDC in the domain.

This entry applies only to PDCs, and it is used only if a BDC cannot retrieve all the changes to the Security Accounts Manager (SAM) or Local Security Authority (LSA) databases in a single remote procedure call (RPC).

To be considered responsive, the BDC must respond to a pulse and must continue reporting its progress on the replication. If the interval between progress reports exceeds the value of this entry, then the BDC is considered to be unresponsive. An unresponsive BDC is not counted against the limit specified by the PulseConcurrency entry, allowing the PDC to send a pulse to another BDC in the domain.

If the value of this entry is set too high, then a slow BDC, or one whose replication rate is too slow, consumes one of the PulseConcurrency slots. But if the value of this entry is too low, then the PDC's workload increases because of the large number of BDCs doing a partial replication.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- In Windows Server 2003, the Net Logon service manages replication only if the replication involves servers running Windows NT 4.0 or earlier. Net Logon for Windows Server 2003 does not manage replication between two servers running Windows Server 2003.

**See Also**
**Concepts**
Pulse
PulseConcurrency
PulseMaximum
PulseTimeout1
Randomize

# Randomize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–120 (*seconds*) | 1 |

**Description**
Specifies how long a backup domain controller (BDC) waits after receiving a pulse before it responds to the primary domain controller (PDC).

The value of this entry must be smaller than the value of the PulseTimeout1 entry, otherwise the PDC will consider the BDC to be unresponsive.

When calculating an optimal value for this entry, remember that the time required to replicate a change to the Security Accounts Manager (SAM) or Local Security Authority (LSA) databases to all the BDCs in a domain will be greater than:

`[(Randomize/2) * NumberOfBDCsInDomain] /` PulseConcurrency

If this entry does not appear in the registry, then Net Logon determines optimal values for the server based on the domain controller's workload.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- In Windows Server 2003, the Net Logon service manages replication only if the replication involves servers running Windows NT 4.0 or earlier. Net Logon for Windows Server 2003 does not manage replication between two servers running Windows Server 2003.

**See Also**
**Concepts**
Pulse
PulseConcurrency
PulseMaximum
PulseTimeout1
PulseTimeout2

# RefusePasswordChange

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Prevents a domain controller from accepting requests from workstations to change their computer account passwords.

By default, Windows XP Professional systems automatically request a change to their computer account passwords at regular intervals (as specified in the value of the MaximumPasswordAge entry). To direct the system to reject these changes each time they arrive, set the value of **RefusePasswordChange** to 1.

| Value | Meaning |
|---|---|
| 0 | Accept changes to workstation computer account passwords. |
| 1 | Reject changes to workstation computer account passwords. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Tip**

- To prevent a workstation from sending a password change to the domain controller, set the value of the **DisablePasswordChange** entry to 1.

**Caution**

- Disabling automatic password changes can make the system more vulnerable to malicious access. Frequent password changes can be a significant safeguard for your system.

**See Also**
**Concepts**
MaximumPasswordAge

# RegisterDnsARecords

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the domain controller registers Domain Name System (DNS) A (address) records for the domain. If the domain controller is a global catalog resource, then this entry also determines whether the domain controller registers DNS A records for the global catalog.

| Value | Meaning |
|---|---|
| 0 | Does not register DNS A records. Lightweight Directory Access Protocol (LDAP) implementations that do not support SRV records will not be able to locate the LDAP server on this domain controller. |
| 1 | Registers DNS A records. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Notes**

- This entry is used only when it appears in the registry of a domain controller.
- You might consider setting this value to 0 if DNS fails to complete its dynamic updates because it cannot update A records. DNS stops updating when an update attempt fails.

# RequireSignOrSeal

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether the system requires that all secure channel communications be either signed or sealed.

The security specifications for secure channel traffic are determined jointly by the value of this entry and the values of the RequireStrongKey, SignSecureChanneland SealSecureChannel entries.

| Value | Meaning |
|---|---|
| 0 | Channel traffic need not be signed or sealed. Instead, the security of channel traffic is negotiated with the domain controller. System preferences in the negotiation of traffic security are determined by the values of SignSecureChanneland SealSecureChannel. |
| 1 | Outgoing traffic on a secure channel must be either signed or sealed. If the domain controller on the other side of the channel does not support signing or sealing, the system refuses to establish a channel. If the value of SealSecureChannel is 1, traffic must be encrypted; otherwise, it must be signed. |

**Notes**

- The value of this entry should be set to 1 only when all of the trusted domains support signed and sealed communications.
- Windows Server 2003 adds this entry to the registry when you install the system for the first time or when you change the default value. This entry might not appear in the registry, for example, if you upgrade to Windows Server 2003 from a system running Windows NT 4.0 that does not define this entry. If this entry does not appear in the registry, the system behaves as though the value is 0

**See Also**

**Concepts**

SignSecureChannel

SealSecureChannel

# RequireStrongKey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the system requires that all secure channel keys be computed using a strong key. If it does, then the trusted domain controller on the other side of the channel must be able to compute strong keys. For public, non-exportable keys, a strong key is one that is 128 bytes or longer.

All secure channels have keys. The keys are used for authentication, signing, or encryption, depending on the capability and requirements of the systems.

| Value | Meaning |
|---|---|
| 0 | The system does not require that the trusted domain controller be able to compute a strong key. |
| 1 | The system requires that the trusted domain controller be able to compute a strong key. If the domain controller on the other side of the channel does not support strong key encryption, this system refuses to establish a channel. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- This entry should be set to 1 only when all of the trusted domains are able to compute strong keys.

**See Also**
**Concepts**
SignSecureChannel
SealSecureChannel

# Parameters\ScavengeInterval Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 60–172,800 (*seconds*) (*1 minute–2 days*) | 900 (*15 minutes*) |

**Description**

Specifies how often the Net Logon service performs scavenging activities, such as the following:

- Checks for password expirations on secure channels.
- Checks for idle secure channels.
- Discovers domain controllers for all trusted domains in which a domain controller has not yet been discovered. (Domain controllers only.)
- Attempts to add the domain's NetBIOS <1B> name. (Primary domain controllers only.)
- Cleans up sundry system resources that are no longer needed.

The default value is optimal for most systems. However, in some cases, such as when a trusted domain is connected by an expensive connection, you might add this entry to increase the time between operations.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# SealSecureChannel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether outgoing secure channel traffic is encrypted (sealed). This entry is used when negotiating the conditions of a secure channel with a domain controller. If the value of this entry is 1, and the domain controller supports this functionality, the secure channel traffic is signed.

Channel traffic security is determined jointly by the value of this entry and the values of the registry entries RequireStrongKey, RequireSignOrSeal, and SignSecureChannel.

Because encryption is more secure than signing, when the value of this entry is 1, it takes precedence over the value of SignSecureChannel.

| Value | Meaning |
|---|---|
| 0 | Outgoing traffic on a secure channel should not be encrypted. |
| 1 | Outgoing traffic on a secure channel should be encrypted if the domain controller supports this functionality. |

**Note**

- Windows Server 2003 and Windows 2000 add this entry to the registry when you install the system for the first time or when you change the default value. If you upgrade from Windows NT 4.0 or earlier, the entry does not appear in the registry, but it is still in effect on your system.

RequireStrongKey

**See Also**
**Concepts**
RequireSignOrSeal
SignSecureChannel
RequireStrongKey

# SignSecureChannel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether outgoing secure channel traffic is signed. This entry is used when negotiating the conditions of a secure channel with a domain controller. If this value of this entry is 1 and the domain controller supports signing, the secure channel traffic is signed.

Channel traffic security is determined jointly by the value of this entry and the values of the entries RequireStrongKey, RequireSignOrSeal, and SealSecureChannel. Because encryption is more secure than signing, the system behaves as if the value of this entry is 1, even if it is zero, when the value of **sealsecurechannel** is 1.

| Value | Meaning |
|---|---|
| 0 | Outgoing traffic on a secure channel should not be signed. |
| 1 | Outgoing traffic on a secure channel should be signed if the domain controller supports signing. |

**Note**

- Windows Server 2003 and Windows 2000 add this entry to the registry when you install the system for the first time or when you change the default value. If you upgrade from Windows NT 4.0 or earlier, the entry does not appear in the registry, but it is still in effect on your system.

**See Also**
**Concepts**
RequireSignOrSeal
SealSecureChannel
RequireStrongKey

# NTDS

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
Stores configuration data for Active Directory.

In particular, Active Directory stores domain controller information in this subkey that it uses to restart Active Directory independently and to configure the Active Directory database.

# NTDS\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NTDS

**Description**
The Parameters subkey stores configuration data for Active Directory. In particular, Active Directory saves the location of log files and replication information in this subkey.

# Repl topology update delay (secs)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x14–0xFFFFFFFF (*seconds*) | 0x12C (*300 seconds = 5 minutes*) |

**Description**
Specifies the amount of time after Active Directory starts that the Knowledge Consistency Checker (KCC) first evaluates the replication topology.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
Repl topology update period (secs)

# Repl topology update period (secs)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x14–0xFFFFFFFF (*seconds*) | 0x384 (*900 seconds = 15 minutes*) |

**Description**

Specifies how often the Knowledge Consistency Checker (KCC) reevaluates the replication topology. KCC repeatedly evaluates the topology at the specified rate whenever Active Directory is running.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Replicator async inter site packet size (bytes)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x400 - 0xFFFFFFFF | (sized dynamically) |

**Description**

Stores the maximum size of a packet that transports replication data between sites that are not directory partitions. The default value is sized dynamically, based on the amount of memory. This value is sent on a replication request to the source computer. The source computer uses the value it receives in the replication request or its own setting for this value, whichever is smaller.

For replication between sites that are not directory partitions, Active Directory uses the Simple Mail Transfer Protocol (SMTP). This replication is asynchronous.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Notes

- If the value of this entry is not set, then the system limits the packet size to 1/100th of the size of RAM, with a minimum of 17nbsp;MB and a maximum of 10 MB.

- For the system to function correctly, the value of this entry must be at least 0x400 (1 MB), and it must not be unreasonably high.

# Replicator inter site packet size (bytes)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x400 - 0xFFFFFFFF | (sized dynamically) |

**Description**
Stores the maximum size of a packet that transports replication data between sites that are directory partitions. The default value is sized dynamically, based on the amount of memory. This value is sent on a replication request to the source computer. The source computer uses the value it receives in the replication request or its own setting for this value, whichever is smaller.

For replication between sites that are directory partitions, Active Directory uses the Remote Procedure Call (RPC) transport protocol, with compression. This replication is synchronous.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
You can change this entry only by editing the registry.

### Notes

- If the value of this entry is not set, the system limits the packet size to 1/100th of the size of RAM, with a minimum of 1 MB and a maximum of 10 MB.

- For the system to function correctly, the value of this entry must be at least 0x400 (1 MB), and it must not be unreasonably high.

# Replicator intra site packet size (bytes)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x400 - 0xFFFFFFFF | (sized dynamically) |

**Description**

Stores the maximum size of a packet that transports replication data within a site. The default value is sized dynamically, based on the amount of memory. This value is sent on a replication request to the source computer. The source computer uses the value it receives in the replication request, or its own setting for this value, whichever is smaller.

For replication within a site, Active Directory uses the Remote Procedure Call (RPC) transport protocol, without compression. This replication is synchronous.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Notes

- If the value of this entry is not set, then the system limits the packet size to 1/100th of the size of RAM, with a minimum of 1 MB and a maximum of 10 MB.

- For the system to function correctly, the value of this entry must be at least 0x400 (1 MB), and it must not be unreasonably high.

# TCP/IP Port

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–65,535 (*port number*) | (None) |

**Description**

Specifies the port used for Active Directory replication on this server. The remote procedure call (RPC) locator always returns this port number when queried for the Active Directory replication port for this server.

By default, Active Directory uses a dynamically assigned port to transport replication updates within sites. However, you can use this entry to create a fixed port assignment. This setting is designed for routed environments and those using port filtering.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# NtFrs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
Stores configuration data for the File Replication service (FRS). FRS automatically updates and synchronizes the contents of replicated files.

# NtFrs\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NtFrs

**Description**
The Parameters subkey stores configuration data for the File Replication service (FRS). FRS automatically updates and synchronizes the contents of folders managed by Distributed File System (DFS).

The data in this subkey typically specifies settings that are configurable by the administrator; for example, the severity setting for events logged by FRS.

# Debug Disable

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Ntfrs\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Disables the File Replication service (FRS) trace log.

By default, FRS records its actions in trace log files. These log files are named Ntfrs_000*x* and are located in the *Systemroot*\debug directory. These files are typically used to investigate FRS replication problems. You can disable the logs to reduce disk use.

| Value | Meaning |
|---|---|
| 0 | Enables the trace log. |
| 1 | Disables the trace log. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
Debug Log File
Debug Log Files
Debug Log Severity
Debug Maximum Log Messages

# Debug Log File

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Ntfrs\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Path* | *Systemroot*\debug |

**Description**
Specifies where the File Replication service (FRS) stores its trace log files.

If you change the value of this entry, then FRS creates new trace log files in the directory specified, but it does not move existing trace log files to the new location. Also, FRS does not delete or maintain log files in the original location.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make changes to the value of this entry effective, you must stop and restart FRS.

**Tip**

- If you have more than one hard drive, you can improve disk performance by storing the FRS trace log files on a different hard drive than the one that stores the FRS database and replica tree.

**See Also**
**Concepts**
Debug Disable
Debug Log Files
Debug Log Severity
Debug Maximum Log Messages

# Debug Log Files

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Ntfrs\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 5 |

**Description**
Specifies how many File Replication service (FRS) trace log files can be active simultaneously.

FRS creates a set of sequential log files named Ntfrs_000*x*.log. The files are numbered, beginning with 0001, so that the log with the highest number contains the most recently logged entries. The log files are retained until the number of files reaches the value of this entry, then FRS deletes Ntfrs_0001 and renumbers all remaining logs accordingly.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
Debug Log File
Debug Disable
Debug Log Severity
Debug Maximum Log Messages

# Debug Log Severity

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Ntfrs\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–5 | 2 |

**Description**
Specifies the level of detail that the File Replication service (FRS) records in its trace log files (Ntfrs_000*n*.log).

FRS assigns severity levels to different log print statements in the FRS code. The value of this entry determines which severity levels are written to the log. The higher the value of this entry, the more detail the log records.

If the value of this entry is 0, then only the most severe events are recorded in the log. If the value of this entry is 5, then all events are recorded in the log. The default value is set to 2 so that tracking records are displayed, as well as warning and error messages.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
Debug Log File
Debug Disable
Debug Log Files
Debug Maximum Log Messages

# Debug Maximum Log Messages

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Ntfrs\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x4E20 (20,000) |

**Description**
Specifies how many entries can be stored in each File Replication service (FRS) trace log file (Ntfrs_000*x*.log).

When the number of event entries reaches the value of **Debug Maximum Log Messages,** the system closes the current log file and starts a new one. The log files are numbered sequentially, so when the FRS closes Ntfrs_0001.log, it begins Ntfrs_0002.log.

The default value of 10,000 entries (0x2710) uses approximately 1 MB of disk space.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
Debug Log File
Debug Disable
Debug Log Files
Debug Log Severity

# Staging Space Limit in KB

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Ntfrs\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*KB*) | 0xA1220 (*660,000 KB = 660 MB*) |

**Description**
Specifies the maximum amount of disk space that the File Replication service (FRS) can use to store staging files.

FRS uses staging files to encapsulate the data streams and attributes associated with each replicated file or directory object. The staging file is transferred among replica members, which use it to replicate the object.

When the space consumed by staging files reaches the value of this entry, inbound replication is suspended. FRS tries to release space by replicating staging files to all outbound partners.

Typically, staging files reach this limit because FRS cannot offload files to an outbound partner that connects infrequently.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You can add this entry or change its value while FRS is running. You do not need to restart FRS to make the change effective.

# NWCWorkstation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
Stores configuration data for Client Service for NetWare. This service enables a Windows XP Professional workstation to gain access to file and print services on Novell NetWare servers, even if it is not running NetWare client software.

### Note

- This subkey appears in the registry when you install Client Service for NetWare.

# NWCWorkstation\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation

**Description**
Stores configuration data specific to Client Service for NetWare. This service enables a Windows XP Professional workstation to gain access to file and print services on Novell NetWare servers, even if it is not running NetWare client software.

### Note

- This subkey appears in the registry when you install Client Service for NetWare.

# DefaultLocation

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Server-name* or **Tree\Context* | (None) |

**Description**

Specifies the Netware 3.*x* server or NDS tree/context for all new users. This entry applies to users the first time they log on.

Once a user has logged on, administrators can use Client Service for Netware in Control Panel to change the preferred server, or tree and context.

The **DefaultLocation** and DefaultScriptOptions entries let you set default values for the preferred server and logon script client options. When both of these entries appear in the registry with valid values, the system supresses the **Select NetWare Logon** dialog box that requests this information from the user. Otherwise, the dialog box appears each time a user logs on until they select a preferred server.

**Note**

- This entry is used only when the DefaultScriptOptions entry appears in the registry with a valid value. Otherwise, the entry is ignored and the system requires a user to select a server from the **Select NetWare Logon** dialog box.

# DefaultScriptOptions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0x3 \| 0x800 | 0x0 |

**Description**

Specifies whether logon scripts are enabled or disabled by default. This entry applies to users the first time they log on.

Once a user has logged on, administrators can use Client Service for NetWare in Control Panel to enable or disable logon scripts.

| Value | Meaning |
|---|---|
| 0x0 | Logon scripts are disabled. |
| 0x1 | Logon scripts are enabled only on bindery servers. |
| 0x2 | Logon scripts are enabled only on Novell Directory Services (NDS) directories. |
| 0x3 | Logon scripts are enabled on both bindery servers and NDS directories. |
| 0x800 | Logon scripts enabled and are in debugging mode. |

This entry is a bitmask. Bit 1 enables logon scripts on bindery servers. Bit 2 (10 binary) enables logon scripts on NDS servers. To enable logon scripts on a type of server, set the bit for that server type to 1. To turn debugging mode on, set bit 12 (100000000000, or 0x800) to 1.

The DefaultScriptOptions and **DefaultScriptOptions** entries let you set default values for the preferred server and logon script client options. When both of these entries appear in the registry with valid values, the system supresses the **Select NetWare Logon** dialog box that requests this information from the user. Otherwise, the dialog box appears each time a user logs on, until the user selects a preferred server.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Logon

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters

**Description**
Stores data about the current user for the Client Service for NetWare, which uses the data internally.

This subkey contains subkeys that represent users who have logged on to the system. These subkeys are named for the (security identifiers (SIDs)) of the users. The data in each SID Subkey subkey applies only to the user that the SID represents.

### Note

- This subkey appears in the registry when you install Client Service for NetWare.

### Caution

- Do not change the content of this subkey. If you do, Client Service for NetWare might not operate properly, and users might not be able to gain access to resources on NetWare servers.

**See Also**
**Concepts**
SID Subkey

# SID Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters\Logon

**Description**
The *SID* subkey represents particular users.

The *SID* subkey is named for the security identifier (SID) of the user. It stores data about a specific user for the Client Service for NetWare, which uses the data internally. The data in each *SID* subkey applies only to the user that the SID represents.

**Note**

- The *SID* subkey is a placeholder representing the subkeys named for the SIDs in the subkey HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters\Logon. "SID" does not actually appear in the registry. The actual SID subkeys in your registry are specific to your system and its users.

- This subkey appears in the registry when you install Client Service for NetWare.

**Caution**

- Do not change the contents of the Logon subkey. If you do, Client Service for NetWare might not operate properly, and users might not be able to gain access to resources on NetWare servers.

# Option

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters

**Description**
Stores the value of optional settings for users of this computer for the Client Service for NetWare, which uses the data internally.

This subkey contains subkeys that represent users who have logged on to the system. These subkeys are named for the (security identifiers (SIDs)) of the users. The data in each SID Subkey subkey applies only to the user that the SID represents.

# Option\SID

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters\Option

**Description**
The *SID* subkeys represent particular users.

The *SID* subkey is named for the security identifier (SID) of the user. It stores data about a specific user for the Client Service for NetWare, which uses the data internally. The data in each SID subkey applies only to the user that the SID represents.

**Note**

- The *SID* subkey is a variable representing the subkeys names for the SIDs in the HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters\Option subkey. It does not actually appear in the registry. This variable subkey displays the subkeys and entries that are common to many of the subkeys in Logon. The actual SID subkeys in your registry are specific to your system and its users.

- This subkey appears in the registry when you install Client Service for NetWare.

# LogonScript

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters\Option*SID*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–3 | (See below.) |

**Description**
Specifies whether a logon script runs if the user represented by the security identifier (SID) connects directly to a NetWare server.

This entry stores the values of the **Run Login Script** check box in CSNW in Control Panel and the **Run Login Script** check box on the **Select NetWare Logon** dialog box. The **Select NetWare Logon** dialog box appears the first time a user logs on after installing Client Service for NetWare.

| Value | Meaning |
|---|---|
| 0 | Do not run NetWare login scripts. |
| 1 | Run login scripts when connecting to NetWare bindery servers. |
| 2 | Run login scripts when connecting to Novell Directory Services (NDS) servers. |
| 3 | Run login scripts when connecting to any NetWare servers. |

**Change Method**
To change the value of this entry, select or clear the **Run Login Script** check box in CSNW in Control Panel.

### Note

- This subkey appears in the registry when you install Client Service for NetWare.

# NDSPreferredServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters\Option*SID*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Server-name* | (None) |

**Description**
Specifies the Novell Directory Services (NDS) server that the system tries first when it logs the user on to an NDS tree.

If this entry is not in the registry or the specified server is not available, then the system logs the user on the first server in the user's tree and context that responds to the logon request. (The tree and context are stored in the value of the PreferredServer entry.)

Specifying a particular server assures that, when users log on through an NDS tree containing more than one NetWare server, they are authenticated by the same NetWare server that provides the access to NetWare resources. This provides administrators with better control over the use of NetWare connection licenses.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Note

- This entry is used only when the client connects to a network that uses Novell Directory Services.

**See Also**
**Concepts**
PreferredServer

# PreferredServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters\Option*SID*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | ***** *Tree* **** *Context* or *Server-name* | (None) |

**Description**
Specifies the NetWare server or Novell Directory Services (NDS) tree and context to which this user logs on. If the specified server is not available, then the system notifies the user and lets the user select from a list of other available servers. If the value of this entry is blank, then the system does not log the user on to a NetWare server.

An asterisk (*) preceding the string indicates that the value is an NDS tree and context. Otherwise, the value indicates a NetWare 3.*x* server.

This entry stores the values of the **Select Preferred Server**, **Tree**, and **Context** text boxes in CSNW in Control Panel and in the **Select NetWare Logon** dialog box. The **Select NetWare Logon** dialog box appears the first time a user logs on after installing Client Service for NetWare.

Client Services for NetWare adds this entry to the registry with the value that the user enters in the dialog box options. If the user leaves the dialog box options blank, then the value of this entry is blank. If this entry is not in the registry, then Client Service for NetWare assumes that the user has logged on before.

Users can change the value of this entry by using CSNW in Control Panel. If this value is blank when the user logs on again, then the system displays a dialog box offering users a list of available servers. If the user selects a server, then the name of that server is written to the value of this entry. If the user does not select a server, then the user cannot log on to a NetWare server.

**Change Method**
Do not edit this entry in the registry. To change the value of this entry, use CSNW in Control Panel.

### Note

- When this entry stores an NDS tree and context, then by default, the system logs the user on to the first server in that tree and context that responds to a logon request. To indicate a preferred server within the specified NDS tree and context, use the NDSPreferredServer entry.

**See Also**
**Concepts**
NDSPreferredServer

# PrintOption

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\NWCWorkstation\Parameters\Option*SID*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x08 \| 0x10 \| 0x18 \| 0x80 \| 0x88 \| 0x90 \| 0x98 | 0x98 |

**Description**
Specifies options set by the user for printing to NetWare resources.

The value of this entry is a 4-bit value in which three of the bits represent printing options. The remaining bit is not used. A printing option is enabled when its corresponding bit is 1; the option is disabled when its corresponding bit is 0. The follow table shows the significant bits and their meaning.

| Bit | Meaning |
|---|---|
| 0x08 | Suppress form feed. |
| 0x10 | Send print notification message. |
| 0x80 | Print in banner style. |

**Change Method**
To change the value of this entry, use CSNW in Control Panel.

### Note

- This subkey appears in the registry when you install Client Service for NetWare.

**See Also**
**Concepts**
NetWare or Compatible Network

# PerfDisk

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The PerfDisk subkey stores configuration data for Perfdisk.dll, the Windows Server 2003 disk performance driver. Perfdisk.dll collects performance-related data about logical and physical disk drives for display by System Monitor and other display tools.

**Caution**

- Do not modify or delete the contents of this subkey. Doing so might intefere with the collection of data.

# PerfDisk\Performance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\PerfDisk

**Description**
The Performance subkey stores configuration data for Perfdisk.dll, the Windows Server 2003 disk performance driver. Perfdisk.dll collects performance-related data about logical and physical disk drives for display by System Monitor and other display tools.

**Activation Method**
To make changes to entries in this subkey effective, restart the Remote Registry Service or restart the system.

### Tip

- If particular performance counters do not appear in your monitoring tool, check the Application Log in Event Viewer for events recorded by Perflib. When Perflib detects disabling errors in a counter dynamic-link library (DLL), it records an event in the log, and it disables the counters by setting the value of the Disable Performance Counters entry to 1.

# NT4 Instance Names

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\PerfDisk\Performance

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether instances of the PhysicalDisk performance object displayed in performance tools, such as System Monitor, include the names of logical drives residing on the physical drive.

| Value | Meaning |
|---|---|
| 0 | Instances of the PhysicalDisk performance object display the drive number and logical drive letters. |
| 1 | Instances of the PhysicalDisk performance object display only the drive number. (Same as Windows NT 4.0 and earlier.) |

For example, suppose a physical drive, Disk 0, includes three logical disks, C, D, and E. If the value of this entry is 0, then the physical disk instance is represented as 0: C, D, E. If the value of this entry is 1, then the same instance is represented as 0.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
To make changes to this entry effective, restart the Remote Registry Service or restart Windows.

**Tip**

- To test this entry on your computer, in System Monitor, click the **+** button and, in the **Performance object** box, click **PhysicalDisk**. View the names of instances in the **Select instances from list** box. Then, add this entry to the registry, set its value to 1, and view the names again.

# PerfProc

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# PerfProc\Performance

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\PerfProc

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# ThreadNameFormat

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\PerfProc\Performance

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 \| 2 | (None) |

**Description**
Specifies whether System Monitor or the Performance Logs and Alerts service displays and logs thread counter instances with the ID of the associated thread appended.

If the value of this entry is 1, or if this entry is absent from the registry, then Performance Logs and Alerts appends a unique integer to each instance of a thread name to distinguish the multiple instances from each other. If the value of this entry is 2, then Performance Logs and Alerts appends the thread ID for each instance of the thread name.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
The value of this entry is read each time the ReqQueryValueEx API is called. For Performance Logs and Alerts, you must stop and restart the service for changes to this entry to be effective.

# Rasl2tp

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The Rasl2tp subkey stores configuration data for the Layer Two Tunneling Protocol (L2TP) support in Routing and Remote Access.

L2TP is an industry standard tunneling protocol. Unlike Point-to-Point Tunneling Protocol (PPTP), which is based on Transmission Control Protocol (TCP), L2TP is based on User Datagram Protocol (UDP).

**See Also**
**Concepts**
ProhibitIpSec

# RasMan

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The RasMan subkey stores configuration data for Routing and Remote Access. The subkeys in the RasMan subkey vary, according to which services are installed.

# RasMan\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan

**Description**
The Parameters subkey stores configuration data for Routing and Remote Access.

# DisableSavePassword

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the passwords associated with Network Connections phone book entries are saved. If the user has many phone book entries, accumulated saved passwords can cause a slight delay after the user's credentials are typed in the **Connecting To** dialog box.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# IpOutHighWatermark

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x5 |

**Description**
Specifies the maximum number of unused IP outgoing dial-up adapters that Routing and Remote Access can maintain. If the number of these unused adapters exceeds the value of this entry, then Routing and Remote Access deletes the excess adapters.

Routing and Remote Access dynamically creates IP outgoing dial-up adapters as needed, and it deletes the adapters when they are no longer needed. Routing and Remote Access creates reserve adapters in advance to prevent call delays. However, because adapters consume system resources, you should reserve only as many adapters as you will use.

If the value of this entry is too high, then resources are wasted maintaining excess adapters. But if the value of this entry is too low, then users will experience perceptible delays while adapters are being created.

**See Also**
**Concepts**
IpOutLowWatermark

# IpOutLowWatermark

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x1 |

**Description**
Specifies the minimum number of unused IP outgoing dial-up adapters that Routing and Remote Access can maintain. If the number of these unused adapters falls below the value specified in this entry, then Routing and Remote Access creates more adapters.

Routing and Remote Access dynamically creates IP outgoing dial-up adapters as needed, and it deletes the adapters when they are no longer needed. Routing and Remote Access creates reserve adapters in advance to prevent call delays. However, because adapters consume system resources, you should reserve only as many adapters as you will use.

If the value of this entry is too high, then resources are wasted maintaining excess adapters. But if the value of this entry is too low, then users will experience perceptible delays while adapters are being created.

**See Also**
**Concepts**
IpOutHighWatermark

# Parameters\Logging

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

**Caution**

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# ProhibitIpSec

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | |

**Description**

Prevents Routing and Remote Access from conducting IP security (IPSec) negotiation on calls using Layer Two Tunneling Protocol (L2TP).

By default, Routing and Remote Access combines L2TP and IPSec to provide secure remote communication. If the value of this entry is 1, then L2TP calls are not secure.

| Value | Meaning |
|---|---|
| 0 | Routing and Remote Access conducts IPSec negotiation on L2TP calls. |
| 1 | Routing and Remote Access does not conduct IPSec negotiation on L2TP calls. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# PPP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan

**Description**
The PPP subkey contains entries for the Point-to-Point Protocol (PPP) service for Routing and Remote Access.

# BapListenTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*seconds*) | 0x2D (*45 seconds*) |

**Description**

Specifies how long the client waits for a response to a Bandwidth Allocation Protocol (BAP) callback request. While waiting, the port that the client is using is not available for any other use. When the value of this entry expires, the port is released.

This entry is used only when client computers are configured to use BAP. BAP dynamically adds and drops links as they are needed to support the network traffic. Clients send callback requests to ask the server to connect an additional link.

### Notes

- The default value is optimal for most installations. However, you might consider increasing this value if the client is not answering the server's callback. This might indicate that the client stops waiting for a response before the server calls. Increase this value by approximately 15 seconds and test the callback again.

- The remote access client implements any value you enter in this entry, even if the value is impractical. Be sure to test the effect of changes to these values before you go online.

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# DefaultCallbackDelay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–255 (*seconds*) | 12 |

**Description**
Specifies how long the remote access server's modem waits before calling back. You can add this entry to the registry of the remote access client to change the default value calculated by the remote access server.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# DefaultIdleTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFF (*seconds*) | 0x0 |

**Description**

Specifies how long an idle port remains active after its last transmission before it is closed.

| Value | Meaning |
|---|---|
| 0x0 | Idle ports do not time out. An idle port can remain active indefinitely. |
| 0x1–0xFFFFFFFF *seconds* | The number of seconds an idle port can remain active. |

**Note**

- The "Disconnect if Idle for" Routing and Remote Access policy takes precedence over the value of this entry.
- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# DefaultPortLimit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFE (*ports*) \| 0xFFFFFFFF | 0xFFFFFFFF |

**Description**

Specifies the maximum number of ports that can be grouped into a multi-link bundle.

| Value | Meaning |
|---|---|
| 0x0–0xFFFFFFFE | The maximum number of ports in a bundle. |
| 0xFFFFFFFF | There is no limit on the number of ports in a bundle. |

**Note**

- The "Limit Maximum ports" policy on the **Multilink** tab in Routing and Remote Access takes precedence over the value of this entry.
- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# DefaultSessionTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFF (*seconds*) | 0x0 |

**Description**

Specifies how long a port can remain active before it is closed.

| Value | Meaning |
|---|---|
| 0x0 | Ports do not time out. A port can remain active indefinitely. |
| 0x1–0xFFFFFFFF *seconds* | The number of seconds the port remains active. |

**Note**

- The "Restrict maximum session to" policy in Routing and Remote Access takes precedence over the value of this entry.
- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# LowerBandwidthThreshold

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1–100 (*percent*) | 50 |

**Description**
Specifies the minimum bandwidth used for the Bandwidth Allocation Protocol (BAP). If the percentage of bandwidth used falls below this minimum for more than the value specified in the TimeBelowThreshold entry, then Routing and Remote Access closes the connection.

If the value of this entry is 0, then no minimum bandwidth is enforced.

### Note

- This setting can be overridden by a remote access policy. If you have set the **Capacity** option on the **Multilink** tab in Routing and Remote Access to a value other than 50%, then that setting takes precedence over the value of this registry entry.
- This entry is used only when BAP is enabled on the computer.
- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
TimeBelowThreshold

# MaxConfigure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–255 (*packets*) | 10 |

**Description**

Indicates how many Configure-Request packets Routing and Remote Access can send without receiving a valid Configure-ACK,Configure-NAK, or Configure-Reject response. If Routing and Remote Access does not receive a valid response after sending the number of packets indicated by the value of this entry, then it assumes that the peer cannot respond.

# MaxFailure

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–255 (*packets*) | 5 |

**Description**

Specifies how many Config-Nak packets Routing and Remote Access can send while negotiating a Point-to-Point Protocol (PPP) connection. If the negotiation does not succeed after the specified number of Config-Nak packets are sent, then Routing and Remote Access terminates the PPP negotiation and considers it to have failed.

# MaxReject

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–255 (*packets*) | 5 |

**Description**

Indicates how many Config-Reject packets Routing and Remote Access can send while negotiating a Point-to-Point Protocol (PPP) connection. If the negotiation does not succeed after the specific number of rejections are sent, then Routing and Remote Access terminates the PPP negotiation and considers it to have failed.

# MaxTerminate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–255 (*packets*) | 2 |

**Description**
Specifies how many Terminate-Request packets Routing and Remote Access can send. If Routing and Remote Access does not receive a Terminate-ACK response after sending the number of packets indicated by the value of this entry, it assumes that the peer cannot respond.

**Change Method**
You can change this entry only by editing the registry.

# NegotiateTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFF (*seconds*) | 0x96 (*150 seconds*) |

**Description**

Specifies how long Point-to-Point Protocol (PPP) negotiation can continue. If the negotiation does not result in a successful convergence in the time indicated by the value of this entry, then the server disconnects the line.

| Value | Meaning |
|---|---|
| 0x0 | There is no time limit for PPP negotiation. The line will not be disconnected for exceeding a negotiation time limit. |
| 0x1–0xFFFFFFFF seconds | The time limit for PPP negotiation. |

**Change Method**

You can change this entry only by editing the registry.

# RestartTimer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFF (*seconds*) | 0x3 |

**Description**
Specifies how long Routing and Remote Access waits for a peer to respond to the transmission of Configure-Request and Terminate-Request packets before retransmitting the packets.

**Change Method**
You can make changes to this subkey only by using registry editor.

# TimeBelowThreshold

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*seconds*) | 0x78 (*2 minutes*) |

**Description**
Specifies how long a connection using Bandwidth Allocation Protocol (BAP) can remain below the minimum bandwidth use (as specified by the value of the LowerBandwidthThreshold entry) before it is disconnected.

**Note**

- This setting can be overridden by a remote access policy. If you have set the **Capacity...for a period of** option on the **Multilink** tab in Routing and Remote Access to a value other than 120 seconds, then that setting takes precedence over the value of this entry.

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# ControlProtocols

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

**Description**
The ControlProtocols subkey stores configuration data for Point-to-Point Protocol (PPP) control and authentication protocols, such as IPCP, IPXCP, PAP, and EAP.

# BuiltIn

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\ControlProtocols

**Description**
The BuiltIn subkey stores configuration data for Point-to-Point Protocol (PPP) control and authentication protocols that are included in Windows Server 2003.

# AcceptVJCompression

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether Internet Protocol Control Protocol (IPCP) accepts IPCP standard option 0x2, Van Jacobson (VJ) header compression.

| Value | Meaning |
|---|---|
| 0 | IPCP does not accept VJ compression. |
| 1 | IPCP accepts VJ compression. |

**Note**

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# DefaultDomain

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Domain name or computer name* | *Domain name or computer name* |

**Description**

Specifies the domain used to validate a logon attempt when no domain is specified. This entry applies only to networked computers that are running Windows Server 2003. If the computer is not a member of a domain, then this entry specifies the name of the computer.

By default, the system uses the name of the primary domain of the local computer. However, you can add this entry to the registry to override the default value and to force the system to use the specified domain.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# BuiltIn\Path

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path and file name* | %SystemRoot%\System32\rasppp.dll |

**Description**

Stores the path to the dynamic-link library (DLL) for the Routing and Remote Access control protocols. A single DLL supports all control protocols.

# RegisterRoutersWithNameServers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Directs Routing and Remote Access to register all outgoing dial-up routers with the Domain Name System (DNS) and Windows Internet Name Service (WINS) servers of the server router.

| Value | Meaning |
|---|---|
| 0 | Routers dialing out of the network do not register with the server router's DNS and WINS servers. |
| 1 | Routers dialing out of the network register with the server router's DNS and WINS servers. |

**Activation Method**
Changes to this entry take effect the next time a router dials out of the network.

### Note

- This entry affects routers only when they dial out of the network. It does not affect registration of routers under any other conditions.
- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# RequestNameServerAddresses

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether Internet Protocol Control Protocol (IPCP) on remote access clients request Microsoft extension options 0x81, 0x82, 0x83, and 0x84 for Windows Internet Name Service (WINS) and Domain Name System (DNS) server address negotiation.

| Value | Meaning |
|---|---|
| 0 | Prevents IPCP from requesting the Microsoft extension options for WINS and DNS server address negotiation. |
| 1 | IPCP on remote access clients request the Microsoft extension options for WINS and DNS server address negotiation. |

**Note**

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.
- In Windows NT 4.0 and earlier, this entry appeared in the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Rasman\PPP\IPCP subkey.

# RequestVJCompression

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether Internet Protocol Control Protocol (IPCP) on remote access clients request IPCP standard option 0x2, Van Jacobson (VJ) header compression.

| Value | Meaning |
|---|---|
| 0 | Prevents remote access clients from requesting VJ compression. |
| 1 | Remote access clients request VJ compression. |

**Note**

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.
- In Windows NT 4.0 and earlier, this entry appeared in the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Rasman\PPP\IPCP subkey.

# EAP Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP

**Description**
The EAP subkey stores configuration data for the Extensible Authentication Protocol (EAP) service of the Point-to-Point Protocol (PPP). EAP provides high-level security for virtual private network (VPN) and wireless (IEEE 802.1x) connections.

# 13

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP

**Description**
The 13 subkey stores configuration data for the Extensible Authentication Protocol–Transport Level Security (EAP-TLS). EAP-TLS provides highly secure authentication for smart cards.

# IgnoreNoRevocationCheck

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Lets Extensible Authentication Protocol–Transport Level Security (EAP-TLS) clients connect even if the server does not perform or cannot complete a revocation check of the client's certificate chain (excluding the root certificate). Typically, revocation checks fail because the certificate does not include revocation information.

By default, an EAP-TLS client cannot connect unless the server completes a revocation check of the client's certificate chain (including the root certificate) and verifies that none of the certificates has been revoked. However, you can use this entry to override the default behavior.

This entry does not interfere with the revocation check. Also, it does not permit clients to connect if the revocation check reveals that a certificate in the chain has been revoked.

| Value | Meaning |
|---|---|
| 0 | Clients cannot connect unless a revocation check completes successfully. |
| 1 | Clients can connect even if the revocation check cannot be completed. |

You can use this entry to authenticate clients whose certificate does not include certificate revocation list distribution points (CRPs), such as those from third parties and from the Microsoft Certificate Authority prior to Windows 2000.

**Note**

- This entry is effective only when it appears in the registry of a Routing and Remote Access server.
- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
IgnoreRevocationOffline
NoRootRevocationCheck
NoRevocationCheck

# IgnoreRevocationOffline

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Lets Extensible Authentication Protocol–Transport Level Security (EAP-TLS) clients connect even if a server that stores a relevant revocation list is not available on the network.

By default, EAP-TLS does not permit a client to connect unless it can complete a revocation check of the client's certificate chain and verify that none of the certificates has been revoked. If EAP-TLS cannot connect to a server that stores a revocation list, then it considers the certificate to have failed the revocation check. However, you can add this entry to the registry to direct EAP-TLS to ignore errors related to unavailable servers.

This entry prevents a client from being denied permission to connect because poor network conditions prevented the revocation check from completing successfully.

| Value | Meaning |
|---|---|
| 0 | If the server cannot connect to a server that stores a revocation list, then it considers a certificate to have failed the revocation check. |
| 1 | If the server cannot connect to a server that stores a revocation list, then it ignores the error and permits the client to connect. |

**Note**

- This entry is effective only when it appears in the registry of a Routing and Remote Access server.
- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
NoRootRevocationCheck
IgnoreNoRevocationCheck
NoRootRevocationCheck
NoRevocationCheck

# MaxTLSMessageLength

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x5DC–0xFFFFFFFF (*bytes*) | 0x4000 (*16,384 bytes*) |

**Description**
Specifies the maximum length of a Transport Layer Security (TLS) message or a set of fragmented TLS messages. If the client or server encounters a longer message, then it ignores the message.

This entry is designed to prevent malicious users from harming the service by sending excessively long TLS messages.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# NoRevocationCheck

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Prevents Extensible Authentication Protocol–Transport Level Security (EAP-TLS) from performing a revocation check of the EAP-TLS client's public key certificate.

The revocation check verifies that the public key certificate (and the certificates in its certificate chain) have not been revoked.

| Value | Meaning |
|---|---|
| 0 | Enables the revocation check |
| 1 | Disables the revocation check |

**Note**

- This entry is effective only if it appears in the registry of a Routing and Remote Access server.
- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**Tip**

- This entry disables the revocation check of the client's complete certificate chain. To disable the revocation check of the client's root certificate only, use the IgnoreNoRevocationCheck entry.

**See Also**
**Concepts**
NoRootRevocationCheck
IgnoreRevocationOffline
IgnoreNoRevocationCheck

# NoRootRevocationCheck

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Prevents Extensible Authentication Protocol–Transport Level Security (EAP-TLS) from performing a revocation check of the EAP client's root public key certificate.

The revocation check verifies that the public key certificate (and the certificates in its certificate chain) have not been revoked.

This entry prevents only the revocation check of the client's root certificate. A revocation check is still performed on the remainder of the client's certificate chain.

| Value | Meaning |
|---|---|
| 0 | EAP-TLS performs a revocation check on the client's entire certificate chain, including the root certificate. |
| 1 | EAP-TLS does not perform a revocation check on the root certificate. |

You can use this entry to authenticate clients whose certificate does not include certificate revocation list distribution points (CDPs), such as those from third parties and from the Microsoft Certificate Authority prior to Windows 2000. Also, this entry can prevent certification-related delays that occur when a certificate revocation list is offline or is expired.

**Note**

- This entry is effective only when it appears in the registry of a Routing and Remote Access server.
- Windows Server 2003 does not add this entry to the registry. You can add it by editing the registry or by using a program that edits the registry.

**Tip**

- This entry only disables the revocation check of the client's root certificate. To disable the revocation check of the entire certificate chain, use the IgnoreNoRevocationCheck entry.

**See Also**
**Concepts**
IgnoreRevocationOffline
IgnoreNoRevocationCheck

# RemoteAccess

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The RemoteAccess subkey stores configuration data for Routing and Remote Access.

Several other subkeys, such as RasMan, RasAuto, and RasAcd, also store entries for secondary services that support dial-in connections. The RemoteAccess subkey stores data that applies to the entire service.

**Change Method**
Some entries in this subkey can be changed using the Routing and Remote Access Manager. For details, see the description of each entry.

# Accounting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess

**Description**
The Accounting subkey stores configuration data for Routing and Remote Access accounting providers, such as Windows and Remote Authentication Dial-In User Service (RADIUS). Accounting providers record connection information in log files.

This subkey does not contain any entries. It only contains a Accounting\Providers Subkey subkey, which contains entries that apply to all dial-up accounting providers and subkeys that represent each provider and that store provider-specific entries.

# Accounting\Providers Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Accounting

**Description**
The Providers subkey stores configuration data for Routing and Remote Access authentication providers, such as Windows and Remote Authentication Dial-In User Service (RADIUS).

This subkey contains entries that apply to all dial-up accounting providers. It also contains subkeys that represent each provider, which are named by the globally unique identifier (GUID) of the provider and which store entries that apply only to that provider.

Windows Server 2003 includes the following accounting providers:

- {1AA7F840-C7F5-11D0-A376-00C04FC9DA04} represents the RADIUS accounting provider.
- {1AA7F846-C7F5-11D0-A376-00C04FC9DA04} represents the Windows accounting provider.

# ActiveProvider Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Accounting\Providers

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Accounting or authentication provider GUID* | {1AA7F846-C7F5-11D0-A376-00C04FC9DA04} (*GUID of Windows Authentication*) |

**Description**

Specifies the accounting provider used to record dial-up connection data in a log file. This entry stores the globally unique identifier (GUID) of the accounting provider you select by using the Routing and Remote Access snap-in.

Windows is the default accounting provider, but you can use the Routing and Remote Access snap-in to select an alternate accounting provider. You cannot select a third-party authentication or accounting provider.

Windows Server 2003 includes the following accounting providers:

- {1AA7F840-C7F5-11D0-A376-00C04FC9DA04} represents the RADIUS accounting provider.
- {1AA7F846-C7F5-11D0-A376-00C04FC9DA04} represents the Windows accounting provider.

Entries specific to these providers are stored in the following subkeys in the RemoteAccess\Accounting\Providers subkey:

- {1AA7F840-C7F5-11D0-A376-00C04FC9DA04} (RADIUS accounting provider subkey)
- {1AA7F846-C7F5-11D0-A376-00C04FC9DA04} (Windows accounting provider subkey)

**Change Method**

To change the value of this entry, use Routing and Remote Access. Right-click the name of the server, click **Properties**, click the **Security** tab, and then select a provider from the **Accounting provider** drop-down box.

**Caution**

- To change the value of this entry, use Routing and Remote Access. Do not edit the registry.

# {1AA7F840-C7F5-11D0-A376-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Accounting\Providers

**Description**
The {1AA7F840-C7F5-11D0-A376-00C04FC9DA04} subkey stores configuration data for the Remote Authentication Dial-In User Service (RADIUS) accounting provider.

This subkey can also contain a Servers subkey, which contains subkeys representing each RADIUS server used for RADIUS accounting. The server-specific subkeys store entries that apply only to the named server.

**Tip**

- To configure a Routing and Remote Access server to use RADIUS accounting, use the **Accounting Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

(Windows accounting provider subkey)

**See Also**
**Concepts**
{1AA7F846-C7F5-11D0-A376-00C04FC9DA04}

# {1AA7F840-C7F5-11D0-A376-00C04FC9DA04}\Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Accounting\Providers\{1AA7F840-C7F5-11D0-A376-00C04FC9DA04}

**Description**
The Servers subkey stores server-specific configuration data for the Remote Authentication Dial-In User Service (RADIUS) accounting provider.

This subkey does not contain any entries. It contains subkeys representing each RADIUS server used for RADIUS accounting. The server-specific subkeys store entries that are configured separately for each server and that apply only to the named server.

### Note

- Routing and Remote Access adds this subkey to the registry when you add a server to the list in the **RADIUS Accounting** dialog box in Routing and Remote Access.

### Tip

- To configure a Routing and Remote Access server to use RADIUS accounting, use the **Accounting Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

# {1AA7F840-C7F5-11D0-A376-00C04FC9DA04}\Server-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Accounting\Providers\{1AA7F840-C7F5-11D0-A376-00C04FC9DA04}\Servers

**Description**

The *Server-name* subkey stores server-specific configuration data for the Remote Authentication Dial-In User Service (RADIUS) accounting provider.

This subkey is a variable representing the subkeys in the RemoteAccess\Accounting\{1AA7F840-C7F5-11D0-A376-00C04FC9DA04}\Servers subkey. It does not actually appear in the registry. This variable subkey is used to display the entries and subkeys that are common to the subkeys in the {1AA7F840-C7F5-11D0-A376-00C04FC9DA04}\Servers subkey.

The *Server-name* subkeys in your registry represent RADIUS servers used for RADIUS accounting. Each subkey represents a particular server and stores configuration data that applies only to that server.

**Note**

- Routing and Remote Access adds a *Server-name* subkey to the registry when you add a server to the list in the **RADIUS Accounting** dialog box in Routing and Remote Access.

**Tip**

- To configure a Routing and Remote Access server to use RADIUS accounting, use the **Accounting Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

# Server-name\NASIpAddress Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Accounting\Providers\{1AA7F840-C7F5-11D0-A376-00C04FC9DA04}\Servers*Server-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *IP address* | (None) |

**Description**
Identifies the network access server (NAS) to use with Remote Authentication Dial-In User Service (RADIUS) accounting.

The Routing and Remote Access server binds to this address. Also, it uses this address as the value of the NAS_IP_Address RADIUS attribute (Attribute 4).

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
You can change this entry only by editing the registry.

**Tip**

- To configure a Routing and Remote Access server to use RADIUS accounting, use the **Accounting Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

(RADIUS accounting providers)

**See Also**
**Concepts**
Server-name\NASIpAddress Entry

# Server-name\Retries Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Accounting\Providers\{1AA7F840-C7F5-11D0-A376-00C04FC9DA04}\Servers*Server-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x2 |

**Description**
Specifies how many times the Routing and Remote Access server sends an Accounting-Request message to the Remote Authentication Dial-In User Service (RADIUS) servers. Routing and Remote Access repeats the message if none of the RADIUS servers respond.

If a RADIUS server does not respond before the number of retries is exhausted, then Routing and Remote Access terminates the connection and records an error in the System Log in Event Viewer. If the client tries to connect again, then Routing and Remote Access treats it like a new connection attempt.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Tip**

- To configure a Routing and Remote Access server to use RADIUS accounting, use the **Accounting Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

**See Also**
**Concepts**
Server-name\Retries

# {1AA7F846-C7F5-11D0-A376-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Accounting\Providers

**Description**
The {1AA7F846-C7F5-11D0-A376-00C04FC9DA04} subkey stores configuration data for the Windows accounting provider.

**Tip**

- To configure a Routing and Remote Access server to use Windows accounting, use the **Accounting Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

(RADIUS accounting provider subkey)

**See Also**
**Concepts**
{1AA7F840-C7F5-11D0-A376-00C04FC9DA04}

# Authentication Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess

**Description**
The Authentication subkey stores configuration data for Routing and Remote Access authentication providers, such as Windows and Remote Authentication Dial-In User Service (RADIUS).

The Authentication subkey does not contain any entries. It contains a Authentication\Providers Subkey subkey, which contains entries that apply to all dial-up authentication providers and subkeys that represent each provider and that store provider-specific entries.

# Authentication\Providers Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Authentication

**Description**
The Providers subkey stores configuration data for Routing and Remote Access authentication providers, such as Windows and Remote Authentication Dial-In User Service (RADIUS).

This subkey contains entries that apply to all dial-up authentication providers. It also contains subkeys that represent each provider, which are named by the globally unique identifier (GUID) of the provider and store entries that apply only to that provider.

Windows Server 2003 includes the following authentication providers:

- {1AA7F83F-C7F5-11D0-A376-00C04FC9DA04} represents the RADIUS authentication provider.
- {1AA7F841-C7F5-11D0-A376-00C04FC9DA04} represents the Windows authentication provider.

**See Also**
**Concepts**
{1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}
{1AA7F841-C7F5-11D0-A376-00C04FC9DA04}
{1AA7F840-C7F5-11D0-A376-00C04FC9DA04}

# Providers\ActiveProvider

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Authentication\Providers

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | {1AA7F841-C7F5-11D0-A376-00C04FC9DA04} |

**Description**
Specifies the authentication provider used to authenticate users who dial-up to the computer. This entry stores the globally unique identifier (GUID) of the authentication provider you select by using the Routing and Remote Access snap-in.

Windows is the default authentication provider, but you can use the Routing and Remote Access snap-in to select an alternate authentication provider.

Windows Server 2003 includes the following authentication providers:

- {1AA7F83F-C7F5-11D0-A376-00C04FC9DA04} represents the RADIUS authentication provider.
- {1AA7F841-C7F5-11D0-A376-00C04FC9DA04} represents the Windows authentication provider.

Entries specific to these providers are stored in the following subkeys in the RemoteAccess\Authentication\Providers subkey:

- {1AA7F83F-C7F5-11D0-A376-00C04FC9DA04} (RADIUS authentication provider subkey)
- {1AA7F841-C7F5-11D0-A376-00C04FC9DA04} (Windows authentication provider subkey)

**Change Method**
To change the value of this entry, use the Routing and Remote Access snap-in. Right-click the name of the server, click **Properties,** click the **Security** tab, and then select a provider from the **Authentication provider** drop-down box.

**Caution**

- To change the value of this entry, use the Routing and Remote Access snap-in. Do not edit the registry.

# {1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Authentication\Providers

**Description**
The {1AA7F83F-C7F5-11D0-A376-00C04FC9DA04} subkey stores configuration data for the Remote Authentication Dial-In User Service (RADIUS) authentication provider.

This subkey also can contain a Servers subkey, which contains subkeys representing each RADIUS server used for RADIUS authentication. The server-specific subkeys store entries that apply only to the named server.

**Tip**

- To configure a Routing and Remote Access server to use RADIUS authentication, use the **Authentication Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

(Windows authentication provider subkey) (RADIUS accounting provider subkey)

**See Also**
**Concepts**
{1AA7F841-C7F5-11D0-A376-00C04FC9DA04}
{1AA7F840-C7F5-11D0-A376-00C04FC9DA04}

# {1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}\Servers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Authentication\Providers\{1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}

**Description**
The Servers subkey stores server-specific configuration data for a Remote Authentication Dial-In User Service (RADIUS) authentication provider.

This subkey does not contain any entries. It contains subkeys representing each RADIUS server used for RADIUS authentication. The server-specific subkeys store entries that are configured separately for each server and that apply only to the named server.

### Note

- Routing and Remote Access adds this subkey to the registry when you add a server to the list in the **RADIUS Authentication** dialog box in Routing and Remote Access.

### Tip

- To configure a Routing and Remote Access server to use RADIUS authentication, use the **Authentication Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

# Servers\Server-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Authentication\Providers\{1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}\Servers

**Description**

The *Server-name* subkey stores server-specific configuration data for a Remote Authentication Dial-In User Service (RADIUS) authentication provider.

The *Server-name* subkey is a variable representing the subkeys in the RemoteAccess\Authentication\{1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}\Servers subkey. It does not actually appear in the registry. This variable subkey displays the entries and subkeys that are common to the subkeys in the {1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}\Servers subkey.

The *Server-name* subkeys in your registry represent RADIUS servers that are used for RADIUS authentication. Each subkey represents a particular server and stores configuration data that applies only to that server.

**Note**

- Routing and Remote Access adds the Servers subkey to the registry when you add a server to the list in the **RADIUS Authentication** dialog box in Routing and Remote Access.

**Tip**

- To configure a Routing and Remote Access server to use RADIUS authentication, use the **Authentication Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

# Server-name\NASIpAddress

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Authentication\Providers\{1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}\Servers*Server-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *IP address* | (None) |

**Description**
Identifies the network access server (NAS) to use with Remote Authentication Dial-In User Service (RADIUS) authentication.

The Routing and Remote Access server binds to this address. Also, it uses this address as the value of the NAS_IP_Address RADIUS attribute (Attribute 4).

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
You can change this entry only by editing the registry.

**Tip**

- To configure a Routing and Remote Access server to use RADIUS authentication, use the **Authentication Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

(RADIUS accounting providers)

**See Also**
**Concepts**
Server-name\NASIpAddress Entry

# Server-name\Retries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Authentication\Providers\{1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}\Servers*Server-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x2 |

**Description**
Specifies how many times the Routing and Remote Access server sends an Access-Request message to the Remote Authentication Dial-In User Service (RADIUS) servers. Routing and Remote Access repeats the message if none of the RADIUS servers respond.

If a RADIUS server does not respond before the number of retries is exhausted, then Routing and Remote Access terminates the connection, and it records an error in the System Log in Event Viewer. If the client tries to connect again, Routing and Remote Access treats it like a new connection attempt.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Tip**

- To configure a Routing and Remote Access server to use RADIUS authentication, use the **Authentication Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

**See Also**
**Concepts**
Server-name\Retries Entry

# SendSignature

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Authentication\Providers\{1AA7F83F-C7F5-11D0-A376-00C04FC9DA04}\Servers*Server-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether Routing and Remote Access includes a digital signature in all Access-Request messages to this Remote Authentication Dial-In User Service (RADIUS) server.

| Value | Meaning |
|---|---|
| 0 | Routing and Remote Access includes a digital signature only in Extensible Authentication Protocol (EAP) Access-Request messages. |
| 1 | Routing and Remote Access includes a digital signature in all Access-Request messages. |

**Change Method**

To change the value of this entry, use the Routing and Remote Access snap-in. Right-click the name of the server, click **Properties**, and then click the **Security** tab. From the **Authentication Provider** drop-down list, select **RADIUS Authentication**, click **Configure**, click **Edit** and select the **Always use digital signatures** check box.

**Tip**

- To configure a Routing and Remote Access server to use RADIUS authentication, use the **Authentication Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

# {1AA7F841-C7F5-11D0-A376-00C04FC9DA04}

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Authentication\Providers

**Description**
The {1AA7F841-C7F5-11D0-A376-00C04FC9DA04} subkey stores configuration data for the Windows authentication provider.

**Tip**

- To configure a Routing and Remote Access server to use Windows authentication, use the **Authentication Provider** box on the **Security** tab in the Routing and Remote Access snap-in.

# RemoteAccess\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess

**Description**
The Parameters subkey stores configuration data for Routing and Remote Access. The data in this subkey specifies the value of variables that apply to the entire service.

**Activation Method**
To make changes to the values of the entries in this subkey effective, you must stop and restart the remote access service.

# LanOnlyMode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 0 |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# LoggingFlags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–3 | 3 |

**Description**
Specifies which events Routing and Remote Access records in the System Log in Event Viewer. This entry affects logging by Routing and Remote Access only. It does not affect logging by routing components or by transports, such as IP and IPX.

| Value | Meaning |
|---|---|
| 0 | No log. Routing and Remote Access does not record any events in the System Log. |
| 1 | Records errors only. |
| 2 | Records errors and warnings only. |
| 3 | Records all events (errors, warnings, and information). |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use Routing and Remote Access Manager. Right-click the server name, click **Properties**, and then click the **RAS Event Logging** tab. Do not change the value of this entry by editing the registry directly.

**Activation Method**
You must stop and restart the remote access service before changes to the value of this entry take effect.

# RouterType

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–7 | 1 |

**Description**
Specifies the functions of the remote access router.

To enable more than one function, sum the value for each function. For example, to enable a router for remote access clients (value = 1) and for inter-LAN (local area network) communication (value = 2), set the value of this entry to 3.

| Value | Meaning |
|---|---|
| 1 | Remote access server |
| 2 | LAN-to-LAN router |
| 4 | LAN-to-WAN-to-LAN or Demand Dial router |

**Change Method**
To change the value of this entry, use Routing and Remote Access Manager. Right-click the name of the remote access server, click **Properties**, and click the **General** tab. Do not edit the registry.

**Activation Method**
You must stop and restart the remote access service before changes to the value of this entry take effect.

### Note

- This value appears in the registry when you use the Routing and Remote Access Manager snap-in to configure the router.

# SecurityHostTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 20–600 (*seconds*) | 120 |

**Description**
Specifies how long Routing and Remote Access waits for clients to complete third party authentication. If authentication is not complete before the value of this entry expires, then Routing and Remote Access considers authentication to have failed.

The time limit specified in this entry applies only to the third party authentication that occurs before Routing and Remote Access authentication. The time limit for Routing and Remote Access authentication is specified in the value of the AuthenticateTime entry.

**Activation Method**
You must stop and restart the remote access service before changes to the value of this entry take effect.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
AuthenticateTime

# ServerFlags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x3A740 |

**Description**
Determines whether important features of the remote access service are enabled.

The value of this entry is a bitmask. To enable a feature, set the bit associated with the feature to 1, or add the hexadecimal value associated with the feature to the value of this entry. To disable a feature, set the bit associated with the feature to 0, or subtract the hexadecimal value associated with the feature from the value of this entry.

To view the valid values for this entry, see Values for Remote Access ServerFlags.

The default value, 0x3A740, enables the following features for dial-in clients and routers:

- PPPCFG_UseSwCompression (0x00000002)
- PPPCFG_NegotiateSPAP (0x00000040)
- PPPCFG_NegotiateMSCHAP (0x00000100)
- PPPCFG_UseLcpExtensions (0x00000200)
- PPPCFG_NegotiateMultilink (0x00000400)
- PPPCFG_NegotiateBacp (0x00002000)
- PPPCFG_NegotiateEAP (0x00008000)
- PPPCFG_NegotiatePAP (0x00010000)
- PPPCFG_NegotiateMD5CHAP(0x00020000)

**Change Method**
To change the value of this entry, use Routing and Remote Access Manager.

**Activation Method**
You must stop and restart the remote access service before changes to the value of this entry take effect.

# ServiceDll

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %SystemRoot%\System32\mprdim.dll |

**Description**
Specifies the location of the dynamic-link library (DLL) for the remote access service.

**Caution**

- Do not change the value of this entry. If you change this value or move the DLL file, you can prevent the Routing and Remote Access service from starting.

# Stamp

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 | 0 |

**Description**
Ensures that the time stamp in the Parameters subkey is accurate even if the registry entries are copied to and changed within an internal memory cache. Routing and Remote Access uses this entry internally.

**Activation Method**
You must stop and restart the remote access service before changes to the value of this entry take effect.

### Caution

- Do not change the value of this entry. The Routing and Remote Access service uses this entry internally.

# AuthenticateRetries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–10 | 2 |

**Description**
Specifies the maximum number of times the remote access service attempts to authenticate a client.

This entry differs from the CallbackRetries entry, which specifies the maximum number of times the remote access service tries to call the client back for the purpose of authentication.

**Change Method**
To change the value of this entry, use the **Max number of unsuccessful authentication retries** Group Policy setting (Computer Settings\Software Policies\Windows NT Remote Access).

**Activation Method**
You must stop and restart Routing and Remote Access before changes to the value of this entry take effect.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
AuthenticateTime
CallbackRetries

# AuthenticateTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 20–600 (*seconds*) | 120 |

**Description**
Specifies the time within which a user must be successfully authenticated. If the client does not complete the authentication process within this time, then the user is disconnected.

**Change Method**
To change the value of this entry, use the **Max time limit for authentication** Group Policy setting (Computer Settings\Software Policies\Windows NT Remote Access).

**Activation Method**
You must stop and restart Routing and Remote Access before changes to the value of this entry take effect.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
AuthenticateRetries

# AutoDisconnect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 \| 0x1–0xFFFFFFFF (*seconds*) | 0x0 |

**Description**
Specified how long an inactive connection can remain open. If a connection remains inactive longer than the value of this entry specifies, then the connection is closed. Connections are considered to be inactive if they are not transferring NetBIOS session data, such as copying files, accessing network resources, or sending and receiving e-mail.

| Value | Meaning |
|---|---|
| 0x0 | Inactive connections never disconnect automatically. This setting might be appropriate when clients are running NetBIOS datagram applications. |
| 0x1–0xFFFFFFFF *seconds* | The maximum time the system permits an inactive connection to remain open. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To change the value of this entry, use the **Auto Disconnect** Group Policy setting (Computer Settings\Software Policies\Windows NT Remote Access).

**Activation Method**
You must stop and restart Routing and Remote Access before changes to the value of this entry take effect.

# CallbackRetries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x0 |

**Description**
Specifies how many times the server attempts to call the client back if the initial callback fails. This entry is used only if the callback security feature is enabled.

If the value of this entry is 0x0 and the initial callback fails, then the system does not attempt any further callbacks.

This entry is designed for environments where telephone systems are not always reliable and multiple calls are often required to establish a connection.

This entry differs from the AuthenticateRetries entry, which specifies the maximum number of times the remote access service attempts to authenticate the client, not just to contact the client by calling back.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
DefaultCallbackDelay
AuthenticateRetries
CallbackTime

# CallbackTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 2–12 (*seconds*) | 2 |

**Description**
Specifies how long the server waits before calling back the client for authentication. The value of this entry is a default that is used only if the client does not specify an interval between callbacks, and only if the callback feature is enabled.

This entry differs from the AuthenticateTime entry, which specifies the time within which a user must be successfully authenticated, not just the time to call back the client.

**Change Method**
To change the value of this entry, use the **Wait interval for callback** Group Policy setting (Computer Settings\Software Policies\Windows NT Remote Access).

**Activation Method**
You must stop and restart Routing and Remote Access before changes to these values take effect.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
AuthenticateTime
CallbackRetries
SecurityHostTime

# AccountLockout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

**Description**
Stores the database for the *Account Lockout* feature of Routing and Remote Access and of Internet Authentication Service (IAS).

*Account Lockout* prohibits dial-in access to a server after repeated authentication attempts fail. By default, Account Lockout is disabled.

This subkey stores entries that specify the operating criteria for the Account Lockout feature. It also includes domain:user-namesubkeys, which represent user accounts that the feature affects. Each of the domain:user-name subkeys represents one user, and each tracks the number of failed attempts to authenticate that user.

**Note**

- Account Lockout prevents dial-in access to a locked user account, not just to the computer that submitted the failed authentication.

**Tip**

- To disable Account Lockout or to reset all user denial counts to 0 while the server is running, set the value of the MaxDenials entry to 0. If you delete the AccountLockout subkey or the domain:user-name subkey for a user, then the service does not detect the change until the system is restarted.

# MaxDenials

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 0x0 |

**Description**
Specifies how many attempts to authenticate a user account can fail before the user is denied dial-in access to the server. Authentication fails when a user enters an incorrect password for a valid user name.

This entry also enables and disables the Account Lockout feature of the Routing and Remote Access service. If the value of this entry is greater than 0, then the Account Lockout feature is enabled on this computer. Otherwise, it is disabled.

When the value of the Denials entry reaches the value of this entry within the time specified by the ResetTime (mins) entry, the user account is locked. The system does not process any more dial-in authentication attempts until one of the following occurs:

- The value of the **ResetTime (mins)** entry expires.
- An administrator deletes the user's domain:user-name subkey.
- An administrator sets the value of the **Denials** entry for that account either to 0 or to a value less than that stored in the **MaxDenials** entry.

### Note

- Authentication fails when a user enters an incorrect password for a valid user name.

### Tip

- To disable Account Lockout or to reset all user denial counts to 0 while the server is running, set the value of this entry to 0. If you delete the AccountLockout subkey or the domain:user-name subkey for a user, the service does not detect the change until the system is restarted.

**See Also**
**Concepts**
ResetTime (mins)

# ResetTime (mins)

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 2880 (*48 hours*) |

**Description**
Specifies the time interval for the Account Lockout feature of Routing and Remote Access.

When the value of the Denials entry reaches the value of the MaxDenials entry within the time specified by the value of the ResetTime (mins) entry, the user account is locked. The system does not process any more dial-in authentication attempts until one of the following occurs:

- The value of the **ResetTime (mins)** entry expires.
- An administrator deletes the user's domain:user-name subkey.
- An administrator sets the value of this entry either to 0 or to a value less than that stored in the **MaxDenials** entry.

### Note

- Authentication fails when a user enters an incorrect password for a valid user name.

### Tip

- The denial counter for each user is stored in the Denials entry. The denial threshold for the computer is stored in the MaxDenials entry.

**See Also**
**Concepts**
Denials
MaxDenials
ResetTime (mins)

# domain:user-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout

**Description**
The *domain:user-name* subkey stores the *Account Lockout* record for a user.

*Account Lockout* prohibits dial-in access to a server after repeated authentication attempts fail. By default, Account Lockout is disabled.

The system adds this subkey to the registry when a dial-in authentication attempt for the user account fails, that is, when the value of the Denials entry for a user account is 1. The system deletes this subkey when the value of the ResetTime (mins) entry expires.

The *domain:user-name* subkey is a variable representing the subkeys in the AccountLockout subkey. It does not actually appear in the registry. This variable subkey is used to display the entries that can appear in any AccountLockout user subkey.

**Note**

- Authentication fails when a user enters an incorrect password for a valid user name.

**Tip**

- To disable Account Lockout or to reset all user denial counts to 0 while the server is running, set the value of the MaxDenials entry to 0. If you delete the AccountLockout subkey or the *domain:user-name* subkey for a user, the service does not detect the change until the system is restarted.

**See Also**
**Concepts**
AccountLockout
Denials
MaxDenials
ResetTime (mins)

# Denials

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout*domain:user-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | (None) |

**Description**
Stores the current number of failed attempts to authenticate a user account for dial-up access.

This entry stores a counter that increments on each failed attempt to authenticate the account. Authentication fails when a user enters an incorrect password for a valid user name. This entry and the user subkey that contains it are created the first time an authentication attempt fails during the interval specified by the value of the ResetTime (mins) entry. They are deleted when the value of the **ResetTime (mins)** entry expires or when the user account is successfully authenticated.

When the value of this entry reaches the value of the MaxDenials entry within the time specified by the ResetTime (mins) entry, the user account is locked. The system does not process any more dial-in authentication attempts until one of the following occurs:

- The value of the **ResetTime (mins)** entry expires.
- An administrator deletes the user's domain:user-name subkey.
- An administrator sets the value of this entry either to 0 or to a value less than that stored in the **MaxDenials** entry.

**Note**

- Authentication fails when a user enters an incorect password for a valid user name.

# Ip

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

**Description**
The Ip subkey stores entries that use IP for the Routing and Remote Access service.

# AllowClientIpAddresses

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ip

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether dial-up clients can request a connection to a particular Routing and Remote Access server by specifying the server's IP7 address.

| Value | Meaning |
|---|---|
| 0 | Clients cannot request a particular server. |
| 1 | Clients can request a particular server. |

**Note**

- In Windows NT 4.0 and earlier, this entry appeared in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Rasman\PPP\IPCP subkey.
- The "Client may request an IP address" Routing and Remote Access policy takes precedence over the value of this entry.

# DNSNameServers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ip

| Data type | Range | Default value |
|---|---|---|
| REG_MULTI_SZ | IP address [IP address] | (IP address of DNS server) |

**Description**
Specifies which Domain Name System (DNS) servers the Routing and Remote Access client uses.

You can add this entry to the registry of a Routing and Remote Access server to override the automatic assignment of the remote access server's DNS servers to the remote access client.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
WINSNameServer
WINSNameServerBackup

# InitialAddressPoolSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ip

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*addresses*) | 0xA (*10 addresses*) |

**Description**

Limits the number of addresses that the Routing and Remote Access server gets from the Dynamic Host Configuration Protocol (DHCP) server at one time.

The number of addresses in each batch cannot exceed the value of this entry. By default, the number of addresses in each batch equals the number of available Routing and Remote Access ports, plus one. But if the value of this entry is less than the number of ports plus one, then Routing and Remote Access sets the batch size to the value of this entry.

The Routing and Remote Access server gets a batch of addresses when the service starts, and it gets the same number of addresses each time it uses up the previous batch.

# WINSNameServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ip

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *IP address* | ( IP address of WINS server) |

**Description**
Specifies which Windows Internet Name Service (WINS) server the remote access client uses.

You can add this entry to the registry of a remote access server to override the automatic assignment of the remote access server's WINS server to the remote access client.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
DNSNameServers
WINSNameServerBackup

# WINSNameServerBackup

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ip

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *IP address* | (IP address of WINS server) |

**Description**
Specifies which backup Windows Internet Name Service (WINS) name server the remote access client uses.

You can add this entry to the registry of a remote access server to override the automatic assignment of the remote access server's backup WINS server to the remote access client.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
DNSNameServers
WINSNameServer

# StaticAddressPool

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ip

| Data type | Range | Default value |
|---|---|---|
| | | 0 |

**Description**
The StaticAddressPool subkey stores a range of static IP addresses for use on a Routing and Remote Access server.

This subkey contains numbered subkeys, each of which represents a particular static IP address pool. The entries in the numbered subkeys define the range of IP addresses in the pool.

**Change Method**
To create a static IP address pool, use Routing and Remote Access. Right-click a server name, click **Properties**, click the **IP** tab, select **Static address pool**, and then click **Add**.

### Note

- Routing and Remote Access adds this subkey to the registry when you use Routing and Remote Access to add and configure a static address pool.

### Caution

- To create a static IP address pool, use the **IP** tab in Routing and Remote Access. Do not edit the registry.

# Pool-number

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ip\StaticAddressPool

**Description**
The *Pool-number* subkey represents a particular pool of static IP addresses for Routing and Remote Access servers. These subkeys are numbered, beginning with 0, in the order that the pools are created. The entries in the *Pool-number* subkeys define the range of IP addresses in the pool.

The *Pool-number* subkey is a variable representing the subkeys in the RemoteAccess\Parameters\Ip\StaticAddressPool subkey. It does not actually appear in the registry. This variable subkey is used to display the entries and subkeys that are common to the subkeys in the StaticAddressPool subkey.

### Note

- Routing and Remote Access adds a *Pool-number* subkey to the registry when you add a static IP address pool by defining a range in the **New Address Range** dialog box in Routing and Remote Access.

### Tip

- For more information about static address pools, see Windows Server 2003 Help and Support Center or Windows XP Professional Help and Support Center.

### Caution

- To create a static IP address pool, use the **IP** tab in Routing and Remote Access. Do not edit the registry.

# Pool-number\From

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ip\StaticAddressPool*Pool-number*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *IP address (in integers, not in dotted decimal format)* | (None) |

**Description**
Specifies the first IP address in the pool range.

This entry stores the value in the **Start IP address** box on the **New Address Range** dialog box in Routing and Remote Access.

This entry stores the IP address as four contiguous hexadecimal values. For example, the IP address 1.2.3.4 would appear in this entry as 0x1020304, and 255.255.255.255 would appear as 0xFFFFFFFF (255 = 0xFF).

**Change Method**
To create a static IP address pool, use Routing and Remote Access. Right-click a server name, click **Properties**, click the **IP** tab, select **Static address pool**, and then click **Add**.

### Note

- Routing and Remote Access adds the **From** entry to the registry when you add a static IP address pool by defining a range in the **New Address Range** dialog box in Routing and Remote Access.

### Tip

- For more information about static address pools, see Windows Server 2003 Help and Support Center or Windows XP Professional Help and Support Center.

### Caution

- To create a static IP address pool, use the **IP** tab in Routing and Remote Access. Do not edit the registry.

# To

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ip\StaticAddressPool*Pool-number*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *IP address (in integers, not in dotted decimal format)* | (None) |

**Description**
Specifies the last IP address in the pool range.

This entry stores the value in the **End IP address** box on the **New Address Range** dialog box in Routing and Remote Access.

This entry stores the IP address as four contiguous hexadecimal values. For example, the IP address 1.2.3.4 would appear in this entry as 0x1020304 and 255.255.255.255 would appear as 0xFFFFFFFF (255 = 0xFF).

**Change Method**
To create a static IP address pool, use Routing and Remote Access. Right-click a server name, click **Properties**, click the **IP** tab, select **Static address pool**, and then click **Add**.

### Note

- Routing and Remote Access adds the **From** entry to the registry when you add a static IP address pool by defining a range in the **New Address Range** dialog box in Routing and Remote Access.

### Tip

- For more information about static address pools, see Windows Server 2003 Help and Support Center or Windows XP Professional Help and Support Center.

### Caution

- To create a static IP address pool, use the **IP** tab in Routing and Remote Access. Do not edit the registry.

**See Also**
**Concepts**
Pool-number\From

# IPX Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters

**Description**
Stores entries for remote access when it is implemented by using Internetwork Packet Exchange (IPX).

# FirstWanNode

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\Ipx

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *12-digit hexadecimal value representing the node* | (None) |

**Description**

Identifies the base node number assigned to callers.

If callers are using different network numbers, then the system assigns the base node number to all callers. However, if callers are using the same network number, then the system assigns the base node number to the first caller, and then increments the node number by one for each subsequent caller.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Note**

- You must enter the node number in the registry editor Regedit.exe in hexadecimal format. Regedit.exe does not convert text values to a different radix.

# RemoteAccess\Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess

**Description**
Stores policy settings shared by Internet Authentication Service (IAS) and Routing and Remote Access. Remote access policies determine permissions and connection settings.

The entries in this subkey also determine which attributes are used to authenticate users and how they are evaluated.

# Sermouse

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The Sermouse subkey stores configuration data for the Microsoft serial mouse driver in Windows NT 4.0 and earlier. In Windows Server 2003, this data appears in the device-specific subkeys of the Enum Subkey subkey. However, if you upgrade an installation of Windows NT 4.0, some of the values in the Sermouse subkey are copied to the device-specific keys. For details, see the descriptions of the entries in this subkey.

# Sermouse\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Sermouse

**Description**

The Parameters subkey stores configuration data for the Microsoft serial mouse driver in Windows NT 4.0 and earlier. In Windows Server 2003, this data appears in the device-specific subkeys of the Enum Subkey subkey. However, if you upgrade from Windows NT 4.0, some of the values in this subkey are copied to the device-specific keys. For details, see the descriptions of the entries in this subkey.

# Parameters\MouseDataQueueSize Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Sermouse\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x64 (*100*) |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

In the Windows Server 2003 serial mouse driver, the queue size is equal to one packet, and it is not configurable.

**Caution**

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# Parameters\NumberOfButtons Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Sermouse\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF | 0x2 |

**Description**
Stores the number of buttons on the serial mouse that were detected at startup in Windows NT 4.0 and earlier. In Windows Server 2003, the number of mouse buttons is stored in the device-specific subkeys of the Enum Subkey subkey. However, the system uses the value of this entry as source data.

If you have a serial mouse and have upgraded from Windows NT 4.0, the system copies the value of this entry to the device-specific Enum subkeys of all serial mouse devices, and it uses the values of the entries in those subkeys.

**Caution**

- The value of this entry is copied to the device-specific Enum subkeys of all serial mouse devices recognized by the system. If you change the value of this entry, then the new value is used for all new serial mouse devices that the system recognizes.

# Parameters\PointerDeviceBaseName Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Sermouse\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Base port device name* | PointerPort |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# Parameters\SampleRate Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Sermouse\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | *Mouse sampling rate in Hz* | 0x28 (40 Hz) |

**Description**
Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 or earlier. In Windows Server 2003, the serial mouse driver does not sample mouse activity.

### Caution

- Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results.

# Services\SNMP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The SNMP subkey stores configuration data for the Simple Network Management Protocol (SNMP) service.

The SNMP service implements SNMP network standards. This service allows incoming SNMP requests to be serviced by the local computer. It includes agents that monitor activity in network devices and report to the network console workstation.

# SNMP\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\SNMP

**Description**
The Parameters subkey contains subkeys that function like standard registry entries and contain configuration information for the Simple Network Management Protocol (SNMP) service. The SNMP service implements SNMP network standards. This service allows incoming SNMP requests to be serviced by the local computer. It includes agents that monitor activity in network devices and report to the network console workstation.

# EnableAuthenticationTraps

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\SNMP\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Enables and disables authentication traps.

When authentication traps are enabled, Simple Network Management Protocol (SNMP) sends a trap when it receives a request that does not match any community name or host filter in its lists.

| Value | Meaning |
|---|---|
| 0 | The SNMP service does not send authentication traps. |
| 1 | The SNMP service sends an authentication trap when it receives a request that does not match any community name or host filter in its lists. |

**Change Method**
To change the value of this entry, use Computer Management. In the console tree, click **Services and Applications**, click **Services**, and then click **SNMP Service**. From the **Action** menu, click **Properties**, click the **Security** tab, and then select or clear the **Send authentication trap** check box.

### Note

- In Windows NT 4.0, this data was stored in the **Switch** entry in the EnableAuthenticationTraps subkey.

# ExtensionAgents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\SNMP\Parameters

**Description**
The ExtensionAgents subkey contains entries that indicate where the registry stores the configuration data for Simple Network Management Protocol (SNMP) extension agents.

The names of the entries in this subkey are decimal integers, beginning with 0. The value of each entry is the registry path to the data that is loaded by the dynamic-link library (DLL) for the extension agent. All entries in this subkey have the REG_SZ data type.

For example, the following entry indicates that the configuration data for the WWW service is located in the HKEY_LOCAL_MACHINE\Software\Microsoft\W3SVC\CurrentVersion subkey:

| Entry name | Data type | Value |
|---|---|---|
| 0 | REG_SZ | HKLM\Software\Microsoft\W3SVC\CurrentVersion |

# Parameters\PermittedManagers

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\SNMP\Parameters

**Description**
The PermittedManagers subkey stores the computer names of hosts that can send requests to a Simple Network Management Protocol (SNMP) agent.

Each entry in this subkey represents a host that is a permitted manager. The name of the entry is an ordinal number that reflects the order in which you entered the host when you configured SNMP. The value of the entry is the host name. The data type of all entries in this subkey is REG_SZ.

In the following example, the host (Host01) can send requests to an agent:

| Entry name | Data type | Value |
|---|---|---|
| 1 | REG_SZ | Host01 |

**Change Method**
To change the value of this entry, use the Services snap-in. Right-click **SNMP Service**, then click **Properties**.

### Notes

- The default configuration has no entry under this subkey. The SNMP agent accepts SNMP packets from any host.
- Only users who belong to the local Administrators group are able to access this subkey.

# Parameters\TrapConfiguration

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\SNMP\Parameters

**Description**
The TrapConfiguration subkey stores names of hosts that receive trap messages for communities. Trap messages are alerts that Simple Network Management Protocol (SNMP) agents initiate.

This subkey contains a series of subkeys that each represent one community. The entries in each community subkey identify the hosts to which the SNMP agent sends traps for that community. The name of each entry is an ordinal number, beginning with 1. The value of each entry can be a host name, an IP address, or an Internetwork Packet Exchange (IPX) address. The data type of all entries in the TrapConfiguration subkey is REG_SZ.

For example, the TrapConfiguration subkey contains a subkey representing a valid community, Community01. The Community01 subkey contains one or more entries with the following format:

| Entry name | Data type | Value |
|---|---|---|
| 2 | REG_SZ | 138.10.12.86 |

In this example, the SNMP agent sends traps for Community01 to the host at the specified IP address.

**Change Method**
To view, add, or delete hosts from a community subkey, use Services in Computer Management.

# Parameters\ValidCommunities

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\SNMP\Parameters

**Description**
The ValidCommunities subkey stores a list of the communities defined to the Simple Network Management Protocol (SNMP) service.

In SNMP, a community is a group of hosts (servers, workstations, hubs, and routers) that SNMP administers together. An SNMP agent accepts requests only from management systems within the communities it recognizes and only if the permissions established for the community allow the requested operation.

Each entry in this subkey represents one community that you identify to SNMP. The name of the entry is the community name. (A community name is a network password that the hosts in the community share.) The value is a bitmask specifying the permissions of the community. The data type is REG_DWORD.

For example, the entries in this subkey could store the following list of communities:

| Entry Name | Data Type | Value |
|---|---|---|
| ROCommunity | REG_DWORD | 0x4 |
| RWCommunity | REG_DWORD | 0x8 |

**Change Method**
To change the value of this entry, use the Services snap-in. Right-click **SNMP Service**, then click **Properties**.

### Note

- Only users who belong to the local Administrators group are able to access this subkey.

# SysmonLog

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Replace Blanks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\SysmonLog

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *any single character* | (None) |

**Description**
Specifies a single character that is used to replace blanks in the names of collections and file names created by the Performance console (Perfmon.msc) of the Microsoft Management Console.

**Change Method**
This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must stop and restart the Performance console for changes to take effect.

### Note

- In Windows 2000, the Performance console (Perfmon.msc) created log files without spaces in their names. In Windows XP and Windows Server 2003, spaces are allowed in the names of log files. However, you can add this entry to the registry to emulate the behavior of the Performance console for Windows 2000.

# Tcpip

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The Tcpip subkey stores configuration data for the Microsoft implementation of TCP/IP (Transmission Control Protocol/Internet Protocol). TCP/IP is a set of networking protocols that enable communication over diverse, interconnected networks.

This subkey stores some of the configuration data for TCP/IP, primarily the settings for the TCP/IP service. Other subkeys also store entries that affect the operation of the TCP/IP implementation on your computer. These other subkeys include those for network components and for services related to TCP/IP, such as Dynamic Host Configuration Protocol (DHCP), NetBIOS over TCP/IP (NetBT), and Windows Sockets (Winsock).

**Change Method**
To change the values of many of the entries in the Tcpip subkey and its subkeys, use Network and Dial-up Connections. Right-click **Local Area Connection**, click **Properties**, click **Internet Protocol (TCP/IP)**, and then click the **Properties** button.

# Tcpip\Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip

**Description**

The Parameters subkey stores configuration data that applies to the TCP/IP service as a whole. This data can be overridden by settings specific to a network component using TCP/IP, such as a network adapter driver, or by settings specific to a service using TCP/IP, such as Dynamic Host Configuration Protocol (DHCP) or Windows Sockets (Winsock). These component-specific entries reside in subkeys representing the component.

# Parameters\KeepAliveInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF (*milliseconds*) | 0x3E8 (*1,000 milliseconds = 1 second*) |

**Description**
Specifies how often TCP repeats keep-alive transmissions when no response is received. TCP sends keep-alive transmissions to verify that idle connections are still active. This prevents TCP from inadvertently disconnecting active lines.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
TcpMaxDataRetransmissions
KeepAliveTime

# KeepAliveTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF (*milliseconds*) | 0x6DDD00 (*7,200,000 milliseconds = 2 hours*) |

**Description**
Specifies how often TCP sends keep-alive transmissions. TCP sends keep-alive transmissions to verify that an idle connection is still active.

This entry is used when the remote system is responding to TCP. Otherwise, the interval between transmissions is determined by the value of the Parameters\KeepAliveInterval entry.

By default, keep-alive transmissions are not sent. The TCP keep-alive feature must be enabled by a program (such as Telnet), or by an Internet browser (such as Internet Explorer).

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
Parameters\KeepAliveInterval

# MaxUserPort

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 5,000–65,534 (*port number*) | 5000 |

**Description**
Specifies the highest port number that TCP can assign when an application requests an available user port from the system. Typically, *ephemeral* ports (those used briefly) are allocated to port numbers 1024 through 5000.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# NumTcbTablePartitions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFF (*TCB table partitions*) | 0x1 |

**Description**

Specifies the number of partitions in the Transport Control Block table.

Partitioning the Transport Control Block table minimizes contention for table access. This is especially useful on multiprocessor systems.

### Note

- This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

### Caution

- Do not change the value of this entry before carefully studying the effect of different values in a test environment. When testing, do not enter a value greater than two times the number of processors on the computer.

# PPTPTcpMaxDataRetransmissions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFF (*retransmissions*) | 0x5 |

**Description**
Specifies how many times an unacknowledged Point-to-Point Tunneling Protocol (PPTP) packet is retransmitted before the connection is dropped.

This entry lets you configure a limit on PPTP retransmissions separately from the limit for regular TCP retransmissions, which is stored in the value of the TcpMaxDataRetransmissions entry. You can then adjust the TCP setting to prevent denial-of-service attacks (also known as SYN flooding) without affecting PPTP traffic.

**Tip**

- To prevent the TCP dead gateway detection feature from inadvertently disconnecting a congested Internet link, set this value higher than the default value of the TcpMaxDataRetransmissions entry.

**See Also**
**Concepts**
TcpMaxDataRetransmissions

# PrioritizeRecordData

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the DNS client, upon receiving an answer to a DNS query, looks for the address most similar to its own, when there are multiple addresses in the answer to its query. By choosing the address most like its own, the client connects to the server nearest to it on the network, which reduces network traffic. When the value of this entry is 1, the client prioritizes for the nearest address; when the value is zero, the client attempts to connect to the first address in the answer.

### Note

- This scheme attempts to reduce network traffic by connecting clients with local IP addresses. This registry entry controls whether the client prioritizes addresses, but Windows Server 2003 DNS server can also be configured to prioritize the addresses in its answers according to their similiarity to the address of the requesting client. This behavior by Whister DNS server is controlled by the value of LocalNetPriority.

- Subnet prioritization, whether configured on the client or the server, nullifies the load-balancing benefits of *round-robin name resolution*. Round-robin name resolution is a scheme by which the server rotates the order of the addresses in its answers to spread connection loads across all applicable servers. By default, Windows Server 2003 DNS service uses round-robin name resolution. For more information about subnet prioritization and round-robin name resolution, see the *TCP/IP Core Networking Guide* of the *Microsoft Windows Server 2003 Technical Reference* .

**See Also**
**Concepts**
LocalNetPriority

# SackOpts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Enables and disables the Selective Acknowledgment (SACK) feature of Windows Server 2003 TCP/IP. SACK is specified in RFC 2018, TCP Selective Acknowledgement Options.

SACK is an optimizing feature that lets you acknowledge receipt of individual blocks of data in a continuous sequence, rather than just the last sequence number. The recipient can tell the sender that one or more data blocks are missing from the middle of a sequence, and the sender can retransmit only the missing data.

| Value | Meaning |
|---|---|
| 0 | Disables SACK. |
| 1 | Enables SACK. |

# SynAttackProtect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SynAttackProtect**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the SYN flooding attack protection feature of TCP/IP is enabled. SYN flooding attack protection is enabled when the value of this entry is 1 and the value of the **TcpMaxConnectResponseRetransmissions** entry is at least 2 (see note below).

The SYN flooding attack protection feature of TCP detects symptoms of denial-of-service attacks (also known as SYN flooding), and it responds by reducing the time that the server spends on connection requests that it cannot acknowledge.

| Value | Meaning |
|---|---|
| 0 | SYN flooding attack protection is not enabled. |
| 1 | SYN flooding attack protection is enabled. |

**Notes**

- This value is used only when the number of SYN-ACK retransmissions is likely to impair the server, that is, when the value of the **TcpMaxConnectResponseRetransmissions** entry is at least 2.
- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
TcpMaxHalfOpen
TcpMaxHalfOpenRetried

# Tcp1323Opts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tcp1323Opts**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 \| 3 | 3 |

**Description**
Specifies whether TCP uses the timestamping and window scaling features described in RFC 1323, TCP Extensions for High Performance.

Window scaling permits TCP to negotiate a scaling factor for the TCP receive window size, allowing for a very large TCP receive window of up to 1 GB. The TCP receive window is the amount of data that the sending host can send at one time on a connection.

Timestamps help TCP measure round trip time (RTT) accurately in order to adjust retransmission timeouts. The Timestamps option provides two timestamp fields of 4 bytes each in the TCP header, one to record the time the initial transmission is sent and one to record the time on the remote host.

This entry is a 2-bit bitmask. The lower bit determines whether scaling is enabled; the higher bit determines whether timestamps are enabled. To enable a feature, set the bit representing the feature to 1. To disable a feature, set its bit to 0.

| Value | Meaning |
|---|---|
| 0 (00) | Timestamps and window scaling are disabled. |
| 1 (01) | Window scaling is enabled. |
| 2 (10) | Timestamps are enabled. |
| 3 (11) | Timestamps and window scaling are enabled. |

**Note**

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**Tip**

- For more information about RFC 1323, see the *Request for Comments (RFC)* link on the Web Resources page.

# TcpMaxConnectRetransmissions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpMaxConnectRetransmissions**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–255 (*retransmission attempts*) | 2 |

**Description**
Specifies how many times TCP retransmits an unanswered request for a new connection. TCP retransmits new connection requests until they are answered or until this value expires.

TCP/IP adjusts the frequency of retransmissions over time. The delay between the original transmission and the first retransmission for each interface is determined by the value of the **TcpInitialRTT** entry. By default, it is three seconds. This delay doubles after each attempt. After the final attempt, TCP/IP waits for an interval equal to double the last delay, and then it abandons the connection request.

### Notes

- This entry determines how many times TCP retransmits requests for new connections. When sending data on existing connections, the maximum number of retransmissions is determined by the value of the TcpMaxDataRetransmissions entry.

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
TcpMaxDataRetransmissions

# TcpMaxDataRetransmissions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpMaxDataRetransmissions**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*retransmission attempts*) | 0x5 |

**Description**
Specifies how many times TCP retransmits an unacknowledged data segment on an existing connection. TCP retransmits data segments until they are acknowledged or until this value expires.

TCP/IP adjusts the frequency of retransmissions over time. TCP establishes an initial retransmission interval by measuring the round trip time on the connection. The interval doubles with each successive retransmission on a connection, and it is reset to the initial value when responses resume.

This entry is also used in the Windows algorithm for defining non-operational (*dead*) gateways.

A given connection defines a gateway as dead (and switches to the next gateway in the list in stored in the value of the **DefaultGateway** or **DhcpDefaultGateway** entries) when a packet sent to the gateway must be retransmitted more than half of the number of times specified in the value of this entry. The system defines a gateway as dead when more than 25 percent of its connections have switched to the next default gateway in the list.

### Notes

- This entry determines how many times TCP retransmits data segments. The maximum number of retransmissions of requests for new connections is determined by the value of the TcpMaxConnectRetransmissions entry.

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
TcpMaxConnectRetransmissions

# TcpMaxDupAcks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpMaxDupAcks**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–3 | 2 |

**Description**
Specifies how many duplicate ACKs (ACKs for the same sequence numbers) constitute a signal to retransmit a segment. If you set the value of this entry to 1, then the system retransmits a segment when it receives an ACK for a segment with a sequence number that is less than the number of the segment currently being sent.

When data arrives with a sequence number that is greater than expected, the receiver assumes that data with the expected number was dropped, and it immediately sends an ACK with the ACK number set to the expected sequence number. The receiver sends ACKs set to the same missing number each time it receives a TCP segment that has a sequence number greater than expected. The sender recognizes the duplicate ACKs and sends the missing segment.

**Note**

- This entry is used only when the receiver supports the fast retransmit feature. Fast retransmit lets TCP retransmit data before the retransmission timer (as set by the value of the **TcpInitialRtt** entry) expires.

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# TcpMaxHalfOpen

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpMaxHalfOpen**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1–0xFFFFFFFF (*connections*) | Windows Server 2003: 0x1F4 (*500*) Windows XP Professional: 0x64 (*100*) |

**Description**
Specifies how many connections the server can maintain in the half-open(SYN-RCVD) state before TCP/IP initiates SYN flooding attack protection. This entry is used only when SYN flooding attack protection is enabled on this server, that is, when the value of the SynAttackProtect entry is 1 and the value of the **TcpMaxConnectResponseRetransmissions**entry is at least 2.

This entry establishes one of three configurable thresholds that, if exceeded, trigger TCP's SYN attack flooding protection feature. Because SYN flooding often results in many half-open connections, TCP interprets an elevated number of half-open connections to be a symptom of SYN flooding.

The other two thresholds are:

- The number of connections that remain in the half-open(SYN-RCVD) state even after a connection request has been retransmitted exceeds the value of the TcpMaxHalfOpenRetried entry.

- The number of connection requests that the system refuses exceeds the value of the **TcpMaxPortsExhausted** entry. The system must refuse all connection requests when its reserve of open connection ports runs out.

## Notes

- The value of this entry should be greater than the value of the TcpMaxHalfOpenRetried entry.

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
SynAttackProtect
TcpMaxHalfOpenRetried

# TcpMaxHalfOpenRetried

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpMaxHalfOpenRetried**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x50 (*80*)–0xFFFFFFFF | Windows Server 2003: 0x190 (*400*) Windows XP Professional: 0x50 (*80*) |

**Description**
Specifies how many connections the server can maintain in the half-open(SYN-RCVD) state even after a connection request has been retransmitted. If the number of connections exceeds the value of this entry, then TCP/IP initiates SYN flooding attack protection. This entry is used only when SYN flooding attack protection is enabled on this server, that is, when the value of the SynAttackProtect entry is 1 and the value of the **TcpMaxConnectResponseRetransmissions** entry is at least 2.

This entry establishes one of three configurable thresholds that, if exceeded, trigger TCP's SYN attack flooding protection feature. Because SYN flooding often results in many half-open connections, TCP interprets an elevated number of half-open connections to be a symptom of SYN flooding.

The other two thresholds are:

- The total number of connections in the half-open(SYN-RCVD) state exceeds the value of the TcpMaxHalfOpen entry.
- The number of connection requests that the system refuses exceeds the value of the **TcpMaxPortsExhausted** entry. The system must refuse all connection requests when its reserve of open connection ports runs out.

**Notes**

- The value of this entry should be less than the value of the TcpMaxHalfOpen entry.
- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

**See Also**
**Concepts**
SynAttackProtect
TcpMaxHalfOpen

# TcpMaxSendFree

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpMaxSendFree**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 - 65535 | 5000 |

**Description**
Sets the size limit of the TCP header resource. On computers with ample RAM, increasing this limit can improve responsiveness during a SYN flood denial-of-service attack.

**Activation Method**
You must restart the computer for changes to this entry to take effect.

**See Also**
**Concepts**
TcpMaxHalfOpen
SynAttackProtect

# TcpNumConnections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpNumConnections**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x40000–0xFFFFFE | 0xFFFFFE |

**Description**
Specifies the maximum number of connections that TCP can have open simultaneously.

If the value of this entry is 0, then you cannot establish any connections.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# TcpTimedWaitDelay

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpTimedWaitDelay**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0 0x12C (*0–300 seconds*) | 0x78 (*120 seconds = 2 minutes*) |

**Description**
Specifies the time that must elapse before TCP can release a closed connection and reuse its resources. This interval between closure and release is known as the TIME_WAIT state or 2MSL state. During this time, the connection can be reopened at much less cost to the client and server than establishing a new connection.

RFC 793 requires that TCP maintains a closed connection for an interval at least equal to twice the maximum segment lifetime (2MSL) of the network. When a connection is released, its socket pair and TCP control block can be used to support another connection. By default, the maximum segment lifetime is defined to be 120 seconds, and the value of this entry is equal twice that, or 4 minutes. However, you can use this entry to customize the interval.

Reducing the value of this entry allows TCP to release closed connections faster, providing more resources for new connections. However, if the value is too low, then TCP might release connection resources before the connection is complete, requiring the server to use additional resources to reestablish the connection.

### Notes

- Normally, TCP does not release closed connections until the value of this entry expires. However, TCP can release connections before this value expires if it is running out of TCP control blocks. The number of TCP control blocks that the system creates is specified by the value of the **MaxFreeTcbs** entry.

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# TcpUseRFC1122UrgentPointer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpUseRFC1122UrgentPointer**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies which mode TCP uses for urgent data. The two modes interpret the urgent pointer in the TCP header and the length of the urgent data differently. The two modes do not interoperate.

| Value | Meaning |
|---|---|
| 0 | TCP uses the mode used by systems derived from Berkeley Software Distribution (BSD). |
| 1 | TCP uses the specification in RFC 1122, Requirements for Internet hosts—communication layers, for urgent data. |

**Note**

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# TcpWindowSize

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TcpWindowSize**
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0–0x3FFFFFFF (1073741823 decimal) bytes | see description |

**Description**
Sets the maximum size of the TCP receive window. The receive window specifies the number of bytes that a sender can transmit without receiving an acknowledgment. In general, larger receive windows improve performance over high-latency, high-bandwidth networks. For greatest efficiency, the receive window should be an even multiple of the TCP Maximum Segment Size.

The TCP/IP stack of Windows Server 2003 was designed to tune itself in most environments. Instead of using a fixed size for the receive window, TCP negotiates for and adjusts to an even increment of the maximum segment size. Matching the receive window to even increments of the maximum segment size increases the percentage of full-sized TCP segments used during bulk data transmission. (Sizes larger than 64 KB can be achieved only when connecting to other systems that support RFC 1323 *Window Scaling*.)

The default value is the smaller of the following numbers:

- 65535
- Value of registry entry **GlobalMaxTcpWindowSize**
- Four times the MSS on the network
- Even multiple of the MSS larger than 16384

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Activation Method**
You must restart the computer for changes to this entry to take effect.

### Notes

- The default window size can start at 17520 for Ethernet, but may shrink slightly when a connection to another computer is established that supports extended TCP options (such as SACK and TIMESTAMPS) because these options increase the TCP header beyond the usual 20 bytes, which decreases room for data.
- When an entry with the name of this entry is in HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters, it sets the receive window size globally for all TCP interfaces. However, an entry with this name is in HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interface-name* has precedence for the interface named in its registry path, and its value, rather than whatever value might be set globally, is used for that interface.

# DisableDynamicUpdate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Disables the Domain Name System (DNS) dynamic update registration for all interfaces on the system. The default value is zero.

With *dynamic update*, DNS client computers automatically register and update their resource records whenever address changes occur.

| Value | Meaning |
|---|---|
| 0 | Enables dynamic update registration. |
| 1 | Disables dynamic update registration. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
You can also disable dynamic update registration by using Group Policy.

**Activation Method**
To make changes to the value of this entry effective, you must restart the computer.

### Note

- For dynamic update to operate on any interface, it must be enabled at both the system and interface levels. To disable it for a particular interface, add the entry **DisableDynamicUpdate** to an *Interface-name* subkey, and set its value to 1. To disable it on all interfaces of a system, add the **DisableDynamicUpdate** entry to the Tcpip\Parameters subkey and set its value to 1.

# DisableReplaceAddressesInConflicts

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Prevents the Domain Name System (DNS) client from overwriting an existing resource record when it discovers an address conflict during dynamic update.

An address conflict occurs when the DNS client discovers that an existing A (address) record associates its DNS name with the IP address of a different computer.

By default, the DNS client tries to replace the original registration with a record associating the DNS name to its own IP address. However, you can use this entry to direct DNS to back out of the registration process and to record an error in the Event Viewer log instead.

This entry is designed for zones that do not use secure dynamic update. It prevents unauthorized users from changing the IP address registration of a client computer.

| Value | Meaning |
|---|---|
| 0 | The DNS client overwrites the existing A record with an A record for its own IP address. |
| 1 | The DNS client backs out of the registration process and writes an error to the Event Viewer log. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

- To find other registry entries related to DNS dynamic update, type "dynamic update" on the Search tab in this file.

**See Also**
**Concepts**
DisableReverseAddressRegistrations

# DisableReverseAddressRegistrations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Disables the Domain Name System (DNS) dynamic update registration of PTR (pointer) records by this DNS client. PTR records associate an IP address with a computer name.

This entry is designed for enterprises in which the primary DNS server that is authoritative for the reverse lookup zone cannot or is configured not to perform dynamic updates. It reduces unnecessary network traffic and eliminates event log errors that record failed attempts to register PTR records.

| Value | Meaning |
|---|---|
| 0 | Register PTR records. |
| 1 | Do not register PTR records. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

- To find other registry entries related to DNS dynamic update, type "dynamic update" on the Search tab in this file.

**See Also**
**Concepts**
DisableReplaceAddressesInConflicts

# DisableUserTOSSetting

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether individual applications can alter the type of service (TOS) bits in the header of outgoing IP packets. In general, individual applications should not be allowed to manipulate TOS bits, because this can defeat system policy mechanisms.

**Activation Method**
You must restart the computer for changes to this entry to take effect.

### Note

- Every IP datagram has an octet commonly known as the *TOS byte*. The first three bits specify a level of priority. The next four bits, interpreted as a number, specify a type of service. The last bit is currently unused. For more information about the TOS byte, see IETF RFC 1349: *Type of Service in the Internet Protocol Suite*.

# EnableICMPRedirect

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
This parameter controls whether Windows Server 2003 alters its routing table in response to Internet Control Message Protocol (ICMP) messages that instruct it to direct datagrams for the recipient along a different route.

**Activation Method**
You must restart the computer for changes to this entry to take effect.

### Notes

- ICMP provides a means by which a host sending IP datagrams can be informed about delivery issues. ICMP does not guarantee delivery of IP datagrams (that kind of error correction is left to higher level protocols, like TCP), but rather, it allows network devices, like a router, to tell a sending computer about delivery errors, to suggest shorter routes to a destination, and to assist in probing the network. For more information about ICMP, see RFC 792 *Internet Control Message Protocol*, and RFC 1122 *Requirements for Internet Hosts—Communication Layers*.

- Windows Server 2003 accepts redirection messages from any host in the route between this computer and the destination computer, and not just first-hop routers. Accepting redirection from only first-hop routers causes problems in some scenarios involving Routing & Remote Access Service (RAS).

# IGMPLevel

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 2 |

**Description**

Specifies the extent to which the system supports IP multicasting, and participates in the Internet Group Management Protocol.

| Value | Meaning |
|---|---|
| 0 | Only send IP multicast packets. Treated the same as a value of 1. |
| 1 | Only send IP multicast packets. |
| 2 | Send and receive IP multicast packets (participate in IGMP). |

**Note**

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

# Adapters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Description**
The Adapters subkey stores TCP/IP configuration data that can be configured differently for each network adapter on the computer.

This subkey contains one or more *adapter-name* subkeys. The *adapter-name* subkeys each represent one network adapter, and they are named by their network identifiers. The entries in the *adapter-name* subkeys apply only to the network adapter that the subkey represents, and the values in those entries override any conflicting settings in the Tcpip\Parameters Subkey subkey.

# Adapters\interface-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Adapters

**Description**
The *interface-name* subkeys represent network adapters or network adapter drivers, and they are named for the network identifiers. Each subkey represents a particular adapter, and it stores configuration data that applies only to that adapter.

Most of the entries described here as being stored in the *interface-name* subkey can appear in any of the subkeys in the Adapters subkey.

Some entries can appear in both the Tcpip\Parameters Subkey subkey and the *interface-name* subkeys. If an entry appears in the Parameters subkey, then its value applies to all adapters by default. If the entry appears in an *interface-name* subkey, then its value applies only to the adapter that the subkey represents, and it takes precedence over the value in the Parameters subkey when configuring this particular adapter.

**Note**

- The *interface-name* subkey is a placeholder representing the types of subkeys in the HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Adapters subkey; "interface-name" does not actually appear in the registry.

# LLInterface

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Adapters*interface-name*\

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | Any valid device name. | |

**Description**
Specifies the name of a device to which IP should bind on the interface named in the subkey. If this value contains the name of a link-layer protocol, then IP binds to that protocol. If this value is empty, then IP binds to the built-in ARP module. This entry is always present for all adapters.

**Caution**

- This entry is created and maintained by the TCP/IP Notify object, and you should not modify it. It is included here only for reference.

# Interfaces Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Description**
The Interfaces subkey stores configuration data specific to aTCP/IP communications interface.

This subkey stores TCP/IP configuration data that can be configured differently for each interface that the system uses. It contains one or more subkeys, called *interface-name*, each of which represents one interface. The entries in the *interface-name* subkeys apply only to the interface that the subkey represents, and they take precedence over any conflicting settings in the Tcpip\Parameters Subkey subkey.

# Interfaces\interface-name

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces

**Description**
The *interface-name* subkeys representTCP\IP communication interfaces, and they are named for the network name of the interface. Each subkey represents a particular interface, and it stores configuration data that applies only to that interface.

Most of the entries in the *interface-name* subkey described here can appear in any of the subkeys in the Interfaces subkey.

Some entries can appear in the Tcpip\Parameters Subkey subkey, in an *interface-name* subkey, or in both. If an entry appears in the Parameters subkey, then it applies to all interfaces by default. If it appears in an *interfaces-name* subkey, then it applies only to the interface that the subkey represents, and it takes precedence over the entry in the Parameters subkey when configuring that interface.

**Note**

- The *interface-name* subkey is a placeholder representing the types of subkeys in the subkey Tcpip\Parameters\Interfaces; "interface-name" does not actually appear in the registry.

# EnableDHCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interface-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Determines whether the Dynamic Host Configuration Protocol service (DHCP) is enabled. If the value of this entry is 1, then the DHCP client service is used to configure this interface.

| Value | Meaning |
|---|---|
| 0 | Disables DHCP for this interface. |
| 1 | Enables DHCP for this interface. |

**Change Method**
To change the value of this entry, use Network and Dial-up Connections. Right-click **Local Area Connection**, click **Properties**, and click **Internet Protocol (TCP/IP)**. This entry is associated with the **Obtain an IP address automatically** option.

### Notes

- This entry is added to the registry when you use Network Connections in Control Panel to configure TCP/IP.
- When DHCP is enabled, the phrase **DHCP Enabled** appears in the IP address box in place of a static IP address. To view this setting, in Control Panel double-click Network Connections, right-click **Local Area Connection**, click **Properties**, click **Internet Protocol (TCP/IP)**, click **Advanced**, and then click the **IP Settings** tab.

# IPAutoconfigurationAddress

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interface-name*

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *IP address, in dotted decimal format* | 0.0.0.0 |

**Description**

Stores the IP address that the IP autoconfiguration feature assigns to this interface.

Autoconfiguration allows TCP/IP to configure an interface even when it cannot locate a Dynamic Host Configuration Protocol (DHCP) server. If TCP/IP cannot locate DCHP server, then it configures the interface by using IP addresses from the Microsoft reserved Class B network (169.254.0.0), subnet mask 255.255.0.0. During autoconfiguration, TCP/IP continues trying to locate a DHCP server, and it abandons autoconfiguration if it finds one.

**Note**

- This entry is used only when IP autoconfiguration is enabled, that is, when the value of the IPAutoconfigurationEnabled entry is 1.

**Caution**

- Do not change the value of this entry. If you do, the IP address assigned to the interface might not be reliable.

# IPAutoconfigurationEnabled

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interface-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | (None) |

**Description**
Specifies whether the IP autoconfiguration feature is enabled for this interface. If the value of this entry is 1, then autoconfiguration is enabled; if the value of this entry is 0, or if this entry is absent from the registry, then autoconfiguration is not enabled.

Autoconfiguration allows TCP/IP to configure an interface even when it cannot locate a Dynamic Host Configuration Protocol (DHCP) server. If TCP/IP cannot locate a DCHP server, then it configures the interface by using IP addresses from the Microsoft reserved Class B network (169.254.0.0), subnet mask 255.255.0.0. During autoconfiguration, TCP/IP continues trying to locate a DHCP server, and it abandons autoconfiguration if it finds one.

**IPAutoconfigurationEnabled** also appears in the Parameters subkey, and its value applies, by default, to all interfaces. However, if this entry appears in any subkey of the Interfaces subkey, then it takes precedence over the entry in the Parameters subkey when configuring the interface.

### Note

- Windows Server 2003 does not add this entry to the registry. You can add it by using the registry editor Regedit.exe.

### Caution

- Autoconfiguration, if enabled, could interfere with configuration of IP addresses by other methods.

# SubnetMask

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interface-name*

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | (A space-delimited list of dotted decimal subnet masks.) | 0.0.0.0 |

**Description**

Specifies the subnet mask for the IP address specified in the value of the **IPAddress** entry or the **DhcpIPAddress** entry. This entry, which can be changed by the user, overrides the DHCP-configured values for the subnet mask.

The subnet mask is determined by two entries, **DhcpSubnetMask** (which is configured by DHCP), and **SubnetMask** (which you can configure). If the value of **SubnetMask** is not 0.0.0.0, then it overrides the value of **DhcpSubnetMask**.

**Change Method**

To change the value of this entry, use **Network Connections**. Right-click **Local Area Connection** (or the name of another interface, if appropriate), click **Properties**, click **Internet Protocol (TCP/IP)** , and then click **Properties**. At the bottom right corner of the page, click the **Advanced** button. Under **IP addresses** click the **Add** button, and enter the IP address and corresponding subnet mask.

**Activation Method**

You must restart the computer for changes to take effect.

# Atmarpc

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interface-name*

| Data type | Range | Default value |
|---|---|---|
| NULL (subkey) | | |

**Description**

This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# AddressResolutionTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interface-name*\Atmarpc

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 - 60 | 3 seconds |

**Description**
Specifies the time, in seconds, that the address resolution protocol (ARP) client of the Asynchronous Transfer Mode (ATM) transport waits for a response after sending an ARP request for a unicast IP address or a multicast address resolution service (MARS) request for a multicast or broadcast IP address. After the time specified by this entry has elapsed, the ARP client retransmits the request a maximum number of times equal to one fewer than the number specified in the registry entry **MaxResolutionAttempts**.

**Activation Method**
You must restart the computer for changes to take effect.

### Note

- This entry is installed with its default value during installation of the TCP/IP components.

### Tip

- For more information about configuring ATM under Windows Server 2003, see "ATM Technical Reference" in the *Networking Collection* of the *Windows Server 2003 Technical Reference*.

# ARPEntryAgingTimeout

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interface-name*\Atmarpc

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 90 - 1800 | 900 seconds (15 minutes) |

**Description**

Specifies the length of time, in seconds, that the address resolution protocol (ARP) client of the Asynchronous Transfer Mode (ATM) transport retains address resolution information for a unicast IP address before it is invalidated. When the time specified by this entry has elapsed, the ARP client does one of the following things:

- If there are no virtual circuits (VCs) associated with the IP address, it deletes the ARP entry for this IP address.
- If there is at least one permanent virtual circuit (PVC) associated with the IP address, then it uses Inverse ARP on the PVC to revalidate the ARP entry.
- If there is at least one switched virtual circuit (SVC) associated with the IP address, then it sends an ARP request to the ARP server to revalidate the ARP entry.

# Tcpip/Performance Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip

**Description**
See the description for HKLM\SYSTEM\CurrentControlSet\Services*service-name*\Performance Subkey.

**See Also**
**Concepts**
Performance Subkey

# FriendlyNetworkInterfaceNames

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Performance

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether network adapters are listed by name or by number in performance monitoring tools, such as System Monitor.

Network adapters are represented by instances of the Network Interface performance object. By default, System Monitor displays the display name provided by the network adapter driver as the name of the instance.

This entry does not exist in the registry by default. To direct System Monitor to display numbers instead of names for the network adapters, use the registry editor Regedit.exe to add this entry to the registry and set its value to 0.

| Value | Meaning |
|---|---|
| 0 | Numbers. System Monitor displays numbers to represent network adapters. |
| 1 | Names. System Monitor displays the names of network adapters. |

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
To make changes to this entry effective, restart Windows.

# Parameters\W32Time

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The W32Time subkey stores configuration settings for the Windows Time service. The Windows Time service maintains date and time synchronization for all computers running on a network.

# Config Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\

| Data type | Range | Default value |
|---|---|---|
| NULL | | |

**Description**

The Config subkey stores configuration information for the Windows Time service.

### Notes

The Windows Time service maintains date and time synchronization for all computers that are running on a network, and assigns a current clock value or timestamp for network validation and resource access requests.

# Config\AnnounceFlags Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 10 |

**Description**
Specifies whether this computer is a time server or a reliable time server. A computer will not be classified as reliable unless it is also classified as a time server.

This entry is only used by domain controllers that are running the Windows time service. The flag determines how the service is advertised by Netlogon.

| Value | Meaning |
|---|---|
| 0 | Timeserv_Announce_No, Reliable_Timeserv_Announce_No. The domain controller does not advertise time service. |
| 1 | Timeserv_Announce_Yes. The domain controller always advertises time service. |
| 2 | Timeserv_Announce_Auto. The domain controller automatically determines whether it should advertise time service. |
| 4 | Reliable_Timeserv_Announce_Yes. The domain contoller will always advertise reliable time service. |
| 8 | Reliable_Timeserv_Announce_Auto. The domain controller automatically determines whether it should advertise reliable time service. |

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Note**

- This value combines flags using the *bitwise OR* operator. For each position in the binary value, if any flag has the value of 1 in the corresponding position, the combined value does as well. For example, the flags 1 and 8, which are 0001 and 1000 in binary, are 1001 when combined with bitwise OR.

**Tip**

- W32Time has robust failover mechanisms, and manually configuring additional time services as reliable is generally unnecessary. However, there are two exceptions:

  1. You might want to configure a reliable time service in each site in your enterprise to minimize out-of-site traffic. Each service should be running on a domain controller in the root domain of each site. If the root domain of the site is also the root of the domain hierarchy, then the reliable time services should be configured manually to synchronize from external sources that are known to be highly accurate.

  2. You might want to configure a time service as reliable if it is configured to synchronize with a hardware clock.

  The optimal configuration varies with the enterprise. For more information about configuring time services, see "Windows Time Service Technical Reference" in the *Networking Collection* of the *Microsoft Windows Server 2003 Technical Reference*.

# Config\EventLogFlags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1 \| 2 \| 3 | 2 |

**Description**
Specifies the events that the time service logs.

| Value | Meaning |
|---|---|
| 1 | EvtLog_TimeJump. Time service must make a discontinuous change to the clock. |
| 2 | EvtLog_SourceChange. Time service chooses a new source of time information. |
| 3 | EvtLog_SourceNone. Time service has not acquired time samples for a period that is 1.5 times the maximum poll interval specified by Config\MaxPollInterval, and no longer trusts the local clock to be accurate. |

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**See Also**
**Concepts**
FileLogEntries

# FileLogEntries

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | | |

**Description**

Contains a list of flags, specified by number and seperated by commas, that specify the types of information that should be written to the log file specified by **FileLogName**. A range of numbers is valid, in addition to single numbers. For example, **0-100,103,106** is a valid entry for this value. The following table shows common values and the type of data they specify.

| Value | Meaning |
|---|---|
| 2 | Errors associated with polling a time source. |
| 3 | Errors associated with establishing a session with an NTP peer. |
| 4 | Errors associated with domain hierarchy peers. |
| 5 | Errors associated with manual peers. |
| 6 | Errors associated with sending an NTP response. |
| 7 | Errors associated with NTP authentication. |
| 9 | Errors associated with provider failures. |
| 10 | Errors associated with processing NTP time samples. |
| 11 | Errors associated with configuration information. |
| 12 | Errors associated with stopping and starting the service. |
| 14 | Errors associated with NTP packet validation. |
| 52, 111 | Information relevent to automatic detection of a time source within a domain. |
| 65 | Configuration change information. |
| 67 | Configuration change made to network (for example, when a new network adapter is installed). |
| 105, 106 | Network time protocol (NTP) packets. |

**Activation Method**

To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

# FrequencyCorrectRate Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 4 |

**Description**
Affects the rate at which the time service corrects the clock. This value, in units equal to the reciprocal of the rate at which the clock is corrected, is multiplied by the number of clock ticks in 64 seconds, to produce the base gain in the phase-locked loop (PLL) of the clock correction algorithm.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Note**

- A smaller value makes the time service more responsive to time changes, but if this value is too small, then the clock time is erratic. A larger value makes the system time more stable, but if the value is too large, then the time service waits a long time before synchronizing the clock.

# Config\HoldPeriod

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 5 |

**Description**
Specifies the period of time for which spike detection is disabled in order to bring the local clock into synchronization quickly. A spike is a time sample indicating that time is off by a number of seconds, usually received after good time samples have been returned consistently.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Note**

- In this period, phase-locked loop (PLL) synchronization and spike detection are disabled to allow faster convergence.

# Config\LargePhaseOffset

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 50000000 |

**Description**

Specifies the time offset, in $10^{-7}$ seconds, by which times greater than or equal to this value are to be considered suspicious and possibly erroneous.

**Activation Method**

To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Tip**

- Setting this value higher makes it less likely that the system clock will be set directly, and makes clock correction less frequent. This can be useful if you want to avoid discontinuities in system time (for example, while logging important transactions). However, the time service is more likely to regard bad time samples as good, which can lead to less accurate time.

# LastClockRate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**

Stores the clock rate (in $10^{-7}$ seconds) of the CPU when the service was last stopped. The default value depends on the computer.

**Caution**

This entry is maintained automatically, and you should not modify it.

# LocalClockDispersion Entry

Updated: January 28, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008 R2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 10 |

**Description**
This setting, expressed in seconds, is applicable only when the Network Time Protocol (NTP) server is using the time of the local CMOS clock. The LocalClockDispersion value indicates the maximum error in seconds that is reported by the NTP server to clients that are requesting a time sample. The default value is 10 seconds.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Note**

- This entry is used when the local clock acts as a source of time for other computers, but is not synchronized with another network source. When the clock is synchronized from a network source, the dispersion is calculated relative to the source from which it synchronizes.

# MaxClockRate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**

Specifies the largest base clock rate (in $10^{-7}$ seconds) that the clock discipline allows. The default value is 100288 plus the time interval of the system clock interrupt, which varies by computer.

**Activation Method**

To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Caution**

- Increasing this value increases the rate at which time convergence is accomplished, but also increases the potential for instability.

**See Also**

**Concepts**

MinClockRate

# Config\MaxNegPhaseCorrection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF (*seconds*) | 54000 seconds (15 hours) for a client or member server; 0xFFFFFFFF for a domain controller |

**Description**
Specifies the time, in seconds, of the largest negative time correction that the Windows Time service is allowed to make. If the service determines that a change larger than this is required, then the service logs an event instead.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Tip**

- The smaller the value of this entry, the smaller the impact of erroneous time data. Since erroneous time data could be supplied maliciously, there is greater security in having a smaller value. However, a smaller value might also prevent the Windows Time service from making a larger but legitimate correction to the time.

# Config\MaxPollInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 15 |

**Description**

Specifies the longest interval (in units of $2^n$ seconds, where *n* is the value of this entry) that is allowed for system polling. While the system does not request samples less frequently than this, a provider may refuse to produce samples when requested to do so.

**Activation Method**

To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

### Note

- The time service itself is considered unsynchronized after 1.5 times the number of seconds specified by this entry have elapsed. The Network Time Protocol specifies that the maximum clock age is 86,400 seconds, so if the value of this entry is greater than 15, then peers will eventually ignore this server.

### Caution

- The default time providers installed in Windows Server 2003 do not synchronize at an interval greater than 1024 seconds.

**See Also**

**Concepts**

Config\MinPollInterval

# Config\MaxPosPhaseCorrection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 54000 seconds (15 hours) for a client or member server; 0xFFFFFFFF for a domain controller |

**Description**
Specifies the largest positive time correction, in seconds, that the Windows Time service is allowed to make. If the service determines that a change larger than this is required, then the service logs an event instead.

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Tip**

- The smaller the value of this entry, the smaller the impact of erroneous time data. Since erroneous time data could be supplied maliciously, there is greater security in having a smaller value. However, a smaller value might also prevent the Windows Time service from making a larger but legitimate correction to the time.

# MinClockRate

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | |

**Description**

Specifies the smallest base clock rate (in $10^{-7}$ seconds) that the clock discipline allows. The default value is 100288 plus the time interval of the system clock interrupt, which varies by computer.

**Activation Method**

To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Caution**

- Increasing this value increases the rate at which time convergence is accomplished, but also increases the potential for instability.

**See Also**

**Concepts**

MaxClockRate

# Config\MinPollInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 6 (*domain controllers*) 10 (*member servers and standalone computers*) |

**Description**

Specifies the shortest interval (in units of $2^n$ seconds, where *n* is the value of this entry) that is allowed for system polling. While the system does not request samples more frequently than this, a provider may produce samples whenever it wants.

**Activation Method**

To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

### Note

- The lower this value, the more network accesses the computer must make. This value should be lower on domain controllers, where clock accuracy is more important than efficiency.

### Caution

- The default providers installed with Windows Server 2003 do not synchronize more quickly than every 16 seconds, regardless of the value of this entry.

**See Also**
**Concepts**
Config\MaxPollInterval

# Config\PollAdjustFactor

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 5 |

**Description**
Controls the decision to increase or decrease the system-recommended poll interval. Larger values mean that a smaller amount of error will cause the poll interval to be decreased.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

### Note

- The time service increases the poll interval when the computer clock proves stable for several polls.

### Tip

- Consider increasing the value of this entry if the time service is synchronizing under poor network conditions. Under such conditions, the time service might never achieve great enough stability to increase the poll interval, and the service will continue to poll at the minimum poll interval. Shorter poll intervals create more network traffic.

# Config\SpikeWatchPeriod

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 900 |

**Description**
Specifies the amount of time, in seconds, that a suspicious offset must persist before it is accepted as correct.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Tip**

- Decreasing this value causes the time service to respond to time slips more quickly. However, the time service is more likely to regard bad time samples as good, which can lead to less accurate time.

# Config\UpdateInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Config

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | | 100 (*domain controllers*) 30000(*member servers*) 360000(*standalone computers*) |

**Description**
Specifies the number of intervals, .01 seconds in length, between phase correction adjustments.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

### Note

- Increasing the value of this entry increases the efficiency of the time service, but makes the time service less able to account for clock drift. The value should be set low on domain controllers, which must maintain time as accurately as possible, and increased on computers for which efficiency is a higher priority.

# W32Time\Parameters

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# Parameters\NtpServer

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *DNS name or IP address*[,0x1 \| 0x2 \| 0x3] | time.windows.com,0x1 |

**Description**

Stores a list of time sources that this computer uses to synchronize its time. The time sources appear either as a DNS name or an IP address of the time source server. The list is space-delimited, with one address per line.

Each time source may also be followed by a comma and a DWORD containing flags for special handling of that time source. If the flags are not specified, they default to 0x0. The available flags are:

| Value | Meaning |
|---|---|
| 0x1 | Instead of following the NTP specification, wait for the interval specified in the **SpecialPollInterval** entry before attempting to recontact this time source. Setting this flag decreases network usage, but it also decreases accuracy. |
| 0x2 | Use this time source only as a fallback. If all time sources that are not fallbacks have failed, then the system selects one fallback time source at random and uses it. |
| 0x4 | Set the local computer to operate in symmetric active mode in the association with this source. |
| 0x8 | Set the local computer to operate in client mode in the association with this source. |

**Change Method**

You can change the value of this entry by using the **/setsntp** parameter of the **net time** command at the command line. For more information about **net time**, see Help and Support Center for Windows Server 2003.

# Parameters\Type Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | NoSync \| NTP \| Nt5DS \| AllSync | Nt5DS for computers joined to a domain; NTP for others |

**Description**

Specifies the type of time sources, if any, to which the computer synchronizes its time.

| Value | Meaning |
|---|---|
| NoSync | Does not synchronize time. |
| NTP | Synchronizes time to the time sources specified in the Parameters\NtpServer entry. |
| Nt5DS | Synchronizes time to the domain hierarchy. |
| AllSync | Uses all synchronization mechanisms available. |

**Change Method**

To change the value of this entry, use the command-line tool w32tm.exe (specifically, the command **w32tm /config /syncfromflags**). For more information about w32tm.exe, see Windows Server 2003 Help and Support Center.

**Activation Method**

To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

# W32Time\TimeProviders

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# TimeProviders\NtpClient

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# AllowNonstandardModeCombinations Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the system responds to Network Time Protocol (NTP) mode combinations that the NTP specification can understand, even if the combinations are not standard.

| Value | Meaning |
|---|---|
| 0 | The system does not respond to non-standard mode combinations. |
| 1 | The system responds to non-standard mode combinations, if the NTP specification can understand them. |

**Activation Method**
To make changes to the value of this entry effective, from the command prompt, type:

w32tm /config /update

### Note

- If this entry contains a value other than 0 or 1, the system behaves as though the value is 1.

**See Also**
**Concepts**
NtpServer\AllowNonstandardModeCombinations

# CompatibilityFlags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1 \| 0x2 \| 0x40000000 \| 0x80000000 | 0x80000000 |

**Description**

Controls flags for behaviors that are needed for compatibility.

| Value | Meaning |
|---|---|
| 0x1 | The system ignores maximum error bounds (dispersion values), with the result that it cannot determine which of several time sources is the most accurate. |
| 0x2 | The system accepts a time stamp even if the time at which it was transmitted is earlier than the time stamp. |
| 0x40000000 | Reserved. Do not set this flag. |
| 0x80000000 | The system automatically detects Windows 2000 time sources and sets the compatibility flags appropriately. |

**Activation Method**

To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

# NtpClient\CrossSiteSyncFlags

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 \| 2 | 2 |

**Description**
Specifies whether the time service can synchronize with a partner that is outside the machine's site.

| Value | Meaning |
|---|---|
| 0 | The time service cannot synchronize with a partner that is outside the computer's site. |
| 1 | The time service can synchronize only with the primary domain controller. |
| 2 | The time service can synchronize with a partner that is outside the computer's site. |

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

# DllName Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Full path of a DLL file* | C:\WINDOWS\System32\w32time.dll |

**Description**
Specifies the dynamic-link library (DLL) file that contains the time-service provider.

**Change Method**
Changing the file name or file path of the DLL, or changing the value of this entry, is not recommended.

**Activation Method**
You must restart the computer for changes to the value of this entry to take effect.

**See Also**
**Concepts**
NtpServer\DllName

# NtpClient\Enabled Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the time-service provider should be started.

| Value | Meaning |
|---|---|
| 0 | The time-service provider should not be started. |
| 1 | The time-service provider should be started. |

**Activation Method**
To make changes to the value of this entry effective, from the command prompt type:

w32tm /config /update

### Note

- You might change the value of this entry if you install a third-party time provider to provide time samples, and you need to prevent the Windows time provider from supplying time samples that might interfere.

**See Also**
**Concepts**
NtpServer\Enabled Subkey

# NtpClient\EventLogFlags Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the time service client logs an event when its ability to reach the peer computer changes. The *peer computer* is the computer to which the client is synchronizing its time.

| Value | Meaning |
|---|---|
| 0 | Do not log an event if the reachability of the peer computer changes. |
| 1 | Log an event if the reachability of the peer computer changes. |

**Change Method**
You can change this entry only by editing the registry.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Note**

- Configuring the Windows Time service to create a log entry each time it cannot reach a peer might be useful for troubleshooting. However, entries would also be added to the log for transient failures, which you might find undesirable, since they are not likely to represent a meaningful failure.

# InputProvider Entry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the provider is an *input provider*.

An *input provider* acquires time samples and provides them to the time service so that it can synchronize the local clock.

| Value | Meaning |
|---|---|
| 0 | The provider is not an input provider. |
| 1 | The provider is an input provider. |

**Change Method**
Changing the value of this entry is not recommended, though you can change it by editing the registry.

**Activation Method**
You must stop and restart the time service for changes to this entry to take effect.

**See Also**
**Concepts**
NtpServer\InputProvider

# NtpClient\ResolvePeerBackoffMaxTimes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF | 7 |

**Description**
Specifies the number of times that the time service doubles the length of the interval between attempts to find a peer computer. If contact with a particular peer continues to fail, then the time service continues to reattempt contact, but it stops doubling the wait interval if the number of reattempted contacts exceeds the value of this entry.

The initial wait interval is stored in the NtpClient\ResolvePeerBackoffMinutes entry.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Note**

- Consider setting this entry to a lower value if you want to discover failures more quickly, though doing so will increase network traffic and the number of entries in the event log.

# NtpClient\ResolvePeerBackoffMinutes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–0xFFFFFFFF *minutes* | 15 *(15 minutes)* |

**Description**
Specifies the initial length of the interval that the time service waits between failing to receive a response from a peer computer and reattempting to contact that peer.

If contact with a particular peer continues to fail, then the time service continues to reattempt contact; for each reattempt, it doubles the length of the previous interval, until the number of reattempts matches the value specified by the NtpClient\ResolvePeerBackoffMaxTimes entry. If the number of reattempts to contact a peer exceeds the value of **ResolvePeerBackoffMaxTimes**, then the time service continues to reattempt contacting that peer on an interval equal to the last of the intervals that it was allowed to double.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Note**

- Consider setting this entry to a lower value if you want to discover failures more quickly, though doing so will increase network traffic and the number of entries in the event log.

# TimeProviders\NtpServer Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders

**Description**
This subkey contains documentation for registry elements that are listed in the table of contents under this subkey.

# NtpServer\AllowNonstandardModeCombinations

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the system responds to Network Time Protocol (NTP) mode combinations that the NTP specification can understand, even if the combinations are not standard.

| Value | Meaning |
|---|---|
| 0 | The system does not respond to non-standard mode combinations. |
| 1 | The system responds to non-standard mode combinations, if the NTP specification can understand them. |

**Activation Method**
To make changes to the value of this entry effective, from the command prompt type:

w32tm /config /update

### Note

- If this entry contains a value other than 0 or 1, the system behaves as though the value is 1.

**See Also**
**Concepts**
AllowNonstandardModeCombinations Entry

# NtpServer\DllName

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Full path of a DLL file* | C:\WINDOWS\System32\w32time.dll |

**Description**
Specifies the dynamic-link library (DLL) file that contains the time-service provider.

**Change Method**
Changing the file name or file path of the DLL, or changing the value of this entry, is not recommended.

**Activation Method**
You must restart the computer for changes to the value of this entry to take effect.

**See Also**
**Concepts**
DllName Entry

# NtpServer\Enabled Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**
Specifies whether the time-service provider should be started.

| Value | Meaning |
|---|---|
| 0 | The time-service provider should not be started. |
| 1 | The time-service provider should be started. |

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**Note**

- If you do not want this computer to respond to NTP requests, you can set the value of this entry to 0. This prevents the time-service provider from running and, therefore, from responding to requests.

**See Also**
**Concepts**
NtpClient\Enabled Subkey

# NtpServer\InputProvider

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the provider is an *input provider*.An *input provider* acquires time samples and provides them to the time service so that it can synchronize the local clock.

| Value | Meaning |
|---|---|
| 0 | The provider is not an input provider. |
| 1 | The provider is an input provider. |

**Change Method**
Changing the value of this entry is not recommended.

**Activation Method**
To make changes to this entry effective, at the command line, type:

**w32tm /config /update**

**See Also**
**Concepts**
InputProvider Entry

# WINS Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services

**Description**
The WINS subkey stores configuration data for the Windows Internet Name Service (WINS).

WINS associates IP addresses to more familiar DNS and NetBios computer names. WINS enables users to find computer resources by their familiar names, instead of requiring them to use IP addresses.

# WINS\Parameters Subkey

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS

**Description**
The Parameters subkey stores general configuration data for the Windows Internet Name Service (WINS), that is, configuration data that is not related to replicating the WINS database.

WINS associates IP addresses to more familiar DNS and NetBios computer names. WINS enables users to find computer resources by their familiar names, instead of requiring them to use IP addresses.

# BackupDirPath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Path* | (None) |

**Description**
Specifies the directory where the backup of the Windows Internet Name Service (WINS) database is performed.

This entry appears in the registry only when you use the WINS console or the WINS netshell extension to establish a database backup path.

**Change Method**
To change the value of this entry, use the WINS console.

### Caution

- If you administer Windows NT 3.51 servers from a computer running a later version of Windows, the data type of this entry might be set incorrectly to REG_EXPAND_SZ. Contact Microsoft Product Support Services for assistance.

# BurstHandling

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether the Windows Internet Name Service (WINS) server uses burst handling. When using burst handling, the WINS server sends success responses to the clients whose requests are dropped.

If burst handling is enabled, then by default the WINS server stores 500 name registration and refresh queries in its queue before dropping queries.

| Value | Meaning |
|---|---|
| 0 | Disables burst handling. The WINS server does not send success responses to clients whose requests are dropped. |
| 1 | Enables burst handling. The WINS server sends success responses to clients whose requests are dropped. |

Burst handling attains temporary steady state on the WINS server when the WINS server is started with an empty database or when many WINS clients log on for the first time. The resulting name registration and name refresh operations cause a large amount of traffic.

Because success responses have small Time to Live values (TTLs), they slow the refresh rate of new WINS clients and regulate the burst of WINS client traffic. As a result, a steady state is reached much more quickly.

**Change Method**

To change the value of this entry, use the WINS console. In WINS server properties, click the **Advanced** tab. This entry corresponds to the **Enable burst handling** check box.

# DbFileNm

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %SystemRoot%\System32\Wins\wins.mdb |

**Description**
Stores the full path to the Windows Internet Name Service (WINS) database file.

**Change Method**
To change the value of this entry, use the WINS console. In WINS server properties, click the **Advanced** tab. This entry corresponds to the **Database path** box.

### Note

- This entry can also have the REG_SZ data type.

# DoBackupOnTerm

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the Windows Internet Name Service (WINS) database is backed up automatically when the WINS service stops.

| Value | Meaning |
|---|---|
| 0 | The WINS database is not backed up automatically when the WINS service stops. |
| 1 | The WINS database is backed up automatically when the WINS service stops. |

**Change Method**
To change the value of this entry, use the WINS console. In WINS server properties, click the **General** tab. This entry corresponds to the **Back up database during server shutdown** option.

# DoStaticDataInit

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether the Windows Internet Name Service (WINS) server initializes its database.

| Value | Meaning |
|---|---|
| 0 | The WINS server does not initialize its database. |
| 1 | The WINS server initializes its database. |

If the value of this entry is 1, the WINS server initializes its database when the WINS process starts and when an entry in the WINS\Parameters Subkey or Parameters\DataFiles subkeys changes value (unless the change sets the value of this entry to 0). The database is initialized using records specified in the Datafiles subkey.

# LogDetailedEvents

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies the level of detail of events that the Windows Internet Name Service (WINS) server writes to the System Log in Event Viewer.

| Value | Meaning |
|---|---|
| 0 | Errors only. |
| 1 | Errors, warnings, and informational messages. |

**Change Method**
To change the value of this entry, use the WINS console. On the **Action** menu, click **Properties**, and then click the **Advanced** tab. This entry corresponds to the **Log detailed events to the Windows event log** option.

# Parameters\LogFilePath

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_EXPAND_SZ | *Path* | %SystemRoot%\System32\Wins |

**Description**
Specifies whether Windows Internet Name Service (WINS) automatically backs up its database, and specifies the location where the backups are stored.

If you enter a path in the value of this entry, then WINS automatically backs up its database every three hours and stores the backup in the file named Jet.log in the specified path. If you delete this entry or its value, then WINS does not back up the database.

**Change Method**
To change the value of this entry, use the WINS console. On the **Action** menu, click **Properties**, and then click the **Advanced** tab. This entry corresponds to the **Log detailed events to Windows event log** box.

### Note

- This entry can have a data type of REG_SZ or REG_EXPAND_SZ.

# LoggingOn

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 1 |

**Description**

Specifies whether database changes are recorded in a log file.

| Value | Meaning |
|---|---|
| 0 | Database changes are not logged. |
| 1 | Database changes are logged in Jet.log. |

# McastIntvl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x960–0xFFFFFFFE (*seconds*) \| 0xFFFFFFFF | 0x960 (*2,400 seconds = 40 minutes*) |

**Description**
Specifies how often the Windows Internet Name Service (WINS) server sends a multicast message to announce its presence to other WINS servers.

| Value | Meaning |
|---|---|
| 0x960–0xFFFFFFFE seconds | WINS sends multicast messages at this interval. |
| 0xFFFFFFFF | WINS does not send multicast messages. |

**See Also**
**Concepts**
McastTtl
UseSelfFndPnrs

# McastTtl

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1–32 (*hops*) | 2 |

**Description**
Specifies the maximum number of *hops* (steps between subnets) that a Windows Internet Name Service (WINS) multicast announcement can perform. If the multicast announcement attempts to perform more hops than are permitted by the value of this entry, then WINS returns an Internet Control Message Protocol (ICMP) Time Exceeded error.

**Change Method**
To change the value of this entry, use the WINS console. Right-click the **Replication Partners** node, click **Properties**, and then click the **Advanced** tab. This entry corresponds to the **Multicast Time to Live (TTL)** box.

**See Also**
**Concepts**
McastIntvl
UseSelfFndPnrs

# MigrateOn

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies how the Windows Internet Name Service (WINS) server regards static database records that conflict with a new registration or replica.

| Value | Meaning |
|---|---|
| 0 | If unique and multihomed static records conflict with a new registration or replica, then they are still regarded as static records. |
| 1 | If unique and multihomed static records conflict with a new registration or replica, then they are treated as dynamic records. |

**Change Method**
To change the value of this entry, use the WINS console. Right-click the **Replication Partners** node, click **Properties**, and then, from the **General** tab, select the **Overwrite unique static mappings at this server (migrate on)** box.

# PriorityClassHigh

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies the priority class of the Windows Internet Name Service (WINS) server process. Setting the value of this entry to 1 ensures that other processes, such as electronic mail and other applications, do not preempt WINS.

| Value | Meaning |
|---|---|
| 0 | The WINS server runs in a normal priority class. |
| 1 | WINS runs in high-priority class. |

# UseSelfFndPnrs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**
Specifies whether the Windows Internet Name Service (WINS) server detects or ignores multicasts from other WINS servers on the network.

| Value | Meaning |
|---|---|
| 0 | The WINS server ignores multicasts from other WINS servers. |
| 1 | The WINS server detects multicasts from other WINS servers. |

If network routers support multicasting, then the WINS server automatically finds other WINS servers on the network by listening for their multicasts (as defined by RFC 1112, Host extensions for IP multicasting) to the IP address 224.0.1.24. Any WINS servers found by accepting WINS multicasts are automatically configured as push and pull replication partners, with pull replication set to occur every two hours.

If network routers do not support multicasting, then even if the value of this entry is 1, the WINS server finds WINS servers only on its subnet.

By default, WINS maintains partnership data that is discovered through multicasting. A WINS server that is discovered through multicasting is automatically removed as a replication partner when the WINS server shuts down successfully. However, if the WINS console has been used to change or view the found WINS server information, then the partnership data is not removed.

**Tip**

- To disable WINS multicasting, set the value of the McastIntvl entry to 0xFFFFFFFF.

**See Also**
**Concepts**
McastTtl
McastIntvl

# VerifyInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x1FA400–0x7FFFFFFF (*seconds*) | 0x1FA400 (*576 hours = 24 days*), or 3 times the value of the entry **TombstoneInterval**, whichever is greater. |

**Description**

Specifies how long the Windows Internet Name Service (WINS) server can retain a name that it does not own before it must verify that the name is still active.

**Change Method**

To change the value of this entry, use the **Intervals** tab on the WINS ***Server-name*** **Properties** dialog box. This entry corresponds to the **Verification interval** field.

# VersCounterStartVal_HighWord

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x0–7FFFFFFF | 0x0 |

**Description**
Stores the highest numbered *version identifier* used in the Windows Internet Name Service (WINS) database. Each record in the database has a *version identifier*, which is a counter used to determine which WINS database entries must be updated during replication.

### Note

- The *version identifier* is a number that identifies a unique WINS record; it does not refer to the version number of the WINS database program.

**See Also**
**Concepts**
VersCounterStartVal_LowWord

# VersCounterStartVal_LowWord

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

| Data type | Range | Default value |
|---|---|---|
| | 0x0–0x7FFFFFFF | 0x0 |

**Description**
Stores the lowest *version identifier* used in the Windows Internet Name Service (WINS) database. Each record in the WINS database has a *version identifier*, which is a counter used to determine which WINS database entries must be updated during replication. WINS uses this entry to determine the starting record in the WINS database.

### Note

- The *version identifier* is a number that identifies a unique WINS record; it does not refer to the version number of the WINS database program.

**See Also**
**Concepts**
VersCounterStartVal_HighWord

# ConsistencyCheck

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

**Description**
The ConsistencyCheck subkey is an optional subkey that you can add to the registry to direct the Windows Internet Name Service (WINS) to perform periodic database consistency checks. Typically, consistency checks are not performed because they consume substantial resources on both the WINS server and the network.

If you add this subkey, then you should set the values of its entries to minimize the load on the network due to consistency checking. To determine optimal values for these entries, consider the network topology and the number of WINS servers, WINS clients, and wide area network (WAN)/local area network (LAN) links between WINS servers.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To enable periodic automated consistency checking, use the WINS console. In WINS server properties, click the **Database Verification** tab, click the **Verify database consistency check** box, and use the options to configure consistency checking.

**Caution**

- Consistency checking uses significant network resources, and it can consume substantial processor time on WINS servers.

# MaxRecsAtATime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters\ConsistencyCheck

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 1,000–*Total number of records* | 30,000 |

**Description**
Specifies the maximum number of records replicated in each consistency check cycle.

Windows Internet Name Service (WINS) performs consistency checks on the records of one *owner WINS server* at a time. (The *owner WINS server* is the WINS server from which the record was replicated.) After checking the owner WINS server, the local WINS server checks the value of this entry. If the number of records meets or exceeds this value, then the WINS server stops checking records.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**
To enable periodic automated consistency checking, use the WINS console. In the WINS ***Server-name*** **Properties** dialog box, click the **Database Verification** tab, select the **Verify database consistency** check box, and use the options to configure consistency checking.

### Caution

- Consistency checking uses significant network resources, and it can consume substantial processor time on WINS servers.

# SpTime

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters\ConsistencyCheck

| Data type | Range | Default value |
|---|---|---|
| REG_SZ | *Time in hh:mm:ss format* | 2:00:00 (2:00 a.m.) |

**Description**

Specifies the time of day when the first Windows Internet Name Service (WINS) server database consistency check occurs. Thereafter, the WINS database is checked for consistency at the frequency specified in the value of the TimeInterval entry.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**

To change the value of this entry, use the WINS console. In the WINS ***Server-name*** **Properties** dialog box, click the **Database Verification** tab. This entry corresponds to the option **Begin verifying at**.

# TimeInterval

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters\ConsistencyCheck

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0x5460–0xFFFFFFFF (*seconds*) (*6 hours or more*) | 0x15180 (*86,400 seconds = 1 day*) |

**Description**
Specifies how often the Windows Internet Name Service (WINS) server performs database consistency checks.

**Change Method**
To change the value of this entry, use the WINS console. In the WINS ***Server-name*** **Properties** dialog box, click the **Database Verification** tab. This entry corresponds to the **Verify database consistency every [ n ] hours** option.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# UseRplPnrs

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters\ConsistencyCheck

| Data type | Range | Default value |
|---|---|---|
| REG_DWORD | 0 \| 1 | 0 |

**Description**

Specifies whether the Windows Internet Name Service (WINS) contacts only its replication partners when performing consistency checks. Limiting the scope of a consistency check to replication partners is recommended for networks with a large number of WINS servers.

| Value | Meaning |
|---|---|
| 0 | WINS contacts the *owner WINS server* (the WINS server from which the record was replicated). |
| 1 | WINS contacts only its replication partners. If the owner of a record is a WINS replication partner, then WINS contacts the owner WINS server. If the owner is not a replication partner, then WINS contacts a WINS server selected randomly from its list of pull partners. |

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

**Change Method**

To change the value of this entry, use the WINS console. In the WINS ***Server-name*** **Properties** dialog box, click the **Database Verification** tab. This entry corresponds to the **Owner servers** or **Randomly selected partners** options in the **Verify against** section.

# DataFiles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Parameters

**Description**
The DataFiles subkey contains entries representing the files that the Windows Internet Name Service (WINS) server uses to initialize and reinitialize its local database. These files contain static records.

By default, this subkey does not appear in the registry. However, you can add it and its entries to the registry to direct WINS to add static records in the specified data files to its database.

The entries in this subkey have the following format:

| Entry name | Data type | Value |
|---|---|---|
| [Path\] File name | REG_SZ \| REG_EXPAND_SZ | (Blank) |

The path is optional. If you do not enter a path, then WINS assumes the files are located in the directory from which the WINS service runs. The file can have any name and file name extension.

This entry does not exist in the registry by default. You can add it by using the registry editor Regedit.exe.

# Partners

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS

**Description**
The Partners subkey contains entries that specify the properties of the Windows Internet Name Service (WINS) server and its replication partners.

This subkey contains two subkeys, Pull and Push, which store entries for pull and push operations of the WINS server.

**Change Method**
To display and change these values, use the WINS console whenever possible. Changes to these values are effective immediately; you do not have to restart WINS or Windows Server 2003.

**Tip**

To see the changes reflected in your registry editor, refresh the display.

**See Also**
**Concepts**
Push
Pull

# Pull

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Partners

**Description**
The Pull subkey contains entries that specify how and when replicas are pulled from pull partners of the Windows Internet Name Service (WINS) server. These entries are stored at the root level of the Pull subkey and apply to all pull operations.

In addition to these entries, this subkey contains subkeys representing each pull partner from which the WINS server can pull replicas. These subkeys, which are named by the IP address of the pull partner, store data that applies only to that pull partner. (See the IP address of pull partner subkey.)

**Change Method**
To display and change these values, use the WINS console. In the console tree, click **Replication Partners**, on the **Action** menu, click **Properties**, and then click the **Pull Replication** tab.

**Activation Method**
Changes to the entries in this subkey are effective immediately; you do not have to restart WINS or Windows. To see the changes reflected in your registry editor, refresh the display.

**Tip**

To see the changes reflected in your registry editor, refresh the display.

**See Also**
**Concepts**
Push
IP address of pull partner

# IP address of pull partner

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Partners\Pull

**Description**

The *IP address of pull partner* subkey is a variable representing the IP address subkeys in the WINS\Partners\Pull subkey. It does not actually appear in the registry. This variable subkey displays the entries and subkeys that are common to subkeys in the Adapters Subkey subkey.

The Pull subkey contains one subkey for each pull partner from which the Windows Internet Name Service (WINS) server can pull replicas. The subkeys are named for the IP address of the pull partner. They contain data that is configured separately for each pull partner and that applies only to that pull partner.

**Change Method**

To display and change these values, use the WINS console whenever possible. Changes to these values are effective immediately; you do not have to restart WINS or Windows Server 2003. To see the changes reflected in your registry editor, refresh the display.

# Push

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM\CurrentControlSet\Services\WINS\Partners

**Description**
The Push subkey contains entries that specify how and when the Windows Internet Name Service (WINS) server sends update notification messages to its push partners. These entries are stored at the root level of the Push subkey and apply to all push operations.

In addition to these entries, this subkey contains subkeys representing each push partner to which the WINS server sends update notification messages. These subkeys, which are named by the IP address of the push partner, store data that applies only to that push partner. (See the IP address of push partner subkey.)

**Change Method**
To display and change these values, use the WINS console. In the console tree, click **Replication Partners**, on the **Action** menu, click **Properties**, and then click the **Push Replication** tab.

**Activation Method**
Changes to the entries in this subkey are effective immediately; you do not have to restart WINS or Windows Server 2003. To see the changes reflected in your registry editor, refresh the display.

**See Also**
**Concepts**
Pull
IP address of pull partner

# Disk

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM

**Description**
Windows Server 2003 does not use this subkey. The subkey remains in the registry to support programs designed for Windows NT 4.0 and earlier.

In Windows NT 4.0 and earlier, this subkey contained information about all physical disk drives, logical partitions of disk drives, CD-ROM drives, and network drive mappings on the computer. Windows Server 2003 eliminates this subkey and stores similar information in the MountedDevices subkey.

**Caution**

- Do not delete this subkey or its entries from the registry, and do not change the values in the entries. These changes can cause serious, unexpected results.

# MountedDevices

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM

**Description**
The MountedDevices subkey stores the database of mounted devices for the NTFS file system. This database associates persistent volume names with unique internal indentifiers for volumes. The internal identifiers often include the disk signatures of the volumes. This pairing enables NTFS to maintain the identity and persistent name of a volume, even when its variable device name changes.

NTFS uses this subkey to support mounted drives on all types of volumes, including removable drives and CD-ROM drives.

**Note**

- This subkey replaces the Disk subkey in Windows NT 4.0 and earlier. However, its contents and use are changed for Windows Server 2003.

# Select

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

HKLM\SYSTEM

**Description**
The Select subkey contains entries that specify which control set is used to start the system. The entries in this subkey record which control set was used to start the system this time, which control set will be used to start it next time, and which control sets have failed to start the system in the past.

**Caution**

- Do not change the values of the entries in this subkey. This subkey is maintained by Windows Server 2003 and should not be changed by users.

# Deployment Kit Compact Disc

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following contents are included on the *Windows Server 2003 Deployment Kit* companion CD:

- **Windows Server 2003 Deployment Kit.** A searchable online version of the *Windows Server 2003 Deployment Kit*.
- **Resource Kit Tools for Windows Server 2003.** A collection of tools included with the Windows Deployment and Resource Kits that can help you deploy, configure, maintain, and troubleshoot Windows Server 2003.
- **Resource Kit Registry Reference for Windows Server 2003.** A searchable online reference providing detailed descriptions of the Windows Server 2003 registry, including many entries that cannot be edited by using Windows Server 2003 tools or programming interfaces.
- **Resource Kit Performance Counters Reference for Windows Server 2003.** A searchable online reference describing what each performance counter monitors. You can use performance counters to diagnose problems or detect bottlenecks in your system.
- **Deploying Internet Information Services (IIS) 6.0 & Migration Tools.** A searchable online version of *Deploying Internet Information Services (IIS) 6.0* and tools that you can use to migrate to IIS 6.0.
- **Job Aids for the Windows Server 2003 Deployment Kit.** Worksheets and resources that can help you create your deployment plan for Windows Server 2003.
- **Windows Server 2003 Help.** The searchable Help file included with the Windows Server 2003 operating system containing technical content for the IT professional, which can be installed on Microsoft® Windows® XP Professional.
- **Microsoft Office Viewers.** Viewers you can install on your computer if you do not have Microsoft® Office, which allow you to see worksheets and resources on the *Windows Server 2003 Deployment Kit* companion CD.
- **CD-ROM Release Notes.** Late breaking information about the contents of the *Windows Server 2003 Deployment Kit* companion CD.
- **Links to Microsoft Press.** Links to the Microsoft Press Support site, which you can search for Knowledge Base articles, and to the Microsoft Press product registration site, which you can use to register this book online.

# Deployment Guide Document Conventions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following art symbols and text conventions are used throughout this book.

**Flowchart Symbols**
Use the following table of symbols as a resource for understanding the flowcharts included in this guide.

| Symbol | Meaning | Symbol | Meaning |
|---|---|---|---|
| | Step or component process | | Data stored to a database |
| | Predefined process or subroutine | | Flowchart beginning or end |
| | Decision point | 1 | Intra-chart connector: Flow continues to next page |
| | Output to a document or input from a document | 1 | Intra-chart connector: Flow continues from previous page |
| | Data transfer to a file on disk | A | Inter-chart connector: Indicates an exit point to another flowchart |
| | Data transfer to a data store | B | Inter-chart connector: Indicates an entry point from another flowchart |

**Art Symbols**
Use the following table of the art symbols as a resource for understanding the graphics included in this guide.

| Symbol | Meaning | Symbol | Meaning |
|---|---|---|---|
| | Workstation | | Macintosh client |
| | Portable computer | | Tablet computer |
| | Terminal | | Cellular phone |
| | Portable digital assistant (PDA) | | Document |
| | File folder | | E-mail |
| | Chart | | Wireless network adapter |
| | Modem | | Video camera |
| | Network adapter | | Digital camera |
| | Facsimile | | Printer |
| | Telephone | | Scanner |
| | Hard disk | | Tape drive |
| | Database | | Tape |
| | Compact disc | | Security key |
| | Digital certificate | | Padlock |

| | | | |
|---|---|---|---|
| | Padlock | | Uninterruptable power supply |
| | Access token | | Hub |
| | Modem bank | | Automated library |
| | Windows NT–based server | | Generic server |
| | Mainframe computer | | Host |
| | Server farm | | Clustered servers |
| | Router | | Switch |
| | Shadowed router | | Windows 2000–based router |
| | Data jack | | Input/output (I/O) filter |
| | Firewall | | Tunnel |
| Internet | Internet | Intranet | An intranet |
| | Transceiver | | Script |
| | Interface | | Packets |
| | Process or communication failure | “ · ” | DNS root |
| | Directory tree root | | Root |
| | Organization | | Organizational unit |
| | Common name | | Generic node |
| | Active Directory domain | | User group |
| Windows 2000 Domain | Windows 2000 domain | | Site *or* Windows NT 4.0 domain |
| | Active Directory™ | | |

**Reader Alert Conventions**

Reader alerts are used throughout this guide to notify you of both supplementary and essential information. The following table explains the meaning of each alert.

| Reader Alert | Meaning |
|---|---|
| **Tip** | Alerts you to supplementary information that is not essential to the completion of the task at hand. |
| **Note** | Alerts you to supplementary information. |
| **Important** | Alerts you to supplementary information that is essential to the completion of a task. |
| **Caution** | Alerts you to possible data loss, breaches of security, or other more serious problems. |

| | |
|---|---|
| **Warning** | Alerts you that failure to take or avoid a specific action might result in physical harm to you or to the hardware. |

**Command-line Style Conventions**

The following style conventions are used in documenting scripting and command-line tasks throughout this book.

| Element | Meaning |
|---|---|
| **bold font** | Characters that you type exactly as shown, including commands and parameters. User interface elements are also bold. |
| *Italic font* | Variables for which you supply a specific value. For example, *Filename.ext* can refer to any valid file name. |
| `Monospace font` | Code samples. |
| **`Command`** | Command that is typed at the command prompt. |
| `Syntax` | Syntax of script elements. |
| `Output` | Output from running a script. |

# Support Policy

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft does not support the software supplied in the *Windows Server 2003 Deployment Kit*. Microsoft does not guarantee the performance of the scripting examples, job aids, or tools, bug fixes for the tools, or response times for answering questions. However, we do provide a way for customers who purchase the *Windows Server 2003 Deployment Kit* to report any problems with the software and receive feedback for such issues. You can do this by sending e-mail to rkinput@microsoft.com. This e-mail address is only for issues related to the *Windows Server 2003 Deployment Kit*. For issues related to the Windows Server 2003 operating systems, please refer to the support information included with your product.

# Deployment Guide for the Security Configuration Wizard

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This topic discusses the deployment of the Security Configuration Wizard (SCW) and SCW policies.

Security Configuration Wizard (SCW) is a tool for reducing the attack surface of computers running Microsoft® Windows Server™ 2003 with Service Pack 1 (SP1). It determines the minimum functionality required for a server's role or roles, and disables functionality that is not required. Specifically, SCW helps you author and deploy a security policy that:

- Disables unneeded services.
- Blocks unused ports.
- Allows additional address or security restrictions for ports that are left open.
- Prohibits unnecessary Internet Information Services (IIS) Web extensions, if applicable.
- Reduces protocol exposure to server message block (SMB), LAN Manager, and Lightweight Directory Access Protocol (LDAP).
- Defines a high signal-to-noise audit policy.

SCW guides you through the process of creating, editing, applying, or rolling back a security policy based on the selected roles of the server.

This deployment guide guides you through the process of deploying SCW and SCW policies to your computers running Windows Server 2003 with SP1.

**Requirements for Installing and Running SCW**
SCW is included with Windows Server 2003 SP1, and works only with computers running Windows Server 2003 with SP1. It is not for use with Microsoft Small Business Server or client operating systems like Windows XP Professional.

**SCW Deployment Overview**
Before you deploy SCW to your servers, you should become familiar with the SCW user interface and the Scwcmd.exe command-line tool.

Once you are familiar with SCW, perform a test deployment of SCW and SCW policies. Make the test environment as similar as practicable to the production environment but on a smaller scale. Be sure to include the server types and Active Directory organizational unit (OU) structure that are used in your production environment. When you are confident that SCW policy application is working correctly in the lab, you are ready to move to the production environment.

Install Windows Server 2003 SP1 on all servers and then install SCW on each server. If you have to manage many servers, you can perform an unattended installation of Windows Server 2003 SP1 and SCW in one step.

In an Active Directory environment, your next step is to adjust Active Directory OU structure according to server types.

You also need to complete policy prototyping, where you create policies for each type.

Finally, you deploy the policies by using one of the following three ways:

- SCW user interface
- Scwcmd.exe
- Group Policy

At each step, you will have choices to make, such as:

- What to do on each wizard page
- Whether and how to use security templates
- Whether to use Group Policy to deploy policies

**SCW Deployment Steps**

This section lists the steps you take for deploying SCW.

1. Pre-plan, gather information, lab test, and back up your servers.
2. Install SP1 and SCW.
    - If you have only a few servers to manage, install SP1 manually and locally on each server. Then install SCW on each server.
    - If you do have many servers to manage, do an unattended installation of SP1 and SCW simultaneously with SCW=On in the [Components] section of Unattend.txt. Using Systems Management Server (SMS) with Service Pack 1 is another option for deploying SCW.

**SCW Policy Deployment Steps**

After you have deployed SCW, you are ready to deploy SCW policies. This section lists the steps you take for deploying SCW policies.

1. Pre-plan, gather information, lab test, and back up your servers.
2. Back up your current security configuration.
3. If you apply policy to servers by using Active Directory, arrange server roles in organizational units (OUs).
4. Create a prototype policy per server role.
5. If you need to apply settings that SCW cannot create, attach an .inf security template.
6. Internet Information Services (IIS) settings and rollback capability will be lost if you use Group Policy to deploy the settings. Use Group Policy anyway? If so, create a Group Policy object (GPO) through the **scwcmd transform** command and link the GPO by using Group Policy Management Console (GPMC). This completes your SCW policy deployment.
7. If you have arrived at this step you are not using Group Policy. Are you going to target multiple servers? If so, apply policy as .xml through the **scwcmd configure** command, specifying a list of servers as a parameter. This completes your SCW policy deployment.
8. If you have arrived at this step you are targeting a single server. Apply policy as .xml through the Wizard interface of SCW. This completes your SCW policy deployment.

# Planning the SCW Deployment

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This topic discusses planning for SCW deployment.

Safe use of any technology that affects the settings, configuration, or behavior of multiple computers requires planning before deployment. Safe use of SCW includes lab-testing of tools and policies prior to live rollout into a production environment. This and other best practices for SCW are covered in this section.

After a discussion of additional resources that can help you become familiar with SCW, this section includes guidelines for defining exactly which computers support local installation of SCW and the tools included with SCW and which computers are valid targets for policy that you author by using SCW.

**Review SCW resources**

Additional information about SCW is available at the following sources:

- SCW documentation from the Download Center on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=45183)
    - SCW Quick Start Guide
    - SCW Deployment Guide
    - ExtendingSCW
    - SCW Troubleshooting Guide
- SCW online Help

  Each page of SCW contains a link to online help that is specific to that wizard page. You can view the online help file in its entirety from the command line by **hh scwhelp.chm**. There is also a link to the help file on the desktop once SP1 is installed. The help file is available in SP1 regardless of whether SCW itself is installed.

- Scwcmd.exe command-line tool help

  After SCW is installed, help for the command-line tool can be obtained from the command line by typing scwcmd. This displays the following:

  ```
  Syntax: scwcmd [ analyze | configure | register | rollback | transform | view ]
  ```

  Help for each of these six options is obtained by typing scwcmd option_name.

| **Note** |
| --- |
| The **scwcmd view** command allows the use of a template (.xsl file) to format the display in the SCW viewer. The correct syntax is to have the template first, for example **scwcmd view /s:scwpolicy.xsl /x:main.xml**. The command-line help that you get by typing **scwcmd view** incorrectly shows the **/x:** parameter first. |

**Review requirements for local installation of SCW**

You can install SCW only on computers running Windows Server 2003 with Service Pack 1 (SP1), which includes SCW as an optional Windows component. SCW is not installed by default, but you can install it by clicking **Add/Remove Windows Components** in **Add or Remove Programs** in **Control Panel**. To install SCW, you need to be a member of the Administrators group on the server.

SCW should not be used with Windows Small Business Server 2003 or client operating systems like Windows XP Professional.

**Determine valid targets for applying SCW security policy**

The security policies that are created with SCW should be applied only to servers and groups of servers. They should not be applied to client operating systems such as Windows XP Professional, nor should they be applied to systems running Windows Small Business Server 2003.

**Identify valid targets for SCW policy prototyping**
In the context of SCW, "policy prototyping" means using SCW to author security policy for groups of similarly configured computers by targeting SCW at a prototype computer: a model server whose configuration is representative of every computer in the group. Only computers running Windows Server 2003 with SP1 are valid targets for SCW security policy authoring.

When SCW is installed and run locally but targeted at a remote server for authoring policies, it is best for the remote server to also have SCW installed. This is because SCW includes so-called "helper functions" that inspect the computer. If SCW is installed on the remote computer, the locally-run wizard can make these determinations about the remote computer:

- Whether unknown services are listening on a port.
- Whether approved applications are installed.
- The list of named pipes.
- Whether an IPsec policy is applied

**Review SCW Best Practices**
This section explains how to get the most out of SCW by following best practices.

**Identify and target similar servers.**
SCW helps to reduce the attack surface of servers by creating a security policy that is specifically designed for their specific roles. Administrators can simplify policy authoring and distribution by identifying groups of servers that perform the same, or similar, tasks.

**Author one policy for a group of servers.**
SCW authors a security policy based on the roles and functions performed by a server. Others servers that perform the same, or very similar, functions can be configured with the same security policy. Administrators can use SCW once to author a security policy, save it, and apply it to all servers that perform the job function. A single security policy that contains all desired security settings for a server also simplifies configuration, rollback, and analysis.

For simple configuration and rollback, a single security policy for a computer, or set of computers, is much easier to understand and update than a series of policies. The scwcmd analyze command has a p*:/policyname* parameter that only accepts a single SCW native policy file as a parameter.

**Group similar servers in one organizational unit (OU).**
Grouping servers by OUs in Active Directory domains facilitates the application of security policy through Group Policy. The SCW transform operation can create an unlinked security policy containing GPO in Active Directory. To simplify policy distribution, an administrator can group servers that perform similar functions into a single OU, and link the GPO to the OU.

**Create different policies for different platforms.**
For services or ports specific to 64-bit computers, create the policies on a 64-bit computer. Then deploy these policies to other 64-bit computers only (not 32-bit computers) to ensure the services are properly identified and configured.

It is highly recommended that the prototype server from which the security policy will be created resembles the target servers to be configured as closely as possible. Not only are there options in the wizard that are not shown unless the prototype server has those options, the modifications done to services and other items could cause severe interoperability problems unless the prototype resembles the target.

The security policy might disable any service that was not present on the prototype server when the policy was created. Any service defined in the security configuration database but not present on the prototype server will be disabled. The configured state of a service that is not present on the prototype and that is not in the security configuration database is configurable in the wizard. For example, if the DCOM Server Process Launcher service is listed in the Security Configuration Database, but is not present on the prototype server, the security policy created from the prototype server will set the DCOM Server Process Launcher state to disabled. When you apply the security policy to other servers, the DCOM Server Process Launcher service will be disabled on those servers.

**Test new security policies before deployment.**
The settings configured in the new security policies may cause compatibility issues with applications or services. Therefore,

thoroughly test new security policies in a test environment before applying the policies to production servers. Virtual machines are particularly well suited for this purpose.

# Deploying SCW to multiple servers

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This topic explains how to deploy SCW to multiple servers.

Your options for installing Windows Server 2003 SP1 and SCW on multiple servers include:

- Systems Management Server (SMS) with Service Pack 1 (SP1)
- Unattended installation

**SMS with SP1**

Systems Management Server (SMS) is a full-featured change and configuration management, software deployment, and license-tracking system that is part of the Windows Server System. It is available from Microsoft at an additional cost.

SMS can be used to deploy operating system updates including Windows Server 2003 SP1 and SCW to your servers.

For more information about SMS, see What Is Windows Server System? on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48227).

The Group Policy Object Editor and Group Policy Management Console (GPMC) include the ability to assign patches and software updates, but only for Windows Installer packages. SCW is part of Windows Server 2003 SP1 and is not available as a stand-alone Windows Installer package. Therefore you can not deploy SCW itself by using Group Policy.

SCW policies can, however, be deployed by using Group Policy.

**Unattended installation of SCW**

You can create an answer file that installs Windows Server 2003 SP1 and SCW automatically. To do so, follow these steps:

**To install SCW as part of an unattended installation**

- Before initiating the unattended installation, add the following entry to the [Components] section of Unattend.txt:

  **SCW=On**

**To apply an SCW security policy during the unattended installation, also carry out the following steps.**

1. Create the security policy file on a server that already has SCW installed.
2. Either create a Cmdlines.txt file or modify the existing one so that it has a [Commands] section containing the following line:

   **scwcmd configure /p:SCWPolicy.xml**

3. Copy the Cmdlines.txt file and the previously created policy file (SCWPolicy.xml in this case) to the $OEM$ directory.

For detailed instructions about using unattended installation, see Automating and Customizing Installations in the Microsoft Windows Server 2003 Deployment Kit on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=4748).

# Creating SCW security policy files

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This topic tells how to create SCW security policy files.

Before you can deploy an SCW policy, you need to author it. This section covers using the SCW user interface to author a security policy. SCW security policies are always created from prototype servers of the same type as the servers that will receive the policy.

**To create a security policy**

1. Click **Start**, click **Administrative Tools**, and then click **Security Configuration Wizard**.
2. Read the Welcome page and click **Next.**
3. Select **Create a new security policy**, and then click **Next**.
4. Type the name of the prototype server, and then click **Next**.
5. Wait for the Security Configuration Database to be processed, and then click **Next**.
6. For each of the next five wizard pages, just click **Next**:
   - **Role-Based Service Configuration** page.
   - **Select Server roles** page.
   - **Select Client Features** page.
   - **Select Administration and Other Options** page.
   - **Select Additional Services** page.
7. On the **Handling Unspecified Services** page, click **Next**.
8. For each of the next 20 wizard pages, just click **Next**:
   - **Confirm Service Changes** page.
   - **Network Security** page.
   - **Open Ports and Confirm Applications** page.
   - **Confirm Service Changes** page.
   - **Confirm Port Configuration** page.
   - **Registry Settings** page.
   - **Require SMB Security Signatures** page.
   - **Require LDAP Signing** page.
   - **Outbound Authentication Methods** page.
   - **Outbound Authentication Methods using Domain Accounts** page.
   - **Registry Settings Summary** page.

- **Audit Policy** page.
- **System Audit Policy** page.
- **Audit Policy Summary** page.
- **Internet Information Services** page.
- **Select Web Service Extensions for Dynamic Content** page.
- **Select the Virtual Directories to Retain** page.
- **Prevent Anonymous Users from Accessing Content Files** page.
- **IIS Settings Summary** page.
- **Save Security Policy** page.

9. On the **Security Policy File Name** page, type a name for the prototype policy, and then click **Next**.

   Do not name the security policy by using the name of the prototype computer because scwcmd.exe uses *computername.xml* to save analysis results, and you do not want the security policy to have the same name as the analysis results. That would risk confusion or overwriting.

   | **Note** |
   | --- |
   | The security policy settings that you can configure within SCW overlap with those that can be set by using security templates (.inf files). On the Security Policy File Name page, you can include a security template if you want to add settings that cannot be configured directly from SCW. If you attach a security template, and it contains settings that conflict with some SCW-configured settings, the SCW-configured settings have precedence. |

10. On the **Completing the Security Configuration Wizard** page, click **Finish**.

    The policy will now be created.

# Deploying SCW security policy files

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This topic discusses the deployment of SCW security policy files.

Security policy files authored with SCW are created as .XML files that are saved by default in %systemdir%\security\msscw\Policies. They can be deployed by using the SCW user interface or by using the scwcmd.exe command-line tool included with SCW.

**Using the SCW user interface**

This method is best suited for a single server. To apply an SCW security policy to multiple servers, use the command-line method or Group Policy instead.

**To apply an SCW security policy to a server**

1. Click **Start**, click **Administrative Tools**, and then click **Security Configuration Wizard**.
2. Read the **Welcome page**, and click **Next**.
3. On the **Configuration Action** page, select **Apply an existing security policy**, type the full path and file name of the policy, and then click **Next**.
4. On the **Select Server** page, type the name of the server to which the policy will be applied, and then click **Next**.
5. On the **Apply Security Policy** page, click **Next**.
6. On the **Applying Security Policy** page, wait for processing to finish, and then click **Next**.
7. On the **Completing the Security Configuration Wizard** page, click **Finish**.

**Using the Command-line method**

To configure multiple servers with a policy, you can use **scwcmd configure /p:***PolicyFile* **/i:***MachineList* at the command prompt, rather than this SCW UI procedure. Type **scwcmd configure** at the command prompt to learn about the parameters.

# Converting of SCW security policy files into Group Policy objects (GPOs)

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This topic discusses the conversion of SCW security policy files into GPOs.

SCW saves its security policies as .xml files, and the Scwcmd.exe command-line tool allows you to convert these files and save them as GPOs by using the **scwcmd transform** command. The SCW user interface itself does not support GPOs.

**Scwcmd transform**. Creates a new GPO and defines the settings for these Group Policy extensions:

- **Security Settings**. Contains service settings, registry values, audit policy and security template settings that were imported into the SCW XML policy.
- **IP Security Policies**. Contains the IPsec configuration that is defined in the SCW policy.
- **Windows Firewall**. Contains Windows Firewall settings that are defined in the SCW policy.

Any Internet Information Services (IIS) settings that are defined in the SCW policy will be lost during the **scwcmd transform** operation because Group Policy does not support configuration of IIS settings.

After the GPO is created, the administrator must manually link the GPO to the target organizational unit (OU) by using Active Directory Users and Computers, or by using the Group Policy Management Console (GPMC).

**To save an SCW security policy in native Group Policy format**

- At the command prompt, type

  **scwcmd transform /p:** *PathAndPolicyFileName* **/g:** *GPODisplayName*

  where *PathAndPolicyFileName* is the policy you created earlier with SCW, including its .xml file name extension and *GPODisplayName* is the name of the Group Policy object (GPO) as it appears in Group Policy Object Editor or in Group Policy Management Console (GPMC).

  When the **scwcmd transform** command is complete, the GPO gets created in Active Directory, but the policy it contains is not applied until the GPO is linked to an organizational unit. For instructions about linking GPOs, see the GPMC Help.

# Using Group Policy and Active Directory with SCW

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This topic is about using Group Policy and Active Directory with SCW.

SCW is not a replacement for the security settings available in Group Policy; it is a security tool to supplement Active Directory and its policy infrastructure, and to make the familiar Active Directory tools such as the Active Directory Users and Computers snap-in more useful to you.

Active Directory best practices include using Active Directory Users and Computers to group computer objects into organizational units (OUs) for convenient administration. This helps you deploy SCW policies by using Group Policy objects (GPOs). You can create an SCW security policy by prototyping a server with the SCW user interface, convert it into a GPO by using scwcmd.exe, and link the GPO to an OU. If all the functionally similar servers are in the OU, they all receive your SCW-authored security policy.

**Group Policy Object Editor**
The Group Policy Object Editor is the user and computer configuration management tool that is included with Active Directory. Group Policy settings include security settings in the Group Policy Object Editor even though security settings are displayed and processed differently from most of the other Group Policy settings. For example, security settings are persistent in the registry, whereas Group Policy administrative templates settings are rewritten whenever policy is refreshed.

The Group Policy Object Editor is used to edit all Group Policy objects, including those converted from SCW format. So SCW makes the Group Policy Object Editor more useful by supplying GPOs tailored to server types.

**Group Policy Management Console**
Group Policy Management Console (GPMC) answered the need for easier analysis, planning, and backup of policy in an Active Directory environment, where multiple GPOs frequently apply to the same system and custom OU arrangements affect inheritance. GPMC is made available as a free download. It supports all enterprise-wide Group Policy tasks with the exception of editing individual GPOs, which is still done with Group Policy Object Editor.

In particular, GPMC is used for linking GPOs to OUs. Linking is the mechanism by which GPOs are applied to the users and computers within the OUs.

If you use Group Policy Object Editor to edit a GPO that includes SCW security policies, be aware that security settings authored manually in Group Policy Object Editor take precedence over the same settings applied using SCW.

If SCW-authored settings are transformed into a GPO, the decision as to which settings take precedence is determined by ordinary GPO inheritance rules.

**Security Templates and Precedence**
In addition to using SCW to author security policies, you can also apply security settings by using security templates. Security templates are .inf files such as securedc.inf that are located by default in %systemroot%\security\templates\. You can see the security template settings in Group Policy Object Editor and GPMC in the following location:

*GPO_Name*\Computer Configuration\Windows Settings\Security Settings\

The configuration changes that can be made by using the SCW user interface without the use of security templates partially overlap the configuration changes available by using security templates alone. Neither set of configuration changes totally includes the other. For example, the SCW user interface includes IIS settings that are not included in any security template. Conversely, security templates can include such items as Software Restriction policies, which are not configurable through SCW user interface. Some configuration changes, such as IP Security (IPsec) policies, can be set by using the SCW user interface, or by attaching a security template to the native SCW policy file (.xml file), or by using both methods. If you use both methods, however, conflicting policy settings are resolved by using the precedence rules explained later in this section.

In addition, the scwcmd.exe command-line tool with its transform option supports the creation of new, unlinked GPOs from SCW policies. This means that there are, in effect, three ways to use security templates:

- If you are familiar with Group Policy, and have used the Group Policy Object Editor and GPMC, then you are familiar with attaching a security template to a GPO by right-clicking **Security Settings**, clicking **Import Policy**, and then browsing to the .inf file.

- A second way is to include a security template in an SCW policy from the Security Policy File Name page of SCW by

clicking **Include Security Templates**, and then clicking **Add**. As with the Group Policy method, you then browse to an .inf file.

- The third way is to attach a template to a .xml policy file (as just described) and then type at the command line **scwcmd transform /p:***policyfile.xml* **/g:***GPOdisplayname* to create a GPO, which you then link using GPMC.

In an environment utilizing Group Policy, SCW, and multiple security templates, use the following guidelines to anticipate the precedence of security settings:

- Security policy applied through Active Directory-based GPOs has higher precedence than security policy applied through SCW policy files (.xml files).

- Precedence among GPOs is unaffected by whether each GPO was created by scwcmd.exe or not: Only the standard Active Directory inheritance rules (in which local, site, domain, and organization unit GPOs are applied in succession) and link order determine precedence for GPOs.

- Security policy set in the SCW user interface has higher precedence than conflicting policy set in .inf security templates that are attached to the .xml policy file.

- If multiple security templates are attached to the .xml security template, a template that is listed higher in the Include Security Templates dialog box has precedence over a template that appears lower in the list.

Always test security policy in a lab before deploying it to your live production servers.

# Deploying SCW security policy as Group Policy objects

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This topic discusses deploying SCW security policy as Group Policy objects.

Once you have used **scwcmd transform** to convert SCW policies into GPOs, those GPOs are no different than any other Active Directory GPOs from the standpoint of GPMC and Group Policy Object Editor.

Group Policy does not support the rollback of security settings. If you apply SCW security settings by using Group Policy, you cannot roll them back.

See the Group Policy and GPMC Help for instructions about linking GPOs.

# Rollback considerations

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

This topic discusses rolling back SCW policy.

Rolling back SCW policies is not part of routine SCW deployment, but you might face circumstances where you want to roll back an SCW security policy. In those cases, keep the following rollback considerations in mind.

Security settings are persistent in the registry. This behavior is called tattooing. If you apply SCW security settings through Group Policy, the registry will be tattooed, and you will not be able to roll them back.

A best practice for facilitating rollback of SCW policies is to put the policy settings for a server type into just one SCW policy file rather than several.

Rollback can be carried out through the SCW user interface, or alternatively through the command-line tool by using the **scwcmd rollback** command. Type **scwcmd rollback** to see how the rollback parameter is used.

XML files supporting rollback are created and stored by default in %systemdir%\security\msscw\RollbackFiles. SCW and scwcmd.exe search in that location, so you will not generally need to do anything with these files yourself. This is just for your information.

If, after you roll back a security policy, some services do not behave as expected, check whether services that were started manually were stopped. If you apply a security policy that disables manually-started services and then roll back this policy, those services might need to be restarted manually.

There are special rollback considerations for Windows Firewall and Internet Information Services (IIS) settings:

**Windows Firewall**
Keep in mind that SCW policies affect Windows Firewall and IPsec settings based on server roles, approved applications, and port behavior. If you make changes to Windows Firewall and IPsec outside of SCW, .the system might not behave exactly as you might expect after rolling back SCW policies.

**Internet Information Services (IIS)**
IIS settings within SCW policies are lost when the .XML policy files are transformed into GPOs. This means that neither the application nor the rollback of IIS settings are supported by GPMC. To apply or roll back IIS settings, you need to use the SCW user interface or **scwcmd configure** at the command line.

You will only see the IIS section of SCW if you have selected the Web Server role on the **Select Server Roles** page of SCW.

# Network Load Balancing: Security Best Practices for Windows 2000 and Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

**In This White Paper**


- General Assumptions
- Deployment and Operational Management
- Summary of Security Attributes


**General Assumptions**

There are a number of general assumptions that need to hold around the infrastructure and operational best practices to ensure a secure environment in which to run NLB clusters:

1. Servers and storage are in physically secure locations.
2. Practical security implementations such firewalls, network probes and management tools to detect irregular traffic are in place.
3. Best practices and common sense in terms of security are adhered to in areas like administration, storage of logs, backup and restore, etc.
4. Platform level security best practices are adhered to in terms of assigning administrative permissions, Access Control Limiting resources, and other housekeeping roles.
5. The network infrastructure services such as Active Directory, DNS, DHCP, WINS, etc. must be secure. Any compromise of these infrastructure services can lead to a compromise of the NLB itself.
6. Administration tools or other applications that administer NLB clusters can be run from remote workstations. The administrator must ensure that the applications are run from trusted computers. Any compromise on the computers on which the applications are executing (that the cluster administrator runs), can compromise the cluster. For example, if there are untrusted users with elevated privileges on the workstation where the administration tools are run, untrusted or malicious code can be run against the cluster by the cluster administrator without the cluster administrator realizing.
7. NLB creates and maintains a set of objects in the operating system such as files, devices, registry keys, etc. These objects have a default security setting that ensures non-privileged users cannot impact the cluster configuration or the applications running on the cluster. Changing these security settings to less restrictive security settings can lead to the cluster being compromised and application data being corrupted.

# Deployment and Operational Management

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

**Cluster administrators**
Administrators can specify groups or individuals that are allowed to manage the cluster. Because NLB does not implement a cluster-wide administrative account, the cluster administrator must ensure that an appropriate administrative account with Administrator privileges has been enabled on all cluster hosts. This account is used by NLB management tools, such as NLB Manager and the NLB control program, wlbs.exe, to administer individual NLB hosts. To simplify this process and ensure uniform access on all hosts, it is preferable to use a single domain or global group account for this purpose.

**Using the cluster administrative tools through a firewall**
The NLB control program, wlbs.exe, should only be used on NLB hosts or, when remote control operations have been enabled, only on a trusted, internal computer within the firewall. By default, remote control operations are disabled, and they must only be enabled if the NLB cluster has been protected by a firewall that blocks UDP control ports 1717 and 2504. Otherwise, unauthorized remote control packets could be delivered to NLB hosts and thereby cluster operations.

The NLB Manager administrative tool can safely be used on trusted computers outside the firewall. This tool uses a secure method of communication to access NLB hosts. NLB Manager uses WMI interface to manage NLB hosts. Please refer to WMI/DCOM documentation to see what ports to unblock to allow WMI/DCOM access from outside the firewall. Nevertheless, make sure that this tool is only executed on a trusted, remote computer.

**Using Kerberos Authentication in an NLB Cluster**
NLB does not use its own cluster-wide authentication mechanism for cluster administration. Instead, NLB management tools use an appropriate administrative account, which must be enabled with administrative privileges on each NLB host. If a domain account is used for this purpose, it will use Kerberos authentication to control access to NLB hosts.

**Using NLB to support VPN connections**
NLB can be used to load balance virtual private network connections using the PPTP protocol in Windows 2000 and using both the PPTP or IPSec protocols in Windows Server 2003. When using IPSec, some client sessions may be disrupted when a new NLB host joins the cluster. This issue has been corrected in Windows XP SP1 and in Windows Server 2003.

**Network Security**
To coordinate their actions, all NLB hosts communicate with each other by periodically exchanging heartbeat messages on a common subnet, which is also used to receive incoming client requests. This subnet must be physically protected from intrusion. Otherwise, unauthorized heartbeat messages could be delivered to NLB hosts and disrupt cluster operations. Note that NLB heartbeat messages use a uniquely assigned Ethertype (0x886F) which is not normally routed across subnets. However, the cluster administrator must ensure that unauthorized computers and devices which could emit invalid NLB heartbeat packets are not placed directly on the NLB subnet. The effects of these disruptions are described below in the subsection "Rogue Servers".

**Network Flooding**
NLB cluster hosts may be affected by denial of service attacks, which flood the cluster with invalid network packets. These attacks can create additional CPU and network load on the cluster hosts and thereby delay the handling of valid client requests. In addition, they can cause NLB hosts to increase their memory usage up to parameterized limits. These limits are adjustable by setting registry variables on each NLB host, as described in the Windows 2000 Resource Kit. In extreme cases, a high volume of invalid network packets can disrupt extant client connections and require affected clients to reconnect to the cluster. Note that the Windows network protocol stack and server applications, such as Web servers, also may be affected by denial of service attacks.

**Rogue Servers**
Rogue NLB servers on an NLB subnet could emit heartbeat packets that disrupt cluster operations. Disruptions can include impeding NLBs convergence process such that cluster hosts cannot be added and recovery for a failed host cannot be completed. In addition, disruptions can block some or all of NLBs service to clients. NLB subnets must be physically protected from intrusion by unauthorized computers and devices.

**Best practices**
To summarize, best practices for NLB cluster administrators:

- Cluster administrators should establish a common administrative account on all NLB hosts that have Administrator privileges. To simplify cluster management, it is preferable that a single domain or global group account be used for this

purpose.

- The NLB subnet must be physically protected from intrusion by unauthorized computers and devices to avoid interference from unauthorized heartbeat packets.

- NLBs remote control mechanism is disabled by default. It should remain disabled unless the administrator can ensure that the NLB subnet has been physically protected from intrusion and that UDP control ports 1717 and 2504 have been firewalled to prevent unauthorized remote control.

- If remote control has been enabled, remote control operations using the NLB control program wlbs.exe must only be performed from a trusted computer within the firewall.

- NLB Manager provides a secure method for remote cluster management and should be used instead of wlbs.exe whenever possible. NLB Managers optional host list file should only be saved to a file that is accessible by user accounts (such as the NLB cluster administrators account) that have local Administrator privileges.

# Summary of Security Attributes

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section outlines the various objects and their associated security attributes needed to ensure that NLB runs successfully. This section lists the minimum security requirements. If the security attributes are made more restrictive, NLB may not run. If the security attributes are less restrictive, it may result in malicious attacks being able to compromise the cluster or corrupt end-user data. We recommend that the security attributes on the objects not be changed from the default settings.

**Directory and file protection**

| Object | Description | Minimum recommended security attributes |
|---|---|---|
| **<filespec>** | **NLB Manager host list file** | **BUILTIN\Administrators: Full control** |

**System resources**

The cluster service uses other system resources that are protected by default.

| Object | Description | Minimum recommended security attributes |
|---|---|---|
| **HKLM*** | **HKLM registry hive** | **SYSTEM:Full control**<br><br>**BUILTIN\Administrators: Full control** |

# Windows Cluster Technologies: Remote Setup, Unattended Installations and Image-based Installations of Network Load Balancing

Updated: June 1, 2002

Applies To: Windows Server 2003 with SP1

Published: June 1, 2002

**In This White Paper**

Unattended Installation of NLB Clusters

Image-based Installations of NLB Clusters

Generating and Distributing a Virtual Hard Disk

Appendix A - Windows Cluster Technologies: Remote Setup, Unattended Installations and Image-based Installations of Network Load Balancing

**Overview**

You can install Network Load Balancing (NLB) clusters remotely in the following ways:

1. Using WMI to install and configure NLB cluster
2. Unattended installation of NLB clusters during the operating system install
3. Image-based installation of NLB clusters during the operating system install

**Unattended installation of the operating system and the cluster service**

To perform a clean install or to upgrade the operating system, you can either perform a hands-on installation (you must manually provide all of the answers required for setup), which this document will not address, or an unattended installation using a so-called answer file. Server clusters can be installed after the operating system installation is complete, by specifying certain switches in the answer file. Unattended installation of the operating system and server clusters is discussed in detail in the Unattended installation of NLB Clusters section of this document.

**Image-based installation of the operating system and the cluster service**

Another way to perform a clean install of the operating system is to use image-based installations, also known as disk cloning or simply cloning. Cloning is the fastest method of deploying Microsoft Windows Server 2003, Enterprise Server. Image-based installations are useful if you are deploying computers that have similar hardware and you want to perform quick installations with no or minimal manual interaction. You cannot use image-based installation to upgrade an operating system. Image-based installation is discussed in more detail in the Image-based installations of NLB Clusters section of this document.

# Unattended Installation of NLB Clusters

Updated: June 1, 2002

Applies To: Windows Server 2003 with SP1

Unattended installation of NLB clusters consists of the following steps:

1. Planning and Preparation for operating system installation
2. Preparing an answer file for unattended installation of the Windows Server 2003, Enterprise Edition operating system
3. Creating a Master Installation
4. Running winnt32.exe
5. Enabling and configuring Network Load Balancing

**Planning and Preparation for Operating System installation**
For detailed instructions on how to plan and prepare for Windows Server 2003, Enterprise Server operating system installation, please refer to the *Microsoft Windows XP Preinstallation Guide*. This document is located in Deploy.cab on the Windows Server 2003 CD. Although the Guide contains the instructions on preparing for Windows XP installation, the same instructions apply to Windows Server 2003, Enterprise Edition as well.

**Answer file**
To automate the Windows Server 2003 installation completely you must prepare an answer file, usually called unattend.txt, which contains all of the answers for a series of graphical user interface (GUI) dialog boxes. Unattend.txt tells Setup how to interact with the distribution folders and files you create and supplies Setup with all of the information that the end user provides during a typical Windows Setup. For example, Unattend.txt contains a FullName entry in the **[UserData]** section that instructs Windows Setup to prompt the end user to provide a full name.

Although the answer file for Setup is commonly called Unattend.txt, for a network pre-installation you can name the file anything you like. For a CD-based Setup, you must name the answer file Winnt.sif.

If you are using the OPK toolkit (available only to OEMs), you can create or modify a Unattend.txt file by using a tool called Setup Manager. You can also use any text editor to create or modify the answer file.

**Operating System Sections and Entries for Unattended Install**
A fully unattended operating system setup requires the following sections and entries:

| Required section | Required entries in that section | Comments |
|---|---|---|
| [Unattended] | UnattendMode<br><br>TargetPath | |
| [GuiUnattended] | AdminPassword<br><br>TimeZone | |
| [Identification] | JoinWorkgroup | You may instead specify JoinDomain, DomainAdmin, and DomainAdminPassword. |
| [LicenseFilePrintData] | AutoMode<br><br>AutoUsers | Requires **AutoUsers** if AutoMode = PerServer. |
| [Networking] | | Required to configure network protocols. |
| [UserData] | ComputerName | |

**[Unattended]**

This section contains entries for running Setup, such as whether to extend the partition on which you install the Windows operating system, specifies location of files necessary for installation during Setup, etc. You must instruct Setup to accept the EULA if you want your installation to be fully automated:

```
 [Unattended]
OemSkipEula = Yes
```

**[GuiUnattended]**
This section of Sysprep.inf contains entries for preparing the graphical user interface (GUI) for unattended Setup. To completely automate operating system setup you need to include the following entries:

```
[GuiUnattended]
AutoLogon = Yes
AdminPassword = *** <replace the asterisks with the actual password>
OemSkipWelcome = 1
OemSkipRegional = 1
TimeZone = 20 <specify your time zone here>
```

**[Identification]**
This section contains entries for specifying the network identification of a computer. If these entries are not present, Setup adds the computer to the default workgroup called WORKGROUP. If there is not enough information for this entry, Setup will prompt the end user to provide this information.

Prior to adding the first node to a cluster or adding additional nodes, all nodes must have joined a domain (instead of a workgroup) and be members of the same domain. When you create a cluster or join nodes to a cluster, you specify the domain user account under which the Cluster service runs. Therefore, you need to specify in this section which domain a destination computer should join, and provide the valid credentials under which to join

```
[Identification]
DomainAdmin = your_domain\user
DomainAdminPassword = ***** <replace the asterisks with the actual password>
JoinDomain = your_domain
```

**[LicenseFilePrintData]**
This section contains entries for installing licensing information for the Windows Server 2003 family.

```
[LicenseFilePrintData]
AutoMode = PerSeat
```

**[Networking]**
This section contains no entries. However, to configure network protocols during an unattended setup, you must include the **[Networking]** section name in your answer file. In addition, you must include additional sections that describe your network adapters, protocols to be installed, and static IP address information.

If you decide to change your Network Load Balancing model, add, remove or modify virtual clusters during the unattended installation, you must make sure you have all of the required information (new cluster virtual IP addresses, additional static IP addresses, etc.) available. You will need to provide all of that information in the [Networking] section of the answer file.

Depending on the particular scenario in which you use NLB clusters the [Networking] section can contain different sets of information. For example, if you use multiple network adapters, you need to configure them both here; otherwise you only need to configure one network adapter. If you use multiple virtual clusters, you also need to configure the network adapter(s) and bind NLB in this section. Below is an example of a [Networking] section that configures two network adapters that are configured for use with NLB later in the [params.MS_WLBS.Adapter] sections.

For a complete list of instructions for configuring and enabling NLB and NLB best practices, please refer to Windows Advanced Server 2003 Online Help/Availability and Scalability/Network Load Balancing.

Use the following entries in unattend.txt to enumerate your network cards:

```
[NetAdapters]
;list your network adapters here
Adapter01 = params.Adapter01
Adapter02 = params.Adapter02
```

Next, you need to uniquely identify each network adepter. You can either use the INFID entry where you need to specify the Plug and Play IDs of your network adapters or the NetCardAddress entry to specify MAC addresses. You need to run a special tool called pnpids.exe to retrieve the PnP IDs for your network cards. For computers with multiple adapters of the same type (that is, the Plug and Play ID for each is the same), you must specify the **NetCardAddress** or PCI location information. If you do not specify the **NetCardAddress** or PCI location, only the first adapter enumerated/detected that matches the PnP ID receives

the answer file entries.

If you specify the **NetCardAddress** or PCI location, Setup does not use the INFID entry because it is the least specific entry.

Therefore, we recommend that you simply identify your network cards by their MAC addresses as follows:

```
[params.Adapter01]
NetCardAddress = 0x0002A5D0510F
[params.Adapter02]
NetCardAddress = 0x000476D5F57D
;Install Client for MS Networks
[NetClients]
MS_MSClient = params.MS_MSClient
[params.MS_MSClient]
;Specify which network protocols to install
[NetProtocols]
MS_TCPIP = params.MS_TCPIP
;TCP/IP properties
[params.MS_TCPIP]
AdapterSections = params.MS_TCPIP.Adapter01, params.MS_TCPIP.Adapter02
```

Now we need to configure both network cards with static IP address information since NLB does not allow DHCP-assigned addresses.

```
;Adapter01 is used for intra-cluster communication.
 [params.MS_TCPIP.Adapter01]
DHCP = No
IPAddress = 172.24.23.91
SpecificTo = Adapter01
SubnetMask = 255.255.0.0
DNSServerSearchOrder = 157.5.0.1,157.6.0.2
WINS = Yes
WINSServerList = 157.5.0.1,157.5.1.2
;Adapter02 represents a public network
[params.MS_TCPIP.Adapter02]
DefaultGateway = 157.6.8.1
DHCP = No
IPAddress = 157.5.4.1
SpecificTo = Adapter02
SubnetMask = 255.255.252.0
DNSServerSearchOrder = 157.5.0.1,157.6.0.2
WINS = Yes
WINSServerList = 157.5.0.1,157.5.1.2
```

**[UserData]**
This section contains entries for specifying user settings during setup:

```
[UserData]
ComputerName = SV-NODE1
FullName = User Name
OrgName = Organization Name
ProductKey = 11111-22222-33333-44444-55555
```

Of course you need to replace all of the names, passwords and the cluster IP address with real ones.

**NLB-specific Sections and Entries**
**[params.MS_WLBS]**
This section contains parameters for configuring NLB and binding NLB to network adapters. For unique Configuration/Binding, see the [NetBindings] section below.

```
[params.MS_WLBS]
AdapterSections=params.MS_WLBS.Adapter01,params.MS_WLBS.Adapter02
[params.MS_WLBS.Adapter01]
SpecificTo=Adapter01
ClusterIpAddress=172.24.23.10
ClusterNetworkMask=255.255.252.0
IGMPSupport=0
MulticastSupportEnable=1
RemoteControlEnabled=1
HostPriority=1
ClusterModeOnStart=1
DedicatedIPAddress=172.24.23.91
```

```
DedicatedNetworkMask=255.255.255.0
Ports=1,65535,Both,Multiple,Single,Equal
[params.MS_WLBS.Adapter02]
SpecificTo=Adapter02
ClusterIpAddress=172.24.23.10
ClusterName = nlbcluster.domain.com
ClusterNetworkMask=255.255.252.0
IGMPSupport=0
MulticastSupportEnable=1
RemoteControlEnabled=1
HostPriority=1
ClusterModeOnStart=1
DedicatedIPAddress=157.5.4.1
DedicatedNetworkMask=255.255.252.0
Ports= 1,65535,Both,Multiple,Single,Equal
```

What each line means:

ClusterIpAddress - This address is a virtual IP address and it must be set identically for all hosts in the NLB cluster

ClusterName - Specifies a full Internet name for the NLB cluster

ClusterNetworkMask - Specifies the subnet mask for the NLB cluster IP address

IGMPSupport - Specifies whether or not to use IGMP support to prevent switch flooding

MulticastSupportEnable - Specifies whether or not to use a multicast MAC address for NLB cluster operations

RemoteControlEnabled - Specifies whether or not to enable remote-control operations

HostPriority - Specifies hosts unique priority ID in the NLB cluster

ClusterModeOnStart - Specifies whether a host joins the NLB cluster upon startup or waits for an explicit command to do so

DedicatedIPAddress Specifies hosts unique IP address

DedicatedNetworkMask - Specifies the subnet mask for the dedicated IP address

Ports - Specifies rules for distributing traffic among NLB cluster hosts according to protocol and destination ports

For a complete list of all available NLB configuration options in the answer file, please refer to the following document: *Microsoft Windows XP Preinstallation Guide*. This document is located in Deploy.cab on the Windows Server 2003 CD.

**[NetBindings]**
This section contains entries for disabling, enabling, and demoting binding rules for network components such as services, protocols, and adapters.

[NetBindings]

Disable=MS_WLBS,Adapter01

Enable=MS_WLBS,Adapter02

For more information about unattend.txt, please refer to the following document: *Microsoft Windows XP Preinstallation Guide*. This document is located in Deploy.cab on the Windows Server 2003 CD.

**Creating a Master Installation**
There are no cluster-specific items in this step. Please refer to the *Microsoft Windows XP Preinstallation Guide*, located in Deploy.cab on the Windows Server 2003 CD for detailed instructions on how to create a Master Installation.

**Running winnt32.exe**
You can run Winnt32.exe from a command prompt to start Windows Setup.

For a complete list of the Winnt32.exe command-line options, see Winnt32.exe Command-Line Options topic in the *Microsoft Windows XP Preinstallation Guide*. This document is located in Deploy.cab on the Windows Server 2003 CD. Although the Guide contains the instructions on preparing for Windows XP installation, the same instructions apply to Windows Server 2003, Enterprise Edition as well.

To install Windows from a network share, boot the computer using your Windows Server 2003 CD (or any other startup media) and then connect to the desired network location. Placing the distribution share and the configuration set on the

network adds considerable flexibility to the pre-installation process.

To run Setup using Winnt32.exe

Prepare an Unattend.txt file as part of building a configuration set as described above.

1. Start the computer where you want to install Windows.

2. Click **Start**, click **Run**, and then type:

   path_to_winn32**winnt32** /**unattend**:filename

   where:

   path_to_winnt32

   Specifies the path to the location of Winnt32.exe. For example, \\server\share\Lang\Eng\sku\pro\x86\I386.

   filename

   Specifies the name of the answer file, which contains answers to installation questions you want to automate. The value of *filename* is usually Unattend.txt.

**Installing and Configuring NLB**

If you specified all of the information in the [params.MS_WLBS] section of unattend.txt as described in section NLB-specific Sections and Entries above, you need to do nothing further you have already installed and configured an NLB cluster with at least one host in it. You can now use NLB Manager or WMI to add hosts to this cluster, or you can add hosts to the cluster when you perform an unattended install on those computers with the [params.MS_WLBS] section of the answer file configured as described above.

# Image-based Installations of NLB Clusters

Updated: June 1, 2002

Applies To: Windows Server 2003 with SP1

**Image-based Installations of NLB Clusters**
Image-based installations, also known as disk cloning or simply cloning, is the fastest method of deploying Windows Server 2003, Enterprise Server. Image-based installations are useful if you are deploying computers that have similar hardware and you want to perform quick, clean installations with minimal end user interaction

Cloning NLB Clusters consists of the following steps:

1. Building a master installation on a master computer. Building a master installation includes installing and configuring the operating system and any software and drivers you need to include on your virtual hard disk.

2. Preparing the master installation with the Sysprep tool. This includes configuring and running the Sysprep tool on the master computer.

3. Generating a virtual hard disk of the master installation with the disk-imaging tool. This includes saving each virtual hard disk to a permanent storage location.

4. Loading virtual hard disks on destination computers that will make up a cluster.

**Using SYSPREP to prepare for imaging**
**Building a Master Installation**
You need to setup a master computer which you will use to build a master installation and later to create a virtual hard disk.

Building a master installation includes installing and configuring the operating system and any software and drivers you need to include on your virtual hard disk.

**Prepare the master installation.**
After you have prepared a master installation on the master computer, you need to prepare a virtual hard disk for cloning. Microsoft provides a tool called sysprep that allows you to prepare an image for cloning. The tool can be found on **the Windows Advanced Server 2003 CD in the \support\tools directory**. When you run sysprep on the master computer, you can supply one of the following two switches: **-factory or reseal**.

You can use Factory mode to install additional drivers and applications at the stage after the reboot that follows Sysprep. Normally, running Sysprep as the last step in the pre-installation process prepares the computer for delivery. When rebooted, the computer starts Mini-Setup. By running Sysprep with the -factory option, the computer reboots in a network-enabled state without starting Mini-Setup. In this state, Factory.exe processes its answer file, Winbom.ini, and performs the following actions:

1. Copies drivers from a network source to the computer

2. Starts Plug and Play enumeration

3. Stages, installs, and uninstalls applications on the computer from source files located on either the computer or a network source

4. Adds customer data

When finished, run Sysprep with the -reseal option to prepare the virtual hard disk for cloning.

For example, you can create a master installation that contains a minimal set of drivers, run sysprep -factory, create an image of the installation, and copy the image to multiple destination computers. If any destination computers require additional drivers, run sysprep -factory on the destination computer and load those drivers at that time.

For more information about the Factory mode and winbom.ini file, refer to *Microsoft Windows XP Preinstallation Guide*. This document is located in Deploy.cab on the Windows Server 2003 CD.

If you do not need to install any additional drivers or applications on the destination computers, run sysprep in the Reseal mode.

**Sysprep and Sysprep.inf Answer file**
To completely automate NLB cluster installation, you will need to use the sysprep.inf file which is an answer file used by sysprep during and after Mini-Setup.

**Required Sections and Entries in sysprep.inf**
A fully unattended Setup requires the following sections and entries:

| Required section | Required entries in that section | Comments |
|---|---|---|
| [Unattended] | UnattendMode<br><br>TargetPath | |
| [GuiUnattended] | AdminPassword<br><br>TimeZone | |
| [Identification] | JoinWorkgroup | You may instead specify JoinDomain, DomainAdmin, and DomainAdminPassword. |
| [LicenseFilePrintData] | AutoMode<br><br>AutoUsers | Requires **AutoUsers** if AutoMode = PerServer. |
| [Networking] | | Required to configure network protocols. |
| [UserData] | ComputerName<br><br>FullName | |

Please refer to the Operating System Sections and Entries for Unattended Install section of this document for an explanation of what each of these entries is used for.

**NLB-specific Sections and Entries**
See section NLB-specific Sections and Entries for NLB specific configuration options in the sysprep.inf answer file.

# Generating and Distributing a Virtual Hard Disk

Updated: June 1, 2002

Applies To: Windows Server 2003 with SP1

**Creating a Virtual Hard Disk**
After you have prepared a virtual hard disk for cloning, you need to create a virtual hard disk of your master installation with a disk-imaging tool and save the virtual hard disk to a permanent storage location. You can use a third-party disk imaging software or a Microsoft technology called iBIG. If you are using a third-party product, refer to the accompanying documentation on how to create and distribute a virtual hard disk.

**Startup Media**
Before you can load virtual hard disks on destination computers, you need some kind of startup media to boot your computers from. Startup media contains the system files and device drivers that are necessary to start a computer so that the primary hard disk is accessible but not in use. Startup media might also contain network adapter and network drivers, CD and DVD device drivers, disk configuration tools, and scripts or batch files. You can use a floppy, CD, DVD or network boot as your startup media, depending on the capabilities of your destination computers.

If you use third-party disk imaging products, they often provide tools to create different startup media. Otherwise you need to create your own.

Follow these guidelines when creating your startup media:

- Your startup media must provide network support if you are distributing virtual hard disks across a network.
- Your startup media must provide CD or DVD device support if you are distributing virtual hard disks on media and you are using a floppy disk as your startup media.
- Your startup media must support the tools you need to copy a virtual hard disk from a storage location to a destination computer. For example, if your startup media is an MS-DOS startup disk then you need to use MS-DOS tools to copy the virtual hard disk onto the destination computer.

For more information about choosing and creating startup media, refer to the Microsoft Windows Server 2003 Corporate Deployment Tools User's Guide (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating CD.

**Distributing Virtual Hard Disks**
After an image (or images) has been created and placed on a distribution share (or distribution media such as CD or DVD) and you have a startup media to boot your destination computers, you are ready to distribute the images to destination computers.

You need to make sure that your cluster hardware and networks are set up as described in the Windows Advanced Server 2003 Online Help/Availability and Scalability/Cluster Servers. All of your cluster nodes that you will be installing already have to be connected to the shared storage.

You can load virtual hard disks to all of your cluster nodes simultaneously. Many third-party tools support multicast image distribution. You can also use iBIG to distribute virtual hard disks to your cluster nodes.

After you have distributed virtual hard disks to destination computers, sysprep runs Mini-Setup. After Mini-Setup finishes, you should verify that all of the nodes have successfully joined the cluster. Open NLB Manager to see which nodes participate in the cluster, and whether everything is up and running. If it is, your cluster is ready.

# Appendix A - Windows Cluster Technologies: Remote Setup, Unattended Installations and Image-based Installations of Network Load Balancing

Updated: June 1, 2002

Applies To: Windows Server 2003 with SP1

**Sample sysprep.inf file**

This is a sample sysprep.inf file that can be used for unattended installation of a server cluster.

All computer names, cluster name, user data, passwords and IP addresses need to be replaced with real ones before you can use this file

```
[Unattended]
OemSkipEula = Yes
[GuiUnattended]
AutoLogon = Yes
AdminPassword = ****
AdvServerType = Servernt
OemSkipWelcome = 1
OemSkipRegional = 1
TimeZone = 20
[LicenseFilePrintData]
AutoMode=PerSeat
[UserData]
ComputerName = SV-NODE1
FullName = "User"
OrgName = "Company"
ProductKey = 11111-22222-33333-44444-55555
[Networking]
InstallDefaultComponents=Yes
[Identification]
DomainAdmin = domain\user
DomainAdminPassword = password
JoinDomain = domain
;We'll have two network adapters - Adapter01 and Adapter02.
;Note: The adapter specified here as No.1 may not be
;LAN connection No.1 in the user interface.
[NetAdapters]
Adapter01 = params.Adapter01
Adapter02 = params.Adapter02
;Specify the network card used by network adapter No. 1.
[params.Adapter01]
ConnectionName="Intel Adapter"
netcardaddress = 0x0002A5D0510f
;Specify the network card used by network adapter No. 2.
[params.Adapter02]
ConnectionName="3Com Adapter"
netcardaddress = 0x000476D5F57D
;Install client for MS networks.
[NetClients]
MS_MSClient = params.MS_MSClient
[params.MS_MSClient]
;Only install TCP.
[NetProtocols]
MS_TCPIP = params.MS_TCPIP
[params.MS_TCPIP]
;TCP/IP properties
AdapterSections = params.MS_TCPIP.Adapter01,params.MS_TCPIP.Adapter02
;Card 1 using static IP address info
[params.MS_TCPIP.Adapter01]
DHCP = No
IPAddress = 172.24.23.91
SpecificTo = Adapter01
SubnetMask = 255.255.255.0
DNSServerSearchOrder = 157.5.2.2,157.5.23.13
```

```
WINS = Yes
WINSServerList = 157.5.2.22,157.5.7.5
;Card 2 using static info
[params.MS_TCPIP.Adapter02]
DefaultGateway = 157.5.4.1
DHCP = No
IPAddress = 157.5.4.5
SpecificTo = Adapter02
SubnetMask = 255.255.252.0
DNSServerSearchOrder = 157.5.2.2,157.5.23.13
WINS = Yes
WINSServerList = 157.5.2.22,157.5.7.5
[params.MS_WLBS]
AdapterSections=params.MS_WLBS.Adapter01,params.MS_WLBS.Adapter02
[params.MS_WLBS.Adapter01]
SpecificTo=Adapter01
ClusterName=nlbcluster.domain.com
ClusterIpAddress=157.5.4.15
ClusterNetworkMask=255.255.252.0
IGMPSupport=0
MulticastSupportEnable=1
RemoteControlEnabled=1
HostPriority=1
ClusterModeOnStart=1
DedicatedIPAddress=172.24.23.91
DedicatedNetworkMask=255.255.255.0
Ports=1,65535,Both,Multiple,Single,Equal
[params.MS_WLBS.Adapter02]
SpecificTo=Adapter02
ClusterName=nlbcluster.domain.com
ClusterIpAddress=157.5.4.15
ClusterNetworkMask=255.255.252.0
IGMPSupport=0
MulticastSupportEnable=1
RemoteControlEnabled=1
HostPriority=1
ClusterModeOnStart=1
DedicatedIPAddress=157.5.4.5
DedicatedNetworkMask=255.255.252.0
Ports= 1,65535,Both,Multiple,Single,Equal
[NetBindings]
Disable=MS_WLBS,Adapter01
Enable=MS_WLBS,Adapter02
```

# Print Server Upgrade, Migration, and Interoperability

Updated: April 22, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft® Windows® 2000 introduced a dramatic new version of the Microsoft Windows NT® operating system. The printing subsystem was among the many areas of Windows 2000 that experienced substantial changes from previous versions of Windows NT. Even though the benefits of these changes have proven their worth in terms of the stability and robustness of the operating system, they come at a price. During the upgrade and migration to Windows 2000 or Windows Server 2003, companies often find themselves in a mixed environment that includes Windows NT 4.0 as well as Windows 2000, Windows XP, or the Windows Server 2003 family.

Adding to the complexity of the interoperability impact is the fact that network printing involves rendering of print jobs using both local and remote (client and server) driver functions and settings. For this reason, network printing in an environment with Windows NT 4.0 and Windows 2000, Windows XP, or the Windows Server 2003 family can present challenges to corporations.

The purpose of this document is to help explain the procedural and environmental conditions that exist in a network printing environment, especially as they relate to mixed versions of the Windows NT operating system. Additionally, various methods for dealing with these problems are discussed.

Some of the interoperability concerns that arise in a mixed environment of Windows NT 4.0 and later versions of Windows are rooted in changes to the printer driver architectural model. This document examines the move from kernel-mode to user-mode process space, as well as the upgrade of the Microsoft Universal Printer Driver. Other issues related to device limitations in the printer hardware are presented only when interoperability problems arise in the driver model. This document also reviews how device font support affects the transitional period for customers upgrading from Windows NT 4.0.

It is important to note that this discussion focuses on the transitional period that corporations face due to the challenge of upgrading tens of thousands of clients in order to migrate to Windows 2000 or later. The benefits of Windows 2000 and later printing improvements are not fully accessible to mixed or legacy clients. This paper assumes the customers' goal is to achieve a network printing environment in which all clients and servers are running Windows 2000 or later.

**In This White Paper**


Windows 2000 and Windows Server 2003 Printing Architecture

Changes to Printer Drivers in Windows 2000

Upgrading to Windows 2000 Server or the Windows Server 2003 Family

Migrating to Windows 2000

Point and Print for Windows 2000

Customer Actions Moving Forward

Summary (Print Server Upgrade, Migration, and Interoperability)

Related Links (Print Server Upgrade, Migration, and Interoperability)

# Windows 2000 and Windows Server 2003 Printing Architecture

Updated: April 22, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have a mixed environment that includes Windows NT, Windows 2000, or the Windows Server 2003 family, it can be useful to understand the Windows print subsystem to identify interoperability issues that may arise. The printing architecture of Windows 2000 (and later versions of Windows) consists of a print spooler and a set of printer drivers. By calling device-independent Win32® printing and Graphics Device Interface (GDI) functions, applications can create print jobs and send them to any of a variety of devices, including laser printers, vector plotters, raster printers, and fax machines. Printer drivers include a user interface (UI) component that allows users to control a printer's selectable options.

An application's calls to Win32 GDI functions are passed to the GDI graphics engine, which either spools the drawing instructions as an enhanced metafile (EMF) or, in conjunction with a printer driver, renders a printable image that can be sent to the spooler. Spooler components interpret EMF files, and they can insert page layout information and job control instructions into the data stream. The spooler then sends the data stream to the serial, parallel, or network port driver associated with the target printer's I/O port. The device interprets the job, which is formatted in printer description language (PDL) for the device, and outputs the printed document.

Spooler and driver components are designed to be replaceable so that hardware vendors can easily add support for new hardware.

# Changes to Printer Drivers in Windows 2000

Updated: April 22, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Printer drivers are the software that allows application data to be converted into a format that a print device can use to create a printed page. In Windows NT 4.0, all printer drivers ran in kernel mode. These drivers are referred to as version 2 drivers in the Windows printing environment. Windows 2000 and later support user-mode drivers natively and kernel-mode drivers solely for backward compatibility. Microsoft moved to user-mode drivers (also called version 3 drivers) for system reliability and robustness. With a kernel-mode driver, an error can crash the entire server, whereas the effect of an error with a user-mode driver is restricted to the process the driver runs in—usually the spooler. Restarting the spooler is much faster than rebooting the entire system. In addition, a server's additional services remain available to users when the spooler is down. Windows Server 2003 provides for automatic recovery of the Spooler service by default. (A setting in the service control manager)

Moving printer drivers from kernel mode to user mode wasn't the only change introduced with Windows 2000. The Microsoft Universal Printer Driver (Unidrv) was also introduced. This printer driver core replaced the Raster Device Driver (RASDD) interface used in Windows NT 4.0 and earlier versions. The concept behind Unidrv is that independent hardware vendors (IHVs) are not required to write an entire driver, but instead can use the core printing functionality of the Unidrv "engine," and only need to provide device-specific functionality through the use of a minidriver. A minidriver is typically a customization file, or possibly a plug-in, that allows special rendering functions or customized UI options. The core printer control language (PCL) engine (Unidriv.dll) and the PostScript engine (PScript5.dll) both allow the use of minidrivers. The customization file for a Unidrv minidriver is a .gpd file, and for a PostScript minidriver it is a .ppd file.

Windows Server 2003 includes drivers for over 3,800 print devices. Although these drivers are supplied by Microsoft as part of Windows Server 2003, most were provided to Microsoft by the respective printer manufacturer. These so-called inbox drivers ship in the product itself. Although these standard print drivers are adequate for the majority of corporate printing devices, printer manufacturers may choose to make their own drivers available to customers instead of using Unidrv. These drivers are often referred to as monolithic, because they are not implemented as UniDrv minidrivers. (Many manufacturers are beginning to develop all of their drivers using the UniDrv model rather than the traditional monolithic model.)

To summarize, there are several driver types:

- Kernel mode (version 2) or user mode (version 3)
- PostScript (PS) or PCL
- Minidriver or monolithic

**Note**

Drivers can be a combination of these types—for example, kernel-mode PCL driver, user-mode monolithic driver, PostScript minidriver. (Drivers that ship with Windows 2000 or Windows Server 2003 are all user mode.)

All of these driver types are available from the printer manufacture depending on the type of device. The inbox drivers supplied with Windows 2000 and later are also supplied by the printer manufacturer and distributed with their permission. All drivers supplied by Microsoft with Windows 2000, Windows XP, and Windows Server 2003 are Unidrv or PostScript minidrivers and are compatible across these platforms.

Independent hardware manufacturers (IHV) also provide minidrivers that do not ship on the Windows CD.

When you understand the differences among the types of printer drivers, it becomes more clear how mixing and matching can cause problems in a network printing environment. Because Windows NT 4.0 can run only kernel-mode print drivers, and Windows 2000 and Windows Server 2003 can support both user-mode and kernel–mode drivers, resolving an incompatibility sometimes involves modifying the servers running Windows 2000 or Windows Server 2003 so they match the Windows NT 4.0 server.

# Upgrading to Windows 2000 Server or the Windows Server 2003 Family

Updated: April 22, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Prior to upgrading to Windows 2000 or Windows Server 2003 from Windows NT 4 Server, many IHVs recommend uninstalling all IHV printer drivers and reinstalling the latest versions of those drivers for Windows 2000 and the Windows Server 2003 family of products. Microsoft recommends that you use the Fixprnsv.exe utility that is located in the \Printers\Fixprnsv directory on the Windows Server CD. This utility will help identify known bad drivers and replace them with compatible drivers prior to an upgrade. If no replacement is available on the Windows CD, the utility prompts you to obtain an updated driver from the printer manufacturer. Microsoft recommends that you follow the IHV directions to update the print drivers on your system, and then use the Fixprnsv.exe utility to locate any remaining print drivers that could be problematic during the upgrade.

The following sections describe the behavior if drivers are not uninstalled prior to upgrading to Windows 2000 or the Windows Server 2003 family operating system.

**Upgrade Process for Drivers**
When upgrading from Windows NT 4.0 to Windows 2000 or later, there are a number of possible actions with respect to printer drivers. Each situation is described in this section.

- Windows NT 4.0 inbox drivers (PCL and PostScript) that shipped with the operating system are upgraded to the new version of these drivers in Windows 2000 or the Windows Server 2003 family. (The same replacement/update occurs when upgrading Windows 2000 Server to the Windows Server 2003 family.)

- IHV drivers (PCL or PostScript) with exact driver model names matching Windows 2000 or the Windows Server 2003 family driver names are upgraded to inbox drivers with no user intervention.

- Drivers that are blocked for stability reasons are not used. Printers that use these drivers are removed during the upgrade process. If a printer is connected directly to a computer and has a recognized Plug and Play ID, Windows searches for a suitable driver. If one is found, the printer is installed during the upgrade as a new printer.

- Drivers that are not blocked and do not have name matches are carried through the upgrade unchanged.

**Windows NT 4.0 Inbox Drivers**
All Windows NT 4.0 PCL drivers (RASDD) are upgraded to the latest Unidrv drivers as part of the upgrade process. (UNIDRV5 is used for Windows 2000 and UNIDRV5.1 for the Windows Server 2003 family.) PostScript drivers that were shipped with Windows NT 4.0 are upgraded to PostScript version 5.0 in Windows 2000 or PostScript version 5.2 in the Windows Server 2003 family. Both versions 5.0 and 5.2 support PostScript levels 1, 2 and 3.

**Windows 2000 Inbox Drivers**
All Windows 2000 inbox drivers are updated to the Windows Server 2003 version as well.

| **Note** |
| --- |
| Drivers that are changed (updated to a driver with a new name) are no longer associated with the additional drivers that were used on the original server. (Additional drivers are those added through the **Additional Drivers** option in the Printers dialog box.) After a Windows 2000 or later print server is configured, the additional drivers—such as those for Windows NT as well as Windows 95, Windows 98, and Windows Millennium Edition clients—need to be reinstalled accordingly. The association of additional drivers is done per model, so it is only necessary to install these down-level drivers once per model driver, not for each printer queue. |

**IHV Drivers with Matching Names**
IHV Windows NT 4.0 drivers are upgraded to Windows 2000 or the Windows Server 2003 family drivers with no user intervention if the driver model name matches the existing Windows NT 4.0 driver name.

In this situation, the IHV driver is treated exactly the same as the Microsoft driver. Due to interaction problems between Windows 2000 drivers and Windows NT 4.0 drivers in Point and Print environments, many IHVs recommends that their

Windows NT 4.0 driver be reloaded following the upgrade. This issue applies to users who plan to use Point and Print between Windows NT and Windows 2000, Windows XP, and the Windows Server 2003 family of products. (For more information and supported configurations, see "Point and Print for Windows 2000" later in this paper.)

**Blocked IHV Drivers (User Intervention Required)**

Microsoft designed the upgrade path from Windows NT 4.0 so that if the new operating system has a newer driver available, the old driver is replaced. If Windows cannot find a newer driver version, and the driver is not identified with any known problems in Windows 2000, Windows XP, or the Windows Server 2003 family of products, the existing driver is used.

A list of known problematic drivers is maintained in the system file, Printupg.inf. Drivers in this list are not upgraded during the operating system installation. Instead, a newer version of the driver from the IHV must be installed, or an alternate driver selected from the inbox drivers available in Windows 2000 (and later versions of Windows). The information found in Printupg.inf can be helpful in identifying whether a new version from an IHV is blocked or not. Additionally, the Fixprnsrv.exe utility can be used to replace known bad drivers with inbox substitutes if available. If a compatible driver cannot be located and installed, Fixprnsv.exe reports that and refers the user to an IHVs Web site. By running Fixprnsv.exe prior to upgrade, you can identify printers that need to be removed during the upgrade in advance so that sudden service disruptions are avoided.

| **Note** |
| --- |
| Drivers are blocked either because Microsoft has determined, through testing, that they cause substantial instability of the operating system, or because the IHV has requested they be blocked based on their own testing and available updates. |

**Non-Blocked IHV Drivers (No User Intervention Required)**

Certain Windows NT 4.0 drivers are not blocked and do not match the driver name of an inbox driver. These drivers are not altered for replaced during the upgrade process. Microsoft recommends using the Fixprnsv.exe utility to identify drivers that should be replaced. For drivers that are not blocked (either because they are not in the Printupg.inf file or because they have a later date than the date in Printupg.inf), Fixprnsv.exe takes no action.

# Migrating to Windows 2000

Updated: April 22, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Unlike the previous section, in which print server upgrades were discussed, this section deals with the common practice of deploying a new (replacement) Windows 2000 or Windows .NET print server to replace an existing Windows NT 4 or Windows 2000 print server. For the purposes of this discussion, the migration process can be used to either replace a print server (move printers from one server to a new server), or to consolidate print servers (moving printers from multiple print servers to a single print server) Administrators have two main methods for moving printers from one Windows print server to another: manual or automated. It is possible to perform a partially-automated migration of printers as well. This is covered in the following sections also.

**Manual Print Server Migration**

The manual migration process is very straightforward. After you document the current printing environment on the existing Windows print server (that is, the ports, printer names, share names, driver models and versions, and so on), you can recreate each printer manually on the new Windows 2000 print server using the Add Printer Wizard. The advantage to this method is that you can create all of the queues using the latest corporate standards (such as for naming) and the newest driver versions available. The most obvious disadvantage to this method is that it can be very time consuming for large print environments– and tends to be more error-prone than an automated method.

**Automated Print Server Migration**

Microsoft has developed a utility call Print Migrator. In the first two releases of this utility, printers can be moved from one server to another as long as both servers used the same operating system. With the release of version 3, the following key features will be added:

- **Cross-platform support:** Print Migrator 3 lets you move printers from a Windows NT 4 print server to Windows 2000 or Windows. NET Server print servers.
- **Cluster support:** You can migrate printers from a standalone Windows NT 4 or Windows 2000 print server to a clustered print server running Windows 2000 or Windows . NET Server.
- **LPR port conversion:** During backup and restore operations you can convert LPR ports to standard TCP/IP printer ports.

Clearly, it's faster and more accurate to back up and restore these printers automatically. It is, however, important to note that the printers are moved to Windows 2000 or Windows .NET Server without changing their driver versions. That is, drivers previously used on a Windows NT 4.0 print server are created as kernel-mode (version 2) drivers on Windows 2000 or Windows .NET Server. Kernel-mode drivers in use on Windows 2000 are also preserved. As detailed in the "Point and Print" section later in this document, this process provides a much higher level of interoperability with any legacy (Windows NT 4) clients that may still be using print services from a server.

| **Note** |
| --- |
| For more information about the capabilities and limitations of Print Migrator 3.0, see its product documentation.<br><br>Print Migrator 3.0 is available via the Microsoft Print Server web site. See Related Links section below for additional information. |

**Partially-Automated Server Migration Method**

Administrators can migrate printers using a combination of the manual and automated methods. There are several scriptable actions supported by Windows 2000, Windows XP and Windows .NET Server that perform some of the necessary tasks. These scripts are documented in the Windows resource kits.

# Point and Print for Windows 2000

Updated: April 22, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section describes the Point and Print feature, explains methods of installing printer drivers and shared printers for Point and Print, and discusses using Point and Print.

**Using Printer Drivers and Point and Print**

Point and Print refers to the capability of allowing a Windows user to create a connection to a remote printer without providing disks or other installation media. All necessary files and configuration information are automatically downloaded from the print server to the client.

Clients that use the same driver architecture as a print server do not require any additional drivers to be installed on the print server for full support. Other clients require that you install the appropriate driver on the print server so that updates can be downloaded automatically at connection time. For example, a Windows 2000 print server using a version 3 (user-mode) driver for a shared printer does not require any additional drivers to be installed to provide full Point and Print support for Windows 2000 and later clients, because those clients also support the version 3 drivers. However, to vend version 2 drivers for a Windows NT 4.0 client, a Windows 2000 print server requires the appropriate version 2 driver, which can be installed as an using the Additional Drivers option. Other platform architectures, such as Intel x86, Intel IA64 or Alpha, also require specific drivers to be installed. For more information about network printing, including point and Print driver installation, see the *Windows Server Resource Kit*.

Generally, there are four Point and Print methods commonly used by clients to connect to a shared network printer hosted on a Windows 2000 or later print server:

- **UNC path.** From the **Start** menu, choose **Run**, type **\\PrintServer\Printer**, and then click **OK**.
- **Add Printer Wizard.** Select **Network printer** in the wizard, and then specify the path to the shared printer.
- **Drag and Drop.** Use Net View or the **Run** command from the **Start** menu to view shared printers on a remote print server; for example, \\PrintServer. Then drag the desired printer icon into the local workstation's **Printers & Faxes** folder.
- **Double-click.** Use Net View of the Run command as above to view shared printers, and then double-click the shared printer icon to install the printer.

| **Note** |
| --- |
| Windows 2000 and later operating systems also support command line and scriptable interfaces for printer installation and configuration, such as PRNADMIN.DLL or any Microsoft Visual Basic® scripts that use the Windows Management Instrumentation (WMI) print provider. For more information, see the Windows Server resource kits. |

**Point and Print from Windows 2000 Professional to Windows 2000 Server**
**Expected Driver Behavior**

In a Point and Print installation, the printer properties and settings from the server are passed to the client machine. This configuration includes settings for installable options and certain job-processing features. Clients have a local cache of printing preferences, which determine how a job is printed. To perform remote administration on the server, in which the actual device settings or global printing preferences are modified, a user must be logged on with credentials that include administrative rights on the remote server. If the user has administrative rights on the server, the server's printer object can be modified through the Advanced and Device Settings tabs. This functionality does not exist on down-level clients such as Windows 95 or Windows 98.

**Point and Print from Windows NT 4.0 to Windows 2000 Server**
**Expected Driver Behavior**




**Figure 1. Additional Drivers**

Setting up printer drivers on a Windows 2000 server for Point and Print support for Windows NT 4.0 is generally the same as for Windows Server 2003. However, there are a couple of specific differences.

As Figure 1 shows, the **Additional Drivers** dialog box contains three entries under the Intel environment for Windows 2000 drivers, **Windows 2000 or XP**, **Windows 95, 98 and Me** and **Windows NT 4.0 or 2000**. They differ as follows:

If **Windows 2000 or XP** is checked before any additional drivers are loaded, the server is currently running a user-mode driver written specifically for the Windows 2000 or Windows XP platform. That driver cannot be vended to a Windows NT 4.0 client for Point and Print functionality. Microsoft refers to these drivers as version 3 drivers.

If **Windows NT 4.0 or 2000** is checked, the driver running on the server is a Windows NT 4.0 driver that is running under Windows 2000. Microsoft calls these drivers version 2 (or kernel-mode) drivers. In this case, no additional drivers are required to support Point and Print for Windows NT 4.0 clients along with Windows 2000 or XP client computers.

**Server Running Version 3 Driver**

When a version 3 printer driver written specifically for the Windows 2000, Windows XP or Windows Server 2003 platform is installed on a server, a Windows NT 4.0 (version 2) driver with the same model name must be installed in order to support Point and Print for Windows NT 4.0 clients. Changes in the internal data structures that were made to the version 3 Unidrv.dll file might cause a problem in sharing or preserving settings between the clients and the server. This problem can be manifested in two primary ways:

- The Print Server might have new or advanced settings that are simply not visible from a client printer properties dialog box.
- The client properties might not be correctly interpreted by the print server, causing them to be lost.

As a result, the Point and Print connection allows the client to print but might not allow modification of certain device settings or preferences if they do not match those on the server.

| **Note** |
| --- |
| Incompatible data structures can cause certain printer settings—such as Print Text as Graphics or Print Optimization—to be lost between the Windows NT 4.0 and Windows 2000 machines. Additionally, settings such as paper tray assignments and media types might not be preserved between driver models. |

Microsoft and Adobe Systems have developed the latest version of the PostScript driver through a joint development program. Adobe has compiled the latest version of the PostScript driver core for Windows NT 4.0 and later versions of Windows. As a result, the same PSCRIPT core code can be run by Windows NT 4.0, Windows 2000, Windows XP, and the Windows Server 2003 family. By using common PostScript minidrivers (.PPD) files, a higher level of interoperability may be obtained. To increase compatibility, Microsoft recommends that PostScript drivers be considered in mixed environments that include Windows NT with Windows 2000 or the Windows Server 2003 family systems. Another benefit of this situation is that Windows 2000 and later systems run version 3 (user-mode) drivers; only the legacy Windows NT 4.0 machines run the version 2 (kernel-mode) printer drivers. Using PostScript version 5.0 and later is recommended.

**Server Running Version 2 Driver**

When a server has a version 2 driver (Windows NT 4.0) installed as its primary driver, no additional drivers need to be installed to support Point and Print on Windows 2000, Windows XP or Windows NT clients. Further, because the driver being vended to the clients is the same one used by the server, there are virtually no compatibility issues between the Windows Server 2003 server and Windows NT 4.0 clients. After the Version 2 driver is installed on the Windows Server 2003, and the installed printer

is shared through the Sharing tab in the Printer Properties dialog box, Windows 2000 and Windows NT 4.0 clients can connect to the shared printer by either method described above.

On the **Sharing Tab** (accessed through the **Additional Drivers** button), all drivers listed as **Windows NT 4.0 and Windows 2000** are version 2 drivers.

If driver incompatibility or mismatches cause problems and a PostScript solution is not available, Microsoft recommends using these drivers as a solution to support Point and Print between Windows 2000 (and later versions) and Windows NT.

If you are unsure about Point-and-Print compatibility among the different versions of Windows operating systems, contact the manufacturer of your print hardware and ask for drivers that work with a Windows Point-and-Print connection.

**Point and Print from Other Windows Versions to Windows 2003**
**Expected Driver Behavior**
The behavior for Point and Print from Windows Server 2003 to Windows 95, Windows 98 or Windows Millennium Edition differs from that of Windows NT and later versions. No configuration settings are vended from the server to the client. The client user must configure the driver to match the physical device after installation. Additionally, after a driver is downloaded initially, a client running Windows 95, Windows 98, Windows Millennium Edition does not receive driver updates from the print server. The same connection methods are available: drag and drop, the Add Printer Wizard, referencing the UNC path, or double-clicking the shared printer icon.

# Customer Actions Moving Forward

Updated: April 22, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The methods listed above are intended to provide organizations with the maximum interoperability during transitional periods when clients and print servers are running both Windows Server 2003 and Windows NT.

It is strongly recommended to minimize the transitional period during which your environment mixes Windows NT with Windows Server 2003 or Windows XP clients and servers. The highest levels of network printing functionality and interoperability are achieved in environments with clients and servers running Windows 2000 and later.

To help minimize the impact during migrations and upgrades, try the following:

- **Upgrade servers and associated client bases at the same time or close together.** In a very large organization, you might not be able to move the entire installed based to a new operating system as quickly as you can upgrade the servers. By targeting specific servers to the user population that most utilizes them, you can move groups of users and their servers together.

- **Upgrade device font support in older print hardware and maintain updated drivers from the printer manufacturer.** Doing this can help resolve issues related to missing font or character support in many older PCL devices (for example, the Euro symbol).

- **Take advantage of the PostScript printer driver architecture whenever possible.** See your printer manufacture for Point and Print "stacks" that contain all the drivers necessary to support mixed clients. (Windows NT 4.0 plus Windows 2000, Windows XP, or Windows Server 2003).

- **Allow one server to remain at its current operating-system level, and plan for it to be used by clients that have not yet been upgraded.**

- **If the impact of interoperability issues is substantial, use version 2 drivers during the transition.** Microsoft recommends tracking the printers or print servers that are using version 2 drivers so that they can be upgraded to version 3 drivers after the client migration is complete.

- **Use like-model drivers.** If you cannot find a version 2 driver for a particular model, check with the printer manufacturer for a compatible driver that achieves the most functionality and quality. Many IHVs provide drivers for a class, series, or family of devices. Using one of these drivers can help if a manufacturer stops producing version 2 drivers for a specific model.

It is also important to note that the Windows Server 2003 family has a **local policy** that blocks version 2 (kernel-mode) printer drivers from being installed by default. Upon upgrade, existing version 2 drivers are preserved, but an error message appears when an administrator attempts to install new version 2 drivers onto the print server. Installation blocking does not occur when a driver is being installed solely to vend to Windows NT clients. Such a driver is not active on the server but simply shared for the network printing clients, so this action is not blocked. This policy can be disabled using the Local Policy snap-in.

# Summary (Print Server Upgrade, Migration, and Interoperability)

Updated: April 22, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As described in the preceding sections, the Windows printing model provides flexibility to support a wide variety of operating system versions, printer driver languages (PostScript and PCL), IHV driver customization, and automatic downloading. In some scenarios, these features conflict with each other leading to interoperability problems. Using a common PostScript driver core on the clients and servers prevents most of these issues. For PCL printers, the solution for interoperability issues can be the use of version 2 drivers, which are supported by Windows 2000 and the Windows Server 2003 family to provide legacy compatibility.

# Related Links (Print Server Upgrade, Migration, and Interoperability)

Updated: April 22, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

See the following resources for further information:

- For more information about deployment and management of Windows print servers, see the *Windows .NET Server Deployment Guide* (not yet available when this paper was published).
- For more information about the Windows Server 2003 architecture, see the "Display and Print Devices" sections in the *Windows Driver Development Kit* at http://go.microsoft.com/fwlink/?LinkID=123179
- For more information about Print Migrator, see the *Windows .NET Server Resource Kit* at: http://go.microsoft.com/fwlink/?LinkId=458
- For more information about the Windows Server 2003 architecture, see the "Display and Print Devices" sections in the *Windows Driver Development Kit* at http://go.microsoft.com/fwlink/?LinkID=123179
- For more information about connectivity options and support protocols in Windows Server 2003, see Microsoft Windows 2000 Server Resource Kit published by Microsoft Press.
- For additional information about Windows print servers in general, visit the print server Web site at: http://www.microsoft.com/printserver.
- For the latest information about Windows Server 2003, see the Windows 2003 Server Web site at http://www.microsoft.com/windowsserver2003.

# Windows Print Server Scalability and Sizing Technical Overview

Updated: February 7, 2005

Applies To: Windows Server 2003

Print server scalability plays a key role in the purchase and deployment of Microsoft Windows print servers. The maximum load and performance level of a Windows print server is greatly dependent upon key configuration decisions. This paper provides information that helps you size a server to meet an organization's needs by identifying and investigating the main issues and configuration details to maximize print server performance.

**In This White Paper**

Introduction (Windows Print Server Scalability and Sizing Technical Overview)

Factors Influencing Print Server Performance

Performance Monitoring and Baseline Measurement

Reference Systems

Resources (Windows Print Server Scalability and Sizing Technical Overview)

# Introduction (Windows Print Server Scalability and Sizing Technical Overview)

Updated: February 7, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Capacity and scalability are key factors in planning the purchase and deployment of Microsoft Windows print servers. This paper describes several of the key factors that influence the capacity of a given print server configuration. While this paper cannot provide a predictive formula to determine the printing throughput of a given configuration, it does describe several reference systems and their capacity. This paper also presents the information necessary to help the system administrator or capacity planner to estimate and later monitor their server workload.

Some of the factors that influence print server performance and capacity include:

- The number of print server clients
- The operating system version or versions on each client's machine
- The number, size, and frequency of print jobs
- The types of print jobs sent to the print server
- The number and type of printers served by the print server and the types of printer drivers each printer uses

Based on customer feedback, analysis of past print server deployments and performance testing Microsoft has performed in its labs, Microsoft can recommend certain system configurations for specific workloads.

# Factors Influencing Print Server Performance

Updated: February 7, 2005

Applies To: Windows Server 2003

This section examines some of the different factors that can influence print server performance.

**Number of Print Server Clients**
Print server clients are those client computers and applications that have established connections to a print server. These clients may currently be in the process of a submitting a print job or they may have finished submitting a print job and not yet closed their connection to the print server.

Each connection, or named pipe, to the print spooler on a print server uses some limited system resources. Depending on the application used to connect to the print server, some of these connections may remain open, continuing to use resources for many hours or, in some cases, many days. It is possible for a print server to become resource-constrained simply because there are too many open client connections. Because of these lingering connections, there may be more active clients on the print server than you might expect.

The Windows Point and Print Technical Overview presents more information about these connections and provides information about using Named Pipes with remote procedure call (RPC) connections.

**Client Operating System**
Client computers running older versions of Microsoft Windows, such as Windows 95, Windows 98, or Windows Millennium Edition, perform more print processing on the client computer than those clients running Windows 2000 or later versions of Microsoft Windows. As a result, clients running the older versions of Windows place less of a load on the print server than clients running a later version of Windows. Upgrading a client to Windows 2000 or later will shift some of the print processing to the print server which will improve print performance on the client computer but add an additional workload to the server.

The system administrator or capacity planner should consider the additional print server workload resulting from the upgrade of many machines to Windows 2000 or later versions of Windows during the capacity planning process. The overall change in print server performance depends on the percentage of clients originally running older versions of Microsoft Windows—the greater the percentage of older versions of Windows that are upgraded, the greater the change in print server performance.

The Windows Point and Print Technical Overview provides more information about this topic.

**Type of Print Jobs**
Complex print jobs, such as those containing detailed graphics, PDF files, or many different fonts, require more CPU processing than simpler documents. The proportion of this processing performed on the client computer versus the print server depends on the operating system the client computer is running. If the client computer is running Windows 2000 or a later version of Windows, the print server will perform more of this processing workload.

This is an important factor to consider as it will affect both the overall server requirements as well as the change in server workload resulting from the upgrade of older clients to newer versions of Windows.

**Number, Size, and Frequency of Print Jobs**
The available network bandwidth to a print server computer will affect the number, size, and frequency of print jobs that print server will be able to process. The size and type of disk drives installed the print server will determine how many jobs may be queued and processed at one time. For example, in an environment where clients frequently submit small jobs the disk space requirements are different than in one where clients submit larger print jobs, less often. Further, if your organization needs to save print jobs after printing, the number, size, and frequency of print jobs will be important factors when determining the size of the print server hard disk configuration.

**Printer Hardware**
The types of printer devices served by the print server affect the throughput and performance of a Windows print server. Higher-end printers can typically accept more complex data and perform the necessary rendering, processing and production steps internally. Less expensive devices, those having less internal processing and memory capacity, require the Windows print server to perform the necessary rendering, processing and formatting.

The processing speed and internal storage of a printer device also influence the time required for the printer to accept a print job from a Windows print server. The longer it takes to send a print job to the printer, the longer that job remains on the print server's hard disks and the more disk space the job requires on the print server.

When selecting printer hardware for your organization, consider the installable options that may help improve printer and

print server performance. Standardizing printer models and selecting configurations that enable the printer to perform more print-job processing can help improve print server performance.

Printer device and driver reliability are also important considerations for selecting printers. Microsoft recommends using the user-mode printer drivers supplied with Microsoft windows 2000 and Microsoft Windows Server 2003 or drivers provided by printer manufacturers displaying the Designed for Windows logo.

**Types of Printer Drivers**
The number of different driver types supported by a print server has a direct influence on the scalability of the print server. Each additional printer driver requires additional system resources and registry accesses. The fewer system resources used to support a printer, the more of those printers the print server may support and the fewer system resources each print job uses.

For more information about driver architecture, testing, and selection, read Choosing the Right Printer Drivers, found at http://www.microsoft.com/windows2000/technologies/fileandprint/print/default.asp.

**Print Server Configuration**
To achieve maximum performance and capacity for dedicated print servers, consider the following configuration settings.

**Server Service**
Microsoft recommends that the **Maximize performance for file sharing** option be set on dedicated print servers. This setting configures the server to dedicate as many resources as possible to file and print server services.

To view or change server optimization settings:

1. From the **Start** menu, point to **Settings**, and then click **Control Panel**.

2. Double-click **Network and Dial-up Connections**.

3. Double-click **Local Area Connection**, and then click the **Properties** button.

4. In the **Local Area Connection Properties** dialog box, double-click **File and Printer Sharing for Microsoft Networks**.




5. Under **Optimization**, if you are running Windows Server 2003, the **Maximize Data Throughput for File Sharing** option is selected by default.

6. Click **OK**.

The following table shows the values for the **LargeSystemCache** entry with the corresponding options under **File and Printer Sharing for Microsoft Networks Properties**.

| Registry value | User interface option | Description | Notes |
|---|---|---|---|
| 0 | Maximize data throughput for network applications | Optimizes systems for distributed applications that typically do their own memory caching (such as SQL Server, which sets this option by default). This setting is useful for computers providing application services because it favors the working sets of processes over the working set of the file system cache. | This setting is useful for larger server applications and database servers such as SQL Server that need to maximize process working sets over the file system cache working set. |
| 1 | Maximize data throughput for file sharing | Optimizes the system for file and printer sharing resources. This is the default set by Windows 2000 Server Setup.<br><br>If you have at least 128 MB of RAM, this value results in a very high maximum size for the cache working set.<br><br>The Maximize Data Throughput for File Sharing option is useful for computers that typically run the Server service for file sharing | On Windows 2000 the file system cache working set can be increased by 464 MB of additional virtual address space (up to 960 MB) if the system has less than 16 GB of RAM, is not configured to start with the /3GB Boot.ini switch, and the PagedPoolSize entry is set to a value other than 0xFFFFFFFF (or 0 on systems with more than 1 GB of RAM), such as 192000000, the system is not running Terminal Services, and the SystemPages entry is not set to 0xFFFFFFFF. For more information about how various memory-management registry entries interact, see "Optimizing Your Memory Configuration" later in this chapter. |

**Spool Directory**

While accepting jobs and processing them for de-spooling, the print server can perform many disk input/output (I/O) operations. Print server activity also uses memory that can cause excessive page faulting during periods of high print server activity. Overall system performance may degrade because of the system becoming I/O constrained from the combination of high page file and disk file activity.

Microsoft recommends storing spooler files on a physical disk drive or drive array that is different from the physical disk drive used by the operating system for the paging file or the operating system files. This configuration helps to minimize the influence spooler-related disk access has on the dynamic management and allocation of the paging file. The performance of SCSI hard drives can be improved by assigning each hard drive to a different SCSI RAID controller channel.

The spool folder is defined using the Advanced Settings tab on the Server Properties page shown in the following screen shot below.



# Performance Monitoring and Baseline Measurement

Updated: February 7, 2005

Applies To: Windows Server 2003

It is difficult to predict the exact mix of client operating system versions, number of active clients, and the types and sizes of print jobs in advance. The most effective and quantitative approach to server sizing is to observe existing print server performance. There are 13 new performance counters available in Windows 2000 and later versions of Windows that enable administrators to monitor the general load of the print server on the server computer as well as monitor other printer-specific information.

A summary of the performance objects related to print server performance is presented here.

**Performance Objects**
Windows performance counters are grouped into related sets called performance objects. In any performance analysis project, the performance counters of the Processor, Physical Disk, Memory, and Network performance objects should always be examined. They provide a general overview of the system's performance and can also help identify processing bottlenecks and other issues requiring further investigation.

In addition to the performance objects that provide an overview of the system, there are also performance objects that provide information about specific aspects of system and print server operation. Two of these more specific performance objects are the Process object and the Print Queue object.

**Process Object**
The process object reports performance information about each process running on the system. To monitor the performance and load of the print server using System Monitor, select the Spoolsv.exe from the list of performance instances and add the following performance counters:

| Counter Name | Description |
|---|---|
| % Processor Time | The percentage of elapsed time that all process threads used to execute instructions.<br><br>An instruction is the basic unit of execution in a computer, a thread is the object that executes instructions, and a process is the object created when a program is run. This count includes code executed to handle some hardware interrupts and trap conditions. |
| Handle Count | The total number of handles currently open by this process.<br><br>This number is equal to the sum of the handles currently open by each thread in this process. |
| Virtual Bytes | The current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and by using too much, the process can limit its ability to load libraries. |
| Virtual Bytes Peak | The maximum size, in bytes, of virtual address space the process has used at any one time. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. However, virtual space is finite, and the process might limit its ability to load libraries by using too much. |
| Pool Paged Bytes | The size, in bytes, of the paged pool, an area of system memory (physical memory used by the operating system) for objects that can be written to disk when they are not being used. Memory\\Pool Paged Bytes is calculated differently than Process\\Pool Paged Bytes, so it might not equal Process\\Pool Paged Bytes_Total. This counter displays the last observed value only; it is not an average. |
| Pool Nonpaged Bytes | The size, in bytes, of the nonpaged pool, an area of system memory (physical memory used by the operating system) for objects that cannot be written to disk, but must remain in physical memory as long as they are allocated. Memory\\Pool Nonpaged Bytes is calculated differently than Process\\Pool Nonpaged Bytes, so it might not equal Process\\Pool Nonpaged Bytes_Total. This counter displays the last observed value only; it is not an average. |

| | |
|---|---|
| Working Set | The current size, in bytes, of the working set of this process. The working set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the working set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from working sets. If they are needed, they will then be soft-faulted back into the working set before leaving main memory. |
| Working Set Peak | The maximum size, in bytes, in the working set of this process at any point in time. If free memory in the computer is above a certain threshold, pages are left in the working set of a process even if they are not in use. When free memory falls below a certain threshold, pages are trimmed from working sets. If they are needed, they are then soft-faulted back into the working set before they leave main memory. |

**Print Queue Object**

Windows Server 2003 provides the Print Queue performance object for tracking activity in the print server print queues. The counter values are reset when either the print server or the spooler service is restarted.

| **Counter Name** | **Description** |
|---|---|
| Add Network Printer Calls | Shows the total number of calls from other print servers to add shared network printers to this server since last restart. |
| Bytes Printed/sec | Shows the number of bytes per second printed in a print queue and provides a rough indication of how busy a printer is. This counter can be used for bottleneck detection. |
| Enumerate Network Printer Calls | Shows the total number of calls from browse clients to this print server to request network browse lists since last restart. |
| Job Errors | Shows the total number of job errors in a print queue since last restart. Job errors can occur if the connection to the printer has errors due to network problems. |
| Jobs | Shows the current number of jobs in a print queue. Use this counter to identify excessive use. |
| Jobs Spooling | Shows the current number of spooling jobs (incoming or incomplete) in a print queue. |
| Max Jobs Spooling | Shows the maximum number of spooling jobs (incoming or incomplete) in a print queue since last restart. |
| Max References | Shows the peak number of references (open handles) to this printer. |
| Not Ready Errors | Shows the total number of printer-not-ready errors in a print queue since the last restart. |
| Out of Paper Errors | Shows the total number of out-of-paper errors in a print queue since the last restart. |
| References | Shows the current number of references to a print queue. A reference can be a user or a program that is connecting to a printer and opening a print queue. |
| Total Jobs Printed | Shows the total number of jobs printed on a print queue since the last restart. |
| Total Pages Printed | Shows the total number of pages printed through GDI on a print queue since the last restart. Since RAW print jobs do not provide page-count data, they are not included in this counter. (Windows 98 and Windows 95 clients submit RAW jobs. The print volume of a RAW job is not captured by this counter.) |

Performance monitoring helps you tune your deployments and anticipate when to add more or more powerful processors, increase the system memory, or modify the disk configuration.

# Reference Systems

Updated: February 7, 2005

Applies To: Windows Server 2003

Microsoft has studied print server performance at external customer sites and also studied print server performance internally. Two sample configurations from this testing are presented in this section to provide additional reference information that can be used when sizing your print server requirements.

**Standalone Print Server**
Workgroup-class hardware like the HP/Compaq DL560 or DL580 with 1.5GB RAM, ample, high-performance disks and sufficient network throughput can host approximately 1,500 printers, serving 5,000-10,000 customers.

It is important to note that the primary limitation of the print server in this configuration is related to remote management and not the result of constraints imposed by the print server hardware or core spooler component. Specifically, when viewing and sorting columns in the **Printers and Faxes** folder, one may begin to experience delays if the server is hosting more than 1,500 printers.

Another important factor to consider is the number of client connections handled by the server. Often, a single client will have multiple connections to the print server. For example, one for each document, one for monitoring the printer status and possibly more resulting from other applications running on the client. Because of this, the factors traditionally considered in server capacity planning play a less significant role in print server capacity planning than they might when planning more traditional server functions. Factors such as CPU speed, RAM usage, Virtual Memory consumption, and Disk Read/Write times, are not constraining factors until there are more than several thousand printers per server or there are more than 10,000 clients using the server.

The table below shows the hardware configuration used by Microsoft for performance testing and benchmarking. In actual customer deployments, Microsoft expects the system configuration selected and deployed by customers to be superior to the system described in the table below. Specifically, Microsoft believes that faster processors, additional RAM, and higher capacity drives will be deployed by customers.

| | |
|---|---|
| System Manufacturer | Compaq |
| System Model | ProLiant DL580 |
| System Type | X86-based computer |
| Processor | x86 Family 6 Model 10 Stepping 4 GenuineIntel ~902 Mhz |
| Processor | x86 Family 6 Model 10 Stepping 4 GenuineIntel ~902 Mhz |
| BIOS Version/Date | Compaq P20, 4/5/2002 |
| SMBIOS Version | 2.3 |
| Total Physical Memory | 1,024.00 MB |
| Available Physical Memory | 473.55 MB |
| Total Virtual Memory | 3.41 GB |
| Available Virtual Memory | 2.47 GB |
| Page File Space | 2.41 GB |
| Disk | Compaq SCSI Drive Array (4 Ultra3 18GB disks) |

**Clustered Print Server**

Microsoft has studied clustered print server performance at many customer sites. Using this information, a clustered print server consisting of a two-node cluster was selected to test print server performance internally. Each system in the cluster was configured as described in the following table.

| | |
|---|---|
| System Manufacturer | Compaq |
| System Model | ProLiant DL380 G2 |
| System Type | X86-based computer |
| Processor | x86 Family 6 Model 11 Stepping 1 GenuineIntel ~1396 Mhz |
| BIOS Version/Date | Compaq P24, 12/11/2001 |
| SMBIOS Version | 2.3 |
| Total Physical Memory | 256.00 MB |
| Available Physical Memory | 53.66 MB |
| Total Virtual Memory | 875.91 MB |
| Available Virtual Memory | 517.17 MB |
| Page File Space | 620.38 MB |

Based on these tests, Microsoft found that such a clustered system could be used to host up to 2,500 printers and service up to 10,000 clients running a mix of Windows operating system versions.

# Resources (Windows Print Server Scalability and Sizing Technical Overview)

Updated: February 7, 2005

Applies To: Windows Server 2003

**Microsoft Hardware and Driver Central**

Includes Windows Driver Development Kits [DDK], Windows Hardware Compatibility Test [HCT] Kits, and Windows Logo Program requirements http://www.microsoft.com/whdc/default.mspx.

**Windows .NET Server Deployment Resource Kit**

For more information about deployment and management of Windows print servers, see the Windows .NET Server Deployment Resource Kit.

**Windows Driver Development Kit**

For more information about the Windows Server 2003 architecture, see the "Display and Print Devices" sections in the Windows Driver Development Kit at http://go.microsoft.com/fwlink/?LinkID=123179.

**Windows 2000 Server Resource Kit**

For more information about connectivity options and support protocols in Windows Server 2003, see Microsoft Windows 2000 Server Resource Kit published by Microsoft Press.

**Microsoft.com Web site**

For additional information about Windows print servers in general, visit the print server Web site at: http://www.microsoft.com/printserver.

For the latest information about Windows Server 2003, see the Windows 2003 Server Web site at http://www.microsoft.com/windowsserver2003.

# Best Practices for Deploying Printer Location with Active Directory

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

**In This White Paper**


Introduction (Best Practices for Deploying Printer Location with Active Directory)

Creating the Location Schema

Configuring Printer Location and Setting Location Policies

Setting Standards When Printer Location Tracking is Disabled

Summary (Best Practices for Deploying Printer Location with Active Directory)

Related Links (Best Practices for Deploying Printer Location with Active Directory)

# Introduction (Best Practices for Deploying Printer Location with Active Directory)

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

In network environments, most users do not need to know the physical location of files, databases, or Web sites. However, they do need to know the physical location of hardware devices, such as printers. The printer location tracking feature in Microsoft Windows 2003 helps users find printers by storing their actual location in the Microsoft Active Directory service. When this feature is enabled, the location string is displayed automatically for each printer.

To use this feature, your network environment must have the following characteristics:

- A Windows 2000 or Windows 2003 domain controller (which, by definition, runs Active Directory) and preferably more than one subnet.
- A network IP addressing scheme that corresponds roughly to the physical layout of your enterprise (not required but very helpful).
- Client computers that can query the directory service.
- A subnet object for each site. Subnet objects can be created using Active Directory Sites and Services.
- Your own naming convention for printer location
- Administrator privileges on the print server to set the printers location field, and permissions to create sites/subnets on the Domain Controller to populate the location field.

Group Policy lets you configure printer location tracking for a group of computers using the Computer Location and Pre-Populate Printer search location text policies.

| **Note** |
| --- |
| For more information about these policies, see the Explain tab for each policy.<br><br>For the steps to set up printer location tracking, type "printer location tracking" in the index of Windows 2003 Server Help and Support Center.<br><br>The location search ability can be very useful for your clients, as long as all printers in your organization use a consistent schema that allows users to search. The location schema used for printers should be planned out thoroughly before deployment as it will be used throughout your printing organization. This white paper outlines some best practices for creating and maintaining a useful location schema. |

The major steps involved in this effort are:

1. Verify the printer names and their subnets. It's best to keep this information in a spreadsheet or database and update it regularly. This information is helpful when upgrading print servers and for general maintenance as well.
2. Define the location tree structure.
3. Configure the printer locations.
4. Enable the printer location tracking policy.
5. Test the final setup.

# Creating the Location Schema

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

As noted, you must first create and maintain a location schema to allow users to search for location. It is important to think carefully about your location schema before deploying it. You want to make sure that it will be flexible enough to describe all locations where printers can be placed. This schema will be used by your entire organization, so its worth spending the time up front to design it properly.

Because location names are used by end users, the names should be simple and easy to recognize. Avoid using special names known only to facilities management. For better readability, avoid using special characters in a name, and keep names to a maximum of 32 characters to make sure that the whole name string is visible in the user interface. Names should conform to the following characteristics:

- Location names are in the form *name/name/name/name/...* (The slash (/) must be the dividing character.)
- A name can consist of any characters except for the slash (/).
- The number of levels to a name is limited to 256.
- The maximum length of name is 32 characters.
- The maximum length of an entire location name is 260 characters.

In particular, you should make certain that the schema is flexible enough to allow for future changes in your organization. Ideally, the schema can follow the subnet layout of your organization (assuming you have a geographical-based subnet hierarchy), because then it is possible to write scripts that will automatically populate the location field of all of your printers based on the subnet.

As an example, consider an organization with multiple locations in North America and Europe. The naming scheme for the organization might be as follows:

| **Top level** | **Level** | **Level** | **Level** |
|---|---|---|---|
| NorthAmerica | | | |
| | NewYork | | |
| | | HQ | |
| | | | Floor1 |
| | | | Floor2 |
| | | Buffalo | |
| | | | Design |
| | | | Plant |
| | | | Admin |
| | Vancouver | | |
| | | Building1 | |

| | | | |
|---|---|---|---|
| | | Building2 | |
| | Dallas | | |
| Europe | | | |
| | Paris | | |
| | | Design | |
| | London | | |
| | | ... | |
| | Madrid | | |

Note that the tree varies in depth depending on the complexity of the organization and the amount of detail available in the IP network. The naming convention for this example includes more levels than for a smaller organization located in a single city, or a single building. The full name of the design building in Buffalo, New York is **NorthAmerica/NewYork/Buffalo/Design**; for the design building in Paris, the full name is **Europe/Paris/Design**, and the smaller office in Dallas has the name as **NorthAmerica/Dallas**.

# Configuring Printer Location and Setting Location Policies

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

To enable printer location tracking and set the automatic print population policy, do the following:

1. Open **Active Directory Sites and Services**.

2. Create the sites you want, if you haven't done so already.

3. Associate a subnet for this newly created site, if you haven't done so already.

4. Right-click each subnet you created, and then click **Properties**.

5. Click the **Location** tab, type the location of the subnet in the following format based on your location naming convention, and then click **OK**:

   *location/location*

6. Using Group Policy, double-click the **Pre-populate printer search location** policy setting, and then click to enable it.

   The following example demonstrates how to set the policy to a group (organizational unit in Active Directory terms) of users that use that printer. However, you can set the policy to a different object using the appropriate Active Directory utility.

   a. Start **Group Policy** from **Active Directory Users and Computers**.

   b. Right-click the organizational unit you want to apply the policy to, and then click **Properties**.

   c. Click the **Group Policy** tab, and then click **New** to create a new object. Type the name of the new object, and then click **OK**.

   d. Select the newly created object, and then click **Edit**.

   e. Expand **Computer Configuration**, then **Administrative Templates**, and then **Printers**.

   f. Double-click the policy setting **Pre-populate printer search location**, and then click to enable it. When you search the directory service for printers, it fills in your location field using the value specified in the Location tab of the Subnet properties page. If this policy is disabled (default), the location will be empty.

   | **Note** |
   | --- |
   | It is also possible to set location manually per-computer by Group Policy. |

7. Open **Printers and Faxes**.

8. Right-click the printer and then click **Properties**.

9. Type the printer location based on your schema, or click **Browse** to find it. You probably want to be more precise in describing the printer location than the subnet location. For example, for the subnet location NorthAmerica/Dallas, you might type *NorthAmerica/Dallas/Floor42/Room4207*.

10. Repeat the last two steps for each printer in Active Directory.

It is recommended to script Steps 710 by using the prnadmin.dll utility that ships in the resource kit. To use this utility, you must refer to your existing spreadsheet or database of printers and their locations and then, rather than opening each printers properties to set the location, you can use the tool to set the printer locations according to the inputted information.

The following is an example of how one can set the location using prnadmin.dll. Typing **cscript prncfg.vbs /?** in the command prompt will show this example.

```
cscript prncfg.vbs -s -b PrinterName -l "Building A/Floor 100/Office 1"
```

Where PrinterName can be a UNC name.

| **Note** |
| --- |
| To open Active Directory Sites and Services, click **Start**, point to **Programs**, point to **Administrative Tasks**, and then click **Active Directory Sites and Services**. |

# Setting Standards When Printer Location Tracking is Disabled

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

As mentioned previously, you can add information to **Location** in the **General** tab of the printer property page even if printer location tracking is disabled. However, this imposes a potential problem. When users search for printers on the tenth floor, they need to know exactly what to type in the **Find Printers** property page when they click **Start**, then **Search**, and then **For Printers**. Incidentally, wildcards such as * (asterisk) may be used to help find printers when exact or complete location names are unknown. When printer location tracking is enabled, the **Location** field in the **Find printer** property page is filled in automatically.

# Summary (Best Practices for Deploying Printer Location with Active Directory)

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

As described in the preceding sections, the location search ability can be very useful for your clients, as long as all printers in your organization use a consistent schema that allows users to search. Following the recommendations in this white paper will help create a consistent and useful schema that you can deploy throughout your print organization.

# Related Links (Best Practices for Deploying Printer Location with Active Directory)

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

See the following resources for further information:

For more information about using the Active Directory for printing, choose the Active Directory and Printing link at http://www.microsoft.com/windows2000/technologies/fileandprint/print/default.asp

- For more information about Windows Print Servers in general, visit the Print Server Web site at http://www.microsoft.com/windows2000/technologies/fileandprint/print/default.asp

- For the latest information about Windows 2003 Server, see the Windows Server 2003 Web site at http://www.microsoft.com/windowsserver2003/default.mspx.

# Certificate Autoenrollment in Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

**In This White Paper**


Introduction (Certificate Autoenrollment in Windows Server 2003)

How Autoenrollment Works

Configuring an Enterprise CA

Configuring Group Policy

User Autoenrollment

Certificate Renewal

Autoenrollment Functions

Updating Group Policy

Advanced Features

Supported Hardware (Certificate Autoenrollment in Windows Server 2003)

Troubleshooting (Certificate Autoenrollment in Windows Server 2003)

Summary (Certificate Autoenrollment in Windows Server 2003)

Related Links (Certificate Autoenrollment in Windows Server 2003)

# Introduction (Certificate Autoenrollment in Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Microsoft Windows Server 2003, Enterprise Edition introduces the capability to automatically enroll users and computers for certificatesincluding smart cardbased certificates.

Using the autoenrollment feature, organizations can manage the certificate lifecycle for users, which includes:

- Certificate renewal
- Superseding of certificates
- Multiple signature requirements

Certificate autoenrollment is based on the combination of Group Policy settings and version 2 certificate templates. This combination allows the Windows XP Professional or Windows Server 2003 client to enroll users when they log on to their domain, or a machine when it boots, and keeps them periodically updated between these events.

Automatic enrollment of user certificates provides a quick and simple way to issue certificates to users and to enable public key infrastructure (PKI) applications, such as smart card logon, Encrypting File System (EFS), Secure Sockets Layer (SSL), Secure/Multipurpose Internet Mail Extensions (S/MIME), and others, within an Active Directory directory service environment. User autoenrollment minimizes the high cost of normal PKI deployments and reduces the total cost of ownership (TCO) for a PKI implementation when Windows XP Professional clients are configured to use Active Directory.

**Supports Pending Certificate Requests and Certificate Renewal**
User autoenrollment in Windows XP Professional and Windows Server 2003 supports both pending certificate requests and renewal features.

**Requesting a Certificate**
You can manually or automatically request a certificate from a Windows Server 2003 certification authority (CA). This request is held until administrative approval is received or the verification process is completed. Once the certificate has been approved or issued, the autoenrollment process completes and installs your certificate automatically.

**Renewing a Certificate**
The process for renewing an expired user certificate also takes advantage of the autoenrollment mechanism. Certificates are automatically renewed on behalf of the userdependent on the specifications in the certificate template.

**Dependencies**
The autoenrollment feature has several infrastructure requirements. These include:

- Windows Server 2003 schema and Group Policy updates
- Windows 2000 Server domain controllers running Service Pack 3 or later
- Windows XP Professional or Windows Server 2003 clients
- Windows Server 2003, Enterprise Edition or Datacenter Edition running as an Enterprise CA

| **Note** |
| --- |
| Autoenrollment Group Policy settings may only be applied when using a Windows XP Professional or Windows Server 2003 computer. |

**Topics Covered**

- How Autoenrollment Works
- Configuring Certificate Templates
- Configuring an Enterprise CA
- Configuring Group Policy
- User Autoenrollment
- Certificate Renewal
- Autoenrollment Functions
- Updating Group Policy
- Advanced Features
- Supported Hardware
- Troubleshooting

# How Autoenrollment Works

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section discusses how autoenrollment works, including autoenrollment and Winlogon, an analysis of the components of the autoenrollment process, and working with certification authority interfaces.

**Key Points**

Autoenrollment works best in a Windows Server 2003 Enterprise environment where the Windows XP client is integrated with Active Directory.

Only domain-joined machines can use certificate autoenrollment. Although the autoenrollment process does not explicitly look for domain-joined machines, the Winlogon process will not activate userinit.exe unless the machine/user is part of a domain.

**Autoenrollment Timing**

The autoenrollment process is normally triggered by the Winlogon process, and is designed to be activated and managed by a domain-based Group Policy. Both machine-based and user-based Group Policy can activate autoenrollment for machines and users. By default, the Group Policy is applied at reboot for machines, or at logon for users, and is refreshed every eight hours. The refresh interval can be configured using Group Policy. Autoenrollment is also triggered by an internal timer that activates every eight hours after the last time autoenrollment was activated.

For additional information, see Updating Group Policy.

| **Note** |
| --- |
| Unlocking the workstation does not trigger autoenrollmentonly a full interactive logon or a Group Policy refresh will initiate the Winlogon trigger. |

**The Autoenrollment Process**

The autoenrollment feature handles all aspects of certificate enrollment, renewal, and certificate housekeepingexcept in the case where user interaction is explicitly defined on a certificate template in Active Directory. When the autoenrollment process is triggered by Winlogon or a Group Policy refresh interval, the operating system queries Active Directory to download the appropriate certificate stores into the local store on the client machine; for example, root CA certificates, cross-certificates, and the NTAuth container. The autoenrollment process also downloads certificate templates from the forest and caches the list in the registry at the same time. The last step performed by autoenrollment is user-object cleanup (userCertificate attribute) in Active Directory. Revoked, expired, and superseded certificates are removed from the user object automatically; however, expired certificates are not removed unless a new valid certificate is issued at the same time. Certificates in the local user profile or on the user object in Active Directory are only managed if the certificate corresponds to a certificate template in Active Directory. Foreign certificates and certificates that do not contain the template extension are not managed. This is a transparent activity that is processed asynchronously.

**Requirements List**

The autoenrollment process will then process the list of templates and create a requirements list for any templates that have an autoenroll access control entry (ACE) set on the template for the current machine or user. The machine and/or user must also have the Read ACE set on the template or the template will not be enumerated. The users or machines MY (personal) store will also be processed at this time to look for revoked certificates, certificates without private keys, time invalid certificates and so on, and add these certificates to the requirements list. For more information about certificate stores, refer to the Microsoft Platform SDK: http://go.microsoft.com/fwlink/?LinkID=116205

It is very possible that a user may have a certificate in the MY store but not have permissions set on a template access control list (ACL) in Active Directory. These will be processed and added to the list, but enrollment will most likely fail due to the fact that the template permissions do not allow enrollment/renewal at the updated point in time.

Items in the requirements list may be removed if an appropriate valid certificate is found in the MY store. If a certificate template is marked to check Active Directory for an existing certificate, Active Directory will be queried for an existing duplicate certificate on the userCertificate attribute of the user object and the requirement will be removed from the list, if successful.

| **Note** |
| --- |

Checking Active Directory for the presence of an existing certificate associated with the user or machine object can affect performance and may delay autoenrollment processing due to the network and directory requirements for performing this operation. This is because the actual certificates in the userCertificate attribute will be downloaded and examined. When this happens, the directory cannot be queried via Light Weight Directory Access Protocol (LDAP) to simply respond whether a given certificate type exists without downloading and processing the certificates locally.

**Autoenrollment also manages the CryptoAPI REQUEST store for the user.** This process enumerates each pending request in the store and then installs the pending certificate, if possible, from the issuing CA. If a certificate is to be archived or deleted, based on the certificate template rule, it will be processed as follows:

- If a request already exists in the REQUEST store, this certificate will be removed from the summarized requirements list.
- If a request has been pending for more than 60 days, the request will be deleted and the requirements list will remain as-is.

Autoenrollment can be used to retrieve pending requests only for certificates with template information, for example, an initial request involving a certificate template. The autoenroll ACL on the certificate template is not necessary for the autoenrollment process to retrieve a pending certificate request. If the user enrolls via a Web page and the certificate request is pending, autoenrollment will retrieve the pending request for the user.

Template supersede rules will be evaluated and appropriate additions and deletions will be processed for the requirements list. For example, if the template says "X supersedes Y ", it means that if you have been told to enroll for X and Y, you really only need X. If you only have Y, you still must get X. This is the last step in rule processing. After it is done, the requirements list is complete.

**For each template that does not require user interaction**, the autoenrollment process will create the requests in the background and submit them to a CA. Once this is done, the requirements list is updated.

Autoenrollment always performs a revocation check of the entire certificate chain starting with the issuing certification authority to ensure that the CA offering enrollment services is not revoked before performing enrollment. If the CA is revoked, autoenrollment will not send requests to that certification authority. However, autoenrollment will ignore revocation errors if a CDP (CRL Distribution Point) extension does not exist in the CA certificate or if the revocation status is offline.

If a certificate is issued from the CA, it is installed in the users or machines MY store. If the certificate is pended [specified by the CA Manager approval check box in the Certificate Template Microsoft Management Console (MMC) snap-in], the request information is saved in the REQUEST store.

**Balloon User Interface**

For each request that requires user interaction as per the certificate template, the balloon user interface (UI) is invoked.

- Approximately 60 seconds after logon, the balloon UI is displayed. If no user interaction is explicitly defined on the certificate template, no UI will be displayed to the user. This delay is incorporated to allow for speedy application and shell response times during the logon and booting of the client machine.
- If the 60-second delay is not desired, the following registry key may be added on a per-user basis.

HKEY_CURRENT_USER\SOFTWARE\Microsoft\Cryptography\AutoEnrollment\AEExpress

Using this key in a normal production environment is not recommended. If it is used, it must be created on a per-user basis.

| **Important** |
| --- |
| Machine certificates do not support user interaction and should not be configured to require this setting. |

- The balloon UI waits for the user to see the balloon and is activated by a mouse click. Note that after approximately 15 seconds, the balloon pop-up window is replaced in the notification area by a certificate icon that may be activated by a mouse click.
- If no activation occurs within seven hours, the taskbar icon will disappear and the silent thread will re-activate at the next

logon, machine reboot, or Group Policy refresh interval, whichever is first.

- Once the user activates the UI, the REQUEST store is checked first for pending requests.

**Issuing the Template**
Once a certificate template with the proper ACE has been enumerated, the autoenrollment process will search for a Microsoft Enterprise Certification Authority in Active Directory that can issue the template. If more than one Enterprise CA is found, the client will try each CA in the list in random order (for load balancing) until a CA responds and is able to issue a certificate.

The client contacts a CA through a Distributed Component Object Model (DCOM) interface and supplies a security context through DCOM to provide an authenticated request. The default policy module of the Microsoft Enterprise CA enforces certificate profiles and enrollment security as defined by the templates.

If the certification authority is set to pend the request for an administrator or certificate manager to examine and approve, autoenrollment will periodically query the CA during every Group Policy refresh interval for approved requests. Autoenrollment will also re-enroll templates when **Reenroll all certificate holders** has been set in Group Policy. For more information, see Certificate Renewal.

**Certification Authority Interfaces**
The following methods are used by the autoenrollment process for contacting and enrolling against a Microsoft Enterprise CA.

- GetCAProperty
- Submit
- GetLastStatus
- GetRequestId
- GetFullResponseProperty
- GetCertificate
- Release
- RetrievePending

These methods can be found in the Platform SDK at http://msdn.microsoft.com/library/default.aspConfiguring the Certificate Templates

This section covers how to configure certificate templates and provides a step-by-step example of how to create a new template for the autoenrollment of a smart card. Certificate template permissions are also explained.

**Key Point**
A version 2 certificate template must first be created in Active Directory to enable autoenrollment.

**Default Settings**
The following are default settings.

- Only root domain administrators or explicitly delegated users in Active Directory may configure templates in a domain that has been upgraded from Windows 2000 Server.
- Both domain administrators from the root domain and enterprise administrators for fresh installations of Windows Server 2003 domains may configure templates.
- Certificate template ACLs are viewed in the Certificate Templates MMC snap-in.
- Certificate templates can be cloned or edited using the Certificate Templates MMC snap-in.

| **Note** |
| --- |
| Only a domain with the Windows Server 2003 schema will support version 2 templates, and only a Windows Server 2003, Enterprise Edition or Datacenter Edition certification authority may issue a version 2 template certificate. |

**Creating a New Template for the Autoenrollment of a Smart Card**

To create a new template for autoenrollment of a smart card

1. Log on as a domain administrator.
2. Click the **Start** button, and then click **Run**.
3. In the **Run** dialog box, in the **Open** box, type **mmc.exe**, and then click **OK**.
4. On the **File** menu, click **Add/Remove Snap-in**.
5. In the **Add/Remove Snap-in** dialog box, click **Add**.
6. In the **Add Standalone Snap-in** dialog box, click **Certificate Templates**, and then click **Add**.

   | **Note** |
   | --- |
   | The Certificate Templates MMC snap-in is available on the Server version of Windows Server 2003 or on Windows XP Professional through the Administration Tools Pack installation on the Server media. |

7. Click **Close**.
8. Click **OK**.

   | **Note** |
   | --- |
   | The Certificate Templates MMC snap-in may also be invoked using the Certification Authority MMC snap-in by selecting the **Certificate Templates** folder, right-clicking, and then selecting **Manage**. |

9. In the console tree, click **Certificate Templates**.
10. In the details pane, right-click the **Smartcard User** template, and then click **Duplicate Template** (Figure 1).

    The **General** tab of the new template properties dialog box appears.




**Figure 1: Creating a New Template for Autoenrollment of a Smart Card**

11. In the **Template display name** field, type a unique name for the template (Figure 2).




**Figure 2: Naming the Template**

| **Note** |
| --- |
| If the **Do not automatically reenroll if a duplicate certificate exists in Active Directory** is enabled, autoenrollment will not enroll a user for the certificate template, even if a certificate does not exist in the users MY store. Active Directory is queried and determines if the user should be enrolled. This is an extremely valuable feature for users who do not have roaming profiles and log on to multiple machines. Without this setting and without roaming profiles, the user will automatically be enrolled on every machine that is logged on to (including servers). |

12. Click the **Request Handling** tab (Figure 3).

This tab is used to define how the certificate request should be processed, including the cryptographic service providers (CSP) and minimum key sizes that will be used by the enrollment template.




**Figure 3: Defining How the Certificate Request Should Be Processed**

| **Important** |
| --- |
| If more than one smart card CSP is made available on this tab, the user may be prompted for every CSP that is selected when enrolling for this template. The behavior may vary depending on the CSPs available on the client machine. If the user has only one smart card, the prompts for the unavailable CSPs will have to be cancelled. This behavior is by design. It is also important to select a minimum key size that is supported by the selected CSP; otherwise, enrollment will fail. |

13. Select the **Prompt the user during enrollment** check box (Figure 3). (Enrollment for smart card requires user input to succeed.)

| **Important** |
| --- |
| If the certificate template is not going to be used for smart cards or if it is not desired for the user to be prompted to enroll for certificates, this option is not required. Machine certificates should not have this enabled or machine autoenrollment will fail. |

14. Click the **Subject Name** tab.

    This tab is used to define how the subject name and certificate properties will be built. It is recommended to use the default selections when enrolling for a smart card template.

15. Click the **Extensions** tab.

    This tab is used to define how the various extensions will be added to this certificate template during enrollment. It is recommended that you use the default settings.

| **Note** |
| --- |
| Application Policies is a replacement for Extended Key Usage (EKU) in Windows Server 2003, although EKU is still supported for legacy applications and client operating systems. |

16. Click the **Security** tab.

    This tab is used to define which users or groups may enroll or autoenroll for a certificate template. A user or group must have the Read, Enroll, and Autoenroll permissions to automatically be enrolled for a certificate template.

17. Click **OK** when finished.

    The Autoenrollment column should now show **Allowed**.

| **Note** |
| --- |
| The Autoenrollment column will automatically indicate if a template is suitable for autoenrollment. The user will not be able to change the status. If a template has more than one registration authority (RA) signature required or if the subject is supplied by the requester, autoenrollment will not be allowed. The column does not reflect the state of the autoenroll ACL on the template. |

| **Important** |
| --- |
| A smart card that has been issued on Windows XP client or Windows XP Enrollment Station (RA) might only be compatible with a Windows XP client. The compatibility varies between various smart card and CSP vendors. If you need to be able to log on to both Windows 2000 and Windows XP, see your smart card and CSP vendor for additional information. |

**Certificate Template Permissions**

For a user or computer to enroll for a certificate template, it must have appropriate permissions (ACEs) set on the template in Active Directory. A user or computer must have both Enroll and Read permissions to enroll for a selected certificate template.

- **Read permission** allows the template to be discovered by the user.

- **Enroll permission** is enforced by the Enterprise CA when a user requests a certificate for a selected template. The Enterprise CA must also have Read permissions on a template to enumerate the template in the directory and issue certificates based on that template. Normally, the Enterprise CA is included in the Authenticated Users group, which has Read permissions by default on a template.

- **Full Control permission** is given to enterprise administrators and the primary domain administrators group by default when installing a fresh Windows Server 2003 domain. If a domain has been upgraded from Windows 2000, enterprise administrators will not have this permission by default. The Full Control permission allows a user to set or modify the permissions on a selected template.

- **Autoenroll permission** is set on a template when it is desired for a user or computer to automatically enroll for a selected certificate template. The Autoenroll permission is needed in addition to the Enroll permission for a user to enroll for a given certificate template. Only version 2 templates or newly created templates may have the Autoenroll ACE set.

- **Write permission** allows a user to modify the contents of a certificate template. Note that only version 2 certificates with a Windows Server 2003 schema may be modified. Version 1 certificate templates only allow ACLs to be modified.

# Configuring an Enterprise CA

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section shows how a Microsoft Enterprise CA must be configured to issue a certificate template after it has been created.

**How to Configure an Enterprise CA**
To configure an Enterprise CA

1. On the **Administrative Tools** menu on a Windows Server 2003 console, open the **Certification Authority** snap-in.

2. In the console tree, expand the Certification Authority, and then click **Certificate Templates** (Figure 4).




**Figure 4: Selecting Certificate Templates**

3. Right-click **Certificate Templates**, click **New**, and then click **Certificate Template** to Issue.

4. In the **Enable Certificate Templates** dialog box (Figure 5), click **Auto Enroll Smartcard User** (this is the certificate template name that was created previously), and then click **OK**.

   (See Creating a New Template for the Autoenrollment of a Smart Card.)




**Figure 5: Enabling a Certificate Template on a CA**

| **Note** |
| --- |
| Due to replication latency and template caching in the registry, a certification authority may not be able to issue a certificate template immediately. The timing of issuance is dependent on replication latency between domain controllers. |

5. Close the **Certification Authority** MMC snap-in.

# Configuring Group Policy

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section shows how to configure the Group Policy settings for a site, domain, or organizational unit (OU).

**How to Configure Group Policy**
Group Policy settings for a site, domain, or OU must be configured to enable certificate autoenrollment in a domain.

To configure Group Policy

1. Open the Active Directory **Users and Computers** MMC snap-in.

2. Right-click the site, domain, or OU that you want to configure Group Policy for, and then click **Properties**.

3. Click the **Group Policy** tab, and then click **Edit** (Figure 6).




**Figure 6: Selecting Group Policy Configuration Options**

| **Note** |
| --- |
| Machine policy for automatic enrollment of machine and domain controller certificates is configured identically, even though it is controlled through the machine policy of a Group Policy object. |

4. Click **User Configuration**, **Windows Settings**, **Security Settings**, and finally **Public Key Policies**. In **Object Type**, right-click **Autoenrollment Settings** (Figure 7), and then click **Properties**.




**Figure 7: Selecting Autoenrollment Settings**

5. Ensure that **Enroll certificates automatically** is selected as well as the two check boxes under this option (Figure 8). Automatic renewal, certificate cleanup, and publishing in Active Directory are only enabled with all options selected.

| **Note** |
| --- |
| Both machine and user policy must be configured to enable certificate enrollment for both types. |




**Figure 8: Selecting Autoenrollment Settings**

6. Click **OK**.

   Autoenrollment is now enabled.

# User Autoenrollment

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section illustrates manually pulsing autoenrollment and smart card enrollment.

**Key Points**
User autoenrollment for a smart card requires mandatory manual steps and user interaction, unlike other certificate types. Once autoenrollment has been enabled, the user will receive an informational balloon on the taskbar at the next Group Policy pulse interval (default of eight hours) or at the next logon.

**Manually Pulsing Autoenrollment**
Autoenrollment may be pulsed manually through the Certificates MMC snap-in.

To manually trigger autoenrollment

1. Log on to the domain with the appropriate user account.

2. Click the **Start** button, and then click **Run**.

3. Type **mmc.exe**, and press ENTER.

   An empty MMC shell starts.

4. Select the **File** menu, and then select **Add/Remove Snap-In**.

   A dialog box appears with a list of the snap-ins that have been added to the MMC shell.

5. Click **Add**.

   A list of the registered snap-ins on the current machine appears.

6. Double-click the **Certificates** snap-in, select **My User Account**, and then click **Finish**. If enrolling the machine for a certificate, such as a domain controller or a Web server, select **Computer account**.

7. Click **Next**.

8. Select **Local Computer**, and then click **Finish**.

9. In the **Add Standalone Snap-in** dialog box, click **Close**. In the **Add/Remove Snap-in** dialog box, click **OK**.

   The MMC now contains the personal certificate store for the user.

10. Right-click the top of the tree on **CertificateCurrent User**, select **All Tasks** on the context menu, and then select **Automatically Enroll Certificates** (Figure 9).




**Figure 9: Automatically Enrolling Certificates**

| **Note** |
| --- |
| It will take approximately one minute for the Certificate Enrollment balloon to be displayed, unless the registry key mentioned previously has been set. (See Balloon User Interface.) |

**Smart Card Enrollment**

1. Click the balloon or the corresponding certificate icon in the notification area once the **Certificate Enrollment** balloon is displayed. After a short period of time, the balloon will automatically disappear, and only the icon in the notification area will remain. After clicking the balloon, the Autoenrollment UI will start.

| **Note** |
| --- |
| The certificate enrollment balloon and wizard are not only for smart card enrollment but also for self-registration authority. |

2. Click the **Start** button (Figure 10).

   The wizard will begin. (The **Remind Me Later** button will cause the Certificate Enrollment balloon to re-appear at the next Group Policy pulse interval or the next interactive logon.)




**Figure 10: Begin Enrolling Certificates**

3. If a smart card with the required CSP on the certificate is not inserted in the smart card reader, the user will be prompted to insert a smart card. Insert the smart card and click **OK**.

| **Note** |
| --- |
| If the certificate template on the users machine contains more than one CSP, the user may have to cycle through the wizard to reach the desired smart card CSP. |

4. If the displayed smart card CSP is not the desired CSP, click **Cancel**.

   The next CSP listed in the certificate template will be displayed.

5. After inserting an appropriate smart card, click **OK**.

   The wizard will continue.

| **Note** |
| --- |
| If the smart card contains a private key and certificate in Slot0 (default container), the user will be warned about replacing the credentials on the smart card. |

| **Important** |
| --- |
| Due to limitations with the smart card CSPs, smart card logon with both Windows 2000 and Windows XP requires that Slot0, or the default container on the card, be used to hold the certificate and private key used for smart card logon. If the card contains multiple keys and certificates, the last generated key and certificate will be marked as the default container on the card. |

6. If it is desired to replace the credentials on the smart card, click **Yes** (Figure 11).

   The wizard will continue (Figure 12).

   (The following is only one example of the dialog box a user may see. The smart card UI varies depending on the CSP being used.)




**Figure 11: Replacing Credentials Dialog Box**




**Figure 12: Enrolling Certificates**

7. If a PIN is necessary for the smart card, a dialog box will be displayed. Enter the appropriate PIN and click **Enter**.

   Enrollment completes.

| **Note** |
| --- |
| If you enter the PIN incorrectly, the number of times you can retry will be limited. |

The success or failure of the autoenrollment process will be logged in the Application event log on the local computer.

Also, a summary dialog box will appear for failed certificate requests that involved user interaction. If a failure occurs during enrollment, the user will be notified of the failure (Figure 13).




**Figure 13: Notifying the User of Errors While Enrolling Certificates**

| **Note** |
| --- |
| Users are not prompted when enrollment succeeds. |

# Certificate Renewal

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section discusses renewal intervals, forced re-enrollment, and smart card renewal.

**Renewal Intervals**
Windows XP Professional or Windows Server 2003 clients, when combined with a Windows Server 2003, Enterprise Edition certification authority, will perform automatic renewal of certificates as specified on a per-template basis. Renewal intervals are dictated by the certificate template, which is set to six weeks (before expiration) by default. When certificate renewal is performed, the old (previous) certificate enrollment is always archived automatically on the client machine, and the user directory object is updated.

Important certificate renewal criteria include the following:

- Automatic certificate renewal will only occur when 80 percent of the certificate lifetime has passed, or when the renewal interval period specified on the template has been reachedwhichever timeframe is smaller.

- If the renewal period is greater than 20 percent of the certificate lifetime, autoenrollment will not automatically attempt certificate renewal until the 80 percent threshold has been reached.

**Forcing Re-Enrollment**
An administrator may force all users to re-enroll for a given template by updating the major version number of the template. When Active Directory is queried during logon for required certificate templates, the version number is examined. If the version number has incremented, the certificate template is considered to be updated and the user must re-enroll for that template.

To manually force the template version to be updated (thereby forcing re-enrollment)

- Right-click the template and select **Reenroll All Certificate Holders** (Figure 14).

| **Note** |
| --- |
| Templates are not updated automatically. By default, templates are updated at a minimum interval of 10 minutes. |




**Figure 14: Manually Forcing Certificate Re-Enrollment**

**Smart Card Renewal**

The renewal behavior of a smart card may vary depending on the type of smart card CSP being used and the state of the card at the time of renewal. In general, if the smart card being used has available space for an additional enrollment and the CSP supports multiple keys on a single card, autoenrollment will request the card to generate a new key for enrollment. If this succeeds, the certificate is written to the card and the container is marked as default. The default container is the only container that the Winlogon process will enumerate for a smart card logon certificate and key. If the smart card or CSP cannot generate a new key on the card, the existing key will be reused and a new certificate will be forced onto the card. This action will generate an event in the machines application event log.

| **Note** |
| --- |
| Autoenrollment will always use a newly generated key for all enrollment and renewal requests. The only exception to this rule is in the case of some smart card CSPs that cannot support a new key due to storage limitations on the smart card. If a key is reused, an event will be entered in the Client application log. |

**Revoked Certificates and Renewal**

Revoked certificates may not be renewed and may not be used to sign a renewal request. This scenario is explicitly blocked by autoenrollment. In this scenario, a user must perform a new manual enrollment request instead of renewal.

# Autoenrollment Functions

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section discusses various functions performed by the autoenrollment process on Active Directory domain-joined machines.

**Download of Active Directory Certificates and Trust Objects**
Autoenrollment automatically downloads and manages trusted root certificates, cross-certificates, and NTAuth certificates from Active Directory into the local machine registry for domain-joined machines. All users who log on to the machine inherit the trust and downloaded certificates that are downloaded and managed by autoenrollment.

**Deleting Expired and Revoked Certificates**
Autoenrollment deletes expired and revoked certificates in the userCertificate attribute on the user object in Active Directory. This feature can be enabled through user or machine Group Policy to help ensure that only valid and active certificates are used for encryption operations.

The exit module on the Windows Server 2003 CA also helps to manage the user account in Active Directory, but only deletes expired certificatesit does not remove revoked certificates due to performance reasons. In general, there is no value in publishing a signing certificate to the user object in Active Directory, except for purposes of record-keeping.

**Managing User Certificates in the CryptoAPI MY Store**
Certificates in the users local **MY certificate store** may also be managed through the autoenrollment process. On a per-template basis, autoenrollment can be enabled to delete expired and revoked signature certificates. Encryption certificates and keys are never automatically deleted. However, autoenrollment only manages certificates that correspond to certificate templates defined in Active Directory that contain the certificate template extension. This feature is enabled by setting this policy on the Request Handling tab in the Properties of a given certificate template (Figure 15).




**Figure 15: Managing Certificates**

If the Delete revoked or expired certificates check box is not enabled, autoenrollment will archive all expired and revoked certificates in the users MY store. As mentioned previously, this setting does not affect the userCertificate attribute in Active Directory for the user object.

| **Note** |
| --- |
| This feature is only enabled with encryption keys, not signature key types, which are not normally published to Active Directory. |

# Updating Group Policy

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section discusses how to update User and Machine Group Policy.

**User and Machine Group Policy**
User autoenrollment is triggered by the Winlogon process (interactive logon with CTRL+ALT+DELETE keys) or at Group Policy refresh intervals. Normally, User Group Policy is refreshed at logon and Machine Group Policy is refreshed at machine reboot. Group Policy may be manually refreshed using the gpupdate.exe tool that is included in Windows XP.

**Manually Forcing a Group Policy Update**
To manually force a Group Policy update (for a user or a machine)

- Run gpupdate.exe with the -Force option at the command line (Figure 16).




**Figure 16: Forcing a Group Policy Update**

# Advanced Features

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section discusses templates that require certificate manager approval, self-registration authority, and how to supersede a certificate template.

**Requiring Certificate Manager Approval**
A specific certificate template can require that a certificate manager (CA officer) approve the request prior to the CA actually signing and issuing the certificate. This advanced security feature works in conjunction with autoenrollment and is enabled on the Issuance Requirements tab of a given certificate template (Figure 17). This setting overrides any pending setting on the CA itself.

Once certificate manager approval is required, all automatic enrollment requests are not issued until a certificate manager manually approves the request.




**Figure 17: Setting the Requirement for Certificate Manager Approval**

The autoenrollment process will periodically check the CA for approved requests and install the certificates automatically. Smart cards, user certificates, and machine certificates support pending requests. In the case of smart cards, the user will be prompted to insert the smart card when the certificate is issued so that the certificate may be written to the card.

| **Note** |
| --- |
| The autoenrollment process supports a maximum of one signature requirement on the template. This limitation exists to support the self-registration authority feature described in Self-Registration Authority. If multiple signatures are desired for a given certificate enrollment, manual enrollment should be used. |

**Self-Registration Authority**
The self registration authority (Self RA) is an advanced feature of certificate enrollment that may be combined with the autoenrollment process.

Self RA refers to certificate enrollment based on the existence of a previously enrolled certificate in which the users private key is used to sign the new certificate request. The Certificate Management Messages over CMS (CMC) protocol provides for this feature where one or more signatures may be used or required for a given certificate enrollment. Self RA requirements are defined in a certificate template which may be managed using the Certificate Templates MMC snap-in.

- To add an issuance (signature) requirement to a certificate template, open the template and click the **Issuance**

**Requirements** tab.

To add a signature or issuance requirement, select the **This number of authorized signatures** check box and add the appropriate number in the following number field (Figure 18).

Now you may add specific requirements for the signing certificate.



**Figure 18: Setting the Number of Authorized Signatures**

The previous setting is a useful configuration for customers who want to manually enroll users for smart cards with an enrollment station. Then they can supersede the original template with a new template with the previous settings to allow automatic renewal through the autoenrollment process, which will require the user to sign the renewal request with the old certificate.

**Additional Self RA Example:** You could add the Application Policy for a smart card logon certificate that would be used to enroll for an EFS certificate. This would mandate that users sign their request for an autoenrolled EFS certificate with a valid smart card certificate. The user would then be prompted to insert a smart card and enter a PIN when autoenrollment was activated for the EFS certificate.

**Superseding Certificate Templates**

Certificate autoenrollment also supports the concept of superseding a template or a previously enrolled certificate. Superseding a template allows an administrator to re-enroll, change, or combine previously issued certificate enrollments into a new certificate enrollment. Autoenrollment always examines existing certificates in the users store and determines if the template used in the issued certificate has been superseded. If a certificate template has been superseded, the user will automatically be enrolled with the new template, and the old certificates will be deleted or archived depending on the template setting.

**Benefits**

Superseding certificate templates is especially useful in the following scenarios.

- Changing certificate lifetime
- Increasing key size
- Adding extended key usage or application policies
- Correcting enrollment policy errors
- Updating users from version 1 templates to version 2 templates

To create or modify a template to supersede an existing certificate

1. Open the **Properties** of the template to take precedence, click the **Superseded Templates** tab (Figure 20), and then click **Add**.

2. Select the template you want to supersede (Figure 19), and then click **OK**.




**Figure 19: Selecting a Template to Supersede**

3. The template will be added to the **Superseded Templates** tab (Figure 20). If you wish to add additional templates that should be superseded with this new template, click **Add** and repeat. Otherwise, click **OK**.




**Figure 20: Listing Superseded Templates**

| **Note** |
| --- |
| Superseding a certificate always generates a new private key for the user or machine. |

# Supported Hardware (Certificate Autoenrollment in Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section lists the smart card readers and smart cards that are supported in the default installations of Windows 2000 and Windows Server 2003.

**Table 1 Smart Card Readers**

| Name | Type | Common Name | Windows 2000 | Windows XP/Server 2003 |
|---|---|---|---|---|
| BULL SmarTLP3 | Serial | | Yes | Yes |
| GEMPLUS GCR410P | Serial | GemPC410 | Yes | Yes |
| GEMPLUS GPR400 | PCMCIA | GemPC400 | Yes | Yes |
| Litronic 220 | Serial | | Yes | Yes |
| Rainbow Technologies SCR3531 | Serial | | Yes | No |
| Schlumberger Reflex 20 | PCMCIA | | Yes | Yes |
| Schlumberger Reflex 72 | Serial | | Yes | Yes |
| SCM Microsystems SCR120 | PCMCIA | | Yes | Yes |
| SCM Microsystems SCR200 | Serial | | Yes | Yes |
| SCM Microsystems SCR300 | USB | | Yes | Yes |
| GEMPLUS GemPC430 | USB | | No | Yes |
| HP ProtectTools | Serial | | No | Yes |
| Schlumberger Reflex Lite | Serial | | No | Yes |
| SCM Microsystems SCR111 | Serial | | No | Yes |
| Systemneeds External | Serial | | No | Yes |
| Omnikey AG CardMan 2020 | USB | Utimaco CardMan | No | Yes |
| Omnikey AG CardMan 4000 | PCMCIA | Utimaco CardMan | No | Yes |
| Omnikey AG CardMan 2010 | Serial | Utimaco CardMan | No | Yes |

**Table 2 Smart Cards**

| Name | Capacity | Windows 2000 | Windows XP/Server2003 |
|---|---|---|---|

| | | | |
|---|---|---|---|
| Gemplus GemSAFE | 4k | Yes | Yes |
| Schlumberger Cryptoflex | 4k | Yes | Yes |
| Schlumberger Cryptoflex | 8k | Yes | Yes |
| Gemplus GemSAFE | 8k | No | Yes |
| Infineon (formerly Siemens) | | No | Yes |
| Schlumberger Cyberflex Access | 16k | No | Yes |

| **Note** |
|---|
| The GEMPLUS GemPC430 is not supported on Windows Server 2003. |

# Troubleshooting (Certificate Autoenrollment in Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section outlines key scenarios that need to be considered when troubleshooting autoenrollment. It also covers how to prepare for autoenrollment failures and lists event logging messages.

**Key Issues**
The following key issues need to be considered when troubleshooting autoenrollment.

**Infrastructure Requirements**
Windows XP clients and Windows Server 2003 CAs will always request LDAP-signed communications with domain controllers as a security function. Before deploying autoenrollment or a Windows Server 2003 CA, all domain controllers running Windows 2000 should be upgraded to Service Pack 3 or greater.

**Root and Cross-Certificate Download from Active Directory**
Autoenrollment automatically downloads root certificates and cross-certificates from Active Directory whenever a change is detected in the directory or when a different domain controller is contacted. If a third-party root certificate or cross-certificate is deleted from the local machine store, autoenrollment will not download the certificates again until a change occurs in Active Directory or a new domain controller is contacted.

To manually force a new download, delete the following registry key and all subordinate keys on all affected machines.
**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Cryptography\AutoEnrollment\AEDirectoryCache**

**EFS and Autoenrollment**
EFS always attempts to enroll for the Basic EFS template by default. The EFS component driver generates an autoenrollment request that autoenrollment tries to fulfill. For customers who want to ensure that a specific template is used for EFS (such as to include key archival), the new template should supersede the Basic EFS template. The Basic EFS template should also be removed from any Enterprise CA. This will ensure that autoenrollment will not attempt enrollment for the Basic EFS template any more. For customers who wish to replace the Basic EFS template with a certificate and key that is archived through the Windows Server 2003, Enterprise Edition CA, the proper procedure is to supersede the Basic EFS template with a new version 2 certificate template.

**Smartcard Renewal**
The Smartcard Logon and Smartcard User version 1 templates may not be renewed through autoenrollment. To renew a version 1 Smartcard Logon or Smartcard User template, the proper procedure is to supersede these templates with a new version 2 template.

Autoenrollment always attempts to generate a new key when performing certificate renewal. For smart cards with limited space that do not support additional key generation, autoenrollment will attempt to reuse the key; however, additional space will still be required to install the new certificate. If no space is available on the card for these operations, the renewal through autoenrollment may fail.

**Autoenrollment and Strong Private Key Protection**
The version 2 certificate template properties on the Request Handling tab support the ability to require a user password when the private key is used by applications. This is set by selecting the Prompt the user during enrollment option and requires user input when the private key is used. It is important to never use this option for smart card certificates as smart card CSPs also do not support this capability. If this option is chosen, autoenrollment may fail.

**Removal of Certificates on Domain Join/Change Domain**
When a machine is removed from a domain or added to a new domain, all the downloaded certificates from Active Directory will be removed and refreshed if applicable. Certificates that were issued or autoenrolled from a previous forest will not be removed unless the machine is a domain controller. All client machines will automatically update certificates when the domain or machine information changes. When machines or users have certificates that are required for secure network communications, wireless communications, and so on, it may be necessary to delete the old certificates after joining a new domain or forest.

**Autoenrollment Failures**
Autoenrollment will warn the user with a warning dialog box when an autoenrollment failure occurs. This feature is only enabled when user interaction is required on the certificate template.

To enable the warning feature for an autoenrollment failure

1. Open the specified template in the **Certificate Templates** MMC snap-in.

2. Click the **Request Handling** tab.

3. Click **Prompt the user during enrollment** on the Request Handling tab of the certificate template properties.

**Re-Initialized Smart Cards**
If enrollment for a certificate is based on the existence of a smart card certificate and if the smart card has been re-initialized, the smart card Insertion dialog box will ask the user to insert a smart card matching the key container identified by the old certificate. Since the key container has been deleted, the Insertion dialog box will continue to display despite the fact that the user has removed and inserted the card. The only choice is to click Cancel and fail the enrollment.

**Enhanced Event Logging**
By default, autoenrollment logs errors/failures and successful enrollments in the Application event log on the client machine.

To enable enhanced logging of autoenrollment processes to include warning and informational messages, the following registry values must be created.

**User Autoenrollment**
**HKEY_CURRENT_USER\Software\Microsoft\Cryptography\Autoenrollment**: Create a new DWORD value named AEEventLogLevel"; set value to 0.

**Machine Autoenrollment**
**HKEY_LOCAL_MACHINE\Software\Microsoft\Cryptography\Autoenrollment**: Create a new DWORD value named "AEEventLogLevel", set value to 0.

| **Note** |
|---|
| All failures and errors are automatically logged. It is not necessary to enable the registry key to turn on failure logging. |

**Event Log Messages**
The following event log messages only appear when additional event logging is enabled.

**Success Event Log Messages**
The following are samples of successful event log messages.

| | |
|---|---|
| Event Type: | Information |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 2 |
| Date: | 2/26/2001 |
| Time: | 12:52:02 PM |
| User: | N/A |
| Computer: | COMPUTER1 |
| Description: | Automatic certificate enrollment for local system started. |

| | |
|---|---|
| Event Type: | Information |
| Event Source: | AutoEnrollment |
| Event Category: | None |

| | |
|---|---|
| Event ID: | 3 |
| Date: | 2/26/2001 |
| Time: | 12:52:10 PM |
| User: | N/A |
| Computer: | COMPUTER1 |
| Description: | Automatic certificate enrollment for local system completed. |

| | |
|---|---|
| Event Type: | Information |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 27 |
| Date: | 2/26/2001 |
| Time: | 3:26:03 PM |
| User: | CONTOSO\USER1 |
| Computer: | COMPUTER1 |
| Description: | Automatic certificate enrollment for logged on user is cancelled. |

| | |
|---|---|
| Event Type: | None |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 28 |
| Date: | 6/25/2001 |
| Time: | 7:36:16 AM |
| User: | CONTOSO\USER1 |
| Computer: | COMPUTER1 |
| Description: | Automatic certificate enrollment for CONTOSO\User1 successfully installed one AutoEnrollSmart cardEmail certificate when retrieving pending requests. User interaction was required. |

| | |
|---|---|
| Event Type: | None |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 29 |

| | |
|---|---|
| Date: | 7/9/2001 |
| Time: | 6:39:29 AM |
| User: | CONTOSO\USER1 |
| Computer: | COMPUTER1 |
| Description: | Automatic certificate enrollment for CONTOSO\USER1 reused the private key when requesting one AutoEnrollSmart cardUser certificate. |

| | |
|---|---|
| Event Type: | None |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 20 |
| Date: | 7/9/2001 |
| Time: | 6:39:29 AM |
| User: | CONTOSO\USER1 |
| Computer: | COMPUTER1 |
| Description: | Automatic certificate enrollment for CONTOSO\USER1successfully renewed one AutoEnrollSmart cardUser certificate from certification authority TestCA on Server1.contoso.com. |

| | |
|---|---|
| Event Type: | None |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 29 |
| Date: | 7/17/2001 |
| Time: | 9:37:29 AM |
| User: | CONTOSO\user1 |
| Computer: | TESTCA |
| Description: | Automatic certificate enrollment for CONTOSO\user1 reused the private key when requesting one Autoenroll Smart card User certificate. |

| **Note** |
|---|
| This event signifies the fact that the private key was reused during a certificate renewal. |

**Failed Event Log Messages**

The following are samples of failed event log messages.

| | |
|---|---|
| Event Type: | Error |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 15 |
| Date: | 7/8/2001 |
| Time: | 3:09:41 PM |
| User: | N/A |
| Computer: | TEST1 |
| Description: | Automatic certificate enrollment for CONTOSO\User1 failed to contact Active Directory (0x8007054b). The specified domain either does not exist or could not be contacted. Enrollment will not be performed. |

| **Note** |
|---|
| This error most often occurs when a user is logged on to a machine with cached credentials and is offline. Therefore, autoenrollment cannot continue and will be attempted later. |

| | |
|---|---|
| Event Type: | Error |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 15 |
| Date: | 2/24/2001 |
| Time: | 10:36:08 AM |
| User: | N/A |
| Computer: | TEST1 |
| Description: | Automatic certificate enrollment for local system failed to contact a directory server (0x80072751). A socket operation was attempted to an unreachable host. Enrollment will not be performed. |

| **Note** |
|---|
| This error most often occurs when a domain controller is not available or is not accessible by the client. Common causes include network errors, network connectivity, and so on. |

| | |
|---|---|
| Event Type: | Error |
| Event Source: | AutoEnrollment |
| Event Category: | None |

| | |
|---|---|
| Event ID: | 13 |
| Date: | 7/5/2001 |
| Time: | 9:37:44 AM |
| User: | N/A |
| Computer: | TEST1 |
| Description: | Automatic certificate enrollment for local system failed to enroll for one CONTOSO IPSEC certificate (0x800706ba). The RPC server is unavailable. |

| **Note** |
|---|
| This error typically occurs when the certification authority is not available on the network or the service is stopped. |

| | |
|---|---|
| Event Type: | Error |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 13 |
| Date: | 7/5/2001 |
| Time: | 7:41:27 AM |
| User: | N/A |
| Computer: | TEST1 |
| Description: | Automatic certificate enrollment for local system failed to enroll for one CONTOSO IPSEC certificate (0x8009400f). An attempt was made to open a certification authority database session, but there are already too many active sessions. The server may need to be configured to allow additional sessions. |

| **Note** |
|---|
| This is a rare event when the certification authority is under heavy load and cannot respond to the request in a timely manner. Autoenrollment will automatically try again at a later time. |

| | |
|---|---|
| Event Type: | Error |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 16 |
| Date: | 7/5/2001 |
| Time: | 2:53:34 AM |
| User: | N/A |

| | |
|---|---|
| Computer: | TEST1 |
| Description: | Automatic certificate enrollment for local system failed to renew one CONTOSO IPSEC certificate (0x8009400f). An attempt was made to open a Certification Authority database session, but there are already too many active sessions. The server may need to be configured to allow additional sessions. |

| **Note** |
|---|
| This is the same error as the previous one, but it involves a renewal. |

| | |
|---|---|
| Event Type: | Warning |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 7 |
| Date: | 7/24/2001 |
| Time: | 7:48:27 PM |
| User: | CONTOSO\USER1 |
| Computer: | TEST1 |
| Description: | Automatic certificate enrollment for CONTOSO\USER1 could not enroll for Key Recovery Agent certificate template due to one of the following situations. |

- Enrollment access is not allowed to this template.
- Template subject name, signature, or hardware requirements cannot be met.
- No valid certification authority can be found to issue this template.

| **Note** |
|---|
| This is an autoenrollment error that occurs when a user has a certificate and private key installed that corresponds to a given template that is now expiring. Autoenrollment attempts to automatically renew the certificate; however, the user does not have applicable permissions for this template and therefore autoenrollment fails. Autoenrollment is based on certificates in the store as well as certificate template settings. |

| | |
|---|---|
| Event Type: | Error |
| Event Source: | AutoEnrollment |
| Event Category: | None |
| Event ID: | 13 |
| Date: | 7/17/2001 |
| Time: | 9:22:10 AM |

| | |
|---|---|
| User: | CONTOSO\user1 |
| Computer: | TESTCA |
| Description: | Automatic certificate enrollment for CONTOSO\user1 failed to enroll for one Autoenroll smart card user certificate (0x80094812). The e-mail name is unavailable and cannot be added to the Subject or Subject Alternate name. |

| **Note** |
|---|
| This error occurs when the user account in Active Directory does not have a valid e-mail address on the user property page in Active Directory Users and Computers MMC snap-in. Enrollment for certificate templates in Active Directory requires an e-mail address to exist prior to enrollment. |

**Event Log Tools**

A new tool (script) is included on the Windows XP Professional and Windows Server 2003 client to query a local system for various events. This script can be used to identify autoenrollment errors on the client and perform appropriate actions. The command line help for this tool is noted and provided as follows:

```
Z:\>eventquery /?
Microsoft (R) Windows Script Host Version 5.6
Copyright © Microsoft Corporation 1996-2001. All rights reserved.
---------------------------------------------------------------------
EVENTQUERY.vbs [/S system [/U user name [/P password]]] [/FI filter]
        [/FO format] [/R range] [/NH] [/V] [/L logname | *]
Description:
  The EVENTQUERY.vbs script enables an administrator to list
  the events and event properties from one or more event logs.
Parameter List:
  /S   system     Specifies the remote system to connect to.
  /U   [domain\]user  Specifies the user context under which the
             command should execute.
  /P   password    Specifies the password for the given
             user context.
  /V           Specifies that the detailed information
             should be displayed in the output.
  /FI  filter     Specifies the types of events to
             filter in or out of the query.
  /FO  format     Specifies the format in which the output
             is to be displayed.
             Valid formats are "TABLE", "LIST", "CSV".
  /R   range      Specifies the range of events to list.
             Valid Values are:
               'N' - Lists 'N' most recent events.
              '-N' - Lists 'N' oldest events.
             'N1-N2' - Lists the events N1 to N2.
  /NH          Specifies that the "Column Header" should
             not be displayed in the output.
             Valid only for "TABLE" and "CSV" formats.
  /L   logname     Specifies the log(s) to query.
 /?          Displays this help/usage.
  Valid Filters Operators allowed  Valid Values
  ------------- ------------------ ------------
  DATETIME    eq,ne,ge,le,gt,lt  mm/dd/yy(yyyy),hh:mm:ssAM(/PM)
  TYPE      eq,ne        ERROR, INFORMATION, WARNING,
                   SUCCESSAUDIT, FAILUREAUDIT
  ID       eq,ne,ge,le,gt,lt  non-negative integer
  USER      eq,ne        string
  COMPUTER    eq,ne        string
  SOURCE     eq,ne        string
  CATEGORY    eq,ne        string
```

| **Note** |
|---|
| Filter "DATETIME" can be specified as "FromDate-ToDate" |

```
   Only "eq" operator can be used for this format.
Examples:
  EVENTQUERY.vbs
  EVENTQUERY.vbs /L system
  EVENTQUERY.vbs /S system /U user /P password /V /L *
  EVENTQUERY.vbs /R 10 /L Application /NH
  EVENTQUERY.vbs /R -10 /FO LIST /L Security
  EVENTQUERY.vbs /R 5-10 /L "DNS Server service"
  EVENTQUERY.vbs /FI "Type eq Error" /L Application
  EVENTQUERY.vbs /L Application /FI "Datetime eq 06/25/00,03:15:00AM-
06/25/00,03:15:00PM"
  EVENTQUERY.vbs /FI "Datetime gt 08/03/00,06:20:00PM"           /FI "Id gt 700" /FI
"Type eq warning" /L System
  EVENTQUERY.vbs /FI "Type eq error OR Id gt 1000 "
```

# Summary (Certificate Autoenrollment in Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Windows Server 2003, Enterprise Edition through the Certificate Services component provides user certificate autoenrollment. This allows administrators to easily deploy certificates throughout the enterprise while requiring no user interaction. User certificate autoenrollment in the Windows XP Professional and Windows Server 2003 operating systems builds on Microsofts long-established reputation for shipping robust PKI components that have a low TCO. Since PKI is an integral part of the Windows XP Professional operating system, Windows Server 2003 PKI provides some distinct advantages over third-party add-in components. These advantages include:

- No per-certificate fees or per-user PKI licenses
- Centralized user security management
- Integration with normal enterprise management tasks
- Single sign-on capabilities to networks and applications
- Managed trust capabilities
- Support for all applications through CryptoAPI

Keep in mind that almost all third-party PKIs must be purchased separately, and require per-certificate license fees and increased management tasks.

Overall, certificate autoenrollment features in Windows Server 2003 should provide organizations and enterprises with the ability to effortlessly deploy digital certificates and PKI-enabled applications with little or no additional cost to a normal IT operations budget.

# Related Links (Certificate Autoenrollment in Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

- Whats New in Security for Windows XP Professional and Windows XP Home Edition at http://www.microsoft.com/windowsxp/pro/techinfo/planning/security/whatsnew/default.asp
- Windows XP and .NET: An Overview at http://www.microsoft.com/windowsxp/pro/techinfo/planning/dotnet/default.asp
- CMC (Certificate Management Messages over CMS) at http://www.ietf.org/rfc/rfc2797.txt?number=2797
- Data Protection and Recovery in Windows XP at http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/default.mspx
- Securing Mobile Computers with Windows XP Professional at http://www.microsoft.com/windowsxp/pro/techinfo/administration/mobile/default.asp
- PKI Enhancements in Windows XP Professional and Windows Server 2003 at http://www.microsoft.com/technet/prodtechnol/winxppro/plan/default.mspx
- For the latest information about Windows XP, see the Windows XP Web site at http://www.microsoft.com/windowsxp/default.asp
- For the latest information about Windows Server 2003, see the Windows Server 2003 Web site at http://www.microsoft.com/windowsserver2003/default.mspx

# Account Passwords and Policies in Windows Server 2003

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Microsoft Windows NT Server 4.0, Windows 2000, Windows XP, and the Windows Server 2003 Family

Password and account lockout settings are designed to protect accounts and data in your organization by mitigating the threat of brute force guessing of account passwords. Settings in the Account Lockout and Password Policy nodes of the Default Domain policy settings enable account lockout and control how account lockout operates. This white paper describes how these settings affect account lockout and makes some general recommendations for configuring and troubleshooting account lockout issues.

This document describes the reasons why you should take a structured approach to setting the account and password policy features. The document also provides information about the tools and log files that you can use to troubleshoot account lockouts. After you read this document, you should be able to determine from which computer the account lockouts are being generated, as well as the program or service that is generating the lockout.

**In This White Paper**


- Account Lockout and Password Concepts
- Configuring Account Lockout Settings
- Configuring Account Lockout
- Details of Account Lockout Settings and Processes
- Maintaining and Monitoring Account Lockout
- Troubleshooting Account Lockout
- Account Lockout Tools
- Appendix One: Additional References for Account Lockout
- Appendix Two: Gathering Information to Troubleshoot Account Lockout Issues

# Account Lockout and Password Concepts

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Passwords are an important step in a security plan for your network. Users may see passwords as a nuisance; however, the security of your enterprise relies on a combination of password length, password uniqueness, and password lifespan. These three items help defend against dictionary attacks and brute force attacks. A dictionary attack occurs when a malicious user tries known words that are in the dictionary and a number of common password names to try and guess a password. A brute force attack occurs when a malicious user tries all of the possible permutations until one is successful.

Because most users prefer passwords that they can easily remember, dictionary attacks are often an effective method for a malicious user to find a password in significantly less time than they would with brute force attacks. Therefore, the strength of a password depends on how many characters are in the password, how well the password is protected from being revealed by the owner, how well the password is protected if it is intercepted by a malicious user on the network, and how difficult the password is to guess. Even good passwords that are protected by cryptography on the network and that are not subject to dictionary attacks can be discovered by brute force in a few weeks or months by a malicious user who intercepts the password on the network.

Currently, several attack methods are based on guessing weak passwords by using dictionary and brute force attacks. For a few simple ways to help prevent these attacks, see "Protecting from External Lockout Denial of Service Attacks" in this document for ports to block and registry values that you can set to help prevent such attacks.

Frequently, a malicious user will guess a number of passwords during a password-based attack. To help prevent the attacks from being successful, you can configure account lockout settings. The result of this configuration is that the associated account is temporarily disabled after a specified number of incorrect passwords are tried. This helps to prevent a successful attack by preventing the account from being used. However, a legitimate user cannot use that account until it is unlocked. This paper discusses the balance between the benefits and risks of account lockout.

**Understanding Password Complexity**

A complex password that is enforced by the operating system is one of the most effective methods that you can use to deter the opportunity for a successful attack. When you configure both an expiration time and a minimum length for a password, you decrease the time in which a successful attack could occur. For example, when you enforce password complexity with a password length of 6 and set the password to expire in 60 days, a user can choose from a permutation of:

- 26 lowercase characters
- 26 uppercase characters
- 32 special characters
- 10 numbers

This means that:

- 26 + 26 + 32 + 10 = 94 possible characters in a password
- Password length policy = 6
- $94^6$ = 689,869,781,056 unique password permutations

With a 60-day password expiration time, the malicious user would have to make 133,076 password attempts every second to attempt all of the possible passwords during that password's limited lifetime. If it takes only 50 percent of the permutations to guess the password, a malicious user would have to attempt to log on to the computer about 66,538 (133,076 * .50) times every second to discover the password before it expires.

To decrease the chances that a malicious user has to discover the password, you can use a password length of 7. When you set the minimum password length to 7, the possible password permutations exceed 64 trillion (947= 64,847,759,419,264). When you compare the calculations above that have a password length of 6 to the calculations below that have a password length of

7, you will notice that the malicious user would have to log on to the computer about 6,254,606 times for each second that the password is valid in the 60-day expiration time that you set.

The following list describes how increasing password length deters both dictionary and brute force attacks. Note that the examples that are in this list assume that you are have applied a policy that requires users to create complex passwords. When you do this, there are 94 possible characters from which the users can choose their password.

- 6 characters: $94^6$ = 689,869,781,056
- 7 characters: $94^7$ = 64,847,759,419,264
- 8 characters: $94^8$ = 6,095,689,385,410,816
- 9 characters: $94^9$ = 572,994,802,228,616,704
- 10 characters: $94^{10}$ = 53,861,511,409,489,970,176

| **Note** |
| --- |
| A few of these password possibilities are not valid. By default, users cannot choose any part of their user name for their password and they cannot use all of the same characters as a password. Because of this, these password possibilities must be deducted from the total number of possible passwords that are listed above. Because there are very few passwords that apply to these exceptions and because the number of passwords that do apply to these exceptions can vary (based on the number of letters that are in the user's logon name), this document does not account for these exceptions. |

These statistics explain how difficult it is for a malicious user to discover a password when you require the users in your network to use a complex password. Because of this, Microsoft recommends that you enforce a complex password policy that requires users to choose passwords with a specific number of characters for the security needs of your organization. The "Password Policies Settings" section in this document describes the complex password policies and settings for Microsoft® Windows NT® Server 4.0, the Windows® 2000 family, and the Windows Server 2003 family of operating systems.

Microsoft recommends that you use the account lockout feature to help deter malicious users and some types of automated attacks from discovering user passwords. The following section provides more information about how you can use the account lockout feature.

**Authentication**

Authentication is the process of validating a user name and password on a domain controller for:

- The initial logon to either a workstation or domain that uses the CTRL+ALT+DELETE secure logon sequence.
- An attempt to unlock a locked workstation by using the CTRL+ALT+DELETE secure logon sequence.
- An attempt to type a password for a password-protected screen saver.
- A user, script, program, or service that attempts to connect to a network resource by using either a mapped drive or a Universal Naming Convention (UNC) path.

| **Note** |
| --- |
| An account that is locked out may still be able to gain access to some resources if the user has a valid Kerberos ticket to the resource. The ability to access the resource ends when the Kerberos ticket expires. However, neither a user who is locked out nor a computer account can renew the ticket. Kerberos cannot grant a new ticket to the resource because the account is locked out. |

There are two primary authentication protocols used by Windows: NTLM and Kerberos. This paper assumes you are familiar with these authentication protocols and does not focus on authentication details. Instead, the focus is placed on how authentication plays a role in account lockout. For more information about authentication protocols, see online help in Windows XP and the Windows Server 2003 family.

**How Domain Controllers Verify Passwords**
To illustrate the authentication process, the following diagram describes the steps that occur when a logon attempt does not work.

**Figure 1: Process for a Failed Logon Attempt**




1. The client computer presents the user logon information to a domain controller. This includes the users account name and a cryptographic hash of their password. This information can be sent to any domain controller and is typically sent to the domain controller that is identified as the closest domain controller to the client computer.

2. When a domain controller detects that an authentication attempt did not work and a condition of STATUS_WRONG_PASSWORD, STATUS_PASSWORD_EXPIRED, STATUS_PASSWORD_MUST_CHANGE, or STATUS_ACCOUNT_LOCKED_OUT is returned, the domain controller forwards the authentication attempt to the primary domain controller (PDC) emulator operations master. Essentially, the domain controller queries the PDC to authoritatively determine if the password is current. The domain controller queries the PDC for this information because the domain controller may not have the most current password for the user but, by design, the PDC emulator operations master always has the most current password.

3. The authentication request is retried by the PDC emulator operations master to verify that the password is correct. If the PDC emulator operations master rejects the bad password, the PDC emulator operations master increments the **badPwdCount** attribute for that user object. The PDC is the authority on the user's password validity.

4. The failed logon result information is sent by the PDC emulator operations master to the authenticating domain controller.

5. The authenticating domain controller also increments its copy of the **badPwdCount** attribute for the user object.

6. The authenticating domain controller then sends a response to the client computer that notifies the domain controller that the logon attempt did not work.

As long as that user, program, or service continues to send incorrect credentials to the authenticating domain controller, logon attempts that failed because of an incorrect password continue to be forwarded to the PDC until the threshold value for incorrect logon attempts is reached (if you set it in a policy). When this occurs, the account is locked out.

For more information, see "How the Bad Password Count Is Incremented in Windows NT" in the Microsoft Knowledge Base.

**New Features in the Windows Server 2003 Family**
In the Windows Server 2003 family of operating systems, Microsoft has improved the function of the Account Lockout feature on both servers and client computers.

**Computers Running Windows Server 2003 That Act As Network Servers**
To improve the experience for users and to decrease the overall total cost of ownership, Microsoft made the following changes to the behavior of domain controllers in the Windows Server 2003 family:

- **Password history check (N-2):** Before a Windows Server 2003 operating system increments **badPwdCount**, it checks the invalid password against the password history. If the password is the same as one of the last two entries that are in the password history, **badPwdCount** is not incremented for both NTLM and the Kerberos protocol. This change to domain controllers should reduce the number of lockouts that occur because of user error.

- **Single user object on demand replication:** See the "Urgent Replication" section in this document for more information.

- **Optimized replication frequency:** The default frequency for replication between sites is to replicate every 15 minutes with a 3-second offset to stagger the replication interval. This optimization improves the replication of a password change in a site because it decreases the chances that the domain controller would have to contact the PDC operations master.

**Computers Running Windows Server 2003 Family Acting As Network Clients**

Microsoft has added the following features in the Windows Server 2003 family to gather the process ID that is using the credentials that fail authentication:

- **Auditing logon changes:** There are entries for all logon and logoff events (528 and 540, as well as 529 through 539).
- **Auditing of processes encountering authentication failures:** New information is added to the Security event log when authentication failures occur:
  - Caller User Name
  - Caller Domain
  - Caller Logon ID
  - Caller Process ID

| **Note** |
| --- |
| To use the process ID, turn on success auditing for **Audit process tracking** events so that you can obtain the process identifier (PID) for the associated Event 592. If you do not do this, the PID is not useful after the process stops. To view audit process tracking, in the Group Policy Microsoft Management Console (MMC), in the console tree, double-click **Computer Configuration**, double-click **Windows Settings**, double-click **Security Settings**, double-click **Local Policies**, and then double-click **Audit Policy**. |

Microsoft has added the following administrative enhancements to provide more account lockout information than the information that is available in the default configuration of the Windows Server 2003 family:

- **AcctInfo.dll:** The AcctInfo.dll file is a property page extension for user objects in the Active Directory Users and Computers MMC that provides detailed information about user password attributes. An administrator can use the AcctInfo.dll file to reset user account passwords on a domain controller that is in the user's Active Directory site.
- **LockoutStatus.exe:** The LockoutStatus.exe tool displays bad password count and time information from all of the domain controllers that are in a domain. You can run this tool as either a stand-alone tool or as an extension to the AcctInfo.dll file when you place it in the Systemroot\System32 folder on your computer.

More information about both AcctInfo.dll and LockoutStatus.exe is available in "Account Lockout Tools" in this document.

# Configuring Account Lockout Settings

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

To enable the Account Lockout policy settings for both domain and local users, configure the settings that are described in this section.

**Recommended Service Packs and Hotfixes**
Security issues in Windows operating systems are discovered and fixed often. These fixes often have an impact on account lockout and password policy features, as well as their dependent components. Therefore, you should apply the latest service packs and hotfixes to all of the domain controllers, servers, and clients to ensure that the account security settings that you want are applied and to ensure that the domain controllers and operating systems are up-to-date. Note that service packs resolve groups of issues and hotfixes resolve a specific issue.

You should have an ongoing strategy to keep your computers updated and protected against viruses, trojans, and so on, that may use vulnerabilities that are already fixed. For more information about how to create a software update strategy, refer to Help and Support Center in the Windows Server 2003 family and Windows XP, or refer to Help in Windows 2000.

Because these issues are discovered and fixed on an ongoing basis, they are not listed in this document. For more information, see "Service Packs and Hotfixes Available to Resolve Account Lockout Issues" in the Microsoft Knowledge Base.

# Configuring Account Lockout

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The account lockout policy settings are designed to help prevent a brute force attack on user passwords. This section describes where you can configure the setting, as well as some things that you should consider before you use the settings.

**Configuring Account Lockout for Domain Users**
For domain users, set the appropriate values by configuring the default domain policy in the console tree. To configure the default domain policy, open the Group Policy MMC, double-click **Computer Configuration**, double-click **Windows Settings**, double-click **Security Settings**, double-click **Account Policies**, and then double-click **Account Lockout Policy**. For more detailed steps, see "Account Lockout Settings" in this document.

**Configuring Account Lockout for Local Users**
For a stand-alone workstation, set the appropriate registry values by configuring the local policy:

1. Click **Start**, click **Run**, type **gpedit.msc**, and then press ENTER.
2. In the Group Policy Object Editor MMC, double-click **Computer Configuration**, double-click **Windows Settings**, double-click **Security Settings**, double-click **Account Policies**, and then double-click **Account Lockout Policy**.
3. In the details pane, right-click the policy setting that you want, and then click **Properties**.
4. If you are defining this policy setting for the first time, click **Define this policy setting**.
5. Select the options that you want, and then click **OK**.

Microsoft recommends that you do not exempt the privileged accounts from password policies. Your privileged accounts should have complex passwords, an expiration period, and the passwords should be a minimum of fifteen characters in length. Microsoft also recommends that you also protect the local accounts (non-domain clients) by using a local password policy for all users. For all workstations in a domain, set a domain-level Group Policy object (GPO) and filter it to apply to the domain member computer.

**Choosing Account Lockout Settings for Your Deployment**
This section describes the ramifications of changing the various settings and a way to estimate the difficulty of using the brute force method of password guessing with a certain configuration. Like other settings that are associated with passwords, choosing settings for account lockout involves balancing the benefits and drawbacks between security, usability, and cost. The primary consideration is how much risk is acceptable when you configure the password policy of the domain. For example, consider the following two domain configurations:

- User accounts that are in domain A have a minimum password length of 3 characters, no password complexity requirements, and passwords never expire.
- User accounts that are in domain B have a minimum password length of 6 characters, password complexity, and passwords in the domain expire in 42 days.

There is less risk of a malicious user guessing the password of a domain B user; for the same risk tolerance, domain B can have a less stringent account lockout policy. Whether you set the **LockoutDuration** registry value to 0 or not also has an important on the setting that is permitted for both the **ObservationWindow** and **LockoutThreshold** registry values.

As an example, assume that the administrator resets the password when the account is locked with **LockoutDuration** registry value of 0. With the **LockoutDuration** registry value set to 0 and the **LockoutThreshold** registry value set to 20, the malicious user has 20 guesses to use against that password. If the lockout duration is 30 minutes, the malicious user has 20 guesses every 30 minutes against that password until it is changed. This is a very significant difference in the total number of guesses that are available to the malicious user.

In comparison, if the administrator sets the maximum password age to 42 days, the first malicious user has only 20 guesses against any given password, while the second malicious user has 40,320 guesses (20 tries for ever lockout, multiplied by 48

lockouts every day, multiplied by 42 days before the user changes the password). With the default password settings, there are approximately 1012 possible passwords. This means that the malicious user has approximately a .000004 percent (%) chance of guessing the password. With an optimized guessing scheme, this percentage would most likely be a larger percentage.

This example demonstrates that setting the **LockoutDuration** registry value to 0 allows the **LockoutThreshold** registry value to be a significantly higher number for equivalent risk tolerance. If you increase the **LockoutThreshold** registry value, you help to reduce the behavior of a user accidentally locking themselves out of their computer.

When you choose the setting for account lockout, it is important that you consider the inherent denial of service that is associated with a locked out account. Ideally, the only accounts that are locked out are the accounts that are being attacked. However, the computer cannot determine the difference between a user typing an incorrect password or an automated task that is using an incorrect password. As a result, from the users perspective, users who are trying to perform their daily tasks may be suddenly unable to perform their work, which incurs both the cost of support to the domain and the cost of lost work. The lower the value that you set for the **LockoutThreshold** registry value, the more likely this behavior is to occur. The length of the **ObservationWindow** registry value has much less effect on this behavior than the **LockoutThreshold** registry value.

Microsoft recommends that you regularly review the Security event logs of all computers so that you are aware of any patterns that might show an attack or user error. The values that are necessary to identify specific malicious users and targets can be obtained only after you implement the auditing policies that are mentioned in the "Appendix Two: Gathering Information to Troubleshoot Account Lockout Issues" section in this document. Microsoft offers an event log monitoring solution, Microsoft Operations Manager (MOM), that you can script with responses. This tool also has many other built-in actions that you can use. For more information about MOM, see the MOM Web site.

| **Note** |
| --- |
| Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here. |

#### LockoutThreshold vs. ObservationWindow

In general, the **LockoutThreshold** value has more of an effect on how the computer behaves than the **ObservationWindow** value. Most logon attempts that do not work occur during a very short period of time. Because of this, the period of time is inside of the **ObservationWindow** time. Users rarely type a bad password many times in a row, so the **LockoutThreshold** value is rarely exceeded. The exception to this is the environment where the **LockoutThreshold** registry value is set so low that a user could accidentally mistype their password often enough to lock themselves out (for example, if you set the **LockoutThreshold** value to 2). Malicious password attacks are almost always automated. A user typically locks themselves out of their computer when they type a bad password once and try to type that same bad password over and over again.

#### Recommended Account Lockout Settings

The security configuration for an organization is determined by the level of protection that is required in the organization's environment. In some low-security scenarios (such as in a small office where no sensitive information is stored in the system), a simple password policy may be sufficient to protect the resources. However, in a high-security environment (such as in a banking system), much stronger security protection is desired. You should use account lockout and strong password policies in these environments. In all examples, the stronger the security that is implemented, the higher the cost that is associated with maintaining that security.

The following table provides recommended account lockout settings for many different security configurations.

**Table 1 Recommended Account Lockout Settings**

| Security category | Account lockout settings | | | Cost |
| --- | --- | --- | --- | --- |
| | Threshold | Observation window | Lockout duration | |
| **Low** | N/A | N/A | N/A | **Low** |
| **Medium** | 10 | 30 | 30 | **Medium** |
| **High** | 10 | 30 | Infinite/0 | **High** |

| **Note** |
| --- |
| "Cost" includes downtime cost for the user whose account is locked out, as well as support cost for servicing the locked out account. |

#### Recommended Password Policy Settings

The table below provides recommended password policy settings for various security configurations.

**Table 2 Recommended Password Policy Settings**

| Security Category | Password policy settings | | | | | Cost |
|---|---|---|---|---|---|---|
| | Password history | Maximum password age | Minimum password age | Minimum password length | Complexity | |
| **Low** | 3 | 42 | 0 | 0 | disabled | **Low** |
| **Medium** | 24 | 42 | 1 | 7 | enabled | **Medium** |
| **High** | 24 | 42 | 1 | 8 | enabled | **High** |

**General Recommendations for Account Lockout and Password Policy Settings**
In addition to the specific account lockout and password policy settings in the previous tables, there are some other configuration changes that may help you achieve the level of security that you want. These include:

- When you enable account lockout, set the **ForceUnlockLogon** registry value to 1. This setting requires that Windows re-authenticates the user with a domain controller when that user unlocks a computer. This helps to ensure that a user cannot use a previously-cached password to unlock their computer after the account is locked out.

- False lockouts can occur if you set the **LockoutThreshold** registry value to a value that is lower than the default value of 10. This is because users and programs can retry bad passwords frequently enough to lock out the user account. This adds to administrative costs.

- After you unlock an account that is locked out, verify that the **LockoutDuration** value is set. You should do this because the value may have changed during the account unlock process.

- Carefully consider setting the **LockoutDuration** registry value to 0. When you apply this setting, you may incur additional administrative labor by requiring administrators to manually unlock a locked out user account. Although this does increase security, the increased labor drawback may outweigh the security benefit.

- Define account lockout and password policies once in every domain. Ensure that these policies are defined only in the default domain policy. This helps to avoid conflicting and unexpected policy settings.

- Unlock an account from a computer that is in the same Active Directory site as the account. By unlocking the account in the local site, urgent replication occurs in that site which triggers immediate replication of the change. Because of this, the user account should be able to regain access to resources faster than waiting for replication to occur. Note that the AcctInfo.dll tool helps to identify an appropriate domain controller and unlock the account. For more information about AcctInfo.dll, see the "Account Lockout Tools" section in this document.

**Protecting from External Account Lockout Denial of Service Attacks**
It is possible for a malicious user to launch a denial-of-service attack against your enterprise from outside of your network. Because most networks are interconnected, this can be a difficult attack to mitigate. The following techniques technologies are common techniques and technologies that you can use to help mitigate or prevent such attacks:

- **Require complex passwords:** All accounts should have a complex password. All administrator accounts (local and domain) should have a long, complex password and you should change the password regularly.

- **Rename the administrator account:** Because the administrator account cannot be locked out, it is recommended that you rename the account. Although this does not mitigate all of the attacks against the administrator account, it does help mitigate these attacks most of the time. For more information, see "HOW TO: Rename the Administrator and Guest Account in Windows 2000" on the Microsoft Knowledge Base.

- **Protect your environment with firewalls:** To avoid an account lockout denial of service attack, block the TCP and UDP ports 135 through 139 and port 445 on your routers and firewalls. When you do this, you prevent logon attempts that occur outside of your network.

- **Prevent anonymous access:** Set the **RestrictAnonymous** value to 2 on all computers that are exposed to the internet and to the entire domain if all of the computers are running versions of Windows 2000 or later. This stops malicious users from making anonymous connections to resources and may help defeat some types of attacks. Note that some operating systems have limited support for computers that have this setting. Some programs may also have issues with this setting if the programs use an anonymous connection to gain access to resources. For more information, see "How to Use the RestrictAnonymous Registry Value in Windows 2000" on the Microsoft Knowledge Base.

- **Protect site-to-site traffic by using a VPN tunnel:** If communication between domain members in two sites is required, use a site-to-site VPN tunnel to securely connect site networks together. Do not open all NetBIOS ports on the firewall. You can use the Windows 2000 Server Routing and Remote Access service to create site-to-site VPN tunnels. If no VPN devices are available, you should configure the edge firewall or router filters to limit the traffic that is permitted to flow between the Internet Protocol (IP) address ranges that are used by each site. If sites need to use Active Directory replication only across the Internet, then use Internet Protocol security (IPSec) transport mode through the firewalls to secure all traffic between Active Directory servers.

- **Protecting authentication and NetBIOS ports from Internet attack:** On either the firewall or the router that connects your internal network to the Internet, block access to TCP and UDP ports 135 through 139 and port 445. If no edge filtering device is available, you can use IPSec filters to block these ports. To do this, use the configuration that is described in "How to Block Specific Network Protocols and Ports by Using IPSec" on the Microsoft Knowledge Base.

- In the same IPSec policy, you must create an additional rule that adds filters to permit traffic to these ports when the source address is in a subnet that is used by the internal network. To do this, use the configuration that is described in "How to Block Specific Network Protocols and Ports by Using IPSec" on the Microsoft Knowledge Base.

- **Protecting authentication and NetBIOS ports from internal attack:** If you must protect access to both authentication and NetBIOS ports from internal malicious users, you can restrict the computers that are permitted to gain access to these ports to only domain member computers by using the feature in IPSec that allows you to negotiate security. By allowing only trusted computers (domain member computers) to gain access to both authentication and NetBIOS ports, you reduce the number of computers that can perform the attack. This extra protection provides a defense against any breaches in your security perimeter and against malicious users who can connect to the internal network. For information about how to create a custom IPSec policy to use Kerberos authentication when negotiating IPSec security for access to TCP and UDP ports 135 through 139 and port 445 see the "Step-by-Step Guide to Internet Protocol Security (IPSec)" on the Microsoft Web site.

- **Update the server:** Keep all of your servers up-to-date with current versions of antivirus software, firewall software, and Windows security patches. This helps prevent trojan horse programs and viruses from attacking your resources if the malicious user can launch an attack from your internal network instead of from the Internet. These updates are an important part of an in-depth and defensive security strategy.

# Details of Account Lockout Settings and Processes

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

With the account lockout feature enabled, access to the account is denied when the number of logon attempts that did not work exceeds the **LockoutThreshold** registry value (the account lockout threshold) in a specified amount of time. A locked-out account cannot be used until it is reset by an administrator or until the lockout duration for that account expires.

Account Lockout is disabled on default installations of Windows NT Server 4.0, Windows 2000, and Windows Server 2003 domains. Account lockout operation is enabled after the domain administrator enables settings in the default domain policy. The policy settings remain enabled when you upgrade domain controllers to a later version of an operating system.

Although the Group Policy Object Editor appears to support account lockout and password policy in each organizational unit, these settings actually occur across the domain; you must define the settings on the root organizational unit for the domain. Microsoft recommends that you define account lockout and password policies in only one Group Policy object (GPO) for every domain (in the Default Domain policy settings).

**Password Policy Settings**
The first step that you should use to secure your network is to enforce password policy settings. When you implement a secure password policy, you may not need to implement the account lockout feature.

**Password Complexity**
Passwords, by default, can include any combination of characters; passwords can also be blank. Microsoft recommends that you require the use of complex passwords to help ensure that passwords provide the best security possible. These complex passwords are much more resistant to attack than blank or simple passwords.

To enforce password complexity in your organization, you can enable the **Password must meet complexity requirements** security setting. The complexity requirements enforced by this setting are stored in Passfilt.dll. In Windows 2000 operating systems and later, Passfilt.dll is built into the operating system. In Windows NT Server 4.0, you must add the Passfilt.dll file to the operating system to achieve the same results. Passfilt.dll is included in Windows NT Server 4.0 Service Pack 2 and in later service packs.

By default, complex passwords enforced by Passfilt.dll have the following properties:

- Do not contain all or part of the user's account name.
- Contain characters from three of the following four categories:
  - English uppercase characters (A through Z).
  - English lowercase characters (a through z).
  - Base-10 digits (0 through 9).
  - Non-alphanumeric (for example, !, $, #, %). extended ASCII, symbolic, or linguistic characters.

| **Note** |
| --- |
| Depending on your environment, using extended ASCII, symbolic, or linguistic characters in passwords can have unexpected results. It is highly recommended that you test these characters before using them in production. |

When implementing this policy, it is recommended to inform your users of the change in policy so that a smooth transition can take place from a simple password to a complex password. Otherwise, users may be confused by the new password criteria and circumvent security to avoid the difficulty.

You can create and register your own custom password filter if you want to modify the complexity requirements enforced in the security setting.

**Password History**
You can use the Password History setting to prevent users from repeatedly using the same password. When you use the

password history feature, a user is prevented from using passwords that they used in the past, up to the number of passwords that you specify. You can configure Windows to retain between 0 and 24 passwords by using the Password History feature. Microsoft recommends that you set the password history to the maximum value to help ensure the least amount of password reuse by users.

In the Windows 2000 Server family and later, the location of the password history is in the Default Domain policy settings at Computer Configuration\Windows Settings\Security Settings\Account Policies\Enforce Password History.

Valid non-zero values are between 1 and 24. The default value is 24 for domain controllers running a member of the Windows Server 2003 family, 3 for domain controllers running a member of the Windows 2000 family, and 0 for all other Windows operating systems.

**Minimum Password Length**
You can use the **Minimum Password Length** setting to decrease the chances that a password can be discovered by a malicious user. For more information about the **Minimum Password Length** setting, see "Understanding Password Complexity" in this document.

In versions of Windows 2000 operating systems and later, you can change the **Minimum Password Length** setting in the Group Policy MMC, in the **Default Domain** policy settings at Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\Minimum Password Length.

An administrator can set the value between 0 and 14 characters. Each additional character increases the total possible password permutations. However, if you set the value to 0, blank passwords are not permitted.

Valid non-zero values are between 1 and 14, with a default value of zero.

**Maximum Password Age**
You can use the **Maximum Password Age** setting to limit the time for which a given password is valid. This decreases the odds of being able to crack a password. For more information, see the example in the Passwords section in this document.

In the Windows 2000 family and later, the **Maximum Password Age** setting is located in the Group Policy MMC, in the **Default Domain** policy at Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\Maximum Password Age.

This setting determines the period of time (in days) that a user can use their password before the computer requires the user to change it. You can set passwords to expire in between 1 and 999 days, or you can specify that passwords never expire by setting the number of days to 0.

Valid non-zero values are between 1 and 999, with a default value of 42.

**Minimum Password Age**
You can use the **Minimum Password Age** setting to preventing users from repeatedly changing passwords until the user is able to use their original password, if you enforce the **Password History** setting.

When you use the **Minimum Password Age** setting, you prevent the circumvention of password expiration and help to assure unique passwords.

The **Minimum Password Age** setting determines the period of time (in days) that a password must be used before the user can change it. You can set the value to between 1 and 999 days, or allow immediate changes by setting the number of days to 0.

Configure the **Minimum Password Age** setting to be a number that is larger than 0 if you want the **Enforce Password History** setting to be effective. If you do not set a minimum password age, users can repeatedly cycle through passwords until they are able to use an old favorite password. This could allow users to circumvent established password policy.

The **Minimum Password Age** setting is located in the Group Policy MMC, in Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy.

Valid non-zero values are between 1 and 999, with a default value of one for domain controllers and zero for other computers.

**Account Lockout Settings**
You can set the Account Lockout settings in the Active Directory Users and Computers MMC by using the procedure in this section.

| **Note** |
| --- |
| The value that you set for **LockoutDuration** cannot be a value that is less than **OberservationWindow**. |

1. Click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

2. In the console tree, right-click the domain on which you want to set a Group Policy object.

3. Click **Properties**, and then click the **Group Policy** tab.

4. In **Group Policy Object Links**, click **Default Domain Policy** or create and name your Group Policy object, and then click **Edit**.

5. In the console tree, double-click **Computer Configuration**, double-click **Windows Settings**, double-click **Security Settings**, double-click **Account Policies**, and then click **Account Lockout Policy**.

6. In the details pane, right-click the policy setting that you want, and then click **Properties**.

7. If you are defining this policy setting for the first time, click **Define this policy setting**.

8. Click the options that you want, and then click **OK**.

**ObservationWindow**

The **ObservationWindow** setting (also known in Group Policy as the **Reset account lockout counter after** setting) is the number of minutes after which an accounts **badPwdCount** registry value is reset. You can use the **ObservationWindow** setting to help mitigate lockout issues that are initiated by users. When you enable this setting, the bad password attempt is removed from the server after a period of time.

Valid non-zero values are between 1 and 99999, with a default value of 30.

**LockoutDuration**

The **LockoutDuration** setting (also known in Group Policy as the **Account lockout duration** setting) is the amount of time, in minutes, that account lockout is enforced on an account that has exceeded the **LockoutDuration** registry value, measured from the time of lockout. If you set the **LockoutDuration** registry value to 0, the account is permanently locked out until either an administrator or a user who has a delegated account resets the account. If the administrator or a delegated user account does not unlock the account, the operating system unlocks the account after the number of minutes that you set in the **LockoutDuration** registry value. Non-zero values for the **LockoutDuration** registry value reduce the administrative overhead of unlocking accounts by having them unlocked automatically; however, non-zero values do not provide the added security of user validation before the account is restored.

Valid non-zero values are between 1 and 99999, with a default value of 30.

**LockoutThreshold**

The **LockoutThreshold** setting (also known in Group Policy as the **Account lockout threshold** setting) is the number of times that the user, computer, service, or program can send a bad password during logon authentication before the account is locked out. Account lockout occurs when the **badPwdCount** registry value is equal to or exceeds the **LockoutThreshold** value. You can adjust the **LockoutThreshold** value to prevent both brute force and dictionary attacks, but you can set the value too low to capture user error and other non-attack errors. Administrators often set this value too low (3 through 5), which causes a large number of account lockouts because of user error, program caching by service accounts, or issues with networking clients. If you set the **LockoutThreshold** value to 0, no account lockouts occur on the domain.

Valid non-zero values are between 1 and 999, with a default value of zero.

**Account Lockout Values**

Account lockout registry values are described in this section. These values store the information that you need to track account lockout information.

| **Note** |
| --- |
| These values are maintained by the operating system, so you should not manually modify them. |

**badPasswordTime**

The **badPasswordTime** value stores the last time that the user, computer, or service account submitted a password that did

not match the password on the authenticating domain controller This property is stored locally on each domain controller that is in the domain. A value of 0 means that the last incorrect password time is unknown. For an accurate value for the user's last incorrect password time in the domain, you must query each domain controller that is in the domain; the largest one is the accurate value. For more information, see the "LockoutStatus.exe" section in this document.

| **Note** |
| --- |
| The **badPasswordTime** registry value is not replicated between domain controllers. This attribute, however, is reported to the PDC operations master. |

**badPwdCount**

The **badPwdCount** value stores the number of times that the user, computer, or service account tried to log on to the account by using an incorrect password. This value is maintained separately on each domain controller in the domain, except for the PDC operations master of the accounts domain that maintains the total number of incorrect password attempts. A value of 0 indicates that the value is unknown. For an accurate total of the user's incorrect password attempts in the domain, you must query each domain controller and use the sum of the values. For more information, see the "LockoutStatus.exe" section in this document.

| **Note** |
| --- |
| The **badPwdCount** registry value is not replicated between domain controllers. This registry value, however, is reported to the PDC operations master. |

**ntPwdHistory**

The **ntPwdHistory** registry value contains the password history for the user in Windows NT Server 4.0 one-way function (OWF). Both Windows 2000 and the Windows Server 2003 family use the Windows NT Server 4.0 OWF. This property is used by only the operating system. Note that you cannot obtain the password from the password in OWF form.

**Other Settings that Affect Account Lockout**

This section describes another setting that affect account lockout behavior. While the setting is focused on authentication, it is closely tied with account lockout policy.

You can set the authentication settings in the Active Directory Users and Computers MMC by using the procedure in this section.

1. Click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

2. In the console tree, right-click the domain on which you want to set a Group Policy object.

3. Click **Properties**, and then click the **Group Policy** tab.

4. In **Group Policy Object Links**, click **Default Domain Policy** or create and name your Group Policy object, and then click **Edit**.

5. In the console tree, double-click **Computer Configuration**, double-click **Windows Settings**, double-click **Security Settings**, double-click **Local Policies**, and then click **Security Options**.

6. In the details pane, right-click the policy setting that you want, and then click **Properties**.

7. If you are defining this policy setting for the first time, click **Define this policy setting**.

8. Click the options that you want, and then click **OK**.

**ForceUnlockLogon**

The **ForceUnlockLogon** setting (also known as Interactive logon: Require Domain Controller authentication to unlock workstation controls the behavior of a computer running Windows 2000, Windows XP or the Windows Server 2003 family when the computer is unlocked by a user.

- With a value of 1 (or **Enabled**, in Group Policy), unlocking the computer also performs a synchronous logon to the

domain to verify user authenticity. This option is slower than allowing cached authentication because it requires network-based authentication.

- With a value of 0 (or **Disabled**, in Group Policy), cached information is used to verify the users identity. When the verification is successful, the user is logged on. Windows then performs an asynchronous logon to the domain in the background. This means that the user can still unlock a computer when the account is unlocked.

- Valid values are 0 and 1, with a default value of 0.

For additional information about unlocking a workstation, see the following articles:

- "Information About Unlocking a Workstation" in the Microsoft Knowledge Base.

- "Screensaver Password Works Even If Account Is Locked Out" in the Microsoft Knowledge Base.

**Replication and Account Lockout**
Account lockout relies on the replication of lockout information between domain controllers to ensure that all domain controllers are notified of an accounts status. In addition, password changes must be communicated to all domain controllers to ensure that a user's new password is not considered incorrect. This data replication is accomplished by the various replication features of Active Directory and is also discussed in this section.

**Immediate Replication**
When you change a password, it is sent over Netlogon's secure channel to the PDC operations master. Specifically, the domain controller makes a remote procedure call (RPC) to the PDC operations master that includes the user name and new password information. The PDC operations master then locally stores this value.

Immediate replication between Windows 2000 domain controllers is caused by the following events:

- Lockout of an account

- Modification of a Local Security Authority (LSA) secret

- State changes of the Relative ID (RID) Manager

**Urgent Replication**
Active Directory replication occurs between domain controllers when directory data is updated on one domain controller and that update is replicated to all other domain controllers. When a change in directory data occurs, the source domain controller sends out a notice that its directory store now contains updated data. The domain controller's replication partners then send a request to the source domain controller to receive those updates. Typically, the source domain controller sends out a change notification after a delay. This delay is governed by a notification delay. (The Windows 2000 default notification delay is 5 minutes; the Windows Server 2003 default notification delay is 15 minutes.) However, any delay in replication can result in a security risk for certain types of changes. Urgent replication ensures that critical directory changes are immediately replicated, including account lockouts, changes in the account lockout policy, changes in the domain password policy, and changes to the password on a domain controller account. With urgent replication, an update notification is sent out immediately, regardless of the notification delay. This design allows other domain controllers to immediately request and receive the critical updates. Note, however, that the only difference between urgent replication and typical replication is the lack of a delay before the transmission of the change notification. If this does not occur, urgent replication is identical to standard replication. When replication partners request and subsequently receive the urgent changes, they receive, in addition, all pending directory updates from the source domain controller, and not only the urgent updates.

When either an administrator or a delegated user unlocks an account, manually sets password expiration on a user account by clicking **User Must Change Password At Next Logon**, or resets the password on an account, the modified attributes are immediately replicated to the PDC emulator operations master, and then they are urgently replicated to other domain controllers that are in the same site as the PDC emulator. By default, urgent replication does not occur across site boundaries. Because of this, administrators should make manual password changes and account resets on a domain controller that is in that user's site.

The following events are not urgent replications in Windows 2000 domains:

- Changing the account lockout policy

- Changing the domain password policy
- Changing the password on a computer account
- Domain trust passwords

For additional information about urgent and immediate replication, see "Urgent Replication Triggers in Windows 2000" in the Microsoft Knowledge Base.

**Single User Object On Demand Replication**
In the Windows 2000 family, when an administrator resets and immediately expires a user's password on a domain controller in site A (so that the user is given a new password but forced to change it when the user first logs on), the logon may still succeed when the user logs on with that new password in site B. This occurs because the domain controller chains to the PDC operations master during authentication. However, the users password change may not replicate correctly because domain controllers in site B do not yet have the reset password. This occurs because there is replication latency between sites.

An update is available for Windows 2000 that changes this behavior. For more information to help change this behavior by implementing an "on-demand" replication scheme, see "You Cannot Change Your Password After an Administrator Resets It" on the Microsoft Knowledge Base. The updated replication scheme allows the domain controller to contact the PDC operations master to request an update of the user object that failed authentication because of an incorrect password. This helps to ensure that the authenticating domain controller receives the most current user account information as quickly as possible.

**Mixed Environments with Windows NT Server 4.0 and Active Directory Domain Controllers**
If servers that are running Windows NT Server 4.0 and earlier are in the domain, account lockout is not a dependable security feature.

For example, a Windows NT Server 4.0, Enterprise Edition, backup domain controller (BDC) may authenticate a user, even though the account is marked as locked out on a domain controller that is running Windows NT Server 4.0 and earlier. Also, Windows NT Server 4.0, Enterprise Edition, BDCs cannot unlock an account. The server that is running Windows NT Server 4.0, Enterprise Edition, can increment the bad password count when the user logs in with an incorrect password. That server can then report the increment to the other domain controller. However, the Windows NT Server 4.0, Enterprise Edition, BDC does not send this information to the domain controller that is running Windows NT Server 4.0 and earlier if the user logs on with the correct password. Because of this, the bad password count is not reset after the successful logon attempt.

The account lockout feature of Microsoft LAN Manager is not compatible with the account lockout feature on computers that are running Windows NT Server 4.0 and earlier. The domain controller that is running Windows NT Server 4.0 and earlier does not replicate any account lockout information to a LAN Manager BDC. If the account is marked as locked out on the Windows NT Server 4.0 and earlier domain controller, the LAN Manager BDC may still validate the user. The LAN Manager BDC displays the account lockout policy as set to "Never," even in a domain running Windows NT Server 4.0 and earlier where account lockout is enabled.

For these reasons, you should consider ensuring that all domain controllers in your network are running Windows 2000 or the Windows Server 2003 family. This is the only way to ensure that account lockout is enforced consistently across your network.

# Maintaining and Monitoring Account Lockout

Updated: January 5, 2010

Applies To: Windows Server 2003 with SP1

After you configure the account lockout options that you want, set up the computers so that you can capture more data about the accounts that are being locked out. This section describes how to enable auditing, Netlogon logging, and Kerberos logging, as well as which computers to retrieve the logs from. After you configure the logging and capture the appropriate data, this section will show you how to analyze the information so that you can ensure account lockout settings are working and identify attacks.

**Enable Auditing at the Domain Level**
The following sections describe how to enable auditing at the domain level for different operating systems.

To effectively troubleshoot account lockout, enable auditing at the domain level for the following events:

- Account Logon Events – Failure
- Account Management – Success
- Logon Events – Failure

**Windows 2000 and Windows Server 2003 Domains**
The Audit Policy settings are located in the Default Domain policy settings. To view the Auditing policy settings, in the Group Policy MMC, double-click **Computer Configuration**, double-click **Windows Settings**, double-click **Security Settings**, double-click **Local Policies**, and then double-click **Audit Policy**. Enable auditing for the event types listed in the previous section.

**Windows NT Server 4.0 Domain**
Open User Manager, click **Policies**, click **Auditing**, enable **Logon** and **Logoff** auditing for failure events, and then enable User and Group Management auditing for success events.

**Settings for Event Logs**
For troubleshooting purposes, change some of the settings for the Security event logs:

- Set the maximum security log size to 10,000 KB or more. This helps to ensure that important events are not overwritten when the log file becomes large in size.
- Set the event retention method to **Overwrite events as needed** to ensure that the computer is not shut down because there are too many Security event log entries, even when the log file becomes large in size.

For information about how to change the size and retention method of the Security event log, see the Help documentation for the operating system with which you are working.

Because the events can occur on both the client and the server, you can use the following tools to help you gather the information in a single location.

- Use the EventCombMT.exe tool, a multithreaded tool, to gather specific events from event logs from several different computers to one central location and then search those event logs for specific data of interest. Some specific search categories are built into the tool, such as account lockouts, which is already configured to include events 529, 644, 675, 676, and 681. For more information, see the Help file that is included with the tool.
- Use the Eventlog.pl tool to help you manage event logs in Windows 2000. You can use this tool to change the properties of event logs, back up event logs, export event lists to text file, clear event logs, and query the properties of the event logs. For more information, see "HOW TO: Use the Event Log Management Script Tool (Eventlog.pl) to Manage Event Logs in Windows 2000" in the Microsoft Knowledge Base.

**Netlogon Logging**
You can use Netlogon logging to capture Netlogon and NTLM events. It is recommended that you configure Netlogon logging

in a Windows 2000 domain that has Windows 2000 clients.

You must configure Netlogon logging on the primary domain controller (PDC) and on any other domain controllers that are involved in user authentication. To determine the authenticating domain controller, at a command prompt, type **set l** or use the LockoutStatus.exe tool. For more information about the LockoutStatus.exe tool, see the "Account Lockout Tools" section in this document. For enterprises that have less than 10 domain controllers, you should enable Netlogon logging on all domain controllers for each domain.

**Enabling Netlogon Logging on Computers Running Windows 2000 Server**
To enable Netlogon logging on computers that are running Windows 2000 Server, at a command prompt, type **nltest /dbflag:2080ffff**. The log file is created in *Systemroot*\Debug\Netlogon.log. If the log file is not in that location, stop and restart the Netlogon service on that computer. To do this, at a command prompt, type **net stop netlogon & net start netlogon**. For more information, see "Enabling Debug Logging for the Netlogon Service" on the Microsoft Knowledge Base.

If free disk space is low, make sure there is enough space to allow the 40 megabytes (MB) maximum space for the logging. You should also consider the disk space that Netlogon logging uses. When Netlogon.log reaches 20 MB in size, it is renamed to Netlogon.bak and a new Netlogon.log is created with the latest Netlogon data. After that Netlogon.log reaches 20 MB in size, Netlogon.bak is truncated, and the current Netlogon.log file is renamed to Netlogon.bak. Because of this process, the total disk space that is used by Netlogon logging is never more than 40 MB.

| **Note** |
| --- |
| Performance may be slightly degraded by the logging process. Therefore, you should disable Netlogon logging after you have captured the events that you want in the log file. To disable Netlogon logging, at a command prompt, type **nltest /dbflag:0**, press ENTER, type **net start netlogon** and then press ENTER. |

**Kerberos Logging**
If account lockouts involve Kerberos clients that are running a member of the Windows 2000 family or later, you can enable Kerberos logging on those client computers. You would typically perform this step after you have determined that there is an authentication issue that is related to Kerberos.

**To enable Kerberos event logging on a computer**

1. Click **Start**, click **Run**, type **regedit**, and then press ENTER.

2. Add the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters** registry value to the registry key:

   - **Registry value: LogLevel**

   - **Value type: REG_DWORD**

   - **Value data: 0x1**

   If the **Parameters** registry key does not exist, create it.

3. Close Registry Editor and restart the computer.

| **Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

| **Note** |
| --- |
| Performance may be degraded by the logging process. Therefore, you should disable the logging process after you capture the events that you want in the log file. To disable logging, remove the LogLevel registry value, and then restart the computer. |

You can automate this process by using the script that is in the "Account Lockout Tools" section in this document. This script sets the Kerberos logging key in the registry on client computers that are running Windows 2000. If you want to enable logging for groups of computers, you can specify this script as a startup script in an Active Directory Group Policy.

**To disable Kerberos event logging on a computer**

1. Click **Start**, click **Run**, type **regedit**, and then press ENTER.

2. Delete the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters\LogLevel** registry value.

3. Close Registry Editor and restart the computer.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

For more information, see the "HOW TO: Enable Kerberos Event Logging" in the Microsoft Knowledge Base.

**Event and Netlogon Log Retrieval**

After you set the auditing and logging, wait until account lockouts occur. When the account lockout occurs, retrieve both the Security event log and the System event log, as well as the Netlogon logs for all of the computers that are involved with the client's lockout. This includes the PDC emulator operations master, the authenticating domain controller, and all of the client computers that have user sessions for the locked-out user.

To determine the domain controllers that are involved with the lockout, run the LockoutStatus.exe tool and specify the user account that is locked out. This tool gathers and displays information about the specified user account from all the domain controllers in the domain. In addition, the tool displays the user's **badPwdCount** value on each domain controller. The domain controllers that have a **badPwdCount** value that reflects the bad password threshold setting for the domain are the domain controllers that are involved in the lockout. These domain controllers always include the PDC emulator operations master.

The **badPwdCount** value may appear to be higher than the threshold because of the way that passwords are chained to the PDC emulator operations master. When a bad password is presented by a user or program, both the authenticating domain controller and the PDC emulator operations master increment their **badPwdCount** value for that account. When Active Directory replication occurs, this can result in an increased value. However, the end result—the account becoming locked out—remains the same.

You can also use the EventCombMT.exe tool to gather specific event log data from multiple computers to one central location. For more information about both the EventCombMT.exe and LockoutStatus.exe tools, see the "Account Lockout Tools" section in this document.

**Analyzing Log File Information**

The previous section described the processes that you can use to enable log files to record information that is lockout-specific on your computers. This section focuses on analyzing those log files and determining what behavior occurred that created the log files and caused the issue that you are trying to resolve. This section also describes how to resolve the issues that you find when you analyze the log files.

**Analyzing Netlogon Log Files**

Before you start to analyze the Netlogon log files, you should be familiar with the authentication process works from previous sections in this paper. Although this section describes an NTLM authentication process, a similar chain of events occurs during Kerberos authentication.

The following sample scenario discusses what occurs when a user who is on a client computer tries to gain accesses to a resource that is on a file server in the same domain as the user account. In this process:

1. User credentials are passed to the file server. This is displayed in the Network Logon section in the Netlogon.log file.

2. The file server tries to authenticate the user, but the file server has to forward the credentials to the authenticating domain controller for validation because this account is a domain user account. This behavior is displayed as Transitive Network logon in the Netlogon.log file and is commonly referred to as pass-through authentication.

3. If the password is incorrect or if it is not the same as the password that is stored by the authenticating domain controller, the authenticating domain controller chains the credentials to the PDC for validation. This is displayed as Transitive Network Logon in the Netlogon.log file.

**Netlogon Log File Walkthrough**

The following sections provide a sample Netlogon.log file output from the following three computers:

- The PDC operations master for the domain: DC002
- The authenticating domain controller: DC003
- The member server: MEMSERVER01

The sample output sections show the following participants involved in network authentication:

- Domain name: Tailspintoys
- Logon user name: User1
- Logon computer name: Computer-006

**Transitive Network Logon (Pass-Through Authentication)**

Sample from the DC002 PDC emulator Netlogon log file:

```
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via DC003) 0xC000006A
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via DC003) 0xC000006A
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via DC003) 0xC000006A
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via DC003) 0xC000006A
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via DC003) 0xC000006A
29-Mar 14:28:31 Transitive Network logon Tailspintoys\User1
Computer-006 (via DC003) 0xC0000234
```

Sample from the DC003 authentication domain controller Netlogon log file:

```
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via MEMSERVER01) 0xC000006A
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via MEMSERVER01) 0xC000006A
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via MEMSERVER01) 0xC000006A
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via MEMSERVER01) 0xC000006A
29-Mar 14:28:31 Transitive Network logon Tailspintoys\User1
Computer-006 (via MEMSERVER01) 0xC000006A
29-Mar 14:28:31 Transitive Network logon Tailspintoys\User1
Computer-006 (via MEMSERVER01) 0xC0000234
```

Sample from the MEMSERVER01 member server Netlogon log file:

```
29-Mar 14:28:31 Network logon Tailspintoys\User1
Computer-006 0xC000006A
29-Mar 14:28:31 Network logon Tailspintoys\User1
Computer-006 0xC000006A
29-Mar 14:28:32 Network logon Tailspintoys\User1
Computer-006 0xC000006A
29-Mar 14:28:32 Network logon Tailspintoys\User1
Computer-006 0xC000006A
29-Mar 14:28:32 Network logon Tailspintoys\User1
Computer-006 0xC000006A
29-Mar 14:28:32 Network logon Tailspintoys\User1
```

```
Computer-006 0xC0000234
```

These Netlogon.log file samples provide an example of the information contained in the Netlogon logs. This information is used to trace the account lockout from the domain controller back to the member server on which a user or application tried to gain access with the incorrect credentials.

**Stepping Through the Netlogon Log File Sample**
This section describes the standard analysis process of log files when attempting to determine the cause of an account lockout issue.

In most troubleshooting scenarios, you should begin your log file analysis by examining the Netlogon.log file on the PDC operations master. Because this is a transitive network logon, you can find the authenticating domain controller by viewing the "Via" line in the Netlogon.log file for the domain controller that is chaining logon to the PDC operations master.

**Sample from the PDC Emulator (DC002) Netlogon Log File**
On the "Via" line from the PDCs Netlogon.log file in the following example, note that the authentication is being chained from DC003.

```
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer006 (via DC003) 0xC000006A
```

This is an illustration of step 3 of the authentication process that was detailed in "Analyzing Netlogon Log Files" previously in this document.

**Sample from the Authentication Domain Controller (DC003) Netlogon Log File**
In the Netlogon log file on DC003, this authentication is still a transitive network logon, because credentials were sent to DC003 for verification. Because of this, note where the credentials are sent from. In this example, the credentials are being sent via MEMSERVER01":

```
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer006 (via MEMSERVER01) 0xC000006A
```

This is an illustration of step 2 of the authentication process that was detailed in "Analyzing Netlogon Log Files" previously in this document.

The file server tries to authenticate the user, but the file server has to forward the credentials to the authenticating domain controller for validation because this is a domain user account. This is displayed as Transitive Network logon in the FileServername section of the Netlogon.log file.

**Sample from the Member Server (MEMSERVER01) Netlogon Log File**
From the Netlogon.log file on MEMSERVER01 from the same time period, verify the actual client computer name where the original logon or session setup request came from. In this example, the request came from Computer-006:

```
29-Mar 14:28:31 Network logon   Tailspintoys\User1
Computer-006 0xC000006A
```

This is an illustration of step 1 of the authentication process that was detailed in "Analyzing Netlogon Log Files" previously in this document.

User credentials are passed to the file server. This is displayed in the Network Logon section in the Netlogon.log file.

Even though the log file does not display the exact process that is sending the incorrect credentials, Netlogon log files do provide the following information to help you troubleshoot the lockout:

- Netlogon output displays the number of unsuccessful logon attempts (0xC000006A) for a user's account in a certain time period. Logs in which there are several 0xC000006A events in one second indicate that the lockout is most likely caused by a process, program, or script that is sending incorrect credentials.

- Netlogon output provides a complete picture of all computers that are involved in the account lockout. You can narrow down the culprit by determining the common elements, such as programs, among the computers involved. For example, from the Netlogon output above, after you determined that MEMSERVER01 was common to all user lockouts, the troubleshooting focus changed to the particular network services or user accounts that are used by MEMSERVER01.

- In this example, MEMSERVER01 is the Microsoft Exchange server. After you examine the Microsoft Outlook client and Exchange server settings, you may want to use the information that is in the following two articles to help resolve the issue. These articles describe how to remove unnecessary RPC bindings from the Exchange server. For example, remove Named Pipe support if there is no client that requires the named pipes.

- "Outlook Locks Your Account Because of a Directory Service Referral with Exchange 2000 Server" in the Microsoft Knowledge Base.

- "Unexpected Account Lockouts Caused When Logging On to Outlook from an Untrusted Domain" in the Microsoft Knowledge Base.

- If the Netlogon logs from all domain controllers from the time of lockout but do not display data that pertains to any of the locked-out user accounts that you are analyzing, then NTLM authentication is not involved in the lockouts. This normally indicates that the authentication issues are between computers running Windows 2000 or later, because earlier versions of Windows used NTLM authentication exclusively. You should focus on Kerberos authentication troubleshooting by using Kerberos logging and examining the Security event logs.

**Netlogon Log File Error Codes**

Each event in the Netlogon log contains a corresponding error code. The following table describes these error codes.

**Table 3 Netlogon Log Error Codes**

| Log Code | Description |
|---|---|
| 0x0 | Successful login |
| 0xC0000064 | The specified user does not exist |
| 0xC000006A | The value provided as the current password is not correct |
| 0xC000006C | Password policy not met |
| 0xC000006D | The attempted logon is invalid due to a bad user name |
| 0xC000006E | User account restriction has prevented successful login |
| 0xC000006F | The user account has time restrictions and may not be logged onto at this time |
| 0xC0000070 | The user is restricted and may not log on from the source workstation |
| 0xC0000071 | The user account's password has expired |
| 0xC0000072 | The user account is currently disabled |
| 0xC000009A | Insufficient system resources |
| 0xC0000193 | The user's account has expired |
| 0xC0000224 | User must change his password before he logs on the first time |
| 0xC0000234 | The user account has been automatically locked |

**Note**

Many of these codes provide information in the log file that is redundant with the corresponding Netlogon event log. This allows you to analyze the events in a variety of ways.

**Frequently Asked Questions**

This section answers common questions that might be helpful in troubleshooting the Account Lockout feature in the Windows Server 2003 family.

- Do the logon attempts happen seconds apart or are there many invalid logon attempt events (error code 0xC000006A) in the same second?

  The pattern shows whether this is a user error or a process or program that is creating the lockout. Users take several seconds between events, however, a process or program typically register many invalid or unsuccessful logon attempt

events in one second.

| Note |
| --- |
| You can use the NLParse.exe tool to parse Netlogon logs for specific events that are related to account lockout, such as events 0xC000006A and 0xC0000234. The parsed data is sent to a .csv file that you can read by using a program like Microsoft Excel. For more information about NLParse, see the "Account Lockout Tools" section in this document. |

- From which computers are the invalid logon attempt events generated?

  When you review the Netlogon logs and event logs, you can isolate from which computer the user was logged on during the account lockouts. In many situations, you will discover that a user is logged onto multiple computers; after the user changes their password on one computer, the user account is locked out.

- What client computers are displayed in the Netlogon log files?

  If only Windows 98 or Windows 95 clients are locked out, you may need to install the directory service client for those clients. For example, below, the Computer-006 computer generates the invalid logon attempt event:

```
29-Mar 14:28:30 Transitive Network logon Acme\User1
Computer-006 (via DC003) 0xC000006A
```

- Which user accounts are associated with the invalid logon attempt events?

  If privileged accounts (such as the administrator, service accounts, and well-known application account names) are receiving a large number of incorrect password attempts, first review the information for the computers that have made the attempts with the incorrect passwords, and then determine if there is a wrong password for the account. After you do this, if the passwords on all of the accounts are reset and incorrect password attempts persist, perform a trace to determine if the computer is under an attack. You can place an event trigger to stop the trace and determine where the attempt may be coming from. Internal or external computers can be a threat if there are worm viruses or compromises. The following example shows that the 0xC000006A error code is generated from the Tailspintoys\User1 user:

```
29-Mar 14:28:30 Transitive Network logon Tailspintoys\User1
Computer-006 (via DC003) 0xC000006A
```

- Is there an obvious pattern to the invalid logon attempt events and account lockouts?

  **Pattern:** All users on the domain are locked out, including users who did not change their password. There are many unsuccessful authentication attempts per second.

  **Possible Solution:** If you determine that the log files show that most or all of the user accounts are locked out in your domain, you must perform a trace to determine whether the source of the attack is internal or external to your network. If the attack appears to come from a internal computer, examine the processes running on these computers as this likely indicates a common process that uses outdated or incorrect credentials. Attacks from outside your network often indicate denial of service or brute force attacks against your user accounts.

  **Pattern:** Alphabetical list of users in the log files.

  **Possible Solution:** If the log files show that all of the user accounts are locked out in a list that is almost alphabetical, it is most likely that this behavior is caused by an attempt to break passwords or a denial of service attack. You must perform a network trace to the source of the attack.

  **Pattern:** A specific number of logon attempts are made on each locked-out account.

  **Possible Solution:** If you determine that the log files display a specific number of logon attempts for a each user, add the number of occurrences of 0xC000006A and 0xC0000234 errors for the user. In some scenarios, you may see a pattern of a specific number of attempts for one user, and then the same number of attempts for another user, and so on. This behavior may be an attack on the network or a program could be sending a specific set of attempts. 16 or 17

attempts per user is a common figure for these types of attacks.

In most account lockout situations, you must use Netlogon log files to determine which computers are sending bad credentials. When you analyze Netlogon log files, look for the 0xC000006A event code, because this event will help you determine where the bad password attempts began to occur. When you see the 0xC000006A event code and it is followed by a 0xC0000234 event code, the event codes that come after these event codes help you determine what caused the account lockout. If you see patterns in the log files, the patterns can help you determine if the event code was logged because of either a program attack or user error.

**Analyzing Event Logs**

You cannot determine the authentication type that was used when an account is locked out unless you enable Netlogon logging before the account lockout. However, because of differences in authentication, there may be situations in which Netlogon logging does not capture the data that you need to determine which computers were involved in an account lockout. Configuring the appropriate computers to create event logs may provide additional information in these situations.

Before the problems occur, you should enable security auditing and Kerberos logging on all computers that might be involved in the account lockout event. Enabling auditing and Netlogon log files is discussed elsewhere in this document. If the auditing is not configured before the initial error occurs, it can be done afterwards.

Once the account lockout occurs, there are several tasks that should be completed to help identify the cause of the issue:

1. Obtain both the Security and System event logs from all of the computers that are locked out if those computers were logged on when the lockout occurred. Also, obtain these log files from the PDC emulator operations master and all domain controllers that may be involved in the account lockout.

2. Look for Event 675 (Preauthentication Failures) in the Security event log for the domain controllers for the locked-out user account. This event displays the IP address of the client computer from which the incorrect credentials were sent. When you view these events in the Security event log from the PDC, an IP address with Event 675 may be the IP address of another domain controller because of password chaining from other domain controllers. If this is true, obtain the Security event log from that domain controller to see the Event 675. The IP address that is listed in that Event 675 should be the IP address for the client computer that sent the invalid credential.

3. After you know which client computer is sending the invalid credentials, determine the services, programs, and mapped network drives on that computer. If this information does not reveal the source of the account lockout, perform network traces from that client computer to isolate the exact source of the lockout.

| **Note** |
| --- |
| You can use the EventCombMT.exe tool to gather event log dates from different domain controllers at the same time. For more information about EventCombMT.exe, see the Account Lockout Tools section in this document. |

For more information, see the following articles:

- "Windows 2000 Security Event Descriptions (Part 1 of 2)" in the Microsoft Knowledge Base.

- "Windows 2000 Security Event Descriptions (Part 2 of 2) in the Microsoft Knowledge Base.

**Example Event Log Entry: Incorrect Password Processed by Kerberos**

The following example displays a sample Event 675 in the Security event log from the PDC emulator operations master:

```
Event Type: Failure Audit
Event Source: Security
Event Category: Account Logon
Event ID: 675
Date: 12/5/2001
Time: 5:47:26 PM
User: NT AUTHORITY\SYSTEM
Computer: COMPUTER-006
Description:
Pre-authentication failed:
User Name: user1
```

```
User ID: %{S-1-5-21-4235101579-1759906425-16398432-1114}
Service Name: krbtgt/Tailspintoys.com
Pre-Authentication Type: 0x2
Failure Code: 0x18
Client Address: 172.16.1.85
```

In this example, failure code 0x18 is listed because an incorrect user name or password was used. The client address of 172.16.1.85 identifies the network client that caused this failure. The user name "user1" is also included in this event. The client address and user name should provide enough information for you to begin to address the issue, because you know which user is attempting to logon from which computer.

**Example Event Log Entry: Account is Locked Out**
The following example displays a sample of Event 644, which indicates that the account is locked out:

```
Event Type: Success Audit
Event Source: Security
Event Category: Account Management
Event ID: 644
Date: 12/5/2001
Time: 5:47:26 PM
User: Everyone
Computer: COMPUTER-006
Description:
User Account Locked Out:
Target Account Name: user1
Target Account ID:%{S-1-5-21-4235101579-1759906425-16398432-1114}
Caller Machine Name:COMPUTER-006
Caller User Name:USER1$
Caller Domain:TAILSPINTOYS
Caller Logon ID:(0x0,0x3E7)
```

For more information about account lockout events, see "Audit Account Lockout" on the Microsoft TechNet Web site.

**Example Event Log Entry: Logon Failure**
The following example displays a sample of Event 529, which results from an unsuccessful logon attempt due to an invalid user name or password. This event is often useful in identifying the source of the lockout:

```
Event Type: Failure Audit
Event Source: Security
Event Category: Logon/Logoff
Event ID: 529
Date: 12/21/2001
Time: 2:05:20 PM
User: NT AUTHORITY\SYSTEM
Computer: COMPUTER-006
Description:
Logon Failure:
   Reason: Unknown user name or bad password
   User Name: user1
   Domain: Tailspintoys
   Logon Type: 2
   Logon Process: User32
   Authentication Package:    Negotiate
   Workstation Name: COMPUTER-006
```

This event contains several useful elements. It identifies the name of the computer that is attempting authentication, as well as the user and domain name. It also displays the logon type, which is discussed later in this section.

When Event 529 is logged, you should look for patterns in the event. Determine if there are several 529 events logged and determine if they all occur in one second or if they occur at specific time intervals. If so, there is a process or service that is running on the computer that is sending incorrect credentials. Look at the Logon Process and Logon Type entries in the log to determine the type of process that is passing incorrect credentials and to determine how the process is logging on.

**Example Event Log Entry: Account Is Disabled**
When there is an attempt to logon using a disabled account, a specific event is created in the event log. This can help you quickly identify intruders, because normal operations should not allow for the use of locked out accounts. You should analyze and respond quickly to these events.

```
Event Type:    Failure Audit
Event Source: Security
Event Category: Logon/Logoff
```

```
Event ID: 531
Date: 12/21/2001
Time: 2:05:21 PM
User: NT AUTHORITY\SYSTEM
Computer: COMPUTER-006
Description:
   Logon Failure:
   Reason: Account currently disabled
   User Name: user1
   Domain: TAILSPINTOYS
   Logon Type: 2
   Logon Process: User32
   Authentication Package: Negotiate
```

**Kerberos Events Logged During an Account Lockout**

Once Kerberos logging is enabled, certain events will be logged when an account lockout occurs. These events are described in this section.

**Incorrect Password**

This event is logged when an incorrect password is received by Kerberos as part of an authentication request.

```
Event Type:   Error
Event Source: Kerberos
Event Category: None
Event ID: 4
Date: 12/21/2001
Time: 2:02:05 PM
User: N/A
Computer: COMPUTER-006
Description:
   The function LogonUser received a Kerberos Error Message:
   on logon session TAILSPINTOYS\user1
   Client Time:
   Server Time: 19:2:5.0000 12/21/2001 (null)
   Error Code: 0x18 KDC_ERR_PREAUTH_FAILED
   Client Realm:
   Client Name:
   Server Realm: TAILSPINTOYS.COM
   Server Name: krbtgt/TAILSPINTOYS.COM
   Target Name: krbtgt/TAILSPINTOYS@TAILSPINTOYS
   Error Text:
   File:
   Line: Error Data is in record data.
```

**Kerberos Event When a User Account Is Locked Out**

This event is logged when Kerberos is used for authentication and an account lockout occurs.

```
Event Type: Error
Event Source: Kerberos
Event Category: None
Event ID: 4
Date: 12/21/2001
Time: 2:05:21 PM
User: N/A
Computer: COMPUTER-006
Description:
   The function LogonUser received a Kerberos Error Message:
   on logon session TAILSPINTOYS\user1
   Client Time:
   Server Time: 19:5:21.0000 12/21/2001 (null)
   Error Code: 0x12 KDC_ERR_CLIENT_REVOKED
   Client Realm:
   Client Name:
   Server Realm: TAILSPINTOYS.COM
   Server Name: krbtgt/TAILSPINTOYS.COM
   Target Name: krbtgt/TAILSPINTOYS@TAILSPINTOYS
   Error Text:
   File:
   Line: Error Data is in record data
```

**Logon Events**

Many different events can be created by various logon and logoff actions. The following table describes each logon event.

**Table 4 Logon Event IDs**

| Event ID | Description |
|---|---|
| 528 | A user successfully logged on to a computer. For information about the type of logon, see the Logon Types table below. |
| 529 | Logon failure. A logon attempt was made with an unknown user name or a known user name with a bad password. |
| 530 | Logon failure. A logon attempt was made, but the user account tried to log on outside of the allowed time. |
| 531 | Logon failure. A logon attempt was made using a disabled account. |
| 532 | Logon failure. A logon attempt was made using an expired account. |
| 533 | Logon failure. A logon attempt was made by a user who is not allowed to log on at this computer. |
| 534 | Logon failure. The user attempted to log on with a type that is not allowed. |
| 535 | Logon failure. The password for the specified account has expired. |
| 536 | Logon failure. The Netlogon service is not active. |
| 537 | Logon failure. The logon attempt failed for other reasons.<br>**Note**<br>In some cases, the reason for the logon failure may not be known. |
| 538 | The logoff process was completed for a user. |
| 539 | Logon failure. The account was locked out at the time the logon attempt was made. |
| 540 | A user successfully logged on to a network. |
| 541 | Main mode Internet Key Exchange (IKE) authentication was completed between the local computer and the listed peer identity (establishing a security association), or quick mode has established a data channel. |
| 542 | A data channel was terminated. |
| 543 | Main mode was terminated.<br>**Note**<br>This might occur as a result of the time limit on the security association expiring, policy changes, or peer termination. (The default expiration time for security associations is eight hours.) |
| 544 | Main mode authentication failed because the peer did not provide a valid certificate or the signature was not validated. |
| 545 | Main mode authentication failed because of a Kerberos failure or a password that is not valid. |
| 546 | IKE security association establishment failed because the peer sent a proposal that is not valid. A packet was received that contained data that is not valid. |
| 547 | A failure occurred during an IKE handshake. |
| 548 | Logon failure. The security identifier (SID) from a trusted domain does not match the account domain SID of the client. |
| 549 | Logon failure. All SIDs that correspond to untrusted namespaces were filtered out during an authentication across forests. |
| 550 | A denial-of-service attack may have taken place. |
| 551 | A user initiated the logoff process. |

| | |
|---|---|
| 552 | A user successfully logged on to a computer using explicit credentials while already logged on as a different user. |
| 672 | An authentication service (AS) ticket was successfully issued and validated. |
| 673 | A ticket-granting service (TGS) ticket was granted. |
| 674 | A security principal renewed an AS ticket or TGS ticket. |
| 675 | Preauthentication failed. This event is generated on a Key Distribution Center (KDC) when a user types in an incorrect password. |
| 676 | Authentication ticket request failed. This event is not generated in Windows XP or in the Windows Server 2003 family. |
| 677 | A TGS ticket was not granted. This event is not generated in Windows XP or in the Windows Server 2003 family. |
| 678 | An account was successfully mapped to a domain account. |
| 681 | Logon failure. A domain account logon was attempted. This event is not generated in Windows XP or in the Windows Server 2003 family. |
| 682 | A user has reconnected to a disconnected terminal server session. |
| 683 | A user disconnected a terminal server session without logging off.<br>**Note**<br>This event is generated when a user is connected to a terminal server session over the network. It appears on the terminal server. |

**Netlogon Logon Types**

When many Netlogon logon events are logged, a logon type is also listed in the event details. The following table describes each logon type.

**Table 5 Netlogon Logon Types**

| Logon type | Logon title | Description |
|---|---|---|
| 2 | Interactive | A user logged on to this computer. |
| 3 | Network | A user or computer logged on to this computer from the network. |
| 4 | Batch | The batch logon type is used by batch servers, where processes may be executing on behalf of a user without their direct intervention. |
| 5 | Service | A service was started by the Service Control Manager. |
| 7 | Unlock | This workstation was unlocked. |
| 8 | NetworkCleartext | A user logged on to this computer from the network. The user's password was passed to the authentication package in its unhashed form. The built-in authentication packages all hash credentials before sending them across the network. The credentials do not traverse the network in plaintext (also called cleartext). |
| 9 | NewCredentials | A caller cloned its current token and specified new credentials for outbound connections. The new logon session has the same local identity, but uses different credentials for other network connections. |
| 10 | RemoteInteractive | A user logged on to this computer remotely using Terminal Services or Remote Desktop. |
| 11 | CachedInteractive | A user logged on to this computer with network credentials that were stored locally on the computer. The domain controller was not contacted to verify the credentials. |

# Troubleshooting Account Lockout

Updated: June 9, 2008

Applies To: Windows Server 2003 with SP1

In an environment where you set the account lockout feature, you may notice a large number of lockouts that occur. To determine if these lockouts are false lockouts or a real attack:

1. Verify that the domain controllers and client computers are up-to-date with service packs and hotfixes. For more information, see the "Recommended Service Packs and Hotfixes" section in this document.

2. Configure your computers to capture data:

   a. Enable auditing at the domain level.

   b. Enable Netlogon logging.

   c. Enable Kerberos logging.

   For more information, see the "Appendix Two: Gathering Information to Troubleshoot Account Lockout Issues" section in this document.

3. Analyze data from the Security event log files and the Netlogon log files to help you determine where the lockouts are occurring and why.

4. Analyze the event logs on the computer that is generating the account lockouts to determine the cause.

   For more information, see the Account Lockout Tools section in this document.

The following section further describes the account lockout troubleshooting process.

**Common Causes for Account Lockouts**

This section describes some of the common causes for account lockouts The common troubleshooting steps and resolutions for account lockouts are also described in this section.

To avoid false lockouts, check each computer on which a lockout occurred for the following behaviors:

- **Programs:** Many programs cache credentials or keep active threads that retain the credentials after a user changes their password.

- **Service accounts:** Service account passwords are cached by the service control manager on member computers that use the account as well as domain controllers. If you reset the password for a service account and you do not reset the password in the service control manager, account lockouts for the service account occur. This is because the computers that use this account typically retry logon authentication by using the previous password. To determine whether this is occurring, look for a pattern in the Netlogon log files and in the event log files on member computers. You can then configure the service control manager to use the new password and avoid future account lockouts.

- **Bad Password Threshold is set too low:** This is one of the most common misconfiguration issues. Many companies set the Bad Password Threshold registry value to a value lower than the default value of 10. If you set this value too low, false lockouts occur when programs automatically retry passwords that are not valid. Microsoft recommends that you leave this value at its default value of 10. For more information, see "Choosing Account Lockout Settings for Your Deployment" in this document.

- **User logging on to multiple computers:** A user may log onto multiple computers at one time. Programs that are running on those computers may access network resources with the user credentials of that user who is currently logged on. If the user changes their password on one of the computers, programs that are running on the other computers may continue to use the original password. Because those programs authenticate when they request access to network

resources, the old password continues to be used and the users account becomes locked out. To ensure that this behavior does not occur, users should log off of all computers, change the password from a single location, and then log off and back on.

| **Note** |
| --- |
| Computers running Windows XP or a member of the Windows Server 2003 family automatically detect when the users password has changed and prompt the user to lock and unlock the computer to obtain the current password. No logon and logoff is required for users using these computers. |

- **Stored user names and passwords retain redundant credentials:** If any of the saved credentials are the same as the logon credential, you should delete those credentials. The credentials are redundant because Windows tries the logon credentials when explicit credentials are not found. To delete logon credentials, use the Stored User Names and Passwords tool. For more information about Stored User Names and Passwords, see online help in Windows XP and the Windows Server 2003 family.

| **Note** |
| --- |
| Computers that are running Windows 95, Windows 98, or Windows Millennium Edition do not have a Stored User Names and Passwords file. Instead, you should delete the user's .pwl file. This file is named *Username*.pwl, where *Username* is the user's logon name. The file is stored in the Systemroot folder. |

- **Scheduled tasks:** Scheduled processes may be configured to using credentials that have expired.

- **Persistent drive mappings:** Persistent drives may have been established with credentials that subsequently expired. If the user types explicit credentials when they try to connect to a share, the credential is not persistent unless it is explicitly saved by Stored User Names and Passwords. Every time that the user logs off the network, logs on to the network, or restarts the computer, the authentication attempt fails when Windows attempts to restore the connection because there are no stored credentials. To avoid this behavior, configure **net use** so that is does not make persistent connections. To do this, at a command prompt, type **net use /persistent:no**. Alternately, to ensure current credentials are used for persistent drives, disconnect and reconnect the persistent drive.

- **Active Directory replication:** User properties must replicate between domain controllers to ensure that account lockout information is processed properly. You should verify that proper Active Directory replication is occurring.

- **Disconnected Terminal Server sessions:** Disconnected Terminal Server sessions may be running a process that accesses network resources with outdated authentication information. A disconnected session can have the same effect as a user with multiple interactive logons and cause account lockout by using the outdated credentials. The only difference between a disconnected session and a user who is logged onto multiple computers is that the source of the lockout comes from a single computer that is running Terminal Services.

- **Service accounts:** By default, most computer services are configured to start in the security context of the Local System account. However, you can manually configure a service to use a specific user account and password. If you configure a service to start with a specific user account and that accounts password is changed, the service logon property must be updated with the new password or that service may lock out the account.

| **Note** |
| --- |
| You can use the System Information tool to create a list of services and the accounts that were used to start them. To start the System Information tool, click **Start**, click **Run**, type **winmsd**, and then click **OK**. |

#### Other Potential Issues

Some additional considerations regarding account lockout are described in the following sections.

#### Account Lockout for Remote Connections

The account lockout feature that is discussed in this paper is independent of the account lockout feature for remote connections, such as in the Routing and Remote Access service and Microsoft Internet Information Services (IIS). These services

and programs may provide their own unrelated account lockout features.

**Internet Information Services**

By default, IIS uses a token-caching mechanism that locally caches user account authentication information. If lockouts are limited to users who try to gain access to Exchange mailboxes through Outlook Web Access and IIS, you can resolve the lockout by resetting the IIS token cache. For more information, see "Mailbox Access via OWA Depends on IIS Token Cache" in the Microsoft Knowledge Base.

**MSN Messenger and Microsoft Outlook**

If a user changes their domain password through Microsoft Outlook and the computer is running MSN Messenger, the client may become locked out. To resolve this behavior, see "MSN Messenger May Cause Domain Account Lockout After a Password Change" in the Microsoft Knowledge Base.

# Account Lockout Tools

Updated: November 9, 2007

Applies To: Windows Server 2003 with SP1

After you determine the pattern for the account lockouts and narrow down your scope to a specific client computer or member server, you should gather detailed information about all of the programs and services that are running on that computer. Some of the information that you should obtain includes:

- Mapped network drives
- Logon scripts that map network drives
- **RunAs** shortcuts
- Accounts that are used for service account logons
- Processes on the client computers
- Programs that may pass user credentials to a centralized network program or middle-tier application layer

The following sections discuss the tools that you can use to help you gather information from the network environment.

**The LockoutStatus.exe Tool**

The LockoutStatus.exe displays information about a locked out account. It does this by gathering account lockout-specific information from Active Directory. The following list describes the different information that is displayed by the tool:

- **DC Name:** Displays all domain controllers that are in the domain.
- **Site:** Displays the sites in which the domain controllers reside.
- **UserState:** Displays the status of the user and whether that user is locked out of their account.
- **Bad Pwd Count:** Displays the number of bad logon attempts on each domain controller. This value confirms the .domain controllers that were involved in the account lockout.
- **Last Bad Pwd:** Displays the time of the last logon attempt that used a bad password.
- **Pwd Last Set:** Displays the value of the last good password or when the computer was last unlocked.
- **Lockout Time:** Displays the time when the account was locked out.
- **Orig Lock:** Displays the domain controller that locked the account (the domain controller that made the originating write to the **LockoutTime** attribute for that user).

**Where to Obtain the LockoutStatus.exe Tool**

LockoutStatus.exe is included with the ALTools.exe package that is available at "Account Lockout and Management Tools" on the Microsoft Web site.

**How to Install the LockoutStatus.exe Tool**

To install the LockoutStatus.exe tool, install the ALTools package on your domain controller.

**How to Use the LockoutStatus.exe Tool**

To run the LockoutStatus.exe tool and display information about a locked out user account:

1. Double-click **LockoutStatus.exe**.

2. On the **File** menu, click **Select target**.

3. Type the user name whose lockout status on the enterprise's domain controllers you want information about.

The following figure displays an example where two domain controllers have a **badPwdCount** value of 5, which is also the bad password threshold. One domain controller is the PDC operations master, and the other domain controller is the authenticating domain controller. These two domain controllers are displayed because of password chaining from the authenticating domain controller to the PDC.

**Figure 2: The LockoutStatus.exe Tool**




**The ALockout.dll Tool**

The ALockout.dll tool and the Appinit.reg script are included in the ALTools package. ALockout.dll is a logging tool that may help you determine the program or process that is sending the incorrect credentials in an account lockout scenario. The tool attaches itself to a variety of function calls that a process might use for authentication. The tool then saves information about the program or process that is making those calls into the *Systemroot*\Debug\Alockout.txt file. The events are time stamped so that you can match them to the events that are logged in either the Netlogon log files or the Security event log files.

You can use Appinit.reg to initialize the .dll file. This file provides no other functionality.

| **Note** |
| --- |
| Microsoft does not recommend that you use this tool on servers that host network programs or services. You should not enable ALockout.dll on Exchange servers because the ALockout.dll tool may prevent the Exchange store from starting. |

| **Important** |
| --- |
| Before you install the ALockout.dll tool on any mission-critical computer, make a full backup copy of the operating system and any valuable data. |

For more information, see "Errors Installing Exchange Server with CleanSweep" on the Microsoft Knowledge Base.

In most account lockout scenarios, you should install ALockout.dll on client computers. Use the information that is stored in both the Netlogon log file and the Security event log to determine the computers from which the incorrect credentials are being sent that are locking out the user's account. When you install the ALockout.dll tool on the client computer that is sending the incorrect credentials, the tool logs the process that is sending the incorrect credentials.

**Where to Obtain the ALockout.dll Tool**

ALockout.dll is included with the ALTools.exe package that is available at "Account Lockout and Management Tools" on the Microsoft Web site.

**How to Install the ALockout.dll Tool**

There two versions of the ALockout.dll file. One version of the file is for computers that are running a Windows 2000 operating system, and the other version of the file is for computers that are running a Windows XP operating system. View the Readme.txt file that is included with the ALTools package.

**To install ALockout.dll**

1. On the computer that has generated account lockout error messages in the Security event log, copy both the ALockout.dll and Appinit.reg files to the *Systemroot*\System32 folder .

2. Double-click the Appinit.reg file to run the script. When you do this, the ALockout.dll file is registered and can begin providing information.

3. Restart the computer to complete the installation.

**How to Remove the ALockout Tool**

**To remove the ALockout.dll file from the computer**

1. At a command prompt, type **regsvr32 /u alockout.dll**.

2. Delete the Alockout.dll value that is under the following registry key:

   HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\

   Windows AppInit_DLLs

   After you delete the **Alockout.dll** value, the **AppInit_DLLs** registry key is blank.

3. Restart the computer.

4. Delete the ALockout.dll file from the *Systemroot*/System32 folder.

**How to Use the ALockout.dll Tool**

You should use the ALockout.dll tool with Netlogon logging and security auditing. To use the ALockout.dll tool:

1. Wait for an account to lock out on the computer.

2. When an account is locked out, the ALockout.txt file is created in the *Systemroot*\Debug folder.

3. Compare event time stamps in ALockout.txt with the time stamps in both the Netlogon log files and the Security event log files. When you do this, you can determine the process that is causing the lockouts.

You can use the ALockout.dll tool if you have already set up Netlogon logging, as well as Kerberos and logon auditing on the local computer. ALockout.dll does not interfere with any other logging or event generation.

**The ALoInfo.exe Tool**

If account lockouts seem to happen most frequently after a user is forced to change their password, you may want to determine which users' passwords are about to expire. You can use the ALoInfo.exe tool to display all user account names and the password age for those user accounts. This will allow you to use the ALockout.dll tool and other account lockout tools to set up the tools prior to the initial account lockout. You can also obtain a list of all local services and startup account information by using the ALoInfo.exe tool.

| **Note** |
| --- |
| You can also use the SecDump tool to display password expiration information in a Windows NT Server 4.0 domain. You can download this tool from the SystemTools Web site. Note that Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here. |

**Where to Obtain the ALoInfo.exe Tool**

The ALoInfo.exe file is included with the ALTools.exe package that is available at "Account Lockout and Management Tools" on

the Microsoft Web site.

**How to Install the ALoInfo.exe Tool**
To install the ALoInfo.exe tool, install the ALTools package on your domain controller. The ALTools package contains the ALoInfo.exe tool.

**How to Use the ALoInfo.exe Tool**
You can use ALoInfo.exe at a command prompt with either of the following methods:

- To display an accounts password ages from a domain controller, at a command prompt, type the following:

```
aloinfo /expires /server:Domain_Controller_Name
```

- To display all local service startup account information and mapped drive information for a user who is currently logged on, at a command prompt, type the following command:

```
aloinfo /stored /server:Computer_Name
```

You can redirect the output of ALoInfo.exe to a text file and then sort the results to determine which users may be involved in the account lockout. This information can also be stored for later analysis.

**The AcctInfo.dll Tool**
You can use the AcctInfo.dll tool to add new property pages to user objects in the Active Directory Users and Computers MMC Snap-in. You can use these property pages to help isolate and troubleshoot account lockouts and to reset a users password on a domain controller in that user's local site.

AcctInfo.dll displays the following user account information that you may be able to use to identify and resolve account lockout issues:

- Last time the password was set
- When the password will expire
- User Account Control Raw Value and Decode
- Time the account was locked out
- If the account is locked out now, when it will be unlocked
- Security identifier (SID) of the account, and its SIDHistory
- Globally unique identifier (GUID) of the account
- These account properties:
    - Last Logon
    - Last Logoff
    - Last Bad Logon Time
    - Logon Count
    - Bad Password Count

You can also use the AcctInfo.dll tool to obtain the domain password information (expiration, lockout time, and so on). You can type the user's computer name in the tool, and then reset the user's password on a domain controller in that user's site.

**Note**

Because of replication latency, domain controllers may store different information about the same user account. AcctInfo.dll displays information that is retrieved from a single domain controller.

**Where to Obtain the AcctInfo.dll Tool**
The AcctInfo.dll tool is included in the ALTools.exe package that is available at "Account Lockout and Management Tools" on the Microsoft Web site.

**How to Install the AcctInfo.dll Tool**
On the computer where you want to run Active Directory Users and Computers MMC Snap-in:

1. Copy the AcctInfo.dll file to the System32 folder.

2. At a command prompt, type **regsvr32 acctinfo.dll**, and then press ENTER.

The AcctInfo.dll file is registered and is displayed on a user's property sheet in the Active Directory Users and Computers MMC Snap-in after you follow these steps.

To use the **Account Lockout Status** button in the tool, verify that LockoutStatus.exe is in the *Systemroot*\System32 folder. If LockoutStatus.exe is not installed in this location, this button is unavailable.

**How to Use the AcctInfo.dll Tool**
To use the AcctInfo.dll tool, open the Active Directory Users and Computers MMC, right-click a user, click **Properties**, and then click **Additional Account Info**. An example of the information that is provided by AcctInfo.dll is shown in the following figure.

**Figure 3: Main Property Box**




The following figure displays the domain password policy information that you can view to determine the password policy that applies to the domain controller.

**Figure 4: Domain Password Policy**

**Change Password on a Domain Controller in the User's Site**
The AcctInfo.dll tool allows you to increase the functionality of the Active Directory Users and Computers MMC by adding the ability to reset a user's password in that user's local site. When you reset the password in the remote site, you avoid replication delays that can occur before that user logs on.

When you reset the password, you can also unlock the account and set the **User must change password** value. These options are in the **Change Password On a DC In The Users Site** box as displayed in the following figure.

**Figure 5: Change Password On a DC In The Users Site**




**How to Remove the AcctInfo.dll Tool**
To remove the AcctInfo.dll tool, delete the AcctInfo.dll file from the Systemroot/System32 folder, and then type the following command at a command prompt:

**regsvr32 /u acctinfo.dll**

**The EventCombMT.exe Tool**
You can use the EventCombMT.exe tool to gather specific events from event logs from several different computers into one central location. You can configure EventCombMT.exe to search for events and computers. Some specific search categories are built into the tool, such as account lockouts. Note that the account lockouts category is preconfigured to include events 529, 644, 675, 676, and 681.

**Figure 6: The EventCombMT.exe Tool**




**Where to Obtain the EventCombMT.exe Tool**
The EventCombMT.exe tool is included in the ALTools.exe package that is available at "Account Lockout and Management Tools" on the Microsoft Web site.

**How to Install the EventCombMT.exe Tool**
You do not need to install this tool separately. When you install ALTools on the domain controller, EventCombMT.exe is also installed to the directory you specified during setup.

**How to Use the EventCombMT.exe Tool**
To use the EventCombMT.exe tool, open the folder you specified during setup for ALTools, double-click **EventCombMT.exe**, click the **Searches** menu, click **Built in searches**, and then click **Account lockouts**. When you do this, the events that will be

pulled from the event logs are automatically displayed in the tool. These events are from all of the domain controllers in your environment. In addition to 529, 644, 675, and 681, type **12294** in the **Event Ids** box, and then click **Search**. The tool then searches the computers for these events, and then saves them to a .txt file that you specify.

**The NLParse.exe Tool**

Because Netlogon log files may become more than 10 MB in size, you may want to parse the files for the information that you want to view. You can use the NLParse.exe tool to parse Netlogon log files for specific Netlogon return status codes. The output from this tool is saved to a comma-separated values (.csv) file that you can open in Excel to sort further.

| **Note** |
| --- |
| The return codes that are specific to account lockouts are 0xC000006A and 0xC0000234. |

The following figure displays the interface for the NLParse.exe tool.

**Figure 7: Netlogon-Parse Return Status Codes**




**Where to Obtain the NLParse.exe Tool**

The NLParse.exe tool is included in the ALTools.exe package that is available at "Account Lockout and Management Tools" on the Microsoft Web site.

**How to Install the NLParse.exe Tool**

You do not need to install this tool separately; when you install ALTools on the domain controller, NLParse.exe is also installed.

**How to Use the NLParse.exe Tool**

To use the NLParse.exe tool, open the folder you specified during setup for ALTools, double-click **Nlparse.exe**, click **Open** to open the Netlogon.log file that you want to parse, select the check boxes for the status codes that you want to search for, and then click **Extract**. After you do this, view the output from the NLParse.exe tool. Typically, you may want to look at both the 0xC000006A and 0xC0000234 code statuses to determine from where the lockouts are coming.

**The FindStr.exe Tool**

You can also use the FindStr.exe tool to parse Netlogon log files. FindStr.exe is a command-line tool that you can use to parse several Netlogon.log files at the same time. After you gather the Netlogon.log files from several domain controllers, extract information about a specific user account from the files (user1, error code 0xC000006A, or error code0xC0000234). You can use this tool to help you obtain output about a user, computer, or error code in the Netlogon.log files.

**Where to Obtain the FindStr.exe Tool**

The FindStr.exe tool is included in the default installation of Windows 2000, Windows XP, and the Windows Server 2003 family operating systems. No additional installation or configuration is required for the FindStr.exe tool.

**How to Use the FindStr.exe Tool**

To use the FindStr.exe tool, rename the Netlogon.log files, and then save the files to one folder. To parse all of the Netlogon log

files, type the following command at a command prompt:

**FindStr /I User1** **netlogon** **.log >c:** *user1* **.txt**

**The Replmon and Repadmin Tools**
If you have not already verified Active Directory replication on a domain controller, at a command prompt, type **repadmin /showreps** or **replmon** to verify that proper Active Directory replication is occurring. In many scenarios, you may find that you unlock an account but the new credentials do not work. This behavior typically occurs because of replication latency. Change the users password in their local site to avoid replication latency issues.

**Where to Obtain the Relmon and Repadmin Tools**
Both of these tools are included with the support tools on the Windows 2000 CD-ROM.

**How to Install and Configure the Replmon and Repadmin Tools**
For more information about how to obtain and installing Replmon and Repadmin, see the Windows Support Tools documentation.

**Network Monitor**
Network Monitor is a powerful tool that you can use to capture unfiltered network communication.

If the account lockout occurs because of a process or program and an account is already locked out on a specific client computer, gather network traces of all traffic to and from that client computer while the account is still locked out. The program or process most likely will continue to send incorrect credentials while trying to gain access to resources that are on the network. Capturing all traffic to and from the client may help you determine which network resource the process is trying to gain access to. After you determine the network resource, you can determine which program or process is running on that client computer.

If you can narrow your search to a specific computer but the user account is not yet locked out, keep running Network Monitor until the lockout occurs for that user. After the lockout occurs, compare the time stamps of events when the in the Netlogon or Security event logs with the data that was captured in the trace. You should see that the network resource that is being accessed with incorrect credentials.

After you identify a program or service as the cause of the lockout, view the software manufacturers Web site for known resolutions. This behavior typically occurs because the program is running with the currently logged on user's credentials. If a service is causing the lockout, consider creating accounts that are specifically for running services so user account password changes do not affect the services.

**Where to Obtain Network Monitor**
The full version of Network Monitor is included with Microsoft System Management Server (SMS). A limited version of the tool is included with Windows XP and the Windows 2000 and Windows Server 2003 families.

**How to Install Network Monitor on Supported Operating Systems**
This section describes how to install Network Monitor on both the Windows 2000 Server family and Windows XP.

**Windows 2000 Server**
**To install Network Monitor on computers that are running Windows 2000 Server**

1. Right-click My Network Places, and then click Advanced.

2. Click Optional networking components, and then click Management and Monitoring.

For more information, see "HOW TO: Install Network Monitor in Windows 2000" in the Microsoft Knowledge Base.

**Windows XP**
Network Monitor is included with the Windows support tools. For more information about how to install and configure Network Monitor on computers that are running Windows XP, view the following articles:

- "How to Install the Support Tools from the Windows XP CD-ROM" on the Microsoft Knowledge Base.

- Description of the Network Monitor Capture Utility on the Microsoft Knowledge Base.

**How to Use Network Monitor**
For information about how to use Network Monitor to capture information, view the documentation that is provided with the tool or read "How to Capture Network Traffic with Network Monitor" on the Microsoft Knowledge Base.

# Appendix One: Additional References for Account Lockout

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

For more information about how to lock down your environment, as well as for more information about security features that are not addressed in this document, see the following Microsoft Web sites:

- "Microsoft Solution for Securing Windows 2000 Server" on the Microsoft TechNet Web site
- "Checklist: Create Strong Passwords" on the Microsoft Web site
- "New Registry Key to Remove LM Hashes from Active Directory and Security" on the Microsoft Web site
- "HOW TO: Configure Remote Access Client Account Lockout in Windows 2000" on the Microsoft Web site
- "HOW TO: Prevent Users From Changing a Password Except When Required in Windows 2000" on the Microsoft Web site
- "HOW TO: Prevent Users From Submitting Alternate Logon Credentials in Windows 2000" on the Microsoft Web site
- "HOW TO: Manage Stored User Names and Passwords on a Computer in a Domain in Windows XP" on the Microsoft Web site
- "Best Practices for Enterprise Security" on the Microsoft Web site

# Appendix Two: Gathering Information to Troubleshoot Account Lockout Issues

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

You can use the information in this section to help you gather information before you start to troubleshoot account lockout issues. Collect the following information to troubleshoot account lockout issues:

**Client Platform**
If client account lockouts are occurring on a single, common operating system, there may be specific issues with the operating system. Different operating systems use different processes for name resolution and authentication protocols, and they have different levels of security, so there might be an infrastructure or program issue. For more information, see the Microsoft Knowledge Base article "Service Packs and Hotfixes Available to Resolve Account Lockout Issues" on the Microsoft Web site.

You should gather the following information in these situations:

- Do users log on to multiple computers at the same time?
- Are there any common patterns? For example:
    - Do the computers have the same mapped drives?
    - Do the computers have the same mapped printers?
    - Do the computers have the same antivirus software?
    - Do the computers use management software?
    - Is another networking client installed on the computers?
    - Is SMS installed on the computers?
- Does the network include a Wide Area Network (WAN)?
- If the computer is running Windows 95, Windows 98, Windows 98 SE, or Windows Millennium Edition, what is the version of the Vredir.vxd file?
- If the computer is running Windows 95, Windows 98, Windows 98 SE, or Windows Millennium Edition, is the Directory Services Client installed on the computer?

**Domain Platform**
When you know the domain environment, the security boundaries, and how the user is gaining access to the resources that are in other domains, you can better determine the cause of the account lockouts. You should gather the following information:

- The number of domain controllers, including operating system, location, service pack level, and so on.
- Is Active Directory and Netlogon replication occurring?
- What domain does the user log onto?
- List all domain trusts that the user uses.
- Is there a matching account with the same logon name in the trusted domain?
- Are there any third-party SMB servers running in the environment?

**Background**
Look at the client and the network resources that the user might be contacting to help you determine the cause of the lockouts. You should gather the following information:

- When did you first notice the lockouts?
- When did the lockouts start?
- What has changed in the environment (new programs, new network services, and so on)?
- Are there any identifiable patterns:
    - After a password change?
    - When the user logs on?
    - When the user gains access to mapped drives?
    - When the user uses Outlook?
    - When the user uses Outlook Web Access?
    - Are there no identifiable patterns?
- How many user accounts are locked out each day, a small group of users or a large group?
- Is there an account password policy?
    - How many bad attempts are allowed before a lockout occurs?
    - How much time must elapse before the count resets?
    - What is the **LockoutDuration** registry value?

**Gathering Diagnostic Information**
Gather all of the different log files from Netlogon, Kerberos, and the event logs to help you determine the cause of the lockout. Diagnostic information that you gather from the computer from which the lockout is originating may help you determine the cause for the lockout. You should gather the following information:

- Netlogon log files
- Traces
- Event log files from client computers and domain controllers that are involved in the lockout

**Acknowledgements**
Ramesh Chinta, Program Manager, Microsoft Corporation

Mike Danseglio, Technical Writer, Microsoft Corporation

Mike Resnick, Technical Lead, Microsoft Corporation

Vincent Abella, Technical Editor, Microsoft Corporation

Jen Bayer, Technical Writer, Microsoft Corporation

Emily Moon, Technical Editor, S&T OnSite

Joseph Vasil, Support Engineer, Microsoft Corporation

# Guide to Creating and Configuring a Server Cluster under Windows Server 2003 White Paper

Updated: November 11, 2003

Applies To: Windows Server 2003 with SP1

**In This White Paper**


Introduction (Guide to Creating and Configuring a Server Cluster under Windows Server 2003 White Paper)

Checklists for Server Cluster Configuration

Cluster Installation

Configuring the Cluster Service

Post-Installation Configuration

Test Installation

Appendix (Guide to Creating and Configuring a Server Cluster under Windows Server 2003 White Paper)

Related Links (Guide to Creating and Configuring a Server Cluster under Windows Server 2003 White Paper)

# Introduction (Guide to Creating and Configuring a Server Cluster under Windows Server 2003 White Paper)

Updated: November 11, 2003

Applies To: Windows Server 2003 with SP1

A server cluster is a group of independent servers working collectively and running the Microsoft Cluster service (MSCS). Server clusters provide high availability, failback, scalability, and manageability for resources and applications

Server clusters allow client access to applications and resources in the event of failures and planned outages. If one of the servers in the cluster is unavailable because of a failure or maintenance requirements, resources and applications move to other available cluster nodes.

For Windows Clustering solutions, the term “high availability” is used rather than “fault tolerant.” Fault-tolerant technology offers a higher level of resilience and recovery. Fault-tolerant servers typically use a high degree of hardware redundancy plus specialized software to provide near-instantaneous recovery from any single hardware or software fault. These solutions cost significantly more than a Windows Clustering solution because organizations must pay for redundant hardware that waits in an idle state for a fault.

Server clusters do not guarantee non-stop operation, but they do provide sufficient availability for most mission-critical applications. The cluster service can monitor applications and resources and automatically recognize and recover from many failure conditions. This provides flexibility in managing the workload within a cluster. It also improves overall system availability.

Cluster service benefits include:

- **High Availability**: With server clusters, ownership of resources such as disk drives and Internet protocol (IP) addresses is automatically transferred from a failed server to a surviving server. When a system or application in the cluster fails, the cluster software restarts the failed application on a surviving server, or disperses the work from the failed node to the remaining nodes. As a result, users experience only a momentary pause in service.

- **Failback**: The Cluster service will automatically re-assign the workload in a cluster when a failed server comes back online to its predetermined preferred owner. This feature can be configured, but is disabled by default.

- **Manageability**: You can use the Cluster Administrator tool (CluAdmin.exe) to manage a cluster as a single system and to manage applications as if they were running on a single server. You can move applications to different servers within the cluster. Cluster Administrator can be used to manually balance server workloads and to free servers for planned maintenance. You can also monitor the status of the cluster, all nodes, and resources from anywhere on the network.

- **Scalability**: Cluster services can grow to meet increased demand. When the overall load for a cluster-aware application exceeds the cluster’s capabilities, additional nodes can be added.

This document provides instructions for creating and configuring a server cluster with servers connected to a shared cluster storage device and running Windows Server 2003 Enterprise Edition or Windows Server 2003 Datacenter Edition. Intended to guide you through the process of installing a typical cluster, this document does not explain how to install clustered applications. Windows Clustering solutions that implement non-traditional quorum models, such as Majority Node Set (MNS) clusters and geographically dispersed clusters, also are not discussed. For additional information about server cluster concepts as well as installation and configuration procedures, see the Windows Server 2003 Online Help.

# Checklists for Server Cluster Configuration

Updated: November 11, 2003

Applies To: Windows Server 2003 with SP1

This checklist helps you prepare for installation. Step-by-step instructions begin after the checklist.

**Software Requirements**

- Microsoft Windows Server 2003 Enterprise Edition or Windows Server 2003 Datacenter Edition installed on all computers in the cluster.
- A name resolution method such as Domain Name System (DNS), DNS dynamic update protocol, Windows Internet Name Service (WINS), HOSTS, and so on.
- An existing domain model.
- All nodes must be members of the same domain.
- A domain-level account that is a member of the local administrators group on each node. A dedicated account is recommended.

**Hardware Requirements**

- Clustering hardware must be on the cluster service Hardware Compatibility List (HCL). To find the latest version of the cluster service HCL, go to the Windows Hardware Compatibility List at http://www.microsoft.com/whdc/hcl/default.mspx, and then search for cluster. The entire solution must be certified on the HCL, not just the individual components. For additional information, see the following article in the Microsoft Knowledge Base:

  309395 The Microsoft Support Policy for Server Clusters and the Hardware

  | **Note** |
  | --- |
  | If you are installing this cluster on a storage area network (SAN) and plan to have multiple devices and clusters sharing the SAN with a cluster, the solution must also be on the "Cluster/Multi-Cluster Device" Hardware Compatibility List. For additional information, see the following article in the Microsoft Knowledge Base: 304415 Support for Multiple Clusters Attached to the Same SAN Device |

- Two mass storage device controllers—Small Computer System Interface (SCSI) or Fibre Channel. A local system disk for the operating system (OS) to be installed on one controller. A separate peripheral component interconnect (PCI) storage controller for the shared disks.
- Two PCI network adapters on each node in the cluster.
- Storage cables to attach the shared storage device to all computers. Refer to the manufacturers instructions for configuring storage devices. See the appendix that accompanies this article for additional information about specific configuration needs when using SCSI or Fibre Channel.
- All hardware should be identical, slot for slot, card for card, BIOS, firmware revisions, and so on, for all nodes. This makes configuration easier and eliminates compatibility problems.

**Network Requirements**

- A unique NetBIOS name.
- Static IP addresses for all network interfaces on each node.

| Note |
| --- |
| Server Clustering does not support the use of IP addresses assigned from Dynamic Host Configuration Protocol (DHCP) servers. |

- Access to a domain controller. If the cluster service is unable to authenticate the user account used to start the service, it could cause the cluster to fail. It is recommended that you have a domain controller on the same local area network (LAN) as the cluster is on to ensure availability.

- Each node must have at least two network adapters—one for connection to the client public network and the other for the node-to-node private cluster network. A dedicated private network adapter is required for HCL certification.

- All nodes must have two physically independent LANs or virtual LANs for public and private communication.

- If you are using fault-tolerant network cards or network adapter teaming, verify that you are using the most recent firmware and drivers. Check with your network adapter manufacturer for cluster compatibility.

**Shared Disk Requirements:**

- An HCL-approved external disk storage unit connected to all computers. This will be used as the clustered shared disk. Some type of a hardware redundant array of independent disks (RAID) is recommended.

- All shared disks, including the quorum disk, must be physically attached to a shared bus.

| Note |
| --- |
| The requirement above does not hold true for Majority Node Set (MNS) clusters, which are not covered in this guide. |

- Shared disks must be on a different controller then the one used by the system drive.

- Creating multiple logical drives at the hardware level in the RAID configuration is recommended rather than using a single logical disk that is then divided into multiple partitions at the operating system level. This is different from the configuration commonly used for stand-alone servers. However, it enables you to have multiple disk resources and to do Active/Active configurations and manual load balancing across the nodes in the cluster.

- A dedicated disk with a minimum size of 50 megabytes (MB) to use as the quorum device. A partition of at least 500 MB is recommended for optimal NTFS file system performance.

- Verify that disks attached to the shared bus can be seen from all nodes. This can be checked at the host adapter setup level. Refer to the manufacturer's documentation for adapter-specific instructions.

- SCSI devices must be assigned unique SCSI identification numbers and properly terminated according to the manufacturer's instructions. See the appendix with this article for information about installing and terminating SCSI devices.

- All shared disks must be configured as basic disks. For additional information, see the following article in the Microsoft Knowledge Base:

  237853 Dynamic Disk Configuration Unavailable for Server Cluster Disk Resources

- Software fault tolerance is not natively supported on cluster shared disks.

- All shared disks must be configured as master boot record (MBR) disks on systems running the 64-bit versions of Windows Server 2003.

- All partitions on the clustered disks must be formatted as NTFS.
- Hardware fault-tolerant RAID configurations are recommended for all disks.
- A minimum of two logical shared drives is recommended.

# Cluster Installation

Updated: November 11, 2003

Applies To: Windows Server 2003 with SP1

**Installation Overview**
During the installation process, some nodes will be shut down while others are being installed. This step helps guarantee that data on disks attached to the shared bus is not lost or corrupted. This can happen when multiple nodes simultaneously try to write to a disk that is not protected by the cluster software. The default behavior of how new disks are mounted has been changed in Windows 2003 Server from the behavior in the Microsoft® Windows® 2000 operating system. In Windows 2003, logical disks that are not on the same bus as the boot partition will not be automatically mounted and assigned a drive letter. This helps ensure that the server will not mount drives that could possibly belong to another server in a complex SAN environment. Although the drives will not be mounted, it is still recommended that you follow the procedures below to be certain the shared disks will not become corrupted.

Use the table below to determine which nodes and storage devices should be turned on during each step.

The steps in this guide are for a two-node cluster. However, if you are installing a cluster with more than two nodes, the Node 2 column lists the required state of all other nodes.

| **Step** | **Node 1** | **Node 2** | **Storage** | **Comments** |
|---|---|---|---|---|
| Setting up networks | On | On | Off | Verify that all storage devices on the shared bus are turned off. Turn on all nodes. |
| Setting up shared disks | On | Off | On | Shutdown all nodes. Turn on the shared storage, then turn on the first node. |
| Verifying disk configuration | Off | On | On | Turn on the first node, turn on second node. Repeat for nodes 3 and 4 if necessary. |
| Configuring the first node | On | Off | On | Turn off all nodes; turn on the first node. |
| Configuring the second node | On | On | On | Turn on the second node after the first node is successfully configured. Repeat for nodes 3 and 4 as necessary. |
| Post-installation | On | On | On | All nodes should be on. |

Several steps must be taken before configuring the Cluster service software. These steps are:

- Installing Windows Server 2003 Enterprise Edition or Windows Server 2003 Datacenter Edition operating system on each node.
- Setting up networks.
- Setting up disks.

Perform these steps on each cluster node before proceeding with the installation of cluster service on the first node.

To configure the cluster service, you must be logged on with an account that has administrative permissions to all nodes. Each node must be a member of the same domain. If you choose to make one of the nodes a domain controller, have another domain controller available on the same subnet to eliminate a single point of failure and enable maintenance on that node.

**Installing the Windows Server 2003 Operating System**
Refer to the documentation you received with the Windows Server 2003 operating system package to install the system on each node in the cluster.

Before configuring the cluster service, you must be logged on locally with a domain account that is a member of the local administrators group.

**Note**

The installation will fail if you attempt to join a node to a cluster that has a blank password for the local administrator account. For security reasons, Windows Server 2003 prohibits blank administrator passwords.

**Setting Up Networks**

Each cluster node requires at least two network adapters with two or more independent networks, to avoid a single point of failure. One is to connect to a public network, and one is to connect to a private network consisting of cluster nodes only. Servers with multiple network adapters are referred to as "multi-homed." Because multi-homed servers can be problematic, it is critical that you follow the network configuration recommendations outlined in this document.

Microsoft requires that you have two Peripheral Component Interconnect (PCI) network adapters in each node to be certified on the Hardware Compatibility List (HCL) and supported by Microsoft Product Support Services. Configure one of the network adapters on your production network with a static IP address, and configure the other network adapter on a separate network with another static IP address on a different subnet for private cluster communication.

Communication between server cluster nodes is critical for smooth cluster operations. Therefore, you must configure the networks that you use for cluster communication are configured optimally and follow all hardware compatibility list requirements.

The private network adapter is used for node-to-node communication, cluster status information, and cluster management. Each node's public network adapter connects the cluster to the public network where clients reside and should be configured as a backup route for internal cluster communication. To do so, configure the roles of these networks as either "**Internal Cluster Communications Only**" or "**All Communications**" for the Cluster service.

Additionally, each cluster network must fail independently of all other cluster networks. This means that two cluster networks must *not* have a component in common that can cause both to fail simultaneously. For example, the use of a multiport network adapter to attach a node to two cluster networks would not satisfy this requirement in most cases because the ports are not independent.

To eliminate possible communication issues, remove all unnecessary network traffic from the network adapter that is set to **Internal Cluster communications only** (this adapter is also known as the heartbeat or private network adapter).

To verify that all network connections are correct, private network adapters must be on a network that is on a different logical network from the public adapters. This can be accomplished by using a cross-over cable in a two-node configuration or a dedicated hub (not a smart hub) in a configuration of more than two nodes. Do not use a switch, smart hub, or any other routing device for the heartbeat network.

**Note**

Cluster heartbeats cannot be forwarded through a routing device because their Time to Live (TTL) is set to 1. The public network adapters must be only connected to the public network. If you have a virtual LAN, then the latency between the nodes must be less then 500 milliseconds (ms). Also, in Windows Server 2003, heartbeats in Server Clustering have been changed to multicast; therefore, you may want to make a Madcap server available to assign the multicast addresses. For additional information, see the following article in the Microsoft Knowledge Base: 307962 Multicast Support Enabled for the Cluster Heartbeat.

Figure 1 below outlines a four-node cluster configuration.




**Figure 1. Connections for a four-node cluster.**

**General Network Configuration**

| **Note** |
| --- |
| This guide assumes that you are running the default **Start** menu. The steps may be slightly different if you are running the Classic **Start** menu. Also, which network adapter is private and which is public depends upon your wiring. For the purposes of this white paper, the first network adapter (Local Area Connection) is connected to the public network, and the second network adapter (Local Area Connection 2) is connected to the private cluster network. Your network may be different. |

**To rename the local area network icons**

It is recommended that you change the names of the network connections for clarity. For example, you might want to change the name of *Local Area Connection 2* to something like *Private*. Renaming will help you identify a network and correctly assign its role.

1. Click **Start**, point to **Control Panel**, right-click **Network Connections**, and then click **Open**

2. Right-click the **Local Area Connection 2** icon.

3. Click **Rename**.

4. Type **Private** in the textbox, and then press ENTER.

5. Repeat steps 1 through 3, and then rename the public network adapter as **Public**.




**Figure 2. Renamed icons in the Network Connections window.**

6. The renamed icons should look like those in Figure 2 above. Close the Network Connections window. The new connection names will appear in Cluster Administrator and automatically replicate to all other cluster nodes as they are brought online.

**To configure the binding order networks on all nodes**

1. Click **Start**, point to **Control Panel**, right-click **Network Connections**, and then click **Open**

2. On the **Advanced** menu, click **Advanced Settings**.

3. In the **Connections** box, make sure that your bindings are in the following order, and then click **OK**:

    a. Public

    b. Private

    c. Remote Access Connections

**Configuring the Private Network Adapter**

1. Right-click the network connection for your heartbeat adapter, and then click **Properties**.

2. On the **General** tab, make sure that only the **Internet Protocol (TCP/IP)** check box is selected, as shown in Figure 3 below. Click to clear the check boxes for all other clients, services, and protocols.




**Figure 3. Click to select only the Internet Protocol check box in the Private Properties dialog box.**

3. If you have a network adapter that is capable of transmitting at multiple speeds, you should manually specify a speed and duplex mode. Do not use an auto-select setting for speed, because some adapters may drop packets while determining the speed. The speed for the network adapters must be hard set (manually set) to be the same on all nodes according to the card manufacturer's specification. If you are not sure of the supported speed of your card and connecting devices, Microsoft recommends you set all devices on that path to **10 megabits per second** (Mbps) and **Half Duplex**, as shown in Figure 4 below. The amount of information that is traveling across the heartbeat network is small, but latency is critical for communication. This configuration will provide enough bandwidth for reliable communication. All network adapters in a cluster attached to the same network must be configured identically to use the same **Duplex Mode**, **Link Speed, Flow Control**, and so on. Contact your adapter's manufacturer for specific information about appropriate speed and duplex settings for your network adapters.



**Figure 4. Setting the speed and duplex for all adapters**

| **Note** |
| --- |
| Microsoft does not recommended that you use any type of fault-tolerant adapter or "Teaming" for the heartbeat. If you require redundancy for your heartbeat connection, use multiple network adapters set to Internal Communication Only and define their network priority in the Cluster configuration. Issues seen with early multi-ported network adapters, verify that your firmware and driver are at the most current revision if you use this technology. Contact your network adapter manufacturer for information about compatibility on a server cluster. For more information, see the following article in the Microsoft Knowledge Base: 254101 Network Adapter Teaming and Server Clustering. |

4. Click **Internet Protocol (TCP/IP)**, and then click **Properties**.

5. On the **General** tab, verify that you have selected a static IP address that is not on the same subnet or network as any other public network adapter. It is recommended that you put the private network adapter in one of the following private network ranges:

   - 10.0.0.0 through 10.255.255.255 (Class A)
   - 172.16.0.0 through 172.31.255.255 (Class B)
   - 192.168.0.0 through 192.168.255.255 (Class C)

   An example of a good IP address to use for the private adapters is 10.10.10.10 on node 1 and 10.10.10.11 on node 2 with a subnet mask of 255.0.0.0, as shown in Figure 5 below. Be sure that this is a completely different IP address scheme then the one used for the public network.

| **Note** |
| --- |
| For more information about valid IP addressing for a private network, see the following article in the Microsoft Knowledge Base: 142863 Valid IP Addressing for a Private Network. |




**Figure 5. An example of an IP address to use for private adapters.**

6. Verify that there are no values defined in the **Default Gateway** box or under **Use the Following DNS server addresses**.

7. Click the **Advanced** button.

8. On the **DNS** tab, verify that no values are defined. Make sure that the **Register this connection's addresses in DNS** and **Use this connection's DNS suffix in DNS registration** check boxes are cleared.

9. On the **WINS** tab, verify that there are no values defined. Click **Disable NetBIOS over TCP/IP** as shown in Figure 6 on the next page.




**Figure 6. Verify that no values are defined on the WINS tab.**

10. When you close the dialog box, you may receive the following prompt: "*This connection has an empty primary WINS address. Do you want to continue?*" If you receive this prompt, click **Yes**

11. Complete steps 1 through 10 on all other nodes in the cluster with different static IP addresses.

**Configuring the Public Network Adapter**

| **Note** |
| --- |
| If IP addresses are obtained via DHCP, access to cluster nodes may be unavailable if the DHCP server is inaccessible. For this reason, static IP addresses are required for all interfaces on a server cluster. Keep in mind that cluster service will only recognize one network interface per subnet. If you need assistance with TCP/IP addressing in Windows Server 2003, please see the Online Help. |

**Verifying Connectivity and Name Resolution**

To verify that the private and public networks are communicating properly, ping all IP addresses from each node. You should be able to ping all IP addresses, locally and on the remote nodes.

To verify name resolution, ping each node from a client using the node's machine name instead of its IP address. It should only return the IP address for the public network. You may also want to try a **PING –a** command to do a reverse lookup on the IP Addresses.

**Verifying Domain Membership**
All nodes in the cluster must be members of the same domain and be able to access a domain controller and a DNS server. They can be configured as member servers or domain controllers. You should have at least one domain controller on the same network segment as the cluster. For high availability. another domain controller should also be available to remove a single point of failure. In this guide, all nodes are configured as member servers.

There are instances where the nodes may be deployed in an environment where there are no pre-existing Microsoft® Windows NT® 4.0 domain controllers or Windows Server 2003 domain controllers. This scenario requires at least one of the cluster nodes to be configured as a domain controller. However, in a two-node server cluster, if one node is a domain controller, then the other node also must be a domain controller. In a four-node cluster implementation, it is not necessary to configure all four nodes as domain controllers. However, when following a "best practices" model and having at least one backup domain controller, at least one of the remaining three nodes should be configured as a domain controller. A cluster node must be promoted to a domain controller by using the DCPromo tool before the cluster service is configured.

The dependence in Windows Server 2003 on the DNS further requires that every node that is a domain controller also must be a DNS server *if* another DNS server that supports dynamic updates and/or SRV records is not available (Active directory integrated zones recommended).

The following issues should be considered when deploying cluster nodes as domain controllers:

- If one cluster node in a two-node cluster is a domain controller, the other node must be a domain controller
- There is overhead associated with running a domain controller. An idle domain controller can use anywhere between 130 and 140 MB of RAM, which includes having the Cluster service running. There is also increased network traffic from replication, because these domain controllers have to replicate with other domain controllers in the domain and across domains.
- If the cluster nodes are the only domain controllers, then each must be a DNS server as well. They should point to each other for primary DNS resolution and to themselves for secondary resolution.
- The first domain controller in the forest/domain will take on all Operations Master Roles. You can redistribute these roles to any node. However, if a node fails, the Operations Master Roles assumed by that node will be unavailable. Therefore, it is recommended that you do not run Operations Master Roles on any cluster node. This includes Scheme Master, Domain Naming Master, Relative ID Master, PDC Emulator, and Infrastructure Master. These functions cannot be clustered for high availability with failover.
- Clustering other applications such as Microsoft® SQL Server ™ or Microsoft® Exchange Server in a scenario where the nodes are also domain controllers may not be optimal due to resource constraints. This configuration should be thoroughly tested in a lab environment before deployment

Because of the complexity and overhead involved in making cluster-nodes domain controllers, it is recommended that all nodes should be member servers.

**Setting Up a Cluster User Account**
The Cluster service requires a domain user account that is a member of the Local Administrators group on each node, under which the Cluster service can run. Because setup requires a user name and password, this user account must be created before configuring the Cluster service. This user account should be dedicated only to running the Cluster service, and should not belong to an individual.

| **Note** |
| --- |
| The cluster service account does not need to be a member of the Domain Administrators group. For security reasons, granting domain administrator rights to the cluster service account is not recommended. |

The cluster service account requires the following rights to function properly on all nodes in the cluster. The Cluster

Configuration Wizard grants the following rights automatically:

- Act as part of the operating system
- Adjust memory quotas for a process
- Back up files and directories
- Increase scheduling priority
- Log on as a service
- Restore files and directories

For additional information, see the following article in the Microsoft Knowledge Base:

269229 How to Manually Re-Create the Cluster Service Account

**To set up a cluster user account**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
2. Click the plus sign (+) to expand the domain if it is not already expanded.
3. Right-click **Users**, point to **New**, and then click **User**.
4. Type the cluster name, as shown in Figure 7 below, and then click **Next**.



**Figure 7. Type the cluster name.**

5. Set the password settings to **User Cannot Change Password** and **Password Never Expires**. Click **Next,** and then click **Finish** to create this user.

| **Note** |
| --- |
| If your administrative security policy does not allow the use of passwords that never expire, you must renew the password and update the cluster service configuration on each node before password expiration. For additional information, see the following article in the Microsoft Knowledge Base: 305813 How to Change the Cluster Service Account Password. |

6. Right-click **Cluster** in the left pane of the Active Directory Users and Computers snap-in, and then click **Properties** on the shortcut menu.
7. Click **Add Members to a Group**.

8. Click **Administrators,** and then click **OK**. This gives the new user account administrative privileges on this computer.

9. Quit the Active Directory Users and Computers snap-in.

**Setting up Shared Disks**

| **Important** |
| --- |
| To avoid corrupting the cluster disks, make sure that Windows Server 2003 and the Cluster service are installed, configured, and running on at least one node before you start an operating system on another node. It is critical to never have more then one node on until the Cluster service is configured.<br><br>To proceed, turn off all nodes. Turn on the shared storage devices, and then turn on node 1. |

**About the Quorum Disk**

The quorum disk is used to store cluster configuration database checkpoints and log files that help manage the cluster and maintain consistency. The following quorum disk procedures are recommended:

- Create a logical drive with a minimum size of 50 MB to be used as a quorum disk, 500 MB is optimal for NTFS.

- Dedicate a separate disk as a quorum resource.

| **Important** |
| --- |
| A quorum disk failure could cause the entire cluster to fail; therefore, it is strongly recommended that you use a volume on a hardware RAID array. Do not use the quorum disk for anything other than cluster management. |

The quorum resource plays a crucial role in the operation of the cluster. In every cluster, a single resource is designated as the quorum resource. A quorum resource can be any Physical Disk resource with the following functionality:

- It replicates the cluster registry to all other nodes in the server cluster. By default, the cluster registry is stored in the following location on each node: *%SystemRoot%\Cluster\Clusdb*. The cluster registry is then replicated to the *MSCS\Chkxxx.tmp* file on the quorum drive. These files are exact copies of each other. The *MSCS\Quolog.log* file is a transaction log that maintains a record of all changes to the checkpoint file. This means that nodes that were offline can have these changes appended when they rejoin the cluster.

- If there is a loss of communication between cluster nodes, the challenge response protocol is initiated to prevent a "split brain" scenario. In this situation, the owner of the quorum disk resource becomes the only owner of the cluster and all the resources. The owner then makes the resources available for clients. When the node that owns the quorum disk functions incorrectly, the surviving nodes arbitrate to take ownership of the device. For additional information, see the following article in the Microsoft Knowledge Base: 309186 How the Cluster Service Takes Ownership of a Disk on the Shared Bus.

During the cluster service installation, you must provide the drive letter for the quorum disk. The letter **Q** is commonly used as a standard, and Q is used in the example.

**To configure shared disks**

1. Make sure that only one node is turned on.

2. Right click **My Computer**, click **Manage**, and then expand **Storage**.

3. Double-click **Disk Management**.

4. If you connect a new drive, then it automatically starts the Write Signature and Upgrade Disk Wizard. If this happens, click **Next** to step through the wizard.

> **Note**
> The wizard automatically sets the disk to dynamic. To reset the disk to basic, right-click **Disk** *n* (where *n* specifies the disk that you are working with), and then click **Revert to Basic Disk**.

5. Right-click **unallocated disk space**.

6. Click **New Partition**.

7. The New Partition Wizard begins. Click **Next**.

8. Select the **Primary Partition** partition type. Click **Next**.

9. The default is set to maximum size for the partition size. Click **Next**. (Multiple logical disks are recommended over multiple partitions on one disk.)

10. Use the drop-down box to change the drive letter. Use a drive letter that is farther down the alphabet than the default enumerated letters. Commonly, the drive letter **Q** is used for the quorum disk, then R, S,and so on for the data disks. For additional information, see the following article in the Microsoft Knowledge Base:

    318534 Best Practices for Drive-Letter Assignments on a Server Cluster

> **Note**
> If you are planning on using volume mount points, do not assign a drive letter to the disk. For additional information, see the following article in the Microsoft Knowledge Base: 280297 How to Configure Volume Mount Points on a Clustered Server.

11. Format the partition using NTFS. In the **Volume Label** box, type a name for the disk. For example, **Drive Q**, as shown in Figure 8 below. It is critical to assign drive labels for shared disks, because this can dramatically reduce troubleshooting time in the event of a disk recovery situation.




**Figure 8. It is critical to assign drive labels for shared disks.**

If you are installing a 64-bit version of Windows Server 2003, verify that all disks are formatted as MBR. Global Partition Table (GPT) disks are not supported as clustered disks. For additional information, see the following article in the Microsoft Knowledge Base:

284134 Server Clusters Do Not Support GPT Shared Disks

Verify that all shared disks are formatted as NTFS and designated as MBR Basic.

**To verify disk access and functionality**

1. Start Windows Explorer.

2. Right-click one of the shared disks (such as **Drive Q:**), click **New,** and then click **Text Document**.

3. Verify that you can successfully write to the disk and that the file was created.

4. Select the file, and then press the **Del** key to delete it from the clustered disk.

5. Repeat steps 1 through 4 for all clustered disks to verify they can be correctly accessed from the first node.

6. Turn off the first node, turn on the second node, and repeat steps 1 through 4 to verify disk access and functionality. Assign drive letters to match the corresponding drive labels. Repeat again for any additional nodes. Verify that all nodes can read and write from the disks, turn off all nodes except the first one, and then continue with this white paper.

# Configuring the Cluster Service

Updated: November 11, 2003

Applies To: Windows Server 2003 with SP1

You must supply all initial cluster configuration information in the first installation phase. This is accomplished using the Cluster Configuration Wizard.

As seen in the flow chart, the form (Create a new Cluster) and the Join (Add nodes) take a couple different paths, but they have a few of the same pages. Namely, Credential Login, Analyze, and Re-Analyze and Start Service are the same. There are minor differences in the following pages: Welcome, Select Computer, and Cluster Service Account. In the next two sections of this lesson, you will step through the wizard pages presented on each of these configuration paths. In the third section, after you follow the step-through sections, this white paper describes in detail the Analyze, ,Re-Analyze and Start Service pages, and what the information provided in these screens means.




| Note |
| --- |
| During Cluster service configuration on node 1, you must turn off all other nodes. All shared storage devices should be turned on. |

**To configure the first node**

1. Click **Start**, click **All Programs**, click **Administrative Tools**, and then click **Cluster Administrator**.

2. When prompted by the Open Connection to Cluster Wizard, click **Create new cluster** in the **Action** drop-down list, as shown in Figure 9 below.

**Figure 9. The Action drop-down list.**

3. Verify that you have the necessary prerequisites to configure the cluster, as shown in Figure 10 below. Click **Next**.

**Figure 10. A list of prerequisites is part of the New Server Cluster Wizard Welcome page.**

4. Type a unique NetBIOS name for the cluster (up to 15 characters), and then click **Next**. In the example shown in Figure 11 below, the cluster is named **MyCluster**.) Adherence to DNS naming rules is recommended. For additional information, see the following articles in the Microsoft Knowledge Base:

   163409 NetBIOS Suffixes (16th Character of the NetBIOS Name)

   254680 DNS Namespace Planning




**Figure 11. Adherence to DNS naming rules is recommended when naming the cluster.**

5. If you are logged on locally with an account that is not a **Domain Account with Local Administrative privileges**, the wizard will prompt you to specify an account. This is not the account the Cluster service will use to start.

| **Note** |
| --- |
| If you have appropriate credentials, the prompt mentioned in step 5 and shown in Figure 12 below may not appear. |




**Figure 12. The New Server Cluster Wizard prompts you to specify an account.**

6. Because it is possible to configure clusters remotely, you must verify or type the name of the server that is going to be used as the first node to create the cluster, as shown in Figure 13 below. Click **Next**.




**Figure 13. Select the name of the computer that will be the first node in the cluster.**

**Note**

The Install wizard verifies that all nodes can see the shared disks the same. In a complex storage area network the target identifiers (TIDs) for the disks may sometimes be different, and the Setup program may incorrectly detect that the disk configuration is not valid for Setup. To work around this issue you can click the **Advanced** button, and then click **Advanced (minimum) configuration**. For additional information, see the following article in the Microsoft Knowledge Base: 331801 Cluster Setup May Not Work When You Add Nodes




7. Figure 14 below illustrates that the Setup process will now analyze the node for possible hardware or software problems that may cause problems with the installation. Review any warnings or error messages. You can also click the **Details** button to get detailed information about each one.




**Figure 14. The Setup process analyzes the node for possible hardware or software problems.**

8. Type the unique cluster IP address (in this example 172.26.204.10), and then click **Next**.

   As shown in Figure 15 below, the New Server Cluster Wizard automatically associates the cluster IP address with one of the public networks by using the subnet mask to select the correct network. The cluster IP address should be used for administrative purposes only, and not for client connections.




**Figure 15. The New Server Cluster Wizard automatically associates the cluster IP address with one of the public networks.**

9. Type the **user name** and **password** of the cluster service account that was created during pre-installation. (In the example in Figure 16 below, the user name is "Cluster"). Select the domain name in the **Domain** drop-down list, and then click **Next**.

   At this point, the **Cluster Configuration** Wizard validates the user account and password.




**Figure 16. The wizard prompts you to provide the account that was created during pre-installation.**

10. Review the **Summary** page, shown in Figure 17 below, to verify that all the information that is about to be used to create the cluster is correct. If desired, you can use the quorum button to change the quorum disk designation from the default auto-selected disk.

    The summary information displayed on this screen can be used to reconfigure the cluster in the event of a disaster recovery situation. It is recommended that you save and print a hard copy to keep with the change management log at the server.

| **Note** |
| --- |
| The **Quorum** button can also be used to specify a Majority Node Set (MNS) quorum model. This is one of the major configuration differences when you create an MNS cluster |




**Figure 17. The Proposed Cluster Configuration page.**

11. Review any warnings or errors encountered during cluster creation. To do this, click the plus signs to see more, and then click **Next**. Warnings and errors appear in the Creating the Cluster page as shown in Figure 18.




**Figure 18. Warnings and errors appear on the Creating the Cluster page.**

12. Click **Finish** to complete the installation. Figure 19 below illustrates the final step.

**Figure 19. The final step in setting up a new server cluster.**

| **Note** |
| --- |
| To view a detailed summary, click the **View Log** button or view the text file stored in the following location:<br><br>%SystemRoot%\System32\LogFiles\Cluster\ClCfgSrv.Log |

**Validating the Cluster Installation**

Use the Cluster Administrator (CluAdmin.exe) to validate the cluster service installation on node 1.

**To validate the cluster installation**

1. Click **Start**, click **Programs**, click **Administrative Tools**, and then click **Cluster Administrator**.

2. Verify that all resources came online successfully, as shown in Figure 20 below.




**Figure 20. The Cluster Administer verifies that all resources came online successfully.**

| **Note** |
| --- |
| As general rules, do not put anything in the cluster group, do not take anything out of the cluster group, and do not use anything in the cluster group for anything other than cluster administration. |

**Configuring the Second Node**

Installing the cluster service on the other nodes requires less time than on the first node. Setup configures the cluster service network settings on the second node based on the configuration of the first node. You can also add multiple nodes to the cluster at the same time, and remotely.

| **Note** |
| --- |
| For this section, leave node 1 and all shared disks turned on. Then turn on all other nodes. The cluster service will control access to the shared disks at this point to eliminate any chance of corrupting the volume. |

1. Open **Cluster Administrator** on Node 1.
2. Click **File**, click **New**, and then click **Node**.
3. The Add Cluster Computers Wizard will start. Click **Next**.
4. If you are not logged on with appropriate credentials, you will be asked to specify a domain account that has administrative rights over all nodes in the cluster.
5. Enter the machine name for the node you want to add to the cluster. Click **Add**. Repeat this step, shown in Figure 21 below, to add all other nodes that you want. When you have added all nodes, click **Next**.




**Figure 21. Adding nodes to the cluster.**

6. The Setup wizard will perform an analysis of all the nodes to verify that they are configured properly.

7. Type the password for the account used to start the cluster service.

8. Review the summary information that is displayed for accuracy. The summary information will be used to configure the other nodes when they join the cluster.

9. Review any warnings or errors encountered during cluster creation, and then click **Next**.

10. Click **Finish** to complete the installation.

# Post-Installation Configuration

Updated: November 11, 2003

Applies To: Windows Server 2003 with SP1

**Heartbeat Configuration**

Now that the networks have been configured correctly on each node and the Cluster service has been configured, you need to configure the network roles to define their functionality within the cluster. Here is a list of the network configuration options in Cluster Administrator:

- **Enable for cluster use**: If this check box is selected, the cluster service uses this network. This check box is selected by default for all networks.
- **Client access only (public network)**: Select this option if you want the cluster service to use this network adapter only for external communication with other clients. No node-to-node communication will take place on this network adapter.
- **Internal cluster communications only (private network)**: Select this option if you want the cluster service to use this network only for node-to-node communication.
- **All communications (mixed network)**: Select this option if you want the cluster service to use the network adapter for node-to-node communication and for communication with external clients. This option is selected by default for all networks.

This white paper assumes that only two networks are in use. It explains how to configure these networks as one mixed network and one private network. This is the most common configuration. If you have available resources, two dedicated redundant networks for internal-only cluster communication are recommended.

**To configure the heartbeat**

1. Start Cluster Administrator.
2. In the left pane, click **Cluster Configuration,** click **Networks**, right-click **Private,** and then click **Properties**.
3. Click **Internal cluster communications only (private network)**, as shown in Figure 22 on the next page.




**Figure 22. Using Cluster Administrator to configure the heartbeat.**

4. Click **OK**.
5. Right-click **Public**, and then click **Properties** (shown in Figure 23 below).
6. Click to select the **Enable this network for cluster use** check box.

7. Click the **All communications (mixed network)** option, and then click **OK**.




**Figure 23. The Public Properties dialog box.**

**Heartbeat Adapter Prioritization**

After configuring the role of how the cluster service will use the network adapters, the next step is to prioritize the order in which they will be used for intra-cluster communication. This is applicable only if two or more networks were configured for node-to-node communication. Priority arrows on the right side of the screen specify the order in which the cluster service will use the network adapters for communication between nodes. The cluster service always attempts to use the first network adapter listed for remote procedure call (RPC) communication between the nodes. Cluster service uses the next network adapter in the list only if it cannot communicate by using the first network adapter.

1. Start Cluster Administrator.

2. In the left pane, right-click the cluster name (in the upper left corner), and then click **Properties**.

3. Click the **Network Priority** tab, as shown in Figure 24 below.




**Figure 24. The Network Priority tab in Cluster Administrator.**

4. Verify that the **Private** network is listed at the top. Use the **Move Up** or **Move Down** buttons to change the priority order.

5. Click **OK**.

**Configuring Cluster Disks**

- Start Cluster Administrator, right-click any disks that you want to remove from the cluster, and then click **Delete**.

**Note**

By default, all disks not residing on the same bus as the system disk will have Physical Disk Resources created for them, and will be clustered. Therefore, if the node has multiple buses, some disks may be listed that will not be used as shared storage, for example, an internal SCSI drive. Such disks should be removed from the cluster configuration. If you plan to implement Volume Mount points for some disks, you may want to delete the current disk resources for those disks, delete the drive letters, and then create a new disk resource without a drive letter assignment.

**Quorum Disk Configuration**

The Cluster Configuration Wizard automatically selects the drive that is to be used as the quorum device. It will use the smallest partition that is larger then 50 MB. You may want to change the automatically selected disk to a dedicated disk that you have designated for use as the quorum.

**To configure the quorum disk**

1. Start Cluster Administrator (CluAdmin.exe).

2. Right-click the cluster name in the upper-left corner, and then click **Properties**.

3. Click the **Quorum** tab.

4. In the **Quorum resource** list box, select a different disk resource. In Figure 25 below, Disk Q is selected in the **Quorum resource** list box.



**Figure 25. The Quorum resource list box.**

5. If the disk has more than one partition, click the partition where you want the cluster-specific data to be kept, and then click **OK**.

For additional information, see the following article in the Microsoft Knowledge Base:

Q280353 How to Change Quorum Disk Designation

**Creating a Boot Delay**

In a situation where all the cluster nodes boot up and attempt to attach to the quorum resource at the same time, the Cluster service may fail to start. For example, this may occur when power is restored to all nodes at the exact same time after a power failure. To avoid such a situation, increase or decrease the **Time to Display list of operating systems** setting. To find this setting, click **Start**, point to **My Computer**, right-click **My Computer**, and then click **Properties**. Click the **Advanced** tab, and then click **Settings** under **Startup And Recovery**.

# Test Installation

Updated: November 11, 2003

Applies To: Windows Server 2003 with SP1

There are several methods for verifying a cluster service installation after the Setup process is complete. These include:

- **Cluster Administrator:** If installation was completed only on node 1, start Cluster Administrator, and then attempt to connect to the cluster. If a second node was installed, start Cluster Administrator on either node, connect to the cluster, and then verify that the second node is listed.
- **Services Applet:** Use the services snap-in to verify that the cluster service is listed and started.
- **Event Log:** Use the Event Viewer to check for **ClusSvc** entries in the system log. You should see entries confirming that the cluster service successfully formed or joined a cluster.
- **Cluster service registry entries:** Verify that the cluster service installation process wrote the correct entries to the registry. You can find many of the registry settings under HKEY_LOCAL_MACHINE\Cluster
- Click **Start**, click **Run**, and then type the Virtual Server name. Verify that you can connect and see resources.

**Test Failover**
**To verify that resources will failover**

1. Click **Start**, click **Programs**, click **Administrative Tools**, and then click **Cluster Administrator**, as shown in Figure 26 below.



**Figure 26. The Cluster Administrator window.**

2. Right-click the **Disk Group 1** group, and then click **Move Group**. The group and all its resources will be moved to another node. After a short period of time, the **Disk F: G:** will be brought online on the second node. Watch the window to see this shift. Quit Cluster Administrator.

Congratulations! You have completed the configuration of the cluster service on all nodes. The server cluster is fully operational. You are now ready to install cluster resources such as file shares, printer spoolers, cluster aware services like Distributed Transaction Coordinator, DHCP, WINS, or cluster-aware programs such as Exchange Server or SQL Server.

# Appendix (Guide to Creating and Configuring a Server Cluster under Windows Server 2003 White Paper)

Updated: November 11, 2003

Applies To: Windows Server 2003 with SP1

**Advanced Testing**

Now that you have configured your cluster and verified basic functionality and failover, you may want to conduct a series of failure scenario tests that will demonstrate expected results and ensure the cluster will respond correctly when a failure occurs. This level of testing is not required for every implementation, but may be insightful if you are new to clustering technology and are unfamiliar how the cluster will respond or if you are implementing a new hardware platform in your environment. The expected results listed are for a clean configuration of the cluster with default settings, this does not take into consideration any user customization of the failover logic. This is not a complete list of all tests, nor should successfully completing these tests be considered "certified" or ready for production. This is simply a sample list of some tests that can be conducted. For additional information, see the following article in the Microsoft Knowledge Base:

197047 Failover/Failback Policies on Microsoft Cluster Server

**Test:** Start Cluster Administrator, right-click a resource, and then click "**Initiate Failure**". The resource should go into an failed state, and then it will be restarted and brought back into an online state on that node.

*Expected Result:* Resources should come back online on the same node

**Test:** Conduct the above "**Initiate Failure**" test three more times on that same resource. On the fourth failure, the resources should all failover to another node in the cluster.

*Expected Result:* Resources should failover to another node in the cluster

**Test:** Move all resources to one node. Start Computer Management, and then click **Services** under **Services and Applications**. Stop the Cluster service. Start Cluster Administrator on another node and verify that all resources failover and come online on another node correctly.

*Expected Result:* Resources should failover to another node in the cluster

**Test:** Move all resources to one node. On that node, click **Start,** and then click **Shutdown**. This will turn off that node. Start Cluster Administrator on another node, and then verify that all resources failover and come online on another node correctly.

*Expected Result:* Resources should failover to another node in the cluster

**Test:** Move all resources to one node, and then press the power button on the front of that server to turn it off. If you have an ACPI compliant server, the server will perform an "Emergency Shutdown" and turn off the server. Start Cluster Administrator on another node and verify that all resources failover and come online on another node correctly. For additional information about an Emergency Shutdown, see the following articles in the Microsoft Knowledge Base:

325343 HOW TO: Perform an Emergency Shutdown in Windows Server 2003

297150 Power Button on ACPI Computer May Force an Emergency Shutdown

*Expected Result:* Resources should failover to another node in the cluster

| **Important** |
| --- |
| Performing the Emergency Shutdown test may cause data corruption and data loss. Do **not** conduct this test on a production server |

**Test:** Move all resources to one node, and then pull the power cables from that server to simulate a hard failure. Start Cluster Administrator on another node, and then verify that all resources failover and come online on another node correctly

*Expected Result:* Resources should failover to another node in the cluster

| **Important** |
| --- |

Performing the hard failure test may cause data corruption and data loss. This is an extreme test. Make sure you have a backup of all critical data, and then conduct the test at your own risk. Do **not** conduct this test on a production server

**Test:** Move all resources to one node, and then remove the public network cable from that node. The IP Address resources should fail, and the groups will all failover to another node in the cluster. For additional information, see the following articles in the Microsoft Knowledge Base:

286342 Network Failure Detection and Recovery in Windows Server 2003 Clusters

*Expected Result:* Resources should failover to another node in the cluster

**Test:** Remove the network cable for the Private heartbeat network. The heartbeat traffic will failover to the public network, and no failover should occur. If failover occurs, please see the "Configuring the Private Network Adapter" section in earlier in this document

*Expected Result:* There should be no failures or resource failovers

**SCSI Drive Installations**

This appendix is provided as a generic set of instructions for SCSI drive installations. If the SCSI hard disk vendor's instructions conflict with the instructions here, always follow the instructions supplied by the vendor.

The SCSI bus listed in the hardware requirements must be configured prior to cluster service installation. Configuration applies to:

- The SCSI devices.
- The SCSI controllers and the hard disks so that they work properly on a shared SCSI bus.
- Properly terminating the bus. The shared SCSI bus must have a terminator at each end of the bus. It is possible to have multiple shared SCSI buses between the nodes of a cluster.

In addition to the information on the following pages, refer to documentation from the manufacturer of your SCSI device or to the SCSI specifications, which can be ordered from the American National Standards Institute (ANSI). The ANSI Web site includes a catalog that can be searched for the SCSI specifications.

**Configuring the SCSI Devices**

Each device on the shared SCSI bus must have a unique SCSI identification number. Because most SCSI controllers default to SCSI ID 7, configuring the shared SCSI bus includes changing the SCSI ID number on one controller to a different number, such as SCSI ID 6. If there is more than one disk that will be on the shared SCSI bus, each disk must have a unique SCSI ID number.

**Terminating the Shared SCSI Bus**

There are several methods for terminating the shared SCSI bus. They include:

- **SCSI controllers**

  SCSI controllers have internal soft termination that can be used to terminate the bus, however this method is not recommended with the cluster server. If a node is turned off with this configuration, the SCSI bus will be terminated improperly and will not operate correctly.

- **Storage enclosures**

  Storage enclosures also have internal termination, which can be used to terminate the SCSI bus if the enclosure is at the end of the SCSI bus. This should be turned off.

- **Y cables**

  Y cables can be connected to devices if the device is at the end of the SCSI bus. An external active terminator can then be attached to one branch of the Y cable in order to terminate the SCSI bus. This method of termination requires either disabling or removing any internal terminators that the device may have.

  Figure 27 outlines how a SCSI cluster should be physically connected.




**Figure 27. A diagram of a SCSI cluster hardware configuration.**

| **Note** |
| --- |
| Any devices that are not at the end of the shared bus must have their internal termination disabled. Y cables and active terminator connectors are the recommended termination methods because they will provide termination even when a node is not online. |

**Storage Area Network Considerations**

There are two supported methods of Fibre Channel-based storage in a Windows Server 2003 server cluster: arbitrated loops and switched fabric.

| **Important** |
| --- |
| When evaluating both types of Fibre Channel implementation, read the vendor's documentation and be sure you understand the specific features and restrictions of each. |

Although the term *Fibre Channel* implies the use of fiber-optic technology, copper coaxial cable is also allowed for interconnects.

**Arbitrated Loops (FC-AL)**

A Fibre Channel arbitrated loop (FC-AL) is a set of nodes and devices connected into a single loop. FC-AL provides a cost-effective way to connect up to 126 devices into a single network. As with SCSI, a maximum of two nodes is supported in an FC-AL server cluster configured with a hub. An FC-AL is illustrated in Figure 28.

**Figure 28 FC-AL Connection**

**Figure Error! Reference source not found.1 FC-AL Connection**




FC-ALs provide a solution for two nodes and a small number of devices in relatively static configurations. All devices on the loop share the media, and any packet traveling from one device to another must pass through all intermediate devices.

If your high-availability needs can be met with a two-node server cluster, an FC-AL deployment has several advantages:

- The cost is relatively low.
- Loops can be expanded to add storage (although nodes cannot be added).
- Loops are easy for Fibre Channel vendors to develop.

The disadvantage is that loops can be difficult to deploy in an organization. Because every device on the loop shares the media, overall bandwidth in the cluster is lowered. Some organizations might also be unduly restricted by the 126-device limit.

**Switched Fabric (FC-SW)**
For any cluster larger than two nodes, a switched Fibre Channel switched fabric (FC-SW) is the only supported storage technology. In an FC-SW, devices are connected in a many-to-many topology using Fibre Channel switches (illustrated in Figure 29).

**Figure 29 FC-SW Connection**




When a node or device communicates with another node or device in an FC-SW, the source and target set up a point-to-point connection (similar to a virtual circuit) and communicate directly with each other. The fabric itself routes data from the source to the target. In an FC-SW, the media is not shared. Any device can communicate with any other device, and communication occurs at full bus speed. This is a fully scalable enterprise solution and, as such, is highly recommended for deployment with server clusters.

FC-SW is the primary technology employed in SANs. Other advantages of FC-SW include ease of deployment, the ability to support millions of devices, and switches that provide fault isolation and rerouting. Also, there is no shared media as there is in FC-AL, allowing for faster communication. However, be aware that FC-SWs can be difficult for vendors to develop, and the switches can be expensive. Vendors also have to account for interoperability issues between components from different vendors or manufacturers.

**Using SANs with Server Clusters**
For any large-scale cluster deployment, it is recommended that you use a SAN for data storage. Smaller SCSI and stand-alone Fibre Channel storage devices work with server clusters, but SANs provide superior fault tolerance.

A SAN is a set of interconnected devices (such as disks and tapes) and servers that are connected to a common communication and data transfer infrastructure (FC-SW, in the case of Windows Server 2003 clusters). A SAN allows multiple server access to a pool of storage in which any server can potentially access any storage unit.

The information in this section provides an overview of using SAN technology with your Windows Server 2003 clusters. For additional information about deploying server clusters on SANs, see the Windows Clustering: Storage Area Networks link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources.

| **Note** |
| --- |
| Vendors that provide SAN fabric components and software management tools have a wide range of tools for setting up, configuring, monitoring, and managing the SAN fabric. Contact your SAN vendor for details about your particular SAN solution. |

**SCSI Resets**
Earlier versions of Windows server clusters presumed that all communications to the shared disk should be treated as an isolated SCSI bus. This behavior may be somewhat disruptive, and it does not take advantage of the more advanced features of Fibre Channel to both improve arbitration performance and reduce disruption.

One key enhancement in Windows Server 2003 is that the Cluster service issues a command to break a RESERVATION, and the StorPort driver can do a targeted or device reset for disks that are on a Fibre Channel topology. In Windows 2000 server

clusters, an entire bus-wide SCSI RESET is issued. This causes all devices on the bus to be disconnected. When a SCSI RESET is issued, a lot of time is spent resetting devices that may not need to be reset, such as disks that the CHALLENGER node may already own.

Resets in Windows 2003 occur in the following order:

1. Targeted logical unit number (LUN)
2. Targeted SCSI ID
3. Entire bus-wide SCSI RESET

| **Note** |
| --- |
| Targeted resets require functionality in the host bus adapter (HBA) drivers. The driver must be written for StorPort and not SCSIPort. Drivers that use SCSIPort will use the Challenge and Defense the same as it is currently in Windows 2000. Contact the manufacturer of the HBA to determine if it supports StorPort. |

**SCSI Commands**
The Cluster service uses the following SCSI commands:

- SCSI reserve: This command is issued by a host bus adapter or controller to maintain ownership of a SCSI device. A device that is reserved refuses all commands from all other host bus adapters except the one that initially reserved it, the initiator. If a bus-wide SCSI reset command is issued, loss of reservation occurs.
- SCSI release: This command is issued by the owning host bus adapter; it frees a SCSI device for another host bus adapter to reserve.
- SCSI reset: This command breaks the reservation on a target device. This command is sometimes referred to globally as a "bus reset."

The same control codes are used for Fibre Channel as well. These parameters are defined in this partner article:

309186 How the Cluster Service Takes Ownership of a Disk on the Shared Bus

317162 Supported Fibre Channel Configurations

The following sections provide an overview of SAN concepts that directly affect a server cluster deployment.

**HBAs**
Host bus adapters (HBAs) are the interface cards that connect a cluster node to a SAN, similar to the way that a network adapter connects a server to a typical Ethernet network. HBAs, however, are more difficult to configure than network adapters (unless the HBAs are preconfigured by the SAN vendor). All HBAs in all nodes should be identical and be at the same driver and firmware revision

**Zoning and LUN Masking**
Zoning and LUN masking are fundamental to SAN deployments, particularly as they relate to a Windows Server 2003 cluster deployment.

**Zoning**
Many devices and nodes can be attached to a SAN. With data stored in a single *cloud*, or storage entity, it is important to control which hosts have access to specific devices. Zoning allows administrators to partition devices in logical volumes and thereby reserve the devices in a volume for a server cluster. That means that all interactions between cluster nodes and devices in the logical storage volumes are isolated within the boundaries of the zone; other noncluster members of the SAN are not affected by cluster activity.

Figure 30 is a logical depiction of two SAN zones (Zone A and Zone B), each containing a storage controller (S1and S2, respectively).

**Figure 30 Zoning**




In this implementation, Node A and Node B can access data from the storage controller S1, but Node C cannot. Node C can access data from storage controller S2.

Zoning needs to be implemented at the hardware level (with the controller or switch) and not through software. The primary reason is that zoning is also a security mechanism for a SAN-based cluster, because unauthorized servers cannot access devices inside the zone (access control is implemented by the switches in the fabric, so a host adapter cannot gain access to a device for which it has not been configured). With software zoning, the cluster would be left unsecured if the software component failed.

In addition to providing cluster security, zoning also limits the traffic flow within a given SAN environment. Traffic between ports is routed only to segments of the fabric that are in the same zone.

**LUN Masking**

A LUN is a logical disk defined within a SAN. Server clusters see LUNs and think they are physical disks. LUN masking, performed at the controller level, allows you to define relationships between LUNs and cluster nodes. Storage controllers usually provide the means for creating LUN-level access controls that allow access to a given LUN to one or more hosts. By providing this access control at the storage controller, the controller itself can enforce access policies to the devices.

LUN masking provides more specific security than zoning, because LUNs provide a means for zoning at the port level. For example, many SAN switches allow overlapping zones, which enable a storage controller to reside in multiple zones. Multiple clusters in multiple zones can share the data on those controllers. Figure 31 illustrates such a scenario.

**Figure 31 Storage Controller in Multiple Zones**




LUNs used by Cluster A can be masked, or hidden, from Cluster B so that only authorized users can access data on a shared storage controller.

**Requirements for Deploying SANs with Windows Server 2003 Clusters**

The following list highlights the deployment requirements you need to follow when using a SAN storage solution with your server cluster. For a white paper that provides more complete information about using SANs with server clusters, see the Windows Clustering: Storage Area Networks link on the Web Resources page at

http://www.microsoft.com/windows/reskits/webresources.

Each cluster on a SAN must be deployed in its own zone. The mechanism the cluster uses to protect access to the disks can have an adverse effect on other clusters that are in the same zone. By using zoning to separate the cluster traffic from other cluster or noncluster traffic, there is no chance of interference.

All HBAs in a single cluster must be the same type and have the same firmware version. Many storage and switch vendors require that all HBAs on the same zone—and, in some cases, the same fabric—share these characteristics.

All storage device drivers and HBA device drivers in a cluster must have the same software version.

Never allow multiple nodes access to the same storage devices unless they are in the same cluster.

Never put tape devices into the same zone as cluster disk storage devices. A tape device could misinterpret a bus rest and rewind at inappropriate times, such as during a large backup.

**Guidelines for Deploying SANs with Windows Server 2003 Server Clusters**
In addition to the SAN requirements discussed in the previous section, the following practices are highly recommended for server cluster deployment:

In a highly available storage fabric, you need to deploy clustered servers with multiple HBAs. In these cases, always load the multipath driver software. If the I/O subsystem sees two HBAs, it assumes they are different buses and enumerates all the devices as though they were different devices on each bus. The host, meanwhile, is seeing multiple paths to the same disks. Failure to load the multipath driver will disable the second device because the operating system sees what it thinks are two independent disks with the same signature.

Do not expose a hardware snapshot of a clustered disk back to a node in the same cluster. Hardware snapshots must go to a server outside the server cluster. Many controllers provide snapshots at the controller level that can be exposed to the cluster as a completely separate LUN. Cluster performance is degraded when multiple devices have the same signature. If the snapshot is exposed back to the node with the original disk online, the I/O subsystem attempts to rewrite the signature. However, if the snapshot is exposed to another node in the cluster, the Cluster service does not recognize it as a different disk and the result could be data corruption. Although this is not specifically a SAN issue, the controllers that provide this functionality are typically deployed in a SAN environment.

For additional information, see the following articles in the Microsoft Knowledge Base:

301647 Cluster Service Improvements for Storage Area Networks

304415 Support for Multiple Clusters Attached to the Same SAN Device

280743 Windows Clustering and Geographically Separate Sites

# Related Links (Guide to Creating and Configuring a Server Cluster under Windows Server 2003 White Paper)

Updated: November 11, 2003

Applies To: Windows Server 2003 with SP1

See the following resources for further information:

- Microsoft Cluster Service Installation Resources at http://support.microsoft.com/?id=259267
- Quorum Drive Configuration Information at http://support.microsoft.com/?id=280345
- Recommended Private "Heartbeat" Configuration on Cluster Server at http://support.microsoft.com/?id=258750
- Network Failure Detection and Recovery in a Server Cluster at http://support.microsoft.com/?id=242600
- How to Change Quorum Disk Designation at http://support.microsoft.com/?id=280353
- Microsoft Windows Clustering: Storage Area Networks at http://www.microsoft.com/windows.netserver/techinfo/overview/san.mspx
- Geographically Dispersed Clusters in Windows Server 2003 at http://www.microsoft.com/windows.netserver/techinfo/overview/clustergeo.mspx

For the latest information about Windows Server 2003, see the Windows 2003 Server Web site at http://www.microsoft.com/windowsserver2003/default.mspx

# Deploying SQL Server Logging with Windows Server 2003 Internet Authentication Service (IAS)

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In This White Paper**


Introduction (Deploying SQL Server Logging with Windows Server 2003 Internet Authentication Service (IAS))

Key concepts for IAS SQL Server logging

Local IAS SQL Server logging

IAS SQL Server logging with a central database

IAS SQL Server logging using MSDE 2000

Related Links (Deploying SQL Server Logging with Windows Server 2003 Internet Authentication Service (IAS))

# Introduction (Deploying SQL Server Logging with Windows Server 2003 Internet Authentication Service (IAS))

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

SQL Server logging is a new feature for Internet Authentication Service (IAS) in Microsoft® Windows® Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

IAS is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy. As a RADIUS server, IAS performs centralized connection authentication, authorization, and accounting for many types of network access including wireless, 802.1X authenticating switch, and remote access dial-up and virtual private network (VPN) connections. As a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers. RADIUS is an Internet Engineering Task Force (IETF) standard.

Figure 1 shows the components of a RADIUS infrastructure when the IAS server is configured to log to a computer running SQL Server.




**Figure 1: Components of a RADIUS infrastructure**

A RADIUS server has access to user account information and can check network access authentication credentials. If the user's credentials are authenticated and the connection attempt is authorized, the RADIUS server authorizes the user's access on the basis of specified conditions and logs the network access connection in an accounting log. The use of RADIUS allows the network access user authentication, authorization, and accounting data to be collected and maintained in a central location, rather than on each access server.

The centralization of accounting data simplifies both troubleshooting and the log file backup process. In addition, as networks and IAS deployments scale up and out, you can deploy IAS SQL Server logging to centralize logging for multiple IAS servers.

**Using accounting logs**

IAS accounting logs can be used for a variety of purposes, including the following:

- Connection troubleshooting
- Billing and auditing
- Report creation

- Session state information

IAS provides the following logging options to address these purposes:

**Event logging for IAS.** IAS errors are automatically recorded in the system event logs of Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. IAS authentication events are recorded in the system event log on the basis of event log settings, which you can customize for your logging needs. You can log rejected authentication requests, successful authentication requests, or both. Event logging is used primarily for auditing and troubleshooting connection attempts.

**To configure event logging for IAS**

1. Open the Internet Authentication Service Microsoft Management Console (MMC) snap-in.
2. Right-click **Internet Authentication Service (IAS)**, and then click **Properties**.
3. On the **General** tab, select each required option, and then click **OK**.

For more information, see "Related Links" later in this white paper.

**Logging user authentication and accounting requests to local text files.** Recording user authentication and accounting requests in log files is used primarily for connection analysis and billing purposes, and is also useful as a security investigation tool, providing you with a method of tracking the activity of a malicious user after an attack. You can configure IAS to log in IAS format or database-import format, both of which are text file formats.

**To select requests to be logged in IAS format or database-import format**

1. Open the Internet Authentication Service snap-in.
2. In the console tree, click **Remote access logging**.
3. In the details pane, right-click **Local File**, and then click **Properties**.
4. On the **Settings** tab, select one or more check boxes for recording authentication and accounting requests in the IAS log files:
    - To capture accounting requests and responses, select **Accounting requests**.
    - To capture authentication requests, Access-Accept messages, and Access-Reject messages, select **Authentication requests**.
    - To capture periodic status updates, such as interim accounting packets, select **Periodic status**.

For more information, see "Related Links" later in this white paper.

**IAS SQL Server logging.** IAS SQL Server logging is used when you need session state information, for report creation and data analysis purposes, and to centralize and simplify management of your accounting data. Windows Server 2003 Internet Authentication Service provides the ability to use SQL Server logging to record user authentication and accounting requests received from one or more network access servers (NASs) to a data source on a computer running the Microsoft SQL Server Desktop Engine (MSDE 2000) or SQL Server 2000. Accounting data is passed from IAS in XML format to a stored procedure in an MSDE 2000 or SQL Server 2000 database, which supports both structured query language (SQL) and XML (SQLXML). Recording user authentication and accounting requests in an XML-compliant SQL Server database enables multiple IAS servers to share one data source, and provides the advantages of using a relational database.

#### Introduction to SQL Server 2000

Microsoft SQL Server is a relational database management and analysis product for e-commerce, line-of-business, and data warehousing solutions. SQL Server 2000 includes support for XML and HTTP, performance and availability features to partition loads and ensure uptime, and advanced management and tuning functionality to automate routine tasks and lower total cost of ownership. For more information, see "Related Links" later in this white paper.

#### Advantages of SQL Server logging

Logging to a SQL Server 2000 relational database instead of a standard text file (that is, either IAS format or database-compatible format) has the following advantages:

- Relationships between data tables enable the flexible creation of dynamic data views by using queries and reports.
- SQL Server 2000 has high-speed optimizations that can effectively support very large, terabyte-sized databases.
- Multiple bulk copy operations can be performed concurrently against a single table in SQL Server to speed data entry.
- The Transact-SQL BACKUP and RESTORE statements are optimized to read through a database serially and write in parallel to multiple backup devices. In addition, SQL Server 2000 supports differential backups, which replicate only the data changed after the last backup.
- You can configure SQL Server logging to provide failover and redundancy, and multiple IAS servers can log to the same database, providing ease of administration.
- Data from multiple IAS servers can be stored in the same database, providing a centralized data source for many IAS servers and the ability to achieve a more comprehensive data view for applications that query the database, such as Help Desk applications.

In addition, you can improve logging performance and, in some circumstances, reduce the cost of deploying IAS SQL Server logging by using MSDE 2000 databases installed on each IAS server.

**Stored procedures**
All three of the scenarios depicted in this white paper require the use of SQL Server 2000 stored procedures. A stored procedure is a group of Transact-SQL statements compiled into a single execution plan, or program. A stored procedure centralizes business rules, tasks, and processes within the server. When you use IAS SQL Server logging with MSDE 2000 or SQL Server 2000, the Transact-SQL programming language is the primary programming interface between your IAS server and the MSDE 2000 or SQL Server 2000 database.

When you use Transact-SQL programs, you must store the programs as stored procedures in MSDE 2000 and SQL Server 2000. IAS is configured to send data in XML format to a stored procedure called report_event, which then processes the results. Stored procedures in SQL Server 2000 and MSDE 2000 are similar to procedures and functions in other programming languages in that they can:

- Accept input parameters and return multiple values in the form of output parameters to the calling procedure or batch.
- Contain programming statements that perform operations in the database, including calling other procedures.
- Return a status value to a calling procedure or batch to indicate success or failure (and the reason for failure).

For more information, see “Related Links” later in this white paper.

**Note**

The stored procedure in the SQL Server 2000 or MSDE 2000 database that is called by IAS must be named report_event, or IAS SQL Server logging will fail.

Installing, designing, coding, and administering SQL Server 2000 databases requires extensive knowledge of relational database design and administration, as well as familiarity with SQL Server 2000 or later. Do not attempt to deploy IAS SQL Server logging in a production environment without previous experience with SQL Server 2000.

You can use IAS with any RFC-compliant RADIUS client, including dial-up servers, virtual private network (VPN) servers, wireless access points, and 802.1X authenticating switches. For more information about RADIUS, see RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)."

You can configure IAS on a computer running Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of two remote RADIUS server groups. You can define a RADIUS client by using a fully qualified domain name (FQDN) or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the FQDN of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS on a computer running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range. IAS is not included on computers running the Windows Server 2003, Web Edition, operating system.

# Key concepts for IAS SQL Server logging

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following sections describe key logging concepts such as requests logged by IAS, IAS log file contents, and how IAS creates an XML document from accounting and authentication data.

**Data logged by IAS**

By default, IAS does not log any data until you configure it to do so. When you configure SQL Server logging, all required attributes, accounting, and authentication data that is normally logged in either IAS or database-compatible format is logged to the SQL Server 2000 or MSDE 2000 database.

Initially, it is recommended that you enable the logging of accounting and user authentication requests. You can refine your logging settings after you determine your required data.

**Requests logged by IAS**

You can log the following information in a SQL Server database:

- **Accounting-on requests**, which are sent by the access server to indicate that it is online and ready to accept connections.
- **Accounting-off requests**, which are sent by the access server to indicate that it is going offline.
- **Accounting-start requests**, which are sent by the access server (after the user is authenticated and authorized by the IAS server) to indicate the start of a user session.
- **Accounting-stop requests**, which are sent by the access server to indicate the end of a user session.
- **Accounting interim requests**, which are sent periodically by some access servers during a user session, and which can be logged by IAS. This type of request can be used when the Acct-Interim-Interval RADIUS attribute is configured to support periodic requests in the remote access profile on the IAS server. The access server must support the use of accounting interim requests if you want the interim requests to be logged on the IAS server. If the access server does not send accounting interim requests, they are not logged.
- **Authentication requests**, which are sent by the access server on behalf of the connecting user. These entries in the log contain only incoming attributes.
- **Authentication accepts and rejects**, which are sent by IAS to the access server, indicating whether the user should be accepted or rejected. These entries contain only outgoing attributes.

**IAS log file contents**

Unlike database-import log files, which use a fixed sequence of attributes, the sequence of the attributes in a SQL Server 2000 or MSDE 2000 log file depends upon the sequence used by the access server. For more information about the sequence of these records, see the documentation for your access server.

For lists of attributes that are logged and are not logged in the MSDE 2000 and SQL Server databases, see “Interpreting IAS-formatted log files” in “Related Links” later in this white paper.

For information about vendor-specific attributes (VSAs), see “Interpreting IAS IDs for vendor-specific attributes” in “Related Links” later in this white paper.

**Required database fields for session correlation**

When you create the fields in the tables of your MSDE 2000 and SQL Server databases, you must provide all fields that allow applications to query the database, correlate related fields, and return a cohesive view of each session in the query results. This is called session correlation. At a minimum, to provide *session correlation*, you must log the following IAS accounting data in the MSDE 2000 and SQL Server 2000 databases:

- NAS-IP-Address

- NAS-Identifier (You need both NAS-IP-Address and NAS-Identifier because the network access server, or NAS, can send one or the other.)
- Class
- Acct-Session-Id
- Acct-Multi-Session-Id
- Packet-Type
- Acct-Status-Type
- Acct-Interim-Interval
- NAS-Port
- Event-Timestamp

**Configuring accounting to provide the best session correlation**
When you use an application to query your SQL Server 2000 database, it is important that sufficient data is logged to provide the application with the ability to correlate related fields and information into a cohesive view of any particular user session. To provide the best session correlation, take the following actions:

- Use NASs that send the Class attribute in all accounting-requests. The Class attribute is sent to the RADIUS client in an Access-Accept message, and is useful for correlating Accounting-Request messages with authentication sessions.
- If the Class attribute is sent by the NAS in the accounting request messages, it can be used to match the accounting and authentication records. The combination of the attributes Unique-Serial-Number, Service-Reboot-Time, and Server-Address must be a unique identification for each authentication that the server accepts.
- Use NASs that support interim accounting.
- Use NASs that send Accounting-on and Accounting-off messages.
- Use NASs that support the storing and forwarding of accounting data. NASs that support this feature can store accounting data in circumstances when it cannot communicate with the IAS server. When the IAS server is available, the NAS forwards the stored records to the IAS server, providing increased reliability in accounting over NASs that do not provide this feature.
- Always configure the Acct-Interim-Interval attribute in remote access policies. The Acct-Interim-Interval attribute sets the interval (in seconds) between each interim update that the NAS sends. You can add the Acct-Interim-Interval attribute on the Advanced tab on the profile settings of the remote access policy. According to RFC 2869, the value of the Acct-Interim-Interval attribute must not be smaller than 60 seconds, or one minute, and should not be smaller than 600 seconds, or 10 minutes. For more information, see RFC 2869, "RADIUS Extensions."
- Ensure that logging of periodic status is enabled on your IAS servers.

**How IAS creates an XML document from accounting and authentication data**
If you select SQL Server logging in the IAS snap-in, attribute-value pairs are converted to XML format.

When IAS creates an XML document from attributes, accounting, and authentication data, the XML document includes the attribute ID or name, the attribute value, and the data type of the attribute value. With SQL Server logging, there are five data types for attribute values:

- Non-negative integers (data_type=0)

- Strings (data_type=1)
- Hexadecimal numbers (data_type=2)
- IPv4 addresses (data_type=3)
- Date and time (data_type=4)

The following example XML document created by IAS includes the attribute name (User-Name), the attribute value (DOMAIN*username*), and the data type (1), indicating a value data type of string:

```
<Event>
    <User-Name data_type="1">DOMAIN\username</User-Name>
</Event>
```

The next example XML document created by IAS includes the attribute name (NAS-IP-Address), the attribute value (192.168.0.1), and the data type (3), indicating a value data type of IP address:

```
<Event>
    <NAS-IP-Address data_type="3">192.168.0.1</NAS-IP-Address>
</Event>
```

The next example XML document created by IAS includes the attribute name (Provider-Type), the attribute value (1), and the data type (0), indicating a value data type of non-negative integers:

```
<Event>
    <Provider-Type data_type="0">1</Provider-Type>
</Event>
```

The last example XML document, which is created by IAS, includes the three previous examples combined into one XML document:

```
<Event>
    <User-Name data_type="1">DOMAIN\username</User-Name>
    <NAS-IP-Address data_type="3">192.168.0.1</NAS-IP-Address>
    <Provider-Type data_type="0">1</Provider-Type>
</Event>
```

Following is an example of a typical XML document sent by IAS configured for SQL Server logging:

```
    <Event>
<Computer-Name data_type="1">MYNAS</Computer-Name>
<Event-Source data_type="1">IAS</Event-Source>
<Acct-Session-Id data_type="1">10</Acct-Session-Id>
<NAS-IP-Address data_type="3">10.10.1.1</NAS-IP-Address>
<Service-Type data_type="0">2</Service-Type>
<Framed-Protocol data_type="0">1</Framed-Protocol>
<NAS-Port data_type="0">7</NAS-Port>
<NAS-Port-Type data_type="0">5</NAS-Port-Type>
<Tunnel-Type data_type="0">1</Tunnel-Type>
<Tunnel-Medium-Type data_type="0">1</Tunnel-Medium-Type>
<Calling-Station-Id data_type="1">10.10.1.2</Calling-Station-Id>
<Tunnel-Client-Endpt data_type="1">10.10.1.2</Tunnel-Client-Endpt>
<User-Name data_type="1">MYDOMAIN\Administrator</User-Name>
<Client-IP-Address data_type="3">10.10.1.1</Client-IP-Address>
<Client-Vendor data_type="0">0</Client-Vendor>
<Client-Friendly-Name data_type="1">MYNAS</Client-Friendly-Name>
<MS-RAS-Vendor data_type="0">311</MS-RAS-Vendor>
<MS-RAS-Version data_type="1">MSRASV5.20</MS-RAS-Version>
<MS-RAS-Client-Version data_type="1">MSRASV5.20</MS-RAS-Client-Version>
<MS-RAS-Client-Name data_type="1">MSRAS-0-MYCLIENT</MS-RAS-Client-Name>
<Provider-Type data_type="0">1</Provider-Type>
<Class data_type="1">311 1 192.168.0.123 02/20/2003 19:03:02 9</Class>
<SAM-Account-Name data_type="1">MYDOMAIN\Administrator</SAM-Account-Name>
<Fully-Qualified-User-Name data_type="1">MYDOMAIN\Administrator</Fully-Qualified-User-
Name>
<Authentication-Type data_type="0">4</Authentication-Type>
```

```
<Packet-Type data_type="0">1</Packet-Type>
<Reason-Code data_type="0">0</Reason-Code>
    </Event>
```

**Example stored procedure**

Following is an example stored procedure that creates an IAS database within the SQL Server 2000 database environment and processes XML documents sent to the computer running SQL Server by IAS servers configured to log to this specific computer running SQL Server. If you want to use this stored procedure in a test environment, you can copy the text below into Notepad, then save the document with the file name extension *.sql. Import the stored procedure into a database in SQL Server 2000, and run the stored procedure.

| **Note** |
| --- |
| This stored procedure is provided solely as an example, and is unsupported. |

```
IF EXISTS (SELECT name FROM master.dbo.sysdatabases WHERE name = N'IASODBC')
    DROP DATABASE [IASODBC]
GO

CREATE DATABASE [IASODBC]  ON (NAME = N'IASODBC_Data', FILENAME = N'C:\Program
Files\Microsoft SQL Server\MSSQL\data\IASODBC_Data.MDF' , SIZE = 1, FILEGROWTH = 10%)
LOG ON (NAME = N'IASODBC_Log', FILENAME = N'C:\Program Files\Microsoft SQL
Server\MSSQL\data\IASODBC_Log.LDF' , SIZE = 1, FILEGROWTH = 10%)
 COLLATE SQL_Latin1_General_CP1_CI_AS
GO

exec sp_dboption N'IASODBC', N'autoclose', N'false'
GO

exec sp_dboption N'IASODBC', N'bulkcopy', N'false'
GO

exec sp_dboption N'IASODBC', N'trunc. log', N'false'
GO

exec sp_dboption N'IASODBC', N'torn page detection', N'true'
GO

exec sp_dboption N'IASODBC', N'read only', N'false'
GO

exec sp_dboption N'IASODBC', N'dbo use', N'false'
GO

exec sp_dboption N'IASODBC', N'single', N'false'
GO

exec sp_dboption N'IASODBC', N'autoshrink', N'false'
GO

exec sp_dboption N'IASODBC', N'ANSI null default', N'false'
GO

exec sp_dboption N'IASODBC', N'recursive triggers', N'false'
GO

exec sp_dboption N'IASODBC', N'ANSI nulls', N'false'
GO

exec sp_dboption N'IASODBC', N'concat null yields null', N'false'
GO

exec sp_dboption N'IASODBC', N'cursor close on commit', N'false'
GO

exec sp_dboption N'IASODBC', N'default to local cursor', N'false'
GO

exec sp_dboption N'IASODBC', N'quoted identifier', N'false'
GO
```

```
exec sp_dboption N'IASODBC', N'ANSI warnings', N'false'
GO

exec sp_dboption N'IASODBC', N'auto create statistics', N'true'
GO

exec sp_dboption N'IASODBC', N'auto update statistics', N'true'
GO

if( ( (@@microsoftversion / power(2, 24) = 8) and (@@microsoftversion & 0xffff >= 724)
) or ( (@@microsoftversion / power(2, 24) = 7) and (@@microsoftversion & 0xffff >=
1082) ) )
    exec sp_dboption N'IASODBC', N'db chaining', N'false'
GO

use [IASODBC]
GO

if exists (select * from dbo.sysobjects where id = object_id(N'[dbo].[report_event]')
and OBJECTPROPERTY(id, N'IsProcedure') = 1)
drop procedure [dbo].[report_event]
GO

if exists (select * from dbo.sysobjects where id = object_id(N'[dbo].
[accounting_data]') and OBJECTPROPERTY(id, N'IsUserTable') = 1)
drop table [dbo].[accounting_data]
GO

if exists (select * from dbo.systypes where name = N'ipaddress')
exec sp_droptype N'ipaddress'
GO

setuser
GO

EXEC sp_addtype N'ipaddress', N'nvarchar (15)', N'not null'
GO

setuser
GO

CREATE TABLE [dbo].[accounting_data] (
    [id] [int] IDENTITY (1, 1) NOT NULL ,
    [timestamp] [datetime] NOT NULL ,
    [Computer_Name] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NOT NULL ,
    [Packet_Type] [int] NOT NULL ,
    [User_Name] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [F_Q_User_Name] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Called_Station_Id] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Calling_Station_Id] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Callback_Number] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Framed_IP_Address] [ipaddress] NULL ,
    [NAS_Identifier] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [NAS_IP_Address] [ipaddress] NULL ,
    [NAS_Port] [int] NULL ,
    [Client_Vendor] [int] NULL ,
    [Client_IP_Address] [ipaddress] NULL ,
    [Client_Friendly_Name] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Event_Timestamp] [datetime] NULL ,
    [Port_Limit] [int] NULL ,
    [NAS_Port_Type] [int] NULL ,
    [Connect_Info] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Framed_Protocol] [int] NULL ,
    [Service_Type] [int] NULL ,
    [Authentication_Type] [int] NULL ,
    [NP_Policy_Name] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Reason_Code] [int] NULL ,
    [Class] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Session_Timeout] [int] NULL ,
    [Idle_Timeout] [int] NULL ,
    [Termination_Action] [int] NULL ,
```

```
    [EAP_Friendly_Name] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Acct_Status_Type] [int] NULL ,
    [Acct_Delay_Time] [int] NULL ,
    [Acct_Input_Octets] [int] NULL ,
    [Acct_Output_Octets] [int] NULL ,
    [Acct_Session_Id] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Acct_Authentic] [int] NULL ,
    [Acct_Session_Time] [int] NULL ,
    [Acct_Input_Packets] [int] NULL ,
    [Acct_Output_Packets] [int] NULL ,
    [Acct_Terminate_Cause] [int] NULL ,
    [Acct_Multi_Session_Id] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL
,
    [Acct_Link_Count] [int] NULL ,
    [Acct_Interim_Interval] [int] NULL ,
    [Tunnel_Type] [int] NULL ,
    [Tunnel_Medium_Type] [int] NULL ,
    [Tunnel_Client_Endpoint] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL
,
    [Tunnel_Server_Endpoint] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL
,
    [Acct_Tunnel_Connection] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL
,
    [Tunnel_Pvt_Group_Id] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Tunnel_Assignment_Id] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Tunnel_Preference] [int] NULL ,
    [MS_Acct_Auth_Type] [int] NULL ,
    [MS_Acct_EAP_Type] [int] NULL ,
    [MS_RAS_Version] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [MS_RAS_Vendor] [int] NULL ,
    [MS_CHAP_Error] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [MS_CHAP_Domain] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [MS_MPPE_Encryption_Types] [int] NULL ,
    [MS_MPPE_Encryption_Policy] [int] NULL ,
    [Proxy_Policy_Name] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Provider_Type] [int] NULL ,
    [Provider_Name] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [Remote_Server_Address] [ipaddress] NULL ,
    [MS_RAS_Client_Name] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL ,
    [MS_RAS_Client_Version] [nvarchar] (255) COLLATE SQL_Latin1_General_CP1_CI_AS NULL
) ON [PRIMARY]
GO

SET QUOTED_IDENTIFIER ON
GO
SET ANSI_NULLS OFF
GO

CREATE PROCEDURE dbo.report_event
    @doc ntext
AS

SET NOCOUNT ON

DECLARE @idoc int
EXEC sp_xml_preparedocument @idoc OUTPUT, @doc

/*
    All RADIUS attributes written to the ODBC format logfile are declared here.  One
additional attribute is added: @record_timestamp.
    The value of @record_timestamp is the UTC time the record was inserted in the
database.

    Refer to IAS ODBC Formatted Log Files in Online Help for information on
interpreting these values.
*/

DECLARE @record_timestamp datetime

SET @record_timestamp = GETUTCDATE()

INSERT accounting_data
```

```
SELECT
    @record_timestamp,
    Computer_Name,
    Packet_Type,
    [User_Name],
    F_Q_User_Name,
    Called_Station_Id,
    Calling_Station_Id,
    Callback_Number,
    Framed_IP_Address,
    NAS_Identifier,
    NAS_IP_Address,
    NAS_Port,
    Client_Vendor,
    Client_IP_Address,
    Client_Friendly_Name,
    Event_Timestamp,
    Port_Limit,
    NAS_Port_Type,
    Connect_Info,
    Framed_Protocol,
    Service_Type,
    Authentication_Type,
    NP_Policy_Name,
    Reason_Code,
    Class,
    Session_Timeout,
    Idle_Timeout,
    Termination_Action,
    EAP_Friendly_Name,
    Acct_Status_Type,
    Acct_Delay_Time,
    Acct_Input_Octets,
    Acct_Output_Octets,
    Acct_Session_Id,
    Acct_Authentic,
    Acct_Session_Time,
    Acct_Input_Packets,
    Acct_Output_Packets,
    Acct_Terminate_Cause,
    Acct_Multi_Session_Id,
    Acct_Link_Count,
    Acct_Interim_Interval,
    Tunnel_Type,
    Tunnel_Medium_Type,
    Tunnel_Client_Endpoint,
    Tunnel_Server_Endpoint,
    Acct_Tunnel_Connection,
    Tunnel_Pvt_Group_Id,
    Tunnel_Assignment_Id,
    Tunnel_Preference,
    MS_Acct_Auth_Type,
    MS_Acct_EAP_Type,
    MS_RAS_Version,
    MS_RAS_Vendor,
    MS_CHAP_Error,
    MS_CHAP_Domain,
    MS_MPPE_Encryption_Types,
    MS_MPPE_Encryption_Policy,
    Proxy_Policy_Name,
    Provider_Type,
    Provider_Name,
    Remote_Server_Address,
    MS_RAS_Client_Name,
    MS_RAS_Client_Version
FROM OPENXML(@idoc, '/Event')
WITH (
    Computer_Name nvarchar(255) './Computer-Name',
    Packet_Type int './Packet-Type',
    [User_Name] nvarchar(255) './User-Name',
    F_Q_User_Name nvarchar(255) './Fully-Qualifed-User-Name',
    Called_Station_Id nvarchar(255) './Called-Station-Id',
```

```
    Calling_Station_Id nvarchar(255) './Calling-Station-Id',
    Callback_Number nvarchar(255) './Callback-Number',
    Framed_IP_Address nvarchar(15) './Framed-IP-Address',
    NAS_Identifier nvarchar(255) './NAS-Identifier',
    NAS_IP_Address nvarchar(15) './NAS-IP-Address',
    NAS_Port int './NAS-Port',
    Client_Vendor int './Client-Vendor',
    Client_IP_Address nvarchar(15) './Client-IP-Address',
    Client_Friendly_Name nvarchar(255) './Client-Friendly-Name',
    Event_Timestamp datetime './Event-Timestamp',
    Port_Limit int './Port-Limit',
    NAS_Port_Type int './NAS-Port-Type',
    Connect_Info nvarchar(255) './Connect-Info',
    Framed_Protocol int './Framed-Protocol',
    Service_Type int './Service-Type',
    Authentication_Type int './Authentication-Type',
    NP_Policy_Name nvarchar(255) './NP-Policy-Name',
    Reason_Code int './Reason-Code',
    Class nvarchar(255) './Class',
    Session_Timeout int './Session-Timeout',
    Idle_Timeout int './Idle-Timeout',
    Termination_Action int './Termination-Action',
    EAP_Friendly_Name nvarchar(255) './EAP-Friendly-Name',
    Acct_Status_Type int './Acct-Status-Type',
    Acct_Delay_Time int './Acct-Delay-Time',
    Acct_Input_Octets int './Acct-Input-Octets',
    Acct_Output_Octets int './Acct-Output-Octets',
    Acct_Session_Id nvarchar(255) './Acct-Session-Id',
    Acct_Authentic int './Acct-Authentic',
    Acct_Session_Time int './Acct-Session-Time',
    Acct_Input_Packets int './Acct-Input-Packets',
    Acct_Output_Packets int './Acct-Output-Packets',
    Acct_Terminate_Cause int './Acct-Terminate-Cause',
    Acct_Multi_Session_Id nvarchar(255) './Acct-Multi-Session-Id',
    Acct_Link_Count int './Acct-Link-Count',
    Acct_Interim_Interval int './Acct-Interim-Interval',
    Tunnel_Type int './Tunnel-Type',
    Tunnel_Medium_Type int './Tunnel-Medium-Type',
    Tunnel_Client_Endpoint nvarchar(255) './Tunnel-Client-Endpt',
    Tunnel_Server_Endpoint nvarchar(255) './Tunnel-Server-Endpt',
    Acct_Tunnel_Connection nvarchar(255) './Acct-Tunnel-Connection',
    Tunnel_Pvt_Group_Id nvarchar(255) './Tunnel-Pvt-Group-Id',
    Tunnel_Assignment_Id nvarchar(255) './Tunnel-Assignment-Id',
    Tunnel_Preference int './Tunnel-Preference',
    MS_Acct_Auth_Type int './MS-Acct-Auth-Type',
    MS_Acct_EAP_Type int './MS-Acct-EAP-Type',
    MS_RAS_Version nvarchar(255) './MS-RAS-Version',
    MS_RAS_Vendor int './MS-RAS-Vendor',
    MS_CHAP_Error nvarchar(255) './MS-CHAP-Error',
    MS_CHAP_Domain nvarchar(255) './MS-CHAP-Domain',
    MS_MPPE_Encryption_Types int './MS-MPPE-Encryption-Types',
    MS_MPPE_Encryption_Policy int './MS-MPPE-Encryption-Policy',
    Proxy_Policy_Name nvarchar(255) './Proxy-Policy-Name',
    Provider_Type int './Provider-Type',
    Provider_Name nvarchar(255) './Provider-Name',
    Remote_Server_Address nvarchar(15) './Remote-Server-Address',
    MS_RAS_Client_Name nvarchar(255) './MS-RAS-Client-Name',
    MS_RAS_Client_Version nvarchar(255) './MS-RAS-Client-Version'
    )

EXEC sp_xml_removedocument @idoc

SET NOCOUNT OFF
GO
SET QUOTED_IDENTIFIER OFF
GO
SET ANSI_NULLS ON
GO
```

**IAS SQL Server logging scenarios**

The following three deployment scenarios provide SQL Server logging solutions tuned to organizations with specific logging requirements:

- **Local IAS SQL Server logging.** This scenario depicts IAS servers that log to a local SQL Server 2000 database. Distributed queries can be run against databases on multiple IAS servers by using SQL Query Analyzer or a third-party product. This scenario is optimal for medium-sized organizations that need local collection of accounting data for reporting and billing.

- **IAS SQL Server logging with a central database.** This scenario depicts IAS servers that log to a local SQL Server 2000 database. SQL Server 2000 replication is deployed to publish records from the SQL Server database on each IAS server to a central SQL Server 2000 database. This scenario is optimal for large-sized organizations that need centralized collection of accounting data for reporting and billing.

- **IAS SQL Server logging using MSDE 2000.** This scenario depicts IAS servers that log to a local MSDE 2000 database. A custom accounting-forwarder component installed on each IAS server manages the local databases and then publishes records to a central SQL Server 2000 database. This scenario is optimal for enterprises.

In the sets of instructions that follow, these assumptions are made:

1. Your organization has already deployed Windows Server 2003 Internet Authentication Service on your network, and your network access solution is functioning properly with another form of logging (either IAS format or database-compatible format). If you have not already deployed IAS on your network, see "Deploying IAS as a RADIUS server" and "Deploying IAS as a RADIUS proxy" in Help and Support Center for Windows Server 2003.

2. You have SQL Server 2000 relational database development experience and you understand how to use SQL Server 2000 to create, modify, administer, and manage SQL Server 2000 databases. For more information, see "Related Links" later in this white paper.

3. For **the SQL Server logging using MSDE 2000** scenario, you have software development experience that allows you to create the custom accounting-forwarder component that is installed and run on each IAS server configured for SQL Server logging.

**Deployment in stages**

Before you deploy any of the three scenarios depicted in this white paper, it is recommended that you first deploy the scenario in a test lab environment.

After your test lab solution is functioning properly and you are prepared to deploy SQL Server logging in a production environment, you can deploy SQL Server logging in stages, as follows:

**Stage 1**: Use a text file format for logging. Configure IAS accounting to use either IAS format or database-compatible format and then verify that it is functioning properly.

**Stage 2**: Use both a text file format and SQL Server logging. With all instances of SQL Server 2000 or MSDE 2000 configured and in place in the production environment, enable SQL Server logging by using the IAS snap-in of each IAS server without disabling logging in IAS format or database-compatible format. Log in both formats for a duration that allows you to verify that your SQL Server logging solution is functioning properly. In this circumstance, if either SQL Server logging or text logging fails, IAS will halt processing of network connection attempts, and service to users will be interrupted.

**Stage 3**: Use SQL Server logging format. Disable text file logging after you have verified that your SQL Server logging solution is functioning properly.

# Local IAS SQL Server logging

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following scenario depicts local IAS SQL Server logging for medium-sized organizations. In this scenario, SQL Server 2000 is installed on each IAS server. Each IAS server is then configured to log to the local SQL Server 2000 database. Queries against multiple databases are performed using SQL Query Analyzer or a third-party product with similar capabilities.

**Components of local IAS SQL Server logging**
Figure 2 shows the components of IAS SQL Server logging.




**Figure 2: Components of local IAS SQL Server logging**

This configuration consists of the following components:

**IAS server running SQL Server**
A computer running Windows Server 2003 and Internet Authentication Service (IAS) configured to use SQL Server logging. The IAS server is also running SQL Server 2000. In the IAS snap-in, SQL Server logging is enabled, with a data link configured between IAS and SQL Server 2000 on the local computer.

Within the SQL Server 2000 database environment on the IAS server, the following components are configured:

- **IAS database.** A database configured to store IAS accounting data.
- **Stored procedure.** A program that runs inside the IAS database that accepts and processes the XML documents sent by IAS, and then stores accounting data in SQL Server 2000 tables.

**How local SQL Server logging works**
Local IAS SQL Server logging, including data collection, storage, processing, and querying for data views, works in three basic stages:

1. Accounting data is produced by a client connection attempt
2. IAS logs accounting data to the stored procedure in the local SQL Server 2000 database
3. Applications query the IAS databases

**1. Accounting data is produced by a client connection attempt**
IAS performs centralized connection authentication, authorization, and accounting for many types of network access and router-to-router connections. When an access client attempts to connect to the network, it does so by using a network access server, such as a dial-up server, VPN server, wireless access point, or 802.1X authenticating switch. When these access servers are configured as RADIUS clients to an IAS-RADIUS server, IAS attempts to authenticate the user or computer with the Active Directory user accounts database for the domain, and then IAS performs authorization (with user account properties and/or remote access policies configured on the IAS server) to discover whether the user or computer has the necessary permissions to access network resources.

**2. IAS logs accounting data to the stored procedure in the local SQL Server 2000 database**
Data from the access server and data produced by IAS are sent in XML format to a stored procedure in the SQL Server 2000 database on the local computer.

When IAS logs to a local SQL Server 2000 database, the following interactions occur:

1. The IAS server receives or produces the data to be logged to a SQL Server 2000 database. For example, a network access server configured as a RADIUS client to the IAS server sends the IAS server incoming attributes for an authentication request by a user or computer.

2. The IAS server packages the data as an XML document.

3. The IAS server initiates a connection with SQL Server 2000. If the IAS server cannot establish a connection with SQL Server 2000, the IAS server stops processing authentication and accounting requests and users cannot log on to the network.

    - When you configure IAS to log to a SQL Server 2000 database installed on the local computer, select **Use Windows NT integrated security** on the IAS server **Connection** tab in the **Data Link Properties** dialog box. The IAS server then attempts to authenticate the connection by using the Local System account. For more information, see "Configure log file properties" in Help and Support Center for Windows Server 2003.

4. If a connection with SQL Server 2000 is successful, the IAS server signals SQL Server 2000 to run a stored procedure named report_event. IAS then sends the accounting data (an XML document) to the report_event stored procedure in the SQL Server 2000 database. The stored procedure that is called by IAS must be named report_event, or IAS SQL Server logging will fail. If the IAS server is unable to write to the SQL Server 2000 database, IAS SQL Server logging will fail and users cannot log on.

5. The report_event stored procedure uses only one parameter. This parameter is a Unicode string of type *ntext*. The Unicode string contains an XML document which in turn contains the accounting data to be logged.

6. The report_event stored procedure in the SQL Server 2000 database processes the XML document and stores the processed data in SQL Server 2000 database tables.

7. SQL Server 2000 notifies the IAS server that the transaction is successful.

8. The IAS server is free to process additional data, repeating this process as required.

**3. Applications query the IAS databases**

You can create a custom application, such as a Help Desk application, to view data stored in the SQL Server databases on each IAS server. You can write a stand-alone application with a software development program such as the Microsoft® Visual Basic® development system, or you can use Microsoft® Access 2002, SQL Server 2000 Enterprise Manager, or SQL Query Analyzer to view the data.

Figure 3 depicts a query against SQL Server databases on two IAS servers by using SQL Query Analyzer on a portable computer.




**Figure 3: SQL Query Analyzer runs a SQL query**

For more information, see "SQL Server Connection Basics" in "Related Links" later in this white paper.

**To deploy local IAS SQL Server logging**

1. Install SQL Server 2000 on each IAS server that you want to configure for SQL Server logging. For more information, see "Overview of Installing SQL Server 2000" in "Related Links" later in this white paper.

2. By using SQL Server 2000 tools, you can create the database structure in the SQL Server 2000 database for storage of IAS accounting log data. For more information, see "Creating a database" in "Related Links" later in this white paper.

3. Create a stored procedure in the SQL Server 2000 database called report_event that will process the data passed to it by IAS in XML format. For more information, see "Create Procedure" in "Related Links" later in this white paper.

4. Configure log file properties and the data link for each IAS server. Configure the data link on each IAS server to log to the local SQL Server 2000 database.

To configure log file properties and the data link, do the following on each IAS server:

1. Open the IAS snap-in.

2. In the details pane, right-click **SQL Server**, and then click **Properties**.

3. In **Log the following information**, select the information that you want to log:

   a. To log all accounting requests, click **Accounting requests**.

   b. To log authentication requests, click **Authentication requests**.

   c. To log periodic status, such as interim accounting requests, click **Periodic status**.

   d. To configure the number of concurrent sessions allowed between the IAS server and the SQL Server database, type a number in **Maximum number of concurrent sessions**.

   e. To configure the SQL Server data source, click **Configure**, and then, on the **Connection** tab, specify the following:

      To specify the name of the server on which the database is stored, type or select a name in **Select or enter a server name**.

      To specify the authentication method with which to log on to the server, click **Use Windows NT integrated security**. Or, click **Use a specific user name and password**, and then type credentials in **User name** and **Password**.

      To allow a blank password, click **Blank password**.

      To store the password, click **Allow saving password**.

      To specify which database to connect to on the computer running SQL Server, click **Select the database on the server**, and then select a database name from the list.

To test the connection between the IAS server and the computer running SQL Server, click **Test Connection**.

For more information, see "Configure log file properties" in Help and Support Center for Windows Server 2003.

# IAS SQL Server logging with a central database

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following scenario depicts IAS SQL Server logging with a central database for large-sized organizations. In this scenario, SQL Server 2000 is installed on each IAS server. Each IAS server is then configured to log to the local SQL Server 2000 database. In addition, a computer or computer cluster running SQL Server 2000 is deployed as a central database. All records from IAS servers are periodically published to the central SQL Server database. Queries against the central SQL Server database are performed using SQL Query Analyzer or a third-party product with similar capabilities.

**Components of IAS SQL Server logging with a central database**

Figure 4 depicts the components of SQL Server logging with a central database.




**Figure 4: The components of IAS SQL Server logging with a central database**

This configuration consists of the following components:

**IAS servers running SQL Server**

Two or more computers running Windows Server 2003 and Internet Authentication Service (IAS) configured to use SQL Server logging. The IAS servers are also running SQL Server 2000. In the IAS snap-in, SQL Server logging is enabled, with a data link configured between IAS and SQL Server 2000 on the local computer.

Within the SQL Server 2000 database environment on each IAS server, the following components are configured:

- **IAS database.** A database configured to store IAS accounting data.
- **Stored procedure.** A program that runs inside the IAS database. The stored procedure accepts and processes the XML documents sent by IAS and stores accounting data in SQL Server 2000 tables.

**Central computer running SQL Server**

Within the SQL Server 2000 database, the following components are configured:

- **Central IAS database.** A database that stores IAS accounting data that is published by the IAS servers running SQL Server.

**How IAS SQL Server logging with a central database works**

IAS SQL Server logging, including data collection, storage, processing, and querying for data views, works in five basic stages:

1. Accounting data is produced by a client connection attempt

2. IAS logs accounting data to the stored procedure in the local SQL Server 2000 database

3. IAS servers running SQL Server publish records to the central SQL Server database

4. The central SQL Server database warehouses data from one or more IAS servers

5. Applications query the IAS database on the computer running SQL Server

**1. Accounting data is produced by a client connection attempt**
IAS performs centralized connection authentication, authorization, and accounting for many types of network access and router-to-router connections. When an access client attempts to connect to the network, it does so by using a network access server, such as a dial-up server, VPN server, wireless access point, or 802.1X authenticating switch. When these access servers are configured as RADIUS clients to an IAS-RADIUS server, IAS attempts to authenticate the user or computer with the Active Directory user accounts database for the domain, and then IAS performs authorization (with user account properties and/or remote access policies configured on the IAS server) to discover whether the user or computer has the necessary permissions to access network resources.

**2. IAS logs accounting data to the stored procedure in the local SQL Server 2000 database**
Data from the access server and data produced by IAS is sent in XML format to a stored procedure in the SQL Server 2000 database on the local computer. The access server is configured at the IAS server as a RADIUS client.

When IAS logs to a local SQL Server 2000 database, the following interactions occur:

1. The IAS server receives or produces the data to be logged to a SQL Server 2000 database. For example, a network access server configured as a RADIUS client to the IAS server sends the IAS server incoming attributes for an authentication request by a user or computer.

2. The IAS server packages the data as an XML document.

3. The IAS server initiates a connection with SQL Server 2000. If the IAS server cannot establish a connection with SQL Server 2000, the IAS server stops processing authentication and accounting requests and users cannot log on to the network.

    - When you configure IAS to log to a SQL Server 2000 database installed on the local computer, select **Use Windows NT integrated security** on the IAS server **Connection** tab in the **Data Link Properties** dialog box. The IAS server then attempts to authenticate the connection by using the Local System account. For more information, see "Configure log file properties" in Help and Support Center for Windows Server 2003.

4. If a connection with SQL Server 2000 is successful, the IAS server signals SQL Server 2000 to run a stored procedure named report_event. IAS then sends the accounting data (an XML document) to the report_event stored procedure in the SQL Server 2000 database. The stored procedure that is called by IAS must be named report_event, or IAS SQL Server logging will fail. If the IAS server is unable to write to the SQL Server 2000 database, IAS SQL Server logging will fail and users cannot log on.

5. The report_event stored procedure uses only one parameter. This parameter is a Unicode string of type *ntext*. The Unicode string contains an XML document which in turn contains the accounting data to be logged.

6. The report_event stored procedure in the SQL Server 2000 database processes the XML document and stores the processed data in SQL Server 2000 database tables.

7. SQL Server 2000 notifies the IAS server that the transaction is successful.

8. The IAS server is free to process additional data, repeating this process as required.

**3. IAS servers running SQL Server publish records to the central SQL Server database**
After a time period that you specify, the IAS servers running SQL server act as publishers and replicate data to the central SQL Server database, which acts as the subscriber. SQL Server 2000 replication is a set of technologies for copying and distributing

data and database objects from one database to another and then synchronizing between databases for consistency. You can also configure the SQL Server database on IAS servers to delete previously published records. For more information, see "Replication Overview" in "Related Links" later in this white paper.

**4. The central SQL Server database warehouses data from two or more IAS servers**
The IAS database on the computer running SQL Server contains all records from all IAS servers configured with the SQL Server logging components and configured to log to the computer running SQL Server. The central IAS database provides a comprehensive view of accounting data for network access administrators and others who need to view the IAS accounting data.

**5. Applications query the IAS database on the central SQL Server**
You can create a custom application, such as a Help Desk application, to view data stored in the SQL Server database. You can write a stand-alone application with a software development program such as Visual Basic, or you can use Access 2002, SQL Server 2000 Enterprise Manager, or SQL Query Analyzer to view the data. For more information, see "Creating a Data Warehouse" in "Related Links" later in this white paper.

**To deploy IAS SQL Server logging with a central database**

1. Install SQL Server 2000 on each IAS server that you want to configure for SQL Server logging.

2. By using SQL Server 2000 tools, you can create the database structure in the SQL Server 2000 database for storage of IAS accounting log data.

3. Create a stored procedure in the SQL Server 2000 database called report_event that will process the data passed to it by IAS in XML format.

4. Create the central IAS database. On the computer or computer cluster running SQL Server 2000 that will function as the central SQL Server, create an IAS database to hold records from all IAS servers configured to use SQL Server logging.

5. Configure replication from the IAS servers running SQL Server to the central SQL Server database.

6. Configure log file properties and the data link for each IAS server.

To configure log file properties and the data link, do the following on each IAS server:

1. Open the IAS snap-in.

2. In the details pane, right-click **SQL Server**, and then click **Properties**.

3. In **Log the following information**, select the information that you want to log:

    a. To log all accounting requests, click **Accounting requests**.

    b. To log authentication requests, click **Authentication requests**.

    c. To log periodic status, such as interim accounting requests, click **Periodic status**.

    d. To configure the number of concurrent sessions allowed between the IAS server and the SQL Server database, type a number in **Maximum number of concurrent sessions**.

    e. To configure the SQL Server data source, click **Configure**, and then, on the **Connection** tab, specify the following:

    To specify the name of the server on which the database is stored, type or select a name in **Select or enter a server name**.

    To specify the authentication method with which to log on to the server, click **Use Windows NT integrated security**. Or, click **Use a specific user name and password**, and then type credentials in **User name** and **Password**.

To allow a blank password, click **Blank password**.

To store the password, click **Allow saving password**.

To specify which database to connect to on the computer running SQL Server, click **Select the database on the server**, and then select a database name from the list.

To test the connection between each IAS server and the central computer running SQL Server, click **Test Connection**.

For more information, see "Configure log file properties" in Help and Support Center for Windows Server 2003.

# IAS SQL Server logging using MSDE 2000

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following scenario depicts IAS SQL Server logging for the enterprise. This scenario uses both SQL Server 2000 and MSDE 2000, as well as a custom accounting-forwarder component that manages the MSDE 2000 databases installed on each IAS server.

| **Important** |
| --- |
| **To deploy this scenario you must have software development experience.** You must create the custom accounting-forwarder component described in this scenario. This component is installed and run on each IAS server configured for SQL Server logging. |

**Introduction to MSDE 2000**

The Microsoft SQL Server 2000 Desktop Engine (MSDE 2000) is a data engine built and based on core SQL Server technology. With support for single- and dual-processor computers, MSDE 2000 is a reliable storage engine and query processor for desktop extensions of enterprise applications. The common technology base shared between SQL Server and MSDE 2000 enables developers to build applications that can scale from portable computers to multiprocessor clusters. This common technology base also allows IAS to use either a MSDE 2000 database or a database configured on SQL Server 2000 as storage for accounting data.

Designed to run in the background, supporting transactional desktop applications, MSDE 2000 does not have its own user interface (UI). Users typically interact with MSDE 2000 through the application in which it is embedded. MSDE 2000 is packaged in a self-extracting archive for ease of distribution and embedding. The only tool that is provided with MSDE 2000 is the **osql** utility. The executable file, Sql.exe located in the MSSQL\Binn folder, provides a default instance of MSDE 2000.

When you use MSDE 2000 in conjunction with SQL Server 2000, you can use the client tools included with SQL Server 2000, such as SQL Server Enterprise Manager and SQL  Query Analyzer, to configure and administer MSDE 2000 databases.

For more information about MSDE, see "Related Links" later in this white paper.

**Components of IAS SQL Server logging using MSDE 2000**

Figure 5 shows the components of IAS SQL Server logging using MSDE 2000.




**Figure 5: Components of IAS SQL Server logging using MSDE 2000**

This configuration consists of the following components:

**IAS server**

A computer running Windows Server 2003 and Internet Authentication Service (IAS) configured to use SQL Server logging.

To implement SQL Server logging, the following components are required on the IAS server:

- **MSDE 2000 database.** MSDE 2000 installed on the computer running Windows Server 2003 and IAS. If there are multiple IAS servers configured for SQL Server logging, each server must have MSDE 2000 locally installed and configured for logging.

- **Stored procedure.** A program that runs inside the MSDE 2000 database on the computer running Windows Server 2003 and IAS. The stored procedure must be named report_event. Report_event accepts and processes data from the IAS server and stores it in MSDE 2000 tables.

- **Custom accounting-forwarder component.** A custom accounting-forwarder component, service, or application installed on the computer running Windows Server 2003 and IAS. The custom accounting-forwarder component periodically publishes records from the MSDE database on the IAS server to the IAS database on the computer running SQL Server. The custom component also flags local records as published, and deletes previously published records from the MSDE 2000 database during off-peak hours. If there are multiple IAS servers configured for SQL Server logging, each server must have the custom component installed to process records in the local MSDE 2000 database and forward records to the SQL Server database.

**Computer running SQL Server**
A computer or computer cluster running SQL Server 2000. If you have a small IAS deployment, you might need only one computer running SQL Server; however if you have a large IAS deployment in an enterprise network environment, you might need more than one computer or cluster running SQL Server.

To implement SQL Server logging, the following component is required on the computer running SQL Server:

- **IAS database.** A database on the computer running SQL Server that is dedicated to storing the records from all IAS servers.

**How IAS SQL Server logging works with MSDE 2000**
IAS SQL Server logging, including data collection, storage, processing, and querying for data views, works in five basic stages:

1. Accounting data is produced by a client connection attempt
2. IAS logs accounting data to the stored procedure in the local MSDE 2000 database
3. A custom accounting-forwarder component manages the local MSDE 2000 database
4. SQL Server warehouses data from one or more IAS servers
5. Applications query the IAS master database on the computer running SQL Server

**1. Accounting data is produced by a client connection attempt**
IAS performs centralized connection authentication, authorization, and accounting for many types of network access and router-to-router connections. When an access client attempts to connect to the network, it does so by using a network access server, such as a dial-up server, VPN server, wireless access point, or 802.1X authenticating switch. When these access servers are configured as RADIUS clients to an IAS-RADIUS server, IAS attempts to authenticate the user or computer with the Active Directory user accounts database for the domain, and IAS performs authorization (with user account properties and/or remote access policies configured on the IAS server) to discover whether the user or computer has the necessary permissions to access network resources.

If the RADIUS server is configured to provide accounting, or logging of session information, accounting data is produced that can be logged to an XML-compliant database on a computer running SQL Server 2000.

**2. IAS logs accounting data to the stored procedure in the local MSDE 2000 database**
Data from the access server and data produced by IAS are sent in XML format to a stored procedure called report_event in the MSDE 2000 database. The access server is configured at the IAS server as a RADIUS client.

Figure 6 depicts how RADIUS accounting data flows from an access server to the MSDE 2000 database on the IAS server.




**Figure 6: RADIUS accounting data flow from access server to the MSDE 2000 database**

When IAS SQL Server logging logs to a local MSDE 2000 database, the following interactions occur:

1. The IAS server receives or produces the data to be logged to a MSDE 2000 database. For example, a network access server configured as a RADIUS client to the IAS server sends the IAS server incoming attributes for an authentication request by a user or computer.

2. The IAS server packages the data as an XML document.

3. The IAS server initiates a connection with MSDE 2000. If the IAS server cannot establish a connection with MSDE 2000, the IAS server stops processing authentication and accounting requests and users cannot log on to the network.

    - When you configure IAS to log to an MSDE 2000 database installed on the local computer, select **Use Windows NT integrated security** on the IAS server **Connection** tab in the **Data Link Properties** dialog box. The IAS server then attempts to authenticate the connection by using the Local System account. For more information, see "Configure log file properties" in Help and Support Center for Windows Server 2003.

4. If a connection with MSDE 2000 is successful, the IAS server signals MSDE 2000 to run a stored procedure named report_event. IAS then sends the accounting data (an XML document) to the report_event stored procedure in the MSDE 2000 database. The stored procedure that is called by IAS must be named report_event, or IAS SQL Server logging will fail. If the IAS server is unable to write to the MSDE 2000 database (for example, if the database maximum size limit of 2 gigabytes (GBs) is reached), IAS SQL Server logging will fail and users cannot log on.

5. The report_event stored procedure uses only one parameter. This parameter is a Unicode string of type *ntext*. The Unicode string contains an XML document which in turn contains the accounting data to be logged.

6. The report_event stored procedure in the MSDE 2000 database processes the XML document and stores the processed data in MSDE 2000 database tables.

7. MSDE 2000 notifies the IAS server that the transaction is successful.

8. The IAS server is free to process additional data, repeating this process as required.

**3. A custom accounting-forwarder component manages the local MSDE 2000 database**

For the IAS SQL Server logging solution to work, you must write a custom component that manages the MSDE 2000 database on IAS servers. Due to the 2-GB limit of MSDE 2000, records must be periodically published by a custom accounting-forwarder component from MSDE 2000 to a designated database on the computer running SQL Server. The accounting-forwarder component can be installed as a service, a scheduled task in MSDE 2000, or as an application.

The custom accounting-forwarder component has two primary functions:

- Publish records from the MSDE 2000 database to the SQL Server database

- Delete previously published records from the MSDE 2000 database

The custom accounting-forwarder component performs these actions to protect against possible data loss and to save system resources during peak hours. Differential backups from the MSDE 2000 database to the computer running SQL Server provide the ability to publish small amounts of records at periodic intervals. Deletion of previously published records during off-peak hours saves resources during peak hours.

You can design your custom accounting-forwarder component to perform in the following manner:

- **Connect to a SQL Server 2000 database**. The component creates a data link with the computer running SQL Server and publishes records to the designated SQL Server database. The connection is opened when the custom component starts. The connection between the custom component and the computer running SQL Server is left open unless the custom component is stopped.

- **Perform differential backups to a SQL Server database**. At a specific interval, the component performs a differential

backup, publishing new records from the MSDE 2000 database to the computer running SQL Server. (*New records* means IAS accounting records that are added to the MSDE 2000 database after the previous differential backup.) You can design the component to publish new records at an interval that best suits your IAS deployment and your degree of network access traffic, however the interval must be calculated to publish records before the MSDE database reaches its 2-GB storage limit. If you have a large number of network access requests processed and logged by the IAS server, the 2-GB limit of the MSDE 2000 database will be reached much more quickly than on a network with few network access requests to the IAS server.

- **Flag published records in the MSDE 2000 database**. To save hardware resources during peak hours, records published to the computer running SQL Server are not immediately deleted in the MSDE 2000 database, they are flagged for deletion during off-peak hours.

- **Delete flagged records during off-peak hours**. To free up space in the MSDE 2000 database, records that are flagged as previously published are deleted during off-peak hours. You can design the component to delete flagged records at an interval that best suits your IAS deployment, however the interval must be calculated to delete flagged records before the MSDE database reaches its 2-GB storage limit.

- **Close the connection to the computer running SQL Server**. The component must be capable of closing the data link connection to the computer running SQL Server when the component is stopped.

**Example of the performance of a custom accounting-forwarder service**

Following is an example of the performance of the custom accounting-forwarder component within the data flow process, on a network with a large volume of network access traffic. In the example, IAS-01 is an IAS server configured to use SQL Server logging with a local MSDE 2000 database. The custom accounting-forwarder component, called the Accounting Forwarder service, is also installed on IAS-01. In addition, SQL-01 is a computer running SQL Server 2000 that is configured with a database dedicated to storing records from the IAS server. The database on the computer running SQL Server is named IAS_db.

In this example, authentication, authorization and accounting for network access is provided by IAS-01. When the Accounting Forwarder service on IAS-01 is started, it creates a connection to SQL-01 using a valid user name and password. The connection is open between the service and the computer running SQL Server for this example of the recommended data flow process:

1. IAS-01 logs accounting data to the local MSDE 2000 database.

2. The Accounting Forwarder service on IAS-01 reads the MSDE 2000 database, seeking only new records that are not flagged as previously published.

3. The Accounting Forwarder service on IAS-01 publishes all new, unflagged records by executing a stored procedure on SQL-01 with the data read from the MSDE 2000 database on IAS-01. The Accounting Forwarder service publishes the IAS accounting data at a 15-second interval.

4. The stored procedure on SQL-01 processes and stores the records as they are received from the Accounting Forwarder service on IAS-01. The records are stored in the IAS_db database by the stored procedure.

5. The Accounting Forwarder service on IAS-01 flags the records in the MSDE 2000 database as published.

6. At the off-peak time of 2 A.M., the Accounting Forwarder service on IAS-01 deletes all flagged records from the MSDE 2000 database.

**4. SQL Server warehouses data from one or more IAS servers**

The IAS database on the computer running SQL Server contains all records from all IAS servers configured with the SQL Server logging components and configured to log to the computer running SQL Server, providing a comprehensive view of accounting data for network access administrators and others who need to view the IAS accounting data.

**5. Applications query the IAS database on the computer running SQL Server**

You can create a custom application, such as a Help Desk application, to view data stored in the SQL Server database. You can write a stand-alone application with a software development program such as Visual Basic, or you can use Access 2002, SQL Server 2000 Enterprise Manager, or SQL Query Analyzer to view the data.

**To deploy IAS SQL Server logging with MSDE 2000**

When you deploy IAS SQL Server logging, you can configure IAS to log directly to a computer running SQL Server 2000 or you can log to an MSDE 2000 database on the local computer.

Logging directly to a computer running SQL Server presents the following issues:

- Heavy network traffic can slow logging performance, possibly resulting in a timeout between the IAS server and the computer running SQL Server.
- Network hardware failure between the IAS server and the computer running SQL Server can cause IAS logging to fail. For example, if a router or switch between the IAS server and the computer running SQL Server fails, IAS logging will fail and users cannot log on.
- SQL Server hardware or software failure can cause IAS logging to fail and prevent users from logging on.

If IAS logging stops for any reason, users attempting to authenticate against the RADIUS server or RADIUS proxy server will fail, and users will be unable to log on to the network and access network resources. For this reason, install MSDE 2000 on each IAS server used for logging, and then configure a data link between the IAS server and the local MSDE 2000 database. Because MSDE 2000 has a 2-GB database size limit, local storage of IAS logging data is temporary, and you can instead periodically publish the data to a central computer running SQL Server using a custom component, service, or application.

Figure 7 shows the recommended deployment model for IAS SQL Server logging with MSDE 2000.




**Figure 7: Deployment model for IAS SQL Server logging with MSDE 2000**

To deploy IAS SQL Server logging, the basic steps are as follows:

1. Install MSDE 2000 on each IAS server that you want to configure for SQL Server logging.
2. By using SQL Server 2000 tools, you can create the database structure in the MSDE 2000 database for storage of IAS accounting log data.
3. Create a stored procedure in the MSDE 2000 database called report_event that will process the data passed to it by IAS in XML format.
4. Create a custom accounting-forwarder component, application, or service that performs the functions described earlier in this paper.
5. On the computer running SQL Server 2000, create an IAS server database to hold records from each IAS server configured to use SQL Server logging. Create a stored procedure in the database to process the data passed to it by the custom accounting-forwarder component on each IAS server that the database serves.
6. Install the custom accounting-forwarder component, application, or service on each IAS server.
7. Configure each IAS server for SQL Server 2000 logging, including configuring the data you want to log and the data link

between IAS and the local MSDE 2000 database.

To configure log file properties and the data link, do the following on each IAS server:

1. Open the IAS snap-in.

2. In the details pane, right-click **SQL Server**, and then click **Properties**.

3. In **Log the following information**, select the information that you want to log:

   a. To log all accounting requests, click **Accounting requests**.

   b. To log authentication requests, click **Authentication requests**.

   c. To log periodic status, such as interim accounting requests, click **Periodic status**.

   d. To configure the number of concurrent sessions allowed between the IAS server and the SQL Server database, type a number in **Maximum number of concurrent sessions**.

   e. To configure the SQL Server data source, click **Configure**, and then, on the **Connection** tab, specify the following:

      To specify the name of the server on which the database is stored, type or select a name in **Select or enter a server name**.

      To specify the authentication method with which to log on to the server, click **Use Windows NT integrated security**. Or, click **Use a specific user name and password**, and then type credentials in **User name** and **Password**.

      To allow a blank password, click **Blank password**.

      To store the password, click **Allow saving password**.

      To specify which database to connect to on the computer running SQL Server, click **Select the database on the server**, and then select a database name from the list.

To test the connection between each IAS server and the computer running SQL Server, click **Test Connection**.

For more information, see “Configure log file properties” in Help and Support Center for Windows Server 2003.

# Related Links (Deploying SQL Server Logging with Windows Server 2003 Internet Authentication Service (IAS))

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

See the following resources for further information.

**IAS resources**

- "Internet Authentication Service" in Help and Support Center for Windows Server 2003, or on the Web at http://go.microsoft.com/fwlink/?LinkId=20133.
- "Event logging for IAS" in Help and Support Center for Windows Server 2003, or on the Web at http://go.microsoft.com/fwlink/?LinkId=30283.
- "Logging user authentication and accounting requests" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=30284.
- "Interpreting IAS-formatted log files" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=30285.
- "Interpreting IAS IDs for vendor-specific attributes" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=30286.

If you have questions about IAS, you can post them to the Microsoft RADIUS Usenet newsgroup microsoft.public.internet.radius.

**SQL Server 2000 resources**

- SQL Server product Web site at http://go.microsoft.com/fwlink/?LinkId=20014.
- SQL Server Product Documentation, including SQL Server 2000 Books Online, at http://go.microsoft.com/fwlink/?LinkId=8628.
- "SQL Stored Procedures" on MSDN at http://go.microsoft.com/fwlink/?LinkId=8628.
- "SQL Server Connection Basics" on MSDN at http://go.microsoft.com/fwlink/?LinkId=8628.
- "Overview of Installing SQL Server 2000" on MSDN at http://go.microsoft.com/fwlink/?LinkId=30293.
- "Creating and maintaining databases overview" on MSDN at http://go.microsoft.com/fwlink/?LinkId=30303.
- "Create Procedure" on MSDN at http://go.microsoft.com/fwlink/?LinkId=30305.
- "Replication Overview" on MSDN at http://go.microsoft.com/fwlink/?LinkId=30308.
- "Creating a Data Warehouse" on MSDN at http://go.microsoft.com/fwlink/?LinkId=30309.

**MSDE 2000 resources**

- MSDE 2000 home page at http://go.microsoft.com/fwlink/?LinkId=9140.
- Microsoft Knowledge Base article 325003, "How to: Manage the SQL Server Desktop Engine (MSDE 2000) by Using the Osql Utility" at http://go.microsoft.com/fwlink/?LinkId=30311.

- "SQLXML" on MSDN at http://go.microsoft.com/fwlink/?LinkId=30317.

For the latest information about IAS and Windows Server 2003, see the Windows Server 2003 Web site at http://go.microsoft.com/fwlink/?LinkId=5221.

# Deploying Windows Server 2003 Internet Authentication Service (IAS) with Virtual Local Area Networks (VLANs)

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This white paper describes how to configure remote access policy in Internet Authentication Service (IAS) for use with virtual local area networks (VLANs). When you configure the profile of an IAS remote access policy for use with VLANs, you must configure the attributes Tunnel-Medium-Type, Tunnel-Pvt-Group-ID, Tunnel-Type, and Tunnel-Tag. VLANs allow network architects and administrators to logically group network resources (such as servers, printers, and client computers) even when they are not on the same physical subnet.

**In This White Paper**

Introduction (Deploying Windows Server 2003 Internet Authentication Service (IAS) with Virtual Local Area Networks (VLANs))

Configuring Active Directory

Configuring IAS

Summary (Deploying Windows Server 2003 Internet Authentication Service (IAS) with Virtual Local Area Networks (VLANs))

Related Links (Deploying Windows Server 2003 Internet Authentication Service (IAS) with Virtual Local Area Networks (VLANs))

# Introduction (Deploying Windows Server 2003 Internet Authentication Service (IAS) with Virtual Local Area Networks (VLANs))

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using VLAN-aware network access servers and Internet Authentication Service (IAS) in Microsoft® Windows Server™ 2003, you can provide groups of users with access only to the network resources that are appropriate for their security permissions. For example, you can provide visitors with wireless access to the Internet without allowing them access to your organization network.

In addition, VLANs allow you to logically group network resources that exist in different physical locations or on different physical subnets. For example, members of your sales department and their network resources, such as client computers, servers, and printers, might be located in several different buildings at your organization, but you can place all of these resources on one VLAN using the same IP address range; the VLAN then functions, from the end-user perspective, as a single subnet.

You can also use VLANs when you want to segregate a network between different groups of users. After you have determined how you want to define your groups, you can create security groups in Active Directory and add members to the groups. You can define groups of users in several ways:

- **By role.** You can create groups for the sales team, the finance department, and other departments.
- **By position.** You can create groups for knowledge workers, managers, executives, and other positions.
- **By access level.** You can create groups for visitors, partners, full-time employees, and other categories with different access levels.

After you have created groups in Active Directory, you can open the IAS Microsoft Management Console (MMC) snap-in and create a remote access policy for each group. Within the remote access policy configuration process, you can define the VLAN to which the group will be assigned.

Configuration of VLAN-aware network hardware, such as VLAN-aware routers, switches, wireless access points, and access controllers, is beyond the scope of this white paper. When you have one or more of these devices configured as a Remote Authentication Dial-In User Service (RADIUS) client to your IAS server, however, you can use IAS to designate which VLAN the connecting user is placed on.

For information about how to configure your VLAN-aware network access server, see your access server documentation.

IAS is included in the following products:

- Windows Server 2003, Standard Edition
- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition
- Microsoft Windows Server 2003, 64-Bit Enterprise Edition
- Microsoft Windows Server 2003, 64-Bit Datacenter Edition
- Microsoft Windows Small Business Server 2003, Standard Edition
- Microsoft Windows Small Business Server 2003, Premium Edition

**Requirements**

To deploy IAS with VLANs as described in this white paper, the following components are required:

- A computer running Windows Server 2003 and IAS
- An Active Directory® directory service user accounts database
- RADIUS clients that are VLAN-aware, such as wireless access points, switches, or access controllers

**Remote access policy**

When you use Active Directory as your user accounts database, IAS performs network access authentication and authorization using Active Directory user account dial-in properties and remote access policies, which are configured in the IAS snap-in.

*Authentication* is the process of verifying identity, while *authorization* is the process of verifying that the user or device connecting to the network has permissions to do so.

A *remote access policy* is an ordered set of rules that define how a connection is either authorized or rejected by IAS. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting.

If a connection is authorized, the remote access policy profile settings specify a set of connection restrictions that can include the assignment of the connection to a VLAN.

**Authorization by group**

It is recommended that you manage authorization by security group rather than by individual user. Managing authorization by group provides the ability to create one remote access policy, or rule, for network connection attempts by all members of the group. For example, if you have a sales department that you want to place on a VLAN, you can create a security group named Sales in the Active Directory Users and Computers snap-in, and then you can add all members of the sales department as members of the group. When a member of the Sales group attempts to connect to the network, the connection attempt is processed by IAS with the remote access policy for the group.

If the user is authenticated and authorized, IAS applies connection restrictions, and can instruct a VLAN-aware access server to place the member of the Sales group onto the VLAN for the sales department.

IAS authorizes connection attempts with both the dial-in properties of the user account, which are configured in the Active Directory Users and Computers snap-in, and remote access policies, which are configured in the IAS snap-in.

**User account remote access permission**

One of the dial-in properties of user accounts in Active Directory is **Remote Access Permission (Dial-in or VPN)**. You can use this property to set remote access permission to be explicitly allowed, denied, or determined through remote access policies. In all cases, remote access policies are used to authorize the connection attempt.

If access is explicitly allowed, remote access policy conditions, user account properties, or profile properties can still deny the connection attempt. The **Control access through Remote Access Policy** option is only available on user accounts in a Windows 2000 native domain, a Windows Server 2003 domain, or for local accounts on stand-alone servers running Windows 2000; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition.

New accounts that are created on a stand-alone server or in a Windows 2000 native domain are set to **Control access through Remote Access Policy**. New accounts that are created in a Windows 2000 mixed domain are set to **Deny access**.

When you manage authorization by security group, the remote access permission setting in user account dial-in properties can be set to one of the following:

- **Allow access**. Only apply this setting for users who are members of groups to whom you want to grant remote access permission. You can use this setting when your domain functional level is Windows 2000 mixed. Windows 2000 mixed supports Windows NT 4.0, Windows 2000, and Windows Server 2003 family domain controllers. With user account remote access permission set to **Allow access**, authorization is performed by IAS in circumstances where you have created a policy that matches the conditions of the connection.
- **Control access through Remote Access Policy**. This setting is recommended for domains with a functional level of Windows 2000 native or Windows Server 2003.

**Important**

To set user account dial-in properties to **Control access through remote access policy**, the domain functional level must be Windows 2000 native or higher. For more information, see "Domain and forest functionality" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=30600.

**Ignoring the user account dial-in properties**

You can configure IAS to ignore the dial-in properties of user accounts. This setting is useful for circumstances where you want IAS remote access policy to determine authorization for all connections to your network. For example, you can configure IAS to ignore the dial-in properties of user accounts when:

- Your domain functional level is Windows 2000 mixed.
- Changing the current settings on user accounts is not cost-effective.

For more information about configuring IAS to ignore the dial-in properties of user accounts, see "Dial-in properties of a user account" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=30601.

For more information about remote access policies, see "Introduction to remote access policies" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=30602.

# Configuring Active Directory

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you want to create security groups in Active Directory or raise the domain functional level, perform the following procedures.

| **Important** |
| --- |
| If you have or will have any domain controllers running Windows NT 4.0 and earlier, then do not raise the domain functional level to Windows 2000 native. After the domain functional level is set to Windows 2000 native, it cannot be changed back to Windows 2000 mixed.<br><br>If you have or will have any domain controllers running Windows NT 4.0 and earlier or Windows 2000, then do not raise the domain functional level to Windows Server 2003. After the domain functional level is set to Windows Server 2003, it cannot be changed back to Windows 2000 mixed or Windows 2000 native. |

If you want to raise the domain functional level to either Windows 2000 native or Windows Server 2003, do the following:

1. Open the **Active Directory Domains and Trusts** snap-in. To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.

2. In the console tree, right-click the domain for which you want to raise the domain functional level, and then click **Raise Domain Functional Level**.

3. In **Select an available domain functional level**, do one of the following:

   - To raise the domain functional level to Windows 2000 native, click **Windows 2000 native**, and then click **Raise**.

   - To raise the domain functional level to Windows Server 2003, click **Windows Server 2003**, and then click **Raise**.

| **Note** |
| --- |
| To perform this procedure, you must be a member of the Domain Admins group in the domain for which you want to raise functionality or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority.<br><br>The current domain functional level is displayed under **Current domain functional level** in the **Raise Domain Functional Level** dialog box. |

To create a group and add members to the group, do the following:

1. Open the **Active Directory Users and Computers** snap-in.

2. In the console tree, right-click the folder in which you want to add a new group.

3. Point to **New**, and then click **Group**.

4. Type the name of the new group. By default, the name you type is automatically entered in **Group name (pre-Windows 2000)**.

5. In **Group scope**, click one of the options.

6. In **Group type**, click one of the options, and then click **OK**.

7. In the details pane, right-click the group you just created, and then click **Properties**.

8. On the **Members** tab, click **Add**.

9. In **Enter the object names to select**, type the name of the user, group, or computer that you want to add to the group, and then click **OK**.

For more information, see "To create a new group" at http://go.microsoft.com/fwlink/?LinkId=20018 and "Assign user rights to a group in Active Directory" at http://go.microsoft.com/fwlink/?LinkId=20019.

# Configuring IAS

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure IAS for use with VLANs, do the following:

1. Install Internet Authentication Service on a computer running Windows Server 2003. For more information, see "To install IAS" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20028.

2. Register IAS in Active Directory. In order for IAS to have permission to read user accounts in Active Directory, IAS must be registered with Active Directory. For more information, see "To enable the IAS server to read user accounts in Active Directory" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20030.

3. Add RADIUS clients. In the IAS snap-in, right-click **RADIUS Clients**, and then click **New RADIUS Client**. Use the **New RADIUS Client Wizard** to both add and configure your network access servers as RADIUS clients.

4. Delete the default remote access policies. To delete the policies, open the IAS snap-in, and then click **Remote Access Policies**. Select each existing policy, right-click the policy, and then click **Delete**.

5. Create a new remote access policy. In the console tree of the IAS snap-in, right-click **Remote Access Policies**, and then click **New Remote Access Policy**. Use the **New Remote Access Policy Wizard** to create a policy.

As an example policy, you can choose the following:

- For **How do you want to set up this policy?** select **Use the wizard to set up a typical policy for a common scenario**.

- For **Policy name**, type a name for your policy. For example, type **Sales policy**.

- For **Select the method of access for which you want to create a policy**, select the appropriate type of access, such as **Wireless** or **Ethernet**.

- For **Grant access based on the following**, click **Group**, and then click **Add**. In **Enter the object name to select**, type the name of a security group that you defined when configuring Active Directory. For example, if you created a group named Sales, type **Sales**, and then click **OK**.

- In **Authentication Methods**, select the authentication method that you would like to enforce for users who will be placed on this VLAN. Your choices will differ based upon the access method you have chosen for the policy, such as Wireless or VPN. When you have completed configuring an authentication method, click **Finish**.

After you have completed creating the policy and have closed the wizard, you need to configure additional items for the remote access policy. In the IAS snap-in, click **Remote Access Policies**, and then double-click the policy you just created. Make the following configuration changes to the policy:

1. In the policy **Properties** dialog box, for **Policy conditions**, click **Add**.

2. In **Attribute Types**, click **Day-And-Time-Restrictions**, and then click **Add**. In **Time of day restraints**, select **Permitted**, configure the days and times that access is permitted, and then click **OK**.

3. In the policy **Properties** dialog box, click **Grant remote access permission**.

4. Click **Edit Profile**, and then click the **Advanced** tab. By default, the **Service-Type** attribute appears in **Attributes** with a value of **Framed**. By default, for policies with access methods of VPN and dial-up, the **Framed-Protocol** attribute

appears in **Attributes** with a value of **PPP**. To specify additional connection attributes required for VLANs, click **Add**, and then add the following attributes:

- **Tunnel-Medium-Type**. Select a value appropriate to the previous selections you have made. For example, if the remote access policy you are configuring is a wireless policy, select: **Value: 802 (Includes all 802 media plus Ethernet canonical format)**.
- **Tunnel-Pvt-Group-ID**. Enter the integer that represents the VLAN number to which group members will be assigned. For example, if your Sales VLAN is on VLAN 4, type the number **4**.
- **Tunnel-Type**. Select the value **Virtual LANs (VLAN)**.
- **Tunnel-Tag**. Obtain this value from your hardware documentation.

5. Configure IAS connection request policies as you require.

| **Important** |
| --- |
| IAS evaluates remote access policies in the order in which they appear in the IAS snap-in under **Remote Access Policies**. |

For more information about remote access policies, see "Elements of a remote access policy" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=30605.

For more information about connection request processing, see "Introduction to connection request processing" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=30607.

# Summary (Deploying Windows Server 2003 Internet Authentication Service (IAS) with Virtual Local Area Networks (VLANs))

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you use VLAN-aware network hardware, such as routers, switches, and access controllers, you can configure remote access policy in Internet Authentication Service in Windows Server 2003 to instruct the access servers to place members of Active Directory groups on VLANs. When you configure the profile of an IAS remote access policy for use with VLANs, you must configure the attributes Tunnel-Medium-Type, Tunnel-Pvt-Group-ID, Tunnel-Type, and Tunnel-Tag. This ability to group network resources logically with VLANs provides flexibility when designing and implementing network solutions.

# Related Links (Deploying Windows Server 2003 Internet Authentication Service (IAS) with Virtual Local Area Networks (VLANs))

Updated: June 30, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more information about IAS, see the following:

- The Internet Authentication Service home page at http://go.microsoft.com/fwlink/?LinkId=16171.
- "Deploying IAS" in the *Windows Server 2003 Deployment Kit* at http://go.microsoft.com/fwlink/?LinkId=30603.
- "Internet Authentication Service" in Windows Server 2003 Help on the Web at http://go.microsoft.com/fwlink/?LinkId=30604.

# Kerberos authentication for load balanced web sites

Updated: July 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: July 1, 2003

**In This White Paper**

Introduction (Kerberos authentication for load balanced web sites)

Setting up Kerberos Authentication against the cluster name Service Principal Name

# Introduction (Kerberos authentication for load balanced web sites)

Updated: July 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Problem Statement**

Kerberos is an authentication mechanism used to verify user or host identity. This document explains how to configure Kerberos to allow applications to take advantage of its authentication capabilities while also being load balanced with NLB.

| **Note** |
| --- |
| The procedures described in this document apply only to products in the Windows Server 2003 family. Using Kerberos in an application load balanced with NLB is not supported in Windows 2000. |

**Outline of Solution**

Follow the steps in this white paper to register the cluster name service principal name (SPN) on a domain user account and configure IIS 6.0 to run in Worker Process Isolation Mode under the identity of that user account.

**Prerequisite Steps**

Complete these steps before implementing the Kerberos Authentication solution.

1. Join the Windows Server 2003 computers that would participate in the network load balancing (NLB) cluster to a Windows domain (not a workgroup).

2. Set up the NLB cluster.

3. Assign a name for the cluster (this is the **cluster name**) and map it to the NLB cluster IP address on DNS/WINS. This name does not have to be the **Full Internet Name** entered on the **Cluster Parameters** tab in the **Network Load Balancing Properties** dialog box.

# Setting up Kerberos Authentication against the cluster name Service Principal Name

Updated: October 6, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are three phases to setting up Kerberos Authentication against the cluster name SPN, each targeted at a different entity. A sample configuration is suggested in the following procedures, but your configuration may be different.

**Phase 1: Administration of Domain Controller**
On the Windows Server 2003 domain controller, perform the following steps. You must have administrator credentials to follow this procedure.

1. Create a domain user account for IIS 6.0

   For example, on the domain controller, run net user /add <iisid> <nlb> where <*iisid*> is the user name and <*nlb*> is the password.

2. If there isn't already one, create a domain user account for the client.

   For example, on the domain controller, run net user /add <iiststt> <nlb> where <*iiststt*> is the user name and <*nlb*> is the password.

3. Register the service principal name (SPN) that will be used by the client to identify the service on the domain user account created in step 1 using the setspn.exe tool.

   For example, if Internet Explorer is the client and testweb is the cluster name, on the domain controller, run setspn -a HOST/testweb nlb.net\iisid to register the SPN of HOST/testweb on the domain user account of nlb.net\iisid.

| **Note** |
| --- |
| The SPN will be formatted as <Service-Class>/<Host>. The value of <Service-Class> depends on the client. While Internet Explorer uses the string HOST, other clients might use http or some other string for <Service-Class>. The value of <Host> is the cluster name. |

**Phase 2: Administration of Servers**
Perform each of these steps on every Windows Server 2003 computer that is part of the NLB cluster.

1. Configure IIS to run in Worker Process Isolation Mode. Start IIS Manager, expand **Local Computer**, right-click **Web Sites**, and then click **Properties**. Click the **Service** tab, clear the **Run WWW service in IIS 5.0 isolation mode** check box, and then click **OK**.




If the switch to Worker Process Isolation Mode is successful, a folder named **Application Pools** appears in IIS Manager listing for your local computer as shown in the following diagram:




2. Assign the domain user account (created in Phase 1, step one) identity to the application pool that you wish to use to run the Web site. In this example, DefaultAppPool is the application pool used. Right-click the relevant application pool folder, select **Properties**, click the **Identity** tab, then under **Application pool identity**, select **Configurable** and enter the user name and password created earlier.




3. Add the domain user account created in Phase 1, step one to the local IIS_WPG group. This is the IIS Worker Process Group. This can be accomplished using the Computer Management dialog box or by running this command: net localgroup iis_wpg add <iisid>.

4. Create your Web site. Please refer to the IIS documentation for help with this step.

5. Assign the Web site to the application pool operated on in step 2. Right-click the Web site folder, select **Properties**, click the **Home Directory** tab, and then, in **Application settings**, in the **Application pool** drop-down list, select the appropriate application pool.

6. If using host header names, right-click the Web site folder, select **Properties**, click the **Web Site** tab, and then, in **Web site identification**, click **Advanced** and enter the cluster name for the **Host header value**.

7. Right-click the Web site folder, select **Properties**, click the **Directory Security** tab, and then, in **Authentication and access control**, click **Edit**, and then under **Authenticated access**, clear the **Basic authentication** check box and select **Integrated Windows Authentication**.

8. To force precedence of Kerberos over NTLM, navigate to the Inetpub\AdminScripts directory, where there is a script called adsutil.vbs. Run the following command: adsutil set w3svc/ntauthenticationproviders Negotiate, NTLM.

**Phase 3: Administration of Client**

There are really no special steps to be performed at the client. The client computer has to use the SPN registered on the domain user account created in Phase 1, step 1. Also, if another domain user account was created in step 2, the client should use it to log on. In our sample configuration, Internet Explorer was the client and we were logged on as Administrator. So we forced Internet Explorer to accept new credentials (created in Phase 1, step 2) by selecting **Prompt for user name and password** in Tools/Internet Options/Security/Custom Level/User Authentication/Log on.

**Solution Execution and Verification**

After you have completed the three phases, the configuration is ready to be verified. The client should be able to access the Web site using Kerberos for authentication, In our sample configuration, we accessed http://testweb/test.htm using Internet Explorer. To verify the Kerberos (instead of NTLM) was indeed used for authentication, please look for an event like the following to be logged in the **Security** folder of the servers Windows Event Log.


Event Properties
Event
Date: 12/4/2002
Source: Security
Time: 4:49:40 PM
Category: Logon/Logoff
Type: Success A
Event ID: 540
User: NLB\iistst
Computer: NLB116-1X-NS
Description:
Successful Network Logon:
User Name: iistst
Domain: NLB
Logon ID: (0x0,0x2211F7)
Logon Type: 3
Logon Process: Kerberos
Authentication Package: Kerberos
Workstation Name:
Logon GUID: {f3731da7-c6a9-3712-a6bd-a6d3554c57f3}
Data: Bytes Words
OK
Cancel
Apply

# Server Clusters: Remote Setup, Unattended Installations and Image-based Installations

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

**In This White Paper**


Overview (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Command Line Setup

Unattended Installation of Server Clusters

Image-based Installations of Server Clusters

Configuring Your Server Clusters

Image-based Installation of the Operating System and the Cluster Service Using Automated Deployment Services (ADS)

Appendix A (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Appendix B (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Appendix C (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Appendix D (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Appendix E (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Appendix F (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Appendix G (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

# Overview (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

You can install Server clusters remotely in the following ways:

Using command line to install Server clusters

Unattended installation of Server clusters during the operating system install

Image-based installation of Server clusters during the operating system install

**Using Command Line to Install Server Clusters**
If Microsoft Windows Server 2003 Enterprise Server is already installed and configured on your cluster nodes, you can simply use the command line utility called cluster.exe to install and configure the cluster service on those nodes remotely.

**Unattended Installation of the Operating System and the Cluster Service**
To perform a clean install or to upgrade the operating system, you can either perform a hands-on installation (you must manually provide all the answers required for setup), which this document will not address, or an unattended installation using a so-called answer file. Server clusters can be installed after the operating system installation completes, by specifying certain switches in the answer file. Unattended installation of the operating system and Server clusters is discussed in details in the Unattended installation of Server Clusters section of this document.

**Image-based Installation of the Operating System and the Cluster Service**
Another way to perform a clean install of the operating system is to use image-based installations, also known as disk cloning or simply cloning. Cloning is the fastest method of deploying Windows Server 2003 Enterprise Server. Image-based installations are useful if you are deploying computers that have similar hardware and you want to perform quick installations with no or minimal manual interaction. You cannot use image-based installation to upgrade an operating system. Image-based installation is discussed in more detail in the Image-based installations of Server Clusters section of this document.

**Image-based Installation of the Operating System and the Cluster Service Using Automated Deployment Services (ADS)**
Another way to perform an operating system installation and configuration is with ADS. With ADS, you can deploy an operating system onto one or more destination devices or can repurpose those computers. ADS uses above mentioned cloning mechanism to distribute the image on to a set of servers. By automating the deployment and repurposing processes, ADS reduces the overhead of maintaining diverse systems in your data center unlike plain cloning where it is useful if you are deploying computers that have similar hardware. ADS based operating system Image installation is discussed in more detail in the ADS based Image installations of Server Clusters section of this document.

# Command Line Setup

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

You do not have to log into the cluster nodes to install the cluster service, but you must make sure that the account you use to create a cluster and add nodes to the cluster belongs to the local administrators group on all of the nodes prior to installation of the cluster service. Once you have added a cluster administrator account to the local administrators group (see section Adding a cluster administrator account to the local administrators group for details), you can run the cluster.exe (located in the **%windir%\system32** directory) command line utility to create a cluster and add nodes to the cluster. Cluster.exe automatically adds the Cluster Service Account (CSA) to the local administrators group if it is not already there.

Once you have installed a cluster, you can use cluster.exe to create and configure cluster resources as described in the Windows Server 2003 Enterprise Server help.

**Adding a Cluster Administrator Account to the Local Administrators Group**
You can add an account to the local administrators group by:

logging onto each computer and using the Local users and groups snap-in

connecting to a remote computer using the Remote Desktops snap-in or

connecting to a remote computer using telnet and running the following command:

```
net localgroup administrators /add domain\account
```

You can include this line in the answer file if you want a cluster administrator account to be added to the local administrators group automatically during the operating system setup by using the [GuiRunOnce] entry in the answer file. See sections Answer File and Required Sections and Entries below for more information about GuiRunOnce and other answer file entries.

For more information about these three options please refer to the Windows Server 2003 online help.

**Creating a Cluster**
To create a cluster from the command line, use cluster.exe utility as shown in this example:

```
%windir%\cluster /cluster: CLUSTER_NAME /CREATE /NODE: NODE_NAME1
/USER:your_domain\user /PASS:****** /IPADDR:xxx.yyy.zzz.aaa
```

**Adding Nodes to the Cluster**
To add nodes to a cluster from the command line, use cluster.exe utility as shown in this example:

```
%windir%\ cluster /cluster:CLUSTER_NAME /ADDNODES: NODE_NAME1, NODE_NAME2, NODE_NAME3
/PASSWORD:*********
```

You can combine steps Creating a Cluster and Adding nodes to the cluster into one batch file, if desired.

# Unattended Installation of Server Clusters

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Unattended installation of Server clusters consists of the following steps:

Planning and Preparation for operating system installation

Preparing answer file for unattended installation of the Windows Server 2003 Enterprise Edition operating system

Creating a Master Installation

Running winnt32.exe

Installing the cluster service

**Planning and Preparation for Operating System Installation**
For detailed instructions on how to plan and prepare for Windows Server 2003 Enterprise Server operating system installation, please refer to the following document: *Microsoft Windows XP Preinstallation Guide*. This document is located in the Deploy.cab file on the Windows Server 2003 CD. Although the Guide contains the instructions on preparing for Windows XP installation, the same instructions apply to Windows Server 2003 Enterprise Edition as well.

**Answer File**
To automate the Windows Server 2003 installation completely, you must prepare an answer file, usually called unattend.txt that contains all of the answers for a series of graphical user interface (GUI) dialog boxes. Unattend.txt tells Setup how to interact with the distribution folders and files you create and supplies Setup with all of the information that the end user provides during a typical Windows Setup. For example, Unattend.txt contains a FullName entry in the **[UserData]** section that instructs Windows Setup to prompt the end user to provide a full name.

Although the answer file for Setup is commonly called Unattend.txt, for a network pre-installation you can name the file anything you like. For a CD-based Setup, you must name the answer file Winnt.sif.

If you are using the OPK toolkit (this toolkit is only available for OEMs), you can create or modify a Unattend.txt file by using a tool called Setup Manager. You can also use any text editor to create or modify the answer file.

**Required Sections and Entries**
A fully unattended operating system Setup requires the following sections and entries:

| Required section | Required entries in that section | Comments |
|---|---|---|
| [Unattended] | UnattendMode | |
| | TargetPath | |
| [GuiUnattended] | AdminPassword | |
| | TimeZone | |
| [Identification] | JoinWorkgroup | You may instead specify JoinDomain, DomainAdmin, and DomainAdminPassword. |
| [LicenseFilePrintData] | AutoMode | |
| | AutoUsers | Requires **AutoUsers** if AutoMode = PerServer. |
| [Networking] | | Required to configure network protocols. |
| [UserData] | ComputerName | |
| | FullName | |
| [GuiRunOnce] | | Required to install the cluster service. |

**[Unattended]**

This section contains entries for running Setup (such as whether to extend the partition on which you install the Windows operating system), and specifies location of files necessary for installation during Setup etc. You must instruct Setup to accept the EULA if you want your installation to be fully automated:

```
[Unattended]
OemSkipEula = Yes
```

**[GuiUnattended]**
This section of Sysprep.inf contains entries for preparing the graphical user interface (GUI) for unattended Setup. To completely automate operating system setup you need to include the following entries:

```
[GuiUnattended]
AutoLogon = Yes
AdminPassword = *** <replace the asterisks with the actual password>
OemSkipWelcome = 1
OemSkipRegional = 1
TimeZone = 20 <specify your time zone here>
```

**[Identification]**
This section contains entries for specifying the network identification of a computer. If these entries are not present, Setup adds the computer to the default workgroup called WORKGROUP. If there is not enough information for this entry, Setup will prompt the end user to provide this information.

Prior to adding the first node to a cluster or adding additional nodes, all nodes must have joined a domain (instead of a workgroup) and be members of the same domain. When you create a cluster or join nodes to a cluster, you specify the domain user account under which the cluster service runs. Therefore, you need to specify in this section which domain a destination computer should join, and provide the valid credentials under which to join:

```
[Identification]
DomainAdmin = your_domain\user
DomainAdminPassword = ***** <replace the asterisks with the actual password>
JoinDomain = your_domain
```

**[LicenseFilePrintData]**
This section contains entries for installing licensing information for the Windows Server 2003 family:

```
[LicenseFilePrintData]
AutoMode = PerSeat
```

**[Networking]**
This section contains no entries. However, to configure network protocols during an unattended setup, you must include the **[Networking]** section name in your answer file. In addition, you must include additional sections that describe your network adapters, protocols to be installed, and static IP address information.

To adhere to Server Clusters Best Practices, you must follow the following guidelines when configuring networking:

Use at least two interconnects. Although a server cluster can function with only one interconnect, at least two interconnects are necessary to eliminate a single point of failure and are required for the verification of original equipment manufacturer (OEM) clusters.

Reserve one network exclusively for internal node-to-node communication (the private network).

Do not use teaming network adapters on the private networks.

Use static IP addresses for each network adapter on each node.

For a complete list of instructions for Networking Best Practices, please refer to Windows Server 2003 Enterprise Server Online Help/Availability and Scalability/Cluster Servers.

Use the following entries in unattend.txt to implement the Networking Best Practices recommendations:

```
[NetAdapters]
;list your network adapters here
Adapter01 = params.Adapter01
Adapter02 = params.Adapter02
```

Next, you need to uniquely identify each network adapter. You can either use the INFID entry where you need to specify the Plug and Play IDs of your network adapters, or the NetCardAddress entry to specify MAC addresses. For computers with multiple adapters of the same type (that is, the Plug and Play ID for each is the same), you must specify the **NetCardAddress**

or PCI location information. If you do not specify the **NetCardAddress** or PCI location, only the first adapter enumerated/detected that matches the PnP ID receives the answer file entries.

If you specify the **NetCardAddress** or PCI location, Setup does not use the INFID entry because it is the least specific entry.

Therefore, we recommend that you simply identify your network cards by their MAC addresses as follows:

```
[params.Adapter01]
NetCardAddress = 0x0002A5D0510F
[params.Adapter02]
NetCardAddress = 0x000476D5F57D
;Install Client for MS Networks
[NetClients]
MS_MSClient = params.MS_MSClient
[params.MS_MSClient]
;Specify which network protocols to install
[NetProtocols]
MS_TCPIP = params.MS_TCPIP
;TCP/IP properties
[params.MS_TCPIP]
AdapterSections = params.MS_TCPIP.Adapter01, params.MS_TCPIP.Adapter02
```

Now we need to configure both network cards with static IP address information. Following Cluster Networking Best Practices recommendations, for private networks, we need to define the TCP/IP properties for static IP addresses by specifying a class A, B, or C private address and a subnet mask.

For a public cluster network, we need to specify a static IP address, subnet mask, default gateway, DNS and WINS entries.

```
;Adapter01 is used for intra-cluster communication.
[params.MS_TCPIP.Adapter01]
DHCP = No
IPAddress = 172.24.23.91
SpecificTo = Adapter01
SubnetMask = 255.255.0.0
WINS = No
;Adapter02 represents a public network
[params.MS_TCPIP.Adapter02]
DefaultGateway = 157.6.8.1
DHCP = No
IPAddress = 157.5.4.1
SpecificTo = Adapter02
SubnetMask = 255.255.252.0
DNSServerSearchOrder = 157.5.0.1,157.6.0.2
WINS = Yes
WINSServerList = 157.5.0.1,157.5.1.2
```

**[UserData]**
This section contains entries for specifying user settings during Setup:

```
[UserData]
ComputerName = SV-NODE1
FullName = User Name
OrgName = Organization Name
ProductKey = 11111-22222-33333-44444-55555
```

Of course you need to replace all of the names, passwords and the cluster IP address with real ones.

**Additional Answer File entries**
**Vendor-Specific Devices**
Sometimes you need to install vendor-specific device drivers or third-party device drivers that are not available with Windows Server 2003 Enterprise Server (for example, storage controller drivers).

To install drivers that are not shipped as a part of the Deploy.cab on the Windows Server 2003 CD, use the following key under the [Unattended] section in the unattend.txt:

```
[Unattended]
OemPnpDriversPath = folder_1_on_system_drive[;folder_2_on_system_drive]...
```

Specify this setting for the drivers for all of your secondary mass storage devices.

If you need to load drivers for bootable mass-storage devices, you need to add two additional sections to your answer file:

**[MassStorageDrivers]** and **[OEMBootFiles]**

For more information about **[MassStorageDrivers]** and **[OEMBootFiles]**, refer to the *Microsoft Windows XP Preinstallation Guide* located in the Deploy.cab file on the Windows Server 2003 CD. Although the Guide contains instructions on preparing for Windows XP installation, the same instructions apply to Windows Server 2003 Enterprise Edition as well.

**Creating a Master Installation**
There are no cluster-specific items in this step. Please refer to the *Microsoft Windows XP Preinstallation Guide* for detailed instructions on how to create a Master Installation. This document is located in Deploy.cab on the Windows Server 2003 CD.

**Running Winnt32.exe**
You can run Winnt32.exe from a command prompt to start Windows Setup.

For a complete list of the Winnt32.exe command-line options, see the Winnt32.exe Command-Line Options topic in the *Microsoft Windows XP Preinstallation Guide*. This document is located in Deploy.cab on the Windows Server 2003 CD. Although the Guide contains the instructions on preparing for Windows XP installation, the same instructions apply to Windows Server 2003 Enterprise Edition as well.

To install Windows from a network share, boot the computer using your Windows Server 2003 CD (or any other startup media) and then connect to the desired network location. Placing the distribution share and the configuration set on the network adds considerable flexibility to the pre-installation process.

To run Setup using Winnt32.exe:

Prepare an Unattend.txt file as part of building a configuration set as described above.

Start the computer where you want to install Windows.

Click **Start**, click **Run**, and then type: *path_to_winn32***\winnt32 /unattend:***filename*

where:

*path_to_winnt32*

Specifies the path to the location of Winnt32.exe. For example, \\server\share\Lang\Eng\sku\pro\x86\I386.

*filename*

Specifies the name of the answer file, which contains answers to installation questions you want to automate. The value of *filename* is usually Unattend.txt.

**Installing the Cluster Service**
To install the cluster service you must add the [GuiRunOnce] entry to your unttend.txt as follows:

**[GuiRunOnce]**
To create a cluster, add the following entry to the **GuiRunOnce** section:

```
[GuiRunOnce]
Command0 = %windir%\cluster /cluster:SV-CLUSTER /CREATE /NODE:SV-
NODE1 /USER:your_domain\user /PASS:****** /IPADDR:157.5.5.50
```

To add nodes to an existing cluster, add the following entry to the **GuiRunOnce** section:

```
[GuiRunOnce]
Command0 = %windir%\ cluster /cluster:SV-CLUSTER /ADDNODES:SV-NODE2 /PASSWORD:*********
```

# Image-based Installations of Server Clusters

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Image-based installations, also known as disk cloning or simply cloning, are the fastest method of deploying Windows Server 2003 Enterprise Server. Image-based installations are useful if you are deploying computers that have similar hardware and you want to perform quick, clean installations with minimal end user interaction

Cloning Server clusters consists of the following steps:

Building a master installation on a master computer. Building a master installation includes installing and configuring the operating system and any software and drivers you need to include on your disk image.

Preparing the master installation with the Sysprep tool. This includes configuring and running the Sysprep tool on the master computer.

Generating a disk image of the master installation with the disk-imaging tool. This includes saving each disk image to a permanent storage location.

Loading disk images on destination computers that will make up a cluster.

You cannot clone a cluster node with cluster service installed. You must de-install the cluster service or use a specially prepared master computer without cluster services installed to create a disk image.

**Using SYSPREP to Prepare for Imaging**

**Building a Master Installation**

You need to setup a master computer which you will use to build a master installation and later to create a disk image. Make sure your master computer does not have cluster service installed and is not attached to any shared storage.

Building a master installation includes installing and configuring the operating system and any software and drivers you need to include on your disk image.

**Prepare the Master Installation**

After you have prepared a master installation on the master computer, you need to prepare a disk image for cloning. Microsoft provides a tool called sysprep that allows you to prepare an image for cloning. The tool can be found on the **Windows Server 2003 Enterprise Server CD in the \support\tools directory**. When you run sysprep on the master computer, you can supply one of the following two switches: **-factory or reseal**.

You can use Factory mode to install additional drivers and applications at the stage after the reboot that follows Sysprep. Normally, running Sysprep as the last step in the pre-installation process prepares the computer for delivery. When rebooted, the computer starts Mini-Setup. By running Sysprep with the -factory option, the computer reboots in a network-enabled state without starting Mini-Setup. In this state, Factory.exe processes its answer file, Winbom.ini, and performs the following actions:

Copies drivers from a network source to the computer

Starts Plug and Play enumeration

Stages, installs, and uninstalls applications on the computer from source files located on either the computer or a network source

Adds customer data

When finished, run Sysprep with the -reseal option to prepare the disk image for cloning.

For example, you can create a master installation that contains a minimal set of drivers, run sysprep -factory, create an image of the installation, and copy the image to multiple destination computers. If any destination computers require additional drivers, run sysprep -factory on the destination computer and load those drivers at that time.

For more information about the Factory mode and winbom.ini file, refer to the Microsoft Windows XP Preinstallation Guide which can be found on the Windows Server 2003 Enterprise Server CD in the \support\tools\deploy.cab file.

If you do not need to install any additional drivers or applications on the destination computers, run sysprep in the Reseal mode.

**Sysprep and Sysprep.inf Answer File**

To completely automate cluster server installation, you will need to use the sysprep.inf file which is an answer file used by

sysprep during and after Mini-Setup.

**Required Sections and Entries in sysprep.inf**
A fully unattended Setup requires the following sections and entries:

| Required section | Required entries in that section | Comments |
|---|---|---|
| [Unattended] | UnattendMode | |
| | TargetPath | |
| [GuiUnattended] | AdminPassword | |
| | TimeZone | |
| [Identification] | JoinWorkgroup | You may instead specify JoinDomain, DomainAdmin, and DomainAdminPassword. |
| [LicenseFilePrintData] | AutoMode | |
| | AutoUsers | Requires **AutoUsers** if AutoMode = PerServer. |
| [Networking] | | Required to configure network protocols. |
| [UserData] | ComputerName | |
| | FullName | |

Please refer to the Required Sections and Entries section of this document for an explanation of what each of these entries is used for.

**Additional sysprep.inf Entries**
**Vendor-Specific Devices**
Sometimes you need to install vendor-specific device drivers or third-party device drivers that are not available with Windows Server 2003 Enterprise Server, for example storage controller drivers.

One way to do it is to install the needed drivers on a master computer and then distribute them to destination computers as a part of a disk image if the storage controllers are identical on all cluster nodes. If the storage controllers are different on different cluster nodes, you will need to:

create separate disk images that contain those device drivers or

follow the guidelines below that describe how to use the same image for computers with different storage controllers.

**[SysprepMassStorage] Entry in sysprep.inf**
Before running Sysprep on the master installation and creating your image, identify the different mass-storage controllers that you want to install on each destination computer in the [SysprepMassStorage] section of the Sysprep.inf file. Sysprep then prepopulates the necessary driver information so that Windows can load the correct drivers when the operating system boots on a computer that uses one of the predefined mass-storage controllers.

You do not need to create the entries manually in the **[SysprepMassStorage]** section. Instead, you can use one of the following techniques:

Include the **BuildMassStorageSection** entry in the [Sysprep] section in the Sysprep.inf file. Sysprep will automatically generate the entries in **[SysprepMassStorage]** for all mass-storage controllers specified by the Plug and Play hardware IDs in Machine.inf, Scsi.inf, Pnpscsi.inf and Mshdc.inf and will install those mass-storage controllers.

--OR--

Do not include the **BuildMassStorageSection** entry in the [Sysprep] section in the Sysprep.inf file. Run the command **Sysprep -bmsd**. Sysprep will automatically generate the entries in **[SysprepMassStorage]** for all mass-storage controllers specified by the Plug and Play hardware IDs in Machine.inf, Scsi.inf, Pnpscsi.inf, and Mshdc.inf, but will not install those controllers. You can then delete items from this section before running **Sysprep -reseal** or **Sysprep -factory** on this installation. Installing a smaller number of items in the critical device database reduces the time required for this image to reboot into the operating system.

| Note |
| --- |
| When you run **Sysprep -reseal** or **Sysprep -factory**, Sysprep reads the **[SysprepMassStorage]** section of the Sysprep.inf file and writes the mass storage information into the registry before shutdown. The next time the computer is started, all mass storage controllers are available. |

**Hardware_ID**
Specifies the Plug and Play ID for the device as specified in the devices .inf file.

| **Syntax** | **path_to_device_inf[, disk_directory[, disk_description[, disk_tag]]]** |
| --- | --- |
| Values | path_to_device_inf<br><br>Specifies the path to the .inf file that contains the Plug and Play ID of the controller to install. Enclose *path_to_device_inf* in quotes if it is a long file name.<br><br>disk_directory<br><br>Specifies the name of the directory on the floppy disk provided by the third party that contains the copy of the mass-storage driver.<br><br>disk_description<br><br>Specifies the description of the floppy disk as specified in the Txtsetup.oem file provided by the third party.<br><br>disk_tag<br><br>Specifies the disk tag of the floppy disk as specified in the Txtsetup.oem file provided by the third party. |
| Example | To support a new Qlogic driver, copy the qlogic.inf file to the Sysprep folder and include the following section in the Sysprep.inf file:<br><br>[SysprepMassStorage]<br><br>PCI\VEN_1077&DEV_1080 = "C:\Sysprep\qlogic\qlogic.inf", "\nt", "Qlogic Software Disk", "\qlogic" |

For more information about sysprep, sysprep.inf or **[SysprepMassStorage]** entry refer to the *Microsoft Windows XP Preinstallation Guide* which can be found on the Windows Server 2003 Enterprise Server CD in the \support\tools\deploy.cab file.

If you are using the same image for computers with different storage controllers you will need to include the [SysprepMassStorage] and [BuildMassStorageSection] as described above.

For a complete list of sysprep.inf entries please refer to the *Microsoft Windows XP Preinstallation Guide* which can be found on the Windows Server 2003 Enterprise Server CD in the \support\tools\deploy.cab file.

Refer to Appendix A for a sample sysprep.inf file.

**Assigning Drive Letters to Shared Storage Drives (Optional)**
You may want to control how drive letters are assigned to shared storage drives. You can assign the drive letters using the diskpart.exe utility as shown in an example below. First you need to create a text file with parameters that the diskpart.exe will use as command line parameters, in the example below it is called diskpartparams.txt. The batch file itself is very short, it only calls the diskpart.exe utility with the parameters supplied in the diskpartparams.txt file.

Include the following line in sysprep.inf:

```
[GuiRunOnce]
Command1=%systemdrive%\scripts\AssisgnDriveLetters.bat
```

Make sure you create a directory called scripts (or any other name you like) on the system drive and place the following batch files there:

AssignDriveLetters.bat

```
diskpart /s diskpartparams.txt
pause
diskpartparams.txt
select volume 0
```

```
assign letter Q
select volume 1
assign letter R
select volume 2
assign letter S
select volume 3
assign letter T
exit
```

**Adding Cluster Service Account to the Local Administrators Group**
If you are not installing the cluster service on your destination computers at the time Mini-Setup runs but plan on creating a server cluster later using those nodes, you may want to add a cluster service account to the local administrators group. If you do that, you can later follow the instructions in the Command Line Setup chapter of this document to create a cluster and/or add nodes to an existing cluster.

To add a cluster service account to the Local Administrators group, include the following line in the sysprep.inf file:

```
[GuiRunOnce]
Command0=net localgroup administrators /add domain\clussvc_acct
```

**Generating and Distributing a Disk Image**
**Creating a Disk Image**
After you have prepared a disk image for cloning, you need to create a disk image of your master installation with a disk-imaging tool and save the disk image to a permanent storage location. You can either use a third-party disk imaging software or Microsoft software called iBIG.

If you are using a third-party product, refer to the accompanying documentation on how to create and distribute a disk image.

**Startup Media**
Before you can load disk images on destination computers, you need some kind of startup media to boot your computers from. Startup media contains the system files and device drivers that are necessary to start a computer so that the primary hard disk is accessible but not in use. Startup media might also contain network adapter and network drivers, CD and DVD device drivers, disk configuration tools, and scripts or batch files. You can use a floppy, CD, DVD or network boot as your startup media, depending on capabilities of your destination computers.

If you use third-party disk imaging products, they often provide tools to create different startup media. Otherwise you need to create your own.

Follow these guidelines when creating your startup media:

- Your startup media must provide network support if you are distributing disk images across a network.
- Your startup media must provide CD or DVD device support if you are distributing disk images on media and you are using a floppy disk as your startup media.
- Your startup media must support the tools you need to copy a disk image from a storage location to a destination computer. For example, if your startup media is an MS-DOS boot disk then you need to use MS-DOS tools to copy the disk image onto the destination computer.

For more information about choosing and creating startup media, refer to the Microsoft Windows Server 2003 Corporate Deployment Tools User's Guide (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the support folder on the Windows Server 2003 operating CD.

**Distributing Disk Images**
After an image (or images) has been created and placed on a distribution share (or distribution media such as CD or DVD) and you have a startup media to boot your destination computers, you are ready to distribute the images to destination computers.

You need to make sure that your cluster hardware and networks are set up as described in the Windows Server 2003 Enterprise Server Online Help/Availability and Scalability/Cluster Servers. All of your cluster nodes that you will be installing already have to be connected to the shared storage.

You can load disk images to all of your cluster nodes simultaneously. Many third-party tools support multicast image distribution. You can also use BIG to distribute disk images to your cluster nodes.

When all disk images are loaded on destination computers, do not start or restart them all at once. When creating a cluster, you must be careful to install the first cluster node separately from others. Once the first node is up and running, verify that the

cluster service is up and running as well. Then you can install the rest of the cluster nodes. Since Server clusters use shared storage, corruption of data on shared storage may occur if shared storage is not protected by the cluster service.

After you have distributed disk images to destination computers, sysprep runs Mini-Setup. After Mini-Setup finishes, you should verify that all of the nodes have successfully joined the cluster. Open Cluster Administrator to see which nodes participate in the cluster, and whether everything is up and running. If it is, your cluster is ready.

# Configuring Your Server Clusters

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

After you have installed your Server clusters, most likely you will need to configure it, install additional cluster resources and applications. It is not in the scope of this document to explain how to do these tasks. Please refer to Windows Server 2003 Enterprise Server Online Help/Availability and Scalability/Cluster Servers for more information about configuring Server clusters and cluster applications.

**Using Scripts**
**Using MSClus and/or WMI to Create Resources**
Here we provide two examples of how you can further automate tasks during unattended installation of clusters.

**Using Cluster Automation Server and Windows Script Host**
By adding the following command to the sysprep.inf file, you can create a file share after you have created a cluster or added a node to an existing cluster.

```
[GuiRunOnce]
Command1 = "%systemdrive%\ClusterInstallFiles\createfs.vbs SV-
CLUSTER ClusterGroup SVFileShareResource E:\ SVFileShare"
```

When this command executes, a file share called SVFileShare will be created on the E: drive on SV-CLUSTER.

You must place a Visual Basic script called createfs.vbs in a directory ClusterInstallFiles on your system drive. Please refer to Appendix B for a sample code for such script.

In a similar fashion, you can automate the creation of various resources such as IP address resources, network name, print shares, generic resources etc. For more information about using Cluster Automation Server to create Cluster resources, refer to the Cluster Automation Server Reference in the Windows Platform SDK.

**Using the Command Line**
See Appendix C for a sample batch file that can be used to create a File Share resource with cluster.exe. Instruct Mini-Setup to run this batch file (here called createfs.bat) after Mini-Setup finishes by using the GuiRunOnce entry:

```
[GuiRunOnce]
Command2=%systemdrive%\scripts\createfs.bat
```

# Image-based Installation of the Operating System and the Cluster Service Using Automated Deployment Services (ADS)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Automated Deployment Services (ADS) enables you to personalize the operating system image to your needs. Imaging/Cloning requires you to modify the sysprep.inf (file that contains the configuration information for operating system) every time you need to deploy the image to a different server or deploy a different configuration of the operating system. ADS reduces this complexity by allowing the user to personalize the sysperp.inf file. For more information about sysprep.inf and imaging please refer to section Image based installation of Server Cluster.

In addition, ADS enables you to use a single server to manage servers in your data center. This server, or "Controller," together with the all of the other ADS services, enables you to deploy operating system images onto devices without an operating system or to repurpose existing devices with new operating system images. In short, you can use ADS to:

Remotely purpose a device that has no operating system to a useful state or repurpose a device from one state to another state.

Operate a scale-out data center by providing ability to run extensible and configurable operations, such as scripts, on one or more systems from a single point of control.

Automated Deployment Services consists of the following three main services:

- Controller service

- Image Distribution service

- Network Boot Services

*Refer to ADS help and documentation for more information about ADS and ADS setup and installation*. This section assumes that you have ADS installed and running.

**Following are the steps to deploying Server Clusters using ADS**

1. Add devices (systems that you want to install cluster service to) to ADS Controller

2. Create a Master Image

3. Upload the image to ADS controller

4. Modify sysprep.inf to include Variables

5. Create task sequence file for Image deployment

6. Create a job template for the sequence file

7. Create and store variables associated value in ADS database for the desired device

8. Execute Job Template against the desired device

**Caution**

If you are using shared storage device then it is vital that only one node have access to the shared disk. Otherwise the shared disk can become corrupted. To prevent corruption either shutdown all but one node or use some other techinique(LUN, zoning etc) to isolate the datapath for the nodes. Once a single node cluster is formed then you can add and join other nodes to this cluster. Here are the key things you need be aware and plan accordingly before deploying clusters.

- The disks must be preformatted and ready to mount when the machine boots up in to operating system.

- No other machine should be in full operating system at the time when the first node is creating the cluster ( this is to avoid any shared disk corruption).

- Join Node(adding another node to the cluster) task sequence must be run only after a cluster has been formed . This is to avoid nodes joining simultaneously when the cluster creation is still in progress.

**1. Add Devices (Systems that You Want to Install Cluster Service to) to ADS Controller:**
After installation and configuration of ADS, you would need to add device(system where you want to deploy operating system and configure Server Cluster) to the controller. You can add the device through UI or through command line interface. *Refer to ADS Help section Manage Devices for more information*. You need to make sure the device is set to Network boot in the BIOS. Once the device boots up, it communicates with Controller and boots into deployment agent. This device will show-up in the ADS controller and you should take control of the device. You can control the device through UI (ads.msc) or through command-line interface.

To add a device through command line interface:

**adsdevice** [**/s** *controllername* [**/u** *username* **/w** *password*]] **/add** *devicename* [**/description** "*text*"] [ **/ip** *ipaddress*] [**/adminmac** *macaddress*] [**/guid** *smbiosguid*] [**/assettag** *assettag*] [**/jobtemplate** *templatename*] [**/?**]

or

**adsdevice** [**/s***controllername* [**/u***username***/w** *password*]] **/add @***filename*

E.g.

**Adsdevice** /**add** mnhp6n11 /**description** "server cluster 1st node" /**ip** 172.24.11.207 /**adminmac** 00306E1215b8 /**guid** 239d1df6-1dd6-11b2-f0e7-00306e12 /**jobtemplate** servercluster.xml

**To add a device through ADS.msc (UI interface):**

Open Automated Deployment Services.




In the console tree, right-click the **Devices** folder and click **Add**.

In the **Add Device** dialog box, type the device name in the **Name** box.




If you want to add more information to help identify the device, type the information in the remaining boxes, and then click **OK**.

**2. Create a Master Image:**

Creating Master image consists of following steps:

- Building a master installation on a master computer. Building a master installation includes installing and configuring the operating system and any software and drivers you need to include on your disk image.

- Preparing the master installation with the Sysprep tool. This includes configuring and running the Sysprep tool on the master computer.

- Generating a disk image of the master installation with the disk-imaging tool. This includes saving each disk image to a permanent storage location.

| **Note** |
| --- |
| You cannot clone a cluster node with cluster service installed. You must de-install the cluster service or use a specially prepared master computer without cluster services installed to create a disk image. |

*Refer to Image based installation of Server Cluster for more information about creating master image.*

**3. Upload the Image to ADS Controller**

After Capturing the image of the operating system you would need to add the captured image to the controller. Create directories named images imagebackup and sequences from the root of the controller. Make sure the volume where you create these directories have enough free disk space to hole at least 2 times the image size. Copy the image to images directory.

Next step is to add the image to the ADS controller database. You can add the image either through User Interface or command line. Following is the syntax for command line to add the image

**adsimage** [**/s** *controllername* [**/u** *username* **/w** *password*]] **/add** *imagename* **/path** *imagepath* [**/description**]

After adding the image, you would need to move the image to the image directory. You can check if the image is been properly added by running **adsimage.exe** again.

*Refer to managing Images section in ADS help for more information about managing image.*

**To add the image through UI Interface**

open ADS.MSC (Automated deployment service)




Right click **Images** and click **Add**




Fill out the Image File Path and Image Name, click **OK**

**4. Modify sysprep.inf to include Variables**

Sysprep.inf file is used by the mini-setup to install and configure the operating system. This files contains the configuration information details for the image that you just added to the ADS controller database. This configuration information is static to the image. ADS allows you to customize sysprep.inf file by allowing you to insert variables in the sysprep.inf file. This way, the configuration information that will change from system to system can be removed from the sysprep.inf file. Instead variable configuration is moved to the ADS controller.

Before modifying the sysprep.inf file from the image, you would need to mount the image to a drive. ADS has tools that will allow you this functionality. Change directory to the images and mount the image using below command line syntax

**imgmount** /mount /w {imagefilename} /d: driveletter

Change drive to the above drive letter and change directory to sysprep directory. There 2 ways to create or modify the sysprep.inf either through **setupmgr.exe** or through **notepad.exe**. Refer to Imaged-Base server cluster installation section for more information about sysprep.inf.

Open the sysprep.inf file in notepad.exe. Locate variables for your installation that changes from a system to system. Define them as variables instead of putting static values in sysprep.inf file. For example, In the following snippet of sysprep.inf file, items marked with * at the beginning of the line can be defined as variable. These settings differ for each system.

```
[params.MS_TCPIP.Adapter01]
*     DHCP="No"
*     IPAddress="10.11.26.11,,172.24.11.141"
    SpecificTo=Adapter01
*     SubnetMask="255.255.0.0,255.255.255.0"
*     WINS="No"
;Adapter02 is used for public network
[params.MS_TCPIP.Adapter02]
    SpecificTo=Adapter02
*     DefaultGateway="172.24.11.1"
*     DHCP="No"
*     IPAddress="172.24.11.205"
```

```
*    SubnetMask="255.255.255.0"
*    DNSServerSearchOrder="172.24.10.2,172.24.0.2"
*    WINS="Yes"
*    WINSServerList="157.55.254.201,157.55.254.203"
[GuiRunOnce]
;see section 2.2.1 in this document for a sample text of AssignDriveLetters.bat file
;see configuring cluster section and appendix B for createfs.vbs
*Command0=%systemdrive%\scripts\AssisgnDriveLetters.bat
*Command1 = "%windir%\system32\cluster.exe /cluster:SV-CLUSTER /CREATE /NODE:SV-
NODE1 /USER:domain\user /PASS:***** /IPADDR: 157.5.5.50"
*Command2 = "%systemdrive%\ClusterInstallFiles\createfs.vbs SV-
CLUSTER ClusterGroup SVFileShareResource E:\ SVFileShare"
```

After replacing these with variables, the snippets will look like the following. Notice that each variable name starts with ^ and ends with ^. They must be embedded within quotes (). If not ADS will not install and configure the operating system properly. Note: In the guirunonce section, above sample snippet replaces the cluster configuration commands to variables also. You can run your own scripts in this section also. After successful installation, Windows will run commands in guirunonce section on the first boot for configuring cluster services or any other services specified.

```
[params.MS_TCPIP.Adapter01]
    DHCP="^DHCP1^"
    IPAddress="^IPAddress1^"
    SpecificTo=Adapter01
    SubnetMask="^SubnetMask1^"
    WINS="^WINS1^"
;Adapter02 is used for public network
[params.MS_TCPIP.Adapter02]
    DefaultGateway="^DefaultGateway2^"
    SpecificTo=Adapter02
    DHCP="^DHCP2^"
    IPAddress="^IPAddress2^"
    SubnetMask="^SubnetMask2^"
    DNSServerSearchOrder="^DNSServerSearchOrder2^"
    WINS="^WINS2^"
    WINSServerList="^WINSServerList2^"
[GuiRunOnce]
    ;Mount all volumes before form/join the cluster
    Command0="^CLUSTER_COMMAND_MOUNT^"
    Command1=^AssignDriveLetters^
    Command2="^CLUSTER_COMMAND^"
    Command3=^FileShare^
```

Once you have modified the sysprep.in file, unmount image using the below command line interface.

**Imgmount** u drive:

ADS also allows you to define the variables through scripts. *Refer to ADS help on managing Images for more information.*

| **Note** |
| --- |
| Appendix F contains a modified and complete sample sysprep.inf. |

You can also use **setupmgr.exe** to create, modify unattended setup file, syspre.inf file with variables. **Run setupmgr.exe through cmd a window.**




choose create new




choose sysprep setup




choose the operating system type that you want to deploy







Above allows you to enter values for the setup. Here instead of providing the actual values, you can provide the variable names that you will define in ADS. For example ^Name^ is a variable whose value will be replaced through ADS.




After completing above, click Run Once and enter the cluster configuration command. In the above example, ^cluster_create_command^ variable is inserted in the sysprep.inf file whose actually value will be defined in ADS. Once the sysprep.inf is created with the variables, save it and copy this sysprep.inf file into the image by mounting the image to a drive.

**5. Create task sequence file for image deployment**

Next step is to create a task sequence file for ADS controller. It is a XML file containing sequence of task for controller to perform against a device or a set of devices. ADS ships with few sample sequence files. You can use any of the sample XML file and modify to your needs. In this file you will define the personalize variables for controller to replace in the sysprep.inf file. In

addition, you will define what other sequence of task controller must run against the device. For example, partition the disk, copy the image, replace the variables in the sysprep.inf file, reboot and check the state of the device etc.

You would need to modify at least the following sections in the sample XML files for it to work properly.

```
<!-- STEP 1 Create a single 4999MB partition on the disk -->
    <task description="Partition the disk">
        <command>/bmonitor/bmpart.exe</command>
        <parameters>
            <parameter>\device\harddisk0</parameter> <!-- selects harddisk0 -->
            <parameter>/init</parameter> <!-- erases all partitions on harddisk0 -->
            <parameter>/C:4999</parameter> <!-- creates a new partition (#1) of size
4999MB -->
            <parameter>/A</parameter>
                 <!-- activate the newly created partition (#1) -->
        </parameters>
    </task>
In the above section (step 1) you need to define the partition size for the disk. In
the above example it will create a partition of 4999MB on drive 0.
    <!-- STEP 2 download images -->
    <task description="Download image">
        <command>/imaging/imgbmdeploy.exe</command>
        <parameters>
            <parameter>3718Plain</parameter> <!-- name of the image to be deployed-->
            <parameter>\device\harddisk0\partition1</parameter> <!-- deploy the image
to partition1  -->
            <parameter>-r</parameter> <!-- specifies deploy mode  -->
            <parameter>-client</parameter> <!-- required parameter  -->
        </parameters>
    </task>
In the above section (step 2) you need specify the image name that you used to add to
controller. In the above 3718Plain used for the image.
    <!-- STEP 3 Personalize the sysprep.inf file -->
    <task description="Set sysprep custom info in the sysprep.inf file">
        <command>/bmonitor/bmstrrep.exe</command>
        <parameters>
            <parameter>\device\harddisk0\partition1\sysprep\sysprep.inf</parameter>
            <parameter>^ProductKey^</parameter> <!--  key(ProductKey to be searched in
sysprep.inf file
            <parameter>"$ProductKey$"</parameter> <!value to be replaced with. Make
sure to Embed it in quotes  -->
            <parameter>^OEMDuplicatorstring^</parameter>
            <parameter>"$OEMDuplicatorstring$"</parameter>
            <parameter>"$CLUSTER_COMMAND_MOUNT$"</parameter>
            <parameter>^CLUSTER_COMMAND^</parameter>
            <parameter>"$CLUSTER_COMMAND$"</parameter>
            <parameter>^AssignDriveLetters^</parameter>
            <parameter>"$AssignDriveLetters$"</parameter>
            <parameter>^FileShare^</parameter>
            <parameter>"$FileShare$"</parameter>
        </parameters>
    </task>
```

In the above section (step 3) you need to specify the variable names that you used in sysprep.inf. Notice that the variables are embedded in ^xxx^ and its accompanied value to be replaced with is in quotes ($xxxx$).

*A sample XML file for deploying a single node clusters is included in the appendix D. This sample file will also create file shares on the disk.*

**6. Create a job template for the sequence file**

Job templates in ADS provides a way to define and store instructions for task that you plan to run against a device (or set of devices) more than once. You can use either the **adsjobtemplate** command-line tool or the ADS snap-in to create a job template. You would need to associate this job template with the XML file that you just created for deploying operating system and configuring server cluster. After you create a job template, you can use it to run a job as often as you want against any device with freedom to modify the variable values. For instance, you can first create a single node cluster using the cluster command with create switch. After the cluster creation, modify this command to addnode to the next system you want to join this cluster after operating system deployment. For this you would only need to change the value associated with IPaddress and cluster command variables.

**To add the jobtemplate using ADS UI interface.** Open ADS.MSC




Right click **Job Templat**e and click **Add**. This will invoke the adding job template wizard that will guide you through the process.

**7. Create and store variables associated value in ADS database for the desired device**

Now you would need to define and store the variables for each device in controller database. These values get replaced in sysprep.inf file when the job is executed against the device. You can define variables through ADS snap-in or command-line interface. To add through ADS Snap-in, double-click the device and on the resulting pop-up click user. Click Add to add variable name that is used in the sequence file and sysprep.inf and associate it with a value.

Open ADS.MSC and double click the appropriate device




Click **User** and then click **Add .**




Enter the variable name and then enter the value for this variable and click ok.

You can also use command-line interface as listed below to add variables to the device database on the controller.

Adsdevice /edit device-name /setvar variable-name "value"

E.g.;

**Adsdevice** /edit mnhp11n3 /setvar DNSServerSearchOrder2 "172.24.10.2,172.24.0.2"

Mnhp11n3 is a device name, DNSServerSearchOrder2 is a variable.

You can create a batch script to add all the variables for a device. Appendix G includes a sample batch script file that defines values for the variables from appendix D sysprep.inf file.

**8. Execute Job Template Against the Desired Device**

Next step is to execute the job templates against the device. To assign a job template to devices manually, use either the **ADS snap-in** or **adsController** command-line tool.

Once the device is connected to the controller, right click the device and click run job. This will invoke a wizard that will take you through the steps of selecting the right jobtemplate to execute.

Once the execution starts, you can monitor the job progress by clicking on running job option on the left pane of ADS snap-in. It lists all the task that you have defined in the sequence XML file. After cluster creation you may want to configure the cluster based on your need. You can use any script to configure cluster in the [guirunonce] section of sysprep.inf file.

You can use the same sequence file, sysprep.inf, image for installing and configuring second node with little modification. You would need to modify the defined values for the variables for the second node(device). For example, cluster command to add a node to this cluster, IP address for the node etc.

# Appendix A (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This is a sample sysprep.inf file that can be used for unattended installation of a server cluster. By using this configuration file, Mini-Setup will create a cluster called SV-CLUSTER with one node SV-NODE1, will assign drive letters to shared storage drives and create a cluster file share called SVFileShare.

All computer names, cluster name, user data, passwords and IP addresses need to be replaced with real ones before you can use this file.

```
[Unattended]
OemSkipEula = Yes
[GuiUnattended]
AutoLogon = Yes
AdminPassword = ***
AdvServerType = Servernt
OemSkipWelcome = 1
OemSkipRegional = 1
TimeZone = 20
[LicenseFilePrintData]
AutoMode=PerSeat
[UserData]
ComputerName = SV-NODE1
FullName = "User"
OrgName = "Organization"
ProductKey = 11111-22222-33333-44444-55555
[Networking]
[Identification]
DomainAdmin = domain\user
DomainAdminPassword = *****
JoinDomain = domain
Two network adapters will be available - Adapter01 and Adapter02.
[NetAdapters]
Adapter01 = params.Adapter01
Adapter02 = params.Adapter02
;Specify the network card used by network adapter No. 1.
[params.Adapter01]
netcardaddress = 0x0002A5D0510f
;Specify the network card used by network adapter No. 2.
[params.Adapter02]
netcardaddress = 0x000476D5F57D
;Install client for MS networks.
[NetClients]
MS_MSClient = params.MS_MSClient
[params.MS_MSClient]
;Only install TCP.
[NetProtocols]
MS_TCPIP = params.MS_TCPIP
[params.MS_TCPIP]
;TCP/IP properties
AdapterSections = params.MS_TCPIP.Adapter01,params.MS_TCPIP.Adapter02
;Adapter01 is used for intra-cluster communication.
[params.MS_TCPIP.Adapter01]
DHCP = No
IPAddress = 172.24.23.91
SpecificTo = Adapter01
SubnetMask = 255.255.0.0
WINS = No
;Adapter02 represents a public network
[params.MS_TCPIP.Adapter02]
DefaultGateway = 157.5.8.1
DHCP = No
IPAddress = 157.5.4.1
SpecificTo = Adapter02
SubnetMask = 255.255.252.0
```

```
DNSServerSearchOrder = 157.5.0.1,157.6.0.2
WINS = Yes
WINSServerList = 157.5.0.1,157.5.1.2
[GuiRunOnce]
;see section 2.2.1 in this document for a sample text of AssignDriveLetters.bat file
Command0=%systemdrive%\scripts\AssisgnDriveLetters.bat
Command1 = "%windir%\system32\cluster.exe /cluster:SV-CLUSTER /CREATE /NODE:SV-
NODE1 /USER:domain\user /PASS:***** /IPADDR: 157.5.5.50"
Command2 = "%systemdrive%\ClusterInstallFiles\createfs.vbs SV-
CLUSTER ClusterGroup SVFileShareResource E:\ SVFileShare"
```

# Appendix B (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

```
CreateFS.vbs - Visual Basic script that creates a cluster file share
Const ResourceType = "File Share"
Const UsageString = "Usage: createfs cluster group resourcename path sharename
<remarks> <sharesubdirs>"
Const BasicArgs = 3 ' First three arguments will always be cluster, group, resourcename
Const MinArgs = 5 ' BasicArgs + path + sharename
Const MaxArgs = 7 ' BasicArgs + path + sharename + <remarks> + <sharesubdirs>
'---------------------------------------------------------------------------
'Define the private properties we'll be working with in PrivArgNameArray
' and define the default values in PrivArgArray
'---------------------------------------------------------------------------
PrivArgNameArray = Array( "Path", "ShareName", "Remark", "ShareSubDirs" )
PrivArgArray = Array( "", "", "", 1 )
'---------------------------------------------------------------------------
' Define all of the resource types that we are dependent upon.  We'll only
' add one of each type in the array, if present, and the first one we
' encounter in our enumeration.  If none are found for one of the types
' no errors are thrown.  It is up to the user to verify that everything
' worked fine (e.g. path was to a shared disk we're dependent upon)
'---------------------------------------------------------------------------
Const NumberOfDeps = 3
DependencyArray = Array( "Physical Disk", "IP Address", "Network Name" )
Set oArgs = WScript.Arguments
If oArgs.Count < MinArgs Or oArgs.Count > MaxArgs Then
    '-----------------------------------------------------------------------
    'Usage
    '-----------------------------------------------------------------------
    Wscript.Echo UsageString
Else
    '-----------------------------------------------------------------------
    'Parse the arguments
    '-----------------------------------------------------------------------
    oClusterName = oArgs(0)
    oGroupName = oArgs(1)
    oResourceName = oArgs(2)
    For iCounter = 0 To oArgs.Count - BasicArgs - 1
        If True = IsNumeric( oArgs( iCounter + BasicArgs )) Then
            PrivArgArray(iCounter) = CLng(oArgs(iCounter + BasicArgs))
        Else
            PrivArgArray(iCounter) = oArgs(iCounter + BasicArgs)
        End If
    Next
    '-----------------------------------------------------------------------
    'Open the Cluster
    '-----------------------------------------------------------------------
    Set oCluster = CreateObject("MSCluster.Cluster")
    oCluster.Open (oClusterName)
    '-----------------------------------------------------------------------
    'Open the resgroup and add the new resource
    '-----------------------------------------------------------------------
    Set oGroup = oCluster.ResourceGroups.Item(oGroupName)
    Set oNewResource = oGroup.Resources.CreateItem(oResourceName, ResourceType, 0)
    '-----------------------------------------------------------------------
    'Add any necessary dependencies
    '-----------------------------------------------------------------------
    Set coldeps = oNewResource.Dependencies
    For iCounter = 0 To NumberOfDeps - 1
        For Each oResource In oGroup.Resources
            If oResource.Type.Name = DependencyArray(iCounter) Then
                coldeps.AddItem oResource
                Exit For
            End If
```

```
        Next
    Next
    '-----------------------------------------------------------------------------
    'Set the private properties for the new resource
    '-----------------------------------------------------------------------------
    Set colProperties = oNewResource.PrivateProperties
    For iCounter = 0 To MaxArgs - BasicArgs - 1
        colProperties.Item(PrivArgNameArray(iCounter)).Value = PrivArgArray(iCounter)
    Next
    colProperties.SaveChanges
    '-----------------------------------------------------------------------------
    'Cleanup
    '-----------------------------------------------------------------------------
    Set oCluster = Nothing
    Set oGroup = Nothing
    Set oResourceName = Nothing
    Set iCounter = Nothing
    Set oNewResource = Nothing
    Set colProperties = Nothing
    Set coldeps = Nothing
End If
'-----------------------------------------------------------------------------
'Cleanup
'-----------------------------------------------------------------------------
Set oArgs = Nothing
Set PrivArgArray = Nothing
Set PrivArgNameArray = Nothing
Set DependencyArray = Nothing
```

# Appendix C (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This is a sample batch file that can be used to create a File Share resource with cluster.exe. We assume that there is a group that already contains the Network Name, and Physical Disk resources. Below you will see an explanation of each line. The example below assumes that you have a group called "Disk Group 1" with a "Disk X:".

```
Creatfs.bat
X:
MD ClusterFS
cluster. res "ClusterFS" /create /group:"Disk Group 1" /type:"File Share"
cluster . res "ClusterFS" /priv path="X:\ClusterFS"
cluster . res "ClusterFS" /priv Sharename=ClusFileShare
cluster . res "ClusterFS" /priv Remark="This is a File Share"
cluster . res "ClusterFS" /prop Description="This is a Clustered Share"
cluster . res "ClusterFS" /priv security=Domain\User,grant,c:security
cluster . res "ClusterFS" /priv ShareSubDirs=1
cluster . res "ClusterFS" /AddDep:"Disk X:"
cluster . res "ClusterFS" /AddDep:"Network Name"
cluster . res "ClusterFS" /On
```

What each line means:

**cluster . res "ClusterFS" /Create /Group:"Disk Group 1" /Type:"File Share"**

This creates the File Share resource in the cluster group named "Disk Group 1".

| **Note** |
| --- |
| You do not want to put anything into, or out of the default Cluster Group. |

**cluster . res "ClusterFS" /priv Path="X:\ClusterFS"**

This sets the local path to the file share to X:\ClusterFS. Verify that the path exists, and that it is accessible.

**cluster . res "ClusterFS" /priv ShareName=ClusFileShare**

This sets the share name for the resource, and is the name that clients will use to connect to.

**cluster . res "ClusterFS" /priv Remark="This is a File Share"**

This sets the remarks for the file share. The remark is what the client will see when browsing the virtual server.

**cluster . res "ClusterFS" /prop Description="This is a Clustered Share"**

This gives the resource a description in Cluster Administrator. This name is only for identification in CluAdmin.

**cluster . res "ClusterFS" /priv security=Domain\User,grant,c:security**

This gives the user account "User" from domain "Domain" change share level rights to the File Share resource.

Access types are grant, deny, set, or revoke. An access mask is required for the types grant, deny, or set. Access masks are "R" for read, "C" for change, and "F" for full.

This can also be used to grant multiple users rights at the same time, here is the syntax for the security switch:

Security=Domain\User1,grant,RC,Domain\User2,revoke,everyone,set,F:security

**cluster . res "ClusterFS" /priv ShareSubDirs=1**

This sets the file share to dynamic (which translates to "share all subdirectories"). You will need to omit this line if you are creating a standard file share.

**cluster . res "ClusterFS" /AddDep:"Disk X:"**

This makes the File Share resource dependent on the Physical Disk resource for disk X:\, on which the file share resides.

**cluster . res "ClusterFS" /AddDep:"Network Name"**

This makes the File Share resource dependent on a Network Name resource. This needs to be an already existing Network Name resource.

**cluster . res "ClusterFS" /On**

This brings the newly created File Share resource online.

# Appendix D (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This is a sample ADS Sequence file that can be used for installation of a server cluster. By using this sequence file in conjunction with [sysprep.inf] in appendix F , Mini-Setup will create a cluster called mnhp11 with one node mnhp11n3, will assign drive letters (will use Creatfs.bat from appendix C) to shared storage drives and create a cluster file share called SVFileShare (will use CreateFS.vbs from appendix B). All variables (anything that is between $xxxx$ such as Product key, domainname, computer names etc) in the <!-- STEP 3 Personalize the sysprep.inf file --> section of the below file needs to be defined as described in Appendix G. In addition, you would need to define cluster name, user data, passwords and IP addresses need to be replaced in the appendix G with real ones before you can use this file.

FileName: [da-deploy-image-form-servercluster.xml]

```
<?xml version="1.0" encoding="utf-8"?>
<!--
    This file is part of the Microsoft ADS Samples
    Copyright (c) Microsoft Corporation. All rights reserved.
    Sample sequence to purpose a device that is currently running the
    Deployment Agent and configure it connect it to a domain.
    The image to be provisioned must have been captured with a sysprep.inf
    containing the strings ^ADS_WINDOWS_PRODUCT_KEY^ , ^DomainAdminPassword^
    etc as in the personalize step below.
    Device variables must have been created for the devices as in the
    personalize step below.
    To use: boot a target into the Deployment Agent, then run this sequence
    against that device.
-->
<!-- start sequence -->
<sequence command="da-deploy-image-form-servercluster.xml" description="Deploy an image
and configure the machine to join domain and form a cluster"
    xmlns="http://schemas.microsoft.com/ads/2003/sequence" version="1">
    <!-- STEP 1 Create a single 4999MB partition on the disk -->
    <task description="Partition the disk">
        <command>/bmonitor/bmpart.exe</command>
        <parameters>
            <parameter>\device\harddisk0\partition0</parameter>
            <parameter>/init</parameter>
            <parameter>/C:4999</parameter>
            <parameter>/A</parameter>
        </parameters>
    </task>
    <!-- STEP 2 download images -->
    <task description="Download image">
        <command>/imaging/imgbmdeploy.exe</command>
        <parameters>
            <parameter>3718Plain</parameter>
            <parameter>\device\harddisk0\partition1</parameter>
            <parameter>-r</parameter>
            <parameter>-client</parameter>
        </parameters>
    </task>
    <!-- STEP 3 Personalize the sysprep.inf file -->
    <task description="Set sysprep custom info in the sysprep.inf file">
        <command>/bmonitor/bmstrrep.exe</command>
        <parameters>
            <parameter>\device\harddisk0\partition1\sysprep\sysprep.inf</parameter>
            <parameter>^ProductKey^</parameter>
            <parameter>"$ProductKey$"</parameter>
            <parameter>^OEMDuplicatorstring^</parameter>
            <parameter>"$OEMDuplicatorstring$"</parameter>
            <parameter>^FullName^</parameter>
            <parameter>"$FullName$"</parameter>
            <parameter>^OrgName^</parameter>
            <parameter>"$OrgName$"</parameter>
```

```
            <parameter>^DomainAdmin^</parameter>
            <parameter>"$DomainAdmin$"</parameter>
            <parameter>^DomainAdminPassword^</parameter>
            <parameter>"$DomainAdminPassword$"</parameter>
            <parameter>^DomainName^</parameter>
            <parameter>"$DomainName$"</parameter>
            <parameter>^System_Admin_password^</parameter>
            <parameter>"$System_Admin_password$"</parameter>
            <parameter>^COMPUTER_NAME^</parameter>
            <parameter>"$Device.System.Name$"</parameter>
            <parameter>^ConnectionName1^</parameter>
            <parameter>"$ConnectionName1$"</parameter>
            <parameter>^netcardaddress1^</parameter>
            <parameter>"$netcardaddress1$"</parameter>
            <parameter>^DHCP1^</parameter>
            <parameter>"$DHCP1$"</parameter>
            <parameter>^IPAddress1^</parameter>
            <parameter>"$IPAddress1$"</parameter>
            <parameter>^SubnetMask1^</parameter>
            <parameter>"$SubnetMask1$"</parameter>
            <parameter>^WINS1^</parameter>
            <parameter>"$WINS1$"</parameter>
            <parameter>^ConnectionName2^</parameter>
            <parameter>"$ConnectionName2$"</parameter>
            <parameter>^netcardaddress2^</parameter>
            <parameter>"$netcardaddress2$"</parameter>
            <parameter>^DHCP2^</parameter>
            <parameter>"$DHCP2$"</parameter>
            <parameter>^IPAddress2^</parameter>
            <parameter>"$IPAddress2$"</parameter>
            <parameter>^SubnetMask2^</parameter>
            <parameter>"$SubnetMask2$"</parameter>
            <parameter>^DefaultGateway2^</parameter>
            <parameter>"$DefaultGateway2$"</parameter>
            <parameter>^DNSServerSearchOrder2^</parameter>
            <parameter>"$DNSServerSearchOrder2$"</parameter>
            <parameter>^WINSServerList2^</parameter>
            <parameter>"$WINSServerList2$"</parameter>
            <parameter>^WINS2^</parameter>
            <parameter>"$WINS2$"</parameter>
            <parameter>^CLUSTER_COMMAND_MOUNT^</parameter>
            <parameter>"$CLUSTER_COMMAND_MOUNT$"</parameter>
            <parameter>^CLUSTER_COMMAND^</parameter>
            <parameter>"$CLUSTER_COMMAND$"</parameter>
            <!-- <parameter>^AssignDriveLetters^</parameter> -->
            <!-- <parameter>"$AssignDriveLetters$"</parameter> -->
            <!-- <parameter>^FileShare^</parameter> -->
            <!-- <parameter>"$FileShare$"</parameter> -->
        </parameters>
    </task>
    <!-- STEP 4 set Controller IP -->
    <task description="Set controller IP address">
        <command>/bmonitor/bmsetreg.exe</command>
        <parameters>
            <parameter>-
h:\device\harddisk0\partition1\windows\system32\config\system</parameter>
            <parameter>controlset001\control\bmss</parameter>
            <parameter>controlleripaddress</parameter>
            <parameter>reg_multi_sz</parameter>
            <parameter>$controller.system.adminipaddr$</parameter>
        </parameters>
    </task>
    <!-- STEP 5 set BMDP port number -->
    <task description="Set BMDP port number">
        <command>/bmonitor/bmsetreg.exe</command>
        <parameters>
            <parameter>-
h:\device\harddisk0\partition1\windows\system32\config\system</parameter>
            <parameter>controlset001\control\bmss</parameter>
            <parameter>bmdpipport</parameter>
            <parameter>reg_dword</parameter>
            <parameter>$controller.system.bmdpport$</parameter>
```

```
        </parameters>
    </task>
    <!-- STEP 6 set BMCP port number -->
    <task description="Set BMCP port number">
        <command>/bmonitor/bmsetreg.exe</command>
        <parameters>
            <parameter>-
h:\device\harddisk0\partition1\windows\system32\config\system</parameter>
            <parameter>controlset001\control\bmss</parameter>
            <parameter>bmcpipport</parameter>
            <parameter>reg_dword</parameter>
            <parameter>$controller.system.bmcpport$</parameter>
        </parameters>
    </task>
    <!-- STEP 7 set management MAC address -->
    <task description="Set management MAC address">
        <command>/bmonitor/bmsetreg.exe</command>
        <parameters>
            <parameter>-
h:\device\harddisk0\partition1\windows\system32\config\system</parameter>
            <parameter>controlset001\control\bmss</parameter>
            <parameter>bindexcept</parameter>
            <parameter>reg_multi_sz</parameter>
            <parameter>$controller.system.adminmac$</parameter>
        </parameters>
    </task>
    <!-- STEP 8 copy public key to target -->
    <task description="Copy public key certificate file to target">
        <command>/bmonitor/bmfilexfer.exe</command>
        <parameters>
            <parameter>-d</parameter>
            <parameter>"c:\program files\Microsoft ADS\certificate\adsroot.cer"
</parameter>
            <parameter>\device\harddisk0\partition1\windows\temp\adsroot.cer</parameter
>
        </parameters>
    </task>
    <!-- STEP 9 set BMSS public key certificate  -->
    <task description="Set the BMSS public key certificate">
        <command>/bmonitor/bmsetreg.exe</command>
        <parameters>
            <parameter>-
h:\device\harddisk0\partition1\windows\system32\config\system</parameter>
            <parameter>-f</parameter>
            <parameter>-r</parameter>
            <parameter>controlset001\control\bmss\bmcpcertificates</parameter>
            <parameter>AgentCert</parameter>
            <parameter>reg_binary</parameter>
            <parameter>\device\harddisk0\partition1\windows\temp\adsroot.cer</parameter
>
        </parameters>
    </task>
    <!-- STEP 10 Reboot the machine -->
    <task doesReboot="true" timeout= "600" description="Reboot">
        <command>/bmonitor/reboot</command>
    </task>
    <!-- STEP 11 When PXE request comes in, boot to hard disk (starts mini-setup) -->
    <task doesReboot="true"  timeout= "600" description="Boot to hard disk">
        <command>/pxe/boot-hd</command>
    </task>
    <!-- STEP 12 At end of mini-setup, machine will reboot, so when PXE request comes
in
    boot to the hard disk (start the newly installed operating system)
    -->
    <task timeout= "600" description="Boot to hard disk">
        <command>/pxe/boot-hd</command>
    </task>
    <!-- STEP 13 update device record to always boot to hard disk in the future -->
    <!-- ADS Install folder needs to be updated, if you have installed ADS to a folder
-->
    <!-- other than "C:\program files\Microsoft ads" -->
    <task description="Set default job template as boot to hard disk">
```

```
        <command target="controller">C:\Program Files\Microsoft
ADS\tools\adsdevice.wsf</command>
        <parameters>
            <parameter>-e</parameter>
            <parameter>$Device.System.Name$</parameter>
            <parameter>-jobtemplate</parameter>
            <parameter>boot-to-hd</parameter>
        </parameters>
    </task>
</sequence>
```

# Appendix E (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

```
[boot-to-hd.xml]
<?xml version="1.0" encoding="utf-8"?>
<!--
    This file is part of the Microsoft ADS Samples
    Copyright (c) Microsoft Corporation. All rights reserved.
    Use this in a device job template to cause the device to
    boot exit PXE and boot from the next boot device
    (typically the hard disk) on next PXE boot
-->
<!-- start sequence -->
<sequence command="boot-to-hd.xml" description="PXE boot into hard disk"
  version="1"
  xmlns="http://schemas.microsoft.com/ads/2003/sequence">
    <!-- STEP 1 boot to hard disk -->
    <task description="Boot to hard disk">
        <command>/pxe/boot-hd</command>
    </task>
</sequence>
```

# Appendix F (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Following is the sysprep.inf file that you can use to install Server cluster using ADS. You can use with job sequence defined in appendix D and E and along with defined variables in appendix G. Note: You will have to modify the variables for your need.

```
[sysprep.inf]
;SetupMgrTag
[Unattended]
    OemSkipEula=Yes
    InstallFilesPath=C:\sysprep\i386
    TargetPath=\WINDOWS
[GuiUnattended]
    AdminPassword="^System_Admin_password^"
    EncryptedAdminPassword=NO
    OEMSkipRegional=1
    OEMDuplicatorstring="^OEMDuplicatorstring^"
    TimeZone=4
    OemSkipWelcome=1
    AutoLogon=Yes
    AutoLogonCount=1
[UserData]
    FullName="^FullName^"
    OrgName="^OrgName^"
    ComputerName="^COMPUTER_NAME^"
    ProductKey="^ProductKey^"
[LicenseFilePrintData]
    AutoMode=PerSeat
[Identification]
    DomainAdmin="^DomainAdmin^"
    DomainAdminPassword="^DomainAdminPassword^"
    JoinDomain="^DomainName^"
[Networking]
;    InstallDefaultComponents=Yes
;Two network cards will be available -- Adapte401 and Adapetr02
[NetAdapters]
    Adapter01=params.Adapter01
    Adapter02=params.Adapter02
;Specify the network card used by network adapter # 1
[params.Adapter01]
    ConnectionName="^ConnectionName1^"
    netcardaddress="^netcardaddress1^"
;Specify the network card used by network adapter # 2
[params.Adapter02]
    ConnectionName="^ConnectionName2^"
    netcardaddress="^netcardaddress2^"
;Install client for MS Networks
[NetClients]
    MS_MSClient=params.MS_MSClient
[params.MS_MSClient]
;only TCPIP
[NetProtocols]
    MS_TCPIP=params.MS_TCPIP
[params.MS_TCPIP]
;TCPIP Properties
    AdapterSections=params.MS_TCPIP.Adapter01,params.MS_TCPIP.Adapter02
;Adapter01 is used for intra-cluster communication
[params.MS_TCPIP.Adapter01]
    DHCP="^DHCP1^"
    IPAddress="^IPAddress1^"
    SpecificTo=Adapter01
    SubnetMask="^SubnetMask1^"
    WINS="^WINS1^"
;Adapter02 is used for public network
[params.MS_TCPIP.Adapter02]
```

```
    DefaultGateway="^DefaultGateway2^"
    SpecificTo=Adapter02
    DHCP="^DHCP2^"
    IPAddress="^IPAddress2^"
    SubnetMask="^SubnetMask2^"
    DNSServerSearchOrder="^DNSServerSearchOrder2^"
    WINS="^WINS2^"
    WINSServerList="^WINSServerList2^"
[TapiLocation]
    CountryCode=1
    Dialing=Tone
    AreaCode=425
[Branding]
    BrandIEUsingUnattended=Yes
[Proxy]
    Proxy_Enable=0
    Use_Same_Proxy=0
[GuiRunOnce]
    ;Mount all volumes before form/join the cluster
    Command0="^CLUSTER_COMMAND_MOUNT^"
    Command1=^AssignDriveLetters^
    Command2="^CLUSTER_COMMAND^"
    Command3=^FileShare^
[sysprepcleanup]
```

# Appendix G (Server Clusters: Remote Setup, Unattended Installations and Image-based Installations)

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

ADS requires the user to enter these variables using the ADS snap-in. A user need to defined all the below variable for the device (system where the image will be deployed) before executing the XML sequence job template for cluster installation. You can also create a batch file to create these variables using the below listed command. Execute the batch file before running the sequence template.

```
List of Variables with their associated values:OEMDuplicatorstring= Fresh install of
Windows 2003 server with server cluster configured
ProductKey=xxxxx-yyyyy-xxxxx-yyyyy-xxxxx
FullName="wolfpack"
OrgName="Microsoft"
DomainAdmin=<DOMAIN>\yyyyyy
DomainAdminPassword=xxxxx
DomainName=<DOMAIN>.com
System_Admin_password=xxxxxx
ConnectionName1="Private"
Netcardaddress1=0x00306E121879
DHCP1=No
IPAddress1=10.11.25.14
SubnetMask1=255.255.0.0
WINS1=No
ConnectionName2="Public"
Netcardaddress2=0x00306E121878
DefaultGateway2=172.24.11.1
DHCP2=No
IPAddress2=172.24.11.159
SubnetMask2=255.255.255.0
DNSServerSearchOrder2=172.24.10.2,172.24.0.2
WINS2=Yes
WINSServerList2=157.55.254.201,157.55.254.203
CLUSTER_COMMAND_MOUNT= c:\windows\system32\mountvol /e
CLUSTER_COMMAND=c:\windows\system32\cluster.exe /cluster:mnhp11 /create /node:mnhp11n3
/user:ct\wolfpack /pass:wolfpack /ipaddr:172.24.11.141
AssignDriveLetters=%systemdrive%\scripts\AssisgnDriveLetters.bat
FileShare=%systemdrive%\scripts\createfs.bat Batch File:
This batch file uses Adsdevice command to create the variables. Syntax
Adsdevice edit devicename setvar variablename valueforVariable
[Variables.bat]
Adsdevice /edit mnhp11n3 /setvar OEMDuplicatorstring "Fresh install of Windows 2003
server with server cluster configured"
Adsdevice /edit mnhp11n3 /setvar ProductKey "xxxxx-yyyyy-xxxxx-yyyyy-xxxxx
Adsdevice /edit mnhp11n3 /setvar FullName "wolfpack"
Adsdevice /edit mnhp11n3 /setvar OrgName "Microsoft"
Adsdevice /edit mnhp11n3 /setvar DomainAdmin "y\xxxxx"
Adsdevice /edit mnhp11n3 /setvar DomainAdminPassword "xxxxx"
Adsdevice /edit mnhp11n3 /setvar DomainName "<DOMAIN>.com"
Adsdevice /edit mnhp11n3 /setvar System_Admin_password "xxxxxx"
Adsdevice /edit mnhp11n3 /setvar ConnectionName1 "Private"
Adsdevice /edit mnhp11n3 /setvar Netcardaddress1 "0x00306E121879"
Adsdevice /edit mnhp11n3 /setvar DHCP1 "No"
Adsdevice /edit mnhp11n3 /setvar IPAddress1 "10.11.26.11"
Adsdevice /edit mnhp11n3 /setvar SubnetMask1 "255.255.0.0"
Adsdevice /edit mnhp11n3 /setvar WINS1 "No"
Adsdevice /edit mnhp11n3 /setvar ConnectionName2 "Public"
Adsdevice /edit mnhp11n3 /setvar Netcardaddress2 "0x00306E121878"
Adsdevice /edit mnhp11n3 /setvar DefaultGateway2 "172.24.11.1"
Adsdevice /edit mnhp11n3 /setvar DHCP2 "No"
Adsdevice /edit mnhp11n3 /setvar IPAddress2 "172.24.11.205"
Adsdevice /edit mnhp11n3 /setvar SubnetMask2 "255.255.255.0"
Adsdevice /edit mnhp11n3 /setvar DNSServerSearchOrder2 "172.24.10.2,172.24.0.2"
Adsdevice /edit mnhp11n3 /setvar WINS2 "Yes"
Adsdevice /edit mnhp11n3 /setvar WINSServerList2 "157.55.254.201,157.55.254.203"
Adsdevice /edit mnhp11n3 /setvar CLUSTER_COMMAND_MOUNT "c:\windows\system32\mountvol
```

```
/e"
Adsdevice /edit mnhp11n3 /setvar CLUSTER_COMMAND "c:\windows\system32\cluster.exe
/cluster:mnhp11 /create /node:mnhp11n3 /user:ct\wolfpack /pass:wolfpack
/ipaddr:172.24.11.141"
Rem ;see section 2.2.1 in this document for a sample text of AssignDriveLetters.bat
file
Adsdevice /edit mnhp11n3 /setvar AssignDriveLetters
"%systemdrive%\scripts\AssisgnDriveLetters.bat"
Rem ;see configuring cluster section and appendix B for createfs.vbs
Adsdevice /edit mnhp11n3 /setvar FileShare "%systemdrive%\scripts\createfs.bat"
```

# Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: January 1, 2003

**In This White Paper**

Introduction (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

Benefits

Requirements (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

Limitations

How Rolling Upgrades Work (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

Rolling Upgrade Walkthrough

Known Issues

Troubleshooting (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

Conclusion (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

# Introduction (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

One of the features of Microsoft Windows Server 2003 Enterprise Edition and Datacenter Edition operating system is the ability to perform a rolling upgrade of the operating system on a server cluster. A rolling upgrade is a process of upgrading cluster nodes, one node at a time, in such a way that services and resources offered by the cluster are always available through out the upgrade process. The system downtime associated with the upgrade is reduced to a few minutesthe time needed to move resources from one node to anotheras compared to the few hours that are usually needed to upgrade a Windows based server.

An administrators can perform rolling upgrade from Windows 2000 Advance Server to Windows Server 2003 Enterprise Edition or from Windows 2000 Datacenter Sever to Windows Server 2003 Datacenter Edition. Rolling upgrade from Windows NT version 4.0 Enterprise edition directly to Windows Server 2003 Enterprise Edition is not supported. This means if you want to upgrade from Windows NT version 4.0 to Server 2003 than you have two options:

- You can maintain cluster availability by performing an rolling upgrade to Windows 2000 Advanced Server first, then to Windows Server 2003 Enterprise Edition.

- You can perform a non-rolling upgrade directly from Windows NT 4.0 to Windows Server 2003 Enterprise Edition, which will not allow you to maintain cluster availability.

A mixed-version cluster, a cluster composed of nodes running different versions of the operating system, offers the same level of availability as a homogeneous cluster. Resources that support rolling upgrades can be moved, as well as fail over and fail back, between nodes in a mixed mode cluster.

The same procedure can be used to do a rolling upgrade of system hardware or applications if they support it. However, this paper focuses only on the rolling upgrade of the operating system. If you are interested in doing a rolling upgrade of an application, check application documentation to see if the application supports rolling upgrade.

This paper assumes that the reader has a basic knowledge of clustering and is familiar with Microsoft Server Cluster in Windows 2000 operating system and in Windows Server 2003. For more information about these components, consult the Windows Server 2003 Advanced Server Online Help.

**Terminology**
The node running an older version of the operation system, application or hardware is called a *down-level* node.

The node running a newer version of the operation system, application or hardware is called an *up-level* node.

The cluster composed of nodes running different revisions of the operating system, application, or hardware, is called a *mixed-version cluster*.

# Benefits

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A rolling upgrade offers multiple benefits and can be an alternative for any mission-critical system that requires high availability.

Consider two examples: Jack and Jill. Jack runs a mission-critical database application on a stand-alone Windows Server 2003based server. He regularly applies Windows Server 2003 operating system service packs (every quarter), upgrades his application once a year and performs a maintenance hardware upgrade once a year as well. From past experience, he can tell that the service pack installation takes on average of 60 minutes. An application upgrade generally takes four hours. Hardware upgrades take anywhere from 30 minutes to four hours. He can also testify that once a year, something goes wrong and an upgrade takes four times as long as planned. Total downtime per year of this system averages:

4 x 60 + 240 + 240 + 240 = 16 hours, which equals 99.8 percent uptime.

Jill runs a mission critical e-mail application on a two-node cluster. She applies the same rules as Jack, but instead of upgrading both nodes at once, she performs *rolling upgrades*. She measured that it takes five minutes on average to move this application from one node to another; thus, the downtime associated with each upgrade is limited to five minutes. Total downtime per year for his system averages:

4 x 5 + 5 + 5 = 30 minutes, which equals 99.99 percent uptime.

Rolling upgrades are advantageous in that they:

- **Minimize downtime.** Rolling upgrades minimize downtime associated with software or hardware upgrades.

- **Minimize risk.** Rolling upgrades minimize the risk of losing the service in case the upgrade fails. When an upgrade of one node fails, the other node can still provide the service, giving the system administrator the choice to repair or replace a failed node without incurring any additional downtime.

- **Increase flexibility.** The nearly negligible system downtime caused by a rolling upgrade means that administrators could decide to perform a rolling upgrade during a working day instead of performing it late at night or on weekends.

# Requirements (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Rolling upgrade from Windows 2000 to Windows Server 2003**

- To perform a rolling upgrade of the operating system to Windows Server 2003 operating system on cluster nodes, you must start with Windows 2000 Advanced server (including any service packs for example SP1, SP2 etc) or Datacenter Server. Windows Server 2003 operating system does not support rolling upgrade from earlier Windows operating system, for example Windows NT 4.0.

- You can perform rolling upgrade from Windows 2000 Advanced Server to Windows Server 2003 Enterprise Edition.

- You can perform rolling upgrade from Windows 2000 Datacenter Server to Windows Server 2003 Datacenter Edition.

**Upgrade from Windows NT 4.0 to Windows Server 2003**

- Perform rolling upgrade to Windows 2000 Advanced Server first, then to Windows Server 2003 Enterprise Edition.

- Perform a non-rolling upgrade directly from Windows NT 4.0 to Windows Server 2003 Enterprise Edition, which will not allow you to maintain cluster availability.

**Upgrade from Windows NT 4.0 to Windows 2000**

- A rolling upgrade can be used to upgrade Windows NT 4.0 Server, Enterprise Edition SP4 to SP5, SP5 to SP6, and so forth. It can also be used to upgrade any Enterprise Edition service pack later than SP3 to Windows 2000 Advanced Server.

- Upgrades from Windows NT 4.0 to Windows 2000 Datacenter Server are not supported. Windows 2000 Datacenter Server can be installed in a clean install process only.

**Upgrade from Windows Server 2003 (64-bit) Enterprise Edition LE release to Windows Server 2003 (64-bit) Enterprise Edition**

- A rolling upgrade from any version of Windows Server 2003 (64-bit) Enterprise Edition LE release to Windows Server 2003 (64-bit) Enterprise Edition is not supported. You will have to perform either a fresh install or upgrade of your system to Windows Server 2003 (64-bit) Enterprise Edition. This mean the cluster will not be available during upgrade/installation.

# Limitations

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A rolling upgrade is not as disruptive as a regular upgrade, but you should bear in mind that it requires the applications to be moved between nodes and therefore does cause some minimal disruption in services. When an application is moved from one node to another, it must be stopped. Once the application is stopped, Cluster service moves all the resources the application uses, such as disks, Internet Protocol (IP) addresses or network names to another node, and restarts the application on another node. Any sessions between clients and the server application are cancelled during this process; database transactions are aborted and file handles are invalidated. Client applications can retry and eventually reconnect and recover once the server application is restarted on the second node. While the impact may be minimal, it should not be ignored.

The availability of a two-node cluster while one node is being upgraded is limited. Any failure of the second node while the other node is being upgraded will cause the cluster to fail.

During a rolling upgrade of a 2-node cluster, all resource groups are moved to one node. In a cluster where resources are distributed among all nodes in the cluster, this means that the node hosting all resources can run at maximum capacity while the other node is being upgraded. This can affect the application response time.

For a rolling upgrade of Windows 2000 Datacenter Server cluster with more than 2 nodes, a user will have different options to move the resources from the node that is upgrade to other nodes. Depending on the configuration and node capacity, you may want to move resources to a node that can handle the additional load from the node that is being upgraded. You can also distribute the resources from upgraded node to different nodes to reduce the performance impact. Chart out which resources will be moved to which node/nodes during each node upgrade before proceeding with rolling upgrade.

| **Note** |
| --- |
| Windows Server 2003 Datacenter Edition can only be obtained from an OEM. Please contact your OEM vendor for Windows Server 2003 Datacenter Edition availability. |

It is critical that you plan rolling upgrade of your cluster up front. You may have to plan differently if you decide to run mixed-version cluster for couple of days before upgrading the remaining downlevel nodes to Windows Server 2003. Please list out services and applications that you run on the cluster. Figure out what services or applications can be failed over (and failed back) to Windows Server 2003 from Windows 2000 operating system and vice versa. This information will be useful if you have to failover or move resources between nodes that are running different version of operating system. Check application documentation for the mixed-version cluster support requirements.

Cluster service guarantees that an up-level node can always join a cluster formed by a down-level node. It also guarantees that resources can fail over and fail back to and from a down-level node. However, it does not guarantee that applications can also support rolling upgrades of the operating system. Table 1 below summarizes the behavior of cluster-aware resources supported by Cluster service.

**Table 1 Resources supported/not supported during rolling upgrades.**

| Resource | Note |
| --- | --- |
| **File Share** | Supported during rolling upgrades. |
| **IP Address** | Supported during rolling upgrades. |
| **Network Name** | Supported during rolling upgrades. |
| **Physical Disk** | Supported during rolling upgrades. |

| | |
|---|---|
| **WINS** | WINS is not supported during rolling upgrades. You must follow the below instructions during rolling upgrade for WINS resource.<br>1. Confirm that your systems are running Windows 2000 and right Service packs if there is any.<br>2. Move WINS resource to a single node by double clicking the Groups folder in Cluster Administrator's console tree. Click the group that contains the resource and then in the **File** menu, click **Move Group**.<br>3. Upgrade the operating system on each of the nodes that is not hosting the WINS resource moved in the previous step.<br>4. Move the WINS resource that you moved in the earlier steps from the node that has not been upgraded to one of the newly upgraded nodes.<br>**Notes:**<br>● After the upgrade, the WINS resource can not failback to the Windows 2000 node.<br>● Following the procedure as described above will typically give you the highest availability of the resources that do not support rolling upgrades. |
| **DHCP** | DHCP is not supported during rolling upgrades. You must follow the below instructions during rolling upgrade for DHCP resource.<br>1. Confirm that your systems are running Windows 2000 and right Service packs if there is any.<br>2. Move DHCP resource to a single node by double clicking the Groups folder in Cluster Administrator's console tree. Click the group that contains the resource and then in the **File** menu, click **Move Group**.<br>3. Upgrade the operating system on each of the nodes that is not hosting the DHCP resource moved in the previous step.<br>4. Move the DHCP resource that you moved in the earlier steps from the node that has not been upgraded to one of the newly upgraded nodes.<br>**Notes:**<br>● After the upgrade, the DHCP resource can not failback to the Windows 2000 node.<br>● Following the procedure as described above will typically give you the highest availability of the resources that do not support rolling upgrades. |
| **Distributed Transaction Coordinator** | Supported during rolling upgrades. However, be aware that you should not create or delete the MSDTC resource while you are operating as a mixed mode (mixed version) cluster. You should instead create the MSDTC resource before starting the upgrade on any node, or wait until the upgrade is complete on all nodes. |
| **Internet Information Services** | See section Rolling upgrade Windows 2000 Internet Information Services to Windows Server 2003. |
| **NNTP Service** | NNTP Service Instance is not supported in Windows Server 2003 Enterprise Edition, therefore you must remove any NNTP Service Instance resources prior to upgrading. |

| | |
|---|---|
| **SMTP Service** | SMTP Service is not supported in Windows Server 2003 Enterprise Edition, therefore you must remove any SMTP Service resources prior to upgrading.<br>**Note**<br>This refers only to the built-in clustering resource for Windows SMTP. SMTP resources based on the Exchange 2003 resource DLL (exres.dll) are not affected by this upgrade limitation and should not be removed. |
| **Time Service** | Time Service is not supported in Windows Server 2003 Enterprise Edition, therefore you must remove any Time Service resources prior to upgrading. |
| **Local Quorum** | Supported during rolling upgrade. |
| **Message Queuing** | Primary Enterprise Services, Primary Site Services, and Backup Site Services are not supported during rolling upgrades. All other Message Queuing Services configurations are supported during rolling upgrades.<br>1. Take all of the MSMQ cluster resources offline.<br>2. For each MSMQ resource, take the Network Name resource that depends on it offline also.<br>3. Upgrade all of the nodes from Windows 2000 to Windows Server 2003 Enterprise Edition.<br>4. Bring the Network Name resources back online. Ensure that the Distributed Transaction Coordinator (DTC) resource is online, then bring the MSMQ resources online. |
| **Print Spooler** | The only Print Spooler resources supported during a rolling upgrade are those on line printer remote (LPR) ports. For more information, refer to Print documentation on Windows Server 2003 Enterprise Edition CD or Windows Server 2003 Enterprise Edition resource kit. |
| **Other resource types** | See the product documentation supplied with the application or resource. |

Before performing a rolling upgrade, identify the resources on your cluster that do not support rolling upgrades. This will help in determining which upgrade procedure you should use.

**Internet Information Services Rolling upgrade**

Windows Server 2003 provides a script that allows you to upgrade IIS. You can choose to either use the script or do the upgrade manually. In Windows Server 2003 operating system IIS is managed in a cluster using generic script resource.

Below are two methods by which you can upgrade a cluster that has IIS installed.

1. To upgrade from Windows 2000 on a cluster containing IIS using the provided script.

   a. If you do not already have a Distributed Transaction Coordinator (DTC) resource on the cluster that you are upgrading, create a DTC resource on this cluster.

   b. Upgrade or perform rolling upgrade (whichever is suitable for your needs) all nodes except the node containing the IIS instance resource from Windows 2000 to Windows Server 2003 Enterprise Edition.

   c. On an upgraded node, from command prompt navigate to the %systemroot%\system32\inetsrv directory and run

      **iis_switch** *cluster name [group name]*

| Value | Description |
|---|---|

| | |
|---|---|
| Cluster name | Cluster name or virtual server name that is hosted on an upgraded node. |
| Group name | By default the script will upgrade all IIS web and FTP resources in the cluster. A user has the option to specify a particular resource group also. |

**Note**

The script will bring the group with the newly created Generic Script resource online. However the process of bringing resource online after upgrade will fail if all resources in the group were offline or were only partially online at the time of script run.

The script will not upgrade IIS resources that are in the cluster group.

Once the script completes, using IIS, start the web site on the upgraded node. At this stage IIS can only be hosted on a Server 2003 nodes as it uses generic script that does not exist in Windows 2000 server cluster. Upgrade the final node now.

2. To upgrade from Windows 2000 on a cluster containing IIS manually.

   a. If you do not already have a Distributed Transaction Coordinator (DTC) resource on the cluster that you are upgrading, create a DTC resource on this cluster.

   b. Upgrade or perform rolling upgrade (whichever is suitable for your needs) all nodes except the node containing the IIS instance resource from Windows 2000 to Windows Server 2003 Enterprise Edition. Bring the IIS instance resource offline on Windows 2000 node after noting the dependencies of the IIS instance resource. You should note resources that depend on the IIS resource and also note what resources IIS itself is dependent on.

   c. Delete the IIS resource after removing any dependencies on the IIS instance resource.

   d. Move the group to a Windows Server 2003 Enterprise Edition node and create a Generic Script resource.

   **Note**

   From your notes of Step c, recreate the dependencies of the Generic Script resource identical to those of the now deleted IIS resource. You should make everything that was dependent on the IIS resource dependent instead on the Generic Script resource. You should also make the Generic Script resource dependent on everything that IIS was dependent on. Make sure the upgraded resource is the owner of the Generic Script Resource for IIS.

   e. Bring the group containing the new Generic Script resource online.

   f. Using IIS, start the web site on the upgraded node.

   g. Upgrade the final node.

**Application Considerations**

As stated earlier, Cluster service does not guarantee that applications can support rolling upgrades. Applications can, however, support rolling upgrades provided that they do *not*:

- Store program files on the clustered disk.
- Change the resource DLL name or location.
- Delete application registry keys either in the system registry or in the cluster configuration database.
- Change application on-disk data structures.

An application that does any of the above will not support rolling upgrades.

# How Rolling Upgrades Work (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phase 1: Preliminary Configuration**
In a common scenario, each node runs Windows 2000 Advance Server, with the following software installed:

- Microsoft Server Cluster.
- File Server -- group 3 and 4
- SQL Database -- group 1 and 2

At this point, the cluster is configured as in Figure 1 below so that each node handles client requests.




**Figure 1: Configuration where both nodes handle clients**

**Phase 2: Upgrade Node 1**
Node 1 is paused, as Figure 2 below shows. All resource groups on Node 1 are moved to Node 2. Because Node 1 is paused, no new groups can be created or moved to this node by users. At this point, Node 2 handles all cluster resource groups. Node 1 is idle and can be upgraded.




**Figure 2: Node 1 idle/Node 2 active**

At this point, you can start the installation of a service pack or an upgrade to Windows Server 2003 Enterprise Edition (see Figure 3 below). Once the upgrade of Node 1 is complete, you can perform a test to verify that the operating system is fully functional.




**Figure 3: Node 1 upgrade**

Cluster service maintains each nodes version of the operating system and the Cluster service itself, as well as the aggregate version of the cluster. It uses these version numbers to determine if a node that runs a different operating system version can join a cluster.

**Phase 3: Upgrade Node 2**

Up-level Node 1 rejoins the cluster. Cluster service guarantees that the up-level node understands the down-level nodes protocols and can join a down-level cluster.

Before upgrading Node 2, confirm that the cluster operates correctly. A simple test is to:

1. Select a resource group that is not critical, and move it to Node 1. If this step succeeds, it proves that Cluster service is working.

2. Move this resource group back to Node 2. This step proves that Node 1 did not tamper with the resources in this resource group, and that they will be able to fail back in case Node 1 fails.

Node 2 is paused, as we see in Figure 4 below. All resource groups on Node 2 are moved to Node 1. Because Node 2 is paused, no new groups can be created or moved to this node by other users. At this point Node 1 handles all cluster resource groups. Node 2 is idle and can be upgraded.




**Figure 4: Node 1 upgraded/Node 2 idle**

You can now start the installation of a service pack or an upgrade to Windows Server 2003 Enterprise Edition on Node 2 (Figure 5 below). Once the upgrade of Node 2 is complete, you can perform a test to verify that the operating system is fully functional. At this point, the up-level Node 2 should successfully join the up-level Node 1.




**Figure 5: Node 2 upgraded**

**Phase 4: Final**

Node 2 rejoins the cluster, and you redistribute the resource groups back to both nodes of the cluster configuration, as represented in Figure 6 below.




**Figure 6: Windows Server 2003 upgraded Configuration where both nodes handle clients**

# Rolling Upgrade Walkthrough

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This walkthrough covers the rolling upgrade of a Windows 2000 Advanced Server based cluster to Windows Server 2003 Enterprise Edition. For the latest information about rolling upgrades, please refer to the on-line help and the release notes on the Windows Server 2003 Enterprise Edition CD.

1. Prepare a list of all resources.

2. Identify resources that do not support rolling upgrades.

   If you have a resource that does not support rolling upgrades, you have three choices:

   a. Take the resource offline and continue with the rolling upgrade. The resource will be unavailable during the entire rolling upgrade process. You will have to bring this resource online once all nodes have been upgraded.

   b. Remove the resource. The resource will be unavailable during the entire rolling upgrade process. You will have to reinstall this resource once all nodes have been upgraded.

   c. If the majority of the resources do not support rolling upgrades, you may consider upgrading all nodes at once, or performing a clean install. In both cases, all resources will be unavailable during the upgrade/install process.

3. To start Cluster Administrator (see Figure 7 below), click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Cluster Administrator**.



**Figure 7: Cluster Administrator**

   All nodes and resource groups are up and online.

4. Click the first node (in our example, I4DELL1N1).

5. Click the **File** menu and click **Pause Node**. The status of I4DELL1N1 changes to **Paused**. (See Figure 8 below.)




**Figure 8: I4DELL1N1 is paused.**

6. In the left pane, click the + next to **I4DELL1N1** to expand it. Double-click **Active Groups**. All groups currently hosted on **I4DELL1N1** are displayed in the right pane (Figure 9 below).




**Figure 9: I4DELL1N1 Active groups.**

7. Click **Disk Group 1** in the right pane. Click the **File** menu and click **Move Group**. Repeat this step for each group listed.

   The services will be interrupted during the time that the services in each group are being moved to the other node and restarted. After all groups are moved, I4DELL1N2 hosts all groups and handles all client requests. I4DELL1N1 is idle (Figure 10 below).



**Figure 10: I4DELL1N2 has all groups.**

You will now need to review all resource types installed on your cluster.

8. Click **Resource Types** in the left pane (Figure 11 below).



**Figure 11: Resource types**

9. Determine which resource types do not support rolling upgrades and take offline all resources that are not supported. (Refer to Table 1 above.)

10. After ensuring that the latest released Service Pack has been applied, use Windows Server 2003 Enterprise Edition Setup to upgrade I4DELL1N1.

| **Note** |
| --- |
| To make sure that Windows Server 2003 Setup does not place temporary files on one of the clustered disks, use the installation option **/tempdrive:X**. (For more information, refer to the *Known Issues* section of this paper.) |

Setup detects Cluster Server on I4DELL1N1 and automatically installs Cluster service for Windows Server 2003

Enterprise Edition. I4DELL1N1 automatically rejoins the cluster at the end of the upgrade process, but is still paused and does not handle any cluster-related work.

You will now need to perform validation tests on I4DELL1N1 to confirm functionality.

11. Start Cluster Administrator on I4DELL1N1. You should see both nodes: I4DELL1N1 in **Paused** state, and I4DELL1N1 in **Up** state.

12. Click **I4DELL1N1** and then click **Resume Node**.

13. Repeat steps 4, 5, 6, 7, 8, 10, 11, and 12 on the second node (I4DELL1N2).

14. If necessary, upgrade any applications that do not support rolling upgrades.

15. Test the cluster by moving groups between nodes. All resources should come online on any node.

# Known Issues

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Location of Temporary Files**
The upgrade program, Winnt32.exe, selects a disk with the maximum free space for the storage of all temporary files. In a cluster, this will most likely be a clustered disk. Use the **/tempdrive:X** option to make sure that the upgrade program uses the appropriate drive (*not* a clustered disk) for the storage of temporary files.

**Print Spooler Resource**
The Print Spooler resource supports rolling upgrades but only for the LPR port and standard TCP/IP port types. Other port types are not supported on Windows 2000based server clusters.

You should not attempt to modify the printer spooler or printer configuration while running a mixed-version cluster.

**Changing Configuration in a Mixed-Version Cluster**
Changing the cluster configuration while running a mixed-version cluster is not recommended nor supported. In particular, you should not attempt to create new resources of the type that are only available on up-level Windows Server 2003based nodes.

# Troubleshooting (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource does not support rolling upgrades.**
Follow the instructions in step 2 above.

**Resource does not come online after an upgrade of one node.**
There are a few reasons for a resource to fail to come online or fail after it was brought online on the up-level node.

1. *The resource doesn't support rolling upgrades*.

   Move the resource back to the down-level node and take it offline. You will not be able to bring it online until all nodes are upgraded

2. *The resource is not compatible with Windows Server 2003*.

3. *Another resource that your application depends on failed to come online*.

   Make sure that the other resource does not fall into categories 1 or 2 above.

**Node upgrade fails.**
When a node upgrade fails, you have the following options:

1. Repair the node.

2. Restore the node from a backup.

3. Perform a clean install of the node. In this case, you will have to evict the node from the cluster, and rejoin the cluster once the installation is complete. This process may affect other applications running on the cluster. You may need to reinstall these applications. Please refer to Windows Server 2003 Enterprise Edition or Datacenter Edition online help section on installing server cluster.

# Conclusion (Server Clusters: Rolling Upgrades. Upgrading to Windows Server 2003)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This document has presented information that should help prepare you to perform a rolling upgrade of your cluster. A rolling upgrade will save you time and money, will provide you with the comfort of knowing that even if something goes wrong you will still have a working server, and will increase the satisfaction of clients using services offered by your cluster.

# Server Clusters: Majority Node Set Quorum

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP2

Published: January 1, 2004

The goal of this document is to provide a plain English explanation of the new quorum type available in Windows Server 2003 clusters-majority node set (MNS) clusters. Currently, the standard quorum that is used by the Microsoft Cluster service (MSCS) in Microsoft Windows NT Server 4.0 and the Windows 2000 Server family is sufficient for the majority of situations users will encounter. With the introduction of Windows Server 2003, Microsoft has introduced the MNS quorum, which is primarily for advanced users and other software vendors to create more sophisticated end-to-end solutions.

**In This White Paper**

Motivation (Server Clusters: Majority Node Set Quorum)

What Is a Majority Node Set? (Server Clusters: Majority Node Set Quorum)

How to Set Up a Majority Node Set (Server Clusters: Majority Node Set Quorum)

Using the Cluster Command to Setup a Majority Node Set (Server Clusters: Majority Node Set Quorum)

Sizing a Majority Node Set Cluster (Server Clusters: Majority Node Set Quorum)

What to Do If You Lose Quorum (Server Clusters: Majority Node Set Quorum)

Other Considerations (Server Clusters: Majority Node Set Quorum)

# Motivation (Server Clusters: Majority Node Set Quorum)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP2

The Microsoft Cluster Server (MSCS) architecture requires a single quorum resource in the cluster that is used as the tie-breaker to avoid split-brain scenarios. A split-brain scenario happens when all of the network communication links between two or more cluster nodes fail. In these cases, the cluster may be split into two or more partitions[1] that cannot communicate with each other. The cluster service guarantees that even in these cases, a resource is brought online only on one node. If the different partitions of the cluster each brought a given resource online, it would violate the cluster guarantees and potentially cause data corruption. When the cluster is partitioned, the quorum resource is used as an arbiter. The partition that owns[2] the quorum resource is allowed to continue. The other partitions of the cluster are said to have lost quorum, and the cluster service and any resources hosted on the nodes not part of the partition that has quorum, are terminated.

The quorum resource is a storage-class resource and, in addition to being the arbiter in a split-brain scenario, is used to store the definitive version of the cluster configuration. To ensure that the cluster always has an up-to-date copy of the latest configuration information, the quorum resource must itself be highly available. In Windows NT Server 4.0 and Windows 2000 Server, the quorum device is typically a shared disk or physical disk resource type.[3]

This architecture limits the types of cluster configurations that can be deployed.

In Windows NT 4.0 and Windows 2000, MSCS does not support non-shared disk cluster configurations; that is, there must always be at least one shared disk in the cluster to act as the quorum resource. (Note: In this document, a shared disk is defined as any disk that is physically attached to multiple computers; it does not imply anything about how data is accessed on those disks).[4]

Although the quorum resource may be a highly available resource (such as a mirrored disk or a RAID set), there is a perception that this architecture has a single point of failure. There is a single resource and if it becomes unavailable, the cluster cannot function.

In Windows Server 2003, the goal is to enhance the types of configuration that can be deployed using MSCS. In particular, the following scenarios are addressed:

- Geographically dispersed cluster-In a cluster that spans multiple sites, the notion of quorum as a single shared-disk resource means that the storage subsystem has to interact with the cluster infrastructure to provide the illusion of a single storage device with very strict semantics.

- Clusters with no shared disks-There are some specialized configurations that need tightly consistent cluster features without having shared disks.

- Clusters that host applications that can failover, but where there is an application-specific way to keep data consistent between nodes. Two possible examples are: database log shipping for keeping database status up to date and file replication for relatively static data.

- Clusters that host applications that have no persistent data, but need to cooperate in a tightly coupled way to provide consistent volatile state.

- Clusters that want the semantic of using a persistent store for applications but do not want to add to the complexity of doing the same for the quorum device. In such cases the application data will be stored on the shared disk but the quorum information could be located a local disk of every node.

# What Is a Majority Node Set? (Server Clusters: Majority Node Set Quorum)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP2

A majority node set is a single quorum resource, from a server cluster perspective; however, the data is actually stored on multiple disks across the cluster. Each cluster node stores the configuration on a local disk it can have access to when it starts up. By default, the location is pointed to %systemroot%\cluster\ResourceGUID. It is recommended that users user the default location unless there is an overriding reason not to.

The cluster service ensures that the cluster configuration data stored on the majority node set is kept consistent across each cluster node. This allows cluster topologies as follows:




If the configuration of the cluster changes, that change is replicated across the different disks. The change is only considered to have been committed, that is, made persistent, if that change is made to a majority of nodes that have formed the cluster. Mathematically this equates to:

(<Number of nodes configured in the cluster>/2) + 1

This ensures that a majority of the nodes have an up-to-date copy of the data. This also means that there is a change in how the cluster service starts up on each node, since the cluster service is set to automatic startup on boot. On the first node in the cluster, the cluster service will try to join a cluster by default. Since there is no cluster available at this stage it will not be able to and will stop. The default retry period is set to one minute so the cluster service will try to restart after every minute. This will successfully happen once a majority of the cluster nodes are up and running and can talk to each other through either the public or the private network. Network priority is honored, and each node will try to talk across that network that has been flagged to be used first for network communication. This behavior is the same as in vanilla clusters (clusters that use a device on the share storage interconnect for quorum).

Once the cluster service has started up cleanly on a majority of nodes, it will then bring all the resources online. If at any point in time the cluster loses a majority of nodes, the cluster services will cleanly shut down on all nodes leading to all the resources being terminated as well. This is a very important semantic that is different from the behavior customers are used to in vanilla clusters where the cluster service is running as long as there is at least one node up and running and it can access the quorum device.

In the case of a failure or a split-brain caused by a failure in communications between cluster nodes, all resources hosted on the partition of a cluster that has lost quorum are terminated to ensure that the resources are being run on only one node in a cluster . If a cluster becomes partitioned (that is, there is a communications failure between two sets of nodes in the cluster), any partitions that do not have a majority of the configured nodes (that is, have less than (n/2) + 1 nodes) are said to have lost quorum. This will lead to the cluster service and all resources hosted on the nodes of that partition to be terminated. This ensures that if there is a partition running that contains a majority of the nodes, it can safely start up any resources that are not running on that partition, safe in the knowledge that it can be the only partition in the cluster that is running resources (since all other partitions must have lost quorum). The diagram below depicts the behavior one will see. In the diagram, Nodes 4 and 5 have lost communication with Node 1, 2, and 3. This will cause the cluster services on Nodes 4 and 5 to terminate, causing the resources they hosted to move to Nodes 1 and 3.




In a vanilla server cluster, a cluster can continue as long as **at least one of the nodes** can access and own the quorum disk.

Any node that is a cluster member and can communicate (via heartbeats) with the quorum owner are part of the cluster and can host resources. Any other nodes that are configured to be in the cluster but cannot communicate with the quorum owner are said to have lost quorum, and thus any resources that they are hosting are terminated. When a node starts up, the cluster services can join an existing cluster (if there is one up and running), or create a cluster right away (if it can access the quorum device).

A cluster running with a majority node set quorum resource, on the other hand, will only start up or continue running **if a majority of the nodes configured for the cluster are up and running** and **can all communicate with each other**. The failure semantics of the cluster are different from a vanilla MSCS cluster. Because of this difference in how the cluster survives start up and on how the cluster behaves when nodes fail or get partitioned, care must be taken when deciding to host an application on a vanilla server cluster using a physical disk as quorum resource, or on a server cluster that uses a majority node set as a quorum resource.

The following sections cover how the quorum disk resource can be setup and what happens in the event that quorum is lost.

# How to Set Up a Majority Node Set (Server Clusters: Majority Node Set Quorum)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP2

The majority node set resource is an optional quorum resource that we ship with the MSCS. By default, the physical disk resource is setup as a quorum resource.

Setting up a MNS cluster is no different to setting up a standard cluster. The only difference will be the location of the quorum and the fact that the user can choose not to nominate a disk located on a shared storage interconnect as a quorum disk.

If there is a shared storage, set up can be done using standard mechanisms. The quorum disk can be located on the shared interconnect or as a local quorum disk. Once installation is complete, open Cluster Administrator and create a new MNS resource in the cluster group. Bring it online. Then go to the cluster properties and change the cluster quorum to the newly created MNS resource.

If there is no shared interconnect, the cluster configuration wizard will create a single-node local quorum disk cluster by default. The procedure for creating an MNS quorum resource and using it is the same as above.

Once a single node cluster is set up, the quorum resource can be switched to the majority node set resource using cluster administrator as follows:

- Open cluster administrator and connect to the appropriate cluster.
- Create a new resource of type **Majority node set**.
- Change the cluster quorum resource to the newly created majority node set resource. The Quorum tab on the cluster properties sheet shows the current quorum resource and it can be modified by selecting the appropriate resource.

Once a single node cluster is set up using the majority node set resource, other nodes can be added to the cluster. As each node is joined, the quorum calculation is updated. In other words, if two nodes are added to build a three-node cluster, quorum is adjusted so that two out of the three nodes must be running for the cluster to be considered as having quorum.

The majority node set quorum resource ensures that the quorum data is kept consistent across all nodes of the cluster and it is responsible for making changes to the quorum data stored on all nodes in the cluster.

Each node in the cluster stores its quorum data in the following directory on the system disk:

%SystemRoot%\Cluster\MNS.%ResourceGUID%$\%ResourceGUID%$\MSCS

The contents of this directory should not be modified; this directory contains cluster configuration information. Modifying the data in this directory may mean that the cluster service cannot start on the node.

The quorum resource accesses the data on the other nodes in the cluster using a network share. When a node is added to a cluster using the majority node set resource, a file share is created to expose the cluster quorum data to the majority node set resource. The share name is defined as follows:

\\%NodeName%\%ResourceGUID%$




The share is created with local system access. Since the cluster service is running under a domain account that is also a member of the local system group, the majority node set resource has access to the share.

To avoid potential security issues and issues where the cluster service cannot access the share, the permissions or security attributes associated with this file share should not be modified. It is recommended that default setting be used and no changes be made to ACLs governing the share created to be used by the Cluster service.

Since we use SMB shares as a MNS quorum device the cluster service now have to rely on an "external" component to provide a critical infrastructure for cluster service to be functional. SMB has two components, LanManServer and LanManWorkstation, which install as separate services in Server 2003. Both these services need to be running otherwise the cluster nodes will not be able to publish shares or communicate with the shares on other cluster nodes.

# Using the Cluster Command to Setup a Majority Node Set (Server Clusters: Majority Node Set Quorum)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP2

Rather than using cluster administrator as described above, the cluster command can be used to change the quorum resource for a single node cluster to a majority node set resource as follows:

```
cluster resource "MyQuorum" /create /group:"Cluster Group" /type:"Majority Node Set"
cluster resource "MyQuorum" /online /wait
cluster /quorum:"MyQuorum" /path:"MSCS\"
```

# Sizing a Majority Node Set Cluster (Server Clusters: Majority Node Set Quorum)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP2

As described above, if a partition of a cluster loses quorum, all of the resources on that partition of the cluster are terminated. This is done to guarantee the cluster semantics that only a single instance of the resource will be brought online (and thus serving or modifying data). Loss of quorum not only happens when a cluster is partitioned, but in the event of node failures, a cluster can lose quorum simply because there are insufficient nodes up and running. The number of nodes in a majority node set cluster should be carefully considered when setting up a cluster. For example, in a two-node cluster, if either node fails for any reason, the other node cannot continue because it does not have quorum. In an eight-node cluster, up to three nodes can fail before the cluster loses quorum. The following table shows the number of node failures that a given quorums disk set cluster can tolerate and continue:

| Number of nodes configured in the cluster | Number of node failures |
|---|---|
| 1 | 0 |
| 2 | 0 |
| 3 | 1 |
| 4 | 1 |
| 5 | 2 |
| 6 | 2 |
| 7 | 3 |
| 8 | 3 |

In a situation where a cluster spans multiple sites, it is recommended that at least one site has a majority over the other. That way in case there is a site link failure you can be assured that the cluster is up and running on at least one site. If you do lose the majority site, the nodes in the minority can be restarted manually to form a cluster.

# What to Do If You Lose Quorum (Server Clusters: Majority Node Set Quorum)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP2

There are cases where a cluster must be allowed to continue even if it does not have quorum. Consider the case of a geographically dispersed cluster with four nodes at the "primary" site and three nodes at the "secondary" site. While there are no failures, the cluster is a 7-node cluster where resources can be hosted on any node, on any site (depending on business needs). If there is a communications failure between the sites or if the secondary site is taken offline (or fails), the primary site can continue since it will still have quorum. All resources will be re-hosted and brought online at the primary site.

In the event of a catastrophic failure of the primary site, however, the secondary site will lose quorum and therefore all resources will be terminated at that site. One of the primary purposes for having a multi-site cluster is to survive a disaster at the primary site; however, the cluster software itself cannot make a determination about the state of the primary site. The cluster software cannot differentiate between a communications failure between the sites and a disaster at the primary site. That must be done by manual intervention. In other words, the secondary site can be forced to continue even though the cluster software believes it does not have quorum. This is known as forcing quorum.

Because this mechanism is effectively breaking the semantics associated with the quorum replica set, it must only be done under controlled conditions. In the example above, if the secondary site and primary site lose communication and an administrator forces quorum at the secondary site, resources will be brought online at BOTH sites, thus allowing the potential for inconsistent data and/or data corruption in the cluster.

Forcing quorum is a manual process that requires the following steps

Stop the cluster service **ON ALL** of the remaining nodes using cluster administrator.

The cluster service must be told which nodes should be considered as having quorum. This can be done in one of two ways:

Setup ForceQuorum registry key **ON ALL** remaining nodes in the cluster under

HKLM\SYSTEM\CurrentControlSet\Services\ClusSvc\Parameters\ForceQuorum

This is a REG_SZ key that should be setup to contain a **comma separated list of the names of the nodes that are to have quorum.** The key is case insensitive. So, in the above example, if the secondary site contains "Node5", "Node6" and "Node7", then the ForceQuorum registry key should be setup as

"Node5,Node6,Node7"

| **Note** |
| --- |
| There should be no spaces in the key (except where there are spaces in the node names themselves). |

Once the registry keys are set on all nodes, the cluster service can be started on those nodes.

Setup the cluster service startup parameters **ON ALL** remaining nodes in the cluster. This is done by starting up the services control panel, selecting the cluster service and entering the following into the "start parameters" option:

/forcequorum <node list>

In the above example, if the secondary site contains "Node5", "Node6" and "Node7", then the cluster service start parameter should be set to:

/forcequorum Node5,Node6,Node7

The cluster service MUST be started by clicking the START button on the service control panel, you must not hit OK or Apply first as this does not preserve the parameters.

| **Note** |
| --- |
| Any command line parameters over-ride the registry setting, however, the command line parameters do NOT persist a reboot, and therefore, setting the registry key is the preferred mechanism for forcing quorum. |

The cluster service will now start up on those nodes that are considered part of the quorum set and resources will be brought online.

Special care must be taken if and when the primary site comes back since the nodes are configured as part of the cluster.

Do **NOT** reboot the cluster nodes at the primary site

Stop the cluster service **ON ALL** of the cluster nodes

Remove the registry key setting or the cluster service startup parameters set to force quorum

Startup the cluster service on all of the nodes at the secondary site

Boot the nodes at the primary site

| **Note** |
| --- |
| The cluster service on all nodes NOT in the force quorum node list must remain stopped until the force quorum information is removed. Failure to do so can lead to data inconsistencies OR data corruption. |

While a cluster is running in the force quorum state, it is fully functional. For example, nodes can be added or removed from the cluster; new resources, groups etc. can be defined.

# Other Considerations (Server Clusters: Majority Node Set Quorum)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP2

There is a single quorum of node resources in the cluster and it is brought online on one node at any one time, just like any other cluster resource. The majority node set resource is responsible for ensuring that the quorum data is kept consistent on all disks around the cluster. When a cluster is setup to have a majority node set quorum resource or when a node is added to a cluster with a majority node set, a file share is created on that node. As described above, each node in a majority node set cluster has a file share that exports the quorum directory, so that, regardless of where the majority node set resource is hosted, it can write to all of the members of the majority node set.

The share is NOT under cluster server control (since this would lead to a chicken and egg situation: the cluster needs the share, the share needs the cluster). In other words, the share is available and visible when the node is booted, NOT when the cluster service starts. The share will be used in quorum calculations when it is visible, NOT when the node joins the cluster, so a node that is configured to be in the cluster is contributing to the quorum calculation for the cluster even if the cluster service on that node is stopped. This is VERY important to consider and it can have an impact on operational procedures, particularly when changing the number of nodes in a cluster.

Consider the following case of a 4-node cluster:

Node 1 is booted but the cluster service is not running

Nodes 2, 3, and 4 are fully operational and running in a cluster.

The administrator then evicts nodes 2 and 3 from the cluster, thus the cluster now contains two nodes (1 and 4). Node 4 still has quorum since it has its own share and also has access to the share on Node 1 which is configured as part of the cluster but not running the cluster. If the cluster service on node 1 starts at this point, it will join the existing cluster (containing Node 4) just fine, Node 1's local quorum information stored in the registry will be updated to reflect the new membership. If on the other hand, the cluster service on Node 4 is stopped, and then the cluster service on Node 1 is started, Node 1 will NOT be able to startup and form a cluster. This is because Node 1 tries to form a cluster based on the quorum information it has in its registry (NOT the quorum database on the majority node set). It still believes that Nodes 1, 2, 3, and 4 are members and therefore it will not start until it has quorum.

This situation can be corrected by using the force quorum mechanism above. The force quorum mechanism allows the cluster service on Node 1 to startup at which point it will refresh its local registry state with that from the majority node set.

**Using MNS in single node clusters**-Using MNS as a quorum device on single node clusters is fully supported. There are some startup implications a user needs to consider when choosing this model. As explained before, MNS depends on LanManServer and LanManWorkstation to initialize. These dependencies have not been explicitly created in Service Control Manager. As such when Cluster service initializes if either of the two services are not ready Cluster services will fail and an entry denoting the same will be logged in Event Log. However the default recovery behavior for Cluster services has been set to 1 minute. This means at it will try and restart every minute and will eventually succeed when the dependant services come online. This might mean that there could be a lag time in Cluster services and clustered application to come online even after the operating system has booted successfully. This behavior is slightly different as compared to a single node cluster that has a local quorum device or a quorum device located on a share disk. This will be particularly noticed if the cluster is hosting file share resources that contain a large number of shares (2000 or more).

**Using MNS as a quorum device in File and Print clusters**-The conditions mentioned above are also applicable to larger clusters (greater than two) that choose to use MNS as a quorum device. In a cluster that hosts a large number of shares we will continue to have similar startup and failover issues. Although LanManServer has been optimized to a large extent in Windows Server 2003, it will still take some time (anywhere from two to 16 minutes) if a large number of file shares (2000 to 20,000) are hosted in the cluster. The delay will be seen on the node that will be hosting the group containing the file share resource and it the node also acts as cannon for a majority then cluster functioning could be temporarily affected.

[1] A partition is defined as a set of cluster nodes that can all communicate with each other.

[2] How a partition becomes the owner of the quorum resource is beyond the scope of this discussion.

[3] This is the only resource type supported as a quorum resource in the product shipped as part of Windows 2000. Other vendors have supplied alternative quorum resource types, but they are still typically associated with disks on a shared storage bus.

[4] MSCS does have the notion of local quorum, which can be used as a debugging tool; however, it is not recommended as a

production quorum resource (see KB article 245626).

# Using Windows Server 2003 in a Managed Environment: Controlling Communication with the Internet

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This white paper provides information about the communication that flows between components in Windows Server 2003 and sites on the Internet, and it describes steps to take to limit, control, or prevent that communication in an organization with many users.

**In This White Paper**

Introduction to Controlling Communication with the Internet (Windows Server 2003)

Activation, Registration, and Internet Communication (Windows Server 2003)

Application Help and Internet Communication (Windows Server 2003)

Certificate Support and Internet Communication (Windows Server 2003)

Device Manager and Internet Communication (Windows Server 2003)

Driver Protection and Internet Communication (Windows Server 2003)

Dynamic Update and Internet Communication (Windows Server 2003)

Event Viewer and Internet Communication (Windows Server 2003)

File Association Web Service and Internet Communication (Windows Server 2003)

Help and Support Center: Internet Communication (Windows Server 2003)

HyperTerminal and Internet Communication (Windows Server 2003)

Internet Explorer 6.0 and Internet Communication (Windows Server 2003)

Internet Information Services and Internet Communication (Windows Server 2003)

Internet Protocol Version 6 (IPv6) and Internet Communication (Windows Server 2003)

NetMeeting and Internet Communication (Windows Server 2003)

Online Device Help and Internet Communication (Windows Server 2003)

Outlook Express 6.0 and Internet Communication (Windows Server 2003)

Plug and Play and Internet Communication (Windows Server 2003)

Program Compatibility Wizard and Internet Communication (Windows Server 2003)

Remote Assistance and Internet Communication (Windows Server 2003)

Search Companion and Internet Communication (Windows Server 2003)

Terminal Server Licensing and Internet Communication (Windows Server 2003)

Windows Error Reporting and Internet Communication (Windows Server 2003)

Windows Media Player and Internet Communication (Windows Server 2003)

Windows Media Services and Internet Communication (Windows Server 2003)

Windows Time Service and Internet Communication (Windows Server 2003)

Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003)

Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003)

Appendix B: Resources for Learning About Group Policy (Windows Server 2003)

Appendix C: Message Queuing and Internet Communication (Windows Server 2003)

Appendix D: Connection Manager and Internet Communication (Windows Server 2003)

Appendix E: Passport Manager Administration (Windows Server 2003)

Appendix F: Internet Connection Sharing and Related Networking Features (Windows Server 2003)

Appendix G: Add Network Place Wizard (Windows Server 2003)

Appendix H: New Connection Wizard (Windows Server 2003)

Related Links for Controlling Communication with the Internet (Windows Server 2003)

# Introduction to Controlling Communication with the Internet (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Products in the Microsoft® Windows Server™ 2003 family include a variety of technologies that communicate with the Internet to provide increased ease of use. Browser and e-mail technologies are obvious examples, but there are also technologies such as Automatic Updates that help you obtain the latest software and product information, including bug fixes and security patches. These technologies provide many benefits, but they also involve communication with Internet sites, which administrators might want to control.

Control of this communication can be achieved through a variety of options built into individual components, into the operating system as a whole, and into server components designed for managing configurations across your organization. For example, as an administrator you can use Group Policy to control the way some components communicate, or for some components, you can direct all communication to the organization's own internal Web site instead of to an external site on the Internet.

This white paper provides information about the communication that flows between components in Windows Server 2003 and sites on the Internet, and it describes steps to take to limit, control, or prevent that communication in an organization with many users. The white paper is designed to assist you, the administrator, in planning strategies for deploying and maintaining products in the Windows Server 2003 family in a way that helps provide an appropriate level of security and privacy for your organization's networked assets.

This white paper provides guidelines for controlling components in the following operating systems:

- Windows Server 2003, Web Edition
- Windows Server 2003, Standard Edition
- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition

The white paper is organized around individual components found in the Windows Server 2003 family, so that you can easily find detailed information for any component you are interested in.

**What This White Paper Covers and What It Does Not Cover**
The subsections that follow describe:

- Types of components covered in this white paper
- Types of components not covered in this white paper
- Security basics that are beyond the scope of this white paper, with listings of some other sources of information about these security basics

**Types of components covered in this white paper**
This white paper provides:

- Information about components that in the normal course of operation send information to or receive information from one or more sites on the Internet. An example of this type of component is Windows® Error Reporting; if you choose to use this component, it sends information to a site on the Internet.
- Information about components that routinely display buttons or links that make it easy for you to initiate communication with one or more sites on the Internet. An example of this type of component is Event Viewer; if you open an event in Event Viewer and click a link, you are prompted with a message box that says, "Event Viewer will send the following information across the Internet. Is this OK?" If you click OK, information about the event is sent to a Web site, which

replies with any additional information that might be available about that event.

- Brief descriptions of components like Microsoft Internet Explorer and Internet Information Services (IIS) that are designed to communicate with the Internet. It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running servers that communicate across the Internet. This white paper does, however, provide basic information about how components such as Internet Information Services work, and it provides suggestions for other sources of information about balancing your organization's requirements for communication across the Internet with your organization's requirements for protection of networked assets.

**Types of components not covered in this white paper**
This white paper does not provide:

- Information about managing or working with applications, scripts, utilities, Web interfaces, Microsoft ActiveX® controls, extensible user interfaces, the .NET Framework, and application programming interfaces (APIs). These are either applications, or are layers that support applications, and as such provide extensions that go beyond the operating system itself.

  Windows Installer is not covered in this white paper, although Windows Installer includes some technology that (if you choose) you can use for installing drivers or other software from the Internet. Such Windows Installer packages are not described here because they are like a script or utility that is created specifically for communication across the Internet.

  Note that among the applications not covered in this white paper are Web-based and server-based applications, for example, server-based applications for databases, e-mail, or instant messaging. You must work with your application software provider to learn what you can do to mitigate any risks that are part of using particular applications (including Web-based or server-based applications), scripts, utilities, and other types of software that run on products in the Windows Server 2003 family.

- Information about components that store local logs that could potentially be sent to someone or could potentially be made available to support personnel. This information is similar to any other type of information that can be sent through e-mail or across the Internet in other ways. You must work with your support staff to provide guidelines about the handling of logs and any other similar information you might want to protect.

**Security basics that are beyond the scope of this white paper**
This white paper is designed to assist you, the administrator, in planning strategies for deploying and maintaining products in the Windows Server 2003 family in a way that helps provide an appropriate level of security and privacy for your organization's networked assets. The white paper does not describe security basics, that is, strategies and risk-management methods that provide a foundation for security across your organization. It is assumed you are actively evaluating and studying these security basics as a standard part of network administration.

Some of the security basics that are a standard part of network administration include:

- Monitoring. This includes using a variety of software tools, including tools to assess which ports are open on servers and clients.

- Virus-protection software.

- The principle of least privilege (for example, not logging on as an administrator if logging on as a user is just as effective).

- The principle of running only the services and software that are necessary, that is, stopping unnecessary services and keeping computers (especially servers) free of unnecessary software.

- Strong passwords, that is, requiring all users and administrators to choose passwords that are not easily deciphered.

- Risk assessment as a basic element in creating and implementing security plans.

- Software deployment and maintenance routines to help ensure that your organization's software is running with the latest security updates and patches.

- Defense-in-depth. In this context, defense-in-depth (also referred to as in-depth defense) means redundancy in security systems, for example, using firewall settings together with Group Policy to control a particular type of communication with the Internet.

**Other sources of information about security basics**

The following books and Web sites are a few of the many sources of information about the security basics described previously:

- Howard, Michael, et al. *Designing Secure Web-Based Applications for Microsoft Windows 2000*. Redmond, WA: Microsoft Press, 2000.

- Howard, Michael, and David LeBlanc. *Writing Secure Code*. Redmond, WA: Microsoft Press, 2002.

- Kaufman, C., R. Perlman, and M. Speciner. *Network Security: Private Communication in a Public World*. Upper Saddle River, New Jersey: Prentice-Hall Inc., 2002.

- Smith, B., B. Komar, and the Microsoft Security Team. *Microsoft® Windows® Security Resource Kit*. Redmond, WA: Microsoft Press, 2003.

  For more information, see the Microsoft Press Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29168

- The Security Guidance Center on the Microsoft Web site at:

  http://www.microsoft.com/security/guidance/

- The Prescriptive Architecture Guides on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=29413

- The Web page focused on security for Windows Server 2003 on the Microsoft Windows Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46003

- The Web page focused on security on the Microsoft Developer Network (MSDN®) Web site at:

  http://msdn.microsoft.com/security/

- The Web pages focused on security on the Microsoft TechNet Web site at:

  http://www.microsoft.com/technet/security/

  http://go.microsoft.com/fwlink/?LinkId=46004

# Activation, Registration, and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The purposes of activation and registration connected with a new installation or an upgrade
- How a computer running a product in the Microsoft Windows Server 2003 family communicates with sites on the Internet during activation and registration
- Choosing volume licensing so that product activation need not take place (to limit the flow of information to and from Internet sites)

**Purposes of Activation and Registration for a New Installation or an Upgrade**
This subsection briefly describes the differences between product activation and registration, and then describes the purpose of each.

Product registration involves the provision of personally identifiable information, such as an e-mail address, to Microsoft for the purpose of receiving information about product updates and special offers. Registration is usually done on a per-product basis and is not required. If registration is completed, all registration information is stored using a variety of security technologies and is never loaned or sold outside Microsoft.

Product activation involves the authentication with Microsoft of non-personally identifiable information, including the product identifier and a hardware hash representing the computer, for the purpose of reducing software piracy. (A hardware hash is a non-unique number generated from the computer's hardware configuration.) Activation of products in the Windows Server 2003 family is required in situations where the product is not purchased through a volume licensing program such as Microsoft Select License, Microsoft Enterprise Agreement, or Microsoft Open License. Many computer manufacturers can bypass activation on software preinstalled on a new computer by binding the software to the computer's basic input/output system (BIOS). In this situation, no activation of that software is required. Detailed information about product activation can be found on the following Web site:

http://go.microsoft.com/fwlink/?LinkId=29484

For more information about volume licensing, see "Choosing Volume Licensing So That Individual Product Activation Need Not Take Place," later in this section.

Activation is aimed at reducing software piracy as well as ensuring that Microsoft customers are receiving the product quality that they expect. Activation means that a specific product key becomes associated with the computer (the hardware) it is installed on. After that happens, that product key cannot be used for activation on other computers (unless you are enrolled in a special program that permits additional activations, for example, a program through the Microsoft Developer Network [MSDN]).

**Overview: Activation and Registration in a Managed Environment**
Product activation is an anti-piracy technology designed to verify that software products have been legitimately licensed. If you have software re-imaging rights granted under a Microsoft volume license agreement, and if you obtained a product in the Windows Server 2003 family through a retail channel or preinstalled by the computer manufacturer, you can re-image it with the product that you licensed through one of the Microsoft volume licensing programs. With volume licensing, there is no need to perform product activation.

**How a Computer Communicates with Sites on the Internet During Activation and Registration**
A product in the Windows Server 2003 family can be activated through the Internet or by phone. When it is activated through the Internet, the operating system communicates with Web sites as follows:

- **Specific information sent or received**: During activation, the following information is sent to the activation server at Microsoft:
    - Request information, that is, information necessary for successfully establishing communication with the

activation server.

- Product key information in the form of the product ID, plus the product key itself.

- A hardware hash (a non-unique number generated from the computer's hardware configuration). The hardware hash does not represent any personal information or anything about the software. It is based on the MD5 message-digest hash algorithm, and consists of a combination of partial MD5 hash values of various computer components. The hardware hash cannot be used to determine the make or model of the computer, nor can it be backward-calculated to determine the raw computer information.

- Date and time.

- The language being used on the system (so that any error message that is sent back can be in the correct language).

- The operating system being activated (and the version number of the activation software).

Depending on your preference, the preceding information is either sent over the Internet to the activation system at Microsoft, or the product key information and hardware hash (combined into one number) are called in by phone.

- **Default setting and ability to disable**: Product activation can only be disabled by installing the operating system with software acquired through one of the Microsoft volume licensing programs. Product activation can be bypassed by many computer manufacturers if they bind the product to the computer's BIOS instead. In all other cases, product activation cannot be disabled.

- **Trigger and notification for activation**: When activation is required, the operating system provides a reminder each time you log on and at common intervals until the end of the activation grace period stated in the End-User License Agreement (thirty days is the typical grace period). With software acquired through one of the Microsoft volume licensing programs, there is no need for activation, and therefore there are no reminders that appear about activation.

- **Trigger and notification for registration**: Registration is optional. You can register at the same time you activate by choosing appropriate options on the Windows Product Activation interface. As an alternative, you can type regwiz /r to start the Registration Wizard for Windows Server 2003. Before the wizard starts and in the first page of the wizard, brief explanations notify you that completing the wizard will cause the product to be registered.

- **Logging**: Entries that track the progress of activation and registration (for example, return codes and error codes) are logged into a text file, *systemroot*\setuplog.txt. This file can be used for troubleshooting if activation (or any part of Setup) fails. If you choose to register the product, two entries are made in the text file. One entry records the country or region that was chosen for the operating system. A second entry records whether you chose to have Microsoft (or the computer manufacturer) send information about product updates and special offers. No other registration data is logged.

- **Privacy, encryption, and storage for activation data**: Customer privacy was a paramount design goal in building the product activation technology. No personally identifiable information is collected as part of activation. The data is encrypted (HTTPS) during transmission, and it is stored on servers located in controlled facilities at Microsoft. The data is accessible to a small number of server and program support personnel who oversee and maintain the activation servers and the product activation program.

  To review the Microsoft online privacy statement on activation, see the following Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29923

- **Privacy, encryption, and storage for registration data**: When you register at the same time as you activate (through the Windows Product Activation interface), registration data is encrypted (HTTPS) during transmission. When you register by using the Registration Wizard (which you start by typing regwiz /r), registration data is encrypted (HTTPS) during transmission unless the wizard is unable to establish an HTTPS connection through port 443 with the Microsoft registration server. In this situation, registration data will be sent unencrypted, using HTTP through port 80.

Registration data, which contains information that you choose to send to Microsoft, is stored on servers with restricted access that are located in controlled facilities. The data can be seen by customer service representatives and marketing personnel. To review the Microsoft online privacy statement on registration, see the following Web site at:

http://go.microsoft.com/fwlink/?LinkId=29508

- **Transmission protocol and port**:

  - **For Windows Product Activation**: When the operating system is activated through the Internet and a modem is not used, the first transmission uses HTTP through port 80 and goes to wpa.one.microsoft.com/ to check the HTTP response code. A response code of less than 500 indicates that a product activation server is available. (With a modem, there is only a check to see whether the modem can currently be used to make a connection to the Internet.) If the product activation server can be reached (or for a modem, if a connection to the Internet can be made), any activation or registration data that is sent by Windows Product Activation uses HTTPS through port 443.

  - **For the Registration Wizard**: When you register by using the Registration Wizard (which you start by typing **regwiz /r**), HTTPS is used through port 443 unless the wizard is unable to establish an HTTPS connection through port 443 with the Microsoft registration server. In this situation, HTTP is used through port 80.

**Choosing Volume Licensing So That Individual Product Activation Need Not Take Place**

If you use the rights granted under a volume licensing agreement to purchase or re-image software, you cannot and need not perform activation on the individual computers that are installed under the volume license. Qualifying as a volume licensing customer is not difficult. Customers can qualify for the Microsoft Open Licensing program by purchasing as few as five licenses. For more information, see the Microsoft licensing Web site at:

http://go.microsoft.com/fwlink/?LinkId=29878

# Application Help and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Application Help
- How Application Help communicates with sites on the Internet
- How to control Application Help to prevent the flow of information to and from the Internet

**Benefits and Purposes of Application Help**
Application Help is one of the application compatibility technologies that support the installation and operation of applications on Microsoft Windows Server 2003 family operating systems. Because some applications that work on earlier versions of Windows might not function properly on Windows Server 2003 family operating systems, the application compatibility technologies were developed to solve these potential problems and enable a better user experience.

Application Help is most commonly used to block low-level applications—such as antivirus and disk-access utilities—that were not written for or intended for use on Windows Server 2003 family operating systems. By blocking the installation of these applications, this feature serves to avert serious problems that could compromise system integrity.

**Overview: Using Application Help in a Managed Environment**
Despite testing applications before you deploy Windows Server 2003 family operating systems, your organization may still use some applications that can cause system instability.

Application Help is the last line of defense against users attempting to install incompatible applications, and it is invoked only in rare instances. When a user tries to run an application for which there is no compatibility fix, Application Help is invoked by default. The operating system uses information in a local database to determine if a user is about to run an incompatible application. Compatibility fixes are contained in a database file named SYSMAIN.SDB. The warning information used when an application cannot be run successfully is contained in another database file, APPHELP.SDB. The operating system uses matching information in SYSMAIN.SDB, which in turn determines what messages to draw from APPHELP.SDB to block the operation of applications with known compatibility problems and to inform users about them. The list of incompatible applications is updated through Windows Update.

Application Help generates a message that is presented to the user when a problematic process is about to initiate. A dialog box appears that contains a brief message about the problem, with the severity indicated by an icon:

- If the icon is a yellow triangle with an exclamation mark, then the application is *not blocked*, which means that the user is still able to run the application.
- If the icon is a red stop sign, then the application is *blocked*, which means that the user cannot run the application.

The way these Application Help messages lead the user to interaction with the Internet is described in the following subsection.

While Application Help provides a valuable function, administrators in a highly managed environment might want to block the installation of applications that would automatically invoke Application Help and thereby allow a user to access the Internet. You can create custom Application Help messages that redirect the user to an internal site for more information. This is described in greater detail in the subsection, "Controlling Application Help to Prevent the Flow of Information to and from the Internet."

**How Application Help Communicates with Sites on the Internet**
In the Application Help dialog box, the user can click the Details button, in which case additional information is displayed in Help and Support Center. The Help content comes from either Microsoft.com if the computer is online, or from a local HTML Help file.

The following list describes how interaction with the Internet takes place when Application Help is invoked:

- **Specific information sent or received**: When the Details button is selected, a specific Web page from Microsoft.com is displayed; the Web page provides information about the problem application in the language of the operating system, for example, English (United States). The page that is displayed may provide a link to a non-Microsoft Web site, depending on the application. The URL provided for non-Microsoft Web sites is unique to each application. The user is not uniquely identified.

- **Default and recommended settings**: Application Help is enabled by default. Recommended settings are presented in the following topic, "Controlling Application Help to Prevent the Flow of Information to and from the Internet."

- **Triggers**: A user tries to run an application that is not compatible with Windows Server 2003 family operating systems.

- **User notification**: When the user selects the Details button there is no indication of whether the information is coming from an internal or external site.

- **Logging**: By default events related to Application Help are not logged; however, you can enable Application Help event logging. For this procedure see "To enable event logging for Application Help," later in this section.

- **Encryption**: The query that causes the display of an appropriate Web page (described in "Specific information sent or received," earlier in this list) is not encrypted.

- **Access**: No information from the use of Application Help is retained at Microsoft.

- **Privacy statement**: Application Help is covered by the same privacy statement that covers Windows Update.

- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is HTTP 80.

- **Ability to disable**: You can prevent Application Help from sending the user to the Internet by creating custom Application Help messages.

**Controlling Application Help to Prevent the Flow of Information to and from the Internet**

You can block an application with known compatibility problems, such as antivirus programs. You can also create custom Application Help messages that describe the problem and redirect users to an intranet site rather than sending them to the Internet for more information. To do this you use the Compatibility Administrator tool which is part of the Application Compatibility Toolkit.

The Application Compatibility Toolkit is a collection of tools and documents that can help you resolve application compatibility problems. For information about downloading the toolkit, see "Using the Application Compatibility Toolkit," later in this section.

**How creating custom Application Help messages can affect users and applications**

The user experience with Application Help does not change if you block applications with known compatibility problems and create custom Application Help messages. The only difference is that when users click the Details button they are sent to an internal site for more information instead of to the Internet. Not only can you prevent users from accessing the Internet in this way, but you can also avoid having users try to install incompatible applications.

**Procedures for Configuring Application Help**

This section presents three procedures. The first two procedures help you get started with the Application Compatibility Toolkit. The last procedure explains how to enable event logging.

**Using the Application Compatibility Toolkit**

You need to first download the Application Compatibility Toolkit, and then you can use the Compatibility Administrator tool to create custom Application Help messages and to block specific applications from running.

**To install the Application Compatibility Toolkit**

1. Read about the toolkit and find links for downloading it from the Windows Web site at:

   http://go.microsoft.com/fwlink/?LinkId=29880

2. Follow the installation instructions. Once you have installed the toolkit you can view the Windows Application

Compatibility 3.0 Reference and you can run the Compatibility Administrator tool to make the changes you need.

**To create custom Application Help messages**

1. Make sure that the Application Compatibility Toolkit is installed by using the previous procedure.

2. Click **Start**, point to **Programs** or **All Programs**, point to **Microsoft Windows Application Compatibility Toolkit**, and then click **Compatibility Administrator Tool 3.0**.

3. In the console tree, click **Custom Databases**, and then click **New Database**.

4. On the toolbar, click **AppHelp**. The **Create a custom AppHelp message** dialog box appears.

5. Enter information as prompted in the dialog box.

6. Save the new database file.

| **Note** |
| --- |
| When you have completed your entries and saved the file, you can deploy your changes to multiple computers running Windows Server 2003 and Windows XP. See "Deploying Compatibility Fixes," in Compatibility Administrator Help. |

**Application Help event logging**

Events related to Application Help are not logged by default. You can enable event logging in the Application log (Control Panel\Administrative Tools\Event Viewer) by using Group Policy.

**To enable event logging for Application Help**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Application Compatibility**.

3. In the details pane, double-click **Turn On Application Help Log Events**, and then select **Enabled**.

**Related Links**

For more information about application compatibility resources, see “Windows Application Compatibility” at:

http://go.microsoft.com/fwlink/?LinkId=29880

# Certificate Support and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The following subsections provide information about:

- The benefits of the certificate functionality built into operating systems in the Microsoft Windows Server 2003 family, including the benefits of Update Root Certificates
- How Update Root Certificates communicates with sites on the Internet
- How to control Update Root Certificates to limit the flow of information to and from the Internet

**Benefits and Purposes of Certificate Functionality**

Certificates, and the public key infrastructure of which they are a part, support authentication and encrypted exchange of information on open networks, such as the Internet, extranets, and intranets. A certificate securely binds a public key to the entity that holds the corresponding private key. With certificates, host computers on the Internet no longer have to maintain a set of passwords for individual subjects who need to be authenticated as a prerequisite to access. Instead, the host merely establishes trust in a certification authority that certifies individuals and resources that hold private keys. The host can establish this trust through a certificate hierarchy that is ultimately based on a root certificate, that is, a certificate from an authority that is trusted without assurances from any other certification authority.

Examples of times that a certificate is used are when you:

- Use a browser to engage in a Secure Sockets Layer (SSL) session
- Accept a certificate as part of installing software
- Accept a certificate when receiving an encrypted or digitally signed e-mail message

When learning about public key infrastructure, it is important to learn not only about how certificates are issued, but how certificates are revoked, and how information about those revocations is made available to clients. This is because certificate revocation information is crucial for an application that is seeking to verify that a particular certificate is currently (not just formerly) considered trustworthy. Certificate revocation information is often stored in the form of a certificate revocation list, although this is not the only form it can take. Applications that have been presented with a certificate might contact a site on an intranet or the Internet for information not only about certification authorities, but also for certificate revocation information.

In an organization where servers run products in the Microsoft Windows Server 2003 family, you have a variety of options in the way certificates and certification revocation lists (or other forms of certificate revocation information) are handled. For more information about these options, see the references listed in the next subsection, "Overview: Using Certificate Components in a Managed Environment."

The Update Root Certificates component in the Windows Server 2003 family is designed to automatically check the list of trusted authorities on the Microsoft Windows Update Web site when this check is needed by an application. Specifically, if the application is presented with a certificate issued by a certification authority that is not directly trusted, the Update Root Certificates component (if present) will contact the Microsoft Windows Update Web site to see if Microsoft has added the certification authority to its list of trusted authorities. If the certification authority has been added to the Microsoft list of trusted authorities, its certificate will automatically be added to the trusted certificate store on the computer. Note that the Update Root Certificates component is optional, that is, it can be removed or excluded from installation on a computer running a product in the Windows Server 2003 family.

**Overview: Using Certificate Components in a Managed Environment**

In an organization where servers run products in the Windows Server 2003 family, you have a variety of options in the way certificates are handled. For example, you can establish a trusted root authority, also known as a root certification authority, inside your organization. The first step in establishing a trusted root authority is to install the Certificate Services component. Another step that might be appropriate is to configure the publication of certificate revocation information to the Active

Directory® directory service. When implementing public key infrastructure, we recommend that you also learn about Group Policy as it applies to certificates. Procedures for these steps are provided in the resources listed at the end of this subsection.

When you configure a certification authority inside your organization, the certificates it issues can specify a location of your choosing for retrieval of additional evidence for validation. That location can be a Web server or a directory within your organization. Because it is beyond the scope of this white paper to provide full details about working with certification authorities, root certificates, certificate revocation, and other aspects of public key infrastructure, this section provides a list of conceptual information and a list of resources to help you learn about certificates.

Some of the concepts to study when learning about certificates include:

- Certificates and the X.509 V3 standard (the most widely used standard for defining digital certificates)
- Standard protocols that relate to certificates, for example, Transport Layer Security (TLS), Secure Sockets Layer (SSL), and Secure Multipurpose Internet Mail Extensions (S/MIME)
- Encryption keys and how they are generated
- Certification authorities, including the concept of a certification authority hierarchy and the concept of an offline root certification authority
- Certificate revocation
- Ways that Active Directory and Group Policy can work with certificates

The following list of resources can help you as you plan or modify your implementation of certificates and public key infrastructure:

- Help for products in the Windows Server 2003 family.

  You can view Help for products in the Windows Server 2003 family on the Web at:

  http://go.microsoft.com/fwlink/?linkid=29881

- The *Microsoft Windows Server 2003 Deployment Kit* and the *Microsoft Windows Server 2003 Technical Reference*.

  You can view links on the Windows Deployment and Resource Kits Web site at:

  http://www.microsoft.com/windows/reskits/

- "Troubleshooting Certificate Status and Revocation," a white paper on the Microsoft TechNet Web site at:

  http://go.microsoft.com/fwlink/?LinkId=27081

- Links to information about public key infrastructure for Windows Server 2003 on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=29886

**How Update Root Certificates Communicates with Sites on the Internet**

This subsection focuses on how the Update Root Certificates component communicates with sites on the Internet. The previous subsection, "Overview: Using Certificate Components in a Managed Environment," provides references for the configuration choices that control the way other certificate components communicate with sites on the Internet.

If the Update Root Certificates component is installed on a server, and an application is presented with a certificate issued by a root authority that is not directly trusted, the Update Root Certificates component communicates across the Internet as follows:

- **Specific information sent or received**: Update Root Certificates sends a request to the Windows Update Web site, asking for the current list of root certification authorities in the Microsoft Root Certificate Program. If the untrusted certificate is named in the list, Update Root Certificates obtains that certificate from Windows Update and places it in the

trusted certificate store on the server. No user authentication or unique user identification is used in this exchange.

The Windows Update Web site is located at:

http://windowsupdate.microsoft.com/

- **Default setting and ability to disable**: Update Root Certificates is installed by default with products in the Windows Server 2003 family. You can remove or exclude this component from installation on a server.

- **Trigger and user notification**: Update Root Certificates is triggered when the server is presented with a certificate issued by a root authority that is not directly trusted. There is no user notification.

- **Logging**: Events containing information such as the following will be logged:

  - Event ID: 7

    Description: Successful auto update retrieval of third-party root list sequence number from: *URL_for_Windows_Update_Web_Site*

  - Event ID: 8

    Description: Failed auto update retrieval of third-party root list sequence number from: *URL_for_Windows_Update_Web_Site* with error: *hexadecimal_error_value*

- **Encryption, privacy, and storage**: When requests or certificates are sent to or from Update Root Certificates, no encryption is used. Information about Update Root Certificates activity is not stored on any server at Microsoft.

- **Transmission protocol and port**: The transmission protocol is HTTP and the port is 80.

**Controlling the Update Root Certificates Component to Prevent the Flow of Information to and from the Internet**
If you want to prevent the Update Root Certificates component in the Windows Server 2003 family from communicating automatically with the Windows Update Web site, you can remove or exclude this component from installation on servers. You can do this during deployment by using standard methods for unattended installation or remote installation, as described in Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003). If you are using an answer file, the entry is as follows:

```
[Components]
Rootautoupdate = Off
```

**How removing or excluding Update Root Certificates from servers can affect applications**
If a server is presented with a certificate issued by a root authority that is not directly trusted, and the Update Root Certificates component is not installed on that server, you (or your application) will be prevented from completing the action that required authentication. For example, you might be prevented from installing software, viewing an encrypted or digitally signed e-mail message, or using a browser to engage in an SSL session.

If you choose to remove or exclude Update Root Certificates from servers, consider providing instructions that tell administrators what to do if they receive untrusted certificates.

**Procedures for Excluding or Removing the Update Root Certificates Component from an Individual Computer**
The following procedures describe:

- How to use Control Panel to remove the Update Root Certificates component from an individual computer running a product in the Windows Server 2003 family.

- How to exclude the Update Root Certificates component during unattended installation of a product in the Windows Server 2003 family by using an answer file.

**To remove the Update Root Certificates component from an individual computer running a product in the Windows Server 2003 family**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Add or Remove Programs**.

3. Click **Add or Remove Windows Components** (on the left).

4. Scroll down the list of components to Update Root Certificates, and make sure the check box for that component is cleared.

5. Follow the instructions to complete the Windows Components Wizard.

**To exclude the Update Root Certificates component during unattended installation by using an answer file**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file. For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003).

2. In the [Components] section of the answer file, include the following entry:

   **Rootautoupdate = Off**

# Device Manager and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Device Manager
- How Device Manager communicates with sites on the Internet
- How to control Device Manager to limit the flow of information to and from the Internet

Device Manager interacts with the Internet through the hardware wizards by contacting Windows Update when you install or update hardware. For procedures related to disabling Windows Update, see Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003) in this white paper.

**Benefits and Purposes of Device Manager**

Device Manager provides users and administrators with information about how the hardware on their computers is installed and configured, and how the hardware interacts with the computer's applications. With Device Manager, administrators can update the drivers (or software) for hardware devices, modify hardware settings, and troubleshoot problems.

| **Note** |
| --- |
| Only administrators or users with administrative credentials can install or update device drivers. |

**Overview: Using Device Manager in a Managed Environment**

In the Microsoft Windows Server 2003 family, you access Device Manager through Administrative Tools\Computer Management. Administrators or users with administrative credentials will typically use Device Manager to check the status of hardware and to update device drivers. Administrators who have a thorough understanding of computer hardware might also use Device Manager's diagnostic features to resolve device conflicts and change resource settings.

Device Manager works in conjunction with Windows Update to deliver updated drivers for installed hardware. When you are upgrading a driver or installing new hardware, and your computer has access to the Internet, your computer will automatically check Windows Update for an appropriate device driver. As an IT administrator in a highly managed environment you might want to block certain administrators from downloading drivers through Device Manager. You can do this by configuring Group Policy to disable Windows Update. If you disable Windows Update then Device Manager cannot communicate with the Internet. The following subsection provides details about how Device Manager interacts with the Internet.

**How Device Manager Communicates with Sites on the Internet**

Device Manager communicates with the Internet when an administrator uses it to update a driver. After you open Device Manager you select a hardware device and click **Update Driver** on the **Action** menu (or right-click a hardware device). This activates the Hardware Update Wizard. The way Device Manager communicates with the Internet is based on its interaction with Windows Update through the Hardware Update Wizard; therefore much of the information in this subsection is the same as for Windows Update. Additional details are described as follows:

- **Specific information sent or received**: See the section Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003) in this white paper.
- **Default and recommended settings**: Device Manager is enabled by default. See the subsection "Controlling Device Manager to Limit the Flow of Information to and from the Internet," for recommended settings.
- **Triggers**: Through Device Manager an administrator starts the Hardware Update Wizard, or adds new hardware to a computer.
- **User notification**: See Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003).

- **Logging**: Errors that result from problems installing hardware devices without drivers are logged to the event log.

- **Encryption, access, privacy statement, transmission protocol, and port**: See Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003).

- **Ability to disable**: You cannot disable Device Manager directly. You can, however, prevent interaction with the Internet through Device Manager by disabling Windows Update.

**Controlling Device Manager to Limit the Flow of Information to and from the Internet**

You can prevent the Internet from being accessed through Device Manager by disabling Windows Update or by configuring where computers search for drivers. You use Group Policy settings to perform both of these procedures.

If you disable automatic access to Windows Update, you can have selected administrators update drivers by manually downloading updates from the Windows Update Catalog, or from an intranet server, whereby they can be distributed on your managed network as needed.

For more information about the Windows Update Catalog, see the Windows Update Web site at:

http://windowsupdate.microsoft.com/

**Procedure for Controlling How Drivers Are Updated Through Device Manager**

Group Policy settings for Windows Update and the automatic updating function are in Computer Configuration\Administrative Templates\Windows Components and User Configuration\Administrative Templates\Windows Components. For the procedure to disable Windows Update or configure automatic updating, see the section Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003) in this white paper. The procedure to eliminate Windows Update as a driver search location using Group Policy is included here.

**To disable Windows Update as a driver search location**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **User Configuration**, click **Administrative Templates**, and then click **System**.

3. In the details pane, double-click **Configure Driver Search Locations**, and then click **Enabled**.

4. Select **Don't Search Windows Update**.

# Driver Protection and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Driver Protection
- How Driver Protection communicates with sites on the Internet
- How to control Driver Protection to limit the flow of information to and from the Internet

**Benefits and Purposes of Driver Protection**
The Driver Protection feature in the Microsoft Windows Server 2003 family prevents the operating system from loading drivers that are known to cause stability problems (for example, preventing the operating system from booting). These drivers are listed in a Driver Protection List database included with the operating system. Driver Protection checks this database during operating system upgrades and at run time. These checks are performed to determine whether to load a driver under one of the operating systems in the Windows Server 2003 family.

Driver Protection also displays up-to-date content about these driver problems in Help and Support Center, including links to sites where users can find a solution. Driver Protection relies on Windows Update and Dynamic Update to update the database files so that users are presented with the most current information available on protected drivers. Users cannot directly disable Driver Protection.

Drivers are added to the Driver Protection List based on feedback from end users about problems that can be reproduced and confirmed at Microsoft. The main reasons a driver is added to this list are:

- A Windows Server 2003 family operating system cannot boot with this driver loaded.
- Setup cannot be completed with this driver loaded.
- End users experience data corruption when this driver is loaded.

Decisions to add drivers to this list are made in consultation with the vendors who produce and distribute these drivers. Microsoft engages and informs these vendors before adding a driver to the Driver Protection List.

A listing of the content in the Driver Protection List for the Windows Server 2003 family is available as part of a white paper that provides additional information about Driver Protection on the Windows Platform Development Web site at:

http://go.microsoft.com/fwlink/?LinkId=33957

This section of the white paper explains how to control Driver Protection in a managed environment.

**Overview: Using Driver Protection in a Managed Environment**
Users have no direct control over whether to download files required by Driver Protection for updating the Driver Protection List. In a managed environment it is unlikely that users will be allowed to send and receive driver information freely; this function would normally be controlled in some fashion by the IT department. You can indirectly block Driver Protection from downloading files by disabling Windows Update or by avoiding the use of Dynamic Update. Details on the methods and procedures for controlling Driver Protection are described in the following subsections.

**How Driver Protection Communicates with Sites on the Internet**
This subsection summarizes the communication process:

- **Specific information sent or received**: No information is sent to the Internet about the user's system. Driver Protection downloads updated versions of the following files:
    - drvmain.sdb, apphelp.chm, apphelp.sdb, and apphelp.dll.

- **Default and recommended settings**: Driver Protection is enabled by default. Recommended settings are described in the next subsection, "Controlling Driver Protection to Limit the Flow of Information to and from the Internet."

- **Triggers**: Driver Protection is triggered if the device driver is on the Driver Protection List when the operating system starts, when a new application or device is installed, or during the installation or upgrade of the operating system.

- **User notification**: The notification that the user receives when Driver Protection is triggered differs according to when the driver load request occurs:

    - If a driver on the Driver Protection List is matched when the operating system starts, the operating system displays a pop-up Help balloon titled "Devices or Applications disabled," in the taskbar notification area when the user logs on. If the user clicks that Help balloon, additional driver information and links to solutions for that problem are displayed in Help and Support Center.

    - If a driver on the Driver Protection List is matched during the setup of a Windows Server 2003 family operating system (for an upgrade from Windows NT® 4.0 or Windows 2000), a message will appear in the Report System Compatibility window before the operating system upgrade is completed.

      Users have two options at this point:

      They can cancel Setup and find an alternate driver solution before installing the new operating system. If the driver that users install solves the problem, Setup will continue normally.

      They can continue the upgrade process without first installing a driver that solves the problem. In this case, Setup may disable the driver in order to be completed. When users later log on, the operating system displays the pop-up Help balloon described in the previous case.

      If a driver on the Driver Protection List is matched during installation of a new application or device, and that driver uses system installation services (SetupAPI), the operating system displays a notification during installation and blocks the installation of that driver.

      If a driver is not installed using system installation services, the operating system cannot block the installation of that driver. It can, however, block the driver from loading. If the driver is blocked, a notification will appear every time the operating system attempts to load that driver under an operating system in the Windows Server 2003 family. For example, if a CD writing program that does not use system installation services installs a driver that is included on the Driver Protection List, the Windows Server 2003 family operating systems will display the pop-up Help balloon mentioned previously after the setup for that program is completed.

- **Logging**: If Driver Protection finds a match for a driver in the Driver Protection List, operating systems in the Windows Server 2003 family log an event in the event log.

- **Encryption**: The data packages downloaded to the user's system by Microsoft are digitally signed.

- **Access**: No data is uploaded from the user's computer.

- **Privacy statement**: Driver Protection is covered by the same privacy statement that covers Windows Update.

- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is 80.

- **Ability to disable**: You cannot disable Driver Protection directly. Disabling Windows Update or avoiding the use of Dynamic Update will, however, block Driver Protection from updating the database files for the Driver Protection List on the server. (Of course you can also block the updating of Driver Protection database files by preventing access to the Internet, or by blocking HTTP over port 80.)

**Controlling Driver Protection to Limit the Flow of Information to and from the Internet**

You cannot disable Driver Protection directly. To block the downloading of updates for the Driver Protection database files, you

can disable the settings for Windows Update and (during setup) avoid the use of Dynamic Update. (Of course you can also block downloading by preventing access to the Internet, or by blocking HTTP over port 80.)

**How controlling Driver Protection can affect users and applications**

Driver Protection blocks known problem drivers from loading, but it does not block any associated applications that depend on those drivers. Therefore, the behavior of applications that depend on drivers that are blocked varies depending on the implementation of the application. Some applications, such as antivirus programs, install drivers in order to provide their core functionality. For these applications, Driver Protection may cause the application not to work at all. Other applications, such as CD-burning programs, use drivers for portions of their feature set. For these applications, only those features that do not depend on the driver may work.

If you decide to disable Driver Protection from pulling down updated versions of the Driver Protection List database, drivers that affect system stability will continue to be blocked. The operating system, however, will use the version of the Driver Protection List database that comes with the operating system to identify the drivers to block, instead of a more accurate, up-to-date version of the list.

**Alternate Methods for Controlling Driver Protection**

A more drastic measure to take would be to disable the Upload Manager service (uploadmgr) that manages synchronous and asynchronous file transfers between clients and servers on the network. Disabling this service will block the upload of the anonymous hardware profile data (although users will still be able to complete the Hardware Wizard). The operating system will, however, use the version of the Driver Protection List database that comes with the operating system to identify the drivers to block, instead of a more accurate, up-to-date version of the list. The following subsection gives the procedure for this method.

**Procedure for Disabling How Driver Protection Communicates over the Internet**

You cannot disable Driver Protection directly but can do so indirectly by controlling its ability to connect to the Internet by disabling Windows Update or avoiding the use of Dynamic Update. See the sections in this white paper titled Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003) and Dynamic Update and Internet Communication (Windows Server 2003) for more information about these methods.

As mentioned in the previous subsection, a more drastic method for disabling Driver Protection is to disable the Upload Manager service.

**To disable how Driver Protection communicates over the Internet by disabling the Upload Manager service**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Administrative Tools**, and then double-click **Services**.
3. In the details pane, right-click **Upload Manager**, and then click **Properties**.
4. Click the **Log On** tab, then click the hardware profile that you want to configure, and then click **Disable**.

| **Important** |
| --- |
| If this service is disabled, any services that explicitly depend on it will fail to start. |

# Dynamic Update and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Dynamic Update
- How Dynamic Update communicates with sites on the Internet
- How to control Dynamic Update to limit the flow of information to and from the Internet

**Benefits and Purposes of Dynamic Update**

With Dynamic Update, Setup for the Microsoft Windows Server 2003 family can check the Windows Update Web site for new Setup files, including drivers and other files, while the server operating system is being installed. In an interactive installation (in contrast to an unattended installation), the person installing a product in the Windows Server 2003 family chooses whether to allow Dynamic Update.

In a managed environment, if you are using Setup (Winnt32.exe) for unattended installation, you can create a shared folder on a server and deliver Dynamic Update files to destination computers from that shared folder (instead of downloading the files directly from the Windows Update Web site to the computer being installed). For additional information about how to do this, see "How Dynamic Update Communicates with Sites on the Internet," and "Controlling Dynamic Update to Limit the Flow of Information to and from the Internet," later in this section.

Whenever an important update is made to any crucial Setup file, that update is made available through Dynamic Update functionality built into the Windows Update Web site. Some of the updated files will be replacements (for example, an updated Setup file) and some will be additions (for example, a driver not available at the time that the Setup CD was created). All files on the Dynamic Update section of the Windows Update Web site are carefully tested, and only files that are important in ensuring that Setup runs well are made available.

Using Dynamic Update reduces the need to apply patches to recently installed systems, and makes it easier to run Setup with hardware that requires a driver that was recently added or updated on Windows Update. For example, if a new video adapter requires a driver that was recently added to Windows Update, with Dynamic Update, the video adapter is recognized and supported during Setup.

Dynamic Update downloads only the files that are required for a particular computer, which means that the Dynamic Update software briefly examines the computer hardware. No personal information is collected, and no information is saved. The only purpose for examining the hardware is to select appropriate drivers for it. This keeps the download time as short as possible and ensures that only necessary drivers are downloaded to the hard disk.

**Overview: Using Dynamic Update in a Managed Environment**

If you do not want Dynamic Update to connect to the Windows Update Web site during the installation of a product in the Windows Server 2003 family, you have two options:

- **Creating a shared folder and delivering Dynamic Update files to destination computers from that shared folder**: You can ensure that when Setup (Winnt32.exe) for the Windows Server 2003 family is run in your organization, Dynamic Update does not connect to the Internet but instead uses the files you place on a server. To do this, you create a shared folder on a server in your organization, download Dynamic Update files to that server, and run unattended installations using Winnt32.exe with appropriate options.
- **Avoiding Dynamic Update**: You can avoid using Dynamic Update, which means Setup will use only the files and drivers provided on the CD for products in the Windows Server 2003 family. For more information, see "Avoiding Dynamic Update," later in this section.

The subsections that follow provide more information about these options.

For additional sources of information about performing unattended installation, see Appendix A: Resources for Learning About

Automated Installation and Deployment (Windows Server 2003).

**How Dynamic Update Communicates with Sites on the Internet**

This subsection focuses on the communication that occurs between Dynamic Update and the Windows Update Web site during an interactive installation (or a preinstallation compatibility check) when the computer has access to the Internet. This subsection also provides some description of the default behavior of Dynamic Update with unattended Setup.

For information about how you can control the behavior of Dynamic Update during unattended installation, see "Controlling Dynamic Update to Limit the Flow of Information to and from the Internet," later in this section.

- **Specific information sent or received**: When Dynamic Update contacts the Windows Update Web site, it sends only the information necessary for appropriate drivers to be selected. In other words, it collects only necessary information about the hardware (devices) on that particular computer. No personal information is collected.

  The Setup files and drivers downloaded by Dynamic Update consist only of files that are important in ensuring that Setup runs successfully. Files with minor updates that will not significantly affect Setup are not made available through the Dynamic Update section of the Windows Update Web site. Some of the updated files will be replacements (for example, an updated Setup file) and some will be additions (for example, a driver not available at the time that the Setup CD was created).

- **Default behavior and triggers**: Dynamic Update may connect to the Internet, depending on how Setup is run. The following table provides details.

  **Choices for running Setup and effects on Dynamic Update**

| **Choice** | **Steps to take and effect on Dynamic Update** | **Does Dynamic Update connect to the Internet?** |
|---|---|---|
| Running a preinstallation compatibility check | Insert the Setup CD and choose the appropriate options for checking system compatibility. You are offered the choice of running or skipping Dynamic Update. | Yes, if you choose to run Dynamic Update. |
| Interactive installation | Start Setup from the CD or a network and run it interactively. You are offered the choice of running or skipping Dynamic Update. | Yes, if you choose to run Dynamic Update. |
| Unattended Setup without an answer file and without the use of any options that affect Dynamic Update | Run the Winnt32.exe command with the **/unattend** option, but do not provide the name of an answer file and do not specify **/dudisable** or any other options that affect Dynamic Update. Dynamic Update is triggered under these conditions for both unattended installation and unattended upgrade. | Yes. |
| Unattended Setup with the **/dudisable** option | Run the Winnt32.exe command with the **/unattend** option and also with the **/dudisable** option. If the **/dudisable** option is used, Dynamic Update is not triggered, regardless of whether an answer file is used. | No. |
| Unattended Setup with an answer file that specifies that Dynamic Update should not be disabled | Create an answer file that includes an **[Unattended]** section with an entry that specifies **dudisable = No**. Run the Winnt32.exe command with the **/unattend:***answer_file* option. Dynamic Update is triggered (although see the previous entry in this table). | Yes. |

| | | |
|---|---|---|
| Unattended Setup with an answer file that does not specify any options that affect Dynamic Update | Run the Winnt32.exe command with the **/unattend:***answer_file* option. By default, if the answer file does not specify any options that affect Dynamic Update, Dynamic Update is disabled. | No. |
| Unattended Setup without an answer file and with the **/dushare** option | Prepare a shared folder as outlined in "Creating a shared folder and delivering Dynamic Update files to destination computers from that shared folder," later in this section. When you run Winnt32.exe, run it with the **/dushare =** *path_to_downloaded_files* option. Dynamic Update uses the folder specified in the **/dushare** option and does not connect to the Internet. | No. Dynamic Update uses the files in the shared folder that you created. |
| Unattended Setup with an answer file that contains the DUShare entry | Prepare a shared folder as outlined in "Creating a shared folder and delivering Dynamic Update files to destination computers from that shared folder," later in this section. Create an answer file that includes an **[Unattended]** section with an entry that specifies **dushare =** *path_to_downloaded_files*. Run the Winnt32.exe command with the **/unattend:***answer_file* option. Dynamic Update uses the folder specified in the DUShare entry and does not connect to the Internet. | No. Dynamic Update uses the files in the shared folder that you created. |

- **User notification**: During an interactive installation, the user is notified when the choice of whether to run Dynamic Update is offered. During an unattended installation, there is no notification (unattended installation by definition means that no user interaction is required).

- **Logging**: By default, the progress of Setup is logged in *systemroot*\Winnt32.log. By using command options for the Winnt32.exe command, you can control the name of the log and the level of detail it contains.

- **Encryption**: The data is transferred from Microsoft using HTTPS.

- **Access**: No information about the hardware (devices) on a particular computer is saved or stored, so no one can access this information. The information is used only to select appropriate drivers.

- **Privacy statement**: Dynamic Update is covered by the same privacy statement that covers Windows Update. To view the privacy statement for Windows Update, go to the Web site and click **Microsoft Update Privacy Statement**:

  http://windowsupdate.microsoft.com/

- **Transmission protocol and port**: The transmission protocol is HTTPS and the port is 443.

- **Ability to disable**: You can control the behavior of Dynamic Update by running Setup in specific ways, as shown in the previous table. (You can of course disable Dynamic Update by preventing access to the Internet, or by blocking HTTPS over port 443.)

  If you do not want to disable Dynamic Update but only want to prevent it from communicating with an Internet site, as noted earlier, you can create a shared folder on a server and deliver Dynamic Update files to destination computers from that shared folder.

**Controlling Dynamic Update to Limit the Flow of Information to and from the Internet**

As summarized in "Overview: Using Dynamic Update in a Managed Environment," earlier in this section, if you do not want Dynamic Update to connect to the Windows Update Web site during the installation of a product in the Windows Server 2003 family, you have two options. With the appropriate methods for unattended installation, you can create a shared folder on a server and deliver Dynamic Update files to destination computers from that shared folder. Another alternative is to completely avoid using Dynamic Update.

**Creating a shared folder and delivering Dynamic Update files to destination computers from that shared folder**

This subsection briefly describes the steps for creating a shared folder on a server and delivering Dynamic Update files to destination computers from that shared folder. The subsection also provides links to more detailed information. The steps can be summarized as follows:

- Step 1: Determine what packages you need to download from the Windows Update Web site.

- Step 2: Download the packages and prepare them and the folder they are in for use with Dynamic Update. This step includes extracting files and placing them in folders, as well as running the **/duprepare** option with Winnt32.exe, which creates subfolders and copies appropriate files to those subfolders. This step also requires other actions, such as sharing the folder and setting permissions.

- Step 3: Configure the answer file and Winnt32.exe settings for Dynamic Update (and for any other configuration options you want).

- Step 4: Run the unattended installations.

For more detailed information about performing the preceding steps, see the *Microsoft Windows Server 2003 Deployment Kit*, specifically the book titled *Automating and Customizing Installations*. To view the *Microsoft Windows Server 2003 Deployment Kit*, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?linkid=29887

Similar information is available in the Dynamic Update article on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=4925

For additional sources of information about performing unattended installation, see Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003).

**Avoiding Dynamic Update**

You can avoid using Dynamic Update, which means Setup will use only the files and drivers provided on the CD for products in the Windows Server 2003 family. The method by which you avoid using Dynamic Update depends on how you are performing the installation:

- **Interactive installation**: During interactive installation (not unattended installation), you can select No when offered the option to use Dynamic Update. As an alternative, you can ensure that the computer does not have Internet access.

- **Unattended Setup**: Dynamic Update is disabled when you run Setup in specific ways, as shown in the table in "How Dynamic Update Communicates with Sites on the Internet," earlier in this section. As the table shows, the simplest way to ensure that Dynamic Update does not run during unattended Setup is to use the **/dudisable** option in the command line. This ensures that Dynamic Update will not occur during the installation.

**How avoiding Dynamic Update or directing Dynamic Update to a server on your network can affect users and applications**

Regardless of whether you use Dynamic Update, you can obtain updated system and driver files after installations are complete (for example, through Windows Update or a service pack). Allowing Dynamic Update to run during Setup, however, helps ensure Setup success.

If you create a shared folder on a server and deliver Dynamic Update files to destination computers from that shared folder (instead of downloading the files directly from Windows Update to the computers), you can control the exact set of updated files to be installed. By contrast, when you download the current set of Dynamic Update files directly from the Windows Update Web site to users' computers, you might introduce inconsistencies among your destination computers because the Windows Web Site is periodically updated, and you cannot control when these updates occur.

**Procedures for Controlling Dynamic Update**

For detailed descriptions of Dynamic Update and procedures for controlling it, see the *Microsoft Windows Server 2003 Deployment Kit*, specifically the book titled *Automating and Customizing Installations*. To view the *Microsoft Windows Server 2003 Deployment Kit*, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?linkid=29887

Similar information is available in the Dynamic Update article on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=4925

# Event Viewer and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Event Viewer
- How Event Viewer communicates with sites on the Internet
- How to control Event Viewer to prevent the flow of information to and from the Internet

**Benefits and Purposes of Event Viewer**

Using Event Viewer, administrators can view and set logging options for event logs in order to gather information about hardware, software, and system problems. By default, a computer running an operating system in the Microsoft Windows Server 2003 family records events in three kinds of logs:

- **Application log**: The application log contains events logged by applications or programs. For example, a database program might record a file error in the application log. Application developers decide which events to log.
- **Security log**: The security log records events such as valid and invalid logon attempts, as well as events related to resource use such as creating, opening, or deleting files or other objects. For example, if logon auditing is enabled, attempts to log on to the system are recorded in the security log.
- **System log**: The system log contains events logged by Windows system components. For example, the failure of a driver or other system component to load during startup is recorded in the system log. The event types logged by system components are predetermined by the server.

A computer running a Windows Server 2003 family operating system which is configured as a domain controller records events in two additional logs:

- **Directory service log**: The directory service log contains events logged by the Windows Active Directory directory service. For example, connection problems between the server and the global catalog are recorded in the directory service log.
- **File Replication service log**: The File Replication service log contains events logged by the Windows File Replication service. For example, file replication failures and events that occur while domain controllers are being updated with information about system volume changes are recorded in the file replication log.

A computer running a Windows Server 2003 operating system configured as a Domain Name System (DNS) server records events in an additional log. The DNS server log contains Windows DNS service events.

Other types of events and event logs might be available on a computer, depending on what services are installed.

**Overview: Using Event Viewer in a Managed Environment**

The Event Log service starts automatically when you start the operating system. Administrators access event logs for a server through Control Panel\Administrative Tools\Event Viewer. They can obtain detailed information about a particular event by either double-clicking the event, or by selecting the event and clicking **Properties** on the **Action** menu. The dialog box gives a description of the event, which can contain one or more links to Help.

Links can either be to servers at Microsoft, or to servers managed by the software vendor for the component that generated the event. On products in the Windows Server 2003 family, most events that originate from Microsoft products will have standard text containing a URL at the end of the description ("For more information, see Help and Support Center at go.microsoft.com/fwlink/events.asp").

When you click the link, you are asked to confirm that the information presented can be sent over the Internet. If you click Yes, the information listed will be sent to the Web site named in the link. The parameters in the original URL will be replaced by a standard list of parameters whose contents are detailed in the confirmation dialog box. This list is provided in the next subsection under "Specific information sent or received."

In a highly managed environment, IT administrators might want to prevent users and administrators from sending this information over the Internet through this link and accessing a Web site. In the Windows Server 2003 family, this information flow is governed by a registry key. Administrators can edit this registry key to prevent users and administrators from accessing the Internet through Event Viewer.

**How Event Viewer Communicates with Sites on the Internet**

In order to access the relevant Help information provided by the link in the Event Properties dialog box, you must send the information listed about the event. The collected data is confined to what is needed to retrieve more information about the event from the Microsoft Knowledge Base. User names and e-mail addresses, names of files unrelated to the logged event, computer addresses, and any other forms of personally identifiable information are not collected.

The exchange of information that takes place over the Internet is described as follows:

- **Specific information sent or received**: Information about the event sent over the Internet includes the following:
    - Company name (software vendor)
    - Date and time
    - Event ID (for example, 1704)
    - File name and version (for example, userenv.dll, 5.1.2600.1106)
    - Product name and version (for example, Microsoft Windows Operating System, 5.1.2600.1106)
    - Registry source (for example, userenv)
    - Type of event message (for example, Error)

  The information the user receives is from the Web site named in the link.

- **Default settings**: Access to Event Viewer is enabled by default.
- **Triggers**: The user chooses to send information about the event over the Internet in order to view Help.
- **User notification**: When a user clicks the link, a dialog box listing the information that will be sent is provided.
- **Logging**: This is a feature of Event Viewer.
- **Encryption**: The information may or may not be encrypted, depending on whether it is an HTTP or HTTPS link.
- **Access**: No information is stored.
- **Privacy statement**: See the Windows Server 2003 family Help for a privacy statement. (In Help and Support, type **Linking to Microsoft for Help and Support**.)
- **Transmission protocol and port**: Communication occurs over the standard port for the protocol in the URL, either HTTP or HTTPS.
- **Ability to disable**: The ability to send information over the Internet or to be linked to a Web site can be prevented by editing the registry.

**Controlling Event Viewer to Prevent the Flow of Information to and from the Internet**

You can prevent users and administrators from sending information across the Internet and accessing Internet sites through Event Viewer by editing the registry. When you edit the registry as described in the following subsection, clicking Yes as previously described will still start Help, but it will not access the Internet for information specific to the event.

The Windows Server 2003 family computer registry values listed in this subsection are located in the following registry key:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\Current Version\Event Viewer

The following list describes how this registry key controls the flow of information to and from the Internet.

- **MicrosoftRedirectionProgram**

  **Default value**: %SystemRoot%\PCHealth\HelpCtr\Binaries\HelpCtr.exe

  **Usage**: This program is started with command-line parameters from MicrosoftRedirectionProgramCommandLineParameters

- **MicrosoftRedirectionProgramCommandLineParameters**

  **Default value**: -url hcp://services/centers/support?topic=%s

  **Usage**: "%s" is replaced with the URL in the link

- **MicrosoftRedirectionURL**

  **Default value**: http://go.microsoft.com/fwlink/events.asp

  **Usage**: Governs the text of the standard link for Microsoft events

| **Note** |
| --- |
| If any of these registry values is missing or empty, the link will be started directly using ShellExecute; deleting these values is not a method for preventing information from reaching the Internet. |

**Procedures for Preventing the Flow of Information to and from the Internet Through Event Viewer**
To prevent the flow of information to and from the Internet through Event Viewer you need to edit the registry. You can then apply the registry change to computers in a domain using Group Policy.

**Editing the registry**
Edit HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\Current Version\Event Viewer as follows:

To prevent the user from accessing the Internet when they click the link, delete the final "%s" from the value of MicrosoftRedirectionProgramCommandLineParameters (see the list in the previous subsection). With this change, clicking the link and clicking Yes will still start Help, but it will not access the Internet for information specific to this event.

For more information about the registry, see the *Registry Reference for Windows Server 2003* on the *Microsoft Windows Server 2003 Deployment Kit* companion CD, or on the Windows Server 2003 Web site at:

http://go.microsoft.com/fwlink/?LinkId=428

| **Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied. |

**Distributing the registry change using Group Policy**
You can distribute this registry change to computers in a domain by configuring a Group Policy object (GPO). You first need to create a template using the Event Viewer snap-in as described in the following procedure.

**To enable the Event Viewer Group Policy snap-in**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **User Configuration**, click **Administrative Templates**, and then click **Windows Components**.

3. Click **Microsoft Management Console** and then click **Restricted/Permitted snap-ins**.

4. In the details pane under Setting, double-click **Event Viewer**.

5. In the Event Viewer Properties dialog box, select **Enabled**.

# File Association Web Service and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of the file association Web service
- How the file association Web service communicates with sites on the Internet
- How to control the file association Web service to limit the flow of information to and from the Internet

**Benefits and Purposes of the File Association Web Service**
In products in the Microsoft Windows Server 2003 family, the file association Web service extends the scope of information stored locally by the operating system about file name extensions, file types, and the applications or components to use when opening a particular file type. Both the locally stored information and the file association Web service are intended to provide you with the ability to open (double-click) a file without having to specify which application or component to open it with. The operating system associates the file name extension (for example, .txt or .jpg) with a file type, and it opens each file type with the application or component specified for that file type. For example, file name extensions .htm and .html are both "HTML Document" file types.

The operating system first checks for the file association information locally. If no local information is available about the file name extension and its associated file type, the operating system offers you the option of looking for more information on a Microsoft Web site. For details about the URL for this Web site, see "How the File Association Web Service Communicates with Sites on the Internet," later in this section.

**Overview: Using the File Association Web Service in a Managed Environment**
To limit the flow of information from the file association Web service to the Internet, you have several options. You can use firewall settings, you can disable the file association Web service by setting a registry key, and you can configure automatic server-based software installation through Group Policy. You can also use scripts to limit the file types that can be stored, viewed, or used on computers in your organization, which will limit the likelihood that anyone will need to obtain information about those file types.

**How the File Association Web Service Communicates with Sites on the Internet**
The file association Web service communicates with sites on the Internet as follows:

- **Specific information sent or received**: If the operating system does not find local information about a file name extension, it offers you the option of sending a query to look for more information on a Microsoft Web site. The site is language-specific; the file name extension that you double-click is appended to the query. The query takes the following form:

  **http://shell.windows.com/fileassoc/***nnnn***/xml/redir.asp?Ext=***AAA*

  where *nnnn* is a hexadecimal value used in the Windows Server 2003 family to map to a language identifier (that is, to an RFC1766 identifier), and *AAA* is the file name extension for which information is needed. An example of a hexadecimal value and its corresponding language identifier is 0409 for en-us, English (United States).

  For more information about these hexadecimal values, see information about the multiple language (MLang) registry settings on the MSDN Web site at:

  http://go.microsoft.com/fwlink/?linkid=29165

  To search for information about MLang registry settings or the Microsoft Internet Explorer Multiple Language application programming interface (MLang API), use the Search tool on the MSDN Web site at:

  http://msdn.microsoft.com/

- **Default setting and ability to disable**: The service is enabled by default. It can be disabled by setting a registry key, as described in "Disabling the file association Web service," later in this section.

  There are ways of reducing the likelihood that a person will trigger the file association Web service. One basic way is to configure automatic, server-based software installation based on Group Policy settings. For more information, see "Finding information about the Software Installation extension of Group Policy," later in this section.

- **Trigger and notification**: When you try to open a file (for example, by double-clicking the file), and there is no local information about the correct application or component to use when opening the file, the operating system offers the option either to "Use the Web service to find the appropriate program" or to "Select the program from a list."

- **Logging**: No events are logged by the file association Web service.

- **Encryption, storage, access, and privacy**: The file name extension sent in a query to the Internet is not encrypted. Nothing in the query identifies the person who triggered the query. If the local computer's browser is configured to store information about recently visited Internet sites, the browser will store the query containing the file name extension. Otherwise, the query containing the file name extension is not stored anywhere.

- **Transmission protocol and port**: The transmission protocol is HTTP and the port is 80.

**Controlling the File Association Web Service to Limit the Flow of Information to and from the Internet**

If you want to limit the flow of information from the file association Web service to the Internet, you can use one or more of the following methods:

- Use your firewall to block access to any Web site that contains the following string:

  http://shell.windows.com/fileassoc/

- Disable the file association Web service by setting a registry key, as described in "Disabling the file association Web service," later in this section.

- Configure automatic, server-based software installation. To do this, configure one or more servers with the Software Installation extension of Group Policy. When you do this, if someone tries to open a file for which the corresponding application is not installed locally, a copy of the application (stored on another server) is installed automatically. In this situation, the file association Web service will not be triggered. To learn more about the Software Installation extension, see "Finding information about the Software Installation extension of Group Policy," later in this section.

- Familiarize yourself and other administrators (if you are not already aware) with using Control Panel, Folder Options, and the File Types tab in Folder Options to associate a file name extension with a file type, and a file type with an application. Also, if a message box appears offering the two options, "Use the Web service to find the appropriate program" or "Select the program from a list," always click "Select the program from a list."

- Use scripts to scan your organization's computers for file types that you do not want stored, viewed, or used on your organization's computers. Take actions to ensure that these files do not remain on individual computers' hard disks. Reducing the number of file types on hard disks reduces the likelihood that the file association Web service will be triggered.

**Procedures that Limit Internet Communication Generated by the File Association Web Service**

This subsection contains the following information:

- A procedure for disabling the file association Web service by setting a registry key.

- A link to information about configuring automatic, server-based software installation through the Software Installation extension of Group Policy.

- Procedures for using the File Types tab in Folder Options to associate a file name extension with a file type, and a file type

with an application.

**Disabling the file association Web service**

The following procedure explains how to disable the file association Web service by setting a registry key.

**To disable the file association Web service by setting a registry key**

1. Open Registry Editor by clicking **Start**, clicking **Run**, and then typing **regedit**.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied. |

2. Navigate to the following registry key:

   HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies\System

3. Click the **Edit** menu, point to **New**, and then click **DWORD** Value.

4. Type the following name:

   **NoInternetOpenWith**

5. Click the new entry (**NoInternetOpenWith**), and then select **Modify** in the **Edit** menu.

6. Ensure that **Hexadecimal** is selected, and then for **Value data**, type:

   **1**

7. Close Registry Editor.

**Finding information about the Software Installation extension of Group Policy**

If you are not already familiar with using the Software Installation extension of Group Policy, use one of the following procedures to learn more. For additional information about Group Policy, see Appendix B: Resources for Learning About Group Policy (Windows Server 2003).

**To find information about the Software Installation extension with the Group Policy Management Console**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then click **Group Policy Management**.

3. Expand items in the Group Policy Management Console until the object to which you want to apply Group Policy is displayed. For more information, see the resources listed in Appendix B: Resources for Learning About Group Policy (Windows Server 2003).

4. Right-click the appropriate Group Policy object and then click **Edit**.

5. In the Group Policy Object Editor, in the **Help** menu, click **Help Topics**.

6. Click the **Contents** tab, click **Group Policy Management**, click **Concepts**, click **Group Policy Object Editor Extensions**, and then click **Software installation**.

**To find information about the Software Installation extension without the Group Policy Management Console**

1. On a computer running a product in the Windows Server 2003 family, click **Start** and then click **Help and Support**.

2. Under **Help Contents**, click Common Administrative Tasks, and then click **Deploying and upgrading software**.

   You can also view Help on the Web at:

   http://go.microsoft.com/fwlink/?LinkId=29881

**Specifying associations between file name extensions, file types, and applications or components**

When you associate a file name extension with a file type and an application or component, the result is that the file association Web service cannot be triggered by that file name extension on that computer.

**To associate a file name extension with a file type**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Folder Options**, and then click the **File Types** tab.

3. Click **New**.

4. Type a new or existing file name extension, and then click **Advanced**.

5. In **Associated File Type**, do one of the following:

   - Type or select **New** to create a file type to associate with the file name extension.

   - Type or select a different file type to associate with the extension.

| **Note** |
| --- |
| When you type a file name extension in the Create New Extension dialog box, the Associated File Type list displays the file type that is associated with that extension. To select New, scroll to the top of the list. |

**To associate a file name extension with an application**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Folder Options**, and then click the **File Types** tab.

3. Under **Registered file types**, click a file type.

4. Click **Change**, and then choose the application you want to use to open this file.

# Help and Support Center: Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

**Help and Support Center: The Headlines and Online Search Features**
This section provides information about:

- The benefits of the Headlines and Online Search features in Help and Support Center
- How the Headlines and Online Search features communicate with sites on the Internet
- How to control the Headlines and Online Search features to limit the flow of information to and from the Internet

**Benefits and Purposes of Headlines and Online Search**
Help and Support Center is a self-help portal that was first included in Microsoft Windows Millennium Edition. It is also included in all versions of the Windows Server 2003 family. You can access Help and Support Center in a number of ways, including:

- Selecting Help and Support from the Start menu.
- Selecting Help and Support from the Help menus for Control Panel, Windows Explorer, My Network Places, My Pictures, My Computer, or My Documents.

**Headlines**
A useful feature of Help and Support Center is the Headlines area. This area is typically titled "Top Issues" and is usually located in the lower-right corner of the main window, unless the window has been customized by the OEM or modified for certain languages. A page in Help and Support Center with more Headlines is exposed when you click a "View more Headlines" hyperlink at the bottom of the "Top Issues" area. Headlines provides a dynamic source of content that you can visit to find help and support on current issues as well as those that were known at the time the operating system was released. For example, it may display links to topics about new security bulletins, software updates, or new Help content.

**Online Search**
Online Search, another useful feature of Help and Support Center, enables you to query online Web sites automatically when performing a search. By default, the Microsoft Knowledge Base is designated as one of the Web sites for online searches. OEMs often customize the Online Search feature by, for example, adding a check box to the search window to enable the search engine to query their OEM-specific Web sites for results. To produce the most informative results when querying the Microsoft Knowledge Base, certain information such as the version of the product installed is collected from the computer and uploaded to the servers hosting the Microsoft Knowledge Base.

**Overview: Using Headlines and Online Search in a Managed Environment**
By creating a system registry key, or by performing unattended installation with an appropriate entry in your answer file, or using other tools available in the operating system user interface, you can control the extent to which the Headlines and Online Search features access the Internet. More details on the methods and procedures for controlling these features are described in the following subsections.

**How Headlines and Online Search Communicate with Sites on the Internet**
**Headlines**
The Headlines area is updated only when there is Internet connectivity. You are not required or prompted to connect to the Internet. Help and Support Center uses information contained in the NewsSet.xml file (stored in the *systemroot*\pchealth\helpctr\Config folder) to determine:

- Whether or not to update the Headlines area
- How frequently to update the Headlines area
- Where on the Internet to obtain the Headlines updates

This subsection summarizes the communication process:

- **Specific information sent or received**: If there is Internet connectivity, when you start Help and Support Center the Help and Support service (helpsvc) compares the current date to the date specified by the TIMESTAMP attribute in the NewsSet.xml file and calculates the total number of days that have elapsed since the last time Headlines was successfully updated.

  Then, if the number of elapsed days is greater than the number of days specified by the FREQUENCY attribute in NewsSet.xml, the Help and Support service connects to the Web site specified by the URL attribute and downloads an updated version of the file NewsVer.xml to the *systemroot*\pchealth\helpctr\Config\News folder. The administrator (or other person at the server) is not uniquely identified.

  | **Note** |
  | --- |
  | For Headlines supplied by Microsoft, the URL attribute is:<br><br>http://go.microsoft.com/fwlink/?LinkID=11 |

  The downloaded NewsVer.xml file contains links to the news content files (known as news blocks) for Windows Server 2003 operating systems and the installed language. These news blocks contain the information used to update the Headlines area, that is, links to and descriptions of the latest information from Help and Support Center, Windows, or support-related articles posted on Microsoft Web sites.

  | **Note** |
  | --- |
  | If the OEM has customized the Headlines feature, then the OEM-supplied Headlines may have links to the OEM's Web site. |

  If there is no Internet connectivity, Help and Support Center displays an offline message in the Headlines area similar to the following:

  When you are connected to the Internet, this area will display links to timely help and support information. If you want to connect to the Internet now, start the New Connection Wizard and see how to establish a Web connection through an Internet service provider.

- **Default and recommended settings**: The Headlines feature is enabled by default. Recommended settings are described in the next subsection, "Controlling Headlines and Online Search to Limit the Flow of Information to and from the Internet."

- **Triggers**: The Headlines feature is automatically triggered if there is Internet connectivity when you start Help and Support Center.

- **User notification**: You are not given the choice to select whether to update the Headlines area before an update is performed. An "Updating ..." status indicator is displayed in the Headlines area, however, to indicate when an update is being performed. Once Help and Support Center has completed checking for new headlines, the Headlines area is labeled "Updated: *date*," where *date* is the current date.

- **Logging:** There is no information related to Headlines entered into the event log.

- **Encryption:** The data transferred to Microsoft is not encrypted.

- **Access**: The only data generated on servers at Microsoft from the process of updates to the Headlines area is a single number telling how many times a connection has been made, by any computer, to the link that supports Headlines updates. No computer is identified in the process of a Headlines update. The data can be viewed by the Microsoft group

that provides support for the link through which Headlines is updated.

- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is 80.

- **Ability to disable**: You can disable Headlines by setting a registry key or during unattended installation. For more information, see "Procedures for Disabling Headlines and Online Search," later in this section.

**Online Search**

Online Search can only query online Web sites like the Microsoft Knowledge Base when there is Internet connectivity; you are neither required nor prompted to connect to the Internet. When you perform a search in Help and Support Center, if you have set the search options to search the Microsoft Knowledge Base or an OEM-designated Web site, the search engine automatically searches the specified site.

This subsection summarizes the communication process:

- **Specific information sent or received**: To produce relevant results when querying the Microsoft Knowledge Base, certain information is collected from the computer and uploaded to a server at Microsoft that hosts the Microsoft Knowledge Base. The administrator (or other person at the server) is not uniquely identified. Following is a list of the information collected:

    - The search text string that was entered

    - The language code of the operating system

    - The product Knowledge Base to be searched (for example, the Windows Server 2003 family or Outlook)

    - The version of the operating system installed (for example, Standard Edition, Enterprise Edition, Datacenter Edition, or Web Edition)

    - The number of results specified for the result set

    - Titles field status (indicates whether or not to search the article title only)

    - Type field status (indicates whether to search using "all" or "any" of the search string)

- **Default and recommended settings**: Online Search is enabled by default. Recommended settings are described in the next subsection, "Controlling Headlines and Online Search to Limit the Flow of Information to and from the Internet."

- **Triggers**: Online Search is automatically triggered if you use the default search options or if you have set the search options to encompass searches on the Internet. (Online Search is also dependent on having Internet connectivity when the search is performed.)

- **User notification**: You are not notified before Help and Support Center performs Online Search. A permanent headline is provided in the Headlines area that tells how to set Online Search options, including how to turn the feature off.

- **Logging:** There is no information related to Online Search entered into the event log.

- **Encryption:** The data transferred to Microsoft is not encrypted.

- **Access**: The data uploaded to the server is aggregated and clustered. Information about the most common queries is later made available to the Windows Product Support Services and Windows User Assistance teams to help in developing new content or in revising existing content.

- **Privacy statement**: Microsoft does not retrieve any personally identifiable information from the computer during an online search. A permanent headline is provided in the Headlines area that tells how to set Online Search options, including how to turn the feature off.

- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is 80.

- **Ability to disable**: You can disable Online Search through the Help and Support Center user interface.

**Controlling Headlines and Online Search to Limit the Flow of Information to and from the Internet**
Using the appropriate system registry key, or performing unattended installation with an appropriate entry in your answer file, you can disable the Headlines feature and eliminate the entire "Top Issues" area in the Help and Support Center user interface. If you perform unattended installation, the answer file entry is as follows:

```
[PCHealth]
Headlines = 0
```

The following table describes this and other configuration options for both Headlines and Online Search. For more information, see "Procedures for Disabling Headlines and Online Search," later in this section.

**Configuration settings for Headlines and Online Search**

| **Headlines: Configuration tool** | **Setting** | **Result** |
|---|---|---|
| Unattended Installation | Include Headlines = 0 under the [PCHealth] section. | Replaces text in the Headlines area with white space. |
| Registry | Set HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\PCHealth\HelpSvc\Headlines equal to 0. | Replaces text in the Headlines area with white space. |
| **Online Search: Configuration tool** | **Setting** | **Result** |
| Help and Support Center user interface (Set Search Options pane) | Clear any check boxes for querying the Microsoft Knowledge Base or OEM Web sites for results. | Disables online searches. The search results window neither displays an area for the online Web sites (such as the Microsoft Knowledge Base) nor returns any search results from online Web sites. |
| **Headlines and Online Search: Configuration tool** | **Setting** | **Result** |
| Firewall | Block HTTP port 80. | Displays offline message text in the Headlines area. The search results window displays an area for the online Web sites (such as the Microsoft Knowledge Base), but it does not return any search results from online Web sites. |

**Alternate Methods for Controlling Headlines and Online Search**
You can configure your firewall to restrict access to the Internet through HTTP port 80. You can also use the firewall to block updates to the Headlines area and to block online searches. In this scenario, for the Headlines feature, Help and Support Center displays the offline message text described in "How Headlines and Online Search Communicate with Sites on the Internet." When traffic through HTTP port 80 is blocked in this way, Help and Support Center searches will only query local Help content. The search results window will display an area for the online Web sites, such as the Microsoft Knowledge Base, but it will not contain any results.

**Procedures for Disabling Headlines and Online Search**
The following procedures explain how to:

- Disable the Headlines feature on individual computers by modifying the system registry.

- Disable the Headlines feature during workstation deployment by using standard methods for unattended installation or remote installation.

- Disable online searches in the Help and Support Center user interface.

**To disable the Headlines feature on individual computers by modifying the system registry**

1. Open Registry Editor by clicking **Start**, clicking **Run**, and then typing **regedit**.

> **Caution**
> Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. In the registry tree (on the left), navigate to the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\PCHealth\HelpSvc\.

3. On the **Edit** menu, point to **New**, and then click **DWORD value**.

4. Type **Headlines** as the name for the new value (type REG_DWORD), and then press ENTER.

> **Note**
> Setting the data value to 0 (or leaving the default for a new REG_DWORD) disables Headlines. If the Headlines REG_DWORD has another value or doesn't exist, then Headlines is enabled.

**To disable the Headlines feature during unattended installation by using an answer file**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file. For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003).

2. In the [PCHealth] section of the answer file, include the following entry:

   **Headlines = 0**

> **Note**
> Headlines = 0 specifies that no information is displayed in the Headlines ("Top Issues") area on the Help and Support Center home page. Headlines = 1 specifies that Headlines is displayed.

**To disable the Online Search feature in the Help and Support Center user interface**

1. Click **Start**, and then click **Help and Support**.

2. Below the Search box, click **Set search options**.

3. Clear the **Microsoft Knowledge Base** check box and any other check boxes below it. (For example, the manufacturer of your computer may have added a check box to enable your search to query their Web site for results.)

4. Close Help and Support Center.

# HyperTerminal and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of HyperTerminal
- How HyperTerminal communicates with sites on the Internet
- How to control HyperTerminal to prevent the flow of information to and from the Internet

**Benefits and Purposes of HyperTerminal**

HyperTerminal is a program that you can use to connect to other computers, Telnet sites, bulletin board systems (BBSs), online services, and host computers. HyperTerminal connections are made using a modem, a null modem cable (used to emulate modern communication), or an Ethernet connection.

HyperTerminal has capabilities beyond making connections to other computers. It can, for example, transfer large files from a computer onto your portable computer using a serial port rather than requiring you to set up your portable computer on a network. It can help debug source code from a remote terminal. It can also communicate with many older, character-based computers.

HyperTerminal records the messages passed to and from the computer or service on the other end of your connection. It can therefore serve as a valuable troubleshooting tool when setting up and using your modem. To make sure that your modem is connected properly or to view your modem's settings, you can send commands through HyperTerminal and check the results. HyperTerminal also has scroll functionality that enables you to view received text that has scrolled off the screen.

| **Note** |
| --- |
| HyperTerminal is designed to be an easy-to-use tool yet it is not meant to replace other full-featured tools. You can use HyperTerminal as described in this subsection, but you should not attempt to use HyperTerminal for more complex communication. For more information about what HyperTerminal does and does not support, see the list of frequently asked questions on the Hilgraeve Web site at http://go.microsoft.com/fwlink/?LinkId=225.<br><br>(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.) |

**Overview: Using HyperTerminal in a Managed Environment**

In a managed environment, providing access to local and remote connection points through the use of HyperTerminal may pose security risks. You can prevent HyperTerminal from being installed on products in the Microsoft Windows Server 2003 family as described later in this section. Following are a few security issues to consider when deciding how to configure HyperTerminal for your organization:

- **Viruses**: Incoming files might contain viruses or malicious programs that could compromise or destroy data on your computer. To reduce this risk, use virus-scanning software and ensure that incoming files are from a reliable and trusted source.
- **ID and password**: HyperTerminal cannot automatically provide your login ID and password when you make a connection. If you provide a password when using HyperTerminal for a Telnet session, be aware that this password will be sent to the remote computer using plaintext (as with all Telnet connections).
- **Automatic download**: The automatic download feature of the Zmodem protocol can pose a security risk by allowing remote users to send files to your computer without your explicit permission. To avoid this risk, you should select a protocol other than Zmodem in the Receive File dialog box or you should clear the **Allow remote host-initiated file transfers** check box on the Settings tab of Connection Properties.

Complete information about concepts and procedures associated with using or configuring HyperTerminal is beyond the scope of this white paper. For more information, access the HyperTerminal Help documentation in Help and Support Center on any computer running a product in the Windows Server 2003 family.

**How HyperTerminal Communicates with Sites on the Internet**

The exchange of information that takes place during the HyperTerminal connection is described as follows:

- **Specific information sent or received**: When using HyperTerminal for Telnet connectivity, the user ID and password are sent in plaintext format (as with all Telnet connections). If files are being transmitted, only the protocol, terminal emulation data, and file-specific binaries are sent. The computer running HyperTerminal is identified by its IP address when the connection type is TCP/IP. The computer is not uniquely identified when the connection type is not TCP/IP.

- **Default settings**: HyperTerminal is not installed by default when a product in the Windows Server 2003 family is installed manually. HyperTerminal is installed by default, however, if the Windows Server 2003 product is installed using an answer file during automated installation. To remove or uninstall HyperTerminal, see "Controlling HyperTerminal to Prevent the Flow of Information to and from the Internet," later in this section.

- **Triggers**: When HyperTerminal is set to automatically answer incoming connections, a file transfer can be initiated if the Zmodem transfer protocol is used. Otherwise, communication through HyperTerminal is only triggered when the user deliberately initiates it.

- **User notification**: After a user starts a HyperTerminal connection session, the status of the connection that is currently open within HyperTerminal is displayed in the HyperTerminal title bar. The status of the file and text transfer is displayed in the HyperTerminal window during the transfer process. HyperTerminal does not display connection or transfer status information when the automatic download feature of the Zmodem protocol is used. For more information about the HyperTerminal automatic download feature, see "Overview: Using HyperTerminal in a Managed Environment," earlier in this section.

- **Encryption:** Information sent or received by HyperTerminal is not encrypted.

- **Privacy statement and access**: No information is uploaded from the server to Microsoft.

- **Transmission protocol and port**: The protocols used are Kermit, Xmodem, 1K Xmodem, Ymodem, Ymodem-G, and Zmodem transmissions protocols on port 23.

- **Ability to disable**: You can remove HyperTerminal if it is already installed, or if you are using unattended installation for your operating system deployment, you can configure the answer file not to install HyperTerminal. To remove HyperTerminal, see "Controlling HyperTerminal to Prevent the Flow of Information to and from the Internet," later in this section.

**Controlling HyperTerminal to Prevent the Flow of Information to and from the Internet**

You can prevent the use of HyperTerminal by disabling it through unattended installation during operating system deployment, or by removing HyperTerminal after installing a product in the Windows Server 2003 family.

The following procedures describe:

- How to specify the HyperTerminal installation options in an answer file that will prevent HyperTerminal from being installed during an unattended installation of a product in the Windows Server 2003 family.

- How to use the Add or Remove Programs utility in Control Panel to remove HyperTerminal after the deployment of a product in the Windows Server 2003 family.

**To prevent HyperTerminal from being installed during unattended installation by using an answer file**

- Using the methods you prefer for unattended installation or remote installation, create an answer file. For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003).

- In the [Components] section of the answer file, include the following entry:

  **hyperterm = Off**

**To remove HyperTerminal on an individual computer running a product in the Windows Server 2003 family**

1. Click **Start**, and then either point to **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components** (on the left).
4. Double-click **Accessories and Utilities,** and then double-click **Communications**.
5. Make sure the check box for the HyperTerminal component is cleared.
6. Follow the instructions to complete the Windows Components Wizard.

| **Note** |
| --- |
| You must have administrative credentials to complete this procedure. |

**Related Links**

- For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003).
- For more information about what HyperTerminal does and does not support, see the HyperTerminal list of frequently asked questions on the Hilgraeve Web site at:

  http://go.microsoft.com/fwlink/?LinkId=225
- The Help documentation for the Windows Server 2003 family of products included on the CD and on the Web includes information about HyperTerminal. You can find the documentation on the Web at:

  http://go.microsoft.com/fwlink/?linkid=29890

# Internet Explorer 6.0 and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

**Internet Explorer 6.0**

This section provides information about:

- The benefits of Microsoft Internet Explorer 6.0 in products in the Windows Server 2003 family.
- A description of Internet Explorer Enhanced Security Configuration, which is enabled by default when you install a product in the Windows Server 2003 family.
- Examples of security-related configuration features offered in Internet Explorer 6.0 in products in the Windows Server 2003 family (as compared to Internet Explorer 5).
- Procedures for working with security-related settings in Internet Explorer.
- Resources for learning about topics related to security in Internet Explorer 6.0.

| **Note** |
| --- |
| This section of the white paper describes Internet Explorer 6.0 in general, but it does not describe Outlook Express (the e-mail component in Internet Explorer 6.0), the New Connection Wizard, or the error reporting tool in Internet Explorer. For information about these components, see the respective sections of this white paper (the error reporting tool in Internet Explorer is described in Windows Error Reporting and Internet Communication (Windows Server 2003) in this white paper).<br><br>Also note that the New Connection Wizard replaces the Network Connection Wizard and the Internet Connection Wizard in Windows 2000. |

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization where you use browsers on servers to connect to Web sites, run software from the Internet, download items from the Internet, and perform similar actions. This section, however, provides overview information as well as suggestions for other sources of information.

**Benefits and Purposes of Internet Explorer 6.0**

Internet Explorer 6.0 is designed to make it easy to browse and interact with sites on an intranet or on the Internet. It differs from many of the other components described in this white paper in that its main function is to communicate with sites on the Internet or an intranet (which contrasts with components that communicate with the Internet in the process of supporting some other activity).

Internet Explorer 6.0 is also designed to be highly configurable, with security and privacy settings that can protect your organization's networked assets while at the same time providing access to useful information and tools. In addition, Internet Explorer Enhanced Security Configuration, which is enabled by default when you install a product in the Windows Server 2003 family, helps make your computer more secure by limiting its exposure to malicious Web sites.

With this enhanced level of security, however, you might find that some Web sites are not displayed correctly in Internet Explorer when you are browsing from a server. Also, you might be prompted to enter your credentials when accessing network resources, such as files in shared folders with Universal Naming Convention (UNC) names. You can easily change the enhanced security settings.

If you want to establish a specific configuration on servers (instead of using Internet Explorer Enhanced Security Configuration), Internet Explorer 6.0 offers more security-related configuration options and settings than were available in Internet Explorer 5. The subsections that follow provide more information about Internet Explorer Enhanced Security Configuration and about the security-related configuration options and settings in Internet Explorer 6.0.

**Internet Explorer Enhanced Security Configuration**

Internet Explorer Enhanced Security Configuration is enabled by default when you install a product in the Windows Server 2003 family. With this configuration, each zone uses a higher security setting than was used by default in Windows 2000. You can easily change the enhanced security settings.

The following table outlines some of the differences that Internet Explorer Enhanced Security Configuration makes in security settings on a server. (For a description of zones, see "Examples of Security-Related Features Offered in Internet Explorer 6.0," later in this section.)

**Security settings with and without Internet Explorer Enhanced Security Configuration**

| Zone | With Internet Explorer Enhanced Security Configuration | Without Internet Explorer Enhanced Security Configuration (the same security levels as Windows 2000) |
|---|---|---|
| **Internet zone** | High security settings | Medium security settings |
| **Trusted sites zone** | Medium security settings | Low security settings |
| **Local intranet zone** | Medium-low security settings (intranet sites are not automatically detected) | Medium-low security settings (intranet sites are automatically detected) |

Also, with Internet Explorer Enhanced Security Configuration, several sites are added automatically to specific zones:

- The Windows Update Web site is added to the Trusted sites zone. This allows you to continue to get important updates for your operating system. For more information about Windows Update, see the Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003) section of this white paper.

- The Windows Error Reporting site is added to the Trusted sites zone. This allows you to report problems you encounter with your operating system and search for fixes. For more information about Windows Error Reporting, see Windows Error Reporting and Internet Communication (Windows Server 2003) in this white paper.

- Several local computer sites (for example, http://localhost, https://localhost, hcp://system) are added to the Local intranet zone. This allows applications and code to work locally so that you can complete common administrative tasks.

You can enable or disable the Internet Explorer Enhanced Security Configuration for administrators, all other user groups, or both. For more information, see "To remove Internet Explorer Enhanced Security Configuration and restore the default Internet Explorer 6.0 security settings," later in this section.

For more information about Internet Explorer Enhanced Security Configuration, see the resources listed in "Learning about Internet Explorer Enhanced Security Configuration," later in this section.

**Examples of Security-Related Features Offered in Internet Explorer 6.0**

This subsection describes enhancements in some of the security-related features in Internet Explorer 6.0, as compared to Internet Explorer 5. These features include:

- A Privacy tab that provides greater flexibility in specifying whether cookies will be blocked from specific sites or types of sites. An example of a type of site that could be blocked is one that does not have a compact policy, that is, a condensed computer-readable privacy statement. (The Privacy tab was not available in Internet Explorer 5.)

- Security settings that specify how Internet Explorer 6.0 handles such higher-risk items as ActiveX controls, downloads, and scripts. You can accept the settings in Internet Explorer Enhanced Security Configuration, you can customize these settings as needed, or you can set them to the predefined levels of high, medium, medium-low, or low. You can specify different settings for a number of zones, the most basic being the four preconfigured zones:

    - Local intranet zone: Normally contains only addresses inside your proxy server. (Note that when Internet Explorer Enhanced Security Configuration is enabled, intranet sites are not automatically detected.)

    - Trusted sites: Includes sites you designate as "trusted."

- Restricted sites: Includes sites you designate as "restricted."
- Internet zone: Includes everything that is not in another zone and is not on the local computer.

You can also specify different settings for customized zones that you add programmatically using the URL security zones application programming interface (API). For more information, search for "URL security zones" on the MSDN Web site at:

http://msdn.microsoft.com/

- Support for content-restricted IFrames (inline floating frames). This type of support enables developers to implement these frames in a way that makes it more difficult for malicious authors to start e-mail or content-based attacks.
- Improvements that increase the overall security and reliability of Internet Explorer 6.0.

For more information about features available in Internet Explorer, see "Resources for Learning About Topics Related to Security in Internet Explorer 6.0," later in this section, as well as the Internet Explorer page on the Microsoft Web site at:

http://www.microsoft.com/windows/ie/

**Procedures for Working with Security-Related Settings in Internet Explorer**
This subsection describes how to carry out the following:

- View security settings for zones in Internet Explorer
- Locate Group Policy objects (GPOs) that affect Internet Explorer, and view related Help
- Determine whether Internet Explorer Enhanced Security Configuration is enabled on a specific server
- Remove Internet Explorer Enhanced Security Configuration and restore the default Internet Explorer 6.0 security settings

**To view security settings for zones in Internet Explorer**

1. On the server on which you want to view settings, start Internet Explorer by your preferred method, for example, by clicking the Internet Explorer icon on the taskbar.
2. On the **Tools** menu, click **Internet Options**.
3. Click the **Security** tab.
4. Select the zone for which you want to view security settings:
    - Internet
    - Local intranet
    - Trusted sites
    - Restricted sites

**To locate Group Policy objects (GPOs) that affect Internet Explorer**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Internet Explorer**.

3. View the available settings.

4. Click **User Configuration**, click **Windows Settings**, and then click **Internet Explorer Maintenance**.

5. View the available settings.

6. Click **User Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Internet Explorer**.

7. View the available settings.

**To determine whether Internet Explorer Enhanced Security Configuration is enabled on a specific server**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Add or Remove Programs**.

3. Click **Add/Remove Windows Components** (on the left).

4. Scroll down to **Internet Explorer Enhanced Security Configuration**. If the check box is selected, it is enabled. If the check box is cleared, it is not enabled.

5. If you want to see whether Internet Explorer Enhanced Security Configuration is enabled for administrator groups, all other user groups, or both, select **Internet Explorer Enhanced Security Configuration**, and then click **Details**.

**To remove Internet Explorer Enhanced Security Configuration and restore the default Internet Explorer 6.0 security settings**

1. Click Start, and then either click Control Panel, or point to Settings and then click Control Panel.

2. Double-click Add or Remove Programs.

3. Click Add/Remove Windows Components (on the left).

4. Click Internet Explorer Enhanced Security Configuration, and then do one of the following:

   - To remove Internet Explorer Enhanced Security Configuration for both administrators and all other users, clear the Internet Explorer Enhanced Security Configuration check box, and then click Next.

   - To remove Internet Explorer Enhanced Security Configuration for administrators only or for users who are not in an administrator group, click Details, clear either the For administrator groups check box or the For all other user groups check box, and then click Next.

5. Follow the instructions to complete the Windows Components Wizard.

**Resources for Learning About Topics Related to Security in Internet Explorer 6.0**

This subsection lists resources that can help you learn about the following topics related to security in Internet Explorer 6.0:

- Internet Explorer Enhanced Security Configuration

- Security and privacy settings available in Internet Explorer 6.0

- Methods for mitigating the risks inherent in Web-based programs and scripts

- Ways to use Group Policy objects that control configuration settings for Internet Explorer 6.0
- The Internet Explorer Administration Kit

In addition, for information about unattended installation, see the resources listed in Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003).

| **Note** |
| --- |
| For information about Internet Explorer on clients running Windows XP Professional with Service Pack 2, that is, for information similar to what is provided in this white paper but focused on clients instead of servers, see "Using Windows XP Professional with Service Pack 2 in a Managed Environment: Controlling Communication with the Internet" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=29133. |

**Learning about Internet Explorer Enhanced Security Configuration**

For more information about Internet Explorer Enhanced Security Configuration, see one of the following:

- The informational pages displayed in Internet Explorer after you install a product in the Windows Server 2003 family. To view these pages, start Internet Explorer after completing the installation.
- Help topics in Internet Explorer. To view these topics, start Internet Explorer, click Help, and then click **Enhanced Security Configuration**.
- Help topics in Help and Support Center. To view these topics, start Internet Explorer, click **Start**, click **Help and Support**, and search for "enhanced security configuration."

**Learning about security and privacy settings in Internet Explorer 6.0**

An important source of detailed information about security and privacy settings in Internet Explorer 6.0 is the *Microsoft Internet Explorer 6 Resource Kit*. To learn about this and other Resource Kits, see the Microsoft TechNet Web site at:

http://go.microsoft.com/fwlink/?linkid=29894

The *Microsoft Internet Explorer 6 Resource Kit* consists of a number of parts that include these titles:

- "Privacy and Security Features"
- "Preparation for Deployment"
- "Customization and Installation"
- "Maintenance and Support," including information about keeping programs updated
- Appendices, including an appendix titled "Setting System Policies and Restrictions"

You can also use the following sources for information about security and privacy settings in Internet Explorer 6.0:

- Help for Internet Explorer (with Internet Explorer open, click the Help menu and select an appropriate option).
- The Internet Explorer page on the Microsoft Web site at:

  http://www.microsoft.com/windows/ie/

**Learning about mitigating the risks inherent in Web-based programs and scripts**

In a network-based and Internet-based environment, programs can take a variety of forms including scripts within documents, scripts within e-mail, or programs or other code objects running within Web pages. These programs can move across the Internet and are sometimes referred to as "mobile code." Configuration settings provide ways for you to control the way Internet Explorer 6.0 responds when someone tries to run a particular code object on a server running a product in the

Windows Server 2003 family. Two examples of the ways you can customize the Internet Explorer configuration are as follows:

- You can control the code (ActiveX controls, scripts, and so on) that administrators or operators can run. You can do this by customizing Authenticode® settings, which can, for example, prevent administrators or operators from running any unsigned code or enable them to only run code signed by specific authors.

- If you want to permit the use of ActiveX controls, but you do not want administrators or operators to download code directly from the Internet, you can specify that when Internet Explorer 6.0 looks for a requested executable, it goes to your own internal Web site instead of the Internet. For more information, see the white paper titled "Managing Mobile Code with Microsoft Technologies," at the end of this list, and search for "CodeBaseSearchPath."

You can use the following sources to learn more about mitigating the risks inherent in Web-based programs and scripts:

- To understand more about how a particular Microsoft programming or scripting language works, see the Microsoft Developer Network Web site at:

  http://msdn.microsoft.com/

- To learn about approaches to mitigating the risks presented by mobile code, see "Managing Mobile Code with Microsoft Technologies," a white paper on the Technet Web site at:

  http://go.microsoft.com/fwlink/?linkid=29170

**Learning about Group Policy objects that control configuration settings for Internet Explorer 6.0**
You can control configuration settings for Internet Explorer 6.0 by using Group Policy objects (GPOs). (You can also control the configuration of Internet Explorer by using the Internet Explorer Administration Kit; for more information, see "Learning about the Internet Explorer Administration Kit," later in this section.) For sources of information about Group Policy, see Appendix B: Resources for Learning About Group Policy (Windows Server 2003).

**Learning about the Internet Explorer Administration Kit**
With the deployment technologies available in the Internet Explorer Administration Kit (IEAK), you can efficiently deploy Internet Explorer and control the configuration of Internet Explorer across your organization. (You can also control the configuration of Internet Explorer by using Group Policy; for more information, see "Learning about Group Policy objects that control configuration settings for Internet Explorer 6.0," earlier in this section.)

A few of the features and resources in the IEAK include:

- Internet Explorer Customization Wizard. Step-by-step screens guide you through the process of creating customized browser packages that can be installed on client desktops.

- IEAK Profile Manager. After you deploy Internet Explorer, you can use the IEAK Profile Manager to change browser settings and restrictions automatically.

- IEAK Toolkit. The IEAK Toolkit contains a variety of helpful tools, programs, and sample files.

- IEAK Help. IEAK Help includes many conceptual and procedural topics that you can view by using the Index, Contents, and Search tabs. You can also print topics from IEAK Help.

For more information about the IEAK, see the IEAK Web site at:

http://go.microsoft.com/fwlink/?linkid=29479

# Internet Information Services and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Internet Information Services (IIS) in products in the Microsoft Windows Server 2003 family
- For servers from which you want to offer content on an intranet or the Internet, descriptions of some of the security-related features offered in IIS 6.0, and suggestions for other sources of information about security and IIS 6.0

| **Note** |
| --- |
| For servers from which you do not want to offer content on an intranet or the Internet, you do not need to remove IIS, since by default it is not installed with most products in the Windows Server 2003 family. The exception is Windows Server 2003, Web Edition, on which IIS is installed by default.<br><br>If you use a server as a Web server and then deploy it for some other purpose, remove IIS from that server. |

- Controlling Internet printing
- Subcomponents that are part of IIS, with instructions for finding out which subcomponents are installed on a given server
- Viewing Help for IIS
- Other sources of information about IIS

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running servers that communicate across the Internet. This section, however, provides overview information as well as suggestions for other sources of information about balancing your organization's requirements for communication across the Internet with your organization's requirements for protection of networked assets.

**Benefits and Purposes of IIS**

IIS 6.0 is one of the optional components in products in the Windows Server 2003 family. IIS is a component that provides an easy way to publish information on the Internet or an intranet. In a managed environment, IIS is usually installed on selected servers only. IIS includes innovative security features and a broad range of administrative features for managing Web sites. By using programmatic features like ASP.NET, which is an enhancement to Active Server Pages (ASP), you can more easily create and deploy scalable, flexible Web applications.

IIS is not installed by default with products in the Windows Server 2003 family other than Windows Server 2003, Web Edition. IIS and related components can be added by using either **Add or Remove Programs** in Control Panel or Manage Your Server. After IIS 6.0 is installed, it is configured by default in a "locked down" state. The locked down state means that IIS 6.0 accepts requests for static files only, until it is configured to serve dynamic content. It also means that all time-outs and settings are set to restrictive defaults. You can enable or disable IIS 6.0 functionality based on the needs of your organization by using IIS Manager. You can also enable IIS 6.0 functionality through programmatic and command-line interfaces.

For more information about IIS features, see the following Web sites:

- The product information page for IIS 6.0 at:

  http://go.microsoft.com/fwlink/?LinkId=46005

- The technical overview document for IIS 6.0 at:

http://go.microsoft.com/fwlink/?linkid=22805

- The IIS page on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=31134

If you have a Web site on which you want to use Microsoft .NET Passport for authentication and you also want to use Passport Manager Administration, a component in the Windows Server 2003 family, see Appendix E: Passport Manager Administration (Windows Server 2003).

**Examples of Security-Related Improvements in IIS 6.0**

IIS 6.0 includes a variety of settings and features related to security, some of which are listed in the following table.

**Examples of security-related settings and features in IIS 6.0**

| Setting or feature | Description |
|---|---|
| Disabling through Group Policy | With the Windows Server 2003 family, domain administrators can prevent users from installing IIS 6.0 on their computers. |
| Running as an account with limited privileges | IIS 6.0 worker processes run in a user context with limited privileges by default. This drastically reduces the attack surface of the Web server. |
| Secure ASP | All functions built into ASP pages always run as an account with limited privileges (anonymous user). |
| Recognized file extensions | IIS 6.0 serves requests only to files that have recognized file extensions and rejects requests to file extensions it doesn't recognize. |
| Command-line tools not accessible to Web users | Attackers often take advantage of command-line tools that are executable through the Web server. In IIS 6.0, the command-line tools cannot be executed by the default Web server identity. |
| Write protection for content | Once attackers get access to a server, they try to deface Web sites. By preventing anonymous Web users from overwriting Web content, these attacks can be mitigated. |
| Time-outs and limits | Product settings are set to aggressive and secure defaults. |
| Upload data limitations | Administrators can limit the size of data that can be uploaded to a server. |
| Buffer overflow protection | The Windows Administration Service in IIS will detect if a worker process had a buffer overflow and will exit that process. |
| File verification | The core server verifies that the requested content exists before it gives the request to a request handler (Internet Server Application Programming Interface [ISAPI] extension). |

For more information about creating Web sites with IIS 6.0 and maintaining appropriate levels of awareness and control over the communication to and from those sites, see the IIS Help. For information about viewing the Help, see "To view Help after installing IIS," later in this section.

**Controlling Internet Printing**

Internet printing makes it possible for clients to use printers located anywhere in the world by sending print jobs using Hypertext Transfer Protocol (HTTP). Additionally, a computer running a product in the Windows Server 2003 family can use IIS to create a Web page that provides information about printers and provides the transport for printing over the Internet.

For Internet printing, it is important to consider both the server and the client:

- **Server**: Internet printing is an optional component (not installed by default) of IIS 6.0. A server running a product in the Windows Server 2003 family can be configured to act as a print server allowing Internet printing. In a managed environment, you might want to ensure that the Internet printing subcomponent of IIS is not installed. For information about how to do this, see "Procedures for Checking or Controlling the Installation of IIS Subcomponents," later in this section.

- **Client**: Clients running Windows XP can install an Internet printer using a Web browser, the Add Printer Wizard, or the

Run dialog box. To control whether clients can support Internet printing, see the section about Internet printing in the white paper titled "Using Windows XP Professional with Service Pack 2 in a Managed Environment: Controlling Communication with the Internet." You can view this white paper on the TechNet Web site at:

http://go.microsoft.com/fwlink/?LinkId=29133

**Note**

The white paper available from the preceding link also includes a procedure for using Group Policy to disable Internet printing on a computer running IIS. That procedure is not applicable for computers running products in the Windows Server 2003 family because the Group Policy object referred to in the procedure is not available in the Windows Server 2003 family. For such servers, use the procedures in "Procedures for Checking or Controlling the Installation of IIS Subcomponents," later in this section.

**Answer File Entries and Registry Keys for IIS Subcomponents**

For reference purposes, the following table shows the syntax for answer file entries associated with IIS in the Windows Server 2003 family as well as the corresponding registry keys. Do not change the registry keys. They are shown for use in a script that could check whether a particular component is installed on a particular server. A registry key value of 0x00000000 means the component is not installed, and a value of 0x00000001 means the component is installed.

**Note**

For more details about answer file entries related to IIS components, follow the steps in "To view Help after installing IIS," later in this section, and then search for the topic called "Installing IIS." In that topic, look for a table showing the answer file entries.

**Answer file entries and registry keys associated with IIS subcomponents for the Windows Server 2003 family**

| IIS subcomponent | Syntax for answer file entry (in the [Components] section) | Registry key (for use in a script that checks whether a component is installed): 0x00000000 means it is not installed; 0x00000001 means it is installed |
|---|---|---|
| IIS common files | iis_common = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\ iis_common |
| Active Server Pages (ASP) for IIS | iis_asp = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\iis_asp |
| File Transfer Protocol (FTP) service | iis_ftp = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\iis_ftp |
| IIS Manager (Microsoft Management Console [MMC] snap-in) | iis_inetmgr = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\iis_inetmgr |

| | | |
|---|---|---|
| Internet Data Connector | iis_internetdataconnector = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\ iis_internetdataconnector |
| Network News Transfer Protocol (NNTP) service | iis_nntp = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\iis_nntp |
| Server-Side Includes | iis_serversideincludes = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\ iis_serversideincludes |
| Simple Mail Transfer Protocol (SMTP) service | iis_smtp = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\iis_smtp |
| Web Distributed Authoring and Versioning (WebDAV) publishing | iis_webdav = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\iis_webdav |
| World Wide Web (WWW) service | iis_www = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\iis_www |
| Remote administration (HTML) | sakit_web = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\sakit_web |
| Internet Server Application Programming Interface (ISAPI) for Background Intelligent Transfer Service (BITS) server extensions | BitsServerExtensionsISAPI = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\ bitsserverextensionsisapi |
| Background Intelligent Transfer Service (BITS) server extensions snap-in | BitsServerExtensionsManager = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\ bitsserverextensionsmanager |
| FrontPage server extensions | fp_extensions = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\Setup\ OC Manager\Subcomponents\fp_extensions |

| | | |
|---|---|---|
| Internet printing | inetprint = On \| Off | HKEY_LOCAL_MACHINE\Software\<br>Microsoft\Windows\CurrentVersion\Setup\<br>OC Manager\Subcomponents\inetprint |
| ActiveX control and sample pages for hosting Terminal Services client connections over the Web | TSWebClient = On \| Off | HKEY_LOCAL_MACHINE\Software\<br>Microsoft\Windows\CurrentVersion\Setup\<br>OC Manager\Subcomponents\TSWebClient |

**Note**

For several of the subcomponents in the previous table, the software for the subcomponent is installed regardless of the answer-file entry, but the subcomponent cannot be used unless the answer-file entry is set to On (or the procedure is followed for installing the subcomponent through **Add or Remove Programs** in Control Panel). These subcomponents are Internet Data Connector, Server-Side Includes, and WebDAV publishing.

**Procedures for Checking or Controlling the Installation of IIS Subcomponents**

The following procedures explain how to:

- View the registry keys listed in the table in the previous subsection
- View or change the IIS components currently installed on a computer running a product in the Windows Server 2003 family
- Specify answer file entries that control whether IIS subcomponents are included during unattended installation

**To view registry keys related to IIS subcomponents**

1. Open Registry Editor by clicking **Start**, clicking **Run**, and then typing **regedit**.

   **Caution**

   Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. Navigate to HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\.
3. View the registry keys listed in the table in the previous subsection, and find the value associated with each key. A value of 0x00000000 means the component is not installed. A value of 0x00000001 means the component is installed.
4. Close **Registry Editor**.

**To view or change the IIS components currently installed on a computer running a product in the Windows Server 2003 family**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components** (on the left).
4. Select **Application Server** and then click **Details**.

5. Find **Internet Information Services (IIS)** in the list, and perform one of the following steps:

   - If IIS is installed and you want to remove it, clear the check box for IIS and complete the wizard.
   - If IIS is not installed and you want to add the default set of IIS subcomponents, select the check box for IIS and complete the wizard.
   - If you want to view or select from the list of IIS subcomponents, after selecting IIS, click **Details**.

   | **Note** |
   | --- |
   | The Internet Printing component is in the list of components that appears when you click **Details**. |

6. Follow the instructions to complete the Windows Components Wizard.

**To specify answer file entries that control whether IIS subcomponents are included during unattended installation**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file.
2. In the [Components] section of the answer file, add the appropriate entries listed in the table in "Answer File Entries and Registry Keys for IIS Subcomponents," earlier in this section. Ensure that the entries specify **Off** for components you do not want to install and **On** for components you want to install.

   If no IIS subcomponents are listed in an answer file for unattended installation of a product in the Windows Server 2003 family, these subcomponents are *not* installed by default.

   | **Note** |
   | --- |
   | For more details about answer file entries related to IIS components, follow the steps in the next procedure, "To view Help after installing IIS," and then search for the Help topic called "Installing IIS." In that topic, look for a table showing the answer file entries. |

**To view Help after installing IIS**

1. After installing IIS (including the IIS Manager subcomponent, which is included in default installations of IIS), click **Start**.
2. Either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
3. Double-click **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
4. Click the **Help** menu and then click **Help Topics**.

**Related Links**

For more information about IIS, see the following Web sites:

- The product information page for IIS 6.0 at:

  http://go.microsoft.com/fwlink/?LinkId=46005

- The technical overview document for IIS 6.0 at:

  http://go.microsoft.com/fwlink/?linkid=22805

- The IIS page on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?linkid=31134

- The page for the IIS Lockdown Tool on the TechNet Web site at:

  http://go.microsoft.com/fwlink/?linkid=29896

# Internet Protocol Version 6 (IPv6) and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

**Internet Protocol Version 6 (IPv6)**
This section provides information about:

- An introduction to the IPv6 protocol
- The benefits of the IPv6 protocol
- How the IPv6 protocol communicates with sites on the Internet
- How to control the IPv6 protocol to limit the flow of information to and from the Internet
- How to monitor and troubleshoot the IPv6 protocol after configuration is complete

**An Introduction to the IPv6 Protocol**
The current version of the Internet Protocol (known as IP version 4 or IPv4) has not been substantially changed since 1981, when the Internet Engineering Task Force (IETF) published the definitive specification of IP (RFC 791). IPv4 has proven to be robust, easily implemented, and interoperable. It has stood the test of scaling an internetwork to a global utility the size of today's Internet, which is a tribute to its initial design.

The initial design, however, did not anticipate the exponential growth of the Internet and the exhaustion of the IPv4 address space, or the effort required to maintain routing information. Because of the way in which IPv4 network IDs are allocated, there are routinely over 70,000 routes in the routing tables of Internet backbone routers. Most current IPv4 implementations are configured either manually or through a stateful address configuration protocol such as the Dynamic Host Configuration Protocol (DHCP). With more computers and devices using IP, there is a need for a simpler and more automatic configuration of addresses and other configuration settings that do not rely on the administration of a DHCP infrastructure.

Another factor driving the development of IPv6 is the need for improved encryption. Private communication over a public medium like the Internet requires encryption services that protect the data sent from being viewed or modified in transit. There is a standard for providing security for IPv4 packets (known as Internet Protocol security or IPSec). In IPv4, however, this standard is optional and proprietary solutions are prevalent.

While standards for quality of service (QoS) exist for IPv4, real-time traffic support relies on the IPv4 Type of Service (TOS) field and the identification of the payload, typically using a User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) port. Unfortunately, the IPv4 TOS field has limited functionality and has different interpretations. In addition, payload identification using a TCP or UDP port is not possible when the IPv4 packet payload is encrypted.

To address these concerns, the IETF has developed a suite of protocols and standards known as IP version 6 (IPv6). This new version, previously named IP-The Next Generation (IPng), incorporates the concepts of many proposed methods for updating the IPv4 protocol. IPv6 is intentionally designed for minimal impact on upper and lower layer protocols by avoiding the arbitrary addition of new features.

For the latest set of RFCs and Internet drafts describing IPv6 and IPv4 coexistence and migration technologies, see the Internet Engineering Task Force (IETF) Web site at:

http://go.microsoft.com/fwlink/?LinkId=29136

(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

**Benefits and Purposes of the IPv6 Protocol**
The IPv6 header has a new format that is designed to minimize header validation and processing. An IPv6 address is four times larger than an IPv4 address. The global addresses used on the IPv6 portion of the Internet are designed to create an efficient, hierarchical, and summarized routing infrastructure that addresses the common occurrence of multiple levels of Internet service providers. On the IPv6 Internet, the backbone routers have an efficient and hierarchical addressing and routing infrastructure that uses smaller routing tables.

IPv6 supports both stateful address configuration, such as address configuration in the presence of a DHCP server, and stateless address configuration, or address configuration in the absence of a DHCP server. The support for IPSec is an IPv6 protocol suite requirement. This requirement provides a standards-based solution for network security needs and promotes interoperability between different IPv6 implementations. The new fields in the IPv6 header define how traffic is handled and identified.

Traffic identification, by using a Flow Label field in the IPv6 header, allows routers to identify and provide special handling for packets that belong to a flow. (A flow is a series of packets between a source and destination.) Because the traffic is identified in the IPv6 header, support for quality of service (QoS) can be easily achieved even when the packet payload is encrypted with IPSec.

The new Neighbor Discovery protocol for neighboring node interaction in IPv6 is a series of messages from the Internet Control Message Protocol for IPv6 (ICMPv6) that manage the interaction of neighboring nodes. Neighbor Discovery replaces Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and ICMPv4 Redirect messages with efficient multicast and unicast messages.

IPv6 can be extended for new features by adding extension headers after the IPv6 header. Unlike the IPv4 header, which can only support 40 bytes of options, the size of IPv6 extension headers is only constrained by the size of the IPv6 packet.

The following table compares the key features of the IPv4 and IPv6 protocols.

**Comparison of features in IPv4 and IPv6**

| **IPv4** | **IPv6** |
|---|---|
| Source and destination addresses are 32 bits (4 bytes) in length. | Source and destination addresses are 128 bits (16 bytes) in length. |
| IPsec support is optional. | IPsec support is required. |
| No identification of packet flow for QoS handling by routers is present within the IPv4 header. | Packet flow identification for QoS handling by routers is included in the IPv6 header using the Flow Label field. |
| Fragmentation is done by both routers and the sending host. | Fragmentation is not done by routers, only by the sending host. |
| Header includes a checksum. | Header does not include a checksum. |
| Header includes options. | All optional data is moved to IPv6 extension headers. |
| The Address Resolution Protocol (ARP) uses broadcast ARP Request frames to resolve an IPv4 address to a link layer address. | ARP Request frames are replaced with multicast Neighbor Solicitation messages. |
| The Internet Group Management Protocol (IGMP) is used to manage local subnet group membership. | IGMP is replaced with Multicast Listener Discovery (MLD) messages. |
| ICMP Router Discovery is used to determine the IPv4 address of the best default gateway and is optional. | ICMP Router Discovery is replaced with ICMPv6 Router Solicitation and Router Advertisement messages and is required. |
| Broadcast addresses are used to send traffic to all nodes on a subnet. | There are no IPv6 broadcast addresses. Instead, a link-local scope all-nodes multicast address is used. |
| Must be configured either manually or through DHCP. | Does not require manual configuration or DHCP. |
| Uses host address (A) resource records in the Domain Name System (DNS) to map host names to IPv4 addresses. | Uses host address (AAAA) resource records in the Domain Name System (DNS) to map host names to IPv6 addresses. |
| Uses pointer (PTR) resource records in the IN-ADDR.ARPA DNS domain to map IPv4 addresses to host names. | Uses pointer (PTR) resource records in the IP6.ARPA DNS domain to map IPv6 addresses to host names. |
| Links must support a 576-byte packet size (possibly fragmented). | Links must support a 1280-byte packet size (without fragmentation). |

For more information about IP version 6, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=29519

**Overview: Using the IPv6 Protocol in a Non-Native IPv6 Environment**
On networks that do not have native support for IPv6 traffic, the IPv6 traffic is transmitted on the network by encapsulating the IPv6 packets within IPv4 packet headers. One such method of transmission is referred to as 6to4 tunneling.

For more information about the 6to4 tunneling technique, see "Connection of IPv6 Domains via IPv4 Clouds," in RFC 3056 on the Internet Engineering Task Force (IETF) Web site at:

http://go.microsoft.com/fwlink/?LinkId=29898

(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

**How the IPv6 Protocol Communicates with Sites on the Internet**
Although there are differences between the two protocol versions IPv4 and IPv6, their differences do not prevent them from coexisting or communicating on the IPv4 network.

If native IPv6 connectivity does not exist, a computer makes a Domain Name System (DNS) query for network relay routers that provide IPv6 service as part of the startup process. By default, this DNS query is presently set to "6to4.ipv6.microsoft.com" and the response contains a well-known IPv4 anycast address. (An anycast address is one that identifies multiple nodes and interfaces.) As more IPv6 relay routers are added in the future, this address will be assigned to more computers that are owned by various Internet service providers (ISPs).

If the DNS query provides multiple addresses, the host selects an appropriate relay router by sending an IPv6 packet to each one and choosing the one that responds first.

| **Note** |
| --- |
| 6to4 tunneling is enabled when IPv6 services are not native to your network and there is a public IPv4 Internet address present on the network access point. |

The use of IPv6 in the Microsoft Windows Server 2003 family is currently supported only when IPv4 is also installed.

**Controlling the IPv6 Protocol to Limit the Flow of Information to and from the Internet**
You can stop the ingress or egress of IPv6 traffic on your network by configuring your network firewall to block all IPv6-specific packets. When the 6to4 tunneling technique is used, you can configure your firewall to block all IPv4 packets that include the IPv6 protocol designation of 41 in the protocol field of the IPv4 packet header.

The default settings for a member of the computer users group do not permit those users to install networking protocols. You should limit who is allowed to install the IPv6 stack on network computers by carefully limiting the number of users that have administrative logon credentials.

You can use the Active Directory directory service and Group Policy to filter and control the user's ability to add new networking protocols, or to modify existing networking configurations. For more information about these configuration methods, see Appendix B: Resources for Learning About Group Policy (Windows Server 2003). For information about installing and uninstalling IPv6, see the list of procedures in the next subsection.

**Procedures for Configuration of the IPv6 Protocol**
Installing and uninstalling the IPv6 protocol stack can be done by using the Network Connections folder or the command prompt.

The following two procedures describe installing and uninstalling the IPv6 protocol stack by using the Network Connections folder.

**To Find Information About Security in Relation to IPv6**

1. View updated product documentation for Windows Server 2003 products on the Web at:

   http://go.microsoft.com/fwlink/?linkid=29881

2. In the table of contents, navigate as follows:

   Network Services\Managing Core Network Services\IP Version 6\IPv6 Concepts\IPv6 Overview

3. View topics below “IPv6 Overview,” especially “Security information for IPv6.”

**To install IPv6 using the Network Connections folder**

1. Click **Start**.
2. Either point to **Control Panel** and then double-click **Network Connections**, or point to **Settings**, click **Control Panel**, and then double-click **Network Connections**.
3. Right-click any local area connection, and then click **Properties**.
4. Click **Install**.
5. In the **Select Network Component Type** dialog box, click **Protocol**, and then click **Add**.
6. In the **Select Network Protocol** dialog box, click **Microsoft TCP/IP version 6**.

**To uninstall IPv6 using the Network Connections folder**

1. Click **Start**.
2. Either point to **Control Panel** and then double-click **Network Connections**, or point to **Settings**, click **Control Panel**, and then double-click **Network Connections**.
3. Right-click any local area connection, and then click **Properties**.
4. Click **Microsoft TCP/IP version 6** in the list of installed components, and then click **Uninstall**.

The following two procedures describe installing and uninstalling the IPv6 protocol stack by using the command prompt.

**To install IPv6 on a computer using the command prompt**

1. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**.
2. Type **netsh interface ipv6 install**, and then press ENTER.

**To uninstall IPv6 from a computer using the command prompt**

1. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**.
2. Type **netsh interface ipv6 uninstall**, and then press ENTER.

| **Note** |
| --- |
| The IPv6 configuration options require that you have administrative credentials on the computer. |

**Monitoring and Troubleshooting the IPv6 Protocol**

The following procedures describe ways to view TCP/IP configurations.

**To display the complete list of TCP/IP interface configurations for a computer using the command prompt**

1. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**.
2. Type **ipconfig /all**, and then press ENTER.

**To display the TCP/IP routing table using the command prompt**

1. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. Type **route print**, and then press ENTER.

| **Note** |
| --- |
| For more information about TCP/IP configurations, see "TCP/IP utilities," in the Help and Support Center index. |

**Troubleshooting a command-line installation error**
The installation of the IPv6 protocol stack requires that you have administrative credentials. The command-line prompt will yield the "Access is denied" error (0x800700005) if you attempt to install the IPv6 protocol from the command-line prompt without having the required account credentials.

**Related Links**
**Online resources**

- For more information about 6to4, see "Connection of IPv6 Domains via IPv4 Clouds," in RFC 3056 on the IETF Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29898

- For more information about IP version 6, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29519

- For information about security in relation to IPv6, see “To Find Information About Security in Relation to IPv6,” earlier in this white paper.

- For more information about IPv6 addressing, see "Internet Protocol Version 6 (IPv6) Addressing Architecture" in RFC 3513 on the IETF Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29135

- For the latest set of RFCs and Internet drafts describing IPv6 standards, see “IP Version 6 Working Group (ipv6)” at the IETF Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29136

(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

**Printed references**
For more information about the IPv6 protocol suite, you can consult the following references:


- Davies, J. *Understanding IPv6*. Redmond, WA: Microsoft Press, 2002.

- Hagen, S. *IPv6 Essentials*. Sebastopol, CA: O'Reilly and Associates, Inc., 2002.

# NetMeeting and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of NetMeeting
- Using NetMeeting in a managed environment
- How NetMeeting communicates with sites on the Internet
- How to control NetMeeting to limit the flow of information to and from the Internet

**Benefits and Purposes of NetMeeting**

NetMeeting® conferencing software, which is included in the Microsoft Windows Server 2003 family, enables real-time communication and collaboration over the Internet or an intranet. From a computer running the Windows 95, Windows 98, Windows NT 4.0, Windows 2000, or Windows XP operating system, users can communicate over a network with real-time voice and video technology. Users can work together on virtually any Windows-based application, exchange or mark up graphics on an electronic whiteboard, transfer files, or use the text-based chat program.

NetMeeting helps small and large organizations take full advantage of their corporate intranet for real-time communication and collaboration. On the Internet, connecting to other NetMeeting users is made easy with Internet Locator Service (ILS), enabling participants to call each other from a dynamic directory within NetMeeting or from a Web page. Features include remote desktop sharing, virtual conferencing using Microsoft Outlook, security features, and the ability to embed the NetMeeting user interface in an organization's intranet Web pages.

| **Note** |
| --- |
| NetMeeting is not available on the 64-bit versions of the Windows Server 2003 family. |

To learn more about the NetMeeting features, see the article on the Microsoft TechNet Web site at:

http://go.microsoft.com/fwlink/?LinkId=29175

**Overview: Using NetMeeting in a Managed Environment**

NetMeeting supports communication standards for audio, video, and data conferencing. NetMeeting users can communicate and collaborate with users of other standards-based, compatible products. They can connect by modem, Integrated Services Digital Network (ISDN), or local area network (LAN) using Transmission Control Protocol/Internet Protocol (TCP/IP). In addition, support for Group Policy in NetMeeting makes it easy for administrators to centrally control and manage the NetMeeting work environment.

You can use Active Directory directory service and Group Policy to configure NetMeeting to help meet your security requirements. You can also control the configuration of NetMeeting by using the NetMeeting Resource Kit. For more information about the NetMeeting Resource Kit, see "Alternate Methods for Controlling NetMeeting," later in this section.

NetMeeting components and features require that several ports be open from the firewall. For more information, see "NetMeeting and firewalls," later in this section.

**How NetMeeting Communicates with Sites on the Internet**

NetMeeting provides an infrastructure for communication between network applications and services. In this infrastructure, NetMeeting is both an application and a platform for other applications or services. The components and services in NetMeeting provide real-time communication and collaboration over the Internet or an organization's intranet.

NetMeeting audio and video conferencing features are based on the H.323 infrastructure, which enables NetMeeting to interoperate with other H.323 standards-based products. (H.323 is a standard approved by the International Telecommunication Union [ITU] that defines how audiovisual conferencing data is transmitted across networks.) NetMeeting data conferencing features are based on the T.120 infrastructure, enabling NetMeeting to interoperate with other T.120

standards-based products. (The T.120 standard is a suite of communication and application protocols developed for real-time, multipoint data connections and conferencing.)

Detailed information about the H.323 and T.120 standards is beyond the scope of this white paper. Further information can be found on the following sites:

- For more information about the H.323 standard and NetMeeting, see Part 3, Chapter 11, "Understanding the H.323 Standard," in the *NetMeeting 3.0 Resource Kit* at:

  http://go.microsoft.com/fwlink/?LinkId=36134

- For more information about the H.323 specification, see the Web sites of the International Telecommunication Union (ITU) and the International Multimedia Telecommunications Consortium (IMTC) at:

  http://go.microsoft.com/fwlink/?LinkId=29510

  http://www.imtc.org/

- To learn more about the T.120 standard and NetMeeting, see Part 3, Chapter 10, "Understanding the T.120 Standard," in the *NetMeeting 3.0 Resource Kit* at:

  http://go.microsoft.com/fwlink/?LinkId=36134

- For more information about the T.120 architecture, see the International Multimedia Teleconferencing Consortium (IMTC) Web site at:

  http://www.imtc.org/

(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

**NetMeeting port assignments**
When you use NetMeeting to call other users over the Internet, several IP ports are required to establish the outbound connection. The following table describes the port numbers, their functions, and the resulting connection.

**Port assignments for NetMeeting**

| Port | Function | Outbound connection |
|---|---|---|
| 389 | Internet Locator Service (ILS) | TCP |
| 522 | User Location Service (ULS) | TCP |
| 1503 | T.120 | TCP |
| 1720 | H.323 call setup | TCP |
| 1731 | Audio call control | TCP |
| 1024 through 65535 (dynamic) | H.323 call control | TCP |
| 1024 through 65535 (dynamic) | H.323 streaming | Real-Time Transfer Protocol (RTP) over User Datagram Protocol (UDP) |

For more information about NetMeeting communication ports and firewall configuration topics, see Part 2, Chapter 4, "Firewall Configuration," in the *Microsoft NetMeeting 3 Resource Kit* at:

http://go.microsoft.com/fwlink/?LinkId=36134

**Controlling NetMeeting to Limit the Flow of Information to and from the Internet**
You can configure NetMeeting by using Group Policy objects (GPOs) on servers running products in the Windows Server 2003 family. (You can also control the configuration of NetMeeting by using the NetMeeting Resource Kit; for more information, see

"Alternate Methods for Controlling NetMeeting," later in this section.)

This subsection includes information about the following topics:

- NetMeeting and Group Policy
- NetMeeting security
- NetMeeting and firewalls
- Establishing a NetMeeting connection with a firewall
- Firewall limitations for NetMeeting

**NetMeeting and Group Policy**

Group Policy can be used to define the default NetMeeting configuration settings that will be automatically applied to users and computers. These settings determine which NetMeeting features and capabilities are available to a particular group of users. The Group Policy configuration settings that are specific to NetMeeting are grouped into two different categories. These category groupings enable you to independently manage NetMeeting configuration settings for computers and users within your organization. Through the use of Group Policy you can enable, disable, or set configuration options for NetMeeting features or capabilities.

For additional information about Group Policy, see Appendix B: Resources for Learning About Group Policy (Windows Server 2003).

You can use Group Policy to manage the following NetMeeting configuration options for users in your organization:

- NetMeeting Group Policy settings for computers
- NetMeeting Group Policy settings for users

**Configuring NetMeeting settings for computers through Group Policy**

You can use Group Policy to determine the NetMeeting features and capabilities that are available to all users of the computers that are affected by the application of the NetMeeting Group Policy settings.

For details about locating the Group Policy objects (GPOs) for NetMeeting, see "Procedures for Configuration of NetMeeting," later in this section. The NetMeeting Group Policy configuration setting that is specific to computers is as follows:

- **Disable remote Desktop Sharing**: You can use Group Policy to set remote desktop sharing choices in NetMeeting for all the users who are affected by the application of this Group Policy setting.

For more information about how to use Group Policy to manage the NetMeeting computer settings, see "To disable the NetMeeting remote desktop sharing feature through Group Policy," later in this section.

| **Note** |
| --- |
| Computer-related Group Policy settings are applied when the operating system starts and during the periodic refresh cycle. |

**Configuring NetMeeting settings through Group Policy**

You can use Group Policy to determine the NetMeeting features and capabilities that are available for a user or a group of users that are affected by the application of the NetMeeting Group Policy settings.

These Group Policy configuration options include the policy settings for NetMeeting, application sharing, audio and video, and the options page.

For more information about how to use Group Policy to manage the NetMeeting user settings, see "Procedures for Configuration of NetMeeting," later in this section.

The NetMeeting Group Policy configuration settings that are specific to users are as follows:

**Configuring NetMeeting settings for users through Group Policy**

For details about locating the Group Policy objects for NetMeeting, see "Procedures for Configuration of NetMeeting," later in this section. You can use Group Policy to set configuration settings for the following NetMeeting features:

- **Enable Automatic Configuration**: Configures NetMeeting to download settings for users each time it starts.
- **Disable Directory services**: Disables the directory feature—users will not log on to a directory server when NetMeeting starts. Users will not be able to view or make calls using the NetMeeting directory.
- **Prevent adding Directory servers**: Prevents the user from adding directory servers to the list of available directory servers they can use for placing calls.
- **Prevent viewing Web directory**: Prevents the user from viewing directories as Web pages in a browser.
- **Set the intranet support Web page**: Sets the Web address that NetMeeting will display when users choose the Online Support command from the NetMeeting Help menu.
- **Set Call Security options**: Sets the level of security for outgoing and incoming NetMeeting calls.
- **Prevent changing Call placement method**: Prevents the user from changing the way calls are placed, either directly or by means of a gatekeeper server.
- **Prevent automatic acceptance of Calls**: Prevents the user from turning on automatic acceptance of incoming calls.
- **Allow persisting automatic acceptance of Calls**: Sets automatic acceptance of incoming calls to be persistent.
- **Prevent sending files**: Prevents users from sending files to others in a conference.
- **Prevent receiving files**: Prevents users from receiving files from others in a conference.
- **Limit the size of sent files**: Sets the maximum file size that can be sent to others in a conference.
- **Disable Chat**: Disables the chat feature of NetMeeting.
- **Disable NetMeeting 2.x Whiteboard**: Disables the NetMeeting 2.x Whiteboard feature. (The 2.x feature provides compatibility with earlier versions of NetMeeting only.)
- **Disable Whiteboard**: Disables the whiteboard feature of NetMeeting.

**Configuring NetMeeting application sharing settings through Group Policy**

For details about locating the Group Policy objects (GPOs) for NetMeeting, see "Procedures for Configuration of NetMeeting," later in this section. You can use Group Policy to set configuration settings for the following elements of the NetMeeting Application Sharing feature:

- **Disable application Sharing**: Disables the NetMeeting application sharing feature completely. Users will not be able to host or view shared applications.
- **Prevent Sharing**: Prevents users from sharing anything themselves. They will still be able to view shared applications or desktops from others.
- **Prevent Desktop Sharing**: Prevents users from sharing their Windows desktop. They will still be able to share individual applications.
- **Prevent Sharing Command Prompts**: Prevents the user from sharing command prompts. Enabling this prevents the user from inadvertently sharing applications, since command prompts can be used to start other applications.
- **Prevent Sharing Explorer windows**: Prevents the user from sharing Windows Explorer windows. Enabling this prevents the user from inadvertently sharing applications, since Windows Explorer windows can be used to start other

applications.

- **Prevent Control**: Prevents users from allowing others in a conference to control what they have shared. Enabling this enforces a read-only mode whereby the other participants cannot change the data in the shared application.

- **Prevent Application Sharing in true color**: Prevents users from sharing applications in true color, which uses more bandwidth.

**Configuring NetMeeting audio and video settings through Group Policy**

For details about locating the Group Policy objects (GPOs) for NetMeeting, see "Procedures for Configuration of NetMeeting," later in this section. You can use Group Policy to set configuration settings for the following audio and video elements in NetMeeting:

- **Limit the bandwidth of Audio and Video**: Configures the maximum bandwidth, specified in kilobytes per second, to be used for audio and video.

- **Disable Audio**: Disables the audio feature of NetMeeting; users will not be able to send or receive audio.

- **Disable full duplex Audio**: Disables the full duplex audio mode. Users will not be able to listen to incoming audio while speaking into the microphone. Older audio hardware may not perform well when full duplex audio is enabled.

- **Prevent changing DirectSound Audio setting**: Prevents the user from changing the DirectSound audio setting. DirectSound has a better audio quality, although earlier audio hardware may not support it.

- **Prevent sending Video**: Prevents the user from sending video. Setting this option does not prevent the user from receiving video.

- **Prevent receiving Video**: Prevents the user from receiving video. Setting this option does not prevent the user from sending video.

**Configuring NetMeeting options settings through Group Policy**

For details about locating the Group Policy objects (GPOs) for NetMeeting, see "Procedures for Configuration of NetMeeting," later in this section. You can use Group Policy to set configuration settings for the following elements of the NetMeeting Options page:

- **Hide the General page**: Removes the General tab on the NetMeeting Options page.

- **Disable the Advanced Calling button**: Disables the Advanced Calling button from the General page.

- **Hide the Security page**: Removes the Security tab on the NetMeeting Options page.

- **Hide the Audio page**: Removes the Audio tab on the NetMeeting Options page.

- **Hide the Video page**: Removes the Video tab on the NetMeeting Options page.

| **Note** |
| --- |
| User-related Group Policy settings are applied when a user logs on to the computer and during the periodic refresh cycle. |

**NetMeeting security**

The NetMeeting security architecture for data conferencing takes advantage of the existing, standards-compliant security features of the Windows Server 2003 family and Microsoft Internet Explorer. The NetMeeting security architecture utilizes a 40-bit encryption technology and has the following security features:

- **Password protection**: This feature enables the user to create or participate in a meeting that requires a password to join. Password protection helps to ensure that only authorized users participate in a password-protected meeting. A password is also required to use the remote desktop sharing feature.

- **User authentication**: This feature provides a way to verify the identity of a caller or meeting participant using a certificate.

- **Data encryption**: This feature helps to protect data exchanged during a meeting so that it is not easily read by any unauthorized parties that may intercept the data. The 40-bit data encryption applies to the whiteboard and chat features, shared applications, and transferred files. Audio and video communications are not encrypted.

NetMeeting security features integrate with security in the Windows Server 2003 family and Internet Explorer in a variety of ways, including the following:

- NetMeeting uses the NetMeeting private certificate store to provide personal certificates for user authentication and data encryption.

- NetMeeting uses the Windows certificate store to maintain NetMeeting certificates.

- NetMeeting uses Security Support Provider Interface (SSPI) functions to generate and process security tokens.

These security features can be implemented by an administrator or a NetMeeting user. Using the NetMeeting Resource Kit Wizard or Group Policy in NetMeeting, the administrator can enforce security settings that apply to all users. If allowed by the administrator, NetMeeting users can also select their own security settings in the NetMeeting user interface (UI) and change security settings for individual calls.

You can use the following sources to learn more about NetMeeting configuration and security topics:

- For more information about the NetMeeting Resource Kit Wizard, see Part 2, Chapter 2, "Resource Kit Wizard" in the *NetMeeting 3.0 Resource Kit* at:

  http://go.microsoft.com/fwlink/?LinkId=36134

- For more information about the security features available in NetMeeting, see Part 2, Chapter 5, “NetMeeting Security” in the *NetMeeting 3.0 Resource Kit* at:

  http://go.microsoft.com/fwlink/?LinkId=36134

**NetMeeting and firewalls**

You can configure firewall components in a variety of ways, depending on your organization's specific security policies and overall operations. While most firewalls are capable of allowing primary (initial) and secondary (subsequent) Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) connections, it is possible that they are configured to support only specific connections based on security considerations. For example, some firewalls support only primary TCP connections, which some professionals view as the most reliable.

For NetMeeting multipoint data conferencing—program sharing, whiteboard, chat, file transfer, and directory access—your firewall only needs to pass through primary TCP connections on assigned ports. NetMeeting audio and video features require secondary TCP and UDP connections on dynamically assigned ports.

| **Note** |
| --- |
| NetMeeting audio and video features require secondary TCP and UDP connections. Therefore, when you establish connections through firewalls that accept only primary TCP connections, you are not able to use the audio or video features of NetMeeting. |

Detailed firewall configuration procedures for NetMeeting are beyond the scope of this white paper. For more information about NetMeeting firewall connections, see Part 2, Chapter 4, "Firewall Configuration" in the *NetMeeting 3.0 Resource Kit*, particularly the section titled, "Establishing a NetMeeting Connection with a Firewall" at:

http://go.microsoft.com/fwlink/?LinkId=36134

Microsoft NetMeeting can be configured to work with an organization's existing firewall security. Because of limitations in

most firewall technology, however, few products are available that enable you to securely transport inbound and outbound NetMeeting calls containing audio, video, and data across a firewall. You should consider carefully the relative security risks of enabling different parts of a NetMeeting call in your firewall product. You must especially consider the security risks involved when modifying your firewall configuration to enable any component of an inbound NetMeeting call.

Some organizations have security or policy concerns that require them to limit how fully they support NetMeeting in their firewall configuration. These concerns are based on network capacity planning or weaknesses in the firewall technology being used. For example, security concerns might prohibit an organization from accepting any inbound or outbound flow of UDP data through the firewall. Because these UDP connections are required for NetMeeting audio and video features, disabling this function excludes audio and video features in NetMeeting for calls through the firewall. The organization can still use NetMeeting data conferencing features such as program sharing, file transfer, whiteboard, and chat for calls through the firewall by allowing only TCP connections on ports 522 and 1503.

For more information about NetMeeting firewall security, see the section titled "Security and Policy Concerns," in the chapter of the NetMeeting Resource Kit from the previous link (scroll through the chapter until you find the section).

**Establishing a NetMeeting connection with a firewall**
When you use NetMeeting to call other users over the Internet, several IP ports are required to establish the outbound connection.

If you use a firewall to connect to the Internet, it must be configured so that the following IP ports are not blocked:

- TCP ports 389, 522, 1503, 1720, and 1731
- TCP and UDP ports (1024 through 65535)

To establish outbound NetMeeting connections through a firewall, the firewall must be configured to do the following:

- Pass through primary TCP connections on ports 389, 522, 1503, 1720, and 1731
- Pass through secondary TCP and UDP connections on dynamically assigned ports (1024 through 65535)

The H.323 call setup protocol dynamically negotiates a TCP port for use by the H.323 call control protocol. Also, both the audio call control protocol and the H.323 call setup protocol dynamically negotiate UDP ports for use by the H.323 streaming protocol, called the Real-Time Transfer Protocol (RTP). In NetMeeting, two UDP ports are designated on each side of the firewall for audio and video streaming, for a total of four ports for inbound and outbound audio and video. These dynamically negotiated ports are selected arbitrarily from all ports that can be assigned dynamically.

NetMeeting directory services require either port 389 or port 522, depending on the type of server you are using. The Microsoft Internet Locator Service (ILS), which supports LDAP for NetMeeting, requires port 389. The Microsoft User Location Service (ULS), developed for NetMeeting 1.0, requires port 522.

**Firewall limitations for NetMeeting**
Some firewalls cannot support an arbitrary number of virtual internal IP addresses, or cannot do so dynamically. With these firewalls, you can establish outbound NetMeeting connections from computers inside the firewall to computers outside the firewall, and you can use the audio and video features of NetMeeting. Users outside the organization cannot, however, establish inbound connections from outside the firewall to computers inside the firewall. Typically, this restriction is due to limitations in the network implementation of the firewall.

| **Note** |
| --- |
| Some firewalls are capable of accepting only certain protocols and cannot handle TCP connections. For example, if your firewall is a Web proxy server with no generic connection-handling mechanism, you will not be able to use NetMeeting through the firewall. |

You can use the following sources to learn more about NetMeeting configuration and firewall topics:

- For more information about NetMeeting firewall connections, see Part 2, Chapter 4, "Firewall Configuration" in the *NetMeeting 3.0 Resource Kit*, particularly the section titled "Establishing a NetMeeting Connection with a Firewall" at:

  http://go.microsoft.com/fwlink/?LinkId=36134

- For more information about using NetMeeting and your firewall, see article 158623, "How to Establish NetMeeting Connections through a Firewall" in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?LinkId=29206

**Alternate Methods for Controlling NetMeeting**

You can create customized installation options for specific users or groups within your organization by using the NetMeeting Resource Kit Wizard. Additionally, you can use the NetMeeting Resource Kit Wizard to control user and computer access rights by creating custom configurations of client settings and specific features that you have selected to restrict or allow. For example, you can control audio and video access, set data throughput limits and network speeds, and choose to display online support. The Resource Kit Wizard can also help you set up various configurations of NetMeeting for different types of users and different levels of security. It can help you save network bandwidth by restricting specific features. You can also use the Resource Kit Wizard to both change registry settings for all NetMeeting users, and to implement such changes globally.

| **Note** |
| --- |
| By selecting certain options in the Resource Kit Wizard, be aware that you may be changing the NetMeeting user interface. For example, if you click **Restrict the Use of Video**, the Video tab doesn't appear in the NetMeeting user's Options dialog box. |

The Resource Kit for NetMeeting has a section that provides detailed information about responding to NetMeeting problems, including problem descriptions, causes, and resolutions. For more information about the *Microsoft NetMeeting 3 Resource Kit*, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=36134

**Procedures for Configuration of NetMeeting**

NetMeeting is designed to enhance the enterprise environment and enable users to communicate internally and externally with other NetMeeting users. You can use Group Policy to develop a NetMeeting feature management policy to support the specific business rules or communication policies that exist within your organization. For example, your organization may not want users to be able to access or use the NetMeeting chat feature from their computers. By using Active Directory and Group Policy, you can disable the chat feature from any or all computers that are affected by the application of the Group Policy configuration settings.

For lists of Group Policy settings that you can use to manage NetMeeting configuration options, see "NetMeeting and Group Policy," earlier in this section.

**Procedures for managing NetMeeting features through Group Policy**

This subsection provides procedures for the following configuration methods:

- Locating the Group Policy objects (GPOs) for NetMeeting configuration settings. These are the settings listed in "NetMeeting and Group Policy," earlier in this section.

- Disabling the NetMeeting remote desktop sharing feature. This prevents users from using this feature.

**To locate the Group Policy objects (GPOs) for NetMeeting user configuration settings**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **User Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **NetMeeting**.

3. View the Group Policy objects that are available. For more information about these objects, see "NetMeeting and Group Policy," earlier in this section.

For more information about the GPOs for NetMeeting user configuration settings, see "Configuring NetMeeting settings for users through Group Policy," earlier in this section.

**To disable the NetMeeting remote desktop sharing feature through Group Policy**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **NetMeeting**.

3. In the details pane, double-click **Disable remote Desktop Sharing**.

4. Select **Enabled**.

| **Note** |
| --- |
| Computer-related Group Policy settings are applied when the operating system starts and during the periodic refresh cycle. |

**Related Links**

You can learn more about NetMeeting from the following online resources:

- To view Help for NetMeeting, start NetMeeting by clicking **Start**, clicking **Run**, and typing **conf**. In NetMeeting, on the **Help** menu, click **Help Topics**.

- For more information about using NetMeeting and your firewall, see article 158623, "How to Establish NetMeeting Connections through a Firewall," in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?LinkId=29206

- For more information about NetMeeting, see Windows NetMeeting on the Microsoft Web site at:

  http://www.microsoft.com/windows/NetMeeting/

- For more information about configuring NetMeeting, see the Windows NetMeeting Resource Kit on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=36134

- To learn more about NetMeeting features, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29175

- To view articles that explain how to use some of the features in NetMeeting, see the Microsoft NetMeeting How-to Guide on Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29207

- For more information about the H.323 specification, search for “H.323” on the ITUT Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29510

- For more information about the T.120 architecture, see the International Multimedia Teleconferencing Consortium (IMTC) Web site at:

  http://www.imtc.org/

(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

**Printed references**

For more information about firewall design, policy, and security considerations for firewall design in general, you can consult the following reference:


- Chapman, D. Brent and Elizabeth D. Zwicky. *Building Internet Firewalls*. O'Reilly & Associates, Inc., 1995.

# Online Device Help and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

**Online Device Help**
This section provides information about:

- The benefits of Online Device Help
- How Online Device Help communicates with sites on the Internet
- How to control Online Device Help to limit the flow of information to and from the Internet

**Benefits and Purposes of Online Device Help**
Online Device Help (also known as the "Get help for my hardware device" wizard) delivers targeted content on problems with hardware and peripheral devices installed on the system. This mitigates the need for users to call support professionals to resolve hardware issues. Users interact with Online Device Help in products in the Microsoft Windows Server 2003 family when installing new hardware (through the Found New Hardware Wizard).

At the conclusion of the Found New Hardware Wizard, when the user's system uses a device driver that is not found on the operating system installation CD or is not available through Windows Update, Online Device Help collects anonymous data on the problem device (including a unique hardware identifier for that device) and sends that information over the Internet to a server at Microsoft. If a match for the device is found, content on that device is then downloaded to the user's system and is displayed in the Help and Support Center user interface. This content may include:

- Information from the independent hardware vendor (IHV) about upcoming and planned device support.
- Links to the product compatibility area in Help and Support Center to enable users to search or browse the Windows Catalog Web site for compatible devices. The Windows Catalog Web site is located at:

  http://go.microsoft.com/fwlink/?linkid=4303

- A link to the Web site of the IHV for that device.

The data provided by Online Device Help enables Microsoft to identify the number and system locale of users experiencing hardware problems due to missing drivers and to identify the most common problem devices. Microsoft works with these hardware vendors to provide targeted troubleshooting content on the most common problem hardware devices.

This section of the white paper explains how to control Online Device Help in a managed environment.

**Overview: Using Online Device Help in a Managed Environment**
Users have control over whether to upload the data required by Online Device Help. In a managed environment, however, it is unlikely that users can choose to install any device; this function would normally be controlled in some fashion by the IT department. You can block Online Device Help at the firewall or through the Services snap-in. The configuration options and procedures for controlling Online Device Help are described later in this section.

**How Online Device Help Communicates with Sites on the Internet**
If no information for a particular hardware device is found on either the installation CD for the operating system or through Windows Update, users are prompted to release anonymous information about their hardware profile through Online Device Help. This subsection summarizes the communication process:

- **Specific information sent or received**: The following information is collected from the user's computer and uploaded to a server at Microsoft. The user is not uniquely identified.
  - The hardware ID, also known as the PnPID (code that indicates the device manufacturer, device name, and version)

    - The time and date that the data was sent
    - Language code of the operating system, and platform and build information

- **Default and recommended settings**: Online Device Help is enabled by default. Recommended settings are described in the next subsection, "Controlling Online Device Help to Limit the Flow of Information to and from the Internet."
- **Triggers**: Online Device Help is triggered if no information for a particular hardware device is found after the user has completed the Found New Hardware Wizard.
- **User notification**: Users are prompted to send anonymous hardware profile data to a server at Microsoft. If users opt to send this information, the privacy statement is displayed. Users can view the contents of the hardware.xml file being uploaded by clicking a link on the privacy statement page.
- **Logging**: Errors that result from problems installing hardware devices without drivers are logged to the event log.
- **Encryption**: The data transferred to Microsoft is not encrypted.
- **Access**: The raw data uploaded to the server is accessible to operations engineers at Microsoft.com and is used in the Windows Hardware Quality Labs (WHQL) to improve Windows-compatible hardware and drivers.
- **Privacy statement**: Online Device Help is covered by its own privacy statement. The privacy statement (located in a file on the user's computer at *systemroot*\pchealth\helpctr\system\dfs\privacy.htm) is displayed when users opt to send the anonymous hardware profile data to Microsoft.
- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is 80.
- **Ability to disable**: You cannot disable Online Device Help directly. Disabling Internet access or HTTP port 80 will, however, block Online Device Help.

**Controlling Online Device Help to Limit the Flow of Information to and from the Internet**

Users have control over whether to upload anonymous hardware profile information about their system. You cannot, however, disable Online Device Help directly. To block Online Device Help, you can restrict Internet access. You can also use a firewall or configure the Services snap-in. The following table describes the result of each option.

**Configuration settings for Online Device Help**

| Configuration tool | Setting | Result |
|---|---|---|
| Firewall | Block HTTP port 80. | Blocks Online Device Help. |
| Services snap-in | Disable the Upload Manager service (uploadmgr). | Blocks Online Device Help. Any other services that depend on uploadmgr will also fail to start. |

**How controlling Online Device Help can affect users and applications**

If you decide to disable Online Device Help, users will not be prompted to upload anonymous hardware profile information and they will not receive up-to-date, targeted self-help content on hardware issues relating to missing or problem drivers.

| **Note** |
|---|
| If you restrict Internet access to block Online Device Help, the feature will queue the data and periodically retry to upload the hardware profile information for some period of time. If an Internet connection becomes available during that period, Online Device Help will upload the queued data. If an Internet connection does not become available, no data will be uploaded. Users will not be impacted in either case. |

**Alternate Methods for Controlling Online Device Help**

You can also control Online Device Help by disabling the Upload Manager service (uploadmgr) that manages synchronous and asynchronous file transfers between clients and servers on the network. Disabling this service will block the upload of the anonymous hardware profile data (although users will still be able to complete the Found New Hardware Wizard). The following subsection gives the procedure for this method.

**Procedure for Controlling Online Device Help**

You cannot disable Online Device Help directly but can do so indirectly by disabling the Upload Manager service in products in the Windows Server 2003 family.

**To disable Online Device Help by disabling the Upload Manager service**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then double-click **Services**.

3. In the details pane, right-click **Upload Manager**, and then click **Properties**.

4. Click the **Log On** tab, then click the hardware profile that you want to configure, and then click **Disable**.

| **Important** |
| --- |
| If the Upload Manager service is disabled, any services that explicitly depend on it will fail to start. |

# Outlook Express 6.0 and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The subsections that follow provide:

- A description of Microsoft Outlook® Express 6.0, which is included in Microsoft Internet Explorer 6.0, and a comparison of Outlook and Outlook Express.
- Descriptions of new security-related features in Outlook Express 6.0 (as compared to Outlook Express 5), with information about how they are configured at the desktop.
- Information about controlling Outlook Express 6.0 through Group Policy to limit the risk associated with e-mail attachments. The Group Policy setting you use for this is **Block attachments that could contain a virus**.

| **Note** |
| --- |
| This section of the white paper describes Outlook Express 6.0, but it does not describe Internet Explorer 6.0 (of which Outlook Express is part), the New Connection Wizard, or the tool that can report errors that occur in Outlook Express. For information about these components, see the respective sections of this white paper; the error reporting tool is described in Windows Error Reporting and Internet Communication (Windows Server 2003).<br><br>Also note that the New Connection Wizard replaces the Network Connection Wizard and the Internet Connection Wizard in Windows 2000. |

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization where users send e-mail, receive e-mail, open attachments in e-mail, and perform similar actions. This section, however, provides information about features and configuration methods in Outlook Express 6.0 that can reduce the inherent risks associated with sending and receiving e-mail.

For more information about Outlook Express, see the following resources:

- Help for Outlook Express (which can be accessed in Outlook Express by clicking the **Help** menu and then selecting an appropriate option).
- The section about Internet Explorer 6.0 in this white paper, which describes security zones in Internet Explorer 6.0. These security zones are also used in Outlook Express 6.0.
- The Internet Explorer page on the Microsoft Web site at:

  http://www.microsoft.com/windows/ie/

- The Resource Kit for Internet Explorer (specifically, the chapter describing what's new in Internet Explorer 6.0). To learn about this and other Resource Kits, see the Microsoft TechNet Web site at:

  http://go.microsoft.com/fwlink/?linkid=29894

**Benefits and Purposes of Outlook Express 6.0**

Outlook Express 6.0 is designed to make it easy to send or receive e-mail and to browse or participate in newsgroups. It differs from most of the other components described in this white paper in that its main function is to communicate through the Internet or an intranet (in contrast to components that communicate with the Internet in the process of supporting some other activity).

Outlook Express is part of Internet Explorer, in contrast to Microsoft Outlook, which is an application included in Microsoft Office. Outlook provides comprehensive e-mail capabilities, including information management and collaboration capabilities,

useful to a wide spectrum of users from home to small business to large enterprise. Outlook Express, included as part of Internet Explorer, offers standard Internet e-mail and news access, useful to many home and small-business users. Outlook Express supports Post Office Protocol 3 (POP3) or Internet Message Access Protocol (IMAP).

Outlook Express 6.0 offers more security-related options and settings than were available in Outlook Express 5, as described in the subsections that follow.

**New Security-Related Features in Outlook Express 6.0**

Outlook Express 6.0 is the e-mail component in Internet Explorer 6.0. This version of Outlook Express includes the following new security-related features. The table that follows this list shows how each option is configured in Outlook Express.

- **Warning about harmful e-mail**. To prevent e-mail messages from being sent without your knowledge, Outlook Express warns you when other programs, such as viruses or harmful attachments, attempt to send messages from your computer. This warning appears only if Outlook Express is configured as the default simple MAPI client, and another program attempts to use simple MAPI to programmatically send e-mail messages without presenting a visible user interface on the computer.

- **Blocking of potentially harmful attachments**. If this option is enabled, Outlook Express 6.0 blocks the opening or saving of specific e-mail attachments that are considered "unsafe." To determine whether an attachment is unsafe, Outlook Express 6.0 uses the Internet Explorer 6.0 unsafe file list, plus some additional file types, plus file types you configure with the **Confirm open after download** setting in Folder Options (on the Files Types tab). Any e-mail attachment with a file type reported as "unsafe" is blocked. This option can be enabled or disabled through Group Policy as well as at the local computer. For more information about using this setting, see the table that follows and "To locate the Group Policy object (GPO) for blocking e-mail attachments in Outlook Express 6.0," later in this section.

  For information about the unsafe file list in Internet Explorer 6.0, you can search the Microsoft Knowledge Base. To do this, follow the instructions for searching on the Web site, and search for the phrase "unsafe file list":

  http://support.microsoft.com/

- **Software Restriction Policies technology**. When running with an operating system in the Microsoft Windows Server 2003 family, Outlook Express 6.0 takes advantage of Software Restriction Policies technology to run potentially harmful attachments in a sandbox, which is an area in memory outside of which the program cannot make calls. When you attempt to run or save attachments, Software Restriction Policies technology determines whether the file formats are blocked. If so, Outlook Express displays a warning, and the program running the attachment has only limited access to the computer's hard disk and registry.

- **Plain text format option for reading of e-mail**. Starting with Outlook Express 6.0, Outlook Express can be configured to read all e-mail messages in plain text format. Some HTML e-mail messages may not appear correctly in plain text, but no active content in the e-mail message is run when this setting is enabled.

The following table shows how each option is configured in Outlook Express 6.0 .

**Options for configuring Outlook Express 6.0**

| Option to configure in Outlook Express 6.0 | Menu to click | Menu item to click | Tab to click |
|---|---|---|---|
| Warning about harmful e-mail | Tools | Options | Security |
| Blocking of potentially harmful attachments<br><br>(also configurable through Group Policy) | Tools | Options | Security |
| Software Restriction Policies technology | Tools | Options | Security |
| Plain text format option for reading of e-mail | Tools | Options | Read (in Outlook Express 6.0 only) |

**Overview: Using Outlook Express 6.0 in a Managed Environment**

Although there are inherent risks associated with sending and receiving e-mail (and e-mail attachments), you can use several different features and configuration methods in Outlook Express 6.0 to reduce the risks:

- You can use the graphical user interface to configure the security-related features in Outlook Express 6.0. For more information, see "New Security-Related Features in Outlook Express 6.0," earlier in this section and "To start Outlook Express 6.0 and view or configure security settings," later in this section.

- You can use a Group Policy setting, **Block attachments that could contain a virus**, to limit the risk associated with e-mail attachments in Outlook Express 6.0. For more information, see "To locate the Group Policy object (GPO) for blocking e-mail attachments in Outlook Express 6.0," later in this section.

**Procedures for Working with Outlook Express 6.0**

This subsection provides procedures for the following:

- Opening the dialog box from which you can configure security settings for Outlook Express 6.0.

- Locating the Group Policy setting, **Block attachments that could contain a virus**.

  You can use this Group Policy setting in situations where you want Outlook Express 6.0 to be available but where you want to limit the risk associated with e-mail attachments. For more information about this policy setting, see "New Security-Related Features in Outlook Express 6.0," earlier in this section.

**To start Outlook Express 6.0 and view or configure security settings**

1. Click **Start**, point to **All Programs** or **Programs**, and then click **Outlook Express**.

2. On the **Tools** menu, click **Options**.

3. Click the **Security** tab and view or configure the settings, including the check boxes for the following two options:

   - **Warn me when other applications try to send mail as me**.

   - **Do not allow attachments to be saved or opened that could potentially be a virus**.

   You can also view or configure the security zones setting. Outlook Express 6.0 uses two of the same security zones that you configure in Internet Explorer 6.0. For more information about security zones, see the section about Internet Explorer 6.0 in this white paper.

4. Click the **Read** tab, and view or configure the settings, including the check box for **Read all messages in plain text**.

**To locate the Group Policy object (GPO) for blocking e-mail attachments in Outlook Express 6.0**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **User Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Internet Explorer**.

3. In the details pane, double-click **Configure Outlook Express**.

4. Select or clear the check box for **Block attachments that could contain a virus**.

# Plug and Play and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Plug and Play
- How Plug and Play communicates with sites on the Internet
- How to control Plug and Play to prevent the flow of information to and from the Internet

**Benefits and Purposes of Plug and Play**

Windows Plug and Play provides ease of support for installing devices on computers in your network. You can simply plug in a Plug and Play device and the operating system does the rest by installing any necessary drivers, updating the system, and allocating resources. After you install a Plug and Play device, the driver is configured and loaded dynamically, typically without requiring user input.

Plug and Play in the Microsoft Windows Server 2003 family provides the following services:

- Detects a Plug and Play device and determines its hardware resource requirements and device identification number (Plug and Play ID).
- Locates an appropriate device driver for newly installed devices.
- Allocates hardware resources.
- Dynamically loads, initializes, and unloads drivers.
- Notifies other drivers and applications when a new device is available.
- In conjunction with power management, handles stop and start processes for devices during hibernation, standby, and startup and shutdown operations.
- Supports a wide range of device types.

**Overview: Using Plug and Play in a Managed Environment**

When you install a Plug and Play device, and you are connected to the Internet, the operating system automatically accesses Windows Update to search for a device driver.

| **Note** |
| --- |
| Some buses, such as Peripheral Component Interconnect (PCI) and universal serial bus (USB), take full advantage of Plug and Play. Older buses, such as Industry Standard Architecture (ISA), do not take full advantage of Plug and Play and require more user interaction to ensure that devices are correctly installed. |

In order to install devices using the hardware wizards, you must be logged on as an administrator or a member of the Administrators group. You then use the hardware wizards, such as the Hardware Update Wizard, to search the Windows Update site for device drivers. All drivers obtained through Windows Update are signed by Windows Hardware Quality Labs (WHQL). The WHQL provides compatibility testing services to test hardware and drivers for Windows operating systems.

As an IT administrator in a highly managed network environment, you want to control the ability of administrators to install new hardware and to thereby access the Internet automatically when the operating system searches for device drivers. For a more secure environment you can control how administrators update and install hardware devices by using Group Policy.

There are also policy settings you can use to disable any access to Windows Update. If you do prevent certain administrators from automatically accessing Windows Update, there is the option for manually downloading the updates from the Windows Update Catalog, whereby they can be distributed on your organization's network as needed.

Using Group Policy to disable automatic updating and access to Windows Update, and to configure driver search locations, is described in the subsection, "Controlling Automatic Device Updating to Prevent the Flow of Information to and from the Internet."

**How Plug and Play Communicates with Sites on the Internet**

There are two instances when a computer running a product in the Windows Server 2003 family operating system will access the Internet as part of Plug and Play:

- When Plug and Play searches for a driver for newly installed hardware
- When an administrator updates the driver for existing hardware

When you connect a new hardware device and there is no driver available on the computer, the operating system will use the Windows Update service to search for available drivers on the Windows Update site. If an appropriate driver is found on the Windows Update site, the operating system copies it and installs it on your computer. If your computer is not connected to the Internet, a message prompts you to connect to the Internet.

As part of Plug and Play, when the operating system searches for a device driver, interaction with the Internet takes place as follows:

- **Specific information sent or received**: The Code Download Manager (CDM) calls Windows Update to find and download device drivers. The CDM also calls Help and Support Center, which logs Plug and Play IDs for devices that Microsoft does not have drivers for. Neither of these communications is under the direct control of Plug and Play. The CDM handles all of the communication between the computer and Windows Update. None of the communication between the computer and the Internet uniquely identifies the user.
- **Default and recommended settings**: Plug and Play is enabled by default. Recommended settings are presented in the following subsection, "Controlling Automatic Device Updating to Prevent the Flow of Information to and from the Internet."
- **Triggers**: When an administrator adds hardware or updates a driver on a computer, and the computer is connected to the Internet, Windows Update is automatically contacted for driver updates.
- **User notification**: When searching for a device driver Windows Update sends a list of available drivers to the user's computer.
- **Logging**: If you use a Plug and Play driver with a non-Plug and Play device, any associated issues or problems are recorded in the event log.
- **Encryption**: Data transfer is based on interaction with Windows Update. The data is transferred using HTTPS.
- **Transmission protocol and ports**: The transmission protocols and ports are HTTP 80 and HTTPS 443.
- **Ability to disable**: Plug and Play cannot be disabled as system instability would result. You can disable access to Windows Update using Group Policy.

**Controlling Automatic Device Updating to Prevent the Flow of Information to and from the Internet**

Windows Server 2003 family operating systems will automatically update device drivers using Plug and Play, and they will even search for compatible drivers for non-Plug and Play devices. You therefore might want to exercise various levels of control over administrators' ability to install new hardware and to update hardware devices and drivers.

Using Group Policy there are several levels of control you can configure in order to prevent Plug and Play and associated hardware wizards from accessing the Internet. You can target search locations for drivers, or you can prevent users and computers from automatically accessing the Windows Update Web site in any instance. You can disable automatic updating for some servers and enable it for others, and then have client computers and servers access an intranet server for selected

updates.

You can use Group Policy to:

- Control whether Windows Update is included when Plug and Play searches for a device driver.

  This procedure is presented in the next subsection.

- Eliminate automatic update calls to Windows Update.

  Policy settings related to automatic updating are located at Computer Configuration\Administrative Templates\Windows Components\Windows Update.

  If you disable **Configure Automatic Updates**, any updates that are available on the Windows Update Web site must be downloaded and installed manually.

- Remove access to Windows Update.

  The policy setting for Windows Update is located at User Configuration\Administrative Templates\Windows Components\Windows Update.

  When you enable the policy setting **Remove access to use all Windows Update features**, you block access to the Windows Update site from the Windows Update hyperlink on the Start menu and also on the Tools menu in Microsoft Internet Explorer. Automatic updating is also disabled; you will neither be notified about nor will you receive critical updates from Windows Update. This policy setting also prevents Device Manager from automatically installing driver updates from the Windows Update Web site.

  The Windows Update site is located at:

  http://go.microsoft.com/fwlink/?LinkId=284

**Procedure for Controlling Where Plug and Play Searches for Drivers**

When you install new hardware, the operating system automatically searches four different locations for drivers in the following order: the hard drive, the floppy drive, the CD-ROM drive, and Windows Update. The default approach is to search all four locations successively until the correct device driver is found; however, you can configure the driver search locations to remove selected locations.

Included here is the procedure for configuring the Group Policy setting **Configure driver search locations**. For additional procedures to configure policy settings for Windows Update, see the section Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003) in this white paper.

**To disable Windows Update as a driver search location for Plug and Play devices**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **User Configuration**, click **Administrative Templates**, and then click **System**.

3. In the details pane, double-click **Configure driver search locations**, and then select **Enabled**.

4. Select **Don't Search Windows Update**.

**Related Links**

For more information about Windows Update, see the Windows Update Web site at:

http://windowsupdate.microsoft.com/

# Program Compatibility Wizard and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of the Program Compatibility Wizard
- How the Program Compatibility Wizard communicates with sites on the Internet
- How to control the Program Compatibility Wizard to prevent the flow of information to the Internet

**Benefits and Purposes of the Program Compatibility Wizard**
Most programs run properly on products in the Microsoft Windows Server 2003 family. The exceptions are some older games and other programs that were written specifically for an earlier version of Windows.

To enable a better user experience, Microsoft has integrated technologies for application compatibility into the Windows Server 2003 family. Application compatibility technologies are applied whenever an application is installed on the operating system, whether in the course of a system upgrade or during regular operations. Some of these technologies work automatically to apply compatibility fixes, while others can be selected by users or administrators.

This section addresses a component that users and administrators can use, the Program Compatibility Wizard. If you have an application compatibility problem, you can use the wizard to make setting adjustments and to run the application successfully.

**Overview: Using the Program Compatibility Wizard in a Managed Environment**
IT administrators who want to get an application to work quickly, without addressing compatibility for the application throughout their organization, may choose to use the Program Compatibility Wizard. You can use the wizard to detect and test compatibility settings. If compatibility problems prevent you from installing a program, you can run the Program Compatibility Wizard on the Setup file for the program.

In the Windows Server 2003 family you can access the Program Compatibility Wizard by default through Programs\Accessories or All Programs\Accessories. At the completion of the wizard you are presented with “Program Compatibility Data” that states, “Microsoft has created temporary files that contain information about the settings you selected and whether the problems were fixed. Sending this information to Microsoft will help us improve program compatibility.” You can then choose to send this information to Microsoft or not. You can also choose to click a link to a Web site to view the data collection policy and you can choose to view the temporary files that will be sent. Allowing users or administrators to do this, however, may present a privacy problem for highly managed organizations.

| **Note** |
| --- |
| As an alternative to running the Program Compatibility Wizard, you can set the compatibility properties manually through the Compatibility tab of a program’s Properties sheet. To do this you right-click the program icon, click **Properties**, click **Compatibility**, and then change the compatibility settings. |

Administrators can use Group Policy to disable access to the Program Compatibility Wizard. Alternatively, if you want to allow use of the wizard you can control where data collected by the Program Compatibility Wizard is sent. You can prevent data transfer to the Internet by using Group Policy settings related to error reporting and you can have data from the wizard sent to a server on your intranet instead of to Microsoft. For more information about these procedures, see "Controlling Program Compatibility Wizard Data to Prevent the Flow of Information to the Internet," later in this section.

**How the Program Compatibility Wizard Communicates with Sites on the Internet**
Although you can control information sent by the Program Compatibility Wizard, it is designed to communicate over the Internet to expedite problem solving. This subsection lists details of the communication process:

- **Specific information sent or received**: The results of the Program Compatibility Wizard data, including settings and problems that were encountered with the application being installed, are sent to Microsoft. The user is not uniquely identified.

- **Default and recommended settings**: Use of the Program Compatibility Wizard is enabled by default. Recommended settings are discussed in the next subsection, "Controlling Program Compatibility Wizard Data to Prevent the Flow of Information to the Internet."

- **Trigger and notification**: In the last dialog box of the wizard, users are asked if they want to send information to Microsoft. Data is not sent automatically.

- **Logging**: There is no information related to the Program Compatibility Wizard entered into the event log.

- **Encryption**: HTTPS is used to perform the data transfer to Microsoft.

- **Access**: The Microsoft product group has access to the raw data only.

- **Privacy statement**: The privacy statement is the same as that associated with Windows Error Reporting (WER) data. A link to the privacy statement on the Web is provided in the wizard.

- **Transmission protocol and port**: The transmission protocol used is HTTPS and the port is 443.

- **Ability to disable**: You can disable the Program Compatibility Wizard by using Group Policy.

For more information about the type of information sent to Microsoft, how the data is used, encryption, and the privacy statement, see the section of this white paper titled Windows Error Reporting and Internet Communication (Windows Server 2003).

**Controlling Program Compatibility Wizard Data to Prevent the Flow of Information to the Internet**
To control use of the Program Compatibility Wizard you can either disable it or control where information collected by the wizard is sent.

**Disabling the wizard**
If you want to prevent the use of the Program Compatibility Wizard you can disable it by using Group Policy. In Computer Configuration\Administrative Templates\Windows Components\Application Compatibility, configure **Turn Off Program Compatibility Wizard**. This policy setting does the following:

- When enabled, this policy setting disables the start page of the wizard in Help and Support Center, and in the Start menu.

- These entry points will still exist, but the first page of Help and Support Center will let the user know that this option has been disabled.

**Routing the data to an intranet server**
Using Group Policy you can configure the **Report Errors** policy setting to prevent data collected by the Program Compatibility Wizard from being sent to Microsoft. By using configuration options within error reporting you can have the data sent to a server on your intranet instead of to Microsoft. When you configure error reporting this way you activate Corporate Error Reporting (CER).

The **Report Errors** policy setting is located in Computer Configuration\Administrative Templates\System\Error Reporting. For more information and procedures for configuring error reporting see the section of this white paper titled Windows Error Reporting and Internet Communication (Windows Server 2003).

If you use this approach for reporting errors, the user experience with the Program Compatibility Wizard does not change. The dialog box that presents the option of sending data to Microsoft is the same. If the user selects Yes, the data is sent to the designated server on your intranet.

**Procedure for Disabling the Program Compatibility Wizard**
Use the following procedure to disable the Program Compatibility Wizard. For the procedure to configure error reporting see the section, Windows Error Reporting and Internet Communication (Windows Server 2003).

To disable the Program Compatibility Wizard

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn

about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Application Compatibility**.

3. In the details pane, double-click **Turn Off Program Compatibility Wizard**, and then click **Enabled**.

# Remote Assistance and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Remote Assistance
- How Remote Assistance communicates with sites on the Internet
- How to control Remote Assistance to prevent the flow of information to and from the Internet

**Benefits and Purposes of Remote Assistance**
With operating systems in the Microsoft Windows Server 2003 family, users and administrators in your organization can use Remote Assistance to get help from a member of your support staff. Users or administrators can also collaborate in other ways through screen sharing. Remote Assistance is a convenient way for support professionals to connect to a computer from another computer running a compatible operating system, such as Windows XP, and to show the users or administrators a solution to their problem.

Using Windows Messenger Service or an e-mail program, such as Microsoft Outlook or Outlook Express, you can provide support to users by connecting to their computer. After you are connected you can view their computer screen, communicate with them in real time about what you both see on their computer, send files, use voice communication, and use your mouse and keyboard to work on their computer.

**Overview: Using Remote Assistance in a Managed Environment**
Remote Assistance is disabled by default. You configure Remote Assistance through Control Panel\System\Remote tab on an individual computer, or through Group Policy for groups of users or computers. When it is enabled users and administrators can access Remote Assistance through Help and Support Center under Support Tasks\Support, or Support Tasks\Tools\Help and Support Center Tools.

While a firewall on your organization's network will likely prevent outsiders from connecting directly to a computer on your intranet, the potential for users or administrators to connect remotely to someone outside your network is available through Remote Assistance. There is also the option of a support person or IT administrator in your organization offering unsolicited assistance. An administrator in the domain or a user explicitly authorized through Group Policy settings may offer assistance to users in the same domain without being asked; however, users can decline the offer.

As an administrator in a highly managed environment you might want to prevent access for groups of users and administrators to this feature. You can do this through Group Policy. Controlling the use of unsolicited as well as solicited Remote Assistance is described further in the subsection "Controlling Remote Assistance to Prevent the Flow of Information to and from the Internet."

**How Remote Assistance Communicates with Sites on the Internet**
When a user (referred to as the "novice") initiates a request for assistance through either the e-mail option or the Save invitation as a file option in Remote Assistance, the operating system starts Help and Support Center. Help and Support Center then passes the information to Remote Assistance.

When the person who is being contacted (the "expert") accepts the invitation from the novice, Remote Assistance calls Help and Support Center application programming interfaces (APIs) to initiate the session. Help and Support Center relies on Terminal Services to negotiate the session. Help and Support Center passes the Remote Assistance invitation (the "ticket") file to Terminal Services. The Remote Assistance session is established using RDP (Remote Desktop Protocol) and port 3389 through Terminal Services on the novice and expert computers.

There are safeguards built into the Remote Assistance feature. All sessions are encrypted and can be password-protected. The novice (user soliciting the assistance) sets the maximum time for the duration of the ticket. Also, firewalls on your organization's network might prevent users from making a connection.

The following information presents additional details on how information transfer over the Internet takes place when a connection is made:

- **Specific information sent or received**: Information that is transmitted in a Remote Assistance ticket includes user name, IP address, and computer name. Information necessary to provide functionality for Remote Assistance (for example, screen sharing, file transfer, and voice) is sent in real time using point-to-point connections.

- **Default and recommended settings**: Anyone with access to Help and Support Center can access the Remote Assistance feature. Users can prevent someone from connecting to their computer by declining an invitation. You can also prevent someone from remotely controlling a server running one of the Windows Server 2003 family operating systems through Control Panel settings or Group Policy.

- **Triggers**: A user or administrator establishes contact with the expert by sending an invitation through e-mail, instant messaging, or by saving an invitation as a file and transferring it manually, such as on a floppy disk, to the expert. Or, an expert offers unsolicited assistance to a user.

- **User notification**: The expert is asked through e-mail or instant messaging to provide help to the novice. A connection is not made unless the expert accepts the invitation or opens the ticket. When users are offered unsolicited assistance, they as the novice have to click Yes to start a connection.

- **Logging**: Events such as a person initiating a connection or a user or administrator accepting or rejecting an invitation are recorded in the event logs.

- **Encryption**: The RDP (Remote Desktop Protocol) encryption algorithm for the main Remote Assistance communication and the RTC (Real-Time Communication) encryption algorithm for voice are used. The RDP encryption algorithm is RC4 128-bit.

- **Access**: No information is stored at Microsoft.

- **Transmission protocol and port**: The port is 3389 and the transmission protocols are RDP and RTC.

- **Ability to disable**: Yes, using Group Policy, and locally through Control Panel.

- **Firewall protection**: Any firewall that blocks port 3389 should not allow a connection to users outside the firewall. This does not prevent users from within the network protected by the firewall from connecting to each other. If you close port 3389, you will block all Remote Desktop and Terminal Services events through it as well. If you want to allow these services but want to limit Remote Assistance requests, use Group Policy. If the port is opened only for outbound traffic, a user can request Remote Assistance by using Windows Messenger.

For more information about the Remote Assistance connection process, see article 300692, "Description of the Remote Assistance Connection Process," in the Microsoft Knowledge Base at:

http://go.microsoft.com/fwlink/?LinkId=29212

**Controlling Remote Assistance to Prevent the Flow of Information to and from the Internet**
Administrators can control the use of Remote Assistance in the following ways:

- Group Policy to disable users or administrators from soliciting or offering Remote Assistance

- Local controls of Remote Assistance through Control Panel

Group Policy settings are described in detail in this subsection. Procedures for disabling Remote Assistance are presented in the next subsection.

There are two Group Policy settings you can configure to control the use of Remote Assistance:

- **Solicited Remote Assistance**

  Use this policy setting to determine whether or not solicited remote assistance is allowed from a computer. In **Solicited Remote Assistance** the user of a computer explicitly requests help from another party.

- **Offer Remote Assistance**

  Use this policy setting to determine whether a support person or IT administrator (expert) can offer remote assistance to a computer without a user explicitly requesting it first, through e-mail, a file, or instant messaging.

These policy settings are located in Computer Configuration\Administrative Templates\System\Remote Assistance. Configuration options for these policy settings are described in the following table.

**Group Policy settings for controlling Remote Assistance**

| Policy setting | Description |
|---|---|
| **Solicited Remote Assistance** (enabled) | When this policy setting is enabled, a user can create a Remote Assistance invitation that a person ("expert") can use at another computer to connect to the user's computer. If given permission, the expert can view the user's screen, mouse, and keyboard activity in real time.<br><br>Additional configuration options are available when you enable this policy setting. |
| **Solicited Remote Assistance** (disabled) | If the status is set to Disabled, users cannot request Remote Assistance and this computer cannot be controlled from another computer. If this policy setting is disabled, the **Offer Remote Assistance** policy setting is also disabled. |
| **Solicited Remote Assistance** (not configured) | If the status is set to Not Configured, all additional configuration is determined by the Control Panel settings. |
| **Offer Remote Assistance** (enabled) | When this policy setting is enabled, a remote user or administrator can offer Remote Assistance to the computer. When you configure this policy setting, you have two choices: you can select either "Allow helpers to only view the computer" or "Allow helpers to remotely control the computer." In addition to making this selection, when you configure this policy setting you also specify the list of users or user groups that will be allowed to offer remote assistance. Administrators can offer remote assistance by default; they do not need to be added to the list. |
| **Offer Remote Assistance** (disabled or not configured) | If you disable or do not configure this policy setting, users or groups cannot offer unsolicited remote assistance to this computer. |

For additional configuration options see the Remote Assistance policy settings in Group Policy. To find more information about editing Group Policy, see Appendix B: Resources for Learning About Group Policy (Windows Server 2003).

**Procedures for Disabling Remote Assistance**

This subsection presents procedures administrators can use for disabling Remote Assistance through Group Policy or Control Panel.

**To disable the use of Remote Assistance using Group Policy**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, and then click **Remote Assistance**.

3. In the details pane, double-click **Solicited Remote Assistance,** and then select **Disabled**.

| Note |
| --- |
| When you disable this policy setting, **Offer Remote Assistance** is also disabled. |

**To disable the use of Remote Assistance through Control Panel**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **System**.

3. In System Properties, click the **Remote** tab.

4. Under Remote Assistance, clear the check box labeled **Turn on Remote Assistance and allow invitations to be sent from this computer**.

# Search Companion and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Search Companion
- How Search Companion communicates with sites on the Internet
- How to control Search Companion to prevent the flow of information to and from the Internet

**Benefits and Purposes of Search Companion**
The Microsoft Search Companion Web service enables users to search for files and folders on their desktop computer, to search for files, people, and other computers on their internal network, and to search for information on the Internet. Search Companion uses Indexing Service to maintain an index of all the files on users' computers, making searches faster.

When employing Search Companion, users can specify several search criteria. For example, they can search for files and folders by name, type, or size. They can find files based on when they were last modified, or search for files containing specific text. When searching for information on the Internet, Search Companion enables users to enter search queries in natural language (meaning informal, or conversational language). It then suggests the best way to conduct the search, and sends the query to Internet services that are most likely to yield positive results.

**Overview: Using Search Companion in a Managed Environment**
When the user searches the Internet using Search Companion, the following information is collected:

- The text of the Internet search query
- Grammatical information about the query
- The list of tasks (suggestions) that the Search Companion Web service recommends to refine the search
- Any tasks the user selected from the recommendation list

Search Companion does not collect:

- Personal information
- Demographic information

Microsoft does not use the information it collects to identify the user individually and does not use such information in conjunction with other data sources that may contain personal data. Microsoft does not collect information when the user searches on the local system, LAN, or intranet.

The Search Companion Web service is designed to upgrade automatically as new features become available. It therefore uses the Internet connection periodically to check for and replace necessary files.

If you want to disable the Search Companion Web service, you can do so by changing to Classic Search for the Internet. Microsoft Windows does not collect any query information when Classic Search is used. You can also disable Search Companion by modifying the registry settings. The procedures for both of these methods are described later in this section of the white paper.

**How Search Companion Communicates with Sites on the Internet**
Search Companion in Microsoft Windows Server 2003 family operating systems improves the search process by consolidating search tasks, optimizing searches for the most common scenarios, and offering suggestions for refining the search.

The form the user creates to collect search criteria will post information to an ASP page that displays the search results. The search pages use a combination of XML and Microsoft Visual Basic® development system, Scripting Edition (VBScript) and Microsoft JScript® development software for accessing the search objects. Because the script is run on the server, users can view the search pages from any browser.

Search Companion uses XML files to define both the user interface (UI) and some functional parameters of its tasks (for example, what list of file extensions constitutes the "Music" category of files). The first time in each Search Companion session that an XML file is referenced, Search Companion checks to see if a later version of that XML file is available from sa.windows.com. The "check" is really a file download request, conditioned on the modified date of the file. If there is a later version of the XML file, Search Companion downloads it and replaces the earlier version. The XML files are located in a language-specific subfolder of *systemroot*\srchasst\, and if the current user does not have administrative credentials, the old XML file cannot be overwritten.

This subsection describes various aspects of the data that is sent to and from the Internet through Search Companion, and how the exchange of information takes place:

- **Specific information sent or received**: When you search the Internet using Search Companion, the following information is collected regarding your use of the service: the text of your Internet search query, grammatical information about the query, the list of tasks that the Search Companion Web service recommends, and any tasks you select from the recommendation list.

- **Default and recommended settings**: Search Companion is enabled by default.

- **Triggers**: The user selects Start\Search and uses search options to search the Internet.

- **User notification**: There is no provision in Search Companion for user review or notification of data sent.

- **Uniquely identify users**: The user is not uniquely identified. Session-based cookies are used to maintain state information, but these randomly assigned GUIDs do not persist across browser sessions.

- **Logging**: No information is collected when you search your local system, LAN, or intranet. The only "storage" is the Internet Information Services (IIS) log of the file request on the server at Microsoft that provides the Search Companion Web service. Search Companion does not record your choice of Internet search engines, and it does not collect or request any personal or demographic information.

- **Encryption**: There is no encryption of data.

- **Access**: No user information is collected. The IIS logs (described in the "Logging" item, earlier in this list) are cycled annually, that is, logs are retained for twelve months, and discarded in the thirteenth month following collection.

- **Privacy statement**: The privacy statement is located on the following Web site at:

  http://sa.windows.com/privacy/

- **Transmission protocol and port**: The transmission protocol is HTTP and the port is 80.

- **Ability to disable**: The feature can be disabled by changing to Classic Search.

**Controlling Search Companion to Prevent the Flow of Information to and from the Internet**

You can disable the Search Companion Web service by changing preferences to Classic Search for the Internet. You can also disable Search Companion by changing the registry settings manually. Procedures for both of these approaches are provided in the following subsection.

**Procedures for Configuration of Search Companion**

Search Companion can be configured in several ways as described previously. This subsection lists the procedures to change or disable the features in accordance with your organization's security policies.

**To change to Classic Search for the Internet**

1. Click **Start**, and then either click **Search**, or point to **Search** and click **On the Internet**.

2. Click **Change preferences**.

3. Click **Change Internet search behavior**.

4. Click **With classic Internet search**.

**To disable Search Companion through the registry key**

1. Close Internet Explorer (all instances).

2. Open Registry Editor by clicking **Start**, clicking **Run**, and then typing **regedit**.

> **Caution**
> Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

3. Navigate to HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\CabinetState

4. On the **Edit** menu, point to **New**, and then click **String value**.

5. Type **Use Search Asst** as the name for the new value (the type is REG_SZ), and then press ENTER.

6. Click the new entry (**Use Search Asst**), and then on the **Edit** menu, click **Modify**.

7. For **Value data**, type:

   **no**

> **Note**
> Type the entry and value exactly as shown, including spaces and capitalization.

# Terminal Server Licensing and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The purposes of Terminal Server Licensing
- How Terminal Server Licensing communicates with sites on the Internet
- How to control Terminal Server Licensing to limit the flow of information to and from the Internet

| **Note** |
| --- |
| In operating systems in the Microsoft Windows Server 2003 family, several components that use technology related to Terminal Server have different names than they had in Windows 2000. Terminal Server was previously known as Terminal Services in Application Server mode. Remote Desktop for Administration was previously known as Terminal Services in Remote Administration mode. |

**Purposes of Terminal Server Licensing**

Terminal Server Licensing is an administrative tool that helps manage your Terminal Server licenses. Terminal Server Licensing can be installed on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. If you install Terminal Server, you must also install Terminal Server Licensing (although not necessarily on the same computer). However, there is a grace period of 120 days during which you can use Terminal Server even though you have not set up license servers yet and have not installed client license key packs.

| **Note** |
| --- |
| Terminal Server and Terminal Server Licensing are not available in Windows Server 2003, Web Edition. Remote Desktop for Administration, however, is available in Windows Server 2003, Web Edition. |

When you set up Terminal Server on a server, users can access applications running on that server, which means they can run the applications you provide even if they are working from hardware that might not otherwise support the applications.

Microsoft has introduced new licensing options to address business customer needs and to complement the technical capabilities of the Windows Server 2003 family of products. For details about these changes, see the Windows Server 2003 licensing overview at:

http://go.microsoft.com/fwlink/?LinkId=46007

**Overview: Using Terminal Server Licensing in a Managed Environment**

You can control the communication that occurs between the Terminal Server Licensing component and sites on the Internet by choosing the server or servers on which to install the Terminal Server Licensing component, and by choosing among three methods for activation. (You also choose among the same three activation methods for obtaining client license key packs, which are digital representations of a group of client access licenses.) The three methods for activation are as follows:

- **Automatic**. This method requires the server running Terminal Server Licensing to have an Internet connection.
- **Web Browser**. This method requires an Internet connection, but the connection can be made from any computer, not just the computer that is running Terminal Server Licensing.
- **Telephone**. This method is used to contact the nearest Customer Support Center and receive an ID number to activate the Terminal Server License Server by phone.

For more details on these methods, see the subsections that follow.

**How Terminal Server Licensing Communicates with Sites on the Internet**

A server running the Terminal Server Licensing component communicates with the Microsoft Clearinghouse (a database for managing licensing) on the Internet only when you activate Terminal Server Licensing or when you initiate subsequent transactions with Microsoft to obtain client license key packs. The following list describes the communication that occurs when you activate or obtain client license key packs directly over the Internet or when you connect to the Microsoft Clearinghouse from a computer other than the one where Terminal Server Licensing is installed.

| **Note** |
| --- |
| The information in the following list applies only to activation or obtaining client license key packs over the Internet. It does not apply when you activate by phone.<br><br>Activation by phone is done as follows: When you activate by phone, you start the licensing wizard, choose a country or region from the list that is displayed, and call the number shown. |

The rest of this subsection describes various aspects of the data that is sent to and from the Internet from Terminal Server Licensing and how the exchange of information takes place.

- **Specific information sent or received**: The information sent to the Microsoft Clearinghouse includes company name, first and last name of the user, license server name, and license server ID. Client license key packs are returned to the Terminal Server License Server.

- **Default and recommended settings**: Terminal Server Licensing is not installed by default.

- **User notification and triggers**: The administrator triggers activating, obtaining client license packs, or deactivating Terminal Server Licensing by performing the steps described in the subsection "Procedures for Configuration of Terminal Server Licensing." When the Terminal Server Licensing component starts, the administrator is notified that activating, obtaining client license packs, or deactivating will initiate communication with Microsoft.

- **Logging**: Terminal Server Licensing logs events in the system log. The events can be viewed through Event Viewer.

- **Encryption**: Terminal Server Licensing uses the HTTP protocol over SSL (Secure Sockets Layer) to communicate on the Internet and the Web.

- **Access**: The Microsoft Clearinghouse is the database Microsoft maintains to activate license servers and to issue client license key packs. Microsoft customer service representatives have access to the licensing information and are able to successfully re-create the information on your Terminal Server License Server if technical problems occur. The information you provide might also be used internally at Microsoft to perform aggregate quality testing of the Terminal Server Licensing program.

- **Privacy statement**: To see the privacy statement for Terminal Server Licensing, from the Terminal Services Licensing Web site, click **Microsoft's Terminal Server Licensing privacy policy** at:

  http://go.microsoft.com/fwlink/?LinkId=46192

- **Transmission protocol and port**: HTTPS over port 443, and remote procedure call (RPC) over port 135.

- **Ability to disable**: Terminal Server Licensing is not installed by default. Once installed, however, it can be disabled by the procedures described later in this section.

**External Connector license**

In place of individual Terminal Server client access licenses (TS CALs), you have the option of purchasing the Terminal Server External Connector license. This license enables external users to access a company's terminal servers without the need to purchase individual TS CALS for them or their devices. Generally speaking, an external user is a person who is not an employee or similar personnel of a company or its affiliates, and who is not a customer of hosted services (for specific information, see the External Connector license).

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running servers that communicate with the Internet. For more information about security and the Internet, see the introduction

to this white paper, or see the Microsoft TechNet Web site at:

http://www.microsoft.com/technet/security/

**Controlling Terminal Server Licensing to Limit the Flow of Information to and from the Internet**
You can control the communication that occurs between the Terminal Server Licensing component and sites on the Internet in the following ways:

- **Install the Terminal Server Licensing component on selected servers only**. This follows the basic principle of stopping unnecessary services and keeping computers (especially servers) free of unnecessary software. For information about choosing which computer or computers on which to install Terminal Server Licensing, see the *Microsoft Windows Server 2003 Deployment Kit*, which is available on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46078

  When searching in documentation about Terminal Server Licensing, use the phrases "domain license server" and "enterprise license server" to discover information about these two options.

- **Review the method you want to use for activating Terminal Server Licensing before starting the licensing wizard**. Automatic activation is the fastest method, but if you prefer, you can activate by the other methods mentioned previously (connecting to a Web site from a computer other than the one where Terminal Server Licensing is installed, or activating by phone).

  You are required to activate a license server before it can issue licenses to Terminal Server clients. You are required to activate a license server only once. When you activate the license server, Microsoft provides the server with a limited-use digital certificate that validates server ownership and identity. Microsoft uses the X.509 industry standard certificate for this purpose. Using this certificate, a license server can make subsequent transactions with Microsoft and receive client license key packs.

**Procedures for Configuration of Terminal Server Licensing**
Terminal Server Licensing servers can be configured in several ways as described previously. This subsection provides procedures for:

- Installing and uninstalling Terminal Server Licensing
- Activating Terminal Server Licensing
- Deactivating Terminal Server Licensing
- Viewing Help for Terminal Server and Terminal Server Licensing

**To install or uninstall Terminal Server Licensing**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components** (on the left).
4. Select **Terminal Server Licensing**, and do one of the following:

   - If Terminal Server Licensing is installed and you want to remove it, clear the check box for **Terminal Server Licensing** and complete the licensing wizard.
   - If Terminal Server Licensing is not installed and you want to add it, select the check box for **Terminal Server Licensing** and then click **Next**.

5. In Terminal Server Licensing setup, do one of the following:

   - If your network includes several domains, click **Your entire enterprise**, and then provide the database location. An enterprise license server can serve terminal servers on any Windows Server 2003 family domain.

   - If you want to maintain a separate license server for each domain, or if your network includes workgroups or Windows NT 4.0 domains, click **Your domain or workgroup**, and then provide the database location.

**To activate a license server**

As mentioned previously in this section, you must activate a Terminal Server License Server before it can issue licenses to Terminal Server clients. Use Terminal Server Licensing to activate a Terminal Server License Server through the Microsoft Clearinghouse.

You can find procedure checklists along with complete instructions for configuring and activating Terminal Server License Servers in Help and Support Center. For more detailed instructions, see the procedure below titled, "To view Help for Terminal Server and Terminal Server Licensing" and "To view Help topics specific to Terminal Server Licensing."

**Deactivating a license server**

If Terminal Server and Terminal Server Licensing are already installed and you want to deactivate the license server, use one of the following procedures. You might need to deactivate a license server when the certificate of the server has expired, when the server becomes corrupted, or when the server is being redeployed. Note that when a license server's registration has expired, you are prompted to reactive the license server (not deactivate it). When you deactivate a license server, you will not be able to license additional clients from this server until the license server is reactivated.

**To deactivate a license using the automatic method**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**.

3. In the console tree, right-click the license server you want to deactivate, point to **Advanced**, and then click **Deactivate Server**. The licensing wizard starts.

4. In **Information Needed**, confirm that your name, phone number (optional), and e-mail address (required if you are using the Internet method) are correct, and then click **Next**.

5. Your request to deactivate the license server is sent to Microsoft where it is processed.

| **Note** |
| --- |
| The information sent to the Microsoft Clearinghouse during deactivation is the same information sent during activation, which includes company name, first and last name of the user, license server name, and license server ID. |

6. Click **Finish**.

**To deactivate a license using the telephone method**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**.

3. In the console tree, right-click the Terminal Server License Server you want to deactivate, point to **Advanced**, and then click **Deactivate Server**. The Terminal Server License Server Wizard starts.

4. Select **Telephone** and then click **Next**.

5. Specify your location and then click **Next**.

6. Call the telephone number displayed in the wizard, and give the customer support representative the product ID that is displayed below the telephone number.

7. Type the 35-digit confirmation code provided by the customer support representative in the boxes in the wizard, and then click **Next**.

8. Click **Finish**.

| **Note** |
| --- |
| Deactivating a license server does not remove the component. As mentioned previously, follow the basic principle of stopping unnecessary services and keeping computers (especially servers) free from unnecessary software by removing the Terminal Server Licensing component if you no longer plan to use it as a license server.<br><br>You cannot deactivate a license server using either the fax or Internet connection methods. |

**To view Help for Terminal Server and Terminal Server Licensing**

1. Click **Start** and then click **Help and Support**.

2. Click **Software Deployment** and then expand **Terminal Services**.

**To view Help topics specific to Terminal Server Licensing**

1. Click **Start** and then click **Help and Support**.

2. Click **Software Deployment** and then expand **Terminal Services**.

3. Expand **Terminal Server** and then expand **Checklists: Setting up Terminal Server**.

   When you are finished viewing appropriate checklists, collapse the items you were viewing.

4. Expand **Terminal Server Licensing**.

5. To get information about the following procedures, expand **How To**:

   - Install Terminal Server Licensing
   - Activate Terminal Server Licensing server
   - Install client license key packs
   - Deactivate a Terminal Server Licensing server
   - Reactivate a Terminal Server Licensing server
   - Repeat the installation of a client license key pack
   - Connect to a Terminal Server Licensing server
   - Change Terminal Server Licensing server properties

**Related Links**

- For details about new licensing options for the Windows Server 2003 family, as well as information about the licenses

necessary with Terminal Server, see the Windows Server 2003 licensing overview at:

http://go.microsoft.com/fwlink/?LinkId=46007

- For descriptions of other ways you can use Terminal Server, see the Terminal Server overviews on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=24427

- For information about controlling the issuance of Terminal Server licenses, see the topic, "To control the issuance of Terminal Server licenses," in Help and Support Center. For details about viewing Help and Support Center topics, see "To view Help for Terminal Server and Terminal Server Licensing," earlier in this section.

# Windows Error Reporting and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Windows Error Reporting
- How Windows Error Reporting communicates with sites on the Internet
- How to control Windows Error Reporting to prevent the flow of information to and from the Internet

**Benefits and Purposes of Windows Error Reporting**
The Windows Error Reporting feature in Microsoft Windows Server 2003 family operating systems provides a service which allows Microsoft to track and address errors relating to the operating system, Windows components, and applications. This service, called the Error Reporting service, gives administrators and users with administrative credentials the opportunity to send data about errors to Microsoft and to receive information about them. Moreover, developers can use the Error Reporting service as a problem-solving tool to address customer problems in a timely manner and to improve the quality of Microsoft products.

In addition to having users or administrators send information to Microsoft, in some cases Microsoft may provide information, such as a way to work around a problem or a link to a Web site for updated drivers, patches, or Microsoft Knowledge Base articles.

**Overview: Using Windows Error Reporting in a Managed Environment**
In Windows Server 2003 family operating systems, error reporting is enabled by default and you can report system and application errors to Microsoft if you choose to. When an error occurs, a dialog box is displayed, giving you the option to report the problem. If you choose to report the problem, technical information about it is collected and then sent to Microsoft over the Internet. No information is sent unless you confirm that the error report be sent to Microsoft.

On Windows Server 2003 family operating systems you can configure or disable error reporting through the Control Panel\System\Advanced tab. You can configure error reporting to send specified information such as system errors only, unplanned shutdowns, or errors for Windows components, such as Windows Explorer, Paint, or Microsoft Internet Explorer. You can also send information for applications, such as Microsoft Word. An operating system error causes the computer to display a Stop error screen with error values. An application or component error causes the application or component to stop working.

The default settings for the Windows Server 2003 family are:

- Enable error reporting for the operating system, unplanned computer shutdowns, and applications.

  For application errors, you can configure error reporting in one of two ways: either have the error reporting dialog box appear as soon as an error occurs for any user, or do not have the dialog box appear until the next time an administrator logs on.

  Windows treats operating system errors and unplanned shutdowns differently from the way it treats application errors. If an operating system error or unplanned shutdown occurs, Windows writes the error information to a log file. The next time an administrator logs on, the error reporting dialog box prompts them to report the error.

- Force queue mode for application errors.

  The queued mode displays the last ten errors the next time the administrator logs on to the computer. Each error is displayed in its own window so the administrator can choose the errors to report to Microsoft. In this mode errors are displayed only to an administrator; if users logged on to the server they would not see the errors.

Since error reporting is a valuable service, we do not recommend that IT administrators disable it, but that they control what

information is reported and where it is sent. For an organization where privacy is a concern, we recommend that the IT department review and filter error reports before they are sent to Microsoft. The best method to use to prevent the automatic flow of error reporting information to and from the Internet is to redirect error reports to a server on your intranet by using Group Policy and to set up Corporate Error Reporting (CER). You can configure error reporting to control various aspects of how errors are reported.

IT administrators can use the Corporate Error Reporting tool to manage error reports that have been redirected to a network server. You use the tool to review the redirected error reports and then filter the reports that are sent to Microsoft based on your policies and the data contained within the error report. The tool is also useful for determining the types of problems users are experiencing most often.

If you have not yet deployed the operating system, you can use unattended installation files to configure error reporting in the same way as in Group Policy. If it is necessary in your organization to completely disable Windows Error Reporting you can do so with the unattended installation file or with Group Policy. For more information about these methods, see "Controlling Error Reporting to Prevent the Flow of Information to and from the Internet," later in this section.

**How Windows Error Reporting Communicates with Sites on the Internet**

The data that Microsoft collects is used strictly for the purpose of tracking down and solving problems that users or administrators are experiencing. The information is stored in a secure database with limited access. This subsection describes various aspects of the data that is sent to and from the Internet during error reporting, and how the exchange of information takes place.

- **Specific information sent or received**: Microsoft collects various types of information related to two types of errors, user mode or application errors, and kernel mode or operating system failures. Some information that uniquely identifies the user might inadvertently be collected as part of the crash report. This information, if present, is never used to contact a user. The specific data collected is described later in this subsection. Also, Microsoft may send information about a problem, including links to Web sites.

- **Default and recommended settings**: Error reporting for application and system errors is enabled by default. For more information about recommended settings, see "Controlling Error Reporting to Prevent the Flow of Information to and from the Internet," later in this section.

- **Triggers**: The opportunity to send an error report is triggered by application or system errors.

- **User notification**: A dialog box appears notifying users that an error has occurred and asks if they want to send an error report to Microsoft. Users can review the data that will be sent.

- **Logging**: Descriptions of system and application errors are recorded in the event log.

- **Encryption**: All data that could include personally identifiable information is encrypted (HTTPS) during transmission. The "crash signature," which includes such information as the application name and version, module name and version, and offset (location) is not encrypted.

- **Access**: Microsoft employees and contingent staff who have submitted a business justification for reviewing the information are granted access to the data.

- **Privacy statement**: The privacy statement for Microsoft Error Reporting is located at the following Web site:

  http://go.microsoft.com/fwlink/?LinkId=825

  Details related to privacy of data are presented in "Types of data collected," later in this section.

- **Transmission protocol and port**: The transmission protocol is HTTP and the ports are HTTP 80 and HTTPS 443.

- **Ability to disable**: The feature can be disabled through Group Policy or by administrators on individual servers.

**Types of errors reported**

There are two types of errors that are reported, user mode and kernel mode.

**User mode reporting**

When a user mode error occurs, such as an application error, the Error Reporting service does the following:

- Displays an alert stating that the operating system detected a problem.

  Users can choose to report the problem or not. If they do report it, they will see that the information is being sent to Microsoft.

- Sends a problem report to Microsoft.

  Users may then be queried for additional computer information and again may choose to send it or not. If they choose to do so, the Error Reporting service sends the error report to Microsoft. Users might be prompted to provide additional information to complete the error report. When the process is complete, users have the option of selecting More Information, which directs them to updated drivers, patches, or Microsoft Knowledge Base articles.

If the error report indicates that one or more non-Microsoft products were involved in causing the problem, Microsoft may send the report to the respective companies. Qualified software or hardware developers (employed by Microsoft or one of its partners) will analyze the fault data and try to identify and correct the problem.

**Kernel mode reporting**
When a kernel mode or system error occurs, Windows displays a Stop message and writes diagnostic information to a memory dump file. When you restart your computer using normal mode or Safe Mode (with networking) and log on to Windows, the Error Reporting service gathers information about the problem and displays a dialog box that gives you the option of sending a report to Microsoft.

**Types of data collected**
The Error Reporting service collects Internet Protocol (IP) addresses, which are not used to identify users. It does not intentionally collect anyone's name, address, e-mail address, computer name, or any other form of personally identifiable information. It is possible that such information may be captured in memory or in the data collected from open files, but Microsoft does not use it to identify users.

In rare cases, such as problems that are especially difficult to solve, Microsoft may request additional data, including sections of memory (which may include memory shared by any or all applications running at the time the problem occurred), some registry settings, and one or more files from the user's computer. The user's current documents may also be included. When additional data is requested, the user can review the data and choose to send the information or not.

The specific type of data that is collected when application errors or kernel failures occur is as follows.

**Application errors**
If you have an application error the Error Reporting service collects the following information:

- The Digital Product ID, which can be used to identify your license.

- Information regarding the condition of the computer and the application at the time the error occurred. This includes data stored in memory and stacks, information about files in the application's directory, as well as the operating system version and the computer hardware in use. This information is packaged into what is called a "minidump." The minidump contains the following:

  - Exception information: This is information regarding the problem that occurred; it tells Microsoft what kind of instruction the application received that caused it to generate an error.

  - System information: This is data about the kind of CPU (processor) you have and what operating system you are running.

  - A list of all the modules that are currently loaded and their version information.

  - A list of all the threads that are currently running. For each thread, the current context and the whole stack are collected.

  - Global data.

  The minidump data is shown as a hexadecimal representation that the user cannot read.

| Note |
| --- |
| For the exact specification of the minidump format, see the Microsoft Platform SDK, which is available on the Microsoft Developers Network (MSDN) Web site. |

**Windows kernel failures**

Windows kernel fault reports contain information about what your operating system was doing when the problem occurred. These event reports contain the minimum information that can help to identify why the operating system stopped unexpectedly. The report includes:

- The operating system name (for example, Microsoft Windows 2000).
- The operating system version (for example, 5.1.2426 0.0).
- The operating system language as represented by the locale identifier (LCID) (for example, 1033 for United States English). This is a standard international numeric abbreviation.
- The loaded and recently unloaded drivers. These identify the modules used by the kernel when the Stop error occurred, and the modules that were used recently.
- The list of drivers in the Drivers folder on your hard disk, that is, *systemroot*\System32\Drivers.
- The file size, date created, version, manufacturer, and full product name for each driver.
- The number of available processors.
- The amount of random access memory (RAM).
- The time stamp that indicates when the Stop error occurred.
- The messages and parameters that describe the Stop error.
- The processor context for the process that stopped. This includes the processor, hardware state, performance counters, multiprocessor packet information, deferred procedure call information, and interrupts (requests from software or devices for processor attention).
- The process information and kernel context for the halted process. This includes the offset (location) of the directory table and the database that maintains the information about every physical page (block of memory) in the operating system.
- The process information and kernel context for the thread that stopped. This information identifies registers (data-storage blocks of memory in the processor) and interrupt request levels, and includes pointers to data structures for operating system data.
- The kernel-mode call stack for the interrupted thread. This is a data structure that consists of a series of memory locations and includes a pointer to the initial location.

**Controlling Error Reporting to Prevent the Flow of Information to and from the Internet**

To prevent the automatic flow of information to and from the Internet when users and administrators report errors, you can configure error reporting in two ways: while deploying the operating system using answer files with unattended or remote installation, or after deployment using Group Policy. There may be some aspects of error reporting you want to configure using answer files, and others you may want to configure using Group Policy. Review the tables in this subsection to determine the configuration options that will work best for your organization.

**Using unattended installation**

You can configure error reporting by using standard methods for unattended or remote installation. You use the [PCHealth]

section of an answer file to make entries for this feature. The following table describes those entries.

**Entries for configuring error reporting in an answer file (for unattended installation)**

| Entry | Description |
|---|---|
| ER_Display_UI | Specifies whether Setup notifies the user that an error has occurred and shows details about the error. When the entry is **ER_Display_UI = 0**, Setup does not notify the user that an error has occurred. |
| ER_Enable_Applications<br><br>ER_Include_EXE(*n*)<br><br>and<br><br>ER_Exclude_EXE(*n*) | **ER_Enable_Applications = All**<br><br>Reports errors for all applications except for those listed in ER_Exclude_EXE(*n*).<br><br>**ER_Enable_Applications = Listed**<br><br>Reports errors only for those applications listed in ER_Include_EXE(*n*). You can automatically include Microsoft applications by using ER_Include_MSApps.<br><br>**ER_Enable_Applications = None**<br><br>Reports no application errors.<br><br>Examples of entries that list included applications are:<br><br>**ER_Include_EXE1 = iexplore.exe**<br><br>**ER_Include_EXE2 = explorer.exe**<br><br>Examples of entries that list excluded applications are:<br><br>**ER_Exclude_EXE1 = calc.exe**<br><br>**ER_Exclude_EXE2 = notepad.exe** |
| ER_Enable_Kernel Errors | Specifies whether Windows reports errors in the Windows kernel. When the entry is **ER_Enable_Kernel Errors = 0**, Windows does not report errors in the Windows kernel. |
| ER_Enable_Reporting | Specifies whether Windows automatically reports errors. When the entry is **ER_Enable_Reporting = 0**, Windows does not report errors. |
| ER_Enable_Windows_ Components | Specifies whether to report errors in Windows components. When the entry is **ER_Enable_Windows_Components = 0**, Windows does not report errors in Windows components. To exclude individual Windows components, use ER_Exclude_EXE(*n*), as described earlier in this table. |
| ER_Force_Queue_Mode | Specifies whether to send all reports in queue mode. When the entry is **ER_Force_Queue_Mode = 0**, Windows does not send reports in queue mode. |
| ER_Include_MSApps | Specifies whether to track and report errors in Microsoft applications. When the entry is **ER_Include_MSApps = 0**, errors in Microsoft applications are not tracked and reported. |
| ER_Include_Shutdown_ Errs | Specifies whether to report shutdown errors. When the entry is **ER_Include_Shutdown_Errs = 0**, shutdown errors are not reported. |

For complete details about the entries for error reporting, see the resources listed in Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003). Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix).

**Using Group Policy**

To enable Corporate Error Reporting, perform these steps:

- Configure the **Error Reporting** policy settings in Group Policy so that error reports go to a server on your intranet.
- Use the Corporate Error Reporting tool to filter reports.

Enable error reporting through Group Policy so you can override actions users or administrators might take, and so you can redirect error reports to a server on your intranet instead of to the Internet. Once you have initiated Corporate Error Reporting,

you can use this tool to manage error reports.

In addition to the **Error Reporting** policy settings, this subsection also includes a list of the **Advanced Error Reporting** policy settings you may want to use for additional configuration options.

**Using Error Reporting policy settings**

To configure servers for Corporate Error Reporting you need first to enable the **Report Errors** policy setting. Once you enable this policy setting, you can enter a file path to a server on your intranet, limit data that is exchanged on the Internet when errors are reported, control how users and administrators interact with the Error Reporting service, and take other steps to control information.

For details about locating the error reporting policy settings, see "Procedures for Configuring Error Reporting," later in this section. The following table describes the policy settings.

**Group Policy settings for configuring error reporting**

| **Policy setting** | **What it does** | **Configuration options** |
|---|---|---|
| **Report Errors** (enabled) | Errors are reported to Microsoft through the Internet or to a server on your intranet. Enabling **Report Errors** will override any settings made using Control Panel for error reporting. Default values will be used for any error reporting settings that are not configured, even if settings were adjusted through Control Panel. | Can select:<br><br>Do not display links to any Microsoft provided "more information" Web sites<br><br>Do not collect additional files<br><br>Do not collect additional computer data<br><br>Force queue mode for application errors (note that this is the default configuration for servers)<br><br>Can enter:<br><br>Corporate file path<br><br>Text to replace instances of the word "Microsoft" |
| **Report Errors** (disabled) | Users will not be given the option to report errors. If **Display Error Notification** is enabled, users will still get a message indicating that a problem occurred, but they will not have the option to report it. Disabling Report Errors is useful for servers that do not have interactive users. | Not applicable |
| **Report Errors** (not configured) | Users will be able to adjust the setting using Control Panel, which is set to "enable reporting" by default. | Not applicable |
| **Display Error Notification** (enabled) | This setting controls whether a user is given the choice to report an error. When enabled, the user will be notified that an error has occurred and will be given access to details about the error. | Not applicable |
| **Display Error Notification** (disabled) | The user is not given the choice of whether to report the error. If **Report Errors** is enabled, the error will be automatically reported, but the user will not be notified that an error has occurred. Disabling this setting is useful for servers that do not have interactive users. (Default setting for servers.) | Not applicable |

| | | |
|---|---|---|
| **Display Error Notification** (not configured) | The user will be able to adjust the setting through Control Panel, which is set to enable notification by default. | Not applicable |

**Using Advanced Error Reporting policy settings**

When you enable error reporting you can choose to specify the types of errors that are reported. In a highly managed environment administrators might want to do this based on the kinds of information included in the error report (see "Types of data collected," in the previous subsection).

With **Advanced Error Reporting** you can configure the following policy settings:

- **Default application reporting settings**
- **List of applications to always report errors for**
- **List of applications to never report errors for**
- **Report operating system errors**
- **Report unplanned shutdown events**

These policy settings are located in Computer Configuration\Administrative Templates\System\Error Reporting. When you configure these policy settings they will override any adjustments to error reporting administrators might make through Control Panel. You can configure these same policy settings in an answer file for unattended installation.

To find more information about editing Group Policy, see Appendix B: Resources for Learning About Group Policy (Windows Server 2003).

**How controlling error reporting can affect administrators**

What administrators will see on a server when an error occurs depends on how you have configured the **Error Reporting** policy settings. You can have certain administrators sending error reports to your intranet server only and others using CER to filter reports and send selected ones on to Microsoft. On some servers, for example, administrators may see only operating system or unplanned computer shutdown error reports and not application errors. Or, on some servers you might choose not to have error notification on.

**Procedures for Configuring Error Reporting**

This subsection presents the recommended procedure for enabling Corporate Error Reporting by configuring the **Report Errors** policy setting in Group Policy, for IT administrators who want to control the information that goes out to the Internet. This subsection also presents steps for configuring error reporting during unattended installation of the operating system by using an answer file.

Use the following procedure to configure the **Report Errors** policy setting so error reports are sent to a server on your intranet instead of to Microsoft.

**To enable Corporate Error Reporting by using Group Policy**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.
2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, and then click **Error Reporting**.
3. In the details pane, double-click **Display Error Notification**, and then select **Enabled**.
4. Click **Next Setting**, and then under Report Errors, select **Enabled**.
5. In the Corporate upload file path box, enter a UNC (Universal Naming Convention) path (*servername**sharename*).

| Note |
| --- |
| Administrators can then filter the error reports using the CER tool described in the previous subsection, "Controlling Error Reporting to Prevent the Flow of Information to and from the Internet." |

**To configure error reporting during unattended installation by using an answer file**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file. For information about unattended installation, and for details about the entries for error reporting, see the resources listed in Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003). Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix).

2. In the [PCHealth] section of the answer file, create entries according to the table in "Using unattended installation," earlier in this section. For example, to disable error reporting the entry is:

```
[PCHealth]
ER_Enable_Reporting = 0
```

**Related Links**

- For more information about Windows Error Reporting, see the article on the Microsoft Developer Network Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29903

- To obtain the Corporate Error Reporting tool, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29517

- To read the Microsoft privacy statement for error reporting, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=825

# Windows Media Player and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Windows Media Player
- How Windows Media Player communicates with sites on the Internet
- Procedures for configuration of Windows Media Player

**Resources describing Windows Media Player configuration options**
This section of the white paper describes Microsoft Windows Media Player (also called the Player) in the context of use on a server. The Player is not commonly used in a server environment, unless it is being used to test a Windows Media server, so this section of the white paper does not provide extensive information about implementing a specific configuration for the Player. For example, this section does not provide information about allowing or preventing downloads of codecs or other software for the Player. (A codec, short for compressor/decompressor, is software that compresses or decompresses audio or video data.) This section of the white paper also does not describe Windows Media Services, which is described in a separate section of the white paper.

For information about implementing a specific desktop configuration for Windows Media Player 9 Series in your organization, see the white paper titled "Using Windows XP Professional with Service Pack 2 in a Managed Environment: Controlling Communication with the Internet." The white paper includes information about allowing or preventing downloads for the Windows Media Player 9 Series (for example, downloads of codecs or other software for the Player). This white paper can be found on the TechNet Web site at:

http://go.microsoft.com/fwlink/?LinkId=29133

**Benefits and Purposes of Windows Media Player**
Windows Media Player 9 Series is the media player included with products in the Microsoft Windows Server 2003 family. The Player enables playing and organizing digital media files on computers and on the Internet. If you choose to use the Player on a server, you can listen to radio stations, search for and organize digital media files, and (with the right hardware) play CDs and DVDs, create custom CDs, and copy files to a portable device.

Windows Media Player 9 Series is described in this white paper because it can access media files on the Internet. In a highly managed network environment, you might want to control access to the Internet, including access gained through Windows Media Player. The rest of this section describes how the Player communicates with the Internet, and how to control this communication.

**Overview: Using Windows Media Player in a Managed Environment**
Windows Media Player 9 Series in the Windows Server 2003 family is an integral component of the operating system and is installed by default. The Windows Media Player interface is accessed by navigating to Programs\Accessories\Entertainment or All Programs\Accessories\Entertainment from the Start menu. The Player is not commonly used in a server environment, unless it is being used to test a Windows Media server. If you do not want the Player accessible you can use Group Policy to block access to the Player. The procedure for this configuration method is described later in this section.

**How Windows Media Player Communicates with Sites on the Internet**
The Windows Media Player has a number of features that connect to sites on the Internet. When the user selects a feature such as Media Guide, Radio Tuner, Premium Services, or Skin Chooser\More Skins from the Player taskbar, Windows Media Player connects to www.WindowsMedia.com through either a local area network (LAN) or a modem connection.

WindowsMedia.com is a Web site operated by Microsoft and is tightly integrated into Windows Media Player. Media Guide and Radio Tuner are Web pages provided by WindowsMedia.com. All the CD audio data, DVD data, radio presets, and the information in the Info Center View area of the Now Playing feature also come directly from WindowsMedia.com. Other services provided by WindowsMedia.com include the Player updates and download support for codecs, skins, and visualizations. (A codec, short for compressor/decompressor, is software that compresses or decompresses audio or video data.)

To support the playback of secure content, Windows Media Player will also contact:

- Non-Microsoft digital rights management (DRM) license servers
- Microsoft DRM upgrade service

The other common Internet connections that Windows Media Player makes are to Windows Media servers run by content providers.

If you want more information about how Windows Media Player 9 Series communicates with the Internet (from the client or the server) and how to configure the Player for use, see "Related Links," later in this section.

**Procedure for Controlling Windows Media Player**

Windows Media Player can be configured as described previously. This subsection provides a procedure for using Group Policy to prevent users or administrators from starting Windows Media Player.

| **Note** |
| --- |
| For information about implementing a specific configuration for Windows Media Player in your organization, see the sources listed in "Related Links," later in this section. |

**To use Group Policy to prevent users and administrators from starting Windows Media Player**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.
2. Click **User Configuration**, click **Administrative Templates**, and then click **System**.
3. In the details pane, double-click **Don't run specified Windows applications**.
4. Select **Enabled**, click **Show**, click **Add**, and then enter the application executable name, Wmplayer.exe.

**Related Links**

- For more information about deploying and managing Windows Media Player in an enterprise environment, see the following pages on the Windows Media Web site at:

  http://www.microsoft.com/Windows/WindowsMedia/

  http://go.microsoft.com/fwlink/?LinkId=46006

- For information about implementing a specific desktop configuration for Windows Media Player 9 Series in your organization, see the white paper titled “Using Windows XP Professional with Service Pack 2 in a Managed Environment: Controlling Communication with the Internet.” This white paper can be found on the TechNet Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29133

- For information about Windows Media Services, see the corresponding section of this white paper.

# Windows Media Services and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about Windows Media® Services.

| Note |
| --- |
| Windows Media Services 9 Series is included in Windows Server 2003, Standard Edition, the 32-bit version of Windows Server 2003, Enterprise Edition, and the 32-bit version of Windows Server 2003, Datacenter Edition. Windows Media Services is not included in Windows Server 2003, Web Edition, or in the 64-bit versions of the Windows Server 2003 family.<br><br>In the 32-bit versions of Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, Windows Media Services delivers advanced streaming functionality such as multicasting, wireless network support, Internet authentication, server plug-ins, and Cache/Proxy APIs. |

This section includes information about the following:

- The benefits of Windows Media® Services on servers running an operating system in the Microsoft Windows Server 2003 family.
- For servers from which you want to offer content that will be streamed to an intranet or the Internet, the following types of information are provided:
    - Examples of features in Windows Media Services 9 Series that help you control communication to and from a Windows Media server. Windows Media Services 9 Series is the version of Windows Media Services included with the Windows Server 2003 family.
    - References to more detailed information about Windows Media Services, including information about ports and security-related topics.
    - Information about installing the Windows Media Services and the Windows Media Services subcomponents, along with instructions for viewing the Help that comes with Windows Media Services.
- For servers from which you do not want to offer content on an intranet or the Internet, information about excluding or removing Windows Media Services.

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running servers that communicate across the Internet. This section, however, provides overview information as well as suggestions for other sources of information about balancing your organization's requirements for communication across the Internet with your organization's requirements for protection of networked assets.

| Note |
| --- |
| This section of the white paper describes Windows Media Services (the server component), but it does not describe Windows Media Player (the client component) or Internet Information Services (IIS), both of which are involved in carrying out communication of multimedia content across the Internet. For information about these components, see the respective sections of this white paper. |

**Benefits and Purposes of Windows Media Services**

Windows Media Services is an optional component in products in the Windows Server 2003 family. With Windows Media Services, you can create, manage, and deliver Windows Media content over an intranet or the Internet. The clients receiving the content can render it as it is being received, that is, without downloading the content first. Streaming greatly reduces the wait time and storage requirements on the client. It also permits presentations of unlimited length, as well as live broadcasts.

For more information about features in Windows Media Services, see the sources in "Related Documentation and Links," later in this section.

**Examples of Features that Help You Control Communication to and from a Windows Media Server**
This subsection provides brief descriptions of some features in Windows Media Services 9 Series that help you control communication to and from a Windows Media server. These features are integrated with two aspects of basic functionality built into the Windows Server 2003 operating system:

- Authentication
- Authorization

**Authentication**
Authentication is a fundamental aspect of security for a server running Windows Media Services. It confirms the identity of any unicast client trying to access resources on your Windows Media server. Windows Media Services includes authentication plug-ins that you can enable in order to validate user credentials for unicast clients. Authentication plug-ins work together with authorization plug-ins: after users are authenticated, authorization plug-ins control access to unicast content.

Windows Media Services authentication plug-ins fall into the following categories:

- **Anonymous authentication**. These are plug-ins that do not exchange challenge and response information between the server and a player, such as the WMS Anonymous User Authentication plug-in.
- **Network authentication**. These are plug-ins that validate unicast clients based on user logon credentials, such as the WMS Negotiate Authentication plug-in.

When you make decisions about how authentication might affect users, consider the following points:

- For multicast streaming with Windows Media Services 9 Series, clients do not establish a connection, and therefore authentication and authorization do not apply for multicasting. (Multicast streaming is only available if you have the 32-bit version of Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.)
- If a player is connected through HTTP, the player disconnects from the server each time the user stops, pauses, fast-forwards, or rewinds the content. If the user tries to continue receiving the content, the authentication and authorization process occurs again.

For more information about authentication and about the specific authentication plug-ins that you can enable for Windows Media Services, see the list in "Related Documentation and Links," later in this section.

**Authorization**
In order to control access to unicast content on your Windows Media server, unless you identify users only by IP address, you must enable one or more authentication plug-ins and also one or more authorization plug-ins. Authentication plug-ins verify the credentials of unicast clients attempting to connect to the server. Authorization plug-ins verify that the unicast client is allowed to connect to the server. Authorization occurs after authentication is successful.

You can enable authorization plug-ins to control the access to content by authenticated users. If you enable an authorization plug-in, with one exception, you must also enable an authentication plug-in for unicast clients to be able to access your publishing points. The exception is the WMS IP Address Authorization plug-in, which does not require an authentication plug-in to authenticate a unicast client.

Note that for multicast streaming with Windows Media Services 9 Series, clients do not establish a connection, and therefore authentication and authorization do not apply for multicasting. (Multicast streaming is only available if you have the 32-bit version of Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.)

During the authorization process, the server checks the user against the set of access permissions for the resource to which the user is trying to connect.

For more information about authorization, see the list in "Related Documentation and Links," later in this section.

**Firewall Information for Windows Media Services**
This subsection provides information about configuring firewalls (or proxy servers or both) for use with Windows Media

Services. For more information about firewalls, see the sources in "Related Documentation and Links," later in this section.

You can configure each control protocol plug-in (Microsoft Media Server [MMS] protocol, Real Time Streaming Protocol [RTSP], and HTTP) to use a specific port to make firewall configuration easier. If opening ports on your firewall is not possible, Windows Media Services can stream content by using the HTTP protocol over port 80.

| **Note** |
| --- |
| Using HTTP to stream content is disabled by default. |

Windows Media Services was formerly known as Microsoft NetShow Services; some firewalls have a preconfigured NetShow setting, which may work for Windows Media Services.

**Configuring firewalls for unicast streaming**

To configure a firewall for unicast streaming, you must open the ports on the firewall that are required for the connection protocols enabled on your server. If you are streaming content by using either the Microsoft Media Server (MMS) protocol or the Real Time Streaming Protocol (RTSP), you need to support both the User Datagram Protocol (UDP) and Transmission Control Protocol (TCP).

To enable Windows Media Player and other clients to use the HTTP, RTSP, or MMS protocols to connect to a Windows Media server that is behind a firewall, open the ports described in the following table.

**Ports to open when clients are connecting using HTTP, RTSP, or MMS protocols**

| Ports | Description |
| --- | --- |
| In: TCP on ports 80, 554, and 1755 | The Windows Media server uses the TCP In ports to accept an incoming HTTP connection (port 80), RTSP connection (port 554), or MMS connection (port 1755) from Windows Media Player and other clients. |
| In: UDP on ports 1755 and 5005 | The Windows Media server uses UDP In port 1755 to receive resend requests from clients streaming by using MMSU (MMS used with UDP), and UDP In port 5005 to receive resend requests from clients streaming by using RTSPU (RTSP used with UDP). |
| Out: UDP on ports 1024 through 5000. | The Windows Media server uses UDP Out ports 1024 through 5000 to send data to Windows Media Player and other clients. |

To enable a distribution server that is behind a firewall to use the HTTP or RTSP protocols to stream content that originates from a server outside the firewall, open the ports described in the following table.

**Ports to open when a distribution server is behind a firewall and uses HTTP or RTSP to stream content that originates from a server outside the firewall**

| Ports | Description |
| --- | --- |
| In: UDP on ports 1024 through 5000 | The Windows Media server uses UDP In ports 1024 through 5000 to receive data from another server. |
| Out: TCP on ports 80 and 554 | The Windows Media server uses the TCP Out ports to establish an HTTP connection (port 80) or RTSP connection (port 554) to another server or encoder. |
| Out: UDP on port 5005 | When RTSPU distribution is used, the Windows Media server uses UDP Out port 5005 to send resend requests to another server. |

| **Note** |
| --- |

If it is not possible to open all the UDP Out ports on a firewall, UDP packets sent by a Windows Media server may be blocked by the firewall and may not be able to reach the clients on the other side of the firewall. If this is the case, clients may still be able to receive a stream by automatically rolling over to a TCP-based protocol, such as HTTP or RTSPT (RTSP used with TCP). However, the rollover will cause a delay for the client receiving the stream.

If you know you will not be able to support UDP streaming through a firewall, you can decrease the rollover delay by clearing the UDP check box in the Unicast Data Writer plug-in Properties dialog box. For more information, see the Help for Windows Media Services. A procedure for viewing Help is included in "Procedures for Installing, Removing, or Excluding Windows Media Services and Its Subcomponents," later in this section.

**Configuring firewalls for multicast streaming**

If you distribute content using multicast streaming, network traffic is directed through the standard Class D IP addresses (224.0.0.0 through 239.255.255.255). For multicast streaming, you must enable multicast-forwarding on your network. The Internet Group Management Protocol (IGMP), supported by Windows Media Services, ensures that multicast traffic passes through your network only when a player requests a multicast connection, so that enabling multicasting on your routers does not flood your network.

The following firewall configuration enables multicast packets to traverse your firewall:

**IP multicast address range**: 224.0.0.1 through 239.255.255.255

To enable IP multicasting, you must allow packets sent to the standard IP multicast address range to come through your firewall. This IP multicast address range must be enabled on both the player and server sides, as well as on every router in between.

**Enabling access to an encoder outside a firewall**

Encoders use HTTP to connect to a server running Windows Media Services. By default, Windows Media Encoder uses port 8080 for HTTP connections; however, the encoder administrator can specify a different port. If a different port is used, you must specify the same port when you identify the encoder connection URL for the Windows Media server and when opening the port on your firewall.

The following example of a firewall configuration allows a computer running Windows Media Encoder outside a firewall to access a Windows Media server behind a firewall by using HTTP:

**In/Out**: Transmission Control Protocol (TCP) on port 8080.

(The In port is the port through which the server accepts connections. The Out port is the port through which the server sends data to clients.)

**Answer File Entries and Registry Keys for Windows Media Services Subcomponents**

For reference purposes, the following table shows the syntax for answer file entries associated with Windows Media Services in the Windows Server 2003 family. The table also shows the corresponding registry keys. Do not change the registry keys. They are shown for use in a script that could check whether a particular subcomponent is installed on a particular server. A registry key value of 0x00000000 means the subcomponent is not installed, and a value of 0x00000001 means the subcomponent is installed.

**Note**

For more details about answer-file entries related to Windows Media Services subcomponents, including information about dependencies between the entries, see the references listed in Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003). Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix).

**Answer file entries and registry keys for Windows Media Services subcomponents for the Windows Server 2003 family**

| Windows Media Services subcomponent | Answer file entry (in the [Components] section) | Registry key (for use in a script that checks whether a subcomponent is installed): 0x00000000 means it is not installed; 0x00000001 mean it is installed |
|---|---|---|
| Core Windows Media server components | wms = On \| Off | No key available (check for subcomponents by using other keys) |

| | | |
|---|---|---|
| Windows Media Services Administrator for the Web | wms_admin_asp = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\ Windows\CurrentVersion\Setup\OC Manager\ Subcomponents\wms_admin_asp |
| Windows Media Services MMC snap-in | wms_admin_mmc = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\ Windows\CurrentVersion\Setup\OC Manager\ Subcomponents\wms_admin_mmc |
| Multicast and Advertisement Logging Agent components | wms_isapi = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\ Windows\CurrentVersion\Setup\OC Manager\ Subcomponents\wms_isapi |
| Windows Media Services server components | wms_server = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\ Windows\CurrentVersion\Setup\OC Manager\ Subcomponents\wms_server |

**Procedures for Installing, Removing, or Excluding Windows Media Services and Its Subcomponents**
The following procedures explain how to:

- Add or remove Windows Media Services on a computer after setup is complete for a product in the Windows Server 2003 family.
- View the Help that comes with Windows Media Services.
- Prevent the installation of Windows Media Services during unattended installation by using an answer file.
- Specify answer file entries that control whether Windows Media Services subcomponents are included during unattended installation.

**Note**

Windows Media Services is included in Windows Server 2003, Standard Edition, the 32-bit version of Windows Server 2003, Enterprise Edition, and the 32-bit version of Windows Server 2003, Datacenter Edition. Windows Media Services is not included in Windows Server 2003, Web Edition or in the 64-bit versions of the Windows Server 2003 family.

In the 32-bit versions of Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, Windows Media Services delivers advanced streaming functionality, such as multicasting, wireless network support, Internet authentication, server plug-ins, and Cache/Proxy APIs.

**To add or remove Windows Media Services on an individual computer after setup is complete**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components** (on the left).
4. Select **Windows Media Services**.
5. Perform one of the following steps:
    - If Windows Media Services is installed and you want to remove it, clear the check box for **Windows Media**

**Services** and complete the wizard.

- If Windows Media Services is not installed and you want to add it, select the check box for **Windows Media Services** and complete the wizard.

- If you want to view the list of Windows Media Services subcomponents, after selecting **Windows Media Services**, click **Details**.

**To view the Help that comes with Windows Media Services**

1. Make sure that Windows Media Services is installed by using the previous procedure.

2. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

3. Double-click **Administrative Tools** and then click **Windows Media Services**.

4. Click the **Help** menu and then click **Help Topics**.

**To specify answer file entries that control whether Windows Media Services subcomponents are included during unattended installation**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file.

2. In the [Components] section of the answer file, add the appropriate entries listed in the table in "Answer File Entries and Registry Keys for Windows Media Services Subcomponents," earlier in this section. Ensure that the entries specify **Off** for components you do not want to install and **On** for components you want to install.

   If no Windows Media Services subcomponents are listed in an answer file for unattended installation of a Windows Server 2003 family operating system, these components are *not* installed by default.

| **Note** |
| --- |
| For more information about unattended installation, and for details about dependencies between answer-file entries related to Windows Media Services subcomponents, see the references listed in Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003). Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix). |

**To prevent the installation of Windows Media Services during unattended installation by using an answer file**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file. For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003).

2. In the [Components] section of the answer file, ensure that there are no entries for the subcomponents listed in the table in "Answer File Entries and Registry Keys for Windows Media Services Subcomponents," earlier in this section. If you want to list any of these subcomponents, ensure that the entries specify **Off**.

   If no Windows Media Services subcomponents are listed in an answer file for unattended installation of a Windows Server 2003 family operating system, these subcomponents are *not* installed by default.

**Related Documentation and Links**

The following list of resources can help you as you plan or modify your implementation of Windows Media Services and Windows Media Player in your organization:

- For technical information about Windows Media, see "Windows Media Technical Resources for the Enterprise," on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=58341

A variety of technical resources are available on the preceding Web site, including:

- The Windows Media 9 Series Deployment Guide.
- The Enterprise Deployment Pack for Windows Media Player 9 Series, a downloadable packaging tool that simplifies the configuration, deployment, and management of Windows Media Player 9 Series.
- Other technical articles.

- For information about deploying over an intranet, see "Deploying Windows Media 9 Series over an Intranet," on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46008

- For conceptual and how-to information about using Windows Media Services, including information about authentication, authorization, ports, and firewall settings, see the Help that comes with Windows Media Services. For information about installing Windows Media Services and viewing Help, see "Procedures for Installing, Removing, or Excluding Windows Media Services and Its Subcomponents," earlier in this section.

- For general information about features, and information about ports and firewall or proxy settings, search for the latest information on the Windows Media Web site at:

  http://www.microsoft.com/Windows/WindowsMedia/

- For information about Windows Media Services 9 Series, for example, information about upgrading, optimizing, features, and logging, see the following page on the Windows Media Web site:

  http://go.microsoft.com/fwlink/?LinkId=58342

- For documentation on the distribution of content, see the Windows Media Web site at:

  http://go.microsoft.com/fwlink/?LinkId=58343

- For information about redistributing Windows Media 9 Series components with a custom application, see the Microsoft Developer Network Web site at:

  http://go.microsoft.com/fwlink/?LinkId=58345

- For information about Windows Media Services Software Development Kits (SDKs), see the Microsoft Developer Network Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46019

**Printed reference**

- Birney, B., Tricia Gill, and members of the Microsoft Windows Media Team. *Microsoft Windows Media Resource Kit*. Redmond, WA: Microsoft Press, 2003.

  You can read a sample chapter and view information about the *Microsoft Windows Media Resource Kit* on the MS Press Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46020

# Windows Time Service and Internet Communication (Windows Server 2003)

Updated: February 16, 2009

Applies To: Windows Server 2003 with SP1

The following sections provide information about:

- The benefits of Windows Time Service
- How Windows Time Service communicates with sites on the Internet
- How to control Windows Time Service to limit the flow of information to and from the Internet
- How to monitor and troubleshoot Windows Time Service after configuration is complete

**Benefits and Purposes of Windows Time Service**

Many components of operating systems in the Microsoft Windows Server 2003 family rely on accurate and synchronized time to function correctly. For example, without clocks that are synchronized to the correct time on all computers, Windows Server 2003 family authentication might falsely interpret logon requests as intrusion attempts and consequently deny access to users.

With time synchronization, you can correlate events on different computers in an enterprise. With synchronized clocks on all of your computers, you ensure that you can correctly analyze events that happen in sequence on multiple computers. Windows Time Service automatically synchronizes a local computer’s time with other computers on a network to improve security and performance in your organization.

**Overview: Using Windows Time Service in a Managed Environment**

Computers keep the time on their internal clocks, which allows them to perform any function that requires the date or time. For scheduling purposes, however, the clocks must be set to the correct date and time, and they must be synchronized with the other clocks in the network. Without some other method in place, these clocks must be set manually.

With time synchronization, computers set their clocks automatically to match another computer's clock. One computer maintains very accurate time, and then all other computers set their clocks to match that computer. In this way, you can set accurate time on all computers.

Windows Time Service is installed by default on all computers running Windows 2000, Windows XP, and products in the Windows Server 2003 family. Windows Time Service uses Coordinated Universal Time (UTC), which is based on an atomic time scale and is therefore independent of time zone. Time zone information is stored in the computer's registry and is added to the system time just before it is displayed to the user.

Windows Time Service starts automatically on computers that are joined to a domain. (For computers that are not joined to a domain, you can start the time service manually.) In a domain, time synchronization takes place when Windows Time Service turns on during system startup. In the default configuration, the Net Logon service looks for a domain controller that can authenticate and synchronize time with the client. When a domain controller is found, the client sends a request for time and waits for a reply from the domain controller. This communication is an exchange of Network Time Protocol (NTP) packets intended to calculate the time offset and roundtrip delay between the two computers.

**How Windows Time Service Communicates with Sites on the Internet**

In the Windows Server 2003 family, Windows Time Service automatically synchronizes the local computer's time with other computers on the network. The time source for this synchronization varies, depending on whether the computer is joined to a domain in the Active Directory directory service or to a workgroup.

**When a server running a product in the Windows Server 2003 family is part of a workgroup**

In this scenario, the default setting for the time synchronization frequency is set to "once per week," and this default setting uses the time.windows.com site as the trusted time synchronization source. This setting will remain until you manually set it otherwise. One or more computers might be identified as a locally reliable time source by configuring Windows Time Service on those computers to use a known accurate time source, either by using special hardware or a time source available on the Internet. All other workgroup computers can be configured manually to synchronize their time with these local time sources.

**When a server running a product in the Windows Server 2003 family is a member of a domain**
On servers in this scenario, Windows Time Service configures itself automatically, using the Windows Time Service that is available on the domain controllers.

Windows Time Service on a domain controller can be configured as either a reliable or an unreliable time source. Windows Time Service running on a client will attempt to synchronize its time source with servers that are indicated as reliable. Windows Time Service can configure a domain controller within its domain as a reliable time source, and it synchronizes itself periodically with this source. These settings can be modified or overwritten, depending on specific needs.

**When a computer running Windows 2000 or a Windows Server 2003 family operating system is not a member of a domain**
Windows Time Service must be manually started for computers running Windows 2000 that are not members of a domain. For computers running a Windows Server 2003 operating system that are not members of a domain, Windows Time Service is configured by default to synchronize its time source with time.windows.com. Windows Time Service starts automatically for computers running a Windows Server 2003 operating system. These computers use the Network Time Protocol (NTP), while computers running Windows 2000 use the Simple Network Time Protocol (SNTP).

The following list describes various aspects of Windows Time Service data that is sent to and from the Internet and how the exchange of information takes place:

- **Specific information sent or received**: The service sends information in the form of a Network Time Protocol (NTP) packet. For more information about Windows Time Service and NTP packets, see the references listed in "Related Documentation and Links," later in this section.

- **Default and recommended settings**: Computers that are members of an Active Directory domain synchronize time with domain controllers by default. Domain controllers synchronize time with their parent domain controller. By default, the root parent domain controller will not synchronize to a time source. The root parent domain controller can be set to either synchronize to a known and trusted Internet-based time source, or a hardware time device that provides an NTP (Network Time Protocol) or SNTP interface. Its time accuracy can also be maintained manually.

- **Triggers**: Windows Time Service is started when the computer starts. Additionally, the service will continue to synchronize time with the designated network time source and adjust the computer time of the local computer when necessary.

- **User notification**: Notification is not sent to the user.

- **Logging**: Information related to the service is stored in the Windows System event log. The time and network address of the time synchronization source is contained in the Windows event log entries. Additionally, warning or error condition information related to the service is stored in the Windows System event log.

- **Encryption**: Encryption is used in the network time synchronization for domain peers.

- **Information storage**: The service does not store information, as all information that results from the time synchronization process is lost when the time synchronization service request is completed.

- **Port**: NTP and SNTP default to using User Datagram Protocol (UDP) port 123. If this port is not open to the Internet, you cannot synchronize your server to Internet SNTP servers.

- **Communication protocol**: The service on Windows 2000 implements SNTP to communicate with other computers on the network. The service on the Windows Server 2003 family implements NTP to communicate with other computers on the network.

- **Ability to disable**: Disabling the service has no direct effect on applications or other services. Applications and services that depend on time synchronization, such as Kerberos V5 authentication protocol, may fail, or they may yield undesirable results if there is a significant time discrepancy among computers. Because most computers' hardware-based clocks are imprecise, the difference between computer clocks on the network usually increases over time.

**Controlling Windows Time Service to Limit the Flow of Information to and from the Internet**
Group Policy can be used to control Windows Time Service for computers that are running a Windows Server 2003 family operating system to limit the flow of information to and from the Internet.

The synchronization type and NTP time server information can be managed and controlled through Group Policy. The Windows Time Service Group Policy object (GPO) contains configuration settings that specify the synchronization type. When the synchronization type is set to Nt5DS, Windows Time Service synchronizes its time resource with the network domain controller. Alternatively, setting the type attribute to NTP configures Windows Time Service to synchronize with a specified NTP time server. The NTP server is specified by either its Domain Name System (DNS) name or its IP address when you select NTP as the synchronization type.

For more information about configuring Windows Time Service during deployment of products in the Windows Server 2003 family, see *Designing and Deploying Directory and Security Services* and *Designing a Managed Environment* in the *Microsoft Windows Server 2003 Deployment Kit* at:

http://go.microsoft.com/fwlink/?linkid=44319

Clients on a managed network can be configured to synchronize computer clock settings to an NTP server on the network to minimize traffic out to the Internet and to ensure that the clients synchronize to a single reliable time source. If you choose to do so, you can disable time synchronization for both non-domain and domain computers running Windows Server 2003 family operating systems by using Group Policy. The procedures for configuring Windows Time Service are given at the end of this section of the white paper.

**How Windows Time Service can affect users and applications**

Windows components and services depend on time synchronization. For example, the Kerberos V5 authentication protocol on a Windows Server 2003 family domain has a default time synchronization threshold of five minutes. Computers that are more than five minutes out of synchronization on the domain will fail to authenticate using the Kerberos protocol. This time value is also configurable, allowing for smaller thresholds. Failure to authenticate using the Kerberos protocol can prevent logons, access to Web sites, file shares, printers, and other resources or services within a domain.

When the local clock offset has been determined, the following adjustments are made to the time:

- If the local clock time of the client is behind the current time received from the server, Windows Time Service will change the local clock time immediately.
- If the local clock time of the client is more than three minutes ahead of the time on the server, the service will change the local clock time immediately.
- If the local clock time of the client is less than three minutes ahead of the time on the server, the service will quarter or halve the clock frequency for long enough to synchronize the clocks.
- If the client is less than 15 seconds ahead, it will halve the frequency; otherwise, it will quarter the frequency. The amount of time the clock spends running at an unusual frequency depends on the size of the offset that is being corrected.

**Configuration Settings for Windows Time Service**

You can set the global configuration settings for Windows Time Service by using Group Policy. For details about locating Windows Time Service policy settings, see "Procedures for Configuring Windows Time Service," later in this section. The following table describes the policy settings.

These Group Policy settings correspond to registry entries of the same name, which are located in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time.

**Group Policy settings for configuring Windows Time Service**

| Policy setting | Effect of policy setting | Default setting |
|---|---|---|
| **FrequencyCorrectRate** | One over the rate at which the clock is corrected. If this value is too small, the clock will be unstable and will overcorrect. If the value is too large, the clock will take a long time to synchronize. | 4 |
| **HoldPeriod** | The period of time for which spike detection is disabled in order to bring the local clock into synchronization quickly. A spike is a time sample indicating that time is off a number of seconds, usually received after good time samples have been returned consistently. | 5 |

| | | |
|---|---|---|
| **LargePhaseOffset** | A time offset greater than or equal to this value is considered suspicious by the time service. This occurrence might be caused by a noise spike. | 1,280,000 |
| **MaxAllowedPhaseOffset** | The maximum offset (in seconds) for which Windows Time Service attempts to adjust the computer clock by using the clock rate. When the offset exceeds this rate, the service sets the computer clock directly. | 300 |
| **MaxNegPhaseCorrection** | The largest negative time correction in seconds that the service will make. If the service determines that a change larger than this is required, it logs an event instead. | 54,000 (15 hrs) |
| **MaxPosPhaseCorrection** | The largest positive time correction in seconds that the service will make. If the service determines a change larger than this is required, it will log an event instead. | 54,000 (15 hrs) |
| **PhaseCorrectRate** | One over how much of the remaining phase error in order to correct this update interval. | 7 |
| **PollAdjustFactor** | Controls the decision to increase or decrease the poll interval for the system. The larger the value, the smaller the amount of error that causes the poll interval to be decreased. | 5 |
| **SpikeWatchPeriod** | The amount of time that a suspicious offset must persist before it is accepted as correct (in seconds). | 90 |
| **UpdateInterval** | The number of clock ticks between phase correction adjustments. | 100 |
| **AnnounceFlags** | Controls whether this computer is marked as a reliable time server. A computer is not marked as reliable unless it is also marked as a time server. | 6 |
| **EventLogFlags** | Controls the events that the time service logs. | 2 |
| **LocalClockDispersion** | The dispersion (in seconds) that you must assume when the only time source is the built-in complementary metal oxide semiconductor (CMOS) clock. | 10 |
| **MaxPollInterval** | The largest interval, in log2 seconds, allowed for the system polling interval. Note that while a system must poll according to the scheduled interval, a provider can refuse to produce samples when requested to do so. | 15 |
| **MinPollInterval** | The smallest interval, in log2 seconds, allowed for the system polling interval. Note that while a system does not request samples more frequently than this, a provider can produce samples at times other than the scheduled interval. | 4 |

You can set the Windows NTP Client configuration settings for Windows Time Service by using Group Policy. For details about locating Windows Time Service policy settings, see "Procedures for Configuring Windows Time Service," later in this section. The following table describes the policy settings.

These Group Policy settings correspond to the registry entries of the same name located in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Parameters.

**Group Policy settings for configuring the Windows Time Service NTP Client for computers running Windows Server 2003**

| Policy setting | Effect of setting | Default setting |
|---|---|---|
| **NtpServer** | Establishes a space-delimited list of peers from which a computer obtains time stamps, consisting of one or more DNS names or IP addresses per line. Computers connected to a domain must synchronize with a more reliable time source, such as the official U.S. time clock. | time.windows.com |

| | | |
|---|---|---|
| **Type** | Indicates which peers to accept synchronization from:<br><br>**NoSync**. The time service does not synchronize with other sources.<br><br>**NTP**. The time service synchronizes from the servers specified in the NtpServer registry entry.<br><br>**NT5DS**. The time service synchronizes from the domain hierarchy.<br><br>**AllSync**. The time service uses all the available synchronization mechanisms. | Default options<br><br>**NTP**. Use on computers that are not joined to a domain.<br><br>**NT5DS**. Use on computers that are joined to a domain. |
| **ServiceDll** | Provides the directory location of the Windows Time service dynamic-link library (DLL). | C:\WINDOWS\system32\w32time.dll |
| **ServiceMain** | Service Control Manager (SCM) calls this value. | SvchostEntry_W32Time |
| **CrossSiteSyncFlags** | Determines whether the service chooses synchronization partners outside the domain of the computer.<br><br>None 0<br><br>PdcOnly 1<br><br>All 2<br><br>This value is ignored if the NT5DS value is not set. | 2 |
| **ResolvePeerBackoffMinutes** | Specifies the initial interval to wait, in minutes, before attempting to locate a peer to synchronize with. | 15 |
| **ResolvePeerBackoffMaxTimes** | Specifies the maximum number of times to double the wait interval when repeated attempts to locate a peer to synchronize with fail. A value of zero means that the wait interval is always the minimum. | 7 |
| **SpecialPollInterval** | Specifies the special poll interval in seconds for peers that have been configured manually. When a special poll is enabled, Windows Time Service will use this poll interval instead of a dynamic one that is determined by synchronization algorithms built into Windows Time Service. | 3600 |
| **EventLogFlags** | Controls the events that the time service logs. | 0 |

**Note**

Group Policy and Active Directory are tools that are available for controlling and managing computers and services within an enterprise or organization. The full description of the rich feature set and methods for using Group Policy are beyond the intended scope of this document. For other sources of information about Group Policy, see Appendix B: Resources for Learning About Group Policy (Windows Server 2003).

For information about configuring the authoritative time server in a domain, see article 884776, "Configuring the Windows Time service against a large time offset" in the Microsoft Knowledge Base at:

http://go.microsoft.com/fwlink/?LinkId=46021

**Procedures for Configuring Windows Time Service**

The following procedures explain how to set Windows Time Service configuration settings through Group Policy on operating systems in the Windows Server 2003 family to achieve the configurations described in the previous subsections.

**To set Group Policy for Windows Time Service global configuration settings**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn

about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, and then click **Windows Time Service**.

3. In the details pane, double-click **Global Configuration Settings**, and then click **Enabled**.

**To configure the Group Policy setting to prevent your computer from synchronizing its computer clock with other NTP servers**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Windows Time Service**, and then click **Time Providers**.

3. In the details pane, double-click **Enable Windows NTP Client** and then select **Disabled**.

**To configure the Group Policy setting to prevent your computer from synchronizing its computer clock from the domain hierarchy or a manually configured NTP server**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Windows Time Service**, and then click **Time Providers**.

3. In the details pane, double-click **Configure Windows NTP Client**, and then select **Disabled**.

**To configure the Group Policy setting to prevent your computer from servicing time synchronization requests from other computers on the network**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Windows Time Service**, and then click **Time Providers**.

3. In the details pane, double-click **Enable Windows NTP Server**, and then select **Disabled**.

**Starting and stopping Windows Time Service**

By default, Windows Time Service starts automatically at system startup. You can, however, start or stop the service manually by accessing services in Administrative Tools or by using the **net** command.

**To manually start Windows Time Service using the graphical interface**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then double-click **Services**.

3. Select **Windows Time** from the list of services.

4. On the **Action** menu, click **Start** to begin the service.

**To manually stop Windows Time Service using the graphical interface**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Administrative Tools**, and then double-click **Services**.
3. Select **Windows Time** from the list of services.
4. On the **Action** menu, click **Stop** to discontinue the service.

**To manually start Windows Time Service using the net command**

1. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**.
2. Type **net start w32time**, and then press ENTER.

**To manually stop Windows Time Service using the net command**

1. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**.
2. Type **net stop w32time**, and then press ENTER.

**Synchronizing computers with time sources**

Use the following procedures to synchronize the internal time server with an external time source, and to synchronize the client time with a time server.

**To synchronize an internal time server with an external time source**

1. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**.
2. Type the following, where *PeerList* is a comma-separated list of Domain Name System (DNS) names or Internet protocol (IP) addresses of the desired time sources:

   **w32tm /config /syncfromflags:manual /manualpeerlist:***PeerList*

   and then press ENTER.
3. Type **w32tm /config /update**, and then press ENTER.

| **Note** |
| --- |
| The most common use of this procedure is to synchronize the internal network's authoritative time source with precise external time source. This procedure can be run on any computer running Windows 2000, Windows XP, or an operating system in the Windows Server 2003 family.<br><br>If the computer cannot reach the servers, the procedure fails and an entry is written to the Windows System event log. |

**To synchronize the client time with a time server**

1. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**.
2. Type **w32tm /resync**, and then press ENTER.

| **Note** |
| --- |

This procedure only works on computers that are joined to a domain.

The W32tm command-line tool is used for diagnosing problems that can occur with Windows Time Service. If you are going to use the tool on a domain controller, it is necessary to stop the service. Running the tool and Windows Time Service at the same time on a domain controller generates an error because both are attempting to use the same UDP port. When you finish using W32tm command-line tool, the service must be restarted.

**Monitoring and Troubleshooting Windows Time Service**

In many cases problems with Windows Time Service can be attributed to network configuration. If the network is not configured correctly computers might not be able to communicate to send time samples back and forth. Viewing the contents of NTP packets can help you to identify exactly where a packet is blocked on a network. An error associated with Windows Time Service might occur when a computer is unable to synchronize with an authoritative source. You can use the W32tm command-line tool to assist you in troubleshooting this and other types of errors associated with Windows Time Service.

The W32tm command-line tool is the preferred command-line tool for configuring, monitoring, and troubleshooting Windows Time Service. All tasks that can be performed by using the **net** command can be accomplished by using this tool or Group Policy. For more information, look up "W32tm" in the Help and Support Center index.

**Procedure to follow when a computer is unable to synchronize**

A computer running Windows Time Service refuses to synchronize with a time source if the computer's time is more than 15 hours off. Such occurrences are rare, and are often caused by configuration setting errors. For example, if a user sets the date on the computer incorrectly, the time does not synchronize. Under these circumstances, most often the time is off by a day or more. Be sure to check the computer's calendar and ensure that the correct date has been set.

**To resynchronize the client time with a time server**

1. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. Type **w32tm /resync /rediscover**, and then press ENTER.

**Note**

When you run the preceding command, it redetects the network configuration and rediscovers network resources, causing resynchronization. This procedure only works on computers that are joined to a domain. You can then view the event log for more information about why the time service does not synchronize. For more information, look up "Monitoring and controlling services on computers," in Help and Support Center.

The W32tm tool is used for diagnosing problems that can occur with Windows Time Service. If you are going to use the W32tm tool on a domain controller, it is necessary to stop the service. Running W32tm and Windows Time Service at the same time on a domain controller generates an error because both are attempting to use the same UDP port. When you finish using W32tm, the service must be restarted.

**Related Documentation and Links**

For more information about configuring Windows Time Service during deployment of products in the Windows Server 2003 family, see *Designing and Deploying Directory and Security Services* and *Designing a Managed Environment* in the *Microsoft Windows Server 2003 Deployment Kit* at:

http://go.microsoft.com/fwlink/?linkid=44319

For information about configuring the authoritative time server in a domain, see article 884776, "Configuring the Windows Time service against a large time offset" in the Microsoft Knowledge Base at:

http://go.microsoft.com/fwlink/?LinkId=46021

# Windows Update, Automatic Updates, and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Windows Update and Automatic Updates
- How Windows Update and Automatic Updates communicate with sites on the Internet
- How to control Windows Update and Automatic Updates to limit the flow of information to and from the Internet

| **Important** |
| --- |
| This section describes methods for controlling the way the Automatic Updates component interacts with the Windows Update Web site. One basic way, however, of controlling whether a particular person can install software (including software updates) on a particular computer is to control the type of account that the person has. If the account does not allow the person to install software (for example, if the account is a user account) the person will not be able to use Automatic Updates to install software while logged on with that account. |

**Benefits and Purposes of Windows Update and Automatic Updates**

**Windows Update**

Windows Update is an online catalog customized for computers running a product in the Microsoft Windows Server 2003 family that consists of items such as drivers, critical updates, Help files, and Internet products. Windows Update scans the user's computer and provides a tailored selection of updates that apply only to the software and hardware on that specific computer. Windows Update then enables users to choose updates for their computer's operating system and hardware. New content is added to the Windows Update Web site regularly, so users can always get the most recent and secure updates and solutions.

Windows Update contains two key components:

- **Content update**: Content updates occur when the user accesses the Windows Update Web site and selects component updates to download and install. The user is fully aware of downloads to the computer. The Windows Update Web site is located at:

  http://windowsupdate.microsoft.com/

- **Web service control update**: The Windows Update Web service includes an ActiveX Web control program that downloads and installs the content updates. The Windows Update team receives feedback from their customers on how to improve their Web service and the Windows Update service control is changed to reflect that feedback. In order to access the new content and services customers need, the Web controls are updated periodically. This service automatically downloads a new version of the Web control program when the user visits the Windows Update site or when any of the other Windows features calls on the Windows Update control. Just like downloading an ActiveX control, the user may receive a security dialog box that a Web control is attempting to be installed. Users may not receive the dialog box if they have selected to always trust Microsoft as a content provider (using their security settings in Microsoft Internet Explorer). If users do not click Yes on the security dialog box, the control will not be updated and they will not be able to access the Windows Update site.

**Automatic Updates**

This option for updating a computer allows for updates without interrupting the user's Web experience. Automatic Updates is not enabled by default; users are prompted to enable this option following setup. When Automatic Updates is enabled, users do not need to visit special Web pages or remember to periodically check for new updates. An icon appears in the notification area each time new updates are available. Updates can be downloaded in the background with minimal impact on the user's network connections. Once the update is downloaded, operating systems in the Windows Server 2003 family prompt the user

to install it. Users can set Automatic Updates options in one of three ways to control how and when they want the operating system to update their computers. They can:

- Choose to have the operating system send a notification before downloading and installing any updates.
- Choose to have the operating system download and install updates automatically on a schedule that they specify.
- Choose to have the operating system send a notification whenever it finds updates available for their computers; the operating system will then download the updates in the background, enabling users to continue working uninterrupted. After the download is complete, an icon in the notification area will prompt users that the updates are ready to be installed.

Users can choose not to install a specific update that has been downloaded; in that case, the operating system will delete those files from the computer. Users can download those deleted files again by opening **System** in Control Panel, clicking the **Automatic Updates** tab, and then clicking **Declined Updates**. If any of the updates users previously declined can still be applied to their computers, they will appear the next time an operating system in the Windows Server 2003 family notifies those users of available updates.

**Alternatives to Windows Update and Automatic Updates**
For managed environments, there are several alternatives to Windows Update:

- Windows Update Catalog Web site.
- Microsoft Windows Server Update Services (WSUS).
- Distribution software, such as Microsoft Systems Management Server, that can be used to distribute software updates. For more information, see the documentation for your distribution software, and see Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003), especially the "Related Documentation and Links" subsection in that appendix.

**Windows Update Catalog Web Site**
You can deploy updates to Windows in a managed environment without requiring users to connect to the Windows Update Web site by using the Windows Update Catalog site. This site provides a comprehensive catalog of updates that can be distributed over a managed network. It provides a single location for Windows Update content and drivers that display the Designed for Windows logo. Administrators can search the site using keywords or predefined search criteria to select the relevant downloads and then to download the updates to a location on their internal network.

An enhancement in products in the Windows Server 2003 family enables you to select updates that you plan to deploy later, which means you can control how and when the updates are deployed. For additional information, see information about Windows Update on the Microsoft Web site at:

http://windowsupdate.microsoft.com/

**Microsoft Windows Server Update Services (WSUS)**
Microsoft Windows Server Update Services (WSUS) is a version of Windows Update designed for installation inside an organization's firewall. This feature is very useful for organizations that:

- Do not want their systems or users connecting to an external Web site
- Want to first test these updates before deploying them throughout their organization

Microsoft Windows Server Update Services enables administrators to quickly and reliably deploy critical updates to their computers running Windows Server 2003 operating systems.

For more information about Windows Server Update Services and updated versions of Windows Server Update Services, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=29906

**Overview: Using Windows Update and Automatic Updates in a Managed Environment**
Users have control over whether to enable Automatic Updates following setup and they also have direct control over accepting

downloaded files from Windows Update. In a managed environment, however, it is unlikely that users will be allowed unlimited access to install updated drivers and other updated files; this function would normally be controlled in some fashion by the IT department. You can use Group Policy to block users from accessing Windows Update in the user interface or to specify an internal server for Windows Update to use when searching for updates. You can also disable Automatic Updates using Control Panel or Group Policy. Details on the methods and procedures for controlling these features are described in the following subsections.

**How Windows Update and Automatic Updates Communicate with Sites on the Internet**
This subsection summarizes the communication process:

- **Specific information sent or received**: Drivers and replacement files (critical updates, Help files, and Internet products) may be downloaded to the user's computer. The computer is uniquely identified and is logged in the download and installation success report, but the user is not uniquely identified.

- **Data storage and access**: Windows Update tracks the total number of unique computers that visit the Windows Update Web site. The success or failure of downloading and installing updates is also recorded but no personally identifiable information is recorded as part of this. This information is stored on servers at Microsoft with limited access that are located in controlled facilities. No other information collected during a Windows Update session is retained past the end of the session.

  For more information, see "Privacy statement," later in this list.

  | **Note** |
  | --- |
  | If you want to block the use of the Windows Update Web site, you can apply Group Policy settings to specify an internal server for updates and for storing upload statistics. For more information see "Procedures for Disabling Windows Update and Automatic Updates." |

- **Default and recommended settings**: By default, operating systems in the Windows Server 2003 family provide access to the Windows Update Web site. Recommended settings are described in the next subsection, "Controlling Windows Update and Automatic Updates to Limit the Flow of Information to and from the Internet."

- **Triggers**: The user controls whether to run Windows Update. If Automatic Updates is enabled following setup, it is triggered about once per day when there is an Internet connection.

- **User notification**:

  - **Windows Update**: Users are notified when Windows Update downloads files to their computer, and they have control over whether to install those downloads.

  - **Automatic Updates**: Administrators can specify one of two notification settings for Automatic Updates:

    Notify users before downloading and installing any updates.

    Download the updates automatically and notify users when they are ready to be installed.

    | **Note** |
    | --- |
    | Administrators can also specify that updates be automatically downloaded and installed following a set schedule without user notification. For more information about these settings, click the **Learn more about automatic updating** link on the Automatic Updates dialog box. |

- **Logging**: Automatic Updates logs events to the event log.

- **Encryption**: The data is transferred using HTTPS. The data packages downloaded to the user's system by Microsoft are digitally signed.

- **Privacy statement**: To view the privacy statement for Windows Update, see the Windows Update Web site, and click **Microsoft Update Privacy Statement**. The Windows Update Web site is located at:

  http://windowsupdate.microsoft.com/

  Automatic Updates is covered by the same privacy statement that covers Windows Update.

- **Transmission protocols and ports**: The transmission protocols and ports used are HTTP 80 and HTTPS 443.

- **Ability to disable**: You can use Group Policy to remove user access to Windows Update in the user interface. You can use Group Policy to specify an internal server to use for Windows Update and block it from searching the Windows Update Web site. You can disable Automatic Updates using Control Panel tools or Group Policy. Procedures for these methods are given at the end of this section.

**Controlling Windows Update and Automatic Updates to Limit the Flow of Information to and from the Internet**

The recommended methods for controlling Windows Update and Automatic Updates or both are as follows:

- You can use Group Policy settings to control Windows Update and Automatic Updates by removing end user access to Windows Update.

- You can block Windows Update from searching the Windows Update Web site by using Group Policy settings to specify an internal server for updates.

- You can use Control Panel or Group Policy settings to selectively disable Automatic Updates.

- You can control both Windows Update and Automatic Updates by blocking HTTP port 80 or HTTPS port 443 or both at the firewall.

See the following table for more information about the configuration options.

**Configuration settings for Windows Update and Automatic Updates**

| **Automatic Updates: Configuration tool** | **Setting** | **Result** |
|---|---|---|
| Control Panel (Automatic Updates tool) | In the Automatic Updates dialog box, clear **Keep my computer up to date**. | Disables Automatic Updates. |
| Group Policy | Disable the **Configure Automatic Updates** policy setting in the Wuau.adm Group Policy template. For more information, see "Procedures for Disabling Windows Update and Automatic Updates," later in this section. | Disables Automatic Updates. |
| **Windows Update and Automatic Updates: Configuration tool** | **Setting** | **Result** |
| Firewall | Block HTTP port 80 or HTTPS port 443 or both. | Blocks Windows Update and Automatic Updates. |

| | | |
|---|---|---|
| Group Policy | Enable the **Remove access to use all Windows Update features** policy setting in the Wuau.adm Group Policy template. For more information, see "Procedures for Disabling Windows Update and Automatic Updates," later in this section. | Blocks the user from accessing Windows Update in operating systems in the Windows Server 2003 family user interface. Also blocks Automatic Updates. |
| Group Policy | Enable the **Specify intranet Microsoft update service location** policy setting in the Wuau.adm Group Policy template. For more information, see "Procedures for Disabling Windows Update and Automatic Updates," later in this section. | Blocks Windows Updates from searching for updates on the http://windowsupdate.microsoft.com Web site. Instead, Windows Update searches for updates on a specified internal server. |

**How controlling Windows Update and Automatic Updates can affect users and applications**

When you remove user access to Windows Update, Windows will still search for and download updates to the local computer. Users will not, however, be prompted to install downloaded updates, nor will they be able to access the Windows Update Web site from any of the following locations:

- The Windows Update option on the Start menu
- The Tools menu in Microsoft Internet Explorer
- The Windows Update button in Add New Programs (Add New Programs is in Control Panel under Add or Remove Programs)

Removing user access to Windows Update also disables Automatic Updates; that is, the user for which this policy setting is enabled will neither be notified about nor will receive critical updates from Windows Update. Removing user access to Windows Update is a user-based, not system-based, policy; other users on the same computer will still receive critical updates unless this policy setting is also enabled for those users individually.

Removing end user access to Windows Update also prevents Device Manager from automatically installing driver updates from the Windows Update Web site. For more information about controlling Device Manager, see the section of this white paper titled "Device Manager."

Blocking Windows Update and Automatic Updates will not block applications from running.

The Windows Update site is located at:

http://windowsupdate.microsoft.com/

**Procedures for Disabling Windows Update and Automatic Updates**

This subsection provides procedures for the following configurations:

- Specifying that Windows Update search an internal server, rather than the Windows Update Web site, for updates.
- Removing user access to Windows Update by using Group Policy, which will also block Automatic Updates.
- Disabling and configuring Automatic Updates by using Group Policy.
- Disabling Automatic Updates by using Control Panel tools.

**To specify an internal server for Windows Update using Group Policy**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.
2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Windows Update**.
3. In the details pane, double-click **Specify intranet Microsoft update service location**, supply the name of the internal

server to function as the update server, and supply the name of the server to store upload statistics.

4. Click **Enabled**.

| **Important** |
| --- |
| The upgrade server and the server you specify to store upload statistics can be the same server. |

**To remove user access to Windows Update using Group Policy**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **User Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Windows Update**.

3. In the details pane, double-click **Remove access to use all Windows Update features**, and then click **Enabled**.

| **Important** |
| --- |
| Removing user access to Windows Update also disables Automatic Updates. |

**To disable Automatic Updates using Group Policy**

1. Use the resources described in Appendix B: Resources for Learning About Group Policy (Windows Server 2003) to learn about Group Policy and the Group Policy Management Console. Apply Group Policy objects (GPOs) to an organizational unit, a domain, or a site, as appropriate for your situation.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Windows Update**.

3. In the details pane, double-click **Configure Automatic Updates**, and then click **Disabled**.

**To disable Automatic Updates using Control Panel tools**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Click **System**, and then click the **Automatic Updates** tab.

3. In the **Automatic Updates** dialog box, clear the **Keep my computer up to date** check box.

# Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The following appendix provides:

- An overview of automated installation and deployment
- Procedures and resources for obtaining more information about automated installation and deployment

**Overview: Automated Installation and Deployment**
In the enterprise environment, it is often not cost-effective to install products in the Microsoft Windows Server 2003 family using the standard interactive setup on each computer. To greatly lower the total cost of ownership (TCO) and ensure configuration uniformity, you can perform an automated installation of any product in the Windows Server 2003 family on multiple computers. By using an automated installation method, you can ensure that certain components and applications are not available on your organization's computers, or that certain components and applications are preconfigured in such a way that helps prevent unwanted communication over the Internet.

**Methods for automating the setup process**
There are several options for automating the setup process. Any or all of the following tools can help ensure that all of your servers and clients are configured to appropriately limit communication over the Internet:

- Unattended setup using Setup (Winnt32.exe)

  Unattended setup enables you to simplify the process of setting up the operating system on multiple computers by running Setup unattended. To do this, you can create and use an answer file, which is a customized script that answers Setup questions automatically. Then you can run Setup (Winnt32.exe) from the command line with the appropriate options for invoking unattended setup.

  Using Winnt32.exe, you can upgrade your previous version of the operating system using all user settings from the previous installation, or you can perform a fresh installation using the answer file that provides Setup with your custom specifications. The latter method is most likely the best option to limit component communication over the Internet, provided you use an appropriate answer file. Details about specific answer file entries are included in the appropriate component sections of this white paper.

- Remote Installation Services (RIS)

  You can use RIS to create installation images of operating systems or of complete computer configurations, including desktop settings and applications. You can then make these installation images available to users at client computers. You can also specify which RIS server will provide installations to a given client computer, or you can allow any RIS server to provide the installation.

- Image-based installation using the System Preparation (Sysprep) tool

  Image-based installation is also a good choice if you need to install an identical configuration on multiple computers. On a master computer, you install the operating system and any applications that you want installed on all of the target computers. Then you run Sysprep and a disk imaging utility. Sysprep prepares the hard disk on the master computer so that the disk imaging utility can transfer an image of the hard disk to the other computers. This method decreases deployment time dramatically compared to standard or scripted installations. You can customize the images so that only the files required for a specific configuration appear on the image, such as additional Plug and Play drivers that might be needed on various systems. The image can also be copied to a CD and distributed to remote sites that have slow links.

- System management software, such as Microsoft Systems Management Server (SMS)

  This type of software assists with the many tasks that are involved when you apply automated procedures to multiple

servers and client computers throughout your organization. These tasks include:

- Selecting computers that are equipped for the operating system and that you are ready to support.
- Distributing the operating system source files to all sites, including remote sites and sites without technical support staff.
- Monitoring the distribution to all sites.
- Providing the appropriate user rights to do the upgrade.
- Automatically initiating the installation of the software package with the possibility of having the user control the timing.
- Resolving problems related to the distributions or installations.
- Reporting on the rate and success of deployment.

Using system management software helps to further ensure that all computers within your organization have received the standardized operating system configuration that helps prevent unwanted communication over the Internet.

**Using scripts for configuring computers**
In addition to the automated installation methods described here, another common method of controlling Internet connections is to use a script to configure Group Policy on each client computer. The script can be sent to each client computer using a tool such as Microsoft Systems Management Server (SMS) and can run remotely using Windows Script Host. Alternatively, Group Policy can be applied to a domain, site, or organizational unit. The policy settings would then automatically be applied to every computer in the domain, site, or organizational unit the first time the computer starts after the operating system is installed. For more information about scripts and Group Policy, see "Related Documentation and Links," at the end of this section.

You can also use scripts to monitor activity on client computers and to take appropriate action if certain restricted activities occur. For example, if a user were to start an unauthorized application, a script could be used to detect this and to immediately stop that application. Similarly, scripts can be used to monitor the setup of each computer in order to, for example, determine which applications are installed and which folders are being shared. Configuring these scripts is beyond the scope of this document; however, you can refer to "Related Documentation and Links," at the end of this section for more information.

**Procedures for Accessing Additional Information About Other Automated Setup Tools**
**Accessing the Windows Server 2003 Help documentation**
Products in the Windows Server 2003 family have Help documentation describing unattended installation, RIS, and image-based installation. You can view this documentation from any computer that has Internet access (regardless of the operating system running on that computer), or from any computer running a product in the Windows Server 2003 family. The following procedure gives the details.

**To access the Help documentation for a computer running a product in the Windows Server 2003 family**

1. Open Help and Support by doing one of the following:

   - On any computer running a product in the Windows Server 2003 family, click **Start**, and then click **Help and Support.**
   - View Help on the Web at:

     http://go.microsoft.com/fwlink/?LinkId=29881

     On this site, click the link for the appropriate product. You can then view the Help documentation by expanding the appropriate product listed under **Product Documentation** in the Help navigation pane, on the left.

2. Locate the specific topics as follows:

- For unattended installation: Navigate to Getting Started\Installing and upgrading the operating system\Concepts\Planning for unattended Setup

- For RIS: Navigate to Software Deployment\Remote Installation Services

- For Winnt32.exe: Navigate to Administration and scripting tools\Command-line reference\Command-line reference A-Z\Winnt32

**Related Documentation and Links**

You can also find additional information about all of the topics described earlier in this appendix in a variety of other locations:

- On the CD for any product in the Windows Server 2003 family, you can find additional information about unattended installations in Deploy.chm in Support\Tools\deploy.cab.

- For extensive information about unattended setup, including information about Systems Management Server, see the *Microsoft Windows Server 2003 Deployment Kit*, especially the book titled *Automating and Customizing Installations*. The *Microsoft Windows Server 2003 Deployment Kit* is available on the Web at:

  http://go.microsoft.com/fwlink/?LinkId=29887

- For general information about Group Policy, see the resources listed in Appendix B: Resources for Learning About Group Policy (Windows Server 2003).

- The Help documentation for the Windows Server 2003 family of products included on the CD and on the Web includes information about Windows Script Host. You can find the documentation on the Web at:

  http://go.microsoft.com/fwlink/?LinkId=29221

# Appendix B: Resources for Learning About Group Policy (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This appendix provides a list of resources for learning about Group Policy, Group Policy settings, and the Group Policy Management Console (GPMC).

**Resources for Learning about Group Policy**
To learn about using Group Policy and the Group Policy Management Console (GPMC), see the following resources.

- Windows Server 2003 Group Policy page on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46022

- The Group Policy Technology Center Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29866

- *Designing a Managed Environment* in the *Microsoft Windows Server 2003 Deployment Kit* at:

  http://go.microsoft.com/fwlink/?linkid=44319

- The Group Policy Collection of the *Windows Server 2003 Technical Reference* at:

  http://go.microsoft.com/fwlink/?LinkId=29907

- Article 816662, "Recommendations for Managing Group Policy Administrative Template (.adm) Files" in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?LinkId=29128

To learn about Group Policy settings that can be used with Windows Server 2003, see the following resources.

- The Group Policy Settings Reference on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29911

**Installing Group Policy Management Console (GPMC)**
Group Policy Management Console (GPMC) unifies the management of all aspects of Group Policy. Using GPMC, you can manage Group Policy objects (GPOs), Windows Management Instrumentation (WMI) filters, and permissions on your network.

GPMC is available for download from the Microsoft Web site. For instructions for downloading and installing GPMC, see the Microsoft Windows Server 2003 Web site at:

http://go.microsoft.com/fwlink/?linkid=29909

To run GPMC, click **Start**, point to **Administrative Tools**, and then click **Group Policy Management**.

# Appendix C: Message Queuing and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The purposes of Message Queuing
- How Message Queuing communicates with sites on the Internet
- Controlling Message Queuing to prevent the flow of information to and from the Internet

**Purposes of Message Queuing 3.0**
Message Queuing (MSMQ) is one of the optional components in the Microsoft Windows Server 2003 family operating systems. Information about Message Queuing is presented here so you as an IT administrator can be aware of the potential capabilities for use within your organization's network, if you determine that you will be installing Message Queuing.

Message Queuing 3.0, the version in products in the Microsoft Windows Server 2003 family, is a message queuing service that enables applications on different systems to communicate with each other. Message Queuing implements asynchronous communication by enabling applications to send messages to and receive messages from other applications that may be running on the same computer or on different computers connected by a network. With Message Queuing, these applications can communicate across the Internet and other heterogeneous networks, and with computers that might be temporarily offline.

**Overview: Using Message Queuing in a Managed Environment**
Applications can use Message Queuing to send messages and to continue running regardless of whether the receiving application is running or reachable over the network. To create a message, an application specifies the message fields or properties and supplies the field values. The application then issues an MSMQ application programming interface (API) call to send the message. The MSMQ queue manager (the designated server) transmits the message to the destination message queue. If the destination location is not connected to the network when the message is sent, the message queuing system stores the message at an interim location. The system forwards the message automatically when a connection is established.

To receive a message, an application issues an API call that reads the message from the queue. When an application receives a request message, it processes the request by reading the contents of the message and acting accordingly. If required, the receiving application can send a response message back to the original requester. While in transit between senders and receivers, MSMQ keeps messages in holding areas called queues. MSMQ queues protect messages from being lost in transit and provide a place to retrieve messages when the receivers are ready to receive them.

Message Queuing 3.0 provides support for sending messages over the Internet. This support includes referencing queues by URLs, sending and receiving XML-based messages that are formatted using Simple Object Access Protocol (SOAP), and providing security for HTTP messaging. MSMQ also provides support for Web farms and firewalls. For more information about sending messages over the Internet, refer to "Related Links," later in this section.

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running applications that communicate across the Internet. This section, however, provides overview information as well as suggestions for other sources of information about balancing your organization's requirements for running Internet applications with your organization's requirements for protection of networked assets.

**Controlling Message Queuing to Prevent the Flow of Information to and from the Internet**
**Installing and uninstalling Message Queuing**
Message Queuing is not installed by default, however, if your organization has determined that MSMQ is an essential component of the business process it can be installed as described in this subsection. Although MSMQ is highly configurable, it is beyond the scope of this white paper to describe all the configuration options available to you. If you will be using MSMQ, you can refer to the "Related Links" subsection following these procedures for more detailed information about deploying Message Queuing.

**To install or uninstall Message Queuing**

1. Click **Start**, point to **Control Panel**, and click **Add or Remove Programs**.

2. Click **Add/Remove Windows Components** (on the left).

3. In Windows Components, select **Application Server**, and then click **Details.**

4. Click **Message Queuing** and do one of the following:

    - If Message Queuing is installed and you want to remove it, clear the check box for **Message Queuing** and complete the Message Queuing wizard.

    - If Message Queuing is not installed and you want to add it, select the check box for **Message Queuing**.

5. If you are installing Message Queuing, select **Details**, select one or more of the options listed, and then click **Next**.

**Accessing the operating system Help documentation for Message Queuing**
The operating system has Help documentation describing the use of Message Queuing. You can view this documentation from any computer that has Internet access (regardless of the operating system running on that computer), or from any server running a product in the Windows Server 2003 family.

**To access the Help documentation for Message Queuing**

- Open Help for the Windows Server 2003 family by doing one of the following:

    - Click **Start**, and then click **Help and Support**. Click **Application and Development Support**, and then click **Using Message Queuing**.

    - View Help on the Web at:

      http://go.microsoft.com/fwlink/?LinkId=29881

**Related Links**
For more information about Message Queuing, see the following references:

- The Message Queuing Web site at:

  http://www.microsoft.com/msmq/

- The Message Queuing and Queued Components (MSMQ) SDK documentation, as described in the following list:

    - "Using Message Queuing"

      http://go.microsoft.com/fwlink/?LinkId=46015

    - "Internet Messaging"

      http://go.microsoft.com/fwlink/?LinkId=46014

    - "HTTP Messages"

      http://go.microsoft.com/fwlink/?LinkId=46013

    - "New Features for Windows XP and the Windows Server 2003 family"

      http://go.microsoft.com/fwlink/?LinkId=46012

# Appendix D: Connection Manager and Internet Communication (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Connection Manager 1.3
- How Connection Manager communicates with sites on the Internet
- Controlling Connection Manager to limit the flow of information to and from the Internet

**Benefits and Purposes of Connection Manager 1.3**

Connection Manager 1.3 is a client dialer and connection software program included with products in the Microsoft Windows Server 2003 family. It provides support for local and remote connections to your network service using a network of access points, such as those available worldwide through Internet service providers (ISPs). Connection Manager 1.3 includes features that both simplify and enhance the implementation of connection support for you and your users.

Connection Manager is implemented using the Connection Manager Administration Kit (CMAK). CMAK is designed for network administrators, information officers, and other team members who are responsible for the design, development, testing, distribution, and support of connection software for users who connect to your Internet or network service.

You can customize Connection Manager so that the set of files you distribute to your users (called a service profile) can be easily installed and run. Using Connection Point Services you install Phone Book Service, which enables the distribution of phone books for a Connection Manager service profile.

In addition to the support for basic dial-up connections, you can also use the CMAK wizard to incorporate support for virtual private network (VPN) connections using tunneling protocols to tunnel through a public network (such as when dialing into an ISP to access an organization's server). Connection Manager can create a VPN connection through a preexisting dial-up session, local area network (LAN), or digital subscriber line (DSL).

For more information about Connection Manager, see the instructions in "Accessing Help documentation for Connection Manager for the Windows Server 2003 family," later in this appendix.

**How Connection Manager Communicates with Sites on the Internet**

Administrators control whether to enable the phonebook download feature and which server the client computer sends the request to. The client computer initiates the connection to the phone book server in order to update its data. The phone book server sends one of the following elements to the Connection Manager client while there is a connection:

- Phone book update file: The server sends a phone book update file when the version of the phone book of the connecting client is earlier than the version of the phone book server.
- Replacement phone book: This is sent when there is a difference of more than five versions between the client and the server phone book file versions.
- Standard HTTP error message

The following list describes various aspects of the Connection Manager information that is sent to and from the Internet and how the exchange of information takes place:

- **Specific information sent or received**: The information sent by Connection Manager in the phone book server query string contains the following elements:
    - **osarch**: Specifies the type of processor that the requesting computer is using.
    - **ostype**: Specifies the version of the Windows operating system of the computer making the request.

    - **cmver**: Specifies the version of Connection Manager of the computer making the request.
    - **lcid**: Specifies the locale identifier (LCID) describing the system default locale information of the computer making the request.
    - **pbver**: Specifies the current version of the phone book being used by the computer making the request.
    - **pb**: Specifies the name of the phone book being used by the computer making the request.

- **Default settings**: By default, Connection Point Services and the Connection Manager Administration Kit are not installed with products in the Windows Server 2003 family. (You can add these components by using Add or Remove Programs in Control Panel.)
- **Triggers**: Connection Manager is triggered when users start it to connect to the Internet or their organization's network.
- **Logging**: Information is logged in the Connection Manager log file but not in the operating system event log. The logging feature for Connection Manager is used for troubleshooting and can be disabled.
- **Encryption**: The information contained in the query string is sent as plaintext and is not encrypted.
- **Storage**: The information that is sent is used by the phone book server to process the request being made, although the information is not stored.
- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is 80.
- **Ability to disable**: On the server, you can prevent Connection Point Services and the Connection Manager Administration Kit from being installed, which prevents Connection Manager from being used. For clients, you can disable Connection Manager by not providing the user with a Connection Manager service profile. Also for clients, the Connection Manager phone book download and synchronization feature can be disabled using a Connection Manager service profile.

For more information about the Connection Manager Administration Kit (CMAK), see "Before you start: Understanding Connection Manager and the Administration Kit," on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=46011

You can also view the same information in Help, as described in "Accessing Help documentation for Connection Manager for the Windows Server 2003 family," later in this appendix.

**Controlling Connection Manager to Limit the Flow of Information to and from the Internet**
Connection Manager client software is installed by default. Administrators determine when to create and distribute a Connection Manager service profile, without which Connection Manager cannot run.

You can limit the use of Connection Manager by creating and distributing a service profile only to users and administrators who need to use the Connection Manager client to initiate local or remote network access point connections. You can use the Connection Manager Administration Kit (CMAK) wizard to create a customized service profile to configure Connection Manager not to update the client phone book information.

**Procedures for Installing the Connection Manager Administration Kit and for Viewing Help**
This subsection lists procedures for the following tasks:

- Installing the Connection Manager Administration Kit on your server
- Accessing Help documentation for Connection Manager for the Windows Server 2003 family

**Installing Connection Manager Administration Kit on your server**
The Connection Manager Administration Kit (CMAK) and Connection Point Services (used to install Phone Book Service) are not installed by default on products in the Windows Server 2003 family. Use the following procedure to add these

components.

**To add the Connection Manager Administration Kit and Connection Point Services after server installation**

1. Click **Start**, point to **Control Panel**, and click **Add or Remove Programs**.

2. Click **Add/Remove Windows Components** (on the left).

3. In the Windows Components Wizard scroll down and double-click **Management and Monitoring Tools**.

4. Select **Connection Manager Administration Kit** and **Connection Point Services** and any other subcomponents you want to install.

5. Follow the instructions to complete the Windows Components Wizard.

**Accessing Help documentation for Connection Manager for the Windows Server 2003 family**

Products in the Windows Server 2003 family have Help documentation describing the Connection Manager Administration Kit. You can view this documentation from any computer that has Internet access (regardless of the operating system running on that computer), or from any server running a product in the Windows Server 2003 family. The following procedure provides the details.

**To access Help documentation for a server running a product in the Windows Server 2003 family**

1. Open Help for a product in the Windows Server 2003 family by doing one of the following:

   - On any computer running a product in the Windows Server 2003 family, click **Start**, and then click **Help and Support**.

   - View Help on the Web at:

     http://go.microsoft.com/fwlink/?LinkId=29881

     As appropriate, navigate from this Web site to the documentation for the server product you are using.

2. To view information about Connection Manager Administration Kit, navigate to Network Services\Managing Remote Connections\Connection Manager Administration Kit.

**Related Links**

For information about the Connection Manager Administration Kit (CMAK), see the following Web site:

http://go.microsoft.com/fwlink/?LinkId=46011

# Appendix E: Passport Manager Administration (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

The following appendix provides information about:

- The benefits and purposes of Passport Manager Administration
- Using Passport Manager Administration in a managed environment
- Resources for learning about Passport Manager Administration

**Benefits and Purposes of Passport Manager Administration**
Passport Manager Administration is a utility designed for use on Web sites on which you want to use Microsoft .NET Passport for authentication. In Microsoft Server 2003 family operating systems, Passport Manager Administration provides a graphical-interface alternative to editing the registry to change Passport Manager object settings. The most common use of the Passport Manager Administration utility is to change the default object settings that Passport Manager methods use if certain optional parameters are omitted when the methods are called.

Although some parameters of Microsoft .NET Passport methods are listed as optional, these methods still generate URLs or results that reflect inherent default values when called. These default parameter values can be used to provide consistent values for your Web site, such as the required time window within which all users must be authenticated.

Primary default object settings that affect Passport Manager implementation include: Time Window, Force Login, and Language ID.

**Using Passport Manager Administration in a Managed Environment**
By configuring options in Passport Manager Administration to work with software on your Web site, you can control how Web browsers are authenticated to your Web site. With Passport Manager Administration, you can configure the settings for cookies, such as choosing to disable the use of cookies. You can also choose to have cookies deleted once users sign out of your Web site. (Cookies are a form of data passed both ways between Web browsers and servers. Cookies sent to a user's browser have some implications for privacy.)

To ensure that cookies are deleted when users sign out of a Web site, make sure the following conditions are met:

- The Site ID value in Passport Manager Administration should be the value given to you when you received your Passport encryption key.
- The path and file name of the page to which cookies are written (specified in Passport Manager Administration) should match the ExpireCookieURL provided when you registered your site. Also, if cookies are written to a page in a subdomain of your own domain, you should precede the domain path with a dot (".").
- The path to the page where cookies are written should be the root directory of the Web site. This ensures that cookies written to all possible paths are deleted. If this is not done, not all the cookies may be deleted.
- If your site uses multiple servers with Passport Manager installed on them, you should verify that all of the servers are writing the cookies to the same domain and path.

**Related Links**
You can find additional information about Passport Manager Administration at these sites:

- For more information about .NET Passport and how to download it, see the SDK documentation on the MSDN Web site at:

  http://go.microsoft.com/fwlink/?linkid=7739

- The Help file for Passport Manager Administration utility provides detailed information about configuring settings for this utility. You can find this Help file at:

  http://go.microsoft.com/fwlink/?linkid=7738

# Appendix F: Internet Connection Sharing and Related Networking Features (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

Internet Connection Sharing, Internet Connection Firewall, and Network Bridge are features designed for home and small office networks. These features are offered in some of the Microsoft Windows Server 2003 family operating systems. Information about these features is presented here so you as an IT administrator can be aware of these potential capabilities within your organization's network.

These home and small office features are included only with Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition. These features are not included with Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; or the 64-bit versions of the Windows Server 2003 family.

This appendix includes the following information:

- An overview of Internet Connection Sharing and related networking features.
- How Internet Connection Sharing and related features can be used in a large organization's network.
- How to control or prevent the use of Internet Connection Sharing and related features.

**Overview: Internet Connection Sharing and Related Networking Features**

The features for implementing and administering small networks are described as follows:

- Internet Connection Sharing (ICS)

  ICS provides Internet access for a home or small office network by using one common connection as the Internet gateway. The ICS host is the only computer that is directly connected to the Internet. Multiple ICS clients simultaneously use the common Internet connection and benefit from Internet services as if the clients were directly connected to the Internet service provider (ISP). Security is enhanced when ICS is enabled because only the ICS host computer is visible to the Internet. The addresses of ICS clients are hidden from the Internet rendering ICS clients invisible to the Internet. In addition, ICS simplifies the configuration of small networks by providing local private network services, such as name resolution and addressing.

  | **Note** |
  | --- |
  | You should not use Internet Connection Sharing in an existing network with Windows 2000 Server domain controllers, Domain Name System (DNS) servers, gateways, Dynamic Host Configuration Protocol (DHCP) servers, or systems configured for static IP addresses. |

- Internet Connection Firewall (ICF)

  With ICF, the firewall checks all communications that cross the connection between your network and the Internet and is selective about which responses from the Internet it allows. ICF protects only the computer on which it is enabled. If ICF is enabled on the Internet Connection Sharing (ICS) host computer, however, ICS clients that use the shared Internet connection for Internet connectivity are protected because they cannot be seen from outside your network. For this reason, you should always enable ICF on the ICS host computer. In addition, if there are clients on your network with direct Internet connections, or if you have a stand-alone computer that is connected to the Internet, then you should enable ICF on those Internet connections as well.

- Network Bridge

  Network Bridge removes the need for routing and bridging hardware in a home or small office network that consists of multiple LAN segments. With Network Bridge, multiple LAN segments become a single IP subnet, even if the LAN

segments are of mixed network media types. Network Bridge automates the configuration and management of the address allocation, routing, and name resolution that is typically required in a network that consists of multiple LAN segments.

| ⚠**Caution** |
| --- |
| If neither ICF nor ICS is enabled on your network, do not set up Network Bridge between the public Internet connection and the private network connection. Setting up Network Bridge between the public Internet connection and the private network connection creates an unprotected link between your network and the Internet, leaving your network vulnerable to external attacks. When either ICF or ICS is enabled, this risk is mitigated. |

**Using Internet Connection Sharing and Related Features in a Managed Environment**

Internet Connection Sharing, Internet Connection Firewall, and Network Bridge are not enabled by default, and Internet Connection Sharing (ICS) is available only on computers that have two or more network connections. An administrator or user with administrative credentials can enable ICS by clicking the **Advanced** tab on network connections (Control Panel\Network Connections). Also, when running the New Connection Wizard, administrators can choose to enable ICS. ICS lets administrators configure a computer as an Internet gateway for a small network, and it provides network services such as name resolution through Domain Name System (DNS). It also provides addressing through Dynamic Host Configuration Protocol (DHCP) to the local private network.

Using Internet Connection Firewall, an administrator can enable a firewall to protect the public connection of a small network or single computer that is connected to the Internet. ICF is considered a "stateful" firewall. A stateful firewall is one that monitors all aspects of the communications that cross its path and inspects the source and destination address of each message that it handles.

Any organization that uses domain controllers, DHCP, DNS, and other elements of network infrastructure should not use ICS and ICF, but can instead use a firewall designed for the entire organization.

The Network Bridge menu command **Bridge Connections** is available only when two or more network adapters are present. By default, Network Bridge is disabled, but administrators can use **Bridge Connections** to enable Network Bridge. In a domain environment you should not allow the capability to enable or configure these features. See the following subsection for information about how to disable them.

It is important to be aware of all the methods users and administrators have for connecting to your networked assets, and to review whether your security measures provide in-depth defense (as contrasted with a single layer of defense, more easily breached).

**Controlling the Use of Internet Connection Sharing and Related Features**

You can block administrators from accessing ICS, ICF, and Network Bridge by using answer files during initial installation and Group Policy post-deployment.

**Using answer files for unattended or remote installation**

Using standard methods for preparing an unattended or remote installation you can make entries in the [Homenet] section of the answer file. This section includes entries for installing home and small office networking settings for network adapters, Internet Connection Sharing, Internet Connection Firewall, and Network Bridge. For example, to prevents users and administrators from enabling Internet Connection Sharing by using an answer file, the entry is as follows:

```
[Homenet]
EnableICS = No
```

For additional configuration options for [Homenet] entries for the answer file, and for more information about unattended installation, see the references listed in Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003). Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix).

**Using Group Policy**

Group Policy settings for disabling small office networking features in your domain environment are described as follows:

- **Prohibit use of Internet Connection Sharing on your DNS domain network**

  This policy setting determines whether administrators can enable and configure the Internet Connection Sharing (ICS) feature on a connection. It also determines if ICS can run on a computer when the computer is connected to the DNS domain in which the policy setting is applied.

- **Prohibit use of Internet Connection Firewall on your DNS domain network**

  This policy setting determines whether administrators can enable and configure the Internet Connection Firewall feature on a connection.

- **Prohibit installation and configuration of Network Bridge on your DNS domain network**

  This policy setting determines whether administrators can enable and configure Network Bridge on your domain.

| **Important** |
| --- |
| These policy settings are dependent on the network context that the computer is in. They apply only when a computer is connected to the same DNS domain network it was connected to when the policy setting was refreshed on that computer. If a computer is connected to a DNS domain network other than the one it was connected to when the policy setting was refreshed, the policy setting does not apply. |

These policy settings are located in Computer Configuration\Administrative Templates\Network\Network Connections. Configuration options are described in the following table.

**Group Policy settings for controlling ICS, ICF, and Network Bridge**

| Policy setting | Description |
| --- | --- |
| **Prohibit Use of Internet Connection Sharing on your DNS domain network** (enabled) | If you enable this policy setting, ICS cannot be enabled or configured by administrators, and the ICS service cannot run on the computer. The Advanced tab in the Properties dialog box for a local area network (LAN) or remote access connection is removed. The Internet Connection Sharing page is removed from the New Connection Wizard. The Network Setup Wizard is disabled. |
| **Prohibit Use of Internet Connection Sharing on your DNS domain network** (disabled or not configured) | If you disable this policy setting or do not configure it and have two or more connections, administrators can enable ICS. The Advanced tab in the Properties dialog box for a LAN or remote access connection is available. In addition, the administrator is presented with the option to enable Internet Connection Sharing in the Network Setup Wizard and Make New Connection Wizard. (The Network Setup Wizard is available only in Windows XP Professional.) |
| **Prohibit Use of Internet Connection Firewall on your DNS domain network** (enabled) | If you enable this policy setting, Internet Connection Firewall cannot be enabled or configured by administrators, and the Internet Connection Firewall service cannot run on the computer. The option to enable the Internet Connection Firewall through the Advanced tab is removed. In addition, the Internet Connection Firewall is not enabled for remote access connections created through the Make New Connection Wizard. The Network Setup Wizard is disabled. |
| **Prohibit Use of Internet Connection Firewall on your DNS domain network** (disabled or not configured) | If you disable this policy setting or do not configure it, the Internet Connection Firewall is disabled when a LAN connection or virtual private network (VPN) connection is created, but administrators can use the Advanced tab in the connection properties to enable it. The Internet Connection Firewall is enabled by default on the connection for which Internet Connection Sharing is enabled. In addition, remote access connections created through the Make New Connection Wizard have Internet Connection Firewall enabled. |
| **Prohibit installation and configuration of Network Bridge on your domain network** (enabled) | When you enable this policy setting administrators cannot create a Network Bridge. Enabling this policy setting does not remove an existing Network Bridge from a computer. |
| **Prohibit installation and configuration of Network Bridge on your domain network** (disabled or not configured) | If you disable this policy setting or do not configure it, an administrator will be able to create and modify the configuration of a Network Bridge. |

For more information about home and small office networking features, see Help and Support Center for the Windows Server 2003 family.

# Appendix G: Add Network Place Wizard (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

In the Microsoft Windows Server 2003 family operating systems, users and administrators can use the Add Network Place Wizard to create shortcuts to shared folders and resources on the network or on Web or File Transfer Protocol (FTP) servers. If you don't have folders on a Web server already, the Add Network Place Wizard helps you create a new folder for storing files online.

The content in this appendix includes the following information:

- An overview of the Add Network Place Wizard
- How to control the use of the Add Network Place Wizard

**Overview: Add Network Place Wizard**

The Add Network Place Wizard is enabled by default for all users. In operating systems in the Windows Server 2003 family you access the wizard through Windows Explorer\My Network Places. After you click **My Network Places** and then double-click **Add Network Place**, the wizard appears.

You can use the wizard to sign up for a service that offers online storage space, or to create a shortcut to a Web site, an FTP site, or other local network connection. To add a shortcut in My Network Places to a folder on a Web server, the Web server must support network places. Support for network places requires the Web Extender Client (WEC) protocol and Microsoft FrontPage® Server Extensions, or the Web Distributed Authoring and Versioning (WebDAV) protocol and Internet Information Services (IIS). You must also have read and write access to the Web server.

In a highly managed network environment administrators might want to prevent users or administrators from storing or accessing folders on a Web server. You can remove access to the Add Network Place Wizard using Group Policy.

For more information about the WEC and WebDAV protocols, see Help and Support Center.

**Controlling the Use of the Add Network Place Wizard**

You can block users and administrators from accessing the Add Network Place Wizard in My Network Places by configuring a Group Policy setting.

Configure the following Group Policy setting in User Configuration\Administrative Templates\Windows Components\Windows Explorer: **Remove "Map Network Drive" and "Disconnect Network Drive**."

When you enable this policy setting, in addition to preventing users from using Windows Explorer or My Network Places to map or disconnect network drives, you also remove the Add Network Place option from My Network Places. Users can still connect to another computer on your intranet by typing the name of a shared folder in the Run dialog box.

| **Note** |
| --- |
| This policy setting was documented incorrectly on the Explain tab in Group Policy for Windows 2000. That Explain tab states incorrectly that this policy setting prevents users from connecting and disconnecting drives. |

# Appendix H: New Connection Wizard (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

In Microsoft Windows Server 2003 family operating systems you use the New Connection Wizard to create Internet and other types of network connections for home and small office networks. While this feature is designed for home and small office use, information about this feature is presented here so IT administrators can be aware of these potential capabilities within your organization's network.

The content in this appendix includes the following information:

- An overview of the New Connection Wizard
- How to control the use of the New Connection Wizard

**Overview: New Connection Wizard**

The New Connection Wizard in Windows Server 2003 family operating systems replaces the Windows 2000 Network Connection Wizard and Internet Connection Wizard. Administrators for a home or small office network can use the New Connection Wizard to create any type of network connection including Internet, incoming, dial-up, virtual private network (VPN), and direct connection.

Administrators can create a new connection through Control Panel\Network Connections. When you click **Create a new connection**, the wizard guides you through this process. In a highly managed network environment you might want to prevent administrators as well as users from creating new connections.

| **Note** |
| --- |
| The Internet Connection Wizard still appears when you first click any program that requires an Internet connection, such as Microsoft Internet Explorer. The Internet Connection Wizard can also be started from Outlook Express, if a link for **Set up a Mail account** or for **Set up a Newsgroups account** is visible and you click the link. You can disable the Internet Connection Wizard in Group Policy through User Configuration\Administrative Templates\Windows Components\Internet Explorer. |

**Controlling the Use of the New Connection Wizard**

You can block administrators and users from using the New Connection Wizard by configuring Group Policy. The following two Group Policy settings are in User Configuration\Administrative Templates\Network\Network Connections:

- **Prohibit Access to the New Connection Wizard**

  This policy setting determines whether users, including administrators, can use the New Connection Wizard, which creates new Internet or intranet connections.

  | **Important** |
  | --- |
  | If the policy setting **Enable Windows 2000 Network Connections settings for Administrators** is disabled or not configured, **Prohibit Access to the New Connection Wizard** will not apply to administrators on computers running Windows XP. |

- **Enable Windows 2000 Network Connections settings for Administrators**

  This policy setting determines whether policy settings that exist in Windows 2000 will apply to administrators. The set of Network Connections policy settings that exists in Windows 2000 also exists in Windows XP. In Windows 2000, all of these policy settings have the ability to prohibit the use of certain features by administrators. By default, Network Connections policy settings in Windows XP do not prohibit the use of features from administrators.

| Note |
| --- |
| This policy setting is intended to be used in a situation in which the Group Policy object (GPO) contains computers running both Windows 2000 and Windows XP and identical Network Connections policy setting behavior is required between those computers. |

Configuration options for these policy settings are presented in the following table.

**Group Policy settings for controlling the use of the New Connection Wizard**

| Policy setting | Description |
| --- | --- |
| **Prohibit Access to the New Connection Wizard** (enabled) | If you enable this policy setting (and enable the **Enable Network Connections settings for Administrators** policy setting), the Make New Connection icon does not appear in the Start menu or in the Network Connections folder. As a result, users (including administrators) cannot start the New Connection Wizard. |
| **Prohibit Access to the New Connection Wizard** (disabled or not configured) | If you disable this policy setting or do not configure it, the Make New Connection icon appears in the Start menu and in the Network Connections folder for all users (except in a workgroup environment where only administrators can access this wizard). Clicking the Make New Connection icon starts the New Connection Wizard.<br><br>Changing this policy setting from enabled to not configured does not restore the Make New Connection icon until the user logs off or on. When other changes to this policy setting are applied, the icon does not appear or disappear in the Network Connections folder until the folder is refreshed. |
| **Enable Windows 2000 Network Connections settings for Administrators** (enabled) | If you enable this policy setting, the Windows XP policy settings that exist in Windows 2000 can prohibit administrators from using certain features (see the policy settings in the Network Connections folder). With this policy setting enabled, policy settings that exist in both Windows 2000 and Windows XP behave the same for administrators. |
| **Enable Windows 2000 Network Connections settings for Administrators** (disabled or not configured) | If you disable this policy setting or do not configure it, Windows XP policy settings that exist in Windows 2000 will not apply to administrators. |

For information about using Group Policy, see Appendix B: Resources for Learning About Group Policy (Windows Server 2003).

# Related Links for Controlling Communication with the Internet (Windows Server 2003)

Updated: July 31, 2004

Applies To: Windows Server 2003 with SP1

This section contains a list of the Web sites found in other sections of this white paper, with some additional Web sites that provide general information about Microsoft products.

**Links to product information, support information, TechNet, Microsoft Developer Network, and information in Resource Kits**
The following sites provide information about products in the Windows Server 2003 family and other Microsoft products. The list includes sites containing product Help as well as other basic sites that provide information about Microsoft operating systems and other Microsoft products:

- Products in the Windows Server 2003 family (including Windows Server 2003 Help on the Web):

  http://go.microsoft.com/fwlink/?LinkId=29881

  For links to specific topics in Help, see the following sections of the white paper: "Hyperterminal"; Appendix A: Resources for Learning About Automated Installation and Deployment (Windows Server 2003); and Appendix D: Connection Manager and Internet Communication (Windows Server 2003).

- Windows Catalog:

  http://go.microsoft.com/fwlink/?linkid=4303

- Microsoft product support:

  http://support.microsoft.com/

- Microsoft TechNet:

  http://www.microsoft.com/technet/

- Microsoft Developer Network:

  http://msdn.microsoft.com/

- Prescriptive Architecture Guides:

  http://go.microsoft.com/fwlink/?LinkId=29413

- *Windows Deployment and Resource Kits* (includes links to the *Windows Server 2003 Deployment Kit* and the *Windows Server 2003 Technical Reference*):

  http://www.microsoft.com/windows/reskits/

- Internet Explorer Administration Kit:

  http://go.microsoft.com/fwlink/?linkid=29479

- *Microsoft NetMeeting 3 Resource Kit* (the NetMeeting section in this white paper provides links to specific chapters in this kit):

  http://go.microsoft.com/fwlink/?LinkId=36134

**Links to information about security, management, and deployment**
The following sites provide information about security, management, and deployment topics:

- Managing mobile code:

  http://go.microsoft.com/fwlink/?linkid=29170

- *Automating and Customizing Installations* in the *Microsoft Windows Server 2003 Deployment Kit*:

  http://go.microsoft.com/fwlink/?LinkId=29887

- Using scripts:

  - http://go.microsoft.com/fwlink/?LinkId=24771
  - http://go.microsoft.com/fwlink/?LinkId=29221

- Security:

  - http://www.microsoft.com/security/guidance/
  - http://go.microsoft.com/fwlink/?LinkId=29168
  - http://go.microsoft.com/fwlink/?LinkId=46003
  - http://msdn.microsoft.com/security/
  - http://www.microsoft.com/technet/security/
  - http://go.microsoft.com/fwlink/?LinkId=46004

**Links to information about components in Windows Server 2003 or tools you can use with Windows Server 2003**
The following sites provide information about some of the components in the Windows Server 2003 family:

- Certificates, certificate status, certificate revocation, and Public Key Infrastructure (PKI):

  - http://go.microsoft.com/fwlink/?linkid=29886
  - http://go.microsoft.com/fwlink/?LinkId=27081

- Driver Protection:

  http://go.microsoft.com/fwlink/?LinkId=33957

- Dynamic Update:

  http://go.microsoft.com/fwlink/?linkid=29887

- Hyperterminal:

  http://go.microsoft.com/fwlink/?LinkId=225

  (Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

- Internet Explorer:

  http://www.microsoft.com/windows/ie/

For more information about Internet Explorer, also see the following:

  - "Microsoft Resource Kits" and "Internet Explorer Administration Kit" in "Links to Product Information, Support Information, TechNet, Microsoft Developer Network, and Information in Resource Kits," earlier in this section.
  - "Language codes" in this list.

- Internet Information Services:
  - http://go.microsoft.com/fwlink/?LinkId=46005
  - http://go.microsoft.com/fwlink/?linkid=22805
  - http://go.microsoft.com/fwlink/?linkid=31134
  - http://go.microsoft.com/fwlink/?linkid=29896

- Internet Protocol version 6:

  http://go.microsoft.com/fwlink/?LinkId=29519

  For more information about IPv6, also see "Links to Sites Maintained by Task Forces and Other Organizations," later in this section.

- Language codes (used in Internet Explorer and the file association Web service when a language is being specified):

  http://go.microsoft.com/fwlink/?linkid=29165

- Message Queuing and the Message Queuing SDK:
  - http://go.microsoft.com/fwlink/?LinkId=46016
  - http://go.microsoft.com/fwlink/?LinkId=46015
  - http://go.microsoft.com/fwlink/?LinkId=46014
  - http://go.microsoft.com/fwlink/?LinkId=46013
  - http://go.microsoft.com/fwlink/?LinkId=46012

- NetMeeting:
  - http://www.microsoft.com/windows/NetMeeting/
  - http://go.microsoft.com/fwlink/?LinkId=29175
  - http://go.microsoft.com/fwlink/?LinkId=29206
  - http://go.microsoft.com/fwlink/?LinkId=29207

  For additional information about NetMeeting, see the *Microsoft NetMeeting 3 Resource Kit* in "Links to product information, support information, TechNet, Microsoft Developer Network, and information in Resource Kits," earlier in this section.

- Passport Manager Administration:

    - http://go.microsoft.com/fwlink/?linkid=7739
    - http://go.microsoft.com/fwlink/?linkid=7738
- Remote Assistance:

  http://go.microsoft.com/fwlink/?LinkId=29212
- Search Companion:

  http://sa.windows.com/privacy/
- Terminal Server Licensing (and Terminal Server):
    - http://go.microsoft.com/fwlink/?LinkId=46192
    - http://go.microsoft.com/fwlink/?LinkId=46078
    - http://go.microsoft.com/fwlink/?LinkId=46007
    - http://go.microsoft.com/fwlink/?linkid=24427
- Windows Error Reporting:
    - http://go.microsoft.com/fwlink/?LinkId=825
    - http://go.microsoft.com/fwlink/?LinkId=29903
    - http://go.microsoft.com/fwlink/?LinkId=29517
- Windows Media Player:
    - http://www.microsoft.com/Windows/WindowsMedia/
    - http://go.microsoft.com/fwlink/?LinkId=46006
    - http://go.microsoft.com/fwlink/?LinkId=46008
- Windows Media Services:
    - http://go.microsoft.com/fwlink/?LinkId=29862
    - http://go.microsoft.com/fwlink/?LinkId=46019
    - http://go.microsoft.com/fwlink/?LinkId=46020
    - http://go.microsoft.com/fwlink/?LinkId=58341
    - http://go.microsoft.com/fwlink/?LinkId=58342
    - http://go.microsoft.com/fwlink/?LinkId=58343
    - http://go.microsoft.com/fwlink/?LinkId=58345

- Windows Time Service:

  http://go.microsoft.com/fwlink/?LinkId=46021

- Windows Update and Automatic Updates:

  - http://windowsupdate.microsoft.com/

  - http://go.microsoft.com/fwlink/?LinkId=29906

The following site provides information about using tools that work with the Windows Server 2003 family:

- Windows Application Compatibility Toolkit:

  http://go.microsoft.com/fwlink/?LinkId=29880

**Links to information about licensing, product activation, and registration**
The following sites provide information about licensing, product activation, and registration:

- http://go.microsoft.com/fwlink/?LinkId=29878

- http://go.microsoft.com/fwlink/?linkid=29225

- http://go.microsoft.com/fwlink/?LinkId=29484

- http://go.microsoft.com/fwlink/?LinkId=29923

- http://go.microsoft.com/fwlink/?linkid=29226

- http://go.microsoft.com/fwlink/?LinkId=29508

**Links to sites maintained by task forces and other organizations**
The following sites are maintained by the Internet Engineering Task Force:

- http://go.microsoft.com/fwlink/?LinkId=29898

- http://go.microsoft.com/fwlink/?LinkId=29135

- http://go.microsoft.com/fwlink/?LinkId=29136

- http://go.microsoft.com/fwlink/?LinkId=29215

- http://go.microsoft.com/fwlink/?LinkId=29924

The following site is maintained by the International Multimedia Telecommunications Consortium:

http://www.imtc.org/

The following site is maintained by the International Telecommunication Union:

http://go.microsoft.com/fwlink/?LinkId=29510

(Web addresses can change, so you might be unable to connect to the Web sites mentioned here.)

**Links to information about Group Policy**
The following sites provide information about topics related to Group Policy:

- Windows Server 2003 Group Policy page on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=46022

- The Group Policy Technology Center Web site at:

  http://www.microsoft.com/grouppolicy/

- *Designing a Managed Environment* in the *Microsoft Windows Server 2003 Deployment Kit* at:

  http://go.microsoft.com/fwlink/?linkid=44319

- The Group Policy Collection of the *Windows Server 2003 Technical Reference* at:

  http://go.microsoft.com/fwlink/?LinkId=29907

- Article 816662, "Recommendations for Managing Group Policy Administrative Template (.adm) Files" in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?LinkId=29128

- Information about the Group Policy Management Console at the Microsoft Windows Server 2003 Web site at:

  http://go.microsoft.com/fwlink/?linkid=29909

- The Group Policy Settings Reference on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29911

# Network Load Balancing: Configuration Best Practices for Windows 2000 and Windows Server 2003

Applies To: Windows Server 2003 with SP1

**In This White Paper**


General Considerations

Security and Manageability

High Availability

Troubleshooting Network Load Balancing

# General Considerations

Applies To: Windows Server 2003 with SP1

- Some routers require a static ARP entry because they do not support the resolution of unicast IP addresses to multicast media access control addresses. For example, Cisco routers require an ARP (address resolution protocol) entry for every virtual IP address. While Network Load Balancing uses Level 2 Multicast for the delivery of packets, Cisco's interpretation of the RFCs is that Multicast is for IP Multicast. So, when the router doesn't see a Multicast IP address, it does not automatically create an ARP entry, and one has to manually have to add it on the router.

- Network Load Balancing can operate in two modes: unicast and multicast. Unicast support is enabled by default, which ensures that it operates properly with all routers. You might elect to enable multicast mode so that a second network adapter is not required for communications within the cluster. If Network Load Balancing clients access a cluster (configured for multicast mode) through a router, be sure that the router accepts an Address Resolution Protocol (ARP) reply for the cluster's (unicast) IP addresses with a multicast media access control address in the payload of the ARP structure. ARP is a TCP/IP protocol that uses limited broadcast to the local network to resolve a logically assigned IP address. Verify that all cluster hosts are operating in unicast or multicast mode, one or the other, but not both.

- If the cluster is operating in unicast mode (default setting), Network Load Balancing cannot distinguish between single adapters on each host. Therefore, any communication among cluster hosts is not possible unless each cluster host has at least two network adapters.

- You can configure Network Load Balancing on more than one network adapter. However, if you do bind NLB to a second network adapter ensure that you are configuring them correctly.

- Use only the TCP/IP network protocol on the adapter NLB is enabled for. Do not add any other protocols (for example, IPX) to this adapter.

- Enable Network Load Balancing Manager logging. You can configure Network Load Balancing manager to log each Network Load Balancing Manager event. This log can be very useful in troubleshooting problems or errors when using Network Load Balancing Manager. Enable Network Load Balancing Manager logging by clicking Log Settings in the Network Load Balancing Manager Options menu. Check the Enable logging box and specify a name and location for the log file.

- Verify that the following is true for cluster parameters, port rules, and host parameters:

    - Cluster parameters and port rules are set identically on all cluster hosts.

    - Port rules are set for all ports used by the load-balanced application. For example, FTP uses port 20, port 21, and ports 102465535).

    - Always click Add after setting a port rule. Otherwise, the port rule will not appear in the list of rules, and the rule will not take effect.

- Ensure that the dedicated IP address is unique and the cluster IP address is added to each cluster host.

- Verify that any given load-balanced application is started on all cluster hosts on which the application is installed. Network Load Balancing is not aware higher level applications and does not start or stop applications.

- Verify that the following is true for the dedicated IP address and the cluster IP address:

    - Except in the case of a virtual private network (VPN), both the dedicated IP address and the cluster IP address must be entered during setup in the Network Load Balancing Properties dialog box and also in the Internet Protocol (TCP/IP) Properties dialog box. Make sure that the addresses are the same in both places.

    - When configuring a VPN load balancing cluster, you should not configure the dedicated IP address. On a VPN,

only the cluster IP address should be present on each of the cluster hosts because clients running Windows 95, Windows 98, or Windows NT 4.0 may be unable to connect to the cluster if the dedicated IP address is configured on the Network Load Balancing cluster hosts. If you omit this step, the cluster will converge and appear to be working properly, but the cluster host will not accept and handle cluster traffic.

  - Ensure that the dedicated IP address is always listed first (before the cluster IP address) in the Internet Protocol (TCP/IP) Properties dialog box. This will ensure that responses to connections originating from a host will return to the same host.

  - Both the dedicated IP address and the cluster IP address must be static IP addresses. They cannot be DHCP addresses.

  - Ensure that all hosts in a cluster belong to the same subnet and that the cluster's clients are able to access this subnet.

- No special cluster interconnect is used by Network Load Balancing. NLB uses the same network interface to maintain cluster state awareness.

- Do not enable Network Load Balancing on a computer that is part of a server-cluster cluster. Microsoft does not support this configuration.

# Security and Manageability

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

- Use Network Load Balancing Manager to configure NLB clusters. You can configure many Network Load Balancing options through either Network Load Balancing Manager or the Network Load Balancing Properties dialog box accessed through Network Connections. However, Network Load Balancing Manager is the preferred method. Using both Network Load Balancing Manager and Network Connections together to change Network Load Balancing properties can lead to unpredictable results. Only Windows Server 2003 NLB clusters can be configured by NLB manager. You can however manage clusters that contain both Windows Server 2003 and Windows 2000 or NT 4.0 servers.

- Ensure that applications that are load balanced are properly secured. The NLB security domain does not extend to applications. As such NLB will be totally unaware if security at the applications level is compromised.

- Use two or more network adapters in each cluster host if you would like to separate management functions from regular operations. Two network adapters, is not however a default requirement.

- Command line tool for managing NLB is "nlb.exe". NLB.exe exposes a mechanism for setting up NLB configuration parameters thru the command line. There are 2 additional configuration points not exposed but can be useful for monitoring NLB state. They are ***queryport*** and ***params Nlb.exe queryport*** retrieves the state of a given port rule using the same syntax as the enable/disable/drain command line options... the information returned includes the state of the port rule, enabled, disabled or draining if the port rule is found or an indication that the port rule was not found... if found, it also returns a count of packets accepted and dropped on that port rule. ***Nlb.exe - params*** retrieves the NLB configuration just the same as "nlb display", but rather than retrieving it from the registry, it queries it directly from the kernel-mode driver - this is the CURRENT state of NLB (the registry shows what the NEXT state of NLB would be if a reload or some other operation causing the driver to read the registry was performed - the registry MAY or MAY NOT be the current state of NLB)

- Enabling remote control has security implications and the user must ensure that the NLB cluster is secure (behind a firewall) if remote control is enables. The remote control mechanism uses the UDP protocol and is assigned port 2504. Remote control datagrams are sent to the clusters primary IP address. Since the Network Load Balancing driver on each cluster host handles them, these datagrams must be routed to the cluster subnet (instead of to a back-end subnet to which the cluster is attached). When remote control commands are issued from within the cluster, they are broadcast on the local subnet. This ensures that all cluster hosts receive them even if the cluster runs in unicast mode. As such the subnet the NLB clusters are hosted on should be secure. If remote control is enabled users can use nlb.exe to remotely manage their clusters.

# High Availability

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

- Network Adapters and adapter teaming: Most vendors today offer redundant or fault tolerant adapters i.e. adapter teaming or adapter fault tolerance (AFT). These are supported with NLB, however refer to KB article 278431 for more information.

- Fault Tolerant/Load balancing Switches: Redundancy at the switch layer can easily be provided by striping the NLB cluster hosts across multiple switches and inter-connecting all the switches that contain a single NLB cluster. Additionally, to prevent switch flooding, only the ports connected to the Primary IP address (where all inbound traffic is sent) can be made hosts of a single VLAN.

- Fault tolerant Routers: Redundant routers are the most easily overcome using a VRP (virtual router protocol) or HSRP (hot router standby protocol). This allows the router to map the cluster's primary IP address and other multi-homed addresses to the corresponding media access control address. If your router does not meet this requirement, you can create a static ARP entry in the router or you can use Network Load balancing in its default unicast mode.

- Multiple network interface controllers in cluster nodes

  Windows 2000:

  - If you have two network interface controllers on different subnets then the network interface controller to which NLB is bound should have default gateway and the routing tables need to be reconfigured to make all traffic go through the NLB network interface controller. Default gateway setting on the other network interface controller should be blank.

  - If you have two network interface controllers on the same subnet you will need to configure the network interface controller to which NLB is bound with the default gateway. The other network interface controller should not have a default gateway configured. No need to hack routing tables.

  - Recommendation is to use one network interface controller in each node unless there is a business need for two network interface controllers

  Windows Server 2003

  - If you have two network interface controllers on different subnets you can choose to bind NLB to either or both network interface controllers without any issues. All traffic will go through the correct network interface controller (subnet)

  - If you have two network interface controllers on the same subnet traffic will be routed accordingly.

# Troubleshooting Network Load Balancing

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

- The following tools can be used to troubleshoot NLB clusters:
    - Event Viewer.
    - NLB.exe Display & Query Commands.
    - Ping.exe.
    - Network Monitor.
    - Network Monitor parser for NLB (part of Windows 2000 Server resource kit)

      Refer to KB article 280503 for more information.
    - Performance Monitor

      CPU Load

      Network Interface: packets/sec

      Web Service: connection attempts/sec.

# Deploying Wireless Provisioning Services (WPS) Technology

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This whitepaper describes a Wireless Internet Service Provider (WISP) network that allows you to provide pay-per-use, monthly service, and long-term Internet access to new and existing customers through wireless access points deployed in public areas, such as airports and shopping malls. With Wireless Provisioning Services (WPS) technology, new and existing customers with computers running Windows XP with Service Pack 2 (SP2) can connect to your Wi-Fi network without manual configuration of the computer or network connection.

**In This White Paper**

Introduction (Deploying Wireless Provisioning Services (WPS) Technology)

Configuring IAS for WPS Technology

WPS Technology for a WISP with VLANs

WPS Technology for the Enterprise

WISP Roaming Agreement Deployments

Server Certificate Requirements

Active Directory Replication

XML Schemas

How to Use IAS with a Third-Party User Accounts Database

PEAP-MS-CHAP v2

How to Create an IAS Extension DLL and a URL PEAP-TLV

Beta Documentation Note

WPS Technology for a WISP with IP Filters

WPS Technology for an HSP with IP Filters

Summary (Deploying Wireless Provisioning Services (WPS) Technology)

Related Links (Deploying Wireless Provisioning Services (WPS) Technology)

# Introduction (Deploying Wireless Provisioning Services (WPS) Technology)

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using Wireless Provisioning Services (WPS) technology, your organization can offer wireless access in any location where you provide connectivity with Wi-Fi hotspots to new and existing customers that have computers running Windows XP with Service Pack 2 (SP2).

Wi-Fi means "wireless fidelity" and is the IEEE 802.11 technical standard for short range wireless data transmissions. Wi-Fi hotspots are areas where you deploy one or more wireless access points. WPS technology allows new customers without a previously established account to securely connect to your network at the Wi-Fi hotspot location, create and pay for an account, and access the Internet.

One of the key challenges of deploying wireless technologies in public places is providing customers with the correct network configuration for their computers. Many customers are unaccustomed to manually configuring network settings, and might forego use of public wireless networks due to the difficulty of these configurations. In addition, erroneous configurations can prevent customers from connecting successfully.

With WPS technology, your customers do not need any knowledge of network technology, and they do not need to manually configure their computers or join a domain to connect to your network. During the account-creation process, customers' computers are automatically and transparently supplied — or provisioned — with all configuration necessary for them to successfully access your network and services.

WPS technology allows you to provide your customers with the following account-creation and computer configuration options:

- Customers can create their account and automatically provision their computer over the Internet in advance of arrival at your Wi-Fi hotspot locations.

- Customers who arrive at your Wi-Fi hotspots without previously creating an account and provisioning their computers can spontaneously connect, provision their computer, create an account, and access the Internet through your Wi-Fi hotspot.

WPS technology is deployed by using the following:

- One or more computers running Windows Server 2003 with Service Pack 1 (SP1) and later, and Internet Authentication Service (IAS). If you use a RADIUS server other than IAS in Windows Server 2003 with SP1, you need to verify with your RADIUS server vendor whether their servers support WPS. A royalty-free license that covers WPS protocol implementation is available for RADIUS and other server vendors through the Published Protocols and Royalty-Free License site on MSDN at http://go.microsoft.com/fwlink/?LinkId=33674.

- One or more computers running Windows XP Home Edition with SP2; Windows XP Professional with SP2; or Windows XP Tablet PC Edition with SP2.

- Additional servers on your network. For example, these servers include a provisioning server, a DHCP server, and a domain controller with a user accounts database. Depending on the deployment method you choose, you might need additional servers as described in this paper.

- Network access servers and other network hardware. This hardware includes wireless access points, routers, and other devices.

- One or more custom applications or databases that you design. For example, a Web application running on an HTTPS-enabled Web server that passes customer data from the customer to the provisioning server.

These hardware and software components of WPS technology are deployable by three types of organizations with two different methods used to provide customers with the ability to create and pay for an account before obtaining access to the

Internet. The following introductory sections describe these organizations and methods, and further introduce the key components of WPS technology.

**Who can use WPS technology**

WPS technology is designed for use by three types of organizations:

- **Hotspot service providers (HSPs)**. HSPs deploy wireless access points in public places, such as shopping malls and airports, but HSPs are not Internet service providers (ISPs). Instead, the HSP contracts with one or more ISPs, and offers customers one or more service plans to choose from when they want to establish an account for Internet access.

- **Wireless Internet service provider (WISPs)**. WISPs are ISPs that either deploy Wi-Fi hotspots in public places or outsource Wi-Fi hotspot services to an HSP.

- **Enterprises**. Enterprises can use WPS technology to provide managed guest access on their networks. The WISP scenarios in this paper apply to enterprises as well as WISPs.

**Key WPS components**

The following section introduces key WPS client and server components.

**WPS technology in Windows XP with SP2**

Wireless Provisioning Services technology is included in Service Pack 2 for Windows XP. WPS enables a wireless client computer running Windows XP Home Edition with SP2, Windows XP Tablet PC Edition with SP2, or Windows XP Professional with SP2, to connect to and download network configuration information from a provisioning server. After the Windows XP with SP2 client has obtained network configuration information, it automatically configures the connection to your network.

Windows XP with SP2 WPS technology consists of the following two components:

- **Network Provisioning Service**. The Network Provisioning Service automatically downloads XML configuration files from provisioning servers. Users and administrators do not need to configure the Network Provisioning Service because the service automatically configures itself.

- **Wireless Zero Configuration service**. The Wireless Zero Configuration service in Windows XP with SP2 has new WPS technology capabilities. Wireless Zero Configuration, also called Wireless Auto Configuration in this paper, interacts with the Network Provisioning Service and IEEE 802.1X authentication on the client computer to provide WPS functionality. Users and administrators do not need to configure the Wireless Auto Configuration service because the service automatically configures itself.

In all uses of WPS technology, computer configuration is performed in the background and is transparent to users, with no client computer configuration needed by administrators.

WPS technology on Windows XP with SP2 also supplies your customers with a sign-up wizard that allows them to create an Internet access account. The wizard passes the customer's personal data, such as credit card information, over a secure connection to a custom application that processes the information and creates the user account on your network.

**The provisioning server**

A provisioning server is a computer running Internet Information Services (IIS) or a third-party Web server that maintains a collection of information files that are used to configure client computers during the connection and account sign-up process. These information files are created using Extensible Markup Language (XML) and WPS XML schemas, and are stored on the provisioning server. The provisioning server supplies the XML files to clients when client computers request provisioning information from the provisioning server.

There are multiple XML schemas for WPS that allow you to create XML data files to define network configuration and other parameters for client computers connecting to your network. Using WPS XML data files, you can customize and define the sign-up experience users will have when connecting to your network by using the sign-up wizard included in Windows XP with SP2. In addition to the network and security settings, this includes branding information (such as your company logo), location information (where your Wi-Fi hotspots are located), plan offering (types of accounts your customers can purchase), and Help content to assist your customers when they need additional information.

There are two methods you can use to create and configure your XML master file and subfiles:

- Use the WPS Authoring Tool to create a WPS project and publish your XML master file and subfiles to the provisioning server. The WPS Authoring Tool has a graphical user interface and is designed to assist you in accurately producing and

managing a collection of XML files for your WPS solution. Using the WPS Authoring Tool to create your XML data files is recommended, as it is based on the XML schema and allows you to validate your XML files against the XML schema before using the XML files in a production environment. Download the WPS Authoring Tool at http://go.microsoft.com/fwlink/?LinkId=40535.

- Use the XML schema to manually create your XML master file and subfiles. After you have created these files, you can enter information specific to your network and deployment parameters. For example, where the location of the provisioning server is required, you can provide an HTTPS URL. In another example, you might need to enter your domain name in several places; you can examine the schemas and example files and determine where to insert your domain name.

For more information, see "XML Schemas" later in this document.

In some scenarios depicted in this paper, the provisioning server maintains an account processing application in addition to storing and providing clients with XML configuration files. Other scenarios place the account processing application on a dedicated server. The account processing application is a Web application that you create based on your business model and the requirements of WPS technology.

The account processing application, whether installed on the provisioning server or a dedicated server, processes XML documents sent from client computers when customers create an Internet access account. The XML documents passed from client to server contain data provided by the customer, such as the customer name, address, and credit card information. When promotion codes are used, these codes are also passed to the account processing application from the client in an XML document. The account processing application processes all of the data provided by the client, and performs the appropriate action. For example, the account processing application can verify promotion codes against a Microsoft® SQL Server™ 2000 database, create a user account in an Active Directory® directory service user accounts database, and perform financial functions such as verify credit card information and charge the customer's credit card.

**The IAS server**
Internet Authentication Service (IAS) is the Microsoft implementation of Remote Authentication Dial-In User Service (RADIUS) server and proxy. IAS in Windows Server 2003 with SP1 is a component of all WPS scenarios depicted in this paper. In some scenarios, an IAS proxy is also required.

**Key WPS processes**
The following section introduces key processes that occur during the provisioning of clients and the network connection process that allows your customers to create an account.

**Pre-provisioning the client computer**
Pre-provisioning occurs when client computers are provisioned before arriving at a Wi-Fi hotspot. There are three possible methods for pre-provisioning a client computer:

- Computer original equipment manufacturers (OEMs) pre-provision clients. OEMs can include promotional offers for WISP connectivity with the sale of their computers. This allows purchasers of the OEM's products to arrive at the Wi-Fi hotspot for the advertised WISP with their computer already configured.

- The Information Technology (IT) department at an organization pre-provisions clients. Before supplying employees with new computers, the IT department can pre-provision the client for employees.

- Customers connect to the WISP and pre-provision their computer. WISP customers have the option of creating their account and downloading network configuration information before arriving at the location where they will wirelessly access the Internet through your network. For example, a customer preparing for a business trip might establish an account online and download network configuration files before leaving their office or home.

**Provisioning the client computer**
Client computers are provisioned with your network configuration at a Wi-Fi hotspot.

When new customers connect at a Wi-Fi hotspot, IAS sends a packet containing the location of the provisioning server to Windows XP on the client computer. Windows XP then downloads network configuration information from your provisioning server, and the client computer is automatically configured to access your network.

**Phased network access**
Phased network access occurs when a new customer arrives at your Wi-Fi hotspot with a computer running Windows XP with

SP2. Phased network access consists of the following two stages:

1. Customers are allowed to connect to your network and authenticate as guest to establish an account. While the customer is connected as guest, they do not have access to the Internet. During this first connection phase, customers can create and pay for a new account. When the account is established, guest access is terminated by WPS.

2. Customers are automatically reauthenticated with the newly established account credentials. When the customer is authenticated and authorized with new account credentials, they are provided with access to both your network and to the Internet.

Phased network access is accomplished by temporarily isolating the client computer from the rest of your network using either a virtual local area network (VLAN)-aware gateway device (for example, an access controller or a VLAN-aware router or switch) or IP filters applied to the connection by IAS and wireless access points that provide this capability.

Client isolation is necessary to provide security and to prevent users from accessing the Internet through your network without first establishing a paid account.

When new customers first connect to your network, they are allowed access to your provisioning server and any other necessary network resources (such as your DHCP server), but their access to the Internet is blocked. After they create and pay for an account and reauthenticate with the new account credentials, they are allowed access to the Internet.

If you have deployed IP filters to isolate client computers from the Internet during the first connection phase, customers are granted Internet access upon reauthentication because the IP filters are not applied to the connection. If you have deployed VLANs to isolate client computers from the Internet during the first connection phase, customers are granted Internet access upon reauthentication because they are placed on a VLAN that provides access to the Internet.

**Scenarios overview**

The method that you use to isolate client computers affects the hardware and software configuration of your network. In the first section of this paper, the following scenarios are depicted:

- **WPS technology for a WISP with VLANs**. This scenario depicts a WISP network using VLANs for client isolation, and is recommended for WISPs and HSPs deploying WPS technology.

- **WPS technology for the Enterprise**. This scenario depicts a secure WISP network using VLANs for client isolation, and is recommended for enterprises deploying WPS technology. This scenario is similar to the first scenario, however it includes a perimeter network between the Wi-Fi hotspots and the enterprise local area network (LAN).

In the last two sections of this paper, two untested beta scenarios using IP filters for client computer isolation are presented in overview:

- A WISP using IP filters for client isolation

- An HSP using IP filters for client isolation

# Configuring IAS for WPS Technology

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For all scenarios in this paper, whether you are deploying IAS as a RADIUS proxy or IAS as a RADIUS server, you must configure IAS to be compatible with WPS technology by creating the **EnableWPSCompatibility** registry entry. When you configure and enable this registry entry, and then open a connection request policy profile in the IAS console, connection request policy user interface elements that allow you to configure Protected Extensible Authentication Protocol (PEAP) become visible.

**To configure EnableWPSCompatibility**

1. Open **Registry Editor**.
2. Browse to the following registry path:

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy**

3. Right-click the **Policy** key, click **New**, and then click **DWORD value**.
4. A new value is added to the details pane, with the default name of the value highlighted for editing. Replace the default name by typing **EnableWPSCompatibility**, then press **Enter**.
5. Right-click **EnableWPSCompatibility**, and then click **Modify**.
6. In **Edit DWORD value**, in **Value data**, change the integer to **1**. The default value is **0** (disabled). All values other than **1** (enabled) are treated as **0** (disabled).

**To verify that EnableWPSCompatibility is enabled**

1. Open the IAS console.
2. Double-click **Connection Request Processing**, and then click **Connection Request Policies**.
3. In the right pane, double-click the default connection request policy, named **Use Windows authentication for all users**, and then click **Edit Profile**. The **Edit Profile** dialog box opens.
4. On the **Authentication** tab, below Authenticate requests on this server, you can see the Protected EAP check box.

**To verify that EnableWPSCompatibility is disabled**

1. Open the IAS console.
2. Double-click **Connection Request Processing**, and then click **Connection Request Policies**.
3. In the right pane, double-click the default connection request policy, named **Use Windows authentication for all users**, and then click **Edit Profile**. The **Edit Profile** dialog box opens.
4. On the **Authentication** tab, below Authenticate requests on this server, you cannot see the Protected EAP check box.

For additional IAS configuration steps, see the deployment scenarios in this paper.

# WPS Technology for a WISP with VLANs

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization is a WISP, you can deploy WPS technology using a VLAN-aware gateway device, such as an access controller, a VLAN-aware router, or a VLAN-aware switch. In this circumstance, network resources such as the provisioning server and the IAS server reside on a Network Resource VLAN, while access to the Internet is provided to customers who have established an account by switching them to an Internet VLAN.

In the sections that follow, the components of a WISP network using a VLAN-aware gateway device are described, how the components work together during a new customer sign-on are detailed, and how to deploy a WISP network with VLANs is explained.

**Components of WPS technology with VLANs**
This deployment scenario, designed for an ISP that deploys Wi-Fi hotspots as a WISP, has the following features:

- A VLAN-aware gateway device is used for client computer isolation during the account sign-up process.
- Customers sign up using a promotion code. Many organizations introduce new services through the use of promotional campaigns in which customers are provided with a special or discounted offer when they sign up for an account using a promotion or pre-paid code. This implementation of WPS technology depicts a WISP that uses promotion or pre-paid codes for a spontaneous customer sign-up at a Wi-Fi hotspot. The codes are stored in a database on a server running SQL Server 2000 or a third-party database application.

The following illustration depicts the components of a WISP network using a VLAN-aware gateway device for client computer isolation.




**Components of a WISP network using VLANs**

**Wi-Fi hotspot components**
Following are the components that comprise the wireless local area network (WLAN):

**Wireless client**
A computer running Windows XP Home Edition with SP2, Windows XP Professional with SP2, or Windows XP Tablet PC Edition with SP2. The computer must be equipped with a wireless network adapter that provides support for IEEE standard 802.11, IEEE standard 802.1X authentication, and Wired Equivalent Privacy (WEP). Support for Wi-Fi Protected Access (WPA) is preferred, but not required.

**Wireless access point (RADIUS client)**
One or more wireless access points deployed on the WISP LAN with a wired connection to the access controller, VLAN-aware switch or router, or other gateway device.

The wireless access point is configured as a RADIUS client to the Internet Authentication Service (IAS) server deployed on the WISP LAN. The wireless access points used for WPS technology must meet the following requirements:

- Support for the use of VLANs.

- Support for the IEEE standard 802.1X authentication.

- Support for Wi-Fi Protected Access (WPA) is preferred. WPA is supported by Windows XP with SP2. To deploy WPA, use wireless network adapters and wireless access points that also support WPA.

- Support for RADIUS authentication and RADIUS accounting, including:

    - Support for the Class attribute as defined in RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)," to allow session correlation for RADIUS authentication and accounting records. For session correlation, when you configure RADIUS accounting at your IAS server or proxy, you must log all accounting data that allow applications (such as billing applications) to query the database, correlate related fields, and return a cohesive view of each session in the query results. At a minimum, to provide session correlation, you must log the following IAS accounting data: NAS-IP-Address; NAS-Identifier (you need both NAS-IP-Address and NAS-Identifier because the access server can send either attribute); Class; Acct-Session-Id; Acct-Multi-Session-Id; Packet-Type; Acct-Status-Type; Acct-Interim-Interval; NAS-Port; and Event-Timestamp.

    - Support for accounting interim requests, which are sent periodically by some access servers during a user session, that can be logged. This type of request can be used when the Acct-Interim-Interval RADIUS attribute is configured to support periodic requests in the remote access profile on the IAS server. The access server, in this case a wireless access point, must support the use of accounting interim requests if you want the interim requests to be logged on the IAS server.

    - Support for IP address range filtering.

    - Support for dynamic retransmit timeout (RTO) estimation or exponential backoff to handle congestion and delays in a wide area network (WAN) environment.

In addition, there are some filtering features that the access points must support to provide enhanced security for the network. These filtering options include:

- **DHCP filtering.** The access point must filter on IP ports to prevent the transmission of DHCP broadcast messages in the instance that the client is a DHCP server. The access point must block the client from sending IP packets from port 68 to the network.

- **DNS filtering.** The access point must filter on IP ports to prevent a client from performing as a DNS server. The access point must block the client from sending IP packets from port 53 to the network.

**WISP LAN components**
Following are the components that comprise the WISP LAN.

**VLAN-aware gateway device**
The VLAN-aware gateway device can be an access controller, a VLAN-aware router, a VLAN-aware switch, or any other device that can be configured to apply IAS-provided parameters to client connections. The VLAN-aware gateway device is configured with two VLANs: a Network Resource VLAN and an Internet VLAN.

The Network Resource VLAN allows all users access to the provisioning server and DHCP server. This VLAN grants access to network resources that allow customers to connect to your network as guest, create an account, receive provisioning information from the provisioning server, and pay for the account.

The Internet VLAN provides access to the Internet. Only customers who have created and paid for accounts are switched to this VLAN and granted Internet access. This process occurs when Windows XP reauthenticates the user with the newly created account information. When the user is authenticated and authorized by your IAS server, the IAS server returns attributes to the VLAN-aware wireless access point; the access point instructs the VLAN-aware gateway device to route client traffic to the Internet VLAN.

**Provisioning server**
The WISP provisioning server is configured with the following components.

**HTTPS Web server**

The Internet Information Services (IIS) or third-party Web server must be deployed with Secure Hypertext Transfer Protocol (HTTPS).

**Web application**
The WISP Web server is configured with an account processing Web application that processes data provided during customer sign-up or account renewal. When a customer uses the sign-up wizard on a client computer to create and pay for a WISP account, the customer enters data, such as name, address, and credit card information that is converted to an XML document on the client. Windows XP sends this XML document to the WISP provisioning server.

The account processing application on the provisioning server must be capable of accepting and processing the XML documents containing the user data. For example, the account processing application must compare the promotion code entered by the user to a promotion code in the SQL Server database, and then dynamically create an account in the Active Directory user accounts database with the properties (domain and security group membership) described in the database. The account processing application must also contain a credit card verification component to process customers' payment information. The account processing Web application must permit new customers to sign up and to permit existing customers to renew their subscriptions for service.

**XML master file and subfiles**
The WISP provisioning server stores the XML master file and subfiles that provide the client with all configuration information needed to access the network, create an account, pay for the account, and ultimately access the Internet. For more information about the XML master file and subfiles, see "XML Schemas" in this paper.

**Server certificate**
For server authentication to client computers, the WISP provisioning server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and be issued by a public certification authority (CA), such as Verisign or Thawte, that is trusted by client computers. The Trusted Root Certification Authorities certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust. For more information, see "Server Certificate Requirements" in this paper.

In test lab deployments of WPS technology, you can deploy your own certification authority in lieu of using a public trusted root CA. In this circumstance you must install the private trusted root CA certificate on all clients used for testing so that the clients will trust the CA and so that your servers, such as your provisioning server and your IAS server, can be successfully authenticated by the clients.

**Domain controller and IAS server**
The WISP domain controller and IAS server is running Windows Server 2003, Standard Edition with SP1; Windows Server 2003, Enterprise Edition with SP1; or Windows Server 2003, Datacenter Edition with SP1, and is configured with the following components.

**Active Directory**
In this scenario, Active Directory is deployed. The user accounts database on the domain controller must be an Active Directory user accounts database or a database that uses Lightweight Directory Access Protocol (LDAP) and supports dynamic creation of user accounts.

When a customer signs up for an account, the account processing Web application on the provisioning server creates a new account in the user accounts database, and adds the user to a group that has clearly defined access privileges that match the type of account the customer purchased when signing up.

If you use an accounts database other than Active Directory, IAS authentication and authorization extensions must be written and installed for this process to function correctly.

**Internet Authentication Service (IAS)**
IAS is the Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy, and is used to authenticate and authorize users connecting to your network. IAS is configured with remote access policies that allow guest authentication for non-domain member computers and users. It is also configured to provide attributes to RADIUS clients (access points) that instruct the gateway device to apply the attributes to client connections. Protected Extensible Authentication Protocol (PEAP) with MS-CHAP v2 is configured on remote access policies as the authentication method used by wireless clients.

**Extension DLL and URL PEAP-TLV**
An IAS extension DLL defining a URL PEAP-TLV provides IAS with the ability to send the location of the provisioning server to client computers.

PEAP-Type-Length-Value (PEAP-TLV) is an Extensible Authentication Protocol (EAP) authentication type that allows the IAS server to pass information to client computers attempting to connect to your network.

In this circumstance, the value contained in the PEAP-TLV is an HTTPS Uniform Resource Locator (URL) that provides client computers running Windows XP with SP2 with the location of the WISP provisioning server. With this URL, Windows XP can download the WISP XML files to the client computer.

In addition to the URL of the provisioning server, the URL PEAP-TLV includes an action parameter. The action parameter directs Windows XP to perform a specific task. The action parameter included in the URL PEAP-TLV defines tasks such as new customer sign-up, existing account renewal, and password change.

For more information, see "How to Create an IAS Extension DLL and a URL PEAP-TLV" in this paper.

**Server certificate**
To authenticate the IAS server to the wireless client computers using PEAP, the WISP IAS server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and be issued by a public CA, such as Verisign or Thawte, that is trusted by client computers. The Trusted Root Certification Authorities certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust. If you install IAS and Active Directory on the same computer, the computer must have a certificate. If you install IAS and Active Directory on different computers, only the IAS server needs a certificate.

**SQL server**
A computer running SQL Server 2000 or another SQL-compatible relational database application.

The promotion code database running on the SQL server is configured with the following fields: promotion code, user name, domain name, security group, and expiration date. With the exception of the user name field, each field for each record is preconfigured with a value. The value for the user name field is assigned by the Web application when a customer creates a user account with a promotion code that matches a value in the promotion code field in the database. By predefining the domain in which the user account is created by the Web application and the Active Directory security group to which the user account is joined as a member, you can assign network access and other permissions for your customers.

**DHCP server**
The DHCP server must be able to assign valid public IP addresses to computers accessing the network through the wireless access points.

**How a WISP works with VLANs**
The Internet connection process using WPS technology with a VLAN-aware gateway device differs depending on whether the customer attempting to connect is a new customer or an existing customer. The following example describes the process for a new customer using a promotion code at a Wi-Fi hotspot location. In addition, the manner in which IAS handles an expired account is explained.

**New customer connection example**
When a new customer connects to a WISP and establishes an account, the following stages occur:

1. The customer discovers the WISP network at a Wi-Fi hotspot
2. The customer authenticates as guest
3. The client is provisioned and the customer establishes an account
4. The customer is authenticated with the new account credentials
5. The user is switched to a VLAN that provides Internet access

In the next section we will look at these stages in more detail.

**1. The customer discovers the WISP network at a Wi-Fi hotspot**
When a customer arrives at the WISP Wi-Fi hotspot with a portable computer running Windows XP Home Edition with SP2, Windows XP Tablet PC Edition with SP2, or Windows XP Professional with SP2, the computer comes within range of the WISP access point beacon.

Wireless auto configuration on the client computer detects the beacon information from the access point, which is configured to broadcast the Service Set Identifier (SSID) of the WISP's network. The SSID is equivalent to the network name.

The customer is informed by Windows XP that a wireless network is available. In this example, the customer possesses a promotion code to use for account establishment, and proceeds by clicking **Connect**.

**2. The customer authenticates as guest**

Wireless auto configuration uses 802.1X and PEAP guest authentication to connect to the WISP network through the access point, automatically passing a null user name and a blank password to the WISP IAS server. The access point is connected to a VLAN-aware gateway device that allows traffic from the client to pass through the Network Resource VLAN, but blocks the client from access to the Internet VLAN.

The IAS server is the PEAP authenticator and TLS endpoint for customers who connect as guest. The TLS tunnel is created between the wireless client and the IAS server. All subsequent messages between wireless client and IAS server pass through this tunnel, which traverses the access point and the gateway device.

Server authentication is performed when the IAS server verifies its identity to the client computer using a certificate that contains the Server Authentication purpose in Enhanced Key Usage (EKU) extensions. This certificate is issued by a public CA that the client computer trusts.

The IAS server authenticates and authorizes the customer as guest. In the Access-Challenge message that the IAS server sends to the client is a URL PEAP-TLV attribute. The URL PEAP-TLV message contains the URL of the provisioning server. This URL provides the client with the location of the XML master file.

The wireless client computer receives an IP address lease from the DHCP server. The address is from an IP address range configured in a scope on the DHCP server. In addition to the IP address, the client receives DHCP options, such as a default gateway and a DNS server IP address.

**3. The client is provisioned and the customer creates an account**

The XML master file on the provisioning server contains pointers to the XML subfiles. The wireless client creates an HTTPS connection with the provisioning server and downloads the XML master file and subfiles. When the XML sign-up file is downloaded, the sign-up wizard is started on the wireless client to allow the customer to create and pay for an account with the WISP.

Using the sign-up wizard on the wireless client computer, the customer steps through the process of signing up for an account. The customer enters the promotion code as well as personal data such as name, address, and credit card number. The data entered by the customer is converted into an XML document.

The XML document containing the customer's sign-up data is sent to the Web application on the WISP provisioning server.

The Web application checks the promotion code entered by the user against the promotion code database on the SQL server. If the promotion code is valid, the Web application continues processing the customer's data.

The Web application processes the customer payment information. After payment is verified and sign-up information is completed successfully, the Web application reads the domain and security group information from the promotion code database on the SQL server and creates a user account in Active Directory and adds the account to the security group. The Web application also enters the new user name in the promotion code database.

An XML document containing the new account credentials is sent from the WISP provisioning server to the client computer. The client computer uses the credentials to configure wireless auto configuration and 802.1X under the name of the WISP.

**4. The customer is authenticated with the new account credentials**

Wireless auto configuration restarts the association to the SSID for the WISP.

Wireless auto configuration finds the correct 802.11 profile, which was downloaded with the other WISP information. Wireless auto configuration re-associates with the access point using the correct profile.

802.1X restarts the authentication process using PEAP-MS-CHAP v2 and the new account credentials.

As the client starts the authentication process with PEAP-MS-CHAP v2 authentication, a TLS channel is created between the customer's client computer and the WISP IAS server.

In the second stage of PEAP-MS-CHAP v2 authentication, the WISP IAS server authenticates and authorizes the connection request against the new account in the Active Directory user accounts database. The IAS server sends an Access-Accept message to the access point. Included in the Access-Accept message are attributes that specify which VLAN the customer can access.

**5. The customer is switched to a VLAN that provides Internet access**

The access point instructs the gateway device to assign the client to the Internet VLAN rather than the Network Resource VLAN. In addition, 802.1X on the access point opens the virtual port to provide Internet access to the client.

The wireless client computer receives an IP address lease from the DHCP server. The address is from an IP address range configured in a scope on the DHCP server. In addition to the IP address, the client receives DHCP options, such as a default

gateway and a DNS server IP address.

The wireless client computer can now access the Internet.

**How IAS handles an expired account**
You can determine the types of account plans that you want to offer your customers. These plans can range from fees based on hourly use to accounts with life spans as long as a day, a month, or longer.

It is important for IAS to determine whether a connecting or connected client computer has a valid account, and to take the appropriate action if the customer's account is expired. The following example illustrates how IAS determines that a 24-hour account is current and how WPS technology behaves when the account expires.

**Twenty-four hour connect option example**
When the customer arrives at the WISP, the customer chooses an access account that has a one-day (24-hour) lifespan. The customer and client computer proceed through the account creation process described above, and then connect to the Internet. The following process occurs:

1. In the Access-Accept message sent by the IAS server, the IAS server sets a session timeout of 60 minutes for the client computer connection to the access point.

2. After 60 minutes, the access point requests that the client reauthenticate. The client reauthenticates successfully and the customer's session is not interrupted.

3. Each 60 minutes thereafter, the access point requests that the client reauthenticate. During each authentication the IAS server checks the current time against the expiry time for the user account to discover whether the customer is authorized to access the network.

4. On the last re-authentication, at hour 23 in the account lifespan and before 24 hours have passed, the IAS authorization check fails and the IAS server sends a URL PEAP-TLV message to the client that contains the account renewal action parameter and an HTTPS URL for an XML master file. The URL PEAP-TLV supplies the customer with the location of the provisioning server where the customer can renew the account.

5. Upon receiving the URL in the URL PEAP-TLV message, 802.1X requests that Windows XP display the account renewal application to the customer.

6. The customer renews the account and 802.1X initiates authentication using the new account credentials.

7. During authentication with the IAS server, the IAS server authenticates and authorizes the customer against the user accounts database, and sends an Access-Accept message containing a session timeout of 60 minutes to the access point.

8. During this process, because the account has not expired, the customer maintains connection to the Internet.

If the customer does not complete the renewal process before the 24-hour account lifespan is reached, authentication fails and customer access to the Internet is terminated. When authentication fails, Windows XP attempts authentication as guest. The VLAN-aware gateway device is configured to allow the connection to the Network Resource VLAN, and the customer is provided with the option of renewing the account for continued access.

**How to deploy WPS technology with VLANs**
In the set of instructions that follow, these assumptions are made:

1. The computer functioning as your IAS/RADIUS server has Windows Server 2003 with SP1 installed, and the **EnableWPSCompatibility** registry key is enabled according to the instructions in "Configuring IAS for WPS technology" in this paper.

2. Client computers are running Windows XP Home Edition with SP2; Windows XP Professional with SP2; or Windows XP Tablet PC Edition with SP2.

3. All of your hardware, including the VLAN-aware gateway device and VLAN-aware wireless access points, meet all of the technical requirements stated in "Components of WPS technology with VLANs" in this paper.

4. You have already deployed a computer running SQL Server 2000 or a third-party database program on your network. For information about SQL Server 2000, see SQL Server at http://go.microsoft.com/fwlink/?LinkId=20014.

5. You have SQL Server 2000 or third-party relational database development experience and you understand how to use SQL Server 2000 or a third-party database program to create, modify, administer, and manage your databases.

6. You have experience deploying an Internet Information Services (IIS) or third-party Web server with HTTPS. If you are deploying IIS, you understand how to use Active Server Pages (ASP), and you can develop applications using Microsoft .NET Framework 1.1.

7. You have software development experience that allows you to create your custom Web application that will be installed and run on the provisioning server.

8. You have software development experience that allows you to create an IAS extension DLL.

| **Note** |
| --- |
| The instructions that follow use four servers upon which various programs and services are installed. If you deploy WPS technology in a test lab, you can reduce or increase the number of servers in a manner appropriate to your available hardware resources. For example, in a test lab environment you might want to deploy Active Directory, IAS, IIS, and DHCP on the same server. |

To deploy WPS technology with VLANs, the basic steps are as follows:

1. Configure the domain controller and IAS server

2. Configure the DHCP server

3. Install and configure your VLAN-aware gateway device

4. Install and configure your wireless access points

5. Configure RADIUS clients in IAS

6. Create your Web application

7. Configure the provisioning server

8. Configure XML master and subfiles

9. Configure the database on the SQL server

10. Configure the Windows XP–based client computer

11. Configure certificates in a test lab environment (optional)

**Configure the domain controller and IAS server**

Install Windows Server 2003, Standard Edition with SP1; Windows Server 2003, Enterprise Edition with SP1; or Windows Server 2003, Datacenter Edition with SP1 on a computer that meets or exceeds the minimum hardware requirements for the respective operating system. After the operating system is installed, you can perform General configuration, Active Directory configuration, IAS configuration, and IAS extension DLL & URL PEAP-TLV configuration.

**General configuration**

1. Assign the server a static IP address. Determine the IP address range for the WISP LAN, and determine how many devices need static IP addresses from that range. For example, the domain controller and IAS server must have a static IP address, and it is generally a good idea to provide static IP addresses to other key servers, such as the DHCP server and WISP

provisioning server. IP addresses that are statically assigned are excluded from distribution by a DHCP server through definition of an exclusion range on the DHCP server. The exclusion range falls at the beginning of the IP address range, so choose one of the addresses from the beginning of the range that you plan on excluding when you configure DHCP, and assign it to the domain controller by using Network Connections. For more information, see "To configure TCP/IP for static addressing." At http://go.microsoft.com/fwlink/?LinkId=20015.

2. Install a server certificate obtained from a public trusted root certification authority, such as a certificate from Verisign.

   For more information, see "Obtaining and Installing a VeriSign WLAN Server Certificate for PEAP-MS-CHAP v2 Wireless Authentication" at http://go.microsoft.com/fwlink/?LinkId=33675.

   When you obtain the certificate, it must conform to the minimum server certificate requirements described earlier in this paper.

   For more information, see "Network access authentication and certificates" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20016.

| **Note** |
| --- |
| If you are deploying WPS technology in a test lab environment, you can install a private CA rather than obtain certificates from a public CA. For more information, see "Configure certificates in a test lab environment (optional)" in this paper. |

**Active Directory configuration**

1. Install Active Directory and DNS. To install Active Directory, open Command Prompt, type **dcpromo**, and then follow the instructions provided in the wizard, entering your network configuration information in the wizard as you progress.

2. Design and create your security groups. When users sign up and create an account, your Web application adds the new user account as the member of a security group that you create in this step. The Web application chooses group membership based on the value of the security group field in the promotion code database on your SQL Server, so you need to match the security group you create to the security group field in the SQL server database. If you have the need for multiple security groups, you can assign permissions to each group individually. For more information, see "To create a new group" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20018 and "Assign user rights to a group in Active Directory" at http://go.microsoft.com/fwlink/?LinkId=20019.

   | **Important** |
   | --- |
   | The security groups you create in Active Directory are named and used in the SQL Server database and in IAS remote access policy. |

3. Enable the Guest account in Active Directory. If guest access is not enabled in Active Directory, new customers cannot access your provisioning server by authenticating as guest. To enable the Guest account, open the Active Directory Users and Computers snap-in, and then double-click **Users**. Right-click the account named **Guest**, and then click **Enable Account**.

To perform the next procedure, you must be a member of either the Domain Admins group in the domain for which you want to raise functionality or the Enterprise Admins group in Active Directory; or you must have been delegated the appropriate authority.

Raise the domain functional level to either Windows 2000 native or Windows Server 2003 by doing the following:

1. Open the Active Directory Domains and Trusts snap-in. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.

2. In the console tree, right-click the domain for which you want to raise functionality, and then click **Raise Domain Functional Level**.

3. In **Select an available domain functional level**, do one of the following:

a. To raise the domain functional level to Windows 2000 native, click **Windows 2000 native**, and then click **Raise**.

b. To raise domain functional level to Windows Server 2003, click **Windows Server 2003**, and then click **Raise**.

| **Important** |
| --- |
| If you have or will have any domain controllers running Windows NT 4.0 and earlier, then do not raise the domain functional level to Windows 2000 native. After the domain functional level is set to Windows 2000 native, it cannot be changed back to Windows 2000 mixed.<br><br>Likewise, if you have or will have any domain controllers running Windows NT 4.0 and earlier or Windows 2000, then do not raise the domain functional level to Windows Server 2003. After the domain functional level is set to Windows Server 2003, it cannot be changed back to Windows 2000 mixed or Windows 2000 native. |

The current domain functional level is displayed under **Current domain functional level** in the **Raise Domain Functional Level** dialog box.

For more information, see "Domain and forest functionality" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?linkid=30600.

For information about configuring Active Directory replication, see "Active Directory Replication" in this paper.

**IAS configuration**
For IAS, the three configuration stages are General configuration, Remote access policy configuration, and Connection request policy configuration. (RADIUS client configuration occurs later in the overall WISP deployment process.)

**General configuration**

1. Install Internet Authentication Service.

   For more information, see "To install IAS" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20028.

2. Register IAS in Active Directory. In order for IAS to have permission to read user accounts in Active Directory, IAS must be registered in Active Directory.

   For more information, see "To enable the IAS server to read user accounts in Active Directory" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20030.

3. Delete the default remote access policies. To delete the policies, open the IAS console, and then click **Remote Access Policies**. Select each existing policy, right-click the policy, and then click **Delete**.

4. Create your IAS extension DLL. For more information, see "How to Create an IAS Extension DLL and a URL PEAP-TLV" in this paper.

5. Install the DLL on your IAS server and configure DLL registry keys according to your needs.

   To install your DLL, do the following:

   - Open Command Prompt and change directories to the folder that contains your DLL.
   - Type the following: **regsvr32***DLL_name.dll*, where *DLL_name.dll* is the name of your DLL file.

**Remote access policy configuration**
There are two remote access policies configured for WPS technology. The Guest access policy provides network parameters and rules for users connecting as guest. The Valid Users access policy provides network parameters and rules for users who have valid WISP accounts.

| Note |
| --- |
| If you have a variety of account types that you offer to customers and these accounts have different properties (such as membership to different security groups), you might find it necessary to create more than two remote access policies on your IAS server. If this is the case, you can use the remote access policies described below to extrapolate how to create additional policies. |

**To configure the Guest access policy**

1. Open the Internet Authentication Service console and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**.

3. Use the **New Remote Access Policy Wizard** to create a policy. For the WISP guest access policy, you can choose the following:

    a. For **How do you want to set up this policy?** select **Use the wizard to set up a typical policy for a common scenario**.

    b. For **Policy name**, type **Guest access** (or type another name for your policy that you prefer).

    c. For **Select the method of access for which you want to create a policy**, click **Wireless**.

    d. For **Grant access based on the following**, click **User**.

    e. In **Select the EAP type for this policy**, select **Protected EAP (PEAP)**, and then click **Configure**.

    f. In **Certificate issued**, select the certificate that you want the IAS server to use to verify its identity to client computers. Also select the **Enable Fast Reconnect** check box.

After you have completed creating the policy and have closed the wizard by clicking **Finish**, you need to perform additional policy configuration.

1. In the IAS console, click **Remote Access Policies**, and then double-click the policy you just created. Make the following configuration changes to the policy:

2. In the policy **Properties** dialog box, for **Policy conditions**, click **Add**.

3. In **Attribute Types**, click **Day-And-Time-Restrictions**, and then click **Add**. In **Time of day restraints**, select **Permitted**, configure the days and times that access is permitted, and then click **OK**.

4. In the policy Properties dialog box, click **Grant remote access permission**.

5. Click **Edit Profile**. On the **Authentication** tab, in **Unauthenticated access**, click **Allow clients to connect without negotiating an authentication method**.

**To configure the Valid Users access policy**

1. Open the Internet Authentication Service snap-in and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**.

3. Use the **New Remote Access Policy Wizard** to create a policy. For the WISP Valid users access policy, you can choose the following:

    a. For **How do you want to set up this policy?** verify that **Use the wizard to set up a typical policy for a common scenario** is selected.

b. For **Policy name**, type **Valid Users** (or type another name for your policy that you prefer).

c. For **Select the method of access for which you want to create a policy**, click **Wireless**.

d. For **Grant access based on the following**, click **Group**, and then click **Add**. In **Enter the object name to select**, type the name of a security group that you defined when configuring Active Directory. For example, if you created a group named **Valid Users**, type **Valid Users**, and then click **OK**.

| **Important** |
| --- |
| The following three items must match: the name of the security group in Active Directory, the value of the **security group** field in the SQL Server database, and the name of the security group configured in the Valid Users access policy in IAS. The Web application uses the value of the SQL Server database **security group** field to determine group membership for new accounts. |

e. In **Select the EAP type for this policy**, select **Protected EAP (PEAP)**, and then click **Configure**.

f. In **Certificate issued**, select the certificate that you want the IAS server to use to verify its identity to client computers. Also check the **Enable Fast Reconnect** check box.

After you have completed creating the policy and have closed the wizard by clicking **Finish**, you need to perform additional policy configuration.

1. In the IAS console, click **Remote Access Policies**, and then double-click the policy you just created. Make the following configuration changes to the policy:

2. In the policy **Properties** dialog box, for **Policy conditions**, click **Add**.

3. In **Attribute Types**, click **Day-And-Time-Restrictions**, and then click **Add**. In **Time of day restraints**, select **Permitted**, configure the days and times that access is permitted, and then click **OK**.

4. In the policy **Properties** dialog box, click **Grant remote access permission**.

5. Click **Edit Profile**, and then click the **Advanced** tab. By default, the **Service-Type** attribute appears in **Attributes** with a value of **Framed**. To specify additional connection attributes required for WPS technology with VLANs, click **Add**, and then add the following attributes:

   - **Framed-Protocol**. Value: **PPP**

   - **Tunnel-Medium-Type**. Value: **802 (Includes all 802 media plus Ethernet canonical format)**

   - **Tunnel-Pvt-Group-ID**. Value: Enter the integer that represents the VLAN number for the Internet VLAN. For example, if your access controller's Internet VLAN is VLAN 4, type **4.**

   - **Tunnel-Type**. Value: **Virtual LANs (VLAN)**

   - **Tunnel-Tag**. Value: Obtain this value from your hardware documentation

| **Important** |
| --- |
| IAS evaluates remote access policies in the order in which they appear in the IAS console under **Remote Access Policies**. The Valid Users access policy must be the first policy in the list of remote access policies or valid user authentication will fail. Because IAS places newly created policies in the first position and the Valid Users access policy was the last policy created, the Valid Users access policy should now appear first in the IAS console, with the Guest access policy appearing second. If this is not the case, move the Valid Users access policy into first position. |

**Connection request policy configuration**

By default, there is one connection request policy predefined in the IAS console, called **Use Windows authentication for all users**. This policy can be used for WPS technology.

1. In the **IAS console**, double-click **Connection Request Processing**, click **Connection Request Policies**, and then double-click the policy **Use Windows authentication for all users**.

2. Click **Edit Profile**. The **Edit Profile** dialog box opens.

3. On the **Authentication** tab, click **Authenticate requests on this server**, and then check the **Protected EAP**check box.

4. Click **Configure Certificate**. Select the certificate you want IAS to use to authenticate to clients, and then click **OK** three times to close all dialog boxes and return to the IAS console.

| **Note** |
| --- |
| If you access the profile of a connection request policy in the IAS console and you cannot see the **Protected EAP** check box or the **Configure Certificate** button, you must first configure IAS for compatibility with WPS technology as described in "Configuring IAS for WPS technology" in this paper. |

**Configure the DHCP server**

On a computer running Windows Server 2003:

1. Install DHCP.

   For more information, see "To install a DHCP server" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20034.

2. In the DHCP console, run **New Scope Wizard** twice. Create two VLAN scopes from which IP addresses will be leased to wireless clients connected to the VLANs. Each scope must define a different IP address range using either a private address range or a public IP address range. If you are using network address translation (NAT), you can use a private IP address range; otherwise, the IP addresses leased to wireless clients must be from a public IP address range.

   For more information, see "To create a new scope" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20123.

3. While running New Scope Wizard, create an exclusion range for the IP addresses you will be assigning statically. For example, if you need to statically assign 10 IP addresses from the address range 10.1.1.1 through 10.1.1.254, your exclusion range is defined as 10.1.1.1 through 10.1.1.10.

4. While running New Scope Wizard, assign scope options. On the **Configure DHCP Options** page, select **Yes, I want to configure these options now**. Scope options are applied only to leases of addresses from within the IP address range that the scope defines, which provides flexibility as your network grows. Define the **DNS server** and **Domain name** options, as well as any other options that are appropriate for your network configuration.

5. While running New Scope Wizard, activate the scope. The option to activate the scope while running the wizard is available only if you have chosen to configure DHCP scope options in the previous steps.

   For more information, see "To activate a scope" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20124.

6. Authorize the DHCP server in Active Directory.

   For more information, see "To authorize a DHCP server in Active Directory" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20125.

The DHCP server is now online and able to provide IP address leases to client computers. In some cases you might want to examine the types and durations of accounts you offer your customers and adjust the DHCP lease duration accordingly. In most cases when deploying WPS technology, the default lease duration of eight days is too long and should be shortened considerably.

**DHCP example**
In this example, you are creating two VLANs and two scopes, one scope for each VLAN.

**VLAN 2**
VLAN 2 is the Network Resource VLAN that provides access to network resources (such as the IAS server and DHCP server) for wireless computers connecting as guest. VLAN 2 blocks access to the Internet, however. The DHCP scope for VLAN 2 is defined with the following example parameters:

- **Address range:** 192.168.1.1 through 192.168.1.254. This is a private IP address range. If you are using NAT, you can use this range on your network. If you are not using NAT, use a public IP address range. In addition, if your wireless deployment is large, select an IP address range that provides more IP addresses for lease to clients.

- **Exclusion range:** 192.168.1.1 through 192.168.1.10. By using this exclusion range, the available address pool for clients is 192.168.1.11 through 192.168.1.254. Ten IP addresses are excluded so that you can statically assign these addresses to computers and devices on your network. For example, the router IP address must be statically assigned on the router.

- **DHCP scope option 003, Router:** 192.168.1.1. The router IP address must be an address that falls within the exclusion range so that the DHCP server does not lease the router IP address to a wireless client computer, thereby creating an address conflict.

- **DHCP scope option 006, DNS server:** the IP address of the Active Directory and DNS server on the WISP LAN.

Note that DHCP scope option 003, Router, provides client computers with the IP address of their default gateway IP address. In this case, the default gateway for wireless clients is the VLAN-aware gateway device, whether it is an access controller, a VLAN-aware router, a VLAN-aware switch, or another compatible device. When you configure your VLAN-aware gateway device, you can specify the IP address that the device uses for each VLAN.

In this example, you must configure the VLAN-aware gateway device so that it uses the IP address 192.168.1.1 on VLAN 2.

For your WPS deployment you can configure additional DHCP options as needed. For example, if you are using a WINS server, you can configure your VLAN scopes with DHCP option 044, WINS/NBNS Servers.

**VLAN 4**
VLAN 4 is the Internet VLAN that provides access to the Internet. Users who have successfully created an account are switched to this VLAN after completing the provisioning and sign-up process. The DHCP scope for VLAN 4 is defined with the following example parameters:

- **Address range:** 192.168.2.1 through 192.168.2.254

- **Exclusion range:** 192.168.2.1 through 192.168.2.10

- **DHCP scope option 003, Router:** 192.168.2.1. The router IP address must be an address that falls within the exclusion range so that the DHCP server does not lease the router IP address to a wireless client computer, thereby creating an address conflict.

- **DHCP scope option 006, DNS server:** the IP address of the Active Directory and DNS server on the WISP LAN.

For VLAN 4 in this example, you must configure the VLAN-aware gateway device so that it uses the IP address 192.168.2.1 on VLAN 4.

**Install and configure your VLAN-aware gateway device**
Configure two VLANs on the gateway device: a Network Resource VLAN that provides access to the WISP LAN, and an Internet VLAN that provides access to the Internet.

The remote access policies you create in IAS determine which VLAN your customers can access:

- The Guest access policy places users on the Network Resource VLAN so that they can create and pay for a valid user account.
- The Valid Users access policy in IAS places customers on the Internet VLAN.

Each VLAN has a different IP address range. When configuring your DHCP server, you created a scope for each VLAN, and you defined DHCP scope option 003, Router. This is the IP address commonly referred to as the "default gateway."

You must configure the VLAN-aware gateway device as the default gateway for each VLAN, using an IP address from the IP address range that you defined on your DHCP server.

**Choosing the IP address for the default gateway for each VLAN**
Your VLAN-aware gateway device is the default gateway for both VLANs and is configured with a different IP address for each VLAN.

After you define an IP address range for a VLAN on your DHCP server, you can use any IP address from that range as the IP address for the default gateway. To prevent an IP address conflict, however, exclude some IP addresses from the range for use by devices that you want to configure with a static IP address. When you create an exclusion range for an IP address range, the DHCP server does not lease IP addresses from the exclusion range to DHCP clients.

| **Note** |
| --- |
| IP address conflicts occur when two devices or computers on the same subnet have the same IP address. This situation can occur if you configure a device or computer with a static IP address that is also an address that the DHCP server can lease to DHCP clients. |

Thus ensure that you assign an IP address from the exclusion range as the IP address for the default gateway. The address you use does not need to be the first address in the exclusion range, but it must be an address contained within the exclusion range.

In accordance with the example provided in the DHCP configuration section of this paper, configure the VLAN-aware gateway device so that it uses an IP address from the exclusion range 192.168.1.1 through 192.168.1.10 for VLAN 2. For example, you can configure the VLAN-aware gateway device so that it uses the IP address 192.168.1.1 on VLAN 2.

| **Important** |
| --- |
| In each scope you configure on the DHCP server, the value you enter for DHCP scope option 003, Router, must match the IP address you assign to the VLAN-aware gateway device for use on each VLAN. For example, if you configure a scope on the DHCP server for VLAN 2 with the IP address range 192.168.1.1 through 192.168.1.254, and you assign the DHCP scope option 003, Router, with the value 192.168.1.1, you must configure the VLAN-aware gateway device to use the IP address 192.168.1.1 on VLAN 2. |

Similarly, as described in "Configure the DHCP server" in this paper, configure the VLAN-aware gateway device so that it uses an IP address from the exclusion range 192.168.2.1 through 192.168.2.10 for VLAN 4. For example, you can configure the VLAN-aware gateway device so that it uses the IP address 192.168.2.1 on VLAN 4.

In some circumstances you might prefer to use your VLAN-aware gateway device as the DHCP server for each VLAN. If this is the case, you must define IP address ranges and exclusion ranges on the VLAN-aware gateway device rather than on a DHCP server.

See the product documentation for your VLAN-aware gateway device for information about configuring your hardware.

| **Important** |
| --- |
| The Internet VLAN integer must match the value you configure for the Tunnel-Pvt-Group-ID attribute in the Valid Users access policy on your IAS server. For example, if VLAN 4 leads to the Internet, the value of the Tunnel-Pvt-Group-ID attribute in the profile of the Valid Users access policy must be 4. |

**Install and configure your wireless access points**
Following is an example of how to configure your wireless access points for use with WPS technology. This example depicts configuration of a Cisco 802.1X–compatible access point. If you use a different 802.1X–compatible access point, follow the directions in your product documentation and in "Using other access points" in this paper.

The Cisco access point used in this example provides a Web-based interface for access point configuration. To use this interface for access point configuration, you must physically connect the access point to the network, log on to a computer that has network connectivity to the access point, and then open a Web browser, such as Microsoft® Internet Explorer.

Type the IP address for your 802.1X–compatible access point in the Web browser address bar, and then press ENTER. The access point configuration page appears in the Web browser.

**To configure your wireless access point**

1. On the configuration page, click the **Setup** tab, and then click **Express Setup**.

2. In the **Radio Service Set ID (SSID)** field, type an SSID for the access point, and then click **OK**.

3. Under **Services**, click **Security**, and then click **Radio Data Encryption (WEP)**.

4. On the **AP Radio Data Encryption** page, select the **Open** check boxes for the **Accept Authentication Type** and the **Require EAP** options, and then clear all other check boxes.

5. In the **Encryption Key** box, type a 32-digit WEP key. In the **Key Size** list, select **128 bit**, and then, to update the page, click **Apply**.

6. In the **Use of Data Encryption by Station is** list, select **Full Encryption**, and then click **OK**.

7. On the **Security Setup** page, click **Authentication Server**. The following table shows the values to set on the **Authenticator Configuration** page.

| Item | Value |
|---|---|
| 802.1X Protocol Version (for EAP Authentication) | Draft 10 |
| Server Name/IP | IAS server IP address |
| Server Type | RADIUS |
| Port | 1812 |
| Shared Secret | Type the shared secret you want to use for RADIUS clients |
| Timeout (in seconds) | 20 |
| User server for | EAP Authentication |

8. When you have completed the configuration of the authentication server, click **OK**.

**Using other access points**

If you are using an access point (AP) other than Cisco, follow the directions in your access point documentation using the following guidelines:

- **Authentication or RADIUS server:** Specify your IAS server by IP address or FQDN, depending on the requirements of the AP.

- **SSID:** Specify a Secure Set Identifier (SSID), which is an alphanumeric string that serves as the network name. This name is broadcast by APs to wireless clients and is visible to users at your Wi-Fi hotspots.

- **RADIUS settings:** Use RADIUS authentication on User Datagram Protocol (UDP) port 1812 and use RADIUS accounting on UDP port 1813.

- **Secret or shared secret:** Use a strong shared secret and configure the IAS server with the same shared secret.

- **EAP:** Configure the AP to require EAP from wireless clients.

- **802.1X and WEP:** Enable IEEE 802.1X authentication and WEP.

| **Note** |
| --- |
| Make sure that you use a strong shared secret. For more information about how to create a strong shared secret, see "Shared secrets" in Windows Server 2003 Help and Support Center or on the Web at http://go.microsoft.com/fwlink/?LinkId=41035. |

**Configure RADIUS clients in IAS**
In the IAS console, add each access point on your network as a RADIUS client. In addition, configure the shared secret used between the access points and the IAS server. For more information, see "To add RADIUS clients" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20031 and "To configure the Message Authenticator attribute and shared secret" at http://go.microsoft.com/fwlink/?LinkId=20032.

**Configuring RADIUS clients by IP address range**
If your IAS server is running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, you can configure RADIUS clients by IP address range. This is a useful feature when you have a large number of access points to deploy; if you deploy your access points on the same subnet or VLAN within the same IP address range, configuration of RADIUS clients in IAS is simplified. Instead of individually configuring each access point as a RADIUS client in IAS, you can configure all access points at once using the IP address range of the subnet or VLAN upon which the APs reside. In this circumstance, use the same shared secret for all access points, and make sure that the shared secret is strong.

By contrast, you can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. With Windows Server 2003, Standard Edition, if the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients.

**Create your Web application**
Your Web application must be capable of performing the following functions:

- Communicating with Wireless Provisioning Services using HTTPS.
- Uploading XML master file and subfiles that are stored on the provisioning server to client computers that request the files.
- Accepting and processing XML documents from client computers that contain customer data, such as promotion code, customer name, customer address, and other information.
- Accepting and processing XML documents from client computers that contain credit card information. This includes verifying the credit card and charging the customer account.
- Reading the promotion code database records to validate promotion codes.
- Reading the promotion code database records to determine the domain in which to create a new user account.
- Reading the promotion code database records to determine the security group membership for a new user account.
- Writing a user name to the user name field in the promotion code database.
- Dynamically creating new accounts in Active Directory (or a third-party LDAP-compliant database) using data provided by customers as well as parameters from the promotion code database.

It is recommended that the design of your Web application provides customers with knowledge of their user name and password. Customers can either be allowed to designate their own user name and password or this information can be provided to them upon completion of the sign-up wizard. Following are some circumstances where customers need to know their password-based credentials:

- For a variety of reasons, user authentication might fail. For example, cached credentials might get corrupted or network connectivity issues might prevent wireless client computers from successfully authenticating.

- The user account expires and the customer wants to renew their account. In this circumstance, IAS sends a URL PEAP-TLV to the wireless client that contains the renewal action parameter and the URL of the provisioning server. After the wireless client is directed to your account renewal application, the customer must have their password-based credentials to renew their account.

If your customers know their user name and password, they can attempt to connect to your network until they are successful. If they do not have this information, they cannot fix the problem without calling your help desk.

You can create your Web application with the Microsoft .NET Framework 1.1 or with other application development software. If you want to create your Web application with the .NET Framework 1.1, you need the Microsoft .NET Framework 1.1 Software Development Kit.

**Microsoft .NET Framework 1.1 Software Development Kit**
Microsoft .NET Framework 1.1 Software Development Kit (SDK) includes .NET Framework 1.1, as well as everything you need to write, build, test, and deploy applications using .NET Framework 1.1. This includes documentation, samples, command-line tools, and compilers.

If you have already installed Microsoft® Visual Studio® .NET, you do not need to install .NET Framework 1.1 SDK separately; Visual Studio .NET includes the SDK. If you want to distribute .NET Framework 1.1 with your application, download .NET Framework 1.1 Redistributable in addition to the SDK.

You can get .NET Framework 1.1 SDK from the Download Center at http://go.microsoft.com/fwlink/?LinkId=17161. You can run .NET Framework 1.1 SDK on the following platforms:

- Windows Server 2003

- Microsoft Windows 2000 (Service Pack 2 is recommended)

- Windows XP Professional or Windows XP Home Edition (Windows XP Professional is required to run ASP.NET)

**Configure the provisioning server**
For the provisioning server, if you are using Internet Information Services (IIS), do the following:

1. Install Internet Information Services and ASP.NET.

   For more information, see "To install Internet Information Services (IIS) 6.0" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20033.

2. Verify that you have .NET Framework 1.1 installed by clicking **Start** on your Windows desktop, selecting **Control Panel**, and then double-clicking the **Add or Remove Programs** icon. When that window appears, scroll through the list of applications. If you see Microsoft .NET Framework 1.1 listed, the latest version is already installed and you do not need to install it again.

   a. If Microsoft .NET Framework 1.1 is already installed: register ASP.NET with Microsoft .NET Framework 1.1 by running the following command at Command Prompt:

      **%WINDIR%\Microsoft.NET\Framework\v1.1.4322\aspnet_regiis.exe –I**

   b. If Microsoft .NET Framework 1.1 is not installed: Install .NET Framework 1.1. You can install .NET Framework 1.1 through Microsoft Windows Update at http://go.microsoft.com/fwlink/?LinkId=284 or you can see the MSDN topic "How to Get the .NET Framework 1.1" at http://go.microsoft.com/fwlink/?LinkId=116211.

3. Create two folders in the Web server root location %systemroot%\Inetpub\wwwroot. You can use one folder to hold your custom Web application and one folder to hold the XML master and subfiles that you will be creating in later steps. For example, you can create a folder named **wpsdeploy** (%systemroot%\Inetpub\wwwroot\wpsdeploy) to contain your Web application, and you can create a folder named **wpsfiles** (%systemroot%\Inetpub\wwwroot\wpsfiles) to contain

your XML files.

4. Install your Web application into the folder you created for the Web application files.

5. Set user permissions for the Web application. Open Windows Explorer and browse to the location where you installed your Web application. For example, if you named your folder wpsdeploy, browse to %systemroot%\Inetpub\wwwroot\wpsdeploy. Right-click the wpsdeploy folder, and then click **Sharing and Security**. On the **Security** tab, click **Add**. In **Enter the object names to select**, type **Everyone**, and then click **OK**. In **Permissions for Everyone**, enable **Write** permissions by checking the **Allow** check box.

6. Enable HTTPS. You must configure IIS to use a certificate for secure Web communications between the provisioning server and clients. To enable HTTPS you must install a server certificate obtained from a public trusted root certification authority, such as a certificate from Verisign or Thawte. When you obtain the certificate, it must conform to the minimum server certificate requirements described in this paper.

   For more information, see "Network access authentication and certificates" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20016.

   To enable Secure Sockets Layer (SSL) and HTTPS, complete the following procedures.

| **Note** |
| --- |
| If you are deploying a certification authority in a test lab environment, you must install and configure the CA before completing the following procedures. |

**To obtain a new server certificate using Web Server Certificate Wizard**

1. In **IIS Manager**, expand the local computer, and then expand the **Web Sites** folder.

2. Right-click the Web site or file that you want, and then click **Properties**.

3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.

4. In **Web Server Certificate Wizard**, click **Create a new certificate**.

5. Complete **Web Server Certificate Wizard**, which will guide you through the process of requesting a new server certificate.

**To install a server certificate using Web Server Certificate Wizard**

1. In **IIS Manager**, expand the local computer, and then expand the Web Sites folder.

2. Right-click the Web site or file that you want, and then click **Properties**.

3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.

4. In **Web Server Certificate Wizard**, click **Assign an existing certificate**.

5. Complete **Web Server Certificate Wizard**, which will guide you through the process of installing a server certificate.

| **Note** |
| --- |
| The **Web Server Certificate Wizard** always denotes this step as "assigning" a certificate to a resource (such as a file, directory, or site), not as "installing." |

**Configure XML master and subfiles**

There are two methods you can use to create and configure your XML master and subfiles:

- Use the WPS Authoring Tool to create a WPS project and publish your XML files. The WPS Authoring Tool has a graphical user interface and is designed to assist you in accurately producing and managing a collection of XML files for your WPS solution. For more information, see "Using the WPS Authoring Tool" at http://go.microsoft.com/fwlink/?LinkId=41067. Using the WPS Authoring Tool to create your XML data files is recommended. You can download the WPS Authoring Tool at http://go.microsoft.com/fwlink/?LinkId=40535.

- Use the XML schemas provided in this paper to create your files. After you have created these files, you can enter information specific to your network and deployment parameters. For example, where the location of the provisioning server is required, you can provide an HTTPS URL. In another example, you may need to enter your domain name in several places; you can examine the schemas and example files and determine where to insert your domain name.

When your XML files are configured with the information for your organization, you can store them on your provisioning server and configure your Web application so that it can find the files when necessary. If you are using Internet Information Services as your Web server, you can install the files at the location %systemroot%\inetpub\wwwroot.

**Configure the database on the SQL server**

On the computer running SQL Server 2000 or a third-party product with similar functionality, create a promotion code database with the following fields, using the appropriate data type for each field:

- **Promotion code.** This field contains promotion codes that you distribute publicly to potential customers. When customers sign up for an account, they provide the promotion code that is matched by the Web application to a value in this field in the database.

- **User name.** This field has no value assigned until a customer creates an account. At this time, the Web application assigns a value to this field.

- **Domain name.** This field contains the domain name where you want the Web application to create the user account when a customer signs up using the promotion code.

- **Security group.** The Web application joins the new user account to the security group defined in this field.

- **Expiration date.** If your promotion lasts for a specific period of time, you can enter the expiration date related to the promotion code in this field. If a user with a promotion code attempts to create an account after the expiration date for the promotion, they cannot do so.

Populate the database with records, providing values for all fields except user name, and enable the data link between your Web application and the SQL server. Also configure security and authentication on the SQL server so that your Web application has permission to access and write to the database. For more information, see SQL Server at http://go.microsoft.com/fwlink/?LinkId=20014.

| **Important** |
| --- |
| When you configure values for records in the SQL Server database, the following three items must match: the name of the security group in Active Directory, the value of the **security group** field in the SQL Server database, and the name of the security group configured in the Valid user remote access policy in IAS. The Web application uses the value of the SQL Server database **security group** field to determine group membership for new accounts. |

**Configure the Windows XP-based client computer**

On the computer running Windows XP Professional, Windows XP Tablet PC Edition, or Windows XP Home Edition, install SP2. The computer must have a wireless network adapter compatible with IEEE 802.11 and 802.1X.

| **Note** |
| --- |
| If you are deploying WPS technology in a test lab environment where you have deployed your own enterprise root CA, you must import the CA certificate into the Trusted Root Certification Authorities certificate store on the client computer. If this certificate is not imported into the client certificate store, the client computer will not trust the enterprise root CA and server authentication to the client will fail. For more information, see "Configure certificates in a test lab environment (optional)" in this paper. |

**Configure certificates in a test lab environment (optional)**

If you are deploying WPS technology in a test lab, you can obtain server certificates from a public trusted root CA or you can deploy Certificate Services for Windows Server 2003 on your test network. To deploy Certificate Services and enroll certificates to domain member computers (such as the servers on your networks), do the following:

1. On a computer running Windows Server 2003, install Certificate Services.

   For more information, see "To install an enterprise root certification authority" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20035.

2. Add a server certificate template to the certification authority and configure the certification authority to allow computers to request a certificate that is based on the template you create.

   For WPS technology in a test lab environment, you must create a server certificate that will be trusted by client computers. The server certificate must meet the requirements stated in "Server Certificate Requirements" in this paper.

   In addition, you must base your certificate on the correct certificate template for the operating system you are running. If you are running an enterprise certification authority (CA) on a computer running Windows Server 2003, Standard Edition, and your IAS server is a domain member, base your certificate on the Computer certificate template. If you are running an enterprise certification authority (CA) on a computer running Windows Server 2003, Enterprise Edition, and your IAS server is a domain member, base your certificate on the **RAS and IAS Server** certificate template.

   For more information, see "To create a certificate template" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20036.

3. Autoenroll certificates to domain member computers. Autoenrollment allows domain member computers to automatically obtain, or enroll, certificates. For WPS technology, autoenroll certificates to servers using the certificate template you created in the previous step.

   For more information, see "Planning for autoenrollment deployment" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20037 and "Checklist: Configuring certificate autoenrollment" at http://go.microsoft.com/fwlink/?LinkId=20038.

4. Install the enterprise CA certificate in the Trusted Root Certification Authority certificate store on client computers being used for testing purposes. You can request the enterprise root CA certificate by using Web enrollment services, which is installed with Certificate Services, or you can export the server certificate to a floppy disk, and then import the certificate into the Trusted Root Certification Authority certificate store on the client.

   For more information, see "To export a certificate" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20039 and "To import a certificate" at http://go.microsoft.com/fwlink/?LinkId=20040.

# WPS Technology for the Enterprise

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Enterprises can deploy WPS technology to simplify the process of providing Internet access to organization visitors. You can provide business partners and other visitors with promotion codes, which visitors use to gain Internet access through your Wi-Fi hotspots.

In the sections that follow, the components of WPS technology for the enterprise are described, how the components work together during a user sign-up is detailed, and how to deploy WPS technology for the enterprise is explained.

**Components of WPS technology for the enterprise**

This deployment scenario, designed for the enterprise that deploys Wi-Fi hotspots, has the following features:

- A VLAN-aware gateway device is used for client computer isolation during the account sign-up process.
- Users gain network access by using a promotion code. The codes are stored in a database on a server running SQL Server 2000 or a third-party database application.
- A perimeter network stands between wireless access points and the Enterprise LAN.
- On the perimeter network, an IAS proxy server forwards RADIUS messages between RADIUS clients (access points) and the IAS server on the Enterprise LAN. In addition, IPsec is deployed between the IAS proxy and IAS server to provide security for RADIUS traffic.
- The provisioning server on the perimeter network does not process user information or create accounts in the user account database. The provisioning server on the perimeter network maintains a custom XML-forwarding application that forwards XML files containing customer sign-up information to the account processing application on the account processing server. In addition, IPsec is deployed between the provisioning server and the account processing server on the enterprise LAN to provide security for traffic between the XML-forwarding application and the account processing application.

The account processing server is located on the enterprise LAN. The account processing server verifies promotion codes against the SQL Server database and creates accounts in the user accounts database for business partners or other visitors that have valid promotion codes and access to your Wi-Fi hotspots.

The following illustration depicts the components of an enterprise WISP network using a perimeter network and a VLAN-aware gateway device for client computer isolation.




## Components of WPS technology for the enterprise

In the following sections are descriptions of the components of an enterprise WISP network, including components of the WLAN, the perimeter network, and the enterprise LAN.

### Wi-Fi hotspot components

Following are the components that comprise the wireless local area network (WLAN), or Wi-Fi hotspot:

### Wireless client

A computer running Windows XP Home Edition with SP2, Windows XP Professional with SP2, or Windows XP Tablet PC Edition with SP2. The computer must be equipped with a wireless network adapter that provides support for IEEE standard 802.11, IEEE standard 802.1X authentication, and Wired Equivalent Privacy (WEP). Support for Wi-Fi Protected Access (WPA) is preferred, but not required.

### Wireless access point (RADIUS client)

One or more wireless access points deployed with a wired connection to the access controller, VLAN-aware switch or router, or other gateway device.

The wireless access point is configured as a RADIUS client to the Internet Authentication Service (IAS) proxy deployed on the perimeter network. The wireless access points used for WPS technology must have the following required features:

- Support for the use of VLANs.
- Support for the IEEE standard 802.1X authentication.
- Support for Wi-Fi Protected Access (WPA) is preferred.
- Support for RADIUS authentication and RADIUS accounting, including:
    - Support for the Class attribute as defined in RFC 2865, “Remote Authentication Dial-in User Service (RADIUS),” to allow session correlation for RADIUS authentication and accounting records. For session correlation, when you configure RADIUS accounting at your IAS server or proxy, you must log all accounting data that allow applications (such as billing applications) to query the database, correlate related fields, and return a cohesive view of each session in the query results. At a minimum, to provide session correlation, you must log the

following IAS accounting data: NAS-IP-Address; NAS-Identifier (you need both NAS-IP-Address and NAS-Identifier because the access server can send either attribute); Class; Acct-Session-Id; Acct-Multi-Session-Id; Packet-Type; Acct-Status-Type; Acct-Interim-Interval; NAS-Port; and Event-Timestamp.

- Support for accounting interim requests, which are sent periodically by some access servers during a user session, that can be logged. This type of request can be used when the Acct-Interim-Interval RADIUS attribute is configured to support periodic requests in the remote access profile on the IAS server. The access server, in this case a wireless access point, must support the use of accounting interim requests if you want the interim requests to be logged on the IAS server.

- Support for IP address range filtering.

- Support for dynamic RTO estimation or exponential backoff to handle congestion and delays in a wide area network (WAN) environment.

In addition, there are some filtering features that the access points must support to provide enhanced security for the network. These filtering options include:

- **DHCP filtering.** The access point must filter on IP ports to prevent the transmission of DHCP broadcast messages in the instance that the client is a DHCP server. The access point must block the client from sending IP packets from port 68 to the network.

- **DNS filtering.** The access point must filter on IP ports to prevent a client from performing as a DNS server. The access point must block the client from sending IP packets from port 53 to the network.

**Perimeter network components**

- Following are the components that comprise the perimeter network.

**VLAN-aware gateway device**
The VLAN-aware gateway device can be an access controller, a VLAN-aware router, a VLAN-aware switch, or any other device that can be configured to apply IAS-provided parameters to client connections. The VLAN-aware gateway device is configured with two VLANs: a Network Resource VLAN and an Internet VLAN.

The Network Resource VLAN allows all users access to the provisioning server and DHCP server. This VLAN grants access to network resources that allow users to connect to your network as guest, create an account, and receive provisioning information from the provisioning server.

The Internet VLAN provides access to the Internet. Only users who have promotion codes and have created accounts using the sign-up wizard are switched to this VLAN and granted Internet access. This process occurs when Windows XP reauthenticates the user with the newly created account information. When the user is authenticated and authorized by your IAS server, the IAS server returns attributes to the VLAN-aware wireless access point; the access point instructs the VLAN-aware gateway device to route traffic from the client to the Internet VLAN.

**IAS proxy**
The IAS proxy is running Windows Server 2003, Standard Edition with SP1; Windows Server 2003, Enterprise Edition with SP1; or Windows Server 2003, Datacenter Edition with SP1, and is configured with the following components.

**Internet Authentication Service**
IAS is configured as a proxy to forward RADIUS messages between RADIUS clients, such as access points that are located on the wireless LAN, and the IAS server that is located on the Enterprise LAN.

**IPsec**
An Active Directory-based IPsec policy provides private, secure communications for RADIUS traffic between the IAS proxy and the IAS server.

**Provisioning server**
The enterprise provisioning server is configured with the following components.

**HTTPS Web server**
The Internet Information Services (IIS) or third-party Web server must be deployed with HTTPS.

**XML master and subfiles**
The enterprise provisioning server stores the XML master and subfiles that provide the client with all configuration information needed to access the network, create an account, and ultimately access the Internet. For more information about the XML master file and subfiles, see "XML schemas" in this paper.

**XML-forwarding Web application**
The XML-forwarding Web application is a custom program, application, or service that you create and install on your provisioning server. The XML-forwarding application is used to forward XML documents from client computers to the account processing server on the Enterprise LAN. The XML documents are created on client computers when users run the sign-up wizard and enter all required information, such as their name and promotion code. Windows XP packages the information as an XML document, and then sends the document to the provisioning server. The XML-forwarding application then forwards the document to the account processing server.

**Server certificate**
For server authentication to client computers, the provisioning server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and be issued by a public certification authority (CA), such as Verisign or Thawte, which is trusted by client computers. The trusted root certification authority certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust.

In test lab deployments of WPS technology, you can deploy your own certification authority in lieu of using a public trusted root CA. In this circumstance you must install the private trusted root CA certificate in the Trusted Root Certification Authorities certificate store on all clients. In addition, you must enroll certificates to your provisioning server and IAS server that meet the minimum server certificate requirements. For more information, see "Server Certificate Requirements" in this paper.

**IPsec**
An Active Directory-based IPsec policy provides private, secure communications for all traffic between the account processing server on the Enterprise LAN and the provisioning server on the perimeter network.

**DHCP server**
If you do not deploy network address translation (NAT), the DHCP server must be able to assign valid public IP addresses to computers accessing the network through the wireless access points. However, if you deploy NAT, you can assign IP addresses from a private IP address range to wireless client computers.

**Router**
A router connects the perimeter network and the Enterprise LAN.

**Enterprise LAN components**
Following are the components that comprise the Enterprise LAN.

**Domain controller and IAS server**
The enterprise domain controller and IAS server is running Windows Server 2003, Standard Edition with SP1; Windows Server 2003, Enterprise Edition with SP1; or Windows Server 2003, Datacenter Edition with SP1, and is configured with the following components.

**Active Directory**
In this scenario, Active Directory is deployed. The user accounts database on the domain controller must be an Active Directory user accounts database or a database that uses Lightweight Directory Access Protocol (LDAP) and supports dynamic creation of user accounts.

When a user signs up for an account, the account processing Web application on the account processing server creates a new account in the user accounts database, and adds the user to a group that has clearly defined access privileges.

If you use an accounts database other than Active Directory, IAS extensions must be written and installed for this process to function correctly.

**IP Security (IPsec) policy**
In Active Directory, an IPsec policy for the domain includes the following:

- **IPsec rule for IAS proxy and server.** This rule protects RADIUS traffic between the IAS proxy on the perimeter network and the IAS server on the Enterprise LAN.

- **IPsec rule for HTTPS Web servers.** This rule protects all traffic between the provisioning server on the perimeter network and the account processing server on the Enterprise LAN.

**Internet Authentication Service (IAS)**
IAS is the Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server, and is used to authenticate and authorize users connecting to your network. IAS is configured with remote access policies that allow guest authentication for non-domain member computers and users. It is also configured to provide attributes to RADIUS clients (access points) that instruct the gateway device to apply the attributes to client connections. Protected Extensible Authentication Protocol (PEAP) with MS-CHAP v2 is configured in remote access policies as the authentication method used by clients and server.

**Extension DLL and URL PEAP-TLV**
AN IAS extension DLL defining a URL PEAP-TLV provides IAS with the ability to send the location of the provisioning server to client computers.

PEAP-Type-Length-Value (PEAP-TLV) is an Extensible Authentication Protocol (EAP) authentication type that allows the IAS server to pass information to client computers attempting to connect to your network.

In this circumstance, the value contained in the PEAP-TLV is an HTTPS Uniform Resource Locator (URL) that provides client computers with the location of the WISP provisioning server. With this URL, Windows XP can download the WISP XML files to the client computer.

In addition to the URL of the provisioning server, the URL PEAP-TLV includes an action parameter. The action parameter directs the client to perform a specific task. The action parameter included in the URL PEAP-TLV defines tasks such as new customer sign-up, existing account renewal, and password change.

For more information, see "How to create an IAS extension DLL and a URL PEAP-TLV" in this paper.

**Server certificate**
For server authentication to client computers, the enterprise IAS server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and be issued by a public certification authority (CA), such as Verisign or Thawte, that is trusted by client computers. The Trusted Root Certification Authorities certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust. If you install IAS and Active Directory on the same computer, the computer must have a certificate. If you install IAS and Active Directory on different computers, only the IAS server needs a certificate.

**Account processing server**
The account processing server is configured with the following components:

**HTTPS Web server**
The Internet Information Services (IIS) or third-party Web server must be deployed with HTTPS.

**Account processing Web application**
The account processing Web server is configured with an account processing Web application that processes data provided during user sign-up or account renewal. The Web application accepts and processes XML documents forwarded from the provisioning server on the perimeter network.

When a user completes the sign-up wizard on a client computer, the user enters data, such as name, organization affiliation, and promotion code, that is converted to an XML document on the client. Windows XP sends this XML document to the provisioning server on the perimeter network, which in turn forwards the XML document to the account processing application.

The account processing application must be capable of accepting and processing the XML documents containing the user data. For example, the account processing application must compare the promotion code entered by the user to a promotion code in the SQL Server database, and then dynamically create an account in the Active Directory user accounts database with the properties (domain and security group membership) described in the database.

**IPsec**
An Active Directory-based IPsec policy provides private, secure communications for all traffic between the account processing server on the Enterprise LAN and the provisioning server on the perimeter network.

**SQL server**
A computer running SQL Server 2000 or another SQL-compatible relational database application.

The promotion code database on the SQL server is configured with the following fields: promotion code, user name, domain name, security group, and expiration date. With the exception of the user name field, each field for each record is preconfigured with a value. The value for the user name field is assigned by the Web application when a user creates an account with a promotion code that matches a value in the promotion code field in the database. By predefining the domain in which the user account is created by the Web application and the Active Directory security group to which the user account is joined as a

member, you can assign network access and other permissions for your users.

**How WPS technology works for the enterprise**
The following example describes how the components of a WPS network for the enterprise interact during the connection and account creation process for a new user.

**New user connection example**
When a new user connects and establishes an account, the following four basic stages occur:

1. The user discovers the network at a Wi-Fi hotspot
2. The user authenticates as guest
3. The client is provisioned and the user establishes an account
4. The user is authenticated using the new account credentials

In the next section we will look at these stages in more detail.

**1. The user discovers the network at a Wi-Fi hotspot**
When a user arrives at the Wi-Fi hotspot with a portable computer running Windows XP Home Edition with SP2, Windows XP Tablet PC Edition with SP2, or Windows XP Professional with SP2, the computer comes within range of the access point beacon.

Wireless auto configuration on the client computer detects the beacon information from the access point, which is enabled with broadcast Secure Set Identifier (SSID). The SSID is equivalent to the network name.

The user is informed by Windows XP that a wireless network is available. In this example, the user is employed by a business partner of the enterprise, and is provided by the enterprise with a promotion code to use for account establishment. The user proceeds by clicking **Connect**.

**2. The user authenticates as guest**
Wireless Auto Configuration uses 802.1X and PEAP guest authentication to connect to the enterprise perimeter network through the access point, automatically passing a null user name and a blank password to the IAS proxy, which forwards the message to the IAS server. The access point is connected to a VLAN-aware gateway device that allows traffic from the client to pass through the Network Resource VLAN, but blocks the client from access to the Internet VLAN.

The IAS server is the PEAP authenticator and TLS endpoint for users who connect as guest. The TLS tunnel is created between the client and the IAS server. All subsequent messages between client and server pass through this tunnel, which traverses the access point, the gateway device, and the IAS proxy.

Server authentication is performed when the IAS server verifies its identity to the client computer using a certificate that contains the Server Authentication purpose in Enhanced Key Usage (EKU) extensions. This certificate is issued by a public trusted root CA that the client computer trusts.

The IAS server authenticates and authorizes the customer as guest. In the Access-Challenge message that the IAS server sends to the client is a URL PEAP-TLV message. The URL PEAP-TLV contains the URL of the provisioning server. This URL provides the client with the location of the XML master file.

The client computer receives an IP address lease from the DHCP server. The address is from a public IP address range configured in a scope on the DHCP server. In addition to the IP address, the client receives DHCP options, such as DNS server IP address.

**3. The client is provisioned and the user creates an account**
The XML master file on the provisioning server contains pointers to the XML subfiles. The client downloads the XML master file and subfiles. When the XML sign-up schema is downloaded, the sign-up wizard is started on the client to allow the user to create an account.

Using the sign-up wizard on the client computer, the user steps through the process of signing up for an account. The customer enters the promotion code as well as personal data such as name, employer, and job title. The data entered by the user is converted into an XML document.

The XML document containing the user's sign-up data is sent to the XML-forwarder Web application on the provisioning server.

The XML-forwarder Web application on the provisioning server sends the XML document to the account processing application on the account processing server.

The account processing application checks the promotion code entered by the user against the promotion code database on the SQL server. If the promotion code is valid, the account processing Web application continues processing the user's data.

The account processing Web application reads the domain and security group information from the promotion code database on the SQL server. The account processing application creates a user account in Active Directory and adds the account to the security group. The application also enters the new user name in the promotion code database.

An XML document containing the new account credentials is sent from the account processing server to the XML-forwarder application on the provisioning server; the XML-forwarder application passes the XML document to the client computer. The client computer uses the credentials to configure wireless auto configuration and 802.1X under the name of the enterprise.

**4. The user is authenticated using the new account credentials**

Wireless auto configuration restarts the association to the SSID for the enterprise WLAN.

Wireless auto configuration finds the correct 802.11 profile which was downloaded with the other network information. Wireless auto configuration re-associates with the access point using the correct profile.

802.1X restarts the authentication process using PEAP-MS-CHAP v2 and the new account credentials.

As the client starts the authentication process with PEAP-MS-CHAP v2 authentication, a TLS channel is created between the customer's client computer and the enterprise IAS server.

In the second stage of PEAP-MS-CHAP v2 authentication, the IAS server authenticates and authorizes the connection request against the new account in the Active Directory user accounts database. The IAS server sends an Access-Accept message to the access point. Included in the Access-Accept message are attributes that specify which VLAN the customer can access.

The access point instructs the gateway device to assign the client to the Internet VLAN rather than the Network Resource VLAN.

The gateway device switches the client to the Internet VLAN, and the customer is provided with access to the Internet.

**How to deploy WPS technology for the enterprise**

In the set of instructions that follow, these assumptions are made:

1. The computer functioning as your IAS server has Windows Server 2003 with SP1 installed, and the **EnableWPSCompatibility** registry key is enabled according to the instructions in "Configuring IAS for WPS technology" in this paper.

2. Client computers are running Windows XP Home Edition with SP2; Windows XP Tablet PC Edition with SP2; or Windows XP Professional with SP2.

3. All of your hardware, including the VLAN-aware gateway device and VLAN-aware wireless access points, meet all of the technical requirements stated in "Components of WPS technology for the enterprise" in this paper.

4. You have already deployed a computer running SQL Server 2000 or a third-party database program on your network. For information about SQL Server 2000, see SQL Server at http://go.microsoft.com/fwlink/?LinkId=20014.

5. You have SQL Server 2000 or third-party relational database development experience and you understand how to use SQL Server 2000 or a third-party database program to create, modify, administer, and manage your databases.

6. You have experience deploying an Internet Information Services (IIS) or third-party Web server with HTTPS.

7. You have software development experience that allows you to create two custom Web applications. One Web application runs on the provisioning server and forwards XML documents between the account processing application and client computers. Another Web application runs on the account processing server and processes user data and creates user accounts in Active Directory or an LDAP-compliant third-party user accounts database.

8. You have software development experience that allows you to create an IAS extension DLL.

To deploy WPS technology for the Enterprise, the basic steps are as follows:

| **Note** |
| --- |
| The instructions that follow use four servers upon which various programs and services are installed. If you deploy WPS technology in a test lab, you can reduce or increase the number of servers in a manner appropriate to your available hardware resources. For example, in a test lab environment you might want to deploy Active Directory, IAS, IIS, and DHCP on the same server. |

1. Create your enterprise Web applications
2. Configure the enterprise domain controller and IAS server
3. Configure the enterprise account processing server
4. Configure the database on the enterprise SQL server
5. Install and configure the enterprise router
6. Configure the enterprise XML master file and subfiles
7. Configure the enterprise provisioning server
8. Configure the enterprise DHCP server
9. Install and configure your enterprise VLAN-aware gateway device
10. Install and configure your enterprise wireless access points
11. Install and configure the enterprise IAS proxy
12. Configure RADIUS clients at the enterprise IAS server
13. Configure the enterprise Windows XP-based client computer
14. Configure certificates in an enterprise test lab environment (optional)

**Create your enterprise Web applications**
The XML forwarding Web application to be installed on the provisioning server must be capable of performing the following functions:

- Communicating with client computers using HTTPS.
- Uploading the XML master file and subfiles that are stored on the provisioning server to client computers that request the files.
- Accepting XML documents from client computers that contain user data.
- Forwarding XML documents that contain user data to the account processing server.

The account processing Web application to be installed on the account processing server must be capable of performing the following functions:

- Accepting and processing XML documents from the provisioning server that contain user data, such as promotion code, user name, and other information.

- Reading the promotion code database records to validate promotion codes.
- Reading the promotion code database records to determine the domain in which to create a new user account.
- Reading the promotion code database records to determine the security group membership for a new user account.
- Writing a user name to the user name field in the promotion code database.
- Dynamically creating new accounts in Active Directory (or a third-party LDAP-compliant database) using data provided by users and parameters derived from the promotion code database.

It is recommended that the design of your Web application provides users with knowledge of their password-based credentials (user name and password). Users should either be allowed to designate user name and password or this information should be provided to them upon completion of the sign-up wizard. Following are some circumstances where users will need to know their password-based credentials:

- For a variety of reasons, user authentication might fail. For example, cached credentials might get corrupted or network connectivity issues might prevent wireless client computers from successfully authenticating.
- The user account expires and the user wants to renew their account. In this circumstance, IAS sends a URL PEAP-TLV to the wireless client that contains the renewal action parameter (**#renewal**) and the URL of the provisioning server. After the wireless client is directed to your account renewal application, the user must have their password-based credentials to renew their account.

If your users know their user name and password, they can attempt to connect to your network until they are successful. If they do not have this information, they cannot fix the problem without calling your help desk.

You can create your Web application with the Microsoft .NET Framework 1.1 or with other application development software. If you want to create your Web application with the .NET Framework 1.1, you need the Microsoft .NET Framework 1.1 Software Development Kit.

**Microsoft .NET Framework 1.1 Software Development Kit**
Microsoft .NET Framework 1.1 Software Development Kit (SDK) includes .NET Framework 1.1, as well as everything you need to write, build, test, and deploy applications using .NET Framework 1.1. This includes documentation, samples, command-line tools, and compilers.

If you have already installed Microsoft® Visual Studio® .NET, you do not need to install .NET Framework 1.1 SDK separately; Visual Studio .NET includes the SDK. If you want to distribute .NET Framework 1.1 with your application, download .NET Framework 1.1 Redistributable in addition to the SDK.

You can get .NET Framework 1.1 SDK from the Download Center at http://go.microsoft.com/fwlink/?LinkId=17161. You can run .NET Framework 1.1 SDK on the following platforms:

- Windows Server 2003
- Microsoft Windows 2000 (Service Pack 2 is recommended)
- Windows XP Professional or Windows XP Home Edition (Windows XP Professional is required to run ASP.NET)

**Configure the enterprise domain controller and IAS server**
Install Windows Server 2003, Standard Edition with SP1; Windows Server 2003, Enterprise Edition with SP1; or Windows Server 2003, Datacenter Edition with SP1 on a computer that meets or exceeds the minimum hardware requirements for the respective operating system. After the operating system is installed, you can perform general configuration, Active Directory configuration, IPsec policy configuration, IAS configuration, and IAS extension DLL & URL PEAP-TLV configuration.

**General configuration**

1. Assign the server a static IP address. For more information, see "To configure TCP/IP for static addressing" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20015

2. Install a server certificate obtained from a public trusted root certification authority, such as a certificate from Verisign.

   For more information, see "Obtaining and Installing a VeriSign WLAN Server Certificate for PEAP-MS-CHAP v2 Wireless Authentication" at http://go.microsoft.com/fwlink/?LinkId=33675.

   When you obtain the certificate, it must conform to the minimum server certificate requirements described in this paper.

   For more information, see "Network access authentication and certificates" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20016.

| **Note** |
| --- |
| If you are deploying WPS technology in a test lab environment, you can install a private CA rather than obtain certificates from a public CA. For more information, see "Configure certificates in an enterprise test lab environment (optional)" in this paper. |

**Active Directory configuration**

1. Install Active Directory and DNS. To install Active Directory, open Command Prompt, type **dcpromo**, and then follow the instructions provided in the wizard, entering your network configuration information in the wizard as you progress.

2. Design and create your security group or groups. When users sign up and create an account, your Web application adds the new user account as the member of a security group that you create in this step. The Web application chooses group membership based on the value of the security group field in the promotion code database on your SQL server, so you need to match the security group you create to the security group field in the SQL Server database. If you have the need for multiple security groups, you can assign permissions to each group individually. For more information, see "To create a new group" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20018 and "Assign user rights to a group in Active Directory" at http://go.microsoft.com/fwlink/?LinkId=20019.

   | **Important** |
   | --- |
   | The security groups you create in Active Directory are named and used in the SQL Server database and in IAS remote access policy. |

3. Enable the Guest account in Active Directory. If guest access is not enabled in Active Directory, new customers cannot access your provisioning server by authenticating as guest. To enable the Guest account, open the Active Directory Users and Computers snap-in, and then double-click **Users**. Right-click the account named **Guest**, and then click **Enable Account**.

To perform this procedure, you must be a member of the Domain Admins group in the domain for which you want to raise functionality or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority.

Raise the domain functional level to Windows 2000 native or Windows Server 2003 by doing the following:

1. Open the Active Directory Domains and Trusts snap-in. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.

2. In the console tree, right-click the domain for which you want to raise functionality, and then click **Raise Domain Functional Level**.

3. In **Select an available domain functional level**, do one of the following:

   a. To raise the domain functional level to Windows 2000 native, click **Windows 2000 native**, and then click **Raise**.

   b. To raise domain functional level to Windows Server 2003, click **Windows Server 2003**, and then click **Raise**.

| **Important** |
| --- |

If you have or will have any domain controllers running Windows NT 4.0 and earlier, then do not raise the domain functional level to Windows 2000 native. After the domain functional level is set to Windows 2000 native, it cannot be changed back to Windows 2000 mixed.

Likewise, if you have or will have any domain controllers running Windows NT 4.0 and earlier or Windows 2000, then do not raise the domain functional level to Windows Server 2003. After the domain functional level is set to Windows Server 2003, it cannot be changed back to Windows 2000 mixed or Windows 2000 native.

The current domain functional level is displayed under Current domain functional level in the Raise Domain Functional Level dialog box.

For more information, see "Domain and forest functionality" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=30600.

For information about configuring Active Directory replication, see "Active Directory replication" in this paper.

**IPsec policy configuration**

One IP Security policy can contain rules that accomplish the following:

- Protect RADIUS traffic between the IAS proxy and the IAS server
- Protect all traffic between the provisioning server and the account processing server

**Note**

To successfully complete IP Security policy configuration, you must know the IP addresses of the following servers: the IAS proxy, the provisioning server, and the account processing server.

**To create an IP Security policy in Active Directory using the IP Security Policy Management snap-in, do the following:**

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **IP Security Policy Management**, and then click **Add**.
4. Click the Active Directory domain of which this computer is a member. Click **Finish**, click **Close**, and then click **OK**.
5. To add a new policy, in the console tree, click **IP Security Policies on Active Directory**. On the **Action** menu, click **Create IP Security Policy**. The IP Security Policy Wizard is started. Click **Next**.
6. In **IP Security Policy Name**, type a name for your IPsec policy. Optionally, type a description for the policy. Click **Next**.
7. In **Requests for Secure Communication**, verify that **Activate the default response rule** is selected. Click **Next**.
8. In **Default Response Rule Authentication Method**, verify that **Active Directory default (Kerberos V5 protocol)** is selected. Click **Next**, and then click **Finish**.

   The *PolicyName* **Properties** dialog box opens, where *PolicyName* is the name you typed for your policy. For example, if you named your policy "**WPS policy**," the dialog box that opens is named **WPS policy Properties**.

**To create a rule with RADIUS traffic filters**

1. In *PolicyName* **Properties**, click **Add**. The Create IP Security Rule Wizard is started. Click **Next**.

   **Note**

The Create IP Security Rule Wizard is also referred to as the Security Rule Wizard.

2. In **Tunnel Endpoint**, verify that **This rule does not specify a tunnel** is selected. Click **Next**.

3. In **Network Type**, verify that **All network connections** is selected.

4. In **IP Filter List**, click **Add**. In **Name**, type **RADIUS filters**, and then click **Add**. The **IP Filter Wizard** is started. Click **Next**.

5. In **IP Filter Description and Mirrored property**, type **RADIUS port 1812**. Verify that **Mirrored. Match packets with the exact opposite source and destination addresses** is selected. Click **Next**.

6. In **IP Traffic Source**, verify that **My IP Address is selected** in Source address. Click **Next**.

7. In **IP Traffic Destination**, click **Destination address**, and then select **A specific IP Address**. In **IP address**, type the IP address of the IAS proxy server that is located on the perimeter network. Click **Next**.

8. In **IP Protocol Type**, click **Select a protocol type**, and then select **UDP**. Click **Next**.

9. In **IP Protocol Port**, for **Set the IP protocol port**, verify that **From any port** is selected. Select **To this port**, type **1812**, click **Next**, and then click **Finish**. The **IP Filter List** dialog box is now visible.

10. In **IP filter list**, click **Add**. The **IP Filter Wizard** is started. Click **Next**.

11. In **IP Filter Description and Mirrored property**, type **RADIUS port 1813**. Verify that **Mirrored. Match packets with the exact opposite source and destination addresses** is selected. Click **Next**.

12. In **IP Traffic Source**, verify that **My IP Address** is selected in Source address. Click **Next**.

13. In **IP Traffic Destination**, click **Destination address**, and then select **A specific IP Address**. In **IP address**, type the IP address of the IAS proxy server that is located on the perimeter network. Click **Next**.

14. In **IP Protocol Type**, click **Select a protocol type**, and then select **UDP**. Click **Next**.

15. In **IP Protocol Port**, for **Set the IP protocol port**, verify that **From any port** is selected. Select **To this port**, type **1813**, click **Next**, and then click **Finish**. To close the RADIUS filters **IP Filter List**, click **OK**. The **Security Rule Wizard** is still running and should now be visible. In **IP Filter Lists**, select **RADIUS filters**, and then click **Next**.

16. In **Filter Action**, click **Require Security**. Click **Next**.

17. In **Authentication Method**, verify that **Active Directory default (Kerberos V5 protocol)** is checked, and then click **Next**.

18. In **Completing the IP Filter Wizard**, verify that the **Edit properties** check box is cleared, and then click **Finish**. This returns you to the *PolicyName* **Properties** dialog box.

**To create an IPsec rule for all traffic between the provisioning server and the account processing server**

1. In *PolicyName* **Properties**, click **Add**. The Create IP Security Rule Wizard is started. Click **Next**.

2. In **Tunnel Endpoint**, verify that **This rule does not specify a tunnel** is selected. Click **Next**.

3. In **Network Type**, verify that **All network connections** is selected. Click **Next**.

4. In **IP Filter List**, click **Add**. In **Name**, type **Provisioning server filters**, and then click **Add**. The IP Filter Wizard is started. Click **Next**.

5. In **IP Filter Description and Mirrored property**, type a description for the filter. Verify that **Mirrored. Match packets with the exact opposite source and destination addresses** is selected. Click **Next**.

6. In **IP Traffic Source**, click **Source address**, and then select **A specific IP Address**. In **IP Address**, type the IP address of the provisioning server that is located on the perimeter network. Click **Next**.

7. In **IP Traffic Destination**, click **Destination address**, and then select **A specific IP Address**. In **IP address**, type the IP address of the account processing server that is located on the enterprise LAN. Click **Next**.

8. In **IP Protocol Type**, verify that the value of **Select a protocol type** is Any. Click **Next**.

9. In **Completing the IP Filter Wizard**, verify that the **Edit properties** check box is cleared, and then click **Finish**. This returns you to the **IP Filter List** dialog box. Click **OK**.

10. The **Security Rule Wizard** is still running and should now be visible. In **IP Filter List**, select **Provisioning server filters**, and then click **Next**.

11. In **Filter Action**, click **Require Security**. Click **Next**.

12. In **Authentication Method**, verify that **Active Directory default (Kerberos V5 protocol)** is selected, and then click **Next**.

13. In **Completing the Security Rule Wizard**, verify that the **Edit properties** dialog box is cleared, and then click **Finish**.

14. This returns you to the *PolicyName* **Properties** dialog box. Displayed in **IP Security rules** are two rules, **RADIUS filters** and **Provisioning server filters**. To close the *PolicyName* **Properties** dialog box, click **OK**.

**IAS configuration**

For IAS, the three configuration stages are general configuration, remote access policy configuration, and connection request policy configuration. RADIUS client configuration occurs later in the overall WPS deployment process.

**General configuration**

1. Install Internet Authentication Service.

   For more information, see "To install IAS" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20028.

2. Register IAS in Active Directory. In order for IAS to have permission to read user accounts in Active Directory, IAS must be registered in Active Directory.

   For more information, see "To enable the IAS server to read user accounts in Active Directory" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20030.

3. Delete the default remote access policies. To delete the policies, open the IAS console and click **Remote Access Policies**. Select each existing policy, right-click the policy, and then click **Delete**.

4. Install your RADIUS extension DLL on your IAS server:

   a. Open Command Prompt and change directories to the folder that contains your DLL.

   b. Type the following: **regsvr32***DLL_name.dll*, where *DLL_name.dll* is the name of your DLL file.

   Make sure that you configure DLL registry keys according to your needs.

   For more information, see "How to Create an IAS Extension DLL and a URL PEAP-TLV" in this paper.

**Remote access policy configuration**

There are two remote access policies configured for WPS technology. The Guest access policy provides network parameters and rules for users connecting as a guest. The Valid Users access policy provides network parameters and rules for users who have valid enterprise accounts.

| **Note** |
| --- |
| If you have a variety of account types that you offer to users and these accounts have different properties (such as membership to different security groups), you might find it necessary to create more than two remote access policies on your IAS server. If this is the case, you can use the remote access policies described below to extrapolate how to create additional policies. |

**To configure the Guest access policy**

1. Open the Internet Authentication Service console and, if necessary, double-click **Internet Authentication Service**.
2. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**.
3. The **New Remote Access Policy Wizard** starts. For the Guest access policy, specify the following:
    a. For **How do you want to set up this policy?** verify that **Use the wizard to set up a typical policy for a common scenario** is selected.
    b. For **Policy name**, type **Guest access** (or type another name for your policy that you prefer).
    c. For **Select the method of access for which you want to create a policy**, click **Wireless**.
    d. For **Grant access based on the following**, click **User**.
    e. In **Select the EAP type for this policy**, select **Protected EAP (PEAP)**, and then click **Configure**.
    f. In **Certificate issued**, select the certificate that you want the IAS server to use to verify its identity to client computers. Also select the **Enable Fast Reconnect** check box.

After you have completed creating the policy and have closed the wizard by clicking **Finish**, you need to perform additional policy configuration.

1. In the IAS console, click **Remote Access Policies**, and then double-click the policy you just created.
2. In the policy **Properties** dialog box, for **Policy conditions**, click **Add**.
3. In **Attribute Types**, click **Day-And-Time-Restrictions**, and then click **Add**. In **Time of day restraints**, select **Permitted**, configure the days and times that access is permitted, and then click **OK**.
4. In the policy **Properties** dialog box, click **Grant remote access permission**.
5. Click **Edit Profile**. On the **Authentication** tab, in **Unauthenticated access**, click **Allow clients to connect without negotiating an authentication method**.

**To configure the Valid Users access policy**

1. Open the Internet Authentication Service console and, if necessary, double-click **Internet Authentication Service**.
2. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**.
3. Use **New Remote Access Policy Wizard** to create a policy. For the Valid Users access policy, you can choose the

following:

a. For **How do you want to set up this policy?** verify that **Use the wizard to set up a typical policy for a common scenario** is selected.

b. For **Policy name**, type **Valid Users** (or type another name for your policy that you prefer).

c. For **Select the method of access for which you want to create a policy**, click **Wireless**.

d. For **Grant access based on the following**, click **Group**, and then click **Add**. In **Enter the object name to select**, type the name of a security group that you defined when configuring Active Directory.

| **Important** |
| --- |
| The following three items must match: the name of the security group in Active Directory, the value of the **security group** field in the SQL server database, and the name of the security group configured in the Valid User access policy in IAS. The Web application uses the value of the SQL server database **security group** field to determine group membership for new accounts. |

e. In **Select the EAP type for this policy**, select **Protected EAP (PEAP)**, and then click **Configure**.

f. In **Certificate issued**, select the certificate that you want the IAS server to use to verify its identity to client computers. Also select the **Enable Fast Reconnect** check box.

After you have completed creating the policy and have closed the wizard by clicking **Finish**, you need to perform additional policy configuration.

1. In the IAS console, click **Remote Access Policies**, and then double-click the policy you just created.

2. In the policy **Properties** dialog box, for **Policy conditions**, click **Add**.

3. In **Attribute Types**, click **Day-And-Time-Restrictions**, and then click **Add**. In **Time of day restraints**, select **Permitted**, configure the days and times that access is permitted, and then click **OK**.

4. In the policy **Properties** dialog box, click **Grant remote access permission**.

5. Click **Edit Profile**, and then click the **Advanced** tab. By default, the **Service-Type** attribute appears in **Attributes** with a value of **Framed**. To specify additional connection attributes required for WPS technology with VLANs, click **Add**, and then add the following attributes:

   a. **Framed-Protocol.** Value: **PPP**

   b. **Tunnel-Medium-Type.** Value: **802 (Includes all 802 media plus Ethernet canonical format)**

   c. **Tunnel-Pvt-Group-ID.** Value: Enter the integer that represents the VLAN number for the Internet VLAN. For example, if your access controller's Internet VLAN is VLAN 4, type **4**.

   d. **Tunnel-Type.** Value: **Virtual LANs (VLAN)**

   e. **Tunnel-Tag.** Value: Obtain this value from your hardware documentation

| **Important** |
| --- |

IAS evaluates remote access policies in the order in which they appear in the IAS console under **Remote Access Policies**. The Valid Users access policy must be the first policy in the list of remote access policies or else valid user authentication will fail. Because IAS places newly created policies in the first position and the Valid Users access policy was the last policy created, the Valid Users access policy should now appear first in the IAS console, with the Guest access policy appearing second. If this is not the case, move the Valid Users access policy into first position.

**Connection request policy configuration**

By default, there is one connection request policy predefined in the IAS console, called **Use Windows authentication for all users**. This policy can be used for WPS technology.

**To configure connection request policy**

1. In the IAS console, double-click **Connection Request Processing**, click **Connection Request Policies**, and then double-click the policy **Use Windows authentication for all users**.

2. Click **Edit Profile**. The **Edit Profile** dialog box opens.

3. On the **Authentication** tab, click **Authenticate requests on this server**, and then select the **Protected EAP** check box.

4. Click **Configure Certificate**, select the certificate you want IAS to use to verify its identity to client computers, and then click **OK** three times to close all dialog boxes and return to the IAS console.

**Note**

If you access the profile of a connection request policy in the IAS console and you cannot see the **Protected EAP** check box or the **Configure Certificate** button, you must first configure IAS for compatibility with WPS technology as described in "Configuring IAS for WPS technology" in this paper.

**Configure the enterprise account processing server**

For the account processing server, if you are using Internet Information Services (IIS), do the following:

1. Install Internet Information Services and ASP.NET.

2. For more information, see "To install Internet Information Services (IIS) 6.0" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20033.

3. Verify that you have .NET Framework 1.1 installed by clicking **Start** on your Windows desktop, selecting **Control Panel**, and then double-clicking the **Add or Remove Programs** icon. When that window appears, scroll through the list of programs. If you see Microsoft .NET Framework 1.1 listed, the latest version is already installed and you do not need to install it again.

   a. If Microsoft .NET Framework 1.1 is already installed, register ASP.NET with Microsoft .NET Framework 1.1 by running the following command at Command Prompt:

      **%WINDIR%\Microsoft.NET\Framework\v1.1.4322\aspnet_regiis.exe –I**

   b. If Microsoft .NET Framework 1.1 is not installed, install it by using Microsoft Windows Update at http://go.microsoft.com/fwlink/?LinkId=284.

      For more information, see "How to Get the .NET Framework 1.1" at http://go.microsoft.com/fwlink/?LinkId=116211.

4. Create two folders in the Web server root location %systemroot%\Inetpub\wwwroot. You can use one folder to hold your custom Web application and one folder to hold the XML master and subfiles that you will be creating in later steps. For example, you can create a folder named **wpsdeploy** (%systemroot%\Inetpub\wwwroot\wpsdeploy) to contain your Web application, and you can create a folder named **wpsfiles** (%systemroot%\Inetpub\wwwroot\wpsfiles) to contain your XML files.

5. Install your Web application into the folder you created for the Web application files.

6. Set user permissions for the Web application. Open Windows Explorer and browse to the location where you installed your Web application. For example, if you named your folder wpsdeploy, browse to %systemroot%\Inetpub\wwwroot\wpsdeploy. Right-click the wpsdeploy folder, and then click **Sharing and Security**. On the **Security** tab, click **Add**. In **Enter the object names to select**, type **Everyone**, and then click **OK**. In **Permissions for Everyone**, enable **Write** permissions by selecting the **Allow** check box.

7. Enable HTTPS. You must configure IIS to use a certificate for secure Web communications between the account processing server and the provisioning server. To enable HTTPS you must install a server certificate obtained from a public trusted root certification authority, such as a certificate from Verisign or Thawte. When you obtain the certificate, it must conform to the minimum server certificate requirements described in this paper.

For more information, see "Network access authentication and certificates" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20016.

To enable Secure Sockets Layer (SSL) and HTTPS, complete the following procedures.

| **Note** |
| --- |
| If you are deploying a certification authority in a test lab environment, you must install and configure the CA before completing the following procedures. |

**To obtain a new server certificate using the Web Server Certificate Wizard**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.

2. Right-click the Web site or file that you want, and then click **Properties**.

3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.

4. In the Web Server Certificate Wizard, click **Create a new certificate**.

5. Complete the **Web Server Certificate Wizard**, which will guide you through the process of requesting a new server certificate.

**To install a server certificate using the Web Server Certificate Wizard**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.

2. Right-click the Web site or file that you want, and then click **Properties**.

3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.

4. In the Web Server Certificate Wizard, click **Assign an existing certificate**.

5. Complete the **Web Server Certificate Wizard**, which will guide you through the process of installing a server certificate.

| **Note** |
| --- |
| The Web Server Certificate Wizard always denotes this step as "assigning" a certificate to a resource (such as a file, directory, or site), not as "installing." |

**Configure the database on the enterprise SQL server**

On the computer running SQL Server 2000 or a third-party product with similar functionality, create a promotion code database with the following fields, using the appropriate data type for each field:

- **Promotion code.** This field contains promotion codes that you distribute publicly to potential customers. When customers sign up for an account, they provide the promotion code that is matched by the Web application to a value in this field in the database.

- **User name.** This field has no value assigned until a customer creates an account. At this time, the Web application assigns a value to this field.

- **Domain name.** This field contains the domain name where you want the Web application to create the user account when a customer signs up using the promotion code.

- **Security group.** The Web application joins the new user account to the security group defined in this field.

- **Expiration date.** If your promotion lasts for a specific period of time, you can enter the expiration date related to the promotion code in this field. If a user with a promotion code attempts to create an account after the expiration date for the promotion, they cannot do so.

Populate the database with records, providing values for all fields except user name, and enable the data link between your Web application and the SQL server. Also configure security and authentication on the SQL server so that your Web application has permission to access and write to the database. For more information, see SQL Server at http://go.microsoft.com/fwlink/?LinkId=20014.

| **Important** |
| --- |
| When you configure values for records in the SQL Server database, the following three items must match: the name of the security group in Active Directory, the value of the **security group** field in the SQL Server database, and the name of the security group configured in the Valid User access policy in IAS. The Web application uses the value of the SQL Server database **security group** field to determine group membership for new accounts. |

**Install and configure the enterprise router**

The router must be capable of forwarding traffic between the enterprise LAN and the perimeter network.

**Configure the enterprise XML master file and subfiles**

There are two methods you can use to create and configure your XML master and subfiles:

- Use the WPS Authoring Tool to create a WPS project and publish your XML files. The WPS Authoring Tool has a graphical user interface and is designed to assist you in accurately producing and managing a collection of XML files for your WPS solution. For more information, see "Using the WPS Authoring Tool" at http://go.microsoft.com/fwlink/?LinkId=41067. Using the WPS Authoring Tool to create your XML data files is recommended.

- Use the XML schemas provided in this paper to create your files. Once you have created these files, you can enter information specific to your network and deployment parameters. For example, where the location of the provisioning server is required, you can provide an HTTPS URL. In another example, you might need to enter your domain name in several places; you can examine the schemas and example files and determine where to insert your domain name.

When your XML files are configured with the information for your organization, you can store them on your provisioning server and configure your Web application so that it can find the files when necessary. If you are using Internet Information Services as your Web server, you can install the files at the location *%systemroot%*\inetpub\wwwroot.

**Configure the enterprise provisioning server**

For the provisioning server, do the following:

1. Install Internet Information Services and ASP.NET.

2. For more information, see "To install Internet Information Services (IIS) 6.0" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20033.

3. Verify that you have .NET Framework 1.1 installed by clicking **Start** on your Windows desktop, selecting **Control Panel**, and then double-clicking the **Add or Remove Programs** icon. When that window appears, scroll through the list of programs. If you see Microsoft .NET Framework 1.1 listed, the latest version is already installed and you do not need to

install it again.

a. If Microsoft .NET Framework 1.1 is already installed, register ASP.NET with Microsoft .NET Framework 1.1 by running the following command at Command Prompt:

   **%WINDIR%\Microsoft.NET\Framework\v1.1.4322\aspnet_regiis.exe –I**

b. If Microsoft .NET Framework 1.1 is not installed, install it by using Microsoft Windows Update at http://go.microsoft.com/fwlink/?LinkId=284.

   For more information, see "How to Get the .NET Framework 1.1" at http://go.microsoft.com/fwlink/?LinkId=30841.

4. Create two folders in the Web server root location %systemroot%\Inetpub\wwwroot. You can use one folder to hold your custom Web application and one folder to hold the XML master file and subfiles that you will be creating in later steps. For example, you can create a folder named **wpsdeploy** (%systemroot%\Inetpub\wwwroot\wpsdeploy) to contain your Web application, and you can create a folder named **wpsfiles** (%systemroot%\Inetpub\wwwroot\wpsfiles) to contain your XML files.

5. Install your Web application into the folder you created for the Web application files.

6. Set user permissions for the Web application. Open Windows Explorer and browse to the location where you installed your Web application. For example, if you named your folder wpsdeploy, browse to %systemroot%\Inetpub\wwwroot\wpsdeploy. Right-click the wpsdeploy folder, and then click **Sharing and Security**. On the **Security** tab, click **Add**. In **Enter the object names to select**, type **Everyone**, and then click **OK**. In **Permissions for Everyone**, enable **Write** permissions by selecting the **Allow** check box.

7. Enable HTTPS. You must configure IIS to use a certificate for secure Web communications between the provisioning server and clients. To enable HTTPS you must install a server certificate obtained from a public trusted root certification authority, such as a certificate from Verisign or Thawte. When you obtain the certificate, it must conform to the minimum server certificate requirements described in this paper.

For more information, see "Network access authentication and certificates" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20016.

To enable Secure Sockets Layer (SSL) and HTTPS, complete the following procedures.

| **Note** |
| --- |
| If you are deploying a certification authority in a test lab environment, you must install and configure the CA before completing the following procedures. |

**To obtain a new server certificate using the Web Server Certificate Wizard**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.

2. Right-click the Web site or file that you want, and then click **Properties**.

3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.

4. In the Web Server Certificate Wizard, click **Create a new certificate**.

5. Complete the **Web Server Certificate Wizard**, which will guide you through the process of creating a new server certificate.

**To install a server certificate using the Web Server Certificate Wizard**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.

2. Right-click the Web site or file that you want, and then click **Properties**.

3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.

4. In Web Server Certificate Wizard, click **Assign an existing certificate**.

5. Complete the **Web Server Certificate Wizard**, which will guide you through the process of installing a server certificate.

| **Note** |
| --- |
| Web Server Certificate Wizard always denotes this step as "assigning" a certificate to a resource (such as a file, directory, or site), not as "installing." |

**Configure the enterprise DHCP server**

On a computer running Windows Server 2003:

1. Install DHCP.

   For more information, see "To install a DHCP server" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20034.

2. In the DHCP console, run New Scope Wizard twice. Create two VLAN scopes from which IP addresses will be leased to wireless clients connected to the VLANs. Each scope must define a different IP address range using either a private address range or a public IP address range. If you are using network address translation (NAT), you can use a private IP address range; otherwise, the IP addresses leased to wireless clients must be from a public IP address range.

   For more information, see "To create a new scope" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20123.

3. While running the New Scope Wizard, create an exclusion range for the IP addresses you will be assigning statically. For example, if you need to statically assign ten IP addresses from the address range 10.1.1.1 through 10.1.1.254, your exclusion range is defined as 10.1.1.1 through 10.1.1.10.

4. While running the New Scope Wizard, choose to assign scope options. On the **Configure DHCP Options** page, select **Yes, I want to configure these options now**. Scope options are applied only to leases of addresses from within the IP address range that the scope defines, which provides flexibility as your network grows. Define DNS server and Domain name options, as well as any other options that are appropriate for your network configuration.

5. While running the New Scope Wizard, activate the scope. The option to activate the scope while running the wizard is available only if you have chosen to configure DHCP options in the previous steps.

   For more information, see "To activate a scope" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20124.

6. Authorize the DHCP server in Active Directory.

   For more information, see "To authorize a DHCP server in Active Directory" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20125.

The DHCP server is now online and able to provide IP address leases to client computers. In some cases you might want to examine the types and durations of accounts you offer your users and adjust the DHCP lease duration accordingly. In most cases, when deploying WPS technology the default lease duration of eight days is too long and should be shortened considerably.

**DHCP example**

In this example, you are creating two VLANs and two scopes, one scope for each VLAN.

**VLAN 2**

VLAN 2 is the Network Resource VLAN that provides access to network resources (such as the IAS server and DHCP server) for wireless computers connecting as guest. VLAN 2 blocks access to the Internet, however. The DHCP scope for VLAN 2 is defined with the following example parameters:

- **Address range:** 192.168.1.1 through 192.168.1.254. This is a private IP address range. If you are using NAT, you can use this range on your network. If you are not using NAT, use a public IP address range. In addition, if your wireless deployment is large, select an IP address range that provides more IP addresses for lease to clients.

- **Exclusion range:** 192.168.1.1 through 192.168.1.10. By using this exclusion range, the available address pool for clients is 192.168.1.11 through 192.168.1.254. Ten IP addresses are excluded so that you can statically assign these addresses to computers and devices on your network. For example, the router IP address must be statically assigned on the router.

- **DHCP scope option 003, Router:** 192.168.1.1. The router IP address must be an address that falls within the exclusion range so that the DHCP server does not lease the router IP address to a wireless client computer, thereby creating an address conflict.

- **DHCP scope option 006, DNS server:** the IP address of the Active Directory and DNS server on the enterprise LAN

Note that DHCP scope option 003, Router, provides client computers with the IP address of their default gateway IP address. In this case, the default gateway for wireless clients is the VLAN-aware gateway device, whether it is an access controller, a VLAN-aware router, a VLAN-aware switch, or another compatible device. When you configure your VLAN-aware gateway device, you can specify the IP address that the device uses for each VLAN.

In this example, you must configure the VLAN-aware gateway device so that it uses the IP address 192.168.1.1 on VLAN 2.

For your WPS deployment you can configure additional DHCP options as needed. For example, if you are using a WINS server, you can configure your VLAN scopes with DHCP option 044, WINS/NBNS servers.

**VLAN 4**

VLAN 4 is the Internet VLAN that provides access to the Internet. Users who have successfully created an account are switched to this VLAN after completing the provisioning and sign-up process. The DHCP scope for VLAN 4 is defined with the following example parameters:

- **Address range:** 192.168.2.1 through 192.168.2.254

- **Exclusion range:** 192.168.2.1 through 192.168.2.10

- **DHCP scope option 003, Router:** 192.168.2.1. The router IP address must be an address that falls within the exclusion range so that the DHCP server does not lease the router IP address to a wireless client computer, thereby creating an address conflict.

- **DHCP scope option 006, DNS server:** the IP address of the Active Directory and DNS server on the Enterprise LAN

For VLAN 4 in this example, you must configure the VLAN-aware gateway device so that it uses the IP address 192.168.2.1 on VLAN 4.

**Install and configure your enterprise VLAN-aware gateway device**

Configure two VLANs on the gateway device: a Network Resource VLAN that provides access to the WISP LAN, and an Internet VLAN that provides access to the Internet.

The remote access policies you created in IAS determine which VLAN your customers can access:

- The Guest access policy places users on the Network Resource VLAN so that they can create and pay for a valid user account.

- The Valid Users access policy in IAS places customers on the Internet VLAN.

Each VLAN has a different IP address range. When configuring your DHCP server, you created a scope for each VLAN, and you defined DHCP scope option 003, Router. This is the IP address commonly referred to as the “default gateway.”

You must configure the VLAN-aware gateway device as the default gateway for each VLAN, using an IP address from the IP address range that you defined on your DHCP server.

**Choosing the IP address for the default gateway for each VLAN**
Your VLAN-aware gateway device is the default gateway for both VLANs and is configured with a different IP address for each VLAN.

| **Important** |
| --- |
| In each scope you configure on the DHCP server, the value you enter for DHCP scope option 003, Router, must match the IP address you assign to the VLAN-aware gateway device for use on each VLAN. For example, if you configure a scope on the DHCP server for VLAN 2 with the IP address range 192.168.1.1 through 192.168.1.254, and you assign the DHCP scope option 003, Router, with the value 192.168.1.1, you must configure the VLAN-aware gateway device to use the IP address 192.168.1.1 on VLAN 2. |

Similarly, as described in "Configure the DHCP server" in this paper, configure the VLAN-aware gateway device so that it uses an IP address from the exclusion range 192.168.2.1 through 192.168.2.10 for VLAN 4. For example, you can configure the VLAN-aware gateway device so that it uses the IP address 192.168.2.1 on VLAN 4.

In some circumstances you might prefer to use your VLAN-aware gateway device as the DHCP server for each VLAN. If this is the case, you can define IP address ranges and exclusion ranges on the VLAN-aware gateway device rather than on a DHCP server.

See the product documentation for your VLAN-aware gateway device for information about configuring your hardware.

| **Important** |
| --- |
| The Internet VLAN integer must match the value you configure for the Tunnel-Pvt-Group-ID attribute in the Valid Users access policy on your IAS server. For example, if VLAN 4 provides access to the Internet, the value of the Tunnel-Pvt-Group-ID attribute in the profile of the Valid Users access policy must be 4. |

**Install and configure your enterprise wireless access points**
Configure the access points to use your RADIUS server, including configuration of the server IP address and the shared secret. Follow the directions in your access point documentation for other configuration settings, using the following guidelines:

- **Authentication or RADIUS server:** Specify your IAS server by IP address or FQDN, depending on the requirements of the AP.
- **SSID:** Specify a Secure Set Identifier (SSID), which is an alphanumeric string that serves as the network name. This name is broadcast by APs to wireless clients and is visible to users at your Wi-Fi hotspots.
- **RADIUS settings:** Use RADIUS authentication on User Datagram Protocol (UDP) port 1812 and use RADIUS accounting on UDP port 1813.
- **Secret or shared secret:** Use a strong shared secret and configure the IAS server with the same shared secret.
- **EAP:** Configure the AP to require EAP from wireless clients.
- **802.1X and WEP:** Enable IEEE 802.1X authentication and WEP.

If your IAS server is running Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition, you can configure RADIUS clients by IP address range. This is a useful feature when you have a large number of access points to deploy; if you deploy your access points on the same subnet or VLAN within the same IP address range, configuration of RADIUS clients in IAS is simplified. Instead of individually configuring each access point as a RADIUS client in IAS, you can configure all access points at once using the IP address range of the subnet or VLAN upon which the APs reside. In this circumstance, use the same shared secret for all access points, and make sure that the shared secret is strong.

**Install and configure the enterprise IAS proxy**
For IAS proxy, the three configuration stages are general configuration, RADIUS client configuration, and connection request policy configuration.

**General configuration**

1. Install Internet Authentication Service.

   For more information, see "To install IAS" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20028.

2. Delete the default remote access policies. To delete the policies, open the IAS console, and then click **Remote Access Policies**. Select each existing policy, right-click the policy, and then click **Delete**.

**RADIUS client configuration**

In the IAS console, add the wireless access points as RADIUS clients. In addition, configure the shared secret used between the IAS proxy and the access points. For more information, see "To add RADIUS clients" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20031 and "To configure the Message Authenticator attribute and shared secret" at http://go.microsoft.com/fwlink/?LinkId=20032.

If your IAS server is a computer running Windows Server 2003, Enterprise Edition with SP1, or Windows Server 2003, Datacenter Edition with SP1, and if you have configured your wireless access points with IP addresses from the same IP address range, you can configure the access points as RADIUS clients using the IP address range rather than individually configuring each access point.

| **Note** |
| --- |
| You can configure IAS in Windows Server 2003, Standard Edition with SP1, with a maximum of 50 RADIUS clients. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. With Windows Server 2003, Standard Edition with SP1, if the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition with SP1, and Windows Server 2003, Datacenter Edition with SP1, you can configure an unlimited number of RADIUS clients. In addition, you can configure RADIUS clients by specifying an IP address range. |

**Connection request policy configuration**

When you configure a remote RADIUS server group in the IAS console of the proxy server, you instruct the proxy server to forward messages to the IAS servers included in the remote RADIUS server group. To configure the IAS server on the Enterprise LAN as a member of a remote RADIUS server group, do the following:

1. In the IAS console, double-click **Connection Request Processing**. The folder expands to display Connection Request Policies and Remote RADIUS Server Groups. Right-click **Remote RADIUS Server Groups**, and then select **New Remote RADIUS Server Group**. New Remote RADIUS Server Group Wizard is started. Click **Next**.

2. In **Group Configuration Method**, click **Custom**. Type a value for Group Name. Click **Next**.

3. In **Add Servers**, click **Add**. In **Add RADIUS Server**, type the IP address of the IAS server on the enterprise LAN. To verify connectivity, click **Verify**.

4. Click the **Authentication/Accounting** tab. In **Shared secret** and **Confirm shared secret**, type the shared secret that you will also use on the IAS server. Click **OK**.

5. In **Add Servers**, click **Next**, and then click **Finish**.

6. The **New Connection Request Policy Wizard** is started. Click **Next**. In **Policy Configuration Method**, verify that the type of policy is set to **A typical policy for a common scenario**. In **Policy Name**, type the name for your connection request policy. Click **Next**.

7. In **Request Authentication**, select **Forward connection requests to a remote RADIUS server for authentication**. Click **Next**.

8. In **Realm Name**, type the realm name of the connection requests that will be forwarded. Also verify that the name of the remote RADIUS server group that you created is selected in **Server group**. Click **Next**, and then click **Finish**.

   For more information, see "Realm Names" in Help and Support Center for Windows Server 2003 or on the Web at

http://go.microsoft.com/fwlink/?LinkId=34426.

9. In the IAS console, under **Connection Request Processing**, click **Connection Request Policies**. In the right pane of the IAS console, there are two connection request policies: the policy you just created, and the default connection request policy. The default connection request policy is named **Use Windows authentication for all users**. To delete the default policy, right-click **Use Windows authentication for all users**, and then click **Delete**. To confirm the deletion, click **Yes**.

| **Note** |
| --- |
| If you access the profile of a connection request policy in the IAS console and you cannot see the **Protected EAP** check box or the **Configure Certificate** button, you must first configure IAS for compatibility with WPS technology as described in "Configuring IAS for WPS technology" in this paper. |

**Configure RADIUS clients at the enterprise IAS server**

In the IAS console, add the IAS proxy server as a RADIUS client. In addition, configure the shared secret used between the IAS proxy and the IAS server.

| **Important** |
| --- |
| Because you are using an IAS proxy between the IAS server and the access points, do not configure the access points as RADIUS clients to the IAS server. Access points are configured as RADIUS clients at the IAS proxy only. |

For more information, see "To add RADIUS clients" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20031 and "To configure the Message Authenticator attribute and shared secret" at http://go.microsoft.com/fwlink/?LinkId=20032.

**Configure the enterprise Windows XP-based client computer**

On the computer running Windows XP Professional, Windows XP Tablet PC Edition, or Windows XP Home Edition, install SP2. The computer must have a wireless network adapter compatible with IEEE 802.11 and 802.1X.

| **Note** |
| --- |
| If you are deploying WPS technology in a test lab environment where you have deployed your own enterprise root CA, you must import the CA certificate into the Trusted Root Certification Authorities certificate store on the client computer. If this certificate is not imported into the client certificate store, the client computer will not trust the enterprise root CA and server authentication to the client will fail. For more information, see "Configure certificates in an enterprise test lab environment (optional)" in this paper. |

**Configure certificates in an enterprise test lab environment (optional)**

If you are deploying WPS technology in a test lab, you can obtain server certificates from a public trusted root CA or you can deploy Certificate Services for Windows Server 2003 on your test network. To deploy Certificate Services and enroll certificates to domain member computers (such as the servers on your networks), do the following:

1. On a computer running Windows Server 2003, install Certificate Services.

   For more information, see "To install an enterprise root certification authority" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20035.

2. Add a server certificate template to the certification authority and configure the certification authority to allow computers to request a certificate that is based on the template you create.

   For WPS technology in a test lab environment, you must create a server certificate that will be trusted by client computers. The server certificate must meet the requirements stated in this paper. In addition, you must base your certificate on the correct certificate template for the operating system you are running. If you are running an enterprise certification authority (CA) on a computer running Windows Server 2003, Standard Edition, and your IAS server is a domain member, base your certificate on the **Computer** certificate template. If you are running an enterprise certification authority (CA) on a computer running Windows Server 2003, Enterprise Edition, and your IAS server is a domain member, base your certificate on the **RAS and IAS Server** certificate template.

For more information, see "To create a certificate template" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20036.

3. Autoenroll certificates to domain member computers. Autoenrollment allows domain member computers to automatically obtain, or enroll, certificates. For WPS technology, autoenroll certificates to servers using the certificate template you created in the previous step.

   For more information, see "Planning for autoenrollment deployment" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20037 and "Checklist: Configuring certificate autoenrollment" at http://go.microsoft.com/fwlink/?LinkId=20038.

4. Install the enterprise CA certificate in the Trusted Root Certification Authority certificate store on client computers being used for testing purposes. You can request the enterprise root CA certificate by using Web enrollment services, which is installed with Certificate Services, or you can export the server certificate to a floppy disk, and then import the certificate into the Trusted Root Certification Authority certificate store on the client.

   For more information, see "To export a certificate" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=20039 and "To import a certificate" at http://go.microsoft.com/fwlink/?LinkId=20040.

# WISP Roaming Agreement Deployments

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In some circumstances, two or more wireless ISPs might find it beneficial to create a roaming agreement that allows customers from one WISP to connect to the Internet through the Wi-Fi hotspots owned and operated by the other WISP or WISPs.

For example, one WISP might operate Wi-Fi hotspots in the food court at an airport, while another WISP operates Wi-Fi hotspots at the boarding gates of the same airport. By entering into a roaming agreement, the WISPs provide their customers with the ability to sign up at the food court WISP, roam to the boarding gate WISP, and remain connected to the Internet.

The following two graphics illustrate how a roaming agreement works when two WISPs offer services to customers at the same physical location.

In the illustration below, a new customer connects to WISP_1 at a Wi-Fi hotspot located in the food court of an airport, then roams to the airport boarding gates and connects to an AP at a Wi-Fi hotspot owned and operated by WISP_2.




The following text describes the numbered process in the illustration above:

1. A new customer arrives at the airport food court and connects to WISP_1. The customer runs the sign-up wizard, and creates and pays for an account with WISP_1.

2. The customer is authenticated and authorized by the WISP_1 IAS server and is granted access to the Internet.

3. The WISP_1 customer walks from the airport food court to the airport boarding gates.

4. The WISP_2 AP SSID exists in the WISP_1 SSID XML subfile, and the WISP_1 customer's computer associates with the WISP_2 AP. WPS technology on the client computer automatically begins the authentication process using credentials from the WISP_1 account created at the food court.

5. The WISP_2 IAS proxy is configured to forward access requests containing the WISP_1 realm name to the WISP_1 IAS server. The WISP_1 customer Access-Request message is therefore forwarded to the WISP_1 IAS server.

6. The WISP_1 IAS server authenticates and authorizes the user, and sends an Access-Accept message to the WISP_2 IAS proxy. The WISP_2 IAS proxy forwards the Access-Accept to the AP.

7. The WISP_1 customer is granted access to the Internet through the AP owned and operated by WISP_2.

If you have a WISP that you want to configure for use with WPS technology, you can use the scenarios depicted in this paper to deploy the technology. If you want to establish a roaming agreement with another WISP, you can perform the additional steps that follow.

**Key roaming agreement configuration steps**

If your WISP (WISP_1) has a roaming agreement with another WISP (WISP_2) that allows your customers to connect to the

Internet through WISP_2 APs, perform the following steps:

1. Add SSIDs that are advertised by WISP_2 APs to your SSID XML subfile. For more information, see "Using the WPS Authoring Tool" at http://go.microsoft.com/fwlink/?LinkId=41067.

2. Configure the WISP_2 IAS proxy as a RADIUS client on your IAS server. When your customers connect to WISP_2 APs, the WISP_2 IAS proxy forwards authentication requests to your IAS server for authentication and authorization. For more information, see "To add RADIUS clients" at http://go.microsoft.com/fwlink/?LinkId=20031 and "To configure the Message Authenticator attribute and shared secret" at http://go.microsoft.com/fwlink/?LinkId=20032.

3. Configure IAS logging to provide session correlation so that you can effectively use the IAS logs for billing purposes. For more information, see "Remote Access Logging" at http://go.microsoft.com/fwlink/?LinkId=41038 and "Deploying SQL Server Logging with Windows Server 2003 Internet Authentication Service (IAS)" at http://go.microsoft.com/fwlink/?LinkId=41039.

WISP_2 allows customers of WISP_1 to connect to WISP_2 APs. WISP_2 uses IAS as a RADIUS proxy, including the use of connection request policy and a remote RADIUS server group, to identify WISP_1 customers and forward their access requests to the IAS server at WISP_1 for authentication and authorization.

If your WISP functions as WISP_2 in the roaming agreement scenario, perform the following steps:

1. Create a remote RADIUS server group that contains the IAS servers for WISP_1. In the IAS console, run the New Remote RADIUS Server Group wizard: in **Connection Request Processing**, right-click **Remote RADIUS Server Groups**, and then click **New Remote RADIUS Server Group**. After you have run the New Remote RADIUS Server Group wizard, the New Connection Request Policy wizard is automatically launched.

2. Create a connection request policy that forwards WISP_1 customer connection requests to the WISP_1 IAS server. In the IAS console, run the New Connection Request Policy wizard and create a new connection request policy that is configured to **Forward connection requests to a remote RADIUS server for authentication**. In **Realm name**, type the realm name of WISP_1. For example, if the realm name for WISP_1 is example.com, type **example.com**.

# Server Certificate Requirements

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certificates installed on your servers must be obtained from a public trusted root certification authority (such as Verisign or Thawte) and must meet the minimum server certificate requirements. WPS technology uses the PEAP-MS-CHAP v2 authentication method; with PEAP-MS-CHAP v2, the client accepts the server's authentication attempt when the server certificate meets the following requirements:

- The Subject name contains a value. If you issue a certificate to your IAS server that has a blank Subject, the certificate is not available to authenticate your IAS server.
- The computer certificate on the server chains to a public trusted root CA and does not fail any of the checks that are performed by CryptoAPI and specified in the remote access policy.
- The IAS server computer certificate is configured with the Server Authentication purpose in Enhanced Key Usage (EKU) extensions. (The object identifier for Server Authentication is 1.3.6.1.5.5.7.3.1.)
- The server certificate is configured with a required cryptographic service provider (CSP) value of Microsoft RSA SChannel Cryptographic Provider.
- The Subject Alternative Name (SubjectAltName) extension, if used, must contain the DNS name of the server.

With PEAP and EAP-TLS, servers display a list of all installed certificates in the computer certificate store, with the following exceptions:

- Certificates that do not contain the Server Authentication purpose in EKU extensions are not displayed.
- Certificates that do not contain a Subject name are not displayed.
- Registry-based and smart card-logon certificates.

| **Note** |
| --- |
| When you deploy PEAP-MS-CHAP v2 on a private network, client computers are configured to validate server certificates by using the **Validate server certificate** option. With WPS technology, however, Wireless Provisioning Services automatically configures this option on client computers running Windows XP. |

# Active Directory Replication

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are deploying WPS technology in multiple physical locations, calculate your customers' travel times between these locations and adjust Active Directory replication so that user accounts replicate quickly enough for users to log on at different locations. Latency in Active Directory replication might temporarily affect the ability of a customer or user to log on to the network at a different location.

For example, if you deploy Wi-Fi hotspots at a shopping mall in Helsinki, Finland, and at an airport in Florence, Italy, imagine a customer creating an account at one location and then traveling to the other location. Calculate the minimum amount of time it will take your customer to travel from Helsinki to Florence (or from Florence to Helsinki), and then verify that Active Directory replication between these two locations occurs in a shorter amount of time. Thus if it takes a customer eight hours to travel from Helsinki to Florence, configure Active Directory replication between Helsinki and Florence to occur within seven hours of the creation of the account.

If a user account created at one location has not replicated to another location by the time your customer arrives there, authentication will fail and the customer will not be able to log on to your network using their new account credentials.

For more information, see the "Active Directory Replication Topology Technical Reference" on the Web at http://go.microsoft.com/fwlink/?LinkId=41041.

# XML Schemas

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following XML schemas are supported for WPS technology:

- **Master file schema.** The master file schema is used to provide the locations of XML subfiles to client computers.
- **Wizard schema.** The wizard schema is used by the sign-up and renewal wizards for account creation or renewal.
- **Register schema.** The register schema contains a node with 3 subnodes: one for sign-up, one for renewal, and one for password expired. The text for each sub-node is a URL that points to the corresponding xml file. The signup sub-node is required, however the renewal and password expired sub-nodes are optional.
- **SSID schema.** The SSID schema is used to provide clients with network Secure Set Identifier information.
- **Locations schema.** The Locations schema is used to provide customers with a list of locations from which they can access your network.
- **Branding schema.** The Branding schema is used to identify your company to your customers with artwork, such as your company logo, that is downloaded and displayed in the client computer user interface.
- **Help schema.** The Help schema is used to provide customers with technical information and assistance.
- **EAP schemas.** The EAP schemas are used to configure connection and user properties on client computers.

The next section describes the XML for each schema.

**Master file schema**

The following XML describes the XML master file schema for use by HSPs, WISPs, and enterprises to provide client computers with the locations of the XML subfiles and network information.

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
targetNamespace="http://www.microsoft.com/provisioning/Master"
xmlns="http://www.microsoft.com/provisioning/Master" elementFormDefault="qualified">
    <xs:import namespace="http://www.w3.org/XML/1998/namespace"
schemaLocation="http://www.w3.org/2001/xml.xsd" />
    <xs:element name="Master">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="TTL" type="xs:positiveInteger" />
                <xs:element name="DomainName" type="xs:string" />
                <xs:element name="UpdateFrom" type="anyHttps" />
                <xs:element name="LangSet" minOccurs="0" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="Subfile" maxOccurs="unbounded">
                                <xs:complexType>
                                    <xs:sequence>
                                        <xs:element name="URL" type="anyHttps" />
                                        <xs:element name="Version"
type="xs:positiveInteger" />
                                    </xs:sequence>
                                    <xs:attribute name="Name" type="xs:string"
use="required" />
                                </xs:complexType>
                            </xs:element>
                        </xs:sequence>
                        <xs:attribute ref="xml:lang" use="required" />
                    </xs:complexType>
                </xs:element>
```

```
                <xs:element name="OtherMasters" minOccurs="0">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="MasterURL" type="anyHttps"
maxOccurs="unbounded" />
                        </xs:sequence>
                    </xs:complexType>
                </xs:element>
            </xs:sequence>
        </xs:complexType>
    </xs:element>
    <xs:simpleType name="anyHttps">
        <xs:restriction base="xs:string">
            <xs:pattern value="https://(.)+" />
        </xs:restriction>
    </xs:simpleType>
</xs:schema>
```

**Wizard schema**

The following XML describes the Wizard XML schema used by the sign-up and renewal wizards while customers create or renew an account.

```
<?xml version="1.0"?>
<xs:schema targetNamespace="http://www.microsoft.com/provisioning/Wizard"
xmlns:mstns="http://www.microsoft.com/provisioning/Wizard"
    xmlns="http://www.microsoft.com/provisioning/Wizard"
xmlns:xs="http://www.w3.org/2001/XMLSchema"
    xmlns:msdata="urn:schemas-microsoft-com:xml-msdata"
attributeFormDefault="qualified" elementFormDefault="qualified">
    <xs:element name="wizard">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="purpose" minOccurs="1" maxOccurs="1">
                    <xs:simpleType>
                        <xs:restriction base="xs:string">
                            <xs:enumeration value="signup" />
                            <xs:enumeration value="renewal" />
                            <xs:enumeration value="password" />
                        </xs:restriction>
                    </xs:simpleType>
                </xs:element>
                <xs:element name="postToUrl" type="xs:string" minOccurs="1"
maxOccurs="1" />
                <xs:element name="panel-welcome" minOccurs="1" maxOccurs="1">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="icon" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="branding" type="xs:string" minOccurs="0"
maxOccurs="1" />
                        </xs:sequence>
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-plan" minOccurs="1" maxOccurs="1">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" msdata:Ordinal="0" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" msdata:Ordinal="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" msdata:Ordinal="2" />
                            <xs:element name="plan" minOccurs="1"
maxOccurs="unbounded">
                                <xs:complexType>
                                    <xs:sequence>
```

```
                                        <xs:element name="title" type="xs:string"
minOccurs="0" maxOccurs="1" msdata:Ordinal="0" />
                                        <xs:element name="description" type="xs:string"
minOccurs="0" maxOccurs="1" msdata:Ordinal="1" />
                                        <xs:element name="fulldescription"
minOccurs="0" maxOccurs="1" msdata:Ordinal="2" >
                                            <xs:complexType>
                                                <xs:sequence>
                                                    <xs:element name="MHTML"
type="xs:string" minOccurs="1" maxOccurs="1"/>
                                                </xs:sequence>
                                            </xs:complexType>
                                        </xs:element>
                                        <xs:element name="includePanels" minOccurs="1"
maxOccurs="1">
                                            <xs:complexType>
                                                <xs:sequence>
                                                    <xs:element name="includePanel"
minOccurs="0" maxOccurs="unbounded">
                                                        <xs:complexType>
                                                            <xs:attribute name="name"
form="unqualified" type="xs:string" />
                                                            <xs:attribute name="id"
use="required" form="unqualified" type="xs:string" />
                                                        </xs:complexType>
                                                    </xs:element>
                                                </xs:sequence>
                                            </xs:complexType>
                                        </xs:element>
                                    </xs:sequence>
                                    <xs:attribute name="category" form="unqualified"
type="xs:string" />
                                    <xs:attribute name="selected" form="unqualified"
type="xs:boolean" />
                                </xs:complexType>
                            </xs:element>
                        </xs:sequence>
                        <xs:attribute name="hint" use="optional" form="unqualified"
type="xs:string" />
                        <xs:attribute name="error" use="optional" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-promotion" minOccurs="0" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="promotionalcode" type="editField"
minOccurs="0" maxOccurs="1" nillable="true" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-personal" minOccurs="0" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="salutation" type="comboBox" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="first-name" type="editField"
minOccurs="0" maxOccurs="1" nillable="true" />
```

```
                            <xs:element name="middle-initial" type="editField"
minOccurs="0" maxOccurs="1" nillable="true" />
                            <xs:element name="last-name" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="suffix" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="company" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="jobtitle" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="language-preference" type="comboBox"
minOccurs="0" maxOccurs="1" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-contact" minOccurs="0" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="address1" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="address2" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="city" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="state" type="comboBox" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="country" type="comboBox" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="zipcode" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="telephone1" type="editField"
minOccurs="0" maxOccurs="1" nillable="true" />
                            <xs:element name="telephone2" type="editField"
minOccurs="0" maxOccurs="1" nillable="true" />
                            <xs:element name="email" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-privacy" minOccurs="0" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="check1" type="check box" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="check2" type="check box" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="check3" type="check box" minOccurs="0"
maxOccurs="1" nillable="true" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-creditcard" minOccurs="0"
maxOccurs="unbounded">
                    <xs:complexType>
```

```
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="card-name" type="comboBox" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="card-number" type="nonReadOnlyEditField"
minOccurs="0" maxOccurs="1" nillable="true" />
                            <xs:element name="card-expdate" type="nonReadOnlyEditField"
minOccurs="0" maxOccurs="1" nillable="true" />
                            <xs:element name="card-securitycode"
type="nonReadOnlyEditField" minOccurs="0" maxOccurs="1"
                                nillable="true" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-billing" minOccurs="0" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="first-name" type="editField"
minOccurs="0" maxOccurs="1" nillable="true" />
                            <xs:element name="middle-initial" type="editField"
minOccurs="0" maxOccurs="1" nillable="true" />
                            <xs:element name="last-name" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="suffix" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="address1" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="address2" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="city" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="state" type="comboBox" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="country" type="comboBox" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="zipcode" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="telephone1" type="editField"
minOccurs="0" maxOccurs="1" nillable="true" />
                            <xs:element name="telephone2" type="editField"
minOccurs="0" maxOccurs="1" nillable="true" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-coupon" minOccurs="0" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="coupon" type="editField" minOccurs="1"
maxOccurs="1" nillable="true" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
```

```
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-credentials" minOccurs="0"
maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="username" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                            <xs:element name="password" type="editField" minOccurs="0"
maxOccurs="1" nillable="true" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-terms" minOccurs="1" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="terms" type="xs:string" minOccurs="0"
maxOccurs="1" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-post" minOccurs="1" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
                <xs:element name="panel-finish" minOccurs="1" maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="title" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="subtitle" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="body" type="xs:string" minOccurs="0"
maxOccurs="1" />
                            <xs:element name="info" type="xs:string" minOccurs="0"
maxOccurs="1" />
                        </xs:sequence>
                        <xs:attribute name="id" use="required" form="unqualified"
type="xs:string" />
                    </xs:complexType>
                </xs:element>
            </xs:sequence>
            <xs:attribute name="error" use="optional" form="unqualified"
type="xs:string" />
            <xs:attribute name="hint" use="optional" form="unqualified"
```

```
type="xs:string" />
        </xs:complexType>
    </xs:element>
    <xs:complexType name="editField">
        <xs:simpleContent>
            <xs:extension base="xs:string">
                <xs:attribute name="required" use="optional" form="unqualified"
type="xs:boolean" />
                <xs:attribute name="maxlen" use="optional" form="unqualified"
type="xs:positiveInteger" />
                <xs:attribute name="hint" use="optional" form="unqualified"
type="xs:string" />
                <xs:attribute name="error" use="optional" form="unqualified"
type="xs:string" />
                <xs:attribute name="readonly" use="optional" form="unqualified"
type="xs:boolean" />
            </xs:extension>
        </xs:simpleContent>
    </xs:complexType>
    <xs:complexType name="nonReadOnlyEditField">
        <xs:simpleContent>
            <xs:extension base="xs:string">
                <xs:attribute name="required" use="optional" form="unqualified"
type="xs:boolean" />
                <xs:attribute name="maxlen" use="optional" form="unqualified"
type="xs:positiveInteger" />
                <xs:attribute name="hint" use="optional" form="unqualified"
type="xs:string" />
                <xs:attribute name="error" use="optional" form="unqualified"
type="xs:string" />
            </xs:extension>
        </xs:simpleContent>
    </xs:complexType>
    <xs:complexType name="comboBox">
        <xs:sequence>
            <xs:element name="entry" minOccurs="1" maxOccurs="unbounded"
nillable="true">
                <xs:complexType>
                    <xs:simpleContent>
                        <xs:extension base="xs:string">
                            <xs:attribute name="selected" form="unqualified"
type="xs:boolean" />
                        </xs:extension>
                    </xs:simpleContent>
                </xs:complexType>
            </xs:element>
        </xs:sequence>
        <xs:attribute name="required" use="optional" form="unqualified"
type="xs:boolean" />
        <xs:attribute name="hint" use="optional" form="unqualified" type="xs:string" />
        <xs:attribute name="error" use="optional" form="unqualified" type="xs:string"
/>
        <xs:attribute name="readonly" use="optional" form="unqualified"
type="xs:boolean" />
    </xs:complexType>
    <xs:complexType name="check box">
        <xs:simpleContent>
            <xs:extension base="xs:string">
                <xs:attribute name="check" form="unqualified" type="xs:boolean" />
            </xs:extension>
        </xs:simpleContent>
    </xs:complexType>
</xs:schema>
```

**Register schema**

The following XML describes the Register schema.

```
<?xml version="1.0"?>

<xs:schema id="registerFile"
targetNamespace="http://www.microsoft.com/provisioning/Register"
xmlns="http://www.microsoft.com/provisioning/Register"
```

```
xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:msdata="urn:schemas-microsoft-
com:xml-msdata" attributeFormDefault="qualified" elementFormDefault="qualified">
    <xs:element name="register">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="signup"          type="anyHttps" minOccurs="1"
maxOccurs="1" />
                <xs:element name="renewal"         type="anyHttps" minOccurs="0"
maxOccurs="1" />
                <xs:element name="passwordExpired" type="anyHttps" minOccurs="0"
maxOccurs="1" />
            </xs:sequence>
        </xs:complexType>
    </xs:element>
    <xs:element name="registerFile" msdata:IsDataSet="true"
msdata:EnforceConstraints="False">
        <xs:complexType>
            <xs:choice maxOccurs="1">
                <xs:element ref="register" />
            </xs:choice>
        </xs:complexType>
    </xs:element>
    <xs:simpleType name="anyHttps">
        <xs:restriction base="xs:string">
            <xs:pattern value="https://(.)+"/>
        </xs:restriction>
    </xs:simpleType>
</xs:schema>
```

**SSID schema**

The following XML describes the Secure Set Identifier (SSID) schema.

```
<?xml version="1.0" encoding="utf-8" ?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
                     targetNamespace="http://www.microsoft.com/provisioning/SSID"
                     xmlns="http://www.microsoft.com/provisioning/SSID"
                     elementFormDefault="qualified">
    <xs:element name="SSIDs">
        <xs:complexType>
            <xs:sequence>
                <xs:element maxOccurs="unbounded" name="ssid">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="Priority" type="xs:positiveInteger"
minOccurs="0" />
                            <xs:element name="Connection" minOccurs="0" >
                                <xs:simpleType>
                                    <xs:restriction base="xs:string">
                                        <xs:enumeration value="IBSS" />
                                        <xs:enumeration value="ESS" />
                                    </xs:restriction>
                                </xs:simpleType>
                            </xs:element>
                            <xs:element name="Authentication" minOccurs="0" >
                                <xs:simpleType>
                                    <xs:restriction base="xs:string">
                                        <xs:enumeration value="Open" />
                                        <xs:enumeration value="Shared" />
                                        <xs:enumeration value="WPA" />
                                        <xs:enumeration value="WPAPSK" />
                                    </xs:restriction>
                                </xs:simpleType>
                            </xs:element>
                            <xs:element name="Encryption" minOccurs="0" >
                                <xs:simpleType>
                                    <xs:restriction base="xs:string">
                                        <xs:enumeration value="None" />
                                        <xs:enumeration value="WEP" />
                                        <xs:enumeration value="TKIP" />
                                        <xs:enumeration value="CCMP" />
                                    </xs:restriction>
                                </xs:simpleType>
```

```
                            </xs:element>
                            <xs:element name="KeyIndex" type="xs:positiveInteger"
minOccurs="0" />
                            <xs:element name="IEEE802.1X" type="xs:string"
minOccurs="0" />
                            <xs:element name="Non802.1XURL" type="xs:anyURI"
minOccurs="0" />
                        </xs:sequence>
                        <xs:attribute name="Name" type="xs:string" use="required" />
                    </xs:complexType>
                </xs:element>
            </xs:sequence>
        </xs:complexType>
    </xs:element>
</xs:schema>
```

**Locations schema**

The following XML describes the Locations schema, which provides your customers with the locations of your Wi-Fi hotspots.

```
<?xml version="1.0" ?>
<xs:schema id="hotspotfinder"
targetNamespace="http://www.microsoft.com/provisioning/HotSpotLocation"
xmlns="http://www.microsoft.com/provisioning/HotSpotLocation"
xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:msdata="urn:schemas-microsoft-
com:xml-msdata" attributeFormDefault="qualified" elementFormDefault="qualified">
    <xs:element name="hotspotfinder" msdata:IsDataSet="true"
msdata:EnforceConstraints="False">
        <xs:complexType>
            <xs:choice maxOccurs="unbounded">
                <xs:element name="entry">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="city" type="xs:string" minOccurs="0" />
                            <xs:element name="state" type="xs:string" minOccurs="0" />
                            <xs:element name="country" type="xs:string" minOccurs="0"
/>
                            <xs:element name="zipcode" type="xs:string" minOccurs="0"
/>
                            <xs:element name="operator" type="xs:string" minOccurs="0"
/>
                            <xs:element name="location" type="xs:string" minOccurs="0"
/>
                            <xs:element name="address" type="xs:string" minOccurs="0"
/>
                        </xs:sequence>
                    </xs:complexType>
                </xs:element>
            </xs:choice>
        </xs:complexType>
    </xs:element>
</xs:schema>
```

**Branding schema**

The following XML describes the branding schema, which you can use to identify your company to your customers.

```
<?xml version="1.0" ?>
<xs:schema id="BrandingFile"
targetNamespace="http://www.microsoft.com/provisioning/Branding"
xmlns="http://www.microsoft.com/provisioning/Branding"
xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:msdata="urn:schemas-microsoft-
com:xml-msdata" attributeFormDefault="qualified" elementFormDefault="qualified">
    <xs:element name="Branding">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="FriendlyName" type="xs:string" minOccurs="1"
maxOccurs="1" />
                <xs:element name="PostConnectedAction" type="xs:string" minOccurs="0"
maxOccurs="1"/>
                <xs:element name="ServiceDescription" type="xs:string" minOccurs="1"
maxOccurs="1"/>
                <xs:element name="Logo" type="xs:string" minOccurs="1" maxOccurs="1" />
                <xs:element name="Header" type="xs:string" minOccurs="1"
```

```
maxOccurs="1"/>
            </ xs:sequence>
        </ xs:complexType>
    </ xs:element>

    <xs:element name="BrandingFile" msdata:IsDataSet="true"
msdata:EnforceConstraints="False">
        <xs:complexType>
            <xs:choice maxOccurs="1">
                <xs:element ref="Branding" />
            </xs:choice>
        </xs:complexType>
    </xs:element>

</xs:schema>
```

**Help schema**

The following XML is the Help schema that allows HSPs and WISPs to provide customers with user assistance.

```
<?xml version="1.0" ?>
<xs:schema id="HelpFile" targetNamespace="http://www.microsoft.com/provisioning/Help"
xmlns="http://www.microsoft.com/provisioning/Help"
xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:msdata="urn:schemas-microsoft-
com:xml-msdata" attributeFormDefault="qualified" elementFormDefault="qualified">
    <xs:element name="Help">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="MHTML" type="xs:string" minOccurs="1" maxOccurs="1"/>
            </xs:sequence>
        </xs:complexType>
    </xs:element>

    <xs:element name="HelpFile" msdata:IsDataSet="true"
msdata:EnforceConstraints="False">
        <xs:complexType>
            <xs:choice maxOccurs="1">
                <xs:element ref="Help" />
            </xs:choice>
        </xs:complexType>
    </xs:element>
</xs:schema>
```

**EAP schemas**

There are eight EAP schemas, including the following:

Connection properties schemas:

- Base EAP Connection Properties schema
- EAP Connection Properties schema
- MS-CHAP v2 Connection Properties schema
- Microsoft PEAP Connection Properties schema

User properties schemas:

- Base EAP User Properties schema
- EAP User Properties schema
- MS-CHAP v2 User Properties schema
- Microsoft PEAP User Properties schema

Each of these schemas is described below.

**Connection properties schemas**

Following are the four connection properties schemas.

**Base EAP Connection Properties schema**

The following XML describes the Base EAP Connection Properties schema.

```
<?xml version="1.0" ?>
<xs:schema
      targetNamespace="http://www.microsoft.com/provisioning/BaseEapConnectionPropertie
sV1"
      elementFormDefault="qualified"
      xmlns="http://www.microsoft.com/provisioning/BaseEapConnectionPropertiesV1"
      xmlns:xs="http://www.w3.org/2001/XMLSchema"
      version="1.0"
   >
   <xs:complexType name="BaseEapTypeParameters" abstract="true" />
   <xs:complexType name="BaseEapParameters">
      <xs:sequence>
         <xs:element name="Type" type="xs:integer"/>
         <xs:any
               processContents="lax"
               minOccurs="0"
               maxOccurs="unbounded"
               namespace="##any"
         />
         <!-- One or more elements of the kind as follows should go in here. -->
         <!--
         <xs:element ref="EapType" maxOccurs="unbounded" />
         -->
      </xs:sequence>
   </xs:complexType>
   <xs:element name="EapType" type="BaseEapTypeParameters" abstract="true"/>
   <xs:element name="Eap" type="BaseEapParameters"/>
</xs:schema>
```

**EAP Connection Properties schema**

The following XML describes the EAP Connection Properties schema.

```
<?xml version="1.0" ?>
<xs:schema
      targetNamespace="http://www.microsoft.com/provisioning/EapConnectionPropertiesV1"
      elementFormDefault="qualified"
      xmlns="http://www.microsoft.com/provisioning/EapConnectionPropertiesV1"
      xmlns:xs="http://www.w3.org/2001/XMLSchema"
      xmlns:baseEap="http://www.microsoft.com/provisioning/BaseEapConnectionPropertiesV
1"
      version="1.0"
   >
   <xs:import
         namespace="http://www.microsoft.com/provisioning/BaseEapConnectionPropertiesV1
"
         schemaLocation="BaseEapConnectionPropertiesV1.xsd"
   />
   <xs:element name="Connections">
      <xs:complexType>
         <xs:sequence>
            <xs:element name="Connection" minOccurs="0" maxOccurs="unbounded">
               <xs:complexType>
                  <xs:sequence>
                     <xs:element name="Name" type="xs:string" />
                     <xs:element ref="baseEap:Eap" maxOccurs="unbounded" />
                  </xs:sequence>
               </xs:complexType>
            </xs:element>
         </xs:sequence>
      </xs:complexType>
   </xs:element>
</xs:schema>
```

**MS-CHAP v2 Connection Properties schema**

The following XML describes the MS-CHAP v2 Connection Properties schema.

```
<?xml version="1.0" ?>
<xs:schema
      targetNamespace="http://www.microsoft.com/provisioning/MsChapV2ConnectionProperti
esV1"
      elementFormDefault="qualified"
      xmlns="http://www.microsoft.com/provisioning/MsChapV2ConnectionPropertiesV1"
      xmlns:xs="http://www.w3.org/2001/XMLSchema"
      xmlns:baseEap="http://www.microsoft.com/provisioning/BaseEapConnectionPropertiesV
1"
      version="1.0"
   >
   <xs:import
      namespace="http://www.microsoft.com/provisioning/BaseEapConnectionPropertiesV1"
      schemaLocation="BaseEapConnectionPropertiesV1.xsd"
   />
   <xs:element name="EapType" substitutionGroup="baseEap:EapType">
      <xs:complexType>
         <xs:complexContent>
            <xs:extension base="baseEap:BaseEapTypeParameters"/>
         </xs:complexContent>
      </xs:complexType>
   </xs:element>
</xs:schema>
```

**Microsoft PEAP Connection Properties schema**

When you configure the PEAP Connection Properties XML file, change the value of **DisableUserPromptForServerValidation** from zero (0, or false) to one (1, or true). Also add the Secure Hash Algorithm (SHA-1) hash of your trusted root CA certificate to **TrustedRootCA**. In the following example, the SHA-1 hash of the trusted root CA certificate is the hexadecimal string: **a4 34 89 15 9a 52 0f 0d 93 d0 32 cc af 37 e7 fe 20 a8 b4 19**. When you use a SHA-1 hash in the XML file, you can keep the spaces between the pairs of numbers in the hash - you do not need to delete the spaces.

```
<msPeap:DisableUserPromptForServerValidation>1</msPeap:DisableUserPromptForServerValida
tion>
<msPeap:TrustedRootCA>a4 34 89 15 9a 52 0f 0d 93 d0 32 cc af 37 e7 fe 20 a8 b4
19</msPeap:TrustedRootCA>
```

You can specify more than one trusted root CA certificate by inserting additional lines that contain **TrustedRootCA** and the SHA-1 hashes of the additional certificates. The following example depicts an XML file that designates three trusted root CA certificates:

```
<msPeap:DisableUserPromptForServerValidation>1</msPeap:DisableUserPromptForServerValida
tion>
<msPeap:TrustedRootCA>SHA-1 hash of the CA 1 certificate</msPeap:TrustedRootCA>
<msPeap:TrustedRootCA>SHA-1 hash of the CA 2 certificate</msPeap:TrustedRootCA>
<msPeap:TrustedRootCA>SHA-1 hash of the CA 3 certificate</msPeap:TrustedRootCA>
```

**To obtain the SHA-1 hash of a trusted root CA certificate**

1. Open the Certificates snap-in for the Local Computer certificate store.

2. In the left pane, double-click **Certificates (Local Computer)**, and then double-click the **Trusted Root Certification Authorities** subfolder.

3. The Certificates folder is a subfolder of the Trusted Root Certification Authorities folder. Click the **Certificates** folder.

4. In the details pane, browse to the certificate for your trusted root CA. Double-click the certificate. The **Certificate** dialog box opens.

5. In the **Certificate** dialog box, click the **Details** tab.

6. In the list of fields, select **Thumbprint**.

7. In the lower pane, the hexadecimal string that is the SHA-1 hash of your certificate is displayed. Select the SHA-1 hash and press the Windows keyboard shortcut for the Copy command (CTRL+C) to copy the hash to the Windows clipboard.

For more information about using the Certificates snap-in, see "To manage certificates for a computer" in Windows XP Help

and Support Center or on the Web at http://go.microsoft.com/fwlink/?LinkId=41042.

For more information about Windows keyboard shortcuts, see "Windows keyboard shortcuts overview" in Windows XP Help and Support Center or on the Web at http://go.microsoft.com/fwlink/?LinkId=36303.

The following XML describes the Microsoft PEAP Connection Properties schema.

```
<?xml version="1.0"?>
<xs:schema
      targetNamespace="http://www.microsoft.com/provisioning/MsPeapConnectionProperties
V1"
      elementFormDefault="qualified"
      xmlns="http://www.microsoft.com/provisioning/MsPeapConnectionPropertiesV1"
      xmlns:xs="http://www.w3.org/2001/XMLSchema"
      xmlns:baseEap="http://www.microsoft.com/provisioning/BaseEapConnectionPropertiesV
1"
      version="1.0"
   >
   <xs:import
         namespace="http://www.microsoft.com/provisioning/BaseEapConnectionPropertiesV1
"
         schemaLocation="BaseEapConnectionPropertiesV1.xsd"
   />

   <xs:element name="EapType" substitutionGroup="baseEap:EapType">
      <xs:complexType>
         <xs:complexContent>
            <xs:extension base="baseEap:BaseEapTypeParameters">
               <xs:sequence>
                  <xs:element
                        name="ServerValidation"
                        type="ServerValidationParameters"
                        minOccurs="0"
                  />
                  <xs:element name="FastReconnect" type="xs:boolean" minOccurs="0"/>
                  <xs:element name="InnerEapOptional" type="xs:boolean" minOccurs="0"/>
                  <xs:element ref="baseEap:Eap" minOccurs="0" maxOccurs="unbounded" />
               </xs:sequence>
            </xs:extension>
         </xs:complexContent>
      </xs:complexType>
   </xs:element>

   <xs:complexType name="ServerValidationParameters">
      <xs:sequence>
         <xs:element
               name="DisableUserPromptForServerValidation"
               type="xs:boolean"
               minOccurs="0"
         />
         <!-- A set of server names delimited by semicolons -->
         <!-- each server name can be represented by regular -->
         <!-- expressions -->
         <xs:element name="ServerNames" type="xs:string" minOccurs="0" />
         <!-- The thumbprint of a trusted root CA is -->
         <!-- a hexadecimal string that contains -->
         <!-- the SHA-1 hash of the certificate. -->
         <xs:element
               name="TrustedRootCA"
               type="xs:hexBinary"
               minOccurs="0"
               maxOccurs="unbounded"
         />
      </xs:sequence>
   </xs:complexType>
</xs:schema>
```

**User properties schemas**

Following are the four user properties schemas.

**Base EAP User Properties schema**

The following XML describes the Base EAP User Properties schema.

```
<?xml version="1.0" ?>
<xs:schema
      targetNamespace="http://www.microsoft.com/provisioning/BaseEapUserPropertiesV1"
      elementFormDefault="qualified"
      xmlns="http://www.microsoft.com/provisioning/BaseEapUserPropertiesV1"
      xmlns:xs="http://www.w3.org/2001/XMLSchema"
      version="1.0"
   >
   <xs:complexType name="BaseEapTypeParameters" abstract="true" />
   <xs:complexType name="BaseEapParameters" >
      <xs:sequence>
         <xs:element name="Type" type="xs:integer"/>
         <xs:any
               minOccurs="0"
               processContents="lax"
               maxOccurs="unbounded"
               namespace="##any"
         />
         <!-- One or more elements of the kind as follows should go in here. -->
         <!--
         <xs:element ref="EapType" maxOccurs="unbounded" />
         -->
      </xs:sequence>
   </xs:complexType>
   <xs:element name="EapType" type="BaseEapTypeParameters" abstract="true"/>
   <xs:element name="Identity" type="xs:string" abstract="true" />
   <xs:element name="Eap" type="BaseEapParameters"/>
</xs:schema>
```

**EAP User Properties schema**

The following XML describes the EAP User Properties schema.

```
<?xml version="1.0" ?>
<xs:schema
      targetNamespace="http://www.microsoft.com/provisioning/EapUserPropertiesV1"
      elementFormDefault="qualified"
      xmlns="http://www.microsoft.com/provisioning/EapUserPropertiesV1"
      xmlns:xs="http://www.w3.org/2001/XMLSchema"
      xmlns:baseEap="http://www.microsoft.com/provisioning/BaseEapUserPropertiesV1"
      version="1.0"
   >
   <xs:import
         namespace="http://www.microsoft.com/provisioning/BaseEapUserPropertiesV1"
         schemaLocation="BaseEapUserPropertiesV1.xsd"
   />
   <xs:element name="User">
      <xs:complexType>
         <xs:sequence>
            <xs:element ref="baseEap:Eap" maxOccurs="unbounded"/>
         </xs:sequence>
      </xs:complexType>
   </xs:element>
</xs:schema>
```

**MS-CHAP v2 User Properties schema**

The following XML describes the MS-CHAP v2 User Properties schema.

```
<?xml version="1.0" ?>
<xs:schema
      targetNamespace="http://www.microsoft.com/provisioning/MsChapV2UserPropertiesV1"
      elementFormDefault="qualified"
      xmlns="http://www.microsoft.com/provisioning/MsChapV2UserPropertiesV1"
      xmlns:xs="http://www.w3.org/2001/XMLSchema"
      xmlns:baseEap="http://www.microsoft.com/provisioning/BaseEapUserPropertiesV1"
      version="1.0"
   >
   <xs:import
         namespace="http://www.microsoft.com/provisioning/BaseEapUserPropertiesV1"
         schemaLocation="BaseEapUserPropertiesV1.xsd"
   />
   <xs:element name="Username" substitutionGroup="baseEap:Identity"/>
   <xs:element name="EapType" substitutionGroup="baseEap:EapType">
```

```
      <xs:complexType>
         <xs:complexContent>
            <xs:extension base="baseEap:BaseEapTypeParameters">
               <xs:sequence>
                  <xs:element ref="Username" minOccurs="0" />
                  <xs:element name="Password" type="xs:string" minOccurs="0" />
                  <xs:element name="LogonDomain" type="xs:string" minOccurs="0"/>
               </xs:sequence>
            </xs:extension>
         </xs:complexContent>
      </xs:complexType>
   </xs:element>
</xs:schema>
```

**Microsoft PEAP User Properties schema**

The following XML describes the Microsoft PEAP User Properties schema.

```
<?xml version="1.0" ?>
<xs:schema
      targetNamespace="http://www.microsoft.com/provisioning/MsPeapUserPropertiesV1"
      elementFormDefault="qualified"
      xmlns="http://www.microsoft.com/provisioning/MsPeapUserPropertiesV1"
      xmlns:xs="http://www.w3.org/2001/XMLSchema"
      xmlns:baseEap="http://www.microsoft.com/provisioning/BaseEapUserPropertiesV1"
      version="1.0"
   >
   <xs:import
         namespace="http://www.microsoft.com/provisioning/BaseEapUserPropertiesV1"
         schemaLocation="BaseEapUserPropertiesV1.xsd"
   />
   <xs:element name="RoutingIdentity" substitutionGroup="baseEap:Identity"/>
   <xs:element name="EapType" substitutionGroup="baseEap:EapType">
      <xs:complexType>
         <xs:complexContent>
            <xs:extension base="baseEap:BaseEapTypeParameters">
               <xs:sequence>
                  <xs:element ref="RoutingIdentity" minOccurs="0"/>
                  <xs:element ref="baseEap:Eap" minOccurs="0" maxOccurs="unbounded" />
               </xs:sequence>
            </xs:extension>
         </xs:complexContent>
      </xs:complexType>
   </xs:element>
</xs:schema>
```

# How to Use IAS with a Third-Party User Accounts Database

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you deploy a third-party user accounts database for use with WPS technology, you must create and install two IAS extension DLLs:

- An authorization extension DLL that provides the URL PEAP-TLV if the customer or user account does not exist, is expired, or is disabled.
- An authentication extension DLL that retrieves the customer's or user's plaintext password from the third-party user accounts database and returns the password to IAS.

You can configure IAS for use with a third-party user accounts database by:

- Creating and installing an EAP authentication extension DLL. Because PEAP-MS-CHAP v2 is required for WPS technology, you must write an EAP authentication extension DLL to retrieve the password from the third-party user account database.
- Creating a new user account on your IAS server.
- Configuring a connection request policy on your IAS server that maps all user accounts to one account on the IAS server.
- Configuring a remote access policy in IAS that authorizes accounts mapped to the new account.

**Create an IAS authentication extension DLL**

Your IAS authentication extension DLL can use the following attributes:

- **ratProviderName**. This attribute indicates the remote RADIUS server group to which to forward the authentication request. The **ratProviderType** attribute is read-only. If **ratProviderType** is a RADIUS proxy, the extension DLL can change the value of **ratProviderName** to indicate the remote RADIUS server group to which the request should be forwarded.
- **ratClearTextPassword**. To support third-party user database use with PEAP-MS-CHAP v2, the IAS extension authentication DLL, this attribute retrieves the user password from the third-party user accounts database and sends this information back to IAS.

The IAS authentication extension DLL must also keep track of **ratUniqueId**. After the password is retrieved for the **ratUniqueId**, you do not need to retrieve the password again. If the account does not exist, is disabled, or is expired, the reason code **ratRejectReasonCode** must be sent back to IAS.

For more information, see "How to Create an IAS Extension DLL and a URL PEAP-TLV" in this paper.

**Install the IAS authentication extension DLL on the IAS server**

After you have created your IAS extension DLL, you must install the DLL on your IAS server and configure DLL registry keys according to your needs.

### To install your DLL

1. Open Command Prompt and change directories to the folder that contains your DLL.
2. Type the following: **regsvr32***DLL_name.dll*, where *DLL_name.dll* is the name of your DLL file.

**Create a user account on the IAS server**

You can create user accounts and group accounts in Active Directory to manage domain users. When you are not using Active Directory as your user accounts database, you can create user accounts and group accounts on a local computer to manage users specific to that computer.

To deploy IAS with a third-party user accounts database, you must create one user account on your IAS server to which you can map all user accounts in the third-party database. Your extension DLL performs authentication against the third-party database, while IAS performs authorization with the user account you create on the IAS server.

**To create a user account on the IAS server**

1. Open **Computer Management**. To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

2. In **Computer Management**, under **System Tools**, click **Local Users and Groups**. In the details pane, double-click **Users**.

3. On the **Action** menu, click **New User**. The **New User** dialog box opens.

4. In **User Name**, type a name for the account. In **Password** and **Confirm Password**, type a strong password. Clear the **User must change password at next logon** check box, and then select the **User cannot change password** and **Password never expires** check boxes.

5. Click **Create**, and then click **Close**.

| **Note** |
| --- |
| By default, user accounts created on the local computer in Windows Server 2003 have dial-in properties set to **Control access through Remote Access Policy**. This is the correct setting for the user account you have created and the setting should not be changed. |

For more information, see "Strong passwords" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=34427.

#### Configure IAS connection request policy

**To configure IAS connection request policy**

1. Open **Internet Authentication Service**.

2. In the console tree, double-click **Connection Request Processing**, and then click **Connection Request Policies**.

3. In the details pane, double-click the policy that you want to configure. For example, double-click the default policy, named **Use Windows authentication for all users**.

4. In the **Properties** dialog box, click **Edit Profile**.

5. On the **Authentication** tab, click **Authenticate requests on this server**.

6. On the **Attribute** tab, confirm that the User-Name attribute is selected in **Attribute**. If it is not selected, click **Attribute**, select **User-Name**, and then click **Add**.

7. The **Attribute Manipulation Rule** dialog box opens. In **Find**, type pattern matching syntax that matches all values for User-Name that are passed to IAS by your access servers. For example, if values for the User-Name attribute match the syntax *user@example.com*, type (.*)@.*.

8. In **Replace with**, type the name of the local computer and the name of the user account you created on the IAS server in the following syntax: *computer-name\user*.

9. Click **OK** twice.

| **Note** |
| --- |

If you copy the IAS connection request policy described above to another IAS server, change the attribute manipulation rule by changing the computer name and user name in the rule to the computer name of the new server and the user name of an account on the new server. If you do not change the computer name and user name specified in the rule, attribute manipulation on the new server will not work.

For more information, see "Pattern Matching Syntax" in Help and Support Center for Windows Server 2003 or on the Web at http://go.microsoft.com/fwlink/?LinkId=41049.

#### Configure IAS remote access policy

There are two remote access policies configured for WPS technology. The Guest access policy provides network parameters and rules for users connecting as guest. The Valid users access policy provides network parameters and rules for users who have valid WISP accounts.

**To configure the Guest access policy**

1. Open the Internet Authentication Service snap-in and, if necessary, double-click **Internet Authentication Service**.
2. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**.
3. Use the **New Remote Access Policy Wizard** to create a policy. For the WISP guest access policy, you can choose the following:
    a. For **How do you want to set up this policy?** verify that **Use the wizard to set up a typical policy for a common scenario** is selected.
    b. For **Policy name**, type **Guest access** (or type another name for your policy that you prefer).
    c. For **Select the method of access for which you want to create a policy**, click **Wireless**.
    d. For **Grant access based on the following**, click **User**.
    e. In **Select the EAP type for this policy**, select **Protected EAP (PEAP)**, and then click **Configure**.
    f. In **Certificate issued**, select the certificate that you want the IAS server to use to verify its identity to client computers. Also select the **Enable Fast Reconnect** check box.

After you have completed creating the policy and have closed the wizard by clicking **Finish**, you need to perform additional policy configuration. In the IAS console, click **Remote Access Policies**, and then double-click the policy you just created.

**To complete configuration of the Guest access policy**

1. In the policy **Properties** dialog box, for **Policy conditions**, click **Add**.
2. In **Attribute Types**, click **Day-And-Time-Restrictions**, and then click **Add**. In **Time of day restraints**, select **Permitted**, configure the days and times that access is permitted, and then click **OK**.
3. In the policy **Properties** dialog box, click **Grant remote access permission**.
4. Click **Edit Profile**. On the **Authentication** tab, in **Unauthenticated access**, click **Allow clients to connect without negotiating an authentication method**.

#### Configure the Valid users policy

**To configure the Valid users remote access policy**

1. Open **Internet Authentication Service**.
2. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**.

3. Use the **New Remote Access Policy Wizard** to create a policy. For the remote access policy, you can choose the following:

    a. For **How do you want to set up this policy?** verify that **Use the wizard to set up a typical policy for a common scenario** is selected.

    b. For **Policy name**, type **Valid Users** (or type another name for your policy that you prefer).

    c. For **Select the method of access for which you want to create a policy**, click **Wireless**.

    d. For **Grant access based on the following**, click **User**.

    e. In **Select the EAP type for this policy**, select **Protected EAP (PEAP)**, and then click **Configure**.

    f. In **Certificate issued**, select the certificate that you want the IAS server to use to verify its identity to client computers. Also check the **Enable Fast Reconnect** check box.

After you have completed creating the policy and have closed the wizard by clicking **Finish**, you need to perform additional policy configuration.

**To complete configuration of the Valid users remote access policy**

1. In the IAS console, click **Remote Access Policies**, and then double-click the policy you just created.

2. In **Attribute Types**, click **Day-And-Time-Restrictions**, and then click **Add**. In **Time of day restraints**, select **Permitted**, configure the days and times that access is permitted, and then click **OK**.

3. In the policy **Properties** dialog box, click **Grant remote access permission**.

4. Click **Edit Profile**. On the **Authentication** tab, clear all check boxes except Strongest encryption (MPEE 128 bit).

| **Note** |
| --- |
| If you are isolating client computers by using VLANs, you must also add VLAN attributes to the Valid Users access policy. If you are isolating client computers by using IP filters, you must add IP filters to the access policy. |

# PEAP-MS-CHAP v2

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Protected Extensible Authentication Protocol (PEAP) is a member of the family of Extensible Authentication Protocol (EAP) protocols. PEAP uses Transport Layer Security (TLS) to create an encrypted channel between an authenticating PEAP client, such as a wireless computer, and a PEAP authenticator, such as an Internet Authentication Service (IAS) or Remote Authentication Dial-In User Service (RADIUS) server. PEAP does not specify an authentication method, but provides additional security for other EAP authentication protocols, such as EAP-MS-CHAP v2, that can operate through the TLS encrypted channel provided by PEAP. PEAP is used as an authentication method for 802.1X wireless client computers, but is not supported for virtual private network (VPN) or other remote access clients.

**PEAP authentication process**
There are two stages in the PEAP authentication process between PEAP client and authenticator. The first stage sets up a secure channel between the PEAP client and the authenticating server. The second stage provides EAP authentication between the EAP client and authenticator.

**PEAP stage one: TLS encrypted channel**
The wireless client associates with a wireless access point. An IEEE 802.11-based association provides an Open System or Shared Key authentication before a secure association is created between the client and access point. After the IEEE 802.11-based association is successfully established between the client and access point, the TLS session is negotiated with the access point. After authentication is successfully completed between the wireless client and the server (for example, an IAS server), the TLS session is negotiated between them. The key that is derived during this negotiation is used to encrypt all subsequent communication.

**PEAP stage two: EAP-authenticated communication**
Complete EAP communication, including EAP negotiation, occurs inside the TLS channel created by PEAP during the first stage of the PEAP authentication process. The IAS server authenticates the user or the client computer with the method that is determined by the EAP type and selected for use within PEAP. For deployments of WPS technology, EAP-MS-CHAP v2 is the authentication type used within PEAP. The access point only forwards messages between wireless client and RADIUS server—the access point (or a person monitoring it) cannot decrypt these messages because it is not the TLS end point.

**Packet sequence for a successful authentication attempt with valid credentials**
After PEAP stage one occurs and the TLS channel is created between the IAS server and the 802.1X wireless client, for a successful authentication attempt where the user has supplied valid password-based credentials using WPS technology with PEAP-MS-CHAP v2, the RADIUS message sequence is:

1. The IAS server sends an identity request message to the client: EAP-Request/Identity.

2. The client responds with an identity response message: EAP-Response/Identity.

3. The IAS server sends an MS-CHAP v2 challenge message: EAP-Request/EAP-Type=EAP MS-CHAP-V2 (Challenge).

4. The client responds with an MS-CHAP v2 challenge and response: EAP-Response/EAP-Type=EAP-MS-CHAP-V2 (Response).

5. The IAS server sends back an MS-CHAP v2 success packet when the server has successfully authenticated the client: EAP-Request/EAP-Type=EAP-MS-CHAP-V2 (Success).

6. The client responds with an MS-CHAP v2 success packet when the client has successfully authenticated the server: EAP-Response/EAP-Type=EAP-MS-CHAP-V2 (Success).

7. The IAS server sends an EAP-TLV indicating successful authentication.

8. The client responds with an EAP-TLV status success message.

9. The server completes authentication and sends an EAP-Success message using plaintext. If VLANs are deployed for client isolation, the VLAN attributes are included in this message.

**Packet sequence for a successful guest authentication**
After PEAP stage one occurs and the TLS channel is created between the IAS server and the 802.1X wireless client, for the case of an authentication failure being converted to a successful guest authentication using WPS technology with PEAP-MS-CHAP v2, the RADIUS message sequence is:

1. The IAS server sends an identity request message to the client: EAP-Request/Identity.

2. The client responds with an identity response message: EAP-Response/Identity.

3. The IAS server sends an MS-CHAP v2 challenge message: EAP-Request/EAP-Type=EAP MS-CHAP-V2 (Challenge).

4. The client responds with an MS-CHAP v2 challenge and response: EAP-Response/EAP-Type=EAP-MS-CHAP-V2 (Response). The authentication fails and the WPS extension DLL on the server determines that the user authentication should succeed with limited access.

5. The IAS server sends back an MS-CHAP v2 success packet when the server has successfully authenticated the client: EAP-Request/EAP-Type=EAP-MS-CHAP-V2 (Success).

6. The client responds with an MS-CHAP v2 success packet when the client has successfully authenticated the server: EAP-Response/EAP-Type=EAP-MS-CHAP-V2 (Success).

7. The IAS server sends an EAP-TLV indicating successful authentication, as well as other TLVs, including a URL PEAP-TLV that provides the client with the URL of the provisioning server (or a URL to change password).

8. The client responds with an EAP-TLV status success message.

9. The server completes authentication and sends an EAP-Success message using plaintext. If VLANs are deployed for client isolation, the VLAN attributes are included in this message.

**Packet sequence for other WPS cases (disabled, expired, and unknown accounts)**
After PEAP stage one occurs and the TLS channel is created between the IAS server and the 802.1X wireless client, for other WPS cases (disabled, expired and unknown accounts) with PEAP-MS-CHAP v2, the RADIUS message sequence is:

1. The IAS server sends an identity request message to the client: EAP-Request/Identity.

2. The client responds with an identity response message: EAP-Response/Identity.

3. The IAS server discovers that the account is disabled, expired, or unknown. The WPS extension DLL on the server decides that the user authentication should succeed with limited access.

4. The IAS server sends an EAP-TLV indicating successful authentication, as well as other TLVs, including a URL PEAP-TLV that provides the client with the URL of the provisioning server.

5. The client responds with an EAP-TLV status success message.

6. The server completes authentication and sends an EAP-Success message using plaintext. If VLANs are deployed for client isolation, the VLAN attributes are included in this message.

For more information about PEAP, see "Protected EAP Protocol (PEAP)" at http://go.microsoft.com/fwlink/?LinkId=41301.

For more information about EAP-MS-CHAP v2, see "Microsoft EAP CHAP Extensions" at http://go.microsoft.com/fwlink/?LinkId=41306.

# How to Create an IAS Extension DLL and a URL PEAP-TLV

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are two types of PEAP-TLVs used by IAS and WPS:

- **Status-Result PEAP-TLV with Type=3.** The Status-Result PEAP-TLV is sent by the IAS server and provides client computers with an authentication response (either accept or reject).

- **URL PEAP-TLV with Type=8.** The URL PEAP-TLV is sent by an IAS extension DLL that you create and install on your IAS servers. The URL PEAP-TLV provides client computers with the location of the provisioning server so that the client can download XML files from the server.

IAS always sends a Status-Result PEAP-TLV containing an authentication response. Depending upon circumstances outlined below, your IAS extension DLL might include a URL PEAP-TLV inside the same message as the Status-Result PEAP-TLV.

The process for providing your IAS servers with an extension DLL capable of processing messages is as follows:

- Create an IAS extension DLL that can identify and act upon access requests that should be provisioned and allowed, while passing through all other access requests without any modifications.

- Add the DLL to the Windows registry, so that IAS calls the DLL during the processing of each message. After your DLL identifies an access request upon which it needs to act, it changes the access request into a successful attempt (if the access request was previously marked for rejection), and the DLL will inject the URL PEAP-TLV into the response. In addition, the DLL might also add RADIUS Tunnel attributes or otherwise request that a network restriction be enforced for the current client. The PEAP-TLVs are sent inside the last PEAP Access-Challenge message; the overall authentication result, along with any RADIUS Tunnel attributes, is sent in the final Access-Accept message. The client uses the specified URL PEAP-TLV, which contains an HTTPS URL and an action parameter, to contact a provisioning server and obtain XML files.

**URL PEAP-TLV values**

The value of the URL PEAP-TLV, which is sent in the PEAP channel and is encrypted and integrity-checked, provides the client with the following information:

- The location of the provisioning information, in the form of an HTTPS URL, which is needed by the client to access the provisioning server. An example URL is: http://www.example.com/provisioning/master.xml.

- The action parameter, which represents the action the WISP requires of the client to access the WISP. This action and the pound sign character (#) are appended to the HTTPS URL sent in the URL PEAP-TLV. There are four possible actions that can be configured in a URL PEAP-TLV:

    - **Sign up.** This value is passed by IAS to the client when the customer has authenticated as guest or attempts authentication as an unknown user and does not have a valid user account. An example value for the URL PEAP-TLV is: http://www.example.com/provisioning/master.xml#signup, where **#signup** is the parameter for this action.

    - **Renewal.** This value is passed by IAS to the client when the customer's user account is expired and needs renewal before network access can be granted. An example value for the URL PEAP-TLV is: http://www.example.com/provisioning/master.xml#renewal, where **#renewal** is the parameter for this action.

    - **Password change.** This value is passed by IAS to the client when the customer is required to change the account password. An example value for the URL PEAP-TLV is: http://www.example.com/provisioning/master.xml#passwordchange, where **#passwordchange** is the parameter for this action.

    - **Force update.** This value is passed by IAS to the client when the WISP requires the Windows XP–based client to

download an updated XML master file. This method of updating the XML master file on the client should be used only to correct errors; otherwise, the TTL expiry time in the XML master file is used to provide background updates. An example value for the URL PEAP-TLV is: http://www.example.com/provisioning/master.xml#forceupdate, where **#forceupdate** is the parameter for this action.

**WPS technology response to URL PEAP-TLVs on client computers**

WPS technology on client computers running Windows XP with SP2 evaluates which URL PEAP-TLV values sent by the IAS server are valid based upon whether the client has authenticated as guest or has authenticated with credentials for a valid user account.

When the client has authenticated as guest, WPS technology on the client accepts and utilizes a URL PEAP-TLV that includes the sign-up action parameter (#signup). This allows customers who do not yet have an account to do the following:

- Download the XML files that contain the sign-up wizard
- Run the sign-up wizard
- Create an account

This is the only circumstance under which WPS technology on the client will download and run the sign-up wizard.

When the client has authenticated with credentials for a valid user account, WPS technology on the client ignores a URL PEAP-TLV that contains the sign-up action parameter (#signup). WPS technology on the client will, however, respond to URL PEAP-TLVs that contain the renewal action parameter (#renewal), the password change action parameter (#passwordchange), and the force update action parameter (#forceupdate) by downloading the appropriate XML files and running the wizard specified by the action parameter.

**IAS extension DLL types and functionality**

IAS supports both authentication extension DLLs and authorization extension DLLs. Authentication DLLs validate credentials against a user account database; authorization DLLs are processed after authentication DLLs and are used to modify the result sent back to the client computer. IAS extension DLLs for WPS technology are authorization DLLs.

When IAS starts, it automatically loads all registered extension DLLs. Extension DLLs can then inspect and modify every RADIUS message received and sent by IAS.

When you create your IAS extension DLL, make sure that it can perform the following functions within the IAS processing cycle:

- Recognize particular types of authentication attempts (both accepted and rejected attempts).
- Convert these types of authentication attempts into accepted attempts based on the WPS technology business rules that you have defined in IAS remote access policies and in your sign-up wizard, renewal wizard, and other XML files.
- Add provisioning information, such as a URL PEAP-TLV containing the URL of the provisioning server and an action parameter, to the responses sent to client computers connected to your Wi-Fi hotspots.
- Provide information that allows IAS to isolate clients to a network where they can access resources, such as the provisioning server and the IAS server, used when signing up for accounts.

As with all extension DLLs, a WPS technology extension DLL must first identify and ignore all RADIUS messages that are not relevant to its function. When the DLL recognizes a message that it might need to modify, it must inspect the following:

- The type of message being processed.
- The current authentication result.
- The user name being authenticated.

Then, based on the business logic appropriate to your deployment of WPS technology, the extension DLL should decide

whether the response should:

- Be converted into a successful authentication attempt.
- Have provisioning information added to the message, and/or:
- Indicate to a layer three network device (for example, a wireless access point) that the client computer should be allowed restricted access to the network.

Finally, the WPS technology extension DLL should perform any changes that it determines are appropriate.

| **Important** |
| --- |
| The extension DLL must not return an error to IAS. In addition, the extension DLL should always allow RADIUS messages to continue to be processed; it should not unilaterally decide to drop a packet. |

As previously described, provisioning information consists of a RADIUS ratEAPTLV attribute (type 273), containing a URL PEAP-TLV whose value contains a URL string and an action parameter.

**APIs**

A WPS Extension DLL must implement the following API functions, which are called by IAS:

- **RadiusExtensionInit()** This API allows the DLL to perform one-time initialization actions.
- **RadiusExtensionTerm()** This API allows the DLL to perform one-time clean-up and termination actions.
- **RadiusExtensionProcess2()** This API is called for every RADIUS message received by IAS.

**IAS extension DLL actions**

The following structures, attribute names, and attribute values are defined in the SDK header file authif.h. For more information, see MSDN.

- **pECB** The **RADIUS_EXTENSION_CONTROL_BLOCK** pointer parameter passed into RadiusExtensionProcess2().
- **pRequest** The **RADIUS_ATTRIBUTE_ARRAY** pointer containing the RADIUS attributes that make up the RADIUS request packet from the client host.
- **pResponse(X)** The **RADIUS_ATTRIBUTE_ARRAY** pointer containing the RADIUS attributes that make up the RADIUS response packet of the specified type (e.g., pResponse(REJECT)).

Your WPS technology IAS extension DLL must perform the following actions:

1. **Message Filter**. The WPS extension DLL should ignore all messages, except those that match the following criteria:

    a. The message must be at the authorization stage of IAS processing.

    **pECB->repPoint == repAuthorization**

    b. The inbound message must be an ACCESS_REQUEST.

    **pECB->rcRequestType == rcAccessRequest**

    c. The outbound message must be an ACCESS_CHALLENGE.

    **pECB->rcResponseType == rcAccessChallenge**

    d. The response message must contain a RADIUS ratEAPTLV attribute (type 273).

To perform this check, the DLL must scan all RADIUS attributes in the response to discover whether the attribute ratEAPTLV (type 273) is present.

2. **Message Inspection**. The WPS extension DLL should obtain the following information about the authentication attempt:

   a. RADIUS connection request policy (CRP) used to process this request (RADIUS attribute ratPolicyName = 270 = remote access policy name; ratCRPPolicyName = 275 = connection request policy name).

   b. Reject Reason Code (RADIUS attribute ratRejectReasonCode (type 274))

   c. Username (checking both ratUserName (type 1) and ratFQUserName (type 269))

   d. Overall authentication result, from the RADIUS ratEAPTLV attribute (type 273) containing the Status-Result TLV (type=3) (possible values: Success(1) or Failure(2))

3. **Message Analysis**. The following business logic is recommended, but can be adapted as appropriate:

   a. You can design the DLL to choose to perform completely different business logic based on the RADIUS connection request policy name used to process the request. This allows you to configure the IAS server with remote access policies that identify users based on any supported criteria, then have the extension DLL impose greater or fewer restrictions for different types of accounts based on the remote access policy that was used.

   b. Within a request the Reject Reason Code should take precedence over the other values. If this RADIUS attribute is present, the message should be processed according to the reason code specified. This can have one of the following values:

      **i.** rrrcAccountUnknown(1)

      The authentication attempt is using a user name that does not correspond to any known account. (Suggested action parameter: "signup")

      **ii.** rrrcAccountDisabled(2)

      The authentication attempt is using a user name that corresponds to an account that has been disabled by an administrator. (Suggested action parameter: "signup")

      **iii.** rrrcAccountExpired(3)

      The authentication attempt is using a user name that corresponds to an account that has been expired, either by exceeding its natural expiration lifetime or by administrative action. (Suggested action parameter: "renew")

      **iv.** rrrcAuthenticationFailure(4)

      The authentication attempt is using a user name and password pair that are incorrect for the corresponding account. (Suggested action parameter: "signup")

      You can design the DLL to choose how to respond to each of these types of failures independently, based on the business logic appropriate to your needs.

      The DLL may choose to respond to all types of failures equally, to respond to some types of failures differently than others, or to ignore one or more types of failures. The suggested behavior for the extension DLL is that for each type, the failure is converted to a success and provisioning information is added with the action parameters listed above.

   c. The user name should have the next level of precedence. This is where the DLL can determine whether the current authentication attempt is using a guest account or a valid user account.

      **i.** Guest account attempts should be handled by both the DLL and an IAS connection request policy - the DLL

should add the provisioning information, while the connection request policy should add network restriction information.

**ii.** Normal user authentication attempts must be recognized as such, but the DLL must not make a decision until it has considered the overall authentication result.

d. For Normal user authentication, the overall authentication result must be checked.

**i.** If the attempt was successful, the response may be modified as described above; typically, however, it should not be modified.

Adding provisioning information would cause users who are allowed onto the network to also receive a balloon asking for an account sign-up, renewal, or password change. In situations when these behaviors are desired, then provisioning information may be added.

Adding restricted network information would mean that users who have proven that they have valid credentials would only have access to a limited part of the network. In situations when these behaviors are desired, then restricted network information may be added.

**ii.** If the attempt was unsuccessful, for a reason not already specified by the Reject Reason Code, then the attempt must not be modified. Attempts that are rejected without a corresponding Reject Reason Code are typically rejected for true error conditions; this can include requests attempting to use an EAP type not configured on the IAS server, corrupted messages, and so forth - all true errors where an Accept response should not be sent to the end user.

4. **Message Modification**. The following changes can be performed on the response message:

a. Convert to success.

Find the RADIUS ratEAPTLV attribute (type 273) that contains the Status-Result TLV (type 3), then directly set its data value to Success(1) (modifying the value in place). Because this data value contains a short int in network byte order, you can use the following line of code for the assignment stage:

**pEapTlvIn->Value[1] = MS_PEAP_AVP_VALUE_SUCCESS;**

b. Add provisioning information.

Allocate enough memory for an EAP-TLV structure, set the type to URL PEAP-TLV (type 8), set the length as appropriate, and copy the URL string into the TLV's value buffer. Then, add a RADIUS ratEAPTLV attribute (type 273) whose data value points to the EAP-TLV data structure.

c. Add restricted-network attributes.

This involves adding RADIUS attributes to the response, to indicate to the Wi-Fi hotspot's network hardware that network restrictions should be associated with the given end-user's client computer. These restrictions can be imposed by the 802.11 wireless access point or by a network switch or router between the access point and the rest of the network.

Different network hardware provides different methods for imposing network restrictions; the attributes to be added will depend on the network hardware in use.

For example, if VLANs are deployed for client isolation, then the following RADIUS attributes could be added to the response, to direct the network hardware to place the client on VLAN 0:

ratFramedProtocol (type 7) = PPP(1)

ratTunnelType (type 64) = TUNNEL_TYPE_VLAN(13)

ratMediumType (type 65) = TUNNEL_MEDIUM_TYPE_802(6)

ratTunnelPrivateGroupID (type 81) = "0" (0x30, 0x00)

ratTunnelTag (type 4170) = [obtain value from hardware-specific documentation]

For more information about creating an IAS extension DLL, see the following topics:

- RADIUS_EXTENSION_CONTROL_BLOCK (Internet Authentication Service Extensions: Platform SDK) at http://go.microsoft.com/fwlink/?LinkId=20041. The RADIUS_EXTENSION_CONTROL_BLOCK structure provides information about the current RADIUS request. It also provides functions for obtaining the attributes associated with the request, and for setting the disposition of the request.

- GetRequest (Internet Authentication Service Extensions: Platform SDK) at http://go.microsoft.com/fwlink/?LinkId=20042. The GetRequest function returns the attributes received in the RADIUS request process and any internal attributes describing the request state.

- RadiusExtensionTerm (Internet Authentication Service Extensions: Platform SDK) at http://go.microsoft.com/fwlink/?LinkId=20043. The RadiusExtensionTerm function is called by IAS prior to unloading the Extension DLL. Use RadiusExtensionTerm to perform any clean-up operations for the Extension DLL.

- Internet Authentication Service Extensions at http://msdn.microsoft.com/library/default.asp?url=/library/en-us/ias/ias/ias_start_page.asp

To order a Platform Software Development Kit (SDK) CD, see Microsoft Platform SDK at http://go.microsoft.com/fwlink/?LinkId=20044.

**PEAP-TLV Packets**
The Protected Extensible Authentication Protocol (PEAP), described in RFC 2284, provides a standard mechanism for support of multiple authentication types. Using PEAP-TLV, it is possible for various types of data to be passed directly between the IAS server and the PEAP client, and to provide functionality that is not included in RFC 2284 without defining a multiplicity of new PEAP types. The following table describes PEAP-TLV packet types:

| Packet Types | Packet Description |
|---|---|
| PEAP-TLV URI Packet | Contains a Type-Length-Value object (TLV). |
| PEAP-TLV Result Packet | Provides support for acknowledged success and failure messages |

**Packet details**
The following section provides detailed descriptions of the PEAP-TLV URI Packet and the PEAP-TLV Result Packet.

**PEAP-TLV URI Packet**
The PEAP-TLV URI packet contains a Type-Length-Value (TLV) object.

**Fields**
**MandatoryRequirement**

Data type: **BINARY**

Mandatory TLV packet, set to 0.

| Value | Meaning |
|---|---|
| 0 | Non-mandatory TLV |
| 1 | Mandatory TLV |

**TLVReserved**
Length: 1

Data type: **BINARY**

Reserved. Value is zero.

**TLVType**
Length: 14

Data type: **BINARY**

TLV attribute type.

| Value | Meaning |
|---|---|
| 0 | Reserved |
| 1 | Reserved |
| 2 | Reserved |
| 3 | Acknowledged result |

**TLVValueLength**
Data type: **UCHAR**

Length of **TLVValue** field.

**TLVValue**
Data type: **UCHAR**

URI to a master document.

| Value | Meaning* |
|---|---|
| https://www.example.com/provisioning/master.xml#signup | **1** |
| https://www.example.com/provisioning/master.xml#renewal | **2** |
| https://www.example.com/provisioning/master.xml#passwordchange | **3** |
| https://www.example.com/provisioning/master.xml#forceupdate | **4** |

*__Meaning__ definitions

1. Signup should be returned by the ISP in the TLV when the user name in the authentication is guest or an unknown user name.

2. Renewal should be returned by the ISP in the TLV when the user name's account needs renewal by the user

3. Passwordchange should be returned by the ISP in the TLV when the user should change the password for the account.

4. Forceupdate should be returned by the ISP in the TLV when the ISP needs the client to update the provisioning files as soon as possible. This should not normally be used to update the files; the TTL expiry time in the master file should be used to allow the files to be updated in the background instead. Forceupdate should be used only to correct an error in the XML files.

**PEAP-TLV Result Packet**
The PEAP-TLV Result Packet provides support for acknowledged success and failure messages.

**Fields**
**MandatoryRequirement**

Length: 1

Data type: **BINARY**

Mandatory TLV packet, set to 1.

**TLVReserved**

Length: 1

Data type: **BINARY**

Reserved. Value is zero.

**TLVPacketType**
Length: 14

Data type: **BINARY**

TLV packet type. Value is 3.

**TLVStatusLength**
Data type: **USHORT**

Length of the **TLVStatus** field. Value is 2.

**TLVStatus**
Data type: **USHORT**

TLV status.

| Value | Meaning |
|---|---|
| 1 | Success |
| 2 | Failure |

**Backward compatibility**
PEAP-TLV is a "special case" type, more similar to the Identity and Notification types than to the authentication types such as MD5-Challenge (RFC 2284). PEAP-TLV differs from the Identity and Notification types in that a peer may respond to an EAP-TLV request with a Nak Response. This is allowed for backward compatibility with implementations that do not support the PEAP-TLV type.

# Beta Documentation Note

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following two scenarios — a WISP using IP filters for client isolation and an HSP using IP filters for client isolation — are currently in development and have not been implemented or tested. They are provided as general depictions of possible implementations of WPS technology at a later date.

# WPS Technology for a WISP with IP Filters

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization is an enterprise or a WISP, you can deploy WPS technology by using IP filters for client computer isolation. In this circumstance, client computers are provided access to network resources, such as the provisioning server and IAS servers, and IP filters block access to the Internet until the customer establishes a paid account. After the account is established, the IP filters are lifted, and traffic is allowed between the client computer and the Internet.

In the sections that follow, the components of a WISP network using IP filters are described, and how the components work together during new customer sign-up are detailed.

**Components of WPS technology for a WISP using IP filters**
When an ISP deploys Wi-Fi hotspots as a WISP, WPS technology can be configured with IP filters for client computer isolation during the account sign-up process.

The following illustration depicts the components of a WISP network using IP filters for client computer isolation.




**Components of a WISP network using IP filters**

This configuration consists of the following components.

**Wireless client**
A computer running Windows XP Home Edition with SP2, Windows XP Professional with SP2, or Windows XP Tablet PC Edition with SP2. The wireless client connects to the WISP network, which is in a public location and is accessible to customers with portable computers and wireless network adapters.

**Wireless access point (RADIUS client)**
The wireless access point is configured as a RADIUS client to the IAS server deployed on the WISP LAN. The wireless access points used for WPS technology must meet the following requirements:

- Support for the IEEE standard 802.1X authentication.
- Support for Wi-Fi Protected Access (WPA) is preferred.
- The ability to implement client isolation by applying IP filters to the connection with the client.
- Support for RADIUS authentication and RADIUS accounting, including:
    - Support for the Class attribute as defined in RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)," to allow session correlation for RADIUS authentication and accounting records. For session correlation, when you configure RADIUS accounting at your IAS server or proxy, you must log all accounting data that allow applications (such as billing applications) to query the database, correlate related fields, and return a cohesive view of each session in the query results. At a minimum, to provide session correlation, you must log the following IAS accounting data: NAS-IP-Address; NAS-Identifier (you need both NAS-IP-Address and NAS-Identifier because the access server can send either attribute); Class; Acct-Session-Id; Acct-Multi-Session-Id;

Packet-Type; Acct-Status-Type; Acct-Interim-Interval; NAS-Port; and Event-Timestamp.

- Support for accounting interim requests, which are sent periodically by some access servers during a user session, that can be logged. This type of request can be used when the Acct-Interim-Interval RADIUS attribute is configured to support periodic requests in the remote access profile on the IAS server. The access server, in this case a wireless access point, must support the use of accounting interim requests if you want the interim requests to be logged on the IAS server.

- Support for IP address range filtering.

- Support for dynamic retransmit timeout (RTO) estimation or exponential backoff to handle congestion and delays in a wide area network (WAN) environment.

In addition, there are some filtering features that the access points should support to provide enhanced security for the network. These filtering options include:

- **DHCP filtering.** The access point must filter on IP ports to prevent the transmission of DHCP broadcast messages in the instance that the client is a DHCP server. The access point must block the client from sending IP packets from port 68 to the network.

- **DNS filtering.** The access point must filter on IP ports to prevent a client from performing as a DNS server. The access point must block the client from sending IP packets from port 53 to the network.

**IP filters**
IP filters that isolate client computers connecting as guest are configured in the remote access policy on the WISP IAS server, and are applied to the client connection by the wireless access point. The IP filters must grant the client computer access to the provisioning server and block access to all other locations.

**Router**
The router must provide multiple ports for connection to access points, the WISP LAN, and the Internet.

**Provisioning server**
The WISP provisioning server is configured with the following components.

**HTTPS Web server**
The Internet Information Services (IIS) or third-party Web server must be deployed with HTTPS.

**Web application**
The WISP Web server is configured with an account processing Web application that processes data provided during customer sign-up or account renewal. When a customer uses the sign-up wizard on a client computer to create and pay for a WISP account, the customer enters data, such as name, address, and credit card information, that is converted to an XML document on the client. Windows XP sends this XML document to the WISP provisioning server.

The account processing Web application on the provisioning server must be capable of accepting and processing the XML documents containing the user data. For example, the account processing Web application must dynamically create an account in the Active Directory user accounts database, and must contain a credit card verification component to process customers' payment information. The account processing Web application must permit new customers to sign up and to permit existing customers to renew their subscriptions for service.

**XML master file and subfiles**
The WISP provisioning server stores the XML master file and subfiles that provide the client with all configuration information needed to access the network, create an account, pay for the account, and ultimately access the Internet. For more information about the XML master file and subfiles, see "XML Schemas" in this paper.

**Server certificate**
For server authentication to client computers, the WISP provisioning server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and be issued by a public certification authority (CA), such as Verisign or Thawte, that is trusted by client computers. The Trusted Root Certification Authorities certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust.

**Domain controller and IAS server**
The WISP domain controller and IAS server is configured with the following components.

**Active Directory**
The user accounts database on the domain controller must be an Active Directory user accounts database or a database that uses Lightweight Directory Access Protocol (LDAP) and supports dynamic creation of user accounts. When a customer signs up for an account using the sign-up wizard, the account processing Web application on the provisioning server creates a new account in the user accounts database, and adds the user to a group that has clearly defined access privileges that match the type of account the customer purchased when signing up.

If you use an accounts database other than Active Directory, IAS authentication and authorization extension DLLs must be written and installed for this process to function correctly.

For information about configuring Active Directory replication, see "Active Directory Replication" in this paper.

**Internet Authentication Service (IAS)**
IAS is the Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy, and is used to authenticate and authorize users connecting to your network. IAS is configured with remote access policies that allow guest authentication for non-domain member computers and users. It is also configured to provide IP filters to RADIUS clients (access points) that apply the filters to client connections during guest authentication. IAS also provides unrestricted access to users with valid accounts.

**Extension DLL and URL PEAP-TLV**
An IAS authentication extension DLL defining a URL PEAP-TLV provides IAS with the ability to send the location of the provisioning server to client computers.

PEAP-Type-Length-Value (PEAP-TLV) is an Extensible Authentication Protocol (EAP) authentication type that allows the IAS server to pass information to client computers attempting to connect to your network.

In this circumstance, the value contained in the PEAP-TLV is an HTTPS Uniform Resource Locator (URL) that provides client computers running Windows XP with SP2 with the location of the WISP provisioning server. With this URL and parameter, Windows XP can download the WISP XML files to the client computer.

In addition to the URL of the provisioning server, the URL PEAP-TLV includes an action parameter. The action parameter directs the client to perform a specific task. The action parameter included in the URL PEAP-TLV defines tasks such as new customer sign-up, existing account renewal, and password change.

For more information, see "How to Create an IAS Extension DLL and a URL PEAP-TLV" in this paper.

**Server certificate**
For server authentication to client computers, the WISP IAS server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and is issued by a public certification authority (CA), such as Verisign or Thawte, that is trusted by client computers. The Trusted Root Certification Authorities certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust. If you install IAS and Active Directory on the same computer, the computer must have a certificate. If you install IAS and Active Directory on different computers, only the IAS server needs a certificate.

**DHCP server**
The DHCP server must be able to assign valid public IP addresses to computers accessing the network through the wireless access points.

**How a WISP works using IP filters**
The Internet connection process with WPS technology for a WISP with IP filtering differs depending on whether the customer attempting to connect is a new customer or an existing customer. The following example describes the process for a new customer. In addition, how IAS handles an expired account is explained.

**New customer connection example**
When a new customer connects to a WISP and establishes an account, the following four basic stages occur:

1. The customer discovers the WISP network at a Wi-Fi hotspot

2. The customer authenticates as guest

3. The client is provisioned and the customer establishes an account

4. The customer is authenticated with the new account credentials

5. IP filters are removed and the customer gains access to the Internet

In the next section we will look at these stages in more detail.

**1. The customer discovers the WISP network at a Wi-Fi hotspot**
When a customer arrives at the WISP Wi-Fi hotspot with a portable computer running Windows XP Home Edition with SP1, Windows XP Professional with SP1, or Windows XP Tablet PC Edition with SP1, the computer comes within range of the WISP access point beacon.

The Wireless Auto Configuration service on the client computer detects the beacon information from the access point, which is enabled with broadcast Secure Set Identifier (SSID). The SSID is equivalent to the network name.

The customer is informed by Windows XP that a wireless network is available. The customer views information in Windows XP, and if interested, the customer clicks Connect.

**2. The customer authenticates as guest**
Wireless Auto Configuration uses 802.1X and PEAP guest authentication to connect to the WISP network through the access point, automatically passing a null user name and a blank password to the WISP IAS server.

The IAS server is the PEAP authenticator and TLS endpoint for customers who connect as guest, and the TLS tunnel is created between the client and the IAS server. All subsequent messages between client and server pass through this tunnel.

Server authentication is performed when the IAS server verifies its identity to the client computer using a certificate that contains the Server Authentication purpose in Enhanced Key Usage (EKU) extensions. This certificate is issued by a public trusted root certification authority (CA).

The IAS server authenticates and authorizes the customer as guest. In the Access-Challenge message that the IAS server sends to the client is a URL PEAP-TLV. The URL PEAP-TLV is a container with a value that is the URL of the provisioning server. This URL provides the client with the location of the XML master file.

The IAS server also sends IP filters in the form of Vendor Specific Attributes (VSAs) to the access point. These IP filters are applied to the client connection by the access point and are used to isolate the client; the filters block access to all network resources except the WISP provisioning server.

The customer client computer receives an IP address lease from the DHCP server. The address is from a public IP address range configured in a scope on the DHCP server.

**3. The client is provisioned and the customer creates an account**
The XML master file on the provisioning server contains pointers to the XML subfiles. Windows XP downloads the XML master file and subfiles. When the XML sign-up schema is downloaded, the sign-up wizard is launched on the client to allow the customer to create and pay for an account with the WISP.

Using the sign-up wizard on the client computer, the customer steps through the process of signing up for an account. The data entered by the customer is converted by Windows XP into an XML document.

The XML document containing the customer's sign-up data is sent by Windows XP to the Web application on the WISP provisioning server.

The Web application processes the customer payment information. Once payment is verified and sign-up information is completed successfully, the Web application creates a user account in Active Directory, and permissions are applied to the user account by assigning group membership based on the account type chosen by the customer.

An XML document containing the new account credentials is sent from the WISP provisioning server to the client computer. The client computer uses the credentials to configure Wireless Auto Configuration and 802.1X under the name of the WISP.

**4. The customer is authenticated with the new account credentials**
Wireless Auto Configuration restarts the association to the SSID for the WISP.

Wireless Auto Configuration finds the correct 802.11 profile which was downloaded with the other WISP information. Wireless Auto Configuration re-associates with the access point using the correct profile.

802.1X restarts the authentication process using PEAP-MS-CHAP v2 and the new account credentials.

As the client starts the authentication process with PEAP-MS-CHAP v2 authentication, a TLS channel is created between the

customer's client computer and the WISP IAS server.

In the second stage of PEAP-MS-CHAP v2 authentication, the WISP IAS server authenticates and authorizes the connection request against the new account in the Active Directory user accounts database. The IAS server sends an Access-Accept message to the access point.

**5. IP filters are removed and the customer gains access to the Internet**
Because IP filters are used to isolate the client, the IAS server message causes the access server to remove the IP filters from the client connection, granting the customer access to the Internet.

**How IAS handles an expired account**
You can determine the types of account plans that you want to offer your customers. These plans can range from fees based on hourly use to accounts with life spans as long as a day, a month, or longer.

It is important for IAS to determine whether a connecting or connected client computer has a valid account, and to take the appropriate action if the customer's account is expired. The following example illustrates how IAS determines that a twenty-four hour account is current, and how WPS technology behaves when the account expires.

**Twenty-four hour connect option example**
When the customer arrives at the WISP, the customer chooses an access account that has a one day (24 hour) lifespan. The customer and client computer proceed through the account creation process described above, and then connect to the Internet. The following process occurs:

- In the Access-Accept message sent by the IAS server, the IAS server sets a session timeout of 60 minutes for the client computer connection to the access point.
- After 60 minutes, the access point requests that the client reauthenticate. The client reauthenticates successfully and the customer's session is not interrupted.
- Each 60 minutes thereafter, the access point requests that the client reauthenticate. During each authentication the IAS server checks the current time against the expiry time for the user account to discover whether the customer is authorized to access the network.
- On the last re-authentication, at hour 23 in the account lifespan and before 24 hours have passed, the IAS authorization check fails and the IAS server sends a URL PEAP-TLV to the client that contains the account renewal action parameter and an HTTPS URL for an XML master file. The URL PEAP-TLV supplies the customer with the location of the provisioning server where the customer can renew the account.
- Upon receiving the URL in the URL PEAP-TLV, 802.1X requests that Windows XP display the account renewal application to the customer.
- The customer renews the account and 802.1X initiates authentication with the account credentials.
- During authentication with the IAS server, the IAS server authenticates and authorizes the customer against the user accounts database, and sends an Access-Accept message containing a session timeout of 60 minutes to the access point.
- During this process, because the account has not expired, the customer maintains connection to the Internet.

If the customer does not complete the renewal process before the 24 hour account lifespan is reached, then the access point reapplies IP filters and customer access to the Internet is terminated. The customer is then provided with the option of renewing the account for continued access.

| **Note** |
| --- |
| This scenario is currently in development and has not been implemented or tested. It is provided as a general depiction of a possible implementation of WPS technology. |

# WPS Technology for an HSP with IP Filters

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your organization is an HSP, you can deploy WPS technology by using IP filters for client computer isolation. In this circumstance, client computers are provided access to HSP and WISP network resources, such as provisioning servers and IAS servers, and IP filters block access to the Internet until the customer establishes a paid account. After the account is established, the IP filters are lifted, and traffic is allowed between the client computer and the Internet.

In the sections that follow, the components of an HSP network using IP filters are described, and how the components work together during new customer sign-up are detailed.

**Components of WPS technology for an HSP network using IP filters**
A Hotspot Service Provider (HSP) can deploy WPS technology with IP filters for client computer isolation during the account sign-up process.

The following illustration depicts the components of an HSP network using IP filters for client computer isolation.




**Components of an HSP network using IP filters**

This configuration consists of the following components.

**Wireless client**
A computer running Windows XP Home Edition with SP2, Windows XP Professional with SP2, or Windows XP Tablet PC Edition with SP2. The wireless client connects to the HSP network, which is in a public location and is accessible to customers with portable computers and wireless network adapters.

**Wireless access point (RADIUS client)**
The wireless access point is configured as a RADIUS client to the IAS proxy server deployed on the HSP LAN. The wireless access points used by an HSP for WPS technology must meet the following requirements:

- Support for the IEEE standard 802.1X authentication.
- Support for Wi-Fi Protected Access (WPA) is preferred

- The ability to implement client isolation by applying IP filters to the connection with the client.

- Support for RADIUS authentication and RADIUS accounting, including:

    - Support for the Class attribute as defined in RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)," to allow session correlation for RADIUS authentication and accounting records. For session correlation, when you configure RADIUS accounting at your IAS server or proxy, you must log all accounting data that allow applications (such as billing applications) to query the database, correlate related fields, and return a cohesive view of each session in the query results. At a minimum, to provide session correlation, you must log the following IAS accounting data: NAS-IP-Address; NAS-Identifier (you need both NAS-IP-Address and NAS-Identifier because the access server can send either attribute); Class; Acct-Session-Id; Acct-Multi-Session-Id; Packet-Type; Acct-Status-Type; Acct-Interim-Interval; NAS-Port; and Event-Timestamp.

    - Support for accounting interim requests, which are sent periodically by some access servers during a user session, that can be logged. This type of request can be used when the Acct-Interim-Interval RADIUS attribute is configured to support periodic requests in the remote access profile on the IAS server. The access server, in this case a wireless access point, must support the use of accounting interim requests if you want the interim requests to be logged on the IAS server.

    - Support for IP address range filtering.

    - Support for dynamic retransmit timeout (RTO) estimation or exponential backoff to handle congestion and delays in a wide area network (WAN) environment.

In addition, there are some filtering features that the access points must support to provide enhanced security for the network. These filtering options include:

- **DHCP filtering.** The access point must filter on IP ports to prevent the transmission of DHCP broadcast messages in the instance that the client is a DHCP server. The access point must block the client from sending IP packets from port 68 to the network.

- **DNS filtering.** The access point must filter on IP ports to prevent a client from performing as a DNS server. The access point must block the client from sending IP packets from port 53 to the network.

**IP filters**
IP filters that isolate client computers connecting as guest are configured in the remote access policy on the HSP IAS proxy server, and are applied to the client connection by the wireless access point. The IP filters must grant the client computer access to the provisioning server at the HSP and the WISP, and must block access to all other locations.

**Router**
The router must provide multiple ports for connection to access points, the HSP LAN, and the Internet or a WISP, depending on your configuration.

**HSP LAN**
The HSP local area network (LAN) consists of the following components.

**Provisioning server**
The HSP provisioning server is configured with the following components.

**HTTPS Web server**
The Internet Information Services (IIS) or third-party Web server must be deployed with HTTPS.

**HSP XML master files**
HSP XML master files are stored on the HSP provisioning server. Because an HSP can provide customers with a choice of Internet access plans with a variety of WISPs, one XML master file is configured at the HSP and stored on the HSP provisioning server for each WISP available to HSP customers. Each HSP XML master file contains a master XML schema that describes a list of URLs to other XML documents, which are XML subfiles stored on WISP provisioning servers. The XML master schema also describes properties of the WISP for which the XML master file was created and the Time-To-Live (TTL) for the XML master file before it is updated by the HSP provisioning server from the WISP provisioning server. The XML master file schema is:

- **The domain name of the WISP.** For example, if Microsoft Corporation is the WISP, the domain name is "microsoft.com."

- **The friendly name of the WISP.** For example, Microsoft Corporation.

- **The TTL for the XML master file.** When client computers connect to the HSP provisioning server and download the master file, the TTL value tells Wireless Provisioning Services on the client computer when to request an update of the master file from the HSP provisioning server.

- **The list of URLs that point to XML subfiles.** The XML subfiles are located on the WISP provisioning server for which the XML master file was created. Only HTTPS URLs are supported for use within the XML schema, and by WISP provisioning servers. The name of the XML schema described in the subfile is included, as is the version number for the subfile. Possible subfile names include: Signup, SSID, EAP, Location, and Help. For more information, see "XML Schemas" in this paper.

**Server certificate**
For server authentication to client computers, the provisioning server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and be issued by a public certification authority (CA), such as Verisign or Thawte, that is trusted by client computers connecting to the HSP wireless LAN. The Trusted Root Certification Authorities certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust.

**IAS proxy server**
The HSP IAS proxy server is configured with the following components.

**Internet Authentication Service (IAS)**
The IAS proxy server must be a computer running Windows Server 2003, Standard Edition with Service Pack 1 (SP1); Windows Server 2003, Enterprise Edition with SP1; or Windows Server 2003, Datacenter Edition with SP1; and Internet Authentication Service (IAS). IAS is the Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) proxy and server.

The IAS proxy functions both as an IAS proxy and as an IAS server. The IAS proxy server is configured to locally process connection requests for users that authenticate as guest, and to forward non-guest authentication and accounting requests to the IAS servers located at individual WISPS based on the realm name configuration specified in **Connection Request Policies** in the IAS console.

On the IAS proxy, the IAS servers at each WISP are added to a remote RADIUS server group, with a minimum of one remote RADIUS server group created for each WISP. Because Windows Server 2003, Standard Edition, permits the creation of only two remote RADIUS server groups, an HSP offering Internet connectivity to more than two WISPs must deploy IAS proxies and servers running Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition.

**Server certificate**
For server authentication to client computers, the provisioning server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and be issued by a public certification authority (CA), such as Verisign or Thawte, that is trusted by client computers connecting to the HSP wireless LAN. The Trusted Root Certification Authorities certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust.

**DHCP server**
The DHCP server must be able to assign valid public IP addresses to computers accessing the network through the wireless access points.

**WISP 1 LAN**
The WISP 1 LAN is configured with the following components.

**Provisioning server**
The WISP provisioning server is configured with the following components.

**HTTPS Web server**
Like the HSP Web server, the WISP Web server must use HTTPS. The WISP provisioning server stores XML subfiles that are downloaded over HTTPS by client computers running Windows XP with SP2.

**Web application**

The WISP Web server is configured with an account processing Web application. When a customer uses the sign-up wizard on the client computer to create and pay for an account, the customer enters data, such as name, address, and credit card information, that is converted to an XML document and sent to the WISP provisioning server.

The account processing Web application on the provisioning server must be capable of accepting and processing the XML documents containing the user data. For example, the account processing Web application must dynamically create an account in the Active Directory user accounts database, and must contain a credit card verification component to process customers' payment information. The account processing Web application must permit new customers to sign up and to permit existing customers to renew their subscriptions for service.

**XML subfiles**
The WISP provisioning server maintains the XML subfiles that provide the client with all configuration information needed to access the network, create an account, pay for the account, and ultimately access the Internet. For more information about XML subfiles, see "XML Schemas" in this paper.

**Server certificate**
For server authentication to client computers, the provisioning server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and be issued by a public certification authority (CA), such as Verisign or Thawte, that is trusted by client computers connecting to the HSP wireless LAN. The Trusted Root Certification Authorities certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust.

**Domain controller and IAS server**
The WISP domain controller and IAS server is configured with the following components.

**Active Directory**
The user accounts database on the domain controller must be an Active Directory user accounts database or a database that uses Lightweight Directory Access Protocol (LDAP) and supports dynamic creation of user accounts. When a customer signs up for an account, the account processing Web application on the provisioning server creates a new account in the user accounts database, and adds the user to a group that has clearly defined access privileges that match the type of account the customer purchased when signing up.

If you use an accounts database other than Active Directory, IAS authentication and authorization extensions must be written and installed for this process to function correctly.

For information about configuring Active Directory replication, see "Active Directory Replication" in this paper.

**Internet Authentication Service (IAS)**
IAS is the Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy, and is used to authenticate and authorize users connecting to your network. IAS is configured with remote access policies that allow guest authentication for non-domain member computers and users. It is also configured to provide IP filters to RADIUS clients (access points) that apply the filters to client connections during guest authentication. IAS also provides unrestricted access to users with valid accounts.

**Server certificate**
For server authentication to client computers, the provisioning server must maintain a valid certificate in its certificate store. The certificate must contain the Server Authentication purpose in Enhanced Key Usage (EKU) extensions and be issued by a public certification authority (CA), such as Verisign or Thawte, that is trusted by client computers connecting to the HSP wireless LAN. The Trusted Root Certification Authorities certificate store on computers running Windows XP is installed by default with a multitude of certificates issued by public CAs. You must obtain your server certificate from one of the public CAs for which clients already have trust. If you install IAS and Active Directory on the same computer, the computer must have a certificate. If you install IAS and Active Directory on different computers, only the IAS server needs a certificate.

**Extension DLL and URL PEAP-TLV**
An IAS extension DLL defining a URL PEAP-TLV provides IAS with the ability to send the location of the provisioning server to client computers.

PEAP-Type-Length-Value (PEAP-TLV) is an Extensible Authentication Protocol (EAP) authentication type that allows the IAS server to pass information to client computers attempting to connect to your network.

In this circumstance, the value contained in the PEAP-TLV is an HTTPS Uniform Resource Locator (URL) that provides client computers running Windows XP with the location of the WISP provisioning server. With this URL, Windows XP can download the WISP XML files to the client computer.

In addition to the URL of the provisioning server, the URL PEAP-TLV includes an action parameter. The action parameter directs

the client to perform a specific task. The action parameter included in the URL PEAP-TLV defines tasks such as new customer sign-up, existing account renewal, and password change.

For more information, see "How to Create an IAS Extension DLL and a URL PEAP-TLV" in this paper.

**WISP 2 LAN**
The WISP 2 LAN is identical to the WISP 1 LAN described above, and is depicted in "Components of an HSP network using IP filters" to illustrate that an HSP can provide service offers to customers from multiple WISPs.

**How an HSP works with IP filtering**
The Internet connection process with WPS technology for an HSP with IP filtering differs depending on whether the customer attempting to connect is a new customer or an existing customer. The following example describes the process for a new customer. In addition, how IAS handles an expired account is explained.

**New customer connection example**
When a new customer connects to an HSP and establishes an account with a WISP, the following five basic stages occur:

1. The customer discovers the HSP network at a Wi-Fi hotspot
2. The customer authenticates as guest
3. The client is provisioned with the HSP XML master file
4. The customer selects a WISP and establishes an account
5. The customer is authenticated with the new account credentials

In the next section we will look at these stages in more detail.

**1. The customer discovers the HSP network at a Wi-Fi hotspot**
When a customer arrives at the HSP hotspot with a portable computer running Windows XP Home Edition with SP2, Windows XP Tablet PC Edition with SP2, or Windows XP Professional with SP2, the computer comes within range of the HSP access point beacon.

The Wireless Auto Configuration service on the client computer detects the beacon information from the access point, which is enabled with broadcast Secure Set Identifier (SSID). The SSID is equivalent to the network name.

The customer is informed by Windows XP that a wireless network is available. The customer views information in Windows XP, and if interested, the customer clicks **Connect**.

**2. The customer authenticates as guest**
Wireless Auto Configuration uses 802.1X and PEAP guest authentication to connect to the HSP network through the access point, automatically passing a null user name and a blank password to the HSP IAS proxy.

The IAS proxy is configured both as a proxy and as an IAS server. The IAS proxy/server is configured to locally process users that authenticate as guest, and to forward other messages to the WISP IAS servers based on the value of the User-Name attribute in the Access-Request message.

The HSP IAS proxy server acts as an IAS server for guest authentication, processing these requests locally rather than acting as a proxy and forwarding them. The HSP IAS proxy server is therefore the PEAP authenticator and TLS endpoint for customers who connect as guest, and the TLS tunnel is created between the client and the HSP IAS server. All subsequent messages between client and server pass through this tunnel.

Server authentication is performed when the HSP IAS server verifies its identity to the client computer using a certificate that contains the Server Authentication purpose in Enhanced Key Usage (EKU) extensions. This certificate is issued by a public trusted root CA.

The HSP IAS proxy server authenticates and authorizes the customer as guest. In the Access-Challenge message that the IAS server sends to the client is a URL PEAP-TLV. The URL PEAP-TLV is a container with a value that is the URL of the HSP provisioning server. This URL provides the client with the location of the HSP XML master file.

The HSP IAS proxy server also sends IP filters in the form of Vendor Specific Attributes (VSAs) to the access point. These IP filters are applied to the client connection by the access point and are used to isolate the client; the filters block access to all network resources except the HSP provisioning server and the WISP provisioning servers.

The customer client computer receives an IP address lease from the HSP DHCP server. The address is from a public IP address range configured in a scope on the DHCP server.

**3. The client is provisioned with the HSP XML master file**
The HSP master XML file on the HSP provisioning server contains pointers to multiple WISP master XML files, which in turn contain pointers to each respective WISP's XML subfiles. Windows XP downloads the HSP XML master file.

Windows displays to the customer a list of WISPs whose services are offered by the HSP.

For this example, the customer selects services from WISP 1.

**4. The customer selects a WISP**
After the customer selects the WISP, Windows XP connects to the WISP and, using the Background Intelligent Transfer Service, downloads the WISP XML master file and subfiles. When the XML sign-up schema is downloaded, the sign-up wizard is opened on the client to allow the customer to create and pay for an account with the WISP.

Using the sign-up wizard on the client computer, the customer steps through the process of signing up for an account with WISP 1. The data entered by the customer is converted by Windows XP into an XML document.

The XML document containing the customer's sign-up data is sent by Windows XP to the Web application on the WISP 1 provisioning server.

The Web application processes the customer payment information. Once payment is verified and sign-up information is completed successfully, the Web application creates a user account in Active Directory, and permissions are applied to the user account by assigning group membership based on the account type chosen by the customer.

An XML document containing the new account credentials is sent from the WISP provisioning server to the client computer. The client computer uses the credentials to configure Wireless Auto Configuration and 802.1X under the name of the WISP.

**5. The customer is authenticated using the new account credentials and gains Internet access**
Wireless Auto Configuration restarts the association to the SSID for the HSP.

Wireless Auto Configuration finds the correct 802.11 profile which was downloaded with the other WISP information. Wireless Auto Configuration re-associates using the correct profile.

802.1X restarts the authentication process using PEAP-MS-CHAP v2 and the new account credentials.

As the client starts the authentication process, the HSP IAS proxy forwards messages between the client and the WISP IAS server. The connection request policies on the HSP IAS proxy are configured to forward requests to WISP 1 where the ISP's name exists in the realm portion of the RADIUS User-Name attribute. For example, if "user@wisp1.example" is the user name, the IAS proxy server uses the realm portion of the user name, "wisp1.example," to choose the IAS server that should receive the connection request.

In the first stage of PEAP-MS-CHAP v2 authentication, a TLS channel is created between the customer's client computer and the WISP 1 IAS server.

In the second stage of PEAP-MS-CHAP v2 authentication, the WISP 1 IAS server authenticates and authorizes the connection request against the new account in the user accounts database. The WISP IAS server sends an Access-Accept message that is forwarded by the HSP IAS proxy to the HSP access point.

Because IP filters are used to isolate the client, the IAS server message causes the access server to remove the IP filters from the client connection, granting the customer access to the Internet.

**How IAS handles an expired account**
Each WISP that offers connectivity through an HSP can determine the types of account plans to offer customers. These plans can range from fees based on hourly use to accounts with lifespans as long as a day, a month, or longer.

It is important for IAS to determine whether a connecting or connected client computer has a valid account, and to take the appropriate action if the customer's account is expired. The following example illustrates how IAS determines that a twenty-four hour account is current, and how WPS technology behaves when the account expires.

**Twenty-four hour connect option example**
When the customer arrives at the HSP, the customer chooses an access account with WISP 1 that has a one-day (24-hour) lifespan. The customer and client computer proceed through the account creation process described above, and then connect to the Internet. The following process occurs:

- In the Access-Accept message sent by the WISP 1 IAS server, the IAS server sets a session timeout of 60 minutes for the client computer connection to the access point.

- After 60 minutes, the access point requests that the client reauthenticate. The client reauthenticates successfully and the customer's session is not interrupted.

- Each 60 minutes thereafter, the access point requests that the client reauthenticate. During each authentication the IAS server checks the current time against the expiry time for the user account to discover whether the customer is authorized to access the network.

- On the last re-authentication, at hour 23 in the account lifespan and before 24 hours have passed, the IAS authorization check fails and the IAS server sends a URL PEAP-TLV message to the client that contains the account renewal action parameter and an HTTPS URL for an XML master file. The URL PEAP-TLV supplies the customer with the location of the provisioning server where the customer can renew the account.

- Upon receiving the URL in the URL PEAP-TLV, 802.1X requests that Windows XP display the account renewal application to the customer.

- The customer renews the account and 802.1X initiates authentication using the account credentials.

- During authentication with the WISP 1 IAS server, the IAS server authenticates and authorizes the customer against the user accounts database, and sends an Access-Accept message containing a session timeout of 60 minutes to the access point.

- During this process, because the account has not expired, the customer maintains connection to the Internet.

If the customer does not complete the renewal process before the 24 hour account lifespan is reached, then the access point reapplies IP filters and customer access to the Internet is terminated. The customer is then provided with the option of renewing the account for continued access.

| **Note** |
| --- |
| This scenario is currently in development and has not been implemented or tested. It is provided as a general depiction of a possible implementation of WPS technology. |

# Summary (Deploying Wireless Provisioning Services (WPS) Technology)

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you deploy Wireless Provisioning Services technology, you can provide wireless Internet access to new and existing customers at Wi-Fi hotspots in public locations, such as airports and shopping malls. Customers can connect to your network, create and pay for an account, and connect to the Internet without knowing about your services prior to their arrival at Wi-Fi hotspot locations and without manually configuring their wireless computers. To deploy WPS technology, you must use Windows Server 2003, SP1 and Internet Authentication Service (IAS), and your customers must use computers running Windows XP with Service Pack 2 (SP2). Your IAS server must be configured with an IAS extension DLL and URL PEAP-TLV; these technologies allow the IAS server to inform client computers of the location of your provisioning server in IAS Access-Accept messages. In addition, you must deploy an HTTPS provisioning server that provides client computers with network and other configuration parameters described in a collection of XML schemas and provided to client computers in the form of XML master and subfiles. The provisioning server must also maintain a Web application that parses data entered by customers, verifies credit card information, and can dynamically create a user account in an Active Directory or LDAP compatible user accounts database.

# Related Links (Deploying Wireless Provisioning Services (WPS) Technology)

Updated: March 31, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

See the following resources for further information:

- Wireless Provisioning Services at http://go.microsoft.com/fwlink/?LinkId=42996
- WPS Authoring Tool at http://go.microsoft.com/fwlink/?LinkId=40535
- Using the WPS Authoring Tool at http://go.microsoft.com/fwlink/?LinkId=41067
- Windows Server 2003 Internet Authentication Service at http://go.microsoft.com/fwlink/?LinkId=20133
- Deploying SQL Server Logging with Windows Server 2003 Internet Authentication Service (IAS) at http://go.microsoft.com/fwlink/?LinkId=41039
- Remote Access Logging at http://go.microsoft.com/fwlink/?LinkId=41038
- Internet Authentication Service Extensions at http://go.microsoft.com/fwlink/?LinkId=34431
- Internet Authentication Service Reference at http://go.microsoft.com/fwlink/?LinkId=34429
- Windows Server 2003 Wi-Fi at http://go.microsoft.com/fwlink/?LinkId=5969
- SQL Server at http://go.microsoft.com/fwlink/?LinkId=20014
- Windows XP Professional at http://go.microsoft.com/fwlink/?LinkId=20127
- Windows XP Tablet PC Edition at http://go.microsoft.com/fwlink/?LinkId=20128
- Windows XP Home Edition at http://go.microsoft.com/fwlink/?LinkId=20130

For the latest information about Windows Server 2003, see the Windows Server 2003 Web site at http://go.microsoft.com/fwlink/?LinkId=20045.

# Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

| Note |
| --- |
| To download a copy of this document, see http://go.microsoft.com/fwlink/?LinkId=119655. |

By David B. Cross and Carsten B. Kinder, Microsoft Corporation

**In This White Paper**

About This Document (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Overview of the PKI Design Process (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Integration Into Existing Environments (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Windows Server 2003 PKI and Dependencies (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Deployment Planning (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Creating Certificate Policies and Certificate Practice Statements

Example Scenario for Contoso

Stand-alone Offline Intermediate CA (IntermediateCA1)

Stand-alone Offline Intermediate CA (CorporateSub2CA)

Online Enterprise Issuing CAs (CorporateEnt1CA)

Certification Authority Maintenance

Appendix A: Directory Objects

Contents of \\Localhost\CertConfig and \\Localhost\CertEnroll

Relationship of the Configuration Container and Certificate Store

Default CA Certificate and CRL Storage

Mapping Custom Object Identifiers to Friendly Names

CAPolicy.inf Syntax

CRL Distribution Point Replacement Token

CRL Publishing Properties

AIA Publishing Properties

Sample Script to Configure CorporateRootCA

Sample Script to Configure IntermediateCA

Sample Script to Configure the EnterpriseSubCA

Appendix B: Parameters for a Three-Tier CA Topology

Appendix C: Additional Information

# About This Document (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

This document is a quick start guide that you can use to set up a Windows Server 2003 public key infrastructure (PKI). It provides all the information that you need to deploy a viable PKI that is based on Windows Server 2003 technology.

The document outlines a proven PKI architecture that is applicable for the majority of organizations. It includes tips and decision best practices that have been obtained from customer experiences.

To ensure that configuration steps have been implemented correctly, this document also includes useful verification steps. Where possible, information regarding the configuration and installation of a server running a member of the Windows 2000 Server family is provided for comparison.

**Document Structure**

This document is based on Designing a Public Key Infrastructure, in the *Microsoft Windows Server 2003 Deployment Kit* that is listed in the "Related Information" section in this document. Some issues are addressed only in the *Microsoft Windows Server 2003 Deployment Kit* chapter, while other issues are described only in this document. The similar structure provides easier navigation through the planning and deployment phase if you work with all of these documents.

| **Important** |
| --- |
| This document refers to features included with Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition. These features are not included on computers running Windows Server 2003, Web Edition. |

**Scope**

This document provides implementation guidelines for administrators who are deploying a Windows Server 2003 PKI in their organization.

This white paper is not an introduction to public key technologies, certification authorities, or certificates. It assumes that the reader has a good understanding of PKI and Active Directory concepts.

Because this white paper is focused on technology, it does not outline organizational guidelines and rules that are mandatory for a successful PKI implementation. You should apply organizational requirements and best practices in conjunction with the recommendations in this white paper to ensure a successful deployment.

A number of detailed best practices that are combined with real-world field experience from Microsoft and Hewlett Packard Consulting Services have been incorporated into this white paper.

**Related Information**

This documentation extends the Designing a Public Key Infrastructure chapter in the *Microsoft Windows Server 2003 Deployment Kit*, which contains overall PKI planning and design, and the Windows Server 2003 Help topics, which contains checklists and configuration information. The chapter in the *Microsoft Windows Server 2003 Deployment Guide* focuses on broad deployment considerations.

# Overview of the PKI Design Process (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Designing a PKI involves the following steps which may or may not be performed in this order:

- Outline the business scenario
- Define the application certificate requirements
- Create certificate policies and practices statements
- Design the certification authority (CA) infrastructure
- Create a certificate renewal strategy
- Develop a CA management plan

# Integration Into Existing Environments (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

When you combine client computers running Microsoft Windows 2000 Professional or Windows XP Professional and computers running a member of the Windows Server 2003 family, you have a range of PKI enhancements that let you securely extend your network to employees, partners, customers, and services. It enhances the management and performance features of the Windows 2000 security infrastructure. Windows XP Professional and the Windows Server 2003 family offer many PKI-specific business benefits to organizations that require secure business processes and IT infrastructures.

The base set of features is provided in the Windows Server 2003 family, and enhanced certification authority functionality is provided in Windows Server 2003 Server, Enterprise Edition, and Windows Server 2003 Server, Datacenter Edition. The PKI that is part of the Windows Server 2003 release is an improved version of the Windows 2000 PKI functionality. Nevertheless, you can combine a Windows Server 2003-driven PKI with an existing Windows 2000 Active Directory environment and certification authority (CA) infrastructure.

Client computers running either the Windows 2000 or Windows XP operating systems will benefit the most from a Windows Server 2003 PKI deployment, along with hardware devices that support the Windows environment. For more information about the capabilities of each client, see the Windows Server 2003 Help.

**Determining Secure Application Requirements**

The Windows Server 2003 Standard and Enterprise Editions include a full-featured PKI that delivers the business benefits of current public key cryptography. Users, computers, and services benefit from encryption and signing capabilities.

The Windows Server 2003 PKI supports a broad range of applications, including:

- Secure logon with smart cards
- Confidential and secure e-mail
- Secure code
- Trusted, on-demand access to network resources for remote users and trusted, permanent network connectivity for remote offices with network security, including remote access, virtual private networks (VPN) and wireless authentication
- File protection in the event of stolen or lost portable computers and other storage devices
- Access control and single-identity authorization across a range of Web and application servers
- Digital signatures that enable tamper-proof, legally-binding transactions
- Scalable technology to support millions of users and high-volume digital signature transactions

For more information about how the Windows PKI supports these applications, see the following articles on the Microsoft Web site:

- How PKI Works on the Microsoft TechNet Web site
- Microsoft Windows 2000 Public Key Infrastructure Introduction on the Microsoft TechNet Web site
- Applications Overview on the Microsoft TechNet Web site

The Windows Server 2003 PKI solution has several advantages over commercial third-party PKIs that are not part of the

operating system and must be purchased separately. Users and access control are managed centrally through Active Directory, which simplifies the overall PKI management burden. Further, a Windows Server 2003 PKI does not require either per-certificate or per-user license fees that would raise the total cost of ownership (TCO) of the system. The PKI functionality in the Windows Server 2003 family integrates very well with many other features of the operating system.

# Windows Server 2003 PKI and Dependencies (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

From a technical perspective, a Windows Server 2003 PKI has some requirements before you can deploy it. This section describes fundamental process information and installation details for a successful Windows Server 2003 PKI implementation.

**New Features of a Windows Server 2003 CA**
Windows XP and Windows Server 2003 environments can benefit the most from all of the features of the Windows Server 2003 certification authority (CA), but mixed configurations with earlier versions of Windows are also supported, with a little less functionality.

The following table lists the features that are described later in this section. These features are available through the CA if the CA is installed on a specific version of the operating system.

**Table 1 PKI feature support with a CA that is installed on different versions of the operating system**

| | **Windows Server 2003, Enterprise Edition or Datacenter Edition** | **Windows Server 2003, Standard Edition** | **Windows 2000 Server** |
|---|---|---|---|
| V2 templates | Enterprise CA only | Not supported | Not supported |
| Key archival and recovery | Supported | Not supported | Not supported |
| Auto-enrollment | Both user and computer certificates supported | Computer certificates supported | Computer certificates supported |
| Delta certificate revocation lists (CRLs) | Supported | Supported | Not supported |
| Qualified subordination | Supported | Supported | Not supported |
| Role separation | Supported | Not supported | Not supported |

**Note**

Windows Server 2003, Web Edition does not include certification authority functionality, but may be used as a PKI client.

The following table lists the features that can be used at the client with a given CA infrastructure:

**Table 2 PKI features that are available to clients**

| | Windows XP Professional client | | Windows 2000 Professional client | |
|---|---|---|---|---|
| | Windows 2000 CA | Windows Server 2003 CA | Windows 2000 CA | Windows Server 2003 CA |
| V2 templates | Not supported | Supported | Not supported | By Web enrollment support only |
| Key archival and recovery | Not supported | Supported | Not supported | By Web-enrollment support only |
| Auto-enrollment | User: Not supported<br>Computer: Supported | User: Supported<br>Computer: Supported | User: Not supported<br>Computer: Supported | User: Not supported<br>Computer: Supported |
| Delta CRL | Not supported | Supported | Not supported | Not supported |
| Qualified subordination | Not supported | Supported | Not supported | Not supported |

For additional information, see the following articles:

- What's New in Security for Windows XP Professional and Windows XP Home Edition on the Microsoft Web site
- Data Protection and Recovery in Windows XP on the Microsoft Web site

**Version 2 Templates**

Templates are the building plan for certificates. A template turns a certificate request into a certificate signed with the CAs private key. For example, a template defines the validity time of a certificate and the certificate's subject name.

The most significant difference between the version 1 (V1) and version 2 (V2) templates is that V1 templates are predefined and unchangeable. With V2 templates, a CA administrator is able to configure a wide range of settings that apply during certificate enrollment, such as minimum key length, subject name definition, enrollment requirements like enrollment agent signature, and so on.

V2 templates are available only with an enterprise CA that is running Windows Server 2003, Enterprise Edition or Datacenter Edition. An enterprise CA that is running Windows Server 2003, Standard Edition does not support V2 templates.

Generally, templates are stored in the Active Directory configuration naming context and are usable with any CA in an Active Directory forest if they are assigned to the CA. A single set of templates are available for use by all CAs in the forest. However, V2 templates can only be utilized with Windows Server 2003, Enterprise Edition or Datacenter Edition CAs.

To use V2 templates, the Active Directory schema must be extended to the Windows Server 2003 schema in the forest. If the Active Directory environment consists of Windows Server 2003 domain controllers only, no action is required to benefit from a Windows Server 2003 PKI. If all domain controllers of a forest that hosts the Windows Server 2003 CAs are running Windows 2000 Server, you must also install Microsoft Windows 2000 Service Pack 3 (SP3) or later on all domain controllers, in addition to the new Windows Server 2003 schema definitions. For more information about how to upgrade the schema, see Prepare the Active Directory environment in this document.

- For more detailed information about certificate templates, their usage, and their definition capabilities, see "Certificates" in the Windows Server 2003 Family Help.
- For more information about certificate enrollment, see "Certificate Enrollment" in the MS Windows 2000 Public Key Infrastructure Introduction white paper on the Microsoft TechNet Web site.
- For more information about extending the schema for V2 templates, see the Windows Server 2003 Help topics regarding certificate templates.

**Certificate Enrollment**

You can enroll V2 templates with any computer running Windows XP or later through the default enrollment methods, including the Certificates Microsoft Management Console (MMC), the built-in auto-enrollment mechanism, Web enrollment support, or command-line tools.

A computer running Windows 2000 cannot use the Certificates MMC to enroll a V2 template. However, any client that is running Microsoft Internet Explorer 5.01 or later can use a V2 template to enroll certificates through the Web enrollment methods and a downloaded ActiveX control. To download the ActiveX control on a client computer, it is necessary to log on as an Administrator or Power User on the local computer. In addition, clients running Windows 2000 can enroll V2 templates through a Terminal Server connection running on an appropriate member of the Windows Server 2003 family.

| **Note** |
| --- |
| An enrollment agent that enrolls certificates that are based on V2 templates requires either a Windows XP or Windows Server 2003 enrollment station. There is no support to enroll V2 templates with an enrollment agent on a Windows 2000 enrollment station. Nevertheless, certificates that are based on V2 templates that have been enrolled through an enrollment agent on either a Windows XP or Windows Server 2003 enrollment station can be used on a Windows 2000 client computer. |

If certificates have been enrolled with a Windows 2000 PKI where only V1 templates were available, there is no immediate need to re-enroll or renew these certificates with V2 templates.

The following table lists different enrollment methods that are supported on computers that are running Windows 2000, Windows XP, or the Windows Server 2003 family. You can use scripted enrollment with support of CAPICOM and Xenroll. (CAPICOM and Xenroll documentation, including samples, can be found on the Microsoft Developer Network (MSDN) Web site.)

**Table 3 Certificate Enrollment**

| | **Certificates MMC** | **Web-enrollment** | **Scripted enrollment** |
| --- | --- | --- | --- |
| Self enrollment on a Windows 2000 workstation | V1 template: Yes<br>V2 template: No | V1 template: Yes<br>V2 template: Yes | V1 template: Yes<br>V2 template: Yes |
| Self enrollment through a Windows Server 2003 Terminal Server session | V1 template: Yes<br>V2 template: Yes | V1 template: Yes<br>V2 template: Yes | V1 template: Yes<br>V2 template: Yes |
| Enrollment agent on a Windows 2000 workstation | V1 template: No<br>V2 template: No | V1 template: Yes<br>V2 template: No | V1 template: No<br>V2 template: No |
| Enrollment agent through a Windows Server 2003 Terminal Server session | V1 template: No<br>V2 template: No | V1 template: Yes<br>V2 template: Yes | V1 template: No<br>V2 template: No |

| **Note** |
| --- |
| Because a PKI is a forest resource, the Active Directory site structure is not taken into account when any kind of certificate is requested and issued. An Active Directory–integrated certificate requester enumerates all registered enrollment services in Active Directory and sends its request to a CA that can enroll the certificate type that the user wants. The client does not necessarily choose the closest CA, from a network perspective. Because of this, you should verify that the CA deployment ensures that any client has reliable network connectivity with a CA. |

#### User Certificate Autoenrollment

Autoenrollment provides a quick and simple way to issue user certificates and to benefit from applications that can use PKI. User autoenrollment also minimizes PKI deployment costs.

Certificate autoenrollment also works in a Terminal Server session if you use a Windows Remote Display Protocol (RDP) 5.1 client.

When you use a computer that is running Windows XP, you can automatically enroll users and computers for certificates, including smart card-based certificates. In contrast, the Windows 2000 Server family only supports certificate autoenrollment for computer certificates. User certificate auto-enrollment builds on the standard Windows security model for domain authentication and authorization. This model may not be suitable for all certificate issuance or scenarios.

Using the new autoenrollment feature, organizations can manage the certificate lifecycle through V2 templates for users. This includes:

- Certificate renewal
- Superseding of certificates
- Multiple signature requirements

Depending on the configuration of the template that is used for autoenrollment, the user can be notified when a certificate enrollment or renewal is performed.

Certificate autoenrollment is based on the combination of Group Policy settings and V2 certificate templates. This combination allows certificate enrollment and renewal in the background for computers and users at any time when you apply Group Policy.

To perform autoenrollment, the certificate requester must be registered and authenticated as either a user or computer in Active Directory.

For more information, see Certificate Autoenrollment in Windows Server 2003 on the Microsoft TechNet Web site.

**Certificate Renewal**

When a certificate comes to the end of its lifetime, it must be renewed or replaced to ensure that the certificate owner is able to continue with the certificates purpose.

In certificate renewal, the renewal requester already owns a certificate. The renewal takes the information of the existing certificate into account when the renewal request is submitted. A certificate can either be renewed with a new key or the existing key can be used for the renewed certificate.

If a certificate was enrolled with a V2 template, it cannot be renewed if it was based on a V1 template. However, a certificate that was enrolled with a V1 template can be renewed with a certificate that was made from a V2 template.

**Table 4 Certificate Renewal**

| | **Certificates MMC** | **Web-enrollment** | **Scripted enrollment** |
|---|---|---|---|
| Self renewal on Windows 2000 workstation | V1 template: Yes<br><br>V2 template: No | V1 template: No<br><br>V2 template: No | V1 template: Yes<br><br>V2 template: Yes |
| Self renewal on Windows 2000 workstation with Windows Server 2003 Terminal Server session | V1 template: Yes<br><br>V2 template: Yes | V1 template: No<br><br>V2 template: No | V1 template: Yes<br><br>V2 template: Yes |
| Renewal with enrollment agent on Windows 2000 workstation | V1 template: No<br><br>V2 template: No | V1 template: No<br><br>V2 template: No | V1 template: No<br><br>V2 template: No |
| Renewal with enrollment agent on Windows 2000 workstation with Windows Server 2003 Terminal Server session | V1 template: No<br><br>V2 template: No | V1 template: No<br><br>V2 template: No | V1 template: No<br><br>V2 template: No |

**Key Archival and Recovery**

Key archival and recovery is only available for encryption certificates by V2 templates, because the archival option must be individually set for each template. Key archival is most often used for encryption keys that are used to protect persisted data.

Private keys that are associated with certificates that are used only for digital signature are not archived and are blocked by the certification authority. The archival and recovery function that is available with the Microsoft Exchange 2000 Key Management Server (KMS) has been replaced by the enterprise CA running Windows Server 2003, Enterprise Edition.

The enterprise certification authority on a computer that is running Windows Server 2003, Enterprise Edition, supports migration of the archive database from the Exchange 2000 KMS to ensure a smooth transition of technologies.

Encrypting File System (EFS) will continue to support decentralized data recovery methods as well as key archival on clients that are running Windows XP.

**Delta CRLs**
Delta certificate revocation lists (CRLs) decrease the network traffic that is caused when a new certificate revocation list needs to be downloaded. Without delta CRLs, a client must receive the base CRL that contains all certificates that are revoked by a CA. To decrease the CRL's size and make more frequent updates valuable, delta CRLs only retain the certificates that have been revoked since the last publication of the base CRL.

Some limitations apply to delta CRLs:

- Delta CRLs are issued by Windows Server 2003 stand-alone and enterprise CAs.
- Only clients that are running Windows XP Professional and later are able to check the validity of certificates against delta CRLs.

For more information about this topic, see Troubleshooting Certificate Status and Revocation on the Microsoft TechNet Web site.

**Qualified Subordination**
Qualified subordination allows cross-certification of CA certificates with name constraints and provides for more precise control of certificate trusts. With qualified subordination, an administrator can also include or exclude certificate purposes. For example, qualified subordination might reject Internet Protocol security (IPSec) usage with a third-party certificate, but allows secure e-mail with the same certificate, even if the certificates key usage would allow IPSec and secure e-mail.

Qualified subordination requires a Windows XP or Windows Server 2003 operating system as the certificate requester and a Windows Server 2003, Enterprise Edition CA.

**Simple Certificate Enrollment Protocol**
You can implement secure networking with the Simple Certificate Enrollment Protocol (SCEP). The Microsoft SCEP (MSCEP) component is an Internet Server Application Programming Interface (ISAPI) filter that uses Microsoft Internet Information Services (IIS) and is installed directly on a CA to support the SCEP enrollment protocol with network devices.

For more information, see article 249125, Using Certificates for Windows 2000 and Cisco IOS Interoperation, in the Microsoft Knowledge Base.

**Role Separation**
There are a number of tasks in the PKI process that you should understand:

- **Certificate enrollment**. Sends a certificate request to the CA, and then the CA issues the certificate and then deploys a certificate to the certificate holder.
- **Certificate renewal**. Sends a request to renew an existing certificate to the CA, the CA issues the certificate, and then the CA deploys a certificate to the certificate holder.
- **Certificate revocation**. Revokes certificates and publishes the certificate revocation list (CRL).
- **Recovery**. Provides the certificate holder with both a certificate and a key that are stored in the CA database.

There are also a number of roles that are related to a Windows Server 2003 CA, although you may not need all of these roles:

- The **CA manager** maintains the CA and its configuration.
- The **CA administrator** delegates certificate-management permissions to Certificate Managers.
- The **certificate manager** issues and revokes certificates.
- The **enrollment agent** requests and deploys certificates.

- The **certificate holder** requests self-maintained certificates and is able to use the certificate.
- The recovery agent recovers certificates for specific applications, such as EFS.

The following table shows which role can perform a particular task and what permissions are required to perform that function. The top row lists various roles, and the left column lists the tasks. The table text describes the permission that is required to perform each task.

**Table 5 Certificate Roles and Tasks**

| | **Certificate Holder** | **Enrollment Agent** | **Recovery Agent** | **Certificate Manager** | **CA Manager** |
|---|---|---|---|---|---|
| Maintain | Not applicable | Not applicable | Not applicable | Not applicable | Requires CA Manager permissions on the CA object |
| Request | Requires certificate holders membership on the certificates template ACL | Requires certificate holders membership on the certificates template ACL | Not applicable | Not applicable | Not applicable |
| Renew | Requires certificate holders membership on the certificates template ACL | Requires certificate holders membership on the certificates template ACL | Not applicable | Not applicable | Not applicable |
| Issue | Not applicable | Not applicable | Not applicable | Certificate Manager permission on a template | Not applicable |
| Revoke | Not applicable | Not applicable | Not applicable | Certificate manager permission | Not applicable |
| Recover | Not applicable | Not applicable | Key recovery agent certificate | Not applicable | Not applicable |

**CA Permissions**

For a Windows Server 2003 stand-alone CA that is installed on a server that is not a member of an Active Directory domain, local administrator permissions are mandatory to manage the CA functions.

On a server that is a domain member, the user who installs an enterprise CA must be a member of the Active Directory Root Domain Admins and Enterprise Admins security group. You should ensure that the installation account is a member of both security groups. This set of permissions is required for any enterprise CA installation, and it assumes that the Enterprise Admins group or Domain Admins group also is a member of the local server Administrators group. To install a stand-alone CA, only local administrator privileges are required.

During setup, containers and objects that contain enrollment and CA information are created as part of the configuration container of Active Directory. For a list of object default permissions that are used by a CA, see article 239706, "Default Permission Settings for Enterprise Certificate Authority," in the Microsoft Knowledge Base.

It is recommended that you use only the Certification Authority MMC to change security permissions for the CA. If you use other mechanisms, such as the Active Directory Sites and Services MMC, you may cause an unsupported environment due to a configuration mismatch between Active Directory and the local CA registry. The Certificate Templates MMC ensures consistency in the ACLs. If ACLs are changed manually, specific permissions may be missing and the CA will not function as expected.

**Command-line Administration Tools**

Command-line administration tools are part of the Windows Server 2003 family and may be installed on computers running Windows XP and later through the Windows Server 2003 Administration Tools Pack, which is available on the Windows Server 2003 installation media. Command-line tools that are required for CA administration on Windows 2000 operating systems are

only installed with Windows 2000 Certificate Services. For more information about **Certutil.exe** and **Certreq.exe**, including a description and the necessary syntax, see the Windows Server 2003 Help.

**CA Fault Tolerance**

Generally, you should use certification authority fault tolerance because:

For online CAs, it provides certificate issuance services.

For both online and offline CAs, it provides certificate revocation information.

Neither Windows 2000 nor Windows Server 2003 technology supports native clustering of the CA database or certificate services. Only one CA instance can be installed at a time on a server running a Windows Server 2003 operating system.

An enterprise CA is designed to provide natural fault tolerance in an Active Directory environment. If one enterprise CA does not work or is not available, client services will automatically attempt enrollment with the next available enterprise CA in the forest. No errors are generated and no user interaction is required. For more information, see "Online Enterprise Issuing CAs" later in this document.

If a CA is not available because of a hardware failure, for example, it might still be necessary to publish the CRL on a regular basis. The CRL publication interval depends on the CA configuration. If the CA does not publish the CRL in time, clients cannot verify certificates against the latest version of the CRL.

To publish a CRL on behalf of a CA, you must own the CA private key. If the CA private key has been exported to a file, it is technically possible to resign a CRL on behalf of the CA and extend the lifetime of the CRL.

| **Note** |
| --- |
| Exporting the CA private key could raise a security risk because the owner of the CA's private key is able to act on behalf of the CA. The CA private key must be maintained very carefully and must be stored in a secure vault that is protected through secure and audited processes. |

# Deployment Planning (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Before you can deploy a PKI, you should go through a well-defined planning phase. If you do not do this, the PKI can become valueless after only a short time in operation. To avoid this issue, make sure that the deployment planning covers the following areas.

**Table 6 PKI Planning Considerations**

| Planning area | Possible considerations |
|---|---|
| Business requirements | Defining application requirement<br>Defining solutions goals<br>Choosing appropriate technology |
| CA requirements | Insource the CA infrastructure<br>Outsource the CA infrastructure<br>Interoperability with application requirements<br>PKI trust model |
| Enrollment policy | Certificate practice statements<br>Users and computers<br>Use of certificate templates<br>Service level requirements |
| Revocation policy | CRLs, delta CRLs, Online Certificate Status Protocol (OCSP)<br>Replication latency |

**Designing the CA Infrastructure**

For more information about how to decide what services are provided by the CA types, see the articles on the following Microsoft Web sites:

- "MS Windows 2000 Public Key Infrastructure Introduction" on the Microsoft TechNet Web site
- "An Introduction to the Windows 2000 Public-Key Infrastructure" on the Microsoft Web site
- "Cryptography and Microsoft Public Key Infrastructure" on the Microsoft TechNet Web site
- "Planning Your Public Key Infrastructure" on the Microsoft TechNet Web site

**Considerations**

There are some important parameters that help when a organization starts its PKI planning. From a technical viewpoint, there are a number of key factors that can give you rough estimates:

The number of CAs that you need can be estimated by:

- The size and the geographical spread of the deployment
- The required trust relationship between certificate holders and the CA

- Requirements for different certificate practice statements (CPS)
- Technical requirements based on application demands
- Partner relationships and trust model requirements
- Security requirements, availability, and service levels indicate the depth of the hierarchy and the CA locations.

**Defining CA Types and Roles**
To plan your CA infrastructure, you need to understand the different types of CAs that are available and the roles that the CAs can take.

The following section supplies the most important planning information.

**Choosing Enterprise or Stand-alone CAs**
Certificate Services offers two types of CAs that have different feature sets: enterprise CAs and stand-alone CAs. A Windows Server 2003 PKI may consist of both types of CAs, which is often recommended for the enterprise environment. A comparison of strengths of the stand-alone CA and the enterprise CA may help you decide what CA type is required for which role.

A stand-alone CA should be used if:

- It is an offline root or offline intermediate CA.
- Support of templates that you can customize is not required.
- A strong security and approval model is required.
- Fewer certificates are enrolled and the manual work that you must do to issue certificates is acceptable.
- Clients are heterogeneous and cannot benefit from Active Directory.
- It is combined with a third party Registration Authority solution in a multi-forest or heterogeneous environment
- It issues certificates to routers through the SCEP protocol

An enterprise CA should be used if:

- A large number of certificates should be enrolled and approved automatically.
- Availability and redundancy is mandatory.
- Clients need the benefits of Active Directory integration.
- Features such as autoenrollment or modifiable V2 templates are required.
- Key archival and recovery is required to escrow encryption keys

The following table is an overview about the preferred roles for both CA types. Depending on the CA topology, these roles can be taken by a smaller or larger number of CAs.

**Table 7 CA Types and CA Roles**

| CA type | 3 tier | 2 tier | 1 tier |
|---|---|---|---|
| Offline root CA | Stand-alone CA | Stand-alone CA | |

| | | | |
|---|---|---|---|
| Offline intermediate CA | Stand-alone CA | | |
| Issuing CA | Enterprise CA | Enterprise CA | Enterprise CA |

**Table 8 Comparison of Stand-alone and Enterprise CAs**

| Windows Server 2003 Stand-alone CA | Windows Server 2003 Enterprise CA |
|---|---|
| CA configuration can be published into Active Directory. | CA configuration is always published into Active Directory. |
| CRL and CA certificate must be manually published into Active Directory. | CRL, Delta CRL, CA certificate, and cross certificates are automatically published to the forest where the CA configuration was registered.<br><br>Certificates are automatically published into a directory service if this is specified on a per template level. Certificate publishing may be defined as an attribute on the template in Active Directory. |
| By default, certificate enrollment is available only by using Web enrollment support. | By default, certificate enrollment is possible by using Web enrollment or the Certificates MMC. |
| Certificate request processing is done by using Hypertext Transfer Protocol (HTTP) or Secure Hypertext Transfer Protocol (HTTPS). | Certificate request processing is done primarily by using RPC/Distributed Component Object Model (DCOM) or HTTP and HTTPS protocol. |
| Certificate is based on V1 templates with custom object identifier (also known as OID). | Also issues certificates that can be modified and duplicated, based on V2 templates. |
| User must manually type identification information when the certificate is requested. | User identification information is always automatically retrieved from Active Directory, regardless of whether it is requested through Web enrollment or the Certificates MMC. |
| Enrollment method (automatic or pending) is valid for all templates. You cannot apply a configuration to individual templates. | You can individually set the enrollment method on each template. |
| Certificates are manually approved. | Certificates are manually approved or they are approved through Active Directory authentication and access control. |
| Certificates are not published to a directory location, but to the client or the CA. without a custom-developed policy module. | Depending on the type of certificate, certificate is automatically enrolled into the requester's certificate store and published to Active Directory, based on template definition. |
| Does not support certificate publishing and object management based on Active Directory. | Supports certificate publishing and object management based on Active Directory. |
| Can be installed on a domain controller, member server, or stand-alone server (workgroup member). | Can be installed on a domain controller or member server. (The CA is registered as a forest resource.) It must not be installed on a stand-alone server (workgroup member). |

**Authentication and Authorization**

Stand-alone CAs use local authentication for certificate requests, mainly through the Web enrollment interface. Stand-alone CAs provide an ideal service provider or commercial PKI provider platform for issuing certificates to users outside of an Active Directory environment where the user identity is separately verified and examined before the request is submitted to the CA.

A CA running Windows Server 2003, Enterprise Edition, uses DCOM and Kerberos impersonation for authenticating requesters. It compares the client token against an access control list (ACL) set on the certificate template, as well as the DCOM enrollment interface on the CA itself, when a certificate is requested. A Windows 2000 Server CA uses remote procedure call (RPC) instead of DCOM to authenticate a requester. After the user is authenticated and authorized to gain access to the requested template, the CA can immediately process the request, as long as the user has the appropriate enrollment permissions on the template and if the CAs configuration is set to autoenroll.

**Certificate Request Approval**

When a certificate request reaches a CA that is running a member of the Windows Server 2003 family, both CA types (enterprise and stand-alone) can immediately issue the certificate or put it into a pending state. It is the responsibility of the CA

administrator to configure the enrollment method globally for a CA or on a per-template basis. On a Windows 2000 Server CA, the enrollment method setting is valid only on a CA level: all certificates that are issued take this configuration into account. For a Windows Server 2003, Enterprise Edition enterprise CA, the enrollment method can be set individually for a V2 template.

On a Windows 2000 Server enterprise CA, there is no choice for the enrollment method because it immediately approves and issues certificates. On a Windows 2000 Server stand-alone CA, the enrollment method is applied on the CA level and cannot be set on the template level This occurs because a Windows 2000 CA works only with V1 certificates, which cannot maintain enrollment permissions.

Default configurations of stand-alone CAs rely on administrative action both to verify the requester's identity (known as authentication) and to issue the certificate (known as authorization). Here, the Web enrollment support acts as the registration authority (RA) and the CA acts as the enrollment station. Because of this, it is not recommended that you have a standalone CA automatically issue certificates without administrative approval, because a requester's identity cannot be verified. For additional information about certificate enrollment, see "Allowing for autoenrollment" in Help and Support Center.

**Offline and Online CAs**

Traditionally, the decision of whether to use either an online or offline CAs involves a compromise between availability and usability versus security. The more sensitive that the key material is and the higher the security requirements are, the less accessible the CA should be to users.

| **Important** |
| --- |
| For security reasons, a CA should always run on a separate computer. Do not install an online CA on a domain controller, even if it is technically possible. |

To maintain a CA offline, different approaches may be applied through physical or technical protection techniques as described below:

**Physical Protection**

- Remove the hard disk drive and lock it in a secure location.
- Shut down and power off the system.
- Disconnect the network cable, but keep the system running.
- Protect the system from the network by using either a firewall or a router.

**Technical Protection**

- Keep the system online, but stop the CA service.
- Use a hardware security module (HSM) with an HSM-operator hardware token to limit access to the CA private key. For more information, see "Hardware CSPs," later in this document.

Maintaining a CA either online or offline is a standard functional definition of the CAs operation mode. You can turn an offline stand-alone CA into an online stand-alone CA if you connect it to the network. Any stand-alone CA that runs on a server in a workgroup and is connected to the network can become an offline CA by using one of the approaches that was mentioned earlier.

The stand-alone root CA is usually placed offline because it is the single point of trust for an entire organization or for several organizations. The lifetime of trust depends on the CAs certificate lifetime, but should be planned for the long term. If a CA must be trusted for long periods of time, you should take that CA offline to provide additional security measures. Also, intermediate CAs are typically configured as offline CAs. An intermediate CA is subordinate to a root CA, but also serves as a parent CA to one or more CAs. Those CAs may be issuing CAs or intermediate CAs. In a CA with at least three tiers, an intermediate CA is a mid-tier CA.

Every online CA implies availability and network connectivity. Online CAs are typically issuing CAs, because issuing CAs respond to requests from users, computers, services, and network devices, such as routers. Every enterprise CA must be an online CA, because it requires connectivity to Active Directory at all times to obtain configuration information, validate requests, and publish certificates. An online CA provides more surface area for security attacks.

| Note |
| --- |
| As a best practice, an offline CA server should be placed in a secure vault until a subordinate CA certificate needs to be issued or a new CRL needs to be published. |

**Hardware CSPs**
You might consider using one or more HSMs in your PKI topology. An HSM is a dedicated hardware device that is managed separately from the operating system. These modules work with any Windows Server 2003 CA to provide a secure hardware store for CA keys. From an operating system view through the CryptoAPI interfaces, the HSM is seen as a cryptographic service provider (CSP) device.

The HSM provides highly secure operational management that is protected by multilayered hardware and software tokens, as well as a number of other key features, including:

- Hardware-based, cryptographic operations, such as random number generation, key generation, and digital signatures, as well as key archival and recovery.
- Hardware protection of valuable private keys that are used to secure asymmetric cryptographic operations.
- Secure management of private keys.
- Acceleration of cryptographic operations, which relieves the host server of having to perform processor-intensive, cryptographic calculations.
- Load balancing and failover in hardware modules by using multiple HSMs that are linked together through a daisy chain.

Although HSMs increase security by raising the level of key protection, HSMs increase the complexity and cost of the PKI.

Several vendors offer HSMs that work well on computers that are running either Windows 2000 Server or members of the Windows Server 2003 family. For more information about how to install HSMs that are proven to work with Windows-based CAs, see the section "Installing an HSM on an offline root CA" later in this document.

**Selecting a Trust Model**
Trust is a logical relationship established between domains to allow pass-through authentication, in which a trusting domain honors the logon authentications of a trusted domain. User accounts and global groups that are defined in a trusted domain can be given rights and permissions in a trusting domain, even though the user accounts or groups don't exist in the trusting domain's directory. Because a CA is a certificate holder of a CA certificate and an end entity might be a certificate holder of a user certificate, the trust relationship between the issuing CA and the holder is always the same. In a rooted x.509 PKI hierarchy, the trust relationship is inherited from top to bottom.

You can also control trust relationships through certificate trust lists and through qualified subordination. The selection of an appropriate trust model can determine success for a PKI. An organization must think about the number of tiers that the CA topology requires. The hierarchy can be extended from top to bottom and the number of CAs that are used for one level can grow. Note that deeper structures add complexity to the trust management of the PKI.

For more information, see the article "Trusted Root Certificates That Are Required By Windows 2000" on the Microsoft Knowledge Base.

**Specifying CA Roles**
An ideal PKI hierarchy design divides the responsibility of the CAs. A topology that is designed with requirements that have been carefully considered provides the most flexible and scalable enterprise configuration. In general, CAs are organized in hierarchies. Single tier hierarchies might not provide adequate security compartmentalization, extensibility and flexibility. Hierarchies with more than three tiers might not provide additional value regarding security, extensibility and flexibility.

The most important consideration is protecting the highest instance of trust as much as possible. Single-tier hierarchies are based on the need to compartmentalize risk and reduce the attack surface that is available to users who have malicious intent. A larger hierarchy is much more difficult to administer, with little security benefit.

Depending on the organization's necessities, a PKI should consist of two or three logical levels that link several CAs in a hierarchy. Administrators who understand the design requirements for a three-level topology may also be able to build a two-level topology.

A three-tier CA hierarchy consists of the following components:

- A root CA that is configured as a stand-alone CA without a network connection
- One or more intermediate CAs that are configured as stand-alone CAs without a network connection
- One or more issuing CAs that are configured as enterprise CAs that are connected to the network




**Figure 1: Three-tier CA Hierarchy**

To set up a two-tier topology, apply all of the steps that are described in Example scenario for Contoso Company, later in this document.

If the organization can fulfill its security requirements with a two-tier hierarchy, a three-tier architecture is not required. When you do not have a middle tier, CA management applies to two levels instead of three levels and might lower maintenance cost.

To implement a two-tier topology, use the steps that are outlined in both the "Stand-alone offline root CA" and "Online Enterprise Issuing CAs (CorporateEnt1CA)" sections of this paper.

From a technical perspective, a single level PKI hierarchy can also provide basic PKI services. Leaving out the root and the intermediate tier results in an all-in-one CA. Because the single CA must issue certificates, it cannot be taken offline. Security and flexibility is very limited with this type of implementation. To implement a single-tier topology, apply the steps that are outlined in "Online Enterprise Issuing CAs (CorporateEnt1CA)."

The decision of whether or not to use a separate root CA to issue all certificates in an organization should be a need for security versus a need for cost mitigation and simple administration.

To summarize, a two-tier to four-tier CA topology is the most common deployment. Any organization should be able to deploy a similar PKI architecture to meet any organizational, business, and technical requirement, as well as a respectable level of security.

For reliability and redundancy, improve the availability of the PKI and deploy multiple enterprise CAs instead of extending the depth of the hierarchy.

**Root CA**

A root CA is a self-signed CA. Technically, the root CA runs the same code as an intermediate or issuing CA. The difference between these types of CAs is in the role that the CA takes. The following table displays a list of the characteristics that a root CA should have, depending on the CA topology.

**Table 9 Root CA Characteristics**

| Characteristic | More than 2 tiers | 2 tiers | 1 tier |
|---|---|---|---|
| High level of physical security | Yes | Yes | No |
| Permanently offline | Yes | Yes | No |
| Highly restricted area (vault) | Yes | Yes | Yes |
| Match the level of risk | Yes | Yes | Yes |
| High level of cryptographic security | Yes | Yes | Yes |

| | | | |
|---|---|---|---|
| Largest key size | Yes | Yes | Yes |
| Software CSP<br><br>(FIPS 140-1 level 1) | No | No | Yes |
| Smart cards or PCMCIA tokens with PINs<br><br>(FIPS 140-1 level 2) | No | Yes | Yes, recommended |
| Hardware security modules with operator hardware token<br><br>(FIPS 140-1 level 3 or 4) | Yes | Yes, recommended | Yes, recommended |

Even if an offline root CA might run only when the CA certificate must be renewed or the CRL has to be published, the CA must be installed on reliable hardware. If you are thinking about using a notebook computer to take the role of the root CA, note that it does not meet the requirements for reliability at the time that this document is being published.

In most customer environments, maintaining a root CA requires extraordinary security measures. The level of security requires that the root CA is offline at all times and, preferably, protected in a secure physical environment. In theory, a desktop system with a removable hard drive can be used to protect a root CA.

**Intermediate CAs**

Intermediate CAs are subordinate to the root CA. By definition, if you implement an intermediate CA, the topology consists of a minimum of three tiers. The intermediate layer of a PKI hierarchy often provides useful policy, administrative or operational differentiation. Intermediate CAs are also known as policy CAs because they are often used to manage or dictate different security and operational policies between different geographical areas, business units, or the intranet or extranet for a corporation.

To implement policies without an intermediate CA, you can also assign policies to issuing CAs on a logical basis.

An intermediate CA's security requirements are the same as for the root CA because an intermediate CA provides CA certificates to online issuing CAs. The intermediate CA should be an offline, stand-alone CA.

| **Note** |
|---|
| It is highly recommended that you only issue certificates from an intermediate CA after the administrator manually approves the request. This is the default configuration for a Windows Server 2003 stand-alone CA. |

**Issuing CAs**

Depending on the architecture, an issuing CA is a subordinate of an intermediate CA or a subordinate of the root CA. Enterprise CAs are ideal for issuing large numbers of certificates, because they can automatically validate the user and certificate profile information. The purpose of an issuing CA is to enroll certificates to end-entities and not to subordinate CAs.

| **Note** |
|---|
| You can limit the number of subordinate CA levels in a certificate hierarchy by defining a maximum path length in the basic constraints extension of a CA certificate. A path length of zero will ensure that an issuing CA may only issue certificates to end-enties. You can define the basic constraints extension and path length by using a CApolicy.inf configuration file. |

**Understanding Root Trust**

When a client uses a certificate, it is mandatory that the trust relationship between the certificate and the root CA can be verified. A certificate is trusted if the client that verifies the certificate trusts the root CA certificate that is in the client certificates certificate trust path as well. A client must have the related root CA certificate in its local certificate store to proof a trust-relationship with the root CA. For more information, see "Policies to establish trust of root certification authorities" on the Microsoft Web site.

If Active Directory is available, it is important to understand how clients like users or computers can benefit from Active Directory to establish a trust relationship with the root CA

You can achieve the trust that is obtained from a root certification authority by deploying the root CA certificate through one of the following six methods:

- Enterprise trust in Active Directory

- Group policy in Active Directory
- Certificate Trust Lists (CTLs) in Group Policy
- Manual trust on a local computer
- Manual trust by a user
- Windows Update

Depending on the permissions and the scope of the distribution mechanism, certificates are put into different locations and require different maintenance tools. For more information, see the following table.

**Table 10 Certificate Trust Mechanisms**

| Distribution method | Scope | Uses Group Policy object | Location | Maintained with |
|---|---|---|---|---|
| Enterprise trust | Entire forest | Yes | Services\Public Key Services\CertificationAuthorities | **Certutil.exe** or PKI Health Tool |
| Group policy trust | Domain | Yes | Domain Security Group Policy object | Group Policy MMC |
| NTAuth (for CAs trusted to issue authentication certificates) | Entire forest | Yes | Services\Public Key Services\NTAuth object | **Certutil.exe** or PKI Health Tool |
| Manual trust on the local computer | Local computer and all users that log on to system | No | Registry HKEY_LOCAL_MACHINE | Certificates MMC for the local computer |
| Manual trust by user | Current user | No | Registry HKEY_CURRENT_USER | Certificates MMC for the local computer |
| Windows Update | Local computer and all users that log on to system | No | Registry HKEY_LOCAL_MACHINE | Group Policy MMC or **Add or Remove Programs** in Control Panel |

**Enterprise Trust**

You can use the built-in autoenrollment service to automatically download root CA certificates and certificate trust lists (CTLs) from the Active Directory enterprise trust store on both Windows 2000 and Windows XP clients.

For additional information, see the following articles on the Microsoft Web site:

- Certificate Autoenrollment in Windows Server 2003 on the Microsoft TechNet Web site
- Configure Public Key Group Policy on the Microsoft Web site

**Group Policy Trust**

Group Policy trust is defined and configured by using the Group Policy MMC and the Default Domain Security Group Policy object. Group Policy trust is configured and enforced for the domain where the Group Policy object applies. Because of this, different users in different domains trust different root CAs. It is highly recommended to create a new domain policy and not edit the default domain policies.

| **Note** |
|---|
| Only root CA certificates must be trusted and registered on client computers. Do not add subordinate CA certificates to the Group Policy trust, because intermediate and issuing CAs certificates may not be explicitly trusted. CryptoAPI automatically |

**NTAuth**

The NTAuth store is deployed on all computers in the forest from the configuration partition of the forest in the following directory path:

**CN=NTAuthCertificates,CN=Public Key Services,CN=Services,CN=Configuration,DC=**

| **Important** |
| --- |
| NTAuth CAs are trusted to both issue authentication (logon) certificates for any user in the forest and enable logon for smart cards, Internet Information Services (IIS) mapping, and Extensible Authentication Protocol-Transport Layer Security (EAP-TLS).Precise control of issuing CAs can be achieved through qualified subordination with constraints. |

You can verify the certificates that are currently registered in NTAuth by typing the following at a command prompt where the domain component information is configured with the name of the Active Directory root domain:

**certutil.exe -store ldap:///CN=NTAuthCertificates,CN=Public Key Services,CN=Services,CN=Configuration,DC=contoso,DC=com**

You can see a more visual display of certificates by typing the following at a command prompt:

**certutil -viewstore "ldap:///CN=NTAuthCertificates,CN=Public Key Services,CN=Services,CN=Configuration,DC="**

You can also manually maintain the NTAuth store by typing one of the following commands at a command prompt:

**certutil –addstore**

**certutil –delstore**

**certutil –dspublish***CertificateFile***NTAuth**

For more information about the NTAuth store and smart card logon, see the following articles on the Microsoft Web sites:

- "Step-by-Step Guide to Mapping Certificates to User Accounts" on the Microsoft TechNet Web site
- "How to Import a Third-Party Certificate into the NTAuth Store" on the Microsoft Knowledge Base
- "Enabling Smart Card Logon with Third-Party CAs" on the Microsoft Knowledge Base
- "Requirements for Third-Party CA Domain Controller Certificates" on the Microsoft Knowledge Base

In a Windows Server 2003 Active Directory environment that contains only clients running Windows XP, the NTAuth store is not mandatory for smart card logon and certificate mapping, compared to a Windows Server 2003 mixed environment with Windows 2000 clients. Because Windows Server 2003 Active Directory supports publishing cross-certificates and because clients running Windows XP support name and policy constraints for x.509 certificates, administrators may waive the NTAuth policy in homogenous Windows Server 2003 and Windows XP environments. This option requires and assumes that CAs have defined name constraints instead of being listed in the NTAuth store of the directory. Therefore, domain controllers that process both smart card logon and certificate mapping requests will explicitly trust all CAs that chain to trusted root CAs, assuming that the certificate matches a valid user account in Active Directory.

| **Caution** |
| --- |
| Disabling NTAuth policy verification enables domain controller trust of any CA that issues a valid smart card logon certificate and chains to a trusted root CA in the Active Directory environment. Any CAs, including the default third-party root CAs, should have name constraints defined before disabling the NTAuth policy. If this does not occur, unintended trust and logon access may occur. Use this option with extreme caution and only when root CAs have been properly constrained in the environment. |

For more information about qualified subordination and name constraints, see "Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003" on the Microsoft TechNet Web site.

**Manual Trust on a Local Computer**

Root CA certificates may also be manually trusted on a local computer by using the Certificates MMC snap-in for the local computer. The user must be a local administrator to add root CA certificates to the machine certificate store. All root CA

certificates in the computer's machine certificate store are inherited by all users who log on to that computer. The users trusted root certificate store and the machine trusted root certificate store form a union from a user's perspective.

For more information about certificate stores, see Chapter 13, Public Key Technology on the Windows 2000 Resource Kit Web site.

Confirm that certificates are stored in the correct location. Any root CA certificate that is stored in the local computers certificate store is visible to any user on that computer. If a root CA certificate is registered in the local computer store and if the CA certificate is also manually added by a user, the root CA certificate might appear twice in the Certificates MMC snap-in. If a root certificate is not available in the local computers certificate store but is available in the users store, building a certificate chain also may not work for some applications.

You can maintain the computers certificate store also with the Internet Explorer Administration Kit (IEAK) or CAPICOM. For more information about CAPICOM, see the article "CAPICOM Reference" on the MSDN Web site.

**Manual Trust by User**
It is recommended that only administrators maintain certificate trust and that you store only CA certificates in the local computers certificate store.

**Windows Update**
By default, computers that are running Windows XP and members of the Windows Server 2003 family run a service that will download updated public root CA certificates that have been added to the Microsoft root program. The service is not available in the Windows 2000 family.

Any organization that has a CA that meets the requirements that are outlined in the Microsoft Root Certificate program is able to distribute the CA certificate through Windows Update. For more information, see "Microsoft Root Certificate Program" on the Microsoft TechNet Web site.

Computers that are running either Windows XP or Windows Server 2003 periodically download the current list of Root CA certificates that are added to the Third-Party Root Certification Authority store on the local computer. For more information, see the chapter "Certificate support and the Update Root Certificates component" in "Using Windows XP Professional with Service Pack 1 in a Managed Environment: Controlling Communication with the Internet," which is available for download on the Microsoft Web site.

To install or remove this service, you can use **Add or Remove Programs**. To do this, click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**. In the toolbar, click **Add/Remove Windows Components**, and then, in **Components**, select the **Update Root Certificates** check box.

This service can also be managed through Group Policy in Active Directory.

**Figure 2: Windows Components Wizard**




**Creating an Enrollment Strategy**
A certificate can be used to provide authentication evidence of its owner, encrypt, or sign data. Because of this, the certificate issuer must ensure that certificate holders are known entities. Before certificates are enrolled, you should answer the following questions:

- How will users obtain their certificates?
- What is the process for enrollment identification?

The most secure way to initially enroll user certificates is to do a face-to-face authentication at the registration authority and store user certificates on hardware tokens. This provides the highest level of assurance, but also the highest cost of deployment.

If certificate enrollment with hardware tokens through enrollment agents is not an option, the CA can verify the certificate requester with domain credentials. This authentication method for certificate enrollment is usual when users self-enroll certificates. This scenario assumes that a user is the only person who is able to use the credentials.

It is recommended that you use a combined enrollment strategy that implements a strong initial identity check. Subsequent certificate enrollment and renewal can then be based on the initial certificate.

**Creating a CA Renewal Strategy**
Certificate lifetimes can have an impact on the security of your PKI for the following reasons:

- Over time, encryption keys become more vulnerable to attack. In general, the longer amount of time that a key pair is in use, the greater the risk that the key can be compromised. To mitigate this risk, you must establish the maximum allowable key lifetimes and renew certificates with new key pairs before these limits are exceeded.
- When a CA certificate expires, all subordinate certificates that are issued by this CA for validation also expire. This is known as time nesting and is traditionally enforced by CryptoAPI in the client.
- When a CA certificate is revoked, all certificates that have been issued by the CA must also be re-issued.

End entity certificates expire when the issuing CA certificate reaches the end of its lifetime, unless:

- The end entity certificate is renewed with a new key pair that chains to a CA certificate with a longer lifetime.
- The end entity certificate was revoked before the CA certificate expiration date is reached.

You must plan the CA certificate renewal precisely during the PKI deployment phase. If this important planning step is missed, the entire PKI might stop working when the CA certificate expires, because all of the certificates that depend on the CAs certificate are then no longer usable for both encryption and signing operations. Remember that a certificate is capable of decrypting data, even if it has expired or been revoked.

| **Note** |
| --- |
| It is strongly recommended that you generate new key material when you renew a CA's certificate in order to partition the CRL that is issued by the CA and also prevent ambiguous certificate chaining errors caused by use of the same public key. |

**Determining the Total Number of CAs**
The total number of CAs depends on the organizations security requirements and the organization's size. It is also dependent upon the geographical, political, and business hierarchy of the organization. As outlined earlier in this document, there is a choice of different trust levels that may be applied. After the organization has decided how many tiers should be implemented, it is important to plan the number of CAs that are required at each level. For a PKI topology that uses intermediate CAs, the number of CAs depends on the number of different CA policies that are required to issue CAs. The number of issuing CAs depends on the number of certificates that should be issued, the network connectivity between the requester and the CA, and the number of intermediate CAs.

A three-tier architecture consists of:

- One root-CA
- At least one policy CA (This can be one or many servers.)
- At least two issuing CAs for every policy CA to ensure fault tolerance

A two-tier architecture consists of:

- One root-CA
- At least two issuing CAs to ensure fault tolerance
- A single-tier architecture consists only of a single CA.

Note the following:

- You cannot change the CA type at a later time; you must uninstall the original CA and then reinstall the CA to change it from either a stand-alone CA to an enterprise CA or an enterprise CA to a stand-alone CA.
- You can install only one instance of a CA on a Windows Server 2003 system.
- The certificate distribution point and the CRL publication interval is valid for all certificates that are issued by a CA and cannot be set for individual certificates.

For certificates that are used externally, the naming and information that is part of the certificates should not reveal the internal PKI or network infrastructure, such as the name of a CA or CRL distribution point paths in the issued certificates.

**Hardware Requirements**
This section provides some general guidelines for hardware requirements for a Windows Server 2003 CA. This section should not be used as an authoritative guide for performance characteristics. Specific performance characteristics vary, depending on the implementation and customer environment.

**Hardware Guidelines**
Microsoft performance testing in a lab environment has shown that the signing key length of the CA has the most significant impact on the enrollment rate of the CA. A larger number of certificates can be signed and enrolled in a given time if a smaller key size is used. If a larger key size is used, more CPU time is required to issue certificates.

The total number of issued certificates should not have a significant influence on either server performance or the rate at which the CA issues certificates; the performance of the issuing CA stays nearly the same, whether thousands or millions of certificates have previously been issued. Therefore, the scalability of the CA is considered to be linear, based on the size and performance of the disk arrays that are used to store both the database and log files.

The following table lists configuration factors that may affect performance of the CA.

**Table 11 Resources That Affect CA Performance**

| Resource | Performance notes |
|---|---|
| Number of CPUs | Additional CPUs increase the overall performance of the CA. This is the most critical resource for a Windows Server 2003 CA. |
| Memory | In general, additional memory does not have a significant role in the enrollment performance of the CA. The CA should meet general recommended system requirements (512 MB), however, the minimum amount of memory is 256 MB. |
| Disk size | The capacity of the disk volume that stores the database and log files it the primary limiting factor for the number of certificates that a CA is able to maintain. |
| Disk performance | In general, a short key length (512 KB) generates very little CPU utilization and a very high disk load. Larger key sizes generate more CPU utilization and less disk usage. A high-performance disk subsystem can increase the rate of enrolled certificates. A RAID set is recommended for both performance and fault-tolerance purposes. CA operations are primarily disk-write intensive. |

| | |
|---|---|
| Number of volumes | Using separate disks for the database and log files provides basic performance improvement. In general, the drive that contains the CA database is used more than the drive that contains the log files. The disk write capacity improves if you use more physical drives in a RAID set. |
| RAID stripe size | It is recommended that you use a stripe size larger than 64 KB. |
| Key length | The larger the signature key length, the greater the CPU utilization. Larger keys degrade CA performance. To be CPU-independent, you may want to use hardware acceleration to provide a large number of both key generation and signing operations. |
| Bandwidth | A 100 megabit network connection is suitable to enroll a large number of certificates and causes no performance bottleneck, assuming that the server is running the CA exclusively with no additional applications or network services. |

**Processor Notes**

In general, a computer that has a current processor and 512 MB of memory is considered sufficient for most organizational uses of a Windows Server 2003 CA. The enrollment rate is directly related to the ability of the CA to sign requests that are based on CPU availability. Many hardware, environmental, network, or client factors can affect the performance of a CA.

**Disk Configuration Notes**

Disk space and disk speed also limit the performance and scalability of a CA. Each certificate that is issued uses approximately 16 KB of disk space in the database, and an additional 4 KB is required if the private key is archived. The certificate database must contain all of the issued certificates to be able to revoke certificates and provide a record of operations. Because none of the records are ever automatically tombstoned or automatically deleted, the certificate database continuously increases in size when new certificates are issued. Nevertheless a CA administrator can use the Certutil.exe command-line utility to delete expired records from the CA database.

**Scalability**

The Windows 2000 CA has been tested to issue more than 7 million certificates and the Windows Server 2003 CA has been tested to issue more than 35 million certificates on a single four-processor, x86-based computer. The maximum database size was not reached in either of the test scenarios.

# Creating Certificate Policies and Certificate Practice Statements

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The definition of certificate policies and the certificate practice statement (CPS) is often forgotten by technically-oriented planners. The basis for both the certificate policy and CPS are the organizations security policy.

Creating these documents is usually a joint responsibility of the legal, human resources, and information security departments.




**Figure 3: Relationship Between Certificate Policy and Certificate Practice Statements**

Both the certificate policy and the CPS help the user of a PKI determine the level of trust that those departments can put in the certificates that are issued by a CA. The existence of policies is critical when dealing with a reliable PKI. If certificates are exchanged only within an organization, the creation of a CPS and a security policy might not be mandatory. When this is true, some clauses regarding the use of PKI and certificates in the employee manual may be essential. Note that an organizational CPS is a CPS that covers all CAs in a hierarchy. Generally, a CPS covers only a specific CA.

A CP and CPS may prove to required when certificate holders exchange or use certificates with partners and entities that live outside of the company's network. When external trust is implemented, it is often very important to align PKI policies and practices as part of the external contract terms.

**Security Policy**

The security policy is a high-level document that is created by the corporate IT group. It defines a set of rules about the use and provision of security services in the organization, and should reflect your organization's business and IT strategy. The security policy should answer high-level PKI questions, such as:

- What applications should be secured with certificates?
- What kind of security services should be offered by using certificates?

**Certificate Policy**

A certificate policy focuses on certificates and the CA's responsibilities regarding these certificates. It defines certificate characteristics such as usage, enrollment and issuance procedures, as well as liability issues.

The following references define a certificate policy as a set of rules that determine if a certificate is applicable to either a community or a class of applications that have common security requirements.

For more information about the X.509 standard, see the International Telecommunication Union Web site.

For more information about the European Electronic Signature Standardization Initiative (EESSI) definition, see the EESSI Web site.

A certificate policy typically answers the question about what purposes the certificate serves, and under which policies and procedures the certificate has been issued. A certificate policy typically addresses the following issues:

- How users are authenticated during certificate enrollment
- Legal issues, such as liability, that might arise if the CA becomes compromised or is used for something other than its intended purpose
- The intended purpose of the certificate
- Private key management requirements, such as storage on smart cards or other hardware devices

- Whether the private key can be exported or archived
- Requirements for users of the certificates, including what users must do if their private keys are lost or compromised
- Requirements for certificate enrollment and renewal
- Minimum length for the public key and private key pairs

The certificate policy is typically defined by members of an organization who are known as the *policy authority*. The policy authority typically consists of representatives from different core departments, including management, legal, audit, human resources, and other departments. Overall, the policy authority members will also be members of the group that defined the security policy, which ensures that the certificate policy is in agreement with the security policy.

**Certificate Practice Statement**

The certificate practice statement (CPS) translates certificate policies into operational procedures on the CA level. The certificate policy focuses on a certificate; the CPS focuses on a CA. Both the EESSI and the American Bar Association (ABA) define a CPS as a statement about the way that a CA issues certificates. For more information about the ABA, see the ABA Web site. For more information about the EESSI, see the EESSI Web site.

A CPS might include the following types of information:

- Positive identification of the CA, including the CA name, server name, and Domain Name System (DNS) address
- Certificate policies that are implemented by the CA and the certificate types that are issued
- Policies, procedures, and processes for issuing, renewing, and recovering certificates
- Cryptographic algorithms, cryptographic service providers (CSPs), and the key length that is used for the CA certificate
- Physical, network, and procedural security for the CA
- The certificate lifetime of each certificate that is issued by the CA
- Policies for revoking certificates, including conditions for certificate revocation, such as employee termination and misuse of security privileges
- Policies for certificate revocation lists (CRLs), including where to locate CRL distribution points and how often CRLs are published
- A policy for renewing the CA's certificate before it expires

The CPS should be defined by a team that consists of members of the IT department, people who are operating and administering the IT infrastructure, and the people (often attorneys) that defined the certificate policy. The CPS is a public document that should be published on the Internet. Every certificate that has been issued by a CA that follows a CPS has an URL pointer in the certificate that directs people to the public document. When a certificate has a CPS pointer as part of the certificate, the **Issuer Statement** button becomes available. When you click **Issuer Statement**, the URL that has been specified by the CA administrator is redirected.

| **Important** |
| --- |
| A CPS is always valid for all certificates that are issued subordinate to the CA that contains the qualifier. Make sure that all parameters that are listed in Appendix B are part of the planning process. |

**Revocation Policy**

Before certificates are enrolled, the PKI management team should know how to revoke certificates. Any X.509 V3 certificate (except the root CA certificate itself) should have a pointer to a valid CRL. The CRL distribution point is included in the

certificate's extension and cannot be modified after a certificate is enrolled.

The logical availability of the CRL distribution point that is specified in the certificate allows a PKI-enabled application to verify the certificate's validity against the CRL. The CRL is essential to ensure the quality (status) of certificates that are published by the CA. If the CRL is available and the certificate's serial number is part of the CRL, the certificate is marked as invalid from a clients perspective.

A revoked certificate's serial number is added to the CRL as long as the original certificate lifetime is valid. After the original lifetime of the certificate expires, the serial number of the certificate is added to the CRL for the last time.

| **Note** |
| --- |
| You cannot use revoked certificates for signing or encryption operations anymore. However, you can use revoked certificates for decryption operations, because the revoked certificates are required for decryption. |

If an application is going to verify a certificate against the CRL and no valid CRL is available, the revocation check does not work and the certificate cannot be used for the transaction. If the application has properly implemented CRL checking, no authentication, encryption, or signing is allowed with this certificate until a valid CRL is available again.

For immediate revocation of logon certificates, consider disabling the account in Active Directory instead of revoking the logon certificate. It is more time efficient to delete or disable user accounts if you want to immediately revoke a user's ability to gain access to the logon certificates.

For more information about CRLs, AIA, and chain building refer to the "Troubleshooting Certificate Status and Revocation" white paper on TechNet.

**CRL Best Practices**

When you consider CRL distribution, you should know where and how clients can gain access to a CRL. The CRL distribution mechanism of enterprise and stand-alone CAs is different by default.

It is a common mistake to not modify the default CRL distribution point of an isolated stand-alone CA. Because a root or intermediate CA is typically disconnected from the network, PKI-enabled clients cannot validate the issued certificates against the default CRL distribution point on the CA server. To make a CRL of an offline stand-alone CA publicly available, you must manually publish the CRL or utilize a custom exitmodule or script that publishes the CRL to a predefined location. For more information about custom exit modules, see “Writing Custom Exit Modules” in the Cryptography collection of the Microsoft Platform SDK on MSDN.

An online CA on a computer that is joined to an Active Directory domain or forest automatically publishes the CRL to Active Directory so that it can be accessible through LDAP. Alternatively, the CRL can be made available through an HTTP URL that points to a location on a Web server.

Depending on the certificate types that are issued with a CA, the order of the CRL distribution points is important. For authentication certificates, it is beneficial to have a CRL or fully-qualified LDAP CRL distribution point as the first entry in the list of distribution points. If a relative LDAP CRL distribution point is specified, a client contacts the domain controller that is closest, according to the Active Directory site structure, to get the CRL. Fully-qualified LDAP CRL distribution points eliminate latency issue that may occur until the CRL has been replicated in Active Directory. For non-authentication certificates, you may want to use LDAP because LDAP is more fault-tolerant in an Active Directory environment compared to tolerance in a single-instance HTTP server.

It is also an option to set both an LDAP and HTTP CRL distribution point URL to support clients that are Active Directory-aware, as well as clients that are not running Windows and that are not Active Directory-aware. If you have a mixed client environment or both internal and external clients, it is a best practice to place the HTTP location in the CRL distribution point extension first to avoid network timeouts. Any client that retrieves a CRL on demand during certificate verification caches a copy of the CRL in the Internet Explorer temporary files until the CRL expires.

| **Note** |
| --- |
| It is a best practice to publish a CRL that is available externally through an HTTP location so that users and applications that are outside of the organization may perform certificate validation. It is also a best practice to use paths and naming that do not reveal the internal network infrastructure to external entities. |

The CRL maintains a list of revoked certificates that have been issued by a CA. The CRL does not maintain the validity of certificates that are owned by a subordinate CA, like the CA certificate. The subordinate CA certificates revocation status is maintained by the CAs parent CA, since the parent CA has issued the subordinate CA certificate.

Since a root CA certificate has no parent CA that could maintain the CRL, there is no need to specify a CRL distribution point for the root CA certificate itself. To revoke a root CA, all certificates that have been issued by the root CA must be revoked instead.

Here are some additional planning notes:

- A root CA certificate should have an empty CRL distribution point because the CRL distribution point is defined by the certificate issuer. Since the roots certificate issuer is the root CA, there is no value in including a CRL distribution point for the root CA. In addition, some applications may detect an invalid certificate chain if the root certificate has a CRL distribution point extension set.
- Offline CAs must continue to publish CRLs.
- If certificates are exchanged with external entities, the CRLs must be available at a location that is accessible for all internal and external entities. To satisfy this requirement, in this case the CRL is usually published in the organization's perimeter network.
- To ensure redundancy, make the CRL available through more than one location.
- If the CRL is distributed by using Active Directory, plan for replication latency.
- Plan for CRL publications that cannot be performed as usually scheduled and have contingency operations prepared
- The CRL should be valid for the amount of time that it takes for CA recovery if hardware fails or if software does not work. For example, a one-hour CRL publication period is most likely not adequate time to perform a hardware or software restoration because of the possibility of issues with either the hardware or software.
- The less frequently CRLs are published, the more time you will have for issue resolution.
- If a CA cannot publish the CRL on time, the CRL is not updated and will expire. If a CRL has expired, clients cannot verify certificates that would require the CRL. To prevent certificate misuse, certificates are considered invalid if the CRL has expired or is unavailable.
- A Windows Server 2003 Certification Authority can publish to an IIS cluster using UNC paths for the CRL distribution point URL.

All Windows CAs follow the CRL V2 format that is specified in RFC 2459 and RFC 3280. For additional information about RFC 2459 and RFC 3280, see the Internet Engineering Task Force Web site.

**LDAP CRL Best Practices**

The following best practices apply to LDAP-based CRLs:

- An LDAP CRL is replicated in Active Directory to all domain controllers in the forest. Because of this, it provides fault-tolerance in an environment with more than one domain controller and can be designated as "highly available."
- The Active Directory replication schedule should be taken into account. This is an important consideration since it may take longer than expected for every directory server to receive the latest version of the CRL, depending on the size and replication schedule of the Active Directory environment.
- CRLs should not be published to Active Directory when the CRL publication period is shorter than the replication convergence time for the Active Directory forest.
- Do not use escape characters in LDAP CRL distribution point paths, such as a backslash (\).
- Do not include names that are specific to either internal or organization names in the CRL distribution point. Certificates may be exchanged with external parties and those parties should not be able to obtain information about internal name spaces. To eliminate internal names in the CRL distribution point, also allow internal clients to gain access to the external distribution point, or implement a name mapping mechanism that ensures that internal clients can resolve an external name and gain access to an internal resource.

- If an LDAP CRL distribution point for certificates that are exchanged with external parties is used, do not use the relative LDAP URL that points to the closest domain controller.

The first part of the name is the hosting and distribution point is the LDAP name of the server that is hosting the CRL distribution point; the second part of the name is the complete LDAP path of the directory location where the CRL is stored. The following configuration string (or LDAP reference):

**ldap:///CN=,CN=**

is interpreted as

**ldap:// *ClosestDomainControllerBySite* /CN=,CN=**

When you use this syntax, part one of the LDAP CRL distribution point is left out, but it is automatically inserted when the CRL must be retrieved. The ldap:/// syntax forces a Windows 2000, Windows XP, or Windows Server 2003 client that is joined to an Active Directory domain to find the closest domain controller. Alternatively, a fully qualified domain name (FQDN) or an exact server path with a port value, is also supported.

Not only is the path of the CRL is important when you plan an LDAP CRL distribution point; you must also configure the correct search criteria and append that search criteria to the LDAP path.

LDAP searches support search suffixes to specify attributes, depth, and object-classes. Because the directory service may store several objects of different data types in the same location, it is important to query for the correct data. A search suffix uses the following format:

?attribute?depth?object_class

If the attribute, depth, and object_class search suffixes are missing, the client selects the correct object. Because there are different client implementations, a CRL verification might not work without this extended information.

The following example shows a relative LDAP CRL distribution point (on one line):

ldap:///CN=*CorporateRootCA*,CN=*Root_CA*,CN=CDP,CN=Public Key Services,CN=Services,CN=Configuration,DC=*contoso*,DC=*com*?certificateRevocationList?base?objectClass=cRLDistributionPoint

In the sample line above, CorporateRoot, CA_*Root_CA*, *contoso*, and *com* are placeholders and must be replaced with parameters that are specific to the organizations requirements.

The following example shows an absolute LDAP CRL distribution point (on one line):

ldap://cdp1.contoso.com/CN=*CorporateRootCA*,CN=*Root_CA*,CN=CDP,CN=Public Key Services,CN=Services,CN=Configuration,DC=*contoso*,DC=*com*?certificateRevocationList?base?objectClass=cRLDistributionPoint

In the sample line above, *cdp1.contoso.com*, *CorporateRootCA*, *RootCA*, *contoso*, and *com* are placeholders and must be replaced with parameters that are specific to the organization's requirements.

**HTTP CRL Distribution Point URL Best Practices**

Non-Windows clients might not be able to retrieve CRLs with LDAP URLs based on Active Directory. Because of this, you may need to provide an additional HTTP CRL distribution point location for LDAP-enabled clients. Computers running Windows support LDAP and HTTP URLs. The following are some best practices for HTTP CRL distribution point URLs:

- If you provide an HTTP CRL distribution point location, provide fault tolerance by having either a virtual server name that points to several physical Web servers (round-robin DNS) or a clustered Web server to provide redundancy in the HTTP URL.

- HTTP CRL distribution point locations are ideal for providing accessible CRL locations for clients that are not running the Windows operating system.

- Place HTTP CRL distribution point URLs second in the list of the URLs in the CRL distribution point extension when Active Directory-aware clients are primarily used. This is to decrease network traffic because the client would benefit from intra-site communication with the domain controller.

- Place HTTP CRL distribution point URLs first in the list of the URLs in the CRL distribution point extension when clients cannot connect to the Active Directory to verify certificates. Examples include external Web servers, VPN and remote access servers, and RADIUS (IAS) servers.

- HTTP URLs should contain only valid file name characters.

**Delta CRLs**
In a production environment, the number of certificates that are revoked is in relation to the number of certificates that are issued. The list of revoked certificates will vary in length, depending on the number of certificates that are enrolled by a CA.

Revoked certificates are added to the CRL as a collection of certificate serial numbers. RFC 2459 and RFC 3280 define a method that you can use to reduce base CRL sizes by using delta CRLs. Delta CRLs maintain a list of certificates that have been revoked since the last base CRL publication.

Base CRLs and delta CRLs are cached by Windows clients. To ensure the validity of a certificate, the client uses the locally-cached base and delta CRL until the CRLs validity period expires. If a base CRL expires, the client retrieves a new base CRL from the distribution point that is specified in the certificate. If the base CRL is valid but the cached delta CRL is expired, a Windows client retrieves only the delta CRL. Typically, a delta CRL is much smaller in size than a base CRL because it saves only the certificates that have been revoked after the last base CRL update.

Delta CRL best practices:

- Use delta CRLs with issuing CAs whenever possible.

- Do not use delta CRLs with offline CAs, because there are not as many certificates that require frequent revocation. Offline CAs usually have longer CRL publication cycles than issuing CAs, since it is abnormal to revoke a large number of CA certificates.

To provide clients with the most up-to-date revocation information with smaller network utilization (compared to the network utilization that is required for a base CRL distribution), you can publish the delta CRL on a daily basis and publish the base CRL on a weekly basis. However, if a large number of certificates is revoked and if the number of revoked certificates exceeds the number of revoked certificates that are already part of the base CRL, the size of a delta CRL is larger than the size of a base CRL. Note that this scenario is very unlikely to occur and is not considered to be typical.

Do not publish frequent delta CRLs to Active Directory if replication takes a longer period than the delta CRL is valid.

**Online Certificate Status Protocol Support**
A Windows CA does not provide online certificate status protocol (OCSP) functionality by default. However, you can enable OCSP if you install a revocation provider in CryptoAPI or through a third-party OCSP responder that communicates with the Microsoft Certification Authority. For more information about CryptoAPI or revocation providers, search for CryptoAPI on the MSDN Web site.

**Best Practices for CRL Publication**
The CRL publication interval for CAs that issue certificates to CAs should be a longer period than for CAs that issue certificates to end-entities, because revocation of a CA certificate should be a very rare operation. A recommended creation interval for a new CRL of that type would be in the range of 90 to 180 days.

A CRL for an offline CA should always be published a few days before expiration to allow for unexpected issues. The publication interval for issuing CAs should be set according to the type to issued certificates. Authentication certificates might require a less frequent publication schedule than other certificate types.

For offline CA CRL publication, you should also consider these points:

- For issuing CAs, a short CRL publishing schedule ensures that the CRL is current and that any revocation can be made available as quickly as possible. Note that Windows clients cache a CRL for the validity period.

- For offline CAs, a longer CRL publishing schedule ensures that the CRL does not have to be regenerated and republished through the required manual generation and publishing processes.

**AIA Extensions**

The AIA extension allows the certificate user to obtain a current copy of the CAs current certificate. CA certificates are required when a certificate chain is built. Chain building is performed as part of the certificate verification process.

When you configure AIA extensions, use the same attention to detail that you use when you configure CRL distribution point extensions. See the following example for more information about the AIA extension in a certificate.




**Figure 4: AIA Extensions**

CA certificates are multivalued, Base64-encoded attributes in Active Directory that can store more than one CA certificate. A multivalued AIA attribute is used for every CA because a CA may have more than one valid certificate after CA renewal.

Note that multivalued attributes are limited in the number of values that they can store. You cannot sure more than 1,000 CA certificates in the AIA object.

Compared to an LDAP AIA URL that points to a multivalued object and distinguishes certificates in the same object by the search suffix, an HTTP AIA URL points only to a single file. Because of this, all HTTP URLs must include the certificate suffix (*.cer, *.crt, etc) as the suffix for the file name to distinguish between the multiple certificates that are stored in the same directory on the HTTP server

The following table describes the structure of how certificates are stored. Note that an HTTP location must be unique and only one CRL object (or file) may exist at each explicit URL path. An LDAP location is a single object in Active Directory that supports a multivalued attribute.

**Table 12 Example of Stored Certificate Structure**

| | **AIA HTTP URL (object)** | **AIA LDAP URL (object)** | |
|---|---|---|---|
| Several files at http://www.microsoft.com/ | concorp-ca-00_ CorporateRootCA.crt<br><br>concorp-ca-00_ CorporateRootCA(1).crt<br><br>concorp-ca-00_ CorporateRootCA(2).crt | concorp-ca-00_ CorporateRootCA.crt<br><br>concorp-ca-00_ CorporateRootCA(2).crt<br><br>concorp-ca-00_ CorporateRootCA(2).crt | One multivalued attribute at CN=AIA,CN=Public Key Services,CN= Services,%6%11 |

**Certificate Validity Period and Key Length**

The validity period of certificates depends on the organization's requirements. The following table outlines some recommendations for the validity period for different CA types.

**Table 13 Recommendations for Validity Periods**

| Purpose of Certificate | Certificate Life | Private Key Renewal Strategy |
|---|---|---|
| Stand-alone root CA<br><br>(4096-bit key) | 20 years | Renew at least every 10 years to ensure that policy CA certificates can be issued with lifetimes of 10 years. Renew by using a new key at least every 20 years. |
| Stand-alone policy CAs<br><br>(2048-bit key) | 10 years | Renew at least every 5 years to ensure that child-issuing CAs can be issued for 5 years. Renew by using a new key at least every 10 years. |
| Enterprise issuing CAs for medium-security certificates<br><br>(1024-bit key) | 5 years | Renew at least every 3 years to ensure that certificates can be issued for 2 years. Renew by using a new key at least every 5 years. |
| Enterprise issuing CAs for high-security certificates<br><br>(2048-bit key) | 5 years | Renew with new key at least every 4 years to ensure that certificates can be issued for a year. |
| Enterprise issuing CA for external certificates<br><br>(2048-bit key) | 5 years | Renew at least every 4 years to ensure that certificates can be issued for a year. Renew by using a new key at least every 5 years. |
| Secure e-mail and secure browser certificates | 1 year | Renew by using a new key at least every 2 years. |
| Smart card certificates<br><br>(1024-bit key) | 1 year | Renew by using a new key at least every 2 years. |
| Administrator certificates<br><br>(1024-bit key) | 1 year | Renew by using a new key at least every 2 years. |
| Secure Web server certificates<br><br>(1024-bit key) | 2 years | Renew by using a new key at least every 2 years. |
| Business partners' user certificates for an extranet<br><br>(1024-bit key) | 6 months | Renew by using a new key at least every year. |

**Note**

All these values are suggestions and may be dictated by legal, governmental, or contractual rules that are specific to the organization. Changes of policy and extensions during renewal and rekey of CAs may also require subsequent changes of CPS recertification, audit, and so on.

# Example Scenario for Contoso

Updated: April 13, 2012

Applies To: Windows Server 2003 with SP1

Since a lot of planning considerations and best practice approaches are covered in the previous section, here is a real world example of PKI topology. The example describes the best practices that are mentioned earlier in this document and also describes most of the options of a complex PKI. By leaving out distinct components (like the HSM or the intermediate CA level), you can also adjust the topology to environments for smaller organizations.

The fictitious organizations name is Contoso Company. Contoso is an international company that has already deployed Active Directory and is introducing a Windows Server 2003 PKI. The planning for the project is already finished. Besides other preparation work, the parameters as listed in Appendix B, which will drive the PKI configuration.

Use the following installation instructions to set up a PKI that is similar to the Contoso PKI.

**Platform Decision**

The Contoso Corporation has decided to deploy a Windows Server 2003 PKI hierarchy. The organization recognizes ease of deployment, benefits of strengthened security, security-integrated applications, and the Active Directory-integrated management infrastructure that is in their current Active Directory infrastructure.

Because Contoso wants to benefit from all PKI improvements to the Windows Server 2003 family, they have prepared Active Directory to run in a Windows Server 2003 forest and in Windows Server 2003 domain functional mode. For more information about how to raise the domain functional level on a computer running a member of the Windows Server 2003 family, see the article "How to raise Active Directory domain and forest functional levels" on the Microsoft Knowledge Base.

The Contoso company has a mix of clients that are running a member of the Windows 2000 family or Windows XP Professional and uses a number of integrated applications that can take advantage of the Windows Server 2003 PKI, including Secure/Multipurpose Internet Mail Extensions (S/MIME), encrypting file system (EFS), L2TP or IPsec VPN connections, 802.1x wireless access, and SSL-enabled Web servers. Smart cards are used for user logon.

**PKI Design**

Contoso has decided that a three-tier PKI topology is most suitable for their organization. If an organization wants to benefit from a two-tier PKI topology, the implementation guidelines that are outlined in this documentation can also be applied.

The Contoso PKI consists of different certificate servers. Every CA will be implemented with Certificate Services as implemented in Windows Server 2003, Enterprise Edition. The following figure illustrates the Certificate Services architecture for Contoso Corporation.




**Figure 5: CA Hierarchy for Contoso Corporation**

**Root CA**

The Contoso stand-alone root CA is never connected to a network and remains offline and physically secured. The root CA issues and revokes certificates for intermediate CAs in the hierarchy. To raise the security level of the CAs private key, an HSM extends the root CAs hardware configuration. The CA certificate and the CRL are manually published and made available through an HTTP and an LDAP distribution point.

**Intermediate CAs**
The intermediate CAs are physically secured and operated offline. Contoso has decided to operate two intermediate CAs. The separation of intermediate CAs allows the organization to set different CRL distribution points and CRL publishing intervals. Since intermediate CAs are also treated as sensitive components in the organization's PKI, they are also equipped with HSMs.

**Issuing CAs**
The issuing enterprise CAs are responsible for certificate enrollment to end-entities. These CAs are distributed to different geographic locations to allow local availability. Physical security, however, has to take precedence over close proximity of the servers. These CA servers are online and available to service requests at any time. You can improve availability in the future by deploying more issuing CAs.

**Contoso Environment Summary**
The following list describes the Contoso environment:

- A forest and domain environment with computers that run only members of the Windows Server 2003 family
- Clients that are running either Windows 2000 or Windows XP
- Three-tier PKI hierarchy
- Self-signed stand-alone offline root Windows Server 2003, Enterprise Edition stand-alone CA with HSM to support role-separation
- Intermediate offline stand-alone Windows Server 2003, Enterprise Edition stand-alone CA with HSM to support role separation
- Several online, Active Directory-integrated issuing Windows Server 2003, Enterprise Edition CAs to support auto-enrollment and V2 templates.

**Stand-alone Offline Root CA**
The following section describes the steps that you should use to install an offline root CA by using a computer running Windows Server 2003, Standard Edition. The installation procedure for a Windows 2000 CA is similar to the installation procedure for a computer running Windows Server 2003, Standard Edition, so you can also use the following installation procedure on computers running Windows 2000.

| **Note** |
| --- |
| The stand-alone offline root CA is also referred to as "CorporateRootCA" in this document. |

**Installation Prerequisites**
A server that is running Windows Server 2003, Standard Edition is installed with the base operating system and latest updates. Verify that you have the following information before you start to install Certificate Services:

- The CPS that has all of the parameters that are specific to the organization
- The Windows Server 2003 installation medium, such as the original CD-ROM
- Appropriate hardware with a floppy disk drive
- Both computer and CA naming conventions
- Directory and file paths to be used for CRL distribution point, such as AIA
- Other CA configuration information, such as CRL publication intervals
- HSM, if applicable

**Install the Offline Root CA**

To set up the root CA, use the steps in this section. Before you begin, verify that the following concepts have been reviewed and approved by your organization:

- Public key infrastructure concepts
- Requirements that describe the purpose of certificate usage and enrollment
- Details about CA configuration including the hierarchy of the PKI
- The renewal strategy that you are going to use for the root CA is planed
- Operational security procedures and policies

**Workgroup Membership**

The CorporateRootCA must be a workgroup member because it is not connected to the network and has no link to a domain controller. Nevertheless, the server name must be unique in your organization because the server name is part of the information that will be published in Active Directory.

Verify that the naming information is correct after you log on to the local computer by using the **net config workstation** command at a command prompt. (Note that the values that are in italics may be different for you, according to your configuration.) The following section is an example of how you can use the **net config workstation** command to accomplish this.

```
D:\>net config workstation
Computer name                        \\CONCORP-CA-00
Full Computer name                   concorp-ca-00
User name                             Administrator
Workstation active on
        NetbiosSmb (000000000000)
        NetBT_Tcpip_{7CD8C0C6-02A5-4EB4-8081-5D1977FD0AA5}
(0008C75BDEC0)
Software version                     Microsoft Windows Server 2003
Workstation domain                    WORKGROUP
Logon domain                          CONCORP-CA-00
COM Open Timeout (sec)               0
COM Send Count (byte)                16
COM Send Timeout (msec)              250
The command completed successfully.
```

Verify that the logon domain name is the same as your server name.

For more information, see the following articles on Microsoft Web sites:

- "Checklist: Creating a certification hierarchy with an offline root certification authority" on the Microsoft Web site
- "HOW TO: Install a Windows 2000 Certificate Services Offline Root Certification Authority" on the Microsoft Knowledge Base

**Installing an HSM on an Offline Root CA**

Some organizations may choose to protect the root private key with additional hardware. Before you install and configure Certificate Services, verify that the HSM is correctly set up according to the manufacturer's installation instructions. For information about how to install an HSM on computers that are running either Windows 2000 or Windows Server 2003, see the following Web sites:

- Windows 2000 Server and PKI: Using the nCipher Hardware Security Module on the Microsoft Web site
- Deploying Certificate Services on Windows 2000 and Windows Server with the Chrysalis-ITS Luna CA3 Hardware Security Module on the Microsoft Web site

**Prepare the CAPolicy.inf File for the Root CA**

An issued certificate typically inherits properties (for example, certificate lifetime, the distribution point of the CRL, and other parameters) from a certificate template that is provided by the issuing CA. Since the root CA requires a certificate for itself, the

root CA must self-sign the root CA certificate because there is no parent CA that could issue the CA certificate.

Before the CA certificate is generated, custom configuration of the CA that is relevant to the CA certificate is required. The CAPolicy.inf file has all configuration information that is required to generate the self-signed CA certificate according to the organizations needs.

| **Caution** |
| --- |
| Configuring the CAPolicy.inf file is a very important step that you must finish before you set up a Windows Server 2003 root CA. If you do not use the CAPolicy.inf file for the offline root CA setup procedure, the CRL and AIA distribution points that become part of issued certificates are set to distribution points on the local computer. Because an offline CA is never accessible from the network, clients cannot resolve the CRL or AIA distribution point. To prevent this issue, you must explicitly add both the *CRLDistributionPoint* and *AuthorityInformationAccess* parameters to the CAPolicy.inf file. As noted in the "CRL Best Practices" section in this document, both the CRL and AIA CRL distribution point of a root CA need to be defined as empty. |

To ensure that the CAPolicy.inf file is correctly processed:

- The ASCII text file must be available on the local computer before the CA setup procedure starts or before CA certificate renewal is attempted
- The file is placed in the **%Systemroot%** folder on the local computer on which the CA is installed

The syntax follows the specification that is described in the CAPolicy.inf syntax" section in the appendix of this document.

| **Note** |
| --- |
| After you use CAPolicy.inf on a CA, do not remove it from the computer because configuration parameters, like renewal key length, should be consistent during the life cycle of the CA. |

Also, during the installation procedure, you do not receive a warning message if the CAPolicy.inf file is not in the correct format because there is no syntax or error-checking mechanism. For setup logging and debug information, see the *Systemroot*\Certocm.log file.

To configure the CAPolicy.inf file:

1. Log on to the CorporateRootCA computer as an administrator.
2. Open a text editor, such as Notepad.
3. Copy the sample text file that is in the "Sample CAPolicy.inf file for the CorporateRootCA" section of this document as a template.
4. Paste the text into the file and then save it to **%Systemroot%**\CAPolicy.inf.

**Installing the Offline Root CA Software Components**

| **Important** |
| --- |
| Perform any renaming operations before the CA services become part of the configuration. You cannot change the NetBIOS computer name or the computer's membership in a domain or workgroup after certificate services is installed, because the name is part of the Certification Authority configuration information. |

To install the offline root CA software components, use the following procedure.

1. Log on to the CorporateRootCA computer as an administrator. Note that this account will be permitted as a CA administrator during the CA installation procedure. You can delegate the CA administrator role to other user accounts after the setup and configuration procedures are finished. For more information about CA roles and permission, see Windows Server 2003 Server Help.

2. Use one of the following procedures to open **Add/Remove Windows Components**:

   **To use a command prompt:**

   a. Click **Start**, click **Run**, type **cmd**, and then click **OK**.

   b. At the command prompt, type **sysocmgr /i:sysoc.inf** and then press ENTER.

   **To use Control Panel:**

   a. Click **Start**, point to **Settings**, and then click **Control Panel**.

   b. Click **Add or Remove Programs**.

   c. In **Add or Remove Programs**, click **Add/Remove Windows Components**.

| **Note** |
| --- |
| To run Certificate Services, check for the following software components: |

   - Certificate Services
   - Internet Explorer
   - (Optional) Certificate Services Web enrollment support
   - (Optional) Internet Information Services for Web enrollment support

   It is not recommended that you install any other Windows components on a Windows Server CA. If you install additional components, reliability or security of a root CA may be compromised if a secure configuration is required by the organization.

   An offline Windows 2000 CA requires Internet Information Services to support offline CA enrollment. Unlike Windows 2000 Server, a Windows Server 2003 certification authority can process offline certificate requests within the Certification Authority MMC

3. In the Windows Components Wizard, select the **Certificate Services** check box, and then click **Next**.

4. In **CA Type**, click **Stand-alone root CA**, select the **Use custom settings to generate the key pair and CA certificate** check box, and then click **Next**.

   It is expected that the enterprise root CA and enterprise subordinate CA options are not available because the computer is not a member of an Active Directory domain.

5. Do one of the following:

   - If you installed an HSM, in **CSP**, select the CSP that you installed during the HSM installation procedure.
   - If you did not install an HSM, in **CSP**, click **Microsoft Strong Cryptographic Provider**.

6. **In Hash algorithm**, click the default hash algorithm, **SHA-1**.

   SHA-1 is the most common and interoperable hash algorithm that is used by applications and operating systems. For more information about CSP support on computers that are running Windows 2000, see "Microsoft Enhanced CSP Is Not

Supported for Certificate Services Installations" on the Microsoft Knowledge Base.

7. In **Key length**, click **4096**.

   If you choose a different key length, confirm that the key length is interoperable with organizational applications and other PKI components. There is no verification of the key length that you type into the box. If an HSM or smart card CSP is utilized, the CSP will be required to interact with the desktop.

8. Confirm that both the **Allow this CSP to interact with the desktop** check box and the **Use an existing key** check boxes are cleared, and then click **Next**.

9. On **CA Identifying Information**, in **Common name for this CA**, type a name that will identify the CA to you. In this example, use **CorporateRootCA**.

   As it is specified in the certificate practice statement, you must specify a common name (CN) for the CA. The CN cannot exceed 64 characters in length, however, it is recommended that you use a maximum length of 51 characters to prevent an encoding length rule violation.

10. (Optional) **In Distinguished name suffix**, type **DC=concorp,DC=contoso,DC=com**.

    If you type a name, confirm that you have typed it correctly so that it works in the context of the Active Directory domain name. In the Contoso scenario, the distinguished name is **DC=concorp,DC=contoso,DC=com**. If you install a CA on a computer that is a domain member with Enterprise Administrator privileges, the distinguished name suffix is automatically configured. You can also set the distinguished name suffix at a later time by using the Certutil.exe command.

11. In **Validity period**, select **10 years**.

    Enter the validity time as defined in your organization's certificate practice statement. In this example, a validity period of 10 years is set for CorporateRootCA.

12. If you have uninstalled a CA on this computer already, you receive a warning message that confirms that you want to overwrite the private key from the previous CA installation. It is recommended that you ensure that the private key is never required again. If you make a backup copy of the system, it is more likely that you will not lose any data. (Instead of backing up the entire system, you can also make a backup copy of the private key. To do this, at a command prompt, type **certutil –backupkey -?**) If you are not sure if you want to overwrite the private key, click **No** to cancel the installation procedure. If you click **Yes**, a new key is generated and the new key replaces the existing key. Note that Windows 2000 CAs do not support the distinguished name suffix specification as part of the installation wizard.

    The public and private keys are then generated by the CSP. If you use the default CSP, the keys are written to the local computers key store. If you did not use an HSM, the key is generated by CryptoAPI and is stored in the profile of the system account on the local computer. The length of time that is required to generate the key depends on both the size of the key that is generated and the CPU performance of the local computer. If an HSM is installed and selected, the key is generated in the HSM and stored according to the HSM specific architecture. Since no certificate templates are available on a stand-alone CA server that is a member of a workgroup, the CA certificate needs to be built from configuration information in the registry. The following default key usage extension values are added to the CA certificate:

    - Digital signature
    - Certificate signing
    - Certificate offline CRL signing
    - Certificate CRL signing

    However, if a root CA is installed on a computer that is running either a member of the Windows 2000 family or a

member of the Windows Server 2003 family, and that computer is a domain member, the CA inherits the Enhanced Key Usage extension settings from the CA template in Active Directory, even if the CA is installed as a stand-alone CA. If no Active Directory is available, the Enhanced Key Usage settings are also taken out of the configuration that is available in the registry.

13. In **Certificate database** and **Certificate database log**, enter the locations of the certificate database and the log files for the certificate database.

    The certificate database and the certificate database log must be saved to a local NTFS hard disk.

    On a stand-alone CA that is expected to infrequently issue CA certificates, you can retain both the certificate database and the certificate database log on the local computer's hard disk. To ensure that the CA is reliable and available, schedule backup operations of the computer. Backup may be performed even if the CA service is not running. For more information about CA backup and recovery, see "Certification Authority backup and recovery" in this document.

14. (Optional) To install a CA in the same location as a previously installed CA, select the **Preserve existing certificate database** check box.

    If you select this option, the new CA will use the existing database and preserve the certificates in the database. If you do not select this option, the existing database will be deleted. You should use this option only when you are trying to restore a CA from a backup or for CA migration.

    You can move both the database and log files to a different location. For more information, see HOW TO: Move the Certificate Server Database and Log Files on the Microsoft Knowledge Base.

15. Select the **Store configuration information in a shared folder** check box and, in **Shared folder**, enter a local pathname as the name for the shared folder, such as C:\CAConfig, and then click **Next**.

    The CA setup procedure cannot detect if the computer is supposed to run as either an online or offline CA. For an offline CA, the shared folder is not necessary, but must still be specified. If the CA is connected to the network, clients can gain access to the CA certificate through the shared folder.

    Depending on the name of the shared folder, a new share is created on the CA server computer. The path to the shared folder can be either a universal naming convention (UNC) path such as the default, \\Localhost\CAConfig, or a local path, such as C:\CAConfig. If the server does not have network cards installed or has all network interfaces disabled, you must choose a local path.

    Some information that is stored in the CAs configuration directory must be published to the organization's Active Directory at a later stage. For more information, see "Import parent CA certificates and CRLs into Active Directory" later in this document.

    When the Windows Components Wizard completes the Certificate Services configuration, you may be asked for the server's installation media to finish the installation.

    Because you do not need to install IIS on this computer, you may receive a warning message that states that the Certificate Services Web Enrollment Support is unavailable. If you receive this message, click **OK** and complete the installation procedure.




**Figure 6: IIS Installation Warning**

**Verify the Root CA Configuration**

The next procedure helps you to ensure that the root CA is correctly configured ready for production operations.

**Verify the Root CA Certificate**
Because the CA certificate is mandatory for a reliable certificate validation of all certificates that have been issued by your PKI, it is important to ensure that this certificate has all of the necessary information before proceeding with the installation of a CA hierarchy.

1. On the local root CA computer, at a command prompt, type:

   **certutil –ca.cert corporateRootCA.cer**

2. View the CA certificate to validate the information that is in the CA certificate. To do this, at a command prompt, type:

   **certutil.exe corporateRootCA.cer**

3. Verify that the italicized parameters are the same parameters that you noted in the configuration document in the previous section. In addition, make sure that the certificate lifetime is set to a period of 10 years.

```
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.5  sha1RSA
Issuer:  CN=CorporateRootCA
NotBefore: 6/5/2002 7:47 PM
NotAfter: 6/5/2012 7:54 PM
Subject:  CN=CorporateRootCA
```

**Verify the CorporateRootCA Configuration Information**
Use the steps in this section to verify the CA configuration:

1. At a command prompt, type **certutil –cainfo** and verify the CA type. The result will be similar to the following:

```
CA type: 3 -- Stand-alone Root CA
    ENUM_STANDALONE_ROOTCA -- 3
```

2. At a command prompt, type certutil –getreg | find /I Directory to verify the database settings. Verify the following italicized output:

```
ConfigurationDirectory   REG_SZ =  \\concorp-ca-00\CertConfig
  DBDirectory              REG_SZ =  C:\WINDOWS\system32\CertLog
  DBLogDirectory           REG_SZ =  C:\WINDOWS\system32\CertLog
  DBSystemDirectory        REG_SZ =  C:\WINDOWS\system32\CertLog
  DBTempDirectory          REG_SZ =  C:\WINDOWS\system32\CertLog
```

**Offline Root CA Configuration**
After the stand-alone offline root CA is installed, you must configure the properties of the offline root CA for certificates that are subsequently issued from the CA. These extensions are necessary to ensure correct revocation and chain building.

You can perform all of the steps that are described in this section by using one batch script. For more information, see Sample script to configure CorporateRootCA in this document.

**Map the Namespace of Active Directory to an Offline CA's Registry Configuration**

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your computer. Before making changes to the registry, you should back up any valued data on the computer. |

Because the offline root CA is not connected to the domain and does not automatically publish the CRL to Active Directory, you must set a key in the registry. To do this, at a command prompt, type the following command and then stop and start the CA service:

**certutil.exe –setreg ca\DSConfigDN CN=Configuration,DC=concorp,DC=contoso,DC=com**

where *DC=concorp,DC=contoso,DC=com* is the namespace of the forest root domain. This setting is primarily required for

CRLs and CA certificates (AIA) that are published in Active Directory.

This registry value sets the %6 replacement token that is required for the CRL location extension, as well as the CRL and AIA distribution points that are described in "Configure CorporateRootCA distribution points for CRL and AIA. For more information about the %6 replacement token, see CRL distribution point replacement token in this document.

| **Important** |
| --- |
| After you use this command to change the registry key, a new CRL and any new CA certificates that are issued must be republished. Only new certificates that are issued after you use the previous command will have this information available. It is important to note that you must reissue and republish any subordinate CA certificate if it was issued before you changed the registry key. |

**Configure CorporateRootCA Distribution Points for CRL and AIA**

The CRL and AIA distribution points must be set before any certificates are issued from the new CA.

This configuration step ensures that the correct information is embedded in each of the issued certificates so that the certificate's signature and revocation status can be verified. For additional information about certificate status and chain building, as well as how the AIA and CRL distribution point extensions are used by CryptoAPI, see Certificate Revocation and Status Checking (http://go.microsoft.com/fwlink/?LinkID=27081).

For all CA types (online or offline, root or subordinate, enterprise or stand-alone), the configuration of the AIA extension and the CRL distribution point extension is critical. If they are not configured correctly or if they contain invalid extensions, there may be unexpected problems. For example, smart card login attempts may not work, there may be invalid e-mail digital signatures, or there may be Web sites that are not trusted.

**Configure CorporateRootCA Distribution Points for CRL and AIA by Using the User Interface**

CRL distribution point and AIA extension changes take effect only after the CA is restarted and the extensions appear in certificates issued only after the changes are applied.

On computers that run a member of the Windows 2000 family, the CRL and AIA configuration process is different than on computers that are running a member of the Windows Server 2003 family.

Before changing the CRL configuration, verify the default settings.

Log on to the computer with an account that has Certification Authority Administrator permissions.

Type the following at a command prompt:

**certutil -getreg ca\CRLPublicationURLs**

The following report of the default CRL distribution points is displayed. Note these settings if you need to change the CRL configuration to its original state.

```
CRLPublicationURLs REG_MULTI_SZ =
    0: 65:C:\WINDOWS\system32\CertSrv\CertEnroll\%3%8%9.crl
    CSURL_SERVERPUBLISH -- 1
    CSURL_SERVERPUBLISHDELTA -- 40 (64)

    1: 79:ldap:///CN=%7%8,CN=%2,CN=CDP,CN=Public Key Services,CN=Services,%6%10
    CSURL_SERVERPUBLISH -- 1
    CSURL_ADDTOCERTCDP -- 2
    CSURL_ADDTOFRESHESTCRL -- 4
    CSURL_ADDTOCRLCDP -- 8
    CSURL_SERVERPUBLISHDELTA -- 40 (64)

    2: 6:http://%1/CertEnroll/%3%8%9.crl
    CSURL_ADDTOCERTCDP -- 2
    CSURL_ADDTOFRESHESTCRL -- 4

    3: 0:file://\\%1\CertEnroll\%3%8%9.crl
```

As it is specified in the CPS, you must configure the CRL and AIA distribution point for certificates issued by this CA. To configure these extensions in a Windows Server 2003 CA, perform the following steps:

1. Log on to the computer running certificate services with an account that has Certification Authority Management permissions.

2. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Certification Authority**.

3. In the console tree, right-click the name of the CA that you want to work with, and then click **Properties**.

4. Click the **Extensions** tab.

#### Configure CorporateRootCA Distribution Points for the CRL

1. First, remove all of the CRL distribution point locations, except for the local CRL distribution point.

| **Caution** |
| --- |
| Do not remove the local CRL distribution point location. The local distribution point will look similar to the following path: C:\Windows\System32\CertSrv\CertEnroll*CorporateRootCA*.crl The CA must publish the CRL to the file system because all of the other distribution points are not accessible for this offline CA. The CA uses the local CRL to validate all certificates that are generated before the certificates are issued to users. The local path is not included in the CRL distribution point extension of issued certificates. |

2. On the **Extensions** tab, in **Select extension**, select **CRL Distribution Point (CDP)**.

3. In **Specify location from which users can obtain a certificate revocation list (CRL)**, click the default LDAP location, click **Remove**, and then click **Yes**.

4. Repeat Step 2 for all CRL distribution point locations except for the local CRL distribution point.

After you remove all of the appropriate locations, the remaining list of CRL distribution points will be similar to the following figure.




**Figure 7: CRL Distribution Point Locations**

Now, you will add the desired CRL distribution point locations to the list.

To provide multiple access protocol methods for CRL retrieval, different distribution points are provided to facilitate heterogeneous environments. Note that the LDAP path that is listed in the following table contains information about the organization's Active Directory namespace. If the certificate will be exchanged with external parties, define a neutral namespace. The example in this section uses the following CRL distribution points in the following order.

**Table 14 List of CRL Distribution Points for CorporateRootCA**

| Access protocol | CRL distribution point |
|---|---|
| [local] | C:\WINDOWS\system32\CertSrv\CertEnroll\%3%8%9.crl |
| HTTP | http://www.contoso.com/pki/%3%8%9.crl |
| LDAP | Ldap:///CN=%7%8,CN=%2,CN=CDP,CN=Public Key Services,CN=Services,%6%10 |

The [local] path should be the current Windows installation directory.

The order in which you should choose the access protocols is based on the type of certificates that are issued by a CA. A CRL distribution point with HTTP as the first protocol in the list is recommended for environments where CRL distribution without latency is critical or where most clients are not joined to an Active directory domain. HTTP locations generally do not replicate and do not have latency issues whereas an LDAP distribution point might be located in a distributed directory service, like Active Directory. The CRL distribution points that are listed in the table use the replacement tokens that are described in CRL distribution point Replacement Token in this document. For more information about CRL naming, see the chapter CRL Best Practices.

1. Click **Add**, and in Location, type the appropriate CRL distribution point path from the previous table, and then click **OK**.

   You can also copy and paste the path from the table.

2. Repeat Step 1 for each type of access protocol.

   Next, you must set the configuration parameters that dictate how the CRL will be published by the CA. The properties must be set for every CRL distribution point path.

3. In **Specify locations from which users can obtain a certificate revocation list (CRL)**, click one of the paths.

4. While still on the **Properties** tab, select or clear the check box that is listed in the previous table, depending on the type of path, and then click **Apply**.

5. Repeat steps 3 and 4 for each path.

   **Table 15 CRL Distribution Point Properties**

| CRL distribution point property | File | HTTP | LDAP |
|---|---|---|---|
| Publish CRLs to this location check box | Select | N/A | Clear |
| Include in all CRLs check box | N/A | N/A | Select |
| Include in CRLs check box | N/A | Clear | Select |
| Include in the CDP extension of issued certificates check box | N/A | Select | Select |
| Publish delta CRLs to this location check box | Clear | N/A | Clear |

**Note**

In **Publish CRLs to this location**, since the *CorporateRootCA* computer is not attached to the network, the CA cannot automatically publish the CRL to the LDAP CRL distribution point. By default, this option is chosen on an enterprise CA to automate the CRL publishing to the LDAP CRL distribution point.

In **Publish CRLs to this location**, a UNC file path can be specified to publish to clustered Web servers using IIS for CRL fault tolerance.

If the **Publish Delta CRLs to this location** check box is selected, make sure that the delta CRL is also published. For more information, see "Configure CRL publication interval via the user interface" in this document.

For a description of CRL properties, see CRL publishing properties in this document.

**Configure CorporateRootCA Distribution Points for AIA**

| ⚠Caution |
| --- |
| In the procedure below, do not remove the local AIA path location. The local path will not be contained in the AIA extension of issued certificates. |

1. In *CorporateRootCA***Properties**, on the **Extensions** tab, select **Authority Information Access (AIA)**.

2. In **Specify locations from which users can obtain the certificate for this CA**, click the default LDAP location, click **Remove**, and then click **Yes**.

3. Repeat the previous step to remove all of the entries except for the local entry.

   You can also clear the check boxes for all AIA publishing options instead of removing the AIA path from the list.

4. Click **Add**, and in **Location**, type one of the AIA distribution points from the following table, and then click **OK**.

5. Repeat the previous step for each AIA distribution point in the table below.

   Note that the LDAP path can expose internal namespace information if the certificates will be exchanged with external parties. Change the LDAP CRL distribution point to a permanent and publicly-available distribution point if certificates are exchanged with external parties.

   **Table 16 List of AIA CRL Distribution Points for Contoso**

| Access protocol | AIA Distribution Point |
| --- | --- |
| [local] | D:\WINDOWS\system32\CertSrv\CertEnroll\%1_%3%4.crt |
| HTTP | http://www.contoso.com/pki/%1_%3%4.crt |
| LDAP | ldap:///CN=%7,CN=AIA,CN=Public Key Services,CN=Services,%6%11 |

6. In **Specify locations from which users can obtain the certificate for this CA**, click one of the locations, and then select or clear the check boxes according to the following table.

   These configuration parameters control how the AIA extension is used by the CA in issued certificates. You must set the properties for every AIA path that is specified on the **Extensions** tab.

   **Table 17 AIA Properties**

| AIA property | FILE | HTTP | LDAP |
| --- | --- | --- | --- |
| **Include in the AIA extension of issued certificates** check box | N/A | Select | Select |
| **Include in the online certificate status protocol (OCSP) extension** check box | N/A | Clear | Clear |

7. Repeat the previous step for each location, except for the file locations.

8. Click **OK**, and then click **Yes** to apply the changes you have made and restart the computer.

**Configure the CorporateRootCA CRL and AIA CRL Distribution Point From a Batch File**

The CRL distribution point path is stored as a multivalued extension in the registry. You can also set the appropriate value with the Certutil.exe utility. This procedure is similar to the steps that are outlined in "Configure CorporateRootCA distribution points for CRL and AIA by using the user interface" earlier in this document. To configure both the CRL distribution point and AIA paths for a Windows Server 2003 CA with **Certutil.exe**, follow the steps that are described in this section.

| ◆Important |
| --- |

| |
|---|
| Because percent character (%) variables are handled differently than the configuration UI in batch files and at a command prompt, you must use the %% notation if you want to run the example script in this section as a batch file. If Certutil is called from a command prompt, replace %% with a single %. |

Certutil.exe interprets a multivalued extension when you use **\n** as part of the value string. If a multivalued extension consists of only one value, verify that **\n** is appended as the last character in the value string. Otherwise, you create a string value that might be not recognized by the CA. For more information, type **certutil.exe –setreg -?** at a command prompt.

1. On a server that is running one of the appropriate members of the Windows Server 2003 family, open a text editor, such as Notepad, and then copy the following text as two separate rows into the text editor.

   ```
   certutil -setreg CA\CRLPublicationURLs
   "1:%WINDIR%\system32\CertSrv\CertEnroll\%%3%%8%%9.crl\n2:http://www.
   contoso.com/pki/%%3%%8%%9.crl\n10:LDAP:///CN=%%7%%8,CN=%%2,CN=CDP,CN
   =Public Key Services,CN=Services,%%6%%10"
   certutil -setreg CA\CACertPublicationURLs
   "1:%WINDIR%\system32\CertSrv\CertEnroll\%%1_%%3%%4.crt\n2:LDAP:///CN
   =%%7,CN=AIA,CN=Public Key Services,CN=Services,%%6%%11\n2:http://www
   .contoso.com/pki/%%1_%%3%%4.crt"
   ```

   If a Windows 2000 Server is used, the following commands perform the same function as above:

   ```
   certutil -setreg policy\RevocationCRLURL
   "http://www.contoso.com/pki/%%3%%8.crl\n"
   certutil -setreg policy\LDAPRevocationCRLURL
   "ldap:///CN=%%7%%8,CN=%%2,CN=CDP,CN=Public Key
   Services,CN=Services,%%6?certificateRevocationList?base?objectclass=
   cRLDistributionPoint\n"
   certutil -setreg policy\FileRevocationCRLURL "\n"
   certutil -setreg policy\IssuercertURL
   "http://www.contoso.com/%%1_%%3%%4.crt\n"
   certutil -setreg policy\LDAPIssuercertURL
   "ldap:///CN=%%7,CN=AIA,CN=Public Key
   Services,CN=Services,%%6?cACertificate?base?objectclass=
   certificationAuthority\n"
   certutil -setreg policy\FileIssuercertURL "\n"
   ```

2. Save the text file as **%temp%\MyCRLCDP.cmd**.

3. Close the text editor.

4. At a command prompt, type **%temp%\MyCRLCDP.cmd** to execute the commands from the text file.

5. At a command prompt, type **net stop certsvc** to stop the certificate server, because you must restart the computer to apply the change.

6. At a command prompt, type **net start certsvc** to restart the certificate server.

**Verify the CorporateRootCA CRL and AIA Configuration**

Because the configuration of CRL and AIA distribution points is very important, verify your configuration:

1. Do one of the following:

| If your CA is running | Do this |
|---|---|

| | |
|---|---|
| A member of the Windows 2000 Server family | At a command prompt, type the following commands, pressing ENTER after each line:<br>**certutil –getreg policy\RevocationCRLURL**<br>**certutil –getreg policy\LDAPRevocationCRLURL**<br>**certutil –getreg policy\FileRevocationCRLURL** |
| A member of the Windows 2003 Server family | At a command prompt, type the following command, and press ENTER:<br>**certutil –getreg ca\CRLPublicationURLs** |

2. Verify that the output is similar to the values you have configured in the previous section.

3. Do one of the following:

| If your CA is running | Do this |
|---|---|
| A member of the Windows 2000 Server family | At a command prompt, type the following commands, pressing ENTER after each line:<br>**certutil –getreg policy\IssuerCertURL**<br>**certutil –getreg policy\LDAPIssuerCertURL**<br>**certutil –getreg policy\FileIssuerCertURL** |
| A member of the Windows 2003 Server family | At a command prompt, type the following commands, pressing ENTER after each line:<br>**certutil –getreg ca\CACertPublicationURLs** |

4. Verify that the output resembles the following output and contains the proper organizational naming information.

```
CACertPublicationURLs REG_MULTI_SZ =
    0: 1:D:\WINDOWS\system32\CertSrv\CertEnroll\%1_%3%4.crt
    CSURL_SERVERPUBLISH -- 1
    1: 2:ldap:///CN=%7,CN=AIA,CN=Public Key
Services,CN=Services,%6%11
    CSURL_ADDTOCERTCDP -- 2
    2: 2:http://www.contoso.com/pki/%1_%3%4.crt
    CSURL_ADDTOCERTCDP -- 2
```

**Configure CRL Publication Interval By Using the User Interface**

After the CRL distribution point is set, you must configure the CRL publication interval. To configure the publication schedule, use the following procedure.

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Certification Authority**.

   This opens the Certification Authority MMC Snap-in.

2. In the console tree, right-click **Revoked Certificates**, and then click **Properties**.

3. In CRL publication interval, type a number for the CRL publication interval according to your CPS.

   For planning information, see Best practices for CRL publication in this document. Note that publishing by using minute-based intervals is available only through the registry and is not recommended for most installations.

4. Verify that the **Publish Delta CRLs** check box is not selected.

The **Publish Delta CRLs** setting is not available on computers running a Windows 2000 CA.

**Configure CRL Publication From a Batch File**

You can configure the CRL publication schedule information in the registry by using the **Certutil.exe** utility. To configure the CRL publication schedule information by using **Certutil.exe**, use the steps in this section. Note that these steps work on a Windows 2000 CA.

1. Start a text editor, such as Notepad, and then copy the following text into a text file:

```
certutil -setreg CA\CRLPeriodUnits 180
certutil -setreg CA\CRLPeriod "Days"
net stop certsvc & net start certsvc
```

2. Save the text file as **%temp%\MyCRLPeriod.cmd**, and then close the text editor.

3. At a command prompt, type **%temp%\ MyCRLPeriod.cmd**, and then press ENTER.

Instead, use these steps on a Windows Server 2003 CA.

1. Start a text editor, such as Notepad, and then copy the following text into a text file:

```
certutil -setreg CA\CRLPeriodUnits 180
certutil -setreg CA\CRLPeriod "Days"
certutil -setreg CA\CRLDeltaPeriodUnits 0
net stop certsvc & net start certsvc
```

2. Save the text file as %temp%\MyCRLPeriod.cmd, and then close the text editor.

3. At a command prompt, type %temp%\ MyCRLPeriod.cmd, and then press ENTER.

**Set the Validity Period for Issued Certificates at the Offline Root CA**

The lifetime of certificates that are issued by a Windows stand-alone CA is set to one year by default. Because these values might not match the organization's requirements, you must set a registry key to adjust this default value.

| **Note** |
| --- |
| The validity time of the root CA certificate is controlled at the setup and renewal of the CA certificate through the value that is specified in the CAPolicy.inf file. The registry value that is described in this section does not affect the validity time of the root certificate. |

1. Start a text editor, such as Notepad, and then copy the following to a text file:

```
certutil -setreg ca\ValidityPeriodUnits 10
certutil -setreg ca\ValidityPeriod "Years"
net stop certsvc & net start certsvc
```

2. Save the text file as **%temp%\MyVP.cmd**, and then close the text editor.

3. At a command prompt, type **%temp%\MyVP.cmd**.

For more information about allowing certificate requests to control the certificate validity time, see Control certificate validity time by certificate request in this document.

For more information about how to change the expiration date of certificates that are issued by a Windows 2000 CA, see HOW TO: Change the Expiration Date of Certificates Issued by a Windows 2000 Certification Authority in the Microsoft Knowledge Base.

**Republish the CorporateRootCA CRL**
After the CRL distribution point extensions are updated on the CA, new CRLs must be published to ensure that all of the clients will be able to gain access to the most current revocation list information. The publishing can be done through the MMC or by using **Certutil.exe** at a command prompt with the same results.

**Republish the CRL by Using the MMC**
Use the steps in this section on either a Windows Server 2003 CA or a Windows 2000 CA to republish the CRL. It is important to republish the CRL because adapted configuration parameters such as DSConfigDN are included as extensions in the CRL. Also, CRL properties affect the publication of the CRL.

1. Log onto the CA server with CA Manager permissions.

2. Open the Certification Authority MMC. To do this, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Certification Authority**.

3. Right-click **Revoked Certificates**, point to **All Tasks**, and then click **Publish**.

   A new base CRL is published. A delta CRL is published only if you have also set the CRL delta publication schedule.

4. When you are prompted to confirm the type of CRL that should be published with this request, click **New CRL**.

   Because only base CRLs are published by the offline root CA, only the **New CRL** option is available.

**Republish the CRL from a Command Prompt**
To publish the CRL, at a command prompt, type **certutil -CRL**, and then press ENTER. When you do this, the CRL is published to the location that you configured.

**Verify the Published CRL**
There are two extensions that you should verify after the CRL is published: the publication time and the **Published CRL locations** extension. When you verify the publication time for the CRL, you are also verifying whether the correct CRL publication is set on the configured schedule. You also need to verify that the **DSConfigDN** registry value is set correctly and that the **DSConfigDN** registry value is in the CRL.

**Determine the Name of the Most Current CRL**

1. At a command prompt, type **certutil –dynamicfilelist**, and note that the CRL path name that is displayed.

2. At a command prompt, type **certutil**, and use the CRL path and file name from the previous step as the command-line parameter. For example, you can type the following:

   **certutil %systemroot%\System32\CertSrv\CertEnroll***CorporateRootCA.crl*

   where *CorporateRootCA.crl* is the file name of the current CRL.

3. Verify that the **Effective date** extension has the same time as the expected CRL publishing time.

4. Verify that the **Published CRL Locations** extension does not have a DC=UnavailableConfigDN namespace.

   The **Published CRL Locations** extension is used when publishing the CRL to Active Directory using Certutil -DSPublish. If the namespace is set not defined, the extension is set to **UnavailableConfigDN**, when the CRL is created. Certutil -DSPublish will not work unless you specify the appropriate Configuration container on the command line.

```
Published CRL Locations
 [1]Locations
  Distribution Point Name:
   Full Name:
    URL=ldap:///CN=CorporateRootCA,CN=concorp-ca-
00,CN=CDP,CN=Public%20Key%20Services,CN=Services,DC=Unavailable
ConfigDN?certificateRevocationList?base?objectClass=cRLDistributionPoint
```

If the output has a **DC=UnavailableConfigDN**, to resolve this behavior, see Map the namespace of Active Directory or Republish the CorporateRootCA CRL, earlier in this document.

A CRL that is correctly configured should have the following output:

```
Published CRL Locations
 [1]Locations
  Distribution Point Name:
   Full Name:
     URL=ldap:///CN=CorporateRootCA,CN=concorp-ca-00,CN=CDP,CN
=Public%20Key%20Services,CN=Services,CN=Configuration,DC=concorp,
DC=contoso,DC=c
om?certificateRevocationList?base?objectClass=cRLDistributionPoint
```

| ⚠**Caution** |
| --- |
| Verify that the italicized text is the same as the value that is specified in **DSConfigDN** registry values. You should also verify that **objectClass** component of the LDAP path is correctly defined. |

#### Finalize the CA Configuration

If you use the steps in the previous sections, the CA is operational and ready to issue certificates.

If you install a Windows 2000 CA instead of a Windows Server 2003 CA, it is recommended that you also apply the additional configuration steps that are explained in Disable issuer name and issuer serial number, later in this document.

# Stand-alone Offline Intermediate CA (IntermediateCA1)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The stand-alone offline intermediate CA is also described as IntermediateCA1 in this document.

An offline stand-alone intermediate CA is part of a three-tier topology and is primarily used to add another layer of flexibility in an organization. The IntermediateCA1 is required to issue certificates to enterprise CAs that enroll authentication certificates only according to the CPS.

If you plan to implement a two-tier topology, skip this section and go to "Online Enterprise Issuing CAs (CorporateEnt1CA)," later in this paper.

**Installation Prerequisites**
To correctly install and configure the offline stand-alone intermediate CA, you will need the following:

- The CPS that has all of the parameters that are specific to your organization. For more information, see "Certificate practice statement," earlier in this paper.
- The Windows Server 2003 Server installation media
- Appropriate hardware with a floppy disk drive
- Two floppy disks, one labeled **Transfer-RootCA** and the second labeled **Transfer-IntermediateCA**.

The IntermediateCA1 must be a workgroup member because it is not connected with the network and has no connectivity to a domain controller. It is also important to ensure that the computer name of this server is unique in the organization's network, because the computer name is part of the CA configuration information that is published in Active Directory (For more information, see "Import Parent CA Certificates and CRLs into Active Directory" later in this document.) To ensure that the computer is a workgroup member, log on to the computer that becomes the offline intermediate CA and type the following at a command prompt, and then press ENTER:

**net config workstation**

If necessary, change the domain membership to a workgroup membership.

**Install an HSM on IntermediateCA1**
Before the CA setup procedure starts, verify that the Hardware Security Module (HSM) is set up correctly, according to the manufacturer's installation guide. For more information, see "Installing an HSM on an offline root CA" in this document.

**Prepare the CAPolicy.inf File for IntermediateCA1**
You must provide a CAPolicy.inf file before the CA setup procedure. The most important aspect of the capolicy.inf procedure is to allow all issuance policies at the intermediate level. A root CA always issues a SubCA certificate with all issuance policies allowed. At the intermediate CA level, this attribute must be set explicitly, otherwise it would allow all application policies but no issuing policies. An issuing CA cannot define any issuing policy if the CA certificate does not permit issuing of certificates. For more information, see Chapter 4.2.1.5, "Certificate Policies" at the Internet FAQ Archives Web site.

A stand-alone CA cannot define policies by certificate templates, because this is a feature of customizable V2 templates. A stand-alone CA does not benefit from V2 templates to issue certificates. The CAPolicy.inf file defines the policy that applies to all certificates that are issued by the intermediate CA.

Compared with the CorporateRootCA configuration, the CAPolicy.inf file does not need predefined CRL and AIA extensions because these configuration attributes are inherited from the parent CA which issues the subordinate CA certificate. Remember, that there are prerequisites that a CAPolicy.inf file gets processed properly.

Perform the following steps:

1. Log on to the IntermediateCA1 computer with administrative privileges.
2. Use a text editor, such as Notepad, to prepare the CAPolicy.inf file. For a template, use the sample file that is in "Sample CAPolicy.inf file for the IntermediateCA1 later in this paper.

3. Save the file to **%systemroot%\CAPolicy.inf**

**Obtain the Certificate and Its CRL from CorporateRootCA**

Before you can set up the IntermediateCA1 computer, you must install the root CA certificate and the most current CRL that CorporateRootCA provides, because IntermediateCA1 verifies the root certificate trust during installation.

You may need to manually obtain the parent CAs certificate once. After the parent CA has been renewed, a new CA certificate must be imported into the IntermediateCA1 certificate store again. Because the root CA and the intermediate CA are not normally connected to the network and are offline, you cannot make the root CA certificate available via the network to the intermediate CA.

Compared to the CA certificate that can have a long validity time, such as several years, the importing method that you use for the offline parent CA CRL must be performed at regular intervals that correspond to the CRL publication interval. (For more information, see "Configure CRL publication interval by using the user interface," earlier in this paper.) You have to import the offline parent CA CRL regularly because an offline CA cannot retrieve CRLs automatically through the network. You must install a copy of the latest CRL in the local certificate store of an offline CA.

The CRL and the CA certificate are transferred to the subordinate CA computer on a floppy disk in the following configuration steps. The certificate and the CRL are available in a file format on the computer that is running as CorporateRootCA.

To retrieve the CA certificate and CRL:

1. Log on to the computer that is running the CorporateRootCA as a user.

2. At a command prompt, type the following command to copy the current certificate to the Transfer-RootCA floppy disk:

   **certutil –ca.cert a:\concorp-ca-00**_*CorporateRootCA*_**.crtc > nul**

   | **Note** |
   | --- |
   | If the CA has already been renewed, there might be more than one CA certificate available. In that case, use the copy command to transfer all CA certificates from the file system to the floppy disk. To do this, at a command prompt, type the following and press ENTER. **copy %systemroot%\system32\certsrv\certenroll*.crt a:\.** |

3. To copy the CRL to a floppy disk, at a command prompt, type:

   **certutil -GetCRL a:**_*CorporateRootCA*_**.crl**

4. Remove the **Transfer-RootCA** floppy disk from the drive, and then insert the floppy disk into the subordinate CA computer.

   | **Note** |
   | --- |
   | If the certificate was already renewed, there may be more than one CRL, so it is safest to copy all of the available CRLs from the *Systemroot*\System32\CertSrv\CertEnroll folder to the disk. However, during initial setup of the root CA, there should be only one CRL in the directory. If a .crl file has a + sign at the end of its name, the **Publish Delta CRLs** option has not been switched off, as explained in an earlier configuration step. |

You can also export a certificate through the Certification Authority MMC snap-in. For more information about how to do this, see "Export the offline intermediate certificate at the root CA" in this document.

**Import the Root CA Certificate and CRL to the Intermediate CA**

The root CA certificate is required during the installation of the intermediate CA. It must be installed in the intermediate CAs certificate store before the intermediate CA is set up. Use the Certutil.exe command to import CA certificates into the certificate store, as described later in this section. When you do this, the certificates and CRLs are imported in the correct location.

If the *CorporateRootCA* certificate has been renewed, it is important that you import the entire set of CA certificates and CRLs. A set can be identified by the version number, because the CA certificate and CRL have the same version number.

**Example**

If *CorporateRootCA* is running with a certificate that was generated during installation, then the following files must be imported into the intermediate CA.

**Table 18 Files to Import With a New Certificate**

| File name | Description |
|---|---|
| Concorp-ca-00_CorporateRootCA.crt | CA certificate |
| *CorporateRootCA*.crl | CRL |

As mentioned earlier, you must import the CA certificate and CRL from the *CorporateRootCA* after the Root CA certificate was renewed. In the example below, note that Windows adds an incremental value to the filename if there is more than one CA certificate and CRL. For example, if the *CorporateRootCA* has been renewed twice, you must import the following list of files into IntermediateCA1.

**Table 19 Files to Import With a Previously-Used Certificate**

| File name | Description |
|---|---|
| Concorp-ca-00_CorporateRootCA(2).crt | CA certificate |
| CorporateRootCA(2).crl | CRL |

**Import the Root CA Certificate and CRL to an Intermediate CA Using the MMC**

This section describes how you can import a certificate import by using the Certificates MMC. The following steps about how to import the root CA and CRL to an intermediate CA by using the MMC snap-in are primarily for illustration purposes.

It is easier to import the CA certificates at a command prompt. The procedure for this is given later in this section.

1. To use the Certificates MMC to import a certificate and CRL, first verify that the CA certificate uses both the correct context and container. Log on to the IntermediateCA1 computer as a local administrator. Local admin permissions are required to import certificates or CRLs into the local systems certificate store.

2. Click **Start**, click **Run**, type **mmc.exe**, and then press ENTER.

3. Add the Certificates MMC snap-in:

   a. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.

   b. Click **Certificates**, and then click **Add**.

   c. Click **Computer account**, and then click **Next**.

   d. Click **Local computer**, and then click **Finish**.

   | Note |
   |---|
   | Because there are different certificate stores on a computer, you must select the correct certificate store. The computer account is required because the CA runs as part of the local system security context. Because of this, the CA can gain access to all of the information that is stored in the local computer's certificate store. Only security principals who have administrative permissions on the computer can write certificates to the System certificate store. For detailed description about certificate stores, see the Security chapter of the Windows 2000 Resource Kit on the Microsoft Web site. |

   e. Click **Close**, and then click **OK**.

4. Import the certificate. To do this:

a. Click **Certificates**, and then, on the **View** menu, click **Options**.

b. Select the **Physical certificate stores** check box.

c. In the console tree, double-click **Certificates (Local Computer)**, double-click **Trusted Root Certification Authorities**, and then double-click **Registry**.

d. Right-click **Certificates**, point to **All Tasks**, click **Import**, and then click **Next**.

e. Insert the **Transfer-RootCA** floppy disk into the floppy disk drive, and then click **Browse**.

f. Navigate to the certificate file, and then click **Open**.

g. Decide the location where you want the certificate stored, click **Place all certificates in the following store**, and then click **Next**.

   Note that your current certificate container is the predefined value because the import procedure has been started from there.

h. After you view the report about which options you selected in the import wizard, click **Finish** to import the certificate.

i. After you click **Finish**, you receive a message that confirms the status of operations. Also, the certificate appears in the list of certificates.

5. Import the CRL. To do this:

   a. Double-click **Certificates (Local Computer)**, and then double-click **Trusted Root Certification Authorities**.

   b. Right-click **Registry**, point to **All Tasks**, click **Import**, and then click **Next**.

   c. Insert the **Transfer-RootCA** floppy disk into the floppy disk drive, and then click **Browse**.

   d. Browse to your floppy disk drive, click the CRL file, and then click **Open**.

   e. Click **Place all certificates in the following store to decide in which location the certificate should be stored**, and then click **Next**.

      The registry node that is under Intermediate Certification Authorities is the predefined value because the import procedure has started as an action on the Intermediate Certification Authorities container.

   f. After you review the report that displays the options that you have selected, click Finish to import the certificate.

After you complete this procedure, the CA certificate and the CRL are installed in the local computers certificate store.

| **Note** |
| --- |
| You must repeat the steps in this section to import more CRLs and certificates. You must do this if the CorporateRootCA certificate has been renewed or a new version of the CRL has been published. |

**Find a Certificate in the Certificate Store**

If you imported the certificate into the incorrect certificate store, you may want to use the Find Certificates option.

You can also use the **Find Certificates** option to identify duplicate certificates that exist in several certificate stores. If you correctly set up the certificate, the certificate is kept only one time. If the same certificate appears several times, remove the duplicate certificates and verify that the certificate is stored in the correct container. It is important to know which certificate

belongs in which certificate store. For information about how to verify CA certificates, see the "Relationship of the Configuration Container and Certificate Store" section in this document. For more information regarding root certificates, see the articles "Trusted Root Certificates That Are Required By Windows 2000" on the Microsoft Web site and "How to Remove a Root Certificate from the Trusted Root Store" on the Microsoft Knowledge Base.

To find a certificate in the certificate store:

1. Click **Start**, click **Run**, type **mmc.exe**, and then press ENTER.

2. Right-click **Certificates**, and then click **Find Certificates**.

3. Choose your search criteria, and then click **Find Now**.

If your search is successful, you will see a list of certificates and the certificate's corresponding store that match your search criteria.

**Import the Root CA Certificate and CRL into an Intermediate CA from a Batch File**
To import both the root CA certificate and the CRL from a batch file:

1. Log on to the IntermediateCA1 computer as a local administrator because local administrative permissions are required to import certificates or CRLs into the local computer's certificate store.

2. Click **Start**, click **Run**, in the **Open** box, type **cmd.exe**, and then press ENTER.

3. Insert the floppy disk labeled **Transfer-RootCA** into the floppy disk drive on the intermediate CA computer.

4. At a command prompt, type the following two commands, and then press ENTER.

   **for %C in (***FloppyDrive***:*.crt) do certutil –addstore –f Root %C**

   **for %C in (***FloppyDrive***:*.crl) do certutil –addstore –f Root %C**

   where *FloppyDrive* is the drive letter of the floppy disk drive.

This will install all certificates and the latest CRL to the appropriate CryptoAPI store. (Note that the loop around the **certutil** command simplifies the import procedure because there may be more than one certificate or CRL on the floppy disk that needs to be imported.) The optional **–f** parameter forces an overwrite of the certificate if the certificate has been previously added to the store.

Only valid certificates are imported to the certificate store.

| **Note** |
| --- |
| Because it might be difficult to determine which CA certificate or CRL version is required, it is recommended that, if several CRLs exist, you import all CRLs from the root CA. For more information about the CA certificate and CRL storage, see Relationship of the configuration container and certificate store in this white paper. |

**Verify the Root CA Certificate Import Procedure From a Command Prompt**
After both the root CA certificate and CRL are imported, you can use the **Certutil.exe** utility to confirm that the import procedure was successful. It is important to insure that the certificates have been put into the right certificate stores.

To see a list of certificates that are stored in the root CA certificate store, type the following command at a command prompt:

**certutil –verifystore root**

The version number is shown as part of the output text that will appear. Confirm that the version number of the CA certificate and the CRL match.

**Install the Offline Intermediate CA Software Components**
The installation procedure that you use for a subordinate CA is different from the installation procedure that you use for a root

CA. Use the following steps to set up IntermediateCA1:

1. Log onto IntermediateCA1 as a local administrator.

   During the CA installation procedure, this account becomes a CA administrator, which is a role that can also be delegated to other user accounts. For more information about CA roles and permission, see Windows Server 2003 Server Help.

2. To open the Windows Components Wizard, do one of the following:

| To | Do this |
|---|---|
| Use a command prompt | a. Click **Start**, click **Run**, and in Open, type **cmd**, and then click **OK**.<br>b. At the command prompt, type **sysocmgr /i:sysoc.inf**, and then press ENTER. |
| Use Control Panel | a. Click **Start**, point to **Settings**, point to **Control Panel**, and then click **Add or Remove Programs**.<br>b. In **Add or Remove Programs**, click **Add/Remove Windows Components**. |

3. Select the **Certificate Services** check box, and then click **Next**.

   To correctly run Certificate Services, the following list of software components is required. Web enrollment and IIS are optional components on an offline Windows Server 2003 CA that could be installed with the CA at the same time or at a later date.

| **Note** |
|---|
| As described in "Installing the offline root CA software components" in this document, IIS is not required on an offline CA. However, you can have IIS on the computer in order to enroll certificates through Web enrollment support. IIS is not recommended as a security best practice, but is shown in this document only as an example for the procedures. |

   Certificate Services

   - Certificate Services CA
   - Certificate Services Web enrollment support

   Internet Explorer

   Application Server

   - Enable network COM+ access
   - Internet Information Services (IIS)
   - Common Files
   - Internet Information Services Manager
   - World Wide Web services
   - Active Server Pages
   - World Wide Web services

A Windows 2000 offline CA requires IIS in order to satisfy offline requests. A Windows Server 2003 CA is also able to process offline certificate requests as a function of the Certification Authority MMC. Alternatively, you can submit offline requests from a command prompt by using **Certreq.exe**.

4. When you are prompted to choose the type of installation procedure, click **Stand-alone subordinate CA**, select the **Use custom setting to generate the key pair and CA certificates** check box, and then click **Next**.

   The **Enterprise Root CA** and **Enterprise Subordinate CA** options are not available because the computer is not a member of an Active Directory domain.

5. Do one of the following:

   - If you installed an HSM, in **CSP**, you must select the CSP that you installed during the HSM installation procedure in **CSP**.

   - If you did not install an HSM, in **CSP**, click **Microsoft Strong Cryptographic Provider**.

6. In **Hash algorithm**, click **SHA-1**.

   The default setting, SHA-1, is the most common and interoperable hash algorithm that is used by applications and operating systems. For more information about CSP support on computers that are running Windows 2000, see "Microsoft Enhanced CSP Is Not Supported for Certificate Services Installations" on the Microsoft Knowledge Base.

7. In **Key length**, select **2048**.

   There is no verification of the key length that you type into the box. Because of this, verify that the key length is interoperable with organizational applications and other PKI components.

8. Verify that both the **Allow this CSP to interact with the desktop** and **Use an existing key** check boxes are cleared, and then click **Next**.

9. In **Common name for this CA**, type a common name for the CA. For this example, type **IntermediateCA1**.

   As it is specified in the CPS, you must specify the common name (CN) for this CA. The CN cannot exceed 64 characters in length; however, it is recommended that you use a maximum CN length of 51 characters to prevent encoding length rule violation.

10. (Optional) In **Distinguished name suffix**, type the distinguished name suffix for the CA, and then click **Next**.

    If you type a distinguished name suffix in **Distinguished name suffix**, confirm that you have typed the name correctly so that it works in the context of the Active Directory domain name. In the Contoso scenario, the distinguished name is *DC=concorp,DC=contoso,DC=com*.

11. The CA certificate's validity period for a subordinate CA is always determined by the parent CA. For more information, see "Set the validity period for issued certificates at the offline root CA," earlier in this document.

12. If you have uninstalled a CA on this computer already, you receive a warning message that confirms that you want to overwrite the private key from the previous CA installation. It is recommended that you ensure that the private key is never required again. If you make a backup copy of the system, it is more likely that you will not lose any data. (You can also make a backup copy of the private key as an alternative to a system backup. To do this, at a command prompt, type **certutil –backupkey -?**) If you are not sure if you want to overwrite the private key, click **No** to cancel the installation procedure. If you click **Yes**, a new key is generated and the new key replaces the existing key.

    The key pair is generated by the CSP and written to the local computers key store.

13. On **Certificate Database Settings**, confirm that **Certificate database**, **Certificate database log**, and **Shared folder** are set to the folder that you want to use.

14. (Optional) To install a CA in the same location as a CA that was installed previously, select the **Preserve existing certificate database** check box, and then click **Next**.

15. In **Shared folder**, confirm that the specified folder is set to a local path, such as C:\CAconfig, and then click **Next**.

16. Insert the Transfer-IntermediateCA floppy disk into the disk drive.

17. On **CA Certificate Request**, click **Save the request to a file** and, in **Request file**, type a name for the request file that will be saved to the floppy disk, and then click **Next**.

    The file must have a .req extension, such as a:\IntermediateCA1.req.

18. If you receive a message that IIS must be stopped to continue the installation, click **Yes**.

    The intermediate CA needs to submit the certificate request to its parent offline CA. Because the CorporateRootCA computer is running without a network connection, you must transfer the requested file on a floppy disk.

| **Caution** |
| --- |
| Verify that the floppy disk is available before you proceed. If the storage device is not accessible, you receive an error message, the CA setup procedure stops, and you must reinstall the CA. Before you can reinstall the CA, you must uninstall Certificate Services Web-Enrollment Support if it was supposed to be installed. |

19. The Windows Component Wizard completes the certificate services configuration.

    When the CA certificate has obtained a signed subordinate CA certificate from its parent CA, the wizard displays a message which says that the installation has finished. Make sure that the local storage device is available to save the request file, and then click **OK**.

20. Click **Yes** to enable ASP pages that are required for Web enrollment services.

    IIS is installed for illustration purposes as part of this configuration, but Active Server Pages (ASP) pages are not enabled by default. Because of this, the CA setup procedure provides an option to automatically enable the ASP pages.

    If you click **No**, you can enable ASP by typing **certutil –vroot** at a command prompt at a later time.



**Figure 8: Enable Active Server Pages**

21. After the wizard finishes installing files, click **Finish**, and then click **Close**.

22. Remove the **Transfer-IntermediateCA** floppy disk from the disk drive, and then take the floppy disk to the parent CA (CorporateRootCA).

**Verify the Certificate Request**

Before the certificate request is submitted to the parent CA, verify that the policy identifier that you set in the CA configuration through CAPolicy.inf is correct. If the syntax of the CAPolicy.inf file is incorrect, certain configuration information may be missing from the request file. To verify that all configuration information is properly included in the certificate request, view

and examine the request file. To verify the request file, at a command prompt, type **certutil***RequestFile*, where *RequestFile* is the request file that you save to the floppy disk, including the correct path, and then press ENTER.

The command produces output that is similar to the following output. Verify that the Certificate Policies section is correct as well as all of the other information that is specified in the CAPolicy.inf file.

If the Certificate Policies section does not appear in your certificate request, see Prepare the CAPolicy.inf file for IntermediateCA1, in this document, correct the syntax in the CAPolicy.inf file, and then repeat the subordinate CA installation procedure.

```
Attribute[2]: 1.2.840.113549.1.9.14 (Certificate Extensions)
    Value[2][0]:
    Unknown Attribute type
Certificate Extensions: 6
    1.3.6.1.4.1.311.21.1: Flags = 0, Length = 3
    CA Version
        V0.0
    2.5.29.14: Flags = 0, Length = 16
    Subject Key Identifier
        84 b9 bf 37 a7 9b 0d 75 28 62 00 27 bf 72 da d0 66 a5 79 e8
    2.5.29.32: Flags = 0, Length = 139
    Certificate Policies
        [1]Certificate Policy:
             Policy Identifier=1.3.6.1.4.1.311.21.43
             [1,1]Policy Qualifier Info:
                  Policy Qualifier Id=User Notice
                  Qualifier:
                       Notice Text=Legal policy statement text.
        [2]Certificate Policy:
             Policy Identifier=1.3.6.1.4.1.311.21.47
             [2,1]Policy Qualifier Info:
                  Policy Qualifier Id=CPS
                  Qualifier:
                       http://www.contoso.com/pki/LimitedUsePolicy.htm
             [2,2]Policy Qualifier Info:
                  Policy Qualifier Id=CPS
                  Qualifier:
                       ftp://ftp.contoso.com/pki/LimitedUsePolicy.txt
             [2,3]Policy Qualifier Info:
                  Policy Qualifier Id=User Notice
                  Qualifier:
                       Notice Text=Limited use policy statement text.
```

**Certificate Request Processing with the Root CA through MMC**

The subordinate CA certificate request that is saved on the **Transfer-IntermediateCA** floppy disk must be signed by the parent (CorporateRootCA).

You can submit a request to an offline CA by using either the Certification Authority MMC or the Web Enrollment page that is on the parent Windows Server 2003 CA. You can also submit the request by typing **certreq.exe –submit** at a command prompt. All methods allow you to submit a certificate request that you have saved to a request file (*.req). This section will present the first method, using the Certification Authority MMC. For more information about using the Web Enrollment page, see "Certificate request processing with the offline parent CA (IntermediateCA1) through Web-Enrollment Support," later in this document.

| **Caution** |
| --- |
| If a previous CA setup procedure did not work and you repeat the setup procedure, do not reuse the request file from the earlier CA setup procedure. It has an association with previous key material that will not be associated with the current CA that you are installing. |

If a CA is set up, the key material is generated and the certificate request is submitted to the parent CA. The relationship between key material and certificate is maintained by the **AKI certificate** attribute. To ensure that the association of the CA key pair and certificate request matches, a unique request file must be used when a CA is set up.

1. Log on to the CorporateRootCA computer as a CA administrator.

   Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Certification Authority**.

2. You can also click **Start**, click **Run**, type **certsrv.msc**, and then press ENTER.

3. In the console tree, right-click the certification authority you are working with, point to **All Tasks**, and then click **Submit new request**.

4. Insert the Transfer-IntermediateCA floppy disk into the CorporateRootCA computer's floppy disk drive, browse to the disk drive, click the certificate request file, and then click **Open**.

5. A stand-alone CA typically issues certificates only after a manual issuing process. (You can change the request handling on the Policy Module tab of the CA's Properties.) In the default configuration, you must manually issue the certificate request by the parent CA:

   a. In the Certification Authority MMC console tree, under the name of the CA you are working with, double-click the **Pending Requests** container.

   b. In the details pane, right-click the appropriate pending certificate request that corresponds to the submitted subordinate CA request, point to **All tasks**, point to **View Attributes**, and then click **Extensions**.

   c. Click **Certificate Policies**, and then verify that the information is correct.

      If the certificate policy information that is defined in the Intermediate CA's CAPolicy.inf file does not appear here, deny the request and return to "Prepare the CAPolicy.inf file for IntermediateCA1," earlier in this document.

   d. In the console tree, click **Pending Requests**, and, in the details pane, right-click the pending request, point to **All Tasks**, and then click **Issue**.

      The request is processed and the certificate request is removed from the list.

      By default, a stand-alone Windows 2000 or Windows Server 2003 CA issues certificates with only a two-year lifetime. Because the registry key that has an impact on the validity time of the certificate was previously set, the certificate enrollment continues with the value that you specified. For more information, see "Set the validity period for issued certificates at the offline root CA" in this document.

   e. In the console tree, click expand the **Issued Certificates** container and, in the details pane, right-click the certificate, and then click **Open** to verify the certificate as described in the next step.

**Verify the IntermediateCA1 Certificate**

To ensure that the certificate that was issued for IntermediateCA1 has the correct certificate properties, verify the issued certificate:

1. Because a CA policy is specified for IntermediateCA1, the issued certificate will allow all issuer and all application policies. On the **General** tab, click **Issuer Statement** and verify that the certificate is valid for the following purposes:

   - All issuance policies
   - All application policies.

2. Click **Close** to return to the certificate viewer.

3. Click the **Details** tab, and then verify that the **CRL Distribution Points** and **Authority Information Access** values are the same as the distribution points that are specified . Verify other certificate attributes, as required. If values do not match, see "Configure CorporateRootCA distribution points for CRL and AIA," earlier in this document, to help you correct the configuration.

**Note**

You can also verify the certificate after it has been exported to a file. To view the certificate information from a PKCS #7, .der, or Base64-encoded certificate file, at a command prompt, type the *CertFile* name and hit ENTER. Replace ,*CertFile* with the location and file name of the certificate file.

**Export the Offline Intermediate Certificate at the Root CA**

If the certificate that was issued for the intermediate CA passes the verification steps, export the certificate from the root CA.

**Note**

Because the Certificate Export Wizard can include the complete certificate path with the exported file, you should use this method instead of the binary export method which only exports a single certificate.

1. Open the Certification Authority MMC.

2. In the console tree, under the CA that you want to work with, click **Issued Certificates**.

3. In the details pane, double-click the subordinate CA certificate you want to work with, click the **Details** tab, click **Copy to file**, and then click **Next**.

4. Click Cryptographic Message Syntax Standard – PKCS#7 Certificates (.P7B), select the Include all certificates in the certification path if possible check box, and then click Next.

5. Type a file name without an extension for the export file, and then save the file on the **Transfer-IntermediateCA** floppy disk.

   For example, you could type **A:\IntermediateCA1**. The file is automatically saved on the floppy disk with a .p7b file name extension.

6. Click **Next**, click **Finish**, click **OK**, and then click **OK** again.

   The certificate contains only public information, because the key material that is associated with the certificate was generated and is stored on the IntermediateCA1 computer. There is generally no need to protect the certificate information that is stored on the floppy disk. The CA certificate and the parent CA certificates are always considered to be public information.

7. In the console tree, click **Issued Certificates**.

8. Right-click the issued certificate in the details pane, point to **All Tasks**, and then click **Export Binary Data**.

   In **Columns that contain binary data**, **Binary Certificate** is the default choice.

9. Click **Save binary data to a file**, and then click **OK**.

10. Insert the **Transfer-IntermediateCA** floppy disk into the drive, in **File name**, enter a file name with a .cer extension, and then click **Save**.

    For example, you could type **A:\IntermediateCA1.cer**. The certificate is then saved in the DER-encoded file format.

11. On the **File** menu, click **Exit**, and then log off of the CorporateRootCA computer.

**Install the Certificate on IntermediateCA1**

You have now processed the request that was sent to the root CA and saved it on the **Transfer-IntermediateCA** floppy disk. You must now install the signed subordinate CA certificate that belongs to IntermediateCA1. You can install the CA certificate either by running a command at a command prompt or by using the Certification Authority MMC. The subordinate CA

certificate request will only be accepted by the parent CA if it carries the requesting CAs signature on the request.

**Verify the IntermediateCA1 Certificate Trust Chain**

To prevent unexplained or unintentional behaviors, verify the certificate trust chain. You must complete the trust chain verification procedure from a command prompt because the trust path that is displayed in the Certification Authority MMC Snap-in uses a different implementation for chain-building.

1. Log on to the IntermediateCA1 computer as a local administrator

2. At a command prompt, type

   **certutil –verify***a:\CACertFile***.crt**

   where a:*CACertFile* is the path and name of the file.

3. Press ENTER to view the full certificate verification results.

This command may generate a lot of output. When **dwErrorStatus** is not equal to zero, a certificate verification error has occurred, so you should verify that **dwErrorStatus** is equal to zero (0) on each line that is produced.

You can also use the following command

**certutil –verify** *a:\CACertFile* **.crt | findstr /c:dwErrorStatus**

where a:*CACertFile* is the path and name of the file.

Output that has completed the CA certificate verification without errors looks like the following sample output:

```
CertContext[0][0]: dwInfoStatus=102 dwErrorStatus=0
CertContext[0][1]: dwInfoStatus=10c dwErrorStatus=0
```

The certificate verification process retrieves any CRL that is necessary to verify the certificates. After the verification process, cached copies of the CRLs are available in the temporary Internet Explorer folder on the client.

**Install the Certificate on IntermediateCA1**

After you have verified that the certificate trust chain can be properly built, install the CA certificate.

1. Log on to the IntermediateCA1 computer as either a CA administrator or local administrator.

2. Click **Start**, point to **Administrative Tools**, and then click **Certification Authority** to start the Certification Authority MMC Snap-in.

3. In the console tree, right-click **IntermediateCA1**, point to **All Tasks**, and then click **Install CA Certificate**.

4. Insert the **Transfer-IntermediateCA** floppy disk into the floppy disk drive.

5. Browse to the floppy drive, click **IntermediateCA1.p7b**, and then click **Open**.

6. (Optional) If the parent CA certificate has not been previously trusted, you may receive a message that says that the root certificate is not trusted. Click **OK**, and then install the root CA certificate to the trusted root CA certificate store on the local computer.

   The root CA of the certificate chain must be locally trusted so that the CA service can start. For more information, see "Import the root CA certificate and CRL to the intermediate CA," earlier in this document.

7. In the console tree, right-click the name of the stand-alone offline intermediate CA, point to **All Tasks**, and then click **Start Service**.

   This brings the stand-alone offline intermediate CA into an operational state by starting the CA service. You can also type **net start certsvc** at a command prompt.

Note that, after the CA has been started successfully, the icon that displays the CAs operational state turns into a green check mark.

8. On the **File** menu, click **Exit** to close the Certification Authority MMC.

9. Log off of the IntermediateCA1 computer.

Continue the installation procedure by following the steps in the Installation cleanup section in this document.

**Install the Certificate at IntermediateCA1**
To install the certificate at a command prompt:

1. Log on to the computer as a local administrator with CA Management permissions.

2. At a command prompt, type

   **certutil.exe –installcert A:\IntermediateCA1.p7b**

   | **Note** |
   | --- |
   | If you used a .cer file instead of a p7b file and you receive a warning message at the end of the output such as "A certificate chain was processed, but terminated in a root certificate which is not trusted by the trust provider. 0x800b0109 (-2146762487)," it is possible that the parent CA certificate has not been imported into the local computer certificate store or that the parent CA certificate has been saved to the wrong store. To correct this error, see "Import the root CA certificate and CRL into an intermediate CA from a batch file," later in this document. To resolve this behavior, you can also use a PKCS#7 file (a .p7b file) that includes the entire certificate chain instead of a binary certificate file. |

3. To start the CA service, at a command prompt, type **net start certsvc**.

**Installation Cleanup**
For security reasons, it is recommended that you delete the certificate request file on the **Transfer-IntermediateCA** floppy disk that you used to generate the CA certificate.

**Configure IntermediateCA1**
After you complete the steps in the previous sections to configure the offline CA, you can complete the remaining steps for IntermediateCA1 with a batch file script. The difference between the root CA configuration and the subordinate CA configuration is the validity period for issued certificates. To configure the subordinate CA:

1. Log on to the IntermediateCA1 computer as local or CA administrator.

2. Start a text editor, such as Notepad.

3. In this document, copy the sample text in Sample script to configure IntermediateCASample to a new document in the text editor.

4. Save the text file as **%temp%\subcacfg.cmd**.

5. Close the text editor.

6. At a command prompt, type **%temp%\subcacfg.cmd**, and then press ENTER.

**Include CA Policy in Certificate Requests**
The option around the CA issuer and application policies is a choice at which CA level the policy is applied. If you plan to configure a issuer statement at a CA, you must configure the parent CA to add information about the CA policy to its issued certificates. See "Sample CAPolicy.inf file for the IntermediateCA1 later in this paper.

If this configuration step is skipped, an intermediate CA will not accept or allow CA certificate policies from its subordinate CAs.

If required, you can apply this configuration step at the time when a issuer or application policy needs to be included in a certificate request from a subordinate CA.

To include a policy in issued certificates, enter the following commands at a command prompt:

**certutil -v -setreg policy\EnableRequestExtensionlist "+2.5.29.32"**

**certutil –shudown**

**net start certsvc**

You can disable the setting with **certutil -v -setreg policy\EnableRequestExtensionlist "-2.5.29.32"**

**certutil –shudown**

**net start certsvc**

#### Verify the IntermediateCA1 Configuration

After you use the steps in the previous sections, ensure that the CA is configured properly and ready for production operations. You should apply the verification steps as described in the following sections in this document because they apply to the intermediate CA the same way as for a root CA:

- Verify the root CA configuration
- Verify the CorporateRootCA CRL and AIA configuration
- Verify the published CRL

#### Finalize the CA Configuration

After you apply the steps from the previous sections in this document, the intermediate CA is operational and ready to issue certificates.

If you installed a Windows 2000 CA instead of a Windows Server 2003 CA, you should apply the additional configuration steps that are explained in the Disable issuer name and issuer serial number section in this document.

# Stand-alone Offline Intermediate CA (CorporateSub2CA)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The sample PKI topology design has exemplified two separate offline intermediate stand-alone CAs that provide organizational and security flexibility. The CorporateSub2CA setup is similar to the steps that are outlined for IntermediateCA1. The only differences will be some of the configuration parameters, depending on the guidelines that are specified in the organizational CPS.

To apply the sample design, install IntermediateCA2 using the identical steps for installing IntermediateCA1.

# Online Enterprise Issuing CAs (CorporateEnt1CA)

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The online enterprise issuing CA is also referred to as "CorporateEnt1CA" in this document. The purpose of an enterprise CA is to automate certificate enrollment without making compromises in the security and authentication of issued certificates.

Depending on the PKI topology that you implement, this CA might have at least one parent CA or, in a single tier topology, it might be a self-signed CA.

When you set up either a Windows Server 2003 Server enterprise CA or Windows 2000 Server enterprise CA, it is important to note that domain controllers in an Active Directory environment automatically request certificates when an enterprise CA) becomes available in the forest.

The automatic certificate request behavior is different between computers that are running Windows 2000 or Windows Server 2003 domain controllers. Windows 2000 domain controllers immediately start requesting certificates when the enrollment service is available; however, a Windows Server 2003 domain controller requests certificates according to the Autoenrollment configuration in Group Policy settings. Computers that are running Windows Server 2003 domain controllers and computers that are running Windows XP request certificates, according to the configuration of the Group Policy object (GPO). You must enable the request processing manually, by using the appropriate domain Group Policy. Note that domain controllers have their own GPO settings which are separate from the rest of the computers in the domain. To add an automatic request setting for Windows 2000 domain controllers, create a new request object in the following GPO path using the following procedure.

1. Click **Start**, click **Run**, type **mmc**, and then press ENTER

2. In the console tree, double-click the domain controller policy that you want to work with, double-click **Computer Configuration**, double-click **Windows Settings**, double-click **Security Settings**, double-click **Public Key Policies**, and then click **Automatic Certificate Request Settings**.

3. In the details pane, right-click a blank area, point to **New**, and then click **Automatic Certificate Request**.

4. Follow the instructions in the Automatic Certificate Request Setup Wizard.

For more information about configuring auto-enrollment for Windows Server 2003 domain controllers, see Certificate Autoenrollment in Windows XP on the Microsoft Web site.

**Enterprise CA Installation Prerequisites**

The following items are required to correctly install and configure the online enterprise CA:

- The CPS that has all of the parameters that are specific to your organization. For more information, see "Certificate practice statement," earlier in this paper.

- The Windows Server 2003, Enterprise Edition media

- Appropriate hardware and a floppy disk drive

- A floppy disk that is labelled **Transfer-EnterpriseCA**

- The computer must be joined to a domain in the appropriate Active Directory forest

- Local administrator, Enterprise Administrator and Root Domain Administrator permissions

- File and print sharing is enabled on the CA. This is required to run the certification authority MMC snap-in on the CA.

- If you are using a Windows 2000 Active Directory forest, the domain controllers must have Service Pack 3 (SP3) applied for computers that are running Windows 2000. It is also required to upgrade the schema to Windows Server 2003 functionality as previously described. For more information, see "HOW TO: Raise the Domain Functional Level in Windows Server 2003" on the Microsoft Web site.

You should also ensure that the following tasks have been completed:

- A server running Windows Server 2003, Enterprise Edition should be set up and available to be used as the enterprise CA.
- The server should have the latest service packs available and installed, if appropriate.
- The server that is hosting the CA service must be joined to a domain in the Active Directory forest.

It is possible to install a CA as a multiservice server or as a domain controller, but this is not recommended for security reasons. A CA has high security requirements and should be accessed and maintained only as separate resource.

**Prepare the Active Directory Environment**
You can operate an Windows Server 2003 enterprise CA in a Windows 2000 environment if all domain controllers in the Active Directory are running Windows 2000 SP3 or later. Windows 2000 SP3 domain controllers are the minimum version required to support the schema upgrade for version 2 templates, as described in the previous chapter.

The schema upgrade, which does not cause a full replication in a Windows Server 2003 domain environment, is required to add additional template information, key archival information, cross-certificate objects, and object identifier (also known as OID) objects in the directory.

It is assumed that the schema upgrade is part of the organization change and management process as it can have an overall effect on the production environment.

To upgrade the schema:

1. Log on to the schema master domain controller as the Schema Administrator.

   For more information regarding schema administrator roles, see "HOW TO: Find Servers That Hold Flexible Single Master Operations Roles" in the Microsoft Knowledge Base.

2. Make the Windows Server 2003, Enterprise Edition, installation media available to the server and, at a command prompt, type the following command, and then press ENTER.

   **adprep /forestprep**

   The Adprep.exe file is available from the \i386 directory on the original Windows Server 2003, Enterprise Edition, installation media: After you use this command, you receive an output that contains logging information about the schema-upgrade process. The output ends with the message "Adprep successfully updated the forest-wide information."

To enable the domain to benefit from the Windows Server 2003 schema extensions, you must perform the following procedure on each domain in the forest.

1. Log on to the domain as Domain Administrator.
2. At a command prompt, type:

   **adprep /domainprep**

3. Repeat the previous two steps for each domain in the forest.

Depending on the replication schedule, the time required to apply the schema changes at each domain controller in the forest will vary.

**Domain Membership**
An enterprise CA can enroll any user or computer in the forest. To enable the CA to issue certificates to users and computers that are members of other domains than the domain where the CA is installed, see the Windows Server 2003 Server Help files or the following articles in the Microsoft Knowledge Base:

- Windows 2000 Certification Authority Configuration to Publish Certificates in Active Directory of Trusted Domain on the Microsoft Knowledge Base

- Enterprise CA May Not Publish Certificates from Child Domain or Trusted Domain on the Microsoft Knowledge Base

If the organization's Active Directory consists of several domains, you should plan where to put the enterprise CAs. There are three different approaches that you can take. Note that you must evaluate these approaches according to the requirements and the Active Directory design. Enterprise CAs can be:

- Installed into each production domain.

- Maintained in a separate PKI domain.

- Maintained as members of the forest root domain.

When you run enterprise CAs as members of the forest root domain, isolation and logical grouping is ensured but might not be accepted by the forest root administrator group because of security considerations. All approaches are valid, but must be carefully considered for your environment.

**Retrieve the Certificate and its CRL from CorporateRootCA and IntermediateCA1**

| **Note** |
| --- |
| If you implement a single-tier topology, the steps in this section do not apply to your environment. |

Both the certificate and CRL for all nodes in the PKI hierarchy are required during certificate validation.

Because all parent CAs of an issuing enterprise CA might be disconnected from the network (as in this sample scenario), you cannot automatically retrieve the CA certificates and the latest CRLs through the network when required. Because of this, you must make the CA certificates and the most current CRL available from all parent CAs before you can set up CorporateEnt1CA.

If the certificate of CorporateRootCA is not available on the **Transfer-RootCA** floppy disk, perform the steps that are outlined in "Obtain the certificate and its CRL from CorporateRootCA," earlier in this document.

You can retrieve the certificate and CRL for illustration purposes from IntermediateCA1 through Web enrollment support on the CA. Internet Information Services is not required to retrieve the CA certificate and CRL from the root CA as they may copied from the %systemroot%\system32\certsrv\certEnroll\ path on the local system. To retrieve the CA certificate and CRL through the web pages, perform the following steps:

1. Log on to the IntermediateCA1 computer.

   If the CA is accessible from the network, you can use any other computer to download the CA certificate and the CRL; however, these tasks require an interactive local logon because the CA is disconnected form the network.

2. Click **Start**, and then click **Internet Explorer**.

3. In **Address**, type **http://localhost/certsrv**.

   Localhost is an alias name for your current server. The Welcome page from IntermediateCA1 Web enrollment support is displayed.

4. Click **Download a CA certificate, certificate chain or CRL**

5. Insert the **Transfer-IntermediateCA** floppy disk into the disk drive.

   Next, download the CA certificate chain.

6. Click **Download CA certificate chain**, and then click **Save**.

7. In **Save As**, type a file name (for example, type a:\IntermediateCA1.p7b), and then click **Save**.

Next, download the latest base CRL:

8. Click **Download latest base CRL** and then click **Save**.

9. In **Save As**, type a file name (for example, type a:\IntermediateCA1.crl), and then click **Save**.

   If applicable to your environment, download the latest delta CRL:

10. If applicable, click **Download latest delta CRL**, and then click **Save**.

11. In **Save As**, type a file name (for example, type a:\IntermediateCA1+.crl), and then click **Save**.

12. Remove the **Transfer-IntermediateCA** disk from the disk drive.

**Distribute a Root CA Certificate with Group Policy**

| **Note** |
| --- |
| If you implement a single-tier topology, the steps in this section do not apply to your environment, because an enterprise CA that is configured as the root CA automatically publishes certificates to Active Directory. |

If a single tier topology is implemented, the steps in this section do not apply because an enterprise CA that is configured as the root CA would publish its certificate automatically into Active Directory.

The validation of certificates requires the availability and explicit trust of the root CA certificate that has issued certificates within the certificate trust path. The root CA certificate provides the trust anchor from which PKI hierarchies are derived.

The easiest way to provide clients with the root CA certificate is through group policies. Trust of root certification authorities should be managed and controlled through Group Policy whenever possible. When a root CA certificate has been added to the Trusted Root certification authority's container that is part of the domain security settings, clients that are member of the domain will receive the root certificate automatically.

| **Caution** |
| --- |
| You must not publish subordinate (or intermediate) CA certificates through either the trusted root certification authorities in Group Policy or an enterprise trust. If a subordinate CA certificate is part of this list of certificates that is published with Group Policies, a Windows client will not build a certificate chain correctly. |

The security portion of a domain Group Policy setting distributes the list of trusted root certificates to all computers that are Active Directory-aware and members of a domain. Because the root certificate becomes part of the computer policy, all domain-users inherit the root certificate trust from the computers to which they logon.

| **Note** |
| --- |
| It is recommended that you make a copy of the default domain policy and use a new policy specific to PKI to administer PKI policy in the domain. Modification of any of the default values requires more in-depth planning regarding certificate trusts and name constraints. Group Policy displays a slightly different view in a Windows 2000 environment. but it provides the same functionality. |

To add the root CA certificate of CorporateRootCA through Group Policy to the list of trusted root CA certificates on all computers in the domain, use the following procedure.

1. Log on as the domain administrator to the domain where you want to deploy the root CA certificate.

2. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Domain Security Policy**

   You can also type **Dompol.msc** at a command prompt and then press ENTER.

3. In the console tree, double-click **Default Domain Policy**, double-click **Computer Configuration**, double-click

**Windows Settings**, double-click **Security Settings**, and then click **Public Key Policies**.

4. Right-click Trusted Root Certification Authorities, click Import, and then click Next.

5. Insert the **Transfer-RootCA** floppy disk, and then click **Browse**.

6. Click **Open** to select the certificate file.

7. Click Place all certificates in the following store. After the certificate is placed in the Trusted Root Certification Authorities container, click Next.

8. After a report displays the options that you have selected in the wizard, click **Finish** to import the certificate.

You can configure additional PKI trust properties by using the Trusted Root Certification Authorities Properties page. To set these properties, right-click **Trusted Root Certification Authorities**, and then click **Properties**.

If you disable the default third-party root CAs and enterprise root CAs option, there may be unintended effects when you try to gain access to applications such as SSL secured Web sites on the Internet, and so on.

When only enterprise root certificates are trusted, only CA certificates that are installed in the following container in the configuration partition of Active Directory or through Group Policy are trusted: Domain Name\Configuration\Services\Public Key Services\Certification Authorities.

See the Sites and Services MMC to verify the certificates. You must make the services node visible to see the Services container.

**Import ParentCA Certificates and CRLs into Active Directory**

The following section outlines the procedures for publishing the CA certificates and CRLs of offline CAs (CorporateRootCA, IntermediateCA1, and CorporateSub2CA) into Active Directory.

| **Important** |
| --- |
| Do not use the information in this section if you are implementing a single-tier CA because a single tier enterprise root CA publishes its root Certificate automatically into Active Directory. |

You must import offline CA certificates and their CRLs. Note that this import procedure for the publication of CA certificates must be repeated every time that the offline CA's certificate is renewed.

Also the CRL must be imported according to the CRL publication strategy. Every time that a new CRL is published by the parent CA, you must repeat the publication procedure. To prepare the environment for the enterprise CA installation, you must register the parent CA certificates and CRLs in Active Directory. For information about how to do this, see the following section in this document.

**Get CA Sanitized Name and DNS Name**

The configuration information that is part of the CRL in Windows Server 2003 is different from the configuration information that was included with versions of the Windows 2000 family. A CRL that was published by a Windows 2000 CA does not contain information about the publication location of the CRL. (For more information, see Verify the published CRL in this document.) However, you must know where you want to save the CRL during the import process into Active Directory.

A Windows Server 2003 CA stores this information in the optional **Published CRL Locations** attribute of a CRL if the corresponding CRL property was set. When you perform a CRL import procedure on a computer that is running a Windows 2000 operating system, you must manually set the path of the CRL publication location.

To display the CAs computer name and the CAs sanitized name, at a command prompt, type the following and then press ENTER.

**certutil –cainfo**

You can also use this command on a Windows 2000 CA that should be published to Active Directory. It is important that you note both the sanitized CA short name (DS name) and DNS name for each Windows 2000 CA. A sample output from the command is similar to the following output, where the items in bold are placeholders:

```
Exit module count: 1
```

```
CA name: IntermediateCA1
Sanitized CA short name (DS name): IntermediateCA1
CA type: 4 -- Stand-alone Subordinate CA
    ENUM_STANDALONE_SUBCA -- 4
CA cert count: 1
KRA cert count: 0
KRA cert used count: 0
CA cert[0]: 3 -- Valid
CA cert version[0]: 0 -- V0.0
CA cert verify status[0]: 0
CRL[0]: 3 -- Valid
CRL Publish Status[0]: 5
    CPF_BASE -- 1
    CPF_COMPLETE -- 4
Delta CRL Publish Status[0]: 0xe (14)
    CPF_DELTA -- 2
    CPF_COMPLETE -- 4
    CPF_SHADOW -- 8
DNS Name: connoam-ca-00
Advanced Server: 1
CertUtil: -CAInfo command completed successfully.
```

If either the CRL path that is specified in the CRL or the path where the CRL is physically published is not the same, an issuer distribution point (IDP) intersection error might appear when the CRL is verified by a client.

**ImportCA Certificates and CRLs from CorporateRootCA and CoporateSubxCA**

In order to continue the installation procedure, you must import CA certificates and CRLs from parent offline CAs into Active Directory. In general, this procedure is a common operation when the CA environment is set up. CRL publication and distribution with Active Directory is critical for a successful PKI in your organization. Importing root CA certificates into the Active Directory directly is preferred over using Group Policy to distribute root CA certificates, as this method provides for root CA trust throughout the entire forest instead of individual domains.

Both, the CA certificate and CRL are written by the **Certutil.exe** utility into Active Directory. Note that **Certutil.exe** replaces the **Dsstore.exe** utility that is available in the Windows 2000 Resource Kit. For more information, see "HOW TO: Use the Directory Services Store Tool to Add a Non-Windows 2000 Certification Authority (CA) to the PKI in Windows 2000" on the Microsoft Knowledge Base and "The Dsstore Tool May Not Work If the NetBIOS Name and the DNS Domain Name Are Different" on the Microsoft Knowledge Base.

| ⚠**Caution** |
| --- |
| Do not use the **Dsstore.exe** utility that is included with the Windows 2000 Resource Kit to import CA certificates and CRLs to a Windows Server 2003 environment. |

To import CA certificates and CRLs from parent offline CAs into Active Directory, use the following procedure:

1. Log on to the CorporateEnt1CA computer as **Administrator** of the **Root-Domain** and also a member of the **Enterprise Administrators** group.

2. At a command prompt, use the following sample script to import CA certificates and CRLs from CorporateRootCA and CoporateSubxCA to Active Directory. Note that you must change the script so that both the CA certificates and CRLs correspond to your file names and CA names:

```
:
: Root CA certificates
:
certutil -dspublish -f concorp-ca-00_CorporateRootCA.crt RootCA
:
: Sub CA certificate
:
certutil -dspublish -f connoam-ca-00_IntermediateCA1.crt SubCA
:
: Root CA CRLs
: Since these are .NET CA CRLS that have the publication location as
: part of the CRL, the publication location is optional
```

```
:
:                                                    |-- publication
location ---|
:
certutil -dspublish -f CorporateRootCA.crl     concorp-ca-00
CorporateRootCA
:
: Sub CA CRLs
:
certutil -dspublish -f IntermediateCA1.crl     connoam-ca-00
IntermediateCA1
```

| **Note** |
| --- |
| The **–f** parameter is required because the container structure that is required to store the certificates and CRLs and might not exist. The **–f** parameter is not required if the container structure already exists. |

Deploying the root CA certificate with Group Policy or Active Directory is not the only way that you can provide clients with certificates of trusted root CA certificates. You can also deploy root CA certificates by using the following methods:

- The Internet Explorer Administration Kit (IEAK)
- A CAPICOM script (For more information, see "CAPICOM Reference" on the Microsoft TechNet Web site.)

**Set the Appropriate Permissions for Certificate and CRL access**

The CA administrator must ensure that any client in the PKI can gain access to both the CRL distribution points and CA certificate.

For HTTP URLs, it is recommended that Internet Information Services (IIS) be set up and configured to allow anonymous access to the URLs that are set as CRL distribution points in issued certificates. When you do this, clients, regardless of their primary authentication mechanism, can retrieve the CRL.

For URLs that point to a CRL object in Active Directory, the Everyone group has read access by default. To allow clients that are not Active Directory members to gain access to CRL objects in Active Directory, you must add the Everyone group with read permissions to the domain objects ACL. A more restrictive method that you can use to configure anonymous access is to replace the Everyone built-in account with the Anonymous built-in account.

For more information about how to add Read access permissions to clients that are outside of the forest, see the following Microsoft Web sites:

- "Anonymous Queries" in the Windows 2000 Resource Kit
- "How to Use the RestrictAnonymous Registry Value in Windows 2000" in the Microsoft Knowledge Base

**Publish CA Certificates and CRLs of CorporateRootCA and CoporateSub1CA**

Publishing the CRLs using HTTP requires a Web server, such as IIS. To set up IIS to provide the CRL publishing point using HTTP, use the following procedure:

1. Log on as a local administrator to the computer that has IIS installed.
2. Click **Start**, click **All Programs**, point to **Administrative Tools**, and then click **Internet Information Services (IIS)**.
3. Double-click the IIS server node, and then double-click **Web Sites**.

   The organizational Web site that will host the CRL is displayed.
4. Right-click your Web site, point to **New**, click **Virtual Directory**, and then click **Next**.

   The Virtual Directory Creation Wizard starts.

5. Type an alias name for your Web site (for example, type **PKI**), and then click **Next**.

6. Choose a path where the certificates and CRL will be stored on the a local storage device, and then click **Next**.

   In this example, the CA certificates and CRLs are stored in **C:\PKI**.

7. Select the **Read** check box, clear all of the other check boxes, and then click **Finish**.

   No other permissions are required.

8. Copy the CA certificate and CRLs from the **Transfer-RootCA** floppy disk into the C:\PKI folder. Repeat this step with the CA certificate and CRLs stored on the **Transfer-IntermediateCA** floppy disk.

| **Note** |
| --- |
| Make sure that the file names that are published with HTTP exactly match the CA certificate and CRL distribution point as defined as part of the CA configuration. If the file names do not match, clients will fail to retrieve the CRL with the URL that was specified as the CRL distribution point. |

**Verify that the Domain Controller Has Published the Certificates and CRL into Active Directory**

After you import the list of CA certificates and CRLs into Active Directory, verify that the data is published to the correct location.

| **Note** |
| --- |
| In a distributed environment, a delay occurs before any domain controller has received the certificates and CRLs through Active Directory replication. The delay will vary depending on the Active Directory environment configuration. |

You can verify that the CA certificates and CRLs are published to the correct location through the Active Directory Sites and Services MMC Snap-in.

1. Log on to the CorporateEntxCA computer as the administrator of the root domain and as a member of the Enterprise Administrators group.

2. Open the Active Directory Sites and Services MMC, by doing one of the following

   - To use the Windows interface, click **Start**, click **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.

   - To use a command prompt, at a command prompt, type **dssite.msc**.

3. Click **Active Directory Sites and Services**.

4. On the **View** menu, click **Show Services Node**.

5. Click **Services**, and then click **Public Key Services**.

6. Verify that the AIA container has all parent CA certificates, and then verify that the CRL container has the CRL objects for each CA.

   The AIA container has a flat object structure but the CRL distribution point container stores CRLs in a separate container for each CA.

If the AIA or CRL container do not have the proper objects, it may be necessary to manually republish the objects as described above.

To delete certificates or CRLs that are expired, revoked, or no longer required through the Active Directory Sites and Services MMC, click the certificate or CRL and then press DELETE. Keep in mind that clients will need CA certificates to verify end-entity certificates. Even a revoked or expired CA certificate can be required to verify a certificate chain!

**Verify That the CA Certificate and CRL Are Imported into Active Directory**

To ensure that the correct CRL and AIA information is published to Active Directory, verify the certificate of IntermediateCA1. If clients cannot retrieve a CRL from the CRLs distribution point or do not download a certain CA certificate from Active Directory, the client applications will generate an error when verifying a certificates status.

| **Note** |
| --- |
| If the verification procedure in this section does not work, it is important to investigate the cause of the failure. If the publication location in Active Directory is incorrect, correct the location by using the steps in the Publish CA certificates and CRLs of CorporateRootCA and CoporateSub1CA to a Web site section in this article. |

To verify that the CA certificate and CRL were imported correctly into Active Directory:

1. Log on to a computer that is member of the organization's Active Directory. The user account that you use can be a domain user account because any user should have read permissions to both the CA certificate and CRL in Active Directory.

2. Insert the **Transfer-IntermediateCA** floppy disk into the floppy drive

3. At a command prompt, type the following, and then click **Retrieve**:

   **certutil.exe -url a:\IntermediateCA1.cer**

   The **certutil.exe –url** command works with any X.509V3 certificate. You can also use this command to verify user certificates. You can perform CRL and AIA distribution point verification by using the Certutil.exe utility only on certificate files that are Distinguished Encoding Rules (DER) encoded (*.cer file). If you discover that the CRL distribution point is incorrectly configured, you must correct the issue in the CA configuration to which the CRL belongs. The CRL distribution point is stored as an attribute in every issued certificate. Also, there may be an impact on certificates that are already enrolled if the CRL is not correctly configured.

   When you click **Retrieve**, Windows performs a CRL or AIA download based on the certificate file that has been specified. The certificates name is displayed in **Certificate Subject**.

   Windows lists the certificate's CRL distribution points that are specified in the CRL distribution point extension of the certificate and displays the verification status. The following figure is an example of a successful verification procedure.



**Figure 8: URL Retrieval Tool**

   If the CRL verification procedure does not work, the CRL distribution point extensions may not be valid, or the locations of the CRLs have permissions set that do not allow the current user to retrieve the object.

4. In the **Retrieve** box, click **Certs (from AIA)**, and then click **Retrieve**.

   Since the root CA certificate has no CRL defined in this scenario, you receive a CRL status message saying No CRL. The missing CRL status is expected in this example, since the CRL distribution point was set to Empty in the root CAs CAPolicy.inf file.

**Template Upgrade from Windows 2000**

When you upgrade from Windows 2000 to a member of the Windows Server 2003 family, you must upgrade both the properties and security settings on existing version 1 (V1) templates.

To perform the upgrade to a Windows Server 2003, Enterprise Edition, CA environment, open the Certificate Templates MMC that is included with Windows Server 2003, Enterprise Edition, so that you can install and upgrade the template objects. It detects that templates are available from a previous Windows 2000 CA installation and automatically upgrades the templates.

The upgrade procedure also changes the permissions that are required for template administration. In Windows 2000, you must be a member of both the Enterprise Admins group and Root Domain Admins group to perform this operation. For more information about certificate templates, see "Implementing and Administering Certificate Templates in Windows Server 2003" on the Microsoft TechNet Web site.

**Prepare the CAPolicy.inf File for the Issuing CA**

If you apply an application or issuance policy at the level where end entity certificates are issued, you have precise control over policies. If more than one issuing CA exists in the PKI hierarchy, different policies can be applied to each CA to issue different certificate types. Nevertheless, policies can be applied at a parent CA level as well. Typically, in most deployments, policies are applied at an intermediate CA level than at the leaf node CA. When you define policies at a parent CA level, the policy applies to all of the subordinate CAs. This could have an impact on the CA topology design.

For an example of a CAPolicy.inf file with issuer policy, see "Sample CAPolicy.inf file for the IntermediateCA1" later in this paper.

**Install the Online Issuing Enterprise CA**

If you upgrade a Windows 2000 enterprise CA to a Windows Server 2003 CA, you must log on to the computer with an account that is member of both the root domain administrators group and the enterprise administrators group.

When you perform a clean installation of the first Windows Server 2003, Enterprise Edition CA in a new Windows Server 2003 forest, the installation account requires that you are a member of the Enterprise Admins group and the (root domain) Domain Admins group. After the first CA installation, (root domain) Domain Admins permissions are no longer required.

The installation procedure for an online enterprise CA is different form the installation procedure for the offline parent CAs. Use the following steps to set up CorporateEnt1CA:

1. Log on to the CorporateEnt1CA computer with Local Admin, Enterprise Admin, and (root domain) Domain Admin permissions.

   The installation of an enterprise CA requires that you be able to gain access to the Active Directory configuration container. During the CA installation procedure, the account used to install the CA also becomes a CA administrator account which is a role that can be delegated to other user accounts, as appropriate.

2. Do one of the following:

   - To use the Windows interface, click **Start**, point to **Settings**, click **Control Panel**, double-click **Add or Remove Programs**, and then click **Add/Remove Windows Components**.

   - To use a command prompt, click **Start**, click **Run**, type **cmd**, press ENTER, type **sysocmgr /i:sysoc.inf**, and then press ENTER.

3. Select the **Certificate Services** check box.

   To run Certificate Services, the following software components are required:

   - Certificate Services

- Internet Explorer
- (Optional) Certificate Services Web enrollment support
- (Optional) Internet Information Services for Web enrollment support

It is not recommended that you install any other Windows components on a Windows Server CA. If you install additional components, reliability or security of a root CA may be compromised if a secure configuration is required by the organization.

4. After you select the **Certificate Services** check box, you receive a warning that states that you cannot change the NetBIOS computer name after you install Certificate Services. (Note also that you cannot change the computer's membership to a domain or workgroup.) Click **Yes** to continue with the installation procedure, and then click **Next**.

| **Note** |
| --- |
| IIS is not a required component on an enterprise CA, but it might be necessary to have IIS available for certificate Web enrollment, depending on the enrollment method that is used for clients. Windows 2000 and Windows XP clients typically request certificates by using distributed COM (DCOM) or auto-enrollment instead of HTTP. For more information, see "Authentication and authorization" in this document. |

5. When you are prompted to choose the type of installation, do one of the following:

   - If you want this installation to be in a multi-tier PKI topology, click **Enterprise subordinate CA**.
   - If you want this installation to be an enterprise CA in a single-tier topology, click **Enterprise Root CA**. Also, verify that the **Use custom settings** check box is selected.

   In this scenario, the enterprise CA is installed as a subordinate CA. Select **Enterprise subordinate CA**.

   Note that, if the computer that is supposed to be an enterprise CA is a domain member, both the **Enterprise Root CA** and **Enterprise Subordinate CA** options may be unavailable. To find out why the options are unavailable, see the following table.

   **Table 20 Workarounds for Enterprise Root and Subordinate CA Unavailability**

| Reason | Workaround |
| --- | --- |
| No access to domain controller | The CA cannot gain access to the domain controller through the network. The **nltest.exe /dsgetdc** command might be helpful to ensure connectivity with your domain controller. To use Nltest.exe, install the support tools that are available on the Windows Server 2003 CD-ROM. For more information about nltest see Active Directory support tools in the Windows Server 2003 online help or on the Microsoft Web site. |
| Domain controller has not replicated | See "Verify that the domain controller has published the certificates and CRL in Active Directory" earlier in this article. |
| Configuration container does not exist | See "Verify that the domain controller has published the certificates and CRL in Active Directory" earlier in this article. |

   After you are done, click **Next**.

6. Do one of the following:

- If an HSM was not installed, click **Microsoft Strong Cryptographic Provider**.
- If an HSM was installed, you must click the CSP that was installed during the HSM setup procedure.

7. Select the **SHA-1** hash algorithm.

   SHA-1 is the most common and interoperable hashing algorithm that is used by programs and operating systems.

8. In **Key Length**, select the appropriate setting.

   In this example, set a key length of **2048** for CorporateEnt1CA.

9. Confirm that the **Allow this CSP to interact with the desktop** and the **Use an existing key** check boxes are cleared, and then click **Next**.

   If a CA is already installed on this server, the list of existing keys is created from the system certificates that are stored in the following registry key:

   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\SystemCertificates\My\Certificates**

10. Set the common name for this CA as specified in the CPS, and then click **Next**.

    Because this CA is an enterprise CA, the distinguished name suffix is predefined to use the namespace of the existing Active Directory's (forest) namespace. It is recommended that this value not be changed.

    The key pair is generated by the CSP and written to the local computers key store. If an HSM has been installed and selected, the key is generated in the HSM and stored accordingly. If you do not use an HSM, the key is generated by CryptoAPI and stored in the profile of the system account on the local computer. The length of time that is required to generate the key depends on both the size of key that is being generated and the CPU performance of the local computer.

11. Specify the location of the certificate database and the certificate database's log files, and then click **Next**.

    | **Note** |
    | --- |
    | It is a good practice to place both the certificate database and certificate database log on a separate volume from the partition on which Windows is installed for issuing enterprise CAs. This provides increased disk input and output performance and provides sufficient storage space for the database as it becomes larger. |

    Because this in an online enterprise CA, the shared folder that stores configuration information is optional. The purpose of the shared folder is to serve clients that do not receive the CAs certificates through the Group Policy object (GPO) or that are not able to retrieve the CA certificates from Active Directory or through Web enrollment support. Any client that has access to the shared folder can import the CA certificates into its certificate store. Depending on the shared folders name, a new share will be created on the CA server computer. The default name for the shared folder is \\Localhost\CAConfig. If you do not need to publish the CAs certificate and configuration with a shared directory, do not create a shared folder.

    If you are installing a CA in the same location as a previously installed CA, the **Preserve existing certificate database** option is enabled. Click this option if you want the new CA to use this database and preserve the certificates that are in the database; otherwise, the database will be deleted. Use this option only when you want to restore a CA from backup or for CA migration.

12. Click **Save the request to a file**, type a name for the request file, click **Next**, and then click **OK**.

    For example, you can type **a:\CorporateEnt1CA.req**.

    A subordinate enterprise CA needs to send the certificate request to its parent offline CA. Because the IntermediateCA1

computer is running without a network connection, the request file must be transferred on a floppy disk.

| **Important** |
| --- |
| Make sure that the **Transfer-CorporateEnt1** floppy disk is available before you proceed to the next step. If Windows cannot access the floppy disk, you receive an error message appears and the CA setup process stops. Before you can reinstall the CA, you must uninstall the **Certificate Services Web-Enrollment Support** component if you selected it during setup. |

The Windows Component Wizard continues the certificate services configuration.

When the Windows Component Wizard completes the certificate services configuration, you may need to provide the installation media to finish the installation procedure.

13. When the CA certificate has obtained a signed subordinate CA certificate from its parent CA, the installation wizard displays a final message that indicates that the installation procedure is finished. Before you click **OK**, verify that the **Transfer-CorporateEnt1** floppy disk is available to save the request file.

    Click **OK** to continue the installation.

14. (Optional) If IIS is installed as part of this configuration but ASP pages are not enabled by default, the CA setup procedure asks whether you want to automatically enable the ASP pages. Click **Yes** to enable ASP pages, because they are required for Web enrollment services.

    If you click **No** because you do not want to enable ASP pages, you can enable ASP pages by using the **certutil –vroot** command at a later time

15. Click **Finish** to complete the installation procedure, and then click **Close** to close **Add or Remove Programs**.

16. Remove the **Transfer**-CorporateEnt1 floppy disk from the floppy disk drive, and then take the floppy disk to the parent CA computer (IntermediateCA1).

    If this installation procedure is for an enterprise root CA, continue to "Configure the enterprise online CA" in this document.

**Certificate Request Processing with the Offline Parent CA (IntermediateCA1) Through Web Enrollment Support**

The following procedures are provided as an example and not necessarily as a security best practice.

You do not need IIS on a CA unless you need client support for enrollment on Windows clients earlier than Windows 2000 or non-Windows clients.

In a three-tier topology, the enterprise CA certificate request that is stored on the floppy disk must be issued by the IntermediateCA1 computer. In a two-tier topology, the certificate request must be issued by the root CA.

**To Request a CA Certificate Through Web Enrollment Support**

1. Log on to the CA computer that will issue the enterprise CA certificate with an account that has certificate enrollment permissions.

2. On the offline subordinate CA computer, open Internet Explorer.

3. In **Address**, type **http://localhost/certsrv** and then press ENTER.

4. On the **Welcome** page, click **Request a certificate**.

5. On the **Request a Certificate** page, click **Advanced certificate request**.

6. On the **Advanced Certificate Request page**, click **Submit a certificate request by using a base-64-encoded CMC or PKCS #10 file** or **Submit a renewal request by using a base-64-encoded PKCS #7 file**.

7. Open a text editor, such as Notepad, and open the CorporateEnt1CA request file.

8. On the **Edit** menu, click **Select All**, and then, on the **Edit** menu, click **Copy**.

9. In Internet Explorer, on the **Submit a Certificate Request or Renewal Request** page, in the **Saved Request** box, right-click a blank space, click **Paste**, and then click **Submit**.

   The request is submitted to the CA and remains in the Pending Requests folder. Note the request ID number and the time that are supplied by the Web page. This information is helpful when you use the procedure to approve the request (by request number) and retrieve the certificate (by request time).

10. Click **Home**.

| **Note** |
| --- |
| A request must be retrieved by the same user account on the same computer from which it was submitted. The Web page uses a browser cookie to identify the pending request. If browser cookies are blocked or if you use a different computer, retrieve the certificate directly from the CA by using the Certification Authority MMC snap-in. For more information, see "Export the offline intermediate certificate at the root CA" in this document. |

**To Issue the Pending Request With the Certification Authority MMC**

1. Open the Certification Authority MMC.

2. In the console pane, double-click **Certification Authority (***LocationOfCA***)**, double-click the CA, and then click **Pending Requests**.

3. In the details pane, right-click the request, and then click **Issue**.

**To Issue the Pending Request Through Web Enrollment Support**

1. Open Internet Explorer.

2. In **Address**, type **http://localhost/certsrv** and then press ENTER.

3. On the **Welcome** page, click **View the status of a pending certificate request**.

4. On the **View the Status of a Pending Certificate Request** page, click the request you want to issue.

   If there is more than one certificate available, click the certificate that corresponds to the time that you sent the request to the CA.

5. On the next page, choose a format for the newly issued certificate.

   In a homogenous Windows environment, it is recommended that you use the DER-encoded format.

   Unless the root CA has been previously installed on the computer account where the enterprise CA is installed (Group Policy, and so on), the root CA certificate must be trusted before the enterprise subordinate CA can start.

6. Click **Download certificate chain**, save the output to a .p7b file, and then save the file on a floppy disk.

   For example, you can save the file as CorporateEntCA.p7b.

| **Note** |
| --- |

There is no sensitive data on the floppy disk. A CA certificate is public information, however, the associated CA private key has been generated in the CA (or HSM). The private key does not leave the computer on which the certificate request was created.

**Verify the EnterpriseSub1CA Certificate**

To ensure that the certificate that you issued for the CorporateEnt1CA computer has the correct certificate properties, you should verify the certificate. Open the certificate and examine its certificate properties. To view the certificate, double-click the certificate file in Windows Explorer or use the Certificate Manager MMC. Make sure that the validity time, the key length, and certificate policies are shown correctly.

**Install the Certificate at the CorporateEnt1CA Computer**

The installation procedure for the enterprise CA certificate is very similar to the installation procedure for the IntermediateCA1 computer.

- Insert the floppy disk that has the certificate file for the parent CA certificates into the floppy disk drive, and then, at a command prompt, type the following command, and then press ENTER:

  **certutil.exe –installcert** *CACertFile*.**p7b**

For more information, see "Install the certificate at IntermediateCA1 at a command prompt," later in this paper.

**Verify the CorporateEnt1CA Trust Chain**

Only a valid trust chain ensures that the CA will operate as expected. If the trust chain is not correct or if the issuing CA cannot build the chain and check revocation status of parent certificates, the issuing CA will report and error when attempting to start. If the CA certificate verification fails, the following error message appears: The revocation function was unable to check revocation because the revocation server was offline (0x80092013).

To verify the trust chain of the CorporateEnt1CA computer in advance, at a command prompt, type run **certutil –verify** *CertificateName*.**crt**, and then press ENTER.

It is also recommended that you validate the CA certificate's CRL distribution points. To do this, at a command prompt, type the following command, and then press ENTER:

**certutil –URL certificate-name.crt**

The output of the script looks similar to the following figure:




**Figure 9: Validated CRL Distribution Points**

To view the CRL, in the **Url** column, double-click the CRL URL.

**Note**

You can verify a CRL or AIA according to the URL that is specified in a certificate. If a certificate and a URL are both provided, the CRLs and AIAs for all URLs that are specified in the certificate are requested and, additionally, the information from the location that is specified in **URL to download** is received.

For troubleshooting purposes, you may also be required to download the CRL to the local computer to examine the CRL. This procedure is easy for CRLs that you can gain access to by using HTTP, but this procedure is more difficult for CRLs that do have only an LDAP CRL distribution point.

You can download a binary copy of a CRL from an LDAP CRL distribution point by using the **Certutil.exe** command. To do this, type the following at a command prompt: Note that you should replace the items that are in italic text with your CRL destination and that you should verify that the correct search filter is added after the question mark.

**certutil –store –split ldap:///CN=** *CorporateRootCA* **,CN=** *RootCA* **,CN=CDP,CN=Public Key Services,CN=Services,CN=Configuration,DC=** *contoso* **,DC=com?** *CertificateRevocationList* **?base? objectClass=cRLDistributionPoint**

After you run the command, a new CRL file is available in the directory you used when you ran the **Certutil.exe** command. The file is named Blob0_0.crl. You can open the .crl file through Windows Explorer and display it as you would any other .crl file.

**PKI Health Tool**

You can also use the PKI Health tool (Pkiview.msc) to verify the CRL and AIA location. This tool is available from the Microsoft Windows Deployment and Resource Kits Web site.

The PKI Health tool is of value to CA administrators who maintain the enterprise. The tool enumerates all certification authorities (CAs) that are associated with the current forest, and then the tool displays status properties that are associated with those CAs.

**Configure the Enterprise Online CA**

The CorporateEnt1CA computer configuration is similar to the parent subordinate CA configuration, and you can quickly create this configuration with a batch script. The difference between the parent CA configuration and the CorporateEnt1CA computer configuration is the validity period for issued certificates—they are controlled by templates in the enterprise CA instead of by the registry. Further, delta CRLs are enabled for the enterprise CA. Delta CRL publishing is enabled by default with Windows 2003 Server. To configure the enterprise CA:

1. Open Notepad.
2. In this paper, go to Sample script to configure the EnterpriseSubCA and copy the script to the clipboard.
3. In Notepad, paste the script into a new file.
4. Save the file as **%temp%\entcacfg.cmd**, and then close Notepad.
5. At a command prompt, type **%temp%\entcacfg.cmd** and press ENTER.

| **Important** |
| --- |
| If you study the script that configures the enterprise CA, you may notice that the **LDAP CRL publication** property is different than the properties for the offline CAs. This behavior occurs because an online CA has the capability to publish its CA certificate and the CRL automatically to Active Directory. Therefore, it is recommended that you allow the CA to update the information that is stored at the CRL and AIA distribution point. An offline CA cannot use this setting, because it cannot reach a domain controller by using the network. Because of this, no publication is configured for either CRL or AIA CRL distribution points that are configured with offline CAs. |

**Verify the EnterpriseSubCA Configuration**

The next sections help you ensure that the CA is correctly configured and is ready for production operations. It is recommended that you apply the verification steps that are described in the previous sections in this document:

- Verify the root CA configuration
- Verify the CorporateRootCA CRL and AIA configuration
- Verify the published CRL

An enterprise CA automatically publishes the CA certificate and CRL into Active Directory. Because of this, you should also

verify the CRL's availability in Active Directory. It is recommended that you use the **certutil –URL** command or the PKI Health tool that is described in the previous section.

| **Note** |
| --- |
| The output samples that are mentioned in the previous sections in this document are not the same as the IntermediateCA1 configuration. Verify the appropriate parameters according to the EnterpriseSubxCA configuration. |

# Certification Authority Maintenance

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

CA maintenance and monitoring is an ongoing task after you set up the CA environment.

Some of the most important procedures for CA maintenance are:

- Correct configuration of the CAs
- CRL publication of offline CAs
- CRLs that are not manually published
- Renewal of CA certificates assigned to CA operations
- Backup and recovery

You can maintain the CA that is connected to the network locally or you can maintain the CA that is connected to the network through a remote connection; however, CA maintenance and administration tools are designed to work best for local operations because the CA administration is a sensitive operation and should be kept as secure as possible.

If you want to use the Certificate Services MMC for remote administration, for the appropriate steps to make the CA remotely accessible, see "Users Allowed to Manage the CA Cannot Access It Remotely" in the Microsoft Knowledge Base.

For more information about CA operations and custom configuration, refer to the Windows Server 2003 PKI Operations Guide on the Microsoft Web site.

**Best Practices for CRL Publication**
The following section provides information and best practices for managing and publishing certificate revocation lists.

**CRL Partitioning**
Administrators may renew an issuing CA with a new key to partition the CRL. When a new key and certificate are generated, the CA uses the new key as well as any unexpired previous keys that correspond to previous certificates when generating revocation information. Therefore, a CA may be using multiple keys at the same time and therefore publishes multiple CRLs that correspond to those keys. You can see multiple valid certificates that are assigned to the CA in the Certification Authority MMC if you click the General tab of the CA properties.

You can also determine the renewal status of the CA if you examine the CA certificate. The CA version extension identifies how many times a CA has been renewed and how many times it has been renewed with a new key. The following figure displays a CA certificate that has been renewed three times. Note that a new key was issued with each renewal, which is why **CA Version** is 3.3.




**Figure 10: Renewed CA Certificate**

After a CA is renewed with a new key, only the new key is used when signing new certificates. The unexpired previous keys continue to be used to sign CRLs for certificates that were signed with the previous keys. Therefore, a CA may publish multiple CRLs at the same time, each using a different key. This method of CA renewal may be an ideal method for CRL size control and effective CRL partitioning with the CA.

**Automatic RootCA Cross-Certificate Generation**

On computers that are running a member of the Windows Server 2003 family, Microsoft root CAs can automatically issue and publish cross certificates for a root CA that has been renewed. For example, when a Windows Server 2003 root CA is renewed with a new key, the root CA cross-certifies the renewed root CA certificate and it is considered a qualified subordinate to the earlier root CA certificate. This functionality is important if you have an operational root CA that is trusted by other organizations clients, bridge CAs, or if it is cross-certified by other organizations.

To force the root CA to use the CrossCA certificate template, at a command prompt, type the following command, and then press ENTER. If you do not use this command without this setting configured, the CA does not use the CrossCA certificate template (even if it is available). Instead, it generates a certificate using predefined extensions in the registry:

**certutil -setreg ca\CRLFlags +CRLF_USE_CROSS_CERT_TEMPLATE**

To disable automatic cross CA certificate generation, at a command prompt, type the following command, and then press ENTER:

**certutil -setreg ca\CRLFlags +CRLF_DISABLE_ROOT_CROSS_CERTS**

To force the root CA to use the CAExchange certificate template when generating CA encryption certificates on demand, use the following command. Without this flag, the CA uses the CAExchange certificate template, when available, and generates a certificate without the template by using pre-defined extensions.

**certutil -setreg ca\CRLFlags +CRLF_USE_XCHG_CERT_TEMPLATE**

**Certification Authority Backup and Recovery**

The CA stores information about certificates in its database and the log files that are bound to the database. The CA has to be able to gain access to its certificates and keys that are stored in the local computer's certificate store or on a hardware device. CA configuration information is stored in the registry.

You can back up the database and the log files only by using the **certutil –backup** command. You can individually back up the CA certificate and the keys by using the **certutil –backupkey** command. You can archive the database by using the **certutil –backupdb** command. These backup procedures are appropriate for a restore operation that repairs a damaged CA, assuming that the CA is correctly configured. However, neither of these commands will back up any of the CA configuration or role separation information in the registry. To back up the CA, including its configuration, use the **ntbackup backup systemstate** command.

| Note |
| --- |
| Backing up a CA includes the private key that is owned by the CA. The private key is the most sensitive CA information option. It becomes part of your backup data if you run **certutil -backupkey** or if you perform a system state backup. Handle the backup data with caution at all times. Store it securely if the private CA key is part of your backup. The organization's certificate policy and security policy must cover the handling of backup media. |

A backup procedure may not work as expected, so it is recommended that regularly scheduled tests of restore operations be performed. For more information about CA backup, see the following Microsoft Knowledge Base articles:

- "How to Use Key and Certificate Backup/Restore Utility" on the Microsoft Knowledge Base.
- "Certificate Server Does Not Create Backups of Installed Keys" on the Microsoft Knowledge Base. (This article applies only to Windows 2000 Server.)

**Repair the Certificate Store**

When you restore a CA that maintains its CA certificate with a software CSP on a different computer, you must repair the CA configuration to allow the CA to gain access to the original CA certificate. When the CA software, the database, and CA configuration are restored, you need to install the CA certificate to the local computer's certificate store. You can also use the following command to share an HSM across multiple computers:

**certutil -addstore my CA01.Contoso.com_CARoot.crt**

After you use this command, use the following command to determine the certificate hash:

**cerutil CA01.Contoso.com_CARoot.crt | findstr /c: Key Id Hash(sha1)**

When you use this command, it displays the CACertSHA-1Hash value. Take the hexadecimal string representing the SHA-1 hash (known as the *thumbprint property* of a certificate) and use it as a parameter in the next command. For example, you might type the following command:

**certutil -repairstore my "ea c7 7d 7e e8 cd 84 9b e8 aa 71 6d f4 b7 e5 09 d9 b6 32 1b"**

# Appendix A: Directory Objects

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The various PKI-related containers such as CAs, enrollment services, templates, object identifiers (also known as OIDs), AIA, and CRL distribution points are created when you set up the forest for the first time with the first enterprise CA. The permissions on the objects are also set at that time.

**Directory objects that are created by an enterprise CA**
Installing an enterprise CA creates the following objects:

- Enrollment Services object (includes CA certificate) – under CN=Enrollment Services,CN=Public Key Services,CN=Services,CN=Configuration,DC=
- Trusted root CA object (includes CA certificate) – CN=Certification Authorities,CN=Public Key Services,CN=Services,CN=Configuration,DC=
- AIA object (includes CA certificate) – under CN=AIA,CN=Public Key Services,CN=Services,CN=Configuration,DC=
- KRA object (no significantly sized attributes) (Windows Server 2003 only) – under CN=KRA,CN=Public Key Services,CN=Services,CN=Configuration,DC=
- CRL distribution point container (no significantly sized attributes) – under CN=CDP,CN=Public Key Services,CN=Services,CN=Configuration,DC=
- CRL distribution point object (includes CRL) – under = CN=Computer,CN=CDP,CN=Public Key Services,CN=Services,CN=Configuration,DC=?

The installation procedure also adds the CA certificate to the following existing object to provide trust for logon and authentication certificates:

- Trusted Enterprise CA certificates – CN=NTAuthCertificates,CN=Public Key Services,CN=Services,CN=Configuration,DC=

**Directory Objects That Are Created by the First EnterpriseCA in the Forest**
Installing the first enterprise CA in the forest also installs 29 template objects when running a member of the Windows Server 2003 family or 24 template objects when running Windows 2000 in Active Directory under the following container:

CN=Certificate Templates,CN=Public Key Services,CN=Services,CN=Configuration,DC=

The Windows Server 2003 family adds some additional object identifier containers (also known as OID) to the configuration container. Because object identifiers are not hardcoded in version 2 (V2) templates, object identifier containers are required to work with V2 templates. Only clients running Windows XP and later may resolve object identifiers in Active Directory to friendly names.

CN=OID,CN=Public Key Services,CN=Services,CN=Configuration,DC=

For more information, see article 287547, "Object IDs Associated with Microsoft Cryptography" in the Microsoft Knowledge Base.

# Contents of \\Localhost\CertConfig and \\Localhost\CertEnroll

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Because more than one certificate file exists in the \CertConfig and \CertEnroll share after a period of time, the following table explains the certificate file name extensions and their purpose. If the CA name is used as part of a file name, the sanitized CA name adds additional escape characters in order to accommodate any extended ASCII characters in the file name. The escape characters appear in the file name as %20.

**Table 21 Certificate Paths and File Name Extensions**

| Example of the file name | Description |
|---|---|
| **\\Localhost\CertConfig\Certsrv.txt** | CA configuration file |
| **\\Localhost\CertConfig\Certsrv.bak** | Previous CA configuration file if the CA has been reinstalled |
| **\\Localhost\CertConfig** *CAname* **.req**<br>**\\Localhost\CertConfig** *CAname* **(1).req** | Request file that is used to generate the CA certificate. Request files are used only for subordinate CAs. Request files are generated with the same base file name suffix as certificates. |
| *SystemDriveAndSystemroot**CAname* **.req**<br>*SystemDriveAndSystemroot**CAname* **(1).req** | If no shared folder was created during the CA setup procedure and Active Directory is used to publish the CAs configuration information, request files are written to the Systemroot drive instead of to the \\Localhost\CertConfig file.<br>To verify where the configuration information is published, at a command prompt, type **certutil –getreg CA\UseDS**. If the value is set to 0, the configuration information is written to the shared folder. If the value is set to 1, the configuration is maintained in Active Directory.) |
| **\\Localhost\CertConfig** *CAname* **.crt**<br>**\\Localhost\CertEnroll** *CAname* **.crt** | Original root CA certificate (V0.0) |
| **\\Localhost\CertConfig** *CAname* **(1).crt**<br>**\\Localhost\CertEnroll** *CAname* **(1).crt** | Renewed root CA certificate (V1.0) |
| **\\Localhost\CertConfig** *CAname* **(0-1).crt**<br>**\\Localhost\CertEnroll** *CAname* **(0-1).crt** | Cross certificate for CA certificate V0.0 to V1.0 |
| **\\Localhost\CertConfig** *CAname* **(1-0).crt**<br>**\\Localhost\CertEnroll** *CAname* **(1-0).crt** | Cross certificate for CA certificate V1.0 to V0.0 |

| | |
|---|---|
| **\\Localhost\CertConfig** *CAname* **(2).crt**<br>**\\Localhost\CertEnroll** *CAname* **(2).crt** | renewed root CA cert (V2.0) |
| **\\Localhost\CertEnroll** *CAname* **.crl** | CA base revocation list |
| **\\Localhost\CertEnroll** *CAname* **(1).crl** | CA base revocation list (first instance) |
| **\\Localhost\CertEnroll** *Caname* **+.crl** | Delta CRL |
| **\\Localhost\CertEnroll** *Caname* **(1)+.crl** | Delta CRL (first instance) |

The cross-certificates are automatically generated when the Certificates service starts after renewing a root CA certificate with a new key. Cross-certificates are not created for subordinate CAs, and it does not occur when a root certificate is renewed with the same key. If you upgrade from Windows 2000 Server after renewing a root CA certificate with a new key, the cross certificate is generated the first time that the certificate server service starts after you upgrade to Windows Server 2003.

The following sample is an example of \\Localhost\Certenroll after a clean root CA installation.

```
C:\>dir \\Localhost\certenroll
  Volume in drive \\Localhost\certenroll has no label.
  Volume Serial Number is CC0E-CACB
  Directory of \\Localhost\certenroll
06/12/2002  11:57 AM    <DIR>          .
06/12/2002  11:57 AM    <DIR>          ..
06/12/2002  11:32 AM             1,299 concorp-
ca-00_CorporateRootCA.crt
06/12/2002  11:32 AM               925 CorporateRootCA.crl
06/12/2002  11:32 AM               321 nsrev_CorporateRootCA.asp
               3 File(s)          2,545 bytes
               2 Dir(s)   4,478,095,360 bytes free
```

The following sample is an example of \\Localhost\Certconfig after a clean root CA installation.

```
C:\>dir \\localhost\certconfig
Volume in drive \\localhost\certconfig has no label.
Volume Serial Number is CC0E-CACB
Directory of \\localhost\certconfig
06/12/2002  12:28 PM    <DIR>          .
06/12/2002  12:28 PM    <DIR>          ..
06/12/2002  11:32 AM               105 certsrv.bak
06/12/2002  11:32 AM               216 certsrv.txt
06/12/2002  11:32 AM             1,299 concorp-
ca-00_CorporateRootCA.crt
               3 File(s)          1,620 bytes
               2 Dir(s)   4,478,095,360 bytes free
```

The following sample is an example of \\Localhost\Certenroll after the two key renewals on a CA.

```
C:\>dir \\localhost\certenroll
  Volume in drive \\localhost\certenroll has no label.
  Volume Serial Number is CC0E-CACB
  Directory of \\localhost\certenroll
06/11/2002  07:48 PM    <DIR>           .
06/11/2002  07:48 PM    <DIR>           ..
06/11/2002  05:31 PM             1,338 concorp-
ca-00_CorporateRootCA(1).crt
06/11/2002  05:31 PM             1,928 concorp-ca-00_CorporateRootCA
(0-1).crt
06/11/2002  05:31 PM             1,940 concorp-ca-00_CorporateRootCA
(1-0).crt
06/11/2002  07:48 PM             1,338 concorp-
```

```
ca-00_CorporateRootCA(2).crt
06/11/2002  11:57 AM             1,299 concorp-
ca-00_CorporateRootCA.crt
06/11/2002  05:31 PM               943 CorporateRootCA(1).crl
06/11/2002  05:32 PM               938 CorporateRootCA.crl
06/11/2002  11:57 AM               321 nsrev_CorporateRootCA.asp
               8 File(s)         10,045 bytes
               2 Dir(s)   4,481,171,456 bytes free
```

The following sample is an example of \\Localhost\Certconfig after two key renewals on a CA.

```
C:\>dir \\localhost\certconfig
  Volume in drive \\localhost\certconfig has no label.
  Volume Serial Number is CC0E-CACB
  Directory of \\localhost\certconfig
06/11/2002  07:48 PM    <DIR>           .
06/11/2002  07:48 PM    <DIR>           ..
06/11/2002  11:27 AM           105 certsrv.bak
06/11/2002  11:57 AM           216 certsrv.txt
06/11/2002  05:31 PM             1,928 concorp-ca-00_CorporateRootCA
(0-1).crt
06/11/2002  05:31 PM           1,338 concorp-
ca-00_CorporateRootCA(1).crt
06/11/2002  05:31 PM           1,940 concorp-ca-00_CorporateRootCA
(1-0).crt
06/11/2002  07:48 PM           1,338 concorp-
ca-00_CorporateRootCA(2).crt
06/11/2002  11:57 AM           1,299 concorp-
ca-00_CorporateRootCA.crt
04/24/2002  10:53 AM           1,942 connoam-ca-00_CONNOAM-CA-00.req
           8 File(s)           10,106 bytes
           2 Dir(s)   4,481,171,456 bytes free
```

# Relationship of the Configuration Container and Certificate Store

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The table in this section describes the relationship between the information that is stored in the configuration container of Active Directory and the certificate store. Typically, parts of the configuration information are replicated to the client's certificate store.

The default view of the Certificates MMC does not display the physical structure of the certificate store. To view the physical structure of the certificate store, follow this procedure:

1. Open Certificates.

   To do this, click **Start**, click **Run**, in the **Open** box, type **certmgr.msc**, and then press ENTER.

2. Verify that the local computer's certificates and the current user's certificate are displayed in the console tree.

3. In the console tree, click **Certificate (Local Computer)**.

4. On the **View** menu, click **Options**, and select the **Physical certificates store** check box.

| **Note** |
| --- |
| Any information that is stored in a registry container has an impact on only the local client. Registry containers never receive information from the Active Directory configuration context. The **Intermediate Certificate Authorities – Group Policy** container is not used in the client certificate store. |

Certificates that are stored in the Active Directory Configuration container (Sites and Services) are deployed to all clients across the forest. Certificates that are deployed through domain security are deployed only in the domain. If a certificate is registered in the Configuration container and the Domain Security Group Policy object (GPO), a certificate may occur twice on the client. To prevent confusion with expired or invalid certificates, you must ensure that certificates are correctly published.

You can view the Active Directory configuration context through the Active Directory Sites and Services MMC.

**Table 22 Certificate Containers and Certificate Stores**

| Active Directory Configuration container | Client certificates store |
| --- | --- |
| **Active Directory Sites and Services MMC**<br>In the console tree, navigate to Certification Authorities:<br>*DomainName*\Configuration Services\Public Key Services\Certification Authorities<br>Enterprise CAs are installed and automatically published to this location. CA certificates may also be added manually through the **certutil –dspublish** command. | **Local Computer**<br>In the console tree, navigate to Certificates:<br>Trusted Root Certificate Authorities\Enterprise\Certificates |

| | |
|---|---|
| **Sites and Services MMC**<br><br>In the console tree, navigate to AIA:<br><br>*DomainName*\Configuration\Services\Public Key Services\AIA<br><br>This container also contains qualified subordination certificates (cross-certificates) that are controlled by the template that is used to generate CA certificates. | **Local Computer**<br><br>In the console tree, navigate to Certificates:<br><br>Intermediate Certificate Authorities\Enterprise\Certificates<br><br>Windows 2000, Windows XP, or Windows Server 2003 clients automatically download the content from the configuration container; Windows 2000 clients do not support cross-certificates |
| **Domain Security Settings MMC**<br><br>In the console tree, navigate to Trusted Root Certification Authorities:<br><br>Computer Configuration\Windows settings\Security Settings\Public Key Policies\Trusted Root Certification Authorities | **Local Computer**<br><br>In the console tree, navigate to Certificates:<br><br>Trusted Root Certificate Authorities\Group Policy\Certificates |
| **Domain Security Settings MMC**<br><br>In the console tree, navigate to Enterprise Trust:<br><br>Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Enterprise Trust | **Local Computer**<br><br>In the console tree, navigate to Group Policy:<br><br>Enterprise Trust\Group Policy |

# Default CA Certificate and CRL Storage

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

During the installation of the root CA, the root certificate is saved to the following locations:

- \\Localhost\Certenroll
- \\Locahost\Certconfig
- The Certificates MY store of the local computer
- The Trusted Root Certification Authorities container in the local computer registry

The initial CRL is published in the following locations:

- \\Localhost\Certenroll
- The Intermediate Certification Authorities container in the registry of the local computer

The CA certificate of the stand-alone CA is stored in the following locations:

- \\Localhost\Certenroll
- \\Locahost\Certconfig
- Certificates store of the local computer. To look at the store, do one of the following:
    - In the Certificates MMC, in the console tree, double-click **Certificates (Local Computer)**, double-click **Registry**, and then click **Certificates**.
    - In the Certificates MMC, in the console tree, double-click **Certificates (Local Computer)**, double-click **Intermediate Certification Authorities**, double-click **Registry**, and then click **Certificates**.

The CRL of the root CA should be stored in the following locations:

- Certificates store of the local computer. To look at the store, do the following:
    - In the Certificates MMC, in the console tree, double-click **Certificates (Local Computer)**, double-click **Intermediate Certification Authorities**, double-click **Registry**, and then click **Certificate Revocation List**

The CRL of the stand-alone CA is stored in the following locations:

- File share \\Localhost\Certenroll
- Certificates store of the local computer. To look at the store, do the following:
    - In the Certificates MMC, in the console tree, double-click **Certificates (Local Computer)**, double-click **Intermediate Certification Authorities**, double-click **Registry**, and then click **Certificate Revocation List**

# Mapping Custom Object Identifiers to Friendly Names

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

When a certificate is enrolled and that certificate carries a custom object identifier and the policy information, an enrolled certificate's purpose may display an object identifier instead of a friendly description.

This occurs because the template that is used for certificate enrollment cannot translate the object identifier into a friendly name. Because of this, custom object identifiers are mapped to friendly names through the object identifier (also known as OID) container in the Active Directory. The mapping must be done in the V2 template that will use the custom object identifier. To translate the object identifier into a friendly name:

1. Open the Certificate Templates MMC.

   To do this, click **Start**, click **Run**, in the **Open** box, type **certtmpl.msc**, press ENTER, and then open any V2 template.

2. Click the **Extensions** tab.

3. Click the **Application Policies** extension.

4. Click **Edit**, click **Add**, and then click **New**.

5. Type both the friendly name and related object identifier number, and then click **OK**.

# CAPolicy.inf Syntax

Updated: February 13, 2012

Applies To: Windows Server 2003 with SP1

The purpose of the CAPolicy.inf configuration file and its syntax is described in Windows Server 2003 Server Help.

If a CAPolicy.inf file exists, it supersedes the default configuration that is used to install a CA or renew its CA certificate.

**Sample CAPolicy.inf File for CorporateRootCA**
You can use the samples in this section for the root CAs CAPolicy.inf file. Verify that the parameters in the [Certsrv_Server] section are the same as your requirements, according to the CPS.

| **Note** |
| --- |
| The parameters specified in the [Certsrv_Server] section must be greater or must match the key length and validity period used during CA setup otherwise the value specified in the capolicy.inf will be ignored. |

Unfortunately the keynames are different for Windows 2000 and Windows Server 2003. However, the Windows Server 2003 Certification Server is able to interpret the old Windows 2000 syntax.

A CApolicy.inf file for a Windows 2000 root CA would look like this:

```
[Version]
Signature= "$Windows NT$"
[Certsrv_Server]
RenewalKeyLength=4096
RenewalValidityPeriodUnits=Years
RenewalValidityPeriod=20
[CRLDistributionPoint]
[AuthorityInformationAccess]
```

A CApolicy.inf file for a Windows 2003 Server root CA would look like this:

```
[Version]
Signature= "$Windows NT$"
[Certsrv_Server]
RenewalKeyLength=4096
RenewalValidityPeriod=Years
RenewalValidityPeriodUnits=20
[CRLDistributionPoint]
[AuthorityInformationAccess]
```

As you can see, the RenewalValidityPeriod and RenewalValidityPeriodUnits parameters are switched in Windows Server 2003 to make them more consistent.

| **Important** |
| --- |
| Although you can enter values for CRL Deltas and CRL Overlap, they will not be read during setup. To configure CRL Overlap and CRL Delta periods, you must run commands after installation. See Configure CRL and Delta CRL Overlap Periods (http://technet.microsoft.com/library/cc731104.aspx). |

If you are using a Windows 2000 CA, see article 297528, "CRL Distribution Point Extension Is Not Suppressed by the Capolicy.inf File," in the Microsoft Knowledge Base.

**Sample CAPolicy.inf File for IntermediateCA1**
This section contains a sample for the subordinate CAs CAPolicy.inf file.

The object identifiers and the URLs are provided only as an example. You should replace the object identifier values with object identifiers that belong to your organization and verify that the URLs are pointing to a location that is accessible.

The CAPolicy.inf file syntax in Windows 2000 and Windows Server 2003 are basically the same, except that the [CApolicy] section, which was valid in Windows 2000, is now [PolicyStatementExtension].

On a Windows 2000 CA, the CAPolicy.inf file should look like the following sample, except that the italicized items are

placeholders. The placeholders should be replaced with the information for your specific situation.

```
[Version]
Signature= "$Windows NT$" [CApolicy]
Policies = AllIssuancePolicy
Critical = FALSE
[AllIssuancePolicy]
OID = 2.5.29.32.0
```

For a Windows Server 2003 CA, the CAPolicy.inf file should look like the following sample, where the italicized items are placeholders.

```
[Version]
Signature= "$Windows NT$"
[ PolicyStatementExtension ]
Policies = AllIssuancePolicy
Critical = FALSE
[AllIssuancePolicy]
OID = 2.5.29.32.0
```

**Sample CAPolicy.inf File for CorporateEnt1CA**

### For the Windows 2000 Family

```
[Version]
Signature= "$Windows NT$"
[CApolicy]
Policies = LegalPolicy, LimitedUsePolicy
[LegalPolicy]
OID = 1.1.1.1.1.1.1.1.1
URL = "http://www.contoso.com/pki/Policy/USLegalPolicy.asp"
URL = "ftp://ftp.contoso.com/pki/Policy/USLegalPolicy.txt"
[LimitedUsePolicy]
OID = 2.2.2.2.2.2.2.2.2
URL = "http://www.contoso.com/pki/Policy/USLimitedUsePolicy.asp"
URL = "ftp://ftp.contoso.com/pki/Policy/USLimitedUsePolicy.txt"
```

### For the Windows Server 2003 Family

```
[Version]
Signature= "$Windows NT$"
[PolicyStatementExtension]
Policies = LegalPolicy, LimitedUsePolicy
[LegalPolicy]
OID = 1.1.1.1.1.1.1.1.1
URL = "http://www.contoso.com/pki/Policy/USLegalPolicy.asp"
URL = "ftp://ftp.contoso.com/pki/Policy/USLegalPolicy.txt"
[LimitedUsePolicy]
OID = 2.2.2.2.2.2.2.2.2
URL = "http://www.contoso.com/pki/Policy/USLimitedUsePolicy.asp"
URL = "ftp://ftp.contoso.com/pki/Policy/USLimitedUsePolicy.txt"
```

# CRL Distribution Point Replacement Token

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Replacement tokens are used to retain the configuration of distribution points flexible. You can use replacement tokens in the CAPolicy.inf file and in the Certification Authority MMC in **CA Extensions**.

A replacement token consists of the percent sign and a number. This behavior occurs if you use replacement tokens in the Certificate Services MMC or if you use them in a **certutil** command. If replacement tokens are used in a batch file, and you use the percent sign (%), you must use another escape sign when needed, because the Windows shell typically interprets a percent sign as a command-line parameter.

The mapping of replacement tokens is different in versions of Windows later than Windows 2000 Server. For more information and a list of replacement tokens that are valid on computers that are running Windows 2000, see Article 283119, Error Message: A Replacement Token Entered Does Not Match Any Recognized Token in the Microsoft Knowledge Base.

You can use the following tokens for **CRLDistributionPoint**, **AuthorityInformationAccess**, and **CrossCertificateDistributionPointsExtension** URLs.

**Table 23 CRL Distribution Point Replacement Tokens**

| Token name | Description | Windows 2000 map value | Windows Server 2003 map value |
|---|---|---|---|
| ServerDNSName | The DNS name of the CA server | %1 | %1 |
| ServerShortName | The NetBIOS name of the CA server | %2 | %2 |
| CaName | The name of the CA | %3 | %3 |
| Cert_Suffix | The renewal extension of the CA | %4 | N/A |
| CertificateName | | N/A | %4 |
| Domain_Name | The location of the domain root in Active Directory | %5 | N/A |
| (Not used) | | N/A | %5 |
| ConfigurationContainer | The location of the configuration container in Active Directory | %6 | %6 |
| CATruncatedName | The "sanitized" name of the CA, 32 characters with a hash on the end | %7 | %7 |
| CRLNameSuffix | The renewal extension for the CRL | %8 | %8 |
| DeltaCRLAllowed | | | %9 |
| CDPObjectClass | | | %10 |
| CAObjectClass | | | %11 |

The Certification Server setup process replaces all **%number%** sequences with the appropriate value.

# CRL Publishing Properties

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The **Publish CRLs to this location** flag is used to identify the locations to which the CA should publish (or place) the physical CRLs when the CA publishes a CRL either automatically or manually. This flag specifies only where CRLs are published. It is also used by the **certutil.exe –dspublish** command when you manually publish CRLs to Active Directory. Both the **Publish CRL** and **Publish Delta CRL** flags on the Revoked Certificates Properties page are responsible for turning the publishing activity on and off.

The **Publish CRLs to this location** flag indicates the locations that the CA should attempt to use to publish the CRL. This flag does not configure the server to conduct the publishing activity, but only sets it up so that the CA can determine the appropriate locations to which to publish when publishing occurs. Note that actual publishing activity is governed by the **Revoked Certificates** properties.

The **Include in all CRLs** flag specifies that the Active Directory publication location should be included in the CRL itself. This information is useful for publishing offline CRLs to Active Directory by using the Certutil.exe tool. To use this, at a command prompt, type **certutil –dspublish**, and then press ENTER.

The **Include in CDP extension of issued certificates** flag is used by clients to find the CRL distribution point location for the CRL. You should always specify this flag unless you do not want to use client-side checking or application revocation checking for issued certificates.

The **Include in CRLs. Clients use this to find Delta CRLs** flag is used by clients to determine if a delta CRL exists and where it is located. The location may or may not be the same as the CRL location. The delta CRL location is identified in the CRL by use of the **freshestCRL** extension in the CRL object itself.

You may want to have a base CRL in an LDAP location in Active Directory and a delta CRL at an alternate HTTP location because of the differences in replication. If the delta CRL will be issued at an interval that is shorter than the replication convergence time for your forest, the delta CRL should not be published to Active Directory. In many Active Directory networks, it may take hours for Active Directory objects to fully replicate throughout the network. For delta CRLs that may have a lifetime only of a few hours, the replication latency often means that Active Directory clients receive a delta CRL object that has already expired by the time it reaches the client. You can avoid this latency by publishing the delta CRL to an HTTP location that is serviced by fault-tolerant Web servers, where all clients can immediately retrieve a fresh delta CRL.

**Table 24 CRL Publishing Properties**

| Display name | Description | Decimal value | Hexadecimal value |
|---|---|---|---|
| Publish CRLs to this location | Used by the CA to determine whether to publish base CRLs to this URL | 1 | 0x00000001 |
| Include in the CRL distribution point extension of issued certificates | Used by clients during revocation checking to find base CRL locations | 2 | 0x00000002 |
| Include in [base] CRLs | Used by clients during revocation checking to find delta CRL locations from base CRLs | 4 | 0x00000004 |
| Include in all CRLs | Not used during revocation checking. Specifies where to publish in Active Directory when publishing manually using **certutil -dspublish**. Can be used by an offline CA to specify the LDAP URL for manually publishing CRLs. Must also set the explicit configuration container in the URL or set the **DSConfigDN** value in the registry: **certutil –setreg ca\DSConfigDN CN=** | 8 | 0x00000008 |
| | | 16 | 0x00000010 |

| | | | |
|---|---|---|---|
| | | 32 | 0x00000020 |
| Publish delta CRLs to this location | Used by the CA to determine whether to publish delta CRLs to this URL | 64 | 0x00000040 |

# AIA Publishing Properties

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The table shows the publishing properties for the Authority Information Access (AIA).

**Table 25 AIA Publishing Properties**

| Display name | Decimal value | Hexadecimal value |
|---|---|---|
| Include in the AIA extension of issued certificates | 2 | 0x00000002 |
| Include in the online certificate status protocol (OCSP) extension | 32 | 0x00000020 |

# Sample Script to Configure CorporateRootCA

Updated: February 14, 2012

Applies To: Windows Server 2003 with SP1

The script in this section applies the most important configuration changes to a Windows Server 2003 CA for the CorporateRootCA computer.

| **Important** |
| --- |
| Because percent (%) variables are handled differently in batch files and at a command prompt, you must use two percent signs (%%) if you run this sample script from a batch file, as described. If **certutil** is called from a command prompt and not from a batch file, only use only one percent sign (%), not two percent signs (%%). |

```
REM
REM CA configuration script for a Windows Server 2003 CA
REM
REM The naming context applies to the individual organizations Active
Directory
REM configuration
REM
SET myADnamingcontext=DC=concorp,DC=contoso,DC=com
REM
REM This variable directs to the HTTP publication location that is used
for
REM the CRL and AIA publication
REM
SET myhttpPKIvroot=http://www.contoso.com/pki
REM
certutil.exe -setreg ca\DSConfigDN
"CN=Configuration,%myADnamingcontext%"
REM
REM Configure CRL and AIA CDP
REM
REM By default, Certutil creates a registry value of type REG_SZ if a
string is
REM specified as a parameter. Some registry values are expected as
REG_MULTI_SZ. To
REM create a REG_MULTI_SZ instead of a REG_SZ, add a \n to the end of
any value that
REM becomes part of the REG_MULTI_SZ
REM
certutil -setreg CA\CRLPublicationURLs
"1:%WINDIR%\system32\CertSrv\CertEnroll\%%3%%8%%9.crl\n2:%myhttp
PKIvroot%/%%3%%8%%9.crl\n10:ldap:///CN=%%7%%8,CN=%%2,
CN=CDP,CN=Public Key Services,CN=Services,%%6%%10"
certutil -setreg CA\CACertPublicationURLs
"1:%WINDIR%\system32\CertSrv\CertEnroll\%%1_%%3%%4.crt\n2:%myhttp
PKIvroot%/%%1_%%3%%4.crt\n2:ldap:///CN=%%7,CN=AIA,
CN=Public Key Services,CN=Services,%%6%%11"
REM
REM Configure CRL publication
REM
certutil -setreg CA\CRLPeriodUnits 180
certutil -setreg CA\CRLPeriod "Days"
REM
REM Disable Delta CRL publication
REM
certutil -setreg CA\CRLDeltaPeriodUnits 0
REM
REM Set the validity period for issued certificates
REM
certutil -setreg ca\ValidityPeriodUnits 10
certutil -setreg ca\ValidityPeriod "Years"
REM
REM Restart the CA server service
REM
```

```
net stop certsvc & net start certsvc
REM
REM Repair CA file system shares and IIS virtual roots
REM
certutil -vroot
REM
REM Republish the CRL
REM The CRL publishing may immediately not work
REM after you restart the CA server service. If this behavior
REM occurs, try the certutil -CRL command at a command
REM prompt again.
REM
certutil -CRL
REM
REM Test if CAPolicy.inf file exists
REM
IF EXIST %SYSTEMROOT%\capolicy.inf GOTO ENDCFG
ECHO Warning, no capolicy.inf file used
:ENDCFG
```

The following script applies the same configuration as the previous script, but it configures a Windows 2000 CA. Remember that the delta CRL configuration parameter is not supported in a Windows 2000 CA environment. To perform the **certutil –URL** and **certutil –vroot** commands, you must run the version of **certutil** that is included with Windows Server 2003 on the Windows 2000 CA computer.

```
REM
REM CA configuration script for a Windows 2000 CA
REM
REM This variable directs to the HTTP publication location that is used
for
REM the CRL and AIA publication
REM
SET myhttpPKIvroot=http://www.contoso.com/pki
REM
REM Because CRLs and CA certificates are published in the
organizations Active
REM
REM Configure CRL and AIA CDP
REM
certutil -setreg policy\FileRevocationCRLURL "\n
certutil -setreg policy\RevocationCRLURL
%myhttpPKIvroot%/%%3%%8.crl\n
certutil -setreg policy\LDAPRevocationCRLURL
ldap:///CN=%%7%%8,CN=%%2,CN=CDP,CN=Public Key Services,
CN=Services,%%6?certificateRevocationList?base?objectclass=
cRLDistributionPoint\n"
certutil -setreg policy\FileIssuercertURL  "\n
certutil -setreg policy\IssuercertURL %myhttpPKIvroot%/%%1_%%3%%4.crt"
certutil -setreg policy\LDAPIssuercertURL
ldap:///CN=%%7,CN=AIA,CN=Public  Key
 Services,CN=Services,%%6?cACertificate?base?objectclass=
certificationAuthority
REM
REM Configure CRL publication
REM
certutil -setreg CA\CRLPeriodUnits 180
certutil -setreg CA\CRLPeriod "Days"
REM
REM Set the validity period for issued certificates
REM
certutil -setreg ca\ValidityPeriodUnits 10
certutil -setreg ca\ValidityPeriod "Years"
REM
REM Disable issuer name and issuer serial number
REM
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERNAME
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERSERIAL
REM
REM Restart the CA server service
REM
net stop certsvc & net start certsvc
REM
```

```
REM Repair CA files-system shares and IIS virtual roots
REM
certutil -vroot
REM
REM Publish the CRL with the updated CDP and naming information.
REM It might happen that CRL publishing fails immediately
REM after the CA server service has been restarted. If this
REM is the case, try certutil -CRL at a command prompt again.
REM
certutil -CRL
```

# Sample Script to Configure IntermediateCA

Updated: February 14, 2012

Applies To: Windows Server 2003 with SP1

The following script applies the most important configuration changes to a Windows Server 2003 CA for the IntermediateCA computer.

```
REM
REM CA configuration script for a Windows Server 2003 CA
REM
REM The naming context applies to the individual organizations Active
Directory
REM configuration
REM
SET myADnamingcontext=DC=concorp,DC=contoso,DC=com
REM
REM This variable directs to the HTTP publication location that is used
for
REM the CRL and AIA publication
REM
SET myhttpPKIvroot=http://www.contoso.com/pki
REM
REM Because CRLs and CA certificates are published in the
organizations Active
REM Map the namespace of Active Directory
REM
certutil.exe -setreg ca\DSConfigDN
"CN=Configuration,%myADnamingcontext%"
REM
REM Configure CRL and AIA CDP
REM
REM By default, Certutil creates a registry value of type REG_SZ if a
string is
REM specified as a parameter. Some registry values are expected as
REG_MULTI_SZ.
REM To create a REG_MULTI_SZ instead of a REG_SZ, add a \n to the end
of any value
REM that becomes part of the REG_MULTI_SZ
REM
certutil -setreg CA\CRLPublicationURLs
"1:%WINDIR%\system32\CertSrv\CertEnroll\%%3%%8%%9.crl\n2:%myhttp
PKIvroot%/%%3%%8%%9.crl\n10:ldap:///CN=%%7%%8,CN=%%2,
CN=CDP,CN=Public Key Services,CN=Services,%%6%%10"
certutil -setreg CA\CACertPublicationURLs
"1:%WINDIR%\system32\CertSrv\CertEnroll\%%1_%%3%%4.crt\n2:%myhttp
PKIvroot%/%%1_%%3%%4.crt\n2:ldap:///CN=%%7,CN=AIA,
CN=Public Key Services,CN=Services,%%6%%11"
REM
REM Configure CRL publication
REM
certutil -setreg CA\CRLPeriodUnits 30
certutil -setreg CA\CRLPeriod "Days"
REM
REM Disable Delta CRL publication
REM
certutil -setreg CA\CRLDeltaPeriodUnits 0
REM
REM Set the validity period for issued certificates
REM
certutil -setreg ca\ValidityPeriodUnits 5
certutil -setreg ca\ValidityPeriod "Years"
REM
REM Include certificate policies in certificate request
REM
certutil -v -setreg policy\EnableRequestExtensionlist "+2.5.29.32"
REM
REM
REM Disable issuer name and issuer serial number
```

```
REM
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERNAME
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERSERIAL
REM
REM Restart the CA server service
REM
net stop certsvc & net start certsvc
REM
REM Repair CA files-system shares and IIS virtual roots
REM
certutil -vroot
REM
REM Republish the CRL
REM It might happen that CRL publishing fails immediately
REM after the CA server service has been restarted. If this
REM is the case, try certutil -CRL at a command prompt again.
REM
certutil -CRL
```

The following script applies the same configuration as the previous script but the following script configures a Windows 2000 CA. Remember that the delta CRL configuration parameter is not supported in a Windows 2000 CA environment. To use the **certutil –URL** and **certutil –vroot** command, you must run the Windows Server 2003 version of the **certutil** utility on the Windows 2000 CA computer.

```
REM
REM CA configuration script for a Windows 2000 CA
REM
REM This variable directs to the HTTP publication location that is used
for
REM the CRL and AIA publication
REM
SET myhttpPKIvroot=http://www.contoso.com/pki
REM
REM Configure CRL and AIA CDP
REM
REM By default, certutil creates a registry value of type REG_SZ if a
string is
REM specified as a parameter. Some registry values are expected as
REG_MULTI_SZ. To
REM create a REG_MULTI_SZ value instead of a REG_SZ value, add \n to
the end of any
REM value that becomes part of REG_MULTI_SZ.
REM
certutil -setreg policy\FileRevocationCRLURL \n
certutil -setreg policy\RevocationCRLURL
%myhttpPKIvroot%/%%3%%8.crl\n
certutil -setreg policy\LDAPRevocationCRLURL
ldap:///CN=%%7%%8,CN=%%2,CN=CDP,CN=Public Key Services,
CN=Services,%%6?certificateRevocationList?base?objectclass=
cRLDistributionPoint\n"
certutil -setreg policy\FileIssuercertURL
"%WINDIR%\system32\CertSrv\CertEnroll\%%1_%%3%%4.crt\n
certutil -setreg policy\IssuercertURL %myhttpPKIvroot%/%%1_%%3%%4.crt"
certutil -setreg policy\LDAPIssuercertURL
ldap:///CN=%%7,CN=AIA,CN=Public
Key Services,CN=Services,%%6?cACertificate?base?objectclass=
certificationAuthority
REM
REM Configure CRL publication
REM
certutil -setreg CA\CRLPeriodUnits 30
certutil -setreg CA\CRLPeriod "Days"
REM
REM Set the validity period for issued certificates
REM
certutil -setreg ca\ValidityPeriodUnits 5
certutil -setreg ca\ValidityPeriod "Years"
REM
REM Include certificate policies in certificate request
REM
certutil -v -setreg policy\EnableRequestExtensionlist "+2.5.29.32"
REM
```

```
REM
REM Disable issuer name and issuer serial number
REM
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERNAME
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERSERIAL
REM
REM Restart the CA server service
REM
net stop certsvc & net start certsvc
REM
REM Repair CA files-system shares and IIS virtual roots
REM
certutil -vroot
REM
REM Republish the CRL.
REM It might happen that CRL publishing fails immediately
REM after the CA server service has been restarted. If this
REM is the case try certutil -CRL at a command prompt again.
REM
certutil -CRL
```

# Sample Script to Configure the EnterpriseSubCA

Updated: February 14, 2012

Applies To: Windows Server 2003 with SP1

The following script applies the most important configuration changes to a Windows Server 2003 CA for the EnterpriseSubCA computer.

| **Important** |
| --- |
| Because percent (%) variables are handled differently in batch files and at a command prompt, you must use two percent signs (%%) if you run this sample script from a batch file, as described. If **certutil** is called from a command prompt and not from a batch file, only use one percent sign (%), not two (%%). |

```
REM
REM CA configuration script for a Windows Server 2003 CA
REM
REM This variable directs to the HTTP publication location that is used
for
REM CRL and AIA publication
REM
SET myhttpPKIvroot=http://www.contoso.com/pki
REM
REM Configure CRL and AIA CDP
REM
certutil -setreg CA\CRLPublicationURLs "65:
%WINDIR%\system32\CertSrv\CertEnroll\%%3%%8%%9.crl\n6:%myhttp
PKIvroot%/%%3%%8%%9.crl\n79:ldap:///CN=%%7%%8,CN=%%2,
CN=CDP,CN=Public Key Services,CN=Services,%%6%%10"
certutil -setreg CA\CACertPublicationURLs  "1:
%WINDIR%\system32\CertSrv\CertEnroll\%%1_%%3%%4.crt
n2:%myhttpPKIvroot%/%%1_%%3%%4.crt\n2:ldap:///CN=%%7,
CN=AIA,CN=Public Key Services,CN=Services,%%6%%11\"
REM
REM Configure CRL publication
REM
certutil -setreg CA\CRLPeriodUnits 1
certutil -setreg CA\CRLPeriod "Days"
REM
REM Disable issuer name and issuer serial number
REM
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERNAME
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERSERIAL
REM
REM Restart the CA server service
REM
net stop certsvc & net start certsvc
REM
REM Create Web virtual roots and file shares
REM
certutil.exe -vroot
REM
REM Republish the CRL
REM
certutil -CRL
```

The following script applies the same configuration as the previous script, but it configures a Windows 2000 CA. Remember that the delta CRL configuration parameter is not supported in a Windows 2000 CA environment. To use the **certutil –URL** and **certutil –vroot** commands, you must run the version of the **Certutil.exe** utility that is included with the Windows Server 2003 operating system on the computer serving as the Windows 2000 CA.

```
REM
REM CA configuration script for a Windows 2000 CA
REM
REM This variable directs to the HTTP publication location that is used
for
REM the CRL and AIA publication
```

```
REM
SET myhttpPKIvroot=http://www.contoso.com/pki
REM
REM Configure CRL and AIA CDP
REM
certutil -setreg policy\FileRevocationCRLURL "\n
certutil -setreg policy\RevocationCRLURL
%myhttpPKIvroot%/%%3%%8.crl\n
certutil -setreg policy\LDAPRevocationCRLURL
ldap:///CN=%%7%%8,CN=%%2,CN=CDP,CN=Public Key Services,
CN=Services,%%6?certificateRevocationList?base?objectclass=
cRLDistributionPoint\n"
certutil -setreg policy\FileIssuercertURL
"%WINDIR%\system32\CertSrv\CertEnroll\%%1_%%3%%4.crt\n
certutil -setreg policy\IssuercertURL %myhttpPKIvroot%/%%1_%%3%%4.crt"
certutil -setreg policy\LDAPIssuercertURL
ldap:///CN=%%7,CN=AIA,CN=Public Key
Services,CN=Services,%%6?cACertificate?base?objectclass=
certificationAuthority
REM
REM Configure CRL publication
REM
certutil -setreg CA\CRLPeriodUnits 1
certutil -setreg CA\CRLPeriod "Days"
REM
REM Disable delta CRL publication
REM
certutil -setreg CA\CRLDeltaPeriodUnits 0
REM
REM Disable issuer name and issuer serial number
REM
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERNAME
certutil -setreg policy\EditFlags -EDITF_ENABLEAKIISSUERSERIAL
REM
REM Restart the CA server service
REM
net stop certsvc & net start certsvc
REM
REM Create Web virtual roots and file shares
REM
certutil.exe -vroot
REM
REM Republish the CRL
REM
certutil -CRL
```

# Appendix B: Parameters for a Three-Tier CA Topology

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

This section describes all of the parameters that are required to set up a three-tier CA topology. It is recommended that the values are agreed between the departments in the organization (IT department, legal department, and so on).

The parameters in this section are in the sequence in which they are used during the setup. The heading describes the parameter's name and the table contains detailed information about the parameter.

| **Important** |
|---|
| Make sure that you have predefined all of the parameters in this section, because every value is mandatory. |

**RootCA Configuration Parameters**

This section provides a list of parameters that must be defined during the setup procedure for a stand-alone offline root CA. The sample values are related to the sample configuration that is explained in the previous section.

Registry references follow the syntax that is used by the **certutil** command. To get more information about the registry values, at a command prompt, type **certutil –getreg -?** and press ENTER.

**Renewal Key Length (CA Certificate)**

| | |
|---|---|
| Description | It is recommended that the key length does not exceed 4096 bits because this is the maximum interoperable key length with most programs and PKI providers. The renewal key length must not be shorter than the key length that you chose during the CA installation procedure. |
| Sample value | 4096 |
| Defined at | CAPolicy.inf |
| Stored at | Renewed CA certificate |
| Impacts | The root CA key material |

**Renewal Validity Period (CA Certificate)**

| | |
|---|---|
| Description | Describes the lifetime of a CA certificate that is a renewal of a previous CA certificate. It is recommended that root CAs be configured with a longer lifetime than any other CA in the hierarchy because this configuration reduces the administrative burden that is caused by renewing all certificates that are singed by the CA's certificate. |
| Sample value | 1020 |
| Defined at | CAPolicy.inf |
| Stored at | CA certificate that is related to the date and time when the certificate was enrolled |
| Impacts | The CA root certificate and all certificates that will be signed by the root |

**Renewal Validity Period Units (CA Certificate)**

| | |
|---|---|
| Description | Defines the measurement related to the validity time. Valid values are years, months, or days. For a CA certificate lifetime the usual unit is years. |
| Sample value | Years |

| | |
|---|---|
| Defined at | CAPolicy.inf |
| Stored at | CA certificate related to the date and time when the certificate has been enrolled |
| Impacts | The CA root certificate and all certificates that will be signed by the root |

**Certificate Revocation List (CRL) Distribution Point (CA certificate)**

| | |
|---|---|
| Description | A CRL distribution point must not be configured to be contained in the self-signed root CA certificate. Most applications do not check revocation on root CA certificates; therefore, CRL distribution point extensions are not necessary or recommended. It is also senseless to set a CRL distribution point for a root certificate because there is no higher instance that could revoke the root certificate. |
| Sample value | None |
| Defined at | CAPolicy.inf |
| Stored at | CA certificate |
| Impacts | The attribute setting in the CA root certificate and all applications that verify the root CA's validity |

**Authority Information Access (AIA) (CA certificate)**

| | |
|---|---|
| Description | An AIA must not be specified for a root CA certificate. This is because the AIA points to the location of the certificate that was used for signing this certificate. Since a root CA is self-signed, you do not need to specify an AIA. |
| Sample value | None |
| Defined at | CAPolicy.inf |
| Stored at | CA certificate |
| Impacts | All applications that verify the root CA's validity |

**CSP (CA Certificate)**

| | |
|---|---|
| Description | The CSP is responsible for generating the certificates key material and the certificate generation. |
| Sample value | Microsoft Strong Cryptographic Provider |
| Defined at | CA Installation Wizard |
| Stored at | For the Windows 2000 Server family and the Windows Server 2003 family:<br><br>CA Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration **_CAName_**\CSP\Provider** |
| Impacts | CA certificate |

**Hash Algorithm**

| | |
|---|---|
| Description | Defines the hash algorithm that is used for hashing and signing certificate contents. |
| Sample value | SHA-1 |
| Defined at | CA installation wizard |
| Stored at | CA registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration**_CAName_ **\CSP\HashAlgorithm** |

| | |
|---|---|
| Impacts | CA certificate |

**Key Length (CA Certificate)**

| | |
|---|---|
| Description | Defines the complexity of the key material assigned to the CA certificate. It is recommended that the key length does not exceed 4096 bits because this is the maximum interoperable key length today with most applications and PKI providers. |
| Sample value | 4096 |
| Defined at | CA Installation Wizard |
| Stored at | Certificate request and is only used temporarily |
| Impacts | The Root CA key material that could be stored within a HSM or encrypted on the CAs hard drive |

**Common Name**

| | |
|---|---|
| Description | The common name must not exceed 64 characters in length. It is important to remember that each space in the name will actually use three characters in the total length because of how escape characters are written (%20). |
| Sample value | CorporateRootCA |
| Defined at | CA Installation Wizard |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration **_CAName_**\CommonName** |
| Impacts | The common name becomes part of the certificate issuer name and is also part of the CRL and AIA if replacement tokens are used. The common name is used by several variables that are used to set the CRL and AIA. |

**Distinguished Name Suffix**

| | |
|---|---|
| Description | The name maps to the namespace that is used by the domain where the CA belongs to. Since the Root-CA is configured as a stand-alone CA, the distinguished name should be mapped to the same namespace that will be used for the enterprise CA. |
| Sample value | DC=concorp,DC=contoso,DC=com |
| Defined at | CA configuration that takes place after the installation |
| Stored at | Windows 2000 and Windows 2003<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc \Configuration **_CAName_**\DSConfigDN** |
| Impacts | The distinguished name becomes part of the certificate issuer name and is also part of the CRL and AIA if replacement tokens are used. It is also used by several variables that are used to set the CRL and AIA. |

**Validity Period (CA Certificate)**

| | |
|---|---|
| Description | The parameter defines how long from now the CA certificate will be valid, depending on the validity period units |
| Sample value | 2 |
| Defined at | CA Installation Wizard |
| Stored at | CA certificate related to the date and time when the certificate has been enrolled |
| Impacts | The CA certificate and the validity time of all certificates that are signed by the Root CA certificate. |

**CA Database Path**

| | |
|---|---|
| Description | Defines where the CA's database is located in the root CA's file system. |
| Sample value | C:\Certlog |
| Defined at | CA installation wizard |
| Stored at | Windows 2000 and Windows 2003<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc**<br><br>**\Configuration ** *CAName* **\DBDirectory** |
| Impacts | The CA must be able to get the appropriate path name from the registry when the CA starts up. |

**CA Log File Path**

| | |
|---|---|
| Description | Defines where the CA's transaction log-files are located in the CA's file system. |
| Sample value | C:\Certlog |
| Defined at | CA Installation Wizard |
| Stored at | Windows 2000 and Windows 2003 Server families:<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc**<br><br>**\Configuration ** *CAName* **\DBLogDirectory** |
| Impacts | The CA must be able to get the appropriate path name from the registry when the CA starts up. |

**Shared Folder**

| | |
|---|---|
| Description | Defines where the CA's transaction log-files are located in the root CA's file system. |
| Sample value | \\[{localhost]}\CertConfig |
| Defined at | CA installation wizard |
| Stored at | Windows 2000 and the Windows 2003 Server family<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc**<br><br>**\Configuration ** *CAName* **\ConfigurationDirectory** |
| Impacts | Clients, those are not able to receive the CA certificate through group policies and need to import the certificate manually. |

**Certificate Revocation List (CRL) Distribution Point**

| | |
|---|---|
| Description | Defines the URLs where the client will find the certificate revocation list that is related to the certificate. The CRL distribution point of a root CA should be empty. |
| Sample value | [empty] |
| Defined at | Certification Authority MMC |

| | |
|---|---|
| Stored at | Windows 2000 Server family:<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Policy \FileRevocationCRLURL**<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration \Policy\LDAPRevocationCRLURL**<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration \Policy\RevocationCRLURL**<br><br>Windows Server 2003:<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc \Configuration ***CAName***\CRLPublicationURLs** |
| Impacts | Any user, computer, service, or program that verifies the root certificate |

**Authority Information Access (AIA)**

| | |
|---|---|
| Description | Defines the URLs where the client can locate the certificate's issuer certificate. Because a root CA issues the CA certificate to itself, you do not need to specify an issuer. The AIA of a root CA should be empty. |
| Sample value | [empty] |
| Defined at | Certification Authority MMC |
| Stored at | Windows2000 Server family:<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc \Configuration\Policy \FileIssuerCertURL**<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc \Configuration\Policy \LDAPFileIssuerCertURL**<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc \Configuration\Policy\IssuerCertURL**<br><br>Windows Server 2003:<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc \Configuration ***CAName***\CACertPublicationURLs** |
| Impacts | Any user, computer, service, or program that verifies the root certificate |

**CRL Publication Interval**

| | |
|---|---|
| Description | The value controls the CRL validity time and the CRL publication cycle. According to the value, the CRL is published on a regular basis. Its validity time is set to the publication time and date and the defined value. |
| Sample value | 180 days |
| Defined at | Certification Authority MMC |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc \Configuration***CAName***\CRLperiod**<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc \Configuration***CAName***\CRLperiodUnits** |
| Impacts | CA CRL publication algorithm and any user, computer, service, or program that verifies the CRL. |

**Delta CRL Publication Interval**

| | |
|---|---|
| Description | Defines similar to the CRL publication interval and the publication interval of the delta CRL. For an offline CA, it is recommended that you disable delta CRL publication. |

| | |
|---|---|
| Sample value | 0 (which is equal to disabled delta CRL publication) |
| Defined at | Certification Authority MMC |
| Stored at | Windows 2000:<br><br>Not available<br><br>Windows Server 2003<br><br>Registry:<br><br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CRLDeltaPeriod**<br><br>Registry:<br><br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CRLDeltaPeriodUnits** |
| Impacts | Any client that can verify the certificate validity through delta CRLs |

**Validity period**

| | |
|---|---|
| Description | Defines the period of time that a certificate that was issued by the CA is valid. The validity period cannot extend the certificate validity beyond the certificate of the issuing CA. |
| Sample value | 5 years |
| Defined at | Certification Authority MMC |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration***CA Name***\ValidityPeriodUnits**<br><br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration***CA Name***\ValidityPeriod** |
| Impacts | The validity time of any certificate that will be issued from that stand-alone CA. |

**Intermediate CA Configuration Parameters**

This section provides a list of parameters that must be defined during the setup procedure for a stand-alone offline root CA. The sample values are related to the sample configuration that is explained in the previous section.

**CA Policy**

| | |
|---|---|
| Description | Defines the URL or the text that applies to the CA's policy. The policy describes different types of rules, such as how the CA is operated, which legal policies are valid, and so on. |
| Sample value | OID = 1.1.1.1.1.1.1.1.1<br><br>URL = http://www.contoso.com/pki/Policy/USLegalPolicy.asp<br><br>URL = "ftp://ftp.contoso.com/pki/Policy/USLegalPolicy.txt" |
| Defined at | CAPolicy.inf |
| Stored at | CA certificate |
| Impacts | All certificates that are directly or indirectly signed by this CA certificate |

**CSP (CA Certificate)**

| | |
|---|---|
| Description | Generates the certificate's key material and the certificate generation. |

| | |
|---|---|
| Sample value | Microsoft Strong Cryptographic Provider |
| Defined at | CA installation wizard |
| Stored at | CA Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ***CAName***\CSP\Provider** |
| Impacts | CA certificate |

**Hash Algorithm**

| | |
|---|---|
| Description | Defines the hash algorithm that is used for hashing and signing certificate contents. |
| Sample value | SHA-1 |
| Defined at | CA Installation Wizard |
| Stored at | CA registry: *CAName***\CSP\HashAlgorithm** |
| Impacts | CA certificate |

**Key Length (CA Certificate)**

| | |
|---|---|
| Description | Defines the complexity of the key material that is assigned to the CA certificate. It is recommended that the key length does not exceed 4096 bits, because this is the maximum interoperable key length with most applications and PKI providers. The key length of a subordinate CA is typically shorter than the key length of its parent CA. |
| Sample value | 2048 |
| Defined at | CA Installation Wizard |
| Stored at | Certificate request and is only temporarily used |
| Impacts | The root CA key material that could be stored in an HSM or encrypted on the CAs hard disk |

**Common Name**

| | |
|---|---|
| Description | The common name must not exceed 64 characters in length. It is important to remember that each space in the name uses three characters in the total of the overall length because of the escape character sequence (%20). |
| Sample value | IntermediateCA |
| Defined at | CA Installation Wizard |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ***CAName***\CommonName** |
| Impacts | The common name becomes part of the certificate issuer name and is also part of the CRL and AIA if replacement tokens are used. The common name is used by several variables that are used to set the CRL and AIA. |

**CA Database Path**

| | |
|---|---|
| Description | Defines where the CA's database is located in the CA's file system. |
| Sample value | C:\Certlog |
| Defined at | CA Installation Wizard |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration \DBDirectory** |
| Impacts | The CA must be able to obtain the appropriate path name from the registry when the CA starts. |

**CA Log File Path**

| | |
|---|---|
| Description | Defines where the CA's transaction log files are located in the CA's file system. |
| Sample value | D:\Certlog |
| Defined at | CA Installation Wizard |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration \DBLogDirectory** |
| Impacts | The CA must be able to obtain the appropriate path name from the registry when the CA starts. |

**Shared Folder**

| | |
|---|---|
| Description | Defines where the CA's transaction log files are located in the root CA's file system. |
| Sample value | \\{Localhost}\CertConfig |
| Defined at | CA Installation Wizard |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration \ConfigurationDirectory** |
| Impacts | Clients that cannot receive the CA certificate through group policies and need to manually import the certificate. |

**Distinguished Name Suffix**

| | |
|---|---|
| Description | The name maps to the name space that is used by the domain to which the CA belongs. Because the intermediate CA is configured as a stand-alone CA, the distinguished name should be mapped to the same name space that will be used for the enterprise CA. |
| Sample value | Domain ControllerDC=concorp,DC=contoso,DC=com |
| Defined at | CA configuration that occurs after the installation procedure |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ***CAName***\DSConfigDN** |
| Impacts | The distinguished name becomes part of the certificate issuer name and is also part of the CRL and AIA if replacement tokens are used. It is also used by several variables that are used to set the CRL and AIA. |

**CRL Distribution Point**

| | |
|---|---|
| Description | Defines the URLs where the client can locate the certificate revocation list (CRL) that is related to the certificate. |
| Sample value | http://www.contoso.com/pki/%3%8%9.crl<br><br>ldap:///CN=%7%8,CN=%2,CN=CDP,CN=Public Key Services,CN=Services,%6%10 |
| Defined at | CA MMC |

| | |
|---|---|
| Stored at | In Windows 2000:<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Policy \FileRevocationCRLURL**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Policy \LDAPRevocationCRLURL**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Policy \RevocationCRLURL**<br>In Windows Server 2003:<br>Registry: CAName**\CRLPublicationURLs** |
| Impacts | Any user, computer, service, or program that verifies the root certificate |

**Authority Information Access (AIA)**

| | |
|---|---|
| Description | Defines the URLs where the client can locate the certificate's issuer certificate. Because a root CA issues the CA certificate to itself, no issuer needs to be specified. |
| Sample value | http://www.contoso.com/pki/%1_%3%4.crt<br>ldap:///CN=%7,CN=AIA,CN=Public Key Services,CN=Services,%6%11 |
| Defined at | Certification Authority MMC |
| Stored at | In Windows 2000:<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Policy \FileIssuerCertURL**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Policy \LDAPFileIssuerCertURL**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Policy \IssuerCertURL**<br>In Windows Server 2003:<br>Registry: *CAName***\CACertPublicationURLs** |
| Impacts | Any user, computer, service, or program that verifies the root certificate |

**CRL Publication Interval**

| | |
|---|---|
| Description | Also controls also the CRL validity time |
| Sample value | 180 days |
| Defined at | Certification Authority MMC |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ***CAName***\CRLperiod**<br>Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ***CAName***\CRLperiodUnits** |
| Impacts | CA CRL publication algorithm and any user, computer, service, or program that verifies the CRL. |

**Delta CRL Publication Interval**

| | |
|---|---|
| Description | Defines similar to the CRL publication interval and the publication interval of the delta CRL. For an offline CA, it is recommended that you disable delta CRL publication. |
| Sample value | 0 (which is equal to disabled delta CRL publication) |
| Defined at | Certification Authority MMC |
| Stored at | In Windows 2000:<br>Not available.<br>Windows Server 2003<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CRLDeltaPeriod**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CRLDeltaPeriodUnits** |
| Impacts | Any client that can verify the certificate validity through delta CRLs |

**Validity Period**

| | |
|---|---|
| Description | Defines the period of time that a certificate that was issued by the CA is valid. The validity period cannot extend the certificate validity beyond the certificate of the issuing CA. |
| Sample value | 2 years |
| Defined at | Certification Authority MMC |
| Stored at | Windows 2000 and Windows Server 2003<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\ValidityPeriodUnits**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\ValidityPeriod** |
| Impacts | The validity time of any certificate that will be issued from that stand-alone CA. |

**Issuing CA Configuration Parameters**

**CSP (CA Certificate)**

| | |
|---|---|
| Description | The CSP is responsible for generating the certificate's key material and certificate generation. |
| Sample value | Microsoft Strong Cryptographic Provider |
| Defined at | CA Installation Wizard |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration **_CAName_**\CSP\Provider** |
| Impacts | CA certificate |

**Hash Algorithm**

| | |
|---|---|
| Description | Defines the hash algorithm that is used for hashing and signing certificate contents. |

| | |
|---|---|
| Sample value | SHA-1 |
| Defined at | CA Installation Wizard |
| Stored at | CA registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ***CAName***\CSP\HashAlgorithm** |
| Impacts | CA certificate |

**Key Length (CA Certificate)**

| | |
|---|---|
| Description | Defines the complexity of the key material that is assigned to the CA certificate. It is recommended that the key length does not exceed 4096 bits because this is the maximum interoperable key length with most applications and PKI providers. The key length of a subordinate CA is typically shorter than the key length of its parent CA. |
| Sample value | 2048 |
| Defined at | CA Installation Wizard |
| Stored at | Certificate request and is only used temporarily |
| Impacts | CA key material |

**Common Name**

| | |
|---|---|
| Description | The common name must not exceed 64 characters in length. It is important to remember that each space in the name uses three characters in the total of the overall length because of the escape character sequence (%20). |
| Sample value | CorporateEntCA |
| Defined at | CA Installation Wizard |
| Stored at | Windows 2000 and Windows Server 2003:<br><br>Registry:<br><br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CommonName** |
| Impacts | The common name becomes part of the certificate issuer name and is also part of the CRL and AIA if replacement tokens are used. The common name is used by several variables that are used to set the CRL and AIA. |

**CA Database Path**

| | |
|---|---|
| Description | Defines where the CA's database is located in the CA's file system. |
| Sample value | D:\Certlog |
| Defined at | CA Installation Wizard |
| Stored at | Registry:<br><br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration \DBDirectory** |
| Impacts | The CA must be able to obtain the appropriate path name from the registry when the CA starts. |

**CA Log File Path**

| | |
|---|---|
| Description | Defines where the CA's transaction log files are located in the root CA's file system. |

| | |
|---|---|
| Sample value | D:\Certlog |
| Defined at | CA Installation Wizard |
| Stored at | Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration \DBLogDirectory** |
| Impacts | The CA must be able to obtain the appropriate path name from the registry when the CA starts. |

**Shared folder**

| | |
|---|---|
| Description | Defines where the CA's transaction log files are located in the root CA's file system. The shared folder is not required for an enterprise CA. |
| Sample value | \\Localhost\CertConfig |
| Defined at | CA Installation Wizard |
| Stored at | User-defined location during installation |
| Impacts | Clients that cannot receive the CA certificate through group policies and need to manually import the certificate. |

**Distinguished Name Suffix**

| | |
|---|---|
| Description | The name space is automatically mapped to the Active Directory namespace. The value is predefined because of the domain membership of the CA. |
| Sample value | CN=Configuration,DC=concorp,DC=contoso,DC=com |
| Defined at | Automatically defined |
| Stored at | Registry: **HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ***CAName***\DSConfigDN** |
| Impacts | The distinguished name becomes part of the certificate issuer name and is also part of the CRL and AIA if replacement tokens are used. It is also used by several variables that are used to set the CRL and AIA. |

**CRL Distribution Point**

| | |
|---|---|
| Description | Defines the URLs where the client can locate the certificate revocation list that is related to the certificate. |
| Sample value | http://www.contoso.com/pki/%3%8%9.crl<br><br>ldap:///CN=%7%8,CN=%2,CN=CDP,CN=Public Key Services,CN=Services,%6%10 |
| Defined at | Certification Authority MMC |

| | |
|---|---|
| Stored at | Windows 2000:<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Policy \FileRevocationCRLURL**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Policy \LDAPRevocationCRLURL**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Policy \RevocationCRLURL**<br>Windows Server 2003:<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CRLPublicationURLs** |
| Impacts | Any user, computer, service, or program that verifies the root certificate |

**Authority Information Access (AIA)**

| | |
|---|---|
| Description | Defines the URLs where the client can find the certificate's issuer certificate. |
| Sample value | http://www.contoso.com/pki/%1_%3%4.crt<br>ldap:///CN=%7,CN=AIA,CN=Public Key Services,CN=Services,%6%11 |
| Defined at | Certification Authority MMC |
| Stored at | In Windows 2000:<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc \Configuration\Policy\FileIssuerCertURL**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc \Configuration\Policy\LDAPFileIssuerCertURL**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc \Configuration\Policy\IssuerCertURL**<br>In Windows Server 2003:<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc \Configuration** *CAName* **\CACertPublicationURLs** |
| Impacts | Any user, computer, service, or program that verifies the root certificate |

**CRL Publication Interval**

| | |
|---|---|
| Description | Also controls the CRL validity time |
| Sample value | 7 days |
| Defined at | Certification Authority MMC |

| | |
|---|---|
| Stored at | Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CRLPeriod**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CRLPeriodUnits** |
| Impacts | CA CRL publication algorithm and any user, computer, service, or computer that verifies the CRL. |

**Delta CRL publication interval**

| | |
|---|---|
| Description | Defines similar to the CRL publication interval and the publication interval of the delta CRL. For an offline CA, it is recommended that you disable delta CRL publication. |
| Sample value | 1 day |
| Defined at | Certification Authority MMC |
| Stored at | Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CRLDeltaPeriod**<br>Registry:<br>**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration ** *CAName* **\CRLDeltaPeriodUnits** |
| Impacts | Any client that can verify the certificate validity through delta CRLs |

# Appendix C: Additional Information

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

For additional information, see the following articles.

293819 How to Remove a Root Certificate from the Trusted Root Store on the Microsoft Knowledge Base

Windows XP Home Page on the Microsoft Web site

PKI Enhancements in Windows XP Professional and Windows Server 2003 on the Microsoft Web site

Windows XP Technical Resources on the Microsoft Web site

RFC 2797: Certificate Management Messages over CMS on the Internet Engineering Task Force Web site

White Paper: "Troubleshooting Certificate Status and Revocation" on the Microsoft TechNet Web site

White Paper: Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003 on the Microsoft TechNet Web site

White Paper: Implementing and Administering Certificate Templates in Windows Server 2003 on the Microsoft TechNet Web site

White Paper: Key Archival and Management in Windows Server 2003 on the Microsoft Web site

White Paper: Windows Server 2003 PKI Operations Guide on the Microsoft TechNet Web site

Certificate Autoenrollment in Windows Server 2003 on the Microsoft TechNet Web site

# Server Clusters: Network Configuration Best Practices for Windows 2000 and Windows Server 2003

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Published: January 1, 2003

Microsoft Corporation

This document provides Server cluster requirements and best practices for the network infrastructure for Server clusters running Microsoft Windows 2000 or Windows Server 2003. Requirements must hold for the cluster to function correctly. Best practices are suggestions based on deployment feedback and issues seen in the field.

**In This White Paper**

Cluster Networking Requirements

Cluster Networking Best Practices

Procedures for Implementing Cluster Networking Best Practices

IPSec in Cluster Networking

NetBIOS in Cluster Networking

# Cluster Networking Requirements

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section describes the requirements that Server cluster puts on the network infrastructure. These requirements must be met for the Server cluster solution to function correctly.

**General Requirements**
This section describes the requirements for any Server cluster deployment.

- The complete hardware configuration for a cluster must be selected from the Cluster Hardware Compatibility List (HCL). The network interface controllers along with any other components used in certified cluster configurations must have the Windows logo and appear on the Microsoft Hardware Compatibility List.

  (Windows 2000, Windows Server 2003)

| **Note** |
| --- |
| A cluster built from logged components that do not appear on the cluster HCL is NOT a qualified configuration. |

- Two or more independent networks must connect the nodes of a cluster in order to avoid a single point of failure. The use of two local area networks (LANs) is required; cluster configurations using a single network are not being supported.

  (Windows 2000, Windows Server 2003)

- Each cluster network must fail independently of all other cluster networks. That is, two cluster networks must not have a component in common that can cause both to fail simultaneously. For example, the use of a multi-port network interface controller to attach a node to two cluster networks would not meet this requirement in most cases because the ports are not independent. Likewise, two networks that share a switch could also have a single point of failure. The simplest way to ensure that your cluster meets this requirement is to use physically independent components to construct cluster networks.

  (Windows 2000, Windows Server 2003)

- All of the adapters used to attach nodes to the same cluster network must use the same communication settings – e.g. the same Speed, Duplex Mode, Flow Control, and Media Type. If the adapters are connected to a switch, the port settings of the switch must match those of the adapters.

  (Windows 2000, Windows Server 2003)

- Each cluster network must be configured as a single IP subnet whose subnet number is distinct from those of other cluster networks. For example, a cluster could use two networks configured with the following subnet addresses: 10.1.x.x and 10.2.x.x with mask 255.255.0.0. Addresses may be assigned to the nodes dynamically by DHCP, but manual configuration with static addresses is recommended (see section Cluster Networking Best Practices). The use of Automatic Private IP Addressing (APIPA) to configure cluster networks is not supported. Also APIPA is not designed for use with computers that are attached to multiple networks.

  (Windows 2000, Windows Server 2003)

- At least two of the cluster networks must be configured to support internal communication between cluster nodes in order to avoid a single point of failure. That is, the roles of these networks must be configured as either Internal Cluster Communications Only or All Communications in Cluster Service. Typically, one of these networks is a private interconnect dedicated to internal cluster communication (see section Cluster Networking Best Practices).

  (Windows 2000, Windows Server 2003)

- The use of network interface controller teaming on all cluster networks concurrently is not supported. At least one of the cluster networks that are enabled for internal communication between cluster nodes must not be teamed. Typically, the unteamed network is a private interconnect dedicated to this type of communication. The use of network interface controller teaming on other cluster networks is acceptable; however, if communication problems occur on a teamed network, Microsoft Product Support Services may require that teaming be disabled. If this action resolves the problem or issue, then you must seek further assistance from the manufacturer of the teaming solution.

  (Windows 2000, Windows Server 2003)

- The nodes of a cluster must belong to a single domain. The domain configuration must meet the following requirements in order to avoid a single point of failure in the authentication process:

  - The domain must have at least two domain controllers,

  - If DNS is used to resolve names in the domain, then at least two DNS servers must be deployed as well. The DNS servers should support dynamic updates,

  - Each domain controller and cluster node must be configured with a primary and at least one secondary DNS server. If the domain controllers are also DNS servers, then each should point to itself for primary DNS resolution and to the other DNS servers for secondary resolution,

  - At least two domain controllers must be configured to be Global Catalog Servers.

  (Windows 2000, Windows Server 2003)

**Geographically Dispersed Clusters**

This section describes the additional requirements that geographically dispersed clusters have:

- The nodes in a cluster may be on different physical networks; however, the private and public network connections between cluster nodes must appear as a single, non-routed LAN using technologies such as virtual LANs (VLANs).

  (Windows 2000, Windows Server 2003)

- The round-trip communication latency between any pair of cluster nodes must be no more than 500 milliseconds.

  (Windows 2000, Windows Server 2003)

- As with LANs, each VLAN must fail independently of all other cluster networks.

  (Windows 2000, Windows Server 2003)

- Due to the complexity of geographically dispersed clusters, you need to involve the hardware manufacturer or hardware vendor in any issue. Often, there is third-party software and drivers that are required for the clusters to function. Microsoft Product Support Services may not be aware of how these components interact with Windows Clustering.

  (Windows 2000, Windows Server 2003)

# Cluster Networking Best Practices

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section describes network best practices for deploying a Server cluster.

**Hardware Planning Recommendations**

- Use identical network interface controllers in all cluster nodes; that is, each adapter should be the same make, model, and firmware version.

  (Windows 2000, Windows Server 2003)

- Reserve one network exclusively for internal communication between cluster nodes. This is the private network. Use other networks for communication with clients. These are public networks. Do not use network interface controllers teaming on the private network.

  (Windows 2000, Windows Server 2003)

**Network Interface Controller Configuration Recommendations**

- Manually select the speed and duplex mode of each cluster network interface controller. Do not use automatic detection. Some adapters drop packets while automatically negotiating the settings of the network. All adapters on a network must be configured to use the same speed and duplex mode. If the adapters are connected to a switch, ensure that the port settings of the switch match those of the adapters.

  (Windows 2000, Windows Server 2003)

- Use static IP addresses for all nodes on the private network. Choose the addresses from one of the following ranges.

  - 10.0.0.0 - 10.255.255.255 (Class A network)
  - 172.16.0.0 - 172.31.255.255 (Class B network)
  - 192.168.0.0 - 192.168.255.255 (Class C network)

  (Windows 2000, Windows Server 2003)

- Use static IP addresses for all nodes on all public networks. The use of dynamic configuration via DHCP is not recommended. Failure to renew a lease could disrupt cluster operation.

  (Windows 2000, Windows Server 2003)

- Do not configure DNS servers, WINS servers, or a default gateway on the private network interface controllers.

  (Windows 2000, Windows Server 2003)

- You should configure WINS or DNS servers on the public network interface controllers. If network name resources will be deployed on the public network, then the DNS servers should support dynamic updates; otherwise, name-to-IP address mappings will not be promptly updated during fail over.

  (Windows 2000, Windows Server 2003)

- Do configure a default gateway on the public network interface controllers if cluster nodes use those network interface controllers to communicate with clients or services on remote subnets. Be aware that in clusters with multiple public

networks, configuring nodes with a default gateway on more than one network can cause routing problems.

(Windows 2000, Windows Server 2003)

- On each cluster node, set the network connection order to be:
    - Public network(s) – highest priority
    - Private network
    - Remote access connections – lowest priority

  (Windows 2000, Windows Server 2003)

- Change the default name for each network connection to clearly identify the use of each network. For example, you might change the name of the private network connection from Local Area Connection (x) to Private Cluster Network.

  (Windows 2000, Windows Server 2003)

- The private LAN should be isolated. Only nodes that are part of the cluster should be connected to the private subnet. Where there are several clusters, using the same subnet for the private network for all of the clusters is reasonable. You should not, however, put other network infrastructure such as domain controllers, WINS server, DHCP servers etc. on the private subnet.

  (Windows 2000, Windows Server 2003)

- To create an isolated network segment, you should use a non-intelligent fabric for the heartbeat network to include hubs, unmanaged layer 2 switches (dumb switches), or in the case of a 2-node Server cluster, a crossover cable.

  (Windows 2000, Windows Server 2003)

- You should disable the default media sense policy for TCP/IP to ensure that, if cables are disconnected or media sense is lost, the TCP/IP configuration and corresponding cluster network configuration are not torn down. Add the following registry value to each node:

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters**

  **Value Name: DisableDHCPMediaSense**

  **Data Type: REG_DWORD**

  **Data: 1**

  (Windows 2000)

**Cluster Service Configuration Recommendations**

- Set the private network role to Internal Cluster Communications Only. Verify that the role for each public network is set to All Communications (this is the default value).

  (Windows 2000, Windows Server 2003)

- Configure the private network to be the highest priority for internal cluster communication.

  (Windows 2000, Windows Server 2003)

# Procedures for Implementing Cluster Networking Best Practices

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

This section describes the procedures for implementing best practices:

**Configuring Network Interface Controllers Prior to Configuring Cluster Service**

- Configure the speed of each network interface controller as follows:
    - Open Control Panel. Open Network Connections. Right-click the appropriate connection object and select Properties. Click Configure, and then Advanced.
    - Set the desired network speed using the drop-down list.
    - Ensure that other settings, such as Duplex Mode, are also the same for all adapters on a network.
- Configure the Internet Protocol settings of the private network interface controller as follows:
    - Return to Network Connections. Open the Properties of the appropriate connection object.
    - Ensure that the Internet Protocol (TCP/IP) check box is selected.
    - Highlight Internet Protocol and select Properties.
    - Click the radio-button for Use the Following IP Address and enter a static address.
    - Ensure that there is no default gateway configured for the private network.
    - Verify that there are no values defined in the Use the Following DNS Server Addresses box. Click Advanced. On the DNS tab, verify that there are no values defined. Make sure that the Register this Connection's Addresses in DNS and Use this Connection's DNS Suffix in DNS Registration check boxes are cleared. Note that if the cluster node is a DNS server, then IP address 127.0.0.1 will appear in the list and should remain there.
- Configure the network connection order as follows:
    - Return to Network Connections. Select Advanced. Select Advanced Settings. In the Connections box, order the available network connections as follows:

      Public network(s)

      Private network

      Remote access connections
- Change the default names of the network connections as follows:
    - Return to Network Connections. Right-click a network connection object. Select Rename. Edit the name value.
    - The name used for the connection object that represents a network, such as the private network, should be the same on all nodes. If the names of the connection objects are not the same, Cluster Service will choose one and change the others to match it.

**Configuring Cluster Network Properties after Configuring Cluster Service**

### Windows 2000

When installing the cluster software, the Configuring Cluster Networks dialog box will be presented for each network (in arbitrary order). For public networks, make sure the name and IP address match the network interface for the public network. Check the box "Enable this network for cluster use". Select the option "All communications (mixed network)". For the private network, make sure name and IP address match network interface for private network. Check the box "Enable this network for cluster use". Select the option "Internal cluster communications only".

The default configuration during installation is to configure your public network adapter for "All Communication" and the private (heartbeat) network adapter for "Internal Cluster Communications." Microsoft recommends that you keep this default configuration. For your cluster to install and function properly, you must configure at least one of the networks for "Internal Cluster Communications" or "All Communications."

### Windows Server 2003

The Windows Server 2003 cluster configuration wizard does not provide a way to change the network settings during configuration. The default setting for all networks is to enable All Communications. This will ensure that the cluster can operate correctly. To conform to best practices, you should make one of the networks into a private network and make the private network the highest priority for internal cluster communications as follows:

- Set the private network role to be Internal Cluster Communications as follows:
    - In cluster administrator, double-click the cluster name. You will see a folder Cluster Configuration.
    - Double-click the Cluster Configuration folder and double-click the Networks folder to see all the available cluster networks.
    - Select the network you are configuring for the private cluster communications and chose Properties.
    - In the Properties for the network you will see roles such as Client access only. For the private network make sure that the Enable this network for cluster use check box is checked and that the Internal Communications Only role is selected.
- Configure the private network to be the highest priority for internal cluster communication as follows:
    - In Cluster Administrator, select the cluster and choose Properties.
    - From the Network Priority tab, verify that the private network is listed at the top.
    - If it is not, use the Move Up button to increase its priority.

# IPSec in Cluster Networking

Updated: January 1, 2003

Applies To: Windows Server 2003 with SP1

Although it is possible to use Internet Protocol Security (IPSec) for applications that can failover in a Server cluster, IPSec was not designed for failover situations and we recommend that you do NOT use IPSec for applications in a Server cluster.

The primary issue is that Internet Key Exchange (IKE) Security Associations (SAs) are not transferred from one server to the other if a failover occurs because they are stored in a local database on each node.

In a connection that is protected by IPSec, an IKE SA is created in phase-I negotiations. Two IPSec SAs are created in phase II. A time-out value is associated with the IKE and IPSec SAs. If Master Perfect Forward Secrecy is not used, the IPSec SAs are created by using key material from the IKE SAs. If this is the case, the client must wait for the default time-out or lifetime period for the inbound IPSec SA to expire and then wait for the timeout or lifetime period that is associated with the IKE SA.

The default time-out for the Security Association Idle Timer is 5 minutes, in the event of a failover, clients will not be able to reestablish connections until at least 5 minutes after all resources are online, using IPSec.

Although IPSec is not optimally designed for a clustered environment, it may be used if your business need for secure connectivity outweighs client downtime in the event of a failover.

# NetBIOS in Cluster Networking

Updated: May 5, 2006

Applies To: Windows Server 2003 with SP1

In Windows Server 2003, the cluster service does not require NetBIOS. A basic principle of server security is to disable unneeded services. To determine whether to disable NetBIOS, consider the following:

- Some services and applications other than the cluster service use NetBIOS. Review all the ways a clustered server functions before disabling NetBIOS on it.
- With NetBIOS disabled, you will not be able to use the Browse function in Cluster Administrator when opening a connection to a cluster. Cluster Administrator uses NetBIOS to enumerate all clusters in a domain.
- With NetBIOS disabled, Cluster Administrator does not work if a cluster name is specified. Cluster Administrator calls GetNodeClusterState, which uses the remote registry APIs that, in turn, use named pipes based on the virtual name.

If you want to disable NetBIOS on a clustered server, review the following items:

- By default, when a cluster is configured, NetBIOS is enabled on the cluster IP Address resource. To disable NetBIOS on a new cluster, clear the check box on the parameters page of the Cluster IP Address resource property sheet.
- When you create additional IP Address resources, to disable NetBIOS, clear the NetBIOS check box.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

2002. Microsoft Corporation. All rights reserved. Microsoft, Windows, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

# Virtual Private Networking with Windows Server 2003: Deploying Remote Access VPNs

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

**In This White Paper**


Introduction to Virtual Private Networking with Windows Server 2003: Deploying Remote Access VPNs

Components of Windows Remote Access VPNs

Deploying PPTP-based Remote Access

Deploying L2TP/IPSec-based Remote Access

Appendix A: Configuring Firewalls for a VPN Server Running Windows Server 2003

Appendix B: Alternate Configurations

Appendix C: Setting up a VPN Test Lab

Appendix D: Troubleshooting

Appendix E: Deploying a Certificate Infrastructure

Summary and Related Links



Summary and Related Links

# Introduction to Virtual Private Networking with Windows Server 2003: Deploying Remote Access VPNs

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

A virtual private network (VPN) is the extension of a private network that encompasses links across shared or public networks like the Internet. With a VPN, you can send data between two computers across a shared or public network in a manner that emulates a point-to-point private link (such as a long haul T-Carrier-based wide area network [WAN] link). Virtual private networking is the act of creating and configuring a virtual private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information, which allows the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The link in which the private data is encapsulated and encrypted is a VPN connection.

Figure 1 shows the logical equivalent of a VPN connection.




**Figure 1: The logical equivalent of VPN connections**

Users working at home or on the road can use VPN connections to establish a remote access connection to an organization server by using the infrastructure provided by a public network such as the Internet. From the user's perspective, the VPN connection is a point-to-point connection between the computer (the VPN client) and an organization server (the VPN server). The exact infrastructure of the shared or public network is irrelevant because it appears logically as if the data is sent over a dedicated private link.

Organizations can also use VPN connections to establish routed connections with geographically separate offices or with other organizations over a public network such as the Internet while maintaining secure communications. A routed VPN connection across the Internet logically operates as a dedicated WAN link.

With both remote access and routed connections, an organization can use VPN connections to trade long-distance dial-up or leased lines for local dial-up or leased lines to an Internet service provider (ISP).

**There are two types of remote access VPN technology in the Windows Server 2003 operating system:**

1. Point-to-Point Tunneling Protocol (PPTP)

   PPTP uses user-level Point-to-Point Protocol (PPP) authentication methods and Microsoft Point-to-Point Encryption (MPPE) for data encryption.

2. Layer Two Tunneling Protocol with Internet Protocol security (L2TP/IPSec)

   L2TP uses user-level PPP authentication methods and IPSec for computer-level authentication using certificates and data authentication, integrity, and encryption.

A remote access client (a single user computer) makes a remote access VPN connection that connects to a private network. The

VPN server provides access to the entire network to which the VPN server is attached. The packets sent from the remote client across the VPN connection originate at the remote access client computer.

The remote access client (the VPN client) authenticates itself to the remote access server (the VPN server) and, for mutual authentication, the server authenticates itself to the client.

Computers running Windows Server 2003, Windows XP, Windows 2000, Windows NT version 4.0, Windows Millennium Edition, and Windows 98 operating systems can create remote access VPN connections to a VPN server running Windows Server 2003. VPN clients may also be any non-Microsoft PPTP client or L2TP client using IPSec.

| **Note** |
| --- |
| Using IPSec tunnel mode is not a remote access VPN technology supported by Microsoft VPN clients or servers due to the lack of an industry standard method of performing user authentication and IP address configuration over an IPSec tunnel. IPSec tunnel mode is described in RFCs 2401, 2402, and 2406. |

**For encryption, you can use either link encryption or end-to-end encryption in addition to link encryption:**

- Link encryption encrypts the data only on the link between the VPN client and the VPN server. For PPTP connections, you must use MPPE in conjunction with MS-CHAP, MS-CHAP v2, or EAP-TLS authentication. For L2TP/IPSec connections, IPSec provides encryption on the link between the VPN client and the VPN server.

- End-to-end encryption encrypts the data between the source host and its final destination. You can use IPSec after the VPN connection is made to encrypt data from the source host to the destination host.

# Components of Windows Remote Access VPNs

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

Figure 2 shows the components of Windows remote access virtual private networks.




**Figure 2: Components of Windows remote access VPNs**

**The major components are:**

- VPN clients
- Internet infrastructure
- VPN server
- Intranet infrastructure
- Authentication, authorization, and accounting (AAA) infrastructure
- Certificate infrastructure

**VPN Clients**

The VPN client can be any computer that is capable of creating a PPTP connection using MPPE or L2TP connection using IPSec encryption. Table 1 lists the VPN-capable Microsoft operating systems.

**Table 1 VPN-Capable Microsoft Operating Systems**

| VPN Tunneling Protocol | Microsoft Operating System |
|---|---|
| PPTP | Windows Server 2003, Windows XP, Windows 2000, Windows NT version 4.0, Windows Millennium Edition, or Windows 98 |
| L2TP/IPSec | Windows Server 2003, Windows XP, Windows 2000, and Windows NT 4.0 Workstation, Windows Millennium Edition, and Windows 98 with Microsoft L2TP/IPSec VPN Client |

**Typical VPN clients are:**

- Laptop users who connect to the organization intranet to access e-mail and other resources while traveling.

- Telecommuters who use the Internet to access organization resources from home.
- Remote administrators who use the Internet to connect to an organization network and configure network or application services.

Microsoft VPN clients can configure VPN connections either manually or by using the Connection Manager components available in Windows Server 2003. To manually configure a Windows 2000 VPN client, use Make New Connection in the Network and Dial-up Connections folder to create a VPN connection to the IP address or DNS name of the VPN server on the Internet. To manually configure a Windows XP VPN client, use the New Connection Wizard in the Network Connections folder to create a VPN connection to the IP address or DNS name of the VPN server on the Internet.

**Connection Manager**

When scaling the configuration of VPN connections for an enterprise, there are the following problems:

- The exact procedure to configure a VPN connection varies depending on the version of Windows running on the client computer.
- To prevent configuration errors, it is preferable to have the information technology (IT) staff configure the VPN connection rather than end users.
- A configuration method must be able to scale to hundreds or thousands of client computers in a large organization.
- A VPN connection may need a double-dial configuration, where a user must dial the Internet first before creating a VPN connection with the organization intranet.

The solution to these issues of configuring VPN connections across an enterprise is Connection Manager. Connection Manager consists of the following:

- Connection Manager
- Connection Manager Administration Kit
- Connection Point Services

**Connection Manager**

Connection Manager is a client dialer, included in Windows Server 2003, whose advanced features make it a superset of basic dial-up networking. Windows Server 2003 includes a set of tools that enables a network manager to deliver pre-configured connections to network users. These tools are the Connection Manager Administration Kit (CMAK) and Connection Point Services (CPS).

Connection Manager provides support for local and remote connections to your service using a network of access points, such as those available worldwide through ISPs. If your service requires secure connections over the Internet, you can also use Connection Manager to establish VPN connections to your service.

**Connection Manager Administration Kit**

A network administrator can tailor the appearance and behavior of a connection made with Connection Manager by using CMAK. With CMAK, an administrator can develop client dialer and connection software that allows users to connect to the network by using only the connection features that the administrator defines for them. Connection Manager supports a variety of features that both simplify and enhance implementation of connection support for you and your users, most of which can be incorporated using the Connection Manager Administration Kit Wizard.

CMAK allows you to build profiles customizing the Connection Manager installation package that you deliver to your customers, so that Connection Manager reflects the identity of your organization. It allows you to determine which functions and features you want to include and how Connection Manager appears to your customers. You can do this by using the Connection Manager Administration Kit Wizard to build custom service profiles.

For more information about CMAK and the configuration of connection manager service profiles, see Windows Server 2003 Help and Support.

**Connection Point Services**

Connection Point Services (CPS) enables you to automatically distribute and update custom phone books. These phone books contain one or more Point of Presence (POP) entries, with each POP supplying a telephone number that provides dial-up access to an Internet access point. The phone books give users complete POP information, so when they travel they can connect to different Internet access points rather than being restricted to a single POP.

Without the ability to update phone books (a task CPS handles automatically), users would have to contact their organization's technical support staff to be informed of changes in POP information and to reconfigure their client dialer software.

**CPS has two components:**

1. Phone Book Administrator

   A tool used to create and maintain the phone book database and to publish new phone book information to the Phone Book Service.

2. Phone Book Service

   A Microsoft Internet Information Services (IIS) extension that runs on Windows NT Server 4.0 or later (with IIS). Phone Book Service automatically checks subscribers' or corporate employees' current phone books and, if necessary, downloads a phone book update.

For more information about CPS and the configuration of phone books, see Windows Server 2003 Help and Support.

**Single Sign-on**

Single sign-on is the capability that allows a remote access user to create a remote access connection to an organization and logon to the organization's domain by using the same set of credentials. For a domain-based infrastructure, the user name and password or smart card is used for both authenticating and authorizing a remote access connection and for authenticating and logging on to a Windows domain. Single sign-on is performed by selecting the **Logon by using dial-up networking option** on the Windows XP and Windows 2000 logon dialog box and then selecting a dial-up or VPN connection to use to connect to the organization.

For VPN connections, the user must first connect to the Internet before creating a VPN connection. After the Internet connection is made, the VPN connection and logon to the domain can be accomplished. If there is a separate ISP account that the user uses to connect to the Internet, you can create a dial-up connection with the ISP credentials already configured. Then, configure your VPN connection to dial the ISP connection before attempting the VPN connection. In this configuration, the user will never have to type the ISP credentials when logging on to the domain. This association between the VPN connection and the ISP connection can be configured manually or by using Connection Manager.

**Installing a Certificate on a Client Computer**

If your Windows 2000 or Windows XP VPN clients are either making L2TP connections or using certificates for user-level authentication, certificates must be installed on the VPN client computer. For L2TP connections, a computer certificate must be installed on the VPN client computer to provide authentication for establishing an IPSec security association (SA). For user-level authentication using the Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication protocol, you can either use a user certificate or a smart card.

For user certificate-based authentication, the computer user must request a user certificate from a Windows Server 2003 certification authority (CA) on your intranet. For smart card-based authentication, a network administrator must configure an enrollment station and issue smart cards with certificates that are mapped to individual user accounts.

For more information about installing certificates on VPN client computers, see Certificate Infrastructure in this paper.

**Design Points: Configuring the VPN client**

Consider the following when configuring your VPN clients for remote access VPN connections:

- If you have a small number of VPN clients, perform manual configuration of VPN connections on each computer.

- If you have a large number of VPN clients or they are running different versions of Microsoft operating systems, use the Connection Manager components of Windows Server 2003 to create the custom VPN connection configuration package for distribution and to maintain the phone book database for your POPs.

- If you are using Windows XP, Windows 2000, or Microsoft L2TP/IPSec VPN Client to make L2TP connections, you must install a computer certificate on the VPN client computer.

- If you are using Windows XP or Windows 2000 VPN clients and user-level certificate authentication with EAP-TLS, you must either install a user certificate on the VPN client computer or a user certificate on the smart card used by the VPN client computer.

**Internet Network Infrastructure**
To create a VPN connection to a VPN server across the Internet:

- The VPN server's name must be resolvable.
- The VPN server must be reachable.
- VPN traffic must be allowed to and from the VPN server.

**VPN Server Name Resolvability**
In most cases you want to reference the VPN server by name, rather than an IP address, as names are much easier to remember. You can use a name (for example VPN1.example.microsoft.com) as long as the name can be resolved to an IP address. Therefore, you must ensure that whatever name you are using for your VPN servers when configuring a VPN connection, that name must be able to be resolved to an IP address using the Internet Domain Name System (DNS) infrastructure.

When you use names rather than addresses, you can also take advantage of DNS round robin load balancing if you have multiple VPN servers with the same name. Within DNS, you can create multiple records that resolve a specific name to different IP address. In this situation, DNS servers send back all the addresses in response to a DNS name query and randomize the order of the addresses for successive queries. Because most DNS clients use the first address in the DNS query response, the result is that VPN client connections are on average spread across the VPN servers.

**VPN Server Reachability**
To be reachable, the VPN server must be assigned a public IP address to which packets are forwarded by the routing infrastructure of the Internet. If you have been assigned a static public IP address from an ISP or an Internet registry, this is typically not an issue. In some configurations, the VPN server is actually configured with a private IP address and has a published static IP address by which it is known on the Internet. A device between the Internet and the VPN server translates the published and actual IP addresses of the VPN server in packets to and from the VPN server.

While the routing infrastructure might be in place, the VPN server might be unreachable due to the placement of firewalls, packet filtering routers, network address translators, security gateways, or other types of devices that prevent packets from either being sent to or received from the VPN server computer.

**VPN Servers and Firewall Configuration**
There are two approaches to using a firewall with a VPN server:

1. The VPN server is attached directly to the Internet and the firewall is between the VPN server and the intranet.

   In this configuration, the VPN server must be configured with packet filters that only allow VPN traffic in and out of its Internet interface. The firewall can be configured to allow specific types of remote access traffic.

2. The firewall is attached to the Internet and the VPN server is between the firewall and the intranet.

   In this configuration, both the firewall and the VPN server are attached to a network segment known as the perimeter network (also known as a screened subnet). Both the firewall and the VPN server must be configured with packet filters that allow only VPN traffic to and from the Internet. Figure 2 shows this configuration.

For the details of configuring packet filters for the VPN server and the firewall for both of these configurations, see Appendix A.

**Design Points: VPN Server Accessibility from the Internet**
Consider the following when configuring your Internet infrastructure for remote access VPN connections:

- Ensure that the DNS names of your VPN servers are resolvable from the Internet by either placing an appropriate DNS record in your Internet DNS server or the DNS server of your ISP. Test the resolvability by using the Ping tool to ping the name of each of your VPN server when directly connected to the Internet. Due to packet filtering, the result of the ping command may be "Request timed out", but check to ensure that the name specified was resolved by the Ping tool to the proper IP address.

- Ensure that the IP addresses of your VPN servers are reachable from the Internet by using the Ping tool to ping the name or address of your VPN server with a 5 second timeout (using the -w command line option) when directly connected to the Internet. If you see a "Destination unreachable" error message, the VPN server is not reachable.

- Configure packet filtering for PPTP traffic, L2TP traffic, or both types of traffic on the appropriate firewall and VPN server interfaces connecting to the Internet and the perimeter network. For more information, see Appendix A.

**Authentication Protocols**
To authenticate the user who is attempting to create a PPP connection, Windows Server 2003 supports a wide variety of PPP authentication protocols including:

- Password Authentication Protocol (PAP)

- Challenge-Handshake Authentication Protocol (CHAP)

- Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)

- MS-CHAP version 2 (MS-CHAP v2)

- Extensible Authentication Protocol-Message Digest 5 (EAP-MD5)

- Extensible Authentication Protocol-Transport Level Protocol (EAP-TLS)

For PPTP connections, you must use MS-CHAP, MS-CHAP v2, or EAP-TLS. Only these three authentication protocols provide a mechanism to generate the same encryption key on both the VPN client and the VPN server. MPPE uses this encryption key to encrypt all PPTP data sent on the VPN connection. MS-CHAP and MS-CHAP v2 are password-based authentication protocols.

In the absence of user certificates or smart cards, MS-CHAP v2 is highly recommended as it is a stronger authentication protocol than MS-CHAP and provides mutual authentication. With mutual authentication, the VPN server authenticates the VPN client and the VPN client authenticates the VPN server.

| **Note** |
| --- |
| If you must use a password-based authentication protocol, enforce the use of strong passwords on your network. Strong passwords are long (greater than 8 characters) and contain a random mixture of upper and lower case letters, numbers, and punctuation. An example of a strong password is f3L*02~>xR3w#4o. In an Active Directory service domain, use Group Policy settings to enforce strong user passwords. |

EAP-TLS is used in conjunction with a certificate infrastructure and either user certificates or smart cards. With EAP-TLS, the VPN client sends its user certificate for authentication and the VPN server sends a computer certificate for authentication. This is the strongest authentication method as it does not rely on passwords.

| **Note** |
| --- |
| You can use third-party CAs. For information, see Appendix E. |

For L2TP/IPSec connections, any authentication protocol can be used because the authentication occurs after the VPN client and VPN server have established a secure channel of communication known as an IPSec security association (SA). However, the use of either MS-CHAP v2 or EAP-TLS is recommended to provide strong user authentication.

**Design Point: Which Authentication Protocol to Use?**
Consider the following when choosing an authentication protocol for VPN connections:

- If you are using smart cards or have a certificate infrastructure that issues user certificates, use the EAP-TLS authentication protocol for both PPTP and L2TP connections. Only VPN clients running Windows XP and Windows 2000 support EAP-TLS.

- If you must use a password-based authentication protocol, use MS-CHAP v2 and enforce strong passwords using Group

Policy. MS-CHAP v2 is supported by computers running Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0 with Service Pack 4 and later, Windows Millennium Edition, and Windows 98.

**VPN Protocols**
Windows Server 2003 includes support for two remote access VPN protocols:

1. Point-to-Point Tunneling Protocol

2. Layer Two Tunneling Protocol with IPSec

**Point-to-Point Tunneling Protocol**
Introduced in Windows NT 4.0, PPTP leverages Point-to-Point Protocol (PPP) user authentication and Microsoft Point-to-Point Encryption (MPPE) to encapsulate and encrypt IP traffic. When version 2 of the Microsoft Challenge Handshake Authentication Protocol (MS-CHAP v2) is used with strong passwords, PPTP is a secure VPN technology. For nonpassword-based authentication, Extensible Authentication Protocol-Transport Level Security (EAP-TLS) can be used to support smart cards. PPTP is widely supported, easily deployed, and can be used across network address translators (NATs).

**Layer Two Tunneling Protocol with IPSec**
L2TP leverages PPP user authentication and IPSec encryption to encapsulate and encrypt IP traffic. This combination, known as L2TP/IPSec, uses certificate-based computer identity authentication to create the IPSec session in addition to PPP-based user authentication. L2TP/IPSec provides data integrity and data origin authentication for each packet. However, L2TP/IPSec requires a certificate infrastructure to allocate computer certificates and is supported by Windows Server 2003, Windows XP, Windows 2000, and Microsoft L2TP/IPSec VPN Client L2TP clients.

**Design Point: PPTP or L2TP/IPSec?**
Consider the following when deciding between PPTP and L2TP/IPSec for remote access VPN connections:

- PPTP can be used with a variety of Microsoft clients including Windows Server 2003, Windows XP, Windows 2000, Windows NT version 4.0, Windows Millennium Edition, and Windows 98. PPTP does not require a certificate infrastructure to issue computer certificates.

- PPTP-based VPN connections provide data confidentiality (captured packets cannot be interpreted without the encryption key). PPTP VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data origin authentication (proof that the data was sent by the authorized user).

- PPTP-based VPN clients can be located behind a NAT if the NAT includes a NAT editor that knows how to properly translate PPTP tunneled data. For example, both the Internet connection sharing (ICS) feature of the Network Connections folder and the NAT/Basic Firewall routing protocol component of the Routing and Remote Access service include a NAT editor that translates PPTP traffic to and from PPTP clients located behind the NAT. VPN servers cannot be behind a NAT unless there are multiple public IP addresses and there is a one-to-one mapping of a public IP address to the private IP address of the VPN server or, if there is only one public address, if the NAT is configured to translate and forward the PPTP tunneled data to the VPN server. Most NATs using a single public IP address, including ICS and the NAT/Basic Firewall routing protocol component, can be configured to allow inbound traffic based on IP addresses and TCP and UDP ports. However, PPTP tunneled data does not use TCP or UDP headers. Therefore, a VPN server cannot be located behind a computer using ICS or the NAT routing protocol component when using a single IP address.

- L2TP/IPSec-based VPN clients or servers cannot be behind a NAT unless both the client and server support IPSec NAT Traversal (NAT-T). IPSec NAT-T is supported by Windows Server 2003, Windows XP Service Pack 2 (SP2), Windows XP Service Pack 1 (SP1) and Windows 2000 with L2TP/IPSec NAT-T Update for Windows XP and Windows 2000, and for previous versions of Windows with Microsoft L2TP/IPSec VPN Client. Microsoft recommends that servers, such as VPN servers running Windows Server 2003, not be placed behind NATs. For more information, see IPSec NAT-T is not recommended for Windows Server 2003 computers that are behind network address translators.

  Computers running Windows XP SP2 by default do use IPSec NAT-T to connect to servers that are located behind a NAT. This includes VPN server computers running Windows Server 2003. This default behavior can be modified with a registry setting. For more information, see The default behavior of IPSec NAT traversal (NAT-T) is changed in Windows XP Service Pack 2.

- L2TP/IPSec can be used with Windows Server 2003, Windows XP, Windows 2000, and Microsoft L2TP/IPSec VPN Client

clients and supports computer certificates as the recommended authentication method for IPSec. Computer certificate authentication requires a certificate infrastructure to issue computer certificates to the VPN server computer and all VPN client computers.

- By using IPSec, L2TP/IPSec -based VPN connections provide data confidentiality, data integrity, data origin authentication, and replay protection.

- PPTP and L2TP/IPSec is not an either/or choice. By default, a Windows Server 2003 VPN server supports both PPTP and L2TP/IPSec connections simultaneously. You can use PPTP for some remote access VPN connections (from VPN clients that are not running Windows XP or Windows 2000 and do not have an installed computer certificate) and L2TP/IPSec for other remote access VPN connections (from VPN clients running Windows XP, Windows 2000, or Microsoft L2TP/IPSec VPN Client and have an installed computer certificate).

- If you are using both PPTP and L2TP/IPSec, you can create separate remote access policies that define different connection parameters for PPTP and L2TP/IPSec connections.

**VPN Server**

**A VPN server is a computer running Windows Server 2003 and the Routing and Remote Access service. The VPN server does the following:**

- Listens for PPTP connection attempts and IPSec SA negotiations for L2TP connection attempts.

- Authenticates and authorizes VPN connections before allowing data to flow.

- Acts as a router forwarding data between VPN clients and resources on the intranet.

- Acts as an endpoint of the VPN tunnel from the tunnel client (typically the VPN client).

- Acts as the endpoint of the VPN connection from the VPN client.

The VPN server typically has two or more installed network adapters: one or more network adapters connected to the Internet and one or more network adapters connected to the intranet. The configuration of a VPN server with a single network adapter is discussed in Appendix B.

With Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

When you configure and enable the Routing and Remote Access service, the Routing and Remote Access Server Setup Wizard prompts you to select the role that the computer will fulfill. For VPN servers, you should select the **Remote access (dial-up or VPN)** configuration option.

With the **Remote access (dial-up or VPN)** option, the Routing and Remote Access server operates in the role of a dial-up or VPN server that supports remote access VPN connections. For remote access VPN connections, users run VPN client software and initiate a remote access connection to the server.

When you select the **Remote access (dial-up or VPN)** option in the Routing and Remote Access Server Setup Wizard:

1. You are first prompted to specify whether VPN, dial-up, or both types of access are needed.

2. Next, you are prompted to select the interface that is connected to the Internet. The interface that you select will be automatically configured with packet filters that allow only PPTP and L2TP-related traffic (unless you clear the **Enable security on the selected interface by setting up static packet filters** check box). All other traffic is silently discarded. For example, you will no longer be able to ping the Internet interface of the VPN server. If you want to use the VPN server computer as a network address translator (NAT), Web server, or other function, see Appendix B.

3. Next, if you have multiple network adapters that are connected to the intranet, you are prompted to select an interface

over which DHCP, DNS, and WINS configuration is obtained.

4. Next, you are prompted to determine whether you want to obtain IP addresses to assign to remote access clients by using either DHCP or a specified range of addresses. If you select a specified range of addresses, you are prompted to add one or more address ranges.

5. Next, you are prompted to specify whether you want to use RADIUS as your authentication provider. If you select RADIUS, you are prompted to configure primary and alternate RADIUS servers and the shared secret.

When you select the **Remote access (dial-up or VPN)** option in the Routing and Remote Access Server Setup Wizard, the results are as follows:

1. The Routing and Remote Access service is enabled as both a remote access server and a LAN and demand-dial router, with Windows as the authentication and accounting provider (unless RADIUS was chosen and configured). If there is only one network adapter connected to the intranet, that network adapter is automatically selected as the IP interface from which to obtain DHCP, DNS, and WINS configuration. Otherwise, the network adapter specified in the wizard is selected to obtain DHCP, DNS, and WINS configuration. If specified, the static IP address ranges are configured.

2. Exactly 128 PPTP and 128 L2TP ports are created. All of them are enabled for both inbound remote access connections and inbound and outbound demand-dial connections.

3. The selected Internet interface is configured with input and output IP packet filters that allow only PPTP and L2TP/IPSec traffic.

4. The DHCP Relay Agent component is added with the **Internal** interface. If the VPN server is a DHCP client at the time the wizard is run, the DHCP Relay Agent is automatically configured with the IP address of a DHCP server. Otherwise, you must manually configure the properties of the DHCP Relay Agent with an IP address of a DHCP server on your intranet. The DHCP Relay Agent forwards DHCPInform packets between VPN remote access clients and an intranet DHCP server.

5. The IGMP component is added. The **Internal** interface is configured for IGMP router mode. All other LAN interfaces are configured for IGMP proxy mode. This allows VPN remote access clients to send and receive IP multicast traffic.

#### Design Points: Configuring the VPN Server

**Consider the following before running the Routing and Remote Access Server Setup Wizard:**

- Which connection of the VPN server is connected to the Internet?

  Typical Internet-connected VPN servers have at least two LAN connections: one connected to the Internet (either directly or connected to a perimeter network) and one connected to the organization intranet. To make this distinction easier to see during the Routing and Remote Access Server Setup Wizard, rename the connections with their purpose or role using the Network Connections folder. For example, rename the connection connected to the Internet, default name **Local Area Connection 2**, to **Internet**.

- Can the VPN server be a DHCP client?

  The VPN server must have a manual TCP/IP configuration for its Internet interface. While technically possible, it is not recommended that the VPN server be a DHCP client for its intranet interface(s). Due to the routing requirements of the VPN server, manually configure an IP address, subnet mask, DNS server(s), and WINS server(s), but do not configure a default gateway.

  Note that it is possible for the VPN server to have a manual TCP/IP configuration and still use DHCP to obtain IP addresses for VPN clients.

- How will IP addresses be allocated to remote access VPN clients?

  The VPN server can be configured to obtain IP addresses from DHCP or from a manually configured set of address ranges. Using DHCP to obtain IP addresses simplifies the configuration, however, you must ensure that the DHCP scope

for the subnet to which the intranet connection of the VPN server is attached has enough addresses for all the computers physically connected to the subnet and the maximum number of PPTP and L2TP ports. For example, if the subnet to which the intranet connection of the VPN server is attached contains 50 DHCP clients, then, for the default configuration of the VPN server, the scope must contain at least 307 addresses (50 computers + 128 PPTP clients + 128 L2TP clients + 1 address for the VPN server). If there are not enough IP addresses in the scope, VPN clients that connect after all the addresses in the scope are allocated will be unable to access intranet resources.

If you are configuring a static pool of addresses, there might be additional routing considerations. For more information, see Intranet network infrastructure in this paper.

- What is the authentication and accounting provider?

  The VPN server can use Windows or RADIUS as its authentication or accounting provider.

  When Windows is used as the authentication and accounting provider, the VPN server uses Windows mechanisms to validate the credentials of the VPN client and access the VPN client's user account dial-in properties. Locally configured remote access policies authorize the VPN connection and locally written accounting log files log VPN connection accounting information.

  When RADIUS is used as the authentication and accounting provider, the VPN server uses a configured RADIUS server to validate the credentials of the VPN client, authorize the connection attempt, and store VPN connection accounting information.

- Will there be multiple VPN servers?

  If so, create multiple DNS A records to resolve the same name of the VPN server (for example, vpn.microsoft.com) to the different IP addresses of the separate VPN servers. DNS round robin will distribute the VPN connections across the VPN servers.

Consider the following when changing the default configuration of the VPN server for remote access VPN connections:

- Do you need additional PPTP or L2TP ports?

  By default, the Routing and Remote Access Server Setup Wizard configures 128 PPTP and 128 L2TP ports allowing 128 simultaneous PPTP connections and 128 simultaneous L2TP connections. If this is not sufficient for the maximum number of PPTP or L2TP connections, you can change the number of PPTP and L2TP ports by configuring the **WAN miniport (PPTP)** and **WAN miniport (L2TP)** devices from the properties of the **Ports** object in the Routing and Remote Access snap-in.

- Do you need to install a computer certificate?

  If the VPN server is configured with the Windows authentication provider and is supporting L2TP connections or is authenticating connections using the EAP-TLS authentication protocol, you must install a computer certificate on the VPN server that can be validated by the VPN client and a root certificate that is used to validate the VPN client.

- Do you need custom remote access policies for VPN connections?

  If you configure the VPN server for Windows authentication or for RADIUS authentication and the RADIUS server is a computer running the Internet Authentication Service (IAS), the default remote access policy rejects all types of connection attempts unless the remote access permission of the user accounts dial-in properties is set to **Allow access**. If you want to manage authorization and connection parameters by group or by type of connection, you must configure custom remote access policies. For more information, see Remote Access Policies in this paper.

- Do you want separate authentication and accounting providers?

  The Routing and Remote Access Server Setup Wizard configures both authentication and accounting providers to be the same. After the Wizard is complete, however, you can configure the authentication and accounting providers separately (for example, if you want to use Windows authentication and RADIUS accounting). You can configure authentication and

accounting providers on the **Authentication** tab from the properties of the VPN server in the Routing and Remote Access snap-in.

**Intranet Network Infrastructure**

The network infrastructure of the intranet is an important element of VPN design. Without proper design, VPN clients are unable to obtain proper IP addresses and resolve intranet names, and packets cannot be forwarded between VPN clients and intranet resources.

**Name Resolution**

If you use Domain Name System (DNS) to resolve intranet host names or Windows Internet Name Service (WINS) to resolve intranet NetBIOS names, ensure that the VPN server is configured with the IP addresses of the appropriate DNS and WINS servers. The VPN server can be configured with DNS and WINS server either manually or as a DHCP client. As part of the PPP negotiation process, the VPN clients receive the IP addresses of DNS and WINS server. By default, the VPN clients inherit the DNS and WINS server addresses configured on the VPN server.

After the PPP connection negotiation is complete, Windows XP and Windows 2000 VPN clients send a DHCPInform message to the VPN server. The response is relayed back to the VPN client and contains a DNS domain name, additional DNS server addresses for DNS servers that are checked before the DNS server configured through the PPP negotiation, and WINS server addresses that replace the WINS server addresses configured through the PPP negotiation. This communication is facilitated by the DHCP Relay Agent routing protocol component of the Routing and Remote Access service, which is automatically added by the Routing and Remote Access Server Setup Wizard.

If the VPN server is a DHCP client (the VPN server is using DHCP to configure its intranet interfaces), the VPN server relays the DHCPInform messages to the DHCP server that was in use when the Routing and Remote Access Server Wizard was run. If the VPN server has a manual TCP/IP configuration on its intranet interface (recommended), the DHCP Relay Agent routing protocol component must be configured with the IP address of at least one DHCP server on your intranet. You can add DHCP server IP addresses to the DHCP Relay Agent routing protocol component on the **General** tab from the properties of the **DHCP Relay Agent** object under **IP Routing** in the Routing and Remote Access snap-in.

**Design Points: Name Resolution by VPN Clients for Intranet Resources**

**Consider the following when configuring name resolution for remote access VPN clients:**

- Using the Ping and Net tools, test DNS and WINS name resolution for intranet resources from the VPN server computer. If name resolution does not work from the VPN server, it will not work for VPN clients. Troubleshoot and fix all name resolution problems of the VPN server before testing VPN connections.

- If the VPN server is a DHCP client (the VPN server is using DHCP to configure its intranet interfaces), no other configuration is necessary. The DNS and WINS servers assigned to the VPN server are also assigned to the VPN clients. The default configuration of the Routing and Remote Access Server Setup Wizard adds the DHCP Relay Agent routing protocol component and configures it with the IP address of the VPN server's DHCP server so that DHCPInform messages sent by VPN clients running Windows XP and Windows 2000 and its response are properly relayed between the VPN client and the DHCP server of the VPN server.

  However, configuring the VPN server as a DHCP client is not recommended due to issues with configuring the VPN server's default gateway. Therefore, it is recommended that you manually configure the TCP/IP configuration of the VPN server's intranet interfaces and manually configure the DHCP Relay Agent routing protocol component with the IP address of one or more of your DHCP servers.

- If the VPN server is manually configured with a TCP/IP configuration, verify the DNS and WINS server addresses. In this configuration, the Routing and Remote Access Server Setup Wizard cannot automatically configure the DHCP Relay Agent routing protocol component. You must manually add the IP address of at least one DHCP server on your intranet in order for DHCPInform messages to be replayed between VPN clients running Windows XP and Windows 2000 and the DHCP server. If you do not, DHCPInform messages sent by VPN clients running Windows XP and Windows 2000 are discarded and the VPN clients do not receive the updated DNS and WINS server addresses or the DNS domain name.

- If you have a single-subnet small office/home office (SOHO) with no DHCP, DNS, or WINS server, you must either configure a DNS server or WINS server in order to resolve names for both computers on the SOHO subnet and VPN clients or enable NetBIOS broadcast name resolution, which enables NetBIOS over TCP/IP name resolution between connected VPN clients and computers on the SOHO network. NetBIOS broadcast name resolution can be enabled from the **IP** tab in the properties of a VPN server in the Routing and Remote Access snap-in.

**Routing**
The VPN server is an IP router and as such must be properly configured with the set of routes that makes all locations reachable.

**Specifically, the VPN server needs the following:**

- A default route that points to a firewall or router directly connected to the Internet.

  This route makes all of the locations on the Internet reachable.

- One or more routes that summarize the addresses used on your intranet that point to a neighboring intranet router.

  These routes make all of the locations on your intranet reachable from the VPN server. Without these routes, all intranet hosts not connected to the same subnet as the VPN server are unreachable.

To add a default route that points to the Internet, configure the Internet interface with a default gateway and then manually configure the intranet interface without a default gateway.

**To add intranet routes to the routing table of the VPN server, you can:**

- Add static routes using the Routing and Remote Access snap-in. You do not necessarily have to add a route for each subnet in your intranet. At a minimum, you just need to add the routes that summarize all the possible addresses in your intranet. For example, if your intranet uses portions of the private address space 10.0.0.0/8 to number its subnets and hosts, you do not have to add a route for each subnet. Just add a route for 10.0.0.0 with the subnet mask 255.0.0.0 that points to a neighboring router on the intranet subnet to which your VPN server is attached.

- If you are using the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) routing protocols in your intranet, you can add and configure the RIP or OSPF routing protocol components of the Routing and Remote Access service so that the VPN server participates in the propagation of routing information as a dynamic router.

If your intranet has only a single subnet, no further configuration is required.

Ensuring the reachability of VPN clients from the intranet depends on how you configure the VPN server to obtain IP addresses for VPN clients.

**The IP addresses assigned to VPN clients as they connect can be from:**

- An on-subnet address range, which is an address range of the intranet subnet to which the VPN server is attached.

  An on-subnet address range is used whenever the VPN server is configured to use DHCP to obtain IP addresses for VPN clients and when the manually configured pool(s) of IP addresses are within the range of addresses of the attached subnet.

- An off-subnet address range, which is an address range that represents a different subnet that is logically attached to the VPN server.

  An off-subnet address range is used whenever the VPN server is manually configured with pool(s) of IP addresses for a separate subnet.

If you are using an on-subnet address range, no additional routing configuration is required as the VPN server acts as a proxy for all packets destined for VPN clients. Routers and hosts on the VPN server subnet forward packets destined to VPN clients to the VPN server and the VPN server relays them the appropriate VPN client.

If you are using an off-subnet address range, you must add the route(s) that summarize the off-subnet address range to the intranet routing infrastructure so that traffic destined to VPN clients are forwarded to the VPN server and then sent by the VPN server to the appropriate VPN client. To provide the best summarization of address ranges for routes, choose address ranges that can be expressed using a single prefix and subnet mask. For more information, see the topic titled "Expressing an IP address range with a mask" in Windows Server 2003 Help and Support.

**You can add the routes that summarize the off-subnet address range to the routing infrastructure of the intranet**

**through the following:**

- Add static routes to the neighboring router for the off-subnet address range that point to the VPN server's intranet interface. Configure the neighboring router to propagate this static route to other routers in the intranet using the dynamic routing protocol used in your intranet.

- If the VPN server is using OSPF and participating as a dynamic router, the VPN server must be configured as an autonomous system boundary router (ASBR) so that the static routes of the off-subnet address range are propagated to the other OSPF routers in the intranet.

If your intranet consists of a single subnet, then you must either configure each intranet host for persistent route(s) of the off-subnet address range that point to the VPN server's intranet interface or configure each intranet host with the VPN server as its default gateway. Therefore, it is recommended that you use an on-subnet address pool for a SOHO network consisting of a single subnet.

**VPN Client Routing and Simultaneous Intranet and Internet Access**
By default, when a Windows-based VPN client makes a VPN connection, it automatically adds a new default route for the VPN connection and modifies the existing default route to have a higher metric. Adding the new default route means that all Internet locations except the IP address of the tunnel server and locations based on other routes are not reachable for the duration of the VPN connection.

To prevent the default route from being created, obtain properties of the **Internet Protocol (TCP/IP)** component of the VPN connection. Click **Advanced**. In the **Advanced TCP/IP Settings** dialog box, click the **Genera**l tab, and then clear the **Use default gateway on remote network** check box. When the **Use default gateway on remote network** check box is cleared, a default route is not created, however, a route corresponding to the Internet address class of the assigned IP address is created. For example, if the address assigned during the connection process is 10.0.12.119, the Windows 2000 and Windows XP VPN client creates a route for the class-based network ID 10.0.0.0 with the subnet mask 255.0.0.0.

Based on the **Use default gateway on remote network** setting, one of the following occurs when the VPN connection is active:

- Internet locations are reachable and intranet locations are not reachable except those matching the address class of the assigned IP address (the **Use default gateway on remote network** check box is cleared).

- All intranet locations are reachable and Internet locations are not reachable except the address of the VPN server and location available through other routes (the **Use default gateway on remote network** check box is selected).

For most Internet-connected VPN clients, this behavior does not represent a problem because they are typically engaged in either intranet or Internet communication, not both.

**For VPN clients who want concurrent access to intranet and Internet resources when the VPN connection is active, you can do one of the following:**

- Select the **Use default gateway on remote network** check box (the default setting) and allow Internet access through the organization intranet. Internet traffic between the VPN client and Internet hosts would pass though firewalls or proxy servers as if the VPN client is physically connected to the organization intranet. While there is an impact on performance, this method allows Internet access to be filtered and monitored according to the organization's network policies while the VPN client is connected to the organization network.

- If the addressing within your intranet is based on a single class-based network ID, clear the **Use default gateway on remote network** check box. The best example is when your intranet is using the private IP address space 10.0.0.0/8.

- If the addressing within your intranet is not based on a single class-based network ID, you can use one of the following solutions:

    - The DHCPInform message sent by Windows XP clients includes the requesting of the DHCP Classless Static Routes DHCP option. If configured on a Windows Server 2003 DHCP server, the Classless Static Routes DHCP option contains a set of routes representing the address space of your intranet that are automatically added to the routing table of the requesting client.

- The Connection Manager Administration Kit for Windows Server 2003 allows you to configure specific routes as part of the connection manager profile distributed to VPN users. You can also specify a Uniform Resource Locator (URL) that contains the current set of organization intranet routes or additional routes beyond those configured in the profile.

- Clear the **Use default gateway on remote network** check box and use one of the following a command file (.CMD) on the VPN client with route commands to add static routes for the network IDs of your intranet using your assigned IP address as the gateway IP address after the connection is made.

You can determine your assigned IP address from the display of the Ipconfig command or by double-clicking the VPN connection in the Network Connections folder when the VPN connection is active. In the resulting **Status** dialog box, click the **Details** tab. The VPN client's assigned IP address is listed as **Client IP address**.

**Design Points: Routing Infrastructure**

**Consider the following when configuring the routing infrastructure for remote access VPN connections:**

- Configure the Internet interface of the VPN server with a default gateway. Do not configure the intranet interface of the VPN server with a default gateway.

- Add static IP routes to the VPN server that summarize the addresses used in your intranet. Alternately, if you use either RIP or OSPF for your dynamic routing protocol, configure and enable RIP or OSPF on the VPN server. If you use a routing protocol other than RIP or OSPF, such as Interior Gateway Routing Protocol (IGRP), you might configure the VPN server's neighboring intranet router for RIP or OSPF on the interface connected to subnet to which the VPN server is attached and IGRP on all other interfaces.

- Configure the VPN server with an on-subnet address range by obtaining IP addresses through DHCP or by manually configuring on-subnet address pools.

**Quarantine Resources**

Network Access Quarantine Control, a new feature in the Windows Server 2003 family, delays normal remote access to a private network until the configuration of the remote access computer has been examined and validated by an administrator-provided script. When a remote access computer initiates a connection to a remote access server, the user is authenticated and the remote access computer is assigned an IP address. However, the connection is placed in quarantine mode, with which network access is limited. The administrator-provided script is run on the remote access computer. When the script notifies the remote access server that it has successfully run and the remote access computer complies with current network policies, quarantine mode is removed and the remote access computer is granted normal remote access.

Quarantine resources consists of servers that a remote access client in quarantine mode can access to perform name resolution (DNS servers), obtain the latest version of the script (such as file servers with anonymous access allowed), or instructions and components needed to make the remote access client comply with network policies (such as Web servers with anonymous access allowed).

For more information about Network Access Quarantine Control, see the white paper titled "Windows Server 2003 Network Access Quarantine Control."

**AAA Infrastructure**

The authentication, authorization, and accounting (AAA) infrastructure exists to:

- Authenticate the credentials of VPN clients.

- Authorize the VPN connection.

- Record the VPN connection creation and termination for accounting purposes.

The AAA infrastructure consists of:

- The VPN server computer

- A RADIUS server computer

- A domain controller

As previously discussed, a Windows Server 2003 VPN server can be configured with either Windows or RADIUS as its authentication or accounting provider. RADIUS provides a centralized AAA service when you have multiple VPN servers or a mix of heterogeneous dial-up and VPN equipment.

When you configure Windows as the authentication provider, the VPN server performs the authentication of the VPN connection by communicating with a domain controller using a secure remote procedure call (RPC) channel and authorization of the connection attempt through the dial-in properties of the user account and locally configured remote access policies.

When you configure RADIUS as the authentication provider, the VPN server relies on a RADIUS server to perform both the authentication and authorization. When a VPN connection is attempted, the VPN client credentials and other connection parameters are used to create a RADIUS Access-Request message that is sent to the configured RADIUS server. If the connection attempt is both authenticated and authorized, the RADIUS server sends back a RADIUS Access-Accept message. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends back a RADIUS Access-Reject message.

When you configure Windows as the accounting provider, the VPN server logs VPN connection information in a local log file (*SystemRoot*\System32\Logfiles\Logfile.log by default) based on settings configured on the properties of the **Local File** object in the **Remote Access Logging** folder in the Routing and Remote Access snap-in.

When you configure RADIUS as the authentication provider, the VPN server sends RADIUS accounting messages for VPN connections on a RADIUS server, which records the accounting information.

If you are using RADIUS and a Windows domain as the user account database for which to verify user credentials and obtain dial-in properties, it is recommended to use the Internet Authentication Service (IAS). IAS is a full-featured RADIUS server (for Windows 2000 Server and Windows Server 2003) that is tightly integrated with Active Directory and the Routing and Remote Access service. IAS for Windows Server 2003 also supports a RADIUS proxy.

**When IAS is used as the RADIUS server:**

- IAS performs the authentication of the VPN connection by communicating with a domain controller using a secure RPC channel. IAS performs authorization of the connection attempt through the dial-in properties of the user account and remote access policies configured on the IAS server.

- IAS logs all RADIUS accounting information in a local log file (SystemRoot\System32\Logfiles\Logfile.log by default) based on settings configured on the properties of the **Local File** object in the **Remote Access Logging** folder in the Internet Authentication Service snap-in.

**Remote Access Policies**

Remote access policies are an ordered set of rules that define how connections are either accepted or rejected. For connections that are accepted, remote access policies can also define connection restrictions. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting. Connection attempts are evaluated against the remote access policies in order, trying to determine whether the connection attempt matches all of the conditions of each policy. If the connection attempt does not match all of the conditions of any policy, the connection attempt is rejected.

If a connection matches all the conditions of a remote access policy and is granted remote access permission, the remote access policy profile specifies a set of connection restrictions. The dial-in properties of the user account also provide a set of restrictions. Where applicable, user account connection restrictions override the remote access policy profile connection restrictions.

**Remote access policies consist of the following elements:**

- Conditions

- Permission

- Profile settings

**Conditions**

Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt. If there are multiple conditions, then all of the conditions must match the settings of the connection attempt in order for it to match the policy.

**For VPN connections, you commonly use the following conditions:**

- NAS-Port-Type

  By setting the NAS-Port-Type condition to Virtual (VPN), you can specify all VPN connections.

- Tunnel-Type

  By setting the Tunnel-Type to Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP), you can specify different policies for PPTP and L2TP connections.

- Windows-Groups

  By setting the Windows-Groups to the appropriate groups, you can grant or deny access based on group membership.

**Permission**
You can use the permission setting to either grant to deny remote access for the connection attempt if the remote access permission of the user account is set to **Control access through Remote Access Policy**. Otherwise, the remote access permission setting on the user account determines the remote access permission.

**Profile Settings**
A remote access policy profile is a set of properties that are applied to a connection when it is authorized.

For VPN connections, you can use the following profile settings:

- Dial-in constraints can be used to define how long the connection can exist or be idle before being terminated by the VPN server, among others.

- IP settings can define using IP packet filters the specific types of IP traffic that are allowed for remote access VPN connections. With profile packet filters, you can configure the IP traffic that is allowed from remote access clients (From client filters) or to remote access clients (To client filters) on an exception basis: either all traffic except traffic specified by filters or no traffic except traffic specified by filters. Remote access policy profile filtering applies to all remote access connections that match the remote access policy.

- Authentication settings can define which authentication protocols the VPN client must use in order to send its credentials and the configuration of EAP types, such as EAP-TLS.

- Encryption settings can define whether encryption is required and the encryption strength. For encryption strengths, Windows Server 2003 supports **Basic** (40-bit MPPE for PPTP and 56-bit Data Encryption Standard [DES] for L2TP), **Strong** (56-bit MPPE for PPTP and 56-bit DES for L2TP), or **Strongest** (128-bit MPPE for PPTP and 3DES for L2TP).

For example, you can create a Windows group named VPNUsers whose members are the user accounts of the users creating remote access VPN connections across the Internet. Then, you create a policy with two conditions on the policy: NAS-Port-Type is set to Virtual (VPN) and Windows-Group is set to VPNUsers. Finally, you configure the profile for the policy for a specific authentication method and encryption strength.

**Preventing Traffic Routed from VPN Clients**
Once a VPN client successfully establishes a PPTP or L2TP connection, by default any packet sent over the connection is received by the VPN server and forwarded.

**Packets sent over the connection can include:**

- Packets originated from the remote access client computer

- Packets sent to the remote access client computer by other computers

When the remote access client computer makes the VPN connection, by default it creates a default route so that all traffic that matches the default route is sent over the VPN connection. If other computers are forwarding traffic to the remote access VPN client, treating the remote access client computer as a router, then that traffic is also be forwarded across the VPN connection. This is a security problem because the VPN server has not authenticated the computer that is forwarding traffic to the remote access VPN client. The computer forwarding traffic to the remote access VPN client computer has the same network access as the authenticated remote access VPN client computer.

To prevent the VPN server from receiving traffic across the VPN connection for computers other than authenticated remote access VPN client computers, configure remote access policy packet filters on the remote access policy that is used for your VPN connections.

For the **Input Filters**, set the filter action to **Permit only the packets listed below** and configure a single filter with the settings listed in Table 2.

**Table 2 Input filter settings**

| IP Packet Filter Field | Setting |
|---|---|
| Source Address | User's address |
| Source Network Mask | User's mask |
| Destination Address | Any |
| Destination Network Mask | Any |
| Protocol | Any |

For the **Output Filters**, set the filter action to **Permit only the packets listed below** and configure a single filter with the settings listed in Table 3.

**Table 3 Output filter settings**

| IP Packet Filter Field | Setting |
|---|---|
| Source Address | Any |
| Source Network Mask | Any |
| Destination Address | User's address |
| Destination Network Mask | User's mask |
| Protocol | Any |

**Note**

Although the Routing and Remote Access snap-in displays **User's address** and **User's mask**, the actual filter that is created for each remote access client is for the client's assigned IP address and a subnet mask of 255.255.255.255.

**Note**

The default policy named **Connections to Microsoft Routing and Remote Access server** has the input packet filters previously described already configured. With this set of IP packet filters, the VPN server discards all traffic sent across the VPN connection except traffic that either originated from or is sent to authenticated remote access VPN clients.

#### Windows Domain User Accounts and Groups

Windows NT version 4.0 domains, mixed-mode Active Directory domains, and native-mode Active Directory domains contain the user accounts and groups used by the Routing and Remote Access service and IAS to authenticate and authorize VPN connection attempts.

User accounts contain the user name and a form of the user's password that can be used for validation of the VPN clients user credentials. Additional account properties determine whether the user account is enabled or disabled, locked out, or permitted to logon only during specific hours. If a user account is disabled, locked out, or not permitted to logon during the time of the

VPN connection, the VPN connection attempt is rejected.

**User accounts also contain dial-in settings. The dial-in setting most relevant for VPN connections is the remote access permission setting, which has the following values:**

- Allow access
- Deny access
- Control access through Remote Access Policy

The **Allow access** and **Deny access** settings explicitly allow or deny remote access and are equivalent to the remote access permission setting of Windows NT 4.0 domain accounts. When you use the **Control access through Remote Access Policy** setting, the remote access permission is determined by the remote access permission setting of the matching remote access policy. If the user account is in a mixed-mode domain, the **Control access through Remote Access Policy** setting is not available and you must manage remote access permission on a per-user basis. If the user account is in a native-mode domain, the **Control access through Remote Access Policy** setting is available and you can manage remote access permission on a per-user basis or using groups.

When using groups to manage access, you can use your existing groups and create remote access policies that either allow or reject access or restrict access based on the group name. For example, the Employees group has no VPN remote access restrictions, however, the Contractors group can only create VPN connections during business hours. Alternately, you can create groups based on the type of connection being made. For example, you can create a VPNUsers group and add as members all the user accounts allowed to create VPN connections.

Both the Routing and Remote Access service and IAS can use Active Directory universal principal names (UPNs) and universal groups. In a large domain with thousands of users, create a universal group for all of the users for whom you want to allow access, and then create a remote access policy that grants access for this universal group. Do not put all of your user accounts directly into the universal group, especially if you have a large number of them on your network. Instead, create separate global groups that are members of the universal group, and add users to those global groups.

**Design Points: AAA Infrastructure**

**Consider the following when configuring the AAA infrastructure for remote access VPN connections:**

- If you have multiple VPN servers and you want to centralize AAA service or a heterogeneous mixture of dial-up and VPN equipment, use a RADIUS server and configure the VPN server for the RADIUS authentication and accounting providers.
- If your user account database is a Windows domain, use IAS as your RADIUS server. If you use IAS, install IAS on a domain controller for best performance. Install at least two IAS servers for fail-over and fault tolerance of AAA services.
- Whether configured locally or on an IAS server, use remote access policies to authorize VPN connections and specify connection constraints. For example, use remote access policies to grant access based on group membership, to enforce the use of encryption and a specific encryption strength, to specify the use of EAP-TLS, or to limit traffic using IP packet filtering.
- To prevent VPN clients from forwarding routed traffic, configure remote access policy profile packet filters to discard all traffic on VPN connections except traffic to and from VPN clients.
- For a large Active Directory domain, nest global groups within universal groups to manage access based on group membership.
- Sensitive fields of RADIUS messages, such as the user password and encryption keys, are encrypted with the RADIUS shared secret configured on the VPN server and the RADIUS server. Make the shared secret a long (22 characters or longer), random sequence of letters, numbers, and punctuation. An example of a strong shared secret is 8d#>9fq4bV)H7%a3^jfDe2. To further protect RADIUS traffic, use Windows Server 2003 IPSec policies to provide data confidentiality for all traffic using the RADIUS UDP destination ports (1812 and 1645 for RADIUS authentication traffic and 1813 and 1646 for RADIUS accounting traffic).

**Certificate Infrastructure**

To perform certificate-based authentication for L2TP connections and smart card or user certificate-based authentication for

VPN connections using EAP-TLS, a certificate infrastructure, also known as a public key infrastructure (PKI), must be in place to issue the proper certificates to submit during the authentication process and to validate the certificate being submitted.

**Computer Certificates for L2TP/IPSec**

When you are using the certificate authentication method for L2TP connections, the list of certification authorities (CAs) is not configurable. Instead, each computer in the L2TP connection sends a list of root CAs to its IPSec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Computer A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPSec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPSec peer, Computer B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPSec security negotiation fails.

The VPN client must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN server trusts. Additionally, the VPN server must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN client trusts.

For example, if the VPN client was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies the VPN server during IPSec security negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the VPN server does not have a valid computer certificate issued from a CA that follows a certificate chain to either CertAuth1 or CertAuth2, IPSec security negotiation fails.

A single CA commonly issues computer certificates to all computers in an organization. Because of this, all computers within the organization both have computer certificates from a single CA and request certificates for authentication from the same single CA.

**Deploying computer certificates in your organization consists of the following:**

1. Deploy a certificate infrastructure. For more information, see Appendix D, "Deploying a Certificate Infrastructure".

2. Install a computer certificate on each computer. For more information, see "Deploying L2TP-based Remote Access" in this paper.

**Certificate Infrastructure for Smart Cards**

The use of smart cards for user authentication is the strongest form of user authentication in Windows Server 2003. For remote access VPN connections, you must use the Extensible Authentication Protocol (EAP) with the **Smart card or other certificate (TLS)** EAP type, also known as EAP-Transport Level Security (EAP-TLS).

**Deploying smart cards in your organization consists of the following:**

1. Create a certificate infrastructure using certification authorities.

2. For each domain, set security permissions and delegation for the Smart Card User, Smart Card Logon and Enrollment Agent certificate templates.

3. Configure the CA to issue smart card and Enrollment Agent certificates.

4. Configure an enrollment station, a computer that is used to physically install the smart card certificates on smart cards.

5. Use the enrollment station to create a smart card with a smart card user logon certificate that is installed on the smart card and assigned to a specific user account.

For more information about how to configure smart cards for user logon, see the topic titled Checklist: Deploying smart cards for logging on to Windows in Windows Server 2003 Help and Support.

The individual smart cards are distributed to users who have a computer with a smart card reader. To log in to the computer, the smart card must be inserted into the smart card reader and the smart card personal identification number (PIN) must be typed. When the user attempts a VPN connection, the smart card certificate is sent during the connection negotiation process.

**To configure EAP-TLS for smart cards on the VPN client:**

- The VPN connection must be configured to use EAP with the **Smart Card or other certificate** EAP type.

- In the properties of the **Smart Card or other certificate** EAP type, select **Use my smart card**.

- For Windows 2000 or Windows XP (prior to Service Pack 1) VPN clients, if you want to validate the computer certificate of the VPN or IAS server, select **Validate server certificate**. If you want to ensure that the servers DNS name ends in a specific string, select **Connect only if server name ends with** and type the string. To require the server's computer certificate to have been issued a certificate from a specific trusted root CA, select the CA in **Trusted root certification authority**.

- For Windows XP (Service Pack 1 and later) VPN clients, if you want to validate the computer certificate of the VPN or IAS server, select **Validate server certificate**. If you want to configure the names of the authenticating servers, select **Connect to these servers** and type the server names. To require the server's computer certificate to have been issued a certificate from a specific set of trusted root CAs, select them in **Trusted Root Certification Authorities**.

**To configure EAP-TLS authentication on the VPN server:**

EAP must be enabled as an authentication type on the **Authentication Methods** dialog box available from the **Security** tab in the properties of the VPN server in the Routing and Remote Access snap-in.

**To configure EAP-TLS authentication on the remote access policy:**

- On the remote access policy that is being used for VPN connections, the **Smart Card or other certificate** EAP type must be added to the selected EAP providers from the **Authentication** tab on the policy's profile settings. If the computer on which the remote access policy is being configured has multiple computer certificates installed, configure the properties of the **Smart Card or other certificate** EAP type and select the appropriate computer certificate to submit during the EAP-TLS authentication process.

#### Certificate Infrastructure for User Certificates

The use of registry-based user certificates for user authentication can be used in place of smart cards. However, it is not as strong a form of authentication. With smart cards, the user certificate issued during the authentication process is only made available when the user physically possesses the smart card and has knowledge of the PIN to logon to their computer. With user certificates, the user certificate issued during the authentication process is made available when the user logs on to their computer using a domain-based user name and password.

Just as with smart cards, authentication using user certificates for remote access VPN connections use EAP-TLS as the authentication protocol.

Deploying user certificates in your organization consists of the following:

1. Deploy a certificate infrastructure. For more information, see Appendix D, "Deploying a Certificate Infrastructure".

2. Install a user certificate for each user. For more information, see "Deploying PPTP-based Remote Access" and "Deploying L2TP-based Remote Access" in this paper.

When the user attempts a VPN connection, the user certificate is sent during the connection negotiation process.

**To configure EAP-TLS for user certificates on the VPN client:**

- The VPN connection must be configured to use EAP with the **Smart Card or other certificate** EAP type.

- For Windows 2000 or Windows XP (prior to Service Pack 1) VPN clients, if you want to validate the computer certificate of the VPN or IAS server, select **Validate server certificate**. If you want to ensure that the server's DNS name ends in a specific string, select **Connect only if server name ends with** and type the string. To require the server's computer certificate to have been issued a certificate from a specific trusted root CA, select the CA in **Trusted root certification authority**.

- For Windows XP (Service Pack 1 and later) VPN clients, if you want to validate the computer certificate of the VPN or IAS server, select **Validate server certificate**. If you want to configure the names of the authenticating servers, select **Connect to these servers** and type the server names. To require the server's computer certificate to have been issued a

certificate from a specific set of trusted root CAs, select them in **Trusted Root Certification Authorities**.

**To configure EAP-TLS authentication on the VPN server:**

- EAP must be enabled as an authentication type on the **Authentication Methods** dialog box available from the **Security** tab in the properties of the VPN server in the Routing and Remote Access snap-in.

**To configure EAP-TLS authentication on the remote access policy:**

- On the remote access policy that is being used for VPN connections, the **Smart Card or other certificate** EAP type must be added to the selected EAP providers from the **Authentication** tab on the policy's profile settings. If the computer on which the remote access policy is being configured has multiple computer certificates installed, configure the properties of the **Smart Card or other certificate** EAP type and select the appropriate computer certificate to submit during the EAP-TLS authentication process.

**Design Points: Certificate Infrastructure**

Consider the following when configuring the certificate infrastructure for remote access VPN connections:

- In order to create L2TP/IPSec remote access VPN connections using computer certificate authentication for IPSec, you must install computer certificates, also known as machine certificates, on each VPN client and VPN server. If you are using a Windows Server 2003 enterprise CA as an issuing CA, configure your Active Directory domain for auto-enrollment of computer certificates using Computer Configuration Group Policy. Each computer that is a member of the domain automatically requests a computer certificate when the Computer Configuration Group Policy is updated.

  The computer certificate of the VPN client must be valid and verifiable by the VPN server the VPN server must have a root CA certificate for the CA that issued the computer certificate of the VPN client.

  The computer certificate of the VPN server must be valid and verifiable by the VPN client the VPN client must have a root CA certificate for the CA that issued the computer certificate of the VPN server.

- In order to authenticate VPN connections using a smart card or user certificate with EAP-TLS, the VPN client must have a smart card or registry-based user certificate installed and the authenticating server must have a computer certificate installed. The authenticating server is either the VPN server (if configured for Windows authentication) or the IAS server (if the VPN server is configured for RADIUS authentication and the RADIUS server is a computer running Windows Server 2003 and IAS).

  The smart card or user certificate of the VPN client must be valid and verifiable by the authenticating server the authenticating server trust the root CA for the CA that issued the certificate of the VPN client.

  The computer certificate of the authenticating server must be verifiable by the VPN client the VPN client trust the root CA for the CA that issued the computer certificate of the authenticating server.

- To install a computer or user certificate on a computer across the Internet, make a PPTP connection using a password-based authentication protocol such as MS-CHAP v2. After connecting, use the Certificate Manager snap-in or Internet Explorer to request the appropriate certificates. Once the certificates are installed, disconnect and then reconnect with the appropriate VPN protocol and authentication method. An example of this situation is a laptop computer that is issued to an employee without the certificates needed to make L2TP/IPSec or EAP-TLS-authenticated connections.

# Deploying PPTP-based Remote Access

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

Deploying PPTP-based remote access VPN connections using Windows Server 2003 consists of the following:

- Deploy certificate infrastructure
- Deploy Internet infrastructure
- Deploy AAA infrastructure
- Deploy VPN servers
- Deploy intranet infrastructure
- Deploy VPN clients

**Deploying Certificate Infrastructure**

For PPTP-based VPN connections, a certificate infrastructure is needed only when you are using either smart cards or registry-based user certificates and EAP-TLS authentication. If you are only using a password-based authentication protocol such as MS-CHAP v2, a certificate infrastructure is not required and is not used for the creation of the VPN connection.

If you need a certificate infrastructure for PPTP-based VPN connections, you must install a computer certificate on the authenticating server (the VPN server or the RADIUS server) and either a certificate on each smart card distributed to VPN client users or a user certificate on each VPN client computer.

For information about deploying a certificate infrastructure, see Appendix E, "Deploying a Certificate Infrastructure."

**Installing Computer Certificates**

To install a computer certificate, an issuing CA must be present to issue certificates. Once the issuing CA is configured, you can install a computer certificate in the following ways:

1. By configuring the automatic allocation of computer certificates to computers in an Active Directory domain.

   This method allows a single point of configuration for the entire domain. All members of the domain automatically request the computer certificate through a Group Policy setting. To immediately obtain a computer certificate for the VPN or IAS server that is a member of the domain for which auto-enrollment is configured, restart the computer or type **secedit /refreshpolicy machine_policy** for Windows 2000 or **gpupdate /target:computer** for Windows XP and Windows Server 2003 from a command prompt. To configure an Active Directory domain for automatic enrollment of computer certificates, see the topic titled "Checklist: Configuring certificate autoenrollment" in Windows Server 2003 Help and Support.

2. By using the Certificates snap-in to request a computer certificate.

   If you are using a Windows Server 2003 enterprise CA as an issuing CA, each computer can separately request a computer certificate from the issuing CA using the Certificates snap-in. For more information, see the topics titled "Manage an individual certificate" and "Request a certificate" in Windows Server 2003 Help and Support.

3. By using the Certificates snap-in to import a computer certificate.

   If you have a certificate file that contains the computer certificate, you can import the computer certificate using the Certificates snap-in. For more information about importing a certificate using the Certificates snap-in, see the topic titled Import a certificate in Windows Server 2003 Help and Support.

4. By executing a CAPICOM script that requests a computer certificate.

In this method, each computer that needs a computer certificate must execute a CAPICOM script that requests a computer certificate from the issuing CA. CAPICOM is a COM client, supporting Automation, that performs cryptographic functions (the CryptoAPI) using Microsoft ActiveX and COM objects. CAPICOM can be used via Visual Basic, Visual Basic Scripting Edition, and C++. For more information about CAPICOM, see CAPICOM at http://msdn.microsoft.com/library/en-us/seccrypto/security/using_capicom.asp?frame=true.

**Deploying Smart Cards**

For information about deploying smart cards in Windows Server 2003, see the topic titled "Checklist: Deploying smart cards for logging on to Windows" in Windows Server 2003 Help and Support.

**Installing User Certificates**

To install a user certificate, an issuing CA must be present to issue certificates. Once the issuing CA is configured, you can install a user certificate in the following ways:

1. By configuring the automatic allocation of user certificates in a Windows 2003 Active Directory domain.

   This method allows a single point of configuration for the entire domain. All users that correspond to members of the domain automatically request the user certificate through a Group Policy setting. To immediately obtain a user certificate for a computer running Windows Server 2003 or Windows XP that is a member of the domain for which auto-enrollment is configured, restart the computer or type **gpupdate /target:user** from a command prompt. To configure an Active Directory domain for automatic enrollment of computer certificates, see the topic titled "Checklist: Configuring certificate autoenrollment" in Windows Server 2003 Help and Support. Autoenrollment of user certificates requires Windows Server 2003, Enterprise Edition. Computers running Windows 2000 do not support autoenrollment of user certificates.

2. By using a Web browser to request a user certificate.

   The issuing CA must support Web enrollment of certificates. For example, if you are using a Windows Server 2003 enterprise CA as an issuing CA and the CA computer is also running Internet Information Services (IIS), you can use Web enrollment to request a user certificate. For more information about requesting a user certificate, see the topic titled Submit a user certificate request via the Web in Windows Server 2003 Help and Support.

3. By using the Certificates snap-in to request a user certificate.

   If you are using a Windows Server 2003 enterprise CA as an issuing CA, you can request a user certificate from the Certificates snap-in. For more information about requesting a user certificate using the Certificates snap-in, see the topic titled Request a certificate in Windows Server 2003 Help and Support.

4. By importing a user certificate using the Certificates snap-in.

   If you have a certificate file that contains a user certificate, import the user certificate from the Certificates snap-in. For more information about requesting a user certificate using the Certificates snap-in, see the topic titled Import a certificate in Windows Server 2003 Help and Support.

5. By executing a CAPICOM script that requests a user certificate.

   In this method, each user must execute a CAPICOM script that requests a user certificate from the issuing CA.

**Deploying Internet Infrastructure**

Deploying the Internet infrastructure for remote access VPN connections consists of the following:

- Place VPN servers in perimeter network or on the Internet.
- Install Windows Server 2003 on VPN servers and configure Internet interfaces.
- Add address records to Internet DNS.

**Placing VPN Servers in Perimeter Network or on the Internet**

Decide where to place the VPN servers in relation to your Internet firewall. In the most common configuration, the VPN servers are placed behind the firewall on the perimeter network between your intranet and the Internet. If so, configure packet filters on the firewall to allow PPTP traffic to and from the IP address of the VPN servers' perimeter network interfaces. For more information, see Appendix A.

**Installing Windows Server 2003 on VPN Servers and Configuring Internet Interfaces**

Install Windows Server 2003 on the VPN server computer and connect it to either the Internet or to perimeter network with one network adapter and connect it to the intranet with another network adapter. Without running the Routing and Remote Access Server Setup Wizard, the VPN server computer will not forward IP packets between the Internet and the intranet. For the connection connected to the Internet or the perimeter network, configure the TCP/IP protocol with a public IP address, a subnet mask, and the default gateway of either the firewall (if the VPN server is connected to a perimeter network) or an ISP router (if the VPN server is directly connected to the Internet.) Do not configure the connection with DNS server or WINS server IP addresses.

**Adding Address Records to Internet DNS**

To ensure that the name of the VPN server (for example, vpn.microsoft.com) can be resolved to its proper IP address, either add DNS address (A) records to your DNS server (if you are providing DNS name resolution for Internet users) or have your ISP add DNS address (A) records to their DNS server(s) (if your ISP is providing DNS name resolution for Internet users). Verify that the name of the VPN server can be resolved to its public Internet IP address when connected to the Internet.

**Deploying AAA Infrastructure**

Deploying the AAA infrastructure for remote access VPN connections consists of the following:

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.

**Configuring Active Directory for User Accounts and Groups**

**To configure Active Directory for user accounts and groups, do the following:**

1. Ensure that all users that are making remote access connections have a corresponding user account. This includes employees, contractors, vendors, and business partners.
2. Set the remote access permission on user accounts to **Allow access** or **Deny access** to manage remote access by user. Or, to manage remote access by group, set the remote access permission on user accounts to **Control access through Remote Access Policy**.
3. Organize remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For more information, see the topic titled "Nesting groups" in Windows Server 2003 Help and Support.

**Configuring the Primary IAS Server on a Domain Controller**

**To configure the primary IAS server on a domain controller, do the following:**

1. On the domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows Server 2003 Help and Support.
2. If the IAS server authenticates connection attempts for user accounts in other domains, verify that these domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support. For more information about trust relationships, see the topic titled "Understanding Domains and Forests" in Windows Server 2003 Help and Support.

   If the IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.
3. Enable file logging for accounting and authentication events. For more information, see the topic titled "Configure log file

properties" in Windows Server 2003 Help and Support.

4. Add the VPN server(s) as RADIUS clients of the IAS server. For more information, see the topic titled "Add RADIUS clients" in Windows Server 2003 Help and Support. For the IP address of each VPN server, use the intranet IP address assigned to the VPN server. If you are using names, use the internal name of the VPN server (this is not necessarily the same DNS name used by Internet clients). Use strong shared secrets.

5. Create remote access policies that reflect your remote access usage scenarios.

   For example, to configure a remote access policy that requires PPTP-based VPN connections for members of the Employees group to use EAP-TLS authentication and 128-bit encryption, use the New Remote Access Policy Wizard to create a typical remote access policy with the following settings:

   Policy name: VPN connections (example)

   Access method: VPN

   Group access: Employees group (example)

   Authentication methods: Enable only **Extensible Authentication Protocol (EAP)** and select the **Smart Card or other certificate** EAP type

   Policy encryption level: Select only **Strongest encryption**

If you are using Network Access Quarantine Control, you can use the MS-Quarantine-IPFilter or MS-Quarantine-Session-Timeout RADIUS vendor-specific attributes (VSAs) to specify quarantine settings. Both of these VSAs are configured from the **Advanced** tab in the profile properties of the remote access policy that you create for remote access connections.

You can use the MS-Quarantine-IPFilter attribute to configure input and output packet filters to allow only the following:

- The traffic generated by the remote access client notifier component. If you are using Rqc.exe and its default port, then configure a single input packet filter to allow only traffic from TCP port 7250 and to TCP port 7250.

- The traffic needed to access the quarantine resources. This includes filters that allow the remote access client to access name resolution servers (such as Domain Name System [DNS]), file shares, or Web sites.

You can use the MS-Quarantine-Session-Timeout attribute to specify how long the remote access server must wait to receive the notification that the script has run successfully before terminating the connection.

Because the quarantine VSAs can limit network access and automatically disconnect remote access users, you should configure these attributes only after a quarantine Connection Manager package has been distributed and installed on the remote access client computers of your organization.

For more information about Network Access Quarantine Control, see the white paper titled "Windows Server 2003 Network Access Quarantine Control."

| **Note** |
| --- |
| If IAS is not installed on a domain controller, you must configure the primary IAS server computer to read the properties of user accounts in the domain. For more information, see the "Enable the IAS server to read user accounts in Active Directory" in Windows Server 2003 Help and Support. |

#### Configuring the Secondary IAS server on a Different Domain Controller

**To configure the secondary IAS server on a different domain controller, do the following:**

1. On the other domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows Server 2003 Help and Support.

2. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure

the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support. For more information about trust relationships, see the topic titled "Understanding Domains and Forests" in Windows Server 2003 Help and Support.

If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.

3. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see the topic titled "Copy the IAS configuration to another server" in Windows Server 2003 Help and Support.

If IAS is not installed on a domain controller, you must configure the secondary IAS server computer to read the properties of user accounts in the domain. For more information, see the "Enable the IAS server to read user accounts in Active Directory" in Windows Server 2003 Help and Support.

#### Deploying VPN Servers

**Deploying the VPN servers for remote access VPN connections consists of the following:**

- Configure the VPN server's connection to the intranet.
- Run the Routing and Remote Access Server Setup Wizard.

#### Configuring the VPN Server's Connection to the intranet

For each VPN server, configure the connection connected to the intranet with a manual TCP/IP configuration consisting of IP address, subnet mask, intranet DNS servers, and intranet WINS servers. Note that you must not configure the default gateway on the intranet connection to prevent default route conflicts with the default route pointing to the Internet.

#### Running the Routing and Remote Access Server Setup Wizard

**Run the Routing and Remote Access Server Setup Wizard to configure each Windows Server 2003 VPN server using the following steps:**

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Routing and Remote Access**.
2. Right-click your server name, and then click **Configure and Enable Routing and Remote Access**. Click **Next**.
3. In **Configuration**, click **Remote access (dial-up or VPN)** and then click **Next**. If you additionally want to use the VPN server computer as a network address translator (NAT), Web server, or other function, see Appendix B.
4. In **Remote Access**, select **VPN**. If you also want the VPN server to support dial-up remote access connections, click **Dial-up**. Click **Next**.
5. In **VPN Connection**, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click **Next**.
6. In **IP Address Assignment**, click **Automatically** if the VPN server should use DHCP to obtain IP addresses for remote access VPN clients. Or, click **From a specified range of addresses** to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure in order for the VPN clients to be reachable. When IP address assignment is complete, click **Next**.
7. In **Managing Multiple Remote Access Servers**, if you are using RADIUS for authentication and authorization, click **Yes, set up this server to work with a RADIUS server**, and then click **Next**.
    - In **RADIUS Server Selection**, configure the primary (mandatory) and alternate (optional) RADIUS servers and the shared secret, and then click **Next**.
8. Click **Finish**.

9. If prompted, start the Routing and Remote Access service.

By default, only 128 PPTP ports are configured on the WAN Miniport (PPTP) device. If you need more PPTP ports, configure the **WAN Miniport (PPTP)** device from the properties of the **Ports** object in the Routing and Remote Access snap-in. By default, 128 L2TP ports are also configured.

By default, the MS-CHAP, and MS-CHAP v2, and EAP protocols are enabled.

If you are using Network Access Quarantine Control, install the quarantine listener component on the VPN server. If you are using Rqs.exe, modify the Rqs_setup.bat file to include the correct version string for the version of the network policy compliance script that is being run on the remote access clients. Next, run the Rqs_setup.bat file to install the Remote Access Quarantine Agent service.

For more information about Network Access Quarantine Control, see the white paper titled "Windows Server 2003 Network Access Quarantine Control."

#### Intranet Network Infrastructure

**Deploying the intranet network infrastructure for remote access VPN connections consists of the following:**

- Configure routing on the VPN server.
- Verify name resolution and intranet reachability from the VPN server.
- Configure routing for off-subnet address pools.
- Configure quarantine resources

#### Configuring Routing on the VPN Server

In order for your VPN servers to properly forward traffic to locations on the intranet, you must configure them with either static routes that summarize all the possible addresses used on the intranet or with routing protocols so that the VPN server can participate as a dynamic router and automatically add routes for intranet subnets to its routing table.

To add static routes, see the topic titled "Add a static route" in Windows Server 2003 Help and Support. To configure the VPN server as a RIP router, see the topic titled "Configure RIP for IP". To configure the VPN server as an OSPF router, see the topics titled "OSPF design considerations" and "Configure OSPF".

#### Verifying Name Resolution and Reachability from the VPN Server

From each VPN server, verify that the VPN server can resolve names and successfully communicate with intranet resources by using the Ping command, Internet Explorer, and making drive and printer connections to known intranet servers.

#### Configuring Routing for Off-subnet Address Pools

If you configured any of the VPN servers with manual address pools and any of the pools are an off-subnet pool, you must ensure that the route(s) representing the off-subnet address pool(s) are present in your intranet routing infrastructure. You can ensure this by either adding static route(s) representing the off-subnet address pool(s) as static routes to the neighboring router(s) of the VPN server(s) and then using the routing protocol of your intranet to propagate the route to other routers. When you add the static route(s), you must specify that the gateway or next hop address is the intranet interface of the VPN server.

Alternately, if you are using RIP or OSPF, you can configure the VPN servers using off-subnet address pools as RIP or OSPF routers. For OSPF, you must configure the VPN server as an autonomous system boundary router (ASBR). For more information, see the topic titled "OSPF design considerations" in Windows Server 2003 Help and Support.

#### Quarantine Resources

If you are using Network Access Quarantine Control, designate an internal DNS server, file servers and shares for updated scripts, and Web servers with Web page(s) containing network policy compliance instructions and components.

For more information about Network Access Quarantine Control, see the white paper titled "Windows Server 2003 Network Access Quarantine Control."

#### Deploying VPN Clients

**Deploying VPN clients for remote access VPN connections consists of the following:**

- Manually configure VPN clients
- Configure CM packages with CMAK

**Manually Configuring VPN clients**
If you have a small number of VPN clients, you can manually configure VPN connections for each VPN client. For Windows 2000 VPN clients, use the Make New Connection Wizard to create the Internet and VPN connections and link them together so that when you connect using the VPN connection, the Internet connection is already made. For Windows XP VPN clients, use the New Connection Wizard to create the Internet and VPN connections.

**Configuring CM Packages with CMAK**
For a large number of VPN clients running different versions of Windows, you should use the CMAK to create and distribute customized Connection Manager packages for your users. For more information, see the topic titled "Before you start: Understanding Connection Manager and the Administration Kit" in Windows Server 2003 Help and Support.

If you are using Network Access Quarantine Control, create the Connection Manager package to contain the following:

The CM profile contains the following:

- A post-connect action setting that runs a network policy requirements script.
- A network policy requirements script.

  This script performs validation checks on the remote access client computer to verify that it conforms to network policies. It can be a custom executable file or as simple as a command file (also known as a batch file). When the script has run successfully and the connecting computer has satisfied all of the network policy requirements (as verified by the script), the script runs a notifier component (an executable) with the appropriate parameters and, optionally, copies the latest version of the script from a quarantine resource.

  If the script does not run successfully, it should direct the remote access user to a quarantine resource such as an internal Web page, which describes how to install the components that are required for network policy compliance.
- A notifier component

  The notifier component sends a message that indicates a successful execution of the script to the quarantine-compatible remote access server. You can use your own notifier component or you can use Rqc.exe. If you use Rqc.exe, run it from the script with the correct parameters, including the script version.

For more information, see the white paper titled "Windows Server 2003 Network Access Quarantine Control."

# Deploying L2TP/IPSec-based Remote Access

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

Deploying L2TP-based remote access VPN connections using Windows Server 2003 consists of the following:

- Deploy certificate infrastructure
- Deploy Internet infrastructure
- Deploy AAA infrastructure
- Deploy VPN servers
- Deploy intranet infrastructure
- Deploy VPN clients

**Deploying Certificate Infrastructure**

For L2TP-based VPN connections, a certificate infrastructure is required to issue computer certificates needed to negotiate authentication for IPSec. Additionally, a certificate infrastructure is also needed when you are using either smart cards or user certificates and EAP-TLS for user authentication. You must install a computer certificate on all VPN clients and VPN servers. If you are using EAP-TLS for user authentication, you must install a user certificate on all VPN clients and, if the authenticating server is a RADIUS server, a computer certificate on the RADIUS server.

For information about deploying a certificate infrastructure, see Appendix D, "Deploying a Certificate Infrastructure."

| **Note** |
| --- |
| The Windows Server 2003 Routing and Remote Access service supports the configuration of a preshared key for IPSec authentication of L2TP/IPSec connections. You can type the preshared key from the **Security** tab in the properties of a VPN server in the Routing and Remote Access snap-in. Windows XP and Windows Server 2003-based VPN clients also support the configuration of an IPSec preshared key (click **IPSec Settings** on the **Security** tab in the properties of a VPN connection). However, preshared key authentication for L2TP/IPSec connections is not secure and is not recommended. |

**Deploying Computer Certificates**

**To install a computer certificate, a CA must be present to issue certificates. Once the CA is configured, you can install a computer certificate in the following ways:**

1. By configuring the automatic allocation of computer certificates to computers in an Active Directory domain.

   This method allows a single point of configuration for the entire domain. All members of the domain automatically request the computer certificate through a Group Policy setting. To immediately obtain a computer certificate for a computer running Windows 2000 that is a member of the domain for which auto-enrollment is configured, restart the computer or type **secedit /refreshpolicy machine_policy** from a command prompt. To immediately obtain a computer certificate for a computer running Windows Server 2003 or Windows XP that is a member of the domain for which auto-enrollment is configured, restart the computer or type **gpupdate /target:computer** from a command prompt. To configure an Active Directory domain for automatic enrollment of computer certificates, see the topic titled "Checklist: Configuring certificate autoenrollment" in Windows Server 2003 Help and Support.

2. By using the Certificates snap-in to request a computer certificate.

   If you are using a Windows Server 2003 enterprise CA as an issuing CA, each computer can separately request a computer certificate from the issuing CA using the Certificates snap-in. For more information, see the topics titled "Manage an individual certificate" and "Request a certificate" in Windows Server 2003 Help and Support.

3. By using the Certificates snap-in to import a computer certificate.

If you have a certificate file that contains the computer certificate, you can import the computer certificate using the Certificates snap-in. For more information about importing a certificate using the Certificates snap-in, see the topic titled Import a certificate in Windows Server 2003 Help and Support.

4. By executing a CAPICOM script that requests a computer certificate.

   In this method, each computer must execute a CAPICOM script that requests a computer certificate from the issuing CA. CAPICOM is a COM client, supporting Automation, that performs cryptographic functions (the CryptoAPI) using Microsoft ActiveX and COM objects. CAPICOM can be used via Visual Basic, Visual Basic Scripting Edition, and C++. For more information about CAPICOM, see CAPICOM at http://msdn.microsoft.com/library/en-us/seccrypto/security/using_capicom.asp?frame=true.

**Deploying Smart Cards**

For information about deploying smart cards in Windows Server 2003, see the topic titled "Checklist: Deploying smart cards for logging on to Windows" in Windows Server 2003 Help and Support.

**Deploying User Certificates**

**To install a user certificate, an issuing CA must be present to issue certificates. Once the issuing CA is configured, you can install a user certificate in the following ways:**

1. By configuring the automatic allocation of user certificates to computers in an Active Directory domain.

   This method allows a single point of configuration for the entire domain. All users that correspond to members of the domain automatically request the user certificate through a Group Policy setting. To immediately obtain a user certificate for a computer running Windows Server 2003 or Windows XP that is a member of the domain for which auto-enrollment is configured, restart the computer or type **gpupdate /target:user** from a command prompt. To configure an Active Directory domain for automatic enrollment of computer certificates, see the topic titled "Checklist: Configuring certificate autoenrollment" in Windows Server 2003 Help and Support. Autoenrollment of user certificates requires Windows Server 2003, Enterprise Edition. Computers running Windows 2000 do not support autoenrollment of user certificates.

2. By using a Web browser to request a user certificate.

   For more information about requesting a user certificate, see the topic titled Submit a user certificate request via the Web in Windows Server 2003 Help and Support. The issuing CA must support Web enrollment of certificates. For example, if you are using a Windows Server 2003 enterprise CA as an issuing CA and the CA computer is also running Internet Information Services (IIS), you can use Web enrollment to request a user certificate.

3. By using the Certificates snap-in to request a user certificate.

   If you are using a Windows Server 2003 enterprise CA as an issuing CA, you can request a user certificate using the Certificates snap-in. For more information about requesting a user certificate using the Certificates snap-in, see the topic titled Request a certificate in Windows Server 2003 Help and Support.

4. By importing a user certificate using the Certificates snap-in.

   If you have a certificate file that contains a user certificate, import the user certificate using the Certificates snap-in. For more information about requesting a user certificate using the Certificates snap-in, see the topic titled Import a certificate in Windows Server 2003 Help and Support.

5. By executing a CAPICOM script that requests a user certificate.

   In this method, each user must execute a CAPICOM script that requests a user certificate from the issuing CA.

**Deploying Internet Infrastructure**

**Deploying the Internet infrastructure for remote access VPN connections consists of the following:**

- Place VPN servers in perimeter network or on the Internet.
- Install Windows Server 2003 on VPN servers and configure Internet interfaces.
- Add address records to Internet DNS.

**Placing VPN Servers in Perimeter Network or on the Internet**

Decide where to place the VPN servers in relation to your Internet firewall. In the most common configuration, the VPN servers are placed behind the firewall on the perimeter network between your intranet and the Internet. If so, configure packet filters on the firewall to allow L2TP/IPSec traffic to and from the IP address of the VPN servers' perimeter network interfaces. For more information, see Appendix A.

**Installing Windows Server 2003 on VPN Servers and Configuring Internet Interfaces**

Install Windows Server 2003 on the VPN server computer and connect it to either the Internet or to perimeter network with one network adapter and connect it to the intranet with another network adapter. Without running the Routing and Remote Access Server Setup Wizard, the VPN server computer will not forward IP packets between the Internet and the intranet. For the connection connected to the Internet or the perimeter network, configure the TCP/IP protocol with a public IP address, a subnet mask, and the default gateway of either the firewall (if the VPN server is connected to a perimeter network) or an ISP router (if the VPN server is directly connected to the Internet.) Do not configure the connection with DNS server or WINS server IP addresses.

**Adding Address Records to Internet DNS**

To ensure that the name of the VPN server (for example, vpn.microsoft.com) can be resolved to its proper IP address, either add DNS address (A) records to your DNS server (if you are providing DNS name resolution for Internet users) or have your ISP add DNS address (A) records to their DNS server(s) (if your ISP is providing DNS name resolution for Internet users). Verify that the name of the VPN server can be resolved to its public Internet IP address when connected to the Internet.

**Deploying AAA Infrastructure**

**Deploying the AAA infrastructure for remote access VPN connections consists of the following:**

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.

**Configuring Active Directory for User Accounts and Groups**

**To configure Active Directory for user accounts and groups, do the following:**

1. Ensure that all users that are making remote access connections have a corresponding user account. This includes employees, contractors, vendors, and business partners.
2. Set the remote access permission on user accounts to **Allow access** or **Deny access** to manage remote access by user. Or, to manage remote access by group, set the remote access permission on user accounts to **Control access through Remote Access Policy**.
3. Organize remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For more information, see the topic titled "Nested groups" in Windows Server 2003 Help and Support.

**Configuring the Primary IAS Server on a Domain Controller**

**To configure the primary IAS server on a domain controller, do the following:**

1. On the domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows Server 2003 Help and Support.
2. If the IAS server authenticates connection attempts for user accounts in other domains, verify that these domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server

to read user objects in Active Directory" in Windows Server 2003 Help and Support. For more information about trust relationships, see the topic titled "Understanding Domains and Forests" in Windows Server 2003 Help and Support.

If the IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.

3. Enable file logging for accounting and authentication events. For more information, see the topic titled "Configure log file properties" in Windows Server 2003 Help and Support.

4. Add the VPN server(s) as RADIUS clients of the IAS server. For more information, see the topic titled "Add RADIUS clients" in Windows Server 2003 Help and Support. For the IP address of each VPN server, use the intranet IP address assigned to the VPN server. If you are using names, use the internal name of the VPN server (this is not necessarily the same DNS name used by Internet clients). Use strong shared secrets.

5. Create remote access policies that reflect your remote access usage scenarios.

   For example, to configure a remote access policy that requires L2TP-based VPN connections for members of the Employees group to use MS-CHAP v2 authentication and Triple-DES IPSec encryption, create a typical remote access policy with the following settings:

   Policy name: VPN connections (example)

   Access method: VPN

   Group access: Employees group (example)

   Authentication methods: Enable only **Microsoft Encrypted Authentication version 2 (MD-CHAP v2)**

   Policy encryption level: Select only **Strongest encryption**

If you are using Network Access Quarantine Control, you can use the MS-Quarantine-IPFilter or MS-Quarantine-Session-Timeout RADIUS vendor-specific attributes (VSAs) to specify quarantine settings. Both of these VSAs are configured from the **Advanced** tab in the profile properties of the remote access policy that you create for remote access connections.

You can use the MS-Quarantine-IPFilter attribute to configure input and output packet filters to allow only the following:

- The traffic generated by the remote access client notifier component. If you are using Rqc.exe (from Windows Server 2003 Resource Kit Tools) and its default port, then configure a single input packet filter to allow only traffic from TCP port 7250 and to TCP port 7250.

- The traffic needed to access the quarantine resources. This includes filters that allow the remote access client to access name resolution servers (such as Domain Name System [DNS]), file shares, or Web sites.

You can use the MS-Quarantine-Session-Timeout attribute to specify how long the remote access server must wait to receive the notification that the script has run successfully before terminating the connection.

Because the quarantine VSAs can limit network access and automatically disconnect remote access users, you should configure these attributes only after a quarantine Connection Manager package has been distributed and installed on the remote access client computers of your organization.

For more information about Network Access Quarantine Control, see the white paper titled "Windows Server 2003 Network Access Quarantine Control."

| **Note** |
| --- |
| If IAS is not installed on a domain controller, you must configure the primary IAS server computer to read the properties of user accounts in the domain. For more information, see the "Enable the IAS server to read user accounts in Active Directory" in Windows Server 2003 Help and Support. |

#### Configuring the Secondary IAS Server on a Different Domain Controller

**To configure the secondary IAS server on a different domain controller, do the following:**

1. On the other domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows Server 2003 Help and Support.

2. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support. For more information about trust relationships, see the topic titled "Understanding Domains and Forests" in Windows Server 2003 Help and Support.

   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.

3. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see the topic titled "Copy the IAS configuration to another server" in Windows Server 2003 Help and Support.

If IAS is not installed on a domain controller, you must configure the secondary IAS server computer to read the properties of user accounts in the domain. For more information, see the "Enable the IAS server to read user accounts in Active Directory" in Windows Server 2003 Help and Support.

#### Deploying VPN Servers

**Deploying the VPN servers for remote access VPN connections consists of the following:**

- Configure the VPN server's connection to the intranet.

- Run the Routing and Remote Access Server Setup Wizard.

Windows Server 2003 includes enhanced support for the clustering of L2TP/IPSec VPN servers. For more information, see the topic titled "Checklist: Enabling and configuring Network Load Balancing" in Windows Server 2003 Help and Support.

#### Configuring the VPN Server's Connection to the Intranet

For each VPN server, configure the connection connected to the intranet with a manual TCP/IP configuration consisting of IP address, subnet mask, intranet DNS servers, and intranet WINS servers. Note that you must not configure the default gateway on the intranet connection to prevent default route conflicts with the default route pointing to the Internet.

#### Running the Routing and Remote Access Server Setup Wizard

**Run the Routing and Remote Access Server Setup Wizard to configure each Windows Server 2003 VPN server using the following steps:**

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Routing and Remote Access**.

2. Right-click your server name, and then click **Configure and Enable Routing and Remote Access**. Click **Next**.

3. In **Configuration**, click **Remote access (dial-up or VPN)** and then click **Next**. If you want to use the VPN server computer as a network address translator (NAT), Web server, or other function, see Appendix B.

4. In **Remote Access**, select **VPN**. If you also want the VPN server to support dial-up remote access connections, click **Dial-up**. Click **Next**.

5. In **VPN Connection**, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click **Next**.

6. In **IP Address Assignment**, click **Automatically** if the VPN server should use DHCP to obtain IP addresses for remote access VPN clients. Or, click **From a specified range of addresses** to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure in order

for the VPN clients to be reachable. When IP address assignment is complete, click **Next**.

7. In **Managing Multiple Remote Access Servers**, if you are using RADIUS for authentication and authorization, click **Yes, set up this server to work with a RADIUS server**, and then click **Next**.

    - In **RADIUS Server Selection**, configure the primary (mandatory) and alternate (optional) RADIUS servers and the shared secret, and then click **Next**.

8. Click **Finish**.

9. If prompted, start the Routing and Remote Access service.

By default, only 128 L2TP ports are configured on the WAN Miniport (L2TP) device. If you need more L2TP ports, configure the **WAN Miniport (L2TP)** device from the properties of the **Ports** object in the Routing and Remote Access snap-in. By default, 128 PPTP ports are also configured. If you want to disable the VPN server's ability to accept PPTP connections, set the number of ports on the **WAN Miniport (PPTP)** device to 1 and clear the **Remote access connections (inbound only)** and **Demand-dial connections (inbound and outbound)** check boxes.

By default, the MS-CHAP, and MS-CHAP v2, and EAP protocols are enabled.

If you are using Network Access Quarantine Control, install the quarantine listener component on the VPN server. If you are using Rqs.exe from Windows Server 2003 Resource Kit Tools, modify the Rqs_setup.bat file to include the correct version strings for the version of the network policy compliance script that is being run on the remote access clients. Next, run the Rqs_setup.bat file to install the Remote Access Quarantine Agent service.

For more information about Network Access Quarantine Control, see the white paper titled "Windows Server 2003 Network Access Quarantine Control."

#### Intranet Network Infrastructure

**Deploying the intranet network infrastructure for remote access VPN connections consists of the following:**

- Configure routing on the VPN server.
- Verify name resolution and intranet reachability from the VPN server.
- Configure routing for off-subnet address pools.

#### Configuring Routing on the VPN Server

In order for your VPN servers to properly forward traffic to locations on the intranet, you must configure them with either static routes that summarize all the possible addresses used on the intranet or with routing protocols so that the VPN server can participate as a dynamic router and automatically add routes for intranet subnets to its routing table.

To add static routes, see the topic titled "Add a static route" in Windows Server 2003 Help and Support. To configure the VPN server as a RIP router, see the topic titled "Configure RIP for IP". To configure the VPN server as an OSPF router, see the topics titled "OSPF design considerations" and "Configure OSPF".

#### Verifying Name Resolution and Reachability from the VPN Server

From each VPN server, verify that the VPN server can resolve names and successfully communicate with intranet resources by using the Ping command, Internet Explorer, and making drive and printer connections to known intranet servers.

#### Configuring Routing for Off-subnet Address Pools

If you configured any of the VPN servers with manual address pools and any of the pools are an off-subnet pool, you must ensure that the route(s) representing the off-subnet address pool(s) are present in your intranet routing infrastructure. You can ensure this by either adding static route(s) representing the off-subnet address pool(s) as static routes to the neighboring router(s) of the VPN server(s) and then using the routing protocol of your intranet to propagate the route to other routers. When you add the static route(s), you must specify that the gateway or next hop address is the intranet interface of the VPN server.

Alternately, if you are using RIP or OSPF, you can configure the VPN servers using off-subnet address pools as RIP or OSPF routers. For OSPF, you must configure the VPN server as an autonomous system boundary router (ASBR). For more

information, see the topic titled "OSPF design considerations" in Windows Server 2003 Help and Support.

**Quarantine Resources**
If you are using Network Access Quarantine Control, designate an internal DNS server, file servers and shares for updated scripts, and Web servers with Web page(s) containing network policy compliance instructions and components.

For more information about Network Access Quarantine Control, see the white paper titled "Windows Server 2003 Network Access Quarantine Control."

**Deploying VPN Clients**
Deploying VPN clients for remote access VPN connections consists of the following:

- Manually configure VPN clients
- Configure CM packages with CMAK

**Manually Configuring VPN clients**
If you have a small number of VPN clients, you can manually configure VPN connections for each VPN client. For Windows 2000 VPN clients, use the Make New Connection Wizard to create the Internet and VPN connections and link them together so that when you connect using the VPN connection, the Internet connection is already made. For Windows XP VPN clients, use the New Connection Wizard to create the Internet and VPN connections.

**Configuring CM Packages with CMAK**
For a large number of VPN clients running different versions of Windows, you should use the CMAK to create and distribute customized Connection Manager packages for your users. For more information, see the topic titled "Before you start: Understanding Connection Manager and the Administration Kit" in Windows Server 2003 Help and Support.

If you are using Network Access Quarantine Control, create the Connection Manager package to contain the following:

The CM profile contains the following:

- A post-connect action setting that runs a network policy requirements script.
- A network policy requirements script.

  This script performs validation checks on the remote access client computer to verify that it conforms to network policies. It can be a custom executable file or as simple as a command file (also known as a batch file). When the script has run successfully and the connecting computer has satisfied all of the network policy requirements (as verified by the script), the script runs a notifier component (an executable) with the appropriate parameters and, optionally, copies the latest version of the script from a quarantine resource.

  If the script does not run successfully, it should direct the remote access user to a quarantine resource such as an internal Web page, which describes how to install the components that are required for network policy compliance.
- A notifier component

  The notifier component sends a message that indicates a successful execution of the script to the quarantine-compatible remote access server. You can use your own notifier component or you can use Rqc.exe, which is provided with Windows Server 2003 Resource Kit Tools. If you use Rqc.exe, run it from the script with the correct parameters, including the script version.

For more information, see the white paper titled "Windows Server 2003 Network Access Quarantine Control."

# Appendix A: Configuring Firewalls for a VPN Server Running Windows Server 2003

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

**The following are common configurations of firewalls with a VPN server:**

- The VPN server is attached to the Internet and the firewall is between the VPN server and the intranet.
- The firewall is attached to the Internet and the VPN server is between the firewall and the intranet.
- Two firewalls are used: one between the VPN server and the intranet and one between the VPN server and the Intranet.

**VPN Server in Front of the Firewall**

To secure the VPN server from sending or receiving any traffic on its Internet interface except VPN traffic, you need to configure PPTP or L2TP/IPSec input and output filters on the interface that corresponds to the connection to the Internet. Because IP routing is enabled on the Internet interface, if PPTP or L2TP/IPSec filters are not configured on the Internet interface, then any traffic received on the Internet interface is routed, which may forward unwanted Internet traffic to your intranet.

When the VPN server is in front of the firewall attached to the Internet, you need to add packet filters to the Internet interface that allow only VPN traffic to and from the IP address of the VPN server's Internet interface.

For inbound traffic, when the VPN server decrypts the tunneled data, it is forwarded to the firewall. The firewall in this configuration is acting as a filter for intranet traffic and can prevent specific resources from being accessed, scan data for viruses, perform intrusion detection, and other functions.

Because the only Internet traffic allowed on the intranet must pass through the VPN server, this approach also prevents the sharing of File Transfer Protocol (FTP) or Web intranet resources with non-VPN Internet users.

Figure 3 shows the VPN server in front of the firewall.




**Figure 3: The VPN server in front of the firewall**

The firewall is configured for the appropriate rules for intranet traffic to and from VPN clients according to your network security policies.

For the Internet interface on the VPN server, configure the following input and output filters using the Routing and Remote Access snap-in. These filters are automatically configured when you run the Routing and Remote Access Server Setup Wizard and choose the **Remote access (dial-up or VPN)** option, select the correct interface, and select the **Enable security on the selected interface by setting up packet filters** option on the VPN Connection page (enabled by default).

**Packet Filters for PPTP**

Configure the following input filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP destination port of 1723.
- This filter allows PPTP tunnel management traffic to the VPN server.
- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.

- This filter allows PPTP tunneled data to the VPN server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP [established] source port of 1723.

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site (also known as router-to-router) VPN connection. TCP [established] traffic is accepted only when the VPN server initiated the TCP connection.

Configure the following output filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP source port of 1723.

- This filter allows PPTP tunnel management traffic from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.

- This filter allows PPTP tunneled data from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP [established] destination port of 1723.

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. TCP [established] traffic is sent only when the VPN server initiated the TCP connection.

**Packet Filters for L2TP/IPSec**

Configure the following input filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 500.

  This filter allows Internet Key Exchange (IKE) traffic to the VPN server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 4500.

  This filter allows IPSec NAT-T traffic to the VPN server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 1701.

  This filter allows L2TP traffic to the VPN server.

Configure the following output filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 500.

  This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 5500.

  This filter allows IPSec NAT-T traffic from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 1701.

This filter allows L2TP traffic from the VPN server.

There are no filters required for IPSec Encapsulating Security Protocol (ESP) traffic for the IP protocol of 50. The Routing and Remote Access service filters are applied after the IPSec components remove the ESP header.

**VPN Server Behind the Firewall**

In a more common configuration, the firewall is connected to the Internet and the VPN server is an intranet resource that is connected to the perimeter network, also known as a screened subnet. The perimeter network is an IP network segment that contains resources that are available to Internet users, such as Web and FTP servers. The VPN server has an interface on both the perimeter network and the intranet. In this approach, the firewall must be configured with input and output filters on its Internet interface that allows the passing of tunnel maintenance traffic and tunneled data to the VPN server. Additional filters can allow the passing of traffic to Web, FTP, and other types of servers on the perimeter network. For an added layer of security, the VPN server can also be configured with PPTP or L2TP/IPSec packet filters on its perimeter network interface.

The firewall in this configuration is acting as a filter for Internet traffic and can confine the incoming and outgoing traffic to the specific resources on the perimeter network, perform intrusion attempt detection, prevent denial of service attacks, and other functions.

Because the firewall does not have the encryption keys for each VPN connection, it can only filter on the plaintext headers of the tunneled data. In other words, all tunneled data passes through the firewall. This is not a security concern, however, because the VPN connection requires an authentication process that prevents unauthorized access beyond the VPN server.

Figure 4 shows the VPN server behind the firewall on the perimeter network.




**Figure 4: The VPN server behind the firewall on the perimeter network**

For both the Internet and network perimeter interfaces on the firewall, configure the following input and output filters using the firewall's configuration software.

**Packet Filters for PPTP**

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

**Filters on the Internet Interface**

**Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:**

- Destination IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter allows PPTP tunnel management traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should only be used in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" and configured on the VPN server's network perimeter interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet that use this port.

**Configure the following output filters on the Internet interface of the firewall to allow the following types of traffic:**

- Source IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter allows PPTP tunnel management traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data from the VPN server.

- Source IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should only be used in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" and configured on the VPN server's network perimeter interface. By allowing all traffic from the VPN server to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

#### Filters on the Perimeter Network Interface

**Configure the following input filters on the perimeter network interface of the firewall to allow the following types of traffic:**

- Source IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter allows PPTP tunnel management traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data from the VPN server.

- Source IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should only be used in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" and configured on the VPN server's network perimeter interface. By allowing all traffic from the VPN server to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

**Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:**

- Destination IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter allows PPTP tunnel management traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should only be used in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" and configured on the VPN server's network perimeter interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

#### Packet Filters for L2TP/IPSec

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

#### Filters on the Internet Interface

**Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:**

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4).

  This filter allows IKE traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194).

  This filter allows IPSec NAT-T traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic to the VPN server.

**Configure the following output packet filters on the Internet interface of the firewall to allow the following types of traffic:**

- Source IP address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4).

  This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and UDP source port of 4500 (0x1194).

  This filter allows IPSec NAT-T traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic from the VPN server.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

#### Filters on the Perimeter Network Interface

**Configure the following input packet filters on the perimeter network interface of the firewall to allow the following types of traffic:**

- Source IP address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4).

  This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and UDP source port of 4500 (0x1194).

  This filter allows IPSec NAT-T traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic from the VPN server.

**Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:**

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4).

  This filter allows IKE traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194).

  This filter allows IPSec NAT-T traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

This filter allows IPSec ESP traffic to the VPN server.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

**VPN Server Between Two Firewalls**

Another configuration is when the VPN server computer in placed on the perimeter network between two firewalls. The Internet firewall, the firewall between the Internet and the VPN server, filters all Internet traffic from all Internet clients. The intranet firewall, the firewall between the VPN server and the intranet, filters intranet traffic from VPN clients.

Figure 5 shows the VPN server between two firewalls on the perimeter network.




**Figure 5: The VPN server between two firewalls on the perimeter network**

In this configuration:

- Configure your Internet firewall and VPN server with the packet filters as described in the "VPN Server Behind the Firewall" section.

- Configure your intranet firewall for the appropriate rules for intranet traffic to and from VPN clients according to your network security policies.

# Appendix B: Alternate Configurations

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

This section provides information about common alternate configurations for a Windows Server 2003 VPN server. The most common configuration is described in the "Deploying PPTP-based Remote Access" and "Deploying L2TP-based Remote Access" sections of this paper and whose principal characteristics are the following:

- The VPN server has multiple network adapters: at least one connected to the intranet and at least one connected to the Internet.
- The VPN server has static public IP addresses assigned to its Internet interfaces.
- The VPN server is only acting as a security gateway providing remote access to the intranet. The VPN server is not hosting any other Internet services such as NAT or Web services.

The two other most common configurations are the following:

1. The VPN server computer is performing other functions such as network address translation or Web hosting.
2. The VPN server computer has a single network adapter and its public IP address is published by a firewall.

The following sections detail the changes to make in the deployment of a VPN server to accommodate these additional common configurations.

**Multiple Internet Function VPN Server**
In this configuration, the VPN server's principal characteristics are the following:

- The VPN server has multiple network adapters at least one connected to the intranet and at least one connected to the Internet.
- The VPN server has static public IP addresses assigned to its Internet interfaces.
- The VPN server is acting as a security gateway providing remote access to the intranet and is hosting any other Internet services such as NAT or Web hosting.

In this configuration, you can follow the procedures as described in the "Deploying PPTP-based Remote Access" and "Deploying L2TP-based Remote Access" sections of this paper except that when you run the Routing and Remote Access Server Setup Wizard, clear the **Enable security on the selected interface by setting up static packet filters** check box on the **VPN Connection** page of the Wizard.

When you clear the **Enable security on the selected interface by setting up static packet filters** check box, PPTP and L2TP packet filters are not configured on the Internet interface of the VPN server computer. Whether you have to manually configure these filters depends on whether the VPN server computer is also hosting NAT.

- If NAT is needed on the VPN server computer, do not configure PPTP and L2TP packet filters or packet filters for other types of traffic. If you configure PPTP and L2TP packet filters on the Internet interface, NAT cannot function. Even though you do not configure any packet filters on the Internet interface of the VPN server computer, the function of the NAT discards any traffic from the Internet that does not correspond to traffic requested by intranet clients.
- If NAT is not needed on the VPN server computer, you can configure PPTP and L2TP packet filters and other types of filters for additional services hosted by the VPN server computer. For example, if the VPN server computer is also hosting a Web site, then filters should be added to allow traffic to and from the public IP address of the VPN server computer and TCP port 80.

**Single-Adapter VPN Server**

In this configuration, the VPN server computer has only a single network adapter and VPN clients are accessing services hosted on the VPN server computer. If the VPN server computer has only a single network adapter and is configured with a public IP address, all traffic to and from the services running on the VPN server computer are sent as clear text outside the VPN tunnel.

The only way a single adapter VPN server can work properly is if it is behind a firewall that is providing a publishing and translation service for the VPN server. The firewall publishes or makes known on the Internet a static public IP address for the VPN server. When VPN packets are sent to this published IP address, the firewall translates the address of the packet to a private or other public address by which the VPN server is known on the intranet.

Figure 6 shows an example of the published and actual addresses of a VPN server in this configuration.




**Figure 6: The single-adapter VPN server configuration**

A VPN client uses the Internet DNS to resolve the VPN server's name to its published public IP address. After the VPN connection is made, the intranet DNS and WINS infrastructures resolve the VPN server's name to its actual intranet address. One limitation to this configuration is that only PPTP is supported. Because the firewall is translating addresses, IPSec-protected L2TP traffic cannot traverse the firewall.

The VPN server is configured according to "Deploying PPTP-based Remote Access" in this paper with its intranet interface acting as an Internet interface. The firewall is configured to:

- Publish the name and public IP address of the VPN server on the Internet.
- Translate PPTP traffic sent to the public IP address of the VPN server to the intranet interface of the VPN server computer.
- Discard all traffic except PPTP traffic going to and from the VPN server computer.

# Appendix C: Setting up a VPN Test Lab

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

This section provides detailed information about how you can use five computers to create a test lab with which to configure and test the virtual private network (VPN) features in Windows Server 2003. These instructions are designed to take you through a set of tasks that expose you to VPN connections and their functionality. Beyond the set of tasks, you can use these instructions to create a functioning VPN configuration. You can then use this configuration to experiment with VPN features and functionality in prior to deployment on your production network.

This section covers:

- Setting up the infrastructure
- Virtual private network test lab tasks

**Setting up the Infrastructure**

**The infrastructure for the VPN test lab network consists of five computers performing the following services:**

- A computer running Windows Server 2003 that is acting as a domain controller, a Domain Name System (DNS) server, and a certification authority (CA). This computer is named DC1.
- A computer running Windows Server 2003 that is acting as a Remote Authentication Dial-in User Service (RADIUS) server. This computer is named IAS1.
- A computer running Windows Server 2003 that is acting as a Web server and file sharing server. This computer is named IIS1.
- A computer running Windows Server 2003 that is acting as a VPN server. This computer is named VPN1. VPN1 has two network adapters installed.
- A computer running Windows XP that is acting as a VPN client. This computer is named CLIENT1.

Figure 7 shows the configuration of the VPN test lab.




**Figure 7: Configuration of the VPN test lab**

There is a network segment representing a corporate intranet and a network segment representing the Internet. All computers on the corporate intranet are connected to a common hub or Layer 2 switch. All computers on the Internet are connected to a separate common hub or Layer 2 switch. Private addresses are used throughout the test lab configuration. The private network of 172.16.0.0/24 is used for the intranet. The private network of 10.0.0.0/24 is used for the simulated Internet.

Each computer is manually configured with the appropriate IP address, subnet mask, and DNS server IP address. There are no Dynamic Host Configuration Protocol (DHCP) or Windows Internet Name Service (WINS) servers present.

The following sections describe the configuration for each of the computers in the test lab. To reconstruct this test lab, configure the computers in the order presented.

| Note |
| --- |
| The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network. |

**DC1**

**DC1 is a computer running Windows Server 2003 that is providing the following services:**

- A domain controller for the testlab.microsoft.com domain.
- A DNS server for the testlab.microsoft.com DNS domain.
- The enterprise root certification authority (CA) for the testlab.microsoft.com domain.

To configure DC1 for these services, perform the following steps.

1. Install Windows Server 2003 as a stand-alone server.
2. Configure the TCP/IP protocol with the IP address of 172.16.0.1 and the subnet mask of 255.255.255.0.
3. Run Dcpromo.exe for a new domain called testlab.microsoft.com in a new forest. Install the DNS service when prompted.
4. Install the Certificate Services component as an enterprise root certification authority.
5. Configure the testlab.microsoft.com domain for automatic enrollment of computer certificates.

**IAS1**

IAS1 is a computer running Windows Server 2003 that is providing RADIUS authentication, authorization, and accounting for VPN1 (the VPN server computer).

**To configure IAS1 as a RADIUS server, perform the following steps:**

1. On DC1, add a computer account for the IAS1 computer.
2. Install Windows Server 2003 as a stand-alone server.
3. Configure the TCP/IP protocol with the IP address of 172.16.0.2, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.
4. Join IAS1 to the testlab.microsoft.com domain.
5. Install the Internet Authentication Service.

**IIS1**

IIS1 is a computer running Windows Server 2003 and the Internet Information Services. It is providing Web server services for intranet clients.

**To configure IIS1 as a Web server, perform the following steps:**

1. On DC1, add a computer account for the IIS1 computer.
2. Install Windows Server 2003 as a stand-alone server.
3. Configure the TCP/IP protocol with the IP address of 172.16.0.3, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.

4. Join IIS1 to the testlab.microsoft.com domain.

5. Install the Internet Information Services.

6. To determine whether the Web server is working correctly, run Internet Explorer on IAS1. When prompted by the Internet Connection wizard, configure the wizard for a LAN connection. In Internet Explorer, in **Address**, type **http://IIS1.testlab.microsoft.com/winxp.gif**. You should see a Windows XP graphic.

7. In Windows Explorer, share the root directory of Local Disk (C:)using the share name ROOT to the group **Everyone** with full access.

**VPN1**

VPN1 is a computer running Windows Server 2003 that is providing VPN server services for Internet-based VPN clients.

**To configure VPN1 as a VPN server, perform the following steps:**

1. On DC1, add a computer account for VPN1.

2. Install Windows Server 2003 as a stand-alone server.

3. For the intranet local area connection, configure the TCP/IP protocol with the IP address of 172.16.0.4, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.

4. For the Internet local area connection, configure the TCP/IP protocol with the IP address of 10.0.0.2 and the subnet mask of 255.255.255.0.

5. Join VPN1 to the testlab.microsoft.com domain.

6. Configure and enable the Routing and Remote Access service. In the Routing and Remote Access Server Setup Wizard, select **Remote access (dial-up or VPN)** from the configuration list and enable VPN from the **Remote Access** dialog box. When prompted for IP address assignment, select **From a specified range of addresses** and configure the range 172.16.0.248 to 172.16.0.255. Do not configure RADIUS authentication.

**CLIENT1**

CLIENT1 is a computer running Windows XP that is acting as a VPN client and gaining remote access to intranet resources across the simulated Internet.

**To configure CLIENT1 as a VPN client, perform the following steps:**

1. On DC1, add a computer account for CLIENT1.

2. Connect CLIENT1 to the intranet network segment.

3. On CLIENT1, install Windows XP as a workgroup computer.

4. Configure the TCP/IP protocol with the IP address of 172.16.0.5, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.

5. Join CLIENT1 to the testlab.microsoft.com domain.

6. Configure the TCP/IP protocol with the IP address of 10.0.0.1, the subnet mask of 255.255.255.0, and no DNS server IP address.

7. Shut down the CLIENT1 computer.

8. Disconnect the CLIENT1 computer from the intranet network segment and connect it to the simulated Internet network

segment.

9. Restart the CLIENT computer and log on using the cached credentials of the testlab.microsoft.com administrator account.

**VPN Test Lab Tasks**

The following tasks are designed to take you through the most common elements of remote access VPN support with Windows Server 2003:

- PPTP-based remote access
- L2TP-based remote access
- RADIUS authentication and accounting
- Remote access policies for different types of VPN connections

**PPTP-based Remote Access**

To create a PPTP-based remote access VPN connection between CLIENT1 and VPN1 and test whether intranet resources are available, perform the following steps:

**Create a User Account**

On DC1, use the Active Directory Users and Computers snap-in to create a user account named PPTPUser with a password. Set the remote access permission on the **Dial-in** tab to **Allow access**.

**Create the PPTP Connection**

1. On CLIENT1, use the New Connection Wizard to create a new VPN connection named PPTPtoCorpnet, using the VPN server IP address of 10.0.0.2.
2. Right-click the new **PPTPtoCorpnet** connection, and then click **Properties**.
3. Click the **Networking** tab, and then in **Type of VPN server I am calling**, click **PPTP VPN**.
4. Click **OK** to save changes to the **PPTPtoCorpnet** connection.

**Make the PPTP Connection**

1. On CLIENT1, double-click the **PPTPtoCorpnet** connection.
2. In the **PPTPtoCorpnet** dialog box, type **PPTPUser@testlab.microsoft.com** as the user name, type the password, and then select the **Save this user name and password to use when** check box.
3. Click **Connect**.

**Access Web server and File Share on the Intranet**

1. On CLIENT1, run Internet Explorer.
2. When prompted by the Internet Connection Wizard, configure it for a LAN connection.
3. In Internet Explorer, in **Address**, type **http://IIS1.testlab.microsoft.com/winxp.gif**. You should see a Windows XP graphic.
4. On CLIENT1, click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**. You should see the contents of the Local Drive (C:) on IIS1.

**Disconnect the PPTP Connection**

On CLIENT1, right-click the PPTPtoCorpnet connection and then click **Disconnect**.

**L2TP-based Remote Access**
To create an L2TP-based remote access VPN connection between CLIENT1 and VPN1 and test whether intranet resources are available, perform the following:

**Create a User Account**
On DC1, use Active Directory Users and Computers to create a user account named L2TPUser with a password. Set the remote access permission on the **Dial-in** tab to **Allow access**.

**Create the L2TP Connection**

1. On CLIENT1, use the New Connection Wizard to create a new VPN connection named L2TPtoCorpnet, using the VPN server IP address of 10.0.0.2.

2. Right-click the new **L2TPtoCorpnet** connection, and then click **Properties**.

3. Click the **Networking** tab, and then in **Type of VPN server I am calling**, select **L2TP IPSec VPN**.

4. Click **OK** to save changes to the **L2TPtoCorpnet** connection.

**Make the L2TP Connection**

1. On CLIENT1, double-click the **L2TPtoCorpnet** connection.

2. In the L2TPtoCorpnet dialog box, type **L2TPUser@testlab.microsoft.com** as the user name, type the password, and then select the **Save this user name and password to use when** check box.

3. Click **Connect**.

**Access Web Server and File Share on the Intranet**

1. On CLIENT1, run Internet Explorer.

2. In Internet Explorer, in **Address**, type **http://IIS1.testlab.microsoft.com/winxp.gif**. You should see a Windows XP graphic.

3. On CLIENT1, click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**. You should see the contents of the Local Drive (C:) on IIS1.

**Disconnect the L2TP Connection**
On CLIENT1, right-click the L2TPtoCorpnet connection and then click **Disconnect**.

**RADIUS Authentication and Accounting**
To configure RADIUS authentication and accounting for VPN connections, perform the following:

**Configure IAS1 for VPN1 as a RADIUS Client**
On IAS1, add VPN1 as a RADIUS client using the IP address of 172.16.0.4 and a shared secret. To add a RADIUS client, right-click the **RADIUS Clients** folder and click **New RADIUS Client** in the Internet Authentication Service snap-in.

**Configure IAS1 to Log Authentication Events**
On IAS1, enable the logging of accounting and authentication requests from **Settings** tab in the properties of the **Local File** object in the **Remote Access Logging** folder in the Internet Authentication Service snap-in.

**Configure VPN1 for IAS1 as a RADIUS Server**
On VPN1, add IAS1 as a RADIUS server for both the authentication and accounting provider at the IP address of 172.16.0.2 and the shared secret. To configure the Routing and Remote Access service for RADIUS, obtain properties on the VPN server and click the **Security** tab. For RADIUS authentication, select **RADIUS authentication** as the authentication provider and click **Configure** to add IAS1 as the RADIUS server. For RADIUS accounting, select **RADIUS accounting** as the accounting provider and click **Configure** to add IAS1 as the RADIUS server.

**Make PPTP and L2TP Connections**

1. On CLIENT1, make a PPTP connection with VPN1 using the PPTPtoCorpnet connection.

2. Disconnect the PPTP connection.

3. On CLIENT1, make an L2TP connection with VPN1 using the L2TPtoCorpnet connection.

4. Disconnect the L2TP connection.

**Check the System Event Log for RADIUS Events**
On IAS1, use Event Viewer to view IAS events in the system event log for the PPTP and L2TP connections that were recently created using CLIENT1.

**Check RADIUS Authentication and Accounting Logs**
On IAS1, use Windows Explorer to open the *SystemRoot*\System32\Logfiles\Iaslog.log file. Note the authentication and accounting entries for the PPTP and L2TP connections that were recently created using CLIENT1.

**Remote Access Policies for Different Types of VPN Connections**
To create remote access policies for different types of VPN connections, do the following:

**Create Separate Remote Access Policies for PPTP and L2TP Connections**

1. On IAS1, create a new custom remote access policy with the following settings:

   Policy name: PPTP connections

   Conditions:

   **NAS-Port-Type** matches **Virtual (VPN)**

   **Tunnel-Type** matches **Point-to-Point Tunneling Protocol (PPTP)**

   Permission: Grant remote access permission

   Profile settings, **IP** tab:

   Input packet filter:

   Filter action: **Deny all traffic except those listed below**

   Destination network, IP address: 172.16.0.1

   Destination network, Subnet mask: 255.255.255.255

   Protocol: Any

   Output packet filter:

   Filter action: **Deny all traffic except those listed below**

   Source network, IP address: 172.16.0.1

   Destination network, Subnet mask: 255.255.255.255

   Protocol: Any

2. Create a new custom remote access policy with the following settings:

   Policy name: L2TP connections

   Conditions:

**NAS-Port-Type** matches **Virtual (VPN)**

**Tunnel-Type** matches **Layer Two Tunneling Protocol (L2TP)**

Permission: Grant remote access permission

Profile settings, **IP** tab:

Input packet filter:

Filter action: **Deny all traffic except those listed below**

Destination network, IP address: 172.16.0.2

Destination network, Subnet mask: 255.255.255.255

Protocol: Any

Output packet filter:

Filter action: **Deny all traffic except those listed below**

Source network, IP address: 172.16.0.2

Destination network, Subnet mask: 255.255.255.255

Protocol: Any

**Make a PPTP Connection and Test Connectivity**

1. On CLIENT1, make a VPN connection with VPN1 using the PPTPtoCorpnet connection.

2. Use the ping command to ping DC1 at its IP address of 172.16.0.1.

3. Use the ping command to ping IAS1 at its IP address of 172.16.0.2.

   This command fails because packet filtering for all connections that match the **PPTP connections** policy allows only traffic sent to and from the IP address of 172.16.0.1.

4. Disconnect the PPTPtoCorpnet connection.

**Make an L2TP Connection and Test Connectivity**

1. On CLIENT1, make a VPN connection with VPN1 using the L2TPtoCorpnet connection.

2. Use the ping command to ping IAS1 at its IP address of 172.16.0.2.

3. Use the ping command to ping DC1 at its IP address of 172.16.0.1.

   This command fails because packet filtering for all connections that match the **L2TP connections** policy allows only traffic sent to and from the IP address of 172.16.0.2.

4. Disconnect the L2TPtoCorpnet connection.

**Check the System Event Log for IAS Events**

On IAS1, use Event Viewer to view the IAS events in the system event log for the PPTP and L2TP connections that were recently

created by CLIENT1. Note that the authentication event message text contains the name of the remote access policy that accepted the connection.

# Appendix D: Troubleshooting

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

Windows Server 2003 provides the following tools to troubleshoot VPN connections:

- TCP/IP Troubleshooting Tools
- Authentication and Accounting Logging
- Event Logging
- IAS Event Logging
- PPP Logging
- Tracing
- Oakley Logging
- Network Monitor

**TCP/IP Troubleshooting Tools**

The Ping, Tracert, and Pathping tools use ICMP Echo and Echo Reply messages to verify connectivity, display the path to a destination, and test path integrity. The **route print** command can be used to display the IP routing table. Alternately, on the VPN server, you can use the **netsh routing ip show rtmroutes** command or the Routing and Remote Access snap-in. The Nslookup tool can be used to troubleshoot DNS and name resolution issues.

**Authentication and Accounting Logging**

A VPN server running Windows Server 2003 supports the logging of authentication and accounting information for remote access VPN connections in local logging files when Windows authentication or Windows accounting is enabled. This logging is separate from the events recorded in the system event log. You can use the information that is logged to track remote access usage and authentication attempts. Authentication and accounting logging is especially useful for troubleshooting remote access policy issues. For each authentication attempt, the name of the remote access policy that either accepted or rejected the connection attempt is recorded.

Enable authentication and accounting logging from the **Settings** tab on the properties of the **Local File** object in the **Remote Access Logging** folder in the Routing and Remote Access snap-in (if the Routing and Remote Access service is configured for Windows authentication and accounting) or the Internet Authentication Service snap-in (if the Routing and Remote Access service is configured for RADIUS authentication and accounting and the RADIUS server is an IAS server)

The authentication and accounting information is stored in a configurable log file or files stored in the *SystemRoot*\System32\LogFiles folder. The log files are saved in Internet Authentication Service (IAS) or database-compatible format, meaning that any database program can read the log file directly for analysis.

If the VPN server is configured for RADIUS authentication and accounting and the RADIUS server is a computer running Windows Server 2003 and IAS, the authentication and accounting logs are stored in the *SystemRoot*\System32\LogFiles folder on the IAS server computer.

IAS for Windows Server 2003 can also send authentication and accounting information to a Structured Query Language (SQL) database.

**Event Logging**

On the **Logging** tab in the properties of a VPN server in the Routing and Remote Access snap-in, there are four levels of logging. Select **Log all events**, and then try the connection again. After the connection fails, check the system event log for events logged during the connection process. After you are done viewing remote access events, select the **Log errors and warnings** option on the **Logging** tab to conserve system resources.

**IAS Event Logging**

If your VPN servers are configured for RADIUS authentication and your RADIUS servers are computers running Windows Server 2003 and IAS, check the system event log for IAS events for rejected or accepted connection attempts. IAS system event log entries contain a lot of information about the connection attempt including the name of the remote access policy that accepted or rejected the connection attempt. IAS event logging for rejected or accepted connection attempts is enabled by default and configured from the **Service** tab from the properties of an IAS server in the Internet Authentication Service snap-in.

#### PPP logging

PPP logging records the series of programming functions and PPP control messages during a PPP connection and is a valuable source of information when you are troubleshooting the failure of a PPP connection. To enable PPP logging, select the **Log additional Routing and Remote Access information** option on the **Logging** tab on the properties of a remote access server.

By default, the PPP log is stored as the Ppp.log file in the *SystemRoot*\Tracing folder.

#### Tracing

The Windows Server 2003 Routing and Remote Access service has an extensive tracing capability that you can use to troubleshoot complex network problems. You can enable the components in Windows Server 2003 to log tracing information to files using the Netsh command or through the registry.

#### Enabling Tracing with Netsh

You can use the Netsh command to enable and disable tracing for specific components or for all components.

**To enable and disable tracing for a specific component, use the following syntax:**

**netsh ras set tracing** *Component* **enabled|disabled**

where *Component* is a component in the list of Routing and Remote Access service components found in the Windows Server 2003 registry under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing. For example, to enable tracing for the RASAUTH component, the command is:

**netsh ras set tracing rasauth enabled**

To enable tracing for all components, use the following command:

**netsh ras set tracing * enabled**

#### Enabling Tracing Through the Registry

**You can configure the tracing function by changing settings in the Windows Server 2003 registry under:**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing

You can enable tracing for each Routing and Remote Access service component by setting the registry values described later. You can enable and disable tracing for components while the Routing and Remote Access service is running. Each component is capable of tracing and appears as a subkey under the preceding registry key.

To enable tracing for each component, you can configure the following registry entries for each protocol key:

**EnableFileTracing** REG_DWORD *Flag*

You can enable logging tracing information to a file by setting **EnableFileTracing** to 1. The default value is 0.

**FileDirectory** REG_EXPAND_*SZ Path*

You can change the default location of the tracing files by setting **FileDirectory** to the path you want. The file name for the log file is the name of the component for which tracing is enabled. By default, log files are placed in the *SystemRoot*\Tracing folder.

**FileTracingMask** REG_DWORD *LevelOfTracingInformationLogged*

**FileTracingMask** determines how much tracing information is logged to the file. The default value is 0xFFFF0000.

**MaxFileSize** REG_DWORD *SizeOfLogFile*

You can change the size of the log file by setting different values for **MaxFileSize**. The default value is 0x10000 (64K).

**Note**

Tracing consumes system resources and should be used sparingly to help identify network problems. After the trace is captured or the problem is identified, you should immediately disable tracing. Do not leave tracing enabled on multiprocessor computers.

**Note**

Tracing information can be complex and very detailed. Most of the time this information is useful only to Microsoft support professionals or to network administrators who are very experienced with the Routing and Remote Access service. Tracing information can be saved as files and sent to Microsoft support for analysis.

**Oakley Logging**

You can use the Oakley log to view details about the SA establishment process. The Oakley log is enabled in the registry. It is not enabled by default. To enable the Oakley log, set the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLogging registry setting to 1. The Oakley key does not exist by default and must be created.

After it is enabled, the Oakley log, which is stored in the *SystemRoot*\Debug folder, records all IPSec SA negotiations. A new Oakley.log file is created each time the IPSec Policy Agent is started and the previous version of the Oakley.log file is saved as Oakley.log.sav.

To activate the new EnableLogging registry setting after modifying its value, stop and start the IPSec Policy Agent and related IPSec services by running the following sequence of commands:

1. Stop the Routing and Remote Access service using the **net stop remoteaccess** command.
2. Stop the IPSec services using the **net stop policyagent** command.
3. Start the IPSec services using the **net start policyagent** command.
4. Start the Routing and Remote Access service using the **net start remoteaccess** command.

**Network Monitor**

Use Network Monitor, a packet capture and analysis tool supplied with Windows Server 2003, to capture and view the traffic sent between a VPN server and VPN client during the VPN connection process and during data transfer. You cannot interpret the encrypted portions of VPN traffic with Network Monitor. Network Monitor is installed as an optional networking component.

The proper interpretation of the remote access and VPN traffic with Network Monitor requires an in-depth understanding of PPP, PPTP, IPSec, and other protocols. Network Monitor captures can be saved as files and sent to Microsoft support for analysis.

**Troubleshooting Remote Access VPNs**

**Remote access VPN problems typically fall into the following categories:**

- Connection attempt is rejected when it should be accepted.
- L2TP/IPSec authentication issues.
- EAP-TLS authentication issues.
- Quarantine network access control issues
- Connection attempt is accepted when it should be rejected.
- Unable to reach locations beyond the VPN server.
- Unable to establish a tunnel.

Use the following troubleshooting tips to isolate the configuration or infrastructure problem causing the stated VPN problem.

**Connection Attempt is Rejected When it Should be Accepted**

- Using the Ping command, verify that the host name is being resolved to its correct IP address. The ping itself might not be successful due to packet filtering that is preventing the delivery of ICMP messages to and from the VPN server.

- Verify that the VPN client's credentials, consisting of user name, password, and domain name, are correct and can be validated by the VPN server.

- Verify that the user account of the VPN client is not locked out, expired, disabled, or that the time the connection is being made does not correspond to the configured logon hours. If the password on the account has expired, verify that the remote access VPN client is using MS-CHAP or MS-CHAP v2. MS-CHAP and MS-CHAP v2 are the only authentication protocols provided with Windows Server 2003 that allow you to change an expired password during the connection process.

  For an administrator-level account whose password has expired, reset the password using another administrator-level account.

- Verify that the user account has not been locked out due to remote access account lockout.

- Verify that the Routing and Remote Access service is running on the VPN server.

- Verify that the VPN server is enabled for remote access from the **General** tab on the properties of a VPN server in the Routing and Remote Access snap-in.

- Verify that the **WAN Miniport (PPTP)** and **WAN Miniport (L2TP)** devices are enabled for inbound remote access from the properties of the **Ports** object in the Routing and Remote Access snap-in.

- Verify that the VPN client, the VPN server, and the remote access policy corresponding to VPN connections are configured to use at least one common authentication method.

- Verify that the VPN client and the remote access policy corresponding to VPN connections are configured to use at least one common encryption strength.

- Verify that the parameters of the connection have permission through remote access policies.

  In order for the connection to be accepted, the parameters of the connection attempt must:

  - Match all of the conditions of at least one remote access policy.

  - Be granted remote access permission through the user account (set to **Allow access**), or if the user account has the **Control access through Remote Access Policy** option selected, the remote access permission of the matching remote access policy must have the **Grant remote access permission** option selected.

  - Match all the settings of the profile.

  - Match all the settings of the dial-in properties of the user account.

  To obtain the name of the remote access policy that rejected the connection attempt, scan the accounting log for the entry corresponding to the connection attempt for the policy name. If IAS is being used as a RADIUS server, check the System event log for an entry for the connection attempt.

- If you are logged on using an account with domain administrator permissions when you run the Routing and Remote Access Server Setup Wizard, it automatically adds the computer account of the **RAS and IAS Servers** domain-local security group. This group membership allows the VPN server computer to access user account information. If the VPN

server is unable to access user account information, verify that:

  - The computer account of the VPN server computer is a member of the **RAS and IAS Servers** security group for all the domains that contain user accounts for which the VPN server is authenticating remote access. You can use the **netsh ras show registeredserver** command at the command prompt to view the current registration. You can use the **netsh ras add registeredserver** command to register the server in a domain in which the VPN server is a member or other domains. Alternately, you or your domain administrator can add the computer account of the VPN server computer to the **RAS and IAS Servers** security group of all the domains that contain user accounts for which the VPN server is authenticating remote access.

  - If you add or remove the VPN server computer to the **RAS and IAS Servers** security group, the change does not take effect immediately (due to the way that Windows Server 2003 caches Active Directory information). For the change to take effect immediately, you need to restart the VPN server computer.

- For a VPN server that is a member server in a mixed-mode or native-mode Active Directory domain that is configured for Windows authentication, verify that:

  - The **RAS and IAS Servers** security group exists. If not, then create the group and set the group type to **Security** and the group scope to **Domain local**.

  - The **RAS and IAS Servers** security group has **Read** permission to the RAS and **IAS Servers Access Check** object.

- Verify that IP is enabled for remote access on the VPN server.

- Verify that all of the PPTP or L2TP ports on the VPN server are not already being used. If necessary, change the number of PPTP to L2TP ports from the properties of the **Ports** object in the Routing and Remote Access snap-in to allow more concurrent connections.

- Verify that the VPN server supports the tunneling protocol of the VPN client.

  By default, Windows 2000 remote access VPN clients have the **Automatic** server type option selected, which means that they try to establish an L2TP/IPSec-based VPN connection first, then they try a PPTP-based VPN connection. If either the **Point to Point Tunneling Protocol (PPTP)** or **Layer-2 Tunneling Protocol (L2TP)** server type option is selected, verify that the selected tunneling protocol is supported by the VPN server.

  By default, Windows XP remote access VPN clients have the **Automatic** VPN type option selected, which means that they try to establish a PPTP -based VPN connection first, then they try an L2TP/IPSec-based VPN connection. If either the **PPTP VPNor L2TP IPSec VPN** type is selected, verify that the VPN server supports the selected tunneling protocol.

  Depending on your selections when running the Routing and Remote Access Server Setup Wizard, a Windows Server 2003based computer running the Routing and Remote Access service is a PPTP and L2TP server with five or 128 L2TP ports and five or 128 PPTP ports. To create a PPTP-only server, set the number of L2TP ports to zero. To create an L2TP-only server, set the number of PPTP ports to 1 and disable remote access inbound connections and demand-dial connections for the **WAN Miniport (PPTP)** device from the properties of the **Ports** object in the Routing and Remote Access snap-in.

- For L2TP/IPSec connections, verify that computer certificates, also known as machine certificates, are installed on the VPN client and the VPN server.

- If the VPN server is configured with static IP address pools, verify that there are enough addresses. If all of the addresses in the static pools have been allocated to connected VPN clients, the VPN server is unable to allocate an IP address for TCP/IP-based connections, and the connection attempt is rejected.

- Verify the configuration of the authentication provider. The VPN server can be configured to use either Windows or RADIUS to authenticate the credentials of the VPN client.

  - For RADIUS authentication, verify that the VPN server computer can communicate with the RADIUS server.

- For a VPN server that is a member of a native-mode domain, verify that the VPN server has joined the domain.

- For a Windows NT version 4.0 Service Pack 4 and later VPN server that is a member of a mixed mode domain or a Windows Server 2003 VPN server that is a member of a Windows NT 4.0 domain that is accessing user account properties for a user account in a trusted Active Directory domain, verify that the **Everyone** group is added to the **Pre-Windows 2000 Compatible Access** group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, issue the **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** command on a domain controller computer and then restart the domain controller.

- For a Windows NT version 4.0 Service Pack 3 and earlier VPN server that is a member of a mixed-mode domain, verify that the **Everyone** group has been granted list contents, read all properties, and read permissions to the root node of your domain and all sub-objects of the root domain.

- For PPTP connections using MS-CHAP and attempting to negotiate 40-bit MPPE encryption, verify that the user's password is not larger than 14 characters.

#### L2TP/IPSec Authentication Issues

**The most common problems that cause L2TP/IPSec connections to fail are the following:**

- No certificate

  By default, L2TP/IPSec connections require that the remote access server and remote access client exchange computer certificates for IPSec peer authentication. Check the Local Computer certificate stores of both the remote access client and remote access server using the Certificates snap-in to ensure that a suitable certificate exists.

- Incorrect certificate

  If certificates exist, they must be verifiable. Unlike manually configuring IPSec rules, the list of certification authorities (CAs) for L2TP/IPSec connections is not configurable. Instead, each computer in the L2TP connection sends a list of root CAs to its IPSec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Computer A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPSec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPSec peer, Computer B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPSec security negotiation fails.

  The VPN client must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN server trusts. Additionally, the VPN server must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN client trusts.

  By default, L2TP/IPSec connections require that the remote access server and remote access client exchange computer certificates for IPSec peer authentication. Check the Local Computer certificate stores of both the remote access client and remote access server using the Certificates snap-in to ensure that a suitable certificate exists.

- A NAT between the remote access client and remote access server

  If there is a NAT between a Windows 2000 or Windows XP-based L2TP/IPSec client and a Windows Server 2003 L2TP/IPSec server, you cannot establish an L2TP/IPSec connection unless both the client and server support IPSec NAT-T.

- A firewall between the remote access client and remote access server

  If there is a firewall between a Windows L2TP/IPSec client and a Windows Server 2003 L2TP/IPSec server and you cannot establish an L2TP/IPSec connection, verify that the firewall allows L2TP/IPSec traffic to be forwarded. For more information, see "Integrating VPN Servers and Firewalls" in this article.

One of the best tools for troubleshooting IPSec authentication issues is the Oakley log. For more information, see "Oakley

Logging" in this article.

**EAP-TLS Authentication Issues**

When EAP-TLS is used for authentication, the VPN client submits a user certificate and the authenticating server (the VPN server or the RADIUS server) submits a computer certificate.

**In order for the authenticating server to validate the certificate of the VPN client, the following must be true for each certificate in the certificate chain sent by the VPN client:**

- The current date must be within the validity dates of the certificate.

  When certificates are issued, they are issued with a range of valid dates, before which they cannot be used and after which they are considered expired.

- The certificate has not have been revoked.

  Issued certificates can be revoked at any time. Each issuing CA maintains a list of certificates that should no longer be considered valid by publishing an up-to-date certificate revocation list (CRL). By default, the authenticating server checks all the certificates in the VPN clients certificate chain (the series of certificates from the VPN client certificate to the root CA) for revocation. If any of the certificates in the chain have been revoked, certificate validation fails. This behavior can be modified with registry settings described later in this topic.

  To view the CRL distribution points for a certificate in the Certificates snap-in, obtain the certificate properties, click the **Details** tab, and then click the **CRL Distribution Points** field.

  The certificate revocation validation only works as well as the CRL publishing and distribution system. If the CRL in a certificate is not updated often, a certificate that has been revoked can still be used and considered valid because the published CRL that the authenticating server is checking is out of date.

- The certificate has a valid digital signature.

  CAs digitally sign certificates they issue. The authenticating server verifies the digital signature of each certificate in the chain, with the exception of the root CA certificate, by obtaining the public key from the certificates issuing CA and mathematically validating the digital signature.

The VPN client certificate must also have the Client Authentication certificate purpose (also known as Enhanced Key Usage [EKU]) (Object identifier 1.3.6.1.5.5.7.3.2) and must either contain a UPN of a valid user account or FQDN of valid computer account for the Subject Alternative Name property of the certificate.

To view the EKU for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the **Details** tab, and then click the **Enhanced Key Usage** field. To view the subject alternative name property for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the **Details** tab, and then click the **Subject Alternative Name** field.

Finally, to trust the certificate chain offered by the VPN client, the authenticating server must have the root CA certificate of the issuing CA of the VPN client certificate installed in its Trusted Root Certification Authorities store.

Additionally, the authenticating server verifies that the identity sent in the EAP-Response/Identity message is the same as the name in the Subject Alternative Name property of the certificate. This prevents a malicious user from masquerading as a different user from that specified in the EAP-Response/Identity message.

If the authenticating server is a Windows Server 2003 VPN server or an IAS server, the following registry settings in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13 can modify the behavior of the EAP-TLS when performing certificate revocation:

- IgnoreNoRevocationCheck

  When set to 1, the authenticating server allows EAP-TLS clients to connect even when it does not perform or cannot complete a revocation check of the VPN client's certificate chain (excluding the root certificate). Typically, revocation checks fail because the certificate doesn't include CRL information.

  IgnoreNoRevocationCheck is set to 0 (disabled) by default. An EAP-TLS client cannot connect unless the server completes

a revocation check of the client's certificate chain (including the root certificate) and verifies that none of the certificates have been revoked.

You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties.

- IgnoreRevocationOffline

  When set to 1, the authenticating server allows EAP-TLS clients to connect even when a server that stores a CRL is not available on the network. IgnoreRevocationOffline is set to 0 by default. The authenticating server does not allow clients to connect unless it can complete a revocation check of their certificate chain and verify that none of the certificates has been revoked. When it cannot connect to a server that stores a revocation list, EAP-TLS considers the certificate to have failed the revocation check.

  Setting IgnoreRevocationOffline to 1 prevents certificate validation failure because poor network conditions prevented their revocation check from completing successfully.

- NoRevocationCheck

  When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the VPN client's certificate. The revocation check verifies that the VPN clients certificate and the certificates in its certificate chain have not been revoked. NoRevocationCheck is set to 0 by default.

- NoRootRevocationCheck

  When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the VPN client's root CA certificate. NoRootRevocationCheck is set to 0 by default. This entry only eliminates the revocation check of the client's root CA certificate. A revocation check is still performed on the remainder of the VPN client's certificate chain.

  You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties. Also, this entry can prevent certification-related delays that occur when a certificate revocation list is offline or is expired.

All of these registry settings must be added as a DWORD type and have the valid values of 0 or 1. The VPN client does not use these settings.

In order for the VPN client to validate the certificate of the authenticating server for either EAP-TLS authentication, the following must be true for each certificate in the certificate chain sent by the authenticating server:

- The current date must be within the validity dates of the certificate.

  When certificates are issued, they are issued with a range of valid dates, before which they cannot be used and after which they are considered expired.

- The certificate has a valid digital signature.

  CAs digitally sign certificates they issue. The VPN client verifies the digital signature of each certificate in the chain, with the exception of the root CA certificate, by obtaining the public key from the certificates issuing CA and mathematically validating the digital signature.

Additionally, the authenticating server computer certificate must have the Server Authentication EKU (Object identifier 1.3.6.1.5.5.7.3.1). To view the EKU for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the **Details** tab, and then click the **Enhanced Key Usage** field.

Finally, to trust the certificate chain offered by the authenticating server, the VPN client must have the root CA certificate of the issuing CA of the authenticating server certificate installed in its Trusted Root Certification Authorities store.

Notice that the VPN client does not perform certificate revocation checking for the certificates in the certificate chain of the authenticating servers computer certificate. The assumption is that the VPN client does not yet have a connection to the

network, and therefore cannot access a Web page or other resource in order to check for certificate revocation.

**Connection Attempt is Accepted When it Should be Rejected**

- Verify that the remote access permission on the user account is set to either **Deny access** or **Control accessthrough Remote Access Policy**. If set to the latter, verify that the first matching remote access policy's remote access permission is set to **Deny remote access permission**. To obtain the name of the remote access policy that accepted the connection attempt, scan the accounting log for the entry corresponding to the connection attempt for the policy name.

- If you have created a remote access policy to explicitly reject all connections, verify the policy conditions, remote access permission, and profile settings.

**Unable to Reach Locations Beyond the VPN Server**

- Verify that either the protocol is enabled for routing or that dial-in clients are allowed to access the entire network for LAN protocols being used by the VPN clients.

- Verify the IP address pools of the VPN server.

  If the VPN server is configured to use a static IP address pool, verify that the routes to the range of addresses defined by the static IP address pools are reachable by the hosts and routers of the intranet. If not, then IP route consisting of the VPN server static IP address pools, as defined by the IP address and mask of the range, must be added to the routers of the intranet or enable the routing protocol of your routed infrastructure on the VPN server. If the routes to the remote access VPN client subnets are not present, remote access VPN clients cannot receive traffic from locations on the intranet. Routes for the subnets are implemented either through static routing entries or through a routing protocol, such as Open Shortest Path First (OSPF) or Routing Information Protocol (RIP).

  If the VPN server is configured to use DHCP for IP address allocation and no DHCP server is available, the VPN server assigns addresses from the Automatic Private IP Addressing (APIPA) address range from 169.254.0.1 through 169.254.255.254. Allocating APIPA addresses for remote access clients works only if the network to which the VPN server is attached is also using APIPA addresses.

  If the VPN server is using APIPA addresses when a DHCP server is available, verify that the proper adapter is selected from which to obtain DHCP-allocated IP addresses. The Routing and Remote Access Server Setup Wizard assign the default adapter automatically. You can manually choose a LAN adapter from the **Adapter** list on the **IP** tab on the properties of a VPN server in the Routing and Remote Access snap-in.

  If the static IP address pools are a range of IP addresses that are a subset of the range of IP addresses for the network to which the VPN server is attached, verify that the range of IP addresses in the static IP address pool are not assigned to other TCP/IP nodes, either through static configuration or through DHCP.

- Verify that there are no packet filters on the profile properties of the remote access policy corresponding to VPN connections that are preventing the sending or receiving of traffic.

**Unable to Establish Tunnel**

- Verify that packet filtering on a router interface between the VPN client and the VPN server is not preventing the forwarding of tunneling protocol traffic. See Appendix A for information about the types of traffic that must be allowed for PPTP and L2TP/IPSec traffic.

  On a Windows Server 2003based VPN server, IP packet filtering can be separately configured from the advanced TCP/IP properties and from the Routing and Remote Access snap-in. Check both places for filters that might be excluding VPN connection traffic.

- Verify that the Winsock Proxy client is not currently running on the VPN client.

  When the Winsock Proxy client is active, Winsock API calls such as those used to create tunnels and send tunneled data are intercepted and forwarded to a configured proxy server.

A proxy server based computer allows an organization to access specific types of Internet resources (typically Web and FTP) without directly connecting that organization to the Internet. The organization can instead use private IP network IDs (such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16).

Proxy servers are typically used so that private users in an organization can have access to public Internet resources as if they were directly attached to the Internet. VPN connections are typically used so that authorized public Internet users can gain access to private organization resources as if they were directly attached to the private network. A single computer can act as a proxy server (for private users) and a VPN server (for authorized Internet users) to facilitate both exchanges of information.

# Appendix E: Deploying a Certificate Infrastructure

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

In a typical enterprise deployment, the certificate infrastructure is configured using single root CA in a three-level hierarchy consisting of root CA/intermediate CAs/issuing CAs. Medium-sized organizations should use a two-level hierarchy consisting of root CA/issuing CAs. Small organizations can use a single CA that is both the root CA and the issuing CA.

For VPN connections, issuing CAs are configured to issue computer certificates or user certificates. When the computer or user certificate is installed on the VPN client, the issuing CA certificate, intermediate CA certificates, and the root CA certificate are also installed. When the computer certificate is installed on the authenticating server, the issuing CA certificate, intermediate CA certificates, and the root CA certificate are also installed. The issuing CA for the computer certificate installed on the authenticating server can be different than the issuing CA for the VPN client certificates. In this case, both the VPN client and the authenticating server computer have all the required certificates to perform certificate validation for both IPSec and EAP-TLS authentication.

**When deploying a certificate infrastructure, use the following best practices:**

- Plan your certificate infrastructure before deploying CAs.
- The root CA should be offline and its signing key should be secured by a Hardware Security Module (HSM) and kept in a vault to minimize potential for key compromise.
- Enterprise organizations should not issue certificates to users or computers directly from the root CA but rather should deploy the following:
    - An offline root CA
    - Offline intermediate CAs
    - Online issuing CAs

  This CA infrastructure provides flexibility and insulates the root CA and intermediate CAs from attempts to compromise its private key by malicious users. The offline root and intermediate CAs do not have to be Windows Server 2003 CAs. Issuing CAs can be subordinates of a third-party intermediate CA.
- Backing up the CA database, the CA certificate, and the CA keys is essential to protect against the loss of critical data. The CA should be backed up on a regular basis (daily, weekly, monthly) based on the number of certificates issued over the same interval. The more certificates issued, the more frequently you should back up the CA.
- You should review the concepts of security permissions and access control in Windows, because enterprise certification authorities issue certificates based on the security permissions of the certificate requester.

If you want to take advantage of auto-enrollment for computer certificates and the requesting of certificates using the Certificates snap-in, use Windows Server 2003 Certificate Services and create an enterprise CA at the issuer CA level.

For more information, see the topic titled "Checklist: Creating a certification hierarchy with an offline root certification authority" in Windows Server 2003 Help and Support.

**Certificate Revocation and EAP-TLS Authentication**

By default, the authenticating server checks for certificate revocation for all the certificates in the certificate chain sent by the VPN client during the EAP-TLS authentication process. If certificate revocation fails for any of the certificates in the chain, the connection fails authentication and is rejected. The certificate revocation check for a certificate can fail because of the following:

- **The certificate has been revoked**. The issuer of the certificate has explicitly revoked the certificate.
- **The certificate revocation list (CRL) for the certificate is not reachable or available**. CAs maintain CRLs and publish them to specific CRL distribution points. The CRL distribution points are included in the CRL Distribution Points

property of the certificate. If the CRL distribution points cannot be contacted to check for certificate revocation, then the certificate revocation check fails. Additionally, if there are no CRL distribution points in the certificate, the authenticating server cannot verify that the certificate has not been revoked and the certificate revocation check fails.

- **The publisher of the CRL did not issue the certificate**. Included in the CRL is the publishing CA. If the publishing CA of the CRL does not match the issuing CA for the certificate for which certificate revocation is being checked, then the certificate revocation check fails.

- **The CRL is not current**. Each published CRL has a range of valid dates. If the CRL Next update date has passed, the CRL is considered invalid and the certificate revocation check fails. New CRLs should be published before the expiration date of the last published CRL.

This behavior can be modified using the following registry settings in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13 on the authenticating server:

- **IgnoreNoRevocationCheck** When set to 1, the authenticating server allows EAP-TLS clients to connect even when it does not perform or cannot complete a revocation check of the client's certificate chain (excluding the root certificate). Typically, revocation checks fail because the certificate doesn't include CRL information. IgnoreNoRevocationCheck is set to 0 (disabled) by default. An EAP-TLS client cannot connect unless the authenticating server completes a revocation check of the client's certificate chain (including the root certificate) and verifies that none of the certificates have been revoked. You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties.

- **IgnoreRevocationOffline** When set to 1, the authenticating server allows EAP-TLS clients to connect even when a server that stores a CRL is not available on the network. IgnoreRevocationOffline is set to 0 by default. The authenticating server does not allow clients to connect unless it can complete a revocation check of their certificate chain and verify that none of the certificates has been revoked. When it cannot connect to a server that stores a revocation list, EAP-TLS considers the certificate to have failed the revocation check. Setting IgnoreRevocationOffline to 1 prevents certificate validation failure because poor network conditions prevented their revocation check from completing successfully.

- **NoRevocationCheck** When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the wireless client's certificate. The revocation check verifies that the VPN clients certificate and the certificates in its certificate chain have not been revoked. NoRevocationCheck is set to 0 by default.

- **NoRootRevocationCheck** When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the VPN client's root CA certificate. NoRootRevocationCheck is set to 0 by default. This entry only eliminates the revocation check of the client's root CA certificate. A revocation check is still performed on the remainder of the VPN client's certificate chain. You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties. Also, this entry can prevent certification-related delays that occur when a certificate revocation list is offline or is expired.

All of these registry settings must be added as a DWORD type and have the valid values of 0 or 1. The VPN client does not perform certificate revocation checking of the authenticating server's certificate and does not use these settings.

Because certificate revocation checking can prevent VPN access due to the inaccessibility or expiration of CRLs for each certificate in the certificate chain, design your certificate infrastructure for high availability of CRLs. For instance, configure multiple CRL distribution points for each CA in the certificate hierarchy and configure publication schedules that ensure that the most current CRL is always published and available.

Certificate revocation checking is only as accurate as the last published CRL. For example, if a certificate is revoked, by default the new CRL containing the newly revoked certificate is not automatically published. CRLs are typically published based on a configurable schedule. This means that the revoked certificate can still be used to authenticate because the published CRL is not current; it does not contain the revoked certificate and can therefore still be used to create wireless connections. To prevent this from occurring, the network administrator must manually publish the new CRL with the newly revoked certificate.

By default the authenticating server uses the CRL distribution points in the certificates. However, it is also possible to store a local copy of the CRL on the authenticating server. In this case, the local CRL is used during certificate revocation checking. If a new CRL is manually published to the Active Directory, the local CRL on the authenticating server is not updated. The local CRL is updated when it expires. This can create a situation whereby a certificate is revoked, the CRL is manually published, but the

authenticating server still allows the connection because the local CRL has not yet been updated.

**Using Third-party CAs for EAP-TLS Authentication**
You can use third-party CAs to issue certificates for EAP-TLS authentication as long as the certificates installed can be validated and have the appropriate properties.

**Certificates on the Authenticating Servers**
For the computer certificates installed on the authenticating servers (either the VPN servers or the IAS servers), the following must be true:

- They must be installed in the Local Computer certificate store.
- They must have a corresponding private key.
- The cryptographic service provider for the certificates supports SChannel. If not, the certificate cannot be used and it is not selectable from the properties of the **Smart Card or Other Certificate** EAP type on the **Authentication** tab on the properties of a profile for a remote access policy.
- They must contain the Server Authentication certificate purpose (also known as an Enhanced Key Usage [EKU]). An EKU is identified using an object identifier (also known as an OID). The object identifier for Server Authentication is "1.3.6.1.5.5.7.3.1".
- They must contain the fully qualified domain name (FQDN) of the computer account of the authenticating server in the Subject Alternative Name property of the certificate.

Additionally, the root CA certificates of the CAs that issued the VPN client user certificates must be installed in the Trusted Root Certification Authorities certificate store of the authenticating servers.

**Certificates on VPN Client Computers**
For the user certificates installed on VPN client computers, the following must be true:

- They must have a corresponding private key.
- They must contain the Client Authentication EKU (object identifier "1.3.6.1.5.5.7.3.2").
- They must be installed in the Current User certificate store.
- They must contain the universal principal name (UPN) of the user account in the Subject Alternative Name property of the certificate.

Additionally, the root CA certificates of the CAs that issued the IAS server computer certificates must be installed in the Trusted Root Certification Authorities store of the VPN client computers.

# Summary and Related Links

Updated: October 7, 2005

Applies To: Windows Server 2003 with SP1

This paper described in detail the components and their associated design decisions for a Windows Server 2003-based remote access VPN deployment including VPN clients, Internet infrastructure, authentication protocols, VPN protocols, VPN servers, intranet infrastructure, AAA infrastructure, and certificate infrastructure. This paper also included detailed walkthroughs of PTPP and L2TP-based VPN deployments, details of firewall configuration, how to set up a VPN test lab, and a discussion of VPN troubleshooting tools and common VPN problems with suggested solutions.

**Related Links**
See the following resources for further information:

- Windows VPN Web site at http://www.microsoft.com/windowsserver2003/technologies/networking/vpn/default.mspx
- Microsoft L2TP/IPSec VPN Client at http://www.microsoft.com/windows2000/techinfo/howitworks/communications/remoteaccess/l2tpclientrelnotes.asp
- Internet Authentication Service Web site at http://www.microsoft.com/windowsserver2003/technologies/ias/default.mspx

For the latest information about Windows Server, see the Windows Server 2003 Web site at http://www.microsoft.com/windowsserver2003/default.mspx.

# Virtual Private Networking with Windows Server 2003: Deploying Site-to-Site VPNs

Updated: October 7, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: April 1, 2003 | Updated: October 5, 2004

**In This White Paper**

Introduction (Virtual Private Networking with Windows Server 2003: Deploying Site-to-Site VPNs)

Components of Windows Server 2003 Site-to-Site VPNs

Deploying a PPTP-based Site-to-Site VPN Connection

Deploying an L2TP/IPSec-based Site-to-Site VPN Connection

Appendix A: Configuring Firewalls for a VPN Router Running Windows Server 2003

Appendix B: Alternate Configurations -- Virtual Private Networking

Appendix C: Troubleshooting

Summary -- Virtual Private Networking

# Introduction (Virtual Private Networking with Windows Server 2003: Deploying Site-to-Site VPNs)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A virtual private network (VPN) is the extension of a private network that encompasses links across shared or public networks like the Internet. With a VPN, you can send data between two computers across a shared or public network in a manner that emulates a point-to-point private link (such as a long haul T-Carrier-based wide area network [WAN] link). Virtual private networking is the act of creating and configuring a virtual private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information, which allows the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The logical link over which the private data is encapsulated and encrypted is a VPN connection.

Figure 1 shows the logical equivalent of a VPN connection.




**Figure 1: The logical equivalent of a VPN connection**

Users working at home or on the road can use VPN connections to establish a remote access connection to an organization server by using the infrastructure provided by a public network such as the Internet. From the user's perspective, the VPN connection is a point-to-point connection between the computer (the VPN client) and an organization server (the VPN server). The exact infrastructure of the shared or public network is irrelevant because it appears logically as if the data is sent over a dedicated private link.

Organizations can also use VPN connections to establish site-to-site connections with geographically separate offices or with other organizations over a public network such as the Internet while maintaining secure communications. A site-to-site VPN connection across the Internet logically operates as a dedicated WAN link.

With both remote access and site-to-site connections, an organization can use VPN connections to trade long-distance dial-up or leased lines for local dial-up or leased lines to an Internet service provider (ISP).

There are two types of PPP-based site-to-site VPN technology in the Windows Server 2003 operating system:

1. Point-to-Point Tunneling Protocol (PPTP)

   PPTP uses user-level Point-to-Point Protocol (PPP) authentication methods and Microsoft Point-to-Point Encryption (MPPE) for data encryption.

2. Layer Two Tunneling Protocol (L2TP) with Internet Protocol security (IPSec)

   L2TP with IPSec (L2TP/IPSec) uses user-level PPP authentication methods and IPSec for computer-level authentication using certificates and data authentication, integrity, and encryption.

**Note**

Using IPSec tunnel mode for site-to-site VPN connections is possible using computers running Windows Server 2003. Because the IPSec tunnel is not represented as a logical interface over which packets can be forwarded and received, routing protocols do not operate over IPSec tunnels. Because the configuration of IPSec tunnels for site-to-site VPN connections is vastly different, it is not discussed here. For more information, see article 252735, "How to Configure IPSec Tunneling in Windows 2000," in the Microsoft Knowledge Base.

For encryption, you can use either link encryption or end-to-end encryption in addition to link encryption:

- Link encryption encrypts the data only on the link between the routers. For PPTP connections, you must use MPPE in conjunction with MS-CHAP, MS-CHAP v2, or EAP-TLS authentication. For L2TP/IPSec connections, IPSec provides encryption on the link between the routers.

- End-to-end encryption encrypts the data between the source host and its final destination. You can use IPSec after the site-to-site VPN connection is made to encrypt data from the source host to the destination host.

**Overview of Demand-Dial Routing in Windows Server 2003**

The Windows Server 2003 Routing and Remote Access service includes support for demand-dial routing (also known as dial-on-demand routing) over dial-up connections (such as analog phone lines or ISDN), VPN connections, and PPP over Ethernet (PPPoE) connections. Demand-dial routing is the forwarding of packets across a Point-to-Point Protocol (PPP) link. The PPP link is represented inside the Windows Server 2003 Routing and Remote Access service as a demand-dial interface, which can be used to create on-demand connections across dial-up, non-permanent, or persistent media. Demand-dial connections allow you to use dial-up telephone lines instead of leased lines for low-traffic situations and to leverage the connectivity of the Internet to connect branch offices with VPN connections.

Demand-dial routing is not the same as remote access. While remote access connects a single computer to a network; demand-dial routing connects entire networks. However, both use PPP as the protocol through which to negotiate and authenticate the connection and encapsulate the data sent over it. As implemented in the Windows Server 2003 Routing and Remote Access service, both remote access and demand-dial connections can be enabled separately. However, they still share the same:

- Dial-in properties behavior of user accounts.

- Security (authentication protocols and encryption).

- Remote access policies usage.

- Windows or Remote Authentication Dial-In User Service (RADIUS) usage (for authentication, authorization, and accounting).

- IP address assignment and configuration.

- PPP features usage, such as Microsoft Point-to-Point Compression (MPPC), Multilink PPP, and Bandwidth Allocation Protocol (BAP).

- Troubleshooting facilities, including event logging, Windows or RADIUS authentication and accounting logging, and tracing.

While the concept of demand-dial routing is fairly simple, configuration of demand-dial routing is relatively complex. This complexity is due to the following factors:

- **Connection endpoint addressing.** The connection must be made over public data networks, such as the analog phone system or the Internet. The endpoint of the connection must be identified by a phone number for dial-up connections, and either a fully qualified host name or IP address for VPN connections.

- **Authentication and authorization of the caller**. Anyone calling the router must be authenticated and authorized. Authentication is based on the caller's set of credentials that are passed during the connection establishment process. The credentials that are passed must correspond to an account. Authorization is granted based on the dial-in properties

of the account and remote access policies.

- **Differentiation between remote access clients and calling routers**. Both routing and remote access services coexist on the same computer running Windows Server 2003. Both remote access clients and demand-dial routers can initiate a connection. The computer running Windows Server 2003 that answers a connection attempt must be able to distinguish a remote access client from a demand-dial router.

  If the user name, which is included in the authentication credentials sent by the router that initiates the connection (the calling router), matches the name of a demand-dial interface on the Windows Server 2003 that answers the connection attempt (the answering router), the connection is a demand-dial connection. Otherwise, the incoming connection is a remote access connection.

- **Configuration of both ends of the connection**. Both ends of the connection must be configured, even if only one end of the connection is initiating a demand-dial connection. Configuring only one side of the connection means that packets are successfully routed in only one direction. Communication typically requires that information travel in both directions.

- **Configuration of static routes**. You should not use dynamic routing protocols over temporary demand-dial connections. Therefore, routes for network IDs that are available across the demand-dial interface must be added, as static routes, to the routing tables of the demand-dial routers. You can add static routes manually or by using auto-static updates.

**Demand-dial Routing Updates**

While demand-dial routing can save connection costs, typical routing protocols rely on a periodic advertising process to communicate routing information. For example, RIP for IP advertises the contents of its routing table every 30 seconds on all interfaces. This behavior is not a problem for permanently connected LAN or WAN lines. For usage-sensitive dial-up WAN lines, this type of periodic behavior could cause the router to call another router every 30 seconds, which may result in an undesirable phone bill. Therefore, you should not run routing protocols across temporary dial-up WAN lines.

If you do not use routing protocols to update the routing tables, then you must enter the routes as static routes. The static routes that correspond to the network IDs available across the interface are entered manually or automatically. The automatic entering of static routes for demand-dial interfaces is known as auto-static updates and is supported by the Windows Server 2003 Routing and Remote Access service. Auto-static updates are supported when you use RIP for IP, but not OSPF.

When instructed, a demand-dial interface that is configured for auto-static updates sends a request across an active connection to request all of the routes of the router on the other side of the connection. In response to the request, all of the routes of the requested router are automatically entered as static routes in the routing table of the requesting router. The static routes are persistent; they are kept in the routing table even if the interface becomes disconnected or the router is restarted. An auto-static update is a one-time, one-way exchange of routing information.

You can automate and schedule auto-static updates by executing the update as a Windows Server 2003 scheduled task. For more information, see the topic titled "Scheduling auto-static updates" in Windows Server 2003 Help and Support.

| **Note** |
| --- |
| The "auto" in auto-static refers to the automatic adding of the requested routes as static routes in the routing table. The sending of the request for routes is performed through an explicit action: either through the Routing and Remote Access snap-in or the Netsh utility while the demand-dial interface is in a connected state. Auto-static updates are not automatically performed every time a demand-dial connection is made. |

**Introduction to Site-to-Site VPN connections**

A site-to-site VPN connection is a demand-dial connection that uses a VPN tunneling protocol such as PPTP or L2TP/IPSec to connect two portions of a private network. Each VPN router provides a routed connection to the network to which the VPN router is attached. On a site-to-site VPN connection, the packets sent from either router across the VPN connection typically do not originate at the routers.

The calling router (the VPN client) initiates the connection. The answering router (the VPN server) listens for connection attempts, receives the connection attempt from the calling router, and responds to the request to create a connection. The calling router authenticates itself to the answering router. When using a mutual authentication protocol such as MS-CHAP v2 or EAP-TLS, the answering router also authenticates itself to the calling router.

Table 1 lists the site-to-site VPN-capable Microsoft operating systems.

**Table 1 Site-to-Site VPN-Capable Microsoft Operating Systems**

| VPN Tunneling Protocol | Microsoft Operating System |
|---|---|
| PPTP | Windows Server 2003, Windows 2000 Server, Windows NT version 4.0 with the Routing and Remote Access Service (RRAS) |
| L2TP/IPSec | Windows Server 2003, Windows 2000 Server |

VPN routers can also be any computer that is capable of creating a routed PPTP connection using MPPE or a routed L2TP connection using IPSec encryption.

**On-demand vs. Persistent Connections**

A site-to-site VPN connection can be on-demand or persistent:

- An on-demand site-to-site connection is a connection that is made when traffic must be forwarded across the connection. The connection is made, the traffic is forwarded, and the connection is terminated after a configured amount of idle time. You can configure idle disconnect behavior for the answering router by setting an idle disconnect on the **Dial-in Constraints** tab on the profile properties of the remote access policy that is used for the site-to-site VPN connection. You can configure idle disconnect behavior for the calling router on the **Options** tab on the properties of the demand-dial interface in the Routing and Remote Access snap-in.

- A persistent site-to-site connection is always connected. If the connection is dropped, it is immediately retried. To configure the answering router for connection persistence, clear the **Minutes server can remain idle before it is connected** and **Minutes client can be connected** check boxes on the **Dial-in Constraints** tab on the profile properties of the remote access policy that is used for the site-to-site VPN connection (these settings are disabled by default). To configure the calling router for connection persistence, select **Persistent connection** on the **Options** tab from the properties of the demand-dial interface.

If the calling router connects to the Internet by using a dial-up link such as an analog phone line or ISDN, then you need to configure a dial-up on-demand site-to-site VPN connection consisting of a single demand-dial interface at the answering router and two demand-dial interfaces at the calling router: one to connect to a local Internet service provider (ISP) and one for the site-to-site VPN connection. Dial-up on-demand site-to-site VPN connections also require an additional host route in the IP routing table of the calling router. For more information, see the topic titled "A dial-up router-to-router VPN connection" in Windows Server 2003 Help and Support.

For either on-demand or persistent site-to-site VPN connections, the answering router is permanently connected to the Internet.

**Restricting the Initiation of Demand-dial Connections**

To prevent the calling router from making unnecessary connections, you can restrict the calling router from making on-demand site-to-site VPN connections in the following ways:

- **Demand-dial filtering**. You can use demand-dial filtering to configure either the types of IP traffic that do not cause a demand-dial connection to be made or the types of IP traffic that cause a connection to be made. You can configure demand-dial filtering by right-clicking the demand-dial interface in the **Network Interfaces** node in the Routing and Remote Access snap-in, and then clicking **Set IP Demand-dial Filters**.

- **Dial-out hours**. You can use dial-out hours to configure the hours that a calling router is either permitted or denied to make a site-to-site VPN connection. You can configure dial-out hours by right-clicking the demand-dial interface in the **Network Interfaces** node in the Routing and Remote Access snap-in, and then clicking **Dial-out Hours**.

You can use remote access policies to configure the times when incoming demand-dial routing connections are allowed.

**One-way vs. Two-way Initiated Connections**

With one-way initiated connections, one VPN router is always the calling router and one VPN router is the always the answering router. One-way initiated connections are well suited to a permanent connection spoke-and-hub topology where the branch office router is the only router that initiates the connection. One-way initiated connections require the following:

- The answering router is configured as a LAN and demand-dial router.

- A user account is added for the authentication credentials of the calling router that is accessed and validated by the answering router.

- A demand-dial interface is configured at the answering router with the same name as the user account that is used by the calling router. This demand-dial interface is not used to dial out, therefore it does is not configured with the host name or IP address of the calling router or with valid user credentials.

With two-way initiated connections, either VPN router can be the calling router or answering router depending on who is initiating the connection. Both VPN routers must be configured to both initiate and accept a site-to-site VPN connection. You can use two-way initiated connections when the site-to-site VPN connection is not active 24 hours a day and traffic from either router is used to create an on-demand connection. Two-way initiated site-to-site VPN connections require the following:

- Both routers are connected to the Internet by using a permanent WAN link.

- Both routers are configured as LAN and demand-dial routers.

- User accounts are added for both routers so that the authentication credentials for the calling router are accessed and validated by the answering router.

- Demand-dial interfaces, with the same name as the user account that is used by the calling router, must be fully configured at both routers, including settings for the host name or IP address of the answering router and user account credentials.

Table 2 lists a correct example configuration for two-way initiated demand-dial routing between Router 1, a demand-dial router in the Seattle site, and Router 2, a demand-dial router in the New York site.

**Table 2 Correct example configuration for two-way initiated demand-dial routing**

| Router | Demand-dial interface name | User account name in user credentials |
|---|---|---|
| Router 1 | DD_NewYork | DD_Seattle |
| Router 2 | DD_Seattle | DD_NewYork |

Notice how the user account name in the user credentials of the demand-dial interface of one router matches the name of a demand-dial interface on the other router.

# Components of Windows Server 2003 Site-to-Site VPNs

Updated: October 7, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Figure 2 shows the components of Windows Server 2003 site-to-site virtual private networks.




**Figure 2: Components of Windows Server 2003 site-to-site VPNs**

The major components are:

- VPN routers
- Internet infrastructure
- Intranet infrastructure
- Authentication, authorization, and accounting (AAA) infrastructure
- Certificate infrastructure

**VPN Routers**

VPN routers either initiate or receive VPN-based demand-dial connections and consist of the following components:

- **Routing and Remote Access service.** The Routing and Remote Access service on both the calling and answering router is configured using the Routing and Remote Access Server Setup Wizard.
- **Ports**. A port is a logical or physical communications channel capable of supporting a single PPP connection. Physical ports are based on equipment installed in the VPN router. Virtual private network (VPN) ports are logical ports.
- **Demand-dial interfaces**. A demand-dial interface configured on the calling router represents the PPP connection and contains configuration information such as the type of port to use, the addressing used to create the connection (an IP address or domain name), authentication methods, encryption requirements, and authentication credentials.

  For two-way initiated connections, a demand-dial interface configured on the answering router represents the PPP connection to the calling router. For a one-way initiated connection using static routes on the user account of the calling router, a demand-dial interface on the answering router does not need to be configured.
- **User account.** To authenticate the calling router, the credentials of the calling router must be verified by the properties

of a corresponding user account. If the answering router is configured for Windows authentication, a user account in the authentication credentials of the calling router must be verifiable using Windows security. If the answering router is configured for RADIUS authentication, then the RADIUS server must have access to the user account for the authentication credentials of the calling router.

The user account must have the following settings:

- On the **Dial-in** tab, remote access permission is set to either **Allow access** or **Control access through Remote Access Policy**. When you create user accounts with the Demand-Dial Interface Wizard, the remote access permission is set to **Allow access**.

- On the **General** or **Account** tab, **User must change password at next logon** is disabled and **Password never expires** is enabled. These settings are configured when you create user accounts with the Demand-Dial Interface Wizard.

For a one-way initiated connection, you can configure static IP routes on the Dial-in tab that are added to the answering router's routing table when the demand-dial connection is made.

- **Routes**. To forward traffic across a site-to-site VPN connection, IP routes in the routing tables of the VPN routers is configured to use the correct demand-dial interface.

  For one-way initiated connections, configure the calling router normally. For the answering router, you can configure the user account specified in the authentication credentials of the calling router with static IP routes.

- **Remote access policy**. On the answering router or the Internet Authentication Service (IAS) server that is acting as a RADIUS server to the answering router, to specify connection parameters that are specific to demand-dial connections, create a separate remote access policy that uses the Windows-Groups attribute set to the group that has all of the user accounts for calling routers as members. A separate remote access policy for demand-dial connections is not required.

A calling router does the following:

- Initiates VPN connections based on an administrator action or when a packet being forwarded matches a route using a VPN-based demand-dial interface.

- Waits for authentication and authorization before forwarding packets.

- Acts as a router forwarding packets between nodes in its site and the answering router.

- Acts as an endpoint of the VPN connection.

The answering router does the following:

- Listens for VPN connection attempts.

- Authenticates and authorizes VPN connections before allowing data to flow.

- Acts as a router forwarding packets between nodes in its site and the calling router.

- Acts as an endpoint of the VPN connection.

VPN routers typically have two installed network adapters: one network adapter connected to the Internet and one network adapter connected to the intranet.

When you configure and enable the Routing and Remote Access service, the Routing and Remote Access Server Setup Wizard prompts you to select the role that the computer will fulfill. For VPN routers, you should select the **Remote access (dial-up or VPN)** option. With the **Remote access (dial-up or VPN)** option, the Routing and Remote Access server operates in the role of a VPN server that supports both remote access and site-to-site VPN connections. For remote access VPN connections, users

run VPN client software and initiate a remote access connection to the VPN server. For site-to-site VPN connections, a router initiates a VPN connection to the VPN server. Alternately, the VPN server can initiate a VPN connection to another VPN router.

| **Note** |
| --- |
| Microsoft recommends the choice of **Remote access (dial-up or VPN)** over **Secure connection between two private networks** in the Routing and Remote Access Server Setup Wizard because the **Secure connection between two private networks** option does not prompt you to select an Internet interface over which to automatically configure packet filters, does not prompt you to configure RADIUS servers, and only creates 5 PPTP and 5 L2TP ports. |

When you select the **Remote access (dial-up or VPN)** option in the Routing and Remote Access Server Setup Wizard:

1. You are first prompted to specify whether VPN, dial-up, or both types of access are needed.

2. Next, you are prompted to select the interface that is connected to the Internet. The interface that you select will be automatically configured with packet filters that allow only PPTP and L2TP-related traffic (unless you clear the **Enable security on the selected interface by setting up static packet filters** check box). All other traffic is silently discarded. For example, you will no longer be able to ping the Internet interface of the VPN server. If you want to use the VPN server computer as a network address translator (NAT), Web server, or other function, see Appendix B.

3. Next, if you have multiple network adapters that are connected to the intranet, you are prompted to select an interface over which DHCP, DNS, and WINS configuration is obtained.

4. Next, you are prompted to determine whether you want to obtain IP addresses to assign to remote access clients by using either DHCP or a specified range of addresses. If you select a specified range of addresses, you are prompted to add one or more address ranges.

5. Next, you are prompted to specify whether you want to use RADIUS as your authentication provider. If you select RADIUS, you are prompted to configure primary and alternate RADIUS servers and the shared secret.

When you select the **Remote access (dial-up or VPN)** option in the Routing and Remote Access Server Setup Wizard, the results are as follows:

1. The Routing and Remote Access service is enabled as both a remote access server and a LAN and demand-dial router, with Windows as the authentication and accounting provider (unless RADIUS was chosen and configured). If there is only one network adapter connected to the intranet, that network adapter is automatically selected as the IP interface from which to obtain DHCP, DNS, and WINS configuration. Otherwise, the network adapter specified in the wizard is selected to obtain DHCP, DNS, and WINS configuration. If specified, the static IP address ranges are configured.

2. Exactly 128 PPTP and 128 L2TP ports are created. All of them are enabled for both inbound remote access connections and inbound and outbound demand-dial connections.

3. The selected Internet interface is configured with input and output IP packet filters that allow only PPTP and L2TP/IPSec traffic.

4. The DHCP Relay Agent component is added with the **Internal** interface. If the VPN server is a DHCP client at the time the wizard is run, the DHCP Relay Agent is automatically configured with the IP address of a DHCP server. Otherwise, you must manually configure the properties of the DHCP Relay Agent with an IP address of a DHCP server on your intranet. The DHCP Relay Agent forwards DHCPInform packets between VPN remote access clients and an intranet DHCP server.

5. The IGMP component is added. The **Internal** interface is configured for IGMP router mode. All other LAN interfaces are configured for IGMP proxy mode. This allows VPN remote access clients to send and receive IP multicast traffic.

With Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Installing a Certificate on a VPN Router**
If VPN routers are making L2TP/IPSec connections or using EAP-TLS authentication, certificates must be installed on the VPN router computers. For L2TP/IPSec connections, a computer certificate must be installed on both the calling and answering router computer to provide authentication for establishing an IPSec session. For EAP-TLS authentication, a computer certificate must be installed on the authenticating server (either the answering router or a RADIUS server) and a user certificate must be installed on the calling router.

For more information about installing certificates on calling routers, answering routers, and authentication server computers, see Certificate infrastructure in this paper.

**Design Points: Configuring the VPN Router**
Consider the following before running the Routing and Remote Access Server Setup Wizard:

- Which connection of the VPN router is connected to the Internet?

  Typical Internet-connected VPN routers have at least two LAN connections: one connected to the Internet (either directly or connected to a perimeter network) and one connected to the site. To make this distinction easier to see for the Routing and Remote Access Server Setup Wizard, rename the connections with their purpose or role using Network Connections. For example, rename the connection connected to the Internet, default name **Local Area Connection 2**, to **Internet**.

- Can the VPN router be a DHCP client?

  The VPN router must have a manual TCP/IP configuration for its Internet interface. While technically possible, it is not recommended that the VPN router be a DHCP client for its intranet interface(s). Due to the routing requirements of the VPN router, manually configure an IP address, subnet mask, DNS server(s), and WINS server(s), but do not configure a default gateway.

  Note that it is possible for the VPN router to have a manual TCP/IP configuration and still use DHCP to obtain IP addresses for remote access VPN clients and other calling routers.

- How will IP addresses be allocated to remote access VPN clients and other calling routers?

  The VPN router can be configured to obtain IP addresses from DHCP or from a manually configured set of address ranges. Using DHCP to obtain IP addresses simplifies the configuration, however, you must ensure that the DHCP scope for the subnet to which the site connection of the calling router is attached has enough addresses for all the computers physically connected to the subnet and the maximum number of PPTP and L2TP ports. For example, if the subnet to which the site connection of the VPN router is attached contains 50 DHCP clients, then, for the default configuration of the VPN router, the scope should contain at least 307 addresses (50 computers + 128 PPTP clients + 128 L2TP clients + 1 address for the VPN router). If there are not enough IP addresses in the scope, remote access VPN clients and calling routers that connect after all the addresses in the scope are allocated will be assigned an address in the Automatic Private IP Addressing (APIPA) range of 169.254.0.0/16.

  If you configure a static pool of addresses, ensure that the pool has enough addresses for all your PPTP and L2TP ports, plus an additional address for the VPN router. If there are not enough addresses in your static pool, remote access VPN clients and Windows NT 4.0 RRAS calling routers will not be able to connect. Windows Server 2003 calling routers, however, will still be able to connect. Windows Server 2003 calling and answering routers still request an IP address from each other during the connection establishment process. But if one of the routers does not have an address to assign, both routers continue with the connection establishment process. The logical interface on the point-to-point connection does not have an assigned IP address. This is known as an unnumbered connection. While Windows Server 2003 VPN routers support unnumbered connections, the routing protocols included with Windows Server 2003 do not work over an unnumbered connection.

  If you are configuring a static pool of addresses, there might be additional routing considerations. For more information, see Site network infrastructure in this paper.

- What is the authentication and accounting provider?

  The VPN router can use Windows or RADIUS as its authentication or accounting provider.

  When Windows is used as the authentication and accounting provider, the VPN router uses Windows security to validate

the credentials of a calling router and access the calling router's user account dial-in properties. Locally configured remote access policies authorize the VPN connection and locally written accounting log files log VPN connection accounting information.

When RADIUS is used as the authentication and accounting provider, the VPN router uses a configured RADIUS server to validate the credentials of a calling router, authorize the connection attempt, and store VPN connection accounting information.

- Are you making L2TP/IPSec connections?

  If so, you must install a computer certificate on both the calling router and answering router computers.

- Are you using user-level certificate authentication with EAP-TLS?

  If so, you must install a user certificate on the calling router computer and a computer certificate on the authenticating server (either the answering router computer [if the answering router is configured for the Windows authentication provider] or the RADIUS server [if the answering router is configured for the RADIUS authentication provider]). If the authenticating server is a Windows Server 2003 VPN router or a Windows Server 2003 Internet Authentication Service (IAS) server, EAP-TLS is only available if the authenticating server is a member of an Active Directory service domain.

- For on-demand connections, do you want to prevent connections from occurring during certain times of the day during the week or for certain types of traffic?

  If so, configure dial-out hours or demand-dial filters on the demand-dial interface of the calling router.

- Do you want to match your IP packet filters to the demand-dial filters?

  Demand-dial filters are applied before the connection is made. IP packet filters are applied after the connection is made. To prevent the demand-dial connection from being established for traffic that is discarded by the IP packet filters:

  - If you have configured a set of output IP packet filters with the **Transmit all packets except those that meet the criteria below** option, then configure the same set of filters as demand-dial filters with **Initiate connection** set to **For all traffic except**.

  - If you have configured a set of output IP packet filters with the **Drop all packets except those that meet the criteria below** option, then configure the same set of filters as demand-dial filters with **Initiate connection** set to **Only for the following traffic**.

Consider the following when changing the default configuration of the VPN router for site-to-site VPN connections:

- Do you want to support remote access VPN connections?

  By default, all the PPTP and L2TP ports are configured to allow both remote access connections (inbound only) and demand-dial routing connections (inbound and outbound). To disable remote access connections and create a dedicated site-to-site VPN connection server, clear the **Remote access connections (inbound only)** check box from the properties of the **WAN miniport (PPTP)** and **WAN miniport (L2TP)** devices from the properties of the **Ports** object in the Routing and Remote Access snap-in. Alternately, you can clear the **Remote access server** checkbox on the **General** tab on the properties of the VPN server.

- Do you need to install a computer certificate?

  If the VPN router is supporting L2TP/IPSec connections or is authenticating connections using the EAP-TLS authentication protocol and configured to use the Windows authentication provider, you must install a computer certificate. If the VPN router is a calling router using the EAP-TLS authentication protocol, you must install a user certificate. For more information, see "Certificate infrastructure" in this paper.

- Do you need custom remote access policies for VPN connections?

If you configure the VPN router for Windows authentication or for RADIUS authentication and the RADIUS server is an IAS server, the default remote access policies reject all types of connection attempts unless the remote access permission of the user accounts dial-in properties is set to **Allow access**. If you want to manage authorization and connection parameters by group or by type of connection, you must configure additional remote access policies. For more information, see Remote Access Policies in this paper.

- Do you want separate authentication and accounting providers?

  The Routing and Remote Access Server Setup Wizard configures both authentication and accounting providers to be the same. After the Wizard is complete, however, you can configure the authentication and accounting providers separately (for example, if you want to use Windows authentication and RADIUS accounting). You can configure authentication and accounting providers on the **Authentication** tab from the properties of the VPN router in the Routing and Remote Access snap-in.

After the VPN router is configured, you can begin creating demand-dial interfaces and configuring routes using the Routing and Remote Access snap-in. For more information, see "Deploying a PPTP-based Site-to-Site VPN Connection" and "Deploying an L2TP/IPSec-based Site-to-Site VPN Connection" in this paper.

**Internet Network Infrastructure**
To create a site-to-site VPN connection to an answering router across the Internet:

- The answering router's name must be resolvable.
- The answering router must be reachable.
- VPN traffic must be allowed to and from the answering router.

**Answering Router Name Resolvability**
While it is possible to configure demand-dial interfaces with the names of the answering routers to which a connection is made, it is recommended that you use IP addresses rather than names. If you use a name and the name resolves to the public IP address of the answering router, traffic sent to services running on the VPN router are sent in clear text across the Internet.

**Answering Router Reachability**
To be reachable, the answering router must be assigned a public IP address to which packets are forwarded by the routing infrastructure of the Internet. If you have been assigned a static public IP address from an ISP or an Internet registry, this is typically not an issue. In some configurations, the answering router is actually configured with a private IP address and has a published static IP address by which it is known on the Internet. A device between the Internet and the answering router translates the published and actual IP addresses of the answering router in packets to and from the answering router.

While the routing infrastructure might be in place, the answering router might be unreachable due to the placement of firewalls, packet filtering routers, network address translators, security gateways, or other types of devices that prevent packets from either being sent to or received from the answering router computer.

**VPN Routers and Firewall Configuration**
There are two approaches to using a firewall with a VPN router:

1. The VPN router is attached directly to the Internet and the firewall is between the VPN router and the site.

   In this configuration, the VPN router must be configured with packet filters that only allow VPN traffic in and out of its Internet interface. The firewall can be configured to allow specific types of inter-site traffic.

2. The firewall is attached to the Internet and the VPN router is between the firewall and the site.

   In this configuration, both the firewall and the VPN router are attached to a network segment known as the perimeter network (also known as a screened subnet). Both the firewall and the VPN router must be configured with packet filters that allow only VPN traffic to and from the Internet. Figure 2 shows this configuration.

For the details of configuring packet filters for the VPN router and the firewall for both of these configurations, see Appendix A.

**Design Points: Answering Router Accessibility from the Internet**

Consider the following when configuring your Internet infrastructure for site-to-site VPN connections:

- Wherever possible, configure your demand-dial interfaces with the IP addresses of answering routers. If you are using names, ensure that the DNS names of your answering routers are resolvable by either placing an appropriate DNS record in your Internet DNS server or the DNS server of your ISP. Test the resolvability by using the Ping tool to ping the name of each of your answering routers. Due to packet filtering, the result of the ping command may be "Request timed out", but check to ensure that the name specified was resolved by the Ping tool to the correct IP address.

- Ensure that the IP addresses of your answering routers are reachable from the Internet by using the Ping tool to ping the name or address of your answering router with a 5 second timeout (using the -w command line option) when directly connected to the Internet. If you see a "Destination unreachable" error message, the answering router is not reachable.

- Configure packet filtering for PPTP traffic, L2TP/IPSec traffic, or both types of traffic on the appropriate firewall and answering router interfaces connecting to the Internet and the perimeter network. The correct set of packet filters is automatically configured by the Routing and Remote Access Server Setup Wizard when you select the **Remote access (dial-up or VPN)** configuration. For more information, see Appendix A.

**Authentication Protocols**

To authenticate the calling router who is attempting to create a PPP connection, Windows Server 2003 supports a wide variety of PPP authentication protocols including:

- Password Authentication Protocol (PAP)

- Shiva Password Authentication Protocol (SPAP)

- Challenge-Handshake Authentication Protocol (CHAP)

- Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)

- MS-CHAP version 2 (MS-CHAP v2)

- Extensible Authentication Protocol-Message Digest 5 (EAP-MD5)

- Extensible Authentication Protocol-Transport Level Protocol (EAP-TLS)

For PPTP connections, you must use MS-CHAP, MS-CHAP v2, or EAP-TLS. Only these three authentication protocols provide a mechanism to generate the same encryption key on both the calling router and the answering router. MPPE uses this encryption key to encrypt all PPTP data sent on the VPN connection. MS-CHAP and MS-CHAP v2 are password-based authentication protocols.

In the absence of user certificates, MS-CHAP v2 is highly recommended, as it is a stronger authentication protocol than MS-CHAP and provides mutual authentication. With mutual authentication, the answering router authenticates the calling router and the calling router authenticates the answering router.

| **Note** |
| --- |
| If you must use a password-based authentication protocol, enforce the use of strong passwords on your network. Strong passwords are long (greater than 8 characters) and contain a random mixture of upper and lower case letters, numbers, and punctuation. An example of a strong password is f3L*02~>xR3w#4o. |

EAP-TLS is used in conjunction with a certificate infrastructure and user certificates. With EAP-TLS, the calling router sends a user certificate for authentication and the authenticating server (the answering router or RADIUS server) sends a computer certificate for authentication. This is the strongest authentication method, as it does not rely on passwords. If the authenticating server is a Windows Server 2003 VPN router or an IAS server, EAP-TLS is only available if the authenticating server is a member of an Active Directory domain.

| **Note** |
| --- |

You can use third-party CAs for EAP-TLS certificates. For more information, see Appendix E in the white paper titled "Virtual Private Networking with Windows Server 2003: Deploying Remote Access VPNs."

For L2TP/IPSec connections, any PPP authentication protocol can be used because the user authentication occurs after the calling router and answering router have established a secure channel of communication known as an IPSec security association (SA). However, the use of either MS-CHAP v2 or EAP-TLS is recommended.

**Design Point: Which Authentication Protocol to Use?**
Consider the following when choosing an authentication protocol for VPN connections:

- If you are using a certificate infrastructure that issues user certificates, use the EAP-TLS authentication protocol for both PPTP and L2TP connections. Windows NT 4.0 RRAS routers do not support EAP-TLS.

- If you must use a password-based authentication protocol, use MS-CHAP v2 and enforce strong passwords using Group Policy. MS-CHAP v2 is supported by computers running Windows Server 2003, Windows 2000, and Windows NT 4.0 with RRAS and Service Pack 4 and later.

**VPN Protocols**
Windows Server 2003 includes support for two PPP-based site-to-site VPN protocols:

1. Point-to-Point Tunneling Protocol

2. Layer Two Tunneling Protocol with IPSec

**Point-to-Point Tunneling Protocol**
Introduced in Windows NT 4.0, PPTP leverages Point-to-Point Protocol (PPP) user authentication and Microsoft Point-to-Point Encryption (MPPE) to encapsulate and encrypt IP traffic. When MS-CHAP v2 is used with strong passwords, PPTP is a secure VPN technology. For non-password-based authentication, EAP-TLS can be used in Windows Server 2003 to support user certificates. PPTP is widely supported, easily deployed, and can be used across most network address translators (NATs).

**Layer Two Tunneling Protocol With IPSec**
L2TP leverages PPP user authentication and IPSec Encapsulating Security Payload (ESP) transport mode to encapsulate and encrypt IP traffic. This combination, known as L2TP/IPSec, uses certificate-based computer identity authentication to create the IPSec security association in addition to PPP-based user authentication. L2TP/IPSec provides data integrity and data authentication for each packet. However, L2TP/IPSec requires a certificate infrastructure to allocate computer certificates and is supported by Windows Server 2003 VPN routers and other third-party VPN routers.

**Design Point: PPTP or L2TP?**
Consider the following when deciding between PPTP and L2TP for site-to-site VPN connections:

- PPTP can be used with Windows Server 2003, Windows 2000, and Windows NT version 4.0 with RRAS. PPTP does not require a certificate infrastructure to issue computer certificates.

- PPTP-based VPN connections provide data confidentiality (captured packets cannot be interpreted without the encryption key). PPTP VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data authentication (proof that the data was sent by the authorized computer).

- PPTP-based calling routers can be located behind a NAT because most NATs include a NAT editor that knows how to properly translate PPTP tunneled data. For example, the Internet connection sharing (ICS) feature of Dial-up Connections and the NAT/Basic Firewall routing protocol component of the Windows Server 2003 Routing and Remote Access service include a NAT editor that translates PPTP traffic from PPTP clients located behind the NAT. Answering routers cannot be behind a NAT unless there are multiple public IP addresses and there is a one-to-one mapping of a public IP address to the private IP address of the answering router or, if there is only one public address, if the NAT is configured to translate and forward the PPTP tunneled data to the VPN router. Most NATs using a single public IP address, including ICS and the NAT routing protocol component, can be configured to allow inbound traffic based on IP addresses and TCP and UDP ports. However, PPTP tunneled data does not use TCP or UDP headers. Therefore, an answering router cannot be located behind a computer using ICS or the NAT routing protocol component when using a single IP address.

- L2TP/IPSec-based VPN routers cannot be behind a NAT unless both the calling and answering routers support IPSec NAT

Traversal (NAT-T). IPSec NAT-T for site-to-site VPN connections is only supported by Windows Server 2003. However, Microsoft recommends that answering routers not be placed behind NATs. For more information, see IPSec NAT-T is not recommended for Windows Server 2003 computers that are behind network address translators at http://support.microsoft.com/default.aspx?scid=kb;en-us;885348.

- L2TP/IPSec can only be used with Windows Server 2003, Windows 2000, and third-party VPN routers and supports computer certificates as the default authentication method for IPSec. Computer certificate authentication requires a certificate infrastructure to issue computer certificates to the answering router computer and all calling router computers.

- By using IPSec, L2TP-based VPN connections provide data confidentiality, data integrity, data authentication, and replay protection.

- PPTP and L2TP is not an either/or choice. By default, a Windows Server 2003 VPN router supports both PPTP and L2TP connections simultaneously. You can use PPTP for some site-to-site VPN connections (from calling routers that are running Windows Server 2003, Windows 2000, or Windows NT 4.0 with RRAS and do not have an installed computer certificate) and L2TP for other site-to-site VPN connections (from calling routers running Windows Server 2003 or Windows 2000 that have an installed computer certificate).

- If you are using both PPTP and L2TP, you can create separate remote access policies that define different connection parameters for PPTP and L2TP connections.

**Site Network Infrastructure**
The network infrastructure of the site is an important element of VPN design. Without proper design, calling routers are unable to obtain proper IP addresses, and packets cannot be exchange between computers in the different sites.

**Name Resolution**
If the calling router is configured with the IP addresses of Domain Name System (DNS) or Windows Internet Name Service (WINS) servers, DNS and WINS server IP addresses are not requested from the answering router during the PPP connection negotiation. If the calling router is not configured with the IP addresses of DNS and WINS servers, DNS and WINS servers are requested. The answering router never requests DNS and WINS server IP addresses from the calling router.

Unlike Windows Server 2003, Windows 2000, and Windows XP remote access clients, the calling router does not send a DHCPInform message to the answering router to discover additional TCP/IP configuration information.

By default, the calling router does not register itself with the DNS or WINS servers of the answering router. To change this behavior, set the registry value HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Rasman \PPP\ControlProtocols\BuiltIn\RegisterRoutersWithNameServers to 1.

**Routing**
Each VPN router is an IP router and as such must be properly configured with the set of routes that makes all locations reachable. Specifically, each VPN router needs the following:

- A default route that points to a firewall or router directly connected to the Internet.

  This route makes all of the locations on the Internet reachable.

- One or more routes that summarize the addresses used within the site of the VPN router that point to a neighboring site router.

  These routes make all of the locations within the site of the VPN router reachable from the VPN router. Without these routes, all hosts in the site not connected to the same subnet as the VPN router are unreachable.

To add a default route that points to the Internet, configure the Internet interface with a default gateway and then manually configure the site interface without a default gateway.

To add site routes to the routing table of each VPN router, you can:

- Add static routes using the Routing and Remote Access snap-in. You do not necessarily have to add a route for each subnet in your site. At a minimum, you just need to add the routes that summarize all the possible addresses in your site. For example, if your site uses portions of the private address space 10.0.0.0/8 to number its subnets and hosts, you do

not have to add a route for each subnet. Just add a route for 10.0.0.0 with the subnet mask 255.0.0.0 that points to a neighboring router on the site subnet to which your VPN router is attached.

- If you are using the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) routing protocols in your site, you can add and configure the RIP or OSPF routing protocol components of the Routing and Remote Access service so that the VPN router participates in the propagation of routing information as a dynamic router.

If your site has only a single subnet, no further configuration is required.

When a site-to-site VPN connection is made, each router sends traffic using a logical interface that corresponds to the PPTP or L2TP port of the connection. During the PPP negotiation, IP addresses might be assigned to these logical interfaces. Ensuring the reachability of the logical interfaces of VPN routers depends on how you configure each VPN router to obtain IP addresses for remote access clients and calling routers. The IP addresses assigned to VPN routers as they connect can be from:

- An on-subnet address range, which is an address range of the site subnet to which the VPN router is attached.

  An on-subnet address range is used whenever the VPN router is configured to use DHCP to obtain IP addresses and when the manually configured pool(s) of IP addresses are within the range of addresses of the attached subnet.

- An off-subnet address range, which is an address range that represents a different subnet that is logically attached to the VPN router.

  An off-subnet address range is used whenever the VPN router is manually configured with a pool of IP addresses for a separate subnet.

**On-subnetAddressRange**
If you are using an on-subnet address range, no additional routing configuration is required as the VPN router acts as a proxy for all packets destined to the logical interfaces of the other connected VPN routers. Routers and hosts on the VPN router subnet forward packets destined to the logical interfaces of connected VPN routers to the VPN router and the VPN router relays them the appropriate connected VPN routers.

**Off-subnetAddressRange**
If you are using an off-subnet address range, you must add the route(s) that summarize the off-subnet address range to the site routing infrastructure so that traffic destined to the logical interfaces of connected VPN routers are forwarded to the VPN router and then sent by the VPN router to the appropriate connected VPN router. To provide the best summarization of address ranges for routes, choose address ranges that can be expressed using a single prefix and subnet mask. For more information, see the topic titled "Expressing an IP address range with a mask" in Windows Server 2003 Help and Support.

You can add the routes that summarize the off-subnet address range to the routing infrastructure of the site through the following:

- Add static routes to the neighboring router for the off-subnet address range that point to the VPN routers site interface. Configure the neighboring router to propagate this static route to other routers in the site using the dynamic routing protocol used in your site.

- If the VPN router is using OSPF and participating as a dynamic router, the VPN router must be configured as an autonomous system boundary router (ASBR) so that the static routes of the off-subnet address range are propagated to the other OSPF routers in the site.

If your site consists of a single subnet, then you must either configure each site host for persistent route(s) of the off-subnet address range that point to the VPN routers site interface or configure each site host with the VPN router as its default gateway. Because routing for off-subnet address ranges requires additional host configuration, it is recommended that you use an on-subnet address pool for a small office/home office (SOHO) network consisting of a single subnet.

**Design Points: Routing Infrastructure**
Consider the following when configuring the routing infrastructure for site-to-site VPN connections:

- Configure the Internet interface of the VPN router with a default gateway. Do not configure the site interface of the VPN router with a default gateway.

- Add static IP routes to the VPN router that summarize the addresses used in the site in which the VPN router is located. Alternately, if you use either RIP or OSPF as your dynamic routing protocol, configure and enable RIP or OSPF on the VPN router. If you use a routing protocol other than RIP or OSPF, such as Interior Gateway Routing Protocol (IGRP), configure the neighboring router for RIP or OSPF on the interface connected to the subnet containing the VPN router, and then configure IGRP on all other interfaces.

- Configure the VPN router with an on-subnet address range by obtaining IP addresses through DHCP or by manually configuring on-subnet address pools.

**AAA Infrastructure**

The authentication, authorization, and accounting (AAA) infrastructure exists to:

- Authenticate the credentials of calling routers.

- Authorize the VPN connection.

- Record the VPN connection creation and termination for accounting purposes.

The AAA infrastructure consists of:

- The answering router computer

- RADIUS server computers

- Domain controllers

As previously discussed, a Windows Server 2003 answering router can be configured with either Windows or RADIUS as its authentication or accounting provider. RADIUS provides a centralized AAA service when you have multiple answering routers and remote access VPN servers or a mix of heterogeneous dial-up and VPN equipment.

When you configure Windows as the authentication provider, the answering router performs the authentication of the VPN connection by communicating with a domain controller using a secure remote procedure call (RPC) channel and authorization of the connection attempt through the dial-in properties of the user account and locally configured remote access policies.

When you configure RADIUS as the authentication provider, the answering router relies on a RADIUS server to perform both the authentication and authorization. When a VPN connection is attempted, the answering router sends the calling router credentials and other connection parameters to the configured RADIUS server in a RADIUS Access-Request message. If the connection attempt is both authenticated and authorized, the RADIUS server sends back a RADIUS Access-Accept message. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends back a RADIUS Access-Reject message.

When you configure Windows as the accounting provider, the answering router logs VPN connection information in a local log file (SystemRoot\System32\Logfiles\Logfile.log by default) based on settings on the **Settings** tab of the properties of the **Local File** object in the **Remote Access Logging** folder in the Routing and Remote Access snap-in. By default, all types of logging are disabled. Windows Server 2003 also supports the sending of connection accounting information to a Structured Query Language (SQL) server.

When you configure RADIUS as the authentication provider, the answering router sends RADIUS accounting messages for VPN connections on a RADIUS server, which records the accounting information.

If you are using RADIUS and a Windows domain as the user account database for which to verify user credentials and obtain dial-in properties, Microsoft recommends using the Internet Authentication Service (IAS), included as an optional networking component in Windows Server 2003 and Windows 2000 Server. IAS is a full-featured RADIUS server that is tightly integrated with Active Directory and the Routing and Remote Access service.

When IAS is used as the RADIUS server:

- IAS performs the authentication of the VPN connection by communicating with a domain controller using a secure RPC channel. IAS performs authorization of the connection attempt through the dial-in properties of the user account and remote access policies configured on the IAS server.

- IAS logs all RADIUS accounting information in a local log file (SystemRoot\System32\Logfiles\Logfile.log by default) based on settings configured on the properties of the **Local File** object in the **Remote Access Logging** folder in the Internet Authentication Service snap-in. IAS in Windows Server 2003 also supports the sending of connection accounting information to a SQL server.

IAS for Windows Server 2003 can also be used as a RADIUS proxy. For more information, see Windows Server 2003 Help and Support.

**Remote Access Policies**
Remote access policies are an ordered set of rules that define how connections are either accepted or rejected. For connections that are accepted, remote access policies can also define connection restrictions. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting. Connection attempts are evaluated against the remote access policies in order, trying to determine whether the connection attempt matches all of the conditions of each policy. If the connection attempt does not match all of the conditions of any policy, the connection attempt is rejected.

If a connection matches all the conditions of a remote access policy and is granted remote access permission, the remote access policy profile specifies a set of connection restrictions. The dial-in properties of the user account also provide a set of restrictions. Where applicable, user account connection restrictions override the remote access policy profile connection restrictions.

Remote access policies consist of the following elements:

- Conditions
- Permission
- Profile settings

**Conditions**
Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt. If there are multiple conditions, then all of the conditions must match the settings of the connection attempt in order for it to match the policy. For VPN connections, you commonly use the following conditions:

- **NAS-Port-Type**. By setting the **NAS-Port-Type** condition to **Virtual (VPN)**, you can specify all VPN connections.
- **Tunnel-Type**. By setting the **Tunnel-Type** to **Point-to-Point Tunneling Protocol (PPTP)** or **Layer Two Tunneling Protocol (L2TP)**, you can specify different policies for PPTP and L2TP connections.
- **Windows-Groups**. By setting the **Windows-Groups** to the appropriate groups, you can specify connection parameters based on group membership.

**Permission**
You can use the permission setting to either grant to deny remote access for the connection attempt if the remote access permission of the user account is set to **Control access through Remote Access Policy**. Otherwise, the remote access permission setting on the user account determines the remote access permission.

**Profile Settings**
A remote access policy profile is a set of properties that are applied to a connection when it is authorized. For VPN connections, you can use the following profile settings:

- Dial-in constraints can be used to define how long the connection can exist or be idle before being terminated by the answering router, among others.
- Authentication settings can define which authentication protocols the calling router can use when sending its credentials and the configuration of EAP types, such as EAP-TLS.
- Encryption settings can define whether encryption is required and the encryption strength. For encryption strengths, Windows Server 2003 supports **Basic** (40-bit MPPE for PPTP and 56-bit Data Encryption Standard [DES] for L2TP), **Strong** (56-bit MPPE for PPTP and 56-bit DES for L2TP), or **Strongest** (128-bit MPPE for PPTP and 3DES for L2TP).

For example, you can create a Windows group called VPNRouters whose members are the user accounts of all calling routers. Then, you create a policy using the New Remote Access Policy Wizard that specifies VPN connection using the VPNRouters group. Using the Wizard, you can also select a specific authentication method and encryption strength.

| **Note** |
| --- |
| IP packet filters on the **IP** tab of the profile settings of a remote access policy only apply to remote access VPN connections. They have no effect on demand-dial connections. |

**Windows Domain User Accounts and Groups**

Windows NT version 4.0 domains, mixed-mode Active Directory domains, and native-mode domains contain the user accounts and groups used by the Routing and Remote Access service and IAS to authenticate and authorize VPN connection attempts.

User accounts contain the user name and a form of the users password that can be used for validation of the calling routers user credentials. Additional account properties determine whether the user account is enabled or disabled, locked out, or permitted to logon only during specific hours. If a user account is disabled, locked out, or not permitted to logon during the time of the VPN connection, the site-to-site VPN connection attempt is rejected. Additionally, if the user account of the calling router is configured to change the password at the next login, the site-to-site VPN connection attempt will fail because changing the password while attempting to make the connection is an interactive process. Demand-dial routers need to be able to make connections as needed without requiring human intervention. Therefore, all user accounts for calling routers must be **configured with the User must change password at next logon** checkbox cleared and the **Password never expires** checkbox selected for the account options on the **Account** tab on the properties of the user account. When you create dial-in accounts with the Demand-Dial Interface Wizard, these account settings are automatically configured.

You should use a separate user account for each site that contains a calling router. Each user account should have a name that matches a demand-dial interface configured on the answering router. When you create dial-in accounts with the Demand-Dial Interface Wizard, this one-to-one relationship between user accounts used by calling routers in separate sites and demand-dial interfaces is automatically created.

User accounts also contain dial-in settings. The dial-in setting most relevant for VPN connections is the remote access permission setting, which has the following values:

- Allow access
- Deny access
- Control access through Remote Access Policy

The **Allow access** and **Deny access** settings explicitly allow or deny remote access and are equivalent to the remote access permission setting of Windows NT 4.0 domain accounts. When you use the **Control access through Remote Access Policy** setting, the remote access permission is determined by the remote access permission setting of the matching remote access policy. If the user account is in a mixed-mode domain, the **Control access through Remote Access Policy** setting is not available and you must manage remote access permission on a per-user basis. If the user account is in a native-mode domain, the **Control access through Remote Access Policy** setting is available and you can manage remote access permission on a per-user basis or using groups. When a dial-in account is created with the Demand-Dial Interface Wizard, the remote access permission is set to **Allow access**.

When using groups to manage access, you can use your existing groups and create remote access policies that either allow or reject access or restrict access based on the group name. For example, the Employees group has no VPN remote access restrictions, however, the Contractors group can only create VPN connections during business hours. Alternately, you can create groups based on the type of connection being made. For example, you can create a VPNRouters group and add as members all the user accounts allowed to create VPN connections.

**One-way Initiated Connections and Static Routes on the User Account**

With one-way initiated connections, one router is always the answering router and one router is always the calling router. The answering router accepts the connection and the calling router initiates the connection. One-way initiated connections are well suited to a spoke-and-hub topology where the branch office router is the only router that initiates the connection.

To simplify configuration for one-way initiated connections, user accounts on stand-alone Windows Server 2003 or in a native-mode Active Directory domain support the configuration of static routes. The static routes are automatically added to the routing table of the VPN router when a VPN connection using the user account is made. If the VPN router is participating in

dynamic routing for the site, the routes are propagated to the other routers in the site using routing protocols such as RIP and OSPF. Static routes on user accounts are configured by selecting the **Apply static routes** check box on the **Dial-in** tab on the properties of a user account, and then adding static routes.

To use static routes on the user account, configure the calling router normally. On the answering router, all you have to do is create a user account that is used by the calling router and configure static routes that correspond to the calling router's site. Because there is no demand-dial interface on the answering router with the same name as the user account of the calling router, the incoming VPN connection is determined to be a remote access connection. The static routes of the calling router's user account are added to the VPN router's routing table and all traffic to the locations implied by the static routes is sent across the logical remote access connection to the calling router.

| **Note** |
| --- |
| Static routes on the user account are only applied to the answering router when the incoming connection is a remote access VPN connection (the user name in the credentials of the calling router does not match the name of a demand-dial interface on the answering router). Static routes on the user account are not applied when the incoming connection is a demand-dial connection. |

**Design Points: AAA Infrastructure**

Consider the following when configuring the AAA infrastructure for site-to-site VPN connections:

- If you have multiple VPN routers and you want to centralize AAA service or a heterogeneous mixture of dial-up and VPN equipment, use RADIUS servers and configure the VPN router for the RADIUS authentication and accounting providers.
- If your user account database is a Windows domain, use IAS as your RADIUS server. If you use IAS, install IAS on a domain controller for best performance. Install at least two IAS servers for fail-over and fault tolerance of AAA services.
- Whether configured locally or on an IAS server, use remote access policies to authorize VPN connections and specify connection constraints. For example, use remote access policies to grant access based on group membership, to enforce the use of encryption and a specific encryption strength, or specify the use of EAP-TLS.
- For one-way initiated connections, you can configure the calling router normally and configure the answering router with a user account that contains the static routes of the calling router's site.
- Sensitive fields of RADIUS messages, such as the user password and encryption keys, are encrypted with the RADIUS shared secret configured on the VPN router and the RADIUS server. Make the shared secret a long (22 characters or longer), random sequence of letters, numbers, and punctuation and change it often to protect your RADIUS traffic. An example of a strong shared secret is 8d#>9fq4bV)H7%a3@dW9.>. To further protect RADIUS traffic, use IPSec policies to provide data confidentiality for all traffic using the RADIUS UDP ports (1812 and 1645 for RADIUS authentication traffic and 1813 and 1646 for RADIUS accounting traffic).

**Certificate Infrastructure**

To perform certificate-based authentication for L2TP connections and user certificate-based authentication for site-to-site VPN connections using EAP-TLS, a certificate infrastructure must be in place to issue the proper certificates to submit during the authentication process and to validate the certificate being submitted.

**Computer Certificates for L2TP/IPSec**

If you manually configure the certificate authentication method for a rule of an IPSec policy in Windows Server 2003, you can specify the list of root certification authorities (CAs) from which a certificate is accepted for authentication. For L2TP connections, the IPSec rule for L2TP traffic is automatically configured and the list of root CAs is not configurable. Instead, each computer in the L2TP connection sends a list of root CAs to its IPSec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued certificates that are stored in the computer certificate store. For example, if Computer A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPSec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPSec peer, Computer B, does not have a valid certificate in its computer certificate store issued from either CertAuth1 or CertAuth2, IPSec security negotiation fails.

Ensure one of the following before attempting an L2TP connection:

1. Both the calling router and answering router were issued computer certificates from the same CA.

2. Both the calling router and answering router were issued computer certificates from CAs that follow a valid certificate chain up to the same root CA.

In general, the calling router must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the answering router trusts. Additionally, the answering router must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the calling router trusts.

A single CA commonly issues computer certificates to all computers in an organization. Because of this, all computers within the organization both have computer certificates from a single CA and request certificates for authentication from the same single CA.

For information about installing computer certificates on VPN routers for L2TP connections, see "Deploying an L2TP/IPSec-based Site-to-Site VPN Connection."

| **Note** |
| --- |
| The Windows Server 2003 Routing and Remote Access service supports the configuration of a preshared key for IPSec authentication of L2TP/IPSec connections. To configure the answering router, select **Allow custom IPSec policy for L2TP connection** from the **Security** tab in the properties of a VPN router in the Routing and Remote Access snap-in, and then type the preshared key. To configure the calling router, click **IPSec Settings** on the **Security** tab in the properties of a demand-dial interface, and then type the preshared key. However, preshared key authentication for L2TP/IPSec connections is not secure and is not recommended, except as an interim measure while deploying a certificate infrastructure or to connect to third-party VPN routers that do not support certificate authentication. |

**User and Computer Certificates for EAP-TLS Authentication**

To perform EAP-TLS authentication for a site-to-site VPN connection in Windows Server 2003:

- The calling router must be configured with a user certificate to submit during the EAP-TLS authentication process.
- The authenticating server must be configured with a computer certificate to submit during the EAP-TLS authentication process. The authenticating server is either the answering router (if the answering router is configured to use the Windows authentication provider) or a RADIUS server (if the answering router is configured to use the RADIUS authentication provider).

EAP-TLS authentication is successful when the following conditions are met:

- The calling router submits a valid user certificate that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the answering router trusts.
- The authenticating server submits a valid computer certificate that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the calling router trusts.
- The user certificate of the calling router contains the Client Authentication enhanced key usage (object identifier "1.3.6.1.5.5.7.3.2").
- The computer certificate of the answering router contains the Server Authentication enhanced key usage (object identifier "1.3.6.1.5.5.7.3.1").

For a Windows Server 2003 CA, a Router (Offline request) certificate, a special type of user certificate for demand-dial connections, is created and mapped to an Active Directory user account. When the calling router attempts a VPN connection, the Router (Offline request) certificate is sent during the connection negotiation process. If the Router (Offline request) certificate is valid, it is used to determine the appropriate user account from which dial-in properties are obtained.

For information about configuring user and computer certificates for EAP-TLS authentication, see "Deploying a PPTP-based Site-to-Site VPN Connection" and "Deploying an L2TP-based Site-to-Site VPN Connection."

**Design Points: Certificate Infrastructure**

Consider the following when configuring the certificate infrastructure for site-to-site VPN connections:

- In order to create L2TP/IPSec site-to-site VPN connections using computer certificate authentication for IPSec, you must install a certificate in the computer certificate store of the calling router and the answering router.

- In order to authenticate VPN connections using EAP-TLS, the calling router must have a user certificate installed and the authenticating server (either the answering router or the RADIUS server) must have a computer certificate installed.

- To install a computer or user certificate on a computer across the Internet, make a PPTP connection using a password-based authentication protocol such as MS-CHAP v2. After connecting, use the Certificate Manager snap-in or Internet Explorer to request the appropriate certificates. Once the certificates are installed, disconnect and then reconnect with the appropriate VPN protocol and authentication method. An example of this situation is a computer at a remote branch office without the certificates needed to make L2TP/IPSec or EAP-TLS-authenticated connections.

# Deploying a PPTP-based Site-to-Site VPN Connection

Updated: October 7, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The deployment of PPTP-based site-to-site VPN connections using Windows Server 2003 consists of the following:

- Deploy certificate infrastructure
- Deploy Internet infrastructure
- Deploy the answering router
- Deploy the calling router
- Deploy AAA infrastructure
- Deploy site network infrastructure
- Deploy intersite network infrastructure

**Deploying Certificate Infrastructure**

For PPTP-based VPN connections, a certificate infrastructure is needed only when you are using EAP-TLS authentication. If you are only using a password-based authentication protocol such as MS-CHAP v2, a certificate infrastructure is not required and is not used for the authentication of the VPN connection.

To use EAP-TLS authentication for site-to-site VPN connections, you must:

- Install a user certificate on each calling router computer.
- Configure EAP-TLS on the calling router.
- Install a computer certificate on the authenticating server (the answering router or the RADIUS server).
- Configure EAP-TLS on the answering router and for the remote access policy for site-to-site connections.

**Installing a User Certificate on a Calling Router**

If you are using a Windows Server 2003 CA, a Router (Offline request) certificate is created and mapped to an Active Directory user account. To deploy a Router (Offline request) certificates for a calling router, the network administrator must:

1. Configure the Windows Server 2003 CA to issue Router (Offline request) certificates.
2. Request a Router (Offline request) certificate.
3. Export the Router (Offline request) certificate.
4. Map the certificate to the appropriate user account.
5. Send the Router (Offline request) certificate to the network administrator of the calling router.
6. Import the Router (Offline request) certificate on the calling router.

For more information about deploying Router (Offline request) certificates for demand-dial routing, see the topic titled Branch office demand-dial connection in Windows Server 2003 Help and Support.

For a third-party CA, see the documentation for the CA software for instructions about how to create a user certificate with the

Client Authentication enhanced key usage (object identifier "1.3.6.1.5.5.7.3.2") and export it so that it can be mapped to an Active Directory user account and sent to the network administrator of the calling router. You must also export the root CA certificate, the certificate of the issuing CA, and the certificates of any intermediate CAs and import them to the proper folder of the computer certificate store of the answering router using the Certificate Manager snap-in. For more information, see Appendix E of the Microsoft white paper titled "Virtual Private Networking with Windows Server 2003: Deploying Remote Access VPNs".

**Configuring EAP-TLS on a Calling Router**

To configure EAP-TLS for user certificates on the calling router:

- The demand-dial interface must be configured to use EAP with the **Smart Card or other certificate** EAP type by configuring advanced settings on the **Security** tab on the properties of a demand-dial interface. For the properties of the **Smart Card or other certificate** EAP type, select **Use a certificate on this computer**. If you want to validate the computer certificate of the VPN or IAS server, select **Validate server certificate**. If you want to configure the names of the authenticating servers, select **Connect to these servers** and type the server names. To require the servers computer certificate to have been issued a certificate from a specific trusted root CA, select the CA in the list of **Trusted Root Certification Authorities**.

- Right-click the demand-dial interface and click **Set credentials**. In the **Connect** dialog box, select the correct user or Router (Offline request) certificate in **User name on certificate**, and then click **OK**.

**Installing a Computer Certificate on the Authenticating Server**

To install a computer certificate, a certification authority must be present to issue certificates. If the CA is a Windows Server 2003 CA and the authenticating server is either the answering router or a Windows Server 2003 Internet Authentication Service (IAS) RADIUS server, you can install a certificate in the computer certificate store of the authenticating server in the following different ways:

1. By configuring the automatic allocation of computer certificates to computers in an Active Directory domain.

   This method allows a single point of configuration for the entire domain. All members of the domain automatically receive a computer certificate through Group Policy.

2. By using the Certificate Manager snap-in to request a certificate to store in the Certificates (Local Computer)\Personal folder.

   In this method, each computer must separately request a computer certificate from the CA. You must have administrator permissions to install a certificate using the Certificate Manager snap-in.

3. By using Internet Explorer and web enrollment to request a certificate and store it in local machine store.

   In this method, each computer must separately request a computer certificate from the CA. You must have administrator permissions to install a certificate using Web enrollment.

Based on the certificate policies in your organization, you only need to perform one of these methods.

For more information about using the Windows Server 2003 CA to obtain computer certificates, see the topics titled "Computer certificates for L2TP/IPSec VPN connections" and "Submit an advanced certificate request via the Web" in Windows Server 2003 Help and Support.

For a third-party CA, see the documentation for the CA software for instructions about how to create a certificate with the Server Authentication enhanced key usage (object identifier "1.3.6.1.5.5.7.3.1") and export it so that it can be imported using the Certificate Manager snap-in by an administrator on the answering router. Additionally, the root CA certificate, the certificate of the issuing CA, and the certificates of any intermediate CAs must be exported and imported on the calling router. For additional information, see Appendix E of the Microsoft white paper titled "Virtual Private Networking with Windows Server 2003: Deploying Remote Access VPNs".

**Configuring EAP-TLS on the Answering Router and Remote Access Policy**

To configure EAP-TLS authentication on the answering router:

- EAP must be enabled as an authentication type on the **Authentication Methods** dialog box available from the **Security** tab in the properties of the answering router in the Routing and Remote Access snap-in.

To configure EAP-TLS authentication on the remote access policy for either the answering router or IAS server:

- On the remote access policy that is being used for site-to-site VPN connections, the **Smart Card or other certificate** EAP type must be added to the selected EAP methods from the **Authentication** tab on the policy's profile settings. If the computer on which the remote access policy is being configured has multiple computer certificates installed, configure the properties of the **Smart Card or other certificate** EAP type and select the correct computer certificate to submit during the EAP-TLS authentication process.

If you are using a third-party RADIUS server, see the RADIUS server documentation for information about how to enable EAP-TLS and configure EAP-TLS to use the correct computer certificate.

**Deploying Internet Infrastructure**
Deploying the Internet infrastructure for site-to-site VPN connections consists of the following:

- Place VPN routers in the perimeter network or on the Internet.
- Install Windows Server 2003 on VPN router computers and configure Internet interfaces.

**Placing VPN Routers on the Perimeter Network or on the Internet**
Decide where to place the VPN routers in relation to your Internet firewall. In the most common configuration, the VPN routers are placed behind the firewall on the perimeter network between your site and the Internet. If so, configure packet filters on the firewall to allow PPTP traffic to and from the IP address of the VPN routers' perimeter network interfaces. For more information, see Appendix A.

**Installing Windows Server 2003 on VPN Routers and Configuring Internet Interfaces**
Install Windows Server 2003 on VPN router computers and connect it to either the Internet or to perimeter network with one network adapter and connect it to the site with another network adapter. Without running the Routing and Remote Access Server Setup Wizard, the VPN router computer will not forward IP packets between the Internet and the site. For the connection connected to the Internet or the perimeter network, configure the TCP/IP protocol with a public IP address, a subnet mask, and the default gateway of either the firewall (if the router is connected to a perimeter network) or an ISP router (if the router is directly connected to the Internet). Do not configure the connection with DNS server or WINS server IP addresses.

**Deploying the Answering Router**
Deploying the answering router for a site-to-site VPN connection consists of the following:

- Configure the answering router's connection to the site.
- Run the Routing and Remote Access Server Setup Wizard.
- Configure a demand-dial interface.

**Configuring the Answering Router's Connection to the Site**
Configure the connection connected to the site with a manual TCP/IP configuration consisting of IP address, subnet mask, site DNS servers, and site WINS servers. Note that you must not configure the default gateway on the site connection to prevent default route conflicts with the default route pointing to the Internet.

**Running the Routing and Remote Access Server Setup Wizard**
Run the Routing and Remote Access Server Setup Wizard to configure the Windows Server 2003 answering router using the following steps:

1. Click **Start**, point to Programs, point to **Administrative Tools**, and then click **Routing and Remote Access**.
2. Right-click the answering name, and then click **Configure and Enable Routing and Remote Access**. Click **Next**.
3. In **Configuration**, click **Remote access (dial-up or VPN)** and then click **Next**. If you additionally want to use the answering router as a network address translator (NAT), Web server, or other function, see Appendix B.
4. In **Remote Access**, select **VPN**. If you also want the answering router to support dial-up site-to-site connections, click

**Dial-up**. Click **Next**.

5. In **VPN Connection**, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click **Next**.

6. In **IP Address Assignment**, click **Automatically** if the answering router should use DHCP to obtain IP addresses for remote access VPN clients and calling routers. Or, click **From a specified range of addresses** to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure in order for the virtual interfaces of calling routers to be reachable. When IP address assignment is complete, click **Next**.

7. In **Managing Multiple Remote Access Servers**, if you are using RADIUS for authentication and authorization, click **Yes, set up this server to work with a RADIUS server**, and then click **Next**.

   - In **RADIUS Server Selection**, configure the primary (mandatory) and alternate (optional) RADIUS servers and the shared secret, and then click **Next**.

8. Click **Finish**.

By default, only 128 PPTP ports are configured on the WAN Miniport (PPTP) device. If you need more PPTP ports, configure the **WAN Miniport (PPTP)** device from the properties of the **Ports** object in the Routing and Remote Access snap-in. By default, 128 L2TP ports are also configured.

By default, the MS-CHAP, MS-CHAP v2, and EAP authentication methods are enabled.

**Configuring a Demand-dial Interface**

From the Routing and Remote Access snap-in on the answering router, perform the following steps:

1. In the console tree, right-click **Network Interfaces**, and then click **New Demand-dial Interface**.

2. On the **Welcome to the Demand-Dial Interface Wizard** page, click **Next**.

3. On the **Interface Name** page, type the name of the demand-dial interface, and then click **Next**.

4. On the **Connection Type** page, click **Connect using virtual private networking (VPN)**, and then click **Next**.

5. On the **VPN Type** page, click **Point to Point Tunneling Protocol (PPTP)**, and then click **Next**.

6. On the **Destination Address** page, type the IP address of the calling router.

   For a two-way-initiated router to-router VPN connection, configure the IP address of the calling router. For a one-way initiated site-to-site VPN connection, you can skip this step because the answering router never uses this interface to initiate a connection to the calling router.

7. On the **Protocols and Security** page, select the **Route IP packets on this interface** and **Add a user account so that a remote router can dial in** check boxes, and then click **Next**.

8. On the **Static Routes for Remote Networks** page, click **Add** to add static routes assigned to the demand-dial interface (as needed).

9. On the **Dial In Credentials** page, type the password of the user account used by the calling router in **Password** and **Confirm password**, and then click **Next**. This step automatically creates a user account with the same name as the demand-dial interface that is being created. This is done so that when the calling router initiates a connection to the answering router, it is using a user account name that matches the name of a demand-dial interface. Therefore, the answering router can determine that the incoming connection from the calling router is a demand-dial connection rather than a remote access connection.

10. On the **Dial Out Credentials** page, type the user name in **User name**, the user account domain name in **Domain**, and the user account password in both **Password** and **Confirm password**.

    For a two-way-initiated router to-router VPN connection, configure the name, domain, and password when this router is acting as the calling router. For a one-way initiated site-to-site VPN connection, you can type any name in **User name** and skip the rest of the fields because this router never uses this interface to initiate a connection to the calling router.

11. On the **Completing the Demand-Dial Interface Wizard** page, click **Finish**.

The result of this configuration is a PPTP-based demand-dial interface over which IP routing is enabled. A user account with the same name as the demand-dial interface is automatically added with correct account and dial-in settings.

**Deploying the Calling Router**

Deploying the calling router for a site-to-site VPN connection consists of the following:

- Configure the calling router's connection to the site.
- Run the Routing and Remote Access Server Setup Wizard.
- Configure a demand-dial interface.

**Configuring the Calling Router's Connection to the Site**

Configure the connection connected to the site with a manual TCP/IP configuration consisting of IP address, subnet mask, site DNS servers, and site WINS servers. Note that you must not configure the default gateway on the site connection to prevent default route conflicts with the default route pointing to the Internet.

**Running the Routing and Remote Access Server Setup Wizard**

Run the Routing and Remote Access Server Setup Wizard to configure the Windows Server 2003 calling router using the following steps:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Routing and Remote Access**.
2. Right-click your server name, and then click **Configure and Enable Routing and Remote Access**. Click **Next**.
3. In **Configuration**, click **Remote access (dial-up or VPN)** and then click **Next**. If you additionally want to use the calling router as a network address translator (NAT), Web server, or other function, see Appendix B.
4. In **Remote Access**, select **VPN**. If you also want the VPN router to support dial-up site-to-site connections, click **Dial-up**. Click **Next**.
5. In **VPN Connection**, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click **Next**.
6. In **IP Address Assignment**, click **Automatic** if the calling router should use DHCP to obtain IP addresses for other calling routers when it is acting as an answering router. Or, click **From a specified range of addresses** to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure in order for the virtual interfaces of routers calling this router to be reachable. When IP address assignment is complete, click **Next**.
7. In **Managing Multiple Remote Access Servers**, if you are using RADIUS for authentication and authorization, click **Yes, set up this server to work with a RADIUS server**, and then click **Next**.

    - In **RADIUS Server Selection**, configure the primary (mandatory) and alternate (optional) RADIUS servers and the shared secret, and then click **Next**.

8. Click **Finish**.

**Configuring a Demand-dial Interface**

From the Routing and Remote Access snap-in on the calling router, perform the following steps:

1. In the console tree, right-click **Network Interfaces**, and then click **New Demand-dial Interface**.

2. On the **Welcome to the Demand-Dial Interface Wizard** page, click **Next**.

3. On the **Interface Name** page, type the name of the demand-dial interface. For a two-way initiated connection, this is the same name as the user name in the user credentials used by the answering router when it is acting as a calling router. Click **Next**.

4. On the **Connection Type** page, click **Connect using virtual private networking (VPN)**, and then click **Next**.

5. On the **VPN Type** page, click **Point to Point Tunneling Protocol (PPTP)**, and then click **Next**.

6. On the **Destination Address** page, type the IP address of the answering router.

7. On the **Protocols and Security** page, select the **Route IP packets on this interface** check box. For a two-way initiated connection, select the **Add a user account so that a remote router can dial in** check box. Click **Next**.

8. On the **Static Routes for Remote Networks** page, click **Add** to add static routes assigned to the demand-dial interface (as needed).

9. For a two-way initiated connection, in the **Dial In Credentials** page, type the password of the user account used by the answering router acting as a calling router in **Password** and **Confirm password**, and then click **Next**. This step automatically creates a user account with the same name as the demand-dial interface that is being created. This is done so that when the answering router acting as a calling router initiates a connection to this router, it is using a user account name that matches the name of a demand-dial interface. Therefore, this router can determine that the incoming connection from the answering router acting as a calling router is a demand-dial connection rather than a remote access connection.

10. On the **Dial Out Credentials** page, type the user name in **User name**, the user account domain name in **Domain**, and the user account password in both **Password** and **Confirm password**.

11. On the **Completing the demand-dial interface wizard** page, click **Finish**.

The result of this configuration is a PPTP-based demand-dial interface over which IP routing is enabled. A user account with the same name as the demand-dial interface is automatically added with correct account and dial-in settings (if needed).

**Deploying AAA Infrastructure**

Deploying the AAA infrastructure for site-to-site VPN connections consists of the following:

- Configure Active Directory for user accounts and groups.

- Configure the primary IAS server on a domain controller.

- Configure the secondary IAS server on a different domain controller.

This configuration must be done at each site containing an answering router. For branch offices with few computers and a single answering router, it is easier to configure the Routing and Remote Access service for Windows authentication and use locally configured remote access policies than configuring a separate IAS server computer.

**Configuring Active Directory for User Accounts and Groups**

To configure Active Directory for user accounts and groups, do the following:

1. Ensure that all calling routers have a corresponding user account with the correct account and dial-in settings. This includes calling routers for branch offices and business partners. User accounts with the correct account and dial-in settings are automatically created when you select the **Add a user account so that a remote router can dial in** check

box on the **Protocols and Security** page of the Demand-Dial Interface Wizard.

2. Organize user accounts used by calling routers into the appropriate universal and nested groups to take advantage of group-based remote access policies. For more information, see the topic titled "Nesting groups" in Windows Server 2003 Help and Support.

**Configuring the Primary IAS Server on a Domain Controller**

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows Server 2003 Help and Support.

2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support.

   If the IAS server authenticates connection attempts for user accounts in other domains, verify that these domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support. For more information about trust relationships, see the topic titled "Understanding Domains and Forests" in Windows Server 2003 Help and Support.

3. If the IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.

4. Enable file logging for accounting and authentication events. For more information, see the topic titled "Configure log file properties" in Windows Server 2003 Help and Support.

5. Add the VPN router(s) as RADIUS clients of the IAS server. For more information, see the topic titled "Add RADIUS clients" in Windows Server 2003 Help and Support. For the IP address of each VPN router, use the site IP address assigned to the VPN router. If you are using names, use the internal name of the VPN router. Use strong shared secrets.

6. Create remote access policies that reflect your remote access usage scenarios.

   For example, to configure a remote access policy that requires PPTP-based site-to-site VPN connections for members of the VPNRouters group to use MS-CHAP v2 authentication and 128-bit encryption, use the New Remote Access Policy Wizard to create a typical remote access policy with the following settings:

   Policy name: Site-to-Site VPN connections (example)

   Access method: VPN

   Group access: VPNRouters group (example)

   Authentication methods: **Enable only Microsoft Encrypted Authentication version 2 (MS-CHAP v2)**

   Policy encryption level: Select only **Strongest encryption**

If IAS is not installed on a domain controller, you must configure the primary IAS server computer to read the properties of user accounts in the domain. For more information, see the "Enable the IAS server to read user accounts in Active Directory" in Windows Server 2003 Help and Support.

**Configuring the Secondary IAS Server on a Different Domain Controller**

To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows Server 2003 Help and Support.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support. For more information about trust relationships, see the topic titled "Understanding Domains and Forests" in Windows Server 2003 Help and Support.

   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.

4. To copy the configuration of the primary IAS server to the secondary IAS server, type **netsh aaaa show config** >*path\file.txt* at a command prompt on the primary IAS server. This stores the configuration settings, including registry settings, in a text file. The path can be relative, absolute, or a network path.

5. Copy the file created in step 4 to the secondary IAS server. At a command prompt on the secondary IAS server, type **netsh exec***path\file.txt*. This imports all the settings configured on the primary IAS server to the secondary IAS server.

If IAS is not installed on a domain controller, you must configure the secondary IAS server computer to read the properties of user accounts in the domain. For more information, see the "Enable the IAS server to read user accounts in Active Directory" in Windows Server 2003 Help and Support.

**Deploying Site Network Infrastructure**
Deploying the network infrastructure of a site for site-to-site VPN connections consists of the following:

- Configure routing on the VPN routers.
- Verify reachability from each VPN router.
- Configure routing for off-subnet address pools.

**Configuring Routing on the VPN Routers**
In order for your VPN routers to properly forward traffic to locations within the site in which they are located, you must configure them with either static routes that summarize all the possible addresses used on in the site or with routing protocols so that the VPN router can participate as a dynamic router and automatically add routes for site subnets to its routing table.

To add static routes, see the topic titled "Add a static route" in Windows Server 2003 Help and Support. To configure the VPN router as a RIP router, see the topic titled "Configure RIP for IP". To configure the VPN router as an OSPF router, see the topics titled "OSPF design considerations" and "Configure OSPF" in Windows Server 2003 Help and Support.

**Verifying Reachability from each VPN Router**
From each VPN router, verify that the VPN router computer can resolve names and successfully communicate with resources in the VPN router's site by using the Ping command, Internet Explorer, and making drive and printer connections to known servers within the site.

**Configuring Routing for Off-subnet Address Pools**
If you configured any of the VPN routers with a static address pool and any of the ranges within the pool are an off-subnet range, you must ensure that the route(s) representing the off-subnet address range(s) are present in your site routing infrastructure to reach the virtual interfaces of calling routers. You can ensure this by adding static route(s) representing the off-subnet address range(s) as static routes to the neighboring router(s) of the VPN router(s) and then using the routing protocol of your site to propagate the route to other routers. When you add the static route(s), you must specify that the gateway or next hop address is the site interface of the VPN router.

Alternately, if you are using RIP or OSPF, you can configure the VPN routers using off-subnet address pools as RIP or OSPF routers. For OSPF, you must configure the VPN router as an autonomous system boundary router (ASBR). For more

information, see the topic titled "OSPF design considerations" in Windows Server 2003 Help and Support.

**Deploying Intersite Network Infrastructure**

Deploying the intersite network infrastructure consists of configuring each VPN router with the set of routes for subnets that are available in the other sites (across each site-to-site VPN connection). This can be done in the following ways:

- Manually configure static routes on each VPN router.
- Perform auto-static updates on each VPN router.
- Configure routing protocols to operate over the site-to-site VPN connection.

**Manually Configuring Static Routes on Each VPN Router**

From the Routing and Remote Access snap-in, perform the following steps:

1. In the console tree, click **IP Routing**, and then click **Static Routes**.
2. Right-click **Static Routes**, and then click **New Static Route**.
3. In the **Static Route** dialog box, select the appropriate demand-dial interface name, and type the destination, network mask, and metric.
4. Click **OK** to add the route.
5. For an additional route, go back to step 2.

| **Note** |
| --- |
| Because the demand-dial connection is a point-to-point connection, the **Gateway IP address** field is not configurable. The adding of static routes can also be done during the creation of the demand-dial interface with the New Demand-Dial Interface Wizard. |

**Performing Auto-static Updates on Each VPN Router**

If RIP for IP is enabled on the demand-dial interfaces of both VPN routers, you can use auto-static updates to automatically configure static routes when the VPN connection is in a connected state. A demand-dial interface that is configured for auto-static updates sends a request across an active connection to request all of the routes of the router on the other side of the connection. In response to the request, all of the routes of the requested router are automatically entered as static routes in the routing table of the requesting router.

From the Routing and Remote Access snap-in on a VPN router (assuming the site-to-site VPN connection is active), perform the following steps:

1. In the console tree, click **IP Routing**, and then click **General**.
2. In the details pane, right-click the appropriate demand-dial interface, and then click **Update Routes**.

You can also use the **netsh interface set interface** command to perform an auto-static update. For more information, see the topic titled "Scheduling auto-static updates" in Windows Server 2003 Help and Support.

**Configuring Routing Protocols**

If the site-to-site VPN connection is persistent (always active), you can also configure IP routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) to operate over the VPN connection. For more information, see the topics titled "Setting up a RIP-for-IP routed internetwork" and "Setting up an OSPF routed internetwork" in Windows Server 2003 Help and Support.

# Deploying an L2TP/IPSec-based Site-to-Site VPN Connection

Updated: October 7, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The deployment of L2TP/IPSec-based site-to-site VPN connections using Windows Server 2003 consists of the following:

- Deploy certificate infrastructure
- Deploy Internet infrastructure
- Deploy the answering router
- Deploy the calling router
- Deploy AAA infrastructure
- Deploy site network infrastructure
- Deploy intersite network infrastructure

**Deploying Certificate Infrastructure**
For L2TP-based VPN connections, a certificate infrastructure is required to issue certificates needed to perform IPSec authentication. Additionally, a certificate infrastructure is also needed when you are using EAP-TLS authentication.

**Certificates for L2TP Connections**
To install a computer certificate, a CA must be present to issue certificates. If the CA is a Windows Server 2003 CA, you can install a certificate in the computer certificate store of the VPN router in the following different ways:

1. By configuring the automatic allocation of computer certificates to computers in an Active Directory domain.

   This method allows a single point of configuration for the entire domain. All members of the domain automatically receive a computer certificate through Group Policy.

2. By using the Certificate Manager snap-in to request a certificate to store in the Certificates (Local Computer)\Personal folder.

   In this method, each computer must separately request a computer certificate from the CA. You must have administrator permissions to install a certificate using the Certificate Manager snap-in.

3. By using Internet Explorer and web enrollment to request a certificate and store it in local machine store.

   In this method, each computer must separately request a computer certificate from the CA. You must have administrator permissions to install a certificate using Web enrollment.

Based on the certificate policies in your organization, you only need to perform one of these methods.

For more information about using the Windows Server 2003 CA to obtain computer certificates, see the topics titled "Computer certificates for L2TP/IPSec VPN connections" and "Submit an advanced certificate request via the Web" in Windows Server 2003 Help and Support.

For a third-party CA, see the documentation for the CA software for instructions about how to create a certificate and export it so that it can be imported into the computer certificate store using the Certificate Manager snap-in by an administrator on the calling and answering routers. Additionally, the root CA certificate, the certificate of the issuing CA, and the certificates of any intermediate CAs must be exported and imported on the calling and answering routers. For additional information, see Appendix E of the Microsoft white paper titled "Virtual Private Networking with Windows Server 2003: Deploying Remote Access VPNs".

**Certificates for EAP-TLS Authentication**
To use EAP-TLS authentication of site-to-site VPN connections, you must:

- Install a user certificate on each calling router computer.
- Configure EAP-TLS on the calling router.
- Install a computer certificate on the authenticating server (the answering router or the RADIUS server)
- Configure EAP-TLS on the answering router and remote access policy.

**Installing a User Certificate on a Calling Router**
If you are using a Windows Server 2003 CA, a Router (Offline request) certificate (a special type of user certificate for demand-dial connections), is created and mapped to an Active Directory user account. To deploy a Router (Offline request) certificates for a calling router, the network administrator must:

1. Configure the Windows Server 2003 CA to issue Router (Offline request) certificates.
2. Request a Router (Offline request) certificate.
3. Export the Router (Offline request) certificate.
4. Map the certificate to the appropriate user account.
5. Send the Router (Offline request) certificate to the network administrator of the calling router.
6. Import the Router (Offline request) certificate on the calling router.

For more information about deploying Router (Offline request) certificates for demand-dial routing, see the topic titled Branch office demand-dial connection in Windows Server 2003 Help and Support.

For a third-party CA, see the documentation for the CA software for instructions about how to create a user certificate with the Client Authentication enhanced key usage (object identifier "1.3.6.1.5.5.7.3.2") and export it so that it can be mapped to an Active Directory user account and sent to the network administrator of the calling router. You must also export the root CA certificate, the certificate of the issuing CA, and the certificates of any intermediate CAs and import them to the proper folder of the computer certificate store of the answering router using the Certificate Manager snap-in. For additional information, see Appendix E of the Microsoft white paper titled "Virtual Private Networking with Windows Server 2003: Deploying Remote Access VPNs".

**Configuring EAP-TLS on the Calling Router**
To configure EAP-TLS for user certificates on the calling router:

- The demand-dial interface must be configured to use EAP with the **Smart Card or other certificate** EAP type by configuring advanced settings on the **Security** tab on the properties of a demand-dial interface. In the properties of the **Smart Card or other certificate** EAP type, select **Use a certificate on this computer**. If you want to validate the computer certificate of the VPN or IAS server, select **Validate server certificate**. If you want to configure the names of the authenticating servers, select **Connect to these servers** and type the server names. To require the server's computer certificate to have been issued a certificate from a specific trusted root CA, select the CA in the list of **Trusted Root Certification Authorities**.
- Right-click the demand-dial interface and click **Set credentials**. In the **Connect** dialog box, select the correct user or Router (Offline request) certificate in **User name on certificate**, and then click **OK**.

**Installing a Computer Certificate on the Authenticating Server**
If the authenticating server is the answering router, you can use the same computer certificate that is installed to authenticate L2TP connections for EAP-TLS authentication provided it contains the Server Authentication enhanced key usage (object identifier "1.3.6.1.5.5.7.3.1"). If not, you must install an additional computer certificate that contains the Server Authentication enhanced key usage. If the authenticating server is an IAS server, you must install a computer certificate that contains the Server Authentication enhanced key usage.

If the CA is a Windows Server 2003 CA and the authenticating server is either the answering router or a Windows Server 2003 Internet Authentication Service (IAS) RADIUS server, you can install a certificate in the computer certificate store of the authenticating server using the methods described in the "Certificates for L2TP connections" section of this paper.

For a third-party CA, see the documentation for the CA software for instructions about how to create a certificate with the Server Authentication enhanced key usage (object identifier "1.3.6.1.5.5.7.3.1") and export it so that an administrator on the answering router can import it using the Certificate Manager snap-in. Additionally, the root CA certificate, the certificate of the issuing CA, and the certificates of any intermediate CAs must be exported and imported on the calling router. For additional information, see Appendix E of the Microsoft white paper titled "Virtual Private Networking with Windows Server 2003: Deploying Remote Access VPNs".

**Configuring EAP-TLS on the Answering Router and Remote Access Policy**
To configure EAP-TLS authentication on the answering router:

- EAP must be enabled as an authentication type on the **Authentication Methods** dialog box available from the **Security** tab in the properties of the answering router in the Routing and Remote Access snap-in.

To configure EAP-TLS authentication on the remote access policy of the answering router or IAS server:

- On the remote access policy that is being used for site-to-site VPN connections, EAP must be enabled with the **Smart Card or other certificate** EAP type must be added to the selected EAP methods from the **Authentication** tab on the policy's profile settings. If the computer on which the remote access policy is being configured has multiple computer certificates installed, configure the properties of the **Smart Card or other certificate** EAP type and select the appropriate computer certificate to submit during the EAP-TLS authentication process.

If you are using a third-party RADIUS server, see the RADIUS server documentation for information about how to enable EAP-TLS and configure EAP-TLS to use the correct computer certificate.

**Deploying Internet Infrastructure**
Deploying the Internet infrastructure for site-to-site VPN connections consists of the following:

- Place VPN routers in the perimeter network or on the Internet.
- Install Windows Server 2003 on VPN router computers and configure Internet interfaces.

**Placing VPN Routers on the Perimeter Network or on the INTERNET**
Decide where to place the VPN routers in relation to your Internet firewall. In the most common configuration, the VPN routers are placed behind the firewall on the perimeter network between your site and the Internet. If so, configure packet filters on the firewall to allow L2TP/IPSec traffic to and from the IP address of the VPN routers' perimeter network interfaces. For more information, see Appendix A.

**Installing Windows Server 2003 on VPN Routers and Configuring Internet Interfaces**
Install Windows Server 2003 on VPN router computers and connect them to either the Internet or to the perimeter network with one network adapter and connect it to the site with another network adapter. Without running the Routing and Remote Access Server Setup Wizard, the VPN router computer will not forward IP packets between the Internet and the site. For the connection connected to the Internet or the perimeter network, configure the TCP/IP protocol with a public IP address, a subnet mask, and the default gateway of either the firewall (if the router is connected to a perimeter network) or an ISP router (if the router is directly connected to the Internet). Do not configure the connection with DNS server or WINS server IP addresses.

**Deploying the Answering Router**
Deploying the answering router for a site-to-site VPN connection consists of the following:

- Configure the answering router's connection to the site.
- Run the Routing and Remote Access Server Setup Wizard.
- Configure a demand-dial interface.

**Configuring the Answering Router's Connection to the Site**
Configure the connection connected to the site with a manual TCP/IP configuration consisting of IP address, subnet mask, site

DNS servers, and site WINS servers. Note that you must not configure the default gateway on the site connection to prevent default route conflicts with the default route pointing to the Internet.

**Running the Routing and Remote Access Server Setup Wizard**

Run the Routing and Remote Access Server Setup Wizard to configure the Windows Server 2003 answering router using the following steps:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Routing and Remote Access**.

2. Right-click the answering name, and then click **Configure and Enable Routing and Remote Access**. Click **Next**.

3. In **Configuration**, click **Remote access (dial-up or VPN)** and then click **Next**. If you additionally want to use the answering router as a network address translator (NAT), Web server, or other function, see Appendix B.

4. In **Remote Access**, select **VPN**. If you also want the answering router to support dial-up site-to-site connections, click **Dial-up**. Click **Next**.

5. In **VPN Connection**, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click **Next**.

6. In **IP Address Assignment**, click **Automatically** if the answering router should use DHCP to obtain IP addresses for remote access VPN clients and calling routers. Or, click **From a specified range of addresses** to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure in order for the virtual interfaces of calling routers to be reachable. When IP address assignment is complete, click **Next**.

7. In **Managing Multiple Remote Access Servers**, if you are using RADIUS for authentication and authorization, click **Yes, set up this server to work with a RADIUS server**, and then click **Next**.

   - In **RADIUS Server Selection**, configure the primary (mandatory) and alternate (optional) RADIUS servers and the shared secret, and then click **Next**.

8. Click **Finish**.

By default, only 128 L2TP ports are configured on the WAN Miniport (L2TP) device. If you need more L2TP ports, configure the **WAN Miniport (L2TP)** device from the properties of the **Ports** object in the Routing and Remote Access snap-in. By default, 128 PPTP ports are also configured.

By default, the MS-CHAP, MS-CHAP v2, and EAP authentication methods are enabled.

**Configuring a Demand-dial Interface**

From the Routing and Remote Access snap-in on the answering router, perform the following steps:

1. In the console tree, right-click **Network Interfaces**, and then click **New Demand-dial Interface**.

2. On the **Welcome to the Demand-Dial Interface Wizard** page, click **Next**.

3. On the **Interface Name** page, type the name of the demand-dial interface, and then click **Next**.

4. On the **Connection Type** page, click **Connect using virtual private networking (VPN)**, and then click **Next**.

5. On the **VPN Type** page, click **Layer 2 Tunneling Protocol (L2TP)**, and then click **Next**.

6. On the **Destination Address** page, type the IP address of the calling router.

   For a two-way-initiated router to-router VPN connection, configure the IP address of the calling router. For a one-way initiated site-to-site VPN connection, you can skip this step because the answering router never uses this interface to initiate a connection to the calling router.

7. On the **Protocols and Security** page, select the **Route IP packets on this interface** and **Add a user account so that a remote router can dial in** check boxes, and then click **Next**.

8. On the **Static Routes for Remote Networks** page, click **Add** to add static routes assigned to the demand-dial interface (as needed).

9. On the **Dial In Credentials** page, type the password of the user account used by the calling router in **Password** and **Confirm password**, and then click **Next**. This step automatically creates a user account with the same name as the demand-dial interface that is being created. This is done so that when the calling router initiates a connection to the answering router, it is using a user account name that matches the name of a demand-dial interface. Therefore, the answering router can determine that the incoming connection from the calling router is a demand-dial connection rather than a remote access connection.

10. On the **Dial Out Credentials** page, type the user name in **User name**, the user account domain name in **Domain**, and the user account password in both **Password** and **Confirm password**.

    For a two-way-initiated router to-router VPN connection, configure the name, domain, and password when this router is acting as the calling router. For a one-way initiated site-to-site VPN connection, you can type any name in User name and skip the rest of the fields because this router never uses this interface to initiate a connection to the calling router.

11. On the **Completing the Demand-Dial Interface** Wizard page, click **Finish**.

The result of this configuration is an L2TP-based demand-dial interface over which IP routing is enabled. A user account with the same name as the demand-dial interface is automatically added with correct account and dial-in settings.

**Deploying the Calling Router**

Deploying the calling router for a site-to-site VPN connection consists of the following:

- Configure the calling router's connection to the site.
- Run the Routing and Remote Access Server Setup Wizard.
- Configure a demand-dial interface.

**Configuring the Calling Router's Connection to the Site**

Configure the connection connected to the site with a manual TCP/IP configuration consisting of IP address, subnet mask, site DNS servers, and site WINS servers. Note that you must not configure the default gateway on the site connection to prevent default route conflicts with the default route pointing to the Internet.

**Running the Routing and Remote Access Server Setup Wizard**

Run the Routing and Remote Access Server Setup Wizard to configure the Windows Server 2003 calling router using the following steps:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Routing and Remote Access**.

2. Right-click your server name, and then click **Configure and Enable Routing and Remote Access**. Click **Next**.

3. In **Configuration**, click **Remote access (dial-up or VPN)** and then click **Next**. If you additionally want to use the calling router as a network address translator (NAT), Web server, or other function, see Appendix B.

4. In **Remote Access**, select **VPN**. If you also want the VPN router to support dial-up site-to-site connections, click **Dial-up**. Click **Next**.

5. In **VPN Connection**, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click **Next**.

6. In **IP Address Assignment**, click **Automatic** if the calling router should use DHCP to obtain IP addresses for other

calling routers when it is acting as an answering router. Or, click **From a specified range of addresses** to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure in order for the virtual interfaces of routers calling this router to be reachable. When IP address assignment is complete, click **Next**.

7. In **Managing Multiple Remote Access Servers**, if you are using RADIUS for authentication and authorization, click **Yes, set up this server to work with a RADIUS server**, and then click **Next**.

    - In **RADIUS Server Selection**, configure the primary (mandatory) and alternate (optional) RADIUS servers and the shared secret, and then click **Next**.

8. Click **Finish**.

**Configuring a Demand-dial Interface**

From the Routing and Remote Access snap-in on the calling router, perform the following steps:

1. In the console tree, right-click **Network Interfaces**, and then click **New Demand-dial Interface**.

2. On the **Welcome to the Demand-Dial Interface Wizard** page, click **Next**.

3. On the **Interface Name** page, type the name of the demand-dial interface. For a two-way initiated connection, this is the same name as the user name in the user credentials used by the answering router when it is acting as a calling router. Click **Next**.

4. On the **Connection Type** page, click **Connect using virtual private networking (VPN)**, and then click **Next**.

5. On the **VPN Type** page, click **Layer 2 Tunneling Protocol (L2TP)**, and then click **Next**.

6. On the **Destination Address** page, type the IP address of the answering router.

7. On the **Protocols and Security** page, select the **Route IP packets on this interface** check box. For a two-way initiated connection, select the **Add a user account so that a remote router can dial in** check box. Click **Next**.

8. On the **Static Routes for Remote Networks** page, click **Add** to add static routes assigned to the demand-dial interface (as needed).

9. For a two-way initiated connection, in the **Dial In Credentials** page, type the password of the user account used by the answering router acting as a calling router in **Password** and **Confirm password**, and then click **Next**. This step automatically creates a user account with the same name as the demand-dial interface that is being created. This is done so that when the answering router acting as a calling router initiates a connection to this router, it is using a user account name that matches the name of a demand-dial interface. Therefore, this router can determine that the incoming connection from the answering router acting as a calling router is a demand-dial connection rather than a remote access connection.

10. On the **Dial Out Credentials** page, type the user name in **User name**, the user account domain name in **Domain**, and the user account password in both **Password** and **Confirm password**.

11. On the **Completing the demand-dial interface wizard** page, click **Finish**.

The result of this configuration is an L2TP-based demand-dial interface over which IP routing is enabled. A user account with the same name as the demand-dial interface is automatically added with correct account and dial-in settings.

**Deploying AAA Infrastructure**

Deploying the AAA infrastructure for site-to-site VPN connections consists of the following:

- Configure Active Directory for user accounts and groups.

- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.

This configuration must be done at each site containing an answering router. For branch offices with few computers and a single answering router, it is easier to configure the Routing and Remote Access service for Windows authentication and use locally configured remote access policies than configuring a separate IAS server computer.

**Configuring Active Directory for User Accounts and Groups**

To configure Active Directory for user accounts and groups, do the following:

1. Ensure that all calling routers have a corresponding user account with the correct account and dial-in settings. This includes calling routers for branch offices and business partners. User accounts with the correct account and dial-in settings are automatically created when you select the **Add a user account so that a remote router can dial in** check box on the **Protocols and Security** dialog box in the Demand-Dial Interface Wizard.

2. Organize user accounts used by calling routers into the appropriate universal and nested groups to take advantage of group-based remote access policies. For more information, see the topic titled "Nesting groups" in Windows Server 2003 Help and Support.

**Configuring the Primary IAS Server on a Domain Controller**

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows Server 2003 Help and Support.

2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support.

3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that these domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support. For more information about trust relationships, see the topic titled "Understanding Domains and Forests" in Windows Server 2003 Help and Support.

   If the IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.

4. Enable file logging for accounting and authentication events. For more information, see the topic titled "Configure log file properties" in Windows Server 2003 Help and Support.

5. Add the VPN router(s) as RADIUS clients of the IAS server. For more information, see the topic titled "Add RADIUS clients" in Windows Server 2003 Help and Support. For the IP address of each VPN router, use the site IP address assigned to the VPN router. If you are using names, use the internal name of the VPN router. Use strong shared secrets.

6. Create remote access policies that reflect your remote access usage scenarios.

   For example, to configure a remote access policy that requires L2TP-based site-to-site VPN connections for members of the VPNRouters group to use MS-CHAP v2 authentication and 128-bit encryption, use the New Remote Access Policy Wizard to create a typical remote access policy with the following settings:

   Policy name: Site-to-Site VPN connections (example)

   Access method: VPN

Group access: VPNRouters group (example)

Authentication methods: Enable only **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)**

Policy encryption level: Select only **Strongest encryption**

If IAS is not installed on a domain controller, you must configure the primary IAS server computer to read the properties of user accounts in the domain. For more information, see the "Enable the IAS server to read user accounts in Active Directory" in Windows Server 2003 Help and Support.

**Configuring the Secondary IAS Server on a Different Domain Controller**
To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows Server 2003 Help and Support.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows Server 2003 Help and Support. For more information about trust relationships, see the topic titled "Understanding Domains and Forests" in Windows Server 2003 Help and Support.

   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.

4. To copy the configuration of the primary IAS server to the secondary IAS server, type **netsh aaaa show config** >*path\file.txt* at a command prompt on the primary IAS server. This stores the configuration settings, including registry settings, in a text file. The path can be relative, absolute, or a network path.

5. Copy the file created in step 4 to the secondary IAS server. At a command prompt on the secondary IAS server, type **netsh exec***path\file.txt*. This imports all the settings configured on the primary IAS server to the secondary IAS server.

If IAS is not installed on a domain controller, you must configure the secondary IAS server computer to read the properties of user accounts in the domain. For more information, see the "Enable the IAS server to read user accounts in Active Directory" in Windows Server 2003 Help and Support.

**Deploying Site Network Infrastructure**
Deploying the network infrastructure of a site for site-to-site VPN connections consists of the following:

- Configure routing on the VPN routers.

- Verify reachability from each VPN router.

- Configure routing for off-subnet address pools.

**Configuring Routing on the VPN Routers**
In order for your VPN routers to properly forward traffic to locations within the site in which they are located, you must configure them with either static routes that summarize all the possible addresses used on in the site or with routing protocols so that the VPN router can participate as a dynamic router and automatically add routes for site subnets to its routing table.

To add static routes, see the topic titled "Add a static route" in Windows Server 2003 Help and Support. To configure the VPN router as a RIP router, see the topic titled "Configure RIP for IP". To configure the VPN router as an OSPF router, see the topics titled "OSPF design considerations" and "Configure OSPF" in Windows Server 2003 Help and Support.

**Verifying Reachability from Each VPN Router**
From each VPN router, verify that the VPN router computer can resolve names and successfully communicate with resources in the VPN router's site by using the Ping command, Internet Explorer, and making drive and printer connections to known servers within the site.

**Configuring Routing for Off-Subnet Address Pools**
If you configured any of the VPN routers with a static address pool and any of the ranges within the pool are an off-subnet range, you must ensure that the route(s) representing the off-subnet address range(s) are present in your site routing infrastructure to reach the virtual interfaces of calling routers. You can ensure this by adding static route(s) representing the off-subnet address range(s) as static routes to the neighboring router(s) of the VPN router(s) and then using the routing protocol of your site to propagate the route to other routers. When you add the static route(s), you must specify that the gateway or next hop address is the site interface of the VPN router.

Alternately, if you are using RIP or OSPF, you can configure the VPN routers using off-subnet address pools as RIP or OSPF routers. For OSPF, you must configure the VPN router as an autonomous system boundary router (ASBR). For more information, see the topic titled "OSPF design considerations" in Windows Server 2003 Help and Support.

Deploying Intersite Network Infrastructure

Deploying the intersite network infrastructure consists of configuring each VPN router with the set of routes for subnets that are available in the other sites (across each site-to-site VPN connection). This can be done in the following ways:

- Manually configure static routes on each VPN router.
- Perform auto-static updates on each VPN router.
- Configure routing protocols to operate over the site-to-site VPN connection.

**Manually Configuring Static Routes on Each VPN Router**
From the Routing and Remote Access snap-in, perform the following steps:

1. In the console tree, click **IP Routing**, and then click **Static Routes**.
2. Right-click **Static Routes**, and then click **New Static Route**.
3. In the **Static Route** dialog box, select the appropriate demand-dial interface name, and type the destination, network mask, and metric.
4. Click **OK** to add the route.
5. For an additional route, go back to step 2.

| **Note** |
| --- |
| Because the demand-dial connection is a point-to-point connection, the **Gateway IP address** field is not configurable. The adding of static routes can also be done during the creation of the demand-dial interface with the New Demand-Dial Interface Wizard. |

**Performing Auto-static Updates on Each VPN Router**
If RIP for IP is enabled on the demand-dial interfaces of both VPN routers, you can use auto-static updates to automatically configure static routes when the VPN connection is in a connected state. A demand-dial interface that is configured for auto-static updates sends a request across an active connection to request all of the routes of the router on the other side of the connection. In response to the request, all of the routes of the requested router are automatically entered as static routes in the routing table of the requesting router.

From the Routing and Remote Access snap-in on a VPN router (assuming the site-to-site VPN connection is active), perform the following steps:

1. In the console tree, click **IP Routing**, and then click **General**.
2. In the details pane, right-click the appropriate demand-dial interface, and then click **Update Routes**.

You can also use the **netsh** commands to perform an auto-static update. For more information, see the topic "Scheduling auto-static updates" in Windows Server 2003 Help and Support.

**Configuring Routing Protocols**

If the site-to-site VPN connection is persistent (always active), you can also configure IP routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) to operate over the VPN connection. For more information, see the topics titled "Setting up a RIP-for-IP routed internetwork" and "Setting up an OSPF routed internetwork" in Windows Server 2003 Help and Support.

# Appendix A: Configuring Firewalls for a VPN Router Running Windows Server 2003

Updated: October 7, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following are common configurations of firewalls with a VPN router:

- The VPN router is attached to the Internet and the firewall is between the VPN router and the site.
- The firewall is attached to the Internet and the VPN router is between the firewall and the site.
- Two firewalls are usedone between the VPN router and the site and one between the VPN router and the Internet.

**VPN router in Front of the Firewall**

To secure the VPN router from sending or receiving any traffic on its Internet interface except VPN traffic, you need to configure PPTP or L2TP/IPSec input and output filters on the interface that corresponds to the connection to the Internet. Because IP routing is enabled on the Internet interface, if PPTP or L2TP/IPSec filters are not configured on the Internet interface, then any traffic received on the Internet interface is routed, which may forward unwanted Internet traffic to your site.

When the VPN router is in front of the firewall attached to the Internet, you need to add packet filters to the Internet interface that allow only VPN traffic to and from the IP address of the VPN router's Internet interface.

For inbound traffic, when the tunneled data is decrypted by the VPN router, it is forwarded to the firewall. The firewall in this configuration is acting as a filter for site traffic and can prevent specific resources from being accessed, scan data for viruses, perform intrusion detection, and other functions.

Because the only Internet traffic allowed on the site must pass through the VPN router, this approach also prevents the sharing of File Transfer Protocol (FTP) or Web site resources with non-VPN Internet users.

Figure 3 shows the VPN router in front of the firewall.




**Figure 3: The VPN router in front of the firewall**

The firewall is configured for the appropriate rules for site traffic to and from hosts in other sites according to your network security policies.

For the Internet interface on the VPN router, configure the following input and output filters using the Routing and Remote Access snap-in. These filters are automatically configured when you run the Routing and Remote Access Server Setup Wizard and choose the **Remote access (dial-up or VPN)** option, select the correct interface, and select the **Enable security on the selected interface by setting up packet filters** option on the VPN Connection page (enabled by default).

**Packet Filters for PPTP**

Configure the following input filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Destination IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and TCP destination port of 1723.

  This filter allows PPTP tunnel maintenance traffic to the VPN router.

- Destination IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.

  This filter allows PPTP tunneled data to the VPN router.

- Destination IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and TCP [established] source port of 1723.

  This filter is required only when the VPN router is acting as a calling router. With the TCP [established] filter, traffic is accepted only when the VPN router initiated the TCP connection.

Configure the following output filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Source IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and TCP source port of 1723.

  This filter allows PPTP tunnel maintenance traffic from the VPN router.

- Source IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.

  This filter allows PPTP tunneled data from the VPN router.

- Source IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and TCP destination port of 1723.

  This filter is required only when the VPN router is acting as a calling router.

**Packet Filters for L2TP/IPSec**

Configure the following input filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Destination IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 500.

  This filter allows Internet Key Exchange (IKE) traffic to the VPN router.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 4500.

  This filter allows IPSec NAT-T traffic to the VPN router.

- Destination IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 1701.

  This filter allows L2TP traffic to the VPN router.

Configure the following output filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Source IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 500.

  This filter allows IKE traffic from the VPN router.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 5500.

  This filter allows IPSec NAT-T traffic from the VPN router.

- Source IP address of the VPN router's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 1701.

This filter allows L2TP traffic from the VPN router.

There are no filters required for IPSec Encapsulating Security Protocol (ESP) traffic for the IP protocol of 50. The Routing and Remote Access service filters are applied after the IPSec components remove the ESP header.

**VPN Router Behind the Firewall**

In a more common configuration, the firewall is connected to the Internet and the VPN router is a site resource that is connected to the perimeter network, also known as a screened subnet. The perimeter network is an IP network segment that contains resources that are available to Internet users, such as Web and FTP servers. The VPN router has an interface on both the perimeter network and the site. In this approach, the firewall must be configured with input and output filters on its Internet interface that allow the passing of tunnel maintenance traffic and tunneled data to the VPN router. Additional filters can allow the passing of traffic to Web, FTP, and other types of servers on the perimeter network. For an added layer of security, the VPN router can also be configured with PPTP or L2TP/IPSec packet filters on its perimeter network interface.

The firewall in this configuration is acting as a filter for Internet traffic and can confine the incoming and outgoing traffic to the specific resources on the perimeter network, perform intrusion attempt detection, prevent denial of service attacks, and other functions.

Because the firewall does not have the encryption keys for each VPN connection, it can only filter on the plaintext headers of the tunneled data. In other words, all tunneled data passes through the firewall. This is not a security concern, however, because the VPN connection requires an authentication process that prevents unauthorized access beyond the VPN router.

Figure 4 shows the VPN router behind the firewall on the perimeter network.




**Figure 4: The VPN router behind the firewall on the perimeter network**

For both the Internet and network perimeter interfaces on the firewall, configure the following input and output filters using the firewall's configuration software.

**Packet Filters for PPTP**

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

**Filters on the Internet Interface**

Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN router's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter allows PPTP tunnel maintenance traffic to the VPN router.

- Destination IP address of the VPN router's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data to the VPN router.

- Destination IP address of the VPN router's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter is required only when the VPN router is acting as a calling router. This filter should only be used in conjunction with PPTP packet filters described in "VPN router in front of the firewall" and configured on the VPN router's network

perimeter interface. By allowing all traffic to the VPN router from TCP port 1723, there exists the possibility of network attacks from sources on the Internet that use this port.

Configure the following output filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN router's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter allows PPTP tunnel maintenance traffic from the VPN router.

- Source IP address of the VPN router's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data from the VPN router.

- Source IP address of the VPN router's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter is required only when the VPN router is acting as a calling router. This filter should only be used in conjunction with PPTP packet filters described in "VPN router in front of the firewall" and configured on the VPN router's network perimeter interface. By allowing all traffic from the VPN router to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

**Filters on the Perimeter Network Interface**

Configure the following input filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN router's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter allows PPTP tunnel maintenance traffic from the VPN router.

- Source IP address of the VPN router's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data from the VPN router.

- Source IP address of the VPN router's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter is required only when the VPN router is acting as a calling router. This filter should only be used in conjunction with PPTP packet filters described in "VPN router in front of the firewall" and configured on the VPN router's network perimeter interface. By allowing all traffic from the VPN router to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN router's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter allows PPTP tunnel maintenance traffic to the VPN router.

- Destination IP address of the VPN router's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data to the VPN router.

- Destination IP address of the VPN router's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter is required only when the VPN router is acting as a calling router. This filter should only be used in conjunction with PPTP packet filters described in "VPN router in front of the firewall" and configured on the VPN router's network perimeter interface. By allowing all traffic to the VPN router from TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

**Packet Filters for L2TP/IPSec**

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

**Filters on the Internet Interface**

Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN router's perimeter network interface and UDP destination port of 500 (0x1F4).

  This filter allows IKE traffic to the VPN router.

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194).

  This filter allows IPSec NAT-T traffic to the VPN router.

- Destination IP address of the VPN router's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic to the VPN router.

Configure the following output packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN router's perimeter network interface and UDP source port of 500 (0x1F4).

  This filter allows IKE traffic from the VPN router.

- Source IP address of the VPN router's perimeter network interface and UDP source port of 4500 (0x1194).

  This filter allows IPSec NAT-T from the VPN router.

- Source IP address of the VPN router's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic from the VPN router.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

**Filters on the Perimeter Network Interface**

Configure the following input packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN router's perimeter network interface and UDP source port of 500 (0x1F4).

  This filter allows IKE traffic from the VPN router.

- Source IP address of the VPN router's perimeter network interface and UDP source port of 4500 (0x1194).

  This filter allows IPSec NAT-T traffic from the VPN router.

- Source IP address of the VPN router's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic from the VPN router.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN router's perimeter network interface and UDP destination port of 500 (0x1F4).

  This filter allows IKE traffic to the VPN router.

- Destination IP address of the VPN router's perimeter network interface and UDP destination port of 4500 (0x1194).

This filter allows IPSec NAT-T traffic to the VPN router.

- Destination IP address of the VPN router's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic to the VPN router.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

**VPN Router Between Two Firewalls**

Another configuration is when the VPN router computer in placed on the perimeter network between two firewalls. The Internet firewall, the firewall between the Internet and the VPN router, filters all Internet traffic from all Internet clients. The site firewall, the firewall between the VPN router and the site, filters site traffic from hosts in other sites.

Figure 5 shows the VPN router between two firewalls on the perimeter network.




**Figure 5: The VPN router between two firewalls on the perimeter network**

In this configuration:

- Configure your Internet firewall and VPN router with the packet filters as described in the "VPN router behind the firewall" section.

- Configure your site firewall for the appropriate rules for site traffic to and from calling routers according to your network security policies.

# Appendix B: Alternate Configurations -- Virtual Private Networking

Updated: October 7, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides information about common alternate configurations for a Windows Server 2003 VPN router. The most common configuration is described in the "Deploying an L2TP-based Site-to-Site VPN Connection " and "Deploying an L2TP-based Site-to-Site VPN Connection " sections of this paper and whose principal characteristics are the following:

- The VPN router has multiple network adaptersat least one connected to the site and at least one connected to the Internet.
- The VPN router has static public IP addresses assigned to its Internet interfaces.
- The VPN router is only acting as a security gateway providing a routed connection to the site. The VPN router is not hosting any other Internet services such as NAT or Web services.

The two other most common configurations are the following:

1. The VPN router computer is performing other functions such as network address translation or Web hosting.
2. The VPN router computer has a single network adapter and its public IP address is published by a firewall.

The following sections detail the changes to make in the deployment of a VPN router to accommodate these additional common configurations.

**Multiple Internet Function VPN Router**
In this configuration, the VPN router's principal characteristics are the following:

- The VPN router has multiple network adaptersat least one connected to the site and at least one connected to the Internet.
- The VPN router has static public IP addresses assigned to its Internet interfaces.
- The VPN router is acting as a security gateway providing remote access to the site and is hosting any other Internet services such as NAT or Web hosting.

In this configuration, you can follow the procedures as described in the "Deploying a PPTP-based Site-to-Site VPN Connection" and "Deploying a PPTP-based Site-to-Site VPN Connection" sections of this paper except that when you run the Routing and Remote Access Server Setup Wizard, clear the **Enable security on the selected interface by setting up static packet filters** check box on the **VPN Connection** page of the Wizard.

When you clear the **Enable security on the selected interface by setting up static packet filters** check box, PPTP and L2TP/IPSec packet filters are not configured on the Internet interface of the VPN router computer. Whether you have to manually configure these filters depends on whether the VPN router computer is also hosting NAT.

- If NAT is needed on the VPN router computer, do not configure PPTP and L2TP/IPSec packet filters or packet filters for other types of traffic. If you configure PPTP and L2TP/IPSec packet filters on the Internet interface, NAT cannot function. Even though you do not configure any packet filters on the Internet interface of the VPN router computer, the function of the NAT discards any traffic from the Internet that does not correspond to traffic requested by site clients.
- If NAT is not needed on the VPN router computer, you can configure PPTP and L2TP/IPSec packet filters and other types of filters for additional services hosted by the VPN router computer. For example, if the VPN router computer is also hosting a Web site, then filters should be added to allow traffic to and from the public IP address of the VPN router computer and TCP port 80.

**Single-Adapter VPN Router**

In this configuration, the VPN router computer has only a single network adapter and nodes on the site of the calling router are accessing services hosted on the VPN router computer. If the VPN router computer has only a single network adapter and is configured with a public IP address, all traffic to and from the services running on the VPN router computer are sent as clear text outside the VPN tunnel. For more information about why this happens, see "Routing and multi-use VPN routers" in this paper.

The only way a single adapter VPN router can work properly is if it is behind a firewall that is providing a publishing and translation service for the VPN router. The firewall publishes or makes known on the Internet a static public IP address for the VPN router. When VPN packets are sent to this published IP address, the firewall translates the address of the packet to a private or other public address by which the VPN router is known on the site.

Figure 6 shows an example of the published and actual addresses of a VPN router in this configuration.




**Figure 6: The single-adapter VPN router configuration**

The VPN router is configured according to "Deploying a PPTP-based Site-to-Site VPN Connection" in this paper with its site interface acting as an Internet interface. The firewall is configured to:

- Publish the name and public IP address of the VPN router on the Internet.
- Translate PPTP traffic sent to the public IP address of the VPN router to the site interface of the VPN router computer.
- Discard all traffic except PPTP traffic going to and from the VPN router computer.

# Appendix C: Troubleshooting

Updated: October 7, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting Tools**
Windows Server 2003 provides the following tools to troubleshoot VPN connections:

- TCP/IP Troubleshooting Tools
- Authentication and Accounting Logging
- Unreachability Reason
- Event Logging
- IAS Event Logging
- PPP Logging
- Tracing
- Oakley Logging
- Network Monitor

**TCP/IP Troubleshooting Tools**
The Ping, Tracert, and Pathping tools use ICMP Echo and Echo Reply messages to verify connectivity, display the path to a destination, and test path integrity. The **route print** command can be used to display the IP routing table. Alternately, you can use the **netsh routing ip show rtmroutes** command or the Routing and Remote Access snap-in.

In addition to the normal TCP/IP tools, use the Netdiag tool to test and display your network configuration.

**Authentication and Accounting Logging**
A VPN router running Windows Server 2003 supports the logging of authentication and accounting information for VPN connections in local logging files when Windows authentication or Windows accounting is enabled. This logging is separate from the events recorded in the system event log. You can use the information that is logged to track site-to-site connection usage and authentication attempts. Authentication and accounting logging is especially useful for troubleshooting remote access policy issues. For each authentication attempt, the name of the remote access policy that either accepted or rejected the connection attempt is recorded.

Enable authentication and accounting logging from the **Settings** tab on the properties of the **Local File** object in the **Remote Access Logging** folder in the Routing and Remote Access snap-in (if the Routing and Remote Access service is configured for Windows authentication and accounting) or the Internet Authentication Service snap-in (if the Routing and Remote Access service is configured for RADIUS authentication and accounting and the RADIUS server is an IAS server)

The authentication and accounting information is stored in a configurable log file or files stored in the *SystemRoot*\System32\LogFiles folder. The log files are saved in Internet Authentication Service (IAS) or database-compatible format, meaning that any database program can read the log file directly for analysis.

If the VPN router is configured for RADIUS authentication and accounting and the RADIUS server is a computer running Windows Server 2003 and IAS, the authentication and accounting logs are stored in the *SystemRoot*\System32\LogFiles folder on the IAS server computer.

The Routing and Remote Access service and IAS for Windows Server 2003 can also send authentication and accounting information to a Structured Query Language (SQL) server database.

**Unreachability Reason**
When a demand-dial interface fails to make a connection, the interface is left in an unreachable state and the Routing and

Remote Access service records the reason why the connection attempt failed. To view the unreachable reason in the Routing and Remote Access snap-in, click **Network Interfaces**. In the details pane, right-click the demand-dial interface, and then click **Unreachability Reason**.

#### Event Logging

On the **Logging** tab in the properties of a VPN router in the Routing and Remote Access snap-in, there are four levels of logging. Select **Log all events**, and then try the connection again. After the connection fails, check the system event log for events logged during the connection process. After you are done viewing remote access events, select the **Log errors and warnings** option on the **logging** tab to conserve system resources.

#### IAS Event Logging

If your VPN routers are configured for RADIUS authentication and your RADIUS servers are computers running Windows Server 2003 and IAS, check the system event log for IAS events for rejected or accepted connection attempts. IAS system event log entries contain a lot of information about the connection attempt including the name of the remote access policy that accepted or rejected the connection attempt. IAS event logging for rejected or accepted connection attempts is enabled by default and configured from the **Service** tab from the properties of an IAS server in the Internet Authentication Service snap-in.

#### PPP logging

PPP logging records the series of programming functions and PPP control messages during a PPP connection and is a valuable source of information when you are troubleshooting the failure of a PPP connection. To enable PPP logging, select the **Log additional Routing and Remote Access information** option on the **Logging** tab on the properties of a remote access server.

The PPP log in Windows NT 4.0 has been replaced by the tracing function. To duplicate the PPP log, you need to enable file tracing for the PPP key. By default, the PPP log is stored as the Ppp.log file in the SystemRoot\Tracing folder.

#### Tracing

The Windows Server 2003 Routing and Remote Access service has an extensive tracing capability that you can use to troubleshoot complex network problems. You can enable the components in Windows Server 2003 to log tracing information to files using the Netsh command or through the registry.

#### Enabling Tracing with Netsh

You can use the Netsh command to enable and disable tracing for specific components or for all components. To enable and disable tracing for a specific component, use the following syntax:

**netsh ras set tracing** *Component* **enabled|disabled**

where Component is a component in the list of Routing and Remote Access service components found in the registry under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing. For example, to enable tracing for the RASAUTH component, the command is:

**netsh ras set tracing rasauth enabled**

To enable tracing for all components, use the following command:

**netsh ras set tracing** * **enabled**

#### Enabling Tracing through the Registry

You can also configure the tracing function by changing settings in the registry under:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing

You can enable tracing for each Routing and Remote Access service component by setting the registry values described later. You can enable and disable tracing for components while the Routing and Remote Access service is running. Each component is capable of tracing and appears as a subkey under the preceding registry key.

To enable tracing for each component, you can configure the following registry entries for each protocol key:

**EnableFileTracing** REG_DWORD *Flag*

You can enable logging tracing information to a file by setting **EnableFileTracing** to 1. The default value is 0.

**FileDirectory** REG_EXPAND_SZ *Path*

You can change the default location of the tracing files by setting **FileDirectory** to the path you want. The file name for the log file is the name of the component for which tracing is enabled. By default, log files are placed in the *SystemRoot*\Tracing folder.

**FileTracingMask** REG_DWORD *LevelOfTracingInformationLogged*

**FileTracingMask** determines how much tracing information is logged to the file. The default value is 0xFFFF0000.

**MaxFileSize** REG_DWORD *SizeOfLogFile*

You can change the size of the log file by setting different values for **MaxFileSize**. The default value is 0x10000 (64K).

| **Note** |
| --- |
| Tracing consumes system resources and should be used sparingly to help identify network problems. After the trace is captured or the problem is identified, you should immediately disable tracing. Do not leave tracing enabled on multiprocessor computers. Tracing information can be complex and very detailed. Most of the time this information is useful only to Microsoft support professionals or to network administrators who are very experienced with the Routing and Remote Access service. Tracing information can be saved as files and sent to Microsoft support for analysis. |

**Oakley Logging**

You can use the Oakley log to view details about the IPSec security association (SA) negotiation process. The Oakley log is enabled in the registry. It is not enabled by default. To enable the Oakley log, set the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLogging registry setting to 1. The Oakley key does not exist by default and must be created.

After it is enabled, the Oakley log, which is stored in the *SystemRoot*\Debug folder, records all IPSec SA negotiations. A new Oakley.log file is created each time the IPSec Policy Agent is started and the previous version of the Oakley.log file is saved as Oakley.log.sav.

To activate the new EnableLogging registry setting after modifying its value, stop and start the IPSec Policy Agent and related IPSec services by running the following sequence of commands:

1. Stop the Routing and Remote Access service using the **net stop remoteaccess** command.
2. Stop the IPSec services using the **net stop policyagent** command.
3. Start the IPSec services using the **net start policyagent** command.
4. Start the Routing and Remote Access service using the **net start remoteaccess** command.

**Network Monitor**

Use Network Monitor, a packet capture and analysis tool supplied with Windows Server 2003, to capture and view the traffic sent between VPN routers during the VPN connection process and during data transfer. You cannot interpret the encrypted portions of VPN traffic with Network Monitor. Network Monitor is installed as an optional management and monitoring tool when you select Add/Remove Windows Components from Control Panel-Add or Remove Programs.

The proper interpretation of the VPN traffic with Network Monitor requires an in-depth understanding of PPP, PPTP, IPSec, and other protocols. Network Monitor captures can be saved as files and sent to Microsoft support for analysis.

**Troubleshooting Site-to-Site VPN Connections**

Site-to-site VPN problems typically fall into the following categories:

- Connection attempt is rejected when it should be accepted.
- L2TP/IPSec authentication issues.
- EAP-TLS authentication issues.
- Connection attempt is accepted when it should be rejected.
- Unable to reach locations beyond the VPN router.
- Unable to reach the virtual interfaces of VPN routers.

- On-demand connection is not made automatically.
- Unable to establish tunnel.

Use the following troubleshooting tips to isolate the configuration or infrastructure problem causing the stated VPN problem.

**Connection Attempt is Rejected When it Should be Accepted**

- Verify that the credentials of the calling router, consisting of user name, password, and domain name, are correct and can be validated by the answering router.
- Verify that the user account of the calling router is not locked out, expired, disabled, or that the time the connection is being made does not correspond to the configured logon hours.
- Verify that the user account of the calling router is not configured to change its password at the next logon or if the password has expired. A calling router cannot change an expired password during the connection process and the connection attempt is rejected.
- Verify that the user account has not been locked out due to remote access account lockout. For more information, see the topic titled "Account lockout" in Windows Server 2003 Help and Support.
- Verify that the Routing and Remote Access service is running on the answering router.
- Verify that the answering router is enabled for LAN and demand-dial routing from the **General** tab on the properties of a VPN router in the Routing and Remote Access snap-in.
- Verify that the **WAN Miniport (PPTP)** and **WAN Miniport (L2TP)** devices are enabled for demand-dial routing connections (inbound and outbound) from the properties of the **Ports** object in the Routing and Remote Access snap-in.
- Verify that the calling router, the answering router, and the remote access policy corresponding to site-to-site VPN connections are configured to use at least one common authentication method.
- Verify that the calling router and the remote access policy corresponding to site-to-site VPN connections are configured to use at least one common encryption strength. If the calling router is not capable of 128-bit encryption and the **Strongest encryption** level is required in the remote access policy, the connection attempt is rejected.
- Verify that the parameters of the connection are authorized through remote access policies.

  In order for the connection to be established, the parameters of the connection attempt must:

  - Match all of the conditions of at least one remote access policy.
  - Be granted remote access permission through the user account (set to **Allow access**), or if the user account has the **Control access through Remote Access Policy** option selected, the remote access permission of the matching remote access policy must have the **Grant remote access permission** option selected.
  - Match all the settings of the profile.
  - Match all the settings of the dial-in properties of the user account.

  To obtain the name of the remote access policy that rejected the connection attempt, scan the accounting log for the entry corresponding to the connection attempt for the policy name.

- If you are logged on using an account with domain administrator permissions when you run the Routing and Remote Access Server Setup Wizard, it automatically adds the computer account of the **RAS and IAS Servers** domain-local security group. This group membership allows the answering router computer to access user account information. If the answering router is unable to access user account information, verify that:

- The computer account of the answering router computer is a member of the **RAS and IAS Servers** security group for all the domains that contain user accounts for which the answering router is authenticating remote access. You can use the **netsh ras show registeredserver** command at the command prompt to view the current registration. You can use the netsh ras add registeredserver command to register the server in a domain in which the answering router is a member or other domains. Alternately, you or your domain administrator can add the computer account of the answering router computer to the **RAS and IAS Servers** security group of all the domains that contain user accounts for which the answering router is authenticating site-to-site VPN connections.

- If you add or remove the answering router computer to the **RAS and IAS Servers** security group, the change does not take effect immediately (due to the way that Windows Server 2003 caches Active Directory information). For the change to take effect immediately, you need to restart the answering router computer.

- For demand-dial connections using EAP-TLS and Router (Offline request) certificates, verify that the calling router and answering router are correctly configured.

  On the calling router, verify that EAP is configured as the authentication protocol in the advanced security properties of the demand-dial interface. Verify the settings of the properties of the Smart Card or other Certificate (encryption enabled) EAP type. Verify that the correct Router (Offline request) certificate is selected when configuring the credentials of the demand-dial interface.

  On the answering router, verify that EAP is enabled as an authentication method on the answering router and EAP-TLS is enabled on the matching remote access policy. Verify that the correct computer certificate of the authenticating server (the answering router or IAS server) is selected from the configuration settings of the Smart Card or other Certificate EAP type in the remote access policy for site-to-site VPN connections.

  Verify that the user certificate of the calling router was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the answering router trusts. Additionally, verify that the computer certificate of the answering router was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the calling router trusts.

  By default, an answering router checks the certificate revocation list (CRL) when authenticating the calling router. If the CRL is locally available, it can be checked. In some configurations, the CRL cannot be checked until after the connection is made. The CRL is stored at the root CA and, optionally, in Active Directory. For a branch office router that is acting as an answering router in a site that does not contain the root CA, there are two solutions to this problem:

  1. Publish the CRL in Active Directory. For more information, see the topics titled "Schedule the publication of the certificate revocation list" or "To manually publish the certificate revocation list" in Windows Server 2003 Help and Support. Once the CRL is published in Active Directory, the local domain controller in the site will have the latest CRL after Active Directory synchronization.

  2. On the branch office router, set the following registry value to 1:

     HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RasMan\PPP\EAP\13\IgnoreRevocationOffline

- For an answering router that is a member server in a mixed-mode or native-mode domain that is configured for Windows authentication, verify that:

  - The **RAS** and **IAS Servers** security group exists. If not, then create the group and set the group type to **Security** and the group scope to **Domain local**.

  - The **RAS** and **IAS Servers** security group has Read permission to the **RAS and IAS Servers Access Check** object.

- Verify that IP is enabled for routing on both the calling router and answering router.

- Verify that all of the PPTP or L2TP ports on the calling router and answering router are not already being used. If necessary, change the number of PPTP to L2TP ports from the properties of the **Ports** object in the Routing and Remote Access snap-in to allow more concurrent connections.

- Verify that the answering router supports the tunneling protocol of the calling router.

  By default, a Windows Server 2003 demand-dial interface with the VPN type set to **Automatic** will try to establish a PPTP-based VPN connection first, then try an L2TP/IPSec-based VPN connection. If either the **Point to Point Tunneling Protocol (PPTP)** or **Layer-2 Tunneling Protocol (L2TP)** server type option is selected, verify that the selected tunneling protocol is supported by the answering router.

  Depending on your selections when running the Routing and Remote Access Server Setup Wizard, a Windows Server 2003based computer running the Routing and Remote Access service is a PPTP and L2TP server with five or 128 L2TP ports and five or 128 PPTP ports. To create a PPTP-only server, set the number of L2TP ports to zero. To create an L2TP-only server, set the number of PPTP ports to 1 and disable remote access inbound connections and demand-dial connections for the **WAN Miniport (PPTP)** device from the properties of the **Ports** object in the Routing and Remote Access snap-in.

- Verify the configuration of the authentication provider. The answering router can be configured to use either Windows or RADIUS to authenticate the credentials of the calling router.

  - For RADIUS authentication, verify that the answering router computer can communicate with the RADIUS server.

  - For an answering router that is a member of a native-mode domain, verify that the answering router has joined the domain.

  - For a computer Windows NT version 4.0 Service Pack 4 and later with RRAS server that is a member of a mixed mode domain or a Windows Server 2003 answering router that is a member of a Windows NT 4.0 domain that is accessing user account properties for a user account in a trusted domain, verify that the **Everyone** group is added to the **Pre-Windows 2000 Compatible Access** group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, issue the net **localgroup "Pre-Windows 2000 Compatible Access" everyone /add** command on a domain controller computer and then restart the domain controller.

  - For a Windows NT version 4.0 Service Pack 3 and earlier RRAS server that is a member of a mixed-mode domain, verify that the **Everyone** group has been granted list contents, read all properties, and read permissions to the root node of your domain and all sub-objects of the root domain.

- For PPTP connections using MS-CHAP and attempting to negotiate 40-bit MPPE encryption, verify that the password of the calling router is not larger than 14 characters.

**L2TP/IPSec Authentication Issues**

The most common problems that cause site-to-site L2TP/IPSec connections to fail are the following:

- No certificate

  By default, site-to-site L2TP/IPSec connections require that the calling and answering router exchange computer certificates for IPSec peer authentication. Check the Local Computer certificate stores of both the calling and answering router using the Certificates snap-in to ensure that a suitable certificate exists.

- Incorrect certificate

  If certificates exist, they must be verifiable. Unlike manually configuring IPSec rules, the list of certification authorities (CAs) for L2TP/IPSec connections is not configurable. Instead, each router in the L2TP/IPSec connection sends a list of root CAs to its IPSec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Router A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPSec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPSec peer, Router B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPSec security negotiation fails.

The calling router must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the answering router trusts. Additionally, the answering router must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the calling router trusts.

By default, L2TP/IPSec connections require that the calling and answering routers exchange computer certificates for IPSec peer authentication. Check the Local Computer certificate stores of both the calling and answering routers using the Certificates snap-in to ensure that a suitable certificate exists.

- A NAT between the calling and answering routers

  If either the calling or answering router is behind a NAT, then both the calling and answering routers must support IPSec NAT-T. However, Microsoft recommends that answering routers, such as VPN routers running Windows Server 2003, not be placed behind NATs.

- A firewall between the calling and answering routers

  If there is a firewall between the calling and answering router and you cannot establish an L2TP/IPSec connection, verify that the firewall allows L2TP/IPSec traffic to be forwarded. For more information, see Appendix A.

One of the best tools for troubleshooting IPSec authentication issues is the Oakley log. For more information, see "Oakley Logging" in this article.

**EAP-TLS Authentication Issues**

When EAP-TLS is used for authentication, the calling router submits a Router (Offline request) user certificate and the authenticating server (the answering router or the RADIUS server) submits a computer certificate.

In order for the authenticating server to validate the certificate of the calling router, the following must be true for each certificate in the certificate chain sent by the calling router:

- The current date must be within the validity dates of the certificate.

  When certificates are issued, they are issued with a range of valid dates, before which they cannot be used and after which they are considered expired.

- The certificate has not have been revoked.

  Issued certificates can be revoked at any time. Each issuing CA maintains a list of certificates that should no longer be considered valid by publishing an up-to-date certificate revocation list (CRL). By default, the authenticating server checks all the certificates in the calling routers certificate chain (the series of certificates from the calling router certificate to the root CA) for revocation. If any of the certificates in the chain have been revoked, certificate validation fails. This behavior can be modified with registry settings described later in this topic.

  To view the CRL distribution points for a certificate in the Certificates snap-in, obtain the certificate properties, click the **Details** tab, and then click the **CRL Distribution Points** field.

  The certificate revocation validation only works as well as the CRL publishing and distribution system. If the CRL in a certificate is not updated often, a certificate that has been revoked can still be used and considered valid because the published CRL that the authenticating server is checking is out of date.

- The certificate has a valid digital signature.

  CAs digitally sign certificates they issue. The authenticating server verifies the digital signature of each certificate in the chain, with the exception of the root CA certificate, by obtaining the public key from the certificates issuing CA and mathematically validating the digital signature.

The calling router certificate must also have the Client Authentication certificate purpose (also known as Enhanced Key Usage

[EKU]) (object identifier 1.3.6.1.5.5.7.3.2) and must either contain a UPN of a valid user account or FQDN of valid computer account for the Subject Alternative Name property of the certificate.

To view the EKU for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the **Details** tab, and then click the **Enhanced Key Usage** field. To view the subject alternative name property for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the **Details** tab, and then click the **Subject Alternative Name** field.

Finally, to trust the certificate chain offered by the calling router, the authenticating server must have the root CA certificate of the issuing CA of the calling router certificate installed in its Trusted Root Certification Authorities store.

Additionally, the authenticating server verifies that the identity sent in the EAP-Response/Identity message is the same as the name in the Subject Alternative Name property of the certificate. This prevents a malicious user from masquerading as a different user from that specified in the EAP-Response/Identity message.

If the authenticating server is a Windows Server 2003 answering router or an IAS server, the following registry settings in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13 can modify the behavior of the EAP-TLS when performing certificate revocation:

- IgnoreNoRevocationCheck

  When set to 1, the authenticating server allows EAP-TLS clients to connect even when it does not perform or cannot complete a revocation check of the calling router's certificate chain (excluding the root certificate). Typically, revocation checks fail because the certificate doesn't include CRL information.

  IgnoreNoRevocationCheck is set to 0 (disabled) by default. An EAP-TLS client cannot connect unless the server completes a revocation check of the client's certificate chain (including the root certificate) and verifies that none of the certificates have been revoked.

  You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties.

- IgnoreRevocationOffline

  When set to 1, the authenticating server allows EAP-TLS clients to connect even when a server that stores a CRL is not available on the network. IgnoreRevocationOffline is set to 0 by default. The authenticating server does not allow clients to connect unless it can complete a revocation check of their certificate chain and verify that none of the certificates has been revoked. When it cannot connect to a server that stores a revocation list, EAP-TLS considers the certificate to have failed the revocation check.

  Setting IgnoreRevocationOffline to 1 prevents certificate validation failure because poor network conditions prevented their revocation check from completing successfully.

- NoRevocationCheck

  When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the calling router's certificate. The revocation check verifies that the calling routers certificate and the certificates in its certificate chain have not been revoked. NoRevocationCheck is set to 0 by default.

- NoRootRevocationCheck

  When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the calling router's root CA certificate. NoRootRevocationCheck is set to 0 by default. This entry only eliminates the revocation check of the client's root CA certificate. A revocation check is still performed on the remainder of the calling router's certificate chain.

  You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties. Also, this entry can prevent certification-related delays that occur when a certificate revocation list is offline or is expired.

All of these registry settings must be added as a DWORD type and have the valid values of 0 or 1. The calling router does not

use these settings.

In order for the calling router to validate the certificate of the authenticating server for either EAP-TLS authentication, the following must be true for each certificate in the certificate chain sent by the authenticating server:

- The current date must be within the validity dates of the certificate.

  When certificates are issued, they are issued with a range of valid dates, before which they cannot be used and after which they are considered expired.

- The certificate has a valid digital signature.

  CAs digitally sign certificates they issue. The calling router verifies the digital signature of each certificate in the chain, with the exception of the root CA certificate, by obtaining the public key from the certificates issuing CA and mathematically validating the digital signature.

Additionally, the authenticating server computer certificate must have the Server Authentication EKU (object identifier 1.3.6.1.5.5.7.3.1). To view the EKU for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the **Details** tab, and then click the **Enhanced Key Usage** field.

Finally, to trust the certificate chain offered by the authenticating server, the calling router must have the root CA certificate of the issuing CA of the authenticating server certificate installed in its Trusted Root Certification Authorities store.

Notice that the calling router does not perform certificate revocation checking for the certificates in the certificate chain of the authenticating server's computer certificate. The assumption is that the calling router does not yet have a connection to the network, and therefore might not have access to a Web page or other resource in order to check for certificate revocation.

**Connection Attempt is Accepted When it Should be Rejected**

- Verify that the remote access permission on the user account is set to either **Deny access** or **Control access through Remote Access Policy**. If set to the latter, verify that the first matching remote access policy's remote access permission is set to **Deny remote access permission**. To obtain the name of the remote access policy that accepted the connection attempt, scan the accounting log for the entry corresponding to the connection attempt for the policy name.

- If you have created a remote access policy to explicitly reject all connections, verify the policy conditions, remote access permission, and profile settings.

- For certificate-based authentication when the certificate has been revoked, changes to the CRL of the CA might not yet be published. The authenticating server is using the old CRL that does not have the recently revoked certificate. To minimize the amount of time between the revoking of a certificate and when the authenticating server has knowledge of the revoked certificate, lower the CRL publication time. On a Windows Server 2003 CA, the publication time can be set as low as one hour. For more information, see the topic titled "Revoking certificates and publishing CRLs" in Windows Server 2003 Help and Support.

**Unable to Reach Locations Beyond the VPN Router**

- Verify that IP routing is enabled (on the **IP** tab on the properties of the VPN router).

- Verify that the appropriate demand-dial interface has been added to the protocol being routed.

- Verify that there are routes in the site routers on the calling router and answering router's site so that all locations on both networks are reachable. You can add routes to the routers of each site through static routes or by enabling a routing protocol on the site interface of the calling and answering routers.

  Unlike a remote access connection, a demand-dial connection does not automatically create a default route. You need to create routes on both sides of the demand-dial connection so that traffic can be routed to and from the other side of the demand-dial connection.

  You can manually add static routes to the routing table, or you can add static routes through routing protocols. For persistent demand-dial connections, you can enable Open Shortest Path First (OSPF) or Routing Information Protocol

(RIP) across the demand-dial connection. For on-demand demand-dial connections, you can automatically update routes through an auto-static RIP update.

- For two-way initiated site-to-site VPN connections, verify that the answering router is not interpreting the site-to-site VPN connection as a remote access connection.

  For two-way initiated connections, either router can be the calling router or the answering router. The user names and demand-dial interface names must be properly matched. For example, two-way initiated connections would work under the following configuration:

  Router 1 has a demand-dial interface called NEW-YORK that is configured to use SEATTLE as the user name when sending authentication credentials.

  Router 2 has a demand-dial interface called SEATTLE that is configured to use NEW-YORK as the user name when sending authentication credentials.

  This example assumes that Router 2 can validate the SEATTLE user name and Router 1 can validate the NEW-YORK user name.

  If the incoming caller is a router, the port on which the call was received shows a status of **Active** and the corresponding demand-dial interface is in a **Connected** state. If the user account name in the credentials of the calling router appears under **Remote Access Clients** in the Routing and Remote Access snap-in, then the answering router has interpreted the calling router as a remote access client.

- For a one-way initiated demand-dial connection, verify that the appropriate static routes are enabled on the user account of the calling router and that the answering router is configured with a routing protocol so that when a connection is made, the static routes of the user account of the calling router are advertised to neighboring routers.

- Verify that there are no IP packet filters on the demand-dial interfaces of the calling router and answering router that prevent the sending or receiving of TCP/IP.

  You can configure each demand-dial interface with IP input and output filters to control the exact nature of TCP/IP traffic that is allowed into and out of the demand-dial interface.

**Unable to Reach the Virtual Interfaces of VPN Routers**

- Verify the IP address pools of the calling router and answering router.

  If the VPN router is configured to use a static IP address pool, verify that the routes to the range of addresses defined by the static IP address pools are reachable by the hosts and routers of the site. If not, then IP route consisting of the VPN router static IP address pools, as defined by the IP address and mask of the range, must be added to the routers of the site or enable the routing protocol of your routed infrastructure on the VPN router. If the routes to the address pool subnets are not present, calling router logical interfaces cannot receive traffic from locations on the site. Routes for the subnets are implemented either through static routing entries or through a routing protocol, such as Open Shortest Path First (OSPF) or Routing Information Protocol (RIP).

  If the VPN router is configured to use DHCP for IP address allocation and no DHCP server is available, the VPN router assigns addresses from the Automatic Private IP Addressing (APIPA) address range from 169.254.0.1 through 169.254.255.254. Assigning APIPA addresses to VPN routers works only if the network to which the VPN router is attached is also using APIPA addresses.

  If the VPN router is using APIPA addresses when a DHCP server is available, verify that the proper adapter is selected from which to obtain DHCP-allocated IP addresses. By default, the VPN router chooses the adapter to use to obtain IP addresses through DHCP based on your selections in the Routing and Remote Access Server Setup Wizard. You can manually choose a LAN adapter from the **Adapter** list on the **IP** tab on the properties of the VPN router in the Routing and Remote Access snap-in.

  If the static IP address pools are a range of IP addresses that are a subset of the range of IP addresses for the network to

which the VPN router is attached, verify that the range of IP addresses in the static IP address pool are not assigned to other TCP/IP nodes, either through static configuration or through DHCP.

#### On-demand Connection is Not Made Automatically

- Verify that IP routing is enabled on the IP tab on the properties of the calling router.
- Verify that the correct static routes exist and are configured with the appropriate demand-dial interface.
- For the static routes that use a demand-dial interface, verify that the **Use this route to initiate demand-dial connections** check box on the properties of the demand-dial interface is selected.
- Verify that the demand-dial interface is not in a disabled state.

  To enable a demand-dial interface that is in a disabled state, right-click the demand-dial interface under **Network Interfaces** in the Routing and Remote Access snap-in, and then click **Enable**.
- Verify that the dial-out hours for the demand-dial interface on the calling router are not preventing the connection attempt.

  To configure dial-out hours, right-click the demand-dial interface under **Network Interfaces** in the Routing and Remote Access snap-in, and then click **Dial-out Hours**.
- Verify that the demand-dial filters for the demand-dial interface on the calling router are not preventing the connection attempt.

  To configure demand-dial filters, right-click the demand-dial interface under **Network Interfaces** in the Routing and Remote Access snap-in, and then click **Set IP Demand-dial Filters**.

#### Unable to Establish Tunnel

- Verify that packet filtering on a router interface between the calling router and the answering router is not preventing the forwarding of tunneling protocol traffic. See Appendix A for information about the types of traffic that must be allowed for PPTP and L2TP/IPSec traffic.

  On a Windows Server 2003based VPN router, IP packet filtering can be separately configured from the advanced TCP/IP properties and from the Routing and Remote Access snap-in. Check both places for filters that might be excluding VPN connection traffic.
- Verify that the Winsock Proxy client is not currently running on the VPN router.

  When the Winsock Proxy client is active, Winsock API calls such as those used to create tunnels and send tunneled data are intercepted and forwarded to a configured proxy server.

  A proxy server-based computer allows an organization to access specific types of Internet resources (typically Web and FTP) without directly connecting that organization to the Internet. The organization can instead use private IP network IDs (such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16).

  Proxy servers are typically used so that private users in an organization can have access to public Internet resources as if they were directly attached to the Internet. VPN connections are typically used so that authorized public Internet users can gain access to private organization resources as if they were directly attached to the private network. A single computer can act as a proxy server (for private users) and a VPN server (for authorized Internet users) to facilitate both exchanges of information.

# Summary -- Virtual Private Networking

Updated: October 7, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This paper described in detail the components and their associated design decisions for a Windows Server 2003-based site-to-site VPN deployment including VPN routers, Internet infrastructure, authentication protocols, VPN protocols, site infrastructure, AAA infrastructure, and certificate infrastructure. This paper also included detailed walkthroughs of PTPP and L2TP/IPSec-based site-to-site VPN deployments using computers running Windows Server 2003, details of firewall configuration, and a discussion of VPN troubleshooting tools and common VPN problems with suggested solutions.

**Related Links**
See the following resources for further information:

- Windows VPN Web site at http://www.microsoft.com/windows2000/technologies/communications/vpn/default.asp
- Internet Authentication Service Web site at http://www.microsoft.com/windows2000/technologies/communications/ias/default.mspx

For the latest information about Windows Server, see the Windows Server 2003 Web site at http://www.microsoft.com/windowsserver2003/default.mspx.

# Virtual Private Networking with Windows Server 2003: An Example Deployment

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In This White Paper**


Introduction to Virtual Private Networking with Windows Server 2003: An Example Deployment

Common Configuration for the VPN Server (VPN with Windows Server 2003)

VPN Remote Access for Employees (VPN with Windows Server 2003)

On-Demand Branch Office (VPN with Windows Server 2003)

Persistent Branch Office (VPN with Windows Server 2003)

Extranet for Business Partners (VPN with Windows Server 2003)

Dial-up and VPNs with RADIUS Authentication (VPN with Windows Server 2003)

Summary (VPN with Windows Server 2003)

Related Links for Virtual Private Networking with Windows Server 2003: An Example Deployment

# Introduction to Virtual Private Networking with Windows Server 2003: An Example Deployment

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This white paper describes how to configure common virtual private network connections for a fictional company using the Windows Server 2003 and the Windows XP operating systems. Although your network configuration may be different than those described here, you can still apply the basic concepts of virtual private networking in your network environment.

The use of both public and private networks to create a network connection is called a virtual private network (VPN). A VPN is the extension of a private network that encompasses links across shared or public networks like the Internet. With a VPN, you can send data between two computers across a shared or public network in a manner that emulates a point-to-point private link. Virtual private networking is the act of creating and configuring a virtual private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information, which allows the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The link in which the private data is encapsulated and encrypted is a VPN connection.

Electronic, Inc. is a fictional electronics design and manufacturing company with a main corporate campus in New York and branch offices and distribution business partners throughout the United States. Electronic, Inc. has implemented a VPN solution by using Windows Server 2003 to connect remote access users, branch offices, and business partners.

The VPN server at the corporate office provides both remote access and site-to-site (also known as router-to-router) Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol with Internet Protocol security (L2TP/IPSec) VPN connections. In addition, the VPN server provides the routing of packets to intranet and Internet locations.

Based on the common configuration of the VPN server, the following VPN configurations are described:

- VPN remote access for employees.
- On-demand branch office access.
- Persistent branch office access.
- Extranet for business partners.
- Dial-up and VPNs with RADIUS authentication.

| **Note** |
| --- |
| The example companies, organizations, products, people, and events depicted herein are fictitious. No association with any real company, organization, product, person, or event is intended or should be inferred. |

# Common Configuration for the VPN Server (VPN with Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To deploy a VPN solution for Electronic, Inc., the network administrator performs an analysis and makes design decisions regarding:

- The network configuration.
- The remote access policy configuration.
- The domain configuration.
- The security configuration.

**Network Configuration**

The key elements of the network configuration are:

- The Electronic, Inc. corporate intranet uses the private networks of 172.16.0.0 with a subnet mask of 255.240.0.0 (172.16.0.0/12) and 192.168.0.0 with a subnet mask of 255.255.0.0 (192.168.0.0/16). The corporate campus network segments use subnets of 172.16.0.0 and the branch offices use subnets of 192.168.0.0.
- The VPN server computer is directly attached to the Internet using a T3 (also known as a DS-3) dedicated WAN link.
- The IP address of the WAN adapter on the Internet is 207.46.130.1 as allocated by the Internet service provider (ISP) for Electronic, Inc. The IP address of the WAN adapter is referred to on the Internet by the domain name vpn.electronic.example.com.
- The VPN server computer is directly attached to an intranet network segment that contains a router that connects to the rest of the Electronic, Inc. corporate campus intranet. The intranet network segment has the IP network ID of 172.31.0.0 with the subnet mask of 255.255.0.0.
- The VPN server computer is configured with a static pool of IP addresses to allocate to remote access clients and calling routers that is a subset of the intranet network segment (an on-subnet address pool).

Figure 1 shows the network configuration of the Electronic, Inc. VPN server.




**Figure 1: The network configuration of the Electronic, Inc. VPN server**

Based on the network configuration of the Electronic, Inc. corporate campus intranet, the VPN server computer is configured as follows:

1. Install hardware on the VPN server.

   The network adapter that is used to connect to the intranet segment and the WAN adapter that is used to connect to the Internet are installed according to the adapter manufacturer's instructions. Once drivers are installed and functioning,

both adapters appear as local area connections in Network Connections.

2. Configure TCP/IP on the LAN and WAN adapters.

   For the LAN adapter, an IP address of 172.31.0.1 with a subnet mask 255.255.0.0 is configured. For the WAN adapter, an IP address of 207.46.130.1 with a subnet mask 255.255.255.255 is configured. A default gateway is not configured for either adapter. Domain Name System (DNS) and Windows Internet Name Service (WINS) server addresses are also configured.

3. Configure the Routing and Remote Access service.

   The Routing and Remote Access service is initially configured with the Routing and Remote Access Server Setup Wizard. To run the Wizard, right-click the name of the server in the Routing and Remote Access snap-in, and then click **Configure and Enable Routing and Remote Access**. Configure the VPN server using the following settings:

   **Configuration**: Remote access (dial-up or VPN)

   **Remote Access**: VPN

   **VPN Connection**: Click the connection that corresponds to the interface connected to the Internet

   **IP Address Assignment**: Click **From a specified range of addresses** and create a single range from 172.31.255.1 to 172.31.255.254. This creates a static address pool for up to 253 VPN clients.

   **Managing Multiple Remote Access Servers**: Click **No, use Routing and Remote Access to authenticate connection requests**

   The default method of authenticating remote access and demand-dial connections is to use Windows authentication, which is appropriate in this configuration containing only one VPN server. For information about the use of Remote Authentication Dial-In User Service (RADIUS) authentication for Electronic, Inc., see the Dial-up and VPNs with RADIUS section in this paper. For more information about the use of Windows and RADIUS authentication, see the topic titled Authentication vs. Authorization in Windows Server 2003 Help and Support.

4. Configure the DHCP Relay Agent.

   In the console tree, navigate to IP Routing\DHCP Relay Agent. Right-click **DHCP Relay Agent**, and then click **Properties**. In the **DHCP Relay Agent Properties** dialog box, type the IP address of an intranet Dynamic Host Configuration Protocol (DHCP) server in **Server address**. Click **Add**, and then click **OK**. By configuring the DHCP Relay Agent routing protocol component, VPN remote access clients can receive the correct DNS domain name, DNS server addresses, and WINS server addresses when connecting to the intranet.

5. Configure static routes on the VPN server to reach intranet and Internet locations.

   To reach intranet locations, a static route is configured with the following settings:

   - Interface: The LAN adapter attached to the intranet
   - Destination: 172.16.0.0
   - Network mask: 255.240.0.0
   - Gateway: 172.31.0.2
   - Metric: 1

   This static route simplifies routing by summarizing all destinations on the Electronic, Inc. intranet. This static route is used so that the VPN server does not need to be configured with a routing protocol.

To reach Internet locations, a static route is configured with the following settings:

- Interface: The WAN adapter attached to the Internet
- Destination: 0.0.0.0
- Network mask: 0.0.0.0
- Gateway: 0.0.0.0
- Metric: 1

This static route summarizes all destinations on the Internet. This route allows the VPN server to respond to a remote access client or demand-dial router from anywhere on the Internet.

| **Note** |
| --- |
| Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway. The gateway address of 0.0.0.0 is an example. 0.0.0.0 is the unspecified IP address. |

6. Configure a static route on the intranet router to reach all branch offices.

   To reach branch office locations from the intranet router, a static route is configured with the following settings:

   - Interface: The LAN adapter attached to the intranet
   - Destination: 192.168.0.0
   - Network mask: 255.255.0.0
   - Gateway: 172.31.0.1
   - Metric: 1

   This static route simplifies routing by summarizing all destinations at Electronic, Inc. branch offices. The intranet router advertises this static route to its neighboring routers so that a route to the branch office locations exists on each router of the intranet.

**Remote Access Policy Configuration**

Electronic, Inc. is using a native-mode Active Directory domain and the network administrator for Electronic, Inc. has decided on an access-by-group administrative model. The remote access permission on all user accounts is set to **Control access through Remote Access Policy**. The granting of remote access permission to connection attempts is controlled by the remote access permission setting on the first matching remote access policy. Remote access policies are used to apply different VPN connection settings based on group membership.

For more information, see the topic Introduction to remote access policies in Windows Server 2003 Help and Support.

**Domain Configuration**

To take advantage of the ability to apply different connection settings to different types of VPN connections, the following Active Directory groups are created:

- VPN_Users

  Used for remote access VPN connections

- VPN_Routers

  Used for site-to-site VPN connections from Electronic, Inc. branch offices

- VPN_Partners

  Used for site-to-site VPN connections from Electronic, Inc. business partners

| **Note** |
| --- |
| All users and groups in this example deployment are created in the electronic.example.com Active Directory domain. |

**Security Configuration**

To enable L2TP/IPSec connections, the use of smart cards by remote access clients, and the use of EAP-TLS by routers, the Electronic, Inc. domain is configured to autoenroll computer certificates to all domain members.

For more information, see the topic titled "Checklist: Configuring certificate autoenrollment" in Windows Server 2003 Help and Support.

# VPN Remote Access for Employees (VPN with Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Remote access for Electronic, Inc. employees is deployed by using remote access VPN connections across the Internet based on the settings configured in the Common Configuration for the VPN Server section of this paper and the following additional settings.

Figure 2 shows the Electronic, Inc. VPN server that provides remote access VPN connections.




**Figure 2: The Electronic, Inc. VPN server that provides remote access VPN connections**

**Domain Configuration**
For each employee that is allowed VPN access:

- The remote access permission on the dial-in properties of the user account is set to **Control access through Remote Access Policy**.
- The user account is added to the VPN_Users Active Directory group.

**Remote Access Policy Configuration**
To define the authentication and encryption settings for remote access VPN clients, the following common remote access policy is created:

Policy name: Remote Access VPN Connections

Access method: VPN

User or Group Access: Group with the EXAMPLE\VPN_Users group selected

Authentication Methods: **Extensible Authentication Protocol** with the **Smart card or other Certificate** type, **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)**, and **Microsoft Encrypted Authentication (MS-CHAP)** selected

Policy Encryption Level: **Strong encryption** and **Strongest encryption** selected

**PPTP-based Remote Access Client Configuration**
On the Windows XP remote access client computers, the New Connection Wizard is used to create a VPN connection with the following settings:

**Network Connection Type**: Connect to the network at my workplace

**Network Connection**: Virtual Private Network connection

**Connection Name**: Electronic, Inc.

**VPN Server Selection**: vpn.electronic.example.com

**Connection Availability**: Anyone's use

**L2TP/IPSec-based Remote Access Client Configuration**

The remote access computer logs on to the Electronic, Inc. domain using a local area network (LAN) connection to the Electronic, Inc. intranet and receives a computer certificate through autoenrollment. Then, the New Connection Wizard is used to create VPN connection with the following setting:

**Network Connection Type**: Connect to the network at my workplace

**Network Connection**: Virtual Private Network connection

**Connection Name**: Electronic, Inc.

**VPN Server Selection**: vpn.electronic.example.com

**Connection Availability**: Anyone's use

From the **Connect Electronic, Inc.** dialog box, click **Properties**, and then click the **Networking** tab.

On the **Networking** tab, **Type of VPN** is set to **L2TP/IPSec VPN**. When **Type of VPN** is set to **Automatic**, a PPTP connection is tried first. In this case, the network administrator for Electronic, Inc. does not want remote access clients that are capable of establishing an L2TP/IPSec connection to use PPTP.

# On-Demand Branch Office (VPN with Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Portland and Dallas branch offices of Electronic, Inc. are connected to the corporate office by using on-demand site-to-site VPN connections. Both the Portland and Dallas offices contain a small number of employees who only need occasional connectivity with the corporate office. The Windows Server 2003 routers in the Portland and Dallas offices are equipped with an ISDN adapter that dials a local Internet service provider to gain access to the Internet, and then a site-to-site VPN connection is made across the Internet. When the VPN connection is idle for five minutes, the routers at the branch offices terminate the VPN connection.

The Dallas branch office uses the IP network ID of 192.168.28.0 with a subnet mask of 255.255.255.0 (192.168.28.0/24). The Portland branch office uses the IP network ID of 192.168.4.0 with a subnet mask of 255.255.255.0 (192.168.4.0/24).

To simplify the configuration, the VPN connection is a one-way initiated connection that is always initiated by the branch office router. For more information, see the topic titled One-way Initiated Demand-Dial Connections in Windows Server 2003 Help and Support.

Figure 3 shows the Electronic, Inc. VPN server that provides on-demand branch office connections.




**Figure 3: The Electronic, Inc. VPN server that provides on-demand branch office connections**

To deploy on-demand site-to-site VPN connections to connect the Portland and Dallas branch offices to the corporate office based on the settings configured in the Common Configuration for the VPN server section of this paper, the following additional settings are configured.

**Domain Configuration**

For the VPN connection to the Dallas office, the user account VPN_Dallas is created with the following settings:

- Password of nY7W{q8~=z3.
- For the account properties of the VPN_Dallas account, the **User must change password at next logon** option is cleared and **Password never expire**s option is selected.
- For the dial-in properties on the VPN_Dallas account, the remote access permission is set to **Control access through Remote Access Policy** and the static route 192.168.28.0 with a subnet mask of 255.255.255.0 is added.
- The VPN_Dallas account is added to the VPN_Routers group.

For the VPN connection to the Portland office, the user account VPN_Portland is created with the following settings:

- Password of P*4s=wq!Gx1.
- For the account properties of the VPN_Portland account, the **User must change password at next logon** option is cleared and **Password never expires** option is selected.

- For the dial-in properties on the VPN_Portland account, the remote access permission is set to **Control access through Remote Access Policy** and the static route 192.168.4.0 with a subnet mask of 255.255.255.0 is added.

- The VPN_Portland account is added to the VPN_Routers group.

#### Remote Access Policy Configuration

To define the authentication and encryption settings for the VPN routers, the following remote access policy is created:

**Policy name**: VPN Routers

**Access method**: VPN

**User or Group Access**: Group with the EXAMPLE\VPN_Routers group selected

**Authentication Methods**: **Extensible Authentication Protocol** with the **Smart card or other Certificate** type and **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** selected

**Policy Encryption Level**: **Strong encryption** and **Strongest encryption** selected

The following sections describe a PPTP-based on-demand branch office connection for the Dallas office and an L2TP/IPSec-based on-demand branch office connection for the Portland office.

#### PPTP-based On-Demand Branch Office

The Dallas branch office is a PPTP-based branch office that uses a Windows Server 2003 router to create an on-demand, site-to-site VPN connection with the VPN server in New York as needed. When the connection is made and is idle for five minutes, the connection is terminated.

To deploy a PPTP, one-way initiated, on-demand, site-to-site VPN connection to the corporate office based on the settings configured in the Common Configuration for the VPN Server and On-Demand Branch Office sections of this paper, the following settings are configured on the Dallas router.

#### Demand-Dial Interface for the Connection to the ISP

To connect the Dallas office router to the Internet by using a local ISP, a demand-dial interface is created using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: ISP

- **Connection Type**: Connect using a modem, ISDN adapter, or other physical device

- **Select a Device**: The appropriate ISDN device is selected.

- **Phone Number**: Phone number of the ISP for the Dallas office.

- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  To create the connection to the Dallas ISP when the site-to-site VPN connection needs to be made, the following static route is configured:

  **Destination**: 207.46.130.1

  **Network mask**: 255.255.255.255

  **Metric**: 1

- **Dial Out Credentials**

  **User name**: Dallas office ISP account name

  **Password**: Dallas office ISP account password

**Confirm password**: Dallas office ISP account password

To run the Demand-Dial Interface wizard, right-click **Network Interfaces**, and then click **New Demand-Dial Interface**.

**Demand-Dial Interface for Site-to-Site VPN Connection**

To connect the Dallas office router to the VPN server by using a site-to-site VPN connection over the Internet, a demand-dial interface is created by using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: CorpHQ
- **Connection Type**: Connect using virtual private networking (VPN)
- **VPN Type**: Point to Point Tunneling Protocol (PPTP)
- **Destination Address**: 207.46.130.1
- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.
- **Static Routes for Remote Network**s

  To make all locations on the corporate intranet reachable, the following static route is configured:

  **Destination**: 172.16.0.0

  **Network mask**: 255.240.0.0

  **Metric**: 1

  To make all locations on Electronic, Inc. branch offices reachable, the following static route is configured:

  **Destination**: 192.168.0.0

  **Network mask**: 255.255.0.0

  **Metric**: 1

- **Dial Out Credentials**

  **User name**: VPN_Dallas

  **Domain**: electronic.example.com

  **Password**: nY7W{q8~=z3

  **Confirm password**: nY7W{q8~=z3

**L2TP/IPSec-based On-Demand Branch Office**

The Portland branch office is an L2TP/IPSec-based branch office that uses a Windows Server 2003 router to create an on-demand, site-to-site VPN connection with the VPN server in New York as needed. When the connection is made and is idle for five minutes, the connection is terminated.

To deploy an L2TP/IPSec, one-way initiated, on-demand, site-to-site VPN connection to the corporate office based on the settings configured in the Common Configuration for the VPN Server and On-Demand Branch Office sections of this paper, the following settings are configured on the Portland router:

**Certificate Configuration**

The Portland router was configured by the Electronic, Inc. network administrator while it was physically connected to the Electronic, Inc. intranet and then shipped to the Portland site. While the Portland router was connected to the Electronic, Inc.

intranet, a computer certificate was installed through autoenrollment.

#### Demand-Dial Interface for the Connection to the ISP

To connect the Portland office router to the Internet by using a local ISP, a demand-dial interface is created by using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: ISP
- **Connection Type**: Connect using a modem, ISDN adapter, or other physical device
- **Select a Device**: The appropriate ISDN device is selected.
- **Phone Number**: Phone number of the ISP for the Portland office.
- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.
- **Static Routes for Remote Networks**

  To create the connection to the Portland ISP when the site-to-site VPN connection needs to be made, the following static route is configured:

  **Destination**: 207.46.130.1

  **Network mask**: 255.255.255.255

  **Metric**: 1

- **Dial Out Credentials**

  **User name**: Portland office ISP account name

  **Password**: Portland office ISP account password

  **Confirm password**: Portland office ISP account password

#### Demand-Dial Interface for Site-to-Site VPN Connection

To connect the Portland office router to the VPN server by using a site-to-site VPN connection over the Internet, a demand-dial interface is created by using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: CorpHQ
- **Connection Type**: Connect using virtual private networking (VPN)
- **VPN Type**: Layer 2 Tunneling Protocol (L2TP)
- **Destination Address**: 207.46.130.1
- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.
- **Static Routes for Remote Networks**

  To make all locations on the corporate intranet reachable, the following static route is configured:

  **Destination**: 172.16.0.0

  **Network mask**: 255.240.0.0

  **Metric**: 1

To make all locations on Electronic, Inc. branch offices reachable, the following static route is configured:

**Destination**: 192.168.0.0

**Network mask**: 255.255.0.0

**Metric**: 1

- **Dial Out Credentials**

  **User name**: VPN_Portland

  **Domain**: electronic.example.com

  **Password**: P*4s=wq!Gx1

  **Confirm password**: P*4s=wq!Gx1

# Persistent Branch Office (VPN with Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Chicago and Phoenix branch offices of Electronic, Inc. are connected to the corporate office by using persistent site-to-site VPN connections that stay connected 24 hours a day. The Windows Server 2003 routers in the Chicago and Phoenix offices are equipped with T1 WAN adapters that have a permanent connection to a local Internet service provider to gain access to the Internet.

The Chicago branch office uses the IP network ID of 192.168.9.0 with a subnet mask of 255.255.255.0 (192.168.9.0/24). The Chicago branch office router uses the public IP address of 131.107.0.1 for its Internet interface. The Phoenix branch office uses the IP network ID of 192.168.14.0 with a subnet mask of 255.255.255.0 (192.168.14.0/24). The Phoenix branch office router uses the public IP address of 159.60.0.1 for its Internet interface.

The VPN connection is a two-way initiated connection. The connection is initiated from either the branch office router or the VPN server. Two-way initiated connections require the creation of demand-dial interfaces, remote access policies, IP address pools, and packet filters on the routers on both sides of the connection.

Figure 4 shows the Electronic, Inc. VPN server that provides persistent branch office connections.




**Figure 4: The Electronic, Inc. VPN server that provides persistent branch office connections**

To deploy persistent site-to-site VPN connections to connect the Chicago and Phoenix branch offices to the corporate office based on the settings configured in the Common Configuration for the VPN Server section of this paper, the following additional settings are configured.

**Domain Configuration**

For the Chicago office VPN connection that is initiated by the Chicago router, the user account VPN_Chicago is created with the following settings:

- Password of U9!j5dP(%q1.
- For the account properties of the VPN_Chicago account, the **User must change password at next logon** option is cleared and **Password never expires** option is selected.
- For the dial-in properties on the VPN_Chicago account, the remote access permission is set to **Control access through Remote Access Policy**.
- The VPN_Chicago account is added to the VPN_Routers group.

For the Phoenix office VPN connection that is initiated by the Phoenix router, the user account VPN_Phoenix is created with the following settings:

- Password of z2F%s)bW$4f.
- For the account properties of the VPN_Phoenix account, the **User must change password at next logon** option is

cleared and **Password never expires** option is selected.

- For the dial-in properties on the VPN_Phoenix account, the remote access permission is set to **Control access through Remote Access Policy**.

- The VPN_Phoenix account is added to the VPN_Routers group.

For the Chicago office VPN connection and the Phoenix office VPN connection that are initiated by the VPN server, the user account VPN_CorpHQ is created with the following settings:

- Password of o3\Dn6@`-J4.

- For the dial-in properties on the VPN_CorpHQ account, the remote access permission is set to **Control access through Remote Access Policy**.

- The VPN_CorpHQ account is added to the VPN_Routers group.

**Remote Access Policy Configuration**
Remote access policies must be configured at the VPN server, the Chicago router, and the Phoenix router.

**Remote Access Policy Configuration at the VPN Server**
The remote access policy configuration for the VPN server is the same as described in the On-Demand Branch Office section of this paper.

**Remote Access Policy Configuration at the Chicago Router**
To define the authentication and encryption settings for the VPN connections, the default policy named **Allow access if dial-in permission is enabled** is deleted and the following remote access policy is created:

- **Policy name**: VPN Routers

- **Access method**: VPN

- **User or Group Access**: Group with the VPN_Routers group selected

- **Authentication Methods**: **Extensible Authentication Protocol** with the **Smart card or other Certificate** type and **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** selected

- **Policy Encryption Level**: **Strong encryption** and **Strongest encryption** selected

**Remote Access Policy Configuration at the Phoenix Router**
To define the authentication and encryption settings for the VPN connections, the default policy named **Allow access if dial-in permission is enabled** is deleted and the following remote access policy is created:

- **Policy name**: VPN Routers

- **Access method**: VPN

- **User or Group Access**: Group with the VPN_Routers group selected

- **Authentication Methods**: **Extensible Authentication Protocol** with the **Smart card or other Certificate** type and **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** selected

- **Policy Encryption Level**: **Strong encryption** and **Strongest encryption** selected

**IP Address Pool Configuration**
IP address pools must be configured at the VPN server, the Chicago router, and the Phoenix router.

**IP Address Pool Configuration at the VPN Server**
The IP address pool configuration for the VPN server is the same as described in the Common Configuration for the VPN Server section of this paper.

**IP Address Pool Configuration at the Chicago Router**
A static IP address pool with an IP address of 192.168.9.248 and an ending IP address of 192.168.9.253 is configured. This creates a static address pool for up to five VPN clients.

For more information, see the topic titled Creating a Static IP Address Pool in Windows Server 2003 Help and Support.

**IP Address Pool Configuration at the Phoenix Router**
A static IP address pool with a starting IP address of 192.168.14.248 and an ending IP address of 192.168.14.253 is configured. This creates a static address pool for up to five VPN clients.

For more information, see the topic titled Creating a Static IP Address Pool in Windows Server 2003 Help and Support.

The following sections describe a PPTP-based persistent branch office connection for the Chicago office and an L2TP/IPSec-based persistent branch office connection for the Phoenix office.

**PPTP-based Persistent Branch Office**
The Chicago branch office is a PPTP-based branch office that uses a Windows Server 2003 VPN router to create a persistent, site-to-site VPN connection with the VPN server in New York. The connection is never terminated, even when idle.

To deploy a PPTP, two-way initiated, persistent, site-to-site VPN connection to the corporate office based on the settings configured in the Common Configuration for the VPN Server and Persistent Branch Office sections of this paper, the following settings are configured on the VPN server and Chicago router.

**VPN Server Configuration**
The VPN server is configured with a demand-dial interface, static routes, and PPTP packet filters.

**Demand-Dial Interface for Site-to-Site VPN Connection**
To connect the VPN server to the Chicago router by using a site-to-site VPN connection over the Internet, a demand-dial interface is created by using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: VPN_Chicago
- **Connection Type**: Connect using virtual private networking (VPN)
- **VPN Type**: Point to Point Tunneling Protocol (PPTP)
- **Destination Address**: 131.107.0.1
- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.
- **Static Routes for Remote Networks**

  To make all locations on the Chicago network reachable, the following static route is configured:

  **Destination**: 192.168.9.0

  **Network mask**: 255.255.255.0

  **Metric**: 1

- **Dial Out Credentials**

  **User name**: VPN_CorpHQ

  **Domain**: electronic.example.com

  **Password**: o3\Dn6@`-J4

  **Confirm password**: o3\Dn6@`-J4

Once the demand-dial interface is created, the following change is made:

- For the properties of the demand-dial interface, on the **Options** tab, under **Connection type**, **Persistent connection** is selected.

**Chicago Router Configuration**
The Chicago router is configured with a demand-dial interface and static routes.

**Demand-dial interface for site-to-site VPN connection**
To connect the Chicago office router to the VPN server by using a site-to-site VPN connection over the Internet, a demand-dial interface is created by using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: VPN_CorpHQ
- **Connection Type**: Connect using virtual private networking (VPN)
- **VPN Type**: Layer 2 Tunneling Protocol (L2TP)
- **Destination Address**: 207.46.130.1
- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.
- **Static Routes for Remote Networks**

  To make all locations on the corporate intranet reachable, the following static route is configured:

  **Destination**: 172.16.0.0

  **Network mask**: 255.240.0.0

  **Metric**: 1

  To make all locations on Electronic, Inc. branch offices reachable, the following static route is configured:

  **Destination**: 192.168.0.0

  **Network mask**: 255.255.0.0

  **Metric**: 1

- **Dial Out Credentials**

  **User name**: VPN_Chicago

  **Domain**: electronic.example.com

  **Password**: U9!j5dP(%q1

  **Confirm password**: U9!j5dP(%q1

Once the demand-dial interface is created, the following change is made:

- For the properties of the demand-dial interface, on the **Options** tab, under **Connection type**, **Persistent connection** is selected.

**Static route for the Electronic, Inc. VPN server**

To make the Electronic, Inc. VPN server on the Internet reachable, the following static route is configured:

- **Interface**: The WAN adapter attached to the Internet
- **Destination**: 207.46.130.1
- **Network mask**: 255.255.255.255
- **Gateway**: 0.0.0.0
- **Metric**: 1

| **Note** |
| --- |
| Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway. The gateway address of 0.0.0.0 is an example. 0.0.0.0 is the unspecified IP address. |

**L2TP/IPSec-based Persistent Branch Office**

The Phoenix branch office is an L2TP/IPSec-based branch office that uses a Windows Server 2003 router to create a persistent, site-to-site VPN connection with the VPN server in New York. The connection is never terminated, even when idle.

To deploy an L2TP/IPSec, two-way initiated, persistent, site-to-site VPN connection to the corporate office based on the settings configured in the Common Configuration for the VPN Server and Persistent Branch Office sections of this paper, the following settings are configured on the VPN server and Phoenix router.

**VPN Server Configuration**

The VPN server is configured with a demand-dial interface and a static route.

**Demand-Dial Interface for Site-to-Site VPN Connection**

To connect the VPN server to the Phoenix router by using a site-to-site VPN connection over the Internet, a demand-dial interface is created by using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: VPN_Phoenix
- **Connection Type**: Connect using virtual private networking (VPN)
- **VPN Type**: Layer 2 Tunneling Protocol (L2TP)
- **Destination Address**: 159.60.0.1
- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.
- **Static Routes for Remote Networks**

  To make all locations on the Phoenix network reachable, the following static route is configured:

  **Destination**: 192.168.14.0

  **Network mask**: 255.255.255.0

  **Metric**: 1

- **Dial Out Credentials**

  **User name**: VPN_CorpHQ

  **Domain**: electronic.example.com

**Password**: o3\Dn6@`-J4

**Confirm password**: o3\Dn6@`-J4

After the demand-dial interface is created, the following change is made:

- For the properties of the demand-dial interface, on the **Options** tab, under **Connection type**, **Persistent connection** is selected.

**Phoenix Router Configuration**

The Phoenix router was configured by the Electronic, Inc. network administrator while connected to the Electronic, Inc. intranet and then shipped to the Phoenix site. While the Phoenix router was connected to the Electronic, Inc. intranet, a computer certificate was installed through autoenrollment. Additionally, the Phoenix router computer was configured with a demand-dial interface and a static route.

**Demand-Dial Interface for Site-to-Site VPN Connection**

To connect the Phoenix office router to the VPN server by using a site-to-site VPN connection over the Internet, a demand-dial interface is created by using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: VPN_CorpHQ
- **Connection Type**: Connect using virtual private networking (VPN)
- **VPN Type**: Layer 2 Tunneling Protocol (L2TP)
- **Destination Address**: 207.46.130.1
- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.
- **Static Routes for Remote Networks**

  To make all locations on the corporate intranet reachable, the following static route is configured:

  **Destination**: 172.16.0.0

  **Network mask**: 255.240.0.0

  **Metric**: 1

  To make all locations on Electronic, Inc. branch offices reachable, the following static route is configured:

  **Destination**: 192.168.0.0

  **Network mask**: 255.255.0.0

  **Metric**: 1

- **Dial Out Credentials**

  **User name**: VPN_Phoenix

  **Domain**: electronic.example.com

  **Password**: z2F%s)bW$4f

  **Confirm password**: z2F%s)bW$4f

Once the demand-dial interface is created, the following change is made:

- For the properties of the demand-dial interface, on the **Options** tab, under **Connection type**, **Persistent connection** is selected.

**Static Route for the Electronic, Inc. VPN Server**

To make the Electronic, Inc. VPN server on the Internet reachable, the following static route is configured:

**Interface**: The WAN adapter attached to the Internet

**Destination**: 207.46.130.1

**Network mask**: 255.255.255.255

**Gateway**: 0.0.0.0

**Metric**: 1

| **Note** |
| --- |
| Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway. The gateway address of 0.0.0.0 is an example. 0.0.0.0 is the unspecified IP address. |

# Extranet for Business Partners (VPN with Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The network administrator for Electronic, Inc. has created an extranet, a portion of the Electronic, Inc. private network that is available to business partners through secured VPN connections. The Electronic, Inc. extranet is the network attached to the Electronic, Inc. VPN server and contains a file server and a Web server. Parts distributors Tasmanian Traders and Parnell Aerospace are Electronic, Inc. business partners and connect to the Electronic, Inc. extranet by using on-demand, site-to-site VPN connections. An additional remote access policy is used to ensure that the business partners can only access the extranet file server and Web server.

The file server on the Electronic, Inc. extranet is configured with an IP address of 172.31.0.10 and the Web server is configured with an IP address of 172.31.0.11. Tasmanian Traders uses the public network ID of 131.107.254.0 with a subnet mask of 255.255.255.0. Parnell Aerospace uses the public network ID of 131.107.250.0 with a subnet mask of 255.255.255.0. To ensure that the extranet Web server and file server can reach the business partners, static routes are configured on the file server and Web server for each of the business partner networks that use the gateway address of 172.31.0.1

To simplify configuration, the VPN connection is a one-way initiated connection. The business partner's router always initiates the connection. For more information, see the topic titled One-Way Initiated Demand-Dial Connections in Windows Server 2003 Help and Support.

Figure 5 shows the Electronic, Inc. VPN server that provides extranet connections for business partners.




**Figure 5: The Electronic, Inc. VPN server that provides extranet connections for business partners**

To deploy business partner, on-demand, one-way initiated, site-to-site VPN connections to connect Tasmanian Traders and Parnell Aerospace to the Electronic, Inc. extranet based on the settings configured in the Common Configuration for the VPN Server section of this paper, the following additional settings are configured.

**Domain Configuration**

For the VPN connection to Tasmanian Traders, the user account PTR_Tasmanian is created with the following settings:

- Password of Y8#-vR7?]fl.
- For the account properties of the PTR_Tasmanian account, the **User must change password at next logon** option is cleared and **Password never expires** option is selected.
- For the dial-in properties on the PTR_Tasmanian account, the remote access permission is set to **Control access through Remote Access Policy** and the static route 131.107.254.0 with a subnet mask 255.255.255.0 is added.
- The PTR_Tasmanian account is added to the VPN_Partners group.

For the VPN connection to Parnell Aerospace, the user account PTR_Parnell is created with the following settings:

- Password of W@8c^4r-;2\.

- For the account properties of the PTR_Parnell account, the **User must change password at next logon** option is cleared and **Password never expires** option is selected.

- For the dial-in properties on the PTR_Parnell account, the remote access permission is set to **Control access through Remote Access Policy** and the static route 131.107.250.0 with a subnet mask 255.255.255.0 is added.

- The PTR_Parnell account is added to the VPN_Partners group.

**Remote Access Policy Configuration**
To define the authentication and encryption settings for business partner VPN connections, the following remote access policy is created:

- **Policy name**: VPN Partners

- **Access method**: VPN

- **User or Group Access**: Group with the EXAMPLE\VPN_Partners group selected

- **Authentication Methods**: **Extensible Authentication Protocol** with the **Smart card or other Certificate** type and **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** selected

- **Policy Encryption Level**: **Strong encryption** and **Strongest encryption** selected

After the remote access policy is created, it configuration is modified in the following way:

- On the **IP** tab of the profile settings, the following TCP/IP packet filters are configured:

Input Filters:

- Filter action: Deny all traffic except those listed below

- Filter 1: Destination network IP address of 172.31.0.10 and subnet mask of 255.255.255.255

- Filter 2: Destination network IP address of 172.31.0.11 and subnet mask of 255.255.255.255

Output Filters:

- Filter action: Deny all traffic except those listed below

- Filter 1: Source network IP address of 172.31.0.10 and subnet mask of 255.255.255.255

- Filter 2: Source network IP address of 172.31.0.11 and subnet mask of 255.255.255.255

The following sections describe a PPTP-based extranet for the business partner Tasmanian Traders and an L2TP/IPSec-based extranet for the business partner Parnell Aerospace.

**PPTP-based Extranet for Business Partners**
Tasmanian Traders is a business partner that uses a Windows Server 2003 router to create an on-demand, PPTP-based, site-to-site VPN connection with the Electronic, Inc. VPN server in New York as needed. When the connection is created and is idle for five minutes, the connection is terminated. The Tasmanian Traders router is connected to the Internet by using a permanent WAN connection.

To deploy a PPTP, one-way initiated, on-demand, site-to-site VPN connection to the corporate office based on the settings configured in the Common Configuration for the VPN Server and Extranet for Business Partners sections of this paper, the following settings are configured on the Tasmanian Traders router.

**Demand-Dial Interface for Site-to-Site VPN Connection**

To connect the Tasmanian Traders router to the Electronic, Inc. VPN server by using a site-to-site VPN connection over the Internet, a demand-dial interface is created by using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: Electronic
- **Connection Type**: Connect using virtual private networking (VPN)
- **VPN Type**: Point to Point Tunneling Protocol (PPTP)
- **Destination Address**: 207.46.130.1
- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.
- **Static Routes for Remote Networks**

  To make all locations on the Electronic, Inc. extranet reachable, the following static route is configured:

  **Destination**: 172.31.0.0

  **Network mask**: 255.255.0.0

  **Metric**: 1
- **Dial Out Credentials**

  **User name**: PTR_Tasmanian

  **Domain**: electronic.example.com

  **Password**: Y8#-vR7?]fI

  **Confirm password**: Y8#-vR7?]fI

#### L2TP/IPSec-based Extranet for Business Partners

Parnell Aerospace is a business partner that uses a Windows Server 2003 router to create an on-demand, L2TP/IPSec-based, site-to-site VPN connection with the Electronic, Inc. VPN server in New York as needed. When the connection is created and is idle for five minutes, the connection is terminated. The Parnell Aerospace router is connected to the Internet by using a permanent WAN connection.

To deploy an L2TP/IPSec, one-way initiated, on-demand, site-to-site VPN connection to the corporate office based on the settings configured in the Common Configuration for the VPN Server and Extranet for Business Partners sections of this paper, the following settings are configured on the Parnell Aerospace router:

#### Certificate Configuration

The Parnell Aerospace router was configured by the Electronic, Inc. network administrator while physically connected to the Electronic, Inc. intranet and then shipped to the network administrator at Parnell Aerospace. While the Parnell Aerospace router was connected to the Electronic, Inc. intranet, a computer certificate was installed through auto-enrollment.

#### Demand-Dial Interface for Site-to-Site VPN Connection

To connect the Parnell Aerospace router to the Electronic, Inc. VPN server by using a site-to-site VPN connection over the Internet, a demand-dial interface is created by using the **Demand-Dial Interface** wizard with the following settings:

- **Interface Name**: Electronic
- **Connection Type**: Connect using virtual private networking (VPN)
- **VPN Type**: Layer 2 Tunneling Protocol (L2TP)
- **Destination Address**: 207.46.130.1

- **Protocols and Security**: The **Route IP packets on this interface** check box is selected.
- **Static Routes for Remote Networks**

  To make all locations on the Electronic, Inc. extranet reachable, the following static route is configured:

  **Destination**: 172.31.0.0

  **Network mask**: 255.255.0.0

  **Metric**: 1
- **Dial Out Credentials**:

  **User name**: PTR_Parnell

  **Domain**: electronic.example.com

  **Password**: W@8c^4r-;2\

  **Confirm password**: W@8c^4r-;2\

# Dial-up and VPNs with RADIUS Authentication (VPN with Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In addition to VPN-based remote access, the network administrator for Electronic, Inc. wants to provide modem-based dial-up remote access for employees of the New York office. All employees of the New York office belong to an Active Directory group called NY_Employees. A separate remote access server running Windows Server 2003 provides dial-up remote access at the phone number 555-0111. Rather than administer the remote access policies of both the VPN server and the remote access server separately, the network administrator is using a computer running Windows Server 2003 with the Internet Authentication Service (IAS) as a RADIUS server. The IAS server has an IP address of 172.31.0.9 on the Electronic, Inc. extranet and provides centralized remote access authentication, authorization, and accounting for both the remote access server and the VPN server.

Figure 6 shows the Electronic, Inc. RADIUS server that provides authentication and accounting for the VPN server and the remote access server.




**Figure 6: The Electronic, Inc. RADIUS server that provides authentication and accounting for the VPN server and the remote access server**

**Domain Configuration**
For each New York office employee that is allowed dial-up access, the remote access permission for the dial-in properties of the user account is set to **Control access through Remote Access Policy**.

**Remote Access Policy Configuration**
Remote access policies must be modified in two ways:

1. The existing remote access policies that are configured on the VPN server must be copied to the IAS server.

2. A new remote access policy is added for dial-up remote access clients on the IAS server.

**Copying the Remote Access Policies**
Once the VPN server is configured to use RADIUS authentication, the remote access policies stored on the VPN server are no longer used. Instead, the remote access policies stored on the IAS server are used. Therefore, the current set of remote access policies is copied to the IAS server.

For more information, see the topic titled Copying the IAS Configuration to Another Server in Windows Server 2003 Help and Support.

**Creating a New Remote Access Policy for Dial-up Remote Access Clients**
To define the authentication and encryption settings for dial-up connections by employees of the New York office, the following remote access policy is created on the RADIUS server computer:

- **Policy name**: Dial-Up for New York Employees

- **Access method**: Dial-up

- **User or Group Access**: Group with the EXAMPLE\NY_Employees group selected

- **Authentication Methods**: **Extensible Authentication Protocol** with the **Smart card or other Certificate** type, **Microsoft Encrypted Authentication (MS-CHAP)**, and **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** are selected

- **Policy Encryption Level**: All options selected

**RADIUS Configuration**

To configure RADIUS authentication and accounting, the network administrator for Electronic, Inc. configures the following:

- The RADIUS server is a computer running Windows Server 2003 with the Internet Authentication Service networking component installed. The Internet Authentication Service is configured for two RADIUS clients: the remote access server and the VPN server. For more information, see the topic titled Registering RADIUS Clients in Windows Server 2003 Help and Support.

- The remote access server is configured to use RADIUS authentication and accounting at the IP address of 172.31.0.9 and a shared secret. For more information, see the topics titled Configuring RADIUS authentication and Configuring RADIUS accounting in Windows Server 2003 Help and Support.

- The VPN server is configured to use RADIUS authentication and accounting at the IP address of 172.31.0.9 and a shared secret.

**Dial-up Remote Access Client Configuration**

On the Windows XP remote access client computers, the New Connection Wizard is used to create a dial-up connection with the following settings:

**Network Connection Type**: Connect to the network at my workplace

**Network Connection**: Dial-up connection

**Connection Name**: Electronic, Inc.

**Phone Number**: 555-0111

**Connection Availability**: Anyone's use

# Summary (VPN with Windows Server 2003)

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Electronic, Inc. used VPN technologies included with Windows Server 2003 and Windows XP to leverage the connectivity of the Internet to connect remote users, branch offices, and business partners. Electronic, Inc.'s Windows Server 2003 VPN and dial-up remote access servers, used in conjunction with the Internet Authentication Service, provide centralized authentication, authorization, accounting, and administration of remote access policies for a VPN and dial-up remote access solution.

# Related Links for Virtual Private Networking with Windows Server 2003: An Example Deployment

Updated: January 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

See the following resources for further information:

- Windows VPN Web site (http://go.microsoft.com/fwlink/?LinkID=69993)
- Internet Authentication Service Web site (http://go.microsoft.com/fwlink/?LinkID=67728)

For the latest information about Windows Server 2003, see the Windows Server 2003 Web site at (http://go.microsoft.com/fwlink/?LinkID=67729).

# Using Hotpatching Technology to Reduce Servicing Reboots

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This document describes Hotpatching technology and its use in deploying software updates for the Microsoft® Windows Server® 2003 operating systems. Hotpatching enables you to install software update packages without a system reboot or interruption in services or processes.

In This White Paper

Introduction to Hotpatching

Hotpatch Package Structure

Deployment Strategy

Limitations and Compatibility Issues

Package Installation

# Introduction to Hotpatching

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This white paper provides a general overview of Hotpatching technology and describes the benefits of using it.

The Microsoft® Windows® operating systems and Microsoft business solutions constitute an environment of choice for many large enterprise customers. For such critical systems, these customers demand high availability and assume the industry standard of "five nines" (99.999%) reliability.

It is important to avoid service interruptions when installing updates on the operating system of such critical systems. Hotpatching provides a mechanism to:

- Update system files without rebooting.
- Update system files without stopping services and processes.

**Microsoft reboot reduction initiative**

Hotpatching is part of the Microsoft reboot reduction initiative, which seeks to help minimize the need for a full system reboot after installing updates. Reducing reboots is important because IT departments in many organizations implement a time-consuming test cycle every time an update is installed and the system is rebooted. This results in loss of productivity and revenue to the organization until their system is fully verified and operational. Hotpatching allows customers to deploy important updates and patches in a timely, transparent manner without requiring a full system shutdown and restart. This reduces their rollout time.

The following examples demonstrate possible savings from reboot reduction:

- Of the 22 updates that shipped for Windows Server 2003 RTM between April 2005 and August 2005, 15 of them required a reboot. Eight of these could have been hotpatched. This would have reduced the number of reboots by 53%.
- Of the 14 updates that shipped for Windows Server 2003 Service Pack 1 (SP1) prior to August 2005, ten of them required a reboot. Four of these could have been hotpatched. This would have reduced the number of reboots by 40%.

# Hotpatch Package Structure

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A hotpatch package contains hotpatch and coldpatch binary files for the operating system update.

- The hotpatch binary file contains only the function necessary to address the critical operating system flaw.
- The coldpatch contains the old binary file with the fixed function appended to it and an instruction to jump from the flawed function to the fixed function. Redirecting to the new function ensures that the defective processes in memory are fixed by patching the old function.

Hotpatch and coldpatch binary files are produced by the hotpatch creation tool, as shown in Figure 1.




**Figure 1 Hotpatch and coldpatch binaries**

Figure 1 shows that the hotpatch binary file contains one changed function (Function A') and that the coldpatch is composed of the old binary file with the appended new function. When a hotpatch package is installed, it:

- Replaces the old binary file on disk with the coldpatch binary file.
- Injects the updated function (as a hotpatch binary file) into the loaded image of the defective binary file.
- Inserts a jump instruction above the defective function to redirect all subsequent calls to the updated function.

| **Note** |
| --- |
| Hotpatching must be turned ON. |

Hotpatch application addresses currently running instances of the critical flaws, and the complementing coldpatch secures the new instances of the process and persists the patch beyond reboot. Thus, any process that starts after applying the update will be fixed because the operating system will load the coldpatch binary file. This coldpatch format allows similar representations of the new function on disk and in memory.

The package containing a hotpatch-enabled fix will have two binary files related to the file being serviced. The hotpatch file contains the ".hp" appendix.

# Deployment Strategy

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Hotpatching technology was designed to reduce system downtime and to provide a method for deploying critical updates without a system reboot. Because hotpatches cannot be applied to all update packages, it is common for update cycles to contain updates with and without hotpatches. Figure 2 outlines the recommended hotpatching deployment strategy.




**Figure 2 Update deployment strategy**

# Limitations and Compatibility Issues

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are limitations to hotpatching technology, and because of these limitations, not all updates can be hotpatched. Hotpatching does not support components or applications that contain:

- Additions or changes to exports
- Data structure changes
- Dependency changes
- Global state changes
- Infinite loops and persistent stack frames
- Inline assembly code

**Supported platforms**
Hotpatching is supported on 32-bit versions of Windows Server 2003 operating system with SP1. Support will also be available in future Microsoft operating system releases and platforms.

**Compatibility issues**
Hotpatch installation has been found to cause issues with some non-Microsoft applications. For a current list of application compatibility issues, see article 922612 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=78020).

# Package Installation

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can download hotpatch packages from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkID=19831). The Software Update Installation Wizard will step you through the hotpatch installation process.

Hotpatch installation for non-Microsoft deployment tools (such as IBM Tivoli and Altiris Patch Management) is supported by enabling hotpatching from the command line. To enable hotpatching from the command line, type:

**WindowsServer2003-KB######-x86-LLL.exe /hotpatch:enable**

If the hotpatch installation is successful, a reboot will not be required. A log file is created as a part of the installation process, and it will show if the hotpatch installation was successful.

| **Note** |
| --- |
| For the hotpatch installation to be successful, the existing binary file must be the version that the hotpatch was built for. This criterion is verified before the hotpatch installation, and if the binary file does not match, the coldpatch binary file of the update will install, which will require a reboot. |

For more information about hotpatch installation and installation logs, see article 897341 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=78021).

# Best Practice Guide for Securing Active Directory Installations

Updated: January 23, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Organizations require a network operating system (NOS) that provides secure network access to network data by authorized users and that rejects access by unauthorized users. For a Microsoft® Windows Server® 2003 NOS, the Active Directory® directory service provides many key components for authenticating users and for generating authorization data that controls access to network resources.

**Note**

For a printable .doc version of this guide, see Best Practice Guide for Securing Active Directory Installations.doc at the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=140862).

This guide does not include operations information. For more information about day-to-day service operations for Windows Server 2003 deployments, see the Windows Server 2003 Operations Guide, which is available as a downloadable document on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=63079).

A breach in Active Directory security can result in the loss of access to network resources by legitimate clients or in the inappropriate disclosure of potentially sensitive information. Such information disclosure affects data that is stored on network resources or in Active Directory. To avoid these situations, organizations need more extensive information and support to ensure enhanced security for their NOS environments. This guide addresses this need for organizations that have new, as well as existing, Active Directory deployments.

This guide contains recommendations for protecting domain controllers against known threats, establishing administrative policies and practices to maintain network security, and protecting Domain Name System (DNS) servers from unauthorized updates. It also provides guidelines for maintaining Active Directory security boundaries and securing Active Directory administration.

This guide also includes procedures for enacting these recommendations. For more information, see "Appendix: Procedures" later in this guide.

**Note**

The recommendations and procedures in this guide have been tested in a lab to demonstrate that domain controllers that are built, configured, and administered in accordance with these recommendations can be deployed and operated in an efficient manner that enhances security.

**In This Guide**


- Scope of This Guide (Best Practices for Securing Active Directory Installations)
- Chapter 1: Planning In-Depth Active Directory Security
- Chapter 2: Establishing Secure Active Directory Boundaries
- Chapter 3: Deploying Secure Domain Controllers
- Chapter 4: Strengthening Domain and Domain Controller Policy Settings
- Chapter 5: Establishing Secure Administrative Practices
- Chapter 6: Securing DNS
- Appendix: Procedures
- Appendix: Setting User Account Control Policy for Delegated Administrators

# Scope of This Guide (Best Practices for Securing Active Directory Installations)

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Although NOS security relies on secure design principles and operating practices for all components in the operating system, the scope of this guide is limited to recommendations for deploying and securing Active Directory domain controllers. Other security topics, such as planning for secure network connectivity and secure clients, are not addressed in this guide. For more information about planning for secure network connectivity and secure clients, see "Planning a Secure Environment," "Designing a Public Key Infrastructure," and "Planning a Smart Card Deployment" in Designing and Deploying Directory and Security Services of the Microsoft® Windows® Server 2003 Deployment Kit (or see "Planning a Secure Environment" on the Web at http://go.microsoft.com/fwlink/?LinkId=20374, "Designing a Public Key Infrastructure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4735, and "Planning a Smart Card Deployment" on the Web at http://go.microsoft.com/fwlink/?LinkId=4736).

The process flow that is described in this guide is designed to create a secure domain controller environment by providing guidelines for Active Directory deployment and administrative policies and practices. These guidelines can be applied to both new and existing Active Directory infrastructures. In addition, because Domain Name System (DNS) is an integral component of Active Directory, this guide also includes guidelines for administrative policies and practices for DNS on a domain controller. For more information, see "Securing DNS" later in this guide.

Figure 1 depicts the process flow for the recommendations in this guide.



# General Guidelines (Best Practices for Securing Active Directory Installations)

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The goal of this guide is to assist organizations in enhancing the security of their Active Directory systems. Because any guide that addresses a general audience can provide only guidelines, in some instances these guidelines might conflict with the needs of an organization to lower costs, provide services and line-of-business applications, or maintain an information technology (IT) infrastructure. In such cases, an organization's security planning team can arrive at suitable tradeoffs by evaluating these other needs against the need for security.

In addition, the recommendations in this guide take into consideration how an organization's domain controllers are deployed. Domain controllers can be deployed in datacenters for enterprise intranets, in branch offices, and in datacenters for extranets. In some cases, the guidelines vary in accordance with special circumstances that are encountered in each environment.

# Audience (Best Practices for Securing Active Directory Installations)

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This guide is intended primarily for IT planners, architects, and managers who are responsible for establishing Active Directory deployment and operations practices. As a result, this guide emphasizes the decision-making process, rather than procedures.

# How to Use This Guide (Best Practices for Securing Active Directory Installations)

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The information in this guide is presented as if the reader's organization is in the planning stage of its Active Directory deployment. However, this information can be equally beneficial to an organization that is reviewing its current Active Directory security practices.

The information in this guide focuses on security, and this guide should not be used as the sole reference for planning an Active Directory deployment. For a complete analysis of deployment and upgrade issues and for recommendations for planning, designing, and deploying an Active Directory infrastructure, see "Designing the Active Directory Logical Structure" in Designing and Deploying Directory and Security Services of the Windows Server 2003 Deployment Kit (or see "Designing the Active Directory Logical Structure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4723).

You can proceed through your Active Directory security planning process in the order presented in this guide. Each phase of the Active Directory security planning process, such as securing domain controllers, is contained in its own chapter. Each chapter begins with a discussion of how the security recommendations enhance security, and it also discusses the cost of the recommendations in terms of complexity and performance. If a recommendation is impractical for a specific deployment strategy, that fact is indicated and specific alternatives are recommended, if they exist. Finally, the recommendations in each chapter are summarized in a checklist at the end of the chapter.

You can proceed to the next chapter after completing the checklist of recommendations at the end of the previous chapter. Before proceeding with security planning, make sure that you:

- Set your business goals and practices and understand how they affect Active Directory security.
- Gain executive-level sponsorship to implement a secure, Active Directory–managed Windows network.

Active Directory infrastructure and practices can span both technology and business areas. Therefore, to make progress in security planning, you must articulate the value of security to IT and business decision makers.

# Chapter 1: Planning In-Depth Active Directory Security

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

A layered, in-depth approach to security includes protecting the physical environment, securing administrative policies and practices, and configuring domain controllers with settings that protect them from tampering. This security-in-depth approach provides the general outline for planning Active Directory security.

Active Directory security planning begins with an analysis of the environment in which Active Directory is being deployed. The scope of the environment (intranet, extranet, and branch office) determines the methods that are required to secure the environment.

To best plan for secure operations, assess the security threats against which your Active Directory deployment must be protected, and then implement a deterrent to each threat to address potential problems before they can occur. In addition, put contingency plans in place to ensure appropriate responses to unforeseen disasters.

# Overview of the Role of Active Directory in Secure Access

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

To provide a secure network computing environment, Windows requires verification of every user's identity before it allows access to network resources. The process of verifying the identity of users, also known as authentication, is part of the logon process. The process of granting or denying access to a resource, also known as authorization, occurs when network users attempt to connect to servers and other network devices. Active Directory is the store for authentication and authorization data, and it is also the distributed service that manages the use of this data.

**Authentication: Identifying Network Users**

During the network logon process, all entities that can connect to network resources must identify themselves to the security system on a domain controller, which authenticates their identity by verifying their credentials against identification data that is stored in Active Directory. Active Directory stores security-related identity information for network users in the form of an account object. The system generates a unique security identifier (SID) for every account object that can be authenticated or authorized for access to resources. A SID is a value that uniquely identifies a user, group, or computer account within an Active Directory forest. Access control mechanisms in Windows networks identify security principals by SID rather than by name. Every SID contains the identifier of the domain in which the account object resides.

Objects that have a SID are known as security principals. The following types of objects act as security principals:

- User: a person or a service that requires user credentials
- Computer: a workstation or server that is running a version of Windows other than Microsoft® Windows® 95, Windows 98, or Windows Millennium Edition
- Group: a set of users, computers, or other groups

**Authorization: Allowing Access to Network Resources**

Authorization, also known as access control, is the process of determining the level of access that is allowed to an Active Directory object or file system object. After Active Directory confirms the identity of the user, the security system on the authenticating domain controller generates partial authorization data in the form of the user's primary SID, plus SIDs for groups to which the user belongs that are recognized by all the resources on the Windows network. The remainder of the authorization data is generated at the time that the user requests access to a specific network resource, such as a server, file share, or directory object. The authorization data is used by the computer that houses the network resource to generate an access token. The access token is used to determine the level of access that the user has to the network resource.

The access token contains the following:

- The list of SIDs that represent the user
- All groups (including nested groups) of which the user is a member
- The user's privileges (also called user rights) on the local computer

All objects or resources that are secured have a discretionary access control list (DACL) assigned to them that specifies the access rights of users and groups on that resource. Access to one of these objects or resources is controlled by an access check, in which the security system determines whether the requested access should be granted or denied by evaluating the contents of the access token of the requester against the DACL on the resource.

# Planning for Active Directory Security-in-Depth

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

A security strategy for an enterprise is most effective when data is protected by more than one layer of security. With this type of strategy, the potential losses that might be caused by a security failure in any single layer are minimized by the remaining security layers. For example, when a home is protected by both door locks and a security system, the homeowner is implementing a security-in-depth strategy that is more effective than either of these security features alone.

A security-in-depth strategy first divides all security elements into discrete security layers. In this way, the security effectiveness of each layer can be determined independently, and a security plan can then be implemented. Active Directory security policies and practices can be divided into the following layers:

- Physical security for domain controllers and the network (physical access to domain controllers, backup data, and network components)
- Administrative authority (security management and secure administrative practices)
- End system (domain controller settings, policy settings, and deployment practices)

The following figure shows one way of visualizing the relationships among these security layers for Active Directory.




This guide provides recommendations for security best practices that are based on security-in-depth. Therefore, it is organized along the lines of the security layers for secure deployment of domain controllers, including physical security, secure domain controller policies, and secure administrative practices.

# Deployment Scenarios for Domain Controllers in a Secure Network Operating System

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Evaluating the need for Active Directory security against the other business needs of an organization depends on the type of operating environment in which the domain controllers are deployed. The three most common operating environments in which domain controllers are deployed are the following:

- Intranet datacenter
- Branch office
- Extranet datacenter

The operating environments in which domain controllers are deployed in an organization can, in some cases, affect the feasibility of implementing the Active Directory security guidelines that are presented in this guide. Each set of guidelines includes a discussion, where appropriate, of how the deployment model affects the cost and logistical difficulty of implementing the recommendations.

**Domain Controllers in Intranet Datacenters**
The datacenter for an enterprise intranet is the most common domain controller deployment scenario, and it is the easiest scenario in which to deploy and administer domain controllers. Intranet datacenters normally possess most, if not all, of the following characteristics:

- Fast, high-bandwidth network connectivity among all domain controllers
- Centralized, secure facilities for housing and managing domain controllers
- Facilities for building and configuring domain controllers
- Facilities for testing hardware, operating systems, and policies
- Centralized, dedicated IT staff and security staff with defined administrative roles
- Facilities for monitoring and auditing domain controllers

Larger organizations might contain more than one datacenter to support regional business operations. Each datacenter is likely to have its own IT staff and security staff.

The figure below represents one possible datacenter design for a medium-sized organization with facilities in two regions. Each datacenter services intranet clients and remote clients in its region. Connectivity requirements among datacenters vary in accordance with an organization's business needs and with its Active Directory infrastructure design.




Of the three domain controller deployment scenarios, the intranet datacenter most readily supports the requirements of a managed environment. IT staff members are generally on-site; therefore, access to the datacenter and to administrative tasks can easily be restricted to these IT staff members. In addition, monitoring and troubleshooting of domain controller problems is simplified, and attacks can be detected and handled quickly.

Intranet datacenters also possess a high degree of manageability. Intranet datacenter deployments generally include the following management characteristics:

- Centrally designed and managed domain controller and administrative policies
- Written administrative practices for building, deploying, and managing domains and domain controllers

**Domain Controllers in Branch Offices**

An organization with a large number of locations, each with a limited number of users and servers, might still require local network access to a domain controller for authentication and authorization at each location. Branch office deployments normally possess most, if not all, of the following characteristics:

- Slow, intermittent connectivity to other domain controllers outside the branch office
- Domain controllers sharing a subnet with users and computers
- Domain controllers hosting other applications or infrastructure services (such as Windows Internet Name Service (WINS) and Print Services)
- Lack of dedicated, secure facilities for housing and managing domain controllers
- Lack of a local, dedicated IT staff

The figure below shows an intranet datacenter that has multiple, small, branch office locations, each with at least one domain controller. Branch office connections might include virtual private network (VPN) tunnels, dial-up connections, or leased lines. Local users within the intranet are connected to the intranet datacenter by local area network (LAN) segments.




Of the three domain controller deployment scenarios, the branch office presents the greatest challenges in supporting the requirements of a secure, managed environment. IT staff members are generally off-site; therefore, all administrative tasks cannot easily be restricted to those staff members. In addition, restricting physical access to the domain controllers can be difficult. Detecting and responding to domain controller problems and attacks can be slow and costly.

Another challenge that the branch office model presents is manageability. Branch office deployments generally include the following management characteristics:

- Decentralized domain controller management or remote management of domain controllers by administrators in the datacenter
- Lack of uniformity in the application of domain controller management and administrative practices

**Domain Controllers in Extranet Datacenters**

The scenario for deployment of domain controllers in an extranet datacenter is a variation of the intranet datacenter scenario. Like intranet datacenters, extranet datacenters normally possess most of the following characteristics:

- Fast, high-bandwidth network connectivity among all domain controllers
- Centralized, secure facilities for housing and managing domain controllers
- Facilities for building and configuring domain controllers
- Facilities for testing hardware, operating systems, and policies
- Centralized, dedicated IT staff and security staff with defined administrative roles
- Facilities for monitoring and auditing domain controllers

Domain controllers in extranet datacenters fulfill two primary functions. First, they can be used to manage servers and resources within the extranet. Second, they can be placed in an "outward-facing" position so that they can be used to authenticate customers and partners and to provide customers and partners with access to resources in the extranet. Additional characteristics that most extranet datacenters possess include the following:

- Firewalls that restrict Internet traffic into the datacenter
- Firewalls that prevent users in the datacenter from accessing the corporate network
- Customers and partners tunneling into the extranet from outside the network

The figure below shows an intranet datacenter with an extranet datacenter. The extranet datacenter is protected from the Internet by a firewall, and the intranet datacenter is protected from the extranet datacenter by a firewall. Internet users connect to the extranet Web servers through the Internet, and intranet users connect to the intranet servers both remotely through the Internet and locally from within the intranet.




In addition to the management characteristics of intranet datacenters, extranet datacenters might have administrators managing accounts within the extranet itself, or they might have accounts in the enterprise intranet that are used to manage the extranet.

# Security Planning Through Threat Analysis

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This guide concentrates primarily on recommendations for minimizing known security threats to Active Directory. However, the following sections provide a short summary of how threat analysis can be used to create an overall security plan, based on the security features in Windows Server 2003.

Making Active Directory deployment and administration secure encompasses planning for both mitigation (proactive security) and contingency (reactive security). The proactive portion of the plan protects assets by alleviating any threats to the system that might be caused by user mistakes and by attacks that are based on known threats. The reactive portion of the plan provides contingency plans to implement under the following conditions:

- Threat analysis fails to anticipate a threat.
- It is not possible to completely mitigate a threat.
- Security recommendations cannot be implemented.

The following figure illustrates how threat analysis fits into an overall security plan.

**Identifying Types of Threats**

In their attempts to breach a system's security, all threats exploit weaknesses in the operating system, applications, network design, security policies, or administrative practices. In addition, social engineering phenomena pose a threat to Active Directory.

Threats are commonly categorized according to the goal of the attack. This type of threat analysis is referred to by the acronym STRIDE, which is derived from the first letter of each category of threat, plus the added category of social engineering, as described in the following sections.

**Spoofing**

The goal of a spoofing attack is illicit access to network resources by unauthorized users. Spoofing involves forging the identity of a valid system user or resource to gain access to the system, thereby compromising system security. Spoofing attacks include:

- Changing the identity that is associated with an Active Directory object.
- Subverting a secure logon mechanism.
- Using false credentials.

**Tampering with Data**

The goal of a data-tampering attack is to cause unauthorized modification of data, resulting in a loss of data integrity. This type of attack modifies system or user data, with or without detection, resulting in an unauthorized change to network information, network packets in a communication, sensitive files, or the formatting of a hard disk. Data-tampering attacks include:

- Modifying data that should not be accessible.
- Causing a trusted entity to modify data improperly.
- Creating an elevation-of-privilege attack that allows a user to tamper with data.

**Repudiation**

The goal of a repudiation attack is to perform an authorized or unauthorized action and to eliminate any evidence that could prove the identity of the attacker. Repudiation attacks are associated with users who can deny wrongdoing, without any way to prove otherwise. Repudiation attacks include:

- Circumventing the logging of security events.
- Tampering with the security log to conceal the identity of an attacker.

**Information Disclosure**

An information-disclosure risk exists if a user can gain access to data, intentionally or otherwise, that the user is not authorized to see, resulting in the loss of data privacy or confidentiality or both. Information disclosure attacks include:

- Gaining access to data that is considered private and protected.
- Sniffing data on a network while in transit.
- Using social engineering to improperly reveal user identity or passwords.

**Denial of Service**

The goal of a denial-of-service attack is the loss of access by legitimate users to a server or to services. Generally speaking, denial-of-service attacks occur when a malicious user either disables critical services on a computer or consumes so many resources on a system that no resources are available for legitimate users. The resources that can be exhausted might include CPU cycles, disk space, memory, server connections, or network bandwidth, among others. Denial-of-service attacks include:

- Consuming CPU cycles by infinite or very long programmatic looping.
- Consuming excessive memory or file quotas to block legitimate use.
- Causing a crash, restart, or error mechanism to interfere with normal use.

**Elevation of Privilege**

The goal of an elevation-of-privilege attack is illicit access to network resources or services by unauthorized users. The most severe form of an elevation-of-privilege attack is a situation in which an attacker effectively penetrates all system defenses. The attacker then becomes part of the trusted system itself and can completely compromise or destroy the system. Elevation-of-privilege attacks include:

- Improperly gaining unrestricted rights.
- Running untrusted data as native code in a trusted process.
- Spoofing a more privileged identity to gain elevated privileges.

**Social Engineering**

Social engineering is any type of behavior that can inadvertently or deliberately aid an attacker in gaining access to a user's password. For example, someone in an organization might:

- Write their password and place it in a location where a coworker could find it.

- Coax a fellow worker into revealing their password.
- Befriend a janitor or other worker who has physical access to domain controllers.

**Identifying Sources of Threats**

Identification of the various sources of threats to Active Directory provides a basis for understanding and creating an effective security plan. Active Directory defines groups of users, and, by default, it installs policies that limit user access to network resources and services. Therefore, each user group represents a different source of potential threat, as indicated in Table 1.

**Table 1 Sources of Potential Threat to Active Directory**

| Source | Destination |
|---|---|
| Anonymous users | Represents unauthenticated access to the network that is enabled when Group Policy settings allow anonymous access and when permissions on resources are set to allow access to Anonymous Logon through the Pre–Windows 2000 Compatible Access group.<br><br>Allowing access for anonymous users results in a reduced level of security for Active Directory because unauthorized access to Active Directory information can result in information disclosure. See Establishing Secure Domain Controller Build Practices later in this guide. |
| Authenticated users | Represents any user who has successfully completed the authentication process. Successful authentication implies that the user has an identity in the domain, or in a trusted domain, and has provided valid credentials.<br><br>By default, authenticated users have access to information in the directory and on domain controllers, and they can view system event logs on domain controllers. To enhance Active Directory security against information disclosure, unnecessary access should be eliminated. For more information, see Strengthening Domain Controller Policy Settings. |
| Service administrator | Represents administrative accounts that are used legitimately to control directory service configuration and policies and to have physical access to domain controllers to manage server administration. Also used to control the forest infrastructure by creating or removing domains and domain controllers, managing domain and domain controller configuration, and monitoring domain controller health.<br><br>Service administrators are in a position to launch attacks throughout the forest; therefore, they must be highly trusted. For more information, see Specifying Security and Administrative Boundaries. |
| Data administrator | Similar to Microsoft® Windows® NT administrators in their functions and privileges, this group manages data in Active Directory that does not control the directory service or its configuration. This group supports users and computers in the forest by adding and removing organizational units (OUs), computers, users, and groups and by modifying Group Policy settings.<br><br>Data administrators have delegated rights to manage objects in OUs but not to manage domain controllers or forest configuration. For more information, see Specifying Security and Administrative Boundaries. |
| Users with physical access to domain controllers | Applies to any situation in which an individual has access to the area where domain controllers and administrative workstations reside, or applies to a situation in which an individual steals one of these computers. If an unauthorized individual gains access to computers that contain sensitive data, information disclosure or data tampering is possible. For more information, see Maintaining Physical Security. |

# Chapter 2: Establishing Secure Active Directory Boundaries

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Active Directory provides an infrastructure that can support an organization's need for isolation and autonomy, while enabling collaboration among people and organizations. IT planners must determine precisely their need for isolation, autonomy, and collaboration; understand the security implications of delegating administration; and be aware of the tradeoffs in a directory infrastructure.

The guidelines for making directory infrastructure decisions are overviewed in this chapter as a context for assessing security risks to Active Directory boundaries. However, complete Active Directory deployment information is not within the scope of this guide. For detailed planning, designing, and deployment information for Active Directory in Windows Server 2003, see "Designing the Active Directory Logical Structure" in Designing and Deploying Directory and Security Services of the Windows Server 2003 Deployment Kit (or see "Designing the Active Directory Logical Structure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4723).

# Specifying Security and Administrative Boundaries

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The Active Directory forest represents the outermost boundary within which users, computers, groups, and other objects exist; that is, the forest is the security boundary for Active Directory. Active Directory domains, unlike Windows NT domains, are always part of a forest, and they are not themselves the ultimate security boundary. For Windows 2000 and later networks, though, domains are the boundaries for administration and for certain security policies, such as password complexity and password reuse rules, which cannot be inherited from one domain to another. Each Active Directory domain is authoritative for the identity and credentials of the users, computers, and groups that reside in that domain. However, service administrators have abilities that cross domain boundaries. For this reason, the forest is the ultimate security boundary, not the domain.

| **Important** |
| --- |
| Previously published Active Directory documentation states that a domain is a security boundary, but this documentation does not provide specific details about the level of autonomy and isolation that is possible among domains in a forest. Although a domain is, in fact, a security boundary with regard to the management of security policies for Active Directory, it does not provide complete isolation in the face of possible attacks by service administrators. |

**Delegating Administration**

Organizations typically delegate the administration of all or part of the Active Directory service or the data that is stored in their domains. Table 2 lists the reasons for delegating administration.

**Table 2 Reasons to Delegate Administration**

| Reasons | Explanation |
| --- | --- |
| Organizational | One part of an organization might want to share an infrastructure to reduce costs, but it requires operational independence from the rest of the organization. |
| Operational | One part of an organization or a single application might place unique constraints on the directory service configuration, availability, or security. Examples include:<br><br>Military organizations<br><br>Hosting scenarios<br><br>Extranet-based directory services |
| Legal | Some organizations have legal requirements that compel them to operate in a specific manner, such as restricting information access. Examples include:<br><br>Financial institutions<br><br>Defense contractors<br><br>Government organizations |

For these reasons, an organization might need to delegate control over service management, data management, or both. The goal of delegation is to achieve either autonomy or isolation:

- **Autonomy** is the ability of administrators to manage independently part or all of the Active Directory service or the data in the directory or on member computers.

- **Isolation** is the ability of administrators to prevent other administrators from interfering in or accessing part or all of the Active Directory service or the data in the directory or on member computers.

An organization's requirements for service autonomy, data autonomy, service isolation, and data isolation determine the Active Directory infrastructure that is best suited to its needs. The first step is to define precisely the needs of your organization.

Active Directory boundaries can assist an organization in achieving the level of autonomy and isolation that its business units

require. Table 3 provides a comparison of the characteristics of administrative autonomy and isolation.

**Table 3 Comparison of the Characteristics of Autonomy and Isolation in Administration**

| Administrative Role | Autonomy means to... | Isolation means to... |
|---|---|---|
| Service administrator | Manage independently all or part of service administration (service autonomy). | Prevent other service administrators from controlling or interfering with service administration (service isolation). |
| Data administrator | Manage independently all or part of the data that is stored in Active Directory or on member computers (data autonomy). | Prevent other data administrators from controlling or viewing data in Active Directory or on member computers (data isolation). |

Autonomy is less constrained than isolation. Administrators who require only autonomy accept the fact that other administrators of equal or higher privilege in the system have equivalent or overriding control in the forest. Because autonomy is less constrained than isolation, it is generally less costly and more efficient to delegate administrative autonomy.

Autonomy can be achieved by delegating service or data administration. Isolation requires that a business unit deploy a separate forest to contain its administrators, users, and resources. Multiple forests in an organization require external trusts to support collaboration. These trusts are discussed further in Establishing Secure Collaboration with Other Forests later in this guide.

For a complete discussion of the effects of autonomy and isolation and the service and data administrator roles, see "Designing the Active Directory Logical Structure" in Designing and Deploying Directory and Security Services of the Windows Server 2003 Deployment Kit (or see "Designing the Active Directory Logical Structure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4723).

**Trusting Service Administrators**

Before choosing an Active Directory delegation model, consider the following facts about Active Directory administrative roles:

- Forest owners maintain the right to control domain-level services and to access data that is stored in any domain in the forest.
- Domain owners maintain the right to access data that is stored in the domain or on its member computers.
- Domain owners, although operating autonomously from forest owners and other domain owners, cannot prevent a malicious domain owner from controlling their services and accessing their data.

The need for a high degree of collaboration and trust among the authorities that are responsible for the management of domains in a forest requires that all administrators who manage domains be highly trusted. Therefore, Active Directory domains in a forest should never be deployed with the objective of isolating business units that do not trust each other.

To summarize the facts concerning directory structures and directory structure owners, if an organization joins a forest or domain infrastructure, the organization administrators must agree to trust all service administrators in the forest and in all domains. Trusting service administrators means to:

- Believe that service administrators look out for the organization's best interest. Organizations should not elect to join a forest or domain if the organization fears that the owner of the forest or domain might behave maliciously.
- Believe that service administrators follow best practices for service administrators and for restriction of physical access to domain controllers.
- Understand and accept the risks to the organization that are associated with rogue administrators and coerced administrators.

For more information about service and data administration, see Chapter 5: Establishing Secure Administrative Practices.

# Selecting an Active Directory Structure Based on Delegation Requirements

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

In Active Directory, administrators can delegate both service administration and data administration to achieve either autonomy or isolation between organizations. The combination of service management, data management, autonomy, and isolation requirements of an organization affect the Active Directory containers that are used to delegate administration.

The following steps can assist you in determining if your organization's specific delegation requirements justify delegating control of a separate forest, domain, or OU:

1. Begin by placing all organizations in a single-domain forest.

2. For each business unit with unique administrative requirements, determine the appropriate level of autonomy and isolation, based on the respective characteristics in Table 3 earlier in this guide (see Specifying Security and Administrative Boundaries).

3. When recording the justification for each decision, note whether:

    - Delegation is driven by an organizational, operational, legal, or other requirement (see Specifying Security and Administrative Boundaries, Table 2).

    - The requirement pertains to delegation of service management, data management, or both.

    - The requirement indicates the need for autonomy, isolation, or both.

4. Identify the appropriate Active Directory structure:

    - Single-domain forest with OUs for data autonomy

    - Single forest with multiple domains for domain-level service autonomy

    - Separate forests for service isolation

    - Separate forests for forest-level service autonomy

    - Separate forests for data isolation from service owners

For detailed explanations of these and other scenarios and for processes that help you determine the delegation requirements for your organization and the corresponding Active Directory structure, see "Designing the Active Directory Logical Structure" in Designing and Deploying Directory and Security Services of the Windows Server 2003 Deployment Kit (or see "Designing the Active Directory Logical Structure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4723).

**Implications for Active Directory in Extranet Deployment**

Outward-facing domain controllers in an extranet are in a portion of the network where customers and partners can access network resources through the Internet. As a result of the Internet exposure, any information about the intranet forest that is placed in the extranet is subject to the risks of information disclosure or data tampering by external users. Further, a breach of service administration security in the extranet can be used to breach the intranet, if there is no service isolation boundary.

To mitigate the risks previously discussed, Active Directory implementations in an extranet should maintain complete service isolation from the rest of the organization. Therefore, a separate Active Directory forest should be implemented for the extranet. Any service administrator with responsibilities that span the intranet forest and the extranet forest should have a

separate administrative account in each forest.

For more information about designing and managing a secure extranet, see "Secure Extranet for Business Partners" on the Microsoft TechNet Web site at http://go.microsoft.com/fwlink/?LinkId=18590.

# Establishing Secure Collaboration with Other Forests

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Support for interforest collaboration might become necessary for a variety of reasons, including the following:

- Organizational mergers and acquisitions
- Collaboration between isolated forests within an organization
- Shared applications that are accessed from an application forest

Collaboration between domains in different forests can range from simply sharing e-mail to providing access to data and resources. A number of scenarios, including scenarios in which there is a need for service isolation, can result in an organization or partnership spanning multiple forests. Trust relationships can be used to set up resource sharing between pairs of domains from separate forests or between all domains in two forests. In some cases, plans might call for maintaining separate forests with trust relationships to facilitate collaboration. In other cases, plans might exist to merge forests eventually. In the case of either collaboration or migration, trust relationships that link domains in different forests might be necessary.

Access to applications, data, and other resources that exist in different forests can be enabled through two types of trust relationship:

- An **external trust** is a one-way or two-way, nontransitive trust between two domains in separate Windows 2000 or Windows Server 2003 forests. External trusts can also exist between Windows NT 4.0 domains and domains in Windows 2000 or Windows Server 2003 forests. A domain administrator in one forest can create external trust relationships with domains in other forests.
- A **forest trust** is a one-way or two-way, transitive trust between all domains in two Windows Server 2003 forests. For generalized resource sharing between all domains in both forests, forest administrators can create forest trusts.

Forest and external trust relationships extend the use of security principals that are created in one domain of a forest to a domain or domains in another forest or forests. However, external and forest trusts introduce potential security risks because they allow security principals to cross security boundaries.

When trust relationships exist between domains in separate forests, security implications include the following:

- With external trusts, SIDs from untrusted domains can be added to access tokens, and the SIDs can be accepted by domain controllers in a trusting domain. This risk can be mitigated by SID filtering.
- With forest trusts and external trusts, users from a different forest can be added to administrative groups in the forest root domain or in other domains in the trusting domain. This risk can be mitigated by secure administrative policies and trustworthy high-level administrators.

Avoiding these security risks requires that appropriate administrative policies be in place. In addition, Windows Server 2003 provides default features that mitigate risk from the unauthorized use of SIDs.

**SID History and External Trusts**

The risks that are associated with external trust relationships relate to the SIDs that are included in a trusted user's access token when the user's authentication is accepted by the trusting domain of the external forest.

Because the SID contains the domain affiliation of the user, it is not possible to move the security principal between domains when a user is migrated from one domain to another domain within a forest or between forests. Such a move requires that a new security principal be created in the new domain, with a new user SID that identifies the domain. To ensure that the user does not lose appropriate access to resources, the SIDs that are associated with that user in the source domain are added to the SID history (**sIDHistory**) attribute of the new user account in the target domain. The **sIDHistory** attribute is available on security principal objects only in domains that are operating at one of the following:

- Windows 2000 native mode (Windows 2000 domains)
- Windows 2000 native domain functional level (Windows Server 2003 domains)
- Windows Server 2003 domain functional level (Windows Server 2003 domains)

**Security Risk Posed by SID History**

When a user logs on to a domain, the user's primary SID and the SIDs of groups of which the user is a member are determined by a domain controller in the user's account domain. If the domain of the authenticating domain controller is operating at Windows 2000 native mode or higher, the authenticating domain controller also determines if the user has any values present in the **sIDHistory** attribute, and it includes those SIDs in the authorization data. From the perspective of access control, these SIDs are no different from other group SIDs.

By design, a domain controller in a trusting domain accepts the authorization data of a trusted domain. When the user is from a different forest, accepting all SIDs in the authorization data exposes the trusting domain to potential security threats. Although the authenticating domain may be trusted, a user's **sIDHistory** attribute might contain SIDs that are based on other domains that are not trusted. If a particular SID can be added to a user's **sIDHistory** attribute, that user can gain unauthorized access to resources in the trusting domain, including administrative privileges, which constitutes an elevation-of-privilege threat.

| **Note** |
| --- |
| Exploiting this vulnerability would be a challenge to an attacker. Only the delete operation can be performed using a low-level operation (for example, an LDAP call or ADSI) to directly access the attribute. To add a value to this attribute, an attacker needs, at a minimum, administrative credentials in the trusted domain and the technical ability to modify low-level operating system functions and data structures. Microsoft is not currently aware of any programming interface that allows an attacker — even one with administrative credentials — to introduce a SID into **sIDHistory**. |

**Blocking SIDs from Untrusted Domains**

By default, all new external trusts that are created from domain controllers running Windows Server 2003 or Windows 2000 Service Pack 4 (SP4) enforce SID filtering, which prevents attacks by malicious users who might try to grant elevated user rights to another user account. SID filtering removes any SIDs that are not related to the user's account domain. The trusted domain in this case is said to be quarantined. Quarantining is performed from the trusting domain on a per-domain (and, therefore, per–external trust) basis, and only on domains residing in a different forest.

| **Important** |
| --- |
| Do not attempt to use the SID filtering feature to create a "restricted-access" domain within a forest. Enabling SID filtering between domains in the same forest disrupts the default trust and authentication behavior of a forest. Using SID filtering in this way is likely to lead to problems with programs and with the Active Directory service itself that are difficult to troubleshoot. |

To mitigate the threat posed by **sIDHistory**, all domain controllers in the trusting domain should either be upgraded to Windows Server 2003 or be running at least Windows 2000 SP4 or later before the external trust is created.

To further secure your forest, consider enabling SID filtering on all existing external trusts that are created by domain controllers running Windows 2000 Service Pack 3 (SP3) or earlier. If an external trust was created before you upgraded your Windows 2000 domain to Windows Server 2003, you must either enable SID filtering for that trust or delete the existing trust and create a new one. You can use Netdom.exe (Windows Support Tools) to enable SID filtering on existing external trusts, or you can recreate the external trusts from a domain controller running Windows Server 2003 or Windows 2000 SP 4 or later.

For more information about enabling SID filtering, see Enabling SID Filtering in the Appendix: Procedures section.

For more information about applying SID filtering, see the following articles in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkId=4441:

- 289243, "Forged SID Could Result in Elevated Privileges in Windows 2000"
- 289246, "Forged SID Could Result in Elevated Privileges in Windows NT 4.0"

**Impact of SID Filtering on External Trusts**

SID filtering neutralizes the threat of malicious users in the trusted domain using the sIDHistory attribute to gain elevated privileges in the trusting domain. SID filtering on external trusts can affect your existing Active Directory infrastructure in the following areas:

- SID history data that contains SIDs from any domain other than the trusted domain will be removed from authentication requests that are made from the trusted domain. Removal of these SIDs will result in access being denied to resources that have the user's old SID, as follows:
    - Users that were migrated to the trusted domain lose access to resources in their source domain (if it still exists).
    - Users in the trusted domain who use SID history data for authorization to resources in the trusting domain no longer have access to those resources.
- Your strategy for universal group access control between forests will require changes. If you typically add universal groups from a trusted forest to permissions on shared resources in the trusting domain, SID filtering will have a significant impact on your access control strategy.

For effective SID filtering, follow these guidelines:

- Verify that any universal groups that are assigned to shared resources in the trusting domain were created in the trusted domain. Universal groups must adhere to the same SID filtering guidelines as other security principal objects; that is, the universal group object SID must also contain the domain SID.
- Before assigning access to resources in the trusting domain for users in the trusted domain, confirm that the universal group that contains the trusted domain users was created in the trusted domain. If the universal group in the forest of the trusted domain was not created in the trusted domain, authentication requests that are made from members of that universal group will be filtered and discarded, even though it might contain users from the trusted domain as members.

**Security Risks Associated with Forest Trusts**

When a forest trust is in place, the risk to security comes from high-level administrators in the trusted forest who might have malicious goals. When all domains in one forest are trusted by all domains in another forest, it is possible to add users from the trusted forest to domain local security groups in the trusting forest. Some default administrative security groups have the domain local group type. Adding users from one forest to administrative groups in the forest root domain of another forest compromises the isolation of the forests. The same risk applies for domain-level administrative groups in regional domains within the forests.

By default, forest trusts enforce SID filtering. SID filtering on forest trusts does not prevent migrations to domains within the same forest from using SID history, and it does not affect your strategy for universal group access control within the forest.

**Security Considerations for Resource Access Across Forests**

Even when you create multiple forests for data isolation, service isolation, or both, you might need trust relationships with other forests, or with individual domains in other forests, so that some users can be granted access to resources in a different forest. However, do not include users from other forests in any of the following groups:

- Groups that are responsible for service management or that can manage the membership of service administrator groups
- Groups with administrative control over computers that store protected data
- Groups that have access to protected data or groups that are responsible for the management of user or group objects that have access to protected data

Including users in such groups violates the data and service isolation requirements that created the need for separate forests and exposes the trusting forest to potential data disclosure or service interruption threats.

For information about these administrative groups, see “Designing the Active Directory Logical Structure” in Designing and Deploying Directory and Security Services of the Windows Server 2003 Deployment Kit (or see “Designing the Active Directory Logical Structure” on the Web at http://go.microsoft.com/fwlink/?LinkId=4723).

# Recommendations: Establishing Secure Active Directory Boundaries

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

To help maintain secure Active Directory boundaries, consider implementing the security recommendations that are described earlier in this section. The following lists summarize the security recommendations for establishing and maintaining secure Active Directory boundaries in Windows Server 2003.

**Recommendations for Specifying Security and Administrative Boundaries**
Recommendations for security and administrative boundaries are:

- Determine whether delegation is driven by organizational, operational, or legal requirements.
- Determine whether the requirements indicate the need for autonomy, isolation, or both.
- Assess the level of trust that you have in service administrators (forest owners and domain owners).

**Recommendations for Selecting an Active Directory Structure Based on Delegation Requirements**
Recommendations for selecting an Active Directory Structure are:

- Use the scenarios in "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of the Windows Server 2003 Deployment Kit (or see "Designing the Active Directory Logical Structure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4723) to identify the Active Directory structure that matches your delegation requirements.
- Place outward-facing domain controllers that are deployed in an extranet in a separate forest.

**Recommendations for Establishing Secure Collaboration with Other Forests**
Recommendations for creating secure trusts with domains in other forests and between different Windows Server 2003 forests are:

- Create forest trust relationships between Windows Server 2003 forests only when all forest administrators and all domain administrators are trusted individuals. (A forest functional level of Windows Server 2003 is required in both forests.)
- Before creating an external trust between two domains in isolated forests, be sure that the two domains have a domain functional level of at least Windows 2000 native and that all domain controllers are running at least Windows 2000 Server SP 4, so that SID filtering is enforced by default.

# Chapter 3: Deploying Secure Domain Controllers

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The Active Directory executables and database are stored on the domain controllers in your Active Directory infrastructure. The domain controllers are the servers in your network infrastructure that you must secure to protect Active Directory. If the security of any domain controller in your Active Directory infrastructure is compromised, the security of the entire infrastructure is at risk.

An essential part of deploying your domain controllers is ensuring that they are deployed securely. If you are in the process of deploying your domain controllers, the steps in this section include recommendations for deploying your domain controllers in a manner that enhances their security. If you have already deployed your domain controllers, consider whether to configure your existing domain controllers to reflect the security recommendations in this section.

To deploy secure domain controllers:

1. Secure the domain controller build environment.
2. Establish secure domain controller build practices.
3. Maintain physical security.

# Securing the Domain Controller Build Environment

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The domain controller build environment is the network environment (routers, network segments, switches, and so forth) and physical room (datacenter, secured room, wiring closet, utility closet, and so forth) in which you build your domain controllers. Depending on your IT organizational infrastructure, you might have a centralized datacenter that is secure, both from a network and physical perspective. Alternatively, your IT organization infrastructure might have locations that are not secure from either perspective, such as branch offices.

**Building Domain Controllers in Datacenter Environments**

Whenever possible, build your domain controllers in a secure environment, such as a datacenter. Building your domain controllers in a secure datacenter environment reduces security risks by restricting domain controller access to trusted personnel during the critical build process. This security helps prevent rogue applications, drivers, services, or configurations from being introduced by unauthorized personnel.

If possible, build domain controllers in a datacenter environment, and then ship them to the final location for deployment. This deployment approach is referred to as a staged domain controller deployment.

To help ensure that a domain controller stays secure until deployment, use a trusted shipping method to ship the domain controller to the final location. For example, use a method that requires signatures for the domain controller at the origination and destination locations. Building and shipping the domain controller this way will help protect the integrity of the domain controller.

For more information about building or staging domain controllers, see the Active Directory Branch Office Planning Guide on the Web Resources page at http://go.microsoft.com/fwlink/?LinkId=28523.

**Building Domain Controllers in Branch Office Environments**

If your organization supports branch offices, in some instances you might need to build domain controllers in this relatively insecure environment. For example, you might need to replace a failed domain controller on-site. Table 4 lists recommendations and the corresponding rationales for building domain controllers in branch offices.

**Table 4 Recommendations for Building Domain Controllers in Branch Offices**

| Recommendation | Rationale |
|---|---|
| Limit physical access to domain controllers to trusted personnel only. | To avoid the theft of directory data or the possibility of an altered, less secure domain controller configuration through human intervention. |
| Use an automated method, such as a script, for operating system installation and Active Directory installation. | To reduce the possibility of human intervention, which could result in a vulnerable domain controller configuration. |
| Promote and operate new domain controllers in a restricted access area. | To prevent unauthorized users from compromising the security of the domain controller. |

# Establishing Secure Domain Controller Build Practices

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Secure domain controller build practices are essential to sustained network security. When domain controller builds are planned and implemented according to predictable and repeatable build practices, you can ensure a secure platform on which to run Active Directory. You can ensure predictability by setting a standard order for configurations. You can ensure repeatable builds by automating the build process. Automation adds a large measure of security because it minimizes the possibility of rogue programs, rogue services, and insecure configuration being introduced into the build process through manual intervention.

**Automated Installation Processes**

Automated installation processes can be categorized as image-based and answer-file-based:

- An **image-based installation** is a method of copying, or "cloning," a preconfigured operating system and software applications from a master computer onto destination servers. For the purposes of this chapter, the term image-based installation refers to installations using the System Preparation tool (Sysprep) installation tool.
- An **answer-file-based installation** uses a text file that contains setup instructions. These instructions include the following:
  - Answers to the questions that Windows Setup normally presents during an installation
  - Instructions for configuring operating system settings
  - Instructions for installing applications without user intervention

For the purposes of this guide, the term answer-file-based installation refers to the unattended setup tool Winnt32.exe and the Remote Installation Services (RIS) Setup (Risetup.exe) tool. It is important to note that to use RIS or to perform an installation by using Sysprep or unattended setup from a shared distribution directory, you need reliable, high-bandwidth network connections.

Because image-based installation encompasses all configuration settings, including service and registry settings, this method is recommended over answer-file-based installation methods, which do not ensure consistent and uninterrupted configuration of all services and registry settings.

| **Note** |
|---|
| You cannot use RIS or Sysprep to upgrade an operating system. The only automated installation method that you can use to perform upgrades is unattended installation using Winnt32.exe. |

Sysprep, RIS, and unattended setup are Windows Server 2003 technologies, and Table 5 provides a comparison of their characteristics. Automated setup software from other sources might have combinations of the features that are found in each of these technologies. For more information about automated setup software from other sources, see the documentation that accompanies the automated setup software.

**Table 5 Comparison of Windows Server 2003 Automated Installation Methods**

| Characteristics | Sysprep | RIS | Unattended Setup |
|---|---|---|---|
| Provides image-based installations. | Yes | Yes | No |
| Accommodates a variety of hardware configurations from a single automation script or image. | No | Yes | Yes |
| Requires a high-bandwidth, well-connected network infrastructure. | No | Yes | No |
| Appropriate for datacenter deployments. | Yes | Yes | No |

| | | | |
|---|---|---|---|
| Appropriate for branch office deployments. | Yes | No | Yes |

For more information about using Sysprep, RIS, and unattended setup, see "Overview of Choosing an Automated Installation Method" in Automating and Customizing Installations of the Windows Server 2003 Deployment Kit (or see "Overview of Choosing an Automated Installation Method" on the Web at http://go.microsoft.com/fwlink/?LinkId=4749).

**Using an Image-Based Installation Process**

To automate installation by using an image-based process, you first manually install a single server with all the appropriate settings for your deployment. Then you create an image of that server, which you can use to deploy multiple domain controllers automatically, so that they have the exact same configuration. As part of the image, you can configure setup to start Active Directory installation (Dcpromo.exe) automatically and to use an answer file for the installation.

The general steps for an image-based installation are as follows:

1. Install and configure Windows Server 2003 on a server, and create an image of the installed server. Use Sysprep to define a disk image of the installed server.

2. Use one of the following image-based installation technologies to deploy the image to the target computer:

   - Sysprep
   - RIS
   - Automated setup software from another source

3. After the image is deployed, restart the computer and complete any further setup steps that may be required by the image-based installation technology. For example, complete Sysprep mini setup.

4. After you complete the image-based installation and the computer restarts, log on to the computer. The Active Directory Installation Wizard starts up automatically, and it retrieves information from the answer file. However, because credentials are not provided in the answer file, you are prompted to enter valid credentials to install Active Directory.

5. After Active Directory installation is complete, the computer automatically restarts as a domain controller.

For more information about implementing image-based setup of Windows Server 2003, and before following the instructions in this guide, see "Designing Image-based Installations with Sysprep" and "Designing RIS Installations" in Automating and Customizing Installations of the Windows Server 2003 Deployment Kit (or see "Designing Image-based Installations with Sysprep" on the Web at http://go.microsoft.com/fwlink/?LinkId=4753 and "Designing RIS Installations" on the Web at http://go.microsoft.com/fwlink/?LinkId=4752).

**Using an Answer-File-based Installation Process**

The recommended process for secure automated domain controller installation is the image-based installation process. However, if you must use a non-imaged-based process, unattended setup using an answer file generally presents few security risks and provides consistent server configuration.

For more information about using unattended setup of Windows Server 2003, see "Designing Unattended Installations" in Automating and Customizing Installations of the Windows Server 2003 Deployment Kit (or see "Designing Unattended Installations"on the Web at http://go.microsoft.com/fwlink/?LinkId=4751).

**Creating the Image for Windows Server 2003 Image-Based Installations**

When you use the image-based process to install Windows Server 2003, perform the following tasks to prepare the image computer:

1. Perform a clean installation of Windows Server 2003, and configure server security as specified in the next section, Ensuring Predictable, Repeatable, and Secure Domain Controller Deployments.

2. Configure the computer registry to instruct the computer to run Dcpromo.exe when the computer starts for the first time. Configure the Active Directory installation answer file selections as described in Configuring the Automatic Installation of Active Directory.

3. Create an image of the installed computer. For information about how to create the image, see "Designing Image-based Installations with Sysprep" in Automating and Customizing Installations of the Windows Server 2003 Deployment Kit (or see "Designing Image-based Installations with Sysprep" on the Web at http://go.microsoft.com/fwlink/?LinkId=4753).

**Ensuring Predictable, Repeatable, and Secure Domain Controller Deployments**

Throughout the domain controller deployment process, there are security configuration settings that you need to apply to all domain controllers. To ensure that these security settings are applied uniformly to all domain controllers, implement a predictable, repeatable domain controller deployment process, such as the process specified below:

- Install Windows Server 2003 with secure configuration settings and the latest service packs and hotfixes.
- Create a strong Administrator password.
- Disable NTFS automatic 8.3 name generation.
- Run virus-scanning software on the server.
- Enable only essential services.
- Create a reserve file to enable recovery from disk-space attacks.

**Installing Windows Server 2003 with Secure Configuration Settings, Service Packs, and Hotfixes**

When you create the first master image of a Windows Server 2003 server to be used for installing and promoting multiple domain controllers, apply the configuration settings that are listed in Table 6 to enhance security.

In situations in which your domain controllers already exist, check the recommendations in Table 6 and then modify the domain controller configurations accordingly.

**Table 6 Secure Configuration Settings for Windows Server 2003 Domain Controllers**

| Setting | Rationale |
|---|---|
| Format all partitions as NTFS. | NTFS provides a secure file system. |
| Install only TCP/IP as the transport protocol. | Limiting transport protocols reduces the attack surface on the domain controller. |
| Do not install Simple Mail Transfer Protocol (SMTP) during Windows Server 2003 setup unless you are using it for mail-based, intersite Active Directory replication. | SMTP is not required for normal domain controller operations unless it is used for mail-based, intersite replication. Excluding SMTP reduces the attack surface on the domain controller. By default, SMTP is not selected. |
| Install DNS by selecting it in the Networking category during Windows Server 2003 setup. | Installing DNS during Windows Server 2003 setup ensures that DNS is available in the master image. |

**Note**

In contrast to Windows 2000 Server installation, Internet Information Services (IIS) is not installed by default in Windows Server 2003 installation. IIS is not required on a domain controller, and eliminating IIS reduces the attack surface on the domain controller.

After you complete the installation of Windows Server 2003, obtain any additional security protections that may have been implemented for Windows Server 2003 by applying the most recent service packs and security-related hot fixes from the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkID=7420.

**Creating a Strong Administrator Password**

Use a strong password for the computer local administrator account. The local administrator account becomes the domain Administrator account after the server is promoted to a domain controller. A strong password minimizes the threat of an attacker guessing (cracking) the password and acquiring the credentials of the administrator account (spoofing). A strong password includes all of the following characteristics:

- Contains at least nine characters.
- Does not contain an account name, real name, or company name.
- Does not contain a complete dictionary word.
- Is significantly different from previous passwords. Passwords that increment (Password1, Password2, Password3 ...) are not strong.
- Contains at least one symbol in the first seven characters.
- Contains characters from each of the groups listed in Table 7.

**Table 7 Groups of Characters to Include in Strong Passwords**

| Types of Characters | Examples |
|---|---|
| Uppercase letters | A, B, C ... |
| Lowercase letters | a, b, c ... |
| Numerals | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9 |
| Symbols found on the keyboard (all keyboard characters not defined as letters or numerals) | ` ~ ! @ # $ % ^ & * ( ) _ + - = { } \| [ ] \ : " ; ' < > ? , . / |

Examples of strong passwords are Pf$sw0rR1 and J*p2leO4>F.

**Disabling NTFS Automatic 8.3 Name Generation**

Many viruses and utilities that are used by attackers are 16-bit applications that expect file names to be compatible with 8.3 automatic name generation. Secure domain controllers do not run 16bit applications locally. Therefore, disable 8.3 automatic name generation to prevent these viruses and utilities from compromising security on your domain controllers.

To disable automatic 8.3 name generation, set the following entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem**:

Entry name: **NtfsDisable8dot3NameCreation**

Data type: **REG_DWORD**

Value: **1**

| ⚠**Caution** |
|---|
| Incorrectly editing the registry can have serious, unexpected consequences that might require you to restore the registry or to reinstall Windows. Before you edit the registry, create an emergency startup disk and back up any critical data. For more information, see the Registry Reference for Microsoft Windows Server 2003 at http://go.microsoft.com/fwlink/?LinkId=4543. |

You can create a script that modifies this registry value on all domain controllers in the domain automatically by performing the following tasks:

1. Create the ComputerSearch.vbs script, and copy it to a computer that is a member of the domain.

   For information about how to create the script, see "Identifying Computers to Receive New Registry Settings with ComputerSearch.vbs" in "Appendix B" of the Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations: Part II. To download this document, click the link under "Planning & Deployment Guides" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=18591.

2. Create a list of domain controllers in the domain by typing the following command at a command prompt: **Cscript computersearch.vbs /r:DC**

   Save this list as **ComputerSearch-***date-time***.csv**.

3. Create the ApplyReg.vbs script, and copy it to a computer that is a member of the domain.

   For information about how to create the script, see Applying Registry Settings to a List of Computers with ApplyReg.vbs in the Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations: Part II, Appendix B (http://go.microsoft.com/fwlink/?LinkId=140260).

4. Create a .reg file for the following path, registry entry, and value:

   **[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Filesystem]**

   **"NtfsDisable8dot3NameCreation"=dword:00000001**

   Save the file as Registryfile.reg. For information about how to create the .reg file, see "Creating a .reg File" in Appendix: Procedures section:

   Apply the registry entry to the domain controllers by typing the following command at the command prompt:

   **Cscript ApplyReg.vbs /r:registryfile.reg /f:ComputerSearch-date-time.csv**

   This set of tasks applies the registry changes that are expressed in the file registryfile.reg to all computers that are listed in the file ComputerSearch-*date-time*.csv.

**Running Virus-Scanning Software on the Server**

After you install Windows Server 2003, but before you promote the server to a domain controller, install and run antivirus software on the server to ensure that no viruses have been introduced during the installation of Windows Server 2003.

When you use an image-based process for automating the installation of Windows Server 2003, make sure that you install the virus scanning software as part of the image. For other processes, install and run the virus scanning software immediately after installing Windows Server 2003.

The behavior of some antivirus software can interfere with the proper operation of your domain controllers. Some versions of antivirus software modify the security descriptor while scanning files. When the software scans the files that are stored in the SYSVOL folder, the modification of the security descriptor triggers the File Replication Service (FRS), and which creates high volumes of replication traffic unnecessarily.

For more information about using antivirus software on your domain controller, see Running Virus Scans on Domain Controllers later in this guide.

| **Note** |
| --- |
| If you are starting Dcpromo.exe automatically by using one of the RunOnce methods discussed later in this section, modify your scripts to run the virus scanning first and then start Dcpromo.exe. |

**Enabling Only Essential Services**

Every service that is running on a domain controller provides a potential target for an attack that can compromise domain controller security. To reduce the surface area of attack, it is prudent to enable only those services that are essential to normal domain controller function.

Table 8 lists the services that are installed during Windows Server 2003 installation when you follow the security recommendations in the Installing Windows Server 2003 with Secure Configuration Settings, Service Packs, and Hotfixes section. Table 8 also provides the recommended startup type for each service.

| **Note** |
| --- |
| For new domain controller deployments, configure the startup types for these services before creating the master image for image-based deployments.<br><br>For existing domain controllers, combine Table 8 and Table 9 to determine the recommended startup type of these services. |

**Table 8 Recommended Services to Install on a Windows Server 2003 Domain Controller**

| Service Name | Default Startup Type | Recommended Startup Type | Comment |
|---|---|---|---|
| Alerter | Disabled | (No change) | Notifies selected users and computers of administrative alerts. |
| Application Layer Gateway Service | Manual | (No change) | Provides support for application-level protocol plug-ins, and enables network/protocol connectivity. |
| Application Management | Manual | (See comment) | Provides software installation services for applications that are deployed by using Add or Remove Programs. On dedicated domain controllers, this service can be disabled to prevent unauthorized installation of software. |
| Automatic Updates | Automatic | (See comment) | Provides the download and installation of critical Windows updates, such as security updates or hotfixes. This service can be disabled when automatic updates are not performed on the domain controller. |
| Background Intelligent Transfer Service | Manual | (See comment) | Provides a background file transfer mechanism and queue management, and it is used by Automatic Update to automatically download programs (such as security updates). This service can be disabled when automatic updates are not performed on the domain controller. |
| ClipBook | Disabled | (No change) | Enables the Clipbook Viewer to create and share "pages" of data to be reviewed by remote users. |
| COM+ Event System | Manual | (No change) | Provides automatic distribution of events to Component Object Model (COM)+ components. |
| COM+ System Application | Manual | (No change) | Manages the configuration and tracking of components that are based on COM+. |
| Computer Browser | Automatic | (See comment) | Maintains the list of computers on the network, and supplies the list to clients that request the list. If you do not have earlier versions of applications or clients that use this feature, set the startup type for this service to Manual. |
| Cryptographic Services | Automatic | (No change) | Provides three management services:<br>• Catalog Database Service, which confirms the signatures of Windows files.<br>• Protected Root Service, which adds and removes Trusted Root Certification Authority certificates from this computer.<br>• Key Service, which helps enroll this computer for certificates. |
| DHCP Client | Automatic | (See comment) | Required to update resource records by using dynamic update. If your organization does not use dynamic IP addresses for your domain controllers, the startup type for this service can be set to Manual. |
| Distributed File System | Automatic | (No change) | Manages logical volumes that are distributed across a LAN or wide area network (WAN). This service is required for the Active Directory SYSVOL share. |
| Distributed Link Tracking Client | Automatic | Disabled | Maintains links between NTFS v5 files within the domain controllers and other servers in the domain. Disable this service on dedicated domain controllers. |
| Distributed Link Tracking Server | Disabled<br>*Automatic** | (No change)<br>*Disabled** | Tracks information about files that are moved between NTFS v5 volumes throughout a domain. Disable this service on dedicated domain controllers. |
| Distributed Transaction Coordinator | Automatic | (No change) | Coordinates transactions that span multiple resource managers, such as databases, message queues, and file systems. |

| | | | |
|---|---|---|---|
| DNS Client | Automatic | (No change) | Provides resolution of DNS names. |
| DNS Server | Automatic | (See comment) | Required for Active Directory–integrated DNS zones. If the domain controller is not a DNS server, this service is Disabled. |
| Error Reporting Service | Automatic | (See comment) | Collects, stores, and reports unexpected application crashes to Microsoft. If you do notplan to report critical operating system and application faults to Microsoft for analysis, you can disable this service. |
| Event Log | Automatic | (No change) | Writes event messages that are issued by Windows-based programs and components to the event log files. |
| *Fax** | *Automatic** | *Disabled** | Provides the ability to send and receive faxes through fax resources that are available on the domain controller and on the network. On dedicated domain controllers, this service can be disabled, because sending and receiving faxes is not a normal function of a domain controller. |
| File Replication Service | Automatic | (No change) | Enables files to be copied automatically and maintained simultaneously on multiple computers. This service is used to replicate SYSVOL among all domain controllers. |
| Help and Support | Automatic | (No change) | Enables Help and Support Center to run on this computer. Also used to collect historical machine data on a daily basis for Microsoft Product Support Services incident escalation. |
| HTTP SSL | Manual | (No change) | Implements the secure hypertext transfer protocol (HTTPS) for the HTTP service by using Secure Sockets Layer (SSL). |
| Human Interface Device Access | Disabled | (No change) | Enables generic input access to Human Interface Devices (HID), which activates and maintains the use of predefined hot buttons on keyboards, remote controls, and other multimedia devices. |
| IIS Admin Service* | Automatic* | Disabled* | Provides administration of Web and File Transfer Protocol (FTP) services through the IIS snap-in. IIS is not required for domain controller operations. Eliminating IIS on dedicated domain controllers reduces the attack surface. |
| IMAPI CD-Burning COM Services | Disabled | (No change) | Manages CD recording by using Image Mastering Applications Programming Interface (IMAPI). |
| Indexing Service | Disabled<br>*Manual** | (No change)<br>*Disabled** | Indexes content and properties of files on the domain controller to provide rapid access to the file through a flexible querying language. On dedicated domain controllers, disable this service to prevent users from searching files and file content if sensitive files and folders are inadvertently indexed. |
| Internet Connection Firewall (ICF)/Internet Connection Sharing (ICS) | Disabled<br>*Manual** | (No change)<br>*Disabled** | Provides network address translation (NAT), addressing and name resolution, and intrusion detection when connected through a dial-up or broadband connection. On dedicated domain controllers, disable this service to prevent inadvertent enabling of NAT, which would prevent the domain controller from communicating with the rest of the network. |
| Intersite Messaging | Automatic | (No change) | Required by Active Directory, Distributed File System (DFS), and Netlogon for intersite replication. |
| IPsec Services | Automatic | (No change) | Provides management and coordination of Internet Protocol security (IPsec) policies with the IPsec driver. |
| Kerberos Key Distribution Center | Disabled | (No change) | Provides the ability for users to log on by using the Kerberos V5 authentication protocol. Required on domain controllers. (See Table 9.) |

| | | | |
|---|---|---|---|
| License Logging | Disabled<br>*Automatic** | (No change)<br>*Disabled** | Monitors and records client access licensing for portions of the operating system, such as IIS, Terminal Services, and file and print sharing, and for products that are not a part of the operating system, such as Microsoft® SQL Server 2000 or Microsoft® Exchange 2000 Server. On a dedicated domain controller, this service can be disabled. |
| Logical Disk Manager | Automatic | (No change) | Required to ensure that dynamic disk information is up to date. |
| Logical Disk Manager Administrative Service | Manual | (No change) | Required to perform disk administration. |
| Messenger | Disabled | (No change) | Transmits net sends and Alerter service messages between clients and servers. |
| Microsoft Software Shadow Copy Provider | Manual | (See comment) | Manages software-based volume shadow copies taken by the Volume Shadow Copy service. Ntbackup.exe and other non-Microsoft backup applications can use this service to back up the domain controller by using Windows Server 2003 volume-shadow-copy capabilities. If you do not intend to back up the server, you can disable this service. |
| Net Logon | Manual | (No change) | Maintains a secure channel between the domain controller and other domain controllers, member servers, and workstations in the same domain and in trusting domains. |
| NetMeeting Remote Desktop Sharing | Disabled<br>*Manual** | (No change)<br>*Disabled** | Eliminates a potential security threat by allowing domain controller remote administration through NetMeeting. |
| Network Connections | Manual | (No change) | Manages objects in the Network Connections folder. |
| Network DDE | (Disabled) | (No change) | Provides network transport and security for Dynamic Data Exchange (DDE) for programs that are running on the domain controller. |
| Network DDE DSDM | (Disabled) | (No change) | Used by Network DDE. This service is disabled when Network DDE is disabled. |
| Network Location Awareness (NLA) | Manual | (No change) | Collects and stores network configuration and location information, and notifies applications when this information changes. |
| NTLM Security Support Provider | Manual | (No change) | Provides security to remote procedure call (RPC) programs that use transports other than named pipes, and enables users to log on using the NTLM authentication protocol. |
| Performance Logs and Alerts | Manual | (See comment) | Collects performance data for the domain controller, writes the data to a log, or generates alerts. The startup type for this service can be set to Automatic when you want to log performance data or generate alerts without an administrator being logged on. If you are not collecting performance logs, this service can be disabled. |
| Plug and Play | Automatic | (No change) | Required to automatically recognize and adapt to changes in the domain controller hardware with little or no user input. |
| Portable Media Serial Number Service | Manual | Disabled | Retrieves the serial number of any portable media player that is connected to this computer. A domain controller should not need a portable media player. |

| | | | |
|---|---|---|---|
| Print Spooler | Automatic | (See comment) | Manages all local and network print queues, and controls all print jobs. Can be disabled on dedicated domain controllers where no printing is required. |
| Protected Storage | Automatic | (No change) | Protects storage of sensitive information, such as private keys, and prevents access by unauthorized services, processes, or users. This service is used on domain controllers for the smart card logon process. |
| Remote Access Auto Connection Manager | Manual | (See comment) | Detects unsuccessful attempts to connect to a remote network or computer, and provides alternative methods for connection. This service can be disabled on dedicated domain controllers where no VPN or dial-up connections are initiated. |
| Remote Access Connection Manager | Manual | (See comment) | Manages VPN and dial-up connections from the domain controller to the Internet or other remote networks. This service can be disabled on dedicated domain controllers where no VPN or dial-up connections are initiated. |
| Remote Desktop Help Session Manager | Manual | (No change) | Manages and controls Remote Assistance. If you do not need the Remote Assistance feature on the domain controller, disable this service. |
| Remote Procedure Call (RPC) | Automatic | (No change) | Serves as the RPC endpoint mapper for all applications and services that use RPC communications. |
| Remote Procedure Call (RPC) Locator | Manual | (See comment) | Enables RPC clients that are using the RpcNs* family of application programming interfaces (APIs) to locate RPC servers and manage the RPC name service database. This service can be disabled if no applications use the RpcNs APIs. |
| Remote Registry | Automatic | (No change) | Enables remote users to modify registry settings on the domain controller, provided that the remote users have the required permissions. By default, only members of the Administrators and Backup Operators groups can access the registry remotely. |
| Removable Storage | Manual<br>*Automatic** | (See comment) | Manages and catalogs removable media, and operates automated removable media devices, such as tape autoloaders or CD jukeboxes. You can disable this service if you do notback up to tapes or optical disks that are physically attached to this server or if you do not use backup to create an Automated System Recovery (ASR) set. |
| Resultant Set of Policy Provider | Manual | (No change) | Enables a user to connect to a remote computer, access the Windows Management Instrumentation (WMI) database for that computer, and either verify the current Group Policy settings for that computer or check settings before they are applied. If you do notuse Resultant Set of Policy (RSoP) planning mode to simulate the application of Group Policy settings, you can disable this service. |
| Routing and Remote Access | Disabled | (No change) | Enables LAN-to-LAN, LAN-to-WAN, VPN, and NAT routing services. |
| Secondary Logon | Automatic | (No change) | Allows you to run specific tools and programs with different privileges than your current logon provides. |
| Security Accounts Manager | Automatic | (No change) | A protected subsystem that manages user and group account information. |
| Server | Automatic | (No change) | Provides RPC support for file, print, and named pipe sharing over the network. |
| Shell Hardware Detection | Automatic | Disabled | Provides notifications for Auto Play hardware events. The Windows Image Acquisition (WIA) service depends on this service. Because the WIA service is disabled, this service can also be disabled. |

| | | | |
|---|---|---|---|
| *Simple Mail Transfer Protocol (SMTP)* * | *Automatic** | *Disabled** | Transports electronic mail across the network. In the rare event that you are using mail-based intersite replication, you should enable this service. Otherwise, leave it disabled at all times. |
| Smart Card | Manual | (No change) | Manages and controls access to a smart card that is inserted into a smart card reader that is attached to the domain controller. |
| *Smart Card Helper** | *Manual** | *(No change)** | Provides support for legacy, non-plug-and-play smart card readers. |
| Special Administrator Console Helper | Manual | Disabled | Allows administrators to remotely access a command prompt by using Emergency Management Services (EMS). Enable this service only if you want Windows command prompts for connections over out-of-band management hardware that uses the EMS special administration console. |
| System Event Notification | Automatic | (No change) | Monitors system events and notifies subscribers to the COM+ Event System of these events. |
| Task Scheduler | Automatic | (No change) | Provides the ability to schedule automated tasks on the domain controller. If you do not plan to run batch jobs on the domain controller, you can disable this service. |
| TCP/IP NetBIOS Helper | Automatic | (No change) | Provides support for the NetBIOS over TCP/IP (NetBT) service and network basic input/output system (NetBIOS) name resolution for clients. |
| Telephony | Manual | (See comment) | Provides Telephony API (TAPI) support of client programs that control telephony devices and IP-based voice connections. This service can be disabled on dedicated domain controllers where TAPI is not used by applications. |
| Telnet | Disabled<br>*Manual** | (No change)<br>*Disabled** | Enables a remote user to log on and run applications from a command line on the domain controller. Enable Telnet only when it is used for remote administration for branch offices or remotely administered domain controllers. Terminal Services is the recommended method for remote administration. |
| Terminal Services | (Manual) | (See comment) | Enables multiple remote users to be connected interactively to the domain controller, and provides display of desktops and the ability to run applications. To reduce the surface area of attack, disable Terminal Services unless it is used for remote administration for branch offices or remotely administered domain controllers. |
| Terminal Services Session Directory | Disabled | (No change) | Enables a user connection request to be routed to the appropriate terminal server in a cluster. |
| Themes | Disabled | (No change) | Provides user experience theme management. |
| Uninterruptible Power Supply | (Manual) | (See comments) | Manages an uninterruptible power supply (UPS) that is connected to the domain controller by a serial port. If you don't have a UPS device that is connected to the domain controller, you can disable this service. |
| Upload Manager | Manual | Disabled | Manages the synchronous and asynchronous file transfers between clients and servers on the network. Driver data is anonymously uploaded from these transfers and then used by Microsoft to help users find the drivers they need. |
| *Utility Manager** | *Manual** | *Disabled** | Provides faster access to some accessibility tools, such as Magnifier, Narrator, and On-Screen Keyboard, and also displays the status of the tools or devices that it controls. Disable Utility Manager unless you require these special accessibility tools. |
| Virtual Disk Service | Manual | (No change) | Provides software volume and hardware volume management service |

| | | | |
|---|---|---|---|
| Volume Shadow Copy | Manual | (See comment) | Manages and implements volume shadow copies that are used for backup and other purposes. Ntbackup.exe and other non-Microsoft backup applications use this service to back up this domain controller by using Windows Server 2003 volume shadow copy capabilities. If you do not intend to back up this server, you can disable this service. |
| WebClient | Disabled | (No change) | Enables Windows-based programs to create, access, and modify Internet-based files. |
| Windows Audio | Automatic<br>*Disabled** | Disabled<br>*(No change)** | Manages audio devices for Windows-based programs. |
| Windows Image Acquisition | Disabled | (No change) | Provides image acquisition services for scanners and printers. |
| Windows Installer | Manual | (No change) | Adds, modifies, and removes applications that are provided as a Windows Installer (.msi) package. |
| Windows Management Instrumentation | Automatic | (No change) | Provides a common interface and object model to access management information about the domain controller through the WMI interface. |
| Windows Management Instrumentation Driver Extensions | Manual | (No change) | Monitors all drivers and event trace providers that are configured to publish WMI or event trace information. |
| Windows Time | Automatic | (No change) | Sets the domain controller clock and maintains date and time synchronization on all computers in the network. |
| WinHTTP Web Proxy Auto-Discovery Service | Manual | (No change) | Implements the Web Proxy Auto-Discovery (WPAD) protocol out of process. Advantages of implementing WPAD out of process include isolating the execution of untrusted JavaScript to a low-privileged (Local Service) account, and also avoids loading the COM runtime into the client's process. |
| Wireless Configuration | Automatic | (See comments) | Enables automatic configuration for IEEE 802.11 adapters. If this domain controller does not use a wireless network or does not use the 802.11 authentication protocol, this service can be disabled. |
| WMI Performance Adapter | Manual | (No change) | Provides performance library information from WMI providers to clients on the network. |
| Workstation | Automatic | (No change) | Creates and maintains client network connections to remote servers. |
| *World Wide Web Publishing Service** | *Disabled** | *(No change)** | Provides Web connectivity and administration through the Internet Information Services Manager. |

* Services (or default settings) that are present when you upgrade a domain controller that is running Windows 2000 to Windows Server 2003 but that are not present when you perform a new installation of Windows Server 2003.

| **Note** |
|---|
| The antivirus software installed earlier in the deployment process and described in this section might run as a service in Windows Server 2003. Do not change the default configuration of your antivirus software. |

Table 9 lists the changes to the service startup configuration when a server running Windows Server 2003 is promoted to a domain controller. This table is provided as a reference, and you can combine it with the information in Table 8 to arrive at the final list of services to have running on your domain controller.

**Table 9 Recommended Changes to Service Startup Type After Promotion to Domain Controller**

| Service Name | Default Startup Type | Recommended Startup Type | Comment |
|---|---|---|---|
| Distributed Link Tracking Client | Manual | Disabled | Maintains links between NTFS v5 files within the domain controllers and other servers in the domain. Disable Distributed Link Tracking Client on dedicated domain controllers. |
| Distributed Link Tracking Server | Disabled | (No change) | Tracks information about files that are moved between NTFS v5 volumes throughout a domain. Disable this service on dedicated domain controllers. |
| File Replication Service | Automatic | (No change) | Enables files to be automatically copied and maintained simultaneously on multiple computers. This service is required to replicate SYSVOL between all domain controllers. |
| Intersite Messaging | Automatic | (No change) | Required by Active Directory, DFS, and NETLOGON for intersite replication and automatic site coverage. |
| Net Logon | Automatic | (No change) | Maintains a secure channel between the domain controller and other domain controllers, member servers, and workstations in the same domain and in trusting domains. |
| Remote Procedure Call (RPC) Locator | Manual | (See comment) | Enables RPC clients using the RpcNs* family of APIs to locate RPC servers and manage the RPC name service database. This service can be disabled if no applications use the RpcNs* APIs. |
| Windows Management Instrumentation | Automatic | (No change) | Provides a common interface and object model to access management information about the domain controller through the WMI interface. |
| Windows Time | Automatic | (No change) | Sets the domain controller clock and maintains date and time synchronization on all computers in the network. |

The services that are recommended in Table 8 and Table 9 are only the services that are required to support a dedicated domain controller. Although it is recommended that you have dedicated domain controllers, if your domain controller serves other roles, such as file server or print server, you may need to enable other services for proper operation.

**Creating a Reserve File to Enable Recovery from Disk-Space Attacks**

Many security attacks attempt to consume the system resources of the targeted system. One of the commonly attacked system resources is available disk space. Available disk space can be exhausted by the addition of a large number of objects to the directory by a malicious user or administrator. You can delete the added objects, but Active Directory requires that deleted objects continue to exist in the directory as tombstones for an extended period of time to allow the deletion to replicate. Therefore, the disk space that is consumed by the deleted objects cannot be reclaimed until the tombstone lifetime has expired, which is 60 days by default.

You can mitigate the effects of this type of attack by implementing preventive measures. You can provide for a fast recovery by creating a reserve file on the same disk volume as the Active Directory database (Ntds.dit). A reserve file is simply a large file that takes up available disk space. In the event that an attacker exhausts all disk space by adding a large number of objects to the directory, you can delete the reserve file to quickly restore normal operation while the rogue objects inside Active Directory are identified and removed. For more information about performing this task, see "Creating a Reserve File" in the Appendix: Procedures section.

| **Note** |
|---|
| The ability to assign object ownership quotas in Windows Server 2003 eliminates most of the risk associated with this threat, because administrators can assign a maximum value to the number of objects that a user can create (own). However, this threat is not totally eliminated by setting quotas, because quotas cannot be set to limit the size of an object. In addition, a malicious administrator can create an unlimited number of objects. |

For information about how to recognize and delete illegally added objects, see "Finding and Deleting Rogue Objects" in the Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations: Part II. To download this document, click the link under "Planning & Deployment Guides" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=18591.

**Configuring the Automatic Installation of Active Directory**

You can configure a server to automatically install Active Directory by including instructions to run the Active Directory Installation Wizard (Dcpromo.exe) the first time that the server starts after the automated installation of Windows Server 2003. If you automate Active Directory installation, always install the first domain controller in a domain manually, and then use an automated installation method for additional domain controllers in that domain.

Automatically running Dcpromo.exe ensures that the promotion to a domain controller occurs immediately after the installation of Windows Server 2003. This configuration reduces the potential for introducing unauthorized files or executables before the server is promoted. Just before making your image, use one of the following methods to configure the server to automatically start Dcpromo.exe:

- When using non-Microsoft tools or RIS-based deployments, add the following entry under the registry key **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion**

  **RunOnce:**

  Entry name: **dcpromo**

  Data type: **REG_SZ**

  Value: **dcpromo.exe /answer:**<*answer file path*>

  | ⚠**Caution** |
  | --- |
  | Incorrectly editing the registry can have serious, unexpected consequences that might require you to restore the registry or to reinstall Windows. Before you edit the registry, create an emergency startup disk and back up any critical data. For more information, see the Registry Reference for Microsoft Windows Server 2003 at http://go.microsoft.com/fwlink/?LinkId=4543. |

- When using Sysprep, modify the [GuiRunOnce] section of the Sysprep.inf answer file before running Sysprep.

- The following is an example of a Sysprep.inf answer file that is modified to start Dcpromo.exe automatically the first time the server starts:

  ```
  [GuiRunOnce]
  Command0 = "dcpromo /answer:<answer file path>"
  ```

  If you are creating an image of a server that you plan to use for imaging multiple domain controllers, place the Dcpromo.exe answer file on a floppy disk, so that one image can be used to deploy all domain controllers. Also, make sure that the Dcpromo.exe answer file path designates the floppy disk drive.

**Using Unattended Active Directory Installation**

Automating Active Directory installation generally increases its predictability and consistency by eliminating the need for manual intervention. You can automate Dcpromo.exe by using an answer file. The resulting process is known as unattended installation.

The answer file that you use to install Windows has a **[DCInstall]** section for domain controller installation. If you set the registry to automatically start Dcpromo.exe following the installation of the operating system, or if you modify the **[GuiRunOnce]** section of the Sysprep.inf answer file before running Sysprep, the installation server will run the **dcpromo** command and use the specified answer file. Instructions for creating an answer file for an Active Directory installation are located in the Deploy.cab file in the Support\Tools folder on the Windows Server 2003 operating system CD. Inside the Deploy.cab file, open Ref.chm to access the Unattend.txt file. For the **[DCInstall]** portion of the Unattend.txt file, double-click **Unattend.txt**, and then double-click **domain controllers**. Create a separate file that contains the **[DCInstall]** contents, and save it as the answer file for Active Directory installation.

The command for unattended Active Directory installation is:

**dcpromo /answer:** *answer file path*

Because the answer file is stored on the floppy disk in plaintext, do not include an administrative account name or a password in this file. You must enter this information during the automated Active Directory installation process. The answer file automatically provides all other parameters.

For more information about installing Active Directory from an answer file or from backup media, see "Deploying Windows Server 2003 Regional Domains" in Designing and Deploying Directory and Security Services of the Windows Server 2003 Deployment Kit (or see "Deploying Windows Server 2003 Regional Domains" on the Web at http://go.microsoft.com/fwlink/?LinkId=9691).

#### Selecting Secure Active Directory Configuration Settings

During domain controller promotion, a number of configuration settings should be applied to the server to enhance security. These settings are specified in the Dcpromo.exe answer file, that is, in the **[DCInstall]** section of the Unattend.txt file.

| **Note** |
| --- |
| Many of the configuration settings that are required for domain controller promotion depend on whether you are installing the first domain controller or an additional domain controller in the forest or domain. For this reason, a best practice recommendation is to install the first domain controller in a forest or domain manually. |

#### Specifying Locations for the Active Directory Database, Logs, and SYSVOL

By default, Dcpromo.exe places the database, log files, and SYSVOL folder on the system volume. During the promotion of a server to a domain controller, the Dcpromo.exe answer file can specify an alternative location for the Active Directory database (Ntds.dit), log files, and SYSVOL folder.

Table 10 provides recommended alternative drive locations for these Active Directory components for enhanced security and improved performance.

**Table 10 Recommended Drive Locations for Active Directory Components**

| Recommendation | Security Rationale |
| --- | --- |
| Place the database and SYSVOL folder on the same dedicated physical drive. This drive should not contain the system volume. | The system volume is a common target of disk-space attacks, such as spooling large print jobs. Putting the database on a separate drive makes it immune to damage caused by attacks that target the system volume.<br><br>Placing the database and SYSVOL in a location other than the default location reduces the likelihood of disk-space attacks.<br><br>Attacks can be performed against IIS that allow unauthorized users to navigate the folder structure of its disk volume. Because Web sites are typically stored on the system volume, this increases the risk of information disclosure on this volume. Although dedicated domain controllers are recommended, this cannot always be adhered to in a branch office situation.<br><br>This recommendation improves the reliability of Active Directory operations recovery after a disk-space attack. For more information, see Creating a Reserve File to Enable Recovery from Disk-Space Attacks. |
| Place the log files on a dedicated physical drive that does not contain the system volume. | This recommendation has no security implications, but it will improve performance. The system volume typically includes the paging file, which has high disk utilization. |

On existing domain controllers, move the Active Directory database, log files, and SYSVOL folder to a dedicated physical drive, for the reasons described in Table 10.

For information about:

- Moving the database and log files, see "Moving the Directory Database Files to a Local Drive" in the Active Directory Operations Guide on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=18545.
- Moving the SYSVOL folder, see "Moving SYSVOL Manually" in the Active Directory Operations Guide on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=18545.

#### Configuring Pre–Windows 2000 Compatibility (Anonymous Access to Active Directory) for New Domains

During the promotion of the first domain controller in a domain, one configuration setting that significantly affects Active Directory security is **Permissions compatible with pre–Windows 2000 servers**. This setting enables backward compatibility for certain applications that need to query the directory by using anonymous access.

Anonymous access to Active Directory is disabled by default when you create a new Windows Server 2003 domain. In this case, the default setting is **Permissions compatible only with Windows 2000 or Windows Server 2003 servers**. Unless

your environment has servers or domain controllers that are running Windows NT 4.0 or earlier versions of Windows, this default setting is the recommended selection.

For information about disabling anonymous access to Active Directory, see the next section, Determining the Need for Anonymous Access to Active Directory Data.

**Determining the Need for Anonymous Access to Active Directory Data**

You need Pre–Windows 2000 compatibility if you have applications or services that gain access to the directory anonymously (as opposed to being authenticated) and that require Read permissions on user (or inetOrgPerson), computer, and group objects in the directory. By default, Active Directory does not grant explicit permissions on objects in the directory to the special identity Anonymous Logon, which represents anonymous connections. However, the group Pre–Windows 2000 Compatible Access is assigned Read permissions on the domain root, as well as on user, computer, and group objects. In Windows Server 2003 Active Directory, when you enable Pre–Windows 2000 compatibility, the special identity Anonymous Logon is added as a member of the group Pre–Windows 2000 Compatible Access. Therefore, applications and services that require anonymous access to Active Directory — basically, any user or service with network access — can read these objects.

Applications or services that require anonymous access to Active Directory on a domain controller that is running Windows Server 2003 are those that are running in the following security contexts:

- Local System on a server that is running Windows NT 4.0 within the forest
- Local System on a server that is running Windows NT 4.0, Windows 2000, or Windows Server 2003 in a domain that exists in a different forest and with which a local domain has an external trust relationship
- NetworkService on a server that is running Windows Server 2003 in a domain that exists in a different forest and with which a local domain has an external trust relationship, and the authenticating domain has a domain functional level of Windows 2000 mixed

For example, Microsoft® SQL Server version 6.0 (in Mixed or Windows Integrated security modes) and Routing and Remote Access Service, both running on Windows NT 4.0, require anonymous access to Active Directory to authenticate users and enumerate their group memberships.

Use the following approach to managing anonymous Active Directory access:

1. Perform tests to determine whether any applications are using anonymous access to query Active Directory, as described in Testing for Anonymous Active Directory Access.
2. Eliminate anonymous access requirements, and disable anonymous access, as described in Eliminating the Requirement for Anonymous Active Directory Access.

For more information about pre–Windows 2000 compatibility, see articles 257942, "Error Message: Unable to Browse the Selected Domain Because the Following Error Occurred ..."; 303973, "HOW TO: Add Users to the Pre–Windows 2000 Compatible Access Group"; 240855, "Using Windows NT 4.0 RAS Servers in a Windows 2000 Domain"; 254311, "Enable Windows NT 4.0–Based RAS Servers in a Windows 2000–Based Domain"; and 266712, "SMS: Security Based on Global Groups Fails in Windows 2000 Domains," in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkId=4441.

**Testing for Anonymous Active Directory Access**

If you are installing a new domain and you want to determine whether anonymous access is required, prepare a test environment to determine whether any applications or services require anonymous access to Active Directory, as described below. If you already have a production domain in which pre–Windows 2000 compatibility is enabled, perform steps 3 through 5 to monitor and assess anonymous access to the domain.

1. When you deploy the first domain controller in the test domain, select the option **Permissions compatible with pre–Windows 2000 servers**.
2. In the test domain, include at least one instance of each application that is running in your organization.
3. Enable security monitoring for anonymous access on all domain controllers. For more information about performing this task, see "Enabling Monitoring for Anonymous Active Directory Access" in the Appendix: Procedures section.
4. Analyze the security logs for computers that initiate anonymous access to the directory. For more information about

performing this task, see "Enabling Monitoring for Anonymous Active Directory Access" in the Appendix: Procedures section.

5. Identify the applications or services that are running on the computers from which anonymous access was initiated.

**Eliminating the Requirement for Anonymous Active Directory Access**

After you have determined the applications or services that require anonymous access, identify the applications or services that can be upgraded to versions that do not require anonymous access, and perform the following tasks:

1. Eliminate the requirement for anonymous access to Active Directory by upgrading the applications or services that initiate anonymous access.

2. Monitor the security logs for servers that initiate anonymous access to Active Directory for another 30 days.

3. After reviewing the security logs and ensuring that no application or service is gaining anonymous access to Active Directory, disable anonymous access to Active Directory data as described in "Disabling Anonymous Active Directory Access" below.

Typically, a newer version of the application or service (for example, Routing and Remote Access in Windows 2000 and Windows Server 2003) does not require anonymous access. Whenever possible, upgrade to a newer version of the application or service that is running under Windows 2000 or later operating systems so that you can disable pre–Windows 2000 compatibility.

**Disabling Anonymous Active Directory Access**

After you are certain that no applications or services are establishing anonymous connections to your domain controllers, you can disable anonymous access.

When you create a new Windows Server 2003 domain, simply accept the default option setting **Permissions compatible only with Windows 2000 or Windows Server 2003 servers**.

In an existing Windows Server 2003 domain that has pre–Windows 2000 compatibility enabled, you must remove the Everyone group and Anonymous Logon from the Pre–Windows 2000 Compatible Access group. To ensure that anonymous access is removed in the most efficient and timely manner, perform the following tasks:

1. On the domain controller that holds the primary domain controller (PDC) emulator operations master role, remove Everyone and Anonymous Logon from the Pre–Windows 2000 Compatible Access group, leaving Authenticated Users as the only member.

2. Either wait for changes to replicate to all domain controllers in the domain, or force replication by using the following command (switches are case-sensitive):

   **repadmin /syncall /d/e/P/q** *DNS_name_of_PDC_emulator_operations_master*

3. On every domain controller in the domain, run the following command:

   **net session /delete**

**Running Virus Scans on Domain Controllers**

Even after the domain controller is promoted, continue to run virus scans and to obtain regular virus signature updates from your antivirus software vendor. However, before initiating regular antivirus scanning, be aware that some antivirus software can interfere with the proper operation of domain controllers by:

- Interfering with directory database and log file access by the Extensible Storage Engine (ESE).

- Interfering with FRS database and log file access by ESE.

- Causing excessive replication by FRS.

These issues with running antivirus software on domain controllers are addressed by the recommendations provided in the following three sections.

**Preventing Interference Between Active Directory Database and Log File Access and Virus Scans**
ESE, which underlies Active Directory, opens database and log files in exclusive mode. As a result, if antivirus software opens one of the database or log files, ESE fails when it tries to open the same file. Alternatively, the antivirus software cannot open these files for scanning if they have already been opened by ESE.

Furthermore, these database and log files have internal checksums that, when altered through file changes by the antivirus software, can either cause an access error to occur or cause the database to fail on restart or restore. In all cases, the result is that Active Directory fails on the domain controller.

To run antivirus software on domain controllers, administrators must configure the antivirus software to exclude from scanning the Active Directory database and log files that are specified in Table 11.

**Table 11 Active Directory Files to Exclude from Antivirus Scanning and Registry Values Specifying Their Location**

| **Exclude from Scan** | **In HKLM\System\CurrentControlSet\Services\NTDS\Parameters** |
|---|---|
| Ntds.dit | **DSADatabaseFile** |
| Edb*.log, Edb*.log, Res1.log, and Res2.log | **DatabaseLogFilesPath** |
| Temp.edb and Edb.chk | **DSAWorkingDirectory** |

Excluding the files in Table 11 from regular virus scanning does not increase a domain controller's vulnerability to a virus attack. Viruses tend to attach to files that are executed, such as binaries or scripts, rather than to data files. Dedicated domain controllers with no user-modifiable content are therefore less vulnerable to common virus attacks.

**Preventing Interference Between FRS Database and Log File Access and Virus Scans**
The exclusive mode in which ESE operates causes interference with scanning the FRS database and FRS log files, in the same manner as described for Active Directory database and log files in the Preventing Interference Between Active Directory Database and Log File Access and Virus Scans section.

To run antivirus software on domain controllers, administrators must configure the antivirus software to exclude from regular virus scanning the FRS directory and log files that are specified in Table 12.

**Table 12 FRS Files and Folders to Exclude from Virus Scanning and Registry Values Specifying Their Location**

| **Exclude Files from Scan** | **In HKLM\System\CurrentControlSet\Services\NtFrs\Parameters** |
|---|---|
| jet\sys\Edb.chk, jet\Ntfrs.jdb and jet\log*.log | **Working Directory** |
| Log*log | **DB Log File Directory** |
| **Exclude Folders from Scan** | **In HKLM\Sysem\CurrentControlSet\Services\NtFrs\Parameters\ReplicaSets\GUID** |
| *replica_root* | **Replica Set Root** |
| *staging_directory* | **Replica Set Stage** |
| *preinstall_directory* | *Replica_root* **\DO_NOT_REMOVE_Ntfrs_Preinstall_Directory** |

**Preventing Virus Scans from Triggering Excessive FRS Replication**
Domain controllers that are running Windows Server 2003 use FRS to replicate system policy and logon scripts that reside in the shared SYSVOL folder. Antivirus software can interfere with the normal functioning of FRS by modifying the security descriptor and time stamp of files in the SYSVOL folder. This modification causes FRS to replicate the changes to other domain controllers, which creates a high volume of file replication traffic.

Some antivirus programs modify files and directories in a way that causes FRS replication. Because such software often scans every file and directory in an FRS replicated tree, this interaction amounts to requesting a full synchronization of all files and folders from every domain controller that is running the software. The following symptoms might indicate this problem in the network:

- Files in SYSVOL shares are replicated excessively.
- Network traffic between replication partners consumes excessive bandwidth.

- The number of files in the staging directory constantly grows and then empties sometime after the antivirus scan completes or after FRS replication.

| Note |
| --- |
| Staging directory behavior changes for Windows Server 2003 and for Windows 2000 Server SP3 and later. These operating systems contain an updated version of FRS that leaves staging files in place for one week before removing them. |

You can maintain high security on domain controllers while preventing antivirus software from causing excessive replication by performing the following tasks:

Configure antivirus software to not scan files in the SYSVOL directory tree.

Require script signing on the domain controller.

For more information about running antivirus software on domain controllers, see article 815263, "Antivirus, Backup, and Disk Optimization Programs That Are Compatible with the File Replication Service," in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkId=4441.

**Excluding Antivirus Scanning of SYSVOL**

Antivirus software can interfere with the normal functioning of FRS by modifying the security descriptor and time stamp of files in the SYSVOL folder. These modifications trigger FRS to replicate the changes in the SYSVOL folder to other domain controllers. This interference results in excessive FRS replication between domain controllers. To run antivirus software on domain controllers, configure the antivirus software to exclude the SYSVOL folder and its subtree from scanning.

| Note |
| --- |
| The following antivirus applications do not modify SYSVOL files in a way that triggers excessive replication: |

- eTRUST Antivirus build 96 or later with the NTFS incremental scan feature disabled.
- McAfee/NAI NetShield 4.50 with the NetShield Hot Fix Rollup.
- Norton Antivirus 7.6 or later.

**Requiring Script Signing on Domain Controllers and Administrative Workstations**

In contrast to the directory database and log files, excluding the SYSVOL folder from virus scanning does increase the risk of a virus attack on a domain controller. The reason for the risk is that viruses tend to attach to binaries or scripts. Excluding the SYSVOL folder from virus scans increases the risk of virus attacks on logon scripts and startup scripts in the SYSVOL folder on domain controllers. As a countermeasure, implement script signing to protect the integrity of scripts running on domain controllers and administrative workstations.

| Note |
| --- |
| If you are using one of the antivirus applications listed in the preceding note, you do not need to exclude the SYSVOL folder from being checked by the antivirus application. Therefore, you do not need to implement script signing. |

Windows 2000 Server, Windows XP, and Windows Server 2003 support enforcement of script signing to validate the integrity of scripts and binaries before they are executed. Signing a script guarantees who wrote the script and that it has not been tampered with in any way since it was signed. Establish a practice for your organization that requires that all scripts in the SYSVOL folder be signed. Windows Script Host scripts (.vbs, .js, .wsf, and so on) can be signed or verified by using the same tools that are ordinarily used to sign other executables.

Sign scripts by performing the following tasks:

| Note |
| --- |

At a minimum, you should enforce script signing on domain controllers and administrative workstations. However, it is a best practice recommendation that script signing be enforced on all computers on the network that are running operating systems that support script signing. Such operating systems include the Windows 2000 family of operating systems, Microsoft® Windows® XP, and the Windows Server 2003 family.

- Enroll to acquire a certificate from an internal certification authority (CA).

  Choose an internal CA so that confirming the current validity of the certificate does not require an Internet lookup. To ensure that the certificate has not been revoked, the code checks the CA's certificate revocation list (CRL).

- Secure all scripts in the SYSVOL folder by signing.

**Note**

Batch files cannot be signed; therefore, they are inherently less secure. Convert all batch files to .vbs scripts as soon as possible.

Two alternatives are available for creating signed scripts:

- The Windows Product SDK contains a set of tools for signing scripts (signcode.exe and chktrust.exe).

- Alternatively, for anyone interested in developing their own tool, Windows Script Host version 5.6 ships with a signer object Scripting.Signer to sign or verify scripts.

For the code that is required to implement one of these signing solutions, see "Securing Scripts with Script Signing" in the Appendix: Procedures section.

- Using software restriction policies in Group Policy, you can apply a restriction so that client computers cannot run scripts at all. However, because the goal is that clients be able to run signed scripts, you can further create a certificate rule that allows clients to run scripts if they can verify the signature by using the public key in the certificate.

  This step can be performed only if the client computers are running Windows XP or Windows Server 2003. To require that only signed scripts be run on computers that are running Windows 2000 Server or Windows 2000 Professional, add the following registry entry in **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows Script Host\Settings** on every computer that is running Windows 2000 Server or Windows 2000 Professional:

  Entry name: Trust Policy

  Data type: REG_WORD

  Value: 2

  **Caution**

  Incorrectly editing the registry can have serious, unexpected consequences that might require you to restore the registry or to reinstall Windows. Before you edit the registry, create an emergency startup disk and back up any critical data. For more information, see the Registry Reference for Microsoft Windows Server 2003 at http://go.microsoft.com/fwlink/?LinkId=4543.

For more information about implementing script signing, see:

- "Providing a Secure Experience" in the Scripting Clinic on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=20258.

- "Digital Code Signing Step-by-Step Guide" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?

LinkId=18548.

- "Windows Script Host: New Code-Signing Features Protect Against Malicious Scripts" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=18550.

# Maintaining Physical Security

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

All of the previous domain controller security measures assume that your domain controllers are physically secure. Physical security ensures that unauthorized users cannot turn domain controllers on or off, add or remove hardware, insert or remove removable media, log on by using the domain controllers' keyboards and displays, or remove backup media.

To maintain physical security for your domain controllers:

- Secure domain controllers against physical access.
- Prevent domain controllers from booting into alternate operating systems.
- Protect domain controllers on restart by using SYSKEY.
- Secure backup media against physical access.
- Enhance the security of the network infrastructure.
- Secure the remote restart of domain controllers.

**Securing Domain Controllers Against Physical Access**

The first line of defense in maintaining physical security is to secure domain controllers against any attacks that can be accomplished with physical access to a domain controller. Note that changes in a domain controller's environmental conditions, such as power failures, can also compromise the security of the domain controller.

Take the following common security precautions for restricting physical access to your domain controllers:

- Use UPSs to prevent loss of power.
- Place domain controllers and UPSs in a locked room.
- Require cardkey locks or cipher-locks on the entrances to the locked room.
- Require locks on individual domain controllers or on doors to the racks that house the domain controllers.
- Require specific processes and procedures for any administration or repair of the domain controllers.

These security precautions are intended to prevent an attacker from gaining physical access to your domain controllers. However, in an environment where these recommendations cannot be strictly enforced, such as in branch offices, additional security measures might be required, as described in the following sections.

**Preventing Domain Controllers from Booting into Alternate Operating Systems**

A domain controller can be booted into an alternate operating system. For example, public domain drivers exist for MSDOS that an attacker can use to boot the domain controller and directly access files that are stored on NTFS disk volumes, bypassing existing NTFS permissions. Similar utilities exist for Linux and UNIX operating systems. You can take steps to avoid this type of attack.

To minimize the possibility of domain controllers booting into an alternate operating system:

- Disable or remove the floppy disk drive, unless it is required by SYSKEY. For more information, see the Protecting Domain Controllers on Restart by Using SYSKEY section.
- Disable or remove the CD-ROM or DVD drive.

- Set the [timeout] parameter in the Boot.ini file to 0.
- Disable remote network boot and installation, for example, by RIS or Bootstrap Protocol (BOOTP).
- When you are unable to use SYSKEY with a password or floppy disk, require a basic input/output system (BIOS) password to boot the computer.

**Protecting Domain Controllers on Restart by Using SYSKEY**

Generally, in secure datacenter environments, only authorized personnel can restart domain controllers. However, in an environment where these recommendations cannot be strictly enforced, such as in branch offices, there is increased potential for an unauthorized person to restart a domain controller.

An unplanned or unexpected restart of a domain controller could indicate that an attacker has booted a domain controller with an alternate operating system and compromised the security of the domain controller. On the other hand, the restart might simply be due to a loss of power or to scheduled maintenance on the domain controller.

**Evaluating the Need for SYSKEY**

The system key (SYSKEY) in Windows Server 2003 protects security information (including password information) in the Active Directory database and other Local Security Authority (LSA) secrets against offline attacks by encrypting their storage on a domain controller. SYSKEY can either be derived from a secret password that you specify, or it can be stored on offline media, such as a floppy disk. On a domain controller reboot, either the password or the floppy disk containing SYSKEY must be supplied to successfully restart the computer.

Implementing SYSKEY provides two security advantages:

- Point-in-time control of the domain controller restart, which evaluates the reason for the domain controller restart and determines if security has been compromised.
- Protection for passwords that are stored in the directory database against offline attacks if the domain controller or a disk are stolen.

You can use the system key utility (Syskey.exe), which is installed on the domain controller during Windows Server 2003 installation, to select one of two configurations for the SYSKEY source.

The use of SYSKEY involves the following logistic operational issues:

- Management of SYSKEY passwords or floppy disks can present challenges.

  Requiring a branch manager or local administrative staff to come to the office at 3 A.M. to enter passwords or insert a floppy disk might be difficult.

  Alternatively, allowing centralized IT operations personnel to provide the SYSKEY password remotely requires additional hardware, such as Compaq Remote Insight Lights-out (RILO) or Dell Remote Access Card (DRAC III) boards. For more information about restarting domain controllers remotely, see the Securing the Remote Restart of Domain Controllers section.

- Loss of the SYSKEY password or floppy disk leaves the domain controller in a state in which it cannot be restarted.

  There is no way to recover a domain controller if the SYSKEY password or floppy disk is lost. In this situation, it is necessary to rebuild the domain controller.

**Selecting a Method for Securing Domain Controller Restarts with SYSKEY**

Each method noted in the previous section (namely, manually entering SYSKEY passwords or supplying SYSKEY on a floppy disk) has advantages and difficulties. If you choose to add SYSKEY protection to your domain controllers, you should first evaluate your security environment to determine which method will work best for you.

**Providing SYSKEY Passwords to Secure Domain Controller Restarts**

SYSKEY passwords do not require physical media that could be lost because there are no floppy disks to lose. Trusted personnel must enter a password in the event that the domain controller needs to be restarted. The password should be known to only a small group of trusted administrators, preferably only members of the Domain Admins group. The

disadvantage of using passwords to secure SYSKEY is that trusted personnel are required to memorize passwords and be on-site to enter the passwords.

To support branch offices, you might need to provide SYSKEY passwords remotely through central IT trusted personnel. However, this method requires additional hardware, such as Compaq RILO or Dell DRAC III boards. For more information about restarting domain controllers remotely, see the Securing the Remote Restart of Domain Controllers section.

Because passwords can be compromised, you might be able to increase the security of passwords that are used for SYSKEY restarts by:

- Using strong passwords.
- Storing the passwords in a secure environment, such as a bank safety deposit box.
- Requiring the periodic changing of the SYSKEY passwords.

**Providing SYSKEY Floppy Disks to Secure Domain Controller Restarts**

Using SYSKEY with a password that is stored on a floppy disk does not require that a password be memorized by trusted personnel. However, implementing SYSKEY with a floppy disk does introduce the risk of a lost or damaged floppy disk. Furthermore, trusted personnel are required to insert the floppy disk during domain controller restart. Again, only trusted personnel, preferably members of the Domain Admins group, should have access to the SYSKEY floppy disk.

To support branch offices, you might need to install other hardware devices, such as Compaq RILO or Dell DRAC III boards, so that images of floppy disks can be remotely transferred to the domain controller. Using these devices, central IT trusted personnel can transfer a copy of the SYSKEY disk image to a remote domain controller. After the domain controller is restarted, IT operations personnel can delete the remote image of the SYSKEY floppy disk.

Because the floppy disk contains the cryptographic key for SYSKEY, you should take measures to ensure that the floppy disk is not stolen, lost, destroyed, or copied by an unauthorized person. You can mitigate these risks by doing the following:

- Copy the floppy disk and storing the copy off-site, such as in a bank safe deposit box.
- Store the working copy of the floppy disk in a secure place on-site.
- Remove the floppy disk from the domain controller immediately after it restarts.

**Securing Backup Media Against Physical Access**

As part of your normal operational practices, you should regularly back up your domain controllers and secure the backup media to minimize the risk of data tampering or theft.

Because the backup contains all the information in the Active Directory database, theft of the backup media presents the same risks as theft of a domain controller or theft of a disk drive from the domain controller. An attacker could restore the information elsewhere and illegally access the Active Directory data.

You can help prevent unauthorized users from gaining physical access to backup media by doing the following:

- Store backup media that is used on-site in a secure location where access is audited.
- Store archival backup media securely off-site.
- Establish processes and procedures that require the signatures of authorized administrators when any archival backup media is brought back on-site.
- Ensure that backup media is only in the backup device during the backup or recovery process.

**Enhancing the Security of the Network Infrastructure**

The placement of domain controllers in your network environment directly affects the security of your domain controllers. The primary focus in network security is to isolate the domain controllers from unauthorized users, while providing high-speed, secure access to authorized users. To ensure that domain controllers are properly isolated, secure any cabling rooms and place domain controllers on secured segments in your network.

**Securing Cabling Rooms**
If an attacker can gain access to your cabling rooms, the attacker can use a protocol analyzer to capture network traffic and compromise your network's security, including domain controller security. In intranet and perimeter network datacenters, the possibility of unauthorized personnel gaining access to these cabling rooms is negligible. However, in branch offices, the cabling rooms might be shared with telephony wiring and other utilities.

Typically, your cabling rooms are where some of your routers and switches are located. The routers and switches contain a logical diagram of your network, because they manage and maintain routing information. This routing information is used by Active Directory to determine Internet Protocol (IP) subnets, and it determines the preferred domain controller for computers in the network.

Attackers can gain access to your switches and routers by using Telnet or Web interfaces that are provided by these devices. If attackers gain access to the configuration information of these devices, they can use this information to mount attacks on your domain controllers.

In most environments, the routers and switches all use the same password for reading and configuring information. The management software for routers and switches might support the use of only a single name, or you might want to use a single password to simplify the management of a large number of switches. For example, organizations can use the same Simple Network Management Protocol (SNMP) community name for all these devices to make management easier.

Regardless of the environment, you can improve the security of your cabling rooms by doing the following:

- Require cardkey locks or cipher-locks on the entrances to the cabling rooms.
- Require locks on the racks that house wiring panels, switches, or routers.
- Include UPSs to prevent loss of power.
- Require specific processes and procedures for any administration or repair of cabling, switches, or routers.
- Use strong passwords to secure the configuration of routers and switches.
- Use different passwords for reading and configuring routers and switches.

**Placing Domain Controllers in Secure Network Segments**
The manner in which you place domain controllers in your network can optimize isolation of domain controllers from unauthorized users. At the same time, ensure that users and computers have high-speed connectivity to their respective domain controllers.

**Placing Domain Controllers in Datacenters**
Place domain controllers in the datacenter so that only designated personnel have direct physical contact with the domain controllers. Place your forest root, application, and regional domain controllers for use within your organization's private network in the datacenter. The placement of domain controllers in your perimeter network is discussed in a later section.

Place domain controllers so that firewalls protect domain controllers from Internet users and so that routers, and possibly firewalls, protect domain controllers from users within the organization's private network. Figure 7 illustrates the placement of domain controllers in a datacenter.




**Placing Domain Controllers in Perimeter Networks**

Place domain controllers in your perimeter networks so that the domain controllers are not directly accessible by Internet users. Ensure that only Web servers, database servers, file servers, users within the private network, and computers within the private network have direct access to the domain controllers.

Placing the domain controller behind a stand-alone router helps prevent Internet users from directly accessing the domain controller. To minimize security risks, place the domain controller on the same network segment as the client computers. In addition, use VPN tunnels to ensure that the router communicates exclusively with the organization's hub or central location. Prevent any other requests that originate from the Internet from reaching the network segment of the branch office and, subsequently, the domain controller.

Figure 8 shows a perimeter network in which domain controllers are positioned outside the firewall that protects the internal network.




**Placing Domain Controllers in Branch Offices**

In branch offices, domain controllers often provide other services, such as file or print services, to users in the branch office. These multifunction domain controllers communicate with the rest of the organization through a routed connection, possibly a dial-up modem that is managed by a stand-alone router. The stand-alone router provides the isolation and firewall functionality to protect the domain controller and the users in the branch office.

Place the domain controller behind the stand-alone router to prevent Internet users from directly accessing the domain controller. Place the domain controller on the same network segment as the client computers. Ensure that the router communicates exclusively with the organization's hub or central location by using VPN tunnels. Prevent any other requests that originate from the Internet from reaching the branch office's network segment and, subsequently, the domain controller.

| **Note** |
| --- |
| Whenever possible, use dedicated domain controllers in branch offices to minimize the number of personnel who need logon access to manage other services or applications that might otherwise run on the domain controllers. |

Figure 9 illustrates the placement of a domain controller in a branch office.




In addition, a modem can provide out-of-band management capability for a remote domain controller. The modem directly connects to a COM port on the domain controller, to remote management hardware (such as the Compaq RILO board), or to an intelligent UPS. This out-of-band management capability can provide the ability to perform BIOS configuration, boot process monitoring and selection, and turn the domain controller on and off.

Ensure that the number to the modem is kept secret. Require the modem to call back to a predetermined list of numbers, and require user identification for callback.

**Securing the Remote Restart of Domain Controllers**

In some situations, your domain controllers might require remotely administered management, or they might be placed in branch offices outside your organization's datacenter. In these situations, the restart of a domain controller must be performed remotely. You can use a Terminal Services client on Windows 2000 and earlier computers. On Windows XP and later desktops, Remote Desktop Connection is the Terminal Services client.

Tasks that must be performed during the remote restart of a domain controller include the following:

- Selection of the Windows Server 2003 boot options
- Configuration of the BIOS on the domain controller

Terminal Services is unable to perform these tasks because Windows Server 2003 must be running to support Terminal Services. These tasks require additional hardware to support remote restart functionality. Examples of this type of hardware include the following:

- Smart UPSs
- Remote access hardware that is integrated into the server, such as Compaq RILO or Dell DRAC III boards
- Video switches that connect to the keyboard, mouse, and display that provide services that are similar to Terminal Services

Most of these devices communicate by using RS-232 or Ethernet, and they have only rudimentary security, such as a password. In datacenters, the devices that communicate through RS232 are connected to a terminal concentrator. The terminal concentrator multiplexes a number of RS-232 connections into a single RS-232 or Ethernet connection. Smart UPS and remote access hardware typically communicate through Telnet.

Secure the remote restart of domain controllers by doing the following:

- When the domain controller is in a datacenter, connect the remote restart device's RS232 or Ethernet connection to a network segment that is dedicated to network management and that is isolated from clients.

When the domain controller is in a branch office, connect the remote restart device to a dedicated modem, and require the modem to provide password identification and callback functionality.

# Recommendations: Deploying Secure Domain Controllers

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Following the security recommendations described earlier in this section helps minimize the security risks that may be involved in deploying domain controllers. Of course, as previously mentioned, you should consider the recommendations that are described in other sections when deciding how best to enhance your comprehensive Active Directory security.

In most instances, these recommendations are intended for intranet datacenter, extranet datacenter, and branch office scenarios. However, some of the recommendations depend on the particular scenario. Where the recommendations are scenario specific, notes are included to direct you to the section where the recommendation is discussed.

**Recommendations for Securing the Domain Controller Build Environment**
Recommendations for establishing secure domain controller rollout practices are:

- Whenever possible, build domain controllers in secured environments, such as datacenters.
- When building domain controllers in unsecured environments, ensure that only trusted personnel have physical access.

**Recommendations for Establishing Secure Domain Controller Build Practices**
The following sections provide checklists of recommendations for establishing secure domain controller build practices.

**Recommendations for Using Automated Installation Processes**
Recommendations for using automated installation processes are:

- Use an imaged-based, automated deployment process for installing Windows Server 2003.

| **Note** |
| --- |
| This recommendation may vary or not be feasible, depending on your scenario. For more information, see "Establishing Secure Domain Controller Build Practices" earlier in this guide. |

- Automate the promotion of servers to domain controllers.

**Recommendations for Ensuring Predictable, Repeatable, and Secure Domain Controller Deployments**
Recommendations for ensuring that your domain controller deployments are predictable, repeatable, and secure are:

- Installing Windows Server 2003 with Service Packs and Hotfixes
    - Install Windows Server 2003 with the most recent service packs.
    - Apply all current security-related hotfixes.
    - Format all partitions as NTFS.
    - Create a strong password for the Administrator account.
    - Do not install IIS on domain controllers.
    - Select DNS during installation.
    - Do not install SMTP unless Active Directory replication uses SMTP.
- Disabling NTFS Automatic 8.3 Name Generation

    - Disable NTFS automatic 8.3 name generation.

- Running Virus-Scanning Software on the Server

    - Run virus-scanning software before promoting any server to a domain controller.

    - Ensure that virus-scanning software includes any updates to detect and remove the latest viruses.

    - Configure virus scans appropriately to prevent excessive FRS replication and to prevent interference between virus scans and access to database files and log files by FRS and Active Directory.

    - Enable script signing if you have stopped the virus-scanning software from scanning the SYSVOL folder.

- Enabling Only Essential Services

    - Enable only those services that are required for a computer that is running Windows Server 2003 in the role of a domain controller.

- Creating a Reserve File to Enable Recovery from Disk-Space Attacks

    - Create a reserve file on the same disk volume as Ntds.dit. Ensure that the reserve file is either 250 MB or 1 percent of the available disk space, whichever is larger.

**Recommendations for Configuring the Automatic Installation of Active Directory**
Recommendations for configuring automatic Active Directory installation are:

- Using Unattended Active Directory Installation

    - Set the registry to run Dcpromo.exe automatically.

    - Prepare the DCInstall portion of the Unattend.txt file.

- Selecting Secure Active Directory Configuration Settings

- Place the Active Directory database (Ntds.dit) on a separate physical drive.

    - Place the Active Directory logs on a separate physical drive.

    - Place SYSVOL on the same physical drive as the Active Directory database.

    - Do not enable Pre–Windows 2000 Compatibility unless it is required by earlier applications or services, as described in “Determining the Need for Anonymous Access to Active Directory Data” earlier in this guide.

- Determining the Need for Anonymous Access to Active Directory Data

    - Create a test domain that mirrors your production domain and that allows anonymous access.

    - Monitor the directory for anonymous access according to “Enabling Monitoring for Anonymous Active Directory Access” in “Appendix: Procedures” later in this guide.

    - Identify the applications and services that use anonymous access.

- Eliminating the Requirement for Anonymous Access to Active Directory

    - Upgrade the applications that are identified as using anonymous access.

- Monitor the directory to ensure that no further anonymous access is requested.
- Disable anonymous access, if it is in effect.

**Recommendations for Maintaining Physical Security**

Recommendations for maintaining the physical security of your domain controllers are:

- Securing Domain Controllers Against Physical Access
  - Include UPSs.
  - Place domain controllers and UPSs in locked rooms.
  - Require cardkey locks or cipher-locks on the entrances to the locked rooms.
  - Require locks on individual domain controllers or on doors to the racks that house domain controllers.
  - Require specific processes and procedures for the administration or repair of domain controllers.
- Preventing Domain Controllers from Booting into Alternate Operating Systems
  - Disable or remove the floppy disk drive, unless it is required for SYSKEY.
  - Disable or remove the CD-ROM or DVD drive.
  - Set the [timeout] parameter in boot.ini file to 0.
- Protecting Domain Controllers on Restart by Using SYSKEY
  - Enable SYSKEY.

    | **Note** |
    | --- |
    | This recommendation may vary or not be feasible, depending on your scenario. For more information, see "Protecting Domain Controllers on Restart by Using SYSKEY" earlier in this guide. |

- Securing Backup Media Against Physical Access
  - Store backup media that is used on-site in a locked cabinet or container.
  - Store archival backup media in off-site storage.
  - Establish processes and procedures that require signatures to bring any archival storage back on-site.
  - Ensure that backup media is installed only during backup and that it is in secured storage otherwise.
- Enhancing the Security of the Network Infrastructure
  - Require cardkey locks or cipher-locks on the entrances to the cabling rooms.
  - Require processes and procedures for any administration or repair of cabling, switches, or routers.
  - Use strong passwords to secure the configuration of routers and switches.
  - Use different passwords for reading and configuring your routers and switches.

- Securing the Remote Restart of Domain Controllers
    - When the domain controller is in a datacenter, connect the remote restart devices to the secured management network.

When the domain controller is in a branch office, connect the remote restart device to a dedicated modem, and require the modem to provide password identification and callback functionality.

# Chapter 4: Strengthening Domain and Domain Controller Policy Settings

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Active Directory in Windows Server 2003 contains default security policy settings for the domain and for domain controllers in the form of Group Policy objects (GPOs). These settings are configured by default when Active Directory is installed, and they are applied to all new domains.

| **Important** |
| --- |
| Default Group Policy settings do not overwrite existing Windows 2000 settings. Existing Windows 2000 Group Policy settings are maintained when you: |

- Install Active Directory on a server that is running Windows Server 2003 in an existing Windows 2000 domain.
- Upgrade a domain controller that is running Windows 2000 to Windows Server 2003.

**Default Domain and Domain Controller GPOs**

In a new Windows Server 2003 domain, the following default GPOs protect the domain and all domain controllers:

- Default Domain Policy, which is linked to the domain object and affects all users and computers in the domain (including computers that are domain controllers) through policy inheritance.
- Default Domain Controllers Policy, which is linked to the Domain Controllers OU. This policy generally affects only domain controllers, because by default, computer accounts for domain controllers are kept in the Domain Controllers OU.

You can manage security-specific policies in the Default Domain Policy GPO and in the Default Domain Controllers Policy GPO by using the administrative tools Domain Security Policy and Domain Controller Security Policy, respectively. Security policy settings are organized into the following categories:

- Account Policies, which include:
    - Password Policy: Controls password enforcement and lifetimes on domain accounts.
    - Account Lockout Policy: Determines the circumstances and length of time that an account will be locked out of the system.
    - Kerberos Policy: Determines Kerberos-related settings, such as ticket lifetimes and enforcement.
- Local Policies, which include:
    - Audit Policy: Tracks system security events on computers.
    - User Rights Assignment: Controls user and administrative actions on computers.
    - Security Options: Affects Active Directory, network, file system, and user logon abilities.
- Event Log Policy: Defines attributes that are related to the application, security, and system event logs

**Strengthening Security Policy Settings**

In most cases, default Windows Server 2003 settings are effective in securing domains and domain controllers against various types of threats. However, some default settings can be strengthened to improve the level of protection. In addition to security policies, Active Directory data is protected by default auditing settings on key directory objects. For most settings, changes to default policy are best made to the default GPOs. In cases where new GPOs can be generated without affecting the operation of the default GPOs, this process is recommended and described. For more information about applying security policy settings, see Applying Selected Domain and Domain Controller Policy Settings later in this guide.

Default and recommended changes for strengthening security policy and directory object auditing are presented in the following categories:

- Domain Security Policy settings
- Domain Controller Security Policy settings

In the following sections, default and recommended settings are presented for key settings in each category, even if no changes to default settings are recommended. Security policy settings that require no change are also presented as a means of describing the key security policy settings that are in place by default on Windows Server 2003 domains and domain controllers.

**Applying Security Policy Settings**
Where a security policy setting is applied depends on if it is defined in the Default Domain Policy GPO and inherited from the domain level to the OU level or if it is defined in the Default Domain Controller Policy GPO. Policy settings that are applied at the OU level override policy settings at the domain level. In this way, security policy settings that are specific to domain controllers, but not to all users, groups, and computers in the domain, can be set at the Domain Controllers OU level. For this reason, certain policy settings are discussed in this chapter in relation to the domain and other policy settings are discussed in relation to domain controllers.

When you apply security option settings to affect the Domain Controllers OU, some settings are best applied by adding them to a new GPO, while other settings require application directly to the Default Domain Controllers Policy GPO. Requirements and recommendations for these settings are provided in Applying Selected Domain and Domain Controller Policy Settings later in this guide.

**Auditing Important Active Directory Objects**
In addition to Domain Controller Security auditing settings (Audit Policy), auditing settings are essential for tracking access to Active Directory objects themselves. For this reason, auditing is enabled by default on the topmost object in each directory partition in Active Directory — the domain, schema, and configuration directory partitions — as well as on certain key configuration objects.

Most auditing settings that are recommended for Windows 2000 Server deployments are set by default on important Active Directory objects in Windows Server 2003. These settings are described in Reviewing Audit Settings on Important Active Directory Objects later in this guide.

# Strengthening Domain Policy Settings

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Domain security policy settings provide Active Directory with domain-wide security options for handling authentication and authorization of Active Directory security principals. These policy settings are implemented as security settings within the Default Domain Policy GPO. Domain policy is applied to all security principal accounts in the domain, unless inheritance is specifically blocked or overridden by another policy.

Security Policy settings are applied at the domain level by default for the following categories:

- Account Policies, which include:
    - Password Policy
    - Account Lockout Policy
    - Kerberos Policy

**Modifying Domain Security Policy**

You can make changes to Group Policy by modifying the default GPO or by creating a new GPO. The recommendation for making changes to domain security policy is to always modify the default GPO. The primary reason for this recommendation is that APIs that were developed for earlier versions of the operating system update policy settings in the Default Domain Policy GPO. For this reason, make all changes to domain security policy settings by editing this GPO.

To increase comprehensive security for your domain, apply the Password Policy, Account Lockout Policy, and Kerberos Policy settings that are recommended in this guide.

**Password Policy**

In Windows Server 2003, the most common means of authenticating a user's identity is the use of secret user passwords. After a user has been identified and authenticated, the user can perform any tasks or access any resource for which the user's account is authorized. Strong passwords generally enhance security for Active Directory users. Using strong passwords helps avoid the threat of an unauthorized user guessing (cracking) a weak password and acquiring the credentials of the compromised user account. This benefit applies especially to administrative accounts, because an unauthorized user could obtain administrative credentials and thereby gain elevated privileges.

A complex password that changes regularly reduces the likelihood of a successful spoofing attack. Password Policy settings control the complexity and lifetime for passwords. Table 13 includes the default and recommended Password Policy settings for a domain.

**Table 13 Default and Recommended Password Policy Settings**

| Policy | Default | Recommended | Comments |
|---|---|---|---|
| Enforce password history | 24 passwords remembered | (No change) | Prevents users from reusing passwords. |
| Maximum password age | 42 days | (No change) | N/A |
| Minimum password age | 1 day | (No change) | Prevents users from cycling through their password history to reuse passwords. |
| Minimum password length | 7 characters | (No change) | Sets minimum password length. |
| Password must meet complexity requirements | Enabled | (No change) | For the definition of a complex password, see "Creating a Strong Administrator Password" in the Establishing Secure Domain Controller Build Practices section. |

| | | | |
|---|---|---|---|
| Store password using reversible encryption | Disabled | (No change) | N/A |

**Note**

If possible, use smart cards throughout your organization to ensure that the strongest possible passwords are used on user accounts. Using smart cards causes the system to automatically generate cryptographically strong random passwords for accounts. If you are unable to provide smart cards for all users, require service administrator accounts to use smart cards. For more information about smart cards, see Chapter 5: Establishing Secure Administrative Practices later in this guide.

**Account Lockout Policy**

More than a few unsuccessful password tries during the logon process can represent an attempt by an attacker to determine an account password by trial and error. Windows Server 2003 keeps track of logon attempts, and it can be configured to respond to this type of attack by disabling the account for a preset period of time. This response is referred to as account lockout.

Account Lockout Policy settings control the threshold for this response and the actions to be taken when the threshold is reached. Table 14 includes the default and recommended Account Lockout Policy settings.

**Table 14 Default and Recommended Account Lockout Policy Settings**

| Policy | Default | Recommended | Reason |
|---|---|---|---|
| Account lockout duration | Not defined | 0 minutes | The value 0 means that after account lockout an Administrator is required to reenable the account before account lockout reset has expired. |
| Account lockout threshold | 0 invalid logon attempts | 20 invalid logon attempts | The value 0 means that failed password tries never cause account lockout.<br><br>Because an account lockout duration of 0 minutes (administrator reset) is recommended, a small number for this setting can result in frequent administrator interventions. |
| Reset account lockout counter after | Not defined | 30 minutes | This setting protects against a sustained dictionary attack by imposing a nontrivial delay after 20 unsuccessful attempts. |

**Kerberos Policy**

In Windows Server 2003, Kerberos provides the default mechanism for authentication services, as well as the authorization data that is necessary for a user to access a resource and perform a task with that resource. If the lifetimes of Kerberos tickets are reduced, the risk of having a legitimate user's credentials stolen and used by an attacker diminishes. However, authorization overhead increases. Table 15 includes the default Kerberos policy settings. No changes to these default settings are recommended.

**Table 15 Default and Recommended Kerberos Policy Settings**

| Policy | Default | Recommended | Comments |
|---|---|---|---|
| Enforce user logon restrictions | Enabled | (No change) | N/A |
| Maximum lifetime for service ticket | 600 minutes | (No change) | N/A |
| Maximum lifetime for user ticket | 10 hours | (No change) | N/A |
| Maximum lifetime for user ticket renewal | 7 days | (No change) | N/A |
| Maximum tolerance for computer clock synchronization | 5 minutes | (No change) | Maximum tolerance between the client's and server's clocks. |

# Strengthening Domain Controller Policy Settings

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

In addition to Group Policy settings for domains, Windows Server 2003 Default Domain Controller Group Policy settings also protect domain controllers and Active Directory objects themselves. Domain Controller Security Policy settings apply to the Domain Controllers OU in each domain.

**Domain Controller Security Policy Settings**
Security Policy settings are applied at the Domain Controllers OU level by default for the following categories:

- Local Policies, which include:
    - Audit Policy
    - User Rights Assignment
    - Security Options
- Event Log Policy

To increase security for your domain controllers, apply the User Rights Assignment, Security Options, and Event Log settings that are recommended in the following sections. Audit Policy settings do not require change, but they are presented for your information.

**Changes to Domain Controller Security Policy**
APIs that were developed for earlier versions of the operating system update some security policy settings in the Default Domain Controller Policy GPO, but not others. For this reason, changes to some domain controller security policy settings must be made by editing the default GPO, but others are best implemented by creating a new GPO. For more information about changing the default domain controller security policies as opposed to creating new GPOs for domain controller security policies, see the Applying Selected Domain and Domain Controller Policy Settings section.

**Audit Policy**
On domain controllers that are running Windows Server 2003, auditing is turned on by default to log the success of key security events. Default auditing on domain controllers represents a change from Windows 2000 Server, which does not enable auditing by default. Although no changes are recommended in the default settings, the settings are presented here because they represent significant changes from the Windows 2000 Server default settings.

| **Important** |
| --- |
| There are many possible goals that you can have when you audit a domain for security purposes, such as intrusion detection or forensic analysis of security breaches. The primary goal of the security audit settings is to provide accountability for sensitive directory operations, including any administrative or configuration changes. When auditing for other reasons, such as intrusion detection, additional audit settings might need to be enabled. |

When auditing is enabled on domain controllers, events are recorded in the Security event log. For the default and recommended settings for the maximum size of the Security event log, see the Strengthening Domain Controller Event Log Policy Settings section.

| **Note** |
| --- |
| If you make changes to Audit Policy security policy settings, make all changes by editing the Default Domain Controllers Policy GPO. Security policy settings for this GPO are available in Domain Controller Security Policy in Administrative Tools. |

Table 16 lists the default and recommended settings for domain controller Audit Policy.

**Table 16 Default and Recommended Domain Controller Audit Policy Settings**

| Policy | Default Setting | Recommended Setting | Comments |
|---|---|---|---|
| Audit account logon events | Success | (No change) | Account logon events are generated when a domain user account is authenticated on a domain controller. |
| Audit account management | Success | (No change) | Account management events are generated when security principal accounts are created, modified, or deleted. |
| Audit directory service access | Success | (No change) | Directory services access events are generated when an Active Directory object with a system access control list (SACL) is accessed. |
| Audit logon events | Success | (No change) | Logon events are generated when a domain user interactively logs on to a domain controller or a network logon to a domain controller is performed to retrieve logon scripts and policies. |
| Audit object access | No auditing | (No change) | N/A |
| Audit policy change | Success | (No change) | Policy change events are generated for changes to user rights assignment policies, audit policies, or trust policies. |
| Audit privilege use | No auditing | (No change) | N/A |
| Audit process tracking | No auditing | (No change) | N/A |
| Audit system events | Success | (No change) | System events are generated when a user restarts or shuts down the domain controller or when an event occurs that affects either the system security or the security log. |

**User Rights Assignment**

User rights allow users to log on and perform specific administrative or operations tasks on domain controllers. Ensure that the appropriate user rights are assigned to users in the domain so that users can perform their intended functions without compromising the security of the domain controllers. Establish the policy settings for domain controller user rights assignment to properly limit the users who can log on to the domain controllers and perform the necessary administrative tasks.

Table 17 lists the default and recommended policy settings for domain controller user rights assignment policies. Default Windows Server 2003 settings for all other user rights assignment policies are consistent with security recommendations.

| **Note** |
|---|
| If you make changes to user rights assignment security policy settings, make all changes by editing the Default Domain Controllers Policy GPO. The security policy settings for this GPO are available in Domain Controller Security Policy in Administrative Tools. |

**Table 17 Default and Recommended Domain Controller User Rights Assignment Policy Settings**

| Policy | Default Setting | Recommended Setting | Comments |
|---|---|---|---|
| Allow log on locally | Account Operators<br><br>Administrators<br><br>Backup Operators<br><br>Print Operators<br><br>Server Operators | Administrators<br><br>Backup Operators<br><br>Server Operators | Account Operators and Print Operators have few (if any) reasons to log on locally to a domain controller. |

| | | | |
|---|---|---|---|
| Shut down the system | Account Operators<br><br>Administrators<br><br>Backup Operators<br><br>Print Operators<br><br>Server Operators | Administrators<br><br>Backup Operators<br><br>Server Operators | Account Operators and Print Operators have few (if any) reasons to shut down domain controllers. |

**Note**

Members of the Backup Operators group can log on locally to domain controllers, archive files to backup media, and overwrite system files through restore operations. The members of this group should be limited to those users who perform domain controller backup and restore operations. To reduce the number of users that have these rights, do not grant Backup Operator group membership to users who are responsible only for application backup and restore operations, such as Microsoft SQL Server operators.

**Security Options**

Domain controller Security Options policy settings affect the security-related Windows Server 2003 configuration settings. The domain controller Security Options policy settings affect not only the security configuration settings that are related to Active Directory, but other components in Windows Server 2003 as well, such as the network, file system, and user logon security configuration settings.

**Note**

To implement changes to default Security Options policy, it is recommended that you create a new GPO. This GPO can be added and linked to the Domain Controllers OU above the level of the Default Domain Controllers GPO. In this way, the nondefault settings take precedence over the default settings, and you can also easily revert to default settings by simply deleting this GPO or placing it below the default GPO in the list of linked GPOs. For more information about creating this new GPO, see the Applying Selected Domain and Domain Controller Policy Settings section.

Table 18 lists the default and recommended policy settings for domain controller Security Options. Default settings for all other security options are consistent with security recommendations.

**Table 18 Default and Recommended Domain Controller Security Options Policy Settings**

| Policy | Default Setting | Recommended Setting | Comments |
|---|---|---|---|
| Audit: Audit the access of global system objects | Not defined | Disabled | Disables the creation of a default SACL on system objects, such as mutexes (mutually exclusive), events, semaphores, and MSDOS devices because the default setting is “No auditing.” |
| Audit: Audit the use of Backup and Restore privilege | Not defined | Disabled | Disables auditing for the use of user privileges, including Backup and Restore, when the “Audit privilege use” policy is enabled because this policy is configured for “No auditing.” |
| Audit: Shut down system immediately if unable to log security audits | Not defined | Disabled | Stops the domain controller if a security audit cannot be logged. The auditing goals for domain controllers, described in Reviewing Domain Controller Audit Policy Settings, allow overwriting security audit events as required. |
| Devices: Allow undock without having to log on | Not defined | Disabled | Because a domain controller is most likely not a laptop, undocking should never take place. Therefore, the recommendation is to disable this setting. |
| Devices: Allowed to format and eject removable media | Not defined | Administrators | Allows only Administrators to eject removable NTFS media to protect against the theft of sensitive data. |

| | | | |
|---|---|---|---|
| Devices: Prevent users from installing printer drivers | Not defined | Enabled | Allows only Administrators and Server Operators to install a printer driver when adding a network printer to ensure that users cannot install a printer driver (add a network printer) and perform disk-space attacks by submitting large print jobs. |
| Devices: Restrict CD-ROM access to locally logged-on user only | Not defined | Enabled | Allows only the interactively logged-on service administrator to access removable CD-ROM media to ensure that when no one is logged on interactively, the CD-ROM cannot be accessed over the network. |
| Devices: Restrict floppy access to locally logged-on user only | Not defined | Enabled | Allows only interactively logged-on service administrators to access removable floppy media to ensure that the floppy disk drive cannot be accessed over the network when no one is logged on. |
| Devices: Unsigned driver installation behavior | Not defined | Do not allow installation | Prevents insecure or untrusted device drivers from being installed on domain controllers. |
| Domain controller: Allow server operators to schedule tasks | Not defined | Disabled | Restricts the individuals who can schedule tasks to Administrators because scheduling usually runs as an elevated service. |
| Domain controller: Refuse machine account password changes | Not defined | Disabled | It is more secure to have machine accounts regularly change their password (default: 30 days). Therefore this setting is disabled. |
| Domain member: Digitally encrypt or sign secure channel data (always) | Enabled<br><br>Not defined* | Enabled | Requires Windows NT 4.0 with Service Pack 4 (SP4)or later on all domain controllers in local domains and all trusted domains to ensure that all security fixes have been made. |
| Domain member: Disable machine account password changes | Not defined | Disabled | It is secure to have machine accounts regularly change their passwords. By default, the local security policy on the domain controller disables this setting. |
| Domain member: Maximum machine account password age | Not defined | 30 days | The default local policy value is used by the Default Domain Controller Policy so that it is uniformly applied to all domain controllers. |
| Domain member: Require strong (Windows 2000 or later) session key | Not defined | Enabled | Requires that a secure channel be established with 128-bit encryption to ensure that the key strength is not negotiated but always uses the most secure connection possible with the domain controller. |
| Interactive logon: Do not display last user name | Not defined | Enabled | Removes the name of the last user to successfully log off from the Log On to Windows dialog box to prevent attackers from discovering service account names on domain controllers. |
| Interactive logon: Do not require CTRL+ALT+DEL | Not defined | Disabled | Requires CTRL+ALT+DEL before users log on to ensure that users are communicating by means of a trusted path when entering their passwords. |
| Interactive logon: Number of previous logons to cache (in case domain controller is not available) | Not defined | 0 logons | The value 0 indicates that the domain controller does not cache previous logons and requires authentication at each logon. |
| Interactive logon: Prompt user to change password before expiration | Not defined | 14 days | Notifies users in advance (in days) that their password is about to expire so that the user has time to construct a password that is sufficiently strong. |

| | | | |
|---|---|---|---|
| Interactive logon: Require Domain Controller authentication to unlock workstation | Not defined | Enabled | When cached credentials are used to unlock the console, any changes to the account, such as user rights assignment, group membership changes, or disabling of the account, are not enforced. To ensure that any changes to the account are enforced immediately, require domain controller authentication of the account to unlock the console, instead of cached credentials. |
| Interactive logon: Require smart card | Not defined | (See comments) | It is recommended that you use smart cards for logging on to both domain controllers and administrative workstations. If you have a public key infrastructure (PKI) infrastructure set up to deploy smart cards, set this option to Enabled. |
| Interactive logon: Smart card removal behavior | Not defined | Force logoff | Forces service administrators to keep smart cards inserted while they are logged on interactively on domain controllers to ensure that domain controllers are not left unattended with an active logon. |
| Microsoft network client: Digitally sign communications (always) | Not defined | (See comments) | See SMB Signing on Domain Controllers for requirements. |
| Microsoft network client: Digitally sign communications (if server agrees) | Not defined | (See comments) | See SMB Signing on Domain Controllers for requirements. |
| Microsoft network client: Send unencrypted password to third-party SMB servers | Not defined | Disabled | Prohibits the SMB redirector from sending plaintext passwords to non-Microsoft SMB servers that do not support password encryption. Disable this policy unless your domain controller needs to communicate with non-Microsoft SMB servers. |
| Microsoft network server: Amount of idle time required before suspending session | Not defined | 15 min | Controls when a domain controller suspends an inactive server message block (SMB) session, which has no security implications but which reduces SMB traffic resource usage. |
| Microsoft network server: Digitally sign communications (always) | Enabled<br><br>Not defined* | (See comments) | See SMB Signing on Domain Controllers for requirements. |
| Microsoft network server: Digitally sign communications (if client agrees) | Enabled | (See comments) | See SMB Signing on Domain Controllers for requirements. |
| Microsoft network server: Disconnect clients when logon hours expire | Not defined | Enabled | Forcibly disconnects client sessions with the SMB Service when the user's logon hours expire to ensure that network connections are secured during nonworking hours. |
| Network access: Do not allow storage of credentials or Windows Live ID for network authentication | Not defined | Enabled | A usability feature that is typically not required on domain controllers. |

| | | | |
|---|---|---|---|
| Network access: Restrict anonymous access to Named Pipes and Shares | Not defined | Enabled | Restricts anonymous access to network shared folders and named pipes to those that are enumerated in the following settings:<br><br>Network access: Named pipes that can be accessed anonymously<br><br>Network access: Shares that can be accessed anonymously |
| Network security: Do not store LAN Manager hash value on next password change | Not defined | (See comments) | See Disabling LAN Manager Authentication for other requirements. |
| Network security: LAN Manager authentication level | Send NTLM response only<br><br>Not defined* | (See comments) | See Disabling LAN Manager Authentication for other requirements. |
| Network security: LDAP client signing requirements | Not defined | (See comments) | Set to "Require signing" only if you have domain controllers that are running Windows 2000 SP3 or Windows Server 2003. Otherwise, set to "Negotiate signing." |
| Recovery console: Allow automatic administrative logon | Not defined | Disabled | Requires that an Administrator account password be provided before access is granted to a domain controller to ensure that anyone logging on requires administrator credentials. |
| Recovery console: Allow floppy copy and access to all drives and all folders | Not defined | Disabled | Prevents unauthorized users from gaining access to, copying, and removing the Active Directory database and other secure files from the domain controller. |
| Shutdown: Allow system to be shut down without having to log on | Not defined | Disabled | Requires an authenticated, authorized service account to shut down or restart the domain controller. |
| Shutdown: Clear virtual memory pagefile | Not defined | Enabled | Eliminates process memory data from going into the page file on shutdown in case an unauthorized user manages to directly access the page file. |
| System objects; Strengthen default permissions of internal system objects (e.g. Symbolic Links) | Not defined | Enabled | Allows users who are not administrators to read shared objects but not modify them. Strengthens the default DACL of objects in the global list of shared resources, such as MSDOS device names, mutexes, and semaphores. |
| System settings: Optional subsystems | Not defined | (See comments) | By default, Posix is the only subsystem that is enabled. If you do not need Posix, you can define this policy and remove it from the list (so you that you have a blank list). |
| System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies | Not defined | (See comments) | If you have PKI set up, you can enable this setting to check (CRLs to make sure that the software certificate and signature are valid. |

* Default settings that are present when you upgrade a domain controller that is running Windows 2000 to Windows Server 2003 but that are not present when you perform a new installation of Windows Server 2003.

**Disabling LAN Manager Authentication**

By default, Windows operating systems earlier than Windows 2000 support only the LAN Manager (LM) authentication protocol. To provide compatibility with these earlier versions of Windows, Active Directory stores the account passwords in an LM hash format. Active Directory stores the password for the Windows NT authentication protocol (NTLM) and NTLM version 2 (NTLMv2) protocols in NTLM hash format. In the event that an attacker removes a domain controller or a domain controller

hard disk, it is easier for that attacker to decrypt the passwords in LM hash format. Because the NTLM hash is cryptographically stronger than the LM hash, disable the storage of passwords in LM hash format to provide a higher level of security.

When you use a SYSKEY password or floppy disk, you encrypt the entire Active Directory database and protect any passwords. When you use one of these SYSKEY methods, there is no benefit to disabling the storing of passwords in LM hash format, aside from reducing the size of the Active Directory database.

For more information about allowing only the NTLMv2 authentication protocol, see the following articles in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkId=4441:

- 285901, "Remote Access and VPN Clients Cannot Connect to a Server with NtlmcompatabilityLevel Set to 5"
- 281648, "Error Message: The Account Is Not Authorized to Login from This Station"
- 239869, "How to Enable NTLM 2 Authentication"

You can disable the storing of passwords in LM hash format by performing the following tasks:

1. Upgrade all domain controllers, member servers, and workstations to support the NTLMv2 authentication protocol.

   Table 19 lists the operating systems and the software requirements to support the NTLMv2 authentication protocol.

   **Table 19 Operating System and Software Requirements to Support NTLMv2**

| Operating System | Requires |
|---|---|
| Windows 95, Windows 98, Windows Millennium Edition (Windows Me) | Directory Services Client (Dsclient.exe) in the Clients\Win9x folder on the Windows 2000 Server CD-ROM |
| Windows NT Workstation 4.0 and Windows NT Server 4.0 | Service Pack 4 or later |
| Windows 2000 Professional and Windows 2000 Server | Included as part of the operating system |
| Windows XP Professional | Included as part of the operating system |
| Windows Server 2003 | Included as part of the operating system |

2. Enable the following Security Option in Domain Controller Security Policy:

   **Network security: LAN Manager authentication level to Send NTLMv2 responses/reject LM**

3. Enable the following Security Option in Domain Controller Security Policy:

   **Network security: Do not store LAN Manager hash value on next password change**

   Enabling this setting disables the creation of passwords in LM hash format.

4. Require all users to change their passwords immediately.

   The passwords that are already created in LM hash format are retained until the users change their passwords. Forcing password changes eliminates any passwords that are stored in LM hash format

| **Note** |
|---|
| For the sake of backward compatibility, if you cannot disable storage of your passwords in the LM hash format, you might recommend that your administrators use passwords with more than 14 characters. In the event that the password hashes are stolen, the administrator accounts are protected because accounts with a password of more than 14 characters do not have an LM hash. |

**SMB Signing on Domain Controllers**

On domain controllers that are running Windows Server 2003, default Group Policy settings allow the SMB Service and client to negotiate SMB packet signing. Domain controllers, member servers, and workstations access file shares during the user logon process to access logon scripts and profiles in the Netlogon share. In addition, domain policies are accessed through the SYSVOL share. For these reasons, all domain controllers should take advantage of SMB signing to improve security.

Table 20 lists the Security Options policy settings for SMB signing, and it explains how each setting affects client and server communications.

**Table 20 Security Options Policy Settings for SMB Packet Signing**

| SMB Setting | Explanation |
|---|---|
| Microsoft network client: Digitally sign communications (always) | The domain controller requires SMB signing when initiating SMB requests with other domain controllers, member servers, or workstations. The domain controller refuses to communicate with other systems that do not support SMB signing. For enhanced security, enable this Group Policy setting. |
| Microsoft network client: Digitally sign communications (if server agrees) | The domain controller negotiates SMB signing when initiating SMB requests with other domain controllers, member servers, or workstations. The domain controller requests SMB signing, but it will communicate with other systems that do not support SMB signing. For compatibility with Windows 95 and earlier operating systems, enable this Group Policy setting. |
| Microsoft network server: Digitally sign communications (always) | The domain controller requires SMB signing when receiving SMB requests from other domain controllers, member servers, or workstations. The domain controller refuses to communicate with other systems that do not support SMB signing. For enhanced security, enable this Group Policy setting. |
| Microsoft network server: Digitally sign communications (if client agrees) | The domain controller negotiates SMB signing when receiving SMB requests with other domain controllers, member servers, or workstations. The domain controller requests SMB signing, but it will communicate with other systems that do not support SMB signing. For compatibility with Windows 95 and earlier operating systems, enable this Group Policy setting. |

Enable the Security Option setting **Microsoft network client: Digitally sign communications (if server agrees)** in addition to **Microsoft network server: Digitally sign communications (always)** unless:

- Your network has computers that are running Windows for Workgroups; Windows 95 without the DS Client Pack; Windows NT 4.0 earlier than Service Pack 3.0; or devices, including Microsoft® Windows® Powered Pocket PC 2002 and previous versions, that are based on Microsoft® Windows® CE .NET Version 4.1 or earlier. It is highly recommended that you upgrade your clients rather than disabling this security setting. The DS Client Pack, which is necessary for Windows 95 clients to perform SMB signing, can be obtained from the \clients\win9x subdirectory on the Windows 2000 Server operating system CD.

- Your domain controllers, member servers, and workstations have insufficient available processor resources to support SMB signing. SMB signing generates higher processor utilization on the client side and the server side — an increase of up to 15 percent.

**Event Log Policy**

Because of the default domain controller Audit Policy settings, the maximum size of the security log must be increased to accommodate the increased number of audited events that might be generated.

The recommended Event Log policy settings reflect changes that are necessary for the Security log to support the default Audit Policy. In your environment, you may need to adjust the policy settings for the application or system event logs to support other operational goals.

**Note**

To implement changes to default Event Log policy, it is recommended that you create a new GPO. This GPO can be added and linked to the Domain Controllers OU above the level of the Default Domain Controllers GPO. In this way, the nondefault settings take precedence over the default settings, and you can also easily revert to default settings by simply deleting this GPO or placing it below the default GPO in the list of linked GPOs. For more information about creating this new GPO, see the Applying Selected Domain and Domain Controller Policy Settings section.

As a part of your normal operations tasks, archive the security and system event logs regularly and frequently before they fill up, which can cause events to be missed. The recommended Event Log policy settings allow the events in the security and

system event logs to be overwritten as needed. Back up the logs for future reference before any events can be overwritten.

Table 21 lists the default and recommended policy settings for domain controller Event Log policy settings.

**Table 21 Recommended Domain Controller Event Log Policy Settings**

| Policy | Default Setting | Recommended Setting | Comments |
|---|---|---|---|
| Maximum application log size | Not defined | (No change) | N/A |
| Maximum security log size | Not defined | 131,072 KB | Increased to accommodate security auditing that is enabled in the default domain controller Audit Policy. |
| Maximum system log size | Not defined | (No change) | N/A |
| Prevent local guests group from accessing application log | Not defined | Enabled | Prevents members of the built-in group Guests from reading the application log events. |
| Prevent local guests group from accessing security log | Not defined | Enabled | Prevents members of the built-i group Guests from reading the security log events. |
| Prevent local guests group from accessing system log | Not defined | Enabled | Prevents members of the built-in group Guests from reading the system log events. |
| Retain application log | Not defined | (No change) | N/A |
| Retain security log | Not defined | (No change) | N/A |
| Retain system log | Not defined | (No change) | N/A |
| Retention method for application log | Not defined | (No change) | N/A |
| Retention method for security log | Not defined | Overwrite events as needed | Overwrites the security log when the maximum log size is reached to ensure that the log contains the most recent security events and to ensure that logging continues. |
| Retention method for system log | Not defined | Overwrite events as needed | Overwrites the system log when the maximum log size is reached to ensure that the log contains the most recent security events and to ensure that logging continues. |

# Applying Selected Domain and Domain Controller Policy Settings

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The changed security policy settings that are recommended in this chapter apply to either the Default Domain Policy or the Default Domain Controllers Policy. After you have selected the list of settings that you want to update for your deployment, based on business considerations and the requirements for your environment, the method that you use to make the necessary changes depends on the specific policy or settings that you want to update.

To accommodate APIs from previous versions of the operating system that make changes directly to default GPOs, changes to the following security policy settings must be made directly in the Default Domain Policy GPO or in the Default Domain Controllers Policy GPO:

- Default Domain Security Policy Settings:
    - Password Policy
    - Domain Account Lockout Policy
    - Domain Kerberos Policy
- Default Domain Controller Security Policy Settings:
    - User Rights Assignment Policy
    - Audit Policy

Table 22 lists the Active Directory locations where the default policies are applied, the type of settings in each default policy, and the method for applying the recommended changes to the Group Policy settings.

**Table 22 Methods for Applying Changed Group Policy Settings**

| Location Where Policy is Applied | Policy Settings Being Changed | How to Apply the Changed Settings in Policy |
|---|---|---|
| Domain root | Domain security policy settings (Password Policy, Account Lockout Policy, and Kerberos Policy) | Make all changes to Security Settings in the Default Domain Policy GPO. |
| Domain Controllers OU | User Rights Assignment | Make all changes to Security Settings in the Default Domain Controllers Policy GPO. |
| Domain Controllers OU | Audit | Make all changes to Security Settings in the Default Domain Controllers Policy . |
| Domain Controllers OU | Event Log | Create a new GPO, and link this GPO above the Default Domain Controllers Policy GPO. |
| Domain Controllers OU | Security Options | Create a new GPO, and link this GPO above the Default Domain Controllers Policy GPO. |

When a new Windows Server 2003 domain is created, Active Directory creates a built-in, protected OU, called the Domain Controllers OU, directly under the domain root. Active Directory places the new domain controller computer account in this special OU. As additional domain controllers are promoted, the corresponding computer accounts are also placed in the Domain Controllers OU.

As a best practice, keep all domain controller computer accounts in the default Domain Controllers OU to ensure that domain-controller-specific Group Policy settings are consistently applied to all domain controllers in the domain.

**Modifying the Settings in the Default Domain Policy GPO and the Default Domain Controllers Policy GPO**
Apply changes directly in the Default Domain Policy GPO for Password Policy settings, Account Lockout Policy settings, and Kerberos Policy settings. Apply changes directly in the Default Domain Controllers Policy OU GPO for User Rights Assignment Policy settings and Audit Policy settings.

Update the Default Domain Policy GPO and the Default Domain Controllers Policy GPO by following the procedure "Updating the Default Domain Policy GPO and the Default Domain Controllers Policy GPO" in Appendix: Procedures later in this guide.

**Applying a New Group Policy Object to the Domain Controllers OU**
As a best practice, create and link a new GPO above the default policy when you want to apply changed policy settings to the entire domain or to all domain controllers, with the exception of the policies listed earlier that must be changed directly in the default policy. The advantage of this approach is that if a problem is encountered because of the changed settings, the new GPO can be easily backed out, restoring the original policy settings.

Apply the changed policy settings to the Domain Controllers OU through a new GPO by following the procedure "Creating a New GPO on the Domain Controllers OU and Changing its Precedence" in Appendix: Procedures later in this guide.

# Reviewing Audit Settings on Important Active Directory Objects

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Setting Audit Policy for domain controllers is only part of the task for auditing domain controllers for maximum security. In Windows Server 2003, auditing is also enabled by default on important Active Directory objects.

The audit settings that are described in this section are not enabled by default in Active Directory in Windows 2000. If you upgrade from Windows 2000, the audit settings at the time of the upgrade are maintained and the Windows Server 2003 auditing default settings are not applied. By adding the audit settings that are described in this section, you can bring your Windows Server 2003 Active Directory configuration into alignment with the default audit settings. If you are installing a new Windows Server 2003 domain, no additional audit settings are required.

Audit settings for an Active Directory object affect the permission settings that make up the SACL on the object. To provide auditing of access to the most important Active Directory objects, auditing is enabled by default on the topmost object of each directory partition. Directory partitions are logical divisions of the Active Directory database. In Windows Server 2003, auditing is enabled by default on the following Active Directory objects:

- The schema directory partition. Use ADSI Edit to set auditing on schema directory partition objects. ADSI Edit is a Windows Support Tool that is available as a Microsoft Management Console (MMC) snap-in when you have Windows Support Tools installed. For information about installing and using Windows Support Tools, click Tools in Help and Support Center, and then click Windows Support Tools.

- The configuration directory partition, including specific settings on the Sites and Services containers and on certain child objects. Use ADSI Edit to set auditing on configuration directory partition objects.

- Each domain directory partition, including specific settings on the Domain Controllers OU, Infrastructure object, System container, and certain child objects in the System container. Use Active Directory Users and Computers to set auditing on domain directory partition objects.

For information about how to enable audit settings, see “Enabling Auditing on Important Active Directory Objects” in Appendix: Procedures later in this guide.

Default audit settings for each directory partition are described in the following sections.

**Reviewing Default Audit Settings on the Schema Directory Partition**

The schema-related directory operations that are audited by the default settings in Table 23 include any additions, deletions, or modifications to objects in the Schema container in the schema directory partition, as well as the transfer of the Schema operations master role.

**Table 23 Auditing for CN=Schema,CN=Configuration,DC=ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Modify Permissions<br>Modify Owner<br>Create All Child Objects<br>Delete<br>Delete All Child Objects<br>Delete Subtree | This object only |
| Success | Everyone | Write All Properties | This object and all child objects |
| Success | Everyone | Change Schema Master | This object only |

| | | | |
|---|---|---|---|
| Success | Everyone | Reanimate Tombstones | This object only |
| Success | Administrators | All Extended Rights | This object only |
| Success | Domain Users | All Extended Rights | This object only |

**Reviewing Default Audit Settings on the Configuration Directory Partition**

Default audit settings for the Configuration container and child objects in the configuration directory partition are listed in Table 24, Table 25, Table 26, Table 27, and Table 28.

The directory operations that are audited by the settings in Table 24 include any modifications to the permissions and the **wellKnownObjects** attribute on the configuration directory partition.

**Table 24 Auditing for CN=Configuration,DC=ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Modify Permissions<br>Modify Owner<br>Write All Properties | This object only |
| Success | Everyone | Reanimate Tombstones | This object only |
| Success | Administrators | All Extended Rights | This object only |
| Success | Domain Users | All Extended Rights | This object only |

The directory operations that are audited by the settings in Table 25 include the following:

- Addition and removal of domain controllers in the forest
- Addition and removal of Group Policy settings that are applied to a site
- Association and disassociation of a subnet with a site
- Execution of the following control operations on a domain controller: Do Garbage Collection, Recalculate Hierarchy, Recalculate Security Inheritance, and Check Stale Phantoms
- Addition, removal, and modification of site links
- Addition, removal, and modification of connections

**Table 25 Auditing for CN=Sites,CN=Configuration,DC=ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Create All Child Objects<br>Delete<br>Delete All Child Objects<br>Delete Subtree | This object and all child objects |
| Success | Everyone | All Extended Rights | Domain Controller Settings objects |
| Success | Everyone | Write gPLink (property)<br>Write gPOptions (property) | Site objects |
| Success | Everyone | Write siteObject (property) | Subnet objects |

The directory operations that are audited by the settings in Table 26 include the following:

- Addition and removal of domains (or external directory knowledge references) in the forest

- Modifications to valid UPN Suffixes for the forest
- Transfer of the domain naming operations master role

**Table 26 Auditing for CN=Partitions,CN=Configuration,DC=ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Modify Permissions<br>Modify Owner<br>Write All Properties<br>Create All Child Objects<br>Delete<br>Delete All Child Objects<br>Delete Subtree<br>All Extended Rights | This object and all child objects |

The directory operations that are audited by the settings in Table 27 include changes to the **dSHeuristics** attribute, which controls certain characteristics of forest-wide behavior of the directory service.

**Table 27 Auditing for CN=Directory Service,CN=Windows,CN=Services,CN=Configuration,DC=ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Write dSHeuristics (property) | This object only |

The directory operations that are audited by the settings in Table 28 include changes to forest-wide parameters that govern the behavior of Lightweight Directory Access Protocol (LDAP)–based queries and operations.

**Table 28 Auditing for CN=Default Query Policy,CN=Query-Policies,CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,DC=ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Write lDAPAdminLimits (property) | This object only |

**Reviewing Default Audit Settings on the Domain Directory Partition**

Default audit settings for objects in the domain container of the domain directory partition are listed in Table 29, Table 30, Table 31, Table 32, Table 33, and Table 34.

The directory operations that are audited by the settings in Table 29 include the following:

- Transfer of the PDC emulator operations master role
- Addition and removal of Group Policy settings that are applied to the domain
- Modifications to valid DNS Suffixes for the domain
- Modifications to the permissions and the **wellKnownObjects** attribute on the domain directory partition.
- Migration of SID history

**Table 29 Auditing for DC=domain,DC=ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Modify Permissions<br>Modify Owner<br>Write All Properties | This object only |

| | | | |
|---|---|---|---|
| Success | Administrators | All Extended Rights | This object only |
| Success | Domain Users | All Extended Rights | This object only |
| Success* | Everyone | Write gPLink<br>Write gPOptions | Organizational Unit objects |

* Do not add these policies if you are upgrading from Windows 2000 to Windows Server 2003. Doing so causes the security descriptor on every object in the domain to be edited, which can result in significant database growth and increased background processing to apply these access control entries (ACEs) to all objects.

The directory operations that are audited by the settings in Table 30 include the following:

- Addition and removal of domain controllers for the domain
- Modifications to any properties of domain controller computer accounts

**Table 30 Auditing for OU=Domain Controllers,DC=domain,DC=...ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Modify Permissions<br>Modify Owner<br>Create All Child Objects<br>Delete<br>Delete All Child Objects<br>Delete Subtree | This object only |
| Success | Everyone | Write All Properties | This object and all child objects |

The directory operations that are audited by the settings in Table 31 include the transfer of the infrastructure operations master role.

**Table 31 Settings for CN=Infrastructure,DC=domain,DC=...ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | All Extended Rights<br>Write All Properties | This object only |

**Reviewing Default Audit Settings on the Policies Container**

The directory operations that are audited by the settings in Table 32 include the following:

- Addition and removal of GPOs
- Modifications to GPOs

**Table 32 Settings for CN=Policies,CN=System,DC=domain,DC=...ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Modify Permissions<br>Modify Owner<br>Create groupPolicyContainer Objects<br>Delete<br>Delete groupPolicyContainer Objects<br>Delete Subtree | This object only |

| | | | |
|---|---|---|---|
| Success | Everyone | Modify Permissions<br><br>Write All Properties | groupPolicyContainer objects |

The directory operations that are audited by the settings in Table 33 include modifications to the special security descriptor that protects all service administrator accounts.

**Table 33 Settings for CN=AdminSDHolder,CN=System,DC=domain,DC=...ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | Modify Permissions<br><br>Modify Owner<br><br>Write All Properties | This object only |

The directory operations that are audited by the settings in Table 34 include the transfer of the relative ID (RID) operations master role.

**Table 34 Settings for CN=RID Manager$,CN=System,DC=domain,DC=...ForestRootDomain**

| Type | Name | Access | Apply To |
|---|---|---|---|
| Success | Everyone | All Extended Rights<br><br>Write All Properties | This object only |

# Recommendations: Strengthening Domain and Domain Controller Policy Settings

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

You can create secure domain and domain controller policy settings by following the security recommendations described earlier in this section. Of course, as previously mentioned, your comprehensive Active Directory security plan should also take into consideration the recommendations that are described in the other sections of this guide.

**Recommendations for Strengthening Domain Policy Settings**
Recommendations for ensuring that your domain and domain controller policies are applied in a secure and consistent manner are:

- Strengthening Password Policy Settings for Domains
    - Apply the recommended password policy settings (Table 13) at the domain level.
- Strengthening Account Lockout Policy Settings for Domains
    - Apply the recommended account lockout policy settings (Table 14) at the domain level.
- Reviewing Kerberos Policy Settings for Domains
    - Review the recommended Kerberos policy settings (Table 15) at the domain level.

**Recommendations for Strengthening Domain Controller Policy Settings**
Recommendations for ensuring that your domain controller policy settings are applied in a secure and consistent manner are:

- Reviewing Domain Controller Audit Policy Settings
    - Review default auditing that provides accountability for sensitive operations and intrusion detection.
- Strengthening Domain Controller User Rights Assignment Policy Settings
    - Apply the recommended domain controller user rights assignment policy settings (Table 17).
- Strengthening Domain Controller Security Options Policy Settings
    - Apply the recommended domain controller Security Options policy settings (Table 18).
    - Configure all domain controllers and clients to support NTLMv2, and then use Security Options policy settings to disable LM authentication.
    - Configure all domain controllers and clients to support NTLMv2, and then use Security Options policy settings to disable LM authentication.
    - Enable SMB client and server signing so that SMB signing is always required.
- Strengthening Domain Controller Event Log Policy Settings
    - Apply recommended domain controller event log policy settings (21).

**Recommendations for Applying Selected Domain and Domain Controller Policy Settings**
Recommendations for ensuring that your domain controller policy settings are applied in a secure and consistent manner are:

- Modifying the settings in the Default Domain Policy GPO and the Default Domain Controllers Policy GPO
    - Make all changes to Security Settings in the Default Domain Policy GPO.
    - Make all changes to User Rights Assignment and Audit Policy settings in the Default Domain Controllers Policy GPO.
- Applying a New GPO to the Domain Controllers OU
    - Ensure that all domain controller computer accounts reside in the Domain Controllers OU.
    - Create and link a new GPO above the Default Domain Controllers Policy GPO. This GPO should contain the recommended changes to Security Options and Event Log settings in the security policy.

**Recommendations for Reviewing Audit Settings on Important Active Directory Objects**
The following table provides a checklist of recommendations for ensuring that important Active Directory objects are audited. If you are upgrading a domain from Windows 2000, the default Windows Server 2003 auditing settings for these objects are not applied. In this case, use the settings in this section to enable auditing on important Active Directory objects.

- Reviewing Default Audit Settings on Active Directory Database Objects
    - Review or enable auditing on schema directory partition objects.
    - Review or enable auditing on configuration directory partition objects.
    - Review or enable auditing on domain directory partition objects.

# Chapter 5: Establishing Secure Administrative Practices

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Any user who has administrative access to domain controllers can cause breaches in security. Individuals seeking to damage the system might be unauthorized users who have obtained administrative passwords, or they might be legitimate administrators who are coerced or disgruntled. Furthermore, not all problems are caused with malicious intent. An inexperienced user who is granted administrative access might inadvertently cause problems by failing to understand the ramifications of configuration changes.

You can minimize these problems by carefully controlling the scope of influence that you give to administrative accounts. For the day-to-day management of your environment, avoid using all-powerful administrative accounts that have complete access to every domain controller and full access to the directory. Instead, configure administrative accounts so that their scope of influence is limited to the specific containers in Active Directory that they need to do their jobs. In the event that one of these accounts is misused, the amount of damage that can be done is limited.

For Active Directory in Windows Server 2003, there are two types of administrative responsibility: service administration and data administration:

- Service administrators are responsible for maintaining and delivering the directory service, including domain controller management and directory service configuration.

- Data administrators are responsible for maintaining the data that is stored in the directory service and on domain member servers.

Some information that is needed to manage or configure the Active Directory service is controlled by objects that are stored in the directory itself. Although this information is stored in the directory, it is used and managed by service administrators. For this reason, service administrators can act as data administrators. Due to their limited access, data administrators cannot act as service administrators.

# Establishing Secure Service Administration Practices

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Service administrators have the most widespread power in your network environment. They are responsible for the delivery of the directory service, directory-wide settings, installation and maintenance of software, and application of operating system service packs and hotfixes on domain controllers. They also manage the overall configuration of the directory service, including replication behavior, schema management, and domain creation and deletion. To perform most of these functions, service administrators must have physical access to domain controllers.

Because service administrators have elevated privileges, take appropriate steps to keep their accounts secure:

- Secure service administrator accounts to control how the accounts are used.
- Secure service administrator workstations to control where the accounts can be used.
- Avoid delegating security-sensitive operations.

**Securing Service Administrator Accounts**

As with regular user access, administrative access to resources is controlled by the permissions that are set on the various resources being managed and by the privileges that are granted to the administrative accounts. Active Directory has a number of default service administrator accounts. By default, these accounts are granted access to directory and server resources when Active Directory is installed. Table 35 lists the default service administrator accounts and provides a brief description of each account.

**Table 35 Default Service Administrator Accounts**

| Account Name | Scope | Description |
|---|---|---|
| Enterprise Admins (EA) | Forest | This group is automatically added to the Administrators group in every domain in the forest, providing complete access to the configuration of all domain controllers. This group can modify the membership of all administrative groups. Its own membership can be modified only by the default service administrator groups in the forest root domain. This account is considered to be a service administrator. |
| Schema Admins (SA) | Forest | This group has full administrative access to the schema. The membership of this group can be modified by any of the service administrator groups in the forest root domain. This account is considered to be a service administrator because its members can modify the schema, which governs the structure and content of the entire directory. |
| Administrators (BA) | Domain | This built-in group controls access to all domain controllers in its domain, and it can change the membership of all administrative groups in the domain. Its own membership can be modified by the default service administrator groups BA and DA in the domain, as well as by the EA group. This group has the special privilege of taking ownership of any object in the directory or of any resource on a domain controller. This account is considered to be a service administrator because its members have full access to the domain controllers in the domain. |
| Domain Admins (DA) | Domain | This group controls access to all domain controllers in a domain, and it can modify the membership of all administrative accounts in the domain. Its own membership can be modified by the service administrator groups BA and DA in its domain, as well as by the EA group. This account is a service administrator account because its members have full access to a domain's domain controllers. |
| Server Operators (SO) | Domain | By default, this built-in group has no members, and it has access to server configuration options on domain controllers. Its membership is controlled by the service administrator groups BA and DA in the domain, as well as by the EA group. It cannot change any administrative group memberships. This account is a service administrator account because its members have physical access to domain controllers and they can perform maintenance tasks, such as backup and restore, and they have the ability to change binaries that are installed on the domain controllers. |

| | | |
|---|---|---|
| Account Operators (AO) | Domain | By default, this built-in group has no members, and it can create and manage users and groups in the domain, but it cannot manage service administrator accounts. This account can modify computer accounts, including those accounts that belong to domain controllers. As a best practice, leave the membership of this group empty, and do not use it for any delegated administration. |
| Backup Operators (BO) | Domain | By default, this built-in group has no members, and it can perform backup and restore operations on domain controllers. Its membership can be modified by the default service administrator groups BA and DA in the domain, as well as by the EA group. It cannot modify the membership of any administrative groups. While members of this group cannot change server settings or modify the configuration of the directory, they do have the permissions needed to replace files (including operating system binaries) on domain controllers. For this reason, members of this group are considered service administrators. |
| Administrator | DS Restore Mode | This special account is created during the Active Directory installation process, and it is not the same as the Administrator account in the Active Directory database. This account is only used to start the domain controller in Directory Services Restore Mode. When it is in Directory Services Restore Mode, this account has full access to the directory database, as well as to files (including operating system binaries) on the domain controller. For this reason, this account is considered a service administrator. |

**Limiting the Exposure of Service Administrator Accounts**

Service administrator accounts are highly privileged, making them desirable targets for attack. Therefore, it is especially important to protect the integrity of these accounts. The recommendations in this section are designed to enhance the security of your service administrator accounts. Limiting the exposure of the service administrator accounts gives attackers fewer targets of opportunity. Observe the practices in the following sections to limit the exposure of service administrator accounts.

**Limit the Number of Service Administrator Accounts**

Keep the membership of service administrator accounts to the absolute minimum that is necessary to support your organization. Tasks that are performed by service administrators should be limited to changing the Active Directory service configuration and reconfiguring domain controllers. Do not use service administrator accounts for day-to-day administrative tasks, such as account and member server management; instead, delegate these tasks to data administrators.

**Separate Administrative and User Accounts for Administrative Users**

For users who fill administrative roles, create two accounts: one regular user account to be used for normal, day-to-day tasks and one administrative account to be used only for performing administrative tasks. The administrative account should not be mail enabled or used for running applications that are used every day, such as Microsoft Office, or for browsing the Internet. These precautions reduce the exposure of the accounts to the outside world, and they reduce the amount of time that administrative accounts are logged on to the system.

**Hide the Domain Administrator Account**

Every installation of Active Directory has an account named Administrator in each domain. This is the default administrative account, which is created during domain setup, that you use to access and administer the directory service. This is a special account that the system protects to help ensure that it is available when needed. This account cannot be disabled or locked out.

The fact that this account is always created during domain setup, cannot be deleted, and cannot be disabled means that every malicious user who attempts to break into your system assumes that the account exists and that can it can be used as a target. For this reason, you should rename it to something other than Administrator. When you rename the account, make sure that you also change the text in the **Description** for the account. In addition, you should create a decoy user account called Administrator that has no special permissions or user rights.

To hide the default Administrator account, perform the following tasks:

1. Rename the default Administrator account. For more information about performing this task, see “Renaming the Default Administrator Account” in the Appendix: Procedures section.

2. Create a decoy Administrator account with no special permissions or privileges. For more information about performing this task, see “Creating a Decoy Administrator Account” in the Appendix: Procedures section.

**Managing Service Administrators in a Controlled Subtree**

To help protect highly privileged service administrator accounts, allow only service administrators to manage service administrator accounts. Because these accounts have elevated privileges, data administrators should not be given the authority to modify these accounts. Doing so allows data administrators to elevate their privileges. Service administrator accounts should be accessed and managed in a highly controlled subtree in each domain.

To provide a more controlled environment that facilitates the management of service administrator accounts and workstations, create a controlled OU subtree to manage service administrator accounts in Active Directory. The controlled subtree structure should be created for each domain by a member of the Domain Admins group for that domain, and it should be configured with its recommended security settings.

To create the controlled subtree, perform the following tasks:

1. Create the OU structure for the controlled subtree.

2. Set the permissions on the controlled subtree OUs.

3. Move service administrator groups to the controlled subtree.

4. Add service administrator user accounts to the controlled subtree.

5. Add service administrator workstation accounts to the controlled subtree.

6. Enable auditing on the controlled subtree OU.

By creating your own subtree containing all service administrator accounts and the administrative workstations that they use, you can apply controlled security and policy settings to them to maximize their protection.

**Task 1: Create the OU structure for the controlled subtree.**
Create a high-level OU to hold the groups and user accounts that constitute your service administrators and their workstations. Within that OU, create one OU to hold administrative user and group accounts and another OU for administrative workstations. Figure 10 depicts the recommended OU hierarchy for a controlled subtree to manage service administrator accounts and workstations. It consists of a controlled subtree, rooted at the Service Admins OU, that contains two additional OUs: the Users and Groups OU, to hold the administrative user and group accounts, and the Admin Workstations OU, to hold the computer accounts of the workstations that are used by the service administrators.




**Task 2: Set the permissions on the controlled subtree OUs.**
To limit access to the controlled subtree so that only service administrators can administer the membership of service administrator groups and workstations, do the following:

- Block inheritance of permissions on the Service Admins OU so that changes that are made higher up in the domain tree are not inherited down, altering the locked-down settings

- Set the permissions (ACLs) on the Service Admins OU, Users and Groups OU, and Admin Workstations OU, as indicated in Table 36.

**Table 36 Permission Settings for the Controlled Subtree OUs**

| Type | Name | Access | Applies To |
|---|---|---|---|

| | | | |
|---|---|---|---|
| Allow | Enterprise Admins | Full Control | This object and all child objects |
| Allow | Domain Admins | Full Control | This object and all child objects |
| Allow | Administrators | Full Control | This object and all child objects |
| Allow | Pre–Windows 2000 Compatible Access | List Contents<br>Read All Properties<br>Read Permissions | User and InetOrgPerson objects |
| Allow | Enterprise Domain Controllers | List Contents<br>Read All Properties<br>Read Permissions | This object only |
| Allow | Enterprise Domain Controllers | Read All Properties | User, group, and computer objects |

The following figure shows the Service Admins OU with permissions that control access to the contents of the OU.




**Task 3: Move service administrator groups to the controlled subtree.**

Move the following service administrator groups from their current location in the directory into the Users and Groups OU in your controlled subtree:

- Domain Admins and any nested subgroups
- Enterprise Admins (if this is the root domain of the forest) and any nested subgroups
- Schema Admins (if this is the root domain of the forest) and any nested subgroups
- Any groups that are nested in the domain's Administrators, Server Operators, Backup Operators, or Account Operators groups

- Any group that has delegated rights that effectively grant its users service administrator rights

The built-in groups from Table 35 (Administrators, Server Operators, Account Operators, and Backup Operators) cannot be moved from their default container to the controlled subtree. However, built-in groups are protected by default in Windows Server 2003, as described in Table 38 later in this guide, and they need no additional protection.

**Task 4: Add service administrator user accounts to the controlled subtree.**
Move all administrative user accounts that are members of any of the service administrator groups listed in Table 35 from their current locations in the directory into the Users and Groups OU in your controlled subtree. This set of accounts includes the domain Administrator account.

As recommended, each service administrator should have two accounts: one for administrative duties and one for their normal user access. Place the administrative user accounts in the Users and Groups OU in your controlled subtree. If these accounts already exist elsewhere in the directory, move them into the subtree now. The regular user accounts for those administrators should not be placed in this controlled subtree.

**Task 5: Add administrative workstation accounts to the controlled subtree.**
Designate administrators' computers as administrative workstations. Move the computer accounts for those workstations into the Admin Workstations OU in your controlled subtree.

| **Important** |
|---|
| Do not move any domain controller accounts out of the default Domain Controllers OU, even if some administrators log on to them to perform administrative tasks. Moving these accounts will disrupt the consistent application of domain controller policies to all domains. |

**Task 6: Enable auditing on the controlled subtree.**
It is important to audit and track any additions, deletions, and changes to the service administrator accounts and workstations — and changes in policies that are applied to them — so that improper or unauthorized changes can be detected. Assuming that you have enabled auditing on your domain controllers in accordance with the recommendations in Chapter 4: Strengthening Domain and Domain Controller Policy Settings, set the SACL on the Service Admins OU by setting auditing, as specified in Table 37. Monitor the security audit log for changes to the controlled subtree, so that you can ascertain that the changes are legitimate.

**Table 37 Auditing Settings on OU=Service Admins,DC=domain name**

| Type | Name | Access | Applies To |
|---|---|---|---|
| All | Everyone | Write All Properties | This object and all child objects |
| All | Everyone | Delete | This object and all child objects |
| All | Everyone | Delete Subtree | This object and all child objects |
| All | Everyone | Modify Permissions | This object and all child objects |
| All | Everyone | Modify Owner | This object and all child objects |
| All | Everyone | All Validated Writes | This object and all child objects |
| All | Everyone | All Extended Rights | This object and all child objects |
| All | Everyone | Create All Child Objects | This object and all child objects |
| All | Everyone | Delete All Child Objects | This object and all child objects |

At this point, the controlled subtree for service administrators is set up and ready for use.

**Protecting the Service Administrator Accounts**
To prevent the security descriptors on the key service administrator accounts in each domain from being modified and possibly becoming unusable, a background process runs on the PDC emulator that periodically checks and applies a standard security descriptor on the protected accounts. This process ensures that if a hostile user or other administrator does manage to modify the security descriptor on one of the administrative accounts, the change will be overwritten with the protected settings. This process starts 15 minutes after the system starts, and then it runs once every half hour after that. This refresh interval is not configurable. In Active Directory in Windows Server 2003, all the service administrative groups that are listed in

Table 35 and all their nested member groups and users are protected by this process.

| **Note** |
| --- |
| To protect these accounts during an upgrade to Windows Server 2003, upgrade the PDC emulator first. |

The master security descriptor for service administrator accounts is stored as the security descriptor attribute of the **AdminSDHolder** object, which is located in the system container (CN=AdminSDHolder,CN=System,DC=*DomainName*) of the domain directory partition.

The security descriptor on this object serves two purposes:

- It controls access to the **AdminSDHolder** object itself.
- It acts as the master security descriptor, which is periodically applied to the service administrator groups and their members to ensure that they remain protected.

The default settings in the master security descriptor of the **AdminSDHolder** object are listed in Table 38.

**Table 38 Default Security Descriptor on the AdminSDHolder Object Used to Protect Service Administrator Accounts**

| Type | Name | Permission | Apply To |
| --- | --- | --- | --- |
| Allow | Administrators | List Contents<br>Read All Properties<br>Write All Properties<br>Delete<br>Read Permissions<br>Modify Permissions<br>Modify Owner<br>All Validated Writes<br>All Extended Rights<br>Create All Child Objects<br>Delete All Child Objects | This object only |
| Allow | Authenticated Users | List Contents<br>Read All Properties<br>Read Permissions | This object only |

| | | | |
|---|---|---|---|
| Allow | Domain Admins | List Contents<br>Read All Properties<br>Write All Properties<br>Read Permissions<br>Modify Permissions<br>Modify Owner<br>All Validated Writes<br>All Extended Rights<br>Create All Child Objects<br>Delete All Child Objects | This object only |
| Allow | Enterprise Admins | List Contents<br>Read All Properties<br>Write All Properties<br>Read Permissions<br>Modify Permissions<br>Modify Owner<br>All Validated Writes<br>All Extended Rights<br>Create All Child Objects<br>Delete All Child Objects | This object only |
| Allow | Everyone | Change Password | This object only |
| Allow | Pre–Windows 2000 Compatible Access | List Contents<br>Read All Properties<br>Read Permissions | User and InetOrgPerson objects |
| Allow | SYSTEM | Full Control | This object only |
| Allow | SELF | Change Password | This object only |
| Allow | Cert Publishers | Read userCert<br>Write userCert | This object only |
| Allow | Windows Authorization Access Group | Read tokenGroupsGlobalAndUniversal | This object only |
| Allow | Terminal Server License Servers | Read terminalServer<br>Write terminalServer | This object only |

If you want to modify the security descriptor on one of the service administrator groups or on any of its member accounts, you must modify the security descriptor on the AdminSDHolder object so that it will be applied consistently. Be careful when making these modifications because you are also changing the default settings that will be applied to all of your protected administrative accounts. For more information about changing the security descriptor on AdminSDHolder, see "Changing the Security Descriptor on AdminSDHolder" in the Appendix: Procedures section.

**Hiding the Membership of the Service Administrator Groups**

Because members of the service administrator groups are highly privileged, they constitute an attractive target for attackers.

Therefore, the membership information of these groups should be guarded as much as possible. For maximum security, it is best to hide the membership information for all service administrator groups from regular users. However, the default security descriptor on **AdminSDHolder** that is used to protect service administrator groups allows their membership information to be visible to regular users.

The entry in Table 38 that grants Read access to Authenticated Users allows those users to list the membership information for all service administrator groups. Some applications and services rely on the presence of this permission to enumerate the members of an administrative group. For example, applications such as Microsoft SQL Server, which run under service accounts, might read the list of administrative group members to make its own application-specific access control decisions. Other non-Microsoft applications and line-of-business applications might also rely on this capability. Because there is no way to anticipate every service account that will need this type of Read access, the Authenticated Users entry exists to support all such cases.

It is possible to tighten security further by removing Read access for Authenticated Users from the security descriptor on **AdminSDHolder**, although doing so can cause some of your server-based applications to stop functioning. To ensure the proper functioning of these applications, systematically remove Read access for Authenticated Users by performing the following tasks:

1. If you have not already done so, disable pre–Windows 2000 compatible access for your domain. For more information, see "Disabling Anonymous Active Directory Access" in the Establishing Secure Domain Controller Build Practices earlier in this guide.

2. Create a group called Server Applications, and grant it Read access to **AdminSDHolder** by adding the permissions shown in Table 39.

3. Add the individual service accounts that are used by your applications that require the ability to enumerate group membership of the service administrator groups to the Server Applications group.

4. Remove the Authenticated User and the Pre–Windows 2000 Compatible Access entries from the security descriptor.

**Table 39 Permission Changes on AdminSDHolder for the Server Applications Group**

| Type | Name | Permission | Apply To |
|---|---|---|---|
| Allow | Server Applications | List Contents | This Object Only |
| Allow | Server Applications | Read All Properties | This Object Only |
| Allow | Server Applications | Read Permissions | This Object Only |

At this point, your service administrator accounts should not be visible to regular users on your network. Because it is impossible to predict the impact on every application, closely monitor applications running in your environment and make sure that they still function properly. If you observe application problems, simply add Authenticated Users as a member of the newly created Server Applications group to restore functionality, while you diagnose how to remove the application dependency.

**Managing Group Memberships for Service Administrator Accounts**

To enhance security, limit the membership in each of the service administrator groups to the absolute minimum that your organizational logistics allow, while still enabling you to manage your Active Directory service functions. Limiting these memberships reduces the number of possible administrative accounts that can be compromised by malicious users. The practices in the following sections are recommended for managing the membership of the service administrator groups.

**Assign Trustworthy Personnel**

Service administrators control the configuration and functioning of the directory service. Therefore, this responsibility should be given only to reliable, trusted users who have demonstrated responsible ownership and who fully understand the operation of the directory. They should be completely familiar with your organization's policies regarding security and operations, and they should have demonstrated their willingness to enforce those policies

**Restrict Service Group Membership to Users Within the Forest**

Do not include users or groups from another forest as members of service administrator groups, unless you completely trust the service administrators of the other forest. Because service administrators in the other forest have full control on the user accounts in that forest, they can easily impersonate or authenticate to your forest using the credentials of one of those users. Further, by trusting the remote domain (or forest) in this way, you also trust their security measures, which is something you have no control over.

If it is necessary to have a user from another forest act as a service administrator in your domain, create an account in your domain that he or she can use when performing the administrative role. By using a local account, you eliminate dependence on the security measures of the other forest.

**Limit the Schema Admins Group to Temporary Members**

The Schema Admins group is a special group in the forest root domain that provides administrative access to the Active Directory schema. Members of this group have the necessary user rights to make changes to the schema. In general, because schema changes are rare, it is not necessary for a schema administrator to be available at all times. This account is needed only when a schema update must be processed or if a change must be made to the configuration of the schema master role holder.

To minimize the possibility of an Active Directory attack through a schema administrator account, keep the membership of the Schema Admins group empty. Add a trusted user to the group only when an administrative task must to be performed on the schema. Remove that user after the task is completed.

**Limit Administrator Rights to Those Rights That Are Actually Required**

Active Directory contains a Backup Operators built-in group. Members of this group are considered to be service administrators because members of the group have the privilege to log on locally and restore files, including the system files, on domain controllers. Membership in the Backup Operators group in Active Directory should be limited to those individuals who back up and restore domain controllers.

All member servers also contain a Backup Operators built-in group that is local to each server. Individuals who are responsible for backing up applications on a member server (for example, Microsoft SQL Server) should be made members of the local Backup Operators group on that server. These users should not be members of the Backup Operators group in Active Directory.

On a dedicated domain controller, you can reduce the number of members in the Backup Operators group. When a domain controller is used to run other applications, as might be the case in a branch office, users who are responsible for backing up applications on the domain controller must also be trusted as service administrators because they have the privileges that enable them to restore files, including the system files, on domain controllers.

Avoid using the Account Operators group for strictly delegating a "data administration" task, such as account management. The default directory permissions give this group the ability to modify the computer accounts of domain controllers, including deleting them. By default, there are no members of the Account Operators group, and its membership should be left empty.

**Controlling the Administrative Logon Process**

The members of the Administrators, Enterprise Admins, and Domain Admins groups represent the most powerful accounts in your forest and in each individual domain. To minimize security risks, take the steps in the following sections to enforce strong administrative credentials.

**Require Smart Cards for Administrative Logon**

Require your service administrators to use smart cards for their interactive logons. In addition to forcing the administrative users to have physical possession of the cards to log on, smart cards also ensure the use of randomly generated, cryptographically strong passwords on the user accounts. These strong passwords help to protect against the theft of weak passwords to gain administrative access. To implement this strategy, you must have a PKI available to authenticate the smart cards.

You can enforce the use of smart cards by enabling the **Interactive logon: Require smart card** security option for each administrative account. If you create a subtree, as described in Managing Service Administrators in a Controlled Subtree, set this option on the Users and Groups OU.

For more information about using smart cards authentication, see The Smart Card Deployment Cookbook on the Microsoft TechNet Web site at http://go.microsoft.com/fwlink/?LinkId=18552, and see "Designing a Public Key Infrastructure" and "Planning a Smart Card Deployment" in Designing and Deploying Directory and Security Services of the Windows Server 2003 Deployment Kit (or see "Designing a Public Key Infrastructure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4735 and "Planning a Smart Card Deployment" on the Web at http://go.microsoft.com/fwlink/?LinkId=4736).

**Share Logon Credentials for Sensitive Administrative Accounts**

For each account that is a member of the Enterprise Admins and Domain Admins groups in the forest root domain, assign two users to share that account so that both users must be present to log on successfully with that account. Sharing these accounts provides inherent visual auditing of the use of the accounts, in which one user observes the actions that are performed by the other. It also means that a single user cannot log on privately to access the system as an administrator and compromise its security, either as a rogue administrator or under coerced circumstances.

Shared administrative accounts can be implemented by using either split passwords or split smart cards plus personal identification numbers (PINs). If you are using:

- Password-based credentials for administrative accounts:

  Split the password for the administrative account between the two users who share that account so that each user knows only half of the password. Each user is responsible for maintaining their half of the password. For example, you can create an administrative account called Admin1 and assign two trusted users, Jane and Bob, to share this account. Each of them maintains half of the password. For one of them to log on and use the account, the other must be present to enter the other half of the password.

- Smart-card-based credentials for administrative accounts:

  Split the ownership of the smart card and its PIN between the two users sharing that account so that one user retains physical ownership of the smart card and the other user maintains the PIN for the card. This way, both users must be present to log on to the account.

**Securing Service Administrator Workstations**

In addition to limiting access to resources that are stored on the domain controllers and access to information that is stored in the directory, you can also enhance security by strictly controlling the workstations that are used by service administrators for administrative functions.

**Restricting Service Administrators Logon to Administrative Workstations**

Each administrative account can be restricted so that it is allowed to log on only to specific workstations. If one of your administrative accounts is compromised, limiting the possible workstations limits the number of locations where that account can be used.

As described in described in the Managing Service Administrators in a Controlled Subtree section, the process of creating a controlled subtree of OUs for the management of service administrator accounts can be used for managing administrative workstations as well. Move the administrative workstations into this subtree so that you can apply Group Policy settings that control who can log on at which location.

To limit the locations where the service administrator accounts can log on, perform the following tasks:

1. Modify the User Rights Assignment policy in the Default Domain Policy GPO (Default Domain Security Settings) to **Deny log on locally** to the following groups: Schema Admins, Enterprise Admins, Domain Admins, Server Operators, Backup Operators, and Account Operators.

   For more information about performing this task, see "Denying Logon Access to the Domain" in Appendix: Procedures later in this guide.

2. Define the **User Rights Assignment** policy setting **Deny log on locally**, but do not assign the right to any group. By default, the **Deny log on locally** right on the Admin Workstations OU is not defined. If you leave it that way, the **Deny log on locally** right in the Default Domain Policy GPO applies.

   For more information about performing this task, see "Allowing Logon Access to Administrative Workstations" in Appendix: Procedures later in this guide.

   The figure below shows the application of the **Deny log on locally** right for the various service administrative accounts to the Default Domain Policy and the same **Deny log on locally** right defined on the Admin Workstations OU but unassigned (not applied to any accounts). The policy that is applied at the OU level overrides the domain-level policy.




| Note |
| --- |
| The Default Domain Controllers GPO, which is applied to the Domain Controllers OU, allows administrators to log on to domain controllers. Denying logon access across the entire domain applies to all workstations and member servers. |

**Prohibiting the Use of Cached Credentials in Unlocking an Administrative Console**

When the console on an administrative workstation is locked, either by the action of a user or automatically by a screensaver time-out, the console must be unlocked to regain access to the workstation. The workstation maintains cached credentials for any users that have been authenticated locally. When the console is unlocked, by default the workstation uses these cached credentials, if they exist, for the user who attempts to unlock the console.

When cached credentials are used to unlock the console, any changes to the account, such as user rights assignment, group membership changes, or disabling of the account, are not enforced. For example, if an administrator who is logged on to a workstation console is terminated, he can still unlock the console, even if his account is disabled. To ensure that any changes to the account are enforced immediately, require domain controller authentication of the account to unlock the console, instead of cached credentials.

To ensure that cached credentials cannot be used to unlock an administrative workstation, set the security option Interactive logon: Require Domain Controller authentication to unlock workstation to Enabled in Group Policy for the Admin Workstations OU.

**Avoiding Running Applications in Administrative Contexts**

When a computer becomes a member of a domain, the Domain Admins group is automatically added to the local Administrators group on that computer. This membership gives Domain Admins full access to the member computer. However, if the computer contracts a virus while a domain administrator is logged on, the virus runs in the context of that domain administrator, who is also a local administrator. In this way, the virus is able to use the administrator's privileges to infect the workstation.

To avoid running applications in the context of a service administrator account, remove the Domain Admins group from the local Administrators group on the workstations in the Admin Workstations OU in your controlled subtree. If a service administrator contracts a virus while logged on, the virus will not have administrative access to the computer.

Also, avoid running application or service agents on administrative workstations in the context of a service administrator account. For example, a monitoring agent running in the context of a Domain Admins member might make it possible for a compromised agent to breach the security of the domain and the domain controllers. The local administrator account on the workstation has the necessary privileges to take control of the agent and act as the service administrator. If a user manages to gain access to the system as a local administrator, that user can then assume control of the agent and any information that it is collecting.

**Running Antivirus Software**

Running antivirus software regularly on administrative workstations helps protect the workstations from contracting a virus that can exploit the privileges of a logged-on service administrator to spread itself through the computer and into the domain.

**Securing Traffic Between Administrative Workstations and Domain Controllers**

Table 40 lists the features in Windows Server 2003 that can be used to secure traffic between administrative workstations and

domain controllers, along with the threats that are mitigated by each feature.

**Table 40 Features in Windows Server 2003 That Secure Administrative Traffic**

| Feature | Requirements | Information Disclosure Threat | Data Tampering Threat |
|---|---|---|---|
| LDAP packet encryption and signing | Windows Server 2003 or Windows XP with Adminpak.msi | Mitigated | Mitigated |
| LDAP packet signing | Windows 2000 SP3 and later | N/A | Mitigated |
| Terminal Services (Remote Desktop for Administration) | Windows Server 2003 | Obfuscated | Obfuscated |

**Securing LDAP Traffic Between Administrative Workstations and Domain Controllers**

On administrative workstations that are running Windows XP Professional, you can install the Windows Server 2003 Administration Tools Pack (Adminpak.msi), which can be installed from the i386 directory on the Windows Server 2003 CD. This version of the Administration Tools Pack encrypts and signs LDAP traffic between the administrative tool clients and domain controllers. If your administrative workstations are running Windows Server 2003, the administration tools are installed by default; you do not have to install the Administration Tools Pack.

**Important**

For the Administration Tools Pack to be installed on Windows XP Professional, you must either apply Service Pack 1 to the operating system or install the hotfix that is included with article 329357, "Windows Server 2003 Adminpak.msi Does Not Run on a Windows XP-Base Computer Without Service Pack 1," in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkId=4441.

When LDAP traffic is not encrypted and signed, it is possible for attackers to gain access to sensitive information by intercepting network traffic between administrative workstations and domain controllers. Encrypting and signing LDAP traffic protects against two types of vulnerabilities to the data:

- Information disclosure: A user can gain access to data, intentionally or otherwise, that the user is not authorized to see.
- Data tampering: A user can modify system or user data, with or without detection.

If administrative workstations are running Windows 2000, they cannot run the Administration Tools Pack. However, you can set Group Policy to force signing of all LDAP traffic from the administrative workstation.

**Note**

Windows 2000 Server SP3 and later is updated to support signing and encryption. However, the Windows 2000 Server administrative tools that shipped with Windows 2000 Server do not have the updates. Windows Server 2003 administrative tools support signing and encryption by default.

LDAP packet signing does not encrypt the data in the packet, and the data can still be read by someone who captures the packet. LDAP packet signing does, however, use a digital signature to help prevent the attacker from altering the packet and putting it back on the network. Digital signatures in the packet guarantee the integrity of the data.

To use LDAP packet signing, administrative workstations must be running at least Windows 2000 SP3, and they must have the **Network Security: LDAP client signing requirements** security option set to **Require Signing** in the Admin Workstations OU GPO.

LDAP packet signing and encryption applies only to LDAP traffic between the workstations and domain controllers.

**Using Terminal Services to Perform Procedures Remotely on Domain Controllers**

On domain controllers that are running Windows Server 2003, Remote Desktop for Administration, formerly known as Terminal Services in Remote Administration mode, provides remote access to the desktops of computers that are running any operating system in the Windows Server 2003 family. Remote Desktop for Administration is disabled by default. You must enable Remote Desktop for Administration on all domain controllers that you want to manage over a remote connection.

Terminal Services in Remote Administration mode is the graphical remote administration component in the Windows 2000 Server family. The Terminal Services feature is built into each version of Windows 2000 Server, and Terminal Services in Remote Administration mode is the recommended method for all graphical remote administration tasks that are performed on

computers that are running Windows 2000 Server.

Because LDAP signing and encryption are built into Windows Server 2003 Administration Tools Pack (Adminpak.msi), remote administration by using these tools is the recommended practice for administering Active Directory remotely on domain controllers that are running Windows Server 2003. However, for all other computer management tasks that you need to perform locally on a domain controller, use Remote Desktop for Administration. For example, you can run Ntbackup remotely only by using Remote Desktop for Administration.

The client component that uses Remote Desktop for Administration is Remote Desktop Connection, formerly known as the Terminal Services client. Remote Desktop Connection is available by default on computers that are running Windows XP or Windows Server 2003 operating systems.

For more information about using Remote Desktop for Administration, see Help and Support Center for Windows Server 2003. For more information about using Terminal Services for remote administration of domain controllers that are running Windows 2000 Server, see "Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations: Part I." To download this document, click the link under "Planning & Deployment Guides" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=18591.

**Avoiding the Delegation of Security-Sensitive Operations**

There are some forest-level and domain-level operations that should not be delegated away from the trusted service administrator accounts that they are assigned to by default. Delegation of these operations can lead to an elevation-of-privilege or denial-of-service attack by the user to whom the task is delegated, which could jeopardize the entire forest or domain. Control of such operations should be retained solely within the service administrator groups. The following sections list security-sensitive operations that should not be delegated.

**Restricting the Delegation of Sensitive Forest-Level Operations**

The operations in Table 41 should be performed only by a member of the forest-level service administrator groups (Enterprise Admins or Schema Admins), and they should not be delegated to anyone else. Delegating these operations can lead to an elevation-of-privilege attack by a malicious user.

**Table 41 Forest-Level Operations That Should Not Be Delegated**

| Operation | Reason Why the Operation Should Not Be Delegated |
|---|---|
| Installing the enterprise CA | Any user with the authority to set up the CA has the authority needed to issue certificates to individuals that can result in elevated privileges. The authority to perform this action must remain with the most trusted service administrator account in your organization. |
| Modifying forest LDAP policy settings | Some LDAP policies control the performance of domain controllers. These policies are applied across the forest. Therefore, they can affect every domain controller. If improper settings are distributed, they can result in denial of service. |
| Modifying the schema | Modifications to the schema have a forest-wide impact. Improper modifications can result in the failure of existing applications or possible directory corruption. The ability to modify the schema also makes it possible to modify the default security descriptors of the classes used to protect new objects that are created in the directory. By making changes to default security descriptors, a user can gain the ability to create security principals with elevated privileges. |
| Managing forest-level operations master roles | A user with this ability has the authority to seize operations master roles. If it is performed improperly, this operation can result in a directory that is unusable. Improperly seizing the domain naming master can result in a situation in which domains with duplicate names might be created in the forest. After they are discovered, the process of repairing them might render trusts with other domains in the forest invalid, requiring portions of the forest to be rebuilt. Improperly seizing the schema master role can result in the existence of multiple schemas, which in turn can lead to islands of domain controllers that are associated with each copy of the schema. |
| Managing site topology | A user who has been delegated control over the site topology can launch a denial-of-service attack against every domain in the forest. This attack can be accomplished by breaking all inbound replication connections (that is, allowing no inbound replication). |

| | |
|---|---|
| Managing **crossRef** objects | Modifying **crossRef** objects can result in elevation of privilege. A user with this ability can create a private domain that he or she can specify as trusted. The user can then modify the SID history of an account in the private domain to introduce SIDs that represent the account as a domain administrator in the main domain, giving that account elevated privileges. |

**Restricting the Delegation of Sensitive Domain-Level Operations**

To enhance security, the operations in Table 42 should be performed only by members of the domain-level service administrator groups (Administrators, Domain Admins, Server Operators, or Backup Operators). These operations should not be delegated to anyone else.

**Table 42 Domain-Level Operations That Should Not Be Delegated**

| Operation | Reason Why the Operation Should Not Be Delegated |
|---|---|
| Installation and removal of Active Directory | The account with this authority has the ability to add and remove domains within the forest, including the ability to add and remove domain controllers. Therefore, this account can add and remove nodes of the distributed directory service and the associated security infrastructure. Access to password and account information in those domains can be used for offline password-cracking attacks. This account can also delete entire domains, which is an extreme example of a denial-of-service attack. |
| Software installation on domain controllers | The installer has local administrator access to the domain controller. With this type of access, a user can gain direct access to the directory database file, with the ability to copy it for an offline attack. The user also has the necessary permissions to install software updates and add services to the operating system. Services might include rogue agents that filter passwords or debuggers to monitor the LSA process for sensitive information. |
| Outbound trust management | Users who have the authority to create outbound trusts to external domains or to manipulate trusts with non-Windows Kerberos realms have the authority to establish trust with a dummy target forest of their creation in which they are administrators of the forest root domain, such as domain administrator or enterprise administrator. In the dummy target forest, they can manipulate their group membership in the directory to introduce group SIDs that represent a privileged entity in the source forest. Unless proper SID filtering has been enabled (which is the default setting in Windows Server 2003), creating an external trust can result in elevation of privileges in the source forest. |
| Replication management | Anyone with replication management rights has the ability to direct replication to an outside source, where data from the directory can be collected and then taken offline for a password-cracking attack. In a given domain, if the users who are delegated this right have full control over any computer object in the domain, they can inject arbitrary changes into the domain by creating a fake replication source. They also have the ability to launch denial-of-service attacks by removing replication links and preventing replication from taking place. |
| Domain-level operations master role management | The domain controller with the RID master role controls the domain RID pool as well as the allocation of RIDs to the domain controllers. Improper seizure of this role can result in duplicate RIDs being allocated. As a result, new objects will be given invalid SIDs. |
| Domain controller security policy changes | The user with this ability can control all access to the domain controller. Policy configuration is generally fairly static, and after it is initially set, changing it is an infrequent operation. There is no real need to delegate this operation. |
| Domain security policy changes | Domain security policies affect all domain users, member servers, and workstations that are members of the domain. A user with this ability can apply security policies to the entire domain. Security policies can be used to spread a weak password policy, which in turn can lead to broken passwords. Domain security policy change is an infrequent operation, and there is no real need to delegate it. |
| Backup and restore operations | Users with these rights have the ability to overwrite files on the system and use tools that copy files off the system. Having this ability allows the installation of rogue agents or the removal of files for offline password cracking. |

# Establishing Secure Data Administration Practices

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Data administrators are responsible for managing data that is stored in the directory and on computers that are members of the forest. Data administrators have no control over the configuration and delivery of the directory service itself; they control subsets of objects in the directory. Using permissions on objects that are stored in the directory, it is possible to limit the control of a given administrator account to very specific areas of the directory. Data administrators also manage computers (other than domain controllers) that are members of their domain. They manage local resources, such as print and file shares on local servers, and they also manage the group and user accounts for their own part of the organization. Data administrators can perform all of their responsibilities from management workstations, and they do not need physical access to domain controllers.

Delegation of data administration is accomplished by creating groups, granting the appropriate user rights to those groups, and applying Group Policy settings to the members of those groups. After these steps are complete, delegation is a matter of adding user accounts to the groups that are created. The critical part of this operation is granting proper access and applying the proper policies to maximize security, while still allowing administrators to perform their delegated functions.

**Delegating Data Management**

Because the creation of data administrator accounts and the specific functions that are delegated to them is driven by the needs of individual organizations, it is difficult to list specific recommendations for creating and managing individual accounts. However, this section provides a list of considerations to keep in mind when delegating control to data administrator accounts.

**Restricting Group Policy Application to Trusted Individuals**

Group Policy should be created and applied only by completely trusted individuals. These individuals should be familiar with the organization’s security policies, and they should have demonstrated their willingness to enforce those policies.

Users with accounts that have the ability to create or modify Group Policy settings can elevate the privileges of another user account through those policies. For example, if these users modify a GPO that has been applied to a group of administrative accounts, they might be able to configure a logon script that runs the next time that one of the administrators logs on. The logon script might contain a script that adds their user accounts to the Administrators group. Because the script is launched during the administrator's logon process, it runs in the security context of the administrator. In this case, it acts as if the administrator were adding the users to the group. After this process is complete, the users have administrative access.

**Taking Ownership of a Data Object**

If data administrators have the ability to create objects, make sure that you understand the scope of that ability. When a user creates an object, he or she might also become the owner of that object by virtue of being its creator, and this person is called the Creator Owner. In the discretionary access control model that is used by Windows Server 2003, the owner of an object has Full Control of that object, including the ability to change the permissions on that object. By implication, the owner of an object also has Full Control permissions on any child objects that might be added under that object. The owner has the ability to block permission inheritance from parent objects and to block service administrators’ access to that object.

If a situation arises in which access to an object has been blocked so that even service administrators cannot access it, a member of the Administrators group must take ownership of the object to regain control. Members of the Administrators group always have the right to take ownership of an object, and as the owner they can modify the security descriptor of the object.

**Reserving Ownership of Directory Partition Root Objects for Service Administrators**

Ensure that the appropriate service administrator groups in each domain own the root object for the domain directory partition. Because the owners of these directory partition root objects have the ability to change the security settings of all other objects in the partition through inheritable ACEs, it is vitally important to ensure that the directory partition root object is owned by a highly trusted administrative group. By default, Active Directory partitions are set up so that the directory partition root objects are owned by the administrative groups that are recommended in this guide. The schema directory partition is owned by the Schema Admins group, the configuration directory partition is owned by the Enterprise Admins group, and the domain directory partition is owned by the Builtin\Administrators group.

**Note**

A change to default object ownership rules in Windows Server 2003 affects default object ownership by members of the Builtin\Administrators group. This group is no longer the default owner of any forest-wide objects -- that is, schema directory partition or configuration directory partition objects. If a member of the Builtin\Administrators group creates an object in any directory partition, the individual user who creates the object has default object ownership. This change removes the ability of the Builtin\Administrators group to make forest-wide changes by modifying objects that are created in the schema directory partition and configuration directory partition.

**Preventing Concurrent Group Membership Changes**

When you plan your account management operations, institute operational practices ensuring that changes to group memberships within a delegated scope — for example, within an OU — are performed by a single data administrator within that OU, or that they are tightly coordinated between a few data administrators for that OU. By limiting group administration to a single account or to as few accounts as possible in each administrative entity (domain or OU), you greatly reduce the potential for conflicting changes to group membership. Concern about conflicting membership changes is significant in Windows 2000 as a result of the way that Active Directory in Windows 2000 replicates and handles conflicts in group membership changes between domain controllers. This concern is removed in Windows Server 2003 because group membership changes replicate differently.

**Windows 2000 Replication of Entire Group Membership**

In Windows 2000, when the membership of a group changes, the entire list of members, which is stored as a single multivalued attribute in the group object, replicates as a whole. If a conflict is detected between two concurrent changes to the group membership that originate at two separate domain controllers, one change wins (the change that occurs last) and the other change is lost.

An organization can lose changes if more than one administrator is capable of making changes to group membership. For example, imagine an organization in which a group of administrators can modify group membership. Admin1 is terminated from the organization. Admin2, in an effort to protect organization resources, immediately removes Admin1 from the Administrators group and then clicks **OK** to complete the operation. At the same time, elsewhere on the network, Admin3 modifies the membership of the Administrators group to add a new member and completes that operation at a time later than the completion of Admin2's changes, but before replication of domain changes. Replication of the group membership contains the newly added member, but it still contains Admin1 as a member. The result is that Admin1 is still a member of the Administrators group, with access to the organization's resources.

**Windows Server 2003 Replication of Separate Member Values**

When a Windows Server 2003 forest has a forest functional level of Windows Server 2003 or Windows Server 2003 interim, replication of group membership occurs at the value level, not at the attribute level. The member attribute of a group object is multivalued. Instead of the entire attribute being replicated when group membership changes, only the value that changes is replicated. For a concurrent change to the same membership to occur, two administrators have to change the membership of the same user in the same group at precisely the same time. Although the likelihood of such a concurrent change is small, limiting the number of administrators who can change group membership is still a best practice for secure data administration.

**Setting Object Ownership Quotas**

On domain controllers that are running Windows Server 2003, you can set quotas that limit the number of objects that a security principal (user, group, computer, or service) can own in a domain, configuration, or application directory partition. By default, the security principal that creates an object is the object owner, although ownership can be transferred. Active Directory quotas eliminate the ability to create unlimited numbers of objects in a directory partition, which can be used for denial-of-service attacks.

**Note**

The term "user" in this discussion of security principals indicates both the user and inetOrgPerson classes of objects.

By default, quotas are not set; therefore, there are no limits to the number of objects that any security principal can own. For each target directory partition, you can set different limits for different security principals, you can set limits that apply to all security principals, or you can do both.

You can set quotas per directory partition, except for the schema directory partition, which does not support quotas. You can set quotas to apply to security principals for a directory partition by using the command line. The commands that you can use to manage quotas are as follows:

- **dsadd quota**: Sets quotas for a directory partition.
- **dsmod quota**: Modifies existing quotas.

- **dsmod partition**: Modifies default quota settings for a partition.
- **dsquery quota**: Searches for quota assignments and quota consumption.

For more information about directory quotas and how to use the **dsadd**, **dsmod**, and **dsquery** commands to manage them, see "Directory data store" and the respective topics for each command in Help and Support Center for Windows Server 2003.

**Establishing Other Secure Practices for Delegating Administration**

When you manipulate permissions on objects in the directory, you should be mindful of certain considerations so that you can avoid problems with unpredictable or confusing access control behavior. The following sections describe these considerations.

**Avoiding Use of the Dnprotect Tool**

"Building Enterprise Active Directory Services — Notes from the Field" (Microsoft Press) suggests a utility (Dnprotect.exe) to use for securing objects in the directory. When this tool sets a bit on a directory object, the object can be changed only on a domain controller that is a member of the same domain as the object's owner. Normally, when a user attempts to access an object, the user's access token is evaluated against the access control list (ACL) on the object to determine if the requested access is allowed. If the bit has been set by Dnprotect.exe, the operating system verifies that the user requesting access is a member of the same domain as the object's owner before evaluating the permissions on the object. If the user making the change is not a member of the same domain, access is denied.

Although this tool might provide some security for directory partitions that are widely replicated in the forest across multiple domains, it also introduces an element of confusion into the security model in your environment. It is possible to create a situation in which a user is listed in an object's ACL as having access, but access is denied when the user attempts to access the object.

| **Important** |
| --- |
| Use of the Dnprotect tool is not supported in Windows Server 2003. |

**Avoiding Use of Domain Local Groups for Controlling Read Access to Global Catalog Data**

Do not use domain local groups to control Read permissions on object attributes that are replicated to the global catalog. Doing so can result in unpredictable access control behavior for searches, depending on which global catalog services the search request, as illustrated by the following example. Suppose that a local group called LocGrp1, which is defined in domain Dom1, is granted Read access to an object that replicates to the global catalog. A user who is a member of LocGrp1 initiates a search for that object. When a user initiates a search for an object in the global catalog, the domain controller locator requests and obtains the address of a global catalog server from DNS. The address of the global catalog server that is returned can be one of many global catalog servers in the forest, not necessarily a global catalog server in domain Dom1. If the user binds to a global catalog server in domain Dom1, Read access is granted to the user because the user's authorization data (that is, the list of groups of which the user is a member) that is evaluated on that global catalog server includes the local group LocGrp1. However, if the user binds to a global catalog server in a different domain (Dom2), the user's authorization data that is evaluated on that global catalog server does not include the group LocGrp1. In this case, the user is not granted access to the object. Because the user has no control over which global catalog server is selected, the results of the search are unpredictable.

# Recommendations: Establishing Secure Administrative Practices

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The following lists summarize the recommendations in this chapter for establishing secure administrative practices.

**Recommendations for Establishing Secure Service Administration Practices**
Recommendations for enhancing the security of your service administrator user and group accounts are:

- Securing Service Administrator Accounts
    - Limit the exposure of service administrator accounts.
    - Manage service administrators in a controlled OU subtree.
    - Protect service administrator accounts by hiding membership of service administrator groups.
    - Manage group membership for service administrator accounts by assigning trustworthy personnel from within the forest, assigning Schema Admins temporarily, and restricting the rights of the Backup Operators and Account Operators groups.
    - Control administrative logons by requiring smart cards and sharing logons for sensitive accounts.
- Securing Service Administrators Workstations
    - Restrict service administrator logon to administrative workstations.
    - Prohibit the use of cached credentials in unlocking administrative workstations.
    - Avoid running applications in administrative contexts.
    - Run antivirus software on administrative workstations.
    - Secure LDAP traffic between administrative workstations and domain controllers.
- Avoiding the Delegation of Security-Sensitive Operations
    - Do not delegate the forest-level operations that are described in Table 41.
    - Do not delegate the domain-level operations that are described in Table 42.

**Recommendations for Establishing Secure Data Administration Practices**
Recommendations for enhancing the security of your data administration practices are:

- Delegating Data Management
    - Only allow trusted individuals to create and apply Group Policy.
    - Understand the ramifications of Creator Owner.
    - Ensure that service administrators own directory partition root objects.

    - Avoid group membership conflicts by not delegating control of the same groups to multiple administrators.

- Setting Object Ownership Quotas
    - Set quotas on directory partitions to limit object ownership.

- Establishing Other Secure Practices for Delegating Administration
    - Avoid use of the Dnprotect tool.
    - Avoid use of domain local groups for controlling Read access to global catalog data.

# Chapter 6: Securing DNS

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Active Directory uses DNS to locate domain controllers that host various directory partitions and services. The DNS Server service facilitates replication and client access to the information that is stored in the directory partitions. Because DNS is an integral part of the architecture that is used to access Active Directory, it is important to configure DNS as securely as possible to help prevent unauthorized users from exploiting it. Whether their intent is malicious or innocent, any users who gain access to the DNS infrastructure with anything other than Read access can make changes that might result in the failure of the directory service.

The process of protecting DNS begins during deployment. Awareness of the ways in which DNS can be exploited can help drive decisions that are made during deployment. After deployment, the next step is properly delegating administrative responsibilities to implement those deployment decisions.

DNS data consists of DNS zones and records, zone configuration information, and server configuration information. Zone and server configuration information can be stored in the registry or in Active Directory. Zone data can be stored in zone files on the server or as zone objects in Active Directory. This discussion assumes that you use Active Directory–integrated DNS zones, which is a recommended best practice. When you integrate DNS zones into Active Directory, DNS can take advantage of the security capabilities of Active Directory.

Data that pertains to a single DNS zone that is hosted on the DNS server is stored in Active Directory, according to the scope that you set on zone replication in DNS:

- All domain controllers in the Active Directory domain: the MicrosoftDNS container in the domain directory partition.
- All DNS servers in the Active Directory domain: the DomainDnsZones application directory partition, which has a replica on each DNS server in the domain. Application directory partitions are available only on domain controllers running Windows Server 2003.
- All DNS servers in the Active Directory forest: the ForestDnsZones application directory partition, which has a replica on each DNS server in the forest. This option is available only if all DNS servers are running Windows Server 2003.

When you install a new domain and select Active Directory–integrated DNS, the default is to store Active Directory zone information on every DNS server in the domain, as opposed to on every domain controller in the domain. Storing DNS zone data in application directory partitions provides the following benefits:

- DNS zone data is not stored on, or replicated by, domain controllers that are not DNS servers.
- DNS zone data is not replicated to the global catalog.

Application directory partitions are new in Windows Server 2003. In Windows 2000, Active Directory–integrated DNS zone data is stored only in the domain directory partition; therefore, the data must be replicated to every domain controller in the domain.

# Deploying Secure DNS

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

In a distributed environment, such as Active Directory, multiple servers work together to make the directory available to all users on the network. Advantages of the distributed model include elimination of single points of failure, shared resources for more efficient operation, and an administrative model that can be customized to match existing organizations.

Distributed environments require an infrastructure that allows the servers and clients to locate one another in the environment. In the case of Active Directory, this infrastructure is DNS.

Unauthorized users can attempt to use DNS to affect the directory environment in a variety of ways:

- They might corrupt or block the traffic between a computer that creates a DNS query and the computer that returns the response.
- They might corrupt the data on a DNS server or change it to invalid or incorrect information.
- They might shut down or disable the DNS server, making DNS information unavailable to clients and servers in the enterprise.

These types of attacks can have the results described in Table 43.

**Table 43 Results of Attacks Made Through DNS**

| Result | Description |
|---|---|
| Clients are directed to unauthorized domain controllers. | When a client sends a DNS query that is looking for the address of a domain controller, a compromised DNS server can be instructed to return the address of an unauthorized domain controller, or the DNS traffic in the response might be usurped and spoofed by traffic that directs the client to a rogue domain controller. Then, the client application might pass on secure information to the unauthorized domain controller. |
| Clients and domain controllers cannot locate other domain controllers. | When a computer that makes a DNS query receives a response that contains nonvalid IP addresses, the clients and servers are unable to locate one another. If clients cannot locate domain controllers, they cannot access the directory. If domain controllers cannot locate other domain controllers, the directory stops functioning. |

You can increase your protection of the DNS infrastructure at two levels:

- First, take steps to protect the integrity of a DNS server's responses to client requests.
- Second, improve security for the DNS server — and therefore for the zone data that is stored on the server — so that unauthorized changes cannot be made to the zone data.

**Protecting DNS Servers**

When the integrity of the responses of a DNS server are compromised or corrupted, or when the DNS data is tampered with, clients can be misdirected to unauthorized locations without their knowledge. After the clients start communicating with these unauthorized locations, attempts can be made to gain access to information that is stored on the client computers. Spoofing and cache pollution are examples of this type of attack.

Another type of attack, the denial-of-service attack, attempts to incapacitate a DNS server to make DNS infrastructure unavailable in an enterprise. To protect your DNS servers from these types of attacks:

- Use IPsec between DNS clients and servers.

- Monitor network activity.
- Close all unused firewall ports.

**Implementing IPsec Between DNS Clients and Servers**
IPsec encrypts all traffic over a network connection. Encryption minimizes the risk that data that is sent between the DNS clients and the DNS servers can be scanned for sensitive information or tampered with by anyone attempting to collect information by monitoring traffic on the network. When IPsec is enabled, both ends of a connection are validated before communication begins. A client can be certain that the DNS server with which it is communicating is a valid server. Also, all communication over the connection is encrypted, thereby eliminating the possibility of tampering with client communication. Encryption prevents spoofing attacks, which are false responses to DNS client queries by unauthorized sources that act like a DNS server.

IPsec is not enabled by default. Many organizations choose not to enable IPsec across their entire network because it might have an impact on the performance of the network.

For more information about IPsec, see IPsec (http://go.microsoft.com/fwlink/?LinkID=93520).

**Monitoring Network Activity**
Denial-of-service attacks result in the flooding of a DNS server with so many client requests that the DNS server cannot keep up with the legitimate requests, and it stops responding to DNS queries. To help guard against denial-of-service attacks, watch for unusually high amounts of traffic, and look for anomalies such as a high volume from a single location or high volume for a single type of traffic.

**Closing All Unused Firewall Ports**
Firewalls are a common way to protect servers from outside attack. Firewalls limit the available ports that can be used to communicate between points on your network. DNS servers use User Datagram Protocol (UDP) port 53 to communicate with one another. Open port 53 on firewalls when a firewall is located between:

- DNS servers that perform zone transfers
- DNS servers that delegate zones and the server that is authoritative for the zone
- DNS clients and their DNS servers
- DNS forwarders and servers that are higher in the DNS namespace hierarchy

For more information about managing Active Directory across firewalls, see Active Directory in Networks Segmented by Firewalls (http://go.microsoft.com/fwlink/?LinkId=140089).

**Protecting DNS Data**
If an attacker can gain access to the server itself, the attacker might attempt to alter data that is stored in the DNS zone files. To help keep your zone data secure from a direct attack:

- Use secure dynamic update.
- Set quotas to limit registration of DNS resource records.
- Ensure that DNS administrators are trusted.
- Delegate administration of DNS data.
- Use the appropriate routing mechanism for your environment:
    - Conditional forwarders
    - Secondary zones
    - Root hints

- Use separate internal and external DNS namespaces.
- Disable recursion on DNS servers that do not respond to DNS clients directly and that are not configured with forwarders.

**Using Secure Dynamic Update**

Secure dynamic update is the default method of DNS update in both Windows 2000 and Windows Server 2003. Secure dynamic update prevents the type of DNS attack that involves entering invalid data into the zone files. If an attacker enters invalid data into the zone files, the attacker can:

- Add records that contain invalid IP addresses, causing misdirection of clients for denial of service or misdirection to unauthorized servers.
- Register dummy records for a denial-of-service attack. Possible goals of this type of attack are to:
    - Exhaust the server's disk space by generating a huge zone file filled with dummy records.
    - Slow zone replication by creating a huge number of entries (more than 50,000) in the zone container.

To allow only trusted, authenticated users to register records, use secure dynamic update. Using secure dynamic update guarantees that registration requests are processed only if they are sent from authenticated clients in the forest. This requirement makes it more difficult for a rogue user to launch these types of attacks because the rogue user must be a member of the Authenticated Users group or have more powerful credentials to gain access. To prevent a rogue user from registering bad records, secure dynamic update is used by default in Windows 2000 and Windows Server 2003.

**Using Quotas to Limit the Number of DNS Resource Records That Can Be Registered**

As described in "Setting Object Ownership Quotas" in the Establishing Secure Data Administration Practices section, you can use quotas in combination with security groups to manage the objects that are created in Active Directory. When you use Active Directory–integrated DNS, zone data can be stored in the domain directory partition (if you want DNS data to replicate to all domain controllers) or in application directory partitions (if you want DNS data to replicate only to domain controllers that are DNS servers).

By default, members of the Authenticated Users group have the ability to create resource records on DNS servers that are in the domain of the client computer. This ability enables all computers to dynamically update DNS zone data. A typical authenticated user registers a maximum of 10 records in DNS. To ensure that malicious users or applications do not create inappropriate resource records, you can set a default quota on the application directory partitions ForestDnsZones and DomainDnsZones (if DNS data is replicated to all domain controllers that are DNS servers for the forest or domain — available in Windows Server 2003) or on the domain directory partitions (if DNS data is replicated to all domain controllers for the domain — required in Windows 2000 Server). A default limit of 10 objects ensures that all computers can update DNS appropriately but cannot start denial-of-service attacks.

**Explicit Quotas for Domain Controllers and Other Servers**

The number of resource records that must be registered in DNS is significantly greater for domain controllers and for multihomed computers that register adapter-specific records for all of their connections. To accommodate the resource record requirements of atypical authenticated users such as domain controllers, multihomed computers, and multi-adapter servers, create special groups for these specific types of computers. For example, you might create an Enterprise Domain Controllers group that has all domain controllers as members. To ensure that this group can create the maximum DNS resource records that might be required, set a quota of 300 to 400 objects for this group on the DomainDnsZones and ForestDnsZones application directory partitions or on the domain directory partitions, depending on the replication scope.

The same rule applies to certain servers, such as VPN or Web servers, that are required to register a large number of resource records in DNS. Depending on how many adapter-specific resource records your VPN or Web server must register, you can create a group for those servers and assign the appropriate quota to ensure that the Authenticated Users default quota does not apply.

**Explicit Quotas for DHCP Servers**

To allow DHCP servers to have an essentially unlimited ability to create resource records, when clients depend on DHCP-assigned IP addresses, you can create a group that contains all DHCP servers in the domain or forest and then set an explicit quota assignment for that group on the DomainDnsZones and ForestDnsZones application directory partitions, respectively, or on the domain directory partitions (if DNS data is replicated to all domain controllers). Set a quota that is appropriate to the number of clients for which your DHCP server registers records, multiplied by the upper limit of the number of records that are

registered per computer. This way, DHCP servers can update DNS with client resource records as needed.

For more information about directory quotas and how to use the dsadd, dsmod, and dsquery commands to manage them, see "Directory data store" and the respective topics for each command in Help and Support Center for Windows Server 2003.

**Ensuring That Only Trusted Individuals Are Granted DNS Administrator Privileges**
Members of the DNS Admins group have control over the configuration of the DNS environment. This control means that they can disrupt the functionality of the directory service by using denial-of-service attacks or by misdirecting clients to rogue servers, which puts them in the same category as the service administrator accounts that are mentioned in "Chapter 5: Establishing Secure Administrative Practices" earlier in this guide. Use the same care in determining who is granted membership in the DNS Admins group. Choose only individuals who are aware of your operational policies and who have demonstrated their willingness to enforce those policies. Also, grant access only to the portion of the DNS infrastructure for which each individual is responsible.

**Delegating Administration of DNS Data**
Permissions can be set on the zone containers to limit access by users who attempt to modify them with management tools, such as the DNS snap-in or ADSI Edit. By default, Administrators, Domain Admins, Enterprise Admins, and DNS Admins have Full Control access to all components of DNS. Everyone else has Read access.

**Using Appropriate Routing Mechanisms**
DNS queries can be routed to the appropriate DNS server through the following mechanisms:

- Conditional forwarding, to specify the server to query for a given suffix
- Secondary zones, to provide DNS zone data redundancy
- Root hints, to enable any DNS client to find the DNS server for any name in the namespace by using the delegations that start at the zone for the root domain

If a firewall is deployed between two Active Directory namespaces that are serviced by different DNS servers, conditional forwarding requires that only the forwarder (the server to which queries are being forwarded) be accessible. In the case of secondary zones or delegations that might be used for the same purpose, all ports need to be accessible. In this case, queries are not recursively resolved by the server. Instead, a referral to another DNS server is returned.

If you plan to deploy secondary zones in your DNS infrastructure, consider using forwarders instead. Secondary zone data is not stored in Active Directory; it is stored in a text-based zone file, which makes the data more vulnerable to attack. It is possible to use NTFS permissions to protect the data; however, eliminating the existence of the zone file altogether removes the need to protect such data. Using forwarders accomplishes distribution of the load on the DNS servers; however, forwarders do not require zone files to be stored on every server. They simply forward DNS requests to servers that have the zone data. Using forwarders reduces the potential surface area for attacks on your DNS deployment.

The downside of using forwarders is an increase in traffic on your network and a reduction in the fault tolerance of your DNS infrastructure. Forwarders do not maintain their own copy of the database. When they receive a client query, forwarders first check their own DNS cache to see if the query has been processed before and if the result is in memory. If the result is not in memory, a query is sent to the forwarders' primary DNS server to be processed. Because forwarders must always forward new queries, rather than processing them in a local database, forwarders end up generating more traffic on the network. The reduction in fault tolerance occurs because the use of forwarders means that there are fewer copies of the database available in case a DNS server fails. When a primary or secondary DNS server fails, it means that there is one less DNS server available to process client requests. If you are using forwarders, and the primary DNS server that is used by the forwarders fails, the primary DNS server and all the forwarders become unavailable.

Root hints enable any DNS server to query the DNS servers that host the zone for the root domain. If you have an internal DNS root in your DNS infrastructure, configure the root hints of internal DNS servers to point only to the DNS servers hosting your root domain, not to the DNS servers hosting the Internet root domain. This configuration prevents your internal DNS servers from sending private information over the Internet when they resolve names.

**Using Separate Internal and External DNS Namespaces**
If your organization requires an Internet presence, create separate namespaces for use on the Internet and on your intranet. DNS was designed to distribute information in response to queries from clients. It was not designed to provide secure access to the information it distributes. For this reason, DNS environments are vulnerable to attacks that cause information disclosure.

As security becomes a higher and higher priority for network operations, features have been added to DNS to make it a more secure environment. However, these features are geared more toward preventing the registration of invalid information than preventing information disclosure. A DNS server responds to any DNS client queries that are applicable to the zone for which

the DNS server is authoritative. It does not perform any kind of security check to make sure that a valid user or client is making the query. If you use the same namespace for your Internet presence and your intranet, clients on the Internet can send queries to your DNS servers in an attempt to learn the names of resources on your intranet.

Most organizations have an internal DNS infrastructure — for name resolution of internal resources — and a separate external DNS infrastructure that gives them a presence on the public Internet. This type of DNS infrastructure, called a split namespace, consists of a set of DNS servers for the internal namespace and a separate set of DNS servers for the external namespace. A split namespace ensures that users who query for an external name cannot access or get information about an internal resource.

**Disabling Recursion**

By default, recursion is enabled for the DNS Server service, and clients typically request that the server use recursion to resolve a name when they send a query. A DNS server requires recursion only if it responds to recursive queries from DNS clients or if it is configured with a forwarder. If recursion is disabled, the DNS Server service always uses referral, regardless of the client request.

In general, DNS servers can answer queries for names outside their authoritative zones in two ways:

- **Referrals**: Servers can send referral answers, which are an immediate response to the requesting client of a list of resource records for other DNS servers that can help in resolving the queried name.

- **Recursion**: Servers can use recursion to query other servers on behalf of the requesting client, attempting to fully resolve the name. Recursive lookups continue until the server receives an authoritative answer for the queried name. The server then forwards this answer in response to the original query from the requesting client.

In most cases, disabling recursion on a DNS server is recommended when DNS clients are limited to resolving names that are managed authoritatively on a specific server. This limitation occurs when a DNS server has DNS names data only for an internal network. It also occurs when the DNS server is incapable of resolving external DNS names (such as Internet DNS names) and clients are expected to try another DNS server to resolve these names. When there is no need for a server to answer queries for data that it does not host, ensuring that a DNS server does not query other servers helps prevent attacks in cases where repeated client queries might cause a DNS server to query servers that are located outside an internal, protected network. Disabling recursion minimizes unnecessary load on the server and ensures that extraneous data from external DNS servers is not cached by the server that is providing recursion.

| **Note** |
| --- |
| If you disable recursion on the DNS server, you will not be able to configure the server with forwarders because forwarders use recursion to resolve forwarded queries. |

For more information about securing DNS, see "Checklist: Securing your DNS infrastructure" in Help and Support Center for Windows Server 2003. For more information about configuring DNS, see "Designing the Active Directory Logical Structure" in Designing and Deploying Directory and Security Services and "Deploying DNS" in Deploying Network Services of the Windows Server 2003 Deployment Kit (or see "Designing the Active Directory Logical Structure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4723 and "Deploying DNS" on the Web at http://go.microsoft.com/fwlink/?LinkId=4709).

# Non–Active Directory–Integrated DNS Security

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This guide focuses on using Active Directory–integrated zones, which is the recommended practice for DNS deployment. Although details about file-based DNS deployments (non–Active Directory–integrated deployments) are not presented here, setting NTFS file permissions on zone files is important enough to mention. DNS database files are stored in a plaintext format. Any user who can gain access to the disk drive where these files are stored can open these files with a text editor, such as Notepad, and make changes. You can use NTFS file permissions to prevent this type of access, although you should make sure that any changes do not prevent normal system access.

**Recommendations: Securing DNS**
The following sections summarize the recommendations in this chapter for increasing the security of a DNS environment that is used to support Active Directory.

**Recommendations for Deploying Secure DNS**
Recommendations for making the DNS environment more secure when you are using Active Directory–integrated zones are:-.

- Protecting DNS Servers
    - Use Active Directory–integrated DNS zones.
    - Implement IPsec between DNS clients and servers.
    - Monitor network activity.
    - Close all unused firewall ports.
- Protecting DNS Data
    - Use secure dynamic update.
    - Use quotas to limit the number of DNS resource records that can be registered.
    - Ensure that only trusted individuals are granted DNS administrator privileges.
    - Delegate administration of DNS data.
    - Use appropriate routing mechanisms.
    - Use separate internal and external DNS namespaces.
    - Disable recursion on internal DNS servers.

**Recommendations for Non–Active Directory–Integrated DNS Security**
Recommendations for making your DNS environment more secure when you are not using Active Directory–integrated zones are:

- Non–Active Directory–Integrated DNS Security
    - Set NTFS file permissions on zone files

# Appendix: Procedures

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This appendix contains the procedures that are needed to perform the security updates recommended in this guide. In some instances, a link to another document is provided, rather than a procedure.

# Enabling SID Filtering

Updated: January 23, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, Windows Server 2008 R2

On domain controllers that are running Windows Server 2003 or running Windows 2000 Server SP4 or later, SID filtering is applied by default to an outgoing, external trust to "quarantine" the trusted domain. This feature allows only SIDs from the trusted domain to be included in authorization data.

You might want to enable SID filtering if external trusts are in place that do not have SID filtering applied (that is, that were created on domain controllers running Windows 2000 Server SP3 or earlier).

**Requirements**

- Credentials: Domain Admins for the trusting domain
- Tools: Netdom.exe (Windows Server 2003 Support Tools)

| **Note** |
| --- |
| This procedure assumes that an external trust already exists between the trusting domains and the trusted domains. |

**To enable SID filtering for an outgoing, external trust**

1. Log on to a domain controller in the trusting domain using an account with Domain Admins credentials.
2. At the command line, type:

   **netdom trust** *trusting domain name* **/domain**:*trusted domain name***/userO:user_name [/passwordO:*] /Quarantine:yes**

   If you do not use the character ***** for the password, the password appears in plaintext. If you use *****, you will be prompted for a password that does not display when you type it.

# Enabling Auditing on Important Active Directory Objects

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This procedure configures auditing on actions that are performed against specific important Active Directory objects. The actions to be audited are specified in Table 23 through Table 34 earlier in this guide.

**Requirements**

- Credentials: Domain Admins
- Tools: ADSI Edit (Windows Server 2003 Support Tools)

**To set auditing on important Active Directory objects**

1. Log on to a domain controller in the root domain by using an account with Domain Admins credentials.
2. In the **Run** dialog box, type **adsiedit.msc**, and then click **OK**.
3. In the console tree, double-click the appropriate directory partition, right-click *container_object* — where *container_object* is the topmost object in the domain, configuration, schema, or application directory partition or the domain controller OU in a domain where you want to enable auditing, and then click **Properties**.

   Application directory partitions do not appear in the ADSI Edit console tree by default. If the appropriate directory partition does not appear in the console tree:

   - On the **Action** menu, click **Connect to**.
   - Under **Connection point**, click **Select or type a Distinguished Name or Naming Context**, type or select the distinguished name of the directory partition, and then click **OK**.
4. On the **Security** tab in the **Properties** dialog box, click **Advanced**.
5. On the **Auditing** tab in the **Advanced Security Settings** dialog box, click **Add**.
6. Create auditing settings to match the settings in Table 23 through Table 34 earlier in this guide.

# Securing Scripts with Script Signing

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Two alternatives exist for creating signed scripts. If you want to develop your own script host, the .NET Framework Software Development Kit (SDK) contains a set of tools for signing scripts: the File Signing Tool (Signcode.exe) and the Certificate Verification Tool (Chktrust.exe). When writing your own script host, call the Win32 API **WinVerifyTrust**. This API verifies the trust on a .vbs or .js file.

Alternatively, Windows Script Host version 5.6, which ships with Windows Server 2003, includes a signer object for creating and verifying signed scripts. The following JScript® code creates a signed file:

```
var Signer = new ActiveXObject("Scripting.Signer");
var File  = "c:\\myfile.vbs";
var Cert  = "Jane Q. Programmer";
var Store = "my";
Signer.SignFile(File, Cert, Store);
```

The following sample, in this case as Microsoft® Visual Basic®, Scripting Edition (VBScript) code, verifies the signing on a file:

```
Dim Signer, File, ShowUI, FileOK
Set Signer = CreateObject("Scripting.Signer")
File = "c:\newfile.wsf"
ShowUI = True
FileOK = Signer.VerifyFile(File, ShowUI)
If FileOK Then
    WScript.Echo File & " is trusted."
Else
    WScript.Echo File & " is NOT trusted."
End If
```

For more information about new code-signing features to protect against malicious scripts, see Windows Script Host: New Code-Signing Features Protect Against Malicious Scripts (http://go.microsoft.com/fwlink/?LinkId=140091).

For more information about using Windows Script Host, see Help and Support Center for Windows Server 2003.

# Updating the Default Domain Policy GPO and the Default Domain Controllers Policy GPO

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This procedure modifies either default domain-level Group Policy or the default Domain Controllers OU–level Group Policy.

**Requirements**

- Credentials: Domain Admins
- Tools: Domain Controller Security Policy or Domain Security Policy (Administrative Tools)

| **Important** |
| --- |
| Changes in settings to domain security policy should always be made to the Default Domain Policy GPO. Changes in settings to domain controller security policy for User Rights Assignment and Audit Policy must be made to the default GPO, rather than to a newly created GPO. |

Table 44 contains information about the policies and settings that you can use to update the default GPOs that apply to the domain and to the Domain Controllers OU.

**Table 44 Policy Settings for the Default GPOs**

| Default GPO | Where Policy Is Applied | Recommended Policy Settings |
| --- | --- | --- |
| Default Domain Policy, Account Policies node | Domain root | Password Policy<br>Account Lockout Policy<br>Kerberos Policy |
| Default Domain Controllers Policy, Local Policies node | Domain Controllers OU | Audit Policy<br>User Rights Assignment<br>Security Options<br>Event Log Policy |

**To edit security settings in the Default Domain Policy GPO or the Default Domain Controller Policy GPO**

1. Log on with Domain Admins credentials, and then open either Domain Controller Security Policy or Domain Security Policy.
2. In the console tree, under **Security Settings**, double-click **Account Policies** or **Local Policies**, depending on which node contains the settings that you want to change.
3. Click the policy whose settings you want to change, as specified in Table 44.
4. In the details pane, double-click the policy setting that you want to edit.
5. On the **Security Policy Setting** tab, select or type the recommended value, and then click **OK**.
6. Restart the computer, or run **Secedit /refreshpolicy machine_policy**, to apply the settings to the Default Domain GPO or the Default Domain Controllers GPO.

# Creating a New GPO on the Domain Controllers OU and Changing Its Precedence

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

By performing this procedure, you:

- Create a new GPO and add it to the list of Group Policy objects that are linked to the Domain Controllers OU.
- Move the new GPO to a position in the list that is above the Default Domain Controllers GPO, to ensure that its policy settings take precedence over the Default Domain Controllers GPO settings.

**Requirements**

- Credentials: Domain Admins
- Tools: Active Directory Users and Computers

**To create a new GPO on the Domain Controllers OU and increase its precedence**

1. Log on with Domain Admins credentials, and then open Active Directory Users and Computers.
2. In the console tree, right-click **Domain Controllers**, and then click **Properties**.
3. On the **Group Policy** tab, click **New**.
4. In the **Group Policy Object Links** box, type a name for the new GPO, and then click **Edit**.
5. In the policy tree under **Computer Configuration**, double-click **Windows Settings**, and then double-click **Security Settings**.
6. Double-click the set of policy settings that you want to define for the new GPO: **Local Policies** or **Event Log**.
7. Click the policy whose settings you want to add to the GPO.
8. In the details pane, double-click a setting, type or select values, and then click **OK**.
9. When you have defined all the settings that you want for the new GPO, close the **Group Policy Object Editor** window.
10. To move the new GPO to a position above the Default Domain Controllers Policy GPO, on the **Group Policy** tab, click **Up**, and then click **Close**.
11. Restart the computer, or run **Secedit /refreshpolicy machine_policy** to update the Group Policy settings for the Domain Controllers OU.

# Creating a .reg File

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

To create a scripting solution that changes a registry setting on a computer, you can use the registry editor to add or modify a registry setting and then export the setting to a .reg file. You can use the .reg file in scripts to apply the registry setting to other computers.

**Requirements**

- Credentials: Administrators
- Tools: Regedit.exe

**To export a registry setting to a .reg file**

1. Click **Start**, click **Run**, type **regedit**, and then click **OK**.
    - In the console tree, navigate to the registry entry that you want to include in the file.
    - Add or modify the entry that you want to export to the .reg file.
    - With the entry selected, on the **File** menu, click **Export**.
2. In the **Save in** box, navigate to the location for the .reg file.
3. In the **File name** box, type a name for the .reg file.
4. Under **Export range**, click **Selected branch**, and then click **Save**.

# Creating a Reserve File

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This procedure reserves disk space by creating a file on the same disk volume as the Active Directory database. The reserve file size should be the larger of 250 MB or 1 percent of the size of the logical disk volume where the Active Directory database is stored.

**Requirements**

- Credentials: Domain Admins
- Tools: Fsutil.exe (System32 tool)

**To create a reserve file on a domain controller**

1. Log on to a computer that is running Windows Server 2003 or Windows XP Professional.
2. Map a network drive to the remote administrative share for the disk volume where the Active Directory database is stored, for example, *server_Name*\C$ or D$. When prompted, provide Domain Admins credentials.
3. Open a command prompt on the computer to which you are logged on, and then change to the mapped drive.
4. Create the reserve file by typing the following command:

   **fsutil file createnew** *reserve_file_name reserve_file_size_in_bytes*
5. In Windows Explorer, grant Full Control to Administrators on the reserve file.

# Enabling Monitoring for Anonymous Active Directory Access

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

You can enable monitoring for anonymous access through audit settings on the Active Directory object CN=Server,CN=System,*domain*,DC=...*ForestRootDomain*. The audit settings ensure that a Security audit log event is generated whenever an application or service attempts to list or read domain data in Active Directory with anonymous access. The most common generators of these events are applications or services, such as Routing and Remote Access Service, that are running on Windows NT 4.0 domain controllers or on member servers in the context of Local System.

| **Note** |
| --- |
| In addition to these audit settings, the **Audit directory service access** audit policy must be configured to audit successful access to the directory. Success events for directory service access are audited by default in Windows Server 2003. |

### Requirements

- Credentials: Domain Admins
- Tools: ADSI Edit (Windows Server 2003 Support Tools)

#### To enable monitoring for anonymous Active Directory access

1. Log on to a domain controller in the root domain by using an account with Domain Admins credentials.
2. Click **Start**, click **Run**, type **adsiedit.msc**, and then click **OK**.
3. In the console tree, double-click the domain, and then click the System container.
4. In the details pane, right-click **Server**, and then click **Properties**.
5. On the **Security** tab, click **Advanced**.
6. On the **Auditing** tab, click **Add**.
7. In the **Enter the object name to select** box, type **anonymous**, and then click **OK**.
8. In the **Auditing Entry for Server** dialog box, set the auditing entries shown in Table 45.

**Table 45 Audit Settings for CN=Server,CN=System,DC=** *domain* **,DC=...** *ForestRootDomain*

| Type | Name | Access | Apply Onto |
| --- | --- | --- | --- |
| Success | Anonymous | Read All Properties | This object only |
| Success | Anonymous | Enumerate Entire SAM Domain | This object only |

# Monitoring for Anonymous Active Directory Access

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

In the previous procedure, you enabled security monitoring for anonymous access by setting auditing. The auditing is set on the objects that are accessed anonymously by applications and services from other domain controllers, member servers, or workstations in the domain.

| **Note** |
| --- |
| To complete this procedure, you need software that is capable of aggregating the Security event logs on all domain controllers into a single log. In addition, you need software that can query the Security event log, based on the criteria in this procedure. |

To make the analysis of the aggregated Security event logs easier, export the aggregated Security event logs to a database, such as Microsoft Access. Collect the event logs for about 30 days to ensure that all applications and services have attempted anonymous access, and pay special attention to applications and services that are running on Windows NT 4.0 in the security context of Local System.

**Requirements**

- Credentials: Domain Admins
- Tools: unspecified database product

**To monitor for anonymous Active Directory access**

1. Collect Security event logs for 30 days on all domain controllers.
2. Aggregate the event logs from each domain controller into a database.
3. Identify the anonymous access events in the aggregated event logs by querying the Security event log for any events with Event ID = 565 (directory service access events) and the text “anonymous” (case insensitive) in the event.

   Identify the logon IDs that generated the anonymous access for the events that meet the criteria in the previous step.

   The output of this query contains the list of the logon IDs that generated the anonymous access.
4. Identify the logon and logoff events associated with the logon IDs that you identified in the previous step, by querying the results of the events in the previous step for Event ID = 528 (logon events) or Event ID = 540 (logoff events).
5. Identify the domain controllers, member servers, or workstations where the logon and logoff events that you identified in the previous step originated.
6. Identify the applications or services that are running on the domain controllers, member servers, or workstations where the anonymous access originated.

   Examples of applications or services that require anonymous access include Routing and Remote Access Service running on Windows NT 4.0, Microsoft SQL Server running on Windows NT 4.0 (with Integrated or Mixed SQL logon), and print servers running on Windows NT 4.0.

# Renaming the Default Administrator Account

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This procedure eliminates any information that can alert attackers that this account has elevated privileges. Although an attacker would still need the password, hiding the default Administrator account by renaming it adds an additional layer of protection against password attacks seeking elevation of privilege.

**Requirements**

- Credentials: Domain Admins
- Tools: Active Directory Users and Computers

**To rename the default Administrator account**

1. Log on with Domain Admins credentials, and then open Active Directory Users and Computers.
2. In the console tree, click **Users**.
3. By default, the Administrator account is in the Users container. However, if you have already created a Service Admin subtree, the Administrator account might have been moved to the Users and Groups OU in the new subtree.
4. In the details pane, right-click **Administrator**, and then click **Rename**.
5. Type a fictitious account name, and then press ENTER.
6. In the **Rename User** dialog box, change the **Full name, First name, Last name, Display name, User logon name**, and **User logon name (pre-Windows 2000)** values to match the new account name, and then click **OK**.
7. In the details pane, right-click the renamed account object, and then click **Properties**.
8. On the **General** tab, change the **Description** to resemble other user accounts.
9. On the **Account** tab, change **Account options** if needed.

| **Note** |
| --- |
| This procedure changes only the default Administrator account's logon name and account details, which someone can see if they manage to enumerate a list of accounts on your system. This procedure does not affect the ability to use the Administrator account to boot into Directory Services Restore Mode. |

# Creating a Decoy Administrator Account

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This procedure adds an additional layer of protection when you hide the default Administrator account. An attacker planning a password attack on the Administrator account can be fooled into attacking an account with no special privileges.

**Requirements**

- Credentials: Domain Admins
- Tools: Active Directory Users and Computers

**To create a decoy Administrator account**

1. Log on with Domain Admins credentials, and then open Active Directory Users and Computers.
2. In the console tree, right-click the **Users** container, click **New**, and then click **User.**
3. Type the following information:
   - In **First name** and **User logon name**, type Administrator.
   - Type and confirm a password.

   Your new account appears in the Users container.
4. In the details pane, right-click **Administrator**, and then click **Properties**.
5. On the **General** tab, in the **Description** box, type **Built-in account for administering the computer/domain**, and then click **OK**.

# Changing the Security Descriptor on AdminSDHolder

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The security descriptor on **AdminSDHolder** serves two purposes. First, it controls access to the **AdminSDHolder** object itself. Second, it acts as the master security descriptor that is periodically applied to the service administrator accounts and their members to ensure that they remain protected.

**Requirements**

- Credentials: Domain Admins
- Tools: Active Directory Users and Computers

**To change the security descriptor on AdminSDHolder**

1. Log on with Domain Admins credentials, and then open Active Directory Users and Computers.
2. On the **View** menu, click **Advanced Features**.
3. In the console tree, click the **System** container.
4. In the details pane, right-click **AdminSDHolder**, and then click **Properties**.
5. On the **Security** tab, modify the security descriptor by changing the settings as specified in Table 39.

# Denying Logon Access to the Domain

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

To deny logon access to a domain, limit the locations where the service administrator accounts can log on by denying log on locally to members of the Enterprise Admins, Domain Admins, Server Operators, Backup Operators, and Account Operators groups at the domain level. Doing so prohibits administrators from logging on to any computers in the domain. Also, be sure to follow the procedure in the next section, "Allowing Logon Access to Administrative Workstations," for restoring logon capability to administrators so that they can log on to administrative workstations.

**Requirements**

- Credentials: Domain Admins
- Tools: Active Directory Users and Computers

| ⚠**Caution** |
| --- |
| Do not use this procedure without also following the procedure in the Allowing Logon Access to Administrative Workstations procedure. Failure to perform both procedures can result in your service administrators being unable to log on to any workstations or member servers in the domain. |

**To deny logon access at the domain level to service administrators**

1. Log on with Domain Admins credentials, and then open Active Directory Users and Computers.
2. In the console tree, right-click *domain_name*, and then click **Properties**.
3. On the **Group Policy** tab, click **Default Domain Policy**, and then click **Edit**.
4. Expand the policy tree to Computer Configuration\Windows Settings\Security Settings\Local Policies, and then click **User Rights Assignment**.
5. In the details pane, double-click **Deny logon locally**.
6. Click **Define these policy settings**, and then click **Add**.

   Add all of the service administrator accounts (Administrators, Schema Admins, Enterprise Admins, Domain Admins, Server Operators, Backup Operators, and Account Operators) to the list.

# Allowing Logon Access to Administrative Workstations

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

You can limit the locations where the service administrator accounts can log on by allowing logon locally to Enterprise Admins and Domain Admins only on administrative workstations. Administrative workstations are the computer objects in the Admin Workstations OU.

To ensure that the default **Deny log on locally** right is not applied by the Default Domain Policy GPO, you must also define this policy on the OU without adding any users or groups.

**Requirements**

- Credentials: Domain Admins
- Tools: Active Directory Users and Computers

**To allow logon access to the Admin Workstations OU and override the Deny log on locally right in the Default Domain Policy GPO**

1. Log on with Domain Admins credentials, and then open Active Directory Users and Computers.
2. In the console tree, right-click **Admin Workstations**, and then click **Properties**.
3. On the **Group Policy** tab, click **New**.
4. Type **Service Administrator Policies**, and then click **Edit**.
5. Expand the policy tree to Computer Configuration\Windows Settings\Security Settings\Local Policies, and then click **User Rights Assignment**.
6. In the details pane, double-click **Allow log on locally**.
7. Click **Define these policy settings**, and then click **Add User or Group**.
8. Add any groups to the list that you want to logon locally to the workstations, and then click **OK**.

   When possible, you should create custom groups for administrative purposes and then add members to those groups based on the administrative tasks to be performed. Once you have created those groups, you can add them to the appropriate administrative workstations and grant interactive logon permissions to those workstations. By granting interactive logon permissions to the administrative workstations, you can identify authorized users of those workstations and determine whether they are allowed to log on and access the system.
9. In the details pane, double-click **Deny log on locally**.
10. Click **Define these policy settings**, and then click **OK**.

| **Note** |
| --- |
| By default, **Deny log on locally** is not defined. If **Deny log on locally** is already defined, remove any users and groups from the list. |

# Appendix: Setting User Account Control Policy for Delegated Administrators

Updated: December 2, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

As outlined in Chapter 5: Establishing Secure Administrative Practices, establishing a delegated service administration model through the use of a controlled subtree and organization units (OUs) at the domain level is an effective starting point for setting a secure domain policy. One additional point to consider is that delegated service administrators who are able to create and set user and computer objects can own modification of attributes for those object which can have unintended consequences.

For example, without further administrative control, a delegated administrator might be able to do the following in violation of domain policy for a user account they are delegated the ability to create:

- Set the user account to not require a password.
- Undo expiration of a password on a user account where the current password has aged beyond the allowed duration, thus enabling the user to not have to change or update their password as intended.
- Enable storing of clear text password in a reversibly encrypted format.

**How User Account Policy Violations Occurred in Windows 2000 Server Prior to SP4**

For administrators to accomplish these types of violations, delegated administrators first create user and computers objects under a certain OU. This allows those administrators management over the user and computer objects they are delegated authority to create or own within the controlled OU.

By default, these administrators can modify any attribute on the user or computer objects they own, including the userAccountControl attribute. Flags in this attribute control how a user can change their password and how other aspects of password management operations work at the user object level.

By manipulating these flags, delegated administrators could permit user and computer objects attributes to be modified to allow exceptions to the domain wide password policy. For example, if on a user object the PASSWD_NOTREQD bit flag has been set as part its userAccountControl attribute, that user can change his password to blank even though domain wide password policy might state that the minimum password length is 6 characters.

**How Controlled Access Rights to User Accounts Are Enforced in Windows 2000 Server SP4 and Later**

To enable better enforcement of user password policy to remain at the domain level and limit the ability of delegated service administrators in controlled OUs, several controlled access rights (CARs) checks were added for Windows 2000 Server SP4 and later. These CAR checks are applied at the domain root and specified against the following bit flags for the userAccountControl attribute for all user objects within the domain regardless of OU subdelegation within the domain. The following list below details the specific flags of interest.

| Name | Description |
|---|---|
| PASSWD_NOTREQD | This flag determines whether a password is required or not for the user. |
| DONT_EXPIRE_PASSWORD | This flag determines whether password expiration is enforced for the user. |
| ENCRYPTED_TEXT_PWD_ALLOWED | This flag determines whether a user password is stored in reversibly encrypted form or not. |

Because the ability to reliably enforce password policy at the domain level is essential, additional corresponding security checks were added in Windows 2000 Server SP4 and later versions of Active Directory. Specifically, the following extended control access rights checks were added to assist domain administrators in controlling unintended override to user account attributes that violate or circumvent domain user and password policies. These checks are applied by Active Directory at the domain root object. Each extended control access right corresponds to one of the problems that was specified previously.

Before any attempt to update the bit flags on the userAccountControl attribute (i.e. those specified in the previous table) is

dispatched by the directory service, an access check is triggered. The access check examines whether the current user (or client) is allowed to complete the desired operation or not.

The following table gives the detail of the operations and results of how the directory service performs these access checks when certain operations are performed on the userAccountControl attribute for user objects. For more information, see How to use the userAccountControl flags to manipulate the user account (http://go.microsoft.com/fwlink/?LinkId=140090).

| **Extended Control Access Rights (in Windows 200 Server SP4 and later)** | **Operations subject to access check** | **Operation result** |
|---|---|---|
| Update Password Not Required | Any attempt to modify this flag in userAccountControl attribute on an existing user object. | Operation fails with access denied error if desired control access right is not held. Otherwise, returns success. |
| | User creation | If desired control access right is not held, user account will still be created but with PASSWD_NOTREQD bit set to 0 – off. Otherwise, user account will be created with this bit set to 1. In both case, user creation API will return SUCCESS. |
| Unexpire Password | 1. Any attempt to Unexpire Password, (Note – expiring password is not subject to this DS access check) including:<br>• Set PwdLastSet attribute to max value FFFFFFFF,FFFFFFFF through LDAP or NET API<br>• Password change or set operation, and client also sets password should remain unexpired in the same operation.<br>• Set DONT_EXPIRE_PASSWORD bit in userAccountControl attribute to 1 | Operation fails with access denied error if desired control access right is not held. Otherwise, returns success. |
| | Password change or set operation, but client does not indicate password should remain unexpired. | If desired control access is not held by client, user account password will still be changed or set, but the new password remains expired. Otherwise, newly changed or set password remains unexpired. In both cases, password change or set operation returns success. |
| Enable Per User Reversible Encrypted Password | Any attempt to toggle ENCRYPTED_TEXT_PWD_ALLOWED bit in userAccountControl flags on user and computer objects. | Operation fails with access denied error. |

**Note**

These control access rights (CARs) checks are only performed for domain users when operating in directory service mode and not for local user accounts. Computer objects are not subject to the access checks on the PASSWD_NOTREQD or DONT_EXPIRE_PASSWORD bits because such checks could cause application compatibility issues.

**Modifying the default security setting**

By default, all three CARs are granted to Authenticated Users on the domain root object. Permissions can be restricted to limit the ability to set these rights to either full administrators or delegated administrators for the domain, as described in the following procedure.

**To restrict the ability to set the CARs to domain administrators**

1. Open Active Directory Users and Computers. To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
2. On the **View** menu, click **Advanced Features**.
3. In the console tree, right-click the domain node, and then click **Properties**.

   **Where?**

   - Active Directory Users and Computers/*domain node*

4. Click the **Security** tab.
5. In **Group or user names**, select **Authenticated Users**.
6. In **Permissions for Authenticated Users**, scroll to view each of the following, and then click the **Deny** check box for each:
   - **Enable Per User Reversibly Encrypted Passwords**
   - **Unexpire Password**
   - **Update Password Not Required Bit**
7. In **Group or user names**, select **Domain Admins**.
8. In **Permissions for Domain Admins**, scroll to view the following, and then ensure that the **Allow** check box for each of the following is selected:
   - **Enable Per User Reversibly Encrypted Passwords**
   - **Unexpire Password**
   - **Update Password Not Required Bit**

| **Note** |
| --- |
| As an alternative, you may want to assign these rights to delegated administrators. To do this, perform steps 7 and 8 for the delegated administrators group instead. |

9. Click **OK**.

**Controlling OWF Password Policy**

The Security Accounts Manager (SAM) in Windows Server has an API which provides a method for changing passwords using the One Way Function (OWF) format. Once OWF password hash is sent to server through this API, SAM has no way to enforce domain wide password policy. This is because a clear text password cannot be derived from an OWF formatted password. This feature presents a possible way for users to work around domain password policy.

For Windows 2000 Server SP4 and later, a registry value was introduced to control how the OWF password change API behaves at the server when attempts to set or change passwords are received that are sent and formatted using the OWF password change API.

The details of this registry value are as follows:

**Key**

KEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa

**Value name**
SamRestrictOwfPasswordChange

**Data type**
REG_DWORD

**Allowed values**
0, 1, 2

0 - Client can change password through OWF password change API, and the new password remains unexpired.

1 - Client can change password through OWF password change API (SamrChangePasswordUser), but the password expires immediately. This is the default behavior for Windows 2000 Server SP4 and later.

2 - Client use of the OWF password change API is prohibited at the server. This API (SamrChangePasswordUser) will be totally disabled and return STATUS_ACCESS_DENIED for all clients except for LocalSystem and members of builtin administrators group. This setting can provide some additional security.

**Summary**

- None of the restrictions discussed in this section are applied to the built-in System service account or for members of the built-in Administrators group for the domain.
- If the registry value does not exist or it exists but does not contain a valid value, the default value of 1 is used.
- This registry-based security setting is an independent control. It does not interactive with the control access rights checks discussed in the previous section.
- The feature can work for either domain users or local users depending on the role of the server computer. If set on a domain controller, it applies to domain users. If set on a member server, it applies only to local user accounts on the server computer where it is set.

# Best Practices for Delegating Active Directory Administration (Windows Server 2003)

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2


**Acknowledgements**

Program Manager: Sanjay Tandon

Writer: Mary Hillman

We thank the following people for their contributions in the creation of the Active Directory Delegation Appendices and the Dsrevoke tool:

Umit Akkus, Nona Allison, Colin Brace, Raman Chikkamagalur, Arren Conner, Raju Dantuluri, Dmitry Dukat, Levon Esibov, Dmitri Gavrilov, Don Hacherl, Saif Hasan, Xin He, David Hou, Gokay Hurmali, Khushru Irani, Kamal Janardhan, Gregory Johnson, Ian Jose, Richa Kumar, Klaas Langhout, William Lees, Xiaozhong Luo, Jaeger Mitchell, Nathan Muggli, Arun Nanda, Rich Randall, Ullattil Shaji, Brett Shirley, Scott Turnbull, Andrea Weiss, Jeff Westhead, and BJ Whalen.

We thank the following people for reviewing the guide and providing valuable feedback:

Laurie Brown, John Craddock, Robert DeLuca, Christoph Felix, Eric Fleischman, Guido Grillenmeier, Mike Hickey, David Kayano, Alain Lissoir, Andreas Luther, Astrid McClean, Paul Rich, Joe Richards, and David Trulli.


**Overview**

The Active Directory® directory service is an integral component of network infrastructures that are based on the Microsoft® Windows Server™ Server 2003, Standard Edition; Windows Server™ 2003, Enterprise Edition; Windows Server™ 2003, Datacenter Edition, and Windows® 2000 Server, Windows® 2000 Advanced Server, and Windows® 2000 Datacenter Server operating systems. Successful management of Active Directory environments requires distribution of administrative responsibilities among multiple administrators according to organizational, operational, legal, and administrative requirements. Having the necessary background information, requirements, practices, and recommendations can help you delegate administration to more securely and efficiently manage Active Directory services and data.


**Abstract**

Active Directory provides an enterprise-ready, scalable, distributed directory service that allows organizations to centrally manage and share information about network resources and users, and is at the heart of distributed network security in a Windows Server–based enterprise. Active Directory thus plays a major role in accomplishing the business goals of your organization, and your ability to successfully manage Active Directory has a direct bearing on your ability to accomplish these goals.

Delegation of administration, a key capability of Active Directory, provides a means to successfully manage an Active Directory environment. This document discusses in depth the issues involved in delegating administrative responsibilities, and can help you plan for, implement, and maintain an administrative delegation model that allows secure and efficient management of Active Directory.


**Scope**

This document provides all the information required to create, implement, and maintain a security-conscious and efficient delegation model to manage your Active Directory environments. This information includes an overview of delegation, in-depth explanations of the rationale for delegation, technical descriptions of how delegation works in Active Directory, processes for creating delegation models for both service and data management, the steps needed to implement and maintain the models, and a detailed case study. Appendices to this document provide an exhaustive reference, including a comprehensive list of Active Directory administrative tasks and associated permissions required to delegate every administrative task in Active Directory.

This document does not include Active Directory deployment instructions or recommendations. For information about planning and deploying an Active Directory environment, see Designing and Deploying Directory and Security Services *of the Microsoft® Windows® Server 2003 Deployment Kit* on the Web at http://go.microsoft.com/fwlink/?LinkID=4719.

**Intended Audience**

This document is intended for Information Technology (IT) professionals who are responsible for managing an Active Directory environment. In most IT infrastructures that consist of multiple integrated components and services, the responsibility to deliver a specific component or service is typically entrusted to a component or service owner, who is responsible for the overall delivery of the component or service.

Ownership of Active Directory environments should be entrusted to two specific owners or owner groups, whose roles are typically strategic and managerial – *service owners* and *data owners*. Service owners and data owners have general, overriding responsibility for Active Directory. These usually high-ranking managers are respectively responsible for ensuring reliability and security in the delivery of the directory service and for managing the security of Active Directory content. To that end, they are responsible for delegating and distributing among their administrators responsibility for managing services and content. They do so by creating an administrative delegation model, which documents the distribution of administrative responsibilities among various administrative personnel.

Administrative responsibilities for delegating Active Directory management are divided between:

- **Service owners**, who are responsible for:
    - Planning, deployment, and long-term maintenance of the Active Directory infrastructure.
    - Ensuring that the directory continues to function reliably and at the desired level of security.
    - Ensuring that the goals established in service-level agreements are maintained.
- **Data owners**, who are responsible for maintaining the information that is stored in or protected by the Active Directory directory service, including:
    - Management of user and computer accounts.
    - Management of local resources, such as member servers and workstations and the data they store.
- **Service administrators**, who represent the operational arm of service owners and are responsible for carrying out the tasks that are required to maintain the delivery of the directory service.
- **Data administrators**, who represent the operational arm of data owners and are responsible for carrying out the tasks that are required to manage the content that is stored in or protected by Active Directory.

This document is intended for service and data owners to help them create a security conscious and efficient administrative delegation model that is tailored to the specific requirements of their organization. It is also intended for the service and data administrators who are responsible for implementing the delegation model.

To accommodate the needs of these different stakeholders, the information in this document is divided into four chapters, a case study and an extensive appendices, as follows:

- Chapter 1: Delegation of Administration Overview

  This chapter provides an overview of Active Directory management categories and stakeholders and a roadmap for successfully managing delegation of administration in Active Directory. It is targeted at all stakeholders involved in Active Directory management.

- Chapter 2: How Delegation Works in Active Directory

  This chapter takes an in-depth look at how delegation of administration actually works in Active Directory and presents all the technical aspects involved in delegation of Administration. It contains a wealth of information that will be useful for all stakeholders involved in Active Directory management.

- Chapter 3: Delegating Service Management

  This chapter presents an end-to-end perspective of Active Directory service management, and provides guidance on how to create, implement, and maintain a secure and efficient administrative delegation model for service management. It is targeted at Service Owners and Service Administrators.

- Chapter 4: Delegating Data Management

  This chapter presents an end-to-end perspective of Active Directory data management, and provides guidance on how to

create, implement, and maintain a secure and efficient administrative delegation model for data management. Though it is targeted at Data Owners and Data Administrators, Service Owners and Service Administrators will also benefit from the information in this chapter.

- Case Study: A Delegation Scenario

  The case study walks through the creation, implementation, and maintenance of an administrative delegation model for a fictitious Active Directory environment based on the recommendations presented in Chapters 3 and 4. While it is primarily targeted at Service and Data administrators, service and data owners will also benefit from the case-study.

- Best Practices for Delegating Active Directory Administration: Appendices

  The appendices contain step-by-step procedures to help you administer and maintain Active Directory.

# Chapter 1: Delegation of Administration Overview

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory serves as a central information repository in a Windows-based network infrastructure, It stores identity, authorization, and application-specific information, service-publication and discovery information, and directory service configuration data. It is an inherent part of distributed security and identity management, and plays a key role in the manageability of a Windows Server–based enterprise.

Active Directory administration involves many different operations, which must be distributed among many administrators in a large organization. Delegation of administration, a key capability provided by Active Directory, provides scaleable identity, security, service and data management across the enterprise. Through delegation of administration, organizations can increase security and efficiently manage IT resources while meeting administrative requirements. By increasing administrative efficiency and decentralizing administration, delegation reduces administrative costs and improves manageability of IT infrastructures.

Delegation of administration is the transfer of administrative responsibility for a specific administrative task from a higher authority to a lower authority. From an operational perspective, delegation of administration involves a higher-level administrator conferring upon a lower-level administrator the authority to carry out a specific administrative task. From a technical perspective, delegation of administration involves a higher-level administrator granting a controlled set of permissions to a lower-level administrator in order to carry out a specific administrative task.

# Business Case for Delegating Administration

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Organizations are moving towards a model in which different divisions, structural units, and business units share a common IT infrastructure in order to increase collaboration and reduce the cost of maintaining the IT infrastructure. The IT infrastructure of such organizations often spans multiple organizational and geographic boundaries.

Such an environment might have the following requirements:

- **Organizational structure requirements.** Part of an organization might participate in a shared infrastructure to save costs, but require the ability to operate independently from the rest of the organization.
- **Operational requirements.** An organization or application might place unique constraints on directory service configuration, availability, or security.
- **Legal requirements.** An organization might have legal requirements to operate in a specific manner, such as restricting access to certain information.
- **Administrative requirements.** Different organizations might have different administrative needs, depending on existing and planned IT administration and support models.

The first three requirements express themselves as needs for *autonomy* and *isolation*. Autonomy is the ability of the administrators of an organization to independently manage:

- All or part of service management (*service autonomy*).
- All of part of the data stored in or protected by Active Directory (*data autonomy*).

Isolation is the ability of an administrator or an organization to prevent other administrators from:

- Controlling or interfering with service management (*service isolation*).
- Controlling or viewing a subset of data in Active Directory or on member computers that are joined to Active Directory the directory (*data isolation*).

Strict service or data isolation often requires creating a separate forest or domain. Addressing network architecture design considerations for accommodating autonomy and isolation requirements is beyond the scope of this document. Instead, this document addresses the need and process for delegating administrative authority based on an organization's requirements for administration of its IT resources.

For more information about accommodating autonomy and isolation requirements, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of the Windows Server 2003 Deployment Kit (or see "Designing the Active Directory Logical Structure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4723).

For the purpose of understanding an organization's needs for delegating administrative authority, organizations can be classified in the following categories, based on their size:

- Small organizations, which typically have 25 to 50 workstations and three to five servers.
- Medium organizations, which typically have 50 to 500 workstations and four to five servers.
- Large organizations, which typically have at least 500 workstations and 50 servers.

Small and medium organizations typically have one or a few administrative groups that are responsible for managing all aspects of Active Directory. Small and medium organizations might not need to create an extensive delegation model. Large

organizations usually have a clear need to distribute and delegate administrative authority to various administrative groups, possibly delegating certain aspects of Active Directory management to centralized teams and delegating other aspects to decentralized teams. While large organizations will find the delegation capabilities of Active Directory most useful, small and medium organizations can also achieve enhanced security, increased control, more accountability, and reduced costs by implementing delegation.

# Benefits of Delegation

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By enabling efficient, security-conscious delegation and distribution of administrative responsibilities among various administrative groups that addresses the specific requirements imposed by participating business units, delegation of administration provides the means to successfully manage an Active Directory environment. Delegation of administration provides the following benefits:

- Allows distribution of administrative responsibility among multiple administrative groups, each with a defined scope of authority and a defined set of responsibilities.
- Enables decentralization of administrative authority.
- Allows the security-conscious distribution of administrative responsibility.

In addition, delegation of administration allows organizations to efficiently manage their IT infrastructures and enforce their security precautions by enabling organizations to:

- Distribute administrative responsibilities on the basis of least privilege, which ensures that the individual or group of individuals to whom the task has been delegated can perform only the tasks that are delegated, and cannot perform tasks that have not been explicitly delegated or authorized.
- Increase administrative efficiency by easily and conveniently assigning the responsibilities for managing Active Directory content and the directory service itself.
- Reduce administrative costs by facilitating shared administrative responsibility. For example, administrative responsibility for providing account support to all accounts in the organization can be easily achieved within a matter of minutes.

# Delegation at Work

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A brief example at delegation at work will help you better understand the value and the benefits of delegation that organizations can use to enhance security, decrease TCO and make Active Directory and IT resource management more tractable and efficient.

Contoso Pharmaceuticals, a fictitious company, has recently deployed Active Directory. Contoso Pharmaceuticals is a large organization that has its headquarters in Chicago, Illinois and has operations in five other locations in North America and Europe. The Active Directory infrastructure consists of a single forest, three domains, and six sites. The company has about 16,000 users, 20,000 end-user workstations and about 3000 servers spread across five physical locations. Contoso has four business units that have a presence in each of the five physical locations. These business units include:

- Research and Development
- Production
- Business Management
- IT

With Active Directory delegation, Contoso was able to delegate responsibility in the following areas:

- **Workstation management.** Contoso seamlessly and easily delegated all aspects of workstation management to local administrative groups, one for each physical location.
- **Account management.** Contoso delegated all aspects of account management to the Account Admins of each business unit regardless of the physical location of the managed users, while centrally retaining help desk operations.
- **Security-sensitive operations.** Contoso was able to grant Corporate Security personnel sufficient authority to carry out security-sensitive operations on every user account in the company, such as allowing Corporate Security personnel to disable or lock out any user account in the entire company at the click of a button.
- **Resource management.** Contoso delegated all aspects of resource management to the appropriate resource owners, enabling the resource owners to manage and retain control over their resources. This included the following:
    - A human resources portal on the intranet hosted on a small cluster of three high-performance servers running Internet Information Services (IIS). Contoso was able to delegate full control over the servers to the administrators of this application and grant them the ability to authorize access to their portal.
    - Multiple in-house applications hosted on a set of high-performance servers, with the administration of the servers entrusted to one set of administrators in the data center and the administration of each application entrusted to a separate set of administrators. Using delegation, Contoso could easily delegate to the administrative group responsible for managing the servers the ability to manage the servers while delegating to the administrative group responsible for managing each application the ability to manage their applications.
- **Self-service user accounts.** Contoso enabled self-service on user accounts, and was able to finely control specific information that users could change themselves. With delegation, Contoso could allow their users to modify their phone numbers and personal information while retaining control over the ability to modify sensitive data like the password-not-required flags on user objects. Contoso was also able to grant other stakeholders like Human Resources managers the ability to modify a user's manager and office location information.
- **Active Directory–enabled applications.** Contoso was able to delegate all aspects of Microsoft Exchange mailbox administration to its Exchange administrators, which increased productivity while achieving separation of duties. In addition, Contoso had an in-house Line of Business (LOB) application that stored data in Active Directory. Contoso

administrators were able to delegate complete control over the LOB application's data to the application's administrators, including granting the ability to control access to data by using account and resource groups.

- **Service management.** Contoso was able to distribute Active Directory service management amongst specific administrative groups based on the principle of least privilege. This increased accountability achieved a clear separation of responsibility, and decreased the number of highly-trusted administrators that Contoso needed to hire and maintain.

As demonstrated by the Contoso example, organizations can realize multiple benefits from the delegation capabilities of Active Directory. These benefits include distributing administrative responsibilities on the basis of least privilege, increasing administrative efficiency by easily and conveniently assigning the responsibilities for managing directory content and the directory service itself, and reducing administrative costs by facilitating shared administrative responsibility.

# Active Directory Management

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The objective of delegating administrative authority is to allow organizations to efficiently manage their Active Directory environments and the data stored in or protected by Active Directory in accordance with good security practices. Delegation of administration makes Active Directory management easier and allows organizations to address specific administrative needs. This section discusses the areas of Active Directory management and describes the stakeholders involved.

The administrative responsibilities of managing an Active Directory environment fall into two categories:

- **Service management.** Administrative tasks involved in providing secure and reliable delivery of the directory service.
- **Data management.** Administrative operations involved in managing the content that is stored in or protected by the directory service.

**Service Management**

Service management includes managing all aspects of the directory service that are essential to ensuring the uninterrupted delivery of the directory service across the enterprise. Service management includes, but is not limited to, the following administrative tasks:

- Adding and removing domain controllers.
- Managing and monitoring replication.
- Ensuring the proper assignment and configuration of operations master roles.
- Performing regular backups of the directory database.
- Managing domain and domain controller security policies.
- Configuring directory service parameters, such as setting the functional level of a forest or putting the directory in the special List-Object security mode.

**Data Management**

Data management includes managing the content that is stored in Active Directory, as well as content that is protected by Active Directory. Data management tasks include, but are not limited to, managing the following Active Directory content:

- User accounts, which represent the identities of people who use the network
- Computer accounts, which represent the computers that are joined to domains in the Active Directory forest
- Security groups, which are used to aggregate accounts for the purpose of authorizing access to resources
- Application-specific attributes for Active Directory-enabled and -integrated applications, such as Microsoft Exchange and Microsoft Real-Time Communication service

In addition, Active Directory data management can also facilitate the distribution and delegation of these management tasks:

- Workstation management, which includes managing all aspects of end-user workstations.
- Server management, which includes managing all aspects of all servers joined to any domain in an Active Directory forest.
- Resource management, which includes managing all aspects of services and applications hosted on member servers

joined to any domain in an Active Directory forest, possibly including the server management aspects of the servers on which the application or resource is being hosted.

**Management Stakeholders**

Active Directory plays a central role in a Windows-based IT infrastructure and is an inherent part of distributed security and identity management, touching almost every critical aspect of an organization's infrastructure. Thus, management of an Active Directory environment involves multiple stakeholders, each having specific responsibilities for managing the data, service, or security aspects of Active Directory, and each requiring the ability to perform their assigned administrative responsibilities.

For example, administrative groups who are responsible for managing user identities require the ability to perform account management. Network operators who are responsible for delivering network services that are required for Active Directory to function, such as DNS, require the ability to manage DNS servers and data. Corporate security personnel might require the ability to audit logon events, and Help Desk operators might require administrative rights to perform support operations like resetting passwords for users. Administrators of Active Directory–enabled and –integrated applications require the ability to manage application-specific data stored in the directory.

From a managerial and operational perspective, Active Directory management stakeholders primarily include service and data owners and administrators. However, because Active Directory plays a central role in a Windows Server–based IT infrastructure, it is not uncommon to have other stakeholders, including owners and administrators of other parts of the IT infrastructure who own, manage, or are responsible for aspects of the IT infrastructure related to or dependent on the organization's Active Directory environment.

**Service and Data Owners**

In most IT infrastructures that consist of multiple, integrated components and services, responsibility for delivering a specific component or service is typically entrusted to an owner. This owner is responsible for the overall delivery of the component or service, but does not actually perform the work. Rather, the role of the owner is managerial and strategic, and the responsibility for the day-to-day administrative tasks that are involved in managing Active Directory is assigned by the owner to one or more administrators.

Ownership of all Active Directory environments should be entrusted to two owner groups:

- **Service owners.** Responsible for planning and long-term maintenance of the Active Directory infrastructure, ensuring that the directory service continues to function, and ensuring that goals established in service-level agreements are maintained. While the administrative responsibilities of managing the directory service can be shared among administrative groups from different business units, service ownership is typically not shared, but is held by a centralized service owner.

- **Data owners.** Responsible for maintenance of the data that is stored in Active Directory, as well as content that is protected by Active Directory. This data includes user and computer accounts and local resources, such as member servers and workstations. Because an Active Directory environment might span multiple business units, these business units often require the ability to manage their data autonomously. It is not atypical for data ownership to be shared among multiple business units across the enterprise.

Service and data owners create an administrative model to distribute and delegate administrative responsibilities among administrators, who are responsible for performing Active Directory operations. Service owners employ service administrators to manage the directory service, and data owners employ data administrators to manage the content stored in or protected by the directory service.

**Service and Data Administrators**

The operation of Active Directory involves two types of administrators:

- **Service administrators.** Responsible for performing all administrative tasks that are involved in configuring and administering an Active Directory environment. Service administrators are highly trusted users who are employed by the service owner to manage the directory service. Service administrators implement policy decisions that are made by service owners, and handle the day-to-day tasks that are associated with maintaining the directory service and infrastructure.

- **Data administrators.** Responsible for maintaining data that is stored in Active Directory, such as user and group accounts, and application-specific data. Data administrators are employed by the data owners to manage the content that is stored in or protected by the directory service. Data administrators control subsets of objects within the directory, but have no control over the installation or configuration of the directory service.

**Other Stakeholders**

Active Directory enables directory-integrated applications to store and modify network-centric data in the directory and it also enables network services to publish globally useful information in the directory. User identities are also stored in the directory, as are security and distribution groups. Stakeholders might have business requirements to modify or have access to data that is controlled by data owners.

For example, administrators of Active Directory-enabled applications or human resources personnel might need control over application or user data. Other users might be Active Directory management stakeholders based on the applications they administer. Usually, applications that store or modify data in Active Directory have administrators who are responsible for the management of the respective application-specific data. Additionally, users of these applications might also need to read or modify application-specific data.

Still other users might have job-related reasons for performing administrative tasks in Active Directory. For example, corporate security personnel might require the ability to carry out certain security-sensitive operations, such as auditing logon events. Help desk operators might require administrative rights to perform support operations, such as resetting passwords for users.

Thus, in addition to service and data owners and administrators, other users might be considered management stakeholders based on business needs that require them to have control of or access to data stored in or protected by Active Directory.

Service and data owners (and administrators) are collectively responsible for ensuring that all stakeholders are granted sufficient access rights that are minimally required to perform the administrative tasks that are associated with their business roles in the organization.

# Creating a Successful Active Directory Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Managing a dynamic Active Directory environment typically involves many different administrative tasks that differ in scope, impact, and sensitivity. Administrative responsibility must be distributed and delegated among service and data administrators, taking into account the specific administrative needs of an organization. For example, an organization might have a specific need to delegate administrative responsibility for all account and resource management tasks to decentralized administrative groups while retaining centralized control of account support and operations by delegating responsibility for providing account support to centrally-based administrative groups.

Active Directory service and data owners are responsible for ensuring that the administrative needs of all stakeholders are addressed and might need to authorize administrative access or delegate the ability to carry out specific administrative tasks to these stakeholders, allowing the stakeholders to accomplish their assigned responsibilities based on their business role in the organization.

A structured and methodical approach towards managing an Active Directory environment can greatly enhance the security of the environment, reduce administrative costs involved in managing the environment, and successfully address the administrative needs of all stakeholders, thereby making Active Directory management more tractable seamless, efficient and secure.

A structured and methodical approach to managing an Active Directory environment generally involves the following recommended approach towards planning, creating, and implementing a successful administrative delegation model:

1. Understand all aspects of Active Directory management.

2. Understand the administrative needs of all stakeholders.

3. Create a delegation model that ensures that administrative coverage is provided for all aspects of Active Directory management and that addresses the administrative needs of all stakeholders.

4. Implement the delegation model in a secure and efficient fashion, ensuring that the administrative needs of all delegated administrators and all stakeholders are addressed while also ensuring that all administrative access has been granted on the basis of least privilege.

5. Maintain the implemented delegation model, which involves making modifications to the implemented delegation model in response to changes in administrative requirements or needs.

A particularly effective method for managing an Active Directory environment is to apply the concept of roles-based administration to creating, implementing and maintaining a secure, tractable and efficient administrative delegation model. The following sections introduce the concept of roles-based administration and provide guidance on how this concept can be applied to increase security and efficiency when managing Active Directory environments.

**A Roles-Based Approach to Active Directory Management**

A particularly effective method for creating an efficient administrative delegation model is to use administrative roles. IT infrastructures span geographical, political, and administrative spectrums. A model that allows IT departments to delegate administrative responsibilities on the basis of business functions and administrative scopes lets organizations focus on business processes as opposed to technology processes. Defining these functions and scopes in terms of administrative roles enables business-driven administrative control while affording the ability to securely scale administration.

**Administrative Roles**

A role is a collection of related administrative tasks that can be assigned to a specific set of administrators in a specific scope of influence or authority. For example, account support tasks are similar in nature and are usually assigned to a specific set of administrators for a specific collection of user accounts. By defining a role called Help Desk or Account Support Operators, you can associate with that role a set of administrative tasks typically performed by personnel in that business function. You can then create various instances of this role, each represented by a unique security group, to provide account support in different parts of the organization. For example, consider an organization that is spread across four physical locations and has a decentralized account management and support model. Having defined a role for account support, this organization can simply create four instances of this role, one per physical location, and seamlessly and easily provide account support.

A role definition consists of the following:

- A collection of one or more administrative tasks.
- The collective set of permissions required to perform this set of tasks.
- An instance of a role definition applies to a specific scope of administrative authority. A role instance consists of the following:
- A collection of one or more administrative tasks.
- The collective set of permissions required to perform this set of tasks.
- A specific scope of administrative authority.

To use administrative roles in your delegation model, you do not need to undertake an extensive process of defining roles. Microsoft has engineered role definitions for Active Directory management and recommends the use of these roles for delegating administrative responsibilities for managing Active Directory. For more information about these service and data management role definitions and associated tasks, see Chapter 3: Delegating Service Management and Chapter 4: Delegating Data Management.

**Administrative Tasks**

Management of a dynamic Active Directory environment involves a wide variety of tasks that differ in nature, scope, impact, and sensitivity. For a comprehensive list of administrative tasks involved in Active Directory service and data management, see "Appendix A: Active Directory Administration Tasks" in "Best Practices for Delegating Active Directory Administration: Appendices," which accompanies this document. This list is organized into service and data management categories. These categories are further subdivided on the basis of logical similarities between tasks. For an overview of these categories, see Chapter 3: Delegating Service Management and Chapter 4: Delegating Data Management. Each task in every category maps to an administrative role. In this way, all administrative tasks are covered by the set of Microsoft-recommended administrative roles. The sum of all categories and their related tasks provides full service and data administrative coverage for an Active Directory environment. Thus, the entire realm of Active Directory service and data management tasks can be assigned according to pre-defined management roles. You can also customize the existing roles and define new ones tailored to your organization.

A roles-based approach to Active Directory administration offers multiple benefits. It makes management more tractable and provides the ability to implement uniform administrative coverage that addresses similar administrative needs across the organization. It also provides the ability to easily and reliably delegate responsibility to, and subsequently revoke delegated responsibility from, a set of administrators. This approach eliminates the need to specify multiple sets of permissions across a large set of objects.

The following sections provide guidance on how a roles-based approach can be used to create, implement, and maintain a security-conscious and efficient Active Directory delegation model for managing an Active Directory environment.

**Understanding Active Directory Management**

A good understanding of both service management and data management is essential to creating an administrative delegation model that efficiently distributes administrative responsibility in accordance with your organization's security policies.

Service owners and service administrators are highly encouraged to gain a thorough understanding of all aspects of service management and should gain at least a good understanding of data management.

Data owners are highly encouraged to gain a complete and thorough understanding of all aspects of data management. Data administrators are encouraged to gain a good understanding of all aspects of data management relevant to their administrative role and required to carry out their assigned administrative responsibilities.

For an end-to-end perspective of all aspects of data and service management, see Chapter 3: Delegating Service Management and Chapter 4: Delegating Data Management.

**Understanding the administrative needs of stakeholders**

It is important to identify all stakeholders that have a legitimate business need to access and/or modify data stored in or protected by Active Directory. These stakeholders generally include owners and administrators of other parts of the IT infrastructure, who own, manage or are responsible for aspects of the IT infrastructure that are related to or dependent on the organization's Active Directory environment. Administrators of an Active Directory–enabled or Active Directory–integrated application or Corporate Security personnel are two examples of such stakeholders.

Understanding and documenting the access requirements of these stakeholders will help in determining the minimal set of permissions required to authorize access to or delegate administrative authority to these stakeholders, which will allow you to delegate administrative responsibility based on the principle of least privilege.

**Creating a Delegation Model**

The service and data owners are collectively responsible for the creation of an administrative model that efficiently distributes responsibilities among various administrative groups, within their organization's specific structural, operational, legal, and administrative constraints. After having gained a clear understanding of the full range of administrative tasks and responsibilities involved in managing Active Directory, they should create a delegation model taking into account specific administrative needs related to service and data management.

An Active Directory environment typically has one service owner, and one or more data owners for every participating entity. Because service owners are ultimately responsible for ensuring the secure and reliable delivery of the directory service, they are responsible for creating an administrative delegation model that distributes and delegates all service administrative responsibilities among service administrators.

Similarly, because data owners are ultimately responsible for managing entity content that is stored in or protected by Active Directory, they are responsible for creating an administrative delegation model that distributes and delegates all data administrative responsibilities that are required to manage content for their entity.

In some cases, participating entities might choose to delegate certain aspects of data management to a centralized group of administrators or to a group of administrators that has collective representation from all entities. In such cases, data owners from all participating entities must collaborate in creating the overall administrative delegation model for data management.

**Guidelines for Creating the Delegation Model**

The process of creating a delegation model essentially involves understanding the various aspects of Active Directory that need to be managed, identifying specific administrative needs of all participating business units and stakeholders, and then mapping these needs to a set of administrative roles that together provide administrative coverage for all aspects of your Active Directory environment.

Your organization might choose to define its own custom roles for Active Directory management, which could be based in part on Microsoft recommended roles, or your organization might choose to create a delegation model based purely on Microsoft recommended roles. The creation of the delegation model involves determining how many instances of each role the organization will require (based on the guidance provided in the role description) to provide administrative coverage for all aspects of your Active Directory environment, taking into account specific administrative needs.

A good delegation model has the following attributes:

- Provides coverage for all aspects of Active Directory management.
- Meets unique autonomy and isolation requirements.
- Efficiently distributes administrative responsibilities.
- Delegates administrative responsibilities in a security-conscious manner.

**Implementing the Delegation Model**

Implementation of the delegation model created by owners is entrusted to a small subset of the most highly trusted and skilled administrators. These administrators are also usually responsible for maintaining the model subsequent to its implementation. Typically, these administrators also assist the owners in creating the administrative delegation model. These administrators are responsible for rolling out the delegation model to enable all role instances for service and data management, as defined in the delegation model for service and data management.

The service management delegation model should be implemented by Enterprise Administrators. Once the service management delegation model is in place, the use of administrative accounts that are highly-privileged by default will be minimized, resulting in a much lower possibility of inadvertent damage by administrators who are logged on with higher credentials than are required for the task they are performing. In addition, having fewer service administrators will substantially lower the chance of inadvertent or malicious damage to the directory service.

After the service management delegation model is in place, service administrators should hand off data-management to a small, highly trusted subset of data-administrators. The data owners for each participating entity in the organization designate a small subset of their most highly trusted data administrators to represent the data owners. These data administrators should in turn be responsible for implementing the data management delegation model which will involve delegating administrative

responsibility for data management to lesser privileged data administrators based on administrative roles.

As part of implementing the delegation model, this small subset of data administrators is also responsible for designing an OU structure that will enable them to delegate administrative responsibility, taking into account administrative requirements and Group Policy application requirements. These data administrators design and propose an OU structure to the data owners. After obtaining approval, the data administrators implement the OU structure and implement the various data management role instances as required. This small, highly trusted subset of data administrators who implement the data management delegation model will also responsible for maintaining the implemented delegation model.

Depending on the administrative needs of your organization, delegated administrators might be given the authority to sub-delegate specific aspects of data management within their scope of administrative authority.

Implementing a delegation model requires an understanding of how delegation of administration actually works in active Directory. For more information about the details of Active Directory delegation of administration, see "Chapter Two: How Delegation Works in Active Directory" later in this document.

A well-implemented delegation model ensures that:

- Only delegated administrators can perform the assigned tasks.
- Delegated administrators can perform only the tasks that are assigned and explicitly delegated to them.

**Guidelines for Implementing the Delegation Model**

The following recommendations can be used to guide the creation of a good implementation:

- Represent every instance of every administrative role with a unique security group.
- Use security groups that represent roles for the sole purpose of delegating the roles. Do not use these security groups for other purposes. For example, do not use the security group to authorize access to other resources in the domain that are not associated with the ability to perform the administrative role.
- When delegating data management, as far as possible, delegate permissions only on OUs, which are provided specifically to facilitate delegation of administration. Delegating permissions on OUs enables them to be easily and reliably revoked.
- Unless absolutely required, do not specify permissions on individual objects within an OU.
- To delegate a role, grant permissions sufficient only to perform the set of administrative tasks assigned to the role.

For more information about how to implement a delegation model for service and data administration, see Chapter 3: Delegating Service Management and Chapter 4: Delegating Data Management.

**Maintaining the Delegation Model**

Once the service and data management delegation models have been implemented, all delegated administrators and stakeholders should be able to carry out all assigned administrative tasks.

A well-planned and well-implemented delegation model will minimize the need for much maintenance. However, over time, changes in administrative needs and workforce changes might necessitate appropriate changes to the administrative delegation model. Maintenance of the delegation model involves taking these developments into account and making appropriate changes.

For each management category, service or data administrators might need to:

- Modify and/or re-delegate existing roles. For example, you might need to add a new task to a role definition or to a specific instance of the role.
- Create and delegate customized roles. You might need to create and delegate new roles that complement or supplement existing roles but that were not included in the original delegation model.
- Remove delegation of existing roles. For example, you might need to remove a user from an administrative role, which involves simply removing the user from the administrative group. Alternatively, you might need to revoke a specific capability of the role, or the entire role, which involves removing all of the permissions that are delegated for either that

capability (task) or for the entire role.

| **Note** |
| --- |
| You can use a script or the command-line tool Dsrevoke.exe to remove all permissions for a group or user in the DACLs of all OU objects in a specified scope. For more information about using Dsrevoke.exe, see Appendix G: Active Directory Delegation Tools in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document. |

For more information about maintaining service and data delegation models, see Chapter 3: Delegating Service Management and Chapter 4: Delegating Data Management.

# Chapter 2: How Delegation Works in Active Directory

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This chapter provides an in-depth understanding of how delegation works in Active Directory. Understanding how delegation works helps you to create, implement, and maintain an efficient and security-conscious delegation model. It helps to you address the needs of your stakeholders when you delegate administrative authority to, or authorize access for, any stakeholder. Finally, a better understanding of how delegation works helps you to make better informed design decisions, taking into account the technical capabilities and issues involved in delegating administrative authority.

# Overview of Active Directory Delegation

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Through delegation of administration, responsibility for administrative tasks is transferred to the administrators who must perform the respective tasks, and to no one else. From a technical perspective, delegation of administration involves a higher-level administrator granting a controlled set of permissions to a relatively lower-level administrator to give the lower-level administrator the ability to perform a specific administrative task. In other words, a higher-level administrator *authorizes* the delegated lower-level administrator to carry out a specific administrative task.

Every administrative task involves effecting a change to the state of some data, and changing the state of data involves performing a low-level operation on the object that stores that data. As mentioned earlier in this document, there are two major categories of administrative tasks in Active Directory – data management asks and service management tasks:

- Data management tasks typically include such tasks as creating and managing user and computer accounts, security groups, and application-specific data, all of which are stored in Active Directory. In certain cases, a small subset of tasks might involve the modification of Group Policy settings to affect the configuration state of member computers. Creation of user accounts and modification of group memberships are both examples of data management tasks.

- Service management tasks include tasks that are related to the creation and maintenance of Active Directory configuration data. For example, adding a domain controller to a child domain, associating a new subnet to a site, and extending the Active Directory schema are all Active Directory service management administrative tasks that effect changes to configuration data. A majority of Active Directory configuration data is stored in Active Directory itself. However, certain aspects of Active Directory behavior can or must be configured on a domain controller. The configuration data that is associated with these tasks might be stored in the registry or file system of domain controllers.

Thus, data and service management administration tasks primarily involve effecting the change of data that is stored either in Active Directory, or in some cases on the file system or registry of Domain Controllers and other computers joined to Active Directory.

Every administrative task effects a change to the state of some data, and changing the state of data involves performing a low-level operation on some data represented by an object in the directory (or in some cases by an object on the file system or registry of a computer or Domain Controller). For example, the administrative task of resetting a user's password involves a low-level write operation on an attribute of the user object. Similarly, creating a new site involves creating a corresponding site object. Similarly, modifying the level of detail logged for security events involves modifying a registry key (data) on Domain Controllers. These low-level operations that constitute an administrative task typically involve creation, deletion, access to, modification of, or verification of data.

All data stored in Active Directory is represented by directory objects, each of which can be individually secured. Similarly, all data stored on the file system on computers and in their registries is also represented by individually securable objects. Individually securable refers to the fact that all low-level operations on data (represented by objects) can be individually authorized. In other words, an administrator can specify who can do what on every specific unit of data.

Since every Active Directory administrative task involves affecting the state of some data by means of a low-level operation, and since all data can be individually secured to allow or deny someone the ability to perform a low-level operation on that data, it follows that the ability to carry out every administrative task can be controlled by controlling the appropriate permissions that authorize the execution of the corresponding low-level operation on the target data that is involved in the administrative task.

Thus, by being able to individually authorize every low-level operation on data, administrators can in effect authorize all administrative tasks that involve low-level operations. Since all administrative tasks involve at least one low-level operation, it follows that all administrative tasks can be individually authorized. Since delegation involves a higher-level administrator *authorizing* the delegated lower-level administrator to carry out a specific administrative task, every administrative task in Active Directory can be delegated by controlling the permissions required to perform the low-level operation involved in the administrative task.

In summary:

- Delegation involves granting a controlled set of permissions to someone so that can carry out an administrative task.

- Every administrative task involves performing some low-level operation on data.
- Low-level operations on data can be (individually) authorized.
- By being able to authorize the corresponding low-level task, you can delegate a task.

The following example illustrates the crux of how delegation works in Active Directory. Let us consider a day in the life of David Hamilton, a marketing associate with Contoso Pharmaceuticals.

David Hamilton recently joined the company. When he joined the company, a user account was created for him under the Business Division OU by the Account Admins for the Business Division of the company. The creation of the user account involved a create operation for an object of class User under the parent object, which in this case happened to be the OU object for the Business Division. Since the Account Admins had Create Child permissions on this OU, they were able to create the account.

Once the account was created, David could log on to the Active Directory domain and go about performing his assigned responsibilities. One day, David forgot his password and thus could not log on. He then called the Help Desk for help. Jeff Price, a member of the Help Desk operators group, received the call and reset David's password after verifying David's identity. The reset password operation involved modifying password related attributes on David's user account object and required that the Reset Password extended right on David's user account object be granted to the individual who should be able to reset David's password. During the implementation of the delegation model, the Help Desk operators group, of which Jeff was a member, was granted this extended right on all users in the domain. Thus, Jeff was authorized to reset David's password.

Within six months of employment, David was promoted to the position of senior marketing associate and now reported to a new manager. Michelle Alexander, a Contoso employee in the Human Resources department, needed to appropriately update David's manager information as stored on David's user account in Active Directory. Modifying a user's manager information requires a low-level operation involving the modification of the Manager attribute on user account objects, and requires write-property permissions to the Manager attribute to succeed. Account Admins had granted a group called Human Resources Personnel, of which Michelle was a member, write-property permissions to modify the manager attribute on all user objects. Michelle was thus able to update David's account by appropriately changing David's manager information.

After one year with the company, David was offered a promotion and a new job in the Research and Development (RandD) Division of the company. His user account now needs to be moved to the OU for the RandD Division. Moving an object involves multiple low-level operations, including the deletion of the object from under its current parent object, the creation of a new object under the new parent object, and the modification of the Common-Name and relative distinguished name attributes on the object. (Note that technically the object is not deleted – it only seems that it is deleted and re-created.) Michael Allen, a member of the overall account management team for Contoso (which has membership in the Contoso Account Admins group) was asked to move the object between the two OUs. The Contoso Account Admins group is granted sufficient permissions on both the source and target OUs to enable its members to carry out object moves. Michael thus has the required Delete Child permissions to carry out the low-level delete operation on the Business OU and Create Child permissions to carry out the low-level create operation on the RandD OU, and additionally has write-all-properties permissions on the user object and was thus able to carry out the move operation for David's user account.

# Active Directory Administrative Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As mentioned earlier, Active Directory management can be separated into service and data management.

Data management includes such tasks as creating and managing user and computer accounts, security groups, and application-specific data, all of which are stored in Active Directory. In certain cases, a small subset of tasks might involve the modification of Group Policy settings to affect the configuration state of member computers. Creation of user accounts and modification of group memberships are both examples of data management tasks.

Data management tasks include, but are not limited to, managing the following Active Directory content:

- User accounts, which represent the identities of people who use the network.
- Computer accounts, which represent the computers that are joined to domains in the Active Directory forest.
- Security groups, which are used to aggregate accounts for the purpose of authorizing access to resources.
- Application-specific attributes for Active Directory–enabled and –integrated applications, such as Microsoft Exchange and Microsoft Office Live Communications Server 2003.

Service management tasks are tasks that are related to the creation and maintenance of Active Directory configuration data. For example, adding a domain controller to a child domain, associating a new subnet to a site, and extending the Active Directory schema are all Active Directory service management administrative tasks that effect changes to configuration data. A majority of Active Directory configuration data is stored in Active Directory itself. However, certain aspects of Active Directory behavior can or must be configured on a domain controller. The configuration data that is associated with these tasks might be stored in the registry or file system of domain controllers.

Service management includes, but is not limited to, the following administrative tasks:

- Adding and removing domain controllers.
- Managing and monitoring replication.
- Ensuring the proper assignment and configuration of operations master roles.
- Performing regular backups of the directory database.
- Configuring forest-wide Lightweight Directory Access Protocol (LDAP) settings.
- Managing domain and domain controller security policies.
- Configuring directory service parameters, such as setting the functional level of a forest or putting the directory in the special List-Object security mode.

For a comprehensive list of administrative tasks involved in Active Directory service and data management, see Appendix A: Active Directory Administrative Tasks in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document. For an overview of the categories of administrative tasks, see Chapter 3: Delegating Service Management and Chapter 4: Delegating Data Management later in this document.

# Active Directory Logical Structure and Data Storage

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Effective delegation of Active Directory management tasks requires that you understand the nature of Active Directory data and where the data is stored. Since delegation of administration involves authorizing the ability to perform low-level operations on data, understanding where the data on which these operations is stored is essential to authorizing the ability to perform low-level operations on this data. This section provides descriptions of the major logical components of Active Directory as well as descriptions and locations of the top-level domain and configuration containers.

**Forests and Domains**

An Active Directory forest represents an instance of Active Directory. A forest is a collection of one or more domains with a shared configuration and schema, represented by a single logical global catalog, and connected by a spanning tree of transitive trusts. From a data perspective, Active Directory stores data for an entire forest. "Directory" and "forest" can be considered as synonymous. Although there is a single directory, data storage is distributed among one or more domains. Servers that host the Active Directory directory service are called *domain controllers*.

**Directory Partitions**

In Active Directory, data storage is partitioned into logical segments called *directory partitions*, and each directory partition replicates its changes separately among those domain controllers in the forest that store copies (replicas) of the same directory partitions.

One specific directory partition stores forest-wide configuration information essential to the proper functioning of the forest. Another specific directory partition stores the Active Directory schema. Other directory partitions store information, such as users, groups, and OUs, that is specific to individual domains. Directory partitions that store domain information are replicated to domain controllers in that domain only. Directory partitions that store configuration and schema information are replicated to domain controllers in all domains. In this way, Active Directory provides a data repository that is logically centralized but physically distributed.

| **Note** |
| --- |
| There is a distinction between a directory partition and a database partition. The Active Directory database is not partitioned. Only the directory tree, which is the logical representation of the data that is stored on a domain controller, is partitioned |

Because all domain controllers store forest-wide configuration and schema information, a domain controller in one domain can reference a domain controller in any other domain if the information that a user requests is not stored locally. In addition, domain controllers that are global catalog servers store a full replica of one domain directory partition plus a partial replica of every other domain in the forest. Thus, a domain controller that is a global catalog server can be queried to find any object in the forest.

**Configuration Directory Partition**

There is only one configuration directory partition per forest, and it stores forest-wide configuration data that is required for the proper functioning of Active Directory as a directory service. For example, all information required to ensure the proper functioning of replication is stored in the configuration partition, which also houses information pertaining to the site topology. Information that Active Directory uses to construct the directory tree hierarchy is also stored in the configuration directory partition, as is network-wide, service-specific information that applications use to connect to instances of services in the forest. Every domain controller has one fully writeable copy of the configuration directory partition.

**Schema Directory Partition**

Similarly, there is only one schema directory partition per forest. The schema directory partition contains the definitions of all objects that can be instantiated in Active Directory. It also stores the definitions of all attributes that can be a part of objects in Active Directory. Every domain controller has one fully writeable copy of the schema directory partition, although schema updates are allowed only on the domain controller that is the schema operations master.

The schema directory partition root object contains one child object for each class of objects that can be instantiated in the Active Directory forest and contains one object for each attribute that can be part of an object in the Active Directory forest.

**Domain Directory Partitions**

Every domain is represented by a domain directory partition. The domain directory partition stores users, computers, groups, and other objects for that domain. All domain controllers that are joined to the domain share a full writeable copy of the domain directory partition. Additionally, all domain controllers in the forest that host the global catalog also host a partial read-

only copy of every other domain directory partition in the forest. For the most part, domain directory partitions store domain content – that is, user, group, and computer information. However, some domain-specific configuration data is also stored in the System container of the domain directory partition.

**Active Directory Hierarchy**
Every Active Directory object can be referenced by a unique and unambiguous name known as a *distinguished name* (also known as DN). The distinguished name identifies the complete path through the container hierarchy by which the object is reached.

Every forest has one configuration directory partition, one schema directory partition, and at least one domain directory partition that represents the forest root domain. Every additional domain that belongs to a forest adds one domain directory partition to the forest.

The three default Active Directory partitions have the following distinguished names:

- Forest-root domain directory partition: DC=ForestRootDomain
- Configuration directory partition: CN=configuration,DC=ForestRootDomain
- Schema directory partition: CN=schema,CN=configuration,DC=ForestRootDomain

For more information about the default container hierarchy of all Active Directory partitions, see Appendix I: Default Container Hierarchy for Active Directory Partitions in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

**Data Storage from a Management Perspective**
From a management perspective, it is helpful to identify where configuration data is stored in Active Directory and where domain content is stored. This information is essential in delegating administration because it helps identify where permissions need to be specified to delegate service and data administration tasks.

**Location of Active Directory Service Configuration Data**
Active Directory stores service configuration data primarily in the configuration directory partition. From a service management perspective, schema management is also a part of service management and therefore the schema, which is stored in the schema directory partition, is also a part of configuration data. Additionally, certain domain-specific configuration settings are also stored in the System container in the domain directory partition, and thus this data is also a part of configuration data. Finally, because some domain controller-specific configuration data is stored in the registry and the file system on domain controllers, this data is also a part of service configuration data.

From a management perspective, configuration data can be found in the following places:

- Configuration directory partition
- Schema directory partition
- System container in domain directory partitions
- File system and registry on domain controllers

**Location of Active Directory Content**
Active Directory content is stored in domain directory partitions. Content in directory partitions is typically stored in one of two kinds of containers – objects of a generic object class container and objects of a special-purpose object class called organizationalUnit as defined in the schema. Objects of the generic class container are ordinary containers and are also used in the Configuration partition. Objects of the class OrganizationalUnit are intended to be the primary containers for domain content such as user, computer, and group accounts.

OUs differ from regular containers in that unlike regular container objects, OUs can have Group Policy applied to them. The Group Policy distribution mechanism ensures that a Group Policy object that is applied to an OU applies directly to all users and computers in the OU and, by inheritance, to all users and computers in OUs (and in generic Active Directory containers) further down the Active Directory subtree rooted at an OU.

In addition to OUs, certain default Active Directory containers also contain domain content. These containers are created during installation and contain default administrative groups, users and other required data. The Builtin container stores built-in

groups. The Users container is the default storage area for new user accounts that are created through legacy APIs that are not Active Directory–aware. The Computers container is the default storage area for new computer objects that were originally created through legacy APIs that are not Active Directory–aware

From a management perspective, Active Directory content can be found in the following places:

- OUs in domain directory partitions
- Builtin container in directory domain partitions
- Users container in directory domain partitions
- Computers container in directory domain partitions

# Delegation and Access Control

Updated: July 13, 2009

Applies To: Microsoft Hyper-V Server 2008, Microsoft Hyper-V Server 2008 R2, Windows SBS 2003, Windows SBS 2008, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Access control is the means by which administrators can control, or delegate, the ability of other users to manipulate objects in Active Directory and also to perform actions on domain controllers and file servers. Understanding the access-control model in Active Directory is essential to being able to delegate administration. This section provides an overview of the access control model in Active Directory and describes all relevant aspects of access control that are required to delegate administrative authority.

Access control involves three components:

- The security credentials of the user attempting to access a resource
- Authorization data that protects the resource that is being accessed
- An access check that evaluates whether or not the requested access can be granted

When a user (or a process that is running on behalf of the user) attempts to perform a low-level operation on a securable object, the operation being attempted is subject to an access check. The access check takes into account the user security credentials and the authorization data on the object on which the low-level operation is being requested to determine the abilities of the user in relation to the respective object. If the access check determines that the security credentials of the user requesting the operation and the authorization data on the target object provides sufficient permissions to execute the operation, the operation succeeds. If the user has insufficient permissions to execute the operation that is being requested, the request fails.

The act of delegating Active Directory administrative responsibilities involves identifying the low-level operation that corresponds to the administrative task and the specific data on which it is being performed, and then appropriately modifying authorization settings that protect the data.

**Characteristics of the Access Control Model**
As in Windows NT 4.0, the access-control model in Windows 2000 and Windows Server 2003 has the following characteristics:

**User-based authorization**
Provides the ability of a system to identify a user and match that user to access criteria for a resource. Any resource that can be manipulated by a program or process is represented as a securable object. When a user (or a process running on behalf of the user) attempts to perform an operation on a securable object, the operation being attempted is subject to an access check and is only allowed to succeed if the user performing the operation has been authorized to perform this operation on the object in question.

**Discretionary access to securable objects**
Provides the ability to specify the users who have access to an object, and to specify the access criteria. Every securable object has an owner. The owner is the only security principal who has an inherent right to allow or deny permission to gain access to an object. The first owner of an object is usually the security principal that is associated with the thread that created the object. An object's owner can transfer ownership by giving permission to take ownership to another security principal. Additionally, any security principal who has been granted the privilege to take ownership of files or other objects on the computer can take ownership of any object on the computer.

**Inheritance of permissions**
Allows administrators to easily assign and manage permissions for a large collection of objects. This feature allows administrators to specify permissions that can be applied automatically to all objects contained within a container by having the specified permissions be inherited by all child objects in the container.

**Administrative privileges**
Allow an administrator of a computer system to control which users have the right to perform various administrative functions, or to take any action that affects system-wide resources. There are two types of administrative privileges – user rights and privileges. User rights control the various ways in which users can log on to a system. Privileges control a user's ability to perform a specific task, usually one that affects the entire computer system rather than a particular object.

**Auditing**

Enables the tracking of user or system activity by recording specific (specifiable) types of actions in security logs. Audit data can be used to detect attempts to circumvent protections on resources and to create a record of administrative actions on the computer system.

**Low-level operations**

As mentioned earlier, every administrative task involves some low-level operation on data. An understanding of the various kinds of low-level operations helps to understand the details of how an administrative task works.

The following are the low-level operations can be performed on data:

- **Create an object.** This low-level operation involves the creation of a new object in Active Directory. For example, the administrative task of creating a new user account involves a low-level operation of creating an object of class User under some parent object in Active Directory

- **View an object or all child objects of a specific object.** This low-level operation involves being able to view or see an object in Active Directory. For example, when using one or more Active Directory administration tools, viewing the contents of an OU or a domain involves performing multiple low-level operations (one for each object displayed) to view objects in Active Directory.

- **Read object attributes.** This low-level operation involves accessing and reading the values of one or more attributes on an Active Directory object. For example, any administrative task involving reading some information on a user (for example, phone number or name) involves a low-level read operation on one or more attributes (also referred to as a property) of the user object representing the user's account.

- **Modify object attributes.** This low-level operation involves accessing and modifying the values of one or more attributes on an Active Directory object. For example, any administrative task involving modifying some information on a user object (such as telephone number or office location) involves a low-level modify operation. Changing a user's password involves performing a low-level modify operation modifying the password attribute of the user object corresponding to the user's account. Similarly, modifying the association of a subnet to a specific Active Directory logical site involves a low-level operation modifying the siteobject attribute on the objects representing that specific subnet.

- **Read the security descriptor of an object.** Every securable object is protected by a security descriptor which is stored as an integral part of that object and, for Active Directory objects, is represented by the NTSecurity-Descriptor attribute. The security descriptor contains among other information, a list of access control entries specifying who can perform what low-level operations on an object. Any administrative task involving viewing or changing permissions involves the low-level operation of reading the security descriptor of an object. For example, the administrative task of determining who can modify a group's membership involves a low-level operation that reads the security-descriptor of the corresponding group object.

| **Note** |
| --- |
| An administrative task involving changing permissions also involves the low-level operation of modifying the access-control list stored as part of the security descriptor, as described in "Modify the access control list protecting the object" later in this document. |

- **Modify the access control list protecting the object.** The security descriptor contains, among other information, a list of access control entries specifying who can perform what low-level operation on an object. Any administrative task involving viewing or changing permissions involves the low-level operation of reading and modifying the access control list stored in the security descriptor of an object. For example, the administrative task of delegating another administrative task or authorizing a user access to some object involves the low-level operation of reading the security descriptor and modifying the access-control list stored in the security-descriptor of the corresponding object.

- **Modify the owner of an object.** Every securable object has an owner. The owner of an object has the inherent ability to modify the access control list of the object and to transfer the ownership. The identity of an object's owner is also stored in the security descriptor of the object. The administrative tasks of taking ownership of an object or transferring ownership of an object involve the low-level operation of modifying the owner of the object.

- **Delete an object or delete an entire subtree of objects.** This low-level operation involves the deletion of an existing object from Active Directory. For example, the administrative task of deleting an existing user account involves a low-

level operation of deleting the user class object representing the user's account from Active Directory.

**Active Directory Permissions and Other Access Rights**

Delegating an administrative task involves granting to the delegated administrator the appropriate permissions that are required to perform the low-level operations that the administrative tasks involve. Therefore, understanding the various permissions that control low-level operations is essential to delegating administration.

Active Directory permissions control who can do what on objects in Active Directory. These permissions can be separated into two broad categories – *standard permissions* and *special permissions*. Standard permissions control the standard operations that can be performed on objects in Active Directory, such as creating and deleting child objects, or reading and writing the object properties. However, certain Active Directory administrative operations require access to be controlled in a way that is not supported by the standard permissions. To facilitate these operations, Active Directory allows the standard access control mechanism to be extended through special permissions that control specific Active Directory administrative tasks.

In addition to requiring permissions that control low-level operations on Active Directory data, certain administrative tasks involve the ability to perform specific actions on domain controllers (or in some cases on member workstations). *User rights* can be specified that allow these actions. User rights are different from permissions in that they control the various ways in which a user can log on to a computer. User rights also control the user's ability to perform specific operations that typically have a system-wide effect.

**Standard Permissions**

Permissions regulate which users can gain access to an object (an Active Directory object or a file system folder or file) and in what manner.

| **Note** |
| --- |
| Active Directory applies default security settings that are designed to provide an out-of-the-box security configuration. These security settings grant specific pre-configured permissions to specific security groups that are created by default. |

The following standard permissions control the ability to perform the indicated low-level operations:

- **Create Child.** The right to create children of the object.
- **Delete Child.** The right to delete children of the object.
- **Delete.** The right to delete the object.
- **Delete Tree.** The right to delete all children of the object, regardless of the permissions on the children.
- **Read Permissions.** The right to read data from the security descriptor of the object, not including the data in the system access control list (SACL).
- **Modify Permissions.** The right to modify the discretionary access control list (DACL) in the object security descriptor.
- **Modify Owner.** The right to assume ownership of the object.
- **Read Property.** The right to read properties of the object.
- **Write Property.** The right to write properties on the object.
- **List Child.** The right to list children of the object.
- **List Object.** The right to list a particular object. If the user is not granted such a right, and the user does not have List Child set on the object parent, the object is hidden from the user.

| **Note** |
| --- |

This right is ignored if the third character of the dSHeuristics attribute of the object has a value of 0 or is not set. For more information about this right, see "Controlling Object Visibility" on the Web at http://go.microsoft.com/fwlink/?LinkID=19828.

- **Access System Security.** The right to get or set the SACL in the object security descriptor

For a detailed description of the standard Active Directory permissions, see Appendix C: Active Directory Standard Permissions in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

**Special Permissions**
In addition to the standard set of permissions that are applied in the DACLs of Active Directory objects, other permissions are available to accommodate special requirements that are not supported by the standard permissions. These permissions fall into two categories:

- Validated writes
- Extended rights

**Validated Writes**
Certain administrative operations in Active Directory require that, before writing a value to a property on an Active Directory object, the system perform value checking, or validation, beyond that which is required by the schema. Validated writes are a special type of permission that facilitates the ability to perform validation prior to modifying a property on an Active Directory object. This type of permission ensures that the value that is entered for the property conforms to required semantics, is within a legal range of values, or undergoes some other special checking that would not be performed for a simple low-level write to the property. A validated write is associated with a special permission that is distinct from the Write propertyName permission. The Write propertyName permission allows any value to be written to the property with no value checking.

**Extended Rights**
Extended rights provide the ability to perform access checks on operations that have some special significance in Active Directory and that are not covered by the standard set of access rights. For example, the User class can be granted a Send As right that can be used by a mail application to determine whether a particular user can allow another user to send mail on his or her behalf. To facilitate these special requirements, Active Directory extends the standard access control mechanism through the controlAccessRight class of objects. These objects are called *extended rights*.

Instances of the controlAccessRight class are created by the system. These objects are stored in CN=Extended-Rights,CN=Configuration,DC=*ForestRootDomain*. Rather than being administratively set, properties of these objects automatically identify the type of access in the DACL of the appropriate object.

For a list of all extended rights, see Appendix D: Active Directory Extended Rights in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

**Property Sets and Extended Rights**
Active Directory also uses the controlAccessRight class of objects to represent property sets. A property set is a subset of an object's properties. For example, the Personal-Information property set includes properties such as street address and telephone number, both of which are properties of user objects. Property sets are useful because they provide the means to specify access to a subset of properties in a single ACE.

Existing documentation about control access rights might seem to suggest that property sets are extended rights or are closely related to extended rights. However, property sets are not related to extended rights in any way other than the fact that they are represented in the system by the same class of objects.

**User Rights**
User rights are abilities to perform operations on a computer system. User rights apply on every Windows-based computer, independent of Active Directory membership. User rights are essential to service management on domain controllers.

User rights differ from permissions in that permissions control access to individual objects, whereas user rights control the various ways in which users can log on to a system. They also control specific operations that have a system-wide effect on a computer. In the case of user rights on domain controllers, the effect is on Active Directory, and thus on the entire domain.

User rights are set in the local policy on a computer. User rights are set in Domain Controller Security Policy for rights on domain controllers, and in Domain Security Policy for rights on all other computers in the domain. User rights for computers that are joined to the domain can also be specified by using Group Policy linked to OUs. User rights for all domain controllers

in a domain are specified in the Default Domain Controllers Policy Group Policy object (GPO), which is applied to the Domain Controllers OU. Thus, all domain controllers in a domain share the same set of user rights assignments.

| **Note** |
| --- |
| Although the system permits the specification of different user rights values for different domain controllers in a domain, this configuration is not supported. By design, all domain controllers in the same domain are expected to have identical user rights values |

User rights are of two types:

- **Privilege.** Specifies a user's right to perform a specific task, usually one that affects an entire computer system.
- **Logon right.** Specifies the way in which a user can log on to a system. For example, a user might have the right to log on to a system remotely, but not locally.

Administrators can assign specific rights to security group accounts or to individual user accounts. These rights authorize users to perform specific actions, such as logging on to a system interactively or backing up files and directories.

User rights define capabilities at the local level. Although user rights can apply to individual user accounts, user rights are best administered to group accounts. Doing so ensures that a user who logs on as a member of a group automatically receives the rights that are associated with that group. By assigning user rights to groups rather than to individual users, you simplify the task of managing user rights assignments.

User rights that are assigned to a group are applied to all members of the group as long as they remain members. If a user is a member of multiple groups, the user's rights are cumulative, which means that the user has more than one set of rights. The only time that rights assigned to one group might conflict with those assigned to another group is in the case of certain logon rights. In general, user rights that are assigned to one group do not conflict with the rights assigned to another group. To remove rights from a user, the administrator simply removes the user from the group.

For more information about user rights, see "Appendix C - User Rights and Privileges" in the *Microsoft Windows 2000 Security Configuration Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=19829, and "Chapter 4 - User Rights Assignment" in the *Threats and Countermeasures Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=19830.

**Components of an Access Check**

All data that is stored in Active Directory is represented by objects, each of which can be individually or collectively secured. Being secured means that when a user or a process that is running on behalf of the user attempts to perform an operation on an Active Directory object, the security system checks to determine whether the user or process has legitimate access to the requested object before granting the requested access.

The access check takes into account the following:

- Credentials of the security principal that is requesting the operation.
- Authorization data on the target resource, which is stored in the properties of the securable object.

The access check is performed by the local security subsystem of the computer on which the requested resource is stored. The access check compares the user's security credentials against the authorization data to determine the abilities of the user on the respective object. If the access check determines that the authorization data of the user or process that is requesting the operation includes sufficient permissions to execute the operation, the operation succeeds. If the user has insufficient permissions to execute the operation that is being requested, the operation fails.

**Security Principal Credentials**

Any entity that can be authenticated by the system is referred to as a *security principal*. Thus, a security principal could be a user, a computer, or a thread or process that runs in the security context of a user or a computer. A *security context* is information that describes a particular security principal's identity and capabilities on the computer. In Windows Server 2003, Windows 2000 Server, and Windows NT, all activities take place in a security context. The security subsystem uses the security context to determine what actions a process and its threads of execution can perform on objects on the computer, and the identity of the responsible security principal.

**Security Identifier**

A security principal is represented by a Security Identifier (SID), which is a unique value that identifies that security principal. A

SID is issued to every security principal when it is created. Security groups are also security principals, and therefore are uniquely identified by a SID. A user security principal can be a member of multiple security groups. Consequently, a user's credentials include the SIDs of all groups of which the user is a member.

| **Note** |
| --- |
| The exception to this rule is that the membership of a domain local group that was created in Domain A does not appear in the access token that is generated on a computer that is joined to Domain B. |

**Access Token**

A security principal's security context is represented by an *access token*. When a security principal logs on to a computer, the security subsystem on that computer authenticates the security principal. After authentication, the security subsystem creates an access token for the security principal. The access token is a data structure that contains the SID for a security principal, SIDs for the groups that the security principal belongs to (with the earlier-noted exception of domain local groups from a different domain), and a list of the security principal's privileges (also known as user rights) on the local computer. An access token is created for every security principal who logs on locally at the computer's keyboard, or remotely through a network connection.

The access token provides a security context for the security principal's actions on the computer to which the security principal is logged on. It is important to note that when a user is logged on to a workstation and requests access to data in Active Directory, a logon session is generated on the domain controller to which the user binds, and a security context is generated for this user on the domain controller. The access token that is created on the domain controller represents the credentials of the user who is accessing Active Directory. Therefore, the information in this access token, not information from the workstation to which the user is logged on, is used by the access check.

For more information about SIDs, see Security Identifiers Technical Reference (http://go.microsoft.com/fwlink/?LinkId=157928). For more information about well-known SIDs, see Appendix B: Default Active Directory Security Groups in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

**Security Groups**

Security groups are security principals that you can use to aggregate users into categories and roles for the purpose of applying access control. Security groups have scopes that define where they can be applied and membership requirements that define the selection of members they can have and the groups of which they can be members. The different types of security groups that are available and whether or not they appear in a user's access token is subject to the domain mode (Windows 2000) and domain or forest functional level (Windows Server 2003) that is in effect.

Security groups play a key role in delegation of administration. Security groups are used to represent specific instances of administrative roles and contain as members the user accounts of all delegated administrators who have been assigned to the specific administrative role represented by this security group. An understanding of the various types of security groups is thus beneficial in implementing delegation, as the various security group types differ in their scope, applicability, and uses.

There are four types of security groups, differing primarily in their application, scopes, and membership criteria:

- Local groups (also referred to as Built-in groups)
- Domain Local groups
- Global groups
- Universal group

**Group Scope**

Using security groups effectively for delegation of administration requires a strategy that makes use of the group scopes that are available in Active Directory. The scope of a group identifies the extent (domain or forest) to which the group can be applied:

- **Domain local.** A single domain; that is, domain local groups are visible for administration in only the domain in which the group is created, and thus can be assigned permissions on objects in only that domain. However, a domain local group can have users and groups from its own domain and other domains as members. For this reason, you can best use domain local groups as resource groups to collect other groups that need the same level of access to a resource. Rather than adding multiple groups to an ACL, add all the groups to one domain local group and add only the domain local group to the ACL for the target object.

| Note |
| --- |
| When specifying read access to specific attributes of, or list access to, an Active Directory domain object, do not use a domain local group to set permissions if the attribute or attributes are included in the partial attribute set that is replicated to global catalog servers. Instead, use a global group. |

- **Global.** All domains in a forest; that is, a global group is visible for administration within its domain and all trusted domains and thus can be assigned permissions on objects in all domains in the forest. Because global groups have forest-wide visibility, they are best used to organize users or groups of users into administrative roles.

- **Universal.** All domains in a forest; that is, a universal group is visible for administration within its domain and all trusted domains and thus can be assigned permissions on objects in all domains in the forest. When multiple roles require the same access to a resource, add global groups to a single universal group and add the universal group to a domain local group for the resource.

**Group Scope Availability and Membership Rules**

Group scopes vary according to functional conditions in the domain or forest. Not all scopes are available under all conditions, and group membership is subject to conditions in the domain or forest according to domain modes in Windows 2000 forests and domain functional levels in Windows Server 2003. A domain mode or functional level that indicates "mixed" provides functionality that is consistent with Windows NT 4.0 domain controllers. For this reason, features that are not compatible with Windows NT 4.0 are not available in mixed domains.

The following rules limit the availability of group scopes:

- Domain local groups are always available.

- Global groups are always available.

- Universal groups are not available in:

    - Windows 2000 mixed-mode domains.

    - Windows Server 2003 domains that have a domain functional level of Windows 2000 mixed.

The following rules limit the membership of security groups:

- **Domain local groups:**

    - Can always contain individual user accounts and global groups from any domain in the forest.

    - Can also contain universal groups from any domain in the forest as well as domain local groups from the same domain in Windows 2000 native-mode domains or Windows Server 2003 domains that have a domain functional level of Windows 2000 native or Windows Server 2003.

- **Global groups:**

    - Can always contain users from the same domain.

    - Can also contain global groups from the same domain in Windows 2000 native-mode domains or Windows Server 2003 domains that have a domain functional level of Windows 2000 native or Windows Server 2003.

- **Universal groups**, when available, can always contain user accounts, global groups, and other universal groups from any domain in the forest.

For more information about group scope and membership rules, see Help and Support Center for Windows Server 2003. For a list of the default administrative groups that are created when Active Directory is installed, and their default abilities, see

Appendix N: Default Active Directory Service Administrator Groups in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

**Local Security Groups**
Local groups can be used to grant access to local resources on a computer. As such, each computer (workstation, server, or domain controller) that is running a version of Windows has a set of default local group accounts. On domain controllers, these security groups are domain local accounts that are stored in the Builtin container. The domain controller itself does not have local computer groups. On workstations and servers, these accounts are local to the computer only.

SIDs for these default local accounts are always placed in a user's token, independent of the user's domain and the computer's domain. When a computer (workstation or server) is joined to a domain, by default the Domain Admins group is made a member of the Administrators local group on the computer.

Local groups can contain any security principal. One exception is that on domain controllers, Builtin accounts cannot contain domain local groups.

For information about the abilities that are assigned to local groups by default, see Help and Support Center for your version of Windows.

**Group Membership Changes and Replication**
Delegating administration might involve adding or removing members of large groups. Changing the membership of a large group can result in a significant volume of replication in Windows 2000 forests and in Windows Server 2003 forests that have Windows 2000 domain controllers. The reason for the increase in replication is that group membership is stored in the single, linked, multivalued Members attribute of the group object, and with Windows 2000 forests and domain controllers, the attribute is the smallest value that can be replicated. Therefore, changing the membership of a group results in replication of all group members, not just the changed member value. In addition, changes can be lost if two administrators change the membership of the same group on different domain controllers; only one of the changes can be written to the Active Directory database during the same period of replication latency.

In Windows Server 2003 forests, group replication is improved to eliminate these replication issues, but the improvements have functional level requirements. The improvements are available in Windows Server 2003 forests that have a forest functional level of:

- Windows Server 2003 (all domain controllers are running Windows Server 2003).

  - OR -

- Windows Server 2003 interim (all domain controllers are running either Windows Server 2003 or Windows NT 4.0, but no domain controllers are running Windows 2000).

At these functional levels, the discrete values of linked, multivalued attributes are replicated separately, not the entire attribute. At the Windows 2000 forest functional level in a Windows Server 2003 forest (the default level that accommodates Windows 2000 domain controllers that do not recognize certain new features), or in all Windows 2000 forests, the attribute is still the smallest unit of replication, and updating a large group (greater than 5,000 members) can still result in excessive network bandwidth consumption and lost changes. For this reason, groups that have greater than 5,000 members are not supported in Windows 2000 forests.

For more information about Windows Server 2003 functional levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in *Designing and Deploying Directory and Security Services* of the Microsoft® Windows® Server 2003 Deployment Kit (or see "Enabling Advanced Windows Server 2003 Active Directory Features" on the Web at http://go.microsoft.com/fwlink/?LinkId=6937).

**Authorization Data on Directory Objects**
All objects in Active Directory are securable by means of authorization data that protects the object. This data is stored in the Security Descriptor property of each object.

**Security Descriptors**
The security descriptor includes owner and group identifiers, as well as two types of access control lists (ACLs):

- **Owner.** The security identifier (SID) for the current owner of the object. The object owner is usually, but not always, the creator of the object.

- **Group.** The SID for the owner's primary group. This component of the security descriptor is not used by Active Directory access control, but is available for Portable Operating System Interface for UNIX (POSIX) compliance only.

- **Discretionary access control list (DACL).** A list of zero or more access control entries (ACEs) that specify who has what access to the object.

- **System access control list (SACL).** A list of zero or more access control entries that specify what audits to generate.

**Object Owners and Creators**

Every securable object has an owner. The owner is the only person who has the inherent right to control who has permission to perform operations on the object and in what way. An object's owner can grant or deny permission for different kinds of access to particular users or groups of users. An object's owner can transfer ownership by giving another security principal permission to take ownership.

| **Note** |
| --- |
| Additionally, any security principal who has the Take ownership of objects and other resources privilege on domain controllers (specifiable in the Default Domain Controller Security Policy) can take ownership of any and every Active Directory object and every file and securable resource on domain controllers in that domain. By default, the Builtin Administrators group in Active Directory or on a local computer is assigned a user right that allows this group to take ownership of all objects on the computer. |

It is important to understand how the owner of an object is assigned and how ownership can be transferred. The first owner of an object is usually the security principal that is associated with the thread that created the object. When a user makes a request to create a new object in Active Directory, the information in the user's access token is used by the security subsystem on the domain controller to establish object ownership. The information in the access token is generated by the contacted domain controller. In addition to other information, the access token contains the following information about the user:

- **User.** A field (Token-user) that contains the SID of the user who is logging on.

- **Default object owner.** A field (Default-owner-in-token) that contains the SID of the user or security group that becomes the owner of any object that this user creates.

When the access token is created, the user's SID is inserted into the Default-owner-in-token field unless the user is a member of certain administrative groups. In this case, the SID of the group is inserted into the field and any member of the administrative group can manage the object. If the user is a member of more than one such administrative group, the group with the highest level of authority is the object owner. As the object owner, you cannot change the owner to another user or group; that is, you cannot change the value in the owner field to specify an owner. However, you can apply a user right that allows another user to take ownership.

For more information about how the list of the administrative groups that become default object owners differs in Windows 2000 and Windows Server 2003, see Appendix J: Default Owners of Active Directory Objects in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

**DACL**

DACLs contain ACEs that identify the security principals that have access to the object. Each ACE has fields that specify the abilities (permissions) for the respective security principal. For each security principal that you add to the access control list, you can define a set of permissions that specify the extent to which that user or group of users can manipulate the object. If a user does not appear in an ACE, either individually or as a member of a group, that user has no access to the object.

**SACL**

The SACL is similar to the DACL except that it is used to audit rather than control access to an object. When an audited action occurs, the operating system records the event in the security log. Each ACE in a SACL has the following elements: a header that indicates whether auditing is triggered by success, failure, or both; a SID that specifies a particular user or security group to monitor; and an access mask that lists the operations to audit. The content of the SACL is controlled by security administrators for the local computer. Security administrators are users who have been assigned the Manage auditing and security log privilege. By default, this privilege is assigned to the built-in Administrators group.

**DACLs and ACEs**

For the most part, delegation involves specifying permissions on Active Directory objects in order to authorize delegated users to carry out administrative tasks. As mentioned earlier, these administrative tasks involve low-level operations each of which can be individually authorized. Authorizing a low-level operation involves identifying the object representing the data being targeted in the low-level operation and appropriately modifying permissions in the DACL of the target object's security

descriptor. This authorization is usually expressed in the form of an access control entry (ACE). An understanding of the various aspects of DACLs and ACEs is essential to being able to successfully delegate administration. This section describes DACLs and ACEs in detail.

**Where DACLs come from**

It is important to understand where the ACEs in a DACL come from. When you create an object in Active Directory programmatically, as the object's owner you can create and pass in through the API an explicit security descriptor. Permissions are specified in the Security Descriptor Definition Language (SDDL). If you create an explicit security descriptor, the ACEs in the DACL in this security descriptor are merged with any inheritable ACEs that this object would inherit from its parent object's security descriptor. If you do not specify an explicit security descriptor, a default security descriptor is applied to the new object. Every object class definition in the schema has an attribute called Default-Security-Descriptor. This default security descriptor becomes the security descriptor of the object if an explicit security descriptor is not specified during creation. In this case, the ACEs in the DACL of the default security descriptor are merged with any inheritable ACEs that this object would inherit from its parent object's security descriptor.

Over the lifetime of the object, it continues to inherit any inheritable ACEs that are present in the DACL of the parent object's security descriptor. The addition of new inheritable ACEs on the parent object results in the inheritance of these new ACEs by the object. These inherited ACEs might be effective or ineffective, depending on whether the Inherited-object type field is set and to what value. Similarly, the removal of any inheritable ACEs on the parent object or any object from which the parent inherits will result in the removal of that inheritable ACE on this object.

Inheritance enables the access control information that is defined for a container object in Active Directory to apply to any subordinate objects, including other containers and their objects. Inheritance eliminates the need to apply permissions each time a child object is created.

| **Note** |
| --- |
| If necessary, you can change the default permissions that are specified in the default security descriptor for an object class in the schema. However, as a best practice, you should avoid changing the default permissions in the default security descriptor for an object class unless you need to effect a change in the default security for all instances of that class across all domains in the forest.<br><br>As a best practice, you should also avoid removing any ACEs from the DACL of the domain root object. Doing so could impact the functionality of certain aspects of Active Directory components or applications that depend on, and have access to, Active directory data by virtue of the effective permissions in the DACL of the domain root object. |

**Default Security Descriptor Protection on Service Administrator Accounts**

To prevent the security descriptors on the key service administrator accounts in each domain from being modified and possibly becoming unusable, a background process runs on the primary domain controller (PDC) emulator that periodically checks and applies a standard security descriptor on the protected accounts. This process ensures that if a hostile user or other administrator does manage to modify the security descriptor on one of the administrative accounts, the change will be overwritten with the protected settings. This process starts 15 minutes after the system starts, and then runs once every half hour after that. This refresh interval is not configurable.

The following table lists the accounts and all their nested member groups and users that are protected by this process. The list of protected accounts varies in different versions of Windows Server.

| Windows 2000 Server<br><br>Windows 2000 Server with Service Pack 1 (SP1)<br><br>Windows 2000 Server with Service Pack 2 SP2<br><br>Windows 2000 Server with Service Pack 3 (SP3) | Windows 2000 Server with Service Pack 4 (SP4)<br><br>Windows Server 2003 | Windows Server 2003 with SP1<br><br>Windows Server 2003 with SP2 | Windows Server 2008<br><br>Windows Server 2008 R2 |
| --- | --- | --- | --- |
| Administrators | Account Operators | Account Operators | Account Operators |
| Domain Admins | Administrator | Administrator | Administrator |
| Enterprise Admins | Administrators | Administrators | Administrators |

| | | | |
|---|---|---|---|
| Schema Admins | Backup Operators | Backup Operators | Backup Operators |
| | Cert Publishers | Domain Admins | Domain Admins |
| | Domain Admins | Domain Controllers | Domain Controllers |
| | Domain Controllers | Enterprise Admins | Enterprise Admins |
| | Enterprise Admins | Krbtgt | Krbtgt |
| | Krbtgt | Print Operators | Print Operators |
| | Print Operators | Replicator | Read-only Domain Controllers |
| | Replicator | Schema Admins | Replicator |
| | Schema Admins | Server Operators | Schema Admins |
| | Server Operators | | Server Operators |

The master security descriptor for these accounts is stored as the security descriptor attribute of the AdminSDHolder object, which is located in the system container of the domain directory partition (CN=AdminSDHolder,CN=System,DC=*DomainName*).

The security descriptor on this object serves two purposes:

- It controls access to the AdminSDHolder object itself.
- It acts as the master security descriptor, which is periodically applied to the service administrator groups and their members to ensure that they remain protected.

For the default settings in the master security descriptor of the AdminSDHolder object, see Appendix K: Default Settings in the Master Security Descriptor of the AdminSDHolder Object in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

If you want to modify the security descriptor on one of the service administrator groups or on any of its member accounts, you must modify the security descriptor on the AdminSDHolder object so that it will be applied consistently. Care must be taken when making these modifications because you are also changing the default settings that will be applied to all of your protected administrative accounts. For more information about changing the security descriptor on AdminSDHolder, see Understanding AdminSDHolder and Protected Groups at http://go.microsoft.com/fwlink/?LinkId=157927 and see "Changing the Security Descriptor on AdminSDHolder" in "Appendix: Procedures" in the *Best Practice Guide for Securing Active Directory Installations: Part I* on the Web at http://go.microsoft.com/fwlink/?LinkID=21258.

**Order of ACEs in a DACL**

The DACL contains a sequential list of ACEs. The order of ACEs in the DACL affects the outcome of the access check. By default, ACLs have a preferred order called the *canonical order*. Although a DACL can be in a non-canonical order, most UI tools canonicalize the DACL during an operation on the object.

The access check algorithm checks the ACEs in the DACL in sequence. For each permission that is requested by the user, the algorithm stops at the first ACE that allows or denies the requested access. Thus, the outcome of the access check depends on the relative ordering of ACEs within the DACL.

The following is the canonical order of ACEs in a DACL:

- Explicit ACEs precede inherited ACEs
- Within a given set of explicit or inherited ACEs, Deny ACEs precede Allow ACEs

This effect can be summarized by the following canonicalized DACL:

```
Explicit Deny Ace 1
…
Explicit Deny Ace n
```

```
Explicit Allow Ace 1
…
Explicit Allow Ace n
Inherited Deny Ace 1
…
Inherited Deny Ace n
Inherited Allow Ace 1
…
Inherited Allow Ace n
```

Further, within the set of Deny inherited ACEs and Allow inherited ACEs, inheritable ACEs that are specified on the parent object precede inheritable ACEs that are specified on the parent object's parent object, and so on upward until reaching the ACEs that are specified on the domain root object. Thus, in the event of a conflict of permissions, permissions specified in an inheritable ACE set on the parent object will override the permissions specified in an inheritable ACE set on an object higher up in the object hierarchy.

A security descriptor might contain no DACL if a security descriptor that is passed in programmatically has no DACL specified, and no ACEs are inherited from the parent object. A security descriptor can also have a DACL that contains no ACEs.

| **Note** |
| --- |
| A DACL can be set to null programmatically (that is, no DACL is passed) at creation time or at any later time. |

The effect on security of having no DACL as opposed to having a DACL that has no ACEs is significant:

- No DACL in the security descriptor has the effect of granting unconditional access to Everyone.
- An empty DACL has the effect of not granting access to anyone.

In the case of an empty DACL, the owner of the object always has the inherent right to modify the DACL and grant permissions. If no DACL exists, the owner can create a DACL programmatically and then grant permissions.

For more information about using SDDL to specify permissions programmatically, see the Microsoft Platform Software Development Kit (SDK) link on the Web Resources page at http://go.microsoft.com/fwlink/?LinkId=95521.

**Protected DACLs**
The DACL of an object can be marked as protected. Marking a DACL as protected will block the inheritance of any permissions from the parent object. Any inheritable permissions in the DACL of the parent objects will not be inherited by this object. Additionally, child objects of an object whose DACL is marked protected will also no longer inherit any permissions specified on the protected object's parent. However, they will continue to inherit all inheritable permissions specified in the DACL of the protected object unless the child objects also have their DACLs marked as protected.

**ACEs**
An ACE contains authorization intent for a specific security principal for whom access is allowed, denied, or audited.

Every ACE in the DACL has the following structure:

```
(ACE Type ; ACE Flags ; Permissions ; Object Type ; Inherited Object Type ; Trustee)
```

The following is a brief description of the various fields of an ACE:

- **ACE Type –** This field specifies whether the ACE is an Allow type ACE or a DENY type ACE. An ALLOW type ACE allows access and a DENY type denies access.

  This field can take on one of six values:

  - **Allow –** Indicates that the ACE is of the standard ACCESS ALLOWED type, where the ObjectType and InheritedObjectType fields are NULL.
  - **Deny –** Indicates that the ACE is of the standard system-audit type, where the ObjectType and InheritedObjectType fields are NULL.
  - **System Audit –** Indicates that the ACE is of the standard system type, where the ObjectType and InheritedObjectType fields are NULL.

    - **Object Allow –** Indicates that the ACE grants access to an object or a subobject of the object, such as a property set or property. The ObjectType or InheritedObjectType field or both contain a GUID that identifies a property set, property, extended right, or type of child object.

    - **Object Deny –** Indicates that the ACE denies access to an object or a subobject of the object, such as a property set or property. The ObjectType or InheritedObjectType field or both contain a GUID that identifies a property set, property, extended right, or type of child object.

    - **Object System Audit –** Indicates that the ACE audits access to an object or a subobject of the object, such as a property set or property. The ObjectType or InheritedObjectType field or both contain a GUID that identifies a property set, property, extended right, or type of child object.

- **ACE Flags –** This field contains multiple inheritance related flags that govern the inheritance aspects of the ACE.

  This field can take values including but not limited to one or more of the following:

    - **Container Inherit –** Child objects will inherit this access-control entry (ACE). The inherited ACE is inheritable unless the ADS_ACEFLAG_NO_PROPAGATE_INHERIT_ACE flag is set.

    - **No Propagate –** The system will clear the ADS_ACEFLAG_INHERIT_ACE flag for the inherited ACEs of child objects. This prevents the ACE from being inherited by subsequent generations of objects.

    - **Inherit Only –** Indicates an inherit-only ACE that does not exercise access control on the object to which it is attached. If this flag is not set, the ACE is an effective ACE that exerts access control on the object to which it is attached.

    - **Inherited –** Indicates whether or not the ACE was inherited. The system sets this bit.

- **Permissions –** This field contains the specification of the permissions granted or denied to the trustee (the security principal for whom the permissions specified in this ACE will apply). This field can take one of the following values:

    - RC – Read Control

    - SD – Standard Delete

    - WD – Write DACL

    - WO – Write Owner

    - RP – Read Property

    - WP – Write Property

    - CC – Create Child

    - DC – Delete Child

    - LC – List Child

    - LO – List Object

    - SW – Validated write

    - DT – Delete Tree

- CR – Extended Right

For more information about these rights, see Appendix C: Active Directory Standard Permissions in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

- **Object Type –** This field can contain the GUID representing an object's class. The GUID refers to a property or a property-set if the permissions specified are Read-Property or Write-Property permissions. The GUID specifies an object class when the permissions specified are Create Child or Delete Child permissions

- **Inherited Object Type –** This field can contain the GUID of an object class. When such a GUID is set, the ACE is an effective ACE only on objects of the class referred to by the GUID.

- **Trustee –** This field contains the SID of the security principal for whom the permissions specified in this ACE will apply.

For a detailed description of the various fields of an ACE, see "IADsAccessControlEntry" in the MSDN Library on the Web at http://go.microsoft.com/fwlink/?LinkID=21262.

**Example ACEs**
The following example ACEs illustrate how ACEs specify delegation intent:

| **Note** |
| --- |
| To provide clarity in following examples, the display name of the class or attribute has been shown instead of the actual GUID for the class, and the name of the group has been used instead of the SID for Account Admins. In actual ACEs, GUIDs and SIDs are used instead of the display names shown here. |

- An ACE that grants Account Admins the ability to create user objects in an OU and in all OUs within the subtree rooted at this OU:

  - (OA ; CI ; CC ; User Class ; Organizational Unit class ; Account Admins)

- An ACE that grants Account Admins the ability to reset user passwords in an OU but only on User objects:

  - (OA ; CI ; CR; Reset Password ; User class ; Account Admins)

- An ACE that grants Account Admins the ability to modify the membership of all group objects (but only group objects) in a specific OU (on which this ACE is being applied). This ACE does not give Account Admins the permission to modify the group memberships of any groups that might exist as child objects of other OUs that might be child objects of the OU on which the ACE is being applied.

  - (OA ; ; WP; Member ; Group class ; Account Admins)

    Note the absence of any inheritance flags in this ACE.

**Inheritance and Organizational Units**
Active Directory data is stored in containers. While some containers are general purpose containers, other containers have a specialized purpose in that they are specific in the nature of data that they are intended to store. Generic containers are represented by objects of the generic class Container, while specific containers are represented by specific object classes. An example of a generic container would be the Services container in the Configuration partition, which is intended to contain various kinds of objects representing service specific information. An example of a specific class of container is the Sites container in the Configuration partition which is intended to contain site and subnet related information and is represented by an object of class SitesContainer. While different kinds of configuration information can be stored in different kinds of containers, domain data is typically stored in a special purpose class of objects called Organizational Units (OUs).

OUs are special purpose containers in that they are intended to contain domain data like user accounts, computer accounts, and groups. Additionally, OUs can contain other OUs within them. In this manner, domain data is stored within a hierarchy of OUs which serve the role of containers for domain data. OUs are different from other containers in that Group Policy can be

applied to them. The application of Group Policy to an OU results in the inheritance of the applied Group Policy by user and computer objects, on which that Group Policy becomes effective subject to the precedence order of the application of Group Policy.

OUs are used to form a hierarchy of special purpose containers within a domain for the purpose of storing user and computer accounts and security groups and applying Group Policy onto a collection of user and computer accounts. Organizations create an OU hierarchy structure based on their Group Policy requirements and collectively store user and computer accounts and groups within specific OUs depending on the Group Policy requirements for these user and computer accounts.

Organizational units also play a major role in delegation of administration. Delegation of administration involves granting an administrative group sufficient permissions to allow the group the ability to perform specific administrative tasks. Usually, administrative authority is delegated to manage a set of users. While it is possible to grant an administrative group permissions on the specific user accounts that belong to a set of users, individually modifying the DACL of every single user account in a set is usually unmanageable and impractical. To simplify the application of the permissions required to delegate authority and control access, the inheritance feature of Active Directory enables the application of inheritable permissions on objects.

Instead of having to explicitly specify permissions in the DACL of every managed object, administrators can simply specify the same permission on an OU or any object that contains one ore more objects and mark the specified permissions as inheritable. The inheritance feature of Active Directory will take any permission specified as inheritable and propagate them down to all child objects within the container on which the permission was applied and to all child objects of any container objects within the subtree rooted at the container object on which the inheritable ACE was applied.

Since domain information is typically stored in OUs, inheritable permissions are typically applied on OU objects, and inheritance takes care of propagating all inheritable permissions down to all objects in the subtree rooted at the OU on which these inheritable permissions were applied.

Organizations should take into account two primary requirements when designing their OU structures:

- Group Policy application requirements
- Delegation of Administration requirements

When Active Directory is installed, only one OU, Domain Controlers, is automatically created. Domain Controllers is the default container for new domain controllers.

In addition, a number of other containers are automatically created, two of which, Computers and Users, are of interest from the perspective of domain data storage:

- Computers represents the default storage area for new computer objects that were originally created through legacy APIs that are not Active Directory–aware.
- Users represents the default storage area for new user accounts that are created through legacy APIs that are not Active Directory–aware.

By default, when computers are joined to the domain, the computer accounts for these computers are created in the Computers container. Because these two containers are not OUs, Group Policy cannot be applied to them.

| **Note** |
| --- |
| Even though they are not OUs, inheritable permissions can still be applied to the Computers and Users containers, because the concept of inheritance of permissions applies to other containers as well as to OUs. Inheritance works the same in every directory partition, whether the Configuration partition, the Schema partition, or a domain partition. |

Thus, after the deployment of Active Directory, an organization must typically create an OU structure to store domain data including user and computer accounts and groups. When designing an OU structure, organizations must take into account Group Policy application requirements and requirements around delegation of administration.

For best practices and design considerations for creating an effective OU structure, see Chapter 4: Delegating Data Management.

# Chapter 3: Delegating Service Management

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Service management involves managing all aspects of the Active Directory directory service that are required to provide security and reliability in the delivery of the directory service. Overall responsibility for service management is entrusted to service owners who are responsible for the planning and long-term maintenance of the Active Directory infrastructure, ensuring that the directory continues to function, and ensuring that goals established in service-level agreements are met. To achieve these goals, service owners create an administrative delegation model to distribute the various administrative responsibilities among service administrators. Service administrators thus represent the operational arm of the service owners.

Service owners choose a small and select group of highly trusted and skilled administrators from among their service administrators and confer upon this small set of administrators the highest level of privilege by designating them as Enterprise Administrators of the forest. While service owners are responsible for the creation of a delegation model for service management, service administrators assist them during the creation of the model by providing an operational perspective, including the capabilities and the limitations of the system. Upon the completion of the creation of a well-designed delegation model for service management, Enterprise Administrators implement the delegation model to enable the service management administrators team to carry out its assigned responsibilities.

This chapter provides guidance on how to use delegation to provide administrative coverage for all aspects of service management in an Active Directory environment. The chapter presents an overview of the various categories that comprise service management and provides recommendations for how to efficiently delegate all aspects of service management in a security conscious manner.

To help service owners to create the service delegation model, and to help high-level administrators to implement the model, this chapter provides the following information:

- Overview of the categories of service tasks that require management.
- Overview of the default service administration groups and their levels of privilege.
- Recommendations for how to delegate all service management responsibilities in an efficient and security-conscious manner by using service administration roles.
- Best-practices for implementing the service delegation model by creating service administration groups for each administrative role, assigning appropriate permissions, and populating the groups.

# Level-of-Privilege Considerations in Delegating Service Management

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Closely associated with the concept of delegating administration is the implicit notion that since delegation involves authorizing a lesser privileged administrator the ability to carry out a specific set of administrative tasks, it is usually fair to assume that the nature, scope and impact of a delegated administrative task is less security-sensitive and cannot be used to elevate privilege to the same level as that of the administrator delegating the task.

While this notion is generally true, it does not hold for delegation of service management. Service management involves providing administrative coverage to help ensure security and reliability in the delivery of the Active Directory directory service. Thus, by its very nature, it is fair to assume that most administrative tasks involved in promoting security and reliability can and usually do have a significant impact on the security worthiness and the availability of the directory service.

While some administrative tasks are extremely security-sensitive and can have a forest-wide impact, thus requiring highly elevated privileges, most administrative tasks do not require such high levels of privilege. The administrative tasks that do require elevated levels of privilege should only be entrusted to the most highly trustworthy and skilled set of administrators. While most service management administrative tasks do not require elevated levels of privilege, they do require more privileges than those required for data management tasks and they need to be assigned to sufficiently trustworthy personnel.

The level of skill and trust required of these administrators represents a significant cost to the organization. To minimize the total cost of operation and increase the security worthiness of an Active Directory environment, the number of these most highly trustworthy and skilled set of administrators should be minimized by delegating responsibility for relatively less security-sensitive service management administrative tasks to relatively less highly trusted and skilled administrators.

Additionally, since malicious or inadvertent, careless use of needlessly broad administrative credentials can result in irreversible damage, such as accidental deletion of data, inadvertent assignment of wrong values to sensitive data, and inappropriate configuration of essential services, it makes sense to grant service administrators only the level of administrative authority that is required to carry out the set of administrative tasks assigned and none other. For instance, by default Domain Administrators and Enterprise Administrators are sufficiently privileged to carry out just about every service administration task in Active Directory. A majority of these tasks do not require such high-levels of privilege and can be delegated to lesser-privileged administrators so as to minimize the set of highly trusted and privileged administrators and minimize the chance of inadvertent or malicious use of such high and sweeping levels of privilege.

Finally, delegation can be used to achieve a clear separation of assigned responsibilities, thereby making service management more tractable, and increasing accountability and the security worthiness of an Active Directory environment.

For example, assigning responsibility for all aspects of forest-wide replication (but only forest-wide replication) to a specific administrative group not only ensures that a critical aspect of service has been provided administrative coverage, but also increases accountability and helps ensure service delivery. For example, in the event of a replication issue, there is a clearly identifiable set of administrators to which the issue can be escalated, The same set of administrators can be held accountable should there be suspicion that a replication related administrative task was carried out with malicious intent.

For all these reasons, while delegation of service management does not necessarily lead to lesser-privileged delegated administrators, it does make an Active Directory environment measurably more secure and helps reduce the total cost involved in managing the directory service while continuing to ensure the highest levels of security and availability of service.

Note that it is critical and imperative to understand that each and every service administrator should be equally and highly trusted. Every single service administrator is sufficiently privileged to adversely impact the security and reliability of the delivery of the directory service and can consequently adversely impact the data stored in or protected by Active Directory. The importance of ensuring that all service administrators meet the highest bar of trust cannot be overstated.

# Recommended Approach to Service Management

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As mentioned in Chapter 1: Delegation of Administration Overview earlier in this document, a structured and methodical approach to managing an Active Directory environment can greatly enhance the security of the directory service, reduce administrative costs involved in managing Active Directory, and successfully address the administrative needs of all stakeholders, thereby making Active Directory management more tractable seamless, efficient and secure.

A structured and methodical approach to managing the service aspects of an Active Directory environment generally involves the following recommended steps:

1. Understand all aspects of Active Directory service management.

2. Understand the administrative needs of all stakeholders.

3. Create a service management delegation model that ensures that administrative coverage is provided for all aspects of Active Directory management and that the administrative needs of all stakeholders are addressed.

4. Implement the service management delegation model in an efficient and security-conscious fashion, ensuring that the administrative needs of all delegated administrators and all stakeholders are addressed while also ensuring that all administrative access has been granted on the basis of least privilege.

5. Maintain the implemented delegation model, which involves making modifications to the implemented delegation model in response to changes in administrative requirements or needs.

# Service Management Overview

Updated: August 3, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides an overview of all service management categories that is aimed at helping you gain a better understanding of all aspects of Active Directory service management. It also presents Microsoft recommended roles that should be sufficient for providing administrative coverage for all aspects of Active Directory service management, taking into account the administrative needs of the service owners and administrators, who are the main stakeholders in Active Directory service management.

**Service Management Categories**

Service management includes managing all aspects of the directory service that are essential to ensuring the uninterrupted delivery of the directory service across the enterprise. Service management includes, but is not limited to, the following administrative tasks:

- Adding and removing domains and domain controllers.
- Managing and monitoring replication.
- Ensuring the proper assignment and configuration of operations master roles.
- Performing regular backups of the directory database.
- Managing domain and domain controller security policies.
- Configuring directory service parameters, such as setting the functional level of a forest or putting the directory in the special List-Object security mode.

Active Directory service management can be divided into several categories of essential components that require administration. Each service management category has a set of associated tasks that can be delegated to provide for management of every aspect of that category.

| **Note** |
| --- |
| Several service administrative security groups are created by default when you install Active Directory. Some of these groups are used by the high-level administrators who implement the delegation model. Other default groups can also be used to implement service administration roles. For more information about the default service administrator groups and accounts in active, Directory, see Appendix N: Default Active Directory Service Administrator Groups in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document. |

Active Directory service management can be separated into the following categories:

- Installation management
- Schema management
- Trust management
- Knowledge reference management
- Operations master roles management
- Backup and restore management
- LDAP policy management

- Directory service configuration management
- Replication management
- Functional level management
- Directory database management
- Security policy management
- DNS management
- Domain Controller management

For a complete list of the tasks that map to these categories, see Appendix A: Active Directory Administrative Tasks in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document. Note that this list does not include the list of tasks involved in managing the server management aspects of domain controllers.

**Installation Management**
Active Directory is hosted on domain controllers across the enterprise. A forest is created by installing Active Directory on one server that is running Windows Server 2003 or Windows 2000 Server. After the forest root domain is installed, subsequent installations of Active Directory on member servers result in the creation of either new child domains in the forest, or additional domain controllers in an existing domain.

Administrative tasks in this category include the creation and deletion of child domains and the installation of additional domain controllers in domains. The creation of child domains is typically an infrequent operation. Child domains are usually created during deployment of Active Directory. Subsequent to initial deployment, the addition of child domains does not happen frequently. In contrast, the addition of new domain controllers to an existing domain must be performed as and when required, typically either to address domain controller failures in remote sites that have only one domain controller or to provide domain controllers in new locations as needed.

For more information about managing domain controllers, see the *Active Directory Operations Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=21268.

**Schema Management**
The Active Directory Schema contains definitions for the set of objects that can be stored in Active Directory, and it enforces the rules that govern both the structure and the content of Active Directory. The schema consists of a set of the classes, attributes, and syntaxes that represent an instance of one or more classes in the schema. The schema also specifies the relationships between classes of objects. The base schema that ships with Active Directory contains all class and attribute definitions that are used by Windows 2000 Server and Windows Server 2003 and their components. Schema objects can be modified to accommodate needed object classes or attributes that are not available in the default schema, or to remove existing class or attribute schema objects so that instances cannot be created in Active Directory.

Administrative tasks in this category involve the extension and modification of the Active Directory schema.

For more information about managing the schema, see the *Active Directory Operations Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=21268.

**Trust Management**
A trust relationship is a link that is established between domains to enable users in one domain to be authenticated by a domain controller in the other domain. Trust relationships are authentication pipelines that must be present so that users in one domain can be authorized for access to resources in another domain. All domains in a forest have automatic, two-way, transitive trust relationships.

In addition to these automatic trust relationships, four other types of trusts can be created in an Active Directory environment.

- **Shortcut trust.** Can be set up between two domains within a Windows Server 2003 forest to improve user logon times between the domains.
- **External trust.** Can be set up to provide access to resources that are located in Windows NT 4.0 domains or in domains that are located in a separate Windows 2000 or Windows Server 2003 forest that is not joined by a forest trust.

- **Forest trust.** In a Windows Server 2003 forest, can be used to link two disjoined Windows Server 2003 forests together to form uni- or bi-directional transitive trust relationships. A bi-directional trust relationship across forests forms a transitive trust relationship between every domain in both forests.

- **Non-Windows Kerberos realm trust.** Can be set up between a non-Windows-brand operating system Kerberos version 5 realm, such as a UNIX realm, and an Active Directory domain.

Administrative tasks in this category involve the creation, deletion and management of all aspects of trust relationships in the forest.

For more information about managing trusts, see the *Active Directory Operations Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=21268.

**Knowledge Reference Management**

Active Directory stores information about the existence and location of directory partitions, including the names of all directory partitions, global catalog servers that hold partial, read-only copies of the directory partitions, and domain controllers that hold writable copies of the directory partitions. This information is used by the directory service to generate referrals to other domain controllers in the form of *knowledge references*.

Three kinds of knowledge references can be returned in response to an LDAP query:

- A *subordinate reference*, which is knowledge of a directory partition (or partitions) that are directly adjacent (usually below) in the naming hierarchy to a directory partition that is held by the domain controller.

- A *cross-reference*, which is knowledge of one directory partition and is stored in a cross-reference object in the Partitions container. On a specific domain controller, the combination of all cross-references provides knowledge of all directory partitions in the forest, regardless of their locations in the directory tree.

| **Note** |
| --- |
| The state of cross-reference knowledge at any specific time is subject to the effects of directory replication latency. |

- A *superior reference*, which is knowledge of a specifically designated referral location that is used whenever the domain controller has no knowledge of the search base.

Knowledge references are represented by crossRef objects in Active Directory and are stored in the Configuration partition in the Partitions container. Administrative tasks in this category include pre-creation of cross-references and the specification of superior references.

**Operations Master Roles Management**

To provide conflict prevention, Active Directory enforces the requirement that certain operations can be performed on only a single domain controller per forest or per domain, depending on the operation. The domain controllers that are designated as being able to perform these single-master operations are called operations masters. Each operations master has a role that identifies the type of operations for which it is solely responsible.

Active Directory defines five operations master roles. Two are forest-wide roles and three are domain-wide roles:

- **Schema master.** The only domain controller in a forest that can perform write operations to the schema directory partition.

- **Domain naming master.** The only domain controller in a forest that can add or remove domain or application directory partitions.

- **Relative identifier (RID) master.** The only domain controller in a domain that can assign RIDs to a domain controller. RIDs are required by domain controllers to create new security principals.

- **Primary domain controller emulator.** The only domain controller in a domain that can provide PDC functionality. The PDC is required to interact with Windows NT 4.0 workstations, member servers, and backup domain controllers.

- **Infrastructure master.** The only domain controller in a domain that can update the group-to-user references whenever

the members of groups are renamed or changed within a domain.

| Note |
| --- |
| The schema master and domain naming master are per-forest roles and should be held by domain controllers that belong to the forest root domain. The schema master and domain naming master roles should always be placed on the same domain controller, and this domain controller must be a global catalog server. |

Administrative tasks in this category include the transfer and the seizure of the five operations master roles.

For more information about managing operations masters, see the *Active Directory Operations Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=21268.

**Backup and Restore Management**
Backing up and restoring data is critical to failure recovery management. From a management perspective, the only administrative operations in this category are backing up and restoring Active Directory. Each of these operations is performed on domain controllers and each is individually controlled by separate user rights.

For more information about Active Directory backup and restore, see the *Active Directory Operations Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=21268.

**LDAP Policy Management**
LDAP, an industry standard directory service protocol, is the directory access protocol used to query and retrieve information from Active Directory. LDAP defines how a directory client accesses a directory server, shares directory data, and performs directory operations. LDAP enables clients to query, create, update, and delete information stored in a directory service. Because LDAP is a query protocol, certain behavioral aspects of how queries are handled are configurable. By default, all domain controllers in the forest have the same LDAP policy.

Administrative tasks in this category include the configuration of the default LDAP policy and the creation and configuration of domain-controller specific query policies if needed.

| Note |
| --- |
| Although query policies can be configured on a per–domain controller basis, it is advisable not to specify different LDAP policies for different domain controllers. Also note that some domain–controller specific LDAP policies can lead to a denial of service if configured incorrectly. |

**Directory Service Configuration Management**
Certain configuration aspects of the directory service can be specified on a per–domain controller basis. For example, the directory service on a domain controller can be forced to immediately refresh the security group cache by contacting an available global catalog server. The directory service on a domain controller can also be forced to perform an online defragmentation of the directory database. These configurable parameters affect directory service delivery on that specific domain controller.

Administrative tasks in this category include the configuration of the various parameters that affect directory service delivery on that specific domain controller.

**Replication Management**
Active Directory replication is the process by which the changes that are made on one domain controller are synchronized with all other domain controllers in the domain or forest that store replicas of the same data.

Administrative tasks in this category include the creation and maintenance of all aspects of the replication topology (including sites, site-links, subnets and subnet-to-site associations, bridgehead specification and configuration) and monitoring all operational aspects of replication.

For more information about managing replication, see the *Active Directory Operations Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=21268.

**Functional Level Management**
Active Directory functional levels ensure the proper environment for enabling new Active Directory features. Active Directory features added in successive versions of Windows are not always compatible with earlier versions. Functional levels allow you to safely enable such features when all domain controllers in the domain or forest are running an appropriate version of Windows.

Administrative tasks in this category include setting the appropriate functional levels when your Active Directory environment meets the requirements of a specific functional level and you are ready to activate that functional level.

**Directory Database Management**
The directory database is a self-maintained system wherein growth is managed by periodically compacting data and reorganizing free space into contiguous pages. Other than regular backup, the directory database requires no maintenance during ordinary operation. The only requirement is adequate space to accommodate normal growth.

In Windows 2000 environments, large-scale changes, such as bulk deletions and changes to inherited security descriptors, can cause significant database growth. Improved processing in Windows Server 2003 makes these conditions less problematic. However, database growth must be monitored, and if growth is significant, steps must be taken to reduce the size of the database. Database management includes defragmenting the database offline to release free space to the file system. On domain controllers that are running Windows Server 2003, you can also schedule online defragmentation, which is automatic by default.

Database management tasks also include moving database and log files to different locations when hardware maintenance is required. Service administrators who are responsible for database management also perform tests to troubleshoot database corruption and take steps to restore the database files.

For more information about managing the database, see the *Active Directory Operations Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=21268.

**Security Policy Management**
Default security policy settings for the domain and for domain controllers are specified and stored in the form of Group Policy objects (GPOs). These settings take on default (but configurable) values when Active Directory is installed, and they are applied to all new domains. Group Policy settings apply to either users or computers, depending on the nature of the policy.

In a new Windows Server 2003 domain, the following default GPOs protect the domain and all domain controllers:

- Default Domain Policy, which is linked to the domain object and affects all users and computers in the domain (including computers that are domain controllers) through policy inheritance.

- Default Domain Controllers Policy, which is linked to the Domain Controllers OU and affects only domain controllers, because by default, computer accounts for domain controllers are kept in the Domain Controllers OU.

Security policy settings are organized into the following categories:

- Account Policies, which include:

    - Password Policy, which controls password enforcement and lifetimes on domain accounts

    - Account Lockout Policy, which determines the circumstances and length of time that an account will be locked out of the system

    - Kerberos Policy, which determines Kerberos-related settings, such as ticket lifetimes and enforcement

- Local Policies, which include:

    - Audit Policy, which tracks system security events on computers

    - User Rights Assignment, which controls user and administrative actions on computers

    - Security Options, which affects Active Directory, network, file system, and user logon abilities

- Event Log Policy, which defines attributes that are related to the application, security, and system event logs

The management of these security policies falls under service management because these domain controller security policies control security settings for domain controllers and improper configuration of these policy settings can adversely affect the security of the directory service. Additionally, the domain security policy settings are applied on any Windows 2000 workstation or server that is a member of the domain. An individual could change these settings to weaken the security policy

of the domain so that security could be compromised. For instance, if the account policy dictating password strength is set to allow weak passwords, passwords can be broken and thus the security of the system can be compromised.

Administrative tasks in this category include managing the default domain controller security policy and managing the password, account lockout, and Kerberos account policies.

**Service Admin Groups and Accounts Management**
Clearly, administrators who have control of the membership of the powerful administrative group accounts can easily escalate their privilege to become service administrators or make another user a service administrator, thereby being able to negatively impact many aspects of the directory service. Should such administrators acquire malicious intent or be coerced, they can easily disrupt delivery of the directory service. In the worst case scenario, untrustworthy administrators can cause irreversible damage, including disabling or destroying the forest or exposing content that is stored in or protected by the directory service.

**DNS Management**
Windows Server 2003 and Windows 2000 Server use DNS for name resolution instead of the Windows Internet Name Service (WINS) NetBIOS name resolution method, which is used in Windows NT 4.0–based networks. It is still possible to use WINS for applications that require it; however, Active Directory requires DNS. DNS enables users to use friendly names that are easy to remember to connect to computers and other resources on IP networks.

Active Directory uses the name resolution services that DNS provides to enable clients to locate domain controllers and to enable the domain controllers that host the directory service to communicate with each other.

Active Directory is designed to enable easy integration of the Active Directory namespace into an existing DNS namespace. Features such as Active Directory–integrated zones make it easier to deploy DNS by eliminating the need to set up secondary zones and configure zone transfers.

When DNS servers running on domain controllers store their zones in Active Directory, it is not necessary to configure a separate DNS replication topology that uses ordinary DNS zone transfers. Instead, all zone data is replicated automatically by means of Active Directory replication. Therefore, service administration of Active Directory–integrated DNS principally involves managing DNS data that is stored on domain controllers.

The DNS infrastructure supports the Active Directory logical structure to provide computer name resolution throughout the forest and the Internet. The forest owner assigns an Active Directory DNS owner for the forest. The Active Directory DNS owner has a thorough understanding of the existing DNS infrastructure and the existing namespace of the organization.

The Active Directory DNS owner for the forest is responsible for overseeing the deployment of the Active Directory DNS infrastructure and making sure that, if necessary, domain names are registered with the proper Internet authorities.

The Active Directory DNS owner is responsible for the Active Directory DNS design for the forest. If your organization is currently operating the DNS service, then the DNS designer for the existing DNS service works with the Active Directory DNS owner to delegate the forest root DNS name to DNS servers running on domain controllers.

The Active Directory DNS owner for the forest also maintains contact with the DHCP and DNS teams of the organization and coordinates the plans of any individual DNS owners of each domain in the forest with these groups. The DNS owner for the forest ensures that the DHCP and DNS teams are involved in the Active Directory DNS design process so that each group is aware of the DNS design plan and can provide input early on.

For more information about configuring DNS, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* and "Deploying DNS" in *Deploying Network Services* of the Windows Server 2003 Deployment Kit (or see "Designing the Active Directory Logical Structure" on the Web at http://go.microsoft.com/fwlink/?LinkID=4723 and "Deploying DNS" on the Web at http://go.microsoft.com/fwlink/?LinkID=4709).

**Domain Controller Management**
Domain Controllers are member servers that host the Active Directory directory service. Unavailability of a domain controller directly impacts availability of the Active Directory directory service, and a security breach involving unauthorized access to a Domain Controller can directly undermine the security of an Active Directory environment. Thus, providing adequate administrative coverage for all server management aspects of domain controllers and ensuring that domain controllers are always running with up to date software and proper configurations is of primary importance for ensuring delivery of the directory service.

Administrative tasks in this category include managing all server management aspects of domain controllers including managing the hardware on domain controllers, ensuring adequate file-system space, and making sure that the domain controller is up to date in terms of the application of all security patches and hot-fixes.

For more information about managing domain controllers, see the *Active Directory Operations Guide* on the Web at http://go.microsoft.com/fwlink/?LinkID=21268.

**Recommended Roles for Service Management**
Microsoft has engineered a set of recommended roles for delegating service management. These role recommendations take into account defined sets of logically related administrative tasks and the security sensitivity and impact of these tasks.

The following is the set of recommended roles for delegating service management:

- Forest Configuration Operators
- Domain Configuration Operators
- Security Policy Administrators
- Service Admin Managers
- Domain Controller Administrators
- Backup Operators
- Schema Administrators
- Replication Management Administrators
- Replication Monitoring Operators
- DNS Administrators

**Forest Configuration Operators Role**
A number of administrative operations can affect the security of the entire forest. For example, the creation or removal of child domains can have far-reaching consequences on a forest. In the wrong hands, the ability to create child domains could be used to launch a denial of service attack against the entire forest. Thus, any operation that has a forest-wide impact, in the wrong hands, represents a significant threat to forest security. Microsoft recommends that all administrative tasks that affect the configuration of the forest should be aggregated and assigned to the Forest Configuration Operators role. Administrators in this role are responsible for tasks that include but are not limited to:

- Creating and deleting child domains
- Creating, deleting, and managing all trust relationships for the forest
- Creating, deleting, and managing cross-reference objects
- Transferring and seizing the forest-wide operations master roles
- Modifying forest-wide LDAP settings
- Raising the forest functional level

Microsoft recommends the implementation of one instance of this role for every forest in an Active Directory environment. Microsoft also recommends that no more than two or three administrators be assigned to this role.

**Domain Configuration Operators Role**
A number of administrative operations can affect the security of an entire domain, and possibly by extension affect the entire forest. For example, the addition or removal of a domain controller can have far-reaching consequences on a domain, and by extension, the forest. In the wrong hands, the ability to create child domains could be used to obtain access to the entire domain's data and possibly inject malicious changes into a domain. Thus, any operation that has a domain-wide impact, in the wrong hands, represents a significant threat to domain security and by extension to forest security.

Microsoft recommends that all administrative tasks that affect the configuration of a domain should be aggregated and assigned to the Domain Configuration Operators role. Administrators in this role are responsible for tasks that include but are

not limited to:

- Adding and removing additional domain controllers
- Transferring and seizing the domain-wide operations master roles
- Protecting and managing the default Domain Controllers OU
- Protecting and managing the content stored in the System container
- Restoring Active Directory from backup if and when required

Microsoft recommends the implementation of one instance of this role for each domain in an Active Directory forest. Microsoft also recommends that no more than two or three administrators be assigned to this role.

| **Note** |
| --- |
| It might seem logical to eliminate the Domain Configuration Operators role and to assign these tasks to the Forest Configuration Operators role. However, it is usually advisable to separate the responsibility for a domain from the responsibility for the entire forest. Generally, all service management tasks can have a forest-wide impact; therefore, it can be argued that all such tasks should be assigned to the Forest Configuration Operators role. However, the purpose of the Forest Configuration Operators role is to address highly security-sensitive administrative tasks, each of which has a clearly visible forest-wide impact. The purpose of the Domain Configuration Operators role is to separate powers among different sets of administrators and assign responsibility for managing highly security-sensitive administrative tasks, each of which has a clearly visible domain-wide impact, though it is true that these administrative tasks in the wrong hands could be used to cause forest-wide impact. However, the objective is not simply to grant restricted abilities — the objective is to make service management more tractable. |

**Security Policy Administrators Role**

Security policies applied through Group Policy settings affect every user and computer in a domain and also protect domain controllers.

The Domain Security Policy contains settings that apply to any workstation or server in the domain that is running a version of Windows 2000 Server, the Microsoft® Windows XP® Professional operating system, or Windows Server 2003. Domain Security Policy also includes account policy settings that apply to every user account in the domain. Changes to these settings can weaken the security policy of the domain.

For example, if the Account Policies\Password Policy settings that define password strength are set to allow weak passwords, system security might be compromised. Because all domain accounts are subject to the Domain Security Policy, a weak password policy also compromises service administrator accounts. The consequences of a compromised service administrator account are potentially disastrous to the security of the domain and the forest.

Domain Security Policy tasks include administering:

- Password policy settings
- Account Lockout settings
- Kerberos Policy settings

Similarly, every domain has a Domain Controller Security Policy that applies to all domain controllers in the domain. This security policy includes User Rights Assignment settings, which contain both logon rights and privileges. User Rights Assignment settings that apply to domain controllers are extremely security-sensitive. For example, the set of security principals who can log on to a domain controller is defined by the Allow log on locally User Rights Assignment setting in the Domain Controller Security Policy. Similarly, User Rights Assignment settings also define the set of security principals who can perform various highly security-sensitive operations on domain controllers, such as the Load and unload device drivers user right. Users who have this right can inject code that runs in kernel-mode. The Take ownership of files or other objects User Rights Assignment setting enables an individual to take ownership of any files or objects on the domain controller, which includes any directory partition in Active Directory.

Consequently, any individual who controls the Domain Controller Security Policy of any domain is very highly privileged. Such an individual can gain complete and unrestricted access to every object in Active Directory and every securable resource that is a part of the forest (for example, files and databases on servers that are joined to the domain). This individual can also restrict any and every other individual – including members of the Enterprise Admins, Builtin Administrators, and Domain Admins groups – from being able to log on or gain access to any resource in the forest.

Microsoft thus recommends entrusting responsibility for managing all aspects of Domain Controller Security policy for every domain in the forest, and managing the password, account lockout, and Kerberos aspects of the Domain security policy for every domain in the forest to a single group of administrators who take on the Security Policy Administrators role. Administrators in this role are responsible for the following tasks:

- Managing all aspects of the Domain Controller Security Policy for all domains in the forest
- Managing the following aspects of Domain Security Policy for all domains in the forest:
    - Password Policy
    - Account Lockout
    - Kerberos Policy

Microsoft recommends the implementation of one instance of this role for each Active Directory forest. Microsoft also recommends that no more than two or three administrators be assigned to this role.

| **Note** |
| --- |
| It might seem logical to eliminate the Security Policy Admins role and to assign these tasks to the Forest Configuration Operators role. However, it is usually advisable to separate the responsibility for security policy management from the responsibility for the entire forest. Generally, all service management tasks can have a forest-wide impact; therefore, it can be argued that all such tasks should be assigned to the Forest Configuration Operators role. However, the purpose of the Forest Configuration Operators role is to address highly security-sensitive administrative tasks, each of which has a clearly visible forest-wide impact. The purpose of the Security Policy Admins role is to separate powers among different sets of administrators and assign responsibility for managing all sensitive security policies for the entire forest to a specific set of administrators, though it is true that these administrative tasks in the wrong hands could be used to cause forest-wide impact. However, the objective is not simply to grant restricted abilities — the objective is to make service management more tractable. An organization could very well decide to make the same set of administrators members of more than one service management role. |

**Service Admin Managers Role**

Because of the elevated power that is attached to service administrator accounts, it is advisable to dedicate a role to protecting and managing all service administrator accounts across the forest.

Microsoft recommends the implementation of the Service Admin Managers role to protect the service administrator group memberships across the forest. Administrators assigned to this role must be the only ones in the forest that are allowed to manage and modify service administrator group memberships and manage service administrator accounts. Microsoft thus recommends that there be only one instance of this role in a forest. The size of your Active Directory infrastructure determines the number of service administrator groups and accounts in your environment, and that number determines the number of service administrators to assign to this role. As a best practice, you should strive to keep the number of administrators assigned to this role down to a bare minimum.

To state the responsibilities assigned to this role again, administrators in the Service Admin Managers role should be the only ones responsible for the creation, management, protection, and deletion of all service administration accounts and security groups in an Active directory forest.

| **Note** |
| --- |

It might seem logical to eliminate the Service Admin Managers role and to assign these tasks to the Forest Configuration Operators role. However, it is usually advisable to separate the responsibility for service administrator group and member management from the responsibility for the entire forest. Generally, all service management tasks can have a forest-wide impact; therefore, it can be argued that all tasks should be assigned to the Forest Configuration Operators role. However, the purpose of the Forest Configuration Operators role is to address highly security-sensitive administrative tasks, each of which has a clearly visible forest-wide impact. The purpose of the Service Admin Managers role is to protect and manage all service admin groups and accounts in the entire forest. It is true that these administrative tasks in the wrong hands could be used to cause forest-wide impact. However, the objective is not simply to grant restricted abilities — the objective is to make service management more tractable. An organization could very well decide to make the same set of administrators members of more than one service management role.

**Domain Controller Administrators Role**

Managing Domain Controllers and keeping domain controllers running with current software and proper configurations is of primary importance throughout the forest for ensuring delivery of the directory service. In addition to managing services, administrators who manage domain controllers are responsible for managing the data in the file systems of domain controllers that is also critical to delivery of the directory service.

Microsoft thus recommends entrusting responsibility for managing all aspects of Domain Controllers to the Domain Controller Administrators role. Administrators in this role are responsible for tasks that include but are not limited to the following operations on domain controllers:

- Installing and modifying software
- Installing service packs and hot-fixes
- Configuring directory service settings in the registry
- Maintaining and backing up event logs
- Configuring the Service Control Manager
- Managing directory service files and Sysvol
- Starting and shutting down domain controllers
- Other security-sensitive operations, including directory database offline operations

Domain controllers should always be placed in highly secure locations. Most organizations deploy multiple domain controllers in a central datacenter and also deploy multiple domain controllers in several remote datacenters to provide global coverage.

Some organizations might have a large number of satellite sites for branch offices that have only one or two domain controllers. In this situation, organizations are very highly encouraged to evaluate and deploy a remote administration solution, such as Remote Insight Lights-Out (RILO) to provide complete remote administration and management of domain controllers based in remote locations by a centrally managed administrative team. RILO-based remote server management solutions enable central management of remote servers from the primary operations center. These products are typically hardware-based and allow an administrator to remotely perform almost any function that can be performed locally, other than actually modifying the hardware on the remote servers. Organizations that have a large number of branch office sites should ideally have an administrative group that is located at the primary operations center for the Active Directory environment. The centrally located group is responsible for remotely managing domain controllers in all branch offices.

Use the following criteria to determine how many instances of the Domain Controller Administrators role to assign:

- When datacenters are deployed in different locations, assign one instance of the role for every physical location that houses multiple domain controllers. Assign administrators to this role on the basis of the number of domain controllers that must be managed. It is recommended that the number of administrators assigned to each instance of this role be kept to an absolute minimum. In the situation where domain controllers are housed in a data center that also houses other servers providing various applications and services, if at all possible do not make administrators responsible for managing other servers and services members of this role, and select and appoint only the most highly trusted and skilled administrators to this role.

- When domain controllers are located in a large number of branch office sites, assign an additional instance of the role to carry out branch office administration from the datacenter where the domain controllers for the forest root domain are located. Assign administrators to this role on the basis of the number of domain controllers that must be managed. Again, if at all possible, minimize the number of administrators assigned to this role but ensure that you can provide continues coverage.

For more information about products that enhance remote administration of Windows-based servers, search the Windows Server Catalog home page on the Microsoft Web site.

For more information about planning for remote server management, see "Planning for Remote Server Management" in Planning Server Deployments" of the *Windows Server 2003 Deployment Kit* (or see "Planning for Remote Server Management" on the Web at http://go.microsoft.com/fwlink/?LinkId=15309.

**Note**

It might seem completely logical to eliminate the Domain Controller Admins role and to assign these tasks to the Forest Configuration Operators role. However, the objective is to separate responsibility in order to make service management more tractable. The purpose of the Domain Controller Admins role is to provide adequate coverage to ensure that the server management aspects of Domain Controllers are more than adequately taken care of. If a Domain Controller goes down, its impact is immediately noticeable and usually unaffordable. It is true that these administrative tasks in the wrong hands could be used to cause forest-wide impact. Thus, as is the case with all service management roles, ensure that you assign only extremely trustworthy and highly skilled administrators to all service management roles. An organization could very well decide to make the same set of administrators members of more than one service management role.

**Backup Operators Role**

Backing up data is critical to every computer environment. As part of normal operational practices, you should regularly back up domain controllers and secure the backup media to minimize the risk of data tampering or theft.

Backing up Active Directory is performed by backing up system state. System state includes Active Directory, the domain controller registry, and Sysvol. Because the system state backup contains all directory service–related data, theft of the backup media presents the same risks as theft of the domain controller or theft of a disk drive from the domain controller. The attacker could restore the information elsewhere and gain unauthorized access to Active Directory data.

Administrative tasks that involve backing up Active Directory can be aggregated and assigned to the Backup Operators role. Administrators in this role are responsible for backing up system state on a specified schedule.

For an Active Directory single-forest environment, create one instance of this role for every domain in the forest. For the single instance of this role, assign not more than two or three administrators. In addition, from an operational perspective, it is recommended that a specific domain controller be used to perform regular backup operations; a domain controller located in a main datacenter would make an ideal candidate.

Note that the administrative task of restoring Active Directory has been assigned to the Domain Configuration Operators role. This is because the act of restoring Active Directory, unlike backing up Active Directory, is not a regular or frequent task – it is only performed when an Active Directory partition needs to be restored from backup.

**Note**

It might seem completely logical to eliminate the Domain Controller Admins role and to assign these tasks to the Forest Configuration Operators role. After all, a Backup operator is sufficiently privileged to be able to log on to a Domain Controller and back up Active Directory. Should an administrator in this role turn malicious or be coerced, he or she could modify the backed up data to inject malicious changes. However, the objective is to separate responsibility to make service management more tractable. It is true that these administrative tasks in the wrong hands could be used to cause forest-wide impact. Thus, as is the case with all service management roles, ensure that you assign only extremely trustworthy and highly skilled administrators to all service management roles. An organization could very well decide to make the same set of administrators members of more than one service management role.

**Schema Administrators Role**

Because the schema is common to the entire forest, its modification has a forest-wide impact. Improper modifications can result in disrupting the operation of existing applications or in corrupting the directory. In addition, the ability to modify the schema provides the ability to modify the default security descriptors of classes. Default security descriptors protect newly created objects in the directory. Therefore, through certain changes to the default security descriptor, a malicious administrator

can gain the ability to compromise security by creating user objects that have highly elevated privileges. To isolate this responsibility to a few select administrators, a separate role is assigned for schema management.

Administrators who are assigned to the Schema Administrators role are responsible for all aspects of schema management. This includes making changes that involve the extension of the Active Directory schema as well as making changes that involve deactivating existing classes or attributes and/or specifying that an attribute be replicated to the Global Catalog.

For an Active Directory single-forest environment, create one instance of this role. For the single instance of this role, assign one or a maximum of no more than two administrators.

| **Note** |
| --- |
| It might seem completely logical to eliminate the Schema Admins role and to assign these tasks to the Forest Configuration Operators role. After all, a Schema Administrator could modify default security descriptors and escalate his or her privilege. Should an administrator in this role, turn malicious, or be coerced, he or she could modify the Schema and in the worst case cause Active Directory to stop functioning. However, the objective is to separate responsibility to make service management more tractable. It is true that these administrative tasks in the wrong hands could be used to cause forest-wide impact. Thus, as is the case with all service management roles, ensure that you assign only extremely trustworthy and highly skilled administrators to all service management roles. An organization could very well decide to make the same set of administrators members of more than one service management role. |

**Replication Management Administrators Role**

Active Directory replication automatically synchronizes data among domain controllers in a forest that store replicas of the same directory partitions. Replication is critical to ensure that all domain controllers provide current data. Although replication between domain controllers in the same site requires no manual configuration, replication between domain controllers in different sites must be manually configured to optimize network bandwidth and set replication latency between sites to acceptable levels. While Active Directory replication was designed to require minimal administrative intervention, ensuring healthy, timely and optimal replication of Active Directory requires that the replication topology be set up properly and be appropriately configured when a change is required, and that replication be monitored to ensure uninterrupted and reliable delivery of service

Microsoft recommends that all administrative tasks that involve setting up and configuring replication topology and all tasks that involve managing the operational aspects of replication (changing replication configuration) be assigned to the Replication Management Administrators role. Administrators in this role are responsible for tasks that include but are not limited to:

- Creating, managing, and deleting sites
- Associating subnets to sites
- Creating, managing, and deleting site links and site-link bridges
- Adding and removing sites from site links and site-link bridges
- Modifying the replication schedule and replication interval on site links
- Creating and deleting manually created connections
- Forcing replication between two domain controllers
- Force a full replication synchronization

Note that once the replication topology is initially set up, it will not require frequent changes unless new requirements call for the modification of the existing replication topology.

For an Active Directory single-forest environment, create one instance of this role. For the single instance of this role, depending on the size of the Active Directory environment, this role should have only the minimum number of administrators required at any given time to react to replication related issues.

**Replication Monitoring Operators Role**

Microsoft recommends that all administrative tasks that involve monitoring replication activity be assigned to the Replication

Monitoring Operators role. Administrators in this role are responsible for tasks that include but are not limited to:

- Monitoring replication

For an Active Directory single-forest environment, create one instance of this role. Since the replication monitoring requirements are usually dependent on the size and complexity of an organization's Active Directory environment, the number of administrators to assign to this role should be determined based on the smallest number of administrators that need to be monitoring replication during a single work shift, multiplied by the number of work shifts in the organization.

Responsibility for replication management is thus divided between two roles:

- A role to manage replication topology, including creating and maintaining sites and subnets, site links, and site link bridges, as well as configuring site links and make configuration changes as required.

- A role to monitor replication.

Although all replication tasks could be assigned to just one role, separating administration of tasks that are specific to managing the replication topology from tasks that are specific to monitoring and configuring replication provides separation of responsibilities and increases accountability. It is for this reason that Microsoft recommends two separate roles for complete replication management.

| **Note** |
| --- |
| It might seem logical to eliminate the Replication Management Admins and Replication Monitoring Operators roles and to assign these tasks to the Forest Configuration Operators role. After all, any administrator in either role could disrupt service, thereby causing forest–wide impact. However, the objective is to separate responsibility to make service management more tractable. It is true that these administrative tasks in the wrong hands could be used to cause forest-wide impact. Thus, as is the case with all service management roles, ensure that you assign only extremely trustworthy and highly skilled administrators to all service management roles. An organization could very well decide to make the same set of administrators members of more than one service management role. |

**DNS Administrators Role**

Microsoft recommends that all administrative tasks involved in managing Active Directory–integrated DNS be assigned to the DNS Administrators role. This role provides complete administrative coverage of Active Directory–integrated DNS management.

Administrators in this role are responsible for tasks that include but are not limited to:

- Installing the DNS Server service on domain controllers.

- Configuring domain controllers that are running DNS to use either forwarding or root hints for recursive name resolution, depending on which method the existing DNS service uses.

- Configuring the forest root domain controller to host the DNS zone that corresponds to the forest root DNS name.

- Configuring the domain controllers for each regional domain to host the DNS zone that corresponds to the DNS name of the domain.

- Configuring the zone containing the forest-wide locator records to replicate to every DNS server in the forest by using the forest-wide DNS application partition.

For an Active Directory single-forest environment, create one instance of this role. The number of administrators required for this role depends on the size of the environment and the number of domain controllers that are DNS servers. It is recommended that the number of administrators in this role be limited to an operationally required minimum.

For more information about DNS integration with Active Directory, see "Deploying DNS" in *Deploying Network Services* of the Windows Server 2003 Deployment Kit (or see "Deploying DNS" on the Web at http://go.microsoft.com/fwlink/?LinkId=4709.

| **Note** |
| --- |

Any DNS administrator could disrupt service, thereby causing a forest–wide impact. However, the objective is to separate responsibility to make service management more tractable. It is true that these administrative tasks in the wrong hands could be used to cause forest-wide impact. Thus, as is the case with all service management roles, ensure that you assign only extremely trustworthy and highly skilled administrators to all service management roles. An organization could very well decide to make the same set of administrators members of more than one service management role.

**Summary of Recommended Roles**

The set of 10 Microsoft recommended roles listed earlier in this document is sufficient to provide administrative coverage for all aspects of Active Directory management and does so in an efficient and security-conscious fashion that enhances the security worthiness of your Active Directory service, reduces the risk of inadvertent misuse of high levels of privilege, eliminates requiring a large number of highly-privileged administrators (such as Domain Admins), increases accountability, clearly separates responsibilities, and ensures that all aspects of your Active Directory environment have adequate administrative coverage.

The set of recommended roles together provides coverage for all categories of service management described earlier, as shown in Table 1.

**Table 1 Service Management Categories and Service Management Roles**

| Service Management Category | Service Management Role |
|---|---|
| Installation management | Forest Configuration Operators<br>Domain Configuration Operators |
| Schema management | Schema Administrators |
| Trust management | Forest Configuration Operators |
| Knowledge Reference management | Forest Configuration Operators |
| Operations Master Roles management | Forest Configuration Operators<br>Domain Configuration Operators |
| Backup and Restore management | Backup Operators<br>Domain Configuration Operators |
| LDAP Policy management | Forest Configuration Operators |
| Directory Service Configuration management | Domain Controller Administrators |
| Replication management | Replication Management Administrators<br>Replication Monitoring Operators |
| Functional Levels management | Forest Configuration Operators |
| Directory Database management | Domain Controller Administrators |
| Security Policy management | Security Policy Administrators |
| Service Admin Groups and Accounts management | Service Admin Managers |
| DNS management | DNS Administrators |
| Domain Controller management | Domain Controller Administrators |

Table 2 presents a role-centric view of the various service management categories covered by each role.

**Table 2 Role-Centric view of Service Management Category Coverage**

| Service Management Role | Service Management Category Covered |
|---|---|

| | |
|---|---|
| Forest Configuration Operators | Installation management<br>Trust management<br>Knowledge Reference management<br>Operations Master Roles management<br>LDAP Policy management<br>Functional Levels management |
| Domain Configuration Operators | Installation management<br>Operations Master Roles management<br>Backup and Restore management |
| Security Policy Administrators | Security Policy management |
| Service Admin Managers | Service Admin Groups and Accounts management |
| Domain Controller Administrators | Directory Service Configuration management<br>Directory Database management<br>Domain Controller management |
| Backup Operators | Backup and Restore management |
| Schema Administrators | Schema management |
| Replication Management Administrators | Replication management (topology and operational aspects) |
| Replication Monitoring Operators | Replication management (health monitoring) |
| DNS Administrators | DNS management |

**Default Service Administrative Groups**

Every installation of Active Directory has a number of default service administrative groups, some of which are sufficiently privileged to carry out all service-management administration tasks. For example, members of the Enterprise Admins, Domain Admins, and Built-In Administrators groups can perform all administrative tasks that are involved in managing an Active Directory service environment. For all practical purposes, these groups can be considered to have equal abilities.

By default, these accounts are granted access to directory and server resources when Active Directory is installed. Appendix N: Default Active Directory Service Administrator Groups in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document, lists every default service administrator group and account and provides a brief description of each group and account, including the qualities that make each group a service administrative group.

This deserves mention because once a delegation model based on Microsoft-recommended roles is created and implemented, most of these security groups should theoretically no longer be required. Thus an organization could choose to gradually empty the memberships of most of these groups after carefully ensuring that the groups are no longer required. However, the following groups should ideally not be emptied: Enterprise Admins, Domain Admins, Builtin Admins, and Schema Admins. These groups should be retained either because they might be used to implement a Microsoft recommended role (for example, Backup operators or Schema Admins) or they might be required in the event of an emergency situation that warrants some action to be taken by Enterprise or Domain Admins.

The decision to consider emptying and possibly discontinuing some of these groups is purely at the discretion of service owners and is dependent on whether or not an organization chooses to implement all the Microsoft-recommended roles. Some organizations might choose to customize and delegate some or a few Microsoft-recommended roles while others might choose to stick to the default service administrative group, and still others might completely embrace and implement a service management delegation model based on Microsoft-recommended roles.

# Creating a Service Management Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process of creating an administrative delegation model that addresses service management needs must be guided by security and efficiency considerations. By following the recommendations that are presented in this section, you can create a documented administration model that ensures that:

- Administrative coverage has been provided for all aspects of Active Directory service management.
- Administrative responsibilities have been distributed among appropriate administrative groups in accordance with sound security practices.

Many service administrative tasks require permissions that allow an administrator to make potentially harmful changes to the directory service. Harmful changes might be made in error, by a malicious administrator, or by an unauthorized administrator who has usurped permissions. Avoiding errors is the reason for selecting the most highly skilled administrators to perform security-sensitive administrative tasks. Avoiding malicious changes is the reason for selecting the most highly trusted service administrators. And avoiding the ability of an unauthorized administrator to wrongfully gain permissions is the reason for incorporating security measures into the delegation model itself.

The efficiency of the administrative delegation model requires that tasks be divided among the appropriate number and level of administrators to provide an optimal workload, while also providing complete coverage. Both the efficiency and security of the administrative delegation model are served by assigning the least privileges that are required to perform each task.

| **Note** |
| --- |
| The information in this document applies to Active Directory infrastructures that have an existing forest and domain structure. For information about creating a forest and domain structure, see "Designing the Active Directory Logical Structure" in *Designing and Deploying Directory and Security Services* of the Windows Server 2003 Deployment Kit (or see "Designing the Active Directory Logical Structure" on the Web at http://go.microsoft.com/fwlink/?LinkId=4723). |

**Guidelines for Creating a Service Delegation Model**

Creating an efficient and security-conscious model based on these roles involves the following steps:

1. Understand the nature of the responsibilities assigned to each Microsoft-recommended role and, if needed, customize the definition of these roles by adding or removing assigned tasks from the recommended role definitions.
2. Determine the number of instances of the role that are required for your environment. The number of instances of each Microsoft-recommended role is presented earlier in this document and is summarized in Table 3.

**Table 3 Service Management Roles and Recommended Number of Instances**

| Service Management Role | Recommended Number of Instances |
| --- | --- |
| Forest Configuration Operators | One instance per forest |
| Domain Configuration Operators | One instance per domain |
| Security Policy Administrators | One instance per forest |
| Service Admin Managers | One instance per forest |

| | |
|---|---|
| Domain Controller Administrators | a. One instance for every set of domain controllers physically located in a centrally or remotely located datacenter<br><br>b. One instance for the set of all domain controllers physically located in remote locations but centrally managed via remote management solutions<br><br>c. Optionally, one instance for every domain controller that is locally managed at a branch office, assuming that a remote management solution is not being used. |
| Backup Operators | One instance per domain |
| Schema Administrators | One instance per forest |
| Replication Management Administrators | One instance per forest |
| Replication Monitoring Operators | One instance per forest |
| DNS Administrators | One instance per forest |

3. Based on the size and specifics of your Active Directory environment, determine the number of administrators to assign to each instance of the role.

4. For each instance of each role, identify the specific administrators that shall be assigned to the role. Document these role assignments.

**Documenting Your Service Administrative Delegation Model**

After completing the process for creating the administrative delegation model, service owners should document the model and then maintain this document on an ongoing basis. When changes in service administration are required due to organizational changes or growth, the model must be changed to reflect the new requirements.

The administrative delegation model document should include the following:

- A list of all roles that are being used
- The number of instances of every role
- A list of administrators that are assigned to each role instance
- The administrative responsibilities that are assigned to each role

Later in the process, when the delegation model is implemented, update this document to include the following information:

- The specific security group that is used to represent each role instance
- The specific permissions granted to each security group
- All audit settings that are set to monitor administrative operations

Table 4 is a template you can use to document your role instances.

**Table 4 Template to Document Specific Role Instances (With Example Values)**

| Field | Assignment Information |
|---|---|
| Role Instance Name: | Contoso Sec Pol Admins |

| | |
|---|---|
| Instance of: | Security Policy Admins |
| Instance #: | 1 of 1 |
| Assigned Administrators: | Michael Allen |
| Assigned Tasks: | No customization – based on Microsoft-recommended role-to-task assignments.<br><br>Refer to document located at: document location |
| Security Group | To be supplied during Implementation phase |
| Permissions Assigned | To be supplied during Implementation phase |
| Notes | Members in this role are based in Chicago<br><br>Security Clearance requirements met for all administrators |

# Implementing the Service Management Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Implementing the service administrative delegation model involves the creation of security groups to represent role instances and the granting of sufficient permissions for every security group representing a role instance so as to enable the administrators in each role (represented by membership in corresponding security groups) to carry out their assigned responsibilities. The implementation process is performed by a select few service administrators who are trusted by the service owner to faithfully delegate the administrative roles according to the delegation model. These administrators are usually members of the Enterprise Admins group.

To implement each role, the administrator responsible for implementing the service management delegation model performs the following general tasks:

1. Define a location to store the security groups that represent all role instances.

2. Create a new group to represent the role or designate an existing default administrative group (for example, Schema Admins) for each role instance.

3. Enable each instance of the role by doing the following:

   Assign the collective set of permissions (required to perform the set of all administrative tasks assigned to a role) to the security group representing the specific instance of the role.

   Assign the collective set of user rights (required to perform the set of all administrative tasks assigned to a role) to the security group representing the specific instance of the role for all operations that the administrators must perform on the system (usually the domain controller).

4. Delegate the role that you have enabled by adding the user accounts of the assigned administrators to the appropriate group.

Make sure that the groups that you use to implement the role are used only for that role and contain only the assigned administrators as members.

**Storing Security Groups for Service Management**

Most security groups for service management will be created in, and managed from, the forest root domain. Thus, it is appropriate to create and store the security groups that represent service administration role instances in a specific location so that they can be managed easily.

Create an OU in the forest root domain, directly under the domain object. Name the OU “Service Management” or choose a similar name that adequately conveys the purpose of the OU.

Figure 1 shows a domain hierarchy with the Service Management OU.




**Creating or Assigning Security Groups**

Table 5 shows a recommended list of names to use for the security groups that represent various instances of service management roles.

**Table 5 Recommended Names for Service Management Role Instances**

| Service Management Role | Recommended Names |
|---|---|
| Forest Configuration Operators | Forest Name Forest Config Ops |
| Domain Configuration Operators | Domain Name Domain Config Ops |
| Security Policy Administrators | Forest Name Security Policy Admins |
| Service Admin Managers | Forest Name Service Admin Managers |
| Domain Controller Administrators | Domain Name location DC Admins<br><br>location could refer to:<br><br>1. A set of domain controllers physically located in a centrally or remotely located datacenter<br><br>2. A set of all domain controllers physically located in remote locations but centrally managed via remote management solutions<br><br>3. Every domain controller that is locally managed at a branch office, assuming that a remote management solution is not being used. |
| Backup Operators | Use the pre-defined Backup Operators group for this role |
| Schema Administrators | Use the pre-defined Schema Admins group for this role |
| Replication Management Administrators | Forest Name Repl Mgmt Admins |
| Replication Monitoring Operators | Forest Name Repl Monitoring Ops |
| DNS Administrators | Forest Name DNS Admins |

**Enabling Service Management Administrative Roles Instances**

Appendix L: Implementing Service Management Delegation Roles in Best Practices for Delegating Active Directory

Administration: Appendices, which accompanies this document, provides the guidelines for implementing each service management role and provides the set of permissions that should be granted to implement each role. Implementing each role instance by granting the set of permissions provided for each role enables each role instance.

**Protecting Service Management Administrative Role Instances**

After you have implemented your service management administrative role instances, it is recommended that you adequately protect all security groups representing the various administrative role instances and their members.

One way of doing this is to make every security group representing a role instance a member of one of the default service administrator groups protected by the AdminSDHolder object. Certain administrative roles instances can be represented by existing service administrator groups in Active Directory – these security groups do not need to be members of another service administrator group. This will ensure that only the Service Admin Managers and Enterprise Administrators can modify the memberships of all service admin accounts and of all administrative roles.

It is recommended that you make all these groups members of the Print Operators group in the specific domain to which the security groups representing the roles belong, and take away the Load and unload device drivers privilege from Print Operators (by modifying the default Domain Controller security policy in each domain). This will ensure that all security groups representing administrative roles are protected. Note that there is nothing special about the Print Operators group. The use of this group to protect service admin role instances is only being recommended because it is one of the less used service administrative groups.

**Delegating Role Instances**

To delegate role instances, modify the group memberships of the respective security groups representing each role by adding the user accounts of all administrators that were assigned to each role during the model creation phase.

Do not forget to update your documentation for each role, noting specifics such as the collective set of permissions assigned and the specific security group used to implement each service administration role instance.

# Maintaining the Service Management Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Service delegation models require minimal maintenance under most operating conditions. Tasks that are required to maintain the service delegation model can include:

- Adding (delegating) or removing (undelegating) members to or from service administration role instances
- Adding new instances of roles in the event of infrastructure changes such as, but not limited to:
    - Addition of a new branch site that has domain controllers for the domain of the business unit
    - Removal of a branch site
- Modifying a role by assigning or revoking a new or existing task from the role definition
- Creating new custom roles if the need arises
- Undelegating a role completely by revoking all permissions that are granted to the group that represents the role

**Adding or Removing Members from Role Instances**
To add or remove members from existing role instances, simply modify the membership of the security group representing the role.

**Adding New Role Instances**
To add a new role instance, follow steps 2 through 4 listed in the Implementing the Service Management Delegation Model section.

**Modifying a Role Definition**
To modify an existing role definition, perform the following steps:

1. Modify the documented set of administrative tasks assigned to a specific role definition by adding a new task, removing an existing task.

2. Identify all existing instances of this role and appropriately modify the permissions granted to the role as follows:

   If a new task was added to the role, determine the minimal and precise set of permissions required to delegate the new task by referring to Appendix A: Active Directory Administrative Tasks in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document, and appropriately grant the security group representing the specific instance of this role the set of permissions determined earlier (as indicated in "Appendix A").

   If an existing task was removed from the role definition, determine the minimal and precise set of permissions that are required to delegate this task by referring to Appendix A: Active Directory Administrative Tasks in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document, then ensure that the same setoff permissions (or even one of them) is not required to perform any other tasks in the role definition, and then appropriately revoke assigned permissions for the security group representing the specific instance of this role (as indicated in "Appendix A").

**Creating New Custom Roles**
To create a new custom role, perform the following steps:

1. Understand and document the purpose of the new role.

2. Assign a set of administrative tasks to this new role.

3. Determine the minimal and precise set of permissions required to delegate the set of administrative tasks identified earlier.

4. Document the general scope where permissions must be applied in the directory (or perhaps on the file system or registry or Group Policy on Domain Controllers).

To enable new instances of this role, follow steps 2 through 4 listed in the Implementing the Service Management Delegation Model section.

**Undelegating Role Instances**

To undelegate existing role instances, perform the following steps:

1. Revoke all permissions granted to the security group representing the specific instance of the role. This might involve revoking user rights in the Domain Controller Security policy.

2. Optionally, empty the group membership of the security group representing the specific instance of the role.

3. To ensure that all permissions have been removed, double-check your documentation and make sure that permissions have been removed from all places where they were initially granted.

| **Note** |
| --- |
| Dsrevoke.exe, a command-line tool that automatically revokes permissions, can be used to revoke delegated authority, but only on an OU. Because most of the data that is protected by service management role permissions is stored in the configuration and schema directory partitions in Active Directory, or in the file system or registry of a domain controller, manual intervention is usually required to undelegate a service administration role. |

Creating custom roles involves using the delegation process that described earlier in this chapter to identify tasks and map them to a group, configure permissions and rights, and then populate the group. The same is true for ad hoc roles, which are removed when the need no longer exists.

# Chapter 4: Delegating Data Management

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Data management involves managing all aspects of the data stored in or protected by Active Directory. This content includes domain data, which primarily consists of accounts for users and computers, security groups, and application-specific data.

User accounts represent the identities of an organization's users, and thus data management includes identity management. Security-groups enable aggregation of a set of users for the purpose of authorization, and thus security group management plays a key role in securing an organization's IT resources. Security group management is also a key data management category. Active Directory–integrated and –enabled applications can store application specific data in Active Directory, and management of application-specific data is also a part of data management.

Finally, computers can be viewed as providing resources in an IT infrastructure. From end-user workstations to file servers, and from hosting a Web farm to running enterprise applications, computers provide resources in an organization's IT infrastructure. Thus resource management is yet another data management category.

This chapter provides guidance for how to use delegation to provide administrative coverage for all aspects of data management in an Active Directory environment. The chapter presents an overview of the various categories that data management includes and provides recommendations about how to efficiently delegate all aspects of data management in a security-conscious manner.

# Recommended Approach to Data Management

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As mentioned in Chapter 1: Delegation of Administration Overview, a structured and methodical approach towards managing an Active Directory environment can greatly enhance the security of the environment, reduce administrative costs involved in managing the environment and successfully address the administrative needs of all stakeholders, thereby making Active Directory management more tractable seamless, efficient and secure.

It is not uncommon for multiple business units to participate in a shared Active Directory environment. Each business unit, while participating in a shared Active Directory environment, can have its own domain data and IT resources. Each business unit should be assigned a data owner who should have overall responsibility for all aspects of data management for that business unit. Business unit owners appoint a small set of their most highly trusted administrators to the role of Business Unit Admins – this administrative group represents the operational arm of the data owner.

Since data owners are entrusted with overall accountability for the security-conscious, efficient, and effective management of the content stored in Active Directory and on computers joined to Active Directory, it is their responsibility to create a delegation model under which administrative responsibility for managing business unit data is distributed among and delegated to their data administrators also in an effective, efficient, and security-conscious manner.

It is important to note that this chapter provides best-practices for delegation only on a per domain basis. Thus, if your Active Directory environment has multiple domains, you should apply the best-practices and guidelines presented in this chapter for creating a delegation model to each domain in your Active Directory environment.

A structured and methodical approach towards Active Directory data management generally involves the following recommended steps:

1. Service owners hand-off data management of each business unit to the respective data owners of each business unit.

2. Business Unit data owners in turn create, implement, and maintain a delegation model for managing business unit data.

This approach offers two primary benefits. It clearly separates responsibility for service management and data management. Additionally, it grants responsibility for data management of data that belongs to a specific business unit to the business unit owners themselves. This gives them the flexibility to create, implement and maintain their data in a manner that best addresses the specific administrative needs and requirements of that business unit.

After service owners have handed off responsibility for data management to data owners, the data owners, together with their high-level administrators, should do the following:

1. Understand all aspects of Active Directory data management.

2. Understand the administrative needs of all stakeholders.

3. Create a delegation model that ensures that administrative coverage is provided for all aspects of Active Directory management and that the administrative needs of all stakeholders are addressed.

4. Implement the Active Directory data management delegation model in an efficient and security-conscious fashion, ensuring that the administrative needs of all delegated administrators and all stakeholders are addressed, while also ensuring that all administrative access has been granted on the basis of least privilege.

5. Maintain the implemented delegation model, which involves making modifications to the model in response to changes in administrative requirements or needs.

# Understanding Data Management

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following section provides an overview of all data management categories and is aimed at helping you gain a better understanding of all aspects of Active Directory data management. This section also presents Microsoft-recommended roles that should be sufficient for providing administrative coverage for all aspects of Active Directory data management, taking into account the administrative needs of the service owners and administrators, who are the main stakeholders in Active Directory data management.

**Data Management Categories**
Data management involves managing all aspects of the data stored in or protected by Active Directory. This content includes domain data, which primarily consists of accounts for users and computers, security groups, and application-specific data. As mentioned earlier in this chapter, data management includes account management, security-group management, resource management and application-specific data management.

**Account Management**
Account management entails managing all aspects of user accounts – from creating new accounts to maintaining them and providing account support to disabling and finally deleting accounts that are no longer needed. For example, when a user joins a business unit, either a new account must be created or an existing account must be moved from an existing business unit. During its existence, the account requires maintenance and support. Account support includes administrative operations like resetting passwords and unlocking user accounts. At some point, the user leaves the business unit and the account must be discontinued: Account management includes the administrative tasks for disabling and finally deleting user accounts.

**Security Group Management**
A security group is a collection of security principals. Security-groups enable aggregation of a set of users for the purpose of authorization. Security group management typically includes such administrative tasks as creating security groups, provisioning security groups to grant access to various resources, managing the membership of security groups, and discontinuing security groups.

Security groups typically have two main uses:

- **Accounts Groups –** Security groups can serve as resource groups that are used to grant access to resources.
- **Resource Groups –** Security groups can also serve to aggregate security principals to collectively authorize access to resources, which is usually achieved by making an account group a member of a resource group.

Security groups play a major role in implementing an organization's IT authorization strategy, and thus security group management should not be taken lightly.

**Resource Management**
An organization's IT resources typically include end-user workstations, servers, and applications and resources that are running or hosted on member servers, and in some cases on end-user workstations. All of these collectively constitute an organization's IT resources, and all require administrative coverage. Resource management thus plays an important role in managing an IT infrastructure and Active Directory plays a central role in facilitating resource management.

Resource management involves providing the means for administrators to:

- **Manage end-user workstations –** Provide general IT support, including installation and troubleshooting.
- **Manage individual servers –** Manage the physical computer, including monitoring, hardware maintenance, and installation.
- **Manage resources hosted on servers –** Manage the service aspects of resources and applications (including business applications and portals) that are hosted on one or more servers.

The delegation capabilities of Active Directory can be used to meet the administrative needs of organizations of all sizes. Active Directory can facilitate the distribution and delegation of administrative responsibility for managing all aspects of an organization's IT resources, taking into account administrative needs specific to the organization.

**Workstation Management**
Workstations that are used in day-to-day operations require administrative support. Administrators or operators who are responsible for workstation management need administrative access to these workstations. The access they require is provided as a part of delegating administrative authority for data management.

**Server Management**
In every IT organization, servers play a major role in providing essential services, from serving files to hosting databases and business applications. For example, one set of servers might collectively play the role of file servers, while another set hosts a database application such as Microsoft® SQL Server™. A set of servers might host parts of a distributed application that provides middle-tier functionality, while another set of servers might host internal Web portals for business applications. These servers and the applications and services that they host need to be managed.

As part of resource management, administrators who are responsible for managing computers need appropriate access to these servers, including the ability to monitor the health of the servers, manage event logs, modify hardware and install drivers, install service packs and hot-fixes, and perform other similar tasks.

**Server Resource Management**
Services and applications hosted on a collection of one or more servers also need to be managed and require administrative coverage, often by personnel different from those who manage the physical computer. Managing the service aspects of the applications that are hosted on servers includes administrative tasks such as monitoring the service, performing configuration changes to the service, and ensuring that the service continues to be delivered as specified by service level agreements. These administrators might also require the ability to authorize access to the services that are provided by these servers. For example, administrators for a Web application might require the means to authorize the set of end-users that have access to the application.

**Application-Specific Data Management**
Active Directory–integrated and –enabled applications, such as Microsoft Exchange Server, usually store application-specific data in Active Directory. As a part of managing their applications, administrators for these applications usually require administrative access to their application-specific data stored in Active Directory. Additionally, administrators of these applications might also require the ability to authorize the users of their applications to access and/or modify their application data. Application-specific data management involves facilitating the required access for all stakeholders of an Active Directory–integrated or –enabled application and typically involves delegating management of application-specific data to an application's stakeholders.

**Recommended Roles for Data Management**
A particularly effective method for creating an efficient administrative delegation model is to use administrative roles. A model that allows IT departments to delegate administrative responsibilities on the basis of business functions and administrative scopes lets organizations focus on business processes as opposed to technology processes. Defining these functions and scopes in terms of administrative roles enables business-driven administrative control while affording the ability to securely scale administration.

A roles-based approach to delegation makes management of delegated responsibility more tractable and increases the security of Active Directory. Additionally, it allows you meet additional needs for delegation in a simple yet deterministic fashion. Thus a roles-based approach to delegating data management is highly recommended.

Microsoft has engineered a set of recommended roles for delegating data management. These role recommendations take into account the sets of logically related administrative tasks and the security sensitivity and impact of these tasks.

The following is the set of recommended roles for delegating service management:

- Business Unit Admins
- Account Admins
- Workstation Admins
- Server Operators
- Resource Admins
- Security Group Admins
- Help-Desk Operators

- Application-Specific Admins

Depending on its specific administrative needs, an organization might choose to create and implement a delegation model based on Microsoft-recommended roles or a set of custom roles (which might or might not be based on Microsoft recommended roles) defined by the organization.

**Business Unit Admins Role**
It is not uncommon for multiple business units to participate in a shared Active Directory environment. Each business unit, while participating in a shared Active Directory environment, can have its own domain data and IT resources. Each business unit should be assigned a data owner who should have overall responsibility for all aspects of data management for that business unit.

Each business unit data owner should have a Business Unit Admins administrative role representing the operational arm of data owners. This role should be assigned at least one administrator and no more than a small number of administrators. This administrative group will have complete authority over all business unit data in the directory. Administrators in this role are the business unit's highest-ranking data administrators and are responsible for implementing and maintaining the administrative delegation model for business unit data management. Business Unit Admins also work closely with data owners during the creation of the OU structure to advise about how features such as inheritance of permissions and Group Policy affect the OU structure.

Typically, organizations will choose to first create and implement their service management delegation model, following which the data management model will be implemented. As a part of this process, the service owners should hand over responsibility for Active Directory data management to data owners. During the transfer of responsibility for data management to the data owner, the service owner should create one instance of this role for every business unit. Depending on the size of the business unit, this role should be assigned to no more than a few administrators. Administrators in this role delegate responsibility for the business unit by creating an OU hierarchy and creating and populating administrative groups to manage each OU in the hierarchy.

**Account Admins Role**
Every business unit has user accounts that need to be created, managed, and supported. Microsoft recommends the role of Account Admins for providing administrative coverage to manage user accounts. Responsibilities for user account management include creating accounts, populating account attributes, managing and maintaining accounts, and deleting accounts. Responsibility for account support should ideally be assigned to the Help Desk Operators, thereby removing that burden from the Account Admins role.

The Business Unit administrator of each business unit is responsible for creating one or more instances of the Account Admins role to provide administrative coverage for all aspects of Account management for all business unit user accounts, depending on the OU structure and the distribution of accounts among sub-units that the business unit might contain.

**Workstation Admins Role**
The Workstation Admins role is recommended for managing workstations. Depending on the administrative model, a business unit can have one or more instances of this role. For example, a business unit might be spread across multiple locations and might staff each location with a local administrative group that is responsible for managing these workstations. In this case, each local group must be delegated responsibility for managing all workstations for the business unit in that location. The number of members in each role depends on the requirements for each role instance and is typically a function of the number of workstations that need to be supported by a specific role instance.

**Server Operators Role**
Microsoft recommends the role of Server Operators to provide administrative coverage for managing an organization's server computers. Administrators assigned to the role are typically responsible for managing servers.

**Resource Admins Role**
Microsoft recommends the Resource Admins role for facilitating administrative coverage of a collection of one or more servers and the common services hosted on this set of one or more servers. Administrators who are assigned to this role are responsible for managing the resource (application, database, or files) and the set of servers on which the resource is hosted. For instance, an organization might choose to create a cluster out of a small number of physical servers, and this cluster of servers might serve as a single virtual file server. Administrators who have been granted responsibility for managing this single virtual file server will require the ability to manage each of these servers and manage the virtual file server. One instance of this role should be created for every such resource in the business unit.

In some cases, depending on the administrative model, the group of administrators who manage the servers is different from the group of administrators who manage the services that are hosted on these servers. This role can also be used to facilitate management of only the services hosted on one or more servers. Thus, the meaning and the nature of tasks assigned to

specific instance of this generic role can differ from one instance to another.

**Security Group Admins Role**
Microsoft recommends the role of Security Group Admins for managing all security groups required to meet the authorization needs of a Business Unit. Administrators in this role are responsible for creating and managing both account and resource security groups, and delegating the ability to manage the membership of some of these groups while retaining control over group membership of other security groups, as needed.

Microsoft recommends the implementation of one instance of this role per business unit. The reason that only one instance of this role is recommended is that any individual who can create more than 1015 security groups can misuse his or her authority or be coerced into launching a denial of service attack that could impact the ability of every user in the forest to log on.

Administrators for the various IT resources will typically require resource groups for the purpose of being able to authorize access to their resources and will also typically require account groups for the purpose of being able to aggregate a set of users to whom access can be then granted. IT resource administrators should request the creation and provisioning of required account and resource groups from the administrators in the Security Group Admins role. Administrators in this role are responsible for validating requests and then creating and provisioning the requested security groups so as to enable resource administrators to be able to use the requested groups.

**Application-Specific Admins Roles**
Active Directory–enabled or –integrated applications might store application-specific data in different places in Active Directory, and each such application might differ in its administrative requirements for managing application data stored in the directory. For example, certain applications might extend the class definition of user objects and store their application-specific user information on those extended user objects, while other applications might create their own OUs and store application-specific data in those application-specific OUs.

In either case, owners of such applications will typically request administrative access for their administrators to be able to manage application data. Microsoft recommends the use of Application-Specific Admins role to facilitate the access required by administrators of such applications. Use a separate custom Application-Specific Admins role for every application that stores application-specific data in Active Directory and whose administrators need administrative access to this data.

In addition to the administrators of the application, you might have to provision access for the following accounts which will typically require the ability to perform operations on the application data:

- Instances of the application, running in some security context (LocalSystem or a service account)
- Users of the application

When creating the delegation model, the service owner for each such application makes a request to the data owner to create the role. The service owner specifies the location of the application-specific data, the accounts that require access to the data, and the scope of the access. The data owner relays this information to the Business Unit administrator, who then facilitates this access by implementing instances of this role.

**Custom Roles**
Microsoft-recommended roles are intended to facilitate the creation and implementation of a well managed and efficient Active Directory data management delegation model. They are by no means the only way to address the delegation requirements of an organization.

In addition to the roles recommended by Microsoft, organizations can choose to create custom roles to address unique administrative requirements. Creating a custom role involves the following steps:

1. Understand and document the purpose of the new role.
2. Assign a set of administrative tasks to this new role.
3. Determine the minimal and precise set of permissions required to delegate the set of administrative tasks identified earlier in this chapter.
4. Document the general scope where permissions must be applied in the directory.

After making these decisions, you can determine the number of instances that are required, create the groups to represent each instance of the role, assign permissions to enable the role, and populate the groups with the appropriate administrative

user accounts.

# Determining Data Management Stakeholder Needs

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order to provide administrative coverage for all aspects of data management, you first have to understand your data management administrative needs. This involves evaluating administrative requirements specific to your organization for each aspect of data management.

The following approach is useful in understanding and documenting the administrative needs of all data Identify the number of business units participating in the domain.

1. For each business unit, determine the following management stakeholders:

    - **User Accounts.** Determine the total number of accounts and the number of administrative groups involved in account management or that require the ability to manage accounts.

    - **Workstations.** Determine the total number of workstations, the number of locations, and the number of administrative groups involved in workstation management or that require the ability to manage workstations.

    - **Member Servers.** Determine the total number of servers, the number of locations, and the number of administrative groups involved in server management or that require the ability to manage servers.

    - **Resources.** Resources can include Web, file, application, or database servers or services. Determine the total number of servers and services, the number of locations, and the number of administrative groups involved in resource management or that require the ability to manage resources.

2. For each business unit, determine the Group Policy requirements.

# Creating the Data Management Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The data owners of each business unit should create an administrative delegation model to distribute and delegate administrative responsibility among the data administrators.

Creating a delegation model involves the following steps:

1. Determine your role requirements and determine the number of instances and scope of administrative authority of each data management role.
2. Plan an OU hierarchy based on the administrative delegation and Group Policy requirements of the business unit.
3. Assign specific administrators or administrative groups to the role instances.
4. Document the model.

At the end of this process, business unit data owners should have a documented administrative delegation model that ensures the following:

- Administrative coverage has been provided for all aspects of content management.
- Administrative responsibilities have been distributed according to sound security practices among the instances of the various roles.
- Administrators and administrative groups have been mapped to the various instances of the data administration roles.

Responsibilities for each category of data administration tasks are delegated by using administrative roles that are instantiated as security groups. As mentioned earlier, organizations can choose to create their delegation model around Microsoft recommended roles or around custom created roles which might or might not depend on Microsoft-recommended roles.

**Determining Role Requirements**
The first step in the creation of your delegation model is to determine the following:

- The number of instances and the scope of administrative authority of each data management role.
- The specific set of administrators that will be assigned to each role instance.

**Required Data Management Roles**
The following are the recommended set of roles for Active Directory data management:

- Business Unit Admins
- Account Admins
- Workstation Admins
- Server Operators
- Resource Admins
- Security Group Admins
- Help-Desk Operators

- Application-Specific Admins

**Business Unit Admins**
One instance of this role will be required for every business unit.

**Account Admins**
One instance of this role will be required for every administrative group that needs the ability to manage user accounts. For each such administrative group, document the scope of authority by documenting the set of users that should be managed by this group. Also, for each instance, determine the set of administrative tasks that incumbents of the role should be allowed to perform.

The following is a list of common administrative tasks that belong to this category:

- Create user accounts
- Delete user accounts
- Move user accounts
- Reset a user's password
- Unlock user accounts
- Modify non-security sensitive information, such as phone numbers
- Modify security-sensitive information
- Examples of security-sensitive information include whether or not a password is required on the account, whether the account is trusted for delegation, and modifying the UPN of the user account.
- Authorize access to the account

For each instance, evaluate the needs for and document the set of administrative tasks that incumbents of the role should be allowed to perform. This information will be used during the implementation phase to delegate each of these administrative role instances.

**Workstation Admins**
One instance of this role will be required for every administrative group that needs the ability to manage workstations. Most organizations have one or more administrative groups for every location that houses workstations. For each such administrative group, document the scope of authority by documenting the set of workstations that should be managed by this group.

Workstation management typically includes being able to manage all aspects of workstations; thus, incumbents in this role typically are Local Administrators on member workstations.

**Server Operators**
One instance of this role will be required for every administrative group that requires the ability to manage servers. Most organizations have one or more administrative groups for every location that houses servers. For each such administrative group, document the scope of authority by documenting the set of servers that should be managed by this group.

Server management typically includes being able to manage all aspects of member servers; thus, incumbents in this role typically are Local Administrators on member servers.

**Resource Admins**
One instance of this role will be required for every administrative group that needs the ability to manage resources (services or applications being hosted on a set of one or more servers and collectively managed by a single group of administrators). For each such administrative group, document the scope of authority by documenting the set of servers and the specific service that constitute the resource that should be managed by this group.

Resource management typically includes being able to manage all aspects of the service and manage all aspects of the servers

on which the service is being hosted, although this might not always be the case. In some cases, the responsibility might be restricted to being able to manage only the service that is being hosted on or more servers.

**Security Group Admins**
There should only be one instance of this role for every business unit. Depending on the specific administrative needs of the business unit, the business unit data owners might decide to create more than one instance of this administrative role. For each instance determine the set of administrative tasks that incumbents of the role should be allowed to perform.

The following is a list of common administrative tasks that belong to this category:

- Create security groups
- Modify the membership of security groups
- Modify other properties, such as managed-by information
- Delete security groups

For each instance evaluate the need for and document the set of administrative tasks that incumbents of the role should be allowed to perform. This information will be used during the implementation phase to delegate each of these administrative role instances.

**Help Desk Operators**
One instance of this role will be required for every administrative group that needs the ability to provide account support to a subset of user accounts or to all user accounts. For each such administrative group, document the scope of authority by documenting the set of users for whom this group will provide account support.

The following is a list of common administrative tasks that belong to this category:

- Reset passwords on user accounts
- Unlock user accounts
- Optionally, a business unit might also decide to grant its help desk operators the ability to modify certain attributes on user objects; these attributes will typically be non-security sensitive attributes that users cannot themselves modify

For each instance, evaluate the need for and document the set of administrative tasks incumbents of the role should be allowed to perform. This information will be used during the implementation phase to delegate each of these administrative role instances.

**Application-Specific Admins**
One instance of this role will be required for the administrative group of every Active Directory–enabled or Active Directory–integrated application that requires the ability to manage application-specific data stored in Active Directory. For each such administrative group, document the scope of authority by documenting the set of application-specific data that this group will require the ability to manage.

The nature of administrative tasks in this category is typically a function of the specific Active Directory–enabled or Active Directory–integrated application deployed and the specific administrative requirements of that application. The administrators of these applications should have a good understanding of these administrative requirements. The business unit data owners should understand these requirements and plan for facilitating the required access to enable these administrators to carry out their administrative tasks.

**Assigning Administrators to Role Instances**
For each instance of every data management role required, document the identities of all administrators who will be assigned to these role instances.

**Planning an OU Hierarchy**
Planning an OU hierarchy based on the administrative delegation and Group Policy requirements of the business unit is one of the first steps in creating a delegation model. A well-designed OU hierarchy should adequately address an organization's Group Policy requirements while also facilitating the delegation of administrative authority based on the organization's administrative requirements.

This section provides recommendations for creating a well-designed OU hierarchy.

**An OU for every Business Unit**
In the section "Implementing Your Data Management Delegation Model" later in this document, as part of the handoff process wherein service owners hand off data management delegation to data owners, a single high-level OU named Business Units will be created with the purpose of creating a subtree to contain a single high-level OU for each business unit. That way, all domain data can be contained within this one OU subtree rooted at the Business Units OU instead of at the Domain root. This recommended approach has the following advantages:

- This facilitates delegation of administration of each business unit to the respective business unit administrator while allowing the specification of any permissions that might be required to delegate administrative authority over data across all business units at the Business Units OU instead of at the Domain root. This prevents the inheritance of any ACEs applied in the System container and in the DNS containers, thereby substantially reducing the number of ACEs, which might impact the size of that database significantly in Windows 2000. In Windows 2003, the concept of single-instances makes this a non-issue. For example, if the administrative needs of Contoso Pharmaceuticals required that a single centralized administrative group be delegated the authority to provide Account support for all business units in the domain, their administrators would typically grant the security group representing the operators in the Account support role inheritable permissions expressed in three or four ACEs applied on the domain root object. While this would have the effect of the ACEs flowing down to all user objects in all business units, it would also result in these ACEs flowing down on all DNS record objects in the DNS container; assuming that there are about 5000 records in the DNS container, this approach would result in 5000 additional ACEs in Windows 2000, significantly impacting the database size. On the other hand, if these inheritable ACEs were applied on the Business Units OU in the recommended approach, they would only be inherited by objects stored in the subtree rooted at the Business Units OU.

- This also increases the manageability of domain data by eliminating the need to create multiple OUs that are peers of the OUs that exist by default. It increases manageability because it leaves the top-level OU hierarchy unchanged with the exception of the addition of a single OU, thereby minimizing the possibility of confusion.

Each business unit is then represented by an OU directly under the Business Units OU, and as is presented in the Implementing Your Data Management Delegation Model section, the handoff process also involves granting the Business Unit Admins of each business unit full control over their OUs. Business unit data owners and Business Unit Admins are then free to create their own OU sub-structure to manage business unit content.

The next issue to consider is the creation of an OU structure for the business unit so that responsibility for managing business unit content can be delegated while also taking into account Group Policy requirements. The general approach is to first create an OU sub-structure for the business unit based on delegation of administration requirements, and then apply Group Policy on existing OUs or add additional OUs, possibly in various parts of the OU tree, for the purpose of applying Group Policy. There should be a well-understood reason for the presence of every OU in the OU structure.

**An OU for Storing Delegated Roles**
An organization that plans to create and implement a roles based delegation model will require an OU to store the security groups representing the various roles. Dedicating one OU to storing the security groups allows the Business Unit Admins to easily identify all delegated roles and have a single point of administration to control the set of administrative roles and the set of administrators assigned to each administrative role. Thus it is recommended to create an OU dedicated to storing all administrative roles, as a child object of the OU representing the Business Unit.

**Creating the First Level of OUs**
The following recommendations are aimed at creating an OU structure that meets your organizations needs while remaining simple and flexible, thereby allowing room for change, extensibility, and any modifications that might be needed to address changing administrative needs. To create the first level of OUs, perform the following tasks:

1. Create a single OU under the Business Unit OU for storing all business unit user accounts.

   You should plan on storing all your user accounts in this OU. Recommendations for how to create an OU hierarchy within this OU to address specific administrative requirements for account management are provided later in this section.

2. Create a single OU under the Business Unit OU for storing all business unit security groups

   You should plan on storing all security groups required to meet the authorization needs of your IT resources in this OU. Recommendations on how to create an OU hierarchy within this OU to address specific administrative requirements

around account management are provided later in this section.

3. Create a single OU under the Business Unit OU for storing all business unit user resources

   You should plan on storing computer accounts for all your workstations and servers within this OU. Recommendations on creating an OU hierarchy within this OU to address specific administrative requirements around workstation and server management and facilitate delegation of resource management to specific administrative groups are provided later in this section.

Figure 2 shows a representative OU structure based on the recommendations presented earlier in this section. Note that Product Development and Business Management are the two Business Unit OUs, one for each of the two business units in the fictitious organization represented in the example.




The OU structure thus far can be used to perform the following delegations:

- An organization can delegate account management of all business unit accounts to a single administrative group by granting appropriate permissions on the Accounts OU. This is sufficient for a small organization with a single location and a relatively small number of users.

- An organization can delegate workstation management of all business unit security groups to a single administrative group by granting appropriate permissions on the Security Groups OU. This is sufficient for a small organization with a single location and needing relatively small numbers of security groups to address its authorization requirements.

- An organization can delegate overall resource management of all business unit IT resources to a single administrative group by granting appropriate permissions on the Resources OU. This is sufficient for a small organization with a single location and a relatively small number of IT resources to manage.

- An organization can delegate account support for all business unit user accounts to a single administrative group by granting appropriate permissions on the Accounts OU. This is sufficient to meet the needs of an organization of any size that has a single administrative group that provides centralized account support.

**Creating OUs to Delegate Account Management**

For most small organizations where there is only one group responsible for account management, the OU structure for account management described earlier in this section should be sufficient to delegate account management to that administrative group.

Depending on your administrative needs for delegating account management, an OU subtree can be created under the Accounts OU to facilitate distribution of account management responsibilities among various groups.

For every Account management role identified earlier in this section that requires the ability to manage a specific set of users, create a single OU under the Accounts OU and store in this OU the user accounts of the specific set of users that this group is responsible for managing. For example, if the business unit contains four divisions, and each division requires the ability to manage its own users, create one OU for each division and store all user accounts that should be managed by the administrators of a specific division in the OU corresponding to the division.

The creation of an additional level of OUs will enable the delegation of account management to the Account Admins of each division. In some cases, the Account Admins of a division might want to sub-delegate account management to other Account Admins within the division, one for each sub-division. An organization can create an additional set of OUs under each divisional OU to sub-delegate account management to other Account Admins within the division.

In general, an OU typically corresponds to a unit of delegation, in that it can be used to apply a set of permissions that become effective on all objects in the subtree rooted at the OU. From that perspective, a new OU or a new level of OUs can be added to facilitate delegation of administrative responsibility.

The following are some common scenarios of administrative requirements that can be addressed by the recommendations described earlier in this section:

- Assume that your organization is spread across multiple locations and your IT model requires delegation of account management to locally based administrative groups. In this scenario, you can create an additional set of OUs, one for each location, and then store user accounts in the OU representing the physical location where the user account is based. You can then distribute and delegate account management by specifying permissions on these OUs.

Figure 3 illustrates the three-location OU structure that accommodates user accounts that must be managed separately in Locations A, B, and C.




Each location in Figure 3 has a different set of administrators in the role of Account Admins. To implement this scenario, the Business Unit Admin creates four instances of the Account Admins role, each represented by a unique security group. The Business Unit Admin can grant either Full Control over their respective OUs or specific permissions to delegate specific administrative tasks within the Accounts OU to each instance of Account Admins.

- Assume that your IT model requires that different categories of users be managed by separate administrative groups. For example, suppose your IT model requires that all user accounts for contracting employees are to be managed by one administrative group while all user accounts for full-time employees are to be managed by another administrative group. You can create an additional level of OUs, one representing contracting employees and the second representing full-time employees.

- Taking this one step further, suppose that an additional constraint requires that all user accounts for all company executives are to be separately managed by yet another administrative group. In this case, you can create a third level of

OUs under the OU representing full-time employees, separate regular full-time employees from executives, and accordingly grant permissions on both OUs to delegate account management to different groups.

- You can use the same OU structure to apply a different Group Policy to contractors and full-time employees. Thus, most of the time, you can also derive the additional benefit of being able to meet your administrative requirements and address Group Policy requirements.

- Assume that you have no requirements for further distributing account management. You can still create a second level of OUs to meet Group Policy application needs. For example, assume that corporate policy dictates that a stricter user policy be applied to user accounts of all users who travel extensively. You could create a second level OU structure for the purposes of applying Group Policy to a specific set of users.

These scenarios are only meant to illustrate the creation of an OU structure for addressing administrative and policy enforcement requirements. Other administrative requirements similar to those illustrated here can also be similarly addressed by the creation of an appropriate OU structure.

**Creating OUs to Delegate Security Group Management**
Most organizations can benefit from the creation of an additional level of OUs within the OU assigned for storing and managing security groups. A well-designed OU subtree within the Security Group OU can also make managing security groups for administering IT resources more tractable.

The following is a recommended structure for OUs within the Security Group OU

Create the following OU structure under the Security Group OU:

- An OU for storing common account groups; a suggested name is Common Account Groups.

- An OU for storing Account and Resource groups for specific IT resources; a suggested name is Resource Specific Groups.

Use the Account Groups OU to store and manage common account groups for your business unit. Examples include Business Unit Users, Location A, Division B Users, Division C Managers, Business Unit Executives, and Business Unit Executive Assistants.

- Use the Resource Specific Groups OU to manage resource-specific account and resource groups as follows. Within the Resource Specific Groups OU, perform the following steps to create the OU structure:

- Create one OU for the different kinds of resources in your organization. Examples include Web-Servers, File-Servers, Application servers, and Database servers.

- Within each resource-specific OU, create one OU for each specific resource. For example, within the OU for Web server security groups, create an OU for the Web servers hosting the Human Resources Web application and another OU for the Web servers hosting one of the Intranet Web sites.

Within the OU for a specific resource, create one OU for storing the account groups and another OU for storing the resource groups that resource administrator might require to authorize access to the resource. For example, the administrator of a Web portal will require at least multiple resource groups to specify who can access the Web portal and to what extent. The resource administrator will also require the ability to specify a collection of users that can access the Web portal in a particular way. For instance, the Web portal administrator might use an account group FT to collect all full-time employees and a separate account group CT to collect all contractors. The administrator would then make the FT group a member of the resource group that was used to specify to the Web application all users who have unrestricted access to all parts of the portal, while making the CT group a member of the resource group that was used to specify to the Web application all users who have restricted access to all parts of the portal.

The OU structure recommended earlier in this section can be used to simplify security group management and make it more tractable.

Figure 4 shows an example OU structure that collects account groups in the Account Groups OU and separates resource groups into OU trees according the resources. This OU structure associates the groups with their respective resources for easier management.




Note that the recommendations described earlier in this section are by no means the only way to facilitate management of security groups. Organizations should take their specific administrative needs into account when designing their OU structures. For example, one alternative to storing resource specific security groups in the Security Groups OU is to create an OU to store all security groups for a specific resource under the specific OU representing this resource in the Resource OU subtree, as presented later in this section. This has the additional advantage of storing the security groups associated with a resource along with the servers representing the resource, thereby storing all aspects of a resource in one small subtree.

What is important is to understand how an OU enables delegation of administrative authority and to apply this understanding to the creation of an OU structure that addresses your organization's administrative delegation and Group Policy requirements.

**Creating OUs to Delegate Resource Management**

The benefits of a well-designed OU structure are most evident in the delegation of resource management. This section describes a recommended structure to facilitate delegating resource management.

Note that the recommendations presented here are by no means the only way to facilitate resource management. Organizations should take their specific administrative needs into account when designing their OU structures.

To create the recommended OU structure for delegating resource management, within the Resources OU, perform the following steps to create the OU structure:

- Create one OU to store all workstation accounts; a suggested name is Workstations.
- Create one OU each for every type of resource in your business unit; examples include file servers, Web portals, applications and databases.

**Workstation management**

The Workstations OU can be used to store computer accounts for all workstations in your business unit. An additional level of OUs can be used to address specific administrative requirements. Workstation management is typically delegated to local administrative groups. Thus, an organization could create an additional level of OUs, one for each physical location where the business unit has a presence, and all workstations for a specific physical location can be stored in the OU representing that physical location. Such an OU structure can be used to easily delegate workstation management to all administrative groups. In general, the various instances of the Workstation Admins role will usually correspond to the OUs under the Workstations OU; this is because in general, each OU serves to separate workstation accounts for the purpose of provisioning management of

these accounts by a separate administrative group.

Delegating workstation management typically involves granting complete control over the computer account representing the workstation and using the restricted groups feature of Group Policy to make an administrative group a member of the administrators group on these workstations. Since OUs can facilitate both the granting of full control and the application of Group Policy, they can be used to delegate workstation management.

Figure 5 shows the OU structure that illustrates this decentralized approach to managing workstations.




**Resource Management**

A well-designed OU structure can also be used to manage servers and to delegate management to the respective administrative groups of the various IT resources that are hosted by a collection of servers. The recommendation earlier in this chapter to create one OU for every type of resource under the Resources OU aids in increasing manageability of resources.

Under the OU created for each resource type, introduce an additional level of OUs with each OU in this level representing a specific resource. For example, under the Web Servers OU, create a specific OU for storing the computers of all servers on which the Human Resources Web portal is collectively hosted. Figure 6 illustrates such an OU structure.



The OU hierarchy structure for resource management that is recommended earlier in this chapter offers the following benefits:

- It allows you to delegate administrative authority for all resources of a specific type to a specific administrative group. For example, you could delegate overall management of all file servers to an administrative group responsible for overall file-server management by granting sufficient permissions on the File Servers OU. This group will then be able to manage all file servers within the organization.

- It allows you to delegate administrative authority of specific IT resources to only those administrative groups that have been assigned the responsibility to manage that specific resource. For example, you could specify appropriate permissions on the various OUs representing specific file servers to delegate management of those file servers to specific administrators.

These recommendations are by no means the only way to facilitate resource management. Organizations should take their specific administrative needs into account when designing their OU structures.

After you have created your OU hierarchy structure based on your organization's administrative requirements and needs, you should then consider your organization's Group Policy application requirements.

**Considering Group Policy Application Requirements**
An OU hierarchy structure should be based on administrative delegation requirements and Group Policy application requirements. Group Policy can be associated with OUs and is applied onto all user and computer accounts stored in the OU. Thus, Group Policy application requirements also impact the OU structure.

It is recommended that the OU hierarchy be first designed based on delegation of administration requirements, and then be appropriately modified to address Group Policy application requirements. The following examples illustrate this guideline. Assume that an organization has decided on the OU structure presented in Figure 7.




**First Example**

In the first example, assume that Group Policy requirements for this organization include that a location-specific Group Policy be applied to all computers based on their location, irrespective of the role of the computers. To accommodate this requirement, the OU hierarchy structure in Figure 7 can be modified as follows:

- In the Resources OU, directly below the Resources OU, introduce a single level of OUs, one for each physical location in the business unit. Within each such OU representing a physical location, recreate the same set of OUs as initially designed. Thus, each OU representing a specific location should have one child OU for Workstations and one child OU each for the different types of resources within the business unit. Each of these OUs in turn should continue to have the same sub-structure as was initially created. The primary difference in the modified OU hierarchy structure is that an additional level of OUs has been introduced for the sole purpose of applying Group Policy. Note that these OUs are not being used for delegating administrative authority, although nothing prevents using them for this additional purpose if necessary.

### Second Example

In the second example, assume that Group Policy requirements for this organization include that a location-specific Group Policy be applied to all workstations, based on their location:

- To accommodate this requirement the OU hierarchy structure will not need to be modified, because an organization might already typically have in place a level of OUs separating workstation accounts based on their location. These OUs can then be used to apply Group Policy to workstations based on their location

### Third Example

In the third example, assume that Group Policy requirements for this organization include that a specific Group Policy be applied to all users that belong to a specific division within the business unit:

- In this case, depending on the OU structure within the Accounts OU, an additional level of OUs could be introduced to separate the accounts by division, if they are not already separated, and Group Policy can then be applied to the specific OUs representing each division with the business unit.

These examples illustrate how an OU hierarchy model primarily based on delegation of administration requirements can be suitably modified to address Group Policy application requirements.

Having gained an understanding of how OUs can be used to both delegate administrative authority and meet Group Policy application requirements, you can proceed to create an OU hierarchy based on your administrative delegation requirements and your Group Policy application requirements. After you have created the OU hierarchy, you are ready to document your delegation model and move to the implementation phase.

**Documenting Your Delegation Model**

Before proceeding to the implementation phase, be sure to document your administrative roles requirements, your OU hierarchy, and your Group Policy requirements. Table 6 is a template that you can use to document your role instances.

**Table 6 Template to Document Specific Role Instances (With Example Values)**

| Field | Assignment Information |
|---|---|
| Instance of: | Account Admins |
| Instance Number: | 1 of 4 |
| Assigned Administrators: | Clair Hector |
| Assigned Tasks: | • Create user accounts<br>• Manage all properties<br>• Control access to accounts<br>• Delete user accounts<br>• Sub-delegate<br><br>Refer to document located at: *document location* |
| Security Group | To be filled during Implementation phase |
| Permissions Assigned | To be filled during Implementation phase |
| Notes | Sub-delegation allowed<br><br>Most group members based out of Chicago Group contact: (555)-123-4567 e-mail: DIAccAdm |

In addition, for each OU in the OU hierarchy structure, document the following:

- **Primary purpose –** Delegation and/or Group Policy application.
- **Secondary purpose (if any) –** Also being used for Delegation and/or Group Policy application.
- **Details –** Document all necessary details.
- **Administrative Contact –** Provide administrative contact details.

# Implementing Your Data Management Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After the data owners of each business unit have created their administrative delegation models, Business Unit Administrators should implement their respective delegation models. However, before Business Unit Administrators can delegate their models, service owners need to hand off data management to data owners and provision the Business Unit Admin roles.

Implementing the delegation model involves setting permissions on Active Directory objects. Active Directory ships with a set of delegation tools that can be used to apply and manage permissions in Active Directory. For a list of these administration tools and for details about how to use them, see Appendix G: Active Directory Delegation Tools in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document. It is recommended that you invest the time to gain familiarity with these tools, as these tools are required to manage permissions. One alternative to using these tools is to use scripting to manage permissions.

**Handing Off Data Management to Data Owners**

Typically, organizations will choose to create and implement their service management delegation model, following which the data management model will be implemented. As a part of this process, the service owners should hand over responsibility for Active Directory data management to data owners by handing off data management to the respective business unit administrators, who represent the data owners. These high-level data administrators have overall responsibility for managing business unit content in Active Directory. They represent the data owners in implementing the respective business unit delegation models.

High-level service administrators represent the service owners in performing the tasks that are required to hand off data management to data owners. To hand off data management to business unit data owners, a member of the Domain Configuration Operators role or the Domain Admins group performs the following tasks in each domain:

1. Create a new OU as an immediate child of the domain object. A suggested name for this OU is Business Units. It serves as a high-level container for all domain content.

2. Within the Business Units OU, create one OU for each participating business unit, and name them appropriately following a standard convention.

3. Within the Business Units OU, create an additional OU for storing the security groups that will be used for representing the Business Unit Admins administrative role instances and instances of any other role that will require access to all business unit OUs. A suggested name for this OU is Delegation.

4. Within the Delegation OU, create one security group for every instance of the Business Unit Admins role; that is, create one domain local group for each business unit. Name each security group appropriately: a recommended name format is BusinessUnitName Admins.

5. On each of the security group objects that you create in step 4, grant the following permissions to the security group account itself:

    a. Property permission Read All Properties

    b. Property permission Write Members

6. On each business unit OU object, to the security group that represents the Business Unit Admins role for that business unit, grant Full Control.

7. Create one user account for each Business Unit Admins administrator in his or her respective business unit OU (these accounts will be moved to an appropriate OU within the respective Accounts OU by the business unit administrator after delegation of the business unit OU is implemented).

8. Add each user account to the respective security group that represents the Business Unit Admins role instance for the business unit.

Figure 8 shows a prototype domain-level OU structure that is created by a service administrator for handing off data administration to data owners in a domain that has two business units.




This completes the handoff process of transferring responsibility for data management to data owners, who are represented by respective Business Unit Admins, the operational arm of the data owners for the various business units sharing the domain.

Depending on the size of the organization, it might very well be that there is only one business unit in the entire domain. The handoff described in this section is beneficial in that situation too, as it creates an extensible OU structure which can easily accommodate the addition of new business units to the domain should the need arise.

**Implementing Business Unit Delegation Models**

Handing off the responsibility for data management results in the creation of one high-level OU for every business unit. It also involves enabling the Business Unit Admins roles and leaves Business Unit OUs ready to be managed by their Business Unit Admins. Business Unit Admins now have complete control over their Business Unit OUs and can proceed to implement their delegation models. This section provides recommendations for implementing an efficient and security-conscious delegation model.

The following are the recommended tasks involved in implementing the business unit delegation model:

1. Implement the OU hierarchy structure.

2. Create security groups to represent the various administrative role instances.

3. Enable Administrative Roles.

4. Delegate Administrative Roles.

**Implementing the OU Structure**

The first step in implementing the delegation model is to create the OU structure in the directory. This step is usually simple and straightforward: The Business Unit Admins create all OUs according to the documented OU structure for the business unit.

**Creating Security Groups to Represent Role Instances**

The next step in implementing the OU structure is to create security groups to represent all required administrative role instances as documented during the creation of the model. This step is also very straightforward and involves the following tasks:

1. For every required administrative role instance, the Business Unit Admins create one security group in the Delegation OU, a child of the Business Unit OU. These groups are typically Domain Local groups, as their scope of application is the Domain.

**Note**

> When specifying read access to specific attributes of, or list access to, an Active Directory domain object, avoid using domain local groups to set permissions if the attribute or attributes are included in the partial attribute set that is replicated to global catalog servers. Instead, consider using a global group. Users of domain local groups might not be able to access these attributes if they are outside the domain that hosts the global catalog to which they are bound. Note, however, that this affects only read and list access because Global Catalogs are Read-Only by design.

2. After creating each group, Business Unit Admins can optionally choose to grant these groups the ability to modify their own group memberships, thereby allowing administrators in these roles to add or remove other administrators without Business Unit Admin intervention. This step is purely a function of the administrative requirements of an organization, and an organization might choose to individually assess this option for every role instance.

**Enabling Administrative Roles**

The next step in implementing the delegation model involves enabling all the required instances of the various administrative roles. Enabling a role involves granting the security group representing each role instance the minimal set of permissions required to allow members of this administrative group the ability to perform all administrative tasks assigned to this administrative role instance.

**Enabling account management roles**

Every business unit will typically have at least one instance of the Account Admins role that is granted overall responsibility for managing all aspects of all user accounts. Additionally, other role instances can be delegated account management responsibilities for specific sets of users.

To delegate an instance of the Account Admins role, perform the following tasks:

1. Refer to the documented role instance descriptions and identify the scope of administrative authority. Then map this scope to an OU subtree in the Accounts OU. Based on this information, determine the point of delegation for delegating this role instance. The point of delegation is the usually the object whose DACL will have to be modified to grant the set of permissions required to delegate this role instance. Note that in some cases, the point of delegation might actually map to more than one object depending the nature of the responsibilities being delegated and whether the objects representing the target of delegation are stored in more than one location in the OU hierarchy.

   For example, assume that the Account Admins role that has been granted has overall responsibility for all user accounts in the Business Unit. The point of delegation in this case will be the Accounts OU itself because the subtree of objects rooted at this OU contains the entire scope of authority for this role, which in this case refers to all the user accounts in the business unit.

   As a second example, assume that the Account Admins role that has been granted has account management responsibilities for all accounts in Division A. The scope of administrative authority is thus all accounts in Division A. Mapping this to the OU subtree in the Accounts OU reveals that there is an OU for Division A. This OU contains (or will contain) all user accounts for Division A. Thus, the point of delegation for this role is the Division A OU under the Accounts OU.

2. Based on the documented role instance descriptions, determine the minimal set of permissions required to delegate all assigned administrative tasks by referring to Appendix A: Active Directory Administrative Tasks in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document. This appendix documents the minimal and precise set of permissions required to delegate every Active Directory administrative task.

   For example, assume that the business unit data owners decided to delegate only the following tasks to the Account Admins for the Division A:

   - Create user accounts
   - Modify all attributes that belong to the Personal Information property-set

   According to Appendix A: Active Directory Administrative Tasks, the following are the permissions required to perform these administrative tasks:

- Create user accounts
- Create Child on parent object
- Modify all attributes that belong to the Personal Information property-set
- Write-Property to the Personal-Information property-set

Thus, in order to delegate this role, the security group representing the Division A Account Admins role will need to be granted the following permissions:

- Create Child on parent object
- Write-Property to the Personal-Information property-set

Also, note the specifics of where these permissions are required. Since the Create Child permission is required on the parent object, which in this case happens to the be the Division A OU, there is no need to mark that permission as inheritable. On the other hand, since the Write-Property permission will actually be required on the user objects, this permission will need to be marked as inheritable.

3. Once you have determined the point of delegation and the minimal set of permissions required to delegate the various roles, modify the DACL on the object that corresponds to the point of delegation and grant the security group representing each administrative role the identified set of permissions.

Once these steps have been completed, all administrative roles for account management will have been delegated.

**Enabling security group management roles**
Your business unit should have only one instance of the Security Group Admins role that is responsible for security group management for the entire business unit.

To delegate an instance of the Security Group Admins role, perform the following tasks:

1. Refer to the documented role instance descriptions and identify the scope of administrative authority. Then map this scope to an OU subtree in the Security OU. Based on this information, determine the point of delegation for delegating this role instance.

2. Based on the documented role instance descriptions, determine the minimal set of permissions required to delegate all assigned administrative tasks by referring to Appendix A: Active Directory Administrative Tasks in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document. This appendix documents the minimal and precise set of permissions required to delegate every Active Directory administrative task.

3. Once you have determined the point of delegation and the minimal set of permissions required to delegate the various roles, modify the DACL on the object that corresponds to the point of delegation and grant the security group representing each administrative role the identified set of permissions.

   In the case of the single administrative role that has been assigned overall responsibility for security group management, this point of delegation will correspond to the Security Groups OU. Since this group will typically be assigned responsibility to manage all aspects of security group management, although individual permissions for every administrative task can be granted, this security group could be granted Full Control on all objects of the class Group.

   In some cases, the business unit might decide to delegate security group management to more than one group. In this case, for each role instance, determine the point of delegation and the cumulative set of permissions required to delegate these roles. Then grant the security groups representing these roles the identified set of permissions by modifying DACL on the object that corresponds to the point of delegation.

Once these steps have been completed, all administrative roles for account management will have been delegated.

**Enabling resource management roles**
Your organization will have typically multiple instances of each resource management role.

The implementation aspects of delegating resource management are a little different than those for account or security group management. This is primarily because delegating resource management usually involves using the delegation capabilities of Active Directory to make delegated administrators local administrators on member computers, in addition to granting these administrators the full-control over the computer account representing the computer. In cases where these administrative groups only require the ability to manage the service aspects of a service that they are responsible for, these administrative groups need not be made local administrators on these computers; it is sufficient to make them members of local groups that are authorized to manage these services.

Making administrative groups in Active Directory members of local administrative groups involves the application of a specific feature of Group Policy – restricted groups. The Restricted Groups feature allows an administrator to control membership by creating a list that restricts the members of a group and the groups to which a group can belong. The policy restricts membership by enforcing the list: any members who are not on the list are removed from the group; any security principals in the list that are not members of the group are added to the group.

For more information about restricted groups, see “Restricted Groups” in Help and Support Center for Windows Server 2003. For more information about restricted groups, see article 810076, "Updates to Restricted Groups ("Member of") Behavior of User-Defined Local Groups" in the Microsoft Knowledge base.

The following example illustrates the application of this feature to delegation of the workstation management role.

The Production business unit of Contoso Pharmaceuticals has operations in two physical locations – Chicago and Atlanta. 60 percent of the users of this business unit are based in Chicago and 40 percent in Atlanta. Every user is typically assigned one end-user workstation.

Contoso’s OU hierarchy thus has two OUs within the Workstations OU, one representing Chicago and the other Atlanta. Computer accounts for all workstations located in Chicago are stored in the Chicago OU. Accordingly, computer accounts for all workstations located in Atlanta are stored in the Atlanta OU. Contoso’s data owners have assigned the responsibility of managing all workstations in the Chicago location to a locally based administrative group called Chicago Workstation Admins. Members of this administrative group require the ability to manage all aspects of computer management for these workstations. They might be called upon to perform a variety of administrative tasks including, but not limited to, troubleshooting operating system issues, adding hardware, configuring drivers, and installing patches. Most of these tasks require administrative access to these workstations.

Contoso’s administrators easily facilitated making this administrative group local administrators on all workstations by performing the following steps:

1. Creating a security group called Chicago Workstation Admins to represent this role.

2. Granting this security group inheritable full-control on all computer objects by modifying the DACL of the Chicago OU within the Workstations OU.

3. Using the restricted groups feature of Group Policy to make this group a member of the Local Administrators group on all workstations in Chicago by applying Group Policy on the Chicago OU and specifying that the Chicago Workstation Admins be made a member of the Local Administrators group on all workstations whose computer accounts are present in the Chicago OU.

Enabling resource management administrative roles essentially involves repeating the steps involved in this example for all administrative roles. The following recommendations can be used to enable resource management roles.

Perform the following steps to enable the various instances of the Workstation Admins role:

1. Refer to the documented role description for this instance and identify the scope of administrative authority. Map this scope to an OU subtree in the Workstations OU. Based on this information, determine the point of delegation for delegating this role instance.

2. Grant the security group representing this role instance the following inheritable permission on the point of delegation:

    - Full-Control on Computer objects

3. Use the Restricted Groups feature of Group Policy by applying a GPO to the OU that represents the point of delegation,to make this make the security group representing the specific role instance a member of the local Administrators group on all member workstations that are stored in the OU.

Once these steps have been completed, all administrative roles for workstation management will have been delegated.

For each instance of the Server Operators role, perform the following steps to enable the instance:

1. Refer to the documented role description for this instance and identify the scope of administrative authority. Then map this scope to an OU subtree in the Resources OU. Based on this information, determine the point of delegation for delegating this role instance. Refer to your OU model to determine the specific OU created with the purpose of storing non resource specific servers. You might have more than OU if you created a high-level OU under the Resources OU to store non resource specific servers and possibly created an additional level of OUs within this OU, one for each physical location that has servers located and depends on local administrative groups for management of these servers.

2. Grant the security group representing this role instance the following inheritable permissions on the point of delegation:

   - Full-Control on Computer objects

3. Use the Restricted Groups feature of Group Policy by applying a GPO to the OU that represents the point of delegation, to make this make the security group representing the specific role instance a member of the local Administrators group on all member workstations that are stored in the OU.

Once these steps have been completed for all instances of the Server Operators role, all administrative roles for server management will have been delegated.

For each instance of the Resource Admins role, perform the following steps to enable the instance:

1. Refer to the documented role description for this instance and identify the scope of administrative authority. Map this scope to an OU subtree in the Resources OU. Based on this information, determine the point of delegation for delegating this role instance. Refer to your OU model to determine the specific OU created with the purpose of storing the set of servers that hosts this specific resource. In an OU hierarchy that was based on the recommendations in this document, the specific OU for this resource should be under the corresponding resource-specific OU.

   For example, when delegating resource management for the Human Resources Web portal that is hosted on a set of high-performance Web servers, there should be one OU under the Web-Servers OU (which is under the Resources OU) for the Human Resources Web portal. This OU should contain the set of servers that collectively hosts this Web portal. This OU is the point of delegation for delegating this role.

2. Grant the security group representing this role instance the following inheritable permission on the point of delegation:

   - Full-Control on Computer objects

3. If the administrative requirements for this role instance require that administrators in this role be able to manage all aspects of server management on the servers hosting this resources, use the Restricted Groups feature of Group Policy (by applying a GPO to the OU that represents the point of delegation) to make the security group representing the specific role instance a member of the local Administrators group on all member servers that are stored in the OU.

4. Use the Restricted Groups feature of Group Policy by applying a GPO to the OU that represents the point of delegation. The GPO should make the security group representing the specific role instance a member of the local group on all member servers stored in the OU that is sufficiently authorized to manage all service aspects of the service that administrators in this group will be responsible for managing.

Once these steps have been completed for all instances of the Resource Admins role, all administrative roles for resource management will have been delegated.

**Authorizing access to resources**
Administrators who are responsible for managing specific resources might also require the ability to authorize access to their resources. For this purpose, Resource Admins can request security groups for each resource.

For example, the Portal A Admin requires the ability to specify who can access the portal. To authorize access to resources, the Resource Admin can perform the following tasks:

1. Contact the Security Group Admin and request creation of account and resource groups as follows:

    - A resource group for access to the service
    - Account groups to specify collections of users who can access this service; these groups are added to the resource group

2. Request that the Security Admins group grant to the Resource Admins group the ability to modify the membership of these groups. On each security group object, grant the following permissions to the security group account itself:

    - Property permission Read All Properties
    - Property permission Write Members

The Security Group Admins can evaluate the request and if approved, create the specific security groups in an appropriate location, thereby enabling Resource Admins to authorize access to their resources.

**Enabling help-desk roles**
Every business unit will typically also have at least one instance of the Help Desk Operators role. Administrators in this role are responsible for providing account support for user accounts, and typical helpdesk operations include resetting user passwords and unlocking user accounts. Since most of the administrative tasks assigned to this role involve providing account support, most of these administrative tasks involve low-level operations on user account objects.

To delegate an instance of the Help Desk Operators role, perform the following tasks:

1. Refer to the documented role instance descriptions and identify the scope of administrative authority. Map this scope to an OU subtree in the Accounts OU. Based on this information, determine the point of delegation for delegating this role instance. For most business units, the scope of delegation will typically map to the Accounts OU, since in most organizations a single administrative group is responsible for providing account support.

2. Based on the documented role instance description, determine the minimal set of permissions required to delegate all assigned administrative tasks by referring to Appendix A: Active Directory Administrative Tasks in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document. This appendix documents the minimal and precise set of permissions required to delegate every Active Directory administrative task.

    As mentioned earlier in this section, typical account support tasks include the following:

    - Resetting user passwords
    - Unlocking locked out accounts

    According to Appendix A: Active Directory Administrative Tasks, the following are the permissions required to perform these administrative tasks:

    - Resetting user passwords
    - Reset User Password Extended right on the user account
    - Unlocking locked out accounts

- Write-Property permission to the user object

Thus, in order to delegate this role, the security group representing the Help Desk Operators role for this business unit will need to be granted the following permissions:

- Reset User Password Extended right on the user account
- Write-Property to the user object

Also, note the specifics of where these permissions are required. Since all user accounts are stored in the Accounts OU, the application of inheritable permissions on the Accounts OU should be sufficient to delegate the two administrative tasks that this group is responsible for.

3. Once you have determined the point of delegation and the minimal set of permissions required to delegate this role, modify the DACL on the object that corresponds to the point of delegation and grant the security group representing this administrative role the identified set of permissions.

Once the earlier steps have been completed, the Help Desk Operators administrative role will have been delegated.

Note that in some organizations, a single administrative group might be responsible for providing account support for all Business Units. In this case, the following recommended approach should be used:

1. Domain Administrators should create a security group within the Delegation OU that was used to store the security groups representing the Business Unit Admin role instances.
2. Business Unit Admins of all business units should delegate the Help Desk Operators role as described earlier in this section.
3. Business Unit Admins of all business units should then add the security group created by the Domain Administrators to the security group representing the specific instances of the Help Desk Operators role created in their business units.

**Enabling application-specific data management roles**

It is not uncommon for organizations to have Active Directory–enabled or Active Directory–integrated applications deployed. These applications typically store application specific data in Active Directory. These applications also have administrators that might need to manage this data. Business Unit Admins can create instances of the Application Specific Data Admins roles for the purpose of delegating all administrative tasks that these administrators might need to perform to manage their application data.

Active Directory–enabled or –integrated applications might store application specific data in different places in Active Directory. For example, certain applications might extend the class definition of user objects and store their application specific user information on user objects, while other applications might create their own OUs and store application specific data in these application specific OUs.

Thus, before delegating instances of the Application Specific Data Admins roles, it is important to identify precisely where the application data is being stored. This information will be required to determine the point of delegation and the manner in which application is delegated. For example, if the application stores data in a separate OU that it creates and this OU happens to be outside of any Business Unit OU (for example, if it is a child of the Domain object), then the Domain Admins will be responsible for delegating these roles and the scope of delegation will typically be the specific OU. On the other hand, if this application stores application specific data on user attributes (for example, the application could extend the schema and store user specific information on user objects), then this falls within the scope of administrative authority of the Business Unit Admin and the scope of delegation then becomes the Accounts OU.

To delegate an instance of the Application Specific Data Admins role, perform the following tasks:

1. Refer to the documented role instance descriptions and identify the scope of administrative authority. Map this scope to an OU subtree if it maps to an OU within the business unit. If it maps to an OU outside of the Business Unit, transfer responsibility for delegating this role to Domain Admins. Based on this information, determine the point of delegation for

delegating this role instance.

Based on the documented role instance description, determine the minimal set of permissions required to delegate all assigned administrative tasks by consulting the application's administrators.

2. Once you have determined the point of delegation and the minimal set of permissions required to delegate this role, modify the DACL on the object that corresponds to the point of delegation and grant the security group representing this administrative role the identified set of permissions.

Once these steps have been completed, the Application Specific Data Admins administrative role will have been delegated.

**Delegating self-service on user accounts**
User accounts store information about users in the form of properties, or attributes, of the user object. For example, some of these properties include such information as the user's home phone number, mailing address, manager, cell phone number, and so on. Active Directory defines certain property sets that are collections of such properties. Specifically, a user object has the following property sets:

- Personal-Information
- Email-Information
- Web-Information
- RAS-Information
- User-Account-Restrictions
- Membership
- General Information
- User Logon
- Domain Password

For more information about the membership of these property sets, refer to Appendix E: Active Directory Property Sets in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

The default security descriptor for user objects, as defined in the Active Directory schema, grants users the ability to modify the following property sets:

- Personal-Information
- Email-Information
- Web-Information

User in your business unit might require the ability to modify other properties on user objects in addition to the properties that are covered by these three property sets. In this case, additional administrative steps are required to grant users the required permissions. To grant additional access to users, perform the following tasks:

1. Identify the specific properties that users are to be able to modify.
2. To all users, grant Write-Property permissions (inheritable) to SELF.

**Note**

Alternatively, your organization might choose to revoke the user's ability to modify certain attributes that users can modify by default. Additional administrative steps are also required to revoke specific permissions from users

For the exact set of permissions that are required to modify a specific property or a collection of properties, see Appendix A: Active Directory Administrative Tasks in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

**Delegating Administrative Roles**
The next step in implementing the delegation model involves delegating all the required instances of the various administrative roles. Delegating a role simply involves modifying the group memberships of the security groups of all administrative roles that have been enabled. The act of making an administrator a member of a security group that represents an administrative role results in the role having been delegated to an administrator.

To delegate the various administrative roles that have been enabled, do the following for each role instance:

- Modify the group membership of the security group representing this instance of the administrative role by adding to the membership the user account of every administrator that has been assigned to this role. Refer to the documented description for the role instance to obtain the list of all administrators that have been assigned to this role.

The delegation model for your business unit is now completely implemented and ready to use. All administrators should now be able to perform all administrative tasks assigned to them.

**Best Practices for Implementing Delegation**
The following list summarizes the set of recommended best practices for implementing delegation:

- Use security groups that represent roles for the sole purpose of delegating the roles. Do not use these security groups for other purposes. For example, do not use the security group to authorize access to other resources in the domain that are not associated with the ability to perform the administrative role.

- When delegating data management, as far as possible, delegate permissions only on OUs, which are provided specifically to facilitate delegation of administration. Delegating permissions on OUs enables them to be easily and reliably revoked.

- Unless absolutely required, do not specify permissions on individual objects within an OU.

- To delegate a role, grant permissions sufficient only to perform the set of administrative tasks assigned to an administrative role.

**Using Inheritance**
You can use inheritance of permissions to more easily implement delegation. The use of inheritance of permissions allows administrators to easily apply permissions on a collection of objects without having to explicitly modify the DACL of every object. This capability allows administrators to easily manage and control permissions on large sets of objects by applying and managing permissions on a relatively small set of objects, from which other objects inherit permissions.

When using inheritance to specify permissions, administrators can control the various ways in which an inheritable permission is propagated. The following guidelines can help you use inheritance effectively:

- When specifying inheritable permissions on an object, if these permissions need to apply only to a specific object class, specify the object class instead of letting the permission apply onto all child objects. This will ensure that these permissions, when inherited, are effective only on objects of the specific class on which they are intended to be effective.

- When specifying inheritable permissions on an object, if the permissions being specified are meant for child objects only, specify that these permissions apply to child objects only. This will ensure that the applied permissions are not effective on the object on which these permissions are being specified.

- When specifying inheritable permissions on an object, if these permissions need to be applicable only on immediate child objects, and do not need to be inherited further down, specify that these permissions should be inherited only one level below.

- When using inheritance, it is useful to take into consideration the precedence order of ACEs:

- Any inheritable DENY permissions applied on the parent object will precede any inheritable DENY permissions applied on the grandparent object. Thus, when using DENY permissions, apply inheritable ACEs as close to the target objects as possible and ensure that all target objects of the permission are contained in the subtree rooted at the OU on which the permissions are being applied. In general, as far as possible, avoid the use of DENY permissions.

- Remember that an inheritable permission cannot override an explicit ACE on an object. Thus, before using inheritance, it is recommended that you check to see that there are no explicit ACEs on target objects that will override the inheritable ACE.

| **Note** |
| --- |
| Any user who has the ability to modify the DACL of any OU in the OU hierarchy can use inheritance to grant him or herself or any other user any permission on any object in the subtree rooted at that OU. Therefore, carefully grant permissions to modify the DACL of an OU object. |

**Taking Ownership of Objects**

Some features of object ownership can complicate management of Active Directory data, and you should be aware of them when planning and implementing delegation of administrative authority. Every object in Active Directory has an owner. The owner of an object is the only person who has the inherent right to control who has permission to perform operations on the object and in what way – the owner has the inherent right to modify permissions on an object even if he or she is explicitly denied all access to the object in the object's DACL. An object's owner can also grant or deny permission for different kinds of access to particular users or groups of users. An object's owner can also transfer ownership by giving another security principal permission to take ownership. The creator of an object is usually also the owner of the object; the only exception to this rule is that if an object is being created programmatically, an owner can be explicitly specified. Thus, when a user creates an object, he or she usually becomes the owner of the object.

This aspect of object ownership affects how Active Directory data can be managed. Consider the scenario in which administrative authority has been delegated to a group of administrators by implementing an administrative role. As part of carrying out delegated responsibilities, one of the members of the administrative role creates an object. As the creator of the object, this specific administrator also becomes the owner of the object. At a later date, this administrator might no longer be a member of this administrative role. If this administrator is no longer a member of this administrative role, he or she should ordinarily no longer be able to manage this object. Since this administrator is no longer a member of the administrative group representing the administrative role that he or she was a member of, and since inheritable permissions were used to delegate control of these objects to members of this administrative role, it seems reasonable to conclude that this administrator will no longer have access to this object. However, the administrator will still have access to this object by virtue of the fact that as the creator of the object, he or she still owns the object. This can obviously complicate delegation of administrative authority.

While this issue can be addressed by putting in place business policies to add certain checks when an administrator is removed from an administrative role, such administrative checks will not be fool-proof. A business unit can enforce policy that requires that when an administrator is removed from an administrative role, a check is made to see what objects the administrator owns and that ownership of all objects be appropriately transferred. However, this check is not fool-proof because the administrator might have used his authority as owner to grant himself or herself (or another user) sufficient permissions to modify the DACL of this object, or permissions to perform other operations.

The following precaution is recommended to adequately address this issue with respect to managing user and group objects: Create a script that runs in the security context of a Domain Admin account on a periodic basis (like daily or weekly). This script should perform the following actions:

- Walk through a specifiable subtree of objects, rooted at some OU, and for every object do the following:

    - Check the owner field of the object; if the owner field is not set to Domain Admins, change the owner field to Domain Admins.

    - Additionally, check the DACL of the object to ensure that the only explicit ACEs in the DACL are the ones that are found in the default security descriptor for that object class in the Schema. This check ensures that creator of the object did not modify the DACL explicitly to grant anyone (including himself or herself) any additional permissions. If any explicit ACEs are found that do not exist in the default security descriptor for that object class, report their existence and remove these ACEs.

| Note |
| --- |
| This check will limit an organization's ability to apply explicit ACEs on objects. This limitation can be overcome by having the script report the existence of such ACEs and prompt for manual deletion. This approach, while continuing to allow the use of explicit ACEs, will require manual intervention and might require significant manual intervention if many explicit ACEs are used. However, if explicit ACEs have been used sparingly and minimally, this approach can be very useful. |

The reason this approach will work only for user and group objects is that the default security descriptors of these object classes do not contain any ACEs for the Creator Owner Trustee. The general problem with Creator Owner ACEs on other object classes, however, is that some default security descriptors can contain ACEs for the Creator Owner Trustee, because during object creation, the user field in any ACE for this trustee is replaced with the SID of the user creating the object. Thus, when checking to see that only ACEs that were in the default security descriptor are present, any ACE that was mapped to the SID of the creator will show up as an explicit ACE that was not in the default security descriptor and, if removed, could impact functionality.

| Note |
| --- |
| If your organization should decide to implement the script approach described here, it is strongly advised that this approach be thoroughly tested in non-production environments before being implemented in your production environment. |

# Maintaining Your Data Management Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Over time, your data management administrative delegation model might need to be modified to address changing needs and requirements. The data management tasks required to maintain a changing data management delegation model are the same as those required to maintain a service management delegation model.

Tasks that are required to maintain the service delegation model can include:

- Adding (delegating) or removing (undelegating) members to or from data administration role instances.
- Adding new instances of roles in the event of changes in administrative needs.
- Modifying a role by assigning or revoking a new or existing task from the role definition.
- Creating new custom roles if the need arises.
- Undelegating a role completely by revoking all permissions that are granted to the group that represents the role.

**Modifying the Administrators Assigned to a Data Administration Role**

To address personnel changes that affect the set of administrators who are assigned to a data administration role, you typically need to perform one of the following tasks:

- Add new administrators to an existing role.
- Remove one or more assigned administrators from a role.
- To add one or more data administrators to a role, simply add the user accounts of these administrators to the security group that represents the role. To remove one or more assigned data administrators from a role, simply remove the user accounts of these data administrators from the related security group.

**Adding New Instances of Existing Roles**

Organizational or operational changes might necessitate the introduction of new role instances. For example, the introduction of a new business unit might require creating an entire set of data management role instances to provide administrative coverage. Similarly, the introduction of a new location might require new instances of the Workstation Admins or Resource Admins roles, or both, to facilitate administration of the IT resources at the new location.

To introduce a new role instance, understand the new administrative requirement and decide whether or not this requirement can be addressed by an existing role definition. If this requirement cannot be addressed by an existing role definition, create a custom role to meet this requirement (see Creating New Custom Data Administration Roles).

If the requirement can be addressed by an existing role definition, take the following steps:

1. Document the requirements that call for the creation of a new role instance.
2. Determine the scope of influence of the new instance, which affects how and where this role is granted the required permissions in Active Directory. If needed, create a new OU or move an existing OU to ensure that the target of administration is appropriately located in the OU structure.

   For example, if a new location is being added and the new location will have on-site workstations, you should add a new OU under the Workstations OU in the Business Unit's OU structure. This new OU should contain all of the computer accounts for the workstations that will be physically located at this location. If you have one OU for every unique location, you will have to add a new OU for this location. If you are moving an entire OU from one business unit to another, you could simply move the OU to this new location under the Workstations OU.
3. Create a new security group to represent the new role instance. Depending on your administration model, the Business

Unit administrator might or might not have delegated the authority to sub-delegate administrative authority.

- If this new requirement is being addressed by sub-delegated administrators, this new security group should be created by the sub-delegated administrators in the OU where they have stored existing security groups representing delegated roles.
- If a new role instance is required, the Business Unit administrator can create a security group in the OU that is used by the Business Unit administrator to delegate other roles instances.

4. After the security group has been created, appropriate permissions must be assigned to it so that the administrators in that role can perform the responsibilities of the role. Refer to the documentation that you created during your role definition process and assign this new security group all required permissions in the appropriate administrative scope. For more information about how to implement recommended roles, see Implementing Your Data Management Delegation Model.
5. Add the user accounts that represent the administrators who will be assigned to this role to the corresponding security group.
6. Document the new role instance by using the role instance template to specify the scope of influence, the security group that represents the role instances, the detailed set of permissions that are granted, and the administrators who are assigned to the role.

**Modifying Administrative Tasks Associated With a Data Administration Role**

If the set of administrative tasks that is associated with a role must be modified, you might need to add one or more new tasks to a role or remove one or more assigned tasks from the role.

For example, suppose that there are three existing instances of the Help Desk role, one for each business unit, and that this role is initially assigned the following two tasks:

- Reset Passwords
- Unlock User Accounts

Six months later, it is decided that responsibility for the administrative task of unlocking user accounts will be given solely to the Corporate Security role. Addressing this new requirement involves revoking this ability from all instances of the Help Desk role.

There are two ways to revoke these permissions:

1. Modify the DACLs of the three OUs where the required permissions were granted and remove the permissions that allow the security group representing each role the ability to unlock user accounts.
2. Remove and reassign the role, as follows:
    - Revoke all the permissions granted to the security group that represents the Help Desk role
    - Re-grant the set of permissions that are required to perform the new and updated set of administrative tasks.

Although in this example the approach of modifying the DACLs might appear easy to grasp and implement, it is complicated to use in the following cases:

- The precise set of permissions required for every administrative task cannot be mapped to one specific permission.
- Full Control has been granted.
- A number of permissions were specified for a given role.

In such cases, it is preferable to remove and re-assign the role. The Dsrevoke command-line tool can be used to easily and reliably remove all permissions that are granted to a security principal. For more information about using this tool, see Undelegating a Data Administration Role.

To add a new task to a role for each role instance that needs to be modified, perform the following tasks:

1. Identify the set of permissions that are required to perform the task.

2. Identify the location (specific OU or OU tree) in the OU hierarchy where permissions have been granted to the security group that represents the role instances.

3. Either revoke all permissions and then re-apply the updated set of permissions, or grant this security group the additional minimal set of permissions that are required to perform the new administrative task.

Depending on the requirement, you might need to update the set of administrative tasks that are assigned to a specific instance of the role or to all instances of the role.

**Creating New Custom Data Administration Roles**
Although the set of recommended roles should meet most of your data administration needs, unique requirements or special needs might call for the creation of new custom role definitions. Because the specific requirements can vary, the actual definition of a custom role is necessarily variable.

You can, however, use the following steps to aid in creating and implementing new custom roles:

1. Analyze and obtain a good understanding of the requirements of the role.

2. Create a new role to address these requirements.

3. Define the scope of administrative authority for this role.

4. Define the precise set of administrative tasks for which administrators in this role will be responsible.

5. Identify the minimal and precise set of permissions that are required to delegate the set of administrative tasks.

6. Create a naming convention for the security groups that represent instances of this role.

7. Understand how the administrative scope of the role maps to the OU structure in your organization.

After the data owner creates this model, the Business Unit Admin or a data administrator who has been delegated the ability to sub-delegate should implement the required number of instances of the role.

**Undelegating a Data Administration Role**
If a business unit no longer has a need for a specific instance of some role, the Business Unit Admin needs to discontinue it. For example, an organizational change might result in the discontinuation of business at a physical location, either by moving assets and operations from that location to another location, or by completely stopping operations in that location. In this circumstance, you might no longer need the administrative roles that were initially implemented to manage workstations, resources, and perhaps users in that location.

To completely discontinue the role instance:

1. Revoke all the permissions that are granted to the security group that represents the role.

2. Empty the membership of this security group.

3. Delete the security group.

In another case, your organization might decide to temporarily halt operations at a location for four months due to

construction. During this time, none of the IT resources or users in that location will require administrative coverage, so you might want to disable these roles during that period. In this situation, you will want to only temporarily disable the role of the security group representing the role instance.

**Revoking All Permissions for a Security Principal**

For the first scenario, you can use the Dsrevoke command-line tool to revoke all permissions granted to the security group representing the role instance.

The Dsrevoke command-line tool provides the ability to remove ACEs that have been applied for a security principal on the domain object or on OU objects. By specifying the security group that represents the instantiated role and the OU or domain on which the security group that represents the role has been granted permissions, you can undelegate the role. However, Dsrevoke removes only permissions. If a role has user rights applied, you must manually remove these from the group.

For a specified user or security group, Dsrevoke removes ACEs on the specified root object and on all OU objects below that root in the hierarchy. If a domain root is specified, the tool searches the domain root and all OUs below that root. If an OU root is specified, the tool searches the OU tree, beginning at the specified root. Dsrevoke reports the explicit ACEs that exist for the specified security principal. When you remove the ACEs, for each explicit ACE that Dsrevoke removes, all inheritable ACEs that apply onto child OUs are also removed.

Dsrevoke works only on domain objects and OUs. If you delegate permissions on a container object or if you explicitly set permissions on an object within a container or OU, you must manually remove these ACEs. For this reason, as a best practice, it is recommended that permissions always be applied to OUs rather than to specific objects within OUs, and that permissions be applied to child OUs by using inheritance. For ease of management, always use security groups to apply permissions.

**Using Dsrevoke**

Dsrevoke.exe has the following syntax:

**dsrevoke /report|/remove [/domain:** *domainname*] **[/username:***username*]

[**/password:***password*|*] **[/root:***domain/OU*] **securityprincipal**

Descriptions for each option are as follows:

- **/report:** Reports the explicit ACEs that are currently set for the specified security principal on OU objects in the specified domain or an OU subtree. By default, the command dsrevoke /report starts at the domain root and searches every OU below that root for explicit ACEs that are granted to the specified security principal. If you are sure that the permissions for a security group are set only on or below a specific OU, you can specify the scope of the search by using the /OU switch to make the search more efficient.

- **/remove:** Reports all explicit ACEs and then, after prompting for confirmation, removes the ACEs that are currently set for the security principal, including all inherited ACEs.

- **/domain:** The DNS or NetBIOS name of the domain in which the permissions are to be removed. This value must be specified only when the ACEs that you want to remove are set on OUs in a domain other than the domain of the logged-on user.

- **/username:** The user name of the user who is using the tool. This value is required when:

    - The user is not logged on as an administrator.

    - ACEs are being removed in a domain other than the domain of the logged-on user.

- **/password:** The password of the tool user. If the command is entered with an asterisk (*) in place of a password, the tool prompts the user for a password.

- **/root:** The OU or domain root at which to start the search for ACEs. If no value is specified, the search begins at the root of the specified domain. If no domain is specified, the search begins at the root of the domain of the logged-on user. When specifying a root domain or OU, you must use the distinguished name (for example, /root:OU=BusUnits=DC=DomainA,DC=com). If spaces occur in any part of the distinguished name, enclose the entire option in quotation marks (for example, "/root:OU=Product Development,OU=Delegation,OU=Business Units,DC=DomainA,DC=com").

- **/securityprincipal:** The identity of the user or group in the form DomainName\UserName or DomainName\GroupName. Use the DNS name or NetBIOS name of the domain.

For more information about using Dsrevoke.exe and to obtain the tool, see Appendix G: Active Directory Delegation Tools in Best Practices for Delegating Active Directory Administration: Appendices, which accompanies this document.

# Case Study: A Delegation Scenario

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To help you better understand the various phases of the delegation process, this section provides a walk-through example of the creation and implementation of service and data delegation models by a hypothetical multinational corporation. By following this walkthrough, you can become more familiar with the delegation process.

# Company Overview

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso Pharmaceuticals is a large organization that has its headquarters in Chicago, Illinois and has operations in five other locations in North America and Europe. The Active Directory infrastructure consists of a single forest, three domains, and six sites.

**Physical Locations**

The company has operations in six physical locations, three in North America and three in Europe:

- North America — Chicago, New York, Atlanta
- Europe — London, Paris, Rome

**Business Units**

Contoso has four business units. Table 7 shows the business units and their locations.

**Table 7 Contoso Business Units and Locations**

| Business Units | Locations |
|---|---|
| Research and Development | Chicago |
| Production | Atlanta |
| Business Management | Chicago, New York, London, Paris, Rome |
| IT | Chicago, Atlanta, New York, London, Paris, Rome |

**Users, Workstations, and Servers**

Table 8 shows the distribution of users, workstations, and servers per business unit.

**Table 8 Contoso Business Units and Contents**

| Business Unit | Users | Workstations | Servers |
|---|---|---|---|
| Research and Development | 5,000 | 10,000 | 1,000 |
| Production | 3,000 | 5,000 | 520 |
| Business Management | 5,500 | 7,000 | 250 |
| IT | 2,500 | 5,000 | 100 |

**Servers**

The servers include file servers, Web servers, database servers, and application servers.

Table 9 shows the numbers of servers of each type and their respective business units.

**Table 9 Business Unit Server Numbers and Types, By Business Unit**

| Business Unit | File Servers | Web Servers | Database Servers | Application Servers |
|---|---|---|---|---|
| Research and Development | 800 | 50 | 100 | 50 |
| Production | 300 | 20 | 100 | 100 |
| Business Management | 120 | 45 | 50 | 35 |
| IT | 30 | 30 | 20 | 20 |

# Active Directory Infrastructure

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Contoso has chosen to deploy Active Directory based on the Regional Domain Model. There is one regional domain each for the two major geographic locations (North America and Europe), connected by a Wide Area Network (WAN). The regional domain model enables Contoso to maintain a stable environment over time. This model consists of a forest root domain and two regional domains, as shown in Figure 9.




**Site Topology**

The Contoso site topology consists of six logical Active Directory sites, one for each of the six physical locations, as shown in Figure 10.




**Domain Controller Placement**

Contoso has deployed a total of 17 domain controllers — two domain controllers per site and three additional controllers for the forest root domain.

Table 10 shows the distribution of domain controllers per domain across sites.

**Table 10 Domain Controllers Per Domain in Each Site**

| Site | Concorp.contoso.com | NOAM.concorp.contoso.com | Europe.concorp.contoso.com |
|---|---|---|---|
| Chicago | 3 | 2 | 2 |
| Atlanta | | 2 | |
| New York | | 2 | |
| London | | | 2 |
| Paris | | | 2 |
| Rome | | | 2 |

**Domain Controller Placement**

Domain controllers for each domain are placed in the sites as shown in Table 11. The letter "G" indicates that the domain controller is a global catalog server.

**Table 11 Domain Controllers in Contoso sites**

| Domain Controller | Chicago | Atlanta | New York | London | Paris | Rome |
|---|---|---|---|---|---|---|
| CONTOSO-DC1 | G | | | | | |
| CONTOSO-DC2 | X | | | | | |
| CONTOSO-DC3 | X | | | | | |
| NOAM-DC1 | X | | | | | |
| NOAM-DC2 | G | | | | | |
| NOAM-DC3 | | G | | | | |
| NOAM-DC4 | | X | | | | |
| NOAM-DC5 | | | G | | | |
| NOAM-DC6 | | | X | | | |
| EUROPE-DC1 | X | | | | | |
| EUROPE-DC2 | X | | | | | |
| EUROPE-DC3 | | | | G | | |
| EUROPE-DC4 | | | | X | | |
| EUROPE-DC5 | | | | | G | |
| EUROPE-DC6 | | | | | X | |
| EUROPE-DC7 | | | | | | G |
| EUROPE-DC8 | | | | | | X |

All domain controllers are placed in highly secure physical locations to which only authorized personnel are granted access. All domain controllers in remote locations are equipped with a remote administration solution, such as Remote Insight Lights-Out (RILO), so that administrators can control both hardware and software on domain controllers remotely to manage systems where no IT staff is stationed.

**DNS Server Placement**

To provide fault tolerance, the DNS Server service runs on every domain controller in the Contoso forest. Zones are distributed as follows:

- Domain controllers in the forest root domain host the forest root DNS name.
- Domain controllers for each regional domain host the DNS zone that corresponds to the DNS name of the domain.

The zone containing the forest-wide domain controller locator records replicates to every DNS server in the forest by using the forest-wide DNS application directory partition, ForestDnsZones.

**Global Catalog Server Placement**

All sites in the Contoso environment have at least 100 users. To facilitate user logon requests and forest-wide searches, Contoso follows the general Windows Server 2003 deployment recommendation for placing a global catalog server in any site where there are at least 100 users. Two global catalog servers are placed in Chicago to accommodate the large number of users in that site.

For more information about global catalog server placement, see "Designing the Site Topology" in *Designing and Deploying Directory and Security Services* of the Windows Server 2003 Deployment Kit (or see "Designing the Site Topology" on the Web at http://go.microsoft.com/fwlink/?LinkId=4724).

**Operations Master Roles Placement**

Contoso places the operations master roles according to the best practices that are recommended in the Active Directory Operations Guide Version 1.5, on the Web at http://go.microsoft.com/fwlink/?LinkId=19827.

**Forest-wide Role Placement**
The two forest-wide roles are domain naming master and schema master. By default, all operations master roles (including the domain-wide roles) are placed on the first domain controller that is installed in the forest root domain.

Contoso assigns the two forest-wide operations master roles to the original forest root domain controller, CONTOSO-DC1. Both roles are compatible with a global catalog server.

| **Note** |
| --- |
| In Windows 2000, the domain naming master must be placed on a global catalog server. |

**Domain-wide Role Placement**
The first domain controller that is installed in a domain has the three domain-wide roles by default. Because the concorp.contoso.com domain is the forest root domain, the first domain controller that is installed to create the forest root domain contains the three domain roles and is also a global catalog server. Because the infrastructure master must not be located on a global catalog server, Contoso moves that role, as well as the other two domain-wide roles, to CONTOSO-DC2, which is not a global catalog server.

The first domain controllers that are installed in noam.concorp.contoso.com and in europe.concorp.contoso.com are not global catalog servers. Therefore, the domain-level roles are left on these domain controllers, as shown in Table 12.

**Table 12 Domain-Wide Operations Master Role Placement**

| Domain | PDC Master | Infrastructure Master | RID Master |
| --- | --- | --- | --- |
| concorp.contoso.com | CONTOSO-DC2 | CONTOSO-DC2 | CONTOSO-DC2 |
| noam.concorp.contoso.com | NOAM-DC1 | NOAM-DC1 | NOAM-DC1 |
| europe.concorp.contoso.com | EUROPE-DC1 | EUROPE-DC1 | EUROPE-DC1 |

**Domain Modes and Functional Levels**
The Contoso environment does not use domain controllers that are running Windows NT 4.0 or Windows 2000, and therefore can run all domain controllers in native mode. The following functional levels are in effect:

- **Domain functional level —** Windows Server 2003 in all domains
- **Forest Functional Level —** Windows Server 2003

At these functional levels, all new features in Windows Server 2003 are available throughout the forest. For more information about domain and forest functional levels, see "Enabling Advanced Windows Server 2003 Active Directory Features" in *Designing and Deploying Directory and Security Services* of the Microsoft® Windows® Server 2003 Deployment Kit (or see "Enabling Advanced Windows Server 2003 Active Directory Features" on the Web at http://go.microsoft.com/fwlink/?LinkID=6937).

# Managing Contoso's Active Directory Environment

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To manage the Active Directory environment, Contoso uses the following process:

1. Create the service management administrative delegation model.
2. Implement the service management administrative delegation model.
3. Hand off data management to data administrators.
4. Create the data management administrative delegation model for every business unit:
    a. Create an OU structure for every business unit.
    b. Create an administration model for data management.
5. Implement the data management administrative delegation model for every business unit.

The IT department is responsible for managing the Active Directory environment. The IT department is based in the Chicago headquarters and has the following components and users:

- CIO
    - Core Infrastructure (100)
    - Connectivity (30)
    - Messaging (30)
    - Business Applications (50)
    - IT Security (50)
    - Support Services (75)

**Active Directory Administrative Personnel Overview**
Of the approximately 325 employees in the IT business unit, 135 will be responsible for Active Directory service and data management, including managing the Active Directory directory service, user accounts, and all workstations, servers, and resources in the Contoso organization. They will also provide help desk support to the entire organization. The remaining 190 administrative personnel provide administrative coverage for network connectivity, messaging, and other network services.

Table 13 shows the types and numbers of administrative personnel in each administrative category. This team of administrators will create and implement the service and data delegation models for the organization.

**Table 13 Active Directory Administrative Personnel**

| Role | Number of IT Personnel |
|---|---|
| Management | 10 |
| Service Administrators | 25 |
| Data Administrators | 50 |
| Help Desk Operators | 50 |

# Step 1 — Create the Contoso Service Management Administrative Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The service owners create the service delegation model according to the following criteria:

- **Objective:** Create a delegation model to distribute service management among service administrators by enabling efficient, security-conscious delegation and distribution of administrative responsibilities among various administrative groups.
- **Stakeholders:** Service owners.
- **Approach:** Identify the number of instances of each service management role that are needed and the administrative personnel who are assigned to each role.

**Service Management Roles**

For the service administrative delegation model, Contoso requires the following management roles:

- Forest Configuration Operators
- Domain Configuration Operators
- Schema Admins
- Replication Management Admins
- Replication Monitoring Operators
- DNS Admins
- Security Policy Administrators
- Service Admin Managers
- Domain Controller Admins
- Backup Operators

**Template for Creating the Delegation Model**

Contoso uses the information in Table 14 as a template for assigning each role.

**Table 14 Role Assignment Template**

| Template Fields | Phase During Which to Complete Assignment |
|---|---|
| Role Instance Name | Creation |
| Instance of | Creation |
| Instance Number | Creation |
| Assigned Administrators | Creation |
| Assigned Tasks | Creation |

| | |
|---|---|
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

**Assigning the Roles During Role Creation**
Fields for each role are filled in by using the role template.

**Assigning Forest Configuration Operators**
The service owner follows the recommendation for creating only one instance of this role.

Table 15 shows the model creation entries in the template for this role.

**Table 15 Model Creation Template for Forest Configuration Operators Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Forest Config Ops |
| Instance of | Forest Configuration Operators Role |
| Instance Number | 1 of 1 |
| Assigned Administrators | Joe, Sally, Kevin |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

**Assigning Domain Configuration Operators**
The service owner follows the recommendation for creating only one instance of this role per domain. Because there are three domains, three instances are defined.

Table 16 shows the model creation entries in the template for this role.

**Table 16 Model Creation Template for Domain Configuration Operators Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Root Dom Config Ops |
| Instance of | Domain Configuration Operators Role |
| Instance Number | 1 of 3 |
| Assigned Administrators | Michael, Sally, Gordon |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

| Field | Assignment Information |
|---|---|
| Role Instance Name | NOAM Dom Config Ops |
| Instance of | Domain Configuration Operators Role |
| Instance Number | 2 of 3 |

| | |
|---|---|
| Assigned Administrators | John, Sandra |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

| Field | Assignment Information |
|---|---|
| Role Instance Name | EUROPE Dom Config Ops |
| Instance of | Domain Configuration Operators Role |
| Instance Number | 3 of 3 |
| Assigned Administrators | Christoph, Anna |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in London |

**Assigning Schema Admins**

The service owner follows the recommendation for creating only one instance of this role.

Table 17 shows the model creation entries in the template for this role.

**Table 17 Model creation template for Schema Admins role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Schema Admins |
| Instance of | Schema Admins Role |
| Instance Number | 1 of 1 |
| Assigned Administrators | Joe (also assigned to Forest Config Ops role) |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago<br><br>Only group authorized to perform schema modifications |

**Assigning Replication Management Admins**

The service owner follows the recommendation for creating only one instance of this role.

Table 18 shows the model creation entries in the template for this role.

**Table 18 Model Creation Template for Replication Management Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Repl Mgmt Admins |
| Instance of | Replication Management Admins |
| Instance Number | 1 of 1 |

| | |
|---|---|
| Assigned Administrators | Sally (also assigned to Forest Config Ops role)<br>Kevin (also assigned to Forest Config Ops role) |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

**Assigning Replication Monitoring Operators**

The service owner follows the recommendation for creating only one instance of this role.

Table 19 shows the model creation entries in the template for this role.

**Table 19 Model Creation Template for Replication Monitoring Operators Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Repl Monitoring Ops |
| Instance of | Replication Monitoring Operators |
| Instance Number | 1 of 1 |
| Assigned Administrators | Tom, Russ |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

**Assigning DNS Admins**

The service owner follows the recommendation for creating one instance of this role for the entire forest and one instance for each domain.

Table 20 shows the model creation entries in the template for this role.

**Table 20 Model Creation Template for DNS Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Forest DNS Admins |
| Instance of | DNS Admins |
| Instance Number | 1 of 4 |
| Assigned Administrators | Mark (also assigned to Service Admin Managers, Security Policy Admins) |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso DNS Admins |
| Instance of | DNS Admins |
| Instance Number | 2 of 4 |

| | |
|---|---|
| Assigned Administrators | Andrew |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

| Field | Assignment Information |
|---|---|
| Role Instance Name | NOAM DNS Admins |
| Instance of | DNS Admins |
| Instance Number | 3 of 4 |
| Assigned Administrators | Jay |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

| Field | Assignment Information |
|---|---|
| Role Instance Name | EUROPE DNS Admins |
| Instance of | DNS Admins |
| Instance Number | 4 of 4 |
| Assigned Administrators | Laurie, Samuel |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

**Assigning Security Policy Administrators**

The service owner follows the recommendation for creating one instance of this role.

Table 21 shows the model creation entries in the template for this role.

**Table 21 Model creation template for Security Policy Administrators Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Sec Pol Admins |
| Instance of | Security Policy Admins |
| Instance Number | 1 of 1 |
| Assigned Administrators | Mark (also assigned to Service Admin Managers, DNS Admins) |
| Assigned Tasks | |
| Security Group | |

| | |
|---|---|
| Permissions Assigned | |
| Notes | Based in Chicago |

**Assigning Service Admin Managers**

The service owner follows the recommendation for creating one instance of this role.

Table 22 shows the model creation entries in the template for this role.

**Table 22 Model Creation Template for Service Admin Managers Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Srvc Admin Managers |
| Instance of | Service Admin Managers |
| Instance Number | 1 of 1 |
| Assigned Administrators | Lisa<br><br>Mark (also assigned to Security Policy Admins, DNS Admins) |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

**Assigning Domain Controller Admins**

Contoso implements two instances of the Domain Controller Admins role. One instance is based in Chicago and is responsible for managing all domain controllers that are located in Chicago and for remotely managing all domain controllers that are located in the other two sites in North America. Similarly, one instance of the Domain Controller Admins role is based in London and is responsible for managing all domain controllers in the London site and for remotely managing all domain controllers that are located in the other two sites in Europe.

Contoso has implemented RILOs in remote sites so that all aspects of their domain controllers can be managed remotely from the hub site. The only operation that cannot be performed on RILO-based systems is physically starting and shutting down the domain controllers in a remote location. For this administrative operation, one local administrative group is assigned the responsibility for shutting down and starting domain controllers when needed.

Table 23 shows the model creation entries in the template for this role.

**Table 23 Model Creation Template for Domain Controller Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Root and NOAM DC Admins |
| Instance of | Domain Controller Admins |
| Instance Number | 1 of 2 |
| Assigned Administrators | Paul, Andy |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

| Field | Assignment Information |
|---|---|
| Role Instance Name | Europe DC Admins |

| | |
|---|---|
| Instance of | Domain Controller Admins |
| Instance Number | 2 of 2 |
| Assigned Administrators | James, Jessica |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in London |

**Assigning Backup Operators**

According to the recommendations for this role, there should be one instance for each domain. Although there are three instances of this role in the model, the service owner decides to assign the same group of administrators to each of the three instances. The administrators in this group are all located in Chicago. This one group is responsible for backing up Active Directory for all three domains. Table 24 shows the model creation entries in the template for this role.

**Table 24 Model Creation Template for Backup Operators Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Contoso Root Backup Operators |
| Instance of | Backup Operators |
| Instance Number | 1 of 3 |
| Assigned Administrators | Maria |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

| Field | Assignment Information |
|---|---|
| Role Instance Name | NOAM Backup Operators |
| Instance of | Backup Operators |
| Instance Number | 2 of 3 |
| Assigned Administrators | Kris |
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

| Field | Assignment Information |
|---|---|
| Role Instance Name | Europe Backup Operators |
| Instance of | Backup Operators |
| Instance Number | 3 of 3 |
| Assigned Administrators | Brian |

| | |
|---|---|
| Assigned Tasks | |
| Security Group | |
| Permissions Assigned | |
| Notes | Based in Chicago |

**Service Role Assignment Summary**

Table 25 provides a summary of the roles and instances that have been assigned for service management.

**Table 25 Summary of Assigned Contoso Service Management Roles**

| Role | Instances | Comments |
|---|---|---|
| Forest Configuration Operators | 1 | One role per forest |
| Domain Configuration Operators | 3 | One role per domain |
| Schema Admins | 1 | One role per forest |
| Replication Management Admins | 1 | One role per forest |
| Replication Monitoring Operators | 1 | One role per forest |
| DNS Admins | 4 | One forest-wide role and one role per domain |
| Security Policy Administrators | 1 | One role per forest |
| Service Admin Managers | 1 | One role per forest |
| Domain Controller Admins | 2 | One role per continent (North America and Europe) |
| Backup Operators | 3 | One role per domain |

Now that the creation phase is complete, Contoso has a delegation model that documents the division of responsibility for service management of the Active Directory infrastructure. In the next step, the Enterprise Admins group will implement the Active Directory directory service management model.

# Step 2 — Implement the Administrative Delegation Model for Contoso Service Management

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The service administrators implement the service delegation model according to the following criteria:

- **Objective:** Implement the service management delegation model
- **Stakeholders:** Enterprise Admins
- **Approach:** Implement each role instance according to the delegation model design document
- **Assumption:** The three domains have been installed and are running

**Out-of-Box Container Hierarchy for the Forest Root Domain**
When the first domain controller is installed to create the forest root domain of the Contoso forest, the default set of containers shown in Figure 11 is created:




**Creating an OU to Store Security Groups That Represent Service Roles**
To more easily manage the instances of the Business Unit Admins role, Contoso follows the recommendation to store all instances of the Business Unit Admins roles in one place. A member of the Enterprise Admins group creates an OU called Service Management in the forest root domain, as shown in Figure 12.




**Implementing the Service Administration Role Instances**

Contoso implements the service administration role instances according to the service delegation model.

**Implementing Forest Configuration Operators**
A member of Enterprise Admins creates one instance of this role by carrying out the following steps.

**To implement the Forest Configuration Operators role**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:

   - **Group Name:** Contoso Forest Config Ops.

   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global. "Native mode" has the following meaning, depending on the operating system:

     **Windows Server 2003:** the domain functional level is Windows 2000 native or Windows Server 2003.

     **Windows 2000:** the domain mode is native.

2. Grant the following permissions to Contoso Forest Config Ops:

| **Note** |
| --- |
| The extended rights in these procedures are presented in common-name format, which you can use to script extended rights assignments. To manually assign extended rights, you can use Active Directory Users and Computers to assign rights on domain objects and use ADSI Edit to assign rights on objects in the configuration and schema containers. For the mapping of common names used below to display names that are used in the ADSI Edit UI, see Appendix D: Active Directory Extended Rights. To apply extended rights that are not found in the UI, use Dsacls.exe, which is available in Windows Support Tools. |

**Extended rights:**

- DS-Replication-Get-Changes on CN=Configuration,DC=concorp,DC=contoso,DC=com

- DS-Replication-Get-Changes on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

- DS-Replication-Get-Changes-All on CN=Configuration,DC=concorp,DC=contoso,DC=com

- DS-Replication-Get-Changes-All on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

- DS-Replication-Manage-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com

- DS-Replication-Manage-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

- DS-Replication-Monitor-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com

- DS-Replication-Monitor-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

- DS-Replication-Synchronize on CN=Configuration,DC=concorp,DC=contoso,DC=com

- DS-Replication-Synchronize on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

- Change-Schema-Master on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

- Change-Domain-Master on CN=Partitions,CN=Configuration,DC=concorp,DC=contoso,DC=com

**Permissions on CN=Sites,CN=Configuration,DC=concorp,DC=contoso,DC=com:**

- Inherit-only Full-Control permissions on all objects of class NTDS-Settings

- Inherit-only Create Child permissions to be able to create and delete objects of class Server
- Inherit-only Delete Child permissions to be able to delete objects of class NTDS-Settings
- Inherit-only Delete Child permissions to be able to delete objects of class Server

**Inheritable Full Control permissions on:**

- CN=Partitions,CN=Configuration,DC=concorp,DC=contoso,DC=com
- CN=sites,CN=Configuration,DC=concorp,DC=contoso,DC=com

**Property permissions Read-Property and Write-Property to the FSMO-Role-Owner property on the following objects:**

- CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
- CN=Partitions,CN=Configuration,DC=concorp,DC=contoso,DC=com

**Property permissions Read-Property and Write-Property to the LDAP-Admin-Limits property on the following object:**

- CN=Default Query Policy,CN=Query-Policies,CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,DC=concorp,DC=contoso,DC=com

**Property permissions Read-Property and Write-Property to the ms-DS-Behavior-Version property on the following objects:**

- CN=Partitions,CN=Configuration,DC=concorp,DC=contoso,DC=com

**Inheritable Full Control permissions on the following objects:**

- CN=System,DC=concorp,DC=contoso,DC=com
- CN=System,DC=noam,DC=concorp,DC=contoso,DC=com
- CN=System,DC=europe,DC=concorp,DC=contoso,DC=com

| **Note** |
| --- |
| If the forest root domain is not in native mode, universal groups are not available. In this case, create a new domain local group in each domain and grant that group the permissions specified earlier in this procedure. Then make the Contoso Forest Config Ops global group a member of these domain local groups. |

**Create Child permissions on the following OUs:**

- OU=Domain Controllers,DC=concorp,DC=contoso,DC=com
- OU=Domain Controllers,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Domain Controllers,DC=europe,DC=concorp,DC=contoso,DC=com

3. Grant the Creator Owner trustee Full Control on the following objects:
   - The computer object that represents the server that will be promoted to create a new domain.

- CN=sites,CN=Configuration,DC=concorp,DC=contoso,DC=com

4. Add Contoso Forest Config Ops to the local Administrators group of any member server that is to be promoted to a Domain Controller.

**Implementing Domain Configuration Operators**
A member of Enterprise Admins creates three instances of this role, one for each domain, by performing the following steps.

**To implement the Domain Configuration Operators role, forest root domain instance, for concorp.contoso.co**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:

   - **Group Name:** Contoso Root Dom Config Ops

   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global. "Native mode" has the following meaning, depending on the operating system:

     **Windows Server 2003:** the domain functional level is Windows 2000 native or Windows Server 2003.

     **Windows 2000:** the domain mode is native.

2. Grant the following permissions to Contoso Root Dom Config Ops:

   **Extended rights:**

   - DS-Install-Replica on DC=Contoso,DC=com

   - DS-Replication-Get-Changes on DC=contoso,DC=com

   - DS-Replication-Get-Changes on CN=Configuration,DC=concorp,DC=contoso,DC=com

   - DS-Replication-Get-Changes on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

   - DS-Replication-Get-Changes-All on DC=contoso,DC=com

   - DS-Replication-Get-Changes-All on CN=Configuration,DC=concorp,DC=contoso,DC=com

   - DS-Replication-Get-Changes-All on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

   - DS-Replication-Manage-Topology on DC=contoso,DC=com

   - DS-Replication-Manage-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com

   - DS-Replication-Manage-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

   - DS-Replication-Monitor-Topology on DC=contoso,DC=com

   - DS-Replication-Monitor-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com

   - DS-Replication-Monitor-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

   - DS-Replication-Synchronize on DC=contoso,DC=com

   - DS-Replication-Synchronize on CN=Configuration,DC=concorp,DC=contoso,DC=com

   - DS-Replication-Synchronize on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

- Change-RID-Master on CN=RID Manager$,CN=System,DC=contoso,DC=com
- Change-Infrastructure-Master on DC=contoso,DC=com
- Change-PDC on DC=contoso,DC=com

**Permissions on CN=Sites,CN=Configuration,DC=concorp,DC=contoso,DC=com:**

- Inherit-only Full-Control permissions on all objects of class NTDS-Settings
- Inherit-only Create Child permissions to be able to create and delete objects of class Server
- Inherit-only Delete Child permissions to be able to delete objects of class NTDS-Settings
- Inherit-only Delete Child permissions to be able to delete objects of class Server

**Inheritable Full Control permissions on:**

- CN=System,DC=contoso,DC=com
- CN=System,DC=noam,DC=concorp,DC=contoso,DC=com
- CN=System,DC=europe,DC=concorp,DC=contoso,DC=com

| **Note** |
| --- |
| If the forest root domain is not in native mode, universal groups are not available. In this case, create a new domain local group in each domain and grant that group the permissions specified earlier in this procedure. Then make the Contoso Root Dom Config Ops global group a member of these domain local groups. |

**Create Child permissions on:**

- OU=Domain Controllers,DC=contoso,DC=com
- OU=Domain Controllers,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Domain Controllers,DC=europe,DC=concorp,DC=contoso,DC=com

| **Note** |
| --- |
| If the forest root domain is not in native mode, universal groups are not available. In this case, create a new domain local group in each domain and grant that group the permissions specified earlier in this procedure. Then make the Contoso Root Dom Config Ops global group a member of these domain local groups. |

**Property permissions Read-Property and Write-Property to the FSMO-Role-Owner property on:**

- DC=concorp,DC=contoso,DC=com
- CN=RID Manager$,CN=System,DC=concorp,DC=contoso,DC=com
- CN=Infrastructure,DC=concorp,DC=contoso,DC=com

3. Grant the Creator Owner Trustee Full Control on:
   - The computer object that represents the server that will be promoted to create a new domain

- CN=sites,CN=Configuration,DC=concorp,DC=contoso,DC=com

4. Grant the permission Enable computer and user accounts to be trusted for delegation by adding Contoso Root Dom Config Ops to the User Rights Assignment setting in Domain Controller Security Policy for concorp.contoso.com.
5. Prior to creating a new domain, add Contoso Root Dom Config Ops to the local Administrators group of the member server that is to be promoted to a domain controller.

**To implement the Domain Configuration Operators role, NOAM.contoso.com domain instance**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** NOAM Dom Config Ops
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global
2. Grant the following permissions to NOAM Dom Config Ops:

   **Extended rights:**

   - DS-Replication-Get-Changes on DC=noam,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Get-Changes on CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Get-Changes on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Get-Changes-All on DC=noam,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Get-Changes-All on CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Get-Changes-All on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Manage-Topology on DC=noam,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Manage-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Manage-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on DC=noam,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Synchronize on DC=noam,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Synchronize on CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Synchronize on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com”
   - Change-RID-Master on CN=RID Manager$,CN=System,DC=noam,DC=concorp,DC=contoso,DC=com
   - Change-Infrastructure-Master on DC=noam,DC=concorp,DC=contoso,DC=com
   - Change-PDC on DC=noam,DC=concorp,DC=contoso,DC=com

**Permissions on CN=Sites,CN=Configuration,DC=concorp,DC=contoso,DC=com:**

- Inherit-only Full-Control permissions on all objects of class NTDS-Settings
- Inherit-only Create Child permissions to be able to create and delete objects of class Server
- Inherit-only Delete Child permissions to be able to delete objects of class NTDS-Settings
- Inherit-only Delete Child permissions to be able to delete objects of class Server

**Inheritable Full Control permissions on:**

- CN=System,DC=noam,DC=concorp,DC=contoso,DC=com

| **Note** |
| --- |
| If the forest root domain is not in native mode, universal groups are not available. In this case, create a new domain local group in each domain and grant that group the permissions specified earlier in this procedure. Then make the NOAM Dom Config Ops global group a member of these domain local groups. |

**Create Child permissions on:**

- OU=Domain Controllers,DC=noam,DC=concorp,DC=contoso,DC=com

| **Note** |
| --- |
| If the forest root domain is not in native mode, universal groups are not available. In this case, create a new domain local group in each domain and grant that group the permissions specified earlier in this procedure. Then make the NOAM Dom Config Ops global group a member of these domain local groups. |

**Property permissions Read-Property and Write-Property permissions to the FSMO-Role-Owner property on:**

- DC=noam,DC=concorp,DC=contoso,DC=com
- CN=RID Manager$,CN=System,DC=noam,DC=concorp,DC=contoso.com
- CN=Infrastructure,DC=noam,DC=concorp,DC=contoso,DC=com

3. Grant the Creator Owner trustee Full Control on:
   - The computer object that represents the server that will be promoted to create a new domain
   - CN=sites,CN=Configuration,DC=concorp,DC=contoso,DC=com
4. Grant the permission Enable computer and user accounts to be trusted for delegation by adding NOAM Dom Config Ops to the User Rights Assignment setting in Domain Controller Security Policy for the noam.concorp.contoso.com.
5. Prior to creating a new domain, add NOAM Dom Config Ops to the local Administrators group of the member server that is to be promoted to a domain controller.

**To implement the Domain Configuration Operators role, Europe.contoso.com domain instance**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** EUROPE Dom Config Ops
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global

2. Grant the following permissions to EUROPE Dom Config Ops:

**Extended rights:**

- DS-Replication-Get-Changes on DC=europe,DC=concorp,DC=contoso,DC=com
- DS-Replication-Get-Changes on CN=Configuration,DC=concorp,DC=contoso,DC=com
- DS-Replication-Get-Changes on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
- DS-Replication-Get-Changes-All on DC=europe,DC=concorp,DC=contoso,DC=com
- DS-Replication-Get-Changes-All on CN=Configuration,DC=concorp,DC=contoso,DC=com
- DS-Replication-Get-Changes-All on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
- DS-Replication-Manage-Topology on DC=europe,DC=concorp,DC=contoso,DC=com
- DS-Replication-Manage-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com
- DS-Replication-Manage-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
- DS-Replication-Monitor-Topology on DC=europe,DC=concorp,DC=contoso,DC=com
- DS-Replication-Monitor-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com
- DS-Replication-Monitor-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
- DS-Replication-Synchronize on DC=europe,DC=concorp,DC=contoso,DC=com
- DS-Replication-Synchronize on CN=Configuration,DC=concorp,DC=contoso,DC=com
- DS-Replication-Synchronize on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
- Change-RID-Master on CN=RID Manager$,CN=System,DC=europe,DC=concorp,DC=contoso,DC=com
- Change-Infrastructure-Master on DC=europe,DC=concorp,DC=contoso,DC=com
- Change-PDC on DC=europe,DC=concorp,DC=contoso,DC=com

**Permissions on CN=Sites,CN=Configuration,DC=concorp,DC=contoso,DC=com:**

- Inherit-only Full-Control permissions on all objects of class NTDS-Settings
- Inherit-only Create Child permissions to be able to create and delete objects of class Server
- Inherit-only Delete Child permissions to be able to delete objects of class NTDS-Settings
- Inherit-only Delete Child permissions to be able to delete objects of class Server

**Inheritable Full Control permissions on:**

- CN=System,DC=europe,DC=concorp,DC=contoso,DC=com

| Note |
| --- |
| If the forest root domain is not in native mode, universal groups are not available. In this case, create a new domain local group in each domain and grant that group the permissions specified earlier in this procedure. Then make the EUROPE Dom Config Ops global group a member of these domain local groups. |

**Create Child permissions on:**

- OU=Domain Controllers,DC=europe,DC=concorp,DC=contoso,DC=com


| Note |
| --- |
| If the forest root domain is not in native mode, universal groups are not available. In this case, create a new domain local group in each domain and grant that group the permissions specified earlier in this procedure. Then make the EUROPE Dom Config Ops global group a member of these domain local groups. |


**Property permissions Read-Property and Write-Property permissions to the FSMO-Role-Owner property on:**

- DC=europe,DC=concorp,DC=contoso,DC=com
- CN=RID Manager$,CN=System,DC=europe,DC=concorp,DC=contoso.com
- CN=Infrastructure,DC=europe,DC=concorp,DC=contoso,DC=com

3. Grant the Creator Owner trustee Full Control on:
   - The computer object that represents the server that will be promoted to create a new domain
   - CN=sites,CN=Configuration,DC=concorp,DC=contoso,DC=com
4. Grant **Enable computer and user accounts to be trusted for delegation** by adding EUROPE Dom Config Ops to the User Rights Assignment setting in Domain Controller Security Policy for the europe.concorp.contoso.com.
5. Prior to creating a new domain, add EUROPE Dom Config Ops to the local Administrators group of the member server that is to be promoted to a domain controller.

**Implementing Schema Admins**
A member of Enterprise Admins creates one instance of this role.

**To implement the Schema Admins role**

- Move the default security group Schema Admins from its default location in the Users container to the Service Management OU. By default, the Schema Admins group has sufficient permissions to manage the Active Directory schema.

**Implementing Replication Management Admins**
A member of Enterprise Admins creates one instance of this role.

**To implement the Replication Management Admins role**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** Contoso Repl Mgmt Admins
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global. “Native mode” has the following meaning, depending on the operating system:

     **Windows Server 2003:** the domain functional level is Windows 2000 native or Windows Server 2003.

**Windows 2000:** the domain mode is native.

2. Grant the following permissions to Contoso Repl Mgmt Admins:

   **Extended rights:**

   - DS-Replication-Manage-Topology on DC=concorp,DC=contoso,DC=com
   - DS-Replication-Manage-Topology on DC=noam,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Manage-Topology on DC=europe,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Manage-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Manage-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on DC=noam,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on DC=europe,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com
   - Full Control permissions on CN=Sites,CN=Configuration,DC=concorp,DC=contoso,DC=com:

**Implementing Replication Monitoring Operators**

A member of Enterprise Admins creates one instance of this role.

**To implement the Replication Monitoring Operators role**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:

   - **Group Name:** Contoso Repl Monitoring Ops
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global. “Native mode” has the following meaning, depending on the operating system:

     **Windows Server 2003:** the domain functional level is Windows 2000 native or Windows Server 2003.

     **Windows 2000:** the domain mode is native.

2. Grant the following permissions to Contoso Repl Monitoring Ops:

   **Extended rights:**

   - DS-Replication-Monitor-Topology on DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on DC=noam,DC=concorp,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on DC=europe,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on CN=Configuration,DC=concorp,DC=contoso,DC=com
   - DS-Replication-Monitor-Topology on CN=Schema,CN=Configuration,DC=concorp,DC=contoso,DC=com

| Note |
| --- |
| If your Active Directory environment is a Windows 2000 environment, the Replication-Monitor-Topology extended right is not available. In this case, grant the Replication-Manage-Topology extended right on all of the objects specified in the preceding step. |

- Full Control permissions on CN=Sites,CN=Configuration,DC=concorp,DC=contoso,DC=com:

**Implementing DNS Admins**

A member of Enterprise Admins creates four instances of this role, as follows:

- Contoso Forest DNS Admins for the entire forest
- Contoso DNS Admins for the concorp.contoso.com domain
- NOAM DNS Admins for the noam.concorp.contoso.com domain
- Europe DNS Admins for the europe.concorp.contoso.com domain

**To implement the DNS Admins role, Concorp instance**

1. First the domain instances are implemented, and then the forest root instance, which must be added to the security groups for each domain instance.
2. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** Concorp DNS Admins
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global. "Native mode" has the following meaning, depending on the operating system:

     **Windows Server 2003:** the domain functional level is Windows 2000 native or higher.

     **Windows 2000:** the domain mode is native.
3. Grant the following permissions to Concorp DNS Admins:
   - Full Control on CN=MicrosoftDNS,CN=System,DC=concorp,DC=contoso,DC=com
   - Full Control on CN=MicrosoftDNS,DC=DomainDnsZones,DC=concorp,DC=contoso,DC=com

**To implement the DNS Admins role, NOAM DNS Admins instance**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** NOAM DNS Admins
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global
2. Grant the following permissions to NOAM DNS Admins:
   - Full Control on CN=MicrosoftDNS,CN=System,DC=noam,DC=concorp,DC=contoso,DC=com
   - Full Control over CN=MicrosoftDNS,DC=DomainDnsZones,DC=noam,DC=concorp,DC=contoso,DC=com

**To implement the DNS Admins role, Europe instance**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** Europe DNS Admins
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global
2. Grant the following permissions to Europe DNS Admins:
   - Full Control on CN=MicrosoftDNS,CN=System,DC=europe,DC=concorp,DC=contoso,DC=com
   - Full Control over CN=MicrosoftDNS,DC=DomainDnsZones,DC=europe,DC=concorp,DC=contoso, DC=com

**To implement the DNS Admins role, Contoso Forest instance**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** Contoso Forest DNS Admins
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global
2. Add the Contoso Forest DNS Admins group to the following groups:
   - Concorp DNS Admins
   - NOAM DNS Admins
   - Europe DNS Admins
3. Grant the following permissions to Contoso Forest DNS Admins:
   - Full Control on CN=MicrosoftDNS,DC=ForestDnsZones,DC=concorp,DC=contoso,DC=com

**Implementing Security Policy Admins**
A member of Enterprise Admins creates one instance of this role.

**To implement the Security Policy Admins role**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** Contoso Sec Pol Admins
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global. “Native mode” has the following meaning, depending on the operating system:

     **Windows Server 2003:** the domain functional level is Windows 2000 native or higher.

     **Windows 2000:** the domain mode is native.
2. Grant the following permissions to Contoso Sec Pol Admins on each of the following objects:

   **Objects:**
   - DC=concorp,DC=contoso,DC=com
   - DC=noam,DC=concorp,DC=contoso,DC=com
   - DC=europeDC=concorp,DC=contoso,DC=com

- OU=Domain Controllers,DC=concorp,DC=contoso,DC=com
- OU=Domain Controllers,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Domain Controllers,DC=europe,DC=concorp,DC=contoso,DC=com

**Permissions:**

- Read-property permissions to read the GP-Link property
- Write-property permissions to read the GP-Link property
- Read-property permissions to read the GP-Options property
- Write-property permissions to read the GP-Options property

3. In the properties page for the Domain Controllers OU in each domain, on the **Group Policy** tab, select the **Default Domain Controllers GPO** and on the **Security** tab, grant Full Control to Contoso Sec Pol Admins.

**Implementing Service Admin Managers**

A member of Enterprise Admins creates one instance of this role.

**To implement the Service Admin Managers role**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** Contoso Srvc Admin Managers
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global. "Native mode" has the following meaning, depending on the operating system:

     **Windows Server 2003:** the domain functional level is Windows 2000 native or higher.

     **Windows 2000:** the domain mode is native.

2. In each domain in the forest, modify permissions on the object CN=AdminSDHolder,CN=System,DC=*Domain Name*, as follows:
   - Grant Contoso Srvc Admin Managers Full Control on the object.
   - Revoke permissions granted to the Domain Admins and Builtin Administrators groups.

     | **Note** |
     | --- |
     | Although Builtin Administrators will no longer have control over service administrative accounts because of the security descriptor changes in step 2, Builtin Administrators have the privilege Take ownership of objects and other files (a User Rights Assignment setting in Default Domain Controller Security Policy). This privilege can be exercised to take ownership of (and subsequently modify the DACL of) any object, including the service administrative groups and all accounts that are members of these security groups. |

3. Mark the DACL of the Service Management OU as Protected.
4. Grant Contoso Srvc Admin Managers inheritable Full Control over the Service Management OU.
5. Modify the DACLs of every group that is created for the purpose of representing roles, as follows:
   - Revoke all permissions that are granted to Domain Admins.

- Revoke all permissions that are granted to Account Operators.

**Implementing Domain Controller Admins**
A member of Enterprise Admins creates two instances of the Domain Controller Admins role:

- Contoso Root & NOAM DC Admins
- Europe DC Admins

**To implement the Domain Controller Admins role, Contoso Root & NOAM DC Admins instance**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** Contoso Root & NOAM DC Admins
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global. "Native mode" has the following meaning, depending on the operating system:

     **Windows Server 2003:** the domain functional level is Windows 2000 native or higher.

     **Windows 2000:** the domain mode is native.
2. Add Contoso Root & NOAM DC Admins to the security group CN=Administrators,CN=Builtin,DC=contoso,DC=com
3. Add Contoso NOAM DC Admins to the security group CN=Administrators,CN=Builtin,DC=noam,DC=concorp,DC=contoso,DC=com
4. Make sure no other group is a member of the Builtin Administrators group in noam.concorp.contoso.com.

**To implement the Domain Controller Admins role, Europe DC Admins instance**

1. In the concorp.contoso.com domain, create the following security group under the Service Management OU:
   - **Group Name:** Europe DC Admins
   - **Group Type:** Universal if the forest root domain is in native mode, otherwise global. "Native mode" has the following meaning, depending on the operating system:

     **Windows Server 2003:** the domain functional level is Windows 2000 native or higher.

     **Windows 2000:** the domain mode is native.
2. Add Contoso Europe DC Admins to the security group CN=Administrators,CN=Builtin,DC=europe,DC=concorp,DC=contoso,DC=com
3. Make sure no other group is a member of the Builtin Administrators group in europe.concorp.contoso.com.

**Implementing Backup Operators**
A member of Enterprise Admins creates three instances of this role, one for each domain.

**To implement the Backup Operators role**

1. In the concorp.contoso.com domain, create the following security groups under the Service Management OU:

   **Group Name:**

   - Contoso Root Backup Operators
   - NOAM Backup Operators
   - Europe Backup Operators

**Group Type:** Universal if the forest root domain is in native mode, otherwise global. "Native mode" has the following meaning, depending on the operating system:

**Windows Server 2003:** the domain functional level is Windows 2000 native or higher.

**Windows 2000:** the domain mode is native.

2. In Domain Controller Security Policy in each domain, grant the following user rights under **User Rights Assignment** to the respective instance of the Backup Operators role:
   - Allow log on locally
   - Back up files and directories
   - Shut down the system
3. In Domain Controller Security Policy in each domain, modify the **User Rights Assignment** settings to remove all user rights settings that are granted by default to the Builtin\Backup Operators group, which are:
   - Allow log on locally
   - Back up files and directories
   - Restore files and directories
   - Shut down the system
4. Make sure that no other group is granted these user rights.

# Step 3 — Hand Off Data Management to Contoso Data Administrators

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before the data management administrative delegation model can be implemented, control over data needs to be transferred to data administrators. Conceptually, the service owner hands off ownership and responsibility of content management to data owners. From an implementation perspective, high-level service administrators, who represent the service owners, delegate full control of business unit content to data administrators, who represent data owners.

Service administrators perform the following tasks to hand off data management:

1. Create a Business Units OU in each domain, and one OU for every business unit within the Business Units OU in each domain.

   Contoso has decided not to use the forest root domain for storing business accounts, but to store all business unit data in the NOAM and Europe regional domains.

2. Create a Delegation OU in each domain to store the security groups that represent the instances of Business Unit Admins roles.

3. Create security groups to represent the Business Unit Admins role instances.

4. Grant the security groups full control over the respective business unit OUs.

5. Optionally, grant these security groups the ability to modify their own group memberships.

6. Create user accounts for the Business Unit Admins role holders.

7. Add Business Unit Admins user accounts to the respective security groups that represent instances of this role.

**Creating a Business Units OU Hierarchy in Each Domain**

A member of Domain Configuration Operators in each domain creates one OU called Business Units in the respective domain. Below this OU, the administrator creates one OU for each business unit.

The following steps are performed by a member of NOAM Domain Config Ops and Europe Config Ops in their respective domains.

1. Create Business Units OU roots for each domain by creating the objects in Table 26.

   **Table 26 Business Units OUs for NOAM and Europe Domains**

   | Object DN | Object Class |
   |---|---|
   | OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | organizationalUnit |
   | OU=Business Units,DC= europe,DC=concorp,DC=contoso,DC=com | organizationalUnit |

2. Create one OU for each business unit within the Business Units OU in each domain by creating the objects in Table 27.

   **Table 27 OUs for Each Business Unit in NOAM and Europe Domains**

   | Object DN | Object Class |
   |---|---|
   | OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | organizationalUnit |
   | OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso, DC=com | organizationalUnit |
   | OU=BusMgmt,DC=noam,DC=concorp,DC=contoso,DC=com | organizationalUnit |

| | |
|---|---|
| OU=BusMgmt,DC= europe,DC=concorp,DC=contoso,DC=com | organizationalUnit |
| OU=IT,DC=noam,DC=concorp,DC=contoso,DC=com | organizationalUnit |
| OU=IT,DC=europe,DC=concorp,DC=contoso,DC=com | organizationalUnit |

**Creating a Delegation OU in Each Domain**

In each domain that will store business unit data, a member of the Domain Configuration Operators role for each domain (NOAM Domain Config Ops and Europe Domain Config Ops) creates one OU named Delegation within the Business Units OU for the domain. This OU is used to store the security groups that represent as many instances of the Business Unit Admins role as are needed (usually equal to the number of business units in the domain). It is also used to store any data management roles that need to be granted administrative authority across all business units.

To create the Delegation OUs, a Domain Configuration Operator for each domain creates the objects in Table 28.

**Table 28 Delegation OUs for each Domain in NOAM and Europe**

| Object DN | Object Class |
|---|---|
| OU=Delegation,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | organizationalUnit |
| OU=Delegation,OU=Business Units,DC=europe,DC=concorp,DC=contoso,DC=com | organizationalUnit |

Figure 13 shows the resulting domain and OU structures for NOAM and Europe.




**Creating Security Groups to Represent the Business Unit Admins Role Instances**

The Domain Configuration Operator next creates one security group to represent each instance of the Business Unit Admins role for each business unit in the respective domains.

Table 29 shows the group objects that the Domain Configuration Operators create in the Delegation OU within the Business Units OUs in the respective domains.

**Table 29 OUs for Each Business Unit in NOAM and Europe Domains**

| Object DN | Object Class |
|---|---|
| CN=RandD Bus Unit Admins,OU=Delegation,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | Security group |
| CN=Production Bus Unit Admins,OU=Delegation,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | Security group |
| CN=BusMgmt Bus Unit Admins,OU=Delegation,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | Security group |

| | |
|---|---|
| CN=BusMgmt Bus Unit Admins,OU=Delegation,OU=Business Units,DC=europe,DC=concorp,DC=contoso,DC=com | Security group |
| CN=IT Bus Unit Admins,OU=Delegation,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | Security group |
| CN=IT Bus Unit Admins,OU=Delegation,OU=Business Units,DC=europe,DC=concorp,DC=contoso,DC=com | Security group |

Figure 14 shows the domain and OU structures for each domain, including the security groups within their respective Delegation OUs.




**Granting Business Unit Admins Full-Control Over Business Unit OUs**

The Domain Configuration Operators for each domain grant each instance of the Business Unit Admins role full control over their respective business unit OUs in the domain. To do so, they modify the permissions on each business unit OU and grant the appropriate security group full control over the OU.

Table 30 shows the business unit OUs and the respective security groups that represent the Business Unit Admins role and receive Full Control permissions on the OU.

**Table 30 OUs and Business Unit Admins Groups That Receive Full Control**

| OUs | Allow Full Control To |
|---|---|
| OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | RandD Bus Unit Admins |
| OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | Production Bus Unit Admins |
| OU=BusMgmt,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | BusMgmt Bus Unit Admins |
| OU=IT,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com | IT Bus Admins |
| OU=BusMgmt,OU=Business Units,DC=europe,DC=concorp,DC=contoso,DC=com | BusMgmt Bus Admins |
| OU=IT,OU=Business Units,DC=europe,DC=concorp,DC=contoso,DC=com | IT Bus Admins |

**Granting the Ability to Modify Administrative Group Memberships**
In some cases, the Domain Configuration Operators group might grant security groups that represent the various role instances of the Business Unit Admins role sufficient permissions to modify the membership of these groups. The objective is to allow Business Unit Admins the ability to control their own group membership. Based on the administrative requirements of the Contoso organization, the service owners decide to grant this ability to the groups that represent instances of the Business Unit Admins role.

To enable group members to change the membership of their respective groups, a member of NOAM Domain Config Ops and Europe Domain Config Ops grants to each Business Admins security group the Write property permission to modify the Member attribute on their own group object:

- A member of the NOAM Domain Config Ops group grants each of the following security groups permission to modify the Member attribute on the object that represents the respective security group:
    - RandD BU Admins
    - Production BU Admins
    - BusMgmt BU Admins
    - IT BU Admins
- A member of the Europe Domain Config Ops group grants each of the following security groups permission to modify the Member attribute on the object that represents the respective security group:
    - BusMgmt BU Admins
    - IT BU Admins

At this point, all of the Business Unit Admins roles have been enabled by creating the Business Unit Admins groups and granting them permissions to manage their respective OUs. To delegate the roles, the Domain Configuration Operators next create the user accounts that will perform each role and add them to the appropriate groups.

**Creating User Accounts for Business Unit Admins Groups**
Data owners for each business group have communicated the identities of the users who will serve as the Business Unit Admins to the Domain Configuration Operators. The Domain Configuration Operators create these user accounts in the respective business unit OUs, as shown in Table 31.

**Table 31 Business Unit Administrator Accounts**

| Business Unit/Domain | Business Unit Admins Role Assignments |
|---|---|
| RandD/NOAM | John<br>Chris |
| Production/NOAM | Mary<br>Joe |
| Bus Mgmt/NOAM | Sally |
| IT/NOAM | Kevin |
| Bus Mgmt/Europe | Frank |
| IT/Europe | Anna |

Accordingly, the Domain Config Operators, on behalf of the service owners, create the following user objects:

- CN=John,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- CN=Chris,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- CN=Mary,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- CN=Joe,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- CN=Sally,OU=BusMgmt,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- CN=Frank,OU=BusMgmt,OU=Business Units,DC=europe,DC=concorp,DC=contoso,DC=com
- CN=Kevin,OU=IT,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- CN=Anna,OU=IT,OU=Business Units,DC=europe,DC=concorp,DC=contoso,DC=com

Note that user accounts for Business Unit Admins of the BusMgmt and IT business units are created in different domains in order to spread them across all domains that have the same business unit.

**Adding Business Unit Admins User Accounts to Administrative Security Groups**
To actually delegate the Business Unit Admins role and to complete the data management handoff, the Domain Configuration Operators add the users whose accounts they have created to the security groups that represent the respective Business Unit Admins roles.

Table 32 shows the resulting Business Unit Admins group memberships. At this point, the data management handoff is complete. All Business Unit Admins have full control over their business unit OUs.

**Table 32 Business Unit Admin Role Security Groups and Added Members**

| Group for Business Unit Admins Role | User Accounts | Business Unit OU | Domain |
|---|---|---|---|
| RandD BU Admins | John<br>Chris | RandD | NOAM |
| Production BU Admins | Mary<br>Joe | Production | NOAM |
| Bus Mgmt BU Admins | Frank | Bus Mgmt | NOAM |
| IT BU Admins | Kevin | IT | NOAM |
| Bus Mgmt BU Admins | Sally | Bus Mgmt | Europe |
| IT BU Admins | Anna | IT | Europe |

# Step 4 — Create the Contoso Data Management Administrative Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With the business unit OUs and Business Unit Admins role instances in place, Contoso data owners are ready to create their delegation model according to the following criteria:

- **Objective:** Create a delegation model to distribute data management among data administrators by enabling efficient, security-conscious delegation and distribution of administrative responsibilities among various administrative groups
- **Stakeholders:** Data owners of all business units
- **Approach:** Each business unit data owner creates a delegation model for his or her business unit data, as follows:
  - Understand administrative delegation and Group Policy application requirements.
  - Create an OU structure that meets administrative delegation and Group Policy application requirements.
  - Identify the number of instances of each data management role according to business unit requirements and identify the administrative personnel who are assigned to each role.

Contoso has four business units spread across two domains and six locations. Table 33 shows the locations of the users, workstations, and servers that are managed by each business unit.

**Table 33 Distribution of Users, Workstations, and Servers in Contoso Business Units**

| Business Units | Locations | Users | Workstations | Servers |
|---|---|---|---|---|
| Research and Development | Chicago | 5,000 | 10,000 | 1000 |
| Production | Atlanta | 3,000 | 5,000 | 520 |
| Business Management | Chicago, New York, London, Paris, Rome | 5,500 | 7,000 | 250 |
| IT | Chicago, Atlanta, New York, London, Paris, Rome | 2,500 | 5,000 | 100 |

Table 34 shows the distribution of servers across the Contoso business units by type of server.

**Table 34 Distribution of Servers by Type in Contoso Business Units**

| Business Unit | File Servers | Web Servers | Database Servers | Application Servers |
|---|---|---|---|---|
| RandD | 800 | 50 | 100 | 50 |
| Production | 300 | 20 | 100 | 100 |
| Bus Mgmt | 120 | 45 | 50 | 35 |
| IT | 30 | 30 | 20 | 20 |

To support the delegation model, Contoso assigns the following data management roles:

- Account Admins
- Workstation Admins

- Resource Admins
- Application-specific roles
- Custom roles

**Creating the Delegation Model for the RandD Business Unit**

This business unit is based in Chicago and is the main research and development unit. All of its 500 employees are located in a single building. The workstation-to-employee ratio is roughly 2:1, resulting in a total of 1,000 workstations. Additionally, there are about 1,000 servers, including file, Web, database, and application servers.

Servers play different roles and are administered by different groups of administrators, depending on the resource to which they belong.

Table 35 shows the number of users, workstations, and servers that are stored in the RandD business unit in the Chicago location.

**Table 35 Users, Workstations, and Servers in the RandD Business Unit**

| Location | Users | Workstations | Servers |
|---|---|---|---|
| Chicago | 5,000 | 10,000 | 1,000 |

Table 36 shows the numbers of server types that RandD stores.

**Table 36 Distribution of Server Types in RandD Business Unit**

| Location | File Servers | Web Servers | Database Servers | Application Servers |
|---|---|---|---|---|
| Chicago | 800 | 50 | 100 | 50 |

**RandD Administrative and Group Policy Requirements**

Administrative requirements for the RandD business unit include the management of users, workstations, and resources, as follows:

- **User Account Management.** All users are managed by a single account management team.
- **Workstation Management.** All workstations are managed by a single workstation management team that is located on site.
- **Resource Management.** Four types of resources must be managed: file servers, Web portals, database servers, and applications that are hosted on servers. For every type of resource, there is one administrative group that has overall responsibility for that type of resource. These administrative groups require the ability to sub-delegate resource management to different administrative teams.

Group Policy requirements include the settings that must be applied for different types of user accounts, for workstations (both desktop and portable computers), and for resources according to resource type, as follows:

- **User Accounts.** Requirements for folder redirection and other user configuration settings necessitate that different Group Policy settings be applied for development accounts and research accounts.
- **Workstations.** Requirements for scripts and other computer configuration settings necessitate that different Group Policy settings be applied for desktop and portable computers.
- **Resources.** Computer configuration settings necessitate that different Group Policy settings be applied for different kinds of resources and might require the application of specific Group Policy settings for the various specific resources.

**RandD OU Structure Based on Administrative and Group Policy Requirements**

Figure 15 shows the OU structure for the RandD OU that accommodates its administrative and Group Policy requirements.




Table 37 shows the rationale for the OU structure shown in Figure 15.

**Table 37 Purpose of Each OU in the RandD Business Unit OU Hierarchy**

| Organizational Unit | Purpose |
|---|---|
| User Accounts | Main OU to store user accounts<br><br>Delegation point for Account Admins role |
| User Accounts\Research | Used to apply Group Policy for research user accounts |
| User Accounts\Development | Used to apply Group Policy for development user accounts |
| Workstations | Main OU to store computer accounts for workstations<br><br>Delegation point for Workstation Admins role |
| Workstations\Desktops | Used to apply Group Policy for computer accounts for desktops |
| Workstations\Laptops | Used to apply Group Policy for computer accounts for laptops |
| Resources | Main OU to store computer accounts for servers hosting resources<br><br>Delegation point for Resource Admins role |
| Resources\File Servers | Main OU to store computer accounts for file servers<br><br>Used to apply Group Policy for file servers |
| Resources\Web Servers | Main OU to store computer accounts for Web servers<br><br>Used to apply Group Policy for Web servers |
| Resources\Database Servers | Main OU to store computer accounts for database servers<br><br>Used to apply Group Policy for database servers |
| Resources\Application Servers | Main OU to store computer accounts for application servers<br><br>Used to apply Group Policy for application servers |

**RandD Role Instances Based on Business Unit Requirements**

On the basis of the RandD business unit requirements, the following role instances must be created:

- **User Account management.** Because a single group is responsible for all aspects of user account management, one instance of the Account Admins role should meet the requirements.

- **Workstation management.** Because a single group is responsible for all aspects of workstation account management, one instance of the Workstation Admins role should meet the requirements.

- **Resource management.** Because a single administrative group has overall responsibility for resource account management, one instance of the Resource Admins role is required for this group. Additionally, this group requires the

ability to sub-delegate authority to other administrative groups and thus is responsible for creating as many instances of the Resource Admins role as needed.

Note that the role instance for the single Resource Admins group that has power to sub-delegate encompasses all aspects of resource administration. Although the specific instances that are sub-delegated receive more limited authority, they are still Resource Admin roles. By granting this one Resource Admins role the ability to sub-delegate, the data owners understand that they are effectively granting this role instance the ability create OUs, and hence to create data such as users and workstations. Although the data owners have made these actions technically possible for the administrators to whom the role is delegated, the data owners have implemented specific business policies stating that should these Resource Admins misuse their granted authority, they will be subject to punitive action up to and including termination of employment.

Table 38 shows the model creation template that the data owners fill out to document the RandD Account Admins role.

**Table 38 Model Creation Template for RandD Account Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | RandD Account Admins |
| Instance of | Account Admins |
| Instance Number | 1 of 1 |
| Assigned Administrators | Joe, Mike, Sara, Kevin |
| Assigned Tasks | Manage all aspects of account management |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation/Implementation |

Table 39 shows the model creation template that the data owners fill out to document the RandD Workstation Admins role.

**Table 39 Model Creation Template for RandD Workstation Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | RandD Workstation Admins |
| Instance of | Account Admins |
| Instance Number | 1 of 1 |
| Assigned Administrators | Eugene, Larry, Andy, Peter |
| Assigned Tasks | Manage all aspects of workstation management |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation/Implementation |

Table 40 shows the model creation template that the data owners fill out to document the RandD Resource Admins role.

**Table 40 Model Creation Template for RandD Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name: | RandD Resource Admins |
| Instance of: | Account Admins |
| Instance Number: | 1 of n, where n = as many as needed over time |
| Assigned Administrators: | Deborah, Paul |
| Assigned Tasks: | Overall responsibility for all aspects of resource management |

| | |
|---|---|
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation/Implementation |

**Creating the Delegation Model for the Production Business Unit**

This business unit is based in Atlanta and is responsible for all content for the Atlanta production facility. Approximately 300 employees work for this business unit and approximately 500 workstations are in use at the production facility. Approximately 320 servers provide various required services.

Atlanta has two physical locations that are approximately five miles apart. The servers in these locations play different roles and are managed by different groups of administrators according to resource type.

Because all data for this business unit is physically located in Atlanta, all of its content is stored only in the noam.concorp.contoso.com domain.

Table 41 shows the number of users, workstations, and servers that are stored in the Production business unit in the Atlanta location.

**Table 41 Users, Workstations, and Servers in the Production Business Unit**

| Location | Users | Workstations | Servers |
|---|---|---|---|
| Atlanta | 300 | 500 | 320 |

Table 42 shows the numbers of server types that Production stores.

**Table 42 Distribution of Server Types in the Production Business Unit**

| Location | File Servers | Web Servers | Database Servers | Application Servers |
|---|---|---|---|---|
| Atlanta | 300 | 20 | 100 | 100 |

**Production Administrative and Group Policy Requirements**

Administrative requirements for the Production business unit include the management of users, workstations, and resources, as follows:

- **User Account Management.** All user accounts are managed by a single account management team.
- **Workstation Management.** Each physical location has a separate team that is responsible for managing workstations.
- **Resource Management.** Three production applications are run in the production facility. Each application has its own file, Web, database, and application servers. In addition, some file, Web, database, and application servers are common to the entire facility. Each of the three applications has its own administrators, and one group of administrators is responsible for managing the common servers.

Group Policy requirements include the settings that must be applied for different types of user accounts, for workstations (both desktops and laptops), and for resources according to the resource type, as follows:

- **User Accounts.** One common Group Policy applies to all user accounts in the production business unit.
- **Workstations.** Requirements for scripting and other computer configuration settings necessitate different Group Policy settings for desktop and portable computers.
- **Resources.** Computer configuration settings necessitate that different Group Policy settings be applied for different kinds of resources and might require the application of specific Group Policy settings for the various specific resources.

**Production OU Structure Based on Administrative and Group Policy Requirements**

Figure 16 shows the OU structure for the Production OU that accommodates its administrative and Group Policy requirements.




Table 43 shows the rationale for the OU structure shown in Figure 16.

**Table 43 Purpose of Each OU in the Production Business Unit OU Hierarchy**

| Organizational Unit | Purpose |
|---|---|
| Accounts | Main OU to store accounts<br><br>Delegation point for Account Admins role |
| Workstations | Main OU to store workstation computer accounts<br><br>Delegation point for Workstation Admins role |
| Workstations\Desktops | Used to apply Group Policy for desktops |
| Workstations\Laptops | Used to apply Group Policy for portable computers |
| Resources | Main OU to store servers comprising resources |
| Resources\Production Application 1 | Main OU to store all servers that are part of Application 1 — includes file, Web, database and application servers |
| Resources\Production Application 2 | Main OU to store all servers that are part of Application 2 — includes file, Web, database and application servers |
| Resources\Production Application 3 | Main OU to store all servers that are part of Application 3 — includes file, Web, database and application servers |
| Resources\Common Resources | Main OU to store all servers that are part of the common resource pool |

**Production Role Instances Based on Business Unit Requirements**

On the basis of the Production business unit requirements, the following role instances will be created:

- **Account Management.** Because a single group is responsible for all aspects of account management, one instance of

the Account Admins role should meet the requirements.

- **Workstation Management.** Because a different administrative group is required for each of the two physical locations in Atlanta, two instances of the Workstation Admins role should meet the requirements.

- **Resource Management.** Based on the business unit requirements, a total of four Resource Admins role instances is required — one each for the three production applications and one for the common set of resources.

Table 44 shows the model creation template that the data owners fill out to document the Production Account Admins role.

**Table 44 Model Creation Template for Production Account Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Production Account Admins |
| Instance of | Account Admins |
| Instance Number | 1 of 1 |
| Assigned Administrators | Sandra |
| Assigned Tasks | Manage all aspects of account management |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 45 shows the model creation template that the data owners fill out to document the Production Workstation Admins role in the first Atlanta location.

**Table 45 Model Creation Template for Production Workstation Admins Role in Location 1**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Production Location 1 Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 1 of 2 |
| Assigned Administrators | Michael, Dave |
| Assigned Tasks | Manage all aspects of workstation management for Location 1 |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 46 shows the model creation template that the data owners fill out to document the Production Workstation Admins role in the second Atlanta location.

**Table 46 Model creation template for Production Workstation Admins role in Location 2**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Production Location 2 Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 2 of 2 |
| Assigned Administrators | Adam, Charlotte |
| Assigned Tasks | Manage all aspects of workstation management for Location 2 |

| | |
|---|---|
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 47 shows the model creation template that the data owners fill out to document the Production Resource Admins role for the first application.

**Table 47 Model Creation Template for Production Resource Admins Role for Application 1**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Production Application 1 Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 1 of 4 |
| Assigned Administrators | Nick, Wade |
| Assigned Tasks | Overall responsibility for all aspects of resource management for application 1 |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 48 shows the model creation template that the data owners fill out to document the Production Resource Admins role for the second application.

**Table 48 Model Creation Template for Production Resource Admins Role for Application 2**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Production Application 2 Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 2 of 4 |
| Assigned Administrators | Jennifer, Brad |
| Assigned Tasks | Overall responsibility for all aspects of resource management for application 2 |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 49 shows the model creation template that the data owners fill out to document the Production Resource Admins role for the third application.

**Table 49 Model Creation Template for Production Resource Admins Role for Application 3**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Production Application 3 Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 3 of 4 |
| Assigned Administrators | Scott, Laura |
| Assigned Tasks | Overall responsibility for all aspects of resource management for application 3 |
| Security Group | Implementation |

| | |
|---|---|
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 50 shows the model creation template that the data owners fill out to document the Production Resource Admins role for the common set of resources.

**Table 50 Model Creation Template for Production Resource Admins Role for Common Resources**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Production Common Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 4 of 4 |
| Assigned Administrators | Jim, Justin |
| Assigned Tasks | Overall responsibility for all aspects of resource management common to the Production business unit |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

**Creating the Delegation Model for the Bus Mgmt Business Unit**

This business unit has two main divisions — business management and sales. The business management unit is based in Chicago and includes the product planning, legal, marketing, and other groups. The marketing and legal business management teams and the sales team are spread over different physical locations across North America and Europe.

Approximately 550 employees work for the Bus Mgmt business unit. The sales division is the largest division, with 400 employees. Each sales representative has a portable computer. All users in this business unit have a portable computer and a desktop, making a total of 700 managed workstations. Additionally, about 250 servers provide various required services.

Table 51 shows the number of user, workstation, and server accounts that are stored in the Bus Mgmt business unit in the Chicago and London locations.

**Table 51 Users, Workstations, and Servers in the Bus Mgmt Business Unit**

| Locations | Users | Workstations | Servers |
|---|---|---|---|
| Chicago, London | 5,500 | 7,000 | 250 |

Table 52 shows the number of user, workstation, and server accounts that are stored in the Bus Mgmt business unit in the Chicago and London locations, separated by each division. Because business unit users are based across two different continents, business unit content is distributed across the two domains noam.concorp.contoso.com and europe.concorp.contoso.com.

**Table 52 Users, Workstations, and Servers in the Bus Mgmt Business Unit**

| Division | Locations | Users | Workstations | Servers |
|---|---|---|---|---|
| Planning | Chicago | 300 | 600 | 50 |
| Marketing | Chicago, London | 500 | 1,000 | 40 |
| Legal | Chicago, London | 300 | 600 | 30 |
| Sales | Chicago, New York, London, Paris, Rome | 4,000 | 4,000 | 100 |
| Other | Chicago | 400 | 800 | 30 |

Table 53 shows the numbers of server types that Bus Mgmt stores.

**Table 53 Distribution of Server Types in Bus Mgmt Business Unit**

| Locations | File Servers | Web Servers | Database Servers | Application Servers |
|---|---|---|---|---|

| | | | | |
|---|---|---|---|---|
| Chicago, London | 120 | 45 | 50 | 35 |

**Bus Mgmt Administrative and Group Policy Requirements**

Administrative requirements for the Bus Mgmt business unit include the management of users, workstations, and resources, as follows:

- **User Account Management.** All user accounts in North America are centrally managed by one administrative group based in Chicago. All user accounts in Europe are centrally managed by one administrative group based in London.

- **Workstation Management.** Each physical location has a separate team responsible for managing workstations.

- **Resource Management.** All business applications are hosted on servers in Chicago and managed by one administrative group based in Chicago. Each location has one local administrative group responsible for managing all locally hosted resources.

Group Policy requirements include the settings that must be applied for different types of accounts, for workstations (both desktop and portable computers), and for resources according to the resource type, as follows:

- **User Accounts.** All user accounts in North America need one Group Policy for user configuration settings. Similarly, all accounts in Europe need one Group Policy for user configuration settings. Additionally all users in each division need a common user configuration policy.

- **Workstations.** Requirements for scripts and other computer configuration settings necessitate that different Group Policy settings be applied for desktop and portable computers.

- **Resources.** Computer configuration settings necessitate that different Group Policy settings be applied for different kinds of resources and might require the application of specific Group Policy settings for the various specific resources.

**Bus Mgmt OU Structure Based on Administrative and Group Policy Requirements**

Figure 17 shows the OU structure for the Bus Mgmt OU that accommodates the administrative and Group Policy requirements for the noam.concorp.contoso.com domain.




Table 54 shows the rationale for the OU structure for noam.concorp.contoso.com shown in Figure 17.

**Table 54 Purpose of Each OU in the Bus Mgmt Business Unit OU Hierarchy in noam.concorp.contoso.com**

| Organizational Unit | Purpose |
|---|---|
| User Accounts | Main OU to store user accounts<br><br>Delegation point for Account Admins role |
| User Accounts\Planning | Used to apply Group Policy for all users in Planning |
| User Accounts\Marketing | Used to apply Group Policy for all users in Marketing |
| User Accounts\Legal | Used to apply Group Policy for all users in Legal |
| User Accounts\Sales | Used to apply Group Policy for all users in Sales |
| User Accounts\Other | Used to apply Group Policy for all other users in the business unit |
| Workstations | Main OU to store workstation computer accounts<br><br>Delegation point for Workstation Admins role |

| | |
|---|---|
| Workstations\Chicago | Used to delegate workstation management for all workstations in Chicago to the local administrative group |
| Workstations\New York | Used to delegate workstation management for all workstations in New York to the local administrative group |
| Resources | Main OU to store servers comprising resources |
| Resources\Business Applications | Used to delegate resource management of all servers involved in hosting all business applications to one administrative group |
| Resources\Business Applications\Application 1..n | Used to collectively store and manage all servers that are part of some business application |
| Resources\Chicago | Used to delegate resource management for all other resources in Chicago to the local administrative group |
| Resources\New York | Used to delegate workstation management for all other resources in New York to the local administrative group |
| Resources*X*\File, Web, Database, Application servers | OU Structure to store servers for common resources in location *X* |

Figure 18 shows the OU structure for the Bus Mgmt OU that accommodates the administrative and Group Policy requirements for the europe.concorp.contoso.com domain.




Table 55 shows the rationale for the OU structure for europe.concorp.contoso.com shown in Figure 18.

**Table 55 Purpose of Each OU in the Bus Mgmt Business Unit OU Hierarchy in europe.concorp.contoso.com**

| Organizational Unit | Purpose |
|---|---|
| Accounts | Main OU to store accounts<br><br>Delegation point for Account Admins role |
| Accounts\Legal | Used to apply Group Policy for all users in Legal |
| Accounts\Sales | Used to apply Group Policy for all users in Sales |
| Accounts\Other | Used to apply Group Policy for all other users in the business unit |
| Workstations | Main OU to store workstation computer accounts<br><br>Delegation point for the Workstation Admins role |
| Workstations\London | Used to delegate workstation management for all workstations in London to the local administrative group |
| Workstations\Paris | Used to delegate workstation management for all workstations in Paris to the local administrative group |
| Workstations\Rome | Used to delegate workstation management for all workstations in Rome to the local administrative group |
| Resources | Main OU to store servers comprising resources |
| Resources\Business Applications | Used to delegate resource management of all servers involved in hosting all business applications to one administrative group |
| Resources\Business Applications\Application 1..n | Used to collectively store and manage all servers that are part of some business application |
| Resources\London | Used to delegate workstation management for all other resources in London to the local administrative group |
| Resources\Paris | Used to delegate workstation management for all other resources in Paris to the local administrative group |
| Resources\Rome | Used to delegate workstation management for all other resources in Rome to the local administrative group |
| Resources*X*\ File, Web, Database, Application servers | OU Structure to store servers for common resources in location *X* in an organized fashion |

**Bus Mgmt Role Instances Based on Business Unit Requirements**

On the basis of the Bus Mgmt business unit requirements, the following role instances must be created:

- **User Account Management.** Because one group is responsible for account management for North America and one for Europe, two instances of the Account Admins role will be needed.

| **Note** |
|---|
| One group is instantiated in the noam.concorp.contoso.com domain and one in the europe.concorp.contoso.com domain. |

- **Workstation Management.** Because a different administrative group is required for each of the five physical locations, five instances of the Workstation Admins role are required.

| **Note** |
|---|
| Certain groups are shared across both noam.concorp.contoso.com domain and the europe.concorp.contoso.com domain. |

- **Resource Management.** Based on the business unit requirements, one Resource Admins role instance is required to manage all servers that belong to all the business applications hosted in Chicago and one instance of the Resource Admins role is required for every physical location of this business unit, for a total of six role instances.

Table 56 shows the model creation template that the data owners fill out to document the Bus Mgmt NOAM Account Admins role.

**Table 56 Model Creation Template for Bus Mgmt NOAM Account Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt NOAM Account Admins |
| Instance of | Account Admins |
| Instance Number | 1 of 2 |
| Assigned Administrators | Danielle, Jason |
| Assigned Tasks | Manage all aspects of account management for all accounts in North America |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 57 shows the model creation template that the data owners fill out to document the Bus Mgmt Europe Account Admins role.

**Table 57 Model Creation Template for Bus Mgmt Europe Account Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt Europe Account Admins |
| Instance of | Account Admins |
| Instance Number | 2 of 2 |
| Assigned Administrators | Robert, Michelle |
| Assigned Tasks | Manage all aspects of account management for all accounts in Europe |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 58 shows the model creation template that the data owners fill out to document the Bus Mgmt Chicago Workstation Admins role.

**Table 58 Model creation template for Bus Mgmt Chicago Workstation Admins role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt Chicago Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 1 of 5 |
| Assigned Administrators | Janet, Harold |
| Assigned Tasks | Manage all aspects of workstation management for all workstations in Chicago |
| Security Group | Implementation |
| Permissions Assigned | Implementation |

| | |
|---|---|
| Notes | Creation and implementation |

Table 59 shows the model creation template that the data owners fill out to document the Bus Mgmt London Workstation Admins role.

**Table 59 Model creation template for Bus Mgmt London Workstation Admins role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt London Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 2 of 5 |
| Assigned Administrators | Stuart, Ken |
| Assigned Tasks | Manage all aspects of workstation management for all workstations in London |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 60 shows the model creation template that the data owners fill out to document the Bus Mgmt New York Workstation Admins role.

**Table 60 Model Creation Template for Bus Mgmt New York Workstation Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt New York Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 3 of 5 |
| Assigned Administrators | Linda, Steve |
| Assigned Tasks | Manage all aspects of workstation management for all workstations in New York |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 61 shows the model creation template that the data owners fill out to document the Bus Mgmt Paris Workstation Admins role.

**Table 61 Model Creation Template for Bus Mgmt Paris Workstation Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt Paris Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 4 of 5 |
| Assigned Administrators | Marc, Sara |
| Assigned Tasks | Manage all aspects of workstation management for all workstations in Paris |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 62 shows the model creation template that the data owners fill out to document the Bus Mgmt Rome Workstation

Admins role.

**Table 62 Model Creation Template for Bus Mgmt Rome Workstation Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt Rome Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 5 of 5 |
| Assigned Administrators | Victor, Rosa |
| Assigned Tasks | Manage all aspects of workstation management for all workstations in Rome |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 63 shows the model creation template that the data owners fill out to document the Bus Mgmt Application Admins role.

**Table 63 Model Creation Template for Bus Mgmt Application Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt Application Admins |
| Instance of | Resource Admins |
| Instance Number | 1 of 6 |
| Assigned Administrators | Fred, Glenn |
| Assigned Tasks | Manage all aspects of all business management business unit applications |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 64 shows the model creation template that the data owners fill out to document the Bus Mgmt Chicago Resource Admins role.

**Table 64 Model Creation Template for Bus Mgmt Chicago Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt Chicago Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 2 of 6 |
| Assigned Administrators | Kristen, Terry |
| Assigned Tasks | Manage all aspects of other business management business unit resources in Chicago |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 65 shows the model creation template that the data owners fill out to document the Bus Mgmt London Resource Admins role.

**Table 65 Model Creation Template for Bus Mgmt London Resource Admins Role**

| Field | Assignment Information |
|---|---|

| | |
|---|---|
| Role Instance Name | Business Mgmt London Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 3 of 6 |
| Assigned Administrators | Ron, Allison |
| Assigned Tasks | Manage all aspects of other business management business unit resources in London |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 66 shows the model creation template that the data owners fill out to document the Bus Mgmt New York Resource Admins role.

**Table 66 Model Creation Template for Bus Mgmt New York Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt New York Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 4 of 6 |
| Assigned Administrators | Chris, Julian |
| Assigned Tasks | Manage all aspects of other business management business unit resources in New York |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 67 shows the model creation template that the data owners fill out to document the Bus Mgmt Paris Resource Admins role.

**Table 67 Model Creation Template for Bus Mgmt Paris Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt Paris Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 5 of 6 |
| Assigned Administrators | Albert, Emile |
| Assigned Tasks | Manage all aspects of other business management business unit resources in Paris |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 68 shows the model creation template that the data owners fill out to document the Bus Mgmt Rome Resource Admins role.

**Table 68 Model Creation Template for Bus Mgmt Rome Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | Business Mgmt Rome Resource Admins |

| | |
|---|---|
| Instance of | Resource Admins |
| Instance Number | 6 of 6 |
| Assigned Administrators | David, Thomas |
| Assigned Tasks | Manage all aspects of other business management business unit resources in Rome |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

**Creating the Delegation Model for the IT Business Unit**

The IT business unit includes all the administrative personnel responsible for managing all IT resources and providing IT support for Contoso Pharmaceuticals.

Approximately 250 administrative personnel and 500 workstations are distributed across all physical locations of this business unit. Approximately 100 servers provide various services to users. For example, an internal Web-based application allows end-users to report issues and request assistance. One set of servers provides Web-based documentation. Another set of servers provides an internal tracking and monitoring system.

Table 69 shows the number of user accounts, workstations, and servers that are stored in the IT business unit.

**Table 69 Users, Workstations, and Servers in the IT Business Unit**

| Locations | Users | Workstations | Servers |
|---|---|---|---|
| Chicago, Atlanta, New York, London, Paris, Rome | 2,500 | 5,000 | 100 |

Table 70 shows the numbers of server types that the IT business unit stores.

**Table 70 Distribution of Server Types in the IT Business Unit**

| Business Unit | File Servers | Web Servers | Database Servers | Application Servers |
|---|---|---|---|---|
| IT | 30 | 30 | 20 | 20 |

**IT Administrative and Group Policy Requirements**

Administrative requirements for the IT business unit include the management of accounts, workstations, and resources, as follows:

- **User account management.** All user accounts in the IT business unit in North America are centrally managed by one administrative group based in Chicago. All user accounts in Europe are centrally managed by one administrative group based in London.

- **Workstation Management.** All workstations in the IT business unit are locally managed by separate administrative groups.

- **Resource Management.** All IT applications are hosted on servers in Chicago and managed by one administrative group based in Chicago. Each location has a one local administrative group that is responsible for managing all locally hosted resources.

Group Policy requirements include the settings that must be applied for different types of user accounts, workstations (both desktop and portable computers), and resources according to resource type, as follows:

- **User Accounts.** All user accounts in North America require one Group Policy for user configuration settings. Similarly, all accounts in Europe need one Group Policy for user configuration settings.

- **Workstations.** All workstations in North America require one Group Policy for computer configuration settings and all workstations in Europe need one Group Policy for computer configuration settings.

- **Resources.** Computer configuration settings necessitate that different Group Policy settings be applied for different kinds of resources and might require the application of specific Group Policy settings for the various specific resources.

**IT OU Structure Based on Administrative and Group Policy Requirements**

Figure 19 shows the OU structure for the IT OU that accommodates the administrative and Group Policy requirements for the noam.concorp.contoso.com domain.




Figure 20 shows the OU structure for the IT OU that accommodates the administrative and Group Policy requirements for the europe.concorp.contoso.com domain.




**IT Role Instances Based on Business Unit Requirements**

On the basis of the IT business unit requirements, the following role instances must be created:

- **User Account Management.** Because one group is responsible for account management for North America and one for Europe, two instances of the Account Admins role are needed.

| **Note** |
|---|
| One group is instantiated in the noam.concorp.contoso.com domain and one in the europe.concorp.contoso.com domain. |

- **Workstation Management.** Because a different administrative group is required for each of the five physical locations, five instances of the Workstation Admins role are required.

| **Note** |
|---|
| Certain groups are shared across both noam.concorp.contoso.com domain and the europe.concorp.contoso.com domain. |

- **Resource Management.** Based on the business unit requirements, one instance of the Resource Admins role is required to manage all the servers that belong to all the IT applications in Chicago and one instance of the Resource Admins role is required for every physical location where this business unit is located.

Table 71 shows the model creation template that the data owners fill out to document the IT NOAM Account Admins role.

**Table 71 Model Creation Template for IT NOAM Account Admins Role**

| Field | Assignment Information |
|---|---|

| | |
|---|---|
| Role Instance Name | IT NOAM Account Admins |
| Instance of | Account Admins |
| Instance Number | 1 of 2 |
| Assigned Administrators | Roderick, Francisco |
| Assigned Tasks | Manage all aspects of acct mgmt for all IT accounts in North America |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 72 shows the model creation template that the data owners fill out to document the IT Europe Account Admins role.

**Table 72 Model Creation Template for IT Europe Account Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT Europe Account Admins |
| Instance of | Account Admins |
| Instance Number | 2 of 2 |
| Assigned Administrators | Eric, Valerie |
| Assigned Tasks | Manage all aspects of acct mgmt for all IT accounts in Europe |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 73 shows the model creation template that the data owners fill out to document the IT Chicago Workstation Admins role.

**Table 73 Model Creation Template for IT Chicago Workstation Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT Chicago Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 1 of 5 |
| Assigned Administrators | Larry, John |
| Assigned Tasks | Manage all aspects of workstation mgmt for all IT business unit workstations in Chicago |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 74 shows the model creation template that the data owners fill out to document the IT London Workstation Admins role.

**Table 74 Model Creation Template for IT London Workstation Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT London Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 2 of 5 |

| | |
|---|---|
| Assigned Administrators | Roland, Jonathan |
| Assigned Tasks | Manage all aspects of workstation mgmt for all IT business unit workstations in London |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 75 shows the model creation template that the data owners fill out to document the IT New York Workstation Admins role.

**Table 75 Model Creation Template for IT New York Workstation Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT New York Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 3 of 5 |
| Assigned Administrators | Neil, Paulette |
| Assigned Tasks | Manage all aspects of workstation mgmt for all IT business unit workstations in New York |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 76 shows the model creation template that the data owners fill out to document the IT Paris Workstation Admins role.

**Table 76 Model Creation Template for IT Paris Workstation Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT Paris Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 4 of 5 |
| Assigned Administrators | Elliott, Marie |
| Assigned Tasks | Manage all aspects of workstation mgmt for all IT business unit workstations in Paris |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 77 shows the model creation template that the data owners fill out to document the IT Rome Workstation Admins role.

**Table 77 Model Creation Template for IT Rome Workstation Admins Role**

| Field | Assignment Information |
|---|---|
| Field | Assignment Information |
| Role Instance Name: | IT Rome Workstation Admins |
| Instance of | Workstation Admins |
| Instance Number | 5 of 5 |
| Assigned Administrators | Paul, Eleonora |

| | |
|---|---|
| Assigned Tasks | Manage all aspects of workstation mgmt for all IT business unit workstations in Rome |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 78 shows the model creation template that the data owners fill out to document the IT Application Admins role.

**Table 78 Model Creation Template for IT Application Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name: | IT Application Admins |
| Instance of | Resource Admins |
| Instance Number | 1 of 6 |
| Assigned Administrators | Karen, Luke |
| Assigned Tasks | Manage all aspects of all IT business unit applications |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 79 shows the model creation template that the data owners fill out to document the IT Chicago Resource Admins role.

**Table 79 Model Creation Template for IT Chicago Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT Chicago Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 2 of 6 |
| Assigned Administrators | Perry, Steve |
| Assigned Tasks | Manage all aspects of other IT business unit resources in Chicago |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 80 shows the model creation template that the data owners fill out to document the IT London Resource Admins role.

**Table 80 Model Creation Template for IT London Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT London Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 3 of 6 |
| Assigned Administrators | Jerry, Carol |
| Assigned Tasks | Manage all aspects of other IT business unit resources in London |
| Security Group | Implementation |

| | |
|---|---|
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 81 shows the model creation template that the data owners fill out to document the IT New York Resource Admins role.

**Table 81 Model Creation Template for IT New York Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT New York Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 4 of 6 |
| Assigned Administrators | Jack, Erica |
| Assigned Tasks | Manage all aspects of other IT business unit resources in New York |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 82 shows the model creation template that the data owners fill out to document the IT Paris Resource Admins role.

**Table 82 Model Creation Template for IT Paris Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT Paris Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 5 of 6 |
| Assigned Administrators | Guy, Lise |
| Assigned Tasks | Manage all aspects of other IT business unit resources in Paris |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

Table 83 shows the model creation template that the data owners fill out to document the IT Rome Resource Admins role.

**Table 83 Model Creation Template for IT Rome Resource Admins Role**

| Field | Assignment Information |
|---|---|
| Role Instance Name | IT Rome Resource Admins |
| Instance of | Resource Admins |
| Instance Number | 6 of 6 |
| Assigned Administrators | Enrico, Antonella |
| Assigned Tasks | Manage all aspects of other IT business unit resources in Rome |
| Security Group | Implementation |
| Permissions Assigned | Implementation |
| Notes | Creation and implementation |

# Step 5 – Implement the Contoso Data Management Administrative Delegation Model

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

At this point, one OU exists for every business unit and the Business Unit Admins role for each Business Unit has been delegated Full Control over its respective business unit OU.

Each business unit administrator implements the model for the respective OU according to the following criteria:

- **Stakeholders:** Business Unit Admins from all Business Units
- **Objective:** Implement each role instance according to the delegation model design document
- **Assumption:** The three domains are installed and functioning
- **Approach:** Implement the roles defined in the data management delegation model by performing the following steps:
  - Create the OU structure for the business unit
  - Implement role instances
  - Assign users to roles

**Implementing the Data Delegation Model for the RandD Business Unit**
The RandD business unit administrator is responsible for creating the business unit OU structure and implementing the data management delegation model for this business unit.

**Creating the OU Structure for the RandD Business Unit**
To create the RandD OU structure, the RandD business unit administrator creates the following OU objects:

- OU=User Accounts,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Workstations,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Resources,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Research,OU=User Accounts,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Development,OU=User Accounts,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Desktops,OU=Workstations,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Laptops,OU=Workstations,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=File Servers,OU=Resources,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Web Servers,OU=Resources,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Database Servers,OU=Resources,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Application Servers,OU=Resources,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

**Implementing RandD Role Instances**

To implement the RandD business administrative roles, the RandD business unit administrator creates security groups in the Delegation OU to represent the role instances for the business unit.

In OU=Delegation,OU=RandD,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com, the following groups are created:

- RandD Account Admins
- RandD Workstation Admins
- RandD Resource Admins

The RandD the business unit administrator grants permissions to security groups as follows:

- RandD Account Admins: Full Control on the Accounts OU
- RandD Workstation Admins: Full Control on the Workstations OU
- RandD Resource Admins: Full Control on the Resources OU

Additionally, the Restricted Groups feature of Group Policy is used to add the security groups representing the Workstation Admins and Resource Admins roles to the local Administrators groups on member servers.

**Assigning RandD Administrative Users to Roles**
The RandD business unit administrator adds the accounts of the administrative personnel listed for each role in the delegation model templates to the respective security groups that represent the role instances.

**Implementing the Data Delegation Model for the Production Business Unit**
The Production business unit administrator is responsible for creating the business unit OU structure and implementing the data management delegation model for this business unit.

**Creating the OU Structure for the Production Business Unit:**
To create the Production OU structure, the Production business unit administrator creates the following OU objects:

- OU=User Accounts,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Workstations,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Desktops,OU=Workstations,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Laptops,OU=Workstations,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Production Application 1,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Production Application 2,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Production Application 3,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=Shared Resources,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com
- OU=File Servers,OU=Production Application 1,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Web Servers,OU=Production Application 1,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Database Servers,OU=Production Application 1,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Application Servers,OU=Production Application 1,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=File Servers,OU=Production Application 2,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Web Servers,OU=Production Application 2,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Database Servers,OU=Production Application 2,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Application Servers,OU=Production Application 2,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=File Servers,OU=Production Application 3,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Web Servers,OU=Production Application 3,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Database Servers,OU=Production Application 3,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Application Servers,OU=Production Application 3,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=File Servers,OU=Shared Resources,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Web Servers,OU= Shared Resources,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Database Servers,OU= Shared Resources,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

- OU=Application Servers,OU= Shared Resources,OU=Resources,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com

**Implementing Production Role Instances**

To implement the Production business unit administrative roles, the Production business unit administrator creates security groups in the Delegation OU to represent role instances for the business unit.

In OU=Delegation,OU=Production,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com, the following groups are created:

- Production Account Admins

- Production Location 1 Workstation Admins

- Production Location 2 Workstation Admins

- Production Application 1 Resource Admins
- Production Application 2 Resource Admins
- Production Application 3 Resource Admins
- Production Common Resources Resource Admins

The Production business unit administrator grants permissions to security groups as follows:

- Production Account Admin: Full Control on the Accounts OU
- Production Location 1 Workstation Admins: Full Control on the Location 1 OU within the Workstations OU
- Production Location 2 Workstation Admins: Full Control on the Location 2 OU within the Workstations OU
- Production Application 1 Resource Admins: Full Control on the Application 1 OU within the Resources OU
- Production Application 2 Resource Admins: Full Control on the Application 2 OU within the Resources OU
- Production Application 3 Resource Admins: Full Control on the Application 3 OU within the Resources OU
- Production Common Resources Resource Admins: Full Control on the Common Resources OU within the Resources OU

Additionally, the Restricted Groups feature of Group Policy is used to add the security groups representing Workstation Admins and Resource Admins roles to the local Administrators groups on member servers.

**Assigning Production Administrative Users to Roles**
The Production business unit administrator adds the accounts of the administrative personnel listed for each role in the model documentation (templates) to the respective security groups that represent the role instances.

**Implementing the Data Delegation Model for the Bus Mgmt Business Unit**
The Bus Mgmt business unit administrator is responsible for creating the business unit OU structure and implementing the data management delegation model for this business unit.

**Creating the OU structure for the Bus Mgmt Business Unit**
To create the Bus Mgmt OU structure, the Bus Mgmt business unit administrator creates the OU objects shown in Figures 17 and 18 earlier in this document.

**Implementing Bus Mgmt Role Instances**
To implement the Bus Mgmt business unit administrative roles, the Bus Mgmt business unit administrator creates security groups in the Delegation OU of the respective domain to represent the required role instances.

In OU=Delegation,OU=Bus Mgmt,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com, the following groups are created:

- BusMgmt NOAM Account Admins
- BusMgmt Chicago Workstation Admins
- BusMgmt New York Workstation Admins
- BusMgmt Application Admins
- BusMgmt Chicago Resource Admins
- BusMgmt New York Resource Admins

In OU=Delegation,OU=Bus Mgmt,OU=Business Units,DC=europe,DC=concorp,DC=contoso,DC=com, the following groups are created:

- BusMgmt Europe Account Admins
- BusMgmt London Workstation Admins
- BusMgmt Paris Workstation Admins
- BusMgmt Rome Workstation Admins
- BusMgmt London Resource Admins
- BusMgmt Paris Resource Admins
- BusMgmt Rome Resource Admins

The Bus Mgmt business unit administrator grants permissions to the corresponding security groups as shown in Table 84:

**Table 84 Permissions for Security Groups Representing Bus Mgt Business Unit Role Instances**

| Security Group | Permissions | OU | Domain |
|---|---|---|---|
| BusMgmt NOAM Account Admins | Full Control | Accounts OU | noam.concorp.contoso.com |
| BusMgmt Europe Account Admins | Full Control | Accounts OU | europe.concorp.contoso.com |
| BusMgmt Chicago Workstation Admins | Full Control | Chicago OU within the Workstations OU | noam.concorp.contoso.com |
| BusMgmt London Workstation Admins | Full Control | London OU within the Workstations OU | noam.concorp.contoso.com |
| BusMgmt New York Workstation Admins | Full Control | New York OU within the Workstations OU | noam.concorp.contoso.com |
| BusMgmt Paris Workstation Admins | Full Control | Paris OU within the Workstations OU | europe.concorp.contoso.com |
| BusMgmt Rome Workstation Admins | Full Control | Rome OU within the Workstations OU | europe.concorp.contoso.com |
| BusMgmt Application Admins | Full Control | Applications OU within the Resources OU | noam.concorp.contoso.com |
| BusMgmt Chicago Resource Admins | Full Control | Chicago OU within the Resources OU | noam.concorp.contoso.com |
| BusMgmt London Resource Admins | Full Control | London OU within the Resources OU | europe.concorp.contoso.com |
| BusMgmt New York Resource Admins | Full Control | New York OU within the Resources | noam.concorp.contoso.com |
| BusMgmt Paris Resource Admins | Full Control | Paris OU within the Resources OU | europe.concorp.contoso.com |
| BusMgmt Rome Resource Admins | Full Control | Rome OU within the Resources OU | europe.concorp.contoso.com |

Additionally, the Restricted Groups feature of Group Policy is used to add the security groups representing Workstation Admins and Resource Admins roles to the local Administrators groups on member servers.

**Assigning Bus Mgmt Administrative Users to Roles**

The Bus Mgmt business unit administrator adds the accounts of the administrative personnel listed for each role in the delegation model templates to the respective security groups that represent the role instances.

**Implementing the Data Delegation Model for the IT Business Unit**

The IT business unit administrator is responsible for creating the business unit OU structure and implementing the data management delegation model for this business unit.

**Creating the OU structure for the IT Business Unit**

To create the Bus Mgmt OU structure, the IT business unit administrator creates the OU objects shown in Figures 17 and 18 earlier in this document.

**Implementing IT Role Instances**

To implement the IT business unit administrative roles, the IT business unit administrator creates security groups in the Delegation OU of each domain to represent the required role instances:

- In OU=Delegation,OU=IT,OU=Business Units,DC=noam,DC=concorp,DC=contoso,DC=com, the following groups are created:
    - IT NOAM Account Admins
    - IT Chicago Workstation Admins
    - IT New York Workstation Admins
    - IT Application Admins
    - IT Chicago Resource Admins
    - IT New York Resource Admins
- In OU=Delegation,OU=IT,OU=Business Units,DC=europe,DC=concorp,DC=contoso,DC=com, the following groups are created:
    - IT Europe Account Admins
    - IT London Workstation Admins
    - IT Paris Workstation Admins
    - IT Rome Workstation Admins
    - IT London Resource Admins
    - IT Paris Resource Admins
    - IT Rome Resource Admins

The IT business unit administrator grants permissions to the corresponding security groups as shown in Table 85:

**Table 85 Permissions for Security Groups Representing IT Business Unit Role Instances**

| Security Group | Permissions | OU | Domain |
|---|---|---|---|
| IT NOAM Account Admins | Full Control | Accounts OU | noam.concorp.contoso.com |
| IT Europe Account Admins | Full Control | Accounts OU | europe.concorp.contoso.com |
| IT Chicago Workstation Admins | Full Control | Chicago OU within the Workstations OU | noam.concorp.contoso.com |
| IT London Workstation Admins | Full Control | London OU within the Workstations OU | europe.concorp.contoso.com |
| IT New York Workstation Admins | Full Control | New York OU within the Workstations OU | noam.concorp.contoso.com |
| IT Paris Workstation Admins | Full Control | Paris OU within the Workstations OU | europe.concorp.contoso.com |
| IT Rome Workstation Admins | Full Control | Rome OU within the Workstations OU | europe.concorp.contoso.com |
| IT Application Admins | Full Control | Applications OU within the Resources OU | noam.concorp.contoso.com |
| IT Chicago Resource Admins | Full Control | Chicago OU within the Resources OU | noam.concorp.contoso.com |
| IT London Resource Admins | Full Control | London OU within the Resources OU | europe.concorp.contoso.com |

| | | | |
|---|---|---|---|
| IT New York Resource Admins | Full Control | New York OU within the Resources OU | noam.concorp.contoso.com |
| IT Paris Resource Admins | Full Control | Paris OU within the Resources OU | europe.concorp.contoso.com |
| IT Rome Resource Admins | Full Control | Rome OU within the Resources OU | europe.concorp.contoso.com |

Additionally, the Restricted Groups feature of Group Policy is used to add the security groups representing the Workstation Admins and Resource Admins roles to the local Administrators groups on member servers.

**Assigning IT Administrative Users to Roles**

The IT BU Admin adds the accounts of the administrative personnel listed for each role in the delegation model templates to the respective security groups that represent the role instances.

This completes Contoso's implementation of the delegation model. From this point, Contoso will follow the recommendations for maintaining the delegation model.

# Best Practices for Delegating Active Directory Administration: Appendices

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section contains the following appendices:


- Appendix A: Active Directory Administrative Tasks
- Appendix B: Default Active Directory Security Groups
- Appendix C: Active Directory Standard Permissions
- Appendix D: Active Directory Extended Rights
- Appendix E: Active Directory Property Sets
- Appendix F: Logon Rights and Privileges
- Appendix G: Active Directory Delegation Tools
- Appendix H: Active Directory Display Name Mappings
- Appendix I: Default Container Hierarchy for Active Directory Partitions
- Appendix J: Default Owners of Active Directory Objects
- Appendix K: Default Settings in the Master Security Descriptor of the AdminSDHolder Object
- Appendix L: Implementing Service Management Delegation Roles
- Appendix M: Service Management Delegation Role Definitions
- Appendix N: Default Active Directory Service Administrator Groups
- Appendix O: Active Directory Delegation Wizard File

# Appendix A: Active Directory Administrative Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix is a comprehensive list of Active Directory® directory service administrative tasks for service administration and data administration, along with the permissions required to perform each task.

The permissions in this appendix are referred to by Security Descriptor Definition Language (SDDL) mnemonics. The following table maps the SDDL to the corresponding permissions and to the name by which these permissions are displayed in the ACL Editor.

**SDDL Mappings for Active Directory Permissions**

| SDDL | Permission | Name in ACL Editor |
|---|---|---|
| RC | Read Control | No specific mapping |
| SD | Standard Delete | Delete |
| WD | Write DACL | Modify Permissions |
| WO | Write Owner | Modify Owner |
| RP | Read Property | Read All Properties<br>Read <specific property> |
| WP | Write Property | Write All Properties<br>Write <specific property> |
| CC | Create Child | Create <class of object> |
| DC | Delete Child | Delete <class of object> |
| LC | List Child | List Contents |
| SW | Validated Write | Self-Membership<br>Validated-DNS-Host-Name<br>Validated-SPN |
| LO | List Object | List Object |
| DT | Delete Tree | Delete Subtree |
| CR | Extended Right | Referred to by specific Extended Right name – see Appendix D: Active Directory Extended Rights for more details |

For more information about SDDL, see the Microsoft Platform Software Development Kit (SDK) link on the on at http://msdn2.microsoft.com/EN-US/library/aa379567.aspx.

All classes and attributes in this table are referred to by their Common-Name.

**Active Directory Attributes and Classes**

Active Directory attributes and classes can be referred to by any of three types of names:

- **CN (Common-Name) —** Every object in the DS has a naming attribute from which its Relative Distinguished Name

(also known as RDN) is formed. (The Naming Attribute for most Class-Schema objects is Common-Name.)

- **LDAP-Display-Name —** The name of a class/attribute as known to the LDAP agent for the NTDS.

- **Display-Name —** A relatively more descriptive version of the Common-Name of a given object that the administration tools use to refer to the class/object.

Some attributes have Display Names defined while others do not. The base set of administration tools (ACL Editor) refers to classes/attributes by their Display-Name, if one exists for the class/attribute. If a class/attribute does not have a Display-Name defined, then the UI uses the LDAP-display name.

- To determine the equivalent display names used in the ACL Editor, see Appendix H: Active Directory Display Name Mappings.

- To learn more about how to use the ACL Editor and other Delegation Tools, see Appendix G: Active Directory Delegation Tools.

# Active Directory Service Administration Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section contains the following Active Directory Service Administration Tasks:

- Installation Management Tasks
- Operations Master Roles Management
- LDAP Policy Management Tasks
- Schema Management Tasks
- Trust Management Tasks
- Replication Management Tasks
- Backup and Restore Management Tasks
- Directory Database File Management Tasks
- Directory Service Configuration Management Tasks
- Security Policy Management Tasks
- DNS Management Tasks

# Installation Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following table lists the installation management tasks.

| Task | Permissions Required to Perform Task |
|---|---|
| Create the first domain in a new tree in a new/existing forest | User must be member of Administrators group on member server being promoted |
| Create a child domain in an existing domain tree | User must be member of Administrators group on member server being promoted.<br>The crossRef object under CN=Partitions, CN=Configuration, DC=<forestRootDomain> must be pre-created<br>Inheritable RP on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>Inheritable CC on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>CC on OU=Domain Controllers,DC=<domain> to create Computer objects<br>Full Control on the Computer object for the server that is being promoted<br>Full Control to “Creator Owner” on CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>Extended Right DS-Replication-Get-Changes on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes-All on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes-All on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Manage-Topology on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Manage-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Monitor-Topology on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Monitor-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Synchronize on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Synchronize on CN=Schema, CN=Configuration, DC=<forestRootDomain> |

| | |
|---|---|
| Create a replica (additional Domain Controller) | User must be member of Administrators group on member server being promoted<br>User Right “Enable computer and user accounts to be trusted for delegation”<br>Inheritable RP on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>Inheritable CC on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>CC on OU=Domain Controllers,DC=<domain> to create Computer objects<br>Full Control on the Computer object for the server that is being promoted<br>Full Control to “Creator Owner” on CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>Extended Right DS-Install=Replica on DC=<domain><br>Extended Right DS-Replication-Get-Changes on DC=<domain><br>Extended Right DS-Replication-Get-Changes on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes-All on DC=<domain><br>Extended Right DS-Replication-Get-Changes-All on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes-All on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Manage-Topology on DC=<domain><br>Extended Right DS-Replication-Manage-Topology on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Manage-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Monitor-Topology on DC=<domain><br>Extended Right DS-Replication-Monitor-Topology on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Monitor-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Synchronize on DC=<domain><br>Extended Right DS-Replication-Synchronize on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Synchronize on CN=Schema, CN=Configuration, DC=<forestRootDomain> |

| | |
|---|---|
| Remove a replica | User must be member of Administrators group on member server being promoted<br><br>User must have User Right “Allow Log on Locally”<br><br>Full Control on the NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> where <Server> is the DC being demoted<br><br>Full Control on the Computer object for the server that is being promoted<br><br>Extended Right DS-Replication-Get-Changes on DC=<domain><br><br>Extended Right DS-Replication-Get-Changes on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Get-Changes on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Get-Changes-All on DC=<domain><br><br>Extended Right DS-Replication-Get-Changes-All on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Get-Changes-All on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Manage-Topology on DC=<domain><br><br>Extended Right DS-Replication-Manage-Topology on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Manage-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Monitor-Topology on DC=<domain><br><br>Extended Right DS-Replication-Monitor-Topology on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Monitor-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Synchronize on DC=<domain><br><br>Extended Right DS-Replication-Synchronize on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Synchronize on CN=Schema, CN=Configuration, DC=<forestRootDomain> |

| | |
|---|---|
| Demote the last Domain Controller in a child domain | User must be member of Administrators group on member server being promoted<br><br>User must have User Right “Allow Log on Locally”<br><br>Full Control on CN=<crossRef>,CN=Partitions, CN=Configuration, DC=<forestRootDomain> where <crossRef> is the crossRef for this domain<br><br>Full Control on the NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> where <Server> is the DC being demoted<br><br>Full Control on the Computer object for the server that is being promoted<br><br>Extended Right DS-Replication-Get-Changes on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Get-Changes on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Get-Changes-All on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Get-Changes-All on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Manage-Topology on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Manage-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Monitor-Topology on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Monitor-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Synchronize on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Synchronize on CN=Schema, CN=Configuration, DC=<forestRootDomain> |

| | |
|---|---|
| Demote the last Domain Controller in a tree-root domain | User must be member of Administrators group on member server being promoted<br><br>User must have User Right “Allow Log on Locally”<br><br>Full Control on CN=<crossRef>,CN=Partitions, CN=Configuration, DC=<forestRootDomain> where <crossRef> is the crossRef for this domain<br><br>Full Control on the NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> where <Server> is the DC being demoted<br><br>Full Control on the Computer object for the server that is being promoted<br><br>Extended Right DS-Replication-Get-Changes on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Get-Changes on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Get-Changes-All on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Get-Changes-All on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Manage-Topology on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Manage-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Monitor-Topology on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Monitor-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Synchronize on CN=Configuration, DC=<forestRootDomain><br><br>Extended Right DS-Replication-Synchronize on CN=Schema, CN=Configuration, DC=<forestRootDomain> |
| Demote the last Domain Controller in a forest | User must be member of Administrators group on member server being Promoted |
| Designate a Domain Controller as a Global Catalog | WP on the corresponding NTDS Settings object with distinguished name cn=NTDS Settings, cn=<Computer-Name>, cn=Servers, cn=<SiteName>,cn=Sites,cn=Configuration, dc=<forestRootDomain> to modify the options attribute<br><br>NOTE: These permissions are sufficient to perform the task. However they are insufficient when using the Active Directory UI tools to perform the task. The repadmin tool can be used with these permissions to perform the task. |
| Undesignate a Domain Controller as a Global Catalog | WP on the corresponding NTDS Settings object with distinguished name cn=NTDS Settings, cn=<Computer-Name>, cn=Servers, cn=<SiteName>,cn=Sites,cn=Configuration, dc=<forestRootDomain> to modify the options attribute<br><br>NOTE: These permissions are sufficient to perform the task. However they are insufficient when using the Active Directory UI tools to perform the task. The repadmin tool can be used with these permissions to perform the task. |
| Raise Forest Functionality Level | WP on the object cn=Partitions, cn=Configuration, dc=<forestRootDomain> to modify ms-DS-Behavior-Version attribute |
| Raise Domain Functionality Level | WP on the object dc=<domain> to modify ms-DS-Behavior-Version attribute |
| Migrate SID-History | The extended right Migrate-SID-History is required on dc=<Domain> (root of domain directory partition) |

| | |
|---|---|
| Create the first domain in a new tree in a new/existing forest | User must be member of Administrators group on member server being Promoted |
| Create a child domain in an existing domain tree | User must be member of Administrators group on member server being Promoted<br>The crossRef object under CN=Partitions, CN=Configuration, DC=<forestRootDomain> must be pre-created<br>Inheritable RP on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>Inheritable CC on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>CC on OU=Domain Controllers,DC=<domain> to create Computer objects<br>Full Control on the Computer object for the server that is being Promoted<br>Full Control to "Creator Owner" on CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>Extended Right DS-Replication-Get-Changes on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes-All on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes-All on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Manage-Topology on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Manage-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Monitor-Topology on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Monitor-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Synchronize on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Synchronize on CN=Schema, CN=Configuration, DC=<forestRootDomain> |

| | |
|---|---|
| Create a replica (additional Domain Controller) | User must be member of Administrators group on member server being Promoted<br>User Right “Enable computer and user accounts to be trusted for delegation”<br>Inheritable RP on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>Inheritable CC on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>CC on OU=Domain Controllers,DC=<domain> to create Computer objects<br>Full Control on the Computer object for the server that is being Promoted<br>Full Control to “Creator Owner” on CN=<Site>, CN=Sites, CN=Configuration, DC=<foestRootDomain><br>Extended Right DS-Install=Replica on DC=<domain><br>Extended Right DS-Replication-Get-Changes on DC=<domain><br>Extended Right DS-Replication-Get-Changes on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes-All on DC=<domain><br>Extended Right DS-Replication-Get-Changes-All on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Get-Changes-All on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Manage-Topology on DC=<domain><br>Extended Right DS-Replication-Manage-Topology on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Manage-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Monitor-Topology on DC=<domain><br>Extended Right DS-Replication-Monitor-Topology on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Monitor-Topology on CN=Schema, CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Synchronize on DC=<domain><br>Extended Right DS-Replication-Synchronize on CN=Configuration, DC=<forestRootDomain><br>Extended Right DS-Replication-Synchronize on CN=Schema, CN=Configuration, DC=<forestRootDomain> |

# Schema Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Enable Schema modification on a DC in the enterprise | This permission is not needed in Windows Server 2003.<br><br>In the Microsoft® Windows® 2000 operating system, WP is needed on HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\NTDS\Parameters to modify the Schema Update Allowed entry |
| Change the current Schema Master | Extended Right Change-Schema-Master on cn=Schema, cn=Configuration, dc=<ForestRootDomain> |
| Add a Class definition in the Schema | CC on cn=Schema, cn=Configuration, dc=<ForestRootDomain> (to create objects of class Class-Schema) |
| Add an Attribute definition in the Schema | CC on cn=Schema, cn=Configuration, dc=<ForestRootDomain> (to create objects of class Attribute-Schema) |
| Modify a Class definition in the Schema | WP on the corresponding classSchema object under cn=Schema, cn=Configuration, dc=<ForestRootDomain> |
| Modify an Attribute definition in the Schema | The corresponding attributeSchema object under cn=Schema, cn=Configuration, dc=<ForestRootDomain> is modified |
| Update the Schema cache on demand | WP on the rootDSE object to add and modify the schemaUpdateNow attribute to the object (and set its value equal to 1)<br><br>Extended Right Update-Schema-Cache on cn=Schema, cn=Configuration, dc=<ForestRootDomain> or on the NTDS-Settings object cn=NTDSSettings, cn =<ServerName>, cn=Servers, cn=<SiteName>, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where <ServerName> is the name of the Domain Controller where the operation is being performed |
| Deactivate a Schema Class object /<br><br>Resurrect a deactivated Schema Class object | WP on cn=<Class>, cn=Schema, cn=Configuration, dc=<ForestRootDomain> where <Class> is the classSchema object for the class being deactivated/resurrected to modify the isdefunct attribute |
| Deactivate an Attribute Class object /<br><br>Resurrect a deactivated Schema Attribute object | WP on cn=<Attribute>, cn=Schema, cn=Configuration, dc=<ForestRootDomain> where <Attribute> is the attributeSchema object for the attribute being deactivated/resurrected to modify the isdefunct attribute |
| Make an attribute indexed | WP on the corresponding attributeSchema object cn=<Attribute>, cn=Schema, cn=Configuration, dc=<ForestRootDomain> to modify the searchFlags attribute |
| Add attributes to the ANR Set | WP on the corresponding attributeSchema object cn=<attributeSchema>,cn=Schema, cn=Configuration, dc=<ForestRootDomain> to modify the searchFlags attribute |

| | |
|---|---|
| Designate an attribute as a member of the partial attribute-set that is replicated to the Global Catalog | WP on the corresponding attributeSchema object cn=<Attribute>, cn=Schema, cn=Configuration, dc=<ForestRootDomain> to modify the isMemberOfPartialAttributeSet attribute |
| Remove an attribute from the partial attribute-set that is replicated to the Global Catalog | WP on the corresponding attributeSchema object cn=<Attribute>, cn=Schema, cn=Configuration, dc=<ForestRootDomain> to modify the isMemberOfPartialAttributeSet attribute |

# Operations Master Roles Management

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Transfer the Schema Master Role | Extended Right Change-Schema-Master on cn=Schema, cn=Configuration, dc=<ForestRootDomain> |
| Transfer the Domain Naming Master Role | Extended Right Change-Domain-Master on cn=Partitions, cn=Configuration, dc=<ForestRootDomain> |
| Transfer the RID Master Role | Extended Right Change-RID-Master on cn=RID Manager$, cn=System, dc=<Domain> |
| Transfer the PDC Emulator Master Role | Extended Right Change-PDC on dc=<Domain> (root of domain directory partition) |
| Transfer the Infrastructure Master Role | Extended Right Change-Infrastructure-Master on dc=<Domain> (root of domain directory partition) |
| Seize the Schema Master Role | WP on cn=Schema, cn=Configuration, dc=<ForestRootDomain> to modify the fSMORoleOwner attribute<br>Extended Right Change-Schema-Master on cn=Schema, cn=Configuration, dc=<ForestRootDomain> |
| Seize the Domain Naming Master Role | WP on cn=Partitions, cn=Configuration, dc=<ForestRootDomain> to modify the fSMORoleOwner attribute<br>Extended Right Change-Domain-Master on cn=Partitions, cn=Configuration, dc=<ForestRootDomain> |
| Seize the RID Master Role | WP on cn=RID Manager$, cn=System, dc=<Domain> to modify the fSMORoleOwner attribute<br>Extended Right Change-RID-Master on cn=RID Manager$, cn=System, dc=<Domain> |
| Seize the PDC Emulator Master Role | WP on dc=<Domain> (root of domain directory partition) to modify the fSMORoleOwner attribute<br>Extended Right Change-PDC on dc=<Domain> (root of domain directory partition) |
| Seize the Infrastructure Master Role | WP on cn=Infrastructure, dc=<Domain> to modify the fSMORoleOwner attribute<br>Extended Right Change-Infrastructure-Master on dc=<Domain> (root of domain directory partition) |

# LDAP Policy Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Configure the server to require all LDAP traffic to be signed | The registry entry ldapserverintegrity in HKLM/System/CurrentControlSet/Services/NTDS/Parameters is modified<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Create a new Query Policy object | CC on cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to create objects of class Query-Policy |
| Modify the LDAP admin limits associated with a query policy object | WP on the corresponding Query Policy object under cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the LDAP-Admin-Limits attribute |
| Affect the LDAP query policies associated with a specific DC | WP on the corresponding NTDS-Settings object with distinguished name cn=NTDS Settings, cn=<Computer-Name>, cn=Servers, cn=<SiteName>,cn=Sites,cn=Configuration, dc=<forestRootDomain> to modify the Query-Policy-Object attribute and assign as value the distinguished name of the Query-Policy object that contains the LDAP query policies that should be used for this DC |
| Affect the LDAP query policies associated with all domain controllers in a site | WP on the corresponding NTDS-Site-Settings object with distinguished name cn=NTDS Settings, cn=<Computer-Name>, cn=Servers, cn=<SiteName>,cn=Sites,cn=Configuration, dc=<forestRootDomain> to modify the Query-Policy-object attribute and assign as value the distinguished name of the Query-Policy object that contains the LDAP query policies that should be used for all Domain controllers in this site |
| Specify the maximum time (in seconds) that the server waits for the initial request before the connection closes | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |

| | |
|---|---|
| Specify the maximum number of concurrent LDAP connections allowed on the server | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |
| Specify the maximum amount of time (in seconds) that the client is allowed to be idle before the connection is closed | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |
| Specify the maximum number of concurrent search operations allowed on the server | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |

| | |
|---|---|
| Specify the maximum number of concurrent notification requests allowed per connection on the server | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |
| Specify the maximum number of objects the server will return to any single search request | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |
| Specify the maximum elapsed time (in seconds) allowed for a query to complete | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |

| | |
|---|---|
| Specify the limit (in candidate objects) of the temporary database table the server might create for intermediate results during the course of query | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |
| Specify the total amount of intermediate data that the server will store for the client between the individual searches that make up a paged result search (in order to speed up the next leg of the search) | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |
| Specify the maximum number of threads per processor that can be simultaneously allocated to answer LDAP requests | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br><br>- OR -<br><br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |

| | |
|---|---|
| Specify the maximum size of datagrams that can be received by the server | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br>- OR -<br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br>- OR -<br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |
| Specify the maximum sized LDAP request (in bytes) that the server will attempt to process | WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Settings object is defined and references the <Referenced Query Policy> object<br>- OR -<br>WP on cn=<Referenced Query Policy>,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute of the corresponding NTDS-Site-Settings object is defined and references the <Referenced Query Policy> object<br>- OR -<br>WP on the object cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<forestRootDomain> to modify the ldapAdminLimits attribute IF the Query-Policy attribute is not defined on either the corresponding NTDS-Settings object or the NTDS-Site-Settings object |

# Trust Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Create a shortcut (cross-link) trust relationship | CC on the object cn=System, dc=<trusting domain> (to create objects of class Trusted-Domain) in the trusting domain<br><br>CC on the object cn=System, dc=<trusted domain> (to create objects of class Trusted-Domain) in the trusted domain |
| Create an external trust relationship | CC on the object cn=System, dc=<trusting domain> (to create objects of class Trusted-Domain) in the trusting domain<br><br>CC on the object cn=System, dc=<trusted domain> (to create objects of class Trusted-Domain) in the trusted domain |
| Create a non-Windows Kerberos realm trust relationship | CC on the object cn=System, dc=<trusting domain> (to create objects of class Trusted-Domain) in the domain creating the trust |
| Create an Outbound Forest Trust | CC on the object cn=System, dc=<forestRootDomain> (to create objects of class Trusted-Domain) in the trusting forest<br><br>CC on the object cn=System, dc=<forestRootDomain> (to create objects of class Trusted-Domain) in the trusted forest |
| Create an Inbound Forest Trust | (CC on the object cn=System, dc=<forestRootDomain> (to create objects of class Trusted-Domain) in the trusted forest<br><br>OR Extended right Create-Inbound-Forest-Trust on dc=<Domain> (root of domain directory partition) in the trusted forest)<br><br>AND (CC on the object cn=System, dc=<forestRootDomain> (to create objects of class Trusted-Domain) in the trusting forest) |
| Delete a shortcut (cross-link) trust relationship | SD on the object cn=<trusted domain name>,cn=System, dc=<trusting domain> in the trusting domain<br><br>SD on the object cn=<trusted domain name>,cn=System, dc=<trusted domain> in the trusted domain |
| Delete an external trust relationship | SD on the object cn=<trusted domain name>,cn=System, dc=<trusting domain> in the trusting domain<br><br>SD on the object cn=<trusted domain name>,cn=System, dc=<trusted domain> in the trusted domain |
| Delete a non-Windows Kerberos realm trust relationship | SD on the object cn=<trusted domain name>,cn=System, dc=<trusting domain> in the domain from where the trust is being deleted |
| Delete a forest trust | SD on the object cn=<trusted domain name>,cn=System, dc=<forestRootDomain> in the trusting forest<br><br>SD on the object cn=<trusted domain name>,cn=System, dc=<forestRootDomain> in the trusted forest |

| | |
|---|---|
| Verify that a trust is working properly | Have to be a member of local Administrators group on the machine from where the trust is being verified |
| Change the direction of a trust | WP on the object cn=<Trusted Domain object>, cn=System, dc=<TrustingDomain> where <Trusted Domain object> represents the corresponding trust, to modify the Trust-Direction attribute<br><br>WP on the object cn=<Trusted Domain object>, cn=System, dc=<TrustedDomain> where <Trusted Domain object> represents the corresponding trust, to modify the Trust-Direction attribute |
| Enable Name Suffix Routing (for a given suffix) in a forest | WP on the object cn=<Trusted Domain object>, cn=System, dc=<forestRootDomain> where <Trusted Domain object> represents the corresponding trust, to modify the ms-DS-Trust-Forest-Trust-Info attribute |
| Disable Name Suffix Routing (for a given suffix) in a forest | WP on the object cn=<Trusted Domain object>, cn=System, dc=<forestRootDomain> where <Trusted Domain object> represents the corresponding trust, to modify the ms-DS-Trust-Forest-Trust-Info attribute |
| Add/Remove an exception to a name suffix for a given forest trust | WP on the object cn=<Trusted Domain object>, cn=System, dc=<forestRootDomain> where <Trusted Domain object> represents the corresponding trust, to modify the ms-DS-Trust-Forest-Trust-Info attribute |
| Reset the trust passwords shared by a trust-pair | WP on the object cn=<Trusted Domain object>, cn=System, dc=<trusting and trusted domain> where <Trusted Domain object> represents the corresponding trust, to modify the Initial-Auth-Incoming, Initial-Auth-Outgoing, Trust-Auth-Incoming and Trust-Auth-Outgoing attributes |
| Force the removal of a trust | DC on the object cn=System, dc=<trusting domain> (to delete objects of class Trusted-Domain) in the domain that wants to delete the trust |
| Enable/Disable SID History on an outbound forest trust | WP on the object cn=<Trusted Domain object>, cn=System, dc=<forestRootDomain> where <Trusted Domain object> represents the corresponding trust, to modify the Trust-Attributes attribute |
| Enable/Disable SID Filtering | WP on the object cn=<Trusted Domain object>, cn=System, dc=<Domain> where <Trusted Domain object> represents the corresponding trust, to modify the Trust-Attributes attribute |
| Enable Selective Authentication on an outbound forest/external trust | WP on the object cn=<Trusted Domain object>, cn=System, dc=<Domain> where <Trusted Domain object> represents the corresponding trust, to modify the Trust-Attributes attribute |
| Enable/Disable placing of Name Suffix (Top Level Names) information on a realm trust | WP on the object cn=<Trusted Domain object>, cn=System, dc=<Domain> where <Trusted Domain object> represents the corresponding trust, to modify the Trust-Attributes attribute |
| Add/remove top-level names from a realm trust | WP on the object cn=<Trusted Domain object>, cn=System, dc=<forestRootDomain> where <Trusted Domain object> represents the corresponding trust, to modify the ms-DS-Trust-Forest-Trust-Info attribute |
| Add/remove top-level name exclusions from a realm trust | WP on the object cn=<Trusted Domain object>, cn=System, dc=<forestRootDomain> where <Trusted Domain object> represents the corresponding trust, to modify the ms-DS-Trust-Forest-Trust-Info attribute |
| Modify the transitivity of a realm-trust | WP on the object cn=<Trusted Domain object>, cn=System, dc=<Domain> where <Trusted Domain object> represents the corresponding trust, to modify the Trust-Attributes attribute |

* All trust management tools in Active Directory require that an administrator performing any trust management task using these tools be a member of the BuiltIn Admins group in the domain.

# Replication Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Create a Site / Add a Site | CC on cn=Sites, cn=Configuration, dc=<ForestRootDomain> (to create objects of class Site) |
| Specify the location of a Site | WP on the corresponding site object, cn=<Site>, cn=Sites, cn=Configuration, dc=<forestRootDomain> to modify the Location attribute |
| Associate a Group Policy with a Site | WP on the corresponding site object, cn=<Site>, cn=Sites, cn=Configuration, dc=<forestRootDomain> to modify the GP-Link attribute<br>WP on the corresponding site object, cn=<Site>, cn=Sites, cn=Configuration, dc=<forestRootDomain> to modify the GP-Options attribute |
| Modify Site Group Policy Options | WP on the corresponding site object, cn=<Site>, cn=Sites, cn=Configuration, dc=<forestRootDomain> to modify the GP-Options attribute |
| Disable automatic topology generation for a site | WP on cn=NTDSSiteSettings, cn=<SiteName>, cn=Sites, cn=Configuration, dc=ForestRootDomain where the <SiteName> is the name of the site, to modify the options attribute |
| Disable automatic topology cleanup for a site | WP on cn=NTDSSiteSettings, cn=<SiteName>, cn=Sites, cn=Configuration, dc=ForestRootDomain where the <SiteName> is the name of the site, to modify the options attribute |
| Disable minimum hops topology for a site | WP on cn=NTDSSiteSettings, cn=<SiteName>, cn=Sites, cn=Configuration, dc=ForestRootDomain where the <SiteName> is the name of the site, to modify the options attribute |
| Disable automatic stale server detection for a site | WP on cn=NTDSSiteSettings, cn=<SiteName>, cn=Sites, cn=Configuration, dc=ForestRootDomain where the <SiteName> is the name of the site, to modify the options attribute |
| Disable automatic inter-site topology generation for a site | WP on cn=NTDSSiteSettings, cn=<SiteName>, cn=Sites, cn=Configuration, dc=ForestRootDomain where the <SiteName> is the name of the site, to modify the options attribute |
| Disable Inbound Replication on a DC | WP on the corresponding NTDS Settings object with distinguished name cn=NTDS Settings, cn=<Computer-Name>, cn=Servers, cn=<SiteName>,cn=Sites,cn=Configuration, dc=<forestRootDomain> to modify the options attribute |
| Disable Outbound Replication on a DC | WP on the corresponding NTDS Settings object with distinguished name cn=NTDS Settings, cn=<Computer-Name>, cn=Servers, cn=<SiteName>,cn=Sites,cn=Configuration, dc=<forestRootDomain> to modify the options attribute |

| | |
|---|---|
| Delete a Site | SD on the site object itself OR DC on cn=Sites, cn=Configuration, dc=<ForestRootDomain> (to delete objects of class Site).<br>**Note**<br>DC on parent will grant permission to delete all objects under the parent (and if class is specified, then only all objects of specified class). |
| Create a Subnet / Add a Subnet | CC on cn=Subnets, cn=Sites, cn=Configuration, dc=<ForestRootDomain> (to create objects of type Subnet) |
| Specify the location of a Subnet | WP on the corresponding subnet object, cn=<Subnet>, cn=Subnets, cn=Sites, cn=Configuration, dc=<forestRootDomain> to modify the Location attribute |
| Associate a Subnet with a Site | WP on the corresponding subnet object cn=<SubnetName>, cn=Subnets, cn=Sites, cn=Configuration, dc=<ForestRootDomain> to modify the siteObject attribute |
| Delete a Subnet | DC on cn=Subnets, cn=Sites, cn=Configuration, dc=<ForestRootDomain> (to delete objects of class Subnet) |
| Create a Site Link | CC on cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> OR on cn=SMTP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> (to create objects of class siteLink) |
| Add/Remove sites to/from a Site Link | WP on cn=<siteLink>, cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> OR on cn=<siteLink>, cn=SMTP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where <siteLink> is the site link to/from which a new site is being added/removed, to modify the site-list attribute |
| Modify the cost associated with a site link | WP on the siteLink object cn=<SiteLink>, cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> OR on the siteLink object cn=<SiteLink>, cn=SMTP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where the <SiteLink> identifies the associated site link, to modify the cost attribute. |
| Modify the replication period associated with a site link / Control link availability | WP on the siteLink object cn=<SiteLink>, cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where the <SiteLink> identifies the associated site link, to modify the Repl-Interval attribute |
| Modify the replication schedule for a site link | WP on the corresponding site link object cn=<siteLink>, cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> OR on cn=<siteLink>,cn=SMTP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain>, to modify the schedule attribute |
| Delete a Site Link | DC on cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> OR on cn=SMTP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> (to delete objects of class siteLink) |
| Create a Site Link bridge (object) | CC on cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> OR on cn=SMTP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> (to create objects of class siteLinkBridge) |

| | |
|---|---|
| Add/Remove sites to/from a Site Link Bridge | WP on cn=<siteLinkBridge>, cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> OR on cn=<siteLinkBridge>, cn=SMTP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where <siteLinkBridge> is the site link ridge to/from which a new site is being added/removed, to modify the site-link-list attribute |
| Create a single bridge for the entire network / Turn off the "Bridge all site links" option for IP/SMTP transport | WP on the corresponding (IP/SMTP) interSiteTransport object cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> OR on cn=SMTP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain>, to modify the options attribute |
| Enable Reciprocal Replication between sites (only for IP transport links) | WP on cn=<SiteLink>, cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where <SiteLink> identifies the associated site link, to modify the options attribute |
| Enable Change Notification between sites (only for IP transport links) | WP on cn=<SiteLinkName>, cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where the SiteLinkName identifies the associated site link, to modify the options attribute |
| Delete a Site Link bridge (object) | DC on cn=IP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> OR on cn=SMTP, cn=Inter-Site Transports, cn=Sites, cn=Configuration, dc=<ForestRootDomain> (to delete objects of class siteLinkBridge) |
| Create a Connection (object) | CC on cn=NTDSSettings, cn =<ServerName>, cn=Servers, cn=<SiteName>, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where ServerName is the name of the DC to which the connection in inbound (to create objects of class NTDS-Connection) |
| Take ownership of a KCC-generated connection object | WP on cn=<ConnectionName>, cn=NTDSSettings, cn =<ServerName>, cn=Servers, cn=<SiteName>, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where <ConnectionName> is the name of the KCC-generated connection, to modify the options attribute |
| Manually set a schedule for connection objects | WP on cn=<ConnectionName>, cn=NTDSSettings, cn =<ServerName>, cn=Servers, cn=<SiteName>, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where <ConnectionName> is the name of the KCC-generated connection, to modify the options attribute |
| Enable/disable data compression for intersite replication | WP on cn=<ConnectionName>, cn=NTDSSettings, cn =<ServerName>, cn=Servers, cn=<SiteName>, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where <ConnectionName> is the name of the KCC-generated connection, to modify the options attribute |
| Delete a Connection (object) | DC on cn=NTDSSettings, cn =<ServerName>, cn=Servers, cn=<SiteName>, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where ServerName is the name of the DC to which the connection in inbound (to delete objects of class NTDS-Connection)<br>**Note**<br>An NTDS-Connection object created by the KCC should not be deleted. If it is, the KCC will regenerate it. Only a manually created NTDS-Connection object might be deleted. |
| Change the default setting for the intra-site replication schedule within a site | WP on cn=NTDSSiteSettings, cn=<SiteName>, cn=Sites, cn=Configuration, dc=ForestRootDomain where the <SiteName> is the name of the site, to modify the schedule attribute |

| | |
|---|---|
| Designate / Remove a preferred bridgehead server | WP on cn =<ServerName>, cn=Servers, cn=<SiteName>, cn=Sites, cn=Configuration, dc=<ForestRootDomain> where ServerName is the name of the server being designated as a Preferred Bridgehead server, to modify the Bridgehead-Transport-List attribute |
| Replace a failed Preferred Bridgehead Server | Do one of the following?:<br>Add new domain controllers as preferred bridgehead servers for the corresponding directory partitions, site and transport<br>- OR -<br>Remove all preferred bridgehead designations made for the corresponding site and transport (for the corresponding directory partition), in which case KCC selects new ones automatically; remove them for each domain directory partition and for each transport on a DC in each affected site |
| Specify a fixed-port for RPC-based replication | WP on HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters to modify the registry entry TCP/IP Port |
| Adjust default size of packets that transport Active Directory replication data | The following registry entries (with registry path HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameter) are added/modified (with the REG_DWORD data type):<br>• For RPC replication within a site: Replicator intra site packet size (objects) & Replicator intra site packet size (bytes)<br>• For RPC replication between sites: Replicator inter site packet size (objects) & Replicator inter site packet size (bytes)<br>• For SMTP replication within a site: Replicator async inter site packet size (objects) & Replicator async inter site packet size (bytes)<br>Thus, appropriate permissions required to Create and/or modify these registry keys will be required to delegate the operation |
| Increase the level of detail logged by the KCC in the event log | WP on HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics to modify the 1 Knowledge Consistency Checker entry |
| Modify the interval at which the KCC runs its first replication topology after the DC starts | WP on HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters to modify the registry entry Repl topology update delay (secs) |
| Modify the interval at which the KCC checks the replication topology (after it has run the first time) | WP on HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters to modify the registry entry Repl topology update period (secs) |
| Force Replication Topology Generation | Extended right Manage Replication Topology needed on cn=configuration, dc=<forestRootDomain> |
| Modify the holdback timer that determines the interval between the time a change is made and the time that the source server notifies its replication partners within a site | WP on HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters to modify the registry entry Replicator notify pause after modify (secs) |

| | |
|---|---|
| Modify the default delay between notifications to all the replication partners of a DC | WP on HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters to modify Replicator notify pause between DSAs (secs) |
| Force replication between two servers | Extended right Replication Synchronization needed on cn=configuration, dc= <forestRootDomain> |
| Force a synchronization between two servers | Extended right Replication Synchronization needed on cn=configuration, dc= <forestRootDomain> |
| Set a DC not to contact the PDC emulator if the PDC emulator role owner is not in the current site | WP on HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\ Services\NTDS\Parameters to modify the registry entry AvoidPdcOnWan |
| Modify the thresholds that make the KCC exclude non-responding servers when it recognizes that a DC has failed or is unresponsive | The following registry entries (with registry path HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameter) are added/modified (with the REG_DWORD data type):<br>For replication between sites:<br>• IntersiteFailuresAllowed and MaxFailureTimeForIntersiteLink (secs)<br>• For optimizing connections within a site: NonCriticalLinkFailuresAllowed and MaxFailureTimeForNonCriticalLink<br>• For immediate neighbor connections within a site: CriticalLinkFailuresAllowed and MaxFailureTimeForCriticalLink<br>Thus, appropriate permissions required to Create and/or modify these registry keys will be required to delegate the operation |
| Get Replication Latency Information | In Windows 2000, Extended right Manage Replication Topology needed on domain NC head<br>In Windows Server 2003, Extended right Monitor Replication Topology or Manage Replication Topology needed on domain NC head |
| Get Pending operations on DC ( Queue Length ) | In Windows 2000, Extended right Manage Replication Topology needed on domain NC head<br>In Windows Server 2003, Extended right Monitor Replication Topology or Manage Replication Topology needed on domain NC head |
| Check Replication Status | In Windows 2000, Extended right Manage Replication Topology needed on domain NC head<br>In Windows Server 2003, Extended right Monitor Replication Topology or Manage Replication Topology needed on domain NC head |

# Backup and Restore Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
| --- | --- |
| Back up Active Directory | User-right Backup Files and Directories is required.<br>User performing task needs to be member of Administrators (BuiltIn\Admins) group on DC because NTBACKUP required user to be Administrator to perform System-State backup. |
| Perform an Non-Authoritative restore of Active Directory from Backup Media | In Stand-alone mode:<br>• User-right Restore Files and Directories is required.<br>• User needs to be member of local Administrators group on server because NTBACKUP required user to be Administrator to perform System-State restore.<br>In DS-Restore mode:<br>• User needs to be Administrator in DS Restore mode. |
| Perform an Authoritative restore of Active Directory from Backup Media | Requires user to be Administrator in DS Restore mode. |

# Directory Database File Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| **Task** | **Permissions Required to Perform Task** |
|---|---|
| Perform an online defragmentation of the Ntds.dit database | Add the doOnlineDefrag attribute to the rootDSE object & set its value equal to set to number in seconds the online defrag should run<br><br>Extended Right Do-Garbage-Collection on NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> |
| Perform an offline defragmentation of the Ntds.dit database | Full control on ntds.dit file and on the folder containing the logs (default location: *systemroot*\NTDS\ ) |
| Move the Ntds.dit file to a new location | WP on HKLM/System/CCS/Services/NTDS/Parameters,<br><br>Full control on the ntds.dit file<br><br>Full control on source and target folders |
| Move the directory service log files to a new location | WP on HKLM/System/CCS/Services/NTDS/Parameters,<br><br>Full control on the log folders<br><br>Full control on source and target folders |
| Perform a soft recovery of the database | Full control on ntds.dit file and on the folder containing the logs (default location: *systemroot*\NTDS\ ) |
| Specify the location of the Ntds.dit file | WP on HKLM/System/CCS/Services/NTDS/Parameters<br><br>Read Control and Modify Permissions required on .dit file |
| Specify the location of the log files | WP on HKLM/System/CCS/Services/NTDS/Parameters<br><br>Read Control and Modify Permissions required on the log folder |
| Specify the Active Directory working directory | WP on HKLM/System/CCS/Services/NTDS/Parameters<br><br>Read Control and Modify Permissions required on the working directory folder |
| Restore Database/Subtree of database | Full control on ntds.dit file and on the folder containing the logs (default location: *systemroot*\NTDS\ ) |
| Perform semantic database analysis | Full control on ntds.dit file and on the folder containing the logs (default location: *systemroot*\NTDS\ ) |

# Directory Service Configuration Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Designate a DC as a Global Catalog | WP on the corresponding NTDS Settings object with distinguished name cn=NTDS Settings, cn=<Computer-Name>, cn=Servers, cn=<SiteName>,cn=Sites,cn=Configuration, dc=<forestRootDomain> to modify the options attribute |
| Force the directory service to do garbage collection | Add the doGarbageCollection attribute on the rootDSE object, and set its value equal to any value<br><br>Extended Right Do-Garbage-Collection on NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> |
| Specify the directory service garbage collection period | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the Garbage-Coll-Period attribute |
| Force the directory service to recalculate the Exchange Address Book information hierarchy | Add the recalcHierarchy attribute on the rootDSE object, and set its value equal to any value<br><br>Extended Right Recalculate-Hierarchy on NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> |
| Update the Schema cache on demand | Add and modify the schemaUpdateNow attribute on the rootDSE object, and set its value equal to 1<br><br>Extended Right Update-Schema-Cache on cn=Schema, cn=Configuration, dc=<ForestRootDomain> |
| Force directory service to recompute ACL inheritance on a naming context | Add the recalcHierarchy attribute on the rootDSE object and set its value equal to 'forceupdate'<br><br>Extended Right Recalculate-Security-Inheritance on NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> |
| Force the directory service to check stale phantom objects | On the rootDSE object, add the checkPhantoms attribute to the object and set its value equal to any value<br><br>Extended Right DS-Check-Stale-Phantoms on NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> |
| Force the directory service to immediately refresh the group cache by contacting an available GC | on the rootDSE object, add the updateCachedMemberships attribute to the object and set its value equal to any value<br><br>Extended Right Refresh-Group-Cache on NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> |

| | |
|---|---|
| Force the directory service to remove lingering objects from a Domain Controller | Add the removeLingeringObject attribute on the rootDSE object and set its value equal to the object to remove in the following form: DN_OF_THE_SOURCE_NTDS_SETTINGS_OBJECT:DN_OF_THE_OBJECT_TO_REMOVE<br><br>Extended right Manage Replication Topology needed on the cn=configuration, dc=<forestRootDomain> |
| Reanimate Tombstones | Extended Right Reanimate Tombstones on CN=Schema, CN=Configuration, DC=<forestRootDomain> |
| Force the directory service to perform an online defrag on a Domain Controller | Add the doOnlineDefrag attribute on the rootDSE object and set its value equal to any value.<br><br>Extended Right Do-Garbage-Collection on NTDS-Settings object CN=NTDS Settings, CN=<Server>, CN=<Site>,CN=Sites, CN=Configuration, DC=<forestRootDomain> |
| Specify the default amount of time a dynamic object will exist in the directory | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the ms-DS-Other-Settings attribute |
| Specify the minimum amount of time a dynamic object will exist in the directory | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the ms-DS-Other-Settings attribute |
| Specify the delay between deleting a server object and it being permanently removed from the replication topology | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the Repl-Topology-Stay-Of-Execution attribute |
| Specify the number of days before a deleted object is removed from the directory (tombstone lifetime) | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the Tombstone-Lifetime attribute |
| Adjust ANR searching behavior | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the DS-Heuristics attribute |
| Put the directory in the special “List Object” mode | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the DS-Heuristics attribute |
| Restrict anonymous operations (other than rootDSE searches and binds) through LDAP | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the DS-Heuristics attribute |
| Control the behavior of the *userPassword* attribute | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the DS-Heuristics attribute |
| Specify which SPN types are mapped to “host” | WP on the NTDS-Service object cn=Directory Service,cn=Windows NT,cn=Services,CN=Configuration, DC=<forestRootDomain> to modify the SPN-mappings attribute |
| Increase the level of detail logged by the KCC in the event log | Modify the 1 Knowledge Consistency Checker entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |

| | |
|---|---|
| Modify the level of detail logged for Security Events | Modify the 2 Security Events entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by events related to communication between Active Directory and Exchange clients | Modify the 3 ExDS Interface Events entry and the 4 MAPI Interface Events entry under KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged when objects marked for deletion are actually deleted | Modify the 6 Garbage Collection entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by directory service operations | Modify the 7 Internal Configuration entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by directory access events | Modify the 8 Directory Access entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by internal operation of directory service code | Modify the 9 Internal Processing entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by events related to loading and unloading the NTDS performance object and performance counters | Modify the 10 Performance Counters entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by events related to starting and stopping the directory service | Modify the 11 Initialization/Termination entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by directory service events | Modify the 12 Service Control entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by the events related to address resolution and Active Directory names | Modify the 13 Name Resolution entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |

| | |
|---|---|
| Modify the level of detail logged by the events related to the backup of Active Directory | Modify the 14 Backup entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by events related to LDAP | Modify the 16 LDAP Interface entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by events related to running the Active Directory Installation wizard | Modify the 17 Setup entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by events related to the Global Catalog | Modify the 18 Global Catalog entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Modify the level of detail logged by events the Inter-site messaging service | Modify the 19 Inter-site Messaging entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics<br><br>Thus, appropriate permissions required to modify this registry key will be required to delegate the operation |
| Designate a DC as a Global Catalog | WP on the corresponding NTDS Settings object with distinguished name cn=NTDS Settings, cn=<Computer-Name>, cn=Servers, cn=<SiteName>,cn=Sites,cn=Configuration, dc=<forestRootDomain> to modify the options attribute |

# Security Policy Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
| --- | --- |
| Modify the default Domain Controller Group Policy | Requires “Edit Settings” on the Default Domain Controllers GPO |
| Modify the default Domain Policy | Requires “Edit Settings” on the Default Domain GPO |

# DNS Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Windows 2000, in addition to the permissions specified in the table "DNS Management Tasks", Full Control is required on the container CN=MicrosoftDNS,CN=System,DC=<domain> to perform any DNS administrative task. In Windows Server 2003, the permissions specified in the table "DNS Management Tasks" are sufficient to perform the corresponding DNS administrative task.

| Task | Permissions Required to Perform Task |
|---|---|
| Create a new Active Directory–integrated zone | CC on CN=MicrosoftDNS,CN=System,DC=<domain> OR CN=MicrosoftDNS,DC=<domain DNS Zones application directory partition> to create objects of class dnsZone<br>RP on CN=MicrosoftDNS,CN=System,DC=<domain> |
| Delete an Active Directory–integrated zone | SD on the DNS-Zone object itself OR DC on CN=MicrosoftDNS,CN=System,DC=<domain> OR CN=MicrosoftDNS,DC=<domain DNS Zones application directory partition> to delete objects of class dnsZone<br>RP on CN=MicrosoftDNS,CN=System,DC=<domain> |
| Write Active Directory–integrated zone parameters | WP on dnsZone object under CN=MicrosoftDNS,CN=System,DC=<domain> OR CN=MicrosoftDNS,DC=<domain DNS Zones application directory partition> to modify the dnsProperty attribute<br>RP on CN=MicrosoftDNS,CN=System,DC=<domain> |
| Write the RootHints (stored in Active Directory) | WP on dnsNode objects stored in CN=RootDNSServers under CN=MicrosoftDNS,CN=System,DC=<domain> OR CN=MicrosoftDNS,DC=<domain DNS Zones application directory partition> to modify the dnsRecord attribute<br>RP on CN=MicrosoftDNS,CN=System,DC=<domain> |
| Create a new name in the Active Directory–integrated zone | CC on corresponding dnsZone object CN=<dnsZone>, CN=MicrosoftDNS,CN=System,DC=<domain> OR CN=<dnsZone>, CN=MicrosoftDNS,DC=<domain DNS Zones application directory partition> to create objects of class dnsNode<br>RP on CN=MicrosoftDNS,CN=System,DC=<domain> |
| Write the records in the Active Directory–integrated zone | WP on the corresponding dnsNode object CN=<dnsNode>, CN=<dnsZone>, CN=MicrosoftDNS,CN=System,DC=<domain> OR CN=<dnsNode>,CN=<dnsZone>, CN=MicrosoftDNS,DC=<domain DNS Zones application directory partition> to modify the dnsRecord attribute<br>RP on CN=MicrosoftDNS,CN=System,DC=<domain> |

# Active Directory Data Administration Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section contains the following Active Directory data administration tasks:

- Organizational Unit (OU) Management Tasks
- Group Management Tasks
- Computer Management Tasks
- User Account Management Tasks
- Print Queue Management Tasks
- Service Connection Points Management Tasks
- Generic Data Management Tasks
- Group Policy Management Tasks

# Organizational Unit (OU) Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Create an Organizational Unit | CC on parent object (to create objects of class Organizational-Unit) |
| Delete an Organizational Unit | SD on the object itself OR DC on parent object (to delete objects of class Organizational-Unit)<br>**Note**<br>DC on parent will grant permission to delete all objects under the parent (and if class is specified, then only all objects of specified class). |
| Rename an Organizational Unit | WP on the organizational unit object to modify Organizational-Unit-Name attribute<br>WP on the organizational unit object to modify RDN attribute |
| Move an Organizational Unit | SD on the organizational unit object itself OR DC on current parent (to delete objects of class Organizational-Unit)<br>CC on target parent (to create objects of class Organizational-Unit)<br>WP on the organizational unit object to modify Organizational-Unit-Name attribute<br>WP on the organizational unit object to modify RDN attribute |
| Modify Description of an Organizational Unit | WP on the organizational unit object to modify Description attribute |
| Modify Street of an Organizational Unit | WP on the organizational unit object to modify Street-Address attribute |
| Modify City/Province an Organizational Unit | WP on the organizational unit object to modify Locality-Name attribute |
| Modify State of an Organizational Unit | WP on the organizational unit object to modify State-Or-Province-Name attribute |
| Modify Zip/Postal Code of an Organizational Unit | WP on the organizational unit object to modify Postal-Code attribute |
| Modify Country/Region of an Organizational Unit | WP on the organizational unit object to modify Country-Name attribute |
| Modify Managed-By Information of an OU | WP on the organizational unit object to modify Managed-By attribute |
| Change the COM+ partition set that an Organizational Unit is a member of | WP on the organizational unit object to modify ms-COM-UserPartitionSetLink attribute |
| Modify the Group Policy applied to an Organizational Unit | WP on the organizational unit object to modify the GP-Link attribute<br>WP on the organizational unit object to modify the GP-Options attribute |
| Delegate Control of an Organizational Unit | WD on the organizational unit object |

# Group Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Create a group | CC on parent object (to create objects of class Group) |
| Delete a group | SD on the group object itself OR DC on parent object (to delete objects of class Group)<br>**Note**<br>DC on parent will grant permission to delete all objects under the parent (and if class is specified, then only all objects of specified class). |
| Move a group | SD on the group object itself OR DC on parent object (to delete objects of class Group)<br>CC on target parent (to create objects of class Group)<br>WP on the group object to modify Common-Name attribute<br>WP on the group object to modify RDN attribute |
| Rename a group | WP on the group object to modify Common-Name attribute<br>WP on the group object to modify RDN attribute |
| Specify the Pre-Windows 2000 compatible name for the group | WP on the group object to modify SAM-Account-Name attribute |
| Modify the description of a group | WP on the group object to modify Description attribute |
| Modify the e-mail address for a group | WP on the group object to modify E-mail-Addresses attribute |
| Modify the scope of the group | WP on the group object to modify Group-Type attribute |
| Modify the type of the group | WP on the group object to modify Group-Type attribute |
| Modify notes for a group | WP on the group object to modify Comment attribute |
| Modify group membership | WP on the group object to modify Members attribute |
| Specify Managed-By Information of a Group | WP on the group object to modify Managed-By attribute |

# Computer Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Create a computer account | CC on parent object (to create objects of class Computer) |
| Delete a computer account | SD on the computer object itself OR DC on parent object (to delete objects of class Computer) |
| Rename a computer account | WP on the computer object to modify all attributes<br>**Note**<br>User performing operation must be a Local Administrator on the computer being renamed. |
| Move a computer account | SD on the computer object itself OR DC on parent object (to delete objects of class Computer)<br>CC on target parent (to create objects of class Computer)<br>WP on the computer object to modify Common-Name attribute<br>WP on the computer object to modify RDN attribute |
| Disable a computer account | WP on the computer object to modify User-Account-Control attribute |
| Reset a computer account | The Force-User-Change-Password extended right is required on the computer object<br>**Note**<br>In the UI, this extended right corresponds to Reset Password. |
| Add a computer account to a group | WP on the target group object to modify Member attribute |
| Specify the Pre-Windows 2000 compatible name for a computer | WP on the computer object to modify SAM-Account-Name attribute |
| Set a computer's DNS name | Validated-DNS-Host-Name SW on the computer object |
| Specify a computer's role | WP on the computer object to modify Machine-Role attribute |
| Specify the computer's description | WP on the computer object to modify Description attribute |
| Specify the computer's location | WP on the computer object to modify Location attribute |
| Specify Managed-By information for a computer account | WP on the computer object to modify Managed-By attribute |
| Specify the Operating System running on a computer | WP on the computer object to modify Operating-System attribute |

| | |
|---|---|
| Specify the Operating System Service Pack for a computer | WP on the computer object to modify Operating-System-Service-Pack attribute |
| Specify the Operating System Version for the Computer | WP on the computer object to modify Operating-System-Version attribute |
| Specify a computer's physical location | WP on the computer object to modify Physical-Location-Object attribute |
| Specify that a computer account be trusted for delegation | WP on the computer object to modify User-Account-Control attribute<br>The Enable computer and user accounts to be trusted for delegation user right is required — modified in Default Domain Controller Security Policy |
| Specify whether a computer account can be trusted for delegation to any service (Kerberos only) | User right "Enable User and Computer account to be trusted for Delegation" required (assigned in default Domain Controller Policy) |
| Specify that a computer account be trusted for delegation to specific services only | User right "Enable User and Computer account to be trusted for Delegation" required (assigned in default Domain Controller Policy)<br>WP on the computer object to modify msDS-AllowedToDelegateTo attribute |
| Specify "Use Kerberos Only" | User right "Enable User and Computer account to be trusted for Delegation" required (assigned in default Domain Controller Policy)<br>WP on the computer object to modify msDS-AllowedToDelegateTo attribute |
| Specify "Use any authentication protocol" | User right "Enable User and Computer account to be trusted for Delegation" required (assigned in default Domain Controller Policy)<br>WP on the computer object to modify msDS-AllowedToDelegateTo attribute |
| Add/Remove the services to which a computer account can be present delegated credentials | WP on the computer object to modify msDS-AllowedToDelegateTo attribute |

# User Account Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Create a user account in disabled state | CC on parent object (to create objects of class User) |
| Create a user account | CC on parent object (to create objects of class User)<br>WP on the user object to modify User-Account-Control attribute<br>Extended Right "Reset Password" required on user account |
| Delete a user account | SD on the user object itself OR DC on parent object (to delete objects of class User)<br>**Note**<br>DC on parent will grant permission to delete all objects under the parent (and if class is specified, then only all objects of specified class). |
| Rename a user account | WP on the user object to modify Common-Name attribute<br>WP on the user object to modify RDN attribute<br>WP on the user object to modify Obj-Dist-Name attribute |
| Move a user account | SD on the user object itself OR DC on parent object (to delete objects of class User)<br>CC on target parent (to create objects of class User)<br>WP on the user object to modify Common-Name attribute<br>WP on the user object to modify RDN attribute |
| Disable a user account | WP on the user object to modify User-Account-Control attribute |
| Unlock a user account | WP on the user object to modify the Lockout-Time attribute |
| Enable a disabled user account | WP on the user object to modify User-Account-Control attribute |
| Reset a user account's password | The User-Change-Password extended right is required on the user object |
| Force a user account to change the password at the next logon | WP on the user object to modify User-Password attribute AND The User-Force-Change-Password extended right is required on the User object.<br>ALTERNATIVELY, WP on the user object to modify Pwd-Last-Set attribute also works |
| Modify a user's first name | WP on the user object to modify Given-Name attribute |
| Modify a user's initials | WP on the user object to modify Initials attribute |
| Modify a user's last name | WP on the user object to modify Surname attribute |
| Modify a user's display name | WP on the user object to modify Admin-Display-Name attribute |

| | |
|---|---|
| Modify a user account's description | WP on the user object to modify Description attribute |
| Modify a user's office location | WP on the user object to modify Physical-Delivery-Office-Name attribute |
| Modify a user's telephone number | WP on the user object to modify Telephone-Number attribute |
| Modify the location of a user's primary web page | WP on the user object to modify WWW-Home-Page attribute |
| Modify a user's e-mail address | WP on the ser object to modify E-Mail-Address attribute |
| Modify a user's street address | WP on the user object to modify Street-Address attribute |
| Modify a user's P.O box | WP on the user object to modify Post-Office-Box attribute |
| Modify a user's city/province | WP on the user object to modify Locality-Name attribute |
| Modify a user's state | WP on the user object to modify State-Or-Province-Name attribute |
| Modify a user's zip/postal code | WP on the user object to modify Postal-Code attribute |
| Modify a user's country/region | WP on the user object to modify Country-Name attribute |
| Modify a user's UPN | WP on the user object to modify User-Principal-Name attribute |
| Modify a user's Pre-Windows 2000 user logon name | WP on the user object to modify SAM-Account-Name attribute |
| Modify the hours during which a user can log on | WP on the user object to modify Logon-Hours attribute |
| Specify the computers from which a user can log on | WP on the user object to modify User-Workstations attribute |
| Set *User cannot change password* for a user account | WD on the user object<br>OTE: Granting WD is equivalent of granting Full-Control<br>**Note**<br>The Active Directory Users and Computers Snap-In does not allow this operation to be performed from UI. Use dsacls.exe to perform operation. |
| Set *Password Never Expires* for a user account | WP on the user object to modify User-Account-Control attribute |
| Set *Store Password Using Reversible Encryption* for a user account | WP on the user object to modify User-Account-Control attribute |
| Disable a user account | WP on the user object to modify User-Account-Control attribute |
| Set *Smart card is required for interactive logon* for a user account | WP on the user object to modify User-Account-Control attribute |
| Set *Account is sensitive and cannot be delegated* for a user account | WP on the user object to modify User-Account-Control attribute |

| | |
|---|---|
| Set *Use DES encryption types for this account* for a user account | WP on the user object to modify User-Account-Control attribute |
| Set *Do not require Kerberos pre-authentication* for a user account | WP on the user object to modify User-Account-Control attribute |
| Specify the date when a user account expires | WP on the user object to modify Account-Expires attribute |
| Specify a profile path for a user | WP on the user object to modify Profile-Path attribute |
| Specify a logon script for a user | WP on the user object to modify Script-Path attribute |
| Specify the drive letter to which to map the UNC path specified by the home directory for a user account | WP on the user object to modify Home-Drive attribute<br>WP on the user object to modify Home-Directory attribute (This is not changed, but is needed to enable editing in the ADU&C UI) |
| Specify a user's home folder local path | WP on the user object to modify Home-Directory attribute<br>WP on the user object to modify Home-Drive attribute (This is not changed, but is needed to enable editing in the ADU&C UI) |
| Specify the home folder to connect to for a user account | WP on the user object to modify Home-Drive attribute<br>WP on the user object to modify Home-Directory attribute |
| Specify a user's home telephone number | WP on the user object to modify Phone-Home-Primary attribute |
| Specify the user's other Home Telephone numbers | WP on the user object to modify Phone-Home-Other attribute |
| Specify a user's pager number | WP on the user object to modify Phone-Pager-Primary attribute |
| Specify other pager numbers for a user | WP on the user object to modify Phone-Pager-Other attribute |
| Specify a user's mobile number | WP on the user object to modify Phone-Mobile-Primary attribute |
| Specify other mobile numbers for a user | WP on the user object to modify Phone-Mobile-Other attribute |
| Specify a user's facsimile number | WP on the user object to modify Facsimile-Telephone-Number attribute |
| Specify other facsimile numbers for a user | WP on the user object to modify Phone-Fax-Other attribute |
| Specify a user's IP phone number | WP on the user object to modify Phone-IP-Primary attribute |
| Specify other IP phone numbers for a user | WP on the user object to modify Phone-IP-Other attribute |
| Modify notes for a user account | WP on the user object to modify Comment attribute |
| Specify a user's title | WP on the user object to modify Title attribute |
| Specify a user's department | WP on the user object to modify Department attribute |
| Specify a user's manager | WP on the user object to modify Manager attribute |
| View certificates issued to a user | WP on the user object to modify X-509 Cert attribute |

| | |
|---|---|
| Add certificates from store for a user | WP on the user object to modify X-509 Cert attribute |
| Add certificates from file for a user | WP on the user object to modify X-509 Cert attribute |
| Remove a certificate for a user | WP on the user object to modify X-509 Cert attribute |
| Copy a user's certificate to a file | Create File/Write Data permissions on target parent folder on file-system<br>RP on the user object to modify X-509 Cert attribute<br>WP on the user object to modify X-509 Cert attribute |
| Add a user account to a group | WP on the target Group object to modify Member attribute |
| Remove the user from a group | WP on the target Group object to modify Member attribute |
| Set the Primary Group (used for POSIX Compliance) for a user | WP on the target user object to modify Primary-Group-ID attribute |
| Create a user account in disabled state | CC on parent object (to create objects of class User) |
| Create a user account | CC on parent object (to create objects of class User)<br>WP on the user object to modify User-Account-Control attribute<br>Extended Right "Reset Password" required on user account |
| Delete a user account | SD on the user object itself OR DC on parent object (to delete objects of class User)<br>**Note**<br>DC on parent will grant permission to delete all objects under the parent (and if class is specified, then only all objects of specified class). |
| Rename a user account | WP on the user object to modify Common-Name attribute<br>WP on the user object to modify RDN attribute<br>WP on the user object to modify Obj-Dist-Name attribute |
| Move a user account | SD on the user object itself OR DC on parent object (to delete objects of class User)<br>CC on target parent (to create objects of class User)<br>WP on the user object to modify Common-Name attribute<br>WP on the user object to modify RDN attribute |
| Disable a user account | WP on the user object to modify User-Account-Control attribute |
| Unlock a user account | WP on the user object to modify the Lockout-Time attribute |
| Enable a disabled user account | WP on the user object to modify User-Account-Control attribute |
| Reset a user account's password | The User-Change-Password extended right is required on the user object |

| | |
|---|---|
| Force a user account to change the password at the next logon | WP on the user object to modify User-Password attribute<br>AND The User-Force-Change-Password extended right is required on the User object.<br><br>ALTERNATIVELY, WP on the user object to modify Pwd-Last-Set attribute also works |
| Modify a user's first name | WP on the user object to modify Given-Name attribute |
| Modify a user's initials | WP on the user object to modify Initials attribute |
| Modify a user's last name | WP on the user object to modify Surname attribute |
| Modify a user's display name | WP on the user object to modify Admin-Display-Name attribute |
| Modify a user account's description | WP on the user object to modify Description attribute |
| Modify a user's office location | WP on the user object to modify Physical-Delivery-Office-Name attribute |
| Modify a user's telephone number | WP on the user object to modify Telephone-Number attribute |
| Modify the location of a user's primary web page | WP on the user object to modify WWW-Home-Page attribute |
| Modify a user's e-mail address | WP on the ser object to modify E-Mail-Address attribute |
| Modify a user's street address | WP on the user object to modify Street-Address attribute |
| Modify a user's P.O box | WP on the user object to modify Post-Office-Box attribute |
| Modify a user's city/province | WP on the user object to modify Locality-Name attribute |
| Modify a user's state | WP on the user object to modify State-Or-Province-Name attribute |
| Modify a user's zip/postal code | WP on the user object to modify Postal-Code attribute |
| Modify a user's country/region | WP on the user object to modify Country-Name attribute |
| Modify a user's UPN | WP on the user object to modify User-Principal-Name attribute |
| Modify a user's Pre-Windows 2000 user logon name | WP on the user object to modify SAM-Account-Name attribute |
| Modify the hours during which a user can log on | WP on the user object to modify Logon-Hours attribute |
| Specify the computers from which a user can log on | WP on the user object to modify User-Workstations attribute |

# Print Queue Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Create a print-queue | CC on parent object (to create objects of class Print-Queue) |
| Delete a print-queue | SD on the object itself OR DC on parent object (to delete objects of class Print-Queue)<br>**Note**<br>DC on parent will grant permission to delete all objects under the parent (and if class is specified, then only all objects of specified class). |
| Rename a print-queue | WP on the print queue object to modify Common-Name attribute<br>WP on the print queue object to modify RDN attribute |
| Move a print-queue | DC on current parent (to delete objects of class Print-Queue)<br>CC on target parent (to create objects of class Print-Queue)<br>WP on the print queue object to modify Common-Name attribute<br>WP on the print queue object to modify RDN attribute |
| Specify the display name of an attached printer | WP on the print queue object to modify Printer-Name attribute |
| Specify the server name for a print server | WP on the print queue object to modify Server-Name attribute |
| Specify the Pre-Windows 2000 compatible server name for print servers | WP on the print queue object to modify Short-Server-Name attribute |

# Service Connection Points Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
|---|---|
| Create a service-specific container in the System container | CC on parent object (to create objects of class Container) |
| Publish service-related objects in the System container | CC on parent objects |
| Create a connection-point object | CC on parent object (to create objects of class Service-Connection-Point) |
| Specify the version of the schema an application was based on | WP on the service connection point object to modify App-Schema-Version attribute |
| Specify service-specific binding information for a service | WP on the service connection point object to modify Service-Binding-Information attribute |
| Specify the string name of the service that an administration point represents | WP on the service connection point object to modify Service-Class-Name attribute |
| Specify the type of DNS Record that an application would lookup for a service | WP on the service connection point object to modify Service-DNS-Name-Type attribute |
| Specify an application or other vendor name | WP on the service connection point object to modify Vendor attribute |
| Specify a general purpose version number for a service | WP on the service connection point object to modify Version-Number attribute |
| Specify a general purpose Major version number for a service | WP on the service connection point object to modify Version-Number-Hi attribute |
| Specify a general purpose Minor version number for a service | WP on the service connection point object to modify Version-Number-Lo attribute |
| Create a service-specific container in the System container | CC on parent object (to create objects of class Container) |

# Generic Data Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| Task | Permissions Required to Perform Task |
| --- | --- |
| Add an object with a specific GUID | The Add-GUID extended right is required on the Domain |

# Group Policy Management Tasks

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

All permissions specified here are based on making modifications by using the Group Policy Management Console (GPMC). To download GPMC, see the Microsoft Download Center at http://www.microsoft.com/downloads/details.aspx?familyid=0a6d4c24-8cbd-4b35-9272-dd3cbfc81887.

GPMC simplifies the management of Group Policy by making it easier to understand, deploy, manage, and troubleshoot Group Policy implementations. GPMC runs on computers running the Microsoft® Windows® XP Professional operating system with Service Pack 1 (SP1) and Windows Server 2003 and can manage Group Policy in either Windows 2000 or Windows Server 2003 domains.

| **Task** | **Permissions Required to Perform Task** |
|---|---|
| Associate a Group Policy with a Site | WP on the corresponding site object, cn=<Site>, cn=Sites, cn=Configuration, dc=<forestRootDomain> to modify the GP-Link attribute<br><br>WP on the corresponding site object, cn=<Site>, cn=Sites, cn=Configuration, dc=<forestRootDomain> to modify the GP-Options attribute |
| Modify Site Group Policy Options | WP on the corresponding site object, cn=<Site>, cn=Sites, cn=Configuration, dc=<forestRootDomain> to modify the GP-Options attribute |
| Edit a Group Policy object | Requires “Edit Settings” on the GPO |
| Modify security on a Group Policy object | Requires “Edit Settings, Delete, Modify Security” on the GPO |
| Link a GPO to an OU, domain, or site. | Requires “Link GPOs” permission on the specific site, domain, or OU |
| Perform Group Policy Modeling analysis for objects in a domain or OU | Requires “Perform Group Policy Modeling Analyses” permission |
| Perform Group Policy Results analysis for objects in a domain or OU | Requires “Read Group Policy Results data” permission on the specific domain or OU |

# Appendix B: Default Active Directory Security Groups

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix includes the following groups of tables:

- Default Security Groups in Windows Server 2003
- Well-Known Security Principals in Windows Server 2003

# Default Security Groups in Windows Server 2003

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic contains security groups with attributes and values.

**Account Operators**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-548 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | By default, this built-in group has no members. It can create, manage, and delete users and groups in the Users and Computers containers and other organizational units (OUs), but it cannot manage service administrator accounts. In addition, members of this group cannot create computer accounts in the Domain Controllers OU, although they can modify and delete them. Members of this group can log on locally to domain controllers, but they cannot shut down or restart them. As a best practice, do not add members to this group, and do not use it for any delegated administration. |

**Administrator**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-<domain>-500 |
| Type | User |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | N/A |
| Default member of | Administrators, Domain Admins, Enterprise Administrators, Domain Users* |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | Yes |
| Safe to delegate management of this group to non-Service admins? | No |

| | |
|---|---|
| Notes | This account is the first account created during operating system installation. The account cannot be deleted or locked out. It is a member of the Administrators group and cannot be removed from that group. |

* This membership is due to the fact that the Primary Group ID of all user accounts is Domain Users.

**Administrators**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-544 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | Administrator, Domain Admins, Enterprise Admins |
| Default member of | None |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | The Administrators group has built-in capabilities that give its members full control over the system. This group cannot be renamed, deleted or moved. This built-in group controls access to all the domain controllers in its domain, and it can change the membership of all administrative groups. Its own membership can be modified by the default service administrator groups BA and DA in the domain, as well as the EA group. This group has the special privilege to take ownership of any object in the directory or any resource on a domain controller. This account is considered a service administrator because its members have full access to the domain controllers in the domain. |

**Backup Operators**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-551 |
| Type | Builtin Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | Cannot be moved |

| | |
|---|---|
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | Backup Operators can back up and restore all files on a computer, regardless of the permissions that protect those files. Backup Operators also can log on to the computer and shut it down. This group cannot be renamed, deleted or moved. By default, this built-in group has no members, and it can perform backup and restore operations on domain controllers. Its membership can be modified by the default service administrator groups BA and DA in the domain, as well as the EA group. It cannot modify the membership of any administrative groups. While members of this group cannot change server settings or modify the configuration of the directory, they do have the permissions needed to replace files (including operating system files) on the domain controllers. Because of this, they are considered service administrators. |

**Cert Publishers**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-<domain>-517 |
| Type | Domain Local |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | No |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | A group that includes all computers that are running an enterprise certification authority (CA). Cert Publishers are authorized to publish certificates for User objects in Active Directory. |

**Domain Admins**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-<domain>-512 |
| Type | Domain Global |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | Administrator |
| Default member of | Administrators |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | Yes |

| | |
|---|---|
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | A group whose members are authorized to administer the domain. By default, the Domain Admins group is a member of the Administrators group on all computers that have joined a domain, including the domain controllers. Domain Admins is the default owner of any object that is created in the domain's Active Directory by any member of the group. If members of the group create other objects, such as files, the default owner is the Administrators group. This group controls access to all domain controllers in a domain, and it can modify the membership of all administrative accounts in the domain. Its own membership can be modified by the service administrator groups BA and DA in its domain, as well as the EA group. This is a service administrator account because its members have full access to a domain's domain controllers. |

**Domain Computers**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-<domain>-515 |
| Type | Domain Global |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | All computers joined to the domain, excluding domain controllers |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Yes (but not required) |
| Safe to delegate management of this group to non-Service admins? | Yes |
| Notes | The group includes all computers that have joined the domain, excluding domain controllers. |

**Domain Controllers**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-<domain>-516 |
| Type | Domain Global |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | Computer accounts for all Domain Controllers of the domain |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | No |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | The group includes all domain controllers in the domain. New domain controllers are added to this group automatically. |

**Domain Guests**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-<domain>-514 |
| Type | Domain Global |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | Guest |
| Default member of | Guests |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Can be moved out but it is not recommended |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | This group includes the domain's built-in Guest account. |

**Domain Users**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-<domain>-513 |
| Type | Domain Global |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | All user accounts in the domain |
| Default member of | Users |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Yes |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | This group includes all user accounts in a domain. When you create a user account in a domain, it is added to this group automatically. |

**Enterprise Admins**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-<root domain>-519 |
| Type | Universal (if Domain is in Native-Mode) else Domain Global |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | Administrator |
| Default member of | Administrators |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | Yes |

| | |
|---|---|
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | This group exists only in the root domain of an Active Directory forest of domains. It is a universal group if the domain is in native-mode, a global group if the domain is in mixed-mode. The group is authorized to make forest-wide changes in Active Directory, such as adding child domains. By default, the only member of the group is the Administrator account for the forest root domain. This group is automatically added to the Administrators group in every domain in the forest, providing complete access to the configuration of all domain controllers. This group can modify the membership of all administrative groups. Its own membership can be modified only by the default service administrator groups in the root domain. This account is considered a service administrator. |

**Group Policy Creators Owners**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-<domain>-520 |
| Type | Domain Global |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | Administrator |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | No |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | This group that is authorized to create new Group Policy objects in Active Directory. By default, the only member of the group is Administrator. |

**Guest**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-<domain>-501 |
| Type | User |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | None |
| Default member of | Domain Guests, Guests |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Can be moved out but it is not recommended |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | A user account for people who do not have individual accounts. This user account does not require a password. By default, the Guest account is disabled. |

**Guests**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-32-546 |
| Type | Builtin Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | Guest |
| Default member of | Guest, Domain Guests |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | A built-in group. By default, the only member is the Guest account. The Guests group allows occasional or one-time users to log on with limited privileges to a computer's built-in Guest account. |

**Incoming Forest Trust Builders**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-32-557 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | This group cannot be renamed, deleted or moved. |

**KRBTGT**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-<domain>-502 |
| Type | User |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | None |
| Default member of | Domain Users* |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | Can be moved out but it is not recommended |

| | |
|---|---|
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | A service account that is used by the Key Distribution Center (KDC) service. |

* This membership is due to the fact that the Primary Group ID of all user accounts is Domain Users.

**Network Configuration Operators**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-556 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | Yes |
| Notes | This group cannot be renamed, deleted or moved. |

**Performance Log Users**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-559 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | Yes |
| Notes | This account cannot be renamed, deleted or moved. |

**Performance Monitor Users**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-558 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |

| | |
|---|---|
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | Yes |
| Notes | This group cannot be renamed, deleted or moved. |

**Pre-Windows 2000 Compatible Access**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-554 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | Refer to the Notes column |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | In Windows 2000, if you choose the Pre-Windows 2000 Compatible Permissions mode, Everyone (including Anonymous) is a member, and if you choose the Windows 2000-only permissions mode, the membership is empty. In Windows Server 2003, if you choose the Pre-Windows 2000 Compatible Permissions mode, Everyone and Anonymous are members, and if you choose the Windows 2000-only permissions mode, Authenticated Users are members. |

**RAS and IAS Servers**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-<domain>-553 |
| Type | Domain Local |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Yes |
| Safe to delegate management of this group to non-Service admins? | Yes |

| | |
|---|---|
| Notes | By default, this group has no members. Computers that are running the Routing and Remote Access service are added to the group automatically. Members of this group have access to certain properties of User objects, such as Read Account Restrictions, Read Logon Information, and Read Remote Access Information. |

**Remote Desktop Users**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-555 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | Yes |
| Notes | This group cannot be renamed, deleted or moved. |

**Schema Admins**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-<root domain>-518 |
| Type | Universal (if Domain is in Native-Mode) else Domain Global |
| Default container | CN=Users, DC=<domain>, DC=... |
| Default members | Administrator |
| Default member of | None |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | Yes |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | This group exists only in the root domain of an Active Directory forest of domains. It is a universal group if the domain is in native-mode, a global group if the domain is in mixed-mode. The group is authorized to make schema changes in Active Directory. By default, the only member of the group is the Administrator account for the forest root domain. This group has full administrative access to the schema. The membership of this group can be modified by any of the service administrator groups in the root domain. This account is considered a service administrator because its members can modify the schema, which governs the structure and content of the entire directory. |

**Server Operators**

| Attribute | Value |
|---|---|

| | |
|---|---|
| Well-Known SID/RID | S-1-5-32-549 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | Yes |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | This group exists only on domain controllers. By default, the group has no members. Server Operators can log on to a server interactively; create and delete network shares; start and stop services; back up and restore files; format the hard disk of the computer; and shut down the computer. This group cannot be renamed, deleted or moved. By default, this built-in group has no members, and it has access to server configuration options on domain controllers. Its membership is controlled by the service administrator groups BA and DA in the domain, as well as the EA group. It cannot change any administrative group memberships. This is a service administrator account because its members have physical access to domain controllers and they can perform maintenance tasks (such as backup and restore), and they have the ability to change binaries that are installed on the domain controllers. |

**Terminal Server License Servers**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-561 |
| Type | BuiltIn Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | None |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | Yes |
| Notes | This group cannot be renamed, deleted or moved. |

**Users**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-32-545 |
| Type | Builtin Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |

| | |
|---|---|
| Default members | Authenticated Users, Domain Users |
| Default member of | Domain Users* |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | No |
| Notes | After the initial installation of the operating system, the only member is the Authenticated Users group. When a computer joins a domain, the Domain Users group is added to the Users group on the computer. Users can perform tasks such as running applications, using local and network printers, shutting down the computer, and locking the computer. Users can install applications that only they are allowed to use if the installation program of the application supports per-user installation. This group cannot be renamed, deleted or moved. |

*This membership is due to the fact that the Primary Group ID of all user accounts is Domain Users.

**Windows Authorization Access Group**

| **Attribute** | **Value** |
|---|---|
| Well-Known SID/RID | S-1-5-32-560 |
| Type | Builtin Local |
| Default container | CN=BuiltIn, DC=<domain>, DC=... |
| Default members | Enterprise Domain Controllers |
| Default member of | None |
| Protected by ADMINSDHOLDER? | No |
| Safe to move out of default container? | Cannot be moved |
| Safe to delegate management of this group to non-Service admins? | Yes |
| Notes | This group cannot be renamed, deleted or moved. |

# Well-Known Security Principals in Windows Server 2003

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic contains well-known security principals with attributes and values.

**Anonymous Logon**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-7 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A user who has logged on anonymously. |

**Authenticated User**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-11 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A group that includes all users whose identities were authenticated when they logged on. Membership is controlled by the operating system. |

**Batch**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-3 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A group that implicitly includes all users who have logged on through a batch queue facility such as task scheduler jobs. Membership is controlled by the operating system. |

**Creator Group**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-3-1 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A placeholder in an inheritable ACE. When the ACE is inherited, the system replaces this SID with the SID for the primary group of the object's current owner. The primary group is used only by the POSIX subsystem. |

**Creator Owner**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-3-0 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A placeholder in an inheritable access control entry (ACE). When the ACE is inherited, the system replaces this SID with the SID for the object's current owner. |

**Dialup**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-1 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | |

**Digest Authentication**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-64-21 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | |

**Enterprise Domain Controllers**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-9 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A group that includes all domain controllers an Active Directory™ directory service forest of domains. Membership is controlled by the operating system. |

**Everyone**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-1-0 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A group that includes all users, even anonymous users and guests. Membership is controlled by the operating system. |

**Interactive**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-4 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A group that includes all users who have logged on interactively. Membership is controlled by the operating system. |

**Local Service**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-19 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | |

**LocalSystem**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-18 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A service account that is used by the operating system. |

**Network**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-2 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A group that implicitly includes all users who are logged on through a network connection. Membership is controlled by the operating system. |

**Network Service**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-20 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | |

**NTLM Authentication**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-64-10 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | |

**Other Organization**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-1000 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A group that implicitly includes all users who are logged on to the system through a dial-up connection. Membership is controlled by the operating system. |

**Principal Self**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-10 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A placeholder in an ACE on a user, group, or computer object in Active Directory. When you grant permissions to Principal Self, you grant them to the security principal represented by the object. During an access check, the operating system replaces the SID for Principal Self with the SID for the security principal represented by the object. |

**Proxy**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-8 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | SID not used in Windows 2000. |

**Remote Interactive Logon**

| Attribute | Value |
| --- | --- |
| Well-Known SID/RID | S-1-5-14 |
| Object Class | Foreign Security Principal |

| | |
|---|---|
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | |

**Restricted Code**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-12 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | |

**SChannel Authentication**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-64-14 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | |

**Service**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-6 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A group that includes all security principals that have logged on as a service. Membership is controlled by the operating system. |

**Terminal Server User**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-13 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| Description | A group that includes all users who have logged on to a Terminal Services server. Membership is controlled by the operating system. |

**This Organization**

| Attribute | Value |
|---|---|
| Well-Known SID/RID | S-1-5-15 |
| Object Class | Foreign Security Principal |
| Default Location in Active Directory | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |

| | |
|---|---|
| Description | |

# Appendix C: Active Directory Standard Permissions

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following table contains standard permissions for Active Directory.

| Permission | Definition |
|---|---|
| Create Child (CC) | This permission controls the ability to create child objects under an object. The ObjectType member of an ACE can contain a GUID that identifies the type of child object whose creation is controlled. If **ObjectType** does not contain a GUID, the ACE controls the creation of all child objects. |
| Standard Delete (SD) | This permission controls the ability to delete an object. If a user has this permission on an object, these permissions are sufficient to delete the object; Delete Child permissions are not necessary on the parent object. |
| Delete Child (DC) | This permission controls the ability to delete child objects under an object. The ObjectType member of an ACE can contain a GUID that identifies a type of child object whose deletion is controlled. If **ObjectType** does not contain a GUID, the ACE controls the deletion of all child object types. |
| Delete Tree (DT) | This permission controls the ability to delete all child objects of an object, regardless of the permissions of the children. |
| Read-Property (RP) | This permission controls the ability to read the properties of an object. The ObjectType member of an ACE can contain a GUID that identifies a property set or property. If ObjectType does not contain a GUID, the ACE controls the right to read all of the object properties. |
| Write-Property (WP) | This permission controls the ability to write the properties of an object. The ObjectType member of an ACE can contain a GUID that identifies a property set or property. If ObjectType does not contain a GUID, the ACE controls the right to write all of the object properties. |
| Write Owner (WO) | This permission controls the ability to assume ownership of the object. |
| Write DACL (WD) | This permission controls the ability to modify the discretionary access-control list (DACL) in the object's security descriptor. |
| Read Control (RC) | This permission controls the ability to read data from the security descriptor of the object, not including the data in the SACL. Thus, a user granted this right can read the Owner, Primary Group and DACL fields in the object's security descriptor. |
| List Child (LC) | This permission controls the ability to list children of an object. For more information about this right, see Controlling Object Visibility in the MSDN Library on the Web at http://go.microsoft.com/fwlink/?LinkID=19828. |
| List Object (LO) | This permission controls the ability to list a particular object. If the user is not granted such a right, and the user does not have List Child set on the object parent, the object is hidden from the user. This right is ignored if the third character of the dSHeuristics property is '0' or not set. For more information about this right, see Controlling Object Visibility in the MSDN Library on the Web at http://go.microsoft.com/fwlink/?LinkID=19828. |
| Control Access (CR) | This permission controls the ability to perform an operation controlled by an extended access right. The ObjectType member of an ACE can contain a GUID that identifies the extended right. If ObjectType does not contain a GUID, the ACE controls the right to perform all extended right operations associated with the object. For a detailed description of all Active Directory extended rights, see Appendix D: Active Directory Extended Rights. |

| | |
|---|---|
| Validated Write (SW) | This permission controls the ability to perform an operation controlled by a validated write access right. The ObjectType member of an ACE can contain a GUID that identifies the validated write. If ObjectType does not contain a GUID, the ACE controls the rights to perform all valid write operations associated with the object.<br><br>There are three validated writes:<br><br>• **Self-Membership**<br><br>This permission allows a user to update the membership of a group in terms of adding/removing one's own account. Rights GUID: bf9679c0-0de6-11d0-a285-00aa003049e2.<br><br>**Note**<br>If users have the Write Property permission on the Member attribute of the group object, the Self-Membership permission is not necessary; in that case, users can add or remove anyone (including themselves) from the group.<br><br>• **Validated-DNS-Host-Name**<br><br>This permission allows setting a DNS host name attribute (for a computer) that is compliant with the computer name and domain name. Rights GUID: 72e39547-7b18-11d1-adef-00c04fd8d5cd.<br><br>• **Validated-SPN**<br><br>This permission allows setting a SPN attribute (for a computer) which is compliant to the DNS host name of the computer. Rights GUID: f3a64788-5306-11d1-a9c5-0000f80367c1. |

# Appendix D: Active Directory Extended Rights

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix contains all of the extended rights defined in the Windows 2000 Server and Windows Server 2003 Active Directory schemas. The following table defines the terms used in the tables that describe these rights.

| Item | Description |
|---|---|
| CN (Common Name) | Every object in the DS has a *naming attribute* from which its relative distinguished name is formed. The naming attribute for control-Access-Right objects is Common-Name. |
| Display-Name | The Common-Name of an object might not be descriptive enough for some users; Display-Name provides a more descriptive name. |
| Rights-GUID | The unique ID for identifying a control access right. |
| Applies-to | A list of GUIDs that represent the objects to which this property set applies. For example, e5209ca2-3bba-11d2-90cc-00c04fd91ab1 refers to PKI-Certificate-Template objects. |

**Extended Rights Defined in the Windows 2000 Active Directory Schema**

The tables in this section contain extended rights defined in the Windows 2000 Active Directory schema.

**Abandon-Replication***

| Item | Description |
|---|---|
| Description | Extended right needed to cancel a replication sync. |
| CN | Abandon-Replication |
| Display-Name | Abandon Replication |
| Rights-GUID | ee914b82-0a98-11d1-adbb-00c04fd8d5cd |
| Applies-To | NTDS-DSA |

* This right is not used.

**Add-GUID**

| Item | Description |
|---|---|
| Description | Extended right needed at the NC root to add an object with a specific GUID. |
| CN | Add-GUID |
| Display-Name | Add GUID |
| Rights-GUID | 440820ad-65b4-11d1-a3da-0000f875ae0d |
| Applies-To | Domain-DNS |

**Allocate-RIDs**

| Item | Description |
|---|---|
| Description | Extended right needed to request a pool of RIDs. |
| CN | Allocate-RIDs |

| | |
|---|---|
| Display-Name | Allocate RIDs |
| Rights-GUID | 1abd7cf8-0a99-11d1-adbb-00c04fd8d5cd |
| Applies-To | NTDS-DSA |

**Apply-Group-Policy**

| Item | Description |
|---|---|
| Description | Extended right used by Group Policy engine to determine whether a GPO applies to a particular user or computer. |
| CN | Apply-Group-Policy |
| Display-Name | Apply Group Policy |
| Rights-GUID | edacfd8f-ffb3-11d1-b41d-00a0c968f939 |
| Applies-To | Group-Policy-Container |

**Certificate-Enrollment**

| Item | Description |
|---|---|
| Description | Extended right needed to cause certificate enrollment. |
| CN | Certificate-Enrollment |
| Display-Name | Check Stale Phantoms |
| Rights-GUID | 69ae6200-7f46-11d2-b9ad-00c04f79f805 |
| Applies-To | PKI-Certificate-Template |

**Change-Domain-Master**

| Item | Description |
|---|---|
| Description | Extended right needed to change the Domain Master role owner. |
| CN | Change-Domain-Master |
| Display-Name | Change Domain Master |
| Rights-GUID | 014bf69c-7b3b-11d1-85f6-08002be74fab |
| Applies-To | Cross-Ref-Container |

**Change-Infrastructure-Master**

| Item | Description |
|---|---|
| Description | Extended right needed to change the Infrastructure FSMO role owner. |
| CN | Change-Infrastructure-Master |
| Display-Name | Change Infrastructure Master |
| Rights-GUID | cc17b1fb-33d9-11d2-97d4-00c04fd8d5cd |
| Applies-To | Infrastructure-Update |

**Change-PDC**

| Item | Description |
|---|---|
| Description | Extended right needed to change the PDC Emulator role owner. |

| | |
|---|---|
| CN | Change-PDC |
| Display-Name | Change PDC |
| Rights-GUID | bae50096-4752-11d1-9052-00c04fc2d4cf |
| Applies-To | Domain-DNS |

**Change-RID-Master**

| Item | Description |
|---|---|
| Description | Extended right needed to change the RID-Master role owner. |
| CN | Change-RID-Master |
| Display-Name | Change RID Master |
| Rights-GUID | d58d5f36-0a98-11d1-adbb-00c04fd8d5cd |
| Applies-To | RID-Manager |

**Change-Schema-Master**

| Item | Description |
|---|---|
| Description | Extended right needed to change the Schema Master FSMO role owner. |
| CN | Change-Schema-Master |
| Display-Name | Change Schema Master |
| Rights-GUID | e12b56b6-0a95-11d1-adbb-00c04fd8d5cd |
| Applies-To | DMD |

**Do-Garbage-Collection**

| Item | Description |
|---|---|
| Description | Extended right to force the directory service to do garbage collection. |
| CN | Do-Garbage-Collection |
| Display-Name | Do Garbage Collection |
| Rights-GUID | fec364e0-0a98-11d1-adbb-00c04fd8d5cd |
| Applies-To | NTDS-DSA |

**Domain-Administer-Server**

| Item | Description |
|---|---|
| Description | Legacy SAM right. |
| CN | Domain-Administer-Server |
| Display-Name | Domain Administer Server |
| Rights-GUID | ab721a52-1e2f-11d0-9819-00aa0040529b |
| Applies-To | Sam-Server |

**DS-Check-Stale-Phantoms**

| Item | Description |
|---|---|
| Description | Extended right needed to force the directory service to check stale phantom objects. |

| | |
|---|---|
| CN | DS-Check-Stale-Phantoms |
| Display-Name | Check Stale Phantoms |
| Rights-GUID | 69ae6200-7f46-11d2-b9ad-00c04f79f805 |
| Applies-To | NTDS-DSA |

**DS-Install-Replica**

| Item | Description |
|---|---|
| Description | Extended right needed to do a replica install. |
| CN | DS-Install-Replica |
| Display-Name | Add/Remove Replica In Domain |
| Rights-GUID | 9923a32a-3607-11d2-b9be-0000f87a36b2 |
| Applies-To | Domain-DNS |

**DS-Replication-Get-Changes**

| Item | Description |
|---|---|
| Description | Extended right needed to replicate changes from a given NC.<br>**Note**<br>This extended right was modified in Windows Server 2003 – refer to the Windows Server 2003 Active Directory Schema Extended Rights section for details on how this right works in Windows Server 2003. |
| CN | DS-Replication-Get-Changes |
| Display-Name | Replicating Directory Changes |
| Rights-GUID | 1131f6aa-9c07-11d1-f79f-00c04fc2dcd2 |
| Applies-To | Configuration<br>DMD<br>Domain-DNS |

**DS-Replication-Manage-Topology**

| Item | Description |
|---|---|
| Description | Extended right needed to update the replication topology for a given NC. |
| CN | DS-Replication-Manage-Topology |
| Display-Name | Manage Replication Topology |
| Rights-GUID | 1131f6ac-9c07-11d1-f79f-00c04fc2dcd2 |
| Applies-To | Configuration<br>DMD<br>Domain-DNS |

**DS-Replication-Synchronize**

| Item | Description |
|---|---|

| | |
|---|---|
| Description | Extended right needed to synchronize replication from a given NC. |
| CN | DS-Replication-Synchronize |
| Display-Name | Replication Synchronization |
| Rights-GUID | 1131f6ab-9c07-11d1-f79f-00c04fc2dcd2 |
| Applies-To | Configuration<br>DMD<br>Domain-DNS |

**msmq-Open-Connector**

| Item | Description |
|---|---|
| Description | Allows opening the connector queue. |
| CN | msmq-Open-Connector |
| Display-Name | Open Connector Queue |
| Rights-GUID | b4e60130-df3f-11d1-9c86-006008764d0e |
| Applies-To | Site |

**msmq-Peek**

| Item | Description |
|---|---|
| Description | Allows peeking at messages in the queue. |
| CN | msmq-Peek |
| Display-Name | Peek Message |
| Rights-GUID | 06bd3201-df3e-11d1-9c86-006008764d0e |
| Applies-To | MSMQ-Queue |

**msmq-Peek-computer-Journal**

| Item | Description |
|---|---|
| Description | Allows peeking at messages in the Computer Journal queue. |
| CN | Msmq-Peek-computer-Journal |
| Display-Name | Peek Computer Journal |
| Rights-GUID | 4b6e08c3-df3c-11d1-9c86-006008764d0e |
| Applies-To | MSMQ-Configuration |

**msmq-Peek-Dead-Letter**

| Item | Description |
|---|---|
| Description | Allows peeking at messages in the Dead Letter queue. |
| CN | Msmq-Peek-Dead-Letter |
| Display-Name | Peek Dead Letter |
| Rights-GUID | 4b6e08c1-df3c-11d1-9c86-006008764d0e |

| | |
|---|---|
| Applies-To | MSMQ-Configuration |

**msmq-Receive**

| Item | Description |
|---|---|
| Description | Allows receiving messages from the queue. |
| CN | msmq-Receive |
| Display-Name | Receive Message |
| Rights-GUID | 06bd3200-df3e-11d1-9c86-006008764d0e |
| Applies-To | MSMQ-Queue |

**msmq-Receive-computer-Journal**

| Item | Description |
|---|---|
| Description | Allows receiving messages from the Computer Journal queue. |
| CN | Msmq-Receive-computer-Journal |
| Display-Name | Receive Computer Journal |
| Rights-GUID | 4b6e08c2-df3c-11d1-9c86-006008764d0e |
| Applies-To | MSMQ-Configuration |

**msmq-Receive-Dead-Letter**

| Item | Description |
|---|---|
| Description | Allows receiving messages from the Dead Letter queue. |
| CN | Msmq-Receive-Dead-Letter |
| Display-Name | Receive Dead Letter |
| Rights-GUID | 4b6e08c0-df3c-11d1-9c86-006008764d0e |
| Applies-To | MSMQ-Configuration |

**msmq-Receive-journal**

| Item | Description |
|---|---|
| Description | Allows receiving messages from the queue’s Journal. |
| CN | msmq-Receive-journal |
| Display-Name | Receive Journal |
| Rights-GUID | 06bd3203-df3e-11d1-9c86-006008764d0e |
| Applies-To | MSMQ-Queue |

**msmq-Send**

| Item | Description |
|---|---|
| Description | Allows sending messages to the queue. |
| CN | msmq-Send |
| Display-Name | Send Message |
| Rights-GUID | 06bd3202-df3e-11d1-9c86-006008764d0e |

| | |
|---|---|
| Applies-To | MSMQ-Queue |

**Open-Address-Book**

| Item | Description |
|---|---|
| Description | Extended right checked when opening address book object for address book views. |
| CN | Open-Address-Book |
| Display-Name | Open Address List |
| Rights-GUID | a1990816-4298-11d1-ade2-00c04fd8d5cd |
| Applies-To | Address-Book-Container |

**Recalculate-Hierarchy**

| Item | Description |
|---|---|
| Description | Extended right to force the DS to recalculate the hierarchy. |
| CN | Recalculate-Hierarchy |
| Display-Name | Recalculate Hierarchy |
| Rights-GUID | 0bc1554e-0a99-11d1-adbb-00c04fd8d5cd |
| Applies-To | NTDS-DSA |

**Recalculate-Security-Inheritance**

| Item | Description |
|---|---|
| Description | Extended right needed to force DS to recompute ACL inheritance on a naming context. |
| CN | Recalculate-Security-Inheritance |
| Display-Name | Recalculate Security Inheritance |
| Rights-GUID | 62dd28a8-7f46-11d2-b9ad-00c04f79f805 |
| Applies-To | NTDS-DSA |

**Receive-As**

| Item | Description |
|---|---|
| Description | Exchange right: allows receiving mail as a given mailbox. |
| CN | Receive-As |
| Display-Name | Receive As |
| Rights-GUID | ab721a56-1e2f-11d0-9819-00aa0040529b |
| Applies-To | Computer<br>User |

**Send-As**

| Item | Description |
|---|---|
| Description | Exchange right: allows sending mail as the mailbox. |
| CN | Send-As |
| Display-Name | Send As |

| | |
|---|---|
| Rights-GUID | ab721a54-1e2f-11d0-9819-00aa0040529b |
| Applies-To | Computer<br><br>User |

**Send-To**

| Item | Description |
|---|---|
| Description | Exchange right: allows sending to a mailbox. |
| CN | Send-To |
| Display-Name | Send To |
| Rights-GUID | ab721a55-1e2f-11d0-9819-00aa0040529b |
| Applies-To | Group |

**Update-Schema-Cache**

| Item | Description |
|---|---|
| Description | Extended right to force a schema cache update. |
| CN | Update-Schema-Cache |
| Display-Name | Update Schema Cache |
| Rights-GUID | be2bb760-7f46-11d2-b9ad-00c04f79f805 |
| Applies-To | DMD |

**User-Change-Password**

| Item | Description |
|---|---|
| Description | Permits changing password on user account. |
| CN | User-Change-Password |
| Display-Name | Change Password |
| Rights-GUID | ab721a53-1e2f-11d0-9819-00aa0040529b |
| Applies-To | Computer<br><br>User |

**User-Force-Change-Password**

| Item | Description |
|---|---|
| Description | Permits resetting password on user account. |
| CN | User-Force-Change-Password |
| Display-Name | Reset Password |
| Rights-GUID | 00299570-246d-11d0-a768-00aa006e0529 |
| Applies-To | Computer<br><br>User |

**Windows Server 2003 Active Directory Schema Extended Rights**
**Allowed-To-Authenticate**

| Item | Description |
| --- | --- |
| Description | This extended right controls who can authenticate to a particular machine or service. It is applied on computer, user and InetOrgPerson objects. It is also applicable on the domain object if access is allowed for the entire domain. It can be applied to OUs to permit users to be able to set inheritable ACEs on OUs containing a set of user/computer objects. |
| CN | Allowed-To-Authenticate |
| Display-Name | Allowed to Authenticate |
| Rights-GUID | 68B1D179-0D15-4d4f-AB71-46152E79A7BC |
| Applies-To | Computer<br>inetOrgPerson<br>User |

**Create-Inbound-Forest-Trust**

| Item | Description |
| --- | --- |
| Description | Extended right that enables users to create an inbound-only trust between forests by adding them to the appropriate group. |
| CN | Create-Inbound-Forest-Trust |
| Display-Name | Create Inbound Forest Trust |
| Rights-GUID | e2a36dc9-ae17-47c3-b58b-be34c55ba633 |
| Applies-To | Domain-DNS |

**DS-Execute-Intentions-Script**

| Item | Description |
| --- | --- |
| Description | Extended right, which should be granted to the partitions container, that allows the Rendom.exe or prepare operation to be used in a domain rename. This control access right also appears as an audit-only right when the Redom.exe or execute step operations are performed. |
| CN | DS-Execute-Intentions-Script |
| Display-Name | Execute Forest Update Script |
| Rights-GUID | 2f16c4a5-b98e-432c-952a-cb388ba33f2e |
| Applies-To | Cross-Ref-Container |

**DS-Query-Self-Quota**

| Item | Description |
| --- | --- |
| Description | Control access right which allows a user to query the user's own quotas. |
| CN | DS-Query-Self-Quota |
| Display-Name | Query Self Quota |
| Rights-GUID | 4ecc03fe-ffc0-4947-b630-eb672a8a9dbc |
| Applies-To | ms-DS-Quota-Container |

**DS-Replication-Get-Changes**

| Item | Description |
|---|---|
| Description | Extended right needed to replicate only those changes from a given NC that are also replicated to the Global Catalog (which excludes secret domain data). This constraint is only meaningful for Domain NCs.<br><br>**Note**<br>In Windows 2000, this extended right allowed the replication of all data including secret data. In Windows Server 2003, this right does not allow the replication of secret data. To replicate all data including secret data, the DS-Replication-Get-Changes-All extended right is required (this new extended right was introduced in Windows Server 2003). |
| CN | DS-Replication-Get-Changes |
| Display-Name | Replicating Directory Changes |
| Rights-GUID | 1131f6aa-9c07-11d1-f79f-00c04fc2dcd2 |
| Applies-To | Configuration<br><br>DMD<br><br>Domain-DNS |

**DS-Replication-Get-Changes-All**

| Item | Description |
|---|---|
| Description | Control access right that allows the replication of all data in a given replication NC, including secret domain data.<br><br>**Note**<br>In Windows Server 2003, the Get Replication Changes Extended right has been modified such that it only allows the replication of all data in a given replication NC that is also replicated to the Global Catalog (which excludes secret domain data). |
| CN | DS-Replication-Get-Changes-All |
| Display-Name | Replicating Directory Changes All |
| Rights-GUID | 1131f6ad-9c07-11d1-f79f-00c04fc2dcd2 |
| Applies-To | Configuration<br><br>DMD<br><br>Domain-DNS |

**DS-Replication-Monitor-Topology**

| Item | Description |
|---|---|
| Description | Extended control access right that allows the reading of replication monitoring data, such as replication status and object metadata. |
| CN | DS-Replication-Monitor-Topology |
| Display-Name | Monitor Active Directory Replication |

| | |
|---|---|
| Rights-GUID | f98340fb-7c5b-4cdb-a00b-2ebdfa115a96 |
| Applies-To | Configuration<br><br>DMD<br><br>Domain-DNS |

**Enable-Per-User-Reversibly-Encrypted-Password**

| Item | Description |
|---|---|
| Description | Extended right that allows users to enable or disable the “reversible encrypted password” setting for user and computer objects. |
| CN | Enable-Per-User-Reversibly-Encrypted-Password |
| Display-Name | Enable Per User Reversibly Encrypted Password |
| Rights-GUID | 05c74c5e-4deb-43b4-bd9f-86664c2a7fd5 |
| Applies-To | Domain-DNS |

**Generate-RSoP-Logging**

| Item | Description |
|---|---|
| Description | The user who has this right on an OU/Domain will be able to generate logging mode RSoP data for the users/computers within the OU. |
| CN | Generate-RSoP-Logging |
| Display-Name | Generate Resultant Set of Policy (Logging) |
| Rights-GUID | b7b1b3de-ab09-4242-9e30-9980e5d322f7 |
| Applies-To | Domain-DNS<br><br>Organizational-Unit |

**Generate-RSoP-Planning**

| Item | Description |
|---|---|
| Description | The user who has this right on an OU/Domain will be able to generate planning mode RSoP data for the users/computers within the OU. |
| CN | Generate-RSoP-Planning |
| Display-Name | Generate Resultant Set of Policy (Planning) |
| Rights-GUID | b7b1b3dd-ab09-4242-9e30-9980e5d322f7 |
| Applies-To | Domain-DNS<br><br>Organizational-Unit |

**Migrate-SID-History**

| Item | Description |
|---|---|
| Description | Extended right that enables a user to migrate the SID-History without administrator privileges. |

| | |
|---|---|
| CN | Migrate-SID-History |
| Display-Name | Migrate SID History |
| Rights-GUID | BA33815A-4F93-4c76-87F3-57574BFF8109 |
| Applies-To | Domain-DNS |

**Reanimate-Tombstones**

| Item | Description |
|---|---|
| Description | Extended right that allows deleted schema elements to be restored. |
| CN | Reanimate-Tombstones |
| Display-Name | Reanimate Tombstones |
| Rights-GUID | 45EC5156-DB7E-47bb-B53F-DBEB2D03C40F |
| Applies-To | Configuration<br>DMD<br>Domain-DNS |

**Refresh-Group-Cache**

| Item | Description |
|---|---|
| Description | For Universal group membership caching. Universal group membership caching relies on caching group memberships and this control access right is used to provide administrators/operators with rights to cause an immediate refresh of the cache, contacting an available global catalog server. |
| CN | Refresh-Group-Cache |
| Display-Name | Refresh Group Cache for Logons |
| Rights-GUID | 9432c620-033c-4db7-8b58-14ef6d0bf477 |
| Applies-To | NTDS-DSA |
| Valid-Accesses | 0x100 |
| Localization-Display-ID | 56 |

**SAM-Enumerate-Entire-Domain**

| Item | Description |
|---|---|
| Description | This extended right is used to restrict who can be allowed to use down-level APIs such as NetQueryDisplayInformation and NetUser/GroupEnum and enumerate the entire domain. |
| CN | SAM-Enumerate-Entire-Domain |
| Display-Name | Enumerate Entire SAM Domain |
| Rights-GUID | 91d67418-0135-4acc-8d79-c08e857cfbec |
| Applies-To | Sam-Server |

**Unexpire-Password**

| Item | Description |
| --- | --- |
| Description | Extended right that allows a user to restore an expired password for a user object. |
| CN | Unexpire-Password |
| Display-Name | Unexpire Password |
| Rights-GUID | ccc2dc7d-a6ad-4a7a-8846-c04e3cc53501 |
| Applies-To | Domain-DNS |

**Update-Password-Not-Required-Bit**

| Item | Description |
| --- | --- |
| Description | Extended right that allows a user to enable or disable the “password not required” setting for user objects. |
| CN | Update-Password-Not-Required-Bit |
| Display-Name | Update Password Not Required Bit |
| Rights-GUID | 280f369c-67c7-438e-ae98-1d46f3c6f541 |
| Applies-To | Domain-DNS |

# Appendix E: Active Directory Property Sets

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section contains all of the property sets defined in the Active Directory Schema.

| **Note** |
| --- |
| For descriptions of the individual member properties of each property set refer to the Active Directory schema. Follow the MSDN Library link on the Web Resources page at http://msdn2.microsoft.com/en-us/library/ms683990.aspx. To locate the Active Directory Schema section in the Platform Software Development Kit (SDK), click Networking and Directory Services. |

The following property sets exist in the Windows 2000 Active Directory schema:

- Domain-Password Information
- Email-Information
- General-Information
- Membership
- Personal-Information
- Public-Information
- RAS-Information
- User-Account-Restrictions
- User-Logon
- Web-Information

The following property sets were added in the Windows Server 2003 Active Directory schema:

- DNS-Host-Name-Attributes
- Domain-Other-Parameters

The following table describes the fields used to describe the property-sets in this appendix.

| Term | Description |
| --- | --- |
| Description | A description of the property set. |
| CN (Name) | Every object in Active Directory has a naming attribute from which its relative distinguished name is formed. The Naming Attribute for control Access Right objects is Common Name. |
| Display-Name | The Common Name of an object might not be descriptive enough for some users. Display Name is provided as a more descriptive name. |
| Rights-GUID | The unique ID for identifying a control access right. |

| | |
|---|---|
| Applies-To | Provides the name of the classes that this property set applies to. Includes the Schema ID GUID of the class in parentheses. For example, the PKI-Certificate-Template class (Schema ID GUID e5209ca2-3bba-11d2-90cc-00c04fd91ab1 ). |
| Property Set Members | This lists the set of attributes that belong to this property set. |

**Domain Password Information**
The following table describes the Domain Password property sets.

| Term | Description |
|---|---|
| Description | Property set containing all lockout and password age related attributes on user account |
| CN | Domain-Password |
| Display-Name | Domain Password & Lockout Policies |
| Rights-GUID | c7407360-20bf-11d0-a768-00aa006e0529 |
| Applies-To | Domain (Schema ID GUID: 19195a5a-6da0-11d0-afd3-00c04fd930c9)<br>Domain-DNS (Schema ID GUID: 19195a5b-6da0-11d0-afd3-00c04fd930c9) |
| Property Set Members | Lock-Out-Observation-Window<br>Lockout-Duration<br>Lockout-Threshold<br>Max-Pwd-Age<br>Min-Pwd-Age<br>Min-Pwd-Length<br>Pwd-History-Length<br>Pwd-Properties |

**Email Information**
The following table describes the Email Information property sets.

| Term | Description |
|---|---|
| Description | Property set containing user attributes that describe user e-mail related information |
| CN | Email-Information |
| Display-Name | Phone and Mail Options |
| Rights-GUID | E45795B2-9455-11d1-AEBD-0000F80367C1 |
| Applies-To | Group (Schema ID GUID: bf967a9c-0de6-11d0-a285-00aa003049e2)<br>inetOrgPerson (Schema ID GUID: 4828CC14-1437-45bc-9B07-AD6F015E5F28)<br>User (Schema ID GUID: bf967aba-0de6-11d0-a285-00aa003049e2) |
| Property Set Members | This property does not have any members |

**General-Information**
The following table describes the General Information property sets.

| Term | Description |
|---|---|
| Description | Property set containing a set of user attributes that constitute general user information |
| CN | General-Information |

| | |
|---|---|
| Display-Name | General Information |
| Rights-GUID | 59ba2f42-79a2-11d0-9020-00c04fc2d3cf |
| Applies-To | inetOrgPerson (Schema ID GUID: 4828CC14-1437-45bc-9B07-AD6F015E5F28)<br>User (Schema ID GUID: bf967aba-0de6-11d0-a285-00aa003049e2) |
| Property Set Members | Admin-Description<br>Code-Page<br>Country-Code<br>Display-Name<br>Object-Sid<br>Primary-Group-ID<br>SAM-Account-Name<br>SAM-Account-Type<br>SD-Rights-Effective<br>Show-In-Advanced-View-Only<br>SID-History<br>uid<br>User-Comment |

**Membership**

The following table describes the Membership property sets.

| Term | Description |
|---|---|
| Description | Property set containing user attributes that describe group membership information |
| CN | Membership |
| Display-Name | Group Membership |
| Rights-GUID | bc0ac240-79a9-11d0-9020-00c04fc2d4cf |
| Applies-To | inetOrgPerson (Schema ID GUID: 4828CC14-1437-45bc-9B07-AD6F015E5F28)<br>User (Schema ID GUID: bf967aba-0de6-11d0-a285-00aa003049e2) |
| Property Set Members | Is-Member-Of-DL<br>Member |

**Personal Information**

The following table describes the Personal Information property sets.

| Term | Description |
|---|---|
| Description | Property set containing user attributes that describe personal user information. |
| CN | Personal-Information |
| Display-Name | Personal Information |

| | |
|---|---|
| Rights-GUID | 77B5B886-944A-11d1-AEBD-0000F80367C1 |
| Applies-To | Computer (Schema ID GUID: bf967a86-0de6-11d0-a285-00aa003049e2)<br>Contact (Schema ID GUID: 5cb41ed0-0e4c-11d0-a286-00aa003049e2)<br>inetOrgPerson (Schema ID GUID: 4828CC14-1437-45bc-9B07-AD6F015E5F28)<br>User (Schema ID GUID: bf967aba-0de6-11d0-a285-00aa003049e2) |
| Property Set Members | Address<br>Address-Home<br>Assistant<br>Comment<br>Country-Name<br>Facsimile-Telephone-Number<br>International-ISDN-Number<br>Locality-Name<br>MSMQ-Digests<br>MSMQ-Sign-Certificates<br>Personal-Title<br>Phone-Fax-Other<br>Phone-Home-Other<br>Phone-Home-Primary<br>Phone-Ip-Other<br>Phone-Ip-Primary<br>Phone-ISDN-Primary<br>Phone-Mobile-Other<br>Phone-Mobile-Primary<br>Phone-Office-Other<br>Phone-Pager-Other<br>Phone-Pager-Primary<br>Physical-Delivery-Office-Name<br>Picture<br>Post-Office-Box<br>Postal-Address<br>Postal-Code<br>Preferred-Delivery-Method<br>Registered-Address<br>State-Or-Province-Name<br>Street-Address<br>Telephone-Number |

| | |
|---|---|
| | Teletex-Terminal-Identifier<br>Telex-Number<br>Telex-Primary<br>User-Cert<br>User-Shared-Folder<br>User-Shared-Folder-Other<br>User-SMIME-Certificate<br>X121-Address<br>X509-Cert |

**Public-Information**

The following table describes the Personal Information property sets.

| Term | Description |
|---|---|
| Description | Property set containing user attributes that describe user public information |
| CN | Public-Information |
| Display-Name | Public Information |
| Rights-GUID | e48d0154-bcf8-11d1-8702-00c04fb96050 |
| Applies-To | Computer (Schema ID GUID: bf967a86-0de6-11d0-a285-00aa003049e2)<br>inetOrgPerson (Schema ID GUID: 4828CC14-1437-45bc-9B07-AD6F015E5F28)<br>User (Schema ID GUID: bf967aba-0de6-11d0-a285-00aa003049e2) |
| Property Set Members | Additional-Information<br>Allowed-Attributes<br>Allowed-Attributes-Effective<br>Allowed-Child-Classes<br>Allowed-Child-Classes-Effective<br>Alt-Security-Identities<br>Common-Name<br>Company<br>Department<br>Description<br>Display-Name-Printable<br>Division<br>E-mail-Addresses<br>Given-Name<br>Initials<br>Legacy-Exchange-DN<br>Manager |

| | |
|---|---|
| | ms-DS-Allowed-To-Delegate-To<br>ms-DS-Approx-Immed-Subordinates<br>ms-DS-Auxiliary-Classes<br>Obj-Dist-Name<br>Object-Category<br>Object-Class<br>Object-Guid<br>Organization-Name<br>Organizational-Unit-Name<br>Other-Mailbox<br>Proxy-Addresses<br>RDN<br>Reports<br>Service-Principal-Name<br>Show-In-Address-Book<br>Surname<br>System-Flags<br>Text-Country<br>Title<br>User-Principal-Name |

**RAS-Information**

The following table describes the Personal Information property sets.

| Term | Description |
|---|---|
| Description | Property set containing RAS specific attributes |
| CN | RAS-Information |
| Display-Name | Remote Access Information |
| Rights-GUID | 037088f8-0ae1-11d2-b422-00a0c968f939 |
| Applies-To | inetOrgPerson (Schema ID GUID: 4828CC14-1437-45bc-9B07-AD6F015E5F28)<br>User (Schema ID GUID: bf967aba-0de6-11d0-a285-00aa003049e2) |

| | |
|---|---|
| Property Set Members | msNPAllowDialin<br>msNPCallingStationID<br>msRADIUSCallbackNumber<br>msRADIUSFramedIPAddress<br>msRADIUSFramedRoute<br>msRADIUSServiceType<br>Token-Groups<br>Token-Groups-Global-And-Universal<br>Token-Groups-No-GC-Acceptable |

**User-Account-Restrictions**

The following table describes the Personal Information property sets.

| Term | Description |
|---|---|
| Description | Property set containing user attributes that describe account restrictions. |
| CN | User-Account-Restrictions |
| Display-Name | Account Restrictions |
| Rights-GUID | 4c164200-20c0-11d0-a768-00aa006e0529 |
| Applies-To | Computer (Schema ID GUID: bf967a86-0de6-11d0-a285-00aa003049e2)<br>inetOrgPerson (Schema ID GUID: 4828CC14-1437-45bc-9B07-AD6F015E5F28)<br>User (Schema ID GUID: bf967aba-0de6-11d0-a285-00aa003049e2) |
| Property Set Members | Account-Expires<br>ms-DS-User-Account-Control-Computed<br>Pwd-Last-Set<br>User-Account-Control<br>User-Parameters |

**User-Logon**

The following table describes the Personal Information property sets.

| Term | Description |
|---|---|
| Description | Property set containing user attributes that describe user logon information. |
| CN | User-Logon |
| Display-Name | Logon Information |
| Rights-GUID | 5f202010-79a5-11d0-9020-00c04fc2d4cf |
| Applies-To | inetOrgPerson (Schema ID GUID: 4828CC14-1437-45bc-9B07-AD6F015E5F28)<br>User (Schema ID GUID: bf967aba-0de6-11d0-a285-00aa003049e2) |

| | |
|---|---|
| Property Set Members | Bad-Pwd-Count<br>Home-Directory<br>Home-Drive<br>Last-Logoff<br>Last-Logon<br>Last-Logon-Timestamp<br>Logon-Count<br>Logon-Hours<br>Logon-Workstation<br>Profile-Path<br>Script-Path<br>User-Workstations |

**Web Information**
The following table describes the Personal Information property sets.

| Term | Description |
|---|---|
| Description | Property set containing user attributes that describe user web related information |
| CN | Web Information |
| Display-Name | Web Information |
| Rights-GUID | E45795B3-9455-11d1-AEBD-0000F80367C1 |
| Applies-To | Contact (Schema ID GUID: 5cb41ed0-0e4c-11d0-a286-00aa003049e2)<br>inetOrgPerson (Schema ID GUID: 4828CC14-1437-45bc-9B07-AD6F015E5F28)<br>User (Schema ID GUID: bf967aba-0de6-11d0-a285-00aa003049e2) |
| Property Set Members | WWW-Home-Page<br>WWW-Page-Other |

**DNS-Host-Name-Attributes**
The following table describes the Personal Information property sets.

| Term | Description |
|---|---|
| CN | Contains the DNS-Host-Name and ms-DS-Additional-Dns-Host-Name attributes |
| Display-Name | DNS-Host-Name-Attributes |
| Rights-GUID | DNS Host Name Attributes |
| Applies-To | 72e39547-7b18-11d1-adef-00c04fd8d5cd |
| Property Set Members | Computer (Schema ID GUID: bf967a86-0de6-11d0-a285-00aa003049e2) |
| CN | DNS-Host-Name<br>ms-DS-Additional-DNS-Host-Name |

**Domain-Other-Parameters**
The following table describes the Personal Information property sets.

| Term | Description |
| --- | --- |
| Description | Property set permitting control to a list of domain attributes |
| CN | Domain-Other-Parameters |
| Display-Name | Other Domain Parameters (for use by SAM) |
| Rights-GUID | B8119fd0-04f6-4762-ab7a-4986c76b3f9a |
| Applies-To | Domain-DNS (Schema ID GUID: 19195a5b-6da0-11d0-afd3-00c04fd930c9) |
| Property Set Members | Domain-Replica<br>Force-Logoff<br>Modified-Count<br>OEM-Information<br>Server-Role<br>Server-State<br>UAS-Compat |

# Appendix F: Logon Rights and Privileges

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix consists of tables that list and describe logon rights and privileges.

**Logon Rights**
The default settings listed in this table are for Windows XP Professional and the Windows Server 2003 family.

| Logon Right | Description |
|---|---|
| Access this computer from a network | This user right determines which users and groups are allowed to connect to the computer over the network. Terminal Services are not affected by this user right.<br><br>Default:<br>• On workstations and servers:<br>  • Administrators<br>  • Backup Operators<br>  • Power Users<br>  • Users<br>  • Everyone<br>• On domain controllers:<br>  • Administrators<br>  • Authenticated Users<br>  • Everyone |

| | |
|---|---|
| Allow log on locally | This logon right determines which users can interactively log on to this computer. Logons initiated by pressing CTRL+ALT+DEL on the attached keyboard requires the user to have this logon right. Additionally this logon right might be required by some service or administrative applications that can log on users. If you define this policy for a user or group, you must also give the **Administrators** group this right.<br><br>Default:<br><br>• On workstations and servers:<br>  • Administrators<br>  • Backup Operators<br>  • Power Users<br>  • Users<br>  • Guest<br><br>• On domain controllers:<br>  • Account Operators<br>  • Administrators<br>  • Backup Operators<br>  • Print Operators<br>  • Server Operators |
| Allow log on through Terminal Services | This security setting determines which users or groups have permission to log on as a Terminal Services client.<br><br>Default:<br><br>• On workstation and servers:<br>  • Administrators<br>  • Remote Desktop Users<br><br>• On domain controllers:<br>  • Administrators |
| Deny access to this computer from network | This security setting determines which users are prevented from accessing a computer over the network. This policy setting supersedes the Access this computer from the network policy setting if a user account is subject to both policies.<br><br>Default: No one. |
| Deny log on as a batch job | This security setting determines which accounts are prevented from being able to log on as a batch job. This policy setting supersedes the Log on as a batch job policy setting if a user account is subject to both policies.<br><br>Default: None. |

| | |
|---|---|
| Deny logon as a service | This security setting determines which service accounts are prevented from registering a process as a service. This policy setting supersedes the Log on as a service policy setting if an account is subject to both policies.<br><br>**Note**<br>This security setting does not apply to the System, Local Service, or Network Service accounts.<br><br>Default: None. |
| Deny log on locally | This security setting determines which users are prevented from logging on at the computer. This policy setting supersedes the Allow log on locally policy setting if an account is subject to both policies.<br><br>**Important**<br>If you apply this security policy to the Everyone group, no one will be able to log on locally.<br><br>Default: None. |
| Deny log on through Terminal Services | This security setting determines which users and groups are prohibited from logging on as a Terminal Services client.<br><br>Default: None. |
| Log on as a batch job | This security setting allows a user to be logged on by means of a batch-queue facility.<br><br>For example, when a user submits a job by means of the task scheduler, the task scheduler logs that user on as a batch user rather than as an interactive user.<br><br>**Note**<br>In Windows 2000 Server, Windows 2000 Professional, Windows Server 2003 and Windows XP Professional, the Task Scheduler automatically grants this right as necessary.<br><br>Default: LocalSystem. |
| Log on as a service | This security setting determines which service accounts can register a process as a service.<br><br>Default: None. |

**Privileges**

To ease the task of user account administration, you should assign privileges primarily to group accounts, rather than to individual user accounts. When you assign privileges to a group account, users are automatically assigned those privileges when they become a member of that group. This method of administering privileges is far easier than assigning individual privileges to each user account when the account is created.

Some privileges can override permissions set on an object. For example, a user logged on to a domain account as a member of the Backup Operators group has the right to perform backup operations for all domain servers. However, this requires the ability to read all files on those servers, even files on which their owners have set permissions that explicitly deny access to all users, including members of the Backup Operators group. A user right — in this case, the right to perform a backup — takes precedence over all file and directory permissions.

| **Note** |
|---|
| At a command prompt, you can type whoami /priv to see your privileges. |

The following table lists and describes the privileges that can be granted to a user.

| **Privilege** | **Description** |
|---|---|

| | |
|---|---|
| Act as part of the operating system | This user right allows a process to impersonate any user without authentication. The process can therefore gain access to the same local resources as that user.<br><br>Processes that require this privilege should use the LocalSystem account, which already includes this privilege, rather than using a separate user account with this privilege specially assigned. If your organization only uses servers that are members of the Windows Server 2003 family, you do not need to assign this privilege to your users. However, if your organization uses servers running Windows 2000 or the Microsoft® Windows NT® version 4.0 operating system, you might need to assign this privilege to use applications that exchange passwords in plaintext.<br><br>Default: LocalSystem. |
| Add workstations to a domain | This security setting determines which groups or users can add workstations to a domain.<br><br>This security setting is valid only on domain controllers. By default, any authenticated user has this right and can create up to 10 computer accounts in the domain.<br><br>Adding a computer account to the domain allows the computer to participate in Active Directory–based networking. For example, adding a workstation to a domain enables that workstation to recognize accounts and groups that exist in Active Directory.<br><br>Default: Authenticated Users on domain controllers. |
| Adjust memory quotas for a process | This privilege determines who can change the maximum memory that can be consumed by a process.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>Default: Administrators. |
| Back up files and directories | This user right determines which users can bypass file and directory, registry, and other persistent object permissions for the purposes of backing up the system.<br><br>Default: Administrators and Backup Operators. |
| Bypass traverse checking | This user right determines which users can traverse directory trees even though the user might not have permissions on the traversed directory. This privilege does not allow the user to list the contents of a directory, only to traverse directories.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>Default:<br><br>• On workstations and servers:<br>  • Administrators<br>  • Backup Operators<br>  • Power Users<br>  • Users<br>  • Everyone<br><br>• On domain controllers:<br>  • Administrators<br>  • Authenticated Users |

| | |
|---|---|
| Change the system time | This user right determines which users and groups can change the time and date on the internal clock of the computer. Users that are assigned this user right can affect the appearance of event logs. If the system time is changed, events that are logged will reflect this new time, not the actual time that the events occurred.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>Default:<br><br>• On workstations and servers:<br>  • Administrators<br>  • Power Users<br>• On domain controllers:<br>  • Administrators<br>  • Server Operators |
| Create a pagefile | Allows the user to create and change the size of a pagefile. This is done by specifying a paging file size for a particular drive under **Performance Options** on the **Advanced** tab of **System Properties**.<br><br>Default: Administrators |
| Create a token object | Allows a process to create a token which it can then use to get access to any local resources when the process uses NtCreateToken() or other token-creation APIs.<br><br>It is recommended that processes requiring this privilege use the LocalSystem account, which already includes this privilege, rather than using a separate user account with this privilege specially assigned<br><br>Default: No one |
| Create global objects | This security setting determines which accounts are allowed to create global objects in a terminal services session.<br><br>Default: Administrators and LocalSystem. |
| Create permanent shared objects | Allows a process to create a directory object in the Windows Server 2003 family and Windows XP Professional object manager. This privilege is useful to kernel-mode components that extend the object namespace. Components that are running in kernel mode already have this privilege inherently; it is not necessary to assign them the privilege.<br><br>Default setting: No one |
| Debug programs | This user right determines which users can attach a debugger to any process or to the kernel. Developers who are debugging their own applications to not need to be assigned this user right. Developers who are debugging new system components will need this user right to be able to do so. This user right provides complete access to sensitive and critical operating system components.<br><br>Default :<br><br>• Administrators<br>• LocalSystem |

| | |
|---|---|
| Enable computer and user accounts to be trusted for delegation | This security setting determines which users can set the **Trusted for Delegation** setting on a user or computer object.<br><br>The user or object that is granted this privilege must have write access to the account control flags on the user or computer object. A server process running on a computer (or under a user context) that is trusted for delegation can access resources on another computer using the delegated credentials of a client, as long as the account of the client does not have the **Account cannot be delegated** account control flag set.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>Default on domain controllers:<br><br>• Administrators |
| Force shutdown from a remote system | This security setting determines which users are allowed to shut down a computer from a remote location on the network. Misuse of this user right can result in a denial of service.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>Default:<br><br>• On workstations and servers:<br>  • Administrators<br>• On domain controllers:<br>  • Administrators<br>  • Server Operators |
| Generate security audits | This security setting determines which accounts can be used by a process to add entries to the security log. The security log is used to trace unauthorized system access. Misuse of this user right can result in the generation of many auditing events, potentially hiding evidence of an attack or causing a denial of service if the Audit: Shut down system immediately if unable to log security audits security policy setting is enabled.<br><br>Default: LocalSystem. |
| Impersonate a client after authentication | This security setting determines which accounts are allowed to impersonate other accounts.<br><br>Default: Administrators and Service. |
| Increase scheduling priority | This security setting determines which accounts can use a process with Write Property access to another process to increase the execution priority assigned to the other process. A user with this privilege can change the scheduling priority of a process through the Task Manager user interface.<br><br>Default: Administrators. |
| Load and unload device drivers | This user right determines which users can dynamically load and unload device drivers or other code in to kernel mode. This user right does not apply to Plug and Play device drivers. It is recommended that you do not assign this privilege to other users. Instead, use the StartService() API.<br><br>Default: Administrators. It is recommended that you not assign this privilege to any other user. Device drivers run as trusted (or highly privileged) programs. |
| Lock pages in memory | This security setting determines which accounts can use a process to keep data in physical memory, which prevents the system from paging the data to virtual memory on disk. Exercising this privilege could significantly affect system performance by decreasing the amount of available random access memory (RAM).<br><br>Default: None. Certain system processes have the privilege inherently. |

| | |
|---|---|
| Manage auditing and security log | This security setting determines which users can specify object access auditing options for individual resources, such as files, Active Directory objects, and registry keys.<br><br>This security setting does not allow a user to enable file and object access auditing in general. For such auditing to be enabled, the Audit object access setting in Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policies must be configured.<br><br>You can view audited events in the security log of the Event Viewer. A user with this privilege can also view and clear the security log.<br><br>Default: Administrators. |
| Modify firmware environment values | This security setting determines who can modify firmware environment values. Firmware environment variables are settings stored in the nonvolatile RAM of non-x86-based computers. The effect of the setting depends on the processor.<br><br>On x86-based computers, the only firmware environment value that can be modified by assigning this user right is the Last Known Good Configuration setting, which should only be modified by the system.<br><br>On Itanium-based computers, boot information is stored in nonvolatile RAM. Users must be assigned this user right to run bootcfg.exe and to change the Default Operating System setting on **Startup and Recovery** in **System Properties**.<br><br>On all computers, this user right is required to install or upgrade Windows.<br><br>Default setting:<br><br>• Administrators<br><br>• LocalSystem |
| Profile a single process | This security setting determines which users can use performance monitoring tools to monitor the performance of nonsystem processes.<br><br>Default: Administrators, Power users, LocalSystem. |
| Profile system performance | This security setting determines which users can use performance monitoring tools to monitor the performance of system processes.<br><br>Default:<br><br>• Administrators<br><br>• LocalSystem |
| Remove computer from docking station | This security setting determines whether a user can undock a portable computer from its docking station without logging on.<br><br>If this policy is enabled, the user must log on before removing the portable computer from its docking station. If this policy is disabled, the user might remove the portable computer from its docking station without logging on.<br><br>Default: Disabled. |

| | |
|---|---|
| Replace a process level token | Determines which user accounts can initiate a process to replace the default token associated with a started subprocess.<br><br>This user right is defined in the Default Domain Controller Group Policy object and in the local security policy of workstations and servers.<br><br>Default setting:<br><br>• Local Service<br><br>• Network Service |
| Restore files and directories | This security setting determines which users can bypass file, directory, registry, and other persistent objects permissions when restoring backed up files and directories and determines which users can set any valid security principal as the owner of an object.<br><br>Specifically, this user right is similar to granting the following permissions to the user or group in question on all files and folders on the system:<br><br>• Traverse Folder/Execute File<br><br>• Write<br><br>Default:<br><br>• Workstations and servers:<br><br>  • Administrators<br><br>  • Backup Operators<br><br>• Domain controllers:<br><br>  • Administrators<br><br>  • Backup Operators<br><br>  • Server Operators |

| | |
|---|---|
| Shut down the system | This security setting determines which users who are logged on locally to the computer can shut down the operating system using the Shut Down command. Misuse of this user right can result in a denial of service.<br><br>Default:<br><br>• Workstations:<br>  • Administrators<br>  • Backup Operators<br>  • Power Users<br>  • Users<br>• Servers:<br>  • Administrators<br>  • Backup Operators<br>  • Power Users<br>• Domain controllers:<br>  • Account Operators<br>  • Administrators<br>  • Backup Operators<br>  • Server Operators<br>  • Print Operators |
| Synchronize directory service data | This security setting determines which users and groups have the authority to synchronize all directory service data. This is also known as Active Directory synchronization.<br><br>Defaults: None. |
| Take ownership of files or other objects | This security setting determines which users can take ownership of any securable object in the system, including Active Directory objects, files and folders, printers, registry keys, processes, and threads.<br><br>Default : Administrators |

# Appendix G: Active Directory Delegation Tools

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory ships with all the tools required to delegate administrative tasks. This appendix provides details about the functionality of these tools.

The tools that can be used to delegate administrative authority are:

- Delegation of Control Wizard
- ACL Editor
- Ldp.exe
- Dsacls.exe
- Acldiag.exe
- Dsrevoke.exe

| **Note** |
| --- |
| The screens in this appendix are from Windows 2000, and might differ slightly from the screens in Windows Server 2003. The functionality of the tools, however, is essentially identical. |

**Delegation of Control Wizard**

The Delegation of Control Wizard allows you to delegate administrative tasks to users (or groups) within a specific administrative scope and is primarily used to delegate data administration. This tool is driven by a customizable text file and ships with a base set of common administrative tasks. You can use the information provided earlier in this document in Appendix A: Active Directory Administrative Tasks, to customize this wizard and to increase the number of tasks that can be delegated by using the wizard.

**Using the Delegation of Control Wizard**

The Delegation of Control Wizard can be accessed through Microsoft Management Console (MMC) snap-ins. MMC is an extensible user interface that provides an environment for running management applications (called snap-ins) that are structured as components. Windows 2000 Server and Windows Server 2003 include a base set of MMC snap-ins for managing Active Directory data.

The following MMC snap-ins can be used to view and manage Active Directory objects for which you can delegate control:

- **Active Directory Users and Computers.** This snap-in is the main GUI-based delegation tool for delegating data management. It provides a hierarchical view of domain data. In Active Directory Users and Computers, the Delegation of Control Wizard can be accessed by right-clicking and choosing **Delegate Control** on containers, OUs, and the domain root.
- **Active Directory Sites and Services.** This snap-in provides access to information related to replication, and is used to perform replication management tasks. In Active Directory Sites and Services, the Delegation of Control Wizard can be accessed by right-clicking a container and choosing **Delegate C**.

To delegate administrative authority by using the Delegation Wizard, use the following procedure.

**To delegate administrative authority by using the Delegation Wizard**

1. Right-click a container or OU and select **Delegate Control**. The Delegation of Control Wizard Welcome page is displayed.

2. Click **Next**. The **Users or Groups**page is displayed.

3. On the **Users or Groups** page, click **Add**. The **Select Users, Computers, or Groups** page is displayed.

4. On the **Select Users, Computers, or Groups** page, in the **Enter the object names to select**box, type the name of the user or security group to which you want to delegate tasks. You can add multiple users or security groups. When you are finished entering users or groups, click **OK**.

5. On the **Users or Groups** page, click **Next**. The **Tasks to Delegate** page is displayed:




6. On the **Tasks to Delegate** page, select the check boxes of the tasks that you want to delegate. You can also create a custom task to delegate, as described later in this appendix.

7. Once you have selected the tasks that you want to delegate, click **Next**. The Delegation of Control Wizard displays a summary of the tasks you just delegated.

8. Click **Finish** to complete the delegation.

When the Delegation of Control Wizard completes, the administrative authority you specified is delegated. You can modify the tasks available for delegation (presented on the **Tasks to Delegate** page of the Wizard) by editing Delegwiz.inf, a file stored on the domain controller. For more information about modifying Delegwiz.inf, see article 308404, "HOW TO: Customize the Task List in the Delegation Wizard," in the Microsoft Knowledge Base on the Web at http://go.microsoft.com/fwlink/?LinkId=3202.

You can use the lists in Appendix A: Active Directory Administrative Tasks to obtain the precise set of permissions required to delegate a specific data administration task and customize the Delegwiz.inffile. You can then use the Delegation of Control Wizard to delegate a larger number of tasks. Note that each domain controller stores its own local copy of the file, so you need to copy the modified file to each domain controller where you want to use the Delegation of Control Wizard.

The Delegation of Control Wizard also allows you to delegate a custom task. To delegate a custom task, complete the following steps.

**To create a custom task**

1. On the **Tasks to Delegate** page of the Delegation of Control Wizard, select **Create a custom task to delegate**, and then click **Next**.




2. On the **Active Directory Object Type** page, if you want to delegate control over a specific class of object, select **Only the following objects in the folder** and then select the check boxes for the classes of objects for which you want to delegate control.Otherwise, select **This folder, existing objects in this folder, and creation of new objects in this folder**to delegate control over the folder and all objects in it, and to delegate the ability to create new objects in the folder. When you have finished making selections, click **Next**.




3. Click **Next**. On the **Permissions** page, under **Show these permissions**, select one of the check boxes to display the permissions to delegate:

    a. To delegate general permissions, select **General** and then select the check boxes for the permissions you want to delegate.




b. To delegate access to one or more properties, click **Property Specific** and then select the check boxes for the permissions you want to delegate.




c. To delegate the ability to create and/or delete objects, click **Creation/deletion of specific child objects** and then select the check boxes for the permissions you want to delegate.




4. Click **Next**. The Delegation of Control Wizard presents a summary of the tasks you selected to delegate.

5. Click **Finish**. Once the Delegation Wizard completes, the custom task you specified is delegated

**ACL Editor**
The ACL Editor allows you to view and modify the security descriptors of objects in Active Directory. You can inspect and change security permissions, security auditing, and object ownership information. Although the ACL Editor is primarily for managing access-control permissions, it can also be used to delegate administrative authority, because that involves granting the delegated user sufficient permissions to carry out the low-level operations that map to the delegated administrative task. More significantly, the ACL Editor is currently the only tool that can be used to un-delegate tasks or revoke delegated authority.

**Using the ACL Editor**
You can access theACL Editor from any of the MMC snap-ins listed earlier in this appendix, and from ADSI Edit (Adsiedit.mcs), a graphical Windows Support Tool that is available in the Support\Tools folder on the Windows CD. When Windows Support Tools are installed, you can add ADSI Edit to an MMC snap-in. ADSI Edit provides the ability to edit the attribute values of any object in any directory partition in the forest, and to examine RootDSE attributes.

To use the ACL Editor, start any of these tools:

- Active Directory Users and Computers

- Active Directory Sites and Services

- Adsiedit.mcs

**To delegate permissions by using the ACL Editor**

1. Right click the object on which you want to set permissions and select **Properties**. The **Accounts Properties** dialog box is displayed.




2. To view and change advanced security information and grant or revoke permissions, in the **Accounts Properties** dialog box, on the **Securities** tab, click **Advanced** to open the **Access Control Settings for Accounts**dialog box.




In the**Access Control Settings for Accounts**dialog box, you can view the permission entries in the DACL of the object, specify the operations that should be audited, and access and modify ownership information.

By clearing the check box for **Allow inheritable permissions from parent to propagate to this object**, you can protect the DACL of the object from inheriting permissions from the parent object. When the DACL of an object is protected, the object no longer inherits any permissions from the parent object, and any permissions marked as inheritable on the parent object do not flow down to child objects of the protected object.

3. Click **Add**. The **Select User, Computer, Or Group** dialog box is displayed.

4. In the **Select User, Computer, or Group** dialog box, in the **Name** box, type the name of the user or security group for which to specify permissions ( delegate or undelegate administrative), and then click **OK**. The **Permission Entry for Accounts** dialog box is displayed.




On the **Object** tab, you can allow or deny a user or group generic permissions and all extended rights. In the **Apply-onto** box, you can specify the objects affected by the permissions being applied.

The **Apply these permissions to objects and/or containers within this container only**option controls whether the permissions being applied should be propagated only down to (be inherited by) immediate child objects of this container, and not be inherited further down the hierarchy.

5. To allow or deny access to specific properties, click the **Properties** tab.




The **Properties** tab displays all properties of objects of the class specified in the **Apply onto** box. To grant read or write access to a specific property, select the **Allow** check box for the read or write entry of that property.

Note that the **Permissions** list box on the **Properties** tab might not display every property of the object. To make the list easier to manage, the user interface for access control does not display object and property types.

The list of properties is filtered based on specific settings in a local text file. For more information about filtering of properties, see article 296490 "How to Modify the Filtered Properties of an Object" in the Microsoft Knowledge Base on the Web at http://go.microsoft.com/fwlink/?LinkId=3202.

Also note that every property in Active Directory has three different names:

- **Common-Name.** Every object in Active Directory has a naming attribute from which its relative distinguished name (also known as RDN) is formed. The Naming Attribute for Attribute-Schema objects is Common-Name. The value assigned to Common-Name is the value that the Attribute-Schema has as its Relative Distinguished Name.

- **LDAP-Display-Name.** The name of this attribute is known to the LDAP agent for the Active Directory directory service. This is the name LDAP clients must use to read or write this attribute.

- **Admin-Display-Name.** The Common-Name of a given object might not be descriptive enough for use in administration tools. Admin-Display-Name is available for tools to use as a display name for an attribute when the naming attribute is not appropriate.

Active Directory MMC snap-in tools refer to classes and attributes by the Admin-Display-Nameif one is defined for that class or attribute. For those classes or attributes that do not have a defined Admin-Display-Name, these tools instead use the LDAP-Display Name.

Consequently, if you cannot find a specific attribute in this tab, perform the following checks:

- Check to see that the value in the **Apply onto**box covers the class of objects on which this property exists. For example, if you want to grant Read access to the Organization-Nameproperty on User objects, check to make sure that the **Apply onto** box is set to **User objects**, **This object and all child objects**, or **Child objects only**.

  Additionally, if the object you are specifying permissions on is of class User, the **Apply onto** box can also be set to **This object only**.

  The point is that if the **Apply onto** box is set to an object class that does not include Organization-Nameas a property on the object, the property does not display.

- If, after ensuring that the **Apply onto** box is set to an appropriate value, you still cannot see the property, check to see whether the property has an Admin-Display Name defined.

  For example, you cannot view Organization-Namein the list even if the **Apply onto** box is set to User objects. In this case, refer to Appendix H: "Active Directory Display Name Mappings", later in this document. This appendix provides the corresponding Admin-Display-Name(if one exists)for every class and attribute in the Active Directory Schema. If the attribute has an Admin-Display-Namedefined, search for it instead. In this case, Appendix H shows that Organization-Namedoes not have an Admin-Display-Name defined. Consequently, you need to look for the LDAP-Display-Name, which is "o" for this attribute.

- If you still cannot find the attribute, it might be filtered. In the case of Organization-Name, you will not be able to find "o" in the list because it is filtered. Turn off filtering for that specific attribute by modifying the Dsec.dat attribute as described earlier in this appendix.

  Note that the Dsec.dat file also uses the Admin-Display-Name if one exists, otherwise it uses the LDAP-Display-Name. This will help you look for the correct name in the file if you need to turn-off filtering.

  Alternatively, to avoid this filtering altogether, you can rename the Dssec.dat file so that it is no longer used by the administrative tools user interfaces. Doing so results in less complexity, because you no longer have to deal with filtered attributes. This makes the user interface somewhat more cumbersome to use because all potential attributes are displayed. However, because they are displayed in alphabetical order by default, this is usually not a problem.

Once you turn-off filtering for an attribute, you should finally be able to view the attribute. In the example, you will finally be able to see "o" in the list. You can then allow or deny access to the attribute.

Also be aware of the effect that making different selections in the **Apply onto** box has on the inheritance of permissions.




The **Apply onto** box specifies the class of objects on which the selected permissions will be applied. For example, if you want to grant a delegated user the ability to modify the User-Account-Controlattribute but only on User objects but not on Computer objects, choose **User objects** from the list in the **Apply onto** box.

The following table describes how different **Apply onto** settings affect the setting of inheritance flags for the specified permissions in the ACE of the selected object:

| **Apply-Onto Value** | **Effect on inheritance of permissions** |
|---|---|
| This object only | No inheritance flags set. The resulting ACE will not be an inheritable ACE. |
| This object and all child objects | Sets the Container Inheritflags and leaves the Inherited-Object-Typefield empty. The resulting ACE will be inheritable by all child objects and will be effective (enforced) on this object and all child objects. |
| Child objects only | Sets the Container Inheritand Inherit Onlyflags and leaves the Inherited-Object-Typefield empty. The resulting ACE will be inheritable by all child objects but will be effective on child objects only; it will not be effective on the object on which it is being applied. |
| <*Specific object class*> objects | Sets the Container Inheritand Inherit Onlyflags and sets the value of the Inherited-Object-Typefield to the GUID representing objects of class <*Specific object class*>. The resulting ACE will be inheritable by all child objects, but will only be effective on objects of class <*Specific object class*>. It will not be effective on this object or on child objects of other classes. |

Thus, to set an Inherit-Only ACE, in the **Apply onto** box, choose **Childobjects only or <Specific object class> objects**, as appropriate.

Notethat the selection you make here also applies to the permissions being set on the **Object** tab of the Permission Entry for Accounts dialog box.

6. When you have finished making selections on the **Object** and **Property** tabs, click **OK**. Then, in the **Access Control Settings**dialog box, click **OK** again to apply the settings you selected.

   Administrative authority is now delegated or revoked by having granted or revoked permissions for the selected user or group to perform one or more administrative tasks.

**LDP**

LDP (Ldp.exe) is a graphical tool that allows you to perform Lightweight Directory Access Protocol (LDAP) operations, such as connect, bind, search, modify, add, or delete, against any LDAP-compatible directory, including Active Directory. Many objects stored in Active Directory are not readily displayed using the graphical tools that are installed with Windows. Administrators can install the Ldp.exe tool to view these objects and their metadata, such as security descriptors and replication metadata, to aid in problem determination.

**Using LDP**

Ldp.exe is a support tool that is installed from the Windows CD. For information about how to install Windows support tools, see article 301423, "HOW TO: Install the Windows 2000 Support Tools to a Windows 2000 Server-Based Computer," in the Microsoft Knowledge Base on the Web at http://go.microsoft.com/fwlink/?LinkId=3202. For delegation of administration, you can use LDP to view the security descriptors of Active Directory objects by performing the following tasks. To view the security descriptor of an object, use the following procedure.

**To view the security descriptor of an object by using Ldp.exe**

1. In LDP, on the **Connection** menu, click **Connect** to connect to a domain or a specific domain controller.
2. In the**Connect** dialog box, in the**Server** box, type a server name or leave the entry blank to connect to the local server, and then click**OK**.
3. On the **Connection** menu, click **Bind**.
4. In the **Bind** dialog box, type a user name and password, and then click **OK** to bind to Active Directory.
5. On the **View** menu, click **Tree**. In the **BaseDN** box, either type a specific distinguished name (DN) or leave **BaseDN** blank to view the entire **domain**.
6. To display the object for which you want to view the security descriptor, double-click the domain object in the tree view and then double-click the appropriate container.
7. To view the security descriptor of an object, right click the object in the tree view, select **Advanced**,select **Security Descriptor**, and then in the **Security Descriptor** dialog box, click **OK**.

   The security descriptor of the object is displayed in the details pane. Note that you can scroll to view the Security Descriptor Definition Language (SDDL) version of the security descriptor and to view the security descriptor in text format.

   To analyze the security descriptor in detail, you can either view the information in the details pane, or, if you prefer, right click in the details pane, choose **Select All**, choose **Copy**, and then paste the contents into a text file.

**Dsacls**

Dsacls (Dscals.exe) is a powerful command-line utility that can be used to report and modify permissions on Active Directory objects. Dsacls is included in the Windows support tools that you can install from the Windows CD. For information about how to install Windows support tools, see article 301423, "HOW TO: Install the Windows 2000 Support Tools to a Windows 2000 Server-Based Computer," in the Microsoft Knowledge Base on the Web at http://go.microsoft.com/fwlink/?LinkId=3202. For details about how Dsacls works, run Dsacls.exe to display the help file.

**Acldiag**

Acldiag (Acldiag.exe) is a command-line utility that can be used determine the effective access granted to a specific user or group based on the permissions in the DACL of a specific object. It can also be used to verify that the explicit permissions in the DACL of an object include the permissions specified in the default security settings for objects of that class as defined in the Active Directory Schema. Finally, Acldiag can also be used to determine whether the permissions specified in the object's DACL map to any of the delegation templates in the Delegation of Control Wizard and can be used to repair any delegation templates that are not completely applied.

Acldiag is included in the Windows support tools that you can install from the Windows CD. For information about how to install Windows support tools, see article 301423, "HOW TO: Install the Windows 2000 Support Tools to a Windows 2000 Server-Based Computer," in the Microsoft Knowledge Base on the Web at http://go.microsoft.com/fwlink/?LinkId=3202. For details about how Acldiag works, run Acldiag.exe to display the help file.

**Dsrevoke**

Dsrevoke is a new command-line tool that can be used on domain controllers that are running Windows Server 2003 or Windows 2000 Server to report the existence of all permissions for a specific user or group on a set of OUs in a domain and optionally remove from the DACLs of a set of OUs all permissions specified for a particular user or group. Dsrevoke

complements the functionality provided by Delegation of Control Wizard, which is used to delegate administrative authority, by providing the ability to revoke delegated administrative authority.

**Using Dsrevoke**

You can download Dsrevoke from the Microsoft Web site at http://go.microsoft.com/fwlink/?linkid=12538 and searching by keyword for "dsrevoke."

To maximize the benefits offered by Dsrevoke, follow these guidelines as much as possible when delegating administrative authority:

- Use roles to delegate administrative authority. When delegating roles, be sure to use a unique and specific security group to represent every unique and specific role instance.

- Use inheritance to grant permissions to the security group representing a role instance, and grant permissions on OUs.

Delegating administrative authority by using roles involves the following tasks:

1. Create a specific and unique security group to represent the role.

2. Identify the highest level OU that represents the root of the smallest subtree that contains the subset of all objects the delegated user needs to access and modify in order to perform the delegated tasks.

3. Run the Delegation of Control Wizard on that OU and delegate the required administrative tasks to the unique and specific security group representing the unique and specific role.

If you follow these delegation guidelines, you can use Dsrevoke to easily and reliably undelegate authority. Simply run Dsrevoke in the domain, providing as input the name of the specific security group used to represent the delegated role, and use the **/report** switch to verify the existence of all explicit permissions for that security group that have been set on all OU objects in the domain . Once you have reviewed the reported permissions, you can use the **/remove** switch to revoke all permissions granted to that security group, thereby revoking the delegated authority.

The following is the usage for Dsrevoke:

**dsrevoke /report|/remove** [**/domain:**<*domainname*>] [**/username:**<*username*>] [**/password:**<*password*>|*] [**/root:**<*domain/OU*>] <*securityprincipal*>

- **/report**

  Only reports the ACEs that have been set for the given principal on all domain and OU objects under root.

- **/remove**

  Reports and then removes (after confirmation) the aces for the given principal.

- **/domain**

  DNS or Netbios name of domain (must be specified when <securityprincipal> is in domain other than default or if alternate credentials are provided).

- **/username**

  Username if alternate credentials must be specified.

- **/password**

  Will prompt for password.

- **/root**

Root OU to start search for ACEs. If not specified will default to the specified domain's default naming context (the root domain or OU must be specified using x500 format; if the domain name must include spaces, enclose the option in quotes, e.g., "/root:...").

- *<securityprincipal>*

  Domain\User or Domain\Group for the security principal being looked up.

| **Note** |
| --- |
| Dsrevoke removes only permissions; if a role has user rights applied, you must manually remove them by modifying the appropriate Group Policy. Also, because Dsrevoke works only on domain objects and OUs, you must manually remove ACEs if you set them on a container object or if you explicitly set permissions on an object within a container or OU. For this reason, it is recommended that you always apply permissions to OUs rather than to specific objects within OUs, and that you apply permissions to child OUs by using inheritance. Finally, because Dsrevoke works only on domain objects and OUs, you cannot use it to remove permissions from the Configuration and Schema directory partitions. Consequently, you typically cannot use Dsrevoke to revoke delegation of service management tasks. |

# Appendix H: Active Directory Display Name Mappings

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix shows display name mappings organized in the following tables:

- Classes Display
- Attributes Display
- Attributes Display, Class View

**Classes Display Name Mappings**

| cn | lDAPDisplayName | Attribute Display Name(s) |
|---|---|---|
| account | account | |
| ACS-Policy | aCSPolicy | |
| ACS-Resource-Limits | aCSResourceLimits | |
| ACS-Subnet | aCSSubnet | |
| Address-Book-Container | addressBookContainer | |
| Address-Template | addressTemplate | |
| Application-Entity | applicationEntity | |
| Application-Process | applicationProcess | |
| Application-Settings | applicationSettings | |
| Application-Site-Settings | applicationSiteSettings | |
| Application-Version | applicationVersion | |
| Attribute-Schema | attributeSchema | |
| Builtin-Domain | builtinDomain | |
| Category-Registration | categoryRegistration | |
| Certification-Authority | certificationAuthority | |
| Class-Registration | classRegistration | |
| Class-Schema | classSchema | |
| Class-Store | classStore | |
| Com-Connection-Point | comConnectionPoint | |
| Computer | computer | Computer |
| Configuration | configuration | |
| Connection-Point | connectionPoint | |
| Contact | contact | Contact |
| Container | container | Container |
| Control-Access-Right | controlAccessRight | |
| Country | country | |
| CRL-Distribution-Point | cRLDistributionPoint | |
| Cross-Ref | crossRef | |
| Cross-Ref-Container | crossRefContainer | |
| Device | device | |
| Dfs-Configuration | dfsConfiguration | |
| DHCP-Class | dHCPClass | |

| | | |
|---|---|---|
| Display-Specifier | displaySpecifier | |
| Display-Template | displayTemplate | |
| DMD | dMD | |
| Dns-Node | dnsNode | |
| Dns-Zone | dnsZone | |
| document | document | |
| documentSeries | documentSeries | |
| Domain | domain | |
| Domain-DNS | domainDNS | Domain |
| Domain-Policy | domainPolicy | Domain Policy |
| domainRelatedObject | domainRelatedObject | |
| DSA | dSA | |
| DS-UI-Settings | dSUISettings | |
| Dynamic-Object | dynamicObject | |
| File-Link-Tracking | fileLinkTracking | |
| File-Link-Tracking-Entry | fileLinkTrackingEntry | |
| Foreign-Security-Principal | foreignSecurityPrincipal | Foreign Security Principal |
| friendlyCountry | friendlyCountry | |
| FT-Dfs | fTDfs | |
| Group | group | Group |
| Group-Of-Names | groupOfNames | |
| groupOfUniqueNames | groupOfUniqueNames | |
| Group-Policy-Container | groupPolicyContainer | |
| Index-Server-Catalog | indexServerCatalog | |
| inetOrgPerson | inetOrgPerson | InetOrgPerson |
| Infrastructure-Update | infrastructureUpdate | |
| Intellimirror-Group | intellimirrorGroup | IntelliMirror Group |
| Intellimirror-SCP | intellimirrorSCP | IntelliMirror Service |
| Inter-Site-Transport | interSiteTransport | Inter-Site Transport |
| Inter-Site-Transport-Container | interSiteTransportContainer | Inter-Site Transports Container |
| Ipsec-Base | ipsecBase | |
| Ipsec-Filter | ipsecFilter | |
| Ipsec-ISAKMP-Policy | ipsecISAKMPPolicy | |
| Ipsec-Negotiation-Policy | ipsecNegotiationPolicy | |
| Ipsec-NFA | ipsecNFA | |
| Ipsec-Policy | ipsecPolicy | |
| Leaf | Leaf | |
| Licensing-Site-Settings | licensingSiteSettings | Licensing Site Settings |
| Link-Track-Object-Move-Table | linkTrackObjectMoveTable | |
| Link-Track-OMT-Entry | linkTrackOMTEntry | |
| Link-Track-Vol-Entry | linkTrackVolEntry | |
| Link-Track-Volume-Table | linkTrackVolumeTable | |

| | | |
|---|---|---|
| Locality | locality | |
| Lost-And-Found | lostAndFound | |
| Mail-Recipient | mailRecipient | |
| Meeting | meeting | |
| ms-COM-Partition | msCOM-Partition | |
| ms-COM-PartitionSet | msCOM-PartitionSet | |
| ms-DS-App-Configuration | msDS-App-Configuration | |
| ms-DS-App-Data | msDS-AppData | |
| ms-DS-Az-Admin-Manager | msDS-AzAdminManager | |
| ms-DS-Az-Application | msDS-AzApplication | |
| ms-DS-Az-Operation | msDS-AzOperation | |
| ms-DS-Az-Role | msDS-AzRole | |
| ms-DS-Az-Scope | msDS-AzScope | |
| ms-DS-Az-Task | msDS-AzTask | |
| ms-DS-Quota-Container | msDS-QuotaContainer | |
| ms-DS-Quota-Control | msDS-QuotaControl | |
| ms-Exch-Configuration-Container | msExchConfigurationContainer | |
| ms-ieee-80211-Policy | msieee80211-Policy | |
| MSMQ-Configuration | mSMQConfiguration | MSMQ Configuration |
| MSMQ-Custom-Recipient | msMQ-Custom-Recipient | MSMQ Queue Alias |
| MSMQ-Enterprise-Settings | mSMQEnterpriseSettings | MSMQ Enterprise |
| MSMQ-Group | msMQ-Group | MSMQ Group |
| MSMQ-Migrated-User | mSMQMigratedUser | MSMQ Upgraded User |
| MSMQ-Queue | mSMQQueue | MSMQ Queue |
| MSMQ-Settings | mSMQSettings | MSMQ Settings |
| MSMQ-Site-Link | mSMQSiteLink | MSMQ Routing Link |
| ms-PKI-Enterprise-Oid | msPKI-Enterprise-Oid | |
| ms-PKI-Key-Recovery-Agent | msPKI-Key-Recovery-Agent | |
| ms-PKI-Private-Key-Recovery-Agent | msPKI-PrivateKeyRecoveryAgent | |
| MS-SQL-OLAPCube | mS-SQL-OLAPCube | |
| MS-SQL-OLAPDatabase | mS-SQL-OLAPDatabase | |
| MS-SQL-OLAPServer | mS-SQL-OLAPServer | |
| MS-SQL-SQLDatabase | mS-SQL-SQLDatabase | |
| MS-SQL-SQLPublication | mS-SQL-SQLPublication | |
| MS-SQL-SQLRepository | mS-SQL-SQLRepository | |
| MS-SQL-SQLServer | mS-SQL-SQLServer | |
| ms-TAPI-Rt-Conference | msTAPI-RtConference | |
| ms-TAPI-Rt-Person | msTAPI-RtPerson | |
| ms-WMI-IntRangeParam | msWMI-IntRangeParam | |
| ms-WMI-IntSetParam | msWMI-IntSetParam | |
| ms-WMI-MergeablePolicyTemplate | msWMI-MergeablePolicyTemplate | |

| | | |
|---|---|---|
| ms-WMI-ObjectEncoding | msWMI-ObjectEncoding | |
| ms-WMI-PolicyTemplate | msWMI-PolicyTemplate | |
| ms-WMI-PolicyType | msWMI-PolicyType | |
| ms-WMI-RangeParam | msWMI-RangeParam | |
| ms-WMI-RealRangeParam | msWMI-RealRangeParam | |
| ms-WMI-Rule | msWMI-Rule | |
| ms-WMI-ShadowObject | msWMI-ShadowObject | |
| ms-WMI-SimplePolicyTemplate | msWMI-SimplePolicyTemplate | |
| ms-WMI-Som | msWMI-Som | |
| ms-WMI-StringSetParam | msWMI-StringSetParam | |
| ms-WMI-UintRangeParam | msWMI-UintRangeParam | |
| ms-WMI-UintSetParam | msWMI-UintSetParam | |
| ms-WMI-UnknownRangeParam | msWMI-UnknownRangeParam | |
| ms-WMI-WMIGPO | msWMI-WMIGPO | |
| NTDS-Connection | nTDSConnection | Connection |
| NTDS-DSA | nTDSDSA | Domain Controller Settings |
| NTDS-Service | nTDSService | Active Directory Service |
| NTDS-Site-Settings | nTDSSiteSettings | Site Settings |
| NTFRS-Member | nTFRSMember | FRS Member |
| NTFRS-Replica-Set | nTFRSReplicaSet | FRS Replica Set |
| NTFRS-Settings | nTFRSSettings | FRS Settings |
| NTFRS-Subscriber | nTFRSSubscriber | FRS Subscriber |
| NTFRS-Subscriptions | nTFRSSubscriptions | FRS Subscriptions |
| Organization | organization | |
| Organizational-Person | organizationalPerson | |
| Organizational-Role | organizationalRole | |
| Organizational-Unit | organizationalUnit | Organizational Unit |
| Package-Registration | packageRegistration | |
| Person | person | |
| Physical-Location | physicalLocation | |
| PKI-Certificate-Template | pKICertificateTemplate | Certificate Template |
| PKI-Enrollment-Service | pKIEnrollmentService | |
| Print-Queue | printQueue | Printer |
| Query-Policy | queryPolicy | Query Policy |
| Remote-Mail-Recipient | remoteMailRecipient | |
| Remote-Storage-Service-Point | remoteStorageServicePoint | Remote Storage Service |
| Residential-Person | residentialPerson | |
| rFC822LocalPart | rFC822LocalPart | |
| RID-Manager | rIDManager | |
| RID-Set | rIDSet | |
| room | room | |

| | | |
|---|---|---|
| Rpc-Container | rpcContainer | RPC Services |
| rpc-Entry | rpcEntry | |
| rpc-Group | rpcGroup | |
| rpc-Profile | rpcProfile | |
| rpc-Profile-Element | rpcProfileElement | |
| rpc-Server | rpcServer | |
| rpc-Server-Element | rpcServerElement | |
| RRAS-Administration-Connection-Point | rRASAdministrationConnectionPoint | |
| RRAS-Administration-Dictionary | rRASAdministrationDictionary | |
| Sam-Domain | samDomain | |
| Sam-Domain-Base | samDomainBase | |
| Sam-Server | samServer | |
| Secret | secret | |
| Security-Object | securityObject | |
| Security-Principal | securityPrincipal | |
| Server | server | Server |
| Servers-Container | serversContainer | Servers Container |
| Service-Administration-Point | serviceAdministrationPoint | Service |
| Service-Class | serviceClass | |
| Service-Connection-Point | serviceConnectionPoint | |
| Service-Instance | serviceInstance | |
| simpleSecurityObject | simpleSecurityObject | |
| Site | site | Site |
| Site-Link | siteLink | Site Link |
| Site-Link-Bridge | siteLinkBridge | Site Link Bridge |
| Sites-Container | sitesContainer | Sites Container |
| Storage | storage | |
| Subnet | subnet | Subnet |
| Subnet-Container | subnetContainer | Subnets Container |
| SubSchema | subSchema | |
| Top | top | |
| Trusted-Domain | trustedDomain | Trusted Domain |
| Type-Library | typeLibrary | |
| User | user | User |
| Volume | volume | Shared Folder |
| default* | | |
| DS-UI-Default-Settings* | | |
| localPolicy* | | Local Policy |
| nTDSSettings* | | Settings |

*These items are not defined in the schema, but are provided with classDisplayName or have a placeholder in the display specifier container.

**Attributes Display Name Mappings**

| **cn** | **lDAPDisplayName** | **Attribute Display Name(s)** |
|---|---|---|

| | | |
|---|---|---|
| Account-Expires | accountExpires | |
| Account-Name-History | accountNameHistory | |
| ACS-Aggregate-Token-Rate-Per-User | aCSAggregateTokenRatePerUser | |
| ACS-Allocable-RSVP-Bandwidth | aCSAllocableRSVPBandwidth | |
| ACS-Cache-Timeout | aCSCacheTimeout | |
| ACS-Direction | aCSDirection | |
| ACS-DSBM-DeadTime | aCSDSBMDeadTime | |
| ACS-DSBM-Priority | aCSDSBMPriority | |
| ACS-DSBM-Refresh | aCSDSBMRefresh | |
| ACS-Enable-ACS-Service | aCSEnableACSService | |
| ACS-Enable-RSVP-Accounting | aCSEnableRSVPAccounting | |
| ACS-Enable-RSVP-Message-Logging | aCSEnableRSVPMessageLogging | |
| ACS-Event-Log-Level | aCSEventLogLevel | |
| ACS-Identity-Name | aCSIdentityName | |
| ACS-Max-Aggregate-Peak-Rate-Per-User | aCSMaxAggregatePeakRatePerUser | |
| ACS-Max-Duration-Per-Flow | aCSMaxDurationPerFlow | |
| ACS-Maximum-SDU-Size | aCSMaximumSDUSize | |
| ACS-Max-No-Of-Account-Files | aCSMaxNoOfAccountFiles | |
| ACS-Max-No-Of-Log-Files | aCSMaxNoOfLogFiles | |
| ACS-Max-Peak-Bandwidth | aCSMaxPeakBandwidth | |
| ACS-Max-Peak-Bandwidth-Per-Flow | aCSMaxPeakBandwidthPerFlow | |
| ACS-Max-Size-Of-RSVP-Account-File | aCSMaxSizeOfRSVPAccountFile | |
| ACS-Max-Size-Of-RSVP-Log-File | aCSMaxSizeOfRSVPLogFile | |
| ACS-Max-Token-Bucket-Per-Flow | aCSMaxTokenBucketPerFlow | |
| ACS-Max-Token-Rate-Per-Flow | aCSMaxTokenRatePerFlow | |
| ACS-Minimum-Delay-Variation | aCSMinimumDelayVariation | |
| ACS-Minimum-Latency | aCSMinimumLatency | |
| ACS-Minimum-Policed-Size | aCSMinimumPolicedSize | |
| ACS-Non-Reserved-Max-SDU-Size | aCSNonReservedMaxSDUSize | |
| ACS-Non-Reserved-Min-Policed-Size | aCSNonReservedMinPolicedSize | |
| ACS-Non-Reserved-Peak-Rate | aCSNonReservedPeakRate | |
| ACS-Non-Reserved-Token-Size | aCSNonReservedTokenSize | |
| ACS-Non-Reserved-Tx-Limit | aCSNonReservedTxLimit | |
| ACS-Non-Reserved-Tx-Size | aCSNonReservedTxSize | |
| ACS-Permission-Bits | aCSPermissionBits | |
| ACS-Policy-Name | aCSPolicyName | |
| ACS-Priority | aCSPriority | |
| ACS-RSVP-Account-Files-Location | aCSRSVPAccountFilesLocation | |
| ACS-RSVP-Log-Files-Location | aCSRSVPLogFilesLocation | |
| ACS-Server-List | aCSServerList | |
| ACS-Service-Type | aCSServiceType | |

| | | |
|---|---|---|
| ACS-Time-Of-Day | aCSTimeOfDay | |
| ACS-Total-No-Of-Flows | aCSTotalNoOfFlows | |
| Additional-Trusted-Service-Names | additionalTrustedServiceNames | |
| Address-Book-Roots | addressBookRoots | |
| Address-Entry-Display-Table | addressEntryDisplayTable | |
| Address-Entry-Display-Table-MSDOS | addressEntryDisplayTableMSDOS | |
| Address-Syntax | addressSyntax | |
| Address-Type | addressType | |
| Admin-Context-Menu | adminContextMenu | |
| Admin-Count | adminCount | |
| Admin-Description | adminDescription | |
| Admin-Display-Name | adminDisplayName | |
| Admin-Multiselect-Property-Pages | adminMultiselectPropertyPages | |
| Admin-Property-Pages | adminPropertyPages | |
| Allowed-Attributes | allowedAttributes | |
| Allowed-Attributes-Effective | allowedAttributesEffective | |
| Allowed-Child-Classes | allowedChildClasses | |
| Allowed-Child-Classes-Effective | allowedChildClassesEffective | |
| Alt-Security-Identities | altSecurityIdentities | |
| ANR | aNR | |
| Application-Name | applicationName | |
| Applies-To | appliesTo | |
| App-Schema-Version | appSchemaVersion | |
| Asset-Number | assetNumber | Asset Number |
| Assistant | assistant | Assistant |
| associatedDomain | associatedDomain | |
| associatedName | associatedName | |
| Assoc-NT-Account | assocNTAccount | |
| attributeCertificateAttribute | attributeCertificateAttribute | |
| Attribute-Display-Names | attributeDisplayNames | |
| Attribute-ID | attributeID | |
| Attribute-Security-GUID | attributeSecurityGUID | |
| Attribute-Syntax | attributeSyntax | |
| Attribute-Types | attributeTypes | |
| audio | audio | |
| Auditing-Policy | auditingPolicy | |
| Authentication-Options | authenticationOptions | |
| Authority-Revocation-List | authorityRevocationList | |
| Auxiliary-Class | auxiliaryClass | |
| Bad-Password-Time | badPasswordTime | |
| Bad-Pwd-Count | badPwdCount | |
| Birth-Location | birthLocation | |

| | | |
|---|---|---|
| Bridgehead-Server-List-BL | bridgeheadServerListBL | |
| Bridgehead-Transport-List | bridgeheadTransportList | |
| buildingName | buildingName | |
| Builtin-Creation-Time | builtinCreationTime | |
| Builtin-Modified-Count | builtinModifiedCount | |
| Business-Category | businessCategory | |
| Bytes-Per-Minute | bytesPerMinute | |
| Country-Name | c | Country Abbreviation |
| CA-Certificate | cACertificate | |
| CA-Certificate-DN | cACertificateDN | |
| CA-Connect | cAConnect | |
| Canonical-Name | canonicalName | |
| Can-Upgrade-Script | canUpgradeScript | |
| carLicense | carLicense | |
| Catalogs | catalogs | |
| Categories | categories | |
| Category-Id | categoryId | |
| CA-Usages | cAUsages | |
| CA-WEB-URL | cAWEBURL | |
| Certificate-Authority-Object | certificateAuthorityObject | |
| Certificate-Revocation-List | certificateRevocationList | |
| Certificate-Templates | certificateTemplates | |
| Class-Display-Name | classDisplayName | |
| Common-Name | cn | Name, Directory Service Name |
| Text-Country | co | Country |
| Code-Page | codePage | |
| COM-ClassID | cOMClassID | |
| COM-CLSID | cOMCLSID | |
| COM-InterfaceID | cOMInterfaceID | |
| User-Comment | comment | Comment |
| COM-Other-Prog-Id | cOMOtherProgId | |
| Company | company | Company |
| COM-ProgID | cOMProgID | |
| COM-Treat-As-Class-Id | cOMTreatAsClassId | |
| COM-Typelib-Id | cOMTypelibId | |
| COM-Unique-LIBID | cOMUniqueLIBID | |
| Content-Indexing-Allowed | contentIndexingAllowed | |
| Context-Menu | contextMenu | |
| Control-Access-Rights | controlAccessRights | |
| Cost | cost | |
| Country-Code | countryCode | |
| Create-Dialog | createDialog | |

| | | |
|---|---|---|
| Create-Time-Stamp | createTimeStamp | |
| Create-Wizard-Ext | createWizardExt | |
| Creation-Time | creationTime | |
| Creation-Wizard | creationWizard | |
| Creator | creator | |
| CRL-Object | cRLObject | |
| CRL-Partitioned-Revocation-List | cRLPartitionedRevocationList | |
| Cross-Certificate-Pair | crossCertificatePair | |
| Current-Location | currentLocation | |
| Current-Parent-CA | currentParentCA | |
| Current-Value | currentValue | |
| Curr-Machine-Id | currMachineId | |
| DBCS-Pwd | dBCSPwd | |
| Domain-Component | dc | Name |
| Default-Class-Store | defaultClassStore | |
| Default-Group | defaultGroup | |
| Default-Hiding-Value | defaultHidingValue | |
| Default-Local-Policy-Object | defaultLocalPolicyObject | |
| Default-Object-Category | defaultObjectCategory | |
| Default-Priority | defaultPriority | |
| Default-Security-Descriptor | defaultSecurityDescriptor | |
| Delta-Revocation-List | deltaRevocationList | |
| Department | department | Department |
| departmentNumber | departmentNumber | |
| Description | description | Description, Comment |
| Desktop-Profile | desktopProfile | |
| Destination-Indicator | destinationIndicator | |
| dhcp-Classes | dhcpClasses | |
| dhcp-Flags | dhcpFlags | |
| dhcp-Identification | dhcpIdentification | |
| dhcp-Mask | dhcpMask | |
| dhcp-MaxKey | dhcpMaxKey | |
| dhcp-Obj-Description | dhcpObjDescription | |
| dhcp-Obj-Name | dhcpObjName | |
| dhcp-Options | dhcpOptions | |
| dhcp-Properties | dhcpProperties | |
| dhcp-Ranges | dhcpRanges | |
| dhcp-Reservations | dhcpReservations | |
| dhcp-Servers | dhcpServers | |
| dhcp-Sites | dhcpSites | |
| dhcp-State | dhcpState | |
| dhcp-Subnets | dhcpSubnets | |

| | | |
|---|---|---|
| dhcp-Type | dhcpType | |
| dhcp-Unique-Key | dhcpUniqueKey | |
| dhcp-Update-Time | dhcpUpdateTime | |
| Reports | directReports | Direct Reports |
| Display-Name | displayName | Display Name |
| Display-Name-Printable | displayNamePrintable | |
| Obj-Dist-Name | distinguishedName | X500 Distinguished Name |
| DIT-Content-Rules | dITContentRules | |
| Division | division | Division |
| DMD-Location | dMDLocation | |
| DMD-Name | dmdName | |
| DN-Reference-Update | dNReferenceUpdate | |
| Dns-Allow-Dynamic | dnsAllowDynamic | |
| Dns-Allow-XFR | dnsAllowXFR | |
| DNS-Host-Name | dNSHostName | |
| Dns-Notify-Secondaries | dnsNotifySecondaries | |
| DNS-Property | dNSProperty | |
| Dns-Record | dnsRecord | |
| Dns-Root | dnsRoot | |
| Dns-Secure-Secondaries | dnsSecureSecondaries | |
| DNS-Tombstoned | dNSTombstoned | |
| documentAuthor | documentAuthor | |
| documentIdentifier | documentIdentifier | |
| documentLocation | documentLocation | |
| documentPublisher | documentPublisher | |
| documentTitle | documentTitle | |
| documentVersion | documentVersion | |
| Domain-Certificate-Authorities | domainCAs | |
| Domain-Cross-Ref | domainCrossRef | |
| Domain-ID | domainID | |
| Domain-Identifier | domainIdentifier | |
| Domain-Policy-Object | domainPolicyObject | |
| Domain-Policy-Reference | domainPolicyReference | |
| Domain-Replica | domainReplica | |
| Domain-Wide-Policy | domainWidePolicy | |
| drink | drink | |
| Driver-Name | driverName | Model |
| Driver-Version | driverVersion | |
| DSA-Signature | dSASignature | |
| DS-Core-Propagation-Data | dSCorePropagationData | |
| DS-Heuristics | dSHeuristics | |
| DS-UI-Admin-Maximum | dSUIAdminMaximum | |

| | | |
|---|---|---|
| DS-UI-Admin-Notification | dSUIAdminNotification | |
| DS-UI-Shell-Maximum | dSUIShellMaximum | |
| Dynamic-LDAP-Server | dynamicLDAPServer | |
| EFSPolicy | eFSPolicy | |
| Employee-ID | employeeID | Employee ID |
| Employee-Number | employeeNumber | |
| Employee-Type | employeeType | |
| Enabled | Enabled | |
| Enabled-Connection | enabledConnection | |
| Enrollment-Providers | enrollmentProviders | |
| Entry-TTL | entryTTL | |
| Extended-Attribute-Info | extendedAttributeInfo | |
| Extended-Chars-Allowed | extendedCharsAllowed | |
| Extended-Class-Info | extendedClassInfo | |
| Extension-Name | extensionName | |
| Extra-Columns | extraColumns | |
| Facsimile-Telephone-Number | facsimileTelephoneNumber | Fax Number |
| File-Ext-Priority | fileExtPriority | |
| Flags | flags | |
| Flat-Name | flatName | |
| Force-Logoff | forceLogoff | |
| Foreign-Identifier | foreignIdentifier | |
| Friendly-Names | friendlyNames | |
| From-Entry | fromEntry | |
| From-Server | fromServer | |
| Frs-Computer-Reference | frsComputerReference | |
| Frs-Computer-Reference-BL | frsComputerReferenceBL | |
| FRS-Control-Data-Creation | fRSControlDataCreation | |
| FRS-Control-Inbound-Backlog | fRSControlInboundBacklog | |
| FRS-Control-Outbound-Backlog | fRSControlOutboundBacklog | |
| FRS-Directory-Filter | fRSDirectoryFilter | |
| FRS-DS-Poll | fRSDSPoll | |
| FRS-Extensions | fRSExtensions | |
| FRS-Fault-Condition | fRSFaultCondition | |
| FRS-File-Filter | fRSFileFilter | |
| FRS-Flags | fRSFlags | |
| FRS-Level-Limit | fRSLevelLimit | |
| FRS-Member-Reference | fRSMemberReference | |
| FRS-Member-Reference-BL | fRSMemberReferenceBL | |
| FRS-Partner-Auth-Level | fRSPartnerAuthLevel | |
| FRS-Primary-Member | fRSPrimaryMember | |

| | | |
|---|---|---|
| FRS-Replica-Set-GUID | fRSReplicaSetGUID | |
| FRS-Replica-Set-Type | fRSReplicaSetType | |
| FRS-Root-Path | fRSRootPath | |
| FRS-Root-Security | fRSRootSecurity | |
| FRS-Service-Command | fRSServiceCommand | |
| FRS-Service-Command-Status | fRSServiceCommandStatus | |
| FRS-Staging-Path | fRSStagingPath | |
| FRS-Time-Last-Command | fRSTimeLastCommand | |
| FRS-Time-Last-Config-Change | fRSTimeLastConfigChange | |
| FRS-Update-Timeout | fRSUpdateTimeout | |
| FRS-Version | fRSVersion | |
| FRS-Version-GUID | fRSVersionGUID | |
| FRS-Working-Path | fRSWorkingPath | |
| FSMO-Role-Owner | fSMORoleOwner | |
| Garbage-Coll-Period | garbageCollPeriod | |
| Generated-Connection | generatedConnection | |
| Generation-Qualifier | generationQualifier | Generational Suffix |
| Given-Name | givenName | First Name |
| Global-Address-List | globalAddressList | |
| Governs-ID | governsID | |
| GPC-File-Sys-Path | gPCFileSysPath | |
| GPC-Functionality-Version | gPCFunctionalityVersion | |
| GPC-Machine-Extension-Names | gPCMachineExtensionNames | |
| GPC-User-Extension-Names | gPCUserExtensionNames | |
| GPC-WQL-Filter | gPCWQLFilter | |
| GP-Link | gPLink | |
| GP-Options | gPOptions | |
| Group-Attributes | groupAttributes | |
| Group-Membership-SAM | groupMembershipSAM | |
| Group-Priority | groupPriority | |
| Groups-to-Ignore | groupsToIgnore | |
| Group-Type | groupType | |
| Has-Master-NCs | hasMasterNCs | |
| Has-Partial-Replica-NCs | hasPartialReplicaNCs | |
| Help-Data16 | helpData16 | |
| Help-Data32 | helpData32 | |
| Help-File-Name | helpFileName | |
| Hide-From-AB | hideFromAB | |
| Home-Directory | homeDirectory | Home Folder |
| Home-Drive | homeDrive | Home Drive |
| Phone-Home-Primary | homePhone | Home Phone |
| Address-Home | homePostalAddress | Home Address |

| | | |
|---|---|---|
| host | host | |
| houseIdentifier | houseIdentifier | |
| Icon-Path | iconPath | |
| Implemented-Categories | implementedCategories | |
| IndexedScopes | indexedScopes | |
| Comment | info | Notes |
| Initial-Auth-Incoming | initialAuthIncoming | |
| Initial-Auth-Outgoing | initialAuthOutgoing | |
| Initials | initials | Initials |
| Install-Ui-Level | installUiLevel | |
| Instance-Type | instanceType | |
| International-ISDN-Number | internationalISDNNumber | International ISDN Number (Others) |
| Inter-Site-Topology-Failover | interSiteTopologyFailover | |
| Inter-Site-Topology-Generator | interSiteTopologyGenerator | |
| Inter-Site-Topology-Renew | interSiteTopologyRenew | |
| Invocation-Id | invocationId | |
| Phone-Ip-Primary | ipPhone | IP Phone Number |
| Ipsec-Data | ipsecData | |
| Ipsec-Data-Type | ipsecDataType | |
| Ipsec-Filter-Reference | ipsecFilterReference | |
| Ipsec-ID | ipsecID | |
| Ipsec-ISAKMP-Reference | ipsecISAKMPReference | |
| Ipsec-Name | ipsecName | |
| IPSEC-Negotiation-Policy-Action | iPSECNegotiationPolicyAction | |
| Ipsec-Negotiation-Policy-Reference | ipsecNegotiationPolicyReference | |
| IPSEC-Negotiation-Policy-Type | iPSECNegotiationPolicyType | |
| Ipsec-NFA-Reference | ipsecNFAReference | |
| Ipsec-Owners-Reference | ipsecOwnersReference | |
| Ipsec-Policy-Reference | ipsecPolicyReference | |
| Is-Critical-System-Object | isCriticalSystemObject | |
| Is-Defunct | isDefunct | |
| Is-Deleted | isDeleted | |
| Is-Ephemeral | isEphemeral | |
| Is-Member-Of-Partial-Attribute-Set | isMemberOfPartialAttributeSet | |
| Is-Privilege-Holder | isPrivilegeHolder | |
| Is-Single-Valued | isSingleValued | |
| jpegPhoto | jpegPhoto | |
| Keywords | keywords | Keywords |
| Knowledge-Information | knowledgeInformation | |
| Locality-Name | l | City |
| labeledURI | labeledURI | |

| | | |
|---|---|---|
| Last-Backup-Restoration-Time | lastBackupRestorationTime | |
| Last-Content-Indexed | lastContentIndexed | |
| Last-Known-Parent | lastKnownParent | |
| Last-Logoff | lastLogoff | |
| Last-Logon | lastLogon | |
| Last-Logon-Timestamp | lastLogonTimestamp | |
| Last-Set-Time | lastSetTime | |
| Last-Update-Sequence | lastUpdateSequence | |
| LDAP-Admin-Limits | lDAPAdminLimits | |
| LDAP-Display-Name | lDAPDisplayName | |
| LDAP-IPDeny-List | lDAPIPDenyList | |
| Legacy-Exchange-DN | legacyExchangeDN | |
| Link-ID | linkID | |
| Link-Track-Secret | linkTrackSecret | |
| Lm-Pwd-History | lmPwdHistory | |
| Locale-ID | localeID | |
| Localization-Display-Id | localizationDisplayId | |
| Localized-Description | localizedDescription | |
| Local-Policy-Flags | localPolicyFlags | |
| Local-Policy-Reference | localPolicyReference | |
| Location | location | Location |
| Lockout-Duration | lockoutDuration | |
| Lock-Out-Observation-Window | lockOutObservationWindow | |
| Lockout-Threshold | lockoutThreshold | |
| Lockout-Time | lockoutTime | |
| Logon-Count | logonCount | |
| Logon-Hours | logonHours | |
| Logon-Workstation | logonWorkstation | |
| LSA-Creation-Time | lSACreationTime | |
| LSA-Modified-Count | lSAModifiedCount | |
| Machine-Architecture | machineArchitecture | |
| Machine-Password-Change-Interval | machinePasswordChangeInterval | |
| Machine-Role | machineRole | |
| Machine-Wide-Policy | machineWidePolicy | |
| E-mail-Addresses | mail | E-Mail Address |
| SMTP-Mail-Address | mailAddress | |
| Managed-By | managedBy | Managed By |
| Managed-Objects | managedObjects | |
| Manager | manager | Manager |
| MAPI-ID | mAPIID | |
| Marshalled-Interface | marshalledInterface | |

| | | |
|---|---|---|
| Mastered-By | masteredBy | |
| Max-Pwd-Age | maxPwdAge | |
| Max-Renew-Age | maxRenewAge | |
| Max-Storage | maxStorage | |
| Max-Ticket-Age | maxTicketAge | |
| May-Contain | mayContain | |
| meetingAdvertiseScope | meetingAdvertiseScope | |
| meetingApplication | meetingApplication | |
| meetingBandwidth | meetingBandwidth | |
| meetingBlob | meetingBlob | |
| meetingContactInfo | meetingContactInfo | |
| meetingDescription | meetingDescription | |
| meetingEndTime | meetingEndTime | |
| meetingID | meetingID | |
| meetingIP | meetingIP | |
| meetingIsEncrypted | meetingIsEncrypted | |
| meetingKeyword | meetingKeyword | |
| meetingLanguage | meetingLanguage | |
| meetingLocation | meetingLocation | |
| meetingMaxParticipants | meetingMaxParticipants | |
| meetingName | meetingName | |
| meetingOriginator | meetingOriginator | |
| meetingOwner | meetingOwner | |
| meetingProtocol | meetingProtocol | |
| meetingRating | meetingRating | |
| meetingRecurrence | meetingRecurrence | |
| meetingScope | meetingScope | |
| meetingStartTime | meetingStartTime | |
| meetingType | meetingType | |
| meetingURL | meetingURL | |
| Member | member | Members, Member Queues |
| Is-Member-Of-DL | memberOf | Member Of |
| MHS-OR-Address | mhsORAddress | |
| Other-Name | middleName | Middle Name |
| Min-Pwd-Age | minPwdAge | |
| Min-Pwd-Length | minPwdLength | |
| Min-Ticket-Age | minTicketAge | |
| Phone-Mobile-Primary | mobile | Mobile Number |
| Modified-Count | modifiedCount | |
| Modified-Count-At-Last-Prom | modifiedCountAtLastProm | |
| Modify-Time-Stamp | modifyTimeStamp | |
| Moniker | moniker | |

| | | |
|---|---|---|
| Moniker-Display-Name | monikerDisplayName | |
| Move-Tree-State | moveTreeState | |
| ms-COM-DefaultPartitionLink | msCOM-DefaultPartitionLink | |
| ms-COM-ObjectId | msCOM-ObjectId | |
| ms-COM-PartitionLink | msCOM-PartitionLink | |
| ms-COM-PartitionSetLink | msCOM-PartitionSetLink | |
| ms-COM-UserLink | msCOM-UserLink | |
| ms-COM-UserPartitionSetLink | msCOM-UserPartitionSetLink | |
| Mscope-Id | mscopeId | |
| MS-DRM-Identity-Certificate | msDRM-IdentityCertificate | |
| ms-DS-Additional-Dns-Host-Name | msDS-AdditionalDnsHostName | |
| ms-DS-Additional-Sam-Account-Name | msDS-AdditionalSamAccountName | |
| ms-DS-Allowed-DNS-Suffixes | msDS-AllowedDNSSuffixes | |
| ms-DS-Allowed-To-Delegate-To | msDS-AllowedToDelegateTo | |
| MS-DS-All-Users-Trust-Quota | msDS-AllUsersTrustQuota | |
| ms-DS-Approx-Immed-Subordinates | msDS-Approx-Immed-Subordinates | |
| ms-DS-Auxiliary-Classes | msDS-Auxiliary-Classes | |
| ms-DS-Az-Application-Data | msDS-AzApplicationData | |
| ms-DS-Az-Application-Name | msDS-AzApplicationName | |
| ms-DS-Az-Application-Version | msDS-AzApplicationVersion | |
| ms-DS-Az-Biz-Rule | msDS-AzBizRule | |
| ms-DS-Az-Biz-Rule-Language | msDS-AzBizRuleLanguage | |
| ms-DS-Az-Class-ID | msDS-AzClassId | |
| ms-DS-Az-Domain-Timeout | msDS-AzDomainTimeout | |
| ms-DS-Az-Generate-Audits | msDS-AzGenerateAudits | |
| ms-DS-Az-Last-Imported-Biz-Rule-Path | msDS-AzLastImportedBizRulePath | |
| ms-DS-Az-LDAP-Query | msDS-AzLDAPQuery | |
| ms-DS-Az-Major-Version | msDS-AzMajorVersion | |
| ms-DS-Az-Minor-Version | msDS-AzMinorVersion | |
| ms-DS-Az-Operation-ID | msDS-AzOperationID | |
| ms-DS-Az-Scope-Name | msDS-AzScopeName | |
| ms-DS-Az-Script-Engine-Cache-Max | msDS-AzScriptEngineCacheMax | |
| ms-DS-Az-Script-Timeout | msDS-AzScriptTimeout | |
| ms-DS-Az-Task-Is-Role-Definition | msDS-AzTaskIsRoleDefinition | |
| ms-DS-Behavior-Version | msDS-Behavior-Version | |
| ms-DS-Byte-Array | msDS-ByteArray | |
| ms-DS-Cached-Membership | msDS-Cached-Membership | |
| ms-DS-Cached-Membership-Time-Stamp | msDS-Cached-Membership-Time-Stamp | |
| MS-DS-Consistency-Child-Count | mS-DS-ConsistencyChildCount | |
| MS-DS-Consistency-Guid | mS-DS-ConsistencyGuid | |

| | | |
|---|---|---|
| MS-DS-Creator-SID | mS-DS-CreatorSID | |
| ms-DS-Date-Time | msDS-DateTime | |
| ms-DS-Default-Quota | msDS-DefaultQuota | |
| ms-DS-DnsRootAlias | msDS-DnsRootAlias | |
| ms-DS-Entry-Time-To-Die | msDS-Entry-Time-To-Die | |
| ms-DS-ExecuteScriptPassword | msDS-ExecuteScriptPassword | |
| ms-DS-External-Key | msDS-ExternalKey | |
| ms-DS-External-Store | msDS-ExternalStore | |
| ms-DS-Filter-Containers | msDS-FilterContainers | |
| ms-DS-Has-Domain-NCs | msDS-HasDomainNCs | |
| ms-DS-Has-Instantiated-NCs | msDS-HasInstantiatedNCs | |
| ms-DS-Has-Master-NCs | msDS-hasMasterNCs | |
| ms-DS-Integer | msDS-Integer | |
| ms-DS-IntId | msDS-IntId | |
| ms-DS-KeyVersionNumber | msDS-KeyVersionNumber | |
| ms-DS-Logon-Time-Sync-Interval | msDS-LogonTimeSyncInterval | |
| MS-DS-Machine-Account-Quota | ms-DS-MachineAccountQuota | |
| ms-DS-Mastered-By | msDs-masteredBy | |
| ms-DS-Max-Values | msDs-MaxValues | |
| ms-DS-Members-For-Az-Role | msDS-MembersForAzRole | |
| ms-DS-Members-For-Az-Role-BL | msDS-MembersForAzRoleBL | |
| ms-DS-NC-Repl-Cursors | msDS-NCReplCursors | |
| ms-DS-NC-Replica-Locations | msDS-NC-Replica-Locations | |
| ms-DS-NC-Repl-Inbound-Neighbors | msDS-NCReplInboundNeighbors | |
| ms-DS-NC-Repl-Outbound-Neighbors | msDS-NCReplOutboundNeighbors | |
| ms-DS-Non-Members | msDS-NonMembers | |
| ms-DS-Non-Members-BL | msDS-NonMembersBL | |
| ms-DS-Non-Security-Group-Extra-Classes | msDS-Non-Security-Group-Extra-Classes | |
| ms-DS-Object-Reference | msDS-ObjectReference | |
| ms-DS-Object-Reference-BL | msDS-ObjectReferenceBL | |
| ms-DS-Operations-For-Az-Role | msDS-OperationsForAzRole | |
| ms-DS-Operations-For-Az-Role-BL | msDS-OperationsForAzRoleBL | |
| ms-DS-Operations-For-Az-Task | msDS-OperationsForAzTask | |
| ms-DS-Operations-For-Az-Task-BL | msDS-OperationsForAzTaskBL | |
| ms-DS-Other-Settings | msDS-Other-Settings | |
| MS-DS-Per-User-Trust-Quota | msDS-PerUserTrustQuota | |
| MS-DS-Per-User-Trust-Tombstones-Quota | msDS-PerUserTrustTombstonesQuota | |
| ms-DS-Preferred-GC-Site | msDS-Preferred-GC-Site | |
| ms-DS-Quota-Amount | msDS-QuotaAmount | |
| ms-DS-Quota-Effective | msDS-QuotaEffective | |

| | | |
|---|---|---|
| ms-DS-Quota-Trustee | msDS-QuotaTrustee | |
| ms-DS-Quota-Used | msDS-QuotaUsed | |
| ms-DS-Repl-Attribute-Meta-Data | msDS-ReplAttributeMetaData | |
| MS-DS-Replicates-NC-Reason | mS-DS-ReplicatesNCReason | |
| ms-DS-ReplicationEpoch | msDS-ReplicationEpoch | |
| ms-DS-Replication-Notify-First-DSA-Delay | msDS-Replication-Notify-First-DSA-Delay | |
| ms-DS-Replication-Notify-Subsequent-DSA-Delay | msDS-Replication-Notify-Subsequent-DSA-Delay | |
| ms-DS-Repl-Value-Meta-Data | msDS-ReplValueMetaData | |
| ms-DS-Retired-Repl-NC-Signatures | msDS-RetiredReplNCSignatures | |
| ms-ds-Schema-Extensions | msDs-Schema-Extensions | |
| ms-DS-SD-Reference-Domain | msDS-SDReferenceDomain | |
| ms-DS-Security-Group-Extra-Classes | msDS-Security-Group-Extra-Classes | |
| ms-DS-Settings | msDS-Settings | |
| ms-DS-Site-Affinity | msDS-Site-Affinity | |
| ms-DS-SPN-Suffixes | msDS-SPNSuffixes | |
| ms-DS-Tasks-For-Az-Role | msDS-TasksForAzRole | |
| ms-DS-Tasks-For-Az-Role-BL | msDS-TasksForAzRoleBL | |
| ms-DS-Tasks-For-Az-Task | msDS-TasksForAzTask | |
| ms-DS-Tasks-For-Az-Task-BL | msDS-TasksForAzTaskBL | |
| ms-DS-Tombstone-Quota-Factor | msDS-TombstoneQuotaFactor | |
| ms-DS-Top-Quota-Usage | msDS-TopQuotaUsage | |
| ms-DS-Trust-Forest-Trust-Info | msDS-TrustForestTrustInfo | |
| ms-DS-UpdateScript | msDS-UpdateScript | |
| ms-DS-User-Account-Control-Computed | msDS-User-Account-Control-Computed | |
| ms-Exch-Assistant-Name | msExchAssistantName | |
| ms-Exch-House-Identifier | msExchHouseIdentifier | |
| ms-Exch-LabeledURI | msExchLabeledURI | |
| ms-FRS-Hub-Member | msFRS-Hub-Member | |
| ms-FRS-Topology-Pref | msFRS-Topology-Pref | |
| ms-ieee-80211-Data | msieee80211-Data | |
| ms-ieee-80211-Data-Type | msieee80211-DataType | |
| ms-ieee-80211-ID | msieee80211-ID | |
| Msi-File-List | msiFileList | |
| ms-IIS-FTP-Dir | msIIS-FTPDir | |
| ms-IIS-FTP-Root | msIIS-FTPRoot | |
| Msi-Script | msiScript | |
| Msi-Script-Name | msiScriptName | |
| Msi-Script-Path | msiScriptPath | |
| Msi-Script-Size | msiScriptSize | |
| MSMQ-Authenticate | mSMQAuthenticate | |

| | | |
|---|---|---|
| MSMQ-Base-Priority | mSMQBasePriority | |
| MSMQ-Computer-Type | mSMQComputerType | |
| MSMQ-Computer-Type-Ex | mSMQComputerTypeEx | |
| MSMQ-Cost | mSMQCost | |
| MSMQ-CSP-Name | mSMQCSPName | |
| MSMQ-Dependent-Client-Service | mSMQDependentClientService | |
| MSMQ-Dependent-Client-Services | mSMQDependentClientServices | |
| MSMQ-Digests | mSMQDigests | |
| MSMQ-Digests-Mig | mSMQDigestsMig | |
| MSMQ-Ds-Service | mSMQDsService | |
| MSMQ-Ds-Services | mSMQDsServices | |
| MSMQ-Encrypt-Key | mSMQEncryptKey | |
| MSMQ-Foreign | mSMQForeign | |
| MSMQ-In-Routing-Servers | mSMQInRoutingServers | |
| MSMQ-Interval1 | mSMQInterval1 | |
| MSMQ-Interval2 | mSMQInterval2 | |
| MSMQ-Journal | mSMQJournal | |
| MSMQ-Journal-Quota | mSMQJournalQuota | |
| MSMQ-Label | mSMQLabel | |
| MSMQ-Label-Ex | mSMQLabelEx | |
| MSMQ-Long-Lived | mSMQLongLived | |
| MSMQ-Migrated | mSMQMigrated | |
| MSMQ-Multicast-Address | MSMQ-MulticastAddress | |
| MSMQ-Name-Style | mSMQNameStyle | |
| MSMQ-Nt4-Flags | mSMQNt4Flags | |
| MSMQ-Nt4-Stub | mSMQNt4Stub | |
| MSMQ-OS-Type | mSMQOSType | |
| MSMQ-Out-Routing-Servers | mSMQOutRoutingServers | |
| MSMQ-Owner-ID | mSMQOwnerID | |
| MSMQ-Prev-Site-Gates | mSMQPrevSiteGates | |
| MSMQ-Privacy-Level | mSMQPrivacyLevel | |
| MSMQ-QM-ID | mSMQQMID | |
| MSMQ-Queue-Journal-Quota | mSMQQueueJournalQuota | |
| MSMQ-Queue-Name-Ext | mSMQQueueNameExt | |
| MSMQ-Queue-Quota | mSMQQueueQuota | |
| MSMQ-Queue-Type | mSMQQueueType | |
| MSMQ-Quota | mSMQQuota | |
| MSMQ-Recipient-FormatName | msMQ-Recipient-FormatName | Format Name |
| MSMQ-Routing-Service | mSMQRoutingService | |
| MSMQ-Routing-Services | mSMQRoutingServices | |
| MSMQ-Secured-Source | MSMQ-SecuredSource | |
| MSMQ-Services | mSMQServices | |

| | | |
|---|---|---|
| MSMQ-Service-Type | mSMQServiceType | |
| MSMQ-Sign-Certificates | mSMQSignCertificates | |
| MSMQ-Sign-Certificates-Mig | mSMQSignCertificatesMig | |
| MSMQ-Sign-Key | mSMQSignKey | |
| MSMQ-Site-1 | mSMQSite1 | |
| MSMQ-Site-2 | mSMQSite2 | |
| MSMQ-Site-Foreign | mSMQSiteForeign | |
| MSMQ-Site-Gates | mSMQSiteGates | |
| MSMQ-Site-Gates-Mig | mSMQSiteGatesMig | |
| MSMQ-Site-ID | mSMQSiteID | |
| MSMQ-Site-Name | mSMQSiteName | |
| MSMQ-Site-Name-Ex | mSMQSiteNameEx | |
| MSMQ-Sites | mSMQSites | |
| MSMQ-Transactional | mSMQTransactional | |
| MSMQ-User-Sid | mSMQUserSid | |
| MSMQ-Version | mSMQVersion | |
| msNPAllowDialin | msNPAllowDialin | |
| msNPCalledStationID | msNPCalledStationID | |
| msNPCallingStationID | msNPCallingStationID | |
| msNPSavedCallingStationID | msNPSavedCallingStationID | |
| ms-PKI-Certificate-Application-Policy | msPKI-Certificate-Application-Policy | |
| ms-PKI-Certificate-Name-Flag | msPKI-Certificate-Name-Flag | |
| ms-PKI-Certificate-Policy | msPKI-Certificate-Policy | |
| ms-PKI-Cert-Template-OID | msPKI-Cert-Template-OID | |
| ms-PKI-Enrollment-Flag | msPKI-Enrollment-Flag | |
| ms-PKI-Minimal-Key-Size | msPKI-Minimal-Key-Size | |
| ms-PKI-OID-Attribute | msPKI-OID-Attribute | |
| ms-PKI-OID-CPS | msPKI-OID-CPS | |
| ms-PKI-OID-LocalizedName | msPKI-OIDLocalizedName | |
| ms-PKI-OID-User-Notice | msPKI-OID-User-Notice | |
| ms-PKI-Private-Key-Flag | msPKI-Private-Key-Flag | |
| ms-PKI-RA-Application-Policies | msPKI-RA-Application-Policies | |
| ms-PKI-RA-Policies | msPKI-RA-Policies | |
| ms-PKI-RA-Signature | msPKI-RA-Signature | |
| ms-PKI-Supersede-Templates | msPKI-Supersede-Templates | |
| ms-PKI-Template-Minor-Revision | msPKI-Template-Minor-Revision | |
| ms-PKI-Template-Schema-Version | msPKI-Template-Schema-Version | |
| msRADIUSCallbackNumber | msRADIUSCallbackNumber | |
| msRADIUSFramedIPAddress | msRADIUSFramedIPAddress | |
| msRADIUSFramedRoute | msRADIUSFramedRoute | |
| msRADIUSServiceType | msRADIUSServiceType | |
| msRASSavedCallbackNumber | msRASSavedCallbackNumber | |

| | | |
|---|---|---|
| msRASSavedFramedIPAddress | msRASSavedFramedIPAddress | |
| msRASSavedFramedRoute | msRASSavedFramedRoute | |
| ms-RRAS-Attribute | msRRASAttribute | |
| ms-RRAS-Vendor-Attribute-Entry | msRRASVendorAttributeEntry | |
| MS-SQL-Alias | mS-SQL-Alias | |
| MS-SQL-AllowAnonymousSubscription | mS-SQL-AllowAnonymousSubscription | |
| MS-SQL-AllowImmediateUpdatingSubscription | mS-SQL-AllowImmediateUpdatingSubscription | |
| MS-SQL-AllowKnownPullSubscription | mS-SQL-AllowKnownPullSubscription | |
| MS-SQL-AllowQueuedUpdatingSubscription | mS-SQL-AllowQueuedUpdatingSubscription | |
| MS-SQL-AllowSnapshotFilesFTPDownloading | mS-SQL-AllowSnapshotFilesFTPDownloading | |
| MS-SQL-AppleTalk | mS-SQL-AppleTalk | |
| MS-SQL-Applications | mS-SQL-Applications | |
| MS-SQL-Build | mS-SQL-Build | |
| MS-SQL-CharacterSet | mS-SQL-CharacterSet | |
| MS-SQL-Clustered | mS-SQL-Clustered | |
| MS-SQL-ConnectionURL | mS-SQL-ConnectionURL | |
| MS-SQL-Contact | mS-SQL-Contact | |
| MS-SQL-CreationDate | mS-SQL-CreationDate | |
| MS-SQL-Database | mS-SQL-Database | |
| MS-SQL-Description | mS-SQL-Description | |
| MS-SQL-GPSHeight | mS-SQL-GPSHeight | |
| MS-SQL-GPSLatitude | mS-SQL-GPSLatitude | |
| MS-SQL-GPSLongitude | mS-SQL-GPSLongitude | |
| MS-SQL-InformationDirectory | mS-SQL-InformationDirectory | |
| MS-SQL-InformationURL | mS-SQL-InformationURL | |
| MS-SQL-Keywords | mS-SQL-Keywords | |
| MS-SQL-Language | mS-SQL-Language | |
| MS-SQL-LastBackupDate | mS-SQL-LastBackupDate | |
| MS-SQL-LastDiagnosticDate | mS-SQL-LastDiagnosticDate | |
| MS-SQL-LastUpdatedDate | mS-SQL-LastUpdatedDate | |
| MS-SQL-Location | mS-SQL-Location | |
| MS-SQL-Memory | mS-SQL-Memory | |
| MS-SQL-MultiProtocol | mS-SQL-MultiProtocol | |
| MS-SQL-Name | mS-SQL-Name | |
| MS-SQL-NamedPipe | mS-SQL-NamedPipe | |
| MS-SQL-PublicationURL | mS-SQL-PublicationURL | |
| MS-SQL-Publisher | mS-SQL-Publisher | |
| MS-SQL-RegisteredOwner | mS-SQL-RegisteredOwner | |
| MS-SQL-ServiceAccount | mS-SQL-ServiceAccount | |
| MS-SQL-Size | mS-SQL-Size | |

| | | |
|---|---|---|
| MS-SQL-SortOrder | mS-SQL-SortOrder | |
| MS-SQL-SPX | mS-SQL-SPX | |
| MS-SQL-Status | mS-SQL-Status | |
| MS-SQL-TCPIP | mS-SQL-TCPIP | |
| MS-SQL-ThirdParty | mS-SQL-ThirdParty | |
| MS-SQL-Type | mS-SQL-Type | |
| MS-SQL-UnicodeSortOrder | mS-SQL-UnicodeSortOrder | |
| MS-SQL-Version | mS-SQL-Version | |
| MS-SQL-Vines | mS-SQL-Vines | |
| ms-TAPI-Conference-Blob | msTAPI-ConferenceBlob | |
| ms-TAPI-Ip-Address | msTAPI-IpAddress | |
| ms-TAPI-Protocol-Id | msTAPI-ProtocolId | |
| ms-TAPI-Unique-Identifier | msTAPI-uid | |
| ms-WMI-Author | msWMI-Author | |
| ms-WMI-ChangeDate | msWMI-ChangeDate | |
| ms-WMI-Class | msWMI-Class | |
| ms-WMI-ClassDefinition | msWMI-ClassDefinition | |
| ms-WMI-CreationDate | msWMI-CreationDate | |
| ms-WMI-Genus | msWMI-Genus | |
| ms-WMI-ID | msWMI-ID | |
| ms-WMI-int8Default | msWMI-Int8Default | |
| ms-WMI-int8Max | msWMI-Int8Max | |
| ms-WMI-int8Min | msWMI-Int8Min | |
| ms-WMI-int8ValidValues | msWMI-Int8ValidValues | |
| ms-WMI-intDefault | msWMI-IntDefault | |
| ms-WMI-intFlags1 | msWMI-intFlags1 | |
| ms-WMI-intFlags2 | msWMI-intFlags2 | |
| ms-WMI-intFlags3 | msWMI-intFlags3 | |
| ms-WMI-intFlags4 | msWMI-intFlags4 | |
| ms-WMI-intMax | msWMI-IntMax | |
| ms-WMI-intMin | msWMI-IntMin | |
| ms-WMI-intValidValues | msWMI-IntValidValues | |
| ms-WMI-Mof | msWMI-Mof | |
| ms-WMI-Name | msWMI-Name | |
| ms-WMI-NormalizedClass | msWMI-NormalizedClass | |
| ms-WMI-Parm1 | msWMI-Parm1 | |
| ms-WMI-Parm2 | msWMI-Parm2 | |
| ms-WMI-Parm3 | msWMI-Parm3 | |
| ms-WMI-Parm4 | msWMI-Parm4 | |
| ms-WMI-PropertyName | msWMI-PropertyName | |
| ms-WMI-Query | msWMI-Query | |
| ms-WMI-QueryLanguage | msWMI-QueryLanguage | |

| | | |
|---|---|---|
| ms-WMI-ScopeGuid | msWMI-ScopeGuid | |
| ms-WMI-SourceOrganization | msWMI-SourceOrganization | |
| ms-WMI-stringDefault | msWMI-StringDefault | |
| ms-WMI-stringValidValues | msWMI-StringValidValues | |
| ms-WMI-TargetClass | msWMI-TargetClass | |
| ms-WMI-TargetNameSpace | msWMI-TargetNameSpace | |
| ms-WMI-TargetObject | msWMI-TargetObject | |
| ms-WMI-TargetPath | msWMI-TargetPath | |
| ms-WMI-TargetType | msWMI-TargetType | |
| Must-Contain | mustContain | |
| RDN | name | |
| Name-Service-Flags | nameServiceFlags | |
| NC-Name | nCName | |
| NETBIOS-Name | nETBIOSName | |
| netboot-Allow-New-Clients | netbootAllowNewClients | |
| netboot-Answer-Only-Valid-Clients | netbootAnswerOnlyValidClients | |
| netboot-Answer-Requests | netbootAnswerRequests | |
| netboot-Current-Client-Count | netbootCurrentClientCount | |
| Netboot-GUID | netbootGUID | |
| Netboot-Initialization | netbootInitialization | |
| netboot-IntelliMirror-OSes | netbootIntelliMirrorOSes | |
| netboot-Limit-Clients | netbootLimitClients | |
| netboot-Locally-Installed-OSes | netbootLocallyInstalledOSes | |
| Netboot-Machine-File-Path | netbootMachineFilePath | |
| netboot-Max-Clients | netbootMaxClients | |
| Netboot-Mirror-Data-File | netbootMirrorDataFile | |
| netboot-New-Machine-Naming-Policy | netbootNewMachineNamingPolicy | |
| netboot-New-Machine-OU | netbootNewMachineOU | |
| netboot-SCP-BL | netbootSCPBL | |
| netboot-Server | netbootServer | |
| Netboot-SIF-File | netbootSIFFile | |
| netboot-Tools | netbootTools | |
| Network-Address | networkAddress | |
| Next-Level-Store | nextLevelStore | |
| Next-Rid | nextRid | |
| Non-Security-Member | nonSecurityMember | |
| Non-Security-Member-BL | nonSecurityMemberBL | |
| Additional-Information | notes | |
| Notification-List | notificationList | |
| NT-Group-Members | nTGroupMembers | |

| | | |
|---|---|---|
| NT-Mixed-Domain | nTMixedDomain | |
| Nt-Pwd-History | ntPwdHistory | |
| NT-Security-Descriptor | nTSecurityDescriptor | |
| Organization-Name | o | |
| Object-Category | objectCategory | |
| Object-Class | objectClass | |
| Object-Class-Category | objectClassCategory | |
| Object-Classes | objectClasses | |
| Object-Count | objectCount | |
| Object-Guid | objectGUID | |
| Object-Sid | objectSid | |
| Object-Version | objectVersion | |
| OEM-Information | oEMInformation | |
| OM-Object-Class | oMObjectClass | |
| OM-Syntax | oMSyntax | |
| OMT-Guid | oMTGuid | |
| OMT-Indx-Guid | oMTIndxGuid | |
| Operating-System | operatingSystem | Operating System |
| Operating-System-Hotfix | operatingSystemHotfix | |
| Operating-System-Service-Pack | operatingSystemServicePack | |
| Operating-System-Version | operatingSystemVersion | Operating System Version |
| Operator-Count | operatorCount | |
| Option-Description | optionDescription | |
| Options | options | |
| Options-Location | optionsLocation | |
| organizationalStatus | organizationalStatus | |
| Original-Display-Table | originalDisplayTable | |
| Original-Display-Table-MSDOS | originalDisplayTableMSDOS | |
| Phone-Fax-Other | otherFacsimileTelephoneNumber | Fax Number (Others) |
| Phone-Home-Other | otherHomePhone | Home Phone Number (Others), Home Phone (Others) |
| Phone-Ip-Other | otherIpPhone | IP Phone Number (Others) |
| Other-Login-Workstations | otherLoginWorkstations | |
| Other-Mailbox | otherMailbox | E-Mail Address (Others) |
| Phone-Mobile-Other | otherMobile | Mobile Number (Others) |
| Phone-Pager-Other | otherPager | Pager Number (Others) |
| Phone-Office-Other | otherTelephone | Phone Number (Others) |
| Other-Well-Known-Objects | otherWellKnownObjects | |
| Organizational-Unit-Name | ou | Name |
| Owner | owner | |
| ms-Exch-Owner-BL | ownerBL | |
| Package-Flags | packageFlags | |
| Package-Name | packageName | |

| | | |
|---|---|---|
| Package-Type | packageType | |
| Phone-Pager-Primary | pager | Pager Number |
| Parent-CA | parentCA | |
| Parent-CA-Certificate-Chain | parentCACertificateChain | |
| Parent-GUID | parentGUID | |
| Partial-Attribute-Deletion-List | partialAttributeDeletionList | |
| Partial-Attribute-Set | partialAttributeSet | |
| Pek-Key-Change-Interval | pekKeyChangeInterval | |
| Pek-List | pekList | |
| Pending-CA-Certificates | pendingCACertificates | |
| Pending-Parent-CA | pendingParentCA | |
| Per-Msg-Dialog-Display-Table | perMsgDialogDisplayTable | |
| Per-Recip-Dialog-Display-Table | perRecipDialogDisplayTable | |
| Personal-Title | personalTitle | Title |
| photo | photo | |
| Physical-Delivery-Office-Name | physicalDeliveryOfficeName | Office Location |
| Physical-Location-Object | physicalLocationObject | |
| PKI-Critical-Extensions | pKICriticalExtensions | |
| PKI-Default-CSPs | pKIDefaultCSPs | |
| PKI-Default-Key-Spec | pKIDefaultKeySpec | |
| PKI-Enrollment-Access | pKIEnrollmentAccess | |
| PKI-Expiration-Period | pKIExpirationPeriod | |
| PKI-Extended-Key-Usage | pKIExtendedKeyUsage | |
| PKI-Key-Usage | pKIKeyUsage | |
| PKI-Max-Issuing-Depth | pKIMaxIssuingDepth | |
| PKI-Overlap-Period | pKIOverlapPeriod | |
| PKT | pKT | |
| PKT-Guid | pKTGuid | |
| Policy-Replication-Flags | policyReplicationFlags | |
| Port-Name | portName | Port |
| Possible-Inferiors | possibleInferiors | |
| Poss-Superiors | possSuperiors | |
| Postal-Address | postalAddress | |
| Postal-Code | postalCode | ZIP/Postal Code |
| Post-Office-Box | postOfficeBox | Post Office Box |
| Preferred-Delivery-Method | preferredDeliveryMethod | |
| preferredLanguage | preferredLanguage | |
| Preferred-OU | preferredOU | |
| Prefix-Map | prefixMap | |
| Presentation-Address | presentationAddress | |
| Previous-CA-Certificates | previousCACertificates | |
| Previous-Parent-CA | previousParentCA | |

| | | |
|---|---|---|
| Primary-Group-ID | primaryGroupID | |
| Primary-Group-Token | primaryGroupToken | |
| Phone-ISDN-Primary | primaryInternationalISDNNumber | International ISDN Number |
| Telex-Primary | primaryTelexNumber | Telex Number |
| Print-Attributes | printAttributes | |
| Print-Bin-Names | printBinNames | Input Trays |
| Print-Collate | printCollate | Supports Collation |
| Print-Color | printColor | Supports Color Printing |
| Print-Duplex-Supported | printDuplexSupported | Supports Double-sided Printing |
| Print-End-Time | printEndTime | |
| Printer-Name | printerName | Name |
| Print-Form-Name | printFormName | |
| Print-Keep-Printed-Jobs | printKeepPrintedJobs | |
| Print-Language | printLanguage | Printer Language |
| Print-MAC-Address | printMACAddress | |
| Print-Max-Copies | printMaxCopies | |
| Print-Max-Resolution-Supported | printMaxResolutionSupported | Maximum Resolution |
| Print-Max-X-Extent | printMaxXExtent | |
| Print-Max-Y-Extent | printMaxYExtent | |
| Print-Media-Ready | printMediaReady | Paper Available |
| Print-Media-Supported | printMediaSupported | Paper Types Supported |
| Print-Memory | printMemory | Installed Memory |
| Print-Min-X-Extent | printMinXExtent | |
| Print-Min-Y-Extent | printMinYExtent | |
| Print-Network-Address | printNetworkAddress | |
| Print-Notify | printNotify | |
| Print-Number-Up | printNumberUp | |
| Print-Orientations-Supported | printOrientationsSupported | |
| Print-Owner | printOwner | Owner Name |
| Print-Pages-Per-Minute | printPagesPerMinute | Pages per Minute |
| Print-Rate | printRate | Speed |
| Print-Rate-Unit | printRateUnit | Speed Units |
| Print-Separator-File | printSeparatorFile | |
| Print-Share-Name | printShareName | Share Name |
| Print-Spooling | printSpooling | |
| Print-Stapling-Supported | printStaplingSupported | Supports Stapling |
| Print-Start-Time | printStartTime | |
| Print-Status | printStatus | |
| Priority | priority | |
| Prior-Set-Time | priorSetTime | |
| Prior-Value | priorValue | |
| Private-Key | privateKey | |

| | | |
|---|---|---|
| Privilege-Attributes | privilegeAttributes | |
| Privilege-Display-Name | privilegeDisplayName | |
| Privilege-Holder | privilegeHolder | |
| Privilege-Value | privilegeValue | |
| Product-Code | productCode | |
| Profile-Path | profilePath | |
| Proxied-Object-Name | proxiedObjectName | |
| Proxy-Addresses | proxyAddresses | |
| Proxy-Generation-Enabled | proxyGenerationEnabled | |
| Proxy-Lifetime | proxyLifetime | |
| Public-Key-Policy | publicKeyPolicy | |
| Purported-Search | purportedSearch | |
| Pwd-History-Length | pwdHistoryLength | |
| Pwd-Last-Set | pwdLastSet | |
| Pwd-Properties | pwdProperties | |
| Quality-Of-Service | qualityOfService | |
| Query-Filter | queryFilter | |
| QueryPoint | queryPoint | |
| Query-Policy-BL | queryPolicyBL | |
| Query-Policy-Object | queryPolicyObject | |
| Range-Lower | rangeLower | |
| Range-Upper | rangeUpper | |
| RDN-Att-ID | rDNAttID | |
| Registered-Address | registeredAddress | |
| Remote-Server-Name | remoteServerName | |
| Remote-Source | remoteSource | |
| Remote-Source-Type | remoteSourceType | |
| Remote-Storage-GUID | remoteStorageGUID | |
| Replica-Source | replicaSource | |
| Repl-Interval | replInterval | |
| Repl-Property-Meta-Data | replPropertyMetaData | |
| Repl-Topology-Stay-Of-Execution | replTopologyStayOfExecution | |
| Repl-UpToDate-Vector | replUpToDateVector | |
| Reps-From | repsFrom | |
| Reps-To | repsTo | |
| Required-Categories | requiredCategories | |
| Retired-Repl-DSA-Signatures | retiredReplDSASignatures | |
| Revision | revision | |
| Rid | rid | |
| RID-Allocation-Pool | rIDAllocationPool | |
| RID-Available-Pool | rIDAvailablePool | |
| RID-Manager-Reference | rIDManagerReference | |

| | | |
|---|---|---|
| RID-Next-RID | rIDNextRID | |
| RID-Previous-Allocation-Pool | rIDPreviousAllocationPool | |
| RID-Set-References | rIDSetReferences | |
| RID-Used-Pool | rIDUsedPool | |
| Rights-Guid | rightsGuid | |
| Role-Occupant | roleOccupant | |
| roomNumber | roomNumber | |
| Root-Trust | rootTrust | |
| rpc-Ns-Annotation | rpcNsAnnotation | |
| rpc-Ns-Bindings | rpcNsBindings | |
| rpc-Ns-Codeset | rpcNsCodeset | |
| rpc-Ns-Entry-Flags | rpcNsEntryFlags | |
| rpc-Ns-Group | rpcNsGroup | |
| rpc-Ns-Interface-ID | rpcNsInterfaceID | |
| rpc-Ns-Object-ID | rpcNsObjectID | |
| rpc-Ns-Priority | rpcNsPriority | |
| rpc-Ns-Profile-Entry | rpcNsProfileEntry | |
| rpc-Ns-Transfer-Syntax | rpcNsTransferSyntax | |
| SAM-Account-Name | sAMAccountName | Computer name (pre-Windows 2000), Group name (pre-Windows 2000), Logon Name (pre-Windows 2000) |
| SAM-Account-Type | sAMAccountType | |
| Schedule | schedule | |
| Schema-Flags-Ex | schemaFlagsEx | |
| Schema-ID-GUID | schemaIDGUID | |
| Schema-Info | schemaInfo | |
| Schema-Update | schemaUpdate | |
| Schema-Version | schemaVersion | |
| Scope-Flags | scopeFlags | |
| Script-Path | scriptPath | |
| SD-Rights-Effective | sDRightsEffective | |
| Search-Flags | searchFlags | |
| Search-Guide | searchGuide | |
| secretary | secretary | |
| Security-Identifier | securityIdentifier | |
| See-Also | seeAlso | |
| Seq-Notification | seqNotification | |
| Serial-Number | serialNumber | |
| Server-Name | serverName | Server Name |
| Server-Reference | serverReference | |
| Server-Reference-BL | serverReferenceBL | |
| Server-Role | serverRole | |
| Server-State | serverState | |

| | | |
|---|---|---|
| Service-Binding-Information | serviceBindingInformation | |
| Service-Class-ID | serviceClassID | |
| Service-Class-Info | serviceClassInfo | |
| Service-Class-Name | serviceClassName | |
| Service-DNS-Name | serviceDNSName | |
| Service-DNS-Name-Type | serviceDNSNameType | |
| Service-Instance-Version | serviceInstanceVersion | |
| Service-Principal-Name | servicePrincipalName | |
| Setup-Command | setupCommand | |
| Shell-Context-Menu | shellContextMenu | |
| Shell-Property-Pages | shellPropertyPages | |
| Short-Server-Name | shortServerName | |
| Show-In-Address-Book | showInAddressBook | |
| Show-In-Advanced-View-Only | showInAdvancedViewOnly | |
| SID-History | sIDHistory | |
| Signature-Algorithms | signatureAlgorithms | |
| Site-GUID | siteGUID | |
| Site-Link-List | siteLinkList | |
| Site-List | siteList | |
| Site-Object | siteObject | |
| Site-Object-BL | siteObjectBL | |
| Site-Server | siteServer | |
| Surname | sn | Last Name |
| SPN-Mappings | sPNMappings | |
| State-Or-Province-Name | st | State/Province |
| Street-Address | street | |
| Address | streetAddress | Street Address |
| Structural-Object-Class | structuralObjectClass | |
| Sub-Class-Of | subClassOf | |
| Sub-Refs | subRefs | |
| SubSchemaSubEntry | subSchemaSubEntry | |
| Superior-DNS-Root | superiorDNSRoot | |
| Super-Scope-Description | superScopeDescription | |
| Super-Scopes | superScopes | |
| Supplemental-Credentials | supplementalCredentials | |
| Supported-Application-Context | supportedApplicationContext | |
| Sync-Attributes | syncAttributes | |
| Sync-Membership | syncMembership | |
| Sync-With-Object | syncWithObject | |
| Sync-With-SID | syncWithSID | |
| System-Auxiliary-Class | systemAuxiliaryClass | |
| System-Flags | systemFlags | |

| | | |
|---|---|---|
| System-May-Contain | systemMayContain | |
| System-Must-Contain | systemMustContain | |
| System-Only | systemOnly | |
| System-Poss-Superiors | systemPossSuperiors | |
| Telephone-Number | telephoneNumber | Telephone Number |
| Teletex-Terminal-Identifier | teletexTerminalIdentifier | |
| Telex-Number | telexNumber | Telex Number (Others) |
| Template-Roots | templateRoots | |
| Terminal-Server | terminalServer | |
| Text-Encoded-OR-Address | textEncodedORAddress | |
| Logo | thumbnailLogo | |
| Picture | thumbnailPhoto | |
| Time-Refresh | timeRefresh | |
| Time-Vol-Change | timeVolChange | |
| Title | title | Job Title |
| Token-Groups | tokenGroups | |
| Token-Groups-Global-And-Universal | tokenGroupsGlobalAndUniversal | |
| Token-Groups-No-GC-Acceptable | tokenGroupsNoGCAcceptable | |
| Tombstone-Lifetime | tombstoneLifetime | |
| Transport-Address-Attribute | transportAddressAttribute | |
| Transport-DLL-Name | transportDLLName | |
| Transport-Type | transportType | |
| Treat-As-Leaf | treatAsLeaf | |
| Tree-Name | treeName | |
| Trust-Attributes | trustAttributes | |
| Trust-Auth-Incoming | trustAuthIncoming | |
| Trust-Auth-Outgoing | trustAuthOutgoing | |
| Trust-Direction | trustDirection | |
| Trust-Parent | trustParent | |
| Trust-Partner | trustPartner | |
| Trust-Posix-Offset | trustPosixOffset | |
| Trust-Type | trustType | |
| UAS-Compat | uASCompat | |
| uid | uid | |
| UNC-Name | uNCName | Network Name, Network Path |
| Unicode-Pwd | unicodePwd | |
| uniqueIdentifier | uniqueIdentifier | |
| uniqueMember | uniqueMember | |
| unstructuredAddress | unstructuredAddress | |
| unstructuredName | unstructuredName | |
| Upgrade-Product-Code | upgradeProductCode | |

| | | |
|---|---|---|
| UPN-Suffixes | uPNSuffixes | |
| WWW-Page-Other | url | Web Page Address, Web Page Address (Others) |
| User-Account-Control | userAccountControl | |
| User-Cert | userCert | |
| X509-Cert | userCertificate | |
| userClass | userClass | |
| User-Parameters | userParameters | |
| User-Password | userPassword | |
| userPKCS12 | userPKCS12 | |
| User-Principal-Name | userPrincipalName | Logon Name |
| User-Shared-Folder | userSharedFolder | |
| User-Shared-Folder-Other | userSharedFolderOther | |
| User-SMIME-Certificate | userSMIMECertificate | |
| User-Workstations | userWorkstations | Logon Workstations |
| USN-Changed | uSNChanged | |
| USN-Created | uSNCreated | |
| USN-DSA-Last-Obj-Removed | uSNDSALastObjRemoved | |
| USN-Intersite | USNIntersite | |
| USN-Last-Obj-Rem | uSNLastObjRem | |
| USN-Source | uSNSource | |
| Valid-Accesses | validAccesses | |
| Vendor | vendor | |
| Version-Number | versionNumber | Object Version |
| Version-Number-Hi | versionNumberHi | |
| Version-Number-Lo | versionNumberLo | |
| Vol-Table-GUID | volTableGUID | |
| Vol-Table-Idx-GUID | volTableIdxGUID | |
| Volume-Count | volumeCount | |
| Wbem-Path | wbemPath | |
| Well-Known-Objects | wellKnownObjects | |
| When-Changed | whenChanged | |
| When-Created | whenCreated | |
| Winsock-Addresses | winsockAddresses | |
| WWW-Home-Page | wWWHomePage | Web Page Address |
| X121-Address | x121Address | |
| x500uniqueIdentifier | x500uniqueIdentifier | |

**Attributes Display, Class View Name Mappings1**

| **Class** | **attributeDisplayNames** |
|---|---|
| computer | samAccountName,Computer name (pre-Windows 2000);operatingSystemVersion,Operating System Version;operatingSystem,Operating System;managedBy,Managed By;description,Description;cn,Name; |

| | |
|---|---|
| contact | wWWHomePage,Web Page Address;url,Web Page Address (Others);title,Job Title;telexNumber,Telex Number (Others);telephoneNumber,Telephone Number;streetAddress,Street Address;st,State/Province;sn,Last Name;primaryTelexNumber,Telex Number;primaryInternationalISDNNumber,International ISDN Number;postOfficeBox,Post Office Box;postalCode,ZIP/Postal Code;physicalDeliveryOfficeName,Office Location;personalTitle,Title;pager,Pager Number;otherTelephone,Phone Number (Others);otherPager,Pager Number (Others);otherMobile,Mobile Number (Others);otherMailbox,E-Mail Address (Others);otherIpPhone,IP Phone Number (Others);otherHomePhone,Home Phone Number (Others);otherFacsimileTelephoneNumber,Fax Number (Others);mobile,Mobile Number;middleName,Middle Name;memberOf,Member Of;manager,Manager;mail,E-Mail Address;l,City;ipPhone,IP Phone Number;internationalISDNNumber,International ISDN Number (Others);info,Notes;initials,Initials;homePostalAddress,Home Address;homePhone,Home Phone;givenName,First Name;generationQualifier,Generational Suffix;facsimileTelephoneNumber,Fax Number;employeeID,Employee ID;division,Division;distinguishedName,X500 Distinguished Name;displayName,Display Name;directReports,Direct Reports;description,Description;department,Department;company,Company;comment,Comment;cn,Name;co,Country;c,Country Abbreviation;assistant,Assistant; |
| container | description,Description;cn,Name; |
| domainDNS | description,Description;dc,Name; |
| domainPolicy | |
| foreignSecurityPrincipal | description,Description;cn,Name; |
| group | wWWHomePage,Web Page Address;url,Web Page Address (Others);samAccountName,Group name (pre-Windows 2000);physicalDeliveryOfficeName,Office Location;info,Notes;member,Members;managedBy,Managed By;l,City;distinguishedName,X500 Distinguished Name;description,Description;c,Country Abbreviation;cn,Name; |
| inetOrgPerson | assistant,Assistant;cn,Name;c,Country Abbreviation;co,Country;comment,Comment;company,Company;department,Department;description,Description;directReports,Direct Reports;distinguishedName,X500 Distinguished Name;division,Division;employeeID,Employee ID;facsimileTelephoneNumber,Fax Number;generationQualifier,Generational Suffix;givenName,First Name;homeDirectory,Home Folder;homeDrive,Home Drive;homePhone,Home Phone;homePostalAddress,Home Address;initials,Initials;internationalISDNNumber,International ISDN Number (Others);ipPhone,IP Phone Number;l,City;mail,E-Mail Address;manager,Manager;memberOf,Member Of;middleName,Middle Name;mobile,Mobile Number;info,Notes;otherFacsimileTelephoneNumber,Fax Number (Others);otherHomePhone,Home Phone (Others);otherIpPhone,IP Phone Number (Others);otherMailbox,E-Mail Address (Others);otherMobile,Mobile Number (Others);otherPager,Pager Number (Others);otherTelephone,Phone Number (Others);pager,Pager Number;personalTitle,Title;physicalDeliveryOfficeName,Office Location;postalCode,ZIP/Postal Code;postOfficeBox,Post Office Box;primaryInternationalISDNNumber,International ISDN Number;primaryTelexNumber,Telex Number;samAccountName,Logon Name (pre-Windows 2000);sn,Last Name;st,State/Province;streetAddress,Street Address;telephoneNumber,Telephone Number;telexNumber,Telex Number (Others);title,Job Title;url,Web Page Address (Others);displayName,Display Name;userWorkstations,Logon Workstations;userPrincipalName,Logon Name;wWWHomePage,Web Page Address; |
| IntellimirrorGroup | |
| IntellimirrorSCP | |
| interSiteTransport | |
| interSiteTransportContainer | |
| licensingSiteSettings | |
| lostAndFound | description,Description;cn,Name; |
| msCOM-Partition | |
| msCOM-PartitionSet | |
| msMQ-Custom-Recipient | msMQ-Recipient-FormatName,Format Name;description,Description;cn,Name; |
| msMQ-Group | member,Member Queues;description,Description;cn,Name; |
| mSMQConfiguration | |
| mSMQEnterpriseSettings | |
| mSMQMigratedUser | |
| mSMQQueue | |
| mSMQSettings | |
| mSMQSiteLink | |
| nTDSConnection | |
| nTDSDSA | |
| nTDSService | |

| | |
|---|---|
| nTDSSiteSettings | |
| nTFRSMember | |
| nTFRSReplicaSet | |
| nTFRSSettings | |
| nTFRSSubscriber | |
| nTFRSSubscriptions | |
| organizationalUnit | managedBy,Managed By;description,Description;ou,Name; |
| pKICertificateTemplate | description,Description;cn,Name; |
| printQueue | versionNumber,Object Version;url,Web Page Address;serverName,Server Name;printStaplingSupported,Supports Stapling;printShareName,Share Name;printPagesPerMinute,Pages per Minute;printRateUnit,Speed Units;printRate,Speed;printOwner,Owner Name;printMemory,Installed Memory;printMediaSupported,Paper Types Supported;printMediaReady,Paper Available;printMaxResolutionSupported,Maximum Resolution;printLanguage,Printer Language;printerName,Name;printDuplexSupported,Supports Double-sided Printing;printColor,Supports Color Printing;printCollate,Supports Collation;printBinNames,Input Trays;portName,Port;location,Location;driverName,Model;description,Comment;contactName,Contact;assetNumber,Asset Number;uNCName,Network Name;cn,Directory Service Name; |
| queryPolicy | |
| remoteStorageServicePoint | cn,Name |
| rpcContainer | description,Description;cn,Name; |
| server | |
| serversContainer | |
| serviceAdministrationPoint | |
| site | |
| siteLink | |
| siteLinkBridge | |
| sitesContainer | |
| subnet | |
| subnetContainer | |
| trustedDomain | description,Description;cn,Name; |
| user | wWWHomePage,Web Page Address;userPrincipalName,Logon Name;userWorkstations,Logon Workstations;displayName,Display Name;url,Web Page Address (Others);title,Job Title;telexNumber,Telex Number (Others);telephoneNumber,Telephone Number;streetAddress,Street Address;st,State/Province;sn,Last Name;samAccountName,Logon Name (pre-Windows 2000);primaryTelexNumber,Telex Number;primaryInternationalISDNNumber,International ISDN Number;postOfficeBox,Post Office Box;postalCode,ZIP/Postal Code;physicalDeliveryOfficeName,Office Location;personalTitle,Title;pager,Pager Number;otherTelephone,Phone Number (Others);otherPager,Pager Number (Others);otherMobile,Mobile Number (Others);otherMailbox,E-Mail Address (Others);otherIpPhone,IP Phone Number (Others);otherHomePhone,Home Phone Number (Others);otherFacsimileTelephoneNumber,Fax Number (Others);info,Notes;mobile,Mobile Number;middleName,Middle Name;memberOf,Member Of;manager,Manager;mail,E-Mail Address;l,City;ipPhone,IP Phone Number;internationalISDNNumber,International ISDN Number (Others);initials,Initials;homePostalAddress,Home Address;homePhone,Home Phone;homeDrive,Home Drive;homeDirectory,Home Folder;givenName,First Name;generationQualifier,Generational Suffix;facsimileTelephoneNumber,Fax Number;employeeID,Employee ID;division,Division;distinguishedName,X500 Distinguished Name;directReports,Direct Reports;description,Description;department,Department;company,Company;comment,Comment;co,Country;c,Country Abbreviation;cn,Name;assistant,Assistant; |
| volume | uNCName,Network Path;keywords,Keywords;managedBy,Managed By;description,Description;cn,Name; |
| Default2 | description,Description;cn,Name; |
| DS-UI-Default-Settings2 | |
| localPolicy2 | |
| nTDSSettings2 | |

1 This table lists only classes/items defined in the display specifier container.

2 These items are not defined in the schema, but are provided with classDisplayName or have a place holder in the display specifier container.

# Appendix I: Default Container Hierarchy for Active Directory Partitions

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix contains the default container hierarchy for all Active Directory partitions.

**Configuration Directory Partition**
The configuration directory partition root object has the following child objects:

**DisplaySpecifiers**
Contains the objects that define different user interfaces for each object class in the schema that requires a graphical user interface (for example, right-click menus and property pages).

**Extended-Rights**
Stores objects of class *controlAccessRight* that can be used by applications to extend standard access control.

**ForestUpdates**
Stores operation objects that are generated by forest preparation tasks (when you run **adprep /forestprep**) so that the system can check for the tasks that have and have not been completed when you are upgrading the first domain controller in the forest to Windows Server 2003. The child object CN=Operations contains the objects that represent each update operation. These objects are named for the GUID of the operation. The child object CN=Windows2003Update is created to indicate that all **adprep** operations have run.

**LostAndFoundConfig**
Provides storage for global configuration objects that have been created in or moved to a location that no longer exists after replication.

**NTDS Quotas**
Stores objects (class *msDS-QuotaControl*) that contain object ownership quota assignments for the configuration directory partition. Quotas limit the number of objects that a user (including *inetOrgPerson*), group, computer, or service can own in a domain, configuration, or application directory partition.

| **Note** |
| --- |
| This container only exists in Windows Server 2003 Active Directory. |

**Partitions**
Stores the cross-references to every directory partition in the forest, including the configuration partition, the schema partitions, and all domain directory partitions.

**Physical Locations**
Serves no purpose in Windows 2000 Server or Windows Server 2003. It is reserved for future use.

**Services**
Stores network-wide, service-specific information that applications use to connect to instances of services in the forest.

| **Note** |
| --- |
| The Services node in Active Directory Sites and Services is hidden by default. To reveal the Services node, right-click **Active Directory Sites and Services**, point to **View**, and then click **Show Services Node**. |

**Sites**
Stores all of the site objects in the enterprise network, objects that represent replicating domain controllers in those sites, and objects that define the replication topology.

**Well-Known Security Principals**
Contains the special identities that are defined by the security system, such as Everyone, LocalSystem, Principal Self, Authenticated User, and Creator Owner.

**Schema Directory Partition**
Similarly, there is only one schema directory partition per forest. The schema directory partition contains the definitions of all

objects that can be instantiated in Active Directory. It also stores the definitions of all attributes that can be a part of objects in Active Directory. Every domain controller has one fully writeable copy of the schema directory partition, although schema updates are allowed only on the domain controller that is the schema operations master.

The schema directory partition root object contains one child object for each class of objects that can be instantiated in the Active Directory forest and contains one object for each attribute that can be part of an object in the Active Directory forest.

**Domain Directory Partitions**
The domain directory partition root object has the following child objects:

**Builtin**
Stores built-in groups. All built-in groups have a well-known security identifier (SID).

**Computers**
Default storage area for new computer objects that were originally created through legacy APIs that are not Active Directory–aware.

**Domain Controllers**
Default container for new domain controllers. The Domain Controllers container cannot be renamed.

**ForeignSecurityPrincipals**
Proxy objects for security principals that are from Microsoft® Windows NT® version 4.0 operating system domains or Windows NT® 3.51 domains, or that are from different forests, and that have been added to Windows® 2000 or Windows Server 2003 groups.

**LostAndFound (Advanced Features)**
Storage area for new domain-wide objects whose containers were deleted elsewhere at the same time that the object was created. The LostAndFoundConfig container in the configuration directory partition serves the same purpose for forest-wide objects.

**NTDS Quotas (Advanced Features)**
Storage area for objects of class *msDS-QuotaControl*, which contain object ownership quotas for the domain directory partition. Quotas limit the number of objects that a user, group, computer, or service can create in a directory partition.

| **Note** |
| --- |
| This container only exists in Windows Server 2003 Active Directory. |

**Program Data (Advanced Features)**
Empty container that is available for applications to store application-specific data in the domain directory partition.

**System (Advanced Features)**
Built-in system settings for the various system service containers and objects.

**Users**
Default storage area for new user accounts that are created through legacy APIs that are not Active Directory–aware.

| **Note** |
| --- |
| The Users container, the Computers container, and several other special containers, called “well-known” containers, can be dependably located by applications. |

**Deleted Objects**
A special container, not visible in the user interface (UI), to which objects are moved when they are deleted.

**Infrastructure**
An object of class *infrastructureUpdate* that identifies the NTDS settings object of the domain controller that holds the infrastructure master role for the domain.

| **Note** |
| --- |
| (Advanced Features) above refers to the fact that, when using the Active Directory Users and Computers MMC snap-in, this container is only visible if the snap-in if the Advanced Features option is checked. |

# Appendix J: Default Owners of Active Directory Objects

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix contains information about differences between the administrative groups that become default object owners in Windows 2000 and Windows Server 2003.

**Default Owners in Windows 2000 Server**
In Windows 2000 Server Active Directory, the following rules of ownership apply for new object creation in all directory partition types, including configuration, schema, and domain:

- If the requestor is a member of the Builtin Administrators group, the default object owner is the Administrators group.
- If the requester is a member of the Domain Admins group of the domain of the contacted domain controller, the default object owner is the Domain Admins.
- If the requestor is not a member of the Builtin Administrators group or the Domain Admins group of the domain of the contacted domain controller, the default object owner is the requestor.

**Default Owners in Windows Server 2003**
In Windows Server 2003 Active Directory, the following major differences affect how ownership of new objects is computed:

- Rules of ownership are specific to the target directory partition type.
- The Enterprise Admins group is considered for forest-wide objects, not just the Domain Admins group. If the requestor is a member of the Enterprise Admins group, the default object owner is the group, not the individual requestor.
- The Builtin Administrators group no longer has default group ownership of directory objects.

The following tables show how ownership rules are applied, from highest to lowest precedence, relative to the different directory partition types and requestor membership in administrative groups.

**Directory Partition Type: Configuration**

| Requestor Group | Default Owner |
|---|---|
| Enterprise Admins | Enterprise Admins |
| Domain Admins | Domain Admins group of the domain of the contacted domain controller |
| Builtin Administrators | Token-user |
| Other | Default-owner-in-token |

**Directory Partition Type: Schema**

| Requestor Group | Default Owner |
|---|---|
| Schema Admins | Schema Admins |
| Enterprise Admins | Enterprise Admins |
| Domain Admins | Domain Admins group of the domain of the contacted domain controller |
| Builtin Administrators | Token-user |
| Other | Default-owner-in-token |

**Directory Partition Type: Domain**

| Requestor Group | Default Owner |
|---|---|
| Domain Admins | Domain Admins group of the domain of the contacted domain controller |

| | |
|---|---|
| Enterprise Admins | Enterprise Admins |
| Builtin Administrators | Token-user |
| Other | Default-owner-in-token |

**Directory Partition Type: Application**

| **Requestor Group** | **Default Owner** |
|---|---|
| Domain Admins | Domain Admins of the security descriptor reference domain of the application directory partition |
| Enterprise Admins | Enterprise Admins |
| Builtin Administrators | Token-user |
| Other | Default-owner-in-token |

# Appendix K: Default Settings in the Master Security Descriptor of the AdminSDHolder Object

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix provides the default settings in the master security descriptor of the **AdminSDHolder** object for Windows 2000 and Windows Server 2003.

**Default Setting in the Master Security Descriptor of the AdminSDHolder Object for Windows 2000 Server**

| Type | Name | Permission | Apply To |
|---|---|---|---|
| Allow | Administrators | • List Contents<br>• Read All Properties<br>• Write All Properties<br>• Delete<br>• Read Permissions<br>• Modify Permissions<br>• Modify Owner<br>• All Validated Writes<br>• All Extended Rights<br>• Create All Child Objects<br>• Delete All Child Objects | This object only |
| Allow | Authenticated Users | • List Contents<br>• Read All Properties<br>• Read Permissions | This object only |

| | | | |
|---|---|---|---|
| Allow | Domain Admins | • List Contents<br>• Read All Properties<br>• Write All Properties<br>• Read Permissions<br>• Modify Permissions<br>• Modify Owner<br>• All Validated Writes<br>• All Extended Rights<br>• Create All Child Objects<br>• Delete All Child Objects | This object only |
| Allow | Enterprise Admins | • List Contents<br>• Read All Properties<br>• Write All Properties<br>• Read Permissions<br>• Modify Permissions<br>• Modify Owner<br>• All Validated Writes<br>• All Extended Rights<br>• Create All Child Objects<br>• Delete All Child Objects | This object only |
| Allow | Everyone | • Change Password | This object only |
| Allow | Pre–Windows 2000 Compatible Access | • List Contents<br>• Read All Properties<br>• Read Permissions | User objects |
| Allow | SYSTEM | • Full Control | This object only |

**Default Setting in the Master Security Descriptor of the AdminSDHolder Object for Windows Server 2003**

| Type | Name | Permission | Apply To |
|---|---|---|---|

| | | | |
|---|---|---|---|
| Allow | Administrators | • List Contents<br>• Read All Properties<br>• Write All Properties<br>• Delete<br>• Read Permissions<br>• Modify Permissions<br>• Modify Owner<br>• All Validated Writes<br>• All Extended Rights<br>• Create All Child Objects<br>• Delete All Child Objects | This object only |
| Allow | Authenticated Users | • List Contents<br>• Read All Properties<br>• Read Permissions | This object only |
| Allow | Domain Admins | • List Contents<br>• Read All Properties<br>• Write All Properties<br>• Read Permissions<br>• Modify Permissions<br>• Modify Owner<br>• All Validated Writes<br>• All Extended Rights<br>• Create All Child Objects<br>• Delete All Child Objects | This object only |

| | | | |
|---|---|---|---|
| Allow | Enterprise Admins | • List Contents<br>• Read All Properties<br>• Write All Properties<br>• Read Permissions<br>• Modify Permissions<br>• Modify Owner<br>• All Validated Writes<br>• All Extended Rights<br>• Create All Child Objects<br>• Delete All Child Objects | This object only |
| Allow | Everyone | • Change Password | This object only |
| Allow | Pre–Windows 2000 Compatible Access | • List Contents<br>• Read All Properties<br>• Read Permissions | User and InetOrgPerson objects |
| Allow | SYSTEM | • Full Control | This object only |
| Allow | SELF | • Change Password | This object only |
| Allow | Cert Publisher | • Read user Cert<br>• Write user Cert | This object only |
| Allow | Windows Authorization Access Group | • Read tokenGroups GlobalAndUniversal | This object only |
| Allow | Terminal Server License Servers | • Read terminalServer*<br>• Write terminalServer* | This object only |

The terminalServer property, as well as many other properties, is defined in the schema but filtered from display in the ACL editor UI. The list of filtered properties is stored in the file Dssec.dat that is located in the *systemroot*\System32 folder on all domain controllers. If you need to apply permissions to a property that is not shown in the UI, you can edit the entry in Dssec.dat to display the filtered properties through the UI. For more information about editing this file, see article 296490, "How to Modify the Filtered Properties of an Object (296490)" in the Microsoft Knowledge Base at http://go.microsoft.com/fwlink/?LinkId=4441.

The following table shows the difference in permissions between the default setting in the master security descriptor of the **AdminSDHolder** object for Windows 2000 and Windows Server 2003 by list permissions that are added in Windows

Server 2003.

**New Permissions in Windows Server 2003**

| Type | Name | Permission | Apply To |
|---|---|---|---|
| Allow | Pre–Windows 2000 Compatible Access | • List Contents<br>• Read All Properties<br>• Read Permissions | InetOrgPerson objects |
| Allow | SELF | • Change Password | This object only |
| Allow | Cert Publishers | • Read user Cert<br>• Write user Cert | This object only |
| Allow | Windows Authorization Access Group | • Read tokenGroups GlobalAndUniversal | This object only |
| Allow | Terminal Server License Servers | • Read terminalServer*<br>• Write terminalServer* | This object only |

The following is the DACL (in SDDL) on the ADMINSDHOLDER object in Windows 2000:

```
(A;;RPLCLORC;;;AU)
(A;;RPWPCRLCLOCCDCRCWDWOSWSD;;;BA)
(A;;RPWPCRLCLOCCDCRCWDWOSW;;;EA)
(A;;RPWPCRLCLOCCDCRCWDWOSW;;;DA)
(A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
(OA;;RP;037088f8-0ae1-11d2-b422-00a0c968f939;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RP;59ba2f42-79a2-11d0-9020-00c04fc2d3cf;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RP;bc0ac240-79a9-11d0-9020-00c04fc2d4cf;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RP;4c164200-20c0-11d0-a768-00aa006e0529;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RP;5f202010-79a5-11d0-9020-00c04fc2d4cf;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RPLCLORC;;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;CR;ab721a53-1e2f-11d0-9819-00aa0040529b;;WD)
```

The following is the DACL (in SDDL) on the ADMINSDHOLDER object in Windows Server 2003:

```
(A;;RPLCLORC;;;AU)
(A;;RPWPCRLCLOCCDCRCWDWOSWSD;;;BA)
(A;;RPWPCRLCLOCCDCRCWDWOSW;;;EA)
(A;;RPWPCRLCLOCCDCRCWDWOSW;;;DA)
(A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
(OA;;RP;037088f8-0ae1-11d2-b422-00a0c968f939;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RP;59ba2f42-79a2-11d0-9020-00c04fc2d3cf;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RP;bc0ac240-79a9-11d0-9020-00c04fc2d4cf;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RP;4c164200-20c0-11d0-a768-00aa006e0529;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RP;5f202010-79a5-11d0-9020-00c04fc2d4cf;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;RPLCLORC;;bf967aba-0de6-11d0-a285-00aa003049e2;RU)
(OA;;CR;ab721a53-1e2f-11d0-9819-00aa0040529b;;WD)
(OA;;CR;ab721a53-1e2f-11d0-9819-00aa0040529b;;PS)
(OA;;RPWP;bf967a7f-0de6-11d0-a285-00aa003049e2;;CA)
(OA;;RP;037088f8-0ae1-11d2-b422-00a0c968f939;4828CC14-1437-45bc-9B07-AD6F015E5F28;RU)
(OA;;RP;59ba2f42-79a2-11d0-9020-00c04fc2d3cf;4828CC14-1437-45bc-9B07-AD6F015E5F28;RU)
(OA;;RP;bc0ac240-79a9-11d0-9020-00c04fc2d4cf;4828CC14-1437-45bc-9B07-AD6F015E5F28;RU)
(OA;;RP;4c164200-20c0-11d0-a768-00aa006e0529;4828CC14-1437-45bc-9B07-AD6F015E5F28;RU)
(OA;;RP;5f202010-79a5-11d0-9020-00c04fc2d4cf;4828CC14-1437-45bc-9B07-AD6F015E5F28;RU)
(OA;;RPLCLORC;;4828CC14-1437-45bc-9B07-AD6F015E5F28;RU)
(OA;;RP;46a9b11d-60ae-405a-b7e8-ff8a58d456d2;;S-1-5-32-560)
```

```
(OA;;WPRP;6db69a1c-9422-11d1-aebd-0000f80367c1;;S-1-5-32-561)
```

| **Note** |
| --- |
| In each case, though there are multiple ACEs granting Read-Property permissions to specific property-sets to the permissions to the Pre-Windows 2000 Compatible Access group, the following ACE grants blanket Read-Property permissions, in effect making the other ACEs unnecessary. However, for compatibility reasons, they should not be removed. The tables in this appendix document the access granted by the one ACE granting blanket read-property access. |

# Appendix L: Implementing Service Management Delegation Roles

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix describes how to implement your service management delegation roles. This appendix contains one section for each Microsoft recommended role. Each section contains the following information:

- The security group to create or use to represent the instance of the role
- The set of permissions that this security group should be granted

# Implementing the Forest Configuration Operators Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to implement the forest configuration operators role.

**To implement the one required instance of the Forest Configuration Operators role**

1. Create a Universal Group called <*Domain-Name*> *Domain Config Ops* in this domain's Service Management OU (ou=Service Management, dc=<Domain>).

| **Note** |
| --- |
| If Universal groups are not available, create a Global security group. |

2. Grant this group permissions required to perform assigned Installation Management tasks.

   a. Grant this group the **DS-Replication-Get-Changes** extended right on the following objects:

      - CN=Configuration, DC=<Forest-Root-Domain>
      - CN=Schema, CN=Configuration, DC=<Forest-Root-Domain>

   b. Grant this group the **DS-Replication-Manage Topology** extended right on the following objects:

      - CN=Configuration, DC=<Forest-Root-Domain>
      - CN=Schema, CN=Configuration, DC=<Forest-Root-Domain>

   c. In a Windows 2000 Active Directory environment, additionally grant this group the **DS-Replication-Get-Changes-All** extended right on the following objects:

      - CN=Configuration, DC=<Forest-Root-Domain>
      - CN=Schema, CN=Configuration, DC=<Forest-Root-Domain>

   d. In a Windows 2000 Active Directory environment, additionally grant this group the **DS-Replication-Monitor-Topology** extended right on the following objects:

      - CN=Configuration, DC=<Forest-Root-Domain>
      - CN=Schema, CN=Configuration, DC=<Forest-Root-Domain>

   e. Grant this group the following permissions:

      - Read All Properties on CN=Sites, CN=Configuration, DC=<Forest-Root-Domain> (**Inheritable** – apply onto this object and all child objects)
      - Create All Child Objects on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<Forest-Root-Domain> (**Inheritable** – apply onto this object and all child objects)

- Create Computer objects on OU=Domain Controllers,DC=<domain>

- Full Control to "Creator Owner" on CN=Sites, CN=Configuration, DC=<Forest-Root-Domain> (**Inheritable** – apply onto this object and all child objects)

f. Grant this group the "Enable computer and user accounts to be trusted for delegation" user right by modifying the default domain controller security policy for this domain.

| **Note** |
| --- |
| This is a very powerful user right and in general should be granted with care. |

g. Finally, when a member of this group needs to add a replica DC, he/she must be granted Full Control on the computer object representing the server that is being promoted and must be made a member of the Local Administrators group on that computer.

3. Grant this group permissions required to perform assigned Operations Master Role Management tasks.

   a. Grant this group the **Change-Schema-Master** extended right on cn=Schema, CN=Configuration, DC=<Forest-Root-Domain>

   b. Grant this group the **Change-Domain-Master** extended right on cn=Partitions, CN=Configuration, DC=<Forest-Root-Domain>

   c. Grant this group Write-Property permissions to write the fSMORoleOwner property on cn=Schema, CN=Configuration, DC=<Forest-Root-Domain>

   d. Grant this group Write-Property permissions to write the fSMORoleOwner property on cn=Partitions, CN=Configuration, DC=<Forest-Root-Domain>

4. Grant this group permissions required to protect and manage trusts for the entire forest.

   - In each domain, grant this group the following permissions:

     - Create Trusted-Domain objects on CN=System, DC=<domain>

     - Delete Trusted-Domain objects on CN=System, DC=<domain>

     - Write-Property to all attributes on Trusted-Domain objects CN=System, DC=<domain> (Inheritable ACE, applies to Trust objects)

     - Additionally, if members of this group will use Active Directory trust management tools, make this group a member of the BuiltIn Admins group in the domain

| **Note** |
| --- |
| All trust management tools in Active Directory require that an administrator performing trust management using these tools be a member of the BuiltIn Admins group in the domain. |

5. Grant this group permissions required to perform LDAP policy management:

   a. Grant this group Create Child permissions to create *Query-Policy* objects in the cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<Forest-Root-Domain> container

   b. Grant this group Delete Child permissions to create *Query-Policy* objects in the cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<Forest-Root-Domain> container.

c. Grant this group Write All Properties permissions on *Query-Policy* objects in the cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration, DC=<Forest-Root-Domain> container. **(Inheritable permissions)**

d. To perform advanced LDAP policy administrative tasks such as affecting the LDAP query policies associated with a specific domain controller or with all domain controllers in a specific site, additional permissions might be required – to grant these permissions refer to Appendix A: Active Directory Administrative Tasks.

| **Note** |
| --- |
| Note that as an alternative to the above, the pre-defined *Enterprise Admins* security group (CN=Enterprise Admins, CN=Users, DC=<Forest-Root-Domain>) can be used to represent this role. The *Domain Admins* security group has sufficient permissions required to carry out all the responsibilities assigned to this role. As an alternative, the *Domain Admins* security group could be used to represent the role instance specific to the domain.<br><br>Note that as another alternative to the above, an organization could choose to delegate a customized version of this role. All the permissions required to delegate an Active Directory administrative task are specific in Appendix A: Active Directory Administrative Tasks. |

# Implementing the Domain Configuration Operators Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to implement the domain configuration operators role.

**To implement the one required instance of the Domain Configuration Operators role**

1. Create an OU called Service Management in every domain (CN=Service Management, DC=<domain>)
2. Create a Domain Local Group called <*Domain-Name*> *Domain Config Ops* in this domain's Service Management OU (ou=Service Management, dc=<Domain>).
3. Grant this group permissions required to perform assigned *Installation management* tasks
   a. Grant this group the **Install Replica** extended right on the following objects:
      - DC=<domain>
   b. Grant this group the **DS-Replication-Get-Changes** extended right on the following objects:
      - DC=<domain>
      - CN=Configuration, DC=<Forest-Root-Domain>
      - CN=Schema, CN=Configuration, DC=<Forest-Root-Domain>
   c. Grant this group the **DS-Replication-Manage Topology** extended right on the following objects:
      - DC=<domain>
      - CN=Configuration, DC=<Forest-Root-Domain>
      - CN=Schema, CN=Configuration, DC=<Forest-Root-Domain>
   d. In a Windows 2000 Active Directory environment, additionally grant this group the **DS-Replication-Get-Changes-All** extended right on the following objects:
      - DC=<domain>
      - CN=Configuration, DC=<Forest-Root-Domain>
      - CN=Schema, CN=Configuration, DC=<Forest-Root-Domain>
   e. In an Windows 2000 Active Directory environment, additionally grant this group the **DS-Replication-Monitor-Topology** extended right on the following objects:
      - DC=<domain>
      - CN=Configuration, DC=<Forest-Root-Domain>

- CN=Schema, CN=Configuration, DC=<Forest-Root-Domain>

f. Grant this group the following permissions:

- Read All Properties on CN=Sites, CN=Configuration, DC=<Forest-Root-Domain> (**Inheritable** – apply onto this object and all child objects)
- Create All Child Objects on CN=Servers, CN=<Site>, CN=Sites, CN=Configuration, DC=<Forest-Root-Domain> (**Inheritable** – apply onto this object and all child objects)
- Create Computer objects on OU=Domain Controllers,DC=<domain>
- Full Control to "Creator Owner" on CN=Sites, CN=Configuration, DC=<Forest-Root-Domain> (**Inheritable** – apply onto this object and all child objects)

g. Grant this group the "Enable computer and user accounts to be trusted for delegation" user right by modifying the default domain controller security policy for this domain.

| **Note** |
| --- |
| This is a very powerful user right and in general should be granted with care. |

h. Finally, when a member of this group needs to add a replica DC, that member must be granted full control on the computer object representing the server that is being Promoted and must be made a member of the Local Administrators group on that computer.

4. Grant this group permissions required to perform assigned Operations Master Role management tasks.

   a. Grant this group the **Change-RID-Master** extended right on cn=RID Manager$, cn=System, dc=<Domain>

   b. Grant this group the **Change-PDC** extended right on dc=<Domain>

   c. Grant this group the **Change-Infrastructure-Master** on dc=<Domain>

   d. Grant this group Write-Property permissions to write the fSMORoleOwner property on cn=RID Manager$, cn=System, dc=<Domain>

   e. Grant this group Write-Property permissions to write the fSMORoleOwner property on dc=<Domain>

   f. Grant this group Write-Property permissions to write the fSMORoleOwner property on cn=Infrastructure, dc=<Domain>

5. Grant this group permissions required to protect and manage the Domain Controllers OU

   - Grant the group **Full-Control** over the Domain Controllers OU, CN=Domain Controllers, DC=<domain>

6. Grant this group permissions required to protect and manage content stored in the System container

   - Grant the group **Full-Control** over the System container, CN=System, DC=<domain> (**Inheritable** – apply onto this object and all child objects)

7. Grant this group permissions required to raise the Domain Functionality level:

   - Grant this group Write-Property permissions to write the ms-DS-behavior-version property on dc=<Domain>

8. Grant this group permissions required to restore Active Directory from backup if and when necessitated

   a. Grant this group the "Restore Files and Directories" user right by modifying the default domain controller security policy for this domain.

   b. To perform and authoritative restore, a member of this group will require the credentials of the special Administrator account available only in DS Restore mode.

| **Note** |
| --- |
| As one alternative to the above, the pre-defined Domain Admins security group (CN=Domain Admins, CN=Users, DC=<Forest-Root-Domain>) can be used to represent this role. The *Enterprise Admins* security group has sufficient permissions required to carry out all the responsibilities assigned to this role.<br><br>As another alternative to the above, an organization could choose to delegate a customized version of this role. All the permissions required to delegate an Active Directory administrative task are specific in Appendix A: Active Directory Administrative Tasks. |

# Implementing the Security Policy Admins Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to implement the security policy admins role.

**To implement the one required instance of the Security Policy Admins role**

1. Create a Universal Group called <*Forest-Name*> *Security Policy Admins* in the Service Management OU (ou=Service Management, dc=<Forest Root Domain>).

   | **Note** |
   | --- |
   | If Universal groups are not available, create a Global Security group. |

2. Use the Group Policy Management Console to modify the following permissions on the default Domain Policy object:

   a. Grant the <Forest-Name> Security Policy Admins Edit permissions

   b. Revoke Edit permissions from all other security principals

   | **Note** |
   | --- |
   | To download GPMC, see the Microsoft Download Center at http://www.microsoft.com/downloads/details.aspx?familyid=0a6d4c24-8cbd-4b35-9272-dd3cbfc81887&displaylang=en. GPMC simplifies the management of Group Policy by making it easier to understand, deploy, manage, and troubleshoot Group Policy implementations. GPMC runs on Windows XP Professional SP1 and Windows Server 2003 computers and can manage Group Policy in either Windows 2000 or Windows Server 2003 domains. |

3. Use the Group Policy Management Console to modify the following permissions on the default Domain Controllers Policy object:

   a. Grant the <Forest-Name> Security Policy Admins Edit permissions

   b. Revoke Edit permissions from all other security principals

   | **Note** |
   | --- |
   | Note that revoking *Edit* permissions for all other security principals only the <Forest> Security Policy Admins can edit these Group Policy settings based on current permissions. Any user or group who owns these security policies or has *modify security* permissions can grant themselves the ability to edit these settings. Thus, , revoking *Edit* permissions for all other security principals does not grant the <Forest> Security Policy Admins exclusive ability to modify these settings from a strict security perspective. Thus the <Forest> Security Policy Admins are not the only ones who can modify these settings. An organization should establish business policies that specify that any changes to the default Domain and default Domain Controller policies across all domains should be made only by the <Forest> Security Policy Admins. The purpose of introducing and enabling this role is to make service management more tractable and increase accountability. |

# Implementing the Service Admin Managers Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to implement the service admin managers role.

**To implement the one required instance of the Service Admin Managers role**

1. Create a Universal Group called <*Forest-Name*> *Service Admin Managers* in the forest root domain in the Service Management OU (ou=Service Management, dc=<Forest Root Domain>)

   | **Note** |
   | --- |
   | If Universal groups are not available, create a Global security group. |

2. In each domain, modify the DACL on the ADMINSDHOLDER object in each domain (CN=AdminSDHolder, CN=System, DC=<domain> as follows:

   a. Grant the <Forest-Name> Service Admin Managers group Full-Control permissions on the object.

   b. Grant the Enterprise Admins group Full-Control permissions on the object.

   c. Optionally, to ensure that members of no other security groups can modify the membership of or delete service administrator groups, modify the DACL of the object by remove all permissions granted to Domain Admins, Built-In Admins and Enterprise Admins.

   | **Note** |
   | --- |
   | Note that even if you remove all permissions granted to Built-In Admins, any administrator who has the *Take Ownership of files or other objects* user right granted in the Domain Controller Security Policy can still take ownership of a service administrator security group object or member user object. By default the Built-In Admins are granted this user right and by default Domain Admins and Enterprise Admins are members of the Built-In Admins group. Thus, removing permissions granted to Domain Admins, Built-In Admins and Enterprise Admins in the ADMINSDHOLDER object DACL will not allow a member of these groups to directly modify the membership of these groups or modify a member user, it will not prevent them from taking ownership of the 0bject representing a service administrator security group or member user and then modifying the object. If must ensure that the Service Admin Managers group is the only group that can truly manage all service administrator groups and members, consider taking away this privilege from the above set of administrative groups. |

# Implementing the Domain Controller Administrators Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For each instance of this role, in each domain use the following procedure to implement an instance of the Domain Controller Administrators role.

**To implement an instance of the Domain Controller Administrators role**

1. Create a Domain Local Group called <*Domain-Name*> <*location*> *Domain Controller Admins* in the Service Management OU in the domain (ou=Service Management, dc=<Domain>).
2. Make this group a member of the BuiltIn\Administrators group (CN=Administrators, CN=Builtin, DC=<domain>) in the domain.

# Implementing the Backup Operators Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The pre-defined *Backup Operators* security group (CN=Backup Operators, CN=Builtin, DC=<domain>) has all the permissions required to carry out all the responsibilities assigned to this role.

In each domain, use this domain specific security group to represent the role instance specific to the domain.

# Implementing the Schema Administrators Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The pre-defined *Schema Admins* security group (CN=Enterprise Admins, CN=Users, DC=<Forest-Root-Domain>) has all the permissions required to carry out all the responsibilities assigned to this role. Use this security group to represent this administrative role.

As a best-practice, keep the membership of this group empty and populate the membership of this group as and when a schema management administrative task needs to be performed.

# Implementing the Replication Management Admins Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to implement the replication management admins role.

**To implement the one recommended instance of the Replication Management Admins role**

1. Create a Domain Local Group called <Forest-Name> Replication Management Admins in the Service Management OU (ou=Service Management, dc=<Forest Root Domain>).
2. Grant this group the following extended rights:
   a. Grant the DS-Replication-Manage-Topology (Manage Replication Topology) extended right on CN=Configuration, DC=<Forest Root Domain>
   b. Grant the DS-Replication-Manage-Topology (Manage Replication Topology) extended right on CN=Schema, CN=Configuration, DC=<Forest Root Domain>
   c. Grant the DS-Replication-Manage-Topology (Manage Replication Topology) extended right on all domain partition heads including forest root domain
3. Grant this group the following permissions on the Sites container (CN=Sites, CN=Configuration, DC=<Forest-Root-Domain>:
   - Create Site objects (non-inheritable)
   - Delete Site objects (non-inheritable)
   - Create Connection objects (inheritable)
   - Delete Connection objects (inheritable)
   - Write All Properties on this object and all child objects (inheritable)
4. Grant this group the following permissions on the Subnets container (CN=Subnets, CN=Sites, CN=Configuration, DC=<Forest-Root-Domain>:
   - Create Subnet objects (inheritable)
   - Delete Subnet objects (inheritable)
   - Write All Properties on Subnet objects (inheritable)
5. Grant this group the following permissions on the Inter-Site Transports container (CN= Inter-Site Transports , CN=Sites, CN=Configuration, DC=<Forest-Root-Domain>:
   - Create Site Link objects (non-inheritable)
   - Delete Site Link objects (non-inheritable)
   - Write All Properties on Site Link objects (inheritable)

**Note**

In Active Directory, the creator of an object also becomes the owner of the object. The owner of the object has the implicit permission to give away ownership and to control access to the object. Thus, a delegated administrator who can create objects can grant other users permissions on the object created and can additionally use inheritance to grant him/her self or another user permissions on all child objects in the sub-tree rooted at this object. Administrators delegating responsibility should be aware of this issue. An organization could consider the creation and implementation of script that runs in Enterprise Admin credentials in the Configuration partition and takes ownership of objects that were created by delegated administrators. However the script should be careful enough to ensure that it only takes ownership of those objects that were created by administrators and not take ownership of objects created y the system or by Domain Controller computer accounts.

# Implementing the Replication Monitoring Operators Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to implement the replication monitoring operators role.

**Procedures**
**To implement the one recommended instance of the Replication Monitoring Operators role**

1. Create a Domain Local Group called <Forest-Name> Replication Monitoring Ops in the Service Management OU (ou=Service Management, dc=<Forest Root Domain>).

2. In a Windows 2000 Active Directory environment grant this group the following permissions:

   a. Grant the DS-Replication-Manage-Topology (Manage Replication Topology) extended right on CN=Configuration, DC=<Forest Root Domain>

   b. Grant the DS-Replication-Manage-Topology (Manage Replication Topology) extended right on CN=Schema, CN=Configuration, DC=<Forest Root Domain>

   c. Grant the DS-Replication-Manage-Topology (Manage Replication Topology) extended rights on all domain partition heads including forest root domain

   | **Note** |
   | --- |
   | In Windows 2000, the *Monitor Replication Topology* right does not exist. Thus the *Manage Replication Topology* right needs to be granted to delegate the ability to monitor replication. Note that an individual who is granted the *Manage Replication Topology* extended right is sufficiently privileged to perform many security-sensitive operations including forcing topology regeneration etc. Thus, it is recommended that an organization put in place policies that govern the specific operations that delegated administrators in the *Replication Monitoring Operators* role are legally authorized to carry out. Alternatively, you could just implement the *Replication Management Administrators* role and assign responsibility for managing and monitoring replication to the same set of delegated administrators. |

3. In a Windows Server 2003 Active Directory environment grant this group the following permissions:

   a. Grant the DS-Replication-Monitor-Topology (Monitor Replication Topology) extended right on CN=Configuration, DC=<Forest Root Domain>

   b. Grant the DS-Replication-Monitor-Topology (Monitor Replication Topology) extended right on CN=Schema, CN=Configuration, DC=<Forest Root Domain>

   c. Grant the DS-Replication-Monitor-Topology (Monitor Replication Topology) extended right on all domain partition heads including forest root domain

# Implementing the DNS Admins Role

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to implement the DNS admins role.

**To implement the one recommended instance of the DNS Admins role**

1. Create a Universal Group called <Forest-Name> DNS Admins in the Service Management OU (ou=Service Management, dc=<Forest Root Domain>).

| **Note** |
| --- |
| If Universal groups are not available, create a Global security group. |

2. Grant the <Forest-Name> DNS Admins the following permissions:
   - Full control on CN=MicrosoftDNS, DC=ForestDnsZones, DC=<forest root domain>
3. Create one Global Group called <Domain-Name> DNS Admins in the Service Management OU for each domain (ou=Service Management, dc=<Forest Root Domain>).
4. In each domain and on ever NDNC used by DNS <domain>, grant the respective <Domain-Name> DNS Admins group the following permissions:
   - Full control on CN=MicrosoftDNS, CN=System, DC=<domain>
   - Full control on CN=MicrosoftDNS, DC=DomainDnsZones, DC=<domain>
5. Make the <Forest-Name> DNS Admins a member of the <Domain-Name> DNS Admins group from each domain.

# Appendix M: Service Management Delegation Role Definitions

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix presents the set of administrative tasks assigned to every role. The set of recommended roles for delegating service management is as follows:

- Forest Configuration Operators
- Domain Configuration Operators
- Security Policy Administrators
- Service Admin Managers
- Domain Controller Administrators
- Backup Operators
- Schema Administrators
- Replication Management Administrators
- Replication Monitoring Operators
- DNS Administrators

**Forest Configuration Operators Role**

The following is the set of administrative tasks assigned to this role:

- Create a child domain in an existing domain tree
- Demote the last domain controller in a child domain
- Demote the last domain controller in a tree-root domain
- Raise forest functional level
- Create all types of trusts for all domains
- Delete all types of trusts for all domains
- Change the direction of a trust
- Enable/disable name suffix routing (for a given suffix) in a forest
- Reset the trust passwords shared by a trust-pair
- Force the removal of a trust
- Enable/disable SID History on an outbound forest trust
- Enable/disable SID filtering

- Enable selective authentication on an outbound forest/external trust
- Enable/disable placing of name suffix (top level names) information on a realm trust
- Add/remove top-level names from a realm trust
- Add/remove top-level name exclusions from a realm trust
- Modify the transitivity of a realm-trust
- Transfer the schema master role
- Transfer the domain naming master role
- Seize the schema master role
- Seize the domain naming master role
- Manage all LDAP query policy related administrative tasks

**Domain Configuration Operators Role**
The following is the set of administrative tasks assigned to this role:

- Create a replica (additional domain controller)
- Remove a replica
- Designate a domain controller as a global catalog
- Undesignate a domain controller as a global catalog
- Rename a domain controller
- Raise domain functional level
- Create a replica (additional domain controller)
- Remove a replica
- Transfer the RID master role
- Transfer the PDC emulator master role
- Transfer the infrastructure master role
- Seize the RID master role
- Seize the PDC emulator master role
- Seize the infrastructure master role
- Protect and manage the default domain controllers OU
- Protect and manage the content stored in the System container

- Restore Active Directory from backup

**Security Policy Admins Role**
The following is the set of administrative tasks assigned to this role:

- Manage all aspects of the Domain Controller Security Policy for all domains in the forest
- Manage the following aspects of Domain Security Policy for all domains in the forest:
    - Password Policy
    - Account Lockout
    - Kerberos Policy

**Service Admin Managers Role**
The following is the set of administrative tasks assigned to this role:

- Manage and protect all service administrator security groups in the forest
- Manage and protect all service administrator accounts in the forest

**Domain Controller Administrators Role**
The following is the set of administrative tasks assigned to this role:

- Install and modify software
- Install service packs and hot-fixes
- Configure directory service settings in the registry
- Maintain and backing up event logs
- Configure the Service Control Manager
- Manage directory service files and Sysvol
- Start and shut down domain controllers
- Maintain Active Directory database and log files

**Backup Operators Role**
The following is the set of administrative tasks assigned to this role:

- Perform scheduled backups of Active Directory system state

**Schema Administrators Role**
The following is the set of administrative tasks assigned to this role:

- Enable schema modification on a domain controller in the enterprise
- Change the current schema master
- Add a class definition in the schema

- Add an attribute definition in the schema
- Modify a class definition in the schema
- Modify an attribute definition in the schema
- Update the schema cache on demand
- Deactivate a schema class object or resurrect a deactivated schema class object
- Deactivate an attribute class object or resurrect a deactivated schema attribute object
- Make an attribute indexed
- Add attributes to the ANR set
- Designate an attribute as a member of the partial attribute-set that is replicated to the global catalog
- Remove an attribute from the partial attribute-set that is replicated to the global catalog

**Replication Management Administrators Role**
The following is the set of administrative tasks assigned to this role:

- Create a site and add a site
- Rename a site
- Specify the location of a site
- Delete a site
- Create a subnet and add a subnet
- Specify the location of a subnet
- Associate a subnet with a site
- Delete a subnet
- Create a site link
- Add or remove sites to and from a site link
- Modify the cost associated with a site link
- Modify the replication period associated with a site link
- Modify the replication schedule for a site link
- Delete a site link
- Create a site link bridge (object)
- Add or remove sites to and from a site link bridge

- Create a single bridge for the entire network
- Turn off the “Bridge all site links” option for IP/SMTP transport
- Delete a site link bridge (object)
- Create a connection (only if needed)
- Delete a connection (only if needed)
- Take ownership of a KCC-generated connection object
- Manually set a schedule for connection objects
- Enable and disable data compression for inter-site replication
- Change the default setting for the intra-site replication schedule within a site
- Designate or remove a preferred bridgehead server
- Replace a failed preferred bridgehead server
- Force replication between two servers
- Force a synchronization between two servers
- Disable automatic topology generation for a site
- Disable automatic topology cleanup for a site
- Disable minimum hops topology for a site
- Disable automatic stale server detection for a site
- Disable automatic inter-site topology generation for a site
- Disable inbound replication on a domain controller
- Disable outbound replication on a domain controller
- Enable reciprocal replication between sites (only for IP transport links)
- Enable change notification between sites (only for IP transport links)
- Force replication topology generation

**Replication Monitoring Operators**
The following is the set of administrative tasks assigned to this role:

- Get replication latency information
- Get pending operations on a domain controller
- Get replication summary information

- Check replication status

**DNS Administrators Role**
The following is the set of administrative tasks assigned to this role:

- Install the DNS Server service on domain controllers
- Configure recursive name resolution settings
- Configure the forest root domain controller to host the DNS zone that corresponds to the forest root DNS name
- Configure the domain controllers for each regional domain to host the DNS zone that corresponds to the DNS name of the domain
- Configure the zone containing the forest-wide locator records to replicate to every DNS server in the forest by using the forest-wide DNS application partition

# Appendix N: Default Active Directory Service Administrator Groups

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Every installation of Active Directory has a number of default service administrative groups, some of which have all service management abilities. For example, members of the Enterprise Admins, Domain Admins, and Built-In Administrators groups can perform all administrative tasks that are involved in managing an Active Directory service deployment. For all practical purposes, these groups can be considered to have equal abilities.

By default, these accounts are granted access to directory and server resources when Active Directory is installed. The following table lists the default service administrator accounts and provides a brief description of each account, including the qualities that make each group a service administrative group. For Scope, a scope of Forest means that the group exists only on domain controllers in the forest root domain and that group members have privileges in all domains in the forest as well as in the configuration and schema directory partitions. A scope of Domain means that this group exists on domain controllers in every domain, but group members have privileges only in one domain.

For more information about these groups, see Appendix B: Default Active Directory Security Groups earlier in this document.

**Default Service Administrator Accounts**

| Account Name (Mnemonic) | Scope | Description |
|---|---|---|
| Enterprise Admins (EA) | Forest | This group is automatically added to the Administrators group in every domain in the forest, providing complete access to the configuration of all domain controllers. This group can modify the membership of all administrative groups. Its own membership can be modified only by the default service administrator groups in the root domain. This account is considered a service administrator. |
| Schema Admins (SA) | Forest | This group has full administrative access to the schema. The membership of this group can be modified by any of the service administrator groups in the root domain. This account is considered a service administrator because its members can modify the schema, which governs the structure and content of the entire directory. |
| Administrators (BA) | Domain | This built-in group controls access to all the domain controllers in its domain, and it can change the membership of all administrative groups. Its own membership can be modified by the default service administrator groups BA and DA in the domain, as well as the EA group. This group has the special privilege to take ownership of any object in the directory or any resource on a domain controller. This account is considered a service administrator because its members have full access to the domain controllers in the domain. |
| Domain Admins (DA) | Domain | This group controls access to all domain controllers in a domain, and it can modify the membership of all administrative accounts in the domain. Its own membership can be modified by the service administrator groups BA and DA in its domain, as well as the EA group. This is a service administrator account because its members have full access to a domain's domain controllers. |
| Server Operators (SO) | Domain | By default, this built-in group has no members, and it has access to server configuration options on domain controllers. Its membership is controlled by the service administrator groups BA and DA in the domain, as well as the EA group. It cannot change any administrative group memberships. This is a service administrator account because its members have physical access to domain controllers and they can perform maintenance tasks (such as backup and restore), and they have the ability to change binaries that are installed on the domain controllers. |

| | | |
|---|---|---|
| Account Operators (AO) | Domain | By default, this built-in group has no members. It can create, manage, and delete users and groups in the Users and Computers containers and other organizational units (OUs), but it cannot manage service administrator accounts. In addition, members of this group cannot create computer accounts in the Domain Controllers OU, although they can modify and delete them. Members of this group can log on locally to domain controllers, but they cannot shut down or restart them. As a best practice, do not add members to this group, and do not use it for any delegated administration. |
| Backup Operators (BO) | Domain | By default, this built-in group has no members, and it can perform backup and restore operations on domain controllers. Its membership can be modified by the default service administrator groups BA and DA in the domain, as well as the EA group. It cannot modify the membership of any administrative groups. While members of this group cannot change server settings or modify the configuration of the directory, they do have the permissions needed to replace files (including system files) on the domain controllers. Because of this, they are considered service administrators. |
| Administrator | DS Restore Mode | This special account is created during the Active Directory installation process, and it is not the same as the Administrator account in the Active Directory database. This account is only used to start the domain controller in Active Directory Restore mode. When it is in restore mode, this account has full access to the directory database, as well as files (including system files) on the domain controller. Because of this, it is considered a service administrator. |

# Appendix O: Active Directory Delegation Wizard File

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Delegation of Control Wizard allows you to delegate administrative tasks to users or groups within a specific administrative scope and is primarily used to delegate data administration. This tool is driven by a customizable text file (Delegwiz.inf) and ships with a base set of common administrative tasks.

The list of tasks that can be delegated through the Delegation Wizard is maintained in the Delegwiz.inf file, which is created in the <Windows installation directory>\Inf folder. Administrators can modify this file to add or delete items from the list of tasks that can be delegated.

The Delegwiz.inf file that ships with Windows Server 2003 can be used to delegate about 13 common administrative tasks. This appendix contains an updated version of the file that can be used to delegate more than 70 administrative tasks.

**To use this appendix to modify the Delegwiz.inf file**

1. Navigate to the <*Windows installation directory*>\Inf folder.
2. Back up the existing Delegwiz.inf file by copying it and renaming it to Delegwiz_backup.inf.
3. Open Notepad and create a new empty text file.
4. Copy the following into the text file:

```
[Version]
signature="$CHICAGO$"

[DelegationTemplates]

Templates = template1, template2, template3, template4, template5, template6,
template7, template8, template9, template10, template11, template12, template13,
template14, template15, template16, template17, template18, template19, template20,
template21, template22, template23,template24, template25, template26, template27,
template28, template29, template30, template31, template32, template33,template34,
template35, template36, template37, template38, template39, template40, template41,
template42, template43,template44, template45, template46, template47, template48,
template49, template50, template51, template52, template53,template54, template55,
template56, template57, template58, template59, template60, template61, template62,
template63,template64, template65, template66, template67, template68, template69,
template70
;-----------------------------------------------------------
[template1]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Create, delete, and manage user accounts"

ObjectTypes = SCOPE, user

[template1.SCOPE]
user=CC,DC

[template1.user]
@=GA
;-----------------------------------------------------------

;-----------------------------------------------------------
[template2]
AppliesToClasses=domainDNS,organizationalUnit,container
```

```
Description = "Reset user passwords and force password change at next logon"

ObjectTypes = user

[template2.user]
CONTROLRIGHT= "Reset Password"
pwdLastSet=RP,WP
;-----------------------------------------------------------


;-----------------------------------------------------------
[template3]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Read all user information"

ObjectTypes = user

[template3.user]
@=RP

;-----------------------------------------------------------
[template4]
AppliesToClasses = organizationalUnit,container

Description = "Create, delete and manage groups"

ObjectTypes = SCOPE, group

[template4.SCOPE]
group=CC,DC

[template4.group]
@=GA

;-----------------------------------------------------------


;-----------------------------------------------------------
[template5]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify the membership of a group"

ObjectTypes = group

[template5.group]
member=RP,WP
;-----------------------------------------------------------


;-----------------------------------------------------------
[template6]
AppliesToClasses = domainDNS

Description = "Join a computer to the domain"

ObjectTypes = SCOPE
```

```
[template6.SCOPE]
computer=CC
;-----------------------------------------------------------



;-----------------------------------------------------------
[template7]
AppliesToClasses = domainDNS,organizationalUnit,site

Description = "Manage Group Policy links"

ObjectTypes = SCOPE

[template7.SCOPE]
gPLink=RP,WP
gPOptions=RP,WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template8]
AppliesToClasses=domainDNS,organizationalUnit

Description = "Generate Resultant Set of Policy (Planning)"

ObjectTypes = SCOPE

[template8.SCOPE]
CONTROLRIGHT= "Generate Resultant Set of Policy (Planning)"
;-----------------------------------------------------------

;-----------------------------------------------------------
[template9]
AppliesToClasses=domainDNS,organizationalUnit

Description = "Generate Resultant Set of Policy (Logging)"

ObjectTypes = SCOPE

[template9.SCOPE]
CONTROLRIGHT= "Generate Resultant Set of Policy (Logging)"
;-----------------------------------------------------------

;-----------------------------------------------------------
[template10]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Create, delete, and manage inetOrgPerson accounts"

ObjectTypes = SCOPE, inetOrgPerson

[template10.SCOPE]
inetOrgPerson=CC,DC

[template10.inetOrgPerson]
@=GA
;-----------------------------------------------------------
```

```
;-----------------------------------------------------------
[template11]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Reset inetOrgPerson passwords and force password change at next
logon"

ObjectTypes = inetOrgPerson

[template11.inetOrgPerson]
CONTROLRIGHT= "Reset Password"
pwdLastSet=RP,WP
;-----------------------------------------------------------


;-----------------------------------------------------------
[template12]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Read all inetOrgPerson information"

ObjectTypes = inetOrgPerson

[template12.inetOrgPerson]
@=RP

;-----------------------------------------------------------

;-----------------------------------------------------------
[template13]
AppliesToClasses=container

Description = "Create, Delete, and Manage WMI Filters"

ObjectTypes = SCOPE, msWMI-Som

[template13.SCOPE]
msWMI-Som=CC,DC

[template13.msWMI-Som]
@=GA
;-----------------------------------------------------------

;-----------------------------------------------------------
[template14]
AppliesToClasses=domainDNS,organizationalUnit

Description = "Create an Organizational Unit"

ObjectTypes = SCOPE

[template14.SCOPE]
organizationalUnit=CC
;-----------------------------------------------------------

;-----------------------------------------------------------
[template15]
AppliesToClasses=domainDNS,organizationalUnit
```

```
Description = "Delete a child Organizational Unit"

ObjectTypes = SCOPE

[template15.SCOPE]
organizationalUnit=DC
;-----------------------------------------------------------

;-----------------------------------------------------------
[template16]
AppliesToClasses=organizationalUnit

Description = "Delete this Organizational Unit"

ObjectTypes = organizationalUnit

[template16.organizationalUnit]
@=SD
;-----------------------------------------------------------

;-----------------------------------------------------------
[template17]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Rename an Organizational Unit"

ObjectTypes = organizationalUnit

[template17.organizationalUnit]
ou=WP
name=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template18]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify Description of an Organizational Unit"

ObjectTypes = organizationalUnit

[template18.organizationalUnit]
description=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template19]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify Managed-By Information of an Organizational Unit"

ObjectTypes = organizationalUnit

[template19.organizationalUnit]
managedBy=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
```

```
[template20]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Delegate Control of an Organizational Unit"

ObjectTypes = organizationalUnit

[template20.organizationalUnit]
@=WD
;-----------------------------------------------------------

;-----------------------------------------------------------
[template21]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Create a group"

ObjectTypes = SCOPE

[template21.SCOPE]
group=CC
;-----------------------------------------------------------

;-----------------------------------------------------------
[template22]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Delete a child group"

ObjectTypes = SCOPE

[template22.SCOPE]
group=DC
;-----------------------------------------------------------

;-----------------------------------------------------------
[template23]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Delete this group"

ObjectTypes = group

[template23.group]
@=SD
;-----------------------------------------------------------

;-----------------------------------------------------------
[template24]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Rename a group"

ObjectTypes = group

[template24.group]
cn=WP
name=WP
;-----------------------------------------------------------
```

```
;-----------------------------------------------------------
[template25]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Specify the Pre-Windows 2000 compatible name for the group"

ObjectTypes = group

[template25.group]
sAMAccountName=WP
;-------------------------------------------------------------


;-----------------------------------------------------------
[template26]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify the description of a group"

ObjectTypes = group

[template26.group]
description=WP
;-------------------------------------------------------------

;-----------------------------------------------------------
[template27]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify the scope of the group"

ObjectTypes = group

[template27.group]
groupType=WP
;-------------------------------------------------------------

;-----------------------------------------------------------
[template28]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify the type of the group"

ObjectTypes = group

[template28.group]
groupType=WP
;-------------------------------------------------------------

;-----------------------------------------------------------
[template29]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify notes for a group"

ObjectTypes = group

[template29.group]
```

```
info=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template30]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify group membership"

ObjectTypes = group

[template30.group]
member=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template31]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Specify Managed-By Information of a Group"

ObjectTypes = group

[template31.group]
managedBy=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template32]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Create a computer account"

ObjectTypes = SCOPE

[template32.SCOPE]
computer=CC
;-----------------------------------------------------------

;-----------------------------------------------------------
[template33]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Delete a child computer account"

ObjectTypes = SCOPE

[template33.SCOPE]
computer=DC
;-----------------------------------------------------------

;-----------------------------------------------------------
[template34]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Delete this computer account"

ObjectTypes = computer
```

```
[template34.computer]
@=SD
;-----------------------------------------------------------

;-----------------------------------------------------------
[template35]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Rename a computer account"

ObjectTypes = computer

[template35.computer]
@=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template36]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Disable a computer account"

ObjectTypes = computer

[template36.computer]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template37]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Reset a computer account"

ObjectTypes = computer

[template37.computer]
CONTROLRIGHT= "Reset Password"
;-----------------------------------------------------------

;-----------------------------------------------------------
[template38]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Specify the computer's description"

ObjectTypes = computer

[template38.computer]
description=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template39]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Specify Managed-By information for a computer account"

ObjectTypes = computer
```

```
[template39.computer]
managedBy=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template40]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Specify that a computer account be trusted for delegation"

ObjectTypes = computer

[template40.computer]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template41]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Create a user account in disabled state"

ObjectTypes = SCOPE

[template41.SCOPE]
user=CC
;-----------------------------------------------------------

;-----------------------------------------------------------
[template42]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Create a user account"

ObjectTypes = SCOPE , user

[template42.SCOPE]
user=CC

[template42.user]
userAccountControl=WP
CONTROLRIGHT= "Reset Password"
;-----------------------------------------------------------

;-----------------------------------------------------------
[template43]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Delete a child user account"

ObjectTypes = SCOPE

[template43.SCOPE]
user=DC
;-----------------------------------------------------------

;-----------------------------------------------------------
[template44]
```

```
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Delete this user account"

ObjectTypes = user

[template44.user]
@=SD
;-----------------------------------------------------------

;-----------------------------------------------------------
[template45]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Rename a user account"

ObjectTypes = user

[template45.user]
cn=WP
name=WP
distinguishedName=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template46]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Disable a user account"

ObjectTypes = user

[template46.user]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template47]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Unlock a user account"

ObjectTypes = user

[template47.user]
lockoutTime=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template48]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Enable a disabled user account"

ObjectTypes = user

[template48.user]
userAccountControl=WP
;-----------------------------------------------------------
```

```
;-----------------------------------------------------------
[template49]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Reset a user account's password"

ObjectTypes = user

[template49.user]

CONTROLRIGHT= "Change Password"
;-------------------------------------------------------------

;-----------------------------------------------------------
[template50]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Force a user account to change the password at the next logon"

ObjectTypes = user

[template50.user]
CONTROLRIGHT= "Reset Password"
userPassword=WP
;-------------------------------------------------------------

;-----------------------------------------------------------
[template51]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify a user's display name"

ObjectTypes = user

[template51.user]
adminDisplayName=WP
;-------------------------------------------------------------

;-----------------------------------------------------------
[template52]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify a user account's description"

ObjectTypes = user

[template52.user]
description=WP
;-------------------------------------------------------------

;-----------------------------------------------------------
[template53]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify a user's office location"

ObjectTypes = user
```

```
[template53.user]
physicalDeliveryOfficeName=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template54]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify a user's telephone number"

ObjectTypes = user

[template54.user]
telephoneNumber=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template55]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify the location of a user's primary web page"

ObjectTypes = user

[template55.user]
wWWHomePage=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template56]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify a user's UPN"

ObjectTypes = user

[template56.user]
userPrincipalName=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template57]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify a user's Pre-Windows 2000 user logon name"

ObjectTypes = user

[template57.user]
sAMAccountName=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template58]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Modify the hours during which a user can log on"
```

```
ObjectTypes = user

[template58.user]
logonHours=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template59]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Specify the computers from which a user can log on"

ObjectTypes = user

[template59.user]
userWorkstations=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
;[template60]
;AppliesToClasses=domainDNS,organizationalUnit,container

;Description = "Set User cannot change password for a user account"

;ObjectTypes = user

;[template60.user]

;CONTROLRIGHT= "Change Password"
;-----------------------------------------------------------

;-----------------------------------------------------------
[template61]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Set Password Never Expires for a user account"

ObjectTypes = user

[template61.user]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template62]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Set Store Password Using Reversible Encryption for a user account"

ObjectTypes = user

[template62.user]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template63]
AppliesToClasses=domainDNS,organizationalUnit,container
```

```
Description = "Disable a user account"

ObjectTypes = user

[template63.user]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template64]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Set Smart card is required for interactive logon for a user account"

ObjectTypes = user

[template64.user]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template65]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Set Account is sensitive and cannot be delegated for a user account"

ObjectTypes = user

[template65.user]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template66]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Set Use DES encryption types for this account for a user account"

ObjectTypes = user

[template66.user]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template67]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Set Do not require Kerberos pre-authentication for a user account"

ObjectTypes = user

[template67.user]
userAccountControl=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template68]
AppliesToClasses=domainDNS,organizationalUnit,container
```

```
Description = "Specify the date when a user account expires"

ObjectTypes = user

[template68.user]
accountExpires=WP
;-----------------------------------------------------------

;-----------------------------------------------------------
[template69]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Specify a profile path for a user"

ObjectTypes = user

[template69.user]
profilePath=WP
;-----------------------------------------------------------


;-----------------------------------------------------------
[template70]
AppliesToClasses=domainDNS,organizationalUnit,container

Description = "Specify a logon script for a user"

ObjectTypes = user

[template70.user]
scriptPath=WP
;-----------------------------------------------------------
```

5. Save the text file as delegwiz.inf in the <*Windows installation directory*>\Inf folder.

   The Delegation Wizard will now allow you to delegate more than 70 data management administrative tasks.

| **Note** |
| --- |
| For more information about modifying Delegwiz.inf, see article 308404, "HOW TO: Customize the Task List in the Delegation Wizard," in the Microsoft Knowledge Baseon the Web at http://go.microsoft.com/fwlink/?LinkId=3202. |

# R2: Deployment (R2 only)

Applies To: Windows Server 2003 R2

The content in this section applies to features released in Windows Server 2003 R2.

**In this section**

Installing Windows Server 2003 R2

Step-by-Step Guide for Print Management

ADFS Step-by-Step Guide

Step-by-Step Guide to Deploying ADAM

Step-by-Step Guide to Deploying Password Synchronization

Step-by-Step Guide for Storage Manager for SANs

Step-by-Step Guide for File Server Resource Manager

Step-by-Step Guide to Setting Up Server for NIS

Preserving OEM Preactivation When Reinstalling Windows Server 2003 R2

# Installing Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

This documentation provides information about installing or upgrading the Microsoft® Windows Server™ 2003 R2 operating system and its components.

Windows Server 2003 R2 contains two product discs. Disc 1 contains the latest version of Windows Server 2003. Disc 2 contains the installation files for the Windows Server 2003 R2 optional components.

| **Note** |
| --- |
| Installing Disc 2 does not automatically install the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using Add/Remove Windows Components in Control Panel. You also have the option to add or upgrade server roles by using Manage Your Server. For more information about these options, see Installing Optional Components and Server Roles. |

**In this topic**


- Before You Install
- Installing or Upgrading to Windows Server 2003 R2
- Installing Optional Components and Server Roles
- Administering components from a computer running Windows XP
- Troubleshooting installation issues with Windows Server 2003 R2


© 2006 Microsoft Corporation. All rights reserved.

# Before You Install

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

Before you install Windows Server 2003 R2, do the following:

- **If you are upgrading, fully back up your servers.** Your backup should include all data and configuration information that is necessary for the computer to function. It is important to perform a backup of configuration information for servers, especially those that provide network infrastructure, such as Dynamic Host Configuration Protocol (DHCP). When you perform the backup, be sure to include the boot and system partitions and the System State data. Another way to back up configuration information is to create a backup set for Automated System Recovery.

- **Update the Active Directory schema on domain controllers.** In Windows Server 2003 Active Directory networks, you will need to update the Active Directory schema before:

  - Installing Windows Server 2003 R2 on a domain controller.

  - Configuring a Windows Server 2003 R2 server to be a domain controller by using the Active Directory Installation Wizard (Dcpromo.exe).

  To update the schema, on the schema master domain controller, insert Disc 2 and open the command prompt. Then, change to the Cmpnents\R2\Adprep folder of Disc 2, and type: **adprep.exe /forestprep**

- **Disconnect UPS devices.** If you have an uninterruptible power supply (UPS) connected to your target computer, disconnect the connecting serial cable before running Setup. Setup automatically attempts to detect devices connected to serial ports, and UPS equipment can cause problems with the detection process.

- **Consider the following features in on Disc 1 that might affect security:**

  - **Windows Firewall.** Windows Firewall is turned *off* by default and must be turned on to begin protecting systems. Before enabling Windows Firewall, you should understand the changes that it will make to your environment. If you choose to enable Windows Firewall, you should configure it with Security Configuration Wizard (SCW), or you should base the configuration on SCW security policies.

  - **Security Configuration Wizard (SCW).** SCW is the preferred method for configuring Windows Firewall and creating security policies for servers based on their roles. SCW provides a guided means of configuring Windows Firewall. Before you deploy SCW policies to your servers, you should research SCW, Windows Firewall, and best practices for both. For more information, see "Security Configuration Wizard for Windows Server 2003" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45503).

    In addition, you should conduct pilot deployment testing of SCW and Windows Firewall (or any other firewall you may be using).

  - **Internet Explorer.** If you are upgrading from Windows Server 2003, you should note the many changes to Microsoft Internet Explorer on Disc 1. For example, Internet Explorer has been updated to prevent unintentional downloads of misrepresented, malicious code. It also prevents automatic resizing of browser windows (which is sometimes used to extract sensitive data from computers). For more information about these changes, see "Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47557).

- **Check minimum system requirements.** Review the minimum and recommended system requirements for Windows Server 2003. To do this, see "System Requirements and Hardware Compatibility" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51190). In addition to these requirements, you will need approximately 100 megabytes (MB) of hard disk space for Disc 2.

- **Read the release notes located on both product discs.** The release notes that are located on Disc 1 (Relnotes.htm) contain information about late-breaking issues with Windows Server 2003. The release notes located on Disc 2 (R2relnotes.htm) contain information about late-breaking issues with Windows Server 2003 R2.

- **Provide Mass Storage Drivers.** If you have a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for your hard disk, confirm that the controller is designed for products in the Windows Server 2003 family by checking the hardware and software compatibility information in the Windows Server Catalog on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=37823).

  If your controller is compatible, but you are aware that the manufacturer has supplied a separate driver file for use with products in the Windows Server 2003 family, obtain the file (on a floppy disk or CD) before you begin Setup. During the early part of Setup, a line at the bottom of the screen will prompt you to press F6. Further prompts will guide you in supplying the driver file to Setup so that it can gain access to the mass storage controller.

  If you are not sure whether you must obtain a separate driver file from the manufacturer of your mass storage controller, you can try running Setup. If the controller is not supported by the driver files on the Setup CD and, therefore, requires a driver file that is supplied by the hardware manufacturer, Setup stops and displays a message saying that no disk devices can be found, or it displays an incomplete list of controllers. After you obtain the necessary driver file, restart Setup, and press F6 when you are prompted.

- **Prepare Mirror sets or other disk sets.** If you used the Microsoft Windows NT® 4.0 operating system to create a volume set, mirror set, stripe set, or stripe set with parity, and you want to run Setup for Disc 1, you must prepare the disk set first. For more information, click "Getting Ready for an Upgrade" in Readme.htm located on Disc 1. This issue is documented in the "Working with Volume, Mirror, or Stripe Sets or Stripe Sets with Parity" section.

- **Provide a Custom Hardware Abstraction Layer (HAL) File.** If you have a custom Hardware Abstraction Layer (HAL) file supplied by your computer manufacturer, before you begin Setup for Disc 1, locate the floppy disk or other medium containing the file. During the early part of Setup, a line at the bottom of the screen will prompt you to press F6; however, at this time, press F5 (*not* F6). After you press F5, follow the prompts to include your HAL file in the Setup process.

- **If you are remotely installing Disc 2, enable Terminal Services.** After you install Disc 1,you will need to manually enable Terminal Services on the target computer. This is because Terminal Services is disabled by default. To do this, right click **My Computer**, and then click **Properties**. On the **Remote** tab, select the **Allow users to connect remotely to this computer** check box.

**Important Installation Information**

Before you install Windows Server 2003 R2, consider the following:

- **The Windows Server 2003 R2 optional components are not installed automatically.** When you complete the installation of Disc 2, you are not installing the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using Add/Remove Windows Components in Control Panel. You also have the option to add or upgrade server roles by using Manage Your Server.

| **Note** |
| --- |
| Microsoft Management Console (MMC) 3.0 is installed automatically when you install Windows Server 2003 R2. |

- **You may need to upgrade a component.** If Active Directory Application Mode (ADAM) or Windows SharePoint Services is already installed on your computer, you can upgrade to the new version that is included in Windows Server 2003 R2. For more information about upgrading one of these components, see the Help for the component on the Windows Server 2003 R2 TechCenter Web site (http://go.microsoft.com/fwlink/?LinkId=45560).

- **You cannot uninstall Windows Server 2003 R2.** The Windows Server 2003 R2 source files are permanent, but you will be able to uninstall any optional components that you install.

- **Important information for 64-bit versions.** You can only install 64-bit versions of Disc 2 on computers running the 64-bit versions of Windows Server 2003 that is on Disc 1. In other words, you cannot install the 32-bit version of Windows Server 2003 R2 on a computer running the 64-bit version of Windows Server 2003.

# Installing or Upgrading to Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

Windows Server 2003 R2 contains two product discs. Disc 1 contains Windows Server 2003. Disc 2 contains the installation files for the Windows Server 2003 R2 optional components. Both discs use the same Product Key.

| **Note** |
| --- |
| Installing Disc 2 does not automatically install the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using Add/Remove Windows Components in Control Panel. You also have the option to add or upgrade server roles by using Manage Your Server. For more information about these options, see Installing Optional Components and Server Roles. |

**In this topic**


- Performing a Clean Installation of Windows Server 2003 R2
- Upgrading to Windows Server 2003 R2
- Command-line options for installing Windows Server 2003 R2

# Performing a Clean Installation of Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

Windows Server 2003 R2 contains two product discs. Disc 1 contains the latest version of Windows Server 2003. Disc 2 contains the installation files for the Windows Server 2003 R2 optional components. Both discs use the same Product Key.

**In this topic**


- Install from the product discs
- Install from a shared network resource
- Perform an unattended installation
- Install by using Remote Installation Services

# Install from the product discs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

To perform a clean installation of Windows Server 2003 R2 using the product discs, you will need to install both Disc 1 and Disc 2.

| **Note** |
| --- |
| Installing Windows Server 2003 R2 does not automatically install the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using Add/Remove Windows Components in Control Panel. You also have the option to add or upgrade server roles by using Manage Your Server. For more information about these options, see Installing Optional Components and Server Roles. |

**To install Windows Server 2003 R2 from the product discs**

1. Insert Disc 1 into your CD-ROM drive, and then restart the computer.
2. Follow the instructions in your computer's startup routine to start the computer from the CD.
3. Follow the instructions on your screen. When you are prompted to enter the Product Key, type the Product Key that comes with the two product discs.

   | **Note** |
   | --- |
   | For more information about installing Windows Server 2003 (Disc 1), see the Readme.htm located on the root of Disc 1. |

   | **Note** |
   | --- |
   | When the computer restarts after installation of Disc 1, do not press any key to boot from the CD. |

4. After you install Disc 1, log on to the computer as an administrator.
5. At the prompt, insert Disc 2, and then click **OK**.
6. Follow the instructions on your screen.

# Install from a shared network resource

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

To perform a clean installation of Windows Server 2003 R2 using a shared network resource, you will need to install the files located on both Disc 1 and Disc 2.

| **Note** |
| --- |
| Installing Windows Server 2003 R2 does not autmatically install the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using Add/Remove Windows Components in Control Panel. You also have the option to add or upgrade server roles by using Manage Your Server. For more information about these options, see Installing Optional Components and Server Roles. |

**To install Windows Server 2003 R2 from a shared network resource**

1. On a network server, create a shared resource called *ServerName**FolderName*, where *ServerName* is the name of your server, and *FolderName* is the name of the shared resource that you created.

   | **Important** |
   | --- |
   | Ensure that only system administrators have full access to this folder. Other users should have only Read and Execute permissions. |

2. Copy all of the files and folders from Disc 1 to *ServerName**FolderName*.
3. Copy all of the files and folders from Disc 2 to *ServerName**FolderName*.
4. Click **Yes** to over write the Autorun.inf file.
5. On the computer on which you want to install Windows Server 2003 R2, log on as an administrator and connect to the shared folder that contains the Setup files.
6. To first install Windows Server 2003, in the i386 folder, click **Winnt32.exe**. For 64-bit versions, this file is located in the Amd64 folder.
7. Follow the instructions on your screen. When you are prompted to enter the Product Key, type the Product Key that comes with the two product discs.
8. After Setup for Windows Server 2003 is complete, log on to the computer as an administrator. Setup for Disc 2 should start automatically. If it does not start automatically, connect to the shared folder that contains the Setup files and in the \Cmpnents\R2 folder, click **Setup2.exe**.
9. Follow the instructions on your screen.

# Perform an unattended installation

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

To learn about unattended installation, see the following:

- If you want to perform an unattended installation of both Disc 1 and Disc 2, see "Automating Windows Server 2003 R2 Setup" in Deploy.chm located in the \Docs folder Disc 2. This documentation will help you create and modify the Unattend.txt file.

- If you want to perform an unattended installation of only Disc 2 from the command-line, see Command-line options for installing Windows Server 2003 R2.

- For more information about running unattended Setup (also called automated installation), see the "Windows Server 2003 Deployment Kit: Automating and Customizing Installations" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50476).

- If you are upgrading to a volume licensed version of Windows Server 2003 R2 and you'd like to encrypt your product key, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=55123).

# Install by using Remote Installation Services

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

You can install Windows Server 2003 R2 on remote computers using Remote Installation Services (RIS).

If this is your first time using RIS, you will need to install and configure RIS before performing these procedures. To learn about installing and configuring RIS, see "How To Use Remote Installation Services to Install Windows Server 2003 on Remote Computers" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50199).

**Installing Using RIS**
In the first procedure in this section, you will create an image of Windows Server 2003 (Disc 1). In the second procedure, you will add the Windows Server 2003 R2 installation files (Windows Server 2003 R2 Disc 2) to the image you created.

- To create an image of Windows Server 2003 (Disc 1)
- To add the Windows Server 2003 R2 installation files (Disc 2) to the image

**To create an image of Windows Server 2003 (Disc 1)**

1. On a computer running RIS, log on as a user with administrative privileges.
2. Click **Start**, click **Run**, type **risetup.exe**, and then click **OK** to start the RIS Setup Wizard.
3. When the Welcome screen appears, click **Next**.
4. Select **Add a new OS image to this remote installation server**, and then click **Next**.
5. You are then prompted for the location of the client operating system installation files. Insert Windows Server 2003 R2 Disc 1 in the CD-ROM drive, and then click **Next**.

| **Note** |
| --- |
| Microsoft only supports the use of Microsoft media when creating a client operating system image. The use of non-Microsoft media is not supported. |

6. Type the folder name for the Disc 1 installation files on the RIS server, and then click **Next**.
7. Type a friendly description for the operating system image. For example, you might type **Windows-Server-2003**. This is displayed to users after they start a remote client and run the Client Installation Wizard.
8. Click **Next.**
9. If you have created an operating system image on this server in the past, you will be prompted to specify what Setup should do with the old client installation screens. Select one of the following:
   - If you want to use the old installation screens at a later time, select **Use the new client installation screens, and rename the old ones with a .bak extension**.
   - If you want to delete the old installation screens, select **Use the new client installation screens, and overwrite the old ones**.
   - If you have made modifications to the installation screens that you would like to preserve, select **Use the old client installation screens**.
10. Click **Next**.
11. Click **Finish**, and then click **Done**.

**To add the Windows Server 2003 R2 installation files (Disc 2) to the image**

1. Navigate to the image that you created in the previous procedure. The image will be located in *Drive*:*RemoteInstallFolder*\Setup*Language*\Images*ImageName*. For example, E:\RemoteInstall\Setup\English\Images\Windows-Server-2003.

2. In this directory, create an $OEM$ folder and a Cmpnents folder. For example, you should have the following:

   E:\RemoteInstall\Setup\English\Images\Windows-Server-2003\i386

   E:\RemoteInstall\Setup\English\Images\Windows-Server-2003\$OEM$

   E:\RemoteInstall\Setup\English\Images\Windows-Server-2003\Cmpnents

3. In the $OEM$ folder, create the following directory: $OEM$\$1\R2.

4. In the $OEM$\$1\R2 directory, create a file named R2comp.inf.

5. In R2comp.inf, create a [Version] section and then create a [Components] section that lists the components that you would like to install. For example, the following example installs Active Directory Application Mode (ADAM) and Common Log File System:

   ```
   [Version]
   Signature = "$Windows NT$"
   [Components]
   ADAM=on
   CLFS=on
   ```

   For a complete list of the components that you can install, see Deploy.chm located in the \Docs folder of Windows Server 2003 R2 Disc 2.

6. Copy the content of the Cmpnents folder of Disc 2, into E:\RemoteInstall\Setup\English\Images\Windows-Server-2003\Cmpnents

7. Open the RISstndrd.sif file that is located in:

   E:\RemoteInstall\Setup\English\Images\Windows-Server-2003\i386\templates.

   | **Note** |
   | --- |
   | For x64-based computers, this file is located in the Amd64\Templates folder. |

8. Make the following changes in RISstndrd.sif. This change specifies that Setup should copy the subfolders and files contained in the $OEM$ folder.

   ```
   [Unattended]
   OemPreinstall = yes
   ```

9. In RISstndrd.sif, add a [GuiRunonce] section. The [GuiRunonce] section contains commands to execute the first time an administrator logs on to the computer. In the following example, the /q option specifies that Setup should run in quiet mode, the /cs option creates a shortcut on the desktop to a document that describes the new Windows Server 2003 R2 components, and the /a option specifies that you accept the EULA. In addition, starting the command with "cmd /c" ensures that Windows Server 2003 R2 Setup is complete before executing the next command, and "ProductID" is the Product Key that comes with the two product discs. For a complete list of the options that you can add to this section, see Deploy.chm located in the \Docs folder of Windows Server 2003 R2 Disc 2.

   ```
   [GuiRunonce]
   Command0="cmd /c
   \\RISServer\reminst\setup\Language\ImageName\Cmpnents\R2\setup2.exe /q /cs /a”
   Command1="%windir%\system32\sysocmgr.exe /i:sysoc.inf /u:C:\R2\r2comp.inf"
   [UserData]
   ProductID = "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
   ```

10. If you want the administrator to be able to install a Windows Server 2003 R2 component in the future, copy the contents of Disc 2 into E:\RemoteInstall\Setup\English\Images\Windows-Server-2003\$OEM$\$1\R2. After installation, these files will be available in C:\R2 on the client computer.

# Upgrading to Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Windows Server 2003 R2 contains two product discs. Disc 1 contains Windows Server 2003 with Service Pack 1 (SP1). Disc 2 contains the installation files for the Windows Server 2003 R2 optional components. Both discs use the same Product Key.

If you have Service Pack 1 (SP1) installed, you will only need to install Disc 2. If you are upgrading and you do not have SP1 installed, you will need to install both product discs.

| **Note** |
| --- |
| If you do not yet have Windows Server 2003 SP1 installed, you can also download and install SP1 from the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43106). You might want to do this if you do not want to install the integrated version of SP1 that is included on Disc 1. The integrated version is included within the operating system, and therefore, you cannot uninstall it. |

With a Windows Server 2003 R2 license, you can upgrade to Windows Server 2003 R2 if your computer is running one of the following operating systems:

- Windows Server 2003 with SP1
- Windows Server 2003 without SP1
- Windows 2000 Server
- Windows NT Server 4.0 with Service Pack 5 or later
- Windows NT Server 4.0, Terminal Server Edition, with Service Pack 5 or later

**In this topic**


- Upgrade from the product discs
- Upgrade from a shared network resource

# Upgrade from the product discs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Windows Server 2003 R2 contains two product discs. If you have Windows Server 2003 Service Pack 1 (SP1) installed, you will only need to install Disc 2. If you do not have SP1 installed, you will need to install both product discs. You can upgrade to Windows Server 2003 R2 from computers running the following operating systems:

- Windows Server 2003 with SP1
- Windows Server 2003 without SP1
- Windows 2000 Server
- Windows NT Server 4.0 with Service Pack 5 or later
- Windows NT Server 4.0, Terminal Server Edition, with Service Pack 5 or later

| **Note** |
| --- |
| When you install Windows Server 2003 R2, you are not installing the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using Add/Remove Windows Components in Control Panel. You also have the option to add or upgrade server roles by using Manage Your Server. For more information about these options, see Installing Optional Components and Server Roles. |

**Upgrading by using the product discs**

Choose from the following procedures:

- To upgrade from Windows Server 2003 with SP1 by using the product discs
- To upgrade from Windows NT, Windows 2000, or Windows Server 2003 without SP1 by using the product discs

### To upgrade from Windows Server 2003 with SP1 by using the product discs

1. Log on to the computer as an administrator.
2. Insert Disc 2 into your CD-ROM drive. If the disc does not automatically open, open the disc, and in the \Cmpnents\R2 folder, click **Setup2.exe**.
3. Follow the instructions on your screen. When you enter the Product Key, enter the Product Key that comes with the two product discs.

### To upgrade from Windows NT, Windows 2000, or Windows Server 2003 without SP1 by using the product discs

1. Log on to the computer as an administrator.
2. To first upgrade to Windows Server 2003 with SP1, insert Disc 1 into your CD-ROM drive. If the disc does not automatically open, open the disc, and then click **Setup.exe**.

   | **Note** |
   | --- |
   | If you do not yet have Windows Server 2003 SP1 installed, you can also download and install SP1 from the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43106). You might want to do this if you do not want to install the integrated version of SP1 that is included on Disc 1. The integrated version is included within the operating system, and therefore, you cannot uninstall it. |

3. On the first screen, click the option to install the operating system.

4. In **Installation Type**, select **Upgrade (Recommended)**.
5. Click **Next**, and continue to follow the instructions on your screen. When you enter the Product Key, enter the Product Key that comes with the two product CDs.

| **Note** |
| --- |
| For more information about upgrading to SP1, see Readme.htm located on Disc 1. |

6. After Setup for Disc 1 is complete, log on to the computer as an administrator.
7. When you are prompted, insert Disc 2 and then click **OK**.
8. Follow the instructions on your screen.

# Upgrade from a shared network resource

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

To upgrade to Windows Server 2003 R2 using a shared network resource, you will need to copy the source files to a network share. If you have Windows Server 2003 Service Pack 1 (SP1) installed, you will only need to copy and install the Disc 2 source files. If you do not have SP1 installed, you will need to copy and install the source files of both product discs.

You can upgrade to Windows Server 2003 R2 from computers running the following operating systems:

- Windows Server 2003 with Service Pack 1 (SP1)
- Windows Server 2003 without SP1
- Windows 2000 Server
- Windows NT Server 4.0 with Service Pack 5 or later
- Windows NT Server 4.0, Terminal Server Edition, with Service Pack 5 or later

**Note**

When you install Windows Server 2003 R2, you are not installing the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using Add/Remove Windows Components in Control Panel. You also have the option to add or upgrade server roles by using Manage Your Server. For more information about these options, see Installing Optional Components and Server Roles.

**Note**

If you do not yet have Windows Server 2003 SP1 installed, you can also download and install SP1 from the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43106). You might want to do this if you do not want to install the integrated version of SP1 that is included on Disc 1. The integrated version is included within the operating system, and therefore, you cannot uninstall it.

**Upgrading from a shared network resource**

Choose from the following procedures:

- To upgrade a computer running Windows Server 2003 with SP1
- To upgrade a computer running a Windows Server 2003 (without SP1), Windows NT, or Windows 2000

### To upgrade a computer running Windows Server 2003 with SP1 from a shared network resource

1. Log on to the computer as an administrator.
2. On a network server, create a shared resource called *ServerName**FolderName*, where *ServerName* is the name of your server, and *FolderName* is the name of the shared resource that you created.

   **Important**

   Ensure that only system administrators have full access to this folder. Other users should have only Read and Execute permissions.

3. Copy all of the files and folders from Disc 2 to the shared resource.
4. On the computer that you want to upgrade to Windows Server 2003 R2, connect to the shared folder that contains the Setup files.

5. In the \Cmpnents\R2 folder, click **Setup2.exe**.
6. Follow the instructions on your screen. When you are prompted to enter the Product Key, enter the Product Key that comes with the two product discs.

**To upgrade a computer running Windows Server 2003 (without SP1), Windows NT, or Windows 2000 from a shared network resource**

1. Log on to the computer as an administrator.
2. On a network server, create a shared resource called *ServerName**FolderName*, where *ServerName* is the name of your server, and *FolderName* is the name of the shared resource that you created.

| **Important** |
| --- |
| Ensure that only system administrators have full access to this folder. Other users should have only Read and Execute permissions. |

3. Copy all of the files and folders from Disc 1 to the shared resource.
4. Copy all of the files and folders from Disc 2 to the shared resource.
5. Click **Yes** to overwrite the Autorun.inf file.
6. On the computer that you want to upgrade to Windows Server 2003 R2, connect to the shared folder that contains the Setup files.
7. To first install Windows Server 2003 with SP1, in the i386 folder, click **Winnt32.exe**. For 64-bit versions, this file is located in the Amd64 folder.
8. Follow the instructions on your screen. When you are prompted to enter the Product Key, enter the Product Key that comes with the two product discs.

| **Note** |
| --- |
| For more information about upgrading to SP1, see Readme.htm located on Disc 1. |

9. After Setup for Windows Server 2003 with SP1 is complete, log on to the computer as an administrator. Setup for Disc 2 should start automatically. If it does not start automatically, connect to the shared folder that contains the Setup files and in the \Cmpnents\R2 folder, click **Setup2.exe**.
10. Follow the instructions on your screen.

# Upgrading to Windows Server 2003 R2

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

Windows Server 2003 R2 contains two product discs. Disc 1 contains Windows Server 2003. Disc 2 contains the installation files for the Windows Server 2003 R2 optional components. Both discs use the same product key.

The upgrade scenarios are different depending on the version of this operating system that you have.

- If you have the first release of Windows Server 2003 R2, which contains Service Pack 1 (SP1), see Upgrading when you have the SP1 version of Windows Server 2003 R2. If the discs that you have do not specify which service pack is included, then the discs contain SP1.

- If you have the second release of Windows Server 2003 R2, which contains Service Pack 2, see Upgrading when you have the SP2 version of Windows Server 2003 R2. To verify that you have this version, ensure that "includes Service Pack 2" is labeled on your product discs. If the discs that you have do not specify which service pack is included, then the discs contain SP1.

| **Note** |
| --- |
| When you install Windows Server 2003 R2, you are not installing the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using **Add/Remove Windows Components** in Control Panel. You also have the option to add or upgrade server roles using **Manage Your Server**. For more information about these options, see Installing Optional Components and Server Roles. |

# Upgrading when you have the SP1 version of Windows Server 2003 R2

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

| Important |
| --- |
| **This topic applies only if you have the first release of Windows Server 2003 R2 where Disc 1 contains Service Pack 1 (SP1) . If the product discs that you have do not specify which service pack is included, then the discs contain SP1.** |

You can upgrade to Windows Server 2003 R2 from computers running the following operating systems:

- Windows Server 2003 with a service pack
- Windows Server 2003 without a service pack
- Windows 2000 Server
- Windows NT Server 4.0 with Service Pack 5 or later
- Windows NT Server 4.0 Terminal Server Edition with Service Pack 5 or later

| Note |
| --- |
| When you install Windows Server 2003 R2, you are not installing the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using **Add/Remove Windows Components** in Control Panel. You also have the option to add or upgrade server roles using **Manage Your Server**. For more information about these options, see Installing Optional Components and Server Roles. |

| Note |
| --- |
| If you are running Windows Server 2003 and you do not have the latest service pack installed, you can download and install the service pack from the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=77236). Use this method if you do not want to install the integrated version that is included on Disc 1. The service pack is included within the operating system, and therefore, you cannot uninstall it alone. |

**In this topic**


- Upgrading by using the product discs
- Upgrading from a shared network resource


**Upgrading by using the product discs**

Choose from the following procedures:

- To upgrade from Windows Server 2003 with a service pack by using the product discs
- To upgrade from Windows NT, Windows 2000, or Windows Server 2003 without a service pack by using the product discs

### To upgrade from Windows Server 2003 with a service pack by using the product discs

1. Log on to the computer as an administrator.
2. Insert Disc 2 into your CD-ROM drive. If the disc does not automatically open, type **D:r2auto.exe** at the command line. D

is the drive letter associated with your CD-ROM drive.

3. Follow the instructions on your screen, and enter the product key that comes with the two product discs.

**To upgrade from Windows NT, Windows 2000, or Windows Server 2003 without a service pack by using the product discs**

1. Log on to the computer as an administrator.
2. If you are running Windows NT or Windows 2000, first upgrade to Windows Server 2003 by inserting Disc 1 into your CD-ROM drive. If the disc does not automatically open, open the disc, and then click **Setup.exe**.
3. On the first screen, select the option to install the operating system.
4. In **Installation Type**, select **Upgrade (Recommended)**.
5. Click **Next**, and continue to follow the instructions on your screen. Enter the product key that comes with the two product disks.
6. After Setup for Disc 1 is complete, log on to the computer as an administrator.
7. When you are prompted, insert Disc 2 and then click **OK**.
8. Follow the instructions on your screen.

**Upgrading from a shared network resource**

To upgrade to Windows Server 2003 R2 using a shared network resource, you need to copy the source files to a network share. If you have Windows Server 2003 with a service pack installed, copy and install only the Disc 2 source files. If you do not have a service pack installed, copy and install the source files of both product discs.

Choose from the following procedures:

- To upgrade a computer running Windows Server 2003 with a service pack
- To upgrade a computer running a Windows Server 2003 (without a service pack), Windows NT, or Windows 2000

**To upgrade a computer running Windows Server 2003 with a service pack from a shared network resource**

1. Log on to the computer as an administrator.
2. On a network server, create a shared resource called *ServerName**FolderName*, where *ServerName* is the name of your server, and *FolderName* is the name of the shared resource that you created.

| **Important** |
| --- |
| Ensure that only system administrators have full access to this folder. Other users should have only Read and Execute permissions. |

3. Copy all of the files and folders from only Disc 2 to the shared resource.
4. Open a command prompt on the computer that you want to upgrade to Windows Server 2003 R2, connect to the shared folder that contains the source files, and run **D:r2auto.exe**. D is the drive letter associated with your CD-ROM drive.
5. Follow the instructions on your screen, and enter the product key that comes with the two product discs.

**To upgrade a computer running Windows Server 2003 (without a service pack), Windows NT, or Windows 2000 from a shared network resource**

1. Log on to the computer as an administrator.
2. On a network server, create a shared resource called *ServerName**FolderName*, where *ServerName* is the name of your server, and *FolderName* is the name of the shared resource that you created.

| **Important** |
| --- |

Ensure that only system administrators have full access to this folder. Other users should have only Read and Execute permissions.

3. Copy all of the files and folders from Disc 1 to the shared resource.
4. Copy all of the files and folders from Disc 2 to the shared resource.
5. Click **Yes** to overwrite the Autorun.inf file.
6. On the computer that you want to upgrade to Windows Server 2003 R2, connect to the shared folder that contains the source files.
7. If you are running Windows NT or Windows 2000, first upgrade to Windows Server 2003: In the i386 folder, click **Winnt32.exe**. For 64-bit versions, this file is located in the Amd64 folder.
8. Follow the instructions on your screen, and enter the product key that comes with the two product discs.
9. After Setup for Windows Server 2003 is complete, log on to the computer as an administrator. Setup for Disc 2 should start automatically. If it does not start automatically, connect to the shared folder that contains the source files and run **D:r2auto.exe**. D is the drive letter associated with your CD-ROM drive.
10. Follow the instructions on your screen.

# Upgrading when you have the SP2 version of Windows Server 2003 R2

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

| **Important** |
| --- |
| **This topic applies to you only if you have the second release of Windows Server 2003 R2 which contains Service Pack 2 (SP2). To verify that you have this version, ensure that "includes Service Pack 2" is labeled on your product discs. If the discs that you have do not specify which service pack is included, then the discs contain Service Pack 1 (SP1).** |

You can upgrade to Windows Server 2003 R2 from computers running the following operating systems:

- Windows Server 2003
- Windows Server 2003 with SP1
- Windows 2000 Server
- Windows NT Server 4.0 with Service Pack 5 or later
- Windows NT Server 4.0 Terminal Server Edition with Service Pack 5 or later

| **Note** |
| --- |
| If you are running Windows Server 2003 and you do not have the latest service pack installed, you can download and install it from the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=77236). Use this method if you do not want to install the integrated version that is included on Disc 1. The service pack is included within the operating system, and therefore, you cannot uninstall it alone. |

| **Note** |
| --- |
| When you install Windows Server 2003 R2, you are not installing the Windows Server 2003 R2 optional components. After Setup is complete, you can install the optional components by using **Add/Remove Windows Components** in Control Panel. You also have the option to add or upgrade server roles using **Manage Your Server**. For more information about these options, see Installing Optional Components and Server Roles. |

| **Note** |
| --- |
| If you install Disc 1, you can only install the Disc 2 that came with your version of Windows Server 2003 R2—you cannot install the Disc 2 that came with a previous version of Windows Server 2003 R2. |

**In this topic**


- Upgrading by using the product discs
- Upgrading from a shared network resource
- Upgrade Diagram


**Upgrading by using the product discs**

To upgrade to Windows Server 2003 R2, you need to install both product discs.

**To upgrade to Windows Server 2003 R2**

1. Log on to the computer as an administrator.
2. To first upgrade to the latest version of Windows Server 2003, insert Disc 1 into your CD-ROM drive. If the disc does not automatically open, open the disc, and then click **Setup.exe**.
3. On the first screen, select the option to install the operating system.
4. In **Installation Type**, select **Upgrade (Recommended)**.
5. Click **Next**, and continue to follow the instructions on your screen, and enter the product key that comes with the two product CDs.
6. After Setup for Disc 1 is complete, log on to the computer as an administrator.
7. When you are prompted, insert Disc 2 and then click **OK**.
8. Follow the instructions on your screen.

**Upgrading from a shared network resource**
To upgrade to Windows Server 2003 R2 using a shared network resource, you need to copy the source files from both discs to a network share.

**To upgrade to Windows Server 2003 R2**

1. Log on to the computer as an administrator.
2. On a network server, create a shared resource called *ServerName**FolderName*, where *ServerName* is the name of your server, and *FolderName* is the name of the shared resource that you created.

| **Important** |
| --- |
| Ensure that only system administrators have full access to this folder. Other users should have only Read and Execute permissions. |

3. Copy all of the files and folders from Disc 1 to the shared resource.
4. Copy all of the files and folders from Disc 2 to the shared resource.
5. Click **Yes** to overwrite the Autorun.inf file.
6. On the computer that you want to upgrade to Windows Server 2003 R2, connect to the shared folder that contains the source files.
7. To first install Windows Server 2003, in the i386 folder, click **Winnt32.exe**. For 64-bit versions, this file is located in the Amd64 folder.
8. Follow the instructions on your screen, and enter the product key that comes with the two product discs.
9. After Setup for Windows Server 2003 is complete, log on to the computer as an administrator. Setup for Disc 2 should start automatically. If it does not start automatically, connect to the shared folder that contains the source files and run **D:r2auto.exe**. D is the drive letter associated with your CD-ROM drive.
10. Follow the instructions on your screen.

**Upgrade diagram**
The following diagram shows the upgrade paths that are supported with this version of Windows Server 2003 R2.


Windows NT or Windows 2000
Windows Server 2003 or Windows Server 2003 with SP1
Install CD1
Install SP2 (Web download or CD)
Install CD1
Windows Server 2003 with SP2 update (SP2 can be uninstalled)
Integrated version of Windows Server 2003 with SP2 (SP2 cannot be uninstalled)
Install CD2
Install CD2
Install the SP1 version of CD2 (unsupported)
Windows Server 2003 R2 with SP2

# Command-line options for installing Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

You can define the way Windows Server 2003 R2 is installed by using the following command-line options when you run Setup2.exe. These options are not case-sensitive.

| Option | Description |
|---|---|
| **/q**<br>or<br>**/quiet** | Specifies that Windows Server 2003 R2 will be installed in quiet mode, without a user interface. When you use this option, no prompts appear during the installation process. |
| **/p:** *ProductKey*<br>or<br>**/productkey:** *ProductKey* | Specifies the Product Key that should be used to install Windows Server 2003 R2. This parameter is required if you have not yet entered a Windows Server 2003 R2 Product Key.<br>*ProductKey* must be specified in the *AAAAA-AAAAA-AAAAA-AAAAA-AAAAA* format. |
| **/a**<br>or<br>**/accepteula** | Specifies that you accept the EULA. This parameter is required if you did not use Disc 1 to install Windows Server 2003.<br>By using this option you are agreeing that you have read and accepted the terms of the applicable End User License Agreement (EULA) for Windows Server 2003 R2. You can find the following files in the \Cmpnents\R2 folder of Disc 2.<br>• If you have received this product from an original equipment manufacturer (OEM), the applicable EULA is Eula_oem.txt.<br>• If you have acquired the product in a retail store, the applicable EULA is Eula_retail.txt.<br>• If you have an evaluation copy of Windows Server 2003 R2, the applicable EULA is Eula_eval.txt.<br>**Important**<br>If you are not the end-user of this computer, before you use this option you will need to first verify that the end-user (whether an individual or an organization) has received, read, and accepted the terms of the Windows Server 2003 R2 EULA. |
| **/cs**<br>or<br>**/createshortcut** | Creates a shortcut on the desktop to a document that describes the new Windows Server 2003 R2 components. |
| **/sr**<br>or<br>**/suppressreboot** | Suppresses restart after the installation of Windows Server 2003 R2 is complete. |
| **/?** | Displays the command-line options that you can use to install Windows Server 2003 R2. |

**Note**

If you are using Setup2.exe in a script or answer file, you should start your command with **cmd /c**. This ensures that Windows Server 2003 R2 Setup is complete before executing the next command, for example:**cmd /c setup2.exe /q /a /p:***XXXXX-XXXXX-XXXXX-XXXXX-XXXXX* **/cs**

**Note**

For information about how to create and modify the Unattend.txt file and to perform an unattended installation of Windows Server 2003 R2, see "Automating Windows Server 2003 R2 Setup" in Deploy.chm located in the \Docs folder of Windows Server 2003 R2 Disc 2.

# Installing Optional Components and Server Roles

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

After Setup is complete, you can install the optional components by using Add/Remove Windows Components in Control Panel. You also have the option to add or upgrade server roles by using Manage Your Server.

| **Note** |
| --- |
| For specific instructions on upgrading ADAM or Windows SharePoint Services to the new version that is included in Windows Server 2003 R2, see the Help for the component at the Windows Server 2003 R2 TechCenter Web site (http://go.microsoft.com/fwlink/?LinkId=45560). |

**In this topic**


- Add or upgrade server roles by using Manage Your Server
- Install an optional component by using Add/Remove Windows Components
- Find a component in Add/Remove Windows Components

# Add or upgrade server roles by using Manage Your Server

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

After Setup is complete, you can add or upgrade server roles by using Manage Your Server. The File server role and the Print server role contain updates with Windows Server 2003 R2. Additionally, the Windows® SharePoint® Services role is a new server role with Windows Server 2003 R2.

| **Note** |
| --- |
| For detailed instructions on upgrading ADAM or Windows SharePoint Services to the new version that is included in Windows Server 2003 R2, see the Help for the component at the Windows Server 2003 R2 TechCenter Web site (http://go.microsoft.com/fwlink/?LinkId=45560). |

**To add or upgrade server roles by using Manage Your Server**

1. Log on to the computer as an administrator.
2. Click **Start**, click **Administrative Tools**, and then click **Manage Your Server**.
3. Click **Add or remove a role**.
4. Follow the instructions in Configure Your Server Wizard. When you get to the **Server Role** page, select one of the following roles:

   - **File Server**

     Windows Server 2003 R2 components for file server include File Server Management, DFS Management, DFS Replication Service, File Server Resource Manager, Storage Management for SANs, Microsoft Services for Network File System, and Services for Macintosh.

   - **SharePoint Services**

     Windows Server 2003 R2 components for SharePoint Services includes Windows SharePoint Services**.**

   - **Print Server**

     Windows Server 2003 R2 component for Print Server is Print Management.

   | **Note** |
   | --- |
   | During installation, you might be prompted to insert the Windows Server 2003 R2 Disc 2 (where the Windows Server 2003 R2 Setup files are located) or to specify the location of the Windows Server 2003 R2 Setup files on your hard disk. To avoid receiving this dialog box in the future, see Troubleshooting installation issues with Windows Server 2003 R2. |

5. The Configure Your Server Wizard will install the components associated with the role that you select.

# Install an optional component by using Add/Remove Windows Components

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

After Setup is complete, you can install Windows Server 2003 R2 optional components by using Add/Remove Windows Components in Control Panel. This is a generic procedure and you might need more specific instructions to install your component. For more detailed information about installing a component, see the Help for the component at the Windows Server 2003 R2 TechCenter Web site (http://go.microsoft.com/fwlink/?LinkId=45560).

| **Note** |
| --- |
| For detailed instructions on upgrading ADAM or Windows SharePoint Services to the new version that is included in Windows Server 2003 R2, see the Help for the component at the Windows Server 2003 R2 TechCenter Web site (http://go.microsoft.com/fwlink/?LinkId=45560). |

**To install a component by using Add/Remove Windows Components**

1. Log on to the computer as an administrator.
2. Open **Add or Remove Programs** in Control Panel.
3. Click **Add/Remove Windows Components**.
4. In Windows Components Wizard, locate the component that you want to install. To locate a component, you can click **Details**. For example, to install ADAM, click **Active Directory**, and then click **Details**. If you cannot find a component, see Find a component in Add/Remove Windows Components.
5. Select the component you want to install, and then click **Next**.

   | **Note** |
   | --- |
   | If an item is selected in the Windows Components Wizard, it has already been installed. |

6. If necessary, follow the instructions in the Windows Components Wizard.

   | **Note** |
   | --- |
   | You might be prompted to insert the Windows Server 2003 Disc 2 (where the Windows Server 2003 R2 Setup files are located) or to specify the location of the Windows Server 2003 R2 Setup files on your hard disk. To avoid receiving this dialog box in the future, see Troubleshooting installation issues with Windows Server 2003 R2. |

7. When installation is complete, click **Finish**.

# Find a component in Add/Remove Windows Components

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

The following table shows the relationships between the top-level components and their respective sub-level components located in Add/Remove Windows Components. To find and install a certain sub-level component, locate the associated top-level component, and then click **Details**. For example, to install ADAM, click **Active Directory Services**, and then click **Details**.

| Top-level Components | Sub-level 1 Components | Sub-level 2 and 3 Components |
|---|---|---|
| **Active Directory Services** | • Active Directory Application Mode (ADAM) | Not applicable |
| | • Active Directory Federation Services (ADFS) | • ADFS Web Services Agents<br>  • Claims-aware Applications<br>  • Windows NT token-based applications<br>• Federation Service<br>• Federation Service Proxy |
| | • Identity Management for UNIX | • Administration Components<br>• Password Synchronization<br>• Server For NIS |
| **Distributed File System** | • DFS Management<br>• DFS Replication Service | Not applicable |
| **Management and Monitoring Tools** | • File Server Resource Manager<br>• Hardware Management<br>• Print Management Console Component<br>• Storage Manager for SANs | Not applicable |
| **Microsoft .NET Framework 2.0** | Not applicable | Not applicable |
| **Other Network File and Print Services** | • Common Log File System | Not applicable |

| | | |
|---|---|---|
| | • Microsoft Services for NFS | • Client for NFS<br>• Microsoft Services for NFS Administration<br>• RPC External Data Representation<br>• RPC Port Mapper<br>• Server for NFS<br>• Server For NFS Authentication<br>• User Name Mapping |
| **Subsystem for UNIX-based Applications** | Not applicable | Not applicable |
| **Windows SharePoint Services** | Not applicable | Not applicable |

# Administering components from a computer running Windows XP

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

To make the remote management of your servers easier, you can install the Administration Tools Packs for Windows Server 2003 R2 located on Disc 2 in the \admin directory. The following components each have an Administration Tools Pack:

- Microsoft Management Console (MMC) 3.0
- Print Management
- File Server Resource Manager
- DFS Replication
- Identity Management for UNIX
- File Server Management

In order to administer a component, you must first install the MMC 3.0 Administration Tools Pack. The Administration Tools Packs can be installed on computers running Windows XP Professional with Service Pack 2 (SP2).

**To install or reinstall an Administration Tools Pack for Windows Server 2003 R2**

1. Insert Disc 2 into your CD-ROM drive.
2. Open the command prompt, and type the path to the \admin directory of Disc 2. For example, if the CD drive is drive E, type: **E:\admin**.
3. If you do not have MMC 3.0 installed on your computer, right-click the Administration Tools Pack for MMC 3.0, click **Open**, and then follow the instructions on your screen.
4. Right-click the Administration Tools Pack you would like to install, and then click **Open**.
5. Follow the instructions on your screen to install the administration pack.

| **Note** |
| --- |
| To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. |

| **Note** |
| --- |
| You can reinstall Windows Server 2003 R2 Administration Tools Pack at any time. Reinstallation will overwrite any prior installation. |

# Troubleshooting installation issues with Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP2

The following are specific issues that you may encounter when deploying this operating system. You should review this information to help resolve deployment problems.

**Stop Messages**
This section outlines the steps to troubleshoot Stop messages or unexpected shutdowns, which are the result of a system error in the operating system. Following a system error, your computer stops responding, and a message is displayed on a blue or black background.

You can use the System Information tool to gather important information about system settings, problem devices, and hardware resources. For more information, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51216).

| **Important** |
| --- |
| To ensure system stability, install the latest service pack that is appropriate for your system (follow service pack instructions) and use only signed drivers. For more information, see "Microsoft Update" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=284). |

Follow these steps for any Stop error that does not display identifying text or specific troubleshooting steps:

1. Use the Microsoft Online Crash Analysis tool on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=11699). You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error.

2. Use the File Signature Verification tool (Sigverif.exe), which identifies unsigned drivers and incompatible system files on your computer. The system files and device driver files that are provided with the Windows Server 2003 family have a Microsoft digital signature, which indicates that the files are original, unaltered system files or that they have been approved by Microsoft for use with Windows.

3. If the Stop message is caused by a driver problem, you can use Device Manager to roll back to a previous version of the driver. For more information, see "Roll back to the previous version of a driver" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51192). You may need to disable any newly installed drivers.

4. Using a recent version of virus-protection software, check for viruses on your computer. If a virus is found, perform the steps required to eliminate it from your computer. See your antivirus software documentation for these steps.

5. For updated information about Stop messages, you can search the Support Knowledge Base (KB) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=46046). When typing keywords, use the Stop message number (for example, **stop 0x0000000A**).

   For more information about articles, troubleshooting wizards, and items that you can download, see "Updated technical information" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51193).

6. If you can start the operating system, check the System Log in Event Viewer for additional information that might help determine the device or driver that is causing the problem. Many Stop messages produce specific events that you can search for, such as Event ID 1001 (generated by various Stop messages) and Event ID 6008 (usually caused by an unexpected shutdown or a computer that does not restart correctly). To display a description of an event, in Event Viewer, double-click the event. For more information about Event Viewer, see "Event Viewer" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51195).

For updated information about Event IDs, you can search the Support Knowledge Base (KB) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=46046).

7. If you cannot start your computer, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see "Start the computer in Safe Mode" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51196). For more information about how to start your computer in Last Known Good Configuration, see "Start the computer using the last known good configuration" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51197).

| **Important** |
| --- |
| When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost. |

8. Remove any newly installed hardware (RAM, adapters, hard disks, modems, and so on).

9. Ensure that you have updated drivers for your hardware devices. Also, ensure that you have the latest system BIOS or the correct firmware (for Itanium architecture-based computers). The device or hardware manufacturers can assist you in obtaining these items.

10. Run the system diagnostics supplied by your computer manufacturer, especially the memory check.

11. Confirm that your hardware is designed for the Windows Server 2003 family. For more information, see "System Requirements" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=100120).

12. If Safe Mode and other startup options do not work, you can consider using the Recovery Console. This method is recommended only if you are an advanced user who can use basic commands to identify and locate problem drivers and files. In addition, you must be an administrator to use the Recovery Console. For more information, see "Recovery Console overview" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51199).

For more information, see the following:

- For information about how to obtain product support, see "Windows Server 2003 Family: Support Centers" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=26478).

- For information about troubleshooting specific Stop messages, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51200).

- For articles about specific Stop messages, you can search the Support Knowledge Base (KB) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=46046).

- For information about support options, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51198).

- For information about troubleshooting specific technologies or components (for example, disks and file systems), see the "Windows Server 2003 Technical Reference" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4546).

**The Computer Will Not Start**

Products in the Windows Server 2003 family provide a variety of options you can use when a system will not start:

- **First option.** Try to restart your computer using Safe Mode or related startup options, one of which is Last Known Good Configuration. These options start the system with only the minimal, necessary services. Startup options are valuable for providing access to important system tools such as Device Manager (for example, you could start in Safe Mode and roll back a device driver) or Backup (for example, to restore a System State backup). To start your computer using one of these startup options, restart your computer and press F8 when a text-based or graphical screen appears and then disappears, but before the Windows logo appears. In other words, press F8 immediately after the power-on self test of the basic input/output system (the BIOS POST). There may be a prompt that tells you when to press F8.

  For more information, see "How to troubleshoot startup problems in Windows Server 2003" on the Microsoft Web site

(http://go.microsoft.com/fwlink/?LinkId=100119).

For detailed instructions, see "Recover a System That Will Not Start" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51201).

- **Second option.** An option to consider if Safe Mode does not help is the Recovery Console. This option is recommended only for advanced users or administrators. You can use the Recovery Console by starting the system from the Setup CD. On an x86-based (32-bit) computer, you can also install the Recovery Console as a startup option. The Recovery Console is a command-line interface from which you can perform tasks such as starting and stopping services and accessing the local disk (including disks formatted with NTFS).

| **Important** |
| --- |
| Because the Recovery Console is quite powerful, it is recommended only for advanced users or administrators. |

To start the computer and use the Recovery Console, insert Disc 1 into the CD drive and restart the computer. If prompted to press a key to start the computer from the CD, press a key. When the text-based part of Setup begins, follow the prompts. Choose the Recovery Console option by pressing **R** and continue to follow the instructions on your screen. When you get to the system prompt, type **help** for a list of commands, or **help** *commandname* for help on a specific command.

For more information, see "Recovery Console Technical Reference" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51202).

For detailed instructions, see "Start the computer and use the Recovery Console" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51203).

- **Third option.** If Safe Mode and the Recovery Console do not work in your situation, and if you made appropriate advance preparations (preparing appropriate recovery media with the Automated System Recovery Preparation Wizard in Backup), you can try Automated System Recovery. Automated System Recovery is a last-resort option that is not recommended for use before the other options in this list. To use this option, you start the system with the Setup CD and use the media that you have previously prepared.

  For more information, see "Automated System Recovery (ASR) overview" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51176).

  For detail instructions, see "Create an Automated System Recovery set using Backup" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51178).

- **Fourth option.** If you cannot repair the server using the previous methods in this list, you might be able to use Emergency Management Services to attempt to gain access to the server and examine the problem.

  For sources of information, see "Emergency Management Services" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51179).

#### Cannot Log On

- **Problem: I finished installing Disc 1 and now I cannot log on to the computer by using Terminal Services.**

  **Cause:** Terminal Services is disabled by default when you install Windows Server 2003 on Disc 1. Therefore, after you install Disc 1, you will not be able to connect to the computer without first manually enabling Terminal Services.

  **Solution:** On the target computer, manually enable Terminal Services. To do this, right click **My Computer**, and then click **Properties**. On the **Remote** tab, select the **Allow users to connect remotely to this computer** check box

#### Post Installation Issues

The following are specific problems and solutions that you may encounter when installing this operating system or when

installing a component.

- **Problem: During the installation of a Windows Server 2003 R2 component, I was prompted to insert the Windows Server 2003 R2 Disc 2 or to specify the location of the Setup files on my hard disk.**

  **Cause:** You will receive this dialog box if Windows Server 2003 was installed from a location different from where Windows Server 2003 R2 (Disc 1 and/or Disc 2) was installed. For example, if you installed Windows Server 2003 from a shared network resource, and then you installed Windows Server 2003 R2 by using the product CDs, you will receive this dialog box.

  **Solution:** To avoid receiving this dialog box, install Windows Server 2003 and the Windows Server 2003 R2 from the same location. For example, you can install the operating system by using the Windows Server 2003 R2 product CDs (both Disc 1 and Disc 2) that you obtained from Microsoft. Or, if you are installing from a network share, copy all of the files and folders from Disc 1 and Disc 2 to your shared network resource. When you receive a dialog box that asks if you would like to replace the Autorun.inf file, click **Yes**.

- **Problem: My installation completed successfully, but I was not prompted to insert the Disc 2 after the installation of Disc 1 was complete.**

  **Cause:** This can happen if you enter a Product Key other than the one that was included with the Windows Server 2003 R2 product discs.

  **Solution:** After the installation of Disc 1 is complete, insert Disc 2 to install Windows Server 2003 R2. When you are prompted, enter the Product Key that comes with the two product discs.

- **Problem: I am having problems upgrading a component to the new version that is included in Windows Server 2003 R2.**

  **Solution:** For specific instructions on upgrading an optional component, see the Help for the component on the Windows Server 2003 R2 TechCenter Web site (http://go.microsoft.com/fwlink/?LinkId=45560). The optional components that you can upgrade are Active Directory Application Mode (ADAM) and Windows SharePoint Services.

- **Problem: I installed Windows Server 2003 R2, but the new components are not installed on my computer.**

  **Cause:** When you install Disc 2, you are not installing the Windows Server 2003 R2 optional components.

  **Solution:** After Setup is complete, you can install the optional components by using Add/Remove Windows Components in Control Panel. You also have the option to add or upgrade server roles by using Manage Your Server. For more information about these options, see Installing Optional Components and Server Roles.

| **Note** |
| --- |
| Microsoft Management Console (MMC) 3.0 is installed automatically when you install Windows Server 2003 R2. |

# Windows Server 2003 R2 Volume License Product Key Encryption

Applies To: Windows Server 2003 R2

This information applies to you if you have a volume licensing agreement with Microsoft such as Microsoft Select, Microsoft Enterprise Agreement, and Microsoft Open License.

If you are upgrading to Windows Server 2003 R2 from a computer running Windows Server 2003 with Service Pack 1, you can encrypt your Windows Server 2003 R2 Volume License Product Key. This feature provides an additional layer of protection by encrypting your product key so it won't be visible as plain text. First, you encrypt your product key. Then, you can use the key to run Setup for Windows Server 2003 R2 in unattended mode.

| **Note** |
| --- |
| If you are performing a clean installation of Windows Server 2003 R2, you will need to use the encryption feature available on Windows Server 2003 R2 Disc 1. For more information, see "Windows XP Service Pack 1 Volume License Product Key Encryption" on the Microsoft Web site (http://go.microsoft.com/fwlink?linkId=55330). This document also applies to Windows Server 2003 R2 functionality. |

To use this feature when upgrading you must have the following:

- A computer running a volume licensed version of Windows Server 2003 with Service Pack 1
- Windows Server 2003 R2 Disc 2
- Windows Server 2003 R2 Volume License Product Key

**To encrypt a Volume License Product Key**

1. Insert Disc 2 into a computer running a volume licensed version of Windows Server 2003 with Service Pack 1.
2. At the command prompt, navigate to the drive that contains Disc 2 and run the following command. This will encrypt the product key and save it to the file you specify for *Filename*.

   **cmpnents\r2\Path\setup2.exe /encryptkey:** *xxxxx-xxxxx-xxxxx-xxxxx-xxxxx* **/encryptdays:** *NumberOfDays* **/encryptfile:** *Filename*

   This command uses the following arguments:

   - *xxxxx-xxxxx-xxxxx-xxxxx-xxxxx* - Your Volume License Product Key in plain text.
   - *NumberOfDays* - The number of days (5 to 60) for which the encrypted product key will be valid.
   - *Filename* - The location and name of the file where the encrypted product key will be saved, for example, C:\Myfolder\Myfile.txt. If the file already exists, the encrypted product key will be appended to the end of the file. You should ensure that you have read and write access to this location.

   You can now use the encrypted key to run Setup in unattended mode in the following ways: from a command prompt; in any valid Windows Setup answer file such as Unattend.txt, Sysprep.inf, or Ristndrd.sif; or in a batch file. For example, you can use the following syntax to install Disc 2 in unattended mode using the encrypted key:

   **** *Server* **** *share* **\cmpnents\r2\setup2.exe /q /a /p:** *EncryptedKey*

# ADFS Step-by-Step Guide

Updated: June 20, 2006

Applies To: Windows Server 2003 R2

The Step-by-Step Guide for Active Directory Federation Services provides instructions for setting up Active Directory Federation Services (ADFS) in a small test lab environment. This guide walks you through setup of a claims-aware application and a Windows NT token–based application on an ADFS-enabled Web server.

To download the latest version of the ADFS Step-by-Step guide, go to http://go.microsoft.com/fwlink/?linkid=49531.

# ADFS Deployment Guide

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You can use Active Directory Federation Services (ADFS) in the Microsoft® Windows Server® 2003 R2 operating system to build a federated identity management solution that can extend distributed identification, authentication, and authorization services to Web-based applications across organizational and platform boundaries. By deploying ADFS, you can extend your organization's existing identity management capabilities to the Internet. You can deploy ADFS to:

- Provide your employees or customers with a Web-based, single-sign-on (SSO) experience when they need remote access to internally hosted Web sites.
- Provide your employees or customers with a Web-based, SSO experience when they access cross-organizational Web sites from within the firewalls of your network.
- Provide your employees or customers with seamless access to Web-based resources in any federation partner organization on the Internet without requiring employees or customers to log on more than once.
- Retain complete control over your employee or customer identities without using other sign-on providers (Microsoft .NET Passport, Liberty Alliance, and others).

After you deploy ADFS, you can use it as your organization's optimal SSO solution.

**About this guide**
This guide is intended for use by system administrators and system engineers. It provides detailed guidance for deploying an ADFS design that has been preselected by you or an infrastructure specialist or system architect in your organization.

If a design has not yet been selected, we recommend that you wait to follow the instructions in this guide until after you have reviewed the various design options in the ADFS Design Guide and you have selected the most appropriate design for your organization. For more information about using this guide with a design that has already been selected, see Implementing Your ADFS Design Plan.

After you select your design and you use the worksheets in the design guide to gather the required information about claims, token types, account stores, and other items, you can then use this guide to deploy your ADFS design in your production environment. This guide provides steps for deploying any of the following primary ADFS designs:

- Web SSO
- Federated Web SSO
- Federated Web SSO with Forest Trust

Use the checklists in Implementing Your ADFS Design Plan to determine how best to use the instructions in this guide to deploy your particular design. For information about hardware and software requirements for deploying ADFS, see Appendix A: Reviewing ADFS Requirements in the ADFS Design Guide.

**What this guide does not provide**
This guide does not provide:

- Guidance regarding when and where to place federation servers, federation server proxies, or Web servers in your existing network infrastructure. For this information, see Planning Federation Server Placement, Planning Federation Server Proxy Placement and Planning ADFS-Enabled Web Server Placement in the ADFS Design Guide.
- Guidance for using certification authorities (CAs) to set up ADFS.
- Guidance for setting up or configuring specific Web-based applications.

- Setup instructions that are specific to setting up a test lab environment. For more information about how to configure an ADFS test lab environment, see the Step-by-Step Guide for Active Directory Federation Services (http://go.microsoft.com/fwlink/?linkid=49531).

- Information about how to customize federated logon screens, web.config files, or trust policy files.

- Information about how to modify or remove ADFS settings on specific servers or in the trust policy. For this information, see the ADFS Operations Guide (http://go.microsoft.com/fwlink/?linkid=78683).

**Overview of ADFS**

ADFS is an identity access solution that provides browser-based clients (internal or external to your network) with seamless SSO access to protected Internet-facing applications, even when the user accounts and applications are located in completely different networks or organizations.

When an application is in one network and a user account is in another network, typically the user is prompted for secondary credentials when he or she attempts to access the application. These secondary credentials represent the user's identity in the realm where the application resides. They are usually required by the Web server that hosts the application so that it can make the most appropriate authorization decision.

With ADFS, organizations can bypass requests for secondary credentials by providing trust relationships (federation trusts) that they can use to project a user's digital identity and access rights to trusted partners. In this federated environment, each organization continues to manage its own identities, but each organization can also securely project and accept identities from other organizations.

For more general information about ADFS, see Overview of Active Directory Federation Services (ADFS) in Windows Server 2003 R2 (http://go.microsoft.com/fwlink/?linkid=54650).

# Planning to Deploy ADFS

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

After you collect information about your environment and you decide on an Active Directory Federation Services (ADFS) design by following the guidance in the ADFS Design Guide, you can begin to plan the deployment of your organization's ADFS design. With the completed ADFS design and the information in this topic, you can determine which tasks to perform to deploy ADFS in your organization.

**Reviewing your ADFS Design**

If the design team that constructed the original ADFS design for your organization is different from the deployment team that will actually implement the deployment, make sure that the deployment team reviews the final design with the design team. Review the following points regarding the ADFS design:

- The design team's strategy to determine the best physical topology for the placement of ADFS servers in your corporate network or perimeter network. The deployment team can refer to documentation about this subject by reviewing the following topics in the ADFS Design Guide:

  - Planning Federation Server Placement
  - Planning Federation Server Proxy Placement
  - Planning ADFS-Enabled Web Server Placement

  It is possible that the design team might leave the subject of ADFS server placement for the deployment team. The deployment team is then responsible for documenting and implementing the physical topology of ADFS servers. The deployment team can review the preceding topics, Planning for ADFS Capacity and Appendix A: Reviewing ADFS Requirements to help determine the number of servers and the hardware requirements for the organization.

- The business reasons for your organization's designation as a resource partner, account partner, or both within the scope of the documented ADFS design. Ensure that members of the deployment team understand the reasons why ADFS is being deployed and what other companies or organizations are involved in the federation partnership. Ensure that members of the deployment team also understand the constraints that exist for the other companies or organizations (limited hardware, no extranet environment, and so forth) that may limit the scope of the design in some way. For more information about partner organizations, see Planning Partner Organization Deployments.

After the design teams and deployment teams agree on these issues, they can proceed with the deployment of the ADFS design. For more information, see Implementing Your ADFS Design Plan.

# Implementing Your ADFS Design Plan

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

The following environmental conditions and requirements are important factors in the implementation of your Active Directory Federation Services (ADFS) design plan:

- **Supported partners:** You usually use ADFS to work with partner organizations. To establish identity federation, determine the organizations with which you want to form a partnership. After a baseline ADFS deployment is in place, operating with partners involves adding partners, deleting partners, and updating partner information. Changes to partnerships may occur for a variety of reasons. For example, your ADFS deployment might require partnership updates if your partner changes its business significantly, your organization becomes part of a larger organization or a federation of organizations, or your organization is acquired by a different company. In any scenario in which you federate identities from multiple domains, you will need to know the domains (partners) that you are currently supporting and all the additional domains that represent potential partners.

- **Supported application types:** Some ADFS applications require access to operating system resources, while others are "claims aware." It is important to understand the type of applications that ADFS supports so that you can formulate administration requirements.

- **Logical and physical architectural diagrams or deployment topology:** You will need to know:

  - Whether ADFS servers will function in a set of farmed servers or on a single server.

  - Where your network deploys firewalls and proxies.

  - The location of resources and whether users are accessing resources from within your organization or from outside the organization or both.

**How to implement your ADFS design using this guide**

The next step in implementing your design is to determine in what order each of the deployment tasks must be performed. This guide uses checklists to help you walk through the various server and application deployment tasks that are required to implement your design plan. As shown in the following illustration, parent and child checklists are used as necessary to represent the order in which tasks for a specific ADFS design must be processed:




Use the following parent checklists in this section of the guide to become familiar with the deployment tasks for implementing your organization's ADFS design:

- Checklist: Implementing a Web SSO Design
- Checklist: Implementing a Federated Web SSO Design
- Checklist: Implementing a Federated Web SSO with Forest Trust Design

# Checklist: Implementing a Web SSO Design

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This parent checklist includes cross-reference links to important concepts about the Web Single-Sign-On (SSO) design. It also contains links to subordinate checklists that will help you complete the tasks that are required to implement this design.

| **Note** |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to subordinate checklist, return to this topic after you review the conceptual topic or complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist. |

### Checklist: Implementing a Web SSO Design

| | Task | Reference |
| --- | --- | --- |
| ☐ | Review important concepts and examples for the Web SSO design and determine which Active Directory Federation Services (ADFS) deployment goals you can use to customize this design to meet the needs of your organization.<br>**Note**<br>The Web SSO design does not require that you set up partners in the Federation Service. | Web SSO design<br>Web SSO example<br>Identifying Your ADFS Deployment Goals |
| ☐ | Review the hardware, software, certificate, Domain Name System (DNS), account store, and client requirements for deploying ADFS in your organization. | Appendix A: Reviewing ADFS Requirements |
| ☐ | According to your design plan, install one or more federation servers in the corporate network or in the perimeter network.<br>**Note**<br>The Web SSO design requires only one federation server to function successfully. A single federation server acts in both the account partner role and the resource partner role. | Checklist: Installing a federation server |
| ☐ | (Optional) Determine whether or not your organization needs a federation server proxy in the perimeter network. | Checklist: Installing a federation server proxy |
| ☐ | Install one or more ADFS-enabled Web servers to host your preferred federated application using the appropriate ADFS Web Agent.<br>**Note**<br>The Web SSO design requires at least one ADFS-enabled Web server. | Checklist: Installing an ADFS-enabled Web server |
| ☐ | Depending on your Web SSO design plan and how you intend to use it, install either an Active Directory account store or an Active Directory Application Mode (ADAM) account store in the Federation Service. | Checklist: Configuring the account partner organization |
| ☐ | For SSO, incoming claims must be associated with organization group claims or custom claims that reside in the Federation Service. | Create an organization group or custom claim |

| | | |
|---|---|---|
| ☐ | If you are an administrator in the resource partner organization, install either a claims-aware application or a Windows NT token-based application, or both, using the appropriate checklist.<br><br>**Note**<br>The account partner administrator does not have to complete either of the following checklists. | Checklist: Installing a claims-aware application<br><br>Checklist: Installing a Windows NT token-based application |

# Checklist: Implementing a Federated Web SSO Design

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This parent checklist includes cross-reference links to important concepts about the Federated Web Single-Sign-On (SSO) design. It also contains links to subordinate checklists that will help you complete the tasks that are required to implement this design.

| Note |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist. |

### Checklist: Implementing a Federated Web SSO Design

| | Task | Reference |
| --- | --- | --- |
| ☐ | Review important concepts and examples for the Federated Web SSO design and determine which Active Directory Federation Services (ADFS) deployment goals you can use to customize this design to meet the needs of your organization. | Federated Web SSO design<br>Federated Web SSO example<br>Identifying Your ADFS Deployment Goals |
| ☐ | Review the hardware, software, certificate, Domain Name System (DNS), account store, and client requirements for deploying ADFS in both partner organizations. | Appendix A: Reviewing ADFS Requirements |
| ☐ | According to your design plan, install one or more federation servers in each partner organization.<br>**Note**<br>For the Federated Web SSO design, you need at least one federation server in the account partner organization and at least one federation server in the resource partner organization. | Checklist: Installing a federation server |
| ☐ | (Optional) Determine whether or not your organization needs a federation server proxy. If your design plan calls for a proxy, you can install one or more federation server proxies in each partner organization. | Checklist: Installing a federation server proxy |
| ☐ | If you are an administrator in the resource partner organization, install one or more ADFS-enabled Web servers to host your preferred federated application using the appropriate ADFS Web Agent.<br>**Note**<br>The account partner administrator does not have to complete the following checklist. | Checklist: Installing an ADFS-enabled Web server |
| ☐ | According to your design plan, share certificates, configure clients, and configure the Federation Service in both partner organizations so that they can communicate over a federation trust. | Checklist: Configuring the account partner organization<br>Checklist: Configuring the resource partner organization |

| | | |
|---|---|---|
| ☐ | If you are an administrator in the resource partner organization, install either a claims-aware application or a Windows NT token-based application, or both, using the appropriate checklist.<br><br>**Note**<br>The account partner administrator does not have to complete either of the following checklists. | Checklist: Installing a claims-aware application<br><br>Checklist: Installing a Windows NT token-based application |

# Checklist: Implementing a Federated Web SSO with Forest Trust Design

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This parent checklist includes cross-reference links to important concepts about the Federated Web Single-Sign-On (SSO) with Forest Trust design. It also contains links to subordinate checklists that will help you complete the tasks that are required to implement this design.

| **Note** |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist. |

### Checklist: Implementing a Federated Web SSO with Forest Trust Design

| | Task | Reference |
| --- | --- | --- |
| ☐ | Review important concepts and examples for the Federated Web SSO with Forest Trust design and determine which Active Directory Federation Services (ADFS) deployment goals you can use to customize this design to meet the needs of your organization. | Federated Web SSO with Forest Trust design<br>Federated Web SSO with Forest Trust example<br>Identifying Your ADFS Deployment Goals |
| ☐ | Review the requirements for enabling the Federated Web SSO with Forest Trust design in both the account partner organization and the resource partner organization. | When to enable Windows trusts |
| ☐ | Review the hardware, software, certificate, Domain Name System (DNS), account store, and client requirements for deploying ADFS in both partner organizations. | Appendix A: Reviewing ADFS Requirements |
| ☐ | According to your design plan, install one or more federation servers in each partner organization.<br>**Note**<br>The Federated Web SSO with Forest Trust design requires at least one federation server in the account partner and at least one federation server in the resource partner. | Checklist: Installing a federation server |
| ☐ | (Optional) Determine whether or not your organization needs a federation server proxy. If your design plan calls for a federation server proxy, you can install one or more federation server proxies in each partner organization. | Checklist: Installing a federation server proxy |
| ☐ | If you are an administrator in the resource partner organization, install one or more ADFS-enabled Web servers to host your preferred federated application using the appropriate ADFS Web Agent.<br>**Note**<br>The account partner administrator does not have to complete the following checklist. | Checklist: Installing an ADFS-enabled Web server |

| | | |
|---|---|---|
| ☐ | According to your design plan, share certificates, configure clients, and configure the Federation Service in both partner organizations so that they can communicate over a federation trust. | Checklist: Configuring the account partner organization<br>Checklist: Configuring the resource partner organization |
| ☐ | To enable the Federated Web SSO with Forest Trust scenario, complete the appropriate procedure in the account Federation Service and in the resource Federation Service.<br>**Note**<br>You must enable the Windows trust option in both the account Federation Service and in the resource Federation Service. If only one partner has this option enabled, this design will not function. | Configure an account partner to use Windows trust<br>Configure a resource partner to use Windows trust |
| ☐ | If you are an administrator in the resource partner organization, install either a claims-aware application or a Windows NT token-based application, or both, using the appropriate checklist.<br>**Note**<br>The account partner administrator does not have to complete either of the following checklists. | Checklist: Installing a claims-aware application<br>Checklist: Installing a Windows NT token-based application |

# Deploying Partner Organizations

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

To deploy a new partner organization, complete the tasks in either Checklist: Configuring the resource partner organization or Checklist: Configuring the account partner organization, depending on your Active Directory Federation Services (ADFS) design.

| **Note** |
| --- |
| When you use either of these checklists, we strongly recommend that you first read the references to account partner or resource partner planning guidance (in the ADFS Design Guide) before continuing to the procedures for setting up the new partner organization. Following the checklist in this way will help provide a better understanding of the full ADFS design and deployment story for the account partner or resource partner organization. |

**About account partners**

An account partner is the organization in the federation trust relationship that physically stores user accounts in either an Active Directory store or an Active Directory Application Mode (ADAM) store. The account partner is responsible for collecting and authenticating a user's credentials, building up claims for that user, and packaging the claims into security tokens. These tokens can then be presented across a federation trust to enable access to Web-based resources that are located in the resource partner organization.

In other words, an account partner represents the organization for whose users the account-side Federation Service issues security tokens. The Federation Service in the account partner organization authenticates local users and creates security tokens that are used by the resource partner in making authorization decisions.

With regard to Active Directory, the account partner in ADFS is conceptually equivalent to a single Active Directory forest whose accounts need access to resources that are physically located in another forest. Accounts in this situation forest can access resources in the resource forest only when an external trust or forest trust relationship exists between the two forests and the resources to which the users are trying to gain access have been set with the proper authorization permissions.

**About resource partners**

The resource partner is the organization in an ADFS deployment where ADFS-enabled Web servers are located. The resource partner trusts the account partner to authenticate users. Therefore, to make authorization decisions, the resource partner consumes the claims that are packaged in security tokens coming from users in the account partner.

In other words, a resource partner represents the organization whose Web servers are protected by the resource-side Federation Service. The Federation Service at the resource partner uses the security tokens that are produced by the account partner to make authorization decisions for Web servers in the resource partner.

To function as an ADFS resource, Web servers in the resource partner organization must have the ADFS Web Agent component of ADFS installed. Web servers that function as an ADFS resource can host either claims-aware applications or Windows NT token–based applications. For more information about the two types of applications, see Controlling Access to Web-based Applications (http://go.microsoft.com/fwlink/?LinkId=64040).

# Checklist: Configuring the account partner organization

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This checklist includes tasks for deploying Active Directory Federation Services (ADFS) in the account partner organization. It also includes tasks for configuring the components that are required to establish one-half of a federation partnership.

If you are deploying a Web SSO design, you do not have to follow this checklist. However, you do have to complete the tasks in this checklist to successfully deploy a Federated Web SSO design or Federated Web SSO with Forest Trust design.

| **Important** |
| --- |
| Make sure that the administrator in the resource partner organization follows the guidance in Checklist: Configuring the resource partner organization to ensure that all necessary deployment tasks will be completed to successfully create the second half of the federation partnership. |

| **Note** |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist. |

### Checklist: Configuring the account partner organization

| | Task | Reference |
| --- | --- | --- |
| ☐ | Based on your deployment goals, review information about the components that are required to provide users with access to the federated applications that you host in your organization. | Provide federated access for your employees on the corporate network<br>Provide federated access for your remote employees on the Internet<br>Provide single-sign-on access for customers to your hosted applications |
| ☐ | Determine which ADFS design this account partner organization will be associated with. | Web SSO design<br>Federated Web SSO design<br>Federated Web SSO with Forest Trust design |
| ☐ | To effectively plan and implement the physical topology for the account partner deployment, determine whether your ADFS design requires one or more federation servers or federation server proxies. | Checklist: Installing a federation server<br>Checklist: Installing a federation server proxy |
| ☐ | After you deploy the first federation server in the account partner organization, configure the trust policy. You can do this manually or through a policy file that is provided to you by the administrator of the resource partner organization. | Add a new resource partner by manually configuring the trust policy<br>Add a new resource partner by importing an existing policy file |
| ☐ | Determine whether you want to implement privacy settings on resource partners in the Federation Service. When you enable enhanced identity privacy settings you remove sensitive data from claims that are sent to the resource partner organization. | Review how ADFS may affect privacy<br>Enable enhanced identity privacy |

| | | |
|---|---|---|
| ☐ | If you are deploying the Federation Web Single-Sign-On (SSO) with Forest Trust design, configure the Federation Service for Windows trust. | When to enable Windows trusts<br>Configure a resource partner to use Windows trust |
| ☐ | Incoming claims must be associated with organization group claims or custom claims in the Federation Service. | Create an organization group or custom claim |
| ☐ | Create outgoing claim mappings for each resource partner in the Federation Service so that organizational claims will be transmitted to resource partners. | Create an outgoing group or custom claim mapping |
| ☐ | Determine the type of account store that you need to add to the Federation Service. Then, create the appropriate group or custom claim extraction to map an organization claim to an Active Directory group or user account. | Add an Active Directory account store<br>Add an ADAM account store<br>Map an organization group claim to an Active Directory group (group claim extraction)<br>Map an organization group claim to an ADAM attribute and value (group claim extraction)<br>Map an organization custom claim to an Active Directory or ADAM user attribute (custom claim extraction) |
| ☐ | Prepare client computers for federation by:<br>• Adding the Uniform Resource Locator (URL) for the account federation server to the trusted sites list for the client browser.<br>• Using Group Policy to push the appropriate Secure Sockets Layer (SSL) certificates to client computers. | Prepare client computers for federation<br>Configure client computers to trust the account federation server<br>Distribute certificates to client computers using Group Policy |

# Checklist: Configuring the resource partner organization

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This checklist includes the tasks that are necessary for deploying Active Directory Federation Services (ADFS) in the resource partner organization. It also includes tasks for configuring the components that are required to establish one-half of a federation partnership.

If you are deploying a Web SSO design, you do not have to follow this checklist. However, you do have to complete the tasks in this checklist to successfully deploy a Federated Web SSO design or Federated Web SSO with Forest Trust design.

| **Important** |
| --- |
| Make sure that the administrator of the account partner organization follows the guidance in Checklist: Configuring the account partner organization to ensure that all necessary deployment tasks will be completed to successfully create the second half of the federation partnership |

| **Note** |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist. |

### Checklist: Configuring the resource partner organization

| | Task | Reference |
| --- | --- | --- |
| ☐ | Based on your deployment goals, review information about the components that are required to provide users with access to the federated applications that you host in your organization. | Provide federated access for your hosted applications<br>Provide single-sign-on access for customers to your hosted applications |
| ☐ | Determine which ADFS design this resource partner organization will be associated with. | Web SSO design<br>Federated Web SSO design<br>Federated Web SSO with Forest Trust design |
| ☐ | Review the different application types, and decide which application to deploy. | Identify the type of federated application to deploy<br>Checklist: Installing a claims-aware application<br>Checklist: Installing a Windows NT token-based application |
| ☐ | To effectively plan and implement the physical topology for the resource partner deployment, you determine whether your ADFS design requires one or more federation servers, federation server proxies, or ADFS-enabled Web servers. | Checklist: Installing a federation server<br>Checklist: Installing a federation server proxy<br>Checklist: Installing an ADFS-enabled Web server |

| | | |
|---|---|---|
| ☐ | After the first federation server is deployed in the resource partner organization, you configure the trust policy. You can do this manually or through a policy file that is provided to you by the administrator of the account partner organization. | Add a new account partner by manually configuring the trust policy<br><br>Add a new account partner by importing an existing policy file |
| ☐ | Incoming claims must be associated with organization group claims or custom claims in the Federation Service. | Create an organization group or custom claim |
| ☐ | The resource partner organization will have to create incoming claim mappings for each account partner in the Federation Service so that incoming claims will be mapped to organizational claims in the resource partner. | Create an incoming group claim mapping<br><br>Create an incoming custom claim mapping |
| ☐ | Determine if you need an account store. If you need an account store in the resource partner organization, create the appropriate group claim mappings to map an organization claim to an Active Directory group. | Add an Active Directory account store<br><br>Add an ADAM account store<br><br>Map an organization group claim to an Active Directory group (group claim extraction)<br><br>Map an organization group claim to an ADAM attribute and value (group claim extraction)<br><br>Map an organization custom claim to an Active Directory or ADAM user attribute (custom claim extraction) |
| ☐ | If you are deploying the Federation Web SSO with Forest Trust design, have the administrator for the account partner organization configure the Federation Service for Windows trust. | When to enable Windows trusts<br><br>Configure an account partner to use Windows trust |

# Checklist: Implementing a resource account mapping method

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This checklist includes the tasks for determining which resource account mapping method to implement. Complete the tasks in this checklist to successfully deploy a Windows NT token–based application. You do not have to complete this checklist if you are deploying a claims-aware application.

| **Note** |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist. |

### Checklist: Implementing a resource account mapping method

| | Task | Reference |
| --- | --- | --- |
| ☐ | Review information in the ADFS Design Guide to determine whether to use one or more of the resource account, resource group, or group-to-UPN (user principal name) mapping methods. | Determine your resource account mapping method |
| ☐ | Review information in the ADFS Design Guide about when and how to use resource accounts in the resource forest. | When to use resource accounts |
| ☐ | Review information in the ADFS Design Guide regarding which setting in the resource Federation Service can best optimize resource account mapping methods. | Select the optimal resource account option |
| ☐ | Review information in the ADFS Design Guide about when and how to use resource groups in the resource forest. | When to use resource groups |
| ☐ | Review information in the ADFS Design Guide about whether you should use the group-to-UPN mapping method in the resource forest instead of the resource group mapping method. | When to use group-to-UPN mapping |
| ☐ | If you will be using resource accounts for mapping, create and configure the accounts for use with Windows NT token–based applications.<br>**Note**<br>This procedure is necessary only when you use the resource account mapping method. | Create a resource account in the resource partner forest |
| ☐ | If you will be using resource groups for mapping, associate an existing security group in the resource forest with a specific organization claim in the resource Federation Service.<br>**Note**<br>This procedure is necessary only when you use the resource group mapping method. | Map an organization group claim to a resource group |

# Add a new account partner by manually configuring the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

To add a new account partner in Active Directory Federation Services (ADFS) and manually configure the trust policy, perform the following procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new account partner by manually configuring the trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. Right-click **Account Partners**, point to **New**, and then click **Account Partner**.
4. On the **Welcome to the Add Account Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, click **No**, and then click **Next**.
6. On the **Account Partner Details** page, do the following, and then click **Next**
   - In **Display name**, type the display name of the account partner.
   - In **Federation Service URI**, type the Uniform Resource Identifier (URI) of the account partner Federation Service.
   - In **Federation Service endpoint URL**, type the endpoint Uniform Resource Locator (URL) as follows:

     **With a Federation Service Proxy:** If the Federation Service Proxy is installed on servers in a perimeter network of the account domain, use the fully qualified domain name (FQDN) of any federation server proxy.

     **Without a Federation Service Proxy:** If your deployment does not use federation server proxies, use the FQDN of the account partner Federation Service.
7. On the **Account Partner Verification Certificate** page, type the path to the verification certificate, or browse to it, and then click **Next**.
8. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**, and then go to step 10.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
9. On the **Federated Web SSO with Forest Trust** page, do one of the following, and then click **Next**:
   - To accept users in all domains that are trusted by the account partner, click **All Active Directory domains and forests**. Any user that can authenticate to the account partner will be accepted.
   - To accept user accounts that are located in some of the domains that are trusted by the account partner, click **The following Active Directory domains and forests**. Then, in **New, trusted Active Directory domain or forest**, type the name of a domain or forest, and then click **Add**. Only users from the specified domains will be accepted.
10. On the **Account Partner Identity Claims** page, select one or more identity claims to share with the resource partner,

and then click **Next**:

- If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

| **Important** |
| --- |
| When UPN claims or e-mail claims are used to make authorization decisions, it is essential that each account partner use a unique UPN suffix or e-mail suffix. If two account partners use the same UPN suffix or e-mail suffix, it might not be possible to identify users uniquely. This condition might allow a user from one account partner to receive the permissions that are intended for a user in another account partner. This condition might also introduce a significant security weakness because an administrator can intentionally create user accounts to impersonate users from one of your other account partners. |

| **Note** |
| --- |
| If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable because UPN claims are required for this scenario. |

- If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.

- If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.

11. If you selected **UPN Claim** as an identity claim, on the **Accepted UPN Suffixes** page, type the accepted suffix, click **Add**, and then click **Next**. You may click to accept **All UPN suffixes** if you selected the **Federation Web SSO with Forest Trust** option.
12. If you selected **E-mail Claim** as an identity claim, on the **Accepted E-mail Suffixes** page, type the accepted suffix, click **Add**, and then click **Next**. You may click to accept **All E-mail suffixes** if you selected the **Federation Web SSO with Forest Trust** option.

| **Note** |
| --- |
| Common name claims require no additional information. |

13. On the **Enable this Account Partner** page, if you do not want to enable the account partner now, clear the **Enable this account partner** check box, and then click **Next**.
14. To add the new account partner and close the wizard, click **Finish**.

**See Also**
**Concepts**
Change the Federation Service endpoint URL
View the current verification certificate

# Add a new account partner by importing an existing policy file

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have received an exported Active Directory Federation Services (ADFS) trust policy file from the account partner organization, perform the following procedure on a resource federation server to automatically configure the new account partner by importing the policy file.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new account partner by importing an existing policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, double-click **Trust Policy**, and then double-click **Partner Organizations**.
3. Right-click **Account Partners**, point to **New**, and then click **Account Partner**.
4. On the **Welcome to the Add Account Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, click **Yes**. In **Partner interoperability policy file**, type or browse to the policy file that you want to import, and then click **Open**.
6. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**, and then go to step 8.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
7. On the **Federated Web SSO with Forest Trust** page, do one of the following, and then click **Next**:
   - To accept users in all domains that are trusted by the account partner, click **All Active Directory domains and forests**. Any user that can authenticate to the account partner will be accepted.
   - To accept user accounts that are located in some of the domains that are trusted by the account partner, click **The following Active Directory domains and forests**. Then, in **New, trusted Active Directory domain or forest**, type the name of a domain or forest, and then click **Add**. Only users from the specified domains will be accepted. If you need to remove a domain or forest, click **Remove**.
8. On the **Account Partner Identity Claims** page, select one or more identity claims to share with the resource partner, and then click **Next**:
   - If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

| **Important** |
| --- |
| When UPN claims or e-mail claims are used to make authorization decisions, it is essential that each account partner use a unique UPN suffix or e-mail suffix. If two account partners use the same UPN suffix or e-mail suffix, it might not be possible to uniquely identify users. This condition might allow a user from one account partner to receive the permissions that are intended for a user in another account partner. This condition might also introduce a significant security weakness because an administrator can intentionally create user accounts to impersonate users from one of your other account partners. |

| **Note** |
| --- |

| |
|---|
| If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable because UPN claims are required for this scenario. |

- If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.
- If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.

9. If you selected **UPN Claim** as an identity claim, on the **Accepted UPN Suffixes** page, type the accepted suffix, click **Add**, and then click **Next**. You may click to accept **All UPN suffixes** if you selected the **Federation Web SSO with Forest Trust** option.
10. If you selected **E-mail Claim** as an identity claim, on the **Accepted E-mail Suffixes** page, type the accepted suffix, click **Add**, and then click **Next**. You may click to accept **All E-mail suffixes** if you selected the **Federation Web SSO with Forest Trust** option.

| **Note** |
|---|
| Common name claims require no additional information. |

11. On the **Enable this Account Partner** page, if you do not want to enable the account partner now, clear the **Enable this account partner** check box, and then click **Next**.
12. To add the new account partner and close the wizard, click **Finish**.

**See Also**
**Concepts**
Export an account or resource policy file to a partner organization
Add a new account partner by manually configuring the trust policy

# Add a new resource partner by manually configuring the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

To add a new resource partner in Active Directory Federation Services (ADFS) and manually configure the trust policy, perform the following procedure on an account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new resource partner by manually configuring the trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. In the console tree, right-click **Resource Partners**, point to **New**, and then click **Resource Partner**.
4. On the **Welcome to the Add Resource Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, ensure that **No** is selected, and then click **Next**.
6. On the **Resource Partner Details** page, do the following, and then click **Next**:
   - In **Display name**, type the display name of the resource partner.
   - In **Federation Service URI**, type the Uniform Resource Identifier (URI) for the resource partner Federation Service.
   - In **Federation Service endpoint URL**, type the endpoint Uniform Resource Locator (URL) of the resource partner Federation Service.
7. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
8. On the **Resource Partner Identity Claims** page, select one or more identity claims to share with the resource partner, and then click **Next**:
   - If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

     | **Note** |
     | --- |
     | If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable because UPN claims are required for this scenario. |

   - If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.
   - If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.
9. If you selected **UPN Claim** as an identity claim, on the **Select UPN Suffix** page, do one of the following, and then click

**Next**:

- To pass all UPN suffixes through without replacing them, click **Pass all UPN suffixes through unchanged**.
- To replace all UPN suffixes with a different suffix, click **Replace all UPN suffixes with the following**, and then type the suffix that you want to use to replace all UPN suffixes.

10. If you selected **E-mail Claim** as an identity claim, on the **Select E-mail Suffix** page, do one of the following, and then click **Next**:

    - To pass all e-mail suffixes without replacing them, click **Pass all E-mail suffixes through unchanged**.
    - To replace all e-mail suffixes with a different suffix, click **Replace all E-mail suffixes with**, and then type the suffix that you want to use to replace all e-mail suffixes.

| **Note** |
| --- |
| Common name claims require no additional information. |

11. On the **Enable this Resource Partner** page, if you do not want to enable the resource partner now, clear the **Enable this resource partner** check box, and then click **Next**.
12. To add the new resource partner and close the wizard, click **Finish**.

**See Also**
**Concepts**
Add a new resource partner by importing an existing policy file

# Add a new resource partner by importing an existing policy file

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have received a policy file from your resource partner organization that you can import, perform the following procedure on the account federation server on which you are creating the new resource partner.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new resource partner by importing an existing policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. Right-click **Resource Partners**, point to **New**, and then click **Resource Partner**.
4. On the **Welcome to the Add Resource Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, do the following, and then click **Next**:
   - Click **Yes**.
   - In **Partner interoperability policy file**, browse to or type the location of the resource partner policy file.
6. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
7. On the **Resource Partner Identity Claims** page, select one or more identity claims that the account partner will provide to the resource partner, and then click **Next**:
   - If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

     | **Note** |
     | --- |
     | If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable because UPN claims are required for this scenario. |

   - If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.
   - If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.
8. If you selected **UPN Claim** as an identity claim, on the **Select UPN Suffix** page, select one of the following, and then click **Next**.
   - To pass all UPN suffixes through without replacing them, click **Pass all UPN suffixes through unchanged**.
   - To replace all UPN suffixes with a different suffix, click **Replace all UPN domain suffixes with the following**, type the suffix that you want to use to replace all UPN suffixes, and then click **Add**.

9. If you selected **E-mail Claim** as an identity claim, on the **Select E-mail Suffix** page, do one of the following, and then click **Next**:

   - To pass all e-mail suffixes without replacing them, click **Pass all e-mail suffixes through unchanged**.

   - To replace all UPN suffixes with a different suffix, click **Replace all E-mail suffixes with**, and then type the suffix that you want to use to replace all e-mail suffixes.

| **Note** |
| --- |
| Common name claims require no additional information. |

10. On the **Enable this Resource Partner** page, if you do not want to enable the resource partner now, clear the **Enable this resource partner** check box, and then click **Next**.

11. To add the new resource partner and close the wizard, click **Finish**.

**See Also**
**Concepts**
Export an account or resource policy file to a partner organization
Add a new resource partner by manually configuring the trust policy

# Create an organization group or custom claim

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), an organization group claim is used by the Federation Service to represent a user's membership in a group or role. An organization custom claim is used by the Federation Service to provide custom information, such as an employee identification number, about a user. Group claims and custom claims are created in the same way. You can use the following procedure to create an organization group claim or a custom claim.

Administrators in the account Federation Service create organization group claims and custom claims to represent account users. Administrators in the resource Federation Service create corresponding organization group claims and custom claims to represent groups and users that can be recognized as resource users in the resource Federation Service. Because outgoing claims in the account Federation Service map to incoming claims in the resource Federation Service, the resource Federation Service is able to accept the credentials that are provided by the account Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create an organization group or custom claim**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Organization Claims**, point to **New**, and then click **Organization Claim**.
3. In the **Create a New Organization Claim** dialog box, in **Claim name**, type a unique name that identifies the group or role for a group claim, or type the name of a custom claim.

   | **Note** |
   | --- |
   | You can create a group claim and a custom claim with the same name. |

4. If you want to limit the auditing of the claim to the claim name and omit the value, select the **Limit the auditing of this claim** check box.
5. For **Specify the claim type**, click **Group claim** or **Custom claim**, and then click **OK**.

**See Also**

**Concepts**

Create an incoming group claim mapping
Create an incoming custom claim mapping
Create an outgoing group or custom claim mapping

# Create an incoming group claim mapping

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), incoming group claim mappings are used in the resource Federation Service to transform group claims that are sent by an account partner into organization claims that can be used by the resource partner to make authorization decisions.

For example, an account partner might send a security token for a user that contains the group claim SalesReps. Because the resource partner cannot make authorization decisions based on the account user's membership in the SalesReps group, an incoming group claim mapping is used to map the incoming group claim that is named SalesReps in the account Federation Service to the organization group claim that is named Purchasers in the resource Federation Service. The resource itself provides access to the local security group to which the Purchasers claim is mapped.

Perform this procedure on a resource federation server. To perform this procedure, you must have created an organization group claim to which you can map the incoming claim.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create an incoming group claim mapping**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Account Partners**, right-click your account partner, point to **New**, and then click **Incoming Group Claim Mapping**.
3. In the **Create a New Incoming Group Claim Mapping** dialog box, in **Incoming group claim name**, type the name of the group claim that your account partner sends to you.
4. In **Organization group claim**, select the group claim that you will use in your organization to map the incoming group claim to, and then click **OK**.

**See Also**
**Concepts**
Create an organization group or custom claim
Change the organization claim mapping of an incoming group or custom claim

# Create an incoming custom claim mapping

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), an organization custom claim maps to a user attribute. Incoming custom claim mappings are used in the resource Federation Service to map custom claims that are sent by an account partner to claims that can be used by the resource partner to make authorization decisions.

For example, an account partner might send a security token for a user that contains the custom claim EmployeeID, which maps to a user attribute in the account partner directory database. Because the resource partner cannot make authorization decisions based on the account user's EmployeeID attribute value, an incoming custom claim mapping is used to map the custom organization claim EmployeeID, which is recognized in the account partner, to the organization claim CustomerID, which is recognized in the resource partner.

Perform this procedure on a federation server in the resource Federation Service. To complete this procedure, you must have created an organization custom claim to which you can map the incoming claim.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create an incoming custom claim mapping**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Account Partners**, right-click your account partner, click **New**, and then click **Incoming Custom Claim Mapping**.
3. In the **Create a New Incoming Custom Claim Mapping** dialog box, in **Incoming custom claim name**, type the name of the custom claim that your account partner sends to you.
4. In **Organization custom claim**, select the custom claim that you will use in your organization to map the incoming custom claim to, and then click **OK**.

**See Also**
**Concepts**
Create an organization group or custom claim
Change the organization claim mapping of an incoming group or custom claim

# Create an outgoing group or custom claim mapping

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), an organization claim (group or custom) in the account Federation Service must be mapped to an outgoing claim, which the resource Federation Service will receive when an account organization user is requesting access to a resource. On the resource federation server, this claim is received as an incoming claim, which is likewise configured to map to a local organization claim, which the resource Federation Service uses to make authorization decisions.

Perform this procedure on an account federation server. To perform this procedure, you must have created an organization group or custom claim to which you can map the outgoing claim.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create an outgoing group or custom claim mapping**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Resource Partners**, right-click the resource partner, point to **New**, and then click one of the following:

   If you are mapping a group claim, click **Outgoing Group Claim Mapping**, and then create the mapping as follows:

   a. In the **Create a New Outgoing Group Claim Mapping** dialog box, in **Organization group claims**, select the group claim in the account organization that you want to map to the outgoing claim.

   b. In **Outgoing group claim name**, type the name of the outgoing group claim that you want to send to the resource partner, and then click **OK**.

   If you are mapping a custom claim, click **Outgoing Custom Claim Mapping**, and then create the mapping as follows:

   a. In the **Create a New Outgoing Custom Claim Mapping** dialog box, in **Organization custom claims**, select the custom claim in the account organization that you want to map to the outgoing claim.

   b. In **Outgoing custom claim name**, type the name of the outgoing custom claim that you want to send to the resource partner, and then click **OK**.

**See Also**
**Concepts**
Create an organization group or custom claim
Change the organization claim mapping of an outgoing group or custom claim

# Add an Active Directory account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If user and computer accounts that require access to a resource that is protected by Active Directory Federation Services (ADFS) are stored in Active Directory, you must add Active Directory as an account store on a federation server in the Federation Service that authenticates the accounts.

An Active Directory forest can have only one Active Directory instance. Therefore, you can add only one Active Directory store for the respective Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add an Active Directory account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Account Stores**, point to **New**, and then click **Account Store**.
3. On the **Welcome to the Add Account Store Wizard** page, click **Next**.
4. On the **Account Store Type** page, ensure that **Active Directory** is selected, and then click **Next**.
5. On the **Enable this Account Store** page, ensure that the **Enable this account store** check box is selected, and then click **Next**.
6. On the **Completing the Add Account Store Wizard** page, click **Finish**.

**See Also**
**Concepts**
Add an ADAM account store

# Add an ADAM account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you use multiple Active Directory Application Mode (ADAM) stores for user accounts that require access to one or more Web applications that are protected by Active Directory Federation Services (ADFS), you can add the ADAM account stores to the Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Account Stores**, point to **New**, and then click **Account Store**.
3. On the **Welcome to the Add Account Store Wizard** page, click **Next**.
4. On the **Account Store Type** page, ensure that **Active Directory Application Mode (ADAM)** is selected, and then click **Next**.
5. On the **ADAM Store Details** page, in **Account store display name**, type the name of the ADAM account store as you want it to be displayed in the Active Directory Federation Services snap-in user interface (UI).
6. In **Account store URI**, type the Uniform Resource Identifier (URI) for the ADAM account store, and then click **Next**.

| **Note** |
| --- |
| The account store URI uniquely identifies the ADAM instance among multiple ADAM account stores. |

7. On the **ADAM Server Settings** page, do the following, and then click **Next**:
   a. In **ADAM server name or IP address**, type the name or IP address of the ADAM server.
   b. In **Port number**, type the TCP/IP port number for the account service. Accept the default of 389 unless Active Directory is installed on the same server, in which case you must use a different port.
   c. In **LDAP search base distinguished name**, type the distinguished name of the ADAM instance.
   d. In **User name LDAP attribute**, type the name of the user name attribute that users provide during logon (for example, **userPrincipalName** or **sAMAccountName**).
8. On the **Identity Claims** page, select one or more identity claims that will be provided by the account store, and then click **Next**:
   a. If the account store provides user principal name (UPN) identity claims, select the **User Principal Name (UPN)** check box, and then type the Lightweight Directory Access Protocol (LDAP) attribute name to which UPN identity claims map (the attribute whose value is the user's UPN, usually **userPrincipalName**).
   b. If the account store provides e-mail identity claims, select the **E-mail** check box, and then type the LDAP attribute name to which e-mail identity claims map (the attribute whose value is the user's a-mail name, usually **userPrincipalName**).
   c. If the account store provides a common name identity claim, select the **Common Name** check box, and then type the LDAP attribute name to which the common name identity claim maps (the attribute whose value is the user's common name, usually **displayName**).

9. On the **Enable this Account Store** page, ensure that the **Enable this account store** check box is selected, and then click **Next**.
10. On the **Completing the Add Account Store Wizard** page, click **Finish**.

**See Also**
**Concepts**
Add an Active Directory account store

# Map an organization group claim to an Active Directory group (group claim extraction)

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you use Active Directory as the Active Directory Federation Services (ADFS) account store for an account Federation Service, you map an organization group claim to a security group in Active Directory. This mapping is called a group claim extraction.

Perform this procedure in the account Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To map an organization group claim to an Active Directory group**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Account Stores**, right-click **Active Directory**, point to **New**, and then click **Group Claim Extraction**.
3. In the **Create a New Group Claim Extraction** dialog box, click **Add**, and then select the Active Directory security group that you want to map to a group claim.
4. In **Map to this Organization Claim**, select the group claim to map to the Active Directory security group, and then click **OK**.

**See Also**
**Concepts**
Map an organization group claim to an ADAM attribute and value (group claim extraction)
Map an organization custom claim to an Active Directory or ADAM user attribute (custom claim extraction)

# Map an organization custom claim to an Active Directory or ADAM user attribute (custom claim extraction)

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Whether you use Active Directory or Active Directory Application Mode (ADAM) as the Active Directory Federation Services (ADFS) account store for an account Federation Service, an organization custom claim maps to an administratively assigned Lightweight Directory Access Protocol (LDAP) attribute for the user that the claim identifies. This mapping is called a custom claim extraction.

For example, if the user is to be identified by position, you might create the organization custom claim Position and use the Title attribute to identify the user's position. If the Title attribute is present in the Active Directory or ADAM store, the corresponding organization custom claim is generated with the value of the Title attribute. Suppose that the Title attribute of the user account has the value "Software Engineer." In this case, the organization custom claim Position is generated for this user with the value "Software Engineer." If the Title attribute is not found for the user account, the Position claim is not generated for the user.

Perform this procedure in the account Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To map an organization custom claim to an Active Directory or ADAM user attribute**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Account Stores**, right-click **ADAM** or **Active Directory**, point to **New**, and then click **Custom Claim Extraction**.
3. In the **Create a New Custom Claim Extraction** dialog box, in **Attribute**, type the LDAP attribute name for the user.
4. In **Map to this Organization Claim**, select the organization custom claim to map to the attribute, and then click **OK**.

# Map an organization group claim to an ADAM attribute and value (group claim extraction)

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you use Active Directory Application Mode (ADAM) as the Active Directory Federation Services (ADFS) account store for a Federation Service, an organization group claim maps to a Lightweight Directory Access Protocol (LDAP) attribute and value of the user account in ADAM. This mapping is called a group claim extraction. For example, suppose that the organization group claim Manager is mapped to the ADAM user account attribute **memberOf** and the value **CN=ADAMTestGroup,CN=Users,DC=adatum,DC=com**. In this case, if the ADAM store user account for the logged-on user has the **memberOf** attribute and that attribute has a value of **CN=ADAMTestGroup,CN=Users,DC=adatum,DC=com**, the organization group claim Manager is generated for the user. If both the **memberOf** attribute and the corresponding value that is specified in the group claim extraction are not present on the user account, the organization group claim is not generated.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To map an organization group claim to an ADAM attribute and value**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Account Stores**, right-click **ADAM**, point to **New**, and then click **Group Claim Extraction**.
3. In the **Create a New Group Claim Extraction** dialog box, in **Attribute** and **Value**, type the LDAP attribute and its value, respectively.
4. In **Map to this Organization Claim**, select the organization group claim to map to the ADAM attribute and value, and then click **OK**.

**See Also**
**Concepts**
Map an organization group claim to an Active Directory group (group claim extraction)
Map an organization custom claim to an Active Directory or ADAM user attribute (custom claim extraction)

# Map an organization group claim to a resource group

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Organization group claims in the resource Federation Service can be mapped to resource groups, which reside in Active Directory, so that federated users who have the organization group claim are treated as members of that resource group. Therefore, you must create a security group (which is the resource group) in the resource partner forest that represents the users who have the organization group claim but do not have domain user accounts in the resource partner forest. After you create the organization group claim, map the claim to the resource group that you created in the resource partner forest.

| **Note** |
| --- |
| This procedure is not required if the resource partner forest trusts the account partner forest and the Windows trust option is selected in the Active Directory Federation Services snap-in in both the account Federation Service and the resource Federation Service. |

Perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To map an organization group claim to a resource group**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then click **Organization Claims**.
3. In the details pane, right-click the organization group claim that requires mapping to a local group, click **Properties**, and then click the **Resource Group** tab.
4. Select the **Map this claim to the following local resource group** check box, and then click the **...** button.
5. In **Enter the object name to select**, type the name of the resource group that you want to map to the group claim, and then click **OK**.
6. In the **Group Claim Properties** dialog box, click **OK**.

**See Also**
**Concepts**
When to use resource groups

# Create a resource account in the resource partner forest

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You must use this procedure when you want to use the resource account mapping method so that federated users can access Windows NT token–based applications.

| **Important** |
| --- |
| Before you proceed with the steps in this procedure, make sure that you have configured the resource forest for a new user principal name (UPN) suffix according to Checklist: Configuring the resource partner organization. Creating the UPN suffix in the resource partner forest is a precondition for performing this procedure. |

**To create a resource account in the resource partner forest**

1. If you are creating a resource account on a domain controller, click **Start**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.

   | **Note** |
   | --- |
   | You can also open **Active Directory Users and Computers** on a client computer that has the Windows Administration Tools Pack installed. |

2. In the console tree, double-click the domain that you want to add the resource account to, right-click the container or organizational unit (OU) where you want the resource account to reside, point to **New**, and then click **User**.

3. In **User logon name**, type the user logon name of the federated user that will be impersonated, that is, the exact user name of the user account in the account partner forest. Then, in the drop-down list to the right of **User logon name**, select the UPN suffix that matches the name of the account partner forest.

   For example, say that Darren Parker is an employee of A. Datum Corporation. He regularly logs into the network by using his UPN, DarrenP@Adatum.com. The administrator for Tailspin Toys has been given the task of creating a resource account for the DarrenP account. In this example, Adatum.com is the name of the account partner forest, and Tailspintoys.net is the name of the resource partner forest. After following steps 1 and 2, the administrator of Tailspintoys.net types **DarrenP** in the **User logon name** field and then selects **@adatum.com** in the drop-down list. In this example, **adatum.com** is the only choice in the drop-down list because the administrator added adatum.com as a UPN suffix before completing step 3.

4. In **First name** and **Last name**, type the first name and last name of the resource account. Then, in **User logon name (pre–Windows 2000)**, type any unique name, and then click **Next**.

   | **Note** |
   | --- |
   | Federated applications may require that the resource account name be identical to the name of the federated user's account in the account forest. For this reason, consider typing a user name in this field that is identical to the user name for each account in the account partner. |

5. In **Password** and **Confirm password**, type a long, complex password.

6. Verify that the **User must change password at next logon** check box is not selected and that the **Password never expires** check box is selected, and then click **Next**.

   | **Note** |
   | --- |
   | Resource accounts are not intended to be accessed directly by the federated user whose user account is impersonated. Therefore, the federated user does not need to know about the resource account because he or she does not need to log on with this account directly. Also, the password that is used for the resource account does not have to match the password of the impersonated user account. |

7. If you are prompted to create an Exchange mailbox, do not create one. Mailboxes are not necessary for resource

accounts.

# Configure resource account options

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

By adjusting resource account options in the resource Federation Service, resource partner administrators can refine how access control for federated users is administered and delegated on an account partner–by–account partner basis. Use the following procedure to adjust resource account options .

| **Note** |
| --- |
| We recommend that you first review information about each resource account option before you modify the default resource account setting. For more information about resource account options, see Select the optimal resource account option. |

Perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure resource account options**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners**.
3. Right-click the account partner whose resource account behavior you want to change, and then click **Properties**.
4. Click the **Resource Accounts** tab.
5. Select one of the following options for using resource accounts for this account partner, and then click **OK**:
   - **Resource accounts exist for all users** (Resource group claim mappings are not checked.)
   - **Resource accounts exist for some users (prefer resource account)** (First, check for resource accounts. If they do not exist, use the resource groups in the token.)
   - **Resource accounts exist for some users (prefer groups in token)** (First, process resource groups in tokens. If none exist, check for resource accounts.)
   - **No resource accounts exist for this account partner** (Check only for resource groups in tokens.)

**See Also**
**Concepts**
Map an organization group claim to a resource group

# Enable enhanced identity privacy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Enhanced identity privacy is an optional setting that you can configure on a resource partner in the account Federation Service in an Active Directory Federation Services (ADFS) deployment. This setting hashes the user-name portion of outgoing user principal name (UPN) claims and e-mail claims. It substitutes the common name with a random value. If you select the **Enable enhanced identity privacy** option in ADFS, the resource partner will not be able to correlate identity claims to personally identifiable user information.

The enhanced identity privacy setting affects the information that is sent in identity claims, based on the claim type that is being used to transfer the user identity, as follows:

- UPN and e-mail claim types: The user component of the UPN and e-mail name is hashed, replacing the user component in the identity claim of the security token. In this way, each resource partner can uniquely identify each user without revealing their true identity.

- Common name claim types: The common name identity claim is populated with a randomly generated, globally unique identifier (GUID), which ensures that the identity claim is unique per session with the resource partner and that multiple sessions by the same user cannot be tracked.

Enable this setting if you want to:

- Prevent collusion between partners in correlating identity claims to personally identifiable user information.

- Prevent simple dictionary attacks against the user-name hash.

For more information about the effects of enhanced identity privacy, see Partner organizations (http://go.microsoft.com/fwlink/?LinkId=62227).

Perform this procedure on an account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable enhanced identity privacy on a resource partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Resource Partners**.
3. Right-click the resource partner that will begin using enhanced identity privacy, and then click **Properties**.
4. On the **General** tab, click **Enable enhanced identity privacy**, and then click **OK**.

**See Also**
**Concepts**
Configure a policy page for a Web site

# Configure a claims transform module

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

A claims transform module is custom code that manipulates organization, input (incoming), and output (outgoing) claims. Typically, transform modules use the corporate and input claims to produce additional output claims. However, the claim transform module can enumerate, add, delete, and modify claims in any of the claim sets.

Store the claims transform module in %systemdrive%\adfs\sts\bin. This location provides the following advantages:

- ASP.NET keeps a shadow copy of the dynamic-link library (DLL), which allows the DLL to be replaced without stopping the Federation Service, thereby preventing downtime.
- File security is inherited from the \adfs\sts directory.
- The module can be backed up along with all other Active Directory Federation Services (ADFS) files.

After you deploy the transform module to the federation server, perform the following procedure on the account federation server or resource federation server that is configured with the trust policy whose claims transform module you are configuring.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure a claims transform module**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click the **Trust Policy** node, and then click **Properties**.
3. On the **Transform Module** tab, configure the DLL file and class name for the module as follows:
   a. **DLL file:** Click **Browse** to navigate to the DLL that implements the claim transform module, and then click **Open**.

      | **Note** |
      | --- |
      | This DLL must be a managed-code assembly. |

   b. **Class name:** Type the namespace-qualified class name that implements the claim transform interface (**IClaimTransform**, which is defined in System.Web.Security.SingleSignOn.ClaimTransforms.dll): The namespace qualified name should be of the format namespace.classname.
4. Click **OK** to save the configuration.

# Configure an account partner to use Windows trust

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You must configure the account partner to select the domains that are to be included in the trust relationship. For example, if a Windows Server 2003 forest trust is in place from the resource partner forest to the account partner forest, the trust is transitive to all domains in the trusted account forest.

If you want to grant access to resources in the resource forest to users in only some domains in the account forest, you can specify only those domains. Otherwise, you can allow all trusted domains (all domains in the account forest and any forest that is trusted by the account forest) to be granted access.

Use the following procedure to enable Windows trust for the account partner that will participate in the Federated Web SSO with Forest Trust design scenario.

Perform this procedure on a federation server in the resource partner organization.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure an account partner to use Windows trust**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners**.
3. Right-click the account partner that you want to configure to use Windows trust, and then click **Properties**.
4. On the **Windows Trust** tab, click **Use Windows trust relationship**.
5. In **Trusted Active Directory domains and forests**, do one of the following, and then click **OK**:
   - Click **All Active Directory domains and forests** if you want to allow users in all trusted domains in the account partner forest and in any forest that is trusted by the account partner forest, and then click **OK**.
   - Click **Specified Active Directory domains and forests (press Enter to separate entries)** if you want to name only the domains where you want to allow federated access. Type a domain name, press ENTER, and then repeat this action to add each domain in the account partner forest and in any other trusted forests for the users that you want to grant access to resources.

**See Also**
**Concepts**
Configure a resource partner to use Windows trust
Discontinue Windows trust for an account partner

# Configure a resource partner to use Windows trust

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Use the following procedure to enable Windows trust for a resource partner in an Active Directory Federation Services (ADFS) Federated Web SSO with Forest Trust scenario.

| **Note** |
| --- |
| When you enable the Windows trust option in the account Federation Service, you are sending actual security identifiers (SIDs) to the resource partner organization over the Internet, which may be a security risk. These SIDs are packaged in the ADFS Security Assertion Markup Language (SAML) token. Therefore, enable this option only when you are using the Federated Web SSO with Forest Trust design. This design is meant to establish secure communication within the same organization. |

Perform this procedure on an account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure a resource partner to use Windows trust**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Resource Partners**.
3. Right-click the resource partner for which you want to configure Windows trust, and then click **Properties**.
4. On the **General** tab, click **Use Windows trust relationship for this partner**, and then click **OK**.

**See Also**
**Concepts**
Configure an account partner to use Windows trust
Discontinue Windows trust for a resource partner

# Distribute certificates to client computers using Group Policy

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You can use this procedure to push down the appropriate Secure Sockets Layer (SSL) certificates (or equivalent certificates that chain to a trusted root) for account federation servers, resource federation servers, and Active Directory Federation Services (ADFS)–enabled Web servers to each client computer in the account partner forest by using Group Policy.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory or you must have been delegated the appropriate authority.

**To distribute certificates to client computers using Group Policy**

1. On a domain controller in the forest of the account partner organization, click **Start**, point to **Administrative Tools**, and then click **Domain Security Policy**.
2. In the console tree, double-click **Public Key Policies**, right-click **Trusted Root Certification Authorities**, and then click **Import**.
3. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
4. On the **File to Import** page, type the path to the appropriate certificate files (for example, \\adfsresource\c$\adfsresource.cer), and then click **Next**.
5. On the **Certificate Store** page, click **Place all certificates in the following store**, and then click **Next**.
6. On the **Completing the Certificate Import Wizard** page, verify that the information you provided is accurate, and then click **Finish**.
7. Repeat steps 2 through 6 to add additional certificates for each of the ADFS servers.

# Configure client computers to trust the account federation server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

So that client computers can successfully access federated applications, you must first configure each client's Internet Explorer settings so that the browser trusts the account federation server. You can do this manually or through Group Policy, depending on your administrative preference, by completing one of the following tasks.

**Configuring Internet Explorer settings manually**
Use the following task to manually configure each user's Internet Explorer settings to support federation. If multiple users use a single computer, complete this task multiple times—once for each user profile.

### Administrative credentials

To perform this task, log on as the user who will be accessing federated applications. This is a profile-specific setting. Therefore, it requires that you manually add this setting for each profile that exists on a specific computer.

**To manually configure client computers to trust the account federation server**

1. On the client computer, start Internet Explorer.
2. On the **Tools** menu, click **Internet Options**.
3. On the **Security** tab, click the **Local intranet** icon, and then click **Sites**.
4. Click **Advanced**, and in **Add this Web site to the zone**, type the full Domain Name System (DNS) name of the account federation server (for example, https://adfsaccount.adatum.com), and then click **Add**.
5. Click **OK** three times.

**Configuring Internet Explorer settings using Group Policy**
For most deployments, we recommend that you use Group Policy to push the appropriate Internet Explorer settings to each client computer.

### Administrative credentials

To perform this task, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority.

**To configure client computers to trust the account federation server using Group Policy**

1. Edit the appropriate Group Policy object (GPO) with either the Active Directory User and Computers snap-in or the Group Policy Manager Console snap-in.
2. In Group Policy, click **User Configuration**, click **Windows Settings**, click **Internet Explorer Maintenance**, and then click **Security**.
3. In the details pane, double-click **Security Zones and Content Ratings**.
4. Under **Security Zones and Privacy**, click **Import the current security zones and privacy settings**, and then click **Modify Settings**.
5. Select **Local intranet**, click **Sites**.
6. Click **Advanced**, and in **Add this Web site to the zone**, type the full DNS name of the account federation server (for example, https://adfsaccount.adatum.com), click **Add**, and then click **Close**.
7. Click **OK** three times to apply these changes to Group Policy.

# Deploying Federated Applications

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

To deploy a federated Windows NT token–based application or claims-aware application in the resource partner organization, complete each of the tasks in Checklist: Installing a Windows NT token-based application or Checklist: Installing a claims-aware application.

| **Note** |
| --- |
| When you use this checklist, we strongly recommend that you first read the references to federated application planning in the ADFS Design Guide before continuing to the procedures for configuring the servers. Following the checklist in this way helps provide a better understanding of the full Active Directory Federation Services (ADFS) design and deployment story for federated applications. |

The type of federated application that you choose to deploy will determine the type of ADFS Web Agent and setup tasks that you need to apply to your ADFS-enabled Web server. In other words, if your Web server will host only claims-aware applications, you install only the assemblies of the ADFS Web Agent that are used for claims-aware applications. For more information, see Identify the type of federated application to deploy.

**About Windows NT token–based applications**

A Windows NT token–based application is an Internet Information Services (IIS) application that has been written to use Windows native authorization mechanisms and that is not prepared to consume ADFS claims. Microsoft Windows® SharePoint® Services is one example of a Windows NT token–based application. You can configure Windows SharePoint Services, through ADFS, to provide access to SharePoint sites across the Internet.

**About claims-aware applications**

Claims are statements about a user, such as the user's e-mail address or job function. A claims-aware application is a Microsoft ASP.NET 2.0 application that has been written using the ADFS library and that is fully capable of using claims to make authorization decisions directly. A claims-aware application accepts claims that the Federation Service sends in ADFS security tokens.

# Checklist: Installing a claims-aware application

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This checklist includes the deployment tasks for preparing an Active Directory Federation Services (ADFS)-enabled Web server running Windows Server 2003 R2, Standard Edition, or Windows Server 2003 R2, Enterprise Edition, for installation of a claims-aware application.

**Important**

Make sure that you have set up your ADFS-enabled Web server using the guidance in Checklist: Installing an ADFS-enabled Web server before you proceed with the tasks in this checklist.

**Note**

Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

### Checklist: Installing a claims-aware application

| | Task | Reference |
|---|---|---|
| ☐ | Review information in the Active Directory Federation Services Design Guide about the role that ADFS Web Agents play in deploying federated applications. | Review the role of ADFS Web Agents |
| ☐ | Review the information in the Active Directory Federation Services Design Guide to understand important concepts that are relevant to designing and deploying federated claims-aware applications. | Designing a Federated Application Strategy<br>Identify the type of federated application to deploy |
| ☐ | Review information in the Active Directory Federation Services Design Guide to determine whether to use Public Key Infrastructure (PKI) or Service Principal Name (SPN) for the security token protection method. | Determine your security token protection method |
| ☐ | Review information in the Active Directory Federation Services Design Guide about the benefits of using Authorization Manager to make administration easier by using role-based access control methods. | When to use Authorization Manager |
| ☐ | Using Internet Information Services (IIS), configure your claims-aware application to enable anonymous access and to use ASP.NET 2.0. | Configure IIS to support a federated application |
| ☐ | Edit the web.config file for a claims-aware application so that it can use the claims-aware Web Agent. | Configure web.config to use the claims-aware Web Agent |
| ☐ | Use the Add Application Wizard in the ADFS snap-in to add a new claims-aware application entry to the Federation Service.<br>**Note**<br>Before you begin this task, first configure the trust policy in the resource partner organization. If you have not yet created and configured the trust policy in the resource partner organization, complete the tasks in Checklist: Configuring the resource partner organization. | Add a new claims-aware application to the Federation Service |

| | | |
|---|---|---|
| ☐ | Enable any claims that you want to be sent to the application for authorization purposes.<br>**Note**<br>Before you begin this task, first install the appropriate claims in the resource partner organization. | Enable an organization claim for a federated application in the Federation Service |
| ☐ | (Optional) If your application requires them, you can choose authentication method restrictions for your claims-aware application. By default, an application accepts any authentication method that is provided to it. | Configure authentication methods for a federated application |
| ☐ | (Optional) If your application requires it, you can have the resource federation server sign security tokens using Kerberos or PKI signing schemes. The default token-signing scheme is set to PKI. | Configure the security token protection method for a federated application |
| ☐ | To ensure the likelihood of successfully tracking down issues that may occur with your claims-aware application, configure event logging. | Configure event logging for a claims-aware application |
| ☐ | From a client computer, verify that the claims-aware application can be accessed using Integrated Windows authentication, and verify that anonymous access is enabled. | Verify that an ADFS-enabled Web server is operational |

# Checklist: Installing a Windows NT token-based application

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This checklist includes the deployment tasks for preparing an Active Directory Federation Services (ADFS)-enabled Web server running Windows Server 2003 R2, Standard Edition, or Windows Server 2003 R2, Enterprise Edition, for the installation of a Windows NT token–based application.

**Important**

Make sure that you have set up your ADFS-enabled Web server using the guidance in Checklist: Installing an ADFS-enabled Web server before you proceed with the tasks in this checklist.

**Note**

Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

### Checklist: Installing a Windows NT token–based application

| | Task | Reference |
|---|---|---|
| ☐ | Review information in the Active Directory Federation Services Design Guide about the role that ADFS Web Agents play in deploying federated applications. | Review the role of ADFS Web Agents |
| ☐ | Review the information in the Active Directory Federation Services Design Guide to understand important concepts that are relevant to designing and deploying federated Windows NT token–based applications. | Designing a Federated Application Strategy<br>Identify the type of federated application to deploy |
| ☐ | Review information in the Active Directory Federation Services Design Guide to determine whether to use Public Key Infrastructure (PKI) or Service Principal Name (SPN) for the security token protection method. | Determine your security token protection method |
| ☐ | Based on the needs of your Windows NT token–based application design, determine whether to use the resource account, resource group, or group-to-UPN (user principal name) impersonation method for purposes of authorization. | Checklist: Implementing a resource account mapping method |
| ☐ | Using Internet Information Services (IIS), configure your Windows NT token–based application to enable anonymous access and to use ASP.NET 2.0. | Configure IIS to support a federated application |
| ☐ | Edit the ADFS Web Agent tabs in IIS to enable the Windows NT token–based Web Agent. | Configure the Windows NT token-based Web Agent |
| ☐ | Use the Add Application Wizard in the ADFS snap-in to add a new Windows NT token–based application entry to the Federation Service.<br>**Note**<br>Before you begin this task, first configure the trust policy in the resource partner organization. If you have not yet created and configured the trust policy in the resource partner organization, complete the tasks in Checklist: Configuring the resource partner organization. | Add a new Windows NT token-based application to the Federation Service |

| | | |
|---|---|---|
| ☐ | Enable any claims that you want to be sent to the application for authorization purposes.<br>**Note**<br>Before you begin this task, first install the appropriate claims in the resource partner organization. | Enable an organization claim for a federated application in the Federation Service |
| ☐ | (Optional) If your application requires them, you can choose authentication method restrictions for your claims-aware application. By default, an application accepts any authentication method that is provided to it. | Configure authentication methods for a federated application |
| ☐ | (Optional) If your application requires it, you can have the resource federation server sign security tokens using Kerberos or PKI signing schemes. The default token-signing scheme is set to PKI. | Configure the security token protection method for a federated application |
| ☐ | To ensure the likelihood of successfully tracking down issues that may occur with your Windows NT token–based application, configure event logging. | Configure event logging for a Windows NT token-based application |
| ☐ | From a client computer, verify that the Windows NT token–based application can be accessed using Integrated Windows authentication, and verify that anonymous access is enabled. | Verify that an ADFS-enabled Web server is operational |

# Configure IIS to support a federated application

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Claims-aware applications and Windows NT token–based applications that participate in federation must be configured to use Microsoft ASP.NET 2.0. These applications must also be configured to support sign-in requests from anonymous users. Complete the following procedures so that federated users can successfully access federated applications that are hosted on an ADFS-enabled Web server.

**Configure IIS to use ASP.NET 2.0**

You can use the following procedure to configure ASP.NET 2.0 in Internet Information Services (IIS) for your federated application.

| **Note** |
| --- |
| If you are installing Windows SharePoint Services or SharePoint Portal Server 2003 for federation, do not configure IIS to use ASP.NET 2.0. These applications configure their own virtual directories for ASP.NET 2.0 automatically, and you do not have to configure them manually. |

**To configure IIS to use ASP.NET 2.0**

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, double-click the *ComputerName* folder, double-click **Web Sites**, right-click the *WebSiteName* folder, and then click **Properties**.

   For example, if you want to configure IIS for a Web site serving an application named testapp on the ADFSWeb computer, double-click the **ADFSWEB** folder, double-click **Web Sites**, right-click the **testapp** folder, and then click **Properties**.
3. On the **ASP.NET** tab, in **ASP.NET version**, select **2.0**, and then click **OK**.

| **Note** |
| --- |
| To ensure that ASP.NET 2.0 is enabled properly, verify that **ASP.NET v2.0.50727** is set to **Allowed** under *ComputerName***/Web Service Extensions**. |

**Configure IIS to enable anonymous access**

You can use the following procedure to enable anonymous access in IIS for your federated application.

| **Note** |
| --- |
| When you enable the Windows NT token–based Web Agent in IIS, anonymous access is enabled automatically. Therefore, if you are deploying a Windows NT token-based application on the ADFS-enabled Web server, you can choose to enable anonymous access using this procedure or you can bypass this procedure and let the Windows NT token–based Web Agent do this for you after it is enabled. |

**To configure IIS to enable anonymous access**

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, double-click the *YourComputerName* folder, double-click **Web Sites**, double-click the *YourWebSiteName* folder, right-click your application, and then click **Properties**.
3. On the **Directory Security** tab, under **Authentication and access control** click **Edit**, and then verify that the **Enable anonymous access** check box is selected.
4. Click **OK** twice to apply the settings and exit the property pages.

# Configure web.config to use the claims-aware Web Agent

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You can use the following procedure to configure the web.config file that your claims-aware application uses to work with the claims-aware Web Agent. When you complete the procedure, the claims-aware Web Agent will be able to:

- Locate and collaborate with the resource federation server that is required to authenticate users to the application.
- Identify the claim application's Uniform Resource Locator (URL) that is used by the Federation Service.
- Identify the local path to the claims application on the ADFS-enabled Web server.
- Identify the local path to use to store log files for the application.

The changes that you make to the web.config file are also used to inform the claims-aware application where to load the Active Directory Federation Service (ADFS) assemblies that are necessary to start the ADFS Web Agent Authentication Service.

**To configure the web.config to use the claims-aware Web Agent**

1. Locate the web.config file that is used by your claims-aware application, and then open it with Notepad. This file should be located in \inetpub\wwwroot*virtualdirectory* where your claims-aware application files are stored.
2. If the following code does not exist in the web.config file, paste the code into the file:

```
<?xml version="1.0" encoding="utf-8" ?>

<configuration>

<configSections>

<sectionGroup name="system.web">

<section name="websso" type="System.Web.Security.SingleSignOn.WebSsoConfigurationHandler,
System.Web.Security.SingleSignOn, Version=1.0.0.0, Culture=neutral, PublicKeyToken=31bf3856ad364e35,
Custom=null" />

</sectionGroup>

</configSections>

<system.web>

<compilation defaultLanguage="c#" debug="true">

<assemblies>

<add assembly="System.Web.Security.SingleSignOn, Version=1.0.0.0, Culture=neutral,
PublicKeyToken=31bf3856ad364e35, Custom=null" />

<add assembly="System.Web.Security.SingleSignOn.ClaimTransforms, Version=1.0.0.0, Culture=neutral,
PublicKeyToken=31bf3856ad364e35, Custom=null" />

</assemblies>

</compilation>

<customErrors mode="Off" />

<authentication mode="None" />

<httpModules>

<add name="Identity Federation Services Application Authentication Module"
type="System.Web.Security.SingleSignOn.WebSsoAuthenticationModule, System.Web.Security.SingleSignOn,
```

```
Version=1.0.0.0, Culture=neutral, PublicKeyToken=31bf3856ad364e35, Custom=null" />

</httpModules>

<websso>

<urls>

<returnurl>https://myapp</returnurl>

</urls>

<cookies writecookies="true">

<path>/apppath</path>

<lifetime>240</lifetime>

</cookies>

<fs>https://myfederationserver/adfs/fs/federationserverservice.asmx</fs>

<authenticationrequired>

</authenticationrequired>

<loghttpevent>1</loghttpevent>

<auditlevel>255</auditlevel>

<tokenCacheSize>1</tokenCacheSize>

<tokenCacheEntryLifetime>5</tokenCacheEntryLifetime>

<tokenCacheScavengePeriod>5</tokenCacheScavengePeriod>

</websso>

</system.web>

<system.diagnostics> (this section is optional and is used for debugging)

<switches>

<add name="WebSsoDebugLevel" value="15" />

</switches>

<trace autoflush="true" indentsize="3">

<listeners>

<add name="ADFSLogListener" type="System.Web.Security.SingleSignOn.BoundedSizeLogFileTraceListener,
System.Web.Security.SingleSignOn, Version=1.0.0.0, Culture=neutral, PublicKeyToken=31bf3856ad364e35,
Custom=null" initializeData="applogfile"/>

</listeners>

</trace>

</system.diagnostics>

</configuration>
```

| **Note** |
| --- |
| In the following steps of this procedure, replace the highlighted text with information that is specific to your application and your federation server. |

3. Replace **myapp** in `<returnurl>https://myapp</returnurl>` with the URL address for the claims-aware application that will be loaded. The URL value here must match the **Application URL** value that is specified in the properties page of the claims-aware application (which is located under **Federation Service\Trust Policy\Applications** in the Active Directory Federation Services snap-in) of the resource partner.

4. Replace **`apppath`** in `<path>/apppath</path>` with the path of the virtual directory where you are storing the claims-aware application.
5. Replace **`myresourcefederationserver`** in `<fs>https://myresourcefederationserver/adfs/fs/federationserverservice.asmx</fs>` with a valid resource federation server name.
6. Replace **`applogfile`**in`initializeData="applogfile"` with the local path to the location of the application's log file.
7. In Notepad, on the **File** menu, click **Save**.

# Configure the Windows NT token-based Web Agent

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

The following procedure must be completed on the Web server so that clients in the account partner organization can access Windows NT token–based applications, such as SharePoint sites, that are hosted on the Web server in the resource partner organization.

**To enable the Windows NT token–based Web Agent**

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, double-click *YourComputerName***(local computer)**, right-click **Web Sites**, and then click **Properties**.
3. On the **ADFS Web Agent** tab, in **Federation Service URL**, type the Uniform Resource Locator (URL) of the federation server in the resource partner organization, and then click **OK**.

   For example, if the federation server in the resource partner organization is named fedsrv1 and it is located in the treyresearch.net forest, type **https://fedsrv1.treyresearch.net/adfs/fs/FederationServerService.asmx**.

   | **Note** |
   | --- |
   | The Federation Service URL in Active Directory Federation Services (ADFS) defines the URL that is used for all Web sites and Windows NT token–based applications on a Web server where the ADFS Web Agent is enabled. |

4. Double-click **Web Sites**, right-click the Web site or virtual directory where the application will reside, and then click **Properties**.
5. On the **ADFS Web Agent** tab, select the **Enable Active Directory Federation Services Web Agent** check box.

   | **Note** |
   | --- |
   | This action will enable anonymous access to this Web site. |

6. Modify the following values as necessary, and then click **OK**.
   - Cookie path
   - Cookie domain
   - Return URL

   For detailed information about each of these settings, see Review the role of ADFS Web Agents.

# Add a new claims-aware application to the Federation Service

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you configure your Active Directory Federation Services (ADFS)–enabled Web server to host a claims-aware application, you must add a reference to that application in the trust policy so that you can specify application properties for the Federation Service. You can use the following procedure to add a new claims-aware application to the Federation Service. You perform this procedure on a federation server in the resource partner.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new claims-aware application to the Federation Service**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Applications**, point to **New**, and then click **Application**.
3. On the **Welcome to the Add Application Wizard** page, click **Next**.
4. On the **Application Type** page, click **Claims-aware application**, and then click **Next**.
5. On the **Application Details** page, in **Application display name**, type the name of the application.
6. In **Application URL**, type the Uniform Resource Locator (URL) for your application—for example, https://www.treyresearch.net/*ApplicationName*/—and then click **Next**.

   | **Note** |
   | --- |
   | The value in **Application URL** on this property page must match precisely the Return URL value (in **<returnurl> </returnurl>**) that you specify in the web.config file of the claims-aware application (on the ADFS-enabled Web server). |

7. On the **Accepted Identity Claims** page, select the check boxes for the claim types that will be recognized by your claims-aware application, and then click **Next**.
8. On the **Enable this Application** page, ensure that the **Enable this application** check box is selected, and then click **Next**.
9. On the **Completing the Add Application Wizard** page, click **Finish**.

# Add a new Windows NT token-based application to the Federation Service

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you configure your Active Directory Federation Services (ADFS)–enabled Web server to host a Windows NT token–based application, you must add a reference to that application in the trust policy so that you can specify application properties for the Federation Service. You can use the following procedure to add a new Windows NT token–based application to the Federation Service. You perform this procedure on a federation server in the resource partner.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a Windows NT token–based application to the Federation Service**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Applications**, point to **New**, and then click **Application**.
3. On the **Welcome to the Add Application Wizard** page, click **Next**.
4. On the **Application Type** page, click **Windows NT token–based application**, and then click **Next**.
5. On the **Application Details** page, in **Application display name**, type the name of the application.
6. In **Application URL**, type the Uniform Resource Locator (URL) for your application—for example, https://www.treyresearch.net/*ApplicationName*/—and then click **Next**.

   | **Note** |
   | --- |
   | The value in **Application URL** on this property page must match precisely the **Return URL** value that you specify in the **ADFS Web Agent** tab for the Windows NT token-based application (on the ADFS-enabled Web server). |

7. On the **Accepted Identity Claim** page, click one of the following identity claim types, based on the needs of your application, and then click **Next**:
   - **User principal name (UPN)**
   - **E-mail**
8. On the **Enable this Application** page, ensure that the **Enable this application** check box is selected, and then click **Next**.
9. On the **Completing the Add Application Wizard** page, click **Finish**.

# Enable an organization claim for a federated application in the Federation Service

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

To participate in federation, a resource partner must first create organization identity claims. When an organization claim is created in the resource partner, the claim is not sent to the application unless you first enable the claim after you create an entry in the Federation Service for a new claims-aware application or Windows NT token–based application. You can use the following procedure to enable an organization claim for an application in the Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable an organization claim for a federated application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Applications**, and then click the application where you want to enable the organization claim.
3. In the details pane, right-click the organization claim that you want to enable, and then click **Enable**.

# Configure authentication methods for a federated application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

By default, a Web application that you configure in Active Directory Federation Services (ADFS) accepts any of the following authentication methods:

- Windows Integrated authentication (also referred to as Windows NT Challenge/Response authentication): A secure form of authentication because the user name and password are hashed before they are sent across the network.

- User name and password authentication (also called Basic authentication): A widely used, industry-standard method for transmitting user names and passwords across the network in an unencrypted form.

- Certificate or Secure Sockets Layer / Transport Layer Security (SSL/TLS) client authentication: A form of digital identification for your server and for clients that request information from your server. Its function is similar to that of a passport, or other official identity card, that identifies the person carrying it. Certificates are part of the SSL features in Internet Information Services (IIS) that establish a secure connection for sensitive information.

If you want to use a specific authentication method or methods, you can use the following procedure to select one or more methods that will apply to users of the application. You perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure authentication methods for a federated application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Applications**, right-click the application whose authentication methods you want to change, and then click **Properties**.
3. On the **Authentication Methods** tab, to select from the available authentication methods, do one of the following:
   - To allow all possible authentication methods in the list, select the **Any** check box.
   - To allow one or more of the available authentication methods, clear the **Any** check box, and then select one or more of the available methods.
4. When you finish selecting authentication methods, click **OK**.

# Configure the security token protection method for a federated application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Security tokens that are received by Active Directory Federation Services (ADFS) federation servers are protected during transit by one of two methods:

- Public Key Infrastructure (PKI): A PKI is implemented as a hierarchy of certification authorities that verify identities. When a PKI is in place, a signature is embedded into the token that protects the token from tampering.

- Domain service account: A domain service account, which is identified by a service principal name (SPN), runs under an account that is trusted for delegation and that can impersonate a client to gain access to resources. By default, this account is the Internet Information Services (IIS) application pool identity that hosts a claims-aware application and the identity of the ADFS Web Agent Authentication Service that hosts a Windows NT token–based application. When a token is transferred in a domain service account with this setting, the token contains a binary Kerberos V5 signature for the configured SPN. This signature protects the token from tampering.

You can use the following procedure to change the security token protection method for a federated application on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the security token protection method for a federated application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Trust Policy**, double-click **My Organization**, and then double-click **Applications**.
3. Right-click the application whose security token protection method you want to change, and then click **Properties**.
4. On the **General** tab, under **Security token protection method**, do one of the following, and then click **OK**:
   - If your deployment uses certificates that are issued by a certification authority (CA), select **Public Key Infrastucture (PKI)**.
   - If your deployment does not use certificates that are issued by a CA, select **Domain service account** and then, in **service principal name (SPN) of service account**, type the SPN of the account.

**See Also**
**Concepts**
Determine your security token protection method

# Configure event logging for a claims-aware application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Use the following procedure to specify the level of events that you want to be logged for claims-aware applications in the Application event log on the Web server that is protected by Active Directory Federation Services (ADFS).You set event logging for claims-aware applications in the Web.config file for the application.

You can apply the following event logging settings in the web.config file:

- **Error:** An alert that provides information about a significant problem of which the user should be aware, usually involving a loss of functionality or data.
- **Warning:** An alert that indicates a problem that is not immediately significant but that may signify conditions that could cause future issues.
- **Info:** An alert that provides information about a significant, successful operation.
- **SuccessAudit:** A security event that occurs when an audited access attempt succeeds, for example, a successful logon attempt.
- **FailureAudit:** A security event that occurs when an audited access attempt fails, for example, a failed attempt to open a file.
- **DetailedSuccess:** A success audit event that provides detailed information about each token that is involved in the transaction, including claims information.
- **DetailedFailure:** A failure audit event that provides detailed information about each token that is involved in the transaction, including claims information.
- **Everything:** Enables all logging levels.

### Administrative credentials

To complete this procedure, you must have Read-Write access to the Web.config file.

#### To configure event logging for a claims-aware application

1. In Notepad, open the Web.config file in the directory that stores the claims-aware application.
2. Search the file for <websso>.
3. Add or change the **<auditlevel>** entry under **<websso>**, as follows:

   **<auditlevel>** *Value* **</auditlevel>**

   Where *Value* is one of the following or the combined values of two or more:

   - Error = 1
   - Warning = 2
   - Info = 4
   - SuccessAudit = 16
   - FailureAudit = 32
   - DetailedSuccess = 64

- DetailedFailure = 128
- Everything = 247

4. Save and close the web.config file.

# Configure event logging for a Windows NT token-based application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

On a Web server that is running the Active Directory Federation Services (ADFS) Web Agent for Windows NT token–based applications, you can configure the types of events that you want to be logged for Windows NT token–based applications.

| **Note** |
| --- |
| Event logging is enabled differently for Windows NT token–based applications and claims-aware applications. |

Use the following procedure to specify the types of events that you want to be logged for Windows NT token–based applications in the Application event log on the ADFS-enabled Web server. You set event logging for Windows NT token–based applications in the registry of the Web server.

| **Caution** |
| --- |
| Editing the registry incorrectly may damage your system severely. Before you make changes to the registry, back up any valued data on the computer. |

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

#### To configure event logging for Windows NT token–based applications

1. Open Regedit. Click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Navigate to:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ifssvc\ Parameters**
3. Right-click **Parameters**, click **New**, and then click **DWORD Value**.
4. In the new value file name box, type the following, and then press Enter:

   **ADFSEvent**
5. Double-click the new entry, and then, in **Value data**, provide a value for one of the following levels, or add values to configure multiple levels, and then click **OK**:
   - Warning: 0x01
   - Information: 0x02
   - Success: 0x04
   - Failure: 0x08

# Deploying ADFS-Enabled Web Servers

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

To deploy Active Directory Federation Services (ADFS)-enabled Web servers, complete each of the tasks in Checklist: Installing an ADFS-enabled Web server. After you complete the tasks in this checklist, you can set up the Web server to host claims-aware applications or Windows NT token–based applications in the resource partner organization.

| **Note** |
| --- |
| When you use this checklist, we strongly recommend that you first read the references to Web server planning in the ADFS Design Guide before continuing to the procedures for configuring the servers. Following the checklist in this way helps provide an understanding of the full ADFS design and deployment story for Web servers. |

**About ADFS-enabled Web servers**

In ADFS, Web servers in the resource partner forest host the ADFS Web Agent component to provide secure access to federated Web applications that are hosted on those Web servers. The ADFS Web Agent receives security tokens and authentication cookies that are sent to the Web server from the resource Federation Service. The Web server requires a relationship with a Federation Service in the resource partner, where the forest that the Web server resides in trusts the forest where the resource federation server resides, so that all trusted authentication tokens come from that Federation Service.

The ADFS Web Agent supports two types of applications: claims-aware applications and Windows NT token–based applications. A claims-aware application is a Microsoft ASP.NET 2.0 application that is fully capable of using ADFS claims to make authorization decisions. A Windows NT token–based application is an Internet Information Services (IIS) application that is written to use Windows native authorization mechanisms and that is not capable of consuming ADFS claims.

The type of Web-based applications that your Web server will be hosting determines the type of ADFS Web Agent that you install on the Web server. That is, if your Web server will host only claims-aware applications, you install only the assemblies of the ADFS Web Agent that are used for claims-aware applications. If you have an existing application that uses Windows Integrated authentication, you install only the assemblies of the ADFS Web Agent that are used for Windows NT token–based applications so that the application can use ADFS for authentication.

# Checklist: Installing an ADFS-enabled Web server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This checklist includes the deployment tasks for preparing a server running Windows Server 2003 R2, Standard Edition, or Windows Server 2003 R2, Enterprise Edition, for the Active Directory Federation Services (ADFS)-enabled Web server role.

| **Note** |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist. |

### Checklist: Installing an ADFS-enabled Web server

| | Task | Reference |
| --- | --- | --- |
| ☐ | Review information in the ADFS Design Guide about where to place ADFS-enabled Web servers in your organization. | Planning ADFS-Enabled Web Server Placement<br>Where to place an ADFS-enabled Web server |
| ☐ | Use the information in the ADFS Design Guide to determine whether a single ADFS-enabled Web server or a Web server farm is appropriate for your deployment. | When to create an ADFS-enabled Web server<br>When to create an ADFS-enabled Web server farm |
| ☐ | Review information in the ADFS Design Guide about how ADFS-enabled Web servers require server authentication certificates to authorize client requests securely. | Certificate requirements for ADFS-enabled Web servers |
| ☐ | Review information in the ADFS Design Guide about how to update the perimeter network Domain Name System (DNS) so that successful name resolution between clients and ADFS-enabled Web servers in farms can occur. | Name resolution requirements for ADFS-enabled Web servers |
| ☐ | Join the computer that will become the ADFS-enabled Web server to a domain in the resource partner forest where it will be used to authorize federated clients.<br>**Note**<br>If your ADFS-enabled Web server will be hosting a Windows NT token–based application, the server must be joined to a domain in the same forest, or in a trusting forest, where the resource federation server resides. | Join a computer to a domain |
| ☐ | Create a new resource record in the perimeter network DNS that points the DNS host name of the ADFS-enabled Web server to the IP address of the ADFS-enabled Web server. | Add a host (A) record to perimeter DNS for an ADFS-enabled Web server |
| ☐ | Install prerequisite applications such as, ASP.NET, Internet Information Services (IIS), and Microsoft .NET Framework 2.0 on the computer that will become the ADFS-enabled Web server. | Install prerequisite applications |

| | | |
|---|---|---|
| ☐ | After you obtain a server authentication certificate (or a private key), install it in IIS on the appropriate Web site or virtual directory where your federated application will reside.<br><br>For an example of how to do this using the default Web site, see the link to the right.<br><br>**Note**<br>If you will be adding an ADFS-enabled Web server to an existing ADFS-enabled Web server farm, you must add the same server authentication certificate that you receive from the certification authority (CA) to the appropriate Web site or virtual directory where your federated application will reside on each of the servers that will be participating in the farm. | Import a server authentication certificate to the default Web site |
| ☐ | (Optional) In a scenario in which you want to install the Federation Service on your ADFS-enabled Web server so that the same server will play both the ADFS-enabled Web server role and the federation server role, configure certificates in the following way:<br><br>• Install the server authentication certificate on the appropriate Web site or virtual directory where your application will reside, as indicated in the previous step.<br><br>• Install the server authentication certificate for the federation server. This certificate must be installed in the **Local Computer** certificate store of the ADFS-enabled Web server, and its root certificate or certificates must also be installed in the **Trusted Root** certificate store.<br><br>**Note**<br>Use the Certificate snap-in to install certificates to the appropriate store.<br><br>• Install the token-signing certificate that the federation server will use to sign its tokens. This certificate must be installed in the **Local Computer** certificate store of the ADFS enabled web server, and its root certificate or certificates must also be installed in the **Trusted Root** certificate store.<br><br>**Note**<br>Use the Certificate snap-in to install certificates to the appropriate store. | N/A |
| ☐ | (Optional) As an alternative to obtaining a server authentication certificate from a CA, you can use the SelfSSL.exe tool to create a self-signed certificate for your ADFS-enabled Web server.<br><br>Because the SelfSSL tool generates a self-signed certificate that does not originate from a trusted source, use the SelfSSL tool only in the following scenarios:<br><br>• When you have to create a Secure Sockets Layer (SSL) channel between your server and a limited, known group of users<br><br>• When you have to troubleshoot third-party certificate problems<br><br>**Caution**<br>It is not a security best practice to deploy an ADFS-enabled Web server in a production environment using a self-signed server authentication certificate. | Internet Information Services (IIS) 6.0 Resource Kit Tools (http://go.microsoft.com/fwlink/?LinkId=36285) |

| | | |
|---|---|---|
| ☐ | Install the ADFS Web Agent component on the computer that will become the ADFS-enabled Web server. | Install the ADFS Web Agent component of ADFS |
| ☐ | Install and configure a claims-aware application or a Windows NT token–based application on your new ADFS-enabled Web server. | Checklist: Installing a claims-aware application<br>Checklist: Installing a Windows NT token-based application |
| ☐ | From a client computer, verify that the ADFS-enabled Web server is operational. | Verify that an ADFS-enabled Web server is operational |

# Join a computer to a domain

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

For Active Directory Federation Services (ADFS) to function, each computer that functions as a federation server must be joined to a domain. Federation server proxies may be joined to a domain, but this is not a requirement.

ADFS-enabled Web servers must be joined to any domain in the same forest where the resource federation server resides whenever the Web servers host a Windows NT token–based application. This is a requirement that makes it possible for the Windows NT token-based Web agent to properly retrieve group membership information from Active Directory.

You do not have to join an ADFS-enabled Web server to a domain if the Web server is hosting claims-aware applications only.

**To join a computer to a domain**

1. On the computer that you want to join to a domain, click **Start**, click **Control Panel**, and then double-click **System**.
2. On the **Computer Name** tab, click **Change**.
3. Under **Member of**, click **Domain**, type the name of the domain that this computer will join, and then click **OK**.
4. Click **OK**, and then restart the computer.

**See Also**
**Concepts**
Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server

# Add a host (A) record to perimeter DNS for an ADFS-enabled Web server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

For clients on the Internet to successfully access a federated application that is hosted on an Active Directory Federation Services (ADFS)-enabled Web server, an address (A) resource record must first be created in the perimeter DNS that resolves the fully qualified domain name (FQDN) name of the Web server or server cluster to the IP address of the Web server or server cluster.

For more information about how to configure a cluster IP address or a cluster FQDN using Microsoft Network Load Balancing (NLB) technology, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?LinkID=74651).

**To add a host (A) record to perimeter DNS for an ADFS-enabled Web server**

1. On a DNS server for the perimeter network, open the DNS snap-in.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A)**.
3. In **Name**, type only the computer name of the Web server. For example, type **ws** for the FQDN ws.treyresearch.net.
4. In **IP address**, type the IP address for the new Web server, for example, 131.107.27.67.
5. Click **Add Host**.

**See Also**
**Concepts**
Checklist: Installing an ADFS-enabled Web server
Name resolution requirements for ADFS-enabled Web servers

# Export the private key portion of a server authentication certificate

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Every federation server in an Active Directory Federation Services (ADFS) server farm must have access to the private key of the server authentication certificate. If you are implementing a server farm of federation servers or ADFS-enabled Web servers, you must have a single authentication certificate. This certificate must be issued by an enterprise certification authority (CA), and it must have an exportable private key. The private key of the server authentication certificate must be exportable so that it can be made available to all of the servers in the farm.

**To export the private key of a server authentication certificate**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, double-click **ComputerName**, double-click **Web Sites**, right-click **Default Web Site**, and then click **Properties**.
3. On the **Directory Security** tab, click **View Certificate**, click the **Details** tab, and then click **Copy to File**.
4. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
5. On the **Export Private Key** page, select **Yes, export the private key**, and then click **Next**.
6. On the **Export File Format** page, select **Personal Information Exchange = PKCS #12 (.PFX)**, and then click **Next**.
7. On the **Password** page, type and confirm the password that is required to share the server authentication certificate. You will need this password when you import the exported server authentication certificate into the certificate store of the new federation server.
8. On the **File to Export** page, specify the certificate file, and then click **Next**.
9. On the **Completing the Certificate Export Wizard** page, click **Finish**.
10. Validate the success of your export by confirming that the file you specified is created at the specified location.

| **Important** |
| --- |
| So that this certificate can be imported to the local certificate store on the new server, you must transfer the file to physical media and protect its security during transport to the new server. It is extremely important to guard the security of the private key. If this key is compromised, the security of your entire ADFS deployment (including resources within the organization and resource partner organizations) is compromised. |

11. Import the exported server authentication certificate into the certificate store on the new server before you install the Federation Service. For information about how to import the certificate, see Import a certificate (http://go.microsoft.com/fwlink/?linkid=20040).

# Import a server authentication certificate to the default Web site

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

After you obtain a server authentication certificate from a certification authority (CA), you must manually install that certificate on the default Web site for each federation server or federation server proxy in a server farm.

For ADFS-enabled Web servers, you must manually install the server authentication certificate on the appropriate Web site or virtual directory where your federated application resides.

If you are setting up a farm, be sure to perform this procedure identically (using the exact same settings) on each of the servers in your farm.

**To import a server authentication certificate to the default Web site**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, double-click **Web Sites**, right-click **Default Web Site**, and then click **Properties**.
3. On the **Directory Security** tab, click **Server Certificate**.
4. On the **Welcome to the Web Server Certificate Wizard** page, click **Next**.
5. On the **Server Certificate** page, click **Import a certificate from a .pfx file**, and then click **Next**.
6. On the **Import Certificate** page, click **Browse**, navigate to the location of the .pfx file that you obtained from your CA, click **Open**, and then click **Next**.
7. On the **Import Certificate Password** page, in **Password**, type the password that you entered when you requested the certificate from your CA, and then click **Next**.
8. On the **SSL Port** page, in the **SSL port this web site should use** box, type **443**, and then click **Next**.

   | **Note** |
   | --- |
   | If you chose an alternate port number on one server, be sure to use that same port number on all the other servers in the farm. |

9. On the **Imported Certificate Summary** page, click **Next**.
10. On the **Completing the Web Server Certificate Wizard** page, click **Finish**.

**See Also**

**Concepts**

Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server
Certificate requirements for federation servers
Certificate requirements for federation server proxies
Certificate requirements for ADFS-enabled Web servers

# Install prerequisite applications

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) requires that the following applications be installed on a computer running Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Standard Edition, (only for servers hosting the ADFS Web Agent) so that the computer can host the Federation Service, Federation Service Proxy, or ADFS Web Agent components of ADFS:

- Internet Information Services (IIS)
- Microsoft .NET Framework 2.0
- Microsoft ASP.NET (required for claims-aware applications only)

When you add the Application Server Windows component, these three applications are installed.

**Administrative credentials**

To complete this procedure, you must be a member of the local Administrators group.

**To add the Application Server Windows component**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard**, select the **Application Server** check box, and then click **Next**.
4. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server

# Install the ADFS Web Agent component of ADFS

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can add an Active Directory Federation Services (ADFS) Web Agent to a Web server in a resource partner organization. You use the same procedure to add either the ADFS Web Agent for Windows NT token–based applications or the ADFS Web Agent for claims-aware applications. This action makes that server an ADFS-enabled Web server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add an ADFS Web Agent to a Web server**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard**, select the **Active Directory Services** check box, and then click **Details**.
4. In the **Active Directory Services** dialog box, select the **Active Directory Federation Services (ADFS)** check box, and then click **Details**.
5. In the **Active Directory Federation Services (ADFS)** dialog box, click **ADFS Web Agents**, and then click **Details**.
6. In the **ADFS Web Agents** dialog box, select Web Agent subcomponents as follows, and then click **OK**:
   - To add ADFS Web Agent support for applications that can process ADFS claims directly for authorization, select the **Claims-aware applications** check box.

     | **Note** |
     | --- |
     | If your Web server will host only claims-aware applications, do not select the **Windows NT token–based applications** check box. |

   - To add ADFS Web Agent support for applications that require Windows native authorization mechanisms (that is, for applications that do not directly consume ADFS claims), select the **Windows NT token–based applications** check box.

     | **Note** |
     | --- |
     | If your Web server will host only Windows NT token–based applications, do not select the **Claims-aware applications** check box. |

7. In the **Active Directory Federation Services (ADFS)** dialog box, click **OK**.
8. In the **Active Directory Services** dialog box, click **OK**.
9. In the Windows Components Wizard, click **Next**.
10. If you are prompted for the location of installation files, navigate to *R2 installation files*\cmpnents\r2, and then click **OK**.
11. On the **Completing the Windows Components Wizard** page, click **Finish**.

# Verify that an ADFS-enabled Web server is operational

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

After you set up your Active Directory Federation Services (ADFS)–enabled Web server and you successfully install and configure the applications, you can use one or more of the procedures in this topic to verify that the ADFS-enabled Web server can be reached by a federation server, by a client on the Internet, or by local clients through Windows Integrated authentication.

Depending on whether you currently have federation servers deployed or you want to verify local connectivity, perform one or more of the following tasks:

- If you have a resource federation server deployed in your organization, verify that the ADFS-enabled Web server and the resource federation server can ping one another using their fully qualified domain names (FQDNs) and IP addresses. If the **ping** command fails, use **nslookup** to test Domain Name System (DNS) connectivity.

  For more information about troubleshooting connectivity between ADFS-enabled Web servers and federation servers, see Verify Active Directory Federation Services Computer Settings and Connectivity (http://go.microsoft.com/fwlink/?LinkId=74929).

- Verify that you can access the application with ADFS disabled. Perform the steps in the following procedure when you want to verify basic Windows Integrated authentication connectivity to the application. This procedure can be helpful when you want to test local connectivity to the ADFS-enabled Web server by using a client computer in the same Active Directory forest or a trusting forest. This procedure can also be helpful when you want to verify that the application has been installed correctly even though federation servers have not yet been deployed.

**To verify local access to an application**

1. If you previously enabled Anonymous Access for the Web site or virtual directory where your application resides, temporarily disable Anonymous Access and make sure that Windows Integrated authentication is enabled for the purposes of this verification test.
2. If you previously enabled the ADFS Web Agent by using the Internet Information Services (IIS) Manager UI (for Windows NT token-based agents) or by using the web.config file (for claims-aware agents), temporarily disable the ADFS Web Agent for the purposes of this verification test.
3. Log on to a client computer that is a member of the same forest or trusting forest as the ADFS-enabled Web server.
4. Open a browser window, type the return URL for the federated application that you will attempt to access (for example, https://adfsweb.treyresearch.net/ordering), and then press ENTER.
5. If the test is successful, you should be able to access the application.

- Verify that you can access the application with ADFS enabled. Perform the steps in the following procedure when you want to verify that ADFS components are working as expected and that IIS is publishing the application correctly. The steps that you perform here assume you have the appropriate permissions assigned so that federated users can access your application. These steps are intended only to verify connectivity to the ADFS-enabled Web server.

| **Note** |
| --- |
| If you previously enabled Windows Integrated authentication (for the previous procedure), disable Windows Integrated authentication and then re-enable Anonymous Access again before continuing with this procedure. |

**To verify federated access to an application**

1. Log on to a client computer with Internet access.
2. Open a browser window, type the return URL for the federated application that you will attempt to access (for example, https://adfsweb.treyresearch.net/test), and then press ENTER.

If the Web server is configured properly, a client computer that does not yet trust the server authentication certificates should see a prompt for certificates and then a prompt for the account partner discovery page.

| **Note** |
| --- |
| If the resource Federation Service has only one account partner and no account stores are identified in the trust policy, the client computer will not be prompted fro the account partner discover page. Instead the client will be redirected to the account partner's login page. |

If you do not see either of these prompts, try double-checking that all of the tasks in the checklist for the federation server are all complete, and then try again. If you still do not see these prompts, see Configuring ADFS Servers for Troubleshooting (http://go.microsoft.com/fwlink/?LinkId=74970).

For general information about how to troubleshoot problems with Secure Sockets Layer (SSL)–enabled Web sites, including identifying configuration problems in the IIS metabase, certificates, or certificate stores, see Internet Information Services Diagnostic Tools (http://go.microsoft.com/fwlink/?LinkId=55062).

**See Also**
**Concepts**
Checklist: Installing an ADFS-enabled Web server

# Deploying Federation Servers

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

To deploy federation servers, complete each of the tasks in Checklist: Installing a federation server.

| Note |
| --- |
| When you use this checklist, we strongly recommend that you first read the references to federation server planning guidance (in the ADFS Design Guide) before continuing to the procedures for configuring the servers. Following the checklist in this way will help provide a better understanding of the full design and deployment story for federation servers. |

**About federation servers**

Federation servers are computers running Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Datacenter Edition, that are configured to host the Federation Service component of Active Directory Federation Services (ADFS). Federation servers 'authenticate' or route requests from user accounts in other organizations and from clients that can be located anywhere on the Internet.

The act of installing the Federation Service on a computer makes that computer a federation server. It also makes the Active Directory Federation Services snap-in available on that computer on the **Administrative Tools** menu so that you can specify the following:

- The Federation Service endpoint Uniform Resource Locator (URL) value where partner organizations and applications will send token requests and responses
- The Federation Service Uniform Resource Identifier (URI) value that partner organizations and applications will use to identify the unique name or location of your organization
- The location of the trust policy file that all federation servers that participate in the same server farm will use
- The token-signing certificate that all federation servers in a server farm will use to issue and sign tokens
- The location of customized ASP.NET Web pages for client logon, logoff, and account partner discovery that will enhance the client experience

| Note |
| --- |
| The majority of these core user interface (UI) settings are contained in the web.config file on each federation server. The Federation Service endpoint URL and Federation Service URI values are not specified in the web.config file. |

Federation servers host a security token service that issues tokens that are based on the credentials (for example, user name and password) that are presented to it. A security token is a cryptographically signed data unit that expresses one or more claims. A claim is a statement that a server makes (for example, name, identity, key, group, privilege, or capability) about a client. After the credentials are verified on the federation server (through the user logon process), claims for the user are collected through examination of the user attributes that are stored in Active Directory or Active Directory Application Mode (ADAM).

In Federated Web Single Sign-On (SSO) designs (designs in which two or more organizations are involved), claims can be modified by claim mappings for a specific resource partner. The claims are built into a token that is sent to a federation server in the resource partner organization. After a federation server in the resource partner receives the claims as incoming claims, it maps them into its organization claims. The organization claims are then built into a new token that is sent to the Web server in the resource partner that hosts the ADFS Web Agent.

In the Web SSO design (where only one organization is involved), a single federation server can be used so that employees can log on once and still access multiple applications.

# Checklist: Installing a federation server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This checklist includes the deployment tasks necessary to prepare a server running Windows Server 2003 R2, Enterprise Edition, for the Active Directory Federation Services (ADFS) federation server role.

| Note |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist. |

### Checklist: Installing a federation server

| | Task | Reference |
| --- | --- | --- |
| ☐ | Review information in the Active Directory Federation Services Design Guide about where to place federation servers in your organization | Planning Federation Server Placement<br>Where to place a federation server |
| ☐ | Use the information in the Active Directory Federation Services Design Guide to determine whether a single federation server or federation server farm is preferred for your deployment. | When to create a federation server<br>When to create a federation server farm |
| ☐ | Use the information in the Active Directory Federation Services Design Guide to determine whether this new federation server will be created in the account partner organization or the resource partner organization. | Review the role of the federation server in the account partner organization<br>Review the role of the federation server in the resource partner organization |
| ☐ | Review information in the Active Directory Federation Services Design Guide about how federation servers use server authentication certificates and token-signing certificates to securely authenticate client and federation server proxy requests. | Certificate requirements for federation servers |
| ☐ | Review information in the Active Directory Federation Services Design Guide about how to update the corporate network Domain Name System (DNS) so that successful name resolution to federation servers can occur. | Name resolution requirements for federation servers |
| ☐ | Join the computer that will become the federation server to a domain in the account partner forest or resource partner forest where it will be used to authenticate the users of that forest or from trusting forests.<br>**Note**<br>If you want to create a federation server in the account partner organization, the computer must first be joined to any domain in the forest where your federation server will be used to authenticate users from that forest or from trusting forests. | Join a computer to a domain |

| | | |
|---|---|---|
| ☐ | Create a new resource record in the corporate network DNS that points the DNS host name of the federation server to the IP address of the federation server. | Add a host (A) record to corporate DNS for a federation server |
| ☐ | Install prerequisite applications such as ASP.NET, Internet Information Services (IIS) and Microsoft .NET Framework 2.0 on the computer that will become the federation server. | Install prerequisite applications |
| ☐ | Secure IIS using a server authentication certificate, and configure ADFS with a token-signing certificate. | Checklist: Configuring certificates for a federation server |
| ☐ | Install the Federation Service component on the computer that will become the federation server. Follow this procedure when you want either to create the first federation server in a new farm or to extend an existing farm.<br>**Note**<br>For the Federated Web Single Sign-On (SSO) and Federated Web SSO with Forest Trust scenarios, you must have at least one federation server in the account partner organization and at least one federation server in the resource partner organization. | Install the Federation Service component of ADFS |
| ☐ | If this is the first federation server in your organization, configure the trust policy so that it conforms to your ADFS design. | Checklist: Configuring the account partner organization<br>Checklist: Configuring the resource partner organization |
| ☐ | From a client computer, verify that the federation server is operational. | Verify that a federation server is operational |

# Checklist: Configuring certificates for a federation server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This checklist includes the deployment tasks for configuring certificates on a federation server running Windows Server 2003 R2, Enterprise Edition.

| Note |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist. |

### Checklist: Configuring certificates for a federation server

| | Task | Reference |
| --- | --- | --- |
| ☐ | Before you install the Federation Service component on a computer that will become a federation server, read about the importance of obtaining and (for federation server farms) sharing a server authentication certificate and token-signing certificate across all the servers in the farm. | Certificate requirements for federation servers |
| ☐ | (Optional) As an alternative to obtaining a server authentication certificate from a certification authority (CA), you can use the SelfSSL.exe tool to acquire a sample certificate for your federation server.<br><br>Because the SelfSSL tool generates a self-signed certificate that does not originate from a commonly trusted source, use the SelfSSL tool only in the following scenarios:<br><br>• When you have to create a Secure Sockets Layer (SSL) channel between your server and a limited, known group of users<br><br>• When you have to troubleshoot third-party certificate problems<br><br>**Caution**<br>It is not a security best practice to deploy a federation server in a production environment using a self-signed server authentication certificate. | Internet Information Services (IIS) 6.0 Resource Kit Tools (http://go.microsoft.com/fwlink/?LinkId=36285) |
| ☐ | (Optional) As an alternative to obtaining a token-signing certificate from a CA, you can use the Windows Components Wizard (during the installation of the Federation Service component) to create a self-signed token-signing certificate automatically, or you can use the MakeCert.exe tool to acquire this certificate for your federation server.<br><br>The MakeCert tool generates X.509 root certificates. It is typically used for testing purposes.<br><br>**Caution** It is not a security best practice to deploy a federation server in a production environment using a self-signed token-signing certificate. | Create a self-signed, token-signing certificate |
| ☐ | (Optional) If you will be adding a federation server to a federation server farm, you might have to first export the private key of the existing token-signing certificate (on the first federation server in the farm) and then import it into the personal store of the local federation server computer.<br><br>Exporting the private key is not required when your issued token-signing certificate can be reused by multiple computers (without the need to export) or when you will obtain unique token-signing certificates for each federation server in the farm. | Export the private key portion of a token-signing certificate (http://go.microsoft.com/fwlink/?LinkId=75068)<br><br>Import a certificate (http://go.microsoft.com/fwlink/?linkid=20040) |

| | | |
|---|---|---|
| ☐ | (Optional) If you will be adding a federation server to a federation server farm, you might have to first export the private key of the existing server authentication certificate (on the first federation server in the farm) so that you have a file format of the certificate ready when other federation servers must import the same certificate.<br><br>Exporting the private key is not required when your issued server authentication certificate can be reused by multiple computers (without the need to export) or when you will be obtaining unique server authentication certificates for each federation server in the farm. | Export the private key portion of a server authentication certificate |
| ☐ | After you obtain a server authentication certificate (or private key), you must then import the certificate file to the default Web site for each federation server. | Import a server authentication certificate to the default Web site |
| ☐ | Go back to the main federation server checklist, and proceed to the next task (Install the Federation Service component of ADFS). | Checklist: Installing a federation server |

# Join a computer to a domain

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

For Active Directory Federation Services (ADFS) to function, each computer that functions as a federation server must be joined to a domain. Federation server proxies may be joined to a domain, but this is not a requirement.

ADFS-enabled Web servers must be joined to any domain in the same forest where the resource federation server resides whenever the Web servers host a Windows NT token–based application. This is a requirement that makes it possible for the Windows NT token-based Web agent to properly retrieve group membership information from Active Directory.

You do not have to join an ADFS-enabled Web server to a domain if the Web server is hosting claims-aware applications only.

**To join a computer to a domain**

1. On the computer that you want to join to a domain, click **Start**, click **Control Panel**, and then double-click **System**.
2. On the **Computer Name** tab, click **Change**.
3. Under **Member of**, click **Domain**, type the name of the domain that this computer will join, and then click **OK**.
4. Click **OK**, and then restart the computer.

**See Also**
**Concepts**
Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server

# Add a host (A) record to corporate DNS for a federation server

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

For clients on the corporate network to successfully access a federation server using Windows Integrated authentication, an address (A) resource record must first be created in the corporate Domain Name System (DNS) that resolves the host name of the account federation server (for example, fs.corp.adatum.com) to the IP address of the federation server or federation server cluster. You can use the following procedure to add a host (A) record to corporate DNS for a federation server.

**To add a host (A) record to corporate DNS for a federation server**

1. On a DNS server for the corporate network, open the DNS snap-in.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A)**.
3. In **Name**, type only the computer name of the federation server or federation server cluster (for example, type **fs** for the fully qualified domain name (FQDN) fs.adatum.com).
4. In **IP address**, type the IP address for the federation server or federation server cluster (for example, 192.168.1.4).
5. Click **Add Host**.

**See Also**
**Concepts**
Checklist: Installing a federation server
Name resolution requirements for federation servers

# Export the private key portion of a server authentication certificate

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Every federation server in an Active Directory Federation Services (ADFS) server farm must have access to the private key of the server authentication certificate. If you are implementing a server farm of federation servers or ADFS-enabled Web servers, you must have a single authentication certificate. This certificate must be issued by an enterprise certification authority (CA), and it must have an exportable private key. The private key of the server authentication certificate must be exportable so that it can be made available to all of the servers in the farm.

**To export the private key of a server authentication certificate**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, double-click **ComputerName**, double-click **Web Sites**, right-click **Default Web Site**, and then click **Properties**.
3. On the **Directory Security** tab, click **View Certificate**, click the **Details** tab, and then click **Copy to File**.
4. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
5. On the **Export Private Key** page, select **Yes, export the private key**, and then click **Next**.
6. On the **Export File Format** page, select **Personal Information Exchange = PKCS #12 (.PFX)**, and then click **Next**.
7. On the **Password** page, type and confirm the password that is required to share the server authentication certificate. You will need this password when you import the exported server authentication certificate into the certificate store of the new federation server.
8. On the **File to Export** page, specify the certificate file, and then click **Next**.
9. On the **Completing the Certificate Export Wizard** page, click **Finish**.
10. Validate the success of your export by confirming that the file you specified is created at the specified location.

| **Important** |
| --- |
| So that this certificate can be imported to the local certificate store on the new server, you must transfer the file to physical media and protect its security during transport to the new server. It is extremely important to guard the security of the private key. If this key is compromised, the security of your entire ADFS deployment (including resources within the organization and resource partner organizations) is compromised. |

11. Import the exported server authentication certificate into the certificate store on the new server before you install the Federation Service. For information about how to import the certificate, see Import a certificate (http://go.microsoft.com/fwlink/?linkid=20040).

# Import a server authentication certificate to the default Web site

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

After you obtain a server authentication certificate from a certification authority (CA), you must manually install that certificate on the default Web site for each federation server or federation server proxy in a server farm.

For ADFS-enabled Web servers, you must manually install the server authentication certificate on the appropriate Web site or virtual directory where your federated application resides.

If you are setting up a farm, be sure to perform this procedure identically (using the exact same settings) on each of the servers in your farm.

**To import a server authentication certificate to the default Web site**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, double-click **Web Sites**, right-click **Default Web Site**, and then click **Properties**.
3. On the **Directory Security** tab, click **Server Certificate**.
4. On the **Welcome to the Web Server Certificate Wizard** page, click **Next**.
5. On the **Server Certificate** page, click **Import a certificate from a .pfx file**, and then click **Next**.
6. On the **Import Certificate** page, click **Browse**, navigate to the location of the .pfx file that you obtained from your CA, click **Open**, and then click **Next**.
7. On the **Import Certificate Password** page, in **Password**, type the password that you entered when you requested the certificate from your CA, and then click **Next**.
8. On the **SSL Port** page, in the **SSL port this web site should use** box, type **443**, and then click **Next**.

   | **Note** |
   | --- |
   | If you chose an alternate port number on one server, be sure to use that same port number on all the other servers in the farm. |

9. On the **Imported Certificate Summary** page, click **Next**.
10. On the **Completing the Web Server Certificate Wizard** page, click **Finish**.

**See Also**

**Concepts**

Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server
Certificate requirements for federation servers
Certificate requirements for federation server proxies
Certificate requirements for ADFS-enabled Web servers

# Create a self-signed, token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use the following procedure to create a self-signed, code-signing certificate that also creates and installs a private key. To perform this procedure, use the Makecert.exe utility. Makecert.exe is available in the Microsoft .NET Framework 2.0 Software Development Kit (SDK) (x86) (http://go.microsoft.com/fwlink/?LinkId=79548).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create a self-signed, token-signing certificate**

1. Open a command prompt.
2. Type the appropriate **makecert** syntax.
   - Example command:

**makecert -r -pe -n "CN=CertForADFS" -b 01/01/2006 -e 01/01/2007 -eku 1.3.6.1.5.5.7.3.3 -ss my -sr localMachine -sky exchange -sp "Microsoft RSA SChannel Cryptographic Provider" -sy 12 "CertForADFS.cer"**

| **Note** |
| --- |
| Track certificate expiration dates to make sure that certificates are replaced before they expire. You can do this using the Active Directory Federation Services snap-in for certificates of the current organization and also for partners that you configure in the trust policy. |

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Rolling Over a Client Authentication Certificate

# Install prerequisite applications

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) requires that the following applications be installed on a computer running Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Standard Edition, (only for servers hosting the ADFS Web Agent) so that the computer can host the Federation Service, Federation Service Proxy, or ADFS Web Agent components of ADFS:

- Internet Information Services (IIS)
- Microsoft .NET Framework 2.0
- Microsoft ASP.NET (required for claims-aware applications only)

When you add the Application Server Windows component, these three applications are installed.

**Administrative credentials**

To complete this procedure, you must be a member of the local Administrators group.

**To add the Application Server Windows component**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard**, select the **Application Server** check box, and then click **Next**.
4. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server

# Install the Federation Service component of ADFS

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

Now that you have properly configured a computer with the prerequisite applications and certificates, you are ready to install the Federation Service component of Active Directory Federation Services (ADFS). When you install the Federation Service on a computer, that computer becomes a federation server.

| **Note** |
| --- |
| For the Federated Web Single-Sign-On (SSO) and Federated Web SSO with Forest Trust scenarios, you must have at least one federation server in the account partner organization and at least one federation server in the resource partner organization. For more information, see Where to place a federation server. |

You can use the following procedure to install the Federation Service component of ADFS on a computer that will become the first federation server or on a computer that will become a federation server for an existing federation server farm.

**Prerequisites**

If you will be using a token-signing certificate that is issued by a certification authority (CA), verify that a token-signing certificate with the private key has already been installed or imported into the local certificate store (Personal store) before you start this procedure. As an alternative, you can create a self-signed, token-signing certificate using the Windows Components Wizard, as described in this procedure. For more information about token-signing certificates, see Certificate requirements for federation servers.

If you are adding this new federation server to an existing federation server farm, make sure that the trust policy file is available on the network with the appropriate permissions before starting this procedure. For more information, see When to create a federation server farm.

**To install the Federation Service component of ADFS**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, click **Active Directory Services**, and then click **Details**.
4. In the **Active Directory Services** dialog box, click **Active Directory Federation Services (ADFS)**, and then click **Details**.
5. In the **Active Directory Federation Services (ADFS)** dialog box, select the **Federation Service** check box, and then click **OK**. If ASP.NET 2.0 was not previously enabled, click **Yes** to enable it, and then click **OK**.
6. In the **Active Directory Services** dialog box, click **OK**.
7. In the **Windows Components Wizard**, click **Next**.
8. On the **Federation Service** page, under **Token-signing certificate**, do one of the following:
   - If you want Setup to create a new, self-signed, token-signing certificate for this federation server and add it to the personal store of the local computer, click **Create a self-signed token-signing certificate**, and then go to step 9.
   - If you want to use an existing token-signing certificate that has already been added to the certificate store on the local computer, click **Select token-signing certificate**, click **Select**, highlight the appropriate certificate in the list, click **OK**, and then go to step 9.

     | **Note** |
     | --- |
     | This is the same certificate name that you entered when you created the token-signing certificate request. |

9. On the **Federation Service** page, under **Trust policy**, do one of the following:

- If you are installing the Federation Service component of ADFS on a single federation server or on the first federation server in a server farm, click **Create a new trust policy**, and then click **Next**.

| **Note** |
| --- |
| Every federation server in a server farm will need to use this trust policy file. Therefore, we recommend that you store this file in a protected network shared folder. For more information, see When to create a federation server farm. |

- If you are installing the Federation Service component of ADFS on a computer that will become an additional federation server for an existing federation server farm, click **Use an existing trust policy**, and then click **Browse**. In the **Browse** dialog box, locate the shared TrustPolicy.xml file on the network, highlight it, click **Open**, and then click **Next**.

10. If you are prompted for the location of the installation files, navigate to *R2 Installation Folder*\cmpnents\r2, and then click **OK**.

11. On the **Completing the Windows Components Wizard** page, click **Finish**.

# Configure event logging on a federation server

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Servers that are running the Federation Service component of Active Directory Federation Services (ADFS) log ADFS Federation Service events in the Application event log. These events report information about the operation of the components of the local organization and the partner organizations that are covered by a trust policy.

| **Note** |
| --- |
| When it is manually configured, ADFS also can log debug information. Debug logs are located in *%systemdrive%*\ADFS\logs. For more information about how to configure debug logging, see Configuring ADFS Servers for Troubleshooting (http://go.microsoft.com/fwlink/?LinkId=74970). |

The following types of events are available and enabled by default in ADFS:

- **Error:** Information about a significant problem of which the user should be aware, usually involving a loss of functionality or data.
- **Warning:** Indicates a problem that is not immediately significant, but that may signify conditions that could cause future issues.
- **Info:** Information about a significant, successful operation.
- **Success audit:** Indicates an audited security event when an audited access attempt is successful; for example, a successful logon attempt.
- **Failure audit:** Indicates a security event that occurs when an audited access attempt fails; for example, an inbound token was not valid.
- **Detailed success:** A success audit event with detailed information about each token involved in the transaction, including claims information.
- **Detailed failure:** A failure audit event with detailed information about each token involved in the transaction, including claims information.

You can select the levels that you want to enable and disable.

| **Note** |
| --- |
| Audit object access must be turned on for success or failure to allow the Federation Service to log errors. For more information, see Audit object access (http://go.microsoft.com/fwlink/?LinkId=79749). |

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

#### To change the event types that are logged by ADFS

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click the **Trust Policy** node, and then click **Properties**.
3. Scroll to the **Event Log** tab.
4. Under **Event log level**, click to select and deselect event log types, and then click **OK**.

# Verify that a federation server is operational

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You can use the following procedure to verify that any client on the same network can reach a new federation server.

**To verify that a federation server is operational**

1. To verify that Internet Information Services (IIS) is configured correctly on the federation server, log on to a client computer that is located in the same forest as the federation server.
2. Open a browser window, and then type the Federation Service Uniform Resource Locator (URL) for the new federation server, for example:

   **https://adfsresource.treyresearch.net/adfs/fs/federationserverservice.asmx**

   | **Note** |
   | --- |
   | The Federation Service URL should include the Domain Name System (DNS) host name of the federation server. |

3. Press ENTER. If you see a prompt for certificates, click **Cancel**.

   The FederationServerService page should be displayed along with a list of links that identify the Web methods that are used by the Federation Service. If this page appears, IIS on the federation server is operational and is serving pages successfully.
4. To verify that the federation server is communicating with the Federation Service, obtain access to the new federation server, and then log on to the server.
5. Click **Start**, point to **Administrative Tools**, and then click **Event Viewer**.
6. In the details pane, double-click **Application**.
7. In the **Event** column, look for event ID 674. If the federation server is configured properly, you see a new event, in the Application log of Event Viewer, with the event ID 674. This event verifies that the federation server was able to successfully communicate with the Federation Service.

   | **Note** |
   | --- |
   | Although the text displayed in event ID 674 mentions that the Federation Service Proxy has updated its configuration information, this event also applies to federation servers as well. |

   If you cannot generate event ID 674, see Configuring ADFS Servers for Troubleshooting (http://go.microsoft.com/fwlink/?LinkId=75106).

**See Also**
**Concepts**
Checklist: Installing a federation server

# Deploying Federation Server Proxies

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

To deploy federation server proxies, complete each of the tasks in Checklist: Installing a federation server proxy.

| Note |
| --- |
| When you use this checklist, we strongly recommend that you first read the references to federation server proxy planning guidance (in the ADFS Design Guide) before continuing to the procedures for configuring the servers. Following the checklist in this way will help provide a better understanding of the full design and deployment story for federation server proxies. |

**About federation server proxies**

Federation server proxies are computers running Windows Server 2003 R2 that are configured to host the Federation Service Proxy component of Active Directory Federation Services (ADFS). You can use federation server proxies in your organization to provide intermediary services between an Internet client and a federation server that is behind a firewall on the corporate network.

| Note |
| --- |
| Although the Federation Service and the Federation Service Proxy components cannot be installed on the same computer, a federation server can perform federation server proxy functionality. For more information, see When to create a federation server. |

The act of installing the Federation Service Proxy on a computer makes that computer a federation server proxy. It also makes the Active Directory Federation Services snap-in available on that computer on the **Administrative Tools** menu so that you can specify the following:

- Which Federation Service the federation server proxy should route authentication requests to. You specify this location by typing the path for the Federation Service Uniform Resource Locator (URL).

- The client authentication certificate that makes it possible for the federation server proxy to authenticate to the Federation Service. You can configure which client authentication certificates will be accepted by the Federation Service on the **FSP Certificates** tab in the properties of the Trust Policy Folder.

  After you specify the Federation Service URL and the correct client authentication certificate, the federation server proxy is bound to that Federation Service. As part of this binding, a federation server proxy is trusted by the Federation Service so that the federation server can hand out certain privileged security tokens that originate from the Federation Service.

- Locations of customized ASP.NET Web pages for client logon, logoff, and account partner discovery that enable interaction with a client. These pages provide a convenient way to differentiate the sign-in experience for users in your intranet as opposed to users on the Internet.

  | Note |
  | --- |
  | You can also configure these settings in the Web.config file on each federation server proxy. |

Federation server proxies also send cookies to external clients when necessary to facilitate the single-sign-on (SSO) process. These cookies include authentication cookies, account partner cookies, and sign-out cookies.

A federation server proxy implements the sign-on and sign-out messages that are described in the WS-Federation Passive Requestor Profile (WS-F PRP) protocol specification.

# Checklist: Installing a federation server proxy

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

This checklist includes the deployment tasks for preparing a server running Windows Server 2003 R2, Enterprise Edition, for the federation server proxy role.

| **Note** |
| --- |
| Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist. |

| | Task | Reference |
| --- | --- | --- |
| ☐ | Review information in the Active Directory Federation Services Design Guide about where to place federation server proxies in your organization. | Planning Federation Server Proxy Placement<br>Where to place a federation server proxy |
| ☐ | Use the information in the Active Directory Federation Services Design Guide to determine whether a single federation server proxy or federation server proxy farm is necessary.<br>**Note**<br>Federation servers also perform federation service proxy responsibilities. | When to create a federation server proxy<br>When to create a federation server proxy farm |
| ☐ | Use the information in the Active Directory Federation Services Design Guide to determine whether this new federation server proxy will be created in the perimeter network of the account partner organization or the resource partner organization. | Review the role of the federation server proxy in the account partner organization<br>Review the role of the federation server proxy in the resource partner organization |
| ☐ | Before you install the Federation Service Proxy component on a computer that will become a federation server proxy, read about the importance of obtaining a server authentication certificate and a client authentication certificate and—for federation server proxy farms—adding or sharing certificates across all the servers in a farm. | Certificate requirements for federation server proxies |
| ☐ | Review information in the Active Directory Federation Services Design Guide about how to update the perimeter Domain Name System (DNS) so that successful name resolution for federation servers and federation server proxies can occur. | Name resolution requirements for federation server proxies |
| ☐ | Determine if the federation server proxy must be joined to a domain. Although federation server proxies do not have to be joined to a domain, they are easier to manage with remote administration and Group Policy features when they are joined to a domain. | Join a computer to a domain |

| | | |
|---|---|---|
| ☐ | Depending on how the DNS infrastructure in your perimeter network is configured, complete one of the procedures on the right before you deploy a federation server proxy in your organization.<br><br>**Note**<br>Do not perform both procedures. Read Name resolution requirements for federation server proxies to determine which procedure best suits the requirements of your organization. | Configure name resolution for a federation server proxy in a DNS zone serving only the perimeter network<br><br>Configure name resolution for a federation server proxy in a DNS zone serving both the perimeter network and Internet clients |
| ☐ | (Optional) If you will be adding a federation server proxy to a federation server proxy farm, you may have to first export the private key of the existing server authentication certificate (on the first federation server proxy in the farm) so that you have a file of the certificate ready when other federation server proxies have to import the same certificate.<br><br>This task is not required in scenarios in which your issued server authentication certificate can be reused by multiple computers (without the need to export) or you obtain unique server authentication certificates for each federation server proxy in the farm. | Export the private key portion of a server authentication certificate |
| ☐ | You must install the public key portion of the Federation Service Proxy client authentication certificate on the federation server so that the federation server can authenticate the federation server proxy. Use the following procedure to export the public key portion of the Federation Service Proxy client authentication certificate. | Export the public key portion of a client authentication certificate |
| ☐ | After you obtain a server authentication certificate, you must install it in Internet Information Services (IIS) on the default Web site of the federation server proxy. | Import a server authentication certificate to the default Web site |
| ☐ | After you export the certificate of the federation server proxy, you can use the following procedure to import the certificate into the trust policy of the Federation Service that the proxy will be servicing. | Add a Federation Service Proxy certificate to the trust policy |
| ☐ | (Optional) As an alternative to obtaining a server authentication certificate from a certification authority (CA), you can use SelfSSL.exe to acquire a sample certificate for your federation server proxy.<br><br>Because the SelfSSL tool generates a self-signed certificate that does not originate from a trusted source, use the SelfSSL tool only in the following scenarios:<br><br>• When you have to create a Secure Sockets Layer (SSL) channel between your server and a limited, known group of users<br><br>• When you have to troubleshoot third-party certificate problems<br><br>**Caution**<br>It is not a security best practice to deploy a federation server proxy in a production environment using a self-signed server authentication certificate. | Internet Information Services (IIS) 6.0 Resource Kit Tools (http://go.microsoft.com/fwlink/?LinkId=36285) |
| ☐ | Install prerequisite applications such as ASP.NET, IIS, and Microsoft .NET Framework 2.0 on the computer that will become the federation server proxy. | Install prerequisite applications |
| ☐ | Install the Federation Service Proxy component on the computer that will become the federation server proxy. | Install the Federation Service Proxy component of ADFS |
| ☐ | To ensure successfully tracking of issues that may occur with this federation server proxy, configure event logging. | Configure event logging on a federation server proxy |

| | | |
|---|---|---|
| ☐ | From a client computer, verify that the federation server proxy is operational. | Verify that a federation server proxy is operational |

# Join a computer to a domain

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

For Active Directory Federation Services (ADFS) to function, each computer that functions as a federation server must be joined to a domain. Federation server proxies may be joined to a domain, but this is not a requirement.

ADFS-enabled Web servers must be joined to any domain in the same forest where the resource federation server resides whenever the Web servers host a Windows NT token–based application. This is a requirement that makes it possible for the Windows NT token-based Web agent to properly retrieve group membership information from Active Directory.

You do not have to join an ADFS-enabled Web server to a domain if the Web server is hosting claims-aware applications only.

**To join a computer to a domain**

1. On the computer that you want to join to a domain, click **Start**, click **Control Panel**, and then double-click **System**.
2. On the **Computer Name** tab, click **Change**.
3. Under **Member of**, click **Domain**, type the name of the domain that this computer will join, and then click **OK**.
4. Click **OK**, and then restart the computer.

**See Also**
**Concepts**
Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server

# Configure name resolution for a federation server proxy in a DNS zone serving only the perimeter network

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

So that name resolution can work successfully for a federation server proxy in a scenario in which one or more Domain Name System (DNS) zones serve only the perimeter network, the following tasks must be completed:

- The hosts file on the federation server proxy must be updated to add the IP address of a federation server.
- DNS in the perimeter network must be configured to resolve all client requests for the Federation Service endpoint Uniform Resource Locator (URL) to the federation server proxy. To accomplish this, you add a host address (A) resource record to perimeter DNS for the federation server proxy.

| **Note** |
| --- |
| These procedures assume that a host (A) record for the federation server has already been created in the corporate network DNS. If this record does not yet exist, create this record, and then perform these procedures. For more information about how to create the host (A) record for the federation server, see Add a host (A) record to corporate DNS for a federation server. |

**Adding the IP address of a federation server to the hosts file**

So that a federation server proxy can work as expected in the perimeter network of an account partner, you must add an entry to the hosts file on that federation server proxy that points to a federation server's DNS host name (for example, fs.corp.adatum.com) and IP address (for example, 192.168.1.4) in the corporate network of the account partner. Adding this entry to the hosts file prevents the federation server proxy from contacting itself to resolve a client-initiated call to a federation server in the account partner.

**To add the IP address of a federation server to the hosts file**

1. Navigate to the %systemroot%\Winnt\System32\Drivers directory folder and locate the **hosts** file.
2. Start Notepad, and then open the **hosts** file.
3. Add the IP address and the host name of a federation server in the account partner to the **hosts** file, as shown in the following example:

   **192.168.1.4 fs.adatum.com**

4. Save and close the file.

**Adding a host (A) record to perimeter DNS for a federation server proxy**

So that clients on the Internet can successfully access a federation server through a newly deployed federation server proxy, you must first create a host (A) resource record in the perimeter DNS. This resource record resolves the host name of the account federation server (for example, fs.corp.adatum.com) to the IP address of the account federation server proxy (for example, 131.107.27.68) in the perimeter network.

| **Note** |
| --- |
| It is assumed that you are using a DNS server running Windows Server 2000 or Windows Server 2003 with the DNS Server service to control the perimeter DNS zone. |

**To add a host (A) record to perimeter DNS for a federation server proxy**

1. On a DNS server for the perimeter network, open the DNS snap-in.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A)**.
3. In **Name**, type only the computer name of the federation server. For example, type fs for the fully qualified domain name

(FQDN) fs.adatum.com.

4. In **IP address**, type the IP address for the new federation server proxy (for example, 131.107.27.68).

5. Click **Add Host**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Name resolution requirements for federation server proxies

# Configure name resolution for a federation server proxy in a DNS zone serving both the perimeter network and Internet clients

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

So that name resolution can work successfully for a federation server proxy in a scenario in which one or more Domain Name System (DNS) zones serve both the perimeter network and Internet clients, the following tasks must be completed:

- DNS in the Internet zone that you control must be configured to resolve all Internet client requests for the Federation Service endpoint Uniform Resource Locator (URL) to the federation server proxy. To accomplish this, you add a host (A) record to the Internet DNS zone for the federation server proxy.

- DNS in the perimeter network must be configured to resolve all incoming client requests for the Federation Service endpoint URL to the federation server. To accomplish this, you add a host (A) record to the perimeter DNS zone for the federation server proxy.

| **Note** |
| --- |
| These procedures assume that a host address (A) resource record for the federation server has already been created in the corporate network DNS. If this record does not yet exist, create this record and then perform these procedures. For more information about how to create a host (A) record for the federation server, see Add a host (A) record to corporate DNS for a federation server. |

**Adding a host (A) record to the Internet DNS zone for a federation server proxy**

So that clients on the Internet can successfully access a federation server through a newly deployed federation server proxy, you must first create a host (A) resource record in the Internet DNS zone that you control. This resource record resolves the host name of the account federation server (for example, fs.adatum.com) to the IP address of the account federation server proxy (for example, 131.107.27.68) in the perimeter network.

| **Note** |
| --- |
| It is assumed that you are using a DNS server running Windows Server 2000 or Windows Server 2003 with the DNS Server service to control the Internet DNS zone. |

**To add a host (A) record to the Internet DNS zone for a federation server proxy**

1. On a DNS server for the Internet DNS zone, open the DNS snap-in.

2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A)**.

3. In **Name**, type only the computer name of the federation server. For example, type fs for the fully qualified domain name (FQDN) fs.adatum.com.

4. In **IP address**, type the IP address for the new federation server proxy (for example, 131.107.27.68).

5. Click **Add Host**.

**Adding a host (A) record to the perimeter DNS zone for a federation server proxy**

So that Internet client requests can be successfully processed by the federation server proxy and reach the federation server after they are resolved by the Internet DNS zone, you must create a host (A) resource record in the perimeter DNS zone. This resource record resolves the host name of the account federation server (for example, fs. adatum.com) to the IP address of the account federation server (for example, 192.168.1.4) in the corporate network.

| **Note** |
| --- |

It is assumed that you are using a DNS server running Windows Server 2000 or Windows Server 2003 with the DNS Server service to control the Perimeter DNS zone.

**To add a host (A) record to the perimeter DNS zone for a federation server proxy**

1. On a DNS server for the perimeter network, open the **DNS snap-in**.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A)**.
3. In **Name**, type only the computer name of the federation server. For example, type **fs** for the FQDN fs.adatum.com.
4. In the **IP address** text box, type the IP address for the federation server in the corporate network (for example, 192.168.1.4).
5. Click **Add Host**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Name resolution requirements for federation server proxies

# Import a server authentication certificate to the default Web site

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

After you obtain a server authentication certificate from a certification authority (CA), you must manually install that certificate on the default Web site for each federation server or federation server proxy in a server farm.

For ADFS-enabled Web servers, you must manually install the server authentication certificate on the appropriate Web site or virtual directory where your federated application resides.

If you are setting up a farm, be sure to perform this procedure identically (using the exact same settings) on each of the servers in your farm.

**To import a server authentication certificate to the default Web site**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, double-click **Web Sites**, right-click **Default Web Site**, and then click **Properties**.
3. On the **Directory Security** tab, click **Server Certificate**.
4. On the **Welcome to the Web Server Certificate Wizard** page, click **Next**.
5. On the **Server Certificate** page, click **Import a certificate from a .pfx file**, and then click **Next**.
6. On the **Import Certificate** page, click **Browse**, navigate to the location of the .pfx file that you obtained from your CA, click **Open**, and then click **Next**.
7. On the **Import Certificate Password** page, in **Password**, type the password that you entered when you requested the certificate from your CA, and then click **Next**.
8. On the **SSL Port** page, in the **SSL port this web site should use** box, type **443**, and then click **Next**.

   | **Note** |
   | --- |
   | If you chose an alternate port number on one server, be sure to use that same port number on all the other servers in the farm. |

9. On the **Imported Certificate Summary** page, click **Next**.
10. On the **Completing the Web Server Certificate Wizard** page, click **Finish**.

**See Also**

**Concepts**

Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server
Certificate requirements for federation servers
Certificate requirements for federation server proxies
Certificate requirements for ADFS-enabled Web servers

# Install prerequisite applications

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) requires that the following applications be installed on a computer running Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Standard Edition, (only for servers hosting the ADFS Web Agent) so that the computer can host the Federation Service, Federation Service Proxy, or ADFS Web Agent components of ADFS:

- Internet Information Services (IIS)
- Microsoft .NET Framework 2.0
- Microsoft ASP.NET (required for claims-aware applications only)

When you add the Application Server Windows component, these three applications are installed.

**Administrative credentials**

To complete this procedure, you must be a member of the local Administrators group.

**To add the Application Server Windows component**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard**, select the **Application Server** check box, and then click **Next**.
4. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server

# Install the Federation Service Proxy component of ADFS

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

After you configure a computer with the prerequisite applications and certificates, you are ready to install the Federation Service Proxy component of Active Directory Federation Services (ADFS). You can use the following procedure to install the Federation Service Proxy component. When you install the Federation Service Proxy component on a computer, that computer becomes a federation server proxy.

**To install the Federation Service Proxy component of ADFS**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, click **Active Directory Services**, and then click **Details**.
4. In the **Active Directory Services** dialog box, click **Active Directory Federation Services (ADFS)**, and then click **Details**.
5. In the **Active Directory Federation Services (ADFS)** dialog box, select the **Federation Service Proxy** check box. If ASP.NET 2.0 was not previously enabled, click **Yes** to enable it, and then click **OK**.
6. In the **Active Directory Services** dialog box, click **OK**.
7. In the Windows Components Wizard, click **Next**.
8. On the **Federation Service Proxy** page, click **Select** to browse for the certificate.
9. In the **Select Certificate** dialog box, click **Client Authentication Certificate for** *CompanyName* **FSP**, and then click **OK**.
10. Under **Federation Service Domain Name System (DNS) host name**, type the host name of your federation server.
11. If you are prompted for the location of the installation files, navigate to *R2 Installation Folder*\cmpnents\r2, and then click **OK**.
12. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy

# Export the public key portion of a client authentication certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The public key portion of a client authentication certificate for the federation server proxy must be added to the trust policy on a federation server so that the Federation Service can authenticate the federation server proxy. By exporting the public key portion of the client authentication certificate, you create a file that can be imported into the trust policy. You can use the following procedure on the federation server proxy computer to export the public key portion of its client authentication certificate.

**To export the public key portion of a client authentication certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service Proxy**, and then click **Properties**.
3. On the **General** tab, under **FSP client authentication certificate**, click **View**.
4. In the **Certificate** dialog box, click the **Details** tab, and then click **Copy to File**.
5. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
6. On the **Export Private Key** page, ensure that **No, do not export the private key** is selected, and then click **Next**.
7. On the **Export File Format** page, ensure that **DER encoded binary X.509 (.CER)** is selected, and then click **Next**.
8. On the **File to Export** page, type or browse to the location and file name that you want to use for the exported certificate, and then click **Next**.
9. On the **Completing the Certificate Export Wizard** page, verify that the information that you provided is accurate, and then click **Finish**.
10. In the **Certificate Export Wizard** dialog box, click **OK**.
11. In the **Certificate** dialog box, click **OK**.
12. In the **Federation Service Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Certificate requirements for federation server proxies

# Add a Federation Service Proxy certificate to the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

So that the Federation Service can authenticate the federation server proxy, the public key portion of the client authentication certificate for the federation server proxy must be added to the trust policy on a federation server with which the federation server proxy communicates. You can use the following procedure to add the client authentication certificate for the federation server proxy from a file that you have exported.

| **Note** |
| --- |
| The Trust Policy user interface (UI) in the Active Directory Federation Services snap-in refers to client authentication certificates for federation service proxies as **Federation Service Proxy (FSP) certificates**. |

The Federation Service Proxy certificate should chain to a trusted root in the Federation Service. Perform this procedure on a federation server that hosts the trust policy to which you want to add a Federation Service Proxy certificate.

**To add a Federation Service Proxy certificate to the trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Trust Policy**, and then click **Properties**.
3. On the **FSP Certificates** tab, click **Add**.
4. In the **Browse for Federation Service Proxy Certificate file** dialog box, navigate to the certificate file that you want to add, select the certificate file, and then click **Open**.
5. In the **Trust Policy Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Certificate requirements for federation server proxies

# Configure event logging on a federation server proxy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Federation servers log Active Directory Federation Services (ADFS) Federation Service events in the Application and Security event logs. On a federation server proxy, events in the Application log contain additional information about errors regarding contact with the Federation Service. In addition, when a federation server proxy is in effect, the Federation Service events contain information about the proxy certificates that are used.

Use the following procedure to specify the level of events that you want to be logged on a server that is running the Federation Service Proxy. Event logging for a federation server proxy is set in the web.config file. By default, this file is located in %systemdrive%\ADFS\sts. You can apply the following logging types in the Web.config file:

- **Error:** Information about a significant problem of which the user should be aware, usually involving a loss of functionality or data.
- **Warning:** Indicates a problem that is not immediately significant, but that may signify conditions that could cause future issues.
- **Info:** Information about a significant, successful operation.
- **SuccessAudit:** Indicates an audited security event that occurs when an audited access attempt is successful, for example, a successful logon attempt.
- **FailureAudit:** Indicates a security event that occurs when an audited access attempt fails; for example, authentication failed.
- **DetailedSuccess:** A success audit event with detailed information about each token involved in the transaction, including claims information.
- **DetailedFailure:** A failure audit event with detailed information about each token involved in the transaction, including claims information.
- **Everything:** Enables all logging levels.

Use the following procedure to configure event logging levels on a federation server proxy. Perform this procedure on the federation server proxy.

| **Note** |
| --- |
| For security events to be registered, object access auditing must be enabled in the security policy. For more information, see Configuring ADFS Servers for Troubleshooting (http://go.microsoft.com/fwlink/?LinkId=79242). |

### Administrative credentials

To complete this procedure, you must have Read-Write access to the web.config file.

**To configure event logging for a federation server proxy using the web.config file**

1. In Notepad, open the web.config file in %systemdrive%\ADFS\sts.
2. Search for **<logonserver>**.
3. Add the **<auditlevel>** XML element under **<logonserver>**, as follows:

   **<auditlevel>** *Value* **</auditlevel>**

   Where *Value* is one of the following, or the combined values of two or more:

   - Error = 1

- Warning = 2
- Info = 4
- SuccessAudit = 16
- FailureAudit = 32
- DetailedSuccess = 64
- DetailedFailure = 128
- Everything = 247

4. Save and close the web.config file.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy

# Verify that a federation server proxy is operational

Updated: February 15, 2007

Applies To: Windows Server 2003 R2

You can use the following procedure to verify that the federation server proxy can communicate with the Federation Service.

| **Important** |
| --- |
| The results of this test are a successful generation of a specific event in Event Viewer on the federation server proxy computer. You will notice that the browser fails with an error in step 2 and does not successfully display the clientlogon.aspx page. This is expected behavior for this verification test. |

**To verify that a federation server proxy is operational**

1. Log on to a client computer with Internet access.
2. Open a browser window, and then type the Uniform Resource Locator (URL) for the Federation Service endpoint, along with the path to the Clientlogon.aspx page that is stored on the federation server proxy, for example:

   **https://adfsresource.treyresearch.net/adfs/ls/clientlogon.aspx**

   | **Note** |
   | --- |
   | At this point your browser should display the error message "Server Error in '/adfs' Application." This step is necessary to generate the appropriate event message to verify that the clientlogon.aspx page is being loaded properly by Internet Information Services (IIS). |

3. Press ENTER.
4. Log on to the federation server proxy.
5. Click **Start**, point to **Administrative Tools**, and then click **Event Viewer**.
6. In the details pane, double-click **Application**.
7. In the **Event** column, look for event ID 674.

   If the federation server proxy is configured properly, you see a new event in the Application log of Event Viewer, with the event ID 674. This event verifies that the federation server proxy was able to successfully communicate with the Federation Service.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy

# Finding Additional ADFS Resources

Updated: December 15, 2006

Applies To: Windows Server 2003 R2

You can find the following documentation about Active Directory Federation Services (ADFS) on the Windows Server 2003 R2 TechCenter Web Site:

- Overview of ADFS (http://go.microsoft.com/fwlink/?LinkId=63491)
- Step-by-Step Guide for Active Directory Federation Services (http://go.microsoft.com/fwlink/?LinkId=49531)
- ADFS Deployment Guide
- ADFS Design Guide
- ADFS Operations Guide (http://go.microsoft.com/fwlink/?LinkId=78683)
- Active Directory Federation Services (Product Help) (http://go.microsoft.com/fwlink/?LinkId=57765)
- Active Directory Federation Services (portal page) (http://go.microsoft.com/fwlink/?LinkId=79542)

For other ADFS-related information, see the following Web resources:

- Automate Information Access with Identity Management (http://go.microsoft.com/fwlink/?LinkId=78692)
- Web Services and Other Distributed Technologies (http://go.microsoft.com/fwlink/?LinkId=44189)
- Web Services Protocol Workshops (http://go.microsoft.com/fwlink/?LinkId=44190)
- Web Services Specifications (http://go.microsoft.com/fwlink/?LinkId=44191)
- Web Services Interoperability Organization (WS-I) (http://go.microsoft.com/fwlink/?LinkId=34328)

# Hardware Management Introduction

Applies To: Windows Server 2003 R2

Hardware Management in Microsoft® Windows® Server 2003 R2 comprises a set of features designed to improve management of server hardware. These features enable system administrators to securely manage remote server hardware across a firewall, using a standard Web service-based protocol. Working with Baseboard Management Controller (BMC) hardware attached to a server that supports WS-Management, the Windows Hardware Management components can communicate with the remote system even if the Windows operating system has not yet booted or has failed.

A BMC is a microcontroller that is attached locally to a server. BMCs include an independent network connection that can communicate over the network even if the server is offline. Through the BMC, system administrators can remotely monitor hardware state and errors and control the hardware in response.

Hardware Management is not installed by default when you install Windows Server 2003 R2. You must enable it from the Management and Monitoring Tools section of the Add/Remove Windows Components wizard. (See Enabling Hardware Management.)

System administrators can manage the Hardware Management components through command-line tools and scripting interfaces described below. No Windows graphical user interface is available to manage Hardware Management features in Windows Server 2003 R2.

The three components of Hardware Management are Intelligent Platform Management Interface (IPMI), Windows Remote Management (WinRM), and the Event Collector.

**Intelligent Platform Management Interface (IPMI)**
Hardware Management includes a Windows Management Instrumentation (WMI) provider and a driver for IPMI, an industry standard for BMC hardware architecture. As with all WMI providers, the functionality of the IPMI provider works both locally and remotely using WMI remoting over Distributed Component Object Model (DCOM), or through firewalls using the WS-Management protocol. The IPMI provider and driver support managing hardware "in-band," that is, when the operating system is running.

The IPMI WMI provider exposes several classes that allow scripts and applications to communicate with BMC hardware through the kernel-level IPMI driver. The classes the provider exposes are:

- AdminDomain
- ComputerSystem
- SystemSpecificCollection
- LogRecord
- Sensor
- RecordLog

For details on the IPMI WMI provider classes, properties and methods, see the WinRM Software Development Kit.

The provider is a Microsoft implementation of a subset of the standard IPMI Common Information Model (CIM) Mapping Specification.

**Windows Remote Management (WinRM)**
Windows Remote Management (WinRM) is the Windows implementation of WS-Management, an industry-standard Web services-based protocol. WinRM provides a secure, efficient way for management applications and scripts to communicate with local and remote computers. The Windows service that WinRM installs and uses is also named WinRM.

When a server is connected to a BMC that supports the WS-Management standard, applications and scripts can use WinRM to communicate directly with the BMC, even when the operating system is offline (pre-boot or post-failure).

When a server is not connected to a BMC, WinRM can still be used to connect to WMI remotely in situations where DCOM communication is impeded (for example, across a firewall). This is possible because the WS-Management standard is firewall-

friendly and uses a single port configurable by the system administrator.

WinRM exposes its own application programming interface (API) for scripting, which can be used by scripts written in any Windows Script Host-compatible language. The scripting API communicates with WMI using syntax different from standard WMI scripts. WinRM syntax is documented in the WinRM Software Development Kit. Hardware Management uses a WMI plug-in to expose WMI classes to WinRM. To call these classes, the WMI namespace and class must be converted into a Uniform Resource Identifier (URI). (See Configuration and Security.)

WS-Management is based on the following standard specifications:

- HTTPS
- SOAP over HTTP (WS-I profile)
- SOAP 1.2
- WS-Addressing
- WS-Transfer
- WS-Enumeration
- WS-Eventing

**WinRM Command Line Tool (Winrm.cmd)**
The command-line tool provided as the primary administrative interface for managing WinRM is a batch file (Winrm.cmd) that runs a Visual Basic Scripting Edition (VBScript) script named Winrm.vbs. Because it is a script, you can open it as a text file and view the code to learn how it works. You can also write your own VBScript scripts that take advantage of the WinRM scripting API. Winrm.vbs runs under Cscript.exe, the command-line scripting engine of Windows Script Host.

Winrm.vbs enables system administrators to configure and manage WinRM. Because WS-Management is a Web service that uses XML as its message format, Winrm.vbs output is natively XML as well. The tool provides switches to output more readable XML or plain text.

The command-line usage (/?) of the tool provides detailed syntax and examples for its use. Windows Remote Management Command-Line Tool (Winrm.cmd) covers background and conceptual information on the tool and also provides extensive examples.

**Event Collector**
The Event Collector service is the third component of Hardware Management in Windows Server 2003 R2. This service is a Windows Remote Management client that is used to create subscriptions to WS-Management event providers and store the events received in the Windows System Event Log (SEL). Two scenarios are supported in R2:

- Logging of hardware events from the local SEL when the operating system is running ("in-band"), using the IPMI driver and WMI provider.
- Logging of hardware events from a remote server's BMC using WS-Management (when BMC hardware supports this).

In case of an operating system shutdown or failure, Windows logs IPMI events into the BMC's SEL. The administrator can then access the BMC's SEL using "out-of-band" tools to determine why a machine is not running.

Event Collector is managed with a command-line tool, Windows Event Collector Utility (Wecutil.exe).

BMC hardware has its own SEL in which it logs significant events. Through an Event Log plug-in and Event Collector, installed as part of Hardware Management, Windows Server 2003 R2 can subscribe to events stored in the BMC's SEL so that they appear in Windows Event Viewer.

With Wecutil.exe, system administrators can create and manage event subscriptions to BMC events. Subscriptions require an XML configuration file: a sample, %windir%\system32\WsmSelRg.xml, ships with Windows Server 2003 R2.

The tool provides extensive command-line usage (/?) with examples. Further information is provided in the topic Event Collector.

**Other Windows Remote Management Topics**
The following topics discuss how to configure and use Hardware Management features, and describe the architecture underlying them.

**Enabling Hardware Management**
Because the Hardware Management feature is not enabled by default in Windows Server 2003 R2, you must explicitly enable it with the Add/Remove Windows Components Wizard. This topic explains how. It also discusses Plug and Play support for the IPMI driver.

**Configuration and Security**
Installation of Hardware Management involves configuring WinRM to work with HTTP or HTTPS. To use HTTPS, a server certificate is required. This topic discusses working with certificates, configuring listeners, and other configuration issues.

**Integration with WMI**
This topic discusses how to use the WMI plug-in of WinRM to access WMI classes through the WS-Management protocol. It describes how to convert a WMI namespace to a WS-Management URI.

**Command-Line Interfaces**
Two new command-line interfaces are installed with the Hardware Management component of Windows Server 2003 R2: Winrm.cmd and Wecutil.exe. These tools are used for WinRM management and event collection respectively.

- Windows Remote Management Command-Line Tool (Winrm.cmd)
- Windows Event Collector Utility (Wecutil.exe)

**Architecture**
This topic discusses the evolution of industry efforts to define a common architecture for managing platform hardware. It then describes Microsoft's implementation, Windows Remote Management, and compares it with previous efforts.

**Additional Resources**

- Windows Remote Management Software Development Kit
- IPMI 1.5
- IPMI CIM Mapping Guideline
- WS-Addressing
- WS-Enumeration
- WS-Eventing
- WS-Transfer
- SOAP 1.2
- BP1
- XML Schema, Part 1
- XML Schema, Part 2
- HTTP
- HTTPS

# Enabling Hardware Management

Applies To: Windows Server 2003 R2

The Hardware Management (including WS-Management) feature is not installed by default in Windows Server 2003 R2. This section highlights the steps involved in getting the feature installed and enabled.

**Installation Steps**

In Windows Server 2003 R2, you can install Hardware Management is available from the **Management and Monitoring Tools** section of the **Add/Remove Windows Components Wizard** (See the following figure):




1. In **Control Panel**, select **Add or Remove Programs**.

2. Select **Add/Remove Windows Components**.

3. Select **Management and Monitoring Tools** and press the **Details** button to bring up the following window.




4. Choose the **Hardware Management** option. If a BMC is detected via the SMBIOS Table Type 38h, the following dialog box will be displayed:




5. If no 3rd party IPMI drivers are installed or they have been removed from the system, press the **OK** button to continue. Otherwise, press the **Cancel** button, remove the 3rd party IPMI driver and reinstall the Hardware Management Component.

| **Note** |
| --- |
| The IPMI driver currently cannot detect if a third-party IPMI driver is installed that does not have well-known device identifiers. |

6. Press the **OK** button to select the **Hardware Management Component**, and then press the **Next** button. Hardware Management (including WinRM) will be installed.

**Plug and Play BMC Support**

The IPMI driver supports ACPI (Plug and Play). To fully support this feature, however, the BMC resources must be exposed via ACPI by the system's BIOS. This capability may be supported in new systems or from a BIOS upgrade. If the BMC is detected via Plug and Play, you will notice an Unknown Device appear in Device Manager before the Hardware Management component is installed.

If the BMC resources are exposed properly via ACPI, the Windows Plug and Play infrastructure will detect the hardware and automatically install the IPMI driver. A new device, named Microsoft Generic IPMI Compliant Device, will appear in Device Manager, as shown below:




**Creating the SMBIOS IPMI Device**

If your system does not automatically detect the BMC via Plug and Play and install the driver, but a BMC was detected during the setup process, the BMC device must be manually created.

To create the IPMI BMC device, execute the following command from a command prompt:

```
Rundll32 ipmisetp.dll, AddTheDevice
```

Once this command is executed, the IPMI device will be created and can be seen under the Device Manager as Microsoft Generic IPMI Compliant Device. The device will be removed when the Hardware Management component is uninstalled.

# Configuration and Security

Applies To: Windows Server 2003 R2

This topic discusses how to set up Windows Remote Management (WinRM) to work over HTTPS, a secure protocol that requires a certificate, and over HTTP. It also discusses how to configure a WinRM Listener, which enables communication with WinRM in various scenarios.

**Configuring WinRM to Use HTTPS**

Windows Remote Management (WinRM) in Windows Server 2003 R2 supports two different transports, HTTPS and HTTP.

| **Note** |
| --- |
| The HTTP protocol is included for troubleshooting purposes only and should not be used in production environments. While WinRM will still authenticate the communication, the data will not be encrypted and is sent as clear text on the network. If IPSec is configured between two machines, then it should be safe to use HTTP in that scenario. |

For HTTPS, a secure protocol, WinRM requires a server certificate to operate properly (this certificate is used for encryption of the communication channel). The easiest way to generate these certificates is to use Microsoft Certificate Services.

| **Note** |
| --- |
| This certificate needs to be marked as a Server Authentication Certificate. It must also support Secure Sockets Layer (SSL). No certificate needs to be configured for the WinRM client. The certificate is used only if the WinRM service is enabled for remote access. |

The table below lists some documents covering Microsoft Certificate Services:

| Title | Document Link |
| --- | --- |
| Trigger Automatic Enrollment for Certificates | http://technet2.microsoft.com/WindowsServer/en/Library/c4e272e9-e12d-45eb-b8f2-26a549fd09e51033.mspx |
| Certificate Services Best practices | http://technet2.microsoft.com/WindowsServer/en/Library/b70185ed-93aa-4346-b869-9913282086af1033.mspx |
| Command-line utility for managing certificates: Certutil.exe | http://www.microsoft.com/resources/documentation/WindowsServ/2003/standard/proddocs/en-us/Default.asp?url=/resources/documentation/windowsserv/2003/standard/proddocs/en-us/sag_cs_certutil2.asp |
| Importing and Exporting Certificates | http://www.microsoft.com/resources/documentation/WindowsServ/2003/standard/proddocs/en-us/Default.asp?url=/resources/documentation/WindowsServ/2003/standard/proddocs/en-us/sag_CMimportExport.asp |

**Determining the Certificate Thumbprint**

Once the certificate infrastructure is in place, the thumbprint or hash of the certificate to be used must be configured with the WinRM service for it to function completely. The Certificates MMC snap-in is used to determine the certificate thumbprint:

1. From the **Start** menu, click **Run**.
2. Enter **mmc** in the **Open:** field and press **OK**.
3. Choose the **Add/Remove Snap-in** option from the **File** menu.
4. Click the **Add** button.
5. Choose the **Certificates** snap-in and press the **Add** button:

6. Choose the Computer account option and press the **Next** button:




7. Ensure that the default for Local Computer is selected and press the **Finish** button.

8. Press the **Close** button.

9. Press the **OK** button.

10. Expand the **Certificates (Local Computer)** element (in the left pane).

11. Expand the **Personal** element (in the left pane).

12. Highlight the **Certificates** element (in the left pane):




13. Double-click on the desired certificate from the right pane:

14. Choose the **Details** tab in the **Certificate** property page.

15. Select the **Thumbprint** field.

16. Highlight the thumbprint value and copy it to the Clipboard by pressing Ctrl+C (you will use this value in the next section):




**Configuring WinRM to Use a Certificate**

Configuration of the HTTPS Certificate to be used by WinRM is done using the WinRM Command-Line tool (winrm.vbs). WinRM uses a specific uniform resource identifier (URI) to access its configuration information:

```
wsman:microsoft.com/wsman/2005/06/config
```

The current configuration can be returned using the following command:

```
winrm get wsman:microsoft.com/wsman/2005/06/config
```

The command returns the current configuration in the following form:

```
Config
    MaxEnvelopeSizekb = 50
    MaxTimeoutms = 60000
    MaxBatchItems = 20
    SoapTraceEnabled = false
    MaxProviderRequests = 25
    Client
        NetworkDelayms = 5000
        URLPrefix = wsman
        HTTP
            Port = 80
            Unencrypted
                Basic = false
                Digest = false
                Negotiate = false
        HTTPS
            Port = 443
            Basic = true
            Digest = true
            Negotiate = true
            Mutual = false
    Service
        RootSDDL = O:NSG:BAD:P(A;;GA;;;BA)S:P(AU;FA;GA;;;WD)(AU;SA;GWGX;;;WD)
        MaxConcurrentOperations = 100
        EnumerationTimeoutms = 60000
        MaxClientCertInfoSize = 16384
        MaxConnections = 5
        HTTP
            Unencrypted
                Basic = false
                Negotiate = false
        HTTPS
            Basic = true
            Negotiate = true
            Mutual = false
```

To see the configuration in XML format, the following command can be executed:

```
winrm get wsman:microsoft.com/wsman/2005/06/config -format:#xml
```

This will return the following XML output:

```
<cfg:Config xmlns:cfg="wsman:microsoft.com/wsman/2005/06/config.xsd">
  <cfg:MaxEnvelopeSizekb>50</cfg:MaxEnvelopeSizekb>
  <cfg:MaxTimeoutms>60000</cfg:MaxTimeoutms>
  <cfg:MaxBatchItems>20</cfg:MaxBatchItems>
  <cfg:SoapTraceEnabled>false</cfg:SoapTraceEnabled>
  <cfg:MaxProviderRequests>25</cfg:MaxProviderRequests>
   <cfg:Client>
    <cfg:NetworkDelayms>5000</cfg:NetworkDelayms>
    <cfg:URLPrefix>wsman</cfg:URLPrefix>
    <cfg:HTTP>
      <cfg:Port>80</cfg:Port>
          <cfg:Unencrypted>
        <cfg:Basic>false</cfg:Basic>
        <cfg:Digest>false</cfg:Digest>
        <cfg:Negotiate>false</cfg:Negotiate>
      </cfg:Unencrypted>
    </cfg:HTTP>
    <cfg:HTTPS>
      <cfg:Port>443</cfg:Port>
      <cfg:Basic>true</cfg:Basic>
      <cfg:Digest>true</cfg:Digest>
      <cfg:Negotiate>true</cfg:Negotiate>
      <cfg:Mutual>false</cfg:Mutual>
    </cfg:HTTPS>
  </cfg:Client>
   <cfg:Service>
 <cfg:RootSDDL>O:NSG:BAD:P(A;;GA;;;BA)S:P(AU;FA;GA;;;WD)(AU;SA;GWGX;;;WD)</cfg:RootSDDL>
    <cfg:MaxConcurrentOperations>100</cfg:MaxConcurrentOperations>
    <cfg:EnumerationTimeoutms>60000</cfg:EnumerationTimeoutms>
    <cfg:MaxClientCertInfoSize>16384</cfg:MaxClientCertInfoSize>
    <cfg:MaxConnections>5</cfg:MaxConnections>
    <cfg:HTTP>
      <cfg:Unencrypted>
        <cfg:Basic>false</cfg:Basic>
        <cfg:Negotiate>false</cfg:Negotiate>
      </cfg:Unencrypted>
    </cfg:HTTP>
    <cfg:HTTPS>
      <cfg:Basic>true</cfg:Basic>
      <cfg:Negotiate>true</cfg:Negotiate>
      <cfg:Mutual>false</cfg:Mutual>
    </cfg:HTTPS>
  </cfg:Service>
</cfg:Config>
```

**Creating Listener Instances**

Multiple WinRM Listener instances can be created on a single system. The Listener instances are differentiated by the IP address and the port.

To configure the Listener HTTPS, a Listener instance must be created. The following command can be issued (using the thumbprint copied from the Certificates MMC snap-in in the previous section) to create an instance:

```
winrm create wsman:microsoft.com/wsman/2005/06/config/Listener?IP=192.168.2.1+Port=443
@{Hostname="host";CertificateThumbprint="xxx"}
```

Where:

•**IP** is the TCP/IP address of the Listener system.

•**Port** is the TCP/IP port that the Listener is to listen to requests on.

•**Hostname** is the name of the system as it appears in the Certificate.

•**CertificateThumbprint** is the thumbprint copied from the Certificates MMC Snap-in from the previous section. Note that you must remove the spaces from between the two-digit hex values in the thumbprint as it appears in the Certificate dialog box, for example, **CertificateThumbprint="b7a24c3f42 ..."**.

| **Note** |
| --- |
| In the URI, "Listener" must be capitalized. When using the **Create** or<br><br>**Invoke** commands, the last part of the resource URI must match case-wise the top-level element of the expected XML. |

If the configuration settings are correct, the following response is returned:

```
<wxf:ResourceCreated
xmlns:wxf="http://schemas.xmlsoap.org/ws/2004/09/transfer"xmlns:wsa="http://schemas.xmlsoap.org/ws/
2004/08/addressing" xmlns:wsman="http://schemas.xmlsoap.org/ws/2005/06/management">
<wsa:Address>http://schemas.xmlsoap.org/ws/2004/08/addressing/role/anonymous?ResourceURI=
(wsman:microsoft.com/wsman/2005/06/config/listener)</wsa:Address>
  <wsa:ReferenceParameters>
```

```
    <wsman:SelectorSet>
      <wsman:Selector Name="IP">192.168.2.1</wsman:Selector>
      <wsman:Selector Name="Port">443</wsman:Selector>
    </wsman:SelectorSet>
  </wsa:ReferenceParameters>
</wxf:ResourceCreated>
```

| Note |
| --- |
| The WinRM Command Line interface supports the configuration of Listener instances using IPv6. The IPv4 address is simply replaced by the IPv6 address. |

To configure the Listener for HTTP, the following command can be issued to create an instance:

```
winrm create wsman:microsoft.com/wsman/2005/06/config/Listener?IP=192.168.2.1+Port=80
@{Hostname="host";Transport="http"}
```

Where:

- **IP** is the TCP/IP address of the Listener system.
- **Port** is the TCP/IP port that the Listener is to listen to requests on.
- **Hostname** is the name of the system as it appears in the Certificate.

If the configuration settings are correct, the following response is returned:

```
<wxf:ResourceCreated
xmlns:wxf="http://schemas.xmlsoap.org/ws/2004/09/transfer"xmlns:wsa="http://schemas.xmlsoap.org/ws/
2004/08/addressing" xmlns:wsman="http://schemas.xmlsoap.org/ws/2005/06/management">
<wsa:Address>http://schemas.xmlsoap.org/ws/2004/08/addressing/role/anonymous?ResourceURI=
(wsman:microsoft.com/wsman/2005/06/config/listener)</wsa:Address>
  <wsa:ReferenceParameters>
    <wsman:SelectorSet>
      <wsman:Selector Name="IP">192.168.2.1</wsman:Selector>
      <wsman:Selector Name="Port">443</wsman:Selector>
    </wsman:SelectorSet>
  </wsa:ReferenceParameters>
</wxf:ResourceCreated>
```

**Configuring a Listener Instance Using DHCP**

WinRM Listener instances can be configured on systems using DHCP to get their IP addresses. To create a Listener instance on a DHCP-enabled network interface card (NIC), the MAC address of the NIC is required and the following command can be used:

```
winrm create wsman:microsoft.com/wsman/2005/06/config/Listener?IP=*+Port=443
@{Hostname="host";MACAddress="12-34-56-78-90-AB";CertificateThumbprint="xxx"}
```

This configures the WinRM service to enumerate all the IP addresses on that NIC and listen on the DHCP-assigned address of the NIC.

| Note |
| --- |
| There is a limitation of a single Listener instance for every NIC. For example, on a system with two NICs, only two Listener instances can be created. |

**Co-Existing with Internet Information Server**

WinRM can co-exist with Internet Information Server (IIS) if it is configured for SSL communication. In this instance, WinRM checks to see if the shared certificate configuration of SSL can be reused. If it can, WS-Management will use the same certificate configuration.

| Note |
| --- |
| Deleting a shared IIS configuration from IIS may interfere with WS-Management communications. If an administrator creates a WS-Management listener that shares an HTTPS port with IIS, and later removes the SSL port configuration from IIS, this may cause the WS-Management Listener to stop responding. |

**Automating WinRM Configuration**

To simplify automatic configuration of the WinRM service, the WinRM Configuration Plug-in supports helper methods to retrieve certificates from the machine store. These methods can be performed locally or against a remote machine and can then be used to automatically configure the certificate. Note that Listener configuration can be done only locally.

In this case, a hostname—either the NetBIOS or Fully-Qualified Domain Name (FQDN)—is specified as the input parameter. The first certificate whose CN matches the hostname and whose EKU is set to server authentication is returned.

In this example, the command retrieves configuration from host.fabrikam.com using the FQDN:

```
winrm invoke Retrieve wsman:Microsoft.com/wsman/2005/06/config/Certificate
@{Hostname="host.fabrikam.com"}
```

If successful, a typical result will be returned as shown below with the Certificate Thumbprint embedded within the XML:

```
<cfg:Retrieve_OUTPUT xmlns:cfg="wsman:microsoft.com/wsman/2005/06/config/certificate.xsd">
<cfg:CertificateThumbprint>4907de1950d4a162926d56cb41392d388aac874a</cfg:CertificateThumbprint>
</cfg:Retrieve_OUTPUT>
```

| **Note** |
| --- |
| Once this Certificate Thumbprint is available, a Listener instance can be created using the steps outlined in the previous section. |

If there is no certificate that matches the above criteria, an error is returned:

```
<f:WSManFault xmlns:f="http://schemas.microsoft.com/ws/2005/06/wsmanfault" Code="32898"
Machine="AMD64-R3-10">
  <f:Message>
    <f:ProviderFault providerId="A855B0EE-65C2-4282-830B-7FA181E73BFF">
      <f:WSManFault xmlns:f="http://schemas.microsoft.com/ws/2005/06/wsmanfault" Code="32898"
Machine="">
        <f:Message>The requested certificate was not found.</f:Message>
      </f:WSManFault>
    </f:ProviderFault>
  </f:Message>
</f:WSManFault>
Error number: -2146991998 0x80078082
The requested certificate was not found.
```

Another important piece of information required for the automated configuration of WinRM is whether the WinRM service is installed and running properly.

WMI can be used to check whether the WinRM service is installed. To do this, simply enumerate the **Win32_Service** class and look for service whose name is **winrm**.

In order to check whether the WinRM service is running on a machine, a **Get** request on a specific URI (the URI named "this") can be used. If the WinRM service is running, an XML response is returned with information about the service.

The command to perform this **Get** operation is:

```
winrm get wsman:system/2005/06/this -target:host.acme.com
```

The response would include the version and capabilities of the service:

```
<t:This xmlns:t="http://schemas.xmlsoap.org/ws/2005/06/management/this">
  <t:Vendor>Microsoft Corporation</t:Vendor>
  <t:Version>OS: 5.2.3790 SP: 1.0 Stack: 1.0</t:Version>
  <t:Access>
    <t:Compliance>http://schemas.xmlsoap.org/ws/2005/06/management</t:Compliance>
    <t:Action>http://schemas.xmlsoap.org/ws/2004/09/transfer/Get</t:Action>
    <t:Action>http://schemas.xmlsoap.org/ws/2004/09/transfer/Put</t:Action>
    <t:Action>http://schemas.xmlsoap.org/ws/2004/09/transfer/Create</t:Action>
    <t:Action>http://schemas.xmlsoap.org/ws/2004/09/transfer/Delete</t:Action>
    <t:Action>http://schemas.xmlsoap.org/ws/2004/09/enumeration/Enumerate</t:Action>
  </t:Access>
  <t:Transports>
    <t:Transport Name="http://www.ietf.org/rfc/rfc2818.txt">
      <t:Auth>wsman:auth/basic</t:Auth>
      <t:Auth>wsman:auth/https-mutual</t:Auth>
    </t:Transport>
  </t:Transports>
</t:This>
```

If the WinRM service is not running the following error is returned:

```
Specified destination cannot be connected to.
```

**Viewing Listener Instances**

To view all the Listener instances that exist on the system, the following command can be executed:

```
winrm enumerate wsman:microsoft.com/wsman/2005/06/config/Listener
```

This returns a result similar to this:

```
<wsman:Results xmlns:wsman="http://schemas.microsoft.com/ws/2005/03/wsman/results">
  <cfg:Listener xmlns:cfg="wsman:microsoft.com/wsman/2005/06/config/listener.xsd">
    <cfg:IP>192.168.2.1</cfg:IP>
    <cfg:Port>443</cfg:Port>
    <cfg:Hostname>host.acme.com</cfg:Hostname>
    <cfg:MACAddress/>
```

```
    <cfg:Enabled>true</cfg:Enabled>
    <cfg:URLPrefix>wsman</cfg:URLPrefix>
    <cfg:Transport>HTTPS</cfg:Transport>
    <cfg:CertificateThumbprint>49 07 de 19 50 d4 a1 62 92 6d 56 cb 41 39 2d 38 8a ac 87
4a</cfg:CertificateThumbprint>
  </cfg:Listener>
</wsman:Results>
```

If multiple Listener instances are created, the resulting output XML is similar to this:

```
<wsman:Results xmlns:wsman="http://schemas.microsoft.com/ws/2005/03/wsman/results">
  <cfg:Listener xmlns:cfg="wsman:microsoft.com/wsman/2005/06/config/listener.xsd">
    <cfg:IP>192.168.2.1</cfg:IP>
    <cfg:Port>443</cfg:Port>
    <cfg:Hostname>host.acme.com</cfg:Hostname>
    <cfg:MACAddress/>
    <cfg:Enabled>true</cfg:Enabled>
    <cfg:URLPrefix>wsman</cfg:URLPrefix>
    <cfg:Transport>HTTPS</cfg:Transport>
    <cfg:CertificateThumbprint>49 07 de 19 50 d4 a1 62 92 6d 56 cb 41 39 2d 38 8a ac 87
4a</cfg:CertificateThumbprint>
  </cfg:Listener>
  <cfg:Listener xmlns:cfg="wsman:microsoft.com/wsman/2005/06/config/listener.xsd">
    <cfg:IP>192.168.2.1</cfg:IP>
    <cfg:Port>444</cfg:Port>
    <cfg:Hostname>host.acme.com</cfg:Hostname>
    <cfg:MACAddress/>
    <cfg:Enabled>true</cfg:Enabled>
    <cfg:URLPrefix>wsman</cfg:URLPrefix>
    <cfg:Transport>HTTPS</cfg:Transport>
    <cfg:CertificateThumbprint>49 07 de 19 50 d4 a1 62 92 6d 56 cb 41 39 2d 38 8a ac 87
4a</cfg:CertificateThumbprint>
  </cfg:Listener>
</wsman:Results>
```

If the configuration information is incorrect, Fault XML will be returned followed by the message "The supplied SOAP violates the corresponding XML schema definition."

**Modifying an Existing Listener Instance**

To modify the configuration of an existing Listener instance, use the **Put** action. For example, to disable a Listener instance, execute the following command:

```
winrm put wsman:microsoft.com/wsman/2005/06/config/Listener?IP=192.168.2.1+Port=443
@{Enabled="False"}
```

A successful modification of the Listener instance returns the current configuration:

```
<cfg:Listener xmlns:cfg="wsman:microsoft.com/wsman/2005/06/config/listener.xsd">
  <cfg:IP>192.168.2.1</cfg:IP>
  <cfg:Port>443</cfg:Port>
  <cfg:Hostname>amd64-r3-10.wmx.microsoft.com</cfg:Hostname>
  <cfg:MACAddress/>
  <cfg:Enabled>false</cfg:Enabled>
  <cfg:URLPrefix>wsman</cfg:URLPrefix>
  <cfg:Transport>HTTPS</cfg:Transport>
  <cfg:CertificateThumbprint>49 07 de 19 50 d4 a1 62 92 6d 56 cb 41 39 2d 38 8a ac 87
4a</cfg:CertificateThumbprint>
</cfg:Listener>
```

An incorrect modification of the Listener instance results in Fault XML being returned:

```
<f:WSManFault xmlns:f="http://schemas.microsoft.com/ws/2005/06/wsmanfault" Code="32833"
Machine="AMD64-R3-10">
  <f:Message>
    <f:ProviderFault providerId="A855B0EE-65C2-4282-830B-7FA181E73BFF">
      <f:WSManFault xmlns:f="http://schemas.microsoft.com/ws/2005/06/wsmanfault" Code="32833"
Machine="">
        <f:Message>Input xml contained an unrecognized xml element "enabled".  Double check the
spelling and case of the element name.</f:Message>
      </f:WSManFault>
    </f:ProviderFault>
  </f:Message>
</f:WSManFault>

Error number: -2146992063 0x80078041
The supplied SOAP violates the corresponding XML schema definition.
```

| **Note** |
| --- |
| The URIs listed here are case sensitive. For example, Listener and Enabled must have their first letters capitalized. |

**Deleting a Listener Instance**

To delete a Listener instance that exists on the system, execute the following command:

```
winrm delete wsman:microsoft.com/wsman/2005/06/config/Listener?IP=192.168.2.1+Port=444
```

If the instance is deleted successfully, no response is returned. If an attempt is made to delete an instance that does not exist, the following fault XML is returned:

```
<f:WSManFault xmlns:f="http://schemas.microsoft.com/ws/2005/06/wsmanfault" Code="32768"
Machine="AMD64-R3-10">
  <f:Message>
    <f:ProviderFault providerId="A855B0EE-65C2-4282-830B-7FA181E73BFF">
      <f:WSManFault xmlns:f="http://schemas.microsoft.com/ws/2005/06/wsmanfault" Code="32768"
Machine="">
        <f:Message>The resource identified by the resource uri and selectors could not be found.
</f:Message>
      </f:WSManFault>
    </f:ProviderFault>
  </f:Message>
</f:WSManFault>
Error number: -2146992128 0x80078000
The resource identified by the resource uri and selectors could not be found.
```

# Integration with WMI

Applies To: Windows Server 2003 R2

Another feature of WinRM in Windows Server 2003 R2 is a WMI plug-in that exposes Windows Management Instrumentation (WMI) classes, methods and properties to WinRM. To use WMI with WinRM, you must convert the WMI Common Information Model (CIM) namespace and class to a WS-Management Uniform Resource Identifier (URI). This topic introduces the steps involved in this conversion.

**URI Prefixes**

When accessing WMI classes via WS-Management, a URI is required. The URI consists of the URI prefix and the class to be accessed. Similarly, the URI prefix consists of a hard-coded portion and the CIM namespace. For example, in Windows Server 2003 R2, the URI prefix is:

```
http://schemas.microsoft.com/wsman/2005/06/wmi/root/CIMNamespace
```

This allows the URI prefix to be generated for any CIM namespace accessible via WMI. If, for example, the goal is to access the root\default namespace, the following URI prefix would be used:

```
http://schemas.microsoft.com/wsman/2005/06/wmi/root/default/
```

| **Note** |
| --- |
| Notice that the URI uses forward slashes (/), while the CIM namespace uses backslashes (\). |

In Windows Server 2003 R2, the URI prefix when accessing WMI CIM classes in the root/cimv2 namespace is:

```
http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2/
```

In Windows Server 2003 R2, the URI prefix when accessing the IPMI provider classes is:

```
http://schemas.dmtf.org/wsman/2005/06/wmi/root/cimv2.9/
```

You can also use the following URI to access the IPMI provider classes via WMI:

```
http://schemas.microsoft.com/wsman/2005/06/wmi/root/hardware/
```

| **Note** |
| --- |
| The IPMI provider uses the root/hardware WMI namespace. |

**Generating WMI URIs**

Most WMI classes and operations are currently supported in the WinRM WMI plug-in. (Embedded objects are currently not supported.) Generating the URI for these classes is a fairly simple process. The following section illustrates the process used to generate these URIs.

When generating a URI to a WMI class, the URI is simply made up of the URI prefix (discussed in the previous section), followed by the WMI namespace and the class information:

```
http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2/WMIClass
```

To set the value of a property or to invoke a specific method, a URI such as the following would be used:

```
http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2/WMIClass?
Key1=Value1+Key2=Value2
```

**URI Generation Examples**

This section illustrates some examples of getting WMI information via WS-Management. For these examples, the **Win32_Service** WMI class will be used.

Using WS-Management, WMI classes, methods and properties can be retrieved (**Get**), set (**Put**), listed (**Enumerate**), and executed (**Invoke**).

To enumerate the services within **Win32_Service**, for example, the following URI would be used:

```
http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2/Win32_Service
```

To get the properties of a specific service within **Win32_Service**, the following URI would be used:

```
http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2/Win32_Service?Name=Alerter
```

# Windows Remote Management Command-Line Tool (Winrm.cmd)

Applies To: Windows Server 2003 R2

Windows Remote Management (WinRM) uses a new command-line tool, Winrm.cmd, to perform its operations. These operations include **Get**, **Put**, **Invoke**, and **Enumerate**. The command-line tool also has several switches to allow for remote operations, credentials, and hash tables. Winrm.cmd runs a script, Winrm.vbs, which is written in Visual Basic Scripting Edition.

**Authentication mechanisms**
The following authentication mechanisms are supported:

- Basic
- Digest [Client Only]
- Negotiate

The Negotiate authentication mechanism allows authentication using Kerberos when using domain account, or NTLM when using local accounts.

There are two ways to use this mechanism: If no credentials are specified, then the logged-on credentials are used to authenticate against the remote machine. This allows for a single sign-on experience. If explicit credentials are specified by designating the username and password, then Kerberos is used when credentials for a domain account are specified, and NTLM is used when credentials for a local account on the remote machine are specified.

Negotiate is the default authentication mechanism if no authentication is specified in the command-line tool.

**Specifying a Remote Machine**
A remote machine can be specified using the **–machine** switch. When specifying the remote machine, the following are valid:

- Localhost
- NetBIOS name
- Fully Qualified Domain Name (FQDN)
- IP address

**Using self-signed certificates**
When using self-signed certificates on the server, the client should specify

**–skipCAcheck** in order to skip authentication of the Certificate Authority (CA) that issued the certificate. Note that this should be used only in lab environments due to the security risks.

When using certificates on the server that are issued to an entity other than the server hostname (for example, to the IP address), the client should specify **–skipCNcheck** to skip authentication of the Common Name (CN). Note that this should be used only in lab environments due to the security risks.

| **Note** |
| --- |
| Two switches can be used when communicating with a remote host:<br><br>**-skipCAcheck** Specifies that the certificate issuer need not be a trusted root authority. Used only in remote operations. Note that this option should be used only for trusted machines.<br><br>**-skipCNcheck** Specifies that certificate CN of the server need not match the hostname of the server. Used only in remote operations. Note that this option should be used only for trusted machines. |

Another optional switch that can be used on the command line in conjunction with the **–machine** switch is the **–port** switch.

The **–port** switch allows you to specify on which port the remote WinRM Listener service is listening for requests.

| **Note** |
| --- |
| When using the **–machine** switch, the name used needs to match the CN used in the Server Certificate on the remote system. |

This example uses an FQDN to specify the remote machine:

```
winrm action uri -machine:server.acme.dom -port:443
```

**Handling Hash Tables**

As you'll see in the following sections, both **Put** and **Invoke** operations require input to be supplied for the operation to succeed. There are a couple of different ways to supply this input. The first is by generating the input XML file and letting the command-line tool read the file using the **–file** switch. The second is by using a hash table.

A hash table is a method of supplying property names and their values. The Winrm command-line tool will then generate the input XML automatically.

| **Note** |
| --- |
| More than one key/value pair can be specified by placing a semicolon (;) between each pair (with no space). The quotation marks around the value are required. |

When performing an **Invoke** operation that requires no input, an empty hash table is needed:

```
@{}
```

| **Note** |
| --- |
| The hash table must be the last item on the command line and must be preceded by a space. This is also true with using any other command-line switches, such as **–machine**. If a hash table is used, the switches must come before the hash table. |

**Output Format**

The Winrm command–line tool allows the administrator to choose the type of output that is returned. There are three options:

```
-format:#xml
-format:#pretty
-format:#text
```

The default is text output. An example of each output format option is shown below for an enumeration of **Win32_LogicalDisk**.

| **Note** |
| --- |
| When redirecting the output to a file, Cmd.exe always saves the text as ANSI. When there are Unicode characters in the XML output, they cause the XML to be malformed. To prevent this, open the file in Notepad and save as Unicode. You can also write the output with a script using the XML Document Object Model, which saves the file as Unicode. |

**Raw XML**

In this case, the raw XML is shown without any changes.


```
<wsman:Results>
<p:Win32_LogicalDisk xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:p="http://schemas.microsoft.org/root/cimv2/2004/07"
__uri="wsman:microsoft.operatingsystem.wmi/root/cimv2/Win32_LogicalDisk?DeviceID=C:">
<p:Access xsi:nil="true"/><p:Availability xsi:nil="true"/><p:BlockSize xsi:nil="true"/>
<p:Caption>C:</p:Caption><p:Compressed>false</p:Compressed><p:ConfigManagerErrorCode
xsi:nil="true"/><p:ConfigManagerUserConfig xsi:nil="true"/>
<p:CreationClassName>Win32_LogicalDisk</p:CreationClassName><p:Description>Local Fixed
Disk</p:Description><p:DeviceID>C:</p:DeviceID><p:DriveType>3</p:DriveType>
<p:ErrorCleared xsi:nil="true"/><p:ErrorDescription xsi:nil="true"/><p:ErrorMethodology
xsi:nil="true"/><p:FileSystem>NTFS</p:FileSystem><p:FreeSpace>22935719936</p:FreeSpace>
<p:InstallDate xsi:nil="true"/><p:LastErrorCode xsi:nil="true"/>
<p:MaximumComponentLength>255</p:MaximumComponentLength><p:MediaType>12</p:MediaType>
<p:Name>C:</p:Name><p:NumberOfBlocks xsi:nil="true"/><p:PNPDeviceID xsi:nil="true"/>
```

```
<p:PowerManagementCapabilities xsi:nil="true"/><p:PowerManagementSupported
xsi:nil="true"/><p:ProviderName xsi:nil="true"/><p:Purpose xsi:nil="true"/>
<p:QuotasDisabled>true</p:QuotasDisabled><p:QuotasIncomplete>true</p:QuotasIncomplete>
<p:QuotasRebuilding>false</p:QuotasRebuilding><p:Size>36880187392</p:Size><p:Status
xsi:nil="true"/><p:StatusInfo xsi:nil="true"/>
<p:SupportsDiskQuotas>true</p:SupportsDiskQuotas>
<p:SupportsFileBasedCompression>true</p:SupportsFileBasedCompression>
<p:SystemCreationClassName>Win32_ComputerSystem</p:SystemCreationClassName>
<p:SystemName>MyMachine</p:SystemName><p:VolumeDirty>false</p:VolumeDirty>
<p:VolumeName>IBM_PRELOAD</p:VolumeName>
<p:VolumeSerialNumber>6440000C</p:VolumeSerialNumber></p:Win32_LogicalDisk>
<p:Win32_LogicalDisk xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:p="http://schemas.microsoft.org/root/cimv2/2004/07"
__uri="wsman:microsoft.operatingsystem.wmi/root/cimv2/Win32_LogicalDisk?DeviceID=D:">
<p:Access xsi:nil="true"/><p:Availability xsi:nil="true"/><p:BlockSize xsi:nil="true"/>
<p:Caption>D:</p:Caption><p:Compressed xsi:nil="true"/><p:ConfigManagerErrorCode
xsi:nil="true"/><p:ConfigManagerUserConfig xsi:nil="true"/>
<p:CreationClassName>Win32_LogicalDisk</p:CreationClassName><p:Description>CD-ROM
Disc</p:Description><p:DeviceID>D:</p:DeviceID><p:DriveType>5</p:DriveType>
<p:ErrorCleared xsi:nil="true"/><p:ErrorDescription xsi:nil="true"/><p:ErrorMethodology
xsi:nil="true"/><p:FileSystem xsi:nil="true"/><p:FreeSpace xsi:nil="true"/>
<p:InstallDate xsi:nil="true"/><p:LastErrorCode xsi:nil="true"/>
<p:MaximumComponentLength xsi:nil="true"/><p:MediaType>11</p:MediaType><p:Name>D:
</p:Name><p:NumberOfBlocks xsi:nil="true"/><p:PNPDeviceID xsi:nil="true"/>
<p:PowerManagementCapabilities xsi:nil="true"/><p:PowerManagementSupported
xsi:nil="true"/><p:ProviderName xsi:nil="true"/><p:Purpose xsi:nil="true"/>
<p:QuotasDisabled xsi:nil="true"/><p:QuotasIncomplete xsi:nil="true"/>
<p:QuotasRebuilding xsi:nil="true"/><p:Size xsi:nil="true"/><p:Status xsi:nil="true"/>
<p:StatusInfo xsi:nil="true"/><p:SupportsDiskQuotas xsi:nil="true"/>
<p:SupportsFileBasedCompression xsi:nil="true"/>
<p:SystemCreationClassName>Win32_ComputerSystem</p:SystemCreationClassName>
<p:SystemName> MyMachine </p:SystemName><p:VolumeDirty xsi:nil="true"/><p:VolumeName
xsi:nil="true"/><p:VolumeSerialNumber xsi:nil="true"/></p:Win32_LogicalDisk>
</wsman:Results>
```


**Pretty Formatting**

This option displays XML in indented format. A four-character indentation is used for displaying nested XML elements. Internet Explorer can also be used to display in this format from raw XML output.


```
<cfg:Config xmlns:cfg="wsman:microsoft.com/wsman/2005/02/config.xsd">
    <cfg:MaxEnvelopeSizekb>50</cfg:MaxEnvelopeSizekb>
    <cfg:MaxTimeoutms>60000</cfg:MaxTimeoutms>
    <cfg:MaxBatchItems>5</cfg:MaxBatchItems>
    <cfg:SoapTraceEnabled>true</cfg:SoapTraceEnabled>
    <cfg:MaxProviderRequests>25</cfg:MaxProviderRequests>
    <cfg:Client>
        <cfg:NetworkDelayms>5000</cfg:NetworkDelayms>
        <cfg:HTTPS>
            <cfg:Port>443</cfg:Port>
            <cfg:URLPrefix>wsman</cfg:URLPrefix>
        </cfg:HTTPS>
    </cfg:Client>
    <cfg:Service>
        <cfg:RootSDDL>O:NSG:BAD:P(A;;GA;;;BA)S:P(AU;SAFA;GWGX;;;WD)
        </cfg:RootSDDL>
        <cfg:MaxConcurrentOperations>100</cfg:MaxConcurrentOperations>
        <cfg:EnumerationTimeoutms>60000</cfg:EnumerationTimeoutms>
        <cfg:MaxClientCertInfoSize>16384</cfg:MaxClientCertInfoSize>
        <cfg:HTTPS>
            <cfg:MaxConnections>50</cfg:MaxConnections>
        </cfg:HTTPS>
    </cfg:Service>
</cfg:Config>
```


**Text Formatting**

The command-line tool supports basic text formatting of output XML. A four-character indentation is used for displaying nested elements.


```
Win32Logical_Disk
    DeviceID = C:
    Access = NULL
    Availability = NULL
```

```
    BlockSize = NULL
    Caption = C:
    Compressed = false

Win32Logical_Disk
    DeviceID = D:
    Access = NULL
    Availability = NULL
    BlockSize = NULL
    Caption = D:
    Compressed = false
```

**URI Aliases**

URI aliases can simplify the Winrm command line. The following URI aliases are supported:

```
wmi = http://schemas.microsoft.com/wsman/2005/06/wmi
wsman = wsman:microsoft.com/wsman/2005/06/
cimv2.9 = http://schemas.dmtf.org/wsman/2005/06/cimv2.9
cimv2 = http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2
```

For example, the following command:

```
winrm get http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2/Win32_Service?
Name=WSMan
```

Gets replaced with:

```
winrm get wmi/root/cimv2/Win32_Service?Name=WinRM
```

**Performing a Get Operation**

The WS-Management **Get** operation returns the value of a specific WMI object. In this example, WS-Management retrieves the properties of the **Alerter** service:

```
winrm get http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2/win32_service?
name=Alerter
```

If this operation completes successfully, an XML result similar to the following is returned:

```
<p:win32_service xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:p="http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2">
  <p:AcceptPause>false</p:AcceptPause>
  <p:AcceptStop>false</p:AcceptStop>
  <p:Caption>Alerter</p:Caption>
  <p:CheckPoint>0</p:CheckPoint>
  <p:CreationClassName>Win32_Service</p:CreationClassName>
  <p:Description>Notifies selected users and computers of administrative alerts. If the
service is stopped, programs that use administrative alerts will not receive them. If
this service is disabled, any services that explicitly depend on it will fail to start.
</p:Description>
  <p:DesktopInteract>false</p:DesktopInteract>
  <p:DisplayName>Alerter</p:DisplayName>
  <p:ErrorControl>Normal</p:ErrorControl>
  <p:ExitCode>1077</p:ExitCode>
  <p:InstallDate xsi:nil="true"/>
  <p:Name>Alerter</p:Name>
  <p:PathName>C:\WINDOWS\System32\svchost.exe -k LocalService</p:PathName>
  <p:ProcessId>0</p:ProcessId>
  <p:ServiceSpecificExitCode>0</p:ServiceSpecificExitCode>
  <p:ServiceType>Share Process</p:ServiceType>
  <p:Started>false</p:Started>
  <p:StartMode>Disabled</p:StartMode>
  <p:StartName>NT AUTHORITY\LocalService</p:StartName>
  <p:State>Stopped</p:State>
  <p:Status>OK</p:Status>
  <p:SystemCreationClassName>Win32_ComputerSystem</p:SystemCreationClassName>
  <p:SystemName>Host</p:SystemName>
  <p:TagId>0</p:TagId>
  <p:WaitHint>0</p:WaitHint>
  <cim:Location xmlns:cim="http://schemas.dmtf.org/wsman/2005/02/base">
    <wsa:EndpointReference xmlns:wsa="http://schemas.xmlsoap.org/ws/2004/08/addressing"
xmlns:wsman="http://schemas.xmlsoap.org/ws/2005/02/management">
<wsa:Address>http://schemas.xmlsoap.org/ws/2004/08/addressing/role/anonymous?
```

```
ResourceURI=(http://schemas.microsoft.com/wsman/2005/02/wmi/root/cimv2/Win32_Service)
</wsa:Address>
      <wsa:ReferenceParameters>
        <wsman:SelectorSet>
          <wsman:Selector Name="Name">Alerter</wsman:Selector>
        </wsman:SelectorSet>
      </wsa:ReferenceParameters>
    </wsa:EndpointReference>
  </cim:Location>
</p:win32_service>
```

**Performing a Put Operation**

The WS-Management **Put** operation allows the values of keys to be set. In this example, the value of the **MaxEnvelopeSizekb** key is configured:

```
winrm put wsman:microsoft.com/wsman/2005/06/Config @{MaxEnvelopeSizekb="100"}
```

If the key is modified successfully, the configuration with the new value(s) will be returned in the result XML:

```
<cfg:Config xmlns:cfg="wsman:microsoft.com/wsman/2005/06/config.xsd">
  <cfg:MaxEnvelopeSizekb>100</cfg:MaxEnvelopeSizekb>
  <cfg:MaxTimeoutms>60000</cfg:MaxTimeoutms>
  <cfg:MaxBatchItems>20</cfg:MaxBatchItems>
  <cfg:SoapTraceEnabled>false</cfg:SoapTraceEnabled>
  <cfg:MaxProviderRequests>25</cfg:MaxProviderRequests>
  <cfg:Client>
    <cfg:NetworkDelayms>5000</cfg:NetworkDelayms>
    <cfg:HTTPS>
      <cfg:Port>443</cfg:Port>
      <cfg:URLPrefix>wsman</cfg:URLPrefix>
    </cfg:HTTPS>
  </cfg:Client>
  <cfg:Service>
    <cfg:RootSDDL>O:NSG:BAD:P(A;;GA;;;BA)S:P(AU;SAFA;GWGX;;;WD)</cfg:RootSDDL>
    <cfg:MaxConcurrentOperations>100</cfg:MaxConcurrentOperations>
    <cfg:EnumerationTimeoutms>60000</cfg:EnumerationTimeoutms>
    <cfg:MaxClientCertInfoSize>16384</cfg:MaxClientCertInfoSize>
    <cfg:HTTPS>
      <cfg:MaxConnections>50</cfg:MaxConnections>
    </cfg:HTTPS>
  </cfg:Service>
</cfg:Config>
```

If the key cannot be modified, Fault XML similar to the XML below is returned:

```
<f:WSManFault xmlns:f="http://schemas.microsoft.com/ws/2005/06/wsmanfault" Code="32833"
Machine="AMD64-R3-10">
  <f:Message>
    <f:ProviderFault providerId="A855B0EE-65C2-4282-830B-7FA181E73BFF">
      <f:WSManFault xmlns:f="http://schemas.microsoft.com/ws/2005/02/wsmanfault"
Code="32833" Machine="">
        <f:Message>Input xml contained an unrecognized xml element "MaxEnvelopeSizeKb".
Double check the spelling and case of the element name.</f:Message>
      </f:WSManFault>
    </f:ProviderFault>
  </f:Message>
</f:WSManFault>

Error number: -2146992063 0x80078041
The supplied SOAP violates the corresponding XML schema definition.
```

**Performing an Enumerate Operation**

The WS-Management **Enumerate** operation returns a collection of objects. In this example, the command returns all the services that exist on the computer, using the **Win32_Service** WMI class:

```
winrm enumerate http://schemas.microsoft.com/wsman/2005/02/wmi/root/cimv2/Win32_Service
```

This returns XML data for each service within the enumeration. The result is similar to the **Get** example, except that the information for all the services is returned, rather than a single service. For example, output XML is returned for each service instance in the enumeration. Below is an example of one such instance:

```
<wsman:Results xmlns:wsman="http://schemas.microsoft.com/ws/2005/03/wsman/results">
```

```
  <p:win32_service xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:p="http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2">
    <p:AcceptPause>false</p:AcceptPause>
    <p:AcceptStop>false</p:AcceptStop>
    <p:Caption>Alerter</p:Caption>
    <p:CheckPoint>0</p:CheckPoint>
    <p:CreationClassName>Win32_Service</p:CreationClassName>
    <p:Description>Notifies selected users and computers of administrative alerts. If
the service is stopped, programs that use administrative alerts will not receive them.
If this service is disabled, any services that explicitly depend on it will fail to
start.</p:Description>
    <p:DesktopInteract>false</p:DesktopInteract>
    <p:DisplayName>Alerter</p:DisplayName>
    <p:ErrorControl>Normal</p:ErrorControl>
    <p:ExitCode>1077</p:ExitCode>
    <p:InstallDate xsi:nil="true"/>
    <p:Name>Alerter</p:Name>
    <p:PathName>C:\WINDOWS\System32\svchost.exe -k LocalService</p:PathName>
    <p:ProcessId>0</p:ProcessId>
    <p:ServiceSpecificExitCode>0</p:ServiceSpecificExitCode>
    <p:ServiceType>Share Process</p:ServiceType>
    <p:Started>false</p:Started>
    <p:StartMode>Disabled</p:StartMode>
    <p:StartName>NT AUTHORITY\LocalService</p:StartName>
    <p:State>Stopped</p:State>
    <p:Status>OK</p:Status>
    <p:SystemCreationClassName>Win32_ComputerSystem</p:SystemCreationClassName>
    <p:SystemName>Host</p:SystemName>
    <p:TagId>0</p:TagId>
    <p:WaitHint>0</p:WaitHint>
    <cim:Location xmlns:cim="http://schemas.dmtf.org/wsman/2005/02/base">
      <wsa:EndpointReference
xmlns:wsa="http://schemas.xmlsoap.org/ws/2004/08/addressing"
xmlns:wsman="http://schemas.xmlsoap.org/ws/2005/02/management">
        <wsa:Address>http://schemas.xmlsoap.org/ws/2004/08/addressing/role/anonymous?
ResourceURI=(http://schemas.microsoft.com/wsman/2005/02/wmi/root/cimv2/Win32_Service)
</wsa:Address>
        <wsa:ReferenceParameters>
          <wsman:SelectorSet>
            <wsman:Selector Name="Name">Alerter</wsman:Selector>
          </wsman:SelectorSet>
        </wsa:ReferenceParameters>
      </wsa:EndpointReference>
    </cim:Location>
  </p:win32_service>
</wsman:Results>
```


**Performing an Invoke Operation**

The WS-Management **Invoke** operation executes methods on the target objects. In this example, Winrm starts the **Alerter** service using the **StartService** WMI method:

```
winrm invoke StartService
http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2/Win32_Service?Name=Alerter
@{}
```

Once this operation has executed, the return value is returned in the XML result:


```
<p:StartService_OUTPUT
xmlns:p="http://schemas.microsoft.com/wsman/2005/06/wmi/root/cimv2">
  <p:_ReturnValue>0</p:_ReturnValue>
</p:StartService_OUTPUT>>
```

# Event Collector

Applies To: Windows Server 2003 R2

One of the capabilities of Hardware Management in Windows Server 2003 R2 and Windows Server 2008 is the ability to subscribe to the hardware events stored in the BMC's System Event Log (SEL) and place them in the Windows Event Viewer. This local hardware event collection is automatically configured, so no further configuration is needed.

When you install the Hardware Management component of Windows Server 2003 R2 or Windows Server 2008, two new event logs are created in the Windows Event Viewer, Forwarded Events and Hardware Events, as shown below:




The Event Collector service is a Windows Remote Management client that is used to create subscriptions to WS-Management event providers and store the events received in the Windows SEL. Two scenarios are supported in R2 and Windows Server 2008.

- Logging of hardware events from the local SEL when the operating system is running ("in-band"), using the IPMI driver and WMI provider.

- Logging of hardware events from a remote server's BMC using WS-Management (when BMC hardware supports this).




Windows Server 2003 R2 and Windows Server 2008 use a component called the System Event Log (SEL) plug-in. The SEL plug-in and the Event Collector subscribe to events stored in the BMC's SEL.




In case of an operating system shutdown or failure, Windows logs IPMI events into the BMC's SEL. The administrator can then access the BMC's SEL using "out-of-band" tools to determine why a machine is not running.

**Windows Event Collector Utility (Wecutil.exe)**

This section discusses how to use the command-line tool, Windows Event Collector Utility (Wecutil.exe), to subscribe and unsubscribe to these hardware events.

**Subscribing to Hardware Events from a Windows Server 2008 System**

Subscribing to events is a manual process. To subscribe to SEL events, use the Event Collector command-line tool, Wecutil.exe. Once an event subscription has been established, it is maintained after rebooting.

To generate the subscription, Wecutil requires an XML configuration file.

1. On the BMC server, open an elevated command prompt and run the following:

- Wecutil qc
- Winrm qc
- Winrm set winrm/config/client @{TrustedHosts="localhost"}
- Wecutil cs SELSubscription.xml (as shown below)

```
 <Subscription
xmlns="http://schemas.microsoft.com/2006/03/windows/events/subscription">
<Description>System Event Log event subscription</Description>
<SubscriptionId>SELSubscription</SubscriptionId>
<Uri>http://schemas.microsoft.com/wbem/wsman/1/logrecord/sel</Uri>
<EventSources>
<EventSource>
<Address>localhost</Address>
</EventSource>
</EventSources>
<LogFile>HardwareEvents</LogFile>
<Delivery Mode="pull">
<PushSettings>
<Heartbeat Interval="10000" />
</PushSettings>
</Delivery>
</Subscription>
```

WinRM (specifically the SEL Plugin) should start to populate the SEL events to [Applications and Services Logs\Hardware Events] in Windows Event Log.

**Troubleshooting Subscribing to Hardware Events in Windows Server 2008**
If the following error occurs in a subscription status, the server might not have a BMC or the BIOS does not support BMC detection through ACPI.

[localhost] - Error - Last retry time: 5/22/2007 1:39:32 PM. Code (0x80041001): <f:ProviderFault provider="SEL plugin for Windows Remote Management " path="%systemroot%\system32\WsmSelPl.dll" xmlns:f="http://schemas.microsoft.com/wbem/wsman/1/wsmanfault"> </f:ProviderFault> Next retry time: 5/22/2007 1:44:32 PM.

The following diagnosis test shows if a server has BMC with BIOS supporting BMC detection through ACPI:

1. Open an elevated command prompt and then launch wbemtest.exe from the command prompt.

2. Connect to root\wmi namespace and attempt to enumerate instances of Microsoft_IPMI class. The resulting dialog box should return an instance showing device ID starting with "ACPI". If not, the system BIOS does not support BMC detection through ACPI, which is a requirement for IPMI driver on Windows Server 2008 . To resolve this, upgrade the BIOS.

**Subscribing to Hardware Events from a Windows Server 2003 R2 System**
Likewise in R2, subscribing to events is a manual process.

To generate the subscription, Wecutil requires an XML configuration file. Windows Server 2003 R2 ships with a sample configuration file, %windir%\system32\WsmSelRg.xml:

```
<Subscription xmlns="http://schemas.microsoft.com/2004/05/windows/events/subscription">
  <Description>Subscription 1</Description>
  <SubscriptionId>sub1</SubscriptionId>
  <Uri>wsman:microsoft/logrecord/sel</Uri>
  <EventSources>
    <EventSource Enabled="true">
      <Address>localhost</Address>
    </EventSource>
```

```
  </EventSources>
  <Delivery Mode="pull">
    <Batching>
      <MaxSize>16000</MaxSize>
      <MaxItems>15</MaxItems>
      <MaxLatencyTime>1000</MaxLatencyTime>
    </Batching>
  </Delivery>
  <Heartbeat Interval="10000"/>
  <LogFile>HardwareEvents</LogFile>
</Subscription>
```

| **Note** |
| --- |
| One of the settings configured in the XML file is the log file to be used to store SEL events. |

To create a new subscription to the SEL events using the default configuration file, issue the following command:

```
wecutil cs %windir%\system32\WsSelRg.xml
```

(%windir% is the directory where Windows Server 2003 R2 is installed.)

Shortly after executing the above command, hardware SEL events will start appearing in the Hardware Events log of the Event Viewer (assuming that Hardware Events was the log file configured in the XML file):




Since SEL events are XML-based, the Event Viewer in Windows Server 2003 R2 has been modified to allow for the viewing these SEL events as XML in the Description area of the event. The following shows properly formatted XML for a hardware event in the Event Properties dialog box:




**Enumerating Subscriptions**

Using the Event Collector command-line tool, existing subscriptions can be enumerated. This information can then be used to unsubscribe from BMC events. The following command enumerates the current subscriptions:

```
wecutil es
```

Assuming that the XML configuration file above was used to subscribe to the events, the following is returned:

```
sub1
```

| **Note** |
| --- |
| The name of the subscription returned above is based on the **<SubscriptionID>** tag in the XML file. |

**Getting Information on an Existing Subscription**

Should a system have several subscriptions configured, you can get information about a specific subscription by issuing the following command:

```
wecutil gs sub1
```

(sub1 is the subscription ID of the subscription to be viewed.)

This command returns data based on the original XML configuration file used in the subscription:

```
EventSource[0]:
  Address: localhost
  Enabled: true
Description: Subscription 1
Uri: wsman:microsoft/logrecord/sel
DeliveryMode: pull
DeliveryMaxBufferSize: 16000
DeliveryMaxItems: 15
DeliveryMaxLatencyTime: 1000
HeartbeatInterval: 10000
Locale:
ContentFormat: renderedtext
LogFile: HardwareEvents
```

**Unsubscribing to Hardware Events**

To unsubscribe from the BMC SEL events, the following command is issued:

```
wecutil ds sub1
```

(sub1 is the subscription ID.)

Future hardware events logged to the SEL by the BMC will not appear in the Event Viewer until a new subscription is created.

Any events that have already been logged to the Event Viewer will remain there until manually deleted.

# Hardware Management Architecture

Applies To: Windows Server 2003 R2

This topic discusses the evolution of industry efforts to define a common architecture for managing platform hardware. It then describes Microsoft's implementation, Windows Remote Management, and compares it with previous efforts.

**Platform Hardware Management Today**

Management of platform hardware is a significant part of overall IT administration cost. To address the needs of hardware management, and in particular to enable command and control of hardware components during pre-boot and post-failure of the main operating system, the OEM industry has been converging on a common architecture that includes a specialized Baseboard Management Controller (BMC). The BMC monitors the state of the server hardware and provides remote access to control this hardware, retrieve its state, and receive notifications about critical errors and other hardware state changes.

Several standards have emerged to define the architecture of BMC. The one that is getting the most traction in the industry is Intelligent Platform Management Interface (IPMI). However, in spite of the emergence of this standard, at the system administrator level the management access to server hardware remains highly proprietary to the individual platform implementations and requires use of proprietary management tools supplied by OEMs. This makes the administration experience platform-specific even for very basic operations. The problem is exacerbated by the fact that remote access to BMC is provided via a specialized wire protocol, Remote Management Control Protocol (RMCP), with non-standard security mechanisms.

To achieve uniformity of administration experience across operating system states, various proxy models can be employed. One of the suggested solutions is based on the existing Web-Based Enterprise Management (WBEM) protocol.

The following diagram illustrates that architecture:

**Proxy based architecture**




This architecture has significant problems however:

- The WBEM protocol uses the HTTP transport with special extensions. This protocol has not been broadly adopted, mainly because it is not aligned with the Web services technologies.
- The protocol stack is not supported by the major operating systems' development platforms in the box.
- To use WBEM the customers have to deploy additional components on each server, which increases cost.
- The architecture requires a proxy node that translates the WBEM operations to the wire-level IPMI messages.

**The Microsoft Implementation**

The key element of the Microsoft solution is a standard management protocol: WS-Management. WS-Management is defined in such a way that it scales from a low footprint environment to the operating system-hosted implementations. This introduces uniformity of the operations and a single security model between all operating system states. Because the protocol is defined as a profile of the standard Web services protocols it is also interoperable with non-Microsoft operating systems and development platforms. This creates an opportunity to address the existing customer concerns with high total cost of ownership (TCO) associated with the lack of standardization in managing multi-vendor heterogeneous data centers.

Because the protocol is defined as a composable standard it can be implemented in a resource-constrained environment. This eliminates the need for the low-level wire protocol altogether and thus eliminates the proxy node as a protocol translator. (The specific server architectures may still require a middle-tier node for other purposes.)

The following diagram illustrates the proposed architecture:




For the operating system operations, the architecture requires a few critical components in addition to the protocol stack. Those include a driver and an IPMI provider to expose BMC data and operations to the operating system. Windows Server 2003 R2 provides the customers with the first "industrial strength" implementation of this architecture.

WS-Management can also be used for remote Windows Management Instrumentation (WMI) where Distributed Component Object Model (DCOM) communication is not possible (for example, across a firewall). This is possible since WS-Management uses a single, firewall-friendly port configurable by the system administrator.

# How to Bundle ADAM with Your Application

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory provides scalable LDAP directory services both for directory-enabled applications and for management of Microsoft® Windows® network infrastructure. Active Directory Application Mode (ADAM) is a new mode of Active Directory that provides LDAP directory services specifically for directory-enabled applications — without the infrastructure requirements of Active Directory. As a result, ADAM does not require the use of domains or domain controllers. You can install ADAM on Microsoft® Windows® XP Professional and on member servers and domain controllers running Microsoft® Windows Server™ 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. ADAM runs as a non–operating system service, and it can be installed, stopped, and restarted without requiring a reboot. You can also install multiple ADAM instances on a single computer, with each ADAM instance possessing its own independently managed schema. Because ADAM is built on the same core technology as Active Directory, directory-enabled application developers can use familiar techniques for utilizing ADAM, while reducing the infrastructure dependencies of their applications.

ADAM is an ideal choice for providing dedicated LDAP directory services to a directory-enabled application. By bundling ADAM with your directory-enabled application, you do not need to be concerned about the availability of a directory service before the installation of the application. Instead, you can include ADAM as a part of your application's installation process to ensure that the application has access to a directory service immediately upon installation.

ADAM is distributed as a set of files that are packaged into a single self-executing IExpress package. You can perform a stand-alone installation of ADAM by running the ADAM Setup executable directly in Windows. You can also incorporate ADAM Setup into the Setup of your application. You can do this by including the ADAM IExpress package in your application's Setup and starting ADAM Setup from within your application's Setup.

Creating and installing a directory-enabled application package that includes a bundled copy of ADAM Setup is a three-phase process consisting of the following tasks:

1. Creating the Windows Installer package to install the directory-enabled application.

2. Modifying the install process of the directory-enabled application to optionally install ADAM.

3. Running the directory-enabled application setup to install both the directory-enabled application and ADAM (if necessary), and to create an ADAM instance.

# ADAM Install Components

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before attempting to bundle Active Directory Application Mode (ADAM) in your application's setup program, you should be familiar with the following three main components of the ADAM install:

- Adamsp1.exe
- Adaminstall.exe
- ADAM End User License Agreement (EULA)

**Adamsp1.exe**
ADAM Setup files are compressed and bundled, by using IExpress, into a single executable called Adamsp1.exe. IExpress is a Microsoft utility that helps you create stand-alone installation packages that can be used to extract files — or to temporarily extract installation source files — before running an application's setup program. IExpress creates an output file (an .exe file) that contains all the files that you specify and the instructions to carry out the extraction of these files. The files are compressed and stored in the Microsoft cabinet (.cab) format.

| **Note** |
| --- |
| Adamsp1.exe replaces Adamsetup.exe, which installed the previous version of ADAM. |

Adamsp1.exe, the ADAM Setup file, accepts the parameters that are described in the following table.

| Parameter | Description |
| --- | --- |
| **/q** | Runs **Adamsp1** without displaying progress information on the screen. |
| **/t:** *full path* | Specifies a directory to be used as the temporary working directory for **Adamsp1**. |

**Adaminstall.exe**
After you extract the ADAM Setup files to the working directory, you can run the ADAM installation program, called Adaminstall.exe, to create an ADAM instance. Adaminstall.exe accepts the parameters described in the following table.

| Parameter | Description |
| --- | --- |
| **/adv** | Advanced options. When the install type is Replica, this switch causes the Active Directory Application Mode Setup Wizard to show a page that you can use to install the replica from media (that is, from a restored copy of a backup of an ADAM instance). |
| **/answer:** *drive:\pathname\answer.txt* | Specifies a text answer file containing ADAM installation options. This parameter allows the setup to be run in unattended mode. |
| **/strict** | Specifies that ADAM Setup should treat warnings as errors. |

**ADAM EULA**
Application setup programs that install ADAM as part of their application setup process are required to either display the ADAM EULA or provide a EULA in their application that is not less restrictive than the ADAM EULA. The ADAM EULA is provided as a text file, called AdamEula.txt, and it is one of the files that is extracted when Adamsp1.exe runs. You can extract the ADAM EULA by running Adamsp1.exe and then import the extracted ADAM EULA into the custom application setup.

# Creating a Windows Installer Package

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you write a setup program that incorporates Active Directory Application Mode (ADAM) Setup, you should first create a Windows Installer package for your own application installation. Windows Installer reduces the total cost of ownership (TCO) for your customers by enabling them to install and configure your products and applications efficiently. Windows Installer can also provide your product with new capabilities to advertise features without installing them, install products on demand, and add user customizations.

For more information about Windows Installer, see "Windows Installer" on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=3985.

The basic steps for creating a simple Windows Installer package are as follows:

1. Plan the installation.

   List the general installation layout, including the files to install, the source path, and the target path. In addition, list all registry-related operations. Place the .exe file to install and all supporting files in a specific directory. You can also store supporting files in a hierarchy of subdirectories.

2. Import a blank database.

   To create a Windows Installer package, you copy (or create with a software tool) a Windows Installer database file. A blank installation database, Schema.msi, is provided with the Microsoft Platform Software Development Kit (SDK) components for Windows Installer developers. The SDK also provides a partially blank database, Uisample.msi, which contains the suggested sequence tables and data required for a simple user interface. Copy Uisample.msi into the directory that contains the .exe file to be installed. The installation database file and the source files must be located at the root of the same directory. Otherwise, you will receive Setup errors.

3. Specify the directory structure.

   The installer stores information about the installation directory structure in the Directory table. Use the database editor, Orca — or another editor — to add information to the Directory table.

4. List the constituent components.

   List all components that are part of the installation. A component can be a set of files or resources that are added to the component table of the database.

5. Specify files and file attributes.

   Add all relevant files to the Files table.

6. Enter source media information in the Media table.

   The Media table describes the set of disks that make up the source media for the installation.

7. Define features.

   Add product features to the Feature table. The installer enables users to install and remove portions of an application's functionality, which are called Windows Installer features. When you include ADAM with your application, you can define ADAM as a separate feature or tie it to one of the other existing features. Create a placeholder feature for ADAM if you want to install ADAM as a separate feature.

8. Define feature-to-component relationships.

   Use the FeatureComponents table to define the relationships between features and components. Each feature uses one or more components, and features can share components.

9. Add registry information and shortcut properties.

   The Registry table and related tables of the installation database hold the registry information that must be written in the system registry for the application. The Shortcut table and related tables of the installation database hold information needed to install shortcuts.

10. Specify properties.

    Windows Installer properties are global variables that the installer uses during an installation. You do not have to define all properties in every package; however, a small set of properties is required. The installer sets the values of properties in a particular order of precedence.

11. Populate sequence tables.

    Several sequence tables must be populated for the application setup to run:

    - InstallExecuteSequence
    - InstallUISequence
    - AdminExecuteSequence
    - AdminUISequence
    - AdvtExecuteSequence

12. Add summary information.

    Summary information is not critical for running Setup, but summary information is essential for the package to pass validation. You can use the tool MsiInfo.exe, which is provided with the Windows Installer SDK, to set these properties.

# Bundling ADAM

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can bundle Active Directory Application Mode (ADAM) with your application by having the Windows Installer package that you create for your application optionally install ADAM and then create an ADAM instance.

If your application installation package installs ADAM as part of its installation routine, it must be able to detect which version of Windows is running on the target computer, whether ADAM is installed on that computer, and, if so, what is the version of ADAM that is installed.

If your application is being installed on Microsoft Windows Server 2003 R2, your installation package must determine whether ADAM has been installed on that computer. If it has, then your application installation package need only run %windir%\ADAM\Adaminstall.exe to create the ADAM instance to be used by your application. If on the other hand ADAM has not been installed, your application installation must use the Optional Component Manager (OCM) to install ADAM before running Adaminstall.exe. To install ADAM on a computer running Windows Server 2003 R2, use the following command:

**Sysocmgr /i:sysoc.inf /q /u:script.inf**

The specified script file must contain the following:

```
[components]
adam=on
```

If your application is being installed on a computer running another version of Windows, your installation package must determine whether ADAM SP1 is installed. If it is not, then your installation package must run Adamsp1.exe to install ADAM SP1 before running Adaminstall.exe to create your application's ADAM instance. If ADAM SP1 is already installed, then your application installation package need only run Adaminstall.exe.

To determine whether the target computer is running Windows Server 2003 R2, call the **GetSystemMetrics** API with the *nIndex* parameter set to SM_SERVERR2 (80). If the return value is 0 (zero), then the computer is not running Windows Server 2003 R2.

To determine whether ADAM is installed on the target computer, look for any value in the following system registry key:

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\ADAM_Shared\VersionMajor**

If this value is present, then ADAM is installed.

To determine whether the version of ADAM that is installed is ADAM SP1, look for a value of 1 (one) in the following system registry key:

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\ADAM_SP\SP1\Installed**

You can create an ADAM instance in one of three ways:

- Active Directory Application Mode Setup Wizard
- Unattended ADAM install
- Progress Mode install

**Active Directory Application Mode Setup Wizard**

When Adamsp1.exe is run, it installs the Active Directory Application Mode Setup Wizard. The Active Directory Application Mode Setup Wizard guides the user through creating an ADAM instance by using a series of wizard pages. To start the Active Directory Application Mode Setup Wizard from your application's Windows Installer package, you must start Adamsp1.exe and then run Adaminstall.exe from within the package. You can start Adamsp1.exe and Adaminstall.exe from your application's Windows Installer package through a custom action. The custom action can be scheduled to run at a certain point during application setup, or it can be triggered by means of an event, such as the user clicking a button. Custom actions can be set to return control upon completion, or they can be set to run asynchronously. To ensure that the user completes the ADAM installation before proceeding with the installation of your application, you must set the custom action to return control upon completion.

**Unattended ADAM Install**

An ADAM instance can be installed in an unattended (or "silent") mode that does not require input from the user. An

unattended install of ADAM is accomplished by providing a text answer file to Adaminstall.exe that contains all the necessary configuration options for the ADAM instance that is to be installed. If these options are known at the time that you write the Setup program for your directory-enabled application, a fixed version of the answer file may be included as part of your application's Setup. However, if user input will be used to modify any of the ADAM Setup options during installation of your application, the answer file must be created dynamically and provided to Adaminstall.exe after the options have been specified by the user. A typical ADAM unattended install flow includes the following six steps:

The Windows Installer package writes configuration options to a text answer file.

1. If needed, the Windows Installer package starts Adamsp1.exe and Adamsp1.exe decompresses ADAM files.

2. The Windows Installer package displays the ADAM EULA.

3. The Windows Installer package starts Adaminstall.exe.

4. Adaminstall.exe reads installation options from the answer file.

Invoking Adaminstall.exe with the **/answer** parameter causes ADAM Setup to retrieve the installation settings from the specified answer file, as shown in the following example:

**adaminstall /answer:answer.txt**

The following is a sample answer file that installs a unique ADAM instance:

```
 [ADAMInstall]

; The following line specifies to install a unique ADAM instance.
InstallType=Unique

; The following line specifies the name to be assigned to the new
; instance.
InstanceName=MyAppInstance

; The following line specifies the TCP port to use for LDAP.
LocalLDAPPortToListenOn=1129

; The following line specifies the port to use for SSL.
LocalSSLPortToListenOn=1130

; The following line specifies an application partition to create
NewApplicationPartitionToCreate="CN=Partition4,DC=cohovineyard"

; The following line specifies the directory to use for ADAM data
; files.
DataFilesPath=C:\Program Files\Microsoft ADAM\MyAppInstance \data

; The following line specifies the directory to use for ADAM log files.
LogFilesPath=C:\Program Files\Microsoft ADAM\MyAppInstance \logs
```

**Displaying the ADAM EULA**

The command shown above installs an ADAM instance silently, but it does not allow the application Setup to display the ADAM EULA. To allow the application Setup to display the ADAM EULA, you must have previously incorporated the ADAM EULA into your application's Setup. To obtain the ADAM EULA, the ADAM files must be extracted to a known location, as shown in the following example:

```
adamsp1 /t:"%temp%\adamfiles"
```

After it is extracted, the ADAM EULA can be imported into the application Setup for display. The command shown above can be used during the application Setup to extract the ADAM Setup files. If the user accepts the EULA, ADAM installation can be continued by running Adaminstall.exe.

**Customizing ADAM Options**

If you want the ADAM instance to be customized by the user during the application setup, the answer file must be written to disk with the required options before Adaminstall.exe is run. The following table lists the valid answer file parameters that are accepted by an ADAM install when it is run in unattended mode.

| Parameter | Possible Values |
| --- | --- |
| *InstallType* | Unique — Create a unique instance of ADAM.<br><br>Replica — Create an instance of ADAM by replicating all or part of an existing ADAM instance.<br><br>Tools — Install ADAM tools only.<br><br>Any other value — Return an error message. |
| *ShowOrHideProgressGUI* | Show — ADAM Setup displays progress information during installation.<br><br>Hide — ADAM Setup does not display progress information during installation. |
| *InstanceName* | A valid Domain Name System (DNS) host name label, with a maximum of 63 UCS Transformation Format 8 (UTF-8) bytes.<br><br>"" (empty string) — The instance will be named Instance *n*, where *n* is the lowest number greater than 0 such that Instance *n* is unique on the local computer.<br><br>Any other value — Return an error message. |
| *ApplicationPartitionsToReplicate* | Specifies the distinguished names of the application partitions to replicate from the source ADAM instance. The following example specifies three application partitions to replicate:<br><br>ApplicationPartitionsToReplicate="CN=my,O=partition" "DC=partition2" "CN=embed qu\"ote in DN"<br><br>To replicate all application partitions from the source ADAM instance, specify "*" as the value. Any value that is specified for *ApplicationPartitionsToReplicate* is ignored if *InstallType* is not set to Replica. |
| *ReplicationDataSourcePath* | Specifies the directory path to a restored copy of ADAM data. Any value that is specified for *ReplicationDataSourcePath* is ignored if *InstallType* is not set to Replica or if a value is not also specified for *ReplicationLogSourcePath*. |
| *ReplicationLogSourcePath* | Specifies the directory path to the log file for a restored copy of ADAM data. Any value that is specified for *ReplicationLogSourcePath* is ignored if *InstallType* is not set to Replica or if a value is not also specified for *ReplicationDataSourcePath*. |
| *ReplicateOnlySpecifiedPartitions* | Yes — Replicate only application partitions that are specified by *ApplicationPartitionsToReplicate*.<br><br>No — Replicate all application partitions from the replication source, ignoring any value that is set for *ApplicationPartitionsToReplicate*. |
| *LocalLDAPPortToListenOn* | 389, or any unused port number between 1025 and 65535, inclusive.<br><br>Any other value — Return an error message. |
| *LocalSSLPortToListenOn* | 636, or any unused port number between 1025 and 65535, inclusive.<br><br>Any other value — Return an error message. |
| *SourceServer* | Valid DNS name, network basic input/output system (NetBIOS) name, or Internet Protocol (IP) address.<br><br>Any other value — If *InstallType* is Replica, return an error message. Else, do nothing. |
| *SourceLDAPPort* | Either 389 or a number between 1025 and 65535.<br><br>Any other value — If *InstallType* is Replica, return an error message. Else, do nothing. |

| | |
|---|---|
| *NewApplicationPartitionToCreate* | A valid distinguished name — Creates an application partition with the specified name.<br><br>"" (empty string) — No new application partition will be created.<br><br>Any other value — If *InstallType* is Unique, return an error message. Else, do nothing. |
| *DataFilesPath* | A syntactically correct path name, which may include unresolved environment variables that do not contain existing ADAM files.<br><br>Any other value — Return an error message. |
| *LogFilesPath* | A syntactically correct path name, which may include unresolved environment variables that do not contain existing ADAM files.<br><br>Any other value — Return an error message. |
| *ServiceAccount* | A valid DNS domain name, followed by a backslash and the account or group name.<br><br>A valid NetBIOS domain name, followed by a backslash and the account name.<br><br>A valid user principal name (UPN).<br><br>A valid account name only (generally, not recommended, because resolving an account name that is not accompanied by a domain name requires additional processing).<br><br>Any other value — Return an error message. |
| *AddPermissionsToServiceAccount* | Yes — ADAM Setup attempts to add the **Logon as a service** right to the account that is specified as the service account.<br><br>Any other value — ADAM Setup does not attempt to add the **Logon as a service** right to the account that is specified as the service account. |
| *ServicePassword* | Any string of characters, including the empty string. |
| *Administrator* | A valid DNS domain name, followed by a backslash and the account name.<br><br>**Note**<br>Do not specify built-in groups or built-in accounts, such as DOMAIN\Administrators. Instead, if you want to specify a group, specify a domain group, such as domainname\Domain Admins, where domainname represents the name of your domain.<br><br>A valid NetBIOS domain name, followed by a backslash and the account name.<br><br>A valid UPN.<br><br>A valid account name only — (Generally not recommended, because resolving an account name that is not accompanied by a domain name requires additional processing.)<br><br>Any other value — Return an error message. |
| *ShowInAddRemovePrograms* | Show — List the ADAM instance in Add or Remove Programs.<br><br>Hide — Do not list the ADAM instance in Add or Remove Programs. |
| *ImportLDIFFiles* | The names of one or more optional Microsoft-supplied .ldf files (Ms-User.ldf, Ms-InetOrgPerson.ldf, Ms-UserProxy.ldf, and Ms-azman.ldf) that you want to import into the ADAM schema. The file names must be enclosed in double quotation marks and separated by a space, as shown in this example:<br><br>`"Ms-User.ldf" "Ms-UserProxy.ldf"` |
| *SourceUserName*<br><br>*SourcePassword* | The user name and password of an account that has administrative rights for an existing configuration set. Use these parameters when installing a replica that will be joined to the configuration set. |

| Note |
| --- |
| The *ShowAdamShortcuts* parameter is no longer supported. |

**Progress Mode Install**

When your application Windows Installer package starts ADAM Setup in unattended install mode, the user is unaware that ADAM Setup is running, and he or she may think that there is a problem with the application setup. To avoid this, you can specify **ShowOrHideProgressGUI=Show** in the answer file, as shown in the following example:

```
; The following line in the answer file will cause a progress dialog to
; be displayed during the ADAM setup.
ShowOrHideProgressGUI=Show
```

This parameter causes the ADAM install to display a progress bar during its installation The user has the opportunity to cancel the installation but otherwise can only view the progress of the ADAM installation. After the ADAM installation is complete, control is returned to your application's Windows Installer package, where the application Setup can continue.

# Identifying ADAM Install Errors

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The ADAM Setup Wizard reports errors directly to the user through messages that the wizard displays. When Adaminstall.exe is run in unattended mode, no error messages are displayed. Instead, error information is written to the registry at the following location:

**HKLM\Software\Microsoft\Windows\CurrentVersion\ADAM_Installer_Results**

The registry keys in the following table are written if the install (or uninstall) should fail for any reason.

| Key | Contents |
|---|---|
| **ADAMInstallErrorCode** | The numeric error code that caused the installer to fail. |
| **ADAMInstallErrorMessage** | A message that is associated with the install error code. |
| **ADAMInstallWarnings** | Messages that are associated with the install warnings. |
| **ADAMUninstallErrorCode** | The numeric error code that caused the uninstall operation to fail. |
| **ADAMUninstallErrorMessage** | A message that is associated with the uninstall error code. |
| **ADAMUninstallWarnings** | Messages that are associated with the uninstall warnings. |

These registry keys are not unique to a particular installation of ADAM on a given computer. Each time an ADAM instance is installed, these keys are deleted. Your application installation program should check for the presence of the **ADAMInstallErrorCode** or **ADAMUninstallErrorCode** keys to see if the ADAM install (or uninstall) was successful.

# Managing Dependencies on the ADAM Service

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you write applications as Windows services that use ADAM as a dedicated LDAP directory service, you must ensure that the required ADAM instance is started before the application service is started. The Windows Service Control Manager (SCM) maintains a database of services in the Windows registry at:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services**

Each service has a key at this location. SCM provides a mechanism for controlling the start order of services, using a predefined registry key value, **DependOnService**. **DependOnService** contains a list of the other services that should start before the current service. Your application's setup program can add the name of the required ADAM instance to its **DependOnService** registry location to force Windows to start the ADAM instance before starting the application service. By convention, the name of the ADAM instance service to add to the **DependOnService** is **ADAM**_*instancename*, where *instancename* represents the name of the ADAM instance that you provide through the Active Directory Application Mode Setup Wizard or through a text answer file.

As an alternative, developers who use the Microsoft .NET Framework class library can use the **ServiceProcessInstaller** class to establish the appropriate dependency. The installation utility InstallUtil.exe uses **ServiceProcessInstaller** to write registry values that are associated with the services that you want to install. The following code demonstrates an application that is written as a service establishing a dependency on an ADAM instance called MyAdamDirectory:

```
private void InitializeComponent()
{
this.serviceProcessInstaller1 = new
System.ServiceProcess.ServiceProcessInstaller();
this.serviceInstaller1 = new System.ServiceProcess.ServiceInstaller();
//
// serviceProcessInstaller1
//
this.serviceProcessInstaller1.Account =
System.ServiceProcess.ServiceAccount.LocalSystem;
this.serviceProcessInstaller1.Password = null;
this.serviceProcessInstaller1.Username = null;
//
// serviceInstaller1
//
this.serviceInstaller1.ServiceName = "ApplicationServiceName";
this.serviceInstaller1.ServicesDependedOn =
new string[] {"ADAM_MyAdamDirectory"};
//
// ProjectInstaller
//
this.Installers.AddRange(new System.Configuration.Install.Installer[] {
this.serviceProcessInstaller1,
this.serviceInstaller1});

}
```

Windows applications and services can query the state of the ADAM service instance at any time and issue commands to manipulate its state. The following code checks to see if the ADAM instance is stopped. If the ADAM service is stopped, it will be started.

```
System.ServiceProcess.ServiceController sc2 = new
System.ServiceProcess.ServiceController("ADAM_Instance3");
if (sc2.Status.Equals(System.ServiceProcess.ServiceControllerStatus.Stopped)){
sc2.Start();
}
```

# Uninstalling ADAM

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your application is the only client of the ADAM instance that was installed during your application's installation, your application's uninstall program should remove that instance.

By default, each installed instance of ADAM appears in **Add or Remove Programs** in **Control Panel**. To prevent the user from uninstalling the ADAM instance separately from your application, you can prevent the ADAM instance from appearing in **Add or Remove Programs** by using the **ShowInAddRemovePrograms** option in the text answer file:

```
; The following line in the answer file will cause the ADAM instance to
; not appear in the Add/Remove programs list.
ShowInAddRemovePrograms=Hide
```

In addition to being uninstalled through **Add or Remove Programs**, an instance of ADAM can be uninstalled using Adamuninstall.exe, which is located in the %windir%\ADAM directory. Adamuninstall.exe accepts the parameters in the following table.

| Parameter | Description |
|---|---|
| **/q** | Runs the ADAM uninstall program without displaying progress updates. |
| **/i:** *instancename* | Uninstalls the specified ADAM instance. |
| **/i:tools** | Uninstalls the ADAM administration tools only. |

The following command, when it is issued from your application's uninstall program, removes the ADAM instance Instance1 from the local computer:

**adamuninstall /q /i:** *Instance1*

If your application is being uninstalled from a computer running Microsoft Windows Server 2003 R2, your installation package should remove ADAM from the installed optional components if no more ADAM instances remain on the computer. To uninstall ADAM on a computer running Windows Server 2003 R2, use the following command:

**Sysocmgr /i:sysoc.inf /q /u:** *script* **.inf**

The specified script file must contain the following:

```
[components]
adam=off
```

#### More Information

- For information about Windows Installer, see "Windows Installer" on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=3985.
- For information about Active Directory Application Mode (ADAM), see "Windows Server 2003 Active Directory Application Mode" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkID=20119.
- For information about Windows Server 2003, see "Microsoft Windows Server 2003" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkID=8201.

# Step-by-Step Guide for Print Management

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Welcome to Print Management for the Microsoft® Windows Server™ 2003 R2 operating systems.

**What Is Print Management?**
Print Management is a snap-in in Microsoft Management Console (MMC) that enables you to install, view, and manage all of the printers in your organization from any computer running Windows Server 2003 R2. Print Management provides up-to-the-minute details about the status of printers and print servers on the network. You can use Print Management to install printer connections to a group of client computers simultaneously. Print Management can help you find printers that have an error condition by using filters. It can also send e-mail notifications or run scripts when a printer or print server needs attention. On printer models that provide a Web page, Print Management has access to more data, such as toner and paper levels, which you can manage from remote locations, if needed.

**Who Should Use Print Management?**
This guide is targeted at the following audiences:

- Print Administrators and Help Desk professionals.
- Information Technology (IT) planners and analysts who are evaluating the product.
- Enterprise IT planners and designers.
- Early adopters.

**Benefits of Print Management**
Print Management saves the print administrator a significant amount of time installing printers on client computers and managing and monitoring printers. Tasks that can require up to 10 steps on individual computers now can be accomplished in 2 or 3 steps on multiple computers simultaneously and remotely.

By using Print Management with Group Policy, you can automatically make printer connections available to users and computers in your organization. In addition, Print Management can automatically search for and install network printers on the local subnet of your local print servers.

**In This Guide**


- Requirements for Print Management
- Security Requirements
- Steps for Deploying and Operating Print Management
- Additional Resources


**Requirements for Print Management**
Here are some important notes about the requirements for Print Management and the print servers that you can monitor using Print Management:

- You can install Print Management only on computers running Windows Server 2003 R2.
- You can use Print Management to monitor printers that are on print servers running Microsoft® Windows® 2000 Server, Windows Server 2003, and Windows Server 2003 R2 operating systems.
- You can use Print Management to monitor multiple print servers at a time.

Print Management can display and install printer drivers on computers running the Microsoft® Windows NT® Server 4.0 operating system, but it cannot display printer drivers that are already installed. It is not possible to display forms on computers running Windows NT 4.0.

**Security Requirements**
To take full advantage of Print Management, you must be logged on as an administrator or a member of the Administrators group on the print servers you are managing.

It is good practice for administrators to use an account with restrictive permissions to perform routine, non-administrative tasks and to use an account with broader permissions only when performing specific administrative tasks.

You can open Print Management and monitor any print server and printer without administrative privileges. However, you will be unable to perform certain functions such as adding and deleting printers and printer drivers.

**Steps for Deploying and Operating Print Management**
To deploy and operate Print Management, complete the following tasks:

**Step 1: Install Print Management**
Installing Print Management is accomplished by adding or updating the print server role. Note that the computer on which Print Management is installed does not need to be a print server. Do one of the following:

- If the print server role is not installed, perform the following procedure.

  **To install the print server role**

  1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Manage Your Server**.
  2. Click **Add or remove a role**.
  3. In the Configure Your Server Wizard, click **Next** to get to the **Server Role** page.
  4. On the **Server Role** page, click **Print Server**, and then click **Next** twice. Follow the steps in the wizard to complete the installation.

- If you are running an earlier version of a Windows server operating system that has the print server role and at least one shared printer already installed, perform the following procedure.

  **To update the print server role**

  1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Manage Your Server**.
  2. Click **Update this role**.
  3. In the Print Server Role Wizard, click **Next**.
  4. Follow the steps in the wizard to complete the installation.

**Stopping Trustworthy Site Pop-up Messages**
Some printers are equipped with Web pages that provide access to additional features and functionality. Intermittent warning messages may appear from Internet Explorer as a result of communication with these Web pages. There are two ways to eliminate these messages. One method is to add each printer's Web site to the list of trusted Web sites in Internet Explorer. The second method is to turn off the Enhanced Security Configuration option in Internet Explorer by using Add or Remove Windows Components. For more information about how turning off this option affects your server's security, see "Managing Internet Explorer Enhanced Security Configuration" on the Microsoft Web site (http://go.microsoft.com/fwlink?LinkId=28735).

**Step 2: Open Print Management**
After you install both Windows Server 2003 R2 and Print Management, then open Print Management. If you want to use Print Management on a computer in another location by using Remote Desktop, the remote computer must have Windows Server 2003 R2 and Print Management installed.

**To open Print Management**

- Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Print Management**.

**Step 3: Add and remove print servers**
You can add both a local print server and network print servers to Print Management. Note that the computer on which Print Management is installed does not need to be a print server.

**Adding a Print Server**

If the server on which you use Print Management is also a print server and you want to manage the printers that it hosts by using Print Management, use the following procedure to add the print server to Print Management.

**To add print servers to Print Management**

1. In the Print Management tree, right-click **Print Management**, and then click **Add/Remove Servers**.




2. In the **Add/Remove Servers** dialog box, under **Specify print server**, in **Add server**, do one of the following:
   - Type the name.
   - Click **Browse** to locate and select the print server.
3. Click **Add to List**.
4. Add as many print servers as you want, and then click **OK**.

| **Note** |
| --- |
| You can add the local server on which you are working by clicking **Add the Local Server**. |

**To remove print servers from Print Management**

1. In the Print Management tree, right-click **Print Management**, and then click **Add/Remove Servers**.
2. In the **Add/Remove Servers** dialog box, under **Print servers**, select one or more servers, and click **Remove**.

**Step 4: View Printers**

The Print Management tree contains three places where printer information is stored: **Custom Printer Filters**, **Print Servers**, and **Deployed Printers**. The **Custom Printer Filters** folder contains the **All Printers** object, which contains a dynamic view of all of the printers on all of the servers available for managing by Print Management. All of the custom views, or filters, of printers that you create are stored in **Custom Printer Filters**.

The network printer servers that you add are stored in **Print Servers**. Every printer server automatically is given four objects that serve as filters for information about a server:

- Drivers
- Forms

- Ports
- Printers

To quickly access the **Print Server Properties** dialog box, right-click the **Drivers**, **Forms**, or **Ports** objects, and then click **Manage Drivers**, **Manage Forms**, or **Manage Ports**.

**Deployed Printers** contains a list of all of the printers located in Print Management that are managed by Group Policy objects. For more information about managing printers by using Group Policy, see Deploying Printers to Users or Computers by using Group Policy later in this guide.

**Offline Print Servers**
If a server goes offline, the printer server icon will change. You will be able to unable to manage the Drivers, Forms, Ports, and Printers objects until the server comes back online.

**Printer Details**
The details for each printer display in the results pane of the snap-in console. The results pane contains columns with values such as the printer name, queue status, jobs in the queue, printer location, and driver version. You can add and remove columns to show only certain characteristics of the printers, and you can filter among all the print servers in your organization to display only printers with specific criteria in the columns. Figure 1 shows both a filtered view and the resulting columns.

In any view, you can sort on one criterion by clicking on the heading of one of the columns.




You can also show an extended view, which shows more detail about the queue and provides access to the printer's Web page, if one is available. You can add, remove, and sort columns in extended view.

**Adding and Removing Columns**
When you click the **Printers** object in each printer server or an object in the **Custom Printer Filters** folder, you can view the columns in the results pane that contain printer details.

### To add and remove columns

1. In the Print Management tree, under a print server, right-click **Printers**, point to **View**, and then click **Add/Remove Columns**.
2. In the **Add/Remove Columns** dialog box, select the name of the column, and then click **Add** or **Remove**.

**Showing Extended View**
Extended view displays beneath the columns in the right pane, separated by a splitter bar. Extended view is useful when you want more information about the status of a print job, its owner, the number of pages, the size of the job, when it was submitted, its port, its priority, and so on.

In addition, when the printer has a Web page, extended view displays a **Printer Web Page** tab. The Web page provides details about the physical properties of the printer and specifications, and sometimes allows remote administration. For more information, see Using the printer Web page later in this guide.

### To show extended view

- In the Print Management tree, under any print server, right-click **Printers**, and then click **Show Extended View**.

**Note**

To display or hide columns in extended view, select a printer. On the **Jobs** tab, right-click the column heading row, and then click the name of the column that you want to display or hide.

**Filtering Views**

You can use the filtering feature to create custom views of printers, for example, it might be helpful to filter for printers with certain error conditions or those printers in a group of buildings regardless of the print server they use. Each view is dynamic, so the data is always up to date. All filtered views are stored in the **Custom Printer Filters** folder in the Print Management tree.

An example of a filter that may be useful is shown in Figure 2. This custom view displays all printers on the specified printer server that have a queue status other than Ready and where there are one or more jobs in the print queue.




For every filtered view that you set up, you have the option of setting notifications by e-mail or choosing to run a script. To learn how to use these options, see Troubleshoot Printers later in this guide.

**To set up and save a filtered view**

1. In the Print Management tree, right-click the **Custom Printer Filters** folder, and then click **Add New Printer Filter**. This will launch the **New Printer Filter Wizard**.
2. On the **Printer Filter Name and Description** wizard page, type a name and optional description for the printer filter. The name will appear in the **Custom Printer Filters** folder in the Print Management tree.
3. In **Description**, type an optional description.
4. To display the number of printers that satisfy the conditions of a filter, select the **Display the total number of printers next to the name of the printer filter** check box.
5. Click **Next**.
6. On the **Define a printer filter** wizard page, do the following:
    a. In the **Field** list, click the print queue or printer status characteristic.
    b. In the **Condition** list, click the condition.
    c. In the **Value** box, type a value.
    d. Continue adding criteria until your filter is complete, and then click **Next**.

7. On the **Set Notifications (Optional)** wizard page, do one or both of the following:

   - To set an e-mail notification, select the **Send e-mail notification** check box, and type one or more recipient and sender e-mail addresses. An SMTP server must be specified to route the message. Use the format *account@domain* and semicolons to separate multiple accounts.

   - To set a script to run, select the **Run script** check box, and then type the path where the script file is located. To add more arguments, type them in the **Additional Arguments** box.

8. Click **Finish**.

**Step 5: Saving a custom view**
If you use customized views of Print Management, it might be useful to save one or more views. Print Management files are *.msc files.

The following are Print Management settings that you can customize for viewing:

- Print servers added

- Custom views (filters) and notifications

- Column preferences

**Step 6: Manage Printers**
You can use Print Management to manage all of the printers in your enterprise, including printers in branch offices. You can use the same interface to control how printers are shared, update drivers, and control print queues. It is unnecessary to navigate to the individual folders for each printer on each printer server.

By using Print Management in conjunction with the Configure Your Server Wizard and Terminal Services, you can automatically search for and install network printers on a local print server in branch offices. This is helpful when branch office personnel are not trained in administrative duties. You can deploy a printer to an entire group of people or make a printer available to every computer in a particular room. For more information, see Deploying Printers to Users or Computers by using Group Policy later in this guide.

**Automatically Installing Printers to the Local Print Server**
As the administrator on the local print server, you can add and remove printers. Print Management can automatically detect all printers located on the same subnet as the print server, install the appropriate printer drivers, set up the queues, and share the printers. Unless a printer driver cannot be found, no intervention is needed.

You need to have administrator permissions on the computer where the printers will be added.

If you want to automatically detect printers for a branch office or another location beyond the subnet of your server, use Remote Desktop to log on to the remote print server and start the process from that computer. You need to have administrator permissions on the computer where the printers will be added. Once you connect to the remote server by using Remote Desktop, you can search for network printers by starting Print Management or by running **\windows\pmcsnap\fnprinters.exe** from the command line.

### To add network printers to the local printer server automatically

- In the Print Management tree, right-click *ServerName* **(local)**, click **Automatically Add Network Printers**, and then click **Start** to automatically locate and install printers found on this server's local subnet.

| **Note** |
| --- |
| You may be prompted to specify which driver you want to install for a printer, but otherwise, the process is automatic. |

**Using Print Management for Bulk Printer Management Tasks**
Bulk operations can be performed on all the printers on a particular server or all the printers under a particular filter.

The following is a list of the tasks that are possible to perform on multiple printers simultaneously:

- Pause printing
- Resume printing
- Cancel All Jobs
- List in Directory
- Remove from Directory
- Delete

#### Updating Device Drivers

Printer driver details are shown separately for each server that you have added to the Print Management snap-in console. By default, the driver name, driver version, environment, and provider columns are visible. You can add columns to show more details such as the config file path, data file path, default datatype, driver file path, manufacturer, and so on.

You can export the list of details to a text file for use outside of Print Management.

### To manage device drivers

1. In the Print Management tree, under the appropriate print server, right-click **Drivers**, and then click **Manage Drivers**.
2. Do one or more of the following:
    - To add a printer driver, click **Add**, and then complete the **Add Printer Driver Wizard**.
    - To remove a printer driver, select a driver from the list, and then click **Remove**.
    - To reinstall a printer driver, select a driver, and then click **Reinstall**.
    - To view the details for a printer driver, select the driver, and then click **Properties**.

#### Listing and Removing Printers in Active Directory

Listing printers in the Active Directory® directory service makes it easier for users to locate and install printers. After you install printers on a printer server, you can use Print Management to list them in Active Directory.

In fact, you can list more than one printer simultaneously. You may want to set up a filter to show all of the printers to list or remove, so that you can easily select all of the printers at the same time.

### To list or remove printers in Active Directory

1. In the Print Management tree, under the appropriate print server, click **Printers**.
2. In the results pane, right-click the printer you want to list or remove from the directory service, and then click **List in Directory** or **Remove from Directory**.

#### Exporting a List of Printers and Displayed Columns

You can export a text file containing the data available on any **Printers** object, which includes the currently displayed columns.

### To export a list of printers

1. In the Print Management tree, under the appropriate print server, right-click **Printers**, then click **Export List**.
2. In the **Export List** dialog box, in **File name**, type the name of the file, and then click **Save**.

#### Deploying Printers to Users or Computers by using Group Policy

Print Management can be used with Group Policy to automatically add printer connections to a computer's Printers and Faxes folder.

**Note**

Before the Deploy Printer Connection feature will work, the schema needs to be updated. For more information about how to update the schema, see Upgrading Windows 2000 Domains to Windows Server 2003. For more information about schema updates, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51166).

To do this, you use the **Deploy with Group Policy** dialog box to automatically add a printer connection setting to an existing Group Policy object (GPO) in Active Directory. When Group Policy processing runs on client computers, the printer connection settings are applied to the users or computers associated with the GPO. Printers you deploy by using this method appear in the Deployed Printers object of Print Management tree when the print server they are connected to is being monitored. You will have Read Access on the GPO to which the printers are being deployed.

This method of installing a printer is useful in a laboratory, classroom, or branch office setting where every computer in the room or office needs access to the same printer. It is also useful in large organizations, where computers and printers are often separated by function, workgroup, or department, such as marketing or human resources.

A printer connection that has been installed by using a per-user connection is available to users on any computer the user logs on to in the network. A printer connection that has been installed by using a pushed per-machine connection appears in the printers and faxes folder, ready for use by any user the next time they log on to that computer.

To enable this feature, you must use a utility called PushPrinterConnections.exe.

**Note**

The PushPrinterConnections.exe utility is only needed on computers running Windows 2000, Windows XP, and Windows Server 2003 R2 client computers.

**Important**

In order for the PushPrinterConnections.exe utility to work, you must update your Active Directory schema with the Windows Server 2003 R2 changes. For more information on these schema updates, see the Microsoft Web site (http://go.microsoft.com/fwlink?linkid=51166).

First, add the PushPrinterConnections.exe utility to a machine startup script (for per-machine connections) or to a user logon script (for per-user connections). The utility reads the settings made by the GPO containing the printer setting and adds the printer connection.

It is a good idea to use the same GPO for both the printer connection settings and the PushPrinterConnections.exe computer startup or user logon script. This ensures that only users (or computers) that receive the printer connection settings will run the PushPrinterConnections.exe utility.

**Note**

Only per-user printer connections are supported on computers running Windows 2000. Either Windows XP or Windows Server 2003 is required for per-machine printer connections.

**Important**

You must have write access to the Group Policy object to use it to manage printers. Before you install printers by using Group Policy, you must have a GPO for your printer connections settings that is assigned to the appropriate users and computers. You can use the Group Policy Object Editor or Active Directory Users and Computers to create a GPO.

For more information about using Group Policy, see "Enterprise Management with the Group Policy Management Console" on the Microsoft Web site (http://go.microsoft.com/fwlink?linkid=22814).

**To install printers to groups of users or computers by using Group Policy**

1. In the Print Management tree, under the appropriate print server, click **Printers**.
2. In the results pane, right-click the printer you want to deploy, and then click **Deploy with Group Policy**.
3. In the **Deploy with Group Policy** dialog box, click **Browse**, and then choose a Group Policy object.




4. Click **OK**.

5. To assign the printer connection setting to the GPO, do one or both of the following:

   - As a per-user setting, select the **The users that this GPO applies to (per user)** check box.

   - As a per-machine setting, select the **The computers that this GPO applies to (per machine)** check box.

6. Click **Add**.

7. Repeat steps 3 to 6 to add the printer connection setting to another GPO.

8. Click **OK**.

**To use the PushPrinterConnections.exe file**

1. Using Group Policy Management console (Gpmc.msc), right-click the GPO with your printer connections settings and click **Edit**.

2. In the Group Policy Object Editor tree, navigate to one following locations:

   - If the printer connections are deployed per-machine, go to **Computer Configuration**, **Windows Settings**, **Scripts (Startup/Shutdown)**.

   - If the printer connections are deployed per-user, go to **User Configuration**, **Windows Settings**, **Scripts (Logon/Logoff)**.

3. Right-click **Startup** or **Logon**, and then click **Properties**.

4. In the **Logon Properties** or **Startup Properties** dialog box, click **Show Files**.

5. Copy the PushPrinterConnections.exe file to this location and then close the window.

6. In the **Logon Properties** or **Startup Properties** dialog box, click **Add**.

7. Type **PushPrinterConnections.exe** in the **Script Name** box.

8. If you want to enable logging, type **–log** in the **Script Parameters** box. Log files are written to %windir%\Temp\PpcMachine.log (for per-computer connections) and %temp%\PpcUser.log (for per-user connections) on the computer on which the policy is applied.

9. Click **OK**.

| Note |
| --- |
| For per-computer connections, the printer connections will be added when the client computer restarts. For per-user connections, the printer connections will be added when the user logs on. |

| Note |
| --- |
| If you remove the printer connection settings from the GPO, the PushPrinterConnections.exe will remove the corresponding printers from the client computer on the next restart or user logon. |

**Step 7: Troubleshoot Printers**

Print Management has several features that may help you identify and resolve printer problems—even in remote locations. Setting pre-defined filters lets you easily find all printers that are not in Ready status or that have a backed up queue. Many devices, regardless of manufacturer, provide rich status information, which is readily available to Print Management. By closely monitoring the printers in your organization, you may even be able to resolve problems before they happen, such as identifying when paper or toner is low.

You can set up e-mail notifications to alert you when a printer needs attention. This is especially useful when you have printers at multiple locations with different people responsible for managing them. By using an automated system to notify you when a printer or printer server is down, the problem may be resolved sooner and less attention may need to be focused on the monitoring program.

| Note |
| --- |
| Print Management must be running for notifications to be sent or for scripts to run. |

**Using the Printer Web Page**

Multifunction printers sometimes have a Web page that is accessible though extended view in Print Management. The functionality on the printer Web page will vary depending upon both the model and the manufacturer of the printer. Some multifunction printers provide remote functions, the amount of paper in each tray, and the toner level. You may be able to delete print jobs or upgrade device drivers from the printer Web pages.

| Note |
| --- |
| If the printer is down due to a networking error, the printer Web page cannot be displayed. |

**Setting E-mail Notifications**

When you create a view, or filter for specific printer criteria, you have the option of sending an automatic e-mail notification to someone when the conditions of the filter are met. This is useful for resolving printer problems, particularly in an organization with multiple buildings and administrators.

For example, you can set up a view of all printers managed by a particular print server where the status does not equal Ready. Then, if a printer changes from the Ready status to another status, the administrator could receive a notification e-mail from Print Management.

To send these notifications, you must specify that the SMTP server forward the messages.

### To set e-mail notifications

1. In the Print Management tree, do one of the following:

   - To set a notification on an existing filtered view, right-click the custom printer filter, and then click **Properties**. Click the **Notification** tab.

   - To set a notification on a new filtered view, right-click the **Custom Printer Filters** folder, and then click **Add New Printer Filter**. Follow the instructions until you reach the **Set Notifications** page.

2. Select the **Send e-mail notification** check box.




3. Enter the following information:

   - In **Recipient e-mail address(es)**, type the e-mail address of the recipient. Use the format *account@domain*. Use semicolons to separate multiple accounts.

   - In **Sender e-mail address**, type the e-mail address of the sender. Use the format *account@domain*.

   - In **SMTP Server**, type the host name of the SMTP server that will forward the e-mail notifications.

   - In **Message**, type a text message describing the conditions of the printer problem.

**Setting Server Notifications**

In addition to setting notifications on a custom set of printers, you can set notifications on printer server objects. For example, if the server is offline, or the spooler goes down, an e-mail notification can be sent. To do this, right-click a print server object, click **Set Notifications**, and then follow steps 2 and 3 in "To set e-mail notifications."

**Printing a Test Page**

You can quickly print a test page from any printer in your organization. This is helpful for testing a printer, particularly when you are not in the same building.

### To print a test page

1. In the Print Management tree, click the appropriate print server.

2. In the results pane, right-click the printer for which you want a test page, and then click **Print Test Page**.

**Using Scripts**

When you create a view, or filter for specific printer criteria, you have the option of running a script when the conditions of the filter are met. Script notifications are defined in the **Set Notifications** dialog box. Setting script notifications are useful for resolving printer problems and troubleshooting.

For example, you could automatically run a script to restart a spooler when printers go offline. You could also automatically run a script that prints a test page. The arguments are: *UserSpecifiedArguments***ServerOffline***ServerName*

Scripts can be written in Visual Basic® Script (.vbs) or any scripting language available on the computer. The script has to be on the computer with Print Management. The script should be running with your credentials and you need to have the permissions to do whatever it is that you want the script to do.

An example of a command that you might use in a script is to start the print spooler:

**net start spooler**

**Step 8: Troubleshoot Print Management**

If you are using a firewall with Print Management, some or all of the printers on a network print server may not be displayed. To solve the problem, add File and Printer Sharing to the list of exceptions in the firewall software configuration. For example,

in Windows Firewall the setting is on the **Exceptions** tab.

**To display all of the printers on a network print server when using a firewall**

1. Once File and Printer Sharing has been added, click **Edit**.
2. In the **Edit a Service** dialog box, click **Change scope**.
3. In the **Change Scope** dialog box, select **Any computer (including those on the Internet)**.

**Additional Resources**

The following resources provide additional information about Print Management.

- For information about using scripts, see "Script Repository: Printing" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50001).
- For more information about Print Services, see "Print Services" on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50121).

# Step-by-Step Guide to Deploying ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This document is a step-by-step guide for deploying Active Directory Application Mode (ADAM).

The Active Directory® directory service in Microsoft® Windows® 2000 and in Microsoft® Windows Server™ 2003 is the fastest growing directory service for intranets and extranets, as a result of its rich integration of directory support and security, scalability, and native Lightweight Directory Access Protocol (LDAP) support. Active Directory in Windows Server 2003 builds on that success by supporting a number of new LDAP capabilities that are targeted for information technology (IT) professionals and applications developers. ADAM is one of these new capabilities. Organizations, independent software vendors (ISVs), and developers who want to integrate their applications with a directory service now have an additional capability in Active Directory that provides numerous benefits.

With this document, you will be able to set up ADAM and get it running quickly, so that you can explore some of its new and important features.

Specifically, in this scenario you perform the following tasks for this walk-through:

1. Set up the lab environment, where you install and configure ADAM.
2. Install ADAM.

Next, you explore its features for adding and managing data:

1. Set up application data so that you can use ADAM with an application.
2. Query and retrieve the application data that you imported into your ADAM instance.
3. Practice creating and managing organizational units (OUs), groups, and users in ADAM.
4. Manually add and then delete an application directory partition.
5. Grant and deny user permissions.

Finally, you practice administrative tasks:

1. Practice binding to an ADAM instance in several ways.
2. Practice performing the basic functions of using ADAM such as starting and stopping an instance of ADAM.
3. Back up, remove, and restore an ADAM instance.
4. Install replicas of ADAM.
5. Perform ADAM tasks programmatically.

| **Note** |
| --- |
| It is recommended that you first use the steps provided in this guide in a test lab environment. Step-by-step guides are not necessarily meant to be used to deploy Windows Server features without accompanying documentation. You should use this guide with discretion when using it as a stand-alone document. |

# Requirements for ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Before you start using the procedures in this guide, do the following regarding system requirements:

- Have available at least one test computer on which you can install ADAM. For the purposes of following the exercises in this guide, you can install ADAM on computers running any of the following operating systems:

  - Windows Server 2003 R2, Standard Edition
  - Windows Server 2003 R2, Enterprise Edition
  - Windows Server 2003 R2, Datacenter Edition

  The computer must have 50 MB free disk space.

  | **Note** |
  | --- |
  | You can also run ADAM on computers running Windows XP and earlier versions of Windows Server 2003. The ADAM version that runs on these operating systems is available at "Windows Server 2003 Active Directory Application Mode" on the Microsoft Web site (http://go.microsoft.com/fwlink?linkid=17797). |

- If you have not already done so, download this guide along with lab files that are required to complete all the exercises. This guide and the accompanying lab files are available at ADAM Step-by-Step Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=62861).

- Log on with an administrator account.

- For the purposes of this guide, you can install replica ADAM instances on your first test computer, or you can install them on a second computer, if you have one available.

- If you previously installed an earlier version of ADAM, you must uninstall the earlier version from the computer before installing the new version of ADAM.

# Installing ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can install an ADAM instance either by using the Active Directory Application Mode Setup Wizard or by using the ADAM unattended installation process. In the first exercise, you use the Active Directory Application Mode Setup Wizard to install ADAM. In Managing Configuration Sets, you use an unattended installation to install an ADAM replica.

| **Note** |
| --- |
| To install ADAM, you must log on to your computer using an account that belongs to the local Administrators group. |

In this exercise, you first install ADAM, and then you install an ADAM instance by using the Active Directory Application Mode Setup Wizard.

**To install ADAM**

1. To install ADAM, log on as an administrator, click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. Select the check box next to **Active Directory Services**, and then click **Details**.
4. Select the check box next to **Active Directory Application Mode (ADAM)**, click **OK**, and then click **Next**.
5. Review the message that appears. Based on the contents of message, do one of the following:
   - If the message "You have successfully completed the Windows Component Wizard" appears, click **Finish**.
   - If an error message appears, make a note of the error, click **Finish**, and then review the ADAM event messages in Event Viewer.

**To install an ADAM instance by using the Active Directory Application Mode Setup Wizard**

1. To start the Active Directory Application Mode Setup Wizard, click **Start**, point to **All Programs**, point to **ADAM**, and then click **Create an ADAM instance**. The first page of the Active Directory Application Mode Setup Wizard looks like the following:

2. On the **Welcome to the Active Directory Application Mode Setup Wizard** page, click **Next**.
3. On the **Setup Options** page, you can choose whether to install a unique ADAM instance or join an existing configuration

set. Because you are installing the first ADAM instance, click **A unique instance** (as shown in the following), and then click **Next**. Later, you will create additional ADAM instances and join them in a configuration set.




4. On the **Instance Name** page, provide a name for the ADAM instance that you are installing. This name is used on the local computer to uniquely identify the ADAM instance. For this exercise, simply accept the default name of instance1, and then click **Next**.




5. On the **Ports** page, specify the communications ports that the ADAM instance uses to communicate. ADAM can communicate using both LDAP and Secure Sockets Layer (SSL); therefore, you must provide a value for each port. For this exercise, accept the default values of 389 and 636, and then click **Next**.




**Note**

If you install ADAM on a computer where either of the default ports is in use, the Active Directory Application Mode Setup Wizard automatically locates the first available port, starting at 50000. For example, Active Directory uses ports 389 and 636, as well as ports 3268 and 3269 on global catalog servers. Therefore, if you install ADAM on a domain controller, the Active Directory Application Mode Setup Wizard provides a default value of 50000 for the LDAP port and 50001 for the SSL port.

6. On the **Application Directory Partition** page, you can create an application directory partition (or naming context) by clicking **Yes, create an application directory partition**. Or, you can click **No, do not create an application directory partition**, in which case you must create an application directory partition manually after installation. For this exercise, click **Yes, create an application directory partition**. When you create an application directory partition, you must provide a distinguished name for the new partition. For this exercise, type **o=Microsoft,c=US** as the distinguished name (as shown below), and then click **Next**.




**Note**

ADAM supports both X.500-style and Domain Name System (DNS)-style distinguished names for top-level directory partitions.

7. On the **File Locations** page, you can view and change the installation directories for ADAM data and recovery (log) files. By default, ADAM data and recovery files are installed in %ProgramFiles%\Microsoft ADAM*instancename*\data, where *instancename* represents the ADAM instance name that you specify on the **Instance Name** page. For this exercise, click **Next** to accept the default file locations.




**Important**

When installing ADAM on a computer running Windows XP, you must install these files on the same logical volume. When installing ADAM on Windows Server 2003 and Windows Server 2003 R2 in a production environment, it is recommended that you install the files on separate physical disks.

**Note**

ADAM setup installs program files and administration tools in %windir%\ADAM.

8. On the **Service Account Selection** page, you select an account to be used as the service account for ADAM. The account that you select determines the security context in which the ADAM instance runs. Unless you are installing ADAM on a domain controller, the Active Directory Application Mode Setup Wizard defaults to the Network Service account. For this exercise, click **Next** to accept the **Network service account** default. Or, if you are installing ADAM on a domain controller, click **This account**, and then select a domain user account to use as the ADAM service account.




**Note**

You can change the ADAM service account after ADAM is installed by using the Dsmgmt command-line tool. When you install ADAM on a domain controller, you must select a domain user account as the ADAM service account.

9. On the **ADAM Administrators** page, you select a user or group to become the default administrator for the ADAM instance. The user or group that you select will have full administrative control of the ADAM instance. By default, the Active Directory Application Mode Setup Wizard specifies the currently logged on user. You can change this selection to

any local or domain account or group on your network. For this exercise, click the default value of **Currently logged on user**, and then click **Next**.




10. On the **Importing LDIF Files** page, you can import into the ADAM schema two .ldf files containing **user** class object definitions. Importing these **user** class object definitions is optional. However, these object definitions are required later in this guide so, you should import these definitions now:

    a. Click **Import the selected LDIF files for this instance of ADAM**.

    b. Click **MS-InetOrgPerson.LDF**, and then click **Add**.

    c. Click **MS-User.LDF**, and then click **Add**.

    d. Click **MS-UserProxy.LDF**, click **Add**, and then click **Next**.




11. The **Ready to Install** page gives you an opportunity to review your installation selections. After you click **Next**, the Active Directory Application Mode Setup Wizard begins copying files and setting up ADAM on your computer.




12. When the Active Directory Application Mode Setup Wizard finishes installing ADAM, it displays this message: "You have successfully completed the Active Directory Application Mode Setup Wizard." When the **Completing the Active Directory Application Mode Setup Wizard** page appears, click **Finish** to close the wizard.

| **Note** |
| --- |
| If the Active Directory Application Mode Setup Wizard does not complete successfully, an error message describing the reason for the failure appears on the Summary page. |

If an error occurs in the Active Directory Application Mode Setup Wizard before the **Summary** page, you can review the error message that appears. In addition, you can click **Start**, click **Run**, and type either of the following:

**%windir%\Debug\adamsetup.log**

**%windir%\Debug\adamsetup_loader.log**

The Adamsetup.log and Adamsetup_loader.log files contain information that can help you troubleshoot the cause of an ADAM setup failure.

# Using the ADAM Administration Tools

Updated: April 4, 2006

Applies To: Windows Server 2003 R2

An ADAM instance runs as a standard user service, rather than as a system service, and it can be stopped and started through the Services snap-in in Microsoft Management Console (MMC). In addition, ADAM includes several administration tools for general administration tasks. In the following exercises, you:

- Use the Services snap-in to stop and restart your ADAM instance.
- Use ADAM ADSI Edit (ADAM-adsiedit.msc) to browse your directory.
- Configure the ADAM Schema snap-in.
- Use ADSchemaAnalyzer to produce a file that can be used to extend a schema with elements from another schema.
- Use Active Directory to ADAM Synchronizer to copy data from Active Directory to an ADAM instance.

**Stopping and Restarting an ADAM Instance**
**To stop and restart an ADAM instance by using the Services snap-in**

1. Click **Start**, point to **Administrative Tools**, and then click **Services**.
2. The ADAM instance that you just installed is listed in the details pane of the Services snap-in, along with other services on the computer. ADAM instances are listed in Services by their name, which in this case is instance1. Click the ADAM instance that you installed, as shown in the following:




3. To stop the ADAM instance, on the **Action** menu, click **Stop**.
4. After the ADAM instance is stopped, on the **Action** menu, click **Start** to restart the ADAM instance.

**Using the ADAM ADSI Edit Administration Tool**
The main administration tool for ADAM is ADAM ADSI Edit. In this exercise, you use ADAM ADSI Edit to bind to, view, and browse your ADAM instance.

**To bind to, view, and browse an ADAM instance using ADAM ADSI Edit**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
2. In the console tree, click **ADAM ADSI Edit**. The **ADAM ADSI Edit** snap-in looks like the following:




3. On the **Action** menu, click **Connect to**. The **Connection Settings** dialog box appears.

4. In **Connection name**, you can type a label under which this connection will appear in the console tree of ADAM ADSI Edit. For this connection, type:

   **ADAM demo**

5. In **Server name**, type the host or DNS name of the computer on which the ADAM instance is running.

| **Note** |
| --- |
| Because, in this exercise, ADAM is running on the local computer, you can use localhost as the server name. |

6. In **Port**, type the LDAP or SSL communication ports in use by ADAM. Or, as in this case, accept the default value of 389.

| **Note** |
| --- |
| To list the port numbers used by ADAM instances, click **Start**, point to **All Programs**, point to **ADAM**, click **ADAM Tools Command Prompt**, and then, at the command prompt, type: **dsdbutil "list instances" quit** |

7. Under **Connect to the following node**, you can connect to a well-known naming context, such as the configuration or schema directory partition, or you can specify the distinguished name of a partition to which you want to connect. For this exercise, click **Distinguished name (DN) or naming context**, and type:

   **o=Microsoft,c=US**

   This is the distinguished name of the application partition that you created during setup.

8. Under **Connect using these credentials**, click **The account of the currently logged on user**. The **Connection Settings** dialog box now looks like the following:




9. Click **OK**. The **ADAM ADSI Edit** snap-in looks like the following:




10. In the console tree, double-click **ADAM demo**, and then double-click **O=Microsoft,c=US**. The **ADAM ADSI Edit** snap-in now shows the application directory partition:




11. In the console tree, click any container to view the objects in that container. For example, click **CN=Roles**.

12. To open a different directory partition on the ADAM instance, in the console tree, click **ADAM ADSI Edit**, and then, on the **Action** menu, click **Connect to**.

13. Fill out the **Connection Settings** dialog box as shown, and then click **OK**.




The **Connection Settings** dialog box now looks like the following:




You can now browse the contents of the configuration directory partition of your ADAM instance.

14. To close ADAM ADSI Edit, on the **File** menu, click **Exit**.

**Configuring the ADAM Schema Snap-in Administration Tool**

You can use another ADAM administration tool, the ADAM Schema snap-in, to administer the ADAM schema. If you have ever used the Active Directory Schema snap-in, the ADAM Schema snap-in should look very familiar to you. Before you can use the ADAM Schema snap-in, you need to create an MMC file for it, as described in this procedure.

**To create an MMC file for the ADAM Schema snap-in**

1. Click **Start**, click **Run**, type **mmc /a**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. In **Available Standalone Snap-ins**, click **ADAM Schema**, click **Add**, click **Close**, and then click **OK**.
4. To save this console, on the **File** menu, click **Save**.
5. In **File name**, type the following, and then click **Save**.

   **%windir%\system32\adamschmmgmt.msc**

   The **ADAM Schema** snap-in looks like the following:




6. To connect to your ADAM instance through ADAM Schema, in the console tree, right-click **ADAM Schema**, click **Change ADAM Server**, and then complete the dialog box as follows:




7. Click **OK**. The ADAM Schema snap-in now looks like the following. You can browse and view the ADAM schema classes and attributes:




8. To create a shortcut for the ADAM Schema snap-in on your **Start** menu:

   a. Right-click **Start**, click **Open All Users**, double-click the **Programs** folder, and then double-click the **ADAM** folder.

   b. On the **File** menu, point to **New**, and then click **Shortcut**.

   c. In the **Create Shortcut Wizard**, in **Type the location of the item**, type **adamschmmgmt.msc**, and then click **Next**.

   d. On the **Select a Title for the Program** page, in **Type a name for this shortcut**, type **ADAM Schema**, and then click **Finish**.

**Using ADSchemaAnalyzer**

You can use ADSchemaAnalyzer to help migrate the Active Directory schema to ADAM, from one ADAM instance to another, or from any LDAP-compliant directory to an ADAM instance. You can use ADSchemaAnalyzer to load a target (source) schema, mark the elements you want to migrate, and then export them to the base ADAM schema. You can also compare the two schemas or two LDAP Data Interchange Format (LDIF) files.

| **Important** |
| --- |
| When using ADSchemaAnalyzer to create an LDIF file, you should load both a target and a base schema. Otherwise, the resulting LDIF file might not be usable by the ldifde tool |

**To create an LDIF file with ADSchemaAnalyzer**

1. Click **Start**, point to **All Programs**, point to **ADAM**, click **ADAM Tools Command Prompt**, and then, at the command prompt, type:

   **adschemaanalyzer**

2. To load a target schema, click **File**, and then click **Load target schema**, and then do one of the following:

   - To load the domain Active Directory schema as the target schema, in the dialog box, type your user name, password, and domain, and then click **OK**.

   - To load a different schema (such as the schema of an Active Directory forest or an another LDAP-compliant directory), in the dialog box, type the server name and port of the directory containing the target schema, type your user name ,password, and domain as needed, and then click **OK**.

3. To load the schema of your ADAM instance as the base schema, click **File**, click **Load base schema**, and then in **Server[:port]**, type the server name and port of the ADAM instance.

4. In the dialog box, click **OK**.

5. Click **Tools**, click **Options**, and on the **LDIF generation** tab, click **Update with references to new and present elements**.

| **Important** |
| --- |
| If this option is not selected and you proceed to create an LDIF file with the default option of **Update with references to new elements only**, the resultant LDIF file will not contain all the differences between the schemas. For example, the User class in you Active Directory schema might have Optional Attributes that are not included in the User class in your ADAM schema. If the LDIF file that was created through ADschemaanalyzer does not contain these Optional Attributes and later you attempt to synchronize data in your Active Directory forest and the ADAM configuration set into which this LDIF file has been imported, Adamsync will fail with an object violation error. |

6. In the resulting tree, mark all elements that you want to export to your base schema by right-clicking the element and selecting one of the following options:

   - **Auto** automatically marks an element as included or excluded in the export. If an element is marked as **Auto (included)**, you can right-click that element, and then click **Why auto included?** to see the reverse dependency tree for the element.

   - **Included** marks an element so that it is included in the export. ADSchemaAnalyzer marks all related elements, such as superclasses, auxClasses, must/may contains, defaultObjectCategory, and possSuperiors. ADSchemaAnalyzer includes propsets for included attributes and back-links for links.

   - **Excluded** marks an element so that it is not included in the export. You can block certain paths in the dependency graph. For example, you might want to import domainDns, but not samAccountDomain (which is an auxClass of domainDns). You can exclude a complete element, such as the samAccountDomain class, or you can exclude a relationship; for example, you can remove the auxClass reference from the domainDns class. If you exclude a relationship, any other classes that reference that element continue to include it.

   - **Present** means that the element is present on the target server. By default, the top class is marked as present.

| **Note** |
| --- |
| If later you plan to synchronize data by using Adamsync, click **Schema**, and then click **Mark all non-present elements as included**. |

7. To create the LDIF file, click **File**, and then click **Create LDIF file**. To save the created LDIF file, type in the file name and save it at an appropriate location. For example, *C:\Windows\ADAM\Differences.LDIF*.

8. To import the LDIF file into the AD LDS instance in order to update the AD LDS schema to match the AD DS schema, open the created LDIF file, copy the ldifde command created by the AD DS/LDS Schema Analyzer, (for example, `ldifde –i –u –f differences.ldf –s server:port –b username domain password –j . –c "cn=Configuration, dc=X" #configurationNamingContext`) and paste it into the command prompt. Edit the ldifde command to reflect your AD LDS server name and port, and then press ENTER.

You can use the ldifde command at the ADAM tools command prompt to import the target schema elements in the LDIF file into the base ADAM schema. The beginning of the LDIF file contains complete instructions for performing this task.

**Using Active Directory to ADAM Synchronizer**

Active Directory to ADAM Synchronizer is a command-line tool that synchronizes data from an Active Directory forest to a

configuration set of an ADAM instance. You can also use it to synchronize data between two ADAM instances.

| **Important** |
| --- |
| Active Directory to ADAM Synchronizer does not synchronize user passwords between Active Directory and ADAM. |

There are two prerequisites before Active Directory to ADAM Synchronizer can synchronize data:

- The schema objects in the ADAM instance must match the schema objects in the Windows Server 2003 Active Directory forest.
- The schema in the ADAM instance must be extended for schema objects that are required by Active Directory to ADAM Synchronizer.

| **Note** |
| --- |
| You must use the **-t** *port_number* option with **ldifde** if the ADAM instance uses a port other than the default port 389. |

**To use Active Directory to ADAM Synchronizer for the first time**

1. Click **Start**, point to **All Programs**, click **ADAM**, and then click **ADAM Tools Command Prompt** to open a command window in the ADAM directory.
2. To ensure that your ADAM schema matches the default Windows Server 2003 schema in Active Directory, use ADSchemaanalyzer to create an LDIF file that will contain the target schema elements, and then import this LDIF file into your base ADAM schema by using the **ldifde** command at the ADAM tools command prompt.

   | **Important** |
   | --- |
   | When you create an LDIF file, make sure to select the **Update with references to new and present elements** option on the ADSchemaanalyzer **Tools** menu and the **Mark all non-present elements as included** option on the **Schema** menu.<br><br>For detailed instructions, see Using ADSchemaAnalyzer. |

3. To extend the ADAM schema to include schema objects that are required by Active Directory to ADAM Synchronizer, at the command prompt, type the following command on a single line, and then press ENTER:

   **ldifde -i -s localhost:389 -c CN=Configuration,DC=X #ConfigurationNamingContext -f MS-AdamSyncMetadata.ldf**

4. Modify the configuration file MS-AdamSyncConf.xml with the appropriate parameters:
   - Replace the value of <source-ad-name> with the name of the source Active Directory domain controller, for example, <source-ad-name>SeattleDC1</source-ad-name>.
   - Replace the value of <source-ad-partition> with the distinguished name of the source domain, for example, <source-ad-partition>dc=fabrikam,dc=com</source-ad-partition>.
   - Replace the value of <source-ad-account> with the name of an account in the Domain Admins group of the source domain, for example, <source-ad-account>administrator</source-ad-account>.
   - Replace the value of <account-domain> with the fully qualified name of the source domain, for example, <account-domain>fabrikam.com</account-domain>.
   - Replace the value of <target-dn> with the name of the partition of the target ADAM instance, in this case, <target-dn>DC=microsoft,DC=US</target-dn>.
   - Replace the value of <base-dn> with the base distinguished name of the source domain, for example, <base-dn>dc=fabrikam,dc=com</base-dn>.
   - Modify the query filter (the default being <object-filter>(objectClass=*)</object-filter>), depending on what objects you want to synchronize.

   | **Important** |
   | --- |

Do not delete any unused fields from this file.

**Note**

It is not necessary to synchronize an entire domain naming context. To save disk space and avoid synchronization problems, consider excluding objects and attributes that are not necessary to ADAM (for example, DNS records, FRS subscriptions, and DN-binary values), and edit your MS-AdamSyncConf.xml file appropriately. For more information, see Adamsync Configuration File XML Reference (http://go.microsoft.com/fwlink/?LinkId=119621).

5. Install the configuration file. At a command prompt, type the following command, and then press ENTER:

   **ADAMSync /install localhost:389 %windir%\ADAM\MS-AdamSyncConf.xml**

6. Synchronize the data from the Active Directory forest to the ADAM configuration set. At a command prompt, type the following command, and then press ENTER:

   **ADAMSync /sync localhost:389 "DC=microsoft,DC=US" /log**

   The **/log** option displays detailed information about the status of the synchronization. You can also use ADAM ADSI Edit to verify that the data has been synchronized.

# Setting Up Application Data

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

In most cases, you extend the ADAM schema with object class and attribute definitions for the kinds of data that you want an application to store. Just as with Active Directory, the schema in ADAM is extensible. You can extend the ADAM schema programmatically or with the Ldifde.exe command-line tool.

In the following exercises, you:

- Step 1: Add optional user classes to the ADAM schema.
- Step 2: Extend the ADAM schema to support an application.
- Step 3: Import application data into an ADAM instance.

| **Note** |
| --- |
| You will use the application data that you import in these exercises later with the Windows Address Book application. |

**Step 1: Adding Optional User Classes to the ADAM Schema**

You can add the optional user classes that are provided with ADAM during ADAM setup, or you can add them manually using the Ldifde.exe command-line tool. If you imported the user class definition .ldf files when you ran the Active Directory Application Mode Setup Wizard, you can skip this procedure.

**To manually add optional user classes to the ADAM schema**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
2. At the command prompt, type the following command, and then press ENTER:

   **ldifde -i -f ms-inetorgperson.ldf -s** *servername:portnumber* **-k -j . -c "CN=Schema,CN=Configuration,DC=X" #schemaNamingContext**

   where *servername:portnumber* represents the computer name and LDAP communication port of your ADAM instance. Because the ADAM instance is running on your local computer, you can also use localhost as the computer name.

   | **Note** |
   | --- |
   | Be sure to use the copy of Ldifde.exe that came with Windows Server 2003 R2, rather than a copy that came with an earlier ADAM release or with Windows Support Tools. |

   The **ADAM Tools Command Prompt** window should now look like the following:

   ```
   ADAM Tools Command Prompt
   C:\WINDOWS\ADAM>ldifde -i -f ms-inetorgperson.ldf -s localhost:389 -k -j . -c "c
   n=schema,cn=configuration,dc=x" #schemanamingcontext
   Connecting to "localhost:389"
   Logging in as current user using SSPI
   Importing directory from file "ms-inetorgperson.ldf"
   Loading entries.....................................................................
   ...
   67 entries modified successfully.

   The command has completed successfully

   C:\WINDOWS\ADAM>
   ```

3. Type the following command, and then press ENTER:

   **ldifde -i -f ms-user.ldf -s** *servername:portnumber* **-k -j . -c "CN=Schema,CN=Configuration,DC=X" #schemaNamingContext**

The **ADAM Tools** Command Prompt window now looks like the following:

```
ADAM Tools Command Prompt

C:\WINDOWS\ADAM>ldifde -i -f ms-user.ldf -s localhost:389 -k -j . -c "cn=schema,
cn=configuration,dc=x" #schemanamingcontext
Connecting to "localhost:389"
Logging in as current user using SSPI
Importing directory from file "ms-user.ldf"
Loading entries.....................................................................
...
67 entries modified successfully.

The command has completed successfully

C:\WINDOWS\ADAM>_
```

After you run these commands, the ADAM schema includes the ADAM and inetOrgPerson user object classes. You can verify this by viewing the ADAM schema with the ADAM Schema snap-in.

**Step 2: Extending the ADAM Schema to Support an Application**

In this exercise, you extend the ADAM schema again by adding a contacts object class, again by using the ldifde command-line tool.

**To extend the ADAM schema**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
2. At the command prompt, type the following, and then press **ENTER**:

   **ldifde -i -f** *drive:\path* **\labs_demo\labs\contact.ldf -s** *servername:portnumber***-k -j . -c "CN=Schema,CN=Configuration" #schemaNamingContext**

   where *drive:\path* represents the location where you saved the ADAM download, and *servername:portnumber* represents the computer name and LDAP communications port of your ADAM instance. Because the ADAM instance is running on your local computer, you can also use localhost as the computer name.

| **Note** |
| --- |
| Be sure to use the copy of Ldifde.exe that came with the ADAM release, rather than a copy that came with an earlier ADAM release or with Windows Support Tools. |

After the command runs, the **ADAM Tools Command Prompt** window looks like the following:

```
ADAM Tools Command Prompt

C:\WINDOWS\ADAM>ldifde -i -f c:\labs_demo\labs\contact.ldf -s localhost:389 -k -
j . -c "cn=schema,cn=configuration" #schemanamingcontext
Connecting to "localhost:389"
Logging in as current user using SSPI
Importing directory from file "c:\labs_demo\labs\contact.ldf"
Loading entries.....
4 entries modified successfully.

The command has completed successfully

C:\WINDOWS\ADAM>_
```

Now your ADAM schema also includes the contacts object class, and it is ready for some application data.

**Step 3: Importing Application Data into an ADAM Instance**

In this exercise, you import some sample data into your ADAM instance using the ldifde command-line tool. This data is provided with the ADAM download.

**To import application data**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
2. At the command prompt, type the following:

   **ldifde -i -f** *drive:\path* **\labs_demo\labs\contactimport.ldf -s** *servername:portnumber* **-k -j .**

where *drive:\path* represents the location where your ADAM files are located, and *servername:portnumber* represents the computer name and LDAP communications port of your ADAM instance.

| **Note** |
| --- |
| Be sure to use the copy of Ldifde.exe that came with the ADAM release, rather than a copy that came with an earlier ADAM release or with Windows Support Tools. |

After the command runs, the **ADAM Tools Command Prompt** window looks like the following:




| **Note** |
| --- |
| Ldifde always reports the number of entries that are imported into the directory. In this case, you can see from the ldifde output that contactimport.ldf contains two records. For more information about ldifde, type **ldifde /?** at the command prompt. |

# Using an Application with ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

In this exercise, you use Windows Address Book, an LDAP-based application, to query and retrieve the application data that you imported into your ADAM instance.

**Querying Data with Windows Address Book**
**To query data with Windows Address Book**

1. Click **Start**, click **Run**, type **wab.exe**, and then click **OK**.
2. On the **Tools** menu, click **Accounts**, and then click **Add**. The Internet Connection Wizard appears.
3. In **Internet directory (LDAP) server**, type **localhost**, click **Next** twice, and then click **Finish**.
4. In the **Internet Accounts** dialog box, on the **Directory Service** tab, double-click **localhost**, as shown in the following:




5. Complete the **General** tab in the **localhost Properties** dialog box as shown in the following figure, using the account with which you are currently logged on to the computer. Be sure to select both the **This server requires me to log on** check box and the **Log on using Secure Password Authentication** check box.




6. Click the **Advanced** tab, and then select the **Use simple search filter** check box. In **Search base**, type **o=Microsoft,c=US**as shown in the following, click **OK**, and then click **Close** in the **Internet Accounts** dialog box.




7. In Windows Address Book, click **Find People** on the toolbar. In **Look in**, click **localhost**, and then, in **Name**, type:

   **kim**

   The **localhost Properties** dialog box looks like the following:




8. Click **Find Now**. Your search should return one entry from the ADAM directory, as shown in the following:




9. Double-click the name to view the search result details.

10. On the **Organization** tab, you can view the reporting relationship. Double-click the manager's name to view the associated contact details, which looks like the following:



# Managing OUs, Groups, and Users in ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

ADAM is used most often to store information about users and the organizations and other groups they belong to. In these exercises, you create an organizational unit (OU) called "ADAM users" in the o=Microsoft,c=US application directory partition and add a group in ADAM called "ADAM testers," and you create an ADAM user named Mary Baker with one of the user object classes that you imported earlier. Using ADAM ADSI Edit, you:

- Step 1: Create an OU.
- Step 2: Create a group in the new OU.
- Step 3: Create an ADAM user.
- Step 4: Add an ADAM user to the ADAM users group.

In addition you learn how to enable and disable ADAM user accounts.

**Step 1: Create an OU**
In this exercise, you create an OU.

**To create an OU**

1. If it is not open already, open ADAM ADSI Edit, and then connect to the o=Microsoft,c=US application directory partition, as described in the procedure "To bind to, view, and browse an ADAM instance using ADAM ADSI Edit" in Using the ADAM Administration Tools.
2. In the console tree, right-click **O=Microsoft,c=US**, point to **New**, and then click **Object**. The **Create Object** dialog box looks like the following:




3. In the **Select a class** list, click **organizationalUnit**, and then click **Next**.
4. In **Value**, type **ADAM users**, and then click **Next**.
5. On the next page, you can click **More attributes** to edit additional attributes on the object that you are creating. For this exercise, simply click **Finish**.
6. In the console tree, double-click **O=Microsoft,c=US**. The **ADAM ADSI Edit** snap-in looks like the following:




**Step 2: Create a group**
In this exercise, you create a group in the OU.

**To create a group in an OU**

1. In the console tree, right-click **OU=ADAM Users**, point to **New**, and then click **Object**.
2. In **Select a class**, click **group**, and then click **Next**.
3. In **Value**, type **ADAM testers**, and then click **Next**.
4. In **Value**, type **2147483650** (equivalent to 0x80000002 hexadecimal, which signifies an account group), click **Next**, and then click **Finish**.

| **Note** |
| --- |
| For more information about the **groupType** attribute, see "Group-Type" on the Microsoft Web site (http://go.microsoft.com/fwlink?linkid=51093). |

The **ADAM ADSI Edit** snap-in looks like the following:




**Step 3: Create an ADAM user**
In this exercise, you create an ADAM user in the ADAM Users OU, and then you add the user to the ADAM Testers group.

| **Note** |
| --- |

| |
|---|
| The new user account is disabled by default because it has no associated password. |

**To create an ADAM user**

1. If it is not already open, open ADAM ADSI Edit.
2. Connect and bind to your ADAM instance, as described in the procedure "To bind to, view, and browse an ADAM instance using ADAM ADSI Edit" in Using the ADAM Administration Tools. Then, in the console tree, double-click the ADAM instance.
3. Double-click the **O=Microsoft,c=US** application directory partition.
4. Right-click the **OU=ADAM Users** container that you created previously, point to **New**, and then click **Object**.
5. In **Select a class**, click **user**, and then click **Next**.

| **Note** |
|---|
| If you did not close ADAM ADSI Edit before importing the Adamuser.ldf user class object definitions, you may receive the following warning message during this step: "An invalid directory pathname was passed." |

6. In **Value**, type **Mary Baker** as the common name (cn) for the new user, as shown below, and then click **Next**.




7. Click **Finish**. The **ADAM ADSI Edit** snap-in looks like the following:



**Step 4: Add a user to a group**

You can add both ADAM users and Windows users to ADAM groups, as described in this exercise. First, you add Mary Baker, the user that you just created, to the ADAM testers group.

**To add a user to a group**

1. In the details pane of ADAM ADSI Edit, right-click **CN=ADAM testers**, and then click **Properties**. The **CN=ADAM testers Properties** dialog box looks like the following:




2. In **Attributes**, click **Member**, and then click **Edit**.

3. Click **Add ADAM Account**, type the following as the distinguished name, and then click **OK**:

   **CN=Mary Baker,OU=ADAM users,O=Microsoft,C=US**

   The **Multi-valued Distinguished Name with Security Principal Editor** dialog box looks like the following:




4. You can also add Windows users to an ADAM group. In the **Multi-valued Distinguished Name With Security Principal Editor** dialog box, click **Add Windows Account**. The **Select Users, Computers, or Groups** dialog box looks like the following:




5. In the **Select Users, Computers, or Groups** dialog box, add a Windows user from your computer or domain to the ADAM testers group. In **Enter the object names to select (examples)**, type an account name using the computer\account or domain\account format.

6. Click **OK**. The new user name appears in the **Multi-valued Distinguished Name With Security Principal Editor** dialog box as a member of the group.

7. Click **OK** twice to return to ADAM ADSI Edit.

**Disabling and Enabling ADAM User Accounts**

You can disable and enable ADAM user accounts by using the ADAM ADSI Edit snap-in. In this exercise, you disable the Mary Baker account and then enable it again.

**To enable or disable an ADAM user account**

1. In ADAM ADSI Edit, connect and bind to an ADAM instance as described in the procedure "To bind to, view, and browse an ADAM instance using ADAM ADSI Edit" in Using the ADAM Administration Tools.

2. In the console tree, double-click the **O=Microsoft,c=US** application directory partition.

3. In the console tree, click the **OU=ADAM Users** container.

4. In the details pane, right-click **CN=Mary Baker**, and then click **Properties**.

5. In **Attributes**, click **msDS-UserAccountDisabled**, and then click **Edit**.

6. Click **True**, and then click **OK**. The Mary Baker account is now disabled.

7. To enable the Mary Baker account, edit **msDS-UserAccountDisabled** again, and this time set the attribute to **False**.

# Managing Directory Partitions in ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The following exercises help you become familiar with an additional ADAM administration tool, Ldp.exe. Ldp is installed as part of the ADAM administration tool set. In these exercises, you use Ldp to connect and bind to an ADAM instance, and then you use Ldp to manually add and then delete an application directory partition. (Remember, you can also create an application directory partition by using the ADAM setup wizard.)

**Connecting and Binding to an ADAM Instance Using Ldp.exe**
To begin this exercise, you connect and bind to your ADAM instance using Ldp.exe.

**To connect and bind to an ADAM instance using Ldp.exe**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
2. At the command prompt, type **ldp**, and then press ENTER.
3. On the **Connection** menu, click **Connect**.
4. In **Server**, type the host or DNS name of the computer running ADAM. When the ADAM instance is running locally, you can also type **localhost**.
5. In **Port**, type the LDAP or SSL communication port number for the ADAM instance to which you want to connect, and then click **OK**.




6. On the **Connection** menu, click **Bind**.
7. Do one of the following:
   - To bind using the credentials you logged on with, click **Bind as currently logged on user**.
   - To bind using a domain user account, click **Bind using credentials**,type the user name, password, and domain name (or the computer name, if you are using a local workstation account) of the account that you are using, and then click **OK**, as in the following.




   - To bind using just a user name and password, click **Simple bind**, type the user name and password of the account that you are using, and then click **OK**.
   - To bind using an advanced method (NTLM, DPA, negotiate, or digest), click **Advanced (*method*)**, click **Advanced**, in **Method**, select the desired method, set other options as needed, and then click **OK** twice.

8. When you are finished specifying bind options, click **OK**.

**Adding an Application Directory Partition**
Now, you are ready to add an application directory partition.

**To add an application directory partition using Ldp.exe**

1. On the **Ldp Browse** menu, click **Add child**.
2. In **Dn**, type **cn=test,o=testpartition,c=us** as the distinguished name for the new application directory partition.
3. Under **Edit Entry**, type the following, and then click **Enter**:
   - In **Attribute**, type **ObjectClass**.
   - In **Values**, type **container**.
4. Under **Edit Entry**, type the following, and then click **Enter**:
   - In **Attribute**, type **InstanceType**.
   - In **Values**, type **5**.

   The **Add** dialog box looks like the following:




5. Click **Run**. After the new application directory partition is added, the following result appears in the details pane:

```
***Calling Add...
ldap_add_s(ld, "cn=test,o=testpartition,c=us", [2] attrs)
Added {cn=test,o=testpartition,c=us}.
```

6. Click **Close**.
7. To refresh Ldp and view your new directory partition, you must disconnect and then bind again to the ADAM instance. On the **Connection** menu, click **Disconnect**.
8. Bind to your ADAM instance as you did previously. On the **Connection** menu, click **Bind**.
9. To view the directory tree in Ldp, on the **View** menu, click **Tree**.
10. To view all directory partitions on the ADAM instance, leave **BaseDN** blank, and then click **OK**. The **Ldp** window looks like the following:




11. To view your new directory partition and its default containers and objects, double-click **CN=test,O=testpartition,C=US** in the console tree. The **Ldp** window looks like the following:




**Deleting an Application Directory Partition**

In this exercise, you delete the application directory partition that you just created.

**To delete an application directory partition using Ldp.exe**

1. In the Ldp console tree, double-click the configuration directory partition **CN=Configuration,CN=**{*GUID*}, where *GUID* is the unique identifier that is assigned by ADAM.

2. To view the cross-reference objects for the directory partitions on your ADAM instance, in the console tree, double-click the partitions container **CN=Partitions**. The **Ldp** window looks like the following:




3. In the console tree, under the partitions container **CN=Partitions**, double-click the cross-reference object for which the value of **nCName** (as viewed in the details pane) is equal to CN=test,O=testpartition,C=US, as shown in the following:

| **Note** |
| --- |
| To clear the details pane in Ldp without disturbing your bind or connection, on the **Connection** menu, click **New**. |




4. To delete this cross-reference object (and, therefore, the associated directory partition), in the console tree, right-click the appropriate cross-reference object in the partitions container, click **Delete**, and then click **OK**.

| **Caution** |
| --- |
| You cannot undo a partition deletion after you click **OK**. |

After you delete the cross-reference object, output similar to the following appears in the details pane:

```
ldap_delete_s(ld, "CN=56c5aea2-5cb1-450a-96f0-
5622cd949791,CN=Partitions,CN=Configuration,CN={90BF4692-0FF5-4410-8835-
DCBBEE6E08B1}");
Deleted "CN=56c5aea2-5cb1-450a-96f0-5622cd949791,CN=Partitions,CN=Configuration,CN=
{90BF4692-0FF5-4410-8835-DCBBEE6E08B1}"
```

| **Note** |
| --- |

For more information about Ldp, see ADAM Help. To open ADAM Help, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Help**.

# Managing Authorization in ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Authorization refers to the process of determining which users have access to which directory objects. As with Active Directory, access control lists (ACLs) on each directory object determine which users have access to that object. By default, the only ACLs in ADAM reside in the top-level container of each directory partition. All objects in a given directory partition inherit these ACLs. Using the Dsacls.exe command-line tool, you can view and modify the default ACLs in ADAM, and you can add additional ACLs. In the following exercises, you view and modify ADAM ACLs.

| **Note** |
| --- |
| You may have directory-enabled applications that implement their own custom authorization schemes. These applications generally disregard the ACLs on ADAM directory objects. |

**Viewing Effective Permissions**

In this exercise, you view the effective permissions on the o=Microsoft,c=US directory partition.

**To view effective permissions**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
2. At the command prompt, type the following, and then press ENTER:

   **dsacls ** *servername:portnumber* **\O=Microsoft,C=US**

   where *servername:portnumber* is the computer name and the LDAP communications port of your ADAM instance.

This command lists all the permissions that are currently set on the directory partition object. Your screen should contain output similar to the following:

```
Access list:
Effective Permissions on this object are:
Allow CN=Instances,CN=Roles,CN=Configuration,CN={C98CA450-AC25-4BC1-AC3C-C3BEC88B335E}
                                      SPECIAL ACCESS
                                      READ PERMISSONS
                                      LIST CONTENTS
                                      READ PROPERTY
                                      LIST OBJECT
Allow CN=Readers,CN=Roles,O=Microsoft,C=US
                                      SPECIAL ACCESS
                                      READ PERMISSONS
                                      LIST CONTENTS
                                      READ PROPERTY
                                      LIST OBJECT
Allow CN=Administrators,CN=Roles,O=Microsoft,C=US
                                      FULL CONTROL
Allow CN=Instances,CN=Roles,CN=Configuration,CN={C98CA450-AC25-4BC1-AC3C-C3BEC88B335E}
                                      Replicating Directory Changes
Allow CN=Instances,CN=Roles,CN=Configuration,CN={C98CA450-AC25-4BC1-AC3C-C3BEC88B335E}
                                      Replication Synchronization
Allow CN=Instances,CN=Roles,CN=Configuration,CN={C98CA450-AC25-4BC1-AC3C-C3BEC88B335E}
                                      Manage Replication Topology
Allow CN=Instances,CN=Roles,CN=Configuration,CN={C98CA450-AC25-4BC1-AC3C-C3BEC88B335E}
                                      Replicating Directory Changes All

Permissions inherited to subobjects are:
Inherited to all subobjects
Allow CN=Readers,CN=Roles,O=Microsoft,C=US
                                      SPECIAL ACCESS
                                      READ PERMISSONS
                                      LIST CONTENTS
                                      READ PROPERTY
                                      LIST OBJECT
Allow CN=Administrators,CN=Roles,O=Microsoft,C=US
```

```
                                        FULL CONTROL

The command completed successfully
```

**Granting Permissions**
In this exercise, you grant the Delete permission on the ADAM testers group object to the Mary Baker account.

**To grant the Delete permission**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.

2. At the command prompt, type the following:

   **dsacls "** *servername:portnumber* **\CN=ADAM testers,OU=ADAM users,O=Microsoft,C=US" /G "CN=Mary Baker,OU=ADAM users,O=Microsoft,C=US":SD;;**

   where *servername:portnumber* represents the computer name and LDAP communications port of your ADAM instance. Be sure to use an uppercase G when typing the **/G** parameter, and use quotation marks as shown.

Your screen should contain output similar to the following:

```
Access list:
Effective Permissions on this object are:
Allow CN=Mary Baker,OU=ADAM users,O=Microsoft,C=US
                                        SPECIAL ACCESS
                                        DELETE
Allow CN=Readers,CN=Roles,O=Microsoft,C=US
                                        SPECIAL ACCESS   <Inherited from parent>
                                        READ PERMISSONS
                                        LIST CONTENTS
                                        READ PROPERTY
                                        LIST OBJECT
Allow CN=Administrators,CN=Roles,O=Microsoft,C=US
                                        FULL CONTROL   <Inherited from parent>

Permissions inherited to subobjects are:
Inherited to all subobjects
Allow CN=Readers,CN=Roles,O=Microsoft,C=US
                                        SPECIAL ACCESS   <Inherited from parent>
                                        READ PERMISSONS
                                        LIST CONTENTS
                                        READ PROPERTY
                                        LIST OBJECT
Allow CN=Administrators,CN=Roles,O=Microsoft,C=US
                                        FULL CONTROL   <Inherited from parent>

The command completed successfully
```

**Denying Permissions**
In this exercise, you deny Delete permissions for the currently logged on user in the ADAM testers group. This is done in two phases:

- Deny delete permissions on the parent container of the ADAM testers group

- Deny delete permissions on the group itself

**To deny the Delete permissions on the parent container of a group**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.

2. To deny the Delete, Delete Child, and Delete Tree permissions on the parent container of the ADAM testers group, which is the ADAM users OU. At the command prompt, type the following:

   **dsacls "** *servername:portnumber* **\OU=ADAM users,O=microsoft,C=US" /D** *domain\administrator* **:SDDCDT;;**

   where *servername:portnumber* represents the computer name and LDAP communications port of your ADAM instance, and *domain\administrator* represents the account with which you are currently logged on. Be sure to use an uppercase D when typing the **/D** parameter, and use quotation marks as shown.

Your screen should contain output similar to the following:

```
Access list:
Effective Permissions on this object are:
Deny  domain\account                  SPECIAL ACCESS
                                      DELETE
                                      DELETE CHILD
                                      DELETE TREE
Allow CN=Readers,CN=Roles,O=Microsoft,C=US
                                      SPECIAL ACCESS   <Inherited from parent>
                                      READ PERMISSONS
                                      LIST CONTENTS
                                      READ PROPERTY
                                      LIST OBJECT
Allow CN=Administrators,CN=Roles,O=Microsoft,C=US
                                      FULL CONTROL   <Inherited from parent>

Permissions inherited to subobjects are:
Inherited to all subobjects
Allow CN=Readers,CN=Roles,O=Microsoft,C=US
                                      SPECIAL ACCESS   <Inherited from parent>
                                      READ PERMISSONS
                                      LIST CONTENTS
                                      READ PROPERTY
                                      LIST OBJECT
Allow CN=Administrators,CN=Roles,O=Microsoft,C=US
                                      FULL CONTROL   <Inherited from parent>

The command completed successfully
```

**To deny delete permissions on the group**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
2. To deny the Delete permission on the ADAM testers group for the currently logged on user, at the command prompt, type the following:

   **dsacls "** *servername:portnumber* **\CN=ADAM testers,OU=ADAM users,O=microsoft,C=US" /D domain\administrator:SDDCDT;;**

   where *servername:portnumber* represents the computer name and LDAP communications port of your ADAM instance, and *domain\administrator* represents the account with which you are currently logged on. Be sure to use an uppercase D when typing the **/D** parameter, and use quotation marks as shown.

Your screen should contain output similar to the following:

```
Access list:
Effective Permissions on this object are:
Deny  domain\account                  SPECIAL ACCESS
                                      DELETE
Allow CN=Readers,CN=Roles,O=Microsoft,C=US
                                      SPECIAL ACCESS   <Inherited from parent>
                                      READ PERMISSONS
                                      LIST CONTENTS
                                      READ PROPERTY
                                      LIST OBJECT
Allow CN=Administrators,CN=Roles,O=Microsoft,C=US
                                      FULL CONTROL   <Inherited from parent>

Permissions inherited to subobjects are:
Inherited to all subobjects
Allow CN=Readers,CN=Roles,O=Microsoft,C=US
                                      SPECIAL ACCESS   <Inherited from parent>
                                      READ PERMISSONS
                                      LIST CONTENTS
                                      READ PROPERTY
                                      LIST OBJECT
Allow CN=Administrators,CN=Roles,O=Microsoft,C=US
                                      FULL CONTROL   <Inherited from parent>

The command completed successfully
```

# Managing Authentication in ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

With ADAM, you can bind as a Windows principal, as an ADAM principal, or through an ADAM proxy object. In the following exercises, you:

- Complete a bind as a Windows principal.
- Set a password for the ADAM user account Mary Baker, which you created earlier.
- Complete a bind as an ADAM principal.
- Complete a bind through an ADAM proxy object.

In addition, you test the permissions that you set by using Dsacls.exe command-line tool in the exercises in Managing Authorization in ADAM.

**Binding as a Windows Principal**

In this exercise, you bind to an ADAM instance as a Windows principal and then test the bind.

**To bind as a Windows principal and test the bind**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
2. Using ADAM ADSI Edit, bind to your ADAM instance using the Windows principal that you are logged on as, and connect to the O=Microsoft,c=US directory partition.
3. In the details pane, browse to the ADAM testers group, on which you denied the Delete permission to your current Windows account.
4. Right-click the ADAM testers group, and then click **Delete**. An "Access denied" message appears, confirming that the Delete permission has been successfully denied to your Windows account.

**Setting the Password of an ADAM User**

Before logging on to the ADAM instance with the Mary Baker user account, you first set a password on the account.

| **Note** |
| --- |
| In addition to using Ldp as described in this procedure, you can also use ADAM ADSI Edit to set or modify passwords: right-click the directory object representing the ADAM security principal in ADAM ADSI Edit, and then click Reset Password. |

**To set a password on an ADAM user account**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and click **ADAM Tools Command Prompt**.
2. At the command prompt, type **ldp**, and then press **ENTER**.
3. On the **Connection** menu, click **Connect**, and then connect to your ADAM instance.
4. On the **Options** menu, click **Connection Options**.
5. In **Option Name**, click **LDAP_OPT_SIGN**, type **1** in **Value**, and then click **Set**.
6. In **Option Name**, click **LDAP_OPT_ENCRYPT**, type **1** in **Value**, click **Set**, and then click **Close**.
7. On the **Connection** menu, click **Bind**, and then bind to your ADAM instance.
8. On the **View** menu, click **Tree**, leave **BaseDN** blank, and then click **OK**.
9. In the console tree, locate the O=Microsoft,C=US directory partition. Double-click **O=Microsoft,C=US**, and then double-click **OU=ADAM Users,O=Microsoft,C=US**.

10. Right-click the **CN=Mary Baker** user object, and then click **Modify**. The following dialog box appears:




11. In **Attribute**, type **userpassword**, and then, in **Values**, type a password for the account.

12. Click **Enter**, and then click **Run**. The details pane in Ldp should contain output similar to the following:

```
***Call Modify...

ldap_modify_s(ld, 'CN=Mary Baker,OU=ADAM users,O=Microsoft,C=US',[1] attrs);

Modified "CN=Mary Baker,OU=ADAM users,O=Microsoft,C=US".
```

| **Note** |
| --- |
| When ADAM runs on a computer running Windows Server 2003, it enforces the password policy and account lockout settings of the computer, organizational unit, or domain, whichever is in effect. |

**Binding as an ADAM Principal**

In this exercise, you bind to an ADAM instance as an ADAM principal and then test the bind.

**To bind as an ADAM principal and test the bind**

1. Using Ldp, bind to your ADAM instance using CN=Mary Baker,OU=ADAM users,O=Microsoft,C=US as the account, along with the password that you just assigned to this account.

2. To confirm that you are logged on as Mary Baker and that the Delete permission that you granted earlier is effective, in the Ldp console tree, browse to the ADAM testers group and delete it. To delete the ADAM testers group, right-click the **CN=ADAM testers** object, and then click **Delete**.

| **Note** |
| --- |
| By default, new ADAM users (such as Mary Baker) are granted Read access to the top-level container of a given directory partition, a permission which is inherited by all objects on the partition. But, because you explicitly assigned the Delete permission to Mary Baker on the ADAM testers group object, the delete operation succeeded. For more information about access control and default permissions in ADAM, see ADAM Help. To view ADAM Help, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Help**. |

**Binding Through an ADAM Proxy Object**

In addition to binding as a Windows user or as an ADAM user, you can also bind to an ADAM instance by using ADAM bind redirection. When using bind redirection, ADAM can accept and process bind requests to an ADAM proxy object that contains as one of its attributes the security ID (SID) from an Active Directory security principal. With ADAM, you can use bind redirection to provide Active Directory users with access to both ADAM data and Active Directory data, using Active Directory domain credentials as a single sign-on (SSO). In addition, you can use ADAM proxy objects to store user data that is specific to a particular application in ADAM, while using Active Directory to store more widely used directory data.

Bind redirection enables a user to bind to ADAM by means of a simple bind while still using Active Directory credentials. Other types of binding with Active Directory credentials work without requiring a proxy, but a simple bind does not. Proxy binding works only for a simple bind.

The ADAM .ldf files, which you can import into the ADAM schema during ADAM setup, contain an object definition for the object userProxy, which can be used for bind redirection. This object contains attributes that include a distinguished name and a SID. By creating a userProxy object in ADAM—specifying a distinguished name to be used for binding—and by using a valid SID from an Active Directory user account, you can bind to ADAM using bind redirection. For more information about ADAM authentication, see "Active Directory Application Mode" Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=51640).

For the following exercises, it is assumed that you have already imported the optional user classes into the ADAM schema.

**Binding Security and ADAM Proxy Objects**
By default, binding to ADAM with bind redirection requires an SSL connection. SSL requires the installation and use of certificates on the computer running ADAM and on the computer connecting to ADAM as a client. If you do not have certificates installed in your ADAM test environment, you can, as an alternative, disable the requirement for SSL, as described in the following procedure.

| **Note** |
| --- |
| Disabling the requirement for SSL for bind redirection causes the password of a Windows security principal to be passed to the computer running ADAM, without first being encrypted. Therefore, you should only disable the SSL requirement in a test environment. |

**To disable the SSL requirement for bind redirection**

1. As described earlier in the procedure "To bind to, view, and browse an ADAM instance using ADAM ADSI Edit," connect and bind to your ADAM instance using ADAM ADSI Edit, and then, in the console tree, browse to the following container object in the configuration partition: CN=Directory Service,CN=Windows NT,CN=Services.
2. Right-click **CN=Directory Service**, and then click **Properties**.
3. In **Attributes**, click **msDS-Other-Settings**, and then click **Edit**.
4. In **Values**, click **RequireSecureProxyBind=1**, and then click **Remove**.
5. In **Value to add**, type **RequireSecureProxyBind=0**, click **Add**, and then click **OK**.

**Creating and Binding with an ADAM Proxy Object**
In these exercises, you create a proxy object for an Active Directory user, and you bind to ADAM using the proxy object.

**To bind to ADAM through an ADAM proxy object**

1. As described earlier in the procedure "To connect and bind to an ADAM instance using Ldp.exe," connect and bind to your ADAM instance using Ldp, and then browse to O=Microsoft,C=US.
2. On the Ldp **Browse** menu, click **Add child**.
3. In **Dn**, type **cn=testproxy,o=microsoft,c=us** as the distinguished name for the new userProxy object to be created in the O=Microsoft,C=US container.
4. Under **Edit Entry**, type the following, and then click **Enter**:
   - In **Attribute**, type **ObjectClass**
   - In **Values**, type **userProxy**
5. Again, under **Edit Entry**, type the following, and then click **Enter**:
   - In **Attribute**, type **objectSID**
   - In **Values**, type the valid SID of a user in Active Directory.

     The \LABS_DEMO\LABS\bindredirect directory in the ADAM download contains two commands from the Windows

Server 2003 Administration Tools Pack, Dsquery.exe and Dsget.exe, to help you retrieve the SID of an Active Directory user. You can run these commands on a computer running Windows Server 2003.

To retrieve the SID of an Active Directory user with these commands, type the following (as a single command) at a command prompt:

**dsquery user -samid** *domain\account* **| dsget user -sid**

where *domain\account* represents the user whose SID you want to retrieve. In this command, the results of Dsquery are piped to Dsget.

You can retrieve the SID of the currently logged on user on a computer running Windows Server 2003 by typing the following at a command prompt:

**whoami /user**

(Some versions of **whoami** require the syntax **whoami /user /sid**.)

6. Click **Run**. This adds the userProxy object, with the attributes that you specified, to the ADAM directory store.
7. To disconnect from your ADAM instance, on the **Connection** menu, click **Disconnect**.

Now, you can bind to your ADAM instance using the ADAM proxy object and bind redirection.

**To bind as an ADAM proxy object through bind redirection**

1. On the **Connection** menu, click **Connect**, and then connect to your ADAM instance on a new connection.
2. On the **Options** menu, click **Connection Options**.
3. In **Option Name**, click **LDAP_OPT_SIGN**, type **1** in **Value**, and then click **Set**.
4. In Option Name, click **LDAP_OPT_ENCRYPT**, type **1** in **Value**, click **Set**, and then click **Close**.
5. To bind to your ADAM instance again with Ldp, on the **Connection** menu, click **Bind**.
6. In **User**, type:

   **cn=testproxy,o=Microsoft,c=us**

   This represents the proxy object that you just created.
7. Make sure that the **Domain** option is not selected.
8. In **Password**, type the password that is associated with the Active Directory user that you specified in step 5 in the previous procedure, and then click **OK**.

**Demonstrating ADAM Proxy Object Functionality**

By default, a Windows user binding to an ADAM instance receives membership only in the ADAM groups to which that user has been explicitly added as member. When a user binds to an ADAM instance through a proxy object, the user receives membership in the Users group on each naming context that is held by the ADAM instance.

You can use this difference in group memberships to demonstrate the functional difference between binding to an ADAM instance as a Windows user and binding to an ADAM instance through a proxy object. The following exercise demonstrates this difference.

**To demonstrate binding to ADAM through a proxy object**

1. In the O=Microsoft,C=US directory partition, add the Users group as a member of the Readers group, following the general directions for adding members to groups as described earlier in the procedure "To add a user to a group."
2. Bind to your ADAM instance (using Ldp or ADAM ADSI Edit) as an Active Directory user (other than the ADAM administrator, which receives full access to all partitions by default).
3. Attempt to read any object in the O=Microsoft,C=US directory partition. Your attempt should fail, because the Active

Directory user does not have access to the partition by default.

4. Bind to your ADAM instance (using Ldp or ADAM ADSI Edit) using the proxy object that you created.

5. Attempt to read any object in the O=Microsoft,C=US directory partition. This time, your attempt should succeed; because users who bind to an ADAM instance through a proxy object automatically receive membership in the Users group. And, because you added the Users group to the Readers group in step 1 of this procedure, binding to the ADAM instance through the proxy object enables you to successfully read the partition.

| **Note** |
| --- |
| For more information about bind redirection, see ADAM Help. To view ADAM Help, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Help**. For information about administering proxy objects programmatically, see Administering ADAM Programmatically later in this guide. |

# Backing Up and Restoring Active Directory Application Mode (ADAM)

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

In the following exercises, you back up your ADAM instance. Then, you remove ADAM completely from your computer. Finally, you restore your ADAM instance back to your computer.

**Backing up an ADAM Instance**
In this exercise, you back up your ADAM instance.

**To back up an ADAM instance**

1. Click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
2. In the **Backup or Restore Wizard**, click the **Advanced Mode** link.
3. Click the **Backup** tab, and then, on the **Job** menu, click **New**.
4. Select the check box to the left of your ADAM instance folder, which is, by default, %programfiles%\Microsoft\ADAM\instance1.
5. To back up the ADAM files to a file, in **Backup destination**, click **File**. (If you do not have a tape drive in your computer, **File** is selected by default.) Then, in **Backup media or file name**, type a path and file name for the backup (.bkf) file. Your screen should now appear similar to the following:




6. Click **Start Backup**, and then make any desired changes in the **Backup Job Information** dialog box.
7. If you want to set advanced backup options, such as data verification or hardware compression, in the **Backup Job Information** dialog box, click **Advanced**. When you have finished setting advanced backup options, click **OK**.
8. Click **Start Backup**. After the backup operation is complete, close the backup application.

**Removing an ADAM Instance**
To simulate an accidental loss of an ADAM instance, you can uninstall your ADAM instance, which removes both the ADAM program files and the ADAM data files.

**To uninstall an ADAM instance**

1. Click **Start**, point to **Control Panel**, click **Add or Remove Programs**, and then click **ADAM Instance instance1**. If it is the first item in the list, **ADAM Instance instance1** will already be selected.

2. Click **Remove**, and then click **Yes**. Active Directory Application Mode is now removed from your computer.

**Restoring an ADAM Instance**

In this exercise, you restore your ADAM instance from the backup that you made in the previous exercise.

**To restore an ADAM instance**

1. Create an ADAM instance following the steps in Installing ADAM. Use the same settings as you did during your first ADAM installation except, in this case, do not create an application directory partition during setup. You can restore your original application directory partition from your backup. Therefore, on the **Application Directory Partition** page in the Active Directory Application Mode Setup Wizard, click **No, do not create an application directory partition**.

2. Click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

3. Click the **Advanced Mode** link in the **Backup or Restore Wizard**.

4. Click the **Restore and Manage Media** tab. To select the ADAM instance that you want to restore, in the details pane, select the check box to the left of the instance1 folder, as shown in the following:




5. In **Restore files to**, click **Original location**.

6. On the **Tools** menu, click **Options**, click the **Restore** tab, click **Always replace the file on my computer**, and then click **OK**.

7. Click **Start Restore**.

8. If you receive a message asking if you want to restart your computer, click **Yes**.

9. After the restore is complete, close the Backup application.

10. To confirm that the data from your original ADAM instance is successfully restored, use ADAM ADSI Edit to confirm that the O=Microsoft,C=US directory partition is restored and that the OU=ADAM users OU and the Mary Baker user account exist in the partition.

# Managing Configuration Sets

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

In the following exercises, you create new ADAM instances by replicating your existing ADAM instance. By doing so, you also create an ADAM configuration set. ADAM instances in a configuration set replicate a common schema partition and configuration partition, and they can also replicate application directory partitions (such as O=Microsoft,C=US) to each other.

In the following exercises, you install two replica ADAM instances. You create the first replica instance by using the Active Directory Application Mode Setup Wizard. You create the second replica instance by using unattended installation. You then configure the replication schedule for your configuration set.

| **Note** |
| --- |
| In a production environment, ADAM instances belonging to the same configuration set cannot reside on the same computer. You can have multiple ADAM instances running on a computer, but they must belong to different configuration sets. However, for the purposes of this guide, if you do not have a second computer available, you can install your replica ADAM instances on your first computer. |

**Installing a Replica Using the Active Directory Application Mode Setup Wizard**

You can install an ADAM replica instance by using the Active Directory Application Mode Setup Wizard.

**To install an ADAM instance replica by using the Active Directory Application Mode Setup Wizard**

1. Start the Active Directory Application Mode Setup Wizard, either on your second computer (if you have one) or on your first computer: click **Start**, click **All Programs**, point to **ADAM**, and then click **Create an ADAM instance**. Follow the steps in the wizard until you reach the **Setup Options** page.
2. On the **Setup Options** page, click **A replica of an existing instance**, and then click **Next**.
3. On the **Instance Name** page, accept the default name instance2 (or instance1, if you are installing ADAM on a second computer), and then click **Next**.

   | **Note** |
   | --- |
   | ADAM instance names need to be unique only on a given computer. |

4. On the **Ports** page, accept the default values of 50000 and 50001 (if you are installing onto the first computer) or 389 and 636 (if you are installing onto a second computer), and then click **Next**.
5. On the **Joining a Configuration Set** page, in **Server**, type the host name or DNS name of the computer where the first ADAM instance is installed. Then, type the LDAP port number in use by the first ADAM instance (which is 389 by default), and then click **Next**.

   | **Note** |
   | --- |
   | You must use a valid host name or DNS name, rather than an IP address or localhost when specifying a server on the **Joining a Configuration Set** page of the Active Directory Application Mode Setup Wizard. |

6. On the **Administrative Credentials for the Configuration Set** page, click the account that is used as the ADAM administrator for your first ADAM instance.
7. On the **Copy Application Partition** page, select the application directory partitions that you want to replicate to the new ADAM instance. (The schema and configuration partitions will be replicated automatically.) To select the O=Microsoft,C=US directory partition for replication, in **Available partitions**, click **O=Microsoft,C=US**, and then click **Add**. The **Active Directory Application Mode Setup Wizard** looks like the following:




8. Click **Next**.

9. Accept the default values on the remaining Active Directory Application Mode Setup Wizard pages by clicking **Next** on each page, and then click **Finish** on the Completing the Active Directory Application Mode Setup Wizard page.

10. After the installation is complete, use ADAM ADSI Edit to confirm that the O=Microsoft,C=US directory partition has been replicated to your second ADAM instance.

**Installing a Replica from Media by Using Unattended Installation**

In addition to using the Active Directory Application Mode Setup Wizard, you can also install an ADAM instance replica from media. For this type of installation, you must use a restored copy of an ADAM backup as the media and unattended installation as the installation method.

First, you restore your ADAM instance backup to an alternate location, rather than to the original location, so that you do not overwrite your first ADAM instance. Then you create an answer file and perform the unattended installation.

**To install an ADAM instance replica by using unattended installation**

1. Click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

2. Use the Backup application to restore the backup of your original ADAM instance, as you did previously. But this time, in **Restore files to**, click **Alternate location** instead of **Original location**, and then type or browse to an alternate directory path to which you want to restore the files, as shown in the following.




3. After restoring the backup files, create an answer file for the ADAM unattended installation. An answer file provides the values for the ADAM setup options (the same options that are provided in the Active Directory Application Mode Setup Wizard). Using any text editor, create a text file called Answer.txt, and then add the following contents to the file. Be sure to replace *servername* with the host name or DNS name of the computer on which your first ADAM instance is running. Replace C:\media_install\Program Files\Microsoft\ADAM\instance1\data with the path to your restored copy of the first ADAM instance.

```
[ADAMInstall]
; The following line specifies to install a replica ADAM instance.
InstallType=Replica
; The following line specifies the name to be assigned to the new instance.
InstanceName=instance3
; The following lines specify the communication ports to use for LDAP and SSL.
LocalLDAPPortToListenOn=50002
LocalSSLPortToListenOn=50003
; The following lines specify the directory location of the restored files.
ReplicationDataSourcePath=C:\media_install\Program
Files\Microsoft\ADAM\instance1\data
ReplicationLogSourcePath=C:\media _install\Program
Files\Microsoft\ADAM\instance1\data
; The following lines specify a computer name and ADAM port of an ADAM instance in
the
; configuration set you want to join
; Replace servername with the name of the computer on which your first ADAM
; instance is running
SourceServer=servername
SourceLDAPPort=389
```

4. After saving your Answer.txt file, you are ready to run the unattended installation. At a command prompt, type the following:

   *Drive:\path* **\adamsetup /c:"adaminstall.exe /answer:** *drive:\pathname* **\answer.txt"**

   where the first *drive:\path* is the location of your ADAM download and the second *drive:\pathname* is the location of the Answer.txt file that you created.

5. After running this command, you can confirm by using the Services snap-in that a new ADAM instance is installed and running.2

**Configuring the Replication Schedule**

Now that you have multiple ADAM instances joined in a single configuration set, you can schedule replication. Scheduling replication is optional. As with Active Directory, ADAM always provides a default replication schedule.

**To schedule replication between ADAM instances**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.

2. Connect and bind to one of your ADAM instances.

| **Note** |
| --- |
| Because all of your ADAM instances belong to the same configuration set, you can schedule replication on any one of them. |

3. In the console tree, double-click the configuration partition **CN=Configuration,CN={GUID}**, where **GUID** is the unique identifier assigned during ADAM setup; double-click the sites container, CN=Sites; and then double-click the default sites container, **CN=Default-First-Site-Name**. The **ADAM ADSI Edit** snap-in looks like the following:




| **Note** |
| --- |
| By default, all ADAM instances that you create belong to a single site, Default-First-Site-Name. In this exercise, all your ADAM instances belong to a single site. Therefore, you are scheduling replication within a site, which is called intrasite replication. For more information about ADAM sites, configuration sets, and replication, see the Active Directory Application Mode Administrator's Guide. To view the Active Directory Application Mode Administrator's Guide, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Help**. |

4. In the details pane, right-click **CN=NTDS Site Settings**, and then click **Schedule**.

5. In the **Schedule** dialog box, select the block of time that you want to schedule; click **None**, **Once per Hour**, **Twice per Hour**, or **Four Times per Hour** as the replication frequency; and then click **OK**. The **Schedule** dialog box looks like the following:




**Note**

For intrasite replication, ADAM instances replicate changes through update notifications. The replication frequency schedule only affects intrasite replication when no update notifications occur in the specified time.

**Causing Immediate Replication of a Directory Partition**

The Active Directory Application Mode Setup Wizard installs an ADAM version of Repadmin.exe, which includes the same functionality as the Active Directory version of Repadmin.exe. As with Active Directory, you can cause the immediate synchronization of a directory partition with replication partners by using Repadmin.exe, as described in this exercise.

**To cause immediate replication of a directory partition using Repadmin.exe**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
2. At the command prompt, type the following:

   **repadmin /syncall localhost:389 o=Microsoft,c=us**

**Note**

For more information about **repadmin** syntax, at the ADAM Tools Command Prompt, type **repadmin /?**.

# Administering ADAM Programmatically

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can accomplish programmatically many of the tasks that you can complete manually using the ADAM administration tools. The ADAM download includes several sample scripts and some sample code to help get you started.

**Administering ADAM Programmatically Through Visual Basic Scripts**
The \LABS_DEMO\LABS\VBScript directory in the ADAM download includes sample scripts that are produced in Microsoft® Visual Basic®, Scripting Edition (VBScript), for the following common operations:

- Extend schema to include the contact class. (Adamcontact.vbs)
- Import two contacts. (AdamContactImport.vbs)

The following scripts assume that Adamuser.ldf was imported in a previous exercise:

- Add OU
- Add user
- Add group
- Add user to group
- Delete user
- Get a list of specific objects in a path (Filter_adam.vbs)
- Enumerate users and groups
- Set password

For example, the script for enumerating users and groups contains the following code:

```
'**************************************************
'
' This script enumerates the users and groups in the passed in OU
' To run: cscript member_adam.vbs [OU] [Group]
' Examples: cscript member_adam.vbs ou=testou,c=us testuser
'
'**************************************************
set Args  = Wscript.Arguments
ouName  = Args(0)
' If the application OU DN is "ou=adamou,c=us" and the server is "adamhost" and the
port is 389. Then this parameter should be passed
' as follows:  "LDAP://adamhost:389/ou=adamou,c=us"

set ou = GetObject(ouName )
wscript.echo "Displaying Groups and Group membership..." & vbcrlf

ou.Filter = Array("group")
for each obj in ou
  wscript.echo "Group : "  & obj.Name
  for each member in obj.Members
    wscript.echo "         |"
    wscript.echo "          -- " & member.Name
  Next
  wscript.echo vbcrlf
Next
```

You can run any of these scripts from a command prompt, using the **cscript** command. (For help with **cscript**, at a command prompt, type **cscript /?**.) Each script requires that the distinguished names of both the provider and the host be passed, along with the port specifier.

| **Note** |
| --- |
| The Adamcontact.vbs script only requires *servername*:*portnumber*to be passed, because it extends the schema. You can open the file in Notepad to see the specific syntax. (If you run a script without parameters, the following error message is returned: "Subscript out of range.") |

For example, to run the Member_adam.vbs script to enumerate users and groups of an object with a distinguished name of O=Microsoft,C=US, type:

**cscript member_adam.vbs "LDAP:// *servername* : *portnumber* /o=Microsoft,c=us"**

where *servername*:*portnumber* represents the computer name and LDAP communications port of your ADAM instance.

**Administering ADAM Programmatically Through the System.DirectoryServices API**
The following exercise requires that you have Microsoft® Visual Studio® .NET installed.

**To access ADAM through the System.DirectoryServices Application Programming Interface (API)**

1. Start Visual Studio .NET.
2. On the **File** menu, click **New**, and then click **Project**.
3. In **Project Types**, click a project type (**C#**, **VB.NET**, and so on).
4. In **Templates**, click a project template (**Console**, **Windows**, and so on).
5. In **Name**, type a name for your project.
6. After the project is created, click **Add Reference** on the **Project** menu.
7. In the **Component Name** column, click **System.DirectoryServices.dll**, as shown in the following.




8. Add the following line at the top of your code:

```
C#:
using System.DirectoryServices;
VB.NET:
Imports System.DirectoryServices;
```

| **Note** |
| --- |

Adding the namespace name is not mandatory, but it is easier than typing a long name. For example instead of System.DirectoryServices.DirectoryEntry, use DirectoryEntry.

9. To read an ADAM object, add the following code:

```
int portNumber=1025;  // put the correct port number here.
String serverName="adam01"; // put the correct servername here.
String partitionDir = "O=Fabrikam"; //put the correct partition distinguished name.
DirectoryEntry ent = new
     DirectoryEntry("LDAP://"+serverName+":"+portNumber+"/"+partitionDir);
Console.WriteLine("Hello World, {0}, with Guid {1}", ent.Name, ent.Guid);
```

**Administering ADAM Proxy Objects Programmatically**

The \LABS_DEMO\LABS\bindredirect directory in the ADAM download includes sample code for creating, populating, and testing ADAM proxy objects. In addition, the directory includes a compiled, ready-to-run version of this sample code. This sample code illustrates how you can automate the creation of proxy objects, and it completes the steps in the "To bind to ADAM through an ADAM proxy object" procedure in Using the ADAM Administration Tools.

**Note**

For more information about ADAM bind redirection, see the Active Directory Application Mode Administrator's Guide. To view the Active Directory Application Mode Administrator's Guide, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Help**.

The code in sampleBindRedirect.c completes all of the following operations programmatically:

- Binds to an ADAM instance using a Windows user account that you provide.
- Reads the **tokenGroups** attribute for the Windows user to retrieve the user's SID.
- Binds to an ADAM instance using the ADAM Administrator's account that you provide.
- With the ADAM administrator account, creates a userProxy object for the Windows user.
- Adds the Users group from any given application directory partition to the Readers group of the same partition.
- Binds to an ADAM instance as the Windows user, to demonstrate that the Windows user cannot read the application directory partition.
- Binds to an ADAM instance through the proxy object, to demonstrate that the application directory partition can be read.
- Deletes the userProxy object.

You can run the compiled version of this sample code, BindRedirect.exe, to observe how the sample code works. For help running the BindRedirect.exe sample program, at a command prompt, type **bindredirect /?**.

**Note**

This sample code runs with the following requirements:

To run properly, SSL connections to ADAM must be available (which requires the installation of certificates), or the RequireSecureProxyBind attribute on the msds-Other-Settings attribute of nTDsService object must be set to 0. For more information, see "Binding Security and ADAM Proxy Objects" in Managing Authentication in ADAM.

No foreign security principal object should exist in ADAM for the Windows user that you specify.

When using an SSL connection and binding, you must provide the full DNS name of the computer running ADAM.

# Introduction to Storage Manager for SANs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Storage Manager for SANs is a new Microsoft Management Console (MMC) snap-in that helps you create and manage logical unit numbers (LUNs) on Fibre Channel and iSCSI disk drive subsystems that support Virtual Disk Service (VDS) in your storage area network (SAN).

Before deploying Storage Manager for SANs on a production server, we recommend that you use this guide to first familiarize yourself with the different LUN management options in Storage Manager for SANs by installing and running it on a test server.

# Storage Manager for SANs Requirements

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

To use Storage Manager for SANs, the server and the storage subsystems must meet the following requirements:

- The server must be running a Microsoft® Windows Server™ 2003 R2 operating system with Storage Manager for SANs and Virtual Disk Service (VDS) installed.
- The storage subsystems that you want to administer using Storage Manager for SANs must support VDS. For more information about VDS, see Virtual Disk Service Technical Reference at the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=50517).
- The VDS hardware provider for each storage subsystem must be installed on the server.
- To use Storage Manager for SANs with an iSCSI subsystem, Microsoft iSCSI Software Initiator Version 2.0 must be installed on the server. For information about how to download and install Microsoft iSCSI Software Initiator Version 2.0, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=44352).

# Installing Storage Manager for SANs and Other Required Software

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

To test Storage Manager for SANs, you must install Windows Server 2003 R2 and Storage Manager for SANs. Next, you need to install the hardware provider for each of your storage subsystems. To use Storage Manager for SANs with an iSCSI subsystem, you also need to install Microsoft iSCSI Software Initiator Version 2.0.

To install and use Storage Manager for SANs, you must be logged on under an account that has administrative privileges on the local computer.

**To install Storage Manager for SANs**

1. Install Windows Server 2003 R2.

2. On the same server, install Storage Manager for SANs:

    a. In Control Panel, click **Add or Remove Programs**, and then click **Add/Remove Windows Components**.

    b. From the list of components in the **Windows Components Wizard** dialog box, select **Management and Monitoring Tools**, and click **Details**.

    c. From the list of subcomponents, select **Storage Manager for SANs**, and then click **OK**.

    d. Click **Next**, and after the configuration of the new component has completed, click **Finish**. Microsoft Virtual Disk Service (VDS) version 1.1 is installed with Storage Manager for SANs.

3. If you will be using Storage Manager for SANs to create LUNs on an iSCSI subsystem, install Microsoft iSCSI Software Initiator Version 2.0 on the server. This installs the latest version of the iSCSI initiator that is designed to work with Storage Manager for SANs. For information about how to download and install Microsoft iSCSI Software Initiator Version 2.0, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=44352).

4. On the server, install the VDS hardware providers for each of your storage subsystems.

# Using Storage Manager for SANs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section explains how to open and use Storage Manager for SANs. The following tasks are included:

- Opening Storage Manager for SANs
- Managing storage on a remote computer

**Opening Storage Manager for SANs**

Storage Manager for SANs is a Microsoft Management Console (MMC) snap-in and is available in Administrative Tools.

**To open Storage Manager for SANs from Administrative Tools**

1. On the **Start** menu, click **Programs**, click **Administrative Tools**, and then click **Storage Manager for SANs**.
2. Configure Storage Manager for SANs as appropriate. There are three main nodes:
   - **LUN Management** lists all the LUNs created with Storage Manager for SANs and allows you to create new LUNs, extend the size of a LUN, assign and unassign LUNs, and delete LUNs. In this node, you can also configure the Fibre Channel and iSCSI connections that servers in your SAN will use to access the LUNs that you create.
   - **Subsystems** lists all the storage subsystems currently discovered in your SAN and enables you to change the subsystem name that was returned by VDS.
   - **Drives** lists all the drives in the subsystems discovered in your SAN. To help identify the drives you are working with, you can make the light on a drive blink.

**Managing storage on a remote computer**

To manage storage resources on a remote computer, you can connect to the computer from Storage Manager for SANs. While you are connected, the results pane displays the objects created on the remote computer, which allows you to manage them from MMC.

The remote computer must be running Windows Server 2003 R2 with Storage Manager for SANs installed.

To connect to a remote computer, you must be logged on under a domain account that has administrative rights on the remote computer.

**To connect to a remote computer from Storage Manager for SANs**

1. In the console tree, right-click **Storage Manager for SANs**, and click **Connect to another computer**.
2. In the **Connect to Another Computer** dialog box, select **Another computer**, type the name or browse to the server you want to connect to, and then click **OK**.

# Overview of LUN Management in Storage Manager for SANs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use Storage Manager for SANs to create and manage LUNs on both Fibre Channel and iSCSI disk storage subsystems in your SAN.

A LUN is a logical reference to a portion of a storage subsystem. A LUN can comprise a disk, a section of a disk, a whole disk array, or a section of a disk array in the subsystem. Using LUNs simplifies the management of storage resources in your SAN because they serve as logical identifiers through which you can assign access and control privileges.

Because of hardware, protocol, and security differences, LUN configuration and management on Fibre Channel and iSCSI environments is different. This section explains those differences.

**Managing LUNs in a Fibre Channel environment**
In a Fibre Channel environment, LUNs created on a Fibre Channel disk storage subsystem are assigned directly to a server or cluster, which accesses the LUN through one or more Fibre Channel host bus adapter (HBA) ports. You only need to identify the server or cluster that will access the LUN, and then select which HBA ports on that server or cluster will be used for LUN traffic.

When a server or cluster is identified, Storage Manager for SANs will automatically discover the available Fibre Channel HBA ports on that server or cluster. You can also add ports manually by entering their World Wide Name (WWN).

**Managing LUNs in an iSCSI environment**
Unlike in a Fibre Channel environment, LUNs created on an iSCSI disk storage subsystem are not only assigned to a server or cluster. For iSCSI, LUNs are first assigned to logical entities called targets.

Targets are created in order to manage the connections between an iSCSI device and the servers that need to access it. A target defines the portals (IP addresses) that can be used to connect to the iSCSI device, as well as the security settings (if any) that the iSCSI device requires in order to authenticate the servers requesting access to its resources.

| **Note** |
| --- |
| In most cases, you can create and manage targets yourself. However, some iSCSI storage subsystems only support simple target configurations, where targets are automatically created when you create a LUN. With simple target configurations, you also cannot delete a target or manually assign LUNs to it. LUNs are automatically assigned when they are created. For this type of subsystem, you just need to identify the server or cluster that will access the LUN, and the iSCSI subsystem will enable access from that server or cluster to the LUN. |

To connect to a target, a server in the SAN uses an iSCSI initiator. An iSCSI initiator is a logical entity that enables the server to communicate with the target. The iSCSI initiator first logs on to the target. After access is granted by the target, the server can start reading and writing to all LUNs assigned to that target. Each iSCSI initiator can have one or more network adapters through which communication is established.

As with Fibre Channel environments, you only need to identify the server or cluster that will access the LUN. Storage Manager for SANs automatically discovers the iSCSI initiators on that server or cluster, and lists all the available adapters in those initiators. After the iSCSI adapters have been discovered, you can select which adapters will be used for LUN traffic.

For more information, see Microsoft Storage Technologies - iSCSI at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50522).

# Support for Multipath I/O

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Multipathing increases availability by providing multiple paths (path failover) from a server or cluster to a storage subsystem.

If a server in your SAN supports Microsoft Multipath I/O (MPIO), you can enable more than one read/write path to a logical unit number by assigning multiple Fibre Channel ports or multiple iSCSI adapters on that server to the same LUN. The same is true when accessing the LUN from a cluster.

| **Caution** |
| --- |
| To prevent data loss, ensure that the server or cluster supports MPIO before enabling access through multiple paths to the same LUN. If the server or cluster does not support MPIO, or if you are not sure, assign only a single Fibre Channel port or iSCSI adapter to each LUN. |

For more information, see Microsoft Storage Technologies - Multipath I/O at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50521).

# Deploying LUNs Using Storage Manager for SANs

Updated: April 25, 2007

Applies To: Windows Server 2003 R2, Windows Server 2008

This section walks you through the tasks involved in using Storage Manager for SANs to deploy LUNs on both Fibre Channel and iSCSI subsystems. Figure 1 shows the process for deploying LUNs on both types of subsystems.




If the LUNs will be deployed on a Fibre Channel subsystem, you only need to identify the servers that will access the LUNs, select which HBA ports on those servers will be used for LUN traffic, and finally assign LUNs directly to the servers.

If the LUNs will be deployed on an iSCSI subsystem, the process requires additional steps. After identifying the servers that will access the LUNs and selecting the iSCSI initiator adapters that will be used for LUN traffic, you should do the following:

1. Create and configure targets on the iSCSI subsystems and enable portals on those targets to handle LUN I/O traffic.
2. Optionally configure iSCSI security.
3. Establish logon sessions from the iSCSI initiator on each server to one or more targets.
4. Finally, assign LUNs to the targets.

**In this Section**

Task 1: Configure Server Connections

Task 2: (iSCSI Only) Configure Targets and Enable Portals

Task 3: (iSCSI Only) Configure Security

Task 4: (iSCSI Only) Establish Logon Sessions Between the Initiator and Targets

Task 5: Create LUNs

# Task 1: Configure Server Connections

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The first task in deploying LUNs is to configure the connections that servers in your SAN will use to access the LUNs that you create on your storage subsystems.

Storage Manager for SANs automatically discovers Fibre Channel ports and iSCSI initiator adapters on a server. You only need to identify the server that will access the LUN, and then select which ports and/or initiator adapters on that server will be enabled for LUN access.

| **Note** |
| --- |
| Fibre Channel ports can also be added manually by entering their World Wide Name (WWN). |

**To configure server connections**

1. In **LUN Management**, click **Manage Servers Connections**.
2. To add servers:
   a. In the **Manage Server Connections** dialog box, click **Add**.
   b. In the **Add Server** dialog box, type the name or browse to the server that you want to add and optionally, type a description for it.
   c. Click **OK**. The server will now be listed in the **Manage Server Connections** dialog box.
3. To configure the Fibre Channel ports on a server:
   a. In the **Manage Server Connections** dialog box, select a server from the server list. The **Fibre Channel Ports** tab lists all of the HBA ports that are available on the selected server.
   b. If you want to manually enter a new port, on the **Fibre Channel Ports** tab, click **Add** and enter the WWN of the new port. Optionally, type a description for the new port.
   c. To enable a port, select its check box in the **Enable** column.

      | **Caution** |
      | --- |
      | If you enable more than one Fibre Channel port and the server is not configured to use Multipath I/O, data corruption can occur. For more information about Multipath I/O, see "Support for Multipath I/O" earlier in this guide. |

   d. After selecting all ports that you want to enable, click **OK**.

      | **Note** |
      | --- |
      | If you create LUNs only on a Fibre Channel subsystem, no additional configuration is needed before creating LUNs. You can skip to "Task 5: Create LUNs" later in this guide. |

4. To configure the iSCSI initiator adapters on a server:
   a. In the **Manage Server Connections** dialog box, select a server from the server list. The **iSCSI Initiator Adapters** tab lists all the initiator adapters that are available on the selected server.

b. To enable an initiator adapter, select its check box in the **Enable** column.

| ⚠**Caution** |
| --- |
| If you enable more than one iSCSI initiator adapter and the server is not configured to use Multipath I/O, data corruption can occur. For more information about Multipath I/O, see "Support for Multipath I/O" earlier in this guide. |

c. After selecting all initiator adapters that you want to enable, click **OK**.

# Task 2: (iSCSI Only) Configure Targets and Enable Portals

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Targets are created in order to manage the connections between the iSCSI storage subsystems in your SAN and the servers that need to access them. You assign LUNs to a target, and all servers that log on to that target will have access to the LUNs assigned to it.

When you configure a target, you select the portals (network connections) that the target will provide for LUN access.

| **Note** |
| --- |
| In most cases, you can create and manage targets yourself. However, some iSCSI storage subsystems only support simple target configurations, where targets are automatically created when you create a LUN. With simple target configurations, you also cannot delete a target or manually assign LUNs to it. LUNs are automatically assigned when they are created. For this type of subsystem, you just need to identify the server or cluster that will access the LUN, and the iSCSI subsystem will enable access from that server or cluster to the LUN. |

**To configure a target on an iSCSI storage subsystem**

1. In **LUN Management**, click **Manage iSCSI Targets**.
2. In the subsystem drop-down list, select the iSCSI storage subsystem for which you want to configure a target.
3. To create a target on the selected subsystem, click **Add**.
4. In the **Add Target** dialog box, type a friendly name for the target. A separate iSCSI name will be automatically assigned.
5. In the list of available portals, select all portals that you want to enable for LUN access, and then click **OK**.

   The **Manage iSCSI Targets** dialog box shows the new target, with the portals that you enabled listed at the bottom.
6. To configure more targets, go back to step 2.
7. When you are finished configuring targets, click **OK**.

# Task 3: (iSCSI Only) Configure Security

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

There are several levels of iSCSI security available with Storage Manager for SANs. The basic level is based on the Challenge Handshake Authentication Protocol (CHAP). CHAP is a protocol that is used to authenticate the peer of a connection and is based upon the peers sharing a secret (a security key that is similar to a password). IP security (IPsec) is a protocol that enforces authentication and data encryption at the IP packet layer, which provides an added level of security.

You must choose the security level that best fits the security policies of your organization:

- **One-way CHAP authentication**. With this level of security, only the target authenticates the initiator. The secret is set just for the target and all initiators that want to access that target need to use the same secret to start a logon session with the target.

- **Mutual CHAP authentication**. With this level of security, the target and the initiator authenticate each other. A separate secret is set for each target and for each initiator in the SAN.

- **IPsec**. With this level of security, all IP packets sent during data transfers are encrypted and authenticated. A common key is set on all IP portals, allowing all peers to authenticate each other and negotiate packet encryption. For more information, see IPsec at the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=42985).

| **Caution** |
| --- |
| At a minimum, use one-way CHAP authentication between iSCSI initiators and targets. |

| **Note** |
| --- |
| The level of security that you can set for a storage subsystem depends on the hardware manufacturer. Not all subsystems support all levels of iSCSI security. You should contact your hardware manufacturer to verify what level of security is supported. |

| **Note** |
| --- |
| For more information about iSCSI security and to learn more about iSCSI in general, see Microsoft Storage Technologies - iSCSI at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50522). |

**To configure iSCSI security**

1. In **LUN Management**, click **Manage iSCSI Security**.

2. To configure one-way CHAP authentication, in the **Manage iSCSI Security** dialog box, configure the following settings on the **Targets** tab:

    a. If you want to configure different CHAP secrets for different targets, in the list of targets, select a target that you want to set the CHAP secret for, and click **Set Secret**.

    -Or-

    To use the same CHAP secret for a group of targets, select the targets from the list and click **Set Secret**.

    b. In the **Set Secret** dialog box, type and confirm the target CHAP secret.

    c. Optionally, select **Remember secret on local initiator** if you want to automatically pass the new secret to the local initiator.

   d. To set the new secret, click **OK**.

3. To configure mutual CHAP authentication, you must first configure one-way CHAP authentication by following step 2. Then, enter the following configuration on the **Local Initiator** tab:

   a. Type and confirm the CHAP secret for the local initiator.

   b. Under mutual CHAP authentication, the initiator will only be able to log on to targets that share the initiator secret. To share the initiator secret with the targets that the server needs to access, in the list of targets, select each target that you want to authenticate on the initiator.

   c. To set the new secret for the local initiator and to share it with the selected targets, click **Apply Secret**.

4. To configure IPsec, in the **Manage iSCSI Security** dialog box, configure the following settings on the **Portals** tab:

   a. If you want to use different IPsec keys for different portals, in the list of portals, select a portal and click **Set IPsec Key**.

      -Or-

      To use the same IPsec key for a group of portals, select the portals from the list, and click **Set IPsec Key**.

   b. In the **Set IPsec Key** dialog box, type and confirm a new IPsec key.

   c. Optionally, select **Remember the IPsec key on local initiator** if you want to automatically pass the new key to the local initiator.

   d. To set the new IPsec key, click **OK**.

5. When you are done configuring iSCSI security, click **OK**.

# Task 4: (iSCSI Only) Establish Logon Sessions Between the Initiator and Targets

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

For a server to access LUNs that are assigned to a target on an iSCSI subsystem, the server initiator must first establish a session with the target through an iSCSI logon process.

You can log on to iSCSI targets using Storage Manager for SANs. This starts a logon session that is only between the server that is running Storage Manager for SANs and the selected target. Each server that requires access to the LUNs that are assigned to a target must start a separate logon session with that target.

| **Note** |
| --- |
| You can also start a logon session between a server in your SAN and a target, using the local iSCSI initiator user interface, if available (this depends on the iSCSI initiator software that is installed on that server). |

When you log on to a target, you can select the type of CHAP authentication to use during the session. If you have set an IPsec key for the portals, data encryption is automatic.

**To establish logon sessions between the initiator and iSCSI targets**

1. In **LUN Management**, click **Log On to iSCSI Targets**.
2. In the **Log On to iSCSI Targets** dialog box, select the target that you want to establish a logon session with, and click **Log On**.
3. In the **Log On to Target** dialog box, select the type of CHAP authentication to use for the logon session:
   - **Log on using one-way CHAP authentication**
   - **Log on using mutual CHAP authentication**
4. To establish the logon session, click **OK**.

# Task 5: Create LUNs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

After you complete tasks 1 through 4, you can use the Create LUN Wizard to create a LUN and optionally assign the LUN to a server or target.

**To create and assign a LUN**

1. In **LUN Management**, click **Create LUN**.

2. Follow the steps in the Create LUN Wizard, providing the information in Table 1.

**Table 1 Completing the Create LUN Wizard**

| Wizard Page | Details |
|---|---|
| **Before You Begin** | Before you create a LUN, verify that the following steps have been completed:<br>• At least one subsystem is listed in the **Subsystems** node.<br>• Storage space is available on at least one of the available subsystems.<br>• If you want to use the wizard to assign the LUN to a server, the SAN server connections have been configured. |
| **Subsystem and Type** | Select the subsystem on which to create the LUN and specify the type of LUN to create.<br>• The list of subsystems at the top specifies the total capacity of each available subsystem, the type of subsystem (Fibre Channel or iSCSI), and the percentage of free space.<br>• The list of LUN types at the bottom includes each LUN type that is available on the selected subsystem and specifies the maximum-size LUN for each LUN.<br>• Fault tolerant LUN types require more disk space than simple LUNs. Therefore, the maximum size of the LUN is smaller than for a simple LUN. |
| **Size and Name** | Type a name for the new LUN and specify the size of the LUN:<br>• The LUN name can have as many as 63 characters.<br>• Depending on the version of the VDS provider that the subsystem is running, you may not be able to type a name for the LUN.<br>• Type the LUN size and click the drop-down list to select between megabytes (MB), gigabytes (GB) and terabytes (TB). |
| **Target Access (iSCSI only)** | If iSCSI targets have already been created on the subsystem, you can assign the LUN to one of those targets now. You can also leave it unassigned for later use. |

| | |
|---|---|
| **Server Access** | If you have already configured your server connections, you can assign the LUN to a server or cluster now. You can also leave it unassigned for later use.<br><br>**Note**<br>If you are creating a LUN on an iSCSI subsystem, then this LUN must also be assigned to an iSCSI target. The server or cluster you select on the **Server Access** page will only have access to the LUN after the iSCSI initiator on that server or cluster logs on to the target the LUN has been assigned to. |
| **Create Volume (Fibre Channel only)** | If on the **Server Access** page you chose to assign the LUN to the server you are currently managing with the Storage Manager for SANs snap-in, you can create an NTFS volume for the LUN. If you select to format the volume:<br><br>• As a general rule, leave the **Allocation unit size** option as **Default**. The *allocation unit size* is the smallest amount of disk space that can be allocated to hold a file. Windows selects default sizes based on the size of the volume. These default sizes are selected to reduce the amount of space that is lost and the amount of fragmentation on the volume.<br><br>• The **Quick Format** option speeds the volume creation process on large LUNs. If you want to scan for bad sectors during the formatting process, unselect this option and a full file format will be done.<br><br>**Note**<br>If you prefer, you can select not to format the volume and can use DiskPart to format it later. For more information, see DiskPart. |
| **Create New LUN** | Before creating the LUN, the wizard lists all the properties you have selected for the new LUN. Click **Create LUN** to start the LUN creation process. |

**Monitoring the status of LUN creation**

The process of creating a LUN can take minutes or hours, depending on the size of the LUN. Do not close Storage Manager for SANs while a LUN is being created. If the session ends before the process completes, you might not be able to determine the completion status for the LUN.

To monitor the status of the LUN while it is being created, in the **LUN Management** node, refer to the **Task Status** column on the results pane. While the LUN is being created, the column shows the percentage of the LUN that has been completed. When the operation has finished successfully, the task status changes to **Created**.

If LUN creation fails, you can click the LUN entry in the results pane to display failure information in the description area below it.

# Providing LUN Access to a Server Cluster

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

With Storage Manager for SANs, you can provide LUN access to a group of servers (a server cluster) in your SAN. Server clusters are created by defining them in Storage Manager for SANs and by enabling clustering services on each individual server in the cluster.

| **Caution** |
| --- |
| To prevent data loss, make sure that each server is a member of only one cluster and that clustering services are running on all cluster members. For more information, see Clustering Services at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50520). |

**To provide LUN access to a server cluster**

1. In **LUN Management**, click **Manage Server Connections**.
2. In the **Manage Server Connections** dialog box, click **Manage Clusters**.
3. In the **Manage Clusters** dialog box, click **Add**.
4. In the **Add New Cluster** dialog box, type a name for the cluster and optionally, type a general description that will help you identify the cluster.
5. From the server list, select all the servers that you want to include in the new cluster.

   | **Note** |
   | --- |
   | To add servers to a cluster you must first add them to the server list in **Manage Server Connections**. For more information, see "Task 1: Configure Server Connections" earlier in this guide. |

6. Click **OK** to create the new cluster, and click **OK** again to go back to the **Manage Server Connections** dialog box.
7. The new cluster will now be listed in the **Manage Server Connections** dialog box. To configure the connections of the cluster, configure the Fibre Channel ports or the iSCSI initiator adapters of each individual server included in the cluster by selecting the server from the list. All of the configuration changes you make on the individual servers will take effect on the cluster.
8. After server clusters are created, LUNs can be assigned to them in the same way they are assigned to individual servers.

| **Caution** |
| --- |
| If you are going to enable more than one connection (Fibre Channel ports and/or iSCSI initiator adapters) between a server cluster and a LUN, and the cluster is not configured to use Multipath I/O, data corruption can occur. For more information about Multipath I/O, see "Support for Multipath I/O" earlier in this guide. |

# Introduction to File Server Resource Manager

Updated: April 25, 2007

Applies To: Windows Server 2003 R2

With the increasing demand on storage resources, as organizations rely more heavily on data than ever before, IT administrators face the challenge of overseeing a larger and more complex storage infrastructure, while at the same time, tracking the kind of information available in it. Today, managing storage resources not only includes data size and availability but also the enforcement of company policies and a very good understanding of how existing storage is being used, allowing for sound strategic planning and proper response to organizational changes.

File Server Resource Manager is a suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers. By using File Server Resource Manager, administrators can place quotas on volumes, actively screen files and folders, and generate comprehensive storage reports. This set of advanced instruments not only helps the administrator to efficiently monitor existing storage resources, but it also aids in the planning and implementation of future policy changes.

By using File Server Resource Manager, you can perform the following tasks:

- Create quotas to limit the space allowed for a volume or folder and generate notifications when the quota limits are approached or exceeded.
- Automatically generate and apply quotas to all existing folders and any new subfolders in a volume or folder.
- Create file screens to control the type of files that users can save and send notifications when users attempt to save blocked files.
- Define quota and file screening templates that can be easily applied to new volumes or folders and reused across an organization.
- Schedule periodic storage reports that help identify trends in disk usage.
- Monitor attempts to save unauthorized files for all users or for a selected group of users.
- Generate storage reports instantly, on demand.

# Installing File Server Resource Manager

Updated: April 25, 2007

Applies To: Windows Server 2003 R2

After you complete your server installation, use **Add or Remove Programs** in Control Panel to install the File Server Resource Manager component.

To install File Server Resource Manager, you must be logged on with an account that has administrative rights on the local computer.

| **Note** |
| --- |
| You may need to restart your server after the installation of File Server Resource Manager has completed. |

**To install File Server Resource Manager**

1. Open Control Panel.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard** dialog box, select the **Management and Monitoring Tools** check box, and then click **Details**.
4. Click **Next**, and then click **Finish**.

# Using the File Server Resource Manager Component

Updated: April 25, 2007

Applies To: Windows Server 2003 R2

This section tells you how to open and use File Server Resource Manager. The following tasks are included:

- Open File Server Resource Manager
- Perform storage management tasks
- Manage storage on a remote computer

**Opening File Server Resource Manager**

File Server Resource Manager is a Microsoft Management Console (MMC) snap-in and is located in Administrative Tools on the **Start** menu.

**To open File Server Resource Manager**

- On the **Start** menu, click **Programs**, click **Administrative Tools**, and then click **File Server Resource Manager**.

File Server Resource Manager contains three main nodes:

- **Quota Management.** Use to create *quotas* that place size limits on volumes and folders.
- **File Screening Management.** Use to create *file screens* that block files from volumes and folders.
- **Storage Reports Management.** Use to schedule different types of storage reports and to create reports on demand.

**Performing Storage Management Tasks**

The following is a list of tasks that you can perform from each respective node in File Server Resource Manager:

**Quota Management**

- Create, manage, and obtain information about quotas, which are used to set a space limit on a volume or folder. By defining notification thresholds, you can send e-mail notifications, log an event, run a command or script, or generate reports when users approach or exceed a quota.
- Create and manage quota templates to simplify quota management.
- Create and manage auto quotas.

**File Screening Management**

- Create, manage, and obtain information about file screens, which are used to block selected file types from a volume or folder.
- Create file screening exceptions to override certain file screening rules.
- Create and manage file screen templates to simplify file screening management.
- Create and manage file groups. When used with file screens and file screening exceptions, the file groups determine which files will be blocked and which will be allowed. File groups also are used to select files to include and exclude from the Files by File Group Report and to sort file screens.

**Storage Reports Management**

- Schedule and configure storage reports.
- Generate storage reports on demand.

**Managing Storage Resources on a Remote Computer**

To manage storage resources on a remote computer, you can connect to the computer from File Server Resource Manager. While you are connected, File Server Resource Manager will display the objects created on the remote computer, allowing you to manage them.

The remote computer must be running the same operating system with File Server Resource Manager installed.

To connect to a remote computer, you must be logged on with a domain account that has administrative rights on the remote computer.

**To connect to a remote computer from File Server Resource Manager**

1. In the console tree, right-click **File Server Resource Manager**, and then click **Connect to another computer**.
2. In the dialog box that opens, click **Another computer**, and then select or type the name of the server you want to connect to. Click **OK**.

**Using Command-Line Tools for File Server Resource Manager**

If you prefer to work from the command line, you can use the following tools:

- **Dirquota.exe.** Use to create and manage quotas and quota templates.
- **FileScrn.exe.** Use to create and manage file screens, file screening exceptions, and file groups.
- **StorRept.exe.** Use to configure report parameters and generate storage reports on demand. You can also create report tasks and then use Schtasks.exe to schedule the tasks.

The tools are added to the system path when you install File Server Resource Manager and can be run from the command prompt. They can be used to perform storage management tasks on remote computers that are running the same operating system.

To learn the parameters for a specific tool, include the **/?** switch.

# Configuring E-mail Notifications

Updated: April 25, 2007

Applies To: Windows Server 2003 R2

When you create quotas and file screens, you have the option of sending e-mail notifications to users when their quota limit is approaching or after they have attempted to save files that have been blocked. If you want to routinely notify certain administrators of quota and file screening events, you can configure one or more default recipients.

To enable these notifications, you must specify the Simple Mail Transfer Protocol (SMTP) server to be used to forward the e-mail messages.

**To configure e-mail notifications**

1. In the console tree, right-click **File Server Resource Manager**, and then click **Configure options**. This opens the **File Server Resource Manager Options** dialog box.
2. On the **E-mail Notifications** tab, under **SMTP server name or IP address**, type the host name or the IP address of the SMTP server that will forward e-mail notifications.
3. If you want to routinely notify certain administrators of quota or file screening events, under **Default administrator recipients**, type each e-mail address. Use the format *account@domain*; use semicolons to separate multiple accounts.
4. To test your settings, click **Send Test E-mail**.

# Working with Quotas

Updated: April 25, 2007

Applies To: Windows Server 2003 R2

By creating a *quota* for a volume or folder, you can limit the disk space that is allocated for it. The quota limit applies to the entire folder subtree.

You can create two types of quotas:

- A *hard quota* prevents users from saving files after the space limit is reached and generates notifications when the volume of data reaches each configured threshold.
- A *soft quota* does not enforce the quota limit but generates all configured notifications.

To determine what happens when the quota limit approaches, you can configure *notification thresholds*. For each threshold that you define, you can send e-mail notifications, log an event, run a command or script, or generate storage reports. For example, you might want to notify the administrator as well as the user who saved the file, when a folder reaches 85 percent of its quota limit, and then send another notification when the quota limit is reached. In some cases, you might want to run a script that raises the quota limit automatically when a threshold is reached.

We recommend that you base your quotas on a *quota template*. By using a quota template, you can create a standard quota limit (for example, 200 megabytes (MB) for each user's folder on the server) and define a standard set of notification thresholds that can be applied to many quotas.

If you base your quotas on a quota template, and you later decide to change the quota configuration—for example, you want to allow each user additional space on the storage server—you can simply update the quota template and then choose to automatically update all quotas that are based on this template.

**File Server Resource Manager Quotas vs. NTFS Disk Quotas**

The following table outlines the advantages of using the quota management tools in File Server Resource Manager over NTFS disk quotas.

**Advantages of Using File Server Resource Manager**

| Quota features | File Server Resource Manager | NTFS disk quotas |
|---|---|---|
| Quota tracking | By folder or by volume | Per user on a volume |
| Disk usage calculation | Actual disk space | Logical file size |
| Notification mechanisms | E-mail, custom reports, command execution, event logs | Event logs only |

**Creating a Quota Template**

A *quota template* defines a space limit, the type of quota (hard or soft), and a set of notifications to be generated when the quota limit approaches. If you base your quotas on a template, you can automatically update all quotas that are based on a specific template by editing that template. This feature simplifies the process of updating the properties of quotas by providing one central point where all changes can be made.

For example, you can create a *User Quota* template that you use to place a 200 MB limit on the personal folder of each user. For each user, you would then create a quota based on the User Quota template and assign it to the user's folder. If you later decide to allow each user additional space on the server, you simply change the space limit in the User Quota template and choose to automatically update each quota that is based on that quota template.

File Server Resource Manager provides several quota templates. For example, you can use the **200 MB Limit Reports to User** template to place a hard 200 MB limit on the personal folder of each user. For some folders, you might want to use the **200 MB Limit with 50 MB Extension** template to automatically increase the quota limit when a 200 MB quota limit is reached. Other default templates are designed for monitoring disk usage through soft quotas—for example, the **Monitor 200 GB Volume Usage** template and the **Monitor 500 MB Share** template. To view the default templates, click the **Quota Templates** node in the File Server Resource Manager console tree.

**Working with templates**

File Server Resource Manager gives you flexibility in creating, using, and managing templates, both for quotas and for file screens.

If you create a quota with custom properties, you have the option of saving a quota template with the custom quota settings when you save the new quota. If you choose that option, a new template is saved, and then the template is applied to the new quota so that a link is maintained between the quota and the template.

After creating a template, you can select the template in the results pane and choose the **Create quota from template** action to quickly create a quota based on the template. You only need to specify the volume or folder to which the quota will apply.

**To create a quota template**

1. In **Quota Management**, right-click **Quota Templates**, and click **Create quota template**. This opens the **Create Quota Template** dialog box.

| **Note** |
| --- |
| If you want, you can copy the properties of an existing template by selecting it from the drop-down list box and clicking **Copy**. These properties can be used as a base for the new template, or you can create an entirely new one. |

2. To identify the new template, type a template name and an optional quota label.
3. To configure the **Space Limit** for the quota, under **Limit**, specify the amount of space to allow, and then specify the type of quota to impose:
   - A *hard quota* prevents users from saving files after the space limit is reached and generates notifications when the volume of data reaches each configured threshold.
   - A *soft quota* does not enforce the quota limit but generates all configured notifications.
   - *Notification thresholds* determine what happens when the quota limit approaches. You can send e-mail notifications, log an event, execute a command or a script, or generate reports. By default, no notifications are generated.
4. To configure notifications that will be generated when the quota limit is reached, in the **Notification thresholds** list, click the **Limit (100%)** threshold, and then click **Edit**. This opens the **Limit Properties** dialog box.
5. To configure e-mail notifications, set the following options on the **E-mail Message** tab:
   - Select the **Send e-mail to the following administrators** check box, and enter the names of the administrative accounts that will receive notifications. Use the format *account@domain*; use semicolons to separate multiple accounts.
   - To contact the person who saved the file that exceeded the quota threshold, select the **Send e-mail to the user who exceeded the threshold** check box.
   - Optionally, edit the default **Subject** line and **Message body** that are provided.
   - The text that is in brackets inserts variable information about the quota event that caused the notification. For example, the **[Source Io Owner]** variable inserts the user name of the user or application that saved the file to disk. To insert additional variables in the text, click **Insert Variable**.
   - To configure additional e-mail headers, click **Additional E-mail Headers**. You can add several header fields to the notification e-mail, including “From,” “Cc,” “Blind Cc,” and “Reply-To.”
6. To log an error to the event log, on the **Event Log** tab, select the **Send warning to event log** check box, and optionally edit the default log entry.
7. To run a command or script, on the **Command** tab, select the **Run this command or script** check box, and type the command or browse to the location where the script is stored. Optionally, enter command arguments, select a working directory for the command or script, or modify the command security setting.
8. To specify that one or more storage reports be generated automatically, on the **Report** tab, select the **Generate reports**

check box, and select which reports to generate. You must choose at least one e-mail recipient for the report, either an administrator or the user.

9. After configuring each type of notification that you want to generate, click **OK** to save the threshold.
10. To configure additional notification thresholds, perform these steps:
    a. In the **Create Quota Template** dialog box, under **Notification Thresholds**, click **Add**.
    b. At the top of the **Add Threshold** dialog box, enter the percentage of the quota limit at which to generate the notifications. The default threshold for the first notification is 85 percent.
    c. Use the **E-Mail Message**, **Event Log**, **Command**, and **Reports** tabs to configure each type of notification that you want to generate, as you did in steps 6 through 10. Click **OK** to save the threshold.
11. When you are finished configuring the quota template properties, click **OK** to save the template.

#### Creating a Quota

The following procedure guides you through the process of creating a quota that is based on a template, which is recommended practice. By creating quotas exclusively from templates, you can centrally manage your quotas by updating the templates instead of the individual quotas.

**To create a quota that is based on a template**

1. In the console tree, under **Quota Management**, click **Quota Templates**.
2. In the results pane, right-click the template on which you will base your quota, and then click **Create quota from template**. This opens the **Create Quota** dialog box.
3. Under **Quota path**, select or type the path to the folder to which the quota will apply. Note that the quota properties will also apply to all of its subfolders.
4. To base the quota on a quota template, select the **Derive properties from this quota template** option, and select the template from the drop-down list box. Note that each template properties will be displayed under **Summary of quota properties**.
5. Click **Create**.

| **Note** |
| --- |
| If you do not want to base your quota on a template, or if you want to edit the properties copied from a template, in the **Create Quota** dialog box you can choose **Define custom quota properties**, and then click **Custom Properties**. These properties are like those for quota templates. For instructions, see "To create a quota template" earlier in this guide. Note that if you create quotas with custom properties, you will not have the ability to update your quotas by editing a quota template. |

#### Updating Quotas by Editing a Quota Template

When making changes to a quota template, you have the option of extending those changes to quotas that were created using the original quota template. You can choose to modify only those quotas that still match the original template or all quotas that were derived from the original template, regardless of any modifications made to the quotas since they were created. When you reapply a template, all properties in the template overwrite all properties in the quotas; the updates are not limited to the properties that you changed in the quota.

The same capabilities for creating, applying, and updating templates are available for file screens.

| **Note** |
| --- |
| If you choose to apply the changes to all quotas that were derived from the original template, you will overwrite any custom quota properties that you may have created. |

**To edit a quota template**

1. In the console tree, under **Quota Management**, click **Quota Templates**.

2. In the results pane, right-click the quota template you want to modify, and then click **Edit template properties**, or select **Edit template properties** from the Actions pane, under **Selected Quota Templates**. This opens the **Quota Template Properties** dialog box.

3. Perform all necessary changes. For more information, see "To create a quota template" earlier in this guide. Optionally, you can copy the properties of a different template and modify them for this one.

4. When you are finished editing the template properties, click **OK**. This will open the **Update Quotas Derived from Template** dialog box.

5. Select the type of update that you want to apply:

   - If you have quotas that have been modified since they were created using the original template, and you do not want to change them, select **Apply template only to derived quotas that match the original template**. This option will update only those quotas that have not been previously edited and that still have the original template properties.

   - If you want to modify all existing quotas that were created using the original template, select **Apply template to all derived quotas**.

   - If you want to keep the existing quotas unchanged, select **Do not apply template to derived quotas**.

6. Click **OK**.

**Creating an Auto Quota**

Auto quotas allow you to assign a quota template to a folder and specify that quotas based on that template be automatically generated and applied to the existing subfolders, as well as to any new subfolders created in the future.

For example, you may define an auto quota for folders that are created on demand, for roaming profile users or for new users. Every time a new folder is created, a new quota entry is automatically generated for that folder, using the template defined in the auto quota profile of its parent folder. These automatically generated quota entries can then be viewed in the Results pane, under the **Quotas** node, as individually created quotas, which can be edited and modified separately.

**To create an auto quota**

1. In the console tree, under **Quota Management**, right-click **Quotas**, and then click **Create quota**, or select **Create quota** from the Actions pane. This opens the **Create Quota** dialog box.

2. Under **quota path**, select or type the path to the folder to which the auto quota will apply. The auto quota will be applied to all subfolders (current and future) in this folder. Once the path is entered, select **Auto apply template and create quotas on existing and new subfolders**.

3. Under **Derive properties from this quota template**, select the quota template that you want to apply. You can review the properties of each quota template in the summary list box.

4. Click **Create**.

| **Note** |
| --- |
| You can verify all automatically generated quotas by opening the **Quotas** node; an individual quota for each subfolder will be listed. |

**Editing an Auto Quota**

When making changes to an auto quota, you have the option of extending those changes to existing quotas in the auto quota path. You can choose to modify only those quotas that still match the original auto quota, or all quotas in the auto quota path, regardless of any modifications made to the quotas since they were created. Since auto quotas are based on quota templates, any specific property changes have to be created either in the quota template used for the original auto quota or by creating a new quota template with the desired properties. For more information about how to modify an existing quota template, see "To edit a quota template" earlier in this guide.

| **Note** |
| --- |

| |
|---|
| If you choose to apply the changes to all quotas in the auto quota path, you will overwrite any custom quota properties that you may have created. |

**To edit an auto quota**

1. In the console tree, under **Quota Management**, click **Quotas**.
2. In the results pane, right-click the quota entry, and then click **Edit quota properties**, or select **Edit quota properties** from the Actions pane, under **Selected Quotas**. This opens the **Edit Auto Apply Quota** dialog box.
3. Under **Derive properties from this quota template**, select the new quota template that you want to apply. You can review the properties of each quota template in the summary list box.
4. Click **OK**. This will open the **Update Quotas Derived from Auto Apply Quota** dialog box.
5. Select the type of update you want to apply:
   - If you have quotas that have been modified since they were automatically generated, and you do not want to change them, select **Apply auto apply quota only to derived quotas that match the original auto apply quota**. This option will update only those quotas in the auto apply quota path that have not been edited and maintain the original auto apply quota properties.
   - If you want to modify all existing quotas in the Auto Apply Quota Path, select **Apply auto apply quota to all derived quotas**.
   - If you want to keep the existing quotas unchanged but make the modified auto quota effective for new folders in the Auto Apply Quota Path, select **Do not apply auto apply quota to derived quotas**.
6. Click **OK**.

**Monitoring Quota Use**

In addition to the information included in notifications, you can find out about quota usage by viewing quotas in **Quota Management**, by generating a **Quota Usage** report, or by creating soft quotas for the purpose of monitoring overall disk usage.

**Viewing quota information**

To view quota information in File Server Resource Manager, in the console tree, click **Quota Management**, and then click **Quotas**.

- In the Results pane, you can quickly find out the quota limit, the percentage of the limit that is used, and whether the quota is hard or soft.
  - Use the **View** options at the top of the Results pane to limit the display to the quotas that affect a designated path.
  - Click a column header to group quotas by that attribute.
- To find out additional information about a quota, select (highlight) the quota. The description area lists the notification thresholds and notes the remaining space available on the volume on which the quota was created.
- To view all quotas whose limits apply to a specific folder, right-click the quota of interest (identified by its path) in the Results pane, and click **View quotas affecting folder**. The description area lists the applicable quotas.

  -Or-

  To filter the Results pane display to the quotas that affect the path, in the **View** area at the top of the pane, click **Quotas affecting path**, then enter or browse to the path, and then click **Filter**.

**Quota Usage report**

Use the Quota Usage report to identify quotas that may soon be exceeded so that you can take the appropriate action. For instructions on generating a Quota Usage report, see “Generating Storage Reports” later in this guide.

**Templates for monitoring disk usage**

To monitor overall disk usage, you may also want to create soft quotas for volumes or shares. File Server Resource Manager provides the following default templates that you can use (or adapt) for this purpose: the **Monitor 200 GB Volume Usage** template and the **Monitor 500 MB Share** template.

# Working with File Groups

Updated: April 25, 2007

Applies To: Windows Server 2003 R2

Before you begin working with file screens, you must understand the role of file groups in determining which files are screened. A *file group* is used to define a namespace for a file screen, file screen exception, or storage report.

A file group consists of a set of file name patterns, which are grouped into files to include and files to exclude:

- **Files to include**: files that belong in the group.
- **Files to exclude**: files that do not belong in the group.

For example, an *Audio Files* file group might include the following file name patterns:

- **Files to include: *.mp***: Includes all audio files created in current and future MPEG formats (MP2, MP3, and so forth).
- **Files to exclude: *.mpp**: Excludes files created in Microsoft® Project (.mpp files), which would otherwise be included by the *.mp* inclusion rule.

File Server Resource Manager provides several default file groups. You can define additional file groups, or change the files to include and exclude. Any changes that you make to a file group affect all existing file screens, templates, and reports to which the file group has been added.

| **Note** |
| --- |
| For convenience, you can modify file groups when you edit the properties of file screens, file screen exceptions, and file screen templates. Note that any changes that you make to a file group from these property sheets will affect all items that use that file group. |

**To create a file group**

1. In the console tree, under **File Screen Management**, right-click **File Groups**, and then click **Create file group**.

   -Or-

   While you edit the properties of a file screen, file screen exception, or file screen template, under **Manage file groups**, click **Create**.

   This opens the **Create File Group Properties** dialog box.
2. Type a name for the file group.
3. Add files to include and files to exclude:
   - For each set of files that you want to include in the file group, in **Files to include**, type a file name pattern, and then click **Add**.

     Standard wildcard rules apply. For example, ***.exe** selects all executable files.
   - For each set of files that you want to exclude from the file group, in **Files to exclude**, type a file name pattern, and then click **Add**.
4. Click **OK**.

# Screening Files

Updated: April 25, 2007

Applies To: Windows Server 2003 R2

Create *file screens* to block certain files from being saved on a volume or in a folder tree. A file screen affects all folders in the designated path. For example, you might create a file screen to prevent users from storing audio and video files in their personal folders on the server.

To specify which files to screen, you assign one or more file groups to the file screen. For information about file groups, including a list of the default file groups that are provided with File Server Resource Manager, see "Working with File Groups" earlier in this guide.

To simplify the management of file screens, we recommend that you base your file screens on *file screen templates*. File Server Resource Manager provides several default file screen templates, which you can use to block audio and video files, executable files, image files, e-mail files, and meet some other common administrative needs. Templates set up active or passive screening, and provide default notification text. To view the default templates, select the **File Screen Templates** node in the File Server Resource Manager console tree.

**Creating a File Screen**

In the following procedure, you will create a new file screen, and in the process save a file screen template based on the custom file screen properties that you defined. The new template will be applied to the file screen so that a link is maintained between the file screen and the template. You can use the same method for creating a quota template.

**To create a file screen**

1. In the console tree, under **File Screening Management**, right-click **File Screens**, and then click **Create file screen**. This opens the **Create File Screen** dialog box.

2. Under **File screen path**, type or browse to the folder to which the file screen will apply. The file screen will also apply to all subfolders of the selected folder.

3. Under **How do you want to configure file screen properties**, click the **Define custom file screen properties** , and then click **Custom Properties**. This opens the **File Screen Properties** dialog box.

4. If you want to copy properties from an existing template, select the template that you want to use, and click **Copy**. You can edit the properties for the new file screen.

5. Under **Screening type**, select the type of screening to apply:

   - **Active screening** prevents users from saving files that are members of blocked file groups, and generates notifications when users try to save blocked files. (The file screen does not prevent users and applications from accessing files that were saved to the path before the file screen was created, regardless of whether the files are members of blocked file groups.)

   - **Passive screening** sends notifications, but does not prevent users from saving blocked files.

   To specify which files to screen, under **File Groups**, select each file group that you want to include.

   If you want to view the members and non-members of a file group, click the file group, and then click **Edit**. For instructions on creating a file group, see "Working with File Groups" earlier in this guide.

6. To configure e-mail notifications for the file screen, set the following options on the **E-mail Message** tab:

   - To notify administrators when a user or application attempts to save an unauthorized file, select **Send e-mail to the following administrators**, and enter the administrative accounts that will receive notifications. Use the format *account@domain*; use semicolons to separate multiple accounts.

   - To send a notice to the user who attempted to save an unauthorized file, select **Send e-mail to the user who attempted to save an unauthorized file**.

   - Optionally, edit the default subject line and message body. The text that is in brackets inserts variable information

about the file screen event that caused the notification. For example, the **[Source Io Owner]** variable inserts the user name of the user or application that attempted to write the file to disk. Use the **Insert Variable** button to insert additional variables in the text.

7. If you want to log an error to the event log when a user tries to save a blocked file, on the **Event Log** tab, select the **Send warning to event log** check box, and, optionally, edit the default log entry.

   In addition, you can set options on the **Command** tab to run a command or script when a user attempts a file screen violation, and set options on the **Report** tab to specify one or more storage reports to be generated automatically.

8. When all file screen properties are complete, click **Create** to save the file screen.

   You will be asked if you want to save a file screen template based on the custom file screen properties you just defined. If you plan to use the current settings in other file screens, we recommend that you save them as a template. The template will be applied to the new file screen. This will enable you to update the file screen automatically whenever you update the template.

9. To save a template when you create the file screen, select **Save the custom properties as a template**, enter a name for the template, and click **OK**.

#### Creating a File Screen Exception

Occasionally, you will need to allow exceptions to file screening. For example, you might want to block video files from a file server but you need to allow your training group to save the video files for their computer-based training. To allow files that other file screens are blocking, create a *file screen exception*.

A file screen exception is a special type of file screen that overrides any file screening that would otherwise apply to a folder, and all its subfolders, in a designated exception path. That is, it creates an exception to any rules derived from a parent folder. To determine which files to include in the exception, file groups are assigned.

| **Note** |
| --- |
| You cannot create a file screen exception on a folder for which you already have a file screen. You should assign the exception to a subfolder or make changes to the existing file screen. |

**To create a file screen exception**

1. In the console tree, under **File Screening Management**, right-click **File Screens**, and then click **Create file screen exception**. This opens the **Create File Screen Exception** dialog box.

2. Under **Exception path**, type or browse to the folder to which the exception will apply. The exception will also apply to all subfolders.

3. To determine which files to exclude from file screening, in **File groups**, select each file group that you want to include in the file screen exception.

   If you want to view the members and non-members of a file group, select the file group, and then click **Edit**. For information about creating file groups, see “Working with File Groups” earlier in this guide.

4. Click **OK**.

#### Monitoring File Screening

In addition to the information in your file screen notifications, you can monitor file screening by viewing file screens in the **File Screening** task area and by generating a **File Screening Audit** report.

#### Viewing file screening information

To view file screening information in the File Server Resource Manager console tree, click **File Screening Management**, and then click the node that you want to view (**File Screens**, **File Screen Exceptions**, or **File Groups**).

- For each file screen, the results pane displays the following information: the path that the file screen was created for, the type of file screen (file screen or exception), the file groups included in the file screen, and the template on which the file screen is based.

- For the selected file screen, the description area lists all file groups that are being blocked on the file screen path. This

includes file groups that are blocked by the current file screen as well as file groups blocked by file screens created higher in the file screen path.

**File Screening Audit report**

Use the **File Screening Audit** report to identify individuals or applications that violate file screening policy. For instructions on generating a **File Screening Audit** report, see "Generating Storage Reports" later in this guide.

# Generating Storage Reports

Updated: April 25, 2007

Applies To: Windows Server 2003 R2

From the **Storage Reports Management** node, you can generate reports that will help you understand file use on the storage server. You can use the storage reports to monitor disk usage patterns (by file type or user), identify duplicate files and dormant files, track quota usage, and audit file screening. The following table describes each storage report that is available.

**Storage Reports Available in File Server Resource Manager**

| Report | Description |
|---|---|
| Large Files | Lists files that are larger than a specified size. Use this report to identify files that are consuming the most disk space on the server. |
| Files by Owner | Lists files, grouped by owner. Use this report to analyze usage patterns on the server and to identify users who use large amounts of disk space. |
| Files by File Group | Lists files that belong to specified file groups. Use this report to identify file group usage patterns and to identify file groups that occupy large amounts of disk space. This can help you determine which file screens to configure on the server. |
| Duplicate Files | Lists files that appear to be duplicates (files with the same size and last-modified date). Use this report to identify and reclaim disk space that is wasted because of duplicate files. |
| Least Recently Used Files | Lists files that have not been accessed for a specified number of days. This can help you identify seldom used data that might be archived and removed from the server. |
| Most Recently Used Files | Lists files that have been accessed within a specified number of days. Use this report to identify frequently used data that needs to be highly available. |
| Quota Usage | Lists quotas for which the quota usage is higher than a specified percentage. Use this report to identify quotas with high usage levels so that appropriate action can be taken. (The report includes quotas that were created for volumes and folders in File Server Resource Manager only. It does not include quotas applied to shares in NTFS.) |
| File Screening Audit | Lists file screening violations that have occurred on the server, for a specified number of days. Use this report to identify individuals or applications that violate file screening policy. |

You can create *report tasks*, which are used to schedule one or more periodic reports, or you can generate reports on demand and optionally display the reports immediately. For on-demand reports, as with scheduled reports, current data is gathered before the report is generated.

**Configuring Report Parameters**

Most of the reports have configurable *report parameters*, which determine the content that the report includes. The parameters vary with the type of report. For some reports, report parameters can be used to select the volumes and folders on which to report, set a minimum file size to include, or restrict a report to files owned by specific users.

To configure the default parameters for each report, use the **Configure options** action, which is available when you select the main File Server Resource Manager node. This action opens the **File Server Resource Manager Options** dialog box. The default parameters are used in the incidental reports that are generated automatically during quota and file screen notifications. You can edit the default parameters for scheduled reports and reports generated on demand.

**Important**

Before you run a **File Screen Audit** report, in the **File Server Resource Manager Options** dialog box, on the **File Screen Audit** tab, verify that the **Record file screening activity in the auditing database check** box is selected.

**Saving reports**

Regardless of how you generate a report, or whether you choose to view the report immediately, the report is saved on disk, using the file formats that you specify. By default, reports are saved in Dynamic HTML (DHTML) format. You can also save reports in HTML, XML, DSV, and text formats.

Scheduled reports, on-demand reports, and incident reports are saved in separate folders within a designated report repository. By default, the reports are stored in subdirectories of the **%Systemroot%\StorageReports** folder. To change the default report locations, use the **Configure options** action.

**Scheduling a Set of Reports**

To generate a set of reports on a regular schedule, you schedule a *report task*. The report task specifies which reports to generate and what parameters to use, which volumes and folders to report on, how often to generate the reports, and which file formats to save them in.

**Note**

To minimize the impact of report processing on server performance, generate multiple reports on the same scheduled report task, so that the data is only gathered once. To quickly add reports to existing report tasks, under **Storage Reports Management**, select the **Add or remove reports for a report task** action. This action allows you to quickly add or remove reports from all existing report tasks, and to edit the report parameters. To change schedules or delivery addresses, you must edit individual report tasks.

**The schedule a report task**

1. In the console tree, right-click **Storage Reports Management**, and then click **Schedule a new report task**. This opens the **Storage Reports Task Properties** dialog box.

2. To select a volume or folder on which to report:

    a. In **Scope**, click **Add**.

    b. Browse to the volume or folder that you want to report on, and click **OK**.

    This adds the selected volume or folder to the list. Add as many paths as you want to include.

3. In **Report data**, select each report that you want to generate. (By default, all reports are generated for a scheduled report task.)

4. To edit the parameters of a report:

    a. Highlight the report label, and click **Edit Parameters**.

    b. In the **Report Parameters** dialog box, edit the parameters as needed, and click **OK**.

If you want to deliver copies of the reports to administrators by e-mail, on the **Delivery** tab, click **Send reports to the following administrators**, and enter each e-mail account. Use the format *account@domain*; use semicolons to separate multiple accounts.

1. To schedule the reports, on the **Schedule** tab, click **Create Schedule**. This opens the **Schedule** dialog box.

2. In the **Schedule** dialog box, click **New**. This creates a default schedule set for 9:00 A.M. daily, starting on the next day.

3. To change the default schedule:

    a. Under **Schedule Task**, select the interval at which to generate the report.

    You can schedule daily, weekly, or monthly reports, or generate the reports only once. You can also arrange to generate the reports at system startup or at logon, or specify background printing when the system is idle.

    b. In the schedule options below, enter additional information, as needed, to determine when the reports will be generated.

    The available schedule options vary based on the interval that you choose. For example, for a weekly report, you can specify how many weeks between reports and which days of the week to generate reports on; for a monthly report, you can specify what day of the month and what specific months to generate the reports on.

    c. Under **Start time**, enter the time of day when you want to generate the report.

4. To add more schedules to the report task, repeat steps 7 through 8. Each new schedule will be individually numbered, click the drop-down list box under the **Schedule** tab to see the list of schedules. To delete a schedule, select it in the drop-down list box and click **Delete**.

5. After you are finished adding schedules to the report task, click **OK**.

6. To save the report task, click **OK**.

**Generating Reports on Demand**

During daily operations, you may want to generate reports on demand to analyze different aspects of current disk usage on the server. Use the **Generate reports now** action to generate one or more reports. Before the reports are generated, current

data is gathered.

When you generate reports on demand, the reports are saved in the reports repository, but no report task is created for later use. You can optionally view the reports when they have completed.

| **Note** |
| --- |
| If you choose to open the reports immediately, you must wait while the reports are generated. Processing time varies, depending on the types of reports and the scope of the data. |

**To generate reports immediately**

1. In the console tree, right-click **Storage Reports Management**, and then click **Generate reports now**. This opens the **Storage Reports Task Properties** dialog box.

2. In **Scope**, add each volume or folder that you want to report on.

3. In **Report data**, select and format each report that you want to include:

   a. To add a report to the task, select the check box for the report.

   b. To edit the parameters of a report, highlight the report label and click **Edit Parameters**. When you finish editing the parameters, click **OK**.

4. In **Report formats**, select each file format that you want to save the report in.

5. To deliver the reports by e-mail, on the **Delivery** tab, select the **E-mail report to the following administrators** check box. Then, enter each e-mail address that you want the report delivered to. Use the format *account@domain*; use semicolons to separate multiple accounts.

6. To gather the data and generate the reports, click **OK**.

7. In the dialog box that opens, choose whether to open the reports when they are complete or generate the reports in the background, in which case the reports are not displayed.

8. Both options save the reports in the formats that you selected. If you choose to open the reports after they are generated, each report opens in its own Internet Explorer browser window.

# Test Scenarios for File Server Resource Manager

Updated: April 25, 2007

Applies To: Windows Server 2003 R2

Before deploying File Server Resource Manager on a production server, we recommend that you familiarize yourself with the different quota, file screening and reporting options, by performing the following scenarios on a test server.

| **Important** |
| --- |
| To successfully perform some of the test scenarios presented in this guide, you must first configure the general e-mail notifications parameters. For more information, see "Configuring E-mail Notifications" earlier in this guide. |

**Quota Management Scenarios**

By creating a quota for a volume or folder, you can limit the disk space that is allocated for that volume or folder. The **Quota Management** node of File Server Resource Manager includes all the necessary options to work with quotas.

There are three groups of tasks that can be tested under **Quota Management**:

- Quota tasks
- Auto quota tasks
- Quota template tasks

| **Note** |
| --- |
| For more information about **Quota Management**, the different tasks available in this node, and step-by-step procedures, see "Working with Quotas" earlier in this guide. |

**Testing quotas**

We recommend using the following two scenarios to familiarize yourself with some of the quota tasks in File Server Resource Manager.

**Scenario 1: Quotas for soft monitoring**

In some situations, rather than creating a hard limit on the amount of allocated space, you may want to perform soft monitoring of a volume or a folder, so that you receive notifications when predetermined usage thresholds have been exceeded. This is a very useful technique that alerts administrators before a shortage of storage space occurs.

**Test setup:**

Create a soft quota on one of the volumes in your server (for example, D:\ or E:\) using the following parameters:

- **Quota limit:** total volume capacity
- **Quota type:** soft quota
- **Notifications:**
    - **Limit (100%):** send e-mail to administrator, send warning to the event log
    - **Warning (85%):** send e-mail to administrator, send warning to the event log
    - **Warning (80%):** send e-mail to administrator
    - **Warning (50%):** send e-mail to administrator

**Verification:**

- Verify that it is still possible to save files to the volume.

- Copy or generate enough data in the volume to exceed any of the notification thresholds and then verify that the notifications that you selected have been created and received.

| **Note** |
| --- |
| To generate large files you can use the **fsutil** command with the following parameters: **fsutil file createnew** *file_name file_size*, where *file_name* is the path and name of the file that you want to create and *file_size* is the size of the file in bytes. An example: fsutil file createnew D:\bigfile.txt 1024000. This will create a file called bigfile.txt on the root directory of volume D, with an approximate size of 1 megabyte (MB). |

**Scenario 2: Quotas on group shares**
Several templates for working with quotas are included with File Server Resource Manager. One of these templates is the **Monitor 500 MB Share**, which is intended to be used when administering a shared folder on your server.

**Test setup:**
Using the **Monitor 500 MB Share** quota template, create a quota on a shared folder on your server (for example: E:\Scratch). For more information about how to create a quota from a template, see "Creating a Quota" earlier in this guide.

**Verification:**

- Copy or generate enough data in the shared folder to exceed any of the notification thresholds and then verify that the selected notifications have been created and received.

**Testing auto quotas**
We recommend using the following two scenarios to familiarize yourself with some of the auto quota tasks in File Server Resource Manager.

**Scenario 1: Auto quotas on user folders**
It is common that each user in an organization has a folder hosted on a file server. Usually, all of these user folders are located under a common root and a new folder is created every time a new user is granted access to the server. Using auto quotas can simplify the creation of quotas for user folders, as well as the administration of these quotas.

**Test setup:**
Create a folder called Users on the root directory of a volume. Under this folder, create subfolders, one for each user. For example:

D:\Users\User001

D:\Users\User002

D:\Users\User003

...

...

...

D:\Users\User100

When you have finished creating all of the Users subfolders, create an auto quota on the Users folder, using the **100 MB Limit** template. For more information about how to create an auto quota, see "Creating an Auto Quota" earlier in this guide.

**Verification:**

- Under **Quotas**, in the results pane, verify that an individual quota was automatically generated for each subfolder in the Users folder.

- Verify that no quota was generated for the Users (parent) folder.

- Create a new subfolder in the Users folder (for example: D:\Users\User101) and verify that a new individual quota is automatically generated for the new subfolder.

- Edit the auto quota that you just created and select to derive properties from a different quota template (for example, the **200 MB Limit Reports to User** template). Choose to apply the auto quota changes to all derived quotas and then, under **Quotas**, verify that all of the individual quotas for the Users subfolders have now been updated with the properties of the new quota template.

**Scenario 2: Auto quotas on shared folders under a common root**
Auto quotas can also be applied to a volume root. This scenario is especially useful when, for example, a volume is used to directly host a number of folders shared by the organization.

**Test setup:**
Create several shared folders on the root directory of a volume in your server, for example:

E:\Accounting (shared as: *Server_name*\Accounting)

E:\Finance (shared as: *Server_name*\Finance)

E:\Marketing (shared as: *Server_name*\Marketing)

E:\Payroll (shared as: *Server_name*\Payroll)

E:\Sales (shared as: *Server_name*\Sales)

When you have finished creating the folders and their shares, create an auto quota on the volume root (in this example, on E:\), using the **100 MB Limit** template. For more information about how to create an auto quota, see "Creating an Auto Quota" earlier in this guide.

**Verification:**

- Under **Quotas**, in the results pane, verify that an individual quota was automatically generated for each of the folders in the volume and that a shared path is listed for it.

- From a different server in the network, access any of the shared folders that you created, and copy or generate enough data in that shared folder to exceed the 100 MB limit. Verify that the quota limit is enforced on the shared folder.

**Testing quota templates**
We recommend using the following two scenarios to familiarize yourself with some of the quota template tasks in File Server Resource Manager.

**Scenario 1: Creating a new template from an existing template**
File Server Resource Manager has several default templates that can be easily modified to create templates that are more appropriate for your server environment.

**Test setup:**
Under **Quota Templates**, click the **Create quota template** action and copy the properties from the **100 MB Limit** template. Then modify the following parameters:

- **Template name:** 50 MB Limit

- **Quota limit:** 50 MB

- **Quota type:** soft quota

- **Notifications:**

    - **Limit (100%):** send e-mail to administrator, send warning to the event log

    - **Warning (85%):** send e-mail to administrator, send warning to the event log

    - **Warning (50%):** send e-mail to administrator

When you are finished performing these changes, click **OK** to save the new template.

**Verification:**

- Under **Quota Templates**, in the results pane, verify that the new **50 MB Limit** template is listed.

**Scenario 2: Using a quota template to update quotas**
If you base your quotas and auto quotas on templates, you can automatically update all quotas that are based on a specific template simply by editing that template. This feature simplifies the process of updating quota properties by providing one central point where all changes can be made.

**Test setup:**
If you have not already done so, perform the first scenario for testing auto quotas ("Scenario 1: Auto quotas on user folders") and the first scenario for testing quota templates ("Scenario 1: Create a new template from an existing template").

When you have completed these two scenarios, under **Auto Quotas**, edit the auto quota for the Users folder and select to derive properties from the **50 MB Limit** template. When prompted, choose to apply the auto quota changes to all derived quotas.

**Verification:**

- Under **Quotas**, in the results pane, verify that all the quotas created for the user folders have now been updated with the properties of the **50 MB Limit** template.

- Next, under **Quota Templates**, select to edit the **50 MB Limit** template and modify the following properties:

  - **Quota limit:** 25 MB

  - **Quota type:** hard quota

  When prompted, select to apply the template changes to all derived quotas. Finally, under **Quotas**, verify that all the individual quotas for the Users subfolders have now been updated with the new 25 MB quota limit and that the type of quota has been set to "hard."

**File Screening Management Scenarios**
File screens are used to block certain types of files from being saved on a volume or in a folder tree. To specify which files to screen, you assign one or more file groups to a file screen. The **File Screening Management** node of File Server Resource Manager includes all the necessary options to work with file screens.

There are three groups of tasks that can be tested under **File Screening Management**:

- File group tasks

- File screen tasks

- File screen template tasks

| **Note** |
| --- |
| For more information about **File Screening Management**, the different tasks that are available in this node, and step-by-step procedures, see "Screening Files" earlier in this guide. For more information about file groups, see "Working with File Groups" earlier in this guide. |

**Testing file groups**
We recommend using the following scenario to familiarize yourself with some of the file group tasks in File Server Resource Manager.

**Scenario 1: Using a file group to block specific file names**
File Server Resource Manager includes several default file groups for screening, which include various different file types, based on their file name pattern. But you may want to block specific files that do not follow a specific pattern.

**Test setup:**
Under **File Groups**, click the **Create file group** action to create a new file group with the following properties:

- **File group name:** Specific Files
- **Files to include:** file_name1.exe, file_name2.dll, file_name3.txt

**Verification:**

- Under **File Groups**, in the results pane, verify that the new **Specific Files** file group is listed and includes the different inclusion files.

**Testing file screens**
We recommend using the following two scenarios to familiarize yourself with some of the file screen tasks in File Server Resource Manager.

**Scenario 1: File screens on user folders**
File screens can be used to block specific types of files from being saved on shared storage resources. When a file screen is applied on a specific folder, the screening properties apply to all its subfolders.

**Test setup:**
If you have not already done so, perform the first scenario for testing auto quotas (see "Scenario 1: Auto quotas on user folders" under "Quota Management Scenarios" earlier in this guide).

When you have completed this scenario and have created all of the user folders, create a file screen on the Users folder, using the **Block Executable Files** template. For more information about how to create a file screen based on a file screen template, see "Creating a File Screen" earlier in this guide.

**Verification:**

- Under **File Screens**, in the results pane, verify that the new file screen for the Users folder is listed.
- Verify that you cannot save a file in the Users folder with any of the extensions included in the **Executable Files** file group (for example, .exe, .cmd, .bat, and so on).
- Verify that you cannot save the same type of files in any of the user folders (for example, in: D:\Users\User001).
- Create a new subfolder inside a user folder (for example: D:\Users\User001\Temp) and verify that you cannot save the same type of files in the newly created subfolder.

**Scenario 2: File screen exceptions on specific user folders**
File screen exceptions expand the flexibility of the file screening capabilities in File Server Resource Manager by creating an exception to any screening rules derived from a parent folder.

**Test setup:**
If you have not already done so, perform the first scenario for testing file screens (see "Scenario 1: File screens on user folders" earlier in this guide).

When you have completed this scenario, apply a file screen exception on one of the users folders (for example: D:\Users\User002) by selecting the **Executable Files** file group to be excluded from screening. For more information about how to create a file screen exception, see "Creating a File Screen Exception" earlier in this guide.

**Verification:**

- Under **File Screens**, in the results pane, verify that the new file screen exception is listed for the specific user folder that you selected.
- Verify that you can now save files in this folder with any of the extensions included in the **Executable Files** file group (for example, .exe, .cmd, .bat, and so on).
- Verify that you still cannot save a file in the Users folder (the parent folder of the one you selected for the file screen

exception) with any of the extensions included in the **Executable Files** file group.

**Testing file screen templates**
We recommend using the following two scenarios to familiarize yourself with some of the file screen template tasks in File Server Resource Manager.

**Scenario 1: A file screen to monitor executables and system files saved on a shared folder**
File screens can also be used to generate notifications when specific types of files are saved on a volume or on a folder, without blocking the user from saving them.

**Test setup:**
Using the **Monitor Executable and System Files** template, create a file screen on a shared folder on your server (for example, E:\Scratch). For more information about how to create a file screen based on a file screen template, see "Creating a File Screen" earlier in this guide.

**Verification:**

- Copy or generate a file in the shared folder, with any of the extensions included in the **Executable Files** file group (for example, .exe, .cmd, .bat, and so on) and then verify that the selected notifications have been created and received.

- Copy or generate a file in the shared folder, with any of the extensions included in the **System Files** file group (for example, .dll, .sys, .vxd, and so on) and then verify that the selected notifications have been created and received.

**Scenario 2: Using a file screen template to update file screens**
As with quota templates, if you base your file screens on file screen templates, you can automatically update all file screens that are based on a specific template simply by editing that template.

**Test setup:**
If you have not already done so, perform the first scenario for testing file screens (see "Scenario 1: File screens on user folders" earlier in this guide).

When you have completed this scenario, under **File Screen Templates**, edit the **Block Executable Files** template and select to also block files in the **Backup Files** file group. When prompted, select to apply the template changes to all derived file screens.

**Verification:**

- Under **File Screens**, in the results pane, verify that the file screen for the Users folder now includes **Backup Files** in the list of file groups to block.

- Verify that you cannot save a file in the Users folder or in its subfolders with any of the extensions included in the **Backup Files** file group (for example, .old, .bak, .bck).

- Under **File Groups**, select to edit the **Backup Files** group and add the following parameter in **Files to exclude**: test_file.bak

- Finally, verify that you can now save a file called test_file.bak in the Users folder and in any of its subfolders.

**Storage Reports Management Scenarios**
File Server Resource Manager offers advanced reporting capabilities to help you understand how storage is used and how quota and file screening tasks are behaving. The **Storage Reports Management** node of File Server Resource Manager includes all the necessary options to schedule and generate storage reports.

There are two types of report tasks that can be tested under **Storage Reports Management**:

- Scheduled reports tasks

- On-demand reports tasks

**Note**

For more information about **Storage Reports Management**, the different tasks available in this node, and step-by-step procedures, see "Generating Storage Reports" earlier in this guide.

**Testing scheduled reports**
We recommend using the following two scenarios to familiarize yourself with some of the scheduled reports tasks in File Server Resource Manager.

**Scenario 1: Scheduling a report to monitor quota usage**
Several preconfigured reports are included with File Server Resource Manager. One of these is the **Quota Usage** report. This report lists all quotas that exceed a selected percentage of the quota usage (for example, 85%). You can use the **Quota Usage** report to continuously monitor storage resources on your server and to identify those users that may soon exceed their quota limit.

**Test setup:**
If you have not already done so, perform any of the scenarios for testing quotas or auto quotas (see "Quota Management Scenarios" earlier in this guide).

When you are finished applying some quotas to a volume on your server, under **Storage Reports Management**, click the **Schedule a new report task** action and create a report task with the following parameters:

- **Scope:** the volume in your server where quotas have been applied (for example, D:\)
- **Report data:** Quota Usage (clear all other reports)
- **Edit Parameters:**
    - **Minimum quota usage:** 50%
- **Report formats:** DHTML
- **Delivery:** send the report to your e-mail address
- **Schedule:**
    - **Schedule Task:** Once
    - **Start time:** 5 to 10 minutes after your current time

**Verification**

- Under **Storage Reports Management**, in the results pane, verify that the new report task is listed and that the schedule and next run time values are those that you selected.
- Wait for the scheduled run time and then allow a few minutes for the report to be generated and for the e-mail to be sent and received. When the e-mail has been received, scan through the report information and scroll all the way to the end of the report to see the quota usage by owner data.

**Scenario 2: Adding a report to a scheduled report task**
Reports can be easily added or removed from existing report tasks.

**Test setup:**
If you have not already done so, perform the first scenario for testing scheduled reports (see " Scenario 1: Scheduling a report to monitor quota usage" earlier in this guide).

When you have completed this previous scenario, under **Storage Reports Management**, click the **Add or Remove reports for a report task** action. Edit the parameters of the **Large Files** report to include files with a minimum file size of 20 MB, and then add this report to the report task of the previous scenario.

**Verification**

- Under **Storage Reports Management**, in the results pane, verify that the report task now includes the **Large Files** report.

- Right-click the report task and select **Run report task now**. When asked how you want to proceed, select **Generate reports in the background**. Wait a few minutes for the reports to be generated and for the e-mails to be sent and received. Verify that you receive two e-mails: a separate e-mail for each report in the task.

**Testing on-demand reports**

We recommend using the following two scenarios to familiarize yourself with some of the on-demand reports tasks in File Server Resource Manager.

**Scenario 1: Running a report of all duplicate files in a volume**

A very useful report to run is the **Duplicate Files** report. This report not only lists all duplicate files in a volume or folder but also presents statistical data related to the space these duplicate files occupy on disk. You can use the information collected in this report to quickly reclaim wasted disk space.

**Test setup:**

Under **Storage Reports Management**, click the **Generate reports now** action and then choose the following parameters:

- **Scope:** the volume in your server where the system root is located (for example, C:\)

- **Report data:** Duplicate Files (clear all other reports)

- **Edit Parameters:**

    - **Minimum days since file screening event occurred:** change to 0 (zero)

- **Report formats:** DHTML

When asked how you want to proceed, select **Wait for reports to be generated and then display them**.

**Verification**

- Wait a few minutes for the report to be generated. When completed, it will open automatically in Internet Explorer.

- View the report information and scroll all the way to the end of the report to see the statistical data.

**Scenario 2: Running a report to monitor file screening activity**

File Server Resource Manager allows you to record file screening activity in an auditing database. The **File Screening Audit** report is used to view the information in this database.

| **Important** |
| --- |
| Before you run a **File Screen Audit** report, in the **File Server Resource Manager Options** dialog box, on the **File Screen Audit** tab, verify that the **Record file screening activity in the auditing database** check box is selected. For more information, see "Generating Storage Reports" earlier in this guide. |

**Test setup:**

If you have not already done so, perform any of the scenarios for testing file screening (see "File Screening Management Scenarios" earlier in this guide). Make sure that you generate some file screening activity by attempting to save blocked files.

Under **Storage Reports Management**, click the **Generate reports now** action and then choose the following parameters:

- **Scope:** the volume in your server where file screens have been applied (for example: D:\)

- **Report data:** File Screening Audit (clear all other reports)

- **Report formats:** HTML

When asked how you want to proceed, select **Wait for reports to be generated and then display them**.

**Verification**

- Wait a few minutes for the report to be generated. When completed, it will open automatically in Internet Explorer.
- View the report information and verify that the screening activity that you generated is listed.

**Clustering Scenarios**
The quota, file screening, and reporting capabilities in File Server Resource Manager are also available for clustered environments.

| **Note** |
| --- |
| You only need to perform these scenarios if you will be deploying File Server Resource Manager in a clustered environment. |

**Preliminary configuration guidelines**

- Quota templates, file screen templates, and file groups can be configured by connecting to any node in the cluster but not by using the virtual cluster name.
- Quotas and file screens on a particular volume must be configured by connecting to the node that currently owns that volume.

**Scenario 1: Global configuration**
**Test setup:**
Connect to any node in the cluster and create or make changes to a quota template, a file screen template, or a file group.

**Verification:**

- Connect to a different node in the cluster and verify that the changes you made have propagated to this node.

**Scenario 2: Volume failover**
**Test setup:**
Connect to any node in the cluster and create quotas and file screens on a volume owned by that node. Next, fail over that volume to a different cluster node.

**Verification:**

- Connect to the node that now owns the volume and verify that the quotas and file screens that you created continue to be enforced.

# Step-by-Step Guide to Setting Up Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Server for NIS enables a Microsoft Windows–based Active Directory domain controller to administer UNIX Network Information Service (NIS) networks. This guide describes the tasks you must perform to install and configure Server for NIS on your network.

**In this Guide**


- Task 1: Install Server for NIS on a Domain Controller
- Task 2: Migrate NIS Maps to Active Directory
- Task 3: Set the Frequency of Map Propagation
- Task 4: Set the Encryption Method for NIS Domains


**Requirements for Installing Server for NIS**

You can install Server for NIS on the following Windows Server operating systems:

- Windows Server 2003
- Windows Server 2003 R2
- Windows 2000 Server

Server for NIS is not available for installation except on Active Directory domain controllers.

You must install Server for NIS on a partition that is formatted with the NTFS file system. If you are installing Server for NIS as an upgrade to a previous version that was installed on a partition formatted with the FAT file system, you must convert the FAT partition to NTFS before you can perform the upgrade. File system operations on FAT partitions are not supported.

Server for NIS requires 5 MB of free hard disk space. It is recommended that the computer have at least 16 MB of RAM in addition to the recommended minimum configuration for the operating system.

Server for NIS cannot be run from a network server. All files must be installed on the local computer.

# Task 1: Install Server for NIS on a Domain Controller

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Server for Network Information Service (NIS) must be installed on a Windows Server operating system that is an Active Directory domain controller. The option to install Server for NIS is unavailable on a computer unless it is a domain controller.

This topic contains the following sections:


Verify that your computer is a domain controller

Promote your computer to a domain controller

Install Server for NIS

Continue with Server for NIS Setup


**Verify that your computer is a domain controller**

If you are uncertain whether the computer on which you want to install Server for NIS is a domain controller, click **Start**, click **Administrative Tools**, and then click **Manage Your Server**. The **Manage Your Server** window lists the roles installed on the computer under the text **Your server has been configured with the following roles**, and specifies whether a computer is a domain controller.

Domain controllers also include tools installed with **Administrative Tools** that are not available on other computers, such as **Active Directory Domains and Trusts**, **Active Directory Sites and Services**, and **Active Directory Users and Computers**.

If the computer is not a domain controller, follow the procedure Promote your computer to a domain controller. Otherwise, proceed to Install Server for NIS.

**Promote your computer to a domain controller**

If the computer on which you want to install Server for NIS is not a domain controller, perform the steps in this section. Otherwise, proceed to Install Server for NIS.

| **Important** |
| --- |
| Because promoting your computer to a domain controller requires domain administrator access, and can remove the computer from your network, read Domain controllers and all related topics, in the Active Directory Help before promoting your computer to a domain controller. |

1. Click **Start**, click **Run**, type **dcpromo** in the text box, and then click **OK**.

   The Active Directory Installation Wizard opens.

2. Click **Next**.

3. On the **Operating System Compatibility** page, read the information and then click **Next**.

   If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.

4. On the **Domain Controller Type** page, click to select one of two options:

   - **Additional domain controller for an existing domain**

     This option requires that you are a member of the Domain Admins group for the target domain.

     If you choose this option, keep in mind that if Server for NIS is already running as an NIS master server on a domain controller within the existing domain, then Server for NIS must function as an NIS subordinate (also known as slave) server on this computer.

- **Domain controller for a new domain**

  If you choose this option, you must configure this installation of Server for NIS as the NIS master server, until one or more additional domain controllers are configured within the new domain.

5. Click **Next**.

   - If you chose **Additional domain controller for an existing domain**, follow the procedure Create an additional domain controller in the Windows Server 2003 Help, starting with Step 4.

   - If you chose **Domain controller for a new domain**, go on to the next step in this section.

6. On the **Create New Domain** page, choose one of the following options:

| **Note** |
| --- |
| If you choose to create a new domain, you must configure this installation of Server for NIS as the NIS master server, until one or more additional domain controllers are configured within the new domain. |

- **Child domain in an existing domain tree**

  You must be a member of the **Enterprise Admins** group to continue with this installation.

- **Domain tree in an existing forest**

  You must be a member of the **Enterprise Admins** group to continue with this installation.

- **Domain in a new forest**

  Creating a new forest requires some advance planning. Before creating a new forest, decide on a practical Domain Name System (DNS) name for this computer, as well as a NetBIOS name. For more information, see Namespace planning for DNS in Windows Server 2003 Help.

1. Click **Next**.

2. Complete setup using the Windows Server 2003 Help procedure for the domain option you selected in Step 6 of this section.

   - If you selected **Child domain in an existing domain tree**, follow the steps from Step 5 in Create a new child domain.

   - If you selected **Domain tree in an existing forest**, follow the steps from Step 5 in Create a new domain tree.

   - If you selected **Domain in a new forest**, follow the steps from Step 5 in Create a new forest.

3. When you have completed the Active Directory Installation Wizard and successfully configured your domain controller, proceed to Install Server for NIS.

**Install Server for NIS**

You can install Server for NIS using either the Windows interface, or by using command-line tools.

To install Server for NIS by using the Windows interface

To install Server for NIS by using command-line tools

**To install Server for NIS by using the Windows interface**

Before installing Server for NIS, have your Windows Server 2003 product CD available, or have available the network path to your Windows Server 2003 R2 installation files. The Server for NIS installation process may prompt you to locate required files that are stored on the product CD.

1. Click **Start**, click **Control Panel**, and then click **Add or Remove Programs**.

2. Click **Add/Remove Windows Components**.

3. When the Windows Components Wizard opens, click to fill the **Active Directory Services** check box.

4. With **Active Directory Services** highlighted, click **Details**.

5. When the **Active Directory Services** dialog box opens, click to fill the **Identity Management for UNIX** check box.

6. With **Identity Management for UNIX** highlighted, click **Details**.

7. When the **Identity Management for UNIX** dialog box opens, click to fill the **Server for NIS** check box.

   Because Administration Components are required for Server for NIS operation, this item is automatically selected for installation when you select Server for NIS.

8. Click **OK**.

   The Windows Components Wizard begins installing the Identity Management for UNIX components you selected.

9. If you are prompted to locate files required for Server for NIS installation, insert the Windows Server 2003 R2 product CD, or browse to the network location of your Windows Server 2003 R2 installation files.

10. If NIS map data compatible with Windows Services for UNIX 3.5 is stored in Active Directory, the Server for NIS installation process automatically migrates the data, and uses it to populate Server for NIS maps.

11. When installation completes, restart your computer to begin working with Server for NIS.

**To install Server for NIS by using command-line tools**

You can use an answer file to run the Identity Management for UNIX Installation Wizard without your having to be present for interactive responses. Create the answer file, and then run it at a command prompt on the installation computer.

Have your Windows Server 2003 product CD available for this command-line installation.

Perform the following steps to install Server for NIS by using a command-line environment:

1. Create an unattended answer file in %windir% with the following contents.

   | **Note** |
   | --- |
   | The unattended answer file is a plain text file that Windows Server uses to respond to prompts about your installation preferences. For more information about unattended installations, see Unattended Installation Fundamentals in the Windows Server 2003 Deployment Guide. |

   **[Components]**

   **Snis=on**

   **Psync=on**

**ldmumgmt=on**

2. At a command prompt, type the following, and press **Enter**.

   **synocmgr /i:%***windir***%\inf\sysoc.inf /u:<***answerfile***.txt> /q**

**Continue with Server for NIS Setup**
To continue setting up Server for NIS, go on to Task 2: Migrate NIS Maps to Active Directory.

**See Also**
**Other Resources**
Checklist: Installing a domain controller
Installing and uninstalling Server for NIS
How Unattended Installation Works
Unattended Installation Fundamentals

# Task 2: Migrate NIS Maps to Active Directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Server for Network Information Service (NIS) allows NIS map data to be migrated to Active Directory. After the migration, and upon receiving NIS requests from clients or other servers, Server for NIS searches Active Directory to reply to NIS queries.

This topic contains the following sections:


Planning for NIS Migration

Migrate NIS Maps to a Windows-based Server

Configure UNIX NIS Servers to use Server for NIS as the Master Server

Creating the Structure of Nonstandard Maps

Additional Server for NIS Configuration Steps


**Planning for NIS Migration**
Before beginning NIS migration to Server for NIS, it is strongly recommended that you read Checklist: NIS migration to Active Directory using the NIS Data Migration wizard.

**Steps in a Typical Migration**
Migration consists of the following three procedures:

1. Migrate NIS maps to a Windows-based server.

   - Using a command line

     Server for NIS includes a command-line tool called **nis2ad** to migrate maps from UNIX-based NIS servers to Active Directory–based Server for NIS.

   - Using the Windows interface

     Server for NIS includes a migration wizard that extracts the information necessary to perform the migration. Even when using the migration wizard, however, you must complete steps 2 and 3, which follow.

   The migration wizard and the **nis2ad** command read map data from NIS map source files, which are the plain text files from which the NIS map databases are compiled. These source files must be stored in a location that can be accessed by the domain controller during migration, such as on a disk on the domain controller or in a shared directory accessible by the domain controller.

   If the map you want to migrate is a nonstandard NIS map, create the structure using the procedure Creating the Structure of Nonstandard Maps below.

2. Configure UNIX NIS Servers to Use Server for NIS as the Master Server.

   After the migration, the original UNIX-based NIS server must send an update of maps to all subordinate NIS servers, with the name of the new master server in the maps.

3. Disable the original NIS server.

UNIX-based subordinate NIS servers can continue to work as before; however, they will receive map updates from the Windows-based computer running Server for NIS instead of the UNIX-based computer. Client computers running UNIX-based operating systems can be configured to get NIS maps or data from the new master server.

**Read more about NIS Migration**
Before you begin migrating NIS map data, it is recommended that you read the following conceptual topics that discuss NIS

migration:

Migrating NIS to Active Directory

Migrating standard and nonstandard maps

Resolving migration conflicts

Handling special users during migration to Active Directory

**Migrate NIS Maps to a Windows-based Server**
Perform the following steps to migrate NIS maps to a Windows-based server on which Server for NIS has been installed.

Using the Windows interface

Using a command line

**Using the Windows interface**

1. Open the Identity Management for UNIX management console by doing one of the following:

   - Click **Start**, click **Administrative Tools**, and then click **Identity Management for UNIX**.
   - Click **Start**, click **Run**, type **idmumgmt.msc** in the **Open** text box, then click **OK**.

2. In the hierarchy tree, open the **Microsoft Identity Management for UNIX** node, and click to highlight the **Server for NIS** node.

3. Start the NIS Data Migration Wizard by doing one of the following:

   - Right-click the **Server for NIS** node, and then click **NIS Data Migration Wizard**.
   - With the **Server for NIS** node highlighted, click **NIS Data Migration Wizard** in the **Actions** pane.
   - On the **Actions** menu, click **NIS Data Migration Wizard**.

4. Follow the step-by-step directions in the wizard.

| **Note** |
| --- |
| Password file entries with names longer than eight characters will not be migrated.<br><br>Windows user accounts created as a result of the migration are disabled. After performing the migration, you must enable the accounts. For security reasons, it is recommended that you assign a temporary password to these accounts and instruct the affected users to change their Windows password as soon as possible. |

**Using a command line**

1. Open a **Command Prompt** window in one of the following two ways:

   - Click **Start**, and then click **Command Prompt** on the **Start** menu.
   - Click **Start**, click **Run**, type **cmd** in the **Open** text box, and click **OK**.

2. At a command prompt, type:

   **nis2ad -y** *UNIXNISDomain* **-a** *ActiveDirectoryNISDomain* **[***Options***]** *MapfileToMigrate*

The following arguments are required:

| Argument | Description |
|---|---|
| **-y** *UNIXNISDomain* | Specifies the name of the NIS domain that contains the map to migrate. |
| **-a** *ActiveDirectoryNISDomain* | Specifies the NIS domain name in Active Directory. |
| *MapfileToMigrate* | Specifies the name of the NIS map source file to migrate. NIS map source files are the plain text files from which the NIS map databases are compiled. |

The **nis2ad** command accepts the following options.

| Option | Description |
|---|---|
| **-m** | Perform the migration. If this option is omitted, the program finds and reports conflicts but does not actually perform the migration. |
| **-c** *FileName* | Specifies the file where conflict details are written. Uses a default file (%windir%\idmu\nis\conflicts.log) if not specified. |
| **-t** *TargetContainer* | Specifies the target container name. Applicable only when creating a new NIS domain. If not specified, uses the default or uses the container of the target domain. |
| **-f** *FileName* | Specifies the name of the log file. If not specified, **nis2ad** uses a default file (%windir%\idmu\nis2ad.log). |
| **-r yes\|no** | Replace object in Active Directory with object being migrated. Default is no. |
| **-n** | Resolves conflicts by changing the Windows account name in Active Directory. If objects of different types have the same name, the names of both objects are changed before the data is migrated. |
| **-p** *Password* | If needed and if not specified, the user will be prompted. |
| **-d** *Directory* | Specifies the path of the directory that contains NIS map source files. |
| **-s** *Server* | Specifies the domain controller server hosting Active Directory. Otherwise use the current server. |
| **-u** *User* | Specifies the name of the user having administrator privileges on this computer. If not specified, **nis2ad** uses the current user.<br><br>Even if you specify another user by using the **-u** option, the currently logged-on user must have write permissions for the folder that will contain the log and conflict files. If necessary, modify the permissions on the folder to grant write access to the user who will be running the **nis2ad** utility, before running the utility. |

**Note**

To view the complete syntax for this command, at a command prompt, type **nis2ad /?**

You can migrate only one map at a time using **nis2ad**. To migrate more than one map at a time, use the NIS Data Migration Wizard.

Password file entries with names longer than eight characters will not be migrated.

Windows user accounts created as a result of the migration will be disabled. After performing the migration, you must enable the accounts. For security reasons, it is recommended that you assign a temporary password to these accounts and instruct the affected users to change their Windows password as soon as possible.

**Configure UNIX NIS Servers to use Server for NIS as the Master Server**

To change a UNIX-based NIS server from a master server to a subordinate (also known as slave) server, follow these steps:

1. Migrate NIS maps to a Windows-based computer running Server for NIS.

2. Transfer the maps from the old master server to other subordinate NIS servers by providing the name of the new Server

for NIS for each map. At a command prompt, type:

**ypxfr –h** *newserver mapname*

where *newserver* is the name of the new NIS master server, and *mapname* is the name of the map to be transferred.

3. Run this command for each map on each of the subordinate servers. After this step, the UNIX subordinate servers will recognize the new Server for NIS master server.

**Creating the Structure of Nonstandard Maps**

You can migrate nonstandard maps to Server for NIS using either the Windows-based NIS Data Migration Wizard, or a command-line environment.

Using the NIS Data Migration Wizard

Using the command line

| **Important** |
| --- |
| After the map structure is created using this procedure, there is no way to remove it.<br><br>This structure applies to all NIS domains, so it is important to ensure that the format is consistent across all NIS domains. |

**Using the NIS Data Migration Wizard**

1. Open the Identity Management for UNIX MMC snap-in by doing one of the following:

   - Click **Start**, click **Administrative Tools**, and then click **Identity Management for UNIX**.
   - Click **Start**, click **Run**, type **idmumgmt.msc** in the **Open** text box, then click **OK**.

2. In the hierarchy tree, open the **Microsoft Identity Management for UNIX** node, and click to highlight the **Server for NIS** node.

3. Start the NIS Data Migration Wizard by doing one of the following:

   - Right-click the **Server for NIS** node, and then click **NIS Data Migration Wizard**.
   - With the **Server for NIS** node highlighted, click **NIS Data Migration Wizard** in the **Actions** pane.
   - On the **Actions** menu, click **NIS Data Migration Wizard**.

4. Follow the step-by-step directions in the wizard.

5. On the **NIS Map Selection** panel, click **New**.

6. In the **Add Nonstandard Map** dialog box, do the following:

   - In the **Map name** string box, enter the name of the existing nonstandard map you want to migrate to Server for NIS. The map migration process assigns the same name to a new file containing your map structure.
   - In the **Separator** string box, type the single character you want to use to delimit or separate fields in your map structure. Suggested characters include a semicolon (**;**) or a dash (**-**).
   - In the **Key field** string box, type the number of the column you want to use as the map key. Use Arabic numeral

characters; do not spell out the number.

- Click **Next**.

7. In the **Location of UNIX NIS Map Source Files** window, enter the directory path name in which the map file you created in Step 6 is located, and then click **Next**.

8. Click **Finish** to start migrating map data from the existing nonstandard map to the new map file.

| **Note** |
| --- |
| Password file entries with names longer than eight characters will not be migrated.<br><br>Windows user accounts created as a result of the migration are disabled. After performing the migration, you must enable the accounts. For security reasons, it is recommended that you assign a temporary password to these accounts and instruct the affected users to change their Windows password as soon as possible. |

**Using the command line**

1. Open a **Command Prompt** window in one of the following two ways:

   - Click **Start**, and then click **Command Prompt** on the **Start** menu.

   - Click **Start**, click **Run**, type **cmd** into the **Open** text box, and click **OK**.

2. At a command prompt, type:

   **nismap create –i** *fieldNumber* **–g "***separator***"** *mapName*

| **Note** |
| --- |
| Do not use the hash character (#) as a field separator because this character is used in standard maps to mark the beginning of a comment. |

The following table shows the acceptable arguments for the **nismap create** command.

| **Argument** | **Description** |
| --- | --- |
| *fieldNumber* | The number of the field that contains the key to the map. |
| **"** *separator* **"** | The character used to separate fields, in quotation marks. To specify a space as a separator, enclose the space in double quotation marks (**" "**). For example:<br><br>**nismap create –i 1 –g " " Phones**<br><br>creates a map called Phones in which the key field is the first field and the separator character is a space. Other white-space characters, such as tab, are also accepted. |
| *mapName* | The name of the map. |

| **Note** |
| --- |
| To view the complete syntax for this command, at a command prompt, type:<br><br>**nismap /?** |

**Additional Server for NIS Configuration Steps**

When you have successfully completed NIS data migration, proceed to Task 3: Set the Frequency of Map Propagation to configure the time interval at which you want NIS maps propagated across the domain.

If you encountered problems during migration, see Server for NIS Troubleshooting.

**See Also**

**Other Resources**

Migrating NIS to Active Directory
Migrating standard and nonstandard maps
Remove a nonstandard NIS map
Internet Engineering Task Force Web site

# Task 3: Set the Frequency of Map Propagation

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

After you have completed migrating Network Information Service (NIS) maps to a Windows-based server running Server for NIS, you can ensure that your maps refresh across the network at regular intervals by setting how often maps are propagated to subordinate (also known as slave) servers.

This topic contains the following sections:

Change the Frequency of Map Updates to UNIX Subordinate NIS servers

Propagating Maps Immediately

Completing Server for NIS Configuration

**Change the Frequency of Map Updates to UNIX Subordinate NIS servers**

Perform the following steps to change the frequency of map updates to UNIX-based subordinate (also known as slave) NIS servers.

Using the Windows interface

Using a command line

**Using the Windows interface**

1. Open the Identity Management for UNIX management console by doing one of the following:

   - Click **Start**, click **Administrative Tools**, and then click **Identity Management for UNIX**.
   - Click **Start**, click **Run**, type **idmumgmt.msc** in the **Open** text box, then click **OK**.

2. If necessary, connect to the computer you want to manage by right-clicking the **Identity Management for UNIX** node in the hierarchy pane, and then clicking **Connect to another computer**. Otherwise, go on to the next step.

3. Click **Server for NIS** in the hierarchy pane.

4. Open **Map Updates** by doing one of the following:

   - Right-click the **Server for NIS** node, and then click **Map Updates**.
   - With the **Server for NIS** node highlighted, click **Map Updates** in the **Actions** pane.

5. In the **Server for NIS Properties** dialog box, type the number of days, hours and minutes you want to lapse between map updates.

6. Click **OK** to save your changes.

**Using a command line**

1. Open a **Command Prompt** window in one of the following two ways:

   - Click **Start**, and then click **Command Prompt** on the **Start** menu.
   - Click **Start**, click **Run**, type **cmd** into the **Open** text box, and click **OK**.

2. At a command prompt, type:

   **nisadmin config pushint=**[[*days***:**]*hh***:**]*mm* [**-s** *Server*] [**-u** *User*] [**-p** *Password*]

The following table contains the arguments for the nisadmin command.

| Argument | Description |
|---|---|
| [[*days***:**]*hh***:**]*mm* | Specifies the interval at which the service checks changes to NIS maps in Active Directory and propagates them to secondary NIS servers for all domains, in days, hours, and minutes. If *hh* is specified, *hh* must be in the range 0–23 and *mm* must be in the range 0–59. |
| *Server* | The master server for the domain. |
| *User* | The name of the user who has administrative privileges on the server to be started, if different from the current user. |
| *Password* | The password of the user who has administrative privileges on the server to be started, if different from the current user. If you type a user name but omit the password, you will be prompted for the password. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**nisadmin /?** |

**Propagating Maps Immediately**

You do not have to wait for the **Map Updates** interval to expire to refresh maps. You can propagate maps immediately using one of the following two procedures:

Propagate changed maps now

Propagate selected maps now

**Propagate changed maps now**

You can propagate any changed maps immediately either by using the Windows interface, or in a command-line environment.

Using the Windows interface

Using a command line

**Using the Windows interface**

1. Open the Identity Management for UNIX management console by doing one of the following:

   - Click **Start**, click **Administrative Tools**, and then click **Identity Management for UNIX**.
   - Click **Start**, click **Run**, type **idmumgmt.msc** in the **Open** text box, then click **OK**.

2. If necessary, connect to the computer you want to manage by right-clicking the **Identity Management for UNIX** node in the hierarchy pane, and then clicking **Connect to another computer**. Otherwise, go on to step 3.

3. Click **Server for NIS** in the hierarchy pane.

4. Click **Check for updates now**.

**Using a command line**

1. Open a **Command Prompt** window in one of the following two ways:

- Click **Start**, and then click **Command Prompt** on the **Start** menu.
- Click **Start**, click **Run**, type **cmd** into the **Open** text box, and click **OK**.

2. At a command prompt, type:

   **nisadmin** [*server*] **syncall** [**–u** *user* [**–p** *password*]]

   The following arguments are acceptable with the **nisadmin syncall** command.

| Argument | Description |
|---|---|
| **syncall** | Propagate all maps. |
| *server* | The name of the server where the maps are stored. |
| *user* | The name of the user who has administrator permissions on the server, if different from the current user. |
| *password* | The password of the user who has administrator permissions on the server, if different from the current user. If you type a user name but omit the password, you will be prompted for the password. |

**Note**

The **nisadmin syncall** command propagates maps only on UNIX-based NIS subordinate servers. It does not propagate maps on NIS subordinate servers running Windows operating systems. Active Directory updates Windows-based NIS subordinate servers.

To view the complete syntax for this command, at a command prompt, type:

**nisadmin /?**

**Propagate selected maps now**
**Using the Windows interface**

1. Open the Identity Management for UNIX management console by doing one of the following:
   - Click **Start**, click **Administrative Tools**, and then click **Identity Management for UNIX**.
   - Click **Start**, click **Run**, type **idmumgmt.msc** in the **Open** text box, then click **OK**.
2. If necessary, connect to the computer you want to manage by right-clicking the **Identity Management for UNIX** node in the hierarchy pane, and then clicking **Connect to another computer**. Otherwise, go on to step 3.
3. In the console tree, expand **Server for NIS** and view the list of available domains.
4. Expand the domain of interest, and click the **NIS Maps** object in the hierarchy pane to view NIS maps in the details pane.
5. In the **NIS Maps in this Domain** list, click to select a map you want to update immediately.

**Note**

To select multiple maps at once, click and drag, or press and hold the **Ctrl** key while selecting additional maps.

1. Propagate the selected maps by doing one of the following:

- Click **Propagate** in the **Actions** pane.
- On the **Actions** menu, click **Propagate**.
- Right-click the selected map object(s), then click **Propagate**.

2. On the **Identity Management for UNIX** dialog box, click **Yes** to begin map propagation.

**Using a command line**

1. Open a **Command Prompt** window in one of the following two ways:

   - Click **Start**, and then click **Command Prompt** on the **Start** menu.
   - Click **Start**, click **Run**, type **cmd** into the **Open** text box, and click **OK**.

2. At a command prompt, type:

   **yppush** [**-d** *ActiveDirectoryNISDomain*] [**-q**] [**-t** *Timeout*] [**-h** *Hosts*] *MapName*

   The following table shows the arguments accepted by the **yppush** command.

| Argument | Description |
|---|---|
| **-d** *ActiveDirectoryNISDomain* | NIS domain name in Active Directory. |
| **-q** | Quiet mode. Do not wait for response from subordinate (slave) servers and do not report errors. |
| **-t** *Timeout* | The number of seconds to wait for a response from the subordinate server before sending the next request. Must be greater than zero. The default value is **30**. |
| **-h** *Hosts* | The hosts to notify of changes. Default is all subordinate servers in the domain. Can be used multiple times for more than one computer. |
| *MapName* | The name of the NIS map to be transferred. |

| **Note** |
|---|
| The **yppush** command propagates maps only on UNIX-based NIS subordinate servers. It does not propagate maps on NIS subordinate servers running Windows operating systems. Active Directory updates Windows-based NIS subordinate servers.<br><br>To view the complete syntax for this command, at a command prompt type:<br><br>**yppush /?** |

**Completing Server for NIS Configuration**

The final step in initial setup of Server for NIS is to specify an encryption method for user passwords. To learn how to perform this task, see Task 4: Set the Encryption Method for NIS Domains.

**See Also**

**Other Resources**

Sending periodic map updates to subordinate (slave) NIS servers
Change the frequency of map updates to UNIX subordinate (slave) NIS servers
Manage NIS Maps

# Task 4: Set the Encryption Method for NIS Domains

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Server for Network Information Service (NIS) provides limited support for keeping passwords synchronized between a user's Windows and UNIX accounts. Whenever a user's Windows password is changed, Password Synchronization (which is installed with Server for NIS for this reason) captures the new password, encrypts it, and then stores the password in the passwd map in Active Directory. The new password is propagated to NIS subordinate (also known as slave) servers either during the next scheduled update (if a propagation interval is configured), or by using commands that propagate maps immediately. For more information about configuring map propagation, see Task 3: Set the Frequency of Map Propagation.

When synchronizing passwords, Server for NIS can use either crypt(3) (refers to DES encryption) or Message Digest 5 (MD5) encryption. Server for NIS can support different encryption methods for multiple domains, but all UNIX computers in a particular domain must use the same encryption method.

This topic contains the following sections:

Setting the Encryption Method for a Domain

Completing Server for NIS Configuration Tasks

**Setting the Encryption Method for a Domain**

You can set the encryption method for a domain either by using the Windows interface, or working in a command-line environment.

Using the Windows interface

Using a command line

**Using the Windows interface**

1. Open the Identity Management for UNIX management console by doing one of the following:
   - Click **Start**, click **Administrative Tools**, and then click **Identity Management for UNIX**.
   - Click **Start**, click **Run**, type **idmumgmt.msc** in the **Open** text box, then click **OK**.
2. If necessary, connect to the computer you want to manage by right-clicking the **Identity Management for UNIX** node in the hierarchy pane, and then clicking **Connect to another computer**. Otherwise, go on to Step 3.
3. In the console tree, expand **Server for NIS** and view the list of NIS domains.
4. Select the domain for which you want to set an encryption method.
5. Open the **UNIX Password Encryption Properties** dialog box by doing one of the following:
   - Click **UNIX Password Encryption** in the **Actions** pane.
   - On the **Actions** menu, click **UNIX Password Encryption**.
   - Right-click the selected domain, then click **UNIX Password Encryption**.
6. In the **Encryption Scheme** area, click the drop-down menu to select the encryption method used by all UNIX computers in the domain.

**Note**

You can select the MD5 encryption method for a UNIX domain that consists exclusively of computers running Linux and using MD5 encryption. Domains that contain one or more computers using the **crypt** algorithm or that run any other operating system must use **crypt**. Although Linux versions 6.2 and later support MD5 encryption, Identity Management for UNIX is not supported for versions of Linux prior to version 8.

**Using a command line**

1. Open a **Command Prompt** window in one of the following two ways:

   - Click **Start**, and then click **Command Prompt** on the **Start** menu.
   - Click **Start**, click **Run**, type **cmd** into the **Open** text box, and click **OK**.

2. At a command prompt, type:

   **nisadmin** [*computer*] **encryptiontype -d** *domain* {**crypt** | **md5**} [**-u** *usr* [**-p** *pword*]]

| Argument | Description |
|---|---|
| *computer* | Specifies the remote computer you want to administer. You can specify the computer using a WINS or DNS name, or by Internet Protocol (IP) address. |
| *domain* | Specifies the name of the domain for which the change is being made. |
| *usr* | Specifies the user name of the user whose credentials are to be used. It might be necessary to add the domain name to the user name in the form *domain******username*. |
| *pword* | Specifies the password of the user specified using the **-u** option. If you specify the **-u** option but omit the **-p** option, you are prompted for the user's password. |

**Note**

To view the complete syntax for this command, at a command prompt, type:

**nisadmin /?**

You can select the MD5 encryption method for a UNIX domain that consists exclusively of computers running Linux and using MD5 encryption. Domains that contain one or more computers using **crypt** or that run any other operating system must use **crypt**. Although Linux versions 6.2 and later support MD5 encryption, Identity Management for UNIX is not supported for versions of Linux prior to version 8.

**Completing Server for NIS Configuration Tasks**

You have completed all the tasks for initial setup of Server for NIS. If you want to configure Server for NIS on another computer, refer to the start of the Step-by-Step Guide to Setting Up Server for NIS.

**See Also**

**Other Resources**

Password encryption

Set the encryption method for a domain

# Step-by-Step Guide to Deploying Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Password Synchronization helps integrate Windows and UNIX networks by simplifying the process of maintaining secure passwords in both environments. Users are freed of the difficulty of maintaining separate passwords for their Windows and UNIX accounts or having to remember to change the password wherever it is used.

With Password Synchronization, whenever a user's password is changed on a Windows-based computer or domain, the password can also be automatically changed on every UNIX host for which the user has an account. Password Synchronization can also be configured to change the user's Windows password when the user's UNIX password is changed.

**Password Synchronization Components**
Password Synchronization is a combination of three software components:

- The Password Synchronization service running on one or more Windows-based computers
- The Password Synchronization daemon running on one or more UNIX computers
- The Password Synchronization pluggable authentication module (PAM) installed on one or more UNIX computers.

**Password Synchronization Deployment Scenarios**
You can deploy Password Synchronization in the following scenarios:

- You can use Password Synchronization to synchronize passwords between one or more Windows-based computers and one or more stand-alone UNIX hosts.
- You can use Password Synchronization to synchronize passwords within an NIS domain that has a UNIX-based NIS master server. For more information about synchronizing passwords within NIS domains, see Synchronizing passwords within an NIS domain.
- You can use Password Synchronization to synchronize passwords within an NIS domain that has a Windows-based master server.

**In this Guide**
Task 1: Install Password Synchronization on a Domain Controller

Task 2: Install the Password Synchronization Daemon on UNIX-based Computers

Task 3: Install the Pluggable Authentication Module on UNIX-based Computers

Task 4: Configure Password Synchronization

Task 5: Start Password Synchronization

**Requirements for Installing Password Synchronization**
This Step-by-Step guide is intended to help you install and configure Password Synchronization on computers running Windows Server 2003 R2.

Password Synchronization is not available for installation except on Active Directory domain controllers.

You must be logged on to the computer as a member of both the Schema Administrators and Enterprise Administrators groups to install Password Synchronization. The Password Synchronization installation process verifies that users have both Schema Administrator and Enterprise Administrator access permissions before installation can progress.

You must install Password Synchronization on a partition that is formatted with the NTFS file system. If you are installing Password Synchronization as an upgrade to a previous version that was installed on a partition formatted with the FAT file system, you must convert the FAT partition to NTFS before you can perform the upgrade. File system operations on FAT partitions are not supported.

Password Synchronization requires 3 MB of free hard disk space. It is recommended that the computer have at least 16 MB of RAM in addition to the recommended minimum configuration for the operating system.

Password Synchronization cannot be run from a network server. All files must be installed on the local computer.

Before you install and configure Password Synchronization, it is strongly recommended that you read Best practices for Password Synchronization in the Password Synchronization Help.

**Important Security Considerations for Password Synchronization**
It is strongly recommended that you perform the Windows Server 2003 Service Pack 1 (SP1) compatibility check when selecting **Enable Windows to NIS (AD) Password Sync** on the **Configuration** tab in the **Password Synchronization Properties** dialog box. To protect the security of user account passwords in your enterprise, it is strongly recommended that you allow Password Synchronization to identify all domain controllers in the forest that are not running Windows Server 2003 SP1 or later releases.

Password Synchronization prompts you to allow the compatibility check when you select the **Enable Windows to NIS (AD) Password Sync** option. With Windows Server 2003 SP1 or a later release installed on all the domain controllers in a forest, the risk of exposing user password hashes to unauthorized viewers is greatly reduced. When Windows Server 2003 SP1 is not the minimum functional level of all the domain controllers in a forest, it is possible for any authenticated user on the domain to view the password hash for any UNIX user whose account has been migrated to Active Directory.

In the event an unauthorized user breaks the password hash for a UNIX-based user account in Active Directory, the Windows-based password for the account is no longer secure.

**See Also**
**Other Resources**
Understanding Password Synchronization

# Task 1: Install Password Synchronization on a Domain Controller

Applies To: Windows Server 2003 R2

This topic contains the following sections:

Verify that your computer is a domain controller

Promote your computer to a domain controller

Install Password Synchronization

Continue with Password Synchronization Setup

**Verify that your computer is a domain controller**
If you are uncertain whether the computer on which you want to install Password Synchronization is a domain controller, click **Start**, click **Administrative Tools**, and then click **Manage Your Server**. The **Manage Your Server** window lists the roles installed on the computer under the text **Your server has been configured with the following roles**, and specifies whether a computer is a domain controller.

Domain controllers also include tools installed with **Administrative Tools** that are not available on other computers, such as **Active Directory Domains and Trusts**, **Active Directory Sites and Services**, and **Active Directory Users and Computers**.

If the computer is not a domain controller, follow the procedure Promote your computer to a domain controller. Otherwise, proceed to Install Password Synchronization.

**Promote your computer to a domain controller**
If the computer on which you want to install Password Synchronization is not a domain controller, perform the steps in this section. Otherwise, proceed to Install Password Synchronization.

| **Important** |
| --- |
| Because promoting your computer to a domain controller requires domain administrator access, and can remove the computer from your network, read Domain controllers and all related topics, in the Active Directory Help before promoting your computer to a domain controller. |

1. Click **Start**, click **Run**, type **dcpromo** in the text box, and then click **OK**.

   The Active Directory Installation Wizard opens.

2. Click **Next**.

3. On the **Operating System Compatibility** page, read the information and then click **Next**.

   If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.

4. On the **Domain Controller Type** page, click to select one of two options:

   - **Additional domain controller for an existing domain**

     This option requires that you are a member of the Domain Administrators group for the target domain, or the Enterprise Administrators group in Active Directory.

   - **Domain controller for a new domain**

5. Click **Next**.

- If you chose **Additional domain controller for an existing domain**, follow the procedure Create an additional domain controller in the Windows Server 2003 Help, starting with Step 4.

- If you chose **Domain controller for a new domain**, go on to the next step in this section.

6. On the **Create New Domain** page, choose one of the following options:

   - **Child domain in an existing domain tree**

     You must be a member of the **Enterprise Administrators** group to continue with this installation.

   - **Domain tree in an existing forest**

     You must be a member of the **Enterprise Administrators** group to continue with this installation.

   - **Domain in a new forest**

     Creating a new forest requires some advance planning. Before creating a new forest, decide on a practical Domain Name System (DNS) name for this computer, as well as a NetBIOS name. For more information, see Namespace planning for DNS in Windows Server 2003 Help.

7. Click **Next**.

8. Complete setup using the Windows Server 2003 Help procedure for the domain option you selected in Step 6 of this section.

   - If you selected **Child domain in an existing domain tree**, follow the steps from Step 5 in Create a new child domain.

   - If you selected **Domain tree in an existing forest**, follow the steps from Step 5 in Create a new domain tree.

   - If you selected **Domain in a new forest**, follow the steps from Step 5 in Create a new forest.

9. When you have completed the Active Directory Installation Wizard and successfully configured your domain controller, proceed to Install Password Synchronization.

**Install Password Synchronization**

You can install Password Synchronization using either the Windows interface, or by using command-line tools.

To install Password Synchronization by using the Windows interface

To install Password Synchronization by using command-line tools

**To install Password Synchronization by using the Windows interface**

Before installing Password Synchronization, have your Windows Server 2003 R2 product CD available, or have available the network path to your Windows Server 2003 R2 installation files. The Password Synchronization installation process may prompt you to locate required files that are stored on the product CD.

1. Click **Start**, click **Control Panel**, and then click **Add or Remove Programs**.

2. Click **Add/Remove Windows Components**.

3. When the Windows Components Wizard opens, click to fill the **Active Directory Services** check box.

4. With **Active Directory Services** highlighted, click **Details**.

5. When the **Active Directory Services** dialog box opens, click to fill the **Identity Management for UNIX** check box.

6. With **Identity Management for UNIX** highlighted, click **Details**.

7. When the **Identity Management for UNIX** dialog box opens, click to fill the **Password Synchronization** check box.

   Because Administration Components are required for Password Synchronization operation, this item is automatically selected for installation when you select Password Synchronization.

8. Click **OK**.

   The Windows Components Wizard begins installing the Identity Management for UNIX components you selected.

9. If you are prompted to locate files required for Password Synchronization installation, insert the Windows Server 2003 R2 product CD, or browse to the network location of your Windows Server 2003 R2 installation files.

10. When installation completes, restart your computer to begin working with Password Synchronization.

**To install Password Synchronization by using command-line tools**

You can use an answer file to run the Identity Management for UNIX Installation Wizard without your having to be present for interactive responses. Create the answer file, and then run it at a command prompt on the installation computer.

Have your Windows Server 2003 R2 product CD available for this command-line installation.

Perform the following steps to install Password Synchronization by using a command-line environment:

1. Create an unattended answer file in %*windir*% with the following contents.

   | **Note** |
   | --- |
   | The unattended answer file is a plain text file that Windows Server operating systems use to respond to prompts about your installation preferences. For more information about unattended installations, see Unattended Installation Fundamentals in the "Windows Server 2003 Deployment Guide." |

   **[Components]**

   **Psync=on**

   **Idmumgmt=on**

2. At a command prompt, type the following, and press **Enter**.

   **synocmgr /i:%***windir***%\inf\sysoc.inf /u:<***answerfile***.txt> /q**

**Continue with Password Synchronization Setup**

To continue setting up Password Synchronization, go on to Task 2: Install the Password Synchronization Daemon on UNIX-based Computers.

# Task 2: Install the Password Synchronization Daemon on UNIX-based Computers

Updated: June 27, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To allow synchronization of Windows passwords with UNIX hosts, you must install the Password Synchronization daemon on each UNIX host on which passwords are to be synchronized.

**In this topic**


What is the Password Synchronization Daemon?

What Operating Systems are Supported by the Password Synchronization Daemon?

To install the Password Synchronization daemon

Continue with Password Synchronization Setup


**What is the Password Synchronization Daemon?**

When Password Synchronization receives a request for a password change, it encrypts the password and sends it to all UNIX hosts that are to be synchronized with the Windows-based computer or domain. To process the password change request, the UNIX host must be running the Password Synchronization daemon, also known as the single sign-on daemon (SSOD). This daemon receives the request and changes the password on the UNIX host.

In addition, if the UNIX host is a master Network Information Service (NIS) server, the Password Synchronization daemon runs **make** to rebuild the NIS **passwd** map so it can be replicated to subordinate (slave) servers in the NIS domain. The Password Synchronization daemon performs event logging through the **syslogd** daemon running on the UNIX host.

**What Operating Systems are Supported by the Password Synchronization Daemon?**

Password Synchronization supports synchronization with UNIX computers running any of the following operating systems:

- Hewlett-Packard HP-UX version 11i
- IBM AIX version 5L 5.2

| **Important** |
| --- |
| Only Windows to UNIX password synchronization is supported on AIX. UNIX to Windows password synchronization is not supported. |

- Red Hat Linux versions 8 and 9
- Sun Solaris version 8 running on x86-based computers and Scalable Processor Architecture (SPARC)–based computers, and Solaris version 9 running on SPARC–based computers

**To install the Password Synchronization daemon**

| **Note** |
| --- |
| Because the Password Synchronization daemon is installed on UNIX-based computers, there is no Windows user interface procedure for completing this task. |

1. Copy the appropriate source binary file from \Unix\Bins on Windows Server 2003 R2 CD or DVD-ROM to /usr/bin or /usr/local/bin on the UNIX computer, and change its name to **ssod**. The name of the source binary file depends on the version of UNIX you are using.

    - If the computer is running Hewlett-Packard HP-UX, the source binary file name is **ssod.hpx**.

- If the computer is running Red Hat Linux, the source binary file name is **ssod.rhl**.
- If the computer is running Sun Microsystems Solaris, the source binary file name is **ssod.sol**.
- If the computer is running IBM AIX, the source binary file name is **ssod.aix**.

2. Using a binary file-copy method such as File Transfer Protocol (FTP) to avoid corrupting CR/LF (carriage-return/line-feed) pairs, copy **Sso.cfg** from \Unix\Bins on the UNIX Identity Management CD to /etc on the UNIX computer, and change its name to **sso.conf**.

3. Open sso.conf with a text editor.

4. If you have changed the default encryption key, edit the following line to specify the new default key. This value must match the default key specified on all domain controllers with which this computer will synchronize passwords:

   **ENCRYPT_KEY=***encryptionKey*

5. If you have changed the default port, edit the following line to specify the new port. This value must match the port number specified on all domain controllers with which this computer will synchronize passwords.

   **PORT_NUMBER=***portNumber*

6. Edit the following line to specify one domain controller in each Windows domain with which the computer is to synchronize passwords. If you have specified a nondefault port number or encryption key for the UNIX-based computer when configuring Password Synchronization on the Windows domain controllers, specify that value where indicated; otherwise, leave the value blank:

   **SYNC_HOSTS=(***domainController*[**,** *portNumber* [**,** *encryptionKey*]]**)** ...

   Each entry in the list must be enclosed by parentheses (the "(" and ")" characters) and separated from the next entry by a blank space.

7. If the computer is a Network Information Service (NIS) master server, and if you want passwords to be synchronized throughout the NIS domain, edit the following line as shown to enable NIS synchronization:

   **USE_NIS=1**

   Also, if required, edit the following line to specify the location of the NIS makefile:

   **NIS_UPDATE_PATH=***makefilePath*

8. Set the file permissions of sso.conf to read/write for the root user only, and deny access to all other users.

9. If the computer is running Linux, copy /etc/pam.d/system-auth to /etc/pam.d/ssod.

| **Important** |
| --- |
| The sso.conf file contains encryption keys and other sensitive information. For this reason, it must be accessible only by system administrators. |

| **Note** |
| --- |
| This daemon program must be installed on the computer running UNIX to enable Password Synchronization to change users' passwords on that computer. |

**Continue with Password Synchronization Setup**

To continue setting up Password Synchronization, go on to Task 3: Install the Pluggable Authentication Module on UNIX-based Computers.

# Task 3: Install the Pluggable Authentication Module on UNIX-based Computers

Applies To: Windows Server 2003 R2

Pluggable authentication modules (PAMs) allow a UNIX computer to support multiple authentication technologies. Password Synchronization uses PAMs to provide UNIX-to-Windows password synchronization. To allow passwords on Windows-based computers or domains to be changed when users change their UNIX password, the Password Synchronization PAM module (**pam_sso**) must be installed on each UNIX host where users can change their passwords.

Much like Password Synchronization running on a Windows-based computer, the Password Synchronization PAM on a UNIX computer intercepts the password change request, encrypts the password, and then transmits the request to the appropriate Windows-based computers running Password Synchronization. Like the Password Synchronization daemon, the Password Synchronization PAM performs event logging through the syslogd daemon running on the UNIX host.

This section contains instructions on installing the PAM on computers running any of the following four UNIX-based operating systems:

- AIX
- HP-UX
- Red Hat Linux
- Solaris

**To install the pluggable authentication module (PAM) on AIX**

1. Copy pam_sso.aix from \Unix\Bins on the Windows Server 2003 R2 product CD or DVD-ROM to /usr/lib/ on the UNIX computer, and change its name to pam_sso.aix.1.

2. On the UNIX computer, log on as **root**, and then enter the following commands:

   **chown root /usr/lib/pam_sso.aix.1 chmod 555 /usr/lib/pam_sso.aix.1**

3. If necessary, create the /etc/pam.conf file according to your network requirements, setting the owner to root and the base permissions to 644. For more information about creating the pam.conf file, see "Pluggable Authentication Modules" in *System Management Guides: Security Guide* in your AIX documentation.

   Sample pam.conf file

```
#
# Authentication management
#
OTHER   auth     required       /usr/lib/security/pam_aix

#
# Account management
#
OTHER   account  required       /usr/lib/security/pam_aix

#
# Session management
#
OTHER   session  required       /usr/lib/security/pam_aix
```

4. Open /etc/pam.conf with a text editor.

5. In the **Password management** section, add the following line:

   **passwd password required /usr/lib/security/pam_sso.aix.1**

   Sample pam.conf file with this line added

```
#
# Authentication management
#
OTHER   auth     required       /usr/lib/security/pam_aix

#
# Account management
#
OTHER   account  required       /usr/lib/security/pam_aix

#
# Session management
#
OTHER   session  required       /usr/lib/security/pam_aix

#
# Password management
#
passwd   password required       /usr/lib/security/pam_sso.aix.1
```

6. Open /usr/lib/security/methods.cfg with a text editor and add the following lines at the end of the file:

   **PAM:  program = /usr/lib/security/PAM**

   **PAMfiles:  options = auth=PAM,db=BUILTIN**

7. Open /etc/security/user with a text editor and add authentication information for the specific users whose passwords you want to synchronize. For example:

   user1:  admin = false  SYSTEM = PAMfiles[*] AND "compat"  registry = PAMfiles

| **Note** |
| --- |
| You can choose to change the default section of /etc/security/user to allow all users to synchronize their passwords. In this case, to restrict access to Password Synchronization, you can use the SYNC_USERS attribute in the /etc/sso.conf file to restrict access. For more information, see Using sso.conf to configure Password Synchronization on the UNIX computer.<br><br>To disable UNIX-to-Windows password synchronization, remove the entry in /etc/pam.conf that you added in step 5. |

1. If you are finished installing the PAM on AIX-based computers, go on to Task 4: Configure Password Synchronization.

**To install the pluggable authentication module (PAM) on HP-UX**

1. Copy pam_sso.hpx from \Unix\Bins on the Windows Server 2003 R2 product CD or DVD-ROM to /usr/lib/security on the UNIX computer, change its name to pam_sso.hp.1, and then set its file-mode bits to 544.

2. On the UNIX computer, open /etc/pam.conf with a text editor.

3. In the **Password management** section, locate the following line:

   **other password required /usr/lib/security/libpam_unix.1**

4. Immediately following the line located in the previous step, add the following line:

   **other password required /usr/lib/security/pam_sso.hp.1**

| **Note** |
| --- |
| To disable UNIX-to-Windows password synchronization, remove the entry in /etc/pam.conf that you added in step 4.<br><br>Before installing the pam_sso module, make sure that PAM support is properly installed and configured on the UNIX computer. The following file samples show a typical configuration. Actual contents of these files may differ, depending on your system configuration.<br><br>Sample HP-UX PAM configuration file |

```
#
# PAM configuration
#
# Authentication management
#
login    auth required  /usr/lib/security/libpam_unix.1
su       auth required  /usr/lib/security/libpam_unix.1
dtlogin  auth required  /usr/lib/security/libpam_unix.1
dtaction auth required  /usr/lib/security/libpam_unix.1
ftp      auth required  /usr/lib/security/libpam_unix.1
OTHER    auth required  /usr/lib/security/libpam_unix.1
#
# Account management
#
login    account required       /usr/lib/security/libpam_unix.1
su       account required       /usr/lib/security/libpam_unix.1
dtlogin  account required       /usr/lib/security/libpam_unix.1
dtaction account required       /usr/lib/security/libpam_unix.1
ftp      account required       /usr/lib/security/libpam_unix.1
#
OTHER    account required       /usr/lib/security/libpam_unix.1
#
# Session management
#
login    session required       /usr/lib/security/libpam_unix.1
dtlogin  session required       /usr/lib/security/libpam_unix.1
dtaction session required       /usr/lib/security/libpam_unix.1
OTHER    session required       /usr/lib/security/libpam_unix.1
#
# Password management
#
login    password required      /usr/lib/security/libpam_unix.1
dtlogin  password required      /usr/lib/security/libpam_unix.1
dtaction password required      /usr/lib/security/libpam_unix.1
other    password required      /usr/lib/security/libpam_unix.1
other    password required      /usr/lib/security/pam_sso.hp.1
```

The file-mode bits for pam_sso.hp.1 must be set to 544 (o:r-x,g:r--,w:r--) or it will not function properly.

1. If you are finished installing the PAM on HP-UX-based computers, go on to Task 4: Configure Password Synchronization.

**To install the pluggable authentication module (PAM) on Linux**

1. Copy pam_sso.rhl from \Unix\Bins on the Windows Server 2003 R2 product CD or DVD-ROM to /lib/security on the

UNIX computer, and change its name to pam_sso.so.1.

2. On the UNIX computer, copy /etc/pam.d/system-auth to /etc/pam.d/ssod.

3. Open /etc/pam.d/system-auth with a text editor, and locate the following line:

   **passwordrequired/lib/security/pam_cracklib.soretry=3**

4. After the line in the previous step, add the following line:

   **password required /lib/security/pam_sso.so.1**

5. Locate and delete the following line:

   **passwordrequired/lib/security/pam_deny.so**

6. Save the modified file.

**Note**

These instructions apply to the typical Linux configuration. If you have configured PAM support differently, you might have to adjust these instructions to your specific configuration.

To disable UNIX-to-Windows password synchronization, remove the entry in /etc/pam.d/system-auth that you added in step 4.

Before installing the pam_sso module, make sure that PAM support is properly installed and configured on the UNIX computer. The following file samples show a typical configuration. Actual contents of these files may differ, depending on your system configuration.

Sample Linux PAM configuration file:

```
/etc/pam.d/passwd

#%PAM-1.0
auth       required     /lib/security/pam_stack.so service=system-auth
account    required     /lib/security/pam_stack.so service=system-auth
password   required     /lib/security/pam_stack.so service=system-auth


/etc/pam.d/ssod

#%PAM-1.0
# This file is auto-generated.
# User changes will be destroyed the next time authconfig is run.
auth        required      /lib/security/pam_env.so
auth        sufficient    /lib/security/pam_unix.so likeauth nullok
auth        required      /lib/security/pam_deny.so

account     required      /lib/security/pam_unix.so

password    required      /lib/security/pam_cracklib.so retry=3 type=
password    sufficient    /lib/security/pam_unix.so nullok use_authtok shadow
password    required      /lib/security/pam_deny.so

session     required      /lib/security/pam_limits.so
session     required      /lib/security/pam_unix.so


/etc/pam.d/system-auth

#%PAM-1.0
# This file is auto-generated.
# User changes will be destroyed the next time authconfig is run.
```

```
auth        required      /lib/security/pam_env.so
auth        sufficient    /lib/security/pam_unix.so likeauth nullok
auth        required      /lib/security/pam_deny.so

account     required      /lib/security/pam_unix.so

password    required      /lib/security/pam_cracklib.so retry=3 type=
password    required      /lib/security/pam_sso.so.1
password    sufficient    /lib/security/pam_unix.so nullok use_authtok shadow
password    required      /lib/security/pam_deny.so

session     required      /lib/security/pam_limits.so
session     required      /lib/security/pam_unix.so
```

1. If you are finished installing the PAM on Linux-based computers, go on to Task 4: Configure Password Synchronization.

**To install the pluggable authentication module (PAM) on Solaris**

1. Copy pam_sso.sol from the \Unix\Bins folder on the Windows Server 2003 R2 product CD or DVD-ROM to the /usr/lib/security directory on the UNIX computer, and change its name to pam_sso.so.1.

2. On the UNIX computer, open /etc/pam.conf with a text editor.

3. In the **Password management** section, locate the following line:

   **other password required /usr/lib/security/$ISA/pam_unix.so.1**

4. Immediately following the line located in the previous step, add the following line:

   **other password required /usr/lib/security/$ISA/pam_sso.so.1**

| **Note** |
| --- |
| To disable UNIX-to-Windows password synchronization, remove the entry in /etc/pam.conf that you added in step 4.<br><br>Before installing the pam_sso module, make sure that PAM support is properly installed and configured on the UNIX computer. The following file samples show a typical configuration. Actual contents of these files may differ, depending on your system configuration.<br><br>Sample Solaris PAM configuration file |

```
#
#ident  "@(#)pam.conf   1.14    99/09/16 SMI"
#
# Copyright (c) 1996-1999, Sun Microsystems, Inc.
# All Rights Reserved.
#
# PAM configuration
#
# Authentication management
#
login   auth required   /usr/lib/security/$ISA/pam_unix.so.1
login   auth required   /usr/lib/security/$ISA/pam_dial_auth.so.1
#
rlogin  auth sufficient /usr/lib/security/$ISA/pam_rhosts_auth.so.1
rlogin  auth required   /usr/lib/security/$ISA/pam_unix.so.1
#
dtlogin auth required   /usr/lib/security/$ISA/pam_unix.so.1
#
rsh     auth required   /usr/lib/security/$ISA/pam_rhosts_auth.so.1
other   auth required   /usr/lib/security/$ISA/pam_unix.so.1
#
# Account management
```

```
#
login   account requisite       /usr/lib/security/$ISA/pam_roles.so.1
login   account required        /usr/lib/security/$ISA/pam_unix.so.1
#
dtlogin account requisite       /usr/lib/security/$ISA/pam_roles.so.1
dtlogin account required        /usr/lib/security/$ISA/pam_unix.so.1
#
other   account requisite       /usr/lib/security/$ISA/pam_roles.so.1
other   account required        /usr/lib/security/$ISA/pam_unix.so.1
#
# Session management
#
other   session required        /usr/lib/security/$ISA/pam_unix.so.1
#
# Password management
#

other   password required       /usr/lib/security/$ISA/pam_unix.so.1
other  password required        /usr/lib/security/$ISA/pam_sso.so.1
dtsession auth required /usr/lib/security/$ISA/pam_unix.so.1

#
# Support for Kerberos V5 authentication (uncomment to use Kerberos)
#
#rlogin auth optional   /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
#login  auth optional   /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
#dtlogin        auth optional   /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
#other  auth optional   /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
#dtlogin        account optional /usr/lib/security/$ISA/pam_krb5.so.1
#other  account optional /usr/lib/security/$ISA/pam_krb5.so.1
#other  session optional /usr/lib/security/$ISA/pam_krb5.so.1
#other  password optional /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
```

1. If you are finished installing the PAM on Solaris-based computers, go on to Task 4: Configure Password Synchronization.

**See Also**
**Other Resources**
Using sso.conf to configure Password Synchronization on the UNIX computer

# Task 4: Configure Password Synchronization

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you have installed Password Synchronization on the appropriate Windows–based computer, you can administer it by using the Identity Management for UNIX management console. You can also administer Password Synchronization by using the command-line utility **psadmin**. You must be a member of the Administrators group on the computer you want to administer.

The procedures in this topic show how you can use Password Synchronization administration to:

- Select the server to be administered.
- Set the default settings that apply to the entire configuration of UNIX-based computers defined for the Windows–based computer or domain. The settings determine what the event log displays, the maximum number of times to resend the failed password update, and the length of time the service waits before resending a password update that has failed.
- Add or remove a UNIX-based computer from the list of computers designated to receive password updates.
- Create or modify the configuration for the UNIX-based computer, including the custom settings (the default settings applied to that single computer) and the encryption settings for secure communication. The order in which the names of the UNIX hosts appear in the list determines their order in the registry and the order in which they are processed for password synchronization.

**In this Section**


- Set Default Synchronization
- Set the Default Encryption Key
- Set the Default Port
- Add or Remove Computers for Synchronization
- Set Computer-specific Synchronization Properties
- Configure Password Synchronization Audit Logging
- Continue with Password Synchronization Setup


**Set Default Synchronization**

Set default synchronization by using the Windows interface

Set default synchronization by using the command line

**Set default synchronization by using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**, and then click the **Default** tab.
4. To allow password synchronization from UNIX-based computers to Windows-based computers, click **Synchronize password changes from computers that run UNIX to computers that run Windows**.
5. To allow password synchronization from Windows-based computers to UNIX-based computers, click **Synchronize**

**password changes from computers that run Windows to computers that run UNIX**.

6. To save the new settings, click **Apply**.

| Note |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.<br><br>These settings affect the default synchronization for UNIX hosts when they are added for synchronization. It does not affect computers that have already been added for synchronization. |

**Set default synchronization by using the command line**

- To set the default direction of synchronization, at a command prompt, type

  **psadmin** *computer_name* **-enable** [**WintoUnix** | **UnixToWin** | **BothDir**]

  in which *computer_name* represents the name of the computer for which you want to configure the direction of password synchronization.

The following table describes the values available for the common option **-enable**.

| Value | Description |
| --- | --- |
| **WintoUnix** | Synchronizes password changes from computers that run Windows operating systems to computers that run UNIX operating systems. |
| **UnixToWin** | Synchronizes password changes from computers that run UNIX operating systems to computers that run Windows operating systems. |
| **BothDir** | Enables two-way password synchronization. |

**Set the Default Encryption Key**

Before you complete this procedure, have available the encryption key you want to use to encrypt passwords between UNIX-based and Windows-based computers on your network.

For more information, see Encryption key requirements in the Password Synchronization Help.

Set default encryption by using the Windows interface

Set default encryption by using the command line

**Set default encryption by using the Windows interface**

1. Open Identity Management for UNIX.

2. If necessary, connect to the computer you want to manage.

3. Click **Password Synchronization**, and then click the **Default** tab.

4. In the **Encryption key** text box, type the key you want to use, or, to have the program produce a key for you, click **New Key**.

5. To save the new settings, click **Apply**.

| Note |
| --- |

To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.

This setting affects the default encryption key for UNIX hosts when they are added for synchronization as well as the port used for UNIX-to-Windows synchronization. If you change this setting, you must edit the /etc/sso.conf file to specify the same encryption key on UNIX hosts that are configured for UNIX-to-Windows password synchronization with this computer. For more information, see Using sso.conf to configure Password Synchronization on the UNIX computer.

**Set default encryption by using the command line**

- To set the default encryption method, at a command prompt, type

  **psadmin -comp** *name* **-key** *keyvalue*

  in which *name* represents the name of the computer for which you want to configure the direction of password synchronization, and *keyvalue* represents the encryption key you want to use.

**Note**

To have Password Synchronization assign a key for you, enter **random** as the key value.

**Set the Default Port**

This setting affects the default port number for UNIX hosts when they are added for synchronization, as well as the port used for UNIX-to-Windows synchronization. If you change this setting, you must edit the /etc/sso.conf file to specify the same port on UNIX hosts that are configured for UNIX-to-Windows password synchronization with this computer. For more information, see Using sso.conf to configure Password Synchronization on the UNIX computer.

For maximum security, use a port number other than the default (6677).

Set the default port by using the Windows interface

Set the default port by using the command line

**Set the default port by using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**, and then click the **Default** tab.
4. To use a port other than 6677, in the **Port number** box, type the port number you want.
5. To save the new settings, click **Apply**.

**Note**

To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.

This setting affects the default port number for UNIX hosts when they are added for synchronization, as well as the port used for UNIX-to-Windows synchronization. If you change this setting, you must edit the /etc/sso.conf file to specify the same port on UNIX hosts that are configured for UNIX-to-Windows password synchronization with this computer. For more information, see Using sso.conf to configure Password Synchronization on the UNIX computer.

For maximum security, use a port number other than the default (6677).

**Set the default port by using the command line**

- To set the default port, at a command prompt, type

  **psadmin -comp** *name* **-port** *port_number*

  in which *name* represents the name of the computer for which you want to change the port number, and *port_number* represents the port number you want to use.

**Add or Remove Computers for Synchronization**

Add or remove computers for synchronization by using the Windows interface

Add or remove computers for synchronization by using the command line

**Add or remove computers for synchronization by using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**.
4. Click the **Advanced** tab, and then do one of the following:
5. To add a computer to the list of current computers, in **Computer name**, type the name of the UNIX-based computer you want to add, and then click **Add**.
6. To remove a computer, in the **Current computers** list, click the UNIX-based computer you want to remove, and then click **Remove**.
7. To save the new settings, click **Apply**.

| **Note** |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.<br><br>In addition to adding a UNIX-based computer to the list, if you want to change the user's password on the UNIX computer when the corresponding Windows user's password is changed, you must install the Password Synchronization single sign-on daemon (SSOD) on the UNIX-based computer. For more information about installing the SSOD, see Task 2: Install the Password Synchronization Daemon on UNIX-based Computers.<br><br>If you want to change the Windows user's password when the corresponding UNIX-based computer user's password is changed, you must install the pluggable authentication module (PAM) on the UNIX-based computer. For more information about installing the PAM, see Task 3: Install the Pluggable Authentication Module on UNIX-based Computers. |

**Add or remove computers for synchronization by using the command line**

- To add a computer for synchronization, at a command prompt, type

  **psadmin add** *computer_name*

  in which *computer_name* represents the name of the computer you want to participate in password synchronization.

- To remove a computer from the password synchronization process, type

  **psadmin delete** *computer_name*

  in which *computer_name* represents the name of the computer you want to remove from the password synchronization process.

| Note |
| --- |
| **psadmin list** displays a list of all computers participating in password synchronization. |

**Set Computer-specific Synchronization Properties**

Set computer-specific synchronization properties by using the Windows interface

Set computer-specific synchronization properties by using the command line

**Set computer-specific synchronization properties by using the Windows interface**

1. Open Identity Management for UNIX. If necessary, connect to the computer you want to manage.
2. Click **Password Synchronization**, and then click the **Advanced** tab.
3. In the **Current computers** list, click the one for which you want to set properties, and then click **Configure**.
4. Set the properties you want to apply to the selected computer.
5. To save the new settings, click **Apply**.

| Note |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.<br><br>Before setting computer-specific synchronization properties in Password Synchronization, edit the sso.conf file on the UNIX-based computer to specify the same settings. For information about assigning computer-specific settings in the sso.conf file, see Using sso.conf to configure Password Synchronization on the UNIX computer. |

**Set computer-specific synchronization properties by using the command line**

- To modify computer-specific synchronization properties by using the command line, at a command prompt, type

  **psadmin -comp** *computer_name*[*common_option*] [*common_option_value*]

  in which *computer_name* represents the name of the computer for which you want to modify password synchronization properties, *common_option* represents one of the configurable parameters in the following table, and *common_option_value* represents one of the acceptable values for the parameter, as found in the following table.

| Option | Description |
| --- | --- |
| **-comp** *name* | Computer to which configuration options are applied. If **-comp** is unspecified, Password Synchronization modifies the default configuration settings. If **-comp** is the only option specified, then Password Synchronization configuration of the specified computer is displayed. |
| **-enable** *direction* | Specifies the direction of password synchronization. The variable *direction* can contain one of the following values:<br><br>**WintoUnix**: Synchronize password changes from computers that run Windows operating systems to computers that run UNIX operating systems.<br><br>**UnixToWin**: Synchronize password changes from computers that run UNIX operating systems to computers that run Windows operating systems.<br><br>**BothDir**: Enable two-way password synchronization. |
| **-key** *keyvalue* | Sets the encryption and decryption key for the computer specified by **-comp**. If *keyvalue* is random, Password Synchronization uses a random encryption key. |

| | |
|---|---|
| **-port** *number* | Sets the port number for the specified computer. |
| **-?** | Displays **psadmin** usage and arguments. |

**Configure Password Synchronization Audit Logging**

Set auditing options by using the Windows interface

Set auditing options by using the command line

**Set auditing options by using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**, and then click the **Default** tab.
4. To have information events as well as warnings and errors logged to Event Viewer, click **Enable extensive logging**.
5. To save the new settings, click **Apply**.

| **Note** |
|---|
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.<br><br>When the log file reaches the limit you set, logging stops and Event Viewer displays a note that the log file is full. |

**Set auditing options by using the command line**

- To enable or disable logging by using the command line, at a command prompt, type

  **psadmin -log** [**yes** | **no**]

  in which **yes** enables logging, and **no** disables logging.

| **Note** |
|---|
| The **-log** option is a global setting; it can be used only when **-comp** is not used. |

**Continue with Password Synchronization Setup**

To continue setting up Password Synchronization, go on to Task 5: Start Password Synchronization.

**See Also**

**Other Resources**

Encryption key requirements
Connect to a computer you want to manage
Implementing Password Synchronization
Understanding Password Synchronization
Password encryption
psadmin

# Task 5: Start Password Synchronization

Applies To: Windows Server 2003 R2

After installing and configuring Password Synchronization, you can test your configuration by starting the Password Synchronization service.

**In this Section**

Start the Password Synchronization Daemon

Stop the Password Synchronization Daemon

**Start the Password Synchronization Daemon**

This procedure is carried out in the command-line environment.

- To start Password Synchronization, run the Password Synchronization daemon (SSOD) by entering the following at a command prompt:

  **ssod**

**Stop the Password Synchronization Daemon**

This procedure is carried out in the command-line environment.

- To stop Password Synchronization, stop the SSOD process by entering the following at a command prompt:

  **kill –term** *pid*

  where *pid* represents the process identifier of the Password Synchronization SSOD.

**See Also**

**Concepts**

Task 2: Install the Password Synchronization Daemon on UNIX-based Computers

**Other Resources**

Troubleshooting Password Synchronization

# Services for Network File System in Microsoft Windows Server 2003 R2

Applies To: Windows Server 2003 R2

Services for Network File System (NFS) is a new feature of Microsoft® Windows Server® 2003 R2 that builds on the authentication framework introduced in Windows® Services for UNIX. It also adds additional functionality and authentication mechanisms that utilize Active Directory® Domain Services and the new Identity Management for UNIX to identify which user is requesting access to a resource and to authenticate that request to the appropriate authority. This white paper covers the various mechanisms that are involved in authentication and focuses on the new mechanisms introduced in Windows Server 2003 R2.

**Contents**

Introduction

Concepts

Configuration

Mapping UNIX Users to Windows Users

Conclusion

For More Information

# Introduction

Applies To: Windows Server 2003 R2

Microsoft® Windows Server® 2003 R2 extends the Microsoft commitment to full interoperability between Microsoft Windows® and UNIX by integrating the key components necessary for interoperability directly into the core Windows Server product.

| **Note** |
| --- |
| Throughout this document, we use the term "UNIX" to refer to resources and authentication of UNIX and Linux users and computers. |

The interoperability components in Windows Server 2003 R2 are functionally similar to their equivalent modules in Windows Services for UNIX 3.5. However, there are many important enhancements and updates since the release of Windows Services for UNIX 3.5.

Windows Services for UNIX 3.5 is available for use on 32-bit versions of Microsoft Windows 2000, Microsoft Windows XP Professional Edition, and Microsoft Windows Server 2003. The UNIX interoperability components of Windows Server 2003 R2 are available and supported on all versions of the product.

| **Note** |
| --- |
| The Active Directory® Domain Services schema is updated on your domain controller (DC) during the Windows Server 2003 R2 installation. This schema update includes the necessary extensions for Network File System (NFS) authentication, and an additional schema update is not required to enable the components described in this white paper. |

**In this section**

What Is New in Windows Server 2003 R2?

NFS Components in Windows Server 2003 R2 and Past Releases

# What Is New in Windows Server 2003 R2?

Applies To: Windows Server 2003 R2

Windows Server 2003 R2 includes a suite of UNIX interoperability components as part of the core operating system. These components include:

- Identity Management for UNIX
    - Password synchronization
    - Server for NIS
- Services for Network File System (NFS)
    - Client for NFS
    - Server for NFS
    - User Name Mapping service
    - Server for NFS Authentication
    - RPC port mapper
- Subsystem for UNIX-based applications (SUA)

In addition to these components, the administration tools have been updated and improved. For example, the graphical administration tools include the Services for NFS console (shown in Figure 1) and the Microsoft Identity Management for UNIX console (shown in Figure 2).



**Figure 1. The Microsoft Services for NFS console**




**Figure 2. The Microsoft Identity Management for UNIX console**

Another important feature in the UNIX interoperability components is support for 64-bit operating systems. Windows Services for UNIX 3.5 does not support 64-bit Windows, but Windows Server 2003 R2 fully supports 64-bit interoperability.

| **Note** |
| --- |
| Microsoft Windows XP Professional x64 Edition has no support for any UNIX interoperability components, nor does it allow you to install Windows Services for UNIX 3.5. |

# NFS Components in Windows Server 2003 R2 and Past Releases

Applies To: Windows Server 2003 R2

A number of components are either involved in or dependent on Network File System (NFS) authentication. These include:

- **User Name Mapping service**. This component maps UNIX users to Windows users and vice versa. Even when a user references the same name on both systems, the programs do not recognize them as the same user. User Name Mapping service tells the components in Services for NFS that Windows user jdoe is the same as UNIX user johnd.

- **Server for NFS Authentication**. This component is not a server at all, but rather an authentication component used by Server for NFS. If you use local user accounts, install Server for NFS Authentication on all Windows NFS servers. If you use domain user accounts, raise the domain functional level to Windows Server 2003 or install Server for NFS Authentication on all domain controllers.

- **Client for NFS**. This is the Windows NFS client component of Services for NFS. Client for NFS allows the computer on which it is installed to access and use NFS resources anywhere on the network.

- **Server for NFS**. This is the Windows NFS server component of Services for NFS. Server for NFS allows the machine on which it is installed to provide file system resources to NFS clients anywhere on the network.

Each of these components included in Windows Server 2003 R2 have equivalent components in the earlier Windows Services for UNIX releases. However, two components from Windows Services for UNIX 3.5 are not included in Windows Server 2003 R2: Server for Personal Computer Network File System (PCNFS) and Gateway for NFS.

Windows Services for UNIX 1.0 uses the PCNFS protocol for authentication to NFS servers, and some third-party legacy NFS applications have also used PCNFS. However, since Windows Services for UNIX 2.0, the authentication mechanism has not used PCNFS directly, though the Server for PCNFS was included to support those legacy clients.

The Gateway for NFS component attempted to provide a single server that connected to NFS servers and shared the resulting connection to downstream Windows clients by using the standard Windows networking protocol: server message blocks (SMB). This component, however, was a potential bottleneck in the connection and introduced substantial network overhead. With the availability of Client for NFS, the need for a gateway product has been substantially reduced.

# Concepts

Applies To: Windows Server 2003 R2

Windows and UNIX use very different authentication mechanisms. The Microsoft Identity Management for UNIX components of Windows Server 2003 R2 build on the identity management components originally included in Windows Services for UNIX (User Name Mapping service and Server for NIS). Windows Server 2003 R2 adds direct support for Active Directory Domain Services Lookup while continuing to support the legacy Windows Services for UNIX mechanisms.

**In this section**

Windows and UNIX User and Group Identities

Integrated Identity Management in Active Directory Domain Services

Identity Matching in User Name Mapping Service

# Windows and UNIX User and Group Identities

Applies To: Windows Server 2003 R2

When a UNIX user requests access to a Windows resource, or when a Windows user requests access to a UNIX resource, the computers need to map the UNIX user to the proper Windows user or vice versa. Windows Services for UNIX 3.5 uses User Name Mapping service (still used by Windows Server 2003 R2). However, Windows Server 2003 R2 adds direct lookup of UNIX User ID (UID) and Group ID (GID) from Active Directory Domain Services, with no User Name Mapping service required.

When a UNIX user requests access to a file or folder ("object") stored on a Windows Network File System (NFS) shared network resource, the first step is for Windows to map the UNIX user to the Windows user. This is accomplished through Active Directory Domain Lookup or User Name Mapping service.

The request for access to the Windows object is authenticated against a Windows domain controller by using the security identifier (SID) of the mapped Windows user. (In the special case of using a local Windows account with a local User Name Mapping service, the request is authenticated against the local Windows computer.)

If the mapped Windows user has the appropriate access permissions, the UNIX user who initiated the request is granted access to the object. Conversely, if the Windows user does not have sufficient permission for the UNIX user to access the object, access is denied to the UNIX user.

Access is not mapped to a Windows user and the permissions of an anonymous user are used if:

- The UNIX user does not have a mapped Windows account.
- The User Name Mapping service can't be reached.
- Active Directory Domain Services does not include the UNIX user's identity data.

When a Windows user tries to access an object stored on a UNIX computer, the same process is used but in reverse.

# Integrated Identity Management in Active Directory Domain Services

Applies To: Windows Server 2003 R2

As part of the product installation, the Windows Server 2003 R2 schema is extended to support direct lookup of UNIX identities in Active Directory Domain Services. You have the choice of running Server for NIS or managing the schema attributes directly.

**UNIX Attributes tab**
Server for NIS adds a **UNIX Attributes** tab in the Active Directory Domain Services Users and Computers console. This addition:

- Provides you with direct graphical user interface (GUI) access to the schema extensions.
- Allows you to directly set attributes for each user (as shown in Figure 3).

The above requires that Server for NIS be installed on a Windows Server 2003 R2 domain controller.




**Figure 3. UNIX Attributes tab for user Charlie Russel**

If access is requested to a Windows resource, the access requires the permissions of a Windows user. Active Directory Domain Services looks up the UID/GID of the UNIX user requesting access and maps the request to the Windows user to set the appropriate permissions.

If access is requested from a Windows user to a UNIX resource, the Windows user's UID/GID pair is provided to the UNIX authentication mechanism to set appropriate permissions.

# Identity Matching in User Name Mapping Service

Applies To: Windows Server 2003 R2

The User Name Mapping service was the core authentication component of Windows Services for UNIX, and it continues to be supported in Services for Network File System (NFS). However, it is important to note that this is the final release of the User Name Mapping service, and it will not be supported in future releases of Services for NFS.

Although User Name Mapping service is at the core of the authentication processes for Windows Services for UNIX 3.5, it does not perform the authentication. The authentication is performed by the:

- Windows domain controller or the Local Security Authority Subsystem Service (LSASS) if access permissions are requested by a Windows user (when Server for NFS shares files with UNIX users).

- UNIX authentication mechanism if access permissions are requested by a UNIX user (when Windows users access UNIX-based shared network resources through Client for NFS).

In other words, if a user requests access to a shared network resource on a Windows computer, that computer uses a Windows authentication mechanism. If a user requests access to a shared network resource on a UNIX computer, that computer uses UNIX authentication.

One other important authentication occurs when a request comes to the User Name Mapping service. The .maphosts file resides in the %SFUDIR%\Mapper subdirectory of the Windows Services for UNIX 3.5 installation directory or the %WINDIR%\msnfs directory in Windows Server 2003 R2. The .maphosts file is checked to see whether the requesting computer is permitted to access the User Name Mapping service. If it is not, the request is denied.

Several steps are involved in any authentication, and any of the steps could cause an authentication failure. It is important to understand the process to troubleshoot it effectively.

# Configuration

Applies To: Windows Server 2003 R2

The Configuration section of this white paper covers the basic installation and configuration of the Windows Server 2003 R2 components used by Services for Network File System (NFS) and Server for NFS Authentication.

**In this section**

Installing Services for Network File System

Installing Identity Management for UNIX

Extending the Active Directory Domain Services Schema

Importing UNIX Identities into Active Directory by Using IDMU

# Installing Services for Network File System

Applies To: Windows Server 2003 R2

**Installing UNIX Interoperability Components**

When you upgrade from Windows Server 2003 Service Pack 1 (SP1) to Windows Server 2003 R2, the installation program does not install the UNIX interoperability components on your server. The installation only adds to the list of components that can be installed. Follow these steps to perform the installation:

1. Open **Add/Remove Programs**, and then click the **Add/Remove Windows Components** button to start the Windows Components Wizard (shown in Figure 4).




**Figure 4. The Windows Components Wizard**

2. Click **Other Network File and Print Services**, and then click **Details**.

3. Click **Highlight Services for NFS**, and then click **Details** to open the Services for NFS dialog box (shown in Figure 5).




**Figure 5. The Windows Components Wizard, Microsoft Services for NFS page**

4. Select the NFS-related services you want, and follow the wizard prompts to finish.

| **Note** |
| --- |
| Installing Services for NFS may require that you restart the computer, so plan your installation accordingly. |

# Installing Identity Management for UNIX

Applies To: Windows Server 2003 R2

Server for NIS is the component of Microsoft Identity Management for UNIX (IDMU) that is of primary importance to Network File System (NFS) authentication, and it can be installed only on a Windows Server 2003 R2 domain controller. Server for NIS cannot be installed on a member server. To install Server for NIS, follow these steps:

1. Open **Add/Remove Programs**, and then click the **Add/Remove Windows Components** button to start the Windows Components Wizard (as shown in Figure 4 in the Installing Services for NFS section of this white paper).

2. Click **Active Directory Services**, and then click **Details**.

3. Click **Highlight Identity Management for UNIX**, and then click **Details** to open the **Identity Management for UNIX** dialog box (shown in Figure 6).

4. Click **Administration Components and Server for NIS** and follow the prompts to install the NIS components of Windows Server 2003 R2. You will need to have the Windows Server 2003 R2 installation media available.

| **Note** |
| --- |
| Installing IDMU may require that you restart the computer, so plan your installation accordingly. |

| **Note** |
| --- |
| If you are using User Name Mapping service, you do not need to install Server for NIS, nor do you need to populate Active Directory Domain Services with UNIX identity information. |




**Figure 6. Microsoft Identity Management for UNIX dialog box**

# Extending the Active Directory Domain Services Schema

Applies To: Windows Server 2003 R2

The typical Windows Server 2003 R2 installation on a domain controller provides the necessary extension of the Active Directory Domain Services schema. No additional extension is required to authenticate Network File System users. However, unless Server for NIS is installed, the Active Directory Domain Services tabs for editing UNIX attributes are not visible. You can still assign UNIX attributes to users programmatically or by using Active Directory Security Interfaces (ADSI) Edit.

# Importing UNIX Identities into Active Directory by Using IDMU

Applies To: Windows Server 2003 R2

If your organization has an existing Network Information Service (NIS) domain environment, you can import your NIS users into Active Directory Domain Services by using the NIS Data Migration Wizard. The NIS Data Migration Wizard can be started from the Microsoft Identity Management for UNIX (IDMU) console or by using the **nis2ad** utility.

**For more information**

For information about how to use the NIS Data Migration Wizard, see:

- The IDMU Help files
- Microsoft TechNet article about migrating to Active Directory Domain Services using the NIS Data Migration Wizard:

  http://go.microsoft.com/fwlink/?LinkId=75009

For information about the **nis2ad** utility, see:

- Microsoft Support "How to" article about migrating existing NIS Maps to Server for NIS: http://go.microsoft.com/fwlink/?LinkId=75007
- TechNet article about **nis2ad**: http://go.microsoft.com/fwlink/?LinkId=75008

# Mapping UNIX Users to Windows Users

Applies To: Windows Server 2003 R2

Services for Network File System (NFS) supports:

- Windows Services for UNIX 3.5 User Name Mapping service
- Direct lookup of UNIX UID/GID pairs for Windows users (through Active Directory Domain Services)

Note that this is the final release that will include the User Name Mapping service as part of Services for NFS. Customers should migrate to support direct lookup of UNIX identities in Active Directory Domain Services.

**In this section**

Configuring Services for Network File System to Use Active Directory Domain Lookup

Configuring Services for Network File System to Use User Name Mapping Service

Open Ports on Firewalls

Creating a Network File System Shared Network Resource

# Configuring Services for Network File System to Use Active Directory Domain Lookup

Applies To: Windows Server 2003 R2

The preferred method to map Windows users to UNIX users is to use Active Directory Domain Lookup. Both Client for NFS and Server for NFS use one of two lookup methods: Active Directory Domain Lookup or User Name Mapping service.

**Configuring Services for NFS to use Active Directory Domain Lookup**
To configure Services for NFS to use Active Directory Domain Lookup, follow these steps:

1. Open the Services for NFS console.
2. In the left pane, at the top of the tree, right-click **Services for NFS**.
3. On the **Action** menu, click **Properties** to open the dialog box (shown in Figure 7).
4. Select the **Active Directory Domain Lookup** check box, and then enter the Active Directory Domain Services domain to which to connect.
5. Click **OK**.

| **Note** |
| --- |
| Configuring Services for NFS to use Active Directory Domain Lookup may require that you restart your computer or Services for NFS before the change is enabled. |




**Figure 7. Services for NFS Properties**

# Configuring Services for Network File System to Use User Name Mapping Service

Applies To: Windows Server 2003 R2

User Name Mapping service can be used instead of Active Directory Domain Lookup to map UNIX users to Windows users and Windows users to UNIX users. Any version of User Name Mapping service from Windows Services for UNIX 3.0 and later can be used by any version of Windows Services for UNIX or the Services for NFS from Windows Server 2003 R2. Windows Server 2003 R2 includes a version of the User Name Mapping service.

**Configuring Services for NFS to use User Name Mapping service**
To configure Services for NFS to use User Name Mapping service, follow these steps:

1. Open the Services for NFS console.
2. In the left pane, at the top of the tree, right-click **Services for NFS**.
3. From the **Action** menu, click **Properties** to open the dialog box shown in Figure 7 in the Configuring Services for Network File System to Use Active Directory Domain Lookup section of this white paper.
4. Enter the name or IP address of the User Name Mapping service.
5. Click **OK**.

| **Note** |
| --- |
| Configuring Services for NFS to use User Name Mapping service may require that you restart your computer or Services for NFS before the change is enabled. |

**Domain accounts vs. local accounts**
If you are running in a domain environment, you should generally use only domain accounts for your mapping. Local accounts are valid only on the single, local computer on which they are located, so you would need to use multiple User Name Mapping services (one for each local computer) if you intended to access resources on more than one Windows server. Obviously, if you are using Active Directory Domain Lookup, you are required to you use domain accounts.

**Simple maps**
Simple maps are maps between accounts with the exact same name in UNIX and Windows. Keep in mind that although the accounts have the same name and the exact same spelling, they are not the same account. So until you actually create the maps, the accounts will not line up. For the purposes of User Name Mapping service, the account name in Windows is not case sensitive, so Charlie will align with the UNIX account Charlie, charlie, or CHARLIE.

**Advanced maps**
Advanced maps are used to align accounts that have different names. For example, the Windows user account for Jane H. Doe might be jdoe, but the UNIX account could be jhd. Create an advanced map to align the two accounts. You can also use advanced maps to map users from different Windows domains and explicitly map accounts that would generally be mapped by simple maps.

**Authorized hosts**
To prevent unauthorized access to the User Name Mapping service, the file .maphosts is used to explicitly specify which computers may access the server. This file is installed when the User Name Mapping service is installed. In Windows Services for UNIX 3.5, the file resides in the **%SFUDIR%\Mapper** directory, and with Windows Server 2003 R2, the file resides in **%WINDIR%\msnfs**.

By default, the only computer that can access the User Name Mapping service is the local computer. You need to edit the .maphosts file with a pure ASCII text editor, such as Notepad or gvim, and add the host names or IP addresses of all the hosts that are permitted to access User Name Mapping service. Remember to include any Windows computer on which you have installed Client for NFS or Server for NFS.

During initial setup and configuration of the Server for NFS Authentication, it is a good idea to configure .maphosts to allow all computers to connect to User Name Mapping service. This can be changed after authentication is working correctly. To enable

all computers to access User Name Mapping service, type a plus sign (+) on an empty line at the bottom of the file. (For more information about the syntax of entries in .maphosts, see the comments in your .maposts file.)

**Configure domain authorization**
In addition to configuring User Name Mapping service, Server for NFS requires one additional step before UNIX users can be properly authenticated to access shared network resources on Windows servers if any of the Windows domain controllers (DCs) are not running Windows Server 2003 or if your domain is running at less than Windows Server 2003 domain functional level.

The Server for NFS Authentication must be installed on all Windows 2000 Servers and pre–Windows Server 2003 R2 DCs that has mapped users (available from the Microsoft Download Center: http://go.microsoft.com/fwlink/?linkid=53938).

**Verifying installation**
The easiest way to verify that the Server for NFS Authentication installation is complete is to check for the presence of the file %WINDIR%\system32\nfssa.dll. Server for NFS Authentication does not install any new service on a computer, and its presence is not easily discovered on the **Add/Remove Programs** list.

**Installing only Server for NFS Authentication**
In Windows Services for UNIX 3.0 and 3.5, you can easily install only Server for NFS Authentication by using the command line. This allows for easy scripting across multiple DCs. The command line to silently install Server for NFS Authentication is:

```
msiexec /I D:\<path>\sfusetup.msi /q addlocal="NFSServerAuth"
```

`D:\<path>` is the location of the Windows Services for UNIX 3.5 installation file.

Note that the `addlocal` parameter:

- Is case sensitive.
- Will remove any other components of Windows Services for UNIX 3.0 except Server for NFS Authentication.

If you are installing multiple components from the command line, you must use a comma-separated list within the quotes of the `addlocal` parameter. For example, to install both Server for NFS and Server for NFS Authentication, use:

```
msiexec /I D:\sfusetup.msi /q addlocal="NFSServer,NFSServerAuth"
```

| **Note** |
| --- |
| If you are using a version of Windows Services for UNIX that is not the free 3.5 download, you need to add `pidkey="key"` to this command line, where `"key"` is the 25-character product key without spaces or hyphens. |

In Windows Server 2003 R2, you can install the Server for NFS Authentication component from CD2. The .inf file is: \CMPNENTS\R2\PACKAGES\NEWBINS\<ARCHITECTURE>\NFSAUTH.INF.

- <ARCHITECTURE> is replaced by the appropriate folder for your server architecture.
- The actual DLL that is installed is: \CMPNENTS\R2\NFSSA.DL_.

| **Note** |
| --- |
| For information about how to add Windows components by using the Sysocmgr.exe utility, see the Microsoft Support article a http://go.microsoft.com/fwlink/?LinkId=58064. |

**Licensing**
The Microsoft Software License Terms for Windows Services for UNIX 3.0 explicitly allow for Server for NFS Authentication to be installed on all DCs without purchasing additional licenses. This is not a concern for Windows Services for UNIX 3.5, which is available as a free download from the Microsoft Download Center: http://go.microsoft.com/fwlink/?LinkId=54483.

# Open Ports on Firewalls

Applies To: Windows Server 2003 R2

To utilize Network File System (NFS), you need to configure the firewalls running on your network. Typically, externally facing firewalls should not be configured to pass NFS or authentication requests. However, internally facing firewalls, such as the Windows Firewall in Windows XP SP2, will need configuration. You need to open ports 111, 1039, 1047, 1048, and 2049 for both User Datagram Protocol (UDP) and TCP. For complete details about ports that you must open in Windows XP Service Pack 2 to enable access to Windows Services for UNIX 3.5, see the Microsoft Support article at: http://go.microsoft.com/fwlink/?LinkId=75006.

# Creating a Network File System Shared Network Resource

Applies To: Windows Server 2003 R2

Creating a Network File System shared network resource is exactly like creating any other shared network resource in Windows.

**How to create an NFS shared network resource**
If you use Windows Explorer to create the shared network resource, follow these steps:

1. Highlight the folder that you want to share.
2. Right-click and select **Sharing and Security** from the **Action** menu.
3. Click the **NFS Sharing** tab (shown in Figure 8).
4. Click **Share this folder** and enter a shared network resource name. The defaults used are:
   - Encoding: ANSI
   - Do not allow anonymous access
   - Read only permissions to all computers, with no root access allowed
5. To change the default permissions, click the **Permissions** button (shown in Figure 9).
6. You can change the behavior for all computers or specify a specific computer or group of computers by clicking the **Add** button.
7. Click **OK** as necessary to complete the wizard and create the NFS shared network resource.

If you want to create an NFS shared network resource from the command line, use **nfsshare**. To see the full syntax of **nfsshare**, type the following in a command window:

```
nfsshare /?
```

The command to export an NFS shared network resource from the E:\Public folder to /public to all computers with read/write privileges and root allowed would be:

```
nfsshare public=e:\public –o rw –o root
```




**Figure 8. The NFS Sharing tab**




**Figure 9. NFS Share Permissions dialog box**

# Conclusion

Applies To: Windows Server 2003 R2

The Windows Server 2003 R2 release builds on the UNIX interoperability components of Windows Services for UNIX 3.5 by thoroughly integrating them into the core operating system and providing improved administrative consoles for them.

The Microsoft Identity Management for UNIX (IDMU) components support the mechanisms used by Windows Services for UNIX 3.5 and extend them directly into Active Directory Domain Lookup.

The Network File System components of Windows Server 2003 R2 support the Windows Services for UNIX 3.5 authentication methods, while simplifying the administration and supporting Active Directory Domain Lookup for UNIX user identifier (UID) and group identifier (GID).

# For More Information

Applies To: Windows Server 2003 R2

- "A description of the ports that you must open in Windows XP SP2 to enable access to Windows Services for UNIX 3.5" (This article was written for Windows Services for UNIX 3.5, but it also applies to Network File System in Windows Server 2003 R2.)

  http://go.microsoft.com/fwlink/?LinkId=75006

- "HOW TO: Migrate Existing NIS Maps to Server for NIS in a UNIX-to-Windows Migration"

  http://go.microsoft.com/fwlink/?LinkId=75007

- "nis2ad"

  http://go.microsoft.com/fwlink/?LinkId=75008

- "Checklist: NIS migration to Active Directory using the NIS Data Migration Wizard"

  http://go.microsoft.com/fwlink/?LinkId=75009

- Windows Services for UNIX 3.5 free download

  http://go.microsoft.com/fwlink/?LinkId=54483

- "Sysocmgr"

  http://go.microsoft.com/fwlink/?LinkId=75011

- "How to Add or Remove Windows Components with Sysocmgr.exe"

  http://go.microsoft.com/fwlink/?LinkId=58064

- The public Microsoft Services for UNIX newsgroup

  news://msnews.microsoft.com/microsoft.public.servicesforunix.general

  (To access the newsgroup through a Web browser: http://go.microsoft.com/fwlink/?LinkId=67433)

# Network File System in Microsoft Unified Data Storage Server 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Services for Network File System (NFS) is a new feature of Microsoft® Windows Server® 2003 R2 that builds on the authentication framework introduced in Windows® Services for UNIX. It also adds additional functionality and authentication mechanisms that utilize Active Directory® Domain Services and the new Identity Management for UNIX to identify which user is requesting access to a resource and to authenticate that request to the appropriate authority. This white paper covers the various mechanisms that are involved in authentication and focuses on the new mechanisms introduced in Windows Unified Data Storage Server 2003.

**Contents**

Introduction

Concepts

Windows Unified Data Storage Server 2003 Configuration

Windows Server 2003 R2 Configuration

Mapping UNIX Users to Windows Users

Open Ports on Firewalls

Creating a Network File System Shared Network Resource

Conclusion

For More Information

# Introduction

Applies To: Windows Server 2003 R2

Microsoft® Windows Server® 2003 R2 extends the Microsoft commitment to full interoperability between Microsoft Windows® and UNIX by integrating the key components necessary for interoperability directly into the core Windows Server product.

| **Note** |
| --- |
| Throughout this document, we use the term "UNIX" to refer to resources and authentication of UNIX-based users and computers. |

The interoperability components in the Windows Server 2003 R2 family of products are functionally similar to their equivalent modules in Windows Services for UNIX 3.5. However, there are many important enhancements and updates since the release of Windows Services for UNIX 3.5.

Windows Services for UNIX 3.5 is available for use on 32-bit versions of Microsoft Windows 2000, Microsoft Windows XP Professional Edition, and Microsoft Windows Server 2003. The UNIX interoperability components of Windows Server 2003 R2 are available and supported on all versions of the product, including Windows Storage Server 2003 R2 and Windows Unified Data Storage Server 2003. All general references to Windows Server 2003 R2 also apply to Windows Storage Server 2003 R2 and Windows Unified Data Storage Server 2003.

| **Note** |
| --- |
| The Active Directory® Domain Services schema is updated on your domain controller (DC) during the Windows Server 2003 R2 installation. This schema update includes the necessary extensions for Network File System (NFS) authentication, and an additional schema update is not required to enable the components described in this white paper. |

**In this section**


What Is New in Windows Server 2003 R2?

NFS Components in Windows Server 2003 R2 and Past Releases

# What Is New in Windows Server 2003 R2?

Applies To: Windows Server 2003 R2

Windows Server 2003 R2 includes a suite of UNIX interoperability components as part of the core operating system. These components include:

- Identity Management for UNIX
    - Password synchronization
    - Server for NIS
- Services for Network File System (NFS)
    - Client for NFS
    - Server for NFS
    - User Name Mapping service
    - Server for NFS Authentication
    - RPC port mapper
- Subsystem for UNIX-based applications (SUA)

In addition to these components, the administration tools have been updated and improved. For example, the graphical administration tools include the Services for NFS console (shown in Figure 1) and the Microsoft Identity Management for UNIX console (shown in Figure 2).



**Figure 1. The Microsoft Services for NFS console**




**Figure 2. The Microsoft Identity Management for UNIX console**

Another important feature in the UNIX interoperability components is support for 64-bit operating systems. Windows Services for UNIX 3.5 does not support 64-bit Windows, but Windows Server 2003 R2 fully supports 64-bit interoperability.

| **Note** |
| --- |
| Microsoft Windows XP Professional x64 Edition has no support for any UNIX interoperability components, nor does it allow you to install Windows Services for UNIX 3.5. |

# NFS Components in Windows Server 2003 R2 and Past Releases

Applies To: Windows Server 2003 R2

A number of components are either involved in or dependent on Network File System (NFS) authentication. These include:

- **User Name Mapping service**. This component maps UNIX users to Windows users and vice versa. Even when a user references the same name on both systems, the programs do not recognize them as the same user. User Name Mapping service tells the components in Services for NFS that Windows user jdoe is the same as UNIX user johnd.

- **Server for NFS Authentication**. This component is not a server at all, but rather an authentication component used by Server for NFS. If you use local user accounts, install Server for NFS Authentication on all Windows NFS servers. If you use domain user accounts, raise the domain functional level to Windows Server 2003 or install Server for NFS Authentication on all domain controllers.

- **Client for NFS**. This is the Windows NFS client component of Services for NFS. Client for NFS allows the computer on which it is installed to access and use NFS resources anywhere on the network.

- **Server for NFS**. This is the Windows NFS server component of Services for NFS. Server for NFS allows the machine on which it is installed to provide file system resources to NFS clients anywhere on the network.

Each of these components included in Windows Server 2003 R2 have equivalent components in the earlier Windows Services for UNIX releases. However, two components from Windows Services for UNIX 3.5 are not included in Windows Server 2003 R2: Server for Personal Computer Network File System (PCNFS) and Gateway for NFS.

Windows Services for UNIX 1.0 uses the PCNFS protocol for authentication to NFS servers, and some third-party legacy NFS applications have also used PCNFS. However, since Windows Services for UNIX 2.0, the authentication mechanism has not used PCNFS directly, though the Server for PCNFS was included to support those legacy clients.

The Gateway for NFS component attempted to provide a single server that connected to NFS servers and shared the resulting connection to downstream Windows clients by using the standard Windows networking protocol: server message blocks (SMB). This component, however, was a potential bottleneck in the connection and introduced substantial network overhead. With the availability of Client for NFS, the need for a gateway product has been substantially reduced.

# Concepts

Applies To: Windows Server 2003 R2

Windows and UNIX use very different authentication mechanisms. The Microsoft Identity Management for UNIX components of Windows Server 2003 R2 build on the identity management components originally included in Windows Services for UNIX (User Name Mapping service and Server for NIS). Windows Server 2003 R2 adds direct support for Active Directory Domain Services Lookup while continuing to support the legacy Windows Services for UNIX mechanisms. The following conceptual material applies to the Windows Server 2003 R2 family, the Windows Storage Server 2003 R2 family, and the Windows Unified Data Storage Server 2003 family.

**In this section**

Windows and UNIX User and Group Identities

Integrated Identity Management in Active Directory Domain Services

Identity Matching in User Name Mapping Service

# Windows and UNIX User and Group Identities

Applies To: Windows Server 2003 R2

When a UNIX user requests access to a Windows resource, or when a Windows user requests access to a UNIX resource, the computers need to map the UNIX user to the proper Windows user or vice versa. Windows Services for UNIX 3.5 uses User Name Mapping service (still used by Windows Server 2003 R2). However, Windows Server 2003 R2 adds direct lookup of UNIX User ID (UID) and Group ID (GID) from Active Directory Domain Services, with no User Name Mapping service required.

When a UNIX user requests access to a file or folder ("object") stored on a Windows Network File System (NFS) shared network resource, the first step is for Windows to map the UNIX user to the Windows user. This is accomplished through Active Directory Domain Lookup or User Name Mapping service.

The request for access to the Windows object is authenticated against a Windows domain controller by using the security identifier (SID) of the mapped Windows user. (In the special case of using a local Windows account with a local User Name Mapping service, the request is authenticated against the local Windows computer.)

If the mapped Windows user has the appropriate access permissions, the UNIX user who initiated the request is granted access to the object. Conversely, if the Windows user does not have sufficient permission for the UNIX user to access the object, access is denied to the UNIX user.

Access is not mapped to a Windows user and the permissions of an anonymous user are used if:

- The UNIX user does not have a mapped Windows account.
- The User Name Mapping service can't be reached.
- Active Directory Domain Services does not include the UNIX user's identity data.

When a Windows user tries to access an object stored on a UNIX computer, the same process is used but in reverse.

# Integrated Identity Management in Active Directory Domain Services

Applies To: Windows Server 2003 R2

As part of the product installation, the Windows Server 2003 R2 schema is extended to support direct lookup of UNIX identities in Active Directory Domain Services. You have the choice of running Server for NIS or managing the schema attributes directly.

**UNIX Attributes tab**
Server for NIS adds a **UNIX Attributes** tab in the Active Directory Domain Services Users and Computers console. This addition:

- Provides you with direct graphical user interface (GUI) access to the schema extensions.
- Allows you to directly set attributes for each user (as shown in Figure 3).

The above requires that Server for NIS be installed on a Windows Server 2003 R2 domain controller.




**Figure 3. UNIX Attributes tab for user Charlie Russel**

If access is requested to a Windows resource, the access requires the permissions of a Windows user. Active Directory Domain Services looks up the UID/GID of the UNIX user requesting access and maps the request to the Windows user to set the appropriate permissions.

If access is requested from a Windows user to a UNIX resource, the Windows user's UID/GID pair is provided to the UNIX authentication mechanism to set appropriate permissions.

# Identity Matching in User Name Mapping Service

Applies To: Windows Server 2003 R2

The User Name Mapping service was the core authentication component of Windows Services for UNIX, and it continues to be supported in Services for Network File System (NFS). However, it is important to note that this is the final release of the User Name Mapping service, and it will not be supported in future releases of Services for NFS.

Although User Name Mapping service is at the core of the authentication processes for Windows Services for UNIX 3.5, it does not perform the authentication. The authentication is performed by the:

- Windows domain controller or the Local Security Authority Subsystem Service (LSASS) if access permissions are requested by a Windows user (when Server for NFS shares files with UNIX users).

- UNIX authentication mechanism if access permissions are requested by a UNIX user (when Windows users access UNIX-based shared network resources through Client for NFS).

In other words, if a user requests access to a shared network resource on a Windows computer, that computer uses a Windows authentication mechanism. If a user requests access to a shared network resource on a UNIX computer, that computer uses UNIX authentication.

One other important authentication occurs when a request comes to the User Name Mapping service. The .maphosts file resides in the %SFUDIR%\Mapper subdirectory of the Windows Services for UNIX 3.5 installation directory or the %WINDIR%\msnfs directory in Windows Server 2003 R2. The .maphosts file is checked to see whether the requesting computer is permitted to access the User Name Mapping service. If it is not, the request is denied.

Several steps are involved in any authentication, and any of the steps could cause an authentication failure. It is important to understand the process to troubleshoot it effectively.

# Windows Unified Data Storage Server 2003 Configuration

Applies To: Windows Server 2003 R2

The content in this section applies only to Microsoft® Windows® Unified Data Storage Server 2003 operating system. It is recommended that you follow these configuration steps. (The steps described in the "Windows Server 2003 R2 Configuration" section are also valid for Windows Unified Data Storage Server 2003.)

**Configuration**

In Windows Unified Data Storage Server 2003, the initial configuration for Microsoft Services for Network File System (NFS) can be completed by using the Microsoft Services for NFS Configuration Guide (shown in Figure 5). This guide includes the Identity Mapping Setup Wizard.

After completing the initial configuration, use the:

- Microsoft Services for NFS console for ongoing maintenance and administration.
- NFSADMIN.EXE command-line tool for advanced tasks.

To configure Microsoft Services for NFS, follow these steps:

1. In the Initial Configuration Tasks window (shown in Figure 4), select **Configure Microsoft Services for Network File System (NFS)**.




**Figure 4 Initial Configuration Tasks window**

2. The Microsoft Services for NFS Configuration Guide opens. This guide will help you configure your server to share files with UNIX-based computers by using Microsoft Services for NFS (shown in Figure 5).




**Figure 5 Microsoft Services for NFS Configuration Guide**

3. In the right pane, click the **Identity Mapping Wizard** button.

4. Active Directory Domain Services and User Name Mapping (UNM) are available for identity mapping. Select **Retrieve identity mappings from Active Directory (recommended)**, **Retrieve identity mappings from User Name Mapping**, or both, and then click **Next**.




**Figure 6 Identity Mapping Wizard**

5. Depending on which option you selected in Step 4, different steps are required.

For Active Directory Domain Services:

| **Note** |
| --- |
| Administrators must ensure that Active Directory Domain Services contains UNIX-style identity data. |

1. In the Active Directory Domain Services window of the Identity Mapping Setup Wizard (shown in Figure 7), type the path to the Active Directory Domain Services or other compatible directory service that you want to use.

2. Click **Next** to review your settings (shown in Figure 8).

3. Select **Configure**, and then select **Close** to finalize your Active Directory Domain Services mapping setup.




**Figure 7 Active Directory Domain Services**




**Figure 8 Review settings and configure identity mapping**




**Figure 9 Confirmation**

For User Name Mapping:

1. In the User Name Mapping Server window of the Identity Mapping Setup Wizard, type the path to the User Name Mapping host (can be set to **localhost**).

2. Click **Next** to review your settings (shown in Figure 11).

3. Select **Configure**, and then select **Close** to finalize your User Name Mapping setup.

4. The Confirmation window (shown in Figure 12) confirms the selected settings.




**Figure 10 Identity Mapping Setup Wizard**




**Figure 11 Review settings and configure identity mapping**




**Figure 12 Confirmation screen**

| **Note** |
| --- |
| If you are using a remote User Name Mapping server, you must ensure that the remote UNM server is configured. If a local UNM server is used, extra configuration steps are required, including: **Configure Password and Group Files**, **Create Simple Maps**, **Create Advanced User Maps**, and **Create Advanced Group Maps**. |

6. To establish a shared resource, you may have to launch the Microsoft Services for NFS Configuration Guide. In the Initial Configuration Tasks window (shown in Figure 4), click **Configure Microsoft Services for Network File System (NFS)**. Select **Use NFS to Share Folders**, and then select **Provision a Shared Folder Wizard**.

1. To specify the folder that you want to share, click **Browse**, and then choose an existing folder with the appropriate characteristics and space. If one does not exist, select **Provision Storage** and create one.




**Figure 13 Provision a Shared Folder Wizard**

2. Establish the NTFS permissions desired for the shared resource, and then click **Next**.

| **Note** |
| --- |
| Ensure that the NTFS permissions match the intended use of the shared resource. |

7. To create an NFS shared resource, enter in the name of the shared resource in the **Share name** field. The path will populate automatically. Click **Next**.Specify permissions for NFS-based access to the shared resource folder, and then click **Next**.




**Figure 14 Provision a Shared Folder Wizard**

| **Note** |
| --- |
| The level of restriction placed on the shared resource will determine the access available to users or groups. |

8. Optionally, a quota policy can be specified to limit the maximum size of the shared resource folder. To apply a quota to the shared resource, select **Apply quota**, and then select either **Create quota on path** or **Auto apply template to create quotas on existing and new subfolders**. Satisfy the quota's need by selecting a quota template, and then click **Next**.




**Figure 15 Quota policy configuration**

9. File screens are used to control the types of files that the folder can contain. To apply a file screen, select **Apply file screen**, and then select the template to satisfy the the needs of the shared resource. Click **Next**.




**Figure 16 File screening configuration**

10. Review the options that you selected in the previous steps. To provision the shared resource with the listed settings, click

**Create**. To make any changes, click **Previous** or select the appropriate page in the orientation pane.




**Figure 17 Review settings and create shared resource**

| **Note** |
| --- |
| When the initial configuration is complete, tasks may be accomplished by repeating the previous steps (or by following the steps listed in the "Windows Server 2003 R2 Configuration" section). |

# Windows Server 2003 R2 Configuration

Applies To: Windows Server 2003 R2

The Configuration section of this white paper covers the basic installation and configuration of the Windows Server 2003 R2 components used by Services for Network File System (NFS) and Server for NFS Authentication. The following steps are also valid on Windows Unified Data Storage Server 2003.

**In this section**


Installing Services for Network File System

Installing Identity Management for UNIX

Extending the Active Directory Domain Services Schema

Importing UNIX Identities into Active Directory by Using IDMU

# Installing Services for Network File System

Applies To: Windows Server 2003 R2

**Installing UNIX Interoperability Components**

When you upgrade from Windows Server 2003 Service Pack 1 (SP1) to Windows Server 2003 R2, the installation program does not install the UNIX interoperability components on your server. The installation only adds to the list of components that can be installed. Follow these steps to perform the installation:

1. Open **Add/Remove Programs**, and then click the **Add/Remove Windows Components** button to start the Windows Components Wizard (shown in Figure 4).




**Figure 18. The Windows Components Wizard**

2. Click **Other Network File and Print Services**, and then click **Details**.

3. Click **Highlight Services for NFS**, and then click **Details** to open the Services for NFS dialog box (shown in Figure 5).




**Figure 19. The Windows Components Wizard, Microsoft Services for NFS page**

4. Select the NFS-related services you want, and follow the wizard prompts to finish.

| **Note** |
| --- |
| Installing Services for NFS may require that you restart the computer, so plan your installation accordingly. |

# Installing Identity Management for UNIX

Applies To: Windows Server 2003 R2

Server for NIS is the component of Microsoft Identity Management for UNIX (IDMU) that is of primary importance to Network File System (NFS) authentication, and it can be installed only on a Windows Server 2003 R2 domain controller. Server for NIS cannot be installed on a member server. To install Server for NIS, follow these steps:

1. Open **Add/Remove Programs**, and then click the **Add/Remove Windows Components** button to start the Windows Components Wizard (as shown in Figure 4 in the Installing Services for NFS section of this white paper).

2. Click **Active Directory Services**, and then click **Details**.

3. Click **Highlight Identity Management for UNIX**, and then click **Details** to open the **Identity Management for UNIX** dialog box (shown in Figure 6).

4. Click **Administration Components and Server for NIS** and follow the prompts to install the NIS components of Windows Server 2003 R2. You will need to have the Windows Server 2003 R2 installation media available.

| **Note** |
| --- |
| Installing IDMU may require that you restart the computer, so plan your installation accordingly. |

| **Note** |
| --- |
| If you are using User Name Mapping service, you do not need to install Server for NIS, nor do you need to populate Active Directory Domain Services with UNIX identity information. |




**Figure 20. Microsoft Identity Management for UNIX dialog box**

# Extending the Active Directory Domain Services Schema

Applies To: Windows Server 2003 R2

The typical Windows Server 2003 R2 installation on a domain controller provides the necessary extension of the Active Directory Domain Services schema. No additional extension is required to authenticate Network File System users. However, unless Server for NIS is installed, the Active Directory Domain Services tabs for editing UNIX attributes are not visible. You can still assign UNIX attributes to users programmatically or by using Active Directory Security Interfaces (ADSI) Edit.

# Importing UNIX Identities into Active Directory by Using IDMU

Applies To: Windows Server 2003 R2

If your organization has an existing Network Information Service (NIS) domain environment, you can import your NIS users into Active Directory Domain Services by using the NIS Data Migration Wizard. The NIS Data Migration Wizard can be started from the Microsoft Identity Management for UNIX (IDMU) console or by using the **nis2ad** utility.

**For more information**

For information about how to use the NIS Data Migration Wizard, see:

- The IDMU Help files
- Microsoft TechNet article about migrating to Active Directory Domain Services using the NIS Data Migration Wizard:

  http://go.microsoft.com/fwlink/?LinkId=75009

For information about the **nis2ad** utility, see:

- Microsoft Support "How to" article about migrating existing NIS Maps to Server for NIS: http://go.microsoft.com/fwlink/?LinkId=75007
- TechNet article about **nis2ad**: http://go.microsoft.com/fwlink/?LinkId=75008

# Mapping UNIX Users to Windows Users

Applies To: Windows Server 2003 R2

Services for Network File System (NFS) supports:

- Windows Services for UNIX 3.5 User Name Mapping service
- Direct lookup of UNIX UID/GID pairs for Windows users (through Active Directory Domain Services)

Note that this is the final release that will include the User Name Mapping service as part of Services for NFS. Customers should migrate to support direct lookup of UNIX identities in Active Directory Domain Services.

**In this section**

Configuring Services for Network File System to Use Active Directory Domain Lookup

Configuring Services for Network File System to Use User Name Mapping Service

# Configuring Services for Network File System to Use Active Directory Domain Lookup

Applies To: Windows Server 2003 R2

The preferred method to map Windows users to UNIX users is to use Active Directory Domain Lookup. Both Client for NFS and Server for NFS use one of two lookup methods: Active Directory Domain Lookup or User Name Mapping service.

**Configuring Services for NFS to use Active Directory Domain Lookup**
To configure Services for NFS to use Active Directory Domain Lookup, follow these steps:

1. Open the Services for NFS console.
2. In the left pane, at the top of the tree, right-click **Services for NFS**.
3. On the **Action** menu, click **Properties** to open the dialog box (shown in Figure 7).
4. Select the **Active Directory Domain Lookup** check box, and then enter the Active Directory Domain Services domain to which to connect.
5. Click **OK**.

| **Note** |
| --- |
| Configuring Services for NFS to use Active Directory Domain Lookup may require that you restart your computer or Services for NFS before the change is enabled. |




**Figure 21. Services for NFS Properties**

# Configuring Services for Network File System to Use User Name Mapping Service

Applies To: Windows Server 2003 R2

User Name Mapping service can be used instead of Active Directory Domain Lookup to map UNIX users to Windows users and Windows users to UNIX users. Any version of User Name Mapping service from Windows Services for UNIX 3.0 and later can be used by any version of Windows Services for UNIX or the Services for NFS from Windows Server 2003 R2. Windows Server 2003 R2 includes a version of the User Name Mapping service.

**Configuring Services for NFS to use User Name Mapping service**
To configure Services for NFS to use User Name Mapping service, follow these steps:

1. Open the Services for NFS console.
2. In the left pane, at the top of the tree, right-click **Services for NFS**.
3. From the **Action** menu, click **Properties** to open the dialog box shown in Figure 7 in the Configuring Services for Network File System to Use Active Directory Domain Lookup section of this white paper.
4. Enter the name or IP address of the User Name Mapping service.
5. Click **OK**.

| **Note** |
| --- |
| Configuring Services for NFS to use User Name Mapping service may require that you restart your computer or Services for NFS before the change is enabled. |

**Domain accounts vs. local accounts**
If you are running in a domain environment, you should generally use only domain accounts for your mapping. Local accounts are valid only on the single, local computer on which they are located, so you would need to use multiple User Name Mapping services (one for each local computer) if you intended to access resources on more than one Windows server. Obviously, if you are using Active Directory Domain Lookup, you are required to you use domain accounts.

**Simple maps**
Simple maps are maps between accounts with the exact same name in UNIX and Windows. Keep in mind that although the accounts have the same name and the exact same spelling, they are not the same account. So until you actually create the maps, the accounts will not line up. For the purposes of User Name Mapping service, the account name in Windows is not case sensitive, so Charlie will align with the UNIX account Charlie, charlie, or CHARLIE.

**Advanced maps**
Advanced maps are used to align accounts that have different names. For example, the Windows user account for Jane H. Doe might be jdoe, but the UNIX account could be jhd. Create an advanced map to align the two accounts. You can also use advanced maps to map users from different Windows domains and explicitly map accounts that would generally be mapped by simple maps.

**Authorized hosts**
To prevent unauthorized access to the User Name Mapping service, the file .maphosts is used to explicitly specify which computers may access the server. This file is installed when the User Name Mapping service is installed. In Windows Services for UNIX 3.5, the file resides in the **%SFUDIR%\Mapper** directory, and with Windows Server 2003 R2, the file resides in **%WINDIR%\msnfs**.

By default, the only computer that can access the User Name Mapping service is the local computer. You need to edit the .maphosts file with a pure ASCII text editor, such as Notepad or gvim, and add the host names or IP addresses of all the hosts that are permitted to access User Name Mapping service. Remember to include any Windows computer on which you have installed Client for NFS or Server for NFS.

During initial setup and configuration of the Server for NFS Authentication, it is a good idea to configure .maphosts to allow all computers to connect to User Name Mapping service. This can be changed after authentication is working correctly. To enable

all computers to access User Name Mapping service, type a plus sign (+) on an empty line at the bottom of the file. (For more information about the syntax of entries in .maphosts, see the comments in your .maposts file.)

**Configure domain authorization**
In addition to configuring User Name Mapping service, Server for NFS requires one additional step before UNIX users can be properly authenticated to access shared network resources on Windows servers if any of the Windows domain controllers (DCs) are not running Windows Server 2003 or if your domain is running at less than Windows Server 2003 domain functional level.

The Server for NFS Authentication must be installed on all Windows 2000 Servers and pre–Windows Server 2003 R2 DCs that has mapped users (available from the Microsoft Download Center: http://go.microsoft.com/fwlink/?linkid=53938).

**Verifying installation**
The easiest way to verify that the Server for NFS Authentication installation is complete is to check for the presence of the file %WINDIR%\system32\nfssa.dll. Server for NFS Authentication does not install any new service on a computer, and its presence is not easily discovered on the **Add/Remove Programs** list.

**Installing only Server for NFS Authentication**
In Windows Services for UNIX 3.0 and 3.5, you can easily install only Server for NFS Authentication by using the command line. This allows for easy scripting across multiple DCs. The command line to silently install Server for NFS Authentication is:

```
msiexec /I D:\<path>\sfusetup.msi /q addlocal="NFSServerAuth"
```

`D:\<path>` is the location of the Windows Services for UNIX 3.5 installation file.

Note that the `addlocal` parameter:

- Is case sensitive.
- Will remove any other components of Windows Services for UNIX 3.0 except Server for NFS Authentication.

If you are installing multiple components from the command line, you must use a comma-separated list within the quotes of the `addlocal` parameter. For example, to install both Server for NFS and Server for NFS Authentication, use:

```
msiexec /I D:\sfusetup.msi /q addlocal="NFSServer,NFSServerAuth"
```

| **Note** |
| --- |
| If you are using a version of Windows Services for UNIX that is not the free 3.5 download, you need to add `pidkey="key"` to this command line, where `"key"` is the 25-character product key without spaces or hyphens. |

In Windows Server 2003 R2, you can install the Server for NFS Authentication component from CD2. The .inf file is: \CMPNENTS\R2\PACKAGES\NEWBINS\<ARCHITECTURE>\NFSAUTH.INF.

- <ARCHITECTURE> is replaced by the appropriate folder for your server architecture.
- The actual DLL that is installed is: \CMPNENTS\R2\NFSSA.DL_.

| **Note** |
| --- |
| For information about how to add Windows components by using the Sysocmgr.exe utility, see the Microsoft Support article a http://go.microsoft.com/fwlink/?LinkId=58064. |

**Licensing**
The Microsoft Software License Terms for Windows Services for UNIX 3.0 explicitly allow for Server for NFS Authentication to be installed on all DCs without purchasing additional licenses. This is not a concern for Windows Services for UNIX 3.5, which is available as a free download from the Microsoft Download Center: http://go.microsoft.com/fwlink/?LinkId=54483.

# Open Ports on Firewalls

Applies To: Windows Server 2003 R2

To utilize Network File System (NFS), you need to configure the firewalls running on your network. Typically, externally facing firewalls should not be configured to pass NFS or authentication requests. However, internally facing firewalls, such as the Windows Firewall in Windows XP SP2, will need configuration. You need to open ports 111, 1039, 1047, 1048, and 2049 for both User Datagram Protocol (UDP) and TCP. For complete details about ports that you must open in Windows XP Service Pack 2 to enable access to Windows Services for UNIX 3.5, see the Microsoft Support article at: http://go.microsoft.com/fwlink/?LinkId=75006.

# Creating a Network File System Shared Network Resource

Applies To: Windows Server 2003 R2

Creating a Network File System shared network resource is exactly like creating any other shared network resource in Windows.

**How to create an NFS shared network resource**
If you use Windows Explorer to create the shared network resource, follow these steps:

1. Highlight the folder that you want to share.
2. Right-click and select **Sharing and Security** from the **Action** menu.
3. Click the **NFS Sharing** tab (shown in Figure 8).
4. Click **Share this folder** and enter a shared network resource name. The defaults used are:
   - Encoding: ANSI
   - Do not allow anonymous access
   - Read only permissions to all computers, with no root access allowed
5. To change the default permissions, click the **Permissions** button (shown in Figure 9).
6. You can change the behavior for all computers or specify a specific computer or group of computers by clicking the **Add** button.
7. Click **OK** as necessary to complete the wizard and create the NFS shared network resource.

If you want to create an NFS shared network resource from the command line, use **nfsshare**. To see the full syntax of **nfsshare**, type the following in a command window:

```
nfsshare /?
```

The command to export an NFS shared network resource from the E:\Public folder to /public to all computers with read/write privileges and root allowed would be:

```
nfsshare public=e:\public –o rw –o root
```




**Figure 22. The NFS Sharing tab**




**Figure 23. NFS Share Permissions dialog box**

# Conclusion

Applies To: Windows Server 2003 R2

The Windows Server 2003 R2 release builds on the UNIX interoperability components of Windows Services for UNIX 3.5 by thoroughly integrating them into the core operating system and providing improved administrative consoles for them.

The Microsoft Identity Management for UNIX (IDMU) components support the mechanisms used by Windows Services for UNIX 3.5 and extend them directly into Active Directory Domain Lookup.

The Network File System components of Windows Server 2003 R2 support the Windows Services for UNIX 3.5 authentication methods, while simplifying the administration and supporting Active Directory Domain Lookup for UNIX user identifier (UID) and group identifier (GID).

# For More Information

Applies To: Windows Server 2003 R2

- "A description of the ports that you must open in Windows XP SP2 to enable access to Windows Services for UNIX 3.5" (This article was written for Windows Services for UNIX 3.5, but it also applies to Network File System in Windows Server 2003 R2.)

  http://go.microsoft.com/fwlink/?LinkId=75006

- "HOW TO: Migrate Existing NIS Maps to Server for NIS in a UNIX-to-Windows Migration"

  http://go.microsoft.com/fwlink/?LinkId=75007

- "nis2ad"

  http://go.microsoft.com/fwlink/?LinkId=75008

- "Checklist: NIS migration to Active Directory using the NIS Data Migration Wizard"

  http://go.microsoft.com/fwlink/?LinkId=75009

- Windows Services for UNIX 3.5 free download

  http://go.microsoft.com/fwlink/?LinkId=54483

- "Sysocmgr"

  http://go.microsoft.com/fwlink/?LinkId=75011

- "How to Add or Remove Windows Components with Sysocmgr.exe"

  http://go.microsoft.com/fwlink/?LinkId=58064

- The public Microsoft Services for UNIX newsgroup

  news://msnews.microsoft.com/microsoft.public.servicesforunix.general

  (To access the newsgroup through a Web browser: http://go.microsoft.com/fwlink/?LinkId=67433)

# Preserving OEM Preactivation When Reinstalling Windows Server 2003 R2

Updated: June 17, 2009

Applies To: Windows Server 2003 R2

Large server manufacturers, known as original equipment manufacturers (OEMs), have the ability to preactivate the Microsoft® Windows Server® 2003 R2 software that they install on new server. As a result, end users are typically not required to activate the software during the set up process by typing in the product key found on the Certificate of Authenticity attached to the server (COA Key).

Due to the form factor and typical installations of server hardware, it is not always easy for the product key to be read during the set up process. To help minimize the potential impact on licensed users, this whitepaper provides steps that licensed end users can take to preserve OEM preactivation when reinstalling the operating system by using any of the three following scenarios: Sysprep, manual, or unattended.

The product keys in the following table should be substituted for XXXXX-XXXXX-XXXXX-XXXXX-XXXXX in the solutions that follow.

| **Note** |
| --- |
| OEM recovery media must be used as the source for the reinstallation. |

| | |
| --- | --- |
| Windows Server 2003 R2 Standard Edition x86 | PWBJC-22697-D4CVH-FCJWW-DTF9J |
| Windows Server 2003 R2 Standard Edition x64 | XCP6P-7WVXP-F8FQ4-JV6CD-6XV28 |
| Windows Server 2003 R2 Enterprise Edition x86 | XHPV3-PTCWJ-7Y94F-Q6BVH-J849J |
| Windows Server 2003 R2 Enterprise Edition x64 | WQ3GW-Y8GQW-8VJYB-JYM43-D24C8 |
| Windows Server 2003 R2 Datacenter Edition x86 | FXYF6-VTXGX-3JPX9-HJ9K4-6TKTW |
| Windows Server 2003 R2 Datacenter Edition x64 | KDX8X-FYW4T-C6D9J-BKM6M-M89TW |
| Windows Server 2003 R2 Web Edition x86 | GM8KD-GB7JY-QGQYP-XRV74-RT728 |

For any of the following solutions, you can verify that the OEM preactivation has been preserved:

Click **Start**, click **Run**, type **%SYSTEMROOT%\system32\oobe\msoobe.exe /A**, and then click **OK**. If the procedure was successful, a message that indicates Windows has already been activated appears.

| **Note** |
| --- |
| Instead of implementing any of the following solutions, Volume License customers can use Volume License media to install the same operating system that was licensed from an OEM without acquiring a volume license for a server that is being re-imaged. For more information about this option, click Re-Imaging Rights on the Microsoft Volume Licensing Web site (http://go.microsoft.com/fwlink/?LinkId=154939). |

**Using the "sysprep –reseal –mini" command**

OEM preactivation can be preserved when Sysprep.exe is run to reseal the system by using the –reseal – mini switches. Follow these steps:

**To preserve OEM preactivation by using Sysprep**

1. Ensure that the following three files are present in the C:\Sysprep folder: sysprep.exe, setupcl.exe, and factory.exe.

These files can be found in the deploy.cab, which can be downloaded from the Microsoft® Download Center at the following locations:

- For x86-based computers:

  System Preparation Tool for Windows Server 2003 Service Pack 2 Deployment

- For x64-based computers:

  System Preparation Tool for Windows Server 2003 Service Pack 2 Deployment (x64)

2. Create a file named **sysprep.inf** by using a text editor (such as Notepad) that contains the following information. Substitute the appropriate product key from the table above.

   [UserData]

   ProductKey ="XXXXX-XXXXX-XXXXX-XXXXX-XXXXX"

3. Save the sysprep.inf file in **c:\sysprep\ folder**.

4. Run **c:\sysprep –reseal –mini**.

For more information on the Sysprep process, see Sysprep Tools and Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=154952).

**Manual install**
OEM preactivation can be preserved when doing a manual installation by inserting the appropriate product key listed above when prompted for a product key by Windows setup.

| **Note** |
| --- |
| If you run sysprep –reseal –mini following an unattended install, you need to follow the steps outlined above. |

**Unattended install**
OEM preactivation can be preserved if you use an unattended answer file (Unattend.txt or Winnt.sif) to automate installation by including the following information in the answer file (substitute the appropriate product key from the table above).

[UserData]

ProductKey ="XXXXX-XXXXX-XXXXX-XXXXX-XXXXX"

| **Note** |
| --- |
| If you run sysprep –reseal –mini following an unattended install, you need to follow the steps outlined above. |

For more information on the Sysprep process, see Sysprep Tools and Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=154952).

# IIS 6.0 Deployment Guide

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Welcome to Deploying Internet Information Services (IIS) 6.0 of the Microsoft® Windows® Server 2003 Deployment Kit.

IIS 6.0 provides the services to support a secure, available, and scalable Web server on which to run your Web sites and applications. This guide provides prescriptive, task-based, and scenario-based guidance to help you design an IIS 6.0 solution that meets the specific needs of your organization. Deployment scenarios include installing a new Web server, upgrading an existing Web server from an earlier version of IIS, and migrating existing Apache or IIS Web sites and applications to a newly installed Web server.

This guide contains the following sections:


Overview of Deploying IIS 6.0

Deploying ASP.NET Applications in IIS 6.0

Securing Web Sites and Applications

Ensuring Application Availability

Upgrading an IIS Server to IIS 6.0

Migrating IIS Web Sites to IIS 6.0

Migrating Apache Web Sites to IIS 6.0

Appendix A: IIS Deployment Procedures

Appendix B: Changes to Metabase Properties in IIS 6.0

Appendix C: Choosing a Cluster Technology for IIS 6.0

Appendix D: Running 32-bit Applications on 64-bit Windows

Appendix E: Unattended Setup

Additional Resources_IIS_SP1_Dep

# Overview of Deploying IIS 6.0

Updated: June 22, 2009

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To deploy Internet Information Services (IIS) 6.0, you must identify the key design and deployment processes that you must complete for your IIS 6.0 solution. This guide provides prescriptive, task-based, and scenario-based guidance to help you design an IIS 6.0 solution and then deploy that solution within your organization. After you deploy and secure your IIS 6.0 solution as a platform for your Web applications, you can then deploy your Web applications. However, the target audience of this guide is Web server administrators, not developers. Modifying Web applications is a developer-related activity and is described in other topics. It is not covered in this guide.

In This Section

Overview of Deploying an IIS 6.0 Web Server

Overview of IIS 6.0

Determining Application Compatibility with IIS 6.0

For information about how to deploy Microsoft Exchange in a cluster environment, see the Exchange Server 2003 Deployment Guide.

For information about how to configure Web server farms that use the same certificate, see How to load balance a Web server farm by using one SSL certificate in IIS 6.0 and in IIS 5.0.

# Overview of Deploying an IIS 6.0 Web Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Organizations and individuals use Web sites and applications every day as a way to do business on the Internet and within their intranets. Internet Information Services (IIS) 6.0 helps you meet your business needs by providing the services to support a secure, available, and scalable Web server on which to run these Web sites and applications.

| **Note** |
| --- |
| You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

This section describes the high-level processes that are presented in this guide for deploying a new IIS 6.0 Web server in your organization's production environment. The other sections in this guide are divided into separate IIS deployment topics that target a specific area of the deployment process including server security, application availability, deploying ASP.NET applications, Web site migration, and server upgrades. For a comprehensive understanding of IIS 6.0 deployment, read all of the section in sequential order. For information about a specific aspect of IIS 6.0 deployment, read the individual section that corresponds to your area of interest.

| **Note** |
| --- |
| If you are considering deploying IIS 6.0 Web servers in a cluster environment, see Appendix C: Choosing a Cluster Technology for IIS 6.0. You must complete the cluster deployment before you begin the IIS 6.0 deployment. |

Everyone deploying IIS 6.0 needs to decide in which application isolation mode IIS should run. This guide highlights the new IIS 6.0 worker process isolation mode because of the security and availability improvements found in this mode. This guide also compares worker process isolation mode to IIS 5.0 isolation mode, which is provided for maximum backward compatibility with existing applications. If your existing Web applications do not require the characteristics of IIS 5.0 isolation mode, you should run IIS in worker process isolation mode.

Finally, while you prepare to deploy IIS 6.0, you must verify that your existing Web sites and applications are compatible with IIS 6.0 and with the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; and Windows® Server 2003, Web Edition operating systems. Verification of the Web site and application compatibility should be done on a test Web server before deploying on a production Web server.

The processes in this guide have been carefully developed and tested to provide a blueprint for an easy and comprehensive deployment of IIS 6.0. Following the recommendations presented during the deployment process will help your Web servers be as secure as possible and highly available.

# Process for Deploying an IIS 6.0 Web Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 deployment process is written for Web server administrators who are responsible for installing and configuring IIS on new or existing servers. The sections in this guide can be divided into two main deployment scenarios:

- Deploying a new Web server running Windows Server 2003 with IIS 6.0
- Upgrading or migrating to a Web server running Windows Server 2003 with IIS 6.0

Figure 1.1 illustrates the sections in this guide that correspond to these scenarios.

| **Note** |
| --- |
| You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

Review the flowchart in this figure and read the descriptions of each section to help identify the deployment tasks that you need to complete for your IIS 6.0 solution and to discover which sections will help you complete those tasks.

**Figure 1.1 IIS 6.0 Deployment Process Steps and the Corresponding Sections**




Each section includes a flowchart that represents the organizing principle of the section content. Each section also includes a quick-start guide to facilitate the fastest possible deployment. You can use these quick-start guides to help identify the steps of the deployment process that you need additional information to complete, and then you can skip the information with which you are already familiar.

# Deploying a New IIS 6.0 Web Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following sections in this guide are about deploying a new server running IIS 6.0 to host your Web sites and applications:

- Deploying ASP.NET Applications in IIS 6.0

  Read this section to understand specific considerations for deploying ASP.NET applications in IIS 6.0. In particular, the section describes how you can run multiple versions of the Microsoft .NET Framework on the same Web server and how you can configure ASP.NET applications to use the appropriate version of the .NET Framework.

  | **Note** |
  | --- |
  | You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

  After completing the process in this section, you must still optimize your server for security and scalability. The remaining sections in this guide focus on topics that are applicable to all Web applications, including ASP.NET applications, noting exceptions where appropriate.

| **Important** |
| --- |
| ASP.NET is not available on the following operating systems: Microsoft® Windows® XP 64-Bit Edition; the 64-bit version of Windows® Server 2003, Enterprise Edition; and the 64-bit version of Windows® Server 2003, Datacenter Edition. For more information, see "Features unavailable on 64-bit versions of the Windows Server 2003 family" in Help and Support Center for Windows Server 2003. |

- Securing Web Sites and Applications

  Read this section to learn how you can further protect the Web sites and applications that are hosted on a Web server running IIS 6.0. This section describes how to secure a Web server and how to secure individual Web sites and applications running on IIS 6.0. In particular, this section describes the process for reducing the attack surface of a Web server, preventing unauthorized access to Web sites and applications, isolating Web sites and applications, configuring user authentication, encrypting confidential data exchanged with clients, and maintaining the security of Web sites and applications.

- Ensuring Application Availability

  Read this section to learn how you can fully utilize the features of IIS 6.0 to enhance the availability of your Web applications. This section describes how you can set realistic application availability goals and then verify that you are attaining those goals. In addition, it describes how you can configure the features of worker process isolation mode for optimum application availability, depending on your business needs and the processing load on your applications. Finally, this section identifies common application incompatibilities with IIS 6.0, and it explains how you can test your applications for compatibility with worker process isolation mode.

# Upgrading and Migrating a Server to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following sections in this guide assume that you have Web sites and applications that are hosted on an existing Web server running Apache or an earlier version of IIS. Select one of the following sections to perform the deployment of IIS 6.0 when you have an existing server:

- Upgrading an IIS Server to IIS 6.0

  Read this section when you are upgrading Web sites and applications that are hosted on a server running the Microsoft Windows NT® Server 4.0 operating system with IIS 4.0 or a server running the Microsoft® Windows® 2000 Server operating system with IIS 5.0, to a server running Windows Server 2003 with IIS 6.0. This section defines upgrading as the process of updating a Web server running IIS 4.0 or IIS 5.0 to run IIS 6.0 when the Web sites and applications hosted by that Web server are not moved to another server.

- Migrating IIS Web Sites to IIS 6.0

  Read this section when you are migrating Web sites that are hosted on a server running Windows NT Server 4.0 with IIS 4.0, or a server running Windows 2000 Server with IIS 5.0, to a newly installed server running Windows Server 2003 with IIS 6.0. This section assumes that you are moving Web sites from an existing Web server running an earlier version of IIS, to a newly installed Web server running IIS 6.0. This section explains how you can use the IIS 6.0 Migration Tool to migrate your IIS Web sites to IIS 6.0. This section also describes a manual process for migrating IIS Web sites to IIS 6.0.

- Migrating Apache Web Sites to IIS 6.0

  Read this section when you are migrating Apache Web sites that are hosted on an existing server to a server running IIS 6.0 and Windows Server 2003. This section assumes that you are moving applications from an existing Web server running a Linux-based operating system with Apache to a newly installed Web server running Windows Server 2003 with IIS 6.0. The process in this section describes how you can use the Apache to IIS 6.0 Migration Tool to migrate your Apache Web sites to IIS 6.0.

# Overview of IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 with Windows Server 2003 provides integrated, reliable, scalable, secure, and manageable Web server capabilities over an intranet or the Internet. IIS is a stable and secure platform for running dynamic network applications. Organizations of all sizes use IIS to host and manage Web sites on the Internet or on their intranets, to host and manage FTP sites, and to route news or mail by using the Network News Transfer Protocol (NNTP) and the Simple Mail Transfer Protocol (SMTP).

IIS 6.0 takes advantage of the latest Web standards like ASP.NET, XML, and Simple Object Access Protocol (SOAP) for the development, implementation, and management of Web applications. IIS 6.0 includes new features that are designed to help organizations, IT professionals, and Web administrators achieve their goals of performance, reliability, scalability, and security for potentially thousands of Web sites hosted on a single server or on multiple servers.

# IIS 6.0 Benefits and Features

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 provides the following benefits and features:

- Reliability. IIS 6.0 uses a new request-processing architecture and application isolation environment that enables individual Web applications to function within a self-contained worker process. This environment prevents one application or Web site from stopping another, and it reduces the amount of time that administrators spend restarting services to correct application-related problems. The new environment also includes proactive health monitoring for application pools. For more information about application reliability in IIS 6.0, see Ensuring Application Availability.

- Scalability. IIS 6.0 introduces a new kernel-mode driver for Hypertext Transfer Protocol (HTTP) parsing and caching that is specifically tuned to increase Web server throughput and scalability of multiprocessor computers. The result is an increase in the following:

  - The number of Web sites that a single IIS 6.0 server can host

  - The number of concurrently active worker processes

  - The performance for startup and shutdown times for the Web server and for individual Web sites

  - The number of simultaneous requests that a Web server can service

  Also, by configuring the startup and shutdown time limits for worker processes, IIS allocates resources to active Web sites instead of keeping resources on idle requests.

- Security. IIS 6.0 provides significantly improved security over IIS 5.0. For example, to reduce the attack surface of systems, IIS 6.0 is not installed by default on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. After installing these products, administrators must manually install IIS 6.0. When IIS 6.0 is installed, it is locked down by default so that it can serve only static content. By using the Web Service Extensions node in IIS Manager, Web site administrators can enable or disable IIS functionality based on the individual needs of their organization.

  IIS 6.0 includes a variety of security features and technologies to help ensure the integrity of your Web and File Transfer Protocol (FTP) site content, as well as the data that is transmitted through your sites. These security features and technologies include Advanced Digest authentication, improved access control, Secure Sockets Layer (SSL) encryption, centralized certificate storage, and detailed auditing capabilities.

  For more information about IIS security, see Securing Web Sites and Applications and Security in IIS 6.0.

- Manageability. To meet the needs of a diverse set of organizations, IIS 6.0 provides a variety of manageability and administration tools. Administrators can configure an IIS 6.0 server by using IIS Manager, by running administration scripts, or by directly editing the IIS metabase. Administrators can also remotely administer IIS servers and Web sites.

- Enhanced Development. Compared to Windows 2000 Server, Windows Server 2003 offers an improved developer experience with ASP.NET and IIS integration. ASP.NET runs most Active Server Pages (ASP) code while providing greater functionality for building enterprise-class Web applications that can work as a part of the .NET Framework. Use ASP.NET to fully utilize the features of the common language runtime, such as type safety, inheritance, language interoperability, and versioning. IIS 6.0 also offers support for the latest Web standards including XML, SOAP, and Internet Protocol version 6 (IPv6).

  For more information about enhanced development of ASP.NET applications, see Deploying ASP.NET Applications in IIS 6.0.

- Application Compatibility. According to feedback from thousands of customers and independent software vendors (ISVs),

IIS 6.0 is compatible with most of their existing Web applications. Also, to ensure maximum compatibility, you can configure IIS 6.0 to run in IIS 5.0 isolation mode.

- HTTP Compression. To make best use of available bandwidth, IIS 6.0 includes the ability to compress both static content and dynamic content. This results in faster transmission time between compression-enabled browsers and IIS regardless of whether your content is served from local storage or a Universal Naming Convention (UNC) resource. The HTTP compression feature is highly customizable to meet the needs of your organization.

# Internet and Intranet Applications on IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 deployment process can be applied to Web sites and applications that are hosted on the Internet or within the intranet of your organization. Throughout the deployment process, explicit references are made for deployment considerations relating to Web servers facing the Internet or within an intranet.

| **Note** |
| --- |
| The deployment processes and recommendations described in this guide can be used to deploy an Internet or intranet Web server, unless otherwise noted. |

# Determining Application Compatibility with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

One of the primary concerns when deploying IIS 6.0 is whether or not your existing applications are compatible with IIS 6.0. Windows Server 2003 and IIS 6.0 are designed to provide maximum application compatibility. In most cases, existing Web sites and applications will run on IIS 6.0 without modification.

The deployment process in this guide provides prescriptive guidance about how to address known Web site and application compatibility issues. For a more detailed explanation of Web site and application compatibility with IIS 6.0, read the following:

- Preparing to Upgrade when you are performing an upgrade of an existing Web server running an earlier version of IIS.
- Preparing for Migration when you are moving existing Web sites and applications that are hosted on a Web server running an earlier version of IIS to a newly installed Web server running IIS 6.0.
- Preparing for Migration when you are moving existing Web sites and applications that are hosted on a Web server running Apache to a newly installed Web server running IIS 6.0.

See Also

**Other Resources**

Planning and Testing for Application Deployment

# Moving from IIS 5.0 Isolation Mode to Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 can run in one of two distinct modes of operation, which are called application isolation modes. *Application isolation* is the separation of applications by process boundaries that prevent the applications from affecting one another, and it is configured differently for each of the two IIS application isolation modes: IIS 5.0 isolation mode and worker process isolation mode.

Before you begin deployment, review the following:

- Differences between IIS 5.0 isolation mode and worker process isolation mode
- Benefits of moving from IIS 5.0 isolation mode to worker process isolation mode

| **Note** |
| --- |
| This guide assumes that IIS 6.0 is running in worker process isolation mode, unless otherwise noted. |

# Reviewing Application Isolation Modes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Worker process isolation mode uses the redesigned architecture for IIS 6.0. This isolation mode runs all application code in an isolated environment. However, unlike earlier versions of IIS, IIS 6.0 provides isolation without a performance penalty because fewer processor instructions are run when switching from one application pool to another. Worker process isolation mode is compatible with most existing Web sites and applications. Whenever possible, run IIS 6.0 in worker process isolation mode to benefit from the enhanced performance and security in IIS 6.0.

IIS 5.0 isolation mode provides compatibility for applications that depend upon the process behavior and memory model of IIS 5.0. Run IIS in this mode only when a Web site or application cannot run in worker process isolation mode, and run it only until the compatibility issues are resolved.

| **Important** |
| --- |
| IIS 6.0 cannot run both application isolation modes simultaneously on the same server. Therefore, on a single server running IIS 6.0, you cannot run some Web applications in worker process isolation mode and others in IIS 5.0 isolation mode. If you have applications that require separate modes, you must run them on separate servers. |

IIS 6.0 defaults to a different application isolation mode based on the type of deployment you select. For new installations and migrations, IIS is configured to run in worker process isolation mode by default. After you perform an upgrade from an earlier version of IIS, IIS is configured to run in IIS 5.0 isolation mode by default.

Before configuring IIS 6.0 to run in worker process isolation mode, evaluate whether your Web sites and applications are compatible with worker process isolation mode. In most cases, IIS hosts your Web sites and applications in worker process isolation mode without any problems. Nevertheless, determine application compatibility in your lab before deploying your IIS solution into production.

For more information about worker process isolation mode, IIS 5.0 isolation mode, and evaluating Web site and application compatibility with worker process isolation mode, see Determining Application Compatibility with Worker Process Isolation Mode.

| **Note** |
| --- |
| Identifying a complete list of potential incompatibilities that applications can experience with worker process isolation mode is beyond the scope of this guide. Even after following the guidelines in this section, you still need to verify in your lab whether your Web sites and applications are compatible with worker process isolation mode. |

# Benefits of Moving to Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Worker process isolation mode provides higher levels of security and availability for Web sites and applications than IIS 5.0 isolation mode. Therefore, it is recommended that you configure IIS 6.0 to run in worker process isolation mode.

In IIS 5.0, applications can be pooled together out-of-process, but in only one application pool. In IIS 6.0, worker process isolation mode supports multiple application pools, where each application pool can have a different configuration, such as a unique recycling configuration. Therefore, you can prevent a single Web site or application that periodically fails from disrupting other Web sites and applications.

In addition, worker process isolation mode provides the following improvements to IIS:

- Security Enhancements
- Performance and Scaling Enhancements
- Availability Enhancements

# Security Enhancements

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 includes a variety of security features and technologies that help ensure the integrity of your Web site content, and of the data that is transmitted through your sites. The following security enhancement is only available when IIS 6.0 is running in worker process isolation mode.

Default process identity for Web sites and applications set to NetworkService

In IIS 5.0 isolation mode, the default process identity is LocalSystem, which enables access to, and the ability to alter, nearly all of the resources on the Web server.

# Performance and Scaling Enhancements

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Future growth in the utilization of your Web sites and applications requires increased performance and scalability of Web servers. By increasing the speed at which HTTP requests can be processed and by allowing more applications and sites to run on one Web server, the number of Web servers that you need to host a site is reduced. The following are a few of the performance improvements included in worker process isolation mode.

Support for processor affinity for worker processes in an application pool

You can configure all of the worker processes in an application pool to have affinity with specific processors in a multiprocessor or server. Processor affinity allows the worker processes to take advantage of more frequent processor caching (Level 1 or Level 2).

Elimination of inactive worker processes and reclamation of unused resources

You can configure application pools to have worker processes request a shutdown if they are idle for a certain amount of time. This can free unused resources for other active worker processes. New worker processes are then started only when they are needed.

Distributing client connections across multiple worker processes

You can configure an application pool to have more than one worker process servicing client connections, also known as a *Web garden*. Because there are multiple worker processes, the incoming client connections are distributed across the worker processes and throughput is not constrained by a single worker process.

Ability to isolate Web sites and applications from each other

You can isolate Web sites and applications without incurring a performance penalty.

# Availability Enhancements

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because worker process boundaries isolate the applications in an application pool from the applications in other application pools, if an application fails, it does not affect the availability of other applications running on the server. Deploying applications in application pools is a primary advantage of running IIS 6.0 in worker process isolation mode.

Reduced number of server restarts that are required when administering Web sites and applications

Many of the common operation tasks do not force the restart of the server or the Web service. These tasks, such as upgrading site content or components, debugging Web applications, or dealing with faulty Web applications, can be performed without affecting service to other Web sites or applications on the server.

A fault-tolerant request-processing model for Web sites and applications

In IIS 5.0 isolation mode, each Web site or application has only one worker process. However, in worker process isolation mode, you can create a *Web garden* by configuring a number of worker processes to share the processing. The benefit of a Web garden is that if one worker process stops responding, other worker processes are available to accept and process requests.

Isolation of failed worker processes from healthy worker processes

In worker process isolation mode, IIS can determine that a worker process has failed and start a new worker process. To minimize the interruption of service, new requests are queued until the new worker process is active. In the case where a worker process has not yet failed but is considered unhealthy, IIS starts a new worker process. When the new worker process is active, IIS shuts down the unhealthy worker process. After IIS creates the new worker process, the failed worker process can be separated, or *orphaned*, from the application pool. The advantage of orphaning a worker process rather than terminating it is that debugging can be performed on the orphaned worker process.

Health monitoring of Web sites and applications

In worker process isolation mode, you can configure an application pool to monitor not only the health of the entire application pool, but also individual worker processes servicing the application pool. Monitoring the health of a worker process allows IIS to detect that a worker process is unable to serve requests and to take corrective action, such as recycling the failed worker process.

In addition, worker process isolation supports other responses when a failed worker process or application pool is detected. For example, IIS can attach a debugger to an orphaned worker process or notify an administrator that an application pool has failed due to rapid-fail protection.

Prevention of Web sites or applications that fail quickly from consuming system resources

In some instances, availability can be affected by Web sites and applications that fail very quickly, are automatically restarted, and then fail quickly again. The endless cycle of failure and restarting can consume system resources, causing other Web sites and applications to experience denial of services because of system resource shortages.

Worker process isolation mode includes *rapid-fail protection* that stops an application pool when too many of the worker processes assigned to an application pool are found to be unhealthy within a specified period of time.

Automatic restart of individual Web sites and applications

Some Web sites and applications have memory leaks, are poorly coded, or have other unidentified problems. In IIS 5.0 isolation mode, these applications can force you to restart the entire Web server. The recycling feature in worker process isolation mode can periodically restart the worker processes in an application pool without affecting service availability. Worker processes can be scheduled to restart based on several options, such as elapsed time or the number of requests served.

# Deploying ASP.NET Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Microsoft® Windows® Server 2003 includes support for ASP.NET applications and the Microsoft .NET Framework version 1.1 with the operating system installation. This chapter describes how to deploy ASP.NET applications on a newly installed server running Internet Information Services (IIS) 6.0. Version 1.1 of the .NET Framework is installed with Windows Server 2003. Most ASP.NET applications run without modification on version 1.1 of the .NET Framework.

You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI.

**In This Section**

Overview of Deploying ASP.NET Applications in IIS 6.0

Deploying the Web Server

Installing ASP.NET Applications

Completing the ASP.NET Application Deployment

Additional Resources_IIS_SP1_Dep

# Overview of Deploying ASP.NET Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP.NET is a unified Web application platform that provides services to help you build and deploy enterprise-class Web applications and XML-based Web services. ASP.NET is supported on the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server2003, Enterprise Edition; Windows® Server2003, Datacenter Edition; and Windows® Server2003, Web Edition operating systems. ASP.NET is installed with the Microsoft .NET Framework version 1.1 as a part of Windows Server 2003. However, to run ASP.NET applications, you must also install IIS 6.0.

| **Important** |
| --- |
| ASP.NET is not available on the following operating systems: Microsoft® Windows® XP 64-Bit Edition; the 64-bit version of Windows® Server 2003, Enterprise Edition; and the 64-bit version of Windows® Server 2003, Datacenter Edition. For more information, see "Features unavailable on 64-bit versions of the Windows Server 2003 family" in Help and Support Center for Microsoft® Windows® Server 2003. |

The deployment process presented in this section describes how to deploy ASP.NET applications on a newly installed IIS 6.0 Web server. Before you begin this process, complete the following steps:

- Install Windows Server 2003, which includes version 1.1 of the .NET Framework, with the default options.
- Install IIS 6.0 with the default settings in **Add or Remove Programs** in Control Panel.

If you need to install ASP.NET applications that were written for IIS 5.0 or version 1.0 of the .NET Framework on a new Web server, see Migrating IIS Web Sites to IIS 6.0. If you want to upgrade a Web server running IIS 5.0 that is hosting existing ASP.NET applications, see Upgrading an IIS Server to IIS 6.0.

When you configure IIS 6.0 to run in IIS 5.0 isolation mode, the settings in the **<processModel>** section of the Machine.config file are configured in the same way as they were in IIS 5.0 — in the Machine.config or Web.config files. For more information about configuring ASP.NET applications when IIS 6.0 is configured to run in IIS 5.0 isolation mode, see Overview of ASP.NET Configuration.

Upon completing the process described in this section, you will have a Web server running IIS 6.0 and hosting your ASP.NET applications. However, you can further configure the Web server to improve the security and availability of your ASP.NET applications. For more information about configuring your Web server to improve the security and availability of your ASP.NET applications, see Securing Web Sites and Applications and Ensuring Application Availability.

| **Note** |
| --- |
| The configuration settings discussed in this section are appropriate for Web sites and applications that are hosted on Web servers on an intranet and the Internet, unless specifically noted. |

# Process for Deploying ASP.NET Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The process for deploying new ASP.NET applications on a newly installed Web server requires no understanding of earlier versions of IIS or the .NET Framework. All the ASP.NET configuration sections in the Machine.config and Web.config files are configured the same way in IIS 6.0, except for the **<processModel>** section of the Machine.config file. When IIS 6.0 is configured to run in worker process isolation mode, some of the attributes in the **<processModel>** section of the Machine.config file are now in equivalent IIS 6.0 metabase properties. For more information about how to migrate attributes in the Machine.config file to their equivalent IIS 6.0 metabase property settings, see Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings.

In addition, if your ASP.NET applications need to retain session state, you must configure IIS 6.0 to use the appropriate ASP.NET application session state method. Depending on the method you select, you might need to configure the ASP.NET state service or Microsoft SQL Server™ to act as the repository for centralized state storage.

The process for deploying ASP.NET applications in IIS 6.0 is shown in Figure 2.1.

**Figure 2.1 Deploying ASP.NET Applications in IIS 6.0**




| **Note** |
| --- |
| Before deploying your ASP.NET applications on a production server, perform the process outlined in this section on a test server that is configured identically to your production server. |

The following quick-start guide provides a detailed overview of the process for deploying ASP.NET applications in IIS 6.0. You can use this guide to help identify the steps of the ASP.NET application deployment process that you need additional information to complete, and to skip the steps with which you are already familiar. In addition, all of the procedures that are required to complete the ASP.NET application deployment process are documented in Appendix A: IIS Deployment Procedures.

**Deploy the Web Server**

1. Install Windows Server 2003.
2. Install and configure IIS 6.0.
3. Enable ASP.NET in the Web service extensions list.

**Install ASP.NET Applications**

1. Create Web sites and virtual directories for each ASP.NET application by doing the following:
   - Create Web sites and home directories.
   - Create virtual directories.
2. Copy ASP.NET application content to the Web server.

3. Enable common storage for ASP.NET session state by completing the following steps:

    - Select the method for maintaining and storing ASP.NET session state.
    - If you decided to maintain session state with the ASP.NET state service, configure out-of-process session state with the ASP.NET state service.
    - If you decided to maintain session state with SQL Server, configure out-of-process session state with SQL Server.
    - Configure encryption and validation keys.
    - Configure ASP.NET to use the appropriate session state.
    - Secure the ASP.NET session state connection string.

**Complete the ASP.NET Application Deployment**

1. Ensure the security and availability of your ASP.NET applications.
2. Verify that the ASP.NET applications were deployed successfully.
3. Back up the Web server.
4. Enable client access to your ASP.NET applications.

# Deploying the Web Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You must install the Web server before you can install your ASP.NET applications. In addition to installing Windows Server 2003, you must install and configure IIS 6.0 on the Web server. You must also enable ASP.NET so that the Web server can run ASP.NET applications.

| **Note** |
| --- |
| You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

Figure 2.2 illustrates the process for deploying the Web server.

**Figure 2.2 Deploying the Web Server**



# Installing Windows Server 2003_IIS_SP1_Dep

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The deployment process presented here assumes that you install Windows Server 2003 with the default options. If you use other methods for installing and configuring Windows Server 2003, such as unattended setup, your configuration settings might be different.

| **Note** |
| --- |
| When you complete the installation of Windows Server 2003, Manage Your Server automatically starts. The deployment process assumes that you quit Manage Your Server, and then further configure the Web server in **Add or Remove Programs**in Control Panel. |

For more information about the essential components and services you should enable in Windows Server 2003, see Enabling Only Essential Windows Server 2003 Components and Services.

# Installing and Configuring IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because IIS 6.0 is not installed during the default installation of Windows Server 2003, the next step in deploying the Web server is to install and configure IIS 6.0. The deployment process presented here assumes that you install IIS 6.0 with the default options in **Add or Remove Programs** in Control Panel. If you use other methods for installing and configuring Windows Server 2003, such as Manage Your Server, the default configuration settings might be different.

| **Note** |
| --- |
| You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

1. Install IIS 6.0 with only the essential components and services.

   As with installing Windows Server 2003, the primary concern when installing and configuring IIS 6.0 is to ensure that the security of the Web server is maintained. Enabling unnecessary components and services increases the attack surface of the Web server. You can help ensure that the Web server is secure by enabling only the essential components and services in IIS 6.0.

   For more information about how to install IIS 6.0, see Install IIS 6.0. For more information about the IIS 6.0 components and services that you should enable, see Enabling Only Essential IIS Components and Services

2. If you want to manage the Web site content by using Microsoft® FrontPage®, install FrontPage 2002 Server Extensions from Microsoft on the Web server.

   For information about how to enable FrontPage Server Extensions after installing them on the Web server, see Configure Web Service Extensions.

# Enabling ASP.NET in the Web Service Extensions List

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you install IIS 6.0, you need to enable ASP.NET. You can enable ASP.NET in **Add or Remove Windows Components**, which is accessible from **Add or Remove Programs** in Control Panel. When you enable ASP.NET by using this method, ASP.NET is also enabled in the Web service extensions list. If you enabled ASP.NET in this way, then you can continue to the next step in the deployment process. To continue to the next step in the deployment process, see Installing ASP.NET Applications.

For more information about enabling ASP.NET, see Enable ASP.NET in **Add or Remove Programs.**

**ASP.NET is not Enabled**

ASP.NET might not be enabled in the Web service extensions list if either of the following is true:

- You installed a version of the .NET Framework and ASP.NET (other than version 1.1) from a Web download or as part of an application such as the Microsoft Visual Studio® .NET development tool.

- You disabled ASP.NET in the Web service extensions list because you were not running ASP.NET applications on an existing Web server.

If ASP.NET is not already enabled, view the Web service extensions list in IIS Manager and configure the status of the ASP.NET v1.1.4322 Web service extension to **Allowed**. For more information about enabling ASP.NET in the Web service extensions list see Configure Web Service Extensions.

# Installing ASP.NET Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After the Web server is deployed, you can install your ASP.NET applications. First, you must create a Web site and virtual directories for each ASP.NET application. Then you need to install each ASP.NET application in the corresponding Web site and virtual directory.

| **Note** |
| --- |
| You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

When there are provisioning or setup scripts for your ASP.NET applications, use these scripts to install the ASP.NET applications on the Web server. Because the provisioning and setup scripts create the Web sites and virtual directories while installing ASP.NET applications, you do not need to perform any manual steps to install the ASP.NET applications. In this case, run the provisioning or setup scripts to install and configure the Web sites and applications, and then continue to the next step in the application deployment process. For information on the next step in the process, see Enabling Common Storage for ASP.NET Session State.

Figure 2.3 illustrates the process for installing your ASP.NET applications.

**Figure 2.3 Installing ASP.NET Applications**



# Creating Web Sites and Virtual Directories for each ASP.NET Application

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For each ASP.NET application, you must create a virtual directory in a new or existing Web site. Later in the installation process, you will install your ASP.NET applications into their corresponding Web sites and virtual directories.

Create the Web sites and virtual directories for your ASP.NET applications by completing the following steps:

1. Create Web sites and home directories.

2. Create virtual directories.

# Creating Web Sites and Home Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Each Web site must have one home directory. The home directory is the central location for your published Web pages. It contains a home page or index file that serves as a portal to other pages in your Web site. The home directory is mapped to the domain name of the Web site or to the name of the Web server.

Create a Web site and home directory for an ASP.NET application by completing the following steps:

1. Create the folder that will be the home directory for the Web site on the Web server.

   The folder that is the home directory of the Web site contains all of the content and subdirectories for the Web site. The folder can be created on the same computer as the Web server or on a Universal Naming Convention (UNC)–shared folder on a separate server. At a minimum, create the folder on the following:

   - An NTFS file system partition, which helps ensure proper security.
   - A disk volume other than the system volume, which reduces the potential of an attack on a Web site bringing down the entire Web server and improves performance.
   - In a location that will not require requests for Web site content to contain /bin in the requested URL. As a security measure, ASP.NET returns a 404 error for all requests containing /bin in the requested URL.

   For more information about securing Web sites and applications, see Securing Web Sites and Applications. For more information about creating directories for Web sites, see Create a Web Site.

2. Create the Web site on the server.

   For more information about how to create a Web site, see Create a Web Site.

3. If the Web site is FrontPage extended, then configure the Web site on the Web server to be FrontPage extended.

   For more information about how to configure a Web site to be FrontPage extended, see Configure a Web Site to be FrontPage Extended.

# Creating Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A *virtual directory* is a folder name, used in an address, which corresponds to a physical directory on the Web server or a Universal Naming Convention (UNC) location. This is also sometimes referred to as *URL mapping*. Virtual directories are used to publish Web content from any folder that is not contained in the home directory of the Web site. When clients access content in a virtual directory, the content appears to be in a subdirectory of the home directory, even though it is not.

For security reasons, you might want to move the Web site content to a different disk volume during the application deployment process. You can move the content to another disk volume on the Web server or to a shared folder on a separate server. You can use virtual directories to specify the UNC name for the location where the content is placed, and provide a user name and password for access rights.

For each virtual directory required by the ASP.NET application, create a corresponding virtual directory on the Web server by completing the following steps:

1. Create the folder on the Web server to contain the virtual directory content.

   Ensure that you create the folder in a secure manner that does not compromise the security of the Web server.

   For more information about securing virtual directories, see Preventing Unauthorized Access to Web Sites and Applications.

2. Create the virtual directory under the appropriate Web site on the server.

   For more information about creating virtual directories, see Create a Virtual Directory in IIS 6.0.

# Copying ASP.NET Application Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When no installation program or provisioning scripts exist for your ASP.NET application, you can copy the content of the ASP.NET application to the corresponding Web site and virtual directories that you created on the Web server.

You can copy the ASP.NET application content to the Web server by using one of the following methods:

- Run the **Xcopy** command to copy ASP.NET application content to the Web server on an intranet or internal network.
- Use Microsoft Windows Explorer to copy ASP.NET application content to the Web server on an intranet or internal network.
- Use the **Copy Project** command in Visual Studio .NET to copy ASP.NET application content to the Web server on an intranet or internal network, if the application has been developed by using Visual Studio .NET.

| **Note** |
| --- |
| FrontPage Server Extensions must be installed on the Web server to use the **Copy Project** command. |

- Use the **Publish Web** command in FrontPage to copy ASP.NET application content to the Web server on an intranet or over the Internet, if the Web site that contains the application has been developed using FrontPage.

  For more information about how to publish ASP.NET application content using FrontPage, see Publish Web Site Content with FrontPage in Appendix A: IIS Deployment Procedures.

# Enabling Common Storage for ASP.NET Session State

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP.NET session state lets you share client session data across all of the Web servers in a Web farm or across different worker processes or worker process instances on a single Web server. Clients can access different servers in the Web farm across multiple requests and still have full access to session data.

You can enable common storage for ASP.NET session state by performing the following steps:

1. Select the method for maintaining and storing ASP.NET session state.
2. If you decided to maintain session state with the ASP.NET state service, configure out-of-process session state with the ASP.NET state service.
3. If you decided to maintain session state with SQL Server, configure out-of-process session state with SQL Server.
4. Configure the encryption and validation keys.
5. Configure ASP.NET to use the session state method that you selected in Step 1.
6. Secure the ASP.NET session state connection string in the registry.

# Selecting the Method for Maintaining and Storing ASP.NET Session State

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure the method used by ASP.NET to maintain session state. ASP.NET supports the following methods for maintaining session state:

- **In process**. The process for maintaining session state runs in the same worker process as the ASP.NET application and approximates how Active Server Pages (ASP) applications maintain session state. If the worker process for the ASP.NET application recycles, all session state information is lost.

- **Out of process**. The process for maintaining session state runs in a different worker process, either on the same Web server or another server, than the worker process for the ASP.NET application. If the worker process for the ASP.NET application recycles, session state information is retained.

You can configure ASP.NET session state settings by modifying the **mode** attribute in the <**sessionState**> section of the Machine.config file for all applications or in the Web.config file for specific applications. The session state settings in the IIS 6.0 metabase only apply to ASP applications. ASP.NET and ASP applications cannot share session state.

Select one of the following methods for storing.ASP.NET session state:

- **In process**, local to the Web server. ASP.NET session state is managed by and stored on the local Web server. This is the default mode for ASP.NET session state management and it approximates how ASP applications manage session state.

  This method allows the session state to be stored across worker process threads in a Web garden. However, the in-process mode does not allow session state to be store across servers in a Web farm. For more information about Web gardens, see "Configuring Web Gardens" in Ensuring Application Availability.

  To use this method to store ASP.NET session state, set the session state **mode** attribute to **InProc**.

- **Out of process**, by using the ASP.NET state service. The ASP.NET state service (aspnet_state.exe) runs as a service on Windows Server 2003. You can run the ASP.NET state service local to a Web server to support Web gardens, or on a separate server to support Web farms.

  To use this method to store ASP.NET session state, set the session state **mode** attribute to **StateServer**.

- **Out of process**, by using a computer running SQL Server. ASP.NET session state can be managed by and stored in a database on a computer running Microsoft SQL Server. Like the ASP.NET state service, this out-of-process mode provides support for Web gardens or Web farms.

  To use this method to store ASP.NET session state, set the session state **mode** attribute to **SQLServer**.

Table 2.1 compares the methods for maintaining ASP.NET session state, listing the advantages and disadvantages of each method.

**Table 2.1 Comparison of Methods for Maintaining ASP.NET Session State**

| Method | Advantages | Disadvantages |
|---|---|---|

| | | |
|---|---|---|
| In-process | • Requires no additional computers.<br>• Provides faster access to session state than the out-of-process modes because state is maintained in memory and is in-process. | • Does not provide centralized storage of session state for Web farms.<br>• Provides no redundancy in the event of a Web server failure.<br>• Does not survive application restarts or process recycles. |
| Out-of-process with the ASP.NET state service running on the local Web server | • Provides centralized storage of session state to support Web gardens.<br>• Requires no additional computers. | • Does not have a mechanism for failover or partitioning.<br>• Out-of-process state requires serializing data and does not perform as well as in-process state.<br>• Provides slower access to session state than in-process mode. |
| Out-of-process with the ASP.NET state service running on a separate server | • Provides centralized storage of session state to support Web gardens or Web farms. | • Requires an additional computer to store the session state.<br>• Does not have a mechanism for failover or partitioning.<br>• Does not support remote authentication, so administrators must control access to the computer with Internet Protocol security (IPsec) and/or firewall rules.<br>• Out-of-process state requires serializing data and does not perform as well as in-process state.<br>• Provides slower access to session state than in-process mode. |
| Out-of-process with session state stored on a separate computer running SQL Server | • Provides centralized storage of session state to support Web gardens or Web farms.<br>• Can store session state on a computer running SQL Server that is used for other purposes and share the costs of the server.<br>• Can cluster the computer running SQL Server to provide higher availability. | • Requires an additional computer to store the session state.<br>• Out-of-process state requires serializing data and does not perform as well as in-process state.<br>• Provides slower access to session state than in-process mode. |

**Note**

Before you can use an out-of-process state method for managing and storing session state, you must ensure that the objects defined by and used by your ASP.NET application are *serializable*. Making an object serializable is usually a matter of adding the object class with the **Serializable** attribute. Consult with the ASP.NET application developer to ensure that all of the objects stored in session state by your ASP.NET applications are serializable.

# Configuring Out-of-Process Session State with the ASP.NET State Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you decide to manage session state by using the ASP.NET state service, you must determine whether you are going to maintain session state for a Web garden or a Web farm. Then you need to ensure that the ASP.NET state service (Aspnet_state.exe) is running and, that it is configured to start automatically.

Configure out-of-process session state with the ASP.NET state service by completing one of the following tasks:

- If you are configuring session state for a Web garden, then configure the ASP.NET state service on the local Web server.

  For more information about configuring the ASP.NET state service on the local Web server, see Configure the State Service on the ASP.NET State Server.

- If you are configuring session state for a Web farm, then complete the following steps:

  1. Install Windows Server 2003 on separate computer that will run the ASP.NET state service.

     The ASP.NET state service is installed as part of the .NET Framework with Windows Server 2003. Install Windows Server 2003 by using the process described in Installing Windows Server 2003_IIS_SP1_Dep earlier in this section.

  2. Configure the ASP.NET state service on the newly installed Web server.

     For more information about configuring the ASP.NET state service on the new server, see Configure the State Service on the ASP.NET State Server.

# Configuring Out-of-Process Session State with SQL Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you decide to store the session state in SQL Server, you need to create the session state database on the computer running SQL Server that is used by the ASP.NET state service. You can create the database by running the InstallSqlState.sql script on the computer running Microsoft SQL Server that is going to be storing the session data. This script creates a database called ASPState, which includes several stored procedures and adds the ASPStateTempApplications and ASPStateTempSessions tables to the TempDB database.

For more information about creating the SQL Server database for storing ASP.NET session state, see Create A SQL Server Database for Storing ASP.NET Session State.

# Configuring Encryption and Validation Keys

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP.NET uses a key to help protect data so that session state data is only accessible from the Web server that created the data. In a Web garden, all of the worker processes use the same Machine.config or Web.config file, so no additional configuration is necessary. If you are configuring ASP.NET session state for a Web garden, continue to the next step in the deployment process. To continue to the next step in the ASP.NET application deployment process, see Configuring ASP.NET Applications to Use the Appropriate Session State.

In a Web farm, each Web server in the farm has a separate Machine.config or Web.config file. As a result, you need to manually configure each Web server in the farm to share the same encryption and validation keys so that they can share session state data. This allows one server to decrypt the session state data created by another server in the Web farm. You can configure the encryption and validation keys in the <**machineKey**> section of the Machine.config file.

| **Tip** |
| --- |
| When you want all of the ASP.NET applications on a Web server to use the same encryption and validation keys, make the modifications in the Machine.config file. When you want to customize the encryption and validation keys for each ASP.NET application, modify the corresponding Web.config file for the application. |

For each server in the Web farm, identically configure the values of the following attributes in the Machine.config or Web.config file:

- The **validationKey** attribute. This attribute contains the key that is used to validate that data tampering has not occurred. The validation algorithm, which is specified in the **validation** attribute, validates the data by using the key in the **validationKey** attribute. The **validationKey** attribute can range from 40 to 128 hexadecimal characters in length. The default value for the **validationKey** attribute is auto-generated. Configure the value for the **validationKey** attribute to be the same value for all of the servers in a Web farm.

- The **validation** attribute. This attribute is used to configure the validation algorithm used to verify the data. The validation algorithms that you can select include Message Digest 5 (MD5), Secure Hash Algorithm 1 (SHA1), or triple DES (3DES). The Web server validates the data with the key in the **validationKey** attribute and the algorithm specified by the **validation** attribute. Configure the value of the **validation** attribute to be the same value for all of the servers in a Web farm.

- The **decryptionKey** attribute. This attribute contains the key that is used to encrypt data. Valid values are 16 or 48 hexadecimal characters, which corresponds to the Data Encryption Standard (DES) or 3DES algorithm, respectively. The default value for the **decryptionKey** attribute is auto-generated. Configure the value of the **decryptionKey** attribute to be the same value for all of the servers in a Web farm.

For an example that describes how to configure the encryption and validation keys for Microsoft Content Management Server 2002, see Best Practices for Authentication for Web Farms on MSDN. You can use this example to assist you in configuring the encryption and validation keys for your ASP.NET application.

# Configuring ASP.NET Applications to Use the Appropriate Session State

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure ASP.NET session state persistence in the <**sessionState**> section of the Machine.config file for all of the ASP.NET applications on the Web server, or you can configure it in the Web.config file for each ASP.NET application. If the <**sessionState**> section of the Web.config file for an ASP.NET application is empty, the session state configuration for the application is inherited from the <**sessionState**> section of the Machine.config file.

If the <**sessionState**> section does not exist in the Machine.config file or the Web.config file, the following is the default session state behavior:

- The session time-out value for an ASP.NET session state is 20 minutes.
- The session state is maintained within all of the applications running in the same application pool.
- The session state is only maintained on the local Web server (where the session state **mode** attribute is set to **InProc**), not across multiple servers in a Web farm.

In addition to the configuration settings that are specific to the mode you selected for maintaining session state, the following attributes need to be configured:

- The **cookieless** attribute. This attribute determines whether or not the session identifier is transferred between the Web server and the client by using cookies. When the **cookieless** attribute is set to **True**, cookies are not used to convey session identifiers. When the **cookieless** attribute is set to **False**, cookies are used to convey session identifiers. The default setting is **False**.
- The **timeout** attribute. This attribute specifies the number of minutes that a session is considered valid. When the time specified in the **timeout** attribute expires, the session is disconnected. The default time-out limit is 20 minutes.

Configure the session state settings in the <**sessionState**> section for the mode that you selected.

**Session state is maintained in-process**
To maintain session state in-process, you can either delete the <**sessionState**> section or configure the <**sessionState**> section in the Machine.config file or the Web.config file.

The following is an example of the configuration when maintaining session state in-process:

```
<configuration>
  <system.web>
     <sessionState

         cookieless="true"
         timeout="20"
     </sessionState>
  </system.web>
</configuration>
```

**Session state is maintained out-of-process with the ASP.NET state service**
To configure ASP.NET to maintain session state with the ASP.NET state service, modify the following attributes in addition to the **mode** attribute:

- The **stateConnectionString** attribute. This attribute specifies the IP address and port number where the ASP.NET state service is running. The format for this attribute is **tcpip=***server*:*port*, where *server* is the IP address or host name of the server, and *port* is the TCP port number that the ASP.NET state service is configured to use. The default port number is 42424.
- The **stateNetworkTimeout** attribute. This optional attribute specifies the length of time, in seconds, that the TCP/IP

network connection between the Web server and the server running the ASP.NET state service can be idle before the session is abandoned. The default time-out limit is 10 seconds.

The following is an example of the configuration when maintaining session state out-of-process with the ASP.NET state service:

```
<configuration>
  <system.web>
     <sessionState

         cookieless="true"
         timeout="20"

     </sessionState>
  </system.web>
</configuration>
```

**Session state is maintained out-of-process with a computer running Microsoft SQL Server**

To configure ASP.NET to maintain session state with a computer running SQL Server, modify the **sqlConnectionString**attribute in addition to the **mode** attribute. The **sqlConnectionString**attribute specifies the ODBC data connection string used for establishing the connection to the computer running SQL Server.

The format for the **sqlConnectionString**attribute is **data source=***odbc_connection_string*, where *odbc_connection_string* is any valid ODBC data connection string that is valid for the computer running SQL Server. For more information about creating ODBC data connection strings for Microsoft SQL Server, see How to allocate handles and connect to SQL Server (ODBC) on MSDN.

The following is an example of the configuration when maintaining session state out-of-process with a computer running SQL Server:

```
<configuration>
  <system.web>
     <sessionState

         cookieless="true"
         timeout="20"

     </sessionState>
  </system.web>
</configuration>
```

# Securing the ASP.NET Session-State Connection String

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When using either of the out-of-process methods for maintaining session state — by using the ASP.NET state service or Microsoft SQL Server — the ASP.NET session state connection string is stored in the Machine.config or Web.config file in plain text. You can further secure the ASP.NET session state connection strings by placing the session state connection strings in the registry. Then the Machine.config or Web.config files are modified to point to the corresponding registry keys.

For more information about configuring ASP.NET to store the session connection strings in the registry, see Knowledge Base article 329290, HOW TO: Use the ASP.NET Utility to Encrypt Credentials and Session State Connection Strings.

| ⚠**Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference on the Microsoft Windows Server 2003 Deployment Kit companion CD. |

# Completing the ASP.NET Application Deployment

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

At this point in the process, your ASP.NET applications are installed on the Web server and the ASP.NET session state settings have been configured on the Web server. Now you need to ensure that the ASP.NET applications are configured to provide the appropriate levels of security and availability for your organizational needs. Then you can verify that the ASP.NET applications have been deployed successfully, capture the current configuration of the Web server, and enable client access to the ASP.NET applications on your Web server. After you complete these last steps, the deployment of your ASP.NET applications is complete.

Figure 2.4 illustrates the process for finalizing the deployment of your ASP.NET applications in IIS 6.0.

**Figure 2.4 Completing the ASP.NET Application Deployment Process**



# Verifying That the ASP.NET Applications Were Deployed Successfully

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before deploying the Web server into a production environment, verify that the ASP.NET applications were deployed successfully by completing the following steps:

1. Review the system log in Windows Server 2003 on the Web server to determine whether any of the ASP.NET applications did not start.

   IIS 6.0 creates entries in the system log when a Web site fails to start for any reason. Search the system log on the Web server to determine whether any errors occurred. For more information about how to troubleshoot Web sites that fail to start, see Troubleshooting IIS 6.0. For information about how to troubleshoot ASP.NET-specific problems, see Troubleshooting an ASP.NET Installation.

2. Perform functional testing of your ASP.NET applications to ensure that they behave as expected.

   You can help eliminate many obvious causes of ASP.NET application failure by reviewing the Windows logs and the application configuration settings. However, the only way to accurately assess the successful deployment of your ASP.NET applications is to perform functional testing. Functional testing is designed to ensure that the ASP.NET applications are functioning correctly for the most common usage scenarios, such as URLs and inputs. This helps ensure that the ASP.NET applications behave as designed, based on typical user interaction.

   Describing the procedures for performing functional testing of your ASP.NET applications is beyond the scope of this section. For more information about the general subject of testing, go to the MSDN Library and then search for "testing."

# Ensuring the Security and Availability of ASP.NET Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

At this point in the ASP.NET application deployment process, your ASP.NET applications are installed on IIS 6.0, and they are configured to run in the default application pool with the default identity of NetworkService. However, you should take further steps to help ensure that the individual ASP.NET applications are secure. In addition, you can help ensure that if any ASP.NET application fails, the remaining ASP.NET applications remain unaffected.

You can help ensure the security and availability of your ASP.NET applications by completing the following steps:

1. Follow the recommendations and process steps described in Securing Web Sites and Applications to help ensure the security of the ASP.NET applications running on your Web server.

2. Follow the recommendations and process steps described in Ensuring Application Availability to help ensure the availability of the ASP.NET applications running on your Web server.

# Backing Up the Web Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you enable client access to the Web server, perform a complete image backup. The purpose of performing this image backup is to provide a point-in-time snapshot of the Web server. If you need to restore the Web server in the event of a failure, you can use this backup to restore the Web server to a known configuration.

| **Important** |
| --- |
| Do not continue unless you have a successful backup of the entire Web server. Otherwise, you can lose Web sites, applications, or data that you deployed to the Web server. |

For more information about how to back up the Web server, see Back Up and Restore the Web Server to a File or Tape.

# Enabling Client Access

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you have deployed the ASP.NET application to the Web server, you are ready to enable client access to the ASP.NET applications. For a period of time that meets your business needs, monitor the client traffic to ensure that clients are successfully accessing the ASP.NET applications on your Web server and that the clients are experiencing expected response times.

Enable client access to the ASP.NET applications on the Web server by completing the following steps:

1. Create the appropriate Domain Name System (DNS) entries for the ASP.NET applications running on the Web server.

   For more information about how to create DNS entries for your applications, see "Managing resource records" in Help and Support Center for Windows Server 2003.

2. Monitor client traffic to determine whether clients are successfully accessing the ASP.NET applications.

   For more information about how to monitor client traffic to ASP.NET applications on the Web server, see Monitor Active Web and FTP Connections.

3. Establish a monitoring period, such as a few hours or a day, to confirm that clients are accessing the ASP.NET applications on the Web server and that they are experiencing response times and application responses that meet or exceed your requirements.

# Securing Web Sites and Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because the day-to-day operations of your organization depend on the mission-critical applications that are running on Internet Information Services (IIS) 6.0 Web servers, your Web sites and applications need the highest possible security. When you install IIS 6.0, it is installed in a highly secure and locked configuration. Depending on your Web sites and applications, you might need to configure IIS to be less restrictive so that your Web sites and applications can operate correctly. Your Web sites and applications might also need increased security configuration to authenticate users or to restrict the Web sites, applications, and data that can be accessed by users.

**In This Section**


Overview of the Securing Web Sites and Applications Process

Reducing the Attack Surface of the Web Server

Preventing Unauthorized Access to Web Sites and Applications

Isolating Web Sites and Applications

Configuring User Authentication

Encrypting Confidential Data Exchanged with Clients

Maintaining Web Site and Application Security

Additional Resources_IIS_SP1_Dep

# Overview of the Securing Web Sites and Applications Process

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To provide comprehensive security for your Web sites and applications, you must ensure that the entire Web server, including each Web site and application that the server hosts, is protected from unauthorized access. Also, you might have to ensure that the Web sites and applications are protected from other Web sites and applications that are hosted on the same server. Finally, you need to initiate practices to help ensure that your Web sites and applications remain secure.

For security reasons, IIS 6.0 is not installed by default on the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems. When you install IIS 6.0, it is locked down — only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. Features such as Active Server Pages (ASP), ASP.NET, Common Gateway Interface (CGI) scripting, FrontPage® 2002 Server Extensions from Microsoft, and Web Distributed Authoring and Versioning (WebDAV) do not work by default. You can serve dynamic content and enable these features in the Web Service Extensions node in IIS Manager.

Before you begin this process, complete the following steps:

- Install Windows Server 2003 with the default options.
- Install IIS 6.0 with the default settings in **Add or Remove Programs** in Control Panel.

If you use other methods for installing and configuring Windows Server 2003, such as unattended setup, or enabling IIS 6.0 by using Manage Your Server, then the default configuration settings might not be identical.

Upon completing the process outlined in this section, you will have a Web server running IIS 6.0 that fulfills your security requirements. However, to maintain the security of your server, you need to implement continuing security practices such as security monitoring, detection, and response. For more information about maintaining Web server security, see Managing a Secure IIS 6.0 Solution.

| **Note** |
| --- |
| The security settings described in this section are appropriate for Web sites and applications that are hosted on Web servers on an intranet and the Internet, unless specifically noted. |

Although not the focus of this section, you can apply many of the security recommendations described in this section to enhance the security of Web servers that have been upgraded from earlier versions of IIS.

# Process for Securing Web Sites and Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To configure security for Web sites and applications that are hosted on a newly installed Web server, you need to follow certain security practices, such as enabling only the *Web service extensions* that you need. Web service extensions provide content and features beyond serving static Web pages. Any dynamic content that is served by the Web server is done by using Web service extensions, such as content and features that are provided by ASP, ASP.NET, or CGI. In addition, each Web site and application might have specific requirements for security settings. Figure 3.1 shows the process for securing your Web sites and applications.

**Figure 3.1 Securing Web Sites and Applications**




Securing the Web sites and applications requires that the Web server as a whole is secure. The process presented in this section assumes that the network infrastructure connecting the Web servers to the clients and to other servers is secure. The security of the network infrastructure is determined by the placement and configuration of the firewalls, routers, and switches in the network infrastructure.

| **Note** |
| --- |
| The process presented in this section includes all of the steps for securing your Web sites and applications in one of many possible sequences. You can complete these steps in the sequence that is recommended in this section or in another sequence. Regardless of the sequence, it is recommended that you evaluate all of the steps in the process. |

In addition to assuming that the network infrastructure is secure, the process presented here assumes that the server is a *dedicated Web server*. A dedicated Web server is a server that is only being used as a Web server and not for other purposes, such as a file server, print server, or database server running Microsoft SQL Server™.

For more information about securing IIS components other than Internet services, such as Simple Mail Transfer Protocol (SMTP) or Network News Transfer Protocol (NNTP), see SMTP Administration or NNTP Administration. For more information about securing other services on a multipurpose server, see Planning a Secure Environment.

| **Tip** |
| --- |
| To secure the Web sites and applications in a Web farm, use the process described in this section to configure security for each server in the Web farm. |

The following quick-start guide provides a detailed overview of how to configure security for IIS 6.0. You can use this guide to help identify the steps of the security process that you need additional information to complete and skip the information with

which you are already familiar. In addition, all of the procedures that are required to complete the security process are documented in Appendix A: IIS Deployment Procedures.

**Reduce the Attack Surface of the Web Server**

- Enable only essential Windows Server 2003 components and services.
- Enable only essential IIS 6.0 components and services.
- Enable only essential Web service extensions.
- Configure Windows Server 2003 security settings.

**Prevent Unauthorized Access to Web Sites and Applications**

- Store content on a dedicated disk volume.
- Set IIS Web site permissions.
- Set IP address and domain name restrictions.
- Set the NTFS file system permissions.

**Isolate Web Sites and Applications**

1. Evaluate the effects of impersonation on application compatibility:
    - Identify the impersonation behavior for ASP applications.
    - Select the impersonation behavior for ASP.NET applications.
2. Configure Web sites and applications for isolation.

**Configure User Authentication**

1. Configure Web site authentication.
    - Select the Web site authentication method.
    - Configure the Web site authentication method.
2. Configure File Transfer Protocol (FTP) site authentication.

**Encrypt Confidential Data Exchanged with Clients**

- Use Secure Sockets Layer (SSL) to encrypt confidential data.
- Use Internet Protocol security (IPsec) or virtual private network (VPN) with remote administration.

**Maintain Web Site and Application Security**

- Obtain and apply current security patches.
- Enable Windows Server 2003 security logs.

- Enable file access auditing for Web site content.
- Configure IIS logs.
- Review security policies, processes, and procedures.

# Reducing the Attack Surface of the Web Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Immediately after installing Windows Server 2003 and IIS 6.0 with the default settings, the Web server is configured to serve only static content. If your Web sites consist of static content and you do not need any of the other IIS components, then the default configuration of IIS minimizes the attack surface of the server. When your Web sites and applications contain dynamic content, or you require one or more of the additional IIS components, you will need to enable additional features. However, you still want to ensure that you minimize the *attack surface* of the Web server. The attack surface of the Web server is the extent to which the server is exposed to a potential attacker.

However, if you reduce the attack surface of the Web server too much, you can eliminate functionality that is required by the Web sites and applications that the server hosts. You need to ensure that only the functionality that is necessary to support your Web sites and applications is enabled on the server. This ensures that the Web sites and applications will run properly on your Web server, but that the attack surface is minimized.

| **Tip** |
| --- |
| In addition to new installations, you can use the information in this section to reduce the attack surface of existing Web servers. |

Figure 3.2 illustrates the process for reducing the attack surface of the Web server.

**Figure 3.2 Reducing the Attack Surface of the Web Server**




Each additional Windows Server 2003 and IIS 6.0 component is configured with the most restrictive possible security that will allow the component to still function. However, in providing any functionality, there is still an opportunity for potential attackers to exploit any weakness of the component.

For example, enabling the Domain Name System (DNS) component in Windows Server 2003 with the default configuration settings would make the server susceptible to any of the standard attacks common to DNS on Windows, UNIX, Linux, or other operating systems. Additional configuration would be required to further secure DNS, such as requiring zones that are integrated with Microsoft Active Directory® directory service.

In addition, if your primary focus is Web server administration, you might not be familiar with DNS-related security attacks. So reducing the attack surface of the server helps eliminate potential attacks that you cannot predict because of your familiarity with other Windows Server 2003 and IIS 6.0 components.

| **Important** |
| --- |

In addition to enabling only essential Windows Server 2003 and IIS6.0 components, ensure that you configure the components to the highest possible security settings. By enabling nonessential components and services, you can increase the attack surface of your server because you have enabled these components and services without further configuring them to the most restrictive security settings.

# Enabling Only Essential Windows Server 2003 Components and Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The attack surface of the Web server is also affected by the other Windows components and services that are enabled in Windows Server 2003. When you install Windows Server 2003 as a dedicated Web server, the default components and services are configured to provide the smallest possible attack surface. In some cases, you might have installed Windows Server 2003 for other purposes, such as a file server, print server, or computer running SQL Server, so you are installing IIS 6.0 on an existing server. In this situation, you need to reevaluate the components and services that are currently running on the Web server to ensure that only the components and services that you need are enabled.

To enable and disable services, change the startup type of the service. You can configure the startup type of the service to one of the following:

- **Automatic .** The service starts automatically when the operating system starts.
- **Manual .** The service can be started by an administrator, a related operating system service, a system device driver, or an action in the user interface that is dependent on the manual service.
- **Disabled .** The service cannot be started automatically or manually; to start a disabled service, you must change the startup type to Automatic or Manual.

Table 3.1 lists the Windows Server 2003 services, as well as the default startup type, the recommended startup type, and comments about the services.

For each of the Windows Server 2003 services that are listed in Table 3.1, complete the following steps:

1. Review the recommended startup type to determine whether you need to change the default startup type.

- Determine, based on the information provided in the comments, if the recommendation applies to your Web server.
- Configure the startup type for the service based on the decisions made in the previous steps.

For more information about how to change the startup type of Windows Server 2003 services, see Configure Windows Server 2003 Services.

**Table 3.1 Recommended Service Startup Types on a Dedicated Web Server**

| Service Name | Default Startup Type | Recommended Startup Type | Comment |
|---|---|---|---|
| Alerter | **Disabled** | No change | Notifies selected users and computers of administrative alerts. |
| Application Layer Gateway Service | **Manual** | No change | Provides support for application-level plug-ins and enables network and protocol connectivity. |
| Application Management | **Manual** | See comment | Provides software installation services for applications that are deployed in **Add or Remove Programs** in Control Panel.<br>On a dedicated Web server, this service can be disabled to prevent unauthorized installation of software. |

| | | | |
|---|---|---|---|
| Automatic Updates | **Automatic** | See comment | Provides the download and installation of critical Windows updates, such as security patches and hotfixes.<br>This service can be disabled when automatic updates are not performed on the Web server. |
| Background Intelligent Transfer Service | **Manual** | See comment | Provides a background file-transfer mechanism and queue management, and it is used by Automatic Update to automatically download programs (such as security patches).<br>This service can be disabled when automatic updates are not performed on the Web server. |
| ClipBook | **Disabled** | See comment | Enables the Clipbook Viewer to create and share data that can be reviewed by remote users. |
| COM+ Event System | **Manual** | No change | Provides automatic distribution of events to COM+ components. |
| COM+ System Application | **Manual** | No change | Manages the configuration and tracking of COM+-based components. |
| Computer Browser | **Automatic** | No change | Maintains the list of computers on the network, and supplies the list to programs that request the list. |
| Cryptographic Services | **Automatic** | No change | Provides three management services: Catalog Database Service, which confirms the signatures of Windows files; Protected Root Service, which adds and removes Trusted Root Certification Authority certificates from the Web server; and Key Service, which helps in enrolling certificates. |
| DHCP Client | **Automatic** | No change | Required to automatically obtain IP configuration and to dynamically update records in the Domain Name System (DNS). |
| Distributed File System | **Automatic** | Disable | Manages logical volumes that are distributed across a local area network (LAN) or wide area network (WAN).<br>On a dedicated Web server, disable Distributed File System (DFS). |
| Distributed Link Tracking Client | **Automatic** | Disabled | Maintains links between NTFS V5 file system files within the Web server and other servers in the domain.<br>On a dedicated Web server, disable Distributed Link Tracking. |
| Distributed Link Tracking Server | **Manual** | Disabled | Tracks information about files that are moved between NTFS V5 volumes throughout a domain.<br>On a dedicated Web server, disable Distributed Link Tracking. |
| Distributed Transaction Coordinator | **Automatic** | No Change | Coordinates transactions that span multiple resource managers, such as databases, message queues, and file systems. |
| DNS Client | **Automatic** | No change | Allows resolution of DNS names. |
| Error Reporting Service | **Automatic** | See comment | Collects, stores, and reports unexpected application crashes to Microsoft. If this service is stopped, then Error Reporting will occur only for kernel faults.<br>On a dedicated Web server, disable Error Reporting Service. |

| | | | |
|---|---|---|---|
| Event Log | **Automatic** | No change | Writes event log messages that are issued by Windows-based programs and components to the log files. |
| Fax Service | **Manual** | Disabled | Provides the ability to send and receive faxes through fax resources that are available on the Web server and network.<br>On a dedicated Web server, this service can be disabled because sending and receiving faxes is not a typical function of a Web Server. |
| File Replication Service | **Manual** | No change | Enables files to be automatically copied and maintained simultaneously on multiple servers. |
| Help and Support | **Automatic** | No change | Enables Help and Support Center to run on the Web server. |
| HTTP SSL | **Manual** | No change | Implements the Secure Hypertext Transfer Protocol (HTTPS) for the HTTP service by using SSL. HTTP.sys automatically starts this service when any Web sites require SSL. |
| Human Interface Device Access | **Disabled** | No change | Enables generic input to Human Interface Devices (HIDs), which activates and maintains the use of predefined hot buttons on keyboards, remote controls, and other multimedia devices. |
| IMAPI CD-Burning COM Service | **Disabled** | No change | Manages CD recording by using the Image Mastering API (IMAPI). |
| Indexing Service | **Manual** | See comment | Indexes content and properties of files on the Web server to provide rapid access to the file through a flexible query language.<br>On a dedicated Web server, disable this service unless Web sites or applications specifically leverage the Indexing Service for searching site content. |
| Internet Connection Firewall (ICF)/Internet Connection Sharing (ICS) | **Disabled** | No change | Provides network address translation (NAT), addressing and name resolution, and intrusion detection when connected through a dial-up or broadband connection.<br>On a dedicated Web server, disable to prevent inadvertent enabling of NAT, which would prevent the Web server from communicating with the remainder of the network. |
| Intersite Messaging | **Disabled** | No changes | Required by DFS. |
| IPSec Services | **Automatic** | No change | Provides management and coordination of Internet Protocol security (IPsec) policies with the IPsec driver. |
| Kerberos Key Distribution enter | **Disabled** | No change | Provides the ability for users to log on using the Kerberos V5 authentication protocol. |
| License Logging Service | **Disabled** | No change | Monitors and records client access licensing for portions of the operating system, such as IIS, Terminal Services, and file and print sharing, and for products that are not a part of the operating system, such as Microsoft SQL Server or Microsoft Exchange Server.<br>On a dedicated Web server, this service can be disabled. |

| | | | |
|---|---|---|---|
| Logical Disk Manager | **Automatic** | No change | Required to ensure that dynamic disk information is up to date. |
| Logical Disk Manager Administrative Service | **Manual** | No change | Required to perform disk administration. |
| Messenger | **Disabled** | No change | Transmits net sends and Alerter service messages between clients and servers. |
| Microsoft Software Shadow Copy | **Manual** | No change | Manages software-based volume shadow copies taken by the Volume Shadow Copy service.<br>On a dedicated Web server, this service can be disabled when volume shadow copies are not used. |
| Net Logon | **Manual** | No change | Maintains a secure channel between the domain controller, other domain controllers, member servers, and workstations in the same domain and trusted domains. |
| NetMeeting Remote Desktop Sharing | **Manual** | Disabled | Eliminates potential security threats by allowing domain-controller remote administration through NetMeeting. |
| Network Connections | **Manual** | No change | Manages objects in the Network Connections directory. |
| Network Dynamic Data Exchange (DDE) | **Disabled** | No change | Provides network transport and security for Dynamic Data Exchange for programs running on the Web server.<br>This service can be disabled when no DDE applications are running locally on the Web server. |
| Network DDE Distributed Share Database Manager (DSDM) | **Disabled** | No change | Used by Network DDE. This service can be disabled when Network DDE is disabled. |
| Network Location Awareness (NLA) | **Manual** | No change | Collects and stores network configuration and location information, and notifies applications when this information changes. |
| NTLM Security Support Provider | **Manual** | No change | Provides security to RPC programs that use transports other than named pipes, and enables users to log on using the NTLM authentication protocol. |
| Performance Logs and Alerts | **Manual** | See comment | Collects performance data for the domain controller, writes the data to a log, or generates alerts.<br>This service can be set to automatic when you want to log performance data or generate alerts without an administrator being logged on. |
| Plug and Play | **Automatic** | No change | Required to automatically recognize and adapt to changes in the Web server hardware with little or no user input. |

| | | | |
|---|---|---|---|
| Portable Media Serial Number Service | **Manual** | No change | Retrieves the serial number of any portable media player that is connected to the computer. |
| Print Spooler | **Automatic** | See comment | Manages all local and network print queues and controls all print jobs.<br>On a dedicated Web server, this service can be disabled when no printing is required. |
| Protected Storage | **Automatic** | No change | Protects storage of sensitive information, such as private keys, and prevents access by unauthorized services, processes, or users.<br>This service is used on a dedicated Web server for smart card logon. |
| Remote Access Auto Connection Manager | **Manual** | See comment | Detects unsuccessful attempts to connect to a remote network or computer and provides alternative methods for connection.<br>On a dedicated Web server, this service can be disabled when no VPN or dial-up connections are initiated. |
| Remote Access Connection Manager | **Manual** | See comment | Manages VPN and dial-up connection from the Web server to the Internet or other remote networks.<br>On a dedicated Web server, this service can be disabled when no VPN or dial-up connections are initiated. |
| Remote Desktop Help Sessions Manager | **Manual** | Disabled | Manages and controls Remote Assistance.<br>On a dedicated Web server, this service can be disabled. Use Terminal Services instead. |
| Remote Procedure Call (RPC) | **Automatic** | No change | Serves as the RPC endpoint mapper for all applications and services that use RPC communications. |
| Remote Procedure Call (RPC) Locater | **Manual** | See comment | Enables RPC clients using the RpcNs* family of application programming interfaces (APIs) to locate RPC servers and manage the RPC name service database.<br>This service can be disabled if no applications use the RpcNs* APIs. |
| Remote Registry Service | **Automatic** | No change | Enables remote users to modify registry settings on the Web server, provided the remote users have the required permissions. By default, only members of the Administrators and Backup Operators groups can access the registry remotely. |
| Removable Storage | **Manual** | See comment | Manages and catalogs removable media, and operates automated removable media devices, such as tape auto loaders or CD jukeboxes.<br>This service can be disabled when removable media devices are directly connected to the Web server. |
| Resultant Set of Policy Provider | **Manual** | No change | Enables a user to connect to a remote computer, access the Windows Management Instrumentation (WMI) database for that Web server, and then either verify the current Group Policy settings or check the settings before they are applied. |
| Routing and Remote Access | **Disabled** | No change | Enables LAN-to-LAN, LAN-to-WAN, VPN, and NAT routing services. |

| | | | |
|---|---|---|---|
| Secondary Logon | **Automatic** | No change | Allows you to run specific tools and programs with different permissions and user rights than the default permissions and user rights of the account under which you logged on. |
| Security Accounts Manager | **Automatic** | No change | A protected subsystem that manages user and group account information. |
| Server | **Automatic** | No change | Provides RPC support, file sharing, print sharing, and named pipe sharing over the network. |
| Shell Hardware Detection | **Automatic** | No change | Provides notification for AutoPlay hardware events. |
| Smart Card | **Manual** | No change | Manages and controls access to a smart card that is inserted into a smart card reader attached to the Web server. |
| Special Administration Console Helper | **Manual** | No change | Allows administrators to remotely access a command prompt by using Emergency Management Services.<br>This service can be disabled when Emergency Management Services is not being used to remotely manage the Web server. |
| System Event Notification | **Automatic** | No change | Monitors system events and notifies subscribers to the COM+ Event System of these events. |
| Task Scheduler | **Automatic** | No change | Provides the ability to schedule automated tasks on the Web server. |
| TCP/IP NetBIOS Helper Service | **Automatic** | No change | Provides support for the NetBIOS over TCP/IP (NetBT) service and NetBIOS name resolution for clients. |
| Telephony | **Manual** | See comment | Provides Telephony API (TAPI) support of client programs that control telephony devices and IP-based voice connections.<br>On a dedicated Web server, this service can be disabled when TAPI is not used by applications. |
| Telnet | **Manual** | Disabled | Enables a remote user to log on and run applications from a command line on the Web server.<br>To reduce the attack surface, disable Telnet unless it is used for remote administration of branch offices or of Web servers that have no keyboard or monitor directly attached (also known as *headless* Web servers). Because Telnet traffic is plaintext, Terminal Services is the preferred method for remote administration. |
| Terminal Services | **Manual** | See comment | Allows multiple remote users to be connected interactively to the Web server, and provides display of desktops and running applications.<br>To reduce the attack surface, disable Terminal Services unless it is used for remote administration of branch offices or headless Web servers. |
| Terminal Services Session Directory | **Disabled** | No change | Enables a user connection request to be routed to the appropriate terminal server in a cluster. |
| Themes | **Disabled** | No change | Provides user-experience theme management. |

| | | | |
|---|---|---|---|
| Uninterruptible Power Supply (UPS) | **Automatic** | No change | Manages a UPS that is connected to the Web server by a serial port. |
| Upload Managers | **Manual** | See comment | Manages the synchronous and asynchronous file transfers between clients and servers on the network. Driver data is anonymously uploaded from these transfers and then used by Microsoft to help users find the drivers they need. The Driver Feedback Server asks for the permission of the client to upload the hardware profile of the Web server and then search the Internet for information about how to obtain the appropriate drivers or how to get support.<br><br>To reduce the attack surface, disable this service on dedicated Web servers. |
| Virtual Disk Services | **Manual** | No change | Provides software volume and hardware volume management service. |
| Volume Shadow Copy | **Manual** | No change | Manages and implements volume shadow copies that are used for backup and other purposes.<br><br>This service can be disabled when volume shadow copies are used on the Web server. |
| WebClient | **Disabled** | No change | Enables Windows-based programs to create, access, and modify Internet-based files. |
| Windows Audio | **Disabled** | No change | Manages audio devices for Windows-based programs. |
| Windows Image Acquisition (WIA) | **Disabled** | No change | Provides image acquisition services for scanners and cameras. |
| Windows Installer | **Manual** | No change | Adds, modifies, and removes applications that are provided as a Windows Installer (.msi) package. |
| Windows Management Instrumentation (WMI) | **Automatic** | No change | Provides a common interface and object model to access management information about the Web server through the WMI interface. |
| Windows Management Instrumentation Driver Extensions | **Manual** | No change | Monitors all drivers and event trace providers that are configured to publish WMI or event trace information. |
| Windows Time | **Automatic** | No change | Sets the Web server clock, and maintains date and time synchronization for all computers in the network. |
| WinHTTP Web Proxy Auto-Discovery Service | **Manual** | See comment | Implements the Web Proxy Auto-Discovery (WPAD) protocol for Windows HTTP services (WinHTTP) and enables an HTTP client to automatically discover a proxy configuration.<br><br>On dedicated Web servers, this service can be disabled |
| Wireless Configuration | **Automatic** | See comment | Enables automatic configuration for IEEE 802.11 adapters.<br><br>On dedicated Web servers without wireless network adapters, this service can be disabled. |

| | | | |
|---|---|---|---|
| WMI Performance Adapter | **Manual** | See comment | Provides performance library information from WMI providers to clients on the network.<br>On dedicated Web servers that do not use WMI to provide performance library information, this service can be disabled. |
| Workstation | **Automatic** | No change | Creates and maintains client network connections to remote servers. |

# Enabling Only Essential IIS Components and Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 includes other components and services in addition to the WWW service, such as the File Transfer Protocol Service (FTP service) and the Simple Mail Transfer Protocol (SMTP) service. You can install and enable IIS components and services by using the **Application Server** subcomponent, which is found in **Add or Remove Windows Components** in **Add or Remove Programs** in Control Panel. After installing IIS, you need to enable the IIS 6.0 components and services that are required by the Web sites and applications running on your Web server.

Enable only the essential IIS 6.0 components and services that are required by your Web sites and applications. Enabling unnecessary components and services increases the attack surface of the Web server.

When a Web site or application does not function on the Web server and you suspect that an IIS 6.0 component or service might need to be enabled, complete the following steps:

- Enable the individual IIS 6.0 component or service that you believe will allow the Web site or application to function.
- Test the Web site or application for proper operation.
- If the Web site or application functions correctly, further configure the IIS 6.0 component or service to the most restrictive security settings.
- If enabling the IIS 6.0 component or service does not allow the Web site or application to function, disable the IIS 6.0 component or service and continue troubleshooting the problem.

The Web site or application might not function properly because of issues that are not security-related. For example, an Internet Server API (ISAPI) extension that is used by an application might not be installed properly. Although it might appear that the ISAPI extension is disabled, the problem might actually be caused by a faulty installation or configuration setting for the ISAPI extension. For more information about troubleshooting problems related to Web sites and applications that are not functioning, see Troubleshooting in IIS 6.0.

| **Important** |
| --- |
| When you are troubleshooting Web site- and application-related problems, do not enable all of the IIS 6.0 components and services. Enabling all of the IIS 6.0 components and services will unnecessarily increase the attack surface of the Web server. |

For each of the subcomponents of the application server that are listed in Table 3.2 through Table 3.6, complete the following steps:

- Review the recommended settings to determine whether you need to make changes to the default settings.
- Determine, based on the information provided in the comments, if the recommendation applies to your server.
- Enable or disable the component based on the decisions made in the previous steps.

  For more information about how to configure the IIS 6.0 protocols and services, see Configure IIS Components and Services.

**Table 3.2 Subcomponents of the Application Server**

| Subcomponent | Default Setting | Recommended Setting | Comment |
| --- | --- | --- | --- |

| | | | |
|---|---|---|---|
| Application Server Console | **Enabled** | No change | Provides a Microsoft Management Console (MMC) snap-in that includes administration for all of the Web Application Server (WAS) components.<br><br>On a dedicated Web server, this component is not required because only IIS Manager is used. |
| ASP.NET | **Disabled** | See comment | Provides support for ASP.NET applications.<br><br>Enable this component when you need to run ASP.NET applications on the Web server. |
| Enable network COM+ access | **Enabled** | See comment | Allows the Web server to host COM+ components for distributed applications.<br><br>Disable this component unless it is required by your applications. |
| Enable network DTC access | **Disabled** | See comment | Allows the Web server to host applications that participate in network transactions through Distributed Transaction Coordinator (DTC).<br><br>Disable this component unless it is required by your applications. |
| Internet Information Services (IIS) | **Enabled** (See Table 3.3 for subcomponents) | No change | Provides basic Web and FTP services.<br><br>This component is required on a dedicated Web server.<br><br>Note: If this component is not enabled, then all subcomponents are not enabled. |
| Message Queuing | **Disabled** (See Table 3.4 for subcomponents) | See comment | Provides guaranteed messaging, security, and transactional support for applications that communicate through messaging services provided by Message Queuing (also known as MSMQ).<br><br>This component is required when your Web sites and applications use Message Queuing.<br><br>Note: If this component is not enabled, then all subcomponents are not enabled. |

**Table 3.3 Subcomponents of Internet Information Services (IIS)**

| Subcomponent | Default Setting | Recommended Setting | Comment |
|---|---|---|---|
| Background Intelligent Transfer Service (BITS) Server Extensions | **Disabled** | See comment | Allows background file transfers from a BITS client to the Web server.<br><br>Enable this component when it is required by any of your client applications.<br><br>Note: Windows Update, SUS, and Automatic Updates do not require this component; they require the BITS client component, which is not part of IIS.<br><br>For more information, see Obtaining and Applying Current Security Patches later in this section. |
| Common Files | **Enabled** | No change | On a dedicated Web server, these files are required by IIS and must always be enabled. |

| | | | |
|---|---|---|---|
| File Transfer Protocol (FTP) Service | **Disabled** | No change | Allows the Web server to provide FTP services.<br><br>This component is not required on a dedicated Web server. However, you might need to enable FTP on a server that is only used for posting content, to support software such as Microsoft FrontPage® 2002 without enabling FrontPage 2002 Server Extensions.<br><br>Because the FTP credentials are always sent in plaintext, it is recommended you connect to FTP servers through a secured connection, such as those provided by Internet Protocol security (IPsec) or a VPN tunnel.<br><br>For more information, see Using IPsec or VPN with Remote Administration later in this section. |
| FrontPage 2002 Server Extensions | **Disabled** | See comment | Provides FrontPage support for administering and publishing Web sites.<br><br>On a dedicated Web server, disable when no Web sites are using FrontPage Server Extensions. |
| Internet Information Services Manager | **Enabled** | See comment | Administrative interface for IIS.<br><br>Disable when you do not want to administer the Web server locally. |
| Internet Printing | **Disabled** | No change | Provides Web-based printer management and allows printers to be shared by using HTTP.<br><br>This component is not required on a dedicated Web server. |
| NNTP Service | **Disabled** | No change | Distributes, queries, retrieves, and posts Usenet news articles on the Internet.<br><br>This component is not required on a dedicated Web server. |
| SMTP Service | **Enabled** | Disabled | Supports the transfer of electronic mail.<br><br>This component is not required on a dedicated Web server. |
| World Wide Web Service | **Enabled** (See Table 3.6 for subcomponents) | No change | Provides Internet services, such as static and dynamic content, to clients.<br><br>This component is required on a dedicated Web server.<br><br>Note: If this component is not enabled, then all subcomponents are not enabled. |

**Table 3.4 Subcomponents of Message Queuing**

| Subcomponent | Default Setting | Recommended Setting | Comment |
|---|---|---|---|
| Active Directory Integration | **Disabled** | See comment | Provides integration with Active Directory whenever the Web server belongs to a domain. |
| Common | **Disabled** | See comment | Required by Message Queuing. |
| Downlevel Client Support | **Disabled** | See comment | Provides access to Active Directory and site recognition for clients that are not Active Directory-aware. |
| MSMQ HTTP Support | **Disabled** | See comment | Provides the sending and receiving of messages over the HTTP transport. |

| | | | |
|---|---|---|---|
| Routing support | **Disabled** | See comment | Provides store-and-forward messaging as well as efficient routing services for Message Queuing. |
| Triggers | **Disabled** | See comment | Provides support to associate the arrival of incoming messages at a queue with functionality in a COM component or stand-alone program.<br><br>This component is required when your Web sites and applications use Message Queuing and use Message Queuing triggers. |

**Table 3.5 Subcomponents of the Background Intelligent Transfer Service (BITS) Server Extension**

| Subcomponent | Default Setting | Recommended Setting | Comment |
|---|---|---|---|
| BITS management console snap-in | **Disabled** | No change | Installs an MMC snap-in for administering BITS.<br><br>Enable this component when you enable the BITS server extension ISAPI component. |
| BITS server extension ISAPI | **Disabled** | No change | Installs the BITS ISAPI so that the Web server can receive background file transfers from a BITS client.<br><br>Enable this component when it is required by any of your client applications.<br><br>For more information, see Obtaining and Applying Current Security Patches later in this section. |

**Table 3.6 Subcomponents of the World Wide Web Service**

| Subcomponent | Default Setting | Recommended Setting | Comment |
|---|---|---|---|
| Active Server Pages | **Disabled** | See comment | Provides support for Active Server Pages (ASP).<br><br>Disable this component when none of the Web sites or applications on the Web server uses ASP. You can disable this component in Add or Remove Windows Components, which is accessible from **Add or Remove Programs** in Control Panel, or in the Web Service Extensions node in IIS Manager.<br><br>For more information, see Enabling Only Essential Web Service Extensions later in this section. |
| Internet Data Connector | **Disabled** | See comment | Provides support for dynamic content provided through files with .idc extensions.<br><br>Disable this component when none of the Web sites or applications on the Web server include files with .idc extensions. You can disable this component in **Add or Remove Windows Components**, which is accessible from **Add or Remove Programs** in Control Panel, or in the Web Service Extensions node in IIS Manager.<br><br>For more information, see Enabling Only Essential Web Service Extensions later in this section. |
| Remote Administration (HTML) | **Disabled** | No change | Provides an HTML interface for administering IIS.<br><br>Use IIS Manager instead to provide easier administration and to reduce the attack surface of the Web server. This component is not required on a dedicated Web server. |
| Remote Desktop Web Connection | **Disabled** | No change | Includes Microsoft ActiveX® controls and sample pages for hosting Terminal Services client connections.<br><br>Use IIS Manager instead to provide easier administration and to reduce the attack surface of the Web server. This component is not required on a dedicated Web server. |

| | | | |
|---|---|---|---|
| Server-Side Includes | **Disabled** | See comment | Provides support for .shtm, .shtml, and .stm files.<br><br>Disable this component when none of the Web sites or applications on the Web server includes files with these extensions. |
| WebDav Publishing | **Disabled** | Disabled | Web Distributed Authoring and Versioning (WebDAV) extends the HTTP/1.1 protocol to allow clients to publish, lock, and manage resources on the Web.<br><br>Disable this component on a dedicated Web server. You can disable this component in **Add or Remove Windows Components**, which is accessible from **Add or Remove Programs** in Control Panel, or in the Web Service Extensions node in IIS Manager<br><br>For more information, see Enabling Only Essential Web Service Extensions later in this section. |
| World Wide Web Service | **Enabled** | No change | Provides Internet services, such as static and dynamic content, to clients.<br><br>This component is required on a dedicated Web server. |

# Enabling Only Essential Web Service Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If your Web sites and applications that are hosted on IIS 6.0 have extended functionality beyond static Web pages, including the ability to generate dynamic content, any dynamic content served or extended features provided by the Web server are done through Web service extensions.

For security reasons, you can enable or disable individual Web service extensions in IIS 6.0. After a new installation, IIS serves only static content. You can enable dynamic content capabilities, such as ASP.NET, Server-Side Includes, Web Distributed Authoring and Version (WebDAV) publishing, and FrontPage 2002 Server Extensions, in the Web Service Extensions node in IIS Manager.

For example, one of your applications might use a custom Internet Server API (ISAPI) extension to provide access to a proprietary database. First, you need to add the custom ISAPI extension to the Web service extensions list. Then you can set the ISAPI extension that is used by the application to **Allowed**, explicitly granting it permission to run.

Enabling all of the Web service extensions ensures the highest possible compatibility with existing applications, regardless of whether you enable each of the Web service extensions individually or change the status of **All Unknown ISAPI Extensions** to **Allowed**. However, enabling all of the Web service extensions creates a security risk because it increases the attack surface of the Web server by enabling functionality that might be unnecessary for your server.

Web service extensions allow you to enable and disable the serving of dynamic content. MIME types allow you to enable and disable the serving of static content.

| **Tip** |
| --- |
| If the appropriate Web service extension is not enabled, the Web server returns a 404 error to the client when attempting to serve the dynamic content. When the 404 error is returned as a result of a Web service extension not being enabled, a 404.2 error entry is placed in the IIS log. For more information about troubleshooting IIS, see Troubleshooting in IIS 6.0. |

Configure the Web service extensions by completing the following steps:

1. Enable the essential predefined Web service extensions based on the information in Table 3.1.

   **Table 3.7 Predefined Web Service Extensions**

   | Web Service Extension | Description |
   | --- | --- |
   | Active Server Pages | Enable this extension when one or more of the Web sites and applications contains ASP content. |
   | ASP.NET version 1.1.4322 | Enable this extension when one or more of the Web sites and applications contains ASP.NET content. |
   | FrontPage Server Extensions 2002 | Enable this extension when one or more of the Web sites use FrontPage Server Extensions. |
   | Internet Data Connector | Enable this extension when one or more of the Web sites and applications uses the Internet Data Connector (IDC) to display database information (content includes .idc and .idx files). |
   | Server-Side Includes | Enable this extension when one or more of the Web sites use server-side include (SSI) directives to instruct the Web server to insert various types of content into a Web page. |
   | WebDAV | Enable this extension when you want to support WebDAV on the Web server. This Web service extension is not recommended on a dedicated Web server. |

2. For each Web service extension that is used by your Web sites and applications and that is not one of the default Web

service extensions, add a new entry to the Web service extensions list and configure the status of the new entry to **Allowed**. For information about how to add a Web service extension to the list, see Configure Web Service Extensions.

3. Use a Web browser on a client computer to verify that the Web sites and applications function properly on the server.

# Configuring Windows Server 2003 Security Settings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After installing Windows Server 2003, the security settings are configured so that the server is locked down. After installing IIS 6.0, evaluate the default security settings to determine whether they are sufficient for the Web sites and applications that your Web server hosts. You might need more stringent security requirements for Web sites and applications when the following is true:

- Users on the Internet access the Web sites and applications.
- Web sites and applications contain confidential information.

Configure Windows Server 2003 to more restrictive security settings by completing the following steps:

1. Rename the Administrator account.

   The built-in account, Administrator, exists by default on every newly installed Web server. Potential attackers only have to guess the password for this well-known user account to exploit it. You can rename the Administrator user account to help protect your Web server from potential attackers. For more information about how to rename the Administrator user account, see Secure Windows Server 2003 Built-in Accounts.

   | **Important** |
   | --- |
   | During the default installation of Windows Server 2003, the Guest account is disabled. Ensure that the Guest account has not been enabled since the installation. |

2. Format all disk volumes with the NTFS file system.

   From a security perspective, the primary reason for requiring that all disk volumes are formatted with NTFS is that NTFS is the only file system supported by Windows Server 2003 that allows you to secure files and folders. FAT or FAT32 partitions cannot be secured.

   Because the Web sites and applications are stored as files and folders on the Web server, NTFS helps prevent unauthorized users from directly accessing or modifying the files and folders that make up your Web sites and applications. For more information about the benefits of formatting disk volumes as NTFS on Web servers, see NTFS Permissions.

   If any existing disk volumes are FAT or FAT32, convert the disk volumes to NTFS. For more information about how to convert existing disk volumes to NTFS, see Convert Existing Disk Volumes to NTFS.

3. Remove NTFS permissions that are granted to the **Everyone** group on the root folder of all disk volumes.

   By default, the **Everyone** group is granted **Read and Execute** permissions on the root folder of each disk volume. The default permissions can pose a potential security threat for any newly created folders on the volumes because, unless explicitly denied, these permissions are inherited in any new folders. To help prevent this potential security problem, remove all permissions that are granted to the **Everyone** group on the root folder of all disk volumes.

   | **Important** |
   | --- |
   | The **Administrators** group still has full control on the root folder of each disk volume. In Setting NTFS Permissions, later in this section, you will grant access to the Web site users by setting the appropriate NTFS permissions on the Web site content. |

   For more information about how to remove the permissions that are granted to the **Everyone** group on the root folder of each disk volume, see Secure the Root Folder of Each Disk Volume.

4. Remove any compilers or development environments.

   If compilers or development environments are installed on production Web servers, potential attackers can use them to upload source files to a malicious program and then use the Web server to compile the malicious program. In many instances, the source files might not be perceived as a threat, whereas an executable file would be. You can remove any compilers and development environments to help ensure that potential attackers cannot remotely compile a malicious program and then run that malicious program on the Web server.

   Consult the documentation of the compiler or development environment for information about how to remove them.

5. Disable NetBIOS over TCP/IP.

   To prevent attackers from executing the NetBIOS Adapter Status command on a server, and reveal the name of the user who is currently logged on, disable NetBIOS over TCP/IP on public connections of the server.

| **Important** |
| --- |
| Before you disable NetBIOS over TCP/IP, make sure that it doesn't affect the management tools that you use to manage the server and other applications running on the server. You can do this by disabling NetBIOS over TCP/IP on a test server before disabling it on your production servers. |

# Preventing Unauthorized Access to Web Sites and Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Each Web site and application in IIS 6.0 and Windows Server 2003 is stored as a grouping of folders and files. Unauthorized access to, or modification of, these files and folders can present a serious breach of security. You must ensure that only authorized users can access or modify the Web sites and applications that are hosted on your Web server.

To help prevent unauthorized access to Web sites and applications on your Web server, use any combination of the steps illustrated in Figure 3.3. Based on the security requirements of your organization, you might perform a subset of the steps or all of the steps.

**Figure 3.3 Preventing Unauthorized Access to Web Sites and Applications**



# Storing Content on a Dedicated Disk Volume

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Store the files and folders that comprise the content of your Web sites and applications on a dedicated disk volume that does not contain the operating system. Doing this helps prevent *directory transversal attacks*. Directory transversal attacks occur when an attacker attempts to send the Web server a request for a file that is located in another directory structure.

For example, Cmd.exe exists in the *systemroot*\System32 folder. Without the appropriate security settings, an attacker might be able to make a request to *systemroot*\System32\Cmd.exe and invoke the command prompt. If the Web site content is stored on a separate disk volume, such a directory transversal attack cannot work because Cmd.exe does not exist on the same disk volume. The default NTFS permissions for Windows Server 2003 prohibit anonymous users from executing or modifying any files in the systemroot folder and subfolders, so that only an unauthorized authenticated user can perform this type of attack.

In addition to security concerns, placing the content on a disk volume that is dedicated to Web site and application content makes administration tasks, such as backup and restore, easier. In cases where you store the content on a separate physical drive that is dedicated to the content, you will reduce the disk contention on the system volume and improve overall disk access performance. Ensure that the dedicated disk volume is formatted as NTFS.

To help protect your Web sites and applications, store content on dedicated disk volumes by completing the following steps:

1. Create a disk volume, or designate an existing disk volume, where the Web sites and applications will be stored.

2. Configure the NTFS permissions on the root of the disk volume so that:

    - The Administrators group has full control.

    - All other permissions are removed.

3. Create a folder, or designate an existing folder, on the dedicated disk volume to hold the subfolders that will contain the Web sites and applications.

4. Beneath the folder that you created, or designated, in the previous step, create a subfolder for each Web site or application that will be installed on the Web server.

5. Install the Web sites and applications in the subfolders that you created in the previous step.

At this step in the deployment process, only members of the Administrators group have access to the content. You will grant access to the users who will access the Web sites and applications in Setting NTFS Permissions later in this section.

# Setting IIS Web Site Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 6.0, you can set Web site permissions, which allow you to control access to a Web site or virtual directory. IIS examines Web site permissions to determine which type of action can occur, such as accessing the source code of a script or browsing folders.

Use Web site permissions in conjunction with NTFS permissions, not in place of NTFS permissions. You can set Web site permissions for specific sites, directories, and files. Unlike NTFS permissions, Web site permissions affect everyone who tries to access your Web site.

| **Note** |
| --- |
| If Web site permissions conflict with NTFS permissions for a directory or file, the more restrictive settings are applied. |

Table 3.8 lists and describes the Web site permissions that are supported by IIS 6.0.

**Table 3.8 Web Site Permissions That Are Supported by IIS 6.0**

| Permission | Description |
| --- | --- |
| Read | Users can view the content and properties of directories or files. This permission is set by default. This permission is required for Web sites that have static content. If all of your content is scripted, such as a Web site that only uses Active Server Pages (ASP) content, you can remove the Read permission. |
| Write | Users can change content and properties of directories or files. |
| Script Source Access | Users can access source files. If the Read permission is set, then users can read source files; if the Write permission is set, then users can modify the content and properties of the source files. The Script Source Access permission also applies to the source code for scripts. This option is not available if both the Read and Write permissions are not set.<br><br>Set this permission only when using Web Distributed Authoring and Versioning (WebDAV). In addition, make sure that you require authentication for this site and that your file permissions are set correctly.<br><br>**Important**<br>When you set the Script Source Access permission, users might be able to view sensitive information, such as a user name and password. Users might also be able to change source code that runs on your server, and seriously affect the security and performance of your server. |
| Directory browsing | Users can view file lists and collections. |
| Log visits | A log entry is created for each visit to the Web site. As an operational security practice, it is highly recommend that you enable logging. |
| Index this resource | Indexing Service can index this resource. This allows searches to be performed on the resource. |
| Execute | Users have the appropriate level of script execution:<br>• None. Does not allow scripts or executables to run on the server.<br>• Scripts only. Allows only scripts to run on the server.<br>• Scripts and Executables. Allows both scripts and executables to run on the server. |

For information about how to set Web site permissions, see Configure Web Site Permissions.

# Setting IP Address and Domain Name Restrictions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

One method of protecting the Web sites and applications that are hosted on your server is to restrict access from specific IP addresses or domain names. You can explicitly grant or deny access to any combination of IP address ranges or domain names.

By restricting access to Web sites and applications by using IP address ranges or domain names, you can grant or deny access to a specific set of computers or to an organization. The restrictions that you specify affect the entire Web site or application and cannot be configured for individual portions of the Web site or application.

Restrict access to a specific Web site for a specific IP address range or domain name by completing the following steps:

1. Specify the default access that will be given to the majority of users accessing the application by doing one of the following:

   - To allow the majority of users to access the application, enable default access.
   - To allow a limited number of users to access the application, disable default access.

2. For each computer, or group of computers, that you want to grant or deny access, specify the IP address range or domain name for the clients that are exceptions to the default access specified in Step 1.

   Unless you are unable to identify the IP address range for the computers, you must specify the domain name. From a performance perspective, specifying the IP address range is preferred. If you specify a domain name, Domain Name System (DNS) reverse lookups must be done each time a user accesses the application and the performance of your application is degraded.

   Specify the IP address range in the form of a single IP address or a network ID with a corresponding subnet mask.

For more information about setting IP address and domain name restrictions, see Configure IP Address and Domain Name Restrictions.

# Setting NTFS Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

NTFS permissions allow you to set permissions that are observed by IIS and by other Windows Server 2003 components. Windows Server 2003 examines NTFS permissions to determine the types of access a user, or a process, has on a specific file or folder.

Use NTFS permissions in conjunction with Web site permissions, not in place of Web site permissions. NTFS permissions affect only the accounts that have been granted or denied access to the Web site and application content. Web site permissions affect all of the users who access the Web site or application.

| **Note** |
| --- |
| If Web site permissions conflict with NTFS permissions for a directory or file, the more restrictive settings are applied. |

You need to set NTFS permissions to allow the following situations:

- Administrators can manage the content of the Web sites and applications.
- Users can, at a minimum, read the content of the Web sites and applications.
- Application pool identities can, at a minimum, read the content of the Web sites and applications.

Web sites and applications can run under the identity of the following:

**The user who is accessing the Web sites and applications**
When you want to restrict access to resources, such as specific Web pages or database content that is stored in SQL Server, run your Web sites and applications under the identity of the user. For example, Basic authentication can allow Web sites and applications to pass through the identity of the user to other servers, such as a computer running SQL Server. By using this method, you can control the behavior of the Web site or application on a user-by-user basis.

**The application pool identity that is used by the Web sites and applications**
When you want to isolate Web sites or applications that are hosted on the same Web server from one another, run the Web sites or applications under the application pool identity. By using this method, you can prevent Web sites and applications from interfering with one another independent of the users who are accessing the Web sites and applications. For more information about isolating Web sites and applications, see Isolating Web Sites and Applications later in this section.

Regardless of the identity that is used to run the Web site or application, you need to assign the appropriate NTFS permissions to the Web site or application so that it can run under the corresponding identity. Typically, these NTFS permissions are assigned to a group to which a number of users belong. Use this group when setting the permissions on the resources.

The primary disadvantage of restricting access by user accounts and NTFS permissions is that each user must have an account and must use that account to run the Web sites and applications. For your Internet-based Web sites and applications, requiring users to have accounts might be impractical. However, for intranet Web sites and applications you can use the existing accounts of users.

Explicitly deny access to anonymous accounts on Web sites and applications when you want to prevent anonymous access. *Anonymous access* occurs when a user who has no authenticated credentials accesses system resources. Anonymous accounts include the built-in Guest account, the group Guests, and the IIS anonymous accounts.

In addition to explicitly denying access to anonymous accounts, eliminate write access permissions for all users except members of the Administrators group.

| **Tip** |
| --- |
| If IIS denies access to content, you can enable object access auditing to find out the account that was used to access the content. The failed access event is recorded in the Security event log. The event log entry specifies the account that was used in the failed access. After you identify the account used in the failed access, grant the appropriate NTFS permissions to the account. |

# Isolating Web Sites and Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Although some of the Web servers that you deploy host only one Web site or application, you might also need to host multiple Web sites and applications on the same Web server. When a Web server hosts multiple Web sites and applications, each Web site and application requires a certain level of isolation.

For example, an Internet service provider (ISP) might host Web sites and applications for hundreds of organizations, each having a unique Web site. In this situation, the security requirements of each organization require a high degree of isolation between Web sites and applications.

Figure 3.4 illustrates the tasks involved in the process isolating your Web sites and applications.

**Figure 3.4 Isolating Web Sites and Applications**




You need to prevent multiple Web sites and applications that are hosted on the same Web sever from adversely interacting with one another. When IIS 6.0 is running in worker process isolation mode, you can isolate Web sites and applications hosted on the same Web server by specifying that the Web sites and applications belong to separate application pools. An application pool is a grouping of Web sites or applications served by the same worker process. Application pools can be used to help prevent the Web sites and applications running in one application pool from accessing the content contained in another application pool.

For each application pool, you can specify an *application pool identity*, which is a user account that is assigned to an application pool. After specifying the application pool identity, you assign permissions (such as NTFS permissions or SQL database permissions) for each application pool identity. Because individual application pools can use different identities, you can selectively grant or deny resource access to an application pool. The Web sites and applications running in an application pool have the same user rights and resource permissions assigned to the application pool identity.

For more information about setting NTFS permissions for Web sites and applications, see Setting NTFS Permissions earlier in this section.

| **Note** |
| --- |
| Web sites and applications that are running in the same application pool can affect the availability of other Web sites and applications in the same application pool. To enhance the availability of your Web sites and applications, isolate unstable Web sites and applications in a separate application pool. For more information about improving the availability of your Web server through application pools see, Ensuring Application Availability. |

# Evaluating the Effects of Impersonation on Application Compatibility

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Securing your Web sites and applications by isolating them into separate application pools with unique identities can cause application compatibility problems when you are using anything but anonymous access. The application compatibility problems arise from the complexities of *impersonation*.

Impersonation allows a worker process to run under security credentials that are different from its base identity. Because the two are commonly confused, it is important to understand how the worker process identity that is established by the application pool identity is related to the impersonated user.

When a worker process is created by the WWW service, it is created with a process token that is associated with the application pool identity. This establishes the process identity of the worker process. By default, all of the actions taken by the worker process are completed in the context of this worker process identity account. However, when a client request is processed, the thread that processes the request uses a token associated with the client, which is also known as the *authenticated user's token*, during the duration of the request.

Before IIS serves a URL, the authenticated user's token is verified against the access control list (ACL) of the resource that is being requested. Additionally, if the request is for an Internet Server API (ISAPI) extension, such as ASP, the worker process applies the authenticated user's token as an impersonation token to the thread that calls the ISAPI extension. When the ISAPI extension begins processing the request, this impersonation token applies to the actions it takes. Consequently, the actions taken by an ISAPI extension are associated with the authenticated user, not the process identity.

Evaluate the effects of impersonation behavior on compatibility for the following:

- ASP applications
- ASP.NET applications

# Identifying the Impersonation Behavior for ASP Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For Active Server Pages (ASP) applications, the type of authentication that is used by the user automatically determines impersonation behavior. Because the impersonation behavior is automatic, no configuration is required.

The impersonation behavior in an ASP application is as follows:

- If an anonymous user makes a request, the thread token is based on the user account that is configured as the anonymous user identity (by default, this is the IUSR_*machinename* user account). You should ensure that the appropriate access control lists (ACLs) are configured on any content that the anonymous user should not be able to access.

- If an authenticated user makes a request, the thread token is based on the authenticated account of the user.

# Selecting the Impersonation Behavior for ASP.NET Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Unlike ASP applications, you need to configure the impersonation behavior for ASP.NET applications. If you enable impersonation, ASP.NET receives the security token to impersonate from IIS. By specifying a value in the Web.config file of the application, you control the impersonation setting. You have the following three options when specifying this setting.

**Impersonation is disabled**
This is the default setting. In this instance, the ASP.NET thread runs using the process token of the application worker process regardless of which combination of IIS and ASP.NET authentication is used. By default, the process token of the application worker process is NetworkService.

Disable impersonation by modifying the Web.config file of the application to include the following setting:

```
<identity impersonate="false" />
```

**Impersonation is enabled**
In this instance, ASP.NET impersonates the token passed to it by IIS, which is either an authenticated user or the anonymous user account (IUSR_*machinename*). For backward compatibility with ASP, you must enable impersonation.

Enable impersonation by modifying the Web.config file of the application to include the following setting:

```
<identity impersonate="true" />
```

**Impersonation is enabled and a specific impersonation identity is specified**
In this instance, ASP.NET impersonates the token that is generated using the configured identity. In this case, ASP.NET does not use the token of the authenticated client, if applicable, except when performing access checks.

Enable impersonation and specify an impersonation identity by modifying the Web.config file of the application to include the following setting:

```
<identity impersonate="true" name="domain\user" password="password" />
```

# Configuring Web Sites and Applications for Isolation

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Complete the following steps to identify when Web sites and applications require isolation for security reasons:

| **Note** |
| --- |
| Running worker processes under different identities can cause application compatibility problems, especially for Web sites that use user authentication. |

1. Create a list of the Web sites and applications to be hosted on the Web server.

2. Group the Web sites and applications by organization (or business unit within an organization if all of the Web sites and applications hosted on the Web server are owned by one organization).

3. Subdivide each group created in the previous step into smaller groups of Web sites and applications that require the same user rights and resource access.

4. For each group created in the previous step, create a new application pool to be used by the Web sites and applications within the pool.

   For information about how to create application pools, see Isolate Applications in Worker Process Isolation Mode.

5. Assign the Web sites and applications within each group to the corresponding application pool.

   For information about how to assign the Web site to the new application, see Isolate Applications in Worker Process Isolation Mode.

6. For each application pool, create a service account to be used as the application pool identity.

   In IIS, the default identity for newly created application pools is NetworkService. To ensure that you can properly assign permissions to resources, create a new *service account*. A service account is a user account that is created explicitly for the purpose of providing a security context for services running on Windows Server 2003.

   In addition, you must add the service account to the IIS_WPG group to provide the appropriate access to the IIS metabase and content. The IIS_WPG group is granted the appropriate user rights and resource permissions to allow most Web sites and applications to run properly.

   For more information about how to create a service account to be used as an identity for an application pool and how to add the account to the IIS_WPG group, see Create a Service Account.

7. Assign any additional *user rights* to the application pool identities.

   User rights authorize users to perform specific actions, such as logging on to a system interactively or backing up files and directories. User rights are different from permissions because user rights apply to user accounts, whereas permissions are attached to objects.

   The user rights granted to the IIS_WPG group are sufficient for most Web sites applications. When the user rights granted to the IIS_WPG group are insufficient, grant only the user rights to the user account, which is used as the identify for the application pool, that are necessary to ensure the appropriate operation and behavior of the application. Ensure that any nonessential user rights are removed to prevent the Web sites and applications from having elevated user rights.

   For more information about how to grant the appropriate user rights for an application pool identity, see Grant User Rights to a Service Account.

8. Assign the service account identity to the corresponding application pool.

   For more information about how to assign the identity to the corresponding application pool, see Configure Application Pool Identity.

9. Assign the appropriate resource permissions, such as NTFS or SQL database permissions, to the application pool identities.

   Assign only the NTFS file and folder permissions that are necessary to ensure the appropriate operation and behavior of the application. By default, grant only Read permissions to the application pool identity to insure that the Web sites and applications in the application pool cannot modify the Web site content or other files on the Web server. If the applications require Write access to any files and folders, consult the application developers to determine if the application can be modified so that Write access is not required.

# Adding Web Sites to an IIS 6.0 Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS enables you to create multiple Web sites on a single server. To add a site to a server, you must prepare the server and associated network services, and create a unique identity for the site.

**Preparing to Add a Web Site**
Adding a Web site to a server that is already hosting a Web site requires some preparation before you run the Web Site Creation Wizard. Consider the following recommendations in your preparation process:

- Carefully review the methods of hosting multiple Web sites to determine which method of creating unique Web site identities is appropriate for your environment.

- Decide if you will use a unique IP address to identify your Web site. Multiple Web sites can use the same IP address. If you use a unique IP address to identify a new Web site, you must obtain a static IP address from your organization or Internet service provider (ISP) and configure the TCP/IP settings on the server.

- Select a unique host header name, if you use host header names to identify a new Web site, On a private network, the host header can be an intranet site name. On the Internet, the host header must be a publicly available Domain Name System (DNS) name, such as support.microsoft.com. Register the public DNS name with an authorized Internet name authority.

- Update your name resolution system (typically DNS) with a new record that contains the new IP address and name of the site. For more information, see Domain Name Resolution.

- Determine whether you need to use a non-standard TCP port number. Standard Internet services use TCP port 80 by default and it is generally not recommended to use any other port for HTTP services. However, if you use a non-standard TCP port number to identify a new Web site for special situations (private Web sites frequently use alternative TCP port numbers for development and testing purposes), select a TCP port number above 1023 so it does not conflict with well-known port numbers assigned by the Internet Assigned Numbers Authority.

- Use Windows Explorer to create a home directory for the Web site content. Create subdirectories to store HTML pages, image files, and other content as needed. To organize home directories for multiple Web sites on the same server, you can create a top-level directory for storing all home directories, and then create subdirectories for each site. You can create a home directory on the local server, as a Universal Naming Convention (UNC) path on a network share, or as a URL that redirects clients to a different Web server. You can also create virtual directories that map to the physical directories. For more information, see Setting Home Directories.

- Determine whether to generate the Web site identification number incrementally, or from the Web site name. When you create a new site on IIS 6.0, a Web site identification number is randomly generated using the name of the Web site. With IIS 5.1 and earlier versions, site identification numbers were generated incrementally. For example, because the default Web site is created first, its Web site identification number is 1; and the next site created is identified as 2. For more information, see **IncrementalSiteIDCreation** in Global Registry Entries.

- Create a home page that clearly identifies the new Web site.

# Configuring User Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Authenticating users early in the connection process reduces the amount of information that can be gained about your IIS 6.0 solution through anonymous access. In highly secure Web sites and applications, you need to require authentication before the user can access any of the information. For other Web sites and applications, you might require authentication only when the user is going to access portions of the Web site or application that contain confidential information.

Figure 3.5 illustrates the process for configuring user authentication.

**Figure 3.5 Configuring User Authentication**




Based on the types of services provided by the Web server, you might need to provide user authentication for Web sites or File Transfer Protocol (FTP) sites. Web sites and FTP sites support different methods for authenticating users.

# Configuring Web Site Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Web site authentication methods you choose are determined by whether you deploy your Web sites and applications on an intranet or the Internet and the relation of the users accessing the Web sites and applications to your organization.

**Intranet-based Web sites and applications**
With intranet-based Web sites and applications, you can typically mandate the type of authentication that is used because the clients are controlled by your organization. In most instances, use Integrated Windows authentication to provide single sign on for users and the strongest possible protection of user credentials. If you are unable to use Integrated Windows authentication, then select the next strongest authentication method from the authentication methods that are described in Table 3.9 in Selecting a Web Site Authentication Method later in this section.

**Internet-based Web sites and applications**
With Internet-based Web sites and applications, you typically need to support a broad range of client operating systems and browsers because users outside your organization own the clients. Anonymous access and Basic authentication are the most common authentication methods used for Internet-based Web sites and applications. For more information about these authentication methods, see Selecting a Web Site Authentication Method later in this section.

Configure Web site authentication by completing the following steps:

1. Select the authentication methods for each Web site.

2. Configure the authentication methods for each Web site.

# Selecting a Web Site Authentication Method

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The authentication method that you select varies based on the level of protection it provides for user credentials (user account and password information), and on the relationship of the user to your organization. Select the strongest authentication method possible to help ensure that the credentials of your users are protected.

Web site authentication methods can be divided into two categories:

- Methods that do not require or encrypt user credentials
- Methods that encrypt user credentials

**Authentication Methods that Do Not Require or Encrypt User Credentials**

For Internet-based Web sites and applications, the most commonly used authentication methods do not require or encrypt user credentials. Select one of these authentication methods when any combination of the following is true:

- Access to the Web sites and applications needs to be anonymous.
- Access to the Web sites and applications needs to be independent of the client configuration and relationship of the user to your organization.

**Anonymous access**

Anonymous access requires no authentication whatsoever. Anonymous access is used for intranet and Internet Web sites when you want unauthenticated users to be able to access the information provided by the Web sites and applications. The majority of Internet Web sites use anonymous access.

**Authentication that is independent of the client configuration and user**

Use Basic authentication when you want to require authentication to access a Web site or application, but need to use an authentication method that provides any combination of the following:

- No special configuration is required on the client. For example, the client can run any operating system or Web browser.
- The users have no close affiliation with your organization, and they are typically not employees or employees of partner organizations. As a result, you cannot require them to use an authentication method that encrypts user credentials.
- Your network infrastructure does not support encrypted authentication methods. For example, to use Windows Integrated authentication, you need the connection between requests to be maintained. Most proxy servers do not support such keep-alive connections.

Because Basic authentication does not encrypt user credentials, use a Secure Sockets Layer (SSL)-secured channel to encrypt user credentials. If you cannot encrypt Basic authentication traffic by using a SSL-secured channel, use one of the authentication methods that encrypt user credentials. These authentication methods are described in Table 3.9.

**Authentication Methods That Encrypt User Credentials**

Authentication methods that encrypt user credentials typically require some level of control over the client computer, and the user is usually an employee of your organization or a partner organization. Table 3.9 provides a comparison of the Web site authentication methods for Web sites and applications that encrypt user credentials.

**Table 3.9 Web Site Authentication Methods That Encrypt User Credentials**

| Authentication Method | Advantages | Disadvantages |
|---|---|---|

| | | |
|---|---|---|
| **Digest** | • Supports authentication through firewalls and proxies.<br>• Encrypts user credentials.<br>• Requires Active Directory running on Microsoft Windows® 2000 Server or later.<br>• Provides medium security. | • Requires Microsoft Internet Explorer 5.0 or later.<br>• Stores user password unencrypted in Active Directory.<br>• Cannot be used to authenticate local accounts.<br>• Requires the associated Application Pool identity to be configured as LocalSystem. |
| **Advanced digest** | • Supports authentication through firewalls and proxies.<br>• Encrypts user credentials.<br>• Stores hash of the user credentials in Active Directory.<br>• Provides medium security. | • Requires Internet Explorer 5.0 or later.<br>• Requires Active Directory running on Windows Server 2003.<br>• Cannot be used to authenticate local accounts. |
| **Integrated Windows** | • Encrypts user credentials.<br>• Provides high security.<br>• Requires Internet Explorer 2.0 or later. | • Requires Microsoft clients. |
| **Client Certificates** | • For server authentication (certificates stored on the server), your organization obtains certificates from a trusted certification authority.<br>• For client authentication, map certificates to user accounts stored in Active Directory running on Windows 2000 Server or later.<br>• Provides high security. | • For client authentication (certificates stored on the clients), your organization has, or is willing to deploy, a public key infrastructure (PKI).<br>• For client authentication, you have a method of securely distributing the certificates to the clients. |
| **Microsoft .NET Passport** | • Supports authentication through firewalls and proxies.<br>• Encrypts user credentials.<br>• Requires Internet Explorer 4.0 or later and Netscape Navigator 4.0 or later. | • Requires Active Directory when account mapping is used.<br>• Requires your organization to license the .NET Passport authentication service. |

# Configuring the Web Site Authentication Method

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you select the Web site authentication method for each Web site, you need to configure the Web site to use that method. For more information about how to configure Web server authentication, see Configure Web Server Authentication.

# Configuring FTP Site Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Configure FTP site authentication for your FTP server by completing the following steps:

1. Select the FTP site authentication method that fulfills the security requirements of your organization, based on the information in Table 3.10.

   The authentication methods that you select vary, based on the ability of the method to protect user credentials (user account and password information). Select the strongest authentication method possible to help ensure that the credentials of your users are protected. Table 3.10 lists and describes these FTP site authentication methods.

   **Table 3.10 FTP Site Authentication Methods**

| Authentication Method | Description |
|---|---|
| Anonymous FTP authentication | For transferring any confidential information, avoid using Anonymous FTP authentication because no user authentication is performed with Anonymous FTP authentication, and anyone can transfer the information. |
| Basic FTP authentication | Basic FTP authentication sends the user name, password, and data in plaintext and can easily be discovered. |

2. Configure the FTP server to use the FTP site authentication method that you selected in the previous step.

When you are transferring confidential data:

- Use FTP within your intranet. Internet Protocol security (IPsec) is required between the client computers and the FTP server to encrypt the user name, password, and any data transferred.

- Use FTP outside your intranet (for remote users). A virtual private network (VPN) tunnel is required between the client computers and your intranet to encrypt the user name, password, and any data transferred.

For information about how to configure FTP server authentication, see Configure FTP Server Authentication.

# Encrypting Confidential Data Exchanged with Clients

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Your business needs might require that confidential data be exchanged between the client computers and the Web server. You can help ensure that this information is safeguarded on the network by using *encryption*. Encryption is a cryptographic process that helps prevent unauthorized users from viewing the encrypted data.

Figure 3.6 illustrates the process for encrypting confidential data that is exchanged between client computers and the Web server.

**Figure 3.6 Encrypting Confidential Data Exchanged with Clients**




The method you select for encrypting the data exchanged between the client computers and the Web server is based on a number of factors. You can encrypt the data exchanged between clients by using:

- Secure Sockets Layer (SSL) for users accessing the Web sites and applications hosted on the Web server.
- Internet Protocol security (IPsec) or VPNs for administrators who remotely manage the Web sites and applications hosted on the Web server.

# Using SSL to Encrypt Confidential Data

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure Secure Sockets Layer (SSL) security features on your Web server to encrypt network transmissions, which will help ensure the integrity of your data transmission, and to verify the identity of users. SSL can be configured to provide security for any portion of the Web sites or applications on the Web server.

| **Note** |
| --- |
| The process presented here describes how to configure SSL to use *server certificates*. Server certificates are installed on the Web server and typically require no additional configuration on the clients. Server certificates allow the clients to verify the identity of the server. Alternatively, some Web sites and applications might require *client certificates*. Client certificates are installed on the client and allow the server to authenticate the clients. For more information about configuring client certificates, see Enabling Client Certificates in IIS 6.0. |

You can configure SSL to help protect confidential data on a URL-by-URL basis (individual portions of the Web site or application). One portion of the application might require encryption of data transmissions with SSL (by specifying HTTPS in the URL), while another portion of the application might allow unencrypted data transmission (by specifying HTTP in the URL). This flexibility in security configuration allows you to provide encryption of confidential data as required, which is unlike IPsec and VPNs because they require that you encrypt all traffic between the clients and the Web server.

As an example, consider a fictitious organization called Contoso Pharmaceuticals that has an e-commerce Web site on the Internet. The Web site contains both secured and unsecured content. The URL for the unsecured home page is http://www.contoso.com. The URL for the secured e-commerce portion of the Web site is https://purchase.contoso.com. Traffic between clients and the home page is unencrypted; whereas SSL encrypts traffic between clients and the e-commerce portion of the Web site.

To use SSL, you must install a valid server certificate on the Web server for each Web site that you want to use with SSL. Certificates are usually granted to organizations through trusted certification authorities. Part of the information that is contained in a certificate is information about the organization to which the certificate was granted, such as the registered domain name. Thus, Web sites with registered domain names need their own certificate.

You can also host multiple SSL-enabled Web sites, each with a different server certificate, on one IP address. To set up this scenario, configure more than one IP address-port combination for the IP address and assign a different certificate to each combination. When an https:// request comes to the Web server, IIS looks up the IP address-port combination to determine which certificate to use.

Client browsers perform a number of verification checks on the SSL certificate. When a client browser detects an incorrect value in the certificate, the browser displays warning messages. The client browser can verify the following:

- Digital signature of the certificate
- Expiration date of the certificate
- Registered domain name to which the certificate was issued against the URL requested

To enable SSL, complete the following steps for each Web site and application:

1. Request a server certificate for the Web site from a certification authority.

   You can use the Web Server Certificate Wizard either to generate a certificate request file (Certreq.txt, by default) that you send to a certification authority, or to generate a request for an online certification authority, such as Microsoft Certificate Services in Windows Server 2003. Depending on the level of identification assurance offered by your server certificate, you can expect to wait several days to several months for the certification authority to approve your request and send you a certificate file.

   For more information about requesting a server certificate by using the Web Server Certificate Wizard, see Request a Server Certificate.

2. Install the server certificate to be used by the Web site on the Web server.

   For more information about installing the server certificate on the Web server by using the Certificate MMC snap-in, see Install a Server Certificate.

3. Assign the server certificate to the Web site.

   For more information about assigning the server certificate to the Web site, see Assign a Server Certificate to a Web Site.

# Using IPsec or VPN with Remote Administration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use Internet Protocol security (IPsec) and virtual private networks (VPNs) to help secure traffic for Web server administration tasks performed over the network, such as uploading content by using FTP or managing the Web server. In addition to encrypting administrative traffic, IPsec and VPNs provide cryptographically strong authentication methods, such as computer certificates, smart cards, and strong password requirements, to provide improved *identity checking*.

Identity checking is the process of verifying the authenticity of the user credentials. Computer certificates, issued by your organization, ensure that remote administration is performed from specific computers and provide improved identity checking.

If your organization provides hosting of Web sites for other groups, these groups need to have a secure method for posting their Web site content to your Web servers. Because the groups post their Web site content over the Internet, you need to encrypt the traffic and provide enhanced identity checking to help protect confidential information.

For example, you can use FTP to upload Web site content to be published on production Web servers. Because FTP exchanges content and user credentials in plaintext, you need to use IPsec or VPNs to encrypt the traffic.

For more information about designing and deploying VPNs, see Determining Your Ipsec Needs.

# Maintaining Web Site and Application Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After securing the Web sites and applications on your Web server, you need to help ensure that the Web sites and applications stay secure. You need to deploy Web servers that are easy to manage and operate. As you deploy the Web server, keep in mind the operations processes that must be performed after the Web server is deployed.

Figure 3.7 illustrates the process for maintaining the security of your Web sites and applications.

**Figure 3.7 Maintaining Web Site and Application Security**




For more detailed coverage of security operations processes, see Managing a Secure IIS 6.0 Solution.

# Obtaining and Applying Current Security Patches

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You should always evaluate and apply the latest security updates to help ensure that your Web sites and applications remain secure. These security updates are published as service packs or hotfixes. As new security vulnerabilities are discovered, Microsoft publishes updates to help mitigate any security risks they might cause. You need to apply these security updates to help ensure that your Web server is protected from the most current security risks.

Stay current with security updates by completing the following steps:

1. Obtain the current security updates by using any combination of the following:

   - Subscribe to the Microsoft Security Notification Service newsletter. The Microsoft Security Notification Service newsletter is a free subscription-based service that sends notification e-mails about available security updates to administrators.

     To subscribe to the Microsoft Security Notification Service newsletter, see the Microsoft.com profile Center.

   - Run Windows Update on a regular basis. Windows Update is a service that runs on Windows-based computers. Windows Update scans the local computer and identifies any updates that are applicable for the software installed on the computer. Windows Update is installed on Windows Server 2003 by default. You must manually start Windows Update on the Web server from Help and Support Center for Microsoft® Windows® Server 2003.

     For more information about running Windows Update, see "Windows Update" in Help and Support Center for Windows Server 2003.

   - Deploy Microsoft Software Update Services (SUS). SUS is a service that acts as an intermediary between the Windows Update server on Microsoft.com and the Windows-based computers in your organization running Windows Update. By using SUS, you can download the latest updates to a server on your intranet, test the updates on test servers, select the updates that you want to deploy, and then deploy the updates to computers within your organization.

     For more information about deploying SUS, see Software Update Services Deployment White Paper in Designing a Managed Environment in the Windows Server 2003 Deployment Kit.

Table 3.11 lists the options for obtaining security updates, and describes the advantages and disadvantages of each option.

**Table 3.11 Options to Obtain Security Updates**

| Option | Advantages | Disadvantages |
|---|---|---|
| Microsoft Security Notification Service Newsletter | • Does not require Web servers to be directly connected to the Internet.<br>• Does not require a dedicated server.<br>• Is free. | • Is not specific to a particular technology, such as IIS.<br>• Is not specific to a particular operating system version.<br>• Requires administrators to manually review newsletters for recommended updates. |

| | | |
|---|---|---|
| Windows Update | • Provides automatic notification of available updates.<br>• Is free. | • Requires the Web server to have Internet access. |
| SUS | • Provides automatic notification of available updates. | • Requires a dedicated server to run properly.<br>• Requires the SUS server be able to access the Internet.<br>• Requires separate purchase of SUS. |

2. Test the security updates on a Web server in a test environment.

   Before deploying the security updates on your production Web server, use one of the options described in Step 1 to test the security updates on a test Web server that is configured identically to your production Web server. Table 3.12 lists the options for deploying the security updates on a Web server in a test environment.

**Table 3.12 Options for Deploying Security Updates**

| Option | Deployment |
|---|---|
| Microsoft Security Notification Service Newsletter | Manually download the updates and then deploy them manually or automatically by using a software distribution program, such as Microsoft System Management Server. |
| Windows Update | Configure Windows Update to do one of the following:<br>• Inform an administrator that is logged on to the Web server that updates are available and then allow the administrator to install the updates.<br>• Install updates automatically. |
| SUS | Configure the SUS to provide updates to the Web server through an updated version of Windows Update called Automatic Updates. |

You can configure Windows Update and Automatic Updates in SUS to install updates automatically, with or without confirmation, based on the security rating of the update. Table 3.13 lists the security ratings used by Windows Update and Automatic Updates, and provides a description of each rating.

**Table 3.13 Security Ratings Used by Windows Update and Automatic Updates**

| Rating | Description |
|---|---|
| Critical | A vulnerability that, if exploited, might allow the propagation of an Internet worm without user action. |
| Important | A vulnerability that, if exploited, might result in a compromise of the confidentiality, integrity, or availability of user data, or of the integrity or availability of processing resources. |
| Moderate | A vulnerability risk that can be mitigated by factors such as default configuration, auditing, or difficulty to exploit. |

| | |
|---|---|
| Low | A vulnerability that is extremely difficult to exploit, or that has minimal impact. |

3. Deploy the security updates to your production Web server by using the same option that you tested on the test Web server.

# Enabling Windows Server 2003 Security Logs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Collecting information about the security aspects of the Web server is required to help ensure that the Web server stays secure. Windows Server 2003 uses security and system logs to store collected security events. The security and system logs are repositories for all events recorded on the Web server. Many management systems, such as Microsoft Operations Manager, periodically scan these logs and can report security problems to your operations staff.

If you audit or log too many events, the log files might become unmanageable and contain superfluous data. Before enabling the system and security logs, you need to enable auditing for the system log and establish the number of events that you want recorded in the security log. You cannot change the information that is logged in the system log: These events are preprogrammed into Windows Server 2003 services and applications. You can customize system log events by configuring *auditing*. Auditing is the process that tracks the activities of users and processes by recording selected types of events in the security log of the Web server. You can enable auditing based on categories of security events. At a minimum, enable auditing on the following categories of events:

- Any changes to user account and resource permissions
- Any failed attempts for user logon
- Any failed attempts for resource access
- Any modification to the system files

You can customize which types of events are recorded in the security log. The most common security events recorded by the Web server are associated with user accounts and resource permissions.

For more information about how to enable security auditing, see Enable Security Auditing.

# Enabling File Access Auditing for Web Site Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to enabling Windows Server 2003 security logs, enable file access auditing for your Web site content. This is a separate step that must be completed to monitor any changes to the files and directories that contain your application and content.

You can enable auditing on a user-by-user basis for each file and directory. However, at a minimum, enable auditing for all users for any successful or failed attempts to do the following:

- Modify or delete existing content.
- Create new content.

| **Tip** |
| --- |
| Beyond these minimal events, you can audit content for other purposes, such as forensic analysis of intruder detection. |

For more information about how to enable file auditing on Web site content and files, see Enable Web Site Content Auditing.

# Configuring IIS Logs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to the Windows Server 2003 system and security logs, you should configure IIS to log site visits. When users access your server running IIS 6.0, IIS logs the information. The logs provide valuable information that you can use to identify any unauthorized attempts to compromise your Web server.

Depending on the amount of traffic to your Web site, the size of your log file (or the number of log files) can consume valuable disk space, memory resources, and CPU cycles. You might need to balance the gathering of detailed data with the need to limit files to a manageable size and number. If you are planning to put thousands of Web sites on one Web server with high traffic volumes and disk writes, you might want to use centralized binary logging to preserve server resources. Also, consider limiting log size by changing the frequency of log file creation. For more information, see Saving Log Files.

The IIS logs allow you to record events for each application and Web site on the Web server. You can create separate logs for each of your applications and Web sites. Logging information in IIS 6.0 goes beyond the scope of the event logging or performance monitoring features provided by Windows. The IIS logs can include information, such as who has visited your site, what the visitor viewed, and when the information was last viewed. You can use the IIS logs to identify any attempts to gain unauthorized access to your Web server.

IIS 6.0 supports different log formats for the IIS logs that you enable. IIS 6.0 supports the following log formats.

**W3C Extended log file format**
World Wide Web Consortium (W3C) Extended format is a customizable ASCII format with a variety of different properties. You can log properties that are important to you, while limiting log size by omitting unwanted property fields. Properties are separated by spaces. Time is recorded as Universal Time Coordinate (UTC).

For information about customizing this format, see W3C Extended Log File Format. For more information about the W3C Extended format specification, see the W3C Extended Log File Format on the W3C World Wide Web Consortium Web site.

**IIS log file format**
IIS log file format is a *fixed* (meaning that it cannot be customized) ASCII format. This file format records more information than other log file formats, including basic items, such as the IP address of the user, user name, request date and time, service status code, and number of bytes received. In addition, IIS log file format includes detailed items, such as the elapsed time, number of bytes sent, action (for example, a download carried out by a **GET** command), and target file. The IIS log file is an easier format to read than the other ASCII formats because the information is separated by commas, while most other ASCII log file formats use spaces for separators. Time is recorded as local time.

For more information about the IIS log file format, see IIS Log File Format.

**NCSA Common log file format**
National Center for Supercomputing Applications (NCSA) Common log file format is a fixed ASCII format that is available for Web sites, but not for FTP sites. This log file format records basic information about user requests, such as remote host name, user name, date, time, request type, HTTP status code, and the number of bytes sent by the server. Items are separated by spaces. Time is recorded as local time.

For more information about the NCSA Common log file format, see NCSA Common Log File Format.

**ODBC logging**
Open Database Connectivity (ODBC) logging format is a record of a fixed set of data properties in a database that complies with ODBC, such as Microsoft Access or Microsoft SQL Server. Some of the items logged include the IP address of the user, user name, request date and time (recorded as local time), HTTP status code, bytes received, bytes sent, action carried out (for example, a download carried out by a **GET** command), and the target file. With ODBC logging, you must both specify the database to be logged to, and set up the database to receive the data.

| **Note** |
| --- |
| When ODBC logging is enabled, IIS disables the kernel-mode cache. As a result, overall server performance can be degraded. |

For more information about ODBC logging, see ODBC Logging.

**Centralized binary logging**

Centralized binary logging is the process of multiple Web sites writing binary, unformatted log data to a single log file. Each Web server running IIS creates one log file for all of the Web sites on that server. Centralized binary logging preserves valuable memory resources. Depending on your configuration, you can see dramatic performance and scalability gains with centralized binary logging.

For more information about centralized binary logging, see Centralized Binary Logging.

For more information about how to configure IIS logs, see Enable Logging_IIS_SP1_Dep. For more information about logging Web site activity, see Logging Site Activity. For more information about managing IIS logs, see Analyzing Log Files.

# Reviewing Security Policies, Processes, and Procedures

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

As a part of maintaining the security of your Web server, you must perform periodic reviews of the security policies, processes, and procedures in use by your organization. Review your security practices for any changes that might affect the security of the Web server. These changes in security practices can include the following:

- Ensuring that any recent security risks are mitigated:
  As new security risks are identified, such as new viruses, you need to ensure that your security practices help mitigate these risks. If your current security practices do not address the new risks, then modify them to help mitigate the risks.

- Identifying changes in Web server configuration that can compromise security:
  Through the course of normal administration of the Web server, configuration changes are made. During this process, security settings might have been inadvertently changed. You need to periodically review the configuration of the Web server to ensure that it complies with the security requirements of your organization.
  You can categorize these Web server security practices by their function, such as operating system security, security policies, firewall security, and router security. In addition, the frequency with which these processes and procedures are completed varies. Some security practices need to be completed continuously while others might be completed monthly.

Table 3.14, Table 3.15, Table 3.16, and Table 3.17 list examples of security policies, processes, and procedures for an ISP, grouped by categories. These examples are representative of the types of security practices that are required to maintain the security of your Web server. For more information about the security policies, processes, and procedures for your Web server, see Managing a Secure IIS 6.0 Solution.

**Table 3.14 Windows Server 2003 Operating System Security**

| Security Policy, Process, or Procedure | Frequency |
|---|---|
| Limit user rights to only those that are required. | Constant |
| Limit any windows for vulnerabilities that can be exploited when deploying new servers. | Constant |
| Limit Terminal Services access to only necessary accounts. | Constant |
| Run a two-tier DNS structure to protect the identity of internal servers. | Constant |
| Run an intrusion detection system. | Constant |
| Scan the ports in use on your server addresses and addresses assigned to remote users. | Daily |
| Review event and IIS logs. | Weekly |
| Test firewalls from inside and outside by using port scanners and other appropriate tools. | Weekly |

**Table 3.15 Windows Server 2003 Policy Security**

| Security Policy, Process, or Procedure | Frequency |
|---|---|
| Explicitly deny interactive logon user right to all non-administrative accounts. | Constant |
| Explicitly deny "Allow logon through Terminal Services" user right to all non-administrative accounts. | Constant |
| Enable FULL (Success/Failure) auditing on domain **Group Policy** objects. | Constant |
| Send event notification when events like **User added to Domain Administrators** occur. | Constant |

| | |
|---|---|
| Allow only Administrators to have write permissions on all content servers. | Constant |
| Require strong passwords for all users. | Constant |
| Require smart cards for all administrators. | Constant |
| Allow administrators to log on only to specific workstations. | Constant |
| Enable account lockout policies for failed logon attempts. | Constant |
| Audit the domain **Group Policy** object. | Monthly |
| Audit Active Directory user rights. | Monthly |
| Audit all servers to determine if nonessential services are running. | Monthly |

**Table 3.16 Firewall and Router Security**

| Security Policy, Process, or Procedure | Frequency |
|---|---|
| Restrict the network segments where management traffic is allowed. | Constant |
| By default, deny IP traffic and log any failed attempts. | Constant |
| Ensure that the minimal firewall rules are enforced, including:<br>• Explicitly deny all traffic to the following:<br>  • TCP and UDP ports 135-139, 445 (NetBIOS/SMB)<br>  • TCP and UDP ports 3389 (Terminal Services)<br>  • Domain controllers<br>  • Internal DNS servers<br>• Permit traffic to TCP and UDP port 53 (DNS) to external DNS servers. | Constant |

**Table 3.17 Miscellaneous Security**

| Security Policy, Process, or Procedure | Frequency |
|---|---|
| Run virus scans on all servers. | Constant |
| Monitor security distribution lists and newsgroups for potential security issues. | Constant |
| During virus outbreaks, block any suspicious content (such as e-mail attachments). | Constant |
| Monitor the number of Non-Delivery mail reports generated (indicates e-mail spamming). | Weekly |
| Monitor Simple Mail Transfer Protocol (SMTP) relay attempts that are not valid (indicates e-mail spamming). | Weekly |
| Audit accounts to determine the users who are no longer employed at the organization, partner organizations, or customer organizations. | Monthly |

# Additional Resources About Security in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this section.

**Related Information**

- Deploying ASP.NET Applications in IIS 6.0 for information about ASP.NET-specific deployment considerations.
- Ensuring Application Availability for information about balancing application security and availability.
- Using IPsec or VPN with Remote Administration for information about designing and deploying VPN.
- Deploying IPsec for information about designing and deploying IPsec.
- Software Update Services Deployment White Paper for information about deploying SUS.
- Planning a Secure Environment for information about securing other services on a multipurpose server.
- Analyzing Log Files for information about managing IIS logs.
- Managing a Secure IIS 6.0 Solution for information about maintaining Web server security.
- Microsoft.com Profile Center for information about how to subscribe to the Microsoft Security Notification Service newsletter.
- W3C World Wide Web Consortium for information about the W3C Extended format specification.

**Related IIS 6.0 Help Topics**

- Enabling Client Certificates in IIS 6.0 for information about configuring client certificates.
- Best Practices for Securing Files with NTFS Permissions for information about the benefits of formatting disk volumes as NTFS on Web servers.
- Saving Log Files for information about balancing the gathering of detailed data with the need to limit files to a manageable size and number.
- SMTP Administration or NNTP Administration for information about securing SMTP or NNTP.
- Troubleshooting in IIS 6.0 for information about troubleshooting problems related to Web sites and applications that are not functioning.

**Related Windows Server 2003 Help topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- See "Windows Update" in Help and Support Center for Windows Server 2003 for information about using Windows Update.

# Ensuring Application Availability

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Application availability is the readiness of an application (and the service it runs under) to handle customer requests and to return timely and accurate responses. Ensuring consistently high application availability requires that you set availability goals that meet your business needs, configure your Web server to achieve those goals, and measure the compatibility and reliability of your applications. Internet Information Services (IIS) 6.0 running in worker process isolation mode provides the advanced features that you need to achieve and maintain a high degree of application availability.

**In This Section**

Overview of the Ensuring Application Availability Process

Establishing Application Availability Goals

Configuring IIS 6.0 for Optimum Availability

Testing Applications for Compatibility

Additional Resources_IIS_SP1_Dep

# Overview of the Ensuring Application Availability Process

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To help ensure optimum availability of your applications to users who need them, you need to do more than ensure that the network is available and that your server is up and running. You need to set availability goals, configure IIS to meet the demands that users are placing on your applications, and test your applications for functional compatibility with IIS 6.0 worker process isolation mode.

To be certain that the level of application availability on your server meets the needs of your customers; you must first establish availability, service, and request-handling goals that accurately represent those customer needs. Next, you need to create application pools, and then configure IIS features to isolate applications and to tune and monitor your application pools. Finally, you need to assess the tradeoffs between availability and other aspects of running your applications, such as performance.

To help maintain high availability of the applications that your server hosts, you can use two kinds of testing: compatibility testing and comparison testing. Functional compatibility testing verifies that your applications will work with the IIS 6.0 worker process isolation mode features that you have enabled. Comparison testing measures the availability of your applications against your application availability goals, which reveals how closely the day-to-day availability of applications hosted on your server matches your original goals.

Before you begin this process, you must:

- Install the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; or Windows® Server 2003, Web Edition operating system with the default options.

- Install IIS 6.0 with the default settings by using Add or Remove Programs in Control Panel.

If you use other methods for installing and configuring Windows Server 2003, such as unattended setup, or enabling IIS 6.0 by using Manage Your Server, then the default configuration settings might not be identical to the configuration settings described in this section.

When you complete the process for ensuring application availability, you will meet the basic requirements for achieving high availability for your applications. You will set and begin to measure availability goals for your applications based on your business needs. Your applications will be configured to make effective use of IIS 6.0 features. Also, you will know whether your applications are compatible with IIS 6.0 worker process isolation mode.

# Process for Ensuring Application Availability

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The process for ensuring application availability includes setting availability goals, configuring IIS to achieve your availability goals, and testing your applications to ensure that they meet the availability goals that you set. Ensure that your availability goals accommodate the needs of your users for high application reliability and reasonable application response times. Figure 4.1 shows the process for ensuring application availability.

**Figure 4.1 Ensuring Application Availability**




*Application pools* are sets of applications and the worker processes that service them. For you to create and configure application pools and make application pool assignments, IIS 6.0 must be running in worker process isolation mode. Thus, you need to evaluate the compatibility of your applications with IIS 6.0 worker process isolation mode in a test environment before deploying them in your production environment.

| **Important** |
| --- |
| The application-pool functionality is only available in IIS 6.0 worker process isolation mode. |

The following quick-start guide provides a detailed overview of how to help ensure the availability of your applications. You can use this guide to help identify the steps of the ensuring application availability process that you need additional information about to complete so that you can disregard the information with which you are already familiar. In addition, for the procedures that are required to complete this process, see Appendix A: IIS Deployment Procedures.

**Establishing Application Availability Goals**

1. Set service availability goals.
2. Set request-handling reliability goals.

**Configuring IIS 6.0 for Optimum Availability**

1. Isolate applications by completing the following steps:
   - Determine the application isolation needs of your server.
   - Create application pools and assign applications to them.
2. Recycle worker processes in one of the following ways:
   - Recycle by elapsed time.
   - Recycle by number of requests.

- Recycle at scheduled times.
- Recycle on a virtual-memory threshold.
- Recycle on a used-memory threshold.

3. Tune performance by completing the following steps:

    - Configure idle time-out for worker processes.
    - Configure a request queue limit.
    - Enable HTTP compression.
    - Configure Web gardens.
    - Set processor affinity on servers that include multiple CPUs.

4. Manage application pool health by completing the following steps:

    - Configure worker process pinging.
    - Configure rapid-fail protection for worker processes.
    - Configure the startup time limit for worker processes.
    - Configure the shutdown time limit for worker processes.
    - Enable debugging for application pool failures.

5. Configure application pool identity.

**Testing Applications for Compatibility**

1. Test applications for setup compatibility with IIS 6.0.
2. Test applications for functional compatibility with IIS 6.0.

# Establishing Application Availability Goals

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To set effective availability goals for your applications, you need to measure the availability of the World Wide Web Publishing service (WWW service), your Web sites, and the worker processes that service your applications. These availability measurements can be calculated as a percentage by dividing the uptime of the WWW service, Web site, or worker process by the time elapsed. In addition, a good user application availability experience depends upon successful handling of correctly formed requests, which are requests that are complete and error-free.

At times, users cannot access your applications because of these common request-handling problems:

- Correctly formed requests are sent, but they are not correctly served.
- Correctly formed requests are sent, but they are served only after a time interval that is unacceptable to the user.

To ensure that you set application availability goals that represent the needs of your users, you must first define how you will measure the successful serving of requests. Make sure that your goals address the following:

- Service availability, which includes the WWW service and individual Web sites.
- Application availability, which includes request processing reliability and the average time that is needed to process requests.

Figure 4.2 shows the process for establishing application availability goals.

**Figure 4.2 Establishing Application Availability Goals**



# Setting Service Availability Goals

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To set effective service availability goals, you need to measure the percentage of time that the WWW service and your Web sites are available to process client requests. WWW service availability is the time that the WWW service and your Web sites are up and running, expressed as a percentage of a set period of time. For example, an organization might set a goal to have a WWW service availability of 99.999 percent over a 21-day period.

You can achieve high WWW service availability when running IIS 6.0 in worker process isolation mode. This is because worker process isolation mode isolates the WWW service from the impacts of failed applications. Generally, the WWW service is unavailable only when the system is restarted and when IIS is taken offline for major service upgrades, such as the installation of a service pack. When the WWW service is available, your Web sites are available unless you deliberately take them offline.

Test service availability and compare the test results to your goals. First, run the tests in a test environment, and then run the tests periodically in your production environment.

**Measuring Service Availability**

You can measure WWW service availability by completing the following steps:

1. Use Performance Monitor to gather the data for WWW service uptime, which measures uptime for W3svc.exe, and uptime for each Web site on the server. The data for WWW service uptime is displayed in the Total instance of the Service Uptime counter. The data for Web site uptime is displayed in the Sites instances of the same counter. For more information about gathering uptime data for the WWW service and Web sites, see Gather and Display WWW Service Uptime Data.

2. Use Service Monitor to display the data for the Service Uptime counter. The Service Monitor displays the uptime data for the WWW service and for each Web site on the server. For more information about displaying WWW service and Web site uptime data, see Gather and Display WWW Service Uptime Data.

3. Calculate the availability of the WWW service by dividing the time (in seconds) for the Total instance by the elapsed time that the server was running (in seconds) during the test; multiply the result by 100 to express the figure as a percent availability.

4. Calculate the availability of each Web site by dividing the time (in seconds) for the Site instance for each site by the elapsed time (in seconds) that the server was running during the test; multiply the result by 100 to express the figure as a percent availability.

5. Compare the results to your goals.

# Setting Request-Handling Goals

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Request handling measures whether client requests were served correctly and the time needed for the WWW service to process client requests. There are two important measures for request handling: *request-handling reliability* and *request processing time*. Request-handling reliability measures the percentage of correctly formed requests that are returned successfully; request processing time measures the average elapsed time for IIS to process a client request.

**Request-Handling Reliability**
For a range of request rates that your applications receive, set goals to achieve the level of overall request handling reliability that you want. Measure only correct requests; do not consider a request that a client composed incorrectly as a failed request.

For example, you can establish a request processing reliability goal of 99.999% for a rate of 20 correctly formed requests per second. You meet this goal when a Web server that receives requests at a rate of 20 per second successfully serves 99,999 responses out of each 100,000 requests. A response may be considered successful if the request returns the expected message, such as 200 OK, and the correct content.

When testing your applications in a test environment, you can use a request rate such as 20 requests per second. However, when testing applications in a production environment, you might need to set a range of request rates. If you estimate that 20 requests per second is a reasonable goal, you might try a range of 10-30 requests per second. Also, monitor request handling over a long period of time, such as 21 days or a month, to ensure that you cover the periods of heaviest loads on the application.

The following is an example of a full expression of a request reliability goal for a Web site or application:

- Duration: 21 days
- Load: 10-30 requests per second
- Success rate goal: 99.999 percent responses with the expected return message, such as 200 OK, measured against 100,000 requests

Keep in mind that you must adjust your goals to accommodate your needs, taking into account client expectations, application backend processing, and caching.

**Request Processing Time**
When a request is processed so slowly that the user is not willing to wait for a return, the effect on the user is similar that of a request failure. The user experiences a loss of availability.

The following is an example of a request reliability goal that accommodates request-handling execution time:

- Duration: 21 days
- Load: 10-30 requests per second
- Success rate goal: 99.999 percent responses with the expected return message, such as 200 OK, measured against 100,000 requests
- Average response time of 5 seconds for each request

To set your own goals for request processing time, evaluate the availability requirements of your organization and the equipment and software that you are using.

**Testing Request-Handling**
Test request handling and compare the test results to your goals. First run tests in a test environment, and then periodically run tests in your production environment. Measure request-handling performance for an application by completing the following steps:

1. Send correctly formed requests from a remote client to the target application on the server that is hosting the application that you are testing. If you are testing an application on a production server, measure the performance of requests sent by users. If you are testing an application running on a test server, use automated load-testing software, such as WCAT, to send client requests to the application at a rate appropriate to test it. You can use WCAT to automate sending customized request loads and to emulate a wide variety of conditions. You can use Log parser to create request schemes that WCAT can use. For information about WCAT and Log parser, see IIS 6.0 Resource Kit Tools.

2. Measure the rate at which the correctly formed requests are being processed successfully by the target application on your test server (responses that return 200 messages and correct content).

3. Calculate the average elapsed time to process requests. One way to do this is to average the readings of the TimeTaken counter. The readings are in the IIS log.

4. Compare the results to your goals.

# Configuring IIS 6.0 for Optimum Availability

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Worker process isolation mode allows you to optimize the availability of applications installed on a Web server. To configure worker process isolation mode for optimum availability of your IIS 6.0 applications and Web sites, you need to determine availability requirements for each application and isolate applications into application pools based on those needs. You can then configure each application pool for recycling, health monitoring, performance, and identity to meet your availability requirements.

Figure 4.3 shows the process for configuring IIS for optimum application availability.

**Figure 4.3 Configuring IIS 6.0 for Optimum Availability**



# Isolating Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because applications in an application pool are isolated from the applications in other application pools by worker process boundaries, if an application fails, it does not affect the availability of applications that are running in other application pools. Deploying applications in application pools is a primary advantage of running IIS 6.0 in worker process isolation mode because you can customize the application pools to achieve the degree of application isolation that you need.

Configuring application isolation involves completing the following steps:

- Categorizing applications according to their isolation profiles for initial deployment.
- Creating and naming the application pools.
- Assigning applications to the appropriate application pools.

| **Note** |
| --- |
| When you configure application pools for optimum availability, also consider how to configure application pools for application security. For example, you might need to create separate application pools for applications that require a high level of security, while allowing applications that require a lower level of security to share the same application pool. Also, client requests might be running under identities that are different from the identities that are assigned to the applications that serve them, which can make the application unavailable to customers. For more information about securing applications by configuring applications pools, see Securing Web Sites and Applications. |

If your organization includes network configurations that require more than one server for implementation, such as Web farms, this involves clustering and load balancing. For more information about clustering and load balancing and the specific processes for setting up replication and clustering on IIS, search for "clustering" and "load balancing" in Help and Support Center for Windows Server 2003.

# Determining the Application Isolation Needs of Your Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The applications that you deploy on a server might fit into several categories, based on your assessment of their importance, their demand for server resources, or their history of instability. Group your applications according to the following categories, or make up your own set of application categories, before assigning your applications to application pools:

- Typical: The application is stable and is not subject to extremely high user demand.
- High-demand: The application is expected to receive a large number of requests, putting stress on the resources of your server.
- Mission-critical: Application availability is mandatory for business purposes.
- Problem: The application must be isolated from other applications and Web sites due to known or anticipated unstable behavior. If you must run test versions of an application on the same server that runs production applications and Web sites, consider it a problem application that must be isolated from all of the other applications on the server.
- Unique: The application requires a configuration that is different from other applications that you have categorized as typical. For example, an application might require an identity that has access to more system resources than the default identity.

# Creating Application Pools and Assigning Applications to Them

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After establishing similar groups for your applications and Web sites, you can assign them to application pools according to the degree of application isolation they require.

To create application pools and assign applications to them when you are deploying a set of Web sites or applications of different types:

1. Create one application pool for all of the applications that you categorized as typical.

2. Assign all of the applications categorized as typical to the application pool that you created in Step 1.

3. Create one application pool for each high-demand, mission-critical, problem, and unique application. Name each application pool after the application that it will contain.

4. Assign each of your high-demand, mission-critical, problem, and unique applications to the corresponding application pool that you created and named in Step 3.

   You can add application pools by using IIS Manager, by using a command-line procedure, or by running Active Directory® Service Interfaces (ADSI)– or Windows Management Instrumentation (WMI)–compatible scripts. It is recommended that you use IIS Manager because IIS Manager provides clear, complete, and well-organized visibility to application pools and their properties.

For more information about how to create and name application pools, and how to assign applications to application pools through IIS Manager, see Isolate Applications in Worker Process Isolation Mode. For information about how to assign applications to application pools by using the command line or scripts, see the **AppPoolId** property in the Metabase Property Reference.

If you create a large number of application pools running many worker processes concurrently, server performance may suffer. The key factor to consider is that large numbers of worker processes with dynamic applications (ASP or ASP.NET applications) may use up available memory.

Unless you have specific reasons for altering the settings for an application pool, deploy your application pools using the default settings. For more information about specific reasons for using non-default settings, see Recycling Worker Processes, Tuning Performance, Managing Application Pool Health, and Configuring Application Pool Identity later in this section.

# Recycling Worker Processes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If an application contains code that causes problems, and you cannot easily rewrite the code, it might be useful to limit the extent of the problems by periodically recycling the worker process that services the application. *Worker process recycling* is the replacing of the instance of the application in memory. IIS 6.0 can automatically recycle worker processes by restarting the worker process, or worker processes, that are assigned to an application pool. This helps keep problematic applications running smoothly, and minimizes problems such as memory leaks.

You can trigger the recycling of the worker processes assigned to an application pool by using worker process recycling methods that are based on elapsed time, the number of Hypertext Transfer Protocol (HTTP) requests, a set time of day, and two kinds of memory consumption, in addition to recycling on demand. Table 4.1 lists and describes the worker process recycling methods and when to use each.

**Table 4.1 Worker Process Recycling Methods**

| Method | Description | Usage |
|---|---|---|
| Elapsed time | Recycles the worker process based on a timeframe, in minutes, that is specified by the administrator. | Use this method when applications have problems running for extended time periods. |
| Number of requests | Recycles the worker process when the number of HTTP requests exceeds a threshold that is specified by the administrator. | Use this method when applications are failing based on the number of requests they receive. |
| Scheduled times | Recycles the worker process at specified times within a 24-hour period. | Use this method when conditions are similar to the conditions for elapsed time. |
| Virtual memory | Recycles the worker process when the worker process virtual memory reaches a threshold that is specified by the administrator. | Use this method when the memory heap becomes highly fragmented, which occurs when applications reserve memory multiple times. This is indicated by a steady increase in virtual memory. |
| Used memory | Recycles the worker process when the memory that is used by the W3wp.exe process reaches a threshold that is specified by the administrator. | Use this method when some of your applications have memory leaks. |
| On demand | Recycles the worker process when an administrator uses Microsoft Management Console (MMC) or a script to request the recycling of an entire application pool. | Use this method when an application pool is causing problems while the other sites are up and running. Consider recycling the application pool, as opposed to resetting the entire WWW service. |

By setting the values for worker process recycling methods, you can schedule recycling events to manage application problems in a timely manner. Table 4.2 lists recommendations for adjusting the values of worker process recycling methods.

**Table 4.2 Recommended Settings for Values of Worker Process Recycling Methods**

| Recycling Method | Recommended Setting |
|---|---|
| Elapsed time | Set this value to be less than the length of time elapsed before the worker process fails. |
| Number of requests | Set this value to be less than the number of requests processed before the application fails. |
| Scheduled times | Set this value to be at intervals that are frequent enough to prevent application failure. |

| | |
|---|---|
| Virtual memory | Initially, set this value to be less than 70 percent of available virtual memory. Then monitor the memory usage of the worker process, and adjust the setting if memory usage continues to increase. Calculate this value in megabytes. |
| Used memory | Set this value to be less than 60 percent of available physical memory. Calculate this value in megabytes. |
| On demand | Consider using one of the automated methods for recycling instead, because on-demand recycling is typically performed in response to a failure. |

**Recycling and Maintaining State**

If you have an application pool with applications that depend on state data, you must decide whether or not to recycle the worker processes that are assigned to that application pool. When you recycle a worker process, the state data for applications that is maintained in process is lost. In this case, do not use recycling.

One alternative to increasing application availability by recycling worker processes is to maintain state data external to the worker process, such as in a database. However, moving data to an external database to allow recycling can affect server performance in the following two ways:

- Performance is reduced because of the added data management that is needed to move the data between the application and the database.
- Recycling flushes any in-process data caches, so the caches need to be rebuilt.

ASP.NET gives you the option of persisting session state using a session state service or a SQL database. For more information on deploying ASP.NET applications, see Deploying ASP.NET Applications in IIS 6.0.

**Global Locks and Overlapped Recycling**

Using global locks, configuring files for exclusive read, and other methods of exclusive resource ownership prevent the use of *overlapped recycling*, or the loading of an application in two application pools simultaneously. Exclusive ownership requires the use of non-overlapped recycling and the routing of all requests for that application to the same application pool. Also, you can run only one worker process for that application pool, which means that exclusive ownership does not work in a Web garden. One way to work around this is to write multithreaded applications, or use shared reads, which might eliminate the need for exclusively owned resources.

# Recycling by Elapsed Time

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure an application pool to recycle the worker process, or worker processes, that are assigned to it at specified intervals. By default, recycling by elapsed time is enabled in IIS 6.0, with the elapsed time set to 29 hours. This default setting fits the recycling needs of applications that are problem-free, or that have flaws that cause problems gradually. Recycling worker processes at this interval refreshes applications often enough to correct problems that build up slowly, without creating the severe overhead that can be caused by closely spaced recycling.

Reset the elapsed-time interval for recycling a worker process if your tests or log records indicate that the application that the worker process is assigned to is encountering problems after a given period of uptime. Set the elapsed time to a shorter period of time if problems are building up in less than the default time interval of 29 hours.

Microsoft does not recommend recycling ASP.NET applications by elapsed time.

For information about how to use IIS Manager to configure a worker process to be recycled after a set elapsed time, see Configure Application Pool Recycling. For information about how to configure a worker process to be recycled after a set elapsed time by using a command-line procedure or scripts, see PeriodicRestartTime Metabase Property.

# Recycling by Number of Requests

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure an application pool to recycle the worker process, or processes, that are assigned to it whenever the worker process handles a specified number of requests since the last recycle. By default, recycling by number of requests is disabled in IIS 6.0. When this feature is enabled, the default is 35,000 requests.

Enable recycling by number of requests if your log records or tests indicate that an application you are deploying causes significant problems on your server, or encounters significant problems, after the application handles a certain number of requests. Then, for the worker process assigned to the application pool containing the unhealthy application, set the number of requests to be handled before the worker process is recycled.

For information about how to configure a worker process to be recycled after processing a set number of requests through IIS Manager, see Configure Application Pool Recycling. For information about how to configure a worker process to be recycled after processing a set number of requests by using the command-shell or scripts, see PeriodicRestartRequests Metabase Property.

# Recycling at Scheduled Times

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To avoid recycling worker processes when server demand is heaviest, you can configure an application pool to recycle the worker process, or worker processes, that are assigned to it at a specified time of day, which is based on a 24-hour format. By default, recycling at a scheduled time is disabled in IIS 6.0.

You might need to schedule some, or all, of your applications to recycle worker processes at times of the day when the server load is light. Enable recycling at a scheduled time so that worker processes are rescheduled at times that best fit the load. Be careful to stagger the recycling of your application pools because allowing them to recycle at the same time can degrade the performance of your server during the recycle period.

For information about how to use IIS Manager to configure a worker process to be recycled at scheduled times, see Configure Application Pool Recycling. For information about how to configure a worker process to be recycled at scheduled times by using a command-line procedure or script, see PeriodicRestartSchedule Metabase Property.

# Recycling on a Virtual-Memory Threshold

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A problem application might periodically increase virtual memory without subsequently releasing it. In this way, virtual memory can increase indefinitely. To limit the impact of increased virtual memory, you can set an application pool to recycle the worker processes that are assigned to it whenever virtual memory reaches a specified threshold.

By default, recycling on a virtual-memory threshold is disabled in IIS 6.0. When this feature is enabled, the default virtual-memory setting is 500 megabytes (MB). Enable recycling on a virtual-memory threshold and set a virtual memory threshold that prevents unacceptable degradation. Begin with a setting of 70 percent of available virtual memory.

For information about using IIS Manager to configure a worker process to be recycled after consuming a set amount of memory, see Configure Application Pool Recycling. For information about how to configure a worker process to be recycled after consuming a set amount of memory by using a command-line procedure or scripts, see the PeriodicRestartMemory Metabase Property.

# Recycling on a Used-Memory Threshold

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A problem application might indefinitely increase the amount of memory that it uses without releasing the memory. This is also known as an application memory leak. You can set an application pool to recycle the worker processes that are assigned to it whenever used memory, which is also called private memory, reaches a specified threshold. Used memory is compared to the memory limit set for the worker process that determines when to recycle the worker process.

By default, recycling on a used-memory threshold is disabled in IIS 6.0. When it is turned on, the default used-memory threshold is 192 MB. To limit the impact from an application that increases used memory over time, enable recycling on a used-memory threshold and set a threshold that prevents unacceptable degradation. For example, such degradation might be caused by a memory-bound application that is unable to process requests. Initially, try a threshold of 60% of system memory. When running ASP.NET applications, try a threshold of 60% of system memory, or 800 MB, whichever is less. Also, when running ASP.NET applications, if you increase the virtual address space above the 2-GB default, you can then increase the used-memory threshold a like amount.

For information about using IIS Manager to configure a worker process to be recycled after consuming a set amount of memory, see Configure Application Pool Recycling. For information about how to configure a worker process to be recycled after consuming a set amount of memory by using a command-line procedure or scripts, see the PeriodicRestartPrivateMemory Metabase Property.

# Tuning Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When IIS is running in worker process isolation mode, you can tune the performance of your applications to conserve CPU and memory resources, and to add worker processes to application pools that need them. By setting the application pool properties for idle time-out, you can optimize the time that worker processes remain idle for each application pool, removing them from memory. Another way to conserve memory is to shorten the request queue length limit property for the worker processes in each application pool.

You can increase the number of worker processes that are available to an application pool by configuring a *Web garden*, which is an application pool that is served by more than one worker process. When a worker process is engaged, another worker process can continue to process requests, enhancing server performance. Finally, on a computer with more than one CPU, you can establish affinity between worker processes and specific CPUs, ensuring that applications are tied to processors in accordance with your priorities.

**Managing Isolation and Performance**

To help achieve high levels of application availability, you can assign fewer applications to each application pool, as low as one application per application pool. This increases isolation, protecting applications from each other. However, configuring fewer applications per application pool increases the total number of application pools. This results in greater CPU and memory usage, degrading server performance, and resulting in reduced request-processing throughput.

The performance impact of increasing the number of application pools on a computer varies based on the number of application pools that are running worker processes. If there are a large number application pools, but few worker processes are running, the impact is relatively small; if there are a large number application pools and many worker processes are running, worker processes can impact each other as they use up memory and CPU capacity, thereby negatively impacting performance.

The degree to which increasing application isolation causes performance degradation can only be determined by testing. You can test for the optimum level of isolation by doing the following:

1. Set the highest level of isolation that your configuration can support, and test with all application pools running simultaneously. Assess the results after using Performance Monitor to record requests processed per second, CPU usage, and memory usage. If performance is satisfactory, the test is complete; if performance is not satisfactory, proceed to Step 2.

2. If the results in Step 1 are unsatisfactory, reduce the isolation level by reducing the number of application pools, and then run the test again.

3. Continue Step 2 until you achieve satisfactory performance. In addition to reducing the number of application pools, you might need to remove applications from the server.

For information about isolating applications, see Isolate Applications in Worker Process Isolation Mode. For information about assigning applications to application pools by using a command-line procedure or scripts, see the AppPoolId Metabase Property.

# Configuring Idle Time-out for Worker Processes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To conserve system resources, configure the idle time-out interval of an application pool to shut down its worker process, or worker processes, after a specified period of idle time. This reduces memory consumption and CPU overhead. This also helps you manage the resources on a Web server when the processing load is heavy, when identified applications consistently fall into an idle state, or when new processing space is not available. To preserve performance, as the number of application pools that run simultaneously on a server increases, you should decrease the idle timeout setting.

By default, idle time-out is enabled, and the default time-out period is 20 minutes. To automatically shut down worker processes as they become idle, configure a scheduled shutdown of the idle worker process, or worker processes.

For information about how to configure idle time-out by using IIS Manager, see Configure Application Pool Performance. For information about how to configure idle time-out by using a command-line procedure or scripts, see the IdleTimeout Metabase Property.

# Configuring a Request Queue Limit

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When an application pool receives requests faster than it can handle them, the unprocessed requests might consume all of the memory on the server, slowing the server and preventing other application pools from processing requests. This can happen when the queue size limit for the requests is large and legitimate requests are coming in at a rapid rate, or during a denial of service attack.

To prevent requests from consuming all the memory for an application pool, limit the size of the request queue for the application pool. To monitor and manage CPU consumption, you must enable CPU monitoring, gather data on CPU usage, and reduce the request queue limit when CPU consumption exceeds the specified limit.

For example, if you have a Web server that is capable of processing 10,000 requests per minute, and you have an application that can queue up 5,000 requests within that period of time, the application consumes 50 percent of the capacity of the server. You can configure IIS 6.0 to limit the number of requests to 2,000, reducing the potential memory consumption of the application and reducing the Web server capacity used by the application to 20 percent.

By default, request queue limit is enabled in IIS 6.0, and the default value is 1,000 requests. When resetting the default value, remember that a value that is too high can prevent other applications from having an equal share of the server resources and a value that is too low can cause the server to return 503 errors, or a disruption of service for the clients.

For information about how to configure a request queue limit by using IIS Manager, see Configure Application Pool Performance. For information about how to configure a request queue limit by using a command-line procedure or scripts, see the AppPoolQueueLength Metabase Property.

**Managing ASP.NET and Request Queue Limit**

Recycling worker processes that are serving ASP.NET applications requires that you restart under a request processing load. Starting the new worker processes under a load might cause requests to be rejected because the queue is too full. When this occurs, either increase the queue size limit or disable recycling.

For more information about ASP.NET–specific considerations for applications, see Deploying ASP.NET Applications in IIS 6.0.

# Enabling HTTP Compression IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Enabling IIS's HTTP compression feature provides faster transmission time between compression-enabled browsers and IIS regardless of whether your content is served from local storage or a UNC resource. You can compress static files only, application response files only (dynamic compression), or both static files and application response files.

Enabling compression of static files in IIS always results in more efficient utilization of bandwidth and faster Web site performance. Enabling dynamic compression always results in more efficient utilization of bandwidth as well, but if your server's processor utilization is already extremely high, the CPU load imposed by dynamic compression might make your site perform more slowly. It is therefore important to do performance testing of a server both with and without dynamic compression to determine whether dynamic compression is of benefit.

If you already have a network-based compression solution installed, enabling compression in IIS might not result in bandwidth utilization or performance improvements, depending on whether the network-based compression and IIS-based compression are redundant.

For more information about how to enable HTTP compression server-wide, see Using HTTP Compression.

For a detailed discussion of how IIS's HTTP compression works, how to test compression, and detailed compression configuration options, see Optimizing IIS 6.0 Performance.

# Configuring Web Gardens

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To improve availability for a Web site or application, you can increase the number of worker processes servicing the application pool by implementing a *Web garden*, an application pool that uses more than one worker process on a single computer. This might be helpful if you have already configured a single-application application pool and need to improve availability.

Creating a Web garden for an application pool enhances performance by providing the following benefits:

- Robust processing of requests: When a worker process in an application pool is tied up (for example, when a script engine stops responding), other worker processes can accept and process requests for the application pool.
- Reduced contention for resources: When a Web garden reaches a steady state, each new TCP/IP connection is assigned, according to a round-robin scheme, to a worker process in the Web garden. This helps smooth out workloads and reduce contention for resources that are bound to a worker process.

As a best practice, create Web gardens only for Web applications that meet the following criteria:

- The application runs multi-instantiated, so that a different instance of the application can be assigned to each worker process.
- The Web application is not CPU-intensive. If the CPU is the bottleneck, then adding worker processes cannot help improve performance.
- The application is subject to synchronous high latency. For example, if an application calls a back-end database and the response is slow, then a Web garden supports other concurrent connections without waiting for the slow connection to complete.

Because Web gardens enable the use of multiple processes, all processes have their own copy of application state, in-process session state, caches, and static data. Web gardens are not advantageous for applications that depend on application state. Identify the standard relative performance in a Web garden before you decide whether Web gardens are the best solution for your server.

| **Note** |
| --- |
| A Web garden differs from a *Web farm*, which distributes the load across multiple Web servers. You can design Web farms to provide your applications with the highest availability. For more information about Web farms, see Scalability. |

By default, the Web garden feature is disabled in IIS 6.0, and the default number of worker processes assigned to an application pool is set to **1**. You can enable Web garden functionality by setting the number of worker processes assigned to an application pool to a number greater than **1**.

For information about how to configure a Web garden by using IIS Manager, see Configure Application Pool Performance. For information about how to configure a Web garden by using a command-line procedure or scripts, see the MaxProcesses Metabase Property.

# Setting Processor Affinity on Servers with Multiple CPUs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

On a server with multiple CPUs, you can configure application pools to establish affinity between worker processes and CPUs to take advantage of more frequent CPU cache hits. By using processor affinity, you can control processors when applications place demands on system resources. Processor affinity is used in conjunction with the processor affinity mask setting to specify the CPUs.

Use processor affinity on computers with multiple processors when it is necessary to continue processing an application that is tied to a subset of the hardware in such a way that transitions from one processor to another are minimized. By default, processor affinity is disabled in IIS 6.0, and the load is distributed across all available CPUs. On a server that is supporting multiple applications, you can enable processor affinity so that a processor is dedicated to each application.

For information about how to set processor affinity, see Set Processor Affinity, or see the SMPAffinitized Metabase Property and the SMPProcessorAffinityMask Metabase Property.

# Managing Application Pool Health

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 health-monitoring and management features support application availability by detecting problems with worker processes that are assigned to application pools, and by taking action to manage the problems. When IIS 6.0 is running in worker process isolation mode, you can detect the health condition of a specific application pool by enabling the WWW service to ping a worker process that is assigned to it. Also, you can set features to take corrective action when an unhealthy condition occurs.

Generally, a healthy worker process responds to periodic pings that are sent by the WWW service. If a worker process does not respond to the pings within the time interval specified by the **PingResponseTime** property, the WWW service determines the worker process to be unhealthy and triggers other health management features.

Consider using worker process health detection for the following reasons:

- It allows IIS to determine whether each process has at least one thread available to work on requests, and then decide whether to pull new requests from HTTP.sys or to continue processing existing requests.

- It allows the worker process to request recycles when Internet Server Application Programming Interface (ISAPI) extensions such as ASP.NET and ASP are identified as unhealthy. The health detection feature uses the ability of a worker process to respond to a ping message. The worker process responds to a ping by requesting a recycle if one or more of the ISAPI extensions it has loaded is identified as unhealthy.

# Configuring Worker Process Pinging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Application pool health detection* is the ability of the WWW service to detect that a worker process assigned to the application pool is in an unhealthy state by periodically *pinging*, or sending a ping message to, the worker process.

**Optimizing Ping Intervals**

You can specify a value for the ping interval by examining the potential amount of system resources that can be consumed by unhealthy worker processes over the interval (the default is 30 seconds). Then determine how long you can wait for system resources that unhealthy worker processes are using to be freed so that they are available to other worker processes.

If a worker process becomes unhealthy, the client does not receive a response until another worker process is started to replace the unhealthy worker process. This leaves the unhealthy worker process active without releasing the system resources used by the process. Over a period of time, the amount of system resources used by an unhealthy worker process can accumulate and affect the overall performance of the Web server.

For example, you might have an application that causes worker processes to become unstable. Over a five-minute period of time, the amount of system resources consumed by these unstable worker processes can consume Web server resources and affect the performance of other applications that are on the same Web server. You can configure IIS 6.0 to ping worker processes in that application pool to check the health of the worker processes every 30 seconds, which ensures that the unhealthy worker processes are shut down before system resources are consumed.

Set ping intervals to avoid false failure conditions caused when healthy worker processes do not respond within the specified response time. You can check this by testing to ensure that the WWW service does not shut down a worker process that is still running, but is too busy to return the ping in the specified response time. If your tests indicate that your worker processes are shutting down because of false failure conditions, increase the ping response time, or alter the load on the worker process.

By default, worker process pinging is enabled in IIS 6.0. The default ping interval is 30 seconds, and the default ping response time is 60 seconds.

For information about how to configure worker process pinging by using IIS Manager, see Configure Application Pool Health. For information about how to configure worker process pinging by using a command-line procedure or scripts, see the PingingEnabled Metabase Property, PingInterval Metabase Property, and PingResponseTime Metabase Property.

# Configuring Rapid-Fail Protection for Worker Processes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When the WWW service detects that a worker process fails more than a set maximum number of times in a specified time period, it places the application pool that the worker process is serving in a rapid-fail state. Rapid-fail protection reduces processing overhead for unhealthy applications because the requests do not enter user-mode processing. As a result, applications that are running in other application pools are protected from the unhealthy application.

By default, rapid-fail protection is enabled in IIS 6.0, with a maximum of 5 failures in a 10-minute interval. You can specify the threshold for placing an application pool in rapid-fail state by setting the maximum number of worker process failures that you want to allow over a specific period of time. Specify a threshold that allows IIS 6.0 to detect that a significant number of the worker processes in an application pool are failing and prevent system resources from being consumed.

For information about how to configure rapid-fail protection by using IIS Manager, see Configure Application Pool Health. For information about how to configure rapid-fail protection by using a command-line procedure or scripts, see the RapidFailProtection Metabase Property, RapidFailProtectionMaxCrashes Metabase Property, and RapidFailProtectionInterval Metabase Property.

# Configuring the Startup Time Limit for Worker Processes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The startup time limit specifies how long IIS will wait for a worker process to start servicing requests. By default, the worker process startup time limit is enabled in IIS 6.0 with a default setting of 90 seconds. If 90 seconds is not long enough to allow the application to begin processing requests, try a longer time interval.

For information about how to configure the startup time limit for worker processes by using IIS Manager, see Configure Application Pool Health. For information about how to configure the startup time limit for worker processes by using a command-line procedure or scripts, see the StartupTimeLimit Metabase Property.

# Configuring the Shutdown Time Limit for Worker Processes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When IIS 6.0 detects that a worker process is unhealthy, IIS 6.0 marks the worker process for termination. The shutdown time limit specifies how long IIS waits until forcibly terminating the worker process when the worker process does not shut down automatically.

By default, worker process shutdown time limit is enabled in IIS 6.0 with a default of 90 seconds. If you have an accurate profile of requests, specify a shutdown time limit that is long enough to allow the worker process to complete any pending requests before shutting down.

For information about how to configure the shutdown time limit for worker processes by using IIS Manager, see Configure Application Pool Health. For information about how to configure the shutdown time limit for worker processes by using a command-line procedure or scripts, see the ShutdownTimeLimit Metabase Property.

# Enabling Debugging for Application Pool Failures

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Debugging a specific worker process that is serving an application pool helps you build high-availability applications. By default, debugging is disabled in IIS 6.0. When you enable debugging, the WWW service stops managing worker processes that fail to respond to a ping in the specified time period. Rather than terminating the application when a process fails to respond to a ping, IIS executes an action such as a process dump or an e-mail notification to alert you to the problem. You can then review the application to identify the cause of the failure.

| **Note** |
| --- |
| Enable debugging on an application pool only if the application pool continues to fail to respond to pings. The unhealthy process that is left running will continue to consume system resources until the process is terminated by the administrator or until IIS causes it to self-terminate. If you enable debugging, be sure to terminate the worker process after the debugging action is finished. |

For more information about how to enable debugging, see Debug Application Pool Failures, or see the OrphanWorkerProcess Metabase Property, OrphanActionExe Metabase Property, and OrphanActionParams Metabase Property.

# Configuring Application Pool Identity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Application pool identity* is the user account that serves as the process identity for the worker processes that are servicing the application pool. *Process identity* is the account that a process runs under. Every Windows Server 2003 process has a process identity that is used to control access to resources on the system.

Application pool identity can be assigned to a predefined account. Predefined accounts are known as *service-user accounts* and they are created by the operating system. You can use the predefined NetworkService, LocalSystem, or LocalService accounts for the application pool identity.

In IIS 6.0 worker process isolation mode, application pools have a default identity of NetworkService. The NetworkService identity has minimal administrative credentials, which helps reduce the attack surface of your Web server. If you deploy applications that require a fixed worker process identity, you can change the application pool identity. For information about security and application pool identities, see "Configuring Application Pool Identity Settings" in Deploying ASP.NET Applications in IIS 6.0.

For information about how to configure application pool identity by using IIS Manager, see Configure Application Pool Identity. For information about how to configure application pool identity by using a command-line procedure or scripts, see the AppPoolIdentityType Metabase Property.

# Testing Applications for Compatibility

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Whether you are deploying existing applications that were developed for earlier versions of IIS on a server running IIS 6.0 in worker process isolation mode, or you are deploying new applications, it is important to test the applications for compatibility with worker process isolation mode features.

Functional testing consists of installing and configuring applications on a test server, and then sending HTTP requests to the application from a client on another computer. Load testing is functional testing that consists of sending HTTP requests in rapid succession to stress the worker process and the server. In both cases, you must diagnose the results for possible incompatibilities.

Figure 4.5 shows the process for testing application compatibility with worker process isolation mode.

**Figure 4.5 Testing Applications for Compatibility**




**Checking for Known Conflicts with IIS 6.0**

Most Web applications that run on earlier versions of IIS will run on IIS 6.0 in worker process isolation mode. However, in some cases, applications cannot run in worker process isolation mode. In other cases, you might be required to disable certain features to run an application in worker process isolation mode.

Applications that have the following characteristics are incompatible with worker process isolation mode, or must have certain features disabled to run in worker process isolation mode:

- Require Read Raw Data Filters. Applications that require Read Raw Data Filters are not compatible with worker process isolation mode. The Read Raw Data Filters feature is supported only in IIS 5.0 isolation mode.
- Require Digest authentication. You must use LocalSystem as the application pool identity for standard Digest authentication. This restriction does not apply when you are using Advanced Digest authentication.
- Run only as a single instance. When an application runs only as a single instance, you cannot configure Web gardens, you cannot run the application in more than one application pool, and you cannot use overlapped recycling.
- URL length exceeds 16 kilobytes. The default request header limit for HTTP.SYS is 16KB.
- ASP applications that use consecutive dots in an include file. Using two or more consecutive dots in an include file works only if parent paths are enabled. Alternatively, consider not using include files.
- Dependency on Inetinfo.exe. The application requires Inetinfo.exe, but the Inetinfo.exe process does not run in worker process isolation mode. This occurs rarely, if at all.
- Require Dllhost.exe. The application requires Dllhost.exe, but the Dllhost.exe process is not available in worker process isolation mode. This occurs rarely, if at all.

# Testing Applications for Compatibility with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can begin to test the compatibility of an application during the setup process for each application that you install on IIS 6.0.

To perform a setup test, complete the following steps:

1. Install the application by using the setup software provided.

2. Examine the IIS log (in \\Windows\System32\LogFiles\W3SVC*xxx*) for messages that indicate a WWW service failure. Also, look at the system Event log for 500-level errors.

Some common incompatibilities can be identified when you set up applications to run in IIS 6.0 worker process isolation mode. You can recognize these common incompatibilities by their symptoms, and then take the appropriate action to resolve them.

**Version Detection Incompatibility**
During setup, some applications fail to detect the IIS version correctly because they try to read non-standard registry keys that do not exist in IIS 6.0. The supported version information is under the following registry key:

HKLM\SOFTWARE\Microsoft\InetStp

**Symptoms**
Some applications check for a 5 or a 4 in the version information, and then quit setup if the value is 6.

Non-standard registry locations where applications might look for version information include the following:

HKLM\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters

HKLM\SOFTWARE\Microsoft\InetMGR\Parameters

**Resolution**
Add the registry keys that the application is searching for to the registry. Otherwise, you cannot run the applications until you modify them to search for an existing registry key.

**Wwwroot and Scripts Retrieval Incompatibility**
Some applications use a nonexistent registry key to retrieve wwwroot and scripts. When this is true, the following message appears during setup:

```
HKLM/System/CurrentControlSet/W3svc/Parameters/Virtual Roots
"/ "=REZ_SZ:<physical location of wwwroot>
"/Scripts"=REZ_SZ:<physical location of scripts>
```

**Symptom**
Setup shows no default value for the wwwroot and scripts directories.

**Resolution**
Add the registry keys that the application is searching for to the registry. Otherwise, you cannot run the applications until you modify them to search for an existing registry key.

# Testing Applications for Functional Compatibility with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can test the functional compatibility of an application running in worker process isolation mode by installing the application on a test server, and then sending HTTP requests to the application from another computer. You can identify the following functional incompatibilities when you send correctly composed HTTP requests to applications that were developed for earlier versions of IIS, or to new applications yet to be tested.

**ISAPIs Not Enabled in the Web Service Extensions Node**
Some application setups install script maps for their script engine but do not enable the script engine, and the documentation for the application might not mention the need to set the properties (by using IIS Manager, for example).

**Symptoms**
If your log file (*systemroot*\system32\logfiles\w3svc*Filename*.log) returns a 404 error with suberror code 2, then your application is trying to use a disabled ISAPI or Common Gateway Interface (CGI) file.

**Resolution**
Use the Web Service Extensions node in IIS Manager to enable the ISAPI or CGI file. Be sure to enter the exact path entry that is in the scriptmaps setting in the path part of the entry.

**MIME Types Incompatibility**
Your application might use a file name extension that is not recognized by IIS 6.0.

**Symptoms**
If your WWW service log file, which is located in *systemroot*\system32\logfiles\w3svc*Filename*.log, returns a 404 error with suberror code 3, then a file name extension is not recognized by IIS 6.0.

**Resolution**
Manually add the file name extension to the computer-level properties in IIS Manager. For more information about configuring MIME types, see Configure MIME Types.

**Recycling Incompatibility**
The configurable recycling feature in IIS 6.0 might expose some design flaws in your applications. The most common problem is that applications take too long to start up or to shut down. For example, it is not a problem in IIS 5.0 if an application receives an access violation when shutting down. However, an IIS 6.0 administrator might see error messages in the Event log indicating that the application pool is taking too long to shut down. If the application needs to warm up when it starts up, then recycling causes the application to run at low efficiency every time it starts up. If the application does not save its state data when it shuts down, it loses state data often. Test the impact of the recycling period by configuring it to a shorter time and continuing to check the log.

**Symptoms**
You can identify recycling incompatibility problems by testing the worker process with a short recycle period. If problems exist in the applications in the application pool that is served by the worker process, the following symptoms might occur:

- If an application is shutting down slowly, the Event log might contain error messages indicating that the worker process is exceeding its shutdown time limit.
- If the application needs time to warm up when it starts, response times might be very slow at first.

**Resolution**
The application might require extensive modification. While running the application unmodified, you cannot use recycling. For more information about configuring the features for worker process recycling, see Configure Application Pool Recycling.

**Overlapped Recycling and Single Instance Application Incompatibility**
If your application cannot run as multiple instances, overlapped recycling fails. The failure occurs because the new worker process is running before the existing worker process is terminated, resulting in an attempt to run two instances of the same application.

**Symptoms**
The symptoms are synchronization-related, such as an access denied error that appears in the IIS log (in \\Windows\System32\LogFiles\W3SVC*xxx*).

**Resolution**
Do not use overlapped recycling. For more information about configuring the features for recycling worker processes, see Configure Application Pool Recycling.

**Subauthentication Is Not Installed**
Subauthentication is not installed by default. You require subauthentication if both of the following are true:

- You are using Digest authentication on your server running IIS 6.0.
- The domain controller is running IIS 5.0.

**Symptom**
Digest authentication fails.

**Resolution**
Install subauthentication on your server. For information about how to install subauthentication, see Install Subauthentication.

# Additional Resources About Application Availability in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to ensuring IIS 6.0 availability.

**Related Information**

- Deploying ASP.NET Applications in IIS 6.0 for more information about ASP.NET–specific considerations for applications.
- Securing Web Sites and Applications for information about securing IIS applications.
- Migrating IIS Web Sites to IIS 6.0 for information about migrating IIS applications.
- Migrating Apache Web Sites to IIS 6.0 for information about migrating Apache Web sites.
- Appendix A: IIS Deployment Procedures for more information about procedures for ensuring application availability.
- HSE_REQ_REPORT_UNHEALTHY for more information about the server support extension function.

**Related Topics in the IIS 6.0 Operations Guide**

- The Metabase Property Reference for more information about programmatically configuring IIS for optimum availability.
- Scalability for more information about Web farms.

# Upgrading an IIS Server to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can move Web sites and applications to Internet Information Services (IIS) 6.0 in two ways. *Upgrading* is installing the Microsoft® Windows® Server 2003 operating system and IIS 6.0 on an existing server running the Windows NT® Server 4.0 operating system and IIS 4.0 or the Windows® 2000 Server operating system and IIS 5.0. *Migrating*, on the other hand, is installing Windows Server 2003 and IIS 6.0 on a new server and then moving, or reinstalling, existing Web sites and applications to that server. Both processes involve minimal outage of service to users that access the Web sites and applications, while retaining the majority of the original configuration settings and fully preserving the content of the Web sites and applications.

You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI.

**In This Section**


Overview of Upgrading an IIS Server to IIS 6.0

Preparing to Upgrade

Performing the Upgrade

Converting to Worker Process Isolation Mode

Configuring IIS 6.0 Properties

Configuring IIS 6.0 to Host ASP.NET Applications

Completing the Upgrade_IIS_SP1_Dep

Additional Resources About Upgrading to IIS 6.0

# Overview of Upgrading an IIS Server to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The advantages of upgrading the server that hosts your existing IIS Web sites and applications to IIS 6.0 include the following:

- **Reduces the time required to deploy IIS 6.0.** Upgrading requires minimal user interaction.

- **Reduces possible configuration errors.** Because the majority of the current operating system, IIS, and Web site configuration settings are retained during upgrade, fewer configuration errors result.

The disadvantages of upgrading the server that hosts your existing IIS Web sites and applications to IIS 6.0 include the following:

- **Retains previous versions of software.** The upgrade process upgrades only software components identified by Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Web Edition. Any other installed software components, such as applications or application dynamic-link libraries (DLLs), remain unchanged.

  During the IT lifecycle of the server, you can install and subsequently remove many of these additional software components. The software removal process often leaves DLLs, other files, and folders on the system. Over time, these unused components can accumulate, consume available disk space, and potentially cause instability with existing applications.

- **Retains the previous registry configuration.** The upgrade process makes the appropriate modifications to the registry entries identified by Windows Server 2003. Registry entries are created by an application as it is installed, but are not removed when the application is removed. Cumulatively, these registry entries consume disk storage and can make troubleshooting registry-related problems more difficult because the unused registry entries can have similar naming conventions to active registry entries.

| **Note** |
| --- |
| You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

After upgrading a server to Windows Server 2003 and IIS 6.0, the server runs in *IIS 5.0 isolation mode*, which allows IIS 6.0 to emulate the memory and request processing model in IIS 5.0. However, to take advantage of the security and performance benefits of IIS 6.0, you need to configure the server to run in *worker process isolation mode*, which is the new IIS 6.0 memory and request processing model.

When upgrading existing IIS Web sites is not the best solution for your existing server, consider migrating your IIS Web sites to a newly installed server running IIS 6.0. For information about migrating your existing IIS Web sites to IIS 6.0, see Migrating IIS Web Sites to IIS 6.0.

# Process for Upgrading an IIS Server to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The process for upgrading an existing IIS server hosting IIS Web sites and applications includes preparing for the upgrade, as well as performing the upgrade. Upon completing the upgrade process outlined in this section, you can then configure IIS 6.0 to run your applications in worker process isolation mode. Lastly, if the newly upgraded IIS 6.0 server hosts existing ASP.NET applications, you must configure IIS to use the proper version of the Microsoft .NET Framework for these applications.

| **Note** |
| --- |
| You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

Figure 5.1 illustrates the process for upgrading an existing IIS server hosting Web sites and applications to IIS 6.0.

**Figure 5.1 Upgrading an Existing IIS Server to IIS 6.0**




In addition to upgrading the existing operating system and IIS, the upgrade process automatically upgrades any components that are a part Windows Server 2003, such as FrontPage® Server Extensions from Microsoft. As a part of the upgrade process, you need to evaluate the compatibility of the software that is installed on your existing server, including device drivers, other Windows components, and software, before performing the upgrade of the server.

| **Note** |
| --- |
| To upgrade a Web farm, use the process described in this section to upgrade each server in the Web farm. |

Depending on your familiarity with IIS and its upgrade process, you might require less information to complete the upgrade process. To facilitate the fastest possible upgrade, the following quick-start guide is provided. You can use this guide to help identify the steps of the upgrade process that you need additional information to complete, and then you can skip the information with which you are already familiar. In addition, all of the procedures that are required to complete the upgrade process are documented in the IIS 6.0 Operations Guide.

**Prepare to Upgrade**

1. Determine compatibility with Windows Server 2003.
2. Identify and compensate for changes to IIS 6.0:
    - Ensure that the World Wide Web Publishing Service (WWW service) is enabled after upgrade.
    - Compensate for changes to IIS components.
3. Determine application compatibility with worker process isolation mode by completing the following steps:
    - Evaluate the benefits of worker process isolation mode.
    - Evaluate application changes required for worker process isolation mode.
    - Evaluate management and provisioning script changes required for worker process isolation mode.
    - Verify application compatibility with worker process isolation mode in a lab.
4. Determine application compatibility with the .NET Framework.

**Perform the Upgrade**

1. Back up the Web server before upgrade.
2. Verify that clients are not accessing Web sites.
3. Prevent the WWW service from being disabled.
4. Upgrade the Web server to IIS 6.0.
5. Verify that the operating system upgrade was successful.
6. Back up the IIS 6.0 metabase after upgrade.

**Convert to Worker Process Isolation Mode**

1. Document the current application isolation settings.
2. Configure IIS to run in worker process isolation mode.
3. Configure application isolation settings in IIS 6.0.

**Configure IIS 6.0 Properties**

1. Enable the WWW service.
2. Configure Web service extensions.
3. Configure Multipurpose Internet Mail Extensions (MIME) types.
4. Modify references to IIS 6.0 metabase properties that have changed or are no longer supported.

5. Upgrade FrontPage extended Web sites.
6. Determine whether to run the IIS Lockdown Tool and UrlScan.
7. Make the following security-related configuration changes:
    - Enable essential IIS components and services.
    - Remove any unnecessary IIS virtual directories.
    - Configure the anonymous user identity.

**Configure IIS to Host ASP.NET Applications**

1. Configure IIS to use the correct version of the .NET Framework.
2. Configure the .NET Framework.
3. Review how ASP.NET applications run in each application isolation mode.
4. Migrate Machine.config attributes to IIS 6.0 metabase property settings by completing the following steps:
    - Migrate recycling-related attributes.
    - Migrate performance-related attributes.
    - Migrate health-related attributes.
    - Migrate identity-related attributes.

**Complete the Upgrade**

1. Verify that the Web sites and applications run properly.
2. Back up the Web server.
3. Enable client access after upgrade is complete.

# Preparing to Upgrade

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before upgrading your existing server running IIS, ensure that Windows Server 2003 supports all of the hardware devices that are on the server. You must also ensure that your Web sites and applications are compatible with IIS 6.0, Windows Server 2003, and worker process isolation mode. Lastly, you must determine whether any of your existing ASP.NET applications are compatible with version 1.1 of the .NET Framework that is installed with Windows Server 2003.

Figure 5.2 illustrates the process for preparing to upgrade an IIS server to IIS 6.0.

**Figure 5.2 Preparing to Upgrade**



# Determining Compatibility with Windows Server 2003

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

At a minimum, your existing system hardware and software must be compatible with Windows Server 2003 before you upgrade your existing server to IIS 6.0. The upgrade process identifies any software or hardware devices that are incompatible with Windows Server 2003 and allows you to stop the upgrade process after notifying you of any incompatible software or device drivers.

**Determining the Compatibility of Existing Windows Servers with Upgrade**

Your current Windows server operating system can limit which version of Windows Server 2003 that you can upgrade to. The following table lists the existing Windows server operating systems and the versions of Windows Server 2003 that are supported for upgrade.

**Windows Server Upgrades Supported by Windows Server 2003**

| Existing Windows Server | Standard | Enterprise | Web |
|---|---|---|---|
| Windows NT Server 4.0 | y | y | |
| Windows NT Server 4.0, Enterprise Edition | | y | |
| Windows 2000 Server | y | y | |
| Windows 2000 Advanced Server | | y | |
| Windows Server 2003, Standard Edition | | y | |
| Windows Server 2003, Web Edition | | | y |

If you want to move from an existing Windows server operating system to a version of Windows Server 2003 that is not supported by upgrade, migrate the existing Web server to a new Web server. For more information about migrating Web sites to a new Web server running IIS 6.0, see Migrating IIS Web Sites to IIS 6.0.

**Determining the Compatibility of Existing Hardware**

In most cases, your existing hardware will be compatible with Windows Server 2003. The most common hardware incompatibility is a device driver that is no longer supported. When devices are no longer supported, remove the existing devices and install an equivalent device that is supported by Windows Server 2003. It is also important that you have the latest BIOS version that is available from your computer manufacturer.

For example, the Web server might have an ISA network adapter that is no longer supported in Windows Server 2003. In this situation, you can obtain and install a new network adapter that is compatible with Windows Server 2003.

For more information about the hardware devices that are supported on Windows Server 2003, see the Hardware Compatibility List.

**Determining the Compatibility of Existing Software**

Before upgrading, you need to consider the compatibility of your existing applications, or other software that is installed on your server, with Windows Server 2003. This includes software and tools from manufacturers other than Microsoft, as well as Microsoft server products that do not ship with the Windows operating system. Make sure that you have the latest versions of all pre-existing software or service packs that are compatible with Windows Server 2003.

The most common software incompatibilities are caused when your application depends on software from another manufacturer that does not support Windows Server 2003. Or, you might have applications that were designed to run on the Microsoft Windows NT® Server 4.0 or Windows 2000 Server operating systems, which reference application programming interfaces (APIs) that have been changed or removed in Windows Server 2003.

To help you determine the compatibility of your existing software with Windows Server 2003, use the Windows Application Compatibility Toolkit before upgrading your server.

# Identifying and Compensating for Changes to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Although most of the changes to IIS that are made during the upgrade process have little affect, some of the changes can affect the Web sites and applications, as well as the administration of the Web server. Before upgrading your existing IIS server, review the changes that will be made to IIS during upgrade and determine whether these changes can affect your Web sites and applications. Otherwise, after the upgrade, applications might not function as they were originally designed to. In addition, you need to identify the changes made to IIS during upgrade so that you can properly administer and configure the Web server upon completion of the upgrade.

Identify the changes incurred during the upgrade by completing the following steps:

1. Select a method to ensure that the WWW service is enabled after upgrading a Web server running Windows 2000 Server.

2. Identify changes made to IIS components during upgrade.

# Ensuring That the WWW Service is Enabled After Upgrade

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For security reasons, the World Wide Web Publishing Service (WWW service) is disabled when you upgrade from Windows 2000 Server or Windows 2000 Advanced Server to Windows Server 2003 or Windows Server 2003, Enterprise Edition, respectively. The WWW service is disabled on upgrade because IIS 5.0 was installed by default during the Windows 2000 Server installation, and, in many cases, the WWW service was not used. To decrease the attack surface of the server, IIS 6.0 is disabled by default when you upgrade from Windows 2000 Server to Windows Server 2003 unless you do one of the following:

- For manual or unattended upgrades, run the IIS Lockdown Tool on your Windows 2000 Server before starting the upgrade process.
- For manual or unattended upgrades, add the registry entry RetainW3SVCStatus to the registry.
- For unattended upgrades, add the entry "DisableWebServiceOnUpgrade = False" in the unattended install script.
- Continue with the upgrade and enable the WWW service after upgrade. For more information about enabling the WWW service, see Preventing the WWW Service from Being Disabled later in this section.

| **Note** |
| --- |
| Because IIS 4.0 is not installed by default on computers running Windows NT Server 4.0, the Windows Server 2003 upgrade does not disable the WWW service when upgrading from IIS 4.0 to IIS 6.0. |

# Compensating for Changes to IIS Components

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because of increased security measures, which remove any unnecessary Web sites or configuration settings from the Web server, the upgrade process can remove or change the existing IIS configuration. The upgrade process ensures that you take the necessary steps to accommodate these changes so that your applications run as they did before the upgrade.

The following table lists the changes in the behavior, administration, and support of IIS 6.0

**Changes to IIS Components After Upgrade**

| Before Upgrade | After Upgrade |
|---|---|
| The IISAdmin virtual directory exists. | The IISAdmin virtual directory is removed and a default Web page is placed in the directory. [1] |
| Web service extensions are mapped to 404.dll. [2] | All extensions mapped to 404.dll are disabled. All other extensions are enabled. |
| Web Distributed Authoring and Versioning (WebDAV) is disabled by using NTFS file system permissions. 3 | WebDAV permissions are changed to allow proper operation, and WebDAV is added to the Web service extensions list as prohibited. |
| Indexing Service is disabled by using NTFS permissions. 3 | Indexing Service remains disabled with the same NTFS permissions. |
| Web services hosted on IIS 5.0 are enabled. | Web services supported by the WWW service are disabled. [4] |
| FrontPage 2000 Server Extensions included lightweight server support. | FrontPage 2002 Server Extensions does not include lightweight server support. |
| FrontPage 2000 Server Extensions support. | FrontPage 2000 Server Extensions is not supported. Upgrade all Web sites to FrontPage 2002 Server Extensions. |

[1] Directory is removed if the IIS Lockdown Tool is run before upgrade.

[2] The IIS Lockdown Tool maps all Web service extensions, used by IIS, to 404.dll to disable the extensions. Additional Web service extensions installed beyond the default IIS Web service extensions are not mapped to 404.dll.

[3] The IIS Lockdown Tool disables WebDAV and the Indexing Service by changing the permissions.

[4] If the IIS Lockdown Tool runs before upgrade or the appropriate registry modifications are made, the Web services remain enabled after upgrade. For more information about preventing the Web services from being disabled, see Preventing the WWW Service from Being Disabled.

As a part of the architectural changes to IIS 6.0, many of the features that were previously implemented as Internet Server API (ISAPI) filters are now integrated into IIS directly. The following table lists the ISAPI filters that have been removed and have the functionality of the ISAPI filter implemented in IIS 6.0

**ISAPI Filters Removed During Upgrade**

| ISAPI Filter Removed | ISAPI Filter File Name | Implemention in IIS 6.0 |
|---|---|---|
| SSL encryption | Sspifilt.dll | Integrated into IIS 6.0 directly |
| Compression | Compfilt.dll | Integrated into IIS 6.0 directly |
| Kerberos authentication | Md5filt.dll | Integrated into IIS 6.0 directly |

# Determining Application Compatibility with Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 can run in one of two distinct modes of operation, which are called application isolation modes. *Application isolation* is the separation of applications by process boundaries that prevent the applications from affecting one another, and it is configured differently for each of the two IIS application isolation modes: IIS 5.0 isolation mode and worker process isolation mode.

*Worker process isolation mode* uses the redesigned architecture for IIS 6.0. This application isolation mode runs all application code in an isolated environment. However, unlike earlier versions of IIS, IIS 6.0 provides isolation without a performance penalty because fewer processor instructions are run when switching from one application pool to another. Worker process isolation mode is compatible with most existing Web sites and applications. Whenever possible, run IIS 6.0 in worker process isolation mode benefit from the enhanced performance and security in IIS 6.0.

*IIS 5.0 isolation mode* provides compatibility for applications that depend upon the process behavior and memory model of IIS 5.0. Run IIS in this mode only when a Web site or application cannot run in worker process isolation mode, and run it only until the compatibility issues are resolved.

| **Note** |
| --- |
| IIS 6.0 cannot run both application isolation modes simultaneously on the same server. Therefore, on a single server running IIS 6.0, you cannot run some Web applications in worker process isolation mode and others in IIS 5.0 isolation mode. If you have applications that require separate modes, you must run them on separate servers. |

After you perform an upgrade of the operating system, IIS is configured to run in IIS 5.0 isolation mode by default. Before configuring IIS 6.0 to run in worker process isolation mode, you should evaluate whether your Web sites and applications are compatible with worker process isolation mode. Most of the compatibility issues with IIS 6.0 occur when configuring IIS 6.0 to run in worker process isolation mode.

One of the most common reasons for incompatibility with worker process isolation mode is that applications do not recognize custom ISAPI extensions or DLLs that depend on the memory and request processing models used by earlier versions of IIS. Determine application compatibility in your lab before upgrading your existing IIS server, and if you determine that your applications are not compatible with worker process isolation mode, you can run the applications in IIS 5.0 isolation mode.

| **Note** |
| --- |
| Identifying a complete list of potential incompatibilities that applications can experience with worker process isolation mode is beyond the scope of this topic. Even after following the guidelines in this section, you need to verify in a lab whether your Web sites and applications are compatible with worker process isolation mode. |

Determine the compatibility of an application with worker process isolation mode by completing the following steps:

1. Evaluate the benefits of moving to worker process isolation mode.

2. Evaluate the application changes that are required so that the applications can run in worker process isolation mode.

3. Evaluate the management and provisioning script changes that are required to set up programs and provisioning scripts in worker process isolation mode.

4. Verify the compatibility of the application with worker process isolation mode in a lab.

# Evaluating the Benefits of Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Worker process isolation mode provides higher levels of security and availability for Web sites and applications than IIS 5.0 isolation mode. Therefore, it is recommended that you configure IIS 6.0 to run in worker process isolation mode.

Worker process isolation mode provides the following improvements to IIS 6.0.

**Security Enhancements**
IIS 6.0 includes a variety of security features and technologies that help ensure the integrity of your Web site content, and of the data that is transmitted through your sites. The following security enhancement is only available when IIS 6.0 is running in worker process isolation mode.

**Default process identity for Web sites and applications is set to NetworkService**
In IIS 5.0 isolation mode, the default process identity is LocalSystem, which enables access to, and the ability to alter, nearly all of the resources on the Web server. The potential of attacks is reduced in worker process isolation mode because Web sites and applications run under the NetworkService identity. The NetworkService identity is granted less privileges, which helps prevent an attack from compromising the Web server, which is possible with the LocalSystem identity.

**Performance and Scaling Enhancements**
Future growth in the utilization of your Web sites and applications requires increased performance and scalability of Web servers. By increasing the speed at which HTTP requests can be processed and by allowing more applications and sites to run on one Web server, the number of Web servers that you need to host a site is reduced. The following are a few of the performance improvements included in worker process isolation mode.

**Support for processor affinity for worker processes in an application pool**
You can configure all of the worker processes in an application pool to have affinity with specific processors in a multiprocessor or server. Processor affinity allows the worker processes to take advantage of more frequent processor caching (Level 1 or Level 2).

**Elimination of inactive worker processes and reclamation of unused resources**
You can configure application pools to have worker processes request a shutdown if they are idle for a certain amount of time. This can free unused resources for other active worker processes. New worker processes are then started only when they are needed.

**Distributing client connections across multiple worker processes**
You can configure an application pool to have more than one worker process servicing client connections, also known as a *Web garden*. Because there are multiple worker processes, the incoming client connections are distributed across the worker processes and throughput is not constrained by a single worker process.

**Ability to Isolate Web sites and applications from each other**
You can isolate applications without incurring a performance penalty.

**Availability Enhancements**
Because worker process boundaries isolate the applications in an application pool from the applications in other application pools, if an application fails, it does not affect the availability of other applications running on the server. Deploying applications in application pools is a primary advantage of running IIS 6.0 in worker process isolation mode.

**Reduced number of Web server restarts required when administering Web sites and applications**
Many of the common operation tasks do not force the restart of the server or the Web service restart. These tasks, such as upgrading site content or components, debugging Web applications, or dealing with faulty Web applications, can be performed without affecting service to other sites or applications on the server.

**A fault-tolerant request processing model for Web sites and applications**
In IIS 5.0 isolation mode, each Web site or application has only one worker process. However, in worker process isolation mode, you can create a *Web garden* by configuring a number of worker processes to share the processing. The benefit of a Web garden is that if one worker process stops responding, other worker processes are available to accept and process requests.

**Isolation of failed worker processes from healthy worker processes**
In worker process isolation mode, IIS can determine that a worker process is failing and start a new worker process to replace the failing worker process. Because a new worker process is created before the old worker process terminates, users

requesting the Web site or application experience no interruption of service. After IIS creates the new worker process, the failed worker process can be separated, or *orphaned*, from the application pool. The advantage of orphaning a worker process rather than terminating it is that debugging can be performed on the orphaned worker process.

**Health monitoring of Web sites and applications**
In worker process isolation mode, you can configure an application pool to monitor not only the health of the entire application pool, but also individual worker processes servicing the application pool. Monitoring the health of a worker process allows IIS to detect that a worker process is unable to serve requests and to take corrective action, such as recycling the failed worker process.

In addition, worker process isolation supports other responses when a failed worker process or application pool is detected. For example, IIS can attach a debugger to an orphaned worker process or notify an administrator that an application pool has failed due to rapid-fail protection.

**Prevention of Web sites or applications that fail quickly from consuming system resources**
In some cases, availability can be affected by Web sites and applications that fail very quickly, are automatically restarted, and then fail quickly again. The endless cycle of failure and restarting can consume system resources, causing other Web sites and applications to experience denial of services because of system resource shortages.

Worker process isolation mode includes *rapid fail protection* that stops an application pool when too many of the worker processes assigned to an application pool are found to be unhealthy within a specified period of time.

**Automatic restart of poorly performing Web sites and applications**
Some Web sites and applications have memory leaks, are poorly coded, or have other unidentified problems. In IIS 5.0 isolation mode, these applications can force you to restart the entire Web server. The recycling feature in worker process isolation mode can periodically restart the worker processes in an application pool to manage faulty applications. Worker processes can be scheduled to restart based on several options, such as elapsed time or the number of requests served.

# Evaluating Application Changes Required for Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In most cases, the existing Web sites and applications can be hosted in worker process isolation mode without modification. However, the following are known application issues that can create incompatibilities with worker process isolation mode and require you to run IIS 6.0 in IIS 5.0 isolation mode.

- Requires Inetinfo.exe.When the application must run in the same process with Inetinfo.exe, IIS must be configured to run in IIS 5.0 isolation mode because Inetinfo.exe does not process Web requests in worker process isolation mode. In worker process isolation mode, W3wp.exe processes Web requests.

- Requires Dllhost.exe. When the application depends on Dllhost.exe to process Web requests for the application, IIS must be configured to run in IIS 5.0 isolation mode because Dllhost.exe is not available in worker process isolation mode.

- Requires read raw data filters. If the application uses an ISAPI raw data filter, IIS must be configured to run in IIS 5.0 isolation mode.

- Requires version 1.0 of the .NET Framework. When the application requires version 1.0 of the .NET Framework, IIS must be configured to run in IIS 5.0 isolation mode because version 1.0 of the .NET Framework is only supported in IIS 5.0 isolation mode.

- Requires ISAPI filters or extensions that are incompatible with worker process isolation mode. When the application requires ISAPI filters or extensions that are incompatible with worker process isolation mode, IIS must be configured to run in IIS 5.0 isolation mode. ISAPI filters or extensions might be incompatible if the filter or extension has one of the following characteristics:

    - Runs in multiple instances and expects to be recycled by using the recycling provided in IIS 5.0.

    - Expects to have exclusive lock on a resource, such as a log file.

If any of the Web sites and applications running on the existing Web server has one or more of these characteristics, then do one of the following:

- Configure IIS to run in IIS 5.0 isolation mode to ensure Web site and application compatibility. If the Web sites and applications need to be hosted in IIS 5.0 isolation mode, then the upgrade process is complete.

- Modify the existing applications to remove the dependencies.

# Evaluating Management and Provisioning Script Changes Required for Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When management or provisioning scripts exist for your Web sites and applications, you might need to modify them so that they properly set up the Web sites and applications for the application isolation mode that is running on the Web server — IIS 5.0 isolation mode or worker process isolation mode. If you do not modify your management and provisioning scripts as required, you will be unable to use them to install and configure your Web sites and applications on IIS 6.0.

For example, you might have a provisioning script that uses Microsoft Active Directory® Service Interfaces (ADSI) to create a Web site and configure the site for High isolation in IIS 5.0 isolation mode, which does not exist in worker process isolation mode. After upgrade when the Web server is running in worker process isolation mode, you need to modify the script to create application pools instead.

Common modifications that might be necessary to your setup programs and provisioning scripts include:

- When the script installs an ISAPI extension or filter, you might need to modify the script to add an entry for the ISAPI extension or filter to the Web service extensions restriction list and set the status of the entry to **Allowed**.

  For more information about modifying your setup programs or provisioning scripts to install and enable an ISAPI extension or filter, see Configuring Web Service Extensions.

- When the script installs an application that contains dynamic content, you might need to modify the script to set the status of the appropriate Web service extensions to **Allowed**, so that IIS allows the dynamic content to run.

  For more information about modifying your setup programs or provisioning scripts to enable dynamic content, see Configuring Web Service Extensions.

- When the script installs a Web site or application that runs in High isolation in IIS 5.0, you might need to modify the script to create an application pool and configures the application pool with settings that are comparable to the original IIS 5.0 isolation settings.

  For more information about modifying your setup programs or provisioning scripts to create and configure application pools, see Configuring Application Isolation Settings in Worker Process Isolation Mode later in this section.

- When the script references metabase properties, you might need to modify the script if it references metabase properties that have changed or are no longer supported in IIS 6.0.

  For more information about modifying your setup programs or provisioning scripts to reference the proper metabase properties in IIS 6.0, see Modifying References to IIS 6.0 Metabase Properties later in this section.

# Verifying Application Compatibility with Worker Process Isolation Mode in a Lab

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After modifying your Web sites, applications, setup programs, and provisioning scripts to be compatible with worker process isolation mode, you need to test your modifications in a lab. Be sure to test for compatibility before performing the upgrade process on a production Web server.

Verify the compatibility of your Web sites, applications, setup programs, and provisioning scripts with worker process isolation mode in a lab by completing the following steps:

1. Make an image backup of the Web server you are going to upgrade.

2. Restore the backup to a Web server in your lab, referred to hereafter as a *test Web server*. Ensure that the test Web server is not connected to your production network to prevent any problems encountered during the upgrade from affecting your production network.

3. Perform an upgrade on the test Web server.

4. Configure the Web server to run in worker processor isolation mode.

5. Make the necessary modifications to the Web sites, applications, setup programs, and provisioning scripts so that they are compatible with worker process isolation mode.

6. Verify the Web sites, applications, setup programs, and provisioning scripts run correctly on the test Web server.

For more information about setting up a test lab, see Designing a Test Environment in Planning, Testing, and Piloting Deployment Projects of the Microsoft® Windows® Server 2003 Deployment Kit.

# Determining Application Compatibility with the .NET Framework

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For a successful upgrade to Windows Server 2003 and IIS 6.0, you need to determine whether your ASP.NET applications are dependent on specific versions of the .NET framework. Windows Server 2003 includes version 1.1 of the .NET Framework, although most of the .NET applications that were developed before the release of Windows Server 2003 were designed to run on version 1.0 of the .NET Framework.

| **Note** |
| --- |
| Version 1.0 of the .NET Framework is only supported in IIS 5.0 isolation mode. Therefore, you can only run version 1.0 and version 1.1 of the .NET Framework on the same server when IIS 6.0 is configured to run in IIS 5.0 isolation mode. Version 1.1 of the .NET Framework is supported in IIS 5.0 isolation mode or worker process isolation mode. |

Typically, ASP.NET applications running on version 1.0 of the .NET Framework are compatible with version 1.1 of the .NET Framework. However, there might be some incompatibilities, the majority of which are security related and can be corrected by configuring the .NET Framework to be less restrictive.

For example, in the .NET Framework version 1.0, the .NET Framework only examines the **SQLPermission.AllowBlankPassword** attribute if the user includes the password keyword in their connection string. If an administrator or user sets "**SQLPermision.AllowBlankPassword**=False", it is possible to specify a connection string like "server=(localhost);uid=sam" and succeed. In the .NET Framework version 1.1, this connection string fails.

You can use both version 1.0 and version 1.1 of the .NET Framework on the same server running IIS 6.0, also known as *side-by-side configuration*. Side-by-side configuration allows you to run a mixture of applications that require version 1.0 or version 1.1 of the .NET Framework. During your lab testing, determine the version of the .NET Framework that is required by your applications.

For a current list of possible compatibility issues when upgrading from version 1.0 to version 1.1 of the .NET Framework, see Compatibility Considerations and Version Changes on the GotDotNet Web site. For more information about using multiple versions of the .NET Framework in side-by-side configuration, see Configuring IIS 6.0 to Use the Correct Version of the .NET Framework.

| **Note** |
| --- |
| Before upgrading the script maps for your ASP.NET applications, point to version 1.0 of the .NET Framework. Upon completion of the upgrade, version 1.1 of the .NET Framework is installed and the script maps are modified to use version 1.1. When your applications require version 1.0, you need to modify the script maps to use version 1.0, rather than version 1.1. For more information about how to modify the script maps to use version 1.0 of the .NET Framework, see Configuring IIS 6.0 to Use the Correct Version of the .NET Framework. |

# Performing the Upgrade

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before upgrading your existing server to Windows Server 2003 and IIS 6.0, you must back up the server, verify that clients are not accessing Web sites on the server, and optionally prevent the WWW service from being disabled. Then upgrade the Web server to Windows Server 2003 and IIS 6.0. Finally, verify that the upgrade to Windows Server 2003 completed successfully.

| **Note** |
| --- |
| You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

Figure 5.3 illustrates the process for performing to upgrade an IIS server to IIS 6.0.

**Figure 5.3 Performing the Upgrade**



# Backing Up the Server_IIS_SP1_Dep

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you change any of the configuration settings on the existing Web server, perform a complete image backup. The purpose of this image backup is to provide a point-in-time snapshot of the Web server. If unforeseen problems occur during the upgrade, you can use this backup to restore the Web server to a known configuration.

| **Note** |
| --- |
| Do not continue with the upgrade process unless you have a successful backup of the entire Web server or you have another Web server that has the same Web sites and applications. Otherwise, you can lose Web sites, applications, or data stored on the Web server. |

For more information about how to back up the Web server, see Back Up and Restore the Web Server to a File or Tape.

# Verifying That Clients Are Not Accessing Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you upgrade the existing server, ensure that no active client sessions are running. Upgrading the server without doing this can result in abnormally terminated client processes and a loss of information.

Verify that clients are no longer accessing Web or File Transfer Protocol (FTP) sites by completing the following steps:

1. Prevent new clients from accessing the sites by pausing the sites.

   For more information about how to pause Web or FTP sites, see Pause Web or FTP Sites.

2. Enable monitoring of active Web and FTP connections.

   For more information about how to monitor the active Web and FTP connections, see Monitor Active Web and FTP Connections.

3. When the number of active Web and FTP counters is zero, disable the network adapter that clients use to access the Web server.

   For more information about how to disable the network adapter used by clients, see Disable Network Adapters.

# Preventing the WWW Service from Being Disabled

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Earlier in the deployment process, you selected the method for enabling the WWW service after upgrading from a server running Windows 2000 Server and IIS 5.0. If you decided to prevent the WWW service from being disabled after completion of the upgrade, you must do one of the following:

- Modify the registry or unattended setup script.
- Run the IIS Lockdown Tool.

| **Note** |
| --- |
| If you are upgrading a Web server running Windows NT Server 4.0 and IIS 4.0, you do not need to run the IIS Lockdown Tool before upgrade. |

# Modifying the Registry or Unattended Setup Script

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Prevent the WWW service from being disabled during upgrade by doing one of the following:

- Create the registry entry do_not_disable in the subkey HKLM\SYSTEM\CurrentControlSet\Services\W3SVC\RetainW3SVCStatus as data type REG_DWORD with a value of 0x1. For more information about how to configure the registry, see Configure the Registry.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference. |

- Include the entry "**DisableWebServiceOnUpgrade** = False" in the [InternetServer] section in an unattended installation script.

# Running the IIS Lockdown Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS Lockdown Tool is designed to help secure earlier versions of IIS by doing the following:

- Preventing the WWW service from being disabled after upgrade on Web servers that are currently running Windows 2000 Server and IIS 5.0. Disabling the WWW service prevents any Web sites or applications from functioning.

- Helping to secure the existing Web server by disabling or removing unnecessary features that are present in IIS 4.0 and IIS 5.0 installations. These features would otherwise remain on the Web server after upgrading, leaving it vulnerable to attacks.

The IIS Lockdown Tool works by turning off unnecessary features, thereby reducing the attack surface that is available to malicious users. To provide multiple layers of protection for an in-depth defense against potential attackers, the IIS Lockdown Tool includes UrlScan and customized security templates based on supported server roles.

| **Note** |
| --- |
| The WWW service is enabled after the upgrade process is complete on servers running IIS 4.0 because IIS is installed by default in Windows 2000 Server. |

**Server Configurations with the IIS Lockdown Tool**

The IIS Lockdown Tool secures the existing IIS server by performing one or more of the following user-specified transactions:

- Enabling or disabling IIS services such as the WWW service, the FTP service, or the Simple Mail Transfer Protocol (SMTP) service.

- Removing services that are disabled.

- Disabling active Active Server Pages (ASP) applications on the server.

- Disabling optional components, including:

    - Index Server Web interface

    - Server-side includes (SSI)

    - Internet Data Connector (IDC)

    - Internet printing

    - HTR scripting

    - Web Distributed Authoring and Versioning (WebDAV)

- Disabling anonymous user access to the server by denying:

    - Execute permissions on the operating system executables and DLLs

    - Write permissions on all Web site content directories

- Removing unnecessary virtual directories, including:

    - IIS Samples
    - Scripts
    - Microsoft Data Access Components (MDAC)
    - IIS Admin
- Installing UrlScan

| **Tip** |
| --- |
| The IIS Lockdown Tool helps secure IIS. However, to maintain IIS security, install all available security patches and hotfixes to help protect against known security vulnerabilities. |

**Server Roles in the IIS Lockdown Tool**

Depending upon the applications that you are hosting and the software that you are using on the existing IIS server, select the server role that most closely corresponds to the server you are upgrading. The IIS Lockdown Tool uses the specified server role to determine the appropriate actions to configure the existing IIS server.

Regardless of the server role selected, UrlScan is not required for the purposes of upgrade. UrlScan can be installed by using the IIS Lockdown Tool or separately. For more information about determining whether you need to run UrlScan after upgrade, see Determining Whether to Run the IIS Lockdown Tool and UrlScan.

The server roles that are included in the IIS Lockdown Tool include the following:

- Small Business Server for Windows NT Server 4.0
- Small Business Server 2000
- Exchange Server 5.5
- Exchange Server 2000
- Microsoft SharePoint™ Portal Server
- FrontPage Server Extensions
- SharePoint Team Services
- BizTalk® Server 2000
- Commerce Server 2000
- Proxy Server
- Static Web server
- Dynamic Web server (ASP-enabled)
- Server that does not require IIS
- Other (a server that does not match any of the roles in this list)

Each of the server roles in the IIS Lockdown Tool secures the Web server by performing a different combination of the security configuration changes listed earlier. For example, if you select the **Exchange Server 2000 (OWA, PF Management, IM,**

**SMTP, NNTP)** server role, then FTP is disabled, and SMTP and NNTP are enabled. However, if you select the SharePoint Portal Server, server role then FTP, SMTP, and Network News Transfer Protocol (NNTP) are disabled. To determine the configuration performed by each server role, review the contents of the IisLockd.ini file in the same folder that contains the IIS Lockdown Tool.

After running the IIS Lockdown Tool and before upgrading, thoroughly test the server to ensure that your Web sites and applications function as they did before. If you find that the configuration settings adversely affect your applications, run the IIS Lockdown Tool again to undo the changes that were made. If you are uncertain whether the IIS Lockdown Tool has been run on the server, you can run the tool again without adversely affecting the system.

| **Tip** |
| --- |
| When the IIS Lockdown Tool locks down a server, it creates a log file named Oblt-log.log and saves it in the folder that contains IISLockd.exe. This file contains information about every action the IIS Lockdown Tool implemented on the system. |

Administrators can run the IIS Lockdown Tool unattended, allowing consistent configuration across many servers through unattended scripts. For more information about running the IIS Lockdown Tool unattended, see RunLockdUnattended.doc, which is located in the folder that contains the files for the IIS Lockdown Tool.

To download the latest version, see IIS Lockdown Tool.

# Upgrading the Server to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The upgrade process from IIS 4.0 or IIS 5.0 to IIS 6.0 completes with minimal interaction because the majority of the IIS 4.0 or IIS 5.0 settings are retained. After the upgrade is complete, the Web sites and applications typically function as they did before the upgrade.

The upgrade process runs a number of compatibility tests before actually performing the upgrade. The compatibility tests relate directly to IIS 6.0 and determine the following:

- Whether system volume for the server is formatted with the NTFS file system
- Whether the IIS Lockdown Tool has been run on the server
- Whether the existing server is currently a node in a Microsoft server cluster

A dialog in the Windows Server 2003 upgrade notifies you if any potential compatibility issues exist. After reviewing the potential compatibility issues, you can abort the upgrade process and resolve the compatibility issues, or you can continue the process. None of the compatibility issues prevent the upgrade process from completing.

# Verifying That the Operating System Upgrade Was Successful

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In most cases, the upgrade process completes without any difficulties. However, before continuing with the upgrade process, verify that the operating system was upgraded successful by completing the following steps:

1. Open *%systemroot%*\Setuperr.log in Notepad and search for "IIS" to determine if any IIS-related errors occurred.

   During the upgrade process, the Windows Server 2003 upgrade creates an error log file (Setuperr.log) for the entire Windows Server 2003 operating system that records any errors encountered during upgrade. In addition to resolving any IIS related-problems encountered during upgrade, resolve any other upgrade problems listed in Setuperr.log before continuing with the upgrade process.

2. If you find any IIS-related errors, open *%systemroot%*\Iis6.log in Notepad and search the log file for "fail" to determine the source of the errors.

   During the upgrade of the IIS components, the Windows Server 2003 upgrade creates an IIS-specific log file, Iis6.log, in the systemroot folder that records any IIS-specific errors encountered during upgrade. Iis6.log entries with the word "fail" reflect problems encountered during the upgrade process. Review Iis6.log and resolve the problems encountered during upgrade before continuing further in the upgrade process.

   Resolve any operating system upgrade-related problems before continuing further in the upgrade process. Subsequent steps in the upgrade process, such as moving the applications to worker process isolation mode, are dependent upon these issues being resolved.

# Backing Up the IIS 6.0 Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Metabase backups created with IIS 4.0 or IIS 5.0 cannot be restored on IIS 6.0. As a result, you cannot use any existing metabase backups of the Web server after the upgrade. After you have verified that IIS 6.0 hosts Web sites and applications as it did before the upgrade, back up the metabase before continuing with the upgrade process.

The remaining steps in the upgrade process focus on hosting the Web sites and applications in worker process isolation mode. Before changing the IIS configuration to worker process isolation mode, verify that you have a backup of the current IIS configuration by backing up the metabase.

In the event that an unforeseen problem occurs while you are configuring IIS 6.0 to run the applications in worker process isolation mode, you can restore the applications to a known operational state. This will provide a known starting place from which to retry the configuration of the server.

For information about how to back up the IIS 6.0 metabase, see Back Up and Restore the IIS Metabase.

| **Note** |
| --- |
| Upon completion of the upgrade process, a backup of the IIS 6.0 metabase is automatically created. However, changes to the Web service extensions list are not reflected in that backup. |

# Converting to Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When the upgrade process is complete, IIS 6.0 is configured to run in IIS 5.0 isolation mode. Although IIS 5.0 isolation mode provides compatibility with existing applications, it is unable to provide all the security, availability, and performance improvements in worker process isolation mode. To take full advantage of these improvements, you must configure IIS 6.0 to run in worker process isolation mode.

After upgrade, most Web sites and applications function in worker process isolation mode without modification. IIS can run in only one application isolation mode at a time — either IIS 5.0 isolation mode or worker process isolation mode. As a result, configuring IIS to run in worker process isolation mode affects all of the Web sites and applications hosted by IIS. Therefore, before configuring IIS to run in worker process isolation mode, you must determine whether your applications are compatible with this isolation mode. If the existing Web sites and applications are not compatible with worker process isolation mode, you can continue to Configuring IIS 6.0 to Host ASP.NET Applications.

Figure 5.4 illustrates the process for converting the Web server to worker process isolation mode.

**Figure 5.4 Converting to Worker Process Isolation Mode**




For more information about worker process isolation mode and IIS 5.0 isolation mode, see Determining Application Compatibility with Worker Process Isolation Mode.

# Documenting the Current Application Isolation Settings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you configure IIS 6.0 to run in worker process isolation mode, document the existing application isolation settings of the Web sites and applications that are hosted by IIS. Later in the upgrade process, you will use this baseline configuration for configuring the application isolation mode for your Web sites and applications.

For each Web site and application currently running on the server, document the following:

**Application isolation settings**
Earlier versions of IIS can host Web sites and applications in pooled or isolated process configurations. For information about how to view the current application isolation mode, see View Application Isolation Configuration.

If you are running IIS 4.0 on Windows NT Server 4.0, your applications are isolated in one of the following ways:

- In-process (running in-process with Inetinfo.exe)
- Isolated (running under Microsoft Transaction Server [MTS])

If you are running IIS 5.0 on Windows 2000, your applications are isolated in one of the following ways:

- In-process (running in-process with Inetinfo.exe)
- Pooled (running in a pooled COM+ package)
- Isolated (running in an isolated COM+ package)

**Process identity that is used by the Web site or application**
Each Web site or application configured in high isolation, or pooled isolation, uses an *identity*. An identity is a user account that provides a security context for worker process servicing the Web site or application. The identity can be used to secure content, by using NTFS permissions, or data, such as data stored in Microsoft SQL Server™. For more information about how to view the identity for each Web site or application, see View Web Site and Application Process Identities.

| **Note** |
| --- |
| All Web sites and applications without identities run under the security context of LocalSystem. |

# Configuring IIS 6.0 to Run in Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

As previously mentioned, configuring the IIS application isolation mode is a Web server-wide configuration setting that affects all Web sites and applications running on the Web server. The server is currently configured in IIS 5.0 isolation mode.

If you determined that one or more of your Web sites or applications are incompatible with worker process isolation mode, leave the Web server in IIS 5.0 isolation mode. Otherwise, configure IIS 6.0 to run in worker process isolation mode either in IIS Manager, or by setting the IIS metabase property **Iis5IsolationModeEnabled** to a value of **False**.

| **Note** |
| --- |
| If you configure IIS 6.0 to run in IIS 5.0 isolation mode and then decide to change the configuration back to worker process isolation mode, the original worker process isolation mode settings are retained. Similarly, if you configure IIS 6.0 to run in IIS 5.0 isolation mode, change to worker process isolation mode, and then change back to IIS 5.0 isolation mode, the IIS 5.0 isolation mode settings are retained. |

For more information about how to configure IIS to run in worker process isolation mode or in IIS 5.0 isolation mode, see Configure Application Isolation Modes.

For more information about determining compatibility with worker process isolation mode, see Evaluating Application Changes Required for Worker Process Isolation Mode.

# Configuring Application Isolation Settings in Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Immediately after configuring IIS to run in worker process isolation mode, you need to configure the application isolation settings to closely approximate their configuration in IIS 5.0 isolation mode by assigning them to *application pools*. An application pool is a grouping of one or more Web sites or applications served by one or more worker processes. You might need to apply additional configurations so that the applications retain their original isolation settings.

After converting to worker process isolation mode, all applications run in the pre-existing application pool named "DefaultAppPool." If all of the applications run in the same process in the previous version of IIS, then they all are assigned to the default application pool.

However, if any one of the applications in the same application pool fails, the other applications can be adversely affected. For this reason it is recommended that you isolate your applications into separate application pools whenever possible.

Configure Web sites and applications to run in their own application pool by completing the following steps:

**For each Web site or application configured in High isolation in IIS 5.0**

1. Create a new application pool to be used by the Web site or application.

   For information about how to create application pools, see Isolate Applications in Worker Process Isolation Mode.

2. If the Web site or application previously ran under an identity that is still required by the Web site or application, configure the application pool to use that same identity.

   For information about how to configure the identity for an application pool, see Configure Application Pool Identity.

3. Assign the Web site or application to the new application pool.

   For information about how to assign the Web site to the new application pool, see Isolate Applications in Worker Process Isolation Mode.

**For each Web site or application configured in Low or Medium isolation in IIS 5.0**
In earlier versions of IIS, applications ran in-process as DLLs in Inetinfo.exe (Low isolation) and the default process identity (account the application runs as) was LocalSystem. With worker process isolation mode in IIS 6.0, applications never run in Inetinfo.exe. However, any applications that are not explicitly assigned to an application pool are assigned to the default application pool, which runs under the NetworkService process identity by default. Because LocalSystem has an elevated security context, run Web sites and applications under the security context of the NetworkService account.

For each Web site or application that ran in Low or Medium isolation in IIS 5.0, do one of the following:

- When the Web site or application is able to function under the identity of the NetworkService account in the default application pool, continue to host the Web sites or applications in the default application pool, named "DefaultAppPool."

- When the Web site or application is unable to function under the identity of the NetworkService account in the default application pool, perform the following steps:

  1. Create a new application pool.

  2. Create a service account to be used as the identity for the application pool.

     For more information about how to create a service account to be used as an identity for an application pool, see Create a Service Account.

3. Configure the application pool identity to use the service account.

   For more information about how to configure the identity for an application pool, see Configure Application Pool Identity.

4. Place the Web site or application in the new application pool.

# Example: Converting to Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A fictitious organization, Contoso, has an existing IIS 5.0 Web server that hosts four Web applications. The administrator plans to upgrade the Web server to IIS 6.0, and has tested the applications for compatibility with IIS 6.0 worker process isolation mode and Windows Server 2003. Table 5.4 lists the existing configuration of the Web applications before upgrading to IIS 6.0.

**Table 5.4 Configuration Before Upgrade**

| Application Name | Request Processing Model | Identity |
|---|---|---|
| Application-A | Isolated | AppIdent-01 |
| Application-B | In-process | LocalSystem |
| Application-C | Isolated | AppIdent-02 |
| Application-D | In-process | LocalSystem |
| Application-E | In-process | LocalSystem |

After the upgrade, the administrator verified that the Web applications continued to run properly in IIS 5.0 isolation mode. Then the administrator configured IIS 6.0 to run in worker process isolation mode. Table 5.5 lists the configuration of the Web applications immediately after configuring IIS to run in worker process isolation mode.

**Table 5.5 Configuration After Converting to Worker Process Isolation Mode**

| Application Name | Application Pool | Application Pool Identity |
|---|---|---|
| Application-A | Default Application Pool | NetworkService |
| Application-B | Default Application Pool | NetworkService |
| Application-C | Default Application Pool | NetworkService |
| Application-D | Default Application Pool | NetworkService |
| Application-E | Default Application Pool | NetworkService |

To approximate the original configuration of the Web applications in worker process isolation mode, the administrator does the following:

1. Creates a new application pool for each application that was configured for isolation.
2. Configures each application pool with the identity assigned previously to the application configured for isolation.
3. Ensures that the identity assigned to each newly created application pool is added to the IIS_WPG local user group.
4. Assigns each application to the corresponding application pool.
5. Continues hosting all other applications in the default application pool.
6. Verifies that the applications in the default application pool properly run under the NetworkService account identity.

**Table 5.6 Final Configuration of Web Sites and Applications**

| Application Name | Application Pool | Application Pool Identity |
|---|---|---|

| | | |
|---|---|---|
| Application-A | AppPool-01 | AppIdent-01 |
| Application-B | Default | NetworkService |
| Application-C | AppPool-02 | AppIdent-02 |
| Application-D | Default | NetworkService |
| Application-E | Default | NetworkService |

# Configuring IIS 6.0 Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Up to this point in the upgrade process, you have upgraded the operating system and all of the operating system components, including IIS 6.0, on the Web server. However, you might need to further configure the IIS 6.0 properties on the Web server so that the Web sites run as they did before the server was upgraded. In addition, you can configure your Web server to take advantage of the enhanced security and availability capabilities of IIS 6.0.

Figure 5.5 illustrates the process for configuring the IIS 6.0 properties on your Web server.

**Figure 5.5 Configuring IIS 6.0 Properties**



# Enabling the WWW Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you upgrade a Web server running Windows 2000 Server and IIS 5.0, the World Wide Web Publishing Service (WWW service) is disabled unless, before upgrading, you elected to run the IIS Lockdown Tool or make the appropriate changes to the registry. However, if you did not choose either of those methods, you must now enable the WWW service.

| **Note** |
| --- |
| If you are upgrading a Web server that is currently running Windows NT Server 4.0 and IIS 4.0, the WWW service is not disabled. Therefore, you can continue to the next step in the deployment process. |

**See Also**

**Concepts**

Enable the WWW Service After Upgrade

# Configuring Web Service Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Many Web sites and applications hosted on IIS include dynamic content and other enhanced capabilities. Providing dynamic content and other enhanced capabilities requires executable code, such as ASP, ASP.NET, and Internet Server API (ISAPI) extensions. The handlers that extend IIS functionality beyond serving static pages are known as *Web service extensions*.

Because of the enhanced security features in IIS 6.0, you can enable or disable individual Web service extensions. After upgrade, all of the Web service extensions are enabled except for the extensions that are mapped to 404.dll by the IIS Lockdown Tool. If you did not run the IIS Lockdown Tool prior to upgrade, all of the Web service extensions are enabled.

The Windows Server 2003 upgrade creates a permission entry for Web service extensions, which enables all of the Web service extensions that are not explicitly prohibited. Enabling all of the Web service extensions ensures the highest possible compatibility with your Web sites. However, doing this creates a security risk by enabling functionality that might not be necessary for your server, which increases the attack surface of the server.

| **Note** |
| --- |
| *Web service extensions* allow you to enable and disable the serving of dynamic content Multipurpose Internet Mail Extensions (MIME) types allow you to enable and disable the serving of static content. For more information about enabling and disabling the serving of static content, see Configuring MIME Types. |

Configure the Web service extensions after upgrade by completing the following steps:

1. Configure the Web service extensions list so that the following entries, which enable all Web service extensions, are set to **Prohibited**:

   - All Unknown Common Gateway Interface (CGI) Extensions
   - All Unknown ISAPI Extensions

   For information about how to prohibit a Web service extension, see Configure Web Service Extensions.

2. Enable the essential predefined Web service extensions based on the information in Table 5.7.

   **Table 5.7 Predefined Web Service Extensions**

| Web Service Extension | Enable When |
| --- | --- |
| Active Server Pages | One or more of the Web sites or applications contains ASP content. |
| ASP.NET version 1.1.4322 | One or more of the Web sites or applications contains ASP.NET content. |
| FrontPage Server Extensions 2002 | One or more of the Web sites are FrontPage extended. |
| Internet Data Connector (IDC) | One or more of the Web sites or applications uses the IDC to display database information (content includes .idc and .idx files). |
| Server-Side Includes | One or more of the Web sites uses server-side include (SSI) directives to instruct the Web server to insert various types of content into a Web page. |
| Web Distributed Authoring and Versioning (WebDAV) | You want to support WebDAV on the Web server. Not recommended for dedicated Web servers. |

3. For each Web service extension that is used by your applications and is not a one of the default Web service extensions, add a new entry to the Web service extensions list and configure the status of the new entry to **Allowed**.

   For example, one of your applications might use an ISAPI extension to provide access to a proprietary database. Add the ISAPI extension to the Web service extensions list, and then configure the status of the ISAPI extension to **Allowed**.

   For information about how to add a Web service extension and enable the extension, see Configure Web Service Extensions.

4. Use a Web browser on a client computer to verify that the Web sites and applications run on the server.

# Configuring MIME Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS serves only the static files with extensions that are registered in the Multipurpose Internet Mail Extensions (MIME) types list. IIS is preconfigured to recognize a default set of global MIME types, which are recognized by all configured Web sites. MIME types can also be defined at the Web site and directory levels, independent of one another or the types defined globally. IIS also allows you to change, remove, or configure additional MIME types. For any static content file extensions used by the Web sites and applications hosted by IIS that are not defined in the MIME types list, you must create a corresponding MIME type entry.

Configure the MIME types after upgrade by completing the following steps:

1. Remove the entry **.* application/octet-stream**, which enables all MIME types. Removing this entry allows you to restrict the static content served by IIS.

   For information about how to remove a MIME type from the list, see Configure MIME Types.

2. For each static file type that is used by your applications, ensure that an entry exists in the MIME types list.

   When your application uses the standard MIME types included in IIS, no new MIME type entry is required. For information about how to add a MIME type, see Configure MIME Types.

3. Use a Web browser on a client computer to verify that the Web sites and applications run on the server.

# Modifying References to IIS 6.0 Metabase Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

There are metabase properties that were used to configure features in earlier versions of IIS that are no longer supported in IIS 6.0. Because some features are eliminated, or implemented differently in IIS 6.0, the corresponding unused metabase properties are not referenced by any code in IIS 6.0. In cases where the feature is implemented differently in IIS 6.0, new metabase properties have been created to replace the obsolete, or unused, property.

In addition, there is one IIS 5.0 metabase property, **CPUResetInterval**, whose behavior has changed because of architectural changes made to IIS 6.0.

To determine whether any of your Web sites, applications, or setup programs reference these changed or unsupported IIS metabase properties, see Appendix B: Changes to Metabase Properties in IIS 6.0. You can follow the recommendations associated with each changed metabase property to accommodate functionality changes.

| **Tip** |
| --- |
| The metabase properties that are no longer supported in IIS 6.0 are not available in IIS 6.0, even when IIS 6.0 is configured to run in IIS 5.0 isolation mode. |

# Upgrading FrontPage Extended Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you upgrade a Web server that has FrontPage 2000 Server Extensions, the upgrade process automatically installs FrontPage 2002 Server Extensions. After upgrade, both FrontPage 2000 Server Extensions and FrontPage 2002 Server Extensions are installed. After upgrade, you must upgrade each of the FrontPage extended Web sites and configure IIS 6.0 so that FrontPage 2000 Server Extensions is prohibited.

Upgrade your FrontPage extended Web sites by completing the following steps:

1. Set the status of the **FrontPage Server Extensions 2000** entry in the Web service extensions list to **Prohibited**.

   For more information about configuring Web service extensions, see Configure Web Service Extensions.

2. For each FrontPage extended Web site, upgrade the Web site to FrontPage 2002 Server Extensions.

   You can upgrade individual Web sites by using the FrontPage 2002 Server Extensions Server Administration Web page. You can also upgrade individual sites, or multiple sites in a batch, by using Owsadm.exe.

   For more information about how to upgrade FrontPage extended Web sites to FrontPage 2002 Server Extensions, see Upgrade FrontPage Extended Web Sites.

# Determining Whether to Run the IIS Lockdown Tool and UrlScan

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

UrlScan and the IIS Lockdown Tool are IIS security related programs designed for IIS 5.1 and earlier. Each tool provides different types of protection for earlier versions of IIS.

**IIS Lockdown Tool**
The IIS Lockdown Tool is provided to assist administrators in configuring optimal security settings for existing IIS servers. You cannot install the IIS Lockdown Tool after migration because all of the default configuration settings in IIS 6.0 meet or exceed the security configuration settings made by the IIS Lockdown Tool.

**UrlScan**
UrlScan is a tool that is provided to reduce the attack surface of Web servers running earlier versions of IIS. By default, IIS 6.0 has features that significantly improve security by reducing the attack surface of the Web server. UrlScan provides flexible configuration for advanced administrators, while maintaining the improved security in IIS 6.0. When you need this flexibility in configuring your Web server, you can run UrlScan on IIS 6.0.

For more information about determining whether to run UrlScan after migrating your server to IIS 6.0, and to download the latest version of the URLScan tool, see Knowledge Base article 307608, Using UrlScan on IIS.

# Making Security-Related Configuration Changes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After upgrading your server to IIS 6.0, you can make additional security-related configuration changes on the Web server. If you ran the IIS Lockdown Tool before upgrading the Web server, most of these changes are already in place. The IIS Lockdown Tool removes unnecessary IIS components, including virtual directories, to reduce the attack surface available to malicious users. Otherwise, make these security-related configuration changes to help reduce the attack surface and increase the security of the Web server.

Make the security-related configuration changes by completing the following steps:

1. Enable essential IIS components and services.

2. Remove unnecessary IIS virtual directories.

3. Configure the anonymous user identity.

# Enabling Essential IIS Components and Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 includes other components and services in addition to the WWW service, such as the FTP service and SMTP service. IIS components and services are installed and enabled by means of the **Application Server** subcomponent in **Add or Remove Windows Components**. After installing IIS, you must enable the IIS 6.0 components and services required by the Web sites and applications running on the Web server.

Enable only the essential IIS 6.0 components and services required by your Web sites and applications. For more information about enabling the essential IIS protocols and services see Enabling Only Essential IIS Components and Services.

# Removing Unnecessary IIS Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following table lists the virtual directories that can be removed from IIS 6.0. Compare the virtual directories in IIS Manager to the virtual directories in this table, and then remove those directories.

**Virtual Directories that Can Be Removed After Upgrade**

| Virtual Directory | Description |
|---|---|
| IISAdmin | Provides an HTML-based administration tool. It is designed primarily for administrators who administer the Web server through a Web interface, but IIS Manger is the recommended method for Web server administration. The virtual directory is not installed as part of IIS 6.0, but it is not removed when a Web server running an earlier version of IIS is upgraded to IIS 6.0. |
| IISADMPWD | This virtual directory allows you to reset passwords from Windows NT Server 4.0 and Windows 2000 Server. It is designed primarily for intranet scenarios and is not installed as part of IIS 6.0, but it is not removed when a Web server running IIS 4.0 is upgraded to IIS 6.0. For more information about this functionality, see Knowledge Base article 184619, How to Change Windows NT Account Passwords Using Internet Information Server (IIS) 4.0. |
| IISHelp | Provides an HTML version of the IIS documentation. It is designed primarily for application developers and is not installed as part of IIS 6.0, but it is not removed when a Web server running an earlier version of IIS is upgraded to IIS 6.0. |
| MDAC | Contains sample applications that illustrate the use of Microsoft Data Access Components that are not required on production Web servers. It is designed primarily for application developers and is not installed as part of IIS 6.0, but it is not removed when a Web server running an earlier version of IIS is upgraded to IIS 6.0. |
| IISSamples | Contains sample applications that are not required on production Web servers. It is designed primarily for application developers and is not installed as part of IIS 6.0, but it is not removed when a Web server running an earlier version of IIS is upgraded to IIS 6.0. |
| Printers | Provides an HTML-based printer administration tool. It is designed primarily for administrators who administer printers through a 2003 printer administration Web interface, but using the Windows Server interface is the recommended method for printer administration. The virtual directory is not installed as part of IIS 6.0, but it is not removed when a Web server running an earlier version of IIS is upgraded to IIS 6.0. |
| Scripts | Contains scripts that are used for the sample applications in other virtual directories and is required on production Web servers. It is designed primarily for application developers and is not installed as part of IIS 6.0, but it is not removed when a Web server running an earlier version of IIS is upgraded to IIS 6.0. |

For more information about how to remove unnecessary virtual directories after upgrade, see Remove Virtual Directories.

# Configuring the Anonymous User Identity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When the Web sites and applications running on the Web server require anonymous access, IIS is configured with a user account specifically for anonymous access. When a user connects to the Web server anonymously, IIS creates a process token for the user based on the user account that you configured in the anonymous user identity. The user account can be stored in the local account database on the Web server, or in a domain.

If, before upgrade, the anonymous user identity is configured to use a domain-based user account, after upgrade you need to configure the anonymous user account to the same domain-based user account. This is because the upgrade process automatically configures IIS to user the default anonymous account IUSR_*computername*, where *computername* is the name of the computer on which IIS is running. You can configure the anonymous user identity to the domain-based user account in IIS Manager.

For more information about how to configure the anonymous user identity, see Configure Anonymous User Identity.

# Configuring IIS 6.0 to Host ASP.NET Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If the newly upgraded IIS 6.0 server includes existing ASP.NET applications, you need to configure IIS to use the correct version of the .NET Framework, and you must configure the .NET Framework to support your applications. If IIS 6.0 is configured to run in worker process isolation mode, you need to migrate ASP.NET configuration settings from the Machine.config file to the equivalent settings in IIS 6.0.

Figure 5.6 illustrates the process for configuring IIS for ASP.NET applications

**Figure 5.6 Configuring IIS for ASP.NET Applications**



# Configuring IIS 6.0 to Use the Correct Version of the .NET Framework

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before upgrading a Web server that hosts ASP.NET applications, version 1.0 of the .NET Framework is installed on the server and the applications are configured to use that version of the .NET Framework. After upgrade, version 1.1 of the .NET Framework is installed and the applications are configured to use version 1.1 of the .NET Framework. However, version 1.0 of the .NET Framework is still installed, which is referred to as side-by-side support.

Running versions 1.0 and 1.1 of the .NET Framework side-by-side is only supported when IIS is configured to run in IIS 5.0 isolation mode. If you have already configured IIS to run in worker process isolation mode, then you can only use version 1.1 of the .NET Framework. In most cases, ASP.NET applications function correctly with version 1.1 of the .NET Framework. For more information about possible application incompatibilities when upgrading from version 1.0 to version 1.1 of the .NET Framework, see Determining Application Compatibility with the .NET Framework. When your ASP.NET application is incompatible with version 1.1 of the .NET Framework, configure the application to use version 1.0 of the .NET Framework and configure IIS to run in IIS 5.0 isolation mode.

You can configure each ASP.NET application to use a specific version of the .NET Framework by registering a *script map* in IIS for the application. A script map associates a file name extension and HTTP verb with the appropriate Internet Server API (ISAPI) for script handling. For example, when IIS receives a request for a .aspx file, the script map for the corresponding application directs IIS to forward the requested file to the appropriate version of the ASP.NET ISAPI for processing.

The script map for each ASP.NET application can be applied directly to an application, or inherited from a parent application. However, ASP.NET supports only one version of the .NET Framework for each application pool. For more information about how to configure the script map for an ASP.NET application, see Configure an ASP.NET Application for ASP.NET.

# Configuring the .NET Framework

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The configuration method for the .NET Framework is determined by the application isolation mode that you use to configure IIS 6.0. This table lists the methods for configuring the .NET Framework that are associated with each IIS 6.0 application isolation mode.

**Methods for Configuring the .NET Framework**

| Application Isolation Mode | Configuration Method for the .NET Framework |
|---|---|
| IIS 5.0 isolation mode | Configured by making changes to the Machine.config file in the *systemroot*\Microsoft.NET\Framework*VersionNumber*\Config folder. |
| Worker process isolation mode | Configured by making changes to the IIS 6.0 metabase. |

When IIS 6.0 is configured to run in IIS 5.0 isolation mode, the .NET Framework uses the <**processModel**> section of the Machine.config file (in the *systemroot*\Microsoft.NET\Framework*versionNumber*\Config folder) for its configuration and no additional steps are required.

However, if you configured IIS 6.0 to run in worker process isolation mode, the .NET Framework largely ignores the **<processModel>** section of the Machine.config file, and instead gets its process configuration from the IIS 6.0 metabase. Because the upgrade process does not migrate the existing settings in the Machine.config file, you must manually convert any settings required by the ASP.NET applications.

For information about how to convert the Machine.config attribute settings to IIS 6.0 metabase property settings, see Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings. For more information about configuring IIS 6.0 for ASP.NET applications, see Deploying ASP.NET Applications in IIS 6.0.

# Reviewing How ASP.NET Applications Run In Each Application Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When running in worker process isolation mode, ASP.NET applications use the W3wp.exe worker process and application pool properties, which are stored in the IIS 6.0 metabase. When you configure IIS 6.0 to run in IIS 5.0 isolation mode, ASP.NET applications use the ASP.NET request processing model, Aspnet_wp.exe, and configuration settings. These configuration settings are stored in the Machine.config file.

**Behavior of ASP.NET Applications that Are Running in IIS 5.0 Isolation Mode**
By default, ASP.NET applications are configured to run in worker process isolation mode. If your application can only run in the ASP.NET process model, you must configure the server to run in IIS 5.0 isolation mode to be able to run the application on IIS 6.0. When IIS 6.0 is configured to run in IIS 5.0 isolation mode, ASP.NET applications behave as follows:

- The applications run within Aspnet_wp.exe.

  Aspnet_wp.exe is a request processing model similar to worker process isolation mode and contains similar worker process management capabilities as the WWW service in IIS 6.0. Aspnet_wp.exe includes most of the IIS application management features, such as recycling, health detection, and *processor affinity*— the ability to force worker processes to run on specific microprocessors.

- The configuration settings are stored in the Machine.config file.

  When configured in IIS 5.0 isolation mode, the configuration settings for ASP.NET applications are managed by modifying the Machine.config file, not the IIS metabase file. Because there is no administrative console for the Machine.config settings, any configuration settings for ASP.NET must be made directly to the Machine.config file.

| **Important** |
| --- |
| In IIS 5.0 isolation mode, the .NET Framework ignores any configuration changes made in the IIS metabase. Administrative consoles, such as IIS Manager, make changes to the IIS 6.0 metabase, but those changes are not read by the .NET Framework. |

When IIS 6.0 is configured in IIS 5.0 isolation mode, your ASP.NET applications should behave as they did in IIS 5.0. However, incompatibilities can result when running version 1.1 of the .NET Framework. For more information about configuring IIS to support ASP.NET applications that use version 1.0 of the .NET Framework, see Configuring IIS 6.0 to Use the Correct Version of the .NET Framework.

**Behavior of ASP.NET Applications that Are Running in Worker Process Isolation Mode**
When IIS 6.0 is configured to run in worker process isolation mode, ASP.NET applications behave as follows:

- The process model within the ASP.NET Internet Server API (ISAPI) extension is disabled and ASP.NET applications run using worker process isolation mode in IIS 6.0.

  In this configuration, the ASP.NET application runs in worker process isolation mode like any other application, such as an ASP application. In addition, IIS 6.0 provides all of the management features, such as recycling, health detection, and processor affinity.

- The ASP.NET ISAPI extension is configured by a combination of configuration settings that are stored in the IIS metabase (MetaBase.xml) and configuration settings in the Machine.config file.

  When configured in worker process isolation mode, the majority of the configuration settings are stored in the IIS metabase. You can adjust these settings directly in MetaBase.xml or from administrative consoles, such as IIS Manager, or scripts.

However, if there are existing settings in the **<processModel>** section of the Machine.config file, those configuration settings must be converted to the appropriate application pool settings when the Web server is configured to run in worker process isolation mode Additionally, there are other configuration settings that are still made through the Machine.config file, regardless of the application mode. For more information about converting Machine.config attributes to worker process isolation mode settings, see Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings.

When IIS 6.0 is configured in worker process isolation mode, your ASP.NET applications should behave as they did on IIS 5.0. Before deploying your ASP.NET applications on your production Web servers, test compatibility with IIS 6.0 running in worker process isolation mode and version 1.1 of the .NET Framework. For more information about determining application compatibility with IIS 6.0 running in worker process isolation mode, see Determining Application Compatibility with Worker Process Isolation Mode. For more information about determining compatibility with version 1.1 of the .NET Framework see Determining Application Compatibility with the .NET Framework.

If you determine that your ASP.NET applications are incompatible with IIS 6.0 running in worker process isolation mode, reconfigure IIS to run in IIS 5.0 isolation mode. If your ASP.NET applications are incompatible with version 1.1 of the .NET Framework, configure IIS to use version 1.0 of the .NET Framework with your ASP.NET application.

For more information about configuring IIS 6.0 to run in IIS 5.0 isolation mode, see Configure Application Isolation Modes.

# Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you upgrade a server running IIS 5.0 and version 1.0 of the .NET Framework, the ASP.NET applications can depend on specific attribute settings in the **<processModel>** section of the Machine.config file. The **<processModel>** section configures the ASP.NET process model settings and therefore affects all ASP.NET applications running on an IIS Web server. If any of the following Machine.config attributes exist, you must migrate the attribute configuration settings to their equivalent metabase property settings in IIS 6.0.

| **Note** |
| --- |
| Follow the deployment steps in this section only when IIS 6.0 is configured to run in worker process isolation mode. When IIS 6.0 is configured to run in IIS 5.0 isolation mode, IIS 6.0 reads the ASP.NET settings from the Machine.config file, and you do not need to migrate the Machine.config attribute settings. Thus, you can continue to the next step in the upgrade process. |

Migrate the Machine.config attributes by adjusting the configuration settings in the Application Pools node in IIS Manager. The following table lists the Machine.config attributes and where to find more information about how to migrate the attributes.

**Machine.config Attributes and How To Migrate Them**

| Machine.config Attribute | Migrating This Attribute |
| --- | --- |
| **cpuMask** | See Migrating Performance-Related Attributes. |
| **idleTimeout** | See Migrating Performance-Related Attributes. |
| **memoryLimit** | See Migrating Recycling-Related Attributes. |
| **password** | See Migrating Identity-Related Attributes. |
| **pingFrequency** | See Migrating Health-Related Attributes. |
| **pingTimeout** | See Migrating Health-Related Attributes. |
| **requestLimit** | See Migrating Recycling-Related Attributes. |
| **restartQueueLimit** | See Migrating Performance-Related Attributes. |
| **shutdownTimeout** | See Migrating Health-Related Attributes. |
| **timeout** | See Migrating Recycling-Related Attributes. |
| **username** | See Migrating Identity-Related Attributes. |
| **webGarden** | See Migrating Performance-Related Attributes. |

The Machine.config attribute settings in Table 5.10 must be migrated to IIS 6.0 settings. However, you must still configure the following attribute settings in the Machine.config file:

- **maxWorkerThreads**
- **maxIoThreads**
- **responseDeadlockInterval**

When configured to run in worker process isolation mode, IIS 6.0 and the .NET Framework ignore any **<processModel>** section Machine.config attribute settings that are not in the previous list or in the table above.

| **Note** |
| --- |
| The **<processModel>** section of the Machine.config file configures the ASP.NET process model settings on an IIS Web server. The **<processModel>** section provides global configuration settings for all of the ASP.NET applications running on the server. |

# Migrating Recycling-Related Attributes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In worker process isolation mode, you can assign one or more Web sites or applications to an application pool. Each application pool has one or more worker processes that process client requests for the Web sites and applications in the application pool. You can configure IIS to periodically recycle worker processes assigned to an application pool, which in turn recycles the Web sites and applications running in that application pool. Recycling worker processes keeps problematic Web sites and applications running smoothly, especially when it is not feasible to modify the application code. Recycling worker processes within an application pool ensures that the Web sites and applications in the application pools remain healthy, and that system resources can be recovered.

The Machine.config attributes that affect recycling-related metabase properties include:

- **timeout**
- **requestLimit**
- **memoryLimit**

#### Migrating the timeout Attribute

The **timeout** Machine.config attribute, default value infinite, specifies the time limit after which ASP.NET starts a new worker process to take the place of the current one. The attribute is a string value in the format of hr:min:sec. To configure this value in IIS 6.0, you need to convert this format to time in minutes.

Configure the Recycling Worker Processes (in minutes) setting in IIS 6.0 to the same time value, in minutes, that is configured in the **timeout** Machine.config attribute.

For more information about how to configure the Recycling Worker Processes (in minutes) setting, see Configure Application Pool Recycling.

#### Migrating the requestLimit Attribute

The **requestLimit** Machine.config attribute, with a default value of infinite, specifies the number of requests after which ASP.NET starts a new worker process to take the place of the current one.

Configure the **Recycling Worker Processes (in requests)** setting in IIS 6.0 to the same value, in number of requests, that is configured in the **requestLimit** Machine.config attribute.

For more information about how to configure the **Recycling Worker Processes (in requests)** setting, see Configure Application Pool Recycling.

#### Migrating the memoryLimit Attribute

The **memoryLimit** Machine.config attribute, with a default value of 60 percent, specifies the percentage of physical memory that the worker process can consume before ASP.NET starts a new worker process to take the place of the current one.

Recycling based on memory consumption works slightly different in IIS 6.0, where worker processes can be recycled based on both virtual memory and physical memory. Also, the memory limit is not a percentage value, but is specified as a value in megabytes.

Configure the **Maximum used memory (in megabytes)** setting in IIS 6.0 to the amount of physical memory represented by the percentage specified in the **memoryLimit** Machine.config attribute.

Also, adjust the **Maximum virtual memory (in megabytes)** settings in IIS 6.0 if necessary. The **Maximum virtual memory (in megabytes)** is the maximum amount of virtual memory, which includes the used memory plus the reserved memory,

For more information about how to configure the **Maximum used memory (in megabytes)** setting, see Configure Application Pool Recycling.

# Migrating Performance-Related Attributes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In worker process isolation mode, IIS 6.0 can be configured to optimize the performance of an application pool, allowing you optimize the performance of your Web applications. Migrating the performance-related Machine.config attributes helps ensure that your Web applications perform as they did on IIS 5.0.

| **Note** |
| --- |
| In addition to the configuration settings mentioned in this section, IIS 6.0 provides CPU utilization monitoring for worker processes. For more information, see Configuring Application Pools in IIS 6.0. |

The Machine.config attributes that affect performance-related metabase properties include:

- **webGarden** and **cpuMask**
- **idleTimeout**
- **restartQueueLimit**

**Migrating the webGarden and cpuMask Attributes**

The combination of the **webGarden** and **cpuMask** Machine.config attributes provide configuration for ASP.NET on IIS 5.0 for *Web gardens*. Web gardens are created when there is more than one worker process servicing applications. In IIS 5.0, Web gardens are created for use by all applications running on the Web server. In IIS 6.0, Web gardens are created within each application pool.

Web gardens in IIS 5.0 allow ASP.NET to schedule a separate worker process for each microprocessor on a Web server with multiple microprocessors. In contrast, you can create a Web garden on IIS 6.0 by specifying multiple worker processes for an application pool, regardless of the number of microprocessors in the Web server. For more information about Web gardens in IIS 6.0, see Configuring Web Gardens.

The **webGarden** Machine.config attribute, which has a default value of **False**, controls CPU affinity when used in conjunction with the **cpuMask** attribute. When the **webGarden** attribute is set to:

- **True .** ASP.NET creates a separate worker process for each processor specified in the **cpuMask** Machine.config attribute.
- **False .** ASP.NET creates only one worker process and the **cpuMask** Machine.config attribute is ignored.

The **cpuMask** Machine.config attribute, default value 0xffffffff, specifies which microprocessors on a multiprocessor Web server are able to run worker processes initiated by ASP.NET. By default, all processors are able to run worker processes initiated by ASP.NET. The **cpuMask** attribute also contains a bitmask value that indicates the microprocessors that are able to run ASP.NET processes. When a bit in the **cpuMask** bitmask value is set to 1, the corresponding microprocessor is able to run ASP.NET processes. You need to record the number of bitmasks enabled to convert this to the correct number of application pools.

For example, if **cpuMask** is set to a value of 0x0d, the equivalent of the binary bit pattern is 1101. If your Web server has four microprocessors, the bitmask value 1101 indicates that you need three worker processes.

Configure the **Maximum number of worker processes** setting in IIS 6.0 to the number of microprocessors that are enabled in the **cpuMask** Machine.config attribute. The default value for **Maximum number of worker processes** is 1, indicating that Web gardens are not enabled.

For more information about how to configure the **Maximum number of worker processes** setting, see Configure Application Pool Performance.

**Migrating the restartQueueLimit Attribute**

The **restartQueueLimit** Machine.config attribute, which has a default value of 10, specifies the maximum number of requests that will be queued in ASP.NET while waiting for the worker process to recycle after an abnormal termination.

Configure the **Limit the kernel request queue to** setting in IIS 6.0 to the same value that is configured in the **restartQueueLimit** Machine.config attribute.

For more information about how to configure the **Limit the kernal request queue to** setting, see Configure Application Pool Performance.

| **Note** |
| --- |
| When application pool queue length limits are enabled, IIS monitors the number of requests for a designated application pool queue before queuing a new request. If adding a new request to the queue causes it to exceed the queue length limit, the server rejects the request and sends a 503-error response (that cannot be customized) to the client. However, requests that are already queued remain in the queue even if the limit is changed to a value that is less than the current queue length. |

**Migrating the idleTimeout Attribute**

The **idleTimeout** Machine.config attribute, with a default value of infinite, specifies the maximum number of minutes that a worker process can be inactive before it is automatically shut down by ASP.NET. The attribute is a string value in the format of hr:min:sec. To configure this value in IIS 6.0, you need to convert this format to time in minutes.

Configure the **Shutdown worker processes after being idle for** setting in IIS 6.0 to the same time value, in minutes, that is configured in the **idleTimeout** Machine.config attribute.

For more information about how to configure the **Shutdown the worker processes being idle for** setting, see Configure Application Pool Performance.

# Migrating Health-Related Attributes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By detecting the degree of stability (or health) of an ASP.NET application, IIS determines whether corrective action is required. In IIS 5.0, ASP.NET detects an unhealthy worker process. In IIS 6.0, the WWW service detects an unhealthy worker process that has terminated abnormally if all of the available IIS threads in the worker process assigned to the application pool are blocked.

The Machine.config attributes that affect health-related metabase properties include:

- **shutdownTimeout**
- **pingFrequency** and **pingTimeout**

**Migrating the shutdownTimeout Attribute**

The **shutdownTimeout** Machine.config attribute, default value five seconds, specifies the maximum number of seconds allotted for a worker process to shut down before ASP.NET automatically shuts it down. The attribute is a string value in the format of hr:min:sec. To configure this value in IIS 6.0, you need to convert this format to time in seconds.

Configure the Worker process must shutdown setting in IIS 6.0 to the same time value, in minutes, that is configured in the **shutdownTimeout** Machine.config attribute.

For more information about how to configure the Worker process must shutdown setting, see Configure Application Pool Health.

**Migrating the pingFrequency and pingTimeout Attributes**

In IIS 5.0, the combination of the **pingFrequency** and **pingTimeout** Machine.config attributes specifies whether ASP.NET should monitor the health of worker processes. If IIS 6.0 is running in worker process isolation mode, the WWW service periodically monitors the health of a worker process.

The **pingFrequency** Machine.config attribute, default value 30 seconds, specifies the interval at which ASP.NET queries the worker processes to determine if each worker process is still operating correctly. The **pingFrequency** Machine.config attribute is a string value in the format of hr:min:sec. To configure this value in IIS 6.0, you need to convert this format to time in seconds.

The **pingTimeout** Machine.config attribute, default value five seconds, specifies the interval after which ASP.NET restarts a worker process if there was no response to the most recent query. If a worker process does not respond to the "ping" within the value set by the **pingTimeout** Machine.config attribute, the worker process is restarted. The **pingTimeout** Machine.config attribute is a string value in the format of hr:min:sec. To configure this value in IIS 6.0, you need to convert this format to time in seconds.

You can configure the equivalent settings in IIS 6.0 by making changes to the equivalent IIS metabase properties, listed in the following table.

**Equivalent IIS 6.0 Metabase Properties for pingFrequency and pingTimeout**

| Machine.config Attribute | Equivalent IIS Metabase Property |
|---|---|
| **pingFrequency** | **PingInterval** |
| **pingTimeout** | **PingResponseTime** |

| Note |
|---|
| Unless required by your ASP.NET applications, use the default settings in IIS 6.0. In IIS 6.0, the default setting for **PingInterval** is 30 seconds and the default setting for **PingResponseTime** is 90 seconds. |

You can configure the **PingInterval** metabase property through the Ping worker process setting for application pools in IIS 6.0 Manager. For more information about how to configure the Ping worker process setting, see Configure Application Pool Health.

You can only configure **PingResponseTime** by directly modifying the metabase or programmatically through Windows

Management Instrumentation (WMI) or Active Directory® Service Interfaces (ADSI). For more information about directly modifying the metabase, see Modify the IIS Metabase Directly.

# Migrating Identity-Related Attributes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 5.0, specifying an identity allowed a worker process in ASP.NET to use a Windows identity other than that of the default *process identity*. Process identity is an operating system term used to denote the account that a process runs under. Every running process on a Windows Server 2003 has a process identity that is used to control access to resources on the Web server.

| **Note** |
| --- |
| In IIS 6.0, the application pool identity allows worker processes servicing the application pool to use an identity other than the default identity NetworkService. In IIS 5.0, the default identity is LocalSystem. |

The Machine.config attributes that affect identity-related metabase properties include:

- **username**
- **password**

**Migrating the username Attribute**

The **username** Machine.config attribute specifies the user account used by ASP.NET as an identity for worker processes. The attribute is a string value and does not exist in Machine.config by default.

In IIS 5.0, the **username** Machine.config attribute has the following values:

- **No entry**.When the **username** attribute does not exist in the Machine.config file, this causes ASP.NET to run worker processes under the identity of LocalSystem.
- **Machine**. Causes ASP.NET to run worker processes under a user account named ASPNET that is created automatically when ASP.NET is installed. This is the default configuration.
- **System**. Causes ASP.NET to run worker processes under a user account named System that is created automatically when ASP.NET is installed and allows ASP.NET processes to have full administrative privileges. Applications running under the identity of the System account have unconstrained privileges on the Web server. Run applications under the identity of the System account only when required by your applications.
- **Configured account**. When the **username** attribute contains a service account, this causes ASP.NET to run worker processes under the identity of the service account.

Configure the equivalent settings in IIS 6.0 by selecting an IIS 6.0 configuration setting that allows the ASP.NET application to run properly while not compromising the security of the Web server. This selection is done in an order of preference, from most secure to least secure.

Configure the User name settings in IIS 6.0 by selecting one of the following in order of security preference:

- **Option 1.** Configure IIS 6.0 to use NetworkService.
- **Option 2.** Configure IIS 6.0 to use a new service account and grant the account the minimal user rights or group membership to allow the applications to run successfully.
- **Option 3.** Configure IIS 6.0 to use a new service account that belongs to the local Administrators group.

**Option 1 Configure IIS 6.0 to use NetworkService**

This is the default identity for IIS 6.0 and is the recommend identity. Most ASP.NET applications can run by using this configuration. Test your ASP.NET application to ensure proper operation with NetworkService as the identity.

For more information about how to configure the **username** setting, see Configure Application Pool Identity

**Option 2 Configure IIS 6.0 to use a new service account and grant the account the minimal user rights or group membership to allow the applications to run successfully**
Select this option when the ASP.NET application is unable to run under the NetworkService identity. You might need to do this when the username Machine.config attribute does not exist or is set to Machine, System, or to a configured account.

For more information about how to create an account see Create a Service Account. For more information about how to grant user rights, see Grant User Rights to a Service Account. For more information about how to configure the **username** setting, see Configure Application Pool Identity.

**Option 3 Configure IIS 6.0 to use a new service account that belongs to the local Administrators group**
As a last resort, select this option when the ASP.NET application requires an identity that is a member of the local Administrators group. You might need to do this when the **username** Machine.config attribute does not exist, is set to System, or is set to a configured account that is a member of the local Administrators group.

For more information about how to create an account see Create a Service Account. For more information about how to make a service account a member of the local Administrators group, see Make a Service Account a Member of the Local Administrators Group. For more information about how to configure the **username** setting, see Configure Application Pool Identity.

**Migrating the password Attribute**
The **password** Machine.config attribute specifies the password for the user account used by ASP.NET as an identity for worker processes. The attribute is a string value and does not exist in Machine.config by default, which indicates that the worker processes should run under the default identity of LocalSystem in IIS 5.0.

If you selected NetworkService for the identity for the worker processes, then no password is required. If you decided to create a service account, then configure the **password** setting in IIS 6.0 with the password for the service account.

For more information about how to configure the **password** setting, see Configure Application Pool Identity.

# Completing the Upgrade_IIS_SP1_Dep

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

At this point in the process, your Web sites and applications have been upgraded on the server. Also, the IIS 6.0 properties have been configured so that you are now ready to verify that the Web sites and applications are running properly. After you confirm that the Web sites and applications are running properly, you can back up the Web server and enable client access. When you complete these steps, the upgrade process is complete.

Figure 5.7 illustrates the process for completing the upgrade to IIS 6.0.

**Figure 5.7 Completing the Upgrade to IIS 6.0**



# Verifying That the Web Sites and Applications Run Properly

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before deploying your server to a production environment, verify that the Web site content and application configuration information upgraded successfully.

Verify that your Web sites and applications are running properly by completing the following steps:

1. Review the system log in Windows Server 2003 on the Web server to determine if any of the Web sites did not start.

   IIS 6.0 creates entries in the system log when a Web site fails to start for any reason. Search the System log on the target server to determine if any errors occurred. For more information about how to troubleshoot Web sites that fail to start, see Troubleshooting IIS 6.0.

2. Run your existing validation procedures against the Web sites and applications that run on the server.

   Because the network adapter that connects the Web server to the clients is disabled, you might need to directly connect the Web server to a client computer and enable the network adapter to validate the Web sites and applications.

   After running your existing validation procedures, disable the network adapter that connects the Web server to the clients and reconnect the Web server as it was originally.

| **Note** |
| --- |
| See Troubleshooting IIS 6.0 for assistance in resolving any difficulties in running the applications. |

# Backing Up the IIS 6.0 Server before Enabling Client Access

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you enable client access to the Web server, perform a complete image backup. The purpose of performing this image backup is to provide a point-in-time snapshot of the Web server. If you need to restore the target Web server in the event of a failure, you can use this backup to restore the Web server to a known configuration.

| **Important** |
| --- |
| Do not continue to the next step until you have a successful backup of the entire Web server. Otherwise, you can lose Web sites, applications, or data that you changed after upgrading the Web server. |

For more information about how to back up the Web server, see Back Up and Restore the Web Server to a File or Tape.

# Enabling Client Access to the IIS 6.0 Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you have upgraded the Web server, you are ready to enable client access to the Web sites. During the upgrade process, you disabled the network adapter on the Web server to prevent users from accessing the Web server during the upgrade process. Now that you know the upgrade is completed successfully, you can re-enable the network adapters.

Enable client access to the Web server by completing the following steps:

1. Enable the network adapter used by clients to access the Web server.

   For more information about how to enable the network adapter used by clients, see Enable Network Adapters.

2. Monitor client traffic to determine if clients are accessing the Web server.

   For more information about how to monitor client traffic to Web sites on the Web server, see Monitor Active Web and FTP Connections.

3. Establish a monitoring period, such as a few hours or a day, to confirm that clients that are accessing Web sites and applications on the Web server are experiencing the response times and application responses that you expected.

# Additional Resources About Upgrading to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this section.

**Related Information**

- See Deploying ASP.NET Applications in IIS 6.0 for information about deploying ASP.NET applications in IIS.
- Ensuring Application Availability for information about improving the availability of Web sites and applications in IIS.
- See Migrating IIS Web Sites to IIS 6.0 for information about migrating Web sites to a new Web server running IIS 6.0.
- For information about migrating Apache Web sites, see Migrating Apache Web Sites to IIS 6.0.
- See Securing Web Sites and Applications for information about securing Web sites and applications in IIS.
- Designing a Test Environment for more information about planning a test environment.
- Compatibility Considerations and Version Changes for a current list of possible compatibility issues when upgrading from version 1.0 to version 1.1 of the .NET Framework.
- Hardware Driver Quality for information about the hardware devices that are supported on Windows Server 2003.
- Using UrlScan on IIS for information about determining whether to run UrlScan after migrating your server to IIS 6.0.

**Related IIS 6.0 Help Topics**

- Troubleshooting in IIS 6.0 for information about how to troubleshoot Web sites that fail to start.

**Related Tools**

- IIS Lockdown Tool to download the latest version of the IIS Lockdown Tool.
- Windows Application Compatibility to download the latest version of the Windows Application Compatibility Toolkit.

# Migrating IIS Web Sites to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can move Web sites and applications to Internet Information Services (IIS) 6.0 in two ways. *Migrating* is installing the Microsoft® Windows® Server 2003 operating system and IIS 6.0 on a new server and then moving, or reinstalling, existing Web sites and applications to that server. *Upgrading* is installing Windows Server 2003 and IIS 6.0 on an existing server running the Microsoft Windows NT® Server 4.0 operating system and IIS 4.0 or the Microsoft Windows® 2000 Server operating system and IIS 5.0. Both processes involve minimal outage of service to users who access the Web sites and applications, while retaining the majority of the original configuration settings and fully preserving the content of the Web sites and applications.

**In This Section**


Overview of Migrating IIS Web Sites to IIS 6.0

Preparing for Migration

Deploying the Target Server

Migrating Web Sites with the IIS 6.0 Migration Tool

Migrating Web Sites Manually

Configuring IIS 6.0 Properties

Performing Application-Specific Migration Tasks

Completing the Migration_IIS_SP1_Dep

Additional Resources About Migrating IIS Web Sites to IIS 6.0

# Overview of Migrating IIS Web Sites to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You begin the migration process by determining the compatibility of your existing Web sites and applications with IIS 6.0 and the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; or Windows® Server 2003, Web Edition operating system. Next, you install Windows Server 2003 and IIS 6.0 on the target server, which is the server that will host your Web sites after migration. Then, you migrate the Web site content and configuration settings from the source server, which is a server running the Microsoft Windows NT® Server 4.0 operating system and IIS 4.0 or the Windows 2000 Server operating system and IIS 5.0, to the target server.

After the migration of your Web site content, you customize the configuration of IIS 6.0, based on your Web sites and applications. Finally, after you have completed the customization of IIS 6.0, you back up the target server, enable client access to the Web sites and applications on the target server, and decommission the source server.

The process in this section focuses on transferring the Web site content and configuration settings, and not on the details of how to make application code changes in dynamic content. If your Web sites contain only static content, you can most likely complete the migration process in a few simple steps. However, if your IIS Web sites contain dynamic content, such as Active Server Pages (ASP), ASP.NET, or Common Gateway Interface (CGI) scripts, you might need to modify the code in the dynamic content separately. In addition, any provisioning scripts or setup programs for your existing Web sites and applications might need to be modified after the migration process is complete. Ensure that you test any modifications after the migration process is complete. For more information about potential modifications, see Preparing for Migration later in this section.

**Dependencies**

Also, if the existing Web sites and applications depend on software other than the Windows operating system and IIS, the complexity of the migration process increases. For example, the process for migrating a Web server that hosts Web sites and applications that were designed to run on Windows 2000 Server and IIS 5.0 is relatively simple. On the other hand, the process for migrating a Web server that hosts Web sites, applications, and other software — such as Microsoft Commerce Server, Microsoft BizTalk® Server, monitoring software, custom applications, or other non-Microsoft software — is more difficult because all of the software must be compatible with Windows Server 2003 and IIS 6.0.

In some cases, you can simplify the Web site migration process by using the IIS 6.0 Migration Tool. For more information about using the IIS 6.0 Migration Tool, see Preparing for Migration later in this section.

# Process for Migrating IIS Web Sites to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The process for migrating Web sites hosted on IIS consists of preparing for and performing the migration. Before performing the migration, you need to evaluate the compatibility of the software installed on your existing Web server (including software that generates dynamic content, database connections, and any non-Microsoft software) with Windows Server 2003 and IIS 6.0. You can then perform the migration manually or with the IIS 6.0 Migration tool. After the migration is complete, you must further customize the configuration of IIS 6.0. Figure 6.1 illustrates the process for migrating existing IIS Web sites to IIS 6.0.

**Figure 6.1 Migrating IIS Web Sites to IIS 6.0**




| **Tip** |
| --- |
| To migrate a Web farm, use the process described in this section to migrate Web sites from each Web server in the source Web farm. Then, use provisioning or Web site staging software to propagate the migrated content and site configuration to other Web servers in the target Web farm. |

Depending on your familiarity with Windows server operating systems, IIS, and the migration process, you might require less information to complete the IIS 6.0 migration process. To facilitate the fastest possible migration, the following quick-start guide is provided. You can use this guide to help identify the steps of the IIS 6.0 migration process that you need additional information to complete, and then you can skip the information with which you are already familiar. In addition, all of the procedures that are required to complete the IIS migration process are documented in Appendix A: IIS Deployment Procedures.

**Prepare for Migration**

1. Identify which Web site and application components to migrate.

2. Determine compatibility with Windows Server 2003.

3. Determine application compatibility with worker process isolation mode by evaluating the following:

- The benefits of worker process isolation mode
- The application changes that are required for worker process isolation mode
- The management and provisioning script changes that are required for worker process isolation mode
- The results of lab tests that were completed to verify application compatibility with worker process isolation mode

4. Determine application compatibility with the Microsoft .NET Framework on Windows Server 2003.
5. Select one of two methods for migration:
    - Using the IIS 6.0 Migration Tool
    - Completing the migration process manually
6. If you are using the IIS 6.0 Migration Tool, identify the following:
    - Tasks that are automated by the migration tool
    - Subsequent tasks that must be performed manually

**Deploy the Target Sever**

1. Install Windows Server 2003.
2. Install and configure IIS 6.0.

**Migrate Web Sites with the IIS 6.0 Migration Tool**

1. Install the IIS 6.0 Migration Tool.
2. Verify that clients are not accessing Web sites.
3. Run the migration tool.
4. Verify that the migration tool ran successfully.
5. Migrate additional Web site content that is in the following two locations:
    - Outside the home directory of the Web site
    - Inside virtual directories
6. Modify IIS metabase properties that reference the systemroot folder.

**Migrate Web Sites Manually**

1. Verify that clients are not accessing the Web sites.
2. Create Web sites and virtual directories.

3. Migrate Web site content to the target server.
4. Configure Web site application isolation settings by completing the following tasks, if appropriate:
    - Document the current application isolation settings on the source server.
    - Configure application isolation settings in IIS 5.0 isolation mode.
    - Configure application isolation settings in worker process isolation mode.
5. Modify IIS 6.0 metabase properties that reference the systemroot folder.

**Configure IIS 6.0 Properties**

1. Configure IIS 6.0 properties that reference local user accounts.
2. Configure Web service extensions.
3. Configure Multipurpose Internet Mail Extensions (MIME) types.
4. Migrate server certificates for Secure Sockets Layer (SSL).
5. Migrate Microsoft FrontPage® users and roles.
6. Configure IIS 6.0 to host ASP.NET applications by completing the following tasks:
    - Configure IIS to use the correct version of the .NET Framework.
    - Configure the .NET Framework.
    - Review how ASP.NET applications run in each application isolation mode.
    - Migrate Machine.config attributes to their equivalent IIS 6.0 metabase property settings.
7. Determine whether to run the IIS Lockdown Tool and UrlScan.

**Perform Application-Specific Migration Tasks**

1. Modify application code for compatibility with Windows Server 2003 and IIS 6.0 by doing the following:
    - Modify references to Windows platform components and application programming interfaces (APIs) that are no longer supported in Windows Server 2003.
    - Modify references to IIS 6.0 metabase properties that have changed or are no longer supported in IIS 6.0.
    - Modify applications to be compatible with worker process isolation mode.
2. Install additional software required by applications.
3. Migrate Microsoft Transaction Server (MTS) packages, Component Object Model (COM) objects, and COM+ applications that are required by applications.

4. Modify data source names (DSNs) and Open Database Connectivity (ODBC) connections required by applications.

5. Create IP addresses that are used by applications.

6. Create users and groups that are used by applications.

7. Create registry entries required by applications on the target server.

**Complete the Migration**

1. Verify that the Web sites and applications migrated successfully.

2. Back up the target server.

3. Enable client access.

# Preparing for Migration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before migrating your existing IIS Web sites, ensure that the Web sites, applications, and their components are compatible with Windows Server 2003 and IIS 6.0. If you have setup programs or provisioning scripts that you are currently using on the source server and that you intend to continue using after migration, then ensure that the setup programs and provisioning scripts are compatible with Windows Server 2003 and IIS 6.0.

In addition, determine whether the existing Web sites and applications are compatible with worker process isolation mode in IIS 6.0. In cases where the existing Web sites and applications are not compatible with worker process isolation mode, you can still migrate to IIS 6.0 and Windows Server 2003 by configuring IIS to run in IIS 5.0 isolation mode. This allows the Web sites and applications to utilize other Windows Server 2003 and IIS 6.0 improvements

Also, if you are migrating ASP.NET applications, you must ensure that the ASP.NET applications are compatible with the latest version of the .NET Framework. If your ASP.NET applications require version 1.0 of the .NET Framework, you must configure IIS to run in IIS 5.0 isolation mode.

Lastly, you must determine whether you can perform the migration with the IIS 6.0 Migration Tool or if you need to perform the migration manually. You need to be aware of the additional migration steps that you must complete manually, regardless of the migration method that you choose. This allows you to have the appropriate tools and resources available when you are ready to perform the migration.

Figure 6.2 illustrates the process for preparing to migrate Web sites to IIS 6.0.

**Figure 6.2 Preparing for Migration of Web Sites to IIS 6.0**



# Identifying Which Web Site and Application Components to Migrate

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you begin the migration, identify the *components* that comprise the Web sites and applications. In addition to identifying the components, you must determine if there are any special circumstances associated with migrating these components. This Web site and application component migration is in addition to the Web site content and configuration that needs to also be migrated.

If you have setup, installation, or provisioning scripts for these Web sites and applications, you can use them to help you identify the components. If no setup, installation, or provisioning scripts exist, you must identify the Web site and application components manually.

Table 6.1 illustrates common application components that require a specific action when migrating to IIS 6.0.

**Table 6.1 Migration Issues Associated with IIS Web Site Components**

| Web Site Component | Description | Migration Issues |
|---|---|---|
| Microsoft Transaction Server (MTS) packages, COM objects, and COM+ applications | MTS, COM, and COM+ provide object-oriented access to business logic and other data in n-tier applications. | MTS packages and COM+ applications need to be re-created in COM+ on the target server. COM objects that are registered on the source server need to be reregistered on the target server.<br><br>In some cases, you might need to modify the MTS packages, COM objects, and COM+ applications, to make them compatible with Windows Server 2003 and IIS 6.0. |
| Registry entries | These entries are custom registry entries that are required by the application. | Applications might save custom configuration information in the Windows registry. The registry entries must be identified and then re-created on the target server. |
| Data source names (DSN) data connection strings | DSN data connection strings are used to provide connectivity to databases for IIS Web sites. | The DSNs might need to be modified on the target server for each of the Web sites and applications that access databases. |
| Open Database Connectivity (ODBC) data connections | ODBC data connections are used to provide connectivity to databases for IIS Web sites. | The ODBC settings might need to be configured on the target server for each of the Web sites and applications that use ODBC to access databases. |
| Internet Server API (ISAPI) applications | ISAPI applications are DLLs that are called by Web sites and applications. There are ISAPI applications (ISAPI extensions) that are shipped with IIS 6.0, such as the ASP.NET ISAPI extension. In addition, your Web sites and applications might rely on ISAPI applications written by your organization. | The ISAPI application must be installed on the target server and then enabled from the Web services extensions node in IIS Manager. In some cases, you might need to modify the ISAPI application to make it compatible with IIS 6.0 and Windows Server 2003. |
| SSL Server Certificates | SSL server certificates are installed on the server and they enable encrypted communications with clients | Certificates must be copied to, or re-created on, the target computer |
| Custom user and group accounts | User and group accounts that are created specifically for the Web sites and applications. | If the user and group accounts are local to the source server, the user and group accounts need to be created on the target server. |

| | | |
|---|---|---|
| Additional software | Additional software refers to installed commercial software and applications required by the applications on your server. | This additional software must be installed on the target server. This software might require modification to be compatible with Windows Server 2003 and IIS 6.0. |

# Determining Compatibility with Windows Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

At a minimum, your existing system hardware and software must be compatible with the Windows Server 2003 family before migrating the Web sites and applications to IIS 6.0. You must identify any software or hardware devices that are incompatible with Windows Server 2003.

**Compatibility of Existing Hardware**
When you select the computer that will be the target server, ensure that the computer is compatible with Windows Server 2003. The most common hardware incompatibility is a device driver that is no longer supported or is not yet supported in Windows Server 2003. When a device is no longer supported, remove the existing device and then install an equivalent device that is supported by Windows Server 2003. It is also important that you have a current system BIOS from your computer manufacturer.

For example, the target server might have a network adapter that is not included with Windows Server 2003. You can review the manufacturers Web site to obtain a driver that is compatible with Windows Server 2003.

For more information about the hardware devices supported by the Windows Server2003 operating systems, see the Hardware Compatibility List or the Windows Server Catalog.

**Compatibility of Existing Software**
Before migrating, you need to consider the compatibility of your existing applications, or other software that is installed on your server, with Windows Server 2003. This includes software and tools from manufacturers other than Microsoft, as well as Microsoft server products that do not ship with the Windows operating system. Make sure that you have the latest versions of all pre-existing software or service packs that are compatible with Windows Server 2003.

The most common software incompatibilities are caused when your application depends on software from another manufacturer that does not support Windows Server 2003. Or, you might have applications that were designed to run on Windows NT Server 4.0 or Windows 2000 Server operating systems, which reference APIs that have been changed or removed in Windows Server 2003.

To help you determine the compatibility of your existing software with Windows Server 2003, use the Windows Application Compatibility Toolkit before migrating your Web sites and applications to the target server.

# Determining Application Compatibility with IIS 6.0 Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 can run in one of two distinct modes of operation, which are called application isolation modes. *Application isolation* is the separation of applications by process boundaries that prevent the applications from affecting one another, and it is configured differently for each of the two IIS application isolation modes: IIS 5.0 isolation mode and worker process isolation mode.

*Worker process isolation mode* uses the redesigned architecture for IIS 6.0. This application isolation mode runs all application code in an isolated environment. However, unlike earlier versions of IIS, IIS 6.0 provides isolation without a performance penalty because fewer processor instructions are run when switching from one application pool to another. Worker process isolation mode is compatible with most existing Web sites and applications. Whenever possible, run IIS 6.0 in worker process isolation mode benefit from the enhanced performance and security in IIS 6.0.

*IIS 5.0 isolation mode* provides compatibility for applications that depend upon the process behavior and memory model of IIS 5.0. Run IIS in this mode only when a Web site or application cannot run in worker process isolation mode, and run it only until the compatibility issues are resolved.

| **Important** |
| --- |
| IIS 6.0 cannot run both application isolation modes simultaneously on the same server. Therefore, on a single server running IIS 6.0, you cannot run some Web applications in worker process isolation mode and others in IIS 5.0 isolation mode. If you have applications that require separate modes, you must run them on separate servers. |

During the migration process, you install Windows Server 2003 and IIS 6.0 on the target server. IIS 6.0 is configured to run in worker process isolation mode by default. Before you begin migrating your production Web sites and applications, evaluate whether your Web sites and applications are compatible with worker process isolation mode. Most of the compatibility issues with IIS 6.0 occur when configuring IIS 6.0 to run in worker process isolation mode.

One of the most common reasons for incompatibility with worker process isolation mode is that applications do not recognize custom Internet Server API (ISAPI) extensions or dynamic-link libraries (DLLs) that depend on the memory and request processing models used by earlier versions of IIS. Determine application compatibility in your lab before migrating your existing IIS Web sites and application, and if you determine that your applications are not compatible with worker process isolation mode, you can run the applications in IIS 5.0 isolation mode.

| **Note** |
| --- |
| Identifying a complete list of potential incompatibilities that applications can experience with worker process isolation mode is beyond the scope of this section. Even after following the guidelines in this section, you need to verify in a lab whether your Web sites and applications are compatible with worker process isolation mode. |

Determine the compatibility of an application with worker process isolation mode by completing the following steps:

1. Evaluate the benefits of moving to worker process isolation mode.

2. Evaluate the application changes that are required so that the applications can run in worker process isolation mode.

3. Evaluate the management and provisioning script changes that are required to set up programs and provisioning scripts in worker process isolation mode.

4. Verify the compatibility of the application with worker process isolation mode in a lab.

# Evaluating the Benefits of IIS 6.0 Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Worker process isolation mode provides higher levels of security and availability for Web sites and applications than IIS 5.0 isolation mode. Therefore, it is recommended that you configure IIS 6.0 to run in worker process isolation mode.

Worker process isolation mode provides the following improvements to IIS 6.0.

**Security Enhancements**
IIS 6.0 includes a variety of security features and technologies that help ensure the integrity of your Web site content, and of the data that is transmitted through your sites. The following security enhancement is only available when IIS 6.0 is running in worker process isolation mode.

**Default process identity for Web sites and applications is set to NetworkService**
In IIS 5.0 isolation mode, the default process identity is LocalSystem, which enables access to, and the ability to alter, nearly all of the resources on the Web server. The potential of attacks is reduced in worker process isolation mode because Web sites and applications run under the NetworkService identity. The NetworkService identity is granted less privileges, which helps prevent an attack from compromising the Web server, which is possible with the LocalSystem identity.

**Performance and Scaling Enhancements**
Future growth in the utilization of your Web sites and applications requires increased performance and scalability of Web servers. By increasing the speed at which Hypertext Transfer Protocol (HTTP) requests can be processed and by allowing more applications and sites to run on one Web server, the number of Web servers that you need to host a site is reduced. The following are a few of the performance improvements included in worker process isolation mode.

**Support for processor affinity for worker processes in an application pool**
You can configure all of the worker processes in an application pool to have affinity with specific processors in a multiprocessor or server. Processor affinity allows the worker processes to take advantage of more frequent processor caching (Level 1 or Level 2).

**Elimination of inactive worker processes and reclamation of unused resources**
You can configure application pools to have worker processes request a shutdown if they are idle for a certain amount of time. This can free unused resources for other active worker processes. New worker processes are then started only when they are needed.

**Distributing client connections across multiple worker processes**
You can configure an application pool to have more than one worker process servicing client connections, also known as a *Web garden*. Because there are multiple worker processes, the incoming client connections are distributed across the worker processes and throughput is not constrained by a single worker process.

**Ability to Isolate Web sites and applications from each other**
You can isolate Web sites and applications without incurring a performance penalty. This is because the Web site and applications, and their associated ISAPI filters, run in the same process.

**Availability Enhancements**
Because worker process boundaries isolate the applications in an application pool from the applications in other application pools, if an application fails, it does not affect the availability of other applications running on the server. Deploying applications in application pools is a primary advantage of running IIS 6.0 in worker process isolation mode.

**Reduced number of Web server restarts required when administering Web sites and applications**
Many of the common operation tasks do not force the restart of the server or the Web service restart. These tasks, such as upgrading site content or components, debugging Web applications, or dealing with faulty Web applications, can be performed without affecting service to other sites or applications on the server.

**A fault-tolerant request processing model for Web sites and applications**
In IIS 5.0 isolation mode, each Web site or application has only one worker process. However, in worker process isolation mode, you can create a *Web garden* by configuring a number of worker processes to share the processing. The benefit of a Web garden is that if one worker process stops responding, other worker processes are available to accept and process requests.

**Isolation of failed worker processes from healthy worker processes**
In worker process isolation mode, IIS can determine that a worker process is failing and start a new worker process to replace the failing worker process. Because a new worker process is created before the old worker process terminates, users requesting the Web site or application experience no interruption of service. After IIS creates the new worker process, the failed worker process can be separated, or *orphaned*, from the application pool. The advantage of orphaning a worker process rather than terminating it is that debugging can be performed on the orphaned worker process.

**Health monitoring of Web sites and applications**
In worker process isolation mode, you can configure an application pool to monitor not only the health of the entire application pool, but also individual worker processes servicing the application pool. Monitoring the health of a worker process allows IIS to detect that a worker process is unable to serve requests and to take corrective action, such as recycling the failed worker process.

In addition, worker process isolation supports other responses when a failed worker process or application pool is detected. For example, IIS can attach a debugger to an orphaned worker process or notify an administrator that an application pool has failed due to rapid-fail protection.

**Prevention of Web sites or applications that fail quickly from consuming system resources**
In some cases, availability can be affected by Web sites and applications that fail very quickly, are automatically restarted, and then fail quickly again. The endless cycle of failure and restarting can consume system resources, causing other Web sites and applications to experience denial of services because of system resource shortages.

Worker process isolation mode includes *rapid-fail protection* that stops an application pool when too many of the worker processes assigned to an application pool are found to be unhealthy within a specified period of time.

**Automatic restart of poorly performing Web sites and applications**
Some Web sites and applications have memory leaks, are poorly coded, or have other unidentified problems. In IIS 5.0 isolation mode, these applications can force you to restart the entire Web server. The recycling feature in worker process isolation mode can periodically restart the worker processes in an application pool in order to manage faulty applications. Worker processes can be scheduled to restart based on several options, such as elapsed time or the number of requests served.

# Evaluating Application Changes Required for IIS 6.0 Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In most cases, the existing Web sites and applications can be hosted in worker process isolation mode without modification. However, the following are known application issues that can create incompatibilities with worker process isolation mode and require you to run IIS 6.0 in IIS 5.0 isolation mode.

- **Requires Inetinfo.exe.** When the application must run in the same process with Inetinfo.exe, IIS must be configured to run in IIS 5.0 isolation mode because Inetinfo.exe does not process Web requests in worker process isolation mode. In worker process isolation mode, W3wp.exe processes Web requests.

- **Requires Dllhost.exe.** When the application depends on Dllhost.exe to process Web requests for the application, IIS must be configured to run in IIS 5.0 isolation mode because Dllhost.exe is not available in worker process isolation mode.

- **Requires read raw data filters.** If the application uses an ISAPI raw data filter, IIS must be configured to run in IIS 5.0 isolation mode.

- **Requires version 1.0 of the .NET Framework.** When the application requires version 1.0 of the .NET Framework, IIS must be configured to run in IIS 5.0 isolation mode because version 1.0 of the .NET Framework is only supported in IIS 5.0 isolation mode.

- **Requires ISAPI filters or extensions that are incompatible with worker process isolation mode.** When the application requires ISAPI filters or extensions that are incompatible with worker process isolation mode, IIS must be configured to run in IIS 5.0 isolation mode. ISAPI filters or extensions might be incompatible if the filter or extension has one of the following characteristics:

    - Runs in multiple instances and expects to be recycled by using the recycling provided in IIS 5.0.

    - Expects to have exclusive lock on a resource, such as a log file.

If any of the Web sites and applications running on the existing Web server has one or more of these characteristics, then do one of the following:

- Configure IIS 6.0 to run in IIS 5.0 isolation mode to ensure Web site and application compatibility. If the Web sites and applications need to be hosted in IIS 5.0 isolation mode, then the migration process is complete.

- Modify the existing applications to remove the dependencies.

# Evaluating Management and Provisioning Script Changes Required for IIS 6.0 Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When management or provisioning scripts exist for your Web sites and applications, you might need to modify them so that they properly set up the Web sites and applications for the application isolation mode that is running on the Web server — IIS 5.0 isolation mode or worker process isolation mode. If you do not modify your management and provisioning scripts as required, you will be unable to use them to install and configure your Web sites and applications on IIS 6.0.

For example, you might have a provisioning script that uses Microsoft Active Directory® Service Interfaces (ADSI) to create a Web site and configure the site for High isolation in IIS 5.0 isolation mode, which does not exist in worker process isolation mode. After migration and when the Web server is running in worker process isolation mode, you need to modify the script to create application pools instead.

Common modifications that might be necessary to your setup programs and provisioning scripts include:

- When the script installs an ISAPI extension or filter, you might need to modify the script to add an entry for the ISAPI extension or filter to the Web service extensions restriction list and set the status of the entry to **Allowed**.

  For more information about modifying your setup programs or provisioning scripts to install and enable an ISAPI extension or filter, see Configuring Web Service Extensions.

- When the script installs an application that contains dynamic content, you might need to modify the script to set the status of the appropriate Web service extensions to **Allowed**, so that IIS allows the dynamic content to run.

  For more information about modifying your setup programs or provisioning scripts to enable dynamic content, see Configuring Web Service Extensions.

- When the script installs a Web site or application that runs in High isolation in IIS 5.0, you might need to modify the script to create an application pool and configure the application pool with settings that are comparable to the original IIS 5.0 isolation settings.

  For more information about modifying your setup programs or provisioning scripts to create and configure application pools, see Configuring Application Isolation Settings in Worker Process Isolation Mode.

- When the script references metabase properties, you might need to modify the script if it references metabase properties that have changed or are no longer supported in IIS 6.0.

  For more information about modifying your setup programs or provisioning scripts to reference the proper metabase properties in IIS 6.0, see Modifying References to IIS 6.0 Metabase Properties.

# Verifying Application Compatibility with IIS 6.0 Worker Process Isolation Mode in a Lab

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After modifying your Web sites, applications, setup programs, and provisioning scripts to be compatible with worker process isolation mode, you need to test your modifications in a lab. Be sure to test for compatibility before performing the migration process on a production Web server.

Verify the compatibility of your Web sites, applications, setup programs, and provisioning scripts with worker process isolation mode in a lab by completing the following steps:

1. Make an image backup of the source server.

2. Restore the backup to a Web server in your lab, referred to hereafter as a *test source server*. Ensure that the test source server is not connected to your production network, to prevent any problems encountered during the migration from affecting your production network.

3. Perform a migration from the test source server to the test target server.

4. Configure the test target server to run in worker processor isolation mode or IIS 5.0 isolation mode.

5. Make the necessary modifications to the Web sites, applications, setup programs, and provisioning scripts so that they are compatible with worker process isolation mode.

6. Verify that the Web sites, applications, setup programs, and provisioning scripts run correctly on the test target server.

For more information about setting up a test lab, see "Designing a Test Environment" in Planning, Testing, and Piloting Deployment Projects.

# Determining Application Compatibility with .NET Framework

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For a successful migration to Windows Server 2003 and IIS 6.0, you need to determine whether your ASP.NET applications are dependent on specific versions of the .NET framework. Windows Server 2003 ships with version 1.1 of the .NET Framework, while most of the .NET applications that were developed before the release of Windows Server 2003 were designed to run on version 1.0 of the .NET framework.

| **Note** |
| --- |
| Version 1.0 of the .NET Framework is only supported in IIS 5.0 isolation mode. Therefore, you can only run version 1.0 and version 1.1 of the .NET Framework on the same server when IIS 6.0 is configured to run in IIS 5.0 isolation mode. Version 1.1 of the .NET Framework is supported in IIS 5.0 isolation mode or worker process isolation mode. |

Typically, ASP.NET applications running on version 1.0 of the .NET Framework are compatible with version 1.1 of the .NET Framework. However, there might be some incompatibilities, the majority of which are security related and can be corrected by configuring the .NET Framework to be less restrictive.

For example, in the .NET Framework version 1.0, the .NET Framework only examines the **SQLPermission.AllowBlankPassword** attribute if the user actually includes the password keyword in their connection string. If an administrator or user sets the **SQLPermision.AllowBlankPassword** attribute to **False**, it is possible to specify a connection string like "server=(localhost);uid=sam" and succeed. In the .NET Framework version 1.1, this connection string fails.

You can use both version 1.0 and version 1.1 of the .NET Framework on the same server running IIS 6.0, which is also known as *side-by-side configuration*. Side-by-side configuration allows you to run a mixture of applications that require version 1.0 or version 1.1 of the .NET Framework. During your lab testing, determine the version of the .NET Framework that is required by your applications.

For a current list of possible compatibility issues when upgrading from version 1.0 to version 1.1 of the .NET Framework, see Compatibility Considerations and Version Changes on the GotDotNet Web site. For more information about using multiple versions of the .NET Framework in side-by-side configuration, see Configuring IIS to Use the Correct Version of the .NET Framework.

| **Important** |
| --- |
| ASP.NET is not available on the following operating systems: Microsoft® Windows® XP 64-Bit Edition; the 64-bit version of Windows® Server 2003, Enterprise Edition; and the 64-bit version of Windows® Server 2003, Datacenter Edition. For more information, see "Features unavailable on 64-bit versions of the Windows Server 2003 family" in Help and Support Center for Windows Server 2003. |

# Selecting a Migration Method

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before migrating your Web sites and applications to IIS 6.0, you need to determine whether to perform the migration manually or with the IIS 6.0 Migration Tool. It is recommended that you use the migration tool to begin the migration process, except when one of the following is true:

- **You have set up programs, installation scripts, or provisioning scripts for the Web sites and applications that you are migrating.** When the Web sites and applications that you are migrating have setup programs, installation scripts, or provisioning scripts, use those programs or scripts to install the Web sites and applications on the target server. Ensure that the setup programs, installation scripts, and provisioning scripts have been properly modified to install the Web sites and applications on IIS 6.0.

- **The target server is configured to run in IIS 5.0 isolation mode.** When the Web sites and applications that you are migrating require the target server to run in IIS 5.0 isolation mode, you must perform the migration manually. To determine whether your Web sites and applications are compatible with worker process isolation mode, see Determining Application Compatibility with Worker Process Isolation Mode earlier in this section.

- **The source server has a significant number of FrontPage extended Web sites.** When the FrontPage extended Web sites make extensive use of the administrative and publishing security-related settings found in FrontPage 2000 Server Extensions, perform the migration manually. To ensure that these security-related settings are migrated properly to the target server, perform the migration manually and use FrontPage publishing to transfer the Web site to the target server.

- **You want to migrate individual virtual directories.** When you want to migrate individual virtual directories, perform the migration manually. The IIS 6.0 Migration Tool only moves Web site content and configuration settings at the Web site level, which means that all of the virtual directories beneath the Web site are migrated to the target server.

If you determine that you need to perform the migration manually, then you can proceed to the next step in the migration process — Deploying the Target Server. Otherwise, you need to know which steps in the migration process are automated by the IIS 6.0 Migration Tool, and which steps you must complete manually after running the migration Tool. Knowing the role of the IIS 6.0 Migration Tool in the migration process enables you to have the appropriate tools and resources available when you are ready to begin the migration.

# Identifying the Role of the IIS 6.0 Migration Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 Migration Tool is a command-line utility that automates some of the steps in the process for migrating Web sites and applications hosted on IIS 4.0 or IIS 5.0 to IIS 6.0. The migration tool does not provide an end-to-end migration solution, but it automates some of the time-consuming, repetitive migration tasks. Although the migration tool is conceptually similar to the IIS 5.0 Migration Tool, the IIS 6.0 Migration Tool is a completely new tool designed specifically for Windows Server2003 and is not compatible with the previous tool.

# Migration Tasks That Are Automated by the IIS 6.0 Migration Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 Migration Tool automates the following steps in the IIS migration process:

- Transferring the Web site content. All of the files and folders located in the home directory of the Web site and virtual directories (which can be located outside of the home directory of the Web site) are copied from the source server to the target server. The NTFS file system permissions assigned to the files and directories that make up the Web site content on the source server are granted to the corresponding files and directories on the target server. Any content referenced by the Web site or application that is not located in subdirectories of the home directory of the Web site or in virtual directories is not migrated.

- Transferring the Web site configuration in the IIS metabase. The configuration for each Web site and application, which is stored in the IIS metabase properties on the source server, is translated and then the corresponding IIS 6.0 metabase properties are appropriately configured on the target server.

- Translating the application isolation configuration. The target server is running in worker process isolation mode. The application isolation configuration for each Web site and application on the source server is translated into application pool configuration settings on the target server.

- Backing up the IIS metabase configuration to the target server. The IIS metabase configuration is backed up on the target server before migration. You can use this backup to restore the target server to a known state in the event that the migration process is not successful.

# Migration Tasks That Must Be Completed Manually

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following steps in the Web site migration process must be completed after running the IIS 6.0 Migration Tool.

**Migrating Additional Web Site and Application Content**
The IIS 6.0 Migration Tool migrates all of the content that is located in the home directory of the Web site and in any subdirectories contained in that home directory. You can migrate any Web site and application content that is not located in these directories by completing the following steps:

- **Migrate content located outside the home directory and subdirectories of the Web site.** When the Web sites and applications reference content that is located in folders outside of the home directory of the Web site or the virtual directories beneath the home directory, you must migrate this content manually.

- **Migrate content located in a virtual directory.** When the virtual directory content on the source server is stored on a disk volume, such as F:, that does not exist on the target server, you must migrate this content manually.

**Configuring Additional Web Site and Application Properties**
After running the IIS 6.0 Migration Tool, the Web sites are configured comparably to how they were configured on the source server. However, depending on the configuration of the Web sites and applications on the source server, you might need to configure additional Web site and application properties, by completing the following steps:

- **Change the IIS metabase settings to reflect where Windows is installed.** If the Windows Server 2003 systemroot path does not match the Windows systemroot path on the source server, you must modify the metabase settings on the migrated Web sites to reference the correct folder on the target server. For example, if Windows was installed on C:\WINNT on the source server, the IIS metabase entries for **ScriptMaps**, and **HTTPErrors** properties might still reference these paths, and therefore need to be updated on the target server.

- **Configure IIS properties that reference local user accounts.** There are a number of Web site configuration properties on the source server that you can configure to utilize user accounts that are local to the source server. Local user and group accounts are not migrated from the source server to the target server by the IIS 6.0 Migration Tool. As a result, the migrated Web sites reference user accounts that do not exist on the target server. In these cases, you must configure the Web sites to utilize user accounts that are domain-based or local to the target server, and then re-create the file system permissions on migrated content by completing the following steps:

    - **Configure local user NTFS permissions on content.** If NTFS permissions are granted to local user accounts on the source server, you must create new user accounts, or designate existing user accounts, for use on the target server and then grant the corresponding NTFS permissions to the user accounts on the target server.

    - **Configure anonymous account properties for Web sites and virtual directories.** If a Web site or virtual directory is configured to use a user account on the source server for anonymous access (other than the default IUSR_computername account), you must create a new user account, or designate an existing user account, for use on the target server.

    - **Configure IIS 4.0 and IIS 5.0 application isolation identities.** If the Web sites or applications are isolated and are configured to use a local user account on the source server as the application isolation identity, you must create a new user account, or designate an existing user account, for use on the target server. Then you must configure the corresponding application pool, on the target server, to use the newly created account as the identity of the application pool.

    - **Add Web service extensions for dynamic content used by the Web sites.** Any dynamic content types, including ISAPI extensions, ISAPI filters, or CGI applications, which are not automatically migrated by the IIS 6.0 Migration Tool need to be added to the Web service extensions list.

    - **Add MIME types for static content used by the Web sites.** MIME types define the types of static files that are

served by the Web server. You must identify the MIME types defined on the source Web server and then create the same associations of MIME types to file name extensions on the target Web server.

- **Configure SSL certificates.** You must export server certificates for Secure Sockets Layer (SSL)-enabled Web sites from the source server, and then install the certificates on the target server after the migration process is complete.

- **Configure FrontPage Server Extensions users and roles.** If FrontPage Server Extensions is configured to use a local user account on the source server as the FrontPage administrator, you must create a new user account, or designate an existing user account, for use on the target server. In addition, you must assign the user the same FrontPage role as the corresponding user had on the source server.

- **Configure IIS for ASP.NET applications.** If you migrate ASP.NET applications from the source server, you might need to migrate the attribute settings in the Machine.config file that are used to set process-model behavior. When the target server is configured for worker process isolation mode, the attribute settings in the Machine.config file need to be converted to corresponding application pool settings. This is because ASP.NET uses the IIS process model when IIS is running in worker process isolation mode. For more information, see Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings later in this section.

**Performing Application-Specific Migration Tasks**

In addition to the IIS 6.0 configuration changes that are required, you might need to complete any migration tasks that are specific to the applications running on the source server. The application-specific steps that you might need to perform include the following:

- **Modifying application code for compatibility with Windows Server 2003 and IIS 6.0.** When the applications on the source server are not compatible with Windows Server 2003 and IIS 6.0, you need to modify the applications. Most of these modifications are required when applications use application programming interfaces (APIs) no longer supported by Windows Server 2003 and IIS 6.0.

- **Installing additional software required by the applications.** When applications on the source server require additional software that was developed by your organization, by Microsoft, or by other organizations, you need to install that software on the target server. This software can include filters and ISAPI extensions. You must obtain a version of the software that is compatible with Windows Server 2003 and IIS 6.0.

- **Migrating MTS packages, COM objects, and COM+ applications.** When your applications include MTS packages, COM objects, or COM+ applications, you must migrate them to the target server. For MTS packages, you need to rewrite the MTS as a COM+ application.

- **Creating IP addresses that are used to uniquely identify the applications.** Web sites and applications are uniquely identified by a unique IP address, a unique combination of an IP address and a TCP port, or host headers. When Web sites and applications on the source server are uniquely identified by IP addresses, you must create corresponding IP addresses on the target server and then configure the migrated Web sites and applications to use the new IP addresses.

- **Creating users and groups that are used by the applications.** When users that access the applications on the source server have accounts that are local to the source server, you need to create new accounts on the source server and then assign the appropriate NTFS permissions on the target server to the new accounts. When the user accounts are in Active Directory, you only need to assign the appropriate NTFS permissions on the target server to the accounts in Active Directory.

- **Creating registry entries for the applications.** When your applications store configuration information in the registry, you might need to create the same registry entries for the applications on the target server.

**Note**

The IIS 6.0 Migration Tool does not migrate the IIS logs from the source server to the target server. The migration of the IIS logs is not necessary for proper operation of the Web sites and applications on the target server. However, you might want to archive the IIS logs before decommissioning the source server for historical reference to events that occurred on the source server before migration.

For information about using the IIS 6.0 Migration Tool to perform your migration, see Migrating Web Sites with the IIS 6.0 Migration Tool later in this section.

# Deploying the Target Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You must install Windows Server 2003 on the target Web server before you can migrate the Web sites and applications. In addition to installing Windows Server 2003, you must install and configure IIS 6.0 on the target server. To follow the process described in this section, install Windows Server 2003 and IIS 6.0 with the default options.

You must also configure the target server to run in worker process isolation mode or in IIS 5.0 isolation mode. Run IIS in IIS 5.0 isolation mode only when a Web site or application cannot run in worker process isolation mode. If there are no incompatibilities, configure IIS to run in worker process isolation mode to utilize the benefits of IIS 6.0 features.

Figure 6.3 illustrates the process for deploying the target server.

**Figure 6.4 Deploying the Target Server**




If you have already installed Windows Server 2003 on the target server, and you have already installed and configured IIS 6.0 on the target server, you can proceed to Migrating Web Sites with the IIS 6.0 Migration Tool or Migrating Web Sites Manually.

# Installing Server 2003

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The primary concern when installing Windows Server 2003 is to ensure that the security of the target server is maintained. When you install Windows Server 2003 as a dedicated Web server, the default components and services are configured to provide the lowest possible attack surface. You can further reduce the attack surface of the target server by enabling only the essential Windows Server 2003 components and services.

The migration process presented in this section assumes that you install Windows Server 2003 with the default options. If you use other methods for installing and configuring Windows Server 2003, such as unattended setup, your configuration settings might be different.

| **Note** |
| --- |
| When you complete the installation of Windows Server 2003, Manage Your Server automatically starts. The migration process assumes that you quit Manage Your Server and then further configure the Web server in **Add or Remove Programs**in Control Panel. |

For more information about enabling the essential Windows Server 2003 components and services required by your server, see Enabling Only Essential Windows Server 2003 Components and Services.

# Installing and Configuring IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because IIS 6.0 is not installed during the default installation of Windows Server 2003, the next step in deploying the target server is to install and configure IIS 6.0. The migration process presented here assumes that you install IIS 6.0 with the default options in **Add or Remove Programs** in Control Panel. If you use other methods for installing and configuring Windows Server 2003, such as Manage Your Server, the default configuration settings might be different.

Install and configure IIS 6.0 by completing the following steps:

1. Install IIS 6.0 with only the essential components and services.

   As with installing Windows Server 2003, the primary concern when installing and configuring IIS 6.0 is to ensure that the security of the target server is maintained. Enabling unnecessary components and services increases the attack surface of the target server. You can help ensure the target server is secure by enabling only the essential components and services in IIS 6.0.

   For more information about how to install IIS 6.0, see Install IIS 6.0. For more information about the IIS 6.0 protocols and services, see Enabling Only Essential IIS Components and Services.

2. If the source server has Web sites with FrontPage Server Extensions, install FrontPage 2002 Server Extensions from Microsoft on the target server.

   For more information about how to enable FrontPage Server Extensions, see Configure Web Service Extensions.

3. Configure IIS 6.0 to run in IIS 5.0 isolation mode.

   If you determined that one or more of your Web sites or applications are incompatible with worker process isolation mode (the target server is currently configured to run in worker process isolation mode), configure IIS 6.0 to run in IIS 5.0 isolation mode in IIS Manager, or by setting the IIS metabase property **IIs5IsolationModeEnabled** to a value of **True**.

   | **Note** |
   | --- |
   | If you configure IIS 6.0 to run in IIS 5.0 isolation mode and then decide to change the configuration back to worker process isolation mode, the original worker process isolation mode settings are retained. Similarly, if you configure IIS 6.0 to run in IIS 5.0 isolation mode, change to worker process isolation mode, and then change back to IIS 5.0 isolation mode, the IIS 5.0 isolation mode settings are retained. |

   For more information about how to configure IIS 6.0 to run in worker process isolation mode or in IIS 5.0 isolation mode, see Configure Application Isolation Modes. For more information about determining compatibility with worker process isolation mode, see Evaluating Application Changes Required for Worker Process Isolation Mode.

# Migrating Web Sites with the IIS 6.0 Migration Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Earlier in the process, you determined whether to migrate your Web sites by using the IIS 6.0 Migration Tool or to migrate them manually. If you decided to use the IIS 6.0 Migration Tool to complete the migration process, perform the steps that are described in this section.

If you decided to complete the migration process manually, you can skip this step and proceed directly to Migrating Web Sites Manually later in this section. For more information about how to determine whether to perform the migration with the IIS 6.0 Migration Tool, see Selecting a Migration Method earlier in this section.

Figure 6.5 illustrates the process of performing the migration to IIS 6.0 with the IIS 6.0 Migration Tool.

**Figure 6.5 Migrating Web Sites to IIS 6.0 with the IIS 6.0 Migration Tool**



# Installing the IIS 6.0 Migration Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 Migration Tool and its accompanying user documentation are included in the Internet Information Services (IIS) 6.0 Resource Kit Tools,which can be installed by running IIS60rkt.exe from the Microsoft Download Center. There is no setup program for the migration tool; you need only copy the IIS 6.0 Migration Tool to the target server. You do not need to install any software on the source server. To avoid installing the entire contents of the IIS 6.0 Resource Kit Tools onto your target server, install the tools onto a workstation first, and then copy the migration tool to the target server.

Install the IIS 6.0 Migration Tool by completing the following steps:

1. Install the Internet Information Services (IIS) 6.0 Resource Kit Tools by running IIS60rkt.exe from the Microsoft Download Center.

2. Create a folder on the target server to contain the migration tool.

3. Copy the IIS 6.0 Migration Tool (Iismt.exe) from Program Files\IIS Resources\IIS 6.0 Migration Tool, on the computer listed in Step 1, to the folder on the target server that you created in Step 2.

4. Verify that the migration tool operates correctly by typing **iismt** at the command prompt.

   This starts the migration tool and displays Help for the tool only. The migration process does not actually begin.

# Verifying That Clients Are Not Gaining Access to Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 Migration Tool requires only read access to the Web site content and configuration settings on the source server. Therefore, the source server can remain online in your production environment. However, you might need to remove the source server from your production network and move it to a private network segment that has direct network connectivity to the target server under the following circumstances:

- The number of files and amount of configuration information being copied across the network generates a high volume of traffic and slows the production network.

- Firewalls that exist between the source and target servers prevent the migration tool from performing the migration. This often occurs because the migration tool uses DCOM ports that are blocked by the firewalls for communicating with the source and target servers.

- Security-related configuration settings on the source server need to be modified to allow the migration tool to work.

  Examples of these security-related configuration settings that can prevent the migration tool from working include the following:

  - Remote access to disk volumes through administrative shares is prohibited.

    The IIS 6.0 Migration Tool requires access to the disk volume that contains the Web site content to perform the migration. For example, if the Web site content is stored in D:\Inetpub\Wwwroot, the migration tool must access the administrative share (D$) of the disk volume. The administrative shares are often removed to help prevent unauthorized access to the Web server. In order to use the migration tool, you must re-create the appropriate administrative shares.

  - Remote access to the source server must be allowed for members of the local Administrators group on the source server.

    For security reasons, many organizations restrict the members of the local Administrators group so that they can only log on locally, not over the network. However, the migration tool must be able to remotely access the source server over the network, as a member of the local Administrators group.

Verify that clients are no longer accessing Web sites by completing the following steps:

1. Prevent new clients from accessing the sites by pausing the sites.

   For more information about how to pause Web sites, see Pause Web or FTP Sites.

2. Monitor the active Web connections to determine when clients are no longer accessing the source server.

   For more information about how to monitor the active Web connections, see Monitor Active Web and FTP Connections.

3. When the number of active Web counters is zero, stop the WWW service or move the source server to another network segment.

   If you elect to stop the WWW service, ensure that the IIS Admin service is running because the migration tool requires the IIS Admin service. For more information about how to stop the WWW service, see Stop the WWW Service.

# Running the IIS 6.0 Migration Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 Migration tool is a command-line utility that is designed to run on the target server.

Before you run the IIS 6.0 Migration tool, you should test the Remote Procedure Call (RPC) connection. For information about testing RPC connections, see Rpcping.exe: RPC Ping.

To run the migration tool, type **iismt** at the command prompt, and provide the parameters, as listed in Table 2, that are appropriate to the Web sites and applications that you are migrating. For more information about each of the parameters and on how to perform the migration with the IIS 6.0 Migration Tool, see the "IIS 6.0 Migration Tool User Guide" in the IIS 6.0 Resource Kit Tools.

The IIS 6.0 Migration Tool uses the following syntax:

**iismt.exe** *Server Website*[**/user** *Username*] [**/password** *Password*] [**/path** *path*] [**/serverbindings** *ServerBindings String*] [**/siteid** *SiteID* |**Replace**][**/configonly**][**/fpse**][**/verbose**][**/overwrite**][**/noninteractive**]

Table 6.2 lists all the command-line parameters that the migration tool accepts, although not all of the parameters listed are required.

**Table 6.2 Command-Line Parameters Accepted by the IIS 6.0 Migration Tool**

| Parameter | Required or Optional | Description |
|---|---|---|
| *SourceServer* | Required | Identifies the source server by providing the following:<br>• DNS or NetBIOS name for the source server<br>• IP address of the source server |
| *WebSite* | Required | Identifies the site to be migrated by providing the following:<br>• Web site description, such as "Default Web Site"<br>• Metabase key path, such as W3SVC/1 |
| **/user** *UserName* | Optional | Specifies the user name of an account that is a member of the Administrators group on the source server. This parameter is not necessary if you log on with an account that is a member of the Administrators group on both the source server and the target server. |
| **/password** *Password* | Optional | Specifies the password that is associated with the user name. |
| **/path** *Path* | Optional | Specifies a different directory location for the home directory of the Web site on the target server. This parameter is ignored if **/configonly** is included. |
| **/serverbindings** *ServerBindingsString* | Optional | Allows a change to the IP address, host header, or port configuration of the Web site during the migration. |

| | | |
|---|---|---|
| **/siteid** *SiteID*\| **Replace** | Optional | Specifies the site ID on the target server, which can be specified as one of the following:<br>● *SiteID* - Overwrites the site ID on the target server.<br>● Replace - Overwrites the site ID on the target server with the site ID from the source server. |
| **/configonly** | Optional | Migrates only the Web site configuration and not the Web site content. |
| **/fpse** | Optional | Reapplies FrontPage Server Extensions to the migrated site on the target server. This parameter is ignored if **/configonly** is included. |
| **/verbose** | Optional | Displays metabase path copy and file copy operations to the screen during the migration process. |
| **/overwrite** | Optional | Does not display messages that prompt the user to confirm the overwriting of an existing destination folder or file when content is being copied from the source server to the target server. |
| **/noninteractive** | Optional | Does not display messages that prompt the user for input. The migration tool will exit on the first error condition. This is a useful switch for invoking the migration tool from a batch file or script program to perform an unattended migration. |

# Verifying That the IIS 6.0 Migration Tool Ran Successfully

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before continuing with the Web site migration, verify that the IIS 6.0 Migration Tool migrated the Web site content and configuration information successfully. When you run the migration tool, the output displayed by the migration tool indicates the success or failure of the migration. If the output indicates that errors occurred, you can use the IIS 6.0 Migration Tool log file to resolve any errors.

Verify that the IIS 6.0 Migration Tool ran successfully by completing the following steps:

1. Open *%systemroot%*\System32\LogFiles\IISMT\iismt_*date_time*.log in a text editor and determine if any errors occurred (where *date* is the date when the tool ran and *time* is the time the tool started).

2. Review the log file and resolve any problems that occurred during migration before proceeding to the next step in the process.

# Migrating Additional Web Site Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Some Web sites and applications have content that is not located in the home directory of the Web site or in subdirectories that are inside the home directory. The IIS 6.0 Migration Tool only migrates Web site content from the following locations:

- Within the home directory and subdirectories of the Web site.
- In virtual directories whose disk volume letter exists on both the source server and the target server.

If the code in your applications directly references content that is located outside the home directory and subdirectories of the Web site, or if a virtual directory is stored on a disk volume letter that does not exist on the target server, you must migrate this Web site content manually.

# Migrating Content Located Outside the Home Directory of the Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A Web site or application can have content that is referenced by the Web site, but is located outside the home directory and subdirectories of the Web site. This content must be migrated manually because the IIS 6.0 Migration Tool migrates content only in the home directory and subdirectories of the Web site.

**Migrating Content**

Migrate content that is located outside the home directory and subdirectories of the Web site by completing the following steps:

1. Create a folder on the target server to contain the content. Ensure that the folder is in the same relative location to the home directory of the Web site. For example, if a Web site on the source server is in D:\Wwwroot\WebSite and the content to be migrated is in D:\Program Files\SiteContent, create the same folder (D:\Program Files\SiteContent) on the target server.

2. Copy the content from the source server to the target server. Use the process of your choice for copying the content as described in Migrating Web Site Content.

3. If necessary, modify any references to the content that is located in directories that are external to the home directory of the Web site on the target server. The target server might have a different disk configuration than the source server, or you might be migrating the content to a different disk volume letter on the target server than the source server. You need to modify any references in the Web site to content that is located outside the home directory and subdirectory of the Web site.

# Migrating Content Located in Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Virtual directory content on the source server might be stored on a disk volume, such as F:, that does not exist on the target server. This content must be migrated manually because the IIS 6.0 Migration Tool migrates virtual directories only when the disk volume letter on which the virtual directory is located exists on both the source server and the target server.

Migrate content that is located in virtual directories, which are stored on a disk volume that does not exist on the target server, by completing the following steps:

1. Create a folder on the target server to contain the virtual directory and the content.

   Ensure that the folder is in the same relative location on the target servers as on the source server.

2. Copy the content from the source server to the target server.

   Use the process of your choice for copying the content as described in Migrating Web Site Content.

3. Create the virtual directory under the appropriate Web site on the target server that references the folder created in Step 1.

   For more information about how to create virtual directories on the target server, see Create a Virtual Directory in IIS 6.0.

# Modifying IIS 6.0 Metabase Properties That Reference the Systemroot Folder

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 Migration Tool migrates metabase properties that reference the systemroot folder on the source server, but does not update the references to the systemroot folder on the target server. If the location of systemroot folder on the target server does not match the location of the systemroot folder on the source server, you must modify the metabase settings on the migrated Web sites to reference the correct systemroot folder on the target server. Because the default systemroot folder name changed from WINNT to Windows in Windows 2000 Server and later versions, you might need to manually modify the metabase properties that reference the systemroot folder on the target server.

These metabase properties that reference the location of the systemroot folder can include the following:

- **HttpErrors**. The **HttpErrors** metabase property specifies the custom error string sent to clients in response to HTTP 1.1 errors. Each string in the list specifies the HTTP error code and subcode, indicates whether the handler will be a URL or a file, and specifies which URL or file the client will be sent. Each string can be in either a URL or a file format. If you migrate the Default Web Site on Windows NT 4.0 or Windows 2000, you must reset the **HttpErrors** metabase property to contain the systemroot folder for Windows Server 2003, even if the systemroot folder exists on the same drive on both the source and target computer.

- **ScriptMaps**. The **ScriptMaps** metabase property specifies the file name extensions of applications that are used for script processor mappings.This property contains references to paths with default ISAPIs, such as C:\Windows\system32\inetsrv\asp.dll.

Compensate for the differences in the location of the systemroot folder on the target server by completing the following steps:

1. Enable the IIS 6.0 metabase edit-while-running feature.

2. Open the MetaBase.xml file in Microsoft Notepad.

3. Search for any references to the systemroot folder on the source server and replace these references with the systemroot folder path on the target server.

   For example, if the source server had been installed in C:\WINNT and the target server is installed into C:\Windows, you should replace any occurrences of "C:\WINNT" with "C:\Windows".

4. Save the MetaBase.xml file.

5. Disable the metabase edit-while-running feature.

   For more information about the edit-while-running feature, see Enabling Edit-While-Running in IIS 6.0.

# Migrating Web Sites Manually

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Earlier in the process, you decided whether to migrate your Web sites with the IIS 6.0 Migration Tool or to migrate them manually. If you are unable to use the migration tool, you must migrate manually, or use existing provisioning or setup scripts. For more information about how to determine whether you can perform the migration with the IIS 6.0 Migration Tool, see Selecting a Migration Method earlier in this section.

When there are provisioning or setup scripts for your Web sites and applications, use these scripts to install the Web sites and applications on the target server. These scripts might require modification to be compatible with worker process isolation mode, which is discussed in Evaluating Application Changes Required for IIS 6.0 Worker Process Isolation Mode earlier in this section.

Because the scripts install and configure the Web sites and applications, no migration is required. In this case, run the scripts to install and configure the Web sites and applications and then proceed to Configuring IIS 6.0 Properties later in this section to continue with the process.

Figure 6.6 illustrates the process migrating Web sites manually to IIS 6.0.

**Figure 6.6 Performing a Manual Migration to IIS 6.0**



# Verifying That Clients Are Not Accessing Web Sites Before a Manual Migration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before migrating your existing Web sites and applications, ensure that no active client sessions are running. For more information about verifying that clients are not accessing Web sites, see Verifying That Clients Are Not Accessing Web Sites earlier in this section.

# Creating Web Sites and Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For each Web site and virtual directory on the source server, you must create a corresponding Web site and virtual directory on the target server. Later in the migration process, you will copy the content into these Web sites and virtual directories.

Create the Web sites and virtual directories by completing the following steps:

1. Create the Web sites and home directories on the target server.

2. Create the virtual directories.

# Creating Web Sites and Home Directories on the Target Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Each Web site must have one home directory. The home directory is the central location for your published pages. It contains a home page or index file that welcomes visitors and contains links to other pages in your site. The home directory is mapped to the Web site's domain name or to the name of the Web server.

Create a Web site and home directory on the target server by completing the following steps:

1. Create the folder that will be the home directory for the Web site on the target server.

   The folder that is the home directory of the Web site contains all of the content and subdirectories for the Web site. The folder can be created on the Web server or on a UNC-shared folder on a separate server. At a minimum, create the folder on the following:

   - An NTFS file system partition, which helps ensure proper security.
   - A disk volume other than the system volume, which reduces the potential of an attack on a Web site bringing down the entire Web server, and improves performance.

   For more information about securing Web sites and applications see Securing Web Sites and Applications. For more information about creating directories for your Web sites see Create a Web Site.

2. Determine whether to generate the Web site identification number incrementally, or from the Web site name.

   Although site identification numbers were generated incrementally in IIS 5.1 and earlier, when you create a new site on IIS 6.0, a Web site identification number is randomly generated by using the name of the Web site. If you have administration scripts, setup programs, or provisioning scripts that depend upon the IIS 5.1 method of generating site identification numbers, you can force IIS 6.0 to use incremental site identification numbers by creating the IncrementalSiteIDCreation registry entry in the subkey HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\InetMgr\Parameters with the data type REG_DWORD and the value set to 0x1.

   For more information about configuring the Web site identification number see Configure the Web Site Identification Number.

3. Create the Web site on the target server.

   Configure the Web site to have the same configuration as the corresponding Web site on the source server. For more information about how to create a Web site, see Create a Web Site.

4. If the Web site on the source server is FrontPage extended, then configure the Web site on the target server to be FrontPage extended.

   For more information about how to configure a Web site to be FrontPage extended see Configure a Web Site to be FrontPage Extended.

# Creating Virtual Directories on the Target Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For each virtual directory within each Web site on the source server, you must create a corresponding virtual directory on the target server. A *virtual directory* is a folder name, used in an address, which corresponds to a physical directory on the Web server or a Universal Naming Convention (UNC) location. This is also sometimes referred to as *URL mapping*. Virtual directories are used to publish Web content from any folder not located in the home directory of the Web site. When clients access content in a virtual directory, the content appears to be in a subdirectory of the home directory even though it is not.

For security reasons, you might want to move the Web site content to a different disk volume during the migration process. You can move the content to another disk volume on the target server or to a shared folder on a separate server. You can use virtual directories to specify the UNC name for the location where the content is placed, and provide a user name and password for access rights.

For each virtual directory in each Web site on the source server, create a corresponding virtual directory on the target server by completing the following steps:

1. Create the folder on the target server to contain the virtual directory content.

   Create the folder in the same location on the target server unless you are placing the content on a different disk volume than the source server or you are using a UNC share to store the content. Ensure that you create the folder in a secure manner that does not compromise the security of the target server.

   For more information about securing virtual directories, see "Preventing Unauthorized Access to Web Sites and Applications" in Securing Web Sites and Applications.

2. Create the virtual directory under the appropriate Web site on the target server.

   For more information about how to create virtual directories, see Create a Virtual Directory in Appendix A: IIS Deployment Procedures.

# Migrating Web Site Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For each Web site and virtual directory on the source server, you must migrate the content to the corresponding Web site and virtual directory on the target server. You can migrate the content from the source server to the target server by using one of the following methods:

- Run the **Xcopy** command to migrate Web site content to the target server on an intranet or internal network.
- Use Windows Explorer to migrate Web site content to the target server on an intranet or internal network.
- Use the **Copy Project** command in Microsoft Visual Studio® .NET to migrate Web site content to the target server on an intranet or internal network, if the application has been developed by using Visual Studio .NET.

  | **Note** |
  | --- |
  | Front Page Server Extensions must be installed on the Web server to use the **Copy Project** command. |

- Use the **Publish Web** command in FrontPage to migrate Web site content to the target server on an intranet or over the Internet, if the Web site has been developed using FrontPage.

  For more information about how to publish Web site content on the target server by using FrontPage, see Web Site Content with FrontPage in Appendix A: IIS Deployment Procedures.

# Configuring Web Site Application Isolation Settings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Based on the application isolation mode settings of the target server, you need to configure the application isolation settings for each migrated Web site. Configure the Web site application isolation settings so that the Web sites provide the highest possible security and availability.

Configure the Web site applications isolation settings by completing the following steps for each Web site on the source server:

1. Document the current application isolation settings for each Web site on the source server.

2. When the target server is configured for IIS 5.0 isolation mode, configure the target server to use the same the isolation settings as the source server.

3. When the target server is configured for worker process isolation mode, convert the isolation settings on the source server to application pool settings on target server.

# Documenting the Current Application Isolation Settings on the Source Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you configure the application isolation settings, document the existing application isolation settings of the Web sites and applications that are hosted on the source server. Later in the migration process, you will use these settings for configuring the application isolation mode for your Web sites and applications.

For each Web site and application currently running on the server, document the following:

**Application isolation settings**
Earlier versions of IIS can host Web sites and applications in pooled or isolated process configurations. For information about how to view the current application isolation mode, see View Application Isolation Configuration.

If you are running IIS 4.0 on Windows NT Server 4.0, your applications are isolated in one of the following ways:

- In-process (running in-process with Inetinfo.exe)
- Isolated (running under MTS)

If you are running IIS 5.0 on Windows 2000, your applications are isolated in one of the following ways:

- In-process (running in-process with Inetinfo.exe)
- Pooled (running in the pooled COM+ application)
- Isolated (running in an isolated COM+ application)

**Process identity that is used by the Web site or application**
Each Web site or application configured in High isolation, or pooled isolation, uses a configurable *identity*. An identity is a user account that provides a security context for worker process servicing the Web site or application. The identity can be used to secure content, by using NTFS permissions or data, such as data stored in Microsoft SQL Server™. For more information about how to view the identity for each Web site or application, see View Web Site and Application Process Identities.

| **Note** |
| --- |
| All Web sites and applications that are configured to run in the Inetinfo.exe process run under the security context of LocalSystem. |

# Configuring Application Isolation Settings in IIS 5.0 Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When the target Web server is configured to run in IIS 5.0 isolation mode, configure the application isolation settings on the target server identically to the settings on the source server. Web sites and applications on a Web server running in IIS 5.0 isolation mode can be configured with the following application isolation settings:

- Low (in-process).
- Medium (pooled).
- Low (isolated).

If the identity on the source server is an account local to the source server, you need to create a service account.

Configure the application isolation settings when IIS 6.0 is configured to run in IIS 5.0 isolation mode by completing the following steps:

1. Review the application isolation settings on the source server, documented earlier in the migration process.

   For more information about how the application isolation settings were documented, see Documenting the Current Application Isolation Settings on the Source Server.

2. Create any required local service accounts used for application isolation identities on the target server.

   When the application pool identity is a service account that is local to the source server, you need to create a new service account, or designate an existing service account, on the target server. Create the service account in Active Directory to:

   - Provide centralized administration of the account.
   - Provide stronger security because the account is stored in Active Directory rather than locally on the Web server.
   - Allow more than one Web server (for instance, in a Web farm) to use the same service account for the same instance of the application pool on other Web servers.

   For more information about how to create a service account to be used as an identity, see Create a Service Account.

3. Configure the application isolation settings for the Web sites on the target server to be identical to the settings in Step 1.

   For more information about how to configure the application isolation settings for a Web site, see Configure Application Isolation Settings for IIS 5.0 Isolation Mode.

4. Configure the application process identities for the Web sites on the target server to be identical to the settings in Step 1.

   For more information about how to configure the application process identities for the Web site, see Configure Application Identity for IIS 5.0 Isolation Mode.

# Configuring Application Isolation Settings in IIS 6.0 Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When the target server is configured to use worker process isolation mode, you need to configure the application isolation settings to closely approximate their configuration in IIS 5.0 isolation mode by assigning them to *application pools*. An application pool is a grouping of one or more Web sites or applications served by one or more worker processes. You might need to apply additional configurations so that the applications retain their original isolation settings.

After converting to worker process isolation mode, all applications run in the preexisting application pool named "DefaultAppPool." If all of the applications run in the same process in the previous version of IIS, then they all are assigned to the default application pool.

However, if any one of the applications in the same application pool fails, the other applications can be adversely affected. For this reason it is recommended that you isolate your applications into separate application pools whenever possible.

Configure Web sites and applications to run in their own application pool by completing the following steps:

**For each Web site or application configured in High isolation in IIS 5.0**

1. Create a new application pool to be used by the Web site or application.

   For information about how to create application pools, see Isolate Applications in Worker Process Isolation Mode.

2. If the Web site or application previously ran under an identity that is still required by the Web site or application, configure the application pool to use that same identity.

   For information about how to configure the identity for an application pool, see Configure Application Pool Identity.

3. Assign the Web site or application to the new application pool.

   For information about how to assign the Web site to the new application pool, see Isolate Applications in Worker Process Isolation Mode.

**For each Web site or application configured in Low or Medium isolation in IIS 5.0**
In earlier versions of IIS, applications ran in-process as DLLs in Inetinfo.exe (Low isolation) and the default process identity (account that the application runs under) was LocalSystem. With worker process isolation mode in IIS 6.0, applications never run in Inetinfo.exe. However, any applications that are not explicitly assigned to an application pool are assigned to the default application pool, which runs under the NetworkService process identity by default. Because LocalSystem has the same permissions and user rights as a member of the Administrators group, run Web sites and applications under the security context of the NetworkService account.

For each Web site or application that ran in Low or Medium isolation in IIS 5.0, do one of the following:

- When the Web site or application is able to function under the identity of the NetworkService account in the default application pool, continue to host the Web sites or applications in the default application pool, named "DefaultAppPool."

- When the Web site or application is unable to function under the identity of the NetworkService account in the default application pool, perform the following steps:

  1. Create a new application pool.

  2. Create a service account to be used as the identity for the application pool.

     For more information about how to create a service account to be used as an identity for an application pool, see Create a Service Account.

3. Configure the application pool identity to use the service account.

   For more information about how to configure the identity for an application pool, see Configure Application Pool Identity.

4. Place the Web site or application in the new application pool.

# Configuring IIS 6.0 Properties after Migration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Up to this point in the migration process, you have migrated the Web site content and configuration settings from the source server — either manually or with the IIS 6.0 Migration Tool. However, you might need to further configure the IIS 6.0 properties on the target server so that the Web sites run as they did before they were migrated. In addition, you should configure your target server even further to utilize the enhanced security and availability capabilities of IIS 6.0.

Figure 6.7 illustrates the process for configuring the IIS 6.0 properties on the target server.

Figure 6.7 Configuring IIS 6.0 Properties

# Configuring IIS 6.0 Properties That Reference Local User Accounts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The configuration of IIS and the Web sites on the source server can reference user accounts that are stored in the local account database on the source server. These accounts that are stored locally on the Web server are known as *local user accounts*. Local user accounts are valid only on the Web server where they are stored, not on any other Web servers.

As a result, when IIS, or any of the Web sites on the source server, reference local user accounts you must configure IIS 6.0 and the Web sites on the target server to reference:

- Domain-based user accounts that you create
- Local user accounts that you create on the target server

For each configuration that references a local account on the source server, you need to do the following on the target server:

1. Create, or designate, a domain-based or local user account that you can use to configure IIS.

   For more information about creating a service account that is domain-based or local to the target server, see Create a Service Account.

2. Modify the IIS 6.0 properties, Web site properties, or content configuration settings, based on the type of property that you are configuring.

   The types of IIS 6.0 properties, Web site properties, or content configuration settings that can reference or use local user accounts include:

   - **NTFS permissions assigned to Web content**. Grant the same NTFS permissions to the account created in step 1on the target sever as were granted to the local user account on the source server. For more information about granting NTFS permission to content, see Configure NTFS Permissions.
   - **Anonymous accounts for Web sites**. Configure the anonymous account identity for a Web site to use the account created in Step 1 on the target server. For more information about configuring the anonymous account identity, see Configure Anonymous User Identity.
   - **Application isolation settings**. Configure the application isolation identity based on the application isolation mode configured for the server.

     When the target server is configured to run in worker processor isolation mode, configure the identity properties of the application pools. For more information about configuring the application isolation settings when the target server is configured for worker process isolation mode, see Configure Application Pool Identity.

     When the target server is configured to run in IIS 5.0 isolation mode, configure the identity properties of the COM+ applications that correspond to the Web site. For more information about configuring the application isolation settings when the target server is configured for IIS 5.0 isolation mode, see Configure Application Isolation Settings for IIS 5.0 Isolation Mode.

# Configuring IIS 6.0 Web Service Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Many Web sites and applications that are hosted on IIS 6.0 have extended functionality beyond static Web pages, including the generation of dynamic content. Providing dynamic content and other enhanced capabilities requires executable code, such as ASP, ASP.NET, and ISAPI extensions. The handlers that extend IIS functionality beyond serving static pages are known as *Web service extensions*.

If you installed IIS 6.0 as described earlier in this section, all Web service extensions are disabled by default. If you used another method to install IIS 6.0, such as using Manage Your Server, the configuration of IIS might be different.

Enabling all of the Web service extensions ensures the highest possible compatibility with your Web sites. However, enabling all of the Web service extensions creates a security risk because it increases the attack surface of IIS by enabling functionality that might be unnecessary for your server.

Web service extensions allow you to enable and disable the serving of dynamic content. *MIME types*allow you to enable and disable the serving of static content. For more information about enabling and disabling the serving of static content, see Configuring MIME Types later in this section.

| **Tip** |
| --- |
| If the appropriate Web service extension is not enabled, the Web server returns a 404 error to the client when attempting to serve the dynamic content. When the 404 error is returned as a result of a Web service extension not being enabled, a 404.2 error entry is placed in the IIS log. For more information about troubleshooting IIS 6.0, see Troubleshooting in IIS 6.0. |

Configure the Web service extensions by completing the following steps:

1. Enable the essential predefined Web service extensions based on the information in Table 6.3.

   **Table 6.3 Predefined Web Service Extensions**

   | Web Service Extension | Description |
   | --- | --- |
   | Active Server Pages | Enable this extension when one or more of the Web sites or applications contains ASP content. |
   | ASP.NET version 1.1.4322 | Enable this extension when one or more of the Web sites or applications contains ASP.NET content. |
   | FrontPage Server Extensions 2002 | Enable this extension when one or more of the Web sites are FrontPage extended. |
   | Internet Data Connector | Enable this extension when one or more of the Web sites or applications uses the Internet Data Connector (IDC) to display database information (content includes .idc and .idx files). |
   | Server-Side Includes | Enable this extension when one or more of the Web sites uses server-side include (SSI) directives to instruct the Web server to insert various types of content into a Web page. |
   | WebDAV | Enable this extension when you want to support Web Distributed Authoring and Versioning (WebDAV) on the Web server, but it is not recommended for dedicated Web servers. |

2. For each Web service extension that is used by your applications and is not a one of the default Web service extensions, add a new entry to the Web service extensions list and configure the status of the new entry to **Allowed**.

   For example, one of your applications might use an ISAPI extension to provide access to a proprietary database. Set the ISAPI extension used by the application to **Allowed** to explicitly grant it permission to run. For information about how to add a Web service extension to the list, see Configure Web Service Extensions.

3. Use a Web browser on a client computer to verify that the Web sites and applications run on the server.

# Configuring MIME Types in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 serves only the static files with extensions that are registered in the Multipurpose Internet Mail Extensions (MIME) types list. IIS 6.0 is preconfigured to recognize a default set of global MIME types, which are recognized by all configured Web sites. You can define MIME types at the Web site and directory levels, independently of one another or the types defined globally. IIS also allows you to change, remove, or configure additional MIME types. For any static content file extensions used by the Web sites hosted by IIS that are not defined in the MIME types list, you must create a corresponding MIME type entry.

Configure the MIME types after migration by completing the following steps:

1. For each static file type used by your Web site, ensure that an entry exists in the MIME types list.

   When your application uses the standard MIME types that are included in IIS 6.0, no new MIME types entry is required. For information about how to add a MIME type to the MIME types list, see Configure MIME Types.

2. Use a Web browser on a client computer to verify that the Web sites and applications run on the server.

# Migrating Server Certificates for SSL

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you use Secure Sockets Layer to encrypt confidential information exchanged between the Web server and the client, you must migrate the server certificate from the source server to the target server, install the certificate on the target server, and then configure the Web site to use the certificate.

| **Note** |
| --- |
| Server certificates are installed on the Web server and typically require no additional configuration on client servers. Server certificates allow clients to verify the identity of the server. Alternatively, some Web sites and applications might require client certificates. Client certificates are installed on the client servers and allow the server to authenticate the clients. For more information about configuring client certificates, see Certificates_IIS_SP1_Ops. |

Migrate the server certificate for Secure Sockets Layer (SSL) by completing the following steps for each Web site and application that uses SSL:

1. Export the server certificate for the Web site from the source server.

   For more information about exporting a server certificate, see Export a Server Certificate.

2. Install the server certificate to be used by the Web site on the target server.

   For more information about installing the server certificate on the Web server by using the Certificate MMC snap-in, see Install a Server Certificate.

3. Assign the server certificate to the Web site.

   For more information about assigning the server certificate to the Web site, see Assign a Server Certificate to a Web Site.

# Migrating FrontPage Users and Roles

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When the source server has Web sites that are FrontPage extended and FrontPage roles have been assigned to the Web site users, you need to migrate the FrontPage roles to the target server. The FrontPage roles control the types of access that users have on FrontPage extended Web sites. FrontPage 2002 Server Extensions are administered through the Microsoft SharePoint™ Team Services HTML administration tool, which is installed with FrontPage 2002 Server Extensions.

The predefined FrontPage roles include the following:

- Administrator. Users assigned this role can view, add, and change all server content; and manage server settings and accounts.
- Advanced author. Users assigned this role can view, add, and change pages, documents, themes, and borders; and recalculate hyperlinks.
- Author. Users assigned this role can view, add, and change pages and documents.
- Contributor. Users assigned this role can view pages and documents, and view and contribute to discussions.
- Browser. Users assigned this role can view pages and documents.

In addition to the predefined FrontPage roles, custom FrontPage roles might be defined on the source server.

Migrate FrontPage users and roles by completing the following steps:

1. Identify the FrontPage roles on the source server and compare them to the FrontPage roles on the target server.
2. Create any FrontPage roles on the target server that exist on the source server but do not exist on the target server.
3. For each FrontPage user on the source server that is local to the source server, create a corresponding user on the target server, and then assign that user the same FrontPage roles that are assigned to the corresponding user on the source server.
4. For each FrontPage user on the source server that is in Active Directory, assign the user the same FrontPage roles on the target server.

For more information about configuring the FrontPage 2002 Server Extensions users and roles, see Configure FrontPage Server Roles. For more information about administering FrontPage 2002 Server Extensions, see Installation.

# Configuring IIS 6.0 for Hosting ASP.NET Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you migrated any ASP.NET applications to the target server, you need to configure IIS 6.0 to use the correct version of the .NET Framework, and you must configure the .NET Framework to support your applications.

Configure IIS to Host ASP.NET Applications

Configure IIS 6.0 to host ASP.NET applications by completing the following steps:

1. Configure IIS 6.0 to use the correct version of the .NET Framework.

2. Configure the .NET Framework.

3. Review how ASP.NET applications run in each application isolation mode.

4. Migrate Machine.config attributes to IIS 6.0 metabase property settings.

# Configuring IIS to Use the Correct Version of the .NET Framework

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you have migrated ASP.NET applications that were developed with version 1.0 of the .NET Framework, you might have to install version 1.0 of the .NET Framework on the target server to ensure that your applications continue to function properly. After migration, version 1.1 of the .NET Framework is installed on the target server and the applications are configured to use version 1.1 of the .NET Framework. After installing version 1.0 of the .NET Framework, both version 1.0 and 1.1 of the .NET Framework are installed on the target server. This is referred to as *side-by-side support*.

Running versions 1.0 and 1.1 of the .NET Framework side-by-side is only supported when IIS is configured to run in IIS 5.0 isolation mode. If you have already configured IIS to run in worker process isolation mode, then you can only use version 1.1 of the .NET Framework on the target server. In most cases, ASP.NET applications function correctly with version 1.1 of the .NET Framework. For more information about possible application incompatibilities when migrating from version 1.0 to version 1.1 of the .NET Framework, see Determining Application Compatibility with the .NET Framework earlier in this section. When your ASP.NET application is incompatible with version 1.1 of the .NET Framework, configure the application to use version 1.0 of the .NET Framework and configure IIS to run in IIS 5.0 isolation mode.

You can configure each ASP.NET application to use a specific version of the .NET Framework by registering a *script map* in IIS for the application. A script map associates a file name extension and HTTP verb with the appropriate ISAPI for script handling. For example, when IIS receives a request for a .aspx file, the script map for the corresponding application directs IIS to forward the requested file to the appropriate version of the ASP.NET ISAPI for processing.

The script map for each ASP.NET application can be applied directly to an application, or inherited from a parent application. However, ASP.NET supports only one version of the .NET Framework for each application pool. For more information about how to configure the script map for an ASP.NET application, see Configure an ASP.NET Application for ASP.NET.

# Configuring the .NET Framework for IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The configuration method for the .NET Framework is determined by the application isolation mode that you use to configure IIS 6.0. Table 6.4 lists the methods for configuring the .NET Framework that are associated with each IIS 6.0 application isolation mode.

**Table 6.4 Methods for Configuring the .NET Framework**

| Application Isolation Mode | Configuration Method for the .NET Framework |
|---|---|
| IIS 5.0 isolation mode | Configured by making changes to the Machine.config file in the *%systemroot%*\Microsoft.NET\Framework*VersionNumber*\Config folder. |
| Worker process isolation mode | Configured by making changes to the IIS 6.0 metabase. |

When IIS 6.0 is configured to run in IIS 5.0 isolation mode, the .NET Framework uses the **<processModel>** section of the Machine.config file (in the *%systemroot%*\Microsoft.NET\Framework*versionNumber*\Config folder) for its configuration and no additional steps are required.

However, if you configured IIS 6.0 to run in worker process isolation mode, the .NET Framework ignores the **<processModel>** section of the Machine.config file, and instead gets its process configuration from the IIS 6.0 metabase. Because the migration process does not migrate the existing settings in the Machine.config file, you must manually convert any settings required by the ASP.NET applications.

For information about how to convert the Machine.config attribute settings to IIS 6.0 metabase property settings, see Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings later in this section. For more information about configuring IIS 6.0 for ASP.NET applications, see Deploying ASP.NET Applications in IIS 6.0.

# Reviewing How Applications Run in Each Application Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When running IIS 6.0 in worker process isolation mode, ASP.NET applications use the W3wp.exe worker process and application pool properties, which are stored in the IIS 6.0 metabase. When you configure IIS 6.0 to run in IIS 5.0 isolation mode, ASP.NET applications use the ASP.NET request processing model, Aspnet_wp.exe, and configuration settings. These configuration settings are stored in the Machine.config file.

Behavior of ASP.NET Applications That Are Running in IIS 5.0 Isolation Mode

By default, ASP.NET applications are configured to run in worker process isolation mode. If your application can only run in the ASP.NET process model, you must configure the server to run in IIS 5.0 isolation mode to be able to run the application on IIS 6.0. When IIS 6.0 is configured to run in IIS 5.0 isolation mode, ASP.NET applications behave as follows:

- The applications run within Aspnet_wp.exe.

  Aspnet_wp.exe is a request processing model that is similar to worker process isolation mode, and it contains worker process management capabilities similar to the WWW service in IIS 6.0. Aspnet_wp.exe includes most of the IIS application management features, such as recycling, health detection, and *processor affinity*. Processor affinity is the ability to force worker processes to run on specific microprocessors.

- The configuration settings are stored in the Machine.config file.

  When IIS 6.0 is running in IIS 5.0 isolation mode, the configuration settings for ASP.NET applications are managed by modifying the Machine.config file, not the IIS metabase file. Because there is no administrative console for the Machine.config settings, any configuration settings for ASP.NET must be made directly to the Machine.config file.

| **Important** |
| --- |
| In IIS 5.0 isolation mode, the .NET Framework ignores any configuration changes made in the IIS metabase. Administrative consoles, such as IIS Manager, make changes to the IIS 6.0 metabase, but those changes are not read by the .NET Framework. |

When IIS 6.0 is configured to run in IIS 5.0 isolation mode, your ASP.NET applications should behave as they did in IIS 5.0. However, incompatibilities can result when running version 1.1 of the .NET Framework. For more information about configuring IIS to support ASP.NET applications that use version 1.0 of the .NET Framework, see Overview of Deploying ASP.NET Applications in IIS 6.0.

Behavior of ASP.NET Applications That Are Running in Worker Process Isolation Mode

When IIS 6.0 is configured to run in worker process isolation mode, ASP.NET applications behave as follows:

- The process model within the ASP.NET ISAPI extension is disabled, and ASP.NET applications run using worker process isolation mode in IIS 6.0.

  In this configuration, the ASP.NET application runs in worker process isolation mode like any other application, such as an ASP application. In addition, IIS 6.0 provides all of the management features such as recycling, health detection, and processor affinity.

- The ASP.NET ISAPI extension is configured by a combination of configuration settings that are stored in the IIS metabase (MetaBase.xml) and configuration settings in the Machine.config file.

  When IIS 6.0 is running in worker process isolation mode, the majority of the application configuration settings are stored in the IIS metabase. You can adjust these settings directly in MetaBase.xml or from administrative consoles, such as IIS Manager, or scripts.

However, if there are existing settings in the **<processModel>** section of the Machine.config file, those configuration settings must be converted to the appropriate application pool settings when the Web server is configured to run in worker process isolation mode Additionally, there are other configuration settings that are still modified in the Machine.config file, regardless of the application isolation mode. For more information about converting Machine.config attributes to worker process isolation mode settings, see Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings.

When IIS 6.0 is configured in worker process isolation mode, your ASP.NET applications should behave as they did on IIS 5.0. Before deploying your ASP.NET applications on your production Web servers, test compatibility with IIS 6.0 running in worker process isolation mode and version 1.1 of the .NET Framework. For more information about determining application compatibility with IIS 6.0 running in worker process isolation mode, see Determining Application Compatibility with Worker Process Isolation Mode. For more information about determining compatibility with version 1.1 of the .NET Framework see Determining Application Compatibility with the .NET Framework.

If you determine that your ASP.NET applications are incompatible with worker process isolation mode, reconfigure IIS to run in IIS 5.0 isolation mode. If your ASP.NET applications are incompatible with version 1.1 of the .NET Framework, configure IIS to use version 1.0 of the .NET Framework with your ASP.NET application.

For more information about configuring IIS 6.0 to run in IIS 5.0 isolation mode, see Configure Application Isolation Modes. For more information about configuring IIS to use version 1.0 of the .NET Framework with your ASP.NET application, see Overview of Deploying ASP.NET Applications in IIS 6.0.

# Migrating Machine.config Attributes to Metabase Property Settings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When IIS 6.0 is configured to run in IIS 5.0 isolation mode, the .NET Framework uses the **<processModel>** section of the Machine.config file (in the *%systemroot%*\Microsoft.NET\Framework*versionNumber*\Config folder) for its runtime configuration. When the target server is running in IIS 5.0 isolation mode, you do not need to convert the attribute configuration settings on the source server to their equivalent IIS 6.0 metabase property settings on the target server so you can proceed to the next step in the migration process. To proceed to the next step in the migration process, see Determining Whether to Run the IIS Lockdown Tool and UrlScan.

However, if you configured IIS 6.0 to run in worker process isolation mode, the .NET Framework ignores the **<processModel>** section of the Machine.config file, and instead gets its process configuration from the IIS 6.0 metabase. Because the migration process does not migrate the existing settings in the Machine.config file, you must manually migrate any settings that are required by your ASP.NET applications.

For information about how to migrate the Machine.config settings to IIS 6.0 metabase settings, see Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings.

# Running the IIS Lockdown Tool and UrlScan

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS Lockdown Tool and UrlScan are IIS security related programs designed for IIS 5.1 and earlier. Each tool provides different types of protection for earlier versions of IIS. The IIS migration process does not install the IIS Lockdown Tool and UrlScan on the target server.

IIS Lockdown Tool

The IIS Lockdown Tool is provided to assist administrators in configuring optimal security settings for existing IIS servers. You cannot install the IIS Lockdown Tool after migration because all of the default configuration settings in IIS 6.0 meet or exceed the security configuration settings made by the IIS Lockdown Tool.

UrlScan

UrlScan is a tool that is provided to reduce the attack surface of Web servers running earlier versions of IIS. By default, IIS 6.0 has features that significantly improve security by reducing the attack surface of the Web server. UrlScan provides flexible configuration for advanced administrators, while maintaining the improved security in IIS 6.0. When you need this flexibility in configuring your Web server, you can run UrlScan on IIS 6.0.

For more information about determining whether to run UrlScan after migrating your server to IIS 6.0, and to download the latest version of the URLScan tool, see the Knowledge Base article 307608, Using UrlScan on IIS.

# Performing Application-Specific Migration Tasks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When no setup or provisioning scripts exist for your applications, you might need to perform additional application-specific migration tasks, whether you have migrated content and metabase information manually or with the IIS 6.0 Migration Tool. Some of these tasks involve the modification of application code and might require the assistance of the application developers. Depending on your application, you might need to perform any combination of the steps in Figure 6.8.

**Figure 6.8 Performing Application-Specific Migration Tasks**



# Modifying Application Code for Compatibility with Windows Server 2003 and IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If any application code needs to be changed or recompiled to run on Windows Server 2003, you must make those changes manually. The most common application code changes include the following:

- Code that references Windows platform components or APIs no longer supported in Windows Server 2003.
- Code that references IIS metabase properties that have changed or are no longer supported in IIS 6.0.
- Code that is incompatible with worker process isolation mode.

# Modifying References to Windows Platform Components and APIs No Longer Supported in Windows Server 2003

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Applications developed to run on earlier versions of Windows server operating systems and IIS can call APIs that are not supported in Windows Server 2003 and IIS 6.0. Usually these APIs are replaced with newer APIs that provide additional functionality.

If your Web sites and applications use an API that is not supported in Windows Server 2003 and IIS 6.0, you must modify your code to use an API that provides the same functionality and is supported in Windows Server 2003 and IIS 6.0.

For example, the APIs supported by the Collaboration Data Objects for Windows NT Server 4.0 (CDONTS) DLL are not supported in Windows Server 2003. The functionality provided by CDONTS is supported by Collaboration Data Objects for Windows2000 (CDOSYS), which, in turn, is supported by Windows Server 2003. Therefore, if your applications use CDONTS, you can modify your code to use CDOSYS instead. For more information about modifying your application to use CDOSYS, see Collaboration Data Objects Roadmap.

# Modifying References to IIS Metabase Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Some metabase properties that were used to configure features in earlier versions of IIS are no longer supported in IIS 6.0. Because some features are eliminated, or implemented differently in IIS 6.0, the corresponding unused metabase properties are not referenced by any code in IIS 6.0. In cases where the feature is implemented differently in IIS 6.0, new metabase properties have been created to replace the obsolete, or unused, properties.

In addition, there is one IIS 5.0 metabase property — **CPUResetInterval** — whose behavior has changed because of architectural changes made to IIS 6.0.

To determine whether any of your Web sites, applications, or setup programs reference these changed or unsupported IIS metabase properties, see Appendix B: Changes to Metabase Properties in IIS 6.0. You can then follow the recommendations associated with each changed metabase property listed in this appendix to accommodate functionality changes in IIS 6.0.

| **Tip** |
| --- |
| The metabase properties that are no longer supported in IIS 6.0 are not available in IIS 6.0, even when IIS 6.0 is configured to run in IIS 5.0 isolation mode. |

# Modifying Applications To Be Compatible with Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Most applications that were developed to run on earlier versions of IIS run in worker process isolation mode without modification. However, you might need to modify your applications to make them compatible with worker process isolation mode. For more information about the types of modifications that you might need to make to your applications, see Evaluating Application Changes Required for Worker Process Isolation Mode earlier in this section.

Your applications can require other modifications that are not described in this section because of the myriad of approaches to developing applications. For the most current information about modifying your applications to be compatible with worker process isolation mode, see IIS 6.0 Application Isolation Modes.

# Installing Additional Software Required by Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Some of your applications might require or be dependent upon additional software that is not in the same folder with the Web site content. This software can be developed by your organization, by Microsoft, or by other organizations. You must install this software on the target server by running the setup or installation program that accompanies the software.

Examples of this type of software include software that creates reports, integrates databases with dynamic content generation, or provides connectivity with applications running on other servers with non-Microsoft operating systems.

# Migrating MTS Packages, COM Objects, and COM+ Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to installing any software that you installed earlier in the migration process, there can be MTS packages, COM objects, and COM+ applications that need to be migrated. Table 6.5 lists the tasks that you need to complete to migrate MTS packages, COM objects, and COM+ application from earlier version of Windows.

**Table 6.5 Migrating MTS Packages, COM Objects, and COM+ Applications**

| Source Server | Target Server | Migration Tasks |
|---|---|---|
| MTS packages on Windows NT 4.0 Server | COM+ applications | Convert the MTS package to a COM+ application. |
| COM objects | COM objects | Copy COM object files (.dll or .exe) from the source server to the target sever and register the COM object by running the regsrv32 command. |
| COM+ applications on Windows 2000 Server | COM+ applications | Create the COM+ application on the target server and use the same configuration on the target server that existed on the source server. |

For more information about migrating MTS packages, COM objects, and COM+ applications to Windows Server 2003, see the following articles: COM: Delivering on the Promises of Component Technology, and Microsoft .NET/COM Migration and Interoperability.

An example of registering a COM object that is running on the source server is CDONTS. If you decide to not modify your code to us CDOSYS instead of CDONTS, you must copy the Cdonts.dll file from the source server to the target server and then register Cdonts.dll with Windows Server2003. For more information about how to migrate CDONTS, see Migrate CDONTS.

# Modifying ODBC Data Connection Strings and DSNs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If your application establishes database connectivity through an ODBC data connection string or through an ODBC data source name (DSN), you must do any combination of the following:

- Manually modify ODBC DSNs when the ODBC data connection string references a database that is stored on the source server and the database has not been migrated to the target server.
- Manually create a system ODBC DSN on the target server for each system ODBC DSN on the source server.

**Modifying ODBC DSNs**

ODBC data connection strings might need to be modified if they reference the database or are migrated to a different location. For more information about modifying ODBC data connection strings, see Connecting to a Data Source or Driver.

**Creating System ODBC DSNs**

In addition, if you have system ODBC DSNs that are defined on the source server, you need to create corresponding ODBC DSNs on the target server. Table 6.6 lists methods for administering ODBC DSNs on different Windows server operating systems. You can administer ODBC DSNs on the source server and the target server by using the methods listed in this table.

**Table 6.6 Administering ODBC DSNs on Source Servers and Target Servers**

| Operating System | Administered Through | Additional Information |
|---|---|---|
| Windows Server 2003 | Data Sources (ODBC) in Administrative Tools | For more information about using Data Sources (ODBC) in Windows Server 2003, see Help and Support Center for Windows Server 2003. |
| Windows 2000 Server | Data Sources (ODBC) in Administrative Tools | For more information about using Data Sources (ODBC) in Windows 2000 Server, see Windows 2000 Server Help. |
| Windows NT 4.0 Server | ODBC Data Sources in Control Panel | For more information about using ODBC Data Sources in Windows NT 4.0 Server, see Help in the ODBC Data Source Administrator in the ODBC Data Sources, in Control Panel. |

# Creating IP Addresses That Are Used by Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can uniquely identify a Web site or application by associating the Web site or application with a unique IP address, a unique combination of an IP address and a TCP port number, or host headers. For Web sites and applications on the source server that are uniquely identified by IP addresses, you must create corresponding IP addresses on the target server and then configure the Web sites and applications to use the newly created IP addresses.

To create IP addresses that are used by applications, complete the following steps:

1. Determine which Web sites and applications on the source server are uniquely identified by IP addresses.

   For more information about how to determine which Web sites and applications on the source server are uniquely identified by IP addresses see Determine Web Sites Uniquely Identified by IP Addresses.

2. For each Web site and application identified in Step 1, assign new a new IP address to the TCP/IP properties of the network adapter in the target server that the clients use to access the Web sites and applications.

   For more information about how to assign a new IP addresses to the network adapter in the target server that the clients use to access the Web sites and applications see Assign Additional IP Addresses to a Network Adapter.

3. Configure the migrated Web sites and applications on the target server to use the IP addresses assigned in Step 2.

   For more information about how to configure the IP addresses assigned to Web sites and applications see Configure IP Address Assigned to Web Sites.

# Creating Users and Groups That Are Used by Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Web sites and applications on the source server might be accessed by accounts that are local to the source server. The local user and group accounts need to be created on the target server so that the accounts can access the Web sites and applications on the target server.

In cases where the users and groups that are used by the applications exist in Active Directory, no steps are required and you can continue to the next migration step. To continue to the next step in the migration process, see Creating Registry Entries for Applications.

Create the users and groups that are used by applications and are local to the source server by completing the following steps:

1. Identify the users and groups that are local to the source server.

2. For each group on the source server, create a corresponding group on the target server.

3. For each user on the source server, create a corresponding user on the target server.

4. For each user created in Step 3, assign the user to the same groups as the corresponding user on the source server.

   For more information about viewing the users and groups on the source server, see "Creating user and group accounts" in Windows 2000 Server Help and Windows NT Server 4.0 Help. For more information about creating users and groups in Windows Server 2003, see Create a Service Account.

5. For each user created in Step 3, assign the same user rights assigned to the corresponding user on the source server.

   For more information about assigning user rights to a user, see Grant User Rights to a Service Account.

6. Assign the same NTFS permissions for the content on the target server as the NTFS permissions for the content on the source server.

   For more information about configuring NTFS permissions on the target server, see Configure NTFS Permissions.

# Creating Registry Entries for Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Some applications save configuration information in the Windows registry. If the setup program or provisioning script creates the registry entries, run the setup program or provisioning script on the target server. Otherwise, you must manually identify the registry entries and then re-create them on the target server.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and then use the Registry Reference directly or connect to it on the Web at http://www.microsoft.com/reskit. |

Create application registry entries on the target server manually by completing the following steps:

1. Identify the registry entries required by the applications that are currently running on the source server.

   Earlier in the migration process, you identified the registry entries required by the applications on the source server. For more information about identifying registry entries, see Identifying Which Web Site and Application Components to Migrate earlier in this section.

2. Back up the registry entries on the source server that you identified in the previous step by using the registry editor Regedit.exe.

   For more information about backing up copies of registry entries, see Back Up and Restore Registry Entries.

3. Copy the .reg backup file, created in Step 2, to the target server.

4. Modify the .reg backup file created in the previous step to accommodate for changes in disk volume letters, such a F:, or paths to folders by completing the following steps:

   a. In Notepad, open the .reg backup file created in Step 2.

   b. Search for specific references to disk volume letters or paths that you want to change, and update them to reflect the disk volume letters or paths on the target server.

   c. Save the .reg file.

5. Restore the registry backup on the target server.

   For more information about restoring registry entries, see Back Up and Restore Registry Entries.

# Completing the Migration_IIS_SP1_Dep

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

At this point in the process, you have migrated your Web sites and applications to the target server and configured the IIS 6.0 properties to settings on the source server. Now you need to verify that the migration completed successfully, capture the current configuration of the target server, and enable client access to the target server. After you complete these last steps, your migration is complete. Figure 6.9 illustrates the process for completing the migration to IIS 6.0.

**Figure 6.9 Completing the Migration to IIS 6.0**



# Verifying That the Web Sites and Applications Migrated Successfully

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before deploying the target server to a production environment, verify that the Web site content and configuration information migrated successfully by completing the following steps:

1. Review the system log in Windows Server 2003 on the target server to determine whether any of the Web sites did not start.

   IIS 6.0 creates entries in the system log when a Web site fails to start for any reason. Search the System log on the target server for to determine if any errors occurred. For more information about how to troubleshoot Web sites that fail to start, see Troubleshooting in IIS 6.0.

2. Verify that the Web site content migrated to the target server.

   Compare the files and folders for each Web site on the target server with the original files and folders on the source server to determine whether the Web site content has been migrated correctly. Ensure that the number and size of the files and folders on the target server approximates the number and size of the same files on the source server.

3. Verify that the Web site configuration migrated to the target server.

   View the Web site configuration information for a random sampling of the Web sites that were migrated. Compare the Web site configuration on the target server with the corresponding Web site configuration on the source server. Ensure that the configuration for the Web sites has been migrated correctly.

4. Perform functional testing of the migrated Web sites and applications to ensure that the Web sites and applications behave as expected.

   You can find possible causes of application failure by reviewing the Windows Server 2003 and IIS 6.0 Migration Tool logs and application configuration, but the only way to accurately assess whether your Web sites and applications migrated successfully is to perform functional testing. Functional testing is designed to ensure that Web sites and applications are functioning correctly in the most common usage scenarios (for example URLs and inputs).

   Procedures for performing functional testing of Web sites and applications are beyond the scope of this section. For more complete information about general testing, go to the MSDN Library and search for "testing".

# Backing Up the Target Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you enable client access to the target server, perform a complete image backup of the target server. Performing this image backup provides you with a point-in-time snapshot of the Web server. If you need to restore the target server in the event of a failure, you can use this backup to restore the Web server to a known configuration.

| **Important** |
| --- |
| Do not proceed further unless you have a successful backup of the entire target server. Otherwise, you can lose Web sites, applications, or data that you migrated to the target server. |

For more information about how to back up the Web server, see Back Up and Restore the Web Server to a File or Tape.

# Enabling Client Access to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you have migrated your Web sites from the source server to the target server, you are ready to enable client access to the Web sites on the target server while maintaining the DNS entries to the source servers. After a period of time that meets your business needs, you can remove the DNS entries that point to the Web sites on the source server.

Enabling Client Access

Enable client access to the Web sites that are on the target server by completing the following steps:

1. Create the appropriate DNS entries for the Web sites and applications running on the target server. For more information about how to create DNS entries for your Web sites and applications see “Managing resource records” in Help and Support Center for Windows Server 2003.

2. Monitor client traffic to determine whether clients are successfully accessing the target server. For more information about how to monitor client traffic to Web sites on the target server, see Monitor Active Web and FTP Connections.

3. Establish a monitoring period, such as a few hours or a day, to confirm that clients accessing Web sites on the target server are experiencing response times and application responses that meet or exceed your requirements.

4. Remove the DNS entries that are pointing to the Web sites and applications on the source server. For more information about how to remove DNS entries for your Web sites and applications, see “Managing resource records” in Help and Support Center for Windows Server 2003.

5. Prevent new clients from accessing the Web sites on the source server by pausing the Web sites on the source server. For more information about how to pause Web sites on the source server, see Pause Web or FTP Sites.

6. Monitor client traffic to the Web sites on the source server to determine when clients are no longer accessing the source server. For more information about how to monitor client traffic to Web sites on the source server, see Monitor Active Web and FTP Connections.

7. When clients are no longer accessing the Web sites on the source server, decommission the hardware for the source server.

# Additional Resources About Migrating IIS Web Sites to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this section.

**Related Information**

- Deploying ASP.NET Applications in IIS 6.0 for information about configuring IIS 6.0 for ASP.NET applications.
- Appendix A: IIS Deployment Procedures for information about specific procedures for migrating IIS Web sites to IIS 6.0.
- Migrating Apache Web Sites to IIS 6.0 for information about migrating Apache Web sites.
- Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings for information about how to migrate the Machine.config settings to IIS 6.0 metabase settings.
- Securing Web Sites and Applications for information about improving the security of your IIS Web sites after migration.
- "Designing a Test Environment" in Planning, Testing, and Piloting Deployment Projects for information about setting up a test lab.
- COM: Delivering on the Promises of Component Technology for information about migrating MTS packages, COM objects, and COM+ applications to Windows Server 2003.
- Collaboration Data Objects Roadmap for information about modifying your applications to use CDOSYS.
- Compatibility Considerations and Version Changes for a current list of possible compatibility issues when upgrading from version 1.0 to version 1.1 of the .NET Framework.
- Products Designed for Microsoft Windows – Windows Catalog and HCL for information about the hardware devices supported by Windows Server 2003.
- Microsoft .NET/COM Migration and Interoperability for information about migrating MTS packages, COM objects, and COM+ applications to Windows Server 2003.
- Installation in the SharePoint Team Services Administrators Guide for information about administering FrontPage 2002 Server Extensions.
- Knowledge Base article 307608, Using UrlScan on IIS, for information about determining whether to run UrlScan after migrating your server to IIS 6.0.
- Windows Server 2003 for the latest information about compatibility with Windows Server 2003.

**Related IIS 6.0 Operations Guide Topics**

- Certificates_IIS_SP1_Ops for information about configuring client certificates.
- Enabling Edit-While-Running in IIS 6.0 for information about the edit-while-running feature.
- Troubleshooting in IIS 6.0 for information about troubleshooting IIS 6.0.

**Related Windows Server 2003 Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**.

Under **Help Topics**, select the **Search in title only** check box.

- "Features unavailable on 64-bit versions of the Windows Server 2003 family" in Help and Support Center for Windows Server 2003 for information about features, such as ASP.NET, that are not supported on 64-bit versions of Windows Server 2003.

- "Managing resource records" in Help and Support Center for Windows Server 2003 for information about how to create DNS entries for your Web sites and applications.

**Related Tools**

- The "IIS 6.0 Migration Tool User Guide" on the *Windows Server 2003 Deployment Kit* companion CD or the *Internet Information Services (IIS) 6.0 Resource Kit Tools,* which can be downloaded by running IIS60rkt.exe from the Microsoft Download Center.

- Windows Application Compatibility to download the latest version of the Windows Application Compatibility Toolkit.

# Migrating Apache Web Sites to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Migration* is the process of installing a Microsoft® Windows® Server 2003 operating system and Internet Information Services (IIS) 6.0 on a server and then moving or reinstalling existing Apache Web sites and applications to that server. Migrating Apache Web sites and applications to IIS 6.0 allows you to take advantage of the improved security, performance, and management provided by Windows Server 2003 and IIS 6.0. The migration process provides minimal outage of service to the users accessing the Web sites and applications, while ensuring that the existing configuration and content of the Web sites and applications are maintained.

**In This Section**


Overview of Migrating Apache Web Sites to IIS 6.0

Preparing for Migration

Deploying the Target Server

Performing the Migration

Migrating Apache-Specific Extensions

Configuring IIS 6.0

Enabling Client Access After Migration

Additional Resources About Migrating Apache Web Sites to IIS 6.0


**Additional Information**

For information about upgrading IIS Web servers, see Upgrading an IIS Server to IIS 6.0.

For information about migrating Web sites on earlier versions of IIS to IIS 6.0, see Migrating IIS Web Sites to IIS 6.0.

# Overview of Migrating Apache Web Sites to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Apache Web site migration process is completed with the *Apache to IIS 6.0 Migration Tool*. This tool is a set of configurable modules to migrate Apache Web sites — including FrontPage® Server Extensions from Microsoft, .htaccess files, and user directories — to servers running IIS 6.0 on Microsoft Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; and Windows® Server 2003, Web Edition. The tool is a Perl-based command-line utility that migrates both Web content and configuration settings.

The Apache to IIS 6.0 Migration Tool supports Apache versions 1.3.0–1.3.22, as well as the following Linux distributions:

- Red Hat software (Linux versions 6.0, 6.2, 7.0, 7.1, and 7.2)
- Mandrake Linux (versions 8.0, 8.1, and 8.2)
- SuSE (versions 7.3 and 8.0)

You can use the Apache to IIS 6.0 Migration Tool to migrate site content and configuration settings from a *source server* (a server running a supported version of Linux and Apache) to a *target server* (a server running a Windows Server 2003 operating system and IIS 6.0). You can perform the migration from the source server, the target server, or on an *intermediate computer*. The intermediate computer must have File Transfer Protocol (FTP) access to the source and target servers.

The Apache to IIS 6.0 Migration Tool migrates only the Apache Web site content and configuration. If the Apache Web site is made up of only static content, such as Hypertext Markup Language (HTML) pages, the migration can most likely be completed with a few simple steps. However, Apache Web sites containing dynamic content, such as PHP: Hypertext Preprocessor (PHP) or Common Gateway Interface (CGI) scripts, might need modifications to the code in the dynamic content after migration. In addition, any existing provisioning scripts or setup programs for the existing Apache Web sites typically require modifications after the migration process is complete. Because each organization has different provisioning scripts and setup programs, you need to review and make modifications based on input from your developers, and then test the Web site after completing the modifications. For more information about possible necessary modifications to dynamic content, see Determining Web Site Compatibility with IIS 6.0 later in this section.

For more information about the Apache to IIS 6.0 Migration Tool, including how to install it, see Preparing to Migrate to IIS 6.0 later in this section. For more information about the additional migration tasks that you must perform after running the Apache to IIS 6.0 Migration Tool, see Migrating Apache-Specific Extensions later in this section.

# Process for Migrating Apache Web Sites to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to migrating the existing Web site content, the migration process automatically migrates the Web site configuration including FrontPage Server Extensions. Before performing the migration of the Web site, you need to evaluate the compatibility of the Apache Web sites with the software installed on your existing server, including dynamic content, database connections, and *external modules*. External modules are software components that are not shipped as a part of the Linux operating system.

If your Apache Web sites contain dynamic content or you want to further secure the Web sites, you can use the Apache to IIS 6.0 Migration Tool to customize the configuration of IIS 6.0 and migrate the dynamic content. Figure 7.1 illustrates the process for migrating existing Apache Web sites to IIS 6.0.

**Figure 7.1 Migrating Apache Web Sites to IIS 6.0**




| **Tip** |
| --- |
| To upgrade a Web farm, use the process described in this section to upgrade each server in the Web farm. |

Depending on your familiarity with Linux, Windows Server 2003, Apache, IIS 6.0, and the migration process, you might require less guidance through the migration process. The following quick-start guide provides a detailed overview of the Apache to IIS migration process described in this section. You can use this guide to help identify the steps of the migration process where you need additional information and skip the information with which you are already familiar. In addition, all of the procedures required to complete the migration process are documented in Appendix A: IIS Deployment Procedures.

**Prepare for Migration**

1. Determine Web site compatibility with IIS 6.0.
2. Determine application compatibility with worker process isolation mode by evaluating the differences between application isolation modes.
3. Identify the role of the Apache to IIS 6.0 Migration Tool in the migration process.
4. Select one of the following Apache to IIS 6.0 Migration Tool installation options:
    - The source server installation option
    - The target server installation option
    - The intermediate computer installation option

**Deploy the Target Sever**

1. Install Windows Server 2003.
2. Install IIS 6.0.
3. Configure the FTP service.

**Perform the Migration**

1. Install the Apache to IIS 6.0 Migration Tool:
    - Install the migration tool on computers running Linux.
    - Install the migration tool on computers running Windows.
    - Configure the target server for migration.
2. Verify that clients are not accessing the Web sites.
3. Migrate the Web site content.
4. Migrate the Web site configuration.
5. Recover from an interruption in the migration process:
    - Determine the cause of resolve any errors.
    - Restart the migration tool in recovery mode.

**Migrate Apache-Specific Extensions**

1. Migrate dynamic content.
2. Migrate database content and connectivity:
    - Migrate the database content.
    - Migrate the database connectivity.
3. Migrate external modules.

**Configure IIS 6.0 After Migration**

1. Configure Web service extensions.
2. Configure Multipurpose Internet Mail Extensions (MIME) types.
3. Configure Web site properties.
4. Configure server certificates for Secure Sockets Layer (SSL).
5. Back up the target server.

**Enable Client Access After Migration**

Enable client access to the Web sites on the target server.

# Preparing to Migrate to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before initiating the migration, ensure that existing Apache Web sites are compatible with IIS 6.0 and Windows Server 2003. In addition, determine if the existing Web sites are compatible with worker process isolation mode in IIS 6.0. Then, you must identify which steps in the migration process are automated by the Apache to IIS 6.0 Migration Tool and which steps you need to complete after running the tool. Finally, you must decide which option to use to install the tool.

**Figure 7.2 Preparing for Migration of Apache Web Sites to IIS 6.0**



# Determining Web Site Compatibility with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you migrate to IIS, you must determine the compatibility of the Apache Web site with IIS 6.0. Because of architectural differences, Web sites hosted on Apache have different dependencies than Web sites hosted on IIS. For example, some PHP applications running on Apache use MySQL to connect to databases, whereas IIS Web sites use Microsoft ActiveX® Data Objects (ADO) or Open Database Connectivity (ODBC) to connect to databases. After you migrate an Apache Web site that uses MySQL for database connectivity, you must modify the Web site content to use ADO database connectivity.

Table 7.1 lists the Apache Web site components that you must consider before migrating the Apache Web site to IIS 6.0, along with brief descriptions of the compatibility issues, if any, that can result after migration.

**Table 7.1 Apache Web Site Component Migration Compatibility Issues**

| Web Site Component | Description | Compatibility Issues |
|---|---|---|
| Static content | Static content contains no executable code. Examples of static content include HTML, .jpg, and .gif files. | Static content should migrate to IIS 6.0 without any modifications. However, IIS 6.0 MIME types might need to be configured to enable static content file types that are not well known.<br><br>For more information about MIME types, see Configuring File MIME Types on IIS 6.0 later in this section. |
| Dynamic content | Dynamic content contains executable code. Examples of dynamic content include CGI and PHP files. | To run Apache dynamic content on IIS, you must do one of the following:<br>• Run the executable code on an appropriate environment that is compatible with IIS 6.0.<br>• Convert the executable code to an environment that is supported by IIS 6.0.<br><br>For more information, see Migrating Apache-Specific Extensions later in this section. |
| Database connectivity | Database connectivity provides a standardized method of connecting to various database systems. MySQL is the most common database connectivity standard used in Apache Web sites. | To provide database connectivity on IIS 6.0, you must do one of the following:<br>• Obtain a version of MySQL that is compatible with IIS 6.0.<br>• Convert the executable code so that it uses ADO. ADO is the database connectivity standard supported by IIS 6.0.<br><br>For more information, see Migrating Apache-Specific Extensions later in this section. |

| | | |
|---|---|---|
| External modules | External modules provide extended features to Apache Web sites. These components include Cold Fusion and Wpoison. | To use the features provided by these external modules, you must do one of the following:<br>• Obtain a version of the external module that is compatible with IIS 6.0.<br>• Convert the external module to IIS 6.0. For example, Wpoison is written in CGI and can be converted to run on IIS 6.0.<br><br>For more information, see Migrating Apache-Specific Extensions later in this section. |

# Determining Web Site Compatibility with Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 can run in one of two distinct modes of operation, which are called application isolation modes. *Application isolation* is the separation of applications by process boundaries that prevent the applications from affecting one another, and is configured differently for each of the two IIS isolation modes: IIS 5.0 isolation mode and worker process isolation mode.

*Worker process isolation mode* takes advantage of the redesigned architecture for IIS 6.0. This isolation mode runs all application code in an isolated environment; however, unlike earlier versions of IIS, IIS 6.0 provides isolation without a performance penalty because there are no added process hops. Worker process isolation mode is compatible with most existing Web sites and applications. Whenever possible, run IIS 6.0 in worker process isolation mode to take advantage of the enhanced performance and security in IIS 6.0. When IIS 6.0 is configured to use worker process isolation mode, the metabase property **IIs5IsolationModeEnabled** contains a value of **False**.

*IIS 5.0 isolation mode* provides compatibility for applications that depend upon the process behavior and memory model of IIS 5.0. Run IIS in this mode only when a Web site or application has compatibility issues when running in worker process isolation mode and you cannot resolve the problem. When IIS 6.0 is configured to use IIS 5.0 isolation mode, the metabase property **IIs5IsolationModeEnabled** contains a value of **True**.

**Important**

IIS 6.0 cannot run both application isolation modes simultaneously on the same server. Therefore, on a single server running IIS 6.0, you cannot run some Web applications in worker process isolation mode and others in IIS 5.0 isolation mode. If you have applications that require separate modes, you must run them on separate servers.

After you install the target server, IIS is configured to use worker process isolation mode by default. The only reason for configuring IIS to use IIS 5.0 isolation mode is if there are Apache extensions that run only in IIS 5.0 isolation mode. For example, if the preprocessor for the dynamic content in an Apache Web site is incompatible with worker process isolation mode, you must configure IIS to use IIS 5.0 isolation mode. Determine Web site compatibility, in particular Apache-specific extensions to the Web sites, with IIS in your lab before migrating your existing Apache Web sites. In most cases, IIS hosts the Web sites in worker process isolation mode without any problems.

For more information about how to migrate Apache-specific extensions and how to determine compatibility of these extensions with worker process isolation mode, see Migrating Apache-Specific Extensions.

**Note**

Identifying a complete list of potential incompatibilities that Web sites can experience with worker process isolation mode is beyond the scope of the content in this section. Even after following the guidelines in this section, you need to verify in your lab whether your Web sites are compatible with worker process isolation mode.

# Identifying the Role of the Migration Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you migrate Apache Web sites to IIS 6.0, you must know which steps of the migration process are automated by the Apache to IIS 6.0 Migration Tool and which steps you must complete manually. This will allow you to have the appropriate tools and resources available when you are ready to start the migration. In addition, you need to determine the configuration and placement of your Web site after running the tool. The automation provided by the Apache to IIS 6.0 Migration Tool is the same for static and dynamic content. However, depending on the content in your Apache Web site, further configuration might be required after running the tool.

The following steps in the Apache Web site migration process are automated by the Apache to IIS 6.0 Migration Tool:

**Copying of Web site content**
The migration tool creates a copy of the files and folders containing Web site content. The copied files are converted from standard UNIX text file format (end-of-line denoted by carriage return) on the server running Apache to standard Windows text file format (end-of-line denoted by carriage return and line feed characters) on the server running IIS. The Web site content is placed on the server running IIS in a location selected while running the tool.

**Translation of Web site configuration**
The configuration of the Apache Web server and individual Web sites is translated to the corresponding IIS Web server and Web site configuration settings in the IIS metabase.

**Migration of ~user directories**
~user directories are private virtual Web directories that contain individual user Web sites. The tool migrates the contents of these directories, creates corresponding users, and converts permissions found in these virtual Web directories.

**Translation of .htaccess permissions**
The .htaccess file sets permissions for access for the virtual Web directories on an Apache Web server. The permissions in this file are translated by the tool into the corresponding directory permissions in the NTFS file system.

**Migration of MIME types and port numbers**
MIME types define the types of *static* files that are served by the Web server. The tool automatically determines the MIME types defined on the Web server running Apache and then creates the associations of MIME types to extensions on the Web server running IIS.

**Migration of FrontPage Server Extensions**
The tool detects the presence of FrontPage 2000 Server Extensions and FrontPage 2002 Server Extensions on the Apache Web server; enumerates the Web sites and subweb sites; enumerates the FrontPage Server Extensions settings, users, and roles; creates the Web sites and subweb sites; and recalculates the links for each Web site and subweb site.

**Creation of status file**
The tool creates a file that contains the status of the migration steps completed by the tool. The status file can be used for confirming that the migration completed successfully or for troubleshooting any problems encountered while performing the migration with the tool.

The following steps in the Apache Web site migration process need to be completed after running the Apache to IIS 6.0 Migration Tool:

**Convert code within dynamic content**
Dynamic content, such as PHP or CGI, on Apache needs to be modified to run on IIS 6.0.

**Modify database connectivity within dynamic content**
For the Web sites and applications that have database connectivity, the content needs to be modified to use database connectivity provided by IIS 6.0 and Windows Server 2003.

**Migrate external modules used by Web sites**
After migration, you need to migrate external modules that are used by the Web sites and applications.

**Add Web service extensions used by the Web sites**
Any dynamic content types that are not automatically migrated by the tool must be added as Web service extensions. These new Web service extensions might be necessary for any modifications or conversions of dynamic code.

**Configure advanced Web site properties**

After the tool completes its migration tasks, the Web sites are configured comparably to how they were on Apache. You might want to configure the Web site to take advantages of the advanced security and availability features in IIS 6.0.

**Obtain server certificates for SSL**

For each Web site on the Apache Web server that uses SSL, you must either renew the existing server certificates in a format that is compatible with IIS 6.0 or you must request new server certificates in a format that is compatible with IIS 6.0. This is because the server certificates on the Apache Web server are in a format that is incompatible with IIS. Once you have the server certificates, you must import the certificates into IIS and configure your Web sites to use the appropriate certificates.

# Selecting a Migration Tool Installation Option

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Apache to IIS 6.0 Migration Tool supports four options you can choose from when determining how you want to migrate Apache Web sites to IIS. Your familiarity with the Linux or Windows operating systems and the access that you have to the Web servers involved in the migration will determine the process you should follow to install and then run the tool. In some cases, you might not have unrestricted access to the Web servers because another group manages the Apache servers or the Web servers are behind firewalls that prevent FTP traffic.

To install and run the Apache to IIS 6.0 Migration Tool for any of the installation options, you must be able to:

- Log on to the source server with an account that has root-level or super-user access.
- Log on to the target server with an account that is a member of the local Administrators group.

In addition, when you select the intermediate computer installation option, you must be able to:

- Log on to the intermediate computer running Linux with an account that has root-level or super-user access.
- Log on to the intermediate computer running Windows with an account that is a member of the local Administrators group.

Table 7.2 lists and describes the installation options for the Apache to IIS 6.0 Migration Tool. Read this table to select the installation option that best fits your needs.

**Table 7.2 Apache to IIS 6.0 Migration Tool Installation Options**

| Installation Option | Installation Details |
|---|---|
| The source server (See Figure 7.3) | • Install and run the Apache to IIS 6.0 Migration Tool on your Linux-based source server (you do not need to physically access the Windows-based target server).<br>• Install Perl on both the Linux-based source server and the Windows-based target server.<br>• Enable the FTP service on the Windows-based target server only. |
| The target server (See Figure 7.4) | • Install and run the migration Apache to IIS 6.0 Migration Tool on the Windows-based target server (you do not need to physically access the Linux-based source server).<br>• Install Perl on the Windows-based target server only.<br>• Enable the FTP service on the Linux-based source server only. |
| An intermediate computer running Microsoft Windows 2000 or later (See Figure 7.5) | • Install and run the Apache to IIS 6.0 Migration Tool on a Windows-based intermediate computer (you do not need to physically access either the Linux-based source server or the Windows-based target server).<br>• Install Perl on both the Windows-based intermediate computer and the Windows-based target server.<br>• Enable the FTP service on both the Windows-based target server and the Linux-based source server. |

| | |
|---|---|
| An intermediate computer running a supported version of Linux (See Figure 7.5) | • Install and run the Apache to IIS 6.0 Migration Tool on a Linux-based intermediate computer (you do not need to physically access either the Linux-based source server or the Windows-based target server).<br>• Install Perl on both the Linux-based intermediate computer and the Windows-based target server.<br>• Enable the FTP service on both the Windows-based target server and the Linux-based source server. |

The Apache to IIS 6.0 Migration Tool must be installed on the source server, target server, or intermediate computer. The high-level process for installing the tool includes the following:

1. Copying the compressed distribution files (.tar for Linux or .zip for Windows) to the appropriate computer based on the installation option selected.

2. Expanding the compressed distribution files.

3. Installing the Perl interpreter, if it is not already installed.

For more detailed information about the installation option that you have selected for the Apache to IIS 6.0 Migration Tool, see Installing the Migration Tool later in this section.

# Source Server Installation Option

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you select the source server installation option, you install and run the Apache to IIS 6.0 Migration Tool on the source server, which runs Apache. In addition, you must install and configure the FTP service on the target server.

Figure 7.3 illustrates the migration process for the source server installation option.

**Figure 7.3 Migrating from Apache to IIS 6.0 with the Source Server Installation Option**




The process for migrating Apache Web sites when you select the source server installation option is as follows:

1. The Apache to IIS 6.0 Migration Tool prompts you for the appropriate user credentials for the source server and target servers.

2. The tool then transfers the Apache-based Web site content and the configuration information from the source server to the target server by way of FTP.

3. The tool applies the Web site configuration settings to the IIS Web sites on the target server.

# Target Server Installation Option

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you select the target server installation option, you install and run the Apache to IIS 6.0 Migration Tool on the target server, which runs Windows Server 2003 and IIS 6.0. In addition, you must install and configure the FTP service on the source server.

Figure 7.4 illustrates the migration process for the target server installation option.

**Figure 7.4 Migrating from Apache to IIS 6.0 with the Target Server Installation Option**




The process for migrating Apache Web sites when you select the target server installation option is as follows:

1. The Apache to IIS 6.0 Migration Tool prompts you for the appropriate user credentials for the source and target servers.

2. The tool then transfers the Apache-based Web site content and the configuration information from the source server to the target server by way of FTP.

3. The tool applies the Web site configuration settings to the IIS 6.0 Web sites on the target server.

# Intermediate Computer Installation Option

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you select the intermediate computer installation option, you install and run the Apache to IIS 6.0 Migration Tool on the intermediate computer, which runs either Windows or Linux. In addition, you must install and configure the FTP service on the source and target servers.

Figure 7.5 illustrates the migration process for the intermediate computer installation option.

**Figure 7.5 Migrating from Apache to IIS 6.0 with the Intermediate Computer Installation Option**




The process for migrating Apache Web sites when you select the intermediate computer installation option is as follows:

1. The Apache to IIS 6.0 Migration Tool prompts you for the appropriate user credentials for the source and target servers.

2. The tool transfers the Apache-based Web site content and the configuration information from the source server to the intermediate computer by way of FTP.

3. The tool applies the site configuration settings to the IIS Web sites on the target server.

# Deploying the Target IIS 6.0 Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Web server running IIS 6.0 must be installed and configured before you can migrate the Apache Web sites. In addition to installing Windows Server 2003, you must install and configure IIS 6.0 and, based on the installation option you selected, install and configure the FTP service on the target server. The tool uses FTP to transfer Web site content between the source and target servers.

**Figure 7.6 Deploying the Target Server**




If the target server is already installed and configured, you can proceed to Performing the Migration later in this section.

# Installing Windows Server 2003 Before Migrating

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The primary concern when installing Windows Server 2003 is to ensure that the security of the target server is maintained. When you install Windows Server 2003 as a dedicated Web server, the default components and services are configured to provide the lowest possible attack surface. You can further reduce the attack surface of the target server by enabling only the essential Windows Server 2003 components and services.

The migration process presented here assumes that you install Windows Server 2003 with the default options. If you use other methods for installing and configuring Windows Server 2003, such as unattended setup, the default configuration settings might be different.

| **Note** |
| --- |
| When you complete the installation of Windows Server 2003, Manage Your Server automatically starts. The migration process assumes that you close Manage Your Server, and then further configure the Web server in **Add or Remove Programs** in Control Panel. |

For more information about the essential components and services to enable in Windows Server 2003, see Enabling Only Essential Windows Server 2003 Components and Services.

# Installing IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because IIS 6.0 is not installed during the default installation of Windows Server 2003, the next step in deploying the target server is to install and configure IIS 6.0. The migration process presented here assumes that you install IIS 6.0 with the default options in **Add or Remove Programs** in Control Panel. If you use other methods for installing and configuring Windows Server 2003, such as Manage Your Server, the default configuration settings might be different.

Install and configure IIS 6.0 by completing the following steps:

1. Install IIS 6.0 with only the essential components and services.

   As with installing Windows Server 2003, the primary concern when installing and configuring IIS 6.0 is to ensure that the security of the target server is maintained. Enabling unnecessary components and services increases the attack surface of the target server. You can help ensure that the target server is secure by enabling only the necessary components and services in IIS 6.0.

   For more information about how to install IIS 6.0, see Install IIS 6.0. For more information about the necessary IIS 6.0 components and services, see Enabling Only Essential IIS Components and Services.

2. In addition to the necessary components and services recommendations in step 1, enable the File Transfer Protocol (FTP) Service.

   The FTP service migrates content to the target server when you install and run the Apache to IIS 6.0 Migration Tool on the source server or on an intermediate computer.

   For more information about how to enable the FTP service, see Configure IIS Components and Services.

3. If the source server has Web sites that are FrontPage extended, enable FrontPage 2002 Server Extensions on the target server.

   For more information about how to enable FrontPage Server Extensions, see Configure a Web Site to be FrontPage Extended.

# Configuring the FTP Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Apache to IIS 6.0 Migration Tool requires the FTP service to migrate Apache Web sites to the target server. The installation option you previously selected will determine on which computers you need to enable the FTP service. Table 7.3 lists the FTP service requirements for each installation option.

**Table 7.3 Methods for Installing and Configuring the FTP Service**

| If You Install and Run the Migration Tool On | Install and Configure FTP Service On |
|---|---|
| The source server | The target server |
| The target server | The source server |
| An intermediate computer | Both the target and source server |

**Configure the FTP service on the target server**

After you install the FTP service on the target server, no additional configuration is required. The Apache to IIS 6.0 Migration Tool works with the FTP service in the default configuration. Later in the deployment process, the tool creates a default FTP site to be used for Web site migration.

**Configure the FTP service on the source server**

Configuring the FTP service on the source server requires installing and configuring the FTP package from the Linux distribution disc. After the FTP service is installed, you must configure it to allow access to the Web sites.

For more information about how to install and configure the FTP service on the source computer, see the FTP documentation that accompanies the version of Linux that is installed on the source server.

# Performing the Migration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In the next step of the migration process, the Apache to IIS 6.0 Migration Tool copies the Web site content from the source server to the target server. Then, the tool transfers the Apache Web site settings to the corresponding IIS 6.0 metabase properties.

**Figure 7.7 Performing the Migration from Apache to IIS 6.0**



# Installing the Migration Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Now you are ready to install the Apache to IIS 6.0 Migration Tool based on the method you selected in Selecting a Migration Tool Installation Option. You can install the tool on one of the following:

- Source server, which is running Linux
- Target server, which is running Windows Server 2003
- Intermediate computer, which can be running either Linux or Windows

The installation process is dependent on the operating system. So, for example, you can use the same procedure for installing the Apache to IIS 6.0 Migration Tool on the source server or for an intermediate computer running Linux.

Install the Apache to IIS 6.0 Migration Tool by completing the following steps:

1. Install the tool on a computer running Linux. Or, install the tool on a computer running Windows.
2. Configure the target server for migration.

# Installing the Migration Tool on Computers Running Linux

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Apache to IIS 6.0 Migration Tool installation process on the source server and on an intermediate computer running Linux is identical. The versions of Linux supported by the tool are listed in Overview of Migrating Apache Web Sites to IIS 6.0.

Install the Apache to IIS 6.0 Migration Tool on the computer running Linux by completing the following steps:

1. Install the Internet Information Services 6.0 Resource Kit Tools from the Microsoft Download Center onto a computer other than the target server.

2. Install Perl on the computer running Linux.

   Because the Apache to IIS 6.0 Migration Tool is written in Perl, you must install a Perl interpreter before you can run the tool. Install Perl on the computer running Linux and verify that the Perl interpreter is configured properly before installing the tool. To obtain a free copy of Perl, see ActivePerl.

3. Create a folder to contain the Apache to IIS 6.0 Migration Tool on the computer running Linux.

4. Change to the folder created in Step 3 so that the folder is the current working folder.

Copy the file AIISMT_linux.tar.gz from the *systemroot*\Program Files\AIISMT folder on the computer in Step 1 to the folder created in Step 3.

1. At a Linux command line type **tar -zxvf AIISMT_linux.tar.gz.**

   The compressed files are extracted to the folder created in step 3.

# Installing the Migration Tool on Computers Running Windows

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Apache to IIS 6.0 Migration Tool installation process on the target server and on an intermediate computer running Windows is identical. The target server must be running Windows Server 2003. The intermediate computer can be running Microsoft Windows 2000 Professional, Microsoft Windows 2000 Server, Windows 2000 Advanced Server, or Microsoft Windows Server 2003.

Install the Apache to IIS 6.0 Migration Tool on the computer running Windows by completing the following steps:

1. Install the Internet Information Services 6.0 Resource Kit Tools from the Microsoft Download Center onto a computer other than the target server.

2. Install Perl on the computer running Windows.

   Because the Apache to IIS 6.0 Migration Tool is written in Perl, you must install a Perl interpreter before you can run the tool. Install Perl on the computer running Windows and verify that the Perl interpreter is configured properly before installing the tool. To obtain a free copy of Perl, see ActivePerl.

3. Copy the AIISMT_windows.exe from the *systemroot*\Program Files\AIISMT folder on the computer in step 1 to the desktop of the target server.

4. Double-click AIISMT_windows.exe.

   The **WinZip Self-Extractor** dialog box appears.

5. In the **Unzip to folder** box, type the path of the folder where you want to extract the tool or click **Browse**to browse to the folder.

6. Click **Unzip**.

   The WinZip Self-Extractor extracts the Apache to IIS 6.0 Migration Tool into the folder selected in step 5.

# Configuring the Target Server for Migration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You must configure the target server before you can perform the migration process. The configuration process creates virtual directories in the FTP root mapped to the physical drives of the server.

Configure the target server for migration by completing the following steps:

1. At the command line, type **cd ***installation_path*, where *installation_path* is the fully qualified path where you installed the Apache to IIS 6.0 Migration Tool, and then press ENTER.

   This changes the default folder to the folder where you installed the tool. For more information about how the Apache to IIS 6.0 Migration Tool was installed, see Installing the Migration Tool on Computers Running Windows.

2. At the command line, type **aiismt_setup**, and then press ENTER.

3. When prompted to create a virtual directory, type **y** to create a virtual directory on the disk volume you want to use for migration.

   You must select at least one disk volume and create a virtual directory.

   | **Important** |
   | --- |
   | You cannot create a virtual directory on a drive formatted with the FAT file system because access rights to directories on that drive cannot be set. |

The aiismt_setup program creates the virtual directories in the FTP site root and configures FTP to listen on TCP port 21.

# Verifying that Clients Are Not Accessing Web Sites Before Migrating

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Apache to IIS 6.0 Migration Tool requires only read access to the Web site content and configuration settings on the source server; therefore, the source server can remain online in your production environment. However, you might need to remove the source server from your production network and move it to a private network segment that has direct network connectivity to the target server when:

- The amount of files and configuration information being copied across the network can generate a high volume of traffic and slow the production network.
- Firewalls that reside between the source and target servers prevent the migration tool from performing the migration. This is most commonly caused because the migration tool uses FTP for communicating with the source and target servers, and the FTP protocol is blocked by the firewalls.
- Security-related configuration settings on the source server need to be changed to allow the migration tool to work. This can include enabling FTP on production Web servers that are directly connected to the Internet.

# Migrating Web Site Content to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The migration of your Apache Web site involves migrating both the Web site content and its configuration. Based on the installation option you selected, the Apache to IIS 6.0 Migration Tool prompts you for the information required to perform the migration of the Web site content.

To run the tool, do one of the following, depending on the operating system on the computer:

- On Linux-based computers: From the Linux shell, change the current folder to the folder where you installed the tool, and then type **perl aiismt_main.pl**.
- On Windows-based computers: At the command line, change the current folder to the folder where you installed the tool, and then type **aiismt**.

After the Apache to IIS 6.0 Migration Tool starts, it prompts you for information. Table 7.4 lists the information that the migration tool requires to migrate the Web site content.

**Table 7.4 Information Required by Migration Tool During Web Site Content Migration**

| Information Required | Do the Following |
|---|---|
| Session name | Type a:<br>• Unique session name to indicate that the tool should begin a new migration process.<br>• Previous session name if you want to resume a previous migration. Supplying a previous session name causes the tool to run in recovery mode. |
| Source server IP address | If you run the tool on the:<br>• Target server or on an intermediate computer, type the IP address of the source server.<br>• Source server, press ENTER to see the default IP address displayed. If the tool cannot detect the source server's IP address, it will prompt you to confirm that the source server is running on a supported operating system. |
| Root user password | If you run the Apache to IIS 6.0 Migration Tool on the:<br>• Target server or on an intermediate computer, type the root user password to gain access to the Linux-based source server's system files and FTP service.<br>• Source server, you do not need to type the root user password. |
| Target server IP address | If you run the Apache to IIS 6.0 Migration Tool on the:<br>• Source server or on an intermediate computer, type the IP address of the source server.<br>• Target server, type the IP address of the target server or press ENTER to use the default (LocalHost) IP address. |

| | |
|---|---|
| User name and password for the target server | If you run the Apache to IIS 6.0 Migration Tool on the:<br>• Source server or on an intermediate computer, type a user name that is a member of the local Administrators group on the target server, and the associated password.<br>• Target server, you do not need to type a user name and password. |
| Configuration file path | Type the complete path to the httpd.conf file, which is a file that contains configuration and settings for Apache on the source server. You can also press ENTER to see the default path (/Etc/Httpd/Conf/Httpd.conf). |
| Web site directory on target server | Type the path of the directory on the target server where you want the Web site content to be migrated. You must supply a path for each Web site migrated.<br>When a Web site has SSL certificates associated with it, the Apache to IIS 6.0 Migration Tool prompts you to indicate whether you want the certificates to be migrated as well. |
| Listening IP address and port number | When you want to change the listening IP address and port number, type the new values. The new IP address and port numbers are applied to all the sites you selected for migration. |
| Web site IP address | When you want to change the IP address at the site level, type the new IP address, or type to use all unassigned IP addresses. |

**Note**

If you run the Apache to IIS 6.0 Migration Tool on the source server, the target server credentials are authenticated. If you run the tool on the target server, the source server credentials are authenticated.

For more information about how to migrate Web sites from Apache to IIS 6.0, see the "Apache to IIS 6.0 Migration Tool User Guide" in the IIS 6.0 Resource Kit Tools.

If the migration tool is interrupted or fails at any point during the process, you can run it in recovery mode. For more information about the recovery process, see Recovering from an Interruption in the Migration Process.

# Migrating Web Site Configuration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After migrating the Web site content, you must migrate the Web site configuration. Before you continue in the migration process, check the Apache to IIS 6.0 Migration Tool status file, Aiismt_Status.html in the session directory, to ensure that the migration has completed successfully up to this point.

To start Web site configuration migration on the target server, at the command line, change the current folder to the folder where you installed the Apache to IIS 6.0 Migration Tool, and then type **aiismt**.

During the migration of the Web site configuration, the tool prompts you to supply the session name. The files with the configuration settings are located in a folder under the FTP root. The folder uses the name of the current migration session, which you specified when you migrated the Web site content in Migrating Web Site Content to IIS 6.0 earlier in this section.

During the Web site configuration process on the target server, the Apache to IIS 6.0 Migration Tool does the following:

- Migrates Apache Web site directives and maps them to corresponding IIS configuration settings, including security settings.
- Migrates the FrontPage Server Extensions settings, users, and roles if any of the Apache Web sites are FrontPage extended.
- Sets the security settings on the sites by assigning user rights to the directories that contain site-related content.
- Verifies that settings have been correctly applied to the Web sites on the target server when the Web site configuration migration is complete.

# Recovering from an Interruption in the Migration Process

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you received no error messages while migrating the Web site content and configuration, you can proceed to one of the following sections in this section:

- Migrating Apache-Specific Extensions if any of your Web sites contain dynamic content.
- Configuring IIS 6.0 if all of your Web sites contain static content.

If you received an error message while running the Apache to IIS 6.0 Migration Tool, do the following:

1. Determine the cause of the error and resolve the problem.
2. Restart the Apache to IIS 6.0 Migration Tool in recovery mode.

In addition to resolving migration errors, you can use the tool in recovery mode to:

- Perform another migration if the majority of the settings are the same.

  For example, you might run the tool to migrate a specific Web site on a Web server that hosts multiple Web sites. You can migrate subsequent Web sites by running the tool in recovery mode and change only the name of the Web site to be migrated. The remainder of the settings remain unchanged.

- Perform the actual migration of Web site content and configuration after a trial migration.

  You can perform a *trial migration* of the Web site and content by supplying all information except the Web sites to migrate. You can use this method to validate the majority of the migration settings. You can subsequently run the tool in recovery mode.

- Provide a scripted environment for automating the Apache to IIS 6.0 Migration Tool migration process.

  The tool creates an answer file that is used by the recovery process, which you can edit to automate the tool. The answer file has the naming convention of AIISMT_*session name*_Recovery.txt, where *session_name* is the session name for the migration.

# Determining Cause of and Resolving Errors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Apache to IIS 6.0 Migration Tool automatically creates a status file for each migration session in the session folder, located in the FTP root on the target server. The name of the session folder is identical to the migration tool session name. The status file, Aiismt_Status.html, contains information about every phase of the migration process.

If an interruption occurs while the tool is running:

1. Refer to the status file to determine the cause of the failure.

2. Resolve any errors before proceeding with the upgrade process.

# Restarting the Migration Tool in Recovery Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If the Apache to IIS 6.0 Migration Tool is interrupted or fails at any point during the process, you can run it again in recovery mode. Recovery mode allows you to pick up the migration process from the point at which the tool was interrupted, rather than starting over from the beginning. Any of the following events can cause a failure or interruption in the migration tool:

- System failure.
- Network connection terminated.
- FTP access denied.

You can restart the migration tool in recovery mode when:

- An interruption occurs during data entry.

  If the tool terminates as you are entering information about the migration (such as IP addresses or file paths) you can restart it in recovery mode. The migration tool validates your user credentials, and the migration process proceeds from the point where the tool was interrupted. You do not need to retype the information you have already entered.

- An interruption occurs during file transfer.

  If the migration tool terminates during the process of transferring files from the source server to the target server, you can restart it in recovery mode. The tool resumes migration with the file that was being transferred during the interruption.

# Migrating Apache-Specific Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If the Apache Web site contains dynamic content, requires database connectivity, or includes external modules, you must separately migrate these Apache-specific extensions to complete the migration process.

Figure 7.8 illustrates the process for migrating Apache-specific extensions to IIS 6.0.

**Figure 7.8 Migrating Apache Specific Extensions to IIS 6.0**




The Apache-specific extensions on Linux commonly include MySQL and PHP. The combination of Linux, Apache, MySQL, and PHP are often referred to as LAMP. LAMP is the application platform used for Linux systems. The migration of the Apache Web sites must include the migration of LAMP.

For more information about migrating Apache Web sites to IIS, see Migrating from Linux-Apache-MySQL-PHP. Although the white paper refers to IIS 5.0 and Windows 2000 Server, the processes and methods are applicable for IIS 6.0 and Windows Server 2003.

| **Important** |
| --- |
| Test the migration of Apache-specific extensions in your lab environment before starting the Web site migration. |

The Apache-specific extensions that you might need to migrate include:

- Dynamic content.
- Database content and connectivity.
- External modules.

# Migrating Dynamic Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Apache to IIS 6.0 Migration Tool transfers the files that make up the dynamic content to the target server during the Web site content migration. However, the dynamic content typically cannot run on IIS and requires further tailoring. Table 7.5 describes the advantages and disadvantages of the methods for migrating dynamic content.

**Table 7.5 Comparison of Methods for Migrating Dynamic Content**

| Migration Method | Advantages | Disadvantages |
|---|---|---|
| Run dynamic content on a preprocessor that ships with IIS and Windows Server 2003. | • Requires minimal migration effort, depending on the extent of changes to the dynamic content.<br>• Requires no purchase of additional preprocessor. | • Might not run as efficiently as dynamic content written in ASP or ASP.NET. |
| Run dynamic content on a preprocessor external to IIS and Windows Sever 2003. | • Requires minimal migration effort, depending on the extent of changes to the dynamic content. | • Might not run as efficiently as dynamic content written in ASP or ASP.NET.<br>• Requires purchase of an additional preprocessor. |
| Rewrite the dynamic content in ASP or ASP.NET. | • Runs the dynamic content more efficiently than the other methods.<br>• Requires no purchase of additional preprocessor. | • Requires the most migration effort because the dynamic content must be completely rewritten. |

**Important**

When you run dynamic content written in preprocessors, you need to enable the *Web service extensions* (request handlers, such as Perl and CGI, that extend IIS functionality beyond serving static pages) for the preprocessor. For more information, see Configuring Web Service Extensions in IIS 6.0.

**Run Dynamic Content on a Preprocessor that Ships with IIS**

When you select this method for migrating the dynamic content, you must have the source code to the dynamic content and you must use a preprocessor that ships with IIS 6.0 and Windows Server 2003, such as CGI or Perl. You might have to adjust the code to accommodate any differences between preprocessors on Apache and IIS and between Linux and Windows.

Ensure that the version of the preprocessor supported on IIS 6.0 and Windows Server 2003 is compatible with the version of the code in the dynamic content. You will most likely need to modify the code in the dynamic content to compensate for version and platform differences. For example, CGI scripts migrated from Apache Web sites typically require modification to run on IIS 6.0 and Windows Server 2003.

**Run Dynamic Content on a Preprocessor External to IIS**

This method for migrating the external module assumes that you have the source code to the dynamic content and you can purchase or obtain a version of the preprocessor that runs with IIS, such as PHP. You might have to adjust the code to accommodate any differences between preprocessors on Apache and IIS 6.0.

For example, if the dynamic content is written in PHP, you must obtain a version of a PHP preprocessor for IIS to install and configure the PHP preprocessor on IIS 6.0. You must then adjust the PHP code to accommodate any differences between the PHP preprocessors on Apache and IIS 6.0.

**Rewrite the Dynamic Content in ASP or ASP.NET**

When you select this method for migrating the dynamic content, you need to evaluate the function of the original dynamic content, draft a functional specification for the dynamic content, and then develop a version of the dynamic content in ASP or ASP.NET.

For more information about selecting either ASP or ASP.NET, see Active Server Pages and ASP.NET.

# Migrating Database Content and Connectivity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Database connectivity provides a standardized method of connecting to various database systems. Although the Apache to IIS 6.0 Migration Tool migrated the Web site content and configuration, any database content and database connectivity must be migrated separately. Many Apache Web sites and applications use a database server, such as MySQL, the most common database connectivity standard used in Apache Web sites. Dynamic and static content can make requests to databases through MySQL.

For more information about how to migrate database content and connectivity to IIS, see Migrating from Linux-Apache-MySQL-PHP. Although the white paper refers to IIS 5.0 and Windows 2000 Server, the processes and methods are applicable for IIS 6.0 and Windows Server 2003.

# Migrating the Database Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Apache to IIS 6.0 Migration Tool does not transfer the database content to the target computer. You must migrate the database content by other methods, whose advantages and disadvantages are described in Table 7.6.

**Table 7.6 Comparison of Methods for Migrating Database Content**

| Migration Method | Advantages | Disadvantages |
|---|---|---|
| Obtain a copy of the database that is compatible with IIS 6.0 and Windows Server 2003. | • Requires minimal migration effort.<br>• Administrators are familiar with administering the database engine. | • Might not run as efficiently as Microsoft Data Engine (MSDE) or Microsoft SQL Server™ 2003.<br>• Requires purchase of an additional preprocessor. |
| Convert the database to MSDE or SQL Server 2000. | • Might run more efficiently than the original database engine on IIS 6.0 and Windows Server 2003.<br>• Requires no purchase of additional database when using MSDE. | • Requires more extensive conversion.<br>• Administrators might not be familiar with administering MSDE or SQL Server 2000. |

**Obtain a Copy of the Database Engine that is Compatible with IIS**

Most database engines on Linux have compatible versions that are designed for Windows. The easiest method of converting the database is to run the same database engine because little or no modification to the dynamic content is required. For more information, see the Web site or product documentation for the database engine used by the source server to determine if there is a version for Windows.

**Convert the Database to MSDE or SQL Server 2000**

Microsoft SQL Server supports the importing of data stored in other databases into SQL Server 2000. For example, you can import MySQL data directly into SQL Server 2000. You can use ODBC connectivity between MySQL and MSDE or SQL Server 2000 to import data as well. Client applications, such as Microsoft Access XP, can be used to import data through ODBC to MSDE or SQL Server 2000.

For more information about importing data into SQL Server 2000, see SQL Server. You can use the same methods for migrating data into MSDE as SQL Server 2000 because MSDE is based on SQL Server 2000.

# Migrating the Database Connectivity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You need to ensure that the dynamic content migrated to the target computer by the Apache to IIS 6.0 Migration Tool can use a database connectivity supported on IIS 6.0 and Windows Server 2003. Table 7.7 describes the advantages and disadvantages of the methods for migrating database connectivity.

**Table 7.7 Comparison of Methods for Migrating Database Connectivity**

| Migration Method | Advantages | Disadvantages |
|---|---|---|
| Obtain a version of the database that runs on IIS 6.0 and Windows Server 2003. | • Requires minimal migration effort. | • Requires purchase of the database that is compatible with IIS 6.0 and Windows Server 2003. |
| Convert the dynamic content to use ADO or ODBC. | • Requires no purchase of additional database. | • Requires more extensive conversion. |

**Obtain a Version of the Database Connectivity that Runs on IIS 6.0**

The database connectivity for dynamic content is implemented on Apache as libraries that are called by the dynamic content. When you install a version of the database connectivity on IIS 6.0 and Windows Server 2003, the appropriate libraries are also installed. There should be minimal changes to the dynamic content.

| Note |
|---|
| This method assumes that you are installing a version of the database, such as MySQL, on the target server. |

**Convert the Dynamic Content to use Windows Database Connectivity**

When you use this method, the assumption is that you are connecting to MSDE or SQL Server 2000 as the database engine through ODBC or ADO. If you are installing a version of the database engine on IIS 6.0 and Windows Server 2003, use the database connectivity libraries that accompany the database engine to reduce the effort required in the migration.

For example, there are ODBC drivers for MySQL; however, if you install MySQL on the target server, there is little or no conversion if you continue to use the MySQL libraries. In this case, converting your dynamic content to use ODBC has no long-term advantage unless you are planning to migrate to another database engine.

For more information about selecting ODBC, see Open Database Connectivity (ODBC). For more information about selecting ADO, see ActiveX Data Objects.

# Migrating External Modules

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

External modules, such as Cold Fusion and Wpoison, provide extended features to Apache Web sites. The Apache to IIS 6.0 Migration Tool transfers external modules to the target computer if they are in the same directory with the other Web site content. When the external modules are stored in other directories, you must manually transfer the external modules to the target server. After the external modules are transferred to the target server, you must complete additional steps so that the external modules run on IIS and Windows Server 2003. Table 7.8 describes the advantages and disadvantages of the methods for migrating external modules.

**Table 7.8 Comparison of Methods for Migrating External Modules**

| Migration Method | Advantages | Disadvantages |
|---|---|---|
| Obtain a version of the external module that runs on IIS 6.0 and Windows Server 2003. | • Requires minimal migration effort.<br>• Source code for the external module is not required. | • Might not run as efficiently as external modules written in ASP or ASP.<br>• Requires purchase of another version of the external module.<br>• Might require minor modifications to the Web site content to accommodate version differences. |
| Convert the external modules to run on the same preprocessor as IIS 6.0 and Windows Sever 2003. | • Can require minimal migration effort, depending on the extent of changes to the code for the external module. | • Might not run as efficiently as external modules written in ASP or ASP.NET.<br>• Might require the purchase of an additional preprocessor.<br>• Requires the source code of the external module so that the module can be converted. |
| Rewrite the external module in ASP or ASP.NET. | • Runs the dynamic content produced by the original external module more efficiently.<br>• Requires no additional purchases. | • Requires the most migration effort because the external module must be completely rewritten. |

**Obtain Version of External Module that Runs on IIS 6.0 and Windows Server 2003**

When you select this method for migrating the external module, you need to purchase a version of the external module for Windows. Install and configure the external module on the target server using the documentation that accompanies the external module.

For example, if an Apache Web site uses Cold Fusion, you would purchase a version of Cold Fusion for IIS 6.0 and Windows Server 2003. After installing and configuring Cold Fusion on the target server, you might have to adjust the code, such as PHP, CGI, or Perl, that calls the external module to accommodate any difference between Cold Fusion on Apache and IIS.

**Convert External Modules to Run on Same Preprocessor as IIS 6.0 and Windows Sever 2003**

This method for migrating the external module assumes you have the source code to the external modules. In addition, you must have a preprocessor on IIS that supports the code in which the external module is written, such as PHP, CGI, or Perl. You might have to adjust the code to accommodate any differences between preprocessors on Apache and IIS.

| **Important** |
| --- |
| When you run external modules written in preprocessors, you need to enable the Web service extensions for the preprocessor. For more information, see Configuring Web Service Extensions in IIS 6.0. |

**Rewrite the External Module in ASP or ASP.NET**

Select this method if no version of the external module exists for IIS 6.0 and Windows Server 2003 and you do not have the source code to the external module, or if you want to rewrite the external module to take advantage of the enhanced performance offered by ASP or ASP.NET. You need to evaluate the functions of the external module, draft a functional specification for the external module, and then develop a version of the external module in ASP or ASP.NET.

For more information about selecting either ASP or ASP.NET, see Active Server Pages and Microsoft® ASP.NET Developer Center.

| **Important** |
| --- |
| When you write ASP or ASP.NET applications, you must enable Web service extensions for ASP or ASP.NET. For more information, see Configuring Web Service Extensions in IIS 6.0. |

# Configuring IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After migrating the Web site content, configuration, and Apache-specific extensions, you must further configure IIS to allow proper operation of the Web sites. This post-migration configuration of IIS ensures that your Web sites run as they did on Apache before they were migrated to IIS 6.0.

Figure 7.9 illustrates the process for configuring IIS 6.0 after completing the migration of the Web sites from Apache.

**Figure 7.9 Configuring IIS 6.0 After Migration**



# Configuring Web Service Extensions in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Many Web sites and applications that you migrate have extended functionality beyond static pages, including the ability to generate dynamic content. Any request handlers, such as Perl and CGI, that extend IIS functionality beyond serving static pages are known as *Web service extensions*. Web service extensions allow you to enable and disable the serving of dynamic content and easily determine the content types on the Web server that can initiate code.

If you installed IIS 6.0 as described in Installing IIS 6.0, all Web service extensions are disabled by default. If you used another method, such as using Manage Your Server, the configuration of IIS might be different.

Enabling all Web service extensions ensures the highest possible compatibility with your Web sites. However, enabling all Web service extensions creates a security risk because it increases the attack surface of IIS by enabling functionality that might be unnecessary for your server.

Web service extensions allow you to enable and disable the serving of dynamic content. *MIME types* allow you to enable and disable the serving of static content. For more information about enabling and disabling the serving of static content, see Configuring File MIME Types on IIS 6.0.

| **Tip** |
| --- |
| If the appropriate Web service extension is not enabled, the Web server returns a 404 error when attempting to serve the dynamic content. |

Configure the Web service extensions by completing the following steps:

1. Enable the essential predefined Web service extensions based on the information in Table 7.9.

   Table 7.9 Predefined Web Service Extensions

   | Web Service Extension | Enable When |
   | --- | --- |
   | Active Server Pages | Enable this extension when one or more of the Web sites or applications contains ASP content. |
   | ASP.NET version 1.1.4322 | Enable this extension when one or more of the Web sites or applications contains ASP.NET content. |
   | FrontPage Server Extensions 2002 | Enable this extension when one or more of the Web sites are FrontPage extended. |
   | Internet Data Connector | Enable this extension when one or more of the Web sites or applications uses the Internet Data Connector (IDC) to display database information (content includes .idc and .idx files). |
   | Server-Side Includes | Enable this extension when one or more of the Web sites uses server-side include (SSI) directives to instruct the Web server to insert various types of content into a Web page. |
   | WebDav | Enable this extension when you want to support Web Distributed Authoring and Versioning (WebDAV) on the Web server, but it is not recommended for dedicated Web servers. |

   | **Important** |
   | --- |
   | Enable only the Web service extensions that are required by the migrated Apache Web sites. |

2. For each Web service extension used by your applications that is not one of the default Web service extensions, add a new entry to the Web service extensions list and configure the status of the new entry to **Allowed**.

For information about how to add a Web server extension to the list, see Configure Web Service Extensions.

3. Use a Web browser on a client computer to verify that the Web sites and applications run on the server.

# Configuring File MIME Types on IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS serves only static files with extensions registered in the Multipurpose Internet Mail Extensions (MIME) types list. IIS is preconfigured to recognize a default set of global MIME types, which are recognized by all configured Web sites. You can define MIME types at the Web site and directory levels independent of one another or of the types defined globally. IIS also allows you to change, remove, or configure additional MIME types. For any static content file extensions used by the Web sites hosted by IIS that are not defined in the MIME types list, you must create a corresponding MIME type entry.

The Apache to IIS 6.0 Migration Tool migrates the MIME types for the Apache Web sites; however, you must add any MIME types for content that was added earlier in the Apache migration process by completing the following steps:

1. For each static type used by your Web site, ensure that an entry exists in the MIME types list.

   When your application uses the standard MIME types included in IIS 6.0 or the MIME types migrated by the Apache to IIS 6.0 Migration Tool, no new MIME type entries are required. However, if you add new static content earlier in the Apache migration process, you must add the new MIME types. For information about adding a MIME type to the MIME types list, see Configure MIME Types.

2. Use a Web browser on a client computer to verify that the Web sites and applications run on the server.

# Configuring Web Site Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After the Apache Web sites are migrated, they run in a configuration that approximates their configuration on the Apache Web server. All of the Web sites run in the default application pool with the default security identity, NetworkService. In this configuration, the Web sites do not take advantage of the enhanced security and availability of IIS 6.0.

You can further configure the Web sites to provide the following:

- Improved security. In many cases, you can improve the security of the Web sites by isolating and running each one with unique identities. For more information about improving the security of the Web sites, see Securing Web Sites and Applications.

- Enhanced availability. In many cases, you can improve the availability of your Web sites. For more information about improving the availability of your Web sites, see Ensuring Application Availability.

# Configuring Server Certificates for SSL

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Some of the Web sites on the source server might use Secure Sockets Layer (SSL) to encrypt confidential information exchanged between the Web server and the client. For each SSL-encrypted Web site that you migrated from the source server, you need to ensure that you have a certificate for that Web site installed on the target server. You can either renew the existing certificates on the source server in a format that is compatible with IIS 6.0 or you can request new certificates. After you have installed the server certificates on the target server, you can import the certificates into the certificate store on the target server, and then configure your Web sites to use the appropriate certificates.

| **Note** |
| --- |
| *Server certificates* are installed on the Web server and typically require no additional configuration on the client servers. Server certificates allow the clients to verify the identity of the server. Alternatively, some Web sites and applications might require *client certificates*. Client certificates are installed on the client servers and allow the server to authenticate the clients. For more information about configuring client certificates, see Certificates_IIS_SP1_Ops. |

For each Web site and application that uses SSL, configure the server certificate for SSL by completing the following steps:

1. Obtain an SSL server certificate that is compatible with IIS 6.0 by doing one of the following:

   - Renew an existing certificate in a format that is compatible with IIS 6.0.

     The format of the SSL server certificate on the source server is in a format that is incompatible with IIS. You can renew an existing certificate with your certification authority in a format that is compatible with IIS 6.0.

     For more information about renewing existing server certificates in a format that is compatible with IIS 6.0, contact your certification authority.

   - Request a new certificate from a certification authority in a format that is compatible with IIS 6.0.

     You can use the Web Server Certificate Wizard either to generate a certificate request file (Certreq.txt, by default) that you send to a certification authority, or to generate a request for an online certification authority, such as Microsoft Certificate Services in Windows Server 2003. Depending on the level of identification assurance offered by your server certificate, you can expect to wait several days to several months for the certification authority to approve your request and send you a certificate file.

     For more information about requesting a server certificate by using the Web Server Certificate Wizard, see Request a Server Certificate.

2. Install the SSL server certificate to be used by the Web site on the target server.

   After you obtain an SSL server certificate, you need to install the certificate on the target server. Install the SSL server certificate on the target server by using the Certificate MMC snap-in.

   For more information about installing the SSL server certificate on the Web server, see Install a Server Certificate.

3. Assign the SSL server certificate to the Web site.

   For more information about assigning the SSL server certificate to the Web site, see Assign a Server Certificate to a Web Site.

# Backing Up the Target IIS 6.0 Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you enable client access to the target server, perform a complete image backup of the target server. Performing this image backup provides you with a point-in-time snapshot of the Web server. If you need to restore the target server in the event of a failure, you can use this backup to restore the Web server to a known configuration.

| **Important** |
| --- |
| Do not continue with the migration process unless you have a successful backup of the entire target server. Otherwise, you can lose Web sites, applications, or data that you migrated to the target server. |

For more information about how to back up the Web server, see Back Up and Restore the Web Server to a File or Tape.

# Enabling Client Access After Migration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you have migrated your Web sites from the source server to the target server, you are ready to enable client access to the Web sites on the target server, while maintaining the DNS entries to the source servers. After a period of time that meets your business needs, you can remove the DNS entries that point to the Web sites on the source server.

Figure 7.10 illustrates the final step of the process for migrating Web sites from Apache to IIS 6.0.

**Figure 7.10 Enabling Client Access After Migration to IIS 6.0**




Enable client access to the Web sites on the target server by completing the following steps:

1. Create the appropriate DNS entries for the Web sites and applications running on the target server.

   For more information about how to create DNS entries for your Web sites and applications, see "Managing resource records" in Help and Support Center for Windows Server 2003.

2. Monitor client traffic to determine whether clients are successfully accessing the target server.

   For more information about how to monitor client traffic to Web sites on the target server, see Monitor Active Web and FTP Connections.

3. Establish a monitoring period, such as a few hours or a day, to confirm that clients accessing Web sites on the target server are experiencing response times and application responses that meet or exceed your requirements.

4. Remove the DNS entries pointing to the Web sites and applications on the source server.

   For more information about how to remove DNS entries for your applications, see "Managing resource records" in Help and Support Center for Windows Server 2003.

5. Monitor client traffic to the Web sites on the source server to determine when clients are no longer accessing the source server.

   For more information about how to monitor client traffic to Web sites on the source server, see Monitor Active Web and FTP Connections.

6. When clients are no longer accessing the Web sites on the source server, decommission the hardware for the source server.

# Additional Resources About Migrating Apache Web Sites to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to migrating Apache Web sites to IIS 6.0.

**Related Information**

- Securing Web Sites and Applications
- Ensuring Application Availability
- Upgrading an IIS Server to IIS 6.0
- Migrating IIS Web Sites to IIS 6.0
- Appendix A: IIS Deployment Procedures
- Active Server Pages and Microsoft ASP.NET Developer Center for more information about selecting either ASP or ASP.NET to rewrite dynamic content.
- Migrating from Linux-Apache-MySQL-PHP for more information about migrating Apache Web sites to IIS.
- SQL Server for more information about importing data into SQL Server 2000.

**Related topics in the IIS 6.0 Operations Guide**
For more information about configuring client certificates, see Certificates_IIS_SP1_Ops.

**Related Windows Server 2003 Help topics**
For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Managing resource records" in Help and Support Center for Windows Server 2003 for more information about how to create DNS entries for your applications.

# Appendix A: IIS Deployment Procedures

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

**In This Appendix**


Assign Additional IP Addresses to a Network Adapter

Assign a Server Certificate to a Web Site

Back Up and Restore Registry Entries

Back Up and Restore the IIS Metabase

Back Up and Restore the Web Server to a File or Tape

Configure an ASP.NET Application for ASP.NET

Configure Anonymous User Identity

Configure a Web Site to be FrontPage Extended

Configure Application Identity for IIS 5.0 Isolation Mode

Configure Application Isolation Modes

Configure Application Isolation Settings for IIS 5.0 Isolation Mode

Configure Application Pool Health

Configure Application Pool Identity

Configure Application Pool Performance

Configure Application Pool Recycling

Configure FrontPage Server Roles

Configure Compression

Configure FTP Server Authentication

Configure IIS Components and Services

Configure IP Address Assigned to Web Sites

Configure IP Address and Domain Name Restrictions

Configure MIME Types

Configure NTFS Permissions

Configure the State Service on the ASP.NET State Server

Configure the Registry

Configure the Web Site Identification Number

Configure Web Server Authentication

Configure Web Service Extensions

Configure Web Site Permissions

Configure Windows Server 2003 Services

Convert Existing Disk Volumes to NTFS

Create a Service Account

Create A SQL Server Database for Storing ASP.NET Session State

Create a Virtual Directory in IIS 6.0

Create a Web Site

Debug Application Pool Failures

Determine Web Sites Uniquely Identified by IP Addresses

Disable Network Adapters

Enable ASP.NET

Enable Buffering on the Administrative Web Site

Enable Logging_IIS_SP1_Dep

Enable Network Adapters

Enable Security Auditing

Enable the WWW Service After Upgrade

Enable Web Site Content Auditing

Export a Server Certificate

Gather and Display WWW Service Uptime Data

Grant User Rights to a Service Account

Install a Server Certificate

Install IIS 6.0

Install Subauthentication

Isolate Applications in Worker Process Isolation Mode

Make a Service Account a Member of the Local Administrators Group

Migrate CDONTS

Modify the IIS Metabase Directly

Monitor Active Web and FTP Connections

Obtain a Server Certificate for an NLB Cluster

Pause Web or FTP Sites

Publish Web Site Content with FrontPage

Remove Virtual Directories

Request a Server Certificate

Secure the Root Folder of Each Disk Volume

Secure Windows Server 2003 Built-in Accounts

Set Processor Affinity

Stop the WWW Service

Upgrade FrontPage Extended Web Sites

View Application Isolation Configuration

View Web Site and Application Process Identities

# Assign Additional IP Addresses to a Network Adapter

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure a network adapter to have additional IP addresses. When you assign a specific IP address to a network adapter, you can uniquely identify a Web site or application by associating the Web site or application with that unique IP address.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc.**

**Procedures**

**To assign additional IP addresses to a network adapter**

1. In Control Panel, click **Network Connections**, right-click the network adapter that connects to the client computers, and then click **Properties**.
2. On the **General** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
3. On the **General** tab, click **Advanced**.
4. On the **Advanced TCP/IP Settings** dialog box, click **Add**.
5. On the **TCP/IP Address** dialog box, enter the **IP address** and the **Subnet mask**, and then click **Add**.
6. Click **OK** twice, and then click **Close**.

# Assign a Server Certificate to a Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Server certificates contain information about the server that allows the client to positively identify the server before sharing sensitive information. After you obtain a server certificate from a trusted certification authority and install the server certificate on the Web server, you need to assign the server certificate to the Web site.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc.**

**Procedures**
**To assign a server certificate to a Web site**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.
2. Right-click the Web site or file that you want, and then click **Properties**.

   Depending on whether you are configuring a Web site or a file, select either the **Directory Security** or **File Security** tab, and under **Secure communications**, click **Server Certificate**.
3. In the Web Server Certificate Wizard, click **Assign an existing certificate**.
4. Follow the steps in the Web Server Certificate Wizard, which guides you through the process of installing a server certificate.
5. You can view the information about the certificate by clicking the **View Certificate** button on the **Directory Security** or **File Security** tab of the Web sites **Properties** page.

# Back Up and Restore Registry Entries

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Some COM objects, COM+ objects, Internet Server API (ISAPI) extensions, and other components save configuration information in the Windows registry. If no setup program or provisioning script exists to make the appropriate registry updates during migration, you must manually identify the registry entries and then recreate them on the target server. To accomplish this, once you have identified the registry entries required by the applications currently running on the source server, back up the registry entries on the source server by using the registry editor Regedit.exe. Modify the registry entries in the backup files created by the registry editor to reflect changes in disk volume drive letters or paths to directories, and then restore the updated registry backups to the target server. Restoring the registry entries is accomplished by merging the backup files into the registry on the target server.

| **Caution** |
| --- |
| The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the registry reference. |

**Procedures**

**To back up registry entries by using the registry editor**

1. In the **Run** dialog box, type **regedit**, and then click **OK**.
2. In the registry editor, right-click the registry key or subkey that contains the registry entries that you want to back up, and then click **Export**.
3. In the **Export Registry File** dialog box, select a location for the backup registry file. Click **Save** to save the registry file to that location.
4. Close the Registry Editor.

**To restore registry entries**

1. Copy the backup registry file to any location on the target server.
2. In Windows Explorer, navigate to the backup registry file on the target server, and right-click the file name.
3. Click **Merge**. In the **Registry Editor** message box, click **Yes** to add the information in the file to the registry.
4. Click **OK** to finish.

**Related Information**

For more information about the registry, see the Registry Reference in the Windows Server 2003 Deployment Kit.

# Back Up and Restore the IIS Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Metabase backup files are copies of the metabase configuration file (MetaBase.xml) and the matching metabase schema file (MBSchema.xml). The metabase can be restored from the backup files by using the metabase configuration backup and restore feature. You can create backup files by using IIS Manager or the command-line script **iisback.vbs**, which is stored in *systemroot*\System32.

The administrator can create a *portable backup* by providing a password that is used by IIS to encrypt the secure metabase properties in the backup files. Encrypting the secure properties in this way allows the backup files to be restored to another computer by using the password. Once a backup is created, the password within the backup file cannot be changed.

When creating a *non-portable backup*, the administrator does not provide a password. Instead, the machine key is used to encrypt secure metabase properties. Because a machine key is unique to the computer it belongs to, a backup created this way can only be restored to the original computer.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc, Iisback.vbs

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc.**

**Procedures**

**To create a portable backup (password required)**

1. In IIS Manager, right-click the local computer, click **All Tasks**, and then click **Backup/Restore Configuration**.
2. Click **Create Backup**.
3. In the **Configuration backup name** box, type a name for the backup file.
4. Select the **Encrypt backup using password** check box, type a password into the **Password** box, and then type the same password in the **Confirm password** box.
5. Click **OK**, and then click **Close**.

   The IIS metabase is created in the *systemroot***\system32\inetsrv\MetaBack** folder.

**To create a non-portable backup (password not required)**

1. In IIS Manager, right-click the local computer, click **All Tasks**, and click **Backup/Restore Configuration**.
2. Click **Create Backup**.
3. In the **Configuration backup name** box, type a name for the backup file.
4. Click **OK**, and then click **Close**.

   The IIS metabase is created in the *systemroot***\system32\inetsrv\MetaBack** folder.

**To restore the metabase backup**

1. In IIS Manager, right-click the local computer, click **All Tasks**, and click **Backup/Restore Configuration**.
2. In the **Backups** list box, click the version of the **Automatic Backup** file that you want to restore, and click **Restore**. If prompted for a password, type the password you chose to secure the backup.

The IIS metabase is restored to the location from which it was backed up.

# Back Up and Restore the Web Server to a File or Tape

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You should back up a Web server before upgrading it or making configuration changes to it. A complete Web server backup includes all Web sites, applications, and data stored on the Web server. For example, before an upgrade, or before you enable client access to a target Web server, perform a complete image backup before you change any of the configuration settings on the existing Web server. The image backup provides a point-in-time snapshot of the Web server. If unforeseen problems occur during the upgrade or configuration process, you can use this backup to restore the Web server to a known configuration.

A backup file can be saved to a hard disk, a floppy disk, or to any other nonremovable or removable media on which you can save a file. Backup files usually have the extension .bkf, but you can change it to any extension you prefer.

You should back up all boot and system volumes, including the System State, when you back up the Web server.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: ntbackup.exe

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc.**

**Procedures**

### To back up to a file or tape

1. Open **Accessories**, click **System Tools**, and then click **Backup**.
2. Click the **Advanced Mode** link on the Backup or Restore Wizard.
3. Click the **Backup** tab, and then on the **Job** menu, click **New**.
4. In the **Click to select the check box for any drive, folder, or file that you want to back up** text box, click the box next to **System State** and any other items you would like to backup. You should back up all boot and system volumes along with the System State.
5. In **Backup destination**, do one of the following:
   - Choose **File** if you want to back up files and folders to a file. This is selected by default. You will not be able to choose any other option if you do not have a tape drive on your system.

     -or-

   - Click a tape device from the drop-down list if you want to back up files and folders to a tape. You cannot use this option if you do not have a tape device on your system.
6. In **Backup media or file name**, do one of the following:
   - If you are backing up files and folders to a file, type a path and file name for the backup (.bkf) file, or click the **Browse** button to find a file.
   - If you are backing up files and folders to a tape, select the tape you want to use from the drop-down list box.
7. Click **Tools**, and then click **Options** to select any backup options you want, such as the backup type and the log file type. When you have finished selecting backup options in the **Options** dialog box, click **OK**.

8. Click **Start Backup**, and then make any necessary changes to the **Backup Job Information** dialog box.
9. Click the **Advanced** tab if you want to set advanced backup options such as data verification or hardware compression. When you have finished setting advanced backup options, click **OK**.
10. Click **Start Backup** to start the backup operation.

**To restore the Web server from a file or tape**

1. Open **Accessories**, point to **System Tools**, and then click **Backup**.
2. Click the **Advanced Mode** link on the Backup or Restore Wizard.
3. Click the **Restore and Manage Media** tab, and then in **Expand the desired media item, then check the box for the items to restore**, expand the media item that contains the backup file you wish to use, and select the check box next to **System State** for that file. This will restore the System State data along with any other data you have selected for the current restore operation.
4. In **Restore files to**, do one of the following:
   - Click **Original location** if you want the backed up files and folders to be restored to the folder or folders they were in when they were backed up. Skip to step 6.

     -or-

   - Click **Alternate location** if you want the backed up files and folders to be restored to a folder that you designate. This option will preserve the folder structure of the backed up data; all folders and subfolders will appear in the alternate folder you designate.

     -or-

   - Click **Single folder** if you want the backed up files and folders to be restored to a folder that you designate. This option will not preserve the folder structure of the backed up data; the files will appear only in the folder that you designate.
5. If you selected **Alternate location** or **Single folder**, type a path for the folder under **Alternate location**, or click the **Browse** button to find the folder.
6. On the **Tools** menu, click **Options**, click the **Restore** tab, and then do one of the following:
   - Click **Do not replace the file on my computer** if you do not want the restore operation to copy over files that are already on your hard disk.

     -or-

   - Click **Replace the file on disk only if the file on disk is older** if you want the restore operation to replace older files on your disk with newer files from your backup.

     -or-

   - Click **Always replace the file on my computer** if you want the restore operation to replace files on your disk regardless of whether the backup files are newer or older.
7. Click **OK** to accept the restore options you have set.
8. Click **Start Restore**.
9. Click the **Advanced** tab if you want to change any advanced restore options, such as restoring security settings and junction point data. When you are done setting advanced restore options, click **OK**.
10. Click **OK** to start the restore operation.

# Configure an ASP.NET Application for ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By default, when the Microsoft® .NET Framework is installed on a computer with an existing .NET Framework installation, most Microsoft ASP.NET applications are automatically updated to use the newly installed version. The exceptions are applications that are bound to an incompatible or later version of the .NET Framework. Although later versions of the .NET Framework are designed to be backward compatible, you might need to configure an ASP.NET application to use an earlier version, if it cannot run successfully on a later version of the .NET Framework.

When multiple versions of the .NET Framework are executing side-by-side on a single computer, the ASP.NET ISAPI version script that is mapped to an ASP.NET application determines which version of the common language .NET Framework is used for the application. The following sections describe the process for configuring an ASP.NET application to target a specific version of the .NET Framework.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type:

**runas /user:** *administrative_accountname* **"mmc %systemroot%\system32\inetsrv\iis.msc".**

**Procedures**

**Viewing the Script Map for an ASP.NET Application**

When managing a computer with multiple versions of the .NET Framework installed, it is often useful to view the script map for an ASP.NET application to determine which version is used by the application.

**To view the script map for an ASP.NET application**

1. Open IIS Manager, expand the local computer, and navigate to the Web site folder that contains the ASP.NET application.
2. Right-click the Web site folder, and then click **Properties**.
3. In the **Properties** dialog box, on the **Home Directory** tab, click the **Configuration** button.
4. In the **Application Configuration** dialog box, on the **Mappings** tab, select an ASP.NET application extension, such as .asmx or .aspx.

   The **Executable Path** column of the dialog box lists the path to the ASP.NET ISAPI version used by the selected application type. By default, the ASP.NET ISAPI is installed in the following location: *systemroot*\Microsoft.NET\Framework*versionNumber*

   The version number shown in the path indicates the version number of the ASP.NET ISAPI used by the application type. The first two digits of the ASP.NET ISAPI version number are the same as the first two digits of the version number of the .NET Framework used by the application.

**Configuring an ASP.NET Application for ASP.NET**

When the .NET Framework is installed on a computer with an existing .NET Framework installation, by default all ASP.NET application script maps are automatically updated to use the newly installed version. To make it easier to reconfigure the script map for an ASP.NET application, each installation of the .NET Framework comes with an associated version of the ASP.NET IIS Registration tool (Aspnet_regiis.exe). Administrators can use Aspnet_regiis.exe to remap an ASP.NET application to the correct version of ASP.NET.

| **Note** |
| --- |
| Because Aspnet_regiis.exe is linked to a specific version of the .NET Framework, administrators must use the appropriate version of Aspnet_regiis.exe to reconfigure the script map for an ASP.NET application. |

Aspnet_regiis.exe can also be used to display the status of all installed versions of ASP.NET, register the associated version of ASP.NET, create client-script directories, and perform other configuration operations.

**To use Aspnet_regiis.exe to update a script map for an ASP.NET application**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
2. At the command prompt, use the **cd** command to change the directory of the Aspnet_regiis.exe version you want to use. Remember that each version of the .NET Framework comes with its own version of Aspnet_regiis.exe. By default, Aspnet_regiis.exe is located in the following directory: *systemroot*\Microsoft.NET\Framework*versionNumber*

   You can check to see which version you are running by using the procedure "To view the script map for an ASP.NET application" earlier in this section.
3. Run **Aspnet_regiis.exe** using the **-s** or **-sn** option along with the path to the application to set up the script maps. The following command-line example updates the script maps for an application called SampleApp1.

   **Aspnet_regiis.exe -s W3SVC/1/ROOT/SampleApp1**

| **Note** |
| --- |
| The **-s** and **-sn** options allow you to install script maps for ASP.NET at the specified path recursively (-s), which copies the script maps to the subfolders, or nonrecursively (-sn), which does not copy the script maps to the subfolders. |

**To get a list of the options available for Aspnet_regiis.exe**

- At the command prompt in the folder where Aspnet_regiis.exe is installed, type: **Aspnet_regiis ?**.

  A list of options available for use with Aspnet_regiis.exe displays.

# Configure Anonymous User Identity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When the Web sites and applications running on the Web server require anonymous access, IIS must be configured with a user account specifically for anonymous access. This user account can be stored in the local account database on the Web server or in a domain.

If, prior to upgrade, the anonymous user identity is configured to use an account stored in a domain you must configure the anonymous user identity to use the same domain-based account after the upgrade. This is necessary because the upgrade process automatically configures IIS to use the default anonymous user account IUSR_*computername*, where *computername* is the name of the computer on which IIS is running. You can configure the anonymous user identity to use the domain-based user account through IIS Manager. If the anonymous user identity does not use an account stored in a domain, it can be configured to use an account stored in the local account database on the Web server.

**Procedures**

**To configure the account used for Anonymous authentication**

1. In IIS Manager, expand the local computer, right-click the site, directory, or file that you want to configure, and then click **Properties**.
2. Click the **Directory Security** or **File Security** tab, depending on the level for which you are changing the security settings.
3. In the **Authentication and access control** section, click **Edit**.
4. Select the **Enable anonymous access** check box.
5. Type the valid Windows user account you want to use for Anonymous access, or click **Browse** to locate it.
6. Click **OK** twice.

# Configure a Web Site to be FrontPage Extended

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage 2002 Server Extensions from Microsoft provides Web-based and command-line administration for extending Web sites. *Extending* a Web site means enabling various FrontPage 2002 Server Extensions features to improve how you manage the content development and security of your site. Extending Web sites with FrontPage 2002 Server Extensions enables the Web site owner to author the site in the Microsoft® FrontPage® web site creation and management tool and delegate Web site ownership and administration credentials. After a Web site is extended, a FrontPage client can open the Web site and author it in FrontPage.

Before you can use FrontPage 2002 Server Extensions, you must create one or more Web sites to contain them. An extended Web site, formerly known as a virtual server, is an IIS Web site that receives additional functionality when it is extended by FrontPage Server Extensions. By default, IIS provides a working extended Web site called **Default Web Site**. This extended Web site points to the content directory *systemroot*\Inetpub\wwwroot.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc.**

**Procedures**
Before extending a site, check to make sure that the FrontPage 2002 Server Extensions optional component is installed. In Control Panel, click **Add or Remove Programs**, and then click **Windows Components Wizard**. If the FrontPage 2002 Server Extensions component is not installed, add it by using the following procedure.

**To add the FrontPage 2002 Server Extensions component**

1. In the **Windows Components Wizard** select **Application Server**, and then click **Details**.
2. In the **Application Server** subcomponents list, select **Internet Information Services**, and then click **Details**.
3. In the IIS subcomponents list, select the check box for **FrontPage 2002 Server Extensions**, and then click **OK** to complete the Windows Components Wizard and install FrontPage 2002 Server Extensions. You might need your Windows Server 2003 installation CD to complete the installation.

**To extend a site by using the FrontPage Server Extensions 2002 Server Administration tool**

1. Open Administrative Tools, and then click **Microsoft SharePoint Administrator**.
2. If the FrontPage 2002 Server Extensions optional component is not installed IIS returns error 404, "The page cannot be found." In this case, add the FrontPage 2002 Server Extensions component by using the procedure described earlier in this section.
3. On the **FrontPage Server Extensions 2002 Server Administration** site, verify that the extended Web site you created is listed in the **Virtual Servers section**.
4. To extend the Web site, click the **Extend** link next to the extended Web site name.
5. Click **Submit**. The **FrontPage Server Extensions 2002 Server Administration** tool adds FrontPage 2002 Server Extensions template directories to the content directory of your extended Web site, and adds other files that contain metadata.

# Configure Application Identity for IIS 5.0 Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you are running Internet Information Services (IIS) 6.0 in IIS 5.0 isolation mode, you also need to configure application identities for IIS 5.0 isolation mode. On your Microsoft® Windows® Server 2003–based server running in IIS 5.0 isolation mode, configure the application identity to the same settings as were set up on your server running IIS 5.0. You can do this by modifying the identity property of the COM+ object created for the Web site.

**Procedures**

**To configure Web site and application process identities in IIS 5.0**

1. In the **Run** dialog box, type **mmc**, and then click **OK** to start the Microsoft Management Console (MMC).
2. If the Component Services Snap-in is not already loaded, on the **File** menu, click **Add/Remove Snap-in**, and then click **Add**. Under **Available Standalone Snap-ins**, click **Component Services**, and then click **Add**. Click **Close** and then click **OK**.
3. In MMC, in the console tree, expand **Component Services**, expand **Computers**, expand **My Computer**, and then expand **COM+ Applications**.
4. Right-click **IIS Out-of-Process Pooled Applications** and then click **Properties**.
5. Click the **Identity** tab.
6. In **User**, type the user name. In the **Password** and **Confirm password** text boxes, enter the password for the user account. To look up a valid user name, click **Browse**.
7. Click **OK**.

# Configure Application Isolation Modes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can run Internet Information Services (IIS) 6.0 in one of two modes: *worker process isolation mode* or *IIS 5.0 isolation mode*. Worker process isolation mode is the default mode the Web server runs in after a clean installation of the Microsoft® Windows® Server 2003 operating system. IIS 5.0 isolation mode is the default mode when you upgrade to IIS 6.0 from a previous version of IIS. Administrators should only run IIS in IIS 5.0 isolation mode when applications on the server are determined to be incompatible with worker process isolation mode.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: lis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname***mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**Configuring IIS for Worker Process Isolation Mode**
To complete the following procedure, you must restart IIS, which will temporarily interrupt the World Wide Web Publishing Service (WWW service).

### To configure IIS for worker process isolation mode

1. In IIS Manager, expand the local computer, right-click **Web Sites**, and then click **Properties**.
2. Click the **Service** tab, clear the **Run WWW service in IIS 5.0 isolation mode** check box, and then click **OK**.
3. To start the WWW service, click **Yes** when asked if you want to restart IIS now.

If the switch to worker process isolation mode is successful, a folder named Application Pools appears in the IIS Manager listing for your local computer. You can quickly determine which isolation mode IIS is running because the Application Pools folder is present in worker process isolation mode and absent in IIS 5.0 isolation mode.

**Configuring IIS for IIS 5.0 Isolation Mode**
After you complete the following procedure, you must restart the WWW service, which will temporarily interrupt the service.

### To configure IIS for IIS 5.0 isolation mode

1. In IIS Manager, expand the local computer, right-click **Web Sites**, and then click **Properties**.
2. Click the **Service** tab, select the **Run WWW service in IIS 5.0 isolation mode** check box, and then click **OK**.
3. To start the WWW Service, click **Yes** when asked if you want to restart IIS now.

# Configure Application Isolation Settings for IIS 5.0 Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Isolating* applications means configuring them to run in a process (memory space) that is separate from the Web server core (the core components required to run Internet Information Services (IIS), such as IISAdmin, the metabase, and so on) and other applications. You can configure applications into one of three levels of application protection:

- Low (IIS process)
- Medium (pooled)
- High (isolated)

Note that server-side includes (SSI), Internet Database Connector (IDC), and other InProcessISAPIApps applications (special applications that must be run in process) cannot be run in medium or high isolation.

**Requirements**

- Mode: This feature of IIS 6.0 is available only when IIS is running in IIS 5.0 isolation mode.
- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc.**

**Procedures**
**To set or change the level of application protection**

1. In IIS Manager, expand the local computer, right-click the Web site or the starting-point directory for the application you want to configure, and then click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab, depending on whether you are configuring a Web site, a virtual directory, or an application.
3. In the **Application protection** box, click the appropriate level of protection, and then click **OK**.

The Web server finishes processing any current requests for the application before it creates a separate process. At the next request for the application, the application will run in the appropriate memory space.

# Configure Application Pool Health

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Monitoring the health of a worker process includes detecting whether the worker process is able to serve requests and then taking appropriate action. For example, if a worker process fails to respond to a ping request by the World Wide Web Publishing Service (WWW service), the worker process probably does not have threads available for processing incoming requests. You can monitor the health condition of an application pool by pinging the worker processes assigned to it at regular intervals. When a configurable unhealthy condition occurs, you can also set the following features to take corrective action:

- Rapid-fail protection
- Startup time limit
- Shutdown time limit

**Requirements**

- Mode: This feature of Internet Information Services (IIS) 6.0 is available only when IIS is running in worker process isolation mode.
- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

Consider the following when setting the configuration for recycling the worker processes assigned to an application pool:

- Worker process pinging is enabled by default. Adjust the ping interval to get timely information about application pool health without triggering false unhealthy conditions.
- When the WWW service detects that a worker process fails more than a set maximum number of times in a set time period, it places the worker process in a rapid-fail protection state. Placing application pools in rapid-fail protection state protects the availability of other applications on the server.
- When the WWW service detects that a worker process is unhealthy, it marks the worker process for termination. Within configured time limits, IIS terminates the worker process. It is important to tune the shutdown time period to fit the load characteristics of the applications that are being processed.
- After IIS terminates an unhealthy worker process, it starts a new worker process to replace the unhealthy one. It is important to tune the startup time period to fit the overlap desired between the unhealthy worker process and the new worker process.

**To configure worker process pinging**

1. In IIS Manager, expand the local computer, right-click the application pool you want to configure, and then click **Properties**.
2. On the **Health** tab, select the **Enable pinging** check box.
3. In the text box to the right of **Ping worker process every (frequency in seconds)**, type the number of seconds to

elapse before the worker process is pinged, and then click **OK**.

**To configure rapid-fail protection**

1. In IIS Manager, expand the local computer, right-click the application pool you want to configure, and then click **Properties**.
2. On the **Health** tab, select the **Enable rapid-fail protection** check box.
3. In the text box to the right of **Failures**, type the maximum number of failures allowed.
4. In the box to the right of **Time period (in minutes)**, type the elapsed time, in minutes, during which the number of failures is counted, and then click **OK**.

**To configure startup and shutdown time limits**

1. In IIS Manager, expand the local computer, right-click the application pool you want to configure, and then click **Properties**.
2. On the **Health** tab, under **Startup time limit**, in the text box to the right of **Worker process must startup within (time in seconds)**, type the time, in seconds, to elapse before starting up a new worker process.
3. On the **Health** tab, under **Shutdown time limit**, in the text box to the right of **Worker process must shutdown within (time in seconds)**, type the time, in seconds, to elapse before shutting down a failing worker process, and then click **OK**.

# Configure Application Pool Identity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The identity of an application pool is the name of the service account under which the application pools worker process runs. By default, application pools operate under the Network Service user account, which has low-level user access rights. You can configure application pools to run under the Local System user account, which is an account with more user rights than the Network Service or Local Service user accounts. However, be mindful that running an application pool under an account with increased user rights presents a high security risk.

**Requirements**

- Mode: This feature of Internet Information Services (IIS) 6.0 is available only when IIS is running in worker process isolation mode.
- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc.**

**Procedures**
By default, application pools operate under the Network Service user account, which has low-level user access rights. Consequently, this account provides better security against attackers or malicious users who might attempt to take over the computer on which the World Wide Web Publishing Service (WWW service) is running. The Local Service user account has low access rights as well, and is useful for situations that do not require access to resources on remote computers. You can, however, configure application pools to run under the Local System user account, which is an account with more user rights than the Network Service or Local Service user accounts.

### To change the service account that an application pool runs under

1. In IIS Manager, expand the local computer, and expand **Application Pools**.
2. Right-click the application pool you want to configure, and then click **Properties**.
3. Click the **Identity** tab, and click either **Predefined** or **Configurable**.

   **Predefined** refers to standard service accounts, such as Network Service (the default), which has low-level user access rights that can be used for access to resources on remote computers, Local Service, which has low-level access rights, and is used for situations that do not require access to resources on remote computers, or, Local System, which is an account with more user rights than the Network Service or Local Service account.

   **Configurable** refers to registered user names.
4. If you click **Predefined**, click a predefined account in the list box.
5. If you click **Configurable**, in the **User name** and **Password** boxes, type the user name and password of the account under which you want the worker process to operate.
6. Click **OK**.

# Configure Application Pool Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In worker process isolation mode, Internet Information Services (IIS) 6.0 can be configured to optimize the performance of an application pool, allowing you optimize the performance of your Web applications. Some of the ways to accomplish this include configuring an application pool to limit the size of its request queue, enabling CPU monitoring to allow the server to take action when CPU usage exceeds maximum CPU use, configuring the server to shut down a worker process after being idle for a specified number of minutes, and creating a Web garden, which is an application pools with more than one worker process assigned. Setting these configurations helps you manage the resources on a Web server.

**Requirements**

- Mode: This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode.
- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
When performance is degraded by too many requests or by problem applications, and you want to configure the application pool to improve the performance of the worker processes assigned to the pool, consider the following:

- When an application pool receives requests faster than it can handle them, the requests might consume all of the application pools memory. This could happen when the requests queue size limit is large and legitimate requests are coming in at a rapid rate, or during a denial of service attack. To prevent requests from consuming all of an application pools memory, limit the size of the request queue for the application pool.
- When problem applications consume excessive CPU resources, you need a way to manage them. In worker process isolation mode, you can enable CPU monitoring, configure CPU limits, and direct IIS to take action when a worker process exceeds the CPU limit set for it.
- Web gardens are application pools with more than one worker process assigned. Create a Web garden to achieve more robust performance when resources often get tied up with a single worker process, or when you need to smooth out the handling of a workload.

### To configure idle timeout

1. In IIS Manager, expand the local computer, right-click the application pool you want to configure, and then click **Properties**.
2. On the **Performance** tab, under **Idle Timeout**, select the **Shutdown worker process after being idle for (time in minutes)** check box.
3. In the text box to the right of **Shutdown worker process after being idle for (time in minutes)**, type the number of minutes to elapse before the worker process is recycled, and then click **OK**.

### To configure a request queue limit

1. In IIS Manager, expand the local computer, right-click the application pool you want to configure, and then click **Properties**.
2. On the **Performance** tab, select the **Limit the kernel request queue (number of requests)** check box.

3. In the text box to the right of **Limit the kernel request queue (number of requests)**, type the maximum number of requests to allow in the request queue, and then click **OK**.

**To configure CPU monitoring**

1. In IIS Manager, expand the local computer, right-click the application pool you want to configure, and then click **Properties**.
2. On the **Performance** tab, select the **Enable CPU monitoring** check box.
3. In the box to the right of **Maximum CPU use (percentage)**, type the percent maximum CPU threshold.
4. In the text box to the right of **Refresh CPU usage numbers (in minutes)**, type the number of minutes before the CPU usage numbers are refreshed.
5. In the drop-down list box below **Action performed when CPU usage exceeds maximum CPU use**, select the action to take, and then click **OK**.

**To configure a Web garden**

1. In IIS Manager, expand the local computer, right-click the application pool you want to configure, and then click **Properties**.
2. On the **Performance** tab, under **Web garden**, in the **Maximum number of worker processes** textbox, type the number of worker processes to assign to the application pool, and then click **OK**.

# Configure Application Pool Recycling

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) can be configured to periodically restart worker processes assigned to an application pool, which recycles faulty Web applications. Recycling keeps problematic applications running smoothly, especially when it is not feasible to modify the application code. Recycling helps ensure that problematic applications do not cause other applications to fail, and that system resources can be recovered from unhealthy applications. Use the Configure Application Pool Health procedure earlier in this appendix to configure worker process shutdown and startup times to ensure that applications do not experience downtime.

**Requirements**

- Mode: This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode.
- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To recycle a worker process immediately on demand**

1. In IIS Manager, expand the local computer.
2. Expand **Application Pools**.
3. Right-click the appropriate application pool you want to recycle, and then click **Recycle**.

**To configure a worker process to be recycled after a set elapsed time**

1. In IIS Manager, expand the local computer, and expand **Application Pools**.
2. Right-click the application pool you want to configure, and then click **Properties**.
3. On the Recycling tab, select the **Recycle worker processes (in minutes)** check box.
4. In the text box to the right of the **Recycle worker processes (in minutes)**, type the number of minutes that you want to elapse before the worker process is recycled, and then click **OK**.

**To configure a worker process to be recycled after a set number of processing requests**

1. In IIS Manager, expand the local computer, and expand **Application Pools**
2. Right-click the application pool you want to configure, and then click **Properties**.
3. On the **Recycling** tab, select the **Recycle worker process (number of requests)** check box.
4. In the text box to the right of the **Recycle worker process (number of requests)**, type the number of requests to be processed before the worker process is recycled, and then click **OK**.

**To configure a worker process to be recycled at scheduled times**

1. In IIS Manager, expand the local computer, expand **Application Pools**.
2. Right-click the application pool you want to configure, and then click Properties.
3. On the Recycling tab, select the Recycle worker processes at the following times check box.

4. Click **Add to add a time to the list, Remove to delete a time**, or **Edit to change an existing time when the worker process is recycled**, and then click **OK**.

| **Note** |
| --- |
| When recycling is set to occur at scheduled times, it may occur off-schedule if the system time is manually altered. To avoid unintended changes in the time that recycling occurs, recycle the scheduled worker processes soon after the system time is changed. |

**To configure a worker process to be recycled after consuming a set amount of memory**

1. In IIS Manager, expand the local computer, and expand **Application Pools**.
2. Right-click the application pool you want to configure, and then click **Properties**.
3. On the Recycling tab, under Memory recycling, select the Maximum virtual memory (in megabytes) check box.
4. In the text box to the right of **Maximum virtual memory (in megabytes)**, type the maximum amount of virtual memory allowed before the worker process is recycled.
5. Select the Maximum used memory (in megabytes) check box.
6. In the text box to the right of **Maximum used memory (in megabytes)**, type the maximum amount of memory allowed before the worker process is recycled, and then click **OK**.

# Configure Compression

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To more efficiently use available bandwidth, it is advisable to enable IIS's HTTP compression feature. HTTP compression provides faster transmission time between compression-enabled browsers and IIS regardless of whether your content is served from local storage or a UNC resource. You can compress static files only, application response files only, or both static files and application response files. Compressing application response files is usually called dynamic compression.

By default, enabling static file compression compresses .htm, .html, and .txt files. Enabling dynamic compression compresses .asp, .dll, and .exe files. You can customize the file types that are compressed. See Customizing the File Types IIS Compresses for more information.

Enabling compression of static files in IIS always results in more efficient utilization of bandwidth and faster Web site performance. Enabling dynamic compression always results in more efficient utilization of bandwidth, but if your server's processor utilization is already extremely high, the CPU load imposed by dynamic compression might make your site perform more slowly.

If you already have a network-based compression solution installed, enabling compression in IIS might not result in bandwidth utilization or performance improvements, depending on whether the network-based compression and IIS-based compression are redundant. See Optimizing IIS 6.0 Performance for detailed information about the many compression configuration options available.

Requirements

- Mode: This feature of IIS 6.0 is available whether IIS is running in worker process isolation mode or IIS 5 compatibility mode.
- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

Procedures

**To enable server-wide HTTP compression**

1. In IIS Manager, expand the local computer, right-click the Web Sites folder, and then click **Properties**.
2. Click the Service tab, and in the HTTP compression section, select the **Compress application files** check box to enable dynamic compression.
3. Select the **Compress static files** check box to compress static files.
4. In the **Temporary directory** box, type the path to a local directory or click **Browse** to locate a directory. Once a static file is compressed, it is cached in this temporary directory until it expires, or the content changes. The temporary directory must be on a local drive on an NTFS-formatted partition. The directory cannot be compressed, and should not be shared.
5. Under **Maximum temporary directory size**, click a folder size option. If you click the **Limited to (in megabytes)** option and enter a number in the text box next to it, IIS automatically cleans up the temporary directory according to a least recently used rule when the set limit is reached.
6. Click **Apply**, and then click **OK**.

| **Note** |
| --- |
| If you notice a performance loss specific to serving dynamic content, disable compression of dynamic content. |

# Configure FrontPage Server Roles

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The configuration of Internet Information Services (IIS) and the Web sites on the source server might reference user accounts that are stored in the local account database on the source server. These accounts stored locally on the Web server are known as *local user accounts*. Local user accounts are only valid on the Web server where they exist, not on other Web servers. When you migrate your Web site to another server, these local user accounts must be recreated on the target server. Once the user accounts have been created, the roles that were assigned to the user accounts on the source server must be assigned to the user accounts on the target server.

You can manage roles from the Site Administration page for your Web site. On this page you can view a list of roles, change which rights are included in a role, add a new role, and delete a role.

**Procedures**

**To add a user account and assign FrontPage server roles to it**

1. Open **Administrative Tools**, and click **Microsoft SharePoint Administrator**.
2. On the **Server Administration** page, click the name of the extended Web site for which you want to assign user roles.
3. On the **Site Administration** page for the Web site, click **Manage users**.
4. On the **Manage Users** page, click **Add a user**.
5. On the **Add a User** page, in the **User** section, click **Add user or group name (For example, DOMAIN\name)**, and enter a user name in the format *LocalComputerName**UserAccountName*.
6. In the **User Role** section, select the check boxes for all roles that apply to this user account, and then click **Add User**.

**To assign FrontPage server roles to an existing user account**

1. Open **Administrative Tools**, and click **Microsoft SharePoint Administrator**.
2. On the **Server Administration** page, click the name of the extended Web site for which you want to assign user roles.
3. On the **Site Administration** page for the Web site, click **Manage users**.
4. On the **Manage Users** page, click the name of the user for which you need to change the roles.
5. On the **Edit User Role Membership** page, next to **User Role**, select the check box for every role that applies to this user, and then click **Submit**.

# Configure FTP Server Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) supports the following File Transfer Protocol (FTP) authentication methods:

- Anonymous FTP authentication
- Basic FTP authentication

Available authentication settings must be set at the site level for FTP sites. FTP service is not enabled by default in IIS 6.0.

| **Important** |
| --- |
| If you change the security settings for your FTP site or virtual directory, your Web server prompts you for permission to reset the security settings for the child nodes of that site or directory. If you choose to accept these settings, the child nodes inherit the security settings from the parent site or directory. |

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**Enable Anonymous FTP Authentication**
If you select Anonymous FTP authentication to secure FTP resources, all requests for that resource are accepted without prompting the user for a user name or password. For Anonymous authentication, IIS automatically creates a Windows user account called IUSR_*computername*, where *computername* is the name of the server on which IIS is running. If you have both Anonymous FTP authentication and Basic FTP authentication enabled, IIS tries to use the Anonymous FTP authentication user account first.

### To enable the Anonymous FTP authentication method

1. In IIS Manager, right-click the FTP site, directory, virtual directory, or file you want to configure, and click **Properties**.
2. Click the **Security Accounts** tab.
3. Select the **Allow anonymous connections** check box.
4. To allow your users to gain access by Anonymous authentication only, select the **Allow only anonymous connections** check box.
5. In the **User name** and **Password** boxes, enter the Anonymous logon user name and password you want to use, and then click OK

   The user name is the name of the anonymous user account, which is typically designated as **IUSR**_*computername*.

   **Note** If the default **IUSR**_*computername* account will not be used for Anonymous FTP authentication, you must create a Windows user account appropriate for the authentication method. For more information about creating a new user account, see the procedure Create a Service Account in this section.
6. Set the appropriate NTFS permissions for the anonymous account.
7. For more information about setting NTFS permissions, see the procedure Configure NTFS Permissions earlier in this

appendix.

**Enable Basic FTP Authentication**

If you select the Basic FTP authentication method to secure your FTP resources, users must log on with a user name and password corresponding to a valid Windows user account. If the FTP server cannot verify a user's identity, the server returns an error message. Basic FTP authentication provides only low security because the user transmits the user name and password across the network in an unencrypted form.

**To enable the Basic FTP authentication method**

1. Create a Windows user account appropriate for the authentication method. If appropriate, add the account to a Windows user group.

   For more information about creating a new user account, see the procedure Create a Service Account earlier in this appendix.

2. Configure NTFS permissions for the directory or file for which you want to control access.

   For more information about setting NTFS permissions, see the procedure Configure NTFS Permissions earlier in this appendix.

3. In IIS Manager, right-click the FTP site, directory, virtual directory, or file you want to configure, and click **Properties**.

4. Click the **Security Accounts** tab.

5. Clear the **Allow anonymous connections** check box, and then click **OK**.

# Configure IIS Components and Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

**Add or Remove Programs** helps you manage programs and components on Microsoft® Windows® Server 2003 and Internet Information Services (IIS) 6.0. As you enable and disable the IIS protocols and services, you correspondingly increase and decrease the attack surface of the Web server.

| **Tip** |
| --- |
| You can also add Windows components by using the Configure Your Server Wizard. |

**Procedures**

**To configure IIS components and services**

1. In Control Panel, open **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. On the Windows Components Wizard page, under **Components**, click **Application Server**, and then click **Details**.
4. Click **Internet Information Services (IIS)**, and then click **Details**.
5. Enable or disable the appropriate IIS components and services by selecting the check box to add the component or clearing the check box to remove the component.
6. Complete the Windows Components Wizard by following the instructions in the wizard.

# Configure IP Address Assigned to Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure your Web site to use a unique IP address that uniquely identifies the Web site.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

### To configure the IP address assigned to a Web site

1. In IIS Manager, expand the local computer, right-click the Web site you want to configure, and then click **Properties**.
2. Click the **Web Site** tab, and then click the drop-down arrow next to **IP Address**.
3. Click the IP address you want to use from the drop-down list of IP addresses, and then click **OK**.

# Configure IP Address and Domain Name Restrictions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure your Web site to grant or deny specific computers, groups of computers, or domains access to Web sites, directories, or files.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**Configure Restrictions Based on IP Address**
You can use IIS Manager to grant or deny access to Web sites or applications for a computer or group of computers.

**Grant or deny access to resources for a single computer**
You can either deny or grant access for a single computer based upon its IP address.

**To grant access to resources for a computer**

1. In IIS Manager, expand the local computer, right-click a Web site, directory, or file you want to configure, and click **Properties**.
2. Click the **Directory Security** or **File Security** tab. In the IP address and domain name restrictions section, click **Edit**.
3. Click **Granted access**.
4. When you select **Granted access**, you grant access to all computers and domains, except to those that you specifically deny access.
5. Click **Add**.
6. Click **Single computer**.
7. Click **DNS Lookup** to search for computers or domains by name, rather than by IP address.
8. Type the DNS name for the computer. IIS searches on the current domain for the computer, and if found, enters its IP address in the IP address box.

   The following information is important to remember when using the **DNS Lookup** feature:

   - Server performance decreases while DNS addresses are being looked up.
   - A user accessing your Web server through a proxy server will appear to have the IP address of the proxy server.
   - Some user server access problems can be corrected by using the "*.domainname.com" syntax rather than the "domainname.com" syntax.

9. Click **OK** three times.

**To deny access to resources for a computer**

1. In IIS Manager, expand the local computer, right-click a Web site, directory, or file you want to configure, and click **Properties**.

2. Click the **Directory Security** or **File Security** tab. In the **IP address and domain name restrictions** section, click **Edit**.
3. Click **Denied access**.

   When you select **Denied access**, you deny access to all computers and domains, except to those that you specifically grant access.
4. Click **Add**.
5. Click **Single computer**.
6. Click **DNS Lookup** to search for computers or domains by name, rather than by IP address.
7. Type the DNS name for the computer. IIS searches on the current domain for the computer, and if found, enters its IP address in the **IP address** box.

   The following information is important to remember when using the **DNS Lookup** feature:

   - Server performance decreases while DNS addresses are being looked up.
   - A user accessing your Web server through a proxy server will appear to have the IP address of the proxy server.
   - Some user server access problems can be corrected by using the "*.domainname.com" syntax rather than the "domainname.com" syntax.
8. Click **OK** three times.

**Grant or deny access to resources for a group of computers**

A group of computers can be either denied or granted access based upon their network ID and a subnet mask. The network ID is the IP address of a host computer, usually a router for the *subnet*. The subnet mask determines which part of the IP address is a subnet ID, and which part is a host ID. All computers in a subnet have the same subnet ID, but have their own unique host ID. By specifying a network ID and a subnet mask, you can select a group of computers.

**To grant access to resources for a group of computers**

1. In IIS Manager, expand the local computer, right-click a Web site, directory, or file you want to configure, and click **Properties**.
2. Click the **Directory Security** or **File Security** tab. In the **IP address and domain name restrictions** section, click **Edit**.
3. Click **Granted access**.

   When you select **Granted access**, you grant access to all computers and domains, except to those that you specifically deny access.
4. Click **Add**.
5. Click **Group of computers**.
6. In the **Network ID** box, type the IP address of the host computer.
7. In the **Subnet mask** box, type the subnet ID for the computer you want grant or deny access to.
8. Click **OK** three times.

**To deny access to resources for a group of computers**

1. In IIS Manager, expand the local computer, right-click a Web site, directory, or file you want to configure, and click **Properties**.
2. Click the **Directory Security** or **File Security** tab. In the **IP address and domain name restrictions** section, click **Edit**.
3. Click **Denied access**.

   When you select **Denied access**, you deny access to all computers and domains, except to those that you specifically grant access.

4. Click **Add**.
5. Click **Group of computers**.
6. In the **Network ID** box, type the IP address of the host computer.
7. In the **Subnet mask** box, type the subnet ID for the computer you want grant or deny access to.
8. Click **OK** three times.

#### Configure Restrictions Based on Domain

Access to resources for a domain can be granted or denied by using IIS Manager.

**To grant access to resources for a domain**

1. In IIS Manager, expand the local computer, right-click a Web site, directory, or file you want to configure, and click **Properties**.
2. Click the **Directory Security** or **File Security** tab. In the **IP address and domain name restrictions** section, click **Edit**.
3. Click **Granted access**.

   When you select **Granted access**, you grant access for all computers and domains, except for those that you specifically deny access.
4. Click **Add**.
5. Click **Domain name**. You will see a warning message saying that "Restricting access by domain name requires a DNS reverse lookup on each connection. This is a very expensive operation and will dramatically affect server performance." Click **OK** to close the message dialog box.
6. In the **Domain name** box, type the domain name.
7. Click **OK** three times.

**To deny access to resources for a domain**

1. In IIS Manager, expand the local computer, right-click a Web site, directory, or file you want to configure, and click **Properties**.
2. Click the **Directory Security** or **File Security** tab. In the **IP address and domain name restrictions** section, click **Edit**.
3. Click **Denied access**.

   When you select **Denied access**, you deny access for all computers and domains, except for those that you specifically grant access.
4. Click **Add**.
5. Click **Domain name**.
6. In the **Domain name** box, type the domain name.
7. Click **OK** three times.

# Configure MIME Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) serves only static files with extensions registered in the Multipurpose Internet Mail Exchange (MIME) types list. IIS is preconfigured to recognize a default set of global MIME types, and also allows you to configure additional MIME types and change or remove MIME types. These MIME types are recognized by all Web sites you create in IIS.

MIME types can also be defined at the Web site and directory levels, independent of one another or the types defined globally. When you view MIME types at the Web site or directory level, only the types unique to that level are displayed, not all types inherited from the next level up. If you apply a MIME type at the global level after modifying the same MIME type at a lower level, the global-level MIME type overrides the modified MIME type at the lower level.

When IIS delivers a mail message to a mail application, or a Web page to a client Web browser, it also sends the MIME type of the data it is sending. If there is an attached or embedded file in a specific format, IIS tells the client application the MIME type of the embedded or attached file. The client application then knows how to process or display the data being received from IIS. IIS returns error 404.3 if a client request refers to a file name extension that is not defined in the MIME types.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

### To add a global MIME type

1. In IIS Manager, right-click the local computer, and click **Properties**.
2. Click the **MIME Types** button.
3. Click **New**.
4. In the **Extension** box, type the file name extension.
5. In the **MIME type** box, type a description that exactly matches the file type defined on the computer.

| **Note** |
| --- |
| You can also create a MIME type for files without an extension or for undefined MIME types. This is not recommended. |

6. Click **OK**.

### To add a MIME type to a Web site or directory

1. In IIS Manager, right-click the Web site or Web site directory for which you want to add a MIME type, and click **Properties**.
2. Click the **HTTP Headers** tab.
3. Click **Mime Types**.
4. Click **New**.
5. In the **Extension** box, type the file name extension.
6. In the **MIME type** box, type a description that exactly matches the file type defined on the computer. If you define a

MIME type that has already been defined at a higher level, you are prompted to select the level where the MIME type should reside.

7. Click **OK**.

**To remove a MIME type from a Web site or directory**

1. In IIS Manager, right-click the Web site or Web site directory from which you want to remove a MIME type, and click **Properties**.
2. Click the **HTTP Headers** tab.
3. Click **Mime Types**.
4. From the **Registered MIME types** list, click the MIME type you want to remove, and then click **Remove**.
5. Click **OK** three times.

# Configure NTFS Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use NTFS permissions to define the level of access to your directories and files that you want to grant to specific users and groups of users. Proper configuration of file and directory permissions is crucial for preventing unauthorized access to your resources.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**To secure a Web site by using NTFS permissions**

1. In IIS Manager, expand the local computer, right-click the Web site or file you want to configure, and click **Permissions**.
   - To add a group or user that does not appear in the **Group or user names** list box, click **Add**, and in the **Enter the object names to select** text box, type the name of the user or group. Click **OK**.

     -OR-

   - To change or remove permissions from an existing group or user, click the name of the group or user in the **Group or user names** list box.
2. To allow or deny a permission such as **Read & Execute**, **List Folder Contents**, **Read**, or **Write**, in the **Permissions for group or user name** list box, select the **Allow** or **Deny** check box next to the appropriate permission, and then click **OK**.

| **Important** |
| --- |
| Inherited Deny permissions do not prevent access to an object if the object has an explicit Allow permission entry. Explicit permissions take precedence over inherited permissions, including inherited Deny permissions. |

With NTFS permissions, you also have the choice of assigning special permissions to groups or users. Special permissions are permissions on a more detailed level. For better management, you should assign broad-level permissions to users or groups, where it is applicable. For descriptions of permissions, see "Permissions for files and folders" in Help and Support Center for Windows Server 2003.

**To secure a Web site using NTFS special permissions**

1. In IIS Manager, expand the local computer, right-click a Web site or file you want to configure, and click **Permissions**.
2. Click **Advanced**, and then do one of the following on the **Permissions** tab:
   - To set special permissions for an additional group or user, click **Add**, and in the **Enter the object name to select** text box, type the name of the user or group. Click **OK**.
   - To view or change special permissions for an existing group or user, click the name of the group or user, and then click **Edit**.

- To remove an existing group or user and its special permissions, click the name of the group or user and then click **Remove**. If the **Remove** button is unavailable, clear the **Allow inheritable permissions from the parent to propagate to this object and all child objects. Include these with entries exclusively defined here.** check box, and then click **Remove**. Click **OK** and skip steps 3-6 below.

3. To allow or deny a permission such as **Read & Execute**, **List Folder Contents**, **Read**, or **Write**, in the **Permissions** list box, select the **Allow** or **Deny** check box next to the appropriate permission.
4. In the **Apply onto** list box, click the folders or subfolders you want these permissions to be applied to.
5. To prevent the subfolders and files from inheriting these permissions, clear the **Apply these permissions to objects and/or containers within this container only** check box, and then click **OK** three times.

| **Important** |
| --- |
| It is recommended that you assign permissions to the highest-level folders as possible and then apply inheritance to propagate the settings to lower-level subfolders and files. For more information on inheritance, see "How inheritance affects file and folder permissions" in Help and Support Center for Windows Server 2003. |

# Configure the State Service on the ASP.NET State Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The ASP.NET state service is used to manage session state on a computer. The ASP.NET state service is installed by default when Microsoft® Windows® Server 2003 is installed. The file aspnet_state.exe is installed on the remote server that will store session state information; the default location is *systemroot***Microsoft.NET\Framework****version***aspnet_state.exe**.

**Procedures**

**To configure the ASP.NET state service**

1. On the remote server that will store session state information, open **Administrative Tools**, and then click **Services**.
2. In the details pane, right-click **ASP.NET State Service**, and then click **Properties**.
3. On the **General** tab, in the **Startup type** list box, click **Automatic**.
4. Under **Service status**, click **Start**, and then click **OK**. The state service starts automatically when the Web server is restarted.

# Configure the Registry

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If a registry entry must be created or modified to correctly configure the server, you can edit the entry directly using the registry editor Regedit.exe.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Regedit.exe.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

| ⚠Caution |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the registry reference. |

**To create a new registry entry by using the registry editor**

1. In the **Run** dialog box, type **regedit**, and then click **OK**.
2. In the registry editor, navigate to the key or subkey under which you wish to add an entry and select the name of the key or subkey by clicking on it.
3. On the **Edit** Menu, point to **New** and then click the data type for the entry, such as **String Value**, **Binary Value**, or **DWORD Value**.
4. In the details pane, type the name of the registry entry, and then press ENTER to create the entry.
5. To assign a value to the registry entry, right-click the entry and then click **Modify**. If the entry has been defined as **Binary Value**, click **Modify Binary Data** instead.
6. In the **Edit ValueType Value** dialog box, type an appropriate value in the **Value data** text box. Type or select the value of other options, such as the base (hexadecimal or decimal) for DWORD values, and then click **OK**.

**To configure an existing registry entry by using the registry editor**

1. In the **Run** dialog box, type **regedit**, and then click **OK**.
2. In the registry editor, navigate to the registry entry that you want to modify.
3. In the details pane, right-click the entry, and click **Modify**. If the entry has been defined as **Binary Value**, then click **Modify Binary Data** instead.
4. In the **Edit ValueType Value** dialog box, type an appropriate value in **Value data**. Type or select the value of other options, such as the base (hexadecimal or decimal) for DWORD values, and then click **OK**.

**Related Information**

For more information about the registry, see the Registry Reference in the Windows Server 2003 Deployment Kit.

# Configure the Web Site Identification Number

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When a new site is created on Internet Information Services (IIS) 6.0, a Web site identification number is randomly generated based on the name of the Web site. In this way, Web sites of the same name usually generate the same Web site identification for IIS 6.0 servers in a Web farm. With IIS 5.0 and earlier versions, Web site identification numbers were incremental. For example, because the default Web site is created first, its site identification number is 1, and the next site to be created is identified as 2.

If you have administrative scripts that depend upon the IIS 5.0 method of generating Web site identification numbers, you can edit the registry entry IncrementalSiteIDCreation in the Windows registry to force IIS to use the incremental method of generating Web site identification numbers.

| **Note** |
| --- |
| When you remotely administer a server running IIS, the value of the IncrementalSiteIDCreation registry entry on the local server is used to determine the generation of Web site identification numbers on the remote server. |

**Requirements**

- Credentials: Membership in the Administrators group on the local and on the remote computer is required to make changes to the remote computers registry. Network policy settings might prevent you from completing this procedure.
- Tools: Regedit.exe.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the registry reference. |

### To configure the IncrementalSiteIDCreation registry entry

1. To open the registry editor, in the **Run** dialog box, type **regedit**, and then click **OK**.
2. Expand **HKEY_LOCAL_MACHINE**, expand **SOFTWARE**, expand **Microsoft**, expand the **InetMGR** subkey, and then click **Parameters**.
3. Under **Parameters**, in the details pane, click **IncrementalSiteIDCreation**. If the entry does not exist, create it by doing the following:
4. Select **Parameters** by clicking on it.
5. On the **Edit** Menu, point to **New**, and then click **DWORD Value**.
6. Type **IncrementalSiteIDCreation** and press ENTER to create the entry.
7. With the entry **IncrementalSiteIDCreation** selected, on the **Edit** menu, click **Modify**.
8. In the **Value data** box, type **1** to force IIS to use the incremental method of generating Web site identification numbers, and then click **OK**.

**Related Information**

For more information about the registry, see the Registry Reference in the Windows Server 2003 Deployment Kit.

# Configure Web Server Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can set the authentication method for your Web resources with property sheets at the Web site, directory, or file level by using IIS Manager.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To configure Web server authentication**

1. In IIS Manager, right-click the **Web Sites** folder, Web site, directory, virtual directory, or file and click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web Sites folder level can be inherited by all Web sites |

2. Click the **Directory Security** or **File Security** tab, depending upon the level at which you are configuring security settings.
3. In **Authentication and access control**, click **Edit**.
4. To configure Integrated Windows authentication, in **Authenticated access**, select the **Integrated Windows authentication** check box.
5. To configure Digest authentication, in **Authenticated access**, select the **Digest authentication for Windows domain servers** check box.
6. To configure **Advanced Digest authentication**, in the **Realm** text box, type the realm name, or click **Select** to browse for a domain.

   | **Note** |
   | --- |
   | If Basic authentication is enabled for the site, virtual directory, or folder you are configuring, the **Default domain** text box will also be available. However, Realm is only meaningful for Advanced Digest authentication. |

   To configure Basic authentication, In the **Authenticated access** section, select the **Basic authentication (password is sent in clear text)** check box. Because Basic authentication sends passwords over the network unencrypted, a dialog box appears asking if you want to proceed. Click **Yes** to proceed. In the **Default domain** box, either type the domain name you want to use, or click **Select** to browse to a new default logon domain.

7. To configure .NET Passport authentication, select the **.NET Passport Authentication** check box. When .NET Passport authentication is selected, all other authentication methods are unavailable. .NET Passport cannot be used with other authentication methods because it validates user credentials in a fundamentally different way.
8. Click **OK** twice.

# Configure Web Service Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In order to take a more proactive stance against malicious users and attackers, Internet Information Services (IIS) is not installed on members of the Microsoft® Windows® Server 2003 family by default. Furthermore, when you initially install IIS, it is installed in a highly secure mode. By default, IIS serves only static content — features such as ASP, ASP.NET, server-side includes, WebDAV publishing, and FrontPage 2002 Server Extensions from Microsoft do not work unless they are specifically enabled. You can configure these features, also called Web service extensions, through the Web Service Extensions node in IIS Manager.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type: **runas /user:***administrative_accountname* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

### To enable and disable a Web service extension

1. In IIS Manager, expand the local computer, and then click **Web Service Extensions**.
2. In the details pane, click the Web service extension that you want to enable or disable.
3. To enable a disabled Web service extension, click **Allow**.
4. To disable an enabled Web service extension, click **Prohibit**.

   A message box with a list of applications that will be prevented from running on the IIS Web server displays.
5. Click **OK** to disable the Web service extension.

To use, or to deny the use of, an HTTP request handler that is not in the list of Web service extensions, you must first register it by adding the HTTP request handler to the list of Web service extensions.

### To add new Web service extensions

1. In IIS Manager, expand the computer name, and then click **Web Service Extensions**.
2. In the details pane, click **Add a new Web service extension**.
3. In the **Extension name** text box, type the name of the new Web service extension, and then click **Add**.
4. In the **Path to file** text box, type the path or click **Browse** to navigate to any files that the new Web service extension requires, and then click **OK**.
5. Optionally, select the **Set extension status to Allowed** check box to automatically set the status of the new Web service extension to **Allowed**.
6. Click **OK** to add the new Web service extension.

You cannot delete built-in extensions. You can delete only extensions that you have added to the Web extension list.

### To delete Web service extensions

1. In IIS Manager, expand the computer name, and then click **Web Service Extensions**.
2. In the details pane, right click the Web Service Extension you want to delete, and then click **Delete**.

3. Click **Yes** to confirm the deletion.

You can also use IIS Manager to specify the applications that are allowed to call Web service extensions.

**To allow an application to call Web service extensions**

1. In IIS Manager, expand the local computer, and then click **Web Service Extensions**.
2. In the details pane, click **Allow all Web service extensions for a specific application**.
3. From the **Application** list box, click the name of the application.

   The Web service extension that the application is allowed to call appears in the **Extensions to be allowed** box.
4. Click **OK** to allow the application to call Web service extensions.

You can disable all Web service extensions that are registered on the local computer with one setting in the Web Service Extension node of IIS Manager.

**To disable all Web service extensions**

1. In IIS Manager, expand the local computer, and then **click Web Service Extensions**.
2. Click **Prohibit all Web service extensions**. The following message appears:

   If you prohibit all extensions, all Web service extensions in the list will be prohibited. This may prevent applications from running on your IIS Web server. Do you want to prohibit all extensions?
3. To disable all extensions, click **Yes**. To cancel the action, click **No**. For more information, click **Help**.
4. If you click **Yes**, the status of each Web service extension is **Prohibited**.

# Configure Web Site Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure access permissions for specific Web sites, directories, and files. Unlike NTFS permissions, Web permissions affect everyone who tries to access your Web site. Web site permissions are not meant to be used in place of NTFS permissions, but are used in conjunction with them.

| **Note** |
| --- |
| If Web site permissions conflict with NTFS permissions for a directory or file, the more restrictive settings are applied. |

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To set permissions for Web content**

1. In IIS Manager, right-click the **Web Sites** folder, Web site, directory, virtual directory, or file you want to configure, and click **Properties**.
2. On the **Home Directory**, **Virtual Directory**, or **File** property sheet, select or clear any of the following check boxes (if available), depending on the type of access you want to grant or deny:
   - **Script Source Access**. Users can access source files. If **Read** is selected, source can be read, if **Write** is selected, source can be written to. **Script Source Access** includes the source code for scripts. This option is not available if neither **Read** nor **Write** is selected.
   - **Read** (selected by default). Users can view directory or file content and properties.
   - **Write**. Users can change directory or file content and properties.
   - **Directory browsing**. Users can view file lists and collections.
   - **Log visits**. A log entry is created for each visit to the Web site.
   - **Index this resource**. Allows Indexing Service to index this resource. This allows searches to be performed on the resource.
3. In the **Execute Permissions** list box, select the appropriate level of script execution:
   - **None** Do not run scripts or executables on the server.
   - **Scripts only** Run only scripts on the server.
   - **Scripts and Executables** Run both scripts and executables on the server.

4. Click **OK**. If child nodes for a directory have different Web site permissions configured, the **Inheritance Overrides** box appears.

   If a child node belonging to the directory whose Web site permissions you have changed has also set the Web site permissions for a particular option, the permissions in the child node will override those you have set for the directory. If you want the Web site permissions at the directory level to apply to the child nodes, you must select those child nodes in the **Inheritance Overrides** box.

5. If the **Inheritance Overrides** box appears, select the child nodes in the **Child Nodes** list to which you want the directory's Web permissions to apply. You can also click **Select All** to set the property to apply the Web permissions to all child nodes.

6. You might see more than one **Inheritance Overrides** box if more than one property has been defined in the child nodes of the directory. Select the child nodes from the **Child Nodes** list or click **Select All**, and then click **OK** to apply the Web permissions for this property to the child nodes.

# Configure Windows Server 2003 Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use Add or Remove Programs in Control Panel to add or remove Microsoft® Windows® Server 2003 services, and you can use the Services console to enable or disable services. As you add and remove Windows Server 2003 services and enable and disable services, you increase and decrease the attack surface of the Web server correspondingly.

| **Note** |
| --- |
| You can also add Windows Server2003 services by using the Configure Your Server Wizard. |

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To add or remove Windows Server 2003 services**

1. From Control Panel, click **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. On the Windows Components Wizard page, in the **Components** box, select or clear the check box next to the component to enable or disable the appropriate Windows Server 2003 services. You may need to select a higher level component and then click **Details** to locate the component you wish to add or remove.
4. Click **Next** to run the Windows Components Wizard.
5. Follow the instructions in the Windows Components Wizard to complete the addition or removal of the component.

| **Note** |
| --- |
| You may need to have your Windows Server2003 CD-ROM available to complete the installation of the component. |

**To enable or disable Windows Server 2003 services**

1. Open Administrative Tools, and then click **Services**.
2. In the details pane, right-click the Windows Server 2003 service that you want to change, and then click **Properties**.
3. On the **General** tab, in the **Startup type** list box, click one of the following:
4. **Automatic**. The service starts automatically when the Web server is restarted.
5. **Manual**. The service can be started manually by an administrator or by another service.
6. **Disabled**. The service cannot be started by an administrator or by another service unless the startup type is changed to **Automatic** or **Manual**.
7. Click **OK** to save the changes.

# Convert Existing Disk Volumes to NTFS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The command-line tool **Convert.exe** converts FAT and FAT32 volumes to the NTFS file system, leaving existing files and folder intact. Volumes converted to the NTFS file system cannot be converted back to FAT or FAT32.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Convert.exe

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

### To convert FAT and FAT32 volumes to NTFS

1. In the **Run** dialog box, type **convert** followed by the appropriate syntax, and then click **OK**.

**Syntax**

```
convert [Volume] /fs:ntfs
```

**Parameters**

`Volume`

Specify the drive letter (followed by a colon), mount point, or volume name to convert to NTFS. Required if the volume to be converted is not the current volume.

`/fs:ntfs`

Required. Converts the volume to NTFS.

# Create a Service Account

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A service account is a user account that is created explicitly to provide a security context for services running on Microsoft® Windows® Server 2003. Application pools use service accounts to assign permissions to Web sites and applications running on Internet Information Services (IIS). Administrators can manage service accounts individually to determine the level of access for each application pool in a distributed environment.

Use Active Directory Users and Computers to create service accounts in the Active Directory® directory service. Use Computer Management to create local service accounts on a local computer.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Active Directory Users and Computers; Computer Management.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**To create a service account in Active Directory**

1. Open **Administrative Tools**, and then click **Active Directory Users and Computers**.
2. In the console tree, double-click the **Domain** node.
3. In the **Details** pane, right-click the organizational unit where you want to add the service account, point to **New**, and then click **User**.
4. In **First name**, type a first name for the service account.
5. In **Last name**, type a last name for the service account.
6. Modify **Full name** as desired.
7. In **User logon name**, type the name that the service account will log on with and, from the drop-down list, click the UPN suffix that must be appended to the service account logon name (following the @ symbol). Click **Next**.
8. In **Password** and **Confirm password**, type a password for the service account.
9. Select the appropriate password options, and then click **Next**.
10. Click **Finish** to complete creating a service account.

**To create a service account on the local Web server**

1. Open **Administrative Tools**, and then click **Computer Management**.
2. In the console tree, expand **System Tools**, expand **Local Users and Groups**, and then click **Users**.
3. On the **Action** menu, click **New User**.
4. Type a **User name**, **Full name**, and a **Description** of the user account.
5. In **Password** and **Confirm password**, type a password for the user account.
6. Set the user account access by selecting the check box to set the option or clearing the check boxes to remove the option for:

- **User must change password at next logon**
- **User cannot change password**
- **Password never expires**
- **Account is disabled**

7. Click **Create**, and then click **Close**.

**To create a service account for IIS_WPG Group**

1. Open Administrative Tools, and then click **Computer Management**.
2. In the console tree, expand **System Tools**, expand **Local Users and Groups**, and click **Groups**.
3. Click the IIS_WPG group and, on the **Action** menu, click **Add to Group**.
4. Under **Description**, type the name of the account and click **Add**.
5. In the **Select Users** dialog box, click the **Object Types** button, and select or clear the check box for the object types you want to find. Click **OK**.
6. Click the **Locations** button to select the location of the service account. Click **OK**.
7. Enter the name of the object under **Enter the object names to select**.
8. Click **OK**, and then click **OK** again.

# Create A SQL Server Database for Storing ASP.NET Session State

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP.NET SQL state server is used to manage session state on a computer running Microsoft® SQL Server™. All versions of Microsoft ASP.NET that are installed on the same computer share the same SQL state server. The SQL state server version of session state that is used is always the one that is installed with the latest version of ASP.NET. When this version of ASP.NET is uninstalled, the latest remaining version on the computer is then registered and used in its place.

**Requirements**

- Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Create an ASP.NET session state database with SQL Server Enterprise Manager**

1. On the **Start** menu, point to **Programs**, point to **Microsoft SQL Server**, and then click **Enterprise Manager**.
2. In **SQL Server Enterprise Manager**, connect to the server running SQL Server that will store the session state.
3. On the **Tools** menu, click **SQL Query Analyzer**.
4. In SQL Query Analyzer, on the **File** menu, click **Open**, and navigate to **InstallSqlState.sql** (the SQL script that builds the ASP.NET session state database).

   The InstallSqlState.sql file is located on the Web server in *systemroot***\Microsoft.NET\Framework***version* (where *version* is the most recent version number of the .NET Framework installed on the Web server).
5. On the **Query** menu, click **Execute**.
6. Ensure the query completed with no errors by reviewing the status in the lower window of the SQL Query Analyzer.
7. Close SQL Query Analyzer.
8. In **SQL Server Enterprise Manager**, in the console tree, expand the server node, and then click **Databases**.
9. In the details pane, verify that **ASPState** is listed.
10. Close SQL Server Enterprise Manager.

# Create a Virtual Directory in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In most cases, the content you publish to your Web site is located in a home directory on your computer, such as C:\Inetpub\Wwwroot\. However, there might be instances when the content is located in another directory, or even on a remote computer.

To be able to publish content from any directory not contained within your home directory, you must create a *virtual directory*. A virtual directory is a directory that is not contained in the home directory but appears to client browsers as though it were. You can create a virtual directory by using IIS Manager or by using Windows Explorer.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**
**To create a virtual directory by using IIS Manager**

1. In IIS Manager, expand the **Web Sites** folder, and expand the Web site to which you want to add a virtual directory.
2. Right-click the Web site or folder within which you wish to create the virtual directory, select **New**, and then click **Virtual Directory**. The Virtual Directory Creation Wizard appears.
3. Click **Next**.
4. In the **Alias** box, type a name for the virtual directory. This is the name the user types, and should be short and easy to type.
5. Click **Next**.
6. In the **Path** box, type the name of the physical directory or click **Browse** to navigate to the physical directory in which the content for the virtual directory resides.
7. Click **Next**.
8. Under **Allow the following permissions**, select the check boxes for the access permissions you want to assign to your users, and then click **Next**.

| **Note** |
| --- |
| For security purposes, when selecting access permissions, consider allowing only the default **Read** permission. By restricting permissions in this way, you can mitigate potential attacks against your Web site by malicious users. |

1. Click **Finish**. The virtual directory is created below the currently selected folder level.

**To create a virtual directory by using Windows Explorer**

1. Open Windows Explorer.
2. Right-click the folder you want to be a virtual directory, and select **Sharing and Security**.
3. On the **Web Sharing** tab, click **Share this folder**.

4. On the **Edit Alias** dialog box, in the **Alias** box, type the name for the virtual directory.
5. Under **Access permissions**, select the check boxes for the type of access you want to assign to the virtual directory.
6. Under **Application permissions**, select the check boxes for the type of application access you want to assign to the virtual directory.
7. Click **OK**, and then click **OK** again.

# Create a Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

During installation of Internet Information Services (IIS), a default home directory and Web site configuration are created on your hard disk. Similarly, creating a Web site by using IIS Manager does not create content, but merely creates a directory structure and configuration files from which to publish the content. To publish your Web content, you can place content in the default home directory, or you can create a different home directory or virtual directory and place content there.

**Important** If you install ASP.NET with IIS 6.0, structure the content of your Web site so that requests for Web site content do not need to contain /bin in the request URL. As a security measure, ASP .NET returns a 404 error for all requests containing /bin in the requested URL.

**Requirements**

- Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To use the default Web site**

1. In IIS Manager, expand the local computer, and expand the **Web Sites** folder.
2. Right-click **Default Web Site**, and select **Properties**.
3. On the **Web Site** tab, under **Web site identification**, type the name of your Web site in the **Description** box.
4. Click **OK**. The new name of the site appears in IIS Manager.

**To create a new Web site**

1. In IIS Manager, expand the local computer, and right-click the **Web Sites** folder.
2. Select **New**, and then click **Web Site**. The Web Site Creation Wizard appears.
3. Click **Next**.
4. In the **Description** box, type a name for your site, and then click **Next**.
5. Type or select the IP address (the default is **All Unassigned**), TCP port, and host header (for example, *www.mysite.com*) for your site.
6. Click **Next**.
7. In the **Path** box, type the name of the directory or click **Browse** to navigate to the directory that contains, or will contain, the site content.
8. Click **Next**.
9. Select the check boxes for the Web site access permissions you want to assign to your users, and then click **Next**.
10. Click **Finish**.
11. To change these and other settings later, right-click the Web site, and select **Properties**.

# Debug Application Pool Failures

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To enable *orphaning* of a worker process serving an application pool (which keeps failed applications running while your diagnostic tools monitor them), and to attach a debugger to the worker process, you must change the values of three metabase properties. The values of these metabase properties, which are set by running a script, must specified as follows:

- Set the **OrphanWorkerProcess** metabase property to **TRUE** to notify the WWW service to orphan the worker process when it fails.
- Set the **OrphanActionExe** metabase property to specify an executable to run when the worker process is orphaned.
- Set the **OrphanActionParams** metabase property to configure command-line parameters for attaching the debugger to the worker process.

Consider the following guidelines when you enable debugging of application pool failures:

- Debugging requires that you set all three metabase properties outlined above; enabling debugging is ineffective until you attach a debugger.
- When you no longer want to use the debugging feature, disable it. Putting an application in orphan state wastes resources if the diagnostics are not running or will not be used.

**Requirements**

- Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To enable and configure debugging on a worker process**

1. Copy the following ADSI script (written in VBScript) into a text editor or word processor.

```
set appPoolObj=GetObject("IIS://localhost/W3svc/AppPools/app pool name")
' Set the application pool properties:
appPoolObj.Put "OrphanWorkerProcess", TRUE
appPoolObj.Put "OrphanActionExe", "full path\ntsd.exe"
appPoolObj.Put "OrphanActionParams", "-g -p %1%"
' Save the property changes in the metabase:
appPoolObj.SetInfo
WScript.Echo "After: " & appPoolObj.OrphanWorkerProcess & ", " &
appPoolObj.OrphanActionExe & ", " & OrphanActionParams
```

2. Save the file with the extension .vbs.
3. In the **Run** dialog box, type **cmd**, and then click **OK**.
4. At the command prompt, run the script by typing the following command:

   **cscript //nologo** *filename*

   where *filename* is the fully qualified path of the script file.

– or –

You can call the script by using a batch file that contains the following command:

```
"cscript //nologo filename"
```

where *filename* is the fully qualified path of the script file.

# Determine Web Sites Uniquely Identified by IP Addresses

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can determine whether you have any Web sites that are uniquely identified with an IP address by viewing the Properties sheet of the Web sites on your source server.

**Requirements**

- Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To determine Web sites that are uniquely identified by an IP address in IIS 5.0**

1. Open Administrative Tools, and then click **Internet Service Manager**.
2. In IIS Manager, expand the server, right-click the Web site, and then click **Properties**.
3. On the **Web Site** tab, under **IP Address**, note whether there is a specific IP address assigned to this Web site. A specific IP address entry in this box denotes a uniquely identified Web site.

**To determine Web sites that are uniquely identified by an IP address in IIS 4.0**

1. Click **Start**, click **Programs**, click **Windows NT 4.0 Option Pack**, click **Microsoft Internet Information Server**, and then click **Internet Service Manager**.
2. In Internet Service Manager, expand the server, right-click the Web site, and then click **Properties**.
3. On the **Web Site** tab, under **IP Address**, note whether there is a specific IP address assigned to this Web site. A specific IP address entry in this box denotes a uniquely identified Web site.

# Disable Network Adapters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To ensure that clients cannot connect to the Web server while it is being upgraded, disable the network adapter that connects the Web server to the clients before taking it offline for the upgrade.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: none.

**Procedures**

**To disable the network adapter for IIS 5 0**

1. Click **Start**, click **Settings**, and then click **Network Connections**.
2. Locate the network adapter that connects the Web server to the clients.
3. Right-click the network adapter, and then click **Disable**.

**To disable the network adapter for IIS 4 0**

1. In Control Panel, double-click **Network**.
2. On the **Bindings** tab, click the network adapter, click **Disable**, and then click **OK**.

# Enable ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A server running a member of the Microsoft® Windows® Server 2003 family supports application server functionality, with Microsoft ASP.NET as an option that you can enable when configuring the application server role. To deploy ASP.NET Web applications to a production server, you must be sure to enable the ASP.NET and Internet Information Services (IIS) roles on the production server before you distribute the application.

| **Note** |
| --- |
| If you want to install ASP.NET on a domain controller, there are special steps you must take to make the installation work properly. For more information, see Knowledge Base article 315158, Fix: ASP.NET does not work with the Default ASPNET Account on a Domain Controller. |

**Procedures**

ASP.NET, along with the .NET Framework version 1.1, is included with Windows Server 2003. You need to install it by using Add or Remove Programs, or enable it by using the Configure Your Server Wizard.

| **Note** |
| --- |
| When you use the Configure Your Server Wizard or the **Add or Remove Programs** dialog box to install ASP.NET on a server running Windows Server 2003, ASP.NET is automatically enabled in IIS Manager. However, if you install ASP.NET from a Web download or as part of an application such as Microsoft® Visual Studio®.NET, you must enable ASP.NET manually, as described later in this section. |

**To enable ASP.NET on a server running Windows Server 2003 by using the Configure Your Server wizard**

1. Click **Start**, and then click **Manage Your Server**.
2. In the **Manage Your Server** window, click **Add or remove a role**.
3. In the Configure Your Server wizard, click **Next**.
4. In the **Server Role** dialog box, click **Application Server (IIS, ASP.NET)** and then click **Next**.
5. In the **Application Server Options** dialog box, select the **Enable ASP.NET** check box.
6. Click **Next**, and then click **Next** again.
7. If you are prompted to do so, insert your Windows Server 2003 installation CD in the CD-ROM drive, and then click **Next**.
8. When the installation is complete, click **Finish**.

**To enable ASP.NET on a server running Windows Server 2003 by using Add or Remove Programs**

1. In Control Panel, click **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. In the **Components** box in the Windows Components Wizard, select the **Application Server** check box, and then click **Details**.
4. In the **Application Server** dialog box, select the **ASP.NET** check box, and then click **OK**.
5. In the Windows Components Wizard, click **Next** to begin installing ASP.NET.
6. When the Windows Components Wizard has finished configuring Windows Server 2003, click **Finish**.

# Enable Buffering on the Administrative Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you install the Remote Administration (HTML) Tool for IIS 6.0 on a server that has been upgraded to Windows Server 2003 from Microsoft Windows NT 4.0, you may receive an error when you try to view the administrative Web site. This error occurs because, by default, buffering is set to **False** in IIS 4.0. However, in IIS 6.0, buffering is set to **True** by default, and this setting is not changed during upgrade. To view the administrative Web site in IIS 6.0, you must enable buffering.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: lis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**How To**

**To enable buffering on the administrative Web site (for upgrades from Windows NT 4.0)**

1. In IIS Manager, double-click *servername* (local computer).
2. Double-click **Web Sites**.
3. Right-click the Administration Web site and then click **Properties**.
4. Click **Home Directory**, click **Configuration**, and then click **Options**.
5. Check the **Enable buffering** check box, and then click **OK** twice.

# Enable Logging_IIS_SP1_Dep

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can enable logging for individual Web and File Transfer Protocol (FTP) sites. After you enable logging for a Web or FTP site, all traffic to the site (including virtual directories) is written to the corresponding file for each site. You can also enable logging for specific virtual directories.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: lis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To enable logging on a Web or FTP site**

1. In IIS Manager, expand the local computer, expand the **Web or FTP Sites** directory, right-click the Web or FTP Site for which you want to enable logging, and then click **Properties**.
2. On the **Web Site** or **FTP Site** tab, click the **Enable logging** check box.
3. In the **Active log format** list box, click a format. By default, the format is **W3C Extended Log File Format**.

   | **Note** |
   | --- |
   | If you select Open Database Connectivity (ODBC) logging, click **Properties** and type the ODBC Data Source Name (DSN) and the name of the table within the database into the text boxes. If a user name and password are required to access the database, type the necessary credentials, and click **OK**. |

4. Click **Apply**, and then click **OK**.

**To enable logging for a specific virtual directory on a site**

1. In IIS Manager, expand the local computer, expand the **Web Sites** directory, right-click the virtual directory, and click **Properties**.
2. On the **Virtual Directory** or **Directory** tab, select the **Log visits** check box if it is not selected. By default, the check box is selected.
3. Click **Apply**, and then click **OK**.

# Enable Network Adapters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you have upgraded the Web server, you are ready to enable client access to the Web sites. During the upgrade process, you disabled the network adapter on the Web server to prevent users from accessing the Web server during the upgrade process. Now that you know the upgrade completed successfully, you can enable the network adapters.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: none.

**To enable the network adapter for IIS 6.0**

1. In Control Panel, click **Network Connections**, and then click **Local Area Connection**.
2. The network adapter is automatically enabled.

# Enable Security Auditing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Microsoft® Windows® Server 2003 uses security and system logs to store collected security events. Before enabling the system and security logs, you need to enable auditing for the system log and establish the number of events you want recorded in the security log. You customize system log events by configuring *auditing*. Auditing is the process that tracks the activities of users and processes by recording selected types of events in the security log of the Web server. You can enable auditing based on categories of security events such as:

- Any changes to user account and resource permissions.
- Any failed attempts for user logon.
- Any failed attempts for resource access.
- Any modification to the system files.

The most common security events recorded by the Web server are associated with user accounts and resource permissions.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Microsoft Management Console (MMC); Local Security Policy

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**
**To define or modify auditing policy settings for an event category on the local Web server**

1. Open Administrative Tools, and then click **Local Security Policy**.
2. In the console tree, click **Local Policies**, and then click **Audit Policy**.
3. In the details pane, double-click an event category for which you want to change the auditing policy settings.
4. On the **Properties** page for the event category, do one or both of the following:
    - To audit successful attempts, select the **Success** check box.
    - To audit unsuccessful attempts, select the **Failure** check box.
5. Click **OK**.

Perform the following procedure on the domain controller.

**To define or modify auditing policy settings for an event category within a domain or organizational unit, when the Web server is joined to a domain**

1. Open Administrative Tools, and then click **Active Directory Users and Computers**.
2. Right-click the appropriate domain, site, or organizational unit and then click **Properties**.
3. On the **Group Policy** tab, select an existing Group Policy object to edit the policy.

4. In **Group Policy Object Editor**, in the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policy**, and then click **Audit Policy**.
5. In the details pane, double-click an event category for which you want to change the auditing policy settings.
6. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
7. Do one or both of the following:
    - To audit successful attempts, select the **Success** check box.
    - To audit unsuccessful attempts, select the **Failure** check box.
8. Click **OK**.

# Enable the WWW Service After Upgrade

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you upgrade a Web server running Internet Information Services (IIS) 5.0 and the Microsoft® Windows® 2000 operating system, the World Wide Web Publishing Service (WWW service) is disabled unless, prior to upgrade, you elected to run the IIS Lockdown Tool or you added the entry RetainW3SVCStatus to the registry. However, if neither of those was done, then you need to enable the WWW service after the upgrade.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Microsoft Management Console (MMC); Local Security Policy.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To enable the WWW Service after upgrade**

1. Open Administrative Tools, and then click **Services**.
2. In the **Services** pane, double-click **World Wide Web Publishing Service**.
3. On the **Properties** page, on the **General** tab, select **Automatic** from the **Startup Type** drop-down box.
4. Under **Service Status**, click the **Start** button.
5. Click **OK**.

# Enable Web Site Content Auditing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Once you have enabled security auditing, you must also enable auditing on the Web site content (files and folders) in order to track any modification or deletion of the content.

Before you set up auditing for files and folders, you must first enable object access auditing. This security setting determines whether to audit the event of a user accessing an object, such as a file, folder, or printer. Enabling object access auditing is accomplished by defining auditing policy settings for the object access event category of the Audit Policies in Local Security Settings. If you do not enable object access auditing, you receive an error message when you set up auditing for files and folders, and no files or folders are audited. After object access auditing is enabled, you can view the security log in Event Viewer to review the results of your changes. You can then set up Web site content auditing.

| **Tip** |
| --- |
| Because the security log is limited in size, carefully select the files and folders to be audited. In addition, consider the amount of disk space that you want to devote to the security log. The maximum size for the security log is defined in Event Viewer. |

If file or folder auditing has been inherited from the parent folder, you will see the following.

- In the **Auditing Entry for File or Folder** dialog box, in the **Access** box, the check boxes are unavailable.

-or-

- In the **Advanced Security Settings for File or Folder** dialog box, the **Remove** button is unavailable.

**Requirements**

- Credentials: You must be logged on as a member of the Administrators group or you must have been granted the Manage auditing and security log right in Group Policy to perform this procedure.
- Tools: Windows Explorer
- File system: To enable auditing of Web site content, the disk volumes on which the Web site is stored must use the NTFS file system.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

### To enable object access auditing

1. Open Administrative Tools, and then click **Local Security Policy**.
2. Expand **Local Policies**, and then click **Audit Policy**.
3. Right-click **Audit object access**, and then click **Properties**.
4. Enable auditing by clicking one of the following:
    - Click **Success** to generate an audit entry when a user successfully accesses an object.
    - Click **Failure** to generate an audit entry when a user unsuccessfully attempts to access an object.
    - If you clear both check boxes, object access auditing is turned off.

5. Click **OK**.

**To apply or modify auditing policy settings for a local file or folder**

1. Open Accessories, and then click **Windows Explorer**.
2. Right-click the file or folder for which you want to set audit policy settings, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, and then click the **Auditing** tab.
4. Do one of the following:
   - To set up auditing for a new user or group, click **Add**. In **Enter the object name to select**, type the name of the user or group that you want to audit, and then click **OK**.
   - To remove auditing for an existing group or user, click the group or user name, click **Remove**, click **OK**, and then skip the rest of this procedure.
   - To view or change auditing for an existing group or user, click the name of the group or user, and then click **Edit**.
5. In the **Apply onto** box, click the location where you want auditing to take place.
6. In the **Access** box, indicate what actions you want to audit by selecting the appropriate check boxes:
   - To audit successful events, select the **Successful** check box.
   - To stop auditing successful events, clear the **Successful** check box.
   - To audit unsuccessful events, select the **Failed** check box.
   - To stop auditing unsuccessful events, clear the **Failed** check box.
   - To stop auditing all events, click **Clear All**.
7. If you want to prevent subsequent files and subfolders of the original object from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.

# Export a Server Certificate

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Web server certificates contain information about the server that allows the client to positively identify the server over a network before sharing sensitive information, in a process called *authentication*. Secure Sockets Layer (SSL) uses these certificates for authentication, and uses encryption for message integrity and confidentiality. SSL is a public key–based security protocol that is used by Internet services and clients to authenticate each other and to establish message integrity and confidentiality.

If you use SSL to protect confidential information exchanged between the Web server and the client, you must migrate or export the certificates and the associated private keys from the source server to the target server.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

### To export a server certificate

1. In the **Run** dialog box, type **mmc**, and then click **OK**. The Microsoft Management Console (MMC) appears.
2. If you do not have **Certificate Manager** installed in MMC, you need to install it.

   For more information on how to add the Certificate snap-in to an MMC console, see the procedure "To add the Certificates Snap-in to MMC" in Install a Server Certificate this appendix.
3. In the console tree, click the logical store where the certificate you want to export exists. Usually this is in the **Certificates** folder in the **Personal** directory under **Certificates (Local Computer)** on the Console Root.
4. Right-click the certificate you want to export, click **All Tasks**, and click **Export** to start the Certificate Export Wizard.
5. Click **Next**.
6. On **Export Private Key**, click **Yes** to export the private key.

   | **Important** |
   | --- |
   | You must export the private key along with your certificate for it to be valid on your target server. Otherwise, you will have to request a new certificate for the target server. |

7. In the **Export File Format** dialog box, click the format you want for the certificate. If the certificate has already been formatted, that format is selected as the default. Click **Next**.

   Do not select **Delete the private key if export is successful**, because this will disable the SSL site that corresponds to that private key.
8. Continue to follow steps in the wizard, and enter a password for the certificate backup file when prompted. Using a strong password is highly recommended because it ensures that the private key is well protected.
9. Type the name of the file you want to export, or click **Browse** to search for the file. Click **Next**.
10. Click **Finish** to complete the Certificate Export Wizard.

# Gather and Display WWW Service Uptime Data

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use Internet Information Services (IIS) 6.0 Performance Monitor to record and display data about the uptime of World Wide Web Publishing Service (WWW service) and Web sites by using the Service Uptime performance counter. These procedures support the task of determining the availability of the WWW service and Web sites running on IIS.

Use the following procedures to:

- Create a log file to record WWW service uptime data.
- Select a performance counter to generate uptime data for the WWW service and your Web sites.
- Start the Service Uptime performance counter.
- Connect System Monitor to the data in the log file you are using to gather WWW service and Web site uptime data.
- Read the display window in the details panel of System Monitor to display WWW service and Web site uptime data.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Perfmon.msc; Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**Gathering WWW Service Uptime Data**
You can use IIS 6.0 Performance Monitor to create a log file, and to select and start the Service Uptime performance counter.

### To gather uptime data on the WWW service and Web sites

1. Open Administrative Tools, and then click **Performance**.
2. In the console pane, click **Performance Logs and Alerts**, and then click **Counter Logs**.
3. Right-click in the details pane, and select **New Log Settings** from the menu.
4. Enter the name you want to use for the new log, for example, *WWWServiceUptime*, and then click **OK**.

   The property sheets for the new counter log open.
5. On the **General** tab, click **Add Counters**.
6. In the **Add Counters** window, if the performance counters you want to monitor are on the local computer, click **Use local computer counters**. If the performance counters you want to monitor are on a remote computer, click **Select counters from computer**, and then click the remote computer in the drop-down list.
7. Under **Performance object**, click **Web Service** in the drop-down list.
8. Click **Select counters from list**, and then click the **Service Uptime** counter.
9. Click **All instances**.
10. Click **Add**, and then click **Close**.

The **Service Uptime** counter for the computer you selected displays in the **Counters** window on the **General** tab.

11. Click **OK**.
12. To start logging data, select the log file on the **Performance** detail pane and, on the **Action** menu, click **Start**.
13. The counter log you created displays in green on the **Computer Logs** detail pane.
14. To manually stop logging data, select the log file and, on the **Action** menu, click **Stop**.
15. To start logging data again, select the log file and, on the **Action** menu, click **Start**.

**Displaying WWW Service Uptime Data**

Consider the following guidelines for displaying WWW service uptime data:

- This procedure reads a log file that gathers WWW service uptime data. Before you can read the log file, you must create the file and connect it to the performance counter, as described earlier in this section.
- The file logs data for all instances of the Service Uptime performance counter. One instance, named the Total instance, accumulates WWW service data. Each of the other instances accumulates data about a Web site. Each Web site is recorded by a separate instance.

**To display data for WWW service uptime and Web site uptime**

1. Open Administrative Tools, and click **Performance**.
2. In the console pane, click **System Monitor**.
3. Right-click in the details pane, and then click **Properties**.
4. On the **Source** tab, under **Data Source**, click **Log files**, and then click **Add**.
5. In the **Select Log File** window, navigate to the log file that you created to gather WWW service and Web sites uptime data.
6. Select the log file name from the list, and then click **Open**.
7. Click the **Data** tab, and then click the **Add** button. Under **Performance object**, select **Web Service** from the drop-down list.
8. Click the **Select counter from list** button and click **Service Uptime**.
9. Click the **All instances** button.
10. Click **Add**, and then click **Close**.
11. On the **System Monitor Properties** window, click **OK**.

    The Service Uptime counter instances appear in the list in the **Performance** detail pane.
12. Select the instance desired to view the graph and statistics for WWW Service ("_Total"), or for each Web site.

# Grant User Rights to a Service Account

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Typically, the user rights assigned to the IIS_WPG group is sufficient for most Web sites or applications. However, when a Web site or application requires additional user rights to run properly, you must assign the required rights to the service account that is used as the identity for the Web sites and applications.

You grant user rights based on where the account is stored. If the service account is created locally on the Web server, you make changes in user rights through Local Computer Policy by using the Group Policy Object Editor MMC snap-in. When the service account is created in Active Directory, make the changes on the appropriate Group Policy object in Active Directory.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**To add the Group Policy Object Editor to MMC**

1. In the **Run** dialog box, type **mmc**, and then click **OK**.

   The Microsoft Management Console appears.

2. On the **File** menu, click **Add/Remove Snap-in**.
3. On the **Standalone** tab, click **Add**.
4. In the **Available Standalone Snap-ins** list box, click **Group Policy Object Editor**, and then click **Add**.
5. In the **Select Group Policy Object** dialog box, in **Group Policy Object**, select **Local Computer**, and then click **Finish**.
6. Click **Close**, and then click **OK**.

**To grant user rights to a service account when the service account is stored locally on the Web server**

1. In MMC, open the **Group Policy Object Editor**.
2. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **User Rights Assignment**.
3. In the details pane, double-click the **user right** that you want to grant to the service account.
4. In the **user_right Properties** dialog box (where **user_right** is the user right you selected in Step 3), click **Add User or Group**.
5. In the **Select Users, Computers, or Group** dialog box, type the name of the service account, and then click **OK**.
6. In the **user_right Properties** dialog box (where **user_right** is the user right you selected in Step 3), click **OK**.

**To grant user rights to a service account when the service account is stored in Active Directory**

1. In MMC, open **Active Directory Users and Computers**.
2. In the console tree, browse to the organizational unit that contains the Group Policy object that you want to modify, right click the organizational unit, and then click **Properties**.

3. In the **organizational_unit Properties** dialog box (where **organizational_unit** is the organizational unit you selected in Step 2), click the **Group Policy** tab.
4. Select the Group Policy object that you want to modify, and then click **Edit**.

   The Group Policy Object Editor appears.
5. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **User Rights Assignment**.
6. In the details pane, double-click the user right that you want to grant to the service account.
7. In the **user_right Properties** dialog box (where **user_right** is the user right you selected in Step 3), click **Add User or Group**.
8. In the **Select Users, Computers, or Group** dialog box, type the name of the service account, and then click **OK**.
9. In the **user_right Properties** dialog box (where **user_right** is the user right you selected in Step 3), click **OK**.
10. Close the Group Policy Object Editor.
11. Click **OK**.

# Install a Server Certificate

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Web server certificates contain information about the server that allows the client to positively identify the server over a network before sharing sensitive information. This process is called *authentication*. If you use Secure Sockets Layer (SSL) to protect confidential information exchanged between the Web server and the client and you have exported the certificates from the source server to the target server, the server certificate needs to be installed on the Web server before you can assign the server certificate to Web sites for use with SSL

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Certificates MMC snap-in.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

### To add the Certificates Snap-in to MMC

1. In the **Run** dialog box, type **mmc**, and then click **OK**.

   The Microsoft Management Console appears.

2. On the **File** menu, click **Add/Remove Snap-in**.
3. On the **Standalone** tab, click **Add**.
4. In the **Available Standalone Snap-ins** list box, click **Certificates**, and then click **Add**.
5. Click the **Computer account** option, and then click **Next**.
6. Click the **Local computer (the computer this console is running on)** option, and then click **Finish**.
7. Click **Close**, and then click **OK**.

### To install a server certificate on a Web server

1. In MMC, open the **Certificates** snap-in.
2. In the console tree, click the logical store where you want to import the certificate.

   The default location of the logical store for certificates is on the Console Root in the **Certificates (Local Computer)/ Personal/Certificates** folder.

3. On the **Action** menu, point to **All Tasks**, and then click **Import** to start the Certificate Import Wizard.

   | **Important** |
   | --- |
   | You should only import certificates obtained from trusted sources. Importing an altered or unreliable certificate could compromise the security of any system component that uses the imported certificate. |

4. Click **Next**.
5. Type the name of the file that contains the certificate to be imported, or click **Browse** and navigate to the file.

   Certificates can be stored in several different file formats. The most secure format is Public-Key Cryptography Standard (PKCS) #12, an encryption format that requires a password to encrypt the private key. It is recommended that you send

certificates using this format for optimum security.

If the certificate file is in a format other than PKCS #12, skip to step 8.

If the certificate file is in the PKCS #12 format, do the following:

- In the **Password** box, type the password used to encrypt the private key. You must have access to the password that was originally used to secure the file.

- (Optional) If you want to be able to use strong private key protection, select the **Enable strong private key protection** check box, if available.

- (Optional) If you want to back up or transport your keys at a later time, select the **Mark key as exportable** check box.

6. Click **Next**.
7. In the **Certificate Store** dialog box, do one of the following:

   - If the certificate should be automatically placed in a certificate store based on the type of certificate, select **Automatically select the certificate store based on the type of certificate**.

   - If you want to specify where the certificate is stored, select **Place all certificates in the following store**, click **Browse**, and select the certificate store to use.

8. Click **Next**, and then click **Finish**.

The file from which you import certificates remains intact after you have completed importing the certificates. You can use Windows Explorer to delete the file if it is no longer needed.

# Install IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can install Internet Information Services (IIS) 6.0 by using **Add/Remove Windows Components** from Add or Remove Programs in Control Panel.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**To install IIS 6.0 using Control Panel**

1. In Control Panel, double-click **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. In the **Components** list box, click **Application Server**, and then click **Details**.
4. In the **Subcomponents of Application Server** box, click **Internet Information Services (IIS)**.
5. Click **OK** to start the installation of IIS 6.0.

# Install Subauthentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To use Digest authentication in Internet Information Services (IIS) 6.0 when the domain controller is running Microsoft® Windows® 2000, you must enable subauthentication, which is not installed by default on IIS 6.0. There are three requirements for enabling subauthentication:

- Install the subauthentication component, Iissuba.dll.
- Set the **UseDigestSSP** metabase property to **False**.
- Set the identity of the application pool to **Local System**. For more information about setting application pool identity, see Ensuring Application Availability.

Consider the following guidelines for enabling subauthentication:

- The requirement to use subauthentication applies to Digest authentication only. Using Advanced Digest authentication does not require subauthentication.
- When you no longer want to use subauthentication, unregister the sub-authentication component and set the identity of the application pool to Local System.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname***mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To install and register the subauthentication component**

1. In the **Run** dialog box, type **cmd**, and click **OK**.
2. At the command prompt type:

   ```
   rundll32 systemroot\system32\iissuba.dll, RegisterIISSUBA.
   ```

3. Press ENTER.
4. For any application pools that use Digest authentication, set the application pool identity to Local System.

   For more information about Digest authentication, see Configure Web Server Authentication.

   For more information on configuring application pool identity, see Configure Application Pool Identity.

**To unregister the subauthentication component**

1. In the **Run** dialog box, type **cmd**, and click **OK**.
2. At the command prompt, type:

   ```
   rundll32 systemroot\system32\iissuba.dll,UnregisterIISSUBA.
   ```

3. Press ENTER.

# Isolate Applications in Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

An *application pool* is a configuration that links one or more applications to a set of one or more worker processes. Because applications in an application pool are separated from other applications by worker process boundaries, an application in one application pool is not affected by problems caused by applications in other application pools.

By creating application pools and assigning Web sites and applications to them, you can improve the availability and reliability of your Web sites and applications.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To create a new application pool**

1. In IIS Manager, expand the local computer, right-click **Application Pools**, click **New**, and then click **Application Pool**.
2. If the ID that appears in the **Application pool ID** box is not the ID that you want, type a new ID.
3. Under **Application pool settings**, click one of the following options:
   - **Use default settings for the new application pool**

   -or-

   - **Use existing application pool as a template**.
4. If you click the **Use existing application pool as template** option, in the drop-down box, click the **Application pool name** of the application pool you want to use as a template.
5. Click **OK**.

**To assign a Web site or application to an application pool**

1. In IIS Manager, right-click the application that you want to assign to an application pool, and then click **Properties**.
2. Click the **Virtual Directory**, **Directory**, or **Home Directory** tab, depending on the type of application you have selected.
3. If you are assigning a directory or virtual directory, verify that **Application name** is filled in. If the **Application name** box is not filled in, click **Create**, and then type a name.
4. In the **Application pool** list box, click the name of the application pool to which you want to assign the Web site.
5. Click **OK**.

# Make a Service Account a Member of the Local Administrators Group

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If an application requires specific user rights to run successfully, but the service account for the application cannot be assigned the appropriate permissions, make the service account a member of the local administrators group. This should give the service account sufficient user rights to allow the application to run successfully.

**Procedures**

**To add a service account to the local administrators group**

1. In Administrative Tools, click **Computer Management**.
2. In the console tree, expand **Local Users and Groups**, and then click **Groups**.
3. Right-click the **Administrators** group, and then click **Add to Group**.
4. Click **Add**.
5. Click **Look in** to display a list of domains from which users and groups can be added to the group.
6. In **Location**, click the domain containing the users and computers you want to add, and then click **OK**.
7. In **Enter the object names to select**, type the name of the user or group you want to add to the group, and then click **OK**.
8. If you want to validate the user or group names that you are adding, click **Check Names**.

# Migrate CDONTS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Collaboration Data Objects (CDO) for Microsoft® Windows NT® Server (CDONTS) has been removed from Microsoft® Windows® Server 2003. However, if you upgrade to Windows Server 2003, CDONTS remains on your server, because Cdonts.dll is not removed during an upgrade to Windows Server 2003. CDONTS is not installed when you perform a clean installation of Windows Server 2003, but if necessary, you can copy it to the computer running Windows Server 2003 and register it.

Note: If your Web applications use CDONTS, you can modify your code to use Collaboration Data Objects for Windows 2000 (CDOSYS), which is supported by Windows Server 2003.

**Procedures**

**To copy CDONTS to a computer running Windows Server 2003**

1. On the source server, copy Cdonts.dll from the *systemroot***\system32** folder to a floppy disk.
2. On the target server, copy Cdonts.dll from the floppy disk to the folder *systemroot***\system32**.
3. To register CDONTS, on the target server, in the **Run** dialog box, type **cmd**, and then click **OK**.
4. At the command prompt, change to the *systemroot***\system32** directory and then type:

   **regsvr32** *%windir%* **\system32\cdonts.dll**.

   If the process is successful, the following message displays:

   **DllRegisterServer in cdonts.dll succeeded.**
5. Click **OK**.

# Modify the IIS Metabase Directly

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you need to modify the Internet Information Services (IIS) metabase, you can do so directly in the Metabase.xml file by using a text editor. Before you can modify the Metabase .xml file, you must enable the edit-while-running feature of the metabase by using IIS Manager. Once you have made your changes and saved the Metabase.xml file, disable the edit-while-running feature of the metabase. To implement changes to some metabase properties, you might need to restart the server.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

### To enable the edit-while running feature of the metabase by using IIS Manager

1. In IIS Manager, right-click the local computer, and then click **Properties**.
2. Select the **Enable Direct Metabase Edit** check box, and then click **OK**.

### To modify the IIS metabase

1. Open the Metabase.xml file in a text editor. The default path to this file is *systemroot***\system32\inetserv\metabase.xml**
2. Modify the metabase properties that you wish to change in the Metabase.xml file.
3. Save the changes to the file, and close the text editor.

Most changes to metabase properties are automatically recognized by IIS; in some cases, you must restart IIS for the metabase property changes to go into effect.

### To disable the Edit-while running feature of the metabase by using IIS Manager

1. In IIS Manager, right-click the local computer, and then click **Properties**.
2. Clear the **Enable Direct Metabase Edit** check box, and then click **OK**.

# Monitor Active Web and FTP Connections

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To ensure that service to clients is not interrupted; monitor the Web server for any active Web and File Transfer Protocol (FTP) connections before taking the Web server offline. Internet Information Services (IIS) 4.0, IIS 5.0, and IIS 6.0 include performance monitor counters that can be used to monitor the active Web and FTP connections. Monitor the active Web and FTP connections to ensure one of the following is true:

- The number of active Web and FTP connections is zero.
- All active Web or FTP sessions can be accounted for and can be terminated.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Perfmon.msc, Perfmon.exe.

**Procedures**

**To monitor active Web and FTP connections in IIS 6.0**

1. Open Administrative Tools, and then click **Performance**.
2. Right-click the **System Monitor details** pane, and then click **Add Counters**.
3. Do one of the following:
   - To monitor any computer on which the monitoring console is run, click **Use local computer counters**.

   -or-

   - To monitor a specific computer, regardless of where the monitoring console is run, click **Select counters from computer**, and specify a computer name or IP address.
4. In **Performance object**, click:
   - **Web Service** to monitor active Web connections.

   -or-

   - **FTP Service** to monitor active FTP connections.
5. Click **Select counters from list**, and select **Current Connections**.
6. Click **All instances**.
7. Click **Add**, and then click **Close**.

**To monitor active Web and FTP connections in IIS 5.0**

1. Open Administrative Tools, and then click **Performance**.
2. Right-click the **System Monitor details** pane and then click **Add Counters**.
3. Do one of the following:

- To monitor any computer on which the monitoring console is run, click **Use local computer counters**.

-or-

- To monitor a specific computer, regardless of where the monitoring console is run, click **Select counters from computer**, and specify a computer name or IP address.

4. In **Performance object**, click:

   - **Web Service** to monitor active Web connections.

   -or-

   - **FTP Service** to monitor active FTP connections.

5. Click **Select counters from list**, and then click **Current Connections**.
6. Click **All instances**.
7. Click **Add**, and then click **Close**.

**To monitor active Web and FTP connections in IIS 4.0**

1. Open Administrative Tools, and then click **Performance Monitor**.
2. Click **Edit**, and then click **Add to Chart**.
3. The local computer is listed in the Computers box. To monitor any other computer on which the monitoring console is run, click the button to the right of **Computers** and select a computer from the **Select Computer** list, and then click **OK**.
4. In **Performance object**, click:

   - **Web Service** to monitor active Web connections.

   -or-

   - **FTP Service** to monitor active FTP connections.

5. In **Counters**, click **Current Connections**.

| **Note** |
| --- |
| The performance counters are not installed when IIS 4.0 is installed on Microsoft® Windows NT® version 4.0. In this case, if the performance counters have not been manually enabled, they do not appear in the Counters list. For more information about enabling the performance counters, see Knowledge Base article 226512, HOW TO: Reinstall IIS 4.0 Performance Monitor Counters. |

6. Click **Add**, and then click **Done**.

# Obtain a Server Certificate for an NLB Cluster

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use the Web Server Certificate Wizard to obtain a server certificate to be used with all Web servers in a Network Load Balancing (NLB) cluster. To ensure that the name on the certificate matches the name of the server, specify the virtual name of the NLB cluster as the common name on the certificate. After you receive the server certificate, install it on all machines in the NLB cluster.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc, Web Server Certificate Wizard.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**
**To obtain a server certificate for an NLB cluster**

1. In IIS Manager, double-click the local computer, and then double-click the **Web Sites** folder.
2. Right-click the Web site for which you want to obtain a certificate, and then click **Properties**.
3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.
4. In the Web Server Certificate Wizard, click **Create a new certificate**.
5. Follow the steps in the Web Server Certificate Wizard. On the **Your Site's Common Name** page, type the virtual name of the NLB cluster in the **Common name** box, and then click **Next** to continue following the steps in the wizard.

By default, the certificate request file is saved as C:\Certreq.txt, but the wizard allows you to specify a different location.

# Pause Web or FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To ensure that service to clients is not interrupted, pause all Web and File Transfer Protocol (FTP) sites on the Web server before taking the server offline. Pausing a site prevents the site from accepting new connections, but does not affect requests that are already being processed. Stopping a site does not interfere with other Internet services that are running.

You can pause a Web or FTP site by using either IIS Manager or a command-line script.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**Pausing Web Sites**
If you pause a Web site, you can temporarily reduce the load on the computer that hosts the site, and can make changes to the folder or document structure of the Web site.

**To pause a Web site by using Internet Services Manager in IIS 4.0**

1. Click **Start**, click **Programs**, click **Windows NT 4.0 Option Pack**, click **Microsoft Internet Information Server**, and then click **Internet Services Manager**.
2. In Internet Services Manager, right-click the Web site you want to pause, and then click **Pause**.

**To pause a Web site by using Internet Services Manager in IIS 5.0**

1. Open Administrative Tools, and then click **Internet Services Manager**.
2. In Internet Services Manager, right-click the Web site you want to pause, and then click **Pause**.

| **Note** |
| --- |
| If a site stops unexpectedly, Internet Services Manager might not correctly indicate the state of the site. In IIS, right-click the **Web Sites** folder and then click **Refresh** to see the current state of all Web sites. |

**Pausing FTP Sites**
**To pause an FTP site by using Internet Services Manager in IIS 4.0**

1. Click **Start**, click **Programs**, click **Windows NT 4.0 Option Pack**, click **Microsoft Internet Information Server**, and then click **Internet Services Manager**.
2. In Internet Services Manager, right-click the FTP site you want to pause; and then click **Pause**.

**To pause an FTP site by using Internet Services Manager in IIS 5.0**

1. Open Administrative Tools, and then click **Internet Services Manager**.
2. In Internet Services Manager, right-click the FTP site you want to pause, and click **Pause**.

| **Note** |
| --- |

If a site stops unexpectedly, Internet Services Manager might not correctly indicate the state of the site. In IIS, right-click the **FTP Sites** folder and click **Refresh** to see the current state of all FTP sites.

# Publish Web Site Content with FrontPage

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage Server Extensions from Microsoft provides a Web-based administration tool for extending Web sites. *Extending* a Web site means implementing features of FrontPage Server Extensions that enable the site owner to author the site in Microsoft® FrontPage® and delegate Web site ownership and Web site administration credentials. Making content available on an extended Web site is called *publishing* the Web site. A Web site that is not extended cannot be opened or authored in FrontPage.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**Extending a Web site**
You can use the FrontPage 2002 Server Administration tool to extend a Web site.

### To extend a Web site by using the FrontPage 2002 Server Administration tool

1. Open Administrative Tools, and then click **Microsoft SharePoint Administrator**.

| **Note** |
| --- |
| If IIS returns error 404, "The page cannot be found," check to make sure FrontPage 2002 Server Extensions is enabled by using **Add or Remove Programs** in Control Panel. |

1. On the **FrontPage Server Extensions 2002 Server Administration** site, check to see that the Web site you created is listed in the **Virtual Servers** section.
2. To extend the Web site, click the **Extend** link next to the name of the Web site.
3. To specify an administrator account for the extended Web site, type the administrator name in **Administrator user name**.
4. Click **Submit**.
5. The **FrontPage 2002 Server Administration** tool adds FrontPage Server Extensions template directories to the content directory of your extended Web site, and adds other files that contain metadata.

**Publishing a Web site**
If your Internet service provider (ISP) has FrontPage Server Extensions from Microsoft installed, you can publish to the Web server by using HTTP. Otherwise, you can use FrontPage to publish your Web site to a File Transfer Protocol (FTP) server.

| **Note** |
| --- |
| If you publish to a location on your local computer, your Web site will not have the full FrontPage functionality unless your computer is a server that has FrontPage Server Extensions or SharePoint Team Services from Microsoft installed. |

### To publish to a Web server by using HTTP

1. In Microsoft FrontPage, on the **File** menu, click **Publish Web**.
2. In the **Publish Destination** dialog box, do one of the following:
   - Type the location of a Web server.

     -or-

   - Click the drop-down arrow to select a location to which you have already published another Web site.

     -or-

   - Click **Browse** to find the publishing location.

| **Note** |
| --- |
| If you have previously chosen a publishing destination for this Web site, the **Publish Destination** dialog box will not appear; proceed to step 4. |

3. Click **OK**.
4. Specify the pages you want to publish by clicking **Options** in the **Publish Web** dialog.
5. Click the **Publish** tab, and do one or more of the following:
   - Under **Publish**, specify whether you want to publish only pages that have changed, or all pages.
   - Under **Changes**, specify how you want FrontPage to determine which pages have been changed.
   - If you want to create a log file for the changes made during publishing, select that check box.
6. Click **OK**.
7. To publish subwebs, select the **Include subwebs** check box.
8. Click **Publish**.

| **Note** |
| --- |
| If you cancel publishing in the middle of the operation, files that have already been published remain on the destination Web server. |

FrontPage publishes your Web site to the Web server you specified. If you want to verify that your Web site was successfully published, click the hyperlink that is displayed after the Web has been published — your Web browser will open to the Web site you just published.

**To publish to an FTP site**

1. In FrontPage, on the **File** menu, click **Publish Web**.
2. In the **Publish Destination** dialog box, type the location of the FTP server, or click the drop-down arrow to select a location to which you have already published another Web site.

| **Note** |
| --- |
| If you have previously chosen a publishing destination for this Web site, the **Publish Destination** dialog box will not appear; proceed to step 4. |

3. Click **OK**.

4. Specify the pages you want to publish by clicking **Options** in the **Publish Web** dialog.
5. Click the **Publish** tab, and do one or more of the following:
6. Under **Publish**, specify whether you want to publish only pages that have changed, or all pages.
7. Under **Changes**, specify how you want FrontPage to determine which pages have been changed.
8. If you want to create a log file for changes made during publishing, select that check box.
9. Click **OK**.
10. To publish subwebs, select the **Include subwebs** check box.
11. Click **Publish**.

| **Note** |
| --- |
| If you cancel publishing in the middle of the operation, files that have already been published remain on the destination FTP server. |

12. FrontPage publishes your Web site to the FTP server you specified.

# Remove Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Internet Information Services (IIS) Lockdown Tool works by turning off unnecessary IIS features and components to reduce the attack surface of the Web server. To assist in mitigating these potential risks, remove any unnecessary virtual directories.

**Procedures**

**To remove a virtual directory using IIS Manager**

1. In IIS Manager, expand the site containing the virtual directory you want to remove, right-click the virtual directory, and then click **Delete**.
2. Click **Yes**.

**To remove a virtual directory using Windows Explorer**

1. In Windows Explorer, browse to the folder containing the virtual directory you want to remove, right-click the directory, and then click **Sharing and Security**.
2. Click the **Web Sharing** tab, click **Remove**, and then click **Yes**.
3. Click **OK**.

   You can verify that the virtual directory was deleted by starting IIS Manager and expanding the Web site.

| **Note** |
| --- |
| The following method does not work for root virtual directories |

**To delete a virtual directory using the Iisvdir.vbs administration script**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
2. At the command prompt, use the **cd** command to change to the directory where the Iisvdir.vbs script is installed. The default location for this file is *systemroot***/system32/iisvdir.vbs**.
3. At the command prompt, type:

   ```
   cscript iisvdir.vbs /delete "Sample Web Site" VirtualDirectoryName.
   ```

   Substitute your Web site name and virtual directory name as appropriate. If there are spaces in the Web site name, use quotation marks around the Web site name, as shown in the preceding example.
4. Click **OK**.

# Request a Server Certificate

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To enable Secure Sockets Layer (SSL) on a Web site, you must first request a server certificate from a trusted certification authority. After you obtain the certificate, you need to install it on the Web server, and then assign it to one or more Web sites.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**To request a server certificate**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.
2. Right-click the Web site or file that you want, and then click **Properties**.
3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.
4. In the Web Server Certificate Wizard, click **Next**, and then click **Create a new certificate**.
5. Follow the steps in the Web Server Certificate Wizard, which guides you through the process of requesting a new server certificate.

# Secure the Root Folder of Each Disk Volume

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Immediately after a new installation of Microsoft® Windows® Server 2003, the special group Everyone has Read and Execute permissions on the root of the *system volume*, which is the disk volume where Windows Server 2003 is installed.

Any folders created beneath the root of the system volume automatically inherit the permissions assigned to the root of the system volume. This means that the Everyone group will have Read and Execute permissions on any new folders created immediately beneath the root of the system volume. To prevent an accidental breach in security, remove the permissions assigned to the special group "Everyone" on dedicated Web servers.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.
- File System: The system volume must use the NTFS file system if you want to set file and folder permissions.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To secure the root of the system volume by removing permissions**

1. Open Accessories, and then click **Windows Explorer**.
2. In Windows Explorer, locate the root of the system volume.
3. Right-click the root of the system volume, click **Properties**, and then click the **Security** tab.
4. In the **Group or user names** list box, click **Everyone**, and then click **Remove**.
5. Click **OK**.

# Secure Windows Server 2003 Built-in Accounts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After the installation of Microsoft® Windows® Server 2003, the built-in accounts Administrator and Guest exist on the Web server. In some instances, potential attackers can exploit these well known accounts unless they are renamed or disabled.

The Administrator account can be renamed, but cannot be disabled. The Guest account can be renamed and disabled. To help prevent potential attackers from exploiting these accounts, do the following:

- Rename the Administrator account.
- Rename and disable the Guest account.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: lis.msc.

**Recommendation**

As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**

**To rename the Administrator user account**

1. In Control Panel, click **Administrative Tools**, and then click **Computer Management**.
2. In the console tree, expand **Local Users and Groups**, and then click **Users**.
3. In the details pane, right-click **Administrator**, and then click **Rename**.
4. Type the new user name, and then press ENTER.

**To disable and rename the Guest user account**

1. In Control Panel, click **Administrative Tools**, and then click **Computer Management**.
2. In the console tree, expand **Local Users and Groups**, and then click **Users**.
3. In the details pane, right-click **Guest**, and then click **Properties**.
4. In the **Guest Properties** dialog box, on the **General** tab, click the **Account is disabled** check box, and then click **OK**.
5. In the **Details** pane, right-click **Guest**, and then click **Rename**.
6. Type the new user name, and then press ENTER.

# Set Processor Affinity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

On a multi-CPU server, you can configure application pools to establish affinity between worker processes and CPUs in order to take advantage of more frequent CPU cache hits. You can use processor affinity to control processors when applications place demands on system resources.

Processor affinity is used when it is necessary to continue processing an application that is tied to a subset of the hardware, so that transitions from one processor to another are minimized. You can establish affinity on a server that is supporting multiple applications so that a processor is dedicated to each application.

Consider the following when setting processor affinity to enable debugging of application pool failures:

- By default, processor affinity is off, and the load is distributed across all available CPUs.
- Processor affinity is used in conjunction with the processor affinity mask setting that is used to specify the CPUs.
- The metabase property **SMPProcessorAffinityMask** is a hexadecimal number. The maximum value, FFFFFFFF, is equivalent to decimal 32, the largest number of processors currently allowed. The hexadecimal value, translated to binary, gives the processors that are connected to the worker process in the application pool. For example, the hexadecimal number FFFFFFFF translates to the binary number 1111 1111 1111 1111 1111 1111 11111111 . This means that all 32 processors are available to the worker processes in the application pool. If the hexadecimal value is 00000005, the binary value is 0000 0000 0000 0000 0000 0000 0000 0101, which means that processor 0 and processor 2 are connected to the worker processes in the application pool (processors are numbered from the right, beginning with 0).

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.
- Mode: The processor affinity feature of IIS 6.0 is available only when running in worker process isolation mode.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**To set processor affinity**

1. Copy the following ADSI script into a text editor or word processor. The script (written in VBScript) enables and configures processor affinity.

```
set appPoolObj=GetObject("IIS://localhost/W3svc/AppPools/app pool name")
' Set the properties. Enable processor affinity for processors 0,1,2,3:
appPoolObj.Put "SMPAffinitized", TRUE
appPoolObj.Put "SMPProcessorAffinityMask", &H0000000F
' Save the property changes in the metabase:
appPoolObj.SetInfo
WScript.Echo "After: " & appPoolObj.SMPAffinitized & ", " &
appPoolObj.SMPProcessorAffinityMask
```

2. In the text editor, change the value of the metabase property **SMPProcessorAffinityMask** to the hexadecimal value that corresponds to the CPU numbers of the CPUs that are used by the worker processes in the application pool.

3. Save the file with extension **.vbs**.
4. In the **Run** dialog box, type **cmd** and then click **OK**.
5. At the command prompt, run the script by typing: **cscript //nologo** *filename***.vbs** where *filename* is the fully qualified path of the script file.
6. To ascertain that the metabase properties have changed, view the metabase properties in the Metabase.xml file to see if the values have been correctly set. To do so, open the Metabase.xml file in Notepad or another text editor. The default path to this file is *systemroot***\system32\inetserv\metabase.xml**.

# Stop the WWW Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Once you have verified that no users are accessing your Web site, you can stop the World Wide Web Publishing Service (WWW service) on your source computer. You can then migrate your Web sites to the target server.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Recommendation**
As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At the command prompt, type **runas /user:***administrative_accountname* **mmc %systemroot%\system32\inetsrv\iis.msc**.

**Procedures**
**To stop the WWW Service in IIS 5.0**

1. Open Administrative Tools, and then click **Services**.
2. In the **Services** pane, double-click **World Wide Web Publishing Service**.
3. On the **Properties** page, on the **General** tab under **Service Status**, click the **Stop** button, and then click **OK**.

**To stop the WWW Service in IIS 4.0**

1. Open Control Panel, and then double-click **Services**.
2. On the **Services** page, click **World Wide Web Publishing Service**, and then click the **Stop** button, and then click **Close**.

# Upgrade FrontPage Extended Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can upgrade Web sites to FrontPage 2002 Server Extensions from Microsoft from any previous version of FrontPage Server Extensions. During the upgrade process, changes are made to your Web site metadata and background structure, and some content is added to your Web site to enable it to work with FrontPage 2002 Server Extensions.

After you have installed FrontPage 2002 Server Extensions, you can upgrade your extended Web sites from the command-line, or by using Microsoft® SharePoint™ Administrator. To upgrade an extended Web site from the command-line, use the upgrade command.

**Procedures**

**To upgrade an extended Web site from the command line**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
2. At the command prompt, change to the FrontPage Extensions directory by typing:

   ```
   cd %CommonProgramFiles%\Microsoft Shared\Web Server Extensions\50\bin.
   ```

3. Then, at the command prompt, type:

   ```
   cwsadm -o upgrade -p Port -m WebSiteURL
   ```

   where *Port* is the optional port number of the site to be upgraded (defaults to port 80), and *WebSiteURL* is the URL for the Web site you are upgrading.

**To upgrade an extended Web site by using SharePoint Administrator**

1. On the server, open **Administrative Tools**, and then click **Microsoft SharePoint Administrator**.
2. In the list of extended Web sites, next to the extended Web site you want to upgrade, click **Upgrade**.
3. In the **Administrator user name** box, type the user name for the administrator of the extended Web site.
4. Click **Submit**.

# View Application Isolation Configuration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before configuring Internet Information Services (IIS) 6.0 to run in worker process isolation mode, you need to document the configuration for existing Web sites and applications running on the source server so you can transfer the settings to the target server.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Iis.msc.

**Procedures**

**To view application isolation settings in IIS 5.0**

1. Open Administrative Tools, and then click **Internet Services Manager**.
2. Expand the local computer, right-click the Web site for which you are documenting the settings, and then click **Properties**.
3. Click the **Home Directory** tab.
4. Under **Application Settings**, note the settings for **Application Protection**. Available settings are:
   - Low (IIS Process)
   - Medium (Pooled)
   - High (Isolated)
5. Click **OK**.

**To view application isolation settings in IIS 4.0**

1. Click **Start**, click **Programs**, click **Windows NT 4.0 Option Pack**, click **Microsoft Internet Information Server**, and then click **Internet Services Manager**.
2. In Internet Services Manager, expand **Internet Information Server**, expand the local computer, right-click the Web site for which you are documenting the settings, and then click **Properties**.
3. On the **Home Directory** tab, under **Application Settings**, note whether the check box **Run in separate memory space (isolated process)** is selected.
4. Click **OK**.

# View Web Site and Application Process Identities

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before configuring Internet Information Services (IIS) 6.0 to run in worker process isolation mode, you need to document the process identity used by the Web site and application. After configuring IIS 6.0 to run in worker process isolation mode, you will use the documented process identities to configure the application pools in worker process isolation mode.

**Requirements**

- Credentials: Membership in the Administrators group on the local computer.
- Tools: Internet Service Manager (ISM) for IIS 4.0 and IIS 5.0, and IIS Manager for IIS 6.0.

**Procedures**

**To view Web site and application process identities in IIS 6.0 running in IIS 5.0 Isolation Mode**

1. In the **Run** dialog box, type **mmc**, and then click **OK**.
2. If you do not have the Component Services Snap-in installed, in MMC, on the **File** menu, click **Add/Remove Snap-ins**, and then click **Add**. Under **Available Standalone Snap-ins**, click **Component Services**, and then click **Add**.
3. In MMC, in the console tree, expand **Component Services**, expand **Computers**, expand **My Computer**, and then click **COM+ Applications**.
4. Right-click *WebSiteName* and then click **Properties** (where *WebSiteName* is the name of the Web site or application).
5. Click the **Identity** tab.
6. Document the identity for the Web site or application found in the **User** text box.

**To view Web site and application process identities in IIS 5.0**

1. In the **Run** dialog box, type **mmc**, and then click **OK**.
2. If you do not have the Component Services Snap-in installed, in the **File** menu click **Add/Remove Snap-ins**, and then click **Add**. Under **Add Standalone Snap-ins**, click **Component Services**, and then click **Add**. Click **Close** and then click **OK**.
3. In MMC, in the console tree, expand **Computers**, expand **My Computer** and expand **COM+ Applications**.
4. Right-click **IIS-{***WebSiteName***//Root}** and then click **Properties** (where *WebSiteName* is the name of the Web site or application).
5. Click the **Identity** tab.
6. Document the identity for the Web site or application found in the **User** text box.

**To view Web site and application process identities in IIS 4.0**

1. Click **Start**, click **Programs**, click **Windows NT 4.0 Option Pack**, click **Microsoft Internet Information Server**, and then click **Internet Services Manager**.
2. In the console tree, expand **Microsoft Transaction Server**, expand **Computers**, expand **My Computer**, expand **Packages installed**, right-click *WebSiteName* and then click **Properties** (where *WebSiteName* is the name of the Web site or application).
3. Click the **Identity** tab.
4. Document the identity for the Web site or application found in the **User** text box.

# Appendix B: Changes to Metabase Properties in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

There are IIS 5.0 metabase properties that are no longer supported in IIS 6.0. When upgrading or migrating from an earlier version of IIS, these properties are still in the metabase, but IIS 6.0 ignores the settings. On a new installation of IIS 6.0, these properties have been removed from the metabase, and they are not available even when IIS 6.0 is running in IIS 5.0 isolation mode.

There is one IIS 5.0 metabase property — **CPUResetInterval** — that still exists in the IIS 6.0 metabase, but it behaves differently in IIS 6.0.

Table 1.1 lists metabase properties that have changed in IIS 6.0. If your Web sites, applications, or setup programs refer to any of these IIS metabase properties, follow the recommendations described in the table to accommodate these changes and ensure compatibility with IIS 6.0.

**Table 1.1 IIS Metabase Properties That Have Changed Since IIS 5.0**

| IIS Metabase Properties | Recommendation |
|---|---|
| • **AspThreadGateEnabled**<br>• **AspThreadGateLoadHigh**<br>• **AspThreadGateLoadLow**<br>• **AspThreadGateSleepMax**<br>• **AspThreadGateTimeSlice** | In IIS 5.0, these metabase properties are provided for the configuration of performance features.<br>In IIS 6.0, these performance features are provided by other means. As a result, IIS 6.0 no longer uses these metabase properties.<br>Remove any references to these metabase properties from Web sites, applications, or setup programs.<br>For more information about improving performance in IIS 6.0, see Web Server Scalability. |

| | |
|---|---|
| • **CPUAppEnabled**<br>• **CPUCGIEnabled**<br>• **CPUCGILimit**<br>• **CPUEnableActiveProcs**<br>• **CPUEnableAllProcLogging**<br>• **CPUEnableAppLogging**<br>• **CPUEnableCGILogging**<br>• **CPUEnableEvent**<br>• **CPUEnableKernelTime**<br>• **CPUEnablePageFaults**<br>• **CPUEnableProcType**<br>• **CPUEnableTerminatedProcs**<br>• **CPUEnableTotalProcs**<br>• **CPUEnableUserTime**<br>• **CPULimitLogEvent**<br>• **CPULimitPause**<br>• **CPULimitPriority**<br>• **CPULimitProcStop**<br>• **CPULimitsEnabled**<br>• **CPULoggingInterval**<br>• **CPULoggingMask**<br>• **CPULoggingOptions** | In IIS 5.0, these metabase properties are provided for the configuration of processor (CPU) management features.<br><br>In IIS 6.0, processor management (CPU throttling) is provided by other means. As a result, IIS 6.0 no longer uses these metabase properties.<br><br>Remove any references to these IIS metabase properties from Web sites, applications, or setup programs.<br><br>The following metabase properties reference and implement CPU throttling in IIS 6.0:<br>• **CPULimit**<br>• **CPUResetInterval**<br>• **CPUAction**<br><br>For more information about CPU throttling in IIS 6.0, see the CPULimit Metabase Property, CPUResetInterval Metabase Property, and CPUAction Metabase Property metabase properties. |
| **CPUResetInterval** | In IIS 5.0, this property enables the monitoring of processor utilization at the Web site level. IIS 5.0 was unable to monitor individual processes for a Web site, such as multiple processes in an ASP.NET Web garden.<br><br>In IIS 6.0, this property enables the monitoring of processor utilization at the worker process level. You cannot monitor the processor utilization for individual Web sites unless there is only one Web site in the application pool.<br><br>For more information about how to configure this value for worker processes in an application pool, see the CPUResetInterval Metabase Property metabase property. |

| | |
|---|---|
| **DisableSocketPooling** | In IIS 4.0 and IIS 5.0, Windows Sockets (WinSock) listens for HTTP requests that use Transmission Control Protocol (TCP). WinSock uses the concept of *sockets* to provide TCP connectivity. As a result, IIS 4.0 and IIS 5.0 are bound to the scalability constraints of using sockets through WinSock. In IIS 5.0, socket pooling addresses these scalability problems, in cases where a large number of Web sites were configured with individual IP addresses.<br><br>In IIS 5.0, socket pooling allows for socket resources to be shared between multiple Web sites and thus provides significant improvements — up to two to three times the scaling capacity of earlier versions of IIS.<br><br>In IIS 6.0, HTTP.sys is responsible for listening for HTTP requests and provides similar functionality to the **DisableSocketPooling** metabase property.<br><br>To configure HTTP.sys, use Httpcfg.exe.<br><br>**Note**<br>On new installations of Windows Server 2003 and IIS 6.0, the **DisableSocketPooling** metabase property still exists; however, IIS 6.0 ignores the property. |
| **HcMimeType** | In IIS 5.0, this metabase property indicates which Multipurpose Internet Mail Extensions (MIME) types are supported by the compression scheme.<br><br>IIS 6.0, HTTP compression is performed in a different way, and this metabase property is no longer used.<br><br>Remove any references to this metabase property from Web sites, applications, or setup programs. |
| **PutReadSize** | In IIS 5.0, this metabase property is provided to support the Web Distributed Authoring and Versioning (WebDAV) component, but is untested and unsupported.<br><br>In IIS 6.0, this functionality is integrated with the WebDAV component. As a result, IIS 6.0 no longer uses this metabase property.<br><br>Remove any references to this metabase property from Web sites, applications, or setup programs.<br><br>For more information about WebDAV in IIS 6.0, see About WebDAV. |
| **UNCAuthenticationPassthrough** | In IIS 5.0, this metabase property enables pass-through user authentication for Universal Naming Convention (UNC) virtual root access, which is for authentication schemes that support delegation.<br><br>In IIS 6.0, pass-through authentication occurs automatically when the **UNCUserName** and **UNCPassword** metabase properties are not specified. As a result, IIS 6.0 no longer uses this metabase property.<br><br>Remove any references to this metabase property from Web sites, applications, or setup programs.<br><br>For more information, see the UNCUserName Metabase Property and UNCPassword Metabase Property metabase properties. |

# Appendix C: Choosing a Cluster Technology for IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Windows Server 2003 supports three different cluster technologies: Server cluster, which is the Windows Server 2003 name for the Microsoft technology first made available as Microsoft Cluster Server (MSCS) in Windows NT Server 4.0, Enterprise Edition; Network Load Balancing (NLB); and Component Load Balancing (CLB). These Microsoft clustering technologies can be key ingredients in improving availability, reliability and scalability. The following table compares the purposes and uses of the three cluster technologies.

**Table 1.1 Cluster Technologies Supported in Windows Server 2003**

| Cluster technology | Purpose | Suitable applications |
|---|---|---|
| Server cluster | Provide high availability by making application software and data available on several servers linked together in a cluster configuration. If one server stops functioning, a process called *failover* automatically shifts the workload of the failed server to another server in the cluster. The failover process is designed to ensure continuous availability of critical applications and data. | • *Stateful* applications have long-running in-memory state, or they have large, frequently updated data states. A database such as Microsoft® SQL Server™ is an example of a stateful application. |
| NLB | • Address bottlenecks caused by front-end Web services and Web servers.<br>• Support load balancing of Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Generic Routing Encapsulation (GRE) traffic requests. | *Stateless* applications do not have long-running in-memory state. A stateless application treats each client request as an independent operation, and therefore it can load-balance each request independently. Stateless applications often have read-only data or data that changes infrequently. Web front-end servers, virtual private networks (VPNs), and File Transfer Protocol (FTP) servers typically use NLB. Microsoft recommends using NLB with IIS 6.0. |
| CLB | Perform dynamic load balancing, over multiple nodes, of middle-tier application components that use COM+. | Applications that use COM+ components. |

**Limitations of clusters**

While a well-designed solution can guard against application failure, system failure and site failure, cluster technologies do have the following limitations:

- Cluster technologies depend on compatible applications and services to respond appropriately when failure occurs.
- Clusters can be designed to handle failure, but they are not meant to protect user data. The cluster by itself does not guard against loss of a user's work. Typically, the recovery of lost work is handled by the application software; the application software must be designed to recover the user's work, or it must be designed in such a way that the user's session state can be maintained if a failure occurs.
- Cluster technology cannot protect against failures caused by viruses, software corruption or human error. To protect against these types of problems, organizations need solid data protection and recovery plans.

For more information on cluster technologies available in Window Server 2003, see An Overview of Windows Clustering Technologies: Server Clusters and Network Load Balancing.

# Clustering IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 Web servers can be made highly available by failover using a Server cluster; however, Microsoft recommends that you use Network Load Balancing (NLB). In addition to providing high availability for IIS, NLB is a good cluster solution for a Web farm scenario because it also offers load balancing capabilities. NLB is generally less expensive to implement than Server cluster, which has additional hardware requirements. Windows Server 2003 Web Edition supports NLB.

**Upgrading to IIS 60 from an IIS 5.0 Cluster**
In Windows 2000, IIS 5.0 Web sites and FTP services used the IIS Server Instance resource type to achieve high availability with Microsoft Cluster Service (MSCS).

In Windows Server 2003, the IIS Server Instance resource type is now a *Generic Script* resource.This new resource type allows you to implement start/stop and health monitoring functions as scripts rather than as C or C++ programs. To manage IIS in a Windows Server 2003 cluster, you must use this Generic Script resource.

Windows Server 2003 provides a script called Iis_switch.vbs, which is located by default in the *systemroot*\system32\inetsrv directory. Iis_switch.vbs allows you to upgrade to IIS 6.0 from a Windows 2000 Server cluster with IIS 5.0 installed. First, upgrade all nodes except the node containing the IIS instance resource from Windows 2000 to Windows Server 2003, Enterprise Edition. Then, you can either use the script or perform the upgrade manually.

For step-by-step instructions about upgrading IIS from a clustered Windows 2000 Server configuration, see Server Clusters: Rolling Upgrades—Upgrading to Windows Server 2003.

# FTP Cluster Deployments

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Depending on the environment in which the FTP servers reside, you can use either Server clusters or NLB clusters to provide highly available FTP servers. The following table compares Server clusters and NLB.

| Server cluster | NLB |
| --- | --- |
| • Appropriate for FTP sites with high update rates, or in situations where you want to have a single copy of the FTP content.<br>• Uses a shared disk; does not require replication.<br>• Can be configured as active/passive only, which provides high availability but does not provide load balancing. | • Appropriate for FTP sites that are mainly read-only and used mostly for downloading content.<br>• Requires replication.<br>• Single-affinity mode is strongly recommended.<br>• Recommended for clustering FTP servers, instead of Server cluster. |

For more information about the process of implementing FTP in a Server cluster, see Checklist: Creating a clustered IIS Web or FTP service in the Windows Server 2003 Operations Guide. For an example of an FTP cluster configuration, see Example: Controlling the Distribution of Client Traffic within the Cluster.

# Additional Resources for Clustering IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following additional resources can help you configure Network Load Balancing (NLB) for IIS 6.0.

- For an introduction to NLB, see What is Network Load Balancing?
- For comprehensive information about deploying IIS in an NLB configuration, see Designing Network Load Balancing and Deploying Network Load Balancing in the Windows Server 2003 Deployment Kit.
- For information about scalability and performance issues for IIS 6.0 in an NLB configuration, see Improving IIS 6.0 Scalability and Availability with Network Load Balancing.
- For information about using Kerberos authentication with load-balanced Web sites, see Kerberos Authentication for load-balanced Web sites.

# Appendix D: Running 32-bit Applications on 64-bit Windows

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Windows Server 2003™, Service Pack 1 enables IIS 6.0 to run 32-bit Web applications on 64-bit Windows using the Windows-32-on-Windows-64 (WOW64) compatibility layer. IIS 6.0 using WOW64 is intended to run 32-bit personal productivity applications needed by software developers and administrators, including 32-bit Internet Information Services (IIS) Web applications.

On 64-bit Windows, 32-bit processes cannot load 64-bit DLLs, and 64-bit processes cannot load 32-bit DLLs. If you plan to run 32-bit applications on 64-bit Windows, you must configure IIS to create 32-bit worker processes. Once you have configured IIS to create 32-bit worker processes, you can run the following types of IIS applications on 64-bit Windows:

- Internet Server API (ISAPI) extensions
- ISAPI filters
- Active Server Page (ASP) applications (specifically, scripts calling COM objects where the COM object can be 32-bit or 64-bit)
- ASP.NET applications

IIS can, by default, launch Common Gateway Interface (CGI) applications on 64-bit Windows, because CGI applications run in a separate process.

**Configuring IIS to run 32-bit Web applications on 64-bit Windows**
Before you configure IIS to run 32-bit applications on 64-bit Windows, note the following:

- IIS only supports 32bit worker processes in Worker Process Isolation mode on 64-bit Windows.
- On 64-bit Windows, the World Wide Web Publishing service can run 32-bit and 64-bit worker processes. Other IIS services like the IIS Admin service, the SMTP service, the NNTP service, and the FTP service run 64-bit processes only.
- On 64-bit Windows, the World Wide Web Publishing service does not support running 32-bit and 64-bit worker processes concurrently on the same server.

**To enable IIS 6.0 to run 32-bit Web applications on 64-bit Windows**

1. Open a command prompt and navigate to the *%windir%*\Inetpub\AdminScripts directory.
2. Type the following:

```
cscript.exe adsutil.vbs set W3SVC/AppPools/Enable32BitAppOnWin64 1
```

3. Press ENTER.

# Appendix E: Unattended Setup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0 is not installed by default when you install the Microsoft® Windows® Server 2003 operating system. To simplify the setup process on multiple computers, you can set up IIS 6.0 while the server is unattended. To do so, you can either create an answer file that you can use to install IIS 6.0 at the same time you install Windows Server 2003, or you can perform an unattended installation of IIS 6.0 after you install Windows Server 2003.

In This Appendix

Creating an Answer File

Installing IIS 6.0 with the Operating System

Installing IIS 6.0 After the Operating System

Related Resources

- For information about configuring Web, File Transfer Protocol (FTP), Network News Transfer Protocol (NNTP), and Simple Mail Transfer Protocol (SMTP) sites and services, see Configuring Internet Sites and Services.

# Creating an Answer File

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To set up unattended installation for IIS 6.0, you must create an answer file containing parameters with appropriate settings, depending on how you want to customize your installation of IIS. You have two options for unattended installation of IIS 6.0:

- Install and set up IIS 6.0 at the same time that you install the Windows Server 2003 operating system.
- Install IIS 6.0 after the Windows Server 2003 operating system has been installed.

If you install IIS 6.0 simultaneously with the operating system, you need to create a complete answer file and include a [Components] section, as well as an optional [InternetServer] section that contains IIS–related parameters. Use the Winnt32.exe command-line utility to run your answer file. For more information about using Winnt32.exe, see Installing IIS 6.0 with the Operating System. For more information about creating and using a Windows unattended setup answer file, see Planning for Unattended Setup in Help and Support Center for Microsoft Windows Server 2003.

If you install IIS 6.0 after the Windows Server 2003 operating system has been installed, you can create an answer file containing only the [Components] and optional [InternetServer] sections. Use the Sysocmgr.exe command-line utility to run this answer file. For more information about using Sysocmgr.exe, see Installing IIS 6.0 After the Operating System.

Table B.1 lists the unattended installation parameters for IIS that you include in the [Components] section of the answer file. The default setting for all of the following parameters is **off** unless a component is required by some other component that you are installing. For example, if you are installing the World Wide Web Publishing Service (WWW service), then Common Files, IIS Manager, and network COM+ access are turned on by default.

**Table B.1 IIS Unattended Installation Parameters**

| Parameter | Description |
|---|---|
| *aspnet*= on/off | Enable ASP.NET . |
| *complusnetwork* = on/off | Enable network COM+ access. |
| *dtcnetwork* = on/off | Enable network Distributed Transaction Coordinator (DTC) network access. |
| *bitsserverextensionsisapi* = on/off | Enable the BITS Server Extensions ISAPI. |
| *bitsserverextensionsmanager* = on/off | Enable the BITS Server Extensions snap-in. |
| *iis_common* = on/off | Enable Common Files. |
| *iis_ftp* = on/off | Enable the File Transfer Protocol (FTP) service. |
| *fp_extensions* = on/off | Enable FrontPage® 2002 Server Extensions from Microsoft. |
| *iis_inetmgr* = on/off | Enable IIS Manager. |
| *iis_nntp* = on/off | Enable the Network News Transfer Protocol (NNTP) service. |
| *iis_smtp* = on/off | Enable the Simple Mail Transfer Protocol (SMTP) service. |
| *iis_asp* = on/off | Enable Active Server Pages (ASP). |

| | |
|---|---|
| *iis_internetdataconnector* = on/off | Enable Internet Data Connector (IDC). |
| *sakit_web* = on/off | Enable Remote Administration (HTML).<br>You cannot perform an unattended installation of this feature simultaneously with performing an unattended installation of the operating system.<br>You can perform an unattended installation of this feature if you install IIS after you have installed the operating system by using the Sysocmgr.exe command-line utility.<br>This feature is not available on Itanium-based computers running Windows XP 64-Bit Edition; the 64-bit version of Windows Server 2003, Enterprise Edition; or the 64-bit version of Windows Server 2003, Datacenter Edition. |
| *tswebclient* = on/off | Enable remote desktop Web connection. |
| *iis_serversideincludes* = on/off | Enable server-side includes (SSI). |
| *iis_webdav* = on/off | Enable WebDAV publishing. |
| *iis_www* = on/off | Enable the WWW service. |
| *appsrv_console* = on/off | Enable Application Server Console.<br>This feature is not available on Itanium-based computers running Windows XP 64-Bit Edition. |
| *inetprint* = on/off | Enable Internet printing. |

Table B.2 lists the optional IIS unattended installation parameters. You include these parameters in the [InternetServer] section.

**Table B.2 Optional IIS Unattended Installation Parameters**

| Parameter | Description |
|---|---|
| *SvcManualStart* = **www** and/or **ftp** | If WWW or FTP is present, those services will be set to manual start. If you specify both **www** and **ftp**, separate them with a comma (**,**). When installed, the default service state is to automatically start. |
| *PathFTPRoot* = path to FTProot | Alternate path to which the FTP service is installed. The default is *systemroot*\Inetpub\Ftproot. |
| *PathWWWRoot* = path to WWWroot | Alternate path to which the WWW service is installed. The default is *systemroot*\Inetpub\Wwwroot. |
| *ApplicationDependency* = *ApplicationName, GroupID, GroupID* | Adds application and dependencies.<br>For example:<br>ApplicationDependency = CommerceServer, ASP, IndexingService<br>ApplicationDependency = ExchangeServer, ASP60<br>ApplicationDependency = MyApp, ASP.NET v1.1.4322<br>This example makes Commerce Server dependent on ASP and the Indexing Service, Exchange Server dependent on ASP, and MyApp dependent on ASP.NET v1.1.4322. |

| | |
|---|---|
| *ExtensionFile* = <*Enabled Flag*>, <*File Path*>[ ,<*UI Deletable Flag*> [ , <*Group ID*> [ ,<*Description*> ] ] ] | Specifies an extension file. The *UI Deletable Flag*, *Group ID*, and *Description* settings are all optional. However, if you specify any of the optional settings, you must include all of the previous settings.<br><br>The parameters represent the following:<br><br>*Enabled Flag* is set to **0** (zero) for disabled and **1** for enabled.<br><br>*File Path*is the path to the file, and it can contain environment variables.<br><br>*UI Deletable Flag* specifies whether the item can be deleted using IIS Manager. Set to **0** for not deletable and **1** for deletable.<br><br>*Group ID* allows you to group different dynamic-link libraries (DLLs) and Common Gateway Interfaces (CGIs) and have dependencies for the applications.<br><br>*Description* is the description of the extension file.<br><br>For example:<br><br>ExtensionFile = 1, windir\MyISAPIS\Myisapi.dll, 1, MyGroup, My Custom ISAPI<br><br>–Or–<br><br>ExtensionFile = 1, windir\MyCGIs\Mycgi.cgi, 1, MyGroup, My Custom CGI |
| *DisableWebServiceOnUpgrade* = True/False | The default for this parameter is **True** for upgrades from the Microsoft® Windows® 2000 Server operating system, unless you have run the IIS Lockdown Tool or added the registry entry RetainW3SVCStatus to the registry. Please see Caution below.<br><br>For more information about when the WWW service is disabled after an upgrade has been performed, see Upgrading an IIS Server to IIS 6.0. |

**Caution**

Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference in the Microsoft Windows Server 2003 Deployment Kit.

Related Information

For more information about the registry, see the Registry Reference in the Microsoft Windows Server 2003 Deployment Kit.

# Installing IIS 6.0 with the Operating System

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use the Winnt32.exe command-line utility to install IIS 6.0 by using an unattended answer file that includes a [Components] section with the unattended setup parameters listed in Table B.1 and — optionally — an [InternetServer] section with the parameters listed in Table B.2.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyMachine******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**To use the Winnt32.exe command-line utility**

1. At a command prompt, type:

   ```
   winnt32 /unattend:AnsweFile /s:InstallSource
   ```

   where *AnswerFile* is a file that contains answers to questions that should be automated during installation, and *InstallSource* is the location of the Windows 2003 installation files.

**To use an x86-based computer to view the command options available for Winnt32.exe**

1. On a computer running Microsoft® Windows® 95, Windows® 98, Windows® Millennium Edition, Microsoft® Windows NT® , Windows 2000, Windows XP, or Windows Server 2003 operating systems, insert the Setup CD for Windows Server 2003 in the CD-ROM drive.
2. On the **Start** menu, click **Run**.
3. In the **Open** box, type

   ```
   cmd
   ```

4. Click **OK**.
5. At the command prompt, change to the CD-ROM drive, change to the I386 directory, type

   ```
   winnt32 /?
   ```

   and then press ENTER.

**To use an Itanium-based computer to view the command options available for Winnt32.exe**

1. On an Itanium-based computer running Windows XP 64-Bit Edition; the 64-bit version of Windows Server 2003, Enterprise Edition; or the 64-bit version of Windows Server 2003, Datacenter Edition, insert the Setup CD for the 64-bit version of the product into the CD-ROM drive.
2. On the **Start** menu, click **Run**.
3. In the **Open** box, type

   ```
   cmd
   ```

4. Click **OK**.
5. At the command prompt, change to the CD-ROM drive, change to the IA64 directory, type

   ```
   winnt32 /?
   ```

   and then press ENTER.

# Installing IIS 6.0 After the Operating System

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To install IIS by using scripts after you have installed and set up Windows Server 2003, use the Sysocmgr.exe command-line utility with an unattended answer file that includes a [Components] section with the unattended setup parameters listed in Table B.1 and, optionally, an [InternetServer] section with the parameters listed in Table B.2.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User***:***MyMachine******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**To install IIS 6.0 by using scripts after the operating system has been installed**

1. On the **Start** menu, click **Run**.
2. In the **Open** box, type

   `cmd`

3. Click **OK**.
4. At the command prompt, type

   **Sysocmgr.exe /i:sysoc.inf /u:%** *PathToUnattendFile* **%**

   and then press ENTER.

To view Sysocmgr.exe Help, type `Sysocmgr.exe/?` at the command prompt.

and then press ENTER .

# Additional Resources_IIS_SP1_Dep

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this chapter.

**Related Information**

- Ensuring Application Availability for information about configuring your Web server to improve the availability of your ASP.NET applications.
- Migrating Machine.config Attributes to IIS 6.0 Metabase Property Settings for information about how to migrate attributes in the Machine.config file to their equivalent IIS 6.0 metabase property settings.
- Securing Web Sites and Applications for information about configuring your Web server to improve the security of your ASP.NET applications.
- The Microsoft Knowledge Base article 329290, How To: Use the ASP.NET Utility to Encrypt Credentials and Session State Connection Strings, for information about configuring ASP.NET to store the session connection strings in the registry.
- How to allocate handles and connect to SQL Server (ODBC), for information about creating ODBC data connection strings for Microsoft SQL Server.

**Related IIS 6.0 Operations Guide Topics**

- ASP.NET Configuration Files Format for information about configuring ASP.NET applications when IIS 6.0 is configured to run in IIS 5.0 isolation mode.
- Troubleshooting in IIS 6.0 for information about how to troubleshoot Web sites that fail to start.
- Troubleshooting an ASP.NET Installation for information about how to troubleshoot ASP.NET-specific problems.

**Related Windows Server 2003 Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the**Search in title only** check box.

- "Features unavailable on 64-bit versions of the Windows Server 2003 family" in Help and Support Center for Windows Server 2003 for information about features, such as ASP.NET, that are not supported on 64-bit versions of Windows Server 2003.
- "Managing resource records" in Help and Support Center for Windows Server 2003 for information about how to create DNS entries for your applications.

# Windows Server 2003 Operations

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft® Windows Server™ 2003 Operations documentation provides step-by-step, task-oriented information for Windows Server 2003 technologies. This set of Operations documentation is designed to assist information technology (IT) operators and administrators with prescriptive guidance for how to operate, manage, and troubleshoot Windows Server 2003 technologies.

**In this guide**


- Active Directory Operations Guide
- DNS Operations Guide
- Group Policy Operations Guide
- Security Technologies Operations Guide
- Server Cluster Operations Guide
- Terminal Server Licensing Operations Guide
- Windows Firewall Operations Guide

# Active Directory Operations Guide

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Active Directory Operations Guide provides administering and troubleshooting information for Active Directory™ directory service technologies in the Microsoft® Windows Server™ 2003 and Windows Server 2003 with Service Pack 1 (SP1) operating systems.

You cannot install Active Directory on a server running Windows Server 2003, Web Edition, but you can join the server to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition, on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=9253).

| **Note** |
| --- |
| The Windows Server 2003 Active Directory Operations Guide is also available as a downloadable document on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=63079). |

**In this guide**


- Administering Active Directory Operations
- Troubleshooting Active Directory Operations

# Troubleshooting Active Directory Operations

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This Active Directory Troubleshooting guide provides troubleshooting information for Active Directory in the Microsoft Windows Server 2003 and Windows Server 2003 with Service Pack 1 (SP1) operating systems.

**In this guide**


- Configuring a Computer for Troubleshooting Active Directory
- Troubleshooting Active Directory Replication Problems
- Additional Resources for Troubleshooting Active Directory


This initial release of the Active Directory Troubleshooting guide includes troubleshooting recommendations and procedures for diagnosing and fixing problems that may occur with Active Directory replication in a Windows Server 2003 forest. This content focuses primarily on responses to Directory Service event log messages and tool-based error messages that might be reported by the Repadmin.exe and Dcdiag.exe tools, which are available in Windows Support Tools.

Installation of Windows Server 2003 with SP1 is encouraged for improved diagnostic support in both Windows Support Tools, which you must install separately, and the Ntdsutil.exe administrative command-line tool, which is included with the operating system.


**Acknowledgments**

Key Technical Reviewers: Arren Conner, Gregory Johnson, Rob Kochman, Ajit Krishnan, Dave Tesar

# Configuring a Computer for Troubleshooting Active Directory

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can use advanced troubleshooting techniques to identify and fix Active Directory problems, you must configure your computer for troubleshooting and have a basic understanding of Windows Server 2003 troubleshooting concepts, procedures, and tools. For information about monitoring tools for Windows Server 2003, see Monitoring and Status Tools on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=59526).

**Configuration Tasks for Troubleshooting**
To configure your computer for troubleshooting, perform the following tasks:

Install Windows Server 2003 SP1

Install Windows Support Tools

Install Network Monitor

Set logging levels

**Install Windows Server 2003 SP1**
If possible, upgrade domain controllers to Windows Server 2003 Service Pack 1 (SP1). To install this service pack, go to the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=9999) and follow instructions for downloading the service pack.

The advantages of running Windows Server 2003 with SP1 with regard to troubleshooting include enhancements to the Ntdsutil command-line tool. Ntdsutil.exe has new functionality that makes it easier to remove domain controller metadata and to authoritatively restore directory objects.

**Install Windows Support Tools**
For improved diagnostic support, install Windows Support Tools that ship with Windows Server 2003 SP1. The SP1 version of Windows Support Tools includes enhanced versions of the Dcdiag.exe and Repadmin.exe tools. The Dcdiag.exe command-line tool now provides new reporting on the overall health of replication with respect to Active Directory security, as well as new Domain Name System (DNS) diagnostic tests. You can use Repadmin.exe to manage replication consistency settings on multiple domain controllers instead of editing the registry on individual computers.

Make sure that the SP1 version of Windows Support Tools is installed on all domain controllers that are running Windows Server 2003 with SP1.

**Options for Running SP1 Windows Support Tools**
You can run Windows Support Tools that ship with Windows Server 2003 SP1 on computers running the following operating systems:

- Windows Server 2003 with SP1
- Windows Server 2003 without SP1

You can also run some tools, such as Repadmin.exe and Dcdiag.exe, on computers running Windows XP Professional, Windows XP Professional with SP1, or Windows XP Professional with Service Pack 2 (SP2). Options for other tools vary by tool. In this guide, the operating system that is required for running a tool is specified as a prerequisite for each procedure.

**Options for Installing SP1 Windows Support Tools**
The SP1 version of Windows Support Tools can be installed as an .msi package only on computers running Windows Server 2003 with SP1. To run Repadmin and Dcdiag from computers running Windows Server 2003 without SP1 or from computers running Windows XP Professional, you must copy the respective executable files to those computers.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Builtin Administrators group.
- Operating system: Windows Server 2003 with SP1. You cannot use suptools.msi to install the SP1 version of Windows Support Tools on a computer that is not running Windows Server 2003 with SP1.

**To install Windows Support Tools**

1. Insert the Windows CD into your CD-ROM drive.
2. If you are prompted to reinstall Windows, click **No**.
3. When the **Welcome** screen appears, click **Perform additional tasks**, and then click **Browse this CD**.
4. Go to the \Support\Tools folder. For complete setup information, see the Readme.htm file in this folder.
5. Double-click **suptools.msi**.
6. Follow the instructions that appear on your screen.

**Install Network Monitor**

Use Network Monitor to troubleshoot connectivity issues by tracing network traffic between computers. For information about installing and using Network Monitor, see Network Monitor on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42987).

**Set Logging Levels**

If the information that you receive in the Directory Service log in Event Viewer is not sufficient for troubleshooting, raise the logging levels by using the appropriate registry entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics**.

By default, the logging levels for all entries are set to **0**, which provides the minimum amount of information. The highest logging level is **5**. Increasing the level for an entry causes additional events to be logged in the Directory Service event log. The following diagram shows the diagnostic entries that are available.




Use the following procedure to change the logging level for a diagnostic entry.

| **Caution** |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

**Requirements**

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the domain of the domain controller on which you are setting the logging level.

- Tools: Regedit.exe

**To change the logging level for a diagnostic entry**

1. Click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Navigate to the entry for which you want to set logging in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics**.
3. Double-click the entry, and for the **Base** click **Decimal**.
4. In the **Value** data box, type an integer from **0** through **5**, and then click **OK**.

# Troubleshooting Active Directory Replication Problems

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory replication problems can have several different sources. For example, Domain Name System (DNS) problems, networking issues, or security problems can all cause Active Directory replication to fail.

Inbound or outbound replication failure causes Active Directory objects that represent the replication topology, replication schedule, domain controllers, users, computers, passwords, security groups, group memberships, and Group Policy to be inconsistent between domain controllers. Directory inconsistency causes either operational failures or inconsistent results, depending on the domain controller that is contacted for the operation at hand. Active Directory depends on network connectivity, name resolution, authentication and authorization, the directory database, the replication topology, and the replication engine. When the root cause of a replication problem is not immediately obvious, determining the cause among the many possible causes requires systematic elimination of probable causes.

**Event and Tool Solution Recommendations**
Ideally, the red (Error) and yellow (Warning) events in the Directory Service event log suggest the specific constraint that is causing replication failure on the source or destination domain controller. If the event message suggests steps for a solution, try the steps listed in the event. The Repadmin tool and other diagnostic tools also provide information that can help you resolve replication failures.

**Ruling Out the Obvious**
Sometimes replication errors occur because of intentional disruptions. For example, when you troubleshoot Active Directory replication problems, rule out intentional disconnections and hardware failures or upgrades first.

**Intentional Disconnections**
If replication errors are reported by a domain controller that is attempting replication with a domain controller that has been built in a staging site and is currently offline awaiting its deployment in the final production site (remote), you can account for those errors. To avoid separating a domain controller from the replication topology for extended periods, which causes continuous errors until the domain controller is reconnected, consider adding such computers initially as member servers and using the install-from-media method to install Active Directory. You can back up an up-to-date domain controller to removable media (CD/DVD or other media) and ship the media to the destination site. Then, you can use the media to promote the domain controllers at the site, without requiring replication. For more information about installing from media, see Installing a Domain Controller in an Existing Domain Using Restored Backup Media.

**Hardware Failures or Upgrades**
If replication problems occur as a result of hardware failure (for example, failure of the motherboard, disk subsystem, or hard drive), notify the server owner so that the hardware problem can be resolved.

Periodic hardware upgrades can also cause domain controllers to be out of service. Ensure that your server owners have a good system of communicating such outages in advance.

**Correct Response to Any Outdated Server Running Windows 2000 Server**
If a domain controller running Windows 2000 Server has failed for longer than the number of days in the tombstone lifetime, the solution is always the same:

1. Move the server from the corporate network to a private network.
2. Either forcefully remove Active Directory or reinstall the operating system.
3. Remove the server metadata from Active Directory so that the server object cannot be revived.

| **Note** |
| --- |
| By default, NTDS Settings objects that are deleted are revived automatically for a period of 14 days. Therefore, if you do not remove server metadata (use Ntdsutil to perform metadata cleanup), the server metadata is reinstated in the directory, which prompts replication attempts to occur. In this case, errors will be logged persistently as a result of the inability to replicate with the missing domain controller. |

**Root Causes**
If you rule out intentional disconnections, hardware failures, and outdated Windows 2000 domain controllers, the remainder of replication problems almost always have one of the following root causes:

- Network connectivity: The network connection might be unavailable or network settings are not configured properly.
- Name resolution: DNS misconfigurations are a common cause for replication failures.
- Authentication and authorization: Authentication and authorization problems cause "Access denied" errors when a domain controller tries to connect to its replication partner.
- Directory database (store): The directory database might not be able to process transactions fast enough to keep up with replication timeouts.
- Replication engine: If intersite replication schedules are too short, replication queues might be too large to process in the time that is required by the outbound replication schedule. In this case, replication of some changes can be stalled indefinitely — potentially, long enough to exceed the tombstone lifetime.
- Replication topology: Domain controllers must have intersite links in Active Directory that map to real wide area network (WAN) or virtual private network (VPN) connections. If you create objects in Active Directory for the replication topology that are not supported by the actual site topology of your network, replication that requires the misconfigured topology fails.

**General Approach to Fixing Problems**
Use the following general approach to fixing replication problems:

1. Monitor replication health daily, or use Repadmin.exe to retrieve replication status daily.
2. Attempt to resolve any reported failure in a timely manner by using the methods described in event messages and this guide. If software might be causing the problem, uninstall the software before you continue with other solutions.
3. If the problem that is causing replication to fail cannot be resolved by any known methods, remove Active Directory from the server and then reinstall Active Directory. For more information about reinstalling Active Directory, see Decommissioning a Domain Controller.
4. If Active Directory cannot be removed normally while connected to the network, use one of the following methods to resolve the problem:

   - Force Active Directory removal in Directory Services Restore Mode, clean up server metadata, and then reinstall Active Directory.
   - Reinstall the operating system, and rebuild the domain controller.

   For more information about forcing Active Directory removal, see Forcing the Removal of a Domain Controller.

**Monitoring Replication Health**
Monitoring for replication failures is critical to being able to solve replication problems quickly and effectively. Use one of the following methods to monitor replication health:

| **Note** |
| --- |
| For detailed information on how to use Repadmin, see Monitoring and Troubleshooting Active Directory Replication Using Repadmin (http://go.microsoft.com/fwlink/?LinkId=122830). |

- Use a monitoring application that you set to capture and report specific errors and events on a daily basis.

- Use the Repadmin tool to retrieve replication status daily.

**Using a Monitoring Application to Monitor Replication Health**

For all domain controllers in a forest, monitor replication health on a daily basis by using Microsoft Operations Manager (MOM) or an equivalent monitoring application. For information about using MOM to monitor Active Directory, see Active Directory Management Pack Technical Reference for MOM 2005 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41369).

**Using Repadmin to Retrieve Replication Status**

Replication status is an important way for you to evaluate the status of the directory service. If replication is working without errors, you know the domain controllers that are online. You also know that the following systems and services are working:

- DNS infrastructure
- Kerberos
- Windows Time service (W32time)
- Remote procedure call (RPC)
- Network connectivity

Use Repadmin (Windows Support Tools) to monitor replication status daily by running a command that assesses the replication status of all domain controllers in your forest. The procedure generates a .csv file that you can open in Excel and filter for replication failures.

Use the following procedure to retrieve the replication status of all domain controllers in the forest.

**Requirements**

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the forest root domain or the Enterprise Admins group in the forest.
- Tools:

  Repadmin.exe (Windows Support Tools)

  Excel (Microsoft Office)

**To retrieve replication status**

1. Open a command prompt, type the following command, and then press ENTER:

   **repadmin /showrepl * /csv >showrepl.csv**

2. In Excel, on the **File** menu, click **Open.**
3. In **Files of type**, click **Text Files (*.prn;*.txt;*.csv)**.
4. In **Look in**, navigate to showrepl.csv, and then click **Open**.
5. In the Excel spreadsheet, right-click the column heading for **showrepl_COLUMNS** (column A) and then click **Hide**. Repeat for the column labeled **Transport Type**.
6. Select the row just under the column headings, and then, on the **Windows** menu, click **Freeze Pane**.
7. Click the upper-left corner of the spreadsheet to highlight the entire spreadsheet. On the **Data** menu, point to **Filter,** and then click **AutoFilte**r.
8. In the heading of the **Last Success** column, click the down arrow, and then click **Sort Ascending**.

9. In the heading of the **Source DC** column, click the down arrow, and then click **Custom**. In the **Custom AutoFilter** dialog box, complete the custom filter as follows:

    a. Under **Source D**C, click **does not contain**.

    b. In the corresponding text box, type **del** to filter deleted domain controllers from the spreadsheet.

10. In the heading of the **Last Failure** column, click the down arrow, and then click **Custom**. In the **Custom AutoFilter** dialog box, complete the custom filter as follows:

    a. Under **Last Failure**, click **does not equal**.

    b. In the corresponding text box, type **0** to filter for only domain controllers that are experiencing failures.

For every domain controller in the forest, the spreadsheet shows the source replication partner, the time that replication last occurred, and the time that the last replication failure occurred for each naming context (directory partition). By using **Autofilter** in Excel, you can view the replication health for working domain controllers only, failing domain controllers only, or domain controllers that are the least or most current, and you can see the replication partners that are replicating successfully.

**Attempting to Resolve Problems**

Replication problems are reported in event messages and in various error messages that occur when an application or service attempts an operation. Ideally, these messages are collected by your monitoring application or when you retrieve replication status.

Most replication problems are identified in the event messages that are logged in the Directory Service event log. Replication problems might also be identified in the form of error messages in the output of the **repadmin /showrepl** command.

**repadmin /showrepl Error Messages That Indicate Replication Problems**

To identify Active Directory replication problems, use the **repadmin /showrepl** command as described in the previous section. The following table shows error messages that are generated by this command, along with the root causes of the errors and links to topics that provide solutions for the errors.

**repadmin /showrepl Error Messages**

| Repadmin error | Root cause | Solution |
|---|---|---|
| The time since last replication with this server has exceeded the tombstone lifetime. | A domain controller has failed inbound replication with the named source domain controller long enough for a deletion to have been tombstoned, replicated, and garbage-collected from Active Directory. | Event ID 2042: It has been too long since this machine replicated |
| No inbound neighbors. | If no items appear in the "Inbound Neighbors" section of the output that is generated by **repadmin /showrepl**, the domain controller was not able to establish replication links with another domain controller. | Fixing Replication Connectivity Problems (Event ID 1925) |
| Access is denied. | A replication link exists between two domain controllers, but replication cannot be performed properly due to an authentication failure. | Fixing Replication Security Problems |
| Last attempt at <date - time> failed with the "Target account name is incorrect." | This problem can be related to connectivity, DNS, or authentication issues.<br><br>If this is a DNS error, the local domain controller could not resolve the globally unique identifier (GUID)–based DNS name of its replication partner. | Fixing Replication DNS Lookup Problems (Event IDs 1925, 2087, 2088)<br><br>Fixing Replication Security Problems<br><br>Fixing Replication Connectivity Problems (Event ID 1925) |
| LDAP Error 49. | The domain controller computer account might not be synchronized with the Key Distribution Center (KDC). | Fixing Replication Security Problems |

| | | |
|---|---|---|
| Cannot open LDAP connection to local host | The administration tool could not contact Active Directory. | Fixing Replication DNS Lookup Problems (Event IDs 1925, 2087, 2088) |
| Active Directory replication has been preempted. | The progress of inbound replication was interrupted by a higher priority replication request, such as a request generated manually with the **repadmin /sync** command. | Wait for replication to complete. This informational message indicates normal operation. |
| Replication posted, waiting. | The domain controller posted a replication request and is waiting for an answer. Replication is in progress from this source. | Wait for replication to complete. This informational message indicates normal operation. |

**Event Messages That Indicate Active Directory Replication Problems**

The following table lists common events that might indicate problems with Active Directory replication, along with root causes of the problems and links to topics that provide solutions for the problems.

**Events That Indicate Active Directory Replication Problems**

| Event ID and source | Root cause | Solution |
|---|---|---|
| 1311 — NTDS KCC | The replication configuration information in Active Directory does not accurately reflect the physical topology of the network. | Fixing Replication Topology Problems (Event ID 1311) |
| 1388 — NTDS Replication | Strict replication consistency is not in effect, and a lingering object has been replicated to the domain controller. | Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042) |
| 1925 — NTDS KCC | The attempt to establish a replication link for a writable directory partition failed. This event can have different causes, depending on the error. | Fixing Replication Connectivity Problems (Event ID 1925)<br><br>Fixing Replication DNS Lookup Problems (Event IDs 1925, 2087, 2088) |
| 1988 — NTDS Replication | The local domain controller has attempted to replicate an object from a source domain controller that is not present on the local domain controller because it may have been deleted and already garbage-collected. Replication will not proceed for this directory partition with this partner until the situation is resolved. | Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042) |
| 2042 — NTDS Replication | Replication has not occurred with this partner for a tombstone lifetime, and replication cannot proceed. | Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042) |
| 2087 — NTDS Replication | Active Directory could not resolve the DNS host name of the source domain controller to an Internet Protocol (IP) address, and replication failed. | Fixing Replication DNS Lookup Problems (Event IDs 1925, 2087, 2088) |
| 2088 — NTDS Replication | Active Directory could not resolve the DNS host name of the source domain controller to an IP address, but replication succeeded. | Fixing Replication DNS Lookup Problems (Event IDs 1925, 2087, 2088) |
| 2095 — NTDS Replication | Update sequence number (USN) rollback has occurred and replication has been stopped. This error indicates an improper Active Directory restore, possibly of a virtual machine file (.vhd). | For an explanation of this problem and recommendations for solutions, see Running Domain Controllers in Virtual Server 2005 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=38330). |
| 5805 — Net Logon | A machine account failed to authenticate, which is usually caused by either multiple instances of the same computer name or the computer name not replicating to every domain controller. | Fixing Replication Security Problems |

For more information about replication concepts, see “Active Directory Replication Technologies” in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41950).

**In this section**

Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042)

Fixing Replication Security Problems

Fixing Replication DNS Lookup Problems (Event IDs 1925, 2087, 2088)

Fixing Replication Connectivity Problems (Event ID 1925)

Fixing Replication Topology Problems (Event ID 1311)

# Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042)

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If a domain controller does not replicate for a period of time that is longer than the tombstone lifetime and the domain controller is then reconnected to the replication topology, objects that were deleted from Active Directory while the domain controller was offline can remain on the domain controller as lingering objects.

**Tombstone Lifetime and Replication of Deletions**
When an object is deleted, Active Directory replicates the deletion as a tombstone object, which consists of a small subset of the attributes of the deleted object. By inbound-replicating this object, other domain controllers in the domain and forest become aware of the deletion. The tombstone is retained in Active Directory for a specified period called the tombstone lifetime. At the end of the tombstone lifetime, the tombstone is deleted from the directory permanently.

The default value of the tombstone lifetime depends on the version of the operating system that is running on the first domain controller that is installed in a forest, as follows:

- Windows 2000 Server or Windows Server 2003: The default value is 60 days.
- Windows Server 2003 with Service Pack 1 (SP1): The default value is 180 days.

| **Note** |
| --- |
| The tombstone lifetime value that is in effect when a domain controller is upgraded to Windows Server 2003 SP1 is not changed by upgrading. The existing value is maintained until you change it manually. |

After the tombstone is removed permanently, the object deletion can no longer be replicated. Therefore, the tombstone lifetime defines how long domain controllers in the forest retain knowledge of a deleted object and thus the time during which a unique deletion must be received by all direct and transitive replication partners of the originating domain controller.

**How Lingering Objects Occur**
When conditions beyond your control cause a domain controller to be disconnected for a period that is longer than the tombstone lifetime, one or more objects that are deleted from Active Directory on all other domain controllers might remain on the disconnected domain controller. Such objects are called lingering objects. Because the domain controller is offline during the entire time that the tombstone is alive, the domain controller never receives replication of the tombstone.

When it is reconnected to the replication topology, this domain controller acts as a source replication partner that has an object that its destination partner does not have.

Replication problems occur when the object on the source domain controller is updated. In this case, when the destination attempts to inbound-replicate the update, one of the following happens:

- If the destination domain controller has strict replication consistency enabled, it recognizes that it cannot update the object and locally halts inbound replication of the directory partition from that source domain controller.
- If the destination domain controller has strict replication consistency disabled, it requests the full replica of the updated object. In this case, the object is reintroduced into the directory.

Lingering objects can reside in writable or read-only partitions that are potentially replicated between domain controllers in the same or different domains in the same forest.

**Causes of Long Disconnections**
Unexpectedly long disconnections can be caused by the following conditions:

- A domain controller is left in a storage room and forgotten, or shipment of a prestaged domain controller to its remote location takes longer than a tombstone lifetime.

- Replication fails and monitoring is not in place. Failures can occur as follows:

    - A domain controller is started and connected to the corporate intranet but experiences inbound replication failure as a result of an underlying network connectivity failure, name resolution failure, or authentication failure that prevents replication from occurring.

    - A bridgehead server is overloaded, and replication becomes backlogged. Excessively high replication load on a global catalog server, in combination with a short intersite replication interval, can result in updates not being replicated.

| **Note** |
| --- |
| Global catalog servers replicate read-only replicas of all domain directory partitions in the forest. The replication of read-only replicas has a lower priority than the replication of writable replicas. In addition, global catalog servers are often bridgehead servers, which adds to the replication load. If the replication load on global catalog servers acting as bridgehead servers is too high as a result of an extremely short replication interval, excessive numbers of concurrent outbound replication partners, or a combination of both, the replication queue can become backlogged. If the condition persists, read-only replicas can remain in the queue indefinitely. These conditions can result in lingering objects on a global catalog server. |

- Wide area network (WAN) connections are unavailable for long periods. For example, a domain controller onboard a cruise ship might be unable to replicate because the ship is at sea for longer than the tombstone lifetime.

- The reported event is a false positive because an administrator shortened the tombstone lifetime to force tombstone deletion (garbage collection).

- The reported event is a false positive because the system clock on the source or destination domain controller is improperly rolled forward or back in time. Clock skews are most common following a system reboot and can have the following causes:

    - System clock battery or motherboard problems.

    - The time source for a computer is improperly configured, including a time source server configured with Windows Time service (W32time), third-party time servers, and network routers.

    - The system clock is advanced or rolled back by an administrator attempting to extend the useful life of a system state backup or accelerate the garbage collection of deleted objects. Make sure that the system clock reflects the actual time and that event logs do not contain events from the future or invalid past.

**Indications That a Domain Controller Has Lingering Objects**

An outdated domain controller can store lingering objects with no noticeable effect as long as an administrator, application, or service does not update the lingering object or attempt to create an object with the same name in the domain or with the same user principal name (UPN) in the forest. However, the existence of lingering objects can cause problems, especially if the object is a security principal.

**Symptoms Associated with Lingering Objects**

The following symptoms indicate that a domain controller has lingering objects:

- A deleted user or group account remains in the global address list (GAL) on Exchange servers. Therefore, although the account name appears in the GAL, attempts to send e-mail messages result in errors.

- Multiple copies of an object appear in the object picker or GAL for an object that should be unique in the forest. Duplicate objects sometimes appear with altered names, causing confusion on directory searches. For example, if the relative distinguished name of two objects cannot be resolved, conflict resolution appends "*CNF:*GUID*" to the name, where * represents a reserved character, CNF is a constant that indicates a conflict resolution, and *GUID* represents the **objectGUID** attribute value.

- E-mail messages are not delivered to a user whose Active Directory account appears to be current. After an outdated

domain controller or global catalog server becomes reconnected, both instances of the user object appear in the global catalog. Because both objects have the same e-mail address, e-mail messages cannot be delivered.

- A universal group that no longer exists continues to appear in a user's access token. Although the group no longer exists, if a user account still has the group in its security token, the user might have access to a resource that you intended to be unavailable to that user.

- A new object or Exchange mailbox cannot be created, but you do not see the object in Active Directory. An error message reports that the object already exists.

- Searches that use attributes of an existing object incorrectly find multiple copies of an object of the same name. One object has been deleted from the domain, but it remains in an isolated global catalog server.

If an attempt is made to update a lingering object that resides in a writable directory partition, events are logged on the destination domain controller. However, if the only version of a lingering object exists in a read-only directory partition on a global catalog server, the object cannot be updated and this type of event will never be triggered.

**Registry Setting That Determines Whether Lingering Objects Are Replicated**

If a writable lingering object exists in your environment and an attempt is made to update the object, the value in the **strict replication consistency** registry entry (type **REG_DWORD**) in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** determines whether replication proceeds or is stopped, as follows:

- **1** (enabled): Inbound replication of the specified directory partition from the source is stopped on the destination.

- **0** (disabled): The destination requests the full object from the source domain controller, and the lingering object is revived in the directory as a new object.

**Default Settings for Strict Replication Consistency**

The default value for the **strict replication consistency** registry entry is determined by the conditions under which the domain controller was installed into the forest.

| **Note** |
| --- |
| Raising the domain or forest functional level does not change the replication consistency setting on any domain controller. |

**Strict replication consistency enabled**

The value of **strict replication consistency** on domain controllers that are installed into a forest defaults to enabled (**1**) under the following conditions:

- The forest root domain of a new forest is created by upgrading the Windows NT 4.0 primary domain controller (PDC) to Windows Server 2003 by using the Windows Server 2003 version of Winnt32.exe.

- The forest root domain of a new forest is created by installing Active Directory on a server running Windows Server 2003.

**Strict replication consistency disabled**

The value of **strict replication consistency** on domain controllers defaults to disabled (**0**) under the following conditions:

- A domain controller running Windows 2000 Server is upgraded to Windows Server 2003.

- A server running Windows 2000 Server is promoted into a Windows Server 2003 forest.

If you have a domain controller that is running Windows Server 2003 with SP1, you do not need to edit the registry to set strict replication consistency. Instead, you can use Repadmin to set the value for one or all domain controllers in the forest. To set strict replication consistency for specific domain controllers or for all domain controllers, see Event ID 1388 or 1988: A lingering object is detected.

For more information about strict replication consistency, see "How the Active Directory Replication Model Works" in the Windows Server 2003 Technical Reference on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=27636.

**Tool for Removing Lingering Objects**

On domain controllers running Windows Server 2003 or Windows Server 2003 with SP1, use Repadmin.exe (in Windows Support Tools) to remove the lingering object or objects. Windows Support Tools are available on the operating system CD in the Support\Tools folder. The version of Repadmin that ships with Windows Server 2003 provides the option **/removelingeringobjects**, which safely removes instances of lingering objects from both writable directory partitions and read-only directory partitions.

The **repadmin /removelingeringobjects** command does the following:

1. Compares the directory database objects on a reference domain controller with the objects on the target domain controller, which contains (or is suspected to contain) lingering objects. There must be connectivity between the reference domain controller and the target domain controller.

2. Either removes the lingering objects or logs the potential deletions to the Directory Service event log, as follows:

    - If you use the **/advisory_mode** parameter, events are logged in the Directory Service event log for the objects that are found.

    - If you do not use the **/advisory_mode** parameter, the found objects are deleted without replicating the deletions; that is, the deletions occur only on the target domain controller.

Choose the problem that best describes your situation from the following list, and then step through the suggested fix:

- Event ID 1388 or 1988: A lingering object is detected

- A deleted account remains in the Address Book, e-mail is not received, or a duplicate account exists

- Event ID 2042: It has been too long since this machine replicated

**See Also**

**Concepts**

Configuring a Computer for Troubleshooting Active Directory

# Event ID 1388 or 1988: A lingering object is detected

Updated: November 19, 2012

Applies To: Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

If a destination domain controller logs Event ID 1388 or Event ID 1988, a lingering object has been detected and one of two conditions exists on the destination domain controller:

- Event ID 1388: Inbound replication of the lingering object has occurred on the destination domain controller.
- Event ID 1988: Inbound replication of the directory partition of the lingering object has been blocked on the destination domain controller.

**In this topic**


- Event ID 1388
- Event ID 1988
- Diagnosis
- Resolution
  - Use Repadmin to identify the GUID of an authoritative domain controller
  - Use Repadmin to remove lingering objects
- Enable strict replication consistency
  - Use Repadmin to enable strict replication consistency
  - Use Regedit to enable strict replication consistency
- Ensure that strict replication consistency is enabled for newly promoted domain controllers


**Event ID 1388**

This event indicates that a destination domain controller that does not have strict replication consistency enabled received a request to update an object that does not reside in the local copy of the Active Directory database. In response, the destination domain controller requested the full object from the source replication partner. In this way, a lingering object was replicated to the destination domain controller. Therefore, the lingering object was reintroduced into the directory.

| **Important** |
| --- |
| When Event ID 1388 occurs, if either the source domain controller (the replication partner that is outbound-replicating the lingering object) or the destination domain controller (the inbound replication partner that reports Event ID 1388) is running Windows 2000 Server, you cannot use the Repadmin tool to remove lingering objects. For information about how to remove lingering objects in this case, see article 314282 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=41410). The procedures and information in this article apply to the removal of lingering objects from global catalog servers as well as from domain controllers that are not global catalog servers. |

The event text identifies the source domain controller and the outdated (lingering) object. The following is an example of the event text:

```
Log Name: Directory Service
Source: Microsoft-Windows-ActiveDirectory_DomainService
Date: 5/3/2008  3:34:01 PM
```

```
Event ID: 1388
Task Category: Replication
Level: Error
Keywords: Classic
User: ANONYMOUS LOGON
Computer: DC3.contoso.com
Description:
Another domain controller (DC) has attempted to replicate into this DC an
object which is not present in the local Active Directory Domain Services database.
The
object may have been deleted and already garbage collected (a tombstone
lifetime or more has past since the object was deleted) on this DC. The
attribute set included in the update request is not sufficient to create
the object. The object will be re-requested with a full attribute set
and re-created on this DC.

Source DC (Transport-specific network address):
4a8717eb-8e58-456c-995a-c92e4add7e8e._msdcs.contoso.com
Object:
CN=InternalApps,CN=Users,DC=contoso,DC=com
Object GUID:
a21aa6d9-7e8a-4a8f-bebf-c3e38d0b733a
Directory partition:
DC=contoso,DC=com
Destination highest property USN:
20510
User Action:
Verify the continued desire for the existence of this object. To
discontinue re-creation of future similar objects, the following
registry key should be created.

Registry Key:
HKLM\System\CurrentControlSet\Services\NTDS\Parameters\Strict Replication Consistency
```

**Event ID 1988**

This event indicates that a destination domain controller that has strict replication consistency enabled has received a request to update an object that does not exist in its local copy of the Active Directory database. In response, the destination domain controller blocked replication of the directory partition containing that object from that source domain controller. The event text identifies the source domain controller and the outdated (lingering) object. The following is an example of the event text:

```
Log Name: Directory Service
Source: Microsoft-Windows-ActiveDirectory_DomainService
Date: 2/7/2008  8:20:11 AM
Event ID: 1988
Task Category: Replication
Level: Error
Keywords: Classic
User: ANONYMOUS LOGON
Computer: DC5.contoso.com
Description:
Active Directory Domain Services Replication encountered the existence of objects
in the following partition that have been deleted from the local
domain controllers (DCs) Active Directory Domain Services database. Not all direct
or transitive replication partners replicated in the deletion
before the tombstone lifetime number of days passed. Objects that
have been deleted and garbage collected from an Active Directory Domain Services
partition but still exist in the writable partitions of other DCs
in the same domain, or read-only partitions of global catalog servers
in other domains in the forest are known as "lingering objects".

This event is being logged because the source DC contains a lingering
object which does not exist on the local DCs Active Directory Domain Services database.
This replication attempt has been blocked.

The best solution to this problem is to identify and remove all
lingering objects in the forest.

Source DC (Transport-specific network address):
4a8717eb-8e58-456c-995a-c92e4add7e8e._msdcs.contoso.com
Object:
CN=InternalApps,CN=Users,DC=contoso,DC=com
Object GUID:
```

a21aa6d9-7e8a-4a8f-bebf-c3e38d0b733a

**Diagnosis**
An object that has been permanently deleted from AD DS (that is, its tombstone has been garbage-collected on all connected domain controllers) remains on a disconnected domain controller. The domain controller failed to receive direct or transitive replication of the object deletion because it was disconnected (it is offline or experiencing an inbound replication failure) from the replication topology for a period that exceeded a tombstone lifetime. The domain controller is now reconnected to the topology and that object has been updated on the domain controller, causing a replication notification to the replication partner that an update is ready for replication. The replication partner responded according to its replication consistency setting. This notification applies to attempted replication of a writable object. A copy of the writable lingering object might also exist on a global catalog server.

**Resolution**
If replication of a lingering object is detected, you can remove the object from AD DS, along with any read-only replicas of the object, by identifying the domain controllers that might store this object (including global catalog servers) and running a **repadmin** command to remove lingering objects on these servers (**repadmin /removelingeringobjects**). This command is available on domain controllers that are running Windows Server 2008. It is also available on domain controllers that are not running Windows Server 2008 but are running the version of Repadmin.exe that is included with Windows Support Tools in Windows Server 2003.

To remove lingering objects, do the following:

1. Use the event text to identify the following:

    a. The directory partition of the object

    b. The source domain controller that attempted replication of the lingering object

2. Use Repadmin to identify the GUID of an authoritative domain controller

3. Use Repadmin to remove lingering objects

4. Enable strict replication consistency, if necessary

5. Ensure that strict replication consistency is enabled for newly promoted domain controllers, if necessary

**Use Repadmin to identify the GUID of an authoritative domain controller**
To perform the procedure that removes lingering objects, you must identify the globally unique identifier (GUID) of an up-to-date domain controller that has a writable replica of the directory partition that contains the lingering object that has been reported. The directory partition is identified in the event message.

The object GUID of a domain controller is stored in the **objectGUID** attribute of the NTDS Settings object.

### Requirements

- Membership in **Domain Admins** in the domain of the domain controller whose GUID you want to identify is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

- Tool: Repadmin.exe

**To identify the GUID of a domain controller**

1. At a command prompt, type the following command, and then press ENTER:

    ```
    repadmin /showrepl <ServerName>
    ```

| Parameter | Description |
|---|---|

| | |
|---|---|
| /showrepl | Displays the replication status, including when the domain controller that is specified by <ServerName> last attempted inbound replication of Active Directory partitions. Also displays the GUID of the specified domain controller. |
| <ServerName> | The name of the domain controller whose GUID you want to display. |

2. In the first section of the output, locate the **objectGuid** entry. Select and copy the GUID value into a text file so that you can use it elsewhere.

**Use Repadmin to remove lingering objects**

If the destination domain controller and source domain controller are running either Windows Server 2003 or Windows Server 2008, you can use this procedure to remove lingering objects with Repadmin. If either domain controller is running Windows 2000 Server, follow the instructions in article 314282 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=41410).

### Requirements

- Membership in **Domain Admins** in the domain of the domain controller that has lingering objects, or **Enterprise Admins** if the directory partition that has lingering objects is the configuration or schema directory partition, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

- Operating system: Windows Server 2003 or Windows Server 2008 for <ServerName> and <ServerGUID>

- Tool: Repadmin.exe

**To use Repadmin to remove lingering objects**

1. Open a Command Prompt as an administrator: On the **Start** menu, right-click **Command Prompt**, and then click **Run as administrator**. If the **User Account Control** dialog box appears, provide Domain Admins or Enterprise Admins credentials, if required, and then click **Continue**.

2. At the command prompt, type the following command, and then press ENTER:

```
repadmin /removelingeringobjects <ServerName> <ServerGUID> <DirectoryPartition> /advisory_mode
```

| Parameter | Description |
|---|---|
| /removelingeringobjects | Removes lingering objects from the domain controller that is specified by <ServerName> for the directory partition that is specified by <DirectoryPartition>. |
| <ServerName> | The name of the domain controller that has lingering objects, as identified in the event message (Event ID 1388 or Event ID 1988). You can use the Domain Name System (DNS) name or the distinguished name, for example, the distinguished name CN=DC5,OU=Domain Controllers,DC=contoso,DC=com or the DNS name DC5.contoso.com. |
| <ServerGUID> | The GUID of a domain controller that has an up-to-date, writable replica of the directory partition that contains the lingering object |
| <DirectoryPartition> | The distinguished name of the directory partition that is identified in the event message, for example:<br>• For the Sales domain directory partition in the contoso.com forest: DC=sales,DC=contoso,DC=com<br>• For the configuration directory partition in the contoso.com forest: CN=configuration,DC=contoso,DC=com<br>• For the schema directory partition in the contoso.com forest: CN=schema,CN=configuration,DC=contoso,DC=com |

| | |
|---|---|
| /advisory_mode | Logs the lingering objects that will be removed so that you can review them, but does not remove them. |

3. Repeat step 2 without `/advisory_mode` to delete the identified lingering objects from the directory partition.

4. Repeat steps 2 and 3 for every domain controller that might have lingering objects.

**Note**

The <ServerName> parameter uses the `DC_LIST` syntax for `repadmin`, which allows the use of `*` for all domain controllers in the forest and `gc:` for all global catalog servers in the forest. To see the `DC_LIST` syntax, type `repadmin /listhelp`. For information about the syntax of the `/regkey` and `/removelingeringobjects` parameters, type `repadmin /experthelp`.

**Enable strict replication consistency**

To ensure that lingering objects cannot be replicated if they occur, enable strict replication consistency on all domain controllers. The setting for replication consistency is stored in the registry on each domain controller. However, on domain controllers that are running Windows Server 2003 with Service Pack 1 (SP1), Windows Server 2003 with Service Pack 2 (SP2), Windows Server 2003 R2, or Windows Server 2008, you can use Repadmin to enable strict replication consistency on one or all domain controllers.

On domain controllers running Windows Server 2003 without SP1 or running any version of Windows 2000 Server, you must edit the registry to enable the setting.

**Use Repadmin to enable strict replication consistency**

Use this procedure to remove lingering objects on a domain controller that is running Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2003 R2, or Windows Server 2008.

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure on a single domain controller. Membership in **Enterprise Admins**, or equivalent, is the minimum required to complete this procedure on all domain controllers in the forest. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To use Repadmin to enable strict replication consistency**

1. Open a Command Prompt as an administrator: On the **Start** menu, right-click **Command Prompt**, and then click **Run as administrator**. If the **User Account Control** dialog box appears, provide Domain Admins or Enterprise Admins credentials, if required, and then click **Continue**.

2. At the command prompt, type the following command, and then press ENTER:

```
repadmin /regkey <DC_LIST> +strict
```

| Parameter | Description |
|---|---|
| /regkey | Enables (+) and disables (-) the value for the **Strict Replication Consistency** registry entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** |
| <DC_LIST> | The name of a single domain controller, or* to apply the change to all domain controllers in the forest. For the domain controller name, you can use the DNS name, the distinguished name of the domain controller computer object, or the distinguished name of the domain controller server object, for example, the distinguished name CN=DC5,OU=Domain Controllers,DC=contoso,DC=com or the DNS name DC5.contoso.com. |
| +strict | Enables the **Strict Replication Consistency** registry entry. |

3. If you do not use `*` to apply the change to all domain controllers, repeat step 2 for every domain controller on which you want to enable strict replication consistency.

**Note**

For more naming options and information about the syntax of the <DC_LIST> parameter, at the command prompt type `repadmin /listhelp`. For information about the syntax of the `/regkey` and `/removelingeringobjects` parameters, type `repadmin /experthelp`.

#### Use Regedit to enable strict replication consistency

As an alternative to using Repadmin, you can enable strict replication consistency by editing the registry directly. The registry method is required for a domain controller that is running a version of Windows Server that is earlier than Windows Server 2003 with SP1. The setting for replication consistency is stored in the **Strict Replication Consistency** entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

The values for the **Strict Replication Consistency** registry entry are as follows:

- Value: **1** (**0** to disable)
- Default: **1** (enabled) in a new Windows Server 2003 or Windows Server 2008 forest; otherwise **0**.
- Data type: **REG_DWORD**

**Requirements:**

- Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).
- Tool: Regedit.exe

| ⚠Caution |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

##### To use Regedit to enable strict replication consistency

1. Open Regedit as an administrator: Click **Start** and then, in **Start Search**, type **regedit**. At the top of the **Start** menu, right-click **regedit.exe**, and then click **Run as administrator**. In the **User Account Control** dialog box, provide Domain Admins credentials, and then click **OK**.
2. Navigate to the **Strict Replication Consistency** entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.
3. Set the value in the **Strict Replication Consistency** entry to **1.**

#### Ensure that strict replication consistency is enabled for newly promoted domain controllers

If you are upgrading a forest that was originally created using a computer running Windows 2000 Server, you should ensure that the forest is configured to enable strict replication consistency on newly promoted domain controllers to help avoid lingering objects. After you update the forest as described in Upgrading Active Directory Domains to Windows Server 2008 AD DS Domains (http://go.microsoft.com/fwlink/?LinkId=89032), all new domain controllers that you subsequently add to the forest are created with strict replication consistency disabled. However, you can implement a forest configuration change that causes new domain controllers to have strict replication consistency enabled. To ensure that new domain controllers that you add to the forest have strict replication consistency enabled, you can use the Ldifde.exe tool to create an object in the configuration directory partition of the forest. This object is responsible for enabling strict replication consistency on any Windows Server 2003 or Windows Server 2008 domain controller that is promoted into the forest.

The object that you create is an operational GUID with the following name:

CN=94fdebc6-8eeb-4640-80de-ec52b9ca17fa,CN=Operations,CN=ForestUpdates,CN=Configuration,DC=<ForestRootDomain>

You can use the following procedure on any domain controller in the forest to add this object to the configuration directory partition.

Membership in **Enterprise Admins**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups

(http://go.microsoft.com/fwlink/?LinkId=83477).

**To create the object that ensures strict replication consistency on new domain controllers**

1. In a text editor, such as Notepad, create the following text file:

   ```
   dn:
   CN=94fdebc6-8eeb-4640-80de-
   ec52b9ca17fa,CN=Operations,CN=ForestUpdates,CN=Configuration,DC=<ForestRootDomain>
   changetype: add
   objectClass: container
   showInAdvancedViewOnly: TRUE
   name: 94fdebc6-8eeb-4640-80de-ec52b9ca17fa
   objectCategory: CN=Container,CN=Schema,CN=Configuration,DC=<ForestRootDomain>
   ```

   Where <ForestRootDomain> contains all domain components (DC=) of the forest root domain, for example, for the contoso.com forest, DC=contoso,DC=com; for the fineartschool.net forest, DC=fineartschool,DC=net.

2. Open a Command Prompt as an administrator: On the **Start** menu, right-click **Command Prompt**, and then click **Run as administrator**. If the **User Account Control** dialog box appears, provide Enterprise Admins credentials, if required, and then click **Continue**.

3. At the command prompt, type the following command, and then press ENTER:

   ```
   ldifde -i -f <Path>\<FileName>
   ```

| Parameter | Description |
|---|---|
| -i | Specifies the import mode. If the import mode is not specified, the default mode is export. |
| -f | Identifies the import or export file name. |
| <Path>\<FileName> | The path and name of the import file that you created in step 1, for example, C:\ldifde.txt. |

For information about using Ldifde, see LDIFDE (http://go.microsoft.com/fwlink/?LinkId=42656).

# A deleted account remains in the Address Book, e-mail is not received, or a duplicate account exists

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Deleted accounts remaining in the global address list (GAL), e-mail messages for existing accounts not being received, or duplicate objects existing in Active Directory are all symptoms that can indicate a lingering object problem. If you have no error or event that identifies the lingering object and its directory partition (for example, event ID 1388 or event ID 1988), you must search the global catalog for an object that you believe might be causing the problem. When you identify the lingering object and the directory partition of the object, you can perform the procedures to remove it.

**Solution**

Based on these symptoms of a lingering object, you usually have a good idea of the name of the object and you can use the following steps to solve the problem:

- Use this name to identify the object in the global catalog.
- Identify the directory partition of the object.
- Remove all lingering objects from that directory partition on all global catalog servers in the forest.

**Identify the Duplicate (Lingering) Object**

Use the following procedure to identify the duplicate (lingering) object by searching the global catalog for its distinguished name. Use an attribute that uniquely identifies the object for the account that is not receiving e-mail, cannot be created because it already exists, or appears in the Address Book or in access control lists (ACLs) when it has already been deleted.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Users group.
- Tool: Ldp.exe (Windows Support Tools)

**To establish the distinguished name of an object**

1. Click **Start**, click **Run**, type **Ldp**, and then click **OK**.
2. On the **Connection** menu, click **Connect**.
3. In **Port**, type **3268**, and then click **OK**.
4. On the **Connection** menu, click **Bind**.
5. In the **Bind** dialog box, provide credentials for a user account in the forest, and then click **OK**.
6. On the **View** menu, click **Tree**.
7. In the **Tree View** dialog box, in **BaseDN**, type the distinguished name of the forest root domain, and then click **OK**.
8. In the console tree, right-click the forest root domain, and then click **Search**.
9. In the **Search** dialog box, in **Filter**, replace the default filter (objectClass=*) to create a filter of the following form:

   (*attribute*=*value*)

   where *attribute* is the Lightweight Directory Access Protocol (LDAP) name of an attribute and *value* is the value that you know or suspect to be associated with the object that you are searching for. For example, use (userPrincipalName=JanD@contoso.com), (sAMAccountName=JanD), or (sn=Dryml) to locate the duplicate user object Jan Dryml. You can use the asterisk (*) in the *value* field if you want to search all objects.
10. In the **Scope** box, click **Subtree**, and then click **Run**.

11. Click **Close**, and then view the results. You must identify which of the displayed objects should be removed from Active Directory. An indication that you have found a lingering object that exists only on a global catalog server is that the object does not exist in a writable replica of the directory partition.

12. If necessary, repeat steps 8 through 10 to rephrase the query, and then run it again.

**Identify the Directory Partition of the Object**

After you identify the distinguished name of the object that is causing problems, if it is a domain object, identify the domain in which it is located by looking at the DC= part of the distinguished name. For example, if the object you find has the distinguished name CN=Jan Dryml,CN=Users,DC=Region1,DC=Contoso,DC=com, the directory partition name for the user account is DC=Region1,DC=Contoso,DC=com.

**Remove the Lingering Object**

Use the directory partition name in the procedure "To use Repadmin to remove lingering objects" to remove the lingering object from all domain controllers and global catalog servers in the forest as described in "Event ID 1388 or 1988: A lingering object is detected."

# Event ID 2042: It has been too long since this machine replicated

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If a domain controller has not replicated with its partner for longer than a tombstone lifetime, it is possible that a lingering object problem exists on one or both domain controllers. When this condition occurs, inbound replication with the source partner is stopped on the destination domain controller and event ID 2042 is logged in the Directory Services event log. The event identifies the source domain controller and the appropriate steps to take to either remove the outdated domain controller or remove lingering objects and restore replication from the source domain controller.

An example of the event text is as follows:

```
Event Type:Error
Event Source:NTDS Replication
Event Category:Replication
Event ID:2042
Date:3/22/2005
Time:7:28:49 AM
User:NT AUTHORITY\ANONYMOUS LOGON
Computer:DC3
Description:
It has been too long since this machine last replicated with the
named source machine. The time between replications with this source
has exceeded the tombstone lifetime. Replication has been stopped
with this source.
The reason that replication is not allowed to continue is that
the two machine's views of deleted objects may now be different.
The source machine may still have copies of objects that have
been deleted (and garbage collected) on this machine. If they
were allowed to replicate, the source machine might return
objects which have already been deleted.
Time of last successful replication:
2005-01-21 07:16:03
Invocation ID of source:
0397f6c8-f6b8-0397-0100-000000000000
Name of source:
4a8717eb-8e58-456c-995a-c92e4add7e8e._msdcs.contoso.com
Tombstone lifetime (days):
60

The replication operation has failed.

User Action:

Determine which of the two machines was disconnected from the
forest and is now out of date. You have three options:

1. Demote or reinstall the machine(s) that were disconnected.
2. Use the "repadmin /removelingeringobjects" tool to remove
inconsistent deleted objects and then resume replication.
3. Resume replication. Inconsistent deleted objects may be introduced.
You can continue replication by using the following registry key.
Once the systems replicate once, it is recommended that you remove
the key to reinstate the protection.
Registry Key:
HKLM\System\CurrentControlSet\Services\NTDS\Parameters\Allow Replication With Divergent
and Corrupt Partner
```

The **repadmin /showrepl** command also reports error 8614:

```
Source: Default-First-Site-Name\DC1
******* 1502 CONSECUTIVE FAILURES since 2005-01-21 07:16:00
Last error: 8614 (0x21a6):
            The Active Directory cannot replicate with this server
because the time since the last replication with this server has
exceeded the tombstone lifetime.
```

**Solution**
Treat this occurrence as a lingering object condition, and do the following:

- Run the **repadmin /showrepl** command on the domain controller that received the error to determine which domain controller has been disconnected for longer than a tombstone lifetime.

- Check the system time on both the source and destination domain controllers. If there is a time difference, ensure that you correct the time issue before proceeding. You may have to wait a couple of hours to see ensure that a time skew does not occur. If a time skew does occur, you must diagnose why and solve the issue. For example, there could be hardware issues, like system battery failure, or a firmware update might be required, or the forest root PDC might not be configured to use an external time source (as described in Configure the Time Source for the Forest).

- Remove lingering objects. Follow the instructions for removing lingering objects from the source and destination domain controllers as described in Event ID 1388 or 1988: A lingering object is detected.

- Restart replication on the destination domain controller. After you remove lingering objects, you must restart replication on the domain controller that logged the event by editing the registry setting that allows replication with a potentially out-of-date domain controller. You can also perform this procedure if you do not want to wait to remove lingering objects and you want to start replication immediately.

- Reset the registry to protect the domain controller against outdated replication. After replication has resumed on the domain controller that logged the event, reset the registry so that this domain controller continues to log events if replication is attempted with a domain controller where the last successful replication occurred longer than a tombstone lifetime ago.

**Restart Replication Following Event ID 2042**
To restart inbound replication on the destination domain controller following event ID 2042, you must edit the **Allow Replication With Divergent and Corrupt Partner** registry entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

Use the following procedure to change the registry entry value. This procedure does not require a restart of the domain controller to take effect.

| ⚠**Caution** |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the domain of the domain controller.

- Tool: Regedit.exe

**To restart replication following event ID 2042**

1. Click **Start**, click **Run**, type **regedit**, and then click **OK**.

2. Navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**

3. In the details pane, create or edit the registry entry as follows:

   If the registry entry exists in the details pane, modify the entry as follows:

   a. In the details pane, right-click **Allow Replication With Divergent and Corrupt Partner**, and then click **Modify**.

b. In the **Value data** box, type **1**, and then click **OK**.

If the registry entry does not exist, create the entry as follows:

a. Right-click **Parameters**, click **New**, and then click **DWORD Value**.

b. Type the name **Allow Replication With Divergent and Corrupt Partner**, and then press ENTER.

c. Double-click the entry. In the **Value data** box, type **1**, and then click **OK**.

**Reset the Registry to Protect Against Outdated Replication**

When you are satisfied that lingering objects have been removed and replication has occurred successfully from the source domain controller, edit the registry to return the value in **Allow Replication With Divergent and Corrupt Partner** to **0**.

| **Note** |
| --- |
| If you did not remove the lingering objects, attempting replication might result in replication of a lingering object. If strict replication consistency is enabled on the destination domain controller, replication with the source domain controller will be blocked again. |

# Fixing Replication Security Problems

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When security problems cause replication to fail, various event log messages and Repadmin messages contain error codes that identify the problems.

The version of Dcdiag.exe that is included with Windows Support Tools in Windows Server 2003 Service Pack 1 (SP1) provides new functionality that reports on the overall health of replication with respect to Active Directory. Dcdiag is modified to detect common causes of "Access denied" events, "Account unknown" events, and similar events.

The error codes that Dcdiag detects are described in the following table. Error codes that are marked with an asterisk (*) are not always caused by a security problem.

| Error code | Description |
|---|---|
| 5 | Access is denied. |
| 1314* | A required privilege is not held by the client. |
| 1326 | Logon failure: unknown user name or bad password. |
| 1396 | Logon failure: The target account name is incorrect. |
| 1908 | Could not find the domain controller for this domain. |
| 1397* | Mutual authentication failed. The server's password is out of date at the domain controller. |
| 1398* | There is a time and/or date difference between the client and server. |
| 1722* | The remote procedure call (RPC) server is unavailable. |
| 2202* | The specified username is invalid. |
| 8453 | Replication access was denied. |

Use the procedures in An "Access denied" or other security error has caused replication problems to diagnose and fix replication security problems.

# An "Access denied" or other security error has caused replication problems

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Replication problems that have security causes can be tested and diagnosed by using the version of Dcdiag.exe that is included with Windows Support Tools in Windows Server 2003 Service Pack 1 (SP1).

**Cause**
A replication destination domain controller cannot contact its source replication partner to get Active Directory updates as a result of one or more security errors occurring on the connection between the two domain controllers.

**Solution**
Run the replication security error diagnostic test that is available in the version of Dcdiag in Windows Support Tools that is included in Windows Server 2003 SP1.

**Test a Domain Controller for Replication Security Errors**
You can test any or all domain controllers in your forest for security errors.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group to test a domain controller in your domain or a member of the Enterprise Admins group to test a domain controller in another domain.

- Tool: Dcdiag.exe (Windows Support Tools) in Windows Server 2003 SP1

- Operating system:

  - Although you can run the enhanced version of Dcdiag on computers running Windows XP Professional and Windows Server 2003 with no service pack installed, to run the new replication security test (**/test:CheckSecurityError**), you must run Dcdiag on a domain controller running Windows Server 2003 with SP1.

  - You can run the new Dcdiag replication security tests against domain controllers that are running the following operating systems:

    Windows 2000 Server with Service Pack 3 (SP3)

    Windows 2000 Server with Service Pack 4 (SP4)

    Windows Server 2003

    Windows Server 2003 with SP1

**To test a domain controller for replication security errors**

1. At a command prompt, type the following command, and then press ENTER:

   **dcdiag /test:CheckSecurityError /s:** *DomainControllerName*

   where *DomainControllerName* is the Domain Name System (DNS) name, network basic input/output system (NetBIOS) name, or distinguished name of the domain controller on which you want to test.

   If you do not use the **/s:** switch, the test is run against the local domain controller. You can also test all domain controllers in the forest by using **/e:** instead of **/s:**.

2. Copy the report into Notepad or an equivalent text editor

3. Scroll to the Summary table near the bottom of the Dcdiag log file.
4. Note the names of all domain controllers that reported "Warn" or "Fail" status in the Summary table.
5. Find the detailed breakout section for the problem domain controller by searching on the string "DC: *DomainControllerName*".
6. Make the required configuration changes on the domain controllers.

   Rerun **Dcdiag /test:CheckSecurityError** with the **/e:** or **/s:** switch to validate the configuration changes.

**Test the Connection Between Two Domain Controllers for Replication Security Errors**
You can test the connection between two domain controllers in your forest for replication security errors. The domain controller that represents the source of the inbound connection does not have to be an existing source to run this test; that is, a connection object from that domain controller does not have to exist on the destination domain controller. The test is useful in the following scenarios:

- A connection exists between a source and a destination, and you receive a security error.
- A connection should be created automatically by the Knowledge Consistency Checker (KCC) and you want to test why the connection does not exist.
- You are trying to create a connection between two domain controllers and you receive a security error.
- You want to determine whether a connection could be created if you wanted to add one on this destination from the specified source.

**Requirements**

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group to test the connection between domain controllers in your domain or a member of the Enterprise Admins group to test the connection between domain controllers in different domains.
- Tool: Dcdiag.exe (Windows Support Tools) included in Windows Server 2003 SP1
- Operating system:
  - Although you can run the enhanced version of Dcdiag on computers that are running Windows XP Professional and Windows Server 2003 with no service pack installed, to run the new replication security test (**/test:CheckSecurityError**), you must run Dcdiag on a domain controller running Windows Server 2003 with SP1.
  - You can run the new Dcdiag replication security tests against domain controllers running the following operating systems:

    Windows 2000 Server with SP3

    Windows 2000 Server with SP4

    Windows Server 2003

    Windows Server 2003 with SP1

**To test the connection between two domain controllers for replication security errors**

1. At a command prompt, type the following command, and then press ENTER:

   **dcdiag /test:CheckSecurityError /ReplSource:** *SourceDomainControllerName*

   where *SourceDomainControllerName* is the DNS name, NetBIOS name, or distinguished name of the real or potential

"from" server that is represented by a real or potential connection object that you want to test. This command tests the connection between the domain controller on which you run the command and the source domain controller.

2. Copy the report into Notepad or an equivalent text editor.
3. Scroll to the Summary table near the bottom of the Dcdiag log file.
4. Note the names of all domain controllers that reported “Warn” or “Fail” status in the Summary table
5. Find the detailed breakout section for the problem domain controller by searching on the string “DC: *DomainControllerName*”.
6. Make the required configuration changes on the domain controllers.
7. Rerun **Dcdiag /test:CheckSecurityError /ReplSource:***SourceDomainControllerName* to validate configuration changes.

# Fixing Replication DNS Lookup Problems (Event IDs 1925, 2087, 2088)

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain controllers running Windows 2000 Server or Windows Server 2003 cannot replicate Active Directory updates if Domain Name System (DNS) lookup failures prevent a destination domain controller from contacting its source replication partner to request changes. Lookup failures occur when a destination domain controller cannot resolve its source replication partner's globally unique identifier (GUID)-based canonical name (CNAME) resource record to an Internet Protocol (IP) address by using DNS. GUID-based CNAME resource records are always registered in the DNS zone _msdcs.*ForestRootDomain*.

The most common DNS failures occur when DNS client settings are misconfigured on the destination or source domain controller, or the direct and intermediate DNS servers that are used to resolve the query are misconfigured. Network problems or domain controller disconnection problems might also be present. If the problem is due to DNS configuration errors or replication latency, the effect on Active Directory replication is minimized by new improvements to domain controller name resolution in Windows Server 2003 Service Pack 1 (SP1).

**Improvements to Domain Controller Name Resolution in SP1**
Domain controllers running Windows Server 2003 with SP1 have a more robust response to DNS name resolution failures. Rather than failing on the first attempt to resolve the IP address of a source domain controller by using its CNAME resource record, destination domain controllers running Windows Server 2003 with SP1 use alternate names to attempt resolution and also log events that report problems and prescribe solutions.

**DNS Failure Scenarios**
All domain controllers register multiple service location (SRV) resource records in DNS as well as a host address (A) resource record for each IP address of the domain controller, plus an additional host resource record for each IP address if the domain controller is a global catalog server. In addition, each domain controller registers a single CNAME resource record.

The following table shows the DNS resource records that are required for proper Active Directory functionality.

| Mnemonic | Type | DNS resource record |
|---|---|---|
| pdc | SRV | _ldap._tcp.pdc._msdcs.*DnsDomainName* |
| gc | SRV | _ldap_tcp.gc._msdcs.*DnsForestRootDomainName* |
| GclpAddress | A | _gc._msdcs.*DnsForestRootDomainName* |
| DsaCname | CNAME | *DsaGuid*._msdcs.*DnsForestRootDomainName* |
| kdc | SRV | _kerberos._tcp.dc._msdcs.*DnsDomainName* |
| dc | SRV | _ldap._tcp.dc._msdcs.*DnsDomainName* |
| None | A | *DomainControllerFQDN* |

In the CNAME resource record (*DSA_GUID*._msdcs.*ForestRootDNSDomainName*), *DSA_GUID* is the GUID of the NTDS Settings object (also called the Directory System Agent (DSA) object) for the domain controller. *ForestRootDNSDomainName* is the DNS name of the forest where the domain controller is located. Destination domain controllers use the CNAME resource record to identify and locate their replication partners.

The Net Logon service on the domain controller registers all SRV resource records when the operating system starts up and at regular intervals thereafter. The DNS client service on the domain controller registers the DNS host A resource record.

A domain controller uses the following steps to locate its replication partner:

1. The destination domain controller queries its DNS server to look for the CNAME resource record of its replication partner. On domain controllers running Windows 2000 Server or Windows Server 2003 with no service pack applied, if this lookup fails to resolve the CNAME resource record to an IP address, DNS lookup (and replication) fails.

2. On domain controllers running Windows Server 2003 with SP1, if the CNAME lookup is unsuccessful, the domain

controller looks for the DNS A resource record of its replication partner. For example, the domain controller looks for DC03.corp.contoso.com.

3. If the DNS A resource record lookup is unsuccessful, the domain controller performs a network basic input/output system (NetBIOS) broadcast by using the host name of its replication partner. For example, the domain controller uses DC03.

When lookups fail, events that describe the condition are logged in the Directory Service event log.

**DNS Events for Lookup Failure**
Two new events, event ID 2087 and event ID 2088, are logged on destination domain controllers running Windows Server 2003 with SP1:

- If all lookups fail, event ID 2087 is logged.
- If lookup succeeds but either the first or second attempt fails, event ID 2088 is logged.

On domain controllers running Windows 2000 Server or Windows Server 2003 with no service pack applied, the destination domain controller that cannot successfully locate its replication partner in DNS logs event ID 1925.

Regardless of whether replication succeeds or fails, if you receive event ID 1925, event ID 2087, or event ID 2088, you should investigate and correct the cause of the failure because incorrect DNS configuration can affect other essential operations — including logon authentication and access to network resources — on member computers, domain controllers, and application servers. In addition, although fallback name resolution might allow replication to occur, it introduces unnecessary latency and overhead into the replication process.

**DNS Requirements for CNAME Lookup Success**
Although name resolution in Windows Server 2003 with SP1 is more aggressive at ensuring that replication can occur when a CNAME lookup fails, failure of this method indicates that either the DNS clients or DNS servers are not configured properly. It is important to understand the requirements for successful CNAME lookup and to ensure that DNS is functioning accordingly.

Resolving the fully qualified, GUID-based, CNAME resource record of the source domain controller to the current IP address of the source domain controller requires the following DNS configurations:

1. In their respective TCP/IP client settings, the source domain controller and destination domain controller must be configured to resolve DNS names by using only valid DNS servers that directly host, forward, or delegate to the following DNS zones:

   a. _msdcs.*ForestRootDNSDomainName*, to resolve queries for computers in the forest.

   b. The DNS zone that corresponds to the primary DNS suffix of the respective target domain controller, to resolve queries for computers in the domain. (The source domain controller can resolve the domain name of the target domain controller, and the reverse is also true.) The primary DNS suffix is usually the same as the DNS name of the domain to which a computer is joined. You can view the primary DNS suffix in the properties of My Computer.

   If the DNS servers that the source domain controller is configured to use for name resolution do not host these zones directly, the DNS servers that are used must forward or delegate to DNS servers that do host these zones.

2. The source domain controller must have successfully registered the following resource records:

   - GUID-based CNAME resource record in the DNS zone _msdcs.*ForestRootDNSDomainName*
   - Host A resource record in the DNS zone that corresponds to its primary DNS suffix

**Potential Preliminary Failures Due to Replication Latency**
At the time that the destination domain controller queries its DNS servers for the location of its source replication partner, DNS configurations might be correct on both the source and destination domain controllers, but DNS resource record registrations might be in flux as a result of configuration changes on the source domain controller. In this case, DNS lookup can fail as a

result of replication latency, as follows:

- If the source domain controller changes the DNS server on which it registers its CNAME and host A resource records, it is possible that the initial DNS server that the destination domain controller queries to resolve the name of the source domain controller is different than any of the DNS servers on which the CNAME and host A resource records for the source domain controller are currently registered. In this case, DNS replication latency or failures might prevent DNS records that are successfully registered on the DNS servers that the source controller uses from being located by the DNS server that is queried by the destination domain controller.

- If the Active Directory domain of the DNS server that the destination domain controller uses initially has a parent-child relationship with the Active Directory domain of the servers on which the source domain controller registers its resource records, the forwarder and delegation configuration on both the DNS servers that the source domain controller uses and the DNS servers that the destination domain controller uses, as well as any intermediate DNS servers that are used to resolve the DNS query, must be valid. Any required records on those DNS servers might be subject to replication latency and failure.

Understanding these basic requirements for name resolution that locates the source replication partner provides a more meaningful context for working through solutions when you have replication DNS lookup problems. Choose a problem from the following list that best describes your situation, and then step through the suggested fix:

Event ID 1925: Attempt to establish a replication link failed due to DNS lookup problem

Event ID 2087: DNS lookup failure caused replication to fail

Event ID 2088: DNS lookup failure occurred with replication success

# Event ID 1925: Attempt to establish a replication link failed due to DNS lookup problem

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you receive event ID 1925 with the error message that Domain Name System (DNS) lookup failed, inbound replication of a directory partition has failed on the destination domain controller, and you must fix the DNS problem.

An example of the event text is as follows:

```
Event Type:Warning
Event Source:NTDS KCC
Event Category:Knowledge Consistency Checker
Event ID:1925
Date:3/24/2005
Time:9:15:46 AM
User:NT AUTHORITY\ANONYMOUS LOGON
Computer:DC3
Description:
The attempt to establish a replication link for the following
writable directory partition failed.

Directory partition:
CN=Configuration,DC=contoso,DC=com
Source domain controller:
CN=NTDS Settings,CN=DC1,CN=Servers,CN=Default-First-Site-
Name,CN=Sites,CN=Configuration,DC=contoso,DC=com
Source domain controller address:
f8786828-ecf5-4b7d-ad12-8ab60178f7cd._msdcs.contoso.com
Intersite transport (if any):

This domain controller will be unable to replicate with the
source domain controller until this problem is corrected.

User Action
Verify if the source domain controller is accessible or
network connectivity is available.

Additional Data
Error value:
8524 The DSA operation is unable to proceed because of a DNS
lookup failure.
```

**Solution**

Proceed with DNS testing as described in "Event ID 2087: DNS lookup failure caused replication to fail."

# Event ID 2087: DNS lookup failure caused replication to fail

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a destination domain controller running Windows Server 2003 with Service Pack 1 (SP1) receives event ID 2087 in the Directory Service event log, attempts to resolve the globally unique identifier (GUID) in the canonical name (CNAME) resource record, the fully qualified domain name (FQDN), and the network basic input/output system (NetBIOS) name to the Internet Protocol (IP) address of the source domain controller have all failed. Failure to locate the source replication partner prevents replication with that source until the problem is fixed.

An example of the event text is as follows:

```
Event Type:Error
Event Source:NTDS Replication
Event Category:DS RPC Client
Event ID:2087
Date:3/9/2005
Time:11:00:21 AM
User:NT AUTHORITY\ANONYMOUS LOGON
Computer:DC3
Description:
Active Directory could not resolve the following DNS host name of
the source domain controller to an IP address. This error prevents
additions, deletions and changes in Active Directory from replicating
between one or more domain controllers in the forest. Security
groups, group policy, users and computers and their passwords will
be inconsistent between domain controllers until this error is
resolved, potentially affecting logon authentication and access
to network resources.

Source domain controller:
dc2
Failing DNS host name:
b0069e56-b19c-438a-8a1f-64866374dd6e._msdcs.contoso.com

NOTE: By default, only up to 10 DNS failures are shown for any
given 12 hour period, even if more than 10 failures occur.  To
log all individual failure events, set the following diagnostics
registry value to 1:

Registry Path:
HKLM\System\CurrentControlSet\Services\NTDS\Diagnostics\22 DS RPC Client

User Action:

1) If the source domain controller is no longer functioning or
its operating system has been reinstalled with a different
computer name or NTDSDSA object GUID, remove the source domain
controller's metadata with ntdsutil.exe, using the steps outlined
in MSKB article 216498.

2) Confirm that the source domain controller is running Active
Directory and is accessible on the network by typing
"net view \\<source DC name>" or "ping <source DC name>".

3) Verify that the source domain controller is using a valid
DNS server for DNS services, and that the source domain
controller's host record and CNAME record are correctly
registered, using the DNS Enhanced version of DCDIAG.EXE
available on http://www.microsoft.com/dns

dcdiag /test:dns

4) Verify that that this destination domain controller is using
a valid DNS server for DNS services, by running the DNS Enhanced
version of DCDIAG.EXE command on the console of the destination
domain controller, as follows:
```

```
dcdiag /test:dns

5) For further analysis of DNS error failures see KB 824449:
http://support.microsoft.com/?kbid=824449

Additional Data
Error value:
11004 The requested name is valid, but no data of the requested
type was found.
```

**Cause**

Failure to resolve the current CNAME resource record of the source domain controller to an IP address can have the following causes:

- The source domain controller is powered off, is offline, or resides on an isolated network, and Active Directory and Domain Name System (DNS) data for the offline domain controller has not been deleted to indicate that the domain controller is inaccessible.
- One of the following conditions exists:
  - The source domain controller has not registered its resource records in DNS.
  - The destination domain controller is configured to use an invalid DNS server.
  - The source domain controller is configured to use an invalid DNS server.
  - The DNS server that is used by the source domain controller does not host the correct zones or the zones are not configured to accept dynamic updates.
  - The direct DNS servers that are queried by the destination domain controller cannot resolve the IP address of the source domain controller as a result of nonexistent or invalid forwarders or delegations.
- Active Directory has been removed on the source domain controller and then reinstalled with the same IP address, but knowledge of the new NTDS Settings GUID has not reached the destination domain controller.
- Active Directory has been removed on the source domain controller and then reinstalled with a different IP address, but the current host address (A) resource record for the IP address of the source domain controller is either not registered or does not exist on the DNS servers that are queried by the destination domain controller as a result of replication latency or replication error.
- The operating system of the source domain controller has been reinstalled with a different computer name, but its metadata either has not been removed or has been removed and not yet inbound-replicated by the destination domain controller.

**Solution**

First, determine whether the source domain controller is functioning. If the source domain controller is not functioning, remove its remaining metadata from Active Directory.

If the source domain controller is functioning, continue with procedures to diagnose and solve the DNS problem, as needed:

- Use Dcdiag to diagnose DNS problems.
- Register DNS SRV resource records plus host records.
- Synchronize replication between the source and destination domain controllers.
- Verify consistency of the NTDS Settings GUID.

#### Determine Whether a Domain Controller Is Functioning

To determine whether the source domain controller is functioning, use the following test.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Users group in the domain of the domain controller.
- Tools: Net view

#### To determine whether a domain controller is functioning

- To confirm that the domain controller is running Active Directory and is accessible on the network, at a command prompt type the following command, and then press ENTER:

  **net view ** *SourceDomainControllerName*

  where *SourceDomainControllerName* is the NetBIOS name of the domain controller.

This command displays the Netlogon and SYSVOL shares, indicating that the server is functioning as a domain controller. If this test shows that the domain controller is not functioning on the network, determine the nature of the disconnection and whether the domain controller can be recovered or whether its metadata must be removed from Active Directory manually. If the domain controller is not functioning and cannot be restored, use the procedure in the following section, "Clean Up Domain Controller Metadata," to delete the data from Active Directory that is associated with that server.

#### Clean Up Domain Controller Metadata

If tests show that the domain controller is no longer functioning but you still see objects representing the domain controller in Active Directory Sites and Services, replication will continue to be attempted, and you must remove these objects from Active Directory manually. You must use Ntdsutil to clean up (delete) the metadata for the defunct domain controller.

If the defunct domain controller is the last domain controller in the domain, you should also remove the metadata for the domain. Allow sufficient time for all global catalog servers in the forest to inbound-replicate the domain deletion before promoting a new domain with the same name.

The process for cleaning up metadata is improved in the version of Ntdsutil that is included with Windows Server 2003 SP1. Instructions for cleaning up metadata with the Windows Server 2003 version of Ntdsutil and the Windows Server 2003 SP1 version of Ntdsutil are provided in the following procedure.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Enterprise Admins group.
- Tools: Ntdsutil (System32 command-line tool)

#### To clean up server metadata

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **ntdsutil**

3. At the **ntdsutil:** command prompt, type the following command, and then press ENTER:

   **metadata cleanup**

4. Perform metadata cleanup as follows:

| **Note** |
| --- |
| If you are removing domain metadata as well as server metadata, skip the following procedure and use the procedure that begins at step a. |

- If you are performing server metadata cleanup only and you are using the version of Ntdsutil.exe that is included with Windows Server 2003 SP1, at the **metadata cleanup:** command prompt, type the following, and then press ENTER:

  **remove selected server** *ServerName*

  Or

  **remove selected server** *ServerName1***on***ServerName2*

| Value | Description |
|---|---|
| *ServerName*, *ServerName1* | The distinguished name of the domain controller whose metadata you want to remove, in the form cn=*ServerName*,cn=Servers,cn=*SiteName*, cn=Sites,cn=Configuration,dc=*ForestRootDomain* |
| *ServerName2* | The DNS name of the domain controller to which you want to connect and from which you want to remove server metadata |

- If you are performing metadata cleanup by using the version of Ntdsutil.exe that is included with Windows Server 2003 with no service pack, or if you are performing both domain metadata cleanup and server metadata cleanup, perform metadata cleanup as follows:

a. At the **metadata cleanup:** command prompt, type the following command, and then press ENTER:

   **connection**

b. At the **server connections:** command prompt, type the following command, and then press ENTER:

   **connect to server** *Server*

c. At the **connection:** command prompt, type the following command, and then press ENTER:

   **quit**

d. At the **metadata cleanup:** command prompt, type the following command, and then press ENTER:

   **select operation target**

e. At the **select operation target:** command prompt, type the following command, and then press ENTER:

   **list sites**

f. A numbered list of sites appears. Type the following command, and then press ENTER:

   **select site** *SiteNumber*

g. At the **select operation target:** command prompt, type the following command, and then press ENTER:

   **list domains in site**

h. A numbered list of domains in the selected site appears. Type the following command, and then press ENTER:

   **select domain** *DomainNumber*

i. At the **select operation target:** command prompt, type the following command, and then press ENTER:

   **list servers in site**

j. A numbered list of servers in a domain and site is displayed. Type the following command, and then press ENTER:

   **select server** *ServerNumber*

k. At the **select operation target:** command, type the following command, and then press ENTER:

   **quit**

l. At the **metadata cleanup:** command, type the following command, and then press ENTER:

   **remove selected server**

m. If the server whose metadata you have removed is the last domain controller in the domain and you want to remove the domain metadata, at the **metadata cleanup:** command prompt, type the following command, and then press ENTER:

   **remove selected domain**

   Metadata for the domain that you selected in step h is removed.

n. At the **metadata cleanup:** and **ntdsutil:** command prompts, type **quit**, and then press ENTER.

| Value | Description |
|---|---|
| *Server* | The DNS name of a domain controller that you want to connect to |
| *SiteNumber* | The number that is associated with the site of the server that you want to clean up, which appears in the list |
| *DomainNumber* | The number that is associated with the domain of the server that you want to clean up, which appears in the list |
| *ServerNumber* | The number that is associated with the server that you want to clean up, which appears in the list |

**Use Dcdiag to Diagnose DNS Problems**

If the domain controller is functioning online, continue by using Dcdiag to diagnose and fix DNS problems that might be interfering with Active Directory replication.

Use the following procedures to complete this process:

- Verify connectivity and basic DNS functionality.
- Verify registration of the CNAME resource record in DNS.
- Verify and enable secure dynamic updates.

Before you begin these procedures, gather the following information, which is contained in the event ID 2087 message text:

- The FQDN of the source domain controller and destination domain controller
- The IP address of the source domain controller

The updated version of Dcdiag that is included in Windows Support Tools in Windows Server 2003 SP1 contains tests that provide consolidated and improved testing of basic and advanced DNS features. You can use this tool to diagnose basic DNS functionality and dynamic updates.

When you use the enhanced SP1 version of Dcdiag for DNS testing, there are specific requirements that do not apply to all Dcdiag tests.

**Requirements**

- Administrative credentials: To complete the new DNS tests that are available in the SP1 version of Dcdiag, you must be a member of the Enterprise Admins group.
- Tools: Dcdiag.exe
- Operating system:
  - You can run the enhanced version of Dcdiag on computers running the following operating systems:

    Windows XP Professional

    Windows Server 2003

    Windows Server 2003 with SP1
  - You can run the new Dcdiag DNS tests against Microsoft DNS servers that are installed on domain controllers running the following operating systems:

    Windows 2000 Server with Service Pack 3 (SP3)

    Windows 2000 Server with Service Pack 4 (SP4)

    Windows Server 2003

    Windows Server 2003 with SP1

| **Note** |
| --- |
| You can use the **/f:** switch in Dcdiag commands to save the output to a text file. Use **/f:***FileName* to generate the file in the location that is indicated in *FileName*, for example, **/f:c:\Test\DnsTest.txt**. |

**Verify Basic DNS Functionality**

To verify the settings that might interfere with Active Directory replication, you can begin by running the basic DNS test that ensures that DNS is operating properly on the domain controller.

The basic DNS test checks the following:

- Connectivity: The test determines whether domain controllers are registered in DNS, can be contacted by PING, and have Lightweight Directory Access Protocol / remote procedure call (LDAP/RPC) connectivity. If the connectivity test fails on a domain controller, no other tests are run against that domain controller. The connectivity test is performed automatically before any other DNS test is run.
- Essential services: The test confirms that the following services are running and available on the tested domain controller:
  - DNS Client service
  - Net Logon service
  - Key Distribution Center (KDC) service
  - DNS Server service (if DNS is installed on the domain controller)
- DNS client configuration: The test confirms that DNS servers on all adapters are reachable.

- Resource record registrations: The test confirms that the address (A) resource record of each domain controller is registered on at least one of the DNS servers that is configured on the client.

- Zone and server of authority (SOA): If the domain controller is running the DNS Server service, the test confirms that the Active Directory domain zone and SOA record for the Active Directory domain zone are present.

- Root zone: Checks if the root (.) zone is present.

**To verify basic DNS functionality**

1. At a command prompt, type the following command, and then press ENTER:

   **dcdiag /test:dns /s:** *SourceDomainControllerName* **/DnsBasic**

   where *SourceDomainControllerName* is the distinguished name, NetBIOS name, or DNS name of the domain controller.

   As an alternative, you can test all domain controllers in the forest by typing **/e:** instead of **/s:**.

2. Copy the report into Notepad or an equivalent text editor.

3. Scroll to the Summary table near the bottom of the Dcdiag log file.

4. Note the names of all domain controllers that report "Warn" or "Fail" status in the Summary table.

5. Find the detailed breakout section for the problem domain controller by searching on the string "DC: *DomainControllerName*".

6. Make the required configuration changes on DNS clients and DNS servers.

7. To validate the configuration changes, rerun **Dcdiag /test:DNS** with the **/e:** or **/s:** switch.

If the basic DNS test shows no errors, continue by verifying that resource records that are used to locate domain controllers are registered in DNS.

**Verify Resource Record Registration**
The destination domain controller uses the DNS CNAME resource record to locate its source domain controller replication partner. Although domain controllers running Windows Server 2003 with SP1 can locate source replication partners by using FQDNs or, if that fails, NetBIOS names, the presence of the CNAME record is expected and should be verified for proper DNS functioning.

You can use Dcdiag to verify registration of all resource records that are essential for domain controller location by using the **dcdiag /test:dns /DnsRecordRegistration** test. This test verifies registration of the following resource records in DNS:

- CNAME (the GUID-based resource record that locates a replication partner)

- A (the host resource record that contains the IP address of the domain controller)

- LDAP SRV (the service resource records that locate LDAP servers)

- GC SRV (the service resource records that locate global catalog servers)

- PDC SRV (the service resource records that locate primary domain controller (PDC) operations masters)

As an alternative, you can use the following procedure to check for only the CNAME resource record.

**To verify CNAME resource record registration**

1. In the DNS console, locate any domain controller that is running the DNS Server service, where the server hosts the DNS zone with the same name as the Active Directory domain of the domain controller.

2. In the console tree, click the zone that is named _msdcs.*Dns_Domain_Name*.

| Note |
| --- |
| In Windows 2000 Server DNS, _msdcs.*Dns_Domain_Name* is a subdomain of the DNS zone for the Active Directory domain name. In Windows Server 2003 DNS, _msdcs.*Dns_Domain_Name* is a separate zone. |

3. In the details pane, verify that the following resource records are present:
   - A CNAME resource record that is named *Dsa_Guid*._msdcs.*Dns_Domain_Name*
   - A corresponding A resource record for the name of the DNS server

If the CNAME resource record is not registered, verify that dynamic updates are functioning properly. Use the test in the following section.

**Verify Dynamic Updates**

If the basic DNS test shows that resource records do not exist in DNS, use the dynamic update test to diagnose why the Net Logon service did not register the resource records automatically. To verify that the Active Directory domain zone is configured to accept secure dynamic updates and to perform registration of a test record (_dcdiag_test_record), use the following procedure. The test record is deleted automatically after the test.

**To verify dynamic updates**

- At a command prompt, type the following command, and then press ENTER:

  **dcdiag /test:dns /s:** *SourceDomainControllerName* **/DnsDynamicUpdate**

  where *SourceDomainControllerName* is the distinguished name, NetBIOS name, or DNS name of the domain controller.

  As an alternative, you can test all domain controllers by using the **/e:** switch instead of the **/s:** switch.

If secure dynamic update is not configured, use the following procedure to configure it.

**To enable secure dynamic updates**

1. Open the DNS console.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, verify that the zone type is Active Directory–integrated.
4. In **Dynamic Updates**, click **Secure only**.

**Register DNS Resource Records**

If DNS resource records do not appear in DNS for the source domain controller, you have verified dynamic updates, and you want to register DNS resource records immediately, you can force the registration manually by using the following procedure. The Net Logon service on a domain controller registers the DNS resource records that are required for the domain controller to be located on the network. The DNS Client service registers the host A resource record that the CNAME record points to.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the forest root domain or a member of the Enterprise Admins group.
- Tools: **net stop/start**, **ipconfig**

**To register DNS resource records manually**

1. To initiate registration of domain controller Locator resource records manually on the source domain controller, at a command prompt type the following commands, and then press ENTER after each command:

   **net stop net logon & net start net logon**

2. To initiate registration of the host A resource record manually, at a command prompt type the following command, and then press ENTER:

**ipconfig /flushdns & ipconfig /registerdns**

3. Wait 15 minutes, and then review events in Event Viewer to ensure proper registration of the resource records.

Repeat the procedure in the "Verify Resource Record Registration" section earlier in this guide to verify that the resource records appear in DNS.

#### Synchronize Replication Between the Source and Destination Domain Controllers

After you complete DNS testing, use the following procedure to synchronize replication on the inbound connection from the source domain controller to the destination domain controller.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the domain of the destination domain controller.

- Tool: Active Directory Sites and Services

#### To synchronize replication from a source domain controller

1. Open Active Directory Sites and Services.

2. In the console tree, double-click the **Sites** container, double-click the site of the domain controller to which you want to synchronize replication, double-click the **Servers** container, double-click the server object of the domain controller, and then click **NTDS Settings**.

3. In the details pane, in the **From Server** column, locate the connection object that shows the name of the source domain controller.

4. Right-click the appropriate connection object, and then click **Replicate Now**.

5. Click **OK**.

If replication does not succeed, use the procedure in the following section to verify consistency of the NTDS Settings GUID.

#### Verify Consistency of the NTDS Settings GUID

If you have performed all DNS tests and other tests and replication does not succeed, use the following procedure to verify that the GUID of the NTDS Settings object that the destination domain controller is using to locate its replication partner matches the GUID that is currently in effect on the source domain controller itself. To perform this test, you view the object GUID as it appears in the local directories of each domain controller.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the domain of the destination domain controller.

- Tool: Ldp (Windows Support Tools)

#### To verify consistency of the NTDS Settings GUID

1. Click **Start**, click **Run**, type **Ldp**, and then click **OK**.

2. On the **Connection** menu, click **Connect**.

3. In the **Connect** dialog box, leave the **Server** box empty.

4. In **Port**, type **389**, and then click **OK**.

5. On the **Connection** menu, click **Bind**.

6. In the **Bind** dialog box, provide Enterprise Admins credentials. If it is not already selected, click **Domain**.

7. In **Domain**, type the name of the forest root domain, and then click **OK**.

8. On the **View** menu, click **Tree**.

9. In the **Tree View** dialog box, type:

   **CN=Configuration,DC=Forest_Root_Domain**

   and then click **OK**.

10. Navigate to the object CN=NTDS Settings,CN=*SourceServerName*,CN=Servers,CN=*SiteName*, CN=Sites,CN=configuration,DC=*ForestRootDomain*.

11. Double-click the **NTDS Settings** object and, in the details pane, view the value for the attribute **objectGUID**. Right-click that value, and then copy it to Notepad.

12. On the **Connection** menu, click **Disconnect**.

13. Repeat steps 2 through 11, but in step 3, type the name of the source domain controller, for example, DC03.

14. In Notepad, compare the values of the two GUIDs.

15. If the values do not match, the destination domain controller must receive replication of the valid GUID. Check the GUID value on other domain controllers and attempt replication on the destination domain controller with a different domain controller that has the correct GUID.

16. If the values match, verify that the GUID matches the GUID in the *Dsa_Guid*._msdcs.*Dns_Domain_Name* resource record for the source domain controller, as follows:

    a. Note the primary DNS servers that each domain controller identifies in the TCP/IP properties in their Network Settings. All the DNS servers that are listed in the respective TCP/IP properties should be able to indirectly or directly resolve this CNAME resource record.

    b. From the servers that are listed, identify the authoritative name server or servers for this domain zone by looking at the server names that are listed for the name server (NS) resource records at the root of the zone. (In the DNS console, select the forward lookup zone for the root domain, and view the NS records in the details pane).

    c. On the name server or servers obtained in step b, open the DNS console, and double-click the forward lookup zone for the forest root domain name. Double-click the _msdcs folder, and note the CNAME resource records that exist for your server name.

    d. If there are no records present or the records are incorrect, see article 241505, SRV Records Missing After Implementing Active Directory and Domain Name System, on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=69994).

# Event ID 2088: DNS lookup failure occurred with replication success

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a destination domain controller running Windows Server 2003 with Service Pack 1 (SP1) receives event ID 2088 in the Directory Service event log, attempts to resolve the globally unique identifier (GUID) in the canonical name (CNAME) resource record to an Internet Protocol (IP) address for the source domain controller failed. However, the destination domain controller tried other means to resolve the name and succeeded by using either the fully qualified domain name (FQDN) or the network basic input/output system (NetBIOS) name of the source domain controller. Although replication was successful, the DNS problem should be diagnosed and resolved.

An example of the event text is as follows:

```
Event Type:Warning
Event Source:NTDS Replication
Event Category:DS RPC Client
Event ID:2088
Date:3/21/2005
Time:2:29:34 PM
User:NT AUTHORITY\ANONYMOUS LOGON
Computer:DC3
Description:
Active Directory could not use DNS to resolve the IP address of the
source domain controller listed below. To maintain the consistency
of Security groups, group policy, users and computers and their passwords,
Active Directory successfully replicated using the NetBIOS or fully
qualified computer name of the source domain controller.

Invalid DNS configuration may be affecting other essential operations on
member computers, domain controllers or application servers in this
Active Directory forest, including logon authentication or access to network
resources.

You should immediately resolve this DNS configuration error so that
this domain controller can resolve the IP address of the source
domain controller using DNS.

Alternate server name:
dc1
Failing DNS host name:
4a8717eb-8e58-456c-995a-c92e4add7e8e._msdcs.contoso.com

NOTE: By default, only up to 10 DNS failures are shown for any given
12 hour period, even if more than 10 failures occur.  To log all
individual failure events, set the following diagnostics registry
value to 1:

Registry Path:
HKLM\System\CurrentControlSet\Services\NTDS\Diagnostics\22 DS RPC Client

User Action:

1) If the source domain controller is no longer functioning or its
operating system has been reinstalled with a different computer
name or NTDSDSA object GUID, remove the source domain controller's
metadata with ntdsutil.exe, using the steps outlined in MSKB article 216498.

2) Confirm that the source domain controller is running Active Directory
and is accessible on the network by typing "net view \\<source DC name>"
or "ping <source DC name>".

3) Verify that the source domain controller is using a valid DNS server
for DNS services, and that the source domain controller's host record
and CNAME record are correctly registered, using the DNS Enhanced
version of DCDIAG.EXE available on http://www.microsoft.com/dns
```

```
dcdiag /test:dns

4) Verify that that this destination domain controller is using a
valid DNS server for DNS services, by running the DNS Enhanced
version of DCDIAG.EXE command on the console of the destination
domain controller, as follows:

dcdiag /test:dns

5) For further analysis of DNS error failures see KB 824449:
http://support.microsoft.com/?kbid=824449

Additional Data
Error value:
11004 The requested name is valid, but no data of the requested
type was found
```

**Cause**

Failure to resolve the source domain controller name by using the CNAME resource record in DNS can be due to DNS misconfigurations or delays in DNS data propagation.

**Solution**

Proceed with DNS testing as described in "Event ID 2087: DNS lookup failure caused replication to fail."

# Fixing Replication Connectivity Problems (Event ID 1925)

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Network connectivity problems can make it impossible for domain controllers to form replication partnerships. Various events and errors can indicate a problem with network connectivity that is preventing replication from occurring.

Use the procedures in Event ID 1925: Attempt to establish a replication link failed due to connectivity problem to diagnose and fix replication connectivity problems.

# Event ID 1925: Attempt to establish a replication link failed due to connectivity problem

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The description text in event ID 1925 reports that the attempt to establish a replication link for the following writable directory partition failed, and the description text provides the distinguished name of the directory partition that the destination is attempting to replicate from the source. The error code in the event gives more specific information about the cause of the problem.

An example of the event text is as follows:

```
Event Type:Warning
Event Source:NTDS KCC
Event Category:Knowledge Consistency Checker
Event ID:1925
Date:3/24/2005
Time:9:15:46 AM
User:NT AUTHORITY\ANONYMOUS LOGON
Computer:DC3
Description:
The attempt to establish a replication link for the following
writable directory partition failed.

Directory partition:
CN=Configuration,DC=contoso,DC=com
Source domain controller:
CN=NTDS Settings,CN=DC1,CN=Servers,CN=Default-First-Site-
Name,CN=Sites,CN=Configuration,DC=contoso,DC=com
Source domain controller address:
f8786828-ecf5-4b7d-ad12-8ab60178f7cd._msdcs.contoso.com
Intersite transport (if any):
CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=contoso,DC=com
This domain controller will be unable to replicate with the source
domain controller until this problem is corrected.

User Action
Verify if the source domain controller is accessible or network
connectivity is available.

Additional Data
Error value:
1908 Could not find the domain controller for this domain.
```

**Cause**

When event ID 1925 contains error 1908, "Could not find the domain controller for this domain," Active Directory replication has failed as a result of a connectivity problem between the domain controller that reported the error and the source domain controller that is named in the event text.

**Solution**

Use the following tests to solve this problem:

- Verify wide area network (WAN) connectivity.
- Determine the maximum packet size, and change it if necessary.
- Force replication, and capture replication traffic in Network Monitor.
- Analyze network traces to see if any traffic is not reaching the source domain controller.

**Verify WAN Connectivity**

Verify that there are no basic connectivity problems with the underlying network between the domain controllers, especially if they are separated by a WAN link or firewalls. For information about testing this type of problem, see article 310099,

Description of the Portquery.exe command-line utility, on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=69995), and see article 159211, Diagnoses and treatment of black hole routers, on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=69996).

**Determine Maximum Packet Size**

By default, the Kerberos authentication protocol in Windows 2000, Windows XP, and Windows Server 2003 uses the User Datagram Protocol (UDP) when the data can be fit in packets of less than 2,000 bytes. Any data above this value uses TCP to carry the packets. Packets of more than 1,500 bytes are often dropped by a device such as a firewall on the network.

To avoid this problem, you can determine the size of packet that your network can accommodate. Then, you can edit the registry so that the maximum number of bytes for using UDP is set to the lowest value that you receive, less 8 bytes to account for header size.

Use the **ping** command to test the size of packets that the network can accommodate.

**Requirements**

- Administrative credentials: To complete this procedure, you must be a member of the Domain Users group and have the **Log on locally** right on the domain controller.
- Tool: PING

**To determine the lowest common packet size**

1. From the destination domain controller, ping the source domain controller by its Internet Protocol (IP) address. At a command prompt, type the following command, and then press ENTER:

   **ping** *IP_address* **-f -l 1472**

2. From the source domain controller, use the command in step 1 to ping the destination domain controller by its IP address.
3. If the **ping** command completes in both directions, no additional modification is required.
4. If the **ping** command fails in either direction, monotonically lower the number that you use in the **-l** parameter until you find the lowest common packet size that works between the source and destination domain controllers.

| **Note** |
| --- |
| The version of Dcdiag that is included with Windows Server 2003 SP1 Support Tools provides the following method to perform this test:<br><br>**dcdiag /test:CheckSecurityError /s:***SourceDomainControllerName* |

You can edit the registry to set the maximum size of packets to the value that you determined by the PING method, less 8 bytes to account for header size. As an alternative, you can edit the registry so that the maximum number of bytes for using UDP is always exceeded and Kerberos always uses TCP.

You can change the default value of 2,000 bytes by modifying the registry entry **MaxPacketSize** in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\ Kerberos\Parameters**. Use the following procedure to change this registry setting.

| **Caution** |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

**Requirements**

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the domain of the domain controller.

- Tool: Regedit.exe change the maximum packet size

**To change the maximum packet size**

1. Click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\ Kerberos\Parameters**.
3. Edit or, if it does not exist in the details pane, create the entry **MaxPacketSize** as follows:
   - To edit the entry if it exists in the details pane:

     Right-click **MaxPacketSize**; click **Modify**; and then, in the **Value data** box, type **1** to force Kerberos to use TCP, or type the value that you established to lower the value to the appropriate maximum size.

   - To create the entry if it does not exist in the details pane:

     Right-click **Parameters**, click **New DWORD Value**, type the name **MaxPacketSize**, and then go to step 3a to edit the entry.

4. Click **OK**.
5. You must restart the domain controller for this change to take effect.

For information about importing an Administrative Template into Group Policy so that this value can be set for all the Windows 2000–based, Windows Server 2003-based, or Windows XP-based computers in the enterprise, see article 244474, "How to force Kerberos to use TCP instead of UDP in Windows Server 2003, in Windows XP, and in Windows 2000," on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=69997).

**Capture Network Traces During Replication**
Use Network Monitor to capture simultaneous traces on both source and destination domain controllers while attempting to replicate (all traffic to and from each domain controller; plus, set the capture buffer to a sufficiently large value). You must select the addresses of the domain controllers from the address database and add them to the capture filter. Start the capture, and then start replication between the two domain controllers. Look for Kerberos fragmentation, out-of-order packets, latency, or network traffic that originates on one side of the connection and does not arrive at the other side.

For information about installing Network Monitor, see Network Monitor on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42987).

**Create an Address Database**
To use address pairs in a capture filter, you must first build an address database. After the database is created, you can use the addresses that are listed in the database to specify address pairs in a capture filter.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the forest root domain or the Enterprise Admins group.

- Tool: Network Monitor

**To create an address database**

1. Open Network Monitor.
2. If you are prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Start**.
4. On the **Capture** menu, click **Stop and View**.
5. On the **Display** menu, click **Find All Names**.

6. In the **Find All Names** dialog box, click **OK**. All addresses are added to the address database.
7. On the **Windows** menu, click the local connection.

You can use the names in the addresses database to specify address pairs in the capture filter.

**Capture Network Frames**

To capture frames that are sent from a specific computer on your network to your computer or that are sent from your computer to a specific computer on your network, specify one or more address pairs in a capture filter. You can monitor up to four address pairs simultaneously.

An address pair consists of:

- The addresses of the two computers between which you want to monitor traffic.
- Arrows that specify the traffic direction that you want to monitor.
- The INCLUDE or EXCLUDE keywords, which indicate how Network Monitor should respond to a frame that meets a filter's specifications.

**Requirements**

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the forest root domain or the Enterprise Admins Group in the forest.
- Tool: Network Monitor

**To capture network frames**

1. Open Network Monitor.
2. If you are prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Buffer Settings.**
4. **In the Capture Buffer Settings** dialog box, set the buffer and frame size as appropriate, and then click **OK**.
5. On the **Capture** menu, click **Filter**.
6. In the **Capture Filter** dialog box, double-click **Address Pairs**.
7. In the **Address Expressions** dialog box, select an address in Station 1 and an address in Station 2 for the computers whose traffic you want to capture.
8. In the **Direction** box, select one of the symbols:

   <--> to monitor the traffic that passes in either direction between the addresses that you have selected.

   --> or <-- to monitor only the traffic that passes in one direction between the computers..
9. Click **OK** twice.
10. On the **Capture** menu, click **Start**.

**Force Replication**

When you have Network Monitor started to capture traffic between the two domain controllers, use the following procedure to force synchronization between the computers so that you can capture the replication traffic in Network Monitor.

**Requirements**

- Credentials: To complete this procedure, you must be a member of the Domain Admins group in the forest root domain or the Enterprise Admins group in the forest.
- Tools: Active Directory Sites and Services (Administrative Tools)

**To synchronize replication from a source domain controller**

1. Open Active Directory Sites and Services.
2. Double-click the **Sites** container, double-click the site of the domain controller to which you want to synchronize replication, double-click the **Servers** container, double-click the server object of the domain controller, and then click **NTDS Settings**.
3. In the **From Server** column in the details pane, locate the connection object that shows the name of the source domain controller.
4. Right-click the appropriate connection object, and then click **Replicate Now**.
5. Click **OK**.

Analyze the traces from both domain controllers to see if there is any traffic that is not getting to the other domain controller. For information about using Network Monitor, see Network Monitor overview on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41936).

# Fixing Replication Topology Problems (Event ID 1311)

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Knowledge Consistency Checker (KCC) constructs and maintains the Active Directory replication topology automatically. Every 15 minutes, the KCC examines the sum of all directory partition replicas that reside on domain controllers in the forest, as well as administrator-defined settings for connections, sites, and site links.

Although generation of the replication topology occurs automatically, administrative configuration errors can result in an Active Directory replication topology that is inconsistent with the physical connections that are available. In Active Directory it is possible to create objects for which there is no physical network support. For example, Active Directory Sites and Services allows you to create a site object and assign subnet addresses that do not exist. The KCC will attempt to use these objects to create connections between domain controllers, but replication cannot occur because the network does not exist to support the replication topology as it is configured.

Event ID 1311 is logged in the Directory Service event log when the replication configuration information in Active Directory does not accurately reflect the physical topology of the network. Use the procedures in Event ID 1311: Replication configuration does not reflect the physical network to identify and fix topology problems.

# Event ID 1311: Replication configuration does not reflect the physical network

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Event ID 1311 is logged in the Directory Service log when configuration errors or unavailable domain controllers prevent replication of a directory partition between domain controllers in different sites.

An example of the event text is as follows:

```
Event Type:Error
Event Source:NTDS KCC
Event Category:Knowledge Consistency Checker
Event ID:1311
Date:3/9/2005
Time:6:39:58 PM
User:NT AUTHORITY\ANONYMOUS LOGON
Computer:DC3
Description:
The Knowledge Consistency Checker (KCC) has detected problems with
the following directory partition.

Directory partition:
CN=Configuration,DC=contoso,DC=com

There is insufficient site connectivity information in Active Directory
Sites and Services for the KCC to create a spanning tree replication
topology. Or, one or more domain controllers with this directory
partition are unable to replicate the directory partition information.
This is probably due to inaccessible domain controllers.

User Action
Use Active Directory Sites and Services to perform one of the
following actions:
- Publish sufficient site connectivity information so that the
KCC can determine a route by which this directory partition can
reach this site. This is the preferred option.
- Add a Connection object to a domain controller that contains
the directory partition in this site from a domain controller
that contains the same directory partition in another site.

If neither of the Active Directory Sites and Services tasks correct
this condition, see previous events logged by the KCC that identify
the inaccessible domain controllers.
```

**Cause**

This problem can have the following causes:

- Site link bridging is enabled on a network that does not support physical network connectivity between two domain controllers in different sites that are connected by a site link.
- **Bridge all site links** is enabled in Active Directory Sites and Services, but the network does not allow network connectivity between any two domain controllers in the forest.
- One or more sites are not contained in a site link.
- Site links contain all sites, but the site links are not interconnected. This condition is known as disjointed site links.
- One or more domain controllers are offline.
- Bridgehead domain controllers are online, but errors occur when they try to replicate a required directory partition between Active Directory sites.

- Administrator-defined preferred bridgehead servers are online, but they do not host the required directory partition. The most common misconfiguration is to define non–global catalog servers as bridgehead servers.

- Preferred bridgeheads are defined correctly by the administrator, but they are currently offline.

- The bridgehead server is overloaded because the server is undersized, too many branch sites are trying to replicate changes from the same hub domain controller, or the replication schedules on site links or connection objects are too frequent.

- The Knowledge Consistency Checker (KCC) has built an alternate path around an intersite connection failure, but it continues to retry the failing connection every 15 minutes.

**Solution**

Use the following procedures for troubleshooting event ID 1311:

- Identify the scope of the problem.

- Check site link bridging.

- Determine whether the network is fully routed.

- Verify that all sites are connected.

- Check preferred bridgehead servers.

**Identify the Scope of the Problem**

Identify the scope of the problem by determining whether event ID 1311 is being logged on all domain controllers in the forest that hold the intersite topology generator (ISTG) role or just on site-specific domain controllers.

First, use the following procedure to locate the ISTG role holders for all sites.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in a domain in the forest.

- Tool: Ldp (Windows Support Tools)

**To locate the ISTG role holders for all sites**

1. Click **Start**, click **Run**, type **Ldp**, and then click **OK**.
2. On the **Connection** menu, click **Connect**.
3. In the **Connect** dialog box, leave the **Server** box empty.
4. In **Port**, type **389**, and then click **OK**.
5. On the **Connection** menu, click **Bind**.
6. In the **Bind** dialog box, provide Enterprise Admins credentials. Click **Domain** if it is not already selected.
7. In **Domain**, type the name of the forest root domain, and then click **OK**.
8. On the **Browse** menu, click **Search**.
9. In **Base dn**, type:

   **CN=Sites,CN=Configuration,DC=** *Forest_Root_Domain*

10. In **Filter**, type:

    **(CN=NTDS Site Settings)**

11. For **Scope**, click **Subtree**.

12. Click **Options**, and in the **Attributes** box, scroll to the end of the list, type:

    **;interSiteTopologyGenerator**

    and then click **OK**.

13. In the **Search** dialog box, click **Run**.

14. Review the **interSiteTopologyGenerator** entries in the output, and make a note of the domain controller names.

Determine the scope of the event by checking the Directory Service event logs of all ISTG role holders in the forest, or check at least a significant number of ISTG role holders.

If event ID 1311 continues to be logged on ISTG role holders, continue with the next step.

#### Check Site Link Bridging

Use the following procedure to determine if site link bridging is enabled.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in a domain in the forest.

- Tool: Active Directory Sites and Services (Administrative Tools)

#### Determine if site link bridging is enabled

1. Open Active Directory Sites and Services.

2. In the console tree, double-click the **Sites** container, and then double-click the **Inter-Site Transports** container.

3. Right-click the **IP** container. If **Bridge all site links** is selected, site link bridging is enabled.

The **Bridge all site links** setting requires a fully routed network. If the network is not fully routed, you must create site link bridges manually.

#### Determine Whether the Network Is Fully Routed

Determine whether a fully routed network connection exists between two sites.

If the network is fully routed, continue by verifying that the sites are connected.

If the network is not fully routed and site link bridging is enabled, either make the network fully routed, or disable site link bridging and then create the necessary site links and site link bridges. For information about creating site links, see Linking Sites for Replication.

| **Note** |
| --- |
| Site link bridging is enabled by default. As a best practice, leave site link bridging enabled for fully routed networks. |

#### Disable Site Link Bridging

If the network is not fully routed and site link bridging is enabled, use the following procedure to disable site link bridging.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the forest root domain or a member of the Enterprise Admins group.

- Tool: Active Directory Sites and Services (Administrative Tools)

#### Determine if site link bridging is enabled

1. Open Active Directory Sites and Services.
2. In the console tree, double-click the **Sites** container, and then double-click the **Inter-Site Transports** container.
3. Right-click the **IP** container. If **Bridge all site links** is selected, click it to disable it.

#### Create a Site Link Bridge

If the network is not fully routed, be sure that you have created site links to connect all sites. When all site links are created, use the following procedure to create a site link bridge.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the forest root domain or a member of the Enterprise Admins group.
- Tool: Active Directory Sites and Services (Administrative Tools)

#### To create a site link bridge

1. Open Active Directory Sites and Services.
2. In the console tree, double-click the **Sites** container, and then expand the **Inter-Site Transports** container.
3. Right-click the **IP** container, and then click **New Site Link Bridge**.
4. In **Name**, type a name for the site link bridge.
5. Click two or more site links to be bridged, and then click **Add**.

Wait for a period of time that is twice as long as the longest replication interval in the forest. If event ID 1311 continues to be logged on ISTG role holders, continue with the next step.

#### Verify That All Sites Are Connected

If the network is fully routed, use the Repadmin command-line tool to view site links to ensure that intersite replication can occur between domain controllers in different sites.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Enterprise Admins group or the Domain Admins group in the forest root domain.
- Tool: Repadmin.exe (Windows Support Tools)

#### To view site links

1. At a command prompt, type the following command, and then press ENTER:

   **repadmin /showism "CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=** *Forest_Root_Domain* **"**

   where *Forest_Root_Domain* is the name of the forest root domain.

2. In the output, review the information for the sites that are listed. For each site, the output of the command shows a string of three numbers separated by colons. The numbers represent <cost>:<replication interval>:<options>. Strings with a value of "-1:0:0" indicate a possible missing site link.

#### Check Preferred Bridgehead Servers

If you have designated preferred bridgehead servers, the ISTG selects bridgehead servers only from that list of servers. If no servers in the list for the site can replicate a domain directory partition that has domain controllers in other sites, the ISTG selects a bridgehead server that can replicate the domain, if one is available in the site. However, if at least one server in the list can replicate a domain but the server is unavailable, the ISTG does not select an alternate bridgehead server and replication of updates to that domain does not occur in the site. In this case, you might have domain controllers that are capable of replicating the domain, but replication does not occur because preferred bridgehead servers have been selected and none is available for the domain.

Check the list of preferred bridgehead servers in the site, and ensure that preferred bridgehead servers for the domain in question are available. Use the following procedure to check the list of preferred bridgehead servers.

To see all servers that have been selected as preferred bridgehead servers in a forest, you can use ADSI Edit to view the **bridgeheadServerListBL** attribute on the IP container object.

**Requirements**

- Administrative credentials: To complete this procedure, you must be a member of the Domain Users group in a domain in the forest.
- Tool: Adsiedit.msc (Windows Support Tools)

**To view the list of preferred bridgehead servers**

1. Click **Start**, click **Run**, type **adsiedit.msc**, and then click **OK**.
2. In the console tree, double-click **Configuration Container**, and then double-click **CN=Configuration,DC=ForestRootDomainName**, **CN=Sites**, and **CN=Inter-Site Transports**.
3. Right-click **CN=IP**, and then click **Properties**.
4. In **Attributes**, double-click **bridgeheadServerListBL**.
5. If any preferred bridgehead servers are selected in any site in the forest, the **Values** box displays the distinguished name for each server object that is currently selected as a preferred bridgehead server.

Verify that all domain controllers in the list are online and functioning as domain controllers.

**Requirements**

- Administrative credentials: To complete this procedure, you must be a member of the Domain Users group in the domain of the domain controller.
- Tool: Net view

**To determine whether a domain controller is functioning**

- To confirm that a domain controller is running Active Directory and is accessible on the network, at a command prompt type the following command, and then press ENTER:

  **net view ** *DomainControllerName*

  where *DomainControllerName* is the network basic input/output system (NetBIOS) name of the domain controller.

This command displays the Netlogon and SYSVOL shares, indicating that the server is functioning as a domain controller. If this test shows that the domain controller is not functioning on the network, determine the nature of the disconnection and whether the domain controller can be recovered.

If a domain controller that is selected as a preferred bridgehead server is not available, use the following procedure to select another preferred bridgehead server in the site that can replicate the domain.

**Requirements**

- Administrative credentials: To complete this procedure, you must be a member of the Domain Admins group in the domain of the selected domain controller or a member of the Enterprise Admins group.
- Tool: Active Directory Sites and Services (Administrative Tools)

**To designate a preferred bridgehead server**

1. Open Active Directory Sites and Services.
2. In the console tree, double-click the **Sites** container, and then expand the **Servers** container.

3. Right-click the server object for the domain controller that you want to make a preferred bridgehead server, and then click **Properties**.

4. On the **General** tab, click the intersite transport or transports for which this server will be a preferred bridgehead server, and then click **Add**.

# Additional Resources for Troubleshooting Active Directory

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For specific information about troubleshooting Active Directory problems, see the following resources:

- Active Directory Management Pack Technical Reference for MOM 2005 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41369)

For information about Active Directory known issues and best practices, see the following resources:

- Known Issues for Creating Domain and Forest Trusts
- Best Practices for Domain and Forest Trusts

For general information about how Active Directory works and how to manage and configure Active Directory, see the following resources:

- Administering Active Directory Operations
- Active Directory Collection on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=34157)

# Monitoring and Troubleshooting Active Directory Replication Using Repadmin

Updated: March 12, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This document describes how to use the Repadmin.exe tool to monitor, diagnose, and troubleshoot common replication problems in your Active Directory® environment. All the information in this document applies to computers running the Microsoft® Windows® 2000 Server and Windows Server® 2003 operation systems. This document includes the following topics:

- Repadmin Introduction and Technology Overview
- Repadmin Requirements, Syntax, and Parameter Descriptions
- Repadmin Usage Scenarios
- Repadmin for Experts

To obtain a copy of this guide in .doc file format, see Troubleshooting replication with repadmin on the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkID=129020).

**Publication and revision history**
The following table summarizes the revision history for this guide, including its original publication on Microsoft TechNet.

| Date | Revision |
| --- | --- |
| June 2008 | Original publication on TechNet. |
| March 2010 | Updated with new commands for managing read-only domain controllers in Windows Server 2008 and Windows Server 2008 R2. For more information, see Administer Passwords and Password Replication Policy for Read-Only Domain Controllers with Repadmin.exe. |

# Repadmin Introduction and Technology Overview

Updated: September 17, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Repadmin.exe is a command line tool that is designed to assist administrators in diagnosing, monitoring, and troubleshooting Active Directory replication problems.

**Active Directory replication dependencies**

Active Directory replication has the following dependencies:

- **Routable IP infrastructure.** The replication topology depends on a routable IP infrastructure from which you can map IP subnet address ranges to site objects. This mapping generates the information that client workstations use to communicate with domain controllers that are close by—when there is a choice—rather than with domain controllers that are located across wide area network (WAN) links.

- **DNS.** The Domain Name System (DNS) that resolves DNS names to IP addresses. Active Directory requires that DNS is properly designed and deployed so that domain controllers can correctly resolve the DNS names of replication partners.

- **Remote procedure call (RPC).** Active Directory replication requires IP connectivity and the remote procedure call (RPC) to transfer updates between replication partners.

- **Kerberos version 5 (V5) authentication.** The authentication protocol for both authentication and encryption that is required for all Active Directory RPC replication.

- **Lightweight Directory Services Protocol (LDAP).** The primary access protocol for Active Directory. Replication of an entire replica of an Active Directory domain, as occurs when Active Directory is installed on an additional domain controller in an existing domain, uses LDAP communication rather than RPC.

- **NetLogon.** NetLogon dynamically registers the globally unique identifier (GUID) CNAME in DNS that a domain controller uses to resolve its partner's host name and IP address for Active Directory replication.

- **Intersite Messaging.** Intersite Messaging is required for Simple Mail Transfer Protocol (SMTP) intersite replication and for site coverage calculations. If the forest functional level is Windows 2000, Intersite Messaging is also required for intersite topology generation.




**Glossary of replication terms**

The following table lists terms that are commonly used in discussions about Active Directory replication.

| Term | Definition |
|---|---|

| | |
|---|---|
| Active Directory replication | Active Directory is a distributed directory service, in which not all objects in the directory are stored on every domain controller. In addition, all domain controllers in a domain can be updated directly, not just one primary domain controller. Active Directory replication is the means by which changes that are made on one domain controller are synchronized with all other appropriate domain controllers in the domain or forest that store copies of the same information. Data integrity is maintained by tracking changes on each domain controller and updating other domain controllers in a systematic way. Replication uses a connection topology that is created automatically to make optimal use of beneficial network connections. |
| Active Directory replication topology | Replication topology is the current set of Active Directory connections by which domain controllers in a forest communicate over local area networks (LANs) and WANs to synchronize the directory partition replicas that the domain controllers have in common. Replication topology generation is usually dynamic. It adapts to the network conditions and availability of domain controllers. As a result of how much we rely and depend on directory services today, it is very important to ensure that a directory replication topology is fine-tuned to maintain and deliver the expected level of performance. |
| Active Directory sites | A site is a part of the network with high bandwidth connectivity. By definition, it is a collection of well-connected computers, based on IP subnets. You can use the Active Directory Sites and Services snap-in to administer sites. Because sites control how replication occurs, changes that you make with this snap-in affect how efficiently domain controllers within a domain (but separated by great distances) will coalesce. |
| Knowledge Consistency Checker (KCC) | A part of the ISTG role in Active directory. The KCC checks and, as an option, re-creates topology information for the Active Directory domain. |
| Intersite Topology Generator (ISTG) | This is a role that one domain controller in an Active Directory site must perform.The ISTG designates one or more bridgehead servers to perform replication between sites. |
| Multimaster replication | Every domain controller can receive originating updates to data for which it is authoritative, rather than having a single domain controller that receives all original updates (also known as single-master replication, such as Microsoft Windows NT® 4.0 replication). |
| Pull replication | Domain controllers request (pull) changes rather than send (push) changes that might not be necessary. |
| Store-and-forward replication | Each domain controller communicates with a subset of domain controllers to transfer replication changes, rather than one domain controller being responsible for communicating with every other domain controller that requires the change. |
| High water mark | High water mark is a value that the destination domain controller maintains to keep track of the most recent changes that it has received from a specific source domain controller for an object in a specific partition. High water mark prevents irrelevant objects from being considered by the source domain controller with respect to a single destination. |
| Up-to-dateness vector | The up-to-dateness vector is a value that the destination domain controller maintains for tracking the originating updates that are received from all source domain controllers. When a destination domain controller requests changes for a directory partition, it provides its up-to-dateness vector to the source domain controller. The source domain controller then uses this value to reduce the set of attributes that it sends to the destination domain controller. |

**Glossary of other replication-related terms**

The following table lists terms that are related to other technologies that depend on Active Directory replication topology.

| Term | Definition |
|---|---|
| File Replication Service (FRS) | The replication service in Windows 2000 Server and Windows Server 2003 that is used to replicate the SYSVOL shared folder. |

| | |
|---|---|
| Replica set | The collection of servers that are all replicating a given set of directories is called a replica set. With an appropriate topology design and sufficient network support, a Windows 2000 or Windows Server 2003 FRS replica set can span thousands of computers. It is also possible for a single computer to be a member of multiple replica sets. |
| Topology | Topology defines the set of connections that are used to send updates between members of a replica set. The topology definition includes both the connections and the properties of those connections, such as the schedule, enabled and disabled flags, and so on. |
| Disconnected operation | FRS can operate even if some or all member computers are disconnected from each other for periods of time. Changes can be accepted by any computer, and changes are replicated to other member computers when connectivity is reestablished. |
| Authenticated RPC with encryption | To provide secure communications, FRS uses the Kerberos authentication protocol for authenticated RPC to encrypt and tamper-proof the data that is sent between replication partners. |

# Repadmin Requirements, Syntax, and Parameter Descriptions

Updated: March 12, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

You can use the **repadmin** command to perform replication tasks and to manage and modify the replication topology, force replication events, and display replication metadata and up-to-dateness vectors. This topic covers:


- System requirements
- File requirements
- Repadmin command-line options
- Repadmin subcommands
- Repadmin /listhelp
- CSV format


**System requirements**

The following are the system requirements for **repadmin**:

- Windows XP Professional, Windows Vista®, Windows Server 2003, or Windows Server 2008
- Administrator rights on the domain controller:
    - Required replication rights can be delegated
    - Some commands do not require Administrator rights

**File requirements**

Repadmin.exe is included in the Windows Server 2003 Service Pack 1 (SP1) Support tools. You must install the Support tools before you can use them. For more information about how to install the Support tools, see Windows Server 2003 SP1 Support Tools in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=44321).

To obtain the Support tools if you do not have the Windows Server 2003 operating system disc, see Windows Server 2003 SP1 32-bit Support Tools on the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkID=70775).

Previous versions of **repadmin** have similar functionality, but they have some limitations regarding the workstations that they can be run on and which functions they can perform. The following table lists the versions of **repadmin**, which operating systems they can be run on, and which domain controllers they can target.

| Version | Client operating system | Target operating system | Important feature sets |
|---|---|---|---|
| Windows 2000 | Windows 2000 and later | All Active Directory versions | /sync<br>/propcheck<br>/showreps<br>/showvector<br>/showmeta |

| | | | |
|---|---|---|---|
| Windows Server 2003 | Windows XP Professional and Windows Server 2003 | All Active Directory versions | /notifyopt<br>/replsummary<br>/replicate<br>/replsingleobj<br>/removelingeringobjects<br>/rehost and /unhost<br>/showmsg<br>/showattr<br>/syncall<br>/viewlist<br>*DC_LIST* |
| Windows Server 2003 with SP 1 | Windows XP Professional and Windows Server 2003 | All Active Directory versions<br>Rehost requires Windows 2000 Server SP4 and later<br>Remove lingering objects requires Windows Server 2003 | /showbackup<br>/rehost bug fix<br>/regkey |
| Active Directory Application Mode (ADAM) | Windows XP Professional and Windows Server 2003 | All Active Directory versions | /setattr<br>/listhelp |

| Deprecated subcommands (from Windows 2000 Server) | Equivalent or improved subcommands in Windows Server 2003 |
|---|---|
| /sync | /repl or /replicate |
| /propcheck | /showchanges |
| /showreps | /showrepl |
| /showvector | /showutdvec |
| /showmeta | /showobjmeta |

**Repadmin command-line options**

Repadmin is executed at the command prompt, and it contains several subcommands, which are described in detail in the following section.

**Syntax**

```
repadmin <subcommand> [<dsa>] [/u: <UserName>] [/pw: {<Password> | *}] [/rpc] [/ldap] [/homeserver:
<dsaname>]
```

**Parameters**

| Parameter | Description |
|---|---|
| <subcommand> | One of the **repadmin** subcommands that is described in the subcommands section. |
| <Dsa> | Directory System Agent (DSA) represents the domain controller to be targeted by the repadmin subcommand.<br>Not all **repadmin** subcommands require the *dsa* parameter<br>Type **repadmin /listhelp** at the command line for additional information about the *dsa* parameter. |

| | |
|---|---|
| /u:<UserName> | Specifies the account name to use for binding to the directory. By default, **/u** uses the account name with which the user is currently logged on. You can use any of the following formats to specify an account name:<br>• account name (for example, Bob)<br>• domain\account name (for example, contoso\Bob)<br>• user principal name (UPN) (for example, Bob@contoso.com) |
| /pw {<Password> \| *} | Specifies the password to use for authentication. If you type *****, you are prompted for a password. |
| /rpc | Forces **repadmin** to communicate by using a remote procedure call (RPC) session. |
| /ldap | Forces **repadmin** to communicate by using a Lightweight Directory Access Protocol (LDAP) session. If LDAP communication fails, **repadmin** attempts to communicate by using RPC. LDAP is the default communication method for **repadmin**. |
| /homeserver: <dsaname> | Forces **repadmin** to run against a specific domain controller, which is determined by the forest membership of the directory server that is represented by <*dsaname*>.<br>You can specify <*dsaname*> in the following formats:<br><*Computername*>, <*Dnsname*>, <*Dsaguid*>, *, ., "site:<*site*>", "fsmo_dnm:", or "fsmo_schema:". |

**Repadmin subcommands**

| Subcommand | Syntax and description |
|---|---|
| bind | repadmin /bind [*dsa*]<br>Connects to and displays the replication features for a directory server. |
| bridgeheads | repadmin /bridgeheads [*dsa*]<br>Lists the directory servers that act as bridgehead servers for a specified site. |
| checkprop | repadmin /checkprop [*dsa*] *Naming ContextOriginatingDCInvocationIDOriginatingUSN*<br>Compares the properties of specified directory servers to determine if they are up to date with each other. The source directory server contains the original information that must be checked. The data on the destination directory server is compared to the data on the source directory server. |
| dsaguid | repadmin /dsaguid [*dsa*] [*GUID*]<br>Returns a server name when given a globally unique identifier (GUID). |
| failcache | repadmin /failcache [*dsa*]<br>Displays a list of failed replication links that are detected by the Knowledge Consistency Checker (KCC). |
| istg | repadmin /istg [*dsa*] [/verbose]<br>Returns the computer name of the Intersite Topology Generator (ISTG) server for a specified site. |
| kcc | repadmin /kcc [*dsa*] [/async]<br>Forces KCC to calculate replication topology for a specified directory server. By default, this calculation occurs every 15 minutes. |

| | |
|---|---|
| latency | repadmin /latency [*dsa*] [/verbose]<br><br>Displays the amount of time between replications, by using the ISTG Keep Alive time stamp. The ISTG Keep Alive time stamp is not used in forests that are set to the Windows Server 2003 forest functional level. Instead, in those environments, use **repadmin /showutdvec /latency**. |
| notifyopt | repadmin /notifyopt [*dsa*] *Naming Context* [/first:*Value*] [/subs:*Value*]<br><br>Displays or sets the notification timing settings for replication of a specified directory partition. |
| queue | repadmin /queue [*dsa*]<br><br>Displays tasks that are waiting in the replication queue. |
| prp | Repadmin /prp [*operation*] *RODC* [*additional arguments*]<br><br>Displays or modified the Password Replication Policy for a read-only domain controller (RODC).<br><br>This command is available only for versions of Repadmin that are included in Windows Server 2008, Windows Server 2008 R2, or Remote Server Administration Tools.<br><br>The *operation* can be **view**, **add**, **delete**, or **move**. For **view**, **add**, and **delete**, *RODC* can be either *RODC_Name* or *****. For **move**, *RODC* must be *RODC_name*. |
| querysites | repadmin /querysites *FromSiteRDNToSite1RDN* [*ToSite2RDN*...]<br><br>Uses routing information to determine the cost of a route from a specified site to another specified site or sites. The **querysites** parameter does not allow the use of alternate credentials. The relative distinguished names that are used in this command are case sensitive. |
| replicate | Syntax 1<br><br>repadmin /replicate *destination_dsasource_dsa* [/force] [/async] [/full] [/addref]<br><br>Syntax 2<br><br>repadmin /replicate *destination_dsa* [/force] [/async] [/full] [/addref] /allsources<br><br>Starts a replication event for the specified directory partition between the source and destination directory servers. You can determine the source GUID when you view the replication partners by using **showrepl**. |
| replsingleobj | repadmin /replsingleobject *dsaDsaSourceGUIDObjectDN*<br><br>Replicates a single object between any two directory servers that have partitions in common. The two directory servers do not have a replication agreement. You can show replication agreements by using the **repadmin /showrepl** command. |
| replsummary | repadmin /replsummary [*dsa*] [/bysrc] [/bydest] [/errorsonly][/sort:{delta\|partners\|failures\|error\|percent}]<br><br>Summarizes the replication state and relative health of an Active Directory forest. |
| rodcpwdrepl | repadmin /rodcpwdrepl [*DSA_list*] *Hub DCUser1 DN* [*User2 DNUser3 DN*]<br><br>Triggers replication of passwords for the specified users from the source *Hub DC* to one or more RODCs.<br><br>This command is available only for versions of Repadmin that are included in Windows Server 2008, Windows Server 2008 R2, or Remote Server Administration Tools. |
| showattr | repadmin /showattr *dsa* [*OBJ_LIST*] [*OBJ_LIST_OPTIONS*] [/attr\|/attrs: *attributeattribute* ...] [/allvalues] [/long] [/nolongblob] [/nolongblob] [/nolongfriendly] [/dumpallblob]<br><br>The **/showattr** operation displays the attributes and contents of an object. |
| showcert | repadmin /showcert *dsa*<br><br>Displays the certificates (used with Simple Mail Transfer Protocol (SMTP)–based replication) that are loaded on a specified directory server. |

| | |
|---|---|
| showchanges | Syntax 1<br><br>repadmin /showchanges *source_dsaNaming Context* [/cookie: *File*] [/atts: *attribute1*,*attribute2*,...]<br><br>Syntax 2<br><br>repadmin /showchanges *dest_dsaSourcedsaObjectGUIDNaming Context* [/verbose] [/statistics] [/noincremental] [/objectsecurity] [/ancestors] [/atts: *attribute1*,*attribute2*,...] [/filter: *ldap filter*]<br><br>Displays changes from a specified directory partition or changes to a specified object. "Syntax 1" saves changes to a directory partition. If this information is saved to a file, you can run the **getchanges** operation again for comparison. "Syntax 2" lists changes to a specified object. For this command to run properly, the account under which the command is run must possess the **replication get changes** right on the specified directory partition. |
| showconn | repadmin /showconn [*dsa*] [*ServerRDN* \| *ContainerDN* \| *dsa_GUID*] [/From:*ServerRDN*] [/intersite]<br><br>Displays the connection objects for a specified directory server. The default is local site. |
| showctx | repadmin /showctx [*dsa*] [/nocache]<br><br>Displays a list of computers that have opened sessions with a specified directory server. |
| showism | repadmin /showism [*TransportDN*] [/verbose]<br><br>Queries the Intersite Messaging Service (ISM) for site routes. This operation cannot be executed remotely. |
| showmsg | repadmin /showmsg {*Win32Error* \| *DSEventID* \| *NTDSMSG*}<br><br>Displays the error message for a given error number. |
| showncsig | repadmin /showncsig [*dsa*]<br><br>Each directory server maintains a directory partition signature list. This command displays a list of the removed application partition GUIDs. You can configure an application directory partition to be held or not held on a particular directory server by using **ntdsutil** (for Active Directory). |
| showobjmeta | repadmin /showobjmeta [*dsa*] *ObjectDN* [/nocache] [/linked]<br><br>Displays the replication metadata for a specified object that is stored in the directory, including attribute ID, version number, originating and local update sequence number (USN), and originating server's GUID and Date and Time stamp. When you compare the replication metadata for the same object on different directory servers, you can determine whether replication has occurred. |
| showoutcalls | repadmin /showoutcalls [*dsa*]<br><br>Displays calls that have been made by the specified directory server to other directory servers but not yet answered. |
| showproxy | Syntax 1<br><br>repadmin /showproxy [*dsa*] [*Naming Context*] [matchstring]<br><br>Syntax 2<br><br>repadmin /showproxy [*dsa*] [*ObjectDN*] [matchstring] /movedobject<br><br>Lists cross-domain move proxy objects. When an object is moved from one domain to another, a marker remains in the original domain. This marker is called a proxy. |
| showrepl | repadmin /showrepl [*dsa*] [*SourceDCObjectGUID*] [*Naming Context*] [/verbose] [/nocache] [/repsto] [/conn] [/csv] [/all] [/errorsonly] [/intersite]<br><br>Displays replication information. Inbound replica links are displayed by default. Outbound links can also be shown, as well as connections corresponding to those links. The command also displays errors that correspond to replica links that cannot be created by KCC. This helps an administrator build a visual representation of the replication topology and see the role of each directory server in the replication process. |

| | |
|---|---|
| showcig | repadmin /showsig [*dsa*]<br><br>Displays the retired invocation IDs on a directory server. A directory server changes its invocation ID when it is restored or when it rehosts an application partition. |
| showtime | repadmin /showtime [*DSTimeValue*]<br><br>Converts a directory service time value to string format for both the local and the UTC time zones. |
| showtrust | repadmin /showtrust [*dsa*]<br><br>Lists all Active Directory domains that are trusted by a specified Active Directory domain. |
| showutdvec | repadmin /showutdvec *dsaNaming Context* [/nocache] [/latency]<br><br>Displays the highest USN for the specified directory server. This information shows how up to date a replica is with its replication partners. |
| showvalue | repadmin /showvalue [*dsa*] *ObjectDN* [*AttributeName*] [*ValueDN*] [/nocache]<br><br>Displays the values of the type, last modified time, originating directory server, and distinguished name of a specified object. |
| syncall | repadmin /syncall *dsa* [*Naming Context*] [*Flags*]<br><br>Synchronizes a specified directory server with all replication partners. This command contains several subcommands, which are described in the usage scenarios.<br><br>By default, if no directory partition is provided in the *NamingContext* parameter, the command performs its operations on the configuration directory partition. |
| viewlist | repadmin /viewlist [*dsa*] [*OBJ_LIST*]<br><br>Displays a list of directory servers. |
| oldhelp | Displays a list of the operations that have been deprecated in this version of **repadmin**. |

**Repadmin /listhelp**

| Arguments | Values | Description |
|---|---|---|
| DC_LIST | "*" | All domain controllers in the enterprise |
| | DC_Name | See under DC_NAME argument |
| | Part-server_name* | Would pick "part_server_name_dc_01" and "part_server_name_dc_02" but not server "part_server_diff_name". |
| | Site:*site_name* | All domain controllers in the specified site. |
| | Gc: | All global catalog servers in the enterprise. |
| | Fsmo_*fsmo_type*:*fsmo_dn* | See under FSMO_TYPE |
| FSMO_TYPE | | Types of operations master (also known as flexible single master operations or FSMO) role holders require different base distinguished names or relative distinguished names. |
| | Fsmo_dnm: | Enterprise-wide FSMO; does not take any distinguished name (also known as DN). |
| | Fsmo_schema: | Enterprise-wide FSMO; does not take any distinguished name. |
| | Fsmo_pdc: | Domain-specific FSMO; takes the distinguished name of the domain that the user specifies. |

| | | |
|---|---|---|
| | Fsmo_rid: | Domain-specific FSMO; takes the distinguished name of the domain that the user specifies. |
| | Fsmo_im: | Domain-specific FSMO; takes the distinguished name of the domain that the user specifies. |
| | Fsmo_istg: | Site-specific quasi-FSMO; takes the relative distinguished name of the site. |
| DC_NAME | | |
| | "." | Tells **repadmin** to try to pick a domain controller for you. |
| | Server_dns | Specifies a server by DNS. |
| | Dc_dsa_guid | Specifies a specific server by its Directory System Agent (DSA) GUID. |
| | Server_obj_rdn | Specifies a server by its server object relative distinguished name (usually the same as its NetBios name). |
| | Dsa_dn | Specifies a server by the distinguished name of its DSA object. |
| OBJ_LIST | | |
| | Ncobj:NC_NAME | Specifies the use of the distinguished name of NC Head that is specified in NC_NAME. |
| | Dsaobj: | Specifies the use of the distinguished name of the DSA that **repadmin** is connected to. |
| NC_NAME | Config: | Configuration directory partition. |
| | Schema: | Schema directory partition. |
| | Domain: | Domain directory partition for the domain of the domain controller that **repadmin** is running against. |
| OBJ_LIST OPTIONS | {/onelevel \| /subtree} /filter:{*ldap_filter*} | With these options, you can use the **showattr** and **viewlist** commands to cover a list of objects, instead of just a single object. |

**CSV format**

The output that **repadmin /showrepl** returns can be difficult to navigate when you are troubleshooting replication errors or viewing replication topology in a large enterprise. There is a new feature (**/CSV**) that you can use to force **/showrepl** output to print in a tightly constrained comma-separated-value (CSV) format for programmatic manipulation or quick import and correlation in Excel.

The CSV format is also an effective way to exchange **repadmin** outputs because it is not prone to user errors.

To generate output as a .csv (comma-delimited) file, perform the following steps:

1. Open a command prompt, type the following command, and then press ENTER:

   ```
   repadmin /showrepl <DC_NAME> /csv > Repl.csv
   ```

2. Open Repl.csv, and then delete or hide column **A** and both **RPC** and **SMTP** columns.

3. Select row **2**. Click **View**, and then click **Freeze Panes**.

4. Highlight the column heading row. Click **Data**, point to **Filter**, and then click **AutoFilter**.

5. Click the drop-down arrow to display replication status based on your situation.

REPLICATION STATUS.csv

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | Destination Site | Destination DC | Naming Context | Source DC Site | Source DC | Total Failures | Last Failure | Last Success | Last Failure |
| 2 | HUB | BRANCH-HUB-BH | CN=Configuration,DC=Contoso,DC=c | BRANCH1 | BRANCH1 | 0 | 0 | 9/18/2005 1:58 | 0 |
| 3 | HUB | BRANCH-HUB-BH | CN=Schema,CN=Configuration,DC=C | BRANCH1 | BRANCH1 | 0 | 0 | 9/18/2005 1:58 | 0 |
| 4 | HUB | BRANCH-HUB-BH | DC=Child,DC=Contoso,DC=com | BRANCH1 | BRANCH1 | 0 | 0 | 9/18/2005 1:58 | 0 |
| 5 | HUB | BRANCH-HUB-BH | CN=Configuration,DC=Contoso,DC=c | BRANCH2 | BRANCH2 | 0 | 0 | 9/18/2005 1:58 | 1722 |
| 6 | HUB | BRANCH-HUB-BH | CN=Schema,CN=Configuration,DC=C | BRANCH2 | BRANCH2 | 0 | 0 | 9/18/2005 1:58 | 1722 |
| 7 | HUB | BRANCH-HUB-BH | DC=Child,DC=Contoso,DC=com | BRANCH2 | BRANCH2 | 0 | 0 | 9/18/2005 1:58 | 1722 |

REPLICATION STATUS

# Repadmin Usage Scenarios

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This section includes explanations and examples for the following usage scenarios:

- Monitor Forest-Wide Replication
- Display Replication Partners and Status of a Domain Controller
- Replication Latency
- View Replication Metadata of an Object
- Display the Attributes of a Specific Object
- How Up to Date Are My Domain Controllers?
- Can I Look at My Connection Objects and Schedule Details?
- Fine-Tuning Change Notification Values
- Forcing Replication
- Keeping Track of Changes That Have Occurred Over a Period of Time
- Usage of Repadmin When Troubleshooting Event ID 1311
- Subcommands Not Covered Under the Previous Scenarios
- Oldhelp

# Monitor Forest-Wide Replication

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Maintaining the health of enterprise-wide directory replication is very important so that the users, services, machines, and applications that rely on it can operate successfully. The Windows Server 2003 version of repadmin has enhanced functionality that makes it easier to monitor forest-wide directory replication and it is compatible with Windows 2000 domains.

**Repadmin /replsummary** summarizes the replication state and relative health of an Active Directory forest by inventorying and contacting every domain controller in the forest, collecting information such as replication deltas and replication failures.

It will also identify any domain controllers that could not be contacted and would report the failure reason (for an example, see Figure 3.1.4).

**Syntax**

```
Repadmin /replsummary <DC_LIST> [/bysrc] [/bydest] [/errorsonly] [/sort:{delta | partners | failures | error
| percent}]
```

| Parameters | Definition |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. For details about <*DC_LIST*>, see **repadmin /listhelp**. |
| **/bysrc** | Shows the output of **repadmin /replsummary**, from the perspective of the replication source (outbound domain controller), in the form of a table. This means that a given source directory server is "pulled on" by multiple client domain controllers. The table is sorted in order of the source domain controllers that are having the most problems, across all the clients in the configuration set. This parameter does not display the destination domain controller. |
| **/bydest** | Shows the output of **repadmin /replsummary**, from the perspective of the replication destination, in the form of a table. This means that a given replication destination (inbound domain controller) is pulling the changes from one or more replication source(s). The table shows the inbound domain controllers and what problems they are having with their partners. The table is sorted in order of the inbound domain controllers that are having the most problem with inbound replication, across all the possible partners in the configuration set. This parameter does not display the source domain controller. |
| **/errorsonly** | Shows only the domain controllers where the partner error is not zero. |
| **/sort:{delta \| partners \| failures \| error \| percent}** | Sorts the replsummary table by the specified column heading. |

**Note**

The **/bysrc** and **/bydest** parameters may be specified at the same time. If they are specified at the same time, repadmin displays the **/bysrc** table first and the **/bydest** table next. If the parameters **/bysrc** and **/bydest** are both absent, repadmin picks the best one and displays the one with the least number of partner errors.

**Simple usage of repadmin /replsummary**

```
C:\>repadmin /replsummary
Replication Summary Start Time: 2005-01-04 22:36:30

Beginning data collection for replication summary, this may take awhile:
  .........


Source DC           largest delta  fails/total  %%  error
 BRANCH-HUB-BH             46m:46s    0 /  18    0
 BRANCH1                   16m:44s    0 /   3    0
 BRANCH2                   16m:44s    0 /   3    0
 BRANCH3                   16m:43s    0 /   5    0
 ROOTDC01                  49m:21s    0 /   6    0
 ROOTDNS                   49m:21s    0 /   7    0


Destination DC    largest delta    fails/total  %%  error
 BRANCH-HUB-BH             48m:21s    0 /  16    0
 BRANCH1                   09m:24s    0 /   3    0
 BRANCH2                   11m:56s    0 /   3    0
 BRANCH3                   02m:30s    0 /   5    0
 ROOTDC01                  45m:47s    0 /   7    0
 ROOTDNS                   42m:46s    0 /   8    0
```

**How to interpret the output**

The output of **repadmin /replsummary** is organized by destination and source domain controllers. You should focus on the destination domain controllers first, because the replication model is pull-based. Replication between domain controllers does not use a "push" mechanism. If the replication is within a site, a domain controller (DC1) notifies another domain controller (DC2) that it has updates, and then the DC2 pulls the updates from DC1. If the replication is between sites, a domain controller requests updates at a scheduled time and if updates are available, the domain controller pulls the updates from a domain controller in the other site.

| Fields of interest | Definition |
|---|---|
| ….. | Each dot after the first three represents a domain controller, with not more than 50 dots per line. So, if you have two lines full of dots, it indicates 97 domain controllers (100-3).<br><br>In figure 3.1.1, there are nine dots, which relates to six domain controllers (9-3). |
| Destination DC | Replication destination. A single destination might be pulling data from one or more sources.<br><br>In figure 3.1.1, we are focusing on ROOTDC01. |
| Source DC | Replication source. Multiple destinations might be pulling from a single source.<br><br>In figure 3.1.1, we do not yet know the source domain controller. |
| Largest delta | Denotes the longest replication gap amongst all replication links for a particular domain controller.<br><br>In figure 3.1.1, the largest delta is 45m:47s. |
| Total | Replica links for a particular domain controller (one for each naming context on each domain controller). Please note that this is not the connection objects or replication partners per domain controller.<br><br>In figure 3.1.1, we have seven replication links. |
| Fails | Total number of replica links failing to replicate for one reason or the other. This will never be greater than the Total field.<br><br>No failures in our example. |
| Percentage | Percentage of failures in relation to the total replica links on the domain controller. |

**How to make more sense of some of the fields**

We ran **repadmin /showrepl** against ROOTDC01 to get detailed replication status. Always focus on inbound neighbors because replication is inbound.

If you notice Figure 3.1.1, the time replsummary taken was 22:36:30. Now, if you look at the schema naming context replication time, 21:49:44 from figure 3.1.2, the difference is 45m:47s, which relates to the largest delta.

- Interestingly, 45 minutes is relatively high in our example because our partners belong to the same site. This is because the default periodic replication frequency is once per hour within a site and because the schema naming context did not have any changes, periodic replication took place only at 21:49:44 as opposed to other partitions that replicated in response to change notifications from its partners.

- We also see seven replica links, one for each naming context on each domain controller.

```
C:\>repadmin /showreps rootdc01
HUB\ROOTDC01
DC Options: IS_GC
Site Options: (none)
DC object GUID: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c
DC invocationID: 195f73bb-5c65-4d14-a1fd-39d0adec3d0c

==== INBOUND NEIGHBORS ======================================

DC=contoso,DC=com
    HUB\ROOTDNS via RPC
        DC object GUID: 076cd5dd-e25e-4897-acd2-7c8691621522
        Last attempt @ 2005-01-04 22:35:15 was successful.

CN=Configuration,DC=contoso,DC=com
    HUB\ROOTDNS via RPC
        DC object GUID: 076cd5dd-e25e-4897-acd2-7c8691621522
        Last attempt @ 2005-01-04 22:31:21 was successful.
    HUB\BRANCH-HUB-BH via RPC
        DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53
        Last attempt @ 2005-01-04 22:30:27 was successful.

CN=Schema,CN=Configuration,DC=contoso,DC=com
    HUB\BRANCH-HUB-BH via RPC
        DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53
        Last attempt @ 2005-01-04 21:49:44 was successful.
    HUB\ROOTDNS via RPC
        DC object GUID: 076cd5dd-e25e-4897-acd2-7c8691621522
        Last attempt @ 2005-01-04 21:49:44 was successful.

DC=research,DC=contoso,DC=com
    HUB\BRANCH-HUB-BH via RPC
        DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53
        Last attempt @ 2005-01-04 22:18:56 was successful.
    HUB\ROOTDNS via RPC
        DC object GUID: 076cd5dd-e25e-4897-acd2-7c8691621522
        Last attempt @ 2005-01-04 22:19:17 was successful.

C:\>_
```

**Common factors that influence the largest delta field**

- Periodic intrasite replication frequency.
- Intersite replication schedule and frequency.
- Redundant replication paths with staggered replication schedules.
- Intrasite and intersite change notifications; first and subsequent replication notification delay values.

**Where do REPADMIN /REPLSUMMARY read replication status information?**

- Similar to **/showrepl**, **Repadmin /replsummary** gathers this information from the **Reps-from** and **Repsto multivalued** attributes stored at the root of each directory partition replica (also known as naming contexts) stored on the domain controller. It is local to the domain controller and not replicated.
- The **Repsfrom** attribute contains configuration and persistent state information associated with inbound replication from each source replica of that directory partition.
- The **Repsto** attribute contains outbound change notification partners. Typically this list would be your intrasite partners.

**Wild card and other parameter usage**

The following example uses a wildcard character to show the replication summary for all of the domain controllers in the forest that have a name that begins with 'ROOT'.

```
C:\>repadmin /replsummary root* /bysrc /bydest /sort:delta
Replication Summary Start Time: 2005-01-04 22:50:10

Beginning data collection for replication summary, this may take awhile:
  .....


Source DC           largest delta  fails/total  %%  error
 BRANCH-HUB-BH             56m:26s    0 /   7    0
 ROOTDC01                  13m:41s    0 /   4    0
 ROOTDNS                   01m:36s    0 /   4    0


Destination DC    largest delta    fails/total  %%  error
 ROOTDNS                   56m:26s    0 /   8    0
 ROOTDC01                     :36s    0 /   7    0


C:\>_
```

**Important**

- If there are no inbound partners for a given domain controller, none would be listed under the Destination DC list. Similarly, if there are no outbound partners for a given domain controller, none would be listed under the Source DC list.

- So it is important to tally the total number of domain controllers in the forest and compare that against the Destination DC and Source DC lists to achieve an accurate view. **Repadmin /viewlist *** should list all the domain controllers in the forest.

- At this time of writing, the total number of replication links that would be reported in the replsummary output is limited to 999.

- If the replication destination has never replicated from the source, the largest delta would report as **unknown**.

**Replsummary reporting failures**

The following example reports replication failure and a domain controller that could not be reached, with the error codes and reasons.

```
C:\>repadmin /replsummary /bysrc /bydest /sort:delta /homeserver:rootdns
Replication Summary Start Time: 2005-01-05 00:04:54

Beginning data collection for replication summary, this may take awhile:
  .........


Source DC           largest delta  fails/total  %%  error
 BRANCH2               01h:25m:48s    3 /   3  100  (1722) The RPC server is unavailable.
 ROOTDNS                   56m:34s    0 /   7    0
 ROOTDC01                  56m:34s    0 /   6    0
 BRANCH1                   23m:05s    0 /   3    0
 BRANCH3                   22m:24s    0 /   5    0
 BRANCH-HUB-BH             18m:01s    0 /  15    0


Destination DC    largest delta    fails/total  %%  error
 BRANCH-HUB-BH         01h:24m:33s    3 /  16   18  (1722) The RPC server is unavailable.
 BRANCH1                   16m:55s    0 /   3    0
 ROOTDC01                  15m:20s    0 /   7    0
 ROOTDNS                   11m:20s    0 /   8    0
 BRANCH3                      :42s    0 /   5    0


Experienced the following operational errors trying to retrieve replication information:
          58 - BRANCH2.research.contoso.com

C:\>_
```

```
C:\>net helpmsg 58

The specified server cannot perform the returned operation.

C:\>net helpmsg 1722

The RPC server is unavailable.
```

In our example, the following occurred:

- We could not reach BRANCH2 and hence the error 58.

- "RPC server is unavailable" being reported by BRANCH-HUB-BH co-relates to the above finding. It could mean that BRANCH2 domain controller is either down or not reachable due to communication link problem.

- We also used **/homeserver:rootdns** to demonstrate that sometimes you have to specify a server (/homeserver: <domain controller name>) if you are not running the command from a domain controller.

# Display Replication Partners and Status of a Domain Controller

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

When troubleshooting replication errors, it is helpful to know who are the replication partners of a specific domain controller and the status of replication with each of those partners.

**Repadmin /showrepl** displays the replication partners (**RepsFrom** and **RepsTo**) for each naming context that is held on the specified domain controller. By enumerating each **RepsFrom** and each **RepsTo** for each domain controller, you can visualize the replication topology for each naming context.

It also indicates whether the domain controller is also a global catalog server. Inbound replica links are displayed by default. Outbound links can also be shown, as well as connections that correspond to those links. The command also displays errors that correspond to replica links that cannot be created by the Knowledge Consistency Checker (KCC). This helps the administrator build a visual representation of the replication topology and see the role of each directory server in the replication process.

**Syntax**

```
Repadmin /showrepl <DC_LIST> <SourceDCObjectGUID> [NamingContext] [/verbose] [/nocache] [/repsto] [/conn]
[/csv] [/all] [/errorsonly] [/intersite]
```

| Parameters | Definition |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax. For details about <*DC_LIST*>, see **repadmin /listhelp**. |
| *SourceDCObjectGUID* | Specifies the unique hexadecimal number that identifies the object whose replication events will be listed. |
| **NamingContext** | Specifies the distinguished name of the directory partition. |
| **/verbose** | Lists detailed information. |
| **/nocache** | Specifies that globally unique identifiers (GUIDs) are left in hexadecimal form. By default, GUIDs are translated into strings. |
| **/repsto** | Lists the directory servers that pull replication information from the specified directory partition. To see the outbound neighbors, specify **/repsto** or **/all**. |
| **/conn** | Displays the connection objects that are associated with each link. |
| **/csv** | Displays the output of the **repadmin /showrepl** operation in a Comma Separated Variable (CSV) format for viewing and analysis in Microsoft Excel. Repadmin supports redirection of screen output to a file. |
| **/all** | Displays all replication partners. |
| **/errorsonly** | Only shows the partnership if it has an error associated with it. |
| **/intersite** | Only shows this partnership if the source server belongs to a different site than the site of the server on which the command is being run. |

**Show replication partners and replication status**

The following example uses the **showrepl** operation of repadmin to display the replication status of ROOTDNS in relation to its partners. In our example, there are no problems reported because replication is running properly. There is lot of information one could gather from this output and please read the comments next to each line explaining what it means.

Figure 3.2.1

```
C:\>repadmin /showrepl rootdns
```

```
HUB\ROOTDNS (Site name and domain controller name)

DC Options: IS_GC (DC Options)

Site Options: (none) (Site options)

DC object GUID: 076cd5dd-e25e-4897-acd2-7c8691621522 (GUID of NTDS settings)

DC invocationID: 076cd5dd-e25e-4897-acd2-7c8691621522 (Database signature)

==== INBOUND NEIGHBORS ======================================================

DC=contoso,DC=com (Naming Context)

HUB\ROOTDC01 via RPC (Replication link)

DC object GUID: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c (GUID of replication partner)

Last attempt @ 2005-01-05 01:04:34 was successful. (Status of last replication)

CN=Configuration,DC=contoso,DC=com (Naming Context)

HUB\ROOTDC01 via RPC (Replication link)

DC object GUID: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c (GUID of replication partner)

Last attempt @ 2005-01-05 01:01:31 was successful. (Status of last replication)

HUB\BRANCH-HUB-BH via RPC (Replication link)

DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53

Last attempt @ 2005-01-05 01:01:44 was successful.

CN=Schema,CN=Configuration,DC=contoso,DC=com (Naming Context)

HUB\BRANCH-HUB-BH via RPC (Replication link)

DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53

Last attempt @ 2005-01-05 00:53:34 was successful.

HUB\ROOTDC01 via RPC (Replication link)

DC object GUID: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c

Last attempt @ 2005-01-05 00:53:34 was successful.
```

In the output under INBOUND NEIGHBORS, repadmin.exe shows the Lightweight Directory Access Protocol (LDAP) distinguished name of each directory partition for which inbound directory replication has been attempted, the site and name of the source domain controller, and whether it succeeded or not, as follows:

- Last attempt @ YYYY-MM-DD HH:MM.SS was successful.
- Last attempt @ [Never] was successful.

If repadmin.exe reports any of the following conditions, further investigation is required:

- The last successful inter-site replication was prior to the last scheduled replication.
- The last intra-site replication was longer than one hour ago.
- Replication was never successful.

DC Object GUID is a reference point used in the Active Directory and Domain Name System (DNS) to locate a domain controller primarily for the purposes of replication. This GUID is automatically generated for each domain controller, is unique when created, and will not be duplicated.

DC invocationID – Active Directory database has its own GUID, which the Directory System Agent (DSA) uses to identify the

database instance (version of the database). The database GUID is stored in the **invocationId** attribute on the nTDSDSA object. Although the DSA GUID never changes for the lifetime of the domain controller, the Active Directory database GUID (also known as the invocation ID or database signature) is changed during the Active Directory restore process to ensure the consistency of the replication process. In Windows Server 2003, it changes when application directory partitions are removed or added to the domain controller.

**Using repadmin /showrepl to display detailed and precise information**

The following **showrepl** output is returned by combining <*Naming Context*> and **/verbose**.

```
C:\>repadmin /showrepl rootdns DC=contoso,DC=com /verbose
HUB\ROOTDNS
DC Options: IS_GC
Site Options: (none)
DC object GUID: 076cd5dd-e25e-4897-acd2-7c8691621522
DC invocationID: 076cd5dd-e25e-4897-acd2-7c8691621522

==== INBOUND NEIGHBORS ======================================

DC=contoso,DC=com
    HUB\ROOTDC01 via RPC
        DC object GUID: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c
        Address: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c._msdcs.contoso.com
        DC invocationID: 195f73bb-5c65-4d14-a1fd-39d0adec3d0c
        SYNC_ON_STARTUP DO_SCHEDULED_SYNCS WRITEABLE
        USNs: 941080/OU, 941080/PU
        Last attempt @ 2005-01-05 01:04:34 was successful.

C:\>_
```

For two domain controllers to engage in replication, they have to first resolve each other's GUID CNAME to a host name and the host name to an IP address, such as the following:

```
C:\>ping 2a92f776-6c0f-4cb4-a111-f5dcd447af6c._msdcs.contoso.com

Pinging rootdc01.contoso.com [192.168.0.101] with 32 bytes of data:

Reply from 192.168.0.101: bytes=32 time<1ms TTL=128
```

**2a92f776-6c0f-4cb4-a111-f5dcd447af6c._msdcs.contoso.com** is the GUID **CNAME** registration in DNS.




**High-watermark value**

The high-watermark value is not required for any administrative task. However, it can help you deduce the state of progress on that replication link. You can see the high-watermark in the output of the **repadmin /showrepl /verbose** command in Figure 3.2.2. Look for lines that begin with USNs:. The high-watermark USN is the number that is followed by **/OU**.

The object update (OU) USN saves the position when in the middle of a replication cycle. It stays the same as the property update (PU) when replication is not occurring, and increases during a replication cycle. At the end of the cycle, the final USN replicated becomes the PU value and the OU is left to match. Thus, the OU indicates progress within a cycle, and the PU indicates the last update seen at the conclusion of a successful cycle. A PU of zero means that the link has never completed a successful cycle, as is the case when performing its first synchronization on a new domain controller connection. If the OU and PU are not equal, it means a replication cycle is in progress.

The following table lists **nbrflagoptions**, which are flags that define expected replication actions with its partner.

| Nbrflagoptions | Meaning |
|---|---|
| WRITEABLE | The local copy of the naming context is writable. |
| SYNC_ON_STARTUP | Replication of this naming context from this source is attempted when the destination server is booted. This normally only applies to intrasite neighbors. |

| | |
|---|---|
| DO_SCHEDULED_SYNCS | Perform replication on a schedule. This flag is normally set unless the schedule for this naming context/source is "never," that is, the empty schedule. |

**Showing outbound neighbors**

By default, **repadmin /showrepl** does not display outbound neighbors, as with previous versions. The **/repsto** parameter provides this feature, as shown in Figure 3.2.5.

```
C:\>repadmin /showrepl rootdns CN=Configuration,DC=contoso,DC=com  /repsto
HUB\ROOTDNS
DC Options: IS_GC
Site Options: (none)
DC object GUID: 076cd5dd-e25e-4897-acd2-7c8691621522
DC invocationID: 076cd5dd-e25e-4897-acd2-7c8691621522

==== INBOUND NEIGHBORS ======================================

CN=Configuration,DC=contoso,DC=com
    HUB\ROOTDC01 via RPC
        DC object GUID: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c
        Last attempt @ 2005-01-05 01:01:31 was successful.
    HUB\BRANCH-HUB-BH via RPC
        DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53
        Last attempt @ 2005-01-05 01:01:44 was successful.

==== OUTBOUND NEIGHBORS FOR CHANGE NOTIFICATIONS ============

CN=Configuration,DC=contoso,DC=com
    HUB\ROOTDC01 via RPC
        DC object GUID: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c
        Last attempt @ 2005-01-05 00:16:12 was successful.
    HUB\BRANCH-HUB-BH via RPC
        DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53
        Last attempt @ 2005-01-05 00:50:35 was successful.

C:\>_
```

**Some of the repadmin /showrepl Error Messages and their root cause**

The following table lists some **repadmin /showrepl** errors and their root cause. The next sections after the table explain some errors in more detail.

| Repadmin error | Root cause |
|---|---|
| No inbound neighbors | If no items appear in the "Inbound Neighbors" section of the output generated by the **repadmin /showrepl** command, the domain controller could not establish replication links with another domain controller. |
| Access denied | A replication link exists between two domain controllers, but replication cannot be properly performed. |
| Last attempt at <date - time> failed with the "Target account name is incorrect." | This problem can be related to connectivity, DNS, or authentication issues.<br><br>If it is a DNS error, the local domain controller could not resolve the GUID–based DNS name of its replication partner. |
| No more end point | This can be caused because no more end-points are available to establish the TCP session with the replication partner.<br><br>This error can also result when the replication partner can be contacted, but its RPC interface is not registered. This usually indicates that the domain controller's DNS name is registered but with the wrong IP address. |
| LDAP Error 49 | The domain controller computer account might not be synchronized with the Key Distribution Center (KDC). |
| Cannot open LDAP connection to local host. | The administration tool could not contact Active Directory. |
| Active Directory replication has been preempted | An inbound replication in progress was interrupted by a higher priority replication request, such as a request generated manually by using the **repadmin /syncall** command. |
| Replication posted, waiting. | The domain controller posted a replication request and is waiting for an answer. Replication is in progress from this source. |

| | |
|---|---|
| Last attempt @ never was successful | The KCC successfully created the replication link between the local domain controller and its replication partner, but because of the schedule or possible bridgehead overload, replication has not occurred.<br><br>A large backlog of inbound replication must be performed on this domain controller. |

**No inbound neighbors**

A "no inbound neighbor" error appears in the **repadmin /showrepl** command output when one or more of the following conditions exists:

- No connection object exists to indicate which domain controller(s) this domain controller should replicate from. These connection objects are typically created by the KCC. However, in some environments, administrators have turned off the part of KCC (Intersite) that creates connection objects for inbound replication from domain controllers in other sites, relying on manual connections instead.

- One or more connection objects exist, but the domain controller cannot contact the source domain controller to create the replication links. In this case, the KCC logs events each time it runs (by default, every 15 minutes) detailing the error that occurred when it attempted to add the replication links.

- Existing replication links has been inadvertently deleted in between KCC executions.

Repadmin in this scenario could be used only to diagnose. The following table explains subcommand usage that can help you diagnose the problems leading to this situation.

| Subcommand | Description |
|---|---|
| **Repadmin /showrepl** | Verify replication status. |
| **Repadmin /showconn** | Verify whether a valid connection object exists between the source and destination. |
| **Repadmin /failcache** | Resolve the underlying connection translation problems. For more information about using **Repadmin /failcache**, see Repadmin /failcache. |
| **Repadmin /KCC** | Ensure that a connection object (Automatic or Manual) has been created properly between the domain controller and its replication partner. And then force the KCC to run so that the connection object is translated to an appropriate replication link. |

**Active Directory replication has been preempted**

When Active Directory replication has been preempted, an inbound replication in progress was interrupted by a higher priority replication request. An example of a higher priority replication request is a request generated manually by using the **repadmin /sync** command.

Repadmin in this scenario could be used only to diagnose. The following table explains subcommand usage that can help you diagnose the problems leading to this situation.

| Subcommand | Description |
|---|---|
| **Repadmin /showrepl** | Verify replication status. |
| **Repadmin /queue** | Check how many inbound synchronizations are in the queue. |

**Last attempt @ never was successful**

Last attempt @ never was successful error typically indicates that KCC successfully created the replication link between the local domain controller and its replication partner, but because of the schedule or possible bridgehead overload, replication has not occurred.

Repadmin in this scenario may be used for both diagnosis and resolution. The following table explains subcommand usage that can help you diagnose or solve the problems.

| Subcommand | Description |
| --- | --- |
| **Repadmin /showrepl** | Verify replication status. |
| **Repadmin /queue** | Check how many inbound synchronizations are in the queue. |
| **Repadmin /sync** | Synchronize replication from a source domain controller. |

**Access denied**

This error indicates that the local domain controller failed to authenticate against its replication partner when creating the replication link or when trying to replicate over an existing link. This typically happens when the domain controller has been disconnected from the rest of the network for a long time and its computer account password is not synchronized with the computer account password that is stored in the Active Directory of its replication partner.

# Replication Latency

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

There are two mechanisms each specific to the underlying operating system functionality to measure replication latencies. Repadmin could be used against both environments based on the following table.

| Windows 2000 functionality | Windows Server 2003 functionality |
|---|---|
| **/latency** provides you replication latency report by measuring how recently the Intersite Topology Generator (ISTG) attribute has changed. | **/showutdvec** provides you replication latency report by leveraging a new field stored in the Up-To-Dateness (UTD) vector – "last successful replication timestamp." |
| Note that this report ceases to give meaningful results when the forest functional level is Windows Server 2003 because the **interSiteTopologyGenerator** on the **NTDS site settings object** is not updated at that functional level. | **/showutdvec** provides you replication latency report by leveraging a new field stored in the UTD vector – "last successful replication timestamp."<br>This timestamp records the last time the corresponding domain controller completed a successful replication cycle with its partner. The replication cycle may have occurred directly (direct replication partner) or indirectly (transitive replication partner). |
| Latency is shown for configuration naming context only. | Because this data is recorded on all domain controllers that host the partition, it is possible to identify non-replicating domain controllers from any domain controller in the forest that has a common partition between them. |

**Syntax**

The following command displays the amount of time between replications on a site by site basis from the perspective of the servers listed in <*DC_LIST*>, using the ISTG Keep Alive time stamp.

| **Note** |
|---|
| The ISTG Keep Alive time stamp is the mechanism used in Windows 2000 to determine whether a new ISTG is required for the site. Prior to Windows Server 2003, all ISTGs will record a time stamp every 30 minutes to indicate they are alive. After this gets replicated within the site, all of the domain controllers in the site know whether an ISTG is down or not by verifying this attribute, which is stored in Active Directory. |

**repadmin /latency** <*DC_LIST*>

```
C:\>repadmin /latency rootdns
Disclaimer:
1. Latency is shown for Configuration NC only.
2. Probes are sent once every half hour. Actual replication may occur more frequently.
3. What is normal Intersite replication frequency depends on many factors: site link schedules,
 intervals and bridgehead availability.

Gathering topology from site HUB (rootdns.contoso.com):

Replication Latency for site HUB (rootdns.contoso.com):
        Originating Site    Ver    Time Local Update    Time Orig. Update   Latency  Since Last

      =================  =====  ==================  ==================  ========  ==========

                    HUB   7263  2005-01-05 16:16:08  2005-01-05 16:16:08  00:00:00  00:36:40
                BRANCH3   1593  2005-01-05 16:28:31  2005-01-05 16:22:38  00:05:53  00:24:17
                BRANCH1   2720  2005-01-05 16:28:31  2005-01-05 15:56:14  00:32:17  00:24:17
                BRANCH2    919  2005-01-05 16:28:31  2005-01-05 16:15:50  00:12:41  00:24:17
```

**How to interpret the data**

In this example, the forest has only four sites.

| Field | Explanation |
|---|---|
| Origination site | This column has a row for each site in the forest |
| Ver | Version number for site specific **interSiteTopologyGenerator** |
| Time Local Update | Local time when the remote ISTG attribute change was replicated in. |
| Time Orig. Update | Time when the ISTG attribute was changed on the originating server. |
| Latency | Difference between the Time Orig. Update and Time Local Update |
| Since Last | Difference between the Tool execution time and Time Local Update |

Examining the UTD vector from time to time on one bridgehead server is another good way to ensure that replication is healthy. The (UTD) vector shows the last time that a domain controller has received updates from each replication partner for a particular naming context. The UTD vector is transitive in that one domain controller does not have to talk directly to another domain controller to receive an update from it.

```
repadmin /showutdvec <DC_LIST> <NamingContext> [/nocache][/latency]
```

| Parameters | Definition |
|---|---|
| *<DC_LIST>* | Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. For details about DC_LIST, see **repadmin /listhelp**. |
| *<NamingContext>* | Specifies the distinguished name of the directory partition. |
| **/nocache** | Specifies that globally unique identifier (GUIDs) are left in hexadecimal form. By default, GUIDs are translated into strings. |
| **/latency** | Sorts the information by the time required to complete the replication. By default, the information is sorted by Update Sequence Number (USN). |

```
C:\>repadmin /showutdvec rootdns cn=configuration,dc=contoso,dc=com /latency
Caching GUIDs.
..
a5515acf-6747-494a-9c14-1f1b226c5de0 @ USN     26593 @ Time 2004-04-28 23:00:17
84580b71-8ea1-48ed-9759-0ca5c5c54c2d @ USN     16480 @ Time 2004-04-28 23:09:44
65af87d6-d157-4e84-be28-551a6593844c @ USN     12750 @ Time 2004-05-19 00:51:58
ecd55a8f-9f51-4a67-9e49-c1c02568eea9 @ USN    520166 @ Time 2004-10-27 21:43:43
90fb7378-9726-430e-9469-c5ec7d230b37 @ USN    704936 @ Time 2005-01-03 13:53:58
9f4fd602-00d7-47ad-93e8-7af064384880 @ USN    131082 @ Time 2005-01-03 14:23:35
2f93bcb5-506a-4133-8a56-4bb3509c7ea4 @ USN    831399 @ Time 2005-01-03 19:32:07
8f05d838-29f0-4602-8918-5e203a0b0129 @ USN     12868 @ Time 2005-01-04 11:35:45
BRANCH1\BRANCH1                      @ USN     13520 @ Time 2005-01-05 17:00:04
BRANCH2\BRANCH2                      @ USN     13110 @ Time 2005-01-05 17:00:17
BRANCH3\BRANCH3                      @ USN     13741 @ Time 2005-01-05 17:00:18
HUB\BRANCH-HUB-BH                    @ USN     19904 @ Time 2005-01-05 17:00:50
HUB\ROOTDC01                         @ USN    944724 @ Time 2005-01-05 17:02:14
HUB\ROOTDNS                          @ USN    793222 @ Time 2005-01-05 17:24:21

C:\>_
```

**How to interpret the data**

- In Figure 3.3.2, there are four sites, two domains and six domain controllers in the forest.
- The output is a list of dates and times indicating the last time that inbound replication of the configuration container occurred from each domain controller. If an excessive amount of time has passed since replication last took place it could indicate a problem and there is reason to be concerned.
- The entries are listed by domain controller and the **/latency** parameter sorts the output by date/time.
- As given in the example, occasionally GUID's will be displayed instead of a domain controller's name. It is safe to ignore the GUID entries as these are a result of InvocationID changes or domain controllers being demoted or rebuilt and do not affect the health of the topology.
- HUB\ROOTDNS will always report the current date and time and the highest committed USN. The reason is that a domain controller does not keep itself in its own UTDVEC and always builds its entry on the fly based on the current

state.

- Latency from the perspective ROOTDNS is the difference between its current date/time with respect to other partners (direct or transitive) for the given Naming Context. For example, latency between ROOTDNS and BRANCH1 is 00:24:17.

**Display the latency only for the domain partition**

```
C:\>repadmin /showutdvec rootdns dc=contoso,dc=com /latency
Caching GUIDs.
..
HUB\ROOTDC01                          @ USN    944916 @ Time 2005-01-05 17:53:34
HUB\ROOTDNS                           @ USN    793295 @ Time 2005-01-05 17:54:22

C:\>_
```

In this example, we are only interested in the domain naming context latency. Both the domain controllers are running Windows Server 2003 and reside in the same site; hence the latency is less than a minute. Also please note that we are only displaying the domain members and not the whole forest due to the scope of the naming context.

While it is important to measure replication latencies, it is equally important to understand that intersite replication depends on many factors such as:

- Site link schedules and intervals
- Availability of bridgehead servers and their load
- Whether change notifications are enabled
- LAN/WAN infrastructure

# View Replication Metadata of an Object

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Displays the replication metadata for a specified object stored in Active Directory, such as attribute ID, version number, originating and local Update Sequence Number (USN), and originating server's globally unique identifier (GUID) and date and time stamp. By comparing the replication metadata for the same object on different domain controllers, an administrator can determine whether replication has occurred.

**Syntax**

```
repadmin /showobjmeta <DC_LIST> <ObjectDN> [/nocache] [/linked]
```

| Parameters | Definitions |
|---|---|
| <DC_LIST> | Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. For details about DC_LIST, see **repadmin /listhelp**. |
| <ObjectDN> | Specifies the distinguished name of the object. |
| **/nocache** | Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings. |
| **/linked** | Displays metadata associated with, but not stored with, the specified object. |

**Example 1: Metadata of a group object**

In this example, we are viewing the metadata of a group object (Domain Admins) and therefore the forward links (members) are listed as well.

```
C:\>repadmin /showobjmeta branch-h* "cn=domain admins,cn=users,dc=research,dc=contoso,dc=com"

repadmin running command /showobjmeta against server BRANCH-HUB-BH.research.contoso.com

15 entries.
Loc.USN                          Originating DC   Org.USN  Org.Time/Date        Ver Attribute
=======                          ============== ========= =============        === =========
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 objectClass
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 cn
  12314                         HUB\BRANCH-HUB-BH     12314 2005-01-04 03:46:48    2 description
  20644                         HUB\BRANCH-HUB-BH     20644 2005-01-06 01:01:17    2 member
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 instanceType
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 whenCreated
  12903                         HUB\BRANCH-HUB-BH     12903 2005-01-04 05:04:01    2 nTSecurityDescriptor
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 name
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 objectSid
  12903                         HUB\BRANCH-HUB-BH     12903 2005-01-04 05:04:01    1 adminCount
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 sAMAccountName
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 sAMAccountType
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 groupType
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 objectCategory
  12312                         HUB\BRANCH-HUB-BH     12312 2005-01-04 03:46:48    1 isCriticalSystemObjec
2 entries.
Type    Attribute     Last Mod Time                          Originating DC  Loc.USN Org.USN Ver
======= ===========  ============                          ================ ======= ======= ===
        Distinguished Name
        ==========================
LEGACY         member
        CN=Administrator,CN=Users,DC=research,DC=contoso,DC=com
LEGACY         member
        CN=Lee Oonmen,CN=Users,DC=research,DC=contoso,DC=com
```

**Example 2: Comparing replication metadata of a user object between two domain controllers**

A domain administrator has restricted user Lee's logon hrs. Lee claims he could still log on during restricted hours from BRANCH3 as opposed to other branch offices. The domain administrator could easily figure whether this is related to Active Directory replication latencies by comparing the replication metadata of Lee's account.

```
C:\>repadmin /showobjmeta branch-h* "cn=lee Oommen,cn=users,dc=research,dc=contoso,dc=com"

repadmin running command /showobjmeta against server BRANCH-HUB-BH.research.contoso.com

27 entries.
Loc.USN                          Originating DC    Org.USN  Org.Time/Date         Ver Attribute
=======                          ==============    ======== =============         === =========
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 objectClass
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 cn
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 sn
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 givenName
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 instanceType
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 whenCreated
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 displayName
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 nTSecurityDescriptor
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 name
  20583                         HUB\BRANCH-HUB-BH     20583 2005-01-06 00:52:03     4 userAccountControl
  20578                         HUB\BRANCH-HUB-BH     20578 2005-01-06 00:52:03     1 codePage
  20578                         HUB\BRANCH-HUB-BH     20578 2005-01-06 00:52:03     1 countryCode
  20579                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 dBCSPwd
  20654                         HUB\BRANCH-HUB-BH     20654 2005-01-06 01:19:59     2 logonHours
  20579                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 unicodePwd
  20579                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 ntPwdHistory
  20579                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 pwdLastSet
  20578                         HUB\BRANCH-HUB-BH     20578 2005-01-06 00:52:03     1 primaryGroupID
  20580                         HUB\BRANCH-HUB-BH     20580 2005-01-06 00:52:03     1 supplementalCredentials
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 objectSid
  20578                         HUB\BRANCH-HUB-BH     20578 2005-01-06 00:52:03     1 accountExpires
  20579                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 lmPwdHistory
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 sAMAccountName
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 sAMAccountType
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 userPrincipalName
  20577                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 objectCategory
  20585                         HUB\BRANCH-HUB-BH     20585 2005-01-06 00:53:04     1 mail
0 entries.
```

Figure 3.4.2 is the metadata of Lee from HUB domain controller (where the change was made) and Figure 3.4.3 is the metadata from the BRANCH3 domain controller. The attribute **logonHours** has been highlighted for clarity.

BRANCH-HUB-BH has version 2, last Orig. time/date is 2005-01-06 01:19:59 and Orig.USN as 20654.

BRANCH3 is still on version 1, last Orig. time/date is 2005-01-06 00:52:03 and Orig.USN as 20578 and hence the logon succeeds in BRANCH3 because that domain controller has not yet replicated the update.

```
C:\>repadmin /showobjmeta branch3 "cn=lee Oommen,cn=users,dc=research,dc=contoso,dc=com"

27 entries.
Loc.USN                          Originating DC    Org.USN  Org.Time/Date         Ver Attribute
=======                          ==============    ======== =============         === =========
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 objectClass
  16508                           BRANCH3\BRANCH3     16508 2005-01-06 01:01:40     1 cn
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 sn
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 givenName
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 instanceType
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 whenCreated
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 displayName
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 nTSecurityDescriptor
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 name
  16508                         HUB\BRANCH-HUB-BH     20583 2005-01-06 00:52:03     4 userAccountControl
  16508                         HUB\BRANCH-HUB-BH     20578 2005-01-06 00:52:03     1 codePage
  16508                         HUB\BRANCH-HUB-BH     20578 2005-01-06 00:52:03     1 countryCode
  16508                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 dBCSPwd
  16508                         HUB\BRANCH-HUB-BH     20578 2005-01-06 00:52:03     1 logonHours
  16508                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 unicodePwd
  16508                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 ntPwdHistory
  16508                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 pwdLastSet
  16508                         HUB\BRANCH-HUB-BH     20578 2005-01-06 00:52:03     1 primaryGroupID
  16508                         HUB\BRANCH-HUB-BH     20580 2005-01-06 00:52:03     1 supplementalCredentials
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 objectSid
  16508                         HUB\BRANCH-HUB-BH     20578 2005-01-06 00:52:03     1 accountExpires
  16508                         HUB\BRANCH-HUB-BH     20579 2005-01-06 00:52:03     2 lmPwdHistory
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 sAMAccountName
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 sAMAccountType
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 userPrincipalName
  16508                         HUB\BRANCH-HUB-BH     20577 2005-01-06 00:52:02     1 objectCategory
  16508                         HUB\BRANCH-HUB-BH     20585 2005-01-06 00:53:04     1 mail
0 entries.
```

# Display the Attributes of a Specific Object

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The **/showattr** operation displays the attributes and contents of an object.

**Syntax**

```
repadmin /showattr <DC_LIST> <OBJ_LIST> <OBJ_LIST_OPTIONS> [/atts: <<att1>>,<<att2>>,...] [/allvalues]
[/long] [/dumpallblob]
```

| Parameters | Definition |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. For details about DC_LIST, see **repadmin /listhelp**. |
| <*OBJ_LIST*> | This parameter takes a distinguished name or a special keyword that expands into a distinguished name. The keywords are as follows:<br>• **Ncobj:config:** Distinguished name of the Configuration partition of the domain controller<br>• **Ncobj:schema:** Distinguished name of the Schema partition of the domain controller<br>• **Ncobj:domain:** Distinguished name of the Domain partition of the domain controller<br>• **Dsaobj:** NTDS settings object of the directory server |
| <*OBJ_LIST_OPTIONS*> | The OBJ_LIST_OPTION parameter is required to perform a generic Lightweight Directory Access Protocol (LDAP) search from the command line. The parameter requires a BaseDN, with the ability to use a search modifier option. The valid search modifier options are as follows:<br>• **/filter:**<*ldap_filter*><br>• **/base**<br>• **/subtree**<br>• **/onelevel** |
| **[/atts:** <*att1*>,**<att2>,...** | Returns only the attributes that are specified. Separate each listed attribute with a comma. |
| **/allvalues** | For an attribute, the tool only displays 20 values unless this flag is specified, in which case it shows all values. |
| **/long** | Displays one value per line. |
| **/dumpallblob** | Dumps the BLOB in a default byte-by-byte format if there is not a friendly formatted interpretation available for it. |

**Note**

A BLOB in this context means an attribute that is not a simple type, like a string or an integer. A BLOB is a complex structured type that is stored as binary bytes. To make sense of the BLOB, a program must interpret it and format it. A friendly BLOB is a BLOB that the program knows about and can format in an understandable way. The program has a list of BLOBs that it understands.

**Example: Display select attributes**

Please note how we specify the naming context as ncobj:domain:

```
C:\>repadmin /showattr ROOTDC01 ncobj:domain: /filter:"(&(objectCategory=us
er)(samaccountname=lee))" /subtree /atts:displayname,description,mail,membe
rof /long
DN: CN=lee,CN=Users,DC=contoso,DC=com
    1> description: Infrastructure RRE
    1> displayName: Lee Oommen
    3> memberOf: CN=Enterprise Admins,CN=Users,DC=contoso,DC=com;
          CN=Domain Admins,CN=Users,DC=contoso,DC=com;
          CN=Group Policy Creator Owners,CN=Users,DC=contoso,DC=com
    1> mail: loommen@microsoft.com


C:\>_
```

# How Up to Date Are My Domain Controllers?

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

**Checkprop** compares properties of specified domain controllers to determine if they are up-to-date with each other. The source domain controller contains the original information that needs to be checked. The destination domain controller data is compared to the source domain controller data.

**Syntax**

```
repadmin /checkprop <DC_LIST> <NamingContext> <OriginatingDCInvocationID> <OriginatingUSN>
```

| Parameter | Definition |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |
| <*NamingContext*> | Specifies the distinguished name of the directory partition on the source domain controller. |
| <*OriginatingDCInvocationID*> | Specifies the unique hexadecimal number that identifies an object on a source domain controller. The InvocationID can be retrieved by using the **/showrepl** operation. |
| <*OriginatingUSN*> | Specifies the Update Sequence Number (USN) for the object on the source domain controller. The USN is for the object whose InvocationID is already listed. |

**Example: Checking replication latency on the BRANCH3 domain controller**

Latency output reveals that the highest OriginatingUSN that BRANCH3 has knowledge of for its HUB site bridgehead server, BRANCH-HUB-BH, is **137844**. It is also apparent that the last successful replication attempt with this HUB site bridgehead server was just less than 5 minutes.

```
C:\>repadmin /showutdvec branch3 cn=configuration,dc=contoso,dc=com /latency
Caching GUIDs.
..
a5515acf-6747-494a-9c14-1f1b226c5de0 @ USN     26593 @ Time 2004-04-28 23:00:17
84580b71-8ea1-48ed-9759-0ca5c5c54c2d @ USN     16480 @ Time 2004-04-28 23:09:44
65af87d6-d157-4e84-be28-551a6593844c @ USN     12750 @ Time 2004-05-19 00:51:58
ecd55a8f-9f51-4a67-9e49-c1c02568eea9 @ USN    520166 @ Time 2004-10-27 21:43:43
90fb7378-9726-430e-9469-c5ec7d230b37 @ USN    704936 @ Time 2005-01-03 13:53:58
9f4fd602-00d7-47ad-93e8-7af064384880 @ USN    131082 @ Time 2005-01-03 14:23:35
2f93bcb5-506a-4133-8a56-4bb3509c7ea4 @ USN    831399 @ Time 2005-01-03 19:32:07
8f05d838-29f0-4602-8918-5e203a0b0129 @ USN     12868 @ Time 2005-01-04 11:35:45
BRANCH3\BRANCH3 (retired)            @ USN     63018 @ Time 2005-02-12 07:47:06
BRANCH2\BRANCH2                      @ USN     36062 @ Time 2005-02-16 14:08:05
BRANCH1\BRANCH1                      @ USN     68884 @ Time 2005-02-16 14:08:05
HUB\ROOTDNS                          @ USN    910177 @ Time 2005-02-16 14:12:17
HUB\ROOTDC01                         @ USN   1170137 @ Time 2005-02-16 14:12:32
HUB\BRANCH-HUB-BH                    @ USN    137844 @ Time 2005-02-16 14:18:45
BRANCH3\BRANCH3                      @ USN     64169 @ Time 2005-02-16 14:23:13

C:\>_
```

**Example: Comparing how up-to-date other domain controllers in the enterprise are with respect to the OriginatingUSN**

In Figure 3.6.2, note that **BRANCH2** domain controller is not up-to-date with the rest of the domain controllers.

```
C:\>repadmin /checkprop branch3 cn=configuration,dc=contoso,dc=com 0cc4891a-4cca-49e8-bbc8
-123de8d656ab 137844
HUB\ROOTDNS: yes (USN 137844)
HUB\ROOTDC01: yes (USN 137844)
HUB\BRANCH-HUB-BH: yes (USN 137844)
BRANCH1\BRANCH1: yes (USN 137844)
BRANCH2\BRANCH2: ** NO! ** (USN 137829)
BRANCH3\BRANCH3: yes (USN 137844)

C:\>_
```

**Example: Further investigation from the perspective of the BRANCH2 domain controller**

Latency was calculated for BRANCH2 which revealed that it is not aware of the latest OriginatingUSN from BRANCH-HUB-BH and in fact it is behind by approximately 20 minutes. Because the latency in this example is just less than 20 minutes (replication interval being 30 minutes) it is expected to catch up during the next replication cycle.

```
C:\>repadmin /showutdvec branch2 cn=configuration,dc=contoso,dc=com /latency
Caching GUIDs.
..
a5515acf-6747-494a-9c14-1f1b226c5de0 @ USN     26593 @ Time 2004-04-28 23:00:17
84580b71-8ea1-48ed-9759-0ca5c5c54c2d @ USN     16480 @ Time 2004-04-28 23:09:44
65af87d6-d157-4e84-be28-551a6593844c @ USN     12750 @ Time 2004-05-19 00:51:58
ecd55a8f-9f51-4a67-9e49-c1c02568eea9 @ USN    520166 @ Time 2004-10-27 21:43:43
90fb7378-9726-430e-9469-c5ec7d230b37 @ USN    704936 @ Time 2005-01-03 13:53:58
9f4fd602-00d7-47ad-93e8-7af064384880 @ USN    131082 @ Time 2005-01-03 14:23:35
2f93bcb5-506a-4133-8a56-4bb3509c7ea4 @ USN    831399 @ Time 2005-01-03 19:32:07
8f05d838-29f0-4602-8918-5e203a0b0129 @ USN     12868 @ Time 2005-01-04 11:35:45
BRANCH3\BRANCH3 (retired)            @ USN     63018 @ Time 2005-02-12 07:47:06
BRANCH1\BRANCH1                      @ USN     68884 @ Time 2005-02-16 14:08:05
HUB\BRANCH-HUB-BH                    @ USN    137829 @ Time 2005-02-16 14:08:30
BRANCH3\BRANCH3                      @ USN     64151 @ Time 2005-02-16 14:09:09
HUB\ROOTDNS                          @ USN    910163 @ Time 2005-02-16 14:09:23
HUB\ROOTDC01                         @ USN   1170114 @ Time 2005-02-16 14:09:23
BRANCH2\BRANCH2                      @ USN     36073 @ Time 2005-02-16 14:28:58


C:\>_
```

# Can I Look at My Connection Objects and Schedule Details?

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Every domain controller that is also a member of the SYSVOL replica set has to have at least one inbound connection. Otherwise, Active Directory and File Replication Service (FRS) would not replicate inbound. The **/showconn** subcommand is very useful to verify this especially:

- When you don't have access to the graphical user interface (GUI)

or

- When you find it task-oriented to directly connect to the various domain controllers from the user interface (UI) to look at Active Directory topology from the perspective of that domain controller.

  The **/showconn** subcommand displays the connection objects for a specified domain controller. The default is the local site.

**Syntax**

```
repadmin /showconn <DC_LIST> {<ServerRDN> | <ContainerDN> | <DC_GUID>} [/From: <ServerRDN>] [/intersite]
```

| Parameter | Definition |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller from where to read the configuration, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |
| <*ServerRDN*> | Specifies the relative distinguished name of a server. |
| <*ContainerDN*> | Specifies the distinguished name of a container. |
| <*DC_GUID*> | Specifies the unique hexadecimal number that identifies the domain controller. The globally unique identifier (GUID) can be retrieved by using the **/showreps** operation. |
| **/intersite** | Displays only those connection objects that are between sites. |

**Example: Simple usage of /showconn**

Figure 3.7.1 shows a simple example of output returned by **/showconn**.

```
          C:\>repadmin /showconn branch1
Base DN: CN=BRANCH1,CN=Sites,CN=Configuration,DC=contoso,DC=com
=====  KCC CONNECTION OBJECTS  ================================
Connection --
     Connection name : ed5e0d25-bec3-4556-9f18-f24cf4ea3a57
     Server DNS name : BRANCH1.research.contoso.com
     Server DN  name : CN=NTDS Settings,CN=BRANCH1,CN=Servers,CN=BRANCH1,CN=Sites,CN=C
onfiguration,DC=contoso,DC=com
          Source: HUB\BRANCH-HUB-BH
               No Failures.
          TransportType: IP
          options:  isGenerated overrideNotifyDefault
          ReplicatesNC: DC=DomainDnsZones,DC=research,DC=contoso,DC=com
          Reason:  IntersiteTopology
               Replica link has been added.
          ReplicatesNC: DC=ForestDnsZones,DC=contoso,DC=com
          Reason:  IntersiteTopology
               Replica link has been added.
          ReplicatesNC: CN=Configuration,DC=contoso,DC=com
          Reason:  IntersiteTopology
               Replica link has been added.
```

```
          ReplicatesNC: DC=research,DC=contoso,DC=com
          Reason:  IntersiteTopology
               Replica link has been added.
1 connections found.
```

In the example in figure 3.7.1, there is only one connection object for the BRANCH1 site. It is also automatically created (options: isgenerated). Depending on the number of connection objects, we may have to further qualify our query to just list what we are interested in such as in the following cases:

Figure 3.7.2

```
repadmin /showconn BRANCH1 CN=HUB,CN=Sites,CN=Configuration,DC=contoso,DC=com
/intersite /v
Here repadmin contacts BRANCH1 DC and list all the incoming intersite connections for
HUB site with verbose details.
```

Figure 3.7.3

```
repadmin /showconn BRANCH-HUB-BH BRANCH-HUB-BH /from:BRANCH3
Here repadmin contacts the BRANCH-HUB-BH DC which is also located in the HUB site and
displays just the connection object from BRANCH3 DC to BRANCH-HUB-BH.
```

With the verbose switch, showconn provides you much more information such as the following:

- Connection replication schedule
- Partition Replication Schedule Loading

Figure 3.7.4 Connection replication schedule

```
day: 0123456789ab0123456789ab
Sun: ffffffffffffffffffffffff
Mon: ffffffffffffffffffffffff
Tue: ffffffffffffffffffffffff
Wed: ffffffffffffffffffffffff
Thu: ffffffffffffffffffffffff
Fri: ffffffffffffffffffffffff
Sat: ffffffffffffffffffffffff
```

Every single number of the above represents one hour of the day as a decimal 4-bit value. Each single bit represents 15 minutes of this hour.

So if we have "1" in decimal, then one bit is set in binary (0001) and we replicate once per hour in which case the output will be:

111111111111111111111111

If the decimal value is five (0101 in binary) we replicate twice per hour, for example:

555555555555555555555555

Finally if it is F (1111) we replicate four times per hour:

FFFFFFFFFFFFFFFFFFFFFFFF

So in our example we replicate four times per hour for the entire week.

# Fine-Tuning Change Notification Values

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Replication within a site occurs as a response to changes elsewhere in the site. Replication across sites occurs based on the replication schedule and interval. It is also possible to enable change notifications across sites.

When a change occurs on a domain controller, two configurable intervals determine the delay between the following events:

- Notification to the first partner.
- Notification to each subsequent partner.

The above two intervals serve to:

- Stagger network traffic caused by replication.
- Spreads out the load of responding to replication requests from its partners.

The following table lists the default notification delays:

| Operating system | Notify first partner (sec) | Subsequent partner (sec) | Forest functional level |
|---|---|---|---|
| Windows 2000 | 300 | 30 | Windows 2000 |
| Windows Server 2003 (upgraded from Windows 2000)<br>**Note**<br>If you changed the default values, then those values that you set are retained after you upgrade from Windows 2000 to Windows Server 2003. | 300 | 30 | Windows 2000 |
| Windows Server 2003 | 15 | 3 | Windows 2000 |
| Windows Server 2003 (either upgraded from Windows 2000 or a clean install) | 15 | 3 | Windows Server 2003 |

The following table lists the storage location of notification delay values for each operating system.

| Operating system | Location | Attribute |
|---|---|---|
| Windows 2000 Server | HKLM\SYSTEM\CSS\Services\NTDS\Parameters | • Replicator notify pause after modify (secs)<br>• Replicator notify pause between Directory System Agent (DSAs) (secs) |

| | | |
|---|---|---|
| Windows Server 2003 | Cross-reference object for each directory partition in the configuration partition. | • msDS-Replication-Notify-First-DSA-Delay<br>• msDS-Replication-Notify-Subsequent-DSA-Delay |

**Repadmin /notifyopt** could be used to view or change the notification timing settings of a specified directory partition in Windows Server 2003.

**Syntax**

```
repadmin /notifyopt <DC_LIST> <NamingContext> [/first: Value] [/subs: Value]
```

| Parameter | Definition |
|---|---|
| *<DC_List>* | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about *<DC_LIST>*, see **repadmin /listhelp.** |
| *<NamingContext>* | Specifies the distinguished name of the directory partition on the source domain controller. |
| **/first** | The number of seconds after a change is made before the domain controller notifies its first replication partner that there is a change. |
| **/subs** | Once the first replication partner is notified of a change, the subs parameter specifies the number of seconds to wait before notifying the next replication partner. |

**Example 1: Displaying the default notification delay on the ForestDnsZones partition**

```
C:\>repadmin /notifyopt rootdns DC=ForestDnsZones,DC=contoso,DC=com

Current Notification Options:
Replication-Notify-First-DSA-Delay is not set.
Replication-Notify-Subsequent-DSA-Delay is not set.


C:\>_
```

**Example 2: Changing the defaults to 300/30 on the ForestDnsZones**

```
C:\>repadmin /notifyopt rootdns DC=ForestDnsZones,DC=contoso,DC=com /first:300 /
subs:30

Current Notification Options:
Replication-Notify-First-DSA-Delay is not set.
Replication-Notify-Subsequent-DSA-Delay is not set.
The set operation is allowed on the Domain Naming Master Only.
Your Domain Naming Master is: CN=ROOTDC01,CN=Servers,CN=HUB,CN=Sites,CN=Configur
ation,DC=contoso,DC=com


C:\>repadmin /notifyopt rootdc01 DC=ForestDnsZones,DC=contoso,DC=com /first:300
/subs:30

Current Notification Options:
Replication-Notify-First-DSA-Delay is not set.
Replication-Notify-Subsequent-DSA-Delay is not set.

New Notification Options:
Replication-Notify-First-DSA-Delay: 300
Replication-Notify-Subsequent-DSA-Delay: 30


C:\>_
```

**Note**

In order to make this change, you have to run **/notifyopt** against the Domain Naming Master. See the highlighted text in figure 3.8.2.

# Forcing Replication

Updated: July 14, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Sometimes it becomes necessary to forcefully replicate objects and entire partitions between domain controllers that may or may not have replication agreements

| **Important** |
| --- |
| These are very powerful subcommands and should be used sparingly as they do not follow replication agreements that are in place and have the potential to cause replication storm and break Active Directory if not used properly. |

**Replicate a single object between two domain controllers**

The **repadmin /replsingleobj** command replicates a single object between any two domain controllers that have partitions in common. The two domain controllers do not require a replication agreement between them. Replication agreements can be shown by using the **repadmin /showreps** command.

**Syntax**

```
repadmin /replsingleobj <DC_LIST> <Source DSA_Name> <ObjectDN>
```

| Parameter | Definition |
| --- | --- |
| <*DC_LIST*> | Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. For details about <*DC_LIST*>, see **repadmin /listhelp** |
| <*Source DSA_Name*> | Specifies the name of the source domain controller. You can specify a host name or the unique hexadecimal number that identifies the source domain controller. You can retrieve the objectGUID by using the **/showreps** operation. |
| <*ObjectDN*> | Specifies the distinguished name of the object. |

**Example: Replicate a single object between all the branch domain controllers by using wild card character**

```
C:\>repadmin /replsingleobj branch* 8168c4c2-bf27-4f92-88ce-3d502b596596 "cn=lee Oo
en,cn=users,dc=research,dc=contoso,dc=com"

repadmin running command /replsingleobj against server BRANCH-HUB-BH.research.conto
.com

Successfully replicated object cn=lee Oommen,cn=users,dc=research,dc=contoso,dc=com
o BRANCH-HUB-BH.research.contoso.com from <GUID=8168c4c2-bf27-4f92-88ce-3d502b59659
.

repadmin running command /replsingleobj against server BRANCH1.research.contoso.com

Successfully replicated object cn=lee Oommen,cn=users,dc=research,dc=contoso,dc=com
o BRANCH1.research.contoso.com from <GUID=8168c4c2-bf27-4f92-88ce-3d502b596596>.

repadmin running command /replsingleobj against server BRANCH2.research.contoso.com

Successfully replicated object cn=lee Oommen,cn=users,dc=research,dc=contoso,dc=com
o BRANCH2.research.contoso.com from <GUID=8168c4c2-bf27-4f92-88ce-3d502b596596>.

repadmin running command /replsingleobj against server BRANCH3.research.contoso.com

Successfully replicated object cn=lee Oommen,cn=users,dc=research,dc=contoso,dc=com
o BRANCH3.research.contoso.com from <GUID=8168c4c2-bf27-4f92-88ce-3d502b596596>.

C:\>_
```

**Force a replication event between two partners**

The **repadmin /replicate** command starts a replication event for the specified directory partition between the source and destination domain controllers. The source universally unique identifier (UUID) can be determined when viewing the replication partners by using the **/showreps** operation.

| **Important** |
| --- |
| The **repadmin /replicate** command will not work if the partners do not have the specified partition in common or replication agreement between them. |

**Syntax1**

```
repadmin /replicate <Destination_DC_LIST> <Source_DC_NAME> <Naming Context> [/force] [/async] [/full]
[/addref] [/readonly]
```

**Syntax2**

```
repadmin /replicate <Destination_DC_LIST> <Naming Context> [/allsources] [/force] [/async] [/full] [/addref]
[/readonly]
```

| Parameter | Definition |
|---|---|
| <*Destination_DC_LIST*> | Specifies the host name of the destination domain controller (Directory Server Agent) with which you want to replicate. For details about <*DC_LIST*>, see **repadmin /listhelp.** |
| <*Source_DC_NAME*> | Specifies the host name of the source domain controller with which you want to replicate. This parameter accepts a globally unique identifier (GUID), GUID-based Domain Name System (DNS) name, or the name of a server object. |
| <*Naming Context*> | Specifies the distinguished name of the directory partition. |
| **/force** | This parameter is used to override the Disable Replication option on a server. |
| **/async** | Specifies that the replication will be asynchronous. This means that repadmin starts the replication event, but it does not expect an immediate response from the destination domain controller. Use this parameter when there are slow links between domain controllers. |
| **/full** | Forces a full replication of all objects from the destination domain controller. |
| **/addref** | Directs the source to check for a notification entry on the source. If the source does not have a notification entry for this destination, one is added. |
| **/allsources** | A given destination can have multiple sources for the same naming context. Directs the destination to sync with all sources instead of just one. This parameter cannot be used with <*Destination_DC_LIST*>. |
| **/readonly** | This parameter is ignored by the **/replicate** operation. |

**Example: replicate in domain partition between two specific partners**

In the example in figure 3.9.2.1, we are attempting to replicate in domain partition between two specific partners. But the source domain controller is rejecting replication requests as configured by the administrator for valid reasons.

```
C:\>repadmin /replicate branch1 branch-hub-bh dc=research,dc=contoso,dc=com

DsReplicaSync() failed with status 8456 (0x2108):
    The source server is currently rejecting replication requests.


C:\>_
```

In the next example, we run **repadmin /showrepl** against the source domain controller (BRANCH-HUB-BH) to read the domain controller options. Figure 3.9.2.2 highlights that outbound replication is currently disabled (DISABLE_OUTBOUND_REPL).

```
C:\>repadmin /showrepl branch-hub-bh
HUB\BRANCH-HUB-BH
DC Options: DISABLE_OUTBOUND_REPL
Site Options: (none)
DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53
DC invocationID: 0cc4891a-4cca-49e8-bbc8-123de8d656ab

==== INBOUND NEIGHBORS ======================================

CN=Configuration,DC=contoso,DC=com
    BRANCH2\BRANCH2 via RPC
```

We could work around this by using the **/force** switch as seen in figure 3.9.2.3. However, use caution you when using the force replication feature. The **/force** switch is dangerous because it overrides any precautions that have been implemented by an enterprise administrator to address specific business needs. For example, in a large forest with hundreds of sites connected across unreliable WAN links, use of the **/force** switch to replicate changes across forest might cause a replication storm (depending on the changes) that the WAN could not handle.

```
C:\>repadmin /replicate branch1 branch-hub-bh dc=research,dc=contoso,dc=com
 /force
Sync from branch-hub-bh to branch1 completed successfully.


C:\>_
```

**Force a replication event with all partners**

the **repadmin /syncall** command synchronizes a specified domain controller with all replication partners.

**Syntax**

```
repadmin /syncall <DC> [<NamingContext>] [<Flags>]
```

| Parameter | Definition |
|---|---|
| *<DC>* | Specifies the host name of the domain controller to synchronize with all replication partners. |
| *<NamingContext>* | Specifies the distinguished name of the directory partition. |
| *<Flags>* | Performs specific actions during the replication. |

The following table lists the flags that you can use with **repadmin /syncall**.

| Flag | Description |
|---|---|
| /a | Abort if any server is unavailable. |
| /A | Sync all naming contexts which are held on the home server. |
| /d | Identify servers by distinguished name in messages. |
| /e | Enterprise, cross sites |
| /h | Print this help screen. |
| /i | Iterate indefinitely. |
| /l | Perform showreps on each server pair in path instead of synchronizing. |
| /j | Synchronize adjacent servers only. |
| /p | Pause for possible user abort after every message. |
| /P | Push changes outward from home server. |
| /q | Run in quiet mode, suppress call back messages. |
| /Q | Run in very quiet mode, report fatal errors only. |
| /s | Do not synchronize. |
| /S | Skip initial server response check. |

**Important**

Use this command and the above flags cautiously or you can damage the replication system because this command does not follow replication agreements nor honor any replication restrictions such as DISABLE_INBOUND_REPL or DISABLE_OUTBOUND_REPL

**Example 1: Synchronizing Configuration Partition within the site**

```
C:\>repadmin /syncall branch-hub-bh cn=configuration,dc=contoso,dc=com
Syncing partition: cn=configuration,dc=contoso,dc=com
CALLBACK MESSAGE: The following replication is in progress:
    From: 076cd5dd-e25e-4897-acd2-7c8691621522._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: The following replication completed successfully:
    From: 076cd5dd-e25e-4897-acd2-7c8691621522._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: The following replication is in progress:
    From: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: The following replication completed successfully:
    From: 2a92f776-6c0f-4cb4-a111-f5dcd447af6c._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: SyncAll Finished.
SyncAll terminated with no errors.
```

There will be two callback messages for each partner in figure 3.9.3.1. One reports the progress and the other reports either

success or failure (with explanation). Also notice that domain controllers are denoted by their GUID CNAMES as used in replication.

**Example 2: Crossing site boundaries and other features**

By default, **repadmin /syncall** does not cross site boundaries as depicted in figure 3.9.3.2. BRANCH-HUB-BH does not have any domain members in its own site for domain dc=research,dc=contoso,dc=com. In this case, use **/e**.

```
C:\>repadmin /syncall branch-hub-bh dc=research,dc=contoso,dc=com
Syncing partition: dc=research,dc=contoso,dc=com
CALLBACK MESSAGE: SyncAll Finished.
SyncAll terminated with no errors.

C:\>_
```

In figure 3.9.3.3, we are using three additional flags. The **/d** flag would translate the GUID CNAME to the distinguished name of the domain controller. The **/e** flag is used to cross site boundaries. The **/a** flag is used to abort if any domain controller is unavailable. In this example, the BRANCH2 domain controller was not reachable and therefore, the process was aborted.

```
C:\>repadmin /syncall branch-hub-bh dc=research,dc=contoso,dc=com /d /e /a
Syncing partition: dc=research,dc=contoso,dc=com
CALLBACK MESSAGE: Error contacting server CN=NTDS Settings,CN=BRANCH2,CN=Se
rvers,CN=BRANCH2,CN=Sites,CN=Configuration,DC=contoso,DC=com (network error
): 1722 (0x6ba):
    The RPC server is unavailable.

SyncAll cancelled by user request.

C:\>_
```

In figure 3.9.3.4, **repadmin /syncall** did succeed because the problem with the BRANCH2 domain controller was fixed. Also notice that we omitted the **/d** switch so that the GUID names are not translated.

```
C:\>repadmin /syncall branch-hub-bh dc=research,dc=contoso,dc=com /e /a
Syncing partition: dc=research,dc=contoso,dc=com
CALLBACK MESSAGE: The following replication is in progress:
    From: 1b510929-29d7-4a60-a877-aa64e027bb29._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: The following replication completed successfully:
    From: 1b510929-29d7-4a60-a877-aa64e027bb29._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: The following replication is in progress:
    From: 8168c4c2-bf27-4f92-88ce-3d502b596596._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: The following replication completed successfully:
    From: 8168c4c2-bf27-4f92-88ce-3d502b596596._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: The following replication is in progress:
    From: 7e50725f-29c4-4d8a-80a2-c165e1c94e79._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: The following replication completed successfully:
    From: 7e50725f-29c4-4d8a-80a2-c165e1c94e79._msdcs.contoso.com
    To  : 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
CALLBACK MESSAGE: SyncAll Finished.
SyncAll terminated with no errors.
```

# Keeping Track of Changes That Have Occurred Over a Period of Time

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

There could be multiple occasions where we would be interested in finding out the number of changes that are either pending replication or that have occurred to a specified directory partition over a period of time.

For example:

- You may want to get statistics of all the changes that have occurred to a domain partition over a period of one day or one week so that you can use this data to either support or calculate intersite replication bandwidth requirements.
- This will help with troubleshooting replication issues and reviewing changes that have not replicated between two partners.

The **repadmin /showchanges** command has two syntaxes that can helpful in these situations.

**Syntax1**

```
repadmin /showchanges <SourceDC> <NamingContext> [/cookie: <File>] [/atts: <attribute>,<attribute>,...]
```

**Syntax2**

```
repadmin /showchanges <DestDC> <SourceDCObjectGUID> <NamingContext> [/verbose] [/statistics] [/noincremental]
[/objectsecurity] [/ancestors] [/atts: <attribute1>,<attribute2>,...] [/filter: <ldap filter>]
```

| Parameter | Description |
|---|---|
| <*DestDC*> | Specifies the host name of the destination domain controller from which to enumerate the host domain controllers. |
| <*SourceDC*> | Specifies the host name of the domain controller that hosts the directory partition whose changes you want to view. |
| <*NamingContext*> | Specifies the distinguished name of the directory partition. |
| **/cookie:** <*File*> | Specifies a name for the file to which list changes are saved. |
| **/atts:** <*attribute*>,<*attribute*>,... | Returns only the attributes specified. Separate each listed attribute with a comma. |
| <*SourceDCObjectGUID*> | Specifies the unique hexadecimal number that identifies the object whose changes will be listed. The objectGUID can be retrieved by using the **/showreps** operation. |
| **/verbose** | Lists detailed information. |
| **/statistics** | Displays a summary of information about changes instead of a list of individual changes. |
| **/noincremental** | Returns changes in value change format, which lists current values for attributes as well with attributes that have been added or deleted. If not specified, changes are returned in attribute change format, which shows only the current value of the attribute. |
| **/objectsecurity** | Overrides the need for the GetChanges right to the directory partition. By default, this right is necessary to run the **GetChanges** parameter. However, only changes that the currently logged on user has the rights to view are displayed. |
| **/ancestors** | Returns changes in Update Sequence Number (USN) order. |
| **/filter:** <*ldap filter*> | Returns only those changes that meet the filter requirements. |

Syntax 1 can be used to compare changes that occurred to a specified directory partition over a period of time.

The idea here is to:

1. Create a cookie file that saves changes to a directory partition that could be used for later comparisons. The first time you use the cookie option, it may take a long time (depending on the size of your partition) to create the file. It is important to note that we store only metadata about all the changes that have occurred to this cookie file on the entire set of domain controllers.

2. Later on when you present this cookie file to any domain controller, it will update the cookie file and provide you with just the change deltas since the last time it was updated.

**Example: Compare changes occurred to configuration partition over a period of time**

```
C:\>repadmin /showchanges rootdns cn=configuration,dc=contoso,dc=com /cookie:config
Using cookie from file config (444 bytes)

==== SOURCE DC: rootdns ====

Objects returned: 3
(0) modify CN=NTDS Site Settings,CN=HUB,CN=Sites,CN=Configuration,DC=contoso,DC=com
    1> objectGUID: 34a1f87f-17f0-41c6-b084-f30107916968
    1> instanceType: 0x4 = ( IT_WRITE )
    1> interSiteTopologyGenerator: (GUID=ddd56c075ee29748acd27c8691621522);CN=NTDS Setting
s,CN=ROOTDNS,CN=Servers,CN=HUB,CN=Sites,CN=Configuration,DC=contoso,DC=com
(1) modify CN=NTDS Site Settings,CN=BRANCH2,CN=Sites,CN=Configuration,DC=contoso,DC=com
    1> objectGUID: 22cc7552-61c0-4020-9921-faa859e4d204
    1> instanceType: 0x4 = ( IT_WRITE )
    1> interSiteTopologyGenerator: (GUID=c2c4688127bf924f88ce3d502b596596);CN=NTDS Setting
s,CN=BRANCH2,CN=Servers,CN=BRANCH2,CN=Sites,CN=Configuration,DC=contoso,DC=com
(2) modify CN=NTDS Site Settings,CN=BRANCH1,CN=Sites,CN=Configuration,DC=contoso,DC=com
    1> objectGUID: e29cdf7d-a95d-4fd0-affb-d1b182b3fd3f
    1> instanceType: 0x4 = ( IT_WRITE )
    1> interSiteTopologyGenerator: (GUID=2909511bd729604aa877aa64e027bb29);CN=NTDS Setting
s,CN=BRANCH1,CN=Servers,CN=BRANCH1,CN=Sites,CN=Configuration,DC=contoso,DC=com
New cookie written to file config (444 bytes)


C:\>_
```

**How to interpret the data**

- Prior to running the showchanges, a cookie file was created using the following syntax:

  **repadmin /showchanges . cn=configuration,dc=contoso,dc=com /cookie:config**

- Re ran **repadmin /showchanges** after some time against another domain controller, which not only displayed the changes but also updated the cookie file called **config**.

- Three objects have been changed. In our example, all the changes are pertaining to Intersite Topology Generators (ISTGs). Because the forest functional level is Windows 2000, we still update the ISTG Keep Alive stamp every 30 minutes.

- You could further apply filters to just target the partitions and objects of your interest.

**Display changes not replicated between two partners**
Syntax 2, shown earlier, is used here to display pending replication changes between partners.

**Example: Display pending replication changes (config partition) between two replication partners**
In this example (figure 3.10.1.1), we ran **repadmin /showchanges** to compare the destinations up-to-date vector with the source and determined that there are two outstanding changes for the configuration partition.

```
C:\>repadmin /showchanges branch3 9090b7ce-53a6-4a44-91bf-b50ed232be53 cn=configuration,dc
=contoso,dc=com
Building starting position from destination server branch3

Source Neighbor:
cn=configuration,dc=contoso,dc=com
==== INBOUND NEIGHBORS ======================================

cn=configuration,dc=contoso,dc=com
    HUB\BRANCH-HUB-BH via RPC
        DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53
        Address: 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com
        DC invocationID: 0cc4891a-4cca-49e8-bbc8-123de8d656ab
        DO_SCHEDULED_SYNCS WRITEABLE COMPRESS_CHANGES NO_CHANGE_NOTIFICATIONS
        USNs: 136025/OU, 136025/PU
        Last attempt @ 2005-02-15 23:47:53 was successful.

Destination's up-to-date vector:
076cd5dd-e25e-4897-acd2-7c8691621522 @ USN 908447
0cc4891a-4cca-49e8-bbc8-123de8d656ab @ USN 136064
195f73bb-5c65-4d14-a1fd-39d0adec3d0c @ USN 1166790
23a13880-47de-4602-ac24-212fa6c68fdd @ USN 63557
2f93bcb5-506a-4133-8a56-4bb3509c7ea4 @ USN 831399
65af87d6-d157-4e84-be28-551a6593844c @ USN 12750
84580b71-8ea1-48ed-9759-0ca5c5c54c2d @ USN 16480
8f05d838-29f0-4602-8918-5e203a0b0129 @ USN 12868
90fb7378-9726-430e-9469-c5ec7d230b37 @ USN 704936
9f4fd602-00d7-47ad-93e8-7af064384880 @ USN 131082
a5515acf-6747-494a-9c14-1f1b226c5de0 @ USN 26593
ae8ccd6c-c8ad-4c74-b8a5-4aa6c7dfc017 @ USN 63018
dc31d024-1bd6-4329-8d69-0c6ac26f8433 @ USN 68091
ecd55a8f-9f51-4a67-9e49-c1c02568eea9 @ USN 520166
fc3e85b3-3ced-405f-b232-7d27ae64ea94 @ USN 35675

==== SOURCE DC: 9090b7ce-53a6-4a44-91bf-b50ed232be53._msdcs.contoso.com ====

Objects returned: 2
(0) modify CN=NTDS Site Settings,CN=BRANCH2,CN=Sites,CN=Configuration,DC=contoso,DC=com
    1> objectGUID: 22cc7552-61c0-4020-9921-faa859e4d204
    1> instanceType: 0x4 = ( IT_WRITE )
    1> interSiteTopologyGenerator: <GUID=c2c4688127bf924f88ce3d502b596596>;CN=NTDS Setting
s,CN=BRANCH2,CN=Servers,CN=BRANCH2,CN=Sites,CN=Configuration,DC=contoso,DC=com
(1) modify CN=NTDS Site Settings,CN=BRANCH1,CN=Sites,CN=Configuration,DC=contoso,DC=com
    1> objectGUID: e29cdf7d-a95d-4fd0-affb-d1b182b3fd3f
    1> instanceType: 0x4 = ( IT_WRITE )
    1> interSiteTopologyGenerator: <GUID=2909511bd729604aa877aa64e027bb29>;CN=NTDS Setting
s,CN=BRANCH1,CN=Servers,CN=BRANCH1,CN=Sites,CN=Configuration,DC=contoso,DC=com


C:\>
```

**Example: Usage of a filter**

In the following example (figure 3.10.1.2), we applied a filter (/filter:"(objectclass=sitelink)") to just provide only changes occurred to the sitelink objectclass since the last successful replication.

```
Objects returned: 1
(0) modify CN=HUB-BRANCH3,CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=cont
oso,DC=com
    1> objectGUID: 87c79531-4f3c-4850-9fff-2fdf13594f54
    1> instanceType: 0x4 = ( IT_WRITE )
    1> options: 0x2 = ( 0x2 )


C:\>_
```

**Example: listing only the summary as opposed to individual changes**

In the following example (figure 3.10.1.3), the previous changes are listed as summary obtained by the **/statistics** switch.

```
********* Cumulative packet totals ************
Packets:                1
Objects:                2
Object Additions:       0
Object Modifications:   2
Object Deletions:       0
Object Moves:           0
Attributes:             6
Values:                 6
Dn-valued Attributes:   2
MaxDnVals on any attr:1
ObjectDn with maxattr:C
Attrname with maxattr:i
#dnvals 1-250   251-500 501-750 751-1000 1000+
add     0       0       0       0        0
mod     2       0       0       0        0
***********************************************

********* Grand total *************************
Packets:                1
Objects:                2
Object Additions:       0
Object Modifications:   2
Object Deletions:       0
Object Moves:           0
Attributes:             6
Values:                 6
Dn-valued Attributes:   2
MaxDnVals on any attr:1
ObjectDn with maxattr:C
Attrname with maxattr:i
#dnvals 1-250   251-500 501-750 751-1000 1000+
add     0       0       0       0        0
mod     2       0       0       0        0
***********************************************


C:\>_
```

# Usage of Repadmin When Troubleshooting Event ID 1311

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

By all means, this topic is not about how to troubleshoot events that have Event ID 1311. In this topic, we are attempting to expose the various usage of repadmin while troubleshooting 1311 in Windows 2000 domains based on Microsoft Knowledge Base (KB) article 307593, How to Troubleshoot Event ID 1311 Messages on a Windows 2000 Professional Domain (http://go.microsoft.com/fwlink/?LinkId=121799). Some or all of the repadmin subcommands used here may be used in Windows Server 2003 environments as well.

The KB article RESOLUTION section has the following action plan. This topic examines how to apply the various repadmin subcommands against each action plan. All of the repadmin subcommands listed in this topic have associated examples either in this section or elsewhere in this document.

| Resolution steps from the KB article | Action plan by using repadmin |
|---|---|
| Determine if the event ID 1311 messages are site-specific or forest-wide. | To determine the scope of event ID 1311 messages:<br>1. First, find all the Inter Site Topology Generators (ISTG) in the forest.<br>2. Then, examine the Directory Service logs of all the ISTG domain controllers in the forest.<br>To determine the ISTG's, use Repadmin /ISTG. |
| Determine if site link bridging is turned on and if the network is fully routed. | To determine this, use **repadmin /showattr** (Determine if site link bridging is turned on). |
| Verify that all of the sites are defined in site links. | Every site defined in Active Directory must be hosted or reside in a site link.<br>The **repadmin /showism** command (Verify inter-site cost matrix and orphaned sites) is useful for locating improperly configured sites. |
| Detect and remove preferred bridgeheads. | To search for preferred bridgehead servers use **repadmin /showattr** (Determine if site link bridging is turned on). |
| Resolve Active Directory replication failures in the forest | When you want to discover and troubleshoot replication failures, the following repadmin subcommands can be useful:<br>• **repadmin /replsummary** (Monitor Forest-Wide Replication)<br>• **repadmin /showrepl** (Display Replication Partners and Status of a Domain Controller)<br>• Repadmin /failcache<br>• **repadmin /removelingeringobjects** (Windows Server 2003 only)<br>• Repadmin /KCC |
| Determine if source servers are overloaded. | A domain controller that is overloaded with a large number of direct replication partners or a replication schedule that is overly aggressive can create a backlog in which some partners never receive changes from a hub domain controller. The following subcommands can be useful in this situation:<br>• **repadmin /showrepl** (Display Replication Partners and Status of a Domain Controller)<br>• Repadmin /queue<br>• **repadmin /showctx** (Open sessions with the domain controller) |

| | |
|---|---|
| Determine if site links are disjointed. | "Disjoint site links" is an Active Directory configuration in which the topology is broken into two or more parts in which some sites do not replicate because site definitions and site link definitions are incorrect. Disjoint site links are the most difficult improper configuration to troubleshoot. The following subcommands may be useful in this situation:<br>• Repadmin /querysites<br>• **repadmin /showconn** (Can I Look at My Connection Objects and Schedule Details?)<br>• Repadmin /KCC<br>• **repadmin /showrepl** (Display Replication Partners and Status of a Domain Controller) |
| Delete connections if the KCC is in "Keep Connection Mode." | If the Knowledge Consistency Checker (KCC) builds a different path around a site-to-site connection failure, but it retries the failing connection every 15 minutes because it is in "connection keeping mode," delete all broken connections and let the KCC rebuild them. Wait two times the longest replication schedule in the forest. |

**Determine if site link bridging is turned on**

Site link bridging is enabled in Active Directory if the following conditions are true:

- The Bridge all site links check box is selected for the IP protocol and the SMTP protocol in the Active Directory Sites and Services snap-in.
- The Options attribute for the IP protocol and the SMTP protocol is NULL or set to zero (0) for the following distinguished name (DN) paths:

  CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=root domain of forest

  CN=SMTP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=root domain of forest

```
C:\>repadmin /showattr branch3 ncobj:config: /filter:(objectcategory=intersitetran
DN: CN=SMTP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=contoso,DC=com
    1> name: SMTP
    1> options: 0x1 = ( NTDSTRANSPORT_OPT_IGNORE_SCHEDULES )
DN: CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=contoso,DC=com
    1> name: IP
    1> options: 0x0 = (  )

C:\>_
```

There are two values that we could set from the graphical user interface (GUI): **Ignore Schedules** and **Bridge all site links**. In our example (figure 3.11.1), the IP transport has **Bridge all site links** enabled and SMTP transport has both values selected.

The following table lists the various values that the options attribute take.

| Option value | Description |
|---|---|
| 0X0 | Only **Bridge all site links** is selected from the above |
| 0X1 | Both the values are selected |
| 0X2 | None selected |
| 0X3 | Only **Ignore schedules** is selected |

**Detect preferred bridgeheads**

Preferred bridgeheads are selected when the following condition is true:

**bridgeheadTransportList** attribute is set to either one of the following values or both values:

- CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=<*root domain of forest*>
- CN=SMTP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=<*root domain of forest*>

By using **repadmin /showattr**, we are setting the base at configuration partition and applying a filter for server **ObjectClass** and looking for all of the domain controllers that have this value set to use either IP or SMTP transports.

If the search returns any results, note the name of server in the distinguished name path in which the **bridgeheadTransportList** attribute is populated.

```
C:\>repadmin /showattr branch3 ncobj:config: /filter:(objectclass=server) /subtree
/atts:bridgeheadTransportList
DN: CN=ROOTDNS,CN=Servers,CN=HUB,CN=Sites,CN=Configuration,DC=contoso,DC=com
DN: CN=ROOTDC01,CN=Servers,CN=HUB,CN=Sites,CN=Configuration,DC=contoso,DC=com
    1> bridgeheadTransportList: CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configur
ation,DC=contoso,DC=com
DN: CN=BRANCH-HUB-BH,CN=Servers,CN=HUB,CN=Sites,CN=Configuration,DC=contoso,DC=com
DN: CN=BRANCH1,CN=Servers,CN=BRANCH1,CN=Sites,CN=Configuration,DC=contoso,DC=com
DN: CN=BRANCH2,CN=Servers,CN=BRANCH2,CN=Sites,CN=Configuration,DC=contoso,DC=com
DN: CN=BRANCH3,CN=Servers,CN=BRANCH3,CN=Sites,CN=Configuration,DC=contoso,DC=com


C:\>_
```

In the example in figure 3.11.2, ROOTDC01 is selected as a preferred Bridgehead for IP transport in site HUB.

**Verify inter-site cost matrix and orphaned sites**

**Repadmin /showism** displays intersite messaging routes calculated by the Intersite Messaging Service and is very useful for locating improperly configured sites. This operation cannot be executed remotely.

As the KCC runs through the progressions of analyzing intersite site links and connections, it must query the Intersite Messaging Service (ISM) service to retrieve data about the network configuration to make intelligent decisions about routing changes.

To display cost and frequency configurations of replication between sites, use the following command:

**Syntax**

```
repadmin /showism [<TransportDN>] [/verbose]
```

| Parameter | Description |
|---|---|
| <*TransportDN*> | Specifies whether the mail server is using SMTP or remote procedure call (RPC) to send messages. |
| **/verbose** | Lists detailed information. |

| **Note** |
|---|
| The **repadmin.exe /showism** cannot be executed against a remote domain controller. |

**Example: Display inter-site cost matrix**

```
C:\>repadmin /showism "CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,dc=contoso,dc=com"

==== TRANSPORT CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,dc=contoso,dc=com CONNECTIVITY INFORMATION FOR 5 SITES:
 ====

Site CN=BRANCH4,CN=Sites,CN=Configuration,DC=contoso,DC=com is not connected by this transport.
        0,    1,    2,    3,    4
Site(0) CN=BRANCH1,CN=Sites,CN=Configuration,DC=contoso,DC=com
    0:0:0, 200:30:1, 200:30:1, -1:0:0, 100:15:1
    All DCs in site CN=BRANCH1,CN=Sites,CN=Configuration,DC=contoso,DC=com (with trans & hosting NC) are bridgehead candidates.


Site(1) CN=BRANCH2,CN=Sites,CN=Configuration,DC=contoso,DC=com
    200:30:1, 0:0:0, 200:30:1, -1:0:0, 100:30:1
    All DCs in site CN=BRANCH2,CN=Sites,CN=Configuration,DC=contoso,DC=com (with trans & hosting NC) are bridgehead candidates.


Site(2) CN=BRANCH3,CN=Sites,CN=Configuration,DC=contoso,DC=com
    200:30:1, 200:30:1, 0:0:0, -1:0:0, 100:30:1
    All DCs in site CN=BRANCH3,CN=Sites,CN=Configuration,DC=contoso,DC=com (with trans & hosting NC) are bridgehead candidates.


Site(3) CN=BRANCH4,CN=Sites,CN=Configuration,DC=contoso,DC=com
    -1:0:0, -1:0:0, -1:0:0, 0:0:0, -1:0:0
    All DCs in site CN=BRANCH4,CN=Sites,CN=Configuration,DC=contoso,DC=com (with trans & hosting NC) are bridgehead candidates.


Site(4) CN=HUB,CN=Sites,CN=Configuration,DC=contoso,DC=com
    100:15:1, 100:30:1, 100:30:1, -1:0:0, 0:0:0
    All DCs in site CN=HUB,CN=Sites,CN=Configuration,DC=contoso,DC=com (with trans & hosting NC) are bridgehead candidates.


C:\>_
```

**How to interpret the data**

- Showism was used against the IP transport and hence the output is specific to IP.
- If a specific transport is not specified, the output will contain both IP and SMTP details.
- The numbers in an entry appear in the following order:
    - Cost: Replication interval: Options
- There are four key pieces of information:
    - Text regarding the status of bridgehead servers.
    - Total cost between two sites. The cost value indicates the preference for a network link for replicating directory information between sites.

- Frequency of replication in minutes between the two sites.
- Options for each replication link.

- In the example in figure 3.11.3, we have five sites and Bridge all site links is enabled, which means that site link transitivity is enabled. Therefore, if we see any "-1:0:0" entries for one or more covered Active Directory sites, we must ensure that the affected sites are listed in a site link. In this example, site Branch4 is not included in any site links and therefore disconnected from rest of the sites. Event 1311 will certainly occur here due to this configuration problem.

| Fields of interest | Definition |
|---|---|
| "0:0:0" | Each site matrix contains one "0:0:0" entry that refers to itself. |
| "200:30:1" | An entry that contains positive numbers for the cost value and replication interval value (for example, "200:30:1" or "100:15:1") indicates that the site connection is good. Specifically in our example for Site BRANCH1<br><br>Site(0) CN=BRANCH1,CN=Sites,CN=Configuration,DC=contoso,DC=com<br><br>0:0:0, 200:30:1, 200:30:1, -1:0:0, 100:15:1<br><br>• 200 stands for the cost to replicate from site(1) which is BRANCH2 that is an aggregate cost between two hops (100 + 100) because a direct replication link between the two sites does not exist.<br>• 30 is the replication interval that is common between the two branches<br>• 1 is the option on the site link which denotes “Change Notifications are enabled across the site link”<br><br>And so on the rest of the sites... |
| "-1:0:0" | A "-1:0:0" entry indicates that the site connection is not working. This occurs if one or more of the following conditions are true:<br>• Site is not included in a site link.<br>• Site does not host any domain controllers (this is known as an "uncovered" site).<br>• Replication protocol is not used. For example, if SMTP replication is not configured, the entries in the SMTP portion of the /SHOWISM matrix all appear as "-1:0:0". |

**Note**

- If site link bridging is enabled and the **repadmin /showism** command returns a site with a full complement of "-1:0:0" entries and one "0:0:0" entry is orphaned unless the site is uncovered (no domain controllers reside in that site).
- If site link bridging is disabled, "-1:0:0" entries are less meaningful. If this is the case, you must manually determine if each site is included in a site link. To do so, write down the list of sites and site links, and manually map each site to a site link.

**Repadmin /failcache**

**Repadmin /failcache** displays a list of replication failures that KCC is aware of. Run this command from the console of each ISTG domain controller in the forest to discover replication failures for bridgeheads in the site for that ISTG.

**Syntax**

```
repadmin /failcache <DC_LIST>
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |

**Example: Display replication failures that KCC is aware of**

The example in figure 3.11.4 shows sample output from the **repadmin /failcache** command.

```
C:\>repadmin /failcache rootdc01
==== KCC CONNECTION FAILURES ============================
    BRANCH2\STUDENT4\0ADEL:81ca9fb4-c63f-4849-a614-66b0ccd665ab (deleted DSA)
        DC object GUID: eae07634-e94b-47b6-a1c2-39c2cb44d221
        4006 consecutive failures since 2005-01-04 01:54:39.
        Last error: 8524 (0x214c):
            The DSA operation is unable to proceed because of a DNS lookup failure.
    BRANCH3\BRANCH3\0ADEL:70163a17-2f67-45ee-a7dd-ebdafb1c7f11 (deleted DSA)
        DC object GUID: 267e815a-2809-43fb-97a4-fc65616f0776
        4006 consecutive failures since 2005-01-04 01:54:39.
        Last error: 8524 (0x214c):
            The DSA operation is unable to proceed because of a DNS lookup failure.

==== KCC LINK FAILURES ==================================
    HUB\BRANCH-HUB-BH
        DC object GUID: 9090b7ce-53a6-4a44-91bf-b50ed232be53
        No Failures.
    BRANCH3\STUDENT3\0ADEL:2ed00423-01e1-4545-9d72-c2d3396e7f28 (deleted DSA)
        DC object GUID: 8332869b-4e65-477b-b8a4-c28ffbaa8810
        No Failures.
    BRANCH1\STUDENT1\0ADEL:04c59341-9d78-45c1-aca6-abf7ccb00d39 (deleted DSA)
        DC object GUID: 0aa9b456-301d-4183-9e38-c4369f2bb03f
        No Failures.

C:\>_
```

The output from the **repadmin /failcache** command is divided into two sections explained in the following table.

| | |
|---|---|
| KCC Link Failures | Lists errors for existing connection links. The ISTG domain controller imports showreps ("repsfroms") data for every bridgehead server in its site. However, the ISTG domain controller does not list errors. The link failure cache is emptied at the beginning of every KCC run and refilled during the course of the current run. |
| KCC Connection Failures | Lists unsuccessful attempts to build connection objects between domain controllers ("reps from" or "reps to"). When you run the **repadmin /failcache** command from the ISTG domain controller, it lists entries that are imported from bridgeheads in the site. At the beginning of each KCC run, the KCC examines each entry in the connection failure cache and tries to DsBind to the failing server. If the bind succeeds, the entry is removed. |

In the example in figure 3.11.4.1, the failures are a result of some topology changes from the past and would continue to exist due to the value of the **replTopologyStayOfExecution** attribute, which determines how long domain controller metadata is retained in Active Directory after a domain controller has been removed.

**Example: Output when there are no failures**

When there are no failures, the output should appears as it does in figure 3.11.4.2.

```
C:\>repadmin /failcache

repadmin running command /failcache against server localhost

==== KCC CONNECTION FAILURES ============================
(none)

==== KCC LINK FAILURES ==================================
(none)

C:\>_
```

**Note**

The **repadmin /failcache** command differs from the **repadmin /showrepl** command in two ways:

- The **repadmin /showrepl** command shows the naming context that is failing.
- The **repadmin /failcache** command does not.

**Repadmin /KCC**

**Repadmin /KCC** forces the KCC to recalculate replication topology for a specified domain controller. By default, this recalculation occurs every 15 minutes.

**Syntax**

```
repadmin /kcc <DC_LIST> [/async]
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp.** |
| **/async** | Specifies that replication will be asynchronous. This means that repadmin starts the replication event, but it does not expect an immediate response from the destination domain controller. Use this parameter to start the KCC and not wait for it to finish. |

**Example 1: Running the KCC on the local domain controller**

```
C:\>repadmin /kcc

repadmin running command /kcc against server localhost

Consistency check on localhost successful.

C:\>_
```

**Example 2: Running the KCC against the ISTG of the HUB site**

```
C:\>repadmin /kcc fsmo_istg:hub

repadmin running command /kcc against server rootdns.contoso.com

Consistency check on rootdns.contoso.com successful.

C:\>_
```

**Example 3: Running the KCC against all the global catalog servers in the forest**

```
C:\>repadmin /kcc GC:

repadmin running command /kcc against server 076cd5dd-e25e-4897-acd2-7c8691621522._msdcs.c
ontoso.com

Consistency check on 076cd5dd-e25e-4897-acd2-7c8691621522._msdcs.contoso.com successful.

repadmin running command /kcc against server 2a92f776-6c0f-4cb4-a111-f5dcd447af6c._msdcs.c
ontoso.com

Consistency check on 2a92f776-6c0f-4cb4-a111-f5dcd447af6c._msdcs.contoso.com successful.

C:\>_
```

**Example 4: Running the KCC against all the domain controllers in the BRANCH2 site**

```
C:\>repadmin /kcc site:branch2

repadmin running command /kcc against server 8168c4c2-bf27-4f92-88ce-3d502b596596._msdcs.c
ontoso.com

Consistency check on 8168c4c2-bf27-4f92-88ce-3d502b596596._msdcs.contoso.com successful.

C:\>_
```

**Repadmin /ISTG**

**Repadmin /ISTG** returns the server name of the ISTG server for a specified site.

**Syntax**

```
repadmin /istg <DC_LIST> [/verbose]
```

| Parameters | Descriptions |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |
| **/verbose** | Lists detailed information. |

**Example: Display ISTGs in my environment**

```
C:\>repadmin /istg /verbose

repadmin running command /istg against server localhost

Gathering topology from site BRANCH3 (BRANCH3.research.contoso.com):
                Site                 ISTG
      ==================   =================
               BRANCH3              BRANCH3
               BRANCH1              BRANCH1
                   HUB              ROOTDNS
               BRANCH2              BRANCH2

C:\>_
```

In the example in figure 3.11.6, the ISTGs are listed from the perspective of the local domain controller from which the command was run. It is important to note that this information may be different from the perspective of each domain controller, depending on the forest-wide Active Directory convergence time and replication status.

**Repadmin /querysites**

**Repadmin /querysites** use routing information to determine cost of a route from a specified site to another specified site or sites.

**Syntax**

```
repadmin /querysites <FromSiteRDN> <ToSite1RDN> <ToSite2RDN>...]
```

| Parameter | Description |
|---|---|
| <*FromSiteRDN*> | Specifies the relative distinguished name of the site from which the cost is calculated. |
| <*ToSite1RDN*> | Specifies the relative distinguished name of the site to which the cost is calculated. |

**Example 1: Display cost between BRANCH1 and HUB**

```
C:\>repadmin /querysites BRANCH1 HUB
From site: BRANCH1

To Site Name                                                    Cost
------------                                                    ----
HUB                                                             100

C:\>_
```

**Example 2: Display cost between BRANCH1 and BRANCH2**

Due to site link transitivity, the cost from BRANCH1 to BRANCH2 is aggregated by adding the cost from BRANCH1 to HUB (100) with the cost from HUB to BRANCH2 (100).

```
C:\>repadmin /querysites BRANCH1 BRANCH2
From site: BRANCH1

To Site Name                                                    Cost
------------                                                    ----
BRANCH2                                                         200

C:\>_
```

**Example 3: Display cost between BRANCH1 and Branch2**

Note that the relative distinguished name of the site is case sensitive and hence the error.

```
C:\>repadmin /querysites BRANCH1 Branch2
From site: BRANCH1

To Site Name                                                    Cost
------------                                                    ----
Branch2                                                         Error: 0x208d

C:\>_
```

**Note**

- The relative distinguished name of the site is case sensitive.
- The **repadmin /querysites** parameter does not allow the use of alternate credentials.

**Repadmin /queue**

**Repadmin /queue** displays tasks that are waiting in the replication queue.

**Syntax**

```
repadmin /queue <DC_LIST>
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |

**Example: Display the queue length against the local domain controller**

Under normal circumstances this list should always be empty and the command should be run outside of the replication window when troubleshooting domain controller overload was caused due to replication requests.

```
C:\>repadmin /queue

repadmin running command /queue against server localhost

Queue contains 0 items.

C:\>_
```

**Example: Queue contains one item**

```
repadmin running command /queue against server BRANCH-HUB-BH.research.contoso.com

Queue contains 1 items.
Current task began executing at 2005-02-14 12:24:21.
Task has been executing for -1 minutes, -12 seconds.

[36368] Enqueued 2005-02-14 12:24:21 at priority 590
    SYNC FROM SOURCE
    NC DC=DomainDnsZones,DC=research,DC=contoso,DC=com
    DC BRANCH3\BRANCH3
    DC object GUID 7e50725f-29c4-4d8a-80a2-c165e1c94e79
    DC transport addr 7e50725f-29c4-4d8a-80a2-c165e1c94e79._msdcs.contoso.com
    FORCE
```

**Repadmin /bridgeheads**

**Repadmin /bridgeheads** lists the bridgehead servers for a specified site.

**Syntax**

```
repadmin /bridgeheads [<DC_LIST>] [/verbose]
```

| Parameter | Description |
|---|---|

| | |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |
| **/verbose** | Lists detailed information. |

For clarity:

- The following example shows only bridgeheads only for the HUB site.
- The following example shows the normal and verbose modes to help compare them.
- "The RPC service is unavailable" status is abbreviated as **RPC.**
- "The operation completed successfully" status is abbreviated as **status.**

**Example 1: Repadmin /bridgeheads rootdns**

```
Bridgeheads for site HUB (rootdns.contoso.com):
Source Site    Local Bridge  Trns         Fail. Time    #    Status
===========    ============  ====     ==============   ===   ======
BRANCH2       BRANCH-HUB-BH   IP 2005-02-14 14:18:52    3    RPC.
Configuration research
BRANCH1       BRANCH-HUB-BH   IP             (never)    0    Success.
Configuration ForestDnsZones DomainDnsZones research
BRANCH3       BRANCH-HUB-BH   IP             (never)    0    Success.
Configuration DomainDnsZones ForestDnsZones research
```

**Example 2: Repadmin /bridgeheads rootdns /verbose**

```
Bridgeheads for site HUB (rootdns.contoso.com):
Source Site    Local Bridge  Trns         Fail. Time    #    Status
===========    ============  ====     ==============   ===   ======
BRANCH2        BRANCH-HUB-BH  IP  2005-02-14 14:18:52   3    RPC.
Naming Context    Attempt Time         Success Time  #Fail  Last Result
==============    ============         ============  =====  ===========
Configuration  2005-02-14 14:51:41  2005-02-14 14:18:51  3   RPC.
research       2005-02-14 14:53:15  2005-02-14 14:18:52  2   RPC.
Source Site    Local Bridge  Trns         Fail. Time    #    Status
===========    ============  ====     ==============   ===   ======
BRANCH1        BRANCH-HUB-BH  IP             (never)    0   Success
Naming Context    Attempt Time         Success Time  #Fail  Last Result
==============    ============         ============  =====  ===========
Configuration  2005-02-14 14:51:41  2005-02-14 14:51:41  0   Success.
ForestDnsZones 2005-02-14 14:52:37  2005-02-14 14:52:37  0   Success.
DomainDnsZones 2005-02-14 14:53:15  2005-02-14 14:53:15  0   Success.
research       2005-02-14 14:52:37  2005-02-14 14:52:37  0   Success.
Source Site    Local Bridge  Trns         Fail. Time    #    Status
===========    ============  ====     ==============   ===   ======
BRANCH3        BRANCH-HUB-BH  IP             (never)    0    Success.
Naming Context    Attempt Time         Success Time  #Fail  Last Result
==============    ============         ============  =====  ===========
Configuration  2005-02-14 14:51:42  2005-02-14 14:51:42  0   Success.
DomainDnsZones 2005-02-14 14:53:15  2005-02-14 14:53:15  0   Success.
ForestDnsZones 2005-02-14 14:52:37  2005-02-14 14:52:37  0   Success.
research       2005-02-14 14:53:15  2005-02-14 14:53:15  0   Success.
```

**How to interpret the data**

**Repadmin /bridgeheads** is run remotely against a domain controller in the HUB site and the output is the perspective of the topology for ROOTDNS. In these examples, we are seeing local bridgehead server **BRANCH-HUB-BH** is having replication problems with the remote bridgehead server in the **BRANCH2** site.

| Fields of interest | Explanation |
|---|---|
| Source Site | Source site from where the local bridge head (inbound) is pulling data. Remember replication is always inbound. |
| Local Bridge | Local Bridge head server for the site for which the tool is displaying results. In the example in figure 3.11.9.2, BRANCH-HUB-BH is the bridgehead server of the HUB site. |
| Trns | In the example in figure 3.11.9.2, the transport is IP. |
| Fail time | This is the last successful replication time. |
| # | Number of failures since the last successful replication time. |

| | |
|---|---|
| Status | Replication status. |
| Naming Context | Directory partition. Remember Bridgeheads are partition specific. |
| Attempt time | Last replication attempt time with the remote bridgehead. |
| Success time | Last successful replication time with the remote bridgehead. |
| #Fail | Number of attempts since the failure per partition. |
| Last result | Latest replication status. |

**Note**

Replication is performed for each partition. But sometimes we do not see the Schema partition listed in the previous example as a naming context (partition) and hence there are no bridgeheads listed. This is not a limitation of the tool; it has to do with the how information is stored in the connection object that is queried to determine the bridgehead. If you see the configuration partition in the output, it is implied that schema is also included because the KCC calculates the configuration and schema partitions to have the same replication topology.

**Repadmin /showmsg**

**Repadmin /showmsg** displays the error message for a given error number.

**Syntax**

```
repadmin /showmsg <Win32Error> | <DSEventID> /NTDSMSG}
```

| Parameter | Description |
|---|---|
| <*Win32Error*> | Returns a short description of the given Win32 error code. |
| <*DSEventID*> **/NTDSMSG** | Returns the actual event log text for the specified event ID. |

**Example: Display the error message for the win32error 1722 and DS event ID 1404**

```
C:\>repadmin /showmsg 1722
1722 = 0x6ba = "The RPC server is unavailable."

C:\>repadmin /showmsg 1404 /ntdsmsg
1404 = 0x57c = "The local domain controller is now the intersite topology genera
enerating and maintaining intersite replication topologies for this site."

C:\>_
```

**Repadmin /viewlist**

By default, this subcommand is used to displays a list of domain controllers. It could also be used to form an Lightweight Directory Access Protocol (LDAP) query to list only objects in the directory.

**Syntax**

```
repadmin /viewlist <DC_LIST> <OBJ_LIST>
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |
| <*OBJ_LIST*> | This parameter takes a distinguished name (DN) or a special keyword that expands into a DN. The keywords are:<br>• **Ncobj:config:** This keyword is the Configuration directory partition for the forest.<br>• **Ncobj:schema:** This keyword is the Schema directory partition for the forest.<br>• **Ncobj:domain:** This keyword is the domain partition DN of the home server.<br>• **Dsaobj**: This keyword is the NTDS settings object of the home server. |

**Example 1: Display all the DC's in the forest**

```
C:\>repadmin /viewlist *
DC_LIST[1] = rootdns.contoso.com
DC_LIST[2] = rootdc01.contoso.com
DC_LIST[3] = BRANCH-HUB-BH.research.contoso.com
DC_LIST[4] = BRANCH1.research.contoso.com
DC_LIST[5] = BRANCH2.research.contoso.com
DC_LIST[6] = BRANCH3.research.contoso.com


C:\>_
```

**Example 2: Display all the Group Policy objects in the domain directory partition for the domain of the domain controller that repadmin is running against**

```
C:\>repadmin /viewlist branch3 ncobj:domain: /subtree /filter:"(objectclass=groupPolicyContainer)"
DC_LIST[1] = branch3
    OBJ_LIST[1] = CN={6AC1786C-016F-11D2-945F-00C04fB984F9},CN=Policies,CN=System,DC=research,DC=contoso,DC=com
    OBJ_LIST[2] = CN={31B2F340-016D-11D2-945F-00C04FB984F9},CN=Policies,CN=System,DC=research,DC=contoso,DC=com


C:\>_
```

Note the usage of **OBJ_LIST** and **OBJ_LIST OPTIONS**. For details please refer to **repadmin /listhelp**.

**Open sessions with the domain controller**

The **repadmin /showctx** command displays a list of computers that have opened sessions with a specified domain controller.

**Syntax**

```
repadmin /showctx <DC_LIST> [/nocache]
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |
| **/nocache** | Specifies that globally unique identifiers (GUIDs) are left in hexadecimal form. By default, GUIDs are translated into strings. |

**Example: Show open sessions with a DSA**

```
C:\>repadmin /showctx branch1
Caching GUIDs.
..
3 open context handles.

HUB\BRANCH-HUB-BH @ 192.168.0.102 (PID 444) (Handle 0x270c238)
    bound, refs=1, last used 2005-02-16 15:11:47

NTDSAPI client @ 192.168.0.111 (PID 1492) (Handle 0x2765300)
    bound, refs=1, last used 2005-02-16 15:19:43

NTDSAPI client @ 192.168.0.113 (PID 328) (Handle 0x270c440)
    bound, refs=2, last used 2005-02-16 15:22:25


C:\>_
```

# Subcommands Not Covered Under the Previous Scenarios

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

This topic covers additional subcommands that you can use with repadmin.

**Display replication features**

The **repadmin /bind** command connects to, and displays the replication features for a directory partition on a domain controller.

**Syntax**

```
repadmin /bind <DC_LIST>
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |

**Example: Display replication features on the local domain controller, which is running Windows Server 2003**

Note that the LINKED_VALUE_REPLICATION is set to NO because the forest functional level is set to Windows 2000 instead of Windows Server 2003.

```
C:\>repadmin /bind

repadmin running command /bind against server localhost

Bind to localhost succeeded.
Extensions supported:
    BASE                            : Yes
    ASYNCREPL                       : Yes
    REMOVEAPI                       : Yes
    MOVEREQ_V2                      : Yes
    GETCHG_COMPRESS                 : Yes
    DCINFO_V1                       : Yes
    RESTORE_USN_OPTIMIZATION        : Yes
    KCC_EXECUTE                     : Yes
    ADDENTRY_V2                     : Yes
    LINKED_VALUE_REPLICATION        : No
    DCINFO_V2                       : Yes
    INSTANCE_TYPE_NOT_REQ_ON_MOD    : Yes
    CRYPTO_BIND                     : Yes
    GET_REPL_INFO                   : Yes
    STRONG_ENCRYPTION               : Yes
    DCINFO_VFFFFFFFF                : Yes
    TRANSITIVE_MEMBERSHIP           : Yes
    ADD_SID_HISTORY                 : Yes
    POST_BETA3                      : Yes
    GET_MEMBERSHIPS2                : Yes
    GETCHGREQ_V6 (WHISTLER PREVIEW) : Yes
    NONDOMAIN_NCS                   : Yes
    GETCHGREQ_V8 (WHISTLER BETA 1)  : Yes
    GETCHGREPLY_V5 (WHISTLER BETA 2) : Yes
    GETCHGREPLY_V6 (WHISTLER BETA 2) : Yes
    ADDENTRYREPLY_V3 (WHISTLER BETA 3): Yes
    GETCHGREPLY_V7 (WHISTLER BETA 3) : Yes
    VERIFY_OBJECT (WHISTLER BETA 3)  : Yes
    XPRESS_COMPRESSION              : Yes

Site GUID: 8c1976e6-f038-45a7-ba88-43f9462149eb
Repl epoch: 0

C:\>_
```

**Server object GUID (DSA GUID) & Database GUID**

The **repadmin /dsaguid** command returns a server name when given a globally unique identifier (GUID).

**Syntax**

```
repadmin /dsaguid <DC_LIST> <GUID>
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |

| | |
|---|---|
| *<GUID>* | Specifies the unique hexadecimal number that identifies the domain controller. The globally unique identifier (GUID) can be retrieved by using the **showreps** operation. |

**Example: Display the domain controller name when given a GUID**
Look at the usage of "." here for *<DC_LIST>*.

```
C:\>repadmin /dsaguid . 1b510929-29d7-4a60-a877-aa64e027bb29

repadmin running command /dsaguid against server localhost

Caching GUIDs.
..
"1b510929-29d7-4a60-a877-aa64e027bb29" = BRANCH1\BRANCH1.

C:\>_
```

Please refer to **repadmin /showrepl** for a detailed explanation and difference between DSA GUID and Database GUID.

**Certificates loaded on a domain controller**
The **repadmin /showcert** command displays the server certificates loaded on a specified domain controller.

**Syntax**
```
repadmin /showcert <DC_LIST>
```

| Parameter | Description |
|---|---|
| *<DC_LIST>* | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about *<DC_LIST>*, see **repadmin /listhelp.** |

**Retired Application partition GUIDs (signature)**
Each domain controller has a naming context signature list. The **repadmin /showncsig** command displays a list of the removed application directory partition GUIDs. An application directory partition can be configured to be held or not held on a particular domain controller by using **ntdsutil.**

**Syntax**
```
repadmin /showncsig <DC_LIST>
```

| Parameter | Description |
|---|---|
| *<DC_LIST>* | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about *<DC_LIST>*, see **repadmin /listhelp.** |

**Example: Display the recently retired ForestDnsZone application directory partition on the local domain controller**

```
C:\>repadmin /showncsig

repadmin running command /showncsig against server localhost

BRANCH3\BRANCH3

Current DC invocationID: ae8ccd6c-c8ad-4c74-b8a5-4aa6c7dfc017
Retired NC invocationID: ae8ccd6c-c8ad-4c74-b8a5-4aa6c7dfc017
NC 439b3266-cb29-42cd-a3b3-d524ec63a618 retired on 2005-02-14 17:03:35 at USN 61676

C:\>_
```

The following information is displayed in figure 3.12.4:

- Partition name
- InvocationID at the time of removal
- Highest update sequence number (USN) at the time of removal
- Date of removal

**Unanswered replication calls**
The **repadmin /showoutcalls** command displays calls that have not yet been answered, made by the specified domain controller to other domain controllers.

**Syntax**
```
repadmin /showoutcalls <DC_LIST>
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp.** |

**Example: Hub domain controller waiting for the request to be answered from a spoke domain controller**

```
repadmin running command /showoutcalls against server BRANCH-HUB-BH.research.contoso.
com

BRANCH-HUB-BH.research.contoso.com has 1 outgoing DRS RPC calls in progress:

Call type: DRS_CALL_GET_NC_CHANGES
Target server: 1b510929-29d7-4a60-a877-aa64e027bb29._msdcs.contoso.com
Handle info: bound 1 FromCache 1 InCache 1
Client thread id: 2112
Time call started: 2005-02-16 17:07:34
Call timeout: 5 minutes
Call duration: 0 minutes and 7 seconds
```

**showproxy**
Lists cross domain move proxy objects. When an object is moved to another domain, a marker is left in the old domain indicating that the object used to be there. This is called the proxy.

**Syntax1**
```
repadmin /showproxy <DC_LIST> <NamingContext> [matchstring]
```

**Syntax2**
```
repadmin /showproxy <DC_LIST> <ObjectDN> [matchstring] /movedobject
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp.** |
| <*NamingContext*> | Specifies the distinguished name of the directory partition on the source domain controller. |
| **matchstring** | Specifies the distinguished name of the object. |
| <*ObjectDN*> | Specifies a filter for the output. Type a string of characters that must be present in the distinguished name in order to display the object. |
| **/movedobject** | Displays a history of information from the original domain on a moved object after it has reached the new domain. |

**Retired Database GUIDs (signature)**
The **repadmin /showsig** command displays the retired InvocationIDs on a domain controller. A domain controller changes its InvocationID on being restored or when re-hosting an application partition.

**Syntax**
```
repadmin /showsig <DC_LIST>
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp.** |

**Example 1: Simple usage of no retired signatures**

```
C:\>repadmin /showsig rootdns
HUB\ROOTDNS

Current DC invocationID: 076cd5dd-e25e-4897-acd2-7c8691621522

No retired signatures.

C:\>_
```

**Example 2: Simple usage of retired signature**

```
C:\>repadmin /showsig

repadmin running command /showsig against server localhost

BRANCH3\BRANCH3

Current DC invocationID: 23a13880-47de-4602-ac24-212fa6c68fdd
ae8ccd6c-c8ad-4c74-b8a5-4aa6c7dfc017 retired on 2005-02-15 18:49:09 at USN 63018

C:\>_
```

**Convert directory service time to readable time**

The **repadmin /showtime** command converts a directory service time value to string format for both the local and the Coordinated Universal Time (UTC) time zones.

**Syntax**

```
repadmin /showtime <DSTimeValue>
```

| Parameter | Description |
|---|---|
| <*DSTimeValue*> | Specifies the time value that needs to be converted. |

| **Note** |
|---|
| With parameters omitted, **repadmin /showtime** displays the current system time in both the directory service format and string format. |

**Example 1: Usage with directory service time format**

```
C:\>repadmin /showtime 12753060350
12753060350 = 0x2f82441fe = 05-02-16 20:45.50 UTC = 2005-02-16 15:45:50 local

C:\>_
```

**Example 2: Current system time**

```
C:\>repadmin /showtime
12753060520 = 0x2f82442a8 = 05-02-16 20:48.40 UTC = 2005-02-16 15:48:40 local

C:\>
```

**Active Directory domains trusted by domain controller**

The **repadmin /showtrust** command lists all Active Directory domains (in the same forest) that are trusted by the specified domain controller's domain.

**Syntax**

```
repadmin /showtrust <DC_LIST>
```

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller, or a list of domain controllers, separated by a space. For details about <*DC_LIST*>, see **repadmin /listhelp**. |

**Example: Display Active Directory domains that are trusted by the domain of the local domain controller**

```
C:\>repadmin /showtrust

repadmin running command /showtrust against server localhost

Domain Trust Info:
    TRUSTED   : DC=contoso,DC=com
    TRUSTED   : DC=research,DC=contoso,DC=com

C:\>_
```

**Linked Distinguished Name values**

The **repadmin /showvalue** command is used to list only linked distinguished name values. Linked distinguished name values can also be obtained by the **repadmin /showobjmeta** subcommand with the **/linked** switch.

**Syntax**

```
repadmin /showvalue <DC_LIST> ObjectDN <AttributeName> <ValueDN> [/nocache]
```

| Parameter | Description |
|---|---|

| | |
|---|---|
| *<DC_LIST>* | Specifies the host name of a domain controller, or a list of domain controllers, separated by a space. For details about **DC_LIST**, see **repadmin /listhelp**. |
| *<ObjectDN>* | Specifies the distinguished name of the object. |
| *<AttributeName>* | Specifies a single attribute whose value you want to display. |
| *<ValueDN>* | Specifies the distinguished name of the attribute that is displayed. |
| **/nocache** | Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings. |

**Example: Display members of the Domain Admins group**

Note that showvalue lists value for only forward links. Backward links (such as **memberOf**) are not obtained.

```
C:\>repadmin /showvalue rootdns "CN=Domain Admins,CN=Users,DC=contoso,DC=com" /verbose
3 entries.
Type     Attribute      Last Mod Time                                Originating DC  Loc.USN Or
g.USN Ver
======= ============  =============                              ================= ======= ==
===== ===
         Distinguished Name
         ==============================
LEGACY          member
         CN=Administrator,CN=Users,DC=contoso,DC=com
LEGACY          member
         CN=Test13,OU=Test11,OU=Test1,DC=contoso,DC=com
LEGACY          member
         CN=lee,CN=Users,DC=contoso,DC=com


C:\>_
```

# Oldhelp

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Oldhelp displays a list of the operations that have been deprecated in the Windows Server 2003 version of repadmin.

**sync**
Starts a replication event for the specified directory partition between the source and destination domain controllers. The source universally unique identifier (UUID) can be determined when viewing the replication partners by using the showreps operation.

**Syntax**
```
repadmin /sync <NamingContext> <DestDC> <SourceDCUUID> [/force] [/async] [/full] [/addref] [/allsources]
```

| Parameter | Description |
|---|---|
| <*NamingContext*> | Specifies the distinguished name of the directory partition. |
| <*DestDC*> | Specifies the host name of the domain controller (Directory Server Agent) with which you want to replicate. |
| <*SourceDCUUID*> | Specifies the unique hexadecimal number that identifies the object whose changes will be listed. The objectGUID can be retrieved by using the showreps operation. |
| **/force** | Overrides the normal replication schedule |
| **/async** | Specifies that the replication will be asynchronous. This means that repadmin starts the replication event, but it does not expect an immediate response from the destination domain controller. Use this parameter when there are slow links between domain controllers. |
| **/full** | Forces a full replication of all objects from the destination domain controller. |
| **/addref** | Directs the source to check for a notification entry on the source. If the source does not have a notification entry for this destination, one is added. |
| /**allsources** | A given destination can have multiple sources for the same naming context. Directs the destination to sync with all sources instead of just one. |

**propcheck**
Compares properties of specified domain controllers to determine if they are up-to-date with each other. The source domain controller contains the original information that needs to be checked. The destination domain controller data will be compared to the source domain controller data.

**Syntax**
```
repadmin /propcheck <NamingContext> <OriginatingDCInvocationID> <OriginatingUSN> <DestDC>
```

| Parameter | Description |
|---|---|
| <*NamingContext*> | Specifies the distinguished name of the directory partition on the source domain controller. |
| <*OriginatingDCInvocationID*> | Specifies the unique hexadecimal number that identifies an object on a source domain controller. The InvocationID can be retrieved by using the showreps operation. |
| <*OriginatingUSN*> | Specifies the update sequence number (USN) for the object on the source domain controller. The USN is for the object whose InvocationID is already listed. |
| *DestDC* | Specifies the host name of the destination domain controller from which to enumerate the host domain controllers. |

**getchanges**
Displays changes from a specified directory partition or changes to a specified object. Syntax 1 saves changes to a directory partition. If this information is saved to a file the getchanges operation can be run again for comparison. Syntax 2 lists changes to a specified object.

**Syntax1**
```
repadmin /getchanges <NamingContext> <SourceDC> [/cookie: <File>] [/atts: <attribute1>,<attribute2>,...]
```

**Syntax2**
```
repadmin /getchanges <NamingContext> <DestDC> <SourceDCObjectGUID> [/verbose] [/statistics] [/noincremental]
[/objectsecurity] [/ancestors] [/atts: <attribute1>,<attribute2>,...] [/filter: <ldap filter>]
```

| Parameter | Description |
|---|---|
| <*NamingContext*> | Specifies the distinguished name of the directory partition. |
| <*SourceDC*> | Specifies the host name of the domain controller that hosts the directory partition whose changes you want to view. |
| **cookie:** <*File*> | Specifies a name for the file to which list changes are saved. |
| **atts:** <*attribute1*>,<*attribute2*> | Returns only the attributes specified. Separate each listed attribute with a comma. |
| <*DestDC*> | Specifies the host name of the destination domain controller from which to enumerate the host domain controllers. |
| <*SourceDCObjectGUID*> | Specifies the unique hexadecimal number that identifies the object whose changes will be listed. The objectGUID can be retrieved by using the showreps operation. |
| **/verbose** | Lists detailed information. |
| **/statistics** | Displays a summary of information about changes instead of a list of individual changes. |
| **/noincremental** | Returns changes in value change format, which lists current values for attributes as well as what attributes have been added or deleted. If not specified, changes are returned in attribute change format, which shows only the current value of the attribute. |
| **/objectsecurity** | Overrides the need for the Get Changes right to the directory partition. By default this right is needed to run the **/getchanges** parameter. However, only changes that the currently logged on user has the rights to view are displayed. |
| **/filter:** <*ldap filter*> | Returns only those changes that meet the filter requirements. |
| **/ancestors** | Returns changes in USN order |

| **Note** |
|---|
| The information from Syntax1 can be saved to a file for later comparison. |

**showreps**
Displays the replication partners for each directory partition on the specified domain controller. Helps the administrator build a visual representation of the replication topology and see the role of each domain controller in the replication process.

**Syntax**
```
repadmin /showreps <NamingContext> <DC> <SourceDCObjectGUID> [/verbose] [/nocache] [/repsto] [/conn] [/all]
```

| Parameter | Description |
|---|---|
| <*NamingContext*> | Specifies the distinguished name of the directory partition. |
| *DC* | Specifies the host name of the domain controller. |

| | |
|---|---|
| *<SourceDCObjectGUID>* | Specifies the unique hexadecimal number that identifies the object whose replication events will be listed. |
| **/verbose** | Lists detailed information. |
| **/nocache** | Specifies that globally unique identifier (GUIDs) are left in hexadecimal form. By default, GUIDs are translated into strings. |
| **[/repsto]** | Lists the domain controllers that pull replication information from the specified directory partition. |
| **/conn** | Displays the connection objects associated with each link. |
| **/all** | Displays all replication partners. |

**showvector**
Displays the highest USN for the specified domain controller. This information shows how up-to-date a replica is with its replication partners.

**Syntax**
```
repadmin /showvector <NamingContext> <DC> [/nocache] [/latency]
```

| Parameter | Description |
|---|---|
| *<NamingContext>* | Specifies the distinguished name of the directory partition. |
| *<DC>* | Specifies the host name of the domain controller. |
| **/nocache** | Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings. |
| **/latency** | Sorts the information by the time required to complete the replication. By default the information is sorted by USN. |

**showmeta**
Displays the replication metadata for a specified object stored in Active Directory such as attribute ID, version number, originating and local Update Sequence Number (USN), and originating server's GUID and Date and Time stamp. By comparing the replication metadata for the same object on different domain controllers, an administrator can determine whether replication has taken place.

**Syntax**
```
repadmin /showmeta <ObjectDN> <DC> [/nocache] [/linked]
```

| Parameter | Description |
|---|---|
| *<ObjectDN>* | Specifies the distinguished name of the object. |
| *<DC>* | Specifies the host name of the domain controller that hosts the object. |
| **/nocache** | Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings. |
| **/linked** | Displays metadata associated with, but not stored with the specified object. |

# Administer Passwords and Password Replication Policy for Read-Only Domain Controllers with Repadmin.exe

Published: March 15, 2010

Updated: March 15, 2010

Applies To: Windows Server 2008, Windows Server 2008 R2

This topic describes the following commands that were added to Repadmin.exe in Windows Server 2008 to manage passwords and Password Replication Policy (PRP) for read-only domain controllers (RODCs). RODCs are available in Windows Server 2008 and Windows Server 2008 R2.

Repadmin.exe is built into Windows Server 2008 and Windows Server 2008 R2. It is available if you have the Active Directory Domain Services (AD DS) server role or the Active Directory Lightweight Directory Services (AD LDS) server role installed. It is also available if you install the Active Directory Domain Services Tools that are part of the Remote Server Administration Tools (RSAT). For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

In Windows Server 2008 and Windows Server 2008 R2, you must run command-line-based tools from an elevated command prompt. To open an elevated Command Prompt using the credentials of a Domain Admin, click **Start**. In **Start Search**, type **runas /user:<domainName>\<domainAdminAccountUser> cmd**, and then press ENTER. Replace **<domainName>** with the domain name, and replace **<domainAdminUser>** with the name of a user account that is a member of the Domain Admins group in that domain.

- repadmin /prp
- repadmin /rodcpwdrepl

For more information about managing passwords and the PRP for RODCs, see Administering the Password Replication Policy (http://go.microsoft.com/fwlink/?LinkId=185778).

**repadmin /prp**
You can use this command to view or modify the PRP for an RODC. The PRP determines which account passwords are allowed to be cached on an RODC and which account are denied from being cached.

**Syntax**

```
Repadmin /prp [operation] RODC [additional arguments]
```

**Operations**
The **repadmin /prp** command can perform the following operations:

- Add
- Delete
- Move
- View

Additional arguments are available for each operation.

**Add**
Adds the specified security principal to the **msDS-RevealOnDemandGroup** attribute that is associated with the RODC. (This attribute is also known as the Allowed List.)

**Note**

You cannot use **repadmin /prp** commands to add an account to the Deny List or remove an account from the Deny List. To configure the Deny List, you can use the Active Directory Users and Computers snap-in or you can create a script. For example, if you want to deny members of the group RODC2Admins from caching passwords on RODC2, which is located in the Branch2 organizational unit (OU) of hq.cpandl.com, you can use the following script:

```
'The following items specify to Clear, Update, Append, or Delete a property of an
Active Directory object
Const ADS_PROPERTY_CLEAR = 1
Const ADS_PROPERTY_UPDATE = 2
Const ADS_PROPERTY_APPEND = 3
Const ADS_PROPERTY_DELETE = 4

Const ATT = "msDS-NeverRevealGroup"
'The setting for ATT determines which list will be modified
'msDS-AuthenticatedToAccountlist is for the authenticated to or Auth2 list
'msDS-RevealedList is for the password revealed or cached list
'msDS-RevealOnDemandGroup is for the allowed to authenticate list
'msDS-NeverRevealGroup is for the denied from authenticating list

'PRPObj defines the object that needs to be modified in the PRP list
PRPObj = "CN=RODC2Admins,OU=Branch2,DC=hq,DC=cpandl,DC=com"

'RODCObj defines the RODC for which the PRP should be modified
RODCObj = "LDAP://CN=RODC2,OU=Domain Controllers,DC=hq,DC=cpandl,DC=com"

'Sets the object to modify based on the LDAP path set in RODCObj
Set objComputer = GetObject(RODCObj)

'Implements the change, which, depending on the word after ADS_PROPERTY_, is a CLEAR,
UPDATE, APPEND, or DELETE operation
objComputer.PutEx ADS_PROPERTY_APPEND, ATT, Array(PRPObj)
objComputer.SetInfo

'Confirms that the modification has taken place (this is optional)
wscript.echo "Modified list attributes for object " & PRPObj

'Closes the script
wscript.quit(0)
```

**Syntax**

```
repadmin /prp add <RODC> allow <PRINCIPAL>
```

**Additional parameters**

| Parameter | Definition |
|---|---|
| <RODC> | Specifies the host name of the RODC. You can specify the single-label host name or the fully qualified domain name (FQDN). In addition, you can use an asterisk (*) as a wildcard character to specify multiple RODCs in one domain. |
| <PRINCIPAL> | Specifies the name of the security principal that you want to add to the Allowed List. |

**Delete**

Deletes one or more specified security principals from the **msDS-AuthenticatedToAccountList** attribute or from the **msDS-RevealOnDemandGroup** attribute that is associated with the RODC. (The **AuthenticatedToAccountList** attribute is also known as the Authenticated to List, and the **msDS-RevealOnDemandGroup** attribute is also known as the Allowed List.)

**Syntax**

```
repadmin /prp delete <RODC> allow {<PRINCIPAL>|/all}
repadmin /prp delete <RODC> auth2 /all
```

**Additional parameters**

| Parameter | Definition |
|---|---|

| | |
|---|---|
| <RODC> | Specifies the host name of the RODC. You can specify the single-label host name or the FQDN. In addition, you can use an asterisk (*) as a wildcard character to specify multiple RODCs in one domain. |
| <PRINCIPAL> | Specifies the name of the security principal that you want to delete from the Allowed List. Specify **/all** to have the operation delete all security principals. |
| /all | Specifies all security principals. You cannot delete only one security principal from the **msDS-AuthenticatedToAccountList** attribute. |

**Move**

Moves all the security principals from the **msDS-AuthenticatedToAccountList** attribute to the specified group. If the group does not exist, this command creates the group. If necessary, this command also adds the group to the **msDS-RevealOnDemandGroup** attribute of the RODC. (The **msDS-AuthenticatedToAccountList** attribute is also known as the Authenticated To List, and the **msDS-RevealOnDemandGroup** attribute is also known as the Allowed List.)

**Syntax**

```
repadmin /prp move <RODC> <Group> [/noauth2cleanup] [/users_only | /comps_only]
```

**Additional parameters**

| Parameter | Definition |
|---|---|
| <RODC> | Specifies the host name of the RODC. For this operation, you can specify the single-label host name or the FQDN. |
| <Group> | Specifies the name of the security group to which you want to move the security principals. If the security group does not exist, this command creates the security group in the built-in **Users** container. You can specify the name of the security group but not the distinguished name. |
| /noauth2cleanup | Retains the list of security principals in the **msDS-AuthenticatedToAccountList** attribute after the Move operation is complete. By default, the **msDS-AuthenticatedToAccountList** attribute is cleared. |
| /users_only | Moves only user accounts from the **msDS-AuthenticatedToAccountList** attribute to the specified group. The group is then added to the **msDS-RevealOnDemandGroup** attribute. |
| /comps_only | Moves only computer accounts from the **msDS-AuthenticatedToAccountList** attribute to the specified group. The group is then added to the **msDS-RevealOnDemandGroup** attribute. |

**View**

Displays the security principals in the specified list or displays the current PRP setting (allowed or denied) for a specified user.

**Syntax**

```
repadmin /prp view <RODC> {<List_Name>|<User>}
```

**Additional parameters**

| Parameter | Definition |
|---|---|
| <RODC> | Specifies the host name of the RODC. You can specify the single-label host name or the FQDN. In addition, you can use an asterisk (*) as a wildcard character to specify multiple RODCs in one domain. |
| <List_Name> | Specifies all the security principals that are in the list that you want to view. The valid list names are as follows:<br>• auth2: The list of security principals that the RODC has authenticated.<br>• reveal: The list of security principals for which the RODC has cached passwords.<br>• allow: The list of security principals in the **msDS-RevealOnDemandGroup** attribute. The RODC can cache passwords for this list of security principals only.<br>• deny: The list of security principals in the **msDS-NeverRevealGroup** attribute. The RODC cannot cache passwords for any security principals in this list. |

| | |
|---|---|
| <User> | Specifies the effective PRP setting (allowed or denied) for the specified user. You can specify the user name only or the distinguished name. |

**Example 1: View the PRP of an RODC**
The following examples show how to view the accounts that are configured in the PRP that applies to an RODC with the host name RODC2 in the domain hq.cpandl.com.

To view the accounts that are allowed to have their passwords cached on the RODC, use the following command:

```
repadmin /prp view rodc2.hq.cpandl.com allow
```

To view the accounts that are denied from having their passwords cached on the RODC (also known as the Deny list), use the following command:

```
repadmin /prp view rodc2.hq.cpandl.com deny
```

**Example 2: View accounts that an RODC has authenticated**
To review the list of authenticated accounts for RODC2 in the hq.cpandl.com domain, use the following command:

```
repadmin /prp view rodc2.hq.cpandl.com auth2
```

**Example 3: Clear the list of authenticated accounts**
Note that this command does not actually remove account passwords from an RODC. It only deletes the list of those accounts.

There is no mechanism to erase passwords after they are cached on an RODC. If you want to clear a password that is stored on an RODC, reset the password in the hub site. This way, the password that is cached in the branch will no longer be valid for accessing any resources in the hub site or other branches. In the branch that contains the RODC on which the password may have been compromised, the password will still be valid for authentication purposes until the next replication cycle, at which time its value that is stored on the RODC will be changed to Null. The new password will be cached only after the user authenticates with it—or the new password is prepopulated on the RODC—and if the PRP has not been changed.

To clear the list of authenticated accounts for RODC2, use the following command:

```
repadmin /prp delete rodc2 auth2 /all
```

**Example 4: Configure the PRP**
To add an account named RODC2users from a top-level OU named West in the domain hq.cpandl.com to the Allowed List (or to remove it from the Allowed List) for an RODC computer with a hostname of RODC2, use one of the following commands:

**Note**

To find the Lightweight Directory Access Protocol (LDAP) distinguished name of a directory object from the command line, you can use the **dsquery** command. For example, if you want to find the distinguished name of a group that has "RODC" as part of its name from a computer in the local domain, you can run the command **dsquery group –name *RODC*.** The asterisks around "RODC" indicate that any number of characters can come before or after the letters RODC. If instead you want to find the distinguished name of a computer or user, substitute either the word **computer** or the word **user** (respectively) for the word **group** in the command. For more information about **dsquery** command syntax, see Dsquery (http://go.microsoft.com/fwlink/?LinkId=120196).

To allow the account RODC2users to be cached on RODC2, use the following command:

```
repadmin /prp add rodc2.hq.cpandl.com allow
cn=RODC2users,ou=west,dc=hq,dc=cpandl,dc=com
```

To remove the account from the Allowed List, use the following command:

```
repadmin /prp delete rodc2.hq.cpandl.com allow
cn=RODC2users,ou=west,dc=hq,dc=cpandl,dc=com
```

**Example 5: Move accounts that an RODC has authenticated to the Allowed RODC Password Replication Policy Group**
To move the current list of only the users from RODC2 to the Allowed List, use the following command:

```
Repadmin /prp move rodc2 /users_only
```

**Note**

| |
|---|
| You cannot selectively move entries from the Auth2 list to the Allowed List by using the **repadmin /prp move** command. However, when you have created an appropriate group, you can use Active Directory Users and Computers, Dsadd, and similar tools to add users or computers to that group. |

**Example 6: View accounts with cached passwords on an RODC**

To see the accounts with cached passwords on an RODC with the host name RODC2 in the domain contoso.com, use the following command:

```
repadmin /prp view rodc2.contoso.com reveal
```

| **Important** |
|---|
| If you have a large number of accounts cached, the **repadmin /prp view <hostname> reveal** command might return only a subset of the accounts. For more information, see Repadmin /PRP might return only a subset of accounts (http://go.microsoft.com/fwlink/?LinkId=185775). |

**repadmin /rodcpwdrepl**

Triggers replication of passwords for the specified users from the source (Hub site domain controllers) to one or more RODCs.

For each destination RODC, the ability to cache the user's password is evaluated before the operation succeeds. In other words, the specified user must be in the Allowed RODC Password Replication Group and not be a member of the Denied RODC Password Replication Group for the destination RODC. You can specify passwords for multiple users, but if a user's password is not allowed to be cached for a destination RODC, the request for that specific user and destination RODC will fail.

**Syntax**

```
Repadmin /rodcpwdrepl <hostnameRODC> <hostnameWDC> <User1LdapPath> <Computer1LdapPath> <UserNLdapPath>
<ComputerNLdapPath>
```

| Parameters | Definitions |
|---|---|
| <hostnameRODC> | The host name or FQDN of the target RODC's password cache that you want to prepopulate. If you are running the command from outside the target domain, use the FQDN. |
| <hostnameWDC> | The host name or FQDN of the writable domain controller that is the replication partner of the RODC. If you are running the command from outside the target domain, use the FQDN. |
| <User1LdapPath> | The LDAP distinguished name of a user account password that you want to prepopulate. |
| <Computer1LdapPath> | The LDAP distinguished name of a computer account whose password that you want to populate. You must add the computer accounts of the users or they will not be able to log on. |
| <UserNLdapPath> | The LDAP distinguished name of another user account password that you want to populate. |
| <ComputerNLdapPath> | The LDAP distinguished name of another computer whose account password you want to prepopulate. You must add the computer accounts of the users or they will not be able to log on. |

**Example**

The following command prepopulates the password cache for an RODC named RODC2 in the domain hq.cpandl.com, using the writeable domain controller named WS2008A to transfer the passwords for a user account for Mike Danseglio (MikeDan) and his computer named MDVista1. The MikeDan account is in a top-level OU named B1 Users, and the MDVista1 account is in the default Computers container.

```
repadmin /rodcpwdrepl rodc2.hq.cpandl.com ws2008a.hq.cpandl.com "cn=mikedan,ou=b1
users,dc=hq,dc=cpandl,DC=com" cn=mdvista1,cn=Computers,dc=hq,dc=cpandl,dc=com
```

# Repadmin for Experts

Updated: July 10, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

The previous topics in this guide have looked at how an administrator can use **repadmin** to view the replication topology (sometimes referred to as **Reps-From** and **Reps-To**) as seen from the perspective of each domain controller, monitor forest-wide replication, diagnose replication problems, and perform miscellaneous tasks.

The following sections are used for advanced operations only. These commands have the potential to break your Active Directory installation, and they should be used only under the expert guidance of Microsoft Customer Support Service representative or engineer.

**Add, Modify, or Delete replication links**

During normal operation, the Knowledge Consistency Checker (KCC) automatically manages the replication topology for each naming context held on domain controllers.

Although in normal practice this should not be necessary, repadmin can be used to manually create the replication topology. This topology would be temporary in nature by default and would last until the next time the KCC is run. So we need to engage these steps only during troubleshooting issues related to Active Directory replication.

| **Note** |
| --- |
| During the normal course of operations, there is no requirement for manual creation of the replication topology. Incorrect use of this tool may adversely impact the replication topology. |

**Syntax**

```
Repadmin /add <Naming Context> <Dest DC> <Source DC> [/asyncrep] [/syncdisable] [/dsadn:< Source DC DN>]
[/transportdn:< Transport DN>] [/mail] [/async] [/readonly]

Repadmin /mod <Naming Context> <Dest DC> <Source GUID> [/readonly] [/srcdsaaddr:< dns address>]
[/transportdn:< Transport DN>] [+nbrflagoption] [-nbrflagoption]

Repadmin /delete <Naming Context> <Dest DC> [<Source DC Address>] [/localonly] [/nosource] [/async]
```

The following table lists the purpose for each of the subcommands.

| Subcommand | Purpose |
| --- | --- |
| **add** | The **add** command will create a **RepsFrom** attribute on the destination domain controller for the specified naming context and initiate a replication request. During a normal replication cycle, the destination domain controller will request updates from the source domain controller. |
| **mod** | The **mod** command will modify the **RepsFrom** attribute on the destination domain controller for the specified naming context and initiate a replication request. During a normal replication cycle, the destination domain controller will request updates from the source domain controller. |
| **delete** | The **delete** command will remove a **RepsFrom** attribute on the destination domain controller for the specified naming context. |

The following table lists the parameters that can be used with the subcommands.

| Parameter | Description |
| --- | --- |
| *<Naming Context>* | Specifies the distinguished name of the directory partition. |
| *<Dest DC>* | Domain controller to which the link is created. |
| *<Source DC>* | Domain controller from which to source the partition. |

| | |
|---|---|
| **asyncrep** | Queue the replication event, but do not wait for the replication to complete before you return control to the user. |
| **syncdisable** | Add the **RepsFrom** attribute but do not participate in the replication cycle. To perform replication between the destination and source domain controllers, **repadmin /sync /force** must be used. |
| **/dsadn:**<<*Source DC DN*> | |
| **transportdn** | The distinguished name of the Inter Site Message transport, only used for mail-based replication. |
| **mail** | specify that the replication is mail-based, therefore requires the **/transportdn** option. |
| **async** | Queue the **add/delete** operation without interrupting the current replication cycle and return control to the user. |
| **readonly** | Specify that the partition is read-only. |
| **/srcdsaaddr:**<*dns address*> | |
| **nbrflagoption** | |
| **localonly** | Do not delete the corresponding **RepsTo** attribute on the source Directory System Agent (DSA). |
| **nosource** | When you remove a read-only naming context such as the global catalog, the associated data stored in the directory is removed in blocks of 500 objects. This allows the **/delete** command to be re-executed without having to specify the Source DSA to remove the remaining objects. |

When you create temporary replication links between replication partners, the process could fail if the KCC starts while you are performing the procedure. The KCC will delete any replication links for which no corresponding connection object exists.

Because these commands can take a very long time to complete as they trigger the replication of the corresponding naming context, it is important to ensure that KCC do not disturb the process. This is where you would use +DISABLE_NTDSCONN_XLATE which effectively disables capability for the KCC to translate connection objects to replication links.

**Add, Modify, or Delete outbound replication partners**

Similar to inbound replication (**Reps-From**) partners, outbound replication (**Reps-To**) partners are instantiated from connection objects by a process called "Connection Translation."

Both **Reps-From** and **Reps-To** attributes are for each partition and they are not replicated. **Reps-To** is only needed when the destination requires the source to notify him that there is a change in the partition at the source, and the destination should synchronize. Because **Reps-To** attributes are used for notification, if the destination has a **Reps-From** marked NO_NOTIFY, then the source will not have a **Reps-To**.

Depending on the underlying operating system, sometimes you might see outbound partners lingering. While Windows Server 2003 takes care of this, Windows 2000 would need some help cleaning out lingering outbound partners.

**Syntax**

```
Repadmin /addrepsto <Naming Context> <DC> <Reps-To DC> <Reps-To DC GUID> Repadmin /updrepsto <Naming Context>
<DC> <Reps-To DC> <Reps-To DC GUID> Repadmin /delrepsto <Naming Context> <DC> <Reps-To DC> <Reps-To DC GUID>
```

The following table lists the purpose for each of the subcommands.

| Subcommand | Purpose |
|---|---|
| **addrepsto** | This will create a **Reps-To** attribute on the domain controller for the specified naming context. Ordinarily there is no requirement to perform this command as the KCC will automatically create the **Reps-To** attributes on destination DSAs based on other DSAs **Reps-From** entries. |
| **updrepsto** | This will update the **Reps-To** attribute on the domain controller for the specified naming context. More specifically it updates the network address used by the source DSA to contact the destination DSA. |

| | |
|---|---|
| **delrepsto** | Delrepsto deletes the **Reps-To** attribute on the domain controller for the specified naming context. |

The following table lists the parameters that can be used with the subcommands.

| Parameter | Description |
|---|---|
| <*Naming Context*> | Specifies the distinguished name of the directory partition. |
| <*DC*> | The domain controller on which the **Reps-To** attribute is modified. |
| <*Reps-To DC*> | Outbound replication partner. |
| <*Reps-To DC GUID*> | DSA globally unique identifier (GUID) of outbound replication partner. |

**Hosting and unhosting read-only partitions**

Hosting and unhosting global catalog partitions is convenient, especially when you want to ensure a faster global catalog removal process. As noted in the following table, these subcommands will also facilitate removal of lingering objects from Active Directory.

| | |
|---|---|
| Global catalog removal process | In Windows 2000 versions earlier than Service Pack 4 (SP4), when the IS_GC bit is turned off, the KCC deletes the read-only objects at a rate of only 500 for each time the KCC runs, which allows a maximum of 2000 object removals for each hour. This presents some challenges in large environments. In order to make the global catalog removal faster, you could potentially remove one partition at a time by using the unhost subcommand. |
| Lingering Objects | A lingering object is an object that is present on one replica, but on another replica it has been deleted and removed from the directory by the garbage collection process.<br><br>When lingering object exists only in one or more read-only naming contexts (global catalog), it makes it all the more difficult to delete the object. Clearing the IS_GC bit may not always be appropriate, because it removes all read-only naming contexts from the global catalog server.<br><br>Unhosting and rehosting a read-only naming context is therefore sometimes considered to be a good solution, especially because you could specify the source to be a good replica that does not contain lingering objects. |

**Syntax**

```
Repadmin /rehost <DC_LIST> <Naming Context> <Good Source DC Address> [/application] Repadmin /unhost
<DC_LIST> <Naming Context> Repadmin /removesources <DC_LIST> <Naming Context>
```

The following table lists the purpose for each of the subcommands.

| Subcommand | Purpose |
|---|---|
| **rehost** | Add a specific read-only partition to a global catalog server. |
| **unhost** | Remove a specific read-only partition from a global catalog server. |
| **removesources** | Removes all replication links for a given naming context. This does not delete the connection objects, so the KCC will build new links on it regular cycle as required. |

The following table lists the parameters that can be used with the subcommands.

| Parameter | Description |
|---|---|
| <*DC_LIST*> | Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. For details about <*DC_LIST*>, see **repadmin /listhelp**. |
| <*Naming Context*> | Specifies the distinguished name of the directory partition. |

| | |
|---|---|
| *<Good Source DC Address>* | Specify the source domain controller. |
| **/application** | Application directory partition |

**Detecting and removing lingering objects**

There are multiple methods that are available to detect or remove lingering objects from Active Directory. This depends on the operating system version that the domain controller is running. Repadmin could be used to detect or remove lingering objects from a directory partition when the source and destination domain controllers are running Windows Server 2003 and therefore the scope here is limited to the following:

- Introduction to lingering objects
- Repadmin usage in Windows Server 2003

A lingering object is an object that is present on one replica, but on another replica it has been deleted and removed from the directory by the garbage collection process.

This condition can occur for a variety of reasons including:

- Prolonged misconfigurations (such as those that cause event ID 1311 messages)
- Prolonged errors in name resolution, authentication or the replication engine that block inbound replication.
- Bringing a domain controller online after it has been offline for a period greater than the **TombStone Lifetime** (TSL).
- Advancing system time or reducing TSL values in an attempt to accelerate garbage collection before end-to-end replication has taken place for all naming contexts in the forest.

Symptoms that you may have lingering objects:

- Active Directory replication is prevented from occurring.
- A user account that no longer exists still appears in the Global Address list for E-mail clients.
- A universal group that no longer exists still appears in a user's access token.
- E-mail messages cannot be delivered due to duplicate e-mail address on two different user objects.

Regardless of the reason, a deleted object can remain on a domain controller in either of the following circumstances:

- A domain controller goes offline immediately prior to the deletion of an object on another domain controller, and remains offline for a period that exceeds the tombstone lifetime.
- A domain controller goes offline immediately following the deletion of an object on another domain controller but prior to receiving replication of the tombstone, and remains offline for a period that exceeds the tombstone lifetime.

What to do with a lingering object?

Determining what to do with a lingering object depends on whether or not it was intended.

| Action | Explanation |
|---|---|
| Unintended | Use repadmin to delete the lingering object on a domain controller that is running Windows Server 2003. |

| | |
|---|---|
| Intended | Change the replication consistency on the inbound domain controller (DC). The object will be re-animated on this DC. See strict and loose replication consistency below |

**Strict and loose replication consistency**

If the attributes of a lingering object never change, the object is never considered for replication. However, if an attribute changes, the attribute is considered for outbound replication. The problem with an attribute update for a lingering object is that the receiving domain controller does not hold the object for the attribute being replicated. An update cannot be performed because the entire object does not exist on the receiving domain controller. What happens next depends on the replication consistency set on the domain controller.

| Replication consistency | Explanation |
|---|---|
| Loose | When replication consistency is set to loose, the receiving domain controller detects that it does not have the object for the attribute that is being replicated. The inbound partner requests the entire object from the outbound partner, and reanimates the object on its copy of the directory. The same process repeats on all domain controllers that do not have a copy of the object. This mechanism can be used to cause lingering objects to "reanimate" across the entire forest. If a lingering object is discovered and its presence is intended, then perform any update to the object. As long as replication consistency is set to loose on all domain controllers, the object will be reanimated as it replicates around the forest. "Loose replication consistency" is the default for Windows 2000 domain controllers, with the exception of domain controllers that have the MS01-044 security rollup package installed. For more information about the MS01-044 security rollup package, see article 297860 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=122508). |
| Strict | The default behavior for domain controllers that run Windows Server 2003 (and domain controllers that are upgraded from Windows NT 4.0) is to block inbound replication for each naming context when a domain controller receives an update to an object that it does not have. Replication is halted in the naming context for the object until the lingering object is removed or the replication mode is set to "loose." |

**Storage for Consistency Setting**

The setting for replication consistency is in the registry on each domain controller.

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

Entry name: Strict Replication Consistency

Data type: REG_DWORD

Values: 1 for enabled; 0 for disabled

Default: 1 (enabled)

**Note**

There was a post-SP2 hotfix (also included in the security rollup package from November 2001) that used a different registry value. A setting of 0 will not recreate the missing object (strict), and a setting of 1 will create the missing object. This value is only needed with the November version of the hotfix.

- Value Name: Correct Missing Objects
- Data type: REG_DWORD
- Value data: 1

The repadmin /removelingeringobjects command does the following:

- Designates an up-to-date domain controller as the authority.
- Compares the Active Directory database objects on the authoritative server with the objects that are on the suspected domain controller that contains the lingering objects.

- With **/advisory_mode**, the subcommand logs the potential deletions to the Directory Service log.
- Without **/advisory_mode**, the subcommand removes the lingering objects.

**Syntax**

```
Repadmin /removelingeringobjects <Dest_DC_LIST> <Source DC GUID> <NC> [/ADVISORY_MODE]
```

| Parameter | Description |
|---|---|
| *<Dest_DC_LIST>* | The domain controller that is suspected to have lingering objects. |
| *<Source DC GUID>* | Source domain controller GUID used to compare with the suspected domain controller. |
| *<NC>* | Specifies the distinguished name of the directory partition. |
| **/ADVISORY_MODE** | Read-only mode. |

| **Note** |
|---|
| During lingering object removal, Event ID 1937 is logged to the Directory Service log. This information includes the source domain controller, the objects that are removed, and a total count of all the objects that are removed. |

**Advanced domain controller options**

By using the option subcommand, we could change the **options** attribute stored on the NTDS Settings Object. The **options** attribute determines the following behaviors on a domain controller:

- Global catalog installation and removal
- Enable or disable inbound or outbound replication
- Disable connection translation

Note that disabling inbound or outbound replication is specific to the domain controller where you target the operation. So this does not disable intrasite or intersite replication. It just disables Active Directory replication for that domain controller. If the domain controller happens to be the bridgehead server and the Intersite Topology Generator (ISTG) is disabled, then effectively intersite replication to and from that site is disabled.

**Syntax**

```
Repadmin /options <DC> [{+|-} IS_GC] [{+|-} DISABLE_INBOUND_REPL] [{+|- DISABLE_OUTBOUND_REPL] [{+|-} DISAB
LE_NTDSCONN_XLATE]
```

+|- turns on or off the associated parameter.

| Parameter | Description |
|---|---|
| *<DC>* | Domain controller |
| **IS_GC** | DSA is a global catalog server. |
| **DISABLE_INBOUND_REPL** | Disables inbound replication. |
| **DISABLE_OUTBOUND_REPL** | Disables outbound replication. |
| **DISAB LE_NTDSCONN_XLATE** | Turns off the capability of the KCC to translate connection objects to replication links. |

The following table lists the possible values for the options attribute.

| Value | Description |
|---|---|

| | |
|---|---|
| 1 | Global catalog server |
| 2 | Disable inbound replication |
| 3 | 2 + 1 |
| 4 | Disable outbound replication |
| 5 | 4 + 1 |
| 6 | 4 + 2 |
| 7 | 4 + 2 + 1 |
| 8 | Disable connection translation |

The following table lists the purpose for the possible procedures using the **options** attribute.

| Procedure | Purpose |
|---|---|
| Disable Outbound Replication | Use this procedure to disable Active Directory replication from a domain controller. The domain controller continues to receive inbound replication.<br><br>**Repadmin /options** <*ServerName*> **+disable_outbound_repl** where <*ServerName*> is the name of the domain controller on which you want to disable outbound replication. The tool reports the current options (the options that were in effect prior to pressing ENTER) and the new options (all options that are in effect after pressing ENTER). |
| Disable inbound Replication | Similar to the above step you could disable inbound replication to a server as well.<br><br>**repadmin /options** <*ServerName*> **+disable_inbound_repl** |
| Disable the ability of the KCC to translate connection objects. | When creating temporary replication links between replication partners, the process could fail if the KCC starts while you perform the procedure. The KCC will delete any replication links for which no corresponding connection object exists. |

**Advanced site options**

By using the **siteoptions** subcommand, we could change the options attribute stored on the NTDS Site Settings Object.

**Syntax**

```
Repadmin /siteoptions <DC> /site:< Site> [{+|-}IS_AUTO_TOPOLOGY_DISABLED] [{+|-} IS_TOPL_CLEANUP_DISABLED]
[{+|-} IS_TOPL_MIN_HOPS_DISABLED] [{+|-} IS_TOPL_DETECT_STALE_DISABLED] [{+|-}
IS_INTER_SITE_AUTO_TOPOLOGY_DISABLED] [{+|-} IS_GROUP_CACHING_ENABLED] [{+|-} FORCE_KCC_WHISTLER_BEHAVIOR]
[{+|-} FORCE_KCC_W2K_ELECTION] [{+|-} IS_RAND_BH_SELECTION_DISABLED] [{+|-} IS_SCHEDULE_HASHING_ENABLED]
[{+|-} IS_REDUNDANT_SERVER_TOPOLOGY_ENABLED]
```

| Parameter | Description |
|---|---|
| <*DC*> | Domain controller |
| site: <*Site*> | Site name where the domain controller resides |
| IS_AUTO_TOPOLOGY_DISABLED | Disables the automatic generation of intra-site topology. |
| IS_TOPL_CLEANUP_DISABLED | Disables the cleanup or unneeded connection objects and replication links. |
| IS_TOPL_MIN_HOPS_DISABLED | Disables the KCC rule that all intrasite replication partners should be no more than three hops from any other partner. |
| IS_TOPL_DETECT_STALE_DISABLED | Disables the detection by the KCC of failing replication links and the behavior of the KCC to route around failing links. Use this with the KCC Branch Office mode. |

| | |
|---|---|
| IS_INTER_SITE_AUTO_TOPOLOGY_DISABLED | Disables the automatic generation of the intersite topology. Commonly used for creating manual connections, either by hand or with MKDSX. |
| IS_GROUP_CACHING_ENABLED | Enables group caching for use with “no-GC logon.” This setting is also exposed in the UI of Active Directory Sites and Services. |
| FORCE_KCC_WHISTLER_BEHAVIOR | Forces the KCC to operate using the new spanning tree algorithm. It’s not recommended to manually change this setting. The recommended alternative is to raise the forest functional level to Windows Server 2003. |
| FORCE_KCC_W2K_ELECTION | Forces the Windows 2000 domain controller ISTG election logic. The default is for any Windows Server 2003 domain controller to assume the ISTG role. |
| IS_RAND_BH_SELECTION_DISABLED | Disables the new random bridgehead selection behavior. Reverts to Windows 2000 KCC behavior of using a single bridgehead server. |
| IS_SCHEDULE_HASHING_ENABLED | Creates a random schedule on each new connection object based in hashed value. Helps to balance the load on bridgehead servers. |
| IS_REDUNDANT_SERVER_TOPOLOGY_ENABLED | Creates two inbound connection objects from different domain controllers in a hub site. Reduces impact on FRS (vvjoin) during failover. |

**Miscellaneous**

The following table lists nbrflagoptions.

| Parameter | Definition |
|---|---|
| SYNC_ON_STARTUP | Replication of this naming context from this source is attempted when the destination server is booted. This normally only applies to intra-site neighbors. |
| DO_SCHEDULED_SYNCS | Perform replication on a schedule. This flag is normally set unless the schedule for this naming context and source is "never", that is, the empty schedule. |
| WRITEABLE | The local copy of the naming context is writable. |
| TWO_WAY_SYNC | If set, indicates that when inbound replication is complete, the destination server must tell the source server to synchronize in the reverse direction. This feature is used in dial-up scenarios where only one of the two servers can initiate a dial-up connection. For example, this option would be used in a corporate headquarters and branch office, where the branch office connects to the corporate headquarters over the Internet by means of a dial-up ISP connection. |
| NEVER_SYNCED | Synchronization has never been successfully completed from this source. |
| IGNORE_CHANGE_NOTIFICATIONS | This neighbor is set to disable notification-based synchronizations. Within a site, domain controllers synchronize with each other based on notifications when changes occur. This setting prevents this neighbor from performing synchronizations that are triggered by notifications. The neighbor will still do synchronizations based on its schedule, or in response to manually requested synchronizations. |
| DISABLE_SCHEDULED_SYNC | This neighbor is set to not perform synchronizations based on its schedule. The only way this neighbor will perform synchronizations is in response to change notifications or to manually requested synchronizations. |
| COMPRESS_CHANGES | Changes received from this source are to be compressed. This is normally set if, and only if, the source server is in a different site. |
| NO_CHANGE_NOTIFICATIONS | No change notifications should be received from this source. Normally set if, and only if, the source server is in a different site. |

# Administering Active Directory Operations

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides administering information for Active Directory in the Microsoft Windows Server 2003 and Windows Server 2003 with Service Pack 1 (SP1) operating systems. Information includes detailed procedures for managing domain controllers, sites, trusts, and other components of Active Directory.

**In this guide**


- Introduction to Administering Active Directory
- New in This Guide for Administering Active Directory
- Administering Domain and Forest Trusts
- Administering the Windows Time Service
- Administering SYSVOL
- Administering the Global Catalog
- Administering Operations Master Roles
- Administering Active Directory Backup and Restore
- Administering Intersite Replication
- Administering the Active Directory Database
- Administering Domain Controllers
- Additional Resources for Administering Active Directory


| **Note** |
| --- |
| You cannot install Active Directory on a server running Windows Server 2003, Web Edition, but you can join the server to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition, on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=9253). |


**Acknowledgments**

Key Technical Reviewers: Chris Macaulay, Nigel Cain, Arren Conner, Dmitry Dukat, Levon Esibov, Khushru Irani, Kamal Janardhan, Gregory Johnson, William Lees, Andreas Luther, Kevin Sims, Jeromy Statia, Eric Kool-Brown, J. K. Jaganathan, Mike Resnick, Michael Snyder, Nathan Muggli, Yi Zhao, Christopher Westpoint, Robert Powalka, Rob Kochman

Microsoft Most Valuable Professional (MVP) Reviewers: Joseph Shook, Thomas Bittner, Nuo Yan, Al Mulnick, Tony Murray, Guido Grillenmeier, M. Rajesh, Todd Myrick

# Introduction to Administering Active Directory

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide explains how to administer Microsoft Active Directory. These activities are part of the operating phase of the information technology (IT) life cycle. If you are not familiar with this guide, review the following sections of this introduction.

**When to Use This Guide**

You should use this guide when:

- You want to manage common Active Directory problems that are associated with misconfiguration.
- You want to configure Active Directory to increase network availability.

This guide assumes a basic understanding of what Active Directory is, how it works, and why your organization uses it to access, manage, and secure shared resources across your network. You should also have a thorough understanding of how Active Directory is deployed and managed in your organization. This includes an understanding of the mechanism your organization uses to configure and manage Active Directory settings.

This guide can be used by organizations that have deployed Windows Server 2003 and Windows Server 2003 with Service Pack 1 (SP1). It includes information that is relevant to different roles within an IT organization, including IT operations management and administrators. It contains high-level information that is required to plan an Active Directory operations environment. This information provides management-level knowledge of Active Directory and the IT processes required to operate it.

In addition, this guide contains more detailed procedures that are designed for operators who have varied levels of expertise and experience. Although the procedures provide operator guidance from start to finish, operators must have a basic proficiency with the Microsoft Management Console (MMC) and snap-ins and know how to start administrative programs and access the command line. If operators are not familiar with Active Directory, it might be necessary for IT planners or IT managers to review the relevant operations in this guide and provide the operators with parameters or data that must be entered when the operation is performed.

**How to Use this Guide**

The operations areas are divided into the following types of content:

- **Objectives** are high-level goals for managing, monitoring, optimizing, and securing Active Directory. Each objective consists of one or more high-level tasks that describe how the objective is accomplished.
- **Tasks** are used to group related procedures and provide general guidance for achieving the goals of an objective.
- **Procedures** provide step-by-step instructions for completing the task.

If you are an IT manager who will be delegating tasks to operators within your organization, you will want to:

- Read through the objectives and tasks to determine how to delegate permissions and whether you need to install tools before operators perform the procedures for each task.
- Before assigning tasks to individual operators, ensure that you have all the tools installed where operators can use them.
- When necessary, create “tear sheets” for each task that operators perform within your organization. Cut and paste the task and its related procedures into a separate document and then either print these documents or store them online, depending on the preference of your organization.

# New in This Guide for Administering Active Directory

Updated: May 8, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This operations guide is updated periodically to incorporate new content, customer feedback, and corrections. The following sections provide details about content that is new or updated in this version of the guide.

**New Content**

- August 2005: Performing an Authoritative Restore of Active Directory Objects contains new procedures for regenerating the group memberships of restored user objects and group objects. This functionality is available in the version of Ntdsutil.exe that is included with Windows Server 2003 with Service Pack 1 (SP1).

- February 2006: Enable Remote Desktop contains a new procedure to enable Remote Desktop remotely by using the registry.

- February 2006: Known Issues for Adding Domain Controllers in Remote Sites contains the additional information that moving the Ntds.dit file takes less time than copying the file when you restore a system state backup.

**Updated Content**

- April 2006: Performing an Authoritative Restore of Active Directory Objects contains corrected information about the details of updating back-link attributes.

# Administering Domain and Forest Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides administrators with step-by-step instructions for managing and securing Windows Server 2003 domain and forest trusts. The way that you create or configure trusts plays an important role in operating and securing your network infrastructure. How you create or configure domain and forest trusts in Windows Server 2003 also determines how far network communications extend within a forest or across forests.

| **Note** |
| --- |
| You cannot install Active Directory on a server running Windows Server 2003, Web Edition, but you can join the server to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition, on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=9253). |

**In this guide**


- Introduction to Domain and Forest Trusts
- Best Practices for Domain and Forest Trusts
- Managing Domain and Forest Trusts
- Securing Domain and Forest Trusts
- Appendix: New Trust Wizard Pages



**Acknowledgements**

Produced by: Microsoft Windows Server Security and Directory Services User Assistance team

Project Writer: Nick Pierson

Project Editor: Jim Becker

# Introduction to Domain and Forest Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using Windows Server 2003 domain and forest trusts, service administrators can create or extend collaborative relationships between two or more domains or forests. Windows Server 2003 domains and forests can also trust Kerberos realms and other Windows Server 2003 forests, as well as Microsoft Windows® 2000 domains and Windows NT® 4.0 domains.

When a trust exists between two domains, the authentication mechanisms for each domain trust the authentications coming from the other domain. Trusts help to provide controlled access to shared resources in a resource domain (the trusting domain) by verifying that incoming authentication requests come from a trusted authority (the trusted domain). In this way, trusts act as bridges that allow only validated authentication requests to travel between domains.

How a specific trust passes authentication requests depends on how it is configured. Trust relationships can be one-way, providing access from the trusted domain to resources in the trusting domain, or two-way, providing access from each domain to resources in the other domain. Trusts are also either nontransitive, in which case a trust exists only between the two trust partner domains, or transitive, in which case a trust automatically extends to any other domains that either of the partners trusts.

In some cases, trust relationships are established automatically when domains are created; in other cases, administrators must choose a type of trust and explicitly establish the appropriate relationships. The specific types of trusts that are used and the structure of the resulting trust relationships in a given trust implementation depend on such factors as how Active Directory is organized and whether different versions of Windows coexist on the network.

# Best Practices for Domain and Forest Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following best practices are proven to increase availability, ensure trouble-free operations, or ease administration when you use them to administer domain and forest trusts:

**When your forest contains domain trees with many child domains and you observe noticeable user authentication delays between the child domains, you can optimize the user authentication process between the child domains by creating shortcut trusts to mid-level domains in the domain tree hierarchy.**

For more information, see When to create a shortcut trust on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42644).

**Keep a current list of trust relationships for future reference.**

You can use the Nltest.exe tool to display and record a list of these trusts. For more information, see "Nltest.exe: NLTest Overview" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42642).

**Perform regular backups of domain controllers to preserve all trust relationships within a particular domain.**

For more information, see Back up system state.

# Managing Domain and Forest Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is necessary to manage domain and forest trusts when your organization needs to collaborate with users or resources that are located in other domains, realms, or forests in your organization and in other organizations. To set up an environment that takes advantage of trusts, you must first create and configure the appropriate trusts that will enable your organization to communicate effectively with users or resources in other places.

The following objectives are part of managing domain and forest trusts:

- Creating Domain and Forest Trusts
- Configuring Domain and Forest Trusts

# Creating Domain and Forest Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Windows Server 2003, there are four trust types that must be created manually. External trusts, realm trusts, and forest trusts help provide interoperability with domains outside your forest or with realms. Shortcut trusts optimize access to resources and logons that are made between domain trees in the same forest.

The following tasks for creating domain and forest trusts are described in this objective:

- Creating External Trusts
- Creating Shortcut Trusts
- Creating Forest Trusts
- Creating Realm Trusts

| **Note** |
| --- |
| A trust does not inherently allow users in a trusted domain to have access to resources in a trusting domain. Users have access when they are assigned the appropriate permissions. In some cases, users in trusted domains may have implicit access if the resources are assigned to Authenticated Users. |

**New Trust Wizard Terminology**

You create trusts in Windows Server 2003 with the New Trust Wizard. Before you use the New Trust Wizard, review the following terminology. Each highlighted term represents the exact term as it is used in the wizard:

- **This domain:** The domain from which you launch the New Trust Wizard. When you start the wizard, it immediately verifies your administrative credentials in the domain for which you are the administrator. Therefore, the wizard uses the term “this domain” to represent the domain that you are currently logged on to.
- **Local domain / Local forest:** The domain or forest where you start the New Trust Wizard.
- **Specified domain / Specified forest:** The other domain or forest that this local domain or local forest will trust. Although the New Trust Wizard is aware of the domain context in which it is running, it does not have knowledge of the other domain that you want to create the relationship with. After you type the name of the other domain or forest in the **Trust Name** page, that name is used whenever the wizard refers to the specified domain or specified forest.
- **Two-way trust:** A trust relationship between two domains in which both domains trust each other. For example, domain A trusts domain B, and domain B trusts domain A. All parent-child trusts are two-way trusts.
- **One-way: incoming trust:** A one-way trust relationship between two domains in which the direction of the trust points toward the domain from which you start the New Trust Wizard (and which is identified in the wizard as **This domain**). When the direction of the trust points toward your domain, users in your domain can access resources in the specified domain. For example, if you are the domain administrator in domain A and you create a one-way, incoming trust to domain B, this provides a relationship through which users who are located in domain A can access resources in domain B. Because this relationship is one-way, users in domain B cannot access resources in domain A.
- **One-way: outgoing trust:** A one-way trust relationship between two domains in which the direction of the trust points toward the domain that is identified as **Specified domain** in the New Trust Wizard. When the direction of trust points toward the specified domain, users in the specified domain can access resources in your domain. For example, if you are the domain administrator in domain A and you create a one-way, outgoing trust to domain B, this provides a relationship through which users who are located in domain B can access resources in domain A. Because this relationship is one way, users in domain A cannot access resources in domain B.

- **Both sides of the trust:** When you create external trusts, shortcut trusts, or forest trusts, you have the option to create each side of the trust separately or both sides of the trust simultaneously. If you choose to create each side of the trust separately, you must run the New Trust Wizard twice — once for each domain. When you create trusts separately, you must supply the same trust password for each domain. As a security best practice, all trust passwords should be strong passwords.

- **Domain-wide authentication:** An authentication setting that permits unrestricted access by any users in the specified domain to all available shared resources that are located in the local domain. This is the default authentication setting for external trusts.

- **Forest-wide authentication:** An authentication setting that permits unrestricted access by any users in the specified forest to all available shared resources that are located in any of the domains in the local forest. This is the default authentication setting for forest trusts.

- **Selective authentication:** An authentication setting that restricts access over an external trust or forest trust to only those users in a specified domain or specified forest who have been explicitly given authentication permissions to computer objects (resource computers) that reside in the local domain or the local forest. This authentication setting must be enabled manually.

- **Trust password:** An option in which both domains in a trust relationship share a password, which is stored in the trusted domain object (TDO) object in Active Directory. When you choose this option, a strong trust password is generated automatically for you. You must use the same password when you create a trust relationship in the specified domain. If you choose to create both sides of the trust simultaneously, you run the New Trust Wizard once.

# DNS and NetBIOS Name Resolution to Create External, Realm and Forest Trusts

Published: July 23, 2009

Updated: July 24, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Before you create a domain or forest trust, there are name resolution requirements that must be met so that you can establish the trust successfully. These requirements ensure that the forest or domain can resolve the address of the domain controllers in the target forest or domain to support the creation and ongoing operation of the trust.

For all trusts except trusts across forests, the New Trust Wizard and its command-line equivalents allow you to refer to domains that participate in a trust by their Domain Name System (DNS) or NetBIOS domain names. For trusts across forests, you must use a DNS name.

Although Windows Internet Name Service (WINS) has been used to support trusts in the past, the focus of this topic is configuring DNS to create and operate a trust.

**DNS Name Resolution**

You can use DNS to resolve names for a forest or domain for which you are attempting to establish a trust. Most DNS names are fully qualified, which means that there is a dot (.) in the name between the forest root domain and a top-level domain, such as .com, .gov, and .local. There are rare instances in which Active Directory domains are created with a single-label DNS domain name, such as “contoso”. When you type in the name of a remote domain in the New Trust wizard that you are expecting to be resolved by DNS, then you should enter the DNS name of the remote domain (not its NetBIOS or WINS name).

The DNS servers that the domain controllers in one forest use may not be able to resolve the DNS names for domain controllers in another forest, and the other way around. If the DNS names for trusting domains are not registered with an Internet Register, then the DNS Servers used by DCs in the trusted domain cannot use forwarders to resolve the DNS queries needed to support the trust. There are two common DNS configurations that you can choose from when you establish DNS name resolution between domains and forests:

- Conditional forwarders
- Secondary zones with zone transfers enabled

| **Note** |
| --- |
| Stub zones can be used to resolve DNS names, but are less commonly used than secondary zones and conditional forwarders; therefore, they are not covered in this document. |

**Using Conditional Forwarders vs. Using Secondary Zones**

There are several conditions you should take into account when choosing between conditional forwarders or secondary zones with zone transfers (to resolve DNS names that your DNS Servers are not authoritative for). When you decide whether to use conditional forwarders or secondary zones with zone transfers enabled, there are several considerations you should take into account before choosing either configuration.

The benefit of using a conditional forwarder is that it is much easier to configure and troubleshoot than a zone transfer. The process of configuring a conditional forwarder is straightforward: all you need to know is the DNS domain name of the domain that houses the DNS server that you are configuring to forward requests and the IP address of the target DNS server.

However, a conditional forwarder is not an efficient way to keep a DNS server that hosts a parent zone aware of the authoritative DNS servers for a child zone. If you use a conditional forwarder, whenever the authoritative DNS servers for the child zone change, the conditional forwarder setting on the DNS server that hosts the parent zone must be configured manually with the IP address for each new authoritative DNS server for the child zone.

Using a secondary zone with zone transfers enabled is beneficial because this configuration maintains a list of all the authoritative DNS servers for the secondary copy of the zone, and the list is updated as DNS servers are added and removed from the target forest or domain. Secondary zones also host a full copy of the DNS zone.

The drawbacks to using secondary zones with zone transfers enabled are that this configuration is much more complicated to configure and maintain and you do not have the direct, point-to-point contact with a DNS server in the target forest or domain as you do with a conditional forwarder. In addition, with secondary zones you expose hosts to IP address mappings for all hosts in the zone. This can expose the domain or forest to security risks due to unauthorized access.

**Conditional Forwarders**

You can use conditional forwarders to route queries between the domains or forests for which you are establishing a forest or domain trust. Conditional forwarders route names to a specified DNS server in the target domain or forest for queries that the conditional forwarders cannot resolve.

For example, if you establish a two-way, forest trust between the fabrikam.com and contoso.com forests, you configure a conditional forwarder in the fabrikam.com forest to route queries that it cannot resolve to a name server in the contoso.com forest. You then configure a conditional forwarder in the contoso.com forest to route names that it cannot resolve to the fabrikam.com forest.

For more information about conditional forwarders, see the following:

- Using forwarders (http://go.microsoft.com/fwlink/?LinkId=157391)
- How to Configure Conditional Forwarders in Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkId=157393). This post not only explains how to configure a conditional forwarder in Windows Server 2008, it also explains how to configure a conditional forwarder in Windows Server 2003.

The following flowchart explains how to troubleshoot instances in which conditional forwarders are not correctly performing name resolution to the target forest or domain. The flowchart shows a one-way, transitive, outgoing, forest trust between forest A (trusting) and forest B (trusted). Simply put, users in forest B have access to specified resources in forest A. To complete the steps in the following flowchart, use the credentials of a member of the Domain Admins group. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).




**Secondary Zones with Zone Transfers Enabled**

Zone transfers replicate DNS databases across a set of DNS servers. If you want to establish a trust between two forests, you can configure the DNS server in the trusting domain or forest to host a secondary zone of the trusted forest, with zone transfers enabled.

For example, if you want to establish a two-way trust between the fabrikam.com and contoso.com forests, you configure the DNS server in the fabrikam.com forest to host a secondary zone of the contoso.com forest's namespace, and the reverse. After

you configure the secondary zones in each forest, you then allow zone transfers between the two forests. For more information about secondary zones or zone transfers, see the following:

- Adding a Secondary DNS Server to a Zone (http://go.microsoft.com/fwlink/?LinkId=157397)
- Understanding zones and zone transfers (http://go.microsoft.com/fwlink/?LinkId=157398)
- Understanding Zone Types (http://go.microsoft.com/fwlink/?LinkId=157399)

**DNS Resource Records That Are Required for Secondary Zones**

There are two DNS resource records that must be registered properly on the DNS server that hosts the secondary copy of the trusted domain or forest:

- Service (SRV) resource record (_ldap._tcp.dc._msdcs.<*computer account domain*>)
- Host (A) resource record

These records must be in place and registered properly before you establish a domain or forest trust.

The following flowchart explains how to troubleshoot instances in which secondary zones with zone transfers enabled are not correctly performing name resolution to the target domain or forest. The flowchart shows a one-way, transitive, outgoing, forest trust between forest A (trusting) and forest B (trusted). Simply put, users in forest B have access to specified resources in forest A. To complete the steps in the following flowchart, use the credentials of a member of the Domain Admins group. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).



# Known Issues for Creating Domain and Forest Trusts

Updated: October 1, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Review the following known issues before creating domain and forest trusts in Windows Server 2003:

- You cannot delegate the creation of trusts to any user who is not a member of the Domain Admins or Enterprise Admins groups. Even though you can grant a user the Create TDO (Trusted Domain Object) right or the Delete TDO right in the System container of a domain, the user will not be granted the right to create a trust. This issue occurs because Netlogon and the trust-creation tools (Active Directory Domains and Trusts and Netdom) are designed so that only members of the Domain Admins group and the Enterprise Admins group can create trusts. However, any user who is a member of the Incoming Forest Trust Builders group can create one-way, incoming forest trusts to your forest. For more information about the Incoming Forest Trust Builders group, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

- When you are logged on locally to a domain controller and you try to create a new trust by using Active Directory Domains and Trusts, the operation may be unsuccessful and you may receive the message “Access denied.” This issue occurs only if you are logged on locally to the domain controller as an ordinary user (meaning that the user is not logged on as Administrator or as a member of any administrative groups for the domain). By default, ordinary users are blocked from logging on locally to a domain controller unless Group Policy is modified to permit this.

- When you use Active Directory Domains and Trusts to create a trust, you may receive the message “Operation failed. Parameter incorrect.” This issue may occur if you try to establish a trust relationship when the source domain and the target domain have one or more of the following identifiers that are the same:

  - Security identifier (SID)

  - Domain Name System (DNS) name

  - Network basic input/output system (NetBIOS) name

  To resolve this issue, rename all conflicting identifiers before you try to create the trust.

- The option to create a forest trust does not appear in the New Trust Wizard. This issue typically occurs when one or both of the Windows Server 2003 forests are not set to the Windows Server 2003 forest functional level. For more information about forest functional levels, see Active Directory Functional Levels Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41698).

- There are restrictions in the number and types of trusts that can be created when you target a Microsoft Windows Small Business Server 2003 domain.

# Creating External Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can create an external trust to form a one-way or two-way, nontransitive trust with domains that are outside your forest. External trusts are sometimes necessary when users need access to resources that are located in a Windows NT 4.0 domain or in a domain that is in a separate Active Directory forest that is not joined by a forest trust.

For example, if you have a Windows Server 2003–based domain whose users want to gain access to resources that are stored in a Windows NT–based domain, you must create a trust relationship in which the Windows NT–based domain trusts the users from the Windows Server 2003–based domain. In this case, the Windows NT–based domain is the trusting domain, and the Windows Server 2003–based domain is the trusted domain.

For more information about external trusts, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

| **Note** |
| --- |
| Trusts that are created between Windows NT 4.0 domains and Active Directory domains are one-way and nontransitive, and they require network basic input/output system (NetBIOS) name resolution. |

**Task requirements**

You can use either of the following tools to perform the procedures for this task:

- Active Directory Domains and Trusts
- Netdom.exe

For more information about how to use the Netdom command-line tool to create an external trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

| **Note** |
| --- |
| If you have the appropriate administrative credentials for each domain, you can create both sides of an external trust at the same time. To create both sides of the trust, simultaneously, follow the appropriate procedure below that contains the words “both sides of the trust” in the procedure title. For example, the procedure “Create a one-way, incoming, external trust for both sides of the trust” provides the steps to follow when you have the administrative credentials for both domains and you want to use the New Trust Wizard to create an incoming, external trust in one operation. For more information about how the “both sides of the trust” option works, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages. |

You can create an external trust by using any of the following procedures, depending on the requirements of your organization and the administrative credentials that you have when you create the trust:

- Create a one-way, incoming, external trust for one side of the trust
- Create a one-way, incoming, external trust for both sides of the trust
- Create a one-way, outgoing, external trust for one side of the trust
- Create a one-way, outgoing, external trust for both sides of the trust
- Create a two-way, external trust for one side of the trust
- Create a two-way, external trust for both sides of the trust

# Create a one-way, incoming, external trust for one side of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates one side of a one-way, incoming, external trust. Although one side of a trust will be created successfully, the new trust will not function until the administrator for the reciprocal domain uses his or her credentials to create the second side of the trust. If you have administrative credentials for both domains that are involved in the trust, you can use the procedure Create a one-way, incoming, external trust for both sides of the trust to create both sides of the trust in one simultaneous operation.

A one-way, incoming, external trust allows users in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to access resources in another Active Directory domain (outside your forest) or in a Windows NT 4.0 domain. For example, if you are the administrator of sales.wingtiptoys.com and users in that domain need to access resources in the marketing.tailspintoys.com domain (which is located in another forest), you can use this procedure (in conjunction with another procedure, which is executed by the administrator in the other forest) to establish one side of the relationship so that users in your domain can access resources in the marketing.tailspintoys.com domain.

You can create this external trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create an external trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, incoming, external trust for one side of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: incoming**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **This domain only**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Trust Password** page, type the trust password twice, and then click **Next**.

   With the administrator of the other domain, agree on a secure channel password to be used in establishing the trust.
9. On the **Trust Selections Complete**page, review the results, and then click **Next**.
10. On the **Trust Creation Complete** page, review the results, and then click **Next**.
11. On the **Confirm Incoming Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.
    - If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

12. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the domain administrator for the specified domain or specified forest must follow the procedure Create a one-way, outgoing, external trust for one side of the trust, using his or her administrative credentials and the exact same trust password that was used during this procedure. |

# Create a one-way, incoming, external trust for both sides of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates both sides of a one-way, incoming, external trust, and it requires you to have administrative credentials for your domain as well for the reciprocal domain. If you have administrative credentials only for your domain, you can use the procedure Create a one-way, incoming, external trust for one side of the trust to create your side of the trust. Then, have the administrator for the reciprocal domain create a one-way, outgoing, external trust from his or her domain.

A one-way, incoming, external trust allows users in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to access resources in another Active Directory domain (outside your forest) or in a Windows NT 4.0 domain. For example, if you are the administrator of sales.wingtiptoys.com and users in that domain need to access resources in the marketing.tailspintoys.com domain (which is located in another forest) you can use this procedure to establish a relationship so that users in your domain can access resources in the marketing.tailspintoys.com domain.

You can create this external trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create an external trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, incoming, external trust for both sides of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: incoming**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **User Name and Password** page, type the user name and password for the appropriate administrator in the specified domain.
9. On the **Outgoing Trust Authentication Level--Specified Domain** page, do one of the following, and then click **Next**:
   - Click **Domain-wide authentication**.
   - Click **Selective authentication**.
10. On the **Trust Selections Complete** page, review the results, and then click **Next**.
11. On the **Trust Creation Complete** page, review the results, and then click **Next**.
12. On the **Confirm Incoming Trust** page, do one of the following:

- If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.
- If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

13. On the **Completing the New Trust Wizard** page, click **Finish**.

# Create a one-way, outgoing, external trust for one side of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates one side of a one-way, outgoing, external trust. Although one side of a trust will be created successfully, the new trust will not function until the administrator for the reciprocal domain uses his or her credentials to create the second side of the trust. If you have administrative credentials for both domains that are involved in the trust, you can use the procedure Create a one-way, outgoing, external trust for both sides of the trust to create both sides of the trust in one simultaneous operation.

A one-way, outgoing, external trust will allow resources in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to be accessed by users in a different Active Directory domain (outside your forest) or in a Windows NT 4.0 domain. For example, if you are the administrator of sales.wingtiptoys.com and you have resources in that domain that need to be accessed by users in the marketing.tailspintoys.com domain (which is located in another forest), you can use this procedure to establish one side of the relationship so that users in the marketing.tailspintoys.com domain can access the resources in sales.wingtiptoys.com.

You can create this external trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create an external trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, outgoing, external trust for one side of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: outgoing**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **This domain only**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Outgoing Trust Authentication Level** page, do one of the following, and then click **Next**:
   - Click **Domain-wide authentication**.
   - Click **Selective authentication**.
9. On the **Trust Password** page, type the trust password twice, and then click **Next**.
10. On the **Trust Selections Complete**page, review the results, and then click **Next**.
11. On the **Trust Creation Complete** page, review the results, and then click **Next**.
12. On the **Confirm Outgoing Trust** page, do one of the following:

- If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.

- If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the **specified domain.**

13. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the domain administrator for the specified domain or specified forest must follow the procedure Create a one-way, incoming, external trust for one side of the trust, using his or her administrative credentials and the exact same trust passwordthat was used during this procedure. |

# Create a one-way, outgoing, external trust for both sides of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates both sides of a one-way, outgoing, external trust, and it requires you to have administrative credentials for your domain as well as for the reciprocal domain. If you have administrative credentials only for your domain, you can use the procedure Create a one-way, outgoing, external trust for one side of the trust to create your side of the trust. Then, have the administrator for the reciprocal domain create a one-way, incoming, external trust from his or her domain.

A one-way, outgoing, external trust allows resources in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to be accessed by users in a different Active Directory domain (outside your forest) or in a Windows NT 4.0 domain. For example, if you are the administrator of sales.wingtiptoys.com and you have resources in that domain that need to be accessed by users in the marketing.tailspintoys.com domain (which is located in another forest), you can use this procedure to establish one side of the relationship so that users in the marketing.tailspintoys.com domain can access the resources in sales.wingtiptoys.com.

You can create this external trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create an external trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, outgoing, external trust for both sides of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: outgoing**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **User Name and Password** page, type the user name and password for the appropriate administrator in the specified domain.
9. On the **Outgoing Trust Authentication Level--Local Domain** page, do one of the following, and then click **Next**:
   - Click **Domain-wide authentication**.
   - Click **Selective authentication**.
10. On the **Trust Selections Complete** page, review the results, and then click **Next**.
11. On the **Trust Creation Complete** page, review the results, and then click **Next**.
12. On the **Confirm Outgoing Trust** page, do one of the following:

- If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.

- If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the specified domain.

13. On the **Completing the New Trust Wizard** page, click **Finish**.

# Create a two-way, external trust for one side of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates one side of a two-way, external trust. Although one side of a trust will be created successfully, the new trust will not function until the administrator for the reciprocal domain uses his or her credentials to create the second side of the trust. If you have administrative credentials for both domains that are involved in the trust, you can use the procedure Create a two-way, external trust for both sides of the trust to create both sides of the trust in one simultaneous operation.

A two-way, external trust allows users in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) and users in the reciprocal domain to access resources in either of the two domains.

You can create this external trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create an external trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a two-way, external trust for one side of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **Two-way**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **This domain only**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Outgoing Trust Authentication Level** page, do one of the following, and then click **Next**:
   - Click **Domain-wide authentication**.
   - Click **Selective authentication**.
9. On the **Trust Password** page, type the trust password twice, and then click **Next**.
10. On the **Trust Selections Complete**page, review the results, and then click **Next**.
11. On the **Trust Creation Complete** page, review the results, and then click **Next**.
12. On the **Confirm Outgoing Trust** page, do one of the following:
    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.
    - If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate

administrative credentials from the specified domain.

13. On the **Confirm Incoming Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.

    - If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

14. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the domain administrator for the specified domain or specified forest must follow this same procedure, using his or her administrative credentials and the exact same trust password that was used during this procedure. |

# Create a two-way, external trust for both sides of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates both sides of a two-way, external trust, and it requires you to have administrative credentials for your domain as well as for the reciprocal domain. If you have administrative credentials only for your domain, you can use the procedure Create a two-way, external trust for one side of the trust to create your side of the trust. Then, have the administrator for the reciprocal domain create a one-way, outgoing, external trust from his or her domain.

A two-way, external trust allows users in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) and users in the reciprocal domain to access resources in either of the two domains.

You can create this external trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create an external trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a two-way, external trust for both sides of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **Two-way**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **User Name and Password** page, type the user name and password for the appropriate administrator in the specified domain.
9. On the **Outgoing Trust Authentication Level--Local Domain** page, do one of the following, and then click **Next**:
   - Click **Domain-wide authentication**.
   - Click **Selective authentication**.
10. On the **Outgoing Trust Authentication Level--Specified Domain** page, do one of the following, and then click **Next**:
    - Click **Domain-wide authentication**.
    - Click **Selective authentication**.
11. On the **Trust Selections Complete** page, review the results, and then click **Next**.
12. On the **Trust Creation Complete** page, review the results, and then click **Next**.

13. On the **Confirm Outgoing Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.

    - If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the specified domain.

14. On the **Confirm Incoming Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.

    - If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

15. On the **Completing the New Trust Wizard** page, click **Finish**.

# Creating Shortcut Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A shortcut trust is a manually created trust that shortens the trust path to improve the speed at which authentications, which are made between domain trees, are processed. This can result in faster logon times and faster access to resources. A trust path is a chain of multiple trusts that enables trust between domains that are not adjacent in the domain namespace. For example, if users in domain A need to gain access to resources in domain C, you can create a direct link from domain A to domain C through a shortcut trust relationship, bypassing domain B in the trust path.

For more information about shortcut trusts, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

**Task requirements**

You can use either of the following tools to perform the procedures for this task:

- Active Directory Domains and Trusts
- Netdom.exe

For more information about how to use the Netdom command-line tool to create a shortcut trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

| **Note** |
| --- |
| If you have the appropriate administrative credentials for each domain, you can create both sides of a shortcut trust at the same time. To create both sides of the trust, follow the appropriate procedure below that contains the words “for both sides of the trust” in the title. For example, the procedure “Create a one-way, incoming, shortcut trust for both sides of the trust” explains how to configure both sides of a shortcut trust. For more information about how the “both sides of the trust” option works, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages. |

You can create a shortcut trust by using any of the following procedures, depending on the requirements of your organization and the administrative credentials that you have when you create the trust:

- Create a one-way, incoming, shortcut trust for one side of the trust
- Create a one-way, incoming, shortcut trust for both sides of the trust
- Create a one-way, outgoing, shortcut trust for one side of the trust
- Create a one-way, outgoing, shortcut trust for both sides of the trust
- Create a two-way, shortcut trust for one side of the trust
- Create a two-way, shortcut trust for both sides of the trust

# Create a one-way, incoming, shortcut trust for one side of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates one side of a one-way, incoming, shortcut trust. Although one side of a trust will be created successfully, the new trust will not function until the administrator for the reciprocal domain uses his or her credentials to create the second side of the trust. If you have administrative credentials for both domains that are involved in the trust, you can use the procedure Create a one-way, incoming, shortcut trust for both sides of the trust to create both sides in one simultaneous operation.

A one-way, incoming, shortcut trust allows users in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to more quickly access resources in another domain (which is nested within another domain tree) in your forest. For example, if you are the administrator of sales.wingtiptoys.com and users in that domain need to access resources in the marketing.tailspintoys.com domain (which is a child domain of the tailspintoys.com tree root domain), you can use this procedure to establish one side of the relationship so that users in your domain can more quickly access resources in the marketing.tailspintoys.com domain.

You can create this shortcut trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a shortcut trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, incoming, shortcut trust for one side of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: incoming**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **This domain only**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Trust Password** page, type the trust password twice, and then click **Next**.
9. On the **Trust Selections Complete** page, review the results, and then click **Next**.
10. On the **Trust Creation Complete** page, review the results, and then click **Next**.
11. On the **Confirm Incoming Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.
    - If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

12. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the domain administrator for the specified domain or specified forest must follow the procedure Create a one-way, outgoing, shortcut trust for one side of the trust, using his or her administrative credentials and the exact same trust password that was used during this procedure. |

# Create a one-way, incoming, shortcut trust for both sides of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates both sides of a one-way, incoming, shortcut trust, and it requires you to have administrative credentials for your domain as well for the reciprocal domain. If you have administrative credentials only for your domain, you can use the procedure Create a one-way, incoming, shortcut trust for one side of the trust to create your side of the trust. Then, have the administrator for the reciprocal domain create a one-way, outgoing, shortcut trust from his or her domain.

A one-way, incoming, shortcut trust allows users in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to more quickly access resources in another domain (which is nested within another domain tree) in your forest. For example, if you are the administrator of sales.wingtiptoys.com and users in that domain need to access resources in the marketing.tailspintoys.com domain (which is a child domain of the tailspintoys.com tree root domain), you can use this procedure to establish one side of the relationship so that users in your domain can more quickly access resources in the marketing.tailspintoys.com domain.

You can create this shortcut trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a shortcut trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, incoming, shortcut trust for both sides of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: incoming**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **User Name and Password** page, type the user name and password for the appropriate administrator in the specified domain.
9. On the **Trust Selections Complete** page, review the results, and then click **Next**.
10. On the **Trust Creation Complete** page, review the results, and then click **Next**.
11. On the **Confirm Incoming Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.
    - If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

12. On the **Completing the New Trust Wizard** page, click **Finish**.

# Create a one-way, outgoing, shortcut trust for one side of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates one side of a one-way, outgoing, shortcut trust. Although one side of a trust will be created successfully, the new trust will not function until the administrator for the reciprocal domain uses his or her credentials to create the second side of the trust. If you have administrative credentials for both domains that are involved in the trust, you can use the procedure Create a one-way, outgoing, shortcut trust for both sides of the trust to create both sides of the trust in one simultaneous operation.

A one-way, outgoing, shortcut trust allows resources in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to be accessed more quickly by users in another domain (which is nested within another domain tree) in your forest. For example, if you are the administrator of marketing.tailspintoys.com and resources in that domain need to be accessed by users in the sales.wingtiptoys.com domain (which is a child domain of the wingtiptoys.com tree root domain), you can use this procedure to establish one side of the relationship so that users in the sales.wingtiptoys.com domain can more quickly access resources in the marketing.tailspintoys.com domain.

You can create this shortcut trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a shortcut trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, outgoing, shortcut trust for one side of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: outgoing**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **This domain only**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Trust Password** page, type the trust password twice, and then click **Next**.
9. On the **Trust Selections Complete** page, review the results, and then click **Next**.
10. On the **Trust Creation Complete** page, review the results, and then click **Next**.
11. On the **Confirm Outgoing Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.

    - If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the specified domain.

12. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the domain administrator for the specified domain or specified forest must follow the procedure Create a one-way, incoming, shortcut trust for one side of the trust, using his or her administrative credentials and the exact same trust password that was used during this procedure. |

# Create a one-way, outgoing, shortcut trust for both sides of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates both sides of a one-way, outgoing, shortcut trust, and it requires that you have administrative credentials for your domain as well as for the reciprocal domain. If you have administrative credentials only for your domain, you can use the procedure Create a one-way, outgoing, shortcut trust for one side of the trust to create your side of the trust. Then, have the administrator for the reciprocal domain create a one-way, incoming, shortcut trust from his or her domain.

A one-way, outgoing, shortcut trust allows resources in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to be accessed more quickly by users in another domain (which is nested within another domain tree) in your forest. For example, if you are the administrator of marketing.tailspintoys.com and resources in that domain need to be accessed by users in the sales.wingtiptoys.com domain (which is a child domain of the wingtiptoys.com tree root domain), you can use this procedure to establish one side of the relationship so that users in the sales.wingtiptoys.com domain can more quickly access resources in the marketing.tailspintoys.com domain.

You can create this shortcut trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a shortcut trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, outgoing, shortcut trust for both sides of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: outgoing**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **User Name and Password** page, type the user name and password for the appropriate administrator in the specified domain.
9. On the **Trust Selections Complete** page, review the results, and then click **Next**.
10. On the **Trust Creation Complete** page, review the results, and then click **Next**.
11. On the **Confirm Outgoing Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.

    - If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate

administrative credentials from the specified domain.

12. On the **Completing the New Trust Wizard** page, click **Finish**.

# Create a two-way, shortcut trust for one side of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates one side of a two-way, shortcut trust. Although one side of a trust will be created successfully, the new trust will not function until the administrator for the reciprocal domain uses his or her credentials to create the second side of the trust. If you have administrative credentials for both domains that are involved in the trust, you can use the procedure Create a two-way, shortcut trust for both sides of the trust to create both sides of the trust in one simultaneous operation.

A two-way, shortcut trust allows users in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) and users in the reciprocal domain to more quickly access resources in either domain (when both domains are separated by a domain tree) in your forest.

You can create this shortcut trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a shortcut trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a two-way, shortcut trust for one side of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **Two-way**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **This domain only**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Trust Password** page, type the trust password twice, and then click **Next**.
9. On the **Trust Selections Complete**page, review the results, and then click **Next**.
10. On the **Trust Creation Complete** page, review the results, and then click **Next**.
11. On the **Confirm Outgoing Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.

    - If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the specified domain.

12. On the **Confirm Incoming Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.

- If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

13. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the domain administrator for the specified domain must follow this same procedure using his or her administrative credentials and the exact same trust passwordthat was used during this procedure. |

# Create a two-way, shortcut trust for both sides of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates both sides of a two-way, shortcut trust, and it requires you to have administrative credentials for your domain as well as for the reciprocal domain. If you have administrative credentials only for your domain, you can use the procedure Create a two-way, shortcut trust for one side of the trust to create your side of the trust. Then, have the administrator for the reciprocal domain create a two-way, shortcut trust from his or her domain.

A two-way, shortcut trust allows users in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) and users in the reciprocal domain to more quickly access resources in either domain (when both domains are separated by a domain tree) in your forest.

You can create this shortcut trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a shortcut trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a two-way, shortcut trust for both sides of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **Two-way**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **User Name and Password** page, type the user name and password for the appropriate administrator in the specified domain.
9. On the **Trust Selections Complete**page, review the results, and then click **Next**.
10. On the **Trust Creation Complete** page, review the results, and then click **Next**.
11. On the **Confirm Outgoing Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.

    - If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the specified domain.

12. On the **Confirm Incoming Trust** page, do one of the following:

- If you do not want to confirm this trust, click **No, do not confirm the incoming trust.**

- If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

13. On the **Completing the New Trust Wizard** page, click **Finish**.

# Creating Forest Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a Windows Server 2003 forest, you can link two disjoined Windows Server 2003 forests together to form a one-way or two-way, transitive trust relationship. You can use a two-way, forest trust to form a transitive trust relationship between every domain in both forests.

For more information about forest trusts, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

**Task requirements**

The following requirements, features, or settings are necessary to create forest trusts successfully:

- You can create a forest trust only between two Windows Server 2003 forests; forest trusts cannot be extended implicitly to a third Windows Server 2003 forest.
- To create a forest trust, you must set the forest functional level for both of the Windows Server 2003 forests that are involved in the trust relationship to Windows Server 2003. For more information about functional levels, see Active Directory Functional Levels Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41698).
- To create a forest trust successfully, you must set up your Domain Name System (DNS) environment properly. If there is a root DNS server that you can make the root DNS server for the DNS namespaces of both forests, make it the root server by ensuring that the root zone contains delegations for each of the DNS namespaces. Also, update the root hints of all DNS servers with the new root DNS server.
- If there is no shared root DNS server and the root DNS servers for each forest DNS namespace are running a member of the Windows Server 2003 family, configure DNS conditional forwarders in each DNS namespace to route queries for names in the other namespace.
- If there is no shared root DNS server and the root DNS servers for each forest DNS namespace are not running a member of the Windows Server 2003 family, configure DNS secondary zones in each DNS namespace to route queries for names in the other namespace. For more information about configuring DNS to work with Active Directory, see DNS Support for Active Directory Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41699).

You can use either of the following tools to perform the procedures for this task:

- Active Directory Domains and Trusts
- Netdom.exe

For more information about how to use the Netdom command-line tool to create a forest trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

| **Note** |
| --- |
| If you have the appropriate administrative credentials for each forest, you can create both sides of a forest trust at the same time. To create both sides of the forest trust, follow the appropriate procedure below that contains the words “for both sides of the trust” in the title. For example, the procedure “Create a one-way, incoming, forest trust for both sides of the trust” explains how to configure both sides of the trust. For more information about how the “both sides of the trust” option works, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages. |

You can create a forest trust by using any one of the following procedures, depending on the requirements of your

organization and the administrative credentials that you have when you create the trust:

- Create a one-way, incoming, forest trust for one side of the trust
- Create a one-way, incoming, forest trust for both sides of the trust
- Create a one-way, outgoing, forest trust for one side of the trust
- Create a one-way, outgoing, forest trust for both sides of the trust
- Create a two-way, forest trust for one side of the trust
- Create a two-way, forest trust for both sides of the trust

# Create a one-way, incoming, forest trust for one side of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates one side of a one-way, incoming, forest trust. Although one side of a trust will be created successfully, the new trust will not function until the administrator for the reciprocal forest uses his or her credentials to create the second side of the trust. If you have administrative credentials for both forests that are involved in the trust, you can use the procedure Create a one-way, incoming, forest trust for both sides of the trust to create both sides of the trust in one simultaneous operation.

A one-way, incoming, forest trust allows users in your Windows Server 2003 forest (the forest that you are logged on to at the time that you run the New Trust Wizard) to access resources in another Windows Server 2003 forest. For example, if you are the administrator of the wingtiptoys.com forest and users in that forest need to access resources in the tailspintoys.com forest, you can use this procedure to establish one side of the relationship so that users in your forest can access resources in any of the domains that make up the tailspintoys.com forest.

You can create this forest trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a forest trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory. If you are a member of the Incoming Forest Trust Builders group, you can create one-way, incoming, forest trusts to your forest. For more information about the Incoming Forest Trust Builders group, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

**To create a one-way, incoming, forest trust for one side of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **Forest trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: incoming**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **This domain only**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Trust Password** page, type the trust password twice, and then click **Next**.
9. On the **Trust Selections Complete** page, review the results, and then click **Next**.
10. On the **Trust Creation Complete** page, review the results, and then click **Next**.
11. On the **Confirm Incoming Trust** page, do one of the following:
    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.
    - If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate

administrative credentials from the specified domain.

12. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the domain administrator for the specified domain (the forest root domain in the specified forest) must complete the procedure Create a one-way, outgoing, forest trust for one side of the trust, using their administrative credentials and the exact same trust passwordthat was used during this procedure. |

# Create a one-way, incoming, forest trust for both sides of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates both sides of a one-way, incoming, forest trust, and it requires you to have administrative credentials for your forest as well as for the reciprocal forest. If you have administrative credentials only for your forest, you can use the procedure Create a one-way, incoming, forest trust for one side of the trust to create your side of the trust. Then, have the administrator for the reciprocal forest create a one-way, outgoing forest trust from his or her domain.

A one-way, incoming, forest trust allows users in your Windows Server 2003 forest (the forest that you are logged on to at the time that you run the New Trust Wizard) to access resources in another Windows Server 2003 forest. For example, if you are the administrator of the wingtiptoys.com forest and users in that forest need to access resources in the tailspintoys.com forest, you can use this procedure to establish one side of the relationship so that users in your forest can access resources in any of the domains that make up the tailspintoys.com forest.

You can create this forest trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a forest trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory. If you are a member of the Incoming Forest Trust Builders group, you can create one-way, incoming, forest trusts to your forest. For more information about the Incoming Forest Trust Builders group, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

**To create a one-way, incoming, forest trust for both sides of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **Forest trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: incoming**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **User Name and Password** page, type the user name and password for the appropriate administrator in the specified domain.
9. On the **Outgoing Trust Authentication Level--Specified Forest** page, do one of the following, and then click **Next**:
   - Click **Forest-wide authentication**.
   - Click **Selective authentication**.
10. On the **Trust Selections Complete** page, review the results, and then click **Next**.
11. On the **Trust Creation Complete** page, review the results, and then click **Next**.

12. On the **Confirm Incoming Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.

    - If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

13. On the **Completing the New Trust Wizard** page, click **Finish**.

# Create a one-way, outgoing, forest trust for one side of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates one side of a one-way, outgoing, forest trust. Although one side of a trust will be created successfully, the new trust will not function until the administrator for the reciprocal forest uses his or her credentials to create the second side of the trust. If you have administrative credentials for both forests that are involved in the trust, you can use the procedure Create a one-way, outgoing, forest trust for both sides of the trust to create both sides of the trust in one simultaneous operation.

A one-way, outgoing, forest trust allows resources in your Windows Server 2003 forest (the forest that you are logged on to at the time that you run the New Trust Wizard) to be accessed by users in another Windows Server 2003 forest. For example, if you are the administrator of the wingtiptoys.com forest and resources in that forest need to be accessed by users in the tailspintoys.com forest, you can use this procedure to establish one side of the relationship so that users in the tailspintoys.com forest can access resources in any of the domains that make up the wingtiptoys.com forest.

You can create this forest trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a forest trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory. If you are a member of the Incoming Forest Trust Builders group, you can create one-way, incoming, forest trusts to your forest. For more information about the Incoming Forest Trust Builders group, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

**To create a one-way, outgoing, forest trust for one side of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **Forest trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: outgoing**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **This domain only**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Outgoing Trust Authentication Level** page, do one of the following, and then click **Next**:
   - Click **Forest-wide authentication**.
   - Click **Selective authentication**.
9. On the **Trust Password** page, type the trust password twice, and then click **Next**.
10. On the **Trust Selections Complete** page, review the results, and then click **Next**.

11. On the **Trust Creation Complete** page, review the results, and then click **Next**.
12. On the **Confirm Outgoing Trust** page, do one of the following:
    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.
    - If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the specified domain.
13. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the domain administrator for the specified domain (the forest root domain in the specified forest) must follow the procedure Create a one-way, incoming, forest trust for one side of the trust, using his or her administrative credentials and the exact same trust password that was used during this procedure. |

# Create a one-way, outgoing, forest trust for both sides of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates both sides of a one-way, outgoing, forest trust, and it requires you to have administrative credentials for your forest as well as for the reciprocal forest. If you have administrative credentials only for your domain, you can use the procedure Create a one-way, outgoing, forest trust for one side of the trust to create your side of the trust. Then, have the administrator for the reciprocal forest create a one-way, incoming, external trust from his or her forest.

A one-way, outgoing, forest trust allows resources in your Windows Server 2003 forest (the forest that you are logged on to at the time that you run the New Trust Wizard) to be accessed by users in another Windows Server 2003 forest. For example, if you are the administrator of the wingtiptoys.com forest and resources in that forest need to be accessed by users in the tailspintoys.com forest, you can use this procedure to establish one side of the relationship so that users in the tailspintoys.com forest can access resources in any of the domains that make up the wingtiptoys.com forest.

You can create this forest trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a forest trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory. If you are a member of the Incoming Forest Trust Builders group, you can create one-way, incoming, forest trusts to your forest. For more information about the Incoming Forest Trust Builders group, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

**To create a one-way, outgoing, forest trust for both sides of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **Forest trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **One-way: outgoing**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **User Name and Password** page, type the user name and password for the appropriate administrator in the specified domain.
9. On the **Outgoing Trust Authentication Level--Local Forest** page, do one of the following, and then click **Next**:
   - Click **Forest-wide authentication**.
   - Click **Selective authentication**.
10. On the **Trust Selections Complete**page, review the results, and then click **Next**.
11. On the **Trust Creation Complete** page, review the results, and then click **Next**.

12. On the **Confirm Outgoing Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust.** Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.

    - If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the specified domain.

13. On the **Completing the New Trust Wizard** page, click **Finish**.

# Create a two-way, forest trust for one side of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates one side of a two-way, forest trust. Although one side of a trust will be created successfully, the new trust will not function until the administrator for the reciprocal forest uses his or her credentials to create the second side of the trust. If you have administrative credentials for both forests that are involved in the trust, you can use the procedure Create a two-way, forest trust for both sides of the trust to create both sides of the trust in one simultaneous operation.

A two-way, forest trust allows users in your forest (the forest that you are logged on to at the time that you run the New Trust Wizard) and users in the reciprocal forest to access resources in any of the domains in either of the two forests.

You can create this forest trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a forest trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory. If you are a member of the Incoming Forest Trust Builders group, you can create one-way, incoming, forest trusts to your forest. For more information about the Incoming Forest Trust Builders group, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

**To create a two-way, forest trust for one side of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **Forest trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **Two-way**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **This domain only**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Outgoing Trust Authentication Level** page, do one of the following, and then click **Next**:
   - Click **Forest-wide authentication**.
   - Click **Selective authentication**.
9. On the **Trust Password** page, type the trust password twice, and then click **Next**.
10. On the **Trust Selections Complete**page, review the results, and then click **Next**.
11. On the **Trust Creation Complete** page, review the results, and then click **Next**.
12. On the **Confirm Outgoing Trust** page, do one of the following:
    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.

- If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the specified domain.

13. On the **Confirm Incoming Trust** page, do one of the following:

    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust.**

    - If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.

14. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the domain administrator in the specified forest must follow this same procedure, using his or her administrative credentials and the exact same trust passwordthat was used during this procedure. |

# Create a two-way, forest trust for both sides of the trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure creates both sides of a two-way, forest trust, and it requires you to have administrative credentials for your forest as well as for the reciprocal forest. If you have administrative credentials only for your forest, you can use the procedure Create a two-way, forest trust for one side of the trust to create your side of the trust. Then, have the administrator for the reciprocal forest create a one-way, outgoing forest trust from his or her forest.

A two-way, forest trust allows users in your forest (the forest that you are logged on to at the time that you run the New Trust Wizard) and users in the reciprocal forest to access resources in any of the domains in either of the two forests.

You can create this forest trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a forest trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory. If you are a member of the Incoming Forest Trust Builders group, you can create one-way, incoming, forest trusts to your forest. For more information about the Incoming Forest Trust Builders group, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

**To create a two-way, forest trust for both sides of the trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **Forest trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **Two-way**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.

   For more information about the selections that are available on the **Sides of Trust** page, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages.
8. On the **User Name and Password** page, type the user name and password for the appropriate administrator in the specified domain.
9. On the **Outgoing Trust Authentication Level--Local Forest** page, do one of the following, and then click **Next**:
   - Click **Forest-wide authentication**.
   - Click **Selective authentication**.
10. On the **Outgoing Trust Authentication Level--Specified Forest** page, do one of the following, and then click **Next**:
    - Click **Forest-wide authentication**.
    - Click **Selective authentication**.

11. On the **Trust Selections Complete** page, review the results, and then click **Next**.
12. On the **Trust Creation Complete** page, review the results, and then click **Next**.
13. On the **Confirm Outgoing Trust** page, do one of the following:
    - If you do not want to confirm this trust, click **No, do not confirm the outgoing trust**. Note that if you do not confirm the trust at this stage, the secure channel will not be established until the first time the trust is used by users.
    - If you want to confirm this trust, click **Yes, confirm the outgoing trust**, and then supply the appropriate administrative credentials from the specified domain.
14. On the **Confirm Incoming Trust** page, do one of the following:
    - If you do not want to confirm this trust, click **No, do not confirm the incoming trust**.
    - If you want to confirm this trust, click **Yes, confirm the incoming trust**, and then supply the appropriate administrative credentials from the specified domain.
15. On the **Completing the New Trust Wizard** page, click **Finish**.

# Creating Realm Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can create a realm trust to form a one-way or two-way, nontransitive or transitive trust with non-Windows Kerberos realms in your organization. You can create the trust when you log on to the domain, or you can use the **Run as** command to create the trust for a different domain.

For more information about realm trusts, see "How Domain and Forest Trusts Work" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

**Task requirements**

You can use either of the following tools to perform the procedures for this task:

- Active Directory Domains and Trusts
- Netdom.exe

For more information about how to use the Netdom command-line tool to create a realm trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

| **Note** |
| --- |
| The New Trust Wizard in Active Directory Domains and Trusts does not support the creation of both sides of a realm trust at the same time. For more information about how the “both sides of the trust” option works, see the section "Sides of Trust" in Appendix: New Trust Wizard Pages. |

You can create a realm trust by using any of the following procedures, depending on the requirements of your organization and the administrative credentials that you have when you create the trust:

- Create a one-way, incoming, realm trust
- Create a one-way, outgoing, realm trust
- Create a two-way, realm trust

# Create a one-way, incoming, realm trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A one-way, incoming realm trust allows users in your Windows Server 2003 domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to access resources in the Kerberos realm. For example, if you are the administrator of the sales.wingtiptoys.com domain and users in that domain need to access resources in the Kerberos realm, you can use this procedure to establish a relationship so that users in the sales.wingtiptoys.com domain can access resources in the Kerberos realm.

You can create this realm trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a realm trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, incoming, realm trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **Realm trust**, and then click **Next**.
6. On the **Transitivity of Trust** page, do one of the following:
   - To form a trust relationship with the domain and the specified realm only, click **Nontransitive**, and then click **Next**.
   - To form a trust relationship with the domain and the specified realm and all trusted realms, click **Transitive**, and then click **Next**.
7. On the **Direction of Trust** page, click **One-way: incoming**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Trust Password** page, type the trust password twice, and then click **Next**.
9. On the **Trust Selections Complete** page, review the results, and then click **Next**.
10. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the administrator of the realm must complete the trust, using his or her administrative credentials and the exact same trust password that was used during this procedure. |

# Create a one-way, outgoing, realm trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A one-way, outgoing realm trust allows resources in your Windows Server 2003 domain (the domain that you are logged on to at the time that you run the New Trust Wizard) to be accessed by users in the Kerberos realm. For example, if you are the administrator of the sales.wingtiptoys.com domain and resources in that domain need to be accessed by users in the Kerberos realm, you can use this procedure to establish a relationship so that users in the Kerberos realm can access resources in the sales.wingtiptoys.com domain.

You can create this realm trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a realm trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a one-way, outgoing, realm trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **Realm trust**, and then click **Next**.
6. On the **Transitivity of Trust** page, do one of the following:
   - To form a trust relationship with the domain and the specified realm only, click **Nontransitive**, and then click **Next**.
   - To form a trust relationship with the domain and the specified realm and all trusted realms, click **Transitive**, and then click **Next**.
7. On the **Direction of Trust** page, click **One-way: outgoing**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Trust Password** page, type the trust password twice, and then click **Next**.
9. On the **Trust Selections Complete** page, review the results, and then click **Next**.
10. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the administrator of the realm must complete the trust, using his or her administrative credentials and the exact same trust password that was used during this procedure. |

# Create a two-way, realm trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A two-way, realm trust allows users in your domain (the domain that you are logged on to at the time that you run the New Trust Wizard) and users in the specified Kerberos realm to access resources in either the domain or the Kerberos realm.

You can create this realm trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a realm trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To create a two-way, realm trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **Realm trust**, and then click **Next**.
6. On the **Transitivity of Trust** page, do one of the following:
   - To form a trust relationship with the domain and the specified realm only, click **Nontransitive**, and then click **Next**.
   - To form a trust relationship with the domain and the specified realm and all trusted realms, click **Transitive**, and then click **Next**.
7. On the **Direction of Trust** page, click **Two-way**, and then click **Next**.

   For more information about the selections that are available on the **Direction of Trust** page, see the section "Direction of Trust" in Appendix: New Trust Wizard Pages.
8. On the **Trust Password** page, type the trust password twice, and then click **Next**.
9. On the **Trust Selections Complete** page, review the results, and then click **Next**.
10. On the **Completing the New Trust Wizard** page, click **Finish**.

| **Note** |
| --- |
| For this trust to function, the administrator of the realm must complete the trust, using his or her administrative credentials and the exact same trust password that was used during this procedure. |

# Configuring Domain and Forest Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can remove manually created trusts, but you cannot remove the default, two-way, transitive trusts between domains in a forest. If you remove manually created trusts, it is particularly important to verify that you successfully removed the trusts if you are planning to re-create them.

The following tasks for removing a manually created trust are described in this objective:

- Validating and removing trusts
- Modifying Name Suffix Routing Settings

# Validating and removing trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After a trust has been established, you may need to verify that it is working as designed, or that communications over the trust are working, by using Active Directory tools to validate the trust's connectivity. It may also be necessary to remove an existing, manually created trust when connectivity between two domains is no longer necessary.

**Task requirements**

You can use either of the following tools to perform the procedures for this task:

- Active Directory Domains and Trusts
- Netdom.exe

For more information about how to use the Netdom command-line tool to create a realm trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

To complete this task, perform the following procedures:

- Validate a trust
- Remove a manually created trust

# Validate a trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can validate all trusts that are made between domains, but you cannot validate realm trusts.

You can validate a trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to create a realm trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To validate a trust**

- Using the Windows interface
- Using the command line

**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that contains the trust that you want to validate, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, click the trust to be validated, and then click **Properties**.
4. Click **Validate**.
5. Do one of the following, and then click **OK**:
   - Click **No, do not validate the incoming trust**.

     If you click this option, it is recommended that you repeat this procedure for the reciprocal domain.
   - Click **Yes, validate the incoming trust**.

     If you click this option, you must type a user account and password with administrative credentials for the reciprocal domain.

**Using the command line**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **netdom trust** *TrustingDomainName* **/d:** *TrustedDomainName* **/verify**

| Term | Definition |
|---|---|
| *TrustingDomainName* | Specifies the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the trusting domain in the trust that is being created. |
| *TrustedDomainName* | Specifies the DNS name (or NetBIOS name) of the domain that will be trusted in the trust that is being created. |

# Remove a manually created trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is possible to remove manually created shortcut trusts, external trusts, realm trusts, or forest trusts. It is not possible to remove default, two-way, transitive trusts between domains in a forest. It is particularly important to verify that you successfully remove trusts if you are planning to re-create them.

You can remove a manually created trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about the Netdom command-line tool, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To remove a manually created trust**

- Using the Windows interface
- Using a command prompt

**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that contains the trust that you want to remove, and then click **Properties**.
3. Click the **Trusts** tab.
4. In either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, click the trust to be removed, and then click **Remove**.
5. Do one of the following, and then click **OK**:
   - Click **No, remove the trust from the local domain only**.

     If you click this option, it is recommended that you repeat this procedure for the reciprocal domain.
   - Click **Yes, remove the trust from both the local domain and the other domain**.

     If you click this option, you must type a user account and password with administrative credentials for the reciprocal domain.

**Using the command line**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **netdom trust** *TrustingDomainName* **/d:** *TrustedDomainName* **/remove/UserD:***User* **/PasswordD:***

| Term | Definition |
|---|---|
| *TrustingDomainName* | Specifies the Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the trusting domain in the trust that is being created. |

| | |
|---|---|
| *TrustedDomainName* | Specifies the DNS name (or NetBIOS name) of the domain that will be trusted in the trust that is being created. |

**Note**

If you are using Netdom to remove a realm trust, you must add the **/force** option to the end of the command (after **/remove**) to remove the trust successfully.

# Modifying Name Suffix Routing Settings

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Name suffix routing is a mechanism that is used to manage how authentication requests are routed across Windows Server 2003 forests that are joined by forest trusts. To simplify the administration of authentication requests, when a forest trust is created, all unique name suffixes are routed by default. A unique name suffix is a name suffix within a forest, such as a user principal name (UPN) suffix, service principal name (SPN) suffix, or Domain Name System (DNS) forest or domain tree name that is not subordinate to any other name suffix. For example, the DNS forest name fabrikam.com is a unique name suffix within the fabrikam.com forest. For more information about name suffix routing, see Routing name suffixes across forests on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35414).

| **Note** |
| --- |
| You cannot enable a name suffix that is in conflict. If the conflict is with a local UPN name suffix, you must remove the local UPN name suffix before you can enable the routing name. If the conflict is with a name that is claimed by another trust partner, you must disable the name in the other trust before it can be enabled for this trust. |

**Task requirements**

You can use either of the following tools to perform the procedures for this task:

- Active Directory Domains and Trusts
- Netdom.exe

For more information about how to use the Netdom command-line tool to create a realm trust, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

To complete this task, perform any of the following procedures:

- Modify the routing status of a name suffix
- Enable or disable an existing name suffix for routing
- Exclude name suffixes from routing to local forests

# Modify the routing status of a name suffix

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can change the routing status (enabled or disabled) of a name suffix by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to modify name suffix routing settings, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To modify the routing status of a name suffix**
**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain that you want to administer, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)**or **Domains that trust this domain (incoming trusts)**, click the forest trust that you want to administer, and then click **Properties**.
4. On the **Name Suffix Routing** tab, under **Name suffixes in the x.x forest**, click the suffix for which you want to modify routing status, and then click **Edit**.
5. In **Existing name suffixes in x.x**, click the suffix that you want to modify, and then click **Enable** or **Disable**.

# Enable or disable an existing name suffix for routing

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to prevent authentication requests for specific name suffixes from being routed to a forest, or you can use this procedure to allow authentication requests for specific name suffixes to be routed to a forest. You can enable or disable an existing name suffix for routing by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to modify name suffix routing settings, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

| **Note** |
| --- |
| When you disable a name suffix, all children of that Domain Name System (DNS) name will also be disabled. |

### Administrative credentials

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To enable or disable an existing name suffix for routing**
**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to administer, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, click the forest trust that you want to administer, and then click **Properties**.
4. Click the **Name Suffix Routing** tab, and then, under **Name suffixes in the x.x forest,** do one of the following:
   - To enable a name suffix, click the suffix that you want to enable, and then click **Enable**. If the **Enable** button is unavailable, the name suffix is already enabled.
   - To disable a name suffix, click the suffix that you want to disable, and then click **Disable**. If the **Disable** button is unavailable, the name suffix is already disabled.

# Exclude name suffixes from routing to local forests

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can exclude existing name suffixes from routing to local forests by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to modify name suffix routing settings, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

| **Note** |
| --- |
| When you exclude a name suffix, all children of that Domain Name System (DNS) name will also be excluded. |

### Administrative credentials

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To exclude name suffixes from routing to local forests**
**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to administer, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, click the forest trust that you want to administer, and then click **Properties**.
4. On the **Name Suffix Routing** tab, under **Name suffixes in the x.x forest**, click the unique name suffix to exclude from routing, and then click **Edit**.
5. In **Name suffixes to exclude from routing to x.x**, click **Add**, type a DNS name suffix that is subordinate to the unique name suffix, and then click **OK**.

# Securing Domain and Forest Trusts

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you create a new trust in an existing Active Directory forest, all communications over that trust are tightly secured. However, when you create a trust between your domain and another domain outside your forest, there are certain security issues involved. For example, you might need to configure security identifier (SID) filtering to deny one domain the right to provide credentials for another domain. You can enable or disable SID filtering for external trusts or forest trusts.

The following tasks for securing domain and forest trusts are described in this objective:

- Configuring SID Filtering Settings
- Configuring Selective Authentication Settings

For more information about how the security settings for domain and forest trusts work, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

# Configuring SID Filtering Settings

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Security principals in Active Directory have an attribute, called SID history, to which domain administrators can add users' old security identifiers (SIDs). This is useful during Active Directory migrations because administrators do not need to modify access control lists (ACLs) on large numbers of resources and users can use their old SIDs to access resources. However, under some circumstances it is possible for attackers or rogue administrators that have compromised a domain controller in a trusted domain to use the SID history attribute (sIDHistory) to associate SIDs with new user accounts, granting themselves unauthorized rights. To help prevent this type of attack, Windows Server 2003 automatically enables SID filter quarantining on all external trusts that are created by a Windows Server 2003 domain controller. External trusts that are created using domain controllers running Windows 2000 Server with Service Pack 3 (SP3) or earlier must be manually configured to enable SID filter quarantining.

| **Note** |
| --- |
| You cannot turn off the default behavior in Windows Server 2003 that enables SID filter quarantining for newly created external trusts. External trusts that are created from domain controllers running Windows 2000 Server with SP3 or earlier do not enforce SID filter quarantining by default. |

You can use SID filter quarantining to filter out migrated SIDs that are stored in SID history from specific domains. For example, where an external trust relationship exists such that the one domain, Contoso (running Windows 2000 Server domain controllers), trusts another domain, Cpandl (also running Windows 2000 Server domain controllers), an administrator of the Contoso domain can manually apply SID filter quarantining to the Cpandl domain, which allows all SIDs with a domain SID from the Cpandl domain to pass but all other SIDs (such as those from migrated SIDs that are stored in SID history) to be discarded.

| **Note** |
| --- |
| Do not apply SID filter quarantining to trusts within a forest that is not using the Windows Server 2003 forest functional level, because doing so removes SIDs that are required for Active Directory replication. If the forest functional level is Windows Server 2003 and quarantining is applied between two domains within a forest, a user in the quarantined domain with universal group memberships in other domains in the forest may not be able to access resources in nonquarantined domains, because the group memberships from those domains are filtered when resources are accessed across the trust relationship. Likewise, SID filter quarantining should not be applied to forest trusts. |

To further secure your forest, consider enabling SID filter quarantining on all existing external trusts that are created by domain controllers running Windows 2000 Server SP3 or earlier. You can do this by using Netdom.exe to enable SID filter quarantining on existing external trusts or by recreating these external trusts from a domain controller running Windows Server 2003 or Windows 2000 Server with Service Pack 4 (SP4) or later. For more information about how to enable SID filtering on trusts that are created by Windows 2000 Server domain controllers, see the Windows 2000 Active Directory Operations Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=18545).

For more information about how SID filtering works, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

**Task requirements**

You can use either of the following tools to perform the procedures for this task:

- Active Directory Domains and Trusts
- Netdom.exe

For more information about how to use the Netdom command-line tool to configure SID filtering settings, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

To complete this task, perform the following procedures:

- Disable SID filter quarantining
- Reapply SID filter quarantining

# Verify SID filter quarantining

Updated: November 4, 2009

Applies To: Windows Server 2003, Windows Server 2008, Windows Server 2008 R2

You can use this procedure to verify the status of security identifier (SID) filter quarantining for an external trust or a forest trust using Netdom.exe.

For more information about the **netdom** command-line tool, see Netdom Overview (http://go.microsoft.com/fwlink/?LinkId=111537).

For more information about how SID filtering works, see Security Considerations for Trusts (http://go.microsoft.com/fwlink/?LinkID=111846)

Membership in **Domain Admins** or **Enterprise Admins** in Active Directory Domain Services (AD DS), or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To verify SID filter quarantining from the trusting domain**

1. Open a command prompt.
2. At the command prompt, type the following syntax, and then press ENTER:

```
netdom trust <TrustingDomainName> /domain: <TrustedDomainName> /quarantine /userD:
<domainadministratorAcct> /passwordD: <domainadminpwd>
```

| Note |
| --- |
| For a description of the placeholder values, see Placeholder reference table. |

**To verify SID filter quarantining from the trusting forest**

1. Open a command prompt.
2. At the command prompt, type the following syntax, and then press ENTER:

```
netdom trust <TrustingDomainName> /domain: <TrustedDomainName> /enablesidhistory /userD:
<domainadministratorAcct> /passwordD: <domainadminpwd>
```

| Note |
| --- |
| For a description of the placeholder values, see Placeholder reference table. |

**To verify SID filter quarantining from the trusted domain**

1. Open a command prompt.
2. At the command prompt, type the following syntax, and then press ENTER:

```
netdom trust <TrustingDomainName> /userO:<domainadministratorAcct> /passwordO:<domainadminpwd> /domain:
<TrustedDomainName> /quarantine
```

| Note |
| --- |
| For a description of the placeholder values, see Placeholder reference table. |

**To verify SID filter quarantining from the trusted forest**

1. Open a command prompt.
2. At the command prompt, type the following syntax, and then press ENTER:

```
netdom trust <TrustingDomainName> /userO:<domainadministratorAcct> /passwordO:<domainadminpwd> /domain:
<TrustedDomainName> /enablesidhistory
```

| Note |
| --- |
| For a description of the placeholder values, see Placeholder reference table. |

**Placeholder reference table**

The following table describes the placeholder values in the syntax for the `netdom` command.

| Value | Description |
| --- | --- |
| <TrustingDomainName> | The Domain Name System (DNS) name (or NetBIOS name) of the trusting domain in the trust that is being created. |
| <TrustedDomainName> | The DNS name (or NetBIOS name) of the domain that will be trusted in the trust that is being created. |
| <domainadministratorAcct> | The user account name with the appropriate administrator credentials to modify the trust. |
| <domainadminpwd> | The password of the user account in *domainadministratorAcct*. |

# Disable SID filter quarantining

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although it is not recommended, you can disable security identifier (SID) filter quarantining for an external trust by using the Netdom.exe tool. You should consider disabling SID filter quarantining only in the following situations:

- You have an equally high level of confidence in the administrators who have physical access to domain controllers in the trusted domain and the administrators with such access in the trusting domain.
- You have a strict requirement to assign universal groups to resources in the trusting domain, even when those groups were not created in the trusted domain.
- Users have been migrated to the trusted domain with their SID histories preserved, and you want to grant those users access to resources in the trusting domain based on the **SIDHistory** attribute.

For more information about how SID filtering works, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

You can disable SID filter quarantining by using the Netdom command-line tool. For more information about the Netdom command-line tool, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To disable SID filter quarantining**

1. To disable SID filter quarantining for the trusting domain, open a Command Prompt.
2. Type the following command, and then press ENTER:

   **Netdom trust** *TrustingDomainName* **/domain:** *TrustedDomainName* **/quarantine:No /userD:** *domainadministratorAcct* **/passwordD:** *domainadminpwd*

| Value | Description |
|---|---|
| *TrustingDomainName* | The Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the trusting domain in the trust that is being created. |
| *TrustedDomainName* | The DNS name (or NetBIOS name) of the domain that will be trusted in the trust that is being created. |
| *domainadministratorAcct* | The user account name with the appropriate administrator credentials to modify the trust. |
| *domainadminpwd* | The password of the user account in *domainadministratorAcct*. |

**Note**

You can enable or disable SID filter quarantining only from the trusting side of the trust. If the trust is a two-way trust, you can also disable SID filter quarantining in the trusted domain by using the domain administrator's credentials for the trusted domain and reversing the *TrustingDomainName* and *TrustedDomainName* values in the command-line syntax.

# Reapply SID filter quarantining

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can reapply security identifier (SID) filter quarantining to an external or forest trust that has had SID filter quarantining disabled. By default, Windows Server 2003 automatically enables SID filter quarantining on all external trusts that are created by a Windows Server 2003 domain controller. For more information about how SID filter quarantining works, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

You can reapply SID filter quarantining by using the Netdom command-line tool. For more information about the Netdom command-line tool, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

### Administrative credentials

To complete this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

#### To reapply SID filter quarantining

1. To reapply SID filter quarantining for the trusting domain, open a Command Prompt.
2. Type the following syntax, and then press ENTER:

   **Netdom trust** *TrustingDomainName* **/domain:** *TrustedDomainName* **/quarantine:Yes /userD:** *domainadministratorAcct* **/passwordD:** *domainadminpwd*

| Term | Definition |
|---|---|
| *TrustingDomainName* | The Domain Name System (DNS) name (or network basic input/output system (NetBIOS) name) of the trusting domain in the trust that is being created. |
| *TrustedDomainName* | The DNS name (or NetBIOS name) of the domain that will be trusted in the trust that is being created. |
| *domainadministratorAcct* | The user account name with the appropriate administrator credentials to modify the trust. |
| *domainadminpwd* | The password of the user account in *domainadministratorAcct*. |

# Configuring Selective Authentication Settings

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Trusts that are created between Windows Server 2003 forests can use legacy authentication settings (settings that were used in Windows 2000 Server) or selective authentication. Selective authentication is a security setting that can be enabled on external trusts and forest trusts between Windows Server 2003 forests. Selective authentication provides Active Directory administrators who manage a trusting forest more control over which groups of users in a trusted forest can access shared resources in the trusting forest. Because creating an external trust or forest trust provides a pathway for all authentication requests between the forests, this increased control is especially important when administrators need to grant access to shared resources in their organization's forest to a limited set of users in another organization's forest.

For more information about how selective authentication settings work, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

**Task requirements**

You can use either of the following tools to perform the procedures for this task:

- Active Directory Domains and Trusts
- Netdom.exe

For more information about how to use the Netdom command-line tool to configure selective authentication settings, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

To complete this task, perform the following procedures:

- Enable selective authentication over an external trust
- Enable selective authentication over a forest trust
- Enable domain-wide authentication over an external trust
- Enable forest-wide authentication over a forest trust
- Grant the Allowed to Authenticate permission on computers in the trusting domain or forest

# Enable selective authentication over an external trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Selective authentication over an external trust restricts access to only those users in a trusted domain who have been explicitly given authentication permissions to computer objects (resource computers) that reside in the trusting domain. To explicitly give authentication permissions to computer objects in the trusting domain to certain users, administrators must grant those users the **Allowed to Authenticate** permission in Active Directory. For more information, see Grant the Allowed to Authenticate permission on computers in the trusting domain or forest. For more information about how selective authentication works, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

You can enable selective authentication over an external trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to configure selective authentication settings, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To enable selective authentication over an external trust**
**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to administer, and then click **Properties**.
3. On the **Trusts** tab, under **Domains trusted by this domain (outgoing trusts)**, click the external trust that you want to administer, and then click **Properties**.
4. On the **Authentication** tab, click **Selective authentication**, and then click **OK**.

| **Note** |
| --- |
| Only the authentication settings for the outgoing trust are displayed when you click **Properties** and then click the **Authentication** tab in Active Directory Domains and Trusts. To view the correct authentication settings for the incoming side of a two-way, external trust, connect to a domain controller in the trusted domain, and then use Active Directory Domains and Trusts to view the authentication settings for the outgoing side of the same trust. |

# Enable selective authentication over a forest trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Selective authentication over a forest trust restricts access to only those users in a trusted forest who have been explicitly given authentication permissions to computer objects (resource computers) that reside in the trusting forest. To explicitly give authentication permissions to computer objects in the trusting forest to certain users, Administrators must grant those users the **Allowed to Authenticate** permission in Active Directory. For more information, see Grant the Allowed to Authenticate permission on computers in the trusting domain or forest. For more information about how selective authentication works, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

You can enable selective authentication over a forest trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to configure selective authentication settings, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory.

**To enable selective authentication over a forest trust**
**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the forest root domain, and then click **Properties**.
3. On the **Trusts** tab, under **Domains trusted by this domain (outgoing trusts)**, click the forest trust that you want to administer, and then click **Properties**.
4. On the **Authentication** tab, click **Selective authentication**, and then click **OK**.

| **Note** |
| --- |
| Only the authentication settings for the outgoing trust are displayed when you click **Properties** and then click the **Authentication** tab in Active Directory Domains and Trusts. To view the correct authentication settings for the incoming side of a two-way, forest trust, connect to a domain controller in the forest root domain of the trusted forest, and then use Active Directory Domains and Trusts to view the authentication settings for the outgoing side of the same trust. |

# Enable domain-wide authentication over an external trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The domain-wide authentication setting permits unrestricted access by any users in the trusted domain to all available shared resources in the trusting domain. This is the default authentication setting for external trusts, and it is representative of the way authentications were routed — without restriction — over Windows 2000 Server trusts. For more information about the domain-wide authentication setting, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

You can enable domain-wide authentication over an external trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to configure selective authentication settings, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To enable domain-wide authentication over an external trust**
**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that you want to administer, and then click **Properties**.
3. On the **Trusts** tab, under **Domains trusted by this domain (outgoing trusts)**, click the external trust that you want to administer, and then click **Properties**.
4. On the **Authentication** tab, click **Domain-wide authentication**, and then click **OK**.

| **Note** |
| --- |
| Only the authentication settings for the outgoing trust are displayed when you click **Properties** and then click the **Authentication** tab in Active Directory Domains and Trusts. To view the correct authentication settings for the incoming side of a two-way, external trust, connect to a domain controller in the trusted domain, and then use Active Directory Domains and Trusts to view the authentication settings for the outgoing side of the same trust. |

# Enable forest-wide authentication over a forest trust

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The forest-wide authentication setting permits unrestricted access by any users in the trusted forest to all available shared resources in any of the domains in the trusting forest. This is the default authentication setting for forest trusts, and it is representative of the way authentications were routed — without restriction — over Windows 2000 Server trusts. For more information about the forest-wide authentication setting, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

You can enable forest-wide authentication over a forest trust by using the New Trust Wizard in Active Directory Domains and Trusts or by using the Netdom command-line tool. For more information about how to use the Netdom command-line tool to configure selective authentication settings, see "Netdom.exe: Windows Domain Manager" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=41700).

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To enable forest-wide authentication over a forest trust**
**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the forest root domain, and then click **Properties**.
3. On the **Trusts** tab, under **Domains trusted by this domain (outgoing trusts)**, click the forest trust that you want to administer, and then click **Properties**.
4. On the **Authentication** tab, click **Forest-wide authentication**, and then click **OK**.

| **Note** |
| --- |
| Only the authentication settings for the outgoing trust are displayed when you click **Properties** and then click the **Authentication** tab in Active Directory Domains and Trusts. To view the correct authentication settings for the incoming side of a two-way, forest trust, connect to a domain controller in the trusted domain (the forest root domain in the other forest), and then use Active Directory Domains and Trusts to view the authentication settings for the outgoing side of the same trust. |

# Grant the Allowed to Authenticate permission on computers in the trusting domain or forest

Updated: June 21, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For users in a trusted Windows Server 2003 domain or forest to be able to access resources in a trusting Windows Server 2003 domain or forest, where the trust authentication setting has been set to selective authentication, each user must be explicitly granted the **Allowed to Authenticate** permission on the security descriptor of the computer objects (resource computers) that reside in the trusting domain or forest. For more information about how the **Allowed to Authenticate** permission works, see "Security Considerations for Trusts" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35413).

| **Note** |
| --- |
| The **Allowed to Authenticate** permission can be set on computer objects that represent member servers running Windows NT Server 4.0, Windows 2000 Server, and Windows Server 2003. |

| **Note** |
| --- |
| By default, only members of the Account Operators, Administrators, Domain Admins, Enterprise Admins, and SYSTEM security groups that are located in the trusting domain can modify the **Allowed to Authenticate** permission. |

To enable access to resources over an external trust or forest trust that is set to selective authentication, complete the following procedure by using Active Directory Users and Computers from the trusting domain.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To grant the Allowed to Authenticate permission on computers in the trusting domain or forest**
**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Computers** container or the container where your computer objects reside.
3. Right-click the computer object that you want users in the trusted domain or forest to access, and then click **Properties**.
4. On the **Security** tab, do one of the following:
   - In **Group or user names**, click the user names or group names for which you want to grant access to this computer, select the **Allow** check box next to the **Allowed to Authenticate** permission, and then click **OK**.
   - Click **Add**. In **Enter the object names to select**, type the name of the user object or group object for which you want to grant access to this resource computer, and then click **OK**. Select the **Allow** check box next to the **Allowed to Authenticate** permission, and then click **OK**.

# Appendix: New Trust Wizard Pages

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Understanding how user input is handled during the trust creation process will help you provide information when it is most necessary and help you better prepare for your specific procedure. This section explains the two most complex pages in the New Trust Wizard:

- **Direction of Trust**
- **Sides of Trust**

**Direction of Trust**

The **Direction of Trust** page in the New Trust Wizard is configured by an administrator in one domain to determine whether authentication requests should be routed from this domain to a specified domain, from the specified domain to this domain, or freely between both domains. The following options are available on the **Direction of Trust** page:

- **Two-way**: A two-way trust allows authentication requests that are sent by users in either domain or forest to be routed successfully to resources in either of the two domains or forests.
- **One-way: incoming**: A one-way, incoming trust allows authentication requests that are sent by users in your domain or forest (the domain or forest where you started the New Trust Wizard) to be routed successfully to resources in the other domain or forest.
- **One-way: outgoing**: A one-way, outgoing trust allows authentication requests that are sent by users in the other domain (the domain or forest that you are indicating in the New Trust Wizard as the specified domain or forest) to be routed successfully to resources in your domain or forest.

These options are explained in the following sections.

**Wizard Option: Two-way**

Use this option when you want to share resources equally between two domains or forests for all of the users that reside in both domains or forests. A two-way trust allows authentication requests that are sent by users in a trusted domain or forest to be routed successfully to the trusting domain or forest.

| **Note** |
| --- |
| Traditionally, documentation about domain and forest trusts have used the terms “trusting” and “trusted” to help administrators pinpoint the direction of the trust. Although this terminology is still used today to define and conceptualize how trusts work, it varies from the terminology that is used in the New Trust Wizard to help administrators determine the direction of trust. |

**Wizard Option: One-way: incoming**

Use this option when you want to allow authentication requests to be routed from your domain or forest (referred to as “this domain” or “this forest” in the wizard) to resources residing in a second domain or forest (referred to as “specified domain” or “specified forest” in the wizard). “One-way” in **One-way: incoming** means that this selection will create a one-way trust that can route authentications to resources in only one direction, while user access to those resources flows in the other direction. “Incoming” in **One-way: incoming** refers to the direction of the trust itself, not the direction in which authentication requests will flow. In other words, as shown in the following illustration, a "one-way incoming trust" means that your domain or forest will be the domain or forest that receives access to the resources in the other domain.




**Wizard Option: One-way: outgoing**

Use this option when you want to allow authentication requests to be routed to your domain or forest (referred to as “this domain” or “this forest” in the wizard) from users residing in a second domain or forest (referred to as “specified domain” or “specified forest” in the wizard). “One-way” in **One-way: outgoing** means that this selection will create a one-way trust that can route authentications to resources in only one direction, while user access to those resources flows in the other direction. “Outgoing” in **One-way: outgoing** refers to the direction of the trust itself, not the direction in which authentication requests will flow. In other words, as shown in the following illustration, a "one-way, outgoing trust" means that your domain or forest will provide access to resources that are located in your domain to users who are located in the other domain or forest.



**Sides of Trust**
In Windows NT 4.0 and Windows 2000, the only way to create trusts using the graphical user interface (GUI) was incrementally — one side of the trust at a time. When you create external trusts, shortcut trusts, realm trusts, or forest trusts in Windows Server 2003, you now have the option to create each side of the trust separately or both sides of the trust simultaneously.

**Wizard Option: This domain only**
Use this option when you want to create each side of the trust separately, which means that you must run the New Trust Wizard twice — once for each domain in the trust. Although the New Trust Wizard presents a different experience than previous version of Windows Server operating systems, this option provides behavior that is similar to the way that trusts were created in Windows NT 4.0 and Windows 2000. When you create trusts using this method, you must supply the same trust password for each domain. As a security best practice, all trust passwords should be strong passwords.

**Wizard Option: Both this domain and the specified domain**
This option provides administrators who possess the appropriate domain credentials for both domains in the trust relationship with the option to quickly create both sides of a trust by completing a single instance of the New Trust Wizard. When you select this option, a strong trust password is automatically generated for you. For this selection to be successful, the administrator running the wizard must acquire the appropriate administrative credentials for each domain in the trust relationship

# Administering the Windows Time Service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Time synchronization is critical for the proper operation of many Windows services and line-of-business applications. The Windows® Time service uses the Network Time Protocol (NTP) to synchronize computer clocks on the network so that an accurate clock value, or time stamp, can be assigned to network validation and resource access requests.

This guide provides information for administering the Windows Time service in the Microsoft Windows Server 2003 operating system.

**In this guide**


- Introduction to Administering the Windows Time Service
- Managing the Windows Time Service



**Acknowledgements**

Published: March 2005

Applies to: Windows Server 2003

Produced by: Microsoft Windows Server User Assistance team

Writer: Shala Brandolini

Editor: Justin Hall

# Introduction to Administering the Windows Time Service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Microsoft® Windows Server™ 2003 Windows Time service, also known as W32Time, synchronizes the date and time for all computers running on a Windows Server 2003 network. The service integrates NTP and time providers, making it a reliable and scalable time service for enterprise administrators.

The purpose of the Windows Time service is to make sure that all computers that are running versions of Windows 2000 Server, Windows Server 2003, or Windows XP in an organization use a common time. To guarantee appropriate common time usage, the Windows Time service uses a hierarchical relationship that controls authority and does not permit loops. A domain controller at the top of the hierarchy provides authoritative time to all other domain controllers, and domain clients use domain controllers as their time source. By default, the domain controller at the top of the hierarchy is the primary domain controller (PDC) operations master (also known as flexible single master operations or FSMO) in the forest root domain.

**Windows time source selection**

By default, Windows-based computers use the following sources for time synchronization:

- For computers that are joined to a domain, the first query is to a time source in the parent domain.

| **Note** |
| --- |
| Computers that are not joined to a domain and are running Windows XP or Windows XP Home Edition are configured to synchronize with time.windows.com by default. |

- If the time client is in a single-domain forest, the first query is to the PDC emulator in the domain.

- All PDC emulator operations masters follow the hierarchy of domains in the selection of their inbound time partner. A PDC emulator can synchronize its time from the PDC emulator in the parent domain or any domain controller in the parent domain.

For more information about time source selection, see How Windows Time Service Works on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=117753).

The authoritative time source at the root of the forest can acquire its time either by connecting to an installed hardware clock on the internal network, or by connecting to an external Network Time Protocol (NTP) server, which is connected to a hardware device. If no domain controller is configured as the authoritative time source in the forest root domain, the domain controller that holds the PDC emulator operations master role uses its internal clock to provide time to forest computers.

**External NTP time servers**

Many external NTP servers are available over the Internet. Use the following information to select an NTP server:

- The National Institute of Standards and Technology (NIST) in Boulder, Colorado, which is used as the external time provider by the Microsoft time server (time.windows.com). NIST provides the Automated Computer Time Service (ACTS), which can set a computer clock with an uncertainty of less than 10 milliseconds. For more information about NTP and for a list of external time servers, see Set Your Computer Clock Via the Internet: NIST Internet Time Service (ITS) (http://go.microsoft.com/fwlink/?LinkId=112035).

- The U.S. Naval Observatory (USNO) Time Service Department in Washington, DC, is another reliable source for accurate time synchronization in the United States. To see a list of USNO servers and their descriptions, see USNO Network Time Servers (http://go.microsoft.com/fwlink/?LinkId=112036).

- Many other sites throughout the world can be used for time synchronization. For more NTP server lists and search criteria, see the NTP Servers Web site (http://go.microsoft.com/fwlink/?LinkId=116972).

For the most highly accurate time synchronization, configure a hardware clock, such as a radio or Global Positioning System

(GPS) device, as the time source for the PDC. There are many consumer and enterprise devices that use NTP, which makes it possible for you to install the device on an internal network for use with the PDC.

You use the **w32tm** command-line tool to configure Windows Time Service. For a detailed technical reference of the Windows Time service, that includes complete documentation of the **w32tm** command-line tool and time service registry settings, see the Windows Time Service Technical Reference (http://go.microsoft.com/fwlink/?LinkID=100940).

**W32tm and net time**

The **net time** commands are predecessors of **w32tm** commands and should not be used to configure the Windows Time Service or to set the time on a computer while the Windows Time Service is actively running. We recommend that you configure and display the Windows Time Service information for Windows XP and Windows Server 2003 operating systems by using the**w32tm** commands.

Although the command **net time /querysntp** may appear to display the Simple Network Time Protocol (SNTP) server for Windows XP and Windows Server 2003 operating systems, it does not display complete time configuration information. You can use the command **w32tm /dumpreg /subkey:parameters** to determine whether the computer is configured to synchronize time from the domain hierarchy or from a manual list of time servers. The command output includes a line labeled **Type** that identifies the time synchronization method that is in use by the client. The following **Type** line outputs are possible for the time client:

- **NoSync**: The client does not synchronize time.
- **NTP**: The client synchronizes time from an external time source. Review the values in the **NtpServer** line in the output to see the name of the server or servers that the client uses for time synchronization.
- **NT5DS**: The client is configured to use the domain hierarchy for its time synchronization.
- **AllSync**: The client synchronizes time from any available time source, including domain hierarchy and external time sources.

# Managing the Windows Time Service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You initially configure the Windows 2003 Time service (W32Time) when you deploy your Active Directory forest root domain. Thereafter, it requires little day-to-day management. After you make changes on your network however, including when you add certain client computers, move the PDC emulator operations master role, or simply change the time source for you network, you might need to perform some of the following tasks:

- Configuring a time source for the forest
- Configuring Windows-based clients to synchronize time
- Restoring Windows Time service to default settings

# Configuring a time source for the forest

Updated: September 23, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first domain controller that you deploy in a domain holds the primary domain controller (PDC) emulator operations master (also known as flexible single master operations or FSMO) role for the domain. By default, the domain controller that holds the PDC emulator master role in the forest root domain is the reliable time source at the top of the time-source domain hierarchy for the forest. As soon as you install the first domain controller in the forest, set the PDC emulator in the forest root domain to synchronize from a valid Network Time Protocol (NTP) source or from a hardware clock installed on the network. If no time source is configured on the PDC emulator or any other domain controller in the forest root domain, the PDC emulator advertises as a reliable time source and uses its internal clock as the source for forest synchronization. In this case, no manual configuration is required.

After initial deployment of your network, you typically reconfigure the time service on the PDC emulator in the forest root domain in only two situations:

- You move the PDC emulator role to a different computer. In this case, you must configure the Windows Time service for the new PDC emulator master role holder and reconfigure the original PDC emulator master role holder to synchronize from the domain and not from an external or internal time source.

- You change the time source for the PDC emulator. For example, you change from synchronizing with an external source to synchronizing with an internal hardware device.

In some environments, one or more domain controllers are configured to act as standby PDC emulator role holders. If the current PDC emulator fails or is otherwise unavailable, the role can quickly be transferred to the standby. If you anticipate moving the PDC emulator role and want to avoid reconfiguring the new and old PDC emulator every time the role is moved, you can configure a domain controller in the forest root domain that is not the PDC emulator as the reliable time source for the forest. In this way, the root of the time service stays the same and remains properly configured.

| **Note** |
| --- |
| Make sure that the domain controller that you configure to be the forest time source is highly available and, if it is not the PDC emulator, that it does not hold other operations master roles that might need to be transferred. |

Use the following recommendations for configuring the time source for the forest root domain, in this order of preference:

1. Install a hardware clock, such as a radio or Global Positioning System (GPS) device, as the time source for the forest root domain and configure Windows Time service (W32time) on the PDC emulator or other domain controller to synchronize with this device. Many consumer and enterprise devices are available that use NTP. You can install the device on an internal network and configure the PDC emulator to use it as its time source.

   Hardware clocks have the following advantages:

   - More secure. You do not have to connect to the Internet.

   - Highest accuracy, although the accuracy level of NTP servers is as high as that of Windows Time Service (that is, the effect of the higher accuracy is not appreciated).

   Hardware clocks have the following disadvantage:

   - Expense and maintenance. You must purchase and install a hardware clock, whereas you can connect to a public time server at no cost and without hardware installation.

2. Configure the Windows Time service on the PDC emulator or other domain controller to synchronize with an external time server. Computer clocks synchronize with external time servers by using the NTP protocol over an IPnetwork. You

can manually configure the PDC emulator in the forest root domain to synchronize with the external time source.

External time servers have the following advantages:

- Low cost or no cost. Cost is usually limited to bandwidth.
- Good accuracy. Although hardware clocks have the highest accuracy, the accuracy of a hardware clock can actually exceed the accuracy of Windows Time Service, so the comparison of accuracy is not relevant.

External time servers have the following disadvantage:

- Security risk. NTP synchronization with an external time source is not authenticated and is therefore less secure than if time is sourced from inside the network.

Many GPS receivers and other time devices can function as NTP servers on a network. You can configure your Active Directory forest to synchronize time from these external hardware devices only if they are also acting as NTP servers on your network. To do so, configure the domain controller functioning as the PDC emulator in your forest root to synchronize with the NTP server provided by the GPS device. To do so, see Configure the Windows Time service on the PDC emulator in the Forest Root Domain (http://go.microsoft.com/fwlink/?LinkId=91969).

If you are using an external time source, you can use the following sites to select an NTP server:

- USNO NTP Network Time Servers (http://go.microsoft.com/fwlink/?LinkId=112036)
- Set Your Computer Clock Via the Internet: NIST Internet Time Service (ITS) (http://go.microsoft.com/fwlink/?LinkId=112035)
- The NTP Public Services Project (http://go.microsoft.com/fwlink/?LinkId=180528/bin/view/Servers/WebHome)

If you choose to implement an NTP time synchronization product other than the Windows Time service, you must disable the Windows Time service on the forest root domain reliable time source. All NTP servers need access to UDP port 123. If the Windows Time service is running on a Windows Server 2003–based computer, port 123 will remain occupied for the Windows Time service.

**Task requirements**

The following tools are required to perform the procedures for this task:

- W32tm.exe
- Services snap-in if you need to disable the Windows Time service

Perform the following procedures as needed to configure a time source for your forest:

1. Configure the Windows Time service on the PDC emulator in the Forest Root Domain
2. If you move the role of the PDC emulator to a new domain controller, Change the Windows Time service configuration on the previous PDC emulator.
3. If you are implementing a time synchronization product other than the Windows Time service in your environment that uses the NTP protocol, Disable the Windows Time service to free UDP port 123 on the network.

# Configure the Windows Time service on the PDC emulator in the Forest Root Domain

Updated: March 17, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure the Windows Time service (W32time) on the primary domain controller (PDC) emulator operations master (also known as flexible single master operations or FSMO) in the forest root domain when you deploy a new forest root domain or when you move the role of the PDC emulator in the forest root domain to a new domain controller. If you move the role of the PDC emulator to a new domain controller, you must also change the configuration of the Windows Time service on the previous PDC emulator. For more information, see Change the Windows Time service configuration on the previous PDC emulator.

| **Important** |
| --- |
| If the PDC emulator for your forest root domain is not configured or if it is unable to synchronize time from an external source, the PDC emulator for the forest logs W32time Event ID 12 in the System log of Event Viewer. For additional troubleshooting information, see Microsoft Knowledge Base article 816042 (http://go.microsoft.com/fwlink/?LinkID=60402) |

Before you configure the Windows Time service on the PDC emulator, you can determine the time difference between it and the source as a means to test basic Network Time Protocol (NTP) communication. After completing the configuration on the PDC emulator, be sure to monitor the System log in Event Viewer for W32time errors.

| **Note** |
| --- |
| For more information about the **w32tm** command, type **w32tm /?** at a command prompt or see Windows Time Service Tools and Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42984). |

### Administrative Credentials

To perform this procedure locally on the PDC emulator, you must be a member of the Administrators group. To perform this procedure from a remote computer, you must be a member of the Domain Admins group.

**To configure the Windows Time service on the PDC emulator**

1. Open a Command Prompt.
2. Type the following command to display the time difference between the local computer and a target computer, and then press ENTER:

   **w32tm /stripchart /computer:** *target* **/samples:** *n* **/dataonly**

   | Value | Definition |
   | --- | --- |
   | *target* | Specifies the Domain Name System (DNS) name or IP address of the NTP server that you are comparing the local computer's time against, for example, time.windows.com. |
   | *n* | Specifies the number of time samples that will be returned from the target computer to test basic NTP communication. |

3. Open User Datagram Protocol (UDP) port 123 for outgoing traffic if needed.
4. Open UDP port 123 (or a different port that you have selected) for incoming NTP traffic.
5. Type the following command to configure the PDC emulator, and then press ENTER:

   **w32tm /config /manualpeerlist:** *peers* **/syncfromflags:manual /reliable:yes /update**

   where *peers* specifies the list of DNS names or IP addresses of the NTP time source that the PDC emulator synchronizes from. For example, you can specify time.windows.com. When specifying multiple peers, use a space as the delimiter and enclose them in quotation marks. For more information about the NTP servers that you can use, see Microsoft

Knowledge Base article 262680 (http://go.microsoft.com/fwlink/?LinkID=60401).

For example, to configure your PDC emulator to use the following list of fictional time servers:

a. ntp1.fabrikam.com

b. ntp.contoso.com

c. time.fineartschool.net

Run the following command:

```
w32tm /config /manualpeerlist:"ntp1.fabrikam.com ntp.contoso.com time.fineartschool.net" /reliable:yes
/update
```

**See Also**
**Other Resources**
Windows Time Service Technical Reference
Event ID 12 — Domain Hierarchy Time Source Acquisition
Videos about the Windows Time Service

# Change the Windows Time service configuration on the previous PDC emulator

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to change the Windows Time service configuration on the previous PDC emulator after you transfer the role to a new domain controller. The previous PDC emulator will now automatically synchronize time with the domain hierarchy, getting it's time from the new reliable time source.

| **Note** |
| --- |
| For more information about the **w32tm** command, type **w32tm /?** at a command prompt or see Windows Time Service Tools and Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42984). |

### Administrative Credentials

To perform this procedure locally on the PDC emulator, you must be a member of the Administrators group. To perform this procedure from a remote computer, you must be a member of the Domain Admins group.

**To change the Windows Time service configuration on the previous PDC emulator**

1. Open a Command Prompt.
2. Type the following command and then press ENTER:

   **w32tm /config /syncfromflags:domhier /reliable:no /update**
3. Type the following command and then press ENTER:

   **net stop w32time**
4. Type the following command and then press ENTER:

   **net start w32time**

# Configure a domain controller in the parent domain as a reliable time source

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to configure a domain controller as a reliable time source if you anticipate moving the PDC emulator role and do not want to reconfigure the Windows Time service afterward. If you are configuring a domain controller as a reliable time source for the forest root domain, perform the procedure Change the Windows Time service configuration on the previous PDC emulator in addition to this procedure. Although you have not moved the PDC emulator role, you must still configure the existing PDC emulator to no longer be the reliable time source for the domain.

| **Note** |
| --- |
| For more information about the **w32tm** command, type **w32tm /?** at a command prompt or see Windows Time Service Tools and Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42984). |

### Administrative Credentials

To perform this procedure locally on the domain controller, you must be a member of the Administrators group. To perform this procedure from a remote computer, you must be a member of the Domain Admins group.

**To configure a domain controller in the parent domain as a reliable time source**

1. Open a Command Prompt.
2. Type the following command and press ENTER:

   **W32tm /config /reliable:yes /update**

# Configure the PDC emulator to synchronize from its internal hardware clock

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to configure the PDC emulator in the forest root domain to synchronize from its internal hardware clock and remain the reliable time source in the forest root domain.

| **Note** |
| --- |
| For more information about the **w32tm** command, type **w32tm /?** at a command prompt or see Windows Time Service Tools and Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42984). |

### Administrative Credentials

To perform this procedure locally on the PDC emulator, you must be a member of the Administrators group. To perform this procedure from a remote computer, you must be a member of the Domain Admins group.

**To configure the PDC emulator to synchronize from its internal hardware clock**

1. Open a Command Prompt.
2. Type the following command and then press ENTER:

   **w32tm /config /syncfromflags:domhier /reliable:yes /update**
3. Type the following command and then press ENTER:

   **net stop w32time**
4. Type the following command and then press ENTER:

   **net start w32time**

# Disable the Windows Time service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to disable the Windows Time service if you choose to implement another time synchronization product that uses the NTP protocol.

**Administrative Credentials**

To perform this procedure on the local computer, you must be a local Administrator on the PDC emulator. To perform this procedure on a remote computer, you must be a member of the Domain Admins group.

**To disable the Windows Time service**

1. Open the Services snap-in.
2. Right-click **Windows Time**, and select **Properties**. The **Windows Time Properties** dialog box appears.
3. In the **Startup type** box, select **Disabled** from the drop-down menu.
4. Click **OK**. Verify that the Startup Type for the time service appears as **Disabled**.

# Configuring Windows-based clients to synchronize time

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certain Windows-based client computers do not automatically synchronize their time with the Active Directory domain. The following client computers do not automatically synchronize to the domain time by using the Windows Time service:

- Client computers that run in a pre-Windows 2000 domain environment.
- Client computers that run in a UNIX environment.
- Computers that are not joined to a domain.

Configure these computers to request time from a particular source, such as a domain controller in the domain. If you do not specify a source that is synchronized with the domain, each computer's internal hardware clock governs its time.

**Task requirements**

The following tool is required to perform the procedures for this task:

- W32tm

Use either of the following procedures to configure your Windows-based clients to synchronize time:

- Configure a manual time source for a selected client computer

  -or-

- Configure a client computer for automatic domain time synchronization

# Configure a manual time source for a selected client computer

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you configure a manual time source for a client computer, you can determine the time difference between it and the source as a means to test basic NTP communication. After completing the configuration on the selected client computer, be sure to monitor the System log in Event Viewer for W32Time errors.

| **Note** |
| --- |
| For more information about the **w32tm** command, type **w32tm /?** at a command prompt or see Windows Time Service Tools and Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42984). |

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer. To perform this procedure from a remote computer, you must be a member of the Domain Admins group.

**To configure a manual time source for a selected client computer**

1. Open a Command Prompt.
2. Type the following command to display the time difference between the local computer and a target computer, and then press ENTER:

   **w32tm /stripchart /computer:** *target* **/samples:** *n* **/dataonly**

   | Value | Definition |
   | --- | --- |
   | *target* | Specifies the DNS name or IP address of the NTP server that you comparing the local computer's time against. |
   | *n* | Specifies the number of time samples that will be returned from the target computer to test basic NTP communication. |

3. Open UDP port 123 for outgoing traffic on firewall if needed.
4. Open UDP port 123 (or a different port you have selected) for incoming NTP traffic.
5. Type the following command to configure a manual time source for the selected computer and then press ENTER:

   **w32tm /config /manualpeerlist:** *peers* **/syncfromflags:manual /update**

   where *peers* specifies the list of DNS names or IP addresses of the NTP time source(s) that the selected computer will synchronize from. When specifying multiple peers, use a space as the delimiter and enclose them in quotation marks.

| **Tip** |
| --- |
| For links to SNTP server lists, see article 262680 in the Microsoft Knowledge Base (http://support.microsoft.com/kb/262680). |

# Configure a client computer for automatic domain time synchronization

Updated: September 28, 2009

Applies To: Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Vista

Some computers that are joined to a domain are configured to synchronize from a manual time source. Use the following procedure to configure a client computer that is currently synchronizing with a manually specified computer, to automatically synchronize time with the domain hierarchy.

| **Note** |
| --- |
| For more information about the **w32tm** command, type **w32tm /?** at a command prompt or see Windows Time Service Tools and Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42984). |

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer. To perform this procedure from a remote computer, you must be a member of the Domain Admins group.

**To configure a client computer for automatic domain time synchronization**

1. Open a Command Prompt.
2. Type the following command and then press ENTER:

   **w32tm /config /syncfromflags:domhier /update**
3. Type the following command and then press ENTER:

   **net stop w32time**
4. Type the following command and then press ENTER:

   **net start w32time**

To resolve this problem automatically, click the Fix this problem link. Then, click Run in the File Download dialog box.

Fix this problem

| **Important** |
| --- |
| This automated fix only applies to computers that are joined to a domain. |

| **Note** |
| --- |
| This wizard may be in English only; however, the automatic fix also works for other language versions of Windows. |

| **Note** |
| --- |
| If you are not on the computer that has the problem, save the automatic fix to a flash drive or to a CD so that you can run it on the computer that has the problem. If you do not want to use the automatic fix, you can do it yourself by performing the steps in the procedure above. |

# Restoring Windows Time service to default settings

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If the local Windows Time service settings are incorrectly configured, restoring the Windows Time service to its default settings might be more efficient than troubleshooting the problem.

**Task requirements**

The following tools are required to perform the procedures for this task:

- W32tm.exe

Perform the following procedure to restore local Windows Time service to the default settings:

- Restore Windows Time service on local computer to default settings

# Restore Windows Time service on local computer to default settings

Updated: September 28, 2009

Applies To: Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Vista

Use this procedure to restore the Windows Time service on the local computer to the default settings.

| Note |
| --- |
| For more information about the **w32tm** command, type **w32tm /?** at a command prompt or see Windows Time Service Tools and Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42984). |

### Administrative Credentials

To perform this procedure on the local computer, you must be a member of the Administrators group. To perform this procedure on a remote computer, you must be a member of the Domain Admins group.

**To restore Windows Time service on local computer to default settings**

1. Open a Command Prompt.
2. Type the following command and then press ENTER:

   **net stop w32time**
3. Type the following command and then press ENTER:

   **w32tm /unregister**
4. Type the following command and then press ENTER:

   **w32tm /register**
5. Type the following command and then press ENTER:

   **net start w32time**

To resolve this problem automatically, click the Fix this problem link. Then, click Run in the File Download dialog box.

Fix this problem

| Note |
| --- |
| This wizard may be in English only; however, the automatic fix also works for other language versions of Windows. |

| Note |
| --- |
| If you are not on the computer that has the problem, save the automatic fix to a flash drive or to a CD so that you can run it on the computer that has the problem. If you do not want to use the automatic fix, you can do it yourself by performing the steps in the procedure above. |

# Administering SYSVOL

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This SYSVOL Administering guide provides administering information for the Active Directory SYSVOL shared folder in the Microsoft Windows Server 2003 operating system.

**In this guide**


- Introduction to Administering SYSVOL
- Managing SYSVOL



**Acknowledgements**

Published: March 2005

Updated:

Applies to: Windows Server 2003

Produced by: Microsoft Windows Server User Assistance team

Writer: Mary Hillman

Editor: Jim Becker

# Introduction to Administering SYSVOL

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Server 2003 System Volume (SYSVOL) is a collection of folders and reparse points in the file systems that exist on each domain controller in a domain. SYSVOL provides a standard location to store important elements of Group Policy objects (GPOs) and scripts so that the File Replication service (FRS) can distribute them to other domain controllers within that domain.

| **Note** |
| --- |
| Only the Group Policy template (GPT) is replicated by SYSVOL. The Group Policy container (GPC) is replicated through Active Directory replication. To be effective, both parts must be available on a domain controller. |

FRS monitors SYSVOL and, if a change occurs to any file stored on SYSVOL, then FRS automatically replicates the changed file to the SYSVOL folders on the other domain controllers in the domain.

The day-to-day operation of SYSVOL is an automated process that does not require any human intervention other than watching for alerts from the monitoring system. Occasionally, you might perform some system maintenance as you change your network.

This objective describes the basic tasks required for managing SYSVOL in order to maintain capacity and performance of SYSVOL, for hardware maintenance, or for data organization.

**Key considerations for administering SYSVOL**

To manage SYSVOL, ensure that FRS properly replicates the SYSVOL data and that enough space is provided to store SYSVOL. Implement a monitoring system to detect low disk space and potential FRS disruptions so that you can address those issues before the system stops replicating. You can use the Ultrasound utility to monitor FRS. To download Ultrasound, see Ultrasound - Monitoring and Troubleshooting Tool for File Replication Service (FRS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=23439).

Other key considerations for managing SYSVOL are:

- Capacity.

Depending upon the configuration of your domain, SYSVOL can require a significant amount of disk space to function properly. During the initial deployment, SYSVOL might be allocated adequate disk space to function. However, as your Active Directory grows in size and complexity, the required capacity can exceed the available disk space.

If you receive indications that disk space is low, determine if the cause is due to inadequate physical space on the disk or a registry setting that limits the size of the staging area. By modifying a setting in the registry, you can allocate more staging area space, rather than relocating SYSVOL or the staging area. Increasing the space allocation in the registry is much faster and easier than relocation

- Performance.

Any changes made to SYSVOL are automatically replicated to the other domain controllers in the domain. If the files stored in SYSVOL change frequently, the replication increases the input and output for the volume where SYSVOL is located. For example, editing a GPO can potentially force a GPO-level replication. If the volume is also host to other system files, such as the directory database or the pagefile, then the increased input and output for the volume can impact the performance of the server.

- Hardware maintenance.

System maintenance, such as removal of a disk drive, can require you to relocate SYSVOL. Even if the maintenance occurs on a different disk drive, verify that that maintenance does not affect the system volume. Logical drive letters could change after you add and remove disks. FRS locates SYSVOL by using pointers stored in the directory and the registry. If drive letters change after you add or remove disk drives, be aware that these pointers are not automatically updated.

- Backing up Group Policy objects (GPOs).

The successful operation of Group Policy is heavily dependant on the reliable operation of SYSVOL. Key components of the GPO exist in the SYSVOL (in the policies subdirectory) and it is essential that these remain in sync with related components in Active Directory. Therefore, backing up only the SYSVOL component does not represent a full and complete backup of your GPOs. The Group Policy Management Console (GPMC) provides both UI-based and scriptable methods for backing up GPOs. It is important that you back up GPOs as part of your regular backup/disaster recovery processes. Soon after installation of a new domain, the default domain and default domain controllers' GPOs should be backed up. They should also be backed up after any subsequent changes are made.

**Understanding the SYSVOL folder structure**

Before you attempt to relocate all or portions of the system volume, you must clearly understand the folder structure and the relationships between the folders and the path information that is stored in the registry and the directory itself. When folders are relocated, any associated parameters that are stored in the registry and the directory must be updated to match the new location. The folder structure contains junctions that might also require updating when folders get moved to a new location.

Maintaining the relationship between the folders, junctions, and stored parameters is important when you must relocate all or portions of SYSVOL. Failure to do so can result in files being replicated to or from the wrong location. It can also result in files failing to replicate, yet FRS will not report any errors. Due to the configuration error, FRS looks in the wrong location for the files that you want to replicate.

The folder structure used by the system volume uses a feature called a junction point. Junction points look like folders and behave like folders (in Windows Explorer you cannot distinguish them from regular folders), but they are not folders. A junction point contains a link to another folder. When a program opens it, the junction point automatically redirects the program to the folder to which the junction point is linked. The redirection is completely transparent to the user and the application.

For example if you create two folders, C:\Folder1 and C:\Folder2, and create a junction called C:\Folder3, and then link the junction back to Folder1, Windows Explorer displays three folders:

\Folder1

\Folder2

\Folder3

If you open Folder3, Windows Explorer is redirected to Folder1 and displays the contents of Folder1. You receive no indication of the redirection because it is transparent to the user and to Windows Explorer. If you look at the contents of Folder1, you see that it is exactly the same as the contents displayed when you open Folder3. If you open a command prompt and list a directory, all three folders appear in the output. The first two are type <DIR> and Folder3 is type <JUNCTION>. If you list a directory of Folder3, you see the contents of Folder1.

| **Note** |
| --- |
| To create or update junctions, you need the Linkd.exe tool supplied with the Windows 2000 Server Resource Kit. Linkd allows you to create, delete, update, and view the links that are stored in junction points. |

By default, the system volume is contained in the %systemroot%\SYSVOL folder. The tree of folders contained within this folder can be extensive, depending on how your network uses FRS. When relocating folders in the system volume, ensure that you move all folders (including any hidden folders) and ensure that the relationships of the folders do not change unintentionally. When you relocate folders, you need to be concerned with the first three levels of subdirectories in order to properly update the parameters used by FRS. These levels are affected by junction points and parameter settings. These folders include:

- %systemroot%\SYSVOL
- %systemroot%\SYSVOL\Domain
- %systemroot%\SYSVOL\Domain\DO_NOT_REMOVE_Ntfrs_ Preinstalled_Directory
- %systemroot%\SYSVOL\Domain\Policies
- %systemroot%\SYSVOL\Domain\Scripts

- %systemroot%\SYSVOL\Staging
- %systemroot%\SYSVOL\Staging\Domain
- %systemroot%\SYSVOL\Staging Areas
- %systemroot%\SYSVOL\Staging Areas FQDN
- %systemroot%\SYSVOL\Sysvol
- %systemroot%\SYSVOL\Sysvol FQDN

(where FQDN is the fully qualified domain name of the domain that this domain controller hosts.)

| **Note** |
| --- |
| If any of the folders do not appear in Windows Explorer, click **Tools** and then click **Folder Options**. On the **View** tab, select **Show hidden files and folders**. |

If you use Windows Explorer to view these folders, they appear to be typical folders. If you open a command prompt and type dir to list these folders, you will notice two special folders are listed as <JUNCTION>. Both folders labeled FQDN are junction points. The junction in %systemroot%\SYSVOL\Sysvol links to %systemroot%\SYSVOL\Domain. The junction in %systemroot%\SYSVOL\Staging Areas is linked to %systemroot%\SYSVOL\Staging\Domain. If you change the path to the folders to which the junctions are linked, you must also update the junctions, including drive letter changes and folder changes.

Besides junction points linking to folders within the system volume tree, the registry and the directory also store references to folders. These references contain paths that you must update if you change the location of the folder. FRS uses two values that are stored in the directory. The first value, fRSRootPath, points to the location of the policies and scripts that are stored in SYSVOL. By default, this location is the %systemroot%\SYSVOL\Domain folder. The second value, fRSStagingPath, points to the location of the folders used as the staging area. By default, this location is the %systemroot%\SYSVOL\Staging\Domain folder. The Net Logon service uses a parameter stored in the registry to identify the location of the folder that it uses to create the SYSVOL and NETLOGON share points. By default, this path is %systemroot%\SYSVOL\Sysvol. If you change the paths to these folders, you must update these values.

When relocating SYSVOL, you first move the entire folder structure to a new location; then you update all the junction points and the parameters that are stored in the registry and the directory in order to maintain the relationships between the parameters, the folders, and the junctions. Optionally, you can relocate the staging area and leave the rest of the system volume at its original location. In this case, you must update the fRSStagingPath parameter in the directory and the junction point stored at %systemroot%\SYSVOL\staging areas.

# Managing SYSVOL

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following tasks for managing SYSVOL are described in this objective:

- Changing the Space Allocated to the Staging Area
- Relocating the Staging Area
- Relocating SYSVOL Manually
- Updating the System Volume Path
- Restoring and Rebuilding SYSVOL

# Changing the Space Allocated to the Staging Area

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The staging area stores files prior to being replicated and stores files that it has just received through replication. Although FRS compresses the data and attributes of the replicated files to save space in the Staging Area folder and reduce the time that is needed to replicate the files, this method requires making and storing a copy of every file prior to replication and can require a substantial amount of disk space.

The default size of the staging area is 660 megabytes (MB). The minimum size is 10 MB and the maximum size is 2 terabytes. You can adjust the size limit of the Staging Folder by setting the value in kilobytes (KB) of the Staging Space Limit registry entry in HKEY_Local_Machine\System\CurrentControlSet\Services\NtFrs\Parameters. For more information about setting the Staging Space Limit in the registry, see KB article 329491 in the Microsoft Knowledge Base.

**Task requirements**

The following tools are required to perform the procedures for this task:

- Net.exe
- Regedit.exe
- Event Viewer

To complete this task, perform the following procedures in order:

1. Stop the File Replication service
2. Change the space allocated to the Staging Area folder
3. Start the File Replication service

# Stop the File Replication service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to stop the File Replication service.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To stop the File Replication service**

1. Open a Command Prompt.
2. Type the following command and then press Enter:

   **net stop ntfrs**

# Change the space allocated to the Staging Area folder

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure outlines the steps needed to modify the registry entry that restricts the amount of disk space allocated to the staging area in SYSVOL.

| ⚠Caution |
| --- |
| The Registry Editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back up system state first. For information about backing up system state, see Administering Active Directory Backup and Restore. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To change the space allocated to the Staging Area folder**

1. Click **Start**, and then click **Run**.
2. In the Run dialog box, type **regedit** and then press Enter.
3. In the Registry Editor, navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters**.
4. Double-click **Staging Space Limit in KB** to open the **Edit DWord Value** dialog box.
5. In the **Base frame**, select **Decimal**.
6. For **Value Data** enter a value from 10000 through 2000000000. Do not use commas. Click **OK**.
7. Close the Registry Editor.

# Start the File Replication service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to restart the File Replication service and review the FRS event log to ensure that the restart succeeded.

**Administrative Credentials**

To perform this procedure you must be a member of the Domain Admins group in Active Directory.

**To start the File Replication service**

1. Open a Command Prompt.
2. Type the following command, and then press Enter:

   **net start ntfrs**

3. You can use Event Viewer to verify that NTFRS restarted correctly. Event ID 13501 indicates that the service restarted. Look for event ID 13516 to verify that the domain controller is running and ready for service. If you moved SYSVOL to a new location or relocated the Staging Area folder, look for event IDs 13553 and 13556, which indicate success.

# Relocating the Staging Area

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, the Active Directory Installation Wizard installs the Staging Area folder within the SYSVOL. The Active Directory Installation Wizard creates two folders—Staging and Staging Area—which FRS uses for the staging process. When you relocate the staging area, you can change the name. Ensure that you identify the proper area in case it is renamed in your environment.

Two parameters determine the location of the staging area. One parameter, *fRSStagingPath*, is stored in the directory and contains the path to the actual location that FRS uses to stage files. The other parameter is a junction point stored in the Staging Area folder in SYSVOL that links to the actual location that FRS uses to stage files. When relocating the staging area, you must update these two parameters to point to the new location.

Except where noted, perform these procedures on the domain controller that contains the Staging Area folder that you want to relocate.

**Task requirements**

To perform this task it is necessary that you understand the folder structure used by the system volume. For more information, see Introduction to Administering SYSVOL.

The following tools are required to perform the procedures for this task:

- Active Directory Sites and Services (Administrative Tool)
- Event Viewer
- Net.exe
- Dcdiag.exe (Windows Support Tools)
- Regedit.exe
- ADSI Edit.msc (Windows Support Tools)
- Linkd.exe (Windows Server 2003 Resource Kit Tools)

| **Note** |
| --- |
| To create or update junctions, you need the Linkd.exe tool supplied with Windows Server 2003 Resource Kit Tools on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=16544. Linkd allows you to create, delete, update, and view the links that are stored in junction points. |

To complete this task, perform the following procedures in order:

1. Identify replication partners
2. Check the status of the shared SYSVOL

   You do not need to perform the test on every partner, but you need to perform enough tests to be confident that the shared system volumes on the partners are healthy.
3. Verify replication with other domain controllers
4. Gather the SYSVOL path information
5. Reset the File Replication service staging folder to a different logical drive

# Identify replication partners

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to examine the Connection objects for a domain controller and determine its replication partners.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To identify replication partners**

1. Open Active Directory Sites and Services.
2. In the console tree, expand the **Sites** container to display the list of sites.
3. Double-click the site that contains the domain controller for which you want to determine Connection objects.

| **Note** |
| --- |
| If you do not know the site in which the domain controller is located, open a command prompt and type **ipconfig** to get the IP address of the domain controller. Use the IP address to Verify that an IP address maps to a subnet and determine the site association. |

4. Expand the **Servers** folder to display the list of servers in that site.
5. Expand the name of your domain controller to display its NTDS settings.
6. Double-click **NTDSSettings** to display the list of Connection objects in the details pane (these represent inbound connections used for replication). The **From Server** column displays the names of the domain controllers that are the replication partners.

# Check the status of the shared SYSVOL

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure involves checking Event Viewer to make sure that the File Replication service is started properly and then ensuring that the SYSVOL and Net Logon shared folders are created.

| **Note** |
| --- |
| You do not need to perform this procedure on every replication partner, but you need to perform it enough times to be confident that the shared system volumes on the replication partners are healthy. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To check the status of the shared SYSVOL**

1. Open Event Viewer.
2. In the Event Viewer tree, click **File Replication Service** to display the FRS events.
3. Look for an event 13516 with a date and time stamp that corresponds with the recent restart. It can take 15 minutes or more to appear. An event 13508 indicates that FRS is in the process of starting the service. An event 13509 indicates that the service has started successfully. Event 13516 indicates that the service is started, the folders are shared, and the domain controller is functional.
4. To verify the shared folder is created, open a command prompt and type **net share** to display a list of the shared folders on this domain controller, including Net Logon and SYSVOL.
5. At a command prompt, type **dcdiag /test:netlogons** and press ENTER.
6. Look for a message that states *computername* **passed test NetLogons** where *computername* is the name of the domain controller. If you do not see the test passed message, some problem will prevent replication from functioning. This test verifies that the proper logon privileges are set to allow replication to occur. If this test fails, verify the permissions set on the Net Logon and SYSVOL shared folders.

# Verify replication with other domain controllers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The tests performed in this procedure verify that different aspects of the replication topology are working properly. They check to see that objects are replicating and they verify that the proper logon permissions are set to allow replication to occur.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To verify replication is functioning**

1. Open a Command Prompt.
2. Type the following command, and then press Enter:

   **dcdiag /test:replications**

   | **Note** |
   | --- |
   | For this set of tests, the **/v** option is available. However, it does not display any significant additional information. Messages indicate that the connectivity and replications tests passed. |

3. To verify that the proper permissions are set for replication, type the following command and then press Enter:

   **dcdiag /test:netlogons**

   Messages indicate that the connectivity and netlogons tests passed.

# Gather the SYSVOL path information

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When relocating SYSVOL, you first move the entire folder structure to a new location; then you update all the junction points and the parameters that are stored in the registry and the directory in order to maintain the relationships between the parameters, the folders, and the junctions. Optionally, you can relocate the staging area and leave the rest of the system volume at its original location. In this case, you must update the fRSStagingPath parameter in the directory and the junction point stored at *%systemroot%*\SYSVOL\staging areas. For more information about the folder structure and the relationships between the folders and the path information stored in the registry and the SYSVOL directory itself see Introduction to Administering SYSVOL.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

Use the procedures below to locate the system volume path information and record the current values in the following table.

To relocate the staging area, record the information for rows 2 and 5.

| **Note** |
|---|
| To restore and rebuild SYSVOL, you will need the information from the domain controller that you are repairing recorded in rows 1, 2, and 3. Use the junctions located on the domain controller that you are copying from the SYSVOL folder structure to record the current value for rows 4 and 5. The new values for rows 4 and 5 are based on the domain controller that you are repairing. |

| | Parameter | Current Value | New Value |
|---|---|---|---|
| 1 | fRSRootPath | | |
| 2 | fRSStagingPath | | |
| 3 | Sysvol parameter in registry | | |
| 4 | Sysvol junction | | |
| 5 | Staging junction | | |

**To gather the system volume path information fRSRootPath and fRSStagingPath**

1. Click **Start**, click **Run**, type **adsiedit.msc**, and then press Enter.
2. Double-click **Domain [***computername***]** (where *computername* is the name of this domain controller). Verify that the Domain expands to display the domain component (DC=) folder.
3. Click the domain component to display the containers and OUs in the details pane.
4. Double-click **OU=Domain Controllers** to display the containers that represent the domain controllers.
5. Double-click the container that represents this domain controller (CN=*computername*) to display more containers.
6. Click the **CN=NTFRS Subscriptions** container.
7. In the details pane, right-click **CN=Domain System Volume**, and then click **Properties**.
8. Ensure that **Show mandatory attributes** is selected. Select it if it is not.
9. In **Attributes**, locate **fRSRootPath** and **fRSStagingPath** and record the current values in the table above.

10. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

11. Click **Cancel** to close the dialog box.

**SYSVOL parameter in the registry**

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.

2. In Registry Editor, navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**.

3. **Sysvol** appears in the details pane. The current value is listed in the **Data** column.

4. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

5. Exit Registry Editor.

**SYSVOL junction**

1. Open a Command Prompt.

2. Change the directory to *%systemroot%***\SYSVOL\Sysvol**.

| **Note** |
| --- |
| This assumes that the system volume is still in the default location. If it has been relocated, substitute the appropriate paths into these instructions. |

3. At the command prompt, type **dir**. Verify that the fully qualified domain name (FQDN) is listed as type <JUNCTION>.

4. At the command prompt, type **linkd** *fqdn* (where *fqdn* is the domain name listed in the Dir output). This displays the value stored in the junction point. Press ENTER.

5. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

**Staging junction**

1. Open a Command Prompt.

2. Change the directory to *%systemroot%***\SYSVOL\Staging Areas**.

| **Note** |
| --- |
| This assumes that the staging area is still in the default location. If it has been relocated, substitute the appropriate paths into these instructions. |

3. At the command prompt, type **dir**. Verify that the fully qualified domain name (FQDN) is listed as type <JUNCTION>.

4. At the command prompt, type **linkd** *fqdn* (where *fqdn* is the domain name listed in the Dir output). This displays the value stored in the junction point. Press ENTER.

5. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

# Reset the File Replication service staging folder to a different logical drive

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to reset the FRS Staging folder to a different logical drive.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To reset the File Replication service staging folder to a different logical drive**

1. Click **Start**, click **Run**, type **adsiedit.msc**, and then press ENTER.
2. Under **Domain [***computername***]**, locate the NtFrs Subscriber object under the host computer account in Active Directory. The generic path for this attribute is: CN=Replica Set Name,CN=NTFRS Subscriptions,CN=*Computername*,DC=*Domain Name*,DC=COM.

   For example, to reset the staging path for the SYSVOL replica set of domain controller \\DC1 in the contoso.com domain, the distinguished name (also known as DN) path for the FrsStagingPath parameter is:

   CN=Domain System Volume (SYSVOL share), CN=NTFRS Subscriptions,CN=DC1,DC=CONTOSO,DC=COM

   Where (when you read the distinguished name path from right to left):

   DC=CONTOSO,DC=COM is the domain hosting the computer account.

   CN=DC1 is the host computer account in the domain naming context (NC).

   CN=NTFRS Subscriptions is the NtfrsSubscriber object that holds the FrsStagingPath parameter.

   CN=Domain System Volume (SYSVOL share) is the FRS subscriber object.
3. Right-click the **CN=Domain System Volume** container, and click **Properties**.
4. Ensure that the **Show mandatory attributes** check box is selected. Select it if it is not.
5. In **Attributes**, click **fRSStagingPath**, and then click **Edit**. The current value appears in the **Value** box in the **String Attribute Editor** dialog box.
6. Enter the path to the new location for the FRS Staging folder in the **Value** box and click **OK**.
7. Click **OK** to close Properties.
8. To make sure that the staging path has been updated in the registry:
   a. Click **Start**, click **Run**, and type **regedit** on the server where you are changing the staging path and then press ENTER.

   b. Locate the following subkey:

      **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters\Replica Sets**

   c. Double-click **Replica Sets**. All replica sets are displayed as a GUID.

   d. To locate the replica set you are updating the staging area for, click a GUID and, in the details pane, find the **Replica Set Name**. Repeat until you find the correct replica set.

   e. After you locate the correct GUID and replica set name, right-click **Replica Set Stage** and then click **Modify**.

   f. In the **Value data** box, type the new staging area path, and then click **OK**.

      When the service detects a change in the staging path, event ID 13563 is logged with a series of self-explanatory

steps on how to proceed:

Event Type: Warning

Event Source: NtFrs

Event Category: None

Event ID: 13563

Date: 3/6/2005

Time: 7:13:01 PM

User: N/A

Computer: <*Computer name*>

Description: The File Replication service has detected that the staging path for the replica set DOMAIN SYSTEM VOLUME (SYSVOL SHARE) has changed.

Current staging path = E:\Windows\Sysvol\Staging\Domain

New staging path = E:\Frsstage

The service will start using the new staging path after it restarts. The service is set to restart after every restart.

It is recommended that you manually restart the service to prevent loss of data in the Staging folder.

To manually restart the service do the following:

[1] Run "net stop ntfrs" or use the Services snap-in to stop File Replication service.

[2] Move all the staging files corresponding to replica set DOMAIN SYSTEM VOLUME (SYSVOL SHARE) to the new staging location. If more than one replica set are sharing the current staging folder then it is safer to copy the staging files to the new staging folder.

[3] Run "net start ntfrs" or use the Services snap-in to start File Replication service, followed by "net start ntfrs".

9. To perform steps 1 through 3 in the event message, open a Command Prompt.
10. Type the following command and then press ENTER:

    **net stop ntfrs**

11. Move all the staging files corresponding to replica set **DOMAIN SYSTEM VOLUME (SYSVOL SHARE)** to the new staging location. If more than one replica set is sharing the current Staging folder, then it is safer to copy the staging files to the new Staging folder.
12. At a command prompt type the following command and then press ENTER:

    **net start ntfrs**

    Microsoft recommends that you follow step 11 (step 2 in the preceding event message) because the FRS Staging folder may contain thousands or tens of thousands of files in the original Staging folder, all of which may be destined for one or more downstream partners. In Windows Explorer, you can view the files in the staging folder. On the **Folder Options** menu, click the **View** tab, and then click to select the **Show hidden files and folders** check box. Copy the files to the new Staging folder, and then follow the remaining steps in the event log message.

# Relocating SYSVOL Manually

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you must move the entire system volume, not just the Staging Area folder, then you can relocate the system volume manually. Because no utilities can automate this process, you must carefully move all folders and properly maintain the same level of security at the new location.

You can also move SYSVOL with the Active Directory wizard, but this requires that you remove Active Directory from the domain controller and then reinstall Active Directory after SYSVOL has been moved. This should only be considered in extreme cases, and only when the domain controller is not running any other services or applications.

Except where noted, perform these steps on the domain controller that contains the system volume that you want to move.

| **Caution** |
| --- |
| This procedure can alter security settings. After you complete the procedure, the security settings on the new system volume are reset to the default settings that were established when you installed Active Directory. You must reapply any changes to the security settings on the system volume that you made since you installed Active Directory. This will cause additional replication traffic. Note that failure to reset permissions can result in unauthorized access to Group Policy objects and logon and logoff scripts. |

**Task Requirements**

The following tools are required to perform the procedures for this task:

- Active Directory Sites and Services (Administrative Tools)
- Event Viewer
- Windows Explorer
- Dcdiag.exe (Windows Support Tools)
- Regedit.exe
- ADSI Edit.msc (Windows Support Tools)
- Linkd.exe (Windows Server 2003 Resource Kit Tools)
- Net.exe
- Secedit.exe
- Notepad.exe

| **Note** |
| --- |
| To create or update junctions, you need the Linkd.exe tool supplied with Windows Server 2003 Resource Kit Tools on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=16544. Linkd allows you to create, delete, update, and view the links that are stored in junction points. |

To complete this task, perform the following procedures:

1. Identify replication partners
2. Check the status of the shared SYSVOL

3. Verify replication with other domain controllers
4. Gather the SYSVOL path information
5. Stop the File Replication service
6. Create the SYSVOL folder structure
7. Set the SYSVOL path
8. Set the staging area path

   If you have moved the Staging Area folder to a different location already, you do not need to do this step.

9. Prepare a domain controller for nonauthoritative SYSVOL restart
10. Update security on the new SYSVOL
11. Start the File Replication service
12. Check the status of the shared SYSVOL

# Identify replication partners

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to examine the Connection objects for a domain controller and determine its replication partners.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To identify replication partners**

1. Open Active Directory Sites and Services.
2. In the console tree, expand the **Sites** container to display the list of sites.
3. Double-click the site that contains the domain controller for which you want to determine Connection objects.

| **Note** |
| --- |
| If you do not know the site in which the domain controller is located, open a command prompt and type **ipconfig** to get the IP address of the domain controller. Use the IP address to Verify that an IP address maps to a subnet and determine the site association. |

4. Expand the **Servers** folder to display the list of servers in that site.
5. Expand the name of your domain controller to display its NTDS settings.
6. Double-click **NTDSSettings** to display the list of Connection objects in the details pane (these represent inbound connections used for replication). The **From Server** column displays the names of the domain controllers that are the replication partners.

# Check the status of the shared SYSVOL

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure involves checking Event Viewer to make sure that the File Replication service is started properly and then ensuring that the SYSVOL and Net Logon shared folders are created.

| **Note** |
| --- |
| You do not need to perform this procedure on every replication partner, but you need to perform it enough times to be confident that the shared system volumes on the replication partners are healthy. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To check the status of the shared SYSVOL**

1. Open Event Viewer.
2. In the Event Viewer tree, click **File Replication Service** to display the FRS events.
3. Look for an event 13516 with a date and time stamp that corresponds with the recent restart. It can take 15 minutes or more to appear. An event 13508 indicates that FRS is in the process of starting the service. An event 13509 indicates that the service has started successfully. Event 13516 indicates that the service is started, the folders are shared, and the domain controller is functional.
4. To verify the shared folder is created, open a command prompt and type **net share** to display a list of the shared folders on this domain controller, including Net Logon and SYSVOL.
5. At a command prompt, type **dcdiag /test:netlogons** and press ENTER.
6. Look for a message that states *computername* **passed test NetLogons** where *computername* is the name of the domain controller. If you do not see the test passed message, some problem will prevent replication from functioning. This test verifies that the proper logon privileges are set to allow replication to occur. If this test fails, verify the permissions set on the Net Logon and SYSVOL shared folders.

# Verify replication with other domain controllers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The tests performed in this procedure verify that different aspects of the replication topology are working properly. They check to see that objects are replicating and they verify that the proper logon permissions are set to allow replication to occur.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To verify replication is functioning**

1. Open a Command Prompt.
2. Type the following command, and then press Enter:

   **dcdiag /test:replications**

   | **Note** |
   | --- |
   | For this set of tests, the **/v** option is available. However, it does not display any significant additional information. Messages indicate that the connectivity and replications tests passed. |

3. To verify that the proper permissions are set for replication, type the following command and then press Enter:

   **dcdiag /test:netlogons**

   Messages indicate that the connectivity and netlogons tests passed.

# Gather the SYSVOL path information

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When relocating SYSVOL, you first move the entire folder structure to a new location; then you update all the junction points and the parameters that are stored in the registry and the directory in order to maintain the relationships between the parameters, the folders, and the junctions. Optionally, you can relocate the staging area and leave the rest of the system volume at its original location. In this case, you must update the fRSStagingPath parameter in the directory and the junction point stored at *%systemroot%*\SYSVOL\staging areas. For more information about the folder structure and the relationships between the folders and the path information stored in the registry and the SYSVOL directory itself see Introduction to Administering SYSVOL.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

Use the procedures below to locate the system volume path information and record the current values in the following table.

To relocate the staging area, record the information for rows 2 and 5.

| **Note** |
|---|
| To restore and rebuild SYSVOL, you will need the information from the domain controller that you are repairing recorded in rows 1, 2, and 3. Use the junctions located on the domain controller that you are copying from the SYSVOL folder structure to record the current value for rows 4 and 5. The new values for rows 4 and 5 are based on the domain controller that you are repairing. |

| | Parameter | Current Value | New Value |
|---|---|---|---|
| 1 | fRSRootPath | | |
| 2 | fRSStagingPath | | |
| 3 | Sysvol parameter in registry | | |
| 4 | Sysvol junction | | |
| 5 | Staging junction | | |

**To gather the system volume path information fRSRootPath and fRSStagingPath**

1. Click **Start**, click **Run**, type **adsiedit.msc**, and then press Enter.
2. Double-click **Domain [***computername***]** (where *computername* is the name of this domain controller). Verify that the Domain expands to display the domain component (DC=) folder.
3. Click the domain component to display the containers and OUs in the details pane.
4. Double-click **OU=Domain Controllers** to display the containers that represent the domain controllers.
5. Double-click the container that represents this domain controller (CN=*computername*) to display more containers.
6. Click the **CN=NTFRS Subscriptions** container.
7. In the details pane, right-click **CN=Domain System Volume**, and then click **Properties**.
8. Ensure that **Show mandatory attributes** is selected. Select it if it is not.
9. In **Attributes**, locate **fRSRootPath** and **fRSStagingPath** and record the current values in the table above.

10. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

11. Click **Cancel** to close the dialog box.

**SYSVOL parameter in the registry**

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.

2. In Registry Editor, navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**.

3. **Sysvol** appears in the details pane. The current value is listed in the **Data** column.

4. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

5. Exit Registry Editor.

**SYSVOL junction**

1. Open a Command Prompt.

2. Change the directory to *%systemroot%***\SYSVOL\Sysvol**.

| **Note** |
| --- |
| This assumes that the system volume is still in the default location. If it has been relocated, substitute the appropriate paths into these instructions. |

3. At the command prompt, type **dir**. Verify that the fully qualified domain name (FQDN) is listed as type <JUNCTION>.

4. At the command prompt, type **linkd** *fqdn* (where *fqdn* is the domain name listed in the Dir output). This displays the value stored in the junction point. Press ENTER.

5. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

**Staging junction**

1. Open a Command Prompt.

2. Change the directory to *%systemroot%***\SYSVOL\Staging Areas**.

| **Note** |
| --- |
| This assumes that the staging area is still in the default location. If it has been relocated, substitute the appropriate paths into these instructions. |

3. At the command prompt, type **dir**. Verify that the fully qualified domain name (FQDN) is listed as type <JUNCTION>.

4. At the command prompt, type **linkd** *fqdn* (where *fqdn* is the domain name listed in the Dir output). This displays the value stored in the junction point. Press ENTER.

5. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

# Stop the File Replication service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to stop the File Replication service.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To stop the File Replication service**

1. Open a Command Prompt.
2. Type the following command and then press Enter:

   **net stop ntfrs**

# Create the SYSVOL folder structure

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to create the SYSVOL folder structure. The %systemroot%\SYSVOL folder is at the top of the folder tree for the Windows system volume. To properly move SYSVOL, you must move the %systemroot%\SYSVOL folder and its contents. A subfolder of %systemroot%\SYSVOL is also named sysvol. Ensure that you move the proper folder (the %systemroot%\SYSVOL folder) and not the subfolder (%systemroot%\SYSVOL\sysvol). Do not confuse the two folders.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To create the SYSVOL folder structure**

1. In Windows Explorer, navigate to the folder that represents your current Windows system volume. By default, this is the %systemroot%\SYSVOL folder.
2. Right-click the **SYSVOL** folder, and then click **Copy**.
3. In Windows Explorer, navigate to the new location you created in the console tree, right-click the new location, and click **Paste**. You might see a dialog box stating that some files already exist and a prompt asking whether you want to continue copying the folder. At each such prompt, click **No**.
4. Verify that the folder structure was copied correctly. Compare the new folder structure to the original by opening a command prompt, typing the following command and pressing Enter to list the contents of the folders:

   **dir /s**

   Ensure that all folders exist. If any folders are missing at the new location (such as **\scripts**), then recreate them.

# Set the SYSVOL path

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to set the new path to the system volume in the registry.

| **Caution** |
| --- |
| The Registry Editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back up system state first. For information about backing up system state, see Administering Active Directory Backup and Restore. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

#### To set the SYSVOL path

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.
2. Navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**.
3. Right-click **SysVol** and click **Modify**.
4. In the **Value data** box, in the **Edit String** dialog box, enter the new path, including the drive letter and click **OK**.
5. Close Registry Editor.

| **Note** |
| --- |
| The path in the registry points to the SYSVOL folder located inside the SYSVOL folder that is under the root. When updating the path in the registry, ensure that it still points to the SYSVOL folder inside the SYSVOL folder that is under the root. |

# Set the staging area path

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to modify the **fRSStagingPath** parameter for a domain controller in Active Directory in order to change the location of the Staging Area folder on that domain controller. Perform this procedure at the console of the domain controller that is hosting the SYSVOL that you must reconfigure.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To set the staging area path**

1. Double-click **Domain [***computername***]** (where *computername* is the name of this domain controller). Verify that the Domain expands to display the domain component (DC=) folder.
2. Click the domain component to display the containers and OUs in the details pane.
3. Double-click **OU=Domain Controllers** to display the containers that represent the domain controllers.
4. Double-click the container that represents this domain controller (CN=*computername*) to display more containers.
5. Click the **CN=NTFRS Subscriptions** container.
6. In the details pane, right-click **CN=Domain System Volume**, and then click **Properties**.
7. Ensure that **Show mandatory attributes** is selected. Select it if it is not.
8. In **Attributes**, click **fRSStagingPath**, and then click **Edit**. The current value appears in the **Value** box in the **String Attribute Editor** dialog box.
9. In the **Value** box, enter the complete path to the new location where you want to locate the Staging Area folder (the path to the new folder that you created earlier), including the drive letter and click OK.
10. Close ADSI Edit.
11. Open a Command Prompt.
12. Change the directory to %systemroot%\SYSVOL\staging areas.
13. Type the following command to list the contents of the directory and then press ENTER:

    **dir**

    Verify that <JUNCTION> appears in the DIR output.
14. Update the junction so that it points to the new location by typing the following command and then pressing ENTER:

    **linkd junctionname** *newpath*

    where *newpath* specifies the same value that you entered for **fRSStagingPath** earlier.

# Prepare a domain controller for nonauthoritative SYSVOL restart

Updated: April 15, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Initiate a nonauthoritative restore of SYSVOL by modifying the value of the BurFlags (backup/restore flags) registry entry. Changing the value to D2 (hexadecimal) or 210 (decimal) prior to disconnecting a domain controller initiates an automatic nonauthoritative restore of SYSVOL when the domain controller is restarted. This process results in a smaller database size and improved synchronization time because the domain controller does not reconcile and process every deletion and modification that took place from the time of the last SYSVOL update to the time the domain controller is restarted in the new site, thereby improving synchronization time. Instead, it processes only the latest changes.

Separate entries exist for global and replica-set-specific BurFlags, as follows:

- To initiate a nonauthoritative restart of SYSVOL when it is the only replica set that is represented on the domain controller, set the value of the global BurFlags (REG_DWORD) entry under

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters\Backup/Restore\Process at Startup**

- If other replica sets are represented on the domain controller and you want to update only SYSVOL, set the value of the replica-set-specific BurFlags (REG_DWORD) entry under

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters\Cumulative Replica Sets\SYSVOL GUID**

Modifying the replica-set-specific BurFlags entry requires identifying the SYSVOL GUID in the registry.

| ⚠**Caution** |
| --- |
| The Registry Editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back up system state first. For information about backing up system state, see Administering Active Directory Backup and Restore. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To prepare a domain controller for nonauthoritative SYSVOL restart**

1. Click **Start**, click **Run**, type **regedit** and then click **OK**.
2. Navigate to

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters**
3. Expand **Parameters**.
4. Modify one of the BurFlags entries as follows:

   To modify the global BurFlags entry:

   a. Expand **Backup/Restore** and then click **Process at Startup**.

   b. In the details pane, right-click **BurFlags** and click **Modify**.

   c. In the **Value data** box, type **D2** hexadecimal or **210** decimal.

   d. Click **OK** and close Registry Editor.

To modify the replica-set-specific BurFlags entry:

a. Expand both **Cumulative Replica Sets** and **Replica Sets**.

b. Match the GUID under **Replica Sets** to the identical GUID under **Cumulative Replica Sets**, and click the matching GUID under **Cumulative Replica Sets**.

c. In the details pane, right click **BurFlags** and click **Modify**.

d. In the **Value data** box, type **D2** hexadecimal or **210** decimal.

e. Click **OK** and close Registry Editor.

# Update security on the new SYSVOL

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure applies the default security settings to the new SYSVOL folders. The settings will be the equivalent of those set by default during Active Directory installation. If additional security settings have been applied to the system volume since Active Directory was installed, you must reapply those settings after completing this procedure.

| **Caution** |
| --- |
| Failure to reapply security changes made after Active Directory was installed might result in unauthorized access to logon and logoff scripts and Group Policy objects. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

#### To update security on the new SYSVOL

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.
2. Navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\ Services\Netlogon\Parameters**.

   Note the path stored under SysVol.
3. In **Control Panel**, double-click **System**.
4. On the **Advanced** tab, click **Environment Variables**.
5. Under **System Variables**, click **New**.
6. For **Variable name**, type **sysvol**.
7. For **Variable value**, type the path that you noted in step 2.
8. Click **OK** twice. Click **OK** again to close **Properties**.
9. Open Notepad and enter the following information:

   [Unicode]

   Unicode=yes

   [Version]

   signature="$CHICAGO$"

   Revision=1

   [Profile Description]

   Description=default perms for sysvol

   [File Security]

   ;"%SystemRoot%\SYSVOL",0,"D:AR(A;OICI;FA;;;BA)"

   "%Sysvol%",2,"D:P(A;CIOI;GRGX;;;AU)(A;CIOI;GRGX;;;SO) (A;CIOI;GA;;;BA)(A;CIOI;GA;;;SY)(A;CIOI;GA;;;CO)"

   "%Sysvol%\domain\policies",2,"D:P(A;CIOI;GRGX;;;AU) (A;CIOI;GRGX;;;SO)(A;CIOI;GA;;;BA)(A;CIOI;GA;;;SY) (A;CIOI;GA;;;CO) (A;CIOI;GRGWGXSD;;;PA)"

   Use this file to apply the security settings to the new SYSVOL folders. Save this file as Sysvol.inf.

   | **Note** |
   | --- |
   | Do not include a space after (A;CIOI;GRGX;;;SO), (A;CIOI;GRGX;;;AU), or (A;CIOI;GA;;;SY). |

10. Open a new Command Prompt. Do not use an existing command prompt that has been open on your desktop because it will not have the proper environment settings. Change the directory to the folder where you saved the Sysvol.inf file.
11. Type the following command all on one line and then press ENTER:

    **SECEDIT /Configure /cfg** *sectemplatepath* **\sysvol.inf /db** *sectemplatepath* **\sysvol.db /overwrite**

    where s*ectemplatepath* specifies the path to where you saved Sysvol.inf.

# Start the File Replication service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to restart the File Replication service and review the FRS event log to ensure that the restart succeeded.

**Administrative Credentials**

To perform this procedure you must be a member of the Domain Admins group in Active Directory.

**To start the File Replication service**

1. Open a Command Prompt.
2. Type the following command, and then press Enter:

   **net start ntfrs**

3. You can use Event Viewer to verify that NTFRS restarted correctly. Event ID 13501 indicates that the service restarted. Look for event ID 13516 to verify that the domain controller is running and ready for service. If you moved SYSVOL to a new location or relocated the Staging Area folder, look for event IDs 13553 and 13556, which indicate success.

# Updating the System Volume Path

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you add or remove disk drives, the logical drive letters of the other drives on the system can change. If either your SYSVOL or Staging Area folder is located on one of the drives whose letter changes, FRS cannot locate them. You must update the paths that FRS uses to locate these folders in order to solve this problem. To change the path for the system volume, you need to make changes to the registry and in the directory. Changing the staging area path requires a change in the directory. Both changes require that you update the junction points. After updating the path information, you must restart File Replication service so it can reinitialize with the new values.

**Task Requirements**

The following tools are required to perform the procedures for this task:

- ADSI Edit.msc
- Net.exe
- Regedit.exe
- Linkd.exe

To complete this task, perform the following procedures in order:

1. Gather the SYSVOL path information
2. Stop the File Replication service
3. Set the SYSVOL path
4. Set the staging area path
5. Start the File Replication service

# Gather the SYSVOL path information

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When relocating SYSVOL, you first move the entire folder structure to a new location; then you update all the junction points and the parameters that are stored in the registry and the directory in order to maintain the relationships between the parameters, the folders, and the junctions. Optionally, you can relocate the staging area and leave the rest of the system volume at its original location. In this case, you must update the fRSStagingPath parameter in the directory and the junction point stored at *%systemroot%*\SYSVOL\staging areas. For more information about the folder structure and the relationships between the folders and the path information stored in the registry and the SYSVOL directory itself see Introduction to Administering SYSVOL.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

Use the procedures below to locate the system volume path information and record the current values in the following table.

To relocate the staging area, record the information for rows 2 and 5.

| **Note** |
|---|
| To restore and rebuild SYSVOL, you will need the information from the domain controller that you are repairing recorded in rows 1, 2, and 3. Use the junctions located on the domain controller that you are copying from the SYSVOL folder structure to record the current value for rows 4 and 5. The new values for rows 4 and 5 are based on the domain controller that you are repairing. |

| | Parameter | Current Value | New Value |
|---|---|---|---|
| 1 | fRSRootPath | | |
| 2 | fRSStagingPath | | |
| 3 | Sysvol parameter in registry | | |
| 4 | Sysvol junction | | |
| 5 | Staging junction | | |

**To gather the system volume path information fRSRootPath and fRSStagingPath**

1. Click **Start**, click **Run**, type **adsiedit.msc**, and then press Enter.
2. Double-click **Domain [***computername***]** (where *computername* is the name of this domain controller). Verify that the Domain expands to display the domain component (DC=) folder.
3. Click the domain component to display the containers and OUs in the details pane.
4. Double-click **OU=Domain Controllers** to display the containers that represent the domain controllers.
5. Double-click the container that represents this domain controller (CN=*computername*) to display more containers.
6. Click the **CN=NTFRS Subscriptions** container.
7. In the details pane, right-click **CN=Domain System Volume**, and then click **Properties**.
8. Ensure that **Show mandatory attributes** is selected. Select it if it is not.
9. In **Attributes**, locate **fRSRootPath** and **fRSStagingPath** and record the current values in the table above.

10. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

11. Click **Cancel** to close the dialog box.

**SYSVOL parameter in the registry**

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.

2. In Registry Editor, navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**.

3. **Sysvol** appears in the details pane. The current value is listed in the **Data** column.

4. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

5. Exit Registry Editor.

**SYSVOL junction**

1. Open a Command Prompt.

2. Change the directory to *%systemroot%***\SYSVOL\Sysvol**.

| **Note** |
| --- |
| This assumes that the system volume is still in the default location. If it has been relocated, substitute the appropriate paths into these instructions. |

3. At the command prompt, type **dir**. Verify that the fully qualified domain name (FQDN) is listed as type <JUNCTION>.

4. At the command prompt, type **linkd** *fqdn* (where *fqdn* is the domain name listed in the Dir output). This displays the value stored in the junction point. Press ENTER.

5. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

**Staging junction**

1. Open a Command Prompt.

2. Change the directory to *%systemroot%***\SYSVOL\Staging Areas**.

| **Note** |
| --- |
| This assumes that the staging area is still in the default location. If it has been relocated, substitute the appropriate paths into these instructions. |

3. At the command prompt, type **dir**. Verify that the fully qualified domain name (FQDN) is listed as type <JUNCTION>.

4. At the command prompt, type **linkd** *fqdn* (where *fqdn* is the domain name listed in the Dir output). This displays the value stored in the junction point. Press ENTER.

5. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

# Stop the File Replication service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to stop the File Replication service.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To stop the File Replication service**

1. Open a Command Prompt.
2. Type the following command and then press Enter:

   **net stop ntfrs**

# Set the SYSVOL path

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to set the new path to the system volume in the registry.

| **Caution** |
| --- |
| The Registry Editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back up system state first. For information about backing up system state, see Administering Active Directory Backup and Restore. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

#### To set the SYSVOL path

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.
2. Navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**.
3. Right-click **SysVol** and click **Modify**.
4. In the **Value data** box, in the **Edit String** dialog box, enter the new path, including the drive letter and click **OK**.
5. Close Registry Editor.

| **Note** |
| --- |
| The path in the registry points to the SYSVOL folder located inside the SYSVOL folder that is under the root. When updating the path in the registry, ensure that it still points to the SYSVOL folder inside the SYSVOL folder that is under the root. |

# Set the staging area path

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to modify the **fRSStagingPath** parameter for a domain controller in Active Directory in order to change the location of the Staging Area folder on that domain controller. Perform this procedure at the console of the domain controller that is hosting the SYSVOL that you must reconfigure.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To set the staging area path**

1. Double-click **Domain [***computername***]** (where *computername* is the name of this domain controller). Verify that the Domain expands to display the domain component (DC=) folder.
2. Click the domain component to display the containers and OUs in the details pane.
3. Double-click **OU=Domain Controllers** to display the containers that represent the domain controllers.
4. Double-click the container that represents this domain controller (CN=*computername*) to display more containers.
5. Click the **CN=NTFRS Subscriptions** container.
6. In the details pane, right-click **CN=Domain System Volume**, and then click **Properties**.
7. Ensure that **Show mandatory attributes** is selected. Select it if it is not.
8. In **Attributes**, click **fRSStagingPath**, and then click **Edit**. The current value appears in the **Value** box in the **String Attribute Editor** dialog box.
9. In the **Value** box, enter the complete path to the new location where you want to locate the Staging Area folder (the path to the new folder that you created earlier), including the drive letter and click OK.
10. Close ADSI Edit.
11. Open a Command Prompt.
12. Change the directory to %systemroot%\SYSVOL\staging areas.
13. Type the following command to list the contents of the directory and then press ENTER:

    **dir**

    Verify that <JUNCTION> appears in the DIR output.
14. Update the junction so that it points to the new location by typing the following command and then pressing ENTER:

    **linkd junctionname** *newpath*

    where *newpath* specifies the same value that you entered for **fRSStagingPath** earlier.

# Start the File Replication service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to restart the File Replication service and review the FRS event log to ensure that the restart succeeded.

**Administrative Credentials**

To perform this procedure you must be a member of the Domain Admins group in Active Directory.

**To start the File Replication service**

1. Open a Command Prompt.
2. Type the following command, and then press Enter:

   **net start ntfrs**

3. You can use Event Viewer to verify that NTFRS restarted correctly. Event ID 13501 indicates that the service restarted. Look for event ID 13516 to verify that the domain controller is running and ready for service. If you moved SYSVOL to a new location or relocated the Staging Area folder, look for event IDs 13553 and 13556, which indicate success.

# Restoring and Rebuilding SYSVOL

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If your efforts to move SYSVOL or perform certain maintenance tasks fail, you must recreate or rebuild the SYSVOL on a single domain controller. Attempt to rebuild SYSVOL on a single domain controller only when all other domain controllers in the domain have a healthy and functioning SYSVOL. Do not attempt to rebuild SYSVOL until you correct any problems that are occurring with FRS in a domain.

Use these procedures only if you are working on a domain controller that does not have a functional SYSVOL.

**Task Requirements**

The following tools are required to perform the procedures for this task:

- Active Directory Sites and Services
- Event Viewer
- Dcdiag.exe
- ADSIEdit.msc
- Net.exe
- Regedit.exe
- Windows Explorer
- Linkd.exe
- Ultrasound for monitoring

To complete this task, perform the following procedures in order:

1. Identify replication partners
2. Check the status of the shared SYSVOL

   Because you will be copying the system volume from one of the partners, you need to make sure that the system volume you copy from the partner is up to date.
3. Verify replication with other domain controllers
4. Restart the domain controller in Directory Services Restore Mode locally

   If you are sitting at the console of the domain controller, locally restart a domain controller in Directory Services Restore Mode. If you are accessing the domain controller remotely using Terminal Services, remotely restart a domain controller in Directory Services Restore Mode.
5. Gather the SYSVOL path information
6. Stop the File Replication service
7. Prepare a domain controller for nonauthoritative SYSVOL restart

8. Import the SYSVOL folder structure
9. Start the File Replication service
10. Check the status of the shared SYSVOL

# Identify replication partners

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to examine the Connection objects for a domain controller and determine its replication partners.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To identify replication partners**

1. Open Active Directory Sites and Services.
2. In the console tree, expand the **Sites** container to display the list of sites.
3. Double-click the site that contains the domain controller for which you want to determine Connection objects.

| **Note** |
| --- |
| If you do not know the site in which the domain controller is located, open a command prompt and type **ipconfig** to get the IP address of the domain controller. Use the IP address to Verify that an IP address maps to a subnet and determine the site association. |

4. Expand the **Servers** folder to display the list of servers in that site.
5. Expand the name of your domain controller to display its NTDS settings.
6. Double-click **NTDSSettings** to display the list of Connection objects in the details pane (these represent inbound connections used for replication). The **From Server** column displays the names of the domain controllers that are the replication partners.

# Check the status of the shared SYSVOL

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure involves checking Event Viewer to make sure that the File Replication service is started properly and then ensuring that the SYSVOL and Net Logon shared folders are created.

| **Note** |
| --- |
| You do not need to perform this procedure on every replication partner, but you need to perform it enough times to be confident that the shared system volumes on the replication partners are healthy. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To check the status of the shared SYSVOL**

1. Open Event Viewer.
2. In the Event Viewer tree, click **File Replication Service** to display the FRS events.
3. Look for an event 13516 with a date and time stamp that corresponds with the recent restart. It can take 15 minutes or more to appear. An event 13508 indicates that FRS is in the process of starting the service. An event 13509 indicates that the service has started successfully. Event 13516 indicates that the service is started, the folders are shared, and the domain controller is functional.
4. To verify the shared folder is created, open a command prompt and type **net share** to display a list of the shared folders on this domain controller, including Net Logon and SYSVOL.
5. At a command prompt, type **dcdiag /test:netlogons** and press ENTER.
6. Look for a message that states *computername* **passed test NetLogons** where *computername* is the name of the domain controller. If you do not see the test passed message, some problem will prevent replication from functioning. This test verifies that the proper logon privileges are set to allow replication to occur. If this test fails, verify the permissions set on the Net Logon and SYSVOL shared folders.

# Verify replication with other domain controllers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The tests performed in this procedure verify that different aspects of the replication topology are working properly. They check to see that objects are replicating and they verify that the proper logon permissions are set to allow replication to occur.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To verify replication is functioning**

1. Open a Command Prompt.
2. Type the following command, and then press Enter:

   **dcdiag /test:replications**

   | **Note** |
   | --- |
   | For this set of tests, the **/v** option is available. However, it does not display any significant additional information. Messages indicate that the connectivity and replications tests passed. |

3. To verify that the proper permissions are set for replication, type the following command and then press Enter:

   **dcdiag /test:netlogons**

   Messages indicate that the connectivity and netlogons tests passed.

# Restart the domain controller in Directory Services Restore Mode locally

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have physical access to a domain controller, you can restart the domain controller in Directory Services Restore Mode locally. Restarting in Directory Services Restore Mode takes the domain controller offline. In this mode, the server is not functioning as a domain controller.

When you start Windows Server 2003 in Directory Services Restore Mode, the local Administrator account is authenticated by the local Security Accounts Manager (SAM) database. Therefore, logging on requires that you use the local administrator password, not an Active Directory domain password. This password is set during Active Directory installation when you provide the password for Directory Services Restore Mode.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restart the domain controller in Directory Services Restore Mode locally**

1. Restart the domain controller.
2. When the screen for selecting an operating system appears, press F8.
3. On the **Windows Advanced Options** menu, select **Directory Services Restore Mode**.
4. When you are prompted, log on as the local administrator.

**See Also**
**Concepts**
Restart the domain controller in Directory Services Restore Mode Remotely

# Gather the SYSVOL path information

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When relocating SYSVOL, you first move the entire folder structure to a new location; then you update all the junction points and the parameters that are stored in the registry and the directory in order to maintain the relationships between the parameters, the folders, and the junctions. Optionally, you can relocate the staging area and leave the rest of the system volume at its original location. In this case, you must update the fRSStagingPath parameter in the directory and the junction point stored at *%systemroot%*\SYSVOL\staging areas. For more information about the folder structure and the relationships between the folders and the path information stored in the registry and the SYSVOL directory itself see Introduction to Administering SYSVOL.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

Use the procedures below to locate the system volume path information and record the current values in the following table.

To relocate the staging area, record the information for rows 2 and 5.

| **Note** |
|---|
| To restore and rebuild SYSVOL, you will need the information from the domain controller that you are repairing recorded in rows 1, 2, and 3. Use the junctions located on the domain controller that you are copying from the SYSVOL folder structure to record the current value for rows 4 and 5. The new values for rows 4 and 5 are based on the domain controller that you are repairing. |

| | Parameter | Current Value | New Value |
|---|---|---|---|
| 1 | fRSRootPath | | |
| 2 | fRSStagingPath | | |
| 3 | Sysvol parameter in registry | | |
| 4 | Sysvol junction | | |
| 5 | Staging junction | | |

**To gather the system volume path information fRSRootPath and fRSStagingPath**

1. Click **Start**, click **Run**, type **adsiedit.msc**, and then press Enter.
2. Double-click **Domain [***computername***]** (where *computername* is the name of this domain controller). Verify that the Domain expands to display the domain component (DC=) folder.
3. Click the domain component to display the containers and OUs in the details pane.
4. Double-click **OU=Domain Controllers** to display the containers that represent the domain controllers.
5. Double-click the container that represents this domain controller (CN=*computername*) to display more containers.
6. Click the **CN=NTFRS Subscriptions** container.
7. In the details pane, right-click **CN=Domain System Volume**, and then click **Properties**.
8. Ensure that **Show mandatory attributes** is selected. Select it if it is not.
9. In **Attributes**, locate **fRSRootPath** and **fRSStagingPath** and record the current values in the table above.

10. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

11. Click **Cancel** to close the dialog box.

**SYSVOL parameter in the registry**

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.

2. In Registry Editor, navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**.

3. **Sysvol** appears in the details pane. The current value is listed in the **Data** column.

4. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

5. Exit Registry Editor.

**SYSVOL junction**

1. Open a Command Prompt.

2. Change the directory to *%systemroot%***\SYSVOL\Sysvol**.

| **Note** |
| --- |
| This assumes that the system volume is still in the default location. If it has been relocated, substitute the appropriate paths into these instructions. |

3. At the command prompt, type **dir**. Verify that the fully qualified domain name (FQDN) is listed as type <JUNCTION>.

4. At the command prompt, type **linkd** *fqdn* (where *fqdn* is the domain name listed in the Dir output). This displays the value stored in the junction point. Press ENTER.

5. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

**Staging junction**

1. Open a Command Prompt.

2. Change the directory to *%systemroot%***\SYSVOL\Staging Areas**.

| **Note** |
| --- |
| This assumes that the staging area is still in the default location. If it has been relocated, substitute the appropriate paths into these instructions. |

3. At the command prompt, type **dir**. Verify that the fully qualified domain name (FQDN) is listed as type <JUNCTION>.

4. At the command prompt, type **linkd** *fqdn* (where *fqdn* is the domain name listed in the Dir output). This displays the value stored in the junction point. Press ENTER.

5. Record the current value in table above. Based on the folder structure discussed in detail in Introduction to Administering SYSVOL and the new location, record the new path value for this parameter in the table.

# Stop the File Replication service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to stop the File Replication service.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To stop the File Replication service**

1. Open a Command Prompt.
2. Type the following command and then press Enter:

   **net stop ntfrs**

# Prepare a domain controller for nonauthoritative SYSVOL restart

Updated: April 15, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Initiate a nonauthoritative restore of SYSVOL by modifying the value of the BurFlags (backup/restore flags) registry entry. Changing the value to D2 (hexadecimal) or 210 (decimal) prior to disconnecting a domain controller initiates an automatic nonauthoritative restore of SYSVOL when the domain controller is restarted. This process results in a smaller database size and improved synchronization time because the domain controller does not reconcile and process every deletion and modification that took place from the time of the last SYSVOL update to the time the domain controller is restarted in the new site, thereby improving synchronization time. Instead, it processes only the latest changes.

Separate entries exist for global and replica-set-specific BurFlags, as follows:

- To initiate a nonauthoritative restart of SYSVOL when it is the only replica set that is represented on the domain controller, set the value of the global BurFlags (REG_DWORD) entry under

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters\Backup/Restore\Process at Startup**

- If other replica sets are represented on the domain controller and you want to update only SYSVOL, set the value of the replica-set-specific BurFlags (REG_DWORD) entry under

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters\Cumulative Replica Sets\SYSVOL GUID**

Modifying the replica-set-specific BurFlags entry requires identifying the SYSVOL GUID in the registry.

| ⚠**Caution** |
| --- |
| The Registry Editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back up system state first. For information about backing up system state, see Administering Active Directory Backup and Restore. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To prepare a domain controller for nonauthoritative SYSVOL restart**

1. Click **Start**, click **Run**, type **regedit** and then click **OK**.
2. Navigate to

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters**
3. Expand **Parameters**.
4. Modify one of the BurFlags entries as follows:

   To modify the global BurFlags entry:

   a. Expand **Backup/Restore** and then click **Process at Startup**.

   b. In the details pane, right-click **BurFlags** and click **Modify**.

   c. In the **Value data** box, type **D2** hexadecimal or **210** decimal.

   d. Click **OK** and close Registry Editor.

To modify the replica-set-specific BurFlags entry:

a. Expand both **Cumulative Replica Sets** and **Replica Sets**.

b. Match the GUID under **Replica Sets** to the identical GUID under **Cumulative Replica Sets**, and click the matching GUID under **Cumulative Replica Sets**.

c. In the details pane, right click **BurFlags** and click **Modify**.

d. In the **Value data** box, type **D2** hexadecimal or **210** decimal.

e. Click **OK** and close Registry Editor.

# Import the SYSVOL folder structure

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to copy the SYSVOL folder structure from another domain controller. The *%systemroot%*\SYSVOL folder is at the top of the folder tree for the Windows system volume. To properly import SYSVOL, you must copy the *%systemroot%*\SYSVOL folder and its contents.

To use this procedure, the default shared folder Admin$ must exist on the domain controller from which you plan to copy the SYSVOL folder structure. Some organizations remove this shared folder or rename it for security reasons. If this shared folder is not available, you must share the *%systemroot%* folder and name the share point Admin$. If you share the *%systemroot%* folder in order to complete this procedure, ensure that you remove the share point after the procedure is complete in order to maintain any security policies established on your network. If the Admin$ share has been renamed, then use the name assigned by your organization instead of Admin$ while completing this procedure.

| **Caution** |
| --- |
| Never copy information from the system volume on one domain controller to the system volume on another domain controller unless you have stopped the File Replication service and configured SYSVOL for a non-authoritative restore during startup. Failure to do so can cause invalid data to be replicated and cause the system volumes on various domain controllers to become inconsistent. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To import the SYSVOL folder structure**

1. Open Windows Explorer.
2. Navigate to the existing *%systemroot%***\SYSVOL** folder that you are rebuilding and delete it.
3. Connect to the **Admin$** share on the domain controller that you identified earlier as the replication partner from which you plan to copy the SYSVOL folder structure.
4. Once you are connected to the **Admin$** share point, verify that a folder labeled **SYSVOL** appears. Right-click the **SYSVOL** folder, and click **Copy**.
5. In the same directory, find some blank space and right-click. Click **Paste**. You might see a dialog box stating that some files already exist and a prompt asking whether you want to continue copying the folder. At each such prompt, click **No**.
6. Verify that the original **SYSVOL** folder and a new folder labeled **Copy of SYSVOL** both appear. Right-click **Copy of SYSVOL** and click **Rename**. Type **SYSVOL2** and press ENTER.
7. Open a command prompt. Change to the drive letter that represents the connection to the remote domain controller where you created the **SYSVOL2** folder.
8. Change the directory to **SYSVOL2\sysvol**.
9. Type **dir** and press ENTER. Verify that **<JUNCTION>** appears in the Dir output and is followed by the name of the domain.
10. You must update the path in this junction so that it points to the new location. Type the following command:

    **linkd** *junctionname newpath*

    where *newpath* is the new value you recorded in row 4 of the table in Gather the SYSVOL path information. Press ENTER.
11. If the staging area has been relocated and is no longer inside the SYSVOL folder, skip steps 10 and 11 and proceed to step 12. At a command prompt, change the directory to \SYSVOL2\staging areas under the copy of SYSVOL that you created. Type **dir** to list the contents and verify that **<JUNCTION>** appears in the Dir output.
12. Update the junction so that it points to the new location. Type the following command:

    **linkd** *junctionname newpath*

where *newpath* is the new value that you recorded in row 5 of Table 1 while gathering system volume path information. Press ENTER.

13. At the command prompt, change back to the *%systemroot%* for the domain controller that you are repairing.
14. From the command prompt, use the Xcopy command to copy the contents of the \SYSVOL2 folder you created to a new SYSVOL folder on your local drive. Type the following command:

    **xcopy** *drive* **:\sysvol2*.* sysvol*.* /s /e /h /c /y**

    where *drive* is the letter representing the connection to the remote domain controller. Press ENTER.
15. Verify that the folder structure copied correctly. Compare the new folder structure to the **SYSVOL** (not the SYSVOL2) on the remote domain controller. Open a command prompt and type **dir** to list the contents of the folders. Ensure that all folders exist.
16. Remove the **SYSVOL2** folder that you created on the remote domain controller.
17. Disconnect from the remote domain controller. If you had to create a shared folder on that domain controller in order to connect to it, remove the shared folder. Some organizations consider it a security risk to retain shared folders that are not in use.
18. Restart the domain controller in normal mode.

# Start the File Replication service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to restart the File Replication service and review the FRS event log to ensure that the restart succeeded.

**Administrative Credentials**

To perform this procedure you must be a member of the Domain Admins group in Active Directory.

**To start the File Replication service**

1. Open a Command Prompt.
2. Type the following command, and then press Enter:

   **net start ntfrs**

3. You can use Event Viewer to verify that NTFRS restarted correctly. Event ID 13501 indicates that the service restarted. Look for event ID 13516 to verify that the domain controller is running and ready for service. If you moved SYSVOL to a new location or relocated the Staging Area folder, look for event IDs 13553 and 13556, which indicate success.

# Administering the Global Catalog

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides information for administering the Active Directory Global Catalog in the Microsoft Windows Server 2003 operating system.

**In this guide**


- Introduction to Administering the Global Catalog
- Managing the Global Catalog



**Acknowledgements**

Published: March 2005

Applies to: Windows Server 2003 SP1

Produced by: Microsoft Windows Server User Assistance team

Writer: Mary Hillman

Editor: Jim Becker

# Introduction to Administering the Global Catalog

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designate global catalog servers in sites to accommodate forest-wide directory searching and to facilitate domain client logons when universal groups are available (that is, when a domain has a domain functional level of Windows Server 2003 or Windows 2000 native). When universal groups are available in the domain, a domain controller must be able to locate a global catalog server to process a logon request.

**Global Catalog Placement**

Placement of global catalog servers in sites is a deployment task when you initially deploy your forest. As your forest grows, you might need to add more global catalog servers. To improve the speed of logging on and searching, place at least one global catalog server in each site and at least two global catalog servers if the site has multiple domain controllers. As a best practice, configure half of all domain controllers in a site to be global catalog servers if the site contains more than three domain controllers. If your deployment uses only a single domain, configure all domain controllers as global catalog servers. In a single-domain forest, configuring all domain controllers as global catalog servers requires no additional resources.

When placing global catalog servers, primary concerns are:

- Does any site have no global catalog servers?
- What domain controllers are designated as global catalog servers in a particular site?

**Initial Global Catalog Replication**

When you add a global catalog server to a site, the Knowledge Consistency Checker (KCC) updates the replication topology, after which replication of partial domain directory partitions that are available within the site begins. Replication of partial domain directory partitions that are available only from other sites begins at the next scheduled interval.

Adding subsequent global catalog servers within the same site requires only intrasite replication and does not affect network performance. Replication of the global catalog potentially affects network performance only when adding the first global catalog server in the site and the impact varies depending on the following conditions:

- The speed and reliability of the wide area network (WAN) link or links to the site.
- The size of the forest.

For example, in a forest that has a large hub site, five domains, and thirty small branch sites (some of which are connected by only dial-up connections), global catalog replication to the small sites takes considerably longer than replication of one or two domains to a few well-connected sites.

**Global Catalog Readiness**

A global catalog server is available to directory clients when it is locatable as a global catalog server in Domain Name System (DNS). Several conditions must be met before the global catalog server is ready to serve clients. These conditions are divided into seven levels (numbered 0 to 6) of readiness called occupancy levels. At each level, a specific degree of synchronization must be achieved before moving to the next level. By default, domain controllers running Windows Server 2003 require all levels to be reached before the global catalog is ready for use. At level 6, all partial, read-only directory partitions have been successfully replicated to the global catalog server. When the requirements of all occupancy levels have been satisfied, the Net Logon service on the global catalog server registers DNS service (SRV) resource records that identify the domain controller as a global catalog server in the site and in the forest.

In summary, a global catalog server is ready to serve clients when the following events occur, in this order:

- Occupancy level requirements are met by replicating read-only replicas.
- The **isGlobalCatalogReady** rootDSE attribute is set to TRUE.
- The Net Logon service on the domain controller has updated DNS with global catalog–specific SRV resource records.

At this point, the global catalog server begins accepting queries on ports 3268 and 3269.

**Global Catalog Removal**

When you remove the global catalog, the domain controller immediately stops advertising in DNS as a global catalog server. The KCC gradually removes the read-only replicas from the domain controller. On domain controllers running Windows Server 2003, the global catalog partial, read-only directory partitions are removed in the background, receiving a low priority so that high-priority services are not interrupted.

One reason that you might want to remove the global catalog from a domain controller is the availability of universal group membership caching in Windows Server 2003, which might eliminate the requirement for a global catalog server in a particular site.

Minimum hardware requirements for global catalog servers depend upon the numbers of users in the site. For disk space requirements and directory database storage guidelines, see "Assessing Disk Space and Memory Requirements" in Designing and Deploying Directory and Security Services on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45434).

For more information, see "How the Global Catalog Works" in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=44139).

# Managing the Global Catalog

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Designate global catalog servers in sites to accommodate forest-wide directory searching and so that Active Directory can determine universal group membership of native-mode domain clients.

The following tasks for managing the global catalog are described in this objective:

- Configuring a Global Catalog Server
- Determining Global Catalog Readiness
- Removing the Global Catalog

# Configuring a Global Catalog Server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When conditions in a site warrant adding a global catalog server, you can configure a domain controller to be a global catalog server. Selecting the global catalog setting on the NTDS Settings object prompts the KCC to update the topology. After the topology is updated, then read-only partial domain directory partitions are replicated to the designated domain controller. When replication must occur between sites to create the global catalog, the site link schedule determines when replication can occur.

**Task Requirements**

The following tools are required to perform the procedures for this task:

- Active Directory Sites and Services
- Repadmin.exe
- Dcdiag.exe

To complete this task, perform the following procedures:

| **Note** |
| --- |
| Some procedures are performed only when you are configuring the first global catalog server in a site. |

1. Determine whether a domain controller is a global catalog server
2. Designate a domain controller to be a global catalog server
3. Monitor global catalog replication progress
4. Verify successful replication to a domain controller

# Determine whether a domain controller is a global catalog server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the setting on the NTDS Settings object to indicate whether a domain controller is designated as a global catalog server.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group in Active Directory.

**To determine whether a domain controller is a global catalog server**

1. Open Active Directory Sites and Services.
2. In the console tree, expand the **Sites** container, expand the site of the domain controller you want to check, expand the **Servers** container, and then expand the **Server** object.
3. Right-click the **NTDS Settings** object, and then click **Properties**.
4. On the **General** tab, if the **Global Catalog** box is selected, the domain controller is designated as a global catalog server.

# Designate a domain controller to be a global catalog server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Setting the **Global Catalog** check box designates a domain controller as a global catalog server and initiates the process of replicating all domains to the server.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain where you are configuring the domain controller to be a global catalog server.

**To designate a domain controller to be a global catalog server**

1. Open Active Directory Sites and Services.
2. In the console tree expand the **Sites** container, and then expand the site in which you are designating a global catalog server.
3. Expand the **Servers** container and then expand the **Server** object for the domain controller that you want to designate as a global catalog server.
4. Right-click the **NTDS Settings** object for the target server, and then click **Properties**.
5. Select the **Global Catalog** check box, and then click **OK**.

# Monitor global catalog replication progress

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Monitor inbound replication progress to see how many (percentage) of the partial read-only directory partitions in the forest have replicated to a new global catalog server.

| **Note** |
| --- |
| Exchange 2003 servers use the global catalog exclusively when looking up addresses. Therefore, in addition to causing Active Directory client search problems, the condition of a global catalog server being advertised before it receives all partial replicas can cause Address Book lookup and mail delivery problems for Exchange clients.<br><br>The Name Service Provider Interface (NSPI) must be running on a global catalog server to enable MAPI access to Active Directory. To enable NSPI, you must restart the global catalog server after replication of the partial directory partitions is complete, or after occupancy requirements are met. |

### Administrative Credentials

To perform this procedure you must be a member of the Domain Admins group in Active Directory.

**To monitor global catalog replication progress**

1. Open a Command Prompt.
2. Type the following command and then press ENTER:

   **dcdiag /v /s:** *servername***| find "%"**

   | Value | Description |
   | --- | --- |
   | *servername* | Specifies the name of the new global catalog server. |

3. Repeat this command periodically to monitor progress. If the test shows no output, then replication has completed.

# Determining Global Catalog Readiness

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After replication of the partial domain directory partitions, the domain controller advertises as a global catalog server and begins accepting queries. If the domain controller advertises as a global catalog server before it has complete information from all domains in the forest, it might return false information to applications that begin using the server for forest-wide searches. Determine if a domain controller is ready to begin advertising itself as a global catalog server using the procedures for this task.

**Task Requirements**

The following tools are required to perform the procedures for this task:

- Ldp.exe
- Nltest.exe
- DNS snap-in

| **Note** |
| --- |
| The global catalog server must be restarted after replication has completed and before readiness is determined. |

To complete this task, perform the following procedures:

1. Verify global catalog readiness
2. Verify global catalog DNS registrations

# Verify global catalog readiness

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When a global catalog server has satisfied replication requirements, the **isGlobalCatalogReady rootDSE** attribute is set to **TRUE** and the global catalog is ready to serve clients.

**Administrative Credentials**

To perform the following procedures you must be a member of the Domain Users group.

**To verify global catalog readiness**

- Using the Windows interface
- Using a command prompt

**Using the Windows interface**

1. Click **Start**, click **Run**, type **Ldp**, and then click **OK**.
2. On the **Connection** menu, click **Connect**.
3. In the **Connect** box, type the name of the server whose global catalog readiness you want to verify.
4. In the **Port** box, if 389 is not showing, type 389.
5. If the **Connectionless** box is selected, clear it, and then click **OK**.
6. In the details pane, verify that the **isGlobalCatalogReady** attribute has a value of **TRUE**.
7. On the **Connection** menu, click **Disconnect**, and then close Ldp.

**Using a command prompt**

1. Open a Command Prompt.
2. Type the following command and then press ENTER:

   **nltest /server:** *servername* **/dsgetdc:** *domainname*

| Value | Description |
|---|---|
| *servername* | Specifies the name of the domain controller you have designated as a global catalog server. |
| *domainname* | Specifies the name of the domain to which the server belongs. |

3. In the **Flags:** line of the output, if GC appears, then the global catalog server has satisfied its replication requirements

# Verify global catalog DNS registrations

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To verify that a server is advertised as a global catalog server, verify the presence of DNS SRV resource records for the server. Restart the global catalog server prior to checking DNS registrations.

**Administrative Credentials**

To perform this procedure you must be a member of the Domain Users group.

**To verify global catalog DNS registrations**

1. Open the DNS snap-in and connect to a domain controller in the forest root domain.
2. Expand **Forward Lookup Zones** and then expand the forest root domain.
3. Click the **_tcp** container.
4. In the details pane, look in the **Name** column for **_gc** and in the **Data** column for the name of the server. The records that begin with **_gc** are global catalog SRV records.

# Removing the Global Catalog

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Removing the global catalog from a domain controller simply requires clearing the Global Catalog check box on the NTDS Settings object properties page. As soon as this is complete, the domain controller stops advertising itself as a global catalog server (Net Logon de-registers the global catalog-related records in DNS) and it immediately stops accepting LDAP requests over ports 3268 and 3269.

**Task Requirements**

The following tool is required to perform the procedures for this task:

- Active Directory Sites and Services

To complete this task, perform the following procedures:

1. Clear the global catalog setting
2. Monitor global catalog removal in Event Viewer

# Clear the global catalog setting

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Clearing the global catalog setting initiates removal of the partial directory partitions from the directory database of the domain controller.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the global catalog server.

**To clear the global catalog setting**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then expand the site from which you are removing a global catalog server.
3. Expand the **Servers** container and then expand the **Server** object for the domain controller that you want to remove as a global catalog server.
4. Right-click the **NTDS Settings** object for the target server, and then click **Properties**.
5. If the **Global Catalog** check box is selected, clear the check box, and then click **OK**.

# Monitor global catalog removal in Event Viewer

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To verify that the global catalog has been removed from a domain controller, monitor Event Viewer. The KCC logs Event 1268, when the global catalog has been successfully removed.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group.

**To monitor global catalog removal in Event Viewer**

1. Click **Start**, point to **Programs**, **Administrative Tools**, and click **Event Viewer**.
2. Right-click **Event Viewer (Local)**, and then click **Connect to another computer**.
3. In the **Select Computer** dialog box, click **Another computer**, type the name of the server from which you removed the global catalog, and then click **OK**.
4. Under **Event Viewer**, click **Directory Services log**.
5. Look for NTDS KCC event ID 1268, which indicates that the global catalog is removed from the local machine.

# Administering Operations Master Roles

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides information for administering Active Directory operations master roles in the Microsoft Windows Server 2003 operating system.

**In this guide**


- Introduction to Administering Operations Master Roles
- Managing Operations Master Roles



**Acknowledgements**

Published: March 2005

Applies to: Windows Server 2003

Produced by: Microsoft Windows Server User Assistance team

Writer: Shala Brandolini

Editor: Jim Becker

# Introduction to Administering Operations Master Roles

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Operations masters keep the directory functioning properly by performing specific tasks that no other domain controllers are permitted to perform.

Three operations master roles exist in each domain:

- The primary domain controller (PDC) emulator. The PDC emulator processes all replication requests from Microsoft Windows NT 4.0 backup domain controllers. It also processes all password updates for clients not running Active Directory–enabled client software, plus any other directory write operations.
- The relative identifier (RID) master. The RID master allocates RID pools to all domain controllers to ensure that new security principals can be created with a unique identifier.
- The infrastructure master. The infrastructure master for a given domain maintains a list of the security principals for any linked-value attributes.

In addition to the three domain-level operations master roles, two operations master roles exist in each forest:

- The schema master, which governs all changes to the schema.
- The domain naming master, which adds and removes domains and application partitions to and from the forest.

To perform these functions, the domain controllers hosting these operations master roles must be consistently available and be located in areas where network reliability is high. Careful placement of your operations masters becomes more important as you add more domains and sites to build your forest.

**Guidelines for Role Placement**

By improperly placing operations master role holders, you might prevent clients from changing their passwords or being able to add domains and new objects, such as Users and Groups. You might also be unable to make changes to the schema. In addition, name changes might not properly appear within group memberships that are displayed in the user interface.

As your environment changes, you must avoid the problems associated with improperly placed operations master role holders. Eventually, you might need to reassign the roles to other domain controllers.

Although you can assign the forest-level and domain-level operations master roles to any domain controller in the forest and domain respectively, improperly placing the infrastructure master role can cause it to function improperly. Other improper configurations can increase administrative overhead.

Although you can assign the operations master roles to any domain controller, follow these guidelines to minimize administrative overhead and ensure the performance of Active Directory. If a domain controller that is hosting operations master roles fails, following these guidelines also simplifies the recovery process. Guidelines for role placement include:

Leave the two forest-level roles on a domain controller in the forest root domain.

- Place the three domain-level roles on the same domain controller.
- Do not place the domain-level roles on a global catalog server.
- Place the domain-level roles on a higher performance domain controller.
- Adjust the workload of the operations master role holder, if necessary.
- Choose an additional domain controller as the standby operations master for the forest-level roles and choose an additional domain controller as the standby for the domain-level roles.

**Forest-level role placement in the forest root domain**

The first domain controller created in the forest is assigned the schema master and domain naming master roles. To ease administration and backup and restore procedures, leave these roles on the original forest root domain controller. Moving the roles to other domain controllers does not improve performance. Separating the roles creates additional administrative overhead when you must identify the standby operations masters and when you implement a backup and restore policy.

Unlike the PDC emulator role, forest-level roles rarely place a significant burden on the domain controller. Keep these roles together to provide easy, predictable management.

**Forest-level role placement on a Global Catalog server**

In addition to hosting the schema master and domain naming master roles, the first domain controller created in a forest also hosts the global catalog.

**Domain-level role placement on the same domain controller**

The three domain-level roles are assigned to the first domain controller created in a new domain. Except for the forest root domain, leave the roles at that location. Keep the roles together unless the workload on your operations master justifies the additional management burden of separating the roles.

Because all clients prior to Active Directory submit updates to the PDC emulator, the domain controller holding that role uses a higher number of RIDs. Place the PDC emulator and RID master roles on the same domain controller so that these two roles interact more efficiently.

If you must separate the roles, you can still use a single standby operations master for all three roles. However, you must ensure that the standby is a replication partner of all three of the role holders.

Backup and restore procedures also become more complex if you separate the roles. Special care must be taken to restore a domain controller that hosted an operations master role. By hosting the roles on a single computer, you minimize the steps that are required to restore a role holder.

**Domain-level role absence on a Global Catalog server**

Do not host the infrastructure master on a domain controller that is acting as a global catalog server.

The infrastructure master updates the names of security principals for any domain-named linked attributes. For example, if a user from one domain is a member of a group in a second domain and the user's name is changed in the first domain, then the second domain is not notified that the user's name must be updated in the group's membership list. Because domain controllers in one domain do not replicate security principals to domain controllers in another domain, the second domain never becomes aware of the change. The infrastructure master constantly monitors group memberships, looking for security principals from other domains. If it finds one, it checks with the security principal's domain to verify that the information is updated. If the information is out of date, the infrastructure master performs the update and then replicates the change to the other domain controllers in its domain.

Two exceptions apply to this rule. First, if all the domain controllers are global catalog servers, the domain controller that hosts the infrastructure master role is insignificant because global catalogs do replicate the updated information regardless of the domain to which they belong. Second, if the forest has only one domain, the domain controller that hosts the infrastructure master role is not needed because security principals from other domains do not exist.

Because it is best to keep the three domain-level roles together, avoid putting any of them on a global catalog server.

**Domain-level role placement on a higher performance domain controller**

Host the PDC emulator role on a powerful and reliable domain controller to ensure that it is available and capable of handling the workload. Of all the operations master roles, the PDC emulator creates the most overhead on the server that is hosting the role. It has the most intensive daily interaction with other systems on the network. The PDC emulator has the greatest potential to affect daily operations of the directory.

Domain controllers can become overloaded while attempting to service client requests on the network, manage their own resources, and handle any specialized tasks such as performing the various operations master roles. This is especially true of the domain controller holding the PDC emulator role. Again, clients prior to Active Directory and domain controllers running Windows NT 4.0 rely more heavily on the PDC emulator than Active Directory clients and Windows 2000 Server domain controllers. If your networking environment has clients and domain controllers prior to Active Directory, you might need to reduce the workload of the PDC emulator.

If a domain controller begins to indicate that it is overloaded and its performance is affected, you can reconfigure the

environment so that some tasks are performed by other, less-used domain controllers. By adjusting the domain controller's weight in the DNS environment, you can configure the domain controller to receive fewer client requests than other domain controllers on your network. Optionally, you can adjust the domain controller's priority in the DNS environment so that it processes client requests only if other DNS servers are unavailable. With fewer DNS client requests to process, the domain controller can use more resources to perform operations master services for the domain.

#### Guidelines for Role Transfer

Role transfer is the preferred method to move an operations master role from one domain controller to another. During a role transfer, the two domain controllers replicate to ensure that no information is lost. After the transfer completes, the previous role holder reconfigures itself so that it no longer attempts to perform as the operations master while the new domain controller assumes those duties. This prevents the possibility of duplicate operations masters existing on the network at the same time, which can lead to corruption in the directory.

Reasons for moving the operations master role(s) include inadequate service performance, failure or decommission of a domain controller hosting an operations master role, or if dictated by configuration changes made by an administrator.

#### Inadequate level of service

The PDC emulator is the operations master role that most impacts the performance of a domain controller. For clients that do not run Active Directory client software, the PDC emulator processes requests for password changes, replication, and user authentication. While providing support for these clients, the domain controller continues to perform its normal services, such as authenticating Active Directory–enabled clients. As the network grows, the volume of client requests can increase the workload for the domain controller that hosts the PDC emulator role and its performance can suffer. To solve this problem, you can transfer all or some of the master operations roles to another, more powerful domain controller. Alternately, you may choose to transfer the role to another domain controller, upgrade the hardware on the original domain controller, and then transfer the role back again.

#### Master operations role holder failure

In the event of a failure, you must decide if you need to relocate the operations master roles to another domain controller or wait for the domain controller to be returned to service. Base that determination on the role that the domain controller hosts and the expected downtime.

#### Decommissioning of the domain controller

Before permanently taking a domain controller offline, transfer any operations master roles held by the domain controller to another domain controller.

When you use the Active Directory Installation Wizard to decommission a domain controller that currently hosts one or more operations master roles, the wizard reassigns the roles to a different domain controller. When the wizard is run, it determines whether the domain controller currently hosts any operations master roles. If it detects any operations master roles, it queries the directory for other eligible domain controllers and transfers the roles to a new domain controller. A domain controller is eligible to host the domain-level roles if it is a member of the same domain. A domain controller is eligible to host a forest-level role if it is a member of the same forest.

#### Configuration changes

Configuration changes to domain controllers or the network topology can result in the need to transfer master operations roles. Except for the infrastructure master, you can assign operations master roles to any domain controller regardless of any other tasks that the domain controller performs. Do not host the infrastructure master role on a domain controller that is also acting as a global catalog server unless all of the domain controllers in the domain are global catalog servers or unless only one domain is in the forest. If the domain controller hosting the infrastructure master role is configured to be a global catalog server, you must transfer the infrastructure master role to another domain controller. Changes to the network topology can result in the need to transfer operations master roles in order to keep them in a particular site.

You can reassign an operations master role by transfer or, as a last resort, by seizure.

| **Important** |
| --- |
| If you must seize an operations master role, never reattach the previous role holder to the network without following the procedures in this guide. Incorrectly reattaching the previous role holder to the network can result in invalid data and corruption of data in the directory. |

# Managing Operations Master Roles

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Operations masters keep the directory functioning properly by performing specific tasks that no other domain controllers are permitted to perform.

The following tasks for managing operations master roles are described in this objective:

- Designating a standby operations master
- Transferring an operations master role
- Seizing an operations master role
- Reducing the workload on the PDC emulator master

# Designating a standby operations master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A standby operations master is a domain controller that you identify as the computer that assumes the operations master role if the original computer fails. A single domain controller can act as the standby operations master for all of the operations master roles in a domain, or you can designate a separate standby for each operations master role.

No utilities or special steps are required to designate a domain controller as a standby operations master. However, the current operations master and the standby should be well connected. This means that the network connection between them must support at least a 10-megabit transmission rate and be available at all times. In addition, configure the current role holder and the standby as direct replication partners by manually creating a Connection object between them.

Configuring a replication partner can save some time if you must reassign any operations master roles to the standby operations master. Before transferring a role from the current role holder to the standby operations master, ensure that replication between the two computers is functioning properly. Because they are replication partners, the new operations master is as updated as the original operations master, thus reducing the time required for the transfer operation.

During role transfer, the two domain controllers exchange any unreplicated information to ensure that no transactions are lost. If the two domain controllers are not direct replication partners, a substantial amount of information might need to be replicated before the domain controllers completely synchronize with each other. The role transfer requires extra time to replicate the outstanding transactions. If the two domain controllers are direct replication partners, fewer outstanding transactions exist and the role transfer operation completes sooner.

Designating a domain controller as a standby also minimizes the risk of role seizure. By making the operations master and the standby direct replication partners, you reduce the chance of data loss in the event of a role seizure, thereby reducing the chances of introducing corruption into the directory.

When you designate a domain controller as the standby, follow all recommendations that are discussed in Guidelines for Role Placement in Introduction to Administering Operations Master Roles. To designate a standby for the forest-level roles, choose a global catalog server so it can interact more efficiently with the domain naming master. To designate a standby for the domain-level roles, ensure that the domain controller is not a global catalog server so that the infrastructure master continues to function properly if you must transfer the roles.

**Task Requirements**

The following tools are required to perform the procedures for this task:

- Active Directory Sites and Services
- Repadmin.exe

To complete this task, perform the following procedures:

1. Determine whether a domain controller is a global catalog server
2. Create a connection object on the current operations master
3. Create a connection object on the standby operations master
4. Verify successful replication to a domain controller

# Determine whether a domain controller is a global catalog server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the setting on the NTDS Settings object to indicate whether a domain controller is designated as a global catalog server.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group in Active Directory.

**To determine whether a domain controller is a global catalog server**

1. Open Active Directory Sites and Services.
2. In the console tree, expand the **Sites** container, expand the site of the domain controller you want to check, expand the **Servers** container, and then expand the **Server** object.
3. Right-click the **NTDS Settings** object, and then click **Properties**.
4. On the **General** tab, if the **Global Catalog** box is selected, the domain controller is designated as a global catalog server.

# Create a connection object on the current operations master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To help ensure that the current role holder and the standby operations master are replication partners, you can manually create a Connection object between the two domain controllers. Even if a Connection object is generated automatically, it is recommended that you manually create one. The system can alter automatically created Connection objects at any time. Manually created connections remain the same until an administrator changes them.

You must know the current operations master role holder to perform the following procedure. To determine the current operations master role holders, see View the current operations master role holders.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To create a connection object on the current operations master**

1. Open **Active Directory Sites and Services**.
2. Expand the site name in which the standby operations master is located to display the **Servers** folder.
3. Expand the **Servers** folder to see a list of the servers in that site.
4. Expand the name of the server that is currently hosting the operations master role to display its NTDS Settings.
5. Right-click **NTDS Settings**, click **New**, and then click **Connection**.
6. In the **Find Domain Controllers** dialog box, select the name of the standby operations master, and then click **OK**.
7. In the **New Object-Connection** dialog box, enter an appropriate name for the Connection object or accept the default name, and click **OK**.

# Create a connection object on the standby operations master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To help ensure that the current role holder and the standby operations master are replication partners, you can manually create a Connection object between the two domain controllers. Even if a Connection object is generated automatically, it is recommended that you manually create one. The system can alter automatically created Connection objects at any time. Manually created connections remain the same until an administrator changes them.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To create a Connection object on the standby operations master**

1. Open **Active Directory Sites and Services**.
2. Expand the site name in which the standby operations master is located to display the **Servers** folder.
3. Expand the **Servers** folder to see a list of the servers in that site.
4. Expand the name of the server that you want to be the standby operations master to display its NTDS Settings.
5. Right-click **NTDS Settings**, click **New**, and then click **Connection**.
6. In the **Find Domain Controllers** dialog box, select the name of the current role holder, and then click **OK**.
7. In the **New Object-Connection** dialog box, enter an appropriate name for the Connection object or accept the default name, and click **OK**.

# Verify successful replication to a domain controller

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **repadmin /showrepl** command to verify successful replication to a specific domain controller. If you are not running Repadmin on the domain controller whose replication you are checking, you can specify a destination domain controller in the command. Repadmin lists INBOUND NEIGHBORS for the current or specified domain controller. INBOUND NEIGHBORS shows the distinguished name of each directory partition for which inbound directory replication has been attempted, the site and name of the source domain controller, and whether replication succeeded or not, as follows:

- Last attempt @ YYYY-MM-DD HH:MM.SS was successful.
- Last attempt @ [Never] was successful.

If @ [Never] appears in the output for a directory partition, replication of that directory partition has never succeeded from the identified source replication partner over the listed connection.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the destination domain controller.

**To verify successful replication to a domain controller**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl** *servername* **/u:** *domainname* **** *username* **/pw:***

   | Term | Definition |
   | --- | --- |
   | *servername* | Specifies the name of the destination domain controller. |
   | *domainname* | Specifies the single-label name of the domain of the destination domain controller. (You do not have to use a fully qualified Domain Name System (DNS) name.) |
   | *username* | Specifies the name of an administrative account in that domain. |

3. When you are prompted for a password, type the password for the user account that you provided, and then press ENTER.

You can also use Repadmin to generate the details of replication to and from all replication partners in a spreadsheet. The spreadsheet displays data in the following columns:

Showrepl_COLUMNS

Destination DC Site

Destination DC

Naming Context

Source DC Site

Source DC

Transport Type

Number of Failures

Last Failure Time

Last Success Time

Last Failure Status

The following procedure shows how to create this spreadsheet and set column headers for improved readability.

**To generate a repadmin /showrepl spreadsheet for all replication partners**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl * /csv >showrepl.csv**
3. Open Microsoft Excel.
4. On the **File** menu, click **Open**, navigate to showrepl.csv, and then click **Open**.
5. Hide or delete column **A** as well as the **Transport Type** column, as follows:
6. Select a column that you want to hide or delete.
   - To hide the column, on the **Format** menu, click **Column**, and then click **Hide**.

     Or

   - To delete the column, right-click the selected column, and then click **Delete**.
7. Select row 1 beneath the column heading row, and then, on the **Window** menu, click **Freeze Panes**.
8. Select the entire spreadsheet. On the **Data** menu, click **Filter**, and then click **AutoFilter**.
9. In the **Last Success Time** column, click the down arrow, and then click **Sort Ascending**.
10. In the **Source DC** column, click the down arrow, and then click **Custom**.
11. In the **Custom AutoFilter** dialog box, under **Show rows where**, click **does not contain**. In the adjacent text box, type **del** to eliminate from view the results for deleted domain controllers.
12. Repeat step 10 for the **Last Failure Time** column, but use the value **does not equal**, and type the value **0**.
13. Resolve replication failures.

The last successful attempt should agree with the replication schedule for intersite replication, or the attempt should be within the last hour for intrasite replication.

If Repadmin reports any of the following conditions, see Troubleshooting Active Directory Replication Problems:

- The last successful intersite replication was prior to the last scheduled replication.
- The last intrasite replication was longer than one hour ago.
- Replication was never successful.

**See Also**
**Concepts**
Troubleshooting Active Directory Replication Problems

# Transferring an operations master role

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you create a new domain, the Active Directory Installation Wizard automatically assigns all of the domain-level operations master roles to the first domain controller that is created in that domain. When you create a new forest, the wizard also assigns the two forest-level operations master roles to the first domain controller. After the domain is created and functioning, you might transfer various operations master roles to different domain controllers to optimize performance and simplify administration.

The transfer of forest-level and domain-level operations master roles is performed as needed and is governed by the guidelines for placing operations master roles. Before you transfer an operations master role, ensure that replication between the current role holder and the domain controller assuming the role is updated.

In addition, you must determine if the domain controller that you intend to assume an operations master role is a global catalog server. However, the infrastructure master for each domain must not host the global catalog.

Do not change the global catalog configuration on the domain controller that you intend to assume an operations master role unless your IT management authorizes that change. Changing the global catalog configuration can cause changes that can take days to complete, and the domain controller might not be available during that period. Instead, transfer the operations master roles to a different domain controller that is already properly configured.

**Transferring to a standby ops master**

By following the recommendations for operations master role placement, the standby operations master is a direct replication partner and is ready to assume the roles. Remember to designate a new standby for the domain controller that assumes the roles.

**Transferring an ops master role when no standby is ready**

If you do not follow the recommendations for role placement and you have not designated a standby operations master, you must properly prepare a domain controller to which you intend to transfer the operations master roles. Preparing the future role holder is the same process as preparing a standby operations master. You must manually create a Connection object to ensure that it is a replication partner with the current role holder and that replication between the two domain controllers is updated.

In addition, you must determine whether the domain controller intended to assume an operations master role is a global catalog server. The infrastructure master for each domain must not host the global catalog.

**Task Requirements**

The following tools are required to perform the procedures for this task:

- Repadmin.exe
- Active Directory Sites and Services
- Active Directory Domains and Trusts
- Active Directory Schema snap-in
- Active Directory Users and Computers
- Ntdsutil.exe

To complete this task, perform the following procedures:

1. Verify successful replication to a domain controller
2. Determine whether a domain controller is a global catalog server

3. Install the Schema snap-in
4. Transfer the schema master
5. Transfer the domain naming master
6. Transfer the domain-level operations master roles
7. View the current operations master role holders

# Verify successful replication to a domain controller

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **repadmin /showrepl** command to verify successful replication to a specific domain controller. If you are not running Repadmin on the domain controller whose replication you are checking, you can specify a destination domain controller in the command. Repadmin lists INBOUND NEIGHBORS for the current or specified domain controller. INBOUND NEIGHBORS shows the distinguished name of each directory partition for which inbound directory replication has been attempted, the site and name of the source domain controller, and whether replication succeeded or not, as follows:

- Last attempt @ YYYY-MM-DD HH:MM.SS was successful.
- Last attempt @ [Never] was successful.

If @ [Never] appears in the output for a directory partition, replication of that directory partition has never succeeded from the identified source replication partner over the listed connection.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the destination domain controller.

**To verify successful replication to a domain controller**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl** *servername* **/u:** *domainname* **** *username* **/pw:***

   | Term | Definition |
   | --- | --- |
   | *servername* | Specifies the name of the destination domain controller. |
   | *domainname* | Specifies the single-label name of the domain of the destination domain controller. (You do not have to use a fully qualified Domain Name System (DNS) name.) |
   | *username* | Specifies the name of an administrative account in that domain. |

3. When you are prompted for a password, type the password for the user account that you provided, and then press ENTER.

You can also use Repadmin to generate the details of replication to and from all replication partners in a spreadsheet. The spreadsheet displays data in the following columns:

Showrepl_COLUMNS

Destination DC Site

Destination DC

Naming Context

Source DC Site

Source DC

Transport Type

Number of Failures

Last Failure Time

Last Success Time

Last Failure Status

The following procedure shows how to create this spreadsheet and set column headers for improved readability.

**To generate a repadmin /showrepl spreadsheet for all replication partners**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl * /csv >showrepl.csv**
3. Open Microsoft Excel.
4. On the **File** menu, click **Open**, navigate to showrepl.csv, and then click **Open**.
5. Hide or delete column **A** as well as the **Transport Type** column, as follows:
6. Select a column that you want to hide or delete.
   - To hide the column, on the **Format** menu, click **Column**, and then click **Hide**.

     Or

   - To delete the column, right-click the selected column, and then click **Delete**.
7. Select row 1 beneath the column heading row, and then, on the **Window** menu, click **Freeze Panes**.
8. Select the entire spreadsheet. On the **Data** menu, click **Filter**, and then click **AutoFilter**.
9. In the **Last Success Time** column, click the down arrow, and then click **Sort Ascending**.
10. In the **Source DC** column, click the down arrow, and then click **Custom**.
11. In the **Custom AutoFilter** dialog box, under **Show rows where**, click **does not contain**. In the adjacent text box, type **del** to eliminate from view the results for deleted domain controllers.
12. Repeat step 10 for the **Last Failure Time** column, but use the value **does not equal**, and type the value **0**.
13. Resolve replication failures.

The last successful attempt should agree with the replication schedule for intersite replication, or the attempt should be within the last hour for intrasite replication.

If Repadmin reports any of the following conditions, see Troubleshooting Active Directory Replication Problems:

- The last successful intersite replication was prior to the last scheduled replication.
- The last intrasite replication was longer than one hour ago.
- Replication was never successful.

**See Also**
**Concepts**
Troubleshooting Active Directory Replication Problems

# Determine whether a domain controller is a global catalog server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the setting on the NTDS Settings object to indicate whether a domain controller is designated as a global catalog server.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group in Active Directory.

**To determine whether a domain controller is a global catalog server**

1. Open Active Directory Sites and Services.
2. In the console tree, expand the **Sites** container, expand the site of the domain controller you want to check, expand the **Servers** container, and then expand the **Server** object.
3. Right-click the **NTDS Settings** object, and then click **Properties**.
4. On the **General** tab, if the **Global Catalog** box is selected, the domain controller is designated as a global catalog server.

# Install the Schema snap-in

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to install the Active Directory Schema snap-in.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory.

**To install the Active Directory Schema snap-in**

1. Open a Command Prompt.
2. Type the following command and then press ENTER:

   **Regsvr32 schmmgmt.dll**

   This command will register schmmgmt.dll on your computer.
3. Click **Start**, click **Run**, type **mmc /a**, and then click **OK**.
4. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
5. Under **Available Standalone Snap-ins**, double-click **Active Directory Schema**, click **Close**, and then click **OK**.
6. To save this console, on the **File** menu, click **Save**.
7. In **Save in**, point to the *systemroot***system32** directory.
8. In **File** name, type **schmmgmt.msc**, and then click **Save**.
9. To create a shortcut on your Start menu:
   a. Right-click **Start**, click **Open All Users**, double-click the **Programs** folder, and then double-click the **Administrative Tools** folder.
   b. On the **File** menu, point to **New**, and then click **Shortcut**.
   c. In the **Create Shortcut Wizard**, in **Type the location of the item**, type schmmgmt.msc, and then click **Next**.
   d. On the **Select a Title for the Program** page, in **Type a name for this shortcut**, type **Active Directory Schema**, and then click **Finish**.

| ⚠Caution |
| --- |
| Modifying the schema is an advanced operation best performed by experienced programmers and system administrators. For detailed information about modifying the schema, see the Active Directory programmer's Guide at the Microsoft Web site. |

# Transfer the schema master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to transfer the schema operations master role. The schema master is a forest-wide operations master role. Before you can use the Active Directory Schema snap-in for the first time, you must register it with the system. If you have not yet prepared the Active Directory Schema snap-in, see Install the Schema snap-in before you begin this procedure.

| **Note** |
| --- |
| This procedure is performed by using the Microsoft Management Console (MMC), although you can also transfer this role by using Ntdsutil.exe. For information about using Ntdsutil.exe to transfer operations master roles, type **?** at the Ntdsutil.exe command prompt. |

### Administrative Credentials

To perform this procedure, you must be a Schema Administrator in Active Directory.

#### Transfer the schema master

1. Open the Active Directory Schema snap-in.
2. In the console tree, right-click **Active Directory Schema**, and click **Change Domain Controller**.
3. In the **Change Domain Controller** dialog box, click **Specify Name**. Then, in the text box, type the name of the server to which you want to transfer the schema master role. Click **OK**.
4. In the console tree, right-click **Active Directory Schema**. Click **Operations Master**. The **Change Schema Master** box displays the name of the server that is currently holding the role. The targeted domain controller is listed in the second box.
5. Click **Change**. Click **Yes** to confirm your choice. The system confirms the operation. Click **OK** again to confirm that the operation succeeded.
6. Click **Close** to close the **Change Schema Master** dialog box.

| **Note** |
| --- |
| Hosting the infrastructure master on a global catalog server is not recommended. If you attempt to transfer the infrastructure master role to a domain controller that is a global catalog, the system displays a warning stating that this is not recommended. |

# Transfer the domain naming master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to transfer the domain naming operations master role. The domain naming master is a forest-wide operations master role.

| **Note** |
| --- |
| This procedure is performed by using the Microsoft Management Console (MMC), although you can also transfer this role by using Ntdsutil.exe. For information about using Ntdsutil.exe to transfer operations master roles, type **?** at the Ntdsutil.exe command prompt. |

### Administrative Credentials

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

#### To transfer the domain naming master

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click **Active Directory Domains and Trusts**, and then click **Connect to Domain Controller**.
3. Ensure that the proper domain name is entered in the **Domain** box.

   The available domain controllers from this domain are listed.
4. In the **Name** column, click the domain controller (to select it) to which you want to transfer the role. Click **OK**.
5. Right-click **Active Directory Domains and Trusts**, and then click **Operations Master**.
6. The name of the current domain naming master appears in the first text box. The server to which you want to transfer the role should appear in the second text box. If this is not the case, repeat steps 1 through 4.
7. Click **Change**. To confirm the role transfer, click **Yes**. Click **OK** again to close the message box indicating the transfer took place. Click **Close** to close the **Change Operations Master** dialog box.

# Transfer the domain-level operations master roles

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to transfer the three domain-level operations master roles: the PDC emulator, the RID master, and the infrastructure master. You can transfer all of these roles by using the Active Directory Users and Computers console.

| **Note** |
| --- |
| These procedures are performed by using MMC, although you can also transfer these roles by using Ntdsutil.exe. For information about using Ntdsutil.exe to transfer the operations master roles, type **?** at the Ntdsutil.exe command prompt. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To transfer a domain-level operations master role**

1. Open Active Directory Users and Computers.
2. At the top of the console tree, right-click **Active Directory Users and Computers**. Click **Connect to Domain Controller**.
3. In the list of **available domain controllers**, click the name of the server to which you want to transfer the role, and then click **OK**.
4. At the top of the console tree, right-click **Active Directory Users and Computers**, point to **All Tasks**, and then click **Operations Masters**.

   The name of the current operations master role holder appears in the **Operations master** box. The name of the server to which you want to transfer the role appears in the lower box.
5. Click the tab for the role you want to transfer: **RID**, **PDC**, or **Infrastructure**. Verify the computer names that appear and then click **Change**. Click **Yes** to transfer the role, and then click **OK**.
6. Repeat steps 4 and 5 for each role that you want to transfer.

# View the current operations master role holders

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Once an operations master role has been transferred, it should be verified that the transfer has occurred successfully throughout the domain. The change must be replicated to all relevant domain members in order to truly take effect.

To view the current operations master role holders, use Ntdsutil.exe with the roles option. This option displays a list of all current role holders.

**Administrative Credentials**

To perform this procedure, you must be logged on as a User or an Administrator.

**To view the current operations master role holder**

1. Click **Start**, click **Run**, type **ntdsutil**, and then press ENTER.
2. At the **ntdsutil:** prompt, type **roles** and press ENTER.
3. At the **fsmo maintenance:** prompt, type **connections** and press ENTER.
4. At the **server connections:** prompt, type **connect to server** *servername* (where *servername* is the name of the domain controller that belongs to the domain containing the operations masters).
5. After receiving confirmation of the connection, type **quit** and press ENTER to exit this menu.
6. At the **fsmo maintenance:** prompt, type **select operation target** and press ENTER.
7. At the **select operations target:** prompt, type **list roles for connected server** and press ENTER.

   The system responds with a list of the current roles and the Lightweight Directory Access Protocol (LDAP) name of the domain controllers currently assigned to host each role.
8. Type **quit** and press ENTER to exit each prompt in Ntdsutil.exe. Type **quit** and press ENTER at the **ntdsutil:** prompt to close the window.

# Seizing an operations master role

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Role seizure is the act of assigning an operations master role to a new domain controller without the cooperation of the current role holder (usually because it is offline due to a hardware failure). During role seizure, a new domain controller assumes the operations master role without communicating with the current role holder.

Role seizure can create two conditions that can cause problems in the directory. It is for this reason that role seizure should be performed only as a last resort. First, the new role holder starts performing its duties based on the data located in its current directory partition. The new role holder might not receive changes that were made to the previous role holder before it went offline if replication did not complete prior to the time when the original role holder went offline. This can cause data loss or introduce data inconsistency into the directory database.

To minimize the risk of losing data to incomplete replication, do not perform a role seizure until enough time has passed to complete at least one complete end-to-end replication cycle across your network. Allowing enough time for complete end-to-end replication ensures that the domain controller that assumes the role is as up-to-date as possible.

Second, the original role holder is not informed that it is no longer the operations master role holder, which is not a problem if the original role holder stays offline. However, if it comes back online (for example, if the hardware is repaired or the server is restored from a backup), it might try to perform the operations master role that it previously owned. This can result in two domain controllers performing the same operations master role simultaneously. Depending on the role that was seized, the severity of duplicate operations master roles varies from no visible effect to potential corruption of the Active Directory database. Seize the operations master role to a domain controller that has the most recent updates from the current role holder to minimize the impact of the role seizure.

**Task Requirements**

- Repadmin.exe
- Ntdsutil.exe

To complete this task, perform the following procedures:

1. Verify successful replication to a domain controller

   This needs to be the domain controller that will be seizing the role.

2. Seize the operations master role

3. View the current operations master role holders

# Verify successful replication to a domain controller

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **repadmin /showrepl** command to verify successful replication to a specific domain controller. If you are not running Repadmin on the domain controller whose replication you are checking, you can specify a destination domain controller in the command. Repadmin lists INBOUND NEIGHBORS for the current or specified domain controller. INBOUND NEIGHBORS shows the distinguished name of each directory partition for which inbound directory replication has been attempted, the site and name of the source domain controller, and whether replication succeeded or not, as follows:

- Last attempt @ YYYY-MM-DD HH:MM.SS was successful.
- Last attempt @ [Never] was successful.

If @ [Never] appears in the output for a directory partition, replication of that directory partition has never succeeded from the identified source replication partner over the listed connection.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the destination domain controller.

**To verify successful replication to a domain controller**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl** *servername* **/u:** *domainname* **** *username* **/pw:***

   | Term | Definition |
   | --- | --- |
   | *servername* | Specifies the name of the destination domain controller. |
   | *domainname* | Specifies the single-label name of the domain of the destination domain controller. (You do not have to use a fully qualified Domain Name System (DNS) name.) |
   | *username* | Specifies the name of an administrative account in that domain. |

3. When you are prompted for a password, type the password for the user account that you provided, and then press ENTER.

You can also use Repadmin to generate the details of replication to and from all replication partners in a spreadsheet. The spreadsheet displays data in the following columns:

Showrepl_COLUMNS

Destination DC Site

Destination DC

Naming Context

Source DC Site

Source DC

Transport Type

Number of Failures

Last Failure Time

Last Success Time

Last Failure Status

The following procedure shows how to create this spreadsheet and set column headers for improved readability.

**To generate a repadmin /showrepl spreadsheet for all replication partners**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl * /csv >showrepl.csv**
3. Open Microsoft Excel.
4. On the **File** menu, click **Open**, navigate to showrepl.csv, and then click **Open**.
5. Hide or delete column **A** as well as the **Transport Type** column, as follows:
6. Select a column that you want to hide or delete.
   - To hide the column, on the **Format** menu, click **Column**, and then click **Hide**.

     Or

   - To delete the column, right-click the selected column, and then click **Delete**.
7. Select row 1 beneath the column heading row, and then, on the **Window** menu, click **Freeze Panes**.
8. Select the entire spreadsheet. On the **Data** menu, click **Filter**, and then click **AutoFilter**.
9. In the **Last Success Time** column, click the down arrow, and then click **Sort Ascending**.
10. In the **Source DC** column, click the down arrow, and then click **Custom**.
11. In the **Custom AutoFilter** dialog box, under **Show rows where**, click **does not contain**. In the adjacent text box, type **del** to eliminate from view the results for deleted domain controllers.
12. Repeat step 10 for the **Last Failure Time** column, but use the value **does not equal**, and type the value **0**.
13. Resolve replication failures.

The last successful attempt should agree with the replication schedule for intersite replication, or the attempt should be within the last hour for intrasite replication.

If Repadmin reports any of the following conditions, see Troubleshooting Active Directory Replication Problems:

- The last successful intersite replication was prior to the last scheduled replication.
- The last intrasite replication was longer than one hour ago.
- Replication was never successful.

**See Also**
**Concepts**
Troubleshooting Active Directory Replication Problems

# Seize the operations master role

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Ntdsutil.exe command-line tool allows you to transfer and seize any operations master role. You must use Ntdsutil.exe to seize the schema master, domain naming master, and RID master roles. When you use Ntdsutil.exe to seize an operations master role, it first attempts a transfer from the current role owner. If the current role owner is unavailable, it performs the seizure.

When using Ntdsutil.exe to seize an operations master role, the procedure is nearly identical for all roles. For more information about using Ntdsutil.exe, type **?** at the Ntdsutil.exe command prompt.

**Administrative Credentials**

To perform this procedure, you must be a member of either the Domain Admins group or the Enterprise Admins group in Active Directory.

**To seize an operations master role**

1. Click **Start**, click **Run**, type **ntdsutil**, and then press ENTER.
2. At the **ntdsutil:** prompt, type **roles** and press ENTER.
3. At the **fsmo maintenance:** prompt, type **connections** and press ENTER.
4. At the **server connections:** prompt, type **connect to server** *servername* (where *servername* is the name of the domain controller that will assume the operations master role), and press ENTER.
5. After you receive confirmation of the connection, type **quit** and press ENTER\.
6. Depending on the role you want to seize, at the **fsmo maintenance:** prompt, type the appropriate command and press ENTER.

| Role | Credentials | Command |
|---|---|---|
| Domain naming master | Enterprise Admins | Seize domain naming master |
| Schema master | Enterprise Admins | Seize schema master |
| Infrastructure master | Domain Admins | Seize infrastructure master |
| PDC emulator | Domain Admins | Seize pdc |
| RID master | Domain Admins | Seize rid master |

The system asks for confirmation. It then attempts to transfer the role. When the transfer fails, some error information appears and the system proceeds with the seizure. After the seizure is complete, a list of the roles and the LDAP name of the server that currently holds each role appears.

During seizure of the RID master, the current role holder attempts to synchronize with its replication partners. If it cannot establish a connection with a replication partner during the seizure operation, it displays a warning and confirms that you want the role seizure to proceed. Click **Yes** to proceed.

7. Type **quit** and press ENTER. Type **quit** again and press ENTER to exit Ntdsutil.exe.

# View the current operations master role holders

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Once an operations master role has been transferred, it should be verified that the transfer has occurred successfully throughout the domain. The change must be replicated to all relevant domain members in order to truly take effect.

To view the current operations master role holders, use Ntdsutil.exe with the roles option. This option displays a list of all current role holders.

**Administrative Credentials**

To perform this procedure, you must be logged on as a User or an Administrator.

**To view the current operations master role holder**

1. Click **Start**, click **Run**, type **ntdsutil**, and then press ENTER.
2. At the **ntdsutil:** prompt, type **roles** and press ENTER.
3. At the **fsmo maintenance:** prompt, type **connections** and press ENTER.
4. At the **server connections:** prompt, type **connect to server** *servername* (where *servername* is the name of the domain controller that belongs to the domain containing the operations masters).
5. After receiving confirmation of the connection, type **quit** and press ENTER to exit this menu.
6. At the **fsmo maintenance:** prompt, type **select operation target** and press ENTER.
7. At the **select operations target:** prompt, type **list roles for connected server** and press ENTER.

   The system responds with a list of the current roles and the Lightweight Directory Access Protocol (LDAP) name of the domain controllers currently assigned to host each role.
8. Type **quit** and press ENTER to exit each prompt in Ntdsutil.exe. Type **quit** and press ENTER at the **ntdsutil:** prompt to close the window.

# Reducing the workload on the PDC emulator master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In addition to processing normal domain controller load from clients, the PDC emulator must also process password changes. In order to mitigate some of the load that is caused by normal domain controller traffic, the PDC can be protected, so the load is distributed to other domain controllers that are capable of processing the requests.

You can configure DNS so that a domain controller is queried less frequently than others. Reducing the number of client requests helps reduce the workload on a domain controller, giving it more time to function as an operations master, and is especially important for the PDC emulator. Of all the operations master roles, the PDC role has the highest impact on the domain controller hosting that role.

To receive information from the domain, a client uses DNS to locate a domain controller and then sends the request to that domain controller. By default, DNS performs rudimentary load balancing and randomizes the distribution of client requests so they are not always sent to the same domain controller. If too many client requests are sent to a domain controller while it attempts to perform other duties, such as those of the PDC emulator, it can become overloaded, which has a negative impact on performance. To reduce the number of client requests that are processed by the PDC emulator, you can adjust its weight or its priority in the DNS environment.

**Adjusting the Weight for DNS SRV Records in the Registry**

Adjusting the weight of a domain controller to a value less than that of other domain controllers reduces the number of clients that DNS refers to that domain controller. The value is stored in the LdapSrvWeight registry entry. The default value is 100, but it can range from 0 through 65535. By reducing this value, DNS refers clients to a domain controller less frequently based on the proportion of this value to the value of other domain controllers. For example, to configure the system so that the domain controller hosting the PDC emulator role receives requests only half as many times as the other domain controllers, configure the weight of the domain controller hosting the PDC emulator role to be 50. DNS determines the weight ratio for that domain controller to be 50/100 (50 for that domain controller and 100 for the other domain controllers). After you reduce this ratio to 1/2, DNS refers clients to the other domain controllers twice as often as it refers to the domain controller with the reduced weight setting. By reducing client referrals, the domain controller receives fewer client requests and has more resources for other tasks, such as performing the role of PDC emulator.

**Adjusting the Priority for DNS SRV Records in the Registry**

Adjusting the priority of the domain controller also reduces the number of client referrals. However, rather than reducing it proportionally to the other domain controllers, changing the priority causes DNS to stop referring all clients to this domain controller unless all domain controllers with a lower priority setting are unavailable.

To prevent clients from sending all requests to a single domain controller, the domain controllers are assigned a priority value. Clients always send requests to the domain controller that has the lowest priority value. If more than one domain controller has the same value, the clients randomly choose from the group of domain controllers with the same value. If no domain controllers with the lowest priority value are available, then the clients send requests to the domain controller with the next highest priority.

A domain controller's priority value is stored in its registry. When the domain controller starts, the Net Logon service registers with the DNS server. The priority value is registered with the rest of its DNS information. When a client uses DNS to discover a domain controller, the priority for a given domain controller is returned to the client with the rest of the DNS information. The client uses the priority value to help determine to which domain controller to send requests.

The value is stored in the **LdapSrvPriority** registry entry. The default value is 0, but it can range from 0 through 65535.

| **Note** |
| --- |
| A lower value that is entered for **LdapSrvPriority** indicates a higher priority. Clients will attempt to use a domain controller with an **LdapSrvPriority** setting of 0 before a domain controller with an **LdapSrvPriority** setting of 1, just as domain controllers with **LdapSrvPriority** settings of 1 would be used before domain controllers with **LdapSrvPriority** settings of 2. |

**Task Requirements**

The following tool is required to perform the procedures for this task:

- Regedit.exe

To complete this task, perform the following procedures:

1. Change the weight for DNS SRV records in the registry

2. Change the priority for DNS SRV records in the registry

# Change the weight for DNS SRV records in the registry

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to reduce the workload on the PDC emulator master by changing the weight for DNS SRV records in the registry.

| ⚠Caution |
| --- |
| The Registry Editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back up system state first. For information about backing up system state, see Administering Active Directory Backup and Restore. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

#### To change the weight for DNS SRV records in the registry

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.
2. In the Registry Editor, navigate to **HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**
3. Click **Edit**, click **New**, and then click **DWORD** value.
4. For the new value name, type **LdapSrvWeight**, and press ENTER.
5. Double-click the value name that you just typed to open the Edit DWORD Value dialog box.
6. Enter a value from 0 through 65535. The default value is 100.
7. Choose **Decimal** as the **Base** option, and then click **OK**.
8. Click **File**, and then click **Exit** to close the Registry Editor.

# Change the priority for DNS SRV records in the registry

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to reduce the workload on the PDC emulator master by changing the priority for DNS SRV records in the registry.

| ⚠Caution |
| --- |
| The Registry Editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back up system state first. For information about backing up system state, see Administering Active Directory Backup and Restore. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To change the priority for DNS SRV records in the registry**

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.
2. In the Registry Editor, navigate to **HKLM\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**
3. Click **Edit**, click **New**, and then click **DWORD** value.
4. For the new value name, type **LdapSrvPriority**, and press ENTER.
5. Double-click the value name that you just typed to open the Edit DWORD Value dialog box.
6. Enter a value from 0 through 65535. The default value is 0.
7. Choose **Decimal** as the **Base** option, and then click **OK**.
8. Click **File**, and then click **Exit** to close the Registry Editor.

# Administering Active Directory Backup and Restore

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This Administering Active Directory Backup and Restore guide provides administering information for Active Directory backup and restore in the Microsoft Windows Server 2003 operating system.

**In this guide**


- Introduction to Administering Active Directory Backup and Restore
- Managing Active Directory Backup and Restore



**Acknowledgements**

Produced by: Microsoft Windows Server User Assistance team

Writer: Mary Hillman

Editor: Jim Becker

# Introduction to Administering Active Directory Backup and Restore

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory backup must be incorporated into your operations schedule for a set of domain controllers that you identify and on which you perform routine backup operations.

Active Directory restore is not performed routinely as an operations task; it is performed only when indicated by a failure or other condition from which a domain controller can recover only by restoring the directory to a previous state.

**System State Components**

Active Directory is backed up as part of system state, a collection of system components that depend on each other. All system state components must be backed up and restored together.

The system state components on a domain controller include the following:

- System startup (boot) files. These files are required for Windows Server 2003 to start.
- System registry
- Class registration database of component services. The Component Object Model (COM) is a binary standard for writing component software in a distributed systems environment.
- System volume (SYSVOL). SYSVOL provides a default location for files that must be shared for common access throughout a domain. The SYSVOL folder on a domain controller contains the following:
    - Net Logon shared folders. These folders usually host user logon scripts and policy settings for network clients that are running pre–Windows 2000 operating systems.
    - User logon scripts for Active Directory–enabled clients
    - System policies
    - Group Policy settings
    - File system junctions
    - File Replication service (FRS) staging directories and files that are required to be available and synchronized between domain controllers
- Active Directory, including the following:
    - The Active Directory database (Ntds.dit)
    - The checkpoint file (Edb.chk)
    - The transaction logs, each 10 megabytes (MB) in size (Edb*.log)
    - Reserved transaction logs (Res1.log and Res2.log)

If you installed Windows Clustering or Certificate Services on your domain controller, they are also backed up as part of system state. Details of these components are not discussed in this guide.

**Purpose of Performing Regular Backups**

You need a current, verified, and reliable backup to:

- Restore Active Directory data that becomes lost. By using an authoritative restore process, you can restore individual objects or sets of objects (containers or directory partitions) from their deleted state.

- Recover a domain controller that cannot start up or operate normally because of software failure or hardware failure.

- Install Active Directory from backup media (using the **dcpromo /adv** command). You can use this installation option of Dcpromo to install Active Directory on a server running Windows Server 2003 to make that server an additional domain controller. Use this method to quickly add a domain controller to a domain that has a large database or that is located in sites that are separated by slow network links.

- Perform a forest recovery if forest-wide failure occurs.

**Restore Requirements and Recommendations**

The tombstone lifetime value in an Active Directory forest defines the default number of days that a domain controller preserves knowledge of deleted objects. This value also defines the useful life of a system state backup that is used for disaster recovery or installation from backup media. Active Directory protects itself from restoring data that is older than the tombstone lifetime by disallowing the restore.

The tombstone lifetime is stored in the **tombstoneLifetime** attribute of the object CN=Directory Service,CN=Windows NT,CN=Services,CN-Configuration,DC=*ForestRootDomain*. In a forest that is created on a domain controller that is running Windows Server 2003 without Service Pack 1 (SP1), the default tombstone lifetime is 60 days. In a forest that is created on a domain controller that is running Windows Server 2003 with SP1, Windows Server 2003 with Service Pack 2 (SP2), or Windows Server 2003 R2, the default tombstone lifetime is 180 days. Because the tombstone lifetime can be changed administratively, do not assume the default value. Be sure that you are aware of the tombstone lifetime that is in effect in your forest.

| **Important** |
| --- |
| You should not modify system clocks in an attempt to improperly extend the useful life of a system state backup. |

System state restore should be undertaken as a last resort, not as primary method of recovering from an error or failure condition.

**Backup Guidelines**

The following guidelines for backup include the performance of appropriate backups to ensure redundancy of Active Directory data:

- Perform normal backup. Normal backup is the only type of backup that is available and supported for Active Directory. The Backup tool in Windows Server 2003 supports multiple types of backup: normal, copy, incremental, differential, and daily. You must use normal backup because Active Directory is backed up as part of system state.

- Perform daily backups of each unique partition on at least two unique domain controllers, with special emphasis on single-domain controller forests, single-domain controller domains, and empty root domains.

- Where partitions exist in only one site, you can ship backup files offsite to a secure location so that no backup file of a unique directory partition exists in only one physical site at any point in time. This provides an extra level of redundancy.

- Make sure your backups are stored in a secure location at all times.

- Back up Domain Name System (DNS) zones. You must be aware of the location of DNS zones and back up DNS servers accordingly. If you use Active Directory-integrated DNS, DNS zone data is captured as part of system state on domain controllers that are also DNS servers.

  If you do not use Active Directory-integrated DNS, you must back up the zone file directories on a representative set of DNS servers for each DNS zone to ensure fault tolerance for the zone.

| **Note** |
| --- |

The DNS server stores settings in the registry, so system state backup is required for DNS regardless of whether the zone data is Active Directory-integrated or stored in the file system.

- If you have application partitions in your forest, make sure that you take a backup of the domain controllers that hold those application partitions.
- Create additional backups in every geographic location where:
  - Mission-critical work is performed.
  - A wide area network (WAN) outage would disrupt business.
  - The elapsed time that it takes to perform either of the following tasks would be cost-prohibitive because of slow link speeds, the size of the directory database, or both:

    To create a domain controller in its intended domain over the network.

    Or

    To copy or transport a system state backup from a site where a backup exists to a site that has no backup, for the purpose of performing an installation from backup media.

**Note**

A backup can be used to restore only the domain controller on which the backup was generated or to create a new additional domain controller in the same domain by installing from backup media. A backup cannot be used to restore a different domain controller or to restore a domain controller onto different hardware. Likewise, a backup that is made on a domain controller running Windows 2000 Server cannot be used to restore a domain controller running Windows Server 2003.

**Backup Frequency**

Backup frequency depends on criteria that vary for individual environments. In most Active Directory environments, users, computers, and administrators make daily changes to directory objects. For example, computer accounts, including domain controller accounts, change their passwords every 30 days by default. Therefore, every day a percentage of computer passwords changes for domain controllers. Rolling the computer password of a domain controller back to a former state affects authentication and replication. A percentage of user passwords might also expire on a daily basis, and if they are lost as a result of domain controller failure, they must be reset manually. Generally, no record of these changes exists except in Active Directory. Therefore, the more frequently you back up domain controllers, the fewer problems you will encounter if you need to restore.

The more Active Directory objects and domain controllers you have, the more frequent your backups should be. For example, in a large organization, to recover from the inadvertent deletion of a large organizational unit (OU) by restoring the domain from a backup that is days or weeks old, you might have to re-create hundreds of accounts that were created in that OU since the backup was taken. To avoid re-creating accounts and potentially performing large numbers of manual password resets, ensure that recent system state backups are always available to recover recent Create, Modify, and Delete operations.

**Frequency Criteria**

Use the following criteria to assess backup frequency:

- Small environments with a single domain controller in the forest, or domains that exist in a single physical location (that is, that have a single point of failure): create backups at least daily.
- Medium (10 to 49 domain controllers) and large environments (50 to 1,000 or more domain controllers): Create backups of each unique directory partition in the forest on two different computers at least daily with an emphasis on backing up application directory partitions, empty root domains, domain partitions in a single geographic site, and sites that have large populations of users or that host mission-critical work.

Make backups with increasing frequency until you are confident that if you were to lose the objects that were created or modified since the last backup, the loss would not create an operational disruption. For this reason, major changes to the environment should always be immediately followed by a new system state backup.

| **Note** |
| --- |
| It is always recommended that you have at least two domain controllers in each domain of your Active Directory forest |

**Immediate Backup**

In addition to regularly scheduled backups, perform an immediate backup when:

- You have moved the Active Directory database, log files, or both to a different location on a disk.
- A domain controller is upgraded from Windows 2000 Server to Windows Server 2003 or there are any other operating system upgrades.
- A Service Pack is installed.
- A hotfix is installed that makes changes to the Active Directory database.
- A current backup is required for installing from backup media for a new domain controller.
- The tombstone lifetime is changed administratively.

**Backup Latency Interval**

On domain controllers running Windows Server 2003 with Service Pack 1 (SP1), Windows Server 2003 with Service Pack 2 (SP2), or Windows Server 2003 R2, event ID 2089 provides the backup status of each directory partition that a domain controller stores, including application directory partitions. Specifically, event ID 2089 is logged in the Directory Service event log when partitions in the Active Directory forest are not backed up within a backup latency interval. The value for the backup latency interval is stored as a **REG_DWORD** value in the **Backup Latency Threshold (days)** entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

| **Note** |
| --- |
| An initial backup must be performed before event ID 2089 will be logged. If no backup is ever performed on a domain controller, this event does not provide warnings of backup latency. |

By default, the value of Backup Latency Threshold (days) is half the value of the tombstone lifetime of the forest. In a forest that is created on a domain controller that is running Windows Server 2003 without Service Pack 1 (SP1), half the default tombstone lifetime is 30 days. In a forest that is created on a domain controller that is running Windows Server 2003 with SP1, Windows Server 2003 with Service Pack 2 (SP2), or Windows Server 2003 R2, half the default tombstone lifetime is 90 days. However, we recommend that you make backups at a much higher frequency than half the tombstone lifetime. By setting a minimum backup frequency, changing **Backup Latency Threshold (days)** to reflect that frequency, and monitoring Event ID 2089, you ensure the backup frequency that is established in your organization.

To set a different **Backup Latency Threshold (days)** value, use Registry Editor (Regedit.exe) to create the entry as a **REG_DWORD** and provide the appropriate number of days.

**See Also**

**Concepts**

Installing a Domain Controller in an Existing Domain Using Restored Backup Media

# Managing Active Directory Backup and Restore

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following tasks for managing Active Directory backup and restore are described in this objective:

- Backing Up Active Directory Components
- Performing a Nonauthoritative Restore of a Domain Controller
- Performing an Authoritative Restore of Active Directory Objects
- Performing an Authoritative Restore of an Application Directory Partition
- Performing an Authoritative Restore of a Group Policy Object
- Restoring a Domain Controller Through Reinstallation and Subsequent Restore from Backup
- Restoring a Domain Controller Through Reinstallation

# Backing Up Active Directory Components

Updated: April 15, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory is backed up as part of Microsoft Windows system state. For more information about all Microsoft Windows system state components and Active Directory components, see Introduction to Administering Active Directory Backup and Restore.

**Naming Backup Files**

To ensure the proper use of backup files, the .bkf file should have a file name that includes the following information:

- The fully qualified computer name that includes the domain name of the domain controller on which the backup was performed
- Whether the backup domain controller is a global catalog server
- Whether the backup domain controller contains MD5 checksum data to source the Sysvol tree
- The date that the backup was performed

For example, you might use a file name format that is similar to the following:

X:*Fully_Qualified_Computer_Name.Build_Number.Service_Pack_Revision*. [No]GC.[No]MD5.*TSL.YYYYMMDD*.bkf

where

- *Fully_Qualified_Computer_Name* is the host name and the domain name of the domain controller. This must be the domain name of the domain where the system state was backed up.
- *Build Number* is the build number of the operating system that was backed up.
- *Service_Pack_Revision* is the service pack build number and the service pack version for the operating system that was backed up.
- [No]GC indicates whether the backup originated from a global catalog or not.
- [No]MD5 indicates whether the system state backup contains MD5 checksum data for the files and folders in the SYSVOL tree. For more information about the need for MD5 data, see Preparing a Server Computer for Shipping and Installation from Backup Media.
- *TSL* is the value in days for the **tombstoneLifetime attribute** when the backup was performed. The **tombstoneLifetime** attribute for the forest determines both the useful life of a system state backup and how frequently garbage collection occurs. (Garbage collection removes tombstones from the directory permanently when their tombstone lifetime expires.)
- *YYYYMMDD* is the year, month, and day that the backup was performed.

For example, suppose that you create a system state backup of a global catalog domain controller on July 1, 2005. The domain controller is in the Contoso.com domain, and its name is DC1. The value of the tombstone lifetime is 60 days, and MD5 data is included in the backup. In this scenario, you might use a file name that is similar to the following:

DC1.CONTOSO.COM.3790.SP0.GC.MD5.60.2005.07.01.BKF

A system state backup that you make of DC1 on July 1, 2005, remains valid until August 29, 2005. For the next 60 days, you can use the backup to restore an existing domain controller or to install an additional domain controller in the Contoso.com domain.

You can save the .bkf file to a local volume or to a network share. The network share can be on a server computer that can be installed later as a domain controller by using the restored backup. For more information about using restored backup media for installing domain controllers, see Installing a Domain Controller in an Existing Domain Using Restored Backup Media.

**Task requirements**

The following tools are required to perform the procedures for this task:

- Backup or Restore Wizard (Ntbackup)
- Tape drive or other backup media

To complete this task, perform one of the following procedures, depending on your backup needs:

- Back up system state
- Back up system state and the system disk

**See Also**
**Concepts**
Installing a Domain Controller in an Existing Domain Using Restored Backup Media
Adding Domain Controllers in Remote Sites

# Back up system state

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Ntbackup.exe provides simple and advanced options for backing up Active Directory components. When you back up system state, you can choose to include or exclude system-protected boot files. System-protected boot files are not used for installations from restored backup media. When the backup file that you create is to be used for additional domain controller installations, you can clear the advanced option to back up system-protected files. Clearing this option decreases the size of the .bkf file, as well as the time required to back up, restore, and copy the system state files.

Use these procedures to back up the system state only. These procedures do not back up the system disk or any other data on the domain controller except for the system-protected files.

Use the first procedure, "To back up system state including system-protected files," for routine system state backup. Use the second procedure, "To back up system state excluding system-protected files," if you want to create a smaller backup that is effective for installing domain controllers from restored backup media.

| **Note** |
| --- |
| To back up system state, you must log on locally to the domain controller, or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

- To perform the following two procedures, you must be a member of the Domain Admins group or a member of the Backup Operators group.

#### To back up system state including system-protected files

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.

   This procedure provides steps for backing up in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard.

2. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
3. Select **Back up files and settings**, and then click **Next**.
4. Select **Let me choose what to back up**, and then click **Next**.
5. In the **Items to Back Up** window, double-click **My Computer**.
6. In the expanded list below **My Computer**, check **System State**, and then click **Next**.
7. Select a location to store the backup:

   - If you are backing up to a file, type the path and file name for the backup (.bkf) file (or click **Browse** to find a folder or file).

   - If you are backing up to a tape unit, choose the tape that you want to use.

     | **Note** |
     | --- |
     | You should not store the backup on the local hard drive. Instead, store it in a location, such as a tape drive, away from the computer that you are backing up. |

8. Type a name for this backup according to the recommendations in Backing Up Active Directory Components, and then click **Next**.
9. On the last page of the wizard, click **Advanced**.

10. Do not change the default options for **Type of Backup**. **Normal** should be selected, and the check box for **Backup migrated remote storage data** should remain cleared. Click **Next**.
11. Select **Verify data after backup**, and then click **Next**.
12. In the **Backup Options** dialog box, select a backup option, and then click **Next**.
13. If you are replacing the existing backups, select the option to allow only the owner and administrator access to the backup data and to any backups that are appended to this medium, and then click **Next**.
14. In the **When to back up** box, select the appropriate option for your needs, and then click **Next**.
15. If you are satisfied with all of the options that are selected, click **Finish** to perform the backup operation according to your selected schedule.

| **Note** |
| --- |
| The system state can also be backed up by using Ntbackup from a command line with appropriate parameters. For more information, at a command prompt type **ntbackup /?**. |

The following procedure produces a smaller .bkf file that does not include system boot files. By using this procedure, you can reduce the time that is required to perform the backup and subsequent restore, as well as the amount of disk space that is required. This method is recommended when the restored backup is to be used for installing additional domain controllers.

**To back up system state excluding system-protected files**

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.
2. On the **Welcome to the Backup or Restore Wizard** page, click **Advanced Mode**, and then click the **Backup** tab.
3. In the console tree, select the **System State** check box.
4. In **Backup media or file name**, type a name for this backup according to the recommendations in Backing Up Active Directory Components.
5. Click **Start Backup**, and then click **Advanced**.
6. Clear the **Automatically back up System Protected Files with the System State** check box, and then click **OK**.
7. Click **Start Backup**.

**See Also**
**Concepts**
Enable Remote Desktop
Create a Remote Desktop Connection

# Back up system state and the system disk

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to back up both the system state and the system disk.

| **Note** |
| --- |
| To back up system state and the system disk, you must log on locally to the domain controller or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

To perform this procedure, you must be a member of the Domain Admins group or a member of the Backup Operators group.

**To back up system state and the system disk**

1. To start the Windows Server 2003 backup utility, click **Start**, point to **Programs**, **Accessories**, **System Tools**, and then click **Backup**.

   This procedure requires Backup Utility Advanced Mode. If the Welcome to the Backup or Restore Wizard appears when you click **Backup** in step 1, clear **Always start in wizard mode**, close the wizard, and then repeat step 1.
2. On the **Welcome to the Backup Utility Advanced Mode** page, click the **Backup Wizard (Advanced)** button.
3. On the **Welcome to the Backup Wizard** page, click **Next**.
4. On the **What to Back Up** page, click **Back up selected files, drives, or network data**, and then click **Next**.
5. In **Items to Back Up**, select the **System State** check box. Then, locate the drive letter containing the system files, click the check box for it, and then click **Next**.
6. In **Backup Type, Destination, and Name**, select the backup media type by choosing one of the following options:
   - In the **Select the backup type** box, click **File** if you want to back up to a file. If you do not have a tape backup unit installed, **File** is selected automatically.

     Or

   - Select a tape device if you want to back up to tape.
7. In the **Choose a place to save your backup** box, select one of the following options, and then click **Next**:
   - If you are backing up to a file, if you want to change the current backup file location, click **Browse** to find a folder or file. If the destination folder or file does not exist, the system creates it.

     Or

   - If you are backing up to a tape unit, select the tape that you want to use.
8. On the **Completing the Backup Wizard** page, click **Advanced**. Do not change the default options for **Type of Backup**. **Normal** should be selected, and the check box should remain cleared for **Backup migrated remote storage data**. Click **Next**.
9. Click **Verify data after backup**, and then click **Next**.
10. In the **Backup Options** dialog box, select a backup option, and then click **Next**.
11. If you are replacing the existing backups, select the option to allow only the owner and administrator access to the backup data and to any backups that are appended to this medium, and then click **Next**.
12. In the **When to back up** box, select the appropriate option for your needs, and then click **Next**.

13. If you are satisfied with all of the options that are selected, click **Finish** to perform the backup operation according to your selected schedule.

**See Also**
**Concepts**
Enable Remote Desktop
Create a Remote Desktop Connection

# Performing a Nonauthoritative Restore of a Domain Controller

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A nonauthoritative restore is the default method for restoring Active Directory. To perform a nonauthoritative restore, you must be able to start the domain controller in Directory Services Restore Mode. After you restore the domain controller from backup, replication partners use the standard replication protocols to update Active Directory and associated information on the restored domain controller.

A nonauthoritative restore returns the domain controller to its state at the time of backup and then allows normal replication to overwrite that state with any changes that occurred after the backup was taken. After you restore the system state, the domain controller queries its replication partners. The replication partners replicate any changes to the restored domain controller, ensuring that the domain controller has an accurate and updated copy of the Active Directory database.

A nonauthoritative restore allows the entire directory to be restored on a domain controller, without reintroducing or changing objects that have been modified since the backup. The most common use of a nonauthoritative restore is to bring an entire domain controller back, often after catastrophic or debilitating hardware failures. It is uncommon for data corruption to drive a nonauthoritative restore, unless the corruption is local and the database cannot be successfully loaded.

If you intend to restore a deleted object (or objects), see Performing an Authoritative Restore of Active Directory Objects.

You can perform a nonauthoritative restore on a Windows Server 2003 system that is a stand-alone server, member server, or domain controller. You must start a server in Directory Services Restore Mode to perform a nonauthoritative restore.

| **Note** |
| --- |
| By performing a nonauthoritative restore on Active Directory, you automatically perform a nonauthoritative restore of the system volume (SYSVOL); no additional steps are required. |

**Task requirements**

The following tool is required to perform the procedures for this task:

- NTBackup.exe

To complete this task, perform the following procedures:

1. Restart the domain controller in Directory Services Restore Mode by using one of the following methods:

   - Restart the domain controller in Directory Services Restore Mode locally
   - Restart the domain controller in Directory Services Restore Mode Remotely

   | **Note** |
   | --- |
   | In cases in which you have to reinstall the operating system, before you restore the directory, you do not have to perform a nonauthoritative restore in Directory Services Restore Mode. After you reinstall the operating system, you can perform a restore after the computer boots normally. |

2. Restore Active Directory from backup

3. Verify Active Directory restore

**See Also**
**Concepts**
Performing an Authoritative Restore of Active Directory Objects
Enable Remote Desktop
Create a Remote Desktop Connection

# Restart the domain controller in Directory Services Restore Mode locally

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have physical access to a domain controller, you can restart the domain controller in Directory Services Restore Mode locally. Restarting in Directory Services Restore Mode takes the domain controller offline. In this mode, the server is not functioning as a domain controller.

When you start Windows Server 2003 in Directory Services Restore Mode, the local Administrator account is authenticated by the local Security Accounts Manager (SAM) database. Therefore, logging on requires that you use the local administrator password, not an Active Directory domain password. This password is set during Active Directory installation when you provide the password for Directory Services Restore Mode.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restart the domain controller in Directory Services Restore Mode locally**

1. Restart the domain controller.
2. When the screen for selecting an operating system appears, press F8.
3. On the **Windows Advanced Options** menu, select **Directory Services Restore Mode**.
4. When you are prompted, log on as the local administrator.

**See Also**
**Concepts**
Restart the domain controller in Directory Services Restore Mode Remotely

# Restart the domain controller in Directory Services Restore Mode Remotely

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If Remote Desktop is enabled on a domain controller, you can use Remote Desktop Connection to connect to the domain controller remotely. Remote Desktop Connection (formerly known as the Terminal Services client) is installed by default on all Windows Server 2003 family operating systems.

If you use Remote Desktop Connection to connect to a domain controller remotely and you want to restart the domain controller in Directory Services Restore Mode, you must first modify the Boot.ini file on the remote server so that you do not lose the connection when the domain controller restarts.

When you start Windows Server 2003 in Directory Services Restore Mode, the local Administrator account is authenticated by the local Security Accounts Manager (SAM) database. Therefore, logging on requires that you use the local administrator password, not an Active Directory domain password. This password is set during Active Directory installation when you provide the password for Directory Services Restore Mode.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restart the domain controller in Directory Services Restore Mode remotely**

1. Use Remote Desktop Connection to connect to the domain controller that you want to restart in Directory Services Restore Mode.
2. Right-click **My Computer**, click **Properties**, and then click the **Advanced** tab.
3. Click **Settings for startup and recovery**.
4. Click the **Edit** button to edit the startup options file.
5. Modify the default entry to include the **/SAFEBOOT:DSREPAIR** switch, as shown in the following example:

   **multi(0)disk(0)rdisk(0)partition(2)\WINNT="W2K DC ** *your server name* **" /fastdetect /SAFEBOOT:DSREPAIR**

   | **Note** |
   | --- |
   | The **/SAFEBOOT:DSREPAIR** switch works for domain controllers running Windows 2000 Server and Windows Server 2003. |

6. Save the modified Boot.ini file, and then close Notepad.
7. On the **Start** menu, click **Shut Down**, and then click **Restart**. During the restart process, the Terminal Services client reports that the session is disconnected.

   | **Caution** |
   | --- |
   | Be sure to click **Restart** and not **Shut Down** at this step. If you click **Shut Down**, you cannot restart the domain controller remotely. |

8. Wait until the restart process completes on the remote domain controller, and then reconnect the client session.
9. When the client session is reconnected, log on as the local administrator.
10. Right-click **My Computer**, click **Properties**, and then click the **Advanced** tab.
11. Click **Settings for startup and recovery**.
12. Click the **Edit** button to edit the startup options file.
13. Delete the **/SAFEBOOT:DSREPAIR** switch from the default entry in the Boot.ini file, save the file, and then close Notepad.

| **Important** |
| --- |
| If you restart the domain controller before you modify the Boot.ini file, the domain controller remains offline. |

The Boot.ini file is now returned to its original state, which starts the domain controller normally.

**See Also**

**Concepts**

Enable Remote Desktop

Create a Remote Desktop Connection

Restart the domain controller in Directory Services Restore Mode locally

# Restore Active Directory from backup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To restore the server, use a good backup containing the system state or the system state and system disk.

| **Note** |
| --- |
| To restore from backup, you must log on locally to the domain controller or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restore Active Directory from backup**

1. Start the computer in Directory Services Restore Mode.
2. To start the Windows Server 2003 backup utility, click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

   This procedure provides steps for restoring from backup in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard.
3. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
4. Click **Restore files and settings**, and then click **Next**.
5. Select **System State**, and then click **Next**.
6. On the **Completing the Backup or Restore Wizard** page, click **Advanced**.
7. In **Restore files to**, click **Original Location**, and then click **Next**.
8. Click **Leave existing files (Recommended)**, and then click **Next**.
9. In **Advanced Restore Options**, select the following check boxes, and then click **Next**:
   - **Restore security settings**
   - **Restore junction points, but not the folders and file data they reference**
   - **Preserve existing volume mount points**
10. For a primary restore of SYSVOL, also select the following check box: **When restoring replicated data sets, mark the restored data as the primary data for all replicas**.

    A primary restore is required only if the domain controller that you are restoring is the only domain controller in the domain. A primary restore is required on the first domain controller that is being restored in a domain if you are restoring the entire domain or forest.
11. Click **Finish**.
12. When the restore process is complete, click **Close**, and then do one of the following:
    - If you do not want to authoritatively restore any objects, click **Yes** to restart the computer. The system will restart and replicate any new information that is received since the last backup with its replication partners.
    - If you want to authoritatively restore any objects or if you want to create an LDAP Data Interchange Format (LDIF) file to restore back-links on this domain controller, click **No** to remain in Directory Services Restore Mode. For information about how to proceed with authoritative restore, see Performing an Authoritative Restore of Active

Directory Objects.

**See Also**
**Concepts**
Restart the domain controller in Directory Services Restore Mode locally
Enable Remote Desktop
Create a Remote Desktop Connection
Restart the domain controller in Directory Services Restore Mode Remotely
Restore system state to an alternate location
Performing an Authoritative Restore of Active Directory Objects

# Verify Active Directory restore

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After the restore is completed, use this procedure to restart the server and perform basic verification.

**Administrative credentials**

To verify Active Directory restore, you must be a member of the Domain Admins group.

**To verify Active Directory restore**

1. After the restore operation completes, restart the computer in **Start Windows Normally** mode. Active Directory and Certificate Services automatically detect that they have been recovered from a backup. They perform an integrity check and re-index the database.
2. After you are able to log on to the system, browse Active Directory. Verify that all of the User objects and Group objects that were present in the directory prior to backup are restored. Similarly, verify that files that were members of a File Replication service (FRS) replica set and certificates that were issued by the Certificate Services are present.

# Performing an Authoritative Restore of Active Directory Objects

Updated: April 10, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

On the domain controller that is being restored, an authoritative restore process returns a designated object or container of objects to its state at the time of the backup. For example, you might need to perform an authoritative restore if an administrator inadvertently deletes an organizational unit (OU) containing a large number of users. If you restore the server from backup, the normal, nonauthoritative restore process does not restore the inadvertently deleted OU because the restored domain controller is updated following the restore process to the current status of its replication partners, which have deleted the OU. Recovering the deleted OU requires authoritative restore. You can use authoritative restore to mark the OU as authoritative and let the replication process restore it to all the other domain controllers in the domain.

| **Note** |
| --- |
| If you can isolate a domain controller in the domain that has not received replication of the deletion, the preliminary, nonauthoritative restore from backup is not necessary. For more information, see Recovering Deletions Without Restoring from Backup later in this topic. |

You can restore objects in domain directory partitions, application directory partitions, and the configuration directory partition, as follows:

- Domain directory partitions: You must restore the objects on a domain controller in the domain.
- Application directory partitions: You must restore the objects on a domain controller that hosts the application directory partition. If you delete an entire application directory partition, you must restore the domain naming operations master to recover the application directory partition.
- Configuration directory partitions: You can restore objects on any domain controller in the forest.

When an object is marked for authoritative restore, its version number is changed so that it is higher than the existing version number of the (deleted) object in the Active Directory replication system. This change ensures that any data that you restore authoritatively is replicated from the restored domain controller to other domain controllers in the forest.

An authoritative restore should not be used to restore an entire domain controller, nor should it be used as part of a change-control infrastructure. Proper delegation of administration and change enforcement will optimize data consistency, integrity, and security.

It is important to ensure successful recovery of the information that is being restored. Group membership is particularly sensitive and can be affected greatly by the procedures that you perform during an authoritative restore.

**Important Preliminary Download Requirement for Windows Server 2003**
Before you perform an authoritative restore on a domain controller that is running Windows Server 2003, you must install an update that fixes problems that can cause unexpected results that follow authoritative restore procedures. The required update is available in Knowledge Base article ID 943576 (http://go.microsoft.com/fwlink/?LinkID=123945). Review the information and use the links in the article to install the hotfix. For more information, see Known Issues for Authoritative Restore.

**Selecting Objects to Restore**
When you select objects that you want to replicate authoritatively, it is important to select the object that is lowest in the directory subtree as possible that you can still use to recover the deleted objects. In this way, you avoid reverting objects that are not related to the deletion. Objects other than the deleted objects might have been modified after the backup was created.

When you restore an OU, any changes that are made up to the time that a backup is restored are rolled back to their values at the time of the backup. For any user accounts, computer accounts, and security groups in the restored OU that were not among the deletions being restored, this rollback might mean the loss of the most recent changes to passwords, home directory, profile path, location and container information, group membership, and any security descriptors that are defined on those objects and attributes.

For example, if an object with a password, such as a user or computer or trust account, is authoritatively restored, the password

value of the restored object reverts to the password value at the time of the backup. In this case, user, computer, and service accounts that have a record of only the current password cannot log on because they have no record of the password that existed when the backup was created. In this way, group membership or other data can also be lost. Updates to the password are blocked because the restored value is authoritative during replication. To minimize the impact of rolling unrelated objects back in time, target as few objects as possible. If you have relatively few deletions to restore, you might be able to restore each object individually.

If you have a relatively large number of deleted objects to restore, use the container object that contains most of the deleted objects. Ideally, the container that you restore will contain all the objects that you need to recover.

**Group Membership Restoration Following Authoritative Restore**

When a user object is inadvertently deleted, you can restore it by marking the user object as authoritative during an authoritative restore procedure. However, depending on the functional level of the forest at the time that any groups to which the user belongs were created (or the forest functional level at the time that the user was added to the group, if they are different), the user's group memberships might not be restored in the process. Multiply this condition by hundreds or thousands of users when an OU is deleted. In this case, additional steps are required to restore the group memberships of user accounts that are restored.

**LVR and Restoration of Group Memberships**

Restoration of group memberships for user objects that are deleted and restored authoritatively differs, depending on whether the group was created (or its membership was updated) before or after the implementation of Windows Server 2003 functionality called linked-value replication (LVR). LVR is a feature that is available when the forest has a functional level of Windows Server 2003 interim or Windows Server 2003. In groups that are created before LVR is in effect, the member attribute is replicated as a single value. Therefore, any change to the group's membership results in replication of the entire member attribute. In groups that are created after LVR is in effect, or in groups that are created before LVR but that are updated after LVR is in effect, updates to the member attribute are replicated separately. In this case, group memberships are restored when you use Ntdsutil to authoritatively restore a user, group, or computer object.

The **memberOf** attribute (or any back-link attribute) is generated only because of its link to the **member** attribute (or any corresponding forward-link attribute). For this reason, restoring the membership on a user object necessarily involves updating the **member** attribute on the group object to include the distinguished name of the restored user.

| **Note** |
| --- |
| Only the forward-link attribute value can be updated and replicated. The back-link attribute value is generated only when it is accessed. It is not stored on the object, and it is not replicated. |

When you use the Ntdsutil command-line tool to authoritatively restore a subtree or single object, the ability of Ntdsutil to restore the group memberships of an object that is authoritatively restored depends on whether the group was created before or after LVR was implemented. For example, if a user object is restored and the user belongs to group G1 that was created before LVR was implemented and the user belongs to group G2 that was created after LVR was implemented (the functional level of the forest was raised to Windows Server 2003 interim or Windows Server 2003), the **member** attribute of G2 is updated during authoritative restore (and, therefore, the **memberOf** attribute of the restored user is updated), but the **member** attribute of G1 is not updated.

However, improvements in the version of Ntdsutil that was introduced in Windows Server 2003 with Service Pack 1 (SP1) provide the ability to also restore the memberships of groups that were created before LVR was implemented.

**Authoritative Restore Improvements Introduced in Windows Server 2003 SP1**

Beginning with Windows Server 2003 with SP1, during an authoritative restore, Ntdsutil creates a text file that identifies the authoritatively restored objects and uses this file to create an LDAP Data Interchange Format (LDIF) file. The LDIF file can be used to regenerate all back-links on the restored objects in a forest where LVR is not in effect.

If you need to restore a large number of users (for example, if you delete an OU) in domain X and your forest also contains domain Y and domain Z, authoritative restore requires the restoration of domain X and then the use of Ntdsutil to generate and run the LDIF file against a domain controller in each additional domain.

In most cases, you begin the authoritative restore process by performing a nonauthoritative restore from backup. If you can identify a global catalog server that has not received replication of the deletions and you can isolate it before replication occurs, you can perform the authoritative restore without first restoring from backup. For information about using a latent global catalog server for authoritative restore, see Recovering deletions without restoring from backup. In either case, you then perform the additional steps to restore group memberships, if necessary. The steps that you perform are different if you are restoring the objects on a domain controller running a version of Windows Server beginning with Windows Server 2003 with SP1.

**Files for Recovering Group Memberships Following Authoritative Restore**
When you perform authoritative restore on a domain controller that is running Windows Server 2003 with SP1, Ntdsutil creates the following files that you can use to recover group memberships:

- ar_*YYYYMMDD-HHMMSS*_links_*Domain*.ldf, which is an LDIF file that is generated for the domain in which you perform the authoritative restore procedure. This file contains back-link information for the restored objects. If you perform the procedure on a global catalog server, a separate .ldf file is created for each domain in the forest. You can use this file with the Ldifde command-line tool to import the back-links to recover universal and global group memberships in environments that include pre-LVR groups. For environments that do not include pre-LVR groups, the Ntdsutil tool recovers group memberships automatically in the recovery domain and in the forest (for universal groups) if the recovery domain controller is a global catalog server. If the restore includes security principals that can have memberships in domain local groups in other domains, a text file that is also generated during authoritative restore is required to restore the memberships in the additional domains.

- ar_*YYYYMMDD-HHMMSS*_objects.txt, which is a text file that contains a list of the authoritatively restored objects. This file is generated for each individual object or container that you mark as authoritative. You can use this file to generate an .ldf file that you can use to recover memberships in domain local groups and universal groups (if you are not restoring a global catalog server) in other domains. This file is created on any domain controller that you authoritatively restore. Global catalog servers do not store the **member** attribute of domain local groups. Therefore, even if you perform the restore on a global catalog server, you must always use this file to generate an .ldf file in any domain where there are domain local groups of which restored security principals might be members. You must create a separate .ldf file for each object or container that you mark as authoritative.

**Using a Global Catalog Server for Authoritative Restore**
If possible, perform the authoritative restore on a global catalog server in the domain where the objects were deleted to recover security principals and group memberships. Global catalog servers store a single, writable domain and a partial, read-only replica of all other domains in the forest. A partial replica means that the global catalog stores all objects, but with a limited set of attributes on each object. Check the properties of the NTDS Settings object of the server object in Active Directory Sites and Services to confirm that a domain controller is a global catalog server.

**Global Catalog and Group Memberships**
In relation to the three types of security groups—global groups, domain local groups, and universal groups—global catalog servers are best suited for recovering group memberships after an authoritative restore procedure because they store memberships of all universal groups in the forest and all global groups in the domain.

Security group memberships are restored on a global catalog server as follows:

- **Global groups**: Security principals (users, groups, and computers) can be members of only the global groups that are created in the same domain. Global catalog servers store a writable domain directory partition. Therefore, they can restore global group memberships for the recovery domain.

- **Universal groups:** Security principals can be members of universal groups that are created in any domain. However, the **member** attribute is among the attributes that are stored on the read-only universal group objects in the global catalog. Therefore, a global catalog server can recover universal group memberships for all domains in the forest. A domain controller that is not a global catalog server stores only universal group objects that are created in its own domain.

- **Domain local groups**: Security principals can be members of domain local groups that are created in any domain. Memberships in domain local groups in the recovery domain are restored automatically during authoritative restore. However, the global catalog does not store the **member** attribute for read-only domain local group objects. Therefore, for restored security principals that have memberships in domain local groups in other domains, you must recover these memberships by performing follow-up procedures in each additional domain. The follow-up procedures are described in Create an LDIF file for recovering back-links for authoritatively restored objects.

**Recovering Deletions Without Restoring from Backup**
If you can isolate a global catalog server (or any domain controller, but preferably a global catalog server) in the domain where the deletion occurred before the server receives replication of the deletion, you might be able to avoid performing a preliminary restore from backup (nonauthoritative restore) and having to extend the restore process to other domains. Use the **repadmin /showrepl** command to determine the date and time of the latest inbound replication of the domain directory partition where the deletions occurred.

Global catalog servers often have greater replication latency than ordinary domain controllers, and they are better restore candidates in general because they store universal group memberships. If you can stop inbound replication on a latent global catalog server, you can perform authoritative restore on the global catalog server to recover the deleted memberships for all groups in the domain and for all universal groups in other domains.

If you want to use a latent global catalog server for restoring deleted objects, you must take steps to stop inbound replication immediately. You can use one of the following methods to stop replication:

- Use the Services snap-in to stop AD DS. In this case, other services continue to operate.
- Take the global catalog service offline by restarting it in Directory Services Restore Mode (DSRM). In this case, all other directory-related services are stopped in addition to AD DS.
- Use Repadmin.exe to stop inbound replication. In this case, the domain controller continues to operate but does not receive replication updates.

**Retention (Merge) of New Group Memberships or Other Attributes Following Authoritative Restore**

The authoritative restore procedure results in a merge of authoritatively restored objects and attributes and existing objects and attributes. For example, do not expect that users that have been added to a group (after the backup that is used to restore the deleted group) will be removed by an authoritative restore of the group object. Instead, new attributes of objects that are specified in the authoritative restore are preserved during replication. Therefore, authoritative restore does not remove group memberships that were added between the time of the backup that is used for authoritative restore and the time of the restore procedure.

Objects and attributes are preserved during authoritative restore as follows:

- If an object exists in the backup, before inbound replication occurs, the post-restore directory partition contains the version of the object that exists in the restored backup.
- If an object was created after the backup was made and there are additional domain controllers that store the directory partition, after inbound replication occurs, the restored directory partition also includes the set of objects that were created after the backup.
- If an object contains new attributes that are not contained in the backup but that exist in the directory partition of an additional domain controller in the domain at the time of the restore, after inbound replication the version of the object and attributes as they existed in the backup—plus any new attributes that were added to the object after the backup—are preserved.

Authoritative restore affects only the objects and attributes that existed at the time of the backup. This functionality applies to objects with linked attributes and nonlinked attributes alike. For example, if you are restoring an object that has attribute A and attribute B in the backup version and has attributes A′, B′, and C in the current directory, attribute C is retained after authoritative restore. Therefore, a group object that has the **member** value of User1 in the backup and has both User1 and User2 in the current directory includes both of those memberships after authoritative restore of the group object. Any post-backup **memberOf** or **member** attribute values that were added to a user or group, respectively, are not affected by replication updates after the restore procedure.

If you want to remove group memberships—or any other unwanted object attribute—complete the following steps:

1. Delete the object whose updates you do not want to retain.
2. Allow the deletion to replicate throughout the forest.
3. Back up a domain controller that has received the deletion.
4. Authoritatively restore the object that you deleted from the backup that does not contain the unwanted values.

**Procedures for Domain Controllers Running Windows Server 2003 with SP1, Windows Server 2003 with SP2, or Windows Server 2003 R2**

These procedures include the use of an LDIF file to restore group memberships following authoritative restore of the objects. If

you are restoring objects that can belong to groups in more than one domain, additional steps are required.

**Task requirements**

The following tools are required to perform the procedures for this task:

- Ntbackup.exe
- Ntdsutil.exe
- Repadmin.exe

To complete this task, perform procedures according to the conditions in your environment:

- Procedures for restoring after deletions have replicated
- Procedures for restoring before deletions have replicated
- Procedures for recovering group memberships (and any other back-link attributes) in other domains

**Procedures for Restoring After Deletions Have Replicated**
If you are performing authoritative restore on a domain controller that has already received replication of the deletions, perform the following procedures on the recovery domain controller:

1. If you do not have a current backup of the recovery domain controller, Back up system state. You can use this backup if your recovery is not successful and you can try again.

2. Restart the domain controller in Directory Services Restore Mode locally

   Or

   Restart the domain controller in Directory Services Restore Mode Remotely

3. Restore Active Directory from backup. Be sure to use a backup that contains the deleted objects.

   Restore system state backup to return the domain controller to its state at the time of the backup so that any groups that are being restored, or whose members are being restored, are present in the directory with their pre-deletion membership intact.

   To ensure that replication does not occur, click **No** at the end of the procedure so that the domain controller does not restart.

4. Mark the object or objects authoritative

   Mark the object or objects that you want to restore so that replication does not overwrite them when you restart the domain controller.

5. Restart the domain controller normally.

6. Synchronize replication with all partners

   For the newly restored object to become available and be instantiated in its restored form on all domain controllers, successful replication must occur between the domain controller that originates the restored changes and its partners.

   Make sure that all domain controllers in the domain and all global catalog servers in the forest have received the restored objects.

7. Use the following procedure to run the LDIF file that was created in step 2 on this domain controller to add the missing

group memberships in the domain that you have just restored:

Run an LDIF file to recover back-links in this domain. This procedure updates the group memberships of a restored security principal object or container of objects in the recovery domain. Perform this procedure for each individual object or container that you marked as authoritative.

8. Back up system state on the recovered domain controller.

9. If the .ldf file shows back-links for objects in other domains, perform the procedures in Procedures for recovering group memberships (and any other back-link attributes) in other domains.

**Procedures for Restoring Before Deletions Have Replicated**

If you have identified a global catalog server or other domain controller that has not received replication of the deletions and for which you have a recent backup, you do not have to perform a preliminary restore from backup. Perform the following procedures on the recovery domain controller:

1. Turn off inbound replication.

2. If you do not have a current backup of the recovery domain controller, Back up system state. You can use this backup if your recovery is not successful and you can try again.

3. Restart the domain controller in Directory Services Restore Mode locally.

   Or

   Restart the domain controller in Directory Services Restore Mode Remotely.

4. Mark the object or objects authoritative.

5. Restart the domain controller normally.

6. Synchronize replication with all partners.

7. Run an LDIF file to recover back-links in this domain.

8. Turn on inbound replication.

9. Back up system state on the recovered domain controller.

10. If the .ldf file shows back-links for objects in other domains, perform the procedures in Procedures for recovering group memberships (and any other back-link attributes) in other domains.

**Procedures for Recovering Group Memberships (and Any Other Back-Link Attributes) in Other Domains**

You can recover group memberships in other domains either by adding the members manually to the respective groups or by using the text file from the original authoritative restore procedure to generate one or more .ldf files that you can use to recover back-links in other domains. Be aware that restored objects might have back-links other than group memberships.

If you have restored security principal objects or other objects that have back-link attributes in a forest that has more than one domain and you do not want to restore the back-links manually, and if the .ldf file generated during authoritative restore shows back-links for objects in other domains, perform the following steps on a domain controller in each additional domain:

| **Note** |
| --- |
| For restored security principals, these steps are required only if the restored security principals have memberships in domain local or universal groups in a different domain from the recovery domain. If you restored the security principals on a global catalog server, you need to recover only domain local group memberships in other domains. In some cases, these accounts might be few enough that you can manually recreate the memberships instead of following these procedures. |

1. Restart the domain controller in Directory Services Restore Mode locally

2. Restore Active Directory from backup.

   When the group members were deleted, the **member** attribute (forward link) on any group of which they were members was removed from the group object. This procedure is required to restore the **member** attribute on group objects for those group members that were deleted. This attribute is required to regenerate the **memberOf** attribute value on the restored group members.

3. While still in Directory Services Restore Mode, use Ntdsutil to Create an LDIF file for recovering back-links for authoritatively restored objects

   In this procedure, you must specify the location of the .txt file that was generated by Ntdsutil during the authoritative restore procedure.

4. Restart the domain controller normally (not in Directory Services Restore Mode).

5. Run an LDIF file to recover back-links in this domain on a domain controller other than the domain controller that you restored from backup and on which you created the LDIF file. Because you have just restored the domain controller on which you created the LDIF file from backup, perform this procedure on a different domain controller to be sure that the group objects you update are current. This procedure updates the group memberships of a restored security principal object or container of objects. Perform this procedure for each individual object or container that you marked as authoritative.

**Procedures for Domain Controllers Running Windows Server 2003 with No Service Pack Installed**
To complete this task, perform the following procedures in order:

| **Note** |
| --- |
| If the objects that were deleted do not include group objects, you do not have to perform steps 3 through 10. In addition, if the groups that were deleted do not have members among the list of deleted objects, you do not have to perform steps 3 through10. |

1. Restore Active Directory from backup

   Restore system state to return the domain controller to its state at the time of the backup. To ensure that replication does not occur, click **No** at the end of the procedure so that the domain controller does not restart.

2. Mark the object or objects authoritative

   Mark the object or objects that you want to restore so that replication does not overwrite them when you restart the domain controller.

3. Restart the computer normally, but in isolation. This step allows you to control replication so that inbound replication does not update any restored object before forcing outbound replication. You cannot turn off inbound replication in Directory Services Restore Mode.

   The most common way to start a computer in isolation is to remove the network connection from the domain controller by physically removing the network cable. Alternative methods may be possible, depending on your network hardware and enterprise practices.

   It is important to prevent the domain controller from communicating with any other domain controller in the domain or forest. You should also isolate the domain controller from any clients that might change an object in the directory.

4. Turn off inbound replication

   This step is required only if the domain or forest functional level is Windows 2000 native or earlier. By turning off

inbound replication, you ensure that no changes replicate in to the domain controller and alter group membership.

5. Reconnect the computer to the network.

   After you turn off inbound replication, it is safe to reconnect the domain controller to the network.

   If you isolated your computer by removing the network cable or by disconnecting the network connection from the domain controller, reconnect it to bring the domain controller back onto the network.

   If you followed other procedures based on your enterprise network equipment, follow the equipment's recommendations for reconnecting the domain controller to the network.

6. Synchronize replication with all partners

   For the newly restored object to become available and be instantiated in its restored form on all domain controllers, successful replication must occur between the domain controller that originates the restored changes and its partners.

   Make sure that all domain controllers in the domain and all global catalog servers in the forest have received the restored objects.

7. Restart the domain controller in Directory Services Restore Mode locally

8. Mark the object or objects authoritative

   One of the challenges of restoring objects, and their group memberships, is the fact that the membership and object may replicate in different orders. If the membership replicates before a user is restored, the receiving domain controller will not update the membership because the user does not exist. To overcome the effects of this behavior, it is necessary to mark the objects that have been restored as authoritative a second time and once again have the information replicated out.

9. Restart the computer normally (not in Directory Services Restore Mode).

   After the authoritative restore of the object or objects has completed a second time, you can restart the domain controller in normal mode.

10. Turn on inbound replication

# Restore Active Directory from backup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To restore the server, use a good backup containing the system state or the system state and system disk.

| **Note** |
| --- |
| To restore from backup, you must log on locally to the domain controller or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restore Active Directory from backup**

1. Start the computer in Directory Services Restore Mode.
2. To start the Windows Server 2003 backup utility, click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

   This procedure provides steps for restoring from backup in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard.
3. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
4. Click **Restore files and settings**, and then click **Next**.
5. Select **System State**, and then click **Next**.
6. On the **Completing the Backup or Restore Wizard** page, click **Advanced**.
7. In **Restore files to**, click **Original Location**, and then click **Next**.
8. Click **Leave existing files (Recommended)**, and then click **Next**.
9. In **Advanced Restore Options**, select the following check boxes, and then click **Next**:
   - **Restore security settings**
   - **Restore junction points, but not the folders and file data they reference**
   - **Preserve existing volume mount points**
10. For a primary restore of SYSVOL, also select the following check box: **When restoring replicated data sets, mark the restored data as the primary data for all replicas**.

    A primary restore is required only if the domain controller that you are restoring is the only domain controller in the domain. A primary restore is required on the first domain controller that is being restored in a domain if you are restoring the entire domain or forest.
11. Click **Finish**.
12. When the restore process is complete, click **Close**, and then do one of the following:
    - If you do not want to authoritatively restore any objects, click **Yes** to restart the computer. The system will restart and replicate any new information that is received since the last backup with its replication partners.
    - If you want to authoritatively restore any objects or if you want to create an LDAP Data Interchange Format (LDIF) file to restore back-links on this domain controller, click **No** to remain in Directory Services Restore Mode. For information about how to proceed with authoritative restore, see Performing an Authoritative Restore of Active

Directory Objects.

**See Also**
**Concepts**
Restart the domain controller in Directory Services Restore Mode locally
Enable Remote Desktop
Create a Remote Desktop Connection
Restart the domain controller in Directory Services Restore Mode Remotely
Restore system state to an alternate location
Performing an Authoritative Restore of Active Directory Objects

# Mark the object or objects authoritative

Updated: May 28, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In this procedure, you select which objects are to be marked authoritative to have them replicate to other domain controllers. You must have completed a nonauthoritative restore procedure, following which the domain controller has not been restarted and remains in Directory Services Restore Mode. To complete this procedure, you must know the full distinguished name of the object or objects that you want to restore.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To mark a subtree or individual object authoritative**

1. In Directory Services Restore Mode, click **Start**, click **Run**, type **ntdsutil**, and then press ENTER.
2. At the **ntdsutil:** prompt, type **authoritative restore**, and then press ENTER.
3. To restore a subtree or individual object, type one of the following commands, as appropriate, and then press ENTER:

   To restore a subtree (for example, an organizational unit and all child objects):

   **restore subtree** *DistinguishedName*

   To restore a single object:

   **restore object** *DistinguishedName*

   ***DistinguishedName***
   The distinguished name of the subtree or object that is to be marked authoritative

4. Click **Yes** in the message box to confirm the command.

   For example, if you want to restore a deleted organizational unit named Marketing NorthAm in the corp.contoso.com domain, type:

   **restore subtree "OU=Marketing NorthAm,DC=corp,DC=contoso,DC=com"**

   (Always enclose the distinguished name in quotes when there is a space or other special characters within the distinguished name.)

   Ntdsutil attempts to mark the object as authoritative. The output message indicates the status of the operation. The most common cause of failure is an incorrectly specified distinguished name or a backup for which the distinguished name does not exist (which occurs if you try to restore a deleted user that was created after the backup).

   | **Note** |
   | --- |
   | If the object is marked authoritative and the restore did not work as expected and then the same backup is used to restore the Active Directory database again, the attribute version numbers of the object to be restored authoritatively must be increased higher than the default of 100000 or the object will not replicate out after the second restore. The syntax below is needed to script an increased version number higher than 100000 (default): ntdsutil "authoritative restore" "restore object \"CN=Smith\, John,OU=sales,DC=contoso,DC=com\" verinc 150000\"" q q<br><br>If the script prompts for confirmation on each object being restored you can turn off the prompts. The syntax to turn off the prompts is: ntdsutil "popups off" "authoritative restore" "restore object \"CN=John Smith,OU=sales,DC=contoso,DC=com\" verinc 150000\"" q q |

   If you are running this command on a domain controller running Windows Server 2003 with Service Pack 1 (SP1), Ntdsutil provides output that indicates whether a restored object has back-links that must be restored. If objects that have back-links are found, Ntdsutil generates a set of files that you can use to restore the back-links in this domain and in other domains, if necessary.

   The following sample output on a domain controller running Windows Server 2003 with SP1 shows that Ntdsutil created a text file (.txt) and an LDAP Data Interchange Format (LDIF) file (.ldf) when the marked object was found to have back-

links:

```
Successfully updated 3 records.

The following text file with a list of authoritatively restored
objects has been created in the current working directory:
ar_20050209-091249_objects.txt

One or more specified objects have back-links in this domain. The
following LDIF files with link restore operations have been created
in the current working directory:
        ar_20050209-091249_links_Test1.com.ldf

Authoritative Restore completed successfully.
```

5. Make a note of the location of the .txt and .ldf files, if any. You will use the .ldf file to restore back-links in this domain. You will use the .txt file to generate an LDIF file to restore back-links in a different domain, if necessary. If you have other domains in which you want to restore back-links for this restored object, make a copy of this .txt file to use on a domain controller in another domain. For more information, see Procedures for Recovering Group Memberships (and Any Other Back-Link Attributes) in Other Domains.

   While still in Directory Services Restore Mode, use Ntdsutil to Create an LDIF file for recovering back-links for authoritatively restored objects. In this procedure, you must specify the location of the .txt file that was generated by Ntdsutil during the authoritative restore procedure.

6. At the **authoritative restore:** and **ntdsutil:** prompts, type **quit**, and then press ENTER.

7. Restart the domain controller in normal operating mode, as follows:

   a. For a domain controller running Windows Server 2003 with no service pack installed: Disconnect the domain controller from the network, and then restart normally. Follow the instructions in "Procedures for Domain Controllers Running Windows Server 2003 with No Service Pack Installed" as described in Performing an Authoritative Restore of Active Directory Objects.

   b. For a domain controller running Windows Server 2003 with SP1: Restart the domain controller normally, and then follow the instructions in "Procedures for Domain Controllers Running Windows Server 2003 with SP1" as described in Performing an Authoritative Restore of Active Directory Objects.

# Synchronize replication with all partners

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to synchronize replication with all replication partners of a domain controller.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the selected domain controller or the Enterprise Admins group in the forest, or you must have been delegated the appropriate authority. If you want to synchronize the configuration and schema directory partitions on a domain controller in a child domain, you must have Domain Admins credentials in the forest root domain or Enterprise Admins credentials in the forest.

**To synchronize replication with all partners**

1. At a command prompt, type the following command, and then press ENTER:

   **repadmin /syncall** *DCName* **/e /d /A /P /q**

| Term | Definition |
|---|---|
| *DCName* | The Domain Name System (DNS) name of the domain controller on which you want synchronize replication with all partners |
| **/e** | Enterprise; includes partners in all sites. |
| **/d** | Identifies servers by distinguished name in messages. |
| **/A** | All; synchronizes all directory partitions that are held on the home server. |
| **/P** | Pushes changes outward from the home server. |
| **/q** | Runs in quiet mode; suppresses callback messages. |

2. Check for replication errors in the output of the command in the previous step. If there are no errors, replication is successful. For replication to complete, any errors must be corrected.

**See Also**
**Concepts**
Verify successful replication to a domain controller

# Run an LDIF file to recover back-links

Updated: April 10, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Ntdsutil in Windows Server 2003 Service Pack 1 (SP1) provides new functionality for performing authoritative restore of objects that have back-links.

The output of the authoritative restore procedure includes the name of an LDAP Data Interchange Format (LDIF) (.ldf) file that contains the forward-links that are required so that the group memberships (back-links) of any restored user, group, or computer objects can be recovered. For each object or subtree that you restore, you must run the LDIF file on a domain controller in each domain that might have group objects that are required to recover back-links on the restored objects.

| **Note** |
| --- |
| This procedure is critical for recovering group memberships for deleted users, groups, or computers, but it applies to any restored objects that have back-link attributes. |

### Administrative credentials

To perform this procedure, you must be a member of the Domain Admins group in the domain of the domain controller on which you run the command.

**To run an LDIF file to recover back-links following authoritative restore**

1. Open a command prompt and change directories, if necessary, to the directory of the .ldf file and its respective log files.
2. At the command prompt, type the following command, and then press ENTER:

   **ldifde -i -k -f** *FileName*

   ***FileName***
   The name of the .ldf file that you want to run, for example, ar_20050609-174604_links_corp.contoso.com.ldf

**See Also**
**Concepts**
Create an LDIF file for recovering back-links for authoritatively restored objects

# Restart the domain controller in Directory Services Restore Mode locally

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have physical access to a domain controller, you can restart the domain controller in Directory Services Restore Mode locally. Restarting in Directory Services Restore Mode takes the domain controller offline. In this mode, the server is not functioning as a domain controller.

When you start Windows Server 2003 in Directory Services Restore Mode, the local Administrator account is authenticated by the local Security Accounts Manager (SAM) database. Therefore, logging on requires that you use the local administrator password, not an Active Directory domain password. This password is set during Active Directory installation when you provide the password for Directory Services Restore Mode.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restart the domain controller in Directory Services Restore Mode locally**

1. Restart the domain controller.
2. When the screen for selecting an operating system appears, press F8.
3. On the **Windows Advanced Options** menu, select **Directory Services Restore Mode**.
4. When you are prompted, log on as the local administrator.

**See Also**
**Concepts**
Restart the domain controller in Directory Services Restore Mode Remotely

# Create an LDIF file for recovering back-links for authoritatively restored objects

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have authoritatively restored objects that have back-links in another domain, you can use this procedure to create an LDAP Data Interchange Format (LDIF) file that you can run against a domain controller in that domain to restore the back-links. Perform this procedure on a domain controller in the domain that has the back-links.

Before you perform this procedure, you must:

- Copy the .txt file that Ntdsutil created during the authoritative restore procedure, which you performed on the first domain controller, to a location on this domain controller or a network share.
- Restore this domain controller from backup media.

After you restore this domain controller from backup media, perform this procedure while the domain controller is still running in Directory Services Restore Mode.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To create an LDIF file for restoring back-links for authoritatively restored objects**

1. In Directory Services Restore Mode, click **Start**, click **Run**, type **ntdsutil**, and then press ENTER.
2. At the **ntdsutil: prompt**, type **authoritative restore**, and then press ENTER.
3. At the **authoritative restore:** prompt, type the following command, and then press ENTER:

   **create ldif files from** *TextFilePath*

   where *TextFilePath* is the location and file name of the .txt file that Ntdsutil created during the initial authoritative restore of the object whose back-links you want to restore, for example, d:\ldif\ar_20050609_091558_objects.txt.

   Ntdsutil displays a message stating that one or more specified objects have back-links in this domain and an LDIF file has been created in the current working directory.
4. At the **authoritative restore:** and **ntdsutil:** prompts, type **quit**.

**See Also**
**Concepts**
Restore Active Directory from backup
Run an LDIF file to recover back-links

# Turn off inbound replication

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to turn off inbound replication so that objects on a domain controller cannot be updated by replication from another domain controller.

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group in the domain of the domain controller whose replication you are disabling, or you must be a member of the Enterprise Admins group.

**To turn off inbound replication**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /options** *ServerName* **+DISABLE_INBOUND_REPL**

   where *ServerName* is the network basic input/output system (NetBIOS) name of the domain controller.
3. Verify that the option is set. The following message should appear:

   **New DC Options: DISABLE_INBOUND_REPL**

**See Also**
**Concepts**
Turn on inbound replication

# Turn on inbound replication

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to turn on inbound replication after it has been turned off manually.

**Administrative credentials**

To complete this procedure, you must be a member of the Domain Admins group in the domain of the domain controller whose replication you are enabling, or you must be a member of the Enterprise Admins group.

**To turn on inbound replication**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /options** *ServerName* **-DISABLE_INBOUND_REPL**

   where *ServerName* is the network basic input/output system (NetBIOS) name of the domain controller.
3. Verify that the option is set. The following message should appear:

   **Current DC options: DISABLE_INBOUND_REPL**

   **New DC Options: <none>**

   Current DC Options displays the conditions that were in effect at the time that you ran the command. New DC Options shows the effect of the command, which is that the option to disable replication is not set.

**See Also**
**Concepts**
Turn off inbound replication

# Performing an Authoritative Restore of an Application Directory Partition

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restoration of an application partition will mark all data that is present in the application partition as authoritative for the replica set. Information that is contained within an application partition will replicate to all domain controllers in the forest that were previously present in the replica set. You should have a current valid backup of the application partition prior to restoring, in the event that particular object changes are lost because of changes since backup.

**Task Requirements**

The following tools are required to perform the procedures for this task:

- Backup.exe
- Ntdsutil.exe

To complete this task, perform the following procedures:

1. Restore Active Directory from backup
2. Mark the application directory partition as authoritative
3. Restart the computer

   Once the authoritative restore of the object or objects has been completed a second time, the domain controller can be restarted in normal mode.

# Restore Active Directory from backup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To restore the server, use a good backup containing the system state or the system state and system disk.

| **Note** |
| --- |
| To restore from backup, you must log on locally to the domain controller or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restore Active Directory from backup**

1. Start the computer in Directory Services Restore Mode.
2. To start the Windows Server 2003 backup utility, click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

   This procedure provides steps for restoring from backup in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard.
3. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
4. Click **Restore files and settings**, and then click **Next**.
5. Select **System State**, and then click **Next**.
6. On the **Completing the Backup or Restore Wizard** page, click **Advanced**.
7. In **Restore files to**, click **Original Location**, and then click **Next**.
8. Click **Leave existing files (Recommended)**, and then click **Next**.
9. In **Advanced Restore Options**, select the following check boxes, and then click **Next**:
   - **Restore security settings**
   - **Restore junction points, but not the folders and file data they reference**
   - **Preserve existing volume mount points**
10. For a primary restore of SYSVOL, also select the following check box: **When restoring replicated data sets, mark the restored data as the primary data for all replicas**.

    A primary restore is required only if the domain controller that you are restoring is the only domain controller in the domain. A primary restore is required on the first domain controller that is being restored in a domain if you are restoring the entire domain or forest.
11. Click **Finish**.
12. When the restore process is complete, click **Close**, and then do one of the following:
    - If you do not want to authoritatively restore any objects, click **Yes** to restart the computer. The system will restart and replicate any new information that is received since the last backup with its replication partners.
    - If you want to authoritatively restore any objects or if you want to create an LDAP Data Interchange Format (LDIF) file to restore back-links on this domain controller, click **No** to remain in Directory Services Restore Mode. For information about how to proceed with authoritative restore, see Performing an Authoritative Restore of Active

Directory Objects.

**See Also**
**Concepts**
Restart the domain controller in Directory Services Restore Mode locally
Enable Remote Desktop
Create a Remote Desktop Connection
Restart the domain controller in Directory Services Restore Mode Remotely
Restore system state to an alternate location
Performing an Authoritative Restore of Active Directory Objects

# Mark the application directory partition as authoritative

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can select which application directory partitions are to be marked authoritative so that you can have them replicated to other domain controllers. To perform this procedure, you must have completed a nonauthoritative restore procedure. After that procedure is complete, the domain controller is not restarted, and it remains in Directory Services Restore Mode.

**Administrative credentials**

To complete this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To mark an application directory partition as authoritative**

1. In Directory Services Restore Mode, open a Command Prompt.
2. Type the following command, and then press ENTER:

   **ntdsutil**

3. At the **ntdsutil:** prompt, type **authoritative restore**, and then press ENTER. For assistance with the Ntdsutil command line-tool, type **help** at any time.
4. Type **List NC CRs**, and then press ENTER.

   Ntdsutil displays a list of the application directory partitions that are available after the restore, along with the associated cross-references. Note the cross-reference distinguished name and application directory partition distinguished name that correspond to the application directory partition that you want to restore.

5. Type **restore subtree** *App Partition DN*, where *App Partition DN* is the distinguished name of the application directory partition that you want to restore.
6. In the confirmation dialog box, click **Yes**.

   The output message indicates the status of the operation. There should be no failures.

7. Type **restore object** *Cross Ref DN* (where *Cross Ref DN* is the distinguished name of the application directory partition cross-reference that you want to restore), and then press ENTER.
8. In the confirmation dialog box, click **Yes**.

   The output message indicates the status of the operation. There should be no failures.

9. Quit the Ntdsutil tool.

# Performing an Authoritative Restore of a Group Policy Object

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restoring a Group Policy Object (GPO) restores the GPO to a previous state. A restore operation can be used in both of the following cases: the GPO was backed up but has since been deleted, or the GPO is live and you want to roll back to a known previous state. A restore operation retains the original GPO GUID even if the restore is recreating a deleted GPO. This is a key difference between the restore operation and the import or copy operations discussed in later sections of this guide.

A restore operation replaces the following components of a GPO:

- GPO settings
- ACLs on the GPO
- WMI filter links (but not the filters themselves)

The restore operation does not restore links to a SOM (Scope of Management). Any existing links will continue to be used—for example, when restoring an existing GPO to a previous state. However, if the user has deleted a GPO and all links to the GPO, the user must recreate these links after restoring the GPO. To facilitate recreating these links, you can view the report in the backup to identify all links in the domain of the GPO.

For more information, see Administering Group Policy with the GPMC on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=17528).

**Task Requirements**

The following tool is required to perform the procedures for this task:

- Group Policy Management Console

To complete this task, perform the following procedure:

- Restore a Group Policy Object

# Restore a Group Policy Object

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to perform an authoritative restore of a Group Policy object.

**Administrative credentials**

To perform this procedure, you must have edit, delete , and modify permissions on the specific Group Policy object.

**To restore a Group Policy object**

1. Open Group Policy Management Console (GPMC).
2. In the console tree, double-click **Domains** to expand the list of domains.
3. Double-click the desired domain to expand the contents of that domain.
4. Right-click **Group Policy Objects**, and select **Manage Backups**.
5. Right-click the object to be restored, and select **Restore from Backup**.
6. Select the backup location, click the policy backup to be restored, and then click **Restore**.
7. Click **OK** to restore the selected GPO backup.

# Restoring a Domain Controller Through Reinstallation and Subsequent Restore from Backup

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you cannot restart a domain controller in Directory Services Restore Mode, you can restore it through reinstallation of the operating system and subsequent restore of Active Directory from backup.

After you reinstall Windows Server 2003, perform a nonauthoritative restore of the system state and the system disk. You do not need to join the computer to the domain before performing the restore procedure. During the restore, the computer account is reestablished automatically.

| **Note** |
| --- |
| The restore procedure must be performed by using the same backup tool with which the backup was made. Procedures in this task describe using Ntbackup to restore Active Directory, but you must use the tool that you used to create the backup file if it is not Ntbackup. |

Restore a domain controller through reinstallation and restore the system state from backup if the following conditions exist:

- A domain controller has failed and you cannot restart in Directory Services Restore Mode. If the failure was caused by a hardware failure, you have resolved the hardware problem (for example, by replacing the disk).
- You have a previous backup for the failed domain controller that is not older than the tombstone lifetime for the forest.
- The domain controller is running other server services such as Exchange, or it contains other data that you must restore from a backup.
- You have the following information about the failed domain controller:
    - Disk configuration. You need a record of the volumes and sizes of the disks and partitions. In the case of a complete disk failure, use this information to recreate the disk configuration. Windows Server 2003 must be reinstalled to the same drive letter and with at least the same amount of physical drive space. Before you restore the system state, you must recreate all disk configurations. Failure to recreate all disk configurations can cause the restore process to fail and can prevent you from starting the domain controller after the restore.
    - Computer name. You need the computer name to restore a domain controller of the same name and avoid changing client configuration settings.
    - Password for the local computer Administrator account. You must know the local Administrator password that was used when the backup was created. The local Administrator password is also required to restore the system state on a domain controller.

**Task requirements**

The following tool is required to perform the procedures for this task:

- Ntbackup.exe

To complete this task, perform the following procedures:

1. Install Windows Server 2003.

| **Note** |
| --- |
| This guide does not provide information for the installation of Windows Server 2003. |

2. Restore Active Directory from backup

    a. Begin with step 2 of this procedure. You cannot start the server in Directory Services Restore Mode because Active Directory is not installed.

    b. This operation requires that you log on as the local Administrator, not the Administrator for Directory Services Restore Mode.

    c. Restore the System State as described, but in normal mode.

    d. When you are prompted to restart the server after you complete the restore operation, click **Yes** to restart the server normally.

3. Verify Active Directory restore

# Restore Active Directory from backup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To restore the server, use a good backup containing the system state or the system state and system disk.

| **Note** |
| --- |
| To restore from backup, you must log on locally to the domain controller or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restore Active Directory from backup**

1. Start the computer in Directory Services Restore Mode.
2. To start the Windows Server 2003 backup utility, click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

   This procedure provides steps for restoring from backup in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard.
3. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
4. Click **Restore files and settings**, and then click **Next**.
5. Select **System State**, and then click **Next**.
6. On the **Completing the Backup or Restore Wizard** page, click **Advanced**.
7. In **Restore files to**, click **Original Location**, and then click **Next**.
8. Click **Leave existing files (Recommended)**, and then click **Next**.
9. In **Advanced Restore Options**, select the following check boxes, and then click **Next**:
   - **Restore security settings**
   - **Restore junction points, but not the folders and file data they reference**
   - **Preserve existing volume mount points**
10. For a primary restore of SYSVOL, also select the following check box: **When restoring replicated data sets, mark the restored data as the primary data for all replicas**.

    A primary restore is required only if the domain controller that you are restoring is the only domain controller in the domain. A primary restore is required on the first domain controller that is being restored in a domain if you are restoring the entire domain or forest.
11. Click **Finish**.
12. When the restore process is complete, click **Close**, and then do one of the following:
    - If you do not want to authoritatively restore any objects, click **Yes** to restart the computer. The system will restart and replicate any new information that is received since the last backup with its replication partners.
    - If you want to authoritatively restore any objects or if you want to create an LDAP Data Interchange Format (LDIF) file to restore back-links on this domain controller, click **No** to remain in Directory Services Restore Mode. For information about how to proceed with authoritative restore, see Performing an Authoritative Restore of Active

Directory Objects.

**See Also**
**Concepts**
Restart the domain controller in Directory Services Restore Mode locally
Enable Remote Desktop
Create a Remote Desktop Connection
Restart the domain controller in Directory Services Restore Mode Remotely
Restore system state to an alternate location
Performing an Authoritative Restore of Active Directory Objects

# Verify Active Directory restore

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After the restore is completed, use this procedure to restart the server and perform basic verification.

**Administrative credentials**

To verify Active Directory restore, you must be a member of the Domain Admins group.

**To verify Active Directory restore**

1. After the restore operation completes, restart the computer in **Start Windows Normally** mode. Active Directory and Certificate Services automatically detect that they have been recovered from a backup. They perform an integrity check and re-index the database.
2. After you are able to log on to the system, browse Active Directory. Verify that all of the User objects and Group objects that were present in the directory prior to backup are restored. Similarly, verify that files that were members of a File Replication service (FRS) replica set and certificates that were issued by the Certificate Services are present.

# Restoring a Domain Controller Through Reinstallation

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Restoring a domain controller through reinstallation is the same process as creating a new domain controller. It does not involve restoring from backup media. This method relies on Active Directory replication to restore a domain controller to a working state, and it is valid only if another healthy domain controller exists in the same domain. This method is normally used on computers that function only as a domain controller.

Restoring through reinstallation is the only method by which a domain controller that is not part of the backup set can be restored. In addition, you might choose to use this method instead of a nonauthoritative restore because backup media is inaccessible or because this method is more convenient. Restoring a domain controller through reinstallation should not be a substitute for regular backup routines.

This method of restoring a domain controller requires a complete reinstallation of the operating system. It is recommended that before you install the operating system, you format the entire system disk, which will remove all information on the system disk. Ensure that any important or relevant data is moved or backed up before you perform these actions.

Bandwidth is the primary consideration for restoring a domain controller through reinstallation. The bandwidth that is required is directly proportional to the size of the Active Directory database and the time in which the domain controller is required to be in a functioning state. Ideally, the existing functional domain controller should be located in the same Active Directory site as the replicating domain controller (the new domain controller) to reduce the impact on the network and the time that the reinstallation takes to complete.

| **Note** |
| --- |
| Before you restore a domain controller through reinstallation, ensure that hardware failure is not the cause of the problem. If faulty hardware is not changed, restoring through reinstallation might not solve the problems with the domain controller. |

**Task requirements**

The following tools are required to perform the procedures for this task:

- Ntdsutil.exe
- Netdiag.exe
- Dcdiag.exe
- Dcpromo.exe

To complete this task, perform the following procedures:

1. If you plan to give the newly reinstalled domain controller the same name as the failed computer, use the following procedure to clean up server metadata to remove the NTDS Settings object of the failed domain controller:

   Clean up server metadata

   If you plan to give the new domain controller a different name, in addition to cleaning up server metadata, perform the following additional procedures:

   Delete a Server object from a site

   Delete a Computer object from the Domain Controllers OU

2. Install Windows Server 2003

   It is assumed that you will perform a fresh installation of Windows Server 2003. Prepare for installation of the operating system by partitioning or reformatting your hard disk drive, if necessary.

3. Verify DNS registration and functionality

4. Verify communication with other domain controllers

5. Verify the availability of the operations masters

6. Install Active Directory

   During the installation process, replication occurs, which ensures that the domain controller has an accurate and up-to-date copy of Active Directory. You have the option to use the same information for this domain controller as the domain controller that it is replacing: site placement, domain controller name, and domain membership should remain the same. If you plan to install the domain controller under a different name, see Installing a Domain Controller in an Existing Domain.

7. Verifying Active Directory Installation

# Clean up server metadata

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You perform the metadata cleanup process by using Ntdsutil.exe, a command-line tool that is automatically installed on all domain controllers. Metadata cleanup removes data from Active Directory that identifies a domain controller to the replication system. This procedure is required only for Active Directory domain controllers that were not successfully demoted using Dcpromo. This procedure does not have to be performed for domain member servers or client computers. On a domain controller that is running Windows Server 2003 with Service Pack 1 (SP1), metadata cleanup also removes File replication service (FRS) connections and attempts to transfer or seize any operations master roles (also known as flexible single master operations or FSMO roles) that the retired domain controller holds. These additional processes are performed automatically.

You can also use a script to clean up server metadata on most Windows operating systems. For information about using this script, see Remove Active Directory Domain Controller Metadata on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=123599).

**Administrative credentials**

- To complete this procedure, you must be a member of the Enterprise Admins group.

**To clean up server metadata**

1. Open a command prompt.
2. Type the following command, and then press ENTER:

   **ntdsutil**

3. At the **ntdsutil:** prompt, type:

   **metadata cleanup**

4. Perform metadata cleanup as follows:

   - If you are performing metadata cleanup by using the version of Ntdsutil.exe that is included with Windows Server 2003 SP1, at the **metadata cleanup:** prompt, type:

     **remove selected server** *ServerName*

     Or

     **remove selected server** *ServerName1* **on** *ServerName2*

     | Value | Definition |
     |---|---|
     | *ServerName, ServerName1* | The distinguished name of the domain controller whose metadata you want to remove, in the form cn=*ServerName*,cn=Servers,cn=*SiteName*, cn=Sites,cn=Configuration,dc=*ForestRootDomain* |
     | *ServerName2* | The DNS name of the domain controller to which you want to connect and from which you want to remove server metadata |

   - If you are performing metadata cleanup by using the version of Ntdsutil.exe that is included with Windows Server 2003 with no service pack, perform metadata cleanup as follows:

     a. At the **metadata cleanup:** prompt, type:

     **connection**

     b. At the **server connections:** prompt, type:

**connect to server** *Server*

c. At the **server connections:** prompt, type:

**quit**

d. At the **metadata cleanup:** prompt, type:

**select operation target**

e. At the **select operation target:** prompt, type:

**list sites**

A numbered list of sites appears.

f. At the **select operation target:** prompt, type:

**select site** *SiteNumber*

g. At the **select operation target:** prompt, type:

**list domains in site**

A numbered list of domains in the selected site appears.

h. At the **select operation target:** prompt, type:

**select domain** *DomainNumber*

i. At the **select operation target:** prompt, type:

**list servers in site**

A numbered list of servers in a domain and site appears.

j. At the **select operation target:** prompt, type:

**select server** *ServerNumber*

k. At the **select operation target:** prompt, type:

**quit**

l. At the **metadata cleanup:** prompt, type:

**remove selected server**

| Value | Description |
|---|---|
| *Server* | The DNS name of a domain controller that you want to connect to |
| *SiteNumber* | The number associated with the site of the server that you want to clean up that appears in the list |

| | |
|---|---|
| *DomainNumber* | The number associated with the domain of the server that you want to clean up that appears in the list |
| *ServerNumber* | The number associated with the server that you want to clean up that appears in the list |

At this point, Active Directory confirms that the domain controller was removed successfully. If you receive an error message that indicates that the object cannot be found, Active Directory might have already removed the domain controller.

5. To verify that the server was removed, type **list servers in site**, and then press ENTER. Ensure that the domain controller that you wanted to be removed is no longer displayed in the command output.
6. At the **metadata cleanup:** and **ntdsutil:** prompts, type **quit**.

| **Important** |
|---|
| If the domain controller that was removed was also a DNS server, ensure that references to it are removed from the **Name Servers** tab in the DNS console. To do this, open the DNS console using another DNS server in the domain (dnsmgmt.msc), and then click the domain name under **Forward Lookup Zones**. Remove any references to the domain controller that was removed from the domain. |

**See Also**
**Other Resources**
Article 216498 in the Microsoft Knowledge Base

# Delete a Server object from a site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When no Child objects are visible below the Server object in Active Directory Sites and Services, you can remove the Server object.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To delete a server object from a site**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then expand the site from which you want to delete a Server object.
3. If no Child objects appear below the Server object, right-click the Server object, and then click **Delete**.

| **Important** |
| --- |
| Do not delete a Server object that has a Child object. If an NTDS Settings or other Child object appears below the Server object you want to delete, either replication on the domain controller on which you are viewing the Configuration container has not occurred, or the server whose Server object you are removing has not been properly decommissioned. |

4. Click **Yes** to confirm your choice.

# Delete a Computer object from the Domain Controllers OU

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to delete the Computer object for a failed domain controller. If a domain controller fails and you cannot use the Dcpromo command to remove Active Directory, you must forcefully remove Active Directory and then clean up server metadata. When you perform Dcpromo normally, server metadata, the Computer object, and the Server object for the domain controller are deleted automatically. After you forcefully remove Active Directory, you must clean up server metadata for the failed domain controller and then delete the Server object and Computer object manually.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the domain controller that you are removing. O

**To delete a Computer object from the Domain Controllers organizational unit (OU)**

1. Open Active Directory Users and Computers.
2. Click the **Domain Controllers OU**.
3. In the details pane, right-click the Computer object that is associated with the failed domain controller, click **Delete**, and then click **Yes**.

**See Also**
**Concepts**
Forcing the Removal of a Domain Controller
Clean up server metadata
Delete a Server object from a site

# Verify DNS registration and functionality

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that Domain Name System (DNS) is functioning so that other domain controllers can be located.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To verify DNS registration and functionality**

1. Open a Command Prompt.
2. On a **Windows Server 2003** computer, type the following command, and then press ENTER:

   **netdiag /test:dns**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   On a **Windows Server 2008** or **Windows Server 2008 R2** computer, type the following command, and then press ENTER:

   **dcdiag /test:dns**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   If DNS is functioning, the last line of the response for all operating system versions is **DNS Test.....: Passed**. The verbose option lists specific information about what was tested. This information can help with troubleshooting if the test fails.

   If the test fails, do not attempt any additional steps until you determine and fix the problem that prevents proper DNS functionality.

# Verify communication with other domain controllers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that domain controllers can be located.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain users group in Active Directory.

**To verify communication with other domain controllers**

1. Open a Command Prompt.
2. On a **Windows Server 2003** computer, type the following command, and then press ENTER:

   **netdiag /test:dsgetdc**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   If domain controllers are successfully located, the last line of the response is **DC discovery test.....: Passed.** The verbose option lists the specific domain controllers that are located.

   On a **Windows Server 2008** or **Windows Server 2008 R2** computer, type the following command, and then press ENTER:

   **nltest /dclist:** <*domain name*>

   If domain controllers are successfully located, the last line of the response is **The command completed successfully**.

   If the tests fail on any of the operating system versions, do not attempt any additional steps until you determine and fix the problem that prevents communication with other domain controllers.

# Verify the availability of the operations masters

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that the operations masters can be located and that they are online and responding.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain users group in Active Directory.

| **Note** |
| --- |
| You can use these tests prior to installing Active Directory as well as afterward. To perform the test prior to installing Active Directory, you must use the **/s** option to indicate the name of a domain controller to use. You do not need the **/s** option to perform the test after installing Active Directory. The test automatically runs on the local domain controller where you are performing the test. The commands listed in this procedure show the **/s** option. If you are performing this test after installing Active Directory, omit the **/s** option. For a more detailed response from this command, you can use the verbose option by adding **/v** to the end of the command to see the detailed response. |

**To verify the availability of the operations masters**

1. Open a Command Prompt.
2. Type the following command to ensure that the operations masters can be located and then press ENTER:

   **dcdiag /s:** *domaincontroller* **/test:knowsofroleholders /verbose**

   where *domaincontroller* is the name of a domain controller in the domain in which you want to add the new domain controller. The verbose option provides a detailed list of the operations masters that were tested. Near the bottom of the screen, a message confirms that the test succeeded. If you use the verbose option, look carefully at the bottom part of the displayed output. The test confirmation message appears immediately after the list of operations masters. Press ENTER.
3. Type the following command to ensure that the operations masters are functioning properly and are available on the network:

   **dcdiag /s:** *domaincontroller* **/test:fsmocheck**

   where *domaincontroller* is the name of a domain controller in the domain in which you want to add the new domain controller. The verbose option provides a detailed list of the operations masters that were tested. Near the bottom of your screen, a message confirms that the test succeeded. Press ENTER.

   If these tests fail, do not attempt any additional steps until you determine and fix the problem that prevents locating operations masters and verifying that they are functioning properly.

# Install Active Directory

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the Active Directory Installation Wizard to install Active Directory on a member server in your domain to create an additional domain controller in an existing domain.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To install Active Directory**

1. Click **Start**, click **Run**, type **dcpromo** and then press ENTER.
2. The Active Directory Installation Wizard appears. At the Welcome screen, click **Next**.
3. For **Domain Controller Type**, select **Additional domain controller for an existing domain**. Click **Next**.
4. For **Network Credentials**, enter the user name, password, and domain for the user account that has permission to add this new domain controller to the domain. Click **Next**.
5. Enter the name of the domain that you want the new domain controller to host. Click **Next**.
6. For **Database and Log Locations**, enter the paths for the locations of the directory database (Ntds.dit) and the log files. For better performance, store the database and log files on separate physical disk drives. Click **Next**.
7. For **Shared System Volume**, enter the path where you want to locate the system volume (SYSVOL). Click **Next**.
8. Under **Directory Services Restore Mode Administrator Password**, enter the password that you want to use when you need to start Directory Services Restore Mode. Click **Next**.
9. The **Summary** screen displays a list of the items you chose. Verify that the information is correct, and then click **Next** to proceed with the installation.
10. The wizard proceeds to install Active Directory. When it finishes, the wizard displays a summary screen listing the domain and site in which the new domain controller is a member. Verify that this information is correct. Click **Finish** to close the wizard.
11. Click **Restart** to restart the domain controller.
12. Let the domain controller restart. If any message indicates that one or more services has failed to start, restart the domain controller one more time. If the initial replication cycles have not had enough time to complete during the first restart on a new domain controller, some services may be unable to start successfully. If the message appears during additional restarts, examine the event logs in Event Viewer to determine the cause of the problem.

# Administering Intersite Replication

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides information for administering intersite replication in the Microsoft Windows Server 2003 operating system.

**In this guide**


- Introduction to Administering Intersite Replication
- Managing Intersite Replication



**Acknowledgements**

Published: March 2005

Applies to: Windows Server 2003

Produced by: Microsoft Windows Server User Assistance team

Writer: Mary Hillman

Editor: Jim Becker

# Introduction to Administering Intersite Replication

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An Active Directory Site object represents a collection of Internet Protocol (IP) subnets, usually constituting a physical local area network (LAN). Multiple sites are connected for replication by Site Link objects.

Sites are used in Active Directory to:

- Enable clients to discover network resources (published shares, domain controllers) that are close to the physical location of the client, reducing network traffic over wide area network (WAN) links.
- Optimize replication between domain controllers.

Managing sites in Active Directory involves adding new subnet, site, and site link objects when the network grows, as well as configuring a schedule and cost for site links. You can modify the site link schedule, cost, or both, to optimize intersite replication. When conditions no longer require replication to a site, or clients no longer require the sites to discover network resources, you can remove the site and associated objects from Active Directory.

| **Note** |
| --- |
| Managing large hub-and-spoke topology or using the SMTP intersite replication transport is beyond the scope of this documentation. |

Managing sites:

- Enables clients to discover network resources (printers, published shares, domain controllers) that are close to the physical location of the client, reducing network traffic over wide area network (WAN) links.
- Optimizes replication between domain controllers.

**The KCC and Replication Topology**

The Knowledge Consistency Checker (KCC) uses site link configuration information to enable and optimize replication traffic by generating a least-cost replication topology. Within a site, for each directory partition, the KCC builds a ring topology that tries to set a maximum number of hops (3) between any two domain controllers. Between sites, the KCC creates a spanning tree of all intersite connections. Therefore, adding sites and domains increases the processing that is required by the KCC. Before adding to the site topology, be sure to consider the guidelines discussed in “Adding a new site” later in this document.

Significant changes to site topology can affect domain controller hardware requirements. For more information about domain controller hardware requirements, see Planning Domain Controller Capacity on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42682).

### Bridgehead Server Selection

By default, bridgehead servers are automatically selected by the intersite topology generator (ISTG) in each site. Alternatively, you can use Active Directory Sites and Services to select preferred bridgehead servers. However, it is recommended for Windows 2000 deployments that you do not select preferred bridgehead servers.

Selecting preferred bridgehead servers limits the bridgehead servers that the KCC can use to those that you have selected. If you use Active Directory Sites and Services to select any preferred bridgehead servers at all in a site, you must select as many as possible and you must select them for all domains that must be replicated to a different site. If you select preferred bridgehead servers for a domain and all preferred bridgehead servers for that domain become unavailable, replication of that domain to and from that site does not occur.

If you have selected one or more bridgehead servers, removing them all from the bridgehead servers list restores the automatic selection functionality to the ISTG.

# Managing Intersite Replication

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following tasks for managing intersite replication are described in this objective:

- Adding a New Site
- Linking Sites for Replication
- Changing Site Link Properties
- Moving a Domain Controller to a Different Site
- Removing a Site

# Adding a New Site

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Design teams or network architects might want to add sites as part of ongoing deployment. Although you typically create subnets to accommodate all address ranges in the network, you do not need to create sites for every location. Generally, sites are required for those locations that have domain controllers or other servers that run applications that depend on site topology, such as Distributed File System (DFS).

When the need for a site arises, the design team typically provides details about the placement and configuration of site links for the new site, as well as subnet assignments or creation if subnets are needed.

If a new range of Internet Protocol (IP) addresses is added to the network, create a Subnet object in Active Directory to correspond to the range of IP addresses. When you create a new Subnet object, you must associate it with a Site object. You can either associate the subnet with an existing site or create a new site first and then create the subnet and associate it with the new site.

**Task requirements**

The following tool is required to perform the procedures for this task:

- Active Directory Sites and Services

To complete this task, perform the following procedures:

1. Create a site object and add it to an existing site link
2. Associate a range of IP addresses with the site by using either of the following methods:
    - Create a subnet object or objects and associate them with the new site
    - Associate an existing subnet object with the new site
3. If you are creating both a new site and a new site link, after you create the new site and add it to an existing site link, Create a site link object and add the appropriate sites. Then, remove the site from the first site link that you added it to when you created the site, if appropriate.
4. Remove the site from the site link

# Create a site object and add it to an existing site link

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To create a new site, you must create a Site object and add it to a site link.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

**To create a site object and add it to an existing site link**

1. Open Active Directory Sites and Services.
2. Right-click the **Sites** container and then click **New Site**.
3. In the **Name** box, type the name of the site.
4. In the **Link Name** list, click a site link for this site, and then click **OK**.
5. In the **Active Directory** message box, read the information, and then click **OK**.

# Create a subnet object or objects and associate them with the new site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Create a Subnet object or objects and associate them with the new site you must have the following information:

- The site to which the subnet is to be associated.
- The network address or any IP address in the range.
- The subnet mask.

Active Directory Sites and Services converts this information into the subnet address.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

**To create a subnet object or objects and associate them with the new site**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, right-click **Subnets**, and then click **New Subnet**.
3. In the **New Object - Subnet** dialog box, in the **Address** box, type the network address or any IP address within the range of IP addresses for the subnet.
4. In the **Mask** box, type the subnet mask.
5. In the **Site Name** box, click the site to which this subnet is being associated, and then click **OK**.

# Associate an existing subnet object with the new site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Associate an existing subnet with a site under the following conditions:

- When you are removing the site to which the subnet was associated.
- When you have temporarily associated the subnet with a different site and want to associate it with its permanent site.

### Administrative Credentials

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

**To associate an existing subnet object with the new site**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then click the **Subnets** container.
3. In the details pane, right-click the subnet with which you want to associate the site, and then click **Properties**.
4. In the **Site** box, click the site with which to associate the subnet, and then click **OK**.

# Create a site link object and add the appropriate sites

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to create a Site Link object in the IP container and add the appropriate sites. To link sites for replication, create a Site Link object in the container for the intersite transport that will replicate the site, and add the sites to it.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group.

**To create a site link object**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and then the **Inter-Site Transports** container.
3. Right-click **IP**, and then click **New Site Link**.
4. In the **Name** box, type a name for the site link.
5. In the **Sites not in this site link** box, click a site that you want to add to the site link. Hold down the SHIFT key to click a second site that is adjacent in the list, or the CTRL key to click a second site that is not adjacent in the list.
6. After selecting all of the sites that you want added to the site link, click **Add**, and then click **OK**.

# Remove the site from the site link

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use Site Link properties to remove a site from a site link.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

**To remove a site from a site link**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and then the **Inter-Site Transports** container.
3. Click **IP**. In the details pane, right-click the site link from which you want to remove a site, and then click **Properties**.
4. In the **Sites in this site link** box, click the site you want to remove from the site link.
5. Click **Remove**, and then click **OK**.

# Linking Sites for Replication

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To link sites for replication, create a Site Link object in the IP transport container and add two or more sites to the link. Use a naming convention that includes the sites that you are linking. For example, if you want to link the site named Seattle to the site named Boston, you might name the site link SEA-BOS.

After you add two or more site names to a Site Link object, the bridgehead servers in the respective sites replicate between the sites according to the replication schedule, cost, and interval settings on the Site Link object. For information about modifying the default settings, see Changing Site Link Properties.

At least two sites must exist when you create a site link. If you are adding a site link to connect a new site to an existing site, create the new site first and then create the site link. For information about creating a site, see Adding a New Site.

**Task Requirements**

The following tool is required to perform the procedures for this task:

- Active Directory Sites and Services

To complete this task, perform the following procedures:

1. Create a site link object and add the appropriate sites
2. By default, the KCC runs every 15 minutes to generate the replication topology. To generate the intersite topology immediately, perform the following two procedures:
    - Determine the ISTG role owner for a site
    - Generate the replication topology on the ISTG

# Create a site link object and add the appropriate sites

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to create a Site Link object in the IP container and add the appropriate sites. To link sites for replication, create a Site Link object in the container for the intersite transport that will replicate the site, and add the sites to it.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group.

**To create a site link object**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and then the **Inter-Site Transports** container.
3. Right-click **IP**, and then click **New Site Link**.
4. In the **Name** box, type a name for the site link.
5. In the **Sites not in this site link** box, click a site that you want to add to the site link. Hold down the SHIFT key to click a second site that is adjacent in the list, or the CTRL key to click a second site that is not adjacent in the list.
6. After selecting all of the sites that you want added to the site link, click **Add**, and then click **OK**.

# Determine the ISTG role owner for a site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to view the NTDS Site Settings object properties and determine the Intersite Topology Generator (ISTG) role owner for the site.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group.

**To determine the ISTG role owner for a site**

1. Open Active Directory Sites and Services.
2. Click the site object whose ISTG you want to determine.
3. In the details pane, right-click the **NTDS Site Settings** object, and then click **Properties**. The current role owner appears in the **Server** box under **Inter-Site Topology Generator**.

# Generate the replication topology on the ISTG

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Knowledge Consistency Checker (KCC) runs by default every 15 minutes. If you want to initiate topology regeneration immediately, you can force the KCC to run as follows:

- To generate the intersite replication topology, run the KCC on the domain controller in the site that holds the ISTG role.
- To generate the intrasite replication topology, run the KCC on any domain controller in the site that does not hold the ISTG role.

| **Note** |
| --- |
| To generate the replication topology on the ISTG, you must first complete the procedure: Determine the ISTG role owner for a site. |

### Administrative Credentials

To perform this procedure, you must be a member of the Enterprise Admins group.

**To generate the replication topology on the ISTG**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then expand the site that contains the server on which you want to run the KCC.
3. Expand the **Servers** container, and then click the **Server** object for the ISTG.
4. In the details pane, right-click **NTDS Settings**, click **All Tasks**, and then click **Check Replication Topology**.
5. In the **Check Replication Topology** message box, click **OK**.

# Changing Site Link Properties

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To control which sites replicate directly with each other and when, use the cost, schedule, and interval properties on the Site Link object.

These settings control intersite replication as follows:

- **Schedule**: The time during which replication can occur (the default setting allows replication at all times).
- **Interval**: The number of minutes between replication polling by intersite replication partners within the open schedule window (default is every 180 minutes).
- **Cost**: The relative priority of the link (default is 100). Lower relative cost increases the priority of the link over other higher-cost links.

Consult your design documentation for information about values to set for site link properties.

**Task Requirements**

The following tool is required to perform the procedures for this task:

- Active Directory Sites and Services

To complete this task, perform the following procedures:

1. Configure the site link schedule to identify times during which intersite replication can occur
2. Configure the site link interval to identify how often replication polling can occur during the schedule window
3. Configure the site link cost to establish a priority for replication routing
4. Generate the intersite topology by performing the following two procedures:
    - Determine the ISTG role owner for a site
    - Generate the replication topology on the ISTG

# Configure the site link schedule to identify times during which intersite replication can occur

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the properties on the Site Link object to define when replication is allowed. Obtain the schedule from your design team.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

**To configure the site link schedule**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and the **Inter-Site Transports** container, and then click the **IP** container.
3. In the details pane, right-click the Site Link object you want to configure, and then click **Properties**.
4. In the **SiteLinkName Properties** dialog box, click **Change Schedule**.
5. In the **Schedule for SiteLinkName** dialog box, select the block of days and hours during which you want replication to occur or not occur (available or not available), and then click the appropriate option.
6. Click **OK** twice.

# Configure the site link interval to identify how often replication polling can occur during the schedule window

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the properties on the Site Link object to determine how often during the available replication schedule you want bridgehead servers to poll their intersite replication partners for changes. Obtain the interval value from your design team.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

**To configure the site link interval**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and the **Inter-Site Transports** container, and then click the **IP** container.
3. In the details pane, right-click the Site Link object you want to configure, and then click **Properties**.
4. In the **Replicate every _____ minutes** box, specify the number of minutes for the intervals at which replication polling occurs during an open schedule, and then click **OK**.

# Configure the site link cost to establish a priority for replication routing

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When creating or modifying site links, use the object properties to configure the relative cost of using the site link.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

**To configure the site link cost**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and the **Inter-Site Transports** container, and then click the **IP** container.
3. In the details pane, right-click the Site Link object you want to configure, and then click **Properties**.
4. In the **Cost** box, specify the number for the comparative cost of using the site link, and then click **OK**.

# Determine the ISTG role owner for a site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to view the NTDS Site Settings object properties and determine the Intersite Topology Generator (ISTG) role owner for the site.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group.

**To determine the ISTG role owner for a site**

1. Open Active Directory Sites and Services.
2. Click the site object whose ISTG you want to determine.
3. In the details pane, right-click the **NTDS Site Settings** object, and then click **Properties**. The current role owner appears in the **Server** box under **Inter-Site Topology Generator**.

# Generate the replication topology on the ISTG

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Knowledge Consistency Checker (KCC) runs by default every 15 minutes. If you want to initiate topology regeneration immediately, you can force the KCC to run as follows:

- To generate the intersite replication topology, run the KCC on the domain controller in the site that holds the ISTG role.
- To generate the intrasite replication topology, run the KCC on any domain controller in the site that does not hold the ISTG role.

| **Note** |
| --- |
| To generate the replication topology on the ISTG, you must first complete the procedure: Determine the ISTG role owner for a site. |

### Administrative Credentials

To perform this procedure, you must be a member of the Enterprise Admins group.

**To generate the replication topology on the ISTG**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then expand the site that contains the server on which you want to run the KCC.
3. Expand the **Servers** container, and then click the **Server** object for the ISTG.
4. In the details pane, right-click **NTDS Settings**, click **All Tasks**, and then click **Check Replication Topology**.
5. In the **Check Replication Topology** message box, click **OK**.

# Moving a Domain Controller to a Different Site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you change the IP address or the subnet-to-site association of a domain controller after Active Directory is installed on the server, the Server object does not change sites automatically. You must move it to the new site manually. When you move the Server object, the Net Logon service on the domain controller registers DNS SRV resource records for the appropriate site.

**TCP/IP Settings**

When you move a domain controller to a different site, if an IP address of the domain controller is statically configured, then you must change the TCP/IP settings accordingly. The IP address of the domain controller must map to a Subnet object that is associated with the site to which you are moving the domain controller. If the IP address of a domain controller does not match the site in which the Server object appears, the domain controller might be forced to communicate over a potentially slow WAN link to locate resources rather than locating resources in its own site.

Prior to moving the domain controller, ensure that the following TCP/IP client values are appropriate for the new location:

- IP address, including the subnet mask and default gateway
- DNS server addresses
- WINS server addresses (if appropriate)

If the domain controller that you are moving is a DNS server, you must also:

- Change the TCP/IP settings on any clients that have static references to the domain controller as the preferred or alternate DNS server.
- Determine whether the parent DNS zone of any zone that is hosted by this DNS server contains a delegation to this DNS server. If yes, update the IP address in all such delegations. For information about creating DNS delegations, see Verifying Active Directory Installation.

**Preferred Bridgehead Server Status**

Before moving any Server object, check the Server object to see whether it is acting as a preferred bridgehead server for the site. This condition has ISTG implications in both sites, as follows:

- Site to which you are moving the server: If you move a preferred bridgehead server to a different site, it becomes a preferred bridgehead server in the new site. If preferred bridgehead servers are not currently in use in this site, the ISTG behavior in this site changes to support preferred bridgehead servers. For this reason, you must either configure the server to not be a preferred bridgehead server (recommended), or select additional preferred bridgehead servers in the site (not recommended).

Site from which you are moving the server: If the server is the last preferred bridgehead server in the original site for its domain, and if other domain controllers for the domain are in the site, the ISTG selects a bridgehead server for the domain. If you use preferred bridgehead servers, always select more than one server as the preferred bridgehead server for the domain. If, after the removal of this domain controller from the site, multiple domain controllers remain that are hosting the same domain and only one of them is configured as a preferred bridgehead server, either configure the server to not be a preferred bridgehead server (recommended), or select additional preferred bridgehead servers hosting the same domain in the site (not recommended).

| **Note** |
| --- |
| If you select preferred bridgehead servers and all selected preferred bridgehead servers for a domain are unavailable in the site, the ISTG does not select a new bridgehead server. In this case, replication of this domain to and from other sites does not occur. However, if no preferred bridgehead server is selected for a domain or transport (through administrator error or as the result of moving the only preferred bridgehead server to a different site), the ISTG automatically selects a preferred bridgehead server for the domain and replication proceeds as scheduled. |

**Task Requirements**

- My Network Places
- DNS snap-in
- Active Directory Sites and Services
- Adsiedit.msc

To complete this task, perform the following procedures in order:

1. Change the static IP address of a domain controller
2. Create a delegation for a domain controller

   If the parent DNS zone of any zone that is hosted by this DNS server contains a delegation to this DNS server, use this procedure to update the IP address in all such delegations.

   If your forest root domain has a parent DNS domain, perform this procedure on a DNS server in the parent domain. If you just added a new domain controller to a child domain, perform this procedure on a DNS server in the DNS parent domain. By following recommended practices, the parent domain is the forest root domain.

3. Verify that an IP address maps to a subnet and determine the site association
4. Determine whether the server is a preferred bridgehead server
5. Configure the server to not be a preferred bridgehead server
6. Move the Server object to the new site

# Change the static IP address of a domain controller

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure includes changing all appropriate TCP/IP values, including preferred and alternate Domain Name System (DNS) servers, as well as Windows Internet Name Service (WINS) servers (if appropriate). Obtain these values from the design team.

| **Note** |
| --- |
| If you change the static IP address of a domain controller, you must also change related TCP/IP settings accordingly. This includes changing the TCP/IP settings for all client computers that rely on the domain controller for related services, such as DNS services. If the affected client computers are configured with static IP addresses, the procedures for changing them should be covered in the operating system help. As an example, see Change TCP/IP settings (http://go.microsoft.com/fwlink/?LinkId=130602). If a Dynamic Host Configuration Protocol (DHCP) server is used to distribute DNS server IP addresses, you may have to modify DHCP server options. For a list of options, see DHCP Options Supported by Clients (http://go.microsoft.com/fwlink/?LinkId=130603). For specific instructions about changing DHCP options, see Modify an option (http://go.microsoft.com/fwlink/?LinkId=130604). |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in the domain of the domain controller whose IP address you are changing.

**To change the static IP address of a domain controller**

1. Log on locally (also known as interactively) to the system console of the domain controller whose IP address you want to change. If you are not able to log on to the domain controller by using the domain, you may have to start the domain controller in Directory Services Restore Mode (DSRM). For more information, see Restart the domain controller in Directory Services Restore Mode locally (http://go.microsoft.com/fwlink/?LinkId=140599).
2. On the desktop, right-click **My Network Places**, and then click **Properties**.
3. In the **Network Connections** dialog box, right-click **Local Area Connection**, and then click **Properties**.
4. In the **Local Area Connection Properties** dialog box, double-click **Internet Protocol (TCP/IP)**.
5. In the **Internet Protocol (TCP/IP) Properties** dialog box, in the **IP address** box, type the new address.
6. In the **Subnet mask** box, type the subnet mask.
7. In the **Default gateway** box, type the default gateway.
8. In the **Preferred DNS server** box, type the address of the DNS server that this computer contacts.
9. In the **Alternate DNS server** box, type the address of the DNS server that this computer contacts if the preferred server is unavailable.
10. If this domain controller uses WINS servers, click **Advanced** and then, in the **Advanced TCP/IP Settings** dialog box, click the **WINS** tab.
11. If an address in the list is no longer appropriate, click the address, and then click **Edit**.
12. In the **TCP/IP WINS Server** dialog box, type the new address, and then click **OK**.
13. Repeat steps 11 and 12 for all addresses that need to be changed, and then click **OK** twice to close the **TCP/IP WINS Server** dialog box and the **Advanced TCP/IP Settings** dialog box.
14. Click **OK** to close the **Internet Protocol (TCP/IP) Properties** dialog box.

After you change the IP address of a domain controller, you should run the **ipconfig /registerdns** command to register the host record and **dcdiag /fix** command to ensure that service records are appropriately registered with DNS. For more information, see Dcdiag Overview and subordinate topics for additional information about the Dcdiag tool (http://go.microsoft.com/fwlink/?LinkId=130605).

Changing the IP settings of a server does not affect the share resources or shared permissions on that server, if the name

resolution structure DNS and WINS settings are correctly configured. However, if network drives or passive connections (connections that are made manually from a command prompt or run line) are mapped using the IP address, an update is required. For example, if a client computer has G: drive mapped using the following command **net use g: \\192.168.0.199\data** and the IP address of the server that hosts the Data shared folder is changed from 192.168.0.199 to 192.168.1.200, the new G: drive mapping command should be changed to **net use g: \\192.168.1.200\data**. A better solution would be to ensure that DNS name resolution is working properly and to use the server name, as opposed to the IP address, in the command. For example, if the server name is DC1, the command to map a G: drive to the Data share on the server is **net use g: \\dc1\data**. It changes only if the server name changes; it is not affected if the IP address of the server changes.

# Create a delegation for a domain controller

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to create a delegation for a new domain controller that is also a DNS server in the parent DNS domain.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To create a delegation for a domain controller**

1. Open the DNS snap-in.
2. Navigate to C*hildDomain* (where C*hildDomain* is the name of the child domain) in the console tree.
3. In the console tree, right-click C*hildDomain*, and then click **Properties**.
4. In the C*hildDomain* **Properties** sheet, on the **Name Servers** tab, click **Add**.
5. In the **New Resource Record** dialog box, in the **Server fully qualified domain name (FQDN)** box, type *ChildDC.ChildDomain.ParentDomain* (where *ChildDC* is the name of the new domain controller, *ChildDomain* is the name of the child domain, and *ParentDomain* is the name of the parent domain).
6. In the **New Resource Record** dialog box, in the **IP address** box, type *IPAddress* (where *IPAddress* is the IP address of the child domain controller), click **Add**, and then click **OK**.

# Verify that an IP address maps to a subnet and determine the site association

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to determine the site to which you want to add a Server object prior to installing Active Directory, or to verify the appropriate site prior to moving a Server object to it.

To be associated with a site, the IP address of a domain controller must map to a Subnet object that is defined in Active Directory. The site to which the subnet is associated is the site of the domain controller.

The subnet address, which is computed from the IP network address and the subnet mask, is the name of a Subnet object in Active Directory. When you know the subnet address, you can locate the Subnet object and determine the site to which the subnet is associated.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group.

**To verify that an IP address maps to a subnet and determine the site association**

1. Log on locally or open a Terminal Services connection to the server for which you want to check the IP address.
2. On the desktop, right-click **My Network Places**, and then click **Properties**.
3. In the **Network Connections** dialog box, right-click **Local Area Connection**, and then click **Properties**.
4. Double-click **Internet Protocol (TCP/IP)**.
5. Use the values in **IP address** and **Subnet mask** to calculate the subnet address and then click **OK**.
6. Click **OK** again and close the **Network Connections** dialog box.
7. Open Active Directory Sites and Services.
8. Expand the **Sites** container, and then click the **Subnets** container.
9. In the **Name** column in the details pane, find the Subnet object that matches the subnet address.
10. In the **Site** column, note the site to which the IP subnet address is associated.

    If the site that appears in the **Site** box is not the appropriate site, contact a supervisor and find out whether the IP address is incorrect or whether to move the Server object to the site indicated by the subnet.

# Determine whether the server is a preferred bridgehead server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Preferred bridgehead servers are distinguished by a property on the Server object that adds the server to the preferred bridgehead server list for the IP transport.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group.

**To view the list of preferred bridgehead servers**

1. Click **Start**, click **Run**, type the following command and then press ENTER:

   **adsiedit.msc**

2. In ADSI Edit, expand the **Configuration Container** and then expand **CN=Configuration,DC=***ForestRootDomainName,* **CN=Sites, and CN=Inter-Site Transports**.
3. Right-click **CN=IP** and then click **Properties**.
4. In the **Attributes** box, double-click **bridgeheadServerListBL**.
5. If any preferred bridgehead servers are selected in any site in the forest, the **Values** box displays the distinguished name for each server object that is currently selected as a preferred bridgehead server.

# Configure the server to not be a preferred bridgehead server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the Server object properties to remove a preferred bridgehead server from the IP transport.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To configure the server to not be a preferred bridgehead server**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then expand the site of the preferred bridgehead server.
3. Expand the **Servers** node to display the list of domain controllers currently configured for that site.
4. Right-click the server you want to remove, and then click **Properties**.
5. If **IP** appears in the list that marks this server as a bridgehead server for the IP transport, click **IP**, click **Remove**, and then click **OK**.

# Move the Server object to the new site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Moving a Server object requires that the IP address of the domain controller maps to the site to which you are moving the Server object. Before performing this procedure, verify that the IP address maps to the target site.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group.

**To move the server object to the new site**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and the site in which the server object resides.
3. Expand the **Servers** container to display the domain controllers that are currently configured for that site.
4. Right-click the Server object you want to move, and then click **Move**.
5. In the **Site Name** box, click the destination site, and then click **OK**.
6. Expand the Site object to which you moved the server, and then expand the **Servers** container.
7. Verify that an object for the server you moved exists.
8. Expand the Server object and verify that an NTDS Settings object exists.

Within an hour, the Net Logon service on the domain controller registers the new site information in DNS. Wait an hour and then open Event Viewer and connect to the domain controller whose Server object you moved. Review the System log for Net Logon errors regarding registration of SRV resource records in DNS that have occurred within the last hour. The absence of errors indicates that Net Logon has updated DNS with site-specific SRV resource records. Net Logon event ID 5774 indicates that the registration of DNS resource records has failed. If this error occurs, contact a supervisor and pursue DNS troubleshooting.

# Removing a Site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If domain controllers are no longer needed in a network location, you can remove them from the site and then delete the Site object. Before deleting the site, you must remove domain controllers from the site either by removing it entirely or by moving it to a new location.

- To remove the domain controller, remove Active Directory from the server and then delete the Server object from the site in Active Directory.
- To retain the domain controller in a different location, move the domain controller to a different site and then move the Server object to the respective site in Active Directory.

Domain controllers can host other applications that depend on site topology and publish objects as Child objects of the respective Server object. For example, when MOM or Message Queuing is running on a domain controller, these applications create Child objects beneath the Server object. In addition, a server running Message Queuing that is not a domain controller and is configured to be a routing server running Message Queuing creates a Server object in the Sites container. Removing the application from the server automatically removes the Child object below the respective Server object. However, the Server object is not removed automatically.

When all applications have been removed from the server (no Child objects appear beneath the Server object), you can remove the Server object. After the application is removed from the server, a replication cycle might be required before Child objects are no longer visible below the Server object.

After you delete or move the Server objects but before you delete the Site object, reconcile the following objects:

IP addresses:

- If the addresses are being reassigned to a different site, associate the Subnet object or objects with that site. Any clients using the addresses for the decommissioned site will thereafter be assigned automatically to the other site.
- If the IP addresses will no longer be used on the network, delete the corresponding Subnet object or objects.

Site Link objects:

- If the site you are removing is added to a site link containing only two sites, delete the Site Link object.
- If the site you are removing is added to a site link that contains more than two sites, do not delete this Site Link object.

Before removing a site, you need to consider the implications. If the site you are removing is added to more than one site link, it might be an interim site between other sites that are added to this site link. Deleting the site might disconnect the outer sites from each other. In this case, the site links must be reconciled according to the instructions of the design team.

**Task Requirements**

The following tool is required to perform the procedures for this task:

- Active Directory Sites and Services

To complete this task, perform the following procedures:

1. Determine whether a Server object has child objects
2. Delete a Server object from a site
3. Delete the Site Link object

4. Associate the subnet or subnets with the appropriate site

5. Delete the Site object

6. To avoid replication errors on bridgehead servers in other sites that received replication from the site that has been removed, generate the intersite topology in those sites by performing the following two procedures:

    - Determine the ISTG role owner for a site

    - Generate the replication topology on the ISTG

# Determine whether a Server object has child objects

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After Active Directory is properly installed on a domain controller, the Server object for the domain controller will have a Child NTDS-Settings object. Other applications that are running on domain controllers can also publish Child objects.

Prior to deleting a Server object from the Servers container for a site, verify that the Server object has no Child objects. If a Child object appears, do not delete the Server object.

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Users group.

**To determine whether a server object has child objects**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and expand the site of the Server object.
3. Expand the **Servers** container, and then expand the Server object to view any Child objects.

# Delete a Server object from a site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When no Child objects are visible below the Server object in Active Directory Sites and Services, you can remove the Server object.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To delete a server object from a site**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then expand the site from which you want to delete a Server object.
3. If no Child objects appear below the Server object, right-click the Server object, and then click **Delete**.

| **Important** |
| --- |
| Do not delete a Server object that has a Child object. If an NTDS Settings or other Child object appears below the Server object you want to delete, either replication on the domain controller on which you are viewing the Configuration container has not occurred, or the server whose Server object you are removing has not been properly decommissioned. |

4. Click **Yes** to confirm your choice.

# Delete the Site Link object

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following procedure to delete a Site Link object.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group.

**To delete a site link object**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and the **Inter-Site Transports container**, and then click the **IP container**.
3. In the details pane, right-click the Site Link object you want to delete, and then click **Delete**.
4. Click **Yes** to confirm your choice.

# Associate the subnet or subnets with the appropriate site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Associate an existing subnet with a site under the following conditions:

- When you are removing the site to which the subnet was associated.
- When you have temporarily associated the subnet with a different site and want to associate it with its permanent site.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group.

**To associate the subnet or subnets with the appropriate site**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then click the **Subnets** container.
3. In the details pane, right-click the subnet with which you want to associate the site, and then click **Properties**.
4. In the **Site** box, click the site with which to associate the subnet, and then click **OK**.

   If the IP addresses are no longer in use, delete the Subnet object or objects with which the addresses are associated.

# Delete the Site object

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Delete a Site object only after you have removed all Server objects from the site and have reassociated the subnets with a different site. The Servers container is deleted when you delete the site.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group.

**To delete the site object**

1. Open Active Directory Sites and Services and click the **Sites** container.
2. In the details pane, right-click the site you want to delete, and then click **Delete**.
3. Click **Yes** to confirm your choice.
4. In the Active Directory message box, read the information, and then click **Yes** to delete the site and its Servers container object.

# Determine the ISTG role owner for a site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to view the NTDS Site Settings object properties and determine the Intersite Topology Generator (ISTG) role owner for the site.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group.

**To determine the ISTG role owner for a site**

1. Open Active Directory Sites and Services.
2. Click the site object whose ISTG you want to determine.
3. In the details pane, right-click the **NTDS Site Settings** object, and then click **Properties**. The current role owner appears in the **Server** box under **Inter-Site Topology Generator**.

# Generate the replication topology on the ISTG

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Knowledge Consistency Checker (KCC) runs by default every 15 minutes. If you want to initiate topology regeneration immediately, you can force the KCC to run as follows:

- To generate the intersite replication topology, run the KCC on the domain controller in the site that holds the ISTG role.
- To generate the intrasite replication topology, run the KCC on any domain controller in the site that does not hold the ISTG role.

| **Note** |
| --- |
| To generate the replication topology on the ISTG, you must first complete the procedure: Determine the ISTG role owner for a site. |

### Administrative Credentials

To perform this procedure, you must be a member of the Enterprise Admins group.

**To generate the replication topology on the ISTG**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then expand the site that contains the server on which you want to run the KCC.
3. Expand the **Servers** container, and then click the **Server** object for the ISTG.
4. In the details pane, right-click **NTDS Settings**, click **All Tasks**, and then click **Check Replication Topology**.
5. In the **Check Replication Topology** message box, click **OK**.

# Administering the Active Directory Database

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides information for administering the Active Directory database in the Microsoft Windows Server 2003 operating system.

**In this guide**


- Introduction to Administering the Active Directory Database
- Managing the Active Directory Database



**Acknowledgements**

Published: March 2005

Applies to: Windows Server 2003

Produced by: Microsoft Windows Server User Assistance team

Writer: Mary Hillman

Editor: Jim Becker

# Introduction to Administering the Active Directory Database

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Active Directory is stored in the Ntds.dit database file. In addition to this file, the directory uses log files, which store transactions prior to committing them to the database file. For best performance, store the log files and the database on separate hard drives.

The Active Directory database is a self-maintained system and requires no daily maintenance, other than regular backup, during ordinary operation. However, it may need to be managed if the following conditions occur:

- Low disk space
- Pending or current hardware failure
- A need to recover physical space following bulk deletion or removal of the global catalog

Monitor free disk space on the partition or partitions that store the directory database and logs. The following are the recommended parameters for free space:

- Ntds.dit partition: The greater of 20 percent of the Ntds.dit file size or 500 megabytes (MB).
- Log file partition: The greater of 20 percent of the combined log files size or 500 MB.
- Ntds.dit and logs on the same volume: The greater of 1 gigabyte (GB) or 20 percent of the combined Ntds.dit and log files sizes.

During ordinary operation, the customer will delete objects from Active Directory. When an object is deleted, it results in white space (or unused space) being created in the database. On a regular basis, the database will consolidate this white space through a process called defragmentation, and this white space will be reused when new objects are added (without adding any size to the file itself). This automatic online defragmentation redistributes and retains white space for use by the database, but does not release it to the file system. Therefore, the database size does not shrink, even though objects might be deleted. In cases where the data is decreased significantly, such as when the global catalog is removed from a domain controller, white space is not automatically returned to the file system. Although this condition does not affect database operation, it does result in large amounts of white space in the database. You can use offline defragmentation to decrease the size of the database file by returning white space from the database file to the file system.

Managing the Active Directory database also allows you to upgrade or replace the disk on which the database or log files are stored or to move the files to a different location, either permanently or temporarily.

Prior to performing any procedures that affect the directory database, be sure that you have a current system state backup. For information about performing system state backup, see Back up system state.

To manage the database file itself, you must take the domain controller offline by restarting in Directory Services Restore Mode, and then use Ntdsutil.exe to manage the file.

| **Note** |
| --- |
| NTFS disk compression is not supported for the database and log files. |

# Managing the Active Directory Database

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following tasks for managing the Active Directory database are described in this objective:

- Relocating Active Directory Database Files
- Returning Unused Disk Space from the Active Directory Database to the File System
- Guarding Against Accidental Bulk Deletions in Active Directory

# Relocating Active Directory Database Files

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following conditions require moving database files:

- Hardware maintenance: If the physical disk on which the database or log files are stored requires upgrading or maintenance, the database files must be moved, either temporarily or permanently.
- Low disk space: When free disk space is low on the logical drive that stores the database file (Ntds.dit), the log files, or both, first verify that no other files are causing the problem. If the database file or log files are the cause of the growth, then provide more disk space by taking one of the following actions:
- Expand the partition on the disk that currently stores the database file, the log files, or both. This procedure does not change the path to the files and does not require updating the registry.
- Use Ntdsutil.exe to move the database file, the log files, or both to a larger existing partition. If you are not using Ntdsutil.exe when moving files to a different partition, you will need to manually update the registry.

If the path to the database file or log files will change as a result of moving the files, be sure that you:

- Use Ntdsutil.exe to move the files (rather than copying them) so that the registry is updated with the new path. Even if you are moving the files only temporarily, use Ntdsutil.exe to move files locally so that the registry remains current.
- Perform a system state backup as soon as the move is complete so that the restore procedure uses the correct path.
- Verify that the correct permissions are applied on the destination folder following the move. Revise permissions to those that are required to protect the database files, if needed.

The registry entries that Ntdsutil.exe updates when you move the database file are as follows:

In **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS**

**Parameters**:

- Database backup path
- Digital Signature Algorithm (DSA) database file
- DSA working directory

The registry entry that Ntdsutil.exe updates when you move the log files is as follows:

In **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS**

**Parameters**:

- Database log files path

**Disk space requirements for relocating Active Directory database files**

**Temporary location**. Free space on the destination drive equivalent to at least the current size of the database file, the combined log files, or both, depending on which files you are moving.

**Permanent location**. Free space on the destination NTFS drive equivalent to at least the size specified below, plus space to accommodate anticipated growth, depending on which file or files you are moving.

**Caution**

| |
|---|
| The drive that is the permanent location of the database file or log files must be formatted as NTFS. |

- Database file only: The size of the database file plus 20 percent of the Ntds.dit file or 500 MB, whichever is greater.
- Log files only: The size of the combined log files plus 20 percent of the combined logs or 500 MB, whichever is greater.
- Database and logs. If the database and log files are stored on the same partition, free space should be at least 20 percent of the combined Ntds.dit and log files, or 1 GB, whichever is greater.

| **Important** |
|---|
| The preceding levels are minimum recommended levels. Therefore, adding additional space according to anticipated growth is recommended. |

**Task Requirements**

The following tools are required to perform the procedures for this task:

- net use
- dir
- xcopy
- Ntdsutil.exe
- Backup software
- Windows Explorer

| **Note** |
|---|
| If you replace or reconfigure a drive that stores the SYSVOL folder, you must first move the SYSVOL folder manually. For information about moving SYSVOL manually, see Relocating SYSVOL Manually. |

To complete this task, perform the following procedures:

| **Note** |
|---|
| The domain controller will not be available during the time in which files are being moved and until the move is verified. Ensure that alternate domain controllers are available during the file relocation to handle the capacity. |

1. Determine the size and location of the Active Directory database by using one of the following procedures:
   - Determine the database size and location online
   - Determine the database size and location offline
2. Compare the size of the directory database files to the volume size
3. Back up system state

System state includes the database file and log files as well as SYSVOL and Net Logon shared folders, among other things. Always ensure that you have a current backup prior to moving database files.

4. Restart the domain controller in Directory Services Restore Mode by using one of the following methods:

   - Restart the domain controller in Directory Services Restore Mode locally
   - Restart the domain controller in Directory Services Restore Mode Remotely

5. Move or copy the directory database and log files by performing one of the following procedures:

   - Move the directory database and log files to a local drive
   - Copy the directory database and log files to a remote share

     The shared folder on a remote drive must have enough free space to hold the database file (Ntds.dit) and log files. Create separate subdirectories for copying the database file and the log files.

6. Back up system state

# Determine the database size and location online

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When determining the database size and location online, the size is reported in bytes. If you must manage the database file, the log files, or both, first determine the location and size of the files. By default, the database file and associated log files are stored in the *%systemroot%*\NTDS directory.

| **Important** |
| --- |
| Be sure to use the same method to check file sizes when you compare them. The size is reported differently, depending on whether the domain controller is online or offline. For information about determining database size offline, see Determine the database size and location offline. |

You can also use the Search command on the Start menu to locate the database file (Ntds.dit) or the edb*.log file for the location of the database and log files, respectively.

If you have set garbage collection logging to report free disk space, then event ID 1646 in the Active Directory service log also reports the size of the database file: "Total allocated hard disk space (megabytes):"

Alternatively, you can determine the size of the database file by listing the contents of the directory that contains the files.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To determine the database size and location online**

1. On the domain controller on which you want to manage database files, open a command prompt and change directories to the directory containing the files you want to manage.
2. Run the dir command to examine the database size. In the following example, Ntds.dit file and the log files are stored in the same directory. In the example, the files take up 58,761,216 bytes of disk space.

   **H:\NTDS>dir**

   **Volume in drive H has no label.**

   **Volume Serial Number is 003D-0E9E**

   **Directory of H:\NTDS**

   **01/29/2002 11:04 AM <DIR> .**

   **01/29/2002 11:04 AM <DIR> ..**

   **01/28/2002 03:03 PM <DIR> Drop**

   **01/29/2002 10:29 AM 8,192 edb.chk**

   **01/29/2002 10:29 AM 10,485,760 edb.log**

   **01/29/2002 10:29 AM 10,485,760 edb00001.log**

   **01/29/2002 10:29 AM 14,696,448 ntds.dit**

   **01/28/2002 02:54 PM 10,485,760 res1.log**

   **01/28/2002 02:54 PM 10,485,760 res2.log**

   **7 File(s) 58,761,216 bytes**

   **3 Dir(s) 779,284,480 bytes free**

# Determine the database size and location offline

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When determining the database size and location offline, the size is reported in megabytes (MB). Use this method if the domain controller is already started in Directory Services Restore Mode.

| **Important** |
| --- |
| Be sure to use the same method to check file sizes when you compare them. The size is reported differently, depending on whether the domain controller is online or offline. For information about determining database size online, see Determine the database size and location online. |

You can also use the Search command on the Start menu to locate the database file (Ntds.dit) or the edb*.log file for the location of the database and log files, respectively.

If you have set garbage collection logging to report free disk space, then event ID 1646 in the Active Directory service log also reports the size of the database file: "Total allocated hard disk space (megabytes):"

Alternatively, you can determine the size of the database file by listing the contents of the directory that contains the files.

**Administrative Credentials**

To perform this procedure, you must provide the Directory Services Restore Mode password for the local administrator account.

**To determine the database size and location offline**

1. With the domain controller in Directory Services Restore Mode, open a command prompt, type **ntdsutil** and then press ENTER
2. At the **ntdsutil:** prompt, type **files** and then press ENTER.
3. At the **file maintenance:** prompt, type **info** and press ENTER.
4. At the **file maintenance:** prompt, type **quit** and then press ENTER. Type **quit** and then press ENTER again to quit Ntdsutil.exe.

# Compare the size of the directory database files to the volume size

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before moving any files in response to low disk space, verify that no other files on the volume are responsible for the condition of low disk space.

You might need to relocate the database file, the log files, or both, if disk space on the volume on which they are stored becomes low. Before moving the database file or log files, examine the size of the database folder, logs folder, or both, if they are stored in the same location, relative to the size of the volume to verify that these files are the cause of low disk space. Include the size of the SYSVOL folder if it is on the same partition.

**Administrative Credentials**

If you are online when comparing the size of the directory database files, you must be a member of the Domain Admins group. If you are offline, you must provide the Directory Services Restore Mode password for the local administrator account.

**To compare the size of the directory database files to the volume size**

1. In Windows Explorer, click **My Computer**.
2. On the **View** menu, click **Details**.
3. In the **Name** column in the details pane, locate the volume. Make a note of the value in the **Total Size** column.
4. Navigate to the folder that stores the database file, the log files, or both.
5. Right-click the folder, and then click **Properties**. Make a note of the value in **Size on disk**.
6. If the volume includes **SYSVOL**, navigate to that folder and repeat step 5.
7. Compare the sizes. If the combined size of the relevant database files and SYSVOL files (if appropriate) is significantly smaller than the volume size, then check the contents of the volume for other files.
8. If other files are present, move those files and reassess the disk space on the volume.

# Back up system state

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Ntbackup.exe provides simple and advanced options for backing up Active Directory components. When you back up system state, you can choose to include or exclude system-protected boot files. System-protected boot files are not used for installations from restored backup media. When the backup file that you create is to be used for additional domain controller installations, you can clear the advanced option to back up system-protected files. Clearing this option decreases the size of the .bkf file, as well as the time required to back up, restore, and copy the system state files.

Use these procedures to back up the system state only. These procedures do not back up the system disk or any other data on the domain controller except for the system-protected files.

Use the first procedure, "To back up system state including system-protected files," for routine system state backup. Use the second procedure, "To back up system state excluding system-protected files," if you want to create a smaller backup that is effective for installing domain controllers from restored backup media.

| **Note** |
| --- |
| To back up system state, you must log on locally to the domain controller, or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

- To perform the following two procedures, you must be a member of the Domain Admins group or a member of the Backup Operators group.

#### To back up system state including system-protected files

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.

   This procedure provides steps for backing up in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard.

2. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
3. Select **Back up files and settings**, and then click **Next**.
4. Select **Let me choose what to back up**, and then click **Next**.
5. In the **Items to Back Up** window, double-click **My Computer**.
6. In the expanded list below **My Computer**, check **System State**, and then click **Next**.
7. Select a location to store the backup:

   - If you are backing up to a file, type the path and file name for the backup (.bkf) file (or click **Browse** to find a folder or file).

   - If you are backing up to a tape unit, choose the tape that you want to use.

     | **Note** |
     | --- |
     | You should not store the backup on the local hard drive. Instead, store it in a location, such as a tape drive, away from the computer that you are backing up. |

8. Type a name for this backup according to the recommendations in Backing Up Active Directory Components, and then click **Next**.
9. On the last page of the wizard, click **Advanced**.

10. Do not change the default options for **Type of Backup**. **Normal** should be selected, and the check box for **Backup migrated remote storage data** should remain cleared. Click **Next**.
11. Select **Verify data after backup**, and then click **Next**.
12. In the **Backup Options** dialog box, select a backup option, and then click **Next**.
13. If you are replacing the existing backups, select the option to allow only the owner and administrator access to the backup data and to any backups that are appended to this medium, and then click **Next**.
14. In the **When to back up** box, select the appropriate option for your needs, and then click **Next**.
15. If you are satisfied with all of the options that are selected, click **Finish** to perform the backup operation according to your selected schedule.

| **Note** |
| --- |
| The system state can also be backed up by using Ntbackup from a command line with appropriate parameters. For more information, at a command prompt type **ntbackup /?**. |

The following procedure produces a smaller .bkf file that does not include system boot files. By using this procedure, you can reduce the time that is required to perform the backup and subsequent restore, as well as the amount of disk space that is required. This method is recommended when the restored backup is to be used for installing additional domain controllers.

**To back up system state excluding system-protected files**

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.
2. On the **Welcome to the Backup or Restore Wizard** page, click **Advanced Mode**, and then click the **Backup** tab.
3. In the console tree, select the **System State** check box.
4. In **Backup media or file name**, type a name for this backup according to the recommendations in Backing Up Active Directory Components.
5. Click **Start Backup**, and then click **Advanced**.
6. Clear the **Automatically back up System Protected Files with the System State** check box, and then click **OK**.
7. Click **Start Backup**.

**See Also**
**Concepts**
Enable Remote Desktop
Create a Remote Desktop Connection

# Restart the domain controller in Directory Services Restore Mode locally

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have physical access to a domain controller, you can restart the domain controller in Directory Services Restore Mode locally. Restarting in Directory Services Restore Mode takes the domain controller offline. In this mode, the server is not functioning as a domain controller.

When you start Windows Server 2003 in Directory Services Restore Mode, the local Administrator account is authenticated by the local Security Accounts Manager (SAM) database. Therefore, logging on requires that you use the local administrator password, not an Active Directory domain password. This password is set during Active Directory installation when you provide the password for Directory Services Restore Mode.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restart the domain controller in Directory Services Restore Mode locally**

1. Restart the domain controller.
2. When the screen for selecting an operating system appears, press F8.
3. On the **Windows Advanced Options** menu, select **Directory Services Restore Mode**.
4. When you are prompted, log on as the local administrator.

**See Also**
**Concepts**
Restart the domain controller in Directory Services Restore Mode Remotely

# Restart the domain controller in Directory Services Restore Mode Remotely

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If Remote Desktop is enabled on a domain controller, you can use Remote Desktop Connection to connect to the domain controller remotely. Remote Desktop Connection (formerly known as the Terminal Services client) is installed by default on all Windows Server 2003 family operating systems.

If you use Remote Desktop Connection to connect to a domain controller remotely and you want to restart the domain controller in Directory Services Restore Mode, you must first modify the Boot.ini file on the remote server so that you do not lose the connection when the domain controller restarts.

When you start Windows Server 2003 in Directory Services Restore Mode, the local Administrator account is authenticated by the local Security Accounts Manager (SAM) database. Therefore, logging on requires that you use the local administrator password, not an Active Directory domain password. This password is set during Active Directory installation when you provide the password for Directory Services Restore Mode.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restart the domain controller in Directory Services Restore Mode remotely**

1. Use Remote Desktop Connection to connect to the domain controller that you want to restart in Directory Services Restore Mode.
2. Right-click **My Computer**, click **Properties**, and then click the **Advanced** tab.
3. Click **Settings for startup and recovery**.
4. Click the **Edit** button to edit the startup options file.
5. Modify the default entry to include the **/SAFEBOOT:DSREPAIR** switch, as shown in the following example:

   **multi(0)disk(0)rdisk(0)partition(2)\WINNT="W2K DC ** *your server name* **" /fastdetect /SAFEBOOT:DSREPAIR**

   | **Note** |
   | --- |
   | The **/SAFEBOOT:DSREPAIR** switch works for domain controllers running Windows 2000 Server and Windows Server 2003. |

6. Save the modified Boot.ini file, and then close Notepad.
7. On the **Start** menu, click **Shut Down**, and then click **Restart**. During the restart process, the Terminal Services client reports that the session is disconnected.

   | **Caution** |
   | --- |
   | Be sure to click **Restart** and not **Shut Down** at this step. If you click **Shut Down**, you cannot restart the domain controller remotely. |

8. Wait until the restart process completes on the remote domain controller, and then reconnect the client session.
9. When the client session is reconnected, log on as the local administrator.
10. Right-click **My Computer**, click **Properties**, and then click the **Advanced** tab.
11. Click **Settings for startup and recovery**.
12. Click the **Edit** button to edit the startup options file.
13. Delete the **/SAFEBOOT:DSREPAIR** switch from the default entry in the Boot.ini file, save the file, and then close Notepad.

| **Important** |
| --- |
| If you restart the domain controller before you modify the Boot.ini file, the domain controller remains offline. |

The Boot.ini file is now returned to its original state, which starts the domain controller normally.

**See Also**

**Concepts**

Enable Remote Desktop

Create a Remote Desktop Connection

Restart the domain controller in Directory Services Restore Mode locally

# Move the directory database and log files to a local drive

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Move the files to a temporary destination if you need to reformat the original location, or to a permanent location if you have additional disk space. Moving the files can be performed locally by using Ntdsutil.exe or remotely (temporarily) by using a file copy.

If you need to reformat the partition that currently stores the database file, the log files, or both, then you must move the files temporarily while you reformat the original drive. After you reformat the drive, use the same procedure to move the files back. Even if you are moving the files only temporarily, use Ntdsutil.exe so that the registry is always current.

**Administrative Credentials**

To perform this procedure, you must provide the Directory Services Restore Mode password for the local administrator account.

**To move the directory database and log files to a local drive**

1. In Directory Services Restore Mode, open a command prompt and change directories to the current location of the directory database file (Ntds.dit) or the log files, whichever you are moving.
2. Run the **dir** command and make a note of the current size and location of the Ntds.dit file.
3. At the command prompt, type **ntdsutil** and then press ENTER.
4. At the **ntdsutil:** prompt, type **files** and then press ENTER.
5. To move the database file, at the file maintenance: prompt, use the following commands:
   - To move the Ntds.dit file, type:

     **move db to** *drive:\directory*

   - To move the log files, type:

     **move logs to** *drive:\directory*

   where *Drive:\directory* specifies the path to the new location. If the directory does not exist, then Ntdsutil.exe creates it.

   | **Note** |
   | --- |
   | If the directory path contains any spaces, the entire path must be surrounded by quotation marks (for example, move db to "*g:\new folder*"). |

6. After the move completes, at the **file maintenance:** prompt, type **quit** and press ENTER. Type **quit** again and press ENTER to quit Ntdsutil.exe.
7. Change to the destination directory and then run the **dir** command to confirm the presence of the files. If you have moved the database file, then check the size of the Ntds.dit file against the file size you noted in step 2 to be sure that you are focused on the correct file.
8. If you are moving the database file or log files permanently, go to step 9.

   If you are moving the database file or log files temporarily, you can now perform any required updates to the original drive. After you update the drive, repeat steps 1 through 7 to move the files back to the original location.

   If the path to the database file or log files has not changed, go to step 10.
9. If the path to the database file or log files has changed from the original location, check permissions on the database folder or logs folder while still in Directory Services Restore Mode, as follows:
   a. In Windows Explorer, right-click the folder to which you have moved the database file or log files, and then click

**Properties**.

b. Click the **Security** tab, and verify that the permissions are:

   **Administrators** group has **Allow Full Control**.

   **System** has **Allow Full Control**.

   **Inheritable permissions are not allowed** (checkbox is cleared).

   **No Deny permissions** are selected.

c. If the permissions in step 9b are in effect, then go to step 10. If permissions other than those described in step 9b are in effect, then perform steps 9d through 9k.

d. If **Allow inheritable permissions from parent to propagate to this object** is selected, click to clear it.

e. When prompted, click **Copy** to copy previously inherited permissions to this object.

f. If **Administrators** or **SYSTEM**, or both, are not in the **Name** list, click **Add**.

g. On the **Select Users or Groups** page, in the **Look in:** box, be sure the name of the local computer is selected.

h. In the **Name** list, click **System** if needed, and then click **Add**. Repeat to add Administrators, if needed, and then click **OK**.

i. On the **Security** tab, click **System** and then in the **Allow** column, click **Full Control**. Repeat for Administrators.

j. In the **Name** box, click any name that is not SYSTEM or Administrators, and then click **Remove**. Repeat until the only remaining accounts are Administrators and SYSTEM, and then click **OK**.

| **Note** |
| --- |
| Some accounts might appear in the form of security identifiers (SIDs). Remove any such accounts. |

k. Click **OK** to close **Properties**.

10. At the command prompt, type **ntdsutil** and then press ENTER.
11. At the **ntdsutil:** prompt, type **files** and then press ENTER.
12. At the **file maintenance:** prompt, type **integrity** and then press ENTER.

    If the integrity check fails, perform semantic database analysis with a fixup record.
13. If the integrity check succeeds, type **quit** and press ENTER to quit the **file maintenance:** prompt. Type **quit** again and press ENTER to quit Ntdsutil.exe.
14. Restart the domain controller normally. If you are performing this procedure remotely over a Terminal Services connection, be sure that you have modified the Boot.ini file for normal restarting before you restart the domain controller.

    If errors appear when you restart the domain controller:

    a. Restart the domain controller in Directory Services Restore Mode.

    b. Check the errors in Event Viewer.

If the following events are logged in Event Viewer on restarting the domain controller, address the events as follows:

- Event ID 1046. “The Active Directory database engine caused an exception with the following parameters.” In this case, Active Directory cannot recover from this error and you must restore from backup media.

- Event ID 1168. “Internal error: An Active Directory error has occurred.” In this case, information is missing from the registry and you must restore from backup media.

# Copy the directory database and log files to a remote share

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When copying any database files from the local computer, always copy both the database file and the log files.

If you need to move the database file or the log files while you reconfigure the drive on which they are currently stored, and you do not have sufficient space to move the files locally, then you can use the xcopy command to copy the files to a remote shared folder temporarily, and then use the same procedure to copy them back to the original drive. You can use this method as long as the path to the files does not change.

| **Important** |
| --- |
| When relocating any database files (the database file or the log files) off the local computer, always copy both the database file and the log files so that all of the files necessary to restore the directory service are maintained. |

### Administrative Credentials

To perform this procedure, you must provide the Directory Services Restore Mode password for the local administrator account.

**To copy the directory database and log files to a remote share and back to the local computer**

1. In Directory Services Restore Mode, open a command prompt and change directories to the current location of the database file (Ntds.dit) or the log files. If the database file and log files are in different locations, perform step 2 for each directory.
2. Run the **dir** command and make a note of the current size and location of the Ntds.dit file and the log files.
3. Establish a network connection to a shared folder, as shown below. Because you are logged on as the local administrator, unless permissions on the shared folder include the built-in Administrator account, you must provide a domain name, user name, and password for an account that has Write permissions on the shared folder.

   In the example below, **\\SERVER1\NTDS** is the name of the shared folder. K: is the drive that you have mapped to the shared folder. Example text that describes information that you type is shown in bold. After typing the first line and pressing ENTER, you are prompted for the password. Type the password and then press ENTER.

   **H:\>net use K: \\SERVER1\NTDS /user:** *domainName\userName* *

   Type the password for **\\SERVER1\NTDS:**

   Drive K: is now connected to **\\SERVER1\NTDS**

   The command completed successfully.
4. Use the xcopy command to copy the database file and log files to the location you established in step 3. In the example where the database file is located in H:\WINNT\NTDS and the share has the subdirectory database, the text you type is shown in bold:

   **H:>xcopy WINNT\NTDS K:\DB**

   The command copies the contents of WINNT\NTDS to the subfolder database in the shared folder described as drive K:. If the database file and log files are in different locations, repeat the xcopy command for the log files, specifying the subfolder for the log files.
5. Change drives to the new location and run the **dir** command to compare the file sizes to those listed in step 2. Use this step to ensure that you copy the correct set of files back to the local computer.
6. At this point, you can safely destroy data on the original local drive.
7. After the destination drive is prepared, re-establish a connection to the network drive as described in step 3, if necessary.
8. Copy the database and log files from the remote shared folder back to the original location on the domain controller.
9. At the command prompt, type **ntdsutil** and then press ENTER.

10. At the **ntdsutil:** prompt, type **files** and then press ENTER.
11. At the **file maintenance:** prompt, type **integrity** and then press ENTER.
12. If the integrity check fails, perform semantic database analysis with a fixup record.
13. If the integrity check succeeds, type **quit** and press ENTER to quit the **file maintenance:** prompt. Type **quit** again and press ENTER to quit Ntdsutil.exe.
14. Restart the domain controller normally. If you are performing this procedure remotely over a Terminal Services connection, be sure that you have modified the Boot.ini file for normal restarting before you restart the domain controller.

If errors appear when you restart the domain controller:

1. Restart the domain controller in Directory Services Restore Mode.

2. Check the errors in Event Viewer.

   If the following events are logged in Event Viewer on restarting the domain controller, respond to the events as follows:

   - Event ID 1046. "The Active Directory database engine caused an exception with the following parameters." In this case, Active Directory cannot recover from this error and you must restore from backup media.

   - Event ID 1168. "Internal error: An Active Directory error has occurred." In this case, information is missing from the registry and you must restore from backup media.

# Returning Unused Disk Space from the Active Directory Database to the File System

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

During ordinary operation, the white space in the Active Directory database file becomes fragmented. Each time garbage collection runs (every 12 hours, by default), white space is automatically defragmented online to optimize its use within the database file. The unused disk space is thereby maintained for the database; it is not returned to the file system.

Only offline defragmentation can return unused disk space from the directory database to the file system. When database contents have decreased considerably through a bulk deletion (for example, you remove the global catalog from a domain controller), or if the size of the database backup is significantly increased due to the white space, use offline defragmentation to reduce the size of the Ntds.dit file.

You can determine how much free disk space is recoverable from the Ntds.dit file by setting the garbage collection logging level in the registry. Changing the garbage collection logging level from the default value of **0** to a value of **1** results in event ID 1646 being logged in the directory service log. This event describes the total amount of disk space used by the database file as well as the amount of free disk space that is recoverable from the Ntds.dit file through offline defragmentation.

At garbage collection logging level 0, only critical events and error events are logged in the directory service log. At level 1, high-level events are logged as well. Events can include one message for each major task that is performed by the service. At level 1, the following events are logged for garbage collection:

- Event IDs 700 and 701: report when online defragmentation begins and ends, respectively.
- Event ID 1646: reports the amount of free space available in the database out of the amount of allocated space.

| **Caution** |
| --- |
| Setting the value of entries in the Diagnostics subkey to greater than 3 can degrade server performance and is not recommended. |

Following offline defragmentation, perform a database integrity check. The integrity command in Ntdsutil.exe detects binary-level database corruption by reading every byte in the database file. The process ensures that the correct headers exist in the database itself and that all of the tables are functioning and consistent. Therefore, depending upon the size of your Ntds.dit file and the domain controller hardware, the process might take considerable time. In testing environments, the speed of 2 GB per hour is considered to be typical. When you run the command, an online graph displays the percentage completed.

**Task requirements**

The following tools are required to perform the procedures for this task:

- Regedit.exe
- Backup software
- Ntdsutil.exe

To complete this task, perform the following procedures:

1. Change the garbage collection logging level to 1
2. Back up system state
3. Use one of the following procedures:
    - Restart the domain controller in Directory Services Restore Mode locally

If you are logged on to the domain controller locally, restart the domain controller in Directory Services Restore Mode.

- Restart the domain controller in Directory Services Restore Mode Remotely

  If you are using Remote Desktop Connection for remote administration, you can restart the domain controller remotely in Directory Services Restore Mode after modifying the Boot.ini file on the remote server.

4. Compact the directory database file (offline defragmentation)

   As part of the offline defragmentation procedure, check directory database integrity.

5. If database integrity check fails, perform semantic database analysis with fixup

# Change the garbage collection logging level to 1

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Check the directory service event log for event ID 1646, which reports the amount of disk space that you can recover by performing offline defragmentation.

The garbage collection logging level is an NTDS diagnostics setting in the registry.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group.

| ⚠**Caution** |
| --- |
| The Registry Editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back up system state first. For information about backing up system state, see Introduction to Administering Active Directory Backup and Restore. |

**To change the garbage collection logging level**

1. Click **Start**, click **Run**, type **regedit** and then press ENTER.
2. In Registry Editor, navigate to the **Garbage Collection** entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics**
3. Double-click **Garbage Collection**, and for the **Base**, click **Decimal**.
4. In the **Value data** box, type **1**, and then click **OK**.

# Back up system state

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Ntbackup.exe provides simple and advanced options for backing up Active Directory components. When you back up system state, you can choose to include or exclude system-protected boot files. System-protected boot files are not used for installations from restored backup media. When the backup file that you create is to be used for additional domain controller installations, you can clear the advanced option to back up system-protected files. Clearing this option decreases the size of the .bkf file, as well as the time required to back up, restore, and copy the system state files.

Use these procedures to back up the system state only. These procedures do not back up the system disk or any other data on the domain controller except for the system-protected files.

Use the first procedure, "To back up system state including system-protected files," for routine system state backup. Use the second procedure, "To back up system state excluding system-protected files," if you want to create a smaller backup that is effective for installing domain controllers from restored backup media.

| **Note** |
| --- |
| To back up system state, you must log on locally to the domain controller, or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

- To perform the following two procedures, you must be a member of the Domain Admins group or a member of the Backup Operators group.

#### To back up system state including system-protected files

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.

   This procedure provides steps for backing up in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard.

2. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
3. Select **Back up files and settings**, and then click **Next**.
4. Select **Let me choose what to back up**, and then click **Next**.
5. In the **Items to Back Up** window, double-click **My Computer**.
6. In the expanded list below **My Computer**, check **System State**, and then click **Next**.
7. Select a location to store the backup:

   - If you are backing up to a file, type the path and file name for the backup (.bkf) file (or click **Browse** to find a folder or file).

   - If you are backing up to a tape unit, choose the tape that you want to use.

     | **Note** |
     | --- |
     | You should not store the backup on the local hard drive. Instead, store it in a location, such as a tape drive, away from the computer that you are backing up. |

8. Type a name for this backup according to the recommendations in Backing Up Active Directory Components, and then click **Next**.
9. On the last page of the wizard, click **Advanced**.

10. Do not change the default options for **Type of Backup**. **Normal** should be selected, and the check box for **Backup migrated remote storage data** should remain cleared. Click **Next**.
11. Select **Verify data after backup**, and then click **Next**.
12. In the **Backup Options** dialog box, select a backup option, and then click **Next**.
13. If you are replacing the existing backups, select the option to allow only the owner and administrator access to the backup data and to any backups that are appended to this medium, and then click **Next**.
14. In the **When to back up** box, select the appropriate option for your needs, and then click **Next**.
15. If you are satisfied with all of the options that are selected, click **Finish** to perform the backup operation according to your selected schedule.

| **Note** |
| --- |
| The system state can also be backed up by using Ntbackup from a command line with appropriate parameters. For more information, at a command prompt type **ntbackup /?**. |

The following procedure produces a smaller .bkf file that does not include system boot files. By using this procedure, you can reduce the time that is required to perform the backup and subsequent restore, as well as the amount of disk space that is required. This method is recommended when the restored backup is to be used for installing additional domain controllers.

**To back up system state excluding system-protected files**

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.
2. On the **Welcome to the Backup or Restore Wizard** page, click **Advanced Mode**, and then click the **Backup** tab.
3. In the console tree, select the **System State** check box.
4. In **Backup media or file name**, type a name for this backup according to the recommendations in Backing Up Active Directory Components.
5. Click **Start Backup**, and then click **Advanced**.
6. Clear the **Automatically back up System Protected Files with the System State** check box, and then click **OK**.
7. Click **Start Backup**.

**See Also**
**Concepts**
Enable Remote Desktop
Create a Remote Desktop Connection

# Restart the domain controller in Directory Services Restore Mode locally

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you have physical access to a domain controller, you can restart the domain controller in Directory Services Restore Mode locally. Restarting in Directory Services Restore Mode takes the domain controller offline. In this mode, the server is not functioning as a domain controller.

When you start Windows Server 2003 in Directory Services Restore Mode, the local Administrator account is authenticated by the local Security Accounts Manager (SAM) database. Therefore, logging on requires that you use the local administrator password, not an Active Directory domain password. This password is set during Active Directory installation when you provide the password for Directory Services Restore Mode.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restart the domain controller in Directory Services Restore Mode locally**

1. Restart the domain controller.
2. When the screen for selecting an operating system appears, press F8.
3. On the **Windows Advanced Options** menu, select **Directory Services Restore Mode**.
4. When you are prompted, log on as the local administrator.

**See Also**
**Concepts**
Restart the domain controller in Directory Services Restore Mode Remotely

# Restart the domain controller in Directory Services Restore Mode Remotely

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If Remote Desktop is enabled on a domain controller, you can use Remote Desktop Connection to connect to the domain controller remotely. Remote Desktop Connection (formerly known as the Terminal Services client) is installed by default on all Windows Server 2003 family operating systems.

If you use Remote Desktop Connection to connect to a domain controller remotely and you want to restart the domain controller in Directory Services Restore Mode, you must first modify the Boot.ini file on the remote server so that you do not lose the connection when the domain controller restarts.

When you start Windows Server 2003 in Directory Services Restore Mode, the local Administrator account is authenticated by the local Security Accounts Manager (SAM) database. Therefore, logging on requires that you use the local administrator password, not an Active Directory domain password. This password is set during Active Directory installation when you provide the password for Directory Services Restore Mode.

**Administrative credentials**

To perform this procedure, you must provide the Administrator password for Directory Services Restore Mode.

**To restart the domain controller in Directory Services Restore Mode remotely**

1. Use Remote Desktop Connection to connect to the domain controller that you want to restart in Directory Services Restore Mode.
2. Right-click **My Computer**, click **Properties**, and then click the **Advanced** tab.
3. Click **Settings for startup and recovery**.
4. Click the **Edit** button to edit the startup options file.
5. Modify the default entry to include the **/SAFEBOOT:DSREPAIR** switch, as shown in the following example:

   **multi(0)disk(0)rdisk(0)partition(2)\WINNT="W2K DC ** *your server name* **" /fastdetect /SAFEBOOT:DSREPAIR**

   | **Note** |
   | --- |
   | The **/SAFEBOOT:DSREPAIR** switch works for domain controllers running Windows 2000 Server and Windows Server 2003. |

6. Save the modified Boot.ini file, and then close Notepad.
7. On the **Start** menu, click **Shut Down**, and then click **Restart**. During the restart process, the Terminal Services client reports that the session is disconnected.

   | **Caution** |
   | --- |
   | Be sure to click **Restart** and not **Shut Down** at this step. If you click **Shut Down**, you cannot restart the domain controller remotely. |

8. Wait until the restart process completes on the remote domain controller, and then reconnect the client session.
9. When the client session is reconnected, log on as the local administrator.
10. Right-click **My Computer**, click **Properties**, and then click the **Advanced** tab.
11. Click **Settings for startup and recovery**.
12. Click the **Edit** button to edit the startup options file.
13. Delete the **/SAFEBOOT:DSREPAIR** switch from the default entry in the Boot.ini file, save the file, and then close Notepad.

| **Important** |
| --- |
| If you restart the domain controller before you modify the Boot.ini file, the domain controller remains offline. |

The Boot.ini file is now returned to its original state, which starts the domain controller normally.

**See Also**

**Concepts**

Enable Remote Desktop

Create a Remote Desktop Connection

Restart the domain controller in Directory Services Restore Mode locally

# Compact the directory database file (offline defragmentation)

Updated: September 16, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| **Note** |
| --- |
| This procedure is for Windows Server 2003. For the Windows Server 2008 version of this procedure, see Compact the Directory Database File (Offline Defragmentation) (http://go.microsoft.com/fwlink/?LinkID=93225). |

As part of the offline defragmentation procedure, check directory database integrity.

Performing offline defragmentation creates a new, compacted version of the database file in a different location. This location can be either on the same computer or a network-mapped drive. However, to avoid potential problems related to network issues, perform this procedure using a local mass storage device. You can use locally attached external mass storage devices, such as Universal Serial Bus (USB), IEEE 1394, and Serial Advanced Technology Attachment (SATA), to provide additional disk space for defragmentation of the database.

After you compact the file to the temporary location, copy the compacted Ntds.dit file back to the original location. If possible, maintain a copy of the original database file that you have either renamed in its current location or copied to an archival location.

| **Note** |
| --- |
| To perform this procedure, the domain controller must be started in Directory Services Restore Mode (DSRM). |

### Administrative Credentials

To perform this procedure, you must provide the Directory Services Restore Mode password for the local administrator account. At the remote location, you must have Read and Write permissions on the destination drive and the shared folder.

### Disk Space

- **Current database drive**. Free space (on the drive that contains the Active Directory database file) equivalent to at least 15 percent of the current size of the database (Ntds.dit) for temporary storage during the index rebuild process.

- **Destination database drive**. Free space equivalent to at least the current size of the database for storage of the compacted database file.

| **Note** |
| --- |
| These disk space requirements mean that if you compress the Active Directory database on a single drive, you should have free space equivalent to at least 115 percent of the space that the current Active Directory database uses on that drive. |

#### To perform offline defragmentation of the directory database

1. In DSRM, compact the database file to a local directory or remote shared folder, as follows:

    - **Local directory:** Go to step 2.

    - **Remote directory:** If you are compacting the database file to a shared folder on a remote computer, establish a network connection to the shared folder as shown below. Because you are logged on as the local administrator, unless permissions on the shared folder include the built-in Administrator account, you must provide a domain name, user name, and password for a domain account that has Write permissions on the shared folder. In the example below, \\SERVER1\NTDS is the name of the shared folder, and K: is the drive that you are mapping to the shared folder. After typing the first line and pressing ENTER, you are prompted for the password. Type the password and then press ENTER.

        **H:\>net use K: \\SERVER1\NTDS /user:***domainName******userName* ***

Type the password for **\\SERVER1\NTDS:**

Drive K: is now connected to **\\SERVER1\NTDS**

The command completed successfully.

2. Type the following command at a command prompt, and then press ENTER:

   **ntdsutil**

3. At the **ntdsutil:** prompt, type **files**, and then press ENTER.
4. At the **file maintenance:** prompt, type **compact to** *drive:\ LocalDirectoryPath* (where *drive:\ LocalDirectoryPath* is the path to a location on the local computer), and then press ENTER.

   If you have mapped a drive to a shared folder on a remote computer, type the drive letter only (for example, **compact to K:**).

   | **Note** |
   | --- |
   | When compacting to a local drive, you must provide a path. If the path contains any spaces, enclose the entire path in quotation marks (for example, compact to "c:\new folder"). If the directory does not exist, Ntdsutil.exe creates it and creates the file named Ntds.dit in that location. |

5. If defragmentation completes successfully, type **quit**, and then press ENTER to quit the **file maintenance:** prompt. Type **quit** again, and then press ENTER to quit Ntdsutil.exe. Go to step 6. If defragmentation completes with errors, go to step 9.

   | **Caution** |
   | --- |
   | Do not overwrite the original Ntds.dit file or delete any log files. |

6. If defragmentation succeeds with no errors, follow the Ntdsutil.exe onscreen instructions to:

   a. Delete all of the log files in the log directory by typing:

      **del** *drive***:***pathToLogFiles****.log**

      | **Note** |
      | --- |
      | You do not have to delete the Edb.chk file. |

   b. You should make a copy of the existing Ntds.dit file if at all possible, even if you have to store that copy to a secured network drive. If the compaction of the database does not work properly, you can then easily restore the database by copying it back to the original location. Do not delete the copy of the Ntds.dit file until you have at least verified that the domain controller starts properly. If space allows, you can rename the original Ntds.dit file to preserve it or else copy it to a different location. Avoid overwriting the original Ntds.dit file.

   c. Manually copy the compacted database file to the original location, as follows:

      **copy** *temporaryDrive***:\ntds.dit** *originalDrive***:***pathToOriginalDatabaseFile***\ntds.dit**

7. Type **ntdsutil** and then press ENTER.
8. At the **ntdsutil:** prompt, type **files** and then press ENTER.
9. At the **file maintenance:** prompt, type **integrity** and then press ENTER.

   If the integrity check fails, the likely cause is that an error occurred during the copy operation in step 6.c. Repeat steps 6.c through step 9. If the integrity check fails again:

   - Contact Microsoft Product Support Services.

-or-

- Copy the original version of the Ntds.dit file that you preserved in step 6.b. to the original database location and repeat the offline defragmentation procedure.

10. If the integrity check succeeds, proceed as follows:

    - If the initial compact to command failed, go back to step 4 and perform steps 4 through 9.

    - If the initial compact to command succeeded, type **quit** and press ENTER to quit the **file maintenance:** prompt, and then type **quit** and press ENTER again to quit Ntdsutil.exe.

11. Restart the domain controller normally. If you are connected remotely through a Terminal Services session, be sure that you have modified the Boot.ini file for normal restarting before you restart the domain controller.

If errors appear when you restart the domain controller:

1. Restart the domain controller in Directory Services Restore Mode.

2. Check the errors in Event Viewer.

   If the following events are logged in Event Viewer on restarting the domain controller, respond to the events as follows:

   - Event ID 1046. "The Active Directory database engine caused an exception with the following parameters." In this case, Active Directory cannot recover from this error and you must restore from backup media.

   - Event ID 1168. "Internal error: An Active Directory error has occurred." In this case, information is missing from the registry and you must restore from backup media.

3. Check database integrity and then proceed as follows:

   If the integrity check fails, try repeating step 6.c through step 9 above, and then repeat the integrity check. If the integrity check fails again:

   - Contact Microsoft Product Support Services.

     -or-

   - Copy the original version of the Ntds.dit file that you preserved in step 6.2. to the original database location and repeat the offline defragmentation procedure.

     If the integrity check succeeds, perform semantic database analysis with fixup.

4. If semantic database analysis with fixup succeeds, quit Ntdsutil.exe and restart the domain controller normally.

If semantic database analysis with fixup fails, contact Microsoft Product Support Services.

# If database integrity check fails, perform semantic database analysis with fixup

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you run semantic database analysis with the Go Fixup command instead of the Go command, errors are written into Dsdit.dmp.xx log files. A progress indicator reports the status of the check.

| **Note** |
| --- |
| To perform this procedure, the domain controller must be started in Directory Services Restore Mode. |

### Administrative Credentials

To perform this procedure, you must provide the Directory Services Restore Mode password for the local administrator account.

**To perform semantic database analysis with fixup**

1. In Directory Services Restore Mode, open a Command Prompt.
2. Type the following command and then press ENTER:

   **ntdsutil:**
3. At the **ntdsutil:** prompt, type **semantic database analysis** and then press ENTER.
4. At the **semantic checker:** prompt, type **verbose on** and then press ENTER.
5. At the **semantic checker:** prompt, type **go fixup** and then press ENTER.

   - If errors are reported during the semantic database analysis Go Fixup phase, perform directory database recovery.

     | **Caution** |
     | --- |
     | Do not confuse the recover command with the repair command. Never use the repair command in Ntdsutil.exe. Forest-wide data loss can occur. |

   - If semantic database analysis with fixup succeeds, type **quit** and then type **quit** again to close Ntdsutil.exe, and then restart the domain controller normally. If you are performing this procedure remotely over a Terminal Services connection, be sure that you have modified the Boot.ini file for normal restarting before you restart the domain controller.

# Guarding Against Accidental Bulk Deletions in Active Directory

Updated: August 25, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

Although bulk deletions are rare, they are disruptive events that you can guard against by removing the Delete and the Delete Subtree permissions in Active Directory. To guard against accidental deletions, you should remove the Delete and Delete Subtree permissions on organizational units (OUs) that contain user accounts, computer accounts, and security groups in Active Directory. You should also remove the Delete All Child Objects permission on the parent container of an OU that you want to protect.

**Types of deletions**

IT administrators have used Active Directory Users and Computers (dsa.msc) and other Active Directory management tools to accidentally delete objects from Active Directory. Two types of deletions are most common:

- Leaf-node deletions, which occur when an administrator or delegated user selects and deletes one or more leaf objects. A leaf object is an object that does not have subordinate objects. User and group objects are examples of leaf objects.

- OU deletion, which occurs when an administrator or delegated user selects and deletes a container that has subordinate objects.

| **Note** |
| --- |
| Although the remainder of this task refers to using dsa.msc, the explanations and guidelines apply to any tool that you use to perform the deletion, including the Active Directory Sites and Services snap-in (dssite.msc) if you use it to perform bulk deletions of configuration partition data. |

A leaf-node deletion has the following characteristics:

- The domain controller that is targeted for the operation checks whether the user has delete permissions on each selected object as dsa.msc tries to delete it.

- Dsa.msc displays a progress bar that is updated as each object is deleted.

- Because of the round-trip communication that takes place between the client and server, you can cancel or close dsa.msc from the workstation session that initiated the operation to stop the deletion while it is still in progress.

An OU deletion has the following characteristics:

- When you delete an OU, dsa.msc checks for subordinate objects and whether the subtree contains any critical objects. A critical object is an object for which the **isSystemCriticalObject** attribute is set to **True**. If critical objects are in the subtree, dsa.msc stops the operation. If there are no critical objects, dsa.msc displays a warning about deleting an object. You must click **Yes** to continue. Dsa.msc then displays a second warning about deleting child objects and containers. You must click **Yes** again to continue.

- The domain controller that is targeted for the operation performs a single-access check on the top-most OU that you are deleting. If the user has the Delete Subtree advanced permission (the symbolic name for this permission is **RIGHT_DS_DELETE_TREE**), the top OU and all subordinate containers are deleted, regardless of what access control lists (ACLs) are defined on subordinate containers. Administrators are implicitly granted the Delete and Delete Subtree advanced permissions if they have Full Control on a container. You can also explicitly grant the Delete and Delete Subtree advanced permissions to any user.

- The domain controller continues its server-side deletion and does not communicate with dsa.msc again until it reaches the maximum threshold of object deletions, which is 16,384 objects. The deletions are not performed in a particular order.

When dsa.msc reaches the maximum threshold of object deletions, it reissues the delete tree operation request until the operation succeeds or some other error is returned. When you cancel dsa.msc or stop it, either by logging off the session or by rebooting the administrative workstation, the deletion ends after the current batch of 16,384 objects is deleted on the server.

You can stop the server-side deletion only by restarting the domain controller that is targeted for the operation or by removing network connectivity from that domain controller if additional domain controllers exist in the domain.

- Dsa.msc does not display a progress bar during OU deletions.

- Because no round-trip communication occurs between the client and server, the OU deletion operation completes more quickly than leaf-node deletion.

For any OU that you want to protect from accidental deletion, you can add explicit Deny Delete and Deny Delete Subtree advanced permissions to the Everyone group on that container object. To further protect the OU, you can add the explicit Deny Delete All Child Objects permission on the parent container of the OU that you want to protect. The explicit Deny access control entries (ACEs) take precedence over any Allow permissions that the user might have on the container, including inherited permissions.

| **Note** |
| --- |
| To prevent accidental deletion in Windows Server 2008, you can select the **Protect object against accidental deletion** check box when you create a new object or you can select this check box on the **Object** tab of the **Properties** property sheet for an existing object. This feature is available for all objects that appear either in Active Directory Users and Computers or Active Directory Sites and Services. |

**Task requirements**

The following tool is required if you want to perform the procedures for this task by using the Windows interface. For more information about performing this task by using Dsacls.exe at the command line or by using a script, see Script to Protect Organizational Units (OUs) from Accidental Deletion (http://go.microsoft.com/fwlink/?LinkId=162623).

- Active Directory Users and Computers (dsa.msc).

To complete this task by using the Windows interface, perform the following procedures as needed:

- Protect an Organizational Unit from Accidental Deletion

- Remove Protection Against Accidental Organizational Unit Deletion

# Protect an Organizational Unit from Accidental Deletion

Updated: January 7, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to add the following access control entries (ACEs):

- On the organizational unit (OU) that you want to protect, add explicit Deny ACEs for the Delete and Delete Subtree advanced permissions for the Everyone group.
- On the parent container of the OU that you want to protect, add an explicit Deny ACE for the Delete All Child Objects permission for the Everyone group.

This protects an OU from accidental deletion. When a user tries to delete the protected object, the operation returns an error that indicates access is denied.

To remove the protection, remove the Deny ACEs that you added for the Everyone group. For more information, see Remove Protection Against Accidental Organizational Unit Deletion.

Although bulk deletions are not common in the configuration partition data that you can view with the Active Directory Sites and Services snap-in, you can apply the same method to protect objects in the configuration partition by using adsiedit.msc.

Membership in the Domain Admins group, or equivalent, is required to complete this procedure.

**To protect an OU from accidental bulk deletion**

1. Log on to the computer as a member of the Domain Admins group.
2. Open Active Directory Users and Computers.
3. Click **View**, and then click **Advanced Features**.
4. First, apply permissions on the OU that you want to protect. To do this, right-click the OU that you want to protect, and then click **Properties**.
5. In *OU* **Properties**, click the **Security** tab, and then click **Advanced**.
6. On **Advanced Security Settings**, click **Add**, type **Everyone**, and then click **OK**.
7. In **Permission Entry**, in **Permissions**, select the **Deny** check boxes for **Delete** and **Delete subtree**.
8. Select the check box for **Apply these permissions to objects and/or containers within this container only**.
9. Click **OK** to close **Permission Entry**.
10. On **Advanced Security Settings**, click **Apply**.
11. Review the Windows Security warning, and then click **Yes** to continue.
12. Click **OK** to close the **Advanced Security Settings**, and then click **OK** to close *OU* **Properties**.
13. Second, apply permissions to the parent container of the OU that you want to protect. To do this, right-click the parent container, and then click **Properties**.
14. In *Container***Properties**, click the **Security** tab.
15. Click **Add**, type **Everyone**, and then click **OK**.
16. In **Permissions for Everyone**, select the **Deny** check box for **Delete All Child Objects**, and then click **Apply**.
17. Review the Windows Security warning, and then click **Yes** to continue.
18. Click **OK** to close *Container* **Properties**.

# Remove Protection Against Accidental Organizational Unit Deletion

Updated: January 7, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to remove the following access control entries (ACEs):

- On the organizational unit (OU) for which you are removing protection, remove explicit Deny ACEs for the Delete and Delete Subtree advanced permissions for the Everyone group.

- On the parent container of the OU for which you are removing protection, remove an explicit Deny ACE for the Delete All Child Objects permission for the Everyone group.

This removes protection that prevents an OU from accidental deletion. Consequently, you may have to remove these permissions to be able to perform bulk deletions of objects in Active Directory.

To add the protection back again, you can reselect the **Deny** check boxes for the Delete and Delete Subtree advanced permissions for the Everyone group for the OU and the Delete All Child Objects permission for the Everyone group on the parent container. For more information, see Protect an Organizational Unit from Accidental Deletion.

Membership in the Domain Admins group, or equivalent, is required to complete this procedure.

**To remove protection that prevents an OU from accidental deletion**

1. Log on to the computer as a member of the Domain Admins group.
2. Open Active Directory Users and Computers.
3. Click **View**, and then click **Advanced Features**.
4. First, clear permissions on the OU for which you want to remove protection. To do this, right-click the OU, and then click **Properties**.
5. In *OU* **Properties**, click the **Security** tab, and then click **Advanced**.
6. In **Permission Entries**, select the **Deny** entry for the **Everyone** group, and then click **Remove**.
7. Click **OK** to close the **Advanced Security Settings**, and then click **OK** to close *OU* **Properties**.
8. Second, clear permissions on the parent container of the OU for which you want to remove protection. To do this, right-click the parent container, and then click **Properties**.
9. In *Container***Properties**, click the **Security** tab.
10. In Group or user names, select the **Everyone** group, and then clear the **Deny** check box for **Delete All Child Objects**, and then click **OK** to close *Container* **Properties**.

# Administering Domain Controllers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides information for administering Active Directory domain controllers in the Microsoft Windows Server 2003 operating system.

**In this guide**


- Introduction to Administering Domain Controllers
- Managing Domain Controllers



**Acknowledgements**

Published: March 2005

Applies to: Windows Server 2003

Produced by: Microsoft Windows Server User Assistance team

Writer: Shala Brandolini

Editor: Jim Becker

# Introduction to Administering Domain Controllers

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although installed domain controllers require little management, your overall operations environment might require change-related tasks such as adding or removing domain controllers, including managing the preparation and shipment of domain controllers to remote sites. During your day-to-day operations, you might need to do some or all of the following:

- Install and remove Active Directory
- Rename domain controllers
- Add domain controllers to remote sites

**Installing and Removing Active Directory**

To create a new domain controller, install Active Directory on a computer that is running Windows Server 2003 or Windows Server 2003 with Service Pack 1 (SP1). Installing domain controllers to create a forest and new domains is a deployment task that you perform when you initially deploy your forest, and it is beyond the scope of this guide. However, as your forest grows, you might need to add more domain controllers to existing domains.

There are several reasons for adding a new domain controller. Additional applications (which are Active Directory–integrated as opposed to running on domain controllers) might be required to meet increased capacity requirements, provide upgrades and fault tolerance, and reduce failures. You might add a new site where users require a domain controller for logging on to the domain. For more information about criteria and best practices for deploying domain controllers, see Designing and Deploying Directory and Security Services on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45801).

When a domain controller is no longer needed, remove Active Directory. The process of removing Active Directory involves steps similar to the steps for installation. You run many of the same tests before you remove the directory as you ran before you installed the directory. These tests ensure that the process occurs without any problems. In the event that a domain controller suffers a hardware failure and you plan to never return it to service, you must use a procedure that forces Active Directory removal and then take additional steps to remove the server object and its metadata from the directory.

**Renaming Domain Controllers**

You often need to rename a domain controller for organizational or administrative reasons or when the computer hardware must be replaced. Renaming a domain controller requires that Domain Name System (DNS) resource records be updated with the new Internet Protocol (IP)-to-host name mappings and that service principal names (SPNs) replicate to all domain controllers in the domain. You must also update File Replication service (FRS) objects.

**Adding Domain Controllers to Remote Sites**

If enough directory users are employed in a remote site, especially in a site that has slow connectivity to the hub site, you might need to add a domain controller to the site to provide directory access for logons and searches. Specifically, you can either install a domain controller in the hub site and ship it to the site or install the domain controller in the remote site. When you install the domain controller in the remote site, Active Directory must be sourced in one of two ways:

- By Active Directory replication over the wide area network (WAN) link
- Directly from restored backup media

Assuming that the remote site is connected to a hub site by a WAN link and does not contain a domain controller for the domain, you might want to avoid the additional time and the performance impact of replicating the full replica of Active Directory over the WAN when you add a new domain controller to the remote site. In this case, you can use backup media to install Active Directory.

If you want to install a domain controller from backup media, both the source of the backup and the target server that is to be promoted to a domain controller must be running Windows Server 2003 or Windows Server 2003 with SP1, and the operating system of the source of the backup and the target server must be the same. The hardware platform (32-bit or 64-bit) of the two computers must also match. Restoring from backup media eliminates the need to use replication to create the Active Directory replica on the new domain controller.

# Managing Domain Controllers

Updated: April 15, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Managing domain controllers involves the following tasks:

- Create an additional domain controller in an existing domain. This task involves preparation steps of gathering information and configuring the TCP/IP and Domain Name System (DNS) client settings. You can use the following methods to install Active Directory to create an additional domain controller in an existing domain:
    - Run the Active Directory Installation Wizard, and use Active Directory replication to create the Active Directory replica and File Replication Service (FRS) replication to create the System Volume (Sysvol) replicas.
    - Run the Active Directory Installation Wizard, and use restored system state backup media to create the Active Directory and Sysvol replicas.
    - Create an answer file and use the Unattend.txt file to provide the information that the Active Directory Installation Wizard requires.
- Perform tests to verify that Active Directory is properly installed and the domain controller is functioning.
- Add domain controllers to remote sites. When you prepare and ship an additional domain controller to a remote site, you can either install the domain controller before shipping or install the domain controller in the remote site.
    - When you install a domain controller in a hub site or staging site before shipment, you must disconnect the domain controller for a period, which requires careful preparation. When you reconnect the domain controller, Active Directory replication brings the domain controller up to date.
    - When you install the domain controller in the remote site, you can use a restored system state backup to avoid having to replicate Active Directory over a wide area network (WAN) link.
- Remove Active Directory from (decommission) a properly functioning domain controller. This task includes first removing operations master roles (also known as flexible single-master operation (FSMO) roles) and the global catalog, if necessary.
- Force the removal of a nonfunctioning domain controller from a domain. If a domain controller is not functioning properly on the network, the Active Directory Installation Wizard cannot contact other domain controllers and DNS servers that are required for Active Directory removal. In this case, a special version of the wizard can be invoked to forcefully remove objects that represent the server as a domain controller from Active Directory.
- Rename a domain controller. You can now rename a domain controller without removing Active Directory. New functionality is available in the Netdom tool when the domain functional level is Windows Server 2003. This new functionality provides better preparation for DNS and service recognition of the new domain controller name. You can also use System Properties, which does not require a domain functional level and does not provide the same preparation, but which relies solely on replication to update the domain controller DNS name and service principal name (SPN). This method can result in a longer delay before clients can use the renamed domain controller.

In addition, to protect domain controllers from infection by viruses that can corrupt directory data or cause software or hardware failure, an integral step in installing any domain controller is to install antivirus software.

**Managing Antivirus Software on Domain Controllers**

Because domain controllers provide critical services to their clients, it is crucial to minimize the risk of disruption of these services caused by malicious code.

Antivirus software is the generally accepted way to mitigate the risk of such malevolent activity. However, one cannot simply install the antivirus software (from any vendor) on a domain controller and tell it to scan everything. Instead, it must be

installed in a manner that mitigates the risk to the highest possible level while not interfering with the performance of the domain controllers in performing their directory service duties.

Installing effective antivirus software on domain controllers minimizes the risk that their activities will be disrupted by malicious code.

**Guidelines for Managing Antivirus Software on Domain Controllers**
Follow the guidelines established by your antivirus software vendor.

| **Note** |
| --- |
| Verify that the antivirus software you are adding is confirmed to work on domain controllers. |

The following recommendations are general and should not be construed as more important than the specific antivirus software vendor's own recommendations. These guidelines must be followed for correct Active Directory and FRS operation.

| **Note** |
| --- |
| Test the chosen antivirus software solution thoroughly in a lab environment to ensure that the software does not compromise the stability of the system. |

- Antivirus software must be installed on all domain controllers in the enterprise. Ideally, such software should also be installed on all other server and client systems that have to interact with the domain controllers. Catching the virus at the earliest point, at the firewall, or the client system on which the virus is first introduced is best—that will prevent the virus from ever reaching the infrastructure systems upon which all clients depend.

- Use a version of antivirus software that is confirmed to work with Active Directory and uses the correct APIs for accessing files on the server. Older versions of most vendors' software inappropriately modified file metadata as it was scanned, causing the FRS replication engine to think the file was changed and to schedule it for replication. Newer versions prevent this problem. For more information about antivirus software versions and FRS, see article 815263, "Antivirus, backup, and disk optimization programs that are compatible with the File Replication Service" in the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=4441), and see the vendor-specific sites for compliant versions.

- Prevent the use of domain controller systems as general workstations. Users should not be using a domain controller to surf the Web or perform any other activities that could allow the introduction of malicious code.

- When possible, do not use the domain controller as a file sharing server. Virus scanning software must be run against all files in those shares and could place an unsatisfactory load on the processor and memory resources of the server.

**Known Issues with McAfee VirusScan 8.0 Antivirus Software**
McAfee VirusScan version 8.0 is known to cause problems when it is installed on a domain controller. Patching levels for McAfee VirusScan might also affect domain controller performance. If you have installed McAfee VirusScan version 8.0 on a domain controller, make the following changes:

1. Disable the **Buffer Overflow Protection** feature in the **VirusScan Console**. Buffer Overflow Protection is intended to be used on workstations only and might cause problems if used on a server role computer with a high IO load.

2. Disable the following driver:

   File name and location: C:\Windows\System32\Drivers\mfetdik.sys

   Company name: McAfee, Inc.

   File description: Anti-Virus Mini-Firewall Driver

   The driver is known to cause system crashes. To reduce complexity and increase system reliability, perform the following steps to disable the Anti-Virus Mini-Firewall Driver MVSDTi5x.SYS on servers:

a. In Registry Editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\naiavtdi1.

b. Set **Start** to a value of **4**, which is disabled.

**Files Not at Risk of Infection**

Exclude the following files and folders from being scanned. These files are not at risk of infection and including them could cause serious performance problems due to file locking and excessive replication between domain controllers. Furthermore, they may cause Active Directory and FRS to work improperly, causing Active Directory or FRS data loss. Where a specific set of files is identified by name, exclude only those files rather than the entire folder. In some cases, the entire folder must be excluded.

Do not exclude any of these based on the file name extension (that is, do not exclude all files with a .dit extension). Microsoft has no control over other files that might choose to use the same extension as those shown here. AV software must not modify any data files in the logs, database, and/or DSA working directories specified below.

- Main NTDS database files. The location of these files is specified in:

  **HKEY_LOCAL_MACHINE \System\CurrentControlSet\Services\NTDS\Parameters\DSA Database File**

  Default location is %windir%\ntds

  The file to exclude is: **NTDS.dit**

- Active Directory transaction log files. The log directory on any given server is specified in:

  **HKEY_LOCAL_MACHINE \System\CurrentControlSet\Services\NTDS\Parameters\Database Log Files Path**

  Default location is %windir%\ntds.

  The specific files to exclude are:

  - EDB*.log (notice the wildcard—there can be several)
  - RES1.log
  - RES2.log

- NTDS Working folder specified in:

  **HKEY_LOCAL_MACHINE \System\CurrentControlSet\Services\NTDS\Parameters\DSA Working Directory**

  Specific files to exclude are:

  - TEMP.edb
  - EDB.chk

- SYSVOL. Use the following table for folders and files to scan or exclude:

| Folder or File | Scan or Exclude |
|---|---|
| %systemroot%\SYSVOL | Exclude |
| %systemroot%\SYSVOL\domain | Scan |
| %systemroot%\SYSVOL\domain\DO_NOT_REMOVE_NtFrs_PreInstall_Directory | Exclude |
| %systemroot%\SYSVOL\domain\policies | Scan |

| | |
|---|---|
| %systemroot%\SYSVOL\domain\scripts | Scan |
| %systemroot%\SYSVOL\staging | Exclude |
| %systemroot%\SYSVOL\staging areas | Exclude |
| %systemroot%\SYSVOL\sysvol | Exclude |

- FRS Working Directory specified in:

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters\Working Directory**

  Files to exclude:

  - <FRS Working Directory>\jet\sys\edb.chk
  - <FRS Working Directory>\jet\ntfrs.jdb
  - <FRS Working Directory>\jet\log*.log

- FRS Database Log files specified in:

  **HKEY_LOCAL_MACHINE\system\currentcontrolset\services\NtFrs\Parameters\DB Log File Directory**

  Default location is %windir%\ntds.

  Files to exclude:

  - <FRS Working Directory>\jet\log*.log (if registry entry is not set)
  - <DB Log File Directory>\log*.log (if registry entry is set)

- FRS Replica_root files specified in:

  **HKEY_LOCAL_MACHINE\system\currentcontrolset\services\NtFrs\Parameters\Replica Sets\GUID\Replica Set Root**

- Staging directory found in:

  **HKEY_LOCAL_MACHINE\system\currentcontrolset\services\NtFrs\Parameters\Replica Sets\GUID\Replica Set Stage**

- FRS Preinstall directory located at:

  **<Replica_root>\DO_NOT_REMOVE_NtFrs_PreInstall_Directory**.

  The Preinstall directory is always open exclusively when FRS is running.

The following tasks for managing domain controllers are described in this objective:

- Preparing for Active Directory Installation
- Installing a Domain Controller in an Existing Domain
- Adding Domain Controllers in Remote Sites
- Installing a Domain Controller in an Existing Domain Using Restored Backup Media

- Performing an Unattended Installation of Active Directory
- Verifying Active Directory Installation
- Renaming a Domain Controller
- Decommissioning a Domain Controller
- Forcing the Removal of a Domain Controller

# Preparing for Active Directory Installation

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Properly preparing for the installation of Active Directory decreases the chances of problems occurring during the installation process and helps you quickly complete the operation.

There are a number of requirements for installing Active Directory on a new domain controller in an existing domain. This task addresses general requirements with respect to Domain Name System (DNS) configuration, placement of the domain controller in a site, and connectivity for the Active Directory Installation Wizard.

After you have gathered all the information that you need to run the Active Directory Installation Wizard and you have performed the tests to verify that all the necessary domain controllers are available, you are ready to install Active Directory on your server and create an additional domain controller in the domain.

Preparation includes installing and configuring DNS and gathering information that you need for the installation.

**Configuring DNS**
The DNS client is always present on a server running Windows Server 2003. A DNS server must be present in the forest that stores DNS data for the server. You should properly configure both the DNS client and the DNS server to ensure that name resolution and related dependencies will function as expected during the installation of Active Directory.

Ensure that any required configuration, forwarders, or zones are present and accessible prior to installation. For more information about DNS configuration best practices, see Designing the Active Directory Logical Structure on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=25466).

**Site Placement**
During installation, the Active Directory Installation Wizard attempts to place the new domain controller in the appropriate site. The appropriate site is determined by the domain controller's IP address and subnet mask. The wizard uses the IP information to calculate the subnet address of the domain controller and checks to see if a Subnet object exists in the directory for that subnet address. If the Subnet object exists, the wizard uses it to place the new Server object in the appropriate site. If not, the wizard places the new Server object in the same site as the domain controller that is being used as a source to replicate the directory database to the new domain controller. Make sure the Subnet object has been created for the desired site prior to running the wizard.

A site is allocated according to the following rules:

1. If you specify a site in the Unattended text file that is used to create the new domain controller, the domain controller will be placed directly into that site when it is built.

2. If no site is specified in the Unattended text file when the new domain controller is built, then by default the domain controller will be placed in a site based on its IP address.

3. If you specify a replica partner in the Unattended text file but do not specify a site, the new domain controller should be placed in the replica partner's site.

4. If the replica partner or site is not specified, then the allocation of the site is random. It will depend on the replica partner selected for initial replication.

**Domain Connectivity**
During the installation process, the Active Directory Installation Wizard needs to communicate with other domain controllers to join the new domain controller to the domain. The wizard needs to communicate with a member of the domain to receive the initial copy of the directory database for the new domain controller. It communicates with the domain naming master for domain installs only, so that the new domain controller can be added to the domain. The wizard also needs to contact the relative ID (RID) master so that the new domain controller can receive its RID pool, and it needs to communicate with another domain controller in order to populate the SYSVOL shared folder on the new domain controller. All of this communication depends on proper DNS installation and configuration. By using Netdiag.exe and Dcdiag.exe, you can test all of these connections prior to starting the Active Directory Installation Wizard.

## Task requirements

During the installation process, the wizard needs to communicate with other domain controllers to add this new domain controller to the domain and get the appropriate information into the Active Directory database. To maintain security, you must provide credentials that allow administrative access to the directory.

Before you begin your installation, the following conditions must exist in your environment:

- Your Active Directory forest root domain must already exist.
- If you are installing a new domain controller in a child domain, there should be at least two properly functioning domain controllers in the forest root domain.
- DNS must be functioning properly. In this guide, it is assumed that you are using Active Directory–integrated DNS zones. You must have configured at least one domain controller as a DNS server.

Creating or removing a domain or forest is beyond the scope of this guide.

The following information and tools are necessary to complete this task:

- The Active Directory Installation Wizard asks for the following specific configuration information before it begins installing Active Directory:
    - A domain administrator's user name and password
    - A location to store the directory database and log files
    - A location to store the shared system volume files (SYSVOL)
    - The password to use for Directory Services Restore Mode
    - The fully qualified DNS name of the domain to which the new domain controller will be added
- My Network Places
- Adsiedit.msc
- Netdiag.exe
- Active Directory Sites and Services
- Dcdiag.exe

To complete this task, perform the following procedures:

1. Install the DNS Server service
2. Verify DNS registration and functionality
3. Verify that an IP address maps to a subnet and determine the site association
4. Verify communication with other domain controllers
5. Verify the availability of the operations masters

| **Caution** |
| --- |
| If any verification test fails, do not continue until you determine what went wrong and fix the problems. If these tests fail, the installation is also likely to fail. |

# Install the DNS Server service

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assign a static IP address, rather than a dynamically-assigned IP address, to any computer that acts as a DNS server. To perform this procedure, your DNS infrastructure must already exist, function properly, and be configured to use Active Directory-integrated zones. This procedure describes the steps to add an additional DNS server into the DNS infrastructure.

**Administrative Credentials**

To perform this procedure, you must be a member of either the Domain Admins group or the Enterprise Admins group.

**To install the DNS server service**

1. Ensure that the computer is using a static IP address by right-clicking **My Network Places** and then clicking **Properties**.
2. In the **Network Connections** dialog box, right-click the connection that represents the connection this computer uses to attach to your network. The default label is Local Area Connection, but this can be changed, so it might not be labeled the same on your computer. Click **Properties**.
3. In the **Local Area Connection Properties** dialog box, click once on **Internet Protocol (TCP/IP)** to highlight it (be sure that you do not clear the check box in front of it), and then click **Properties**.
4. In the **Internet Protocol (TCP/IP) Properties** dialog box, ensure that **Use the following IP address:** is selected and that a valid IP address, subnet mask, and default gateway appear. Click **OK** to close the dialog box. Click **OK** again to return to your desktop.
5. In **Control Panel**, click **Add or Remove Programs**. Click **Add/Remove Windows Components.**
6. Scroll down to **Networking Services**. Highlight it and click **Details**.
7. In the **Networking Services** dialog box, select the check box in front of **Domain Name System (DNS)**. Click **OK**.
8. Click **Next**. Provide the location of the installation files, if necessary. After the installation is complete, click **Finish** to end the wizard, and then click **Close** to exit **Add or Remove Programs**.

# Verify DNS registration and functionality

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that Domain Name System (DNS) is functioning so that other domain controllers can be located.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To verify DNS registration and functionality**

1. Open a Command Prompt.
2. On a **Windows Server 2003** computer, type the following command, and then press ENTER:

   **netdiag /test:dns**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   On a **Windows Server 2008** or **Windows Server 2008 R2** computer, type the following command, and then press ENTER:

   **dcdiag /test:dns**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   If DNS is functioning, the last line of the response for all operating system versions is **DNS Test.....: Passed**. The verbose option lists specific information about what was tested. This information can help with troubleshooting if the test fails.

   If the test fails, do not attempt any additional steps until you determine and fix the problem that prevents proper DNS functionality.

# Verify that an IP address maps to a subnet and determine the site association

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to determine the site to which you want to add a Server object prior to installing Active Directory, or to verify the appropriate site prior to moving a Server object to it.

To be associated with a site, the IP address of a domain controller must map to a Subnet object that is defined in Active Directory. The site to which the subnet is associated is the site of the domain controller.

The subnet address, which is computed from the IP network address and the subnet mask, is the name of a Subnet object in Active Directory. When you know the subnet address, you can locate the Subnet object and determine the site to which the subnet is associated.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group.

**To verify that an IP address maps to a subnet and determine the site association**

1. Log on locally or open a Terminal Services connection to the server for which you want to check the IP address.
2. On the desktop, right-click **My Network Places**, and then click **Properties**.
3. In the **Network Connections** dialog box, right-click **Local Area Connection**, and then click **Properties**.
4. Double-click **Internet Protocol (TCP/IP)**.
5. Use the values in **IP address** and **Subnet mask** to calculate the subnet address and then click **OK**.
6. Click **OK** again and close the **Network Connections** dialog box.
7. Open Active Directory Sites and Services.
8. Expand the **Sites** container, and then click the **Subnets** container.
9. In the **Name** column in the details pane, find the Subnet object that matches the subnet address.
10. In the **Site** column, note the site to which the IP subnet address is associated.

    If the site that appears in the **Site** box is not the appropriate site, contact a supervisor and find out whether the IP address is incorrect or whether to move the Server object to the site indicated by the subnet.

# Verify communication with other domain controllers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that domain controllers can be located.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain users group in Active Directory.

**To verify communication with other domain controllers**

1. Open a Command Prompt.
2. On a **Windows Server 2003** computer, type the following command, and then press ENTER:

   **netdiag /test:dsgetdc**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   If domain controllers are successfully located, the last line of the response is **DC discovery test.....: Passed.** The verbose option lists the specific domain controllers that are located.

   On a **Windows Server 2008** or **Windows Server 2008 R2** computer, type the following command, and then press ENTER:

   **nltest /dclist:** <*domain name*>

   If domain controllers are successfully located, the last line of the response is **The command completed successfully**.

   If the tests fail on any of the operating system versions, do not attempt any additional steps until you determine and fix the problem that prevents communication with other domain controllers.

# Verify the availability of the operations masters

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that the operations masters can be located and that they are online and responding.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain users group in Active Directory.

| **Note** |
| --- |
| You can use these tests prior to installing Active Directory as well as afterward. To perform the test prior to installing Active Directory, you must use the **/s** option to indicate the name of a domain controller to use. You do not need the **/s** option to perform the test after installing Active Directory. The test automatically runs on the local domain controller where you are performing the test. The commands listed in this procedure show the **/s** option. If you are performing this test after installing Active Directory, omit the **/s** option. For a more detailed response from this command, you can use the verbose option by adding **/v** to the end of the command to see the detailed response. |

**To verify the availability of the operations masters**

1. Open a Command Prompt.
2. Type the following command to ensure that the operations masters can be located and then press ENTER:

   **dcdiag /s:** *domaincontroller* **/test:knowsofroleholders /verbose**

   where *domaincontroller* is the name of a domain controller in the domain in which you want to add the new domain controller. The verbose option provides a detailed list of the operations masters that were tested. Near the bottom of the screen, a message confirms that the test succeeded. If you use the verbose option, look carefully at the bottom part of the displayed output. The test confirmation message appears immediately after the list of operations masters. Press ENTER.
3. Type the following command to ensure that the operations masters are functioning properly and are available on the network:

   **dcdiag /s:** *domaincontroller* **/test:fsmocheck**

   where *domaincontroller* is the name of a domain controller in the domain in which you want to add the new domain controller. The verbose option provides a detailed list of the operations masters that were tested. Near the bottom of your screen, a message confirms that the test succeeded. Press ENTER.

   If these tests fail, do not attempt any additional steps until you determine and fix the problem that prevents locating operations masters and verifying that they are functioning properly.

# Installing a Domain Controller in an Existing Domain

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This task covers the installation of Active Directory onto a Windows Server 2003 system that will become a domain controller in an existing Active Directory domain. To ensure successful installation of a new domain controller, you should verify that all critical services that Active Directory depends on are configured following Microsoft best practices. For more information about best practices for planning, testing, and deploying Active Directory, see Designing and Deploying Directory and Security Services on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=27638).

**Task Requirements**

The following tool is required to perform the procedure for this task:

Dcpromo.exe

To complete this task, perform the following procedure:

- Install Active Directory

**Note**

By default, when a domain controller account is added to the existing Active Directory domain, it is assigned an "Account Ops-FC" access control entry (ACE) that gives members of the Account Operators group full control over this domain controller account, which is not a recommended configuration. For example, members of Account Operators group will be able to reset this domain controller's password. Because the Account Operators group has significant power in the domain, we recommend that you add members to it with caution. For a detailed description of the Account Operators group, see Default groups (http://go.microsoft.com/fwlink/?LinkID=131422). To modify permissions for Account Operators on a computer account, you can use the Active Directory Users and Computers snap-in and complete the following steps:

1. To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
2. In the console tree, right-click the affected domain controller account and then click **Properties**.
3. On the **Security** tab, select **Account Operators** in the **Group or user names** list, and then modify permissions according to the specifications of your environment.

You can also install Active Directory from installation media or by performing an unattended installation. For information about completing each of these tasks, see the following:

Installing a Domain Controller in an Existing Domain Using Restored Backup Media

Performing an Unattended Installation of Active Directory

# Install Active Directory

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the Active Directory Installation Wizard to install Active Directory on a member server in your domain to create an additional domain controller in an existing domain.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To install Active Directory**

1. Click **Start**, click **Run**, type **dcpromo** and then press ENTER.
2. The Active Directory Installation Wizard appears. At the Welcome screen, click **Next**.
3. For **Domain Controller Type**, select **Additional domain controller for an existing domain**. Click **Next**.
4. For **Network Credentials**, enter the user name, password, and domain for the user account that has permission to add this new domain controller to the domain. Click **Next**.
5. Enter the name of the domain that you want the new domain controller to host. Click **Next**.
6. For **Database and Log Locations**, enter the paths for the locations of the directory database (Ntds.dit) and the log files. For better performance, store the database and log files on separate physical disk drives. Click **Next**.
7. For **Shared System Volume**, enter the path where you want to locate the system volume (SYSVOL). Click **Next**.
8. Under **Directory Services Restore Mode Administrator Password**, enter the password that you want to use when you need to start Directory Services Restore Mode. Click **Next**.
9. The **Summary** screen displays a list of the items you chose. Verify that the information is correct, and then click **Next** to proceed with the installation.
10. The wizard proceeds to install Active Directory. When it finishes, the wizard displays a summary screen listing the domain and site in which the new domain controller is a member. Verify that this information is correct. Click **Finish** to close the wizard.
11. Click **Restart** to restart the domain controller.
12. Let the domain controller restart. If any message indicates that one or more services has failed to start, restart the domain controller one more time. If the initial replication cycles have not had enough time to complete during the first restart on a new domain controller, some services may be unable to start successfully. If the message appears during additional restarts, examine the event logs in Event Viewer to determine the cause of the problem.

# Installing a Domain Controller in an Existing Domain Using Restored Backup Media

Updated: December 5, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you install Active Directory from restored backup media, you can reduce the replication traffic that is initiated during the installation of an additional domain controller in an Active Directory domain. Reducing the replication traffic reduces the time necessary to install the additional domain controller.

The procedures in this task are particularly useful for installing domain controllers in remote sites. By using these procedures, you can avoid having to either replicate the entire Active Directory replica over a wide area network (WAN) link or disconnect an existing domain controller while it is being shipped to the remote site. If you are installing additional domain controllers in remote sites and you want to minimize the Active Directory and SYSVOL replication that is required during the installation from backup media, use the information in this topic in conjunction with the information in Adding Domain Controllers in Remote Sites.

When the domain controller that you are installing will be a Domain Name System (DNS) server and you are using Active Directory–integrated DNS zones, the DomainDNSZones and ForestDNSZones application directory partitions are not included in the restored backup media by default. If you want to include application directory partitions in the restored backup media that is used to install Active Directory, additional procedures are required to complete the installation task. Follow the instructions for including application directory partitions in the installation media.

**Task requirements**

To begin the task to install a domain controller from restored backup media without application directory partitions, ensure that the following requirements are met:

- A Windows Server 2003–based domain controller must be running in each domain where you will be performing installations from backup media.
- The restored system state backup that is used to create additional domain controllers must be taken from a domain controller in the same domain as the new additional domain controller.
- The server that is being installed as a domain controller must be running Windows Server 2003, and the version must be the same as the domain controller from which the backup was taken. For example, you cannot use backup media from a domain controller running Windows Server 2003 to create a domain controller running Windows Server 2003 with Service Pack 1 (SP1). The reverse is also true.
- The restored system state backup that is used to create additional domain controllers must be taken on a domain controller that matches the processor type of the new domain controller. System state backups that are taken on a domain controller that has a 32-bit processor cannot be used to install a domain controller that has a 64-bit processor. The reverse is also true.
- During Active Directory installation, Dcpromo checks that the value of the tombstone lifetime in the restored system state backup matches the value on an existing domain controller. If you plan to change the value of the tombstone lifetime, change this value before you create the backup.
- If the domain controller that you are creating is to be a global catalog server, the system state backup that you restore must be taken from an existing global catalog server in the domain.

On servers that are running Windows Server 2003 with SP1, you can use restored backup media to install a domain controller that is a DNS server (stores the DomainDNSZones and ForestDNSZones application directory partitions) or that stores other application directory partitions. In addition to the previous requirements, to begin the task to install a domain controller from restored backup media that includes application directory partitions, ensure that the following requirements are met:

- The forest functional level has been raised to Windows Server 2003.
- The domain controller on which you created the system state backup is running Windows Server 2003 with SP1.

- The domain controller on which you created the system state backup contains the application directory partitions that you want to include.
- The server computer that you are installing is running Windows Server 2003 with SP1.
- You have created an answer file that contains the distinguished names (or * for all names) of the application directory partitions that you want to include.

The following tools are required to perform the procedures for this task:

- Ntbackup.exe
- Dcpromo.exe
- Ref.chm or Unattend.txt file, or both, for installations that include application directory partitions.

To complete this task, perform the procedures for the following methods:

1. Back up system state

   Back up the system state of an existing domain controller according to the requirements described above.

2. As an option, before you restore the backup, copy the .bkf file to a CD, DVD, or other removable media from which you will subsequently restore the backup to an alternate location on the local hard drive of the server on which Active Directory is being installed. You can use this media to restore the same backup to any number of servers that will be installed as domain controllers. With this method, you restore the backup for each domain controller that you install. Compare this method to method 3.2, in which you restore the backup only once and copy the restored files to the removable media.

3. Restore system state to an alternate location

   Select the location for the system state backup that you will use to install a new domain controller. Use one of the following locations for restoring the system state backup:

   - Restore the .bkf file to a volume on the server that will be installed as a domain controller. We recommend restoring to a folder named \NTDSRESTORE on the volume that will host the Ntds.dit file when Dcpromo is run, if space permits. Otherwise, restore to a folder named \NTDSRESTORE on a volume that has sufficient free space. For additional criteria regarding the volume on which you restore the backup, see Adding Domain Controllers in Remote Sites.
   - Restore the .bkf file to the local hard drive of any computer, and then burn the expanded restore tree to a CD, DVD, or other removable storage media. Install Active Directory directly from this media. You can use this media to directly install any number of domain controllers. With this method, you restore the backup only once.

4. Install Active Directory from media.

   Install the domain controller from the system state backup that you restored in step 3 by using one of the following methods:

   - Install Active Directory from restored backup media to create a new domain controller that does not include application directory partitions.
   - Include application directory partitions in an Active Directory installation from backup media to create a new domain controller that contains application directory partitions. (See the special requirements described earlier in

this section.) This method uses an answer file to specify the application directory partitions to include in the Active Directory installation. To use this method, you must first Create an answer file for domain controller installation. You can also include instructions to install the DNS Server service in the answer file.

**See Also**
**Concepts**
Adding Domain Controllers in Remote Sites
Create an answer file for domain controller installation

# Back up system state

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Ntbackup.exe provides simple and advanced options for backing up Active Directory components. When you back up system state, you can choose to include or exclude system-protected boot files. System-protected boot files are not used for installations from restored backup media. When the backup file that you create is to be used for additional domain controller installations, you can clear the advanced option to back up system-protected files. Clearing this option decreases the size of the .bkf file, as well as the time required to back up, restore, and copy the system state files.

Use these procedures to back up the system state only. These procedures do not back up the system disk or any other data on the domain controller except for the system-protected files.

Use the first procedure, "To back up system state including system-protected files," for routine system state backup. Use the second procedure, "To back up system state excluding system-protected files," if you want to create a smaller backup that is effective for installing domain controllers from restored backup media.

| **Note** |
| --- |
| To back up system state, you must log on locally to the domain controller, or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

- To perform the following two procedures, you must be a member of the Domain Admins group or a member of the Backup Operators group.

#### To back up system state including system-protected files

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.

   This procedure provides steps for backing up in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard.

2. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
3. Select **Back up files and settings**, and then click **Next**.
4. Select **Let me choose what to back up**, and then click **Next**.
5. In the **Items to Back Up** window, double-click **My Computer**.
6. In the expanded list below **My Computer**, check **System State**, and then click **Next**.
7. Select a location to store the backup:

   - If you are backing up to a file, type the path and file name for the backup (.bkf) file (or click **Browse** to find a folder or file).

   - If you are backing up to a tape unit, choose the tape that you want to use.

     | **Note** |
     | --- |
     | You should not store the backup on the local hard drive. Instead, store it in a location, such as a tape drive, away from the computer that you are backing up. |

8. Type a name for this backup according to the recommendations in Backing Up Active Directory Components, and then click **Next**.
9. On the last page of the wizard, click **Advanced**.

10. Do not change the default options for **Type of Backup**. **Normal** should be selected, and the check box for **Backup migrated remote storage data** should remain cleared. Click **Next**.
11. Select **Verify data after backup**, and then click **Next**.
12. In the **Backup Options** dialog box, select a backup option, and then click **Next**.
13. If you are replacing the existing backups, select the option to allow only the owner and administrator access to the backup data and to any backups that are appended to this medium, and then click **Next**.
14. In the **When to back up** box, select the appropriate option for your needs, and then click **Next**.
15. If you are satisfied with all of the options that are selected, click **Finish** to perform the backup operation according to your selected schedule.

| **Note** |
| --- |
| The system state can also be backed up by using Ntbackup from a command line with appropriate parameters. For more information, at a command prompt type **ntbackup /?**. |

The following procedure produces a smaller .bkf file that does not include system boot files. By using this procedure, you can reduce the time that is required to perform the backup and subsequent restore, as well as the amount of disk space that is required. This method is recommended when the restored backup is to be used for installing additional domain controllers.

**To back up system state excluding system-protected files**

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.
2. On the **Welcome to the Backup or Restore Wizard** page, click **Advanced Mode**, and then click the **Backup** tab.
3. In the console tree, select the **System State** check box.
4. In **Backup media or file name**, type a name for this backup according to the recommendations in Backing Up Active Directory Components.
5. Click **Start Backup**, and then click **Advanced**.
6. Clear the **Automatically back up System Protected Files with the System State** check box, and then click **OK**.
7. Click **Start Backup**.

**See Also**
**Concepts**
Enable Remote Desktop
Create a Remote Desktop Connection

# Restore system state to an alternate location

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Perform this procedure to create media for installing a domain controller from restored backup media or to allow an authoritative restore of SYSVOL. After the domain controller installation is complete, delete the files in the alternate location.

You can restore the system state backup to an alternate location on the domain controller from which the backup was made, a location on another computer, or a location on the computer that you want to install as a domain controller.

**Administrative credentials**

To perform this procedure, you must be a member of the Backup Operators group, as follows:

- Restore system state on a member or workgroup server: Backup Operators group on the local computer
- Restore system state on a domain controller: Backup Operators group in the domain

**To restore system state to an alternate location**

1. Log on to the server that has the alternate location to which you are restoring system state backup files.
2. Click **Start**, click **Run**, type **ntbackup**, and then click **OK**.
3. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
4. Click **Restore Files and settings**, and then click **Next**.
5. On the **What to Restore** page, click **Browse**, and then, in the **Open Backup File** dialog box, click **Browse** again.
6. Navigate to the .bkf file that you want to restore to an alternate location. The .bkf file can be located in a folder on the current computer, in a shared folder on the backup computer or other network computer, or on an external drive that contains removable media.
7. In the **Select file for catalog** dialog box, click the .bkf file that you want to restore, and then click **Open**.
8. In the **Open Backup File** dialog box, click **OK**.
9. In **Items to restore**, double-click **File**, and then double-click the .bkf file that you want to restore.
10. Below the .bkf file that you want to restore, select the **System State** check box, and then click **Next**. (You do not need to restore the system disk to an alternate location.)
11. On the **Completing the Backup or Restore Wizard** page, click **Advanced**.
12. In the **Restore Files to** drop-down list, click **Alternate Location**.
13. In **Alternate Location**, type the path (or browse) to the local folder to which you are restoring the backup, and then click **Next**. We recommend restoring to a folder named NTDSRESTORE, if space permits, on the volume that will host the Ntds.dit file when Dcpromo is run. Otherwise, restore to a folder named \NTDSRESTORE on another volume that has sufficient free space.
14. On the **How to Restore** page, accept the default selection **Leave existing files (Recommended)**, and then click **Next**.
15. On the **Advanced Restore Option**s page, accept the default selections **Restore security settings** and **Preserve existing volume mount points**, and then click **Next**.
16. On the **Completing the Backup or Restore Wizard** page, click **Finish**.

# Install Active Directory from restored backup media

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to install Active Directory from backup media to create an additional domain controller in an existing domain.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain into which you are installing the additional domain controller.

**To install Active Directory from restored backup media**

1. Click **Start**, click **Run**, type **dcpromo /adv**, and then press ENTER.
2. In the Active Directory Installation Wizard, select **Additional domain controller for existing domain**.
3. Select **From these restored backup files**, and point to the same location where you restored the system state data.
4. If the domain controller whose system state backup you are using is a global catalog server, the Active Directory Installation Wizard asks you whether you want this server to also be a global catalog server.
5. Give appropriate credentials for the operation.
6. Enter the domain of the new domain controller. This domain must be the domain of the domain controller whose system state backup you are using.
7. Complete the remaining pages of the Active Directory Installation Wizard.

   Dcpromo.exe will install Active Directory using the data present in the restored files, which eliminates the need to replicate every object from a partner domain controller. However, objects that were modified, added, or deleted since the backup was taken must be replicated. If the backup was recent, the amount of replication required will be considerably less than that required for a regular Active Directory installation.

   After the installation operation completes successfully and the computer is restarted, the folder and subfolders that contain the restored system state can be removed from the local disk.

**See Also**
**Concepts**
Restore system state to an alternate location
Include application directory partitions in an Active Directory installation from backup media

# Include application directory partitions in an Active Directory installation from backup media

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to install Active Directory from restored backup media that includes application directory partitions to create an additional domain controller in an existing domain. In this procedure, you edit a domain controller installation answer file to provide instructions for including application directory partitions in the installation. You must have already created the answer file according to the directions in Create an answer file for domain controller installation.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain into which you are installing the additional domain controller.

**To include application directory partitions in an Active Directory installation from backup media**

1. Open the answer file that you created to install the domain controller.
2. To include application directory partitions that are contained in the system state backup, add the following entry to the end of the answer file:

   **ApplicationPartitionsToReplicate=**

3. Provide a value for **ApplicationPartitionsToReplicate** as follows:

   - If you want to include all application directory partitions, use the value *****.

   - If you want to include specific application directory partitions, type the distinguished name of each directory partition. Separate each distinguished name with a space, and enclose the entire list in quotation marks, as shown in the following example:

     **ApplicationPartitionsToReplicate="dc=app1,dc=contoso,dc=com dc=app2,dc=contoso,dc=com"**

4. In the entry **ReplicationSourcePath=**, type the path to the folder that contains the restored system state backup files on the installation computer.
5. If you do not want Dcpromo to prompt the user for passwords, type the password in the **Password=** entry for the account that you will use to install the domain controller, type the password in the **SafeModeAdminPassword=** entry that you will use to provide access to Directory Services Restore Mode, and then save the answer file.

   | **Note** |
   | --- |
   | Passwords are automatically deleted from the answer file when Dcpromo runs. |

6. Open a command prompt, and then change directories to the location of the answer file.
7. At the command prompt, type the following command, and then press ENTER:

   **dcpromo /adv /answer:"** *Drive* **:** *PathToAnswerFile* **** *AnswerFileName* **"**

   Active Directory installation occurs automatically. If you left passwords blank in the answer file, Dcpromo prompts you for your administrative password and for the Directory Services Restore Mode password. If you specified "no" for the **RebootOnSuccess=** entry in the answer file, Dcpromo prompts you to restart the server after installation.

# Adding Domain Controllers in Remote Sites

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can create an additional domain controller in a domain by installing Active Directory on a server computer. When the additional domain controller is being placed in a remote site, you can install Active Directory on the server either before or after shipping it to the remote site, as follows:

- Ship the computer as a workgroup computer, and install Active Directory in the remote site. Enable Remote Desktop on the computer before you ship it so that you can perform the installation remotely. In the remote site, you can either:
    - Install Active Directory from restored backup media that has been shipped to the site on removable media or that has been restored to a location on the server itself before shipping.
    - Install Active Directory over the network.
- Install Active Directory on the server in a hub or staging site, and ship the installed domain controller to the remote site.

Both methods have advantages and disadvantages, and both methods require care to ensure the secure transfer of Active Directory data, whether it is installed or in the form of backup files that are stored on the server or on removable media.

For information about the advantages and disadvantages of shipping a server to a remote site before or after installing Active Directory, see Known Issues for Adding Domain Controllers in Remote Sites.

For information about how best to manage adding domain controllers to remote sites for the method you are using, see Best Practices for Adding Domain Controllers in Remote Sites.

By following the guidelines in this guide, you can decide the best method for your environment of adding domain controllers in remote sites. By following the instructions in this guide, you can safely and securely install domain controllers in remote sites, either locally or remotely.

The following tasks for adding domain controllers in remote sites are described in this objective:

- Preparing a Server Computer for Shipping and Installation from Backup Media
- Preparing an Existing Domain Controller for Shipping and Long-Term Disconnection
- Reconnecting a Domain Controller After a Long-Term Disconnection

# Known Issues for Adding Domain Controllers in Remote Sites

Updated: February 9, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following two methods to add domain controllers in remote sites:

- Ship the member computer to the remote site, and then install Active Directory by using the **dcpromo /adv** option, which uses restored system state backup media as the source for the Active Directory installation in the remote site.
- Install Active Directory in the hub site by using the normal Dcpromo method, and then ship the installed domain controller to the remote site.

You can use the information in this section to determine the method for adding domain controllers in remote sites that is best for your environment. SYSVOL replication issues potentially affect both methods, and each method has advantages and disadvantages that are discussed in this section.

| **Important** |
| --- |
| Do not attempt to perform actions based only on the recommendations that are discussed in this topic. Step-by-step guidance is provided in the task-based topics for all actions that are recommended in this topic. Follow the See Also links to the related task-based topics. |

**SYSVOL Replication**

SYSVOL is a shared folder that stores files that must be available and synchronized among all domain controllers in a domain. SYSVOL contains the NETLOGON share, Group Policy settings, and File Replication service (FRS) staging directories and files. The SYSVOL share is required for Active Directory to function properly.

The primary focus for both methods of installing additional domain controllers in remote sites is to avoid the replication of Active Directory over a wide area network (WAN) between the remote site and the hub site. Each method accomplishes this goal. However, depending on the size of your SYSVOL, you might also be concerned about replication of SYSVOL files over the network. Unless you follow specific instructions, the SYSVOL tree might be created on the new domain controller through replication of the entire tree from an existing domain controller in the domain. Regardless of which method you use to add domain controllers to remote sites, you might want to take additional steps to manage SYSVOL creation on the new domain controller to avoid replicating the full SYSVOL from another domain controller in the domain.

When you install a domain controller from backup media, preliminary steps are required to ensure that SYSVOL is created from the local copy of the restored backup media. Similarly, preliminary steps are required to avoid full SYSVOL synchronization when you ship an installed domain controller and restart it in the remote site. These requirements are discussed for each method respectively in the following topics:

- Preparing a Server Computer for Shipping and Installation from Backup Media
- Preparing an Existing Domain Controller for Shipping and Long-Term Disconnection

**Using Backup Media to Install Active Directory in a Remote Site**

The ability to use restored backup media to install domain controllers is a new feature in Windows Server 2003. The method for using backup media to install domain controllers includes the following general steps:

1. Back up system state on a domain controller in the domain in which you are adding the new domain controller. If you want the additional domain controller to be a global catalog server, back up a global catalog server. If you want the additional domain controller to be a Domain Name System (DNS) server, back up a DNS server.
2. Restore the backup to an alternate location. You can restore the backup directly to the computer that you want to install as a domain controller, or you can transfer it to removable media.
3. Run Dcpromo with the **/adv** option and indicate the restored backup as the source for the Active Directory installation.

This method of installing domain controllers in remote sites has several advantages. One of the primary advantages of this method is that it substantially reduces the network bandwidth requirement compared to network-based installations. This method also has a few issues that mostly affect deployments that have a large number of remote sites. If you deploy more than 100 remote sites, additional considerations might be necessary. For information about large branch office deployments, see the Windows Server 2003 Active Directory Branch Office Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42506).

**Advantages of Using Backup Media to Install Active Directory in a Remote Site**

The following advantages are associated with using backup media to install a domain controller in a remote site:

- You can install many domain controllers from a single source of removable backup media. Although you can restore backup media directly to an alternate location on the server computer that you are going to install as a domain controller, you can also use that media as the source for any number of domain controllers by either copying or restoring the system state backup to removable media. For more information about the effects of copying — as opposed to restoring — a system state backup to removable media, see Preparing a Server Computer for Shipping and Installation from Backup Media.

- You do not have to disconnect the domain controller from the replication topology. Therefore, you can avoid the disadvantages that are associated with a domain controller that does not replicate. For information about the problems that are associated with domain controller disconnection, see Issues with Installing Domain Controllers Before Shipping Them to the Remote Site.

- You avoid replicating the entire Active Directory over a WAN link, particularly a link that requires a dial-up connection.

- If you enable Remote Desktop on the server before you ship it, you do not have to employ an administrator with Domain Admins credentials in the remote site.

**Issues with Using Backup Media to Install Active Directory in Remote Sites**

The following issues are associated with using backup media to install a domain controller in a remote site:

- **Domain Admins credentials and remote installation.** An administrator must have Domain Admins credentials to install Active Directory. Assuming that you do not employ a service administrator with this level of administrative credentials in each branch site, a domain administrator in the hub site must be able to connect remotely to the server to perform the installation. Therefore, you must be sure to enable Remote Desktop on the server before you ship it to the remote site.

- **Time to restore the system state backup.** The installation media is prepared by restoring a system state backup to an alternate location. Therefore, preparing the media requires taking the backup itself and restoring the backup. These tasks add time to the installation of a single domain controller. However, if you take advantage of the ability to transfer the restored backup files to removable media, you perform the preliminary backup and restore processes only once to install any number of domain controllers. In addition, you can follow instructions to prepare a smaller backup file to further decrease the time that is required for restoring and copying backup media. The volume on which you restore the backup on the target server also affects the speed of the installation. Moving the Ntds.dit file is faster than copying it. If you restore the media to the same location that will be used to host the Active Directory database, the Ntds.dit file will be moved (as opposed to being copied) into the new location, eliminating the additional time required to copy the file. For more information about the criteria that affect how long installation from backup media takes, see Preparing a Server Computer for Shipping and Installation from Backup Media.

- **Backup source for application directory partitions.** When DNS zone data is stored in application directory partitions, the replication impact can be significant if application directory partitions must be replicated over the corporate network. System state data that you restore from backup to an alternate location does not include application directory partitions if the backup is performed on servers running Windows Server 2003 with no service pack installed.

  Including application directory partitions in the backup media has the following requirements:

  - The domain controller that you back up and the computer that you intend to install as a domain controller must both be running Windows Server 2003 with Service Pack 1 (SP1).

  - The forest functional level must be set to Windows Server 2003 because linked-value replication is required to

ensure that cross-references are correctly updated for the application directory partition replica set.

- You must use an answer file to install Active Directory because the Dcpromo user interface (UI) does not provide an option for specifying application directory partitions. Use the answer file to provide the distinguished names of the application directory partitions that you want to include in the installation.

For more information about how to include application directory partitions and create a DNS server, see Preparing a Server Computer for Shipping and Installation from Backup Media.

- **Bridgehead server load balancing.** If backup media are sent to many sites and if enough domain controllers are promoted at the same time, you might experience performance issues with the bridgehead servers that are the source for Active Directory and FRS replication.

| **Note** |
| --- |
| These issues are of concern only in situations in which hundreds of domain controllers might be promoted at the same time and their need for bridgehead server resources is very high. If you are deploying hundreds of domain controllers in branch sites, see the Windows Server 2003 Active Directory Branch Office Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42506). |

  - **Active Directory intersite replication.** You cannot load-balance intersite connections to and from the hub site until the domain controller is installed. If a large number of domain controllers are being installed in remote sites (more than 100), manual rebalancing of connections might be required after the domain controllers are installed. For information about how to use the Active Directory Load Balancing (ADLB) tool to rebalance connections, see the Windows Server 2003 Active Directory Branch Office Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42506).

  - **FRS replication.** Because FRS on the source computer uses CPU, memory, and disk resources, the FRS recommendation is to perform a staged update on no more than 10 branch office domain controllers at a time per source hub domain controller. If a single domain controller functions as the source for SYSVOL replication to more than 10 destination domain controllers, performance on the source domain controller can decrease significantly. To balance source domain controllers, you can use an answer file with Dcpromo to specify the source domain controller.

#### Installing Domain Controllers Before Shipping Them to the Remote Site

When you install a domain controller and then disconnect it from the network for a period of time, you interrupt the normal replication activities of other domain controllers on the network. This naturally creates error conditions that result from various failed operations, such as attempts to replicate with the disconnected domain controller. As long as you are aware of the issues that disconnections cause, you can take steps to ensure smooth disconnection and reconnection.

#### Advantages of Installing Domain Controllers Before Shipping Them to the Remote Site

The following advantages are associated with installing domain controllers before shipping them to the remote site:

- **Standardization.** The process for installing domain controllers can be automated and standardized in the hub or staging site, with the one additional step of packing and shipping the domain controller. If you follow the instructions for safe disconnection and reconnection, restarting the domain controller in the remote site is all that is required.

- **Branch site personnel.** The requirement for personnel with Domain Admins credentials is contained within the hub site; that is, intervention by personnel with Domain Admins credentials is not required at the branch site.

#### Issues with Installing Domain Controllers Before Shipping Them to the Remote Site

The following issues are associated with installing domain controllers and then disconnecting them from the network while they are shipped to the remote site:

- **Disconnection error conditions.** After disconnection, online domain controllers in the domain continue to attempt replication with the disconnected domain controller, causing Active Directory and FRS errors to be generated for as long as the domain controller is disconnected.

- **Additional preparation.** Additional preparation is required to ensure smooth reconnection:
    - Preparing for the nonauthoritative restart of SYSVOL.
    - Ensuring an adequate tombstone lifetime to avoid the possibility of objects remaining on the domain controller that have been permanently deleted from the directory on all other domain controllers. The tombstone lifetime is a forest-wide setting that determines how long an object deletion persists in the directory.
- **Protection of existing accounts and metadata.** You must ensure that computer accounts and metadata for the domain controller are not deleted or improperly modified while the domain controller is disconnected.
- **Risk of lingering objects.** A lingering object is an object that remains on a disconnected domain controller after the object has been permanently deleted from Active Directory on all connected domain controllers. Deletion updates are replicated as tombstone objects. These objects have a limited lifetime in Active Directory, which is defined by the tombstone lifetime. After a tombstone is permanently removed from Active Directory, replication of the deletion it represented is no longer possible. Therefore, if you restart a domain controller on which such an object remains, replication does not recognize that object as a deleted object, and it remains in Active Directory on only the reconnected domain controller and nowhere else. If you plan to disconnect a domain controller for longer than the period of time that a domain controller keeps track of object deletions (the tombstone lifetime), you must take additional steps to ensure directory consistency. For more information about lingering objects and their causes and effects, see Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042).

**Maintaining Directory Consistency When Disconnecting a Domain Controller**

Maintaining consistency of Active Directory data involves several related issues. Review the following known issues before disconnecting an installed domain controller:

- Relationship of tombstone lifetime to Active Directory backup
- Protection against lingering object replication
- Availability of operations master roles in the domain and forest
- Up to dateness of Active Directory replication at the time of disconnection
- SYSVOL consistency on reconnection

For procedures to ensure that all of these issues are addressed, see the following topics:

- Preparing an Existing Domain Controller for Shipping and Long-Term Disconnection
- Reconnecting a Domain Controller After a Long-Term Disconnection

**Tombstone Lifetime and Backup**

Active Directory backups are useful for recovering a domain controller for only as long as the tombstone lifetime. When an object is deleted, Active Directory replicates the object as a tombstone, which consists of a small subset of the attributes of the deleted object. The tombstone is retained in Active Directory for 60 days by default in a Windows Server 2003 forest. This default value is changed in Windows Server 2003 with Service Pack 1 (SP1). The tombstone is retained for 180 days by default in a new forest that is created on a domain controller running Windows Server 2003 with SP1. NTBackup.exe, the Windows Server backup utility, will not restore a system state backup that is older than the tombstone lifetime.

The tombstone lifetime value that is in effect when a domain controller is upgraded to Windows Server 2003 SP1 is not changed by the installation of Windows Server 2003 SP1. The existing value is maintained until you change it manually. You can determine the value of the **tombstoneLifetime** attribute by viewing the properties of CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,DC=*ForestRootDomain* in ADSI Edit (adsiedit.msc), which is available in Windows Support Tools. A value in **tombstoneLifetime** of **<Not Set**> always indicates that the Windows Server 2003 default value of 60 is in effect. If you create a new forest on a domain controller running Windows Server 2003 with SP1, the default **tombstoneLifetime** value of 180 is displayed in the UI.

When the number of days in the tombstone lifetime has passed, the tombstone is permanently removed. Because a domain controller that is disconnected for a period that is longer than the tombstone lifetime cannot receive deletions that occurred before the beginning of the tombstone lifetime, a backup that is older than the tombstone lifetime cannot be used to restore Active Directory.

When conditions beyond your control cause a domain controller to be disconnected for a period that is longer than the tombstone lifetime, one or more objects that have been deleted from the rest of the directory while the domain controller was offline might remain on the disconnected domain controller.

If planned domain controller disconnections are consistently lasting longer than the number of days in the tombstone lifetime, consider extending the tombstone lifetime for the forest prior to disconnecting any domain controllers.

For more information about the causes and effects of lingering objects, see Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042).

**Protection Against Lingering Object Replication**
Domain controllers that have not performed inbound replication in the previous tombstone lifetime number of days are vulnerable to retaining lingering objects. If a domain controller that has one or more lingering objects is reconnected to the replication topology and a lingering object is subsequently updated on that domain controller, the object might be recreated in Active Directory, depending on how the strict replication consistency registry setting is configured.

A lingering object is made known to the replication system only if it is updated on the domain controller that stores it. In this case, the source domain controller attempts replication of an update to an object that the destination does not store. On destination domain controllers running Windows 2000 Server with Service Pack 3 (SP3) or Service Pack 4 (SP4) and Windows Server 2003, the **strict replication consistency** registry entry (type **REG_DWORD**) in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**) determines whether replication is allowed to proceed if the domain controller receives a request for an update to an object that it does not have.

The value in the **strict replication consistency** registry entry determines whether replication proceeds or is stopped, as follows:

- **1** (enabled): Inbound replication of the specified directory partition from the source is stopped on the destination. Replication of the directory partition is halted on both the source and destination domain controllers.

- **0** (disabled): The destination requests the full object from the source domain controller and the destination domain controller reanimates a full copy of an object it has previously deleted and permanently removed through garbage collection.

For more information about how to manage the strict replication consistency setting, including its effects and its default values, see Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042).

**Up to Dateness of Active Directory Replication**
Ensure that a domain controller is updated before you disconnect it. Immediately before you disconnect the domain controller, force replication with all replication partners and verify that each directory partition replicates to the domain controller that you are disconnecting. If replication of any directory partition does not succeed, resolve the replication problem before you disconnect the domain controller. By ensuring that replication is up to date, you can maximize the possible safe disconnection period, which cannot exceed the tombstone lifetime for the forest.

**SYSVOL Consistency**
SYSVOL replication cannot be synchronized manually. For this reason, ensuring that SYSVOL is updated before you disconnect the domain controller is more difficult than simply updating SYSVOL when the domain controller is reconnected. Regardless of the length of the disconnection, to ensure that SYSVOL is synchronized when the domain controller is reconnected, prepare the domain controller to perform a nonauthoritative restart of SYSVOL before you disconnect the domain controller. When the domain controller restarts, nonauthoritative restart of SYSVOL occurs automatically.

**See Also**
**Concepts**
Preparing a Server Computer for Shipping and Installation from Backup Media
Preparing an Existing Domain Controller for Shipping and Long-Term Disconnection
Reconnecting a Domain Controller After a Long-Term Disconnection

# Best Practices for Adding Domain Controllers in Remote Sites

Updated: February 15, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By reviewing the information in Known Issues for Adding Domain Controllers in Remote Sites, you can determine the best method to use for installing domain controllers in your remote sites. You can use this topic to learn about the best practices for the method that you decide to use.

| **Important** |
| --- |
| Do not attempt to perform actions based only on the recommendations that are discussed in this topic. Step-by-step guidance is provided in the task-based topics for all actions that are recommended in this topic. Follow the See Also links to the related task-based topics. |

**Using Backup Media to Install Active Directory in the Remote Site**

The primary purpose of using restored backup media to install a domain controller is to provide a local source for the domain, configuration, and schema directory partitions — and, optionally, global catalog partial, read-only directory partitions and Domain Name System (DNS) application directory partitions. The local source is the restored system state backup files that reside on the server that you are installing. Updates to object attributes that occur since the system state backup was made will replicate over the network from an existing domain controller in the domain or forest. Although SYSVOL is part of the system state backup, under some conditions SYSVOL is not sourced from the backup media. Configuring SYSVOL to be sourced from local backup media is more challenging and might not prove worthwhile. For more information about the conditions that determine the need for SYSVOL replication, see Known Issues for Adding Domain Controllers in Remote Sites.

To use restored backup files for installation of one or more additional domain controllers in a domain, you can either:

- Copy ("burn") either the unrestored .bkf file or the restored backup files onto removable media, such as a portable disk drive, CD, or DVD, which can be shipped with the workgroup computer when it leaves the staging site or shipped separately.

- Restore system state backup to the local hard drive of the workgroup computer before it leaves the staging site.

For information about the advantages and disadvantages of these methods, see Preparing a Server Computer for Shipping and Installation from Backup Media.

The **Dcpromo /adv** option in Windows Server 2003 to install a domain controller from backup media eliminates the Windows 2000 Server requirement to either promote the domain controller before shipping it to the remote site or promote the domain controller in the remote site by replicating the entire directory over a wide area network (WAN) connection when another domain controller for the domain is not present in the site.

The following best practices are recommended to optimize data security and consistency when you add domain controllers in remote sites:

- **Upgrade to Windows Server 2003 with Service Pack 1 (SP1).** If you use Active Directory-integrated DNS or if you want other application directory partitions to be included in the domain controller replica, upgrade the server computer to Windows Server 2003 with SP1 before Active Directory installation. When you use restored backup media to install a computer running Windows Server 2003 with no service pack installed, the replica installation does not include application directory partitions. In the case of DNS application directory partitions, the impact of replicating these directory partitions over the WAN might be significant. When you use restored backup media to install a computer running Windows Server 2003 with SP1, you can use an answer file to include application directory partitions in the replica that you install.

- **Back up the type of domain controller that you want to add.** You must back up the type of domain controller that you want to add. If you want to add a global catalog server in the remote site, back up a global catalog server in the domain. If you want to add a DNS server, back up a DNS server in the domain.

- **Take the same security precautions for shipment of removable backup media or a server computer that contains a restored backup as you would take for shipping an installed domain controller.** For information

about securing domain controllers, see Best Practice Guide for Securing Windows Server Active Directory Installations on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=28521).

- **Minimize the time between the backup and installation.** Minimizing this delay reduces the number of updates that will be required to replicate after installation.

- **Install the operating system before shipping the server to the remote site.** Installing the operating system requires expertise that might not be available at branch sites. Ideally, installation routines are available in the staging site to automate the operating system installation process and ensure uniformity for all domain controllers (partition sizes, drive letter assignments, and so on). As part of the operating system installation, apply a standardized set of hotfixes plus any available service packs to ensure service consistency throughout the forest.

- **Ship the server as a member of a workgroup rather than a member server in a domain.** If the server is joined to a domain and then stolen during shipment, information about domain names, DNS suffixes, and number of domains in the forest can aid attackers in attempts to compromise or steal directory data.

- **Ship computers with properly configured Internet Protocol (IP), subnet mask, and default gateway addresses.** Remember to reconfigure the server with TCP/IP settings that are appropriate to the target site, not the staging site. Specifically, the domain controller must not point to itself for DNS.

- **Enable Remote Desktop on the server computer before shipping.** This best practice assumes that you need to be able to install and manage Active Directory remotely rather than employing an administrator with Domain Admins credentials in each remote site.

#### Installing Domain Controllers Prior to Shipping to the Remote Site

When you install Active Directory in the hub or staging site, disconnect the installed domain controller, and then ship the computer to the remote site, you are disconnecting a viable domain controller from the replication topology. The most significant risk from disconnection is that the domain controller will remain offline long enough to exceed the tombstone lifetime and thereby become capable of retaining objects that have been permanently deleted from the directory on all other domain controllers in the domain. Such objects, called lingering objects, cause directory inconsistency and, under certain conditions, can be reintroduced into the directory. For information about the causes and effects of lingering objects and how to avoid them, see Known Issues for Adding Domain Controllers in Remote Sites.

The following best practices reduce the possibility of Active Directory consistency problems due to lingering objects remaining on domain controllers that are disconnected for long periods of time. Take the following precautions to avoid directory consistency problems when you disconnect an existing domain controller and to ensure that if inadvertent long disconnections occur, lingering objects cannot be replicated.

- **Upgrade all Windows 2000 Server domain controllers to Windows Server 2003.** This process requires upgrading the forest schema by using the **adprep /forestprep** command. Thereafter, you can begin upgrading domain controllers to Windows Server 2003. The Windows Server 2003 schema update adds 25 indexed attributes to the schema directory partition. An update of this size can cause replication delays in a large database. For this reason, domain controllers that are running Windows 2000 Server must be running — at a minimum — Windows 2000 Service Pack 2 (SP2) plus all additional Windows updates. However, it is highly recommended that you install Windows 2000 Service Pack 3 (SP3) on all domain controllers before preparing your infrastructure for upgrade to the Windows Server 2003 operating system. For information about upgrading to Windows Server 2003, see "Upgrading from Windows 2000 Domains to Windows Server 2003 Domains" in the Windows Server 2003 Deployment Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=46082).

- **Enable strict replication consistency on all domain controllers.** This registry setting, which stops inbound replication of a directory partition from a source domain controller that is suspected of having a lingering object, should be set for the forest to prevent the reintroduction of a lingering object into the directory. You can use the **Repadmin /regkey** command, which is available in the version of Windows Support Tools that is included with Windows Server 2003 SP1, to enable this setting on a specific domain controller or on all domain controllers in the forest, which eliminates the need to update the registry manually.

- **Monitor the Knowledge Consistency Checker (KCC) topology and replication to ensure that unintended long disconnections are detected.** By monitoring replication, you can detect disconnections that occur as a result of

network failures, service failures, or configuration errors. Use the Active Directory Management Pack or other monitoring application to implement a monitoring solution for your Active Directory deployment. Event IDs to monitor include 1311, 1388, 1925, 1988, 2042, 2087, and 2088.

- **Ship computers with properly configured IP, subnet mask, and default gateway addresses.** Remember to reconfigure the server with TCP/IP settings that are appropriate to the target site, not the staging site. Specifically, the domain controller must not point to itself for DNS.

- **Configure the tombstone lifetime appropriately.** Ensure that the tombstone lifetime is not lowered below the default. The default tombstone lifetime in a forest that was created on a domain controller running Windows 2000 Server or Windows Server 2003 is 60 days. The default tombstone lifetime in a forest that was created on a server running Windows Server 2003 with SP1 is 180 days. If you must disconnect a domain controller for a period of several weeks or months, before you disconnect the domain controller, do the following:

    - **Estimate the anticipated length of disconnection.**

    - **Determine the value of the tombstone lifetime for the forest.** This value is stored in the **tombstoneLifetime** attribute of CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,DC=*ForestRootDomain*.

    - **Determine the maximum length of time that the domain controller can be safely disconnected.** From the tombstone lifetime number of days, subtract a generous estimate of the number of days that are required for end-to-end replication latency. The resulting amount of time is the maximum period for which the domain controller can safely be disconnected.

    - **Determine whether to extend the tombstone lifetime for the forest.** If you estimate the maximum time of disconnection to be longer than the tombstone lifetime, you must determine whether to extend the tombstone lifetime or perform the procedure to remove lingering objects from the domain controller after it is reconnected. If you extend the tombstone lifetime, you must also make sure that all domain controllers have adequate disk space to store additional tombstones. In addition, make sure that replication of the tombstone lifetime change has reached all potential source domain controllers before you run Dcpromo to install an additional domain controller.

- **Prepare the registry for automatic nonauthoritative restart of SYSVOL when the domain controller restarts.** SYSVOL cannot be updated manually before disconnection. By editing a registry setting, you can ensure that SYSVOL is updated as soon as the domain controller is restarted.

- **Ensure that the domain controller replicates successfully with all replication partners.** Immediately before you disconnect the domain controller, force replication with its partners. Check that replication has succeeded before you disconnect the domain controller.

- **Label the domain controller.** When you disconnect the domain controller, attach a label to the computer that identifies the date and time of disconnection, the destination, and the IP settings.

- **When you reconnect the domain controller, restore SYSVOL as quickly as possible.** The domain controller does not serve as a domain controller until SYSVOL has been updated through replication. If the site has one or more other domain controllers in the same domain, start the domain controller anytime. If the site contains no other domain controller in the same domain, time the restart of the domain controller to coincide with the beginning of intersite replication.

- **Do not allow an outdated Windows 2000 Server domain controller to replicate.** If a domain controller running any version of Windows 2000 Server has been disconnected for longer than the maximum safe time of disconnection (the tombstone lifetime minus end-to-end replication latency), do not allow the domain controller to replicate. Instead, force the removal of Active Directory, perform metadata cleanup, and then reinstall Active Directory. As an alternative, you can reinstall the domain controller with Windows Server 2003. For more information about completing these tasks, see Reconnecting a Domain Controller After a Long-Term Disconnection.

| Note |
| --- |
| This recommendation applies to additional domain controllers in an existing domain. If the outdated domain controller is the only domain controller in the domain, the recommendation is to reconnect the domain controller and follow the instructions to remove lingering objects in article 314282, "Lingering objects may remain after you bring an out-of-date global catalog server back online," in the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=37924). |

- **To avoid time skew issues**, e**nsure that the system clock is synchronized with the domain source on startup.** When you start the domain controller in the remote site, use the following command to set the hardware clock:

  **net time /domain:***DomainName* **/set /y**

**See Also**

**Concepts**

Known Issues for Adding Domain Controllers in Remote Sites
Preparing a Server Computer for Shipping and Installation from Backup Media
Preparing an Existing Domain Controller for Shipping and Long-Term Disconnection
Reconnecting a Domain Controller After a Long-Term Disconnection
Managing SYSVOL

**Other Resources**

How the Active Directory Replication Model Works
Active Directory Management Pack Technical Reference for MOM 2005

# Preparing a Server Computer for Shipping and Installation from Backup Media

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The specific guidelines for installing Active Directory from backup media are provided in the topic Installing a Domain Controller in an Existing Domain Using Restored Backup Media. Be sure to read that topic before performing the procedures that are specified in this topic.

When you want to ship theserver to a remote site and install Active Directory by restoring from backup media in the remote site, you must make certain choices regarding the method that you use to restore the backup. You must also decide whether to use removable media or ship the backup on the server that will become the additional domain controller. You can use the information in this topic to make these decisions and to prepare the server for shipping. Use the information in Installing a Domain Controller in an Existing Domain Using Restored Backup Media to perform the actual backup, restore, and Active Directory installation procedures.

Preparing a computer for installation in a remote site by using restored backup media requires that you perform the following tasks:

- Begin by backing up system state on a domain controller in the domain into which you are installing the domain controller according to the recommendations and requirements in Installing a Domain Controller in an Existing Domain Using Restored Backup Media.

- Determine whether to restore the system state backup onto the computer that will be promoted or use removable media to ship the backup files separately from the computer.

- Determine the volume on which to restore the backup media. If you have a large Ntds.dit file, this decision can affect the amount of time necessary for Active Directory installation. If you have a large SYSVOL, this decision can affect whether full replication of SYSVOL occurs during Active Directory installation. The ability to use the backup media to source SYSVOL depends on various factors. If you want to avoid full replication of SYSVOL, additional preparation is required, as described later in this section.

- Before you ship the server, enable Remote Desktop access on the server so that you can install the domain controller and manage it remotely. You can also enable Remote Desktop remotely by using the registry, but this method should be used only as a fallback measure if, through some oversight, Remote Desktop is not enabled prior to shipping.

- If you are installing a domain controller that is running Windows Server 2003 with Service Pack 1 (SP1) in a forest that has a forest functional level of Windows Server 2003 or Windows Server 2003 interim and you want to include application directory partitions in the installation media, you can do so by creating an answer file that contains the location of the restored backup media and then running an unattended installation of Active Directory.

**Restore the Backup to the Promotion Computer or Ship Removable Media**

When you back up system state for the purpose of creating restored backup media for domain controller installation, you can use various methods to create the media for shipment and installation. You can:

- Before you ship the server, restore the backup directly to a volume on the server that you are shipping. When the server arrives at the remote site, it is ready for installation with no further preparation.

- Copy the .bkf file onto removable media before restoration. Ship the media to the remote site, and then restore the backup from the removable media to an alternate location on each domain controller that you want to install. The advantage of this method is that you retain the potential for SYSVOL to be sourced from the backup media.

- Restore the backup to any location on any server and then copy the restored backup to removable media, such as a CD, DVD, or portable hard drive. The advantage of using this method is that you restore the backup only once; you can install as many domain controllers as necessary from the same media. The disadvantage is that copying the restored files loses the SYSVOL data that is required for sourcing SYSVOL from the restored backup. For more information about ensuring

that SYSVOL is sourced from the restored backup, see "Seeding the SYSVOL tree from restored files during IFM promotion" in article 311078, "How to use the Install from Media feature to promote Windows Server 2003–based domain controllers," on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=37924).

**Determine the Restore Volume**

The volume on which you restore the system state backup has implications for both Active Directory files and SYSVOL files. For faster restore, it is recommended that you restore the backup to the volume that you will designate to host the Ntds.dit file when you run Dcpromo, if space permits. Otherwise, restore the backup to a volume that has sufficient free space. Restoring the backup to the volume that will store Ntds.dit, as opposed to a different volume, affects how files are managed by the system during and after Active Directory installation, as follows:

- **Active Directory files.** The volume to which you restore the Ntds.dit and NTDS log files determines how long installation will take and whether you must delete copied files following installation:

  - If you restore the system state to a location on the same volume (drive letter) that will ultimately host the Ntds.dit and NTDS log files, when you designate the path for the Ntds.dit and NTDS log files during installation, the Active Directory Installation Wizard will move the Ntds.dit and NTDS log files from the restored location to their installed location. Moving the files is much faster than copying the files.

  - If you restore the system state to a different volume than the volume that will ultimately host the Ntds.dit and NTDS log files, the Active Directory Installation Wizard will copy the Ntds.dit and NTDS log files to their final location during installation. In the case of a large Ntds.dit file, the copy process can add significantly to the installation time. In this case, you must manually delete the remaining files and folders in the restored folder after a successful installation. As a best practice, we recommend that you always delete the folder that you use to receive the restored backup, regardless of whether files are copied or moved.

- **SYSVOL replication.** The volume to which you restore the system state backup also determines whether the File Replication service (FRS) can use the restored files as the source for SYSVOL on the new domain controller or whether FRS replicates a new copy of SYSVOL from a different domain controller in the domain. To be able to use the installation media as the source for the SYSVOL data, you must restore the system state backup to the same volume as the drive that you specify in the Active Directory Installation Wizard to host the SYSVOL tree. In addition, you must perform preliminary procedures to pre-seed the SYSVOL data on the installation server. Otherwise, the data will be sourced over the network from a domain controller that is in the same domain as the new domain controller.

  If you store the SYSVOL shared folder on a different volume from the Active Directory files, consider the effect of copying Active Directory files, as described earlier in this topic, as opposed to the effect of replicating the entire contents of the SYSVOL shared folder. If avoiding replication of the SYSVOL shared folder is a goal of the remote installation, restore the backup to a location that is on the same volume as the drive that will contain the SYSVOL share and perform the preliminary SYSVOL pre-seeding procedures.

  In addition to the requirement that you restore the backup to the volume where SYSVOL is to be installed, successful pre-seeding of SYSVOL from the installation media also requires that SYSVOL has been replicated. If only one domain controller is installed in the domain (SYSVOL has not replicated at least once between two domain controllers in the domain), you must configure another domain controller in the domain to prepare the SYSVOL before you perform the system state backup.

| **Note** |
| --- |
| It is recommended that you deploy at least two domain controllers in each domain for redundancy and failover. |

For complete instructions for how to ensure that SYSVOL is sourced from the restored backup, see "Seeding the SYSVOL tree from restored files during IFM promotion" in article 311078, "How to use the Install from Media feature to promote Windows Server 2003–based domain controllers," on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=70809). To assess the effect of replication, as opposed to additional configuration to source SYSVOL from the backup media, test both procedures in a lab environment that mirrors your production environment in terms of wide area network (WAN) speed and replication latency.

**Enable Remote Desktop**

You can use Remote Desktop to connect to the domain controller and manage it as if you were sitting at the console. Remote Desktop is disabled by default in Windows Server 2003 operating systems. To install Active Directory, you must have Domain Admins credentials in the domain into which you are adding the domain controller. This level of service administration might not be available in the remote site. In any case, you will want to be able to install and manage the domain controller from the hub site.

**Create a Domain Controller Installation Answer File**

If you want to include application directory partitions in the restored backup media that you use as the source for an Active Directory installation, you must create a domain controller installation answer file and perform an unattended Active Directory installation. Dcpromo uses the answer file for installation instructions, including the location of restored backup media and instructions to use these files as the source for the installation.

If you are installing a domain controller in a remote site that will also be a DNS server, you might want to include application directory partitions in the installation media rather than replicating them. You can include application directory partitions in the installation media if the following conditions apply:

- The forest has a functional level of Windows Server 2003 or Windows Server 2003 interim.
- The domain controller that you back up and the server that you are installing are both running Windows Server 2003 with SP1.
- For creating a DNS server, your forest uses Active Directory-integrated DNS (DNS zone data is stored in application directory partitions on DNS servers in the forest).
- The domain controller that you back up stores the application directory partitions that you want to include.

Instructions for performing this type of installation are included in this task.

**Task requirements**

The following tools are necessary to complete this task:

- Ntbackup.exe
- System Control Panel
- Dcpromo.exe
- Ref.chm on the Windows Server 2003 installation CD (for unattended installations only)

To complete this task, perform the following procedures:

1. Back up system state on a domain controller in the domain into which you are installing the additional domain controller. The following requirements apply for the backup domain controller and the target server:

    - The backup domain controller and target server must be running the same version of Windows Server 2003. For example, if the domain controller that you back up is running Windows Server 2003 with SP1, you cannot use this backup media to install Active Directory on a server that is running Windows Server 2003 with no service pack installed.
    - The backup domain controller and target server must be running on the same hardware platform (32-bit or 64-bit).
    - To install a domain controller that is a global catalog server, you must back up system state on a global catalog server.
    - To install a domain controller that is a DNS server (that is, a server that stores the DomainDNSZones and ForestDNSZones application directory partitions), you must back up system state on a DNS server that stores these directory partitions.

2. Restore system state to an alternate location. This location can be on the target server or in a different location, from which the backup files can be copied to removable media and then shipped to the remote site separately from the target server. Follow the guidelines described in "Determine the Restore Volume" earlier in this topic.

   As an alternative, you can copy the unrestored .bkf file to removable media and then ship the media to the remote site, where it can be restored to a location on the target server.

   When you restore, you must run Ntbackup on the server that has the alternate location. Therefore, if you are restoring to an alternate location that is not on the server on which the .bkf file is stored, before you run Ntbackup, do the following:

   a. Share the folder that contains the .bkf file.

   b. Map a connection to it from the computer on which you are running Ntbackup.

3. Enable Remote Desktop on the target server.

4. If you are installing a DNS server or a domain controller that will store any application directory partitions, Create an answer file for domain controller installation.

5. Ship the domain controller and any prepared removable media and answer file to the remote site. Ship these items separately and securely.

6. When the server is running in the remote site, install the domain controller as follows:

   - Create a Remote Desktop Connection to the remote server.

   - If you are installing a domain controller that does not require application directory partitions to be included in the installation, Install Active Directory from restored backup media.

   - If you are installing a domain controller that will be a DNS server or that requires other application directory partitions to be included in the installation media, perform the procedure to Include application directory partitions in an Active Directory installation from backup media.

7. If the domain controller is to be a DNS server, Install the DNS Server service after Active Directory has been installed.

**See Also**
**Concepts**
Installing a Domain Controller in an Existing Domain Using Restored Backup Media

# Back up system state

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Ntbackup.exe provides simple and advanced options for backing up Active Directory components. When you back up system state, you can choose to include or exclude system-protected boot files. System-protected boot files are not used for installations from restored backup media. When the backup file that you create is to be used for additional domain controller installations, you can clear the advanced option to back up system-protected files. Clearing this option decreases the size of the .bkf file, as well as the time required to back up, restore, and copy the system state files.

Use these procedures to back up the system state only. These procedures do not back up the system disk or any other data on the domain controller except for the system-protected files.

Use the first procedure, "To back up system state including system-protected files," for routine system state backup. Use the second procedure, "To back up system state excluding system-protected files," if you want to create a smaller backup that is effective for installing domain controllers from restored backup media.

| **Note** |
| --- |
| To back up system state, you must log on locally to the domain controller, or Remote Desktop must be enabled on the remote domain controller so that you can connect remotely. |

### Administrative credentials

- To perform the following two procedures, you must be a member of the Domain Admins group or a member of the Backup Operators group.

#### To back up system state including system-protected files

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.

   This procedure provides steps for backing up in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard.

2. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
3. Select **Back up files and settings**, and then click **Next**.
4. Select **Let me choose what to back up**, and then click **Next**.
5. In the **Items to Back Up** window, double-click **My Computer**.
6. In the expanded list below **My Computer**, check **System State**, and then click **Next**.
7. Select a location to store the backup:

   - If you are backing up to a file, type the path and file name for the backup (.bkf) file (or click **Browse** to find a folder or file).

   - If you are backing up to a tape unit, choose the tape that you want to use.

     | **Note** |
     | --- |
     | You should not store the backup on the local hard drive. Instead, store it in a location, such as a tape drive, away from the computer that you are backing up. |

8. Type a name for this backup according to the recommendations in Backing Up Active Directory Components, and then click **Next**.
9. On the last page of the wizard, click **Advanced**.

10. Do not change the default options for **Type of Backup**. **Normal** should be selected, and the check box for **Backup migrated remote storage data** should remain cleared. Click **Next**.
11. Select **Verify data after backup**, and then click **Next**.
12. In the **Backup Options** dialog box, select a backup option, and then click **Next**.
13. If you are replacing the existing backups, select the option to allow only the owner and administrator access to the backup data and to any backups that are appended to this medium, and then click **Next**.
14. In the **When to back up** box, select the appropriate option for your needs, and then click **Next**.
15. If you are satisfied with all of the options that are selected, click **Finish** to perform the backup operation according to your selected schedule.

| **Note** |
| --- |
| The system state can also be backed up by using Ntbackup from a command line with appropriate parameters. For more information, at a command prompt type **ntbackup /?**. |

The following procedure produces a smaller .bkf file that does not include system boot files. By using this procedure, you can reduce the time that is required to perform the backup and subsequent restore, as well as the amount of disk space that is required. This method is recommended when the restored backup is to be used for installing additional domain controllers.

**To back up system state excluding system-protected files**

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.
2. On the **Welcome to the Backup or Restore Wizard** page, click **Advanced Mode**, and then click the **Backup** tab.
3. In the console tree, select the **System State** check box.
4. In **Backup media or file name**, type a name for this backup according to the recommendations in Backing Up Active Directory Components.
5. Click **Start Backup**, and then click **Advanced**.
6. Clear the **Automatically back up System Protected Files with the System State** check box, and then click **OK**.
7. Click **Start Backup**.

**See Also**
**Concepts**
Enable Remote Desktop
Create a Remote Desktop Connection

# Restore system state to an alternate location

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Perform this procedure to create media for installing a domain controller from restored backup media or to allow an authoritative restore of SYSVOL. After the domain controller installation is complete, delete the files in the alternate location.

You can restore the system state backup to an alternate location on the domain controller from which the backup was made, a location on another computer, or a location on the computer that you want to install as a domain controller.

**Administrative credentials**

To perform this procedure, you must be a member of the Backup Operators group, as follows:

- Restore system state on a member or workgroup server: Backup Operators group on the local computer
- Restore system state on a domain controller: Backup Operators group in the domain

**To restore system state to an alternate location**

1. Log on to the server that has the alternate location to which you are restoring system state backup files.
2. Click **Start**, click **Run**, type **ntbackup**, and then click **OK**.
3. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
4. Click **Restore Files and settings**, and then click **Next**.
5. On the **What to Restore** page, click **Browse**, and then, in the **Open Backup File** dialog box, click **Browse** again.
6. Navigate to the .bkf file that you want to restore to an alternate location. The .bkf file can be located in a folder on the current computer, in a shared folder on the backup computer or other network computer, or on an external drive that contains removable media.
7. In the **Select file for catalog** dialog box, click the .bkf file that you want to restore, and then click **Open**.
8. In the **Open Backup File** dialog box, click **OK**.
9. In **Items to restore**, double-click **File**, and then double-click the .bkf file that you want to restore.
10. Below the .bkf file that you want to restore, select the **System State** check box, and then click **Next**. (You do not need to restore the system disk to an alternate location.)
11. On the **Completing the Backup or Restore Wizard** page, click **Advanced**.
12. In the **Restore Files to** drop-down list, click **Alternate Location**.
13. In **Alternate Location**, type the path (or browse) to the local folder to which you are restoring the backup, and then click **Next**. We recommend restoring to a folder named NTDSRESTORE, if space permits, on the volume that will host the Ntds.dit file when Dcpromo is run. Otherwise, restore to a folder named \NTDSRESTORE on another volume that has sufficient free space.
14. On the **How to Restore** page, accept the default selection **Leave existing files (Recommended)**, and then click **Next**.
15. On the **Advanced Restore Option**s page, accept the default selections **Restore security settings** and **Preserve existing volume mount points**, and then click **Next**.
16. On the **Completing the Backup or Restore Wizard** page, click **Finish**.

# Enable Remote Desktop

Updated: February 9, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can enable Remote Desktop on the server that you are installing as a domain controller so that service administrators can manage the domain controller remotely. Remote Desktop is disabled by default in Windows Server 2003 operating systems.

Use this procedure to enable remote desktop prior to shipping the server that will be installed as a domain controller. If you neglected to perform this procedure prior to shipping the server, use the procedure "To enable Remote Desktop remotely by using the registry," later in this topic

**Administrative credentials**

To complete this procedure, you must be a member of the local Administrators group.

**To enable Remote Desktop**

1. Click **Start**, point to **Control Panel**, and then click **System**.
   Or
   Click **Start**, right-click **My Computer**, and then click **Properties**.
2. On the **Remote** tab, under **Remote Desktop**, select the **Allow users to connect remotely to this computer** check box, and then click **OK**.

| **Note** |
| --- |
| On computers running Windows Server 2003 with Service Pack 1 (SP1), on the **Remote** tab, select the **Enable Remote Desktop on this computer** check box. |

If for any reason you neglected to perform this procedure prior to shipping the server, you can enable Remote Desktop remotely by using the registry.

**Administrative credentials**

To complete this procedure, you must be a member of the local Administrators group.

**To enable Remote Desktop remotely by using the registry**

1. On any computer that is running a version of Windows Server 2003 or Windows XP Professional, click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. On the **File** menu, click **Connect Network Registry**.
3. In the **Select Computer** dialog box, type the computer name and then click **Check Names**.
4. In the **Enter Network Password** dialog box, provide Domain Admins credentials for the domain of the server, and then click **OK**.
5. After the computer name resolves, click **OK**.
6. In the computer node that appears in the Registry Editor, navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Terminal Server**.
7. In the console tree, click **Terminal Server** and then, in the details pane, double-click **fDenyTSConnections**.
8. In the **Edit DWORD Value** box, in **Value data**, type **0**, and then click **OK**.
9. To implement the change, reboot the server remotely, as follows:

   Open a command prompt, type the following, and then press Enter:

   **shutdown -m ** *DomainControllerName* **-r**

# Create an answer file for domain controller installation

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to create a text file that you can use as the answer file for an unattended installation of a domain controller. The answer file contains sensitive information and should be kept in a secure location.

**Administrative credentials**

To perform this procedure, you must be a member of the Authenticated Users group on the local computer on which you create the answer file.

**To create an answer file for domain controller installation**

1. On a local computer, insert the Windows Server 2003 CD-ROM into the CD-ROM drive or DVD-ROM drive. Press and hold down the SHIFT key as you insert the CD to prevent it from starting automatically.
2. Start Windows Explorer, and then open the Support\Tools folder on the Windows Server 2003 CD-ROM.
3. In the console tree, click **Tools**, and then, in the details pane, double-click **Deploy.cab**.
4. In the details pane, right-click **Ref.chm**, and then click **Extract**.
5. In the **Select a Destination** dialog box, navigate to or create a new folder for the expanded Ref.chm file, and then click **Extract**.
6. In its extracted location, open Ref.chm.
7. On the **Contents** tab in the scope pane, double-click **Unattend.txt**, and then click **[DCInstall]**.
8. In the details pane, scroll to **Sample**, select the entire sample, beginning at **[DCInstall]**, and then copy the sample.
9. Open Notepad, paste the sample into the Notepad file, and save the text file.
10. Edit the text file to contain at least the following entries. For additional examples and parameters, see Ref.chm and article 223757 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=155037).

    **[DCINSTALL]**

    **UserName=** *SAM account name that has Domain Admins credentials in the target domain. This account must be used by the administrator who runs the Dcpromo command.*

    **Password=** *Password for the account name. If you leave this blank, Dcpromo prompts the user during installation. Dcpromo deletes this value following installation.*

    **UserDomain=** *Domain name for the user account in* **UserName** .

    **DatabasePath=** *Location of the Ntds.dit file. (The default is %systemroot%*\ntds*.) If you omit this entry, Dcpromo uses the default location.*

    **LogPath=** *Location of the database log files. (The default is %systemroot%*\ntds*.) If you omit this entry, Dcpromo uses the default location.*

    **SYSVOLPath=** *Location of the shared SYSVOL tree. (The default is %systemroot%*\SYSVOL*.)* If you omit this entry, Dcpromo uses the default location.

    **SafeModeAdminPassword=** *Password for the administrator account that must be used to start the domain controller in Directory Services Restore Mode.* If you leave this blank, Dcpromo prompts the user for the password during installation. Dcpromo deletes this value following installation. Passwords are removed from the answer file when Dcpromo is executed.

    **CriticalReplicationOnly=Yes** *or* **no** *, to specify whether to skip noncritical portions of replication and allow Dcpromo to complete before replication is complete.*

    **SiteName=** *The name of the Active Directory site in which this domain controller will be placed.* This site must be created in advance in the Active Directory Sites and Services snap-in.

    **ReplicaOrNewDomain=** *Specify either* **Replica** *for an additional domain controller in an existing domain or*

**NewDomain** *for the first domain controller in a new domain.*

**ReplicaDomainDNSName=** *The fully qualified domain name of the domain of the new domain controller.*

**ReplicationSourceDC=** *The name of an existing domain controller in the domain to use as the source replication partner during installation.* When you install Active Directory from restored backup media, you can use this entry in conjunction with **ReplicateFromMedia** if you want to specify the domain controller from which Active Directory changes and SYSVOL changes are replicated.

**ReplicateFromMedia=Yes** *for installation from media,* **no** *for installation by replication.*

| **Note** |
| --- |
| For more information on this option, see article 311078 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=128859). |

**ReplicationSourcePath=** *When the installation is by replication, the path to the installation CD or network share. When the installation is from restored backup media, the local drive and path to the restored backup files.*

**RebootOnSuccess=Yes** *if you want the domain controller to restart automatically following a successful installation,* **no** *if you want to restart the domain controller manually. If you do not want the domain controller to restart automatically and you do not want to be prompted, use the value* **NoAndNoPromptEither**.

**ApplicationPartitionsToReplicate=** *Comma-separated distinguished names, with the entire string enclosed in quotation marks, of application directory partitions that you want to include when you use restored backup media to install Active Directory (or * to include all application directory partitions).* Using this entry requires Windows Server 2003 with Service Pack 1 (SP1) and Windows Server 2003 forest functional level. For more information about using this entry, see Include application directory partitions in an Active Directory installation from backup media.

| **Note** |
| --- |
| For additional examples and parameters, see Ref.chm and article 223757 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=155037). |

11. Save the answer file to the location on the installation server from which it is to be called by Dcpromo, or save the file to a network shared folder or removable media for distribution.

**See Also**
**Concepts**
Include application directory partitions in an Active Directory installation from backup media

# Create a Remote Desktop Connection

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If Remote Desktop is enabled on a server, you can use Remote Desktop Connection to connect to the server and manage it remotely. Remote Desktop is disabled by default in Windows Server 2003 operating systems.

**Administrative credentials**

To complete this procedure, you must have Remote Desktop permissions by being added to the Remote Desktop Users group or you must be a member of the local Administrators group of the computer to which you are connecting. If the computer is a domain controller, you must have the **Allow Logon Locally** right applied in the Default Domain Controllers Policy.

**To create a new Remote Desktop Connection**

1. On the **Start** menu, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
2. In **Computer**, type a computer name or Internet Protocol (IP) address, and then click **Connect**. The computer can be a terminal server, or it can be a computer running Windows XP Professional or a Windows Server 2003 operating system that has Remote Desktop enabled and for which you have Remote Desktop permissions.
3. In the **Log On to Windows** dialog box, type your user name, password, and domain (if required), and then click **OK**.

**See Also**
**Concepts**
Enable Remote Desktop

# Install Active Directory from restored backup media

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to install Active Directory from backup media to create an additional domain controller in an existing domain.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain into which you are installing the additional domain controller.

**To install Active Directory from restored backup media**

1. Click **Start**, click **Run**, type **dcpromo /adv**, and then press ENTER.
2. In the Active Directory Installation Wizard, select **Additional domain controller for existing domain**.
3. Select **From these restored backup files**, and point to the same location where you restored the system state data.
4. If the domain controller whose system state backup you are using is a global catalog server, the Active Directory Installation Wizard asks you whether you want this server to also be a global catalog server.
5. Give appropriate credentials for the operation.
6. Enter the domain of the new domain controller. This domain must be the domain of the domain controller whose system state backup you are using.
7. Complete the remaining pages of the Active Directory Installation Wizard.

   Dcpromo.exe will install Active Directory using the data present in the restored files, which eliminates the need to replicate every object from a partner domain controller. However, objects that were modified, added, or deleted since the backup was taken must be replicated. If the backup was recent, the amount of replication required will be considerably less than that required for a regular Active Directory installation.

   After the installation operation completes successfully and the computer is restarted, the folder and subfolders that contain the restored system state can be removed from the local disk.

**See Also**
**Concepts**
Restore system state to an alternate location
Include application directory partitions in an Active Directory installation from backup media

# Include application directory partitions in an Active Directory installation from backup media

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to install Active Directory from restored backup media that includes application directory partitions to create an additional domain controller in an existing domain. In this procedure, you edit a domain controller installation answer file to provide instructions for including application directory partitions in the installation. You must have already created the answer file according to the directions in Create an answer file for domain controller installation.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain into which you are installing the additional domain controller.

**To include application directory partitions in an Active Directory installation from backup media**

1. Open the answer file that you created to install the domain controller.
2. To include application directory partitions that are contained in the system state backup, add the following entry to the end of the answer file:

   **ApplicationPartitionsToReplicate=**

3. Provide a value for **ApplicationPartitionsToReplicate** as follows:

   - If you want to include all application directory partitions, use the value *****.

   - If you want to include specific application directory partitions, type the distinguished name of each directory partition. Separate each distinguished name with a space, and enclose the entire list in quotation marks, as shown in the following example:

     **ApplicationPartitionsToReplicate="dc=app1,dc=contoso,dc=com dc=app2,dc=contoso,dc=com"**

4. In the entry **ReplicationSourcePath=**, type the path to the folder that contains the restored system state backup files on the installation computer.
5. If you do not want Dcpromo to prompt the user for passwords, type the password in the **Password=** entry for the account that you will use to install the domain controller, type the password in the **SafeModeAdminPassword=** entry that you will use to provide access to Directory Services Restore Mode, and then save the answer file.

   | **Note** |
   | --- |
   | Passwords are automatically deleted from the answer file when Dcpromo runs. |

6. Open a command prompt, and then change directories to the location of the answer file.
7. At the command prompt, type the following command, and then press ENTER:

   **dcpromo /adv /answer:"** *Drive* **:** *PathToAnswerFile* **** *AnswerFileName* **"**

   Active Directory installation occurs automatically. If you left passwords blank in the answer file, Dcpromo prompts you for your administrative password and for the Directory Services Restore Mode password. If you specified "no" for the **RebootOnSuccess=** entry in the answer file, Dcpromo prompts you to restart the server after installation.

# Preparing an Existing Domain Controller for Shipping and Long-Term Disconnection

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you ship a domain controller to a remote site, you must disconnect it from the network and, consequently, from the replication topology. If a domain controller must be separated from the replication topology for a period of time that might be longer than a tombstone lifetime, you must take preliminary steps to ensure a smooth reconnection. Otherwise, it is possible that a long-term disconnection can result in a deleted object being reintroduced into the directory. Such deleted objects, when they are retained on a domain controller that has been disconnected for a period that is longer than a tombstone lifetime, are called "lingering objects." Lingering objects that are security principals, such as users or groups, can cause problems with Active Directory searches and e-mail delivery. Lingering objects can also jeopardize security if a user is allowed access to a resource by virtue of membership in a group that has been deleted. For more information about lingering objects, see "Maintaining Directory Consistency When Disconnecting a Domain Controller" in Known Issues for Adding Domain Controllers in Remote Sites.

By taking preliminary precautions, you can ensure that long-term disconnections do not result in directory inconsistency from lingering objects.

To complete this task, perform the following procedures:

1. Determine the anticipated length of the disconnection.

2. Determine the tombstone lifetime for the forest.

3. Determine the maximum safe disconnection period by subtracting a generous estimate of the end-to-end replication latency from the tombstone lifetime. Either find the latency estimate in the design documentation for your deployment or request the information from a member of your design or deployment team.

   - If the anticipated time of disconnection exceeds the maximum safe disconnection period, make a decision about whether to extend the tombstone lifetime. To change the tombstone lifetime, see Determine the tombstone lifetime for the forest and change the value in the **tombstoneLifetime** attribute.

   - If the estimated time of disconnection does not exceed the maximum safe disconnection time, proceed with disconnection.

4. View the current operations master role holders to determine whether the domain controller is an operations master role holder.

5. Transfer the domain-level operations master roles, if appropriate.

6. Transfer the schema master, if appropriate.

7. Transfer the domain naming master, if appropriate.

8. Prepare a domain controller for nonauthoritative SYSVOL restart on the domain controller that you are disconnecting. This process prevents the domain controller from having to reconcile and process deletions and modifications that took place from the time of the last SYSVOL update to the time the domain controller is restarted in the new site, and improves synchronization time.

9. Enable strict replication consistency on the domain controller that you are disconnecting. You can use this command-line procedure as an option to enable strict replication consistency on additional other domain controllers or on all domain controllers in the forest.

10. Synchronize replication with all partners. Update the domain controller with the latest changes just before you

disconnect it.

11. Verify successful replication to a domain controller for the domain controller that you are disconnecting.

12. Label the domain controller with the date and time of disconnection and the maximum safe disconnection period.

**See Also**

**Concepts**

Known Issues for Adding Domain Controllers in Remote Sites
Managing Operations Master Roles
Managing SYSVOL
Reconnecting a Domain Controller After a Long-Term Disconnection

**Other Resources**

Windows Server 2003 Technical Reference

# Determine the tombstone lifetime for the forest

Updated: November 12, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 Foundation, Windows Server 2008 R2, Windows Server 2008 R2 Foundation, Windows Server 2012

The tombstone lifetime in an Active Directory forest determines how long a deleted object (called a "tombstone") is retained in Active Directory Domain Services (AD DS). The tombstone lifetime is determined by the value of the **tombstoneLifetime** attribute on the Directory Service object in the configuration directory partition.

You can use this procedure to determine the tombstone lifetime for the forest.

Membership in **Domain Users**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

**To determine the tombstone lifetime for the forest using ADSIEdit**

1. Click **Start**, point to **Administrative Tools**, and then click **ADSI Edit**.
2. In ADSI Edit, right-click **ADSI Edit**, and then click **Connect to**.
3. For **Connection Point**, click **Select a well known Naming Context**, and then click **Configuration**.
4. If you want to connect to a different domain controller, for **Computer**, click **Select or type a domain or server: (Server | Domain [:port])**. Provide the server name or the domain name and Lightweight Directory Access Protocol (LDAP) port (389), and then click **OK**.
5. Double-click **Configuration**, **CN=Configuration,DC=***ForestRootDomainName*, **CN=Services**, and **CN=Windows NT**.
6. Right-click **CN=Directory Service**, and then click **Properties**.
7. In the **Attribute** column, click **tombstoneLifetime**.
8. Note the value in the **Value** column. If the value is **<not set>**, the value is 60 days.

**To determine the tombstone lifetime for the forest using Dsquery**

1. Open a Command Prompt window. To open a command prompt, click **Start**, click **Run**, type **cmd**, and then press ENTER.
2. At the command prompt, type the following command, and then press ENTER:

   ```
   dsquery * "cn=directory service,cn=windows nt,cn=services,cn=configuration,dc=
   <forestDN>" -scope base -attr tombstonelifetime
   ```

   Be sure to replace `<forestDN>` with the actual distinguished name of the forest. For example, if your forest name is corp.proseware.com, type the following, and then press ENTER:

   ```
   dsquery * "cn=directory service,cn=windows
   nt,cn=services,cn=configuration,dc=corp,dc=proseware,dc=com" -scope base -attr
   tombstonelifetime
   ```

# View the current operations master role holders

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Once an operations master role has been transferred, it should be verified that the transfer has occurred successfully throughout the domain. The change must be replicated to all relevant domain members in order to truly take effect.

To view the current operations master role holders, use Ntdsutil.exe with the roles option. This option displays a list of all current role holders.

**Administrative Credentials**

To perform this procedure, you must be logged on as a User or an Administrator.

**To view the current operations master role holder**

1. Click **Start**, click **Run**, type **ntdsutil**, and then press ENTER.
2. At the **ntdsutil:** prompt, type **roles** and press ENTER.
3. At the **fsmo maintenance:** prompt, type **connections** and press ENTER.
4. At the **server connections:** prompt, type **connect to server** *servername* (where *servername* is the name of the domain controller that belongs to the domain containing the operations masters).
5. After receiving confirmation of the connection, type **quit** and press ENTER to exit this menu.
6. At the **fsmo maintenance:** prompt, type **select operation target** and press ENTER.
7. At the **select operations target:** prompt, type **list roles for connected server** and press ENTER.

   The system responds with a list of the current roles and the Lightweight Directory Access Protocol (LDAP) name of the domain controllers currently assigned to host each role.
8. Type **quit** and press ENTER to exit each prompt in Ntdsutil.exe. Type **quit** and press ENTER at the **ntdsutil:** prompt to close the window.

# Transfer the domain-level operations master roles

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to transfer the three domain-level operations master roles: the PDC emulator, the RID master, and the infrastructure master. You can transfer all of these roles by using the Active Directory Users and Computers console.

| **Note** |
| --- |
| These procedures are performed by using MMC, although you can also transfer these roles by using Ntdsutil.exe. For information about using Ntdsutil.exe to transfer the operations master roles, type **?** at the Ntdsutil.exe command prompt. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To transfer a domain-level operations master role**

1. Open Active Directory Users and Computers.
2. At the top of the console tree, right-click **Active Directory Users and Computers**. Click **Connect to Domain Controller**.
3. In the list of **available domain controllers**, click the name of the server to which you want to transfer the role, and then click **OK**.
4. At the top of the console tree, right-click **Active Directory Users and Computers**, point to **All Tasks**, and then click **Operations Masters**.

   The name of the current operations master role holder appears in the **Operations master** box. The name of the server to which you want to transfer the role appears in the lower box.
5. Click the tab for the role you want to transfer: **RID**, **PDC**, or **Infrastructure**. Verify the computer names that appear and then click **Change**. Click **Yes** to transfer the role, and then click **OK**.
6. Repeat steps 4 and 5 for each role that you want to transfer.

# Transfer the schema master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to transfer the schema operations master role. The schema master is a forest-wide operations master role. Before you can use the Active Directory Schema snap-in for the first time, you must register it with the system. If you have not yet prepared the Active Directory Schema snap-in, see Install the Schema snap-in before you begin this procedure.

| **Note** |
| --- |
| This procedure is performed by using the Microsoft Management Console (MMC), although you can also transfer this role by using Ntdsutil.exe. For information about using Ntdsutil.exe to transfer operations master roles, type **?** at the Ntdsutil.exe command prompt. |

### Administrative Credentials

To perform this procedure, you must be a Schema Administrator in Active Directory.

#### Transfer the schema master

1. Open the Active Directory Schema snap-in.
2. In the console tree, right-click **Active Directory Schema**, and click **Change Domain Controller**.
3. In the **Change Domain Controller** dialog box, click **Specify Name**. Then, in the text box, type the name of the server to which you want to transfer the schema master role. Click **OK**.
4. In the console tree, right-click **Active Directory Schema**. Click **Operations Master**. The **Change Schema Master** box displays the name of the server that is currently holding the role. The targeted domain controller is listed in the second box.
5. Click **Change**. Click **Yes** to confirm your choice. The system confirms the operation. Click **OK** again to confirm that the operation succeeded.
6. Click **Close** to close the **Change Schema Master** dialog box.

| **Note** |
| --- |
| Hosting the infrastructure master on a global catalog server is not recommended. If you attempt to transfer the infrastructure master role to a domain controller that is a global catalog, the system displays a warning stating that this is not recommended. |

# Transfer the domain naming master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to transfer the domain naming operations master role. The domain naming master is a forest-wide operations master role.

| **Note** |
| --- |
| This procedure is performed by using the Microsoft Management Console (MMC), although you can also transfer this role by using Ntdsutil.exe. For information about using Ntdsutil.exe to transfer operations master roles, type **?** at the Ntdsutil.exe command prompt. |

### Administrative Credentials

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

#### To transfer the domain naming master

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click **Active Directory Domains and Trusts**, and then click **Connect to Domain Controller**.
3. Ensure that the proper domain name is entered in the **Domain** box.

   The available domain controllers from this domain are listed.
4. In the **Name** column, click the domain controller (to select it) to which you want to transfer the role. Click **OK**.
5. Right-click **Active Directory Domains and Trusts**, and then click **Operations Master**.
6. The name of the current domain naming master appears in the first text box. The server to which you want to transfer the role should appear in the second text box. If this is not the case, repeat steps 1 through 4.
7. Click **Change**. To confirm the role transfer, click **Yes**. Click **OK** again to close the message box indicating the transfer took place. Click **Close** to close the **Change Operations Master** dialog box.

# Prepare a domain controller for nonauthoritative SYSVOL restart

Updated: April 15, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Initiate a nonauthoritative restore of SYSVOL by modifying the value of the BurFlags (backup/restore flags) registry entry. Changing the value to D2 (hexadecimal) or 210 (decimal) prior to disconnecting a domain controller initiates an automatic nonauthoritative restore of SYSVOL when the domain controller is restarted. This process results in a smaller database size and improved synchronization time because the domain controller does not reconcile and process every deletion and modification that took place from the time of the last SYSVOL update to the time the domain controller is restarted in the new site, thereby improving synchronization time. Instead, it processes only the latest changes.

Separate entries exist for global and replica-set-specific BurFlags, as follows:

- To initiate a nonauthoritative restart of SYSVOL when it is the only replica set that is represented on the domain controller, set the value of the global BurFlags (REG_DWORD) entry under

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters\Backup/Restore\Process at Startup**

- If other replica sets are represented on the domain controller and you want to update only SYSVOL, set the value of the replica-set-specific BurFlags (REG_DWORD) entry under

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters\Cumulative Replica Sets\SYSVOL GUID**

Modifying the replica-set-specific BurFlags entry requires identifying the SYSVOL GUID in the registry.

| ⚠**Caution** |
| --- |
| The Registry Editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back up system state first. For information about backing up system state, see Administering Active Directory Backup and Restore. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To prepare a domain controller for nonauthoritative SYSVOL restart**

1. Click **Start**, click **Run**, type **regedit** and then click **OK**.
2. Navigate to

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtFrs\Parameters**
3. Expand **Parameters**.
4. Modify one of the BurFlags entries as follows:

   To modify the global BurFlags entry:

   a. Expand **Backup/Restore** and then click **Process at Startup**.

   b. In the details pane, right-click **BurFlags** and click **Modify**.

   c. In the **Value data** box, type **D2** hexadecimal or **210** decimal.

   d. Click **OK** and close Registry Editor.

To modify the replica-set-specific BurFlags entry:

a. Expand both **Cumulative Replica Sets** and **Replica Sets**.

b. Match the GUID under **Replica Sets** to the identical GUID under **Cumulative Replica Sets**, and click the matching GUID under **Cumulative Replica Sets**.

c. In the details pane, right click **BurFlags** and click **Modify**.

d. In the **Value data** box, type **D2** hexadecimal or **210** decimal.

e. Click **OK** and close Registry Editor.

# Enable strict replication consistency

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The setting for replication consistency is stored in the registry in the **Strict Replication Consistency** entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

Values for this entry are as follows:

- Value: **1** (**0** to disable)
- Default: **1** (enabled) in a new Windows Server 2003 forest; otherwise **0**.
- Data type: **REG_DWORD**

On domain controllers running Windows Server 2003 with Service Pack 1 (SP1), you do not have to edit the registry directly to enable strict replication consistency. It is best to avoid editing the registry directly if possible. You can use a Repadmin command that enables strict replication consistency on one or all domain controllers in the forest. This command is available only in the version of Repadmin that is included with Windows Support Tools in Windows Server 2003 SP1. This command can be applied only on domain controllers running Windows Server 2003 with SP1.

**Administrative credentials**

- To complete this procedure on a single domain controller, you must be a member of the Domain Admins group in the domain.
- To complete this procedure on all domain controllers, you must be a member of the Enterprise Admins group in the forest.

**Requirements:**

- Operating system: Windows Server 2003 with SP1

| **Note** |
| --- |
| To enable strict replication consistency on a domain controller that is not running Windows Server 2003 with SP1, use a registry editor to set the value in the **Strict Replication Consistency** entry to **1**. |

| **Caution** |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

**To enable strict replication consistency**

1. Open a command prompt, type the following command, and then press ENTER:

   **repadmin /regkey** *DC_LIST* **{+|-}** *key*

| Term | Definition |
| --- | --- |
| *DC_LIST* | The name of a single domain controller. (***** applies the change to all domain controllers in the forest.) For the domain controller name, you can use the Domain Name System (DNS) name, the distinguished name of the domain controller computer object, or the distinguished name of the domain controller server object. |

| | |
|---|---|
| **{+\|-}** *key* | **+** to enable and **-** to disable, and *key* is strict. For example, **+strict** enables strict replication consistency; **-strict** disables it. |

2. Repeat step 1 for every domain controller on which you want to enable strict replication consistency.

| **Note** |
|---|
| For more naming options and information about the syntax of the *DC_LIST* parameter, at the command prompt type **repadmin /listhelp**. |

# Synchronize replication with all partners

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to synchronize replication with all replication partners of a domain controller.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the selected domain controller or the Enterprise Admins group in the forest, or you must have been delegated the appropriate authority. If you want to synchronize the configuration and schema directory partitions on a domain controller in a child domain, you must have Domain Admins credentials in the forest root domain or Enterprise Admins credentials in the forest.

**To synchronize replication with all partners**

1. At a command prompt, type the following command, and then press ENTER:

   **repadmin /syncall** *DCName* **/e /d /A /P /q**

| Term | Definition |
|---|---|
| *DCName* | The Domain Name System (DNS) name of the domain controller on which you want synchronize replication with all partners |
| **/e** | Enterprise; includes partners in all sites. |
| **/d** | Identifies servers by distinguished name in messages. |
| **/A** | All; synchronizes all directory partitions that are held on the home server. |
| **/P** | Pushes changes outward from the home server. |
| **/q** | Runs in quiet mode; suppresses callback messages. |

2. Check for replication errors in the output of the command in the previous step. If there are no errors, replication is successful. For replication to complete, any errors must be corrected.

**See Also**
**Concepts**
Verify successful replication to a domain controller

# Verify successful replication to a domain controller

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **repadmin /showrepl** command to verify successful replication to a specific domain controller. If you are not running Repadmin on the domain controller whose replication you are checking, you can specify a destination domain controller in the command. Repadmin lists INBOUND NEIGHBORS for the current or specified domain controller. INBOUND NEIGHBORS shows the distinguished name of each directory partition for which inbound directory replication has been attempted, the site and name of the source domain controller, and whether replication succeeded or not, as follows:

- Last attempt @ YYYY-MM-DD HH:MM.SS was successful.
- Last attempt @ [Never] was successful.

If @ [Never] appears in the output for a directory partition, replication of that directory partition has never succeeded from the identified source replication partner over the listed connection.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the destination domain controller.

**To verify successful replication to a domain controller**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl** *servername* **/u:** *domainname* **** *username* **/pw:***

   | Term | Definition |
   | --- | --- |
   | *servername* | Specifies the name of the destination domain controller. |
   | *domainname* | Specifies the single-label name of the domain of the destination domain controller. (You do not have to use a fully qualified Domain Name System (DNS) name.) |
   | *username* | Specifies the name of an administrative account in that domain. |

3. When you are prompted for a password, type the password for the user account that you provided, and then press ENTER.

You can also use Repadmin to generate the details of replication to and from all replication partners in a spreadsheet. The spreadsheet displays data in the following columns:

Showrepl_COLUMNS

Destination DC Site

Destination DC

Naming Context

Source DC Site

Source DC

Transport Type

Number of Failures

Last Failure Time

Last Success Time

Last Failure Status

The following procedure shows how to create this spreadsheet and set column headers for improved readability.

**To generate a repadmin /showrepl spreadsheet for all replication partners**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl * /csv >showrepl.csv**
3. Open Microsoft Excel.
4. On the **File** menu, click **Open**, navigate to showrepl.csv, and then click **Open**.
5. Hide or delete column **A** as well as the **Transport Type** column, as follows:
6. Select a column that you want to hide or delete.
   - To hide the column, on the **Format** menu, click **Column**, and then click **Hide**.

     Or

   - To delete the column, right-click the selected column, and then click **Delete**.
7. Select row 1 beneath the column heading row, and then, on the **Window** menu, click **Freeze Panes**.
8. Select the entire spreadsheet. On the **Data** menu, click **Filter**, and then click **AutoFilter**.
9. In the **Last Success Time** column, click the down arrow, and then click **Sort Ascending**.
10. In the **Source DC** column, click the down arrow, and then click **Custom**.
11. In the **Custom AutoFilter** dialog box, under **Show rows where**, click **does not contain**. In the adjacent text box, type **del** to eliminate from view the results for deleted domain controllers.
12. Repeat step 10 for the **Last Failure Time** column, but use the value **does not equal**, and type the value **0**.
13. Resolve replication failures.

The last successful attempt should agree with the replication schedule for intersite replication, or the attempt should be within the last hour for intrasite replication.

If Repadmin reports any of the following conditions, see Troubleshooting Active Directory Replication Problems:

- The last successful intersite replication was prior to the last scheduled replication.
- The last intrasite replication was longer than one hour ago.
- Replication was never successful.

**See Also**
**Concepts**
Troubleshooting Active Directory Replication Problems

# Reconnecting a Domain Controller After a Long-Term Disconnection

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Assuming that a domain controller has not been disconnected for longer than the maximum safe period for disconnection (tombstone lifetime minus end-to-end replication latency), reconnecting the domain controller to the replication topology requires no special procedures. By default, the Knowledge Consistency Checker (KCC) on a domain controller runs five minutes after the domain controller starts, automatically incorporating the reconnected domain controller into the replication topology.

**Reconnecting an Outdated Domain Controller**

If you plan appropriately for disconnecting and reconnecting domain controllers, no domain controller will be disconnected from the replication topology for longer than a tombstone lifetime. However, if unexpected events result in a domain controller becoming outdated, reconnect the domain controller as follows:

- The disconnected domain controller is running Windows Server 2003, and an authoritative domain controller running Windows Server 2003 is available in this site or a neighboring site: Reconnect the domain controller, and immediately follow the instructions in Use Repadmin to remove lingering objects.

- The disconnected domain controller is running Windows Server 2003, but no other authoritative domain controller running Windows Server 2003 is available in the domain: Reconnect the domain controller, and follow the instructions in article 314282, "Lingering objects may remain after you bring an out-of-date global catalog server back online," in the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=37924).

- The disconnected domain controller is running Windows 2000 Server, and another domain controller is available in the domain: Do not reconnect the domain controller. Instead, force Active Directory removal on the disconnected domain controller, perform metadata cleanup, and then reinstall Active Directory. To complete these tasks, follow the instructions in Forcing the Removal of a Domain Controller and Installing a Domain Controller in an Existing Domain.

- The disconnected domain controller is running Windows 2000 Server, and no other domain controller is available in the domain: If you want to recover the domain, reconnect the domain controller, and follow the instructions in article 314282, "Lingering objects may remain after you bring an out-of-date global catalog server back online," in the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=37924).

**Updating SYSVOL**

As described in Preparing an Existing Domain Controller for Shipping and Long-Term Disconnection, the recommended practice to ensure consistency of SYSVOL is to modify the registry before disconnecting the domain controller so that SYSVOL is updated automatically when the domain controller is restarted. In addition, if you want to avoid a full synchronization of SYSVOL through intersite replication, you must take preparatory steps before disconnection. For information about how to ensure that SYSVOL is sourced locally and updated over the network only for changes, see "Seeding the SYSVOL tree from restored files during IFM promotion" in article 311078, "How to use the Install from Media feature to promote Windows Server 2003-based domain controllers," in the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=37924). To update SYSVOL as soon as possible after reconnecting a domain controller, plan the time that you restart the domain controller to optimize the replication schedule, as follows:

- If the closest replication partner for the domain is in a different site, view site link properties to determine the replication schedule, and then restart the domain controller as soon as possible after replication is scheduled to start.

- If a replication partner for the domain is available within the site, verify replication success on that partner before restarting the domain controller.

**Important**

Do not use file copy utilities, such as Xcopy or Robocopy, to update an outdated SYSVOL. Copying SYSVOL files is recommended only for recreating a nonfunctioning SYSVOL, which requires several preliminary procedures. Copying SYSVOL files from one domain controller to another without following these procedures causes invalid data to be replicated and causes the system volumes on other domain controllers to become inconsistent. For information about how to recreate a nonfunctioning SYSVOL, see Restoring and Rebuilding SYSVOL.

To complete this task, perform the following procedures:

1. Determine the tombstone lifetime for the forest.

2. Determine whether the maximum safe disconnection time has been exceeded. The maximum safe disconnection time should have been established at the time of disconnection, as follows:

   Subtract a generous estimate of the amount of time for end-to-end replication latency from the tombstone lifetime. Either find the latency estimate in the design documentation for your deployment or request the information from a member of your design or deployment team.

3. If the maximum safe disconnection time has not been exceeded, proceed with the reconnection process as follows:

   - If the site in which you are reconnecting the domain controller has one or more other domain controllers that are authoritative for the domain, start the domain controller anytime.

   - If the site in which you are reconnecting the domain controller has no other domain controllers that are authoritative for the domain, proceed as follows:

     Determine when intersite replication is scheduled to begin by viewing the replication properties on the site link that connects this site to the next closest site that includes a domain controller that is authoritative for this domain.

     As soon as possible after the next replication cycle begins, start the domain controller.

   If the maximum safe disconnection time has been exceeded, proceed in the appropriate manner according to the operating system, as described in "Reconnecting an Outdated Domain Controller" earlier in this topic.

4. After replication is complete, Verify successful replication to a domain controller (the reconnected domain controller) of the domain, configuration, and schema directory partitions. If the domain controller is a global catalog server, check for successful replication of all domain directory partitions.

**See Also**
**Concepts**
Preparing an Existing Domain Controller for Shipping and Long-Term Disconnection

# Determine when intersite replication is scheduled to begin

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the properties on the site link object to determine when intersite replication between sites is scheduled to begin.

**Administrative Credentials**

To complete this procedure, you must be a member of the Domain Users group.

**To determine when intersite replication is scheduled to begin**

1. Open Active Directory Sites and Services.
2. In the console tree, double-click the **Sites** container, double-click the **Inter-Site Transports** container, and then click the **IP** container.
3. In the details pane, right-click the site link object for which you want to view the schedule, and then click **Properties**.
4. In the *SiteLinkName* **Properties** dialog box, click **Change Schedule**. Note the block of days and hours during which replication is allowed (**Replication Available**), and then click **OK**.
5. In the **Replicate every _____ minutes** box, note the number of minutes for the intervals at which replication polling takes place during an open schedule window, and then click **OK**.

# Use Repadmin to remove lingering objects

Updated: July 12, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use Repadmin to remove lingering objects when you reconnect a domain controller that has been offline for longer than a tombstone lifetime and you want to ensure that lingering objects do not exist or, if they do, that they are removed before they are replicated. You can also use this procedure when event ID 1388 or event ID 1988 is logged on a domain controller. In this case, the information that you need to perform the procedure is provided in the event. For information about removing lingering objects when event ID 1388 or event ID 1988 has been logged, see "Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042)."

If you are running the procedure preemptively, you must gather the following information before beginning the procedure:

- Name of the server that has or might have lingering objects. This name can be the Domain Name System (DNS) name, NetBIOS name, or distinguished name of the domain controller.
- Globally unique identifier (GUID) of the NTDS Settings object of a domain controller that is authoritative for the domain of the domain controller from which you want to remove lingering objects.

If necessary, use the following procedure to determine the GUID of a domain controller.

**Administrative Credentials**

To complete this procedure, you must be a member of the Domain Users group in the domain of the domain controller.

**To determine the GUID of a domain controller**

1. At a command prompt, type the following command, and then press ENTER:

   **repadmin /showrepl***DomainControllerName*

   where *DomainControllerName* is the NetBIOS name of the domain controller whose GUID you want to determine.

2. In the top portion of the output, note the value in **DC object GUID:**

If the destination domain controller and source domain controller are both running Windows Server 2003, you can remove lingering objects by using Repadmin. If either domain controller is running Windows 2000 Server, follow instructions in article 314282, "Lingering objects may remain after you bring an out-of-date global catalog server back online," in the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=37924).

**Administrative Credentials**

To complete this procedure, you must be a member of the Domain Admins group in the *DirectoryPartition* domain.

**To use Repadmin to remove lingering objects**

1. At a command prompt, type the following command, and then press ENTER:

| **Important** |
|---|
| There must be connectivity between the server that has or might have lingering objects and the server that has an up-to-date, writable replica of the directory partition. |

**repadmin /removelingeringobjects** *ServerName ServerGUID DirectoryPartition* **/advisory_mode**

| Term | Definition |
|---|---|
| *ServerName* | The DNS name or the distinguished name of the domain controller that has or might have lingering objects. |
| *ServerGUID* | The GUID of a domain controller that has an up-to-date writable replica of the directory partition |

| | |
|---|---|
| *DirectoryPartition* | The distinguished name of the domain directory partition that might have lingering objects. For example, DC=*RegionalDomainName*,DC=*ForestRootDomainName*,DC=com. Also run the command against the configuration directory partition (CN=configuration,DC=*ForestRootDomainName*,DC=com), the schema directory partition (CN=schema,CN=configuration,DC=*ForestRootDomainName*), and any application directory partitions that are hosted on the domain controller you are checking for lingering objects. |

**/advisory_mode** logs the lingering objects that will be removed so that you can review them, but it does not remove them.

2. If lingering objects are found, repeat step 1 without **/advisory_mode** to delete the identified lingering objects from the directory partition.
3. Repeat steps 1 and 2 for every domain controller that might have lingering objects.

| **Note** |
|---|
| The *ServerName* parameter uses the **DC_LIST** syntax for **repadmin**, which allows the use of ***** for all domain controllers in the forest and **gc:** for all global catalog servers in the forest. To see the **DC_LIST** syntax, type **repadmin /listhelp**. |

**See Also**
**Concepts**
Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042)

# Verify successful replication to a domain controller

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **repadmin /showrepl** command to verify successful replication to a specific domain controller. If you are not running Repadmin on the domain controller whose replication you are checking, you can specify a destination domain controller in the command. Repadmin lists INBOUND NEIGHBORS for the current or specified domain controller. INBOUND NEIGHBORS shows the distinguished name of each directory partition for which inbound directory replication has been attempted, the site and name of the source domain controller, and whether replication succeeded or not, as follows:

- Last attempt @ YYYY-MM-DD HH:MM.SS was successful.
- Last attempt @ [Never] was successful.

If @ [Never] appears in the output for a directory partition, replication of that directory partition has never succeeded from the identified source replication partner over the listed connection.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the destination domain controller.

**To verify successful replication to a domain controller**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl** *servername* **/u:** *domainname* **** *username* **/pw:***

   | Term | Definition |
   | --- | --- |
   | *servername* | Specifies the name of the destination domain controller. |
   | *domainname* | Specifies the single-label name of the domain of the destination domain controller. (You do not have to use a fully qualified Domain Name System (DNS) name.) |
   | *username* | Specifies the name of an administrative account in that domain. |

3. When you are prompted for a password, type the password for the user account that you provided, and then press ENTER.

You can also use Repadmin to generate the details of replication to and from all replication partners in a spreadsheet. The spreadsheet displays data in the following columns:

Showrepl_COLUMNS

Destination DC Site

Destination DC

Naming Context

Source DC Site

Source DC

Transport Type

Number of Failures

Last Failure Time

Last Success Time

Last Failure Status

The following procedure shows how to create this spreadsheet and set column headers for improved readability.

**To generate a repadmin /showrepl spreadsheet for all replication partners**

1. Open a Command Prompt.
2. Type the following command, and then press ENTER:

   **repadmin /showrepl * /csv >showrepl.csv**
3. Open Microsoft Excel.
4. On the **File** menu, click **Open**, navigate to showrepl.csv, and then click **Open**.
5. Hide or delete column **A** as well as the **Transport Type** column, as follows:
6. Select a column that you want to hide or delete.
   - To hide the column, on the **Format** menu, click **Column**, and then click **Hide**.

     Or

   - To delete the column, right-click the selected column, and then click **Delete**.
7. Select row 1 beneath the column heading row, and then, on the **Window** menu, click **Freeze Panes**.
8. Select the entire spreadsheet. On the **Data** menu, click **Filter**, and then click **AutoFilter**.
9. In the **Last Success Time** column, click the down arrow, and then click **Sort Ascending**.
10. In the **Source DC** column, click the down arrow, and then click **Custom**.
11. In the **Custom AutoFilter** dialog box, under **Show rows where**, click **does not contain**. In the adjacent text box, type **del** to eliminate from view the results for deleted domain controllers.
12. Repeat step 10 for the **Last Failure Time** column, but use the value **does not equal**, and type the value **0**.
13. Resolve replication failures.

The last successful attempt should agree with the replication schedule for intersite replication, or the attempt should be within the last hour for intrasite replication.

If Repadmin reports any of the following conditions, see Troubleshooting Active Directory Replication Problems:

- The last successful intersite replication was prior to the last scheduled replication.
- The last intrasite replication was longer than one hour ago.
- Replication was never successful.

**See Also**
**Concepts**
Troubleshooting Active Directory Replication Problems

# Performing an Unattended Installation of Active Directory

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Running an unattended install simplifies the process of setting up Active Directory on multiple computers. The unattended install feature uses an "answer file" to provide answers to the questions that are asked during a normal setup. This way, the installation process can proceed from start to completion without user intervention. This method works best when Active Directory is installed with identical options on many computers. This method is required if you want to include application directory partitions in Active Directory installations from restored backup media.

**Task requirements**

The following tool is required to complete this task:

- Dcpromo.exe

The following files are required to complete this task:

- Ref.chm (in the Support\Tools folder on the Windows Server 2003 operating system CD)
- Unattend.txt
- Domain controller answer file

To complete this task, perform the following procedures:

1. Create an answer file for domain controller installation
2. Install Active Directory using an answer file

**See Also**
**Concepts**
Include application directory partitions in an Active Directory installation from backup media

# Create an answer file for domain controller installation

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to create a text file that you can use as the answer file for an unattended installation of a domain controller. The answer file contains sensitive information and should be kept in a secure location.

**Administrative credentials**

To perform this procedure, you must be a member of the Authenticated Users group on the local computer on which you create the answer file.

**To create an answer file for domain controller installation**

1. On a local computer, insert the Windows Server 2003 CD-ROM into the CD-ROM drive or DVD-ROM drive. Press and hold down the SHIFT key as you insert the CD to prevent it from starting automatically.
2. Start Windows Explorer, and then open the Support\Tools folder on the Windows Server 2003 CD-ROM.
3. In the console tree, click **Tools**, and then, in the details pane, double-click **Deploy.cab**.
4. In the details pane, right-click **Ref.chm**, and then click **Extract**.
5. In the **Select a Destination** dialog box, navigate to or create a new folder for the expanded Ref.chm file, and then click **Extract**.
6. In its extracted location, open Ref.chm.
7. On the **Contents** tab in the scope pane, double-click **Unattend.txt**, and then click **[DCInstall]**.
8. In the details pane, scroll to **Sample**, select the entire sample, beginning at **[DCInstall]**, and then copy the sample.
9. Open Notepad, paste the sample into the Notepad file, and save the text file.
10. Edit the text file to contain at least the following entries. For additional examples and parameters, see Ref.chm and article 223757 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=155037).

    **[DCINSTALL]**

    **UserName=** *SAM account name that has Domain Admins credentials in the target domain. This account must be used by the administrator who runs the Dcpromo command.*

    **Password=** *Password for the account name. If you leave this blank, Dcpromo prompts the user during installation. Dcpromo deletes this value following installation.*

    **UserDomain=** *Domain name for the user account in* **UserName** .

    **DatabasePath=** *Location of the Ntds.dit file. (The default is %systemroot%*\ntds*.) If you omit this entry, Dcpromo uses the default location.*

    **LogPath=** *Location of the database log files. (The default is %systemroot%*\ntds*.) If you omit this entry, Dcpromo uses the default location.*

    **SYSVOLPath=** *Location of the shared SYSVOL tree. (The default is %systemroot%*\SYSVOL*.)* If you omit this entry, Dcpromo uses the default location.

    **SafeModeAdminPassword=** *Password for the administrator account that must be used to start the domain controller in Directory Services Restore Mode.* If you leave this blank, Dcpromo prompts the user for the password during installation. Dcpromo deletes this value following installation. Passwords are removed from the answer file when Dcpromo is executed.

    **CriticalReplicationOnly=Yes** *or* **no** *, to specify whether to skip noncritical portions of replication and allow Dcpromo to complete before replication is complete.*

    **SiteName=** *The name of the Active Directory site in which this domain controller will be placed.* This site must be created in advance in the Active Directory Sites and Services snap-in.

    **ReplicaOrNewDomain=** *Specify either* **Replica** *for an additional domain controller in an existing domain or*

**NewDomain** *for the first domain controller in a new domain.*

**ReplicaDomainDNSName=** *The fully qualified domain name of the domain of the new domain controller.*

**ReplicationSourceDC=** *The name of an existing domain controller in the domain to use as the source replication partner during installation.* When you install Active Directory from restored backup media, you can use this entry in conjunction with **ReplicateFromMedia** if you want to specify the domain controller from which Active Directory changes and SYSVOL changes are replicated.

**ReplicateFromMedia=Yes** *for installation from media,* **no** *for installation by replication.*

| **Note** |
| --- |
| For more information on this option, see article 311078 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=128859). |

**ReplicationSourcePath=** *When the installation is by replication, the path to the installation CD or network share. When the installation is from restored backup media, the local drive and path to the restored backup files.*

**RebootOnSuccess=Yes** *if you want the domain controller to restart automatically following a successful installation,* **no** *if you want to restart the domain controller manually. If you do not want the domain controller to restart automatically and you do not want to be prompted, use the value* **NoAndNoPromptEither**.

**ApplicationPartitionsToReplicate=** *Comma-separated distinguished names, with the entire string enclosed in quotation marks, of application directory partitions that you want to include when you use restored backup media to install Active Directory (or * to include all application directory partitions).* Using this entry requires Windows Server 2003 with Service Pack 1 (SP1) and Windows Server 2003 forest functional level. For more information about using this entry, see Include application directory partitions in an Active Directory installation from backup media.

| **Note** |
| --- |
| For additional examples and parameters, see Ref.chm and article 223757 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=155037). |

11. Save the answer file to the location on the installation server from which it is to be called by Dcpromo, or save the file to a network shared folder or removable media for distribution.

**See Also**
**Concepts**
Include application directory partitions in an Active Directory installation from backup media

# Install Active Directory using an answer file

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to perform an unattended Active Directory installation by using an answer file.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group in the domain of the domain controller that you are installing.

**To install Active Directory using an answer file**

1. Click **Start**, click **Run**, and then type **dcpromo /answer:***answerfile*

   ***Answerfile***
   The path to the answer file, including the filename.

2. Click **OK**.

**See Also**
**Concepts**
Create an answer file for domain controller installation
Include application directory partitions in an Active Directory installation from backup media
Verifying Active Directory Installation

# Verifying Active Directory Installation

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are several verification tasks that can be performed on a computer on which Active Directory has been newly installed. Successfully completing the requirements of each verification task will provide a strong indication of a healthy, operational domain controller.

**Task Requirements**

The following tools are recommended to perform the procedures for this task:

- Active Directory Sites and Services
- DNS Manager
- Event Viewer
- Netdiag.exe
- Dcdiag.exe
- Ntdsutil.exe

To complete this task, perform the following procedures:

1. Determine whether a Server object has child objects
2. Verify that an IP address maps to a subnet and determine the site association

   You must ensure that the new domain controller is located in the proper site so that after the installation is complete, the new domain controller can locate replication partners and become part of the replication topology. If the site is not correct, you can use the Active Directory Sites and Services snap-in to move the Server object for the domain controller to the proper site after Active Directory installation is complete.

3. Move the Server object to the new site
4. Configure DNS server forwarders
5. Complete all procedures for the Verifying DNS configuration task.
6. Check the status of the shared SYSVOL
7. Verify DNS registration and functionality
8. Verify domain membership for a new domain controller
9. Verify communication with other domain controllers
10. Verify replication with other domain controllers
11. Verify the availability of the operations masters

# Determine whether a Server object has child objects

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After Active Directory is properly installed on a domain controller, the Server object for the domain controller will have a Child NTDS-Settings object. Other applications that are running on domain controllers can also publish Child objects.

Prior to deleting a Server object from the Servers container for a site, verify that the Server object has no Child objects. If a Child object appears, do not delete the Server object.

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Users group.

**To determine whether a server object has child objects**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and expand the site of the Server object.
3. Expand the **Servers** container, and then expand the Server object to view any Child objects.

# Verify that an IP address maps to a subnet and determine the site association

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to determine the site to which you want to add a Server object prior to installing Active Directory, or to verify the appropriate site prior to moving a Server object to it.

To be associated with a site, the IP address of a domain controller must map to a Subnet object that is defined in Active Directory. The site to which the subnet is associated is the site of the domain controller.

The subnet address, which is computed from the IP network address and the subnet mask, is the name of a Subnet object in Active Directory. When you know the subnet address, you can locate the Subnet object and determine the site to which the subnet is associated.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group.

**To verify that an IP address maps to a subnet and determine the site association**

1. Log on locally or open a Terminal Services connection to the server for which you want to check the IP address.
2. On the desktop, right-click **My Network Places**, and then click **Properties**.
3. In the **Network Connections** dialog box, right-click **Local Area Connection**, and then click **Properties**.
4. Double-click **Internet Protocol (TCP/IP)**.
5. Use the values in **IP address** and **Subnet mask** to calculate the subnet address and then click **OK**.
6. Click **OK** again and close the **Network Connections** dialog box.
7. Open Active Directory Sites and Services.
8. Expand the **Sites** container, and then click the **Subnets** container.
9. In the **Name** column in the details pane, find the Subnet object that matches the subnet address.
10. In the **Site** column, note the site to which the IP subnet address is associated.

    If the site that appears in the **Site** box is not the appropriate site, contact a supervisor and find out whether the IP address is incorrect or whether to move the Server object to the site indicated by the subnet.

# Move the Server object to the new site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Moving a Server object requires that the IP address of the domain controller maps to the site to which you are moving the Server object. Before performing this procedure, verify that the IP address maps to the target site.

**Administrative Credentials**

To perform this procedure, you must be a member of the Enterprise Admins group.

**To move the server object to the new site**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and the site in which the server object resides.
3. Expand the **Servers** container to display the domain controllers that are currently configured for that site.
4. Right-click the Server object you want to move, and then click **Move**.
5. In the **Site Name** box, click the destination site, and then click **OK**.
6. Expand the Site object to which you moved the server, and then expand the **Servers** container.
7. Verify that an object for the server you moved exists.
8. Expand the Server object and verify that an NTDS Settings object exists.

Within an hour, the Net Logon service on the domain controller registers the new site information in DNS. Wait an hour and then open Event Viewer and connect to the domain controller whose Server object you moved. Review the System log for Net Logon errors regarding registration of SRV resource records in DNS that have occurred within the last hour. The absence of errors indicates that Net Logon has updated DNS with site-specific SRV resource records. Net Logon event ID 5774 indicates that the registration of DNS resource records has failed. If this error occurs, contact a supervisor and pursue DNS troubleshooting.

# Configure DNS server forwarders

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure DNS server forwarders based on the forwarders method established on your network.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To configure DNS server forwarders**

1. If your network uses root hints as the forwarders method, you do not need to perform any additional options. Root hints are automatically configured during installation. Do not continue to step 2.
2. If you need to configure forwarders, open the DNS snap-in and continue to step 3.
3. In the console tree, right-click *ComputerName* (where *ComputerName* is the computer name of the domain controller), and then click **Properties**.
4. In the *ComputerName* Properties sheet (where *ComputerName* is the name of the domain controller), on the **Forwarders** tab, select the **Enable forwarders** check box.
5. In the **IP address** box, type *IpAddress* (where *IpAddress* is the IP address of the DNS server or nearest replication partner from which the domain is delegated), click **Add**, and then click **OK**.

# Verifying DNS configuration

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Part of verifying Active Directory installation is verifying that DNS was installed and configured appropriately.

**Task Requirements**

The following tools are required to perform the procedures for this task:

- DNS snap-in
- My Network Places

To complete this task, perform the following procedures:

1. Create a delegation for a domain controller

   If the parent DNS zone of any zone that is hosted by this DNS server contains a delegation to this DNS server, use this procedure to update the IP address in all such delegations.

   If your forest root domain has a parent DNS domain, perform this procedure on a DNS server in the parent domain. If you just added a new domain controller to a child domain, perform this procedure on a DNS server in the DNS parent domain. By following recommended practices, the parent domain is the forest root domain.

2. Configure the DNS client settings

3. Create a secondary zone

# Create a delegation for a domain controller

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to create a delegation for a new domain controller that is also a DNS server in the parent DNS domain.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To create a delegation for a domain controller**

1. Open the DNS snap-in.
2. Navigate to C*hildDomain* (where C*hildDomain* is the name of the child domain) in the console tree.
3. In the console tree, right-click C*hildDomain*, and then click **Properties**.
4. In the C*hildDomain* **Properties** sheet, on the **Name Servers** tab, click **Add**.
5. In the **New Resource Record** dialog box, in the **Server fully qualified domain name (FQDN)** box, type *ChildDC.ChildDomain.ParentDomain* (where *ChildDC* is the name of the new domain controller, *ChildDomain* is the name of the child domain, and *ParentDomain* is the name of the parent domain).
6. In the **New Resource Record** dialog box, in the **IP address** box, type *IPAddress* (where *IPAddress* is the IP address of the child domain controller), click **Add**, and then click **OK**.

# Create a secondary zone

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Perform this procedure only on new domain controllers that are also DNS servers that are located in the child domain, not the forest root domain.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To create a secondary zone**

1. Open the DNS snap-in.
2. In the console tree, right-click the new domain controller and click **New Zone**.
3. In the **New Zone Wizard**, click **Next** to continue.
4. On the **Zone Type** page, select **Secondary zone** and click **Next**.
5. Ensure that **Forward lookup zone** is selected. Click **Next**.
6. For Zone name, type **_msdcs.***forestrootdomain* (where *forestrootdomain* is the fully qualified domain name of the forest root domain), and click **Next**.
7. In the **Master DNS Servers** dialog box, enter the IP addresses of at least two DNS servers in the forest root domain. Click **Next**.
8. Review the settings you defined, and click **Finish** to close the wizard.

# Configure the DNS client settings

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configure the DNS client settings on the new domain controller.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To configure the DNS client settings**

1. On the desktop, right-click **My Network Places** and click **Properties**.
2. In the **Network Connections** dialog box, right-click the connection that represents the connection this computer uses to attach to your network. The default label is **Local Area Connection**, but this can be changed so it might not be labeled the same on your computer. Click **Properties**.
3. In the **Local Area Connection Properties** dialog box, click once on **Internet Protocol (TCP/IP)** to highlight it (be sure you do not clear the check box in front of it), then click **Properties**.
4. In the **Internet Protocol (TCP/IP) Properties** dialog box, verify that **Use the following DNS server addresses:** is selected.
5. If the new domain controller is located in the forest root domain, set the **Preferred DNS server** IP address to that of another DNS server in the forest root domain. Try to choose a server that is located near the new domain controller. Set the **Alternate DNS server** address to the IP address of the new domain controller (so that it is referencing itself).

   If the new domain controller is located in a child domain, set the **Preferred DNS server** IP address to the IP address of the new domain controller (so that it is referencing itself). Set the **Alternate DNS server** address to that of another DNS server in the same domain. Try to choose a server that is located near the new domain controller.
6. Click **OK** to close the dialog box.

# Check the status of the shared SYSVOL

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure involves checking Event Viewer to make sure that the File Replication service is started properly and then ensuring that the SYSVOL and Net Logon shared folders are created.

| **Note** |
| --- |
| You do not need to perform this procedure on every replication partner, but you need to perform it enough times to be confident that the shared system volumes on the replication partners are healthy. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To check the status of the shared SYSVOL**

1. Open Event Viewer.
2. In the Event Viewer tree, click **File Replication Service** to display the FRS events.
3. Look for an event 13516 with a date and time stamp that corresponds with the recent restart. It can take 15 minutes or more to appear. An event 13508 indicates that FRS is in the process of starting the service. An event 13509 indicates that the service has started successfully. Event 13516 indicates that the service is started, the folders are shared, and the domain controller is functional.
4. To verify the shared folder is created, open a command prompt and type **net share** to display a list of the shared folders on this domain controller, including Net Logon and SYSVOL.
5. At a command prompt, type **dcdiag /test:netlogons** and press ENTER.
6. Look for a message that states *computername* **passed test NetLogons** where *computername* is the name of the domain controller. If you do not see the test passed message, some problem will prevent replication from functioning. This test verifies that the proper logon privileges are set to allow replication to occur. If this test fails, verify the permissions set on the Net Logon and SYSVOL shared folders.

# Verify DNS registration and functionality

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that Domain Name System (DNS) is functioning so that other domain controllers can be located.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To verify DNS registration and functionality**

1. Open a Command Prompt.
2. On a **Windows Server 2003** computer, type the following command, and then press ENTER:

   **netdiag /test:dns**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   On a **Windows Server 2008** or **Windows Server 2008 R2** computer, type the following command, and then press ENTER:

   **dcdiag /test:dns**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   If DNS is functioning, the last line of the response for all operating system versions is **DNS Test.....: Passed**. The verbose option lists specific information about what was tested. This information can help with troubleshooting if the test fails.

   If the test fails, do not attempt any additional steps until you determine and fix the problem that prevents proper DNS functionality.

# Verify communication with other domain controllers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that domain controllers can be located.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain users group in Active Directory.

**To verify communication with other domain controllers**

1. Open a Command Prompt.
2. On a **Windows Server 2003** computer, type the following command, and then press ENTER:

   **netdiag /test:dsgetdc**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   If domain controllers are successfully located, the last line of the response is **DC discovery test.....: Passed.** The verbose option lists the specific domain controllers that are located.

   On a **Windows Server 2008** or **Windows Server 2008 R2** computer, type the following command, and then press ENTER:

   **nltest /dclist:** <*domain name*>

   If domain controllers are successfully located, the last line of the response is **The command completed successfully**.

   If the tests fail on any of the operating system versions, do not attempt any additional steps until you determine and fix the problem that prevents communication with other domain controllers.

# Verify replication with other domain controllers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The tests performed in this procedure verify that different aspects of the replication topology are working properly. They check to see that objects are replicating and they verify that the proper logon permissions are set to allow replication to occur.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To verify replication is functioning**

1. Open a Command Prompt.
2. Type the following command, and then press Enter:

   **dcdiag /test:replications**

   | **Note** |
   | --- |
   | For this set of tests, the **/v** option is available. However, it does not display any significant additional information. Messages indicate that the connectivity and replications tests passed. |

3. To verify that the proper permissions are set for replication, type the following command and then press Enter:

   **dcdiag /test:netlogons**

   Messages indicate that the connectivity and netlogons tests passed.

# Verify the availability of the operations masters

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that the operations masters can be located and that they are online and responding.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain users group in Active Directory.

| **Note** |
| --- |
| You can use these tests prior to installing Active Directory as well as afterward. To perform the test prior to installing Active Directory, you must use the **/s** option to indicate the name of a domain controller to use. You do not need the **/s** option to perform the test after installing Active Directory. The test automatically runs on the local domain controller where you are performing the test. The commands listed in this procedure show the **/s** option. If you are performing this test after installing Active Directory, omit the **/s** option. For a more detailed response from this command, you can use the verbose option by adding **/v** to the end of the command to see the detailed response. |

**To verify the availability of the operations masters**

1. Open a Command Prompt.
2. Type the following command to ensure that the operations masters can be located and then press ENTER:

   **dcdiag /s:** *domaincontroller* **/test:knowsofroleholders /verbose**

   where *domaincontroller* is the name of a domain controller in the domain in which you want to add the new domain controller. The verbose option provides a detailed list of the operations masters that were tested. Near the bottom of the screen, a message confirms that the test succeeded. If you use the verbose option, look carefully at the bottom part of the displayed output. The test confirmation message appears immediately after the list of operations masters. Press ENTER.
3. Type the following command to ensure that the operations masters are functioning properly and are available on the network:

   **dcdiag /s:** *domaincontroller* **/test:fsmocheck**

   where *domaincontroller* is the name of a domain controller in the domain in which you want to add the new domain controller. The verbose option provides a detailed list of the operations masters that were tested. Near the bottom of your screen, a message confirms that the test succeeded. Press ENTER.

   If these tests fail, do not attempt any additional steps until you determine and fix the problem that prevents locating operations masters and verifying that they are functioning properly.

# Verify domain membership for a new domain controller

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This test verifies that a new domain controller has successfully become a member of the domain.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group.

**To verify domain membership for a new domain controller**

1. Open a Command Prompt.
2. On a **Windows Server 2003** computer, type the following command, and then press ENTER:

   **netdiag /test:member**

   If the test was successful, you should see the following message: "Domain membership test passed." If you use the **/v** option, it lists the name of the domain controller, its role, the name of the domain, and a number of other statistics about the new domain controller.

   On a **Windows Server 2008** or **Windows Server® 2008 R2** computer, type the following command, and then press ENTER:

   **nltest /dsgetdc:** <*domain name*>

   If the test was successful, you should see the domain name in the list of returned results and the message: "The command completed successfully."

# Renaming a Domain Controller

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The ability to rename domain controllers running Windows Server 2003 (contrary to Windows 2000 Server) provides you with the flexibility to:

- Restructure your network for organizational and business needs.
- Make management and administrative control easier.

Renaming a domain controller is a common operation in many organizations and usually occurs when:

- New hardware is purchased to replace an existing domain controller.
- Domain controllers are decommissioned, or promoted, and renamed to maintain a naming convention.
- Domain controllers are moved or placed in sites.

| **Note** |
| --- |
| It is important to note that domain controller names have a primary impact on administration, rather than client access. Renaming a domain controller is an optional exercise, and the impacts should be well understood prior to renaming. |

Although you can use the System Properties user interface (UI) to rename a domain controller (as you can for any computer), Active Directory and DNS replication latency might temporarily prevent clients from locating or authenticating to the renamed domain controller, or both. To avoid this delay, use the Netdom command-line tool to rename a domain controller.

**Task requirements**

The following tools are required to perform the procedures for this task:

- System Properties or Netdom.exe
- Ldp.exe or Adsiedit.msc

If you want to use Netdom, the domain functional level must be set to Windows Server 2003.

To complete this task, use one of the following two sets of procedures:

1. Rename a domain controller using System Properties
2. Update the FRS member object

Or

1. Rename a domain controller using Netdom
2. Update the FRS member object

# Rename a domain controller using System Properties

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to rename a domain controller by using the System Properties user interface (UI).

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group.

**To rename a domain controller using System Properties**

1. Click **Start**, and then click **Control Panel**.
2. In **Control Panel**, double-click **System Properties**.
3. On the **Computer Name** tab, click **Change**.
4. Click **OK** to acknowledge that renaming the domain controller may cause it to become temporarily unavailable to users and computers. (See note below.)
5. Under **Computer Name**, type the new name.
6. Click **OK** to close the **System Properties** dialog box.
7. If you are prompted, provide the user name and password for an account with Domain Admin or Enterprise Admin credentials.

| **Note** |
| --- |
| Renaming a domain controller in this way may result in Active Directory replication latency, making it more difficult for clients to locate or authenticate the domain controller under its new name. |

**See Also**
**Concepts**
Rename a domain controller using Netdom

# Rename a domain controller using Netdom

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to rename a domain controller by using the Netdom command-line tool.

The **netdom** command updates the service principal name (SPN) attributes in Active Directory for the computer account and registers Domain Name System (DNS) resource records for the new computer name. The SPN value of the computer account must be replicated to all domain controllers in the domain, and the DNS resource records for the new computer name must be distributed to all the authoritative DNS servers for the domain name. If the updates and registrations have not occurred prior to removal of the old computer name, some clients might be unable to locate this computer using the new name or the old name.

**Administrative credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group.

**To rename a domain controller using Netdom**

1. Open a Command Prompt.
2. Type the following command to add the new domain controller name, and then press ENTER:

   **netdom computername** *CurrentComputerName* **/add:** *NewComputerName*

3. Type the following command to designate the new name as the primary computer name, and then press ENTER:

   **netdom computername** *CurrentComputerName* **/makeprimary:** *NewComputerName*

| Term | Definition |
|---|---|
| *CurrentComputerName* | The current, or primary, computer name or Internet Protocol (IP) address of the computer that you are renaming. |
| *NewComputerName* | The new name for the computer. The *NewComputerName* must be a fully qualified domain name (FQDN). The primary DNS suffix that is specified in the FQDN for *NewComputerName* must be the same as the primary DNS suffix of *CurrentComputerName*, or it must match the DNS name of the Active Directory domain that is hosted by this domain controller, or it must be contained in the list of allowed DNS suffixes that is specified in the **msDS-AllowedDNSSuffixes** attribute of the domainDns object. |

4. Restart the computer.
5. After the computer restarts, open a Command Prompt.
6. Type the following command to remove the old domain controller name, and then press ENTER:

   **netdom computername** *NewComputerName* **/remove:** *OldComputerName*

| Term | Definition |
|---|---|
| *NewComputerName* | The new FQDN that you added for the computer in step 2. |
| *OldComputerName* | The old FQDN of the renamed computer. |

**See Also**
**Concepts**
Rename a domain controller using System Properties

# Update the FRS member object

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to update the File Replication service (FRS) member object after renaming a domain controller.

For more information about this procedure, see article 316826 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=82821).

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group.

**To update the FRS member object**

1. Using Ldp.exe (or ADSI edit), find the computer object of the renamed domain controller.
2. Do a recursive search for an object of type **nTFRSSubscriber** with the computer name of "**Domain System Volume (SYSVOL share)**" under the Computer object.
3. The search filter is "(&((cn=Domain System Volume (SYSVOL share))(objectclass=ntfrssubscriber)))".
4. Find the **fRSMemberReference** attribute of the object returned by the search.
5. Find the object whose domain name is in the **fRSMemberReference** attribute. This is the **Ntfrsmember** object corresponding to this domain controller.
6. Change the computer name of this **Ntfrsmember** object from the old name of the domain controller to the new name of the domain controller.

# Decommissioning a Domain Controller

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Decommissioning a domain controller effectively removes all Active Directory and related components and returns the domain controller to a member server role.

**Protecting EFS-encrypted files**

If the domain controller to be decommissioned hosts any Encrypting File System (EFS) encrypted files, you must take precautions to protect the private key for the recovery agent for the local EFS-encrypted documents. It is possible for this key to be lost during the demotion when the Security Accounts Manager (SAM) is recreated on the computer. In this case, you are cannot recover encrypted documents on this computer unless the recovery agent is changed to an existing domain account before encryption. To prevent loss of the private key, you must back up (export) the recovery agent private key before you decommission the domain controller. After you remove Active Directory, re-import the private key.

You must be able to ensure that the domain account that serves as the recovery agent for the certificate remains the same after removing Active Directory. If you cannot guarantee that the account will remain the same after the domain controller is decommissioned, or if you removed Active Directory without backing up the certificate and you cannot recover EFS-encrypted files, see article 276239 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=117370).

**Task Requirements**

The following tools are required to perform the procedures for this task:

- Ntdsutil.exe
- Active Directory Domains and Trusts
- Active Directory Users and Computers
- Active Directory Sites and Services
- Netdiag.exe
- Dcdiag.exe

To complete this task, perform the following procedures:

1. View the current operations master role holders

   To avoid problems, transfer any operations master roles prior to running the Active Directory Installation Wizard to decommission a domain controller so that you can control the operations master role placement. If you need to transfer any roles from a domain controller, understand all the recommendations for role placement before performing the transfer.

   | ⚠Caution |
   | --- |
   | During the decommissioning process, the Active Directory Installation Wizard will attempt to transfer any remaining operations master roles to other domain controllers without any user interaction. However, if a failure occurs, the wizard will continue to uninstall Active Directory and leave your domain without roles. Also, you do not have control over which domain controller receives the roles. The wizard transfers the roles to any available domain controller and does not indicate which domain controller hosts them. |

2. Transfer the schema master
3. Transfer the domain naming master
4. Transfer the domain-level operations master roles

5. Determine whether a domain controller is a global catalog server

   If you remove Active Directory from a domain controller that hosts a global catalog, the Active Directory Installation Wizard confirms that you want to continue with removing Active Directory. This confirmation ensures that you are aware that you are removing a global catalog from your environment. Do not remove the last global catalog server from your environment because users cannot log on without an available global catalog server. If you are not sure, do not proceed with removing Active Directory until you know that at least one other global catalog server is available.

6. Verify DNS registration and functionality

7. Verify communication with other domain controllers

   During the removal of Active Directory, contact with other domain controllers is required to ensure:

   - Any unreplicated changes are replicated to another domain controller.
   - Removal of the domain controller from the directory.
   - Transfer of any remaining operations master roles.

   If the domain controller cannot contact the other domain controllers during Active Directory removal, the decommissioning operation fails. As with the installation process, test the communication infrastructure prior to running the installation wizard. When you remove Active Directory, use the same connectivity tests that you used during the installation of Active Directory.

8. Verify the availability of the operations masters

   | **Important** |
   | --- |
   | If any of the verification tests fail, do not continue until you determine and fix the problems. If these tests fail, the uninstallation is also likely to fail. |

9. If the domain controller hosts encrypted documents, perform the following procedure before you remove Active Directory to ensure that the encrypted files can be recovered after Active Directory is removed:

   Export a certificate with the private key (http://go.microsoft.com/fwlink/?LinkId=20039)

10. Uninstall Active Directory

11. If the domain controller hosts encrypted documents and you backed up the certificate and private key before you remove Active Directory, perform the following procedure to re-import the certificate to the server:

    Import a certificate (http://go.microsoft.com/fwlink/?LinkId=20040)

12. Determine whether a Server object has child objects

13. Delete a Server object from a site

    | **Note** |
    | --- |
    | The administrator may not want to remove the Server object if it hosts something in addition to Active Directory—Microsoft Exchange, for example. |

# View the current operations master role holders

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Once an operations master role has been transferred, it should be verified that the transfer has occurred successfully throughout the domain. The change must be replicated to all relevant domain members in order to truly take effect.

To view the current operations master role holders, use Ntdsutil.exe with the roles option. This option displays a list of all current role holders.

**Administrative Credentials**

To perform this procedure, you must be logged on as a User or an Administrator.

**To view the current operations master role holder**

1. Click **Start**, click **Run**, type **ntdsutil**, and then press ENTER.
2. At the **ntdsutil:** prompt, type **roles** and press ENTER.
3. At the **fsmo maintenance:** prompt, type **connections** and press ENTER.
4. At the **server connections:** prompt, type **connect to server** *servername* (where *servername* is the name of the domain controller that belongs to the domain containing the operations masters).
5. After receiving confirmation of the connection, type **quit** and press ENTER to exit this menu.
6. At the **fsmo maintenance:** prompt, type **select operation target** and press ENTER.
7. At the **select operations target:** prompt, type **list roles for connected server** and press ENTER.

   The system responds with a list of the current roles and the Lightweight Directory Access Protocol (LDAP) name of the domain controllers currently assigned to host each role.
8. Type **quit** and press ENTER to exit each prompt in Ntdsutil.exe. Type **quit** and press ENTER at the **ntdsutil:** prompt to close the window.

# Transfer the schema master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to transfer the schema operations master role. The schema master is a forest-wide operations master role. Before you can use the Active Directory Schema snap-in for the first time, you must register it with the system. If you have not yet prepared the Active Directory Schema snap-in, see Install the Schema snap-in before you begin this procedure.

| **Note** |
| --- |
| This procedure is performed by using the Microsoft Management Console (MMC), although you can also transfer this role by using Ntdsutil.exe. For information about using Ntdsutil.exe to transfer operations master roles, type **?** at the Ntdsutil.exe command prompt. |

### Administrative Credentials

To perform this procedure, you must be a Schema Administrator in Active Directory.

#### Transfer the schema master

1. Open the Active Directory Schema snap-in.
2. In the console tree, right-click **Active Directory Schema**, and click **Change Domain Controller**.
3. In the **Change Domain Controller** dialog box, click **Specify Name**. Then, in the text box, type the name of the server to which you want to transfer the schema master role. Click **OK**.
4. In the console tree, right-click **Active Directory Schema**. Click **Operations Master**. The **Change Schema Master** box displays the name of the server that is currently holding the role. The targeted domain controller is listed in the second box.
5. Click **Change**. Click **Yes** to confirm your choice. The system confirms the operation. Click **OK** again to confirm that the operation succeeded.
6. Click **Close** to close the **Change Schema Master** dialog box.

| **Note** |
| --- |
| Hosting the infrastructure master on a global catalog server is not recommended. If you attempt to transfer the infrastructure master role to a domain controller that is a global catalog, the system displays a warning stating that this is not recommended. |

# Transfer the domain naming master

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to transfer the domain naming operations master role. The domain naming master is a forest-wide operations master role.

| **Note** |
| --- |
| This procedure is performed by using the Microsoft Management Console (MMC), although you can also transfer this role by using Ntdsutil.exe. For information about using Ntdsutil.exe to transfer operations master roles, type **?** at the Ntdsutil.exe command prompt. |

### Administrative Credentials

To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory.

#### To transfer the domain naming master

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click **Active Directory Domains and Trusts**, and then click **Connect to Domain Controller**.
3. Ensure that the proper domain name is entered in the **Domain** box.

   The available domain controllers from this domain are listed.
4. In the **Name** column, click the domain controller (to select it) to which you want to transfer the role. Click **OK**.
5. Right-click **Active Directory Domains and Trusts**, and then click **Operations Master**.
6. The name of the current domain naming master appears in the first text box. The server to which you want to transfer the role should appear in the second text box. If this is not the case, repeat steps 1 through 4.
7. Click **Change**. To confirm the role transfer, click **Yes**. Click **OK** again to close the message box indicating the transfer took place. Click **Close** to close the **Change Operations Master** dialog box.

# Transfer the domain-level operations master roles

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to transfer the three domain-level operations master roles: the PDC emulator, the RID master, and the infrastructure master. You can transfer all of these roles by using the Active Directory Users and Computers console.

| **Note** |
| --- |
| These procedures are performed by using MMC, although you can also transfer these roles by using Ntdsutil.exe. For information about using Ntdsutil.exe to transfer the operations master roles, type **?** at the Ntdsutil.exe command prompt. |

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To transfer a domain-level operations master role**

1. Open Active Directory Users and Computers.
2. At the top of the console tree, right-click **Active Directory Users and Computers**. Click **Connect to Domain Controller**.
3. In the list of **available domain controllers**, click the name of the server to which you want to transfer the role, and then click **OK**.
4. At the top of the console tree, right-click **Active Directory Users and Computers**, point to **All Tasks**, and then click **Operations Masters**.

   The name of the current operations master role holder appears in the **Operations master** box. The name of the server to which you want to transfer the role appears in the lower box.
5. Click the tab for the role you want to transfer: **RID**, **PDC**, or **Infrastructure**. Verify the computer names that appear and then click **Change**. Click **Yes** to transfer the role, and then click **OK**.
6. Repeat steps 4 and 5 for each role that you want to transfer.

# Determine whether a domain controller is a global catalog server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the setting on the NTDS Settings object to indicate whether a domain controller is designated as a global catalog server.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Users group in Active Directory.

**To determine whether a domain controller is a global catalog server**

1. Open Active Directory Sites and Services.
2. In the console tree, expand the **Sites** container, expand the site of the domain controller you want to check, expand the **Servers** container, and then expand the **Server** object.
3. Right-click the **NTDS Settings** object, and then click **Properties**.
4. On the **General** tab, if the **Global Catalog** box is selected, the domain controller is designated as a global catalog server.

# Verify DNS registration and functionality

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that Domain Name System (DNS) is functioning so that other domain controllers can be located.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To verify DNS registration and functionality**

1. Open a Command Prompt.
2. On a **Windows Server 2003** computer, type the following command, and then press ENTER:

   **netdiag /test:dns**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   On a **Windows Server 2008** or **Windows Server 2008 R2** computer, type the following command, and then press ENTER:

   **dcdiag /test:dns**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   If DNS is functioning, the last line of the response for all operating system versions is **DNS Test.....: Passed**. The verbose option lists specific information about what was tested. This information can help with troubleshooting if the test fails.

   If the test fails, do not attempt any additional steps until you determine and fix the problem that prevents proper DNS functionality.

# Verify communication with other domain controllers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that domain controllers can be located.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain users group in Active Directory.

**To verify communication with other domain controllers**

1. Open a Command Prompt.
2. On a **Windows Server 2003** computer, type the following command, and then press ENTER:

   **netdiag /test:dsgetdc**

   | **Note** |
   | --- |
   | For a more detailed response from this command, add **/v** to the end of the command. |

   If domain controllers are successfully located, the last line of the response is **DC discovery test.....: Passed.** The verbose option lists the specific domain controllers that are located.

   On a **Windows Server 2008** or **Windows Server 2008 R2** computer, type the following command, and then press ENTER:

   **nltest /dclist:** <*domain name*>

   If domain controllers are successfully located, the last line of the response is **The command completed successfully**.

   If the tests fail on any of the operating system versions, do not attempt any additional steps until you determine and fix the problem that prevents communication with other domain controllers.

# Verify the availability of the operations masters

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure verifies that the operations masters can be located and that they are online and responding.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain users group in Active Directory.

| **Note** |
| --- |
| You can use these tests prior to installing Active Directory as well as afterward. To perform the test prior to installing Active Directory, you must use the **/s** option to indicate the name of a domain controller to use. You do not need the **/s** option to perform the test after installing Active Directory. The test automatically runs on the local domain controller where you are performing the test. The commands listed in this procedure show the **/s** option. If you are performing this test after installing Active Directory, omit the **/s** option. For a more detailed response from this command, you can use the verbose option by adding **/v** to the end of the command to see the detailed response. |

**To verify the availability of the operations masters**

1. Open a Command Prompt.
2. Type the following command to ensure that the operations masters can be located and then press ENTER:

   **dcdiag /s:** *domaincontroller* **/test:knowsofroleholders /verbose**

   where *domaincontroller* is the name of a domain controller in the domain in which you want to add the new domain controller. The verbose option provides a detailed list of the operations masters that were tested. Near the bottom of the screen, a message confirms that the test succeeded. If you use the verbose option, look carefully at the bottom part of the displayed output. The test confirmation message appears immediately after the list of operations masters. Press ENTER.
3. Type the following command to ensure that the operations masters are functioning properly and are available on the network:

   **dcdiag /s:** *domaincontroller* **/test:fsmocheck**

   where *domaincontroller* is the name of a domain controller in the domain in which you want to add the new domain controller. The verbose option provides a detailed list of the operations masters that were tested. Near the bottom of your screen, a message confirms that the test succeeded. Press ENTER.

   If these tests fail, do not attempt any additional steps until you determine and fix the problem that prevents locating operations masters and verifying that they are functioning properly.

# Uninstall Active Directory

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To use the Active Directory Installation Wizard to remove Active Directory, you must know the password to assign to the local Administrator account of the server after Active Directory is removed.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To uninstall Active Directory**

1. Click **Start**, click **Run**, type **dcpromo** and then click **OK**.
2. The Active Directory Installation Wizard appears. Click **Next** at the Welcome screen.
3. You have an option to select **This server is the last domain controller in the domain**. If you select this option, the wizard attempts to remove the domain from the forest. Do not select this option. Click **Next**.
4. At the **Administrative Password** screen, enter and confirm the password that you want to assign to the local Administrator account after Active Directory is removed. Click **Next**.
5. At the **Summary** screen, verify that the information is correct and then click **Next** to proceed with the removal.
6. The wizard proceeds to remove Active Directory. After it finishes, the wizard displays a completion screen. Click **Finish** to close the wizard.
7. Click **Restart** to restart the domain controller.

# Determine whether a Server object has child objects

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After Active Directory is properly installed on a domain controller, the Server object for the domain controller will have a Child NTDS-Settings object. Other applications that are running on domain controllers can also publish Child objects.

Prior to deleting a Server object from the Servers container for a site, verify that the Server object has no Child objects. If a Child object appears, do not delete the Server object.

### Administrative Credentials

To perform this procedure, you must be a member of the Domain Users group.

**To determine whether a server object has child objects**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container and expand the site of the Server object.
3. Expand the **Servers** container, and then expand the Server object to view any Child objects.

# Delete a Server object from a site

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When no Child objects are visible below the Server object in Active Directory Sites and Services, you can remove the Server object.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To delete a server object from a site**

1. Open Active Directory Sites and Services.
2. Expand the **Sites** container, and then expand the site from which you want to delete a Server object.
3. If no Child objects appear below the Server object, right-click the Server object, and then click **Delete**.

| **Important** |
| --- |
| Do not delete a Server object that has a Child object. If an NTDS Settings or other Child object appears below the Server object you want to delete, either replication on the domain controller on which you are viewing the Configuration container has not occurred, or the server whose Server object you are removing has not been properly decommissioned. |

4. Click **Yes** to confirm your choice.

# Forcing the Removal of a Domain Controller

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Forced removal of a domain controller from Active Directory is intended to be used as a last resort to avoid having to reinstall the operating system on a domain controller that has failed and cannot be recovered. When a domain controller can no longer function in a domain (that is, it is offline), you cannot remove Active Directory in the normal way, which requires connectivity to the domain. Forced removal is not intended to replace the normal Active Directory removal procedure in any way. It is virtually equivalent to permanently disconnecting the domain controller.

Active Directory stores a considerable amount of metadata about a domain controller. During the normal process of uninstalling Active Directory on a domain controller, this metadata is removed from Active Directory through a connection to another domain controller in the domain. A forced removal assumes that there is no connectivity to the domain; therefore, it does not attempt any metadata removal (cleanup).

Consequently, forced removal of Active Directory from a domain controller should always be followed by the metadata cleanup procedure, which removes all references to the domain controller from the domain and forest.

Forced demotion should not be performed on the last domain controller in a domain.

**Task Requirements**

The following tools are required to perform the procedures for this task:

- Active Directory Sites and Services
- Dcpromo.exe
- Ntdsutil.exe

To complete this task, perform the following procedures:

1. Identify replication partners. Connect to one of these domain controllers when you clean up server metadata in procedure 3.
2. Force domain controller removal
3. Clean up server metadata

# Identify replication partners

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to examine the Connection objects for a domain controller and determine its replication partners.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**To identify replication partners**

1. Open Active Directory Sites and Services.
2. In the console tree, expand the **Sites** container to display the list of sites.
3. Double-click the site that contains the domain controller for which you want to determine Connection objects.

| **Note** |
| --- |
| If you do not know the site in which the domain controller is located, open a command prompt and type **ipconfig** to get the IP address of the domain controller. Use the IP address to Verify that an IP address maps to a subnet and determine the site association. |

4. Expand the **Servers** folder to display the list of servers in that site.
5. Expand the name of your domain controller to display its NTDS settings.
6. Double-click **NTDSSettings** to display the list of Connection objects in the details pane (these represent inbound connections used for replication). The **From Server** column displays the names of the domain controllers that are the replication partners.

# Force domain controller removal

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to force the removal of Active Directory from a domain controller.

**Administrative Credentials**

To perform this procedure, you must be a member of the Domain Admins group.

**To force domain controller removal**

1. Click **Start**, click **Run**, type the following command and then press ENTER:

   **Dcpromo /forceremoval**

2. At the **Welcome to the Active Directory Installation Wizard** page, click **Next**.
3. At the **Force the Removal of Active Directory** page, click **Next**.
4. In **Administrator Password**, type the password and confirmed password that you want to assign to the Administrator account of the local SAM database, and then click **Next**.
5. In **Summary**, click **Next**.

# Clean up server metadata

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You perform the metadata cleanup process by using Ntdsutil.exe, a command-line tool that is automatically installed on all domain controllers. Metadata cleanup removes data from Active Directory that identifies a domain controller to the replication system. This procedure is required only for Active Directory domain controllers that were not successfully demoted using Dcpromo. This procedure does not have to be performed for domain member servers or client computers. On a domain controller that is running Windows Server 2003 with Service Pack 1 (SP1), metadata cleanup also removes File replication service (FRS) connections and attempts to transfer or seize any operations master roles (also known as flexible single master operations or FSMO roles) that the retired domain controller holds. These additional processes are performed automatically.

You can also use a script to clean up server metadata on most Windows operating systems. For information about using this script, see Remove Active Directory Domain Controller Metadata on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=123599).

**Administrative credentials**

- To complete this procedure, you must be a member of the Enterprise Admins group.

**To clean up server metadata**

1. Open a command prompt.
2. Type the following command, and then press ENTER:

   **ntdsutil**

3. At the **ntdsutil:** prompt, type:

   **metadata cleanup**

4. Perform metadata cleanup as follows:

   - If you are performing metadata cleanup by using the version of Ntdsutil.exe that is included with Windows Server 2003 SP1, at the **metadata cleanup:** prompt, type:

     **remove selected server** *ServerName*

     Or

     **remove selected server** *ServerName1* **on** *ServerName2*

     | Value | Definition |
     |---|---|
     | *ServerName, ServerName1* | The distinguished name of the domain controller whose metadata you want to remove, in the form cn=*ServerName*,cn=Servers,cn=*SiteName*, cn=Sites,cn=Configuration,dc=*ForestRootDomain* |
     | *ServerName2* | The DNS name of the domain controller to which you want to connect and from which you want to remove server metadata |

   - If you are performing metadata cleanup by using the version of Ntdsutil.exe that is included with Windows Server 2003 with no service pack, perform metadata cleanup as follows:

     a. At the **metadata cleanup:** prompt, type:

     **connection**

     b. At the **server connections:** prompt, type:

**connect to server** *Server*

c. At the **server connections:** prompt, type:

**quit**

d. At the **metadata cleanup:** prompt, type:

**select operation target**

e. At the **select operation target:** prompt, type:

**list sites**

A numbered list of sites appears.

f. At the **select operation target:** prompt, type:

**select site** *SiteNumber*

g. At the **select operation target:** prompt, type:

**list domains in site**

A numbered list of domains in the selected site appears.

h. At the **select operation target:** prompt, type:

**select domain** *DomainNumber*

i. At the **select operation target:** prompt, type:

**list servers in site**

A numbered list of servers in a domain and site appears.

j. At the **select operation target:** prompt, type:

**select server** *ServerNumber*

k. At the **select operation target:** prompt, type:

**quit**

l. At the **metadata cleanup:** prompt, type:

**remove selected server**

| Value | Description |
|---|---|
| *Server* | The DNS name of a domain controller that you want to connect to |
| *SiteNumber* | The number associated with the site of the server that you want to clean up that appears in the list |

| | |
|---|---|
| *DomainNumber* | The number associated with the domain of the server that you want to clean up that appears in the list |
| *ServerNumber* | The number associated with the server that you want to clean up that appears in the list |

At this point, Active Directory confirms that the domain controller was removed successfully. If you receive an error message that indicates that the object cannot be found, Active Directory might have already removed the domain controller.

5. To verify that the server was removed, type **list servers in site**, and then press ENTER. Ensure that the domain controller that you wanted to be removed is no longer displayed in the command output.
6. At the **metadata cleanup:** and **ntdsutil:** prompts, type **quit**.

| **Important** |
|---|
| If the domain controller that was removed was also a DNS server, ensure that references to it are removed from the **Name Servers** tab in the DNS console. To do this, open the DNS console using another DNS server in the domain (dnsmgmt.msc), and then click the domain name under **Forward Lookup Zones**. Remove any references to the domain controller that was removed from the domain. |

**See Also**
**Other Resources**
Article 216498 in the Microsoft Knowledge Base

# Additional Resources for Administering Active Directory

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For general information about how Active Directory works and how to deploy, monitor, and delegate Active Directory, see the following resources:

- Active Directory Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=26094)
- Designing and Deploying Directory and Security Services on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=27638)
- "Monitoring Active Directory Health" in the Active Directory Management Pack Technical Reference for MOM 2005 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43127)
- Best Practices for Delegating Active Directory Administration on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=46579)

For specific information about troubleshooting Active Directory problems, see the following resources:

- Troubleshooting Active Directory Operations
- Identity and Directory Services Community on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=20151)
- Windows Server Active Directory Newsgroup on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43065)

For development information about Active Directory, see the following resources:

- Active Directory Platform SDK on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=142)
- Lightweight Directory Access Protocol Platform SDK on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=2972)
- RFC Pages and Internet-Drafts on the Internet Engineering Task Force Web site (http://go.microsoft.com/fwlink/?LinkID=121)

| **Note** |
| --- |
| Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here. |

# DNS Operations Guide

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The DNS Operations Guide provides administering and troubleshooting information for Domain Name System (DNS) in the Microsoft® Windows Server™ 2003 with Service Pack 1 (SP1) operating system.

**In this guide**


- Administering DNS Operations
- Troubleshooting Domain Name System
- Additional Resources for Domain Name System

# Administering DNS Operations

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This Domain Name System (DNS) Administering guide provides administering information for DNS in the Microsoft Windows Server 2003 with Service Pack 1 (SP1) operating system.

**In this guide**


- Introduction to Administering DNS Operations
- Managing DNS
- Monitoring DNS
- Optimizing DNS
- Securing DNS


This DNS Administering guide provides detailed procedures for managing DNS servers, clients, and resource records. It also provides procedures for monitoring, optimizing, and securing your DNS infrastructure. For most procedures, this guide provides both a user interface (UI) and a command-line method of performing each procedure. In addition, this guide provides sample scripts for the most frequently used, repetitive tasks.

# Introduction to Administering DNS Operations

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide explains how to administer Microsoft Domain Name System (DNS). These activities are part of the operating phase of the information technology (IT) life cycle. If you are not familiar with this guide, review the following sections of this introduction.

**When to Use This Guide**
You should use this guide when:

- You want to manage DNS servers.
- You want to manage DNS clients.

This guide assumes a basic understanding of what DNS is, how it works, and why your organization uses it for name resolution. You should also have a thorough understanding of how DNS is deployed and managed in your organization. This includes an understanding of the mechanism that your organization uses to configure and manage DNS settings.

This guide can be used by organizations that have deployed Windows Server 2003 Service Pack 1 (SP1). It includes information that is relevant to different roles within an IT organization, including IT operations management and administrators. This guide contains high-level information that is required to plan a DNS operations environment, along with management-level knowledge of the DNS and IT processes that are required to operate it.

In addition, this guide contains more detailed procedures that are designed for operators who have varied levels of expertise and experience. Although the procedures provide operator guidance from start to finish, operators must have a basic proficiency with Microsoft Management Console (MMC) and snap-ins and know how to start administrative programs and access the command line. If operators are not familiar with DNS, it might be necessary for IT planners or managers to review the relevant operations in this guide and provide the operators with parameters or data that must be entered when the operations are performed.

**How to Use This Guide**
The operations areas are divided into the following types of content:

- Objectives are high-level goals for managing, monitoring, optimizing, and securing DNS. Each objective consists of one or more high-level tasks that describe how the objective is accomplished. In this guide, Managing Domain Name System Servers is an example of an objective.
- Tasks are used to group related procedures and provide general guidance for achieving the goals of an objective. In this guide, Modifying an Existing DNS Server is an example of a task.
- Procedures provide step-by-step instructions for completing tasks. In this guide, Change the name-checking method of a DNS server is an example of a procedure.

If you are an IT manager who will be delegating tasks to operators in your organization, you will want to:

- Read through the objectives and tasks to determine how to delegate permissions and whether you need to install tools before operators perform the procedures for each task.
- Before assigning tasks to individual operators, ensure that you have all the tools installed where operators can use them.
- When necessary, create “tear sheets” for each task that operators perform in your organization. Cut and paste the task and its related procedures into a separate document and then either print these documents or store them online, depending on the preference of your organization.

# Managing DNS

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide describes processes and procedures for improving the management of Windows Server 2003 Domain Name System (DNS) in your network infrastructure. Ensuring that DNS is functioning properly helps increase system availability for your users.

The following tasks for managing DNS are described in this objective:

- Managing Domain Name System Servers
- Managing Domain Name System Clients
- Managing Domain Name System Zones
- Managing DNS Resource Records

# Managing Domain Name System Servers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following tasks for managing Domain Name System (DNS) servers are described in this objective:

- Adding a Primary DNS Server to an Existing Zone
- Adding a Secondary DNS Server
- Modifying an Existing DNS Server
- Using Forwarders to Manage DNS Servers
- Removing a DNS Server from the Network
- Using DNS Aging and Scavenging

# Adding a Primary DNS Server to an Existing Zone

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are installing Domain Name System (DNS) with Active Directory, use the Active Directory Installation Wizard option to automatically install and configure a local DNS server. This option installs the DNS Server service on the computer where you are running the wizard, and it configures the computer's preferred DNS server setting to use the new local DNS server. Configure any other computers that join this domain to use this DNS server's Internet Protocol (IP) address as their preferred DNS server.

If you are installing DNS on a member server, use the procedures in this task.

It is recommended that you manually configure the computer to use a static IP address. If the DNS server is configured to use Dynamic Host Configuration Protocol (DHCP)–assigned dynamic addresses, when the DHCP server assigns a new IP address to the DNS server, the DNS clients that are configured to use that DNS server's previous IP address will be unable to resolve the previous IP address and locate the DNS server.

After you install a DNS server, you can decide how to administer it and its zones. Although you can use a text editor to make changes to server boot and zone files, this method is not recommended. The DNS console and the DNS command-line tool, Dnscmd, simplify maintenance of these files, and they should be used whenever possible. After you begin managing these files by using the console or the command line, editing them manually is not recommended.

You can administer DNS zones that are stored in Active Directory by using the DNS console or the Dnscmd command-line tool only. These zones cannot be administered by using a text editor.

If you uninstall a DNS server that hosts Active Directory-integrated zones, these zones are saved or deleted according to their storage type. For all storage types, the zone data is stored on other domain controllers or DNS servers. It is not deleted unless the DNS server that you uninstall is the last DNS server hosting that zone.

If you uninstall a DNS server hosting standard DNS zones, the zone files will remain in the systemroot\system32\Dns directory, but they will not be reloaded if the DNS server is reinstalled. If you create a new zone with the same name as an old zone, the old zone file is replaced with the new zone file.

When they write DNS server boot and zone data to text files, DNS servers use the Berkeley Internet Name Domain (BIND) file format that is recognized by legacy BIND 4 servers, not the more recent BIND 8 format.

Complete this task after you determine that you need to add a primary DNS server to your environment. For more information about planning a DNS infrastructure, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

To complete this task, perform one the following procedures:

- Install a new DNS server
- Configure a DNS server

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Install a new DNS server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to install Domain Name System (DNS) on a member server, which makes that server a DNS server.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

#### Installing a new DNS server

### To install a DNS server

1. Open the Windows Components Wizard.
2. In **Components**, select the **Networking Services** check box, and then click **Details**.
3. In **Subcomponents of Networking Services**, select the **Domain Name System (DNS)** check box, click **OK**, and then click **Next**.
4. If you are prompted to do so, in **Copy files from**, type the full path to the installation location, and then click **OK**.

   Required files are copied to your hard disk.

| **Note** |
| --- |
| To open the Windows Components Wizard, click **Start**, point to **Control Panel**, click **Add or Remove Programs**, and then click **Add/Remove Windows Components**. |

# Configure a DNS server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use these procedures to configure a new Domain Name System (DNS) server. When you finish configuring the server, you may need to complete additional tasks, such as enabling dynamic updates for its zones or adding resource records to its zones. See the other tasks in this guide to determine whether they are appropriate for your environment.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd tool at the command line.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Configuring a DNS server**

- Using the Windows interface
- Using a command line

### To configure a DNS server using the Windows interface

1. Open the DNS snap-in.
2. If necessary, add the applicable server to the console and connect to it.
3. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*Applicable DNS server*

4. On the **Action** menu, click **Configure a DNS Server**.
5. Follow the instructions in the Configure a DNS Server Wizard.

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To configure a DNS server using the command line

- At a command prompt,type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config** {*ZoneName*|**..AllZones**} *Property* {**1**|**0**}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/Config** | Specifies the configuration command. |
| {*ZoneName*\|**..AllZones**} | Specifies the name of the zone to be configured. To apply the configuration for all zones that are hosted by the specified DNS server, type **..AllZones**. |

| | |
|---|---|
| *Property* | Specifies the server property or zone property to be configured. There are different properties available for servers and zones. For a list of the available properties, at a command prompt type: **dnscmd/Config /help**. |
| {**1**\|**0**} | Sets configuration options to either **1** (on) or **0** (off). Note that some server and zone properties must be reset as part of a more complex operation. |

| **Note** |
|---|
| To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**. |

# Adding a Secondary DNS Server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain Name System (DNS) design specifications recommend that at least two DNS servers be used to host each zone. For standard, primary zones, a secondary server is required to add and configure the zone so that it appears to other DNS servers in the network. For directory-integrated, primary zones, secondary servers are supported but not required for this purpose. For example, two DNS servers running on domain controllers can be redundant primary servers for a zone. They can provide the same benefits as adding a secondary server while also providing additional benefits.

Secondary servers can be used to offload DNS query traffic in areas of the network where a zone is heavily queried. In addition, if a primary server is unavailable, a secondary server can provide some name resolution in the zone until the primary server is available.

If you add a secondary server, try to locate it as close as possible to clients that have a high demand for names that are used in the zone. Also, consider placing secondary servers across a router, either on other subnets (if you use a routed local area network (LAN)) or across wide area network (WAN) links. This constitutes a good use of a secondary server as a local backup in scenarios in which an intermediate network link becomes the point of failure between DNS servers and clients that use the zone.

Because a primary server always maintains the master copy of updates and changes to the zone, a secondary server relies on DNS zone transfer mechanisms to obtain its information and keep the information current. Issues such as zone transfer methods — using either full or incremental zone transfers — are more applicable when you use secondary servers.

When you consider the impact of zone transfers that are caused by secondary servers, consider their advantage as a backup source of information, and measure this against the added cost that they impose on your network infrastructure. A simple rule is that for each secondary server that you add, network usage (because of added zone replication traffic) increases, and so does the time that is required to synchronize the zone at all secondary servers.

Secondary servers are used most heavily for forward lookup zones. If you are using reverse lookup zones, it is not necessary to add as many secondary servers for those zones. Typically, a secondary server for a reverse lookup zone is not used outside the network and subnet that correspond to the reverse zone.

To complete this task, perform the following procedure:

- Add a secondary server to a zone

# Add a secondary server to a zone

Published: March 2, 2005

Updated: November 18, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To add a secondary server to an existing zone, you must have network access to the server that acts as the master server for this server and its use of the zone. The master server acts as the source for zone data. It is contacted periodically to assist in renewing the zone and to transfer zone updates whenever they are needed.

You can perform this procedure by using the DNS console or by using the Dnscmd command-line tool. This procedure can be performed on the secondary DNS server, or on a computer with permission to manage the secondary DNS server. To add a secondary server to multiple zones, you must repeat this procedure for each zone.

| **Important** |
| --- |
| Before you add a secondary server to a zone, you must allow zone transfers from the primary to the secondary server. For more information, see Modify DNS zone transfer settings. |

### Administrative credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

#### Adding a secondary server to a zone

- Using the Windows interface
- Using the command line

### To add a secondary server to a zone using the Windows interface

1. Click **Start**, point to **Administrative Tools**, and then click **DNS**.
2. In the console tree, click the applicable Domain Name System (DNS) server.
3. On the **Action** menu, click **New Zone**.
4. Follow the instructions in the New Zone Wizard. When you add the zone, select **Secondary zone** as the zone type.

### To add a secondary server to a zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **Dnscmd** *ServerName* **/ZoneAdd** *ZoneName* **/Secondary** *MasterIPaddress***...** [**/file** *FileName*]

| Value | Description |
| --- | --- |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| *ZoneName* | Specifies the fully qualified domain name (FQDN) of the secondary zone that you are adding. The zone name must be the same as the name of the primary zone from which the secondary zone is created. |
| *MasterIPaddress* | Specifies one or more IP addresses for the secondary zone master servers, from which it copies zone data. |
| *FileName* | Specifies the name of the file to use for creating the secondary zone. |

In the following example, zone transfers are first allowed from the primary DNS server *primarydns.contoso.com* at

*10.0.0.2* to the secondary server *secondarydns.contoso.com* at *11.0.0.2*. Next, the secondary DNS server is added to the zone *secondtest.contoso.com*.

**Dnscmd** *primarydns.contoso.com* **/zoneresetsecondaries** *secondtest.contoso.com* **/securelist** *11.0.0.2*

**Dnscmd** *secondarydns.contoso.com* **/zoneadd** *secondtest.contoso.com* **/secondary** *10.0.0.2*

For more information about using dnscmd, see Dnscmd Syntax.

# Modifying an Existing DNS Server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You may need to modify or update the configuration of your Domain Name System (DNS) servers for various reasons. For example, you may need to change the name-checking method of a DNS server to allow the DNS server to resolve non–Request for Comments (RFC)–compliant names. In addition, you may need to modify or update a DNS server in the process of troubleshooting or optimizing it.

**Task requirements**

To begin this task, perform the following requirements:

- Install Dnscmd.

To complete this task, perform one of the following procedures:

- Start, stop, pause, or restart a DNS server
- Manually update DNS server data files
- Clear the DNS server names cache
- Change the boot method of a DNS server
- Change the name-checking method of a DNS server
- Restore DNS server default preferences

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Start, stop, pause, or restart a DNS server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to start, stop, pause, or restart Domain Name System (DNS).

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To start, stop, pause, or restart a DNS server**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.
3. On the **Action** menu, point to **All Tasks**, and then click one of the following:
   - To start the DNS service on this server, click **Start**.
   - To stop the DNS service on this server, click **Stop**.
   - To interrupt the DNS service on this server, click **Pause**.
   - To stop and then automatically restart the DNS service on this server, click **Restart**.

| **Note** |
| --- |
| To open the DNS management console, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

| **Note** |
| --- |
| If you want to resume the service after you pause or stop it, on the **Action** menu, point to **All Tasks**, and then click **Resume** to immediately resume the service. |

# Manually update DNS server data files

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool. Use the DNS snap-in for standard Domain Name System (DNS) zones and the Dnscmd command-line tool for Active Directory–integrated zones.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Manually updating DNS server data files**

- Using the Windows interface
- Using the command line

**To manually update DNS server data files using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.
3. On the **Action** menu, click **Update Server Data Files**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To manually update DNS server data files using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneUpdateFromDs** *ZoneName*

| Value | Description |
| --- | --- |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName* | Specifies the name of the zone to which you want to set aging and scavenging. |

# Clear the DNS server names cache

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Clearing the DNS server names cache**

- Using the Windows interface
- Using the command line

### To clear the DNS server names cache using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, click the applicable Domain Name System (DNS) server.
3. On the **Action** menu, click **Clear Cache**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To clear the DNS server names cache using the command line

- At a command prompt, type the following, and then press ENTER:

  **dnscmd** *ServerName* **/clearcache**

| Value | Description |
| --- | --- |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |

# Change the boot method of a DNS server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, Domain Name System (DNS) servers use information that is stored in the registry to initialize the service and load any zone data for use at the server. In addition, you can configure the DNS server to boot from a file. Or, in Active Directory environments, you can supplement local registry data with zone data that is retrieved for directory-integrated zones that are stored in the Active Directory database. If you use the file method, the file must be a text file named Boot, which is located on the computer in the %Systemroot%\Windows\System32\Dns folder.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To change the boot method of a DNS server**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable DNS server, and then click **Properties**.
3. Click the **Advanced** tab.
4. In the **Load zone data on startup** list, select one of the following:
    - **From registry**
    - **From file**
    - **From Active Directory and registry**

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Change the name-checking method of a DNS server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The DNS Server service supports three different possible methods for checking the names that it receives and processes during normal operations:

- Strict RFC (ANSI): This method strictly enforces Request for Comments (RFC)–compliant naming rules for all Domain Name System (DNS) names that the server processes. Names that are not RFC compliant are treated as erred data by the DNS server.

- Non RFC (ANSI): This method allows names that are not RFC compliant, such as names that use American Standard Code for Information Interchange (ASCII) characters but are not compliant with RFC host naming requirements, to be used with the DNS server.

- Multibyte (UTF8): This method allows names that use the Unicode 8-bit translation encoding scheme, which is a proposed RFC draft, to be used with the DNS server.

By default, the DNS server uses the Multibyte (UTF8) method to check names.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To change the name-checking method of a DNS server**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable DNS server, and then click **Properties**.
3. Click the **Advanced** tab.
4. In the **Name checking** list, click **Strict RFC (ANSI)**, **Non RFC (ANSI)**, **Multibyte (UTF8)**, or **All names**.

   **All names** enables all three name-checking methods.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Restore DNS server default preferences

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to configure the Domain Name System (DNS) server with the initial configuration settings that it had following installation. These initial configuration settings are listed in the following table.

| Property | Settings |
|---|---|
| Disable recursion | Off |
| BIND secondaries | On |
| Fail on load if bad zone data | Off |
| Enable round robin | On |
| Enable netmask ordering | On |
| Secure cache against pollution | On |
| Name checking | Multibyte (UTF8) |
| Load zone data on startup | From Active Directory and registry |
| Enable automatic scavenging of stale records | Off |

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To restore DNS server default preferences**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable DNS server, then click **Properties**.
3. Click the **Advanced** tab.
4. Click **Reset to Default**, and then click **OK**.

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Using Forwarders to Manage DNS Servers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you want to use forwarders to manage the Domain Name System (DNS) traffic between your network and the Internet, configure your network firewall to allow only one DNS server to communicate with the Internet. When you have configured the other DNS servers in your network to forward queries that they cannot resolve locally to that DNS server, it will act as your forwarder.

Consider the following tips for efficient forwarder configuration and use:

- **Keep forwarder configuration uncomplicated.** For every DNS server that is configured with a forwarder, queries can be sent to a number of different places. Each forwarder and each conditional forwarder must be administered for the benefit of DNS client queries, and this process can be time consuming. Use forwarders strategically — where they are needed the most — for example, for resolving offsite queries or for sharing information between namespaces.

- **Avoid chaining your forwarders.** If you have configured a DNS server named server1 to forward queries for wingtiptoys.corp.com to DNS server server2, do not configure server2 to forward queries for wingtiptoys.corp.com to DNS server server3. This is an inefficient resolution process, and it can result in errors if server3 is accidentally configured to forward queries for wingtiptoys.corp.com to server1.

- **Do not concentrate too great a load on forwarders.** The recursive queries that forwarders send to the Internet can require a significant amount of time to answer because of the nature of the Internet. When large numbers of internal DNS servers use these forwarders for Internet queries, the server can experience a substantial concentration of network traffic. If network load is an issue, use more than one forwarder and distribute the load between them.

- **Do not create inefficient resolution by using forwarders.** The DNS server attempts to forward domain names according to the order in which the domain names are configured in the DNS console. For example, a DNS server in Seattle may be incorrectly configured to forward a query to a server in London, instead of another server in Seattle, because the server in London is higher up in the forwarders list. This decreases the efficiency of name resolution on the network. Evaluate your network's forwarding configurations periodically to see if there are similar, inefficient configurations.

To complete this task, perform the following procedure:

- Configure forwarders for a DNS server

# Configure forwarders for a DNS server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you use this procedure to configure a conditional forwarder, note that you cannot use a domain name in a conditional forwarder if the DNS server hosts a primary zone, secondary zone, or stub zone for that domain name. For example, if a DNS server is authoritative for the domain name wingtiptoys.corp.com (that is, it hosts the primary zone for that domain name), you cannot configure that DNS server with a conditional forwarder for wingtiptoys.corp.com.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Configuring forwarders for a DNS server**

- Using the Windows interface
- Using the command line

**To configure forwarders for a DNS server using the Windows graphical user interface**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.
3. On the **Action** menu, click **Properties**.
4. On the **Forwarders** tab, under **DNS domain**, click a domain name.

| **Note** |
| --- |
| To create a new domain name, click **New**, and then, under **DNS domain**, type the domain name. |

5. Under **Selected domain's forwarder IP address list**, type the Internet Protocol (IP) address of a forwarder, and then click **Add**.

| **Note** |
| --- |
| When you specify a conditional forwarder, select a DNS domain name before you enter an IP address. |

6. By default, the DNS server waits five seconds for a response from one forwarder IP address before trying another forwarder IP address. In **Number of seconds before forward queries time out**, you can change the number of seconds that the DNS server waits. If the overall recursion timeout (by default, 15 seconds) is exceeded before all forwarders are exhausted, the DNS server fails the query. If the overall recursion timeout has not been exceeded and the server exhausts all forwarders, it attempts standard recursion.
7. If you want the DNS server to only use forwarders and not attempt any further recursion if the forwarders fail, select the **Do not use recursion for this domain** check box.

| **Note** |
| --- |
| You can disable recursion for the DNS server so that it does not perform recursion on any query. If you disable recursion on the DNS server, you will not be able to use forwarders on the same server. |

| **Note** |
| --- |

To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**.

**To configure forwarders for a DNS server using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneAdd** *ZoneName* **/Forwarder** *MasterIPaddress* **...** [**/TimeOut** *Time*][**/Slave**]

| Value | Description |
|---|---|
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| *ZoneName* | Specifies the fully qualified domain name (FQDN) of the zone. |
| *MasterIPaddress* | Specifies a space-separated list of one or more IP addresses of the DNS servers where queries for ZoneName are forwarded. You can specify a list of space-separated IP addresses. |
| *Time* | Specifies the value for the **/TimeOut** parameter. The value is in seconds. The default timeout is five seconds. |

# Removing a DNS Server from the Network

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To remove a DNS server from the network, perform the following procedures to make changes in zones where the server is configured as an authoritative server for the zone:

1. Use the Delete a resource record procedure to remove the address (A) resource record for the server.

2. Use the Modify an existing resource record procedure to update the name server (NS) records, in zones where the server is configured as authoritative, to no longer include the server by name (as it appeared in the A record that was deleted in procedure 1).

3. If the server is the primary server for a standard zone, use the Modify the SOA record for a zone procedure to revise the owner field of the start of authority (SOA) resource record for the zone to point to the new primary DNS server for the zone. (If the zone is a directory-integrated zone, this procedure is not necessary.)

4. Use the Verify a zone delegation procedure to check the parent zone to ensure that any records (NS or A resource records) that are used for delegation to the zone are revised and that they no longer point to the removed server.

# Delete a resource record

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to delete a resource record from a zone. Pointer (PTR) resource records are deleted automatically if the corresponding address (A) resource record is deleted.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Deleting a resource record**

- Using the Windows interface
- Using the command line

### To delete a resource record using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.
3. In the details pane, right-click the resource record that you want to delete, and then click **Delete**.
4. When you are asked to confirm that you want to delete the selected resource record, click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To delete a resource record using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/RecordDelete** *ZoneName NodeName RRType RRData* [**/f**]

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordDelete** | Required. Deletes a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| *RRType RRData* | Required. Specifies the type of resource record to add, followed by the data to be contained in the resource record. |

| Resource record type | Resource record data |
|---|---|
| A | *IPAddress* |
| NS,CNAME,MB,MD,PTR,MF,MG,MR | *HostName*\|*DomainName* |
| MX,RT,AFSDB | *PreferenceServerName* |
| SRV | *Priority Weight Port HostName* |
| SOA | *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL* |
| AAAA | *Ipv6Address* |
| TXT,X25,HINFO,ISDN | *String* [*String*] |
| MINFO,RP | *MailboxName ErrMailboxName* |
| WKS | *Protocol IPAddress Service...* |
| WINS | *MapFlag LookupTimeout CacheTimeout IPAddress...* |
| WINSR | *MapFlag LookupTimeout CacheTimeout RstDomainName* |

| Value | Description |
|---|---|
| *IPAddress* | Specifies a standard IP address, for example, 255.255.255.255. |
| *ipv6Address* | Specifies a standard IPv6 address, for example, 1:2:3:4:5:6:7:8. |
| *Protocol* | Specifies the transmission protocol: UDP or TCP. |
| *Service* | Specifies a standard service, for example, domain, smtp. |
| *HostName*\|*DomainName* | Specifies the FQDN of a resource record that is located in the DNS namespace. |
| **/f** | Specifies that the command is executed without asking for confirmation. If you omit this parameter, you are prompted to confirm the deletion of the resource record. |

# Modify an existing resource record

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to modify an existing resource record in a zone. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Modifying an existing resource record**

- Using the Windows interface
- Using the command line

**To modify an existing resource record using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.
3. In the details pane, right-click the resource record that you want to modify, and then click **Properties**.
4. In **Properties**, edit the properties that can be modified.

   If necessary, you can view and modify advanced resource record properties with the DNS snap-in. To display advanced properties, on the **View** menu, click **Advanced**.
5. When you have finished modifying the record, click **OK**.

| **Note** |
| --- |
| When advanced view options are enabled, you can modify additional settings for an existing resource record, such as its record-specific Time to Live (TTL). |

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To modify an existing resource record using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName***/RecordAdd***ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] *RRTypeRRData*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Required. Adds a new resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |

| | |
|---|---|
| *RRType RRData* | Required. Specifies the type of resource record to add, followed by the data to be contained in the resource record. |

| **Resource record type** | **Resource record data** |
|---|---|
| A | *IPAddress* |
| NS,CNAME,MB,MD,PTR,MF,MG,MR | *HostName*\|*DomainName* |
| MX,RT,AFSDB | *Preference ServerName* |
| SRV | *Priority Weight Port HostName* |
| SOA | *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL* |
| AAAA | *Ipv6Address* |
| TXT,X25,HINFO,ISDN | *String* [*String*] |
| MINFO,RP | *MailboxName ErrMailboxName* |
| WKS | *Protocol IPAddress Service...* |
| WINS | *MapFlag LookupTimeout CacheTimeout IPAddress...* |
| WINSR | *MapFlag LookupTimeout CacheTimeout RstDomainName* |

| **Value** | **Description** |
|---|---|
| *IPAddress* | Specifies a standard IP address, for example, 255.255.255.255. |
| *ipv6Address* | Specifies a standard IPv6 address, for example, 1:2:3:4:5:6:7:8. |
| *Protocol* | Specifies the transmission protocol: UDP or TCP. |
| *Service* | Specifies a standard service, for example, domain, smtp. |
| *HostName*\|*DomainName* | Specifies the FQDN of a resource record that is located in the DNS namespace. |

# Modify the SOA record for a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to change settings for the start of authority (SOA) resource record for a zone. The settings that are applied for the SOA record affect how zone transfers are made between servers.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Modifying the SOA record for a zone**

- Using the Windows interface
- Using a command line

### To modify the SOA record for a zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. Click the **Start of Authority (SOA)** tab.
4. Modify the properties for the SOA record as needed.

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To modify the SOA record for a zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/RecordAdd** *ZoneName NodeName* [**/Aging**] [*Ttl*] **SOA***PrimSvrAdmin Serial#\ Refresh Retry Expire MinTTL*

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Adds or modifies a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName*, or you can type @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this parameter is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. The default TTL is defined in SOA resource record. |
| **SOA** | Required. Specifies the type of resource record that you are modifying. |

| | |
|---|---|
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *PrimSvr* | Required. Specifies the FQDN name of the server that is the primary source for information about the zone, for example, nameserver.place.sales.wingtiptoys.com.. |
| *Admin* | Required. Specifies the name of the DNS administrator for the zone, for example, postmaster.nameserver.place.sales.wingtiptoys.com.. |
| *Serial#* | Required. Specifies the version information for the zone. |
| *Refresh* | Required. Specifies the refresh interval for the zone. The standard setting is 3600 seconds (one hour). |
| *Retry* | Required. Specifies the retry interval for the zone. The standard setting is 600 seconds (10 minutes). |
| *Expire* | Required. Specifies the expire interval for the zone. The standard setting is 86400 seconds (one day). |
| *MinTTL* | Required. Specifies the minimum TTL value. This is the length of time that is used by other DNS servers to determine how long to cache information for a record in the zone before expiring and discarding it. The standard setting is 3600 seconds (one hour). |

**Note**

To modify any specific SOA record's values using Dnscmd, you must specify all the SOA values (*PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*).

# Verify a zone delegation

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administrative credentials**

You do not need administrative credentials to perform this task. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

**To verify a zone delegation**

1. At a command prompt, type the following command, and then press ENTER:

   **nslookup** *RootServerIpAddress*

2. Type the following command, and then press ENTER:

   **nslookup**

3. At the next prompt, type the following command, and then press ENTER:

   **set norecurse**

4. At the next prompt, type the following command, and then press ENTER:

   **set q=NS**

5. Type the fully qualified domain name (FQDN) for the failed name.

   Use the trailing period (**.**) when you type the name. If zone delegations are set correctly, a list of name server (NS) resource records for delegated servers is returned in the response.

6. If the NS query response contains no names or Internet Protocol (IP) addresses for delegated servers, type **q=ns**, and then query again using the FQDN for the parent zone of the failed name.

   For example, if the failed name that you used in the previous step was sales.wingtiptoys.com, query for wingtiptoys.com.

7. If the response contains NS resource records, but no host address (A) resource records, type **set recurse**, and then query individually for any of the A resource records of the servers that are listed in the NS resource records.

   If, for each NS resource record that you encounter in a zone, you do not find at least one valid IP address in an A resource record, you have a broken delegation.

8. Either fix the broken delegation or retry the delegation test that is described in the previous step and use a different IP address.

   If more than one A resource record or IP address is found, use it to repeat the delegation test described in the previous step. To fix a delegation, add or update an A resource record in the parent zone with a valid IP address for a correct DNS server for the delegated zone.

| Value | Description |
|---|---|
| *RootServerIpAddress* | The IP address of a valid root server for your network. |
| **set norecursion** | Instructs the root server to not perform recursion on your query. |
| **set q=NS** | Sends the query for NS resource records to the root server. |

# Using DNS Aging and Scavenging

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Aging and scavenging of stale resource records are features of Domain Name System (DNS) that are available when you deploy your server with primary zones.

Where aging and scavenging are available, you can use the DNS snap-in to perform the following related tasks for your DNS servers and any directory-integrated zones that they load:

- Enable or disable the use of scavenging at a DNS server
- Enable or disable the use of scavenging for selected zones at the DNS server
- Modify the no-refresh interval, either as a server default or by specifying an overriding value at selected zones
- Modify the refresh interval, either as a server default or by specifying an overriding value at selected zones
- Specify whether periodic scavenging occurs automatically at the DNS server for any of its eligible zones and how often these operations are repeated
- Manually initiate a single scavenging operation for all eligible zones at the DNS server
- View other related properties, such as the time stamp for individual resource records or the start-scavenging time for a specified zone

**Enabling Scavenging of Stale Resource Records**

By default, aging and scavenging features are disabled on all DNS servers and any of their zones. Before using these features, you should configure the following settings for the applicable server and its directory-integrated zones:

- **Server aging and scavenging properties for determining the use of these features on a server-wide basis.** These settings are used to determine the affect of zone-level properties for any directory-integrated zones that are loaded at the server. For more information, see Set aging and scavenging properties for a DNS server.
- **Zone aging and scavenging properties for determining the use of these features on a per zone basis.** When zone-specific properties are set for a selected zone, these settings apply only to the applicable zone and its resource records. Unless these zone-level properties are otherwise configured, they inherit their defaults from comparable settings that are maintained in server aging and scavenging properties. For more information, see Set aging and scavenging properties for a zone.

  **Caution** Enabling aging and scavenging for use with standard primary zones modifies the format of zone files. This change does not affect zone replication to secondary servers, but the modified zone files cannot be loaded by other versions of DNS servers.

**Modifying No-refresh Intervals**

When the no-refresh interval is in effect for a specific resource record, attempts to dynamically refresh its time stamp are suppressed by the DNS server. This aspect of the aging and scavenging mechanism prevents unnecessary refreshes from being processed by the server for aged resource records. These early refresh attempts, if not handled in this way, might otherwise increase Active Directory replication traffic related to processing DNS zone changes.

To ensure that records do not refresh prematurely, keep the no-refresh interval comparable in length to the current refresh interval for each resource record. For example, if you increase the refresh interval to a higher value, you can similarly increase the no-refresh interval.

In most instances, the default interval of seven days is sufficient and does not need to be changed.

**Modifying Refresh Intervals**

When the refresh interval is in effect for a resource record, attempts to dynamically refresh its time stamp are accepted and processed by the DNS server. When you set this interval, it is important that the length of time used be greater than the maximum possible refresh period for any resource records that are contained in the zone. This period is equal to the maximum amount of time that it might take the record to be refreshed under normal network conditions, based on the specific source generating the record refresh.

For example, the following table shows default refresh periods for various services that are known to register and refresh records dynamically in DNS.

| Service | Default refresh period |
|---|---|
| Net logon | 24 hours |
| Clustering | 24 hours |
| DHCP client | 24 hours<br><br>The DHCP Client service sends dynamic updates for the DNS records. This includes both computers that obtain a leased Internet Protocol (IP) address by using Dynamic Host Configuration Protocol (DHCP) and computers that are configured statically for TCP/IP. |
| DHCP server | Four days (half of the lease interval, which is eight days by default).<br><br>Refresh attempts are made only by DHCP servers that are configured to perform DNS dynamic updates on behalf of their clients, for example, Windows 2000 Server DHCP servers and Windows Server 2003 DHCP servers. The period is based on the frequency in which DHCP clients renew their IP address leases with the server. Typically, this occurs when 50 percent of the scope lease time has elapsed. If the DNS default scope lease duration of eight days is used, the maximum refresh period for records that are updated by DHCP servers on behalf of clients is four days. |

By default, the refresh interval is seven days. In most instances, this value is sufficient and does not need to be changed, unless any resource records in the zone are refreshed less often than once every seven days.

**Automated and Manually Initiated Scavenging**

Although scavenging start time and other factors determine when zones and records are actually eligible for scavenging, you can initiate scavenging by using either of two methods:

- **Automatic scavenging.** Automatic scavenging specifies that aging and scavenging of stale records is to be performed automatically by the server for any eligible zones at a recurring interval that is specified as the scavenging period. When you use automatic scavenging, the default scavenging period is one day, and the minimum allowed value that you can use for the scavenging period is one hour. For more information, see Configure automatic scavenging of stale resource records.

- **Manual scavenging.** Manual scavenging specifies that aging and scavenging of stale records is to be performed as a nonrecurring operation for any eligible zones at the server. For more information, see Start scavenging of stale resource records.

**Modifying Time-Stamp Values**

For resource records that are not added dynamically to DNS zone data, a record time-stamp value of zero is applied, which prevents these records from aging or removal during scavenging.

You can, however, reset record properties manually to enable any statically entered records to qualify for the aging and scavenging process. If you do this, the record will be deleted based on the modified time-stamp value, at which point you might need to re-create a record if it is still needed.

For more information, see Reset aging and scavenging properties for a specific resource record.

To complete this task, perform the following procedures:

1. Set aging and scavenging properties for a DNS server

2. Set aging and scavenging properties for a zone

3. Configure automatic scavenging of stale resource records

4. Start scavenging of stale resource records

5. Reset aging and scavenging properties for a specific resource record

# Set aging and scavenging properties for a DNS server

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The settings for server aging and scavenging properties determine the effect of zone-level properties for any directory-integrated zones that are loaded at the server.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Setting aging and scavenging properties for a DNS server**

- Using the Windows interface
- Using the command line

### To set aging and scavenging properties for a DNS server using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable Domain Name System (DNS) server, and then click **Set Aging/Scavenging for All Zones**.
3. Select the **Scavenge stale resource records** check box.
4. Modify other aging and scavenging properties as needed

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To set aging and scavenging properties for a DNS server using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config** {**/ScavengingInterval** *Value*|**/DefaultAgingState** *Value*|**/DefaultNoRefreshInterval** *Value*|**/DefaultRefreshInterval** *Value*}

| Value | Description |
| --- | --- |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *Value* | For **/ScavengingInterval**, type a value in hours. The default is 168 hours (one week). For **/DefaultAgingState**, type **1** to enable aging for new zones when they are created. Type **0** to disable aging for new zones. For **/DefaultNoRefreshInterval**, type a value in hours. The default is 168 hours (one week). For **/DefaultRefreshInterval**, type a value in hours. The default is 168 hours (one week). |

# Set aging and scavenging properties for a zone

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The settings for zone aging and scavenging properties determine the use of these features on a per-zone basis. When you set zone-specific properties for a selected zone, these settings apply only to the applicable zone and its resource records. Unless these zone-level properties are otherwise configured, they inherit their defaults from comparable settings that are maintained in server aging and scavenging properties.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Setting aging and scavenging properties for a zone**

- Using the Windows interface
- Using the command line

### To set aging and scavenging properties for a zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, click **Aging**.
4. Select the **Scavenge stale resource records** check box.
5. Modify other aging and scavenging properties as needed.

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To set aging and scavenging properties for a zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config** {*ZoneName*|**..AllZones**} {**/Aging** *Value*|**/RefreshInterval** *Value*|**/NoRefreshInterval** *Value*}

| Value | Description |
|---|---|
| *ServerName* | Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName*\|**..AllZones** | Specifies the name of the zone to which you want to set aging and scavenging. To apply the operation to all zones, use **..AllZones**. |
| *Value* | For **/Aging**, type **1** to enable aging. Type **0** to disable aging. For **/RefreshInterval**, type a value in hours. The default setting is 168 hours (one week). For **/NoRefreshInterval**, type a value in seconds. The standard setting is 3600 seconds (one hour). |

# Configure automatic scavenging of stale resource records

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To configure automatic scavenging of stale resource records**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable Domain Name System (DNS) server, and then click **Properties**.
3. Click the **Advanced** tab.
4. Select the **Enable automatic scavenging of stale records** check box.
5. To adjust the scavenging period, in **Scavenging period**, select an interval in the drop-down list (either hours or days), and then type a number in the text box.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Start scavenging of stale resource records

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Starting scavenging of stale resource records**

- Using the Windows interface
- Using the command line

**To start scavenging of stale resource records using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable Domain Name System (DNS) server, and then click **Scavenge Stale Resource Records**.
3. When you are prompted to confirm that you want to scavenge all stale resource records on the server, click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To start scavenging of stale resource records using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/StartScavenging**

| Value | Description |
| --- | --- |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |

# Reset aging and scavenging properties for a specific resource record

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is used only for resource records that are registered dynamically. For records that you add to a zone manually, a time-stamp value of zero always applies to the record, which excludes it from the scavenging process.

| **Note** |
| --- |
| Scavenging and aging properties for name server (NS) and start of authority (SOA) resource records are reset in the properties of the zone, not in the properties of the resource record. |

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Resetting aging and scavenging properties for a specific resource record**

- Using the Windows interface
- Using the command line

**To reset aging and scavenging properties for a specific resource record using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.
3. In the details pane, double-click the resource record for which you want to reset scavenging and aging properties.
4. Depending on the how the resource record was originally added to the zone, do one of the following:
   - If the record was added dynamically using dynamic update, clear the **Delete this record when it becomes stale** check box to prevent the record's aging or potential removal during the scavenging process. If dynamic updates to this record continue to occur, the Domain Name System (DNS) server will always reset this check box so that the dynamically updated record can be deleted.
   - If you added the record manually, select the **Delete this record when it becomes stale** check box to permit the record's aging or potential removal during the scavenging process.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To reset aging and scavenging properties for a specific resource record using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config** {*ZoneName*|**..AllZones**} **/ScavengingInterval** *Value*

| **Value** | **Description** |
| --- | --- |

| | |
|---|---|
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName*\|**..AllZones** | Specifies the fully qualified domain name (FQDN) of the zone. To configure all zones that are hosted on the specified DNS server to allow dynamic updates, type **..AllZones**. |
| *Value* | The new value for the scavenging interval, specified in hours. The default is 168 hours (one week). |

# Managing Domain Name System Clients

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following tasks are described in this objective:

Configuring DNS Client Settings for DNS Operations

Managing the DNS Client Resolver Cache

Renewing DNS Client Registration

# Configuring DNS Client Settings for DNS Operations

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain Name System (DNS) configuration involves the following tasks when TCP/IP properties are configured for each computer:

- Setting a DNS computer name or host name for each computer. For example, in the fully qualified domain name (FQDN) wkstn1.sales.wingtiptoys.com., the DNS computer name is wkstn1.
- Setting a primary DNS suffix for the computer, which is placed after the computer name or host name to form the FQDN. Using the previous example, the primary DNS suffix is sales.wingtiptoys.com.
- Setting a list of DNS servers for clients to use when resolving DNS names, such as a preferred DNS server, and any alternate DNS servers to use if the preferred server is not available.
- Setting the DNS suffix search list or search method to be used by the client when it performs DNS query searches for short, unqualified domain names.

These tasks are discussed in more detail in each of the following sections.

**Setting Computer Names**

When you set a computer name for DNS, it is useful to think of the name as the leftmost portion of a fully qualified domain name (FQDN). For example, in wkstn1.sales.wingtiptoys.com., wkstn1 is the computer name.

You can configure all Windows DNS clients with a computer name that is based on any of the standard supported characters that are defined in Request for Comments (RFC) 1123, "Requirements for Internet Hosts -- Application and Support." These characters include the following:

- Uppercase letters: A through Z
- Lowercase letters: a through z
- Numbers: 0 through 9
- Hyphens (-)

If you are supporting both network basic input/output system (NetBIOS) and DNS namespaces on your network, you can use a different computer name in each namespace. However, it is recommended that, wherever possible, you try to use computer names that are 15 characters or less and that you follow the RFC 1123 naming requirements described in the previous paragraph.

By default, the leftmost label in the FQDN for clients equals the NetBIOS computer name, unless this label is 16 or more characters, which is the maximum for NetBIOS names. When the computer name exceeds the maximum length for NetBIOS, the NetBIOS computer name is truncated based on the full label that is specified.

Before you configure computers with varying DNS and NetBIOS names, consider the following issues and their implications for your deployment:

If Windows Internet Name Service (WINS) lookup is enabled for zones that are hosted by your DNS servers, you must use the same name for both NetBIOS and DNS computer naming. Otherwise, the results of clients attempting to query and resolve the names of these computers will be inconsistent.

If you have an investment in using NetBIOS names to support legacy Microsoft networking technology, it is recommended that you revise NetBIOS computer names that are used on your network to prepare for migration to a standard DNS-only environment. This prepares your network well for long-term growth and interoperability with future naming requirements. For example, if you use the same computer name for both NetBIOS and DNS resolution, consider converting any special characters such as the underscore (_) in your current NetBIOS names that do not comply with DNS naming standards. While these

characters are permitted in NetBIOS names, they are more often incompatible with traditional DNS host naming requirements and most existing DNS resolver client software.

| **Note** |
| --- |
| Although the use of the underscore (_) in DNS host names or in host address (A) resource records has traditionally been prohibited by DNS standards, the use of underscores in service-related names — such as those used for service locator (SRV) resource records — has been proposed to avoid naming collisions in the Internet DNS namespace. |

In addition to DNS standard naming conventions, Windows Server 2003 DNS supports the use of extended American Standard Code for Information Interchange (ASCII) and Unicode characters. However, because most resolver software that is written for other platforms (such as UNIX) is based on Internet DNS standards, this enhanced character support can be used only in private networks with computers running Windows 2000 or Windows Server 2003 DNS.

The initial setup of DNS and TCP/IP displays a warning to suggest a standard DNS name if a nonstandard DNS name is entered.

By default, computers and servers use DNS to resolve any name that is greater than 15 characters in length. If the name is less than or equal to 15 characters, both NetBIOS and DNS name resolution can be attempted and used to resolve the name.

**Setting Domain Names**

The domain name is used with the client computer name to form the FQDN, which is also known as the full computer name. In general, the DNS domain name is the remainder of the FQDN that is not used as the unique host name for the computer.

For example, the DNS domain name for a client computer can be defined as the following: If the FQDN is wkstn1.sales.wingtiptoys.com, the domain name is the sales.wingtiptoys.com portion of this name.

DNS domain names have two variations: a DNS name and a NetBIOS name. The full computer name (a fully qualified DNS name) is used during querying and location of named resources on your network. For earlier-version clients, the NetBIOS name is used to locate various types of NetBIOS services that are shared on your network.

The Net Logon service is an example of a service that shows the need for both NetBIOS and DNS names. In Windows Server 2003 DNS, the Net Logon service on a domain controller registers its SRV resource records on a DNS server. For Windows NT Server 4.0 and earlier operating systems, domain controllers register a DomainName entry in WINS to perform the same registration and to advertise their availability for providing authentication service to the network.

When a client computer is started on the network, it uses the DNS resolver to query a DNS server for SRV records for its configured domain name. This query is used to locate domain controllers and provide logon authentication for accessing network resources. A client or a domain controller on the network optionally uses the NetBIOS resolver service to query WINS servers, attempting to locate DomainName [1C] entries to complete the logon process.

Your DNS domain names should follow the same standards and recommended practices that apply to DNS computer naming described in the previous section. In general, acceptable naming conventions for domain names include the use of letters A through Z, numerals 0 through 9, and the hyphen (-). The period (.) in a domain name is always used to separate the discrete parts of a domain name, commonly known as labels. Each label corresponds to an additional level that is defined in the DNS namespace tree.

For most computers, the primary DNS suffix that is configured for the computer can be the same as its Active Directory domain name, although the two values can also be different.

| **Important** |
| --- |
| By default, the primary DNS suffix portion of a computer's FQDN must be the same as the name of the Active Directory domain where the computer is located. To allow different primary DNS suffixes, a domain administrator may establish a restricted list of allowed suffixes by creating the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is created and managed by the domain administrator using Active Directory Service Interfaces (ADSI) or the Lightweight Directory Access Protocol (LDAP). |

**Configuring a DNS Servers List**

For DNS clients to operate effectively, a prioritized list of DNS name servers must be configured for each computer to use when processing queries and resolving DNS names. In most cases, the client computer contacts and uses its preferred DNS server, which is the first DNS server on its locally configured list. The client computer contacts and uses listed alternate DNS servers when the preferred server is not available. For this reason, it is important that the preferred DNS server be appropriate

for continuous client use under normal conditions.

| Note |
| --- |
| For computers running Windows XP, the DNS server list is used by clients only to resolve DNS names. When clients send dynamic updates — for example, when they change their DNS domain name or a configured Internet Protocol (IP) address — they might contact these servers or other DNS servers as needed to update their DNS resource records. |

By default, the DNS client on Windows XP does not attempt dynamic update over a Remote Access Service (RAS) or virtual private network (VPN) connection. To modify this configuration, you can modify the advanced TCP/IP settings of the particular network connection or you can modify the registry.

By default, the DNS client does not attempt dynamic update of top-level domain (TLD) zones. Any zone that is named with a single-label name is considered a TLD zone, for example, com, edu, blank, or my-company. To configure the DNS client to allow the dynamic update of TLD zones, you can use the **Update Top Level Domain Zones** policy setting or you can modify the registry.

When DNS clients are configured dynamically by a Dynamic Host Configuration Protocol (DHCP) server, it is possible to have a larger list of provided DNS servers. To provide an IP address list of DNS servers to your DHCP clients, enable option code 6 on the configured options types that are provided by your DHCP server. For Windows Server 2003 DHCP servers, you can configure a list of up to 25 DNS servers for each client with this option.

To effectively share the load when multiple DNS servers are provided in a DHCP options-specified list, you can configure a separate DHCP scope that rotates the listed order of DNS and WINS servers that are provided to clients.

**Configuring a DNS Suffix Search List**

For DNS clients, you can configure a DNS domain suffix search list that extends or revises their DNS search capabilities. By adding additional suffixes to the list, you can search for short, unqualified computer names in more than one specified DNS domain. Then, if a DNS query fails, the DNS Client service can use this list to append other name suffix endings to your original name and to repeat DNS queries to the DNS server for these alternate FQDNs.

For computers and servers, the following default DNS search behavior is predetermined and used for completing and resolving short, unqualified names.

When the suffix search list is empty or unspecified, the primary DNS suffix of the computer is appended to short, unqualified names, and a DNS query is used to resolve the resultant FQDN. If this query fails, the computer can try additional queries for alternate FQDNs by appending any connection-specific DNS suffix that is configured for network connections.

If no connection-specific suffixes are configured or if queries for these resultant connection-specific FQDNs fail, the client can then begin to retry queries based on systematic reduction of the primary suffix (also known as devolution).

For example, if the primary suffix is sales.wingtiptoys.com, the devolution process is able to retry queries for the short name by searching for it in the wingtiptoys.com and com domains.

When the suffix search list is not empty and it has at least one DNS suffix specified, attempts to qualify and resolve short DNS names are limited to searching only those FQDNs that are made possible by the specified suffix list. If queries for any FQDNs form as a result of appending and trying each suffix in the list, the query process fails, producing a "Name not found" result.

| Note |
| --- |
| If the domain suffix list is used, clients continue to send additional alternate queries based on different DNS domain names when a query is not answered or resolved. After a name is resolved using an entry in the suffix list, unused list entries are not tried. For this reason, it is most efficient to order the list with the most used domain suffixes first. |

**Configuring Multiple Names**

Computers running Windows XP and servers running Windows Server 2003 are given DNS names by default. Each computer can have its DNS names configured using one of two possible methods:

- A primary DNS domain name, which applies as the default, fully qualified, DNS name for the computer and all its configured network connections.

- A connection-specific, DNS domain name, which can be configured as an alternate DNS domain name that applies only for a single network adapter that is installed and configured on the computer.

Although most computers do not need to support or use more than one name in DNS, support for configuring multiple, connection-specific DNS names is sometimes useful. For example, by using multiple names, a user can specify which network connection to use when connecting to a multihomed computer.

To complete these tasks, perform the following procedure:

Configure DNS settings in Network Connections

# Configure DNS settings in Network Connections

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to configure TCP/IP to use Domain Name System (DNS).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer.

**To configure DNS settings in Network Connections**

1. Open Network Connections.
2. Right-click the network connection that you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. If you want to obtain DNS server addresses from a DHCP server, click **Obtain DNS server address automatically**.
5. If you want to configure DNS server addresses manually, click **Use the following DNS server addresses**, and in **Preferred DNS server** and **Alternate DNS server**, type the Internet Protocol (IP) addresses of the preferred DNS server and alternate DNS server.
6. To configure advanced DNS properties, click **Advanced**, click the **DNS** tab, and then do one or more of the following:
   - To configure an additional DNS server IP address:

     a. Under **DNS server addresses**, in order of use, click **Add**.

     b. In **DNS server**, type the IP address of the DNS server, and then click **Add**.

   - To resolve an unqualified name by appending the primary DNS suffix and the DNS suffix of each connection (if configured), click **Append primary and connection specific DNS suffixes**. If you also want to search the parent suffixes of the primary DNS suffix up to the second-level domain, select the **Append parent suffixes of the primary DNS suffix** check box.
   - To resolve an unqualified name by appending the suffixes from a list of configured suffixes, click **Append these DNS suffixes (in order)**, and then click **Add** to add suffixes to the list.
   - To use a DNS dynamic update to register the IP addresses of this connection and the primary domain name of the computer, select the **Register this connection's addresses in DNS** check box. This option is enabled by default. The primary domain name of the computer is the primary DNS suffix appended to the computer name, and it can be viewed as the full computer name on the **Computer Name** tab (which is available in **System** in **Control Panel**).
   - To use a DNS dynamic update to register the IP addresses and the connection-specific domain name of this connection, select the **Use this connection's DNS suffix in DNS registration** check box. This option is disabled by default. The connection-specific domain name of this connection is the DNS suffix for this connection appended to the computer name.
   - To completely disable DNS dynamic update for all names on the computer, clear the **Register this connection's addresses in DNS** and **Use this connection's DNS suffix in DNS registration** check boxes for all connections in Network Connections.

**Note**

To open Network Connections, click **Start**, point to **Control Panel**, and then click **Network Connections**.

# Managing the DNS Client Resolver Cache

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **ipconfig** command to troubleshoot Domain Name System (DNS) problems or to verify DNS settings. To complete this task for troubleshooting or verification, perform the following procedures:

- Preload the DNS client resolver cache
- View a DNS client resolver cache
- Flush and reset a client resolver cache

# Preload the DNS client resolver cache

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Entries that you add with this procedure are always answered first from the local resolver cache. They are not sent to the Domain Name System (DNS) query when queries are made locally to resolve these names to host address (A) resource records.

Every line in the Hosts file contains an Internet Protocol (IP) address, followed by one or more host names. For example, you can add a line, such as the following line, with an IP address (10.0.0.1) that maps to more than one DNS host name:

```
10.0.0.1  host-a  host-a.example.microsoft.com  host-b.example2.microsoft.com
```

Likewise, a single DNS host name can correspond to more than one IP address if each of the addresses is mapped and used in separate lines. For example, you can add lines for the following multihomed or multiaddressable DNS host computer:

```
10.0.0.1  host-a.example.microsoft.com
10.0.0.2  host-a.example.microsoft.com
10.0.0.3  host-a.example.microsoft.com
```

When multiple names or IP addresses are used in the Hosts file, the DNS Client service must be running for all entries to be returned or used in answering queries. If the DNS Client service is not running, only the first entry in the file is used to resolve the query.

**To preload the DNS client resolver cache**

1. At a command prompt, type the following command, and then press ENTER:

   **notepad %systemroot%\system32\drivers\etc\hosts**

2. Using the default entry in the file (a mapping for the local host to the loopback IP address, 127.0.0.1), add additional host name-to-address mappings on separate lines to be preloaded into the resolver cache of the client. For example, you might add:

   10.0.0.1  host-a  host-a.example.microsoft.com

3. On the **File** menu, click **Save**, and then **Exit**.

4. As an option, you can verify that your changes have been updated in the resolver cache by viewing its contents.

# View a DNS client resolver cache

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **ipconfig /displaydns** command to view the contents of the Domain Name System (DNS) client resolver cache, which includes entries that are preloaded from the local Hosts file, as well as any recently obtained resource records for name queries that were resolved by the system. This information is used by the DNS Client service to quickly resolve frequently queried names before it queries its configured DNS servers.

When you use the **ipconfig /displaydns** command to display current resolver cache contents, the resultant output generally includes the local host and loopback Internet Protocol (IP) address (127.0.0.1) mappings. This is because these mappings typically exist in the default (unmodified) contents of the local Hosts file.

**To view a DNS client resolver cache**

- At a command prompt, type the following command, and then press ENTER:

  **ipconfig /displaydns**

# Flush and reset a client resolver cache

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **ipconfig /flushdns** command to flush and reset the contents of the Domain Name System (DNS) client resolver cache. During DNS troubleshooting, if necessary, you can use this procedure to discard negative cache entries from the cache, as well as any other dynamically added entries.

Resetting the cache does not eliminate entries that are preloaded from the local Hosts file. To eliminate those entries from the cache, remove them from this file.

**To flush and reset a client resolver cache**

- At a command prompt, type the following command, and then press ENTER:

  **ipconfig /flushdns**

# Renewing DNS Client Registration

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the **ipconfig /registerdns** command to initiate dynamic registration manually for the Domain Name System (DNS) names and Internet Protocol (IP) addresses that are configured at a computer. This option can assist in troubleshooting a failed DNS name registration or in resolving a dynamic update problem between a client and the DNS server without restarting the client.

By default, the **ipconfig /registerdns** command refreshes all Dynamic Host Configuration Protocol (DHCP) address leases and registers all related DNS names that are configured and used by the client computer.

**To renew DNS client registration**

- At a command prompt, type the following command, and then press ENTER:

  **ipconfig /registerdns**

# Managing Domain Name System Zones

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following tasks for managing Domain Name System (DNS) zones are described in this objective:

Adding and Removing a Zone

Start or pause a zone

Start a zone transfer at a secondary server

Modifying Zone Properties

Configuring Dynamic Updates

Delegating a Zone

Using Stub Zones for DNS Operations

Using WINS Lookup in DNS Zones

# Adding and Removing a Zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A zone starts as a storage database for a single Domain Name System (DNS) domain name. If other domains are added below the domain that is used to create the zone, these domains can either be part of the same zone or belong to another zone. After a subdomain is added, it can then either be:

- Managed and included as part of the original zone records.
- Delegated away to another zone that is created to support the subdomain.

You can use this task after you determine that you need to add or remove a DNS zone from your environment. For more information about planning DNS zones, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

**Task requirements**

To begin this task, perform the following requirement:

- Install Dnscmd.

To complete this task, perform one of the following procedures:

- Delete a DNS zone
- Add a new zone

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Delete a DNS zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Although it can also be used to delete a primary zone, the following procedure is most often used to delete a secondary copy of a zone. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

Deleting a standard primary zone is usually unnecessary, unless you are redesigning your Domain Name System (DNS) namespace and the zone is no longer needed or used. In most cases, you can change the zone type if you only want to modify the zone.

| **Caution** |
| --- |
| Deleting an Active Directory–integrated zone effectively deletes the zone and eliminates its use at all other DNS servers that use the same directory store of zone data. |

### Administrative credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

#### Deleting a DNS zone

- Using the Windows interface
- Using the command line

### To delete a DNS zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.

   Where?

   - DNS/*applicable DNS server*/Forward Lookup Zones (or Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Delete**.
4. When you are asked to confirm that you want to delete the zone, click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To delete a DNS zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneDelete** *ZoneName* [**/DsDel**] [**/f**]

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneDelete** | Required. Specifies the command to delete the zone that is specified by *ZoneName*. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone that you are deleting. |

| | |
|---|---|
| **/DsDel** | Deletes the zone from Active Directory. |
| **/f** | Performs the command without asking for confirmation. If you omit this parameter, you are prompted to confirm the deletion of the resource record. |

# Add a new zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to create a new primary, secondary, stub, or reverse lookup zone. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Adding a new zone**

- Using the Windows interface
- Using a command line

### To add a new zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click a DNS server, and then click **New Zone** to start the New Zone Wizard.
3. Follow the instructions in the wizard to create a new primary, secondary, stub, or reverse lookup zone.

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To add a new zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneAdd** *ZoneName* {**/Primary**|**/DsPrimary**|**/Secondary**|**/Stub**|**/DsStub**} [**/file** *FileName*] [**/load**] [**/a** *AdminEmail*] [**/DP** *FQDN*]

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneAdd** | Required. Adds a zone. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the in-addr.arpa domain for the zone, for example, **20.1.168.192.in-addr.arpa.** |
| **/Primary**\|**/DsPrimary** | Required. Specifies the type of zone. To specify an Active Directory-integrated zone, type **/DsPrimary**. |
| **/file** | Required for **/Primary**. Specifies a file for the new zone. This parameter is invalid for the **/DsPrimary** zone type. |
| *FileName* | Required for **/Primary**. Specifies the name of the zone file. This parameter is invalid for the **/DsPrimary** zone type. |
| **/load** | Loads an existing file for the zone. If this parameter is not specified, default zone records are created automatically. This parameter does not apply to **/DsPrimary**. |
| **/a** | Adds an administrator e-mail address for the zone. |

| | |
|---|---|
| *AdminEmail* | Specifies the administrator e-mail name for the zone. |
| **/DP** | Adds the zone to an application directory partition. You may also use one of the following:<br>• **/DP /domain** for a domain directory partition (replicates to all DNS servers in the domain).<br>• **/DP /forest** for a forest directory partition (replicates to all DNS server in the forest).<br>• **/DP /legacy** for a legacy directory partition (replicates to all domain controllers in the domain). This setting supports domains using legacy Windows 2000 Server domain controllers. |
| *FQDN* | Specifies the FQDN of the directory partition. |

# Start or pause a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to start or pause a zone. By default, zones are started when they are created or loaded at the server. Only zones that have previously been paused need to be restarted.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Starting or pausing a zone**

- Using the Windows interface
- Using the command line

### To start or pause a zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.

   Where?

   - DNS/*applicable DNS server*/Forward Lookup Zones (or Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, click **Start** or **Pause**, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To start or pause a zone using the command line

1. Open a command prompt. To start a zone, type the following command, and then press ENTER:

   **dnscmd** *ServerName* **/ZoneResume** *ZoneName*

2. To pause a zone, type the following command, and then press ENTER:

   **dnscmd** *ServerName* **/ZonePause** *ZoneName*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneResume** | Required. Resumes the hosting of the zone by the DNS server. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone resuming operation. |

# Start a zone transfer at a secondary server

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure checks to see if the start of authority (SOA) resource record in the secondary zone is the most recent version of the SOA resource record in the primary zone. If the SOA resource records are synchronized, there is no zone transfer. If the SOA resource records are not synchronized, there is a zone transfer.

By default, the Domain Name System (DNS) server only allows a zone transfer to authoritative DNS servers that are listed in the name server (NS) resource records for the zone.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Starting a zone transfer at a secondary server**

- Using the Windows interface
- Using the command line

### To start a zone transfer at a secondary server using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Transfer from master**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To start a zone transfer at a secondary server using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneRefresh** *ZoneName*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneRefresh** | Required. Updates the secondary zone. |
| *ZoneName* | Required. Specifies the name of the secondary zone to update. |

# Modifying Zone Properties

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You might decide to modify certain Domain Name System (DNS) zone properties based on a design or performance evaluation of your network. These settings affect zone transfers and zone types. For more information about planning DNS zones, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

**Task requirements**

To begin this task, perform the following requirements:

- Install Dnscmd.

To complete this task, perform one of the following procedures:

1. Change the DNS zone type
2. Change a DNS zone file name
3. Change the zone replication scope
4. Modify the SOA record for a zone
5. Modify DNS zone transfer settings
6. Specify DNS servers as authoritative for a zone
7. Change the master server for a secondary zone
8. Create a notify list for a zone
9. Adjust the refresh, retry, or expire intervals for a zone

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Change the DNS zone type

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to change the Domain Name System (DNS) zone type. You can select a primary, secondary, or stub zone. When you select the secondary or stub zone type, you must specify the Internet Protocol (IP) address of another DNS server to be used as the source for obtaining updated information for the zone.

If the DNS server computer is operating as a domain controller, the option to change the zone type to **Active Directory-Integrated** is available. When you select this zone type, zone data is stored and replicated as part of the Active Directory database.

Changing a zone from secondary to primary can affect other zone activities, including the management of dynamic updates and zone transfers and the use of DNS notify lists to notify other servers about changes in the zone. Changing a zone from stub to primary (or the reverse) is not recommended because of the purpose of stub zones. For more information about planning DNS zones, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Changing the zone type**

- Using the Windows interface
- Using the command line

### To change the zone type using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, note the current zone type, and then click **Change**.
4. In the **Change Zone Type** dialog box, click a zone type other than the current one, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To change the zone type using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneResetType** *ZoneName Property* [*MasterIPaddress...*] [**/file** *FileName*] {**/OverWrite_Mem**|**/OverWrite_Ds**|**/DirectoryPartition** *FQDN*}

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of zone. |

| | |
|---|---|
| *Property* | Required. One of the following zone types:<br>• **/Primary**<br>Standard primary zone. The *FileName* must be required.<br>• **/DsPrimary**<br>Active Directory-integrated primary zone.<br>• **/Secondary**<br>Secondary zone. You must specify at least one *MasterIPaddress...*<br>• **/Stub**<br>Stub zone. You must specify at least one *MasterIPaddress....*<br>• **/DsStub**<br>Active Directory-integrated stub zone. You must specify at least one *MasterIPaddress....* |
| *MasterIPaddress...* | Required for **/Secondary**, **/Stub** and **/DsStub**. Specifies one or more IP addresses for the master servers of the secondary or stub zone, from which it copies zone data. |
| **/file** | Required for **/Primary**. Specifies a file for the new zone. This parameter is invalid for the **/DsPrimary** zone type. |
| *FileName* | Required for **/Primary**. Specifies the name of the zone file. This parameter is invalid for the **/DsPrimary** zone type. |
| **/OverWrite_Mem**\|**/OverWrite_Ds**\|**/DirectoryPartition***FQDN* | **/OverWrite_Mem** overwrites exisiting DNS data using the data in Active Directory.<br>**/OverWrite_Ds** overwrites Active Directory data with data in DNS.<br>**/DirectoryPartition** stores the new zone in the application directory partition that is specified by *FQDN*, such as DomainDnsZones.corp.sales.wingtiptoys.com. |

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Change a DNS zone file name

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you use the following procedure, the name of the zone file changes, not the name of the zone. You can use Windows Explorer to view or verify the new zone file name.

The zone file name is not used for Active Directory-integrated zones because these zones store zone data in the Active Directory database, not a text file on the DNS server computer.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To change a zone file name**

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.

   Where?

   - DNS/*applicable DNS server*/Forward Lookup Zones (or Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, in the **Zone file name** text box, type the new file name for this zone, and then click **OK**.

**Caution**

- If the zone file name is changed, be sure to update the zone file name on other DNS servers that maintain this zone. Otherwise, subsequent zone transfers and updates might fail. This can occur in the following situations:

  - The zone type is primary on this server.

  - The zone type is secondary on this server, and this server acts as a source or master server for this zone to other DNS servers that host secondary copies of this zone.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Change the zone replication scope

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to change the replication scope for a zone. Only Active Directory-integrated primary and stub forward lookup zones can change their replication scope. Secondary forward lookup zones cannot change their replication scope.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Changing zone replication scope**

- Using the Windows interface
- Using the command line

### To change zone replication scope using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, note the current zone replication type, and then click **Change**.
4. Select a replication scope for the zone.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To change zone replication scope using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneChangeDirectoryPartition** *ZoneName NewPartitionName*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneChangeDirectoryPartition** | Required. Changes a zone's replication scope. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NewPartitionName* | Required. The FQDN of the DNS application directory partition where the zone will be stored. |

# Modify the SOA record for a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to change settings for the start of authority (SOA) resource record for a zone. The settings that are applied for the SOA record affect how zone transfers are made between servers.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Modifying the SOA record for a zone**

- Using the Windows interface
- Using a command line

### To modify the SOA record for a zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. Click the **Start of Authority (SOA)** tab.
4. Modify the properties for the SOA record as needed.

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To modify the SOA record for a zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/RecordAdd** *ZoneName NodeName* [**/Aging**] [*Ttl*] **SOA***PrimSvrAdmin Serial#\ Refresh Retry Expire MinTTL*

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Adds or modifies a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName*, or you can type @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this parameter is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. The default TTL is defined in SOA resource record. |
| **SOA** | Required. Specifies the type of resource record that you are modifying. |

| | |
|---|---|
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *PrimSvr* | Required. Specifies the FQDN name of the server that is the primary source for information about the zone, for example, nameserver.place.sales.wingtiptoys.com.. |
| *Admin* | Required. Specifies the name of the DNS administrator for the zone, for example, postmaster.nameserver.place.sales.wingtiptoys.com.. |
| *Serial#* | Required. Specifies the version information for the zone. |
| *Refresh* | Required. Specifies the refresh interval for the zone. The standard setting is 3600 seconds (one hour). |
| *Retry* | Required. Specifies the retry interval for the zone. The standard setting is 600 seconds (10 minutes). |
| *Expire* | Required. Specifies the expire interval for the zone. The standard setting is 86400 seconds (one day). |
| *MinTTL* | Required. Specifies the minimum TTL value. This is the length of time that is used by other DNS servers to determine how long to cache information for a record in the zone before expiring and discarding it. The standard setting is 3600 seconds (one hour). |

**Note**

To modify any specific SOA record's values using Dnscmd, you must specify all the SOA values (*PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*).

# Modify DNS zone transfer settings

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to change Domain Name System (DNS) zone transfer settings. To improve the security of your DNS infrastructure, zone transfers should be allowed only for either the DNS servers in the name server (NS) resource records for a zone or for specified DNS servers. If you allow any DNS server to perform a zone transfer, you are allowing internal network information to be transferred to any host that can contact your DNS server.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Modifying DNS zone transfer settings**

- Using the Windows interface
- Using the command line

### To modify DNS zone transfer settings using the Windows interface

1. Open the DNS snap-in.
2. Right-click a DNS zone, and then click **Properties**.
3. On the **Zone Transfers** tab, do one of the following:
   - To disable zone transfers, clear the **Allow zone transfers** check box.
   - To allow zone transfers, select the **Allow zone transfers** check box.
4. If you allowed zone transfers, do one of the following:
   - To allow zone transfers to any server, click **To any server**.
   - To allow zone transfers only to the DNS servers that are listed on the **Name Servers** tab, click **Only to servers listed on the Name Servers tab**.
   - To allow zone transfers only to specific DNS servers, click **Only to the following servers**, and then add the Internet Protocol (IP) address of one or more DNS servers.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To modify DNS zone transfer settings using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneResetSecondaries** *ZoneName* {**/NoXfr**|**/NonSecure**|**/SecureNs**|**/SecureList**[*SecondaryIPAddress*...]}

| Value | Description |
| --- | --- |

| | |
|---|---|
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of zone. |
| **/NoXfr** | Disables zone transfers for the zone. |
| **/NonSecure** | Permits zone transfers to any DNS server. |
| **/SecureNs** | Permits zone transfers only to DNS servers that are listed in the zone using NS resource records. |
| **/SecureList** | Permits zone transfers only to DNS servers that are specified by *SecondaryIPAddress*. |
| *SecondaryIPAddress* | Required if **/SecureList** is specified. A list of one or more IP addresses for DNS servers that are permitted to obtain zone transfers. |

# Specify DNS servers as authoritative for a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain Name System (DNS) servers that you specify with the following procedure are added to those server Internet Protocol (IP) addresses that are already present for the existing name server (NS) resource record for the zone. Typically, you might only need to perform this procedure at the primary zone when you add DNS servers to act as secondary servers and also to specify that these servers are known to be authoritative when they answer queries for zone data.

DNS servers automatically add and perform initial configuration of the NS resource record for each new primary zone that is added to the server.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Specifying DNS servers as authoritative for a zone**

- Using the Windows interface
- Using the command line

### To specify DNS servers as authoritative for a zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. Click the **Name Servers** tab.
4. Click **Add**.
5. Specify additional DNS servers by their names and IP addresses, and then click **Add** to add them to the list.

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

| **Note** |
|---|
| To add a name server to the list of authoritative servers for the zone, you must specify both the server's IP address and its DNS name. When you enter a name, click **Resolve** to resolve the name to its IP address before adding it to the list. |

### To specify DNS servers as authoritative for a zone using the command line

1. At a command prompt, type the following command, and then press ENTER:

   **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **NS** {*HostName*|*DomainName*}

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Specifies the command to add a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |

| | |
|---|---|
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is updated or removed manually. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. (The default TTL is defined in the start-of-authority (SOA) resource record). |
| **NS** | Required. Specifies that you are adding a name server (NS) resource record to the zone that is specified in *ZoneName*. |
| *HostName*\|*DomainName* | Required. Specifies the host name or FQDN of the new authoritative server. |

# Change the master server for a secondary zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to specify a new master server for a secondary zone. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Changing the master server for a secondary zone**

- Using the Windows interface
- Using the command line

### To change the master server for a secondary zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable secondary zone, and then click **Properties**.
3. On the **General** tab, in **IP address**, specify the Internet Protocol (IP) address for a new master server, and then click **Add** to update the list.

| **Note** |
| --- |
| To open DNS, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To change the master server for a secondary zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneResetMasters** *ZoneName* [**/Local**] *MasterIPaddress...*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneResetMasters** | Required. Updates the master servers for a secondary zone. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone that you are updating. |
| **/Local** | Specifies the local master list for Active Directory-integrated zones. |
| *MasterIPaddress...* | Required. Specifies the IP addresses of the master servers to be used by the DNS server when updating the specified secondary zones. If you do not specify *ServerIPs*, you are requesting the DNS server to reset the value to an empty list. The request may be denied because a zone must always have at least one master server. *MasterIPaddress...* is required to clear the local master list for a zone. |

# Create a notify list for a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to create or change a notify list for a zone. Changes to the notify list properties are available only on primary zones. For secondary zones, these properties are read only.

By default, the DNS server allows a zone transfer only to authoritative DNS servers that are listed in the name server (NS) resource records for the zone.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To create or change a notify list for a zone**

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.

   Where?

   - DNS/*applicable DNS server*/Forward Lookup Zones (or Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Properties**.
4. Click the **Zone Transfers** tab.
5. Click **Notify**.
6. Verify that the **Automatically notify** check box is selected.
7. Select the method to be used for creating a list for notifying other DNS servers when changes to the zone occur. Your options are as follows:

   - Use the default, **Servers listed on the Name Servers tab**, to permit only those servers that appear by Internet Protocol (IP) address on the **Name Servers** tab to be included in the notify list.

   - Select **The following servers** if you want to specify a different notify list to be used instead.

8. If you selected **The following servers** in the previous step, add or remove server IP addresses to form the notify list as needed:

   - To add a server to the notify list, type its IP address in the **IP address** box, and then click **Add**.

   - To remove a server from the notify list, click the server IP address in the list box, and then click **Remove**.

| **Note** |
| --- |
| To open DNS, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Adjust the refresh, retry, or expire intervals for a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to change the following intervals for a Domain Name System (DNS) zone:

- **Refresh interval**. Used to determine how often other DNS servers that load and host the zone must attempt to renew the zone.
- **Retry interval**. Used to determine how often other DNS servers that load and host the zone are to retry a request for update of the zone each time that the refresh interval occurs.
- **Expire interval**. Used by other DNS servers that are configured to load and host the zone to determine when zone data expires if it is not renewed.

The default values for each interval are as follows:

- **Refresh interval**: 15 minutes.
- **Retry interval**: 10 minutes.
- **Expire interval**: one day.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Adjusting the refresh, retry, or expire interval for a zone**

- Using the Windows interface
- Using the command line

**To adjust the refresh, retry, or expire interval for a zone using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, verify that the zone type is either **Primary** or **Active Directory-Integrated**.
4. Click the **Start of Authority (SOA)** tab.
5. In **Refresh interval**, **Retry interval**, or **Expires after**, click a time period in minutes, hours, or days, and type a number in the text box.
6. Click **OK** to save the adjusted interval.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To adjust the refresh, retry, or expire interval for a zone using the command line**

- At a command prompt, type the following command, and then press ENTER:

**dnscmd** *ServerName***/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **SOA** *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Adds or modifies a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName*, or you can type @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this parameter is not used, the resource record remains in the DNS database unless it is updated or removed manually. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. The default TTL is defined in the start-of-authority (SOA) resource record. |
| **SOA** | Required. Specifies the type of resource record that you are modifying. |
| *PrimSvr* | Required. Specifies the FQDN name of the server that is the primary source for information about the zone, for example, nameserver.place.sales.wingtiptoys.com.. |
| *Admin* | Required. Specifies the name of the DNS administrator for the zone, for example, postmaster.nameserver.place.sales.wingtiptoys.com. |
| *Serial#* | Required. Specifies the version information for the zone. |
| *Refresh* | Required. Specifies the refresh interval for the zone. The standard setting is 3600 seconds (one hour). |
| *Retry* | Required. Specifies the retry interval for the zone. The standard setting is 600 seconds (10 minutes). |
| *Expire* | Required. Specifies the expire interval for the zone. The standard setting is 86400 seconds (one day). |
| *MinTTL* | Required. Specifies the minimum TTL value. This is the length of time that is used by other DNS servers to determine how long to cache information for a record in the zone before expiring and discarding it. The standard setting is 3600 seconds (one hour). |

**Note**

To modify any specific SOA resource record's values using Dnscmd, you must specify all the SOA values (*PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*).

# Configuring Dynamic Updates

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Dynamic update enables Domain Name System (DNS) client computers to register and dynamically update their resource records with a DNS server whenever changes occur. This reduces the need for manual administration of zone records, especially for clients that frequently move or change locations and use Dynamic Host Configuration Protocol (DHCP) to obtain an Internet Protocol (IP) address.

The DNS Server service allows dynamic update to be enabled or disabled on a per-zone basis at each server that is configured to load either a standard primary or directory-integrated zone. By default, the DNS Client service dynamically updates host address (A) resource records in DNS when it is configured for TCP/IP.

Secure dynamic update is available only for zones that are integrated into Active Directory. After you directory-integrate a zone, access control list (ACL) editing features are available in the DNS snap-in. You can use these features to add or remove users or groups from the ACL for a specified zone or resource record.

For more information about planning DNS zones and dynamic updates, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

You can use this task to enable or disable dynamic updates or to allow only secure dynamic updates.

**Task requirements**

To begin this task, perform the following requirements:

- Install Dnscmd.

To complete this task, perform one of the following procedures:

- Enable dynamic updates
- Enable secure dynamic updates

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Enable dynamic updates

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to enable dynamic updates for a zone. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Enabling dynamic updates**

- Using the Windows interface
- Using the command line

### To enable dynamic updates using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, verify that the zone **Type** is either **Primary** or **Active Directory-Integrated**.
4. In **Dynamic updates**, click **Nonsecure and secure**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To enable dynamic updates using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config** {*ZoneName*|**..AllZones**} **/AllowUpdate** {**1**|**0**}

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName*\|**..AllZones** | Required. Specifies the fully qualified domain name (FQDN) of the zone. To configure all zones that are hosted on the specified DNS server to allow dynamic updates, type **..AllZones**. |
| **1**\|**0** | Configures dynamic update. To allow dynamic updates, type a value of **1**. To not allow dynamic updates, type a value of **0**. |

# Enable secure dynamic updates

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to allow only secure dynamic updates for a zone. Secure dynamic update is supported only for Active Directory–integrated zones. If the zone type is configured differently, you must change the zone type and directory-integrate the zone before securing it for Domain Name System (DNS) dynamic updates.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Enabling secure dynamic updates**

- Using the Windows interface
- Using the command line

### To enable secure dynamic updates using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, verify that the zone type is **Active Directory-Integrated**.
4. In **Dynamic updates**, click **Secure only**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To enable secure dynamic updates using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config** {*ZoneName*|**..AllZones**} **/AllowUpdate 2**

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName*\|**..AllZones** | Required. Specifies the fully qualified domain name (FQDN) of the zone. To configure all zones that are hosted on the specified DNS server to allow dynamic updates, type **..AllZones**. |
| **2** | Required. Configures the server to allow secure dynamic updates. If you exclude the **2**, the zone is set to perform standard dynamic updates only. |

# Delegating a Zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain Name System (DNS) provides the option of dividing up the namespace into one or more zones, which can then be stored, distributed, and replicated to other DNS servers. When deciding whether to divide your DNS namespace to make additional zones, consider the following reasons to use additional zones:

- You want to delegate management of part of your DNS namespace to another location or department in your organization.
- You want to divide one large zone into smaller zones for distributing traffic loads among multiple servers, improving DNS name resolution performance, or creating a more fault-tolerant DNS environment.
- You want to extend the namespace by adding numerous subdomains at once, for example, to accommodate the opening of a new branch or site.

If, for any of these reasons, your network can benefit from delegating zones, it may make sense to restructure your namespace by adding additional zones. When choosing how to structure zones, use a plan that reflects the structure of your organization.

When you delegate zones within your namespace, be aware that for each new zone you create, you will need delegation records in other zones that point to the authoritative DNS servers for the new zone. This is necessary both to transfer authority and to provide correct referral to other DNS servers and clients of the new servers that are being made authoritative for the new zone.

When a standard primary zone is first created, it is stored as a text file that contains all resource record information on a single DNS server. This server acts as the primary master for the zone. Zone information can be replicated to other DNS servers to improve fault tolerance and server performance.

When you structure your zones, there are several good reasons to use additional DNS servers for zone replication:

- Added DNS servers provide zone redundancy, enabling DNS names in the zone to be resolved for clients if a primary server for the zone stops responding.
- Added DNS servers can be placed so as to reduce DNS network traffic. For example, adding a DNS server to the opposing side of a low-speed, wide area network (WAN) link can be useful in managing and reducing network traffic.
- Additional secondary servers can be used to reduce loads on a primary server for a zone.

For more information about planning DNS zones, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

**Task requirements**

To begin this task, perform the following requirements:

- Install Dnscmd.
- Install Nslookup.

To complete this task, perform one of the following procedures:

- Create a new zone delegation
- Verify a zone delegation

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Create a new zone delegation

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to create a new zone delegation. All domains (or subdomains) that appear as part of the applicable zone delegation must be created in the current zone before you perform delegation as described in this procedure. As necessary, use the DNS snap-in to first add domains to the zone before you perform this procedure. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Creating a new zone delegation**

- Using the Windows interface
- Using the command line

### To create a new zone delegation using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable subdomain, and then click **New Delegation**.
3. Follow the instructions in the New Delegation Wizard to finish creating the new delegated domain.

| **Note** |
|---|
| To open DNS, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To create a new zone delegation using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/RecordAdd** *ZoneName NodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **NS** {*HostName*|*FQDN*}

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the start-of-authority (SOA) record is added. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is updated or removed manually. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. (The default TTL is defined in the SOA resource record). |
| **NS** | Required. Specifies that you are adding a name server (NS) resource record to the zone that is specified in *ZoneName*. |

| | |
|---|---|
| *HostName\|FQDN* | Required. Specifies the host name or FQDN of the new authoritative server. |

# Verify a zone delegation

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administrative credentials**

You do not need administrative credentials to perform this task. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

**To verify a zone delegation**

1. At a command prompt, type the following command, and then press ENTER:

   **nslookup** *RootServerIpAddress*

2. Type the following command, and then press ENTER:

   **nslookup**

3. At the next prompt, type the following command, and then press ENTER:

   **set norecurse**

4. At the next prompt, type the following command, and then press ENTER:

   **set q=NS**

5. Type the fully qualified domain name (FQDN) for the failed name.

   Use the trailing period (**.**) when you type the name. If zone delegations are set correctly, a list of name server (NS) resource records for delegated servers is returned in the response.

6. If the NS query response contains no names or Internet Protocol (IP) addresses for delegated servers, type **q=ns**, and then query again using the FQDN for the parent zone of the failed name.

   For example, if the failed name that you used in the previous step was sales.wingtiptoys.com, query for wingtiptoys.com.

7. If the response contains NS resource records, but no host address (A) resource records, type **set recurse**, and then query individually for any of the A resource records of the servers that are listed in the NS resource records.

   If, for each NS resource record that you encounter in a zone, you do not find at least one valid IP address in an A resource record, you have a broken delegation.

8. Either fix the broken delegation or retry the delegation test that is described in the previous step and use a different IP address.

   If more than one A resource record or IP address is found, use it to repeat the delegation test described in the previous step. To fix a delegation, add or update an A resource record in the parent zone with a valid IP address for a correct DNS server for the delegated zone.

| Value | Description |
|---|---|
| *RootServerIpAddress* | The IP address of a valid root server for your network. |
| **set norecursion** | Instructs the root server to not perform recursion on your query. |
| **set q=NS** | Sends the query for NS resource records to the root server. |

# Using Stub Zones for DNS Operations

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use stub zones to:

- **Keep delegated zone information current.** By updating a stub zone for one of its child zones regularly, the Domain Name System (DNS) server that hosts both the parent zone and the stub zone maintains a current list of authoritative DNS servers for the child zone.

- **Improve name resolution.** Stub zones enable a DNS server to perform recursion by using the stub zone's list of name servers, without needing to query the Internet or the internal root server for the DNS namespace.

- **Simplify DNS administration.** By using stub zones throughout your DNS infrastructure, you can distribute a list of the authoritative DNS servers for a zone without using secondary zones. However, stub zones do not serve the same purpose as secondary zones, and they are not a valid alternative to secondary zones with regard to redundancy and load sharing.

When a DNS server loads a stub zone, it queries the master servers, which can be in different locations, for the necessary resource records of the authoritative servers for the zone. The list of master servers may contain a single server or multiple servers, and the list can be changed anytime.

For more information about planning DNS zones, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

**Task requirements**

To begin this task, perform the following requirements:

- Install Dnscmd.

To complete this task, perform one of the following procedures:

- Reload or transfer stub zones

- Configure a stub zone to use local master servers

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Reload or transfer stub zones

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to reload or transfer stub zones. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Reloading or transferring stub zones**

- Using the Windows interface
- Using the command line

**To reload or transfer stub zones using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable stub zone, and do one of the following:
   - To reload the stub zone from storage, click **Reload**.
   - To have the DNS server determine if the serial number in the stub zone's start-of-authority (SOA) resource record has expired and then perform a zone transfer from the stub zone's master server, click **Transfer from Master**.
   - To perform a zone transfer from the stub zone's master server regardless of the serial number in the stub zone's SOA resource record, click **Reload from Master**.

| **Note** |
| --- |
| To open DNS, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To reload or transfer stub zones using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* {**/ZoneReload**|**/ZoneUpdateFromDs**|**/ZoneRefresh**} *ZoneName*

| **Value** | **Description** |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneReload** | Reloads the stub zone. |
| **/ZoneUpdateFromDs** | Reloads the stub zone from Active Directory. |
| **/ZoneRefresh** | Refreshes the stub zone. The DNS server determines if the serial number in the stub zone's SOA resource record has expired. If the serial number has expired, the DNS server performs a zone transfer from the stub zone's master server. |
| *ZoneName* | Required. Specifies the name of the stub zone that you want to reload or refresh. |

| **Note** |
| --- |

There is no **dnscmd** command to perform a zone transfer regardless of the SOA resource record's expiration date. To perform this operation, use the Windows interface procedure.

# Configure a stub zone to use local master servers

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to specify local master servers that you want the Domain Name System (DNS) server to use when loading and updating the stub zone.

When modifications to the master servers list are made and applied on a domain controller hosting the stub zone, the list of master servers for the stub zone is updated in Active Directory. If the local list of master servers is cleared at a later date, the master servers list from Active Directory is applied and the local list of master servers is deleted.

The DNS server keeps the master servers list from Active Directory stored in memory.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the DnsAdmins or the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using the **Run as** command to perform this procedure.

**Configuring a stub zone to use local master servers**

- Using the Windows interface
- Using the command line

### To configure a stub zone to use local master servers using the Windows interface

1. Open DNS.
2. In the console tree, right-click the stub zone, and then click **Properties**.
3. On the **General** tab, under **IP address**, modify the list to display the Internet Protocol (IP) addresses of the local master servers that you want the DNS server to use when loading and updating the stub zone.

   Ensure that the IP addresses of the local master servers are for only those authoritative DNS servers that should be queried to update the records of the stub zone on this server.
4. Select the **Use the list above as a local list of masters** check box, and then click **OK**.

| **Note** |
| --- |
| To open DNS, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To configure a stub zone to use local master servers using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneResetMasters** *ZoneName* [**/Local**] [*MasterIPaddress...*]

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| **/Local** | Configures the local master list for Active Directory–integrated zones. |

| | |
|---|---|
| *MasterIPaddress...* | List of one or more IP addresses of master servers for this zone. Master servers may include the server hosting the primary zone or servers hosting other secondary copies for the zone. To clear the local list of masters, type the command without entering any IP addresses. Ensure that the IP addresses of the local master servers are for only those authoritative DNS servers that should be queried to update the records of the stub zone on this server. |

# Using WINS Lookup in DNS Zones

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The DNS Server service provides the ability to use Windows Internet Name Service (WINS) servers to look up names that are not found in the Domain Name System (DNS) domain namespace by checking the network basic input/output system (NetBIOS) namespace that is managed by WINS.

For WINS lookup integration, two special resource record types — the WINS and WINS-R resource records — are enabled and added to a zone. When the WINS resource record is used, DNS queries that fail to find a matched host address (A) resource record in the zone are forwarded to WINS servers that are configured in the WINS resource record. For reverse lookup zones, the WINS-R resource record can be enabled and used to provide a similar benefit for further resolving a reverse query that is not answerable in the reverse in-addr.arpa domain.

For example, you can use WINS lookup when you are using a mixed-mode client environment consisting of UNIX clients that use only DNS name resolution and earlier-version Microsoft clients that require NetBIOS naming. In these environments, WINS lookup provides a method for permitting UNIX DNS clients to locate your WINS clients by extending DNS host name resolution into the WINS-managed NetBIOS namespace.

The WINS lookup integration feature is supported only by Windows DNS servers. If you use a mixture of Windows and other DNS servers to host a zone, you should select the **Do not replicate this record** check box option for any primary zones when you use the WINS lookup record. This prevents the WINS lookup record from being included in zone transfers to other DNS servers that do not support or recognize this record. If you do not enable the WINS lookup record to be used only on the local server, it can cause data errors or failed zone transfers at servers running other DNS server implementations that replicate the zone:

- How the caching Time to Live (TTL) and lookup time-out values are configured for use with the WINS and WINS-R records
- The format of the WINS and WINS-R resource records as they are used in zone files that are created by the DNS Server service

**WINS Lookup Interoperability**

Typically, WINS lookup provides the best and most predictable results if only Windows DNS servers are used, and it is only available directly for use at Windows DNS servers. There are ways, however, that you can use and benefit from WINS lookup as an interoperable solution when other DNS servers are deployed.

For example, consider adding a Windows DNS server that hosts a new WINS lookup-enabled zone. When you create and name the zone, use a subdomain that is added to your existing DNS namespace that is used just for WINS-specific referrals that are added to your DNS domain namespace.

For instance, in sales.wingtiptoys.com, call the zone wins.sales.wingtiptoys.com when you create it. You can then use this new WINS referral zone as the root zone for any of your WINS-aware computers that have names that are not found in your other traditional DNS zones.

To use the WINS referral zone, you must specify its domain name (wins.sales.wingtiptoys.com) in a DNS suffix search list for your clients. The suffix list is configurable as part of the TCP/IP properties for a client connection, and it can be updated either manually, by using Dynamic Host Configuration Protocol (DHCP) or by using Group Policy. As long as the name of the WINS referral zone is included in the domain suffix list, any DNS names that are not resolved in traditional zones can be resolved by using the WINS referral subdomain.

Under normal conditions, this should result in recursion from your other DNS servers to the Windows DNS servers that host the WINS-enabled zone. If the queried host names match NetBIOS computer names that are found in the WINS database, the names are resolved to the Internet Protocol (IP) addresses that are mapped in WINS data there.

In our example, the WINS-enabled zone is used only for WINS lookup; therefore, no additional resource records need to be added to it. In general, WINS records can be added to any forward lookup zone.

By using a specific subdomain just for WINS lookup and specifying a static DNS suffix list to be used in resolving and searching for names, you can prevent unusual situations in which DNS queries for different fully qualified domain names (FQDNs) resolve to the same WINS client name and IP address. This might easily occur if you add and configure many zones at each level of your namespace and enable each of them to use WINS lookup integration.

For example, suppose you have two zones, both configured to use WINS lookup. The zones are rooted and originate at the following DNS domain names:

sales1.wingtiptoys.com.

sales2.wingtiptoys.com.

With this configuration, a WINS client named HOST-A can be unintentionally resolved by using either of the following FQDNs:

host-a.sales1.wingtiptoys.com.

host-a.sales2.wingtiptoys.com.

**Advanced Parameters for WINS Lookups**
The two following advanced timing parameters are used with the WINS and WINS-R records:

- The Cache timeout value, which indicates to a DNS server how long it should cache any of the information that is returned in a WINS lookup. By default, this value is set to 15 minutes.

- The Lookup timeout value, which specifies how long to wait before timing out and expiring a WINS lookup that is performed by the DNS Server service. By default, this value is set to two seconds.

You can configure these parameters by using the **Advanced** button in the zone properties dialog box when you configure the zone. This button appears on either the WINS or WINS-R tab, depending on whether the zone that you are configuring is being used for forward lookup or reverse lookup.

If you are using either the WINS or WINS-R resource record, be aware that the minimum TTL that is set in the start-of-authority (SOA) record for the zone is not the default TTL that is used with these records. Instead, when either an IP address or a host name is resolved with WINS lookup, the information is cached on the DNS server for the amount of time that is configured for the WINS cache time-out value. If this address is then ever forwarded to another DNS server, the WINS cache time-out value TTL is what is sent. If your WINS data rarely changes, you can increase the default TTL of 15 minutes.

### Notes

- If you have a zone that is configured for WINS lookup, all DNS servers that are authoritative for that zone need to be capable of WINS lookup or you will have intermittent behavior.

- Because you can specify that the WINS and WINS-R resource records not be replicated to other DNS servers, you can selectively enable and configure WINS lookup at each of your secondary servers for zones where this feature is used. This is not a standard practice for other types of resource records, which are only to be configured at the primary server for the zone.

For more information about planning DNS zones, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

### Task requirements

To begin this task, perform the following requirements:

- Install Dnscmd.

To complete this task, perform the following procedures:

1. Allow DNS to use WINS resolution

2. Verify that WINS is answering a DNS query

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Allow DNS to use WINS resolution

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to enable Domain Name System (DNS) to use Windows Internet Name Service (WINS) name resolution. The specified WINS servers that are configured in this procedure are used for final referral of names that are not found in the applicable zone.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To allow DNS to use WINS resolution**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. Do one of the following:
   - If the applicable zone is a forward lookup zone, on the **WINS** tab, select the **Use WINS forward lookup** check box. In **IP address**, type the Internet Protocol (IP) address of a WINS server to be used for resolution of names not found in DNS, and then click **Add**.
   - If the applicable zone is a reverse lookup zone, on the **WINS-R** tab, select the **Use WINS-R lookup** check box. In **Domain to append to returned name**, type a name.
4. Select the **Do not replicate this record** check box for this WINS record, if applicable.

   If you are replicating this zone between DNS servers that do not recognize the WINS or WINS-R resource records, select this check box. This prevents these records from being replicated to these other servers during zone transfers. If this zone will be used in performing zone transfers to BIND servers, this is a critical option because Berkeley Internet Name Domain (BIND) does not recognize WINS records.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Verify that WINS is answering a DNS query

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to verify that Windows Internet Name Service (WINS) is resolving a Domain Name System (DNS) query.

Normally, when a DNS server answers a query from its authoritative zone data, it uses the set minimum or default Time to Live (TTL) for the zone or the record-specific TTL value (if one is configured). In so doing, TTLs are decreased in answers that the server returns if they are based on nonauthoritative data, such as a cached record at the server.

WINS lookups present an exceptional case, in which an answer that is received back from a WINS server is cached by the DNS server but is also considered to be authoritative data. In this case, the WINS sourced data is returned to clients as authoritative, but it ages while it is in the DNS server names cache, which causes the TTL that is used by the server to decrease over time.

**Administrative credentials**

You do not need administrative credentials to perform this procedure. Therefore, as a security best practice, consider performing this procedure as a user without administrative credentials.

**To verify that WINS is answering a DNS query**

1. At a command prompt, type the following command, and then press ENTER:

   **nslookup**

2. At the **nslookup ("")** prompt, type the following command, and then press ENTER:

   **set debug**

3. Next, either type:

   **set querytype=a**

   if you are testing for a WINS forward lookup, or:

   **set querytype=ptr**

   if you are testing for a WINS-R reverse lookup, and then press ENTER.

   Respectively, these two commands can be used to set the query type to filter either by host address (A) or pointer (PTR) resource records as appropriate for researching either a forward lookup or a reverse lookup.

4. Based on whether you are verifying possible WINS sourcing for either a forward lookup or a reverse lookup, type the appropriate fully qualified domain name (FQDN).

   For example, if the forward lookup that you are tracing is for a domain name host-a.sales.wingtiptoys.com, type:

   **host-a.sales.wingtiptoys.com.**

   If the reverse lookup that you are tracing is for an Internet Protocol (IP) address 10.0.0.1, type:

   **1.0.0.10.in-addr.arpa.**

5. In the response, note whether the server answered authoritatively or nonauthoritatively, and note the TTL value.

6. If the server answered authoritatively, repeat the same query that you performed in step 4.

7. In the response, note whether the TTL value decreased with the second query answer or if it remained consistent with the TTL value that was specified in the first query answer.

   If the TTL value decreased for an authoritatively answered query, the source of the query answer is a WINS server.

8. To leave debug mode and return to the command prompt, type **exit**, and then press ENTER.

| Value | Description |
|---|---|

| | |
|---|---|
| **set debug** | Enables the **nslookup** command to operate in debug mode, providing extended information in the command output.<br>This mode is required to view query response information about whether the source for a query answer is:<br>● Authoritative (from a DNS zone or a WINS server database)<br>● Nonauthoritative (cached data from previous queries made by the DNS server or loaded from root hints) |
| **set querytype** | Changes the type of information query. More information about types can be found in Request for Comments (RFC) 1035. |

# Managing DNS Resource Records

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following tasks for managing Domain Name System (DNS) resource records are described in this objective:

Adding, Changing, and Deleting Resource Records

Disable NS resource record registration

Allow NS resource record creation for domain controllers

Restrict the DNS resource records that are updated by Netlogon

# Adding, Changing, and Deleting Resource Records

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After you create a zone, additional resource records must be added to it. The most common resource records include the following:

- **Host address (A).** Maps a Domain Name System (DNS) domain name to an Internet Protocol (IP) address that is used by a computer.
- **Alias canonical (CNAME).** Maps an alias DNS domain name to another primary name or canonical name.
- **Mail Exchanger (MX).** Maps a DNS domain name to the name of a computer that exchanges or forwards mail.
- **Pointer (PTR).** Maps a reverse DNS domain name based on the IP address of a computer that points to the forward DNS domain name of that computer.
- **Service (SRV).** Maps a DNS domain name to a specified list of DNS host computers that offer a specific type of service, such as Active Directory domain controllers.
- Other resource records, as needed.

**Host A Resource Records**

Host A resource records are used in a zone to associate DNS domain names of computers (or "hosts") to their IP addresses. These resource records can be added to a zone in several ways:

- You can create an A resource record for a static TCP/IP client computer manually by using the DNS snap-in.
- Windows clients and servers use the DHCP Client service to dynamically register and update their own A resource records in DNS when an IP configuration change occurs.
- Dynamic Host Configuration Protocol (DHCP)–enabled client computers running earlier versions of Microsoft operating systems can have their A resource records registered and updated by proxy if they obtain their IP lease from a qualified DHCP server. (Only the Windows 2000 and Windows Server 2003 DHCP Server service currently supports this feature.)

The host A resource record is not required for all computers, but it is required by computers that share resources on a network. Any computer that shares resources and needs to be identified by its DNS domain name must use A resource records to provide DNS name resolution to the IP address for the computer.

Most A resource records that are required in a zone can include other workstations or servers that share resources, other DNS servers, mail servers, and Web servers. These resource records make up the majority of resource records in a zone database.

**Alias CNAME Resource Records**

Alias CNAME resource records are also sometimes called canonical name resource records. With these records, you can use more than one name to point to a single host, making it easy to do such things as host both a File Transfer Protocol (FTP) server and a Web server on the same computer. For example, the well-known server names (ftp, www) are registered by using CNAME resource records that map to the DNS host name — for example, server1 — for the server computer that hosts these services.

CNAME resource records are recommended for use in the following scenarios:

- When a host that is specified in an A resource record in the same zone needs to be renamed
- When a generic name for a well-known server, such as www, must resolve to a group of individual computers (each with individual A resource records) that provide the same service, for example, a group of redundant Web servers

When you rename a computer with an existing A resource record in the zone, you can use a CNAME resource record

temporarily to allow a grace period for users and programs to switch from specifying the old computer name to using the new one. To do this, you need the following:

- For the new DNS domain name of the computer, a new A resource record is added to the zone.
- For the old DNS domain name, a CNAME resource record is added that points to the new A resource record.
- The original A resource record for the old DNS domain name (and its associated PTR resource record, if applicable) is removed from the zone.

When you use a CNAME resource record for aliasing or renaming a computer, set a temporary limit on how long the record is used in the zone before it is removed from DNS. If you forget to delete the CNAME resource record and later its associated A resource record is deleted, the CNAME resource record can waste server resources by trying to resolve queries for a name that is no longer used on the network.

The most common or popular use of a CNAME resource record is to provide a permanent, DNS-aliased domain name for generic name resolution of a service-based name, such as www.sales.wingtiptoys.com, to more than one computer or one IP address that is used in a Web server. For example, the following shows the basic syntax of how a CNAME resource record is used:

*alias_name* **IN CNAME** *primary_canonical_name*

In this example, a computer named host-a.sales.wingtiptoys.com must function as both a Web server named www.sales.wingtiptoys.com. and an FTP server named ftp.sales.wingtiptoys.com. To achieve the intended use for naming this computer, you can add and use the following CNAME entries in the sales.wingtiptoys.com zone:

```
host-a     IN  A      10.0.0.20
ftp        IN  CNAME  host-a
www        IN  CNAME  host-a
```

If you later decide to move the FTP server to another computer, separate from the Web server on host-a, simply change the CNAME resource record in the zone for ftp.sales.wingtiptoys.com and add an additional A resource record to the zone for the new computer hosting the FTP server.

Based on the earlier example, if the new computer is named hostb.sales.wingtiptoys.com, the new and revised A and CNAME resource records are as follows:

```
host-a     IN  A      10.0.0.20
host-b     IN  A      10.0.0.21
ftp        IN  CNAME  host-b
www        IN  CNAME  host-a
```

**MX Resource Records**

The MX resource record is used by e-mail applications to locate a mail server based on a DNS domain name that is used in the destination address for the e-mail recipient of a message. For example, a DNS query for the name sales.wingtiptoys.com can be used to find an MX resource record, which enables an e-mail application to forward or exchange mail to a user with the e-mail address user@wingtiptoys.com.

The MX resource record shows the DNS domain name for the computer or computers that process e-mail for a domain. If multiple MX resource records exist, the DNS Client service attempts to contact e-mail servers in the order of preference from lowest value (highest priority) to highest value (lowest priority). The following shows the basic syntax for use of an MX resource record:

*mail_domain_name* **IN MX** *preference mailserver_host*

By using the MX resource records shown below in the sales.wingtiptoys.com zone, e-mail that is addressed to user@sales.wingtiptoys.com is delivered to user@mailserver0.sales.wingtiptoys.com first, if possible. If this server is unavailable, the resolver client can then use user@mailserver1.sales.wingtiptoys.com instead.

```
@          IN  MX   1    mailserver0
@          IN  MX   2    mailserver1
```

Note that the use of the "at" symbol (@) in the records indicates that the mailer DNS domain name is the same as the name of origin (sales.wingtiptoys.com) for the zone.

**PTR Resource Records**

PTR resource records are used to support the reverse lookup process, based on zones that are created and rooted in the in-addr.arpa domain. These records are used to locate a computer by its IP address and to resolve this information to the DNS domain name for that computer.

PTR resource records can be added to a zone in several ways:

- You can create a PTR resource record for a static TCP/IP client computer manually by using DNS, either as a separate procedure or as part of the procedure for creating an A resource record.

- Computers use the DHCP Client service to dynamically register and update their PTR resource record in DNS when an IP configuration change occurs.

- All other DHCP-enabled client computers can have their PTR resource records registered and updated by the DHCP server if they obtain their IP lease from a qualified server. The Windows 2000 and Windows Server 2003 DHCP Server service provides this capability.

The PTR resource record is used only in reverse lookup zones to support reverse lookup.

**SRV Resource Records**
To locate Active Directory domain controllers, SRV resource records are required. Typically, you can avoid manual administration of the SRV resource record when you install Active Directory.

By default, the Active Directory Installation Wizard attempts to locate a DNS server based on the list of preferred or alternate DNS servers, which are configured in any of its TCP/IP client properties, for any of its active network connections. If a DNS server that can accept dynamic update of the SRV resource record (and other resource records that are related to registering Active Directory as a service in DNS) is contacted, the configuration process is complete.

If, during the installation, a DNS server that can accept updates for the DNS domain name that is used to name your Active Directory domain is not found, the wizard can install a DNS server locally and automatically configure it with a zone to support the Active Directory domain.

For example, if the Active Directory domain that you choose for your first domain in the forest is sales.wingtiptoys.com, a zone that is rooted at the DNS domain name of sales.wingtiptoys.com is added and configured to use with the DNS server that is running on the new domain controller.

Whether or not you install the DNS Server service locally, a file (Netlogon.dns) is written and created during the Active Directory installation process that contains the SRV resource records and other resource records that are necessary to support the use of Active Directory. This file is created in the systemroot\System32\Config folder.

If you are using a DNS server that fits one of the following scenarios, use the records in Netlogon.dns to manually configure the primary zone on that server to support Active Directory:

1. The computer that operates your DNS server is running on another platform, such as UNIX, and it cannot accept or recognize dynamic updates.

2. A DNS server at this computer that does not use the DNS Server service that is provided with Windows Server 2003 is authoritative for the primary zone that corresponds to the DNS domain name for your Active Directory domain.

3. The DNS server supports the SRV resource record, as defined in the Internet draft "A DNS RR specifying the location of services (DNS SRV)," but the DNS server does not support dynamic updates.

   For example, the DNS Server service that is provided with Windows NT Server 4.0, when it is updated to Service Pack 4 or later, fits this description.

In the future, the SRV resource record might also be used to register and look up other well-known TCP/IP services on your network if applications implement and support DNS name queries that specify this record type.

**Other Resource Records**
Other additional resource records are supported by Windows Server 2003 DNS, and they are used less frequently in most zones. You can add these additional types of resource records as needed by using the DNS snap-in.

### Task requirements

To begin this task, perform the following requirements:

- Install Dnscmd.

To complete this task, perform one of the following procedures:

- Add an A resource record to a zone
- Add an MX resource record to a zone
- Add a CNAME resource record to a zone
- Add a PTR resource record to a reverse zone
- Add a resource record to a DNS zone
- Add a domain to a zone
- Modify an existing resource record
- Delete a resource record
- View unsupported resource records

# Add an A resource record to a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to add a host address (A) resource record to a zone. Pointer (PTR) resource records that are created automatically when you add an A resource record to a zone are deleted automatically if the corresponding A resource record is deleted.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Adding an A resource record to a zone**

- Using the Windows interface
- Using the command line

### To add an A resource record to a zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A)**.
3. In **Name**, type the Domain Name System (DNS) computer name for the new host.
4. In **IP address**, type the Internet Protocol (IP) address for the new host.
5. As an option, select the **Create associated pointer (PTR) record** check box to create an additional PTR resource record in a reverse zone for this host, based on the information that you enter in **Name** and **IP address**.
6. Click **Add Host** to add the new host (A) resource record to the zone.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To add an A resource record to a zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName***/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **A** *IPAddress*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Required. Adds a new resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is manually updated or removed. |

| | |
|---|---|
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. The default TTL is defined in the start-of-authority (SOA) resource record. |
| **A** | Required. Specifies the resource record type of the record that you are adding. |
| *IPAddress* | Required. The IP address for the host. |

# Add an MX resource record to a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to add a mail exchanger (MX) resource record to a zone. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Adding an MX resource record to a zone**

- Using the Windows interface
- Using the command line

### To add an MX resource record to a zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable forward lookup zone, and then click **NewMail Exchanger (MX)**.
3. In **Host or child domain**, type the domain name for which this record is to be used to deliver mail.
4. In **Mail server**, type the Domain Name System (DNS) host computer name of the mail exchanger or mail server host that delivers mail for the specified domain name.

   As an option, you can click **Browse** to view the DNS namespace for mail exchanger hosts in this domain that have host (A) records already defined.
5. Adjust the value in **Mail server priority** as needed for this zone.
6. Click **OK** to add the new record to the zone.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To add an MX resource record to a zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName***/RecordAdd** *ZoneNameNodeName* [**/Aging**] [*Ttl*] **MX** *PreferenceMXServerName*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Adds a new resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone in which you will add the new MX resource record. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |

| | |
|---|---|
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. |
| **MX** | Required. Specifies the MX resource record type for the record that you are adding. |
| *Preference* | Required. Specifies a numeric value (between 0 and 65535) that indicates the mail exchange server's priority with respect to the other mail exchange servers. Lower numbers are given greater preference. |
| *MXServerName* | Required. Specifies the FQDN for a mail exchanger. The value entered here must resolve to a corresponding host A resource record in this zone. |

# Add a CNAME resource record to a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to add an alias canonical (CNAME) resource record to a zone. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Adding a CNAME resource record to a zone**

- Using the Windows interface
- Using the command line

### To add a CNAME resource record to a zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Alias (CNAME)**.
3. In **Alias name**, type the alias name.
4. In **Fully qualified domain name (FQDN) for target host**, type the FQDN of the Domain Name System (DNS) host computer for which this alias is to be used.

   As an option, you can click **Browse** to search the DNS namespace for hosts in this domain that have host address (A) records already defined.
5. Click **OK** to add the new record to the zone.

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To add a CNAME resource record to a zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **CNAME** *HostName*|*DomainName*

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Adds a new resource record. |
| *ZoneName* | Required. Specifies the name of the zone where this CNAME resource record will be added. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is aged and scavenged. If this parameter is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |

| | |
|---|---|
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. The default TTL is defined in the start-of-authority (SOA) resource record. |
| **CNAME** | Required. Specifies the resource record type of the record that you are adding. |
| *HostName*\|*DomainName* | Required. Specifies the FQDN of any valid DNS host or domain name in the namespace. For FQDNs, a trailing period (.) is used to fully qualify the name. |

# Add a PTR resource record to a reverse zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to add a pointer (PTR) resource record to a reverse zone. When you create a new address (A) resource record, there is an option to create an associated PTR resource record automatically. PTR resource records that are created automatically during the addition of an A resource record to a zone are deleted automatically if the corresponding A resource record is deleted.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Adding a PTR resource record to a reverse zone**

- Using the Windows interface
- Using the command line

### To add a PTR resource record to a reverse zone using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable reverse lookup zone.
3. On the **Action** menu, click **New Pointer (PTR)**.
4. In the **Host IP number** text box, type the host Internet Protocol (IP) address octet number.
5. In **Host name**, type the fully qualified domain name (FQDN) for the DNS host computer for which this pointer record is to be used to provide reverse lookup (address-to-name resolution).

   As an option, you can click **Browse** to search the Domain Name System (DNS) namespace for hosts in this domain that have host address (A) records already defined.
6. Click **OK** to add the new record to the zone.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To add a PTR resource record to a reverse zone using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName***/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [**Ttl**] **PTR** *HostName*|*DomainName*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Adds a new resource record. |
| *ZoneName* | Required. Specifies the FQDN of the zone where this new PTR resource record will be added. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |

| | |
|---|---|
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is updated or removed manually. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| **Ttl** | Specifies the Time to Live (TTL) setting for the resource record. |
| **PTR** | Required. Specifies the resource record type for the record that you are adding. |
| *HostName*\|*DomainName* | Required. Specifies the FQDN of a resource record that is located in the DNS namespace. The host that you specify is used as the data for answering reverse lookups based on the address information that is specified by this PTR resource record. |

# Add a resource record to a DNS zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to add a resource record to a zone. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Adding a resource record to a zone**

- Using the Windows interface
- Using the command line

**To add a resource record to a zone using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Other New Records**.
3. In **Select a resource record type**, select the type of resource record that you want to add.
4. Click **Create Record**.
5. In **New Resource Record**, enter the information necessary to complete the resource record.
6. After you specify all the necessary information for the resource record, click **OK** to add the new record to the zone.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To add a resource record to a zone using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/RecordAdd***ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [**Ttl**] *RRType RRData*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Required. Adds a new resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is updated or removed manually. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |

| | |
|---|---|
| **Ttl** | Specifies the Time to Live (TTL) setting for the resource record. |
| *RRTypeRRData* | Required. Specifies the type of resource record to add, followed by the data to be contained in the resource record. |

| **Resource record type** | **Resource record data** |
|---|---|
| A | *IPAddress* |
| NS,CNAME,MB,MD,PTR,MF,MG,MR | *HostName*\|*DomainName* |
| MX,RT,AFSDB | *PreferenceServerName* |
| SRV | *PriorityWeightPortHostName* |
| SOA | *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL* |
| AAAA | *Ipv6Address* |
| TXT,X25,HINFO,ISDN | *String* [*String*] |
| MINFO,RP | *MailboxName ErrMailboxName* |
| WKS | *Protocol IPAddress Service...* |
| WINS | *MapFlag LookupTimeout CacheTimeout IPAddress...* |
| WINSR | *MapFlag LookupTimeout CacheTimeout RstDomainName* |

| **Value** | **Description** |
|---|---|
| *IPAddress* | Specifies a standard IP address, for example, 255.255.255.255. |
| *ipv6Address* | Specifies a standard IPv6 address, for example, 1:2:3:4:5:6:7:8. |
| *Protocol* | Specifies the transmission protocol: UDP or TCP. |
| *Service* | Specifies a standard service, for example, domain, smtp. |
| *HostName*\|*DomainName* | Specifies the FQDN of a resource record that is located in the DNS namespace. |

# Add a domain to a zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to add a new domain to a zone.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To add a domain to a zone**

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.

   Where?

   - DNS/*applicable DNS server*/Forward Lookup Zones (or Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **New Domain**, and then type the name of the new domain without using periods.
4. Click **OK** to add the new domain to the zone.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Modify an existing resource record

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to modify an existing resource record in a zone. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Modifying an existing resource record**

- Using the Windows interface
- Using the command line

**To modify an existing resource record using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.
3. In the details pane, right-click the resource record that you want to modify, and then click **Properties**.
4. In **Properties**, edit the properties that can be modified.

   If necessary, you can view and modify advanced resource record properties with the DNS snap-in. To display advanced properties, on the **View** menu, click **Advanced**.
5. When you have finished modifying the record, click **OK**.

| **Note** |
| --- |
| When advanced view options are enabled, you can modify additional settings for an existing resource record, such as its record-specific Time to Live (TTL). |

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To modify an existing resource record using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName***/RecordAdd***ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] *RRTypeRRData*

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Required. Adds a new resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |

| | |
|---|---|
| *RRType RRData* | Required. Specifies the type of resource record to add, followed by the data to be contained in the resource record. |

| **Resource record type** | **Resource record data** |
|---|---|
| A | *IPAddress* |
| NS,CNAME,MB,MD,PTR,MF,MG,MR | *HostName*\|*DomainName* |
| MX,RT,AFSDB | *Preference ServerName* |
| SRV | *Priority Weight Port HostName* |
| SOA | *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL* |
| AAAA | *Ipv6Address* |
| TXT,X25,HINFO,ISDN | *String* [*String*] |
| MINFO,RP | *MailboxName ErrMailboxName* |
| WKS | *Protocol IPAddress Service...* |
| WINS | *MapFlag LookupTimeout CacheTimeout IPAddress...* |
| WINSR | *MapFlag LookupTimeout CacheTimeout RstDomainName* |

| **Value** | **Description** |
|---|---|
| *IPAddress* | Specifies a standard IP address, for example, 255.255.255.255. |
| *ipv6Address* | Specifies a standard IPv6 address, for example, 1:2:3:4:5:6:7:8. |
| *Protocol* | Specifies the transmission protocol: UDP or TCP. |
| *Service* | Specifies a standard service, for example, domain, smtp. |
| *HostName*\|*DomainName* | Specifies the FQDN of a resource record that is located in the DNS namespace. |

# Delete a resource record

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to delete a resource record from a zone. Pointer (PTR) resource records are deleted automatically if the corresponding address (A) resource record is deleted.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Deleting a resource record**

- Using the Windows interface
- Using the command line

### To delete a resource record using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.
3. In the details pane, right-click the resource record that you want to delete, and then click **Delete**.
4. When you are asked to confirm that you want to delete the selected resource record, click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To delete a resource record using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/RecordDelete** *ZoneName NodeName RRType RRData* [**/f**]

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordDelete** | Required. Deletes a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| *RRType RRData* | Required. Specifies the type of resource record to add, followed by the data to be contained in the resource record. |

| Resource record type | Resource record data |
|---|---|
| A | *IPAddress* |
| NS,CNAME,MB,MD,PTR,MF,MG,MR | *HostName*\|*DomainName* |
| MX,RT,AFSDB | *PreferenceServerName* |
| SRV | *Priority Weight Port HostName* |
| SOA | *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL* |
| AAAA | *Ipv6Address* |
| TXT,X25,HINFO,ISDN | *String* [*String*] |
| MINFO,RP | *MailboxName ErrMailboxName* |
| WKS | *Protocol IPAddress Service...* |
| WINS | *MapFlag LookupTimeout CacheTimeout IPAddress...* |
| WINSR | *MapFlag LookupTimeout CacheTimeout RstDomainName* |

| Value | Description |
|---|---|
| *IPAddress* | Specifies a standard IP address, for example, 255.255.255.255. |
| *ipv6Address* | Specifies a standard IPv6 address, for example, 1:2:3:4:5:6:7:8. |
| *Protocol* | Specifies the transmission protocol: UDP or TCP. |
| *Service* | Specifies a standard service, for example, domain, smtp. |
| *HostName*\|*DomainName* | Specifies the FQDN of a resource record that is located in the DNS namespace. |
| **/f** | Specifies that the command is executed without asking for confirmation. If you omit this parameter, you are prompted to confirm the deletion of the resource record. |

# View unsupported resource records

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to view unsupported resource records in a zone.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To view unsupported resource records**

1. Open the DNS snap-in.
2. In the console tree, click the applicable zone.
3. In the details pane, right-click the record that you want to view, and then click **Properties**.
4. In **Properties**, view properties that are specific to this record.
5. When you have finished viewing the record, click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Disable NS resource record registration

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following procedure restricts name server (NS) resource records that are registered for Active Directory domain controllers only. You can perform this procedure by using the Registry Editor or by using the Dnscmd command-line tool.

To configure the Domain Name System (DNS) server to automatically add NS resource records corresponding to itself when loading a zone, you can assign a value of **0x0** to the registry key or enter no value (the default setting). This setting has the same effect as not creating the **DisableNSRecordsAutoCreation** registry entry.

If you configure the registry to restrict the DNS server from registering NS resource records for authoritative zones, any existing NS resource records for the authoritative zones that are located on the DNS server are deleted automatically.

Regardless of the settings of these registry entries, query responses that are sent to DNS clients from the authoritative DNS server will indicate that the responses are from an authoritative DNS server.

The registry key entry that is described in this procedure does not exist by default. It must be created and configured according to this procedure.

| **Caution** |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

### Administrative credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

#### Disabling NS resource record registration

- Using the Windows interface
- Using the command line

### To disable NS resource record registration using the Windows interface

1. Open Registry Editor.
2. In Registry Editor, navigate to the following registry key:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters**

3. Add the following **REG_DWORD** value:

   **DisableNSRecordsAutoCreation**

4. Assign a value of **0x1**.

   The **REG_DWORD** value is a local DNS server setting, and it applies to DNS zones for which this DNS server is authoritative.

| **Note** |
| --- |
| To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**. |

### To disable NS resource record registration using the command line

- At a command prompt, type the following command, and then press ENTER:

**dnscmd** *ServerName* **/Config /DisableNSRecordsAutoCreation0x1**

| Value | Description |
| --- | --- |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/DisableNSRecordsAutoCreation** | Determines the local DNS server configuration for registering NS resource records for authoritative zones. |
| **0x1** | Specifies that the DNS server that is specified in *ServerName* should not add NS resource records for authoritative zones.<br><br>To specify that the DNS server should add NS resource records for all its authoritative zones, type a value of **0x0**. |

# Allow NS resource record creation for domain controllers

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to allow name server (NS) resource record creation for specific domain controllers. This procedure applies to domain controller NS resource records in Active Directory–integrated Domain Name System (DNS) zones that are hosted on DNS servers that are configured to not add these resource records for their authoritative zones.

**Administrative credentials**

To perform this procedure, you must be a member of the DnsAdmins or the Domain Admins group in Active Directory. As a security best practice, consider using the **Run as** command to perform this procedure.

**To allow NS resource record creation for specific domain controllers**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config** *ZoneName* **/AllowNSRecordsAutoCreation** *IpAddresses...*

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| **/AllowNSRecordsAutoCreation** | Required. Specifies that domain controllers that are entered for *Value* add their names to NS resource records for the zone that is specified in *ZoneName*. NS resource records that were previously registered for this zone are not affected. Therefore, you must remove them manually if you do not want them. |
| *IpAddresses...* | Required. Specifies the IP addresses of the domain controllers that add their names in NS resource records for the zone that is specified in *ZoneName*. Type a space-separated list of the IP addresses of the DNS servers, for example, 10.0.0.0 172.16.0.0 192.168.0.0. |

If any domain controllers in the specified zone are not listed for *IpAddresses...*, their names are deleted from the NS resource records for the zone that is specified in *ZoneName*.

To specify that all domain controllers are allowed to add their names to NS resource records for the zone or to clear the list of allowed DNS server IP addresses, type the command and omit *IpAddresses...*:

**dnscmd** *ServerName* **/Config** *ZoneName* **/AllowNSRecordsAutoCreation**

Regardless of the settings that are specified in this command, query responses that are sent to DNS clients from authoritative DNS servers and selected domain controllers will indicate that the responses are from authoritative DNS servers.

# Restrict the DNS resource records that are updated by Netlogon

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following procedure restricts Domain Name System (DNS) resource records that are registered by the Net Logon service for Active Directory domain controllers only.

| ⚠Caution |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To restrict the DNS resource records that are updated by NetlLogon**

1. Open Registry Editor.
2. In Registry Editor, navigate to the following registry key:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**
3. Add the following multistring value **(REG_MULTI_SZ)** value:

   **DnsAvoidRegisterRecords**
4. In this value, specify the list of data corresponding to the DNS resource records that *should not* be registered for this domain controller by the Net Logon service. The following table contains the list of data.

| Data Value | Resource Record Type | DNS Resource Record |
| --- | --- | --- |
| LdapIpAddress | A | *DnsDomainName* |
| Ldap | SRV | _ldap._tcp.*DnsDomainName* |
| LdapAtSite | SRV | _ldap._tcp.*SiteNam*._sites.*DnsDomainName* |
| Pdc | SRV | _ldap._tcp.pdc._msdcs.*DnsDomainName* |
| Gc | SRV | _ldap._tcp.gc._msdcs.*DnsForestName* |
| GcAtSite | SRV | _ldap._tcp.*SiteName*._sites.gc._msdcs.*DnsForestName* |
| DcByGuid | SRV | _ldap._tcp.*DomainGuid*.domains._msdcs.*DnsForestName* |
| GcIpAddress | A | _gc._msdcs.*DnsForestName* |
| DsaCname | CNAME | *DsaGuid*._msdcs.*DnsForestName* |
| Kdc | SRV | _kerberos._tcp.dc._msdcs.*DnsDomainName* |

| | | |
|---|---|---|
| KdcAtSite | SRV | _kerberos._tcp.dc._msdcs.*SiteName*._sites.*DnsDomainName* |
| Dc | SRV | _ldap._tcp.dc._msdcs.*DnsDomainName* |
| DcAtSite | SRV | _ldap._tcp.*SiteName*._sites.dc._msdcs.*DnsDomainName* |
| Rfc1510Kdc | SRV | _kerberos._tcp.*DnsDomainName* |
| Rfc1510KdcAtSite | SRV | _kerberos._tcp.*SiteName*._sites.*DnsDomainName* |
| GenericGc | SRV | _gc._tcp.*DnsForestName* |
| GenericGcAtSite | SRV | _gc._tcp.*SiteName*._sites.*DnsForestName* |
| Rfc1510UdpKdc | SRV | _kerberos._udp.*DnsDomainName* |
| Rfc1510Kpwd | SRV | _kpasswd._tcp.*DnsDomainName* |
| Rfc1510UdpKpwd | SRV | _kpasswd._udp.*DnsDomainName* |

**Notes**

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

- Restart of the Net Logon service is *not* required to make the changes to this value effective. If the **DnsAvoidRegisterRecords** registry key is created or modified while the Net Logon service is stopped or within the first 15 minutes after it is started, appropriate DNS updates may take place with a short delay. However, the delay is no later than 15 minutes after the Net Logon service starts.

# Monitoring DNS

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Monitoring your Domain Name System (DNS) infrastructure on a regular basis and resolving any issues that you find will help to keep your network accessible to your users.

The following tasks for monitoring DNS are described in this objective:

- Check DNS event log
- Verify DNS server responsiveness with Nslookup
- Verify dynamic DNS record updates
- Verify zone transfers
- Check the DNS server debug log file
- Test a query on the DNS server

# Check DNS event log

Published: September 19, 2007

Updated: November 4, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to check for errors in the Domain Name System (DNS) event long for troubleshooting or monitoring purposes.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Checking the DNS event log**
**To check the DNS event log**

1. Open the DNS console.
2. In the console tree, open **Event Viewer** and then click **DNS Events**.
3. Review the list of events that is displayed. The following table provides examples some critical DNS events and procedures you can use to resolve them.

| Event ID | Description |
|---|---|
| 140 | The DNS server could not initialize the Remote Procedure Call (RPC) service. If it is not running, start the RPC service or reboot the computer. For specific error code, see the Record Data page on the Event Viewer.<br><br>In order for DNS to run, the Remote Procedure Call (RPC) service must be running on the DNS server.<br><br>a. Verify that the Remote Procedure Call (RPC) service has been started.<br><br>b. Open Administrative Tools, and double-click Services.<br><br>c. If the service has been started, try restarting the server.<br><br>d. If the error continues, remove and reinstall the **Client for Microsoft Networks** service on the network connection. This will reinstall the Netlogon and RPC locator services. |
| 403 | The DNS server could not create a Transmission Control Protocol (TCP) socket. Restart the DNS server or reboot the computer. For the specific error code, see the Record Data page.<br><br>The Wsock32.dll might be incompatible with a third-party TCP/IP stack. This problem can also occur if the TCP/IP protocol is not bound to the network adapter.<br><br>If you are using a third-party TCP/IP protocol, verify that the protocol is compatible with the Wsock32.dll.<br><br>Check the bindings of the protocol stack. It is a good idea to have TCP/IP bound at the top of the stack. If the error continues, remove and reinstall the TCP/IP protocol, and then try again.<br><br>a. Open Control Panel, and then double-click Network and Dial-up Connections.<br><br>b. Right-click the connection, and then click Properties.<br><br>c. Verify that the bindings for all protocols to network adapters are enabled and that no broken connections exist in the stack. |

| | |
|---|---|
| 407 | DNS server could not bind the main datagram socket. The data is the error.<br><br>This error can occur if there is a mismatch between the configured IP address in the Advanced IP Addressing dialog box and the addresses listed in the Server Properties dialog box for the DNS server. This problem can also occur if the TCP/IP protocol is not bound to the network adapter.<br><br>Verify that the TCP/IP addresses configured in the Advanced IP Addressing dialog box match those configured in the Server Properties dialog box in DNS Manager:<br><br>a. Open Control Panel, and double-click Network.<br><br>b. Click the Protocols tab, and click TCP/IP Protocol in the Network Protocols list.<br><br>c. Click Properties, and then click Advanced.<br><br>Match the IP addresses to those displayed in the DNS server Properties dialog box:<br><br>a. In DNS Manager, right-click the DNS server name, and then click Properties.<br><br>b. Compare the IP addresses with those from the Advanced IP Addressing dialog box. If there are no IP addresses configured in the Advanced IP Addressing dialog box or on the Interfaces tab of the Server Properties dialog box, enter the IP address of your network adapter. Use the ipconfig -all command to obtain your IP address.<br><br>Check the binding of the TCP/IP protocol to the network adapter:<br><br>a. Open Control Panel, and double-click Network.<br><br>b. Click the Bindings tab.<br><br>c. Verify that the bindings for all protocols to network adapters are enabled and that no broken connections exist in the stack. |
| 408 | DNS server could not open socket for address [IP address of server].<br><br>The DNS server could not open a socket with the current TCP/IP and DNS service configurations.<br><br>Verify that this is a valid IP address on this computer.<br><br>If the IP is not valid:<br><br>a. Use the Interfaces dialog under Server Properties in the DNS Manager to remove it from the list of IP interfaces.<br><br>b. Stop and restart the DNS server. (If this was the only IP interface on this computer, the DNS server may not have started as a result of this error. In that case, remove the DNS\Parameters\ListenAddress value in the services section of the registry and restart.)<br><br>If the IP is valid:<br><br>Verify that no other application (for example, another DNS server) is running that would attempt to use the DNS port. |

| | |
|---|---|
| 4000, 4004, 4007, 4014, 4015 | The DNS Server service relies on Active Directory to store and retrieve information for Active Directory–integrated zones. This error indicates that Active Directory is not responding to requests from the DNS Server service. Ensure that Active Directory is functioning properly, troubleshoot any problems, and then restart the DNS Server service.<br><br>For information about troubleshooting Active Directory, see Active Directory Troubleshooting Topics (http://go.microsoft.com/fwlink/?LinkId=95789).<br><br>To restart the DNS Server service:<br><br>a. Open the Services console. To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then click **Services**.<br><br>b. Right-click **DNS Server**, and then click **Restart**.<br><br>If the problem continues, restart the computer, and then use the Services console to verify that the DNS Server service has started. |
| 4001 | The DNS Server service relies on Active Directory to store and retrieve information for Active Directory–integrated zones. This error indicates that Active Directory is not responding to requests from the DNS Server service. Ensure that Active Directory is functioning properly, troubleshoot any problems, and then reload the zone.<br><br>For information about troubleshooting Active Directory, see Active Directory Troubleshooting Topics (http://go.microsoft.com/fwlink/?LinkId=95789).<br><br>To reload a zone:<br><br>a. Open the DNS console.<br><br>b. In the console tree, right-click the applicable zone, and then click **Reload**. |

| | |
|---|---|
| 4016 | The DNS Server service relies on Active Directory to store and retrieve information for Active Directory–integrated zones. This error indicates that Active Directory is not responding to requests from the DNS Server service. Ensure that Active Directory is functioning properly, troubleshoot any problems, and then retry the operation that failed.<br><br>For information about troubleshooting Active Directory, see Active Directory Troubleshooting Topics (http://go.microsoft.com/fwlink/?LinkId=95789).<br><br>Add a zone<br><br>If the event message indicates that an attempt to add a zone failed, you must create the zone after resolving any problems with Active Directory.<br><br>To add a zone:<br><br>a. Open the DNS console.<br><br>b. In the console tree, expand the DNS server, right-click the zone folder for the type of zone that you want to add, and then click **New Zone** to open the New Zone Wizard.<br><br>c. Follow the instructions in the wizard to create the zone.<br><br>Delete a zone<br><br>If the event message indicates that an attempt to remove a zone failed, you must delete the zone after resolving any problems with Active Directory.<br><br>To delete a zone:<br><br>a. Open the DNS console.<br><br>b. In the console tree, expand the DNS server, right-click the zone folder for the type of zone that you want to delete.<br><br>c. Right-click the zone, and then click **Delete**. |

| **Note** |
|---|
| To open the DNS console, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

| **Note** |
|---|
| If the DNS server for which you want to view the log is located on another computer, in the console tree, click **DNS**, and then on the **Action** menu, click **Connect to DNS Server**. Click **The following computer**, and then specify the name or Internet Protocol (IP) address of the remote computer. |

# Verify DNS server responsiveness with Nslookup

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to verify name resolution on a Domain Name System (DNS) server for monitoring or troubleshooting purposes.

**Verifying DNS server responsiveness**
**To verify DNS server responsiveness**

1. Open a command prompt, type the following command, and then press ENTER:

   **nslookup 127.0.0.1** *server_ip_address*

2. If the server is responding, the name "localhost" is returned.

| Value | Description |
|---|---|
| *server_ip_address* | The Internet Protocol (IP) address of the DNS server. For example, if the IP address of your DNS server is 10.0.0.1, type:<br><br>**nslookup 127.0.0.1 10.0.0.1** |

**Note**

In the previous procedure, the syntax for the nslookup command is: **nslookup***[-option] host server* . This command can be entered from any computer that is running a Microsoft Windows operating system and has network connectivity to the DNS server you wish to query. Only the host entry is required for the command. However, if an IP address or hostname for the server is not supplied, then the default DNS server specified in TCP/IP properties will be queried.

When you enter 127.0.0.1 as the host, this IP address will automatically resolve to the name **localhost** if the DNS Server service is running at the IP address that you specify as the *server*.

# Verify dynamic DNS record updates

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is possible to view dynamic update activity of clients by using the DNS server log (dns.log). This log includes server-side information; client-side information is intuitive only.

**Verifying dynamic DNS record updates**

**To verify dynamic DNS record updates**

1. Click **Start**, click **Run**, type **dnsmgmt.msc**, and then click **OK**.
2. In the console tree, click the DNS server for which you want to view dynamic record updates.
3. On the **Action** menu, click **Properties**.
4. Click the **Debug Logging** tab.
5. Select the **Log packets for debugging** check box.
6. Under **Packet direction**, select the **Outgoing** and **Incoming** check boxes.
7. Under **Transport protocol**, select the **UDP** and **TCP** check boxes.
8. Under **Packet contents**, select the **Updates** check box.
9. Under **Packet type**, select the **Request** and **Response** check boxes.
10. Under **Other options**, select the **Details** check box.
11. Click **OK** to begin debug logging.
12. Review the dns.log file.

# Verify zone transfers

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

With Active Directory–integrated Domain Name System (DNS), all DNS servers in the domain can modify the zone and then replicate the changes to other domain controllers. Therefore, the procedures for verifying zone transfers of an Active Directory–integrated DNS server are equivalent to Active Directory replication verification procedures.

Although Active Directory-integrated zones are transferred by using Active Directory replication, you can also perform standard zone transfers to secondary servers in a manner similar to standard DNS zone transfers. You can use the following procedure to verify zone transfers.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Verifying zone transfers**

**To verify zone transfers**

1. Open the DNS snap-in.
2. In the console tree, double-click **Event Viewer**, and then click **DNS Events**.
3. Check for the following critical events:

| Event ID | Description |
|---|---|
| 6527 | Zone expired before it could obtain a successful zone transfer or update from a master server acting as its source for the zone. The zone has been shut down.<br><br>This event ID might appear when the DNS server is configured to host a secondary copy of the zone from another DNS server acting as its source or master server. Verify that this server has network connectivity to its configured master server.<br><br>If the problem continues, consider one or more of the following options:<br><br>a. Delete the zone and recreate it, specifying either a different master server, or an updated and corrected IP address for the same master server.<br><br>b. If zone expiration continues, consider adjusting the expire interval. |
| 6004 | The DNS server received a zone transfer request from %1 for a non-existent or non-authoritative %2. |

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Check the DNS server debug log file

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Checking the DNS server debug log file**

### To check the DNS server debug log file

1. Stop the DNS Server service.
2. Open WordPad.
3. On the **File** menu, click **Open**.
4. In **Open**, for **File name**, specify the path to the Domain Name System (DNS) server debug log file.

   By default, if the applicable DNS server is running locally, the file and path are as follows:

   systemroot\System32\Dns\Dns.log
5. After you specify the correct path and file, click **Open** to view the log file.

# Test a query on the DNS server

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To insure that Domain Name System (DNS) name resolution is functioning according to specifications, you can perform periodic testing with the following procedure.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Testing a query on the DNS server**

**To test a query on the DNS server**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.
3. On the **Action** menu, click **Properties**.
4. Click the **Monitoring** tab.
5. To test a simple query, under **Select a test type**, select the **A simple query against this DNS server** check box.

   Or

   To test a recursive query, under **Select a test type**, select the **A recursive query to other DNS servers** check box.
6. Click **Test Now**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Optimizing DNS

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When Domain Name System (DNS) servers are initialized for service, they use server configuration settings that are taken from the parameters that are stated in a boot information file, the registry, and (possibly) zone information that is provided through Active Directory integration.

In most situations, the installation defaults are acceptable and should not require modification. However, when necessary, you can tune various parameters to accommodate special deployment needs and situations.

The following table describes advanced parameters that you can change to optimize the performance of DNS servers.

| Parameter | Description |
|---|---|
| *Disable recursion* | Determines whether or not the DNS server uses recursion. By default, the DNS Server service is enabled to use recursion. |
| *BIND secondaries* | Determines whether to use fast transfer format for transfer of a zone to DNS servers running legacy Berkeley Internet Name Domain (BIND) implementations.<br><br>By default, all Windows-based DNS servers use a fast zone transfer format. This format uses compression, and it can include multiple records per TCP message during a connected transfer. This format is also compatible with more recent BIND-based DNS servers that run versions 4.9.4 and later. |
| *Fail on load if bad zone data* | Sets the DNS server to parse files strictly.<br><br>By default, the DNS Server service logs data errors, ignores any erred data in zone files, and continues to load a zone. You can reconfigure this option by using the DNS snap-in so that the DNS Server service logs errors and fails to load a zone file containing records data that is found to contain errors. |
| *Enable round robin* | Determines whether the DNS server uses the round robin mechanism to rotate and reorder a list of resource records if multiple resource records exist of the same type that exist for a query answer.<br><br>By default, the DNS Server service uses round robin. |
| *Enable netmask ordering* | Determines whether the DNS server reorders address (A) resource records within the same resource record that is set in the server's response to a query based on the Internet Protocol (IP) address of the source of the query.<br><br>By default, the DNS Server service uses local subnet priority to reorder A resource records. |
| *Secure cache against pollution* | Determines whether the DNS server attempts to clean up responses to avoid cache pollution. This setting is enabled by default.<br><br>By default, DNS servers use a secure response option that eliminates adding unrelated resource records that are included in a referral answer to their cache. In most cases, any names that are added in referral answers are typically cached, and they help expedite the resolution of subsequent DNS queries.<br><br>With this feature, however, the server can determine that referred names are potentially polluting or insecure and then discard them. The server determines whether to cache the name that is offered in a referral on the basis of whether or not it is part of the exact, related, DNS domain name tree for which the original queried name was made.<br><br>For example, if a query is made originally for sales.wingtiptoys.com and a referral answer provides a record for a name outside the wingtiptoys.com domain name tree, such as tailspintoys.com, that name is not cached where this feature is enabled for use. |

For more information about planning DNS, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

To optimize DNS, complete the following procedures:

- Enable or disable fast DNS zone transfers

- Prevent loading of a zone with bad data
- Disable DNS round robin
- Restore Domain Name System server default preferences
- Disable recursion
- Disable local subnet prioritization
- Update root hints
- Secure the server cache against names pollution
- Clear server names cache
- Configure DNSSEC
- Configure EDNS0
- Change UDP message size

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Enable or disable fast DNS zone transfers

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure for optimizing zone transfers only between Windows-based Domain Name System (DNS) servers and other DNS server implementations. Zone transfers between Windows-based DNS servers always use the fast transfer format.

DNS servers running versions of the Berkeley Internet Name Domain (BIND) server implementation earlier than version 4.9.4 do not support the fast transfer format. Enable this option only if you are transferring zones to BIND servers running version 4.9.4 or later.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Enabling or disabling fast DNS zone transfers**

- Using the Windows interface
- Using the command line

### To enable or disable fast DNS zone transfers using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, select the **BIND secondaries** check box, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To enable or disable fast DNS zone transfers using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config /BindSecondaries** {**1**|**0**}

| Value | Description |
| --- | --- |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/BindSecondaries** | Specifies use of the fast transfer format that is used by legacy BIND servers. |
| {**1**\|**0**} | To enable fast transfer format when transferring a zone to legacy BIND DNS servers, type **1** (on). To disable fast transfer format, type **0** (off). |

# Prevent loading of a zone with bad data

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to prevent loading of a zone when bad data is found.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To prevent loading of a zone with bad data**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, select the **Fail on load if bad zone data** check box, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Disable DNS round robin

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to disable round-robin rotation for multihomed names. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Disabling DNS round robin**

- Using the Windows interface
- Using the command line

### To disable DNS round robin using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, clear the **Enable round robin** check box, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To disable DNS round robin using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config /RoundRobin** {**1**|**0**}

| Value | Description |
| --- | --- |
| *ServerName* | Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RoundRobin** | Configures round-robin rotation. |
| {**1**\|**0**} | To enable round robin, type **1** (on). To disable round robin, type **0** (off). |

# Restore Domain Name System server default preferences

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to restore Domain Name System (DNS) server default preferences. The DNS server default preferences are listed in the following table.

| **Property** | **Setting** |
|---|---|
| Disable recursion | Off |
| BIND secondaries | On |
| Fail on load if bad zone data | Off |
| Enable round robin | On |
| Enable netmask ordering | On |
| Secure cache against pollution | On |
| Name checking | Multibyte (UTF8) |
| Load zone data on startup | From Active Directory and registry |
| Enable automatic scavenging of stale records | Off |

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To restore DNS server default preferences**

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable DNS server, and then click **Properties**.
3. Click the **Advanced** tab.
4. Click **Reset to Default**, and then click **OK**.

| **Note** |
|---|
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Disable recursion

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to disable recursion on the Domain Name System (DNS) server.

| **Note** |
| --- |
| If you disable recursion on the DNS server, you will not be able to use forwarders on the same server. |

### Administrative credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

#### Disabling recursion

- Using the Windows interface
- Using the command line

### To disable recursion using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable DNS server, then click **Properties**.
3. Click the **Advanced** tab.
4. In **Server options**, select the **Disable recursion (also disables forwarders)** check box, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, , point to **Administrative Tools**, and then click **DNS**. |

### To disable recursion using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config /NoRecursion** {**1**|**0**}

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/NoRecursion** | Required. Disables recursion. |
| {**1**\|**0**} | Required. To disable recursion, type **1** (off). To enable recursion, type **0** (on). By default, recursion is enabled. |

# Disable local subnet prioritization

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to disable local subnet prioritization for multihomed names. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Disabling local subnet prioritization**

- Using the Windows interface
- Using the command line

### To disable local subnet prioritization using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, click the applicable Domain Name System (DNS) server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, clear the **Enable netmask ordering** check box, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To disable local subnet prioritization using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config /LocalNetPriority** {**1**|**0**}

| Value | Description |
| --- | --- |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/LocalNetPriority** | Configures netmask ordering. |
| {**1**\|**0**} | To enable netmask ordering, type **1** (on). To disable netmask ordering, type **0** (off). |

# Update root hints

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to update root hints on the Domain Name System (DNS) server.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Updating root hints**

**To update root hints**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Root Hints** tab.
5. Modify server root hints as follows:
   - To add a root server to the list, click **Add**, and then specify the name and Internet Protocol (IP) address of the server to be added to the list.
   - To modify a root server in the list, click **Edit**, and then specify the name and IP address of the server to be modified in the list.
   - To remove a root server from the list, select it in the list, and then click **Remove**.
   - To copy root hints from a DNS server, click **Copy from Server**, and then specify the IP address of the DNS server from which you want to copy a list of root servers to use in resolving queries. These root hints will not overwrite any existing root hints.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Secure the server cache against names pollution

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to secure the server cache against names pollution. This setting is enabled by default.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Secure the server cache against names pollution**

**To secure the server cache against names pollution**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, select the **Secure cache against pollution** check box, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Clear server names cache

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to clear the server names cache. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Clearing the server names cache**

- Using the Windows interface
- Using the command line

### To clear the server names cache using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Clear Cache**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To clear the server names cache using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/clearcache**

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/clearcache** | Clears the DNS server cache. |

# Configure DNSSEC

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to modify the configuration of Domain Name System (DNS) Security Extensions (DNSSEC). The value of the registry entry **EnableDnsSec** determines whether the DNS server includes or excludes DNSSEC resource records when it receives queries.

| **Caution** |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**To configure DNSSEC**

1. Open Registry Editor.
2. In Registry Editor, navigate to the following registry subkey:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters**
3. Add the following **DWORD** entry:

   **EnableDnsSec**
4. Do one of the following:
   - To exclude DNSSEC resource records in query responses other than responses to requests for SIG, KEY or NXT resource records, assign a value of **0x0**. Appropriate resource records will be included in responses to requests for SIG, KEY, or NXT resource records only.
   - To include the DNSSEC resource records in all query responses (according to RFC 2535), assign a value of **0x2**.
   - To include DNSSEC resource records only in cases where the original client query contained the OPT resource record (according to RFC 2671), assign a value of **0x1**, or do not create the value at all. The DNS server behaves the same if the value is **0x1** or if the entry does not appear in the registry.

| **Note** |
| --- |
| To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**. |

# Configure EDNS0

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to modify EDNS0 configuration. The value of the registry key **EDNSCacheTimeout** determines how long the Domain Name System (DNS) server keeps information about the extension mechanisms for DNS (EDNS) versions that are supported by other DNS servers that have responded to a query with an OPT resource record.

You can perform this procedure by using Registry Editor or by using the Dnscmd command-line tool.

| **Caution** |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

### Administrative credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

#### Configuring EDNS0

- Using the Windows interface
- Using the command line

### To configure EDNS0 using the Windows interface

1. Open Registry Editor.
2. In Registry Editor, navigate to the following registry subkey:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters**
3. Add the following **DWORD** entry:

   **EDNSCacheTimeout**
4. To change the cache timeout, type a value in seconds between 3600 (1 hour) and 15724800 (182 days).
5. In the same registry subkey (**Parameters**), add the following **DWORD** entry:

   **EnableEDNSProbes**
6. To configure the DNS server to include an OPT resource record only in response to EDNS0 requests containing OPT resource records, type **0x1** (**DWORD**).
7. Restart the DNS server.

| **Note** |
| --- |
| To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**. |

### To modify EDNS0 configuration using the command line

- At a command prompt, type one of the following commands, and then press ENTER:
  - **dnscmd** *ServerName* **/Config /EDNSCacheTimeout** *Value*

- **dnscmd** *ServerName* **/Config /EnableEDNSProbes** *Value*

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/Config** | Required. Specifies the command to configure the DNS server. |
| **/EDNSCacheTimeout** | Required. Specifies the length of time that the DNS server remembers the EDNS parameters remote servers report. |
| **/EnableEdnsProbes** | Required. Specifies whether or not the DNS server probes other DNS servers to determine if they support EDNS. |
| *Value* | Required. For **/EDNSCacheTimeout**, type a value in seconds between 3600 (1 hour) and 15724800 (182 days). For **/EnableEDNSProbes**, type **1** to configure the DNS server to probe other DNS servers and determine if they support EDNS. Type **0** to configure the DNS server to not probe remote servers for EDNS support. If you type **0**, the DNS server will continue to use EDNS if other servers request it. |

# Change UDP message size

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to modify User Datagram Protocol (UDP) message size.

| **Caution** |
| --- |
| When you configure the UDP packet size to be larger than 512 bytes, remember that UDP packets must travel through devices other than UDP hosts, such as routers, and these devices may not support UDP packets larger than 512 bytes. It is recommended that you establish the maximum UDP packet length support for all devices — and the path's maximum transmission unit (MTU), if possible — and configure your UDP hosts according to this maximum. |

| **Caution** |
| --- |
| It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution. |

### Administrative credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

### To change UDP message size

1. Open Registry Editor.
2. In Registry Editor, navigate to the following registry key:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters**
3. Add the following **DWORD** entry:

   **MaximumUdpPacketSize**
4. Type a maximum UDP packet size value in bytes.

   The default value is 1280 bytes. The value must be between 512 and 16384 in decimal format (200 and 4000 in hexadecimal format).
5. Restart the DNS server.

| **Note** |
| --- |
| To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**. |

# Securing DNS

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain Name System (DNS) was originally designed as an open protocol. Therefore, it is vulnerable to attackers. Security features in Windows Server 2003 DNS can help you prevent an attack on your DNS infrastructure. Before considering which of the Windows Server 2003 security features to use, you should be aware of the following:

- Common threats to DNS security
- The level of DNS security in your organization

**DNS Security Threats**

The following are the typical ways in which your DNS infrastructure can be threatened by attackers:

- **Footprinting.** The process by which DNS zone data, including DNS domain names, computer names, and Internet Protocol (IP) addresses for sensitive network resources, is obtained by an attacker. An attacker commonly begins an attack by using this DNS data to diagram, or "footprint," a network. DNS domain names and computer names usually indicate the function or location of a domain or computer to help users remember and identify domains and computers more easily. An attacker takes advantage of this same DNS naming principle to learn the function or location of domains and computers in the network.
- **Denial-of-service attack.** A scenario in which an attacker attempts to deny the availability of network services by flooding one or more DNS servers in the network with recursive queries. As a DNS server is flooded with queries, its CPU usage eventually reaches its maximum, and the DNS Server service becomes unavailable. Without a fully operating DNS server on the network, network services that use DNS are unavailable to network users.
- **Data modification.** An attempt by an attacker that has footprinted a network by using DNS to use valid IP addresses in IP packets that the attacker has created. This gives these packets the appearance of coming from a valid IP address in the network. This process is commonly called IP "spoofing." With a valid IP address — that is, an IP address within the IP address range of a subnet — the attacker can gain access to the network and destroy data or conduct other attacks.
- **Redirection.** A scenario in which an attacker is able to redirect queries for DNS names to servers that are under the control of the attacker. One method of redirection involves an attempt to pollute the DNS cache of a DNS server with erroneous DNS data that may direct future queries to servers that are under the control of the attacker. For example, if a query is made originally for sales.wingtiptoys.com and a referral answer provides a record for a domain name that the attacker has outside the wingtiptoys.com domain, the DNS server uses the cached data for the attacker's domain to resolve a query for that name. Redirection can occur whenever an attacker has writable access to DNS data, for example, in a scenario that includes dynamic updates that are not secure.

**Mitigating DNS Security Threats**

The following sections explain three levels of DNS security that you can apply to your current DNS configuration. You can use these three levels of security to increase the DNS security of your organization.

**Low-Level Security**

Low-level security is a standard DNS deployment without any security precautions configured. You should deploy this level of DNS security only in network environments where there is no concern for the integrity of your DNS data or in a private network where there is no threat of external connectivity:

- The DNS infrastructure of your organization is fully exposed to the Internet.
- Standard DNS resolution is performed by all DNS servers in your network.
- All DNS servers are configured with root hints pointing to the root servers for the Internet.
- All DNS servers permit zone transfers to any server.

- All DNS servers are configured to listen on all of their IP addresses.
- Cache pollution prevention is disabled on all DNS servers.
- Dynamic update is allowed for all DNS zones.
- User Datagram Protocol (UDP) and TCP/IP port 53 is open on the firewall for your network for both source and destination addresses.

**Medium-Level Security**
Medium-level security uses the DNS security features that are available without running DNS servers on domain controllers and storing DNS zones in Active Directory:

- The DNS infrastructure of your organization has limited exposure to the Internet.
- All DNS servers are configured to use forwarders to point to a specific list of internal DNS servers when they cannot resolve names locally.
- All DNS servers limit zone transfers to servers that are listed in the name server (NS) resource records in their zones.
- DNS servers are configured to listen on specified IP addresses.
- Cache pollution prevention is enabled on all DNS servers.
- Dynamic update that is not secure is not allowed for any DNS zones.
- Internal DNS servers communicate with external DNS servers through a firewall with a limited list of allowed source addresses and destination addresses.
- External DNS servers in front of the firewall are configured with root hints that point to the root servers for the Internet.
- All Internet name resolution is performed by using proxy servers and gateways.

**High-Level Security**
High-level security uses the same configuration as medium-level security. It also uses the security features that are available when the DNS Server service is running on a domain controller and DNS zones are stored in Active Directory. In addition, high-level security completely eliminates DNS communication with the Internet. This is not a typical configuration, but it is recommended whenever Internet connectivity is not required:

- The DNS infrastructure of your organization has no Internet communication by means of internal DNS servers.
- Your network uses an internal DNS root and namespace, where all authority for DNS zones is internal.
- DNS servers that are configured with forwarders use internal DNS server IP addresses only.
- All DNS servers limit zone transfers to specified IP addresses.
- DNS servers are configured to listen on specified IP addresses.
- Cache pollution prevention is enabled on all DNS servers.
- Internal DNS servers are configured with root hints that point to the internal DNS servers that host the root zone for your internal namespace.
- All DNS servers are running on domain controllers. A discretionary access control list (DACL) is configured on the DNS Server service to allow only specific individuals to perform administrative tasks on the DNS server.

- All DNS zones are stored in Active Directory. A DACL is configured to allow only specific individuals to create, delete, or modify DNS zones.

- DACLs are configured on DNS resource records to allow only specific individuals to create, delete, or modify DNS data.

- Secure dynamic update is configured for DNS zones except the top-level zones and root zones, which do not allow dynamic updates at all.

**Securing DNS**

The following tasks for securing DNS are described in this objective:

- Securing Domain Name System Zones

- Securing the Domain Name System Server Service

- Securing Domain Name System Clients

# Securing Domain Name System Zones

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To secure the Domain Name System (DNS) zones in your environment, use the following guidelines:

- **Configure secure dynamic updates.** By default, the **Dynamic updates** option is not set to allow dynamic updates. This is the most secure setting because it prevents an attacker from updating DNS zones. However, this setting prevents you from taking advantage of the benefits to administration that dynamic updates provide. To make it possible for computers to update DNS data securely, store DNS zones in Active Directory, and use the secure dynamic update feature. Secure dynamic update restricts DNS zone updates to only the following:

    - Computers that are authenticated and joined to the Active Directory domain where the DNS server is located

    - The specific security settings that are defined in the access control lists (ACLs) for the DNS zone

- Restrict zone transfers. By default, the DNS Server service allows zone information to be transferred only to servers that are listed in the name server (NS) resource records of a zone. This is a secure configuration. However, for increased security this configuration should be changed to enable the option to allow zone transfers to specified Internet Protocol (IP) addresses. Changing this configuration to allow zone transfers to any server at all may expose your DNS data to an attacker who is attempting to footprint your network.

- Understand the compromise involved in zone delegation. When you decide whether to delegate DNS domain names to zones that are hosted on DNS servers that are administered separately, it is important to consider the security implications of giving the ability to administer the DNS data for your network to multiple individuals. DNS zone delegation involves a compromise between the security benefits of having a single authoritative DNS server for all DNS data and the administrative benefits of distributing responsibility for your DNS namespace to separate administrators. This issue is very important when you delegate the top-level domains of a private DNS namespace, because those domains contain very sensitive DNS data.

For more information about planning DNS zones, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

**Task requirements**

To begin this task, perform the following requirements:

- Install Dnscmd.

To complete this task, perform the following procedures:

1. Enable secure dynamic updates

2. Modify DNS zone transfer settings

3. Managing DNS Resource Records

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Enable secure dynamic updates

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to allow only secure dynamic updates for a zone. Secure dynamic update is supported only for Active Directory–integrated zones. If the zone type is configured differently, you must change the zone type and directory-integrate the zone before securing it for Domain Name System (DNS) dynamic updates.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Enabling secure dynamic updates**

- Using the Windows interface
- Using the command line

### To enable secure dynamic updates using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, verify that the zone type is **Active Directory-Integrated**.
4. In **Dynamic updates**, click **Secure only**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To enable secure dynamic updates using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config** {*ZoneName*|**..AllZones**} **/AllowUpdate 2**

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName*\|**..AllZones** | Required. Specifies the fully qualified domain name (FQDN) of the zone. To configure all zones that are hosted on the specified DNS server to allow dynamic updates, type **..AllZones**. |
| **2** | Required. Configures the server to allow secure dynamic updates. If you exclude the **2**, the zone is set to perform standard dynamic updates only. |

# Modify DNS zone transfer settings

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to change Domain Name System (DNS) zone transfer settings. To improve the security of your DNS infrastructure, zone transfers should be allowed only for either the DNS servers in the name server (NS) resource records for a zone or for specified DNS servers. If you allow any DNS server to perform a zone transfer, you are allowing internal network information to be transferred to any host that can contact your DNS server.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Modifying DNS zone transfer settings**

- Using the Windows interface
- Using the command line

### To modify DNS zone transfer settings using the Windows interface

1. Open the DNS snap-in.
2. Right-click a DNS zone, and then click **Properties**.
3. On the **Zone Transfers** tab, do one of the following:
   - To disable zone transfers, clear the **Allow zone transfers** check box.
   - To allow zone transfers, select the **Allow zone transfers** check box.
4. If you allowed zone transfers, do one of the following:
   - To allow zone transfers to any server, click **To any server**.
   - To allow zone transfers only to the DNS servers that are listed on the **Name Servers** tab, click **Only to servers listed on the Name Servers tab**.
   - To allow zone transfers only to specific DNS servers, click **Only to the following servers**, and then add the Internet Protocol (IP) address of one or more DNS servers.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To modify DNS zone transfer settings using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ZoneResetSecondaries** *ZoneName* {**/NoXfr**|**/NonSecure**|**/SecureNs**|**/SecureList**[*SecondaryIPAddress*...]}

| Value | Description |
| --- | --- |

| | |
|---|---|
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of zone. |
| **/NoXfr** | Disables zone transfers for the zone. |
| **/NonSecure** | Permits zone transfers to any DNS server. |
| **/SecureNs** | Permits zone transfers only to DNS servers that are listed in the zone using NS resource records. |
| **/SecureList** | Permits zone transfers only to DNS servers that are specified by *SecondaryIPAddress*. |
| *SecondaryIPAddress* | Required if **/SecureList** is specified. A list of one or more IP addresses for DNS servers that are permitted to obtain zone transfers. |

# Delegating a Zone

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain Name System (DNS) provides the option of dividing up the namespace into one or more zones, which can then be stored, distributed, and replicated to other DNS servers. When deciding whether to divide your DNS namespace to make additional zones, consider the following reasons to use additional zones:

- You want to delegate management of part of your DNS namespace to another location or department in your organization.
- You want to divide one large zone into smaller zones for distributing traffic loads among multiple servers, improving DNS name resolution performance, or creating a more fault-tolerant DNS environment.
- You want to extend the namespace by adding numerous subdomains at once, for example, to accommodate the opening of a new branch or site.

If, for any of these reasons, your network can benefit from delegating zones, it may make sense to restructure your namespace by adding additional zones. When choosing how to structure zones, use a plan that reflects the structure of your organization.

When you delegate zones within your namespace, be aware that for each new zone you create, you will need delegation records in other zones that point to the authoritative DNS servers for the new zone. This is necessary both to transfer authority and to provide correct referral to other DNS servers and clients of the new servers that are being made authoritative for the new zone.

When a standard primary zone is first created, it is stored as a text file that contains all resource record information on a single DNS server. This server acts as the primary master for the zone. Zone information can be replicated to other DNS servers to improve fault tolerance and server performance.

When you structure your zones, there are several good reasons to use additional DNS servers for zone replication:

- Added DNS servers provide zone redundancy, enabling DNS names in the zone to be resolved for clients if a primary server for the zone stops responding.
- Added DNS servers can be placed so as to reduce DNS network traffic. For example, adding a DNS server to the opposing side of a low-speed, wide area network (WAN) link can be useful in managing and reducing network traffic.
- Additional secondary servers can be used to reduce loads on a primary server for a zone.

For more information about planning DNS zones, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

**Task requirements**

To begin this task, perform the following requirements:

- Install Dnscmd.
- Install Nslookup.

To complete this task, perform one of the following procedures:

- Create a new zone delegation
- Verify a zone delegation

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Create a new zone delegation

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to create a new zone delegation. All domains (or subdomains) that appear as part of the applicable zone delegation must be created in the current zone before you perform delegation as described in this procedure. As necessary, use the DNS snap-in to first add domains to the zone before you perform this procedure. You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Creating a new zone delegation**

- Using the Windows interface
- Using the command line

### To create a new zone delegation using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable subdomain, and then click **New Delegation**.
3. Follow the instructions in the New Delegation Wizard to finish creating the new delegated domain.

| **Note** |
|---|
| To open DNS, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

### To create a new zone delegation using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/RecordAdd** *ZoneName NodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **NS** {*HostName*|*FQDN*}

| Value | Description |
|---|---|
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the start-of-authority (SOA) record is added. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is updated or removed manually. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. (The default TTL is defined in the SOA resource record). |
| **NS** | Required. Specifies that you are adding a name server (NS) resource record to the zone that is specified in *ZoneName*. |

| | |
|---|---|
| *HostName\|FQDN* | Required. Specifies the host name or FQDN of the new authoritative server. |

# Verify a zone delegation

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administrative credentials**

You do not need administrative credentials to perform this task. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

**To verify a zone delegation**

1. At a command prompt, type the following command, and then press ENTER:

   **nslookup** *RootServerIpAddress*

2. Type the following command, and then press ENTER:

   **nslookup**

3. At the next prompt, type the following command, and then press ENTER:

   **set norecurse**

4. At the next prompt, type the following command, and then press ENTER:

   **set q=NS**

5. Type the fully qualified domain name (FQDN) for the failed name.

   Use the trailing period (**.**) when you type the name. If zone delegations are set correctly, a list of name server (NS) resource records for delegated servers is returned in the response.

6. If the NS query response contains no names or Internet Protocol (IP) addresses for delegated servers, type **q=ns**, and then query again using the FQDN for the parent zone of the failed name.

   For example, if the failed name that you used in the previous step was sales.wingtiptoys.com, query for wingtiptoys.com.

7. If the response contains NS resource records, but no host address (A) resource records, type **set recurse**, and then query individually for any of the A resource records of the servers that are listed in the NS resource records.

   If, for each NS resource record that you encounter in a zone, you do not find at least one valid IP address in an A resource record, you have a broken delegation.

8. Either fix the broken delegation or retry the delegation test that is described in the previous step and use a different IP address.

   If more than one A resource record or IP address is found, use it to repeat the delegation test described in the previous step. To fix a delegation, add or update an A resource record in the parent zone with a valid IP address for a correct DNS server for the delegated zone.

| Value | Description |
|---|---|
| *RootServerIpAddress* | The IP address of a valid root server for your network. |
| **set norecursion** | Instructs the root server to not perform recursion on your query. |
| **set q=NS** | Sends the query for NS resource records to the root server. |

# Securing the Domain Name System Server Service

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To secure the Domain Name System (DNS) servers in your network, use the following guidelines:

- Limit the Internet Protocol (IP) addresses that the DNS Server service listens on to the IP address that is used by its DNS clients as their preferred DNS server. By default, a DNS Server service that is running on a multihomed computer is configured to listen for DNS queries on all its IP addresses.

- Leave the **Secure cache against pollution** option enabled. By default, the DNS Server service is secured from cache pollution, which occurs when DNS query responses contain nonauthoritative or malicious data. The **Secure cache against pollution** option prevents an attacker from polluting the cache of a DNS server with resource records that were not requested by the DNS server. Changing this default setting reduces the integrity of the responses that are provided by DNS Server service.

- Disable recursion. By default, recursion is not disabled for the DNS Server service. This enables the DNS server to perform recursive queries on behalf of its DNS clients and the DNS servers that have forwarded DNS client queries to it. Recursion can be used by attackers to deny the DNS Server service. Therefore, if a DNS server in your network is not intended to receive recursive queries, it should be disabled.

- If you have an internal DNS root in your DNS infrastructure, configure the root hints of internal DNS servers to point only to the DNS servers that host your root domain, not to the DNS servers that host the Internet root domain. This prevents your internal DNS servers from sending private information over the Internet when they resolve names.

For more information about planning DNS, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

**Task requirements**

To begin this task, perform the following requirements:

- Install Dnscmd.

To complete this task, perform the following procedures:

1. Restrict the DNS server to listen on selected IP addresses

2. Secure the server cache against names pollution.

3. Disable recursion

4. Update root hints

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Restrict the DNS server to listen on selected IP addresses

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to restrict the DNS Server service to listen only on selected Internet Protocol (IP) addresses. By default, the DNS Server service listens for Domain Name System (DNS) message communications on all configured IP addresses for the server computer. Restricting the DNS Server service to listen only on specific IP addresses is an effective security measure because only hosts on the same network subnet — or hosts with a router that connects them to that same segment — have access to the server.

| **Note** |
| --- |
| Server IP addresses that you add with this procedure must be managed statically. If you later change or remove addresses specified here from TCP/IP configurations that are maintained at this server, update this list accordingly. |

After you update or revise the list of restricted interfaces, stop and restart the DNS server to apply the new list.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Restricting the DNS server to listen on selected IP addresses**

- Using the Windows interface
- Using the command line

**To restrict the DNS server to listen on selected IP addresses using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. On the **Interfaces** tab, click **Only the following IP addresses**.
5. In **IP address**, type an IP address for the DNS server to be enabled for use, and then click **Add**.
6. Repeat the previous step as needed to specify other server IP addresses to be enabled for use by this DNS server.

   If you want to remove an IP address from the list, click the IP address, and then click **Remove**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To restrict the DNS server to listen on selected IP addresses using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ResetListenAddresses** [*ListenAddress...*]

| Value | Description |
| --- | --- |

| | |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/ResetListenAddresses** | Required. Resets the IP addresses of the interfaces on which the DNS server listens. |
| *ListenAddress...* | Specifies one or more IP addresses for the interfaces on which you want the DNS server to listen. By default, the DNS Server service listens for DNS message communications on all configured IP addresses for the server computer. |

# Secure the server cache against names pollution

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to secure the server cache against names pollution. This setting is enabled by default.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Secure the server cache against names pollution**

**To secure the server cache against names pollution**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, select the **Secure cache against pollution** check box, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Disable recursion

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to disable recursion on the Domain Name System (DNS) server.

| **Note** |
| --- |
| If you disable recursion on the DNS server, you will not be able to use forwarders on the same server. |

### Administrative credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

#### Disabling recursion

- Using the Windows interface
- Using the command line

### To disable recursion using the Windows interface

1. Open the DNS snap-in.
2. In the console tree, right-click the applicable DNS server, then click **Properties**.
3. Click the **Advanced** tab.
4. In **Server options**, select the **Disable recursion (also disables forwarders)** check box, and then click **OK**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, , point to **Administrative Tools**, and then click **DNS**. |

### To disable recursion using the command line

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/Config /NoRecursion** {**1**|**0**}

| Value | Description |
| --- | --- |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the Internet Protocol (IP) address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/NoRecursion** | Required. Disables recursion. |
| {**1**\|**0**} | Required. To disable recursion, type **1** (off). To enable recursion, type **0** (on). By default, recursion is enabled. |

# Update root hints

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to update root hints on the Domain Name System (DNS) server.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Updating root hints**

**To update root hints**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Root Hints** tab.
5. Modify server root hints as follows:
   - To add a root server to the list, click **Add**, and then specify the name and Internet Protocol (IP) address of the server to be added to the list.
   - To modify a root server in the list, click **Edit**, and then specify the name and IP address of the server to be modified in the list.
   - To remove a root server from the list, select it in the list, and then click **Remove**.
   - To copy root hints from a DNS server, click **Copy from Server**, and then specify the IP address of the DNS server from which you want to copy a list of root servers to use in resolving queries. These root hints will not overwrite any existing root hints.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

# Securing Domain Name System Clients

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following Domain Name System (DNS) client considerations have security implications for DNS clients in a DNS infrastructure:

- Whenever possible, specify static Internet Protocol (IP) addresses for the preferred and alternate DNS servers that are to be used by a DNS client. If a DNS client is configured to obtain its DNS server addresses automatically, it will obtain them from a Dynamic Host Configuration Protocol (DHCP) server. While this method of obtaining DNS server addresses is secure, it is only as secure as the DHCP server. By configuring DNS clients with static IP addresses for the preferred and alternate DNS servers, you eliminate one possible avenue of attack.

- Control which DNS clients have access to the DNS server. If a DNS server is configured to listen only on specific IP addresses, only DNS clients that are configured to use these IP addresses as preferred and alternate DNS servers will contact the DNS server.

For more information about planning DNS, see Deploying Domain Name System (DNS) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45677).

**Task requirements**

To begin this task, perform the following requirements:

- Install Dnscmd.

To complete this task, perform the following procedures:

1. Configure DNS settings in Network Connections.

2. Restrict the DNS server to listen on selected IP addresses.

**See Also**
**Other Resources**
Deploying Domain Name System (DNS)

# Configure DNS settings in Network Connections

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use this procedure to configure TCP/IP to use Domain Name System (DNS).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer.

**To configure DNS settings in Network Connections**

1. Open Network Connections.
2. Right-click the network connection that you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. If you want to obtain DNS server addresses from a DHCP server, click **Obtain DNS server address automatically**.
5. If you want to configure DNS server addresses manually, click **Use the following DNS server addresses**, and in **Preferred DNS server** and **Alternate DNS server**, type the Internet Protocol (IP) addresses of the preferred DNS server and alternate DNS server.
6. To configure advanced DNS properties, click **Advanced**, click the **DNS** tab, and then do one or more of the following:
   - To configure an additional DNS server IP address:

     a. Under **DNS server addresses**, in order of use, click **Add**.

     b. In **DNS server**, type the IP address of the DNS server, and then click **Add**.

   - To resolve an unqualified name by appending the primary DNS suffix and the DNS suffix of each connection (if configured), click **Append primary and connection specific DNS suffixes**. If you also want to search the parent suffixes of the primary DNS suffix up to the second-level domain, select the **Append parent suffixes of the primary DNS suffix** check box.
   - To resolve an unqualified name by appending the suffixes from a list of configured suffixes, click **Append these DNS suffixes (in order)**, and then click **Add** to add suffixes to the list.
   - To use a DNS dynamic update to register the IP addresses of this connection and the primary domain name of the computer, select the **Register this connection's addresses in DNS** check box. This option is enabled by default. The primary domain name of the computer is the primary DNS suffix appended to the computer name, and it can be viewed as the full computer name on the **Computer Name** tab (which is available in **System** in **Control Panel**).
   - To use a DNS dynamic update to register the IP addresses and the connection-specific domain name of this connection, select the **Use this connection's DNS suffix in DNS registration** check box. This option is disabled by default. The connection-specific domain name of this connection is the DNS suffix for this connection appended to the computer name.
   - To completely disable DNS dynamic update for all names on the computer, clear the **Register this connection's addresses in DNS** and **Use this connection's DNS suffix in DNS registration** check boxes for all connections in Network Connections.

**Note**

To open Network Connections, click **Start**, point to **Control Panel**, and then click **Network Connections**.

# Restrict the DNS server to listen on selected IP addresses

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the following procedure to restrict the DNS Server service to listen only on selected Internet Protocol (IP) addresses. By default, the DNS Server service listens for Domain Name System (DNS) message communications on all configured IP addresses for the server computer. Restricting the DNS Server service to listen only on specific IP addresses is an effective security measure because only hosts on the same network subnet — or hosts with a router that connects them to that same segment — have access to the server.

| **Note** |
| --- |
| Server IP addresses that you add with this procedure must be managed statically. If you later change or remove addresses specified here from TCP/IP configurations that are maintained at this server, update this list accordingly. |

After you update or revise the list of restricted interfaces, stop and restart the DNS server to apply the new list.

You can perform this procedure by using the DNS snap-in or by using the Dnscmd command-line tool.

**Administrative credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

**Restricting the DNS server to listen on selected IP addresses**

- Using the Windows interface
- Using the command line

**To restrict the DNS server to listen on selected IP addresses using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, click the applicable DNS server.

   Where?

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. On the **Interfaces** tab, click **Only the following IP addresses**.
5. In **IP address**, type an IP address for the DNS server to be enabled for use, and then click **Add**.
6. Repeat the previous step as needed to specify other server IP addresses to be enabled for use by this DNS server.

   If you want to remove an IP address from the list, click the IP address, and then click **Remove**.

| **Note** |
| --- |
| To open the DNS snap-in, click **Start**, point to **Administrative Tools**, and then click **DNS**. |

**To restrict the DNS server to listen on selected IP addresses using the command line**

- At a command prompt, type the following command, and then press ENTER:

  **dnscmd** *ServerName* **/ResetListenAddresses** [*ListenAddress...*]

| Value | Description |
| --- | --- |

| | |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/ResetListenAddresses** | Required. Resets the IP addresses of the interfaces on which the DNS server listens. |
| *ListenAddress...* | Specifies one or more IP addresses for the interfaces on which you want the DNS server to listen. By default, the DNS Server service listens for DNS message communications on all configured IP addresses for the server computer. |

# Troubleshooting Domain Name System

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides troubleshooting information for Domain Name System (DNS) in Windows Server 2003 SP 1. It is designed to help you identify and resolve problems that may be related to DNS.

**In this guide**


- Introduction to Troubleshooting Domain Name System
- Verifying Computer Settings for Domain Name System
- Configuring a Computer for Troubleshooting Domain Name System
- Troubleshooting Domain Name System Problems
- Additional Resources for Domain Name System



**Acknowledgments**

Produced by: Microsoft Windows Server User Assistance team

Project Writer: Andrea Weiss

Project Editor: Jim Becker

# Introduction to Troubleshooting Domain Name System

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide explains how to troubleshoot Domain Name System (DNS). If you are not familiar with this guide, review the following sections of this introduction.

**When to Use This Guide**
You should use this guide when:

- You have a problem that you believe is related to DNS, but you do not know how to resolve it.
- You configure DNS settings, but DNS does not behave the way that you anticipate.
- A program is not working properly, and you believe that DNS is causing the problem.

Do not use this guide to find out how to perform a task, such as configuring a stub zone or installing DNS. Information about how to perform tasks and configure settings can be found in Administering DNS Operations.

This guide assumes that you have a basic understanding of what DNS is, how it works, and why your organization uses it for name resolution. You should also have a thorough understanding of how DNS is deployed and managed in your organization. This includes an understanding of the mechanism that your organization uses to configure and manage DNS settings.

**How This Guide Is Organized**
This guide is divided into four sections. Each section addresses a type of problem, ranging from less complex to more complex.

Verifying Computer Settings for Domain Name System

This section provides a list of prerequisites and settings that must be verified before you troubleshoot. Read this section first.

Configuring a Computer for Troubleshooting Domain Name System

This section describes how to configure your computer for troubleshooting.

Troubleshooting Domain Name System Problems

This section provides step-by-step diagnostic procedures and possible solutions that help you identify and fix DNS problems.

Additional Resources for Domain Name System

This section provides additional resources that are related to DNS.

# Verifying Computer Settings for Domain Name System

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin troubleshooting, verify that your computer is configured properly and that Domain Name System (DNS) is set up and running properly.

**Settings to Verify Before You Troubleshoot**
Verify all of the following items:

Make sure you have administrative rights on the computer that you are troubleshooting

Install all critical updates and security updates for Windows Server 2003

Verify DNS client settings

**Make sure you have administrative rights on the computer that you are troubleshooting.**
You cannot modify DNS settings unless you are a member of the Administrators group on the computer that you are troubleshooting.

### To verify that you are a member of the Administrators group that you are troubleshooting

1. Open the Computer Management snap-in.
2. In the console tree, double-click **Local Users and Groups**, and then click **Groups**.
3. In the details pane, double-click **Administrators** and verify that your account name or a group to which your account is a member appears in the **Members** list.

**Install all critical updates and security updates for Windows Server 2003.**
Some updates might be required for DNS to function properly.

### To verify that you have all critical updates and security updates for Windows Server 2003

- Click **Start**, click **Windows Update**, and then follow the instructions that appear on your screen.

**Verify DNS client settings.**
Verify that the DNS client does not have an external DNS server, such as a DNS server from an Internet service provider (ISP), in its TCP/IP configuration. In most cases, the client should not use a DNS server from an ISP as either the preferred or alternate DNS server because the DNS server from the ISP is unable to resolve internal names. Using a DNS server from an ISP in the TCP/IP configuration of a client can also cause problems with conflicting internal and external namespaces.

### To verify DNS client settings

1. Log on to the computer with your Administrator account.
2. Click **Start**, click **Control Panel**, and then double-click **Network Connections.**
3. In **Network Connections**, right-click the local area connection that you want, and then click **Properties**.
4. In **Local Area Connection Properties**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
5. If **Obtain an IP address automatically** is selected, type the following at a command prompt:

   **ipconfig /all**

6. Review the DNS server settings and verify that they are correct.

# Configuring a Computer for Troubleshooting Domain Name System

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can use advanced troubleshooting techniques to identify and fix Domain Name System (DNS) problems, you need to configure your computer for troubleshooting. In addition, you need a basic understanding of troubleshooting concepts, procedures, and tools.

**Configuration Tasks for Troubleshooting**
To configure your computer for troubleshooting, perform the following tasks:

Install Windows Server 2003 SP1

Install Windows Support Tools

Install Network Monitor

Enable DNS Debug Logging

**Install Windows Server 2003 SP1**
If possible, upgrade your DNS servers with Windows Server 2003 Service Pack 1 (SP1). To install this service pack, go to the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=9999) and follow the instructions for downloading the service pack.

**Install Windows Support Tools**
For improved diagnostic support, install the Windows Support Tools that ship with Windows Server 2003 SP1. The SP1 version of Windows Support Tools includes enhanced versions of the Dcdiag.exe tool. The Dcdiag.exe command-line tool now provides new reporting on the overall health of replication with respect to Active Directory security as well as new DNS diagnostic tests.

Make sure that the SP1 version of Windows Support Tools is installed on all DNS servers running Windows Server 2003 with SP1.

**Options for Running SP1 Windows Support Tools**
You can run Windows Support Tools that ship with Windows Server 2003 SP1 on computers running the following operating systems:

- Windows Server 2003 with SP1
- Windows Server 2003 without SP1

You can also run Dcdiag.exe on computers running Windows XP Professional, Windows XP Professional with SP1, or Windows XP Professional with Service Pack 2 (SP2). Options for other tools vary by tool.

**Options for Installing SP1 Windows Support Tools**
The SP1 version of Windows Support Tools can be installed as an .msi package only on computers running Windows Server 2003 with SP1. To run Dcdiag from computers running Windows Server 2003 without SP1 or from computers running Windows XP Professional, you must copy the respective executable files to those computers.

### Requirements

- Administrative credentials: To complete this procedure, you must be a member of the Builtin Administrators group.
- Operating system: Windows Server 2003 with SP1. You cannot use Suptools.msi to install the SP1 version of Windows Support Tools on a computer that is not running Windows Server 2003 with SP1.

### To install Windows Support Tools

1. Insert the Windows CD into your CD-ROM drive.

2. If you are prompted to reinstall Windows, click **No**.
3. When the **Welcome** screen appears, click **Perform additional tasks**, and then click **Browse this CD**.
4. Go to the \Support\Tools folder. For complete setup information, see the Readme.htm file in this folder.
5. Double-click **suptools.msi**.
6. Follow the instructions that appear on your screen.

**Install Network Monitor**
You can use Network Monitor to troubleshoot connectivity issues by tracing network traffic between computers. For information about installing and using Network Monitor, see Network Monitor on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42987).

**Enable DNS Debug Logging**
DNS debug logging creates a Dns.log file that contains debug logging activity. By default, this file is located in the C:\Windows\System32\DNS folder. Using debug logging options slows DNS server performance; therefore, all debug logging options are disabled by default. You can use the following procedure to enable DNS debug logging.

### Administrative credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **Run as** command to perform this procedure.

### To enable DNS debug logging

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **DNS**.
2. In the console tree, right-click the applicable DNS server, then click **Properties**.
3. Click the **Debug Logging** tab.
4. Select **Log packets for debugging**, and then select the events that you want the DNS server to record for debug logging

### Notes

- To set the debug logging options, you must first select **Log packets for debugging**.
- To obtain useful debug logging output, select an option under **Packet direction**, an option under **Transport protocol**, and at least one more option.
- In addition to selecting events for the DNS debug log file, you can specify the file name, location, and maximum file size for the file. In most cases, the default selections are adequate. You may want to limit the traffic that the logging captures. If you want to limit the logging traffic to traffic between your server and a specific DNS server, select the **Filter packets by IP address** check box, and then click **Filter** to add the appropriate IP addresses.

# Troubleshooting Domain Name System Problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Domain Name System (DNS) problems can prevent users from accessing the Internet or other computers on your intranet. This section addresses some of the most common DNS-related problems that you might experience.

From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- Secure dynamic updates fail
- DNS server resolves some Internet names incorrectly
- DNS client fails to resolve name
- Zone transfers from a secondary DNS server fail
- Dynamic updates for host records fail

If you suspect that domain controllers are having difficulty resolving DNS names, see Perform DNS Health Check (http://go.microsoft.com/fwlink/?LinkId=111844).

# Secure dynamic updates fail

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem occurs when attempts by hosts fail to update a Domain Name System (DNS) zone that is configured for secure dynamic updates.

**Cause**
The host might not be configured properly to allow secure dynamic DNS updates. It might be configured to use an external DNS server, or it might be experiencing other DNS configuration problems.

**Solution**
First, try to solve this problem by using the troubleshooting information in Dynamic updates for host records fail.

Hosts that perform secure dynamic updates should be members of a Windows 2000 Server or Windows Server 2003 domain, and they should be in a domain that is in the same forest as the DNS server.

Verify that there is no problem with the machine account of the host that is attempting the update. Determine whether other hosts successfully perform secure dynamic updates. If the problem is occurring on only one host, try removing the host from the domain and then rejoining it to the domain.

Verify that a record does not already exist with the same name. By default, records that are created by one host cannot be modified or removed by a different host. If there is an existing record with the same name, delete the existing record and have the host attempt to register again.

**To initiate a dynamic update for host and PTR records**

- At a command prompt, type the following command, and then press ENTER:

  **ipconfig /registerdns**

# DNS server resolves some Internet names incorrectly

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The symptom for this problem is that the Domain Name System (DNS) server stops resolving some Internet names correctly. It might give inaccurate responses or fail to resolve Internet names completely, while other Internet names resolve accurately. This problem may get worse over time. Clearing the DNS server cache, restarting the DNS service, or restarting the DNS server resolves the problem temporarily, but the problem may continue to occur.

**Cause**

The DNS server might be experiencing cache pollution. This is caused by the DNS server receiving and caching an inaccurate start of authority (SOA) record for a portion of the Internet namespace. For example:

1. A DNS server queries a specific name server to resolve host.contoso.com.

2. The DNS server that is being queried gives a response. Along with the response — in the authority section — it gives an incorrect record for the .com namespace.

3. This record is cached and all new queries that contain the top-level domain name of com will not resolve, or they resolve to an incorrect Internet Protocol (IP) address.

**Solution**

Configure the DNS server for protection against cache pollution by completing the following procedure.

**To configure the DNS server for protection against cache pollution**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **DNS**.

2. In the console tree, right-click the DNS server that you want to protect against cache pollution, click **Properties**, and then click the **Advanced** tab.

3. Select the **Secure cache against pollution** check box, and then click **OK**.

# DNS client fails to resolve name

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem may occur when a user is trying to access another computer on the intranet or the Internet. A user typically receives a Domain Name System (DNS) error from the operating system or the browser.

**Cause**

Several problems can cause DNS name resolution to fail. If you have reviewed the topics for other problems in Troubleshooting Domain Name System Problems and they do not seem to be the cause of the problem, DNS settings might be configured incorrectly on the DNS client.

**Solution**

Verify that the client does not have an external DNS server, such as a DNS server from an Internet service provider (ISP), in its TCP/IP configuration. In most cases, the client should not use a DNS server from an ISP as either the preferred or alternate DNS server, because the DNS server at the ISP is unable to resolve internal names. Using a DNS server from an ISP in a client's TCP/IP configuration can also cause problems with conflicting internal and external namespaces.

**To verify DNS configuration in TCP/IP settings**

1. Log on to the DNS client computer with the Administrator account.
2. Click **Start**, click **Control Panel**, and then double-click **Network Connections**.
3. In **Network and Dial-up Connections**, right-click the local area connection that you want, and then click **Properties**.
4. In **Local Area Network Connection Properties**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
5. Ensure that the appropriate DNS server IP addresses are configured in **Preferred DNS server** and **Alternate DNS server**. If **Obtain an IP address automatically** is selected, click the **Alternate Configuration** tab, and then review all the IP settings that are configured there.
6. Type the following at a command prompt, and then press ENTER:

   **ipconfig /all**
7. Review the DNS server settings, and verify that they are correct.
   a. If the DNS server settings are not correct, ensure the appropriate settings are configured on the Dynamic Host Configuration Protocol (DHCP) server.
   b. If your computer has an IP address that begins with 169.254, it is not obtaining an IP address from a DHCP server and likely does not have **Alternate Configuration** enabled. In this case, diagnose the issue with the DHCP server or set an appropriate static IP address either directly or as an alternate configuration.

Next, use the following procedure to verify that the name can be resolved by the DNS server.

**To verify name resolution**

- At a command prompt, type the following command, and then press ENTER:

  **nslookup** *host_name server_IP_address*

Substitute the actual host name that you are trying to resolve for *host_name* and the IP address of the DNS server for *server_IP_address*. For example, if the host name that you are trying to resolve has a fully qualified domain name (FQDN) of server5.contoso.com and the DNS server's IP address is 192.168.0.200, type the following command, and then press ENTER:

**nslookup server5 192.168.0.200**

You can also try using the FQDN:

**nslookup server5.contoso.com 192.168.0.200**

If the host name alone does not resolve, but the FQDN does resolve, confirm that the primary or connection-specific DNS suffix

is configured correctly. You can use the following procedure to add a DNS suffix search list.

**To add a DNS suffix search list**

1. Click **Start**, right-click **My Network Places**, and then click **Properties**.
2. Right-click **Local Area Connection**, and then click **Properties**.
3. Double-click **Internet Protocol (TCP/IP)**, and then click **Advanced**.
4. Click the **DNS** tab, and then click **Append these DNS suffixes (in order)**.
5. Click **Add**, type the domain suffix of the desired domain, and then click **Add**.

If both Nslookup commands fail to resolve the name, the problem is likely with the DNS server records or configuration, or it may be the result of a connectivity issue between the DNS client and DNS server, such as a firewall blocking DNS queries (which are typically offered on TCP port 53). You can use the Portqry tool to test network connectivity between two computers. For more information about downloading and using Portqry, see article 832919 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=111855).

If the connection is not successful, look for a firewall on the DNS client, DNS server, or somewhere between the two that could cause the connection failure. If you are diagnosing a connection failure between two computers, you can try using the Portqry tool to test it. For more information about downloading and using Portqry, see article 832919 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=111855).

If Portqry fails, there may be a connectivity issue between the computers. You can try using the Tracert or Pathping tools to find out where the failure exists. However, these tools are not always reliable because many hosts and servers disable the Internet Control Message Protocol (ICMP) echo functionality on which these tools depend. However, using one of these tools may help locate the source of a problem. For more information about using Tracert, see article 314868 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=111861). For information about using Pathping, see Pathping (http://go.microsoft.com/fwlink/?LinkId=111864)

| **Note** |
| --- |
| If you locate and correct any issues on the DNS server, it is probable that your client computer cached the incorrect information, and your name resolution queries may still fail. To resolve this issue, clear the DNS client resolve cache on the client computer. To clear the cache, type the following command at a command prompt, and then press ENTER:<br><br>**Ipconfig /flushdns** |

For additional information about troubleshooting the DNS client configuration, see Validate DNS Client Settings (http://go.microsoft.com/fwlink/?LinkId=111865).

# Zone transfers from a secondary DNS server fail

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the DNS management console, the secondary zone has a red X in the right pane, along with an error message that reads as follows:

```
The DNS server encountered a problem while attempting to load the zone. The transfer of
zone data from the master server failed.
Correct the problem then either press F5, or on the Action menu, click Refresh.
For more information about troubleshooting DNS zone problems, see Help.
```

**Cause**

The primary DNS server might not be configured properly to allow zone transfers from the secondary DNS server.

**Solution**

Use the following procedure to verify that the primary DNS server is configured to allow zone transfers from the secondary DNS server.

**To verify that the primary DNS server is configured to allow zone transfers from the secondary DNS server**

1. On the primary DNS server, click **Start**, point to **All Programs**, click **Administrative Tools**, and then click **DNS**.
2. In the console tree, double-click the DNS server.
3. In the console tree, double-click **Forward Lookup Zones** or **Reverse Lookup Zones**, as applicable.
4. Right-click the zone, click **Properties**, and then click the **Zone Transfers** tab.
5. Ensure that the **Allow zone transfers** check box is selected.

If zone transfer fails with Event ID 6525 “Zone transfer for secondary zone <zone_name> refused by master server” and the master server allows dynamic updates for the zone, these failures are due to the zone transfer throttling mechanism, and they are expected. This mechanism limits the number of zone transfers to allow regular dynamic updates to take place.

# Dynamic updates for host records fail

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Computers can fail to register the following host records:

- Address (A) records, for a forward lookup zone
- Pointer (PTR) records, for a reverse lookup zone

There might not be any error messages associated with these events. The only symptom of the problem might be the fact that the host's records do not show up in the Domain Name System (DNS) zone.

**Cause**

The host might not be configured properly to allow dynamic DNS updates. It might be configured to use an external DNS server, or it might be experiencing other DNS configuration problems.

**Solution**

You can use the following procedure to verify that the host is configured for dynamic DNS updates.

**To verify that the host is configured for dynamic DNS updates**

1. Log on to the computer with the Administrator account.
2. Click **Start**, click **Control Panel**, and then double-click **Network Connections**.
3. In **Network Connections**, right-click **Local Area Connection**, and then click **Properties**.
4. In **Local Area Network Connection Properties**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
5. In **Internet Protocol (TCP/IP) Properties**, click **Advanced**, and then click the **DNS** tab.
6. Ensure that both of the following check boxes are selected:
   - **Register this connection's addresses in DNS**
   - **Use this connection's DNS suffix in DNS registration**
7. Click **OK**.

You can use the following procedure to verify that the client does not have an external DNS server, such as a DNS server from an Internet service provider (ISP), in its TCP/IP configuration. In most cases, the client should not use a DNS server from an ISP as either the preferred or alternate DNS server, because the DNS server at the ISP is unable to resolve internal names. Using a DNS server from an ISP in a client's TCP/IP configuration can also cause problems with conflicting internal and external namespaces.

**To verify DNS configuration in TCP/IP settings**

1. Log on to the computer with the Administrator account.
2. Click **Start**, click **Control Panel**, and then double-click **Network Connections**.
3. In **Network Connections**, right-click **Local Area Connection**, and then click **Properties**.
4. In **Local Area Network Connection Properties**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
5. If **Obtain an IP address automatically** is selected, type the following at a command prompt:

   **ipconfig /all**
6. Review the DNS server settings and verify that they are correct.

You can use the following procedure to verify that the start-of-authority (SOA) resource record can be resolved by the DNS servers. In this procedure, you use the Nslookup.exe tool to test name resolution for the SOA record for the domain that the

client is attempting to register in. Test this name resolution from each one of the DNS servers that the client is configured to use.

**To verify that the SOA record can be resolved by the DNS servers**

1. At a command prompt, type the following command, and then press ENTER:

   **nslookup**

2. At the **nslookup:** prompt, type the following command, and then press ENTER:

   **set querytype=SOA**

3. At the **nslookup:** prompt, type the full name of the DNS zone that the client should be registering in — and include a terminating dot at the end of the domain name — and then press ENTER.

4. To test another DNS server, at the **nslookup:** prompt, type the following command, and then press ENTER:

   **server** *IP_address*

   Then, type the domain name to be tested, and then press ENTER.

If this query attempt fails from any of your DNS servers, you might need to remove that DNS server from the client's TCP/IP settings.

You can use the following procedure to verify that the DNS zone is enabled for dynamic updates. Open the DNS management console to verify that the zone that the clients need to register in is configured to accept dynamic updates.

**To verify that the DNS zone is enabled for dynamic updates**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **DNS**.

2. In the console tree, double-click the appropriate DNS server name, and then double-click **Forward Lookup Zones**.

3. Right-click the zone, and then click **Properties**.

4. On the **General** tab, view the **Dynamic updates** setting and make sure that it is set to **Nonsecure and secure** or **Secure only**.

5. If the setting is already set to **Secure only** and updates are still failing, try setting the zone to **Nonsecure and secure** for testing. If failures are seen only for **Secure only**, see Secure dynamic updates fail.

You can use the following procedure to verify that the Dynamic Host Configuration Protocol (DHCP) Client service is started. The DHCP Client service is used to perform dynamic updates, and it must be running.

**To verify the status of DHCP or to start DHCP**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Services**.

2. In the details pane, double-click **DHCP Client** and verify that the status of that service is **Started**. If the status is not **Started**, click **Start** to start the service.

You can use the following procedure to verify whether a single-label DNS domain name is being used.

**To verify whether a single-label DNS domain name is being used**

1. At a command prompt, type the following, and then press ENTER:

   **ipconfig /all**

2. View the Primary DNS Suffix and the **Connection-specific DNS Suffix** to make sure that the specified domain name has at least two parts, separated by a dot, for example, fabrikam.com. An example of a single-label domain is fabrikam.

3. If the host is using a single-label domain name, and this is correct for the environment, you have to use the registry setting of **UpdateTopLevelDomainZones**. For more information about this registry setting, see article 300684, "Information about configuring Windows for domains with single-label DNS names," in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=37924).

# Additional Resources for Domain Name System

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For information about how Domain Name System (DNS) works, see the following resources:

- DNS Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48145)
- DNS Support for Active Directory Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48147)

For information about troubleshooting DNS problems, see Troubleshooting Domain Name System.

# Group Policy Operations Guide

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Group Policy Operations Guide provides troubleshooting information for Group Policy in Microsoft Windows Server 2003 with Service Pack 1 (SP1). This guide also covers troubleshooting issues for earlier operating systems including Windows 2000 and Windows XP. It assumes that you are using the Group Policy Management Console (GPMC).

**In this guide**


- Group Policy Troubleshooting

# Group Policy Troubleshooting

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides a framework for troubleshooting Group Policy in an Active Directory environment. It includes checklists for quick fixes, a summary of troubleshooting tools, flowcharts for root cause analysis, recommendations for solving specific Group Policy problems, and guidance for interpreting log files.

**In this guide**


- Introduction to Group Policy Troubleshooting
- Verifying Administrative Workstation Settings for Troubleshooting Group Policy
- Group Policy Quick Fixes
- Configuring Administrative Workstation Settings for Group Policy Troubleshooting
- Troubleshooting Group Policy Problems
- Additional Resources for Group Policy



**Acknowledgments**

Produced by: Microsoft Windows Server User Assistance team

Project Writer: John Kaiser

Project Editor: Kate Gleeson

Contributing Writers: Patty Nicholson, Colin Torretta

Contributing Editors: Carolyn Eller

Technical Reviewers: Tim Thompson, Mike Stephens, Mark Williams

Microsoft Most Valuable Professional (MVP) Reviewers: Darren Mar-Elia, Jeremy Moskowitz, Derek Melber, John Savill

# Introduction to Group Policy Troubleshooting

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide explains how to troubleshoot Group Policy. If you are using this guide for the first time, review the following sections of this introduction.

**When to Use this Guide**
This troubleshooting guide is designed to help you identify and resolve problems related to Group Policy. You should use this guide when:

- You have configured a policy setting but it is not applying as expected.
- You have specified the scope of Group Policy through filtering but it is not applying as expected.
- You are having problems using a specific Group Policy extension, such as Security Settings or Folder Redirection.

**Who Should Use this Guide**
This guide is written for IT professionals in small, medium, and large organizations. It assumes a basic understanding of what Group Policy is, how it works, and why your organization uses it. You should also have a thorough understanding of how Group Policy is deployed and managed in your organization. This includes understanding how Group Policy is configured throughout your organization and the mechanism that your organization uses to configure Group Policy settings.

**How to Use this Guide**
This guide is divided into four sections. Each section addresses a type of problem, ranging from less complex to more complex.

Verifying Administrative Workstation Settings for Troubleshooting Group Policy

This section provides a list of prerequisites and settings that must be verified before you troubleshoot. Read this section first.

Group Policy Quick Fixes

This section provides quick solutions to the most common Group Policy problems. Read this section if you have verified your computers settings, but your problem persists.

Configuring Administrative Workstation Settings for Group Policy Troubleshooting

This section describes how to configure your computer for advanced troubleshooting. Read this section if you cannot find a solution in the Group Policy Quick Fixes section.

Troubleshooting Group Policy Problems

This section provides information and diagnostic procedures that help you identify and fix Group Policy problems.

Additional Resources for Group Policy

This section provides links to additional information that you can use to identify and fix Group Policy problems.

# Verifying Administrative Workstation Settings for Troubleshooting Group Policy

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic provides information about ensuring you have the appropriate permissions to administer and troubleshoot Group Policy and provides other tips for setting up your administrative workstation. Before you begin troubleshooting, verify that your computer is configured properly and that Group Policy is set up and running properly.

**Settings to Verify Before You Troubleshoot**
Verify all of the following items:

Security permissions are configured for administering Group Policy

The correct domain controller is selected in the Group Policy Object Editor and GPMC

**Security permissions are configured for administering Group Policy**
In order to administer Group Policy, you must have the necessary privileges to use GPMC and the Group Policy Object Editor. You also need privileges to create GPOs or to manage links from a specific site, domain, or OU to GPOs. Control of existing GPOs can be delegated to specific users or groups, so it is possible for an administrator to be able to use GPMC to view GPOs, but not be able to modify, delete, or link them. Use the following guidelines to help ensure that permissions are configured appropriately:

- Use Active Directory Users and Computers to verify that the account you are using is a member of a group that has these privileges. (Check the group memberships for the user account, and also verify that the privileges for the group have not been changed.)

- Avoid adding the privileges to an individual user account. If necessary, create a new group with a name that clearly indicates its purpose.

- Changes to security groups' memberships or privileges, or to the permissions on Group Policy objects or actions, need to be replicated to domain controllers throughout the system. Until this replication is completed the changes might be applied unevenly. In rare cases you might want to force replication.

To see or change the access control lists that affect management of a GPO, open the GPO in GPMC and look at the **Delegation** tab. The GPO can only be applied by members of groups that have Read permissions. To change the security filters, click **Advanced**.

**The correct domain controller is selected in the Group Policy Object Editor and GPMC**
Each domain controller has a copy of every GPO in the domain. The default and best practice is to edit GPOs on the primary domain controller (the PDC Emulator), and allow the changes to replicate to other domain controllers.

If that is not practical due to bandwidth or other issues, administrators can change the domain controller focus for the instances of GPMC that they are using.

If administrators in your organization edit GPOs on different domain controllers, set up processes to avoid this sort of conflict. For example, you might delegate editing permissions on individual GPOs to specific users, or to a group that focuses on the same domain controller.

# Group Policy Quick Fixes

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic provides tips for quickly addressing common Group Policy problems. Problems with the application of Group Policy often involve the technologies on which Group Policy depends, or implementation errors.

This section provides a summary of the most common Group Policy problems and the solutions to those problems. You can use the information in this section to resolve problems in the same way you would use a FAQ to find answers to common questions. Read this section before you begin any advanced troubleshooting.

**Common Group Policy Problems**
Listed below are common Group Policy problems and quick fixes:


I want Group Policy changes to be applied immediately.

Some policy areas are missing when I open up the Group Policy Object Editor.

Some settings are missing in the Group Policy Object Editor.

Some policies are not applied over dial-up connections.

Security Options are not applied.

Account Policies are not applied for domain users.


**I want Group Policy changes to be applied immediately.**
Sometimes you may want to see the effect of a Group Policy change without having to wait for the default refresh intervals (90 minutes for domain members and 5 minutes for domain controllers). To see the effect of Group Policy immediately, open a command prompt and run GPUPdate. To force GPOs to reapply even if they have not changed, at a command prompt type **gpupdate /force**. To see other parameters, type **gpupdate /?**.

To run **gpupdate /force** automatically, click the Fix this problem link. Then, click Run in the File Download dialog box.

Fix this problem

| Note |
| --- |
| This wizard may be in English only; however, the automatic fix also works for other language versions of Windows. |

| Note |
| --- |
| If you are not on the computer that has the problem, save the automatic fix to a flash drive or to a CD so that you can run it on the computer that has the problem. If you do not want to use the automatic fix, you can do it yourself by performing the steps in the procedure above. |

**Some policy areas are missing when I open up the Group Policy Object Editor.**
Each area of Group Policy functionality is implemented by a Microsoft Management Console snap-in DLL that is registered by default. For example, Administrative Templates and Scripts use gptext.dll, folder redirection uses fde.dll, and Security Settings uses wsecedit.dll. If these DLLs are un-registered or removed, the underlying Group Policy editing functionality that they implement will not appear in the Group Policy Object Editor.

To resolve this issue, you need to re-register the appropriate MMC snap-in DLL that implements the missing functionality.

**To reregister MMC snap-in DLLs**

1. Open a command prompt.
2. Type the following command: **regsvr32** *snap-inDLLname*

By default, you can find all of the MMC snap-in DLLs related to Group Policy in %systemroot%\system32.

**Some settings are missing in the Group Policy Object Editor.**
Administrative Templates can contain both true policies and preferences, but by default the Group Policy Object Editor exposes only true policies. To show all possible settings, you need to expose preferences in the Group Policy Object Editor.

### To expose preferences in the Group Policy Object Editor

1. Highlight the Administrative Templates node for which you want to see preferences.
2. On the **View** menu, click **Filtering**, and then clear the **Only show policy settings that can be fully managed** check box.

**Some policies are not applied over dial-up connections.**
By default, certain Group Policies are not applied to users who connect to your network over dial-up connections. If you want to make sure that users receive all Group Policies regardless of the speed of their connection, you can change the definition of what counts as a slow link. The default setting for this is 500 Kbps or less, which means that anybody who accesses your network with a connection that is 500 Kbps or slower might not receive all GPOs intended for them. To change the Group Policy settings that define a slow link, modify the following settings:

- **For users:** Click **User Configuration**, click **Administrative Templates**, click **System**, click **Group Policy**, and then click **Group Policy Slow Link Detection**.
- **For computers:** Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Group Policy**, and then click **Group Policy Slow Link Detection**.

| **Note** |
| --- |
| If a user connects to the network with cached credentials, Group Policy is not processed. To ensure that Group Policy is applied over a slow link, the user must select the **Logon using dialup connection** check box in the Windows Logon dialog box. |

**Security Options are not applied.**
The following Security options are only processed if they are applied at the domain level:

- Network Security: Force logoff when logon hours expire
- Accounts: Administrator account status
- Accounts: Guest account status
- Accounts: Rename administrator account
- Accounts: Rename guest account

If these Security Options are configured at the site or OU level they are ignored. You need to configure these Security Options at the domain level.

**Account Policies are not applied for domain users.**
Account Policies are only processed for domain user accounts if they are applied at the Domain level. Account Policies configured elsewhere, such as at the Site or OU level, are ignored. All settings found under Account Policies under Security Settings in Group Policy should only be configured at the domain level. If you are having trouble applying a Security Setting that is found under Account Policies, ensure that you have configured the setting only at the domain level.

# Configuring Administrative Workstation Settings for Group Policy Troubleshooting

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic provides information about the tools you use to troubleshoot Group Policy.

**Configuration Tasks for Troubleshooting**
To configure your administrative workstation for troubleshooting, you will need to perform one or more of the following tasks:

Install tools

Learn about built-in troubleshooting tools

Install tools for troubleshooting external issues

Enable auditing

Enable logging

**Install tools**
If you are administering or troubleshooting Group Policy, you should be using the Group Policy Management Console (GPMC) and other tools as necessary.

**Group Policy Management Console**
Group Policy Management Console is the preferred tool for administering Group Policy and it is an excellent tool for troubleshooting Group Policy. GPMC provides the following reporting functionality:

- Group Policy Modeling reports are used to predict the policies that will be applied at a specific client. A Windows Server 2003 domain controller is required to generate Group Policy Modeling reports.

- Group Policy Results reports gather information directly from the client to show the policies in effect, and include key policy events that have been logged at that client.

Both of these reports include valuable troubleshooting information. For example, you can see a list of the GPOs applied, and also the denied GPOs with the reason for denial. You can see which settings are or would be applied, and the winning GPO that supplied the value for the setting.

**Sample scripts**
GPMC also includes sample scripts that you can use to quickly perform a number of different troubleshooting tasks. If you cannot find a sample script that fits your needs, you can easily modify a sample script, or create your own script. The following sample scripts will help you troubleshoot various issues:

- List All GPOs in a Domain: ListAllGPOs.wsf

- List Disabled GPOs: FindDisabledGPOs.wsf

- List GPO Information: DumpGPOInfo.wsf

- List GPOs at a Backup Location: QueryBackupLocation.wsf

- List GPOs by Policy Extension: FindGPOsByPolicyExtension.wsf

- List GPOs by Security Group: FindGPOsBySecurityGroup.wsf

- List GPOs Orphaned in SYSVOL: FindOrphanGPOsInSYSVOL.wsf

- List GPOs With Duplicate Names: FindDuplicateNamedGPOs.wsf

- List GPOs Without Security Filtering: FindGPOsWithNoSecurityFiltering.wsf
- List SOM Information: DumpSOMInfo.wsf
- List SOMs With Links to GPOs in External Domains: FindSOMsWithExternalGPOLinks.wsf
- List Unlinked GPOs in a Domain: FindUnlinkedGPOs.wsf
- Print the SOM Policy Tree: ListSOMPolicyTree.wsf

For a complete list of available sample scripts, script documentation, and a list of scripting interfaces exposed by GPMC, please see the Group Policy Management Console SDK located at %programfiles%\gpmc\scripts\gpmc.chm on any computer where GPMC has been installed. (The Group Policy Management Console SDK is available in English only.)

For information about downloading and installing GPMC, see Group Policy Management Console with Service Pack 1 on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=21813).

**GPResult.exe**
There are two versions of GPresult.exe.

- The Windows Server 2003 version is included with Windows Server 2003 operating systems. GPresult.exe gathers and reports the RSoP data available from computers running Windows XP or Windows Server 2003. The report is similar to what you would get by generating a Group Policy Results report in GPMC. For more information, see Command-Line Reference for Windows Server 2003 on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=20331).
- The Windows 2000 version is included with the Microsoft Windows® 2000 Server Resource Kit Tools. To download the Windows 2000 version for free, see Windows 2000 Tools and Utilities at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=12920). GPresult.exe estimates the Group Policy settings that would be applied at a specific computer. This version provides information that is not available in the Windows Server 2003 version: It shows the registry value name set by each policy setting and lets you associate the GPO GUID that appears in Sysvol with the friendly name that was given to the policy setting (for example, Remove Run Command). For this reason, you might want to have both versions of Gpresult available. For more information, see the readme file included with the download.

**GPOTool.exe**
GPOTool.exe is a command-line tool to be used in replicated domains—domains that contain more than one domain controller. It traverses all of your domain controllers and checks each for consistency between the Group Policy container (that is, information contained in the directory service) and the Group Policy template (that is, information contained in the SYSVOL share on the domain controller). The tool also determines whether the policies are valid and consistent between all of your domain controllers and displays detailed information about the GPOs that have been replicated between your domain controllers.

- If you suspect that you are having problems with replication of Group Policy information, this tool helps you diagnose and isolate where Group Policy is not being replicated properly.
- The latest version of GPOTool is included in Windows Server 2003 Resource Kit Tools. For more information, see Windows Server 2003 Resource Kit Tools on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=27766). GPOTool.exe originally shipped with the Microsoft Windows 2000 Server Resource Kit. To download GPOTool.exe for free, see Gpotool.exe: Group Policy Verification Tool on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=17911).

**Software Installation Diagnostics Tool (Addiag.exe)**
Windows Server 2003 includes an advanced troubleshooting tool, Software Installation Diagnostics (Addiag.exe) that you can use to gather additional diagnostic information when troubleshooting Software Installation policy issues.

The binary executable for this tool is Addiag.exe. Running **addiag.exe /?** from a command prompt provides the usage syntax.

This tool displays detailed information about the applications visible in Active Directory and installed for the current user, as well as general diagnostic information and related Event Log entries.

Addiag.exe is included in the Windows Server 2003 Support Tools. Support tools not installed with the Windows operating system; you must install them separately from the \Support\Tools folder of the Windows operating system CD.

For more information, see Addiag.exe: Application Deployment Diagnosis on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43037).

**Learn about built-in troubleshooting tools**

This section provides information about using Group Policy Modeling and Group Policy Results in GPMC to generate reports.

**Group Policy Results**

GPMC leverages the Resultant Set of Policy (RsoP) functionality in Windows Server 2003 and Windows XP to provide reports on the way Group Policy is applied at individual clients. Because these reports rely on functionality that is new with Windows XP and Windows Server 2003, the clients for which you generate the reports must be running one of these operating systems.

| **Note** |
| --- |
| Group Policy Results depends on Windows Management Instrumentation (WMI). If WMI is not working, you will not be able to use RSoP or Group Policy Results. In this case, viewing the userenv log file in verbose mode provides you with information you will need to troubleshoot, for example, by showing you which GPOs did not apply. For more information about userenv, see Fixing Group Policy problems by using log files. |

**To generate a Group Policy Results report**

1. In the **Group Policy Management** console, right-click **Group Policy Results**, and then at the bottom of the navigation pane click **Group Policy Results Wizard**.
2. In the wizard, specify the computer or computer/user combination you want to investigate. The report that appears in the details pane provides information about Group Policy application on the client.

**Group Policy Modeling**

Before you implement a GPO, use Group Policy Modeling to validate the effect it will have. The Group Policy Modeling report has the information about the Summary and Settings tabs similar to what you would see for a Group Policy Results report. Group Policy Modeling reports do not collect policy events from the client. Instead of the **Policy Events** tab, a list of the conditions that were applied when creating the model appears on the **Query** tab.

**To generate a Group Policy Modeling report**

1. In the **Group Policy Management** console, right-click **Group Policy Modeling**, near the bottom of the navigation pane, and select **Group Policy Modeling Wizard**.
2. In the **Group Policy Modeling Wizard**, specify the computer or computer/user combination you want to investigate. The report that appears in the details pane provides information about the anticipated Group Policy application on that client.

**Summary Tab**

On the **Summary** tab, the sections that appear under both Computer Configuration and User Configuration headings are listed and described in the following table.

**Table 1 Summary Tab of Group Policy Results Reports**

| Section | Information |
| --- | --- |
| General | Computer name<br><br>The domain and site of which the computer is a member<br><br>The last time Group Policy from the computer's Active Directory hierarchy was applied<br><br>User name (if any)<br><br>The domain and site of which the user is a member<br><br>The last time Group Policy from the user's Active Directory hierarchy was applied |

| | |
|---|---|
| Group Policy Objects Applied | List of GPOs that were applied |
| Group Policy Objects Denied | List of GPOs that were denied, with the reason for the failure. |
| Security Group Membership when Group Policy was applied | Security group memberships in effect when group policies were evaluated |
| WMI Filters | WMI filters that were applied, whether they evaluated as True or False, and what GPO called them |
| Component Status | Success or failure, including errors, of core client Group Policy functionality and client side extensions (CSEs) |

**Settings Tab**

The **Settings** tab lists the actual settings applied. These are sorted by the source of the setting (for example Computer Configuration/Windows Settings or User Configuration/Administrative Templates). The report includes the winning GPO for each setting.

With this information, you can locate the GPO in the navigation pane. The information that is exposed when you click the GPO depends on the privileges granted to your user account. If you have sufficient privileges, you can review or edit the settings and also get a list of the sites, domains, and OUs that link to that GPO.

**Policy Events Tab**

When you use GPMC to generate a report of the resulting set of policy on a client, events that were logged at that client and pertain to Group Policy are listed on the **Policy Events** tab. Sources for these events include the core Group Policy engine on the client (the Userenv process) and the CSEs for Group Policy.

The display on the **Policy Events** tab is similar to the Event Viewer display. This is an especially useful troubleshooting feature of GPMC because it filters only the Group Policy events from the Event Viewer. Table 2 lists and defines the sources.

**Note**

To view the Group Policy events on computers running Windows XP SP1 or Windows Server 2003 you must be a local administrator on that computer. If you have the necessary privileges to generate a Group Policy Results report but you do not have the privileges to view Group Policy events on the client, the Policy tab displays the message "Unable to open event log: Access is Denied" instead of the list of events.

**Table 2 Policy Events Tab of Group Policy Results Reports**

| Source name in Policy Events log | Full name of source | Functionality |
|---|---|---|
| UserEnv | User Environment (Group Policy core engine) | Locates and applies GPOs at start up, log on, or the configured Policy Refresh Interval |
| SceCli | Security CSE | Reads all GPOs that reach the client and determines which policy settings are applied |
| Application Management | Software Installation CSE | Processes Software Installation settings, including installation, upgrades, and removal |
| Folder Redirection | Folder Redirection CSE | Processes Folder Redirection |
| UserInit | Scripts CSE | Implements logon, logoff, startup, and shutdown scripts |

To avoid flooding the client log file, some logging is blocked in certain situations. For example, if a computer is not connected to the network and a user logs on with cached credentials, the Component Status entries on the **Summary** tab of the Group Policy Results report show the failure to access and apply Group Policies. However, the list of associated failure events do not appear in the Application event log on the client or on the **Policy Events** tab.

**Note**

By default, UserEnv logging is not verbose — only errors and warnings are reported, and appear on the **Policy Events** tab. For more information, see Fixing Group Policy problems by using log files.

**Install tools for troubleshooting external issues**
Group Policy has dependencies on several operating system technologies such as Active Directory, Domain Name System (DNS), and File Replication Service (FRS). The following tools can help you troubleshoot problems related to these dependent technologies.

**Sonar.exe**
Sonar is a command-line tool that allows administrators to monitor key statistics and status about members of a file replication service (FRS) replica set. Use Sonar to watch key statistics on a replica set in order to monitor traffic levels, backlogs, and free space. To download Sonar for free, see Sonar.exe: FRS Status Viewer on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=16719).

**Active Directory Support Tools**
Help and Support Center in Windows Server 2003 provides a list of Active Directory support tools in the topic "Active Directory support tools." Use these tools to troubleshoot Active Directory issues.

**Other Windows Server 2003 Command-Line Tools**
Windows Server 2003 includes a number of command line tools including ping.exe, netdiag.exe, and dcdiag.exe. For a complete reference of the tools included with Windows Server 2003, see Windows Server 2003 Technical Reference: Tools and Settings Collection on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4972).

**Enable auditing**
Administrators can use Group Policy to enable auditing of Windows registry keys, or they can use a Security Template to Audit Registry Keys, or they can use Registry Editor for this purpose. Administrators must also configure SACLs on the registry keys that are being audited.

To enable auditing for the Windows registry, you can use the **Audit object access** policy setting in **Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policies**.

The **Audit object access** policy setting determines whether to audit the event of a user accessing an object (for example, a file, folder, registry key, printer, and so on) that has its own system access control list (SACL) specified. Default: No auditing.

To configure permissions for auditing the registry keys, in **Group Policy Object Editor** console, in **Computer Configuration\Windows Settings\Security Settings**, right-click **Registry**, click **Add Key**, select the registry key for which you want to allow auditing, and then click **Advanced**. In the **Database Security** dialog box, click the **Auditing** tab, and then specify the users or groups that you want to give permissions to and set the permissions you want to use.

For more information about enabling auditing of registry keys, see article ID 324739 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=35275).

**Enable logging**
To troubleshoot problems with Group Policy administration or application, you can enable logging based on the functionality you are investigating.

For more information about enabling logging for client-side extensions and server-side logging, see the subsections or Fixing Group Policy problems by using log files.

**See Also**
**Other Resources**
Windows Server 2003 Technical Reference: Tools and Settings Collection

# Enable Logging for Group Policy Management Console

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To troubleshoot problems with administering Group Policies using the Group Policy Management Console, you can enable GPMC logging. This logging covers all Group Policy administration tasks, except editing Group Policy objects.

**When To Enable GPMC Logging**

GPMC logging, by default, is set to only log errors. To avoid using unnecessary disk space, you should leave GPMC logging at the default setting. You can enable GPMC logging to gather verbose troubleshooting information when an administrative task causes errors.

**Enabling Verbose GPMC Logging**

| ⚠Caution |
| --- |
| This section contains information about modifying the registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore the registry if a problem occurs. For information about how to back up, restore, and edit the registry, see Description of the Microsoft Windows Registry on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35500). |

**To enable GPMC logging on the administrative computer**

1. Log on to the administration computer as the administrator.
2. Click **Start**, **Run**, and then run Regedit to add the following registry value (or, , if the value already exists, modify it):

   **Key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics**
3. Enter the following:

   **Value: gpmgmttracelevel**

   **Value Type: REG_DWORD**

   **Value Data: 2 (Hexadecimal)**

**Gpmgmttracelevel** can have the following values:

- **NORMAL 0x00000001**
- **VERBOSE 0x00000002**

The default value is **NORMAL (0x00000001)**.

**To restrict GPMC logging to file output**

1. Log on to the administration computer as the administrator.
2. Click **Start**, **Run**, and then run Regedit to add the following registry value (or, if the value already exists, modify it):

   **Key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics**
3. Enter the following:

   **Value: gpmgmtlogfileonly**

   **Value Type: REG_DWORD**

   **Value Data: 1 (Hexadecimal)**

**Gpmgmtlogfileonly** can have the following values:

- **LOGTOFILEONLY 0x00000001**

- **LOGTOFILEANDDEBUGGER 0x00000002**

The default value is **LOGTOFILEONLY (0x00010001)**.

GPMC writes to the following log files in the **%Temp** folder:

- GPmgmt.log: contains all other GPMC logging information.

- GPmgmtManaged.log: Contains logging information for the reporting functionality.

**See Also**
**Concepts**
Fixing Group Policy problems by using log files

# Enabling Logging for Group Policy Editor

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can enable logging of events generated by the Group Policy Object Editor for its core operations on the administrative computer.

**When To Enable Group Policy Object Editor Logging**
Group Policy Object Editor logging, by default, is set to only log errors. To avoid using unnecessary disk space, you should leave Group Policy Object Editor logging at the default setting. You can enable Group Policy Object Editor logging to gather verbose troubleshooting information when a Group Policy Object editing task causes errors.

**Enabling Group Policy Object Editor Logging**

| ⚠**Caution** |
| --- |
| This section contains information about modifying the registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore the registry if a problem occurs. For information about how to back up, restore, and edit the registry, see Description of the Microsoft Windows Registry on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35500). |

**To enable GPedit logging on the administrative computer**

1. Log on to the client computer as the administrator.
2. Click **Start**, **Run**, and then run Regedit to add the following registry value (or, if the value already exists, modify it):

   **Key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon**
3. Enter the following:

   **Value: GPEditDebugLevel**

   **Value Type: REG_DWORD**

   **Value Data: 10002 (Hexadecimal)**

**GPEditDebugLevel** can have the following values:

- **NONE 0x00000000**
- **NORMAL 0x00000001**
- **VERBOSE 0x00000002**
- **LOGFILE 0x00010000**
- **DEBUGGER 0x00020000**

The default value is **NORMAL|LOGFILE (0x00010001)**.

To disable logging, select **NONE (0X00000000)**.

You can also combine the **GPEditDebugLevel** values. For example, you can combine **VERBOSE 0x00000002** and **LOGFILE 0x00010000** to get 0x00010002. If you set **DGPEditDebugLevel** to 0x00010002, you turn on both **LOGFILE** and **VERBOSE**. By combining these values, you achieve the same effect as using an **OR** statement as follows:

0x00010000 OR 0x00000002 = 0x00010002

The log file is written to the **GPEdit.log** file in the **%Systemroot%\Debug\UserMode** folder.

**See Also**
**Concepts**
Enable Logging for Group Policy Object Editor Client Side Extensions

# Enable Logging for Group Policy Object Editor Client Side Extensions

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can enable logging of events generated by the Group Policy Object Editor for its Client Side Extensions on the administrative computer.

**When To Enable Group Policy Object Editor Client Side Extensions Logging**
Group Policy Object Editor CSE logging, by default, is set to only log errors. To avoid using unnecessary disk space, you should leave CSE logging at the default setting. You can enable CSE logging to gather verbose troubleshooting information when a Group Policy Object CSE editing task causes errors.

**Enabling Group Policy Object Editor Client Side Extensions Logging**

| ⚠Caution |
| --- |
| This section contains information about modifying the registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore the registry if a problem occurs. For information about how to back up, restore, and edit the registry, see Description of the Microsoft Windows Registry on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35500). |

**To enable CSE Extensions logging on the client**

1. Log on to the client computer as the administrator.
2. Click **Start**, **Run**, and then run Regedit to add the following registry value (or modify it, if the value already exists):

   **Key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon**
3. Enter the following:

   **Value: GPTextDebugLevel**

   **Value Type: REG_DWORD**

   **Value Data: 10002 (Hexadecimal)**

**GPTextDebugLevel** can have the following values:

- **NONE 0x00000000**
- **NORMAL 0x00000001**
- **VERBOSE 0x00000002**
- **LOGFILE 0x00010000**
- **DEBUGGER 0x00020000**

The default value is **NORMAL|LOGFILE (0x00010001)**.

To disable logging, select **NONE (0X00000000)**.

You can also combine the **GPTextDebugLevel** values. For example, you can combine **VERBOSE 0x00000002** and **LOGFILE 0x00010000** to get 0x00010002. If you set **GPTextDebugLevel** to 0x00010002, you turn on both **LOGFILE** and **VERBOSE**. By combining these values, you achieve the same effect as using an **OR** statement as follows:

0x00010000 OR 0x00000002 = 0x00010002

The log file is written to the **GPText.log** file in the **%Systemroot%\Debug\UserMode** folder.

**See Also**

**Concepts**

Fixing Group Policy problems by using log files

# Enable Logging for Core Group Policy

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To troubleshoot problems with user profiles and Group Policy processing, you can enable userenv logging.

| ⚠Caution |
| --- |
| This section contains information about modifying the registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore the registry if a problem occurs. For information about how to back up, restore, and edit the registry, see Description of the Microsoft Windows Registry on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35500). |

**To enable userenv logging on the client**

1. Log on to the client computer as the administrator.
2. Click **Start**, **Run**, and then run Regedit to add the following registry value (or, if the value already exists, modify it):

   **Key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon**
3. Enter the following:

   **Value: UserEnvDebugLevel**

   **Value Type: REG_DWORD**

   **Value Data: 10002 (Hexadecimal)**

**UserEnvDebugLevel** can have the following values:

- **NONE 0x00000000**
- **NORMAL 0x00000001**
- **VERBOSE 0x00000002**
- **LOGFILE 0x00010000**
- **DEBUGGER 0x00020000**

The default value is **NORMAL|LOGFILE (0x00010001)**.

To disable logging, select **NONE (0X00000000)**.

You can also combine the **UserEnvDebugLevel** values. For example, you can combine **VERBOSE 0x00000002** and **LOGFILE 0x00010000** to get 0x00010002. If you set **UserEnvDebugLevel** to 0x00010002, you turn on both **LOGFILE** and **VERBOSE**. By combining these values, you achieve the same effect as using an **OR** statement as follows:

0x00010000 OR 0x00000002 = 0x00010002

The log file is written to the **Userenv.log** file in the **%Systemroot%\Debug\UserMode** folder. If the **Userenv.log** file exists and is greater than 300 KB, the existing file is renamed to **Userenv.bak**, and a new log file created.

**See Also**
**Concepts**
Fixing Group Policy problems by using log files

# Enable Logging for Security Settings

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The security settings configuration and analysis tools include scecli.dll, which implements the client-side extension for Group Policy. Scecli.dll provides client-side interfaces to the security configuration engine and does Resultant Set of Policies (RSoP) logging during policy propagation.

Administrators can enable logging for scecli.dll. The log is written to the Winlogon.log file, which is created in the %SYSTEMROOT%\Security\Logs folder.

**To configure logging output for Security Configuration Client**

1. Start Registry Editor (Regedt32.exe).
2. Locate and click the following key in the registry:

   **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtensions\{827D319E-6EAC-11D2-A4EA-00C04F79F83A}**

   | **Note** |
   | --- |
   | The registry key is one path; it has been wrapped for readability. |

3. On the **Edit** menu, click **Add Value**, and then add the following registry value:

   **Value name: ExtensionDebugLevel**

   **Data type: DWORD**

   **Value data: 2**

To reproduce the failure, refresh the policy settings by typing the following at the command prompt: **gpupdate /force**.

For more information, see the Tools and Setting Collection in the Windows Server 2003 Technical Reference.

# Enable Logging for Folder Redirection Client Side Extension

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Folder Redirection can provide a detailed log that you can use to troubleshoot Group Policy issues.

**When To Enable Folder Redirection Logging**
Folder redirection CSE logging, by default, is set to only log errors. To avoid using unnecessary disk space, you should leave folder redirection logging at the default setting. You can enable folder redirection logging to gather verbose troubleshooting information when investigating folder redirection errors on a client computer.

**Enabling Verbose Folder Redirection Logging**

| ⚠**Caution** |
| --- |
| This section contains information about modifying the registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore the registry if a problem occurs. For information about how to back up, restore, and edit the registry, see Description of the Microsoft Windows Registry on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35500). |

### To enable Folder Redirection logging on the client computer

1. Log on to the administration computer as the administrator.
2. Click **Start**, **Run**, and then run Regedit to add the following registry value (or modify it, if the value already exists):

   **Key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics**
3. Enter the following:

   **Value: FdeployDebugLevel**

   **Value Type: REG_DWORD**

   **Value Data: 0F (Hexadecimal)**

**FdeployDebugLevel** can have the following values:

- **NORMAL 0x00000000**
- **VERBOSE 0x0000000F**

The default value is **NORMAL (0x00000000)**.

The log file can be found at: %Systemroot%\debug\usermode\fdeploy.log.

**See Also**
**Concepts**
Fixing Group Policy problems by using log files

# Enable Logging for Software Installation Client Side Extension

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Software installation can provide a detailed log that you can use to troubleshoot Group Policy issues.

**When To Enable Software Installation Logging**
Software Installation CSE logging, by default, is set to only log errors. To avoid using unnecessary disk space, you should leave software installation logging at the default setting. You can enable software installation logging to gather verbose troubleshooting information when investigating software installation errors on a client computer.

**Enabling Verbose Folder Redirection Logging**

| ⚠**Caution** |
| --- |
| This section contains information about modifying the registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore the registry if a problem occurs. For information about how to back up, restore, and edit the registry, see Description of the Microsoft Windows Registry on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35500). |

**To enable software installation logging on the client computer**

1. Log on to the administration computer as the administrator.
2. Click **Start**, **Run**, and then run Regedit to add the following registry value (or modify it, if the value already exists):

   **Key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics**

3. Enter the following:

   **Value: AppmgmtDebugLevel**

   **Value Type: REG_DWORD**

   **Value Data: 9B (Hexadecimal)**

**FdeployDebugLevel** can have the following values:

- **NORMAL 0x00000000**
- **VERBOSE 0x0000009B**

The default value is **NORMAL (0x00000000)**.

The log file can be found at: %Systemroot%\debug\usermode\appmgmt.log.

**See Also**
**Concepts**
Fixing Group Policy problems by using log files

# Enable Logging for Internet Explorer Client Side Extension

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Internet Explorer logs Internet Explorer Maintenance CSE Group Policy processing by default.

Internet Explorer logs Internet Explorer Maintenance CSE Group Policy processing events in the Brndlog.txt file. The Brndlog.txt file can be found in one of the following two locations depending on the version of Internet Explorer:

- %USERPROFILE%\Application Data\Microsoft\Internet Explorer
- %USERPROFILE%\Local Settings\Application Data\Microsoft\Internet Explorer

# Troubleshooting Group Policy Problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic introduces the recommended framework for troubleshooting Group Policy.

**Understanding Group Policy Processing**
Before discussing Group Policy troubleshooting, you need to develop a general understanding of how Group Policy is processed at the client. Group Policy processing has two distinct phases:

- Core Group Policy processing. When a client begins to process Group Policy, it must determine whether it can reach a domain controller, whether any GPOs have changed, and what policy settings (based on client side extension) must be processed. The core Group Policy engine performs the processing of this in the initial phase.

- Client side extension (CSE) processing. Policy settings are grouped into different categories, such as Administrative Templates, Security Settings, Folder Redirection, Disk Quota, and Software Installation. The settings in each category require a specific CSE to process them, and each CSE has its own rules for processing settings. The core Group Policy engine calls the CSEs that are required to process the settings that apply to the client.

**Recommended Process for Troubleshooting**
**High level steps to troubleshoot Group Policy**

1. Check the required infrastructure. Make sure required services and components are running and configured as expected.

2. Check the core configuration. Verify that the computer is connected to the network, joined to the domain, and has the correct system time.

3. Check Group Policy exceptions. Verify that exceptions (scope of management) such as security filtering, WMI filters, block inheritance, enforcement , loopback processing and slow link settings are not affecting normal GPO processing.

4. Use tools like GPResult.exe, GPOTool.exe and the GPMC to ensure that Group Policy settings that are expected to be delivered are actually delivered and that Group Policy objects on domain controllers are consistent and available.

5. Use event logs, userenv logs, and CSE logs to analyze the problem and find a solution.

After checking network connectivity, it's recommended to first verify whether a problem can be traced to core Group Policy. Because CSEs cannot begin to work until core Group Policy processing is completed, the issues described in the topics in Fixing Core Group Policy problems apply regardless of which CSE processes the setting. Therefore, you should make sure that your problem is not a core Group Policy problem before you begin to troubleshoot the other areas listed above.

To help determine what kind of Group Policy problems you have, do the following:

1. Generate a Group Policy Results report using GPMC.

2. Examine the results of the report to find the answers to three questions that are used in navigating the flowchart:

    - Does Group Policy Results list the GPO as applied?

    - Is the setting listed in Group Policy Results Report?

    - Is the GPO listed in the Denied List?

3. Compare the results of the report to the flowchart in Figure 1.

4. Most of the results in the report map to core Group Policy problems in the flowchart. Investigate the core Group Policy problems first, even if the results of your report point to a specific CSE.

5. If the problem persists, assess whether a CSE is involved.

6. If you still haven't located the problem, you might need to look at log files to determine the cause.

7. If you experience a component failure, check the userenv log and see if the Local Security Authority has logged any events.




For more information about running the Group Policy Results Report and interpreting the flowchart, see Fixing Core Group Policy problems.

To troubleshoot Group Policy, you need to address one or more of the following areas:

- Fixing Core Group Policy problems

  This section begins with a flowchart for troubleshooting Group Policy core functionality. It's designed to help you pinpoint likely causes of the problem and quickly eliminate other possible causes. For more information about specific problems, see the following:

  - Fixing Group Policy networking issues
  - Fixing Group Policy processing issues
  - Fixing Group Policy scoping issues
  - Fixing Group Policy structural issues

- Fixing issues relating to a specific CSE or type of policy setting. This section includes ways of troubleshooting that are unique to a specific type of policy. For information about specific problems, see the following:

  - Fixing Administrative Template policy setting problems
  - Fixing Security Settings Problems
  - Fixing Scripts policy settings problems
  - Fixing Software Installation policy setting problems

- Fixing Folder Redirection policy settings problems
- Fixing Disk Quota extension problems

- Fixing Group Policy problems by using log files

  This section helps you interpret many of the errors found in logs such as userenv.log and winlogon.log. Sometimes, you need to enable verbose logging on log files and analyze the results.

**Asynchronous Processing and Logon Optimization in Windows XP**

Group Policy can be applied during startup and logon (synchronous processing) or as a background task after startup or logon has completed (asynchronous processing). Changes received during periodic Group Policy refresh or in response to the gpupdate command are processed asynchronously. On computers running Windows XP, group policies received during logon are also processed asynchronously by default, so that the logon is completed more quickly.

Software Installation and scripts processing must be applied during startup or logon. Folder Redirection assigned to the user must be applied during logon. By default, Windows XP logs a user on in asynchronous mode. Group Policy is then applied in the background after the user is logged on. This results in faster logons. However, when a new GPO setting for Software Installation, Scripts, or Folder Redirection arrives at a computer running Windows XP, the user has already logged on by the time Group Policy has been evaluated. It is too late to apply the Software Installation setting. In this case a flag is set so that the next time computer is rebooted or the user logs on Group Policy will be evaluated and applied before the startup or logon is completed.

In situations where you need for users to receive software, implement folder redirection, or run new scripts in a single logon, apply a GPO with the setting **Always wait for the network at computer startup and logon to the computer**. This setting is located under **Computer Configuration\Administrative Templates\System\Logon** in the Group Policy Object Editor. For this setting to take effect, Group Policy must be refreshed or the computer restarted.

**Table 1 Timing of Synchronous and Asynchronous Processing**

| By default, how is policy processed on the client? | At Startup | At Logon | At Policy Refresh |
|---|---|---|---|
| Windows 2000 | Synchronously | Synchronously | Asynchronously |
| Windows XP Pro | Asynchronously | Asynchronously | Asynchronously |
| Windows Server 2003 | Synchronously | Synchronously | Asynchronously |

| Note |
|---|
| Servers do not perform asynchronous processing. |

For more information about different types of client side extension problems, see the following:

- Fixing Administrative Template policy setting problems
- Fixing Security Settings Problems
- Fixing Scripts policy settings problems
- Fixing Software Installation policy setting problems
- Fixing Folder Redirection policy settings problems
- Fixing Disk Quota extension problems

# Fixing Core Group Policy problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic uses a structured approach and flowcharts to troubleshoot Group Policy core functionality. In Windows XP and Windows Server 2003, you can use Resultant Set of Policy (RSoP) to track the final set of processed policy settings. GPMC provides an easy view into the RSoP data through its Group Policy Results tool. This topic is based on the use of Group Policy Results reports to view and analyze RSoP data. It is recommended that you read through this topic before moving to one of the subsections listed below.

**In this section**

- Fixing Group Policy networking issues
- Fixing Group Policy processing issues
- Fixing Group Policy scoping issues
- Fixing Group Policy structural issues

**Flowchart for Troubleshooting Group Policy Core Functionality**

Use the flowcharts in this topic to quickly identify the likely root causes for unexpected Group Policy behavior, based on three questions that are easily answered from the Group Policy Results report.

**Navigating the Troubleshooting Flowchart**

The troubleshooting flowcharts in this topic focus on core Group Policy processing. Their primary purpose is to help you validate that the underlying infrastructure is in place to support delivery of GPOs to the client, that the user and computer are appropriately targeted to receive the intended GPOs, and that Group Policy processing puts the correct GPOs into effect.




A Group Policy Results report is the primary resource for troubleshooting Group Policy using this flowchart. Specifically, when investigating a problem, the administrator — where possible — should generate a Group Policy Results report for the user and computer combination encountering the problem. The sections of the report contain the information you use to navigate through the flowchart.

**To generate a Group Policy Results report**

1. Open the Group Policy Management Console.

2. In the console tree, double-click the forest in which you want to create a Group Policy Results query, right-click **Group Policy Results**, and then click **Group Policy Results Wizard**.
3. In the **Group Policy Results Wizard**, click **Next** and enter the appropriate information.
4. After completing the wizard, click **Finish**.
5. If you want to print or save the report, right-click the settings report in the results pane, and then do one of the following:
    - Select **Print** to print the report.
    - Select **Save Report** to save the report.

An example of a Group Policy Results report is shown in Figure 2.




This example shows the **Summary** tab of the report with the Group Policy Objects sections under Computer Configuration Summary and User Configuration Summary expanded.

By examining the GPMC Results report, you can find answers to the following questions that are associated with the flowchart:

- Was the GPO applied to the client? The **Summary** tab shows this information.
- Is the policy setting listed in GPMC Results? The **Settings** tab shows this information in a way that reflects the user interface of the Group Policy Object Editor.
- Is the GPO listed as Denied in GPO Results? The **Summary** tab shows this information.

In this scenario the client has successfully received the GPO and the specific policy setting is in effect at the client. This means that the only problem is that the actual value of the policy setting is incorrect. For information about the individual settings that have been applied, see the **Settings** tab of the Group Policy Results report.




**GPO Applied, Policy Setting Listed in Group Policy Results Report**
The scenario used for "Navigating the Troubleshooting Flowchart" applies to this section.

**GPO Inheritance (Setting Listed)**
Although GPOs have been applied and the correct policy setting is listed, Group Policy inheritance might result in an unexpected GPO "winning" and providing a different value from the one expected. The settings are nested by source and type. To expose the settings, click **Show** on the nested rows, and then look at the **Winning GPO** column to discover which GPO defines the value for the policy setting. For more information, see GPO provides unexpected value.

**Replication (Setting Listed)**
After a change is made to either the GPO or the user or computer, that change must be replicated throughout the network. If you expected the winning GPO to supply a value for the setting other than the value that was actually applied, it might be that the GPO was changed recently and the change has not yet been replicated to the domain controller that supplied the GPO to the client. For more information, see Group Policy does not replicate.

**Group Policy Refresh (Setting Listed)**
If Group Policy Refresh has not occurred since the winning GPO was modified and replicated, the old value for the setting is applied. After the changes to a GPO have been replicated to the client's domain controller, they need to be transmitted to the client. This occurs when the client refreshes Group Policy. Until this has occurred the change will not be reflected at the client. You can either wait for a background refresh or force the refresh. For more information, see Group Policy does not refresh.

**Asynchronous Application of Group Policy (Setting Listed)**
Group Policy can be applied after the computer has started and the user has logged on. This is called asynchronous application of Group Policy, in contrast to synchronous processing that occurs as part of startup or logon.

If the problem is with a setting that can only be applied during startup or logon, it might have been detected during asynchronous Group Policy processing (for example, as part of a Group Policy refresh or during the asynchronous processing used for logon optimization in Windows XP).

**Client-Side Extension Issue (Setting Listed)**
After the core Group Policy engine has completed initial processing of the GPOs, it passes specific settings to CSEs to process. If the setting is listed but the value is wrong or the behavior on the client does not reflect the setting value, the failure might have occurred after this setting was passed to a CSE to process. For example, even if a Folder Redirection setting has been successfully passed to the Folder Redirection CSE, the CSE might not be able to complete processing for the setting. For more information, see the appropriate topic for the CSE. For example, for more information about Administrative Templates, see Fixing Administrative Template policy setting problems).

**Loopback Processing (Setting Listed)**
Loopback processing is a way to enforce a set of user settings at a computer regardless of who logs on at that computer. Typically, user settings are applied based on the site and OU membership of the user. If loopback processing is set for a computer, the user settings for anyone logged on to that computer are dependent (partially or fully) on the site and OU membership of the computer. The behavior depends on the mode of loopback processing. In Replace mode, only the user

settings defined in GPOs applied to the computer are used. In Merge mode, user settings from GPOs that would normally apply to the user are used provided they do not conflict with user settings in GPOs that apply to the computer.

Loopback processing only works if the computer and user are both in Windows 2000 or Windows Server 2003 domains. They can be in different domains, and one can be in a Windows 2000 domain while the other is in a Windows Server 2003 domain. You cannot deploy Group Policy to users in a Windows NT 4.0 domain by applying loopback to a computer in a Windows 2000 or Windows Server 2003 domain.

Security filters can affect the way loopback processing is applied. Even when the GPOs associated with the computer are used to define user settings, the user's credentials — not the computer's credentials — are validated against the GPO's security filter. Therefore the user's credentials determine whether the GPO should be applied.

For example, you could create a GPO with a security filter that restricts the GPO to system administrators, and then associate that GPO with a computer that is configured for loopback processing. The settings in that GPO would only be applied when a system administrator is logged on.

To determine whether loopback processing is in effect, look for the User Group Policy loopback processing mode setting on the **Settings** tab of the report, which appears under Computer Configuration\Administrative Templates\System\Group Policy. For more information, see Loopback processing does not work.

**GPO Applied, Policy Setting Not Listed in Group Policy Results Report**

In the Group Policy Results report, the structure of the **Settings** tab is similar to the structure used in the Group Policy Object Editor. Expand the sections on the **Settings** tab by clicking **Show**. If the expected policy setting does not appear at all, either no updated GPO containing the expected setting reached the client, or the setting might not be processed at the client.




**Replication (Setting Not Listed)**

After a setting is added to a GPO, that change must be replicated throughout the network. If the setting is specified in the GPO but is not listed in the Group Policy Results report on the client, it might be that the setting was recently added to the GPO, but the change has not yet been replicated to the domain controller that supplied the GPO to the client. For more information, see Group Policy does not replicate.

**Group Policy Refresh (Setting Not Listed)**

If Group Policy Refresh has not occurred since the winning GPO was modified and replicated, a newly added setting will not be applied. After the changes to the GPO have been replicated to the client's domain controller, they need to be transmitted to the client. This occurs during Group Policy refresh. You can either wait for a background refresh or force the refresh by running gpupdate, by logging off and logging on again (that is, for user configuration), or by restarting the computer (that is, for computer configuration). For more information, see Group Policy does not refresh.

**Lack of Operating System Support (Setting Not Listed)**

Some policy settings are supported on certain operating systems only or they require a minimum service pack to be applied. When a GPO delivers a policy setting to a client computer that does not support that setting, the operating system ignores the setting. For more information, see Policy setting is not supported.

**Slow Link (Setting Not Listed)**
Some policy settings do not apply over a slow link, which is defined by default as 500 kilobits-per-second. These settings include:

- Software Installation
- Scripts
- Folder Redirection settings

For more information, see Group Policy does not apply when connecting remotely over a slow link.

**GPO Not Applied, Listed in Denied List in Group Policy Results Report**
If the GPO successfully reaches the client, it appears either in the list of **Denied GPOs** or in the list of **Applied GPOs**. A GPO can be explicitly denied for any number of reasons. To determine whether a GPO is denied, look on the **Summary** tab or the Group Policy Results report. Under Computer Configuration Summary and again under User Configuration Summary, click **Show** to expand Group Policy Objects, and then show **Denied GPOs**. The reason for the denial is provided for each denied GPO.




**Security Filtering (GPO Denied)**
The user or computer does not have the user rights assigned for the GPO. The required privileges are Read and Apply Group Policy. Alternatively, a GPO might be associated with a Deny ACE, which overrides any other privileges granted to the user or computer. For more information, see Policy settings incorrectly applied or denied due to security filtering.

**Disabled Link (GPO Denied)**
There is a link to the GPO from a site, domain, or OU in the hierarchy of the user or computer, but that link has been explicitly disabled. You can quickly scan the navigation pane of GPMC for disabled links.

**Inaccessible GPO (GPO Denied)**
There is a link to the GPO, but the GPO cannot be accessed. There are several possible reasons for this:

- The permissions on the GPO or on folders in the path to the Group Policy template are insufficient for it to be accessed and read. If this situation occurs the Component Status section of the Group Policy Results report will indicate Failure for the component Group Policy Infrastructure.
- The GPO might have been deleted, but the link to it remains for some reason (such as replication lag).
- Network connectivity problems might prevent access to the GPO.
- The client is unable to contact any domain controller.

**Empty GPO (GPO Denied)**
A GPO will be denied if it has no settings. This occurs when an administrator has configured a GPO and linked to it, but has not set any policy settings within the GPO. Either remove the link to the GPO or add policy settings to the GPO. If there are no remaining links to the GPO, you should consider deleting it.

**WMI Filter (GPO Denied)**
A WMI filter applied to a GPO is essentially a Boolean (true/false) decision as to whether the entire GPO should be applied to the client computer. The filter is evaluated at the client when GPO is applied. Based on the embedded WQL query, the GPO will either be enabled or disabled. For more information, see Policy settings incorrectly applied or denied due to WMI filtering.

**GPO Not Applied, Not Listed in Denied List in Group Policy Results Report**
All the GPOs that reach the client appear on the **Summary** tab in either the **Group Policy Objects Applied** section or the **Group Policy Objects Denied** section. There are four lists altogether: two lists (that is, Applied GPOs and Denied GPOs) under **Computer Configuration Summary** for settings that are delivered from the computer's Active Directory hierarchy, and another two under **User Configuration Summary** for settings that are delivered from the user's Active Directory hierarchy. If the GPO is not listed as either **Applied** or **Denied** under either **Configuration Summary**, it did not reach the client.

Also note whether the GPO is listed in the expected **Configuration Summary** (Computers or Users) because this can affect which settings are actually applied, particularly if loopback processing is in effect.




**Scope of Management (GPO Not at Client)**
One of the most common causes of a GPO not being applied to a user or computer is that the GPO is not linked to a site, domain, or OU of which the computer or user is a member. GPOs are delivered to clients based on the site and OU memberships of the computer and the logged-on user. Group memberships are only used to further restrict application of the GPO. For more information, see GPO provides unexpected value.

**Replication (GPO Not at Client)**
After an administrator has linked a GPO to a site, domain, or OU in the hierarchy of the user or computer, the change must be replicated to the domain controller from which the client retrieved its GPOs. Also, if the user or computer has recently been added to an OU, the GPOs that apply to that OU might not be applied to the client until the change in OU membership has been replicated to the domain controller from which the client retrieves GPOs. For more information, see Group Policy does not replicate.

**Group Policy Refresh (GPO Not at Client)**
After an administrator has linked a GPO to a site, domain, or OU in the hierarchy of the user or computer, and the change has been replicated to the client's domain controller, the GPO still needs to reach the client. This occurs during Group Policy refresh. You can either wait for a background refresh or force the refresh.

**Network Connectivity (GPO Not at Client)**
Group Policy requires a reliable networking infrastructure to ensure appropriate communication between the client computer and a domain controller. This includes TCP/IP, DNS, and other dependent technologies. For more information, see Group Policy

does not refresh.

**See Also**
**Concepts**
Fixing Group Policy networking issues
Fixing Group Policy processing issues
Fixing Group Policy scoping issues
Fixing Group Policy structural issues

# Fixing Group Policy networking issues

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic includes a procedure you can use to verify your network configuration and addresses other network connectivity issues.

From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- Group Policy does not apply when connecting remotely over a slow link
- Unable to access domain controller
- Group Policy does not replicate

**Network Connectivity and Configuration**

For Group Policy to be received at the client, there must be network connectivity between the client and the domain controller. This section includes information about troubleshooting network connectivity.

To test for network connectivity problems, check system event logs on the client computer (that is, look for failed access attempts). You can also use the **ping** or **netdiag** commands to test network connectivity.

**To verify a properly configured connection:**

1. Type the following command at the command prompt:

   **netsh interface ip show config**

2. If you get an error, open **Network Connections**, and then resolve the problem by enabling or repairing the connection.
   - To enable the connection, right-click the connection and select **Enable**.
   - To repair the connection, right-click the connection and select **Repair**.
3. If the network connection is enabled, you should see a list showing values for IP address, subnet mask, default gateway, and so on.
4. If necessary, check each network setting to verify that it is configured properly.

# Group Policy does not apply when connecting remotely over a slow link

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certain policy settings do not apply over a slow link. In addition, Internet Control Message Protocol (ICMP) must be working in order for slow link detection to function. This topic describes how to adjust the default values for slow links and includes registry keys for manually creating registry settings to disable slow links.

**Cause**

By default, Group Policy defines a slow link as 500 kilobits per second or less. The following types of policy settings are always applied over a slow link:

- Security Settings
- Administrative Template settings

The following types of policy settings are not applied over a slow link:

- Software Installation
- Scripts
- Folder Redirection settings

| **Note** |
| --- |
| If a user connects to the network over a slow link with cached credentials, Group Policy will not be processed. |

**Solution**

You can change this setting in the computer configuration, the user configuration, or both. See the policy setting **Group Policy slow link detection**, located in **Computer Configuration\Administrative Templates\System\Group Policy processing**. For example, if you have a 300 kilobit per second network connection, you could specify 150 kilobits per second. This would mean that if bandwidth on the remote connection is less than what would be tested as a 150 kilobit-per-second connection at the time a user logs on, the link will be flagged as a slow link.

To ensure that Group Policy is applied remotely, select the **Logon using dialup connection** check box while using the **Logon** dialog box.

Even if Group Policy settings are configured to run scripts over slow links, the scripts might run so slowly that they exceed the configured time-out period. In this case the script will fail to complete and a UserInit event will be posted. In addition, Gpresult will tell you if a slow link was detected.

**To configure slow link detection**

1. Open the policy setting Group Policy slow link detection, located in **Computer Configuration\Administrative Templates\System\Group Policy processing**
2. In **Connection speed**, type a decimal number between **0** and **4,294,967,200** (0xFFFFFFA0) to indicate a transfer rate in kilobits per second. Any connection slower than this rate is considered to be slow. To disable slow link detection, enter **0** (all connections will be considered to be fast).

If you are trying to disable slow link detection when Group Policy is not applying (and therefore, you cannot configure the policy setting), you can manually create the registry value **"GroupPolicyMinTransferRate" (DWORD)** under the following keys and set each of them to 0.

- **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon**

- **HKEY_CURRENT_USER\Software\Microsoft\Windows NT\CurrentVersion\Winlogon**

**Troubleshooting ICMP**
You can verify that ICMP is not working by looking at the Userenv log, which will indicate that ping computer fails.

**See Also**
**Concepts**
Fixing Group Policy problems by using log files

# Unable to access domain controller

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic provides guidance for troubleshooting connectivity issues that can be traced to Domain Name System (DNS) configuration.

**Cause**

Group Policy application requires clients to access specified servers, including domain controllers and other servers such as share points and install points. Group Policy management also requires access to domain controllers. DNS is used to locate and identify these servers. In Windows Server 2003, Active Directory requires DNS support.

If the network is functioning, but clients or GPMC consoles are unable to locate the servers, there might be a problem with your network's DNS system.

**Solution**

To determine whether there is a problem in DNS, use the Ping tool against the computer using the NetBIOS name, and then use the Ping tool using the fully qualified domain name of the target computer. If the first Ping works but the second does not, this indicates that there probably is a DNS problem. Use Netdiag.exe to research the problem further.

Use Dcdiag.exe to troubleshoot domain controllers and use Netdiag.exe to troubleshoot client computers. These tools can help determine both server and client DNS misconfigurations. To run Dcdiag.exe, type **dcdiag /v** at a command prompt on a domain controller. To run Netdiag.exe, type **netdiag /v** at a command prompt on a workstation or member server.

To download these tools and obtain more information, see DCDiag and NetDiag in Windows 2000 Facilitate Domain Join and DC Creation on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43036).

**See Also**

**Concepts**

Fixing Group Policy problems by using log files

# Group Policy does not replicate

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic explains how to troubleshoot replication issues for Group Policy.

**Causes**

There can be a lag time after a change has been made on one domain controller before the change is replicated to all other domain controllers. The propagation and resolution of these changes throughout the network is an ongoing process called replication convergence.

Until changes to a GPO have been replicated to the domain controller that a client is accessing, that client will receive the earlier version of the GPO during Group Policy refresh. If you suspect both replication and Group Policy refresh issues, address the replication issue first. Then refresh Group Policy at the client.

Changes to the OU memberships of computers and users also need to be replicated before they can be reflected in Group Policy application at the client. For more information see GPO does not apply to a specific user or computer.

In general, it is best to use the same domain controller for all GPO editing or to agree a process — such as delegated administration of GPOs — to minimize the likelihood of the same GPO being edited on a different domain controller. If changes are made to the same GPO at two different domain controllers, the last change wins. In addition, if you delegate control of a specific GPO to a user group, members of that group might be unable to perform the delegated tasks until the permissions have been replicated to their domain controller.

Although Remote Procedure Call (RPC) is not required for general application of policy settings, FRS does require RPC to be working on the domain controller. In addition, anything that calls the Local Security Authority (LSA) requires RPC.

**Solutions**

There are several options for troubleshooting replication issues:

- The Group Policy container and Group Policy template are each assigned version numbers, which are incremented when the GPO is modified. Use GPOTool to verify that the versions are synchronized.
- Use Event Viewer to examine the directory service for event log on the domain controller. Active Directory replication errors appear with source=KCC.
- Use Event Viewer to examine the File Replication service event log on the domain controller. FRS errors appear with source=NTFRS.
- Verify that the SYSVOL share exists on the domain controller. You should be able to find *domain_controller_name*\SYSVOL, where *domain_controller_name* is the fully qualified domain name (not the NetBIOS name) of the domain controller.
- Use Gpotool.exe to determine if there is an inconsistency between Active Directory and SYSVOL versions of the same GPO across peer domain controllers. This information can help you determine if replication latency is causing the Windows client to not receive the correct policy. If you think this is the case, use Replmon.exe and Repadmin.exe (included with Windows 2000 Support Tools) to determine the replication partners for the domain controller that the client used as a source for the Group Policy and determine if replication is succeeding.
- To troubleshoot file replication issues, check the status of the Directory File Service links and targets. Group Policy requires Directory File Service. For more information, see "To check status of a DFS root, DFS link, or target" in Help and Support Center for Windows Server 2003.
- You can use the Sonar.exe tool to check the health of the SYSVOL share.

**See Also**

**Concepts**

Group Policy does not refresh

# Fixing Group Policy processing issues

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section explains Group Policy processing rules and points to troubleshooting tips. Group Policy objects (GPOs) are applied according to a set of standard processing rules, as explained below. You can apply exceptions to these rules and sometimes these exceptions — if not applied carefully — can lead to unexpected behavior.

- Standard processing rules
- Exceptions to processing rules
- Fixes for specific processing issues

**Standard processing rules**

This section provides details about the order in which Group Policy settings for users and computers are processed.

Group Policy settings are processed in the following order:

1. **Local Group Policy object—**Each computer has exactly one Group Policy object that is stored locally. The Local GPO processes for both computer and user Group Policy processing.

2. **Site—**Any GPOs that have been linked to the site that the computer belongs to are processed next. Processing is in the order that is specified by the administrator, on the **Linked Group Policy Objects** tab for the site in Group Policy Management Console (GPMC). The GPO with the lowest **link order** is processed last, and therefore has the highest precedence.

3. **Domain—**Processing of multiple domain-linked GPOs is in the order specified by the administrator, on the **Linked Group Policy Objects** tab for the domain in GPMC. The GPO with the lowest **link order** is processed last, and therefore has the highest precedence.

4. **Organizational units—**GPOs that are linked to the organizational unit that is highest in the Active Directory hierarchy are processed first, then GPOs that are linked to its child organizational unit, and so on. Finally, the GPOs that are linked to the organizational unit that contains the user or computer are processed.

   At the level of each organizational unit in the Active Directory hierarchy, one, many, or no GPOs can be linked. If several GPOs are linked to an organizational unit, their processing is in the order that is specified by the administrator, on the **Linked Group Policy Objects** tab for the organizational unit in GPMC. The GPO with the lowest **link order** is processed last, and therefore has the highest precedence.

This order means that the local GPO is processed first, and GPOs that are linked to the organizational unit of which the computer or user is a direct member are processed last, which overwrites settings in the earlier GPOs if there are conflicts. (If there are no conflicts, the earlier and later settings are merely aggregated.)

**Exceptions to processing rules**

The default order for processing settings is subject to the following exceptions:

- A GPO link may be enforced, or disabled, or both. By default, a GPO link is neither enforced nor disabled. For troubleshooting information, see GPO provides unexpected value.

- A GPO may have its user settings disabled, its computer settings disabled, or all settings disabled. By default, neither user settings nor computer settings are disabled on a GPO.

- An organizational unit or a domain may have **Block Inheritance** set. By default, **Block Inheritance** is not set. For troubleshooting information, see GPO provides unexpected value.

- A computer that is a member of a workgroup processes only the local Group Policy object.
- Loopback may be enabled. For troubleshooting information, see Loopback processing does not work.
- GPOs can be subject to security filtering or WMI filtering to limit the scope of application. For troubleshooting information, see Fixing Group Policy scoping issues.

**Fixes for specific processing issues**

From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- Policy setting is not supported
- GPO provides unexpected value
- Group Policy does not refresh
- Policy settings do not apply for software installation, scripts, or folder redirection
- Group Policy is not applied due to cached credentials
- Machine required resources are not available
- Loopback processing does not work

# Policy setting is not supported

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are trying to apply a policy setting to a computer running an earlier operating system or service pack, the policy setting may not apply.

**Cause**

Many Group Policy settings were introduced for Windows XP, Windows XP with Service Pack 2 (SP2) and Windows Server 2003 that are not supported by earlier operating systems. If a setting is not supported, Group Policy will set the registry key to the specified value, but it is effectively ignored because no application or component will read that registry key. Computers running Windows operating systems that are earlier than Windows 2000 do not support Group Policy.

**Solution**

Verify that the setting is supported by the client's operating system. The Group Policy Object Editor displays the operating systems where each policy setting is supported. For more information about policy settings, see the Group Policy Settings Reference for .adm files and Security Settings included with Windows XP Professional Service Pack 2 on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=15165).

# GPO provides unexpected value

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic explains common causes for unexpected GPO values due to GPO precedence, inheritance, and related scope of management issues.

**Causes**

- Conflict resolution applies to individual settings, not to entire GPOs. It could easily happen that one setting in a GPO encounters a conflict but all other settings in that GPO are applied.
- The GPO with the lowest link number prevails over other GPOs that the same site, domain or OU is linked to. You can use GPMC to change the order of links for a specific site, domain, or OU. (The links are a property of the site, domain, or OU; they are not a property of the GPO.)
- If a Group Policy link is disabled the GPO will not apply to users or computers within the container to which that GPO is linked.
- The **Enforce** setting is a property of the link between an Active Directory container and a GPO. It is used to force that GPO to all Active Directory objects within a container, no matter how deeply they are nested. The settings within a GPO that is enforced override other settings that would prevail because they are applied later. If there are conflicting settings in GPOs that are enforced at two levels of the hierarchy, the setting enforced furthest from the client prevails. This is a reversal of the usual rule, in which the setting from the nearest-linked GPO would prevail.

The actual effect of **Enforce** is to change the order of processing. The settings in an Enforced GPO are processed after all other GPOs settings are processed.

- The **Block Inheritance** setting applies to an entire Active Directory container. It blocks the inheritance of all GPOs except for those for which the link from the parent Active Directory object to the GPO has the **Enforce** setting enabled.
- Administrators who have set **Block Inheritance** on their domain or OU can still make explicit links to GPOs elsewhere in the domain, including GPOs that might otherwise be inherited. (Domains do not inherit GPOs from parent domains.) When **Block Inheritance** is applied at a domain level, it blocks GPOs that are linked to sites.

**Solutions**

Perform the following procedures to resolve this issue.

**To view or change precedence order of GPOs**

1. Open GPMC and click any site, domain, or organizational unit node.
2. Click the **Group Policy Inheritance** tab and examine the precedence order of the GPOs. Within each domain, site, and organizational unit, the link order controls when links are applied.
3. To change the precedence of a link, you can change the link order, moving each link up or down in the list to the appropriate location. The link with the higher order (with 1 being the highest order) has the higher precedence for a given site, domain, or organizational unit. For example, if you add six GPO links and later decide that you want the last one that you added to have highest precedence, you can move the GPO link to the top of the list.

**To check GPO links**

1. In GPMC, select the GPO you are troubleshooting, and then click the **Scope** tab. You will see the containers that are linked to the GPO and the status of those links.
2. To change the status of a link, click the **Details** tab, and then in **GPO Status**, choose an option. You can enable all settings, disable only computer settings, or disable only user settings.

# Group Policy does not refresh

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic describes how Group Policy refresh rules affect processing and explains how to collect refresh information using the Group Policy Management Console.

**Cause**

Group Policy refresh refers to the retrieval of GPOs by a client. During Group Policy refresh, the client contacts an available domain controller. If any GPOs have changed, the domain controller provides a list of all the appropriate GPOs, regardless of whether their version numbers have actually changed.

Replication and Group Policy refresh are both instances of lag-time issues: the system is working properly, but changes have not yet appeared at the client.

By default, GPOs are processed by client side extensions (CSEs) at the computer only if the version number of at least one GPO has changed on the domain controller that the computer is accessing. You can use policy settings to change this behavior.

Some CSEs process unchanged GPOs if the user's group membership has changed. At startup, Group Policy is refreshed, and computer settings are applied. Group Policy is refreshed and computer and user settings are applied in the following instances:

- When a user logs on
- When **gpupdate** is run at the client computer
- At the refresh interval, if one is configured at that computer (by default, domain controllers are refreshed every five minutes, and all other computers are refreshed every 90 minutes, with a random factor of up plus or minus 30 minutes)

**Solution**

To see the last time the GPOs from the computer's OU were processed, on the Group Policy Results report, under **Computer Configuration Summary**, and then under **General**, click the **Summary** tab.

To see the last time the GPOs from the user's OU were processed, on the Group Policy Results report, under **User Configuration Summary**, and then under **General**, click the **Summary** tab.

To collect Group Policy refresh information from clients and store them at a central location, use GPMonitor.exe. For more information and to download GPMonitor.exe, see Windows Server 2003 Resource Kit Tools on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=27766).

| **Note** |
| --- |
| Some types of settings can only be applied when a user logs on. These include Folder Redirection, Roaming Profiles, and Software Installation settings. If these settings are received when Group Policy is refreshed, the settings are evaluated, but they are not applied until the next time the user logs on.<br><br>If the computer is running Windows XP and these settings first reach the computer during logon, they might not be applied until the next time the user logs on. For some extensions, it might take two or three logons for the settings to be applied. |

A simple way to troubleshoot a suspected Group Policy refresh issue is to force the refresh by running **gpupdate** and either restarting the computer, or by logging off and logging on again. If Folder Redirection, roaming profiles, or Software Installation is involved and the computer is running Windows XP, run **gpupdate**, and then log off and log back on. You might need to log off and log back on more than once.

# Policy settings do not apply for software installation, scripts, or folder redirection

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy can be applied during startup and logon (synchronous processing) or as a background task after startup or logon has completed (asynchronous processing). Changes received during periodic Group Policy refresh or in response to the **gpupdate** command are processed asynchronously. On computers running Windows XP, Group Policy settings received during logon are also processed asynchronously by default, so that the logon is completed more quickly. This is also known as "fast boot" logging.

**Cause**

Software Installation and scripts processing must be applied during startup or logon. Folder Redirection assigned to the user must be applied during logon. By default, Windows XP logs a user on in asynchronous mode (that is, fast boot). Group Policy is then applied in the background after the user is logged on. However, when a new GPO setting for Software Installation, Scripts, or Folder Redirection arrives at a computer running Windows XP, the user has already logged on by the time Group Policy has been evaluated. It is too late to apply the Software Installation setting. In this case a flag is set so that the next time computer is rebooted or the user logs on Group Policy will be evaluated and applied before the startup or logon is completed.

**Solution**

In situations where you need for users to receive software, implement folder redirection, or run new scripts in a single logon, apply a GPO with the setting **Always wait for the network at computer startup and logon to the computer**. This setting is located under **Computer Configuration\Administrative Templates\System\Logon** in the Group Policy Object Editor. For this setting to take effect, Group Policy must be refreshed or the computer restarted. To ensure a consistent experience in a mixed desktop environment, consider enabling this setting for all your Windows XP computers.

**Table 1 Timing of Synchronous and Asynchronous Processing**

| By default, how is policy processed on the client? | At Startup | At Logon | At Policy Refresh |
|---|---|---|---|
| Windows 2000 | Synchronously | Synchronously | Asynchronously |
| Windows XP Professional | Asynchronously | Asynchronously | Asynchronously |
| Windows Server 2003 | Synchronously | Synchronously | Asynchronously |

| Note |
|---|
| Servers do not perform asynchronous processing. |

# Group Policy is not applied due to cached credentials

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic explains how cached credentials can prevent Group Policy from applying correctly.

**Cause**

When a user successfully logs on to the network, the credentials for that user can be cached on the local computer. If network connectivity problems prevent the user from being authenticated the next time the user logs on to the same computer, these cached credentials can be used to give the user access to resources on that computer. If the computer successfully connects to the network later, the cached credentials can be used to provide access to network resources, including GPOs that are received at the next Group Policy refresh.

If a domain controller is not available when the user logs on Group Policy cannot be refreshed at logon. In this case, new Group Policy settings will not be applied until a Group Policy refresh occurs while a domain controller is available.

When the logon is done with cached credentials and then a remote access connection is established, Group Policy is not applied during logon. For example, if users connecting through a VPN connection are logging on using cached credentials, folder redirection settings will not be processed because folder redirection policy can only be processed at user logon, not in the background refresh.

| **Note** |
| --- |
| Some user settings can only be applied during logon. These include roaming user profile path, Folder Redirection path, and Software Installation settings. If the user is already logged on when these settings are detected, they will not be applied until the next time the user is logged on. |

**Solution**

To avoid using cached credentials in a remote access connection, users should select the **Logon using dial-up connection** check box on the Windows Logon dialog box. When this occurs, both User and Computer Group Policy is applied, provided the computer is a member of the domain that the remote access server belongs to or trusts. However, computer-based software installation settings are not processed because, typically, computer policy would have been processed before the logon screen, but since no network connection is available until logon, the application of computer policy is done as background refresh at the time of logon.

# Machine required resources are not available

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This error indicates that memory allocation failed, or that some other necessary machine resource could not be obtained. Group Policy processing for the computer or user has failed and will continue to fail until this error is resolved.

**Cause**

This error is usually caused by low memory or low disk space on the machine at the time that Group Policy was processing.

The following Event Numbers trigger this error:

- 1002
- 1035
- 1063
- 1075
- 1078
- 1081
- 1082
- 1094
- 1107

**Solution**

To troubleshoot this problem, try one or all of the following:

- Review the event description to obtain additional details.
- Ensure that an adequate amount of memory is available on the machine. If the machine is very low on memory, take steps to free up the necessary resources.
- Ensure that adequate disk space is available on the machine. If the machine is very low on disk space, take steps to free up the necessary resources.

# Loopback processing does not work

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic explains how loopback processing alters the way GPOs are normally applied and provides tips for troubleshooting.

**Loopback Processing**

Loopback is an advanced Group Policy setting that is useful in scenarios where you are managing computers in closely managed environments, such as servers, kiosks, laboratories, classrooms, and reception areas. Loopback only works when both the user account and the computer account are in a Windows 2000 or later domain. Loopback does not work for computers joined to a workgroup. Setting loopback causes the User Configuration settings in GPOs that apply to the computer to be applied to every user logging on to that computer, instead of (in replace mode) or in addition to (in merge mode) the User Configuration settings of the user. This allows you to ensure that a consistent set of policies is applied to any user logging on to a particular computer, regardless of their location in Active Directory. Loopback is controlled by the **User Group Policy loopback processing mode** setting, which is located under **Computer Configuration\Administrative Templates\System\Group Policy** in Group Policy Object Editor (GPMC).

By default, a user's policy settings come from the set of GPOs that are applied to the user object in Active Directory. During Group Policy processing on the client, the Group Policy engine assembles an ordered list of GPOs from the site, domain, and all organizational units for that user object.

Loopback can be set to Not Configured, Enabled, or Disabled. In the Enabled state, loopback can be set to Merge or Replace. In either case the user only receives user-related policy settings.

- **Loopback with Replace**—In the case of Loopback with Replace, the GPO list for the user is replaced in its entirety by the GPO list that is already obtained for the computer at computer startup. The User Configuration settings from this list are applied to the user.

- **Loopback with Merge**—In the case of Loopback with Merge, the Group Policy object list is a concatenation. The default list of GPOs for the user object is obtained, as normal, but then the list of GPOs for the computer (obtained during computer startup) is appended to this list. Because the computer's GPOs are processed after the user's GPOs, they have precedence if any of the settings conflict.

To find out whether loopback processing was applied when Group Policy was evaluated on the client, look in the Group Policy Results report under **Computer Configuration\Administrative templates\System\Group Policy** on the **Settings** tab.

**Troubleshooting Loopback Processing**

- If the loopback processing is appropriate for this client you might need to educate the users so they know what to expect.

- If loopback is desired and it appears that it is not being applied, first verify the loopback policy setting (which is a computer configuration policy) has been applied to the computer through an appropriate GPO.

- For specific information about troubleshooting loopback replace in a cross-forest environment, see Loopback Replace does not work in cross forest environment.

# Loopback Replace does not work in cross forest environment

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Group Policy Loopback-Replace mode allows you to force a user to use a particular set of user GPOs instead of the ones that the user account is normally tied to. This is helpful in terminal server environments because it guarantees the integrity of your terminal servers. It is also useful in a public computer environment, such as a kiosk or library, where many different people log on to computers.

Cross-forest trust was introduced in Windows Server 2003. With a cross-forest trust you can connect two root domains, which gives both domains access to each other's domain. In addition, any new domains that are created under either root domain is automatically given the same access as the parent root domain. This enables different forests to communicate with each other, and it also enables users to log on to computers in either forests with ease.

Problems with Loopback-Replace in a cross-forest environment can be difficult to troubleshoot, due to the number of factors that can effect it. Even if your Group Policy is set up correctly, you can still have problems getting Loopback-Replace to work in a cross-forest environment.

**Event IDs**

The following event IDs are relevant to this section:

| Source | Event ID | Event Text |
|---|---|---|
| Userenv | 1529 | Loopback Group Policy processing will be applied, using the Group Policy objects scoped to the machine. |
| Userenv | 1109 | A User from a different forest logged onto this machine. Cross Forest Group Policy processing is disabled and loopback processing has been enforced in this forest for this user account. |

**Cause**

To understand how Loopback-Replace works, you need to understand the processing order of GPOs. When a computer is turned on, the computer applies the machine GPOs from the site, domain, and each nested OU that applies to it. When the user logs on to the computer, it processes the user GPOs in the same order. With Loopback-Replace you are changing how the computer processes the GPOs. Instead of the user's GPOs processing the normal way, they get the GPOs that you specify instead; regardless of which domain they are coming from or what GPOs are applied to them normally.

A number of things can cause Loopback-Replace to fail, and due to the complexity of cross-forest environments, the number of potential problems increases when the two technologies are used together. Potential causes of a Loopback-Replace failure could be that loopback was set up inappropriately, Organizational Units do not have the necessary user Group Policy for the Loopback to reference, or RPC is not running on the machine's domain controller.

If you have Loopback-Replace running in a cross-forest environment, you can have some additional problems. The main ones are tied to network health and connectivity. It is important to make sure that the domains are able to communicate with each other and still have full two-way trust. Firewalls between the two domains can also cause problems by blocking access from the machine that the user is attempting to log on to and the domain controller that user belongs to.

**Solution**

Verify that you have setup Loopback-Replace Mode correctly in Group Policy. To turn on Loopback-Replace Mode, create a GPO linked to the server OU you just created. Under **User Configuration**, click **Administrative Templates**, click **System**, click **Group Policy**, and then double-click the **User Group Policy Loopback Processing Mode** policy setting. Activate the policy by clicking **Enabled**. After you have enabled the setting, click **Mode**, click **Replace**, and then click **OK**.

If this does not fix the problem, try troubleshooting network connectivity between the target computer and the domain controller in the user's originating forest. In order for the computer to be able to contact the domain controller of the computer and for the user, RPC must be running. Run the rpings.exe and the rpingc.exe commands to determine the health of RPC on your network. You can also try running a network trace from the target machine to the user's domain controller. If possible, it would also be helpful to run a network trace at the same time from the domain controller to the target machine, which should help you determine the location of any network problems. A network trace will also help you determine if a firewall is causing this issue.

**Supported Operating Systems**

Cross-forest trust and Loopback-Replace are not available on all current Microsoft operating systems. Use the following table to determine whether your network is capable of running this setup.

| **Operating System** | **Supported** |
|---|---|
| Windows Server 2003 | Yes |
| Windows XP Service Pack 2 (SP2) | Yes |
| Windows XP Service Pack 1 (SP1) or earlier | No |
| Windows 2000 Service Pack 4 (SP4) | Yes |
| Windows 2000 Service Pack 3 (SP3) or earlier | No |

# Fixing Group Policy scoping issues

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section explains scoping rules and points to troubleshooting tips.

**Scoping the GPO**
Determining which users and computers will receive the settings in a GPO is referred to as "scoping the GPO." Scoping a GPO is based on three factors:

- The site, domain, or organization unit where the GPO is linked
- The security filtering on the GPO
- The WMI filter on the GPO

You can use the site, domain, and OU links from a GPO as the primary targeting principle for defining which computers and users should receive a GPO. You can then use security filtering and WMI filtering to further reduce the set of computers and users to which the GPO will apply.

If incorrectly configured, security filtering or WMI filtering can lead to unexpected results. This section provides ways to troubleshoot these areas.

From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- Policy settings incorrectly applied or denied due to WMI filtering
- Policy settings incorrectly applied or denied due to security filtering

# Policy settings incorrectly applied or denied due to security filtering

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are using security filtering, you may need to check permissions on the GPO if you are encountering unexpected behavior.

**Cause**

You can use Group Policy to provide or deny access to programs and data in your network, and to enforce policies regarding computer configuration based on assigned privileges and security group memberships. You accomplish this by using the access control functionality built in to Windows 2000 Server and Windows Server 2003 domains and is known as security filtering.

You can restrict application of all the settings in a GPO on the basis of security group memberships by setting a security filter on that GPO. If the computer account or user account does not meet the security filtering criteria, the entire GPO will be denied at that client. For example, you can assign special settings to all the administrators in a portion of the hierarchy by setting the security filter to apply the GPO to all administrators, and then linking the GPO to the highest node in the portion of the hierarchy where you want the settings to apply. All users in that portion of the hierarchy will receive the GPO, but only members of the administrators group will be affected by it.

A security group, user or computer must have both Read and Apply Group Policy permissions for a policy to be applied. By default, all users and computers have these permissions for all new Group Policy Objects. These permissions are inherited from their membership in the implicit group Authenticated Users. An authenticated user is any user (or computer) that has logged on to the domain and been authenticated.

**Solution**

To see the security groups that were in effect when Group Policy was applied to a specific computer, look in the Group Policy Results report for that computer. Under both Computer Configuration Summary and User Configuration Summary, expand Security Group Membership when Group Policy was applied.

**To check security filtering on a GPO**

1. In GPMC, open **Group Policy Objects** node, select the GPO you are troubleshooting, and then in the right pane select the **Scope** tab. The **Security Filtering** and **WMI Filtering** panels show the current filtering configuration.
2. To see the exact set of permissions for users, groups and computers, select the **Delegation** tab and then click **Advanced**. Select the security group, user or computer you want to review. Keep the following in mind:
   - If the policy object should be applied to the security group, user or computer, the minimum permissions should be set to allow **Read** and **Apply Group Policy**.
   - If the policy object should not be applied to the security group, user or computer, the minimum permissions should be set to allow **Read** and deny **Apply Group Policy**.

If a GPO is incorrectly denied or applied due to security filtering because the user or computer had different security group memberships than expected, use Active Directory Users and Computers to check and, if necessary change, the security group memberships.

When restricting the application of a GPO, be sure to remove **Authenticated Users**. Otherwise all users will always be affected by the GPO.

Computers are members of the Authenticated Users group. If you remove **Authenticated Users** from the list on the **Scope** tab and you want the GPO to apply to a computer, you must specifically ensure that the computer belongs to a group that is included in the **Security Filtering** section on the **Scope** tab.

# Policy settings incorrectly applied or denied due to WMI filtering

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

WMI filters can affect whether or not a particular GPO is being processed. If a WMI returns a false value, the GPO that it is linked to will not be processed.

**Cause**
If the GPO is linked to a WMI filter, the queries in the WMI filter are evaluated against the data provided by WMI on the client. Such data can include hardware and software inventory, settings, and configuration information.

- If all of the criteria are true, the GPO is applied.
- If any of the criteria is false, the GPO is denied.

If a WMI filter is deleted, the links to the WMI filter are not automatically deleted. If there is a link to a non-existent WMI filter, the GPO with that link will not be processed until the link is removed or the filter is restored.

**Note** Only Windows XP and Windows Server 2003 support WMI filtering. If the computer is running an earlier operating system (such as Windows 2000), the WMI filter is ignored and the GPO is applied.

**Solution**
GPMC shows any WMI filters linked to GPOs that are processed by the computer and the result of the filter as it was evaluated for that computer (or user).

**To check WMI filters using GPMC**

1. Open GPMC, click the GPO that you want to troubleshoot, and then click the **Scope** tab. The **WMI Filtering** section shows whether a WMI filter is linked to the GPO.
2. If you need to edit the filter, click **Open**, and then click **Edit**.

# Fixing Group Policy structural issues

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This page explains linking rules and points to troubleshooting tips.

**Active Directory and Group Policy**
Active Directory organizes objects by sites, domains, and OUs. Domains and OUs are organized hierarchically, making the containers and the objects within them easy to manage. The settings defined in a GPO can only be applied when the GPO is linked to one or more of these containers.

By linking GPOs to sites, domains, and OUs, you can implement Group Policy settings for as broad or as narrow a portion of the organization as you want. GPO links affect users and computers in the following ways:

- A GPO linked to a site applies to all users and computers in the site.
- A GPO linked to a domain applies directly to all users and computers in the domain and by inheritance to all users and computers in child OUs. Policy is not inherited across domains.
- A GPO linked to an OU applies directly to all users and computers in the OU and by inheritance to all users and computers in child OUs.
- When a GPO is created, it is stored in the domain. When the GPO is linked to an Active Directory container, such as an OU, the link is a component of that container, not a component of the GPO.

**Fixes for specific structural issues**
From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- GPO does not apply to a specific user or computer
- Critical files on domain controllers or client computers are missing or corrupted
- Trust relationships are not working
- Group Policy does not apply after migration

# GPO does not apply to a specific user or computer

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic explains the importance of correct Organizational Unit memberships and GPO linking and shows how to check these issues using GPMC and other tools.

**Cause**
GPOs are applied to a client only if they are linked to a site, domain, or OU to which the computer or the user at that computer belongs.

For troubleshooting purposes, you need a solid understanding of your organization's Active Directory structure and the Group Policy inheritance and filtering rules. With this information and the Resultant Set of Policy (RSoP) functionality in Windows Server 2003 and Windows XP, you can manipulate your Active Directory structure and your Group Policy links and filters to deliver targeted settings to the users and computers in your organization. The same information is needed to troubleshoot situations where these manipulations produce an unexpected result.

**Solution**
Check Active Directory Users and Computers to see what site, domain, and OU the user and the computer are in.

In GPMC, expand the Active Directory containers that contain the affected client. In the navigation pane, scan the list of GPOs for each container for disabled links.

GPOs are filtered according to the Active Directory groups that the users and computers belong to. The Active Directory objects in which you place your Active Directory groups and the ways you group users or computers affect how GPOs can be distributed and applied.

Active Directory and FRS replication lag can affect either part of the GPO.

If you have an OU that contains other OUs and you remove Read permissions to the parent OU, then no policy will be processed by computers or users in that OU hierarchy.

If there are conflicting settings in the GPOs that apply to the client, they are resolved according to the Group Policy inheritance rules.

**Adding a User or Computer to an OU**
When a user or computer is added to an OU, two things need to happen before the GPOs that the new OU links to are applied to the client:

- The new OU assignment must be replicated to the client's domain controller.
- After the replication is complete, you must either log off and log back on again if the user account moved to the new OU, or restart the computer if the computer moved to the new OU.

# Critical files on domain controllers or client computers are missing or corrupted

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy might not apply due to missing or corrupted files.

**Cause**
Group Policy information is contained in files on both the domain controllers and the clients. If any of these files are missing or corrupted, only some or none of the policies can be applied.

**Solution**
Examine files on the domain controller and the client.

**Domain controller**
Use GPOtool.exe to check for the presence and integrity of the following files in the SYSVOL share and its subfolders on the domain controller.

- Files in the Group Policy template.
- Registry.pol (Search %windir%\debug\usermode\UserEnv.log for references to this file). This file is used for processing administrative templates through the registry CSE.

Check the Policy Events tab in the Group Policy Results report for events that point to problems accessing files. When a process is unable to access a file, it generates an event. Whether this event is recorded depends on the event severity and whether verbose logging is enabled for that process.

You can also use Event Viewer on the client to view the Application logs, and you can enable and view verbose logging for UserEnv and for specific CSEs to Group Policy. For more information, see Fixing Group Policy problems by using log files.

**Client**
On the client, check for the presence and integrity of the following files in the %windir%\system32 folder. Replace suspect or files missing from the CD for the client's operating system. The System File Checker (Sfc.exe) can be used to scan all protected files to verify their versions.

- UserEnv.dll. Used for core Group Policy processing including registry CSE and profiles.
- Dskquota.dll. Used for Disk Quota CSE processing.
- Fdeploy.dll. Used for Folder Redirection CSE processing.
- Gptext.dll. Used for Scripts, IP Security, QoS Packet Scheduler, and Wireless settings processing.
- Appmgmts.dll. Used for Software Installation CSE processing.
- Scecli.dll. Used for Security Settings CSE processing.

**See Also**
**Concepts**
Fixing Group Policy problems by using log files

# Trust relationships are not working

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic explains how to check if GPOs are not applying because trust relationships are not working.

**Cause**
You can link GPOs across domains, provided there is a trust relationship between the domains. If the trust relationship is broken, clients will be unable to access the GPO and related files. You might also encounter performance issues with links across domains.

You cannot link a GPO in one forest to a site, domain, or OU in a different forest.

**Solution**
If the GPO cannot be applied due to lack of trust, it will appear in the list of Denied GPOs and the reason given will be Inaccessible.

Use Active Directory Domains and Trusts or nltest.exe to verify the trust relationship, and to if repair it if necessary.

If you are not concerned about the identical GPO being applied in both domains, copy the GPO to the domain with the Active Directory containers you want to link to it.

**See Also**
**Concepts**
Administering Domain and Forest Trusts

# Group Policy does not apply after migration

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A Group Policy object often does not apply as expected after it has been migrated from another domain or forest — from a test environment to a production environment, for example.

**Cause**
Simply backing up and importing the GPO often does not produce the results you want because GPOs include domain-specific information such as security principals and UNC paths.

**Solution**
GPMC includes migration support, including a migration table editor, to address these issues. If a GPO that worked in one forest is not working as expected in the new forest, you might need a migration table. You also might need to back up and import the GPO instead of copying it.

- Check for a trust relationship between the source domain and the target domain. It there is trust, copy the GPO; if not, import it.

- If the GPO has references to security principles or UNC paths, use a migration table. The migration table editor that is included with GPMC provides error checking. However there is still a possibility of mistyped UNC paths.

For more information see the white paper, Migrating GPOs Across Domains with GPMC on the Microsoft Web site at (http://go.microsoft.com/fwlink/?LinkId=14321).

# Fixing Administrative Template policy setting problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic highlights issues related to processing of the Administrative Templates extension, also known as the registry extension.

**Administrative Templates**
Administrative Template settings take the form of **true policies** or **preferences**. Preferences can be deployed using Group Policy, but they cannot be enforced to the same degree as true policies. For this reason they are hidden by default in the Group Policy Object Editor.

Users can change preferences, but they cannot change true policies.

True policies are more secure because they are stored in secured registry hives.

Changes to true policies override but do not overwrite user preferences. If the policy is later removed, the user setting will again prevail.

If the GPO ceases to apply to the user or computer, policies no longer apply but preferences remain. This occurs if the user or computer moves out of the site, domain, or OU that the GPO is linked to, or if the GPO is deleted.

If the setting you are using is a preference rather than a true policy be aware that because preferences can be overwritten but are not removed, the end user might see their behavior as unpredictable. As a result, there might be perceived problems even when preferences are behaving as intended. Troubleshooting preferences requires knowledge about changes to both the GPO and the preferences set by the user.

The registry CSE writes the value for the setting to the registry. Some settings take effect as soon as they are written to the registry, but others take effect only at startup or logon. If you are not seeing the expected results and the Group Policy Results report shows that the policy has been applied, restart the computer or log off and log back on.

# Administrative Templates CSE is processed with errors

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

One or more registry-based policy settings are not applied to the computer or user. As a result some components and applications that are managed through registry-based Group Policy will not function according to administrator defined policies.

**Cause**
This error indicates that there is a missing file in a Group Policy object (GPO). GPOs reside on the domain controller's SYSVOL share, and a local GPO also resides on the local computer's system drive.

The event indicates that it was trying to access a file named "registry.pol" The file registry.pol might be corrupt and the event log describes the location of the corrupted file.

This error can also occur in the event of an improper registry configuration. Group Policy stores registry-based policy settings in the registry. If these registry keys have Access Control Lists (ACLs) that prevent the system from writing to this values, this failure can occur.

The following Userenv event numbers trigger this alert:

- 1019
- 1020 : Windows cannot create registry key <registry key>. (The parameter is incorrect. ).
- 1022
- 1043 : Windows cannot access the registry information at C:\WINDOWS\System32\GroupPolicy\Machine\registry.pol. (Access is denied).
- 1096 : Windows cannot access the registry policy file, <path to .pol file>. (<error description>).

**Solution**
If a problem with the registry.pol is indicated, verify that the file is present at the indicated path on all domain controllers. The Windows Server 2003 Group Policy Infrastructure whitepaper describes the layout of Group Policy information about domain controllers. Domain replication problems are indicated if replicas of the file are missing or if the files cannot be accessed by the targets of the GPO. Troubleshooting the File Replication service (FRS) and the Distributed File System (DFS) of the domain controller's SYSVOL share can diagnose the cause of the replication errors. Replication is not an issue if the file is part of the local GPO.

If the file is present, it might be corrupt. To determine this, use the Group Policy Object Editor snap-in to edit the GPO in which the file resides. If the Administrative Templates node under either User Configuration or Computer Configuration displays an error, you can safely assume that this file is corrupted.

To recover a corrupt or missing registry.pol file, use the Group Policy Management Console (GPMC) to restore the GPO from a backup.

If the file is missing due to replication errors, the replication issue with FRS or DFS must be addressed. After you resolve the replication problems, the error condition will also be resolved. The SONAR tool can be used to diagnose FRS problems.

If an improper registry configuration caused the error, examine the policy registry keys of the trees HKEY_LOCAL_MACHINE (and HKEY_CURRENT_USER if the problem occurred with a specific user logged in). Deleting these subtrees can fix the problem, but be aware that any edits to the registry can cause unexpected problems and must never take place without backup precautions.

# Group Policy Object Editor loads with "[strings] section is too long and has been truncated" message

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Some of the new Administrative templates that are included with the SP2 updated for Windows XP will not fully load on earlier versions of gpedit. If you attempt to load one of the SP2 Administrative templates from an older version of gpedit, you will get multiple error messages about "The following entry in the [strings] section is too long and has been truncated".

**Cause**

This is caused by the earlier versions of gpedit incorrectly handling the `"#if ver >=5"` construct in the new Administrative templates. Once opened through gpedit, you get several error messages. If you click OK on all of the error messages, the Administrative template will load correctly. However, any policy settings that reference the new LISTBOX syntax will not be displayed. To fix this, update the version of gpedit that is installed on the computer. You can also update your computer's operating system to Windows XP SP2.

**Solution**

To fix this problem, update your software to Windows XP SP2 or Windows Server 2003 SP1. To download hot fixes for all operating systems, see "The following entry in the [strings] section is too long and has been truncated" error message when you try to modify or to view GPOs in Windows Server 2003, Windows XP Professional, or Windows 2000 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43056).

# Some policy settings do not show in GPMC reports

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The behavior of the LISTBOX ADDITIVE keyword was modified for the administrative template files in Windows XP with Service Pack 2 (SP2).

There is a problem with this new functionality that occurs when you attempt to run GPMC Reporting for Policy Settings. GPMC Reports does not properly format the information that it receives for policy settings that use the LISTBOX ADDITIVE functionality.

**Cause**

This problem is caused by how GPMC handles the new Administrative templates that were added in Windows XP SP2. It does not recognize entries in the Administrative templates that are enclosed by the `"#if version >= 5" / "#endif"` statement. Because it is unable to recognize those entries, when you run a Reporting for Policy Settings modeling report, GPMC displays all of the data from those entries at the bottom of the document under Extra Registry Settings instead of User Configuration or Computer Configuration. Also, it does not show the display names for these settings, although it does provide information about the name of the registry key and the state of the setting.

**Solution**

This problem affects administrative reports only. It does not effect the functionality of the LISTBOX ADDITIVE keyword. Currently, there is no fix for this problem in GPMC.

# Custom ADM settings do not display in the Group Policy Editor

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certain settings defined in a custom Administrative Template (ADM) file are not displayed in the Group Policy Object Editor by default.

**Cause**

ADM files define registry settings (also referred to as individual policy settings) that are displayed in the Group Policy Object Editor under the Administrative Templates node for both Computer and User Configuration. Two different types of settings can be defined.

In Windows 2000 and later, all policy settings set registry keys and values in one of the following locations:

- HKLM\Software\Policies (preferred location)
- HKLM\Software\Microsoft\Windows\CurrentVersion\Policies
- HKCU\Software\Policies (preferred location)
- HKCU\Software\Microsoft\Windows\CurrentVersion\Policies

Policy settings that are stored in these specific locations of the registry are known as true policies. This type of setting is displayed in the Group Policy Object Editor with a blue icon.

Preferences are settings that are stored outside these specific locations of the registry. If a preference is not reversed before the Group Policy object is disabled, the setting on the client computer remains, a process referred to as "tattooing." These settings are identifiable in the Group Policy Object Editor with a red icon.

The Group Policy Object Editor identifies policy settings that are preferences and does not display them by default. The default policy templates (Conf.adm and System.adm) loaded into the Group Policy Object Editor by default do not contain any preferences, and therefore none are displayed.

**Solution**

**To view preference settings for all ADM files in the Group Policy Object Editor**

1. Right-click any administrative template node and then click **View**.
2. Click **Filtering** and in the filtering dialog box, clear the check box for **Only show policy settings that can be fully managed**.
3. Click **OK**.

For more information, see article ID 228723 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=48209).

# Fixing Security Settings Problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses how to troubleshoot problems related to Security Settings policy.

**Security Settings CSE Background**
The Security Settings CSE provides client side interfaces to the security configuration engine and performs Resultant Set of Policies (RsoP) logging during policy propagation.

The binary file that contains the Security CSE is Scecli.dll. This name will usually appear in Events, error messages, and log entries generated and logged during processing of Security GPOs. Scecli manages application of the Security policy settings that appear under the Security Settings node in the Group Policy Object Editor. Scecli is responsible for the following areas:

- Account Policies
- Local Policies
- Event Log
- Restricted Groups
- System Services
- Registry
- File System

When the Security CSE is notified by the Group Policy engine, it is provided a list of GPOs to apply. The Security CSE then copies the gpttmpl.inf file from the folder structure of each policy in the Sysvol. It copies that file locally to the hidden folder %SYSTEMROOT%\Security\Templates\Policies. The settings are read from the gpttmpl.inf in the Sysvol and written to an intermediary file named tmpgptfl.inf. Once the copy has completed successfully the file is copied off and is named incrementally starting from gpt00000.inf. If the GPO is linked to the domain then the cached template will be named with the .dom extension, otherwise it will be named with the .inf extension.

This is done because some settings are only applied if they are linked to the domain. For more information, see Group Policy Application Rules for Domain Controllers on the Microsoft Web site at (http://go.microsoft.com/fwlink/?LinkId=39978). The templates are then applied from the cached location in order from least to greatest. This means that the gpt00000.inf will be applied before gpt00001.inf and that gpt00001.inf will have a higher precedence in the case of a conflict.

**Fixes for specific Security Settings issues**
From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- Replication of Group Policy settings between domain controllers fails
- Password Policy settings changes do not take effect
- Account Policy templates are applied incorrectly
- Security Settings policies are propagated with warnings
- Scecli.dll errors occur when opening Account Policies or Local Policies
- Default Group Policy objects become corrupted: disaster recovery
- Windows Security Settings remain in effect after removal

# Replication of Group Policy settings between domain controllers fails

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Replication of Group Policy settings between domain controllers fails.

**Cause**
A replication problem may occur if you assign incorrect permissions to the %SystemRoot%\Winnt\Sysvol folder or when you assign improper groups to the **Bypass Traverse Checking** User Rights Assignment. The replication issue also occurs if the sysvol share permissions are too restrictive.

**Solution**
To troubleshoot this issue:

- Ensure that the permissions for the %SystemRoot%\Winnt\Sysvol folder and sub-folders are set as described in Event ID 1000, 1001 Is Logged Every Five Minutes in the Application Event Log on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35272).

- In Group Policy Management Console, right-click the **Default Domain Controllers Policy** GPO (in the Domain Controllers OU), click **Edit**, and verify that the **Bypass traverse checking** security setting (in Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment) contains the following default groups:

    - Authenticated Users

    - Everyone

    - Administrators

- Refresh security policy by running **gpupdate /force**.

**See Also**
**Other Resources**
Event ID 1000, 1001 Is Logged Every Five Minutes in the Application Event Log

# Password Policy settings changes do not take effect

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Password Policy settings do not take effect.

**Cause**

If you have made changes to password policy settings and they do not get applied, this may mean that the policy settings have not been updated.

**Solution**

### To trigger a policy refresh

- Click **Start**, click **Run**, type **gpupdate**, and then click **OK**. This refreshes policy settings, and it updates the changes to password policy on the computer.

**Gpupdate** refreshes local Group Policy settings and Group Policy settings that are stored in Active Directory, including security settings. This command supersedes the now obsolete **/refreshpolicy** option for the **secedit** command.

# Account Policy templates are applied incorrectly

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An administrator applies a security template that sets the account policies to something the administrator does not want to use, such as settings that cause users to change their password when the administrators do not want them to. Setting the Account Policies to Not Defined does not resolve the problem.

**Cause**

Incorrectly applied Account Policy settings. In this case, the administrator applied a security template that had unintended effects.

**Solution**

To resolve misapplied templates issues, do the following:

1. Re-apply the default settings for Account Policies. The default settings in the **Default Domain Policy GPO** for a clean installation of Windows Server 2003 are located in %windir%\inf\dcfirst.inf. The **Dcfirst.inf** security template is the template that is used to populate the security settings in the **Default Domain Policy GPO** when the first domain controller is installed in the domain. Importing this template into the **Default Domain Policy** (and checking the **Clear this database before importing** option) restores the security settings back to a clean install state.

   The same information applies to the **Defdcgpo.inf** security template which contains the default security settings in the **Default Domain Controller Policy GPO** (for a clean installation of Windows Server 2003). **Defdcgpo.inf** is located in **%windir%\inf**.

   | **Note** |
   | --- |
   | For domain accounts, the account policy must be defined in the Default Domain Policy or in a new policy that is linked to the root of the domain and given precedence over the Default Domain Policy, which is enforced by the domain controllers that make up the domain. If more than one GPO containing account policy settings is linked at the domain level, the domain's account policy consists of the cumulative policy settings from all the domain-linked GPOs. A domain controller always obtains the account policy from a GPO linked to the domain, which by default is the **Default Domain Policy GPO**. This behavior occurs even if there is a different account policy applied to the OU that contains the domain controller. By default, workstations and servers joined to a domain (such as member computers) also receive the same account policy for their local accounts. However, local account policies for member computers can be differentiated from the domain account policy by defining an account policy for the OU that contains the member computers. Modifying the Default Domain Policy is not recommended. If you need to set some account policy that differs from that in the **Default Domain Policy GPO**, you can create a new GPO and link to the root of domain, set the policy you want to use, and assign it higher precedence than **Default Domain Policy GPO**. |

2. Wait for replication changes to propagate to the domain controller.

3. Verify that the policy applied to the clients.

**To reapply default security settings using the Windows interface**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. Do one of the following:
   - To create a new console, on the **File** menu, click **New**.
   - To open an existing console, on the **File** menu, click **Open**, click the console that you want to open, and then click **Open**.
3. On the **File** menu, click **Add/Remove Snap-in**, and then, in **Add/Remove Snap-in**, click **Add**.
4. Click **Security Configuration and Analysis**, click **Add**, click **Close**, and then click **OK**.

5. In the console tree, right-click **Security Configuration and Analysis**, and then click **Open Database**.
6. In **File name**, type the file name, and then click **Open**.
7. Do one of the following:
   - For a domain controller, in the console tree, right-click **Security Configuration and Analysis**, click **Import Template**, click **Look in** in the **Import Template** dialog box, and then select **dcfirst.inf** in the **%windir%\inf** folder.
   - For other computers, in the console tree, right-click **Security Configuration and Analysis**, click **Import Template**, and then click **Setup security**.
8. Select the **Clear this database before importing** check box, and then click **Open**.
9. In the console tree, right-click **Security Configuration and Analysis**, and then click **Configure Computer Now**.
10. Do one of the following:
    - To use the default log specified in **Error log file path**, click **OK**.
    - To specify a different log, in **Error log file path**, type a valid path and file name, and then click **OK**.
11. When the configuration is done, right-click **Security Configuration and Analysis**, and then click **View Log File**.

**Important**

Applying the entire setup security template is a drastic measure that should be avoided; it is best to limit the settings being reset to the smallest possible number. Instead, use the **secedit** command-line tool with the **/areas** switch to apply default settings for specific areas.

**Note**

Different permissions are required to perform this procedure, depending on the environment in which you reapply default security settings.

If you reapply default security settings to your local computer: To perform this procedure, you must be an Administrator or a member of the Administrators group on the local computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using **Run as** to perform this procedure.

If you reapply default security settings to a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have a local Administrator account or an account in the local Administrators group.

**To open Security Configuration and Analysis**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. Do one of the following:
   - To create a new console, on the **File** menu, click **New**.
   - To open an existing console, on the **File** menu, click **Open**, click the console that you want to open, and then click **Open**.
3. On the **File** menu, click **Add/Remove Snap-in**, and then, in **Add/Remove Snap-in**, click **Add**.
4. Click **Security Configuration and Analysis**, click **Add**, click **Close**, and then click **OK**.

The default path for the log file is: **systemroot\Documents and Settings\UserAccount\My Documents\Security\Logs**.

When you reapply default security settings, all settings that are defined in **Setup security.inf** are set as the template specifies, but other settings that are not defined in the template may persist.

**To reapply default security settings using a command line**

1. Click **Start**, point to **All Programs**, point to **Accessories**, and then click **Command Prompt**.
2. For a server or workstation, type:

   **secedit /configure /DB FileName /CFG "%windir%\Security\Templates\Setup security.inf" [/overwrite][/areas** *Area1 Area2* **...] [/log** *LogPath* **] [/quiet]**

3. For a domain controller, type:

   **secedit /configure /DB** *FileName***/CFG "%windir%\inf\dcfirst.inf" [/overwrite][/areas** *Area1Area2***...] [/log** *LogPath***] [/quiet]**

The following table describes the **Secedit** arguments used in the procedure.

| Argument | Description |
|---|---|
| **/DB** *FileName* | Required. Provides the path to a database that contains the security template that should be applied. To create a new database, type the database file name and path. |
| **/CFG "%windir%\Security\Templates\Setup security.inf"** | Specifies the Setup Security.inf template that contains the default security settings. |
| **/overwrite** | Specifies that the database should be emptied prior to importing the security template. If this parameter is not specified, the settings in the security template are accumulated into the database. If this parameter is not specified and there are conflicting settings in the database and the template being imported, the template settings win. |
| **/areas** *Area1 Area2* **...** | Specifies the security areas to be applied to the system. If this parameter is not specified, all security settings defined in the database are applied to the system. To configure multiple areas, separate each area by a space. The following security areas are supported:<br>• **SECURITYPOLICY** Includes account policies, audit policies, event log settings, and security options.<br>• **GROUP_MGMT** Includes Restricted Group settings.<br>• **USER_RIGHTS** Includes user rights assignment.<br>• **REGKEYS** Includes registry permissions.<br>• **FILESTORE** Includes file system permissions.<br>• **SERVICES** Includes System Service settings. |
| **/log** *LogPath* | Specifies a file in which to log the status of the configuration process. If not specified, configuration data is logged in the Scesrv.log file, which is located in the %windir%\Security\Logs folder. |
| **/quiet** | Specifies that the configuration process should take place without prompting the user. |

**To view the complete syntax for the secedit command**

- At a command prompt, type: **secedit /?**

# Security Settings policies are propagated with warnings

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Windows 2000, Scecli.dll events 1202 and 1000 may occur every 5 minutes on domain controllers and every 90 minutes on clients in some situations. As a result, security settings propagate with warnings and the Security Settings Group Policy client side extension may fail.

**Cause**

These Scecli.dll events might occur in cases such as the following:

- An administrator has installed a program which creates user accounts and assigns rights to those user accounts. Later, the administrator removes the program, which deletes the user accounts, but does not remove the rights from Group Policy before the accounts are deleted.

- An administrator adds a user account and assigns rights to the account. Later, the administrator deletes the account, but does not remove the account from the user rights policy.

These events can indicate a failure to resolve a security account to a security identifier (SID) or a broken trust relationship.

**Solution**

There are several different issues that make up the Scecli.dll event ID 1202 category of warnings; these are collectively referred to as "1202 events." The error code in the event log warning specifies the type of problem that was encountered.

To see diagnostic and troubleshooting information about the 1202 events:

- Determine the error code from the event log notification (for example 0x4b8).

- Find the error code section in Troubleshooting SCECLI 1202 Events on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=39979), and then follow the troubleshooting steps in that section.

- For information about security policies propagating with warning 0x534, see Error Message: Security Policies Are Propagated with Warning. 0x534 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=39980).

# Scecli.dll errors occur when opening Account Policies or Local Policies

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Scecli.dll errors may occur when you open Account Policies or Local Policies in Local Computer Policy.

**Cause**

When you attempt to open Account Policies or Local Policies in Local Computer Policy, you may receive the following error messages:

- "Windows cannot open the local policy database. An unknown error occurred when attempting to open the database."
- SceCli Event ID 1202: "Security policies are propagated with warning 0x4b8: An extended error has occurred."

This error message can occur if the Local Group Policy log files are corrupted. This usually indicates that the secedit.sdb file is corrupted.

**Solution**

**To verify and restore secedit.sdb**

1. Determine if secedit.sdb is corrupt by running **esentutl /g** %windir%\security\database\secedit.sdb.
2. Try to restore secedit.sdb from a backup before the problem started.
3. If secedit.sdb is corrupt, attempt a soft recovery first by running **esentutl /r edb** in the %windir%\security directory.
4. If soft recovery fails, attempt a repair with **esentutl /p** %windir%\security\Database\secedit.sdb. Then delete the log files %windir%\security*edb* and %windir%\security*log.
5. Rebuild secedit.sdb. Delete the following files: %windir%\security*edb*, %windir%\security*log, and %windir%\security\database\secedit.sdb.
6. Restart the computer.

| **Note** |
| --- |
| If you perform a repair with **esentutl /p** on a computer running Windows 2000, the local security policy settings are lost and you must recreate them.<br><br>If you perform a repair with **esentutl /p** on a computer running Windows Server 2003, persistent security settings (also referred to as tattoo table) are lost so when Group Policy settings are undefined, the server's settings do not revert to their original values. |

For more information, see the following articles which describe other known issues that can cause the 0x4b8 error, see the following:

- Event ID 1000 and 1202 After Configuring Policies on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=39982).
- ESENT Event IDs 1000, 1202, 412, and 454 Are Logged Repeatedly in the Application Event Log on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=39984).

# Default Group Policy objects become corrupted: disaster recovery

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The default domain GPOs become corrupted and there are no GPO backups for the Default Domain Policy GPO and Default Domain Controller Policy GPO.

**Cause**

The default domain GPOs are corrupted (for example, because of misconfiguration) and you do not have backed up versions of the Default Domain Policy GPO or the Default Domain Controller Policy GPO.

**Solution**

If you are in a disaster recovery scenario, you may consider using the Dcgpofix tool. If you use the Dcgpofix tool, it is strongly recommended that as soon as you run it, you review the security settings in these GPOs and manually adjust the security settings to suit your requirements.

Dcgpofix restores the default Group Policy objects to their original default state after initial installation of a domain controller. The Dcgpofix tool recreates the two default Group Policy objects and creates the settings based on the operations that are performed only during Dcpromo. It is important to understand that Dcgpofix does not restore the security settings to the state they were in before you run Dcpromo.

**Important**

The Dcgpofix tool is intended for use only as a last-resort disaster-recovery tool.

To create regular backups of the default domain and all other GPOs, you must use Group Policy Management Console (GPMC).

It is also recommended that you backup the Sysvol directory with a regularly scheduled backup procedure.

**To run Dcgpofix**

- Type the following at the command prompt: **dcgpofix [/ignoreschema][/target: {domain | dc | both}]**

Where:

**/ignoreschema** is an optional parameter. If you set this parameter, the Active Directory schema version number is ignored.

**/target: {domain | dc | both}** is an optional parameter that specifies the target domain, domain controller, or both. If you do not specify **/target**, dcgpofix uses both by default.

**Note**

Dcgpofix.exe is located in the C:\Windows\Repair folder.

You must be a domain or enterprise Administrator to use this tool.

Dcgpofix.exe checks the Active Directory schema version number to ensure compatibility between the version of Dcgpofix you are using and the Active Directory schema configuration. If the versions are not compatible, Dcgpofix.exe does not run.

The following extension settings are maintained in a default Group Policy object: Remote Installation Services (RIS), security settings, and Encrypting File System (EFS). The following extension settings are not maintained or restored in a default Group Policy object: Software Installation, Internet Explorer maintenance, scripts, folder redirection, and administrative templates.

The following changes are not maintained or restored in a default Group Policy object: Security settings made by Exchange 2000 Setup, security settings migrated to default Group Policy during an upgrade from Windows NT to Windows 2000, and policy object changes made through Systems Management Server (SMS).

You can run this tool only on servers running the Windows Server 2003 family.

For more information about using GPMC to back up and restore GPOs, see the Administering Group Policy with the GPMC

white paper on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=17528).

For more information about restoring system state data by using the Backup utility in Windows Server 2003, see Backing Up and Recovering Data on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=22347).

For more information about managing the Sysvol directory, see Best Practices for Sysvol Maintenance on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=39986).

For more information about use of Dcgpofix, see The Dcgpofix tool does not restore security settings in the Default Domain Controller Policy to their original state on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35269).

# Windows Security Settings remain in effect after removal

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**

This problem typically occurs when a security setting is defined through Group Policy and later set to "undefined."

**Solution**

Under some circumstances Windows Security Settings remain in effect after being set to undefined. The flow chart below will help you determine if your security setting is remaining in effect because of this behavior.

**Account Policy, User Rights Assignment and Audit Policy**

If the security setting that is in effect is under Account Policy, Local Policy\User Rights Assignment or Local Policy\Audit Policy use the flow chart below to determine how this particular setting will behave when removed through Group Policy:




**Security Options**

If the security setting that is in effect is under Security Options use the flow chart below to determine how this particular setting will behave when removed through Group Policy:




**Event Log and System Services**

If the security setting that is in effect is under Event Log or System Services use the flow chart below to determine how this particular setting will behave when removed through Group Policy:




**Restricted Groups**

If the security setting that is in effect is under Event Log or System Services use the flow chart below to determine how this particular setting will behave when removed through Group Policy:




**All Other Windows Security Settings**

For Windows 2003 and Windows XP security settings not addressed in the sections above revert to their previous value when removed through Group Policy. For Windows 2000 security settings not addressed in the sections above will revert to their previous value unless the setting is defined in local policy in which case it will reflect the local policy setting.

# Fixing Scripts policy settings problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section discusses troubleshooting issues related to Scripts policy settings.

**Scripts Processing Background**
Startup, logon, logoff, and shutdown scripts can be applied using Group Policy settings. The Scripts CSE processes these script settings. The Scripts CSE updates the registry with the location of one or more script files so that the UserInit process can find those values in the course of its normal processing. When a CSE reports success, it might only mean that the value has been placed in the registry. Even though the setting is in the registry, there could be problems preventing the setting from being applied to the client. For example, if a script specified in a Script setting has an error that prevents it from completing, that Script CSE does not detect error.

Scripts processing contains two steps--Group Policy processes a GPO and stores the script information in the registry:

- HKEY_CURRENT_USER\Software\Policies\Microsoft\Windows\System\Scripts (User Scripts)
- HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\System\Scripts (Machine Scripts)

| **Note** |
| --- |
| Script is run by means of a UserInit process. (By default, scripts that cannot be completed time out after 10 minutes.) |

Only Windows XP, Windows Server 2003, and later operating systems support WMI filtering. If the computer is running an earlier operating system, the WMI filter is ignored and the GPO is applied.

| **Note** |
| --- |
| The time-out is the time allotment for all scripts to run. This can be modified using the Computer policy setting **Maximum wait time for Group Policy Scripts**. |

Common script errors include:

- Incorrect script path.
- Script time-out.
- Access to script is restricted by means of ACLs (typically for startup/shutdown scripts that run as computer, not user).

If a logon script fails, it typically does not affect the other scripts. However, startup scripts are often run synchronously, and a failure of one of these scripts can affect scripts intended to run later.

To investigate, check the Application Event Log for entries with UserInit as the source.

**Fixes for specific Scripts policy settings issues**
From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- Startup scripts do not run as expected on a computer running Windows XP SP1 or Windows XP SP2
- Scripts deployed to computer do not run (Userinit events 1000 and 1001)

# Startup scripts do not run as expected on a computer running Windows XP SP1 or Windows XP SP2

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You start a computer running Windows XP Service Pack 1 (SP1) or Windows XP Service Pack 2 (SP2) on a wired network or on a wireless network that has IEEE 802.11 authentication, and Group Policy startup scripts and logon scripts do not run as expected, or the computer does not join the Active Directory directory service.

**Cause**

The following entries may appear in the log files in Event Viewer:

**Event ID: 1054**

**Source: NT AUTHORITY/SYSTEM**

**Type: Error**

**Description: Windows cannot obtain the domain controller name for your computer network. (The specified domain either does not exist or exist or could not be contacted). Group Policy processing aborted. Data: (unavailable)**

**Event ID: 1000**

**Source: UserInit**

**Type: Error**

**Description: Could not execute the following script AdminPassword.bat. The system cannot find the file specified. Data: (unavailable)**

This problem may occur if the Group Policy engine or Active Directory time out while waiting for the network to start. A race condition may occur between the TCP/IP protocol and the network adaptor driver when they try to register with the Microsoft Network Driver Interface Specification (NDIS). If the TCP/IP protocol registers with NDIS before the network adaptor driver, for a short time it prompts higher user mode networking components that network connectivity is not available. As a result, during this short time, the Group Policy startup script cannot be downloaded.

This problem is more likely to occur on fast networks that use 1 gigahertz (GHz) network adaptors or in teaming environments where the network takes several additional cycles to negotiate link speed.

**Solution**

For computers running Windows XP SP1, you can obtain a supported hotfix from Microsoft to correct this problem. For detailed instructions, see Group Policy startup scripts do not run as expected on a computer that is running Windows XP Service Pack 1 (SP1) or Service Pack 2 (SP2) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43049).

For computers running Windows XP Service Pack 2, you must add the **GpNetworkStartTimeoutPolicyValue** registry entry. This entry defines the number of seconds to wait before trying to run the Group Policy startup script again. For more information, see Group Policy startup scripts do not run as expected on a computer that is running Windows XP Service Pack 1 (SP1) or Service Pack 2 (SP2) on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43049).

# Scripts deployed to computer do not run (Userinit events 1000 and 1001)

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A logon, logoff, machine startup, or machine shutdown script deployed to the computer through Group Policy does not run.

**Cause**
The event logs may show one of the following Userinit events:

**Event ID 1000**

**Message: Could not execute the following script %1. %2.**

**Event ID 1001**

**Message: Could not locate the script command lines in the Group Policy object**.

These Userinit events may be caused by one of the following issues:

- Network problems that prevent the computer from accessing the domain controller or the network location of the script executable files
- Invalid script command lines
- Replication problems on the domain controller

**Solution**
**Network Errors**
Network errors that occur during script execution can cause these errors. In this case, troubleshooting and resolving the network problems will resolve this error. For information about troubleshooting network issues, see Fixing Group Policy networking issues.

A script may also be inaccessible because of file sharing violations. In this case, the administrator of the network share on which the script resides must determine which users or computers are locking the file and close their connection to the file. To perform this diagnosis, you can use file system tools (such as net.exe).

For information about file sharing violations and replication issues, see FRS Encounters "ERROR_SHARING_VIOLATION" Errors When It Tries to Replicate Data That Is Still in Use on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43053).

**Script command lines errors**
**To troubleshoot command line errors**

- Run the command line listed in the event log, using the directory from which the script originates as your default directory. The script directory is either the user\scripts or the machine\scripts directory, depending on whether the script is a logon/logoff script or a machine startup/shutdown script. You can use the Group Policy Management Console (GPMC) to find the GPO containing this path. If the policy is a computer policy, use an account that has access to the network location listed in the script.
- If the command line execution fails, investigate this as you would investigate any error related to running a command. For example, look for typographical mistakes in the path or the specification of a relative path that is not present on the local computer.
- To change an incorrect path, use GPMC to edit the GPO for which you want to modify the path, and type the correct path. To fix this problem, you must be a network administrator with rights to the GPO.

**Domain Controller replication issues**
When scripts are stored on the domain's sysvol share, replication problems can cause the script to not be present or to have incorrect access permissions. For information about troubleshooting replication issues, see Group Policy does not replicate.

# Fixing Software Installation policy setting problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic explains issues related to deploying software using Group Policy software installation.

**Software Installation CSE Background**
A number of special issues affect Group Policy software installation. For example, network connectivity issues can disrupt access to the software installation packages. Software installation processing is never performed while the user is logged on because doing so could disrupt work and result in data loss. This has implications when Group Policy is configured to run asynchronously.

Software Installation is managed by the Software Installation CSE, which appears as the source Application Manager on the Policy Events tab in Group Policy Results reports. If you need additional logging information, enable verbose logging for the Software Installation CSE and Microsoft® Windows® Installer, as described under Troubleshooting Group Policy Using Log Files.

The Software Installation Diagnostics tool (addiag.exe) provides detailed information about the applications visible in Active Directory and installed for the current user, as well as general diagnostic information and related Event Log entries. It is available in the Windows 2000 Server Resource Kit.

**Fixes for specific Software Installation policy setting issues**
From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- Software Installation changes do not apply after logon
- Failure to access software share points
- User uninstalls a deployed application on a computer and application is unavailable to the user on every computer
- Software is not listed in Add or Remove Programs
- Install on demand fails
- Software Installation operation fails

# Software Installation changes do not apply after logon

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When the Fast Logon Optimization feature is set on computers running Windows XP Professional (default), Group Policy software installation changes do not take effect after the user logs on.

**Cause**

In Windows XP Professional, the Fast Logon Optimization feature is set by default for both domain and workgroup members. As a result, Windows XP does not wait for the network to be fully initialized at startup and logon. Existing users are logged on using cached credentials. This results in shorter logon times. Group Policy is applied in the background after the network becomes available. Note that because this is a background refresh, extensions such as Software Installation and Folder Redirection take two logons to apply changes.

Software Installation processing occurs only during computer startup or when the user logs on. This is because processing periodically could cause undesirable results. For example, if an application is no longer assigned, it is removed. If a user were using the application while Group Policy tries to uninstall it or if an assigned application upgrade takes place while someone is using the application, errors would occur.

If the software installation settings are applied through computer configuration, they are applied at startup. If software installation settings are applied through user configuration, they are applied at logon.

**Solution**

If the computer is running Windows XP with logon optimization enabled, the user will need to log on after Group Policy refresh. This can entail logging on two or three times.

Administrators can turn off Fast Logon Optimization, by using the **Always wait for the network at computer startup and logon** policy setting in the **Computer Configuration\Administrative Templates\System\Logon** node of Group Policy Object Editor.

For more information, see Description of the Windows XP Professional Fast Logon Optimization Feature on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43045).

# Failure to access software share points

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An administrator has deployed a software package, but software installation fails. Event logs indicate the package can not be accessed.

**Cause**

One of the reasons that software installation may fail is if the installation source for the software is not available. This may be due to an incorrect path statement for the package or inadequate permissions.

**Solution**

Test for mistyped elements in the path specification by manually following the path. If the path is incorrect, provide the correct path to the share that contains the source package you are deploying. To do this, edit the GPO that you used to deploy the software, and correct the path.

**To correct the default path for packages**

1. In Group Policy Management console, right-click the GPO that contains the software installation policy settings for the package you are deploying, and click **Edit**.
2. Expand either **User Configuration** or **Computer Configuration** (depends on where you configured the software installation policy), click **Software Settings**, right-click **Software installation**, and then click **Properties**.
3. In **Software installation Properties**, on the **General** tab, type the path in the **Default package location** text box.

Verify that the user account or system account has the necessary privileges to traverse the path and access those resources. The account whose OU supplied the GPO is the one that needs to traverse the path.

# User uninstalls a deployed application on a computer and application is unavailable to the user on every computer

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A user with a roaming user profile uninstalls a deployed application on his or her computer, and the user can no longer access that application when the user roams to another computer.

**Cause**

If the user has a roaming user profile and the user uninstalled the application on another computer using **Add or Remove Programs**, the application will be unavailable to that user on every computer (with the possible exception of computers where loopback processing rules apply). For more information about roaming user profiles, see “Using roaming user profiles” in Help and Support Center for Windows Server 2003.

**Solution**

If the application was deployed with the **Uninstall this application when it falls out of the scope of management** option, verify that the computer and user are still in the necessary security groups and that the GPO is still linked to a site, domain, or OU that the user or computer is in.

# Software is not listed in Add or Remove Programs

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An administrator deploys a software package and the software does not appear in users' computers in Add or Remove Programs in Control Panel.

**Cause**
For the software to appear in the Add New Programs section of Add or Remove Programs, it must be Published (not Assigned) in Group Policy.

Other possible causes for this issue include:

- Group Policy did not apply
- The network can not be reached (can not query the domain controller for the user or computer)

**Solution**
Choose the **Published** software deployment method when you deploy the package.

The Windows Installer package must be written such that the software can appear in **Add or Remove Programs**.

To troubleshoot Group Policy application failure, see Fixing Core Group Policy problems.

To troubleshoot network connection issues, see Fixing Group Policy networking issues.

# Install on demand fails

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When users try to install an advertised application by opening a document that is associated with the application, the installation fails. Document activation of a published application does not trigger an application installation.

**Cause**
This issue might occur if the administrator did not set autoinstall options for the software being deployed, or the file name extension for the application was not set correctly.

**Solution**
To troubleshoot this issue:

- Ensure that the autoinstall option is set for the application being deployed.

| **Note** |
| --- |
| To complete these procedures, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group. |

**To set the autoinstall option for an application**

1. In Group Policy Management console, right-click the GPO that contains the software installation settings for the application you are deploying, and click **Edit**.
2. In Group Policy Object Editor console, expand **Computer Configuration** or **User Configuration** (depends on where you configured the software installation policy), click **Software Settings**, and click **Software installation**.
3. In the details pane, right-click the application that you want to be installed automatically when the user tries to open a file with the associated extension.
4. Click **Properties**, and then click the **Deployment** tab.
5. Under **Deployment options**, select the **Auto-install this application by file extension activation** check box, and then click **OK**. This option is recommended, and it is selected by default.

- Check the order of file name extensions specified for the GPO and ensure that the file name extension is associated with the application.

**To specify automatic installation options based on file name extension**

1. In Group Policy Management console, right-click the GPO that contains the software installation settings for the application you are deploying, and click **Edit**.
2. In Group Policy Object Editor console, expand **Computer Configuration** or **User Configuration** (depends on where you configured the software installation policy), click **Software Settings**, right-click **Software installation**, and click **Properties**.
3. In the **Software installation Properties** dialog box, click the **File Extensions** tab.
4. On the **File Extensions** tab, specify the following option:
    - **Select file extension**. Displays a list of all available file name extensions for the managed software. When an extension is selected, all applications in Active Directory that are associated with the selected file name extension are displayed in the **Application precedence** list box. For example, for an .htm extension, many applications can be associated, such as Microsoft Word, Microsoft Internet Explorer, Netscape Communicator, and so on.

- Check the client computer for applications that have been installed locally. If there is a conflict, the locally installed

application is used instead of the deployed application.

# Software Installation operation fails

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An application could not be installed or uninstalled. One or more of these software installation operations fails: installing, uninstalling, upgrading, and reinstalling software.

**Cause**

The cause of the failure depends on the type of operation that failed: an uninstall, install, reinstall, or the domain controller query to determine administrative policy. The event logs indicate that software was assigned in addition to being installed. Assignment is simply the first phase of an installation. For troubleshooting purposes, treat assignment failures in the same manner as installation failures.

Possible causes include:

**Uninstall failure:**

- The Group Policy Software Installation extension might need to uninstall software.
- The system is using a file that needs to be accessible when the application is uninstalled.
- The Windows Installer package that was used to deploy the application was authored incorrectly.

**Install or reinstall failure:**

- The installation source for the software is not available.
- The application is already installed on the computer but has a different set of transforms from the application deployed through Group Policy.
- An existing application on the computer conflicts with the application and is preventing it from being installed.
- The Windows Installer package of the deployed application is incorrectly authored.

**Upgrade failure:**

- The Group Policy Software Installation extension attempted to upgrade an application, but was unable to do so due to an installation or uninstallation error; the event logs provide additional information.

**Domain Controller query failed:**

- The Group Policy Software Installation extension could not access the network, and was not able to query a domain controller for the user or computer.

**Solution**

To troubleshoot software installation issues, try one or more of the following:

- Determine whether any other related events are logged in Event Viewer. Windows Installer usually logs a related event that provides details about why the operation could not be completed.
- Check for errors in any application deployed to the computer and look for Windows Installer events.
- Enable logging for Application Management and Windows Installer.
- Log on to another computer as an administrator, verify that the application can be manually installed and uninstalled on that computer. If this verification fails, the problem is probably a software issue. Contact the software support provider.

- In the case of an install or uninstall failure, manually uninstall the software.

- If the event indicated that the transforms of the deployed application are different from the transforms of an instance of the same application already installed on the computer, deploy a new version of the application that has the same set of transforms as the existing installation, and then remove the deployment that has the conflict.

  Deployment changes can be made by a network administrator with the right to manage Group Policy objects (GPOs) in the Windows domain of the user or computer.

- If the problem is the Group Policy Software Installation extension failing to access network information related to software installation policy, run the **Ping**, **DCdiag.exe**, and the **Nslookup** tools on the computer that is having the problem. This helps determine where on the network the problem is occurring.

- You can use Group Policy to enable logging for Windows Installer by using the **Logging** policy setting in the **Computer Configuration\Administrative Templates\Windows Components\Windows Installer** node of Group Policy Object Editor.

- You can use the **Software Installation Diagnostics** tool (**addiag.exe**) to gather additional diagnostic information. You can run **addiag.exe /?** from a command prompt for information about the syntax.

To learn more about the tools that you can use to force an application uninstall, see the Windows Installer documentation. For more information about Windows Installer, see the MSDN article Roadmap to Windows Installer Documentation on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=21878).

For information about Windows Installer event entries, see the MSDN article Event Logging on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=21877).

For more information about enabling Windows Installer Logging for Windows XP clients, see How to Enable Windows Installer Logging in Windows XP on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35731).

For more information about enabling Windows Installer Logging for Windows 2000 computers, see How to Enable Windows Installer Logging on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35593).

For more information about using **DCDiag**, see DCDiag and NetDiag in Windows 2000 Facilitate Domain Join and DC Creation on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=20908).

# Fixing Folder Redirection policy settings problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic discusses troubleshooting for Folder Redirection issues.

**Folder Redirection CSE Background**
The Folder Redirection client side extension, fdeploy.dll, processes Group Policy settings for Folder Redirection. Folder Redirection also includes the following components:

- {25537BA6-77A8-11D2-9B6C-0000F8080861}.ini: A locally cached redirection information file specific to each user profile.
- {1C08E84D-F112-4252-978B-EC82A225CC20}.ini: A file indicating the status of the previous locations used for redirected folders. This file is used to detect server location changes and user name changes.

Folder Redirection is used to maintain user data in a centralized location. This permits regular backups of the information, and also provides the user with access to the data from any computer in the network.

The following folders can be redirected:

- My Documents
- Application Data
- Desktop
- Start Menu

The Folder Redirection CSE manages folder Redirection. When events from this CSE are listed on the Policy Events tab in Group Policy Results reports, the source is listed as Folder Redirection.

Fdeploy.dll is loaded during the logon process by the Group Policy engine. The Folder Redirection policies are passed to fdeploy.dll, which then examines the policy settings and redirects user folders based on those settings.

The first step in Folder Redirection is to get a list of all Folder Redirection policies and evaluate them. The policies are handed down to the fdeploy.dll by the Group Policy engine through an API call into fdeploy.dll. The policies come in the form of two linked lists. The first list is a list of added policies. This refers to a list of Folder Redirection policies that are still applicable to the user. The second list is a list of deleted policies that refers to policies that used to be applicable to the user but are no longer applicable. The changes that result from Folder Redirection are then updated with the shell. The client-side cache, also referred to as offline files, is also updated to reflect the files moved by Folder Redirection.

**Fixes for specific Folder Redirection policy setting issues**
This section provides information for troubleshooting the following Folder Redirection problem:

- Folder Redirection encounters errors and redirection fails
- Folder Redirection CSE processing is delayed

# Folder Redirection encounters errors and redirection fails

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Folder redirection encounters errors when redirecting the user's folders on the client. As a result, the user's folders are not redirected successfully.

**Cause**

Several conditions must be met in order for users' folders to be redirected successfully. If these conditions are not met, the following Folder Redirection errors may occur:

- Folder redirection initialization fails.
- The Folder redirection path specified by the administrator is invalid, too long, inaccessible, or offline.
- You do not have adequate privileges to create or access the redirected folders on the server, or there is not enough free disk space on the server.
- Folder redirection is unable to set the shell folder registry keys to redirect the folders to the required destination.
- Configuration errors occur in the event that a parent folder was redirected to a subfolder.
- Folder redirection is unable to acquire the fdeploy.ini config file from the GPO.
- Some files to be moved as part of redirection are locked on the client.

**Solution**

To troubleshoot Folder Redirection problems, try one or all of the following:

- Look at the Event Viewer logs for Folder Redirection events. Based on the event, one or a combination of the following steps may be required to resolve the issue.
- If Folder redirection initialization fails, ensure that there are no memory or security issues on the client and that there are no issues with the GPO (this applies only if there was an error extracting the fdeploy.ini file). Ensure that the correct Fdeploy.ini file is available on the domain controller that the client is accessing.
- Verify that the folder redirection path specified by the administrator is valid and accessible. If the target paths are longer than the set limit, you need to make them shorter (look at the event log to see if the path length exceeds permitted number of characters). If the path to the folder does not exist (for example if the path specification is mistyped in the policy setting, if folders in the path have been renamed or removed, or if the server is unavailable), Folder Redirection will fail.
- Ensure that the folder redirection configuration is correct and that you are not redirecting parent folders to subfolders.
- Ensure that no files to be moved are locked by an application or service.
- Ensure that the user has the required ACLS on all the folders and shell folder registry keys.
- Verify that the user has appropriate permissions to create folders at the redirected location and that there is adequate free disk space to create the redirected folders. If disk quotas are applied, adjust the quotas to ensure the user has the required disk space required.

**Note**

In order to use the folder, the file system and share permissions must be set such that the user can navigate the path to the folder, and if the folder exists the user must have ownership privileges on it. (The user is given ownership of the folder by default if you allow the folder to be created automatically.) This is a common cause of confusion with Folder Redirection. When using Folder Redirection Policies, it is best to allow the system to create and set permissions on the folder. This reduces the likelihood of errors due to incorrect security settings. For more information about Folder Redirection permissions, see the "User Data and Settings Management for Windows Server 2003" white paper on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=26719.

If you need to set permissions manually, you must ensure that the user has the appropriate minimum file system and share permissions. For more information about the required permissions, search for "Folder Redirection" in Windows Server 2003 Help and Support.

For more information about disk quotas, search for "Manage Disk Quotas" in Windows Server 2003 Help and Support.

- Folder Redirection, like Software Installation settings, can only be applied during computer startup or user logon. On computers running Windows XP with logon optimization enabled, this can mean that the user needs to log on more than once before the setting takes effect.

# Folder Redirection CSE processing is delayed

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Application of the user's folder redirection policy is delayed until the next logon, because Group Policy logon optimization is in effect. As a result, the user's folders are not redirected during the current logon.

**Cause**

Folder Redirection cannot redirect the user's folders because the Group Policy asynchronous foreground processing setting is turned ON. Folder Redirection requires that the redirection policy be applied synchronously, so that the folders are redirected as part of the login and prior to user access.

The **Always wait for the network at computer startup and logon** policy setting determines whether Windows XP waits for the network during computer startup and user logon. By default, Windows XP does not wait for the network to be fully initialized at startup and logon. Existing users are logged on using cached credentials, which results in shorter logon times. Group Policy is applied in the background once the network becomes available.

If the **Always wait for the network at computer startup and logon** policy is set to **Not configured** or **Disabled**, Windows XP does not wait for the network to be fully initialized at startup and logon. Group Policy is applied asynchronously in the background.

In order for Folder Redirection settings to apply with the first logon, Windows must wait for the network to be fully initialized at startup and logon before applying policy. This means policy must apply synchronously in the foreground. The same applies to Software Installation and roaming user profiles.

**Solution**

To guarantee the application of Folder Redirection in just one logon, enable the **Always wait for the network at computer startup and logon** policy setting to ensure that Windows waits for the network to be available before applying policy. When this setting is enabled, logons are performed in the same way as for Windows 2000 clients, in that Windows XP waits for the network to be fully initialized before users are logged on. Group Policy is applied in the foreground, synchronously.

For more information about this policy, see the **Explain** tab in the **Always wait for the network at computer startup and logon** policy setting.

# Fixing Disk Quota extension problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic discusses troubleshooting for Disk Quota issues.

**Disk Quota CSE Background**
Disk quotas allow administrators to control the usage of hard drive space by their users. The policy is applied on a per user/per drive letter basis. This allows an administrator to keep a user from filling up a hard drive that multiple people share.

If the Disk Quota CSE is unable to read the registry on the target computer, or is unable to access a list of the available hard drives on the target computer, the Disk Quota will not be applied. This could be potentially lead to a user unwittingly filling up all the hard drive space on a server or other unwanted situations.

**Fixes for specific Disk Quota extension issues**
From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- Disk Quota CSE processed with errors

# Disk Quota CSE processed with errors

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This error indicates that the Disk Quota Group Policy client-side extension (CSE) has failed to apply Group Policy settings to one or more volumes on the target computer.

**Cause**

This could be caused by one of the following issues:

- Unable to open a registry key containing the Group Policy information.
- Unable to create the class factory object for ProgID "Microsoft.DiskQuota.1".
- Unable to instantiate the "quota control" object; ProgID "Microsoft.DiskQuota.1".
- Unable to obtain the list of disk drives on the target computer.
- Unable to apply quota settings to a volume on the target computer.
- Insufficient memory to perform an operation.

The following event numbers with a source of DiskQuota can trigger this alert:

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8

**Solution**

To troubleshoot this problem, try one or all of the following:

- Review the Event Viewer for events logged by the Disk Quota client-side extension and use the information in the event description to troubleshoot the problem.
- If the event description indicates that a class factory or quota object could not be created, perform the following tasks:
    - Check for the existence of the file "%SystemRoot%\system32\dskquota.dll"
    - Re-register dskquota.dll using regsvr32.exe
- An event description that indicates that a list of disk drives on the target computer could not be obtained is created when no logical drive(s) are present on the target computer.

- In cases other than these, Disk Quota failure is usually caused by low memory conditions on the target computer.

# Fixing Group Policy problems by using log files

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic provides information related to troubleshooting Group Policy problems by using log files.

**Troubleshooting Group Policy by Using Log Files**
If other tools do not provide the information you need to identify the problems affecting Group Policy application, you can enable verbose logging and examine the resulting log files. Verbose logging can reduce performance and consume significant disk space, so as a best practice enable verbose logging only when necessary.

For example, enabling Userenv.dll logging allows administrators to view information about what occurs in the background when a user logs on to the computer. Userenv.dll is part of the code of the Group Policy engine and is always called whenever policy settings are processed.

**Client Log Files**
Log files can be generated by the core client engine (Userenv) and by every CSE except the Scripts CSE. Scripts processing is logged in the Application log on the client with source=UserInit. Use Event Viewer to view the Application log on the client, or look for these entries on the **Policy Events** tab of the Group Policy Results report.

The **Policy Events** tab in Group Policy Results reports generated in GPMC displays the Group Policy–related events that you would see if you used Event Viewer to view these events in the Application log on the client for which you generated the report.

The following table lists several log files you can generate at the client that relate to Group Policy troubleshooting.

**Client Log Files for Troubleshooting Group Policy**

| Output from: | Is located in this file: | Enable verbose logging by adding this key or value... | ...to this registry key |
|---|---|---|---|
| Group Policy core (UserEnv) and registry CSE | %windir%\debug\usermode\UserEnv.log | UserEnvDebugLevel = REG_DWORD 30002 | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon |
| Security CSE | %windir%\security\logs\winlogon.log | ExtensionDebugLevel = REG_DWORD 0x2 | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon\GpExtensions\{827d319e-6eac-11d2-a4ea-00c04f79f83a}\ |
| Folder Redirection CSE | windir%\debug\usermode\fdeploy.log | FdeployDebugLevel = Reg_DWORD 0x0f | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics |
| Software Installation CSE | %windir%\debug\usermode\appmgmt.log | Appmgmtdebuglevel=dword:0000009b | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics |

| | | | |
|---|---|---|---|
| Windows Installer (deployment-related actions) | %windir%\temp \MSI*.log | Logging = voicewarmup<br><br>Debug = DWORD: 00000003 | HKEY_LOCAL_MACHINE \Software \Policies \Microsoft \Windows \Installer |
| Windows Installer (user-initiated actions) | %temp% \MSI*.log | Logging = voicewarmup<br><br>Debug = DWORD: 00000003 | HKEY_LOCAL_MACHINE \Software \Policies \Microsoft \Windows \Installer |

**Note**

The UserEnv logs entries pertaining to profiles as well as Group Policy core processing and registry (.adm) processing on the client. The entries pertaining to profiles are intermingled with the Group Policy entries and not easily distinguished from them. Use Wilogutl.exe to analyze the Windows Installer log files. For more information, see Wilogutl.exe on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=16156).

**Server Log Files**

You can enable logging of events generated by the Group Policy Object Editor on the Server. There are two different log files, one for events relating to core Group Policy processing and the registry CSE, and another for events relating to all other CSEs.

The following table lists several log files you can generate at the server that relate to Group Policy troubleshooting.

**Server Log Files for Troubleshooting Group Policy**

| **Output from:** | **Is located in this file:** | **Enable verbose logging by adding this keyword...** | **...to this registry key** |
|---|---|---|---|
| GPMC: error logging only | %temp%\gpmgmt.log | gpmgmttracelevel=1 | HKEY_LOCAL_MACHINE \Software \Microsoft \Windows NT \CurrentVersion \Diagnostics |
| GPMC: error and verbose logging | %temp%\gpmgmt.log | gpmgmttracelevel=2 | HKEY_LOCAL_MACHINE \Software \Microsoft \Windows NT \CurrentVersion \Diagnostics |
| GPMC: Output only to log file (not to debugger) | %temp%\gpmgmt.log | gpmgmtlogfileonly=1 | HKEY_LOCAL_MACHINE \Software \Microsoft \Windows NT \CurrentVersion \Diagnostics |
| Group Policy Object Editor: Core-specific entries | %windir%\debug\usermode \gpedit.log | GPEditDebugLevel = REG_DWORD 0x10002 | HKEY_LOCAL_MACHINE \Software \Microsoft \Windows NT \CurrentVersion \Winlogon |

| | | | |
|---|---|---|---|
| Group Policy Object Editor: CSE-specific entries | %windir%\debug\usermode<br><br>\gptext.log | GPTextDebugLevel = REG_DWORD 0x10002 | HKEY_LOCAL_MACHINE<br>\Software<br>\Microsoft<br>\Windows NT<br>\CurrentVersion<br>\Winlogon |

**Userenv Logging**

Because Userenv tracks the Group Policy engine and registry-based Group Policy, it is the most frequently used log file for Group Policy troubleshooting. Userenv is especially useful in a Windows 2000 environment because you don't have the benefit of using Resultant Set of Policy (RSoP). Most of the questions that RSoP answers, are in the userenv log.

To use userenv.log you need to first enable verbose logging.

**To enable verbose logging**

1. Log onto the client computer as the administrator and run Regedit.
2. Locate the following key: **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon**.
3. Right click **Winlogon**, select **New**, and then click **DWORD Value**.
4. Enter the following name for the **DWORD Value**: **UserEnvDebugLevel**.
5. Enter 30002 as the hexadecimal value. This writes the userenv into userenv.log, located in the \%windir%\debug directory.
6. Run "**gpupdate /force**" to ensure a full listing of total Group Policy processing.

**See Also**

**Other Resources**

Microsoft Group Policy Home Page

# Interpreting Userenv log files

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic explains how to interpret userenv log files.

**Userenv Logging**

When you enable userenv logging, you can perform debug logging of the user profile and the system policy processes. Userenv log files also contain information about the status of each Group Policy extension, such as Application Deployment, Security, and Folder Redirection. Userenv log files reveal what is occurring in the background as a user logs on. Userenv log files are especially useful because they be can be used to troubleshoot Windows 2000 operating systems, in which you cannot use Resultant Set of Policy (RSoP). In addition, if Active Directory replication is not working, RSoP will not work, leaving userenv log files as the only troubleshooting option. RSoP also depends on Windows Management Instrumentation (WMI), which might also fail and leave you with no option but to analyze the userenv log files. With verbose logging, information is logged every second; if something fails you can review the userenv logs and pinpoint the likely source of failure.

Userenv log files contains information about the following:

- Group Policy settings that are not processed or not applied as expected
- Folder redirection that does not occur
- Profile or registry hive load, unload, or deletion failures
- Logon script, or script not applied as expected
- Default behaviors occurring because a slow link was detected
- Roaming profile issues
- Slow logon issues
- Whether a given GPO is accessible, and if not, why access was denied.
- The name of the domain controller that is accessing SYSVOL.

**What is in the Userenv log?**

The example of the Userenv log in this section shows the type of details of errors in core Group Policy processing that the Userenv log provides. Reading from left to right, this log shows a process code (for example, cc.500), the time it was processed (note the date is not displayed), the process name, and a short statement of the error. The Userenv log displays Group Policy process failures and warnings.

**Sample Userenv log**

```
USERENV(178.17c) 13:48:51:089 MyRegUnLoadKey:  Failed to unmount hive 00000005
USERENV(178.17c) 13:48:51:089 DumpOpenRegistryHandle: 3 user registry Handles leaked
from \Registry\User\S-1-5-21-2127521184-1604012920-1887927527-2193396
USERENV(178.17c) 13:48:51:089 UnloadUserProfileP: Didn't unload user profile <err = 5>
USERENV(178.2e8) 13:52:00:708 GetGPOInfo:  Local GPO's gpt.ini is not accessible,
assuming default state.
USERENV(178.600) 13:52:09:452 GetGPOInfo:  Local GPO's gpt.ini is not accessible,
assuming default state.
USERENV(178.4d4) 13:52:15:020 PolicyChangedThread: UpdateUser failed with 6.
```

| **Note** |
| --- |
| The Userenv log has a maximum size of 1 megabyte (MB). At system startup, if the log file exceeds 1 MB, the contents is copied into a Userenv.bak file and a new Userenv log is created. If the system remains running and is not restarted, the log file exceeds 1 MB. |

**General Userenv.log Messages**
The following messages have the same meaning no matter where they appear in the log:

- Failed to Allocate...
- Failed to Create...
- Failed to Set...
- Failed to Get...
- These errors normally indicate a problem with system resources at the time when they were called. If you receive multiple errors at the same time, your computer could be having integrity problems. To avoid these problems, try to ensure that the computer is as close to a base install of Windows as possible by removing all third party software programs. If this option is not possible for you, install a new version of Windows and all available service packs on a test machine. If the error is no longer appearing in the log file, slowly install third party software, one program at a time. Monitor the Userenv log after each program is installed to determine which program is causing the error.

The following messages also have the same meaning no matter where they appear in the log:

- Calling function <function failed>
- Calling function <function returned>
- These entries indicate that an underlying function was called and that there was an error in the transaction. Determine which function the entry is causing the error by viewing the error text, and then locate the function that failed in the log. After you have identified the failed function entry in the log, get the return code and start troubleshooting the failed function from there.

**ApplyGroupPolicy Messages**
**ApplyGroupPolicy: Entering. Flags = %x**
You use flags to determine what processing Group Policies need to do and what some of the Client Side Extensions (CSEs) are doing. Flags are hexadecimal bits that might need to be converted if the number is above nine.

| Hexadecimal | Description |
|---|---|
| 0x00000001 | Machine Policy (The absence of a 1 = User Policy) |
| 0x00000002 | Background Refresh |
| 0x00000004 | Apply Directory Services Policy |
| 0x00000008 | Do not Wait on Network Services |

The values listed in the following table are added to the flags after processing.

| Hexadecimal | Description |
|---|---|
| 0x00010000 | Background Thread Processing |
| 0x00020000 | Change in Control Panel Settings |
| 0x00040000 | Slow Network Connection |
| 0x00080000 | Verbose output to the event Log |
| 0x00100000 | Forced Refresh |

| | |
|---|---|
| 0x00800000 | Planning Mode |

**CheckGPOs Messages**
Most of the log messages from this function are self-explanatory. The following is the most common error that users ask about in relation to CheckGPOs log messages.

**CheckGPOs: No GPO changes but couldn't read extension %s's status or policy time.**
Group Policy stores the status and the time of each extension as needed for the machine and every user that logs onto the machine. They are stored in the following registry locations:

MACHINE

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Group Policy\State\Machine\Extension-List

USER

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Group Policy\State\<SID>\Extension-List

If you receive this error, check your registry permissions for any explicit denials on the registry settings. Also, check for any third-party software that might manipulate this key. Some third parties change the Group Policy cache in both the file system and registry to have their product mimic Group Policy.

**EnterCriticalPolicySection Messages**
This is a potentially serious log message. It could indicate that certain portions of the operating system have become corrupt due to improper shut-down or system crashes. It could also indicate a system resource problem, similar to the "Failed to..." errors mentioned earlier in this section. To troubleshoot this issue, remove all third-party software to determine if a non-Microsoft program is causing this error. If that does not fix the problem, start with a clean installation of Windows, install all necessary service packs, and then introduce the clean computer into your environment. Monitor the Userenv log while installing any additional programs to determine what could be causing this error.

**GPOThread**
**GPOThread: Next refresh will happen in %d minutes**
This is an informational message. It is the only place within the operating system that indicates when the next time Group Policy will refresh.

**LeaveCriticalPolicySection**
See the "EnterCriticalPolicySection" section earlier in this document.

**LibMain**
**LibMain: Process Name: %s"**
This is an informational message. This message appears in the userenv.log when an application loads Userenv.dll within its space. Userenv has many exported functions that deal with loading and unloading the user's profile. Therefore, many application need to go through Userenv to get the profile of the currently logged on user for application specific settings. This message indicates which application is loading an instance of Userenv.dll.

**ProcessGPO**
**ProcessGPO: Found flags of: %d"**
This is an informational message. To determine what the flag in the message means, see the following list:

- 0 = Both user and Computer Configuration are Enabled
- 1 = User Configuration is Disabled
- 2 = Computer Configuration is Disabled

**ProcessGPO: Found machine version of: GPC is %d, GPT is %d"**
**ProcessGPO: Found user version of: GPC is %d, GPT is %d"**
These two messages are misrepresented in the Userenv log. It states GPC and GPT but it does not actually read from the GPT. It reads both settings from the GroupPolicyContainer Object (GPC) and not SYSVOL. The Attribute it reads is the Version number, which is stored as an integer value. If you believe you have a SYSVOL problem, run GPOTOOL to determine where errors might exist.

**ProcessGPORegistryPolicy**
**ProcessGPORegistryPolicy: Failed to create archive file with %d"**

The archive file in question is named ntuser.pol and is located in the user's profile. If you receive this error in the Userenv log, determine whether you are having problems with profile loading or unloading. Also, make sure that the user has write permissions to the profile folder because ntuser.pol is created in the root of the user profile.

**ProcessGPOs: Machine role is %d."**
This entry includes a decimal number at the end (%d). To determine what processing is going on within Userenv, see the following list:

- 0 = Standalone machine that is not a member of a Directory Services Domain
- 1 = The machine is a member of a NT4 domain
- 2 = The machine is a member of a domain which supports Directory Services.
- 3 = The machine is a domain controller

This information helps you determine what kind of machine you are troubleshooting, which could indicate why a policy is or is not applying.

**ProcessGPOs: The DC for domain %s is not available. aborting"**
**ProcessGPOs: The DC for domain %s is not available."**
These messages can be caused by network problems or the Windows XP Fast-User logon Optimization option. Try turning off the Fast-User Optimization option. If that does not cause the errors to stop, you might have a network problem. Troubleshoot this as you would any network problem.

**ProcessGPOs: DSGetDCName failed with %d."**
DsGetDCName is a public API that performs Domain Controller discovery. Troubleshooting this most of the time will involve name resolution. You can simulate this call by using the utility nltest.exe. Use the command **ipconfig /flushdns** before using nltest when you are taking a network monitor trace. Windows has a client-side cache, and you want to capture the DNS data in the trace. The syntax for **nltest** is:

**nltest /DSGETDC:** *DomainName*

**ProcessGPOs: GetNetworkName failed with %d."**
This function writes the FQDN of the domain into the Group Policy Cache. Many of the CSEs and Components that are policy implemented look for this name in the cache. This return code should be used wit the net command to determine the error message. Some of the lower-level functions in this name are Winsock library calls. If this is appearing and you are seeing some policies fail, make sure the machine is free from any viruses, worms, or other forms of malicious software.

**ProcessGPOs: Calling GetGPOInfo for normal policy mode"**
This is an informational message. This lets you know that you are processing policy in normal mode.

**ProcessGPOs: Calling GetGPOInfo for replacement user policy mode"**
This is an informational message. This message appears if the computer is in Group Policy loopback mode. Group Policy uses the user policy settings from the location of the computer object, and Group Policy does not honor any of the user policy settings based on the user’s object in Active Directory.

**ProcessGPOs: Calling GetGPOInfo for merging user policy mode"**
This is an informational message. This message appears when the computer is in Group Policy loopback mode. Group Policy merges the user specific policy settings based on where the user’s object is located in Active Directory along with the location of the computer object settings in Active Directory.

**ProcessGPOs: Extension %s skipped with flags 0x%x."**
This message indicates that a client-side extension was skipped. Use the flag values to determine why the extension was skipped. Typically, it is background processing, but make sure it was not caused by an error.

**ProcessGPOs: Extension %s ProcessGroupPolicy failed, status 0x%x."**
This message appears if the extension has actually failed. Use the **net** command with the return status to determine what the possible causes are. Enable additional logging for the extension specified to determine a more accurate error. Remember the status returned to ProcessGPOs is mapped to a more standard Windows Error. More specific errors might be revealed in the extension's debug log.

**ProcessGPOs: Couldn't read status data for %s. Error %d. ignoring..**
This is an informational message. This message appears to be an error, however the Group Policy engine automatically processes the extension because it cannot read the status data, hence the "ignoring" notation at the end of the message.

**ProcessGPOs: Computer Group Policy has been applied.**
**ProcessGPOs: User Group Policy has been applied.**
These are informational messages. These messages indicate that the Group policy has applied policy to the best of its ability. It does not mean that there were not any errors.

**ProcessGPOs: Leaving with %d."**
This is an informational message. If this value is not 0, the policy encountered a problem when it was applied. The fact that it appears in the log as well as the two messages above means the Group Policy was able to handle the error gracefully and move on.

**SearchDSObject**
**SearchDSObject: Searching <%s>"**
This function is used to determine which policy objects are within scope for the logging on user or computer. It should have a Distinguished Name (DN) of the object in Active Directory. The DN should either be a Site object, Domain object, or an Organizational Unit object. This is the location in Active Directory in which the Group Policy engine is searching for linked policy objects.

**SearchDSObject: Too many linked GPOs in search.**
This error message is not actually a Group Policy error but rather an LDAP Error, LDAP_SIZELIMIT_EXCEEDED. The value is determined by the Search Request.

**SearchDSObject: Found GPO(s): <%s>**
Here is an example from a Userenv.log file.

```
SearchDSObject:  Found GPO(s):  <[LDAP://CN={31B2F340-016D-11D2-945F-
00C04FB984F9},CN=Policies,CN=System,DC=corp,DC=mstep,DC=com;0]>
```

Notice there is a `;0]` at the end of the message. This value has meaning and it will appear on every policy that is linked. To determine the value, see the following table :

- 0 = This is the default setting. The policy is enabled
- 1 = This policy is disabled
- 2 = This policy will be forced or cannot be overridden and will win all ties against block policy inheritance
- 3 = Settings 1 and 2 are both enabled hence the policy will be disabled.

**SearchDSObject: <%s> has the Block From Above attribute set"**
The Block From Above attribute prevents a linked policy from being inherited by a Site, Domain, or OU beneath it. Generally, this is only set at an OU level, however the option is available anywhere a policy can be linked. This setting is stored on the actual container object, just as the policies that are linked to the container. This specific value can be found in on the attribute Gpoption. This attribute is an integer value and only has two settings: zero (0) (default) indicates normal propagation of the policy to the container, and one (1) turns on Block Policy Inheritance, which prevents polices higher in the tree from applying to the container or objects in the container. A policy that has No Override or Force, ignores the Block Policy Inheritance flag and still applies.

**RETURN CODES**
This section contains the return codes listed throughout a Userenv.log. You can use this section to determine what the codes mean and how to apply them to the specific function that is returning them.

**General Rules**
Most functions return true, false, or a status code. For this documentation true or successful is represented by a 0. False or unsuccessful is represented by a 1. Status codes have different meanings. This section documents most of the well known status messages.

Most return codes, after they are converted from hexadecimal to decimal, can be described by using the **net** command. The command-line syntax is:

**NET HELPMSG** *Number*

Messages that begin with 0xC or anything similar will probably not work with the **net** command and are listed in the "Return Codes that cannot be converted with the net command " table later in this document.

**Return Codes**

| Return Code | Description |
|---|---|
| 1332 | None Mapped. This usually means you have a SID designated in either Group Policy or Security Policy that is no longer valid on the domain. Typically, this happens when a user or group is created in the domain, assigned to a specific setting, and then later is deleted from the domain without updating the policy to remove the user. This is typically associated with a user or group instead of a computer account. Usually you can look through the event logs and find an event that has a return ID of 1332. This message is also represented as 0x534. To solve the problem, find the entry that has the user or group name or a SID that does not resolve to a user name and remove it. |
| 1355 | No such Domain. Typically, this is a name resolution problem, and most of the time it will be with DNS. There are some functions that return as 1355. If you are using a legacy operating system such as Windows NT, these messages are not as serious as if you are using Windows 2000, Windows XP, or Windows Server 2003. Another possible cause is NetBois. Some environments have DNS lookups look to a WINS server for name resolution. If you are certain you have ruled out Name resolution then start with a network trace to determine where the computer is trying to contact and troubleshoot from there. This error is also seen as 0x54b |
| 1722 | RPC Server is unavailable. This error can mean several different things. Check for name resolution problems first. This is a common cause of this error. The next thing is to use a utility, like Port Query or RPC Ping, to make sure that client and the target server are communicating. Due to Windows dependency on RPC, any failures in RPC can have a repercussions. Make sure the critical services like the RPC Endpoint Mapper and RPCSS are turned on. Also make sure port 135 is not blocked so clients can bind and connect to the endpoint mapper. Finally, you can take a network trace to make sure RPC connectively is valid. Because RPC traffic is usually encrypted, you need someone that knows how to look for RPC in a network trace. |
| 14 | Not enough storage is available to complete this operation. This is usually a resource problem, specifically memory allocation. Verify that the computer has enough resources to complete its tasks. If a process started exceeding its virtual memory limit, you can get this message. Restarting usually clears the problem, however this might be temporary and the problem might continue to occur. Start monitoring the computer's performance after the restart to see if you can determine what is causing the heavy resource usage. Use Performance Monitor to base line and spot memory usage trends. |
| 1908 | Domain Controller could not be found for this domain. Use the same troubleshooting as you would for a 1355 error message |
| 5 | Access Denied. There is a security descriptor that is preventing the current credentials from accessing the resource. This could also be a privilege access that is being denied. For example, the SeNetworkLogonRight allows security principals to connect to the machine and access resources. If a user was not given this privilege they would get the Access Denied error. Check the Permission on the resource. If the function is accessing the registry, find out which key it is attempting to use. If it is attempting to access a file, check the permissions on the file. After you have exhausted the possibility of checking permissions, check the security policy for the proper privileges. |
| 1115 | Shutdown in Progress. The system is shutting down. This event overrides most events that occur in the operating system. If the shutdown is unexpected, check the event logs to determine why the system shut down. |
| 1326 | Unknown user name or bad password. Check the user name and password you are logged on with. If you had a password change log off and log on again with the new password and see if that resolves the problem. |
| 1753 | There are no more endpoints available from the endpoint mapper. Start off by stopping third-party services because this error usually occurs when a RPC application consumes all of the available endpoints or the server is very busy. Follow up with troubleshooting steps provided by the 1722 error. |
| 1317 | The specified user does not exist. This is usually the case where you have a SID that no longer resolves to a user name, either because it has been deleted or because it is a SID from a trusted domain where the trust is not working. It could also be that you have a user name that cannot be converted to a SID. You can use either the name2sid.exe or lua.exe utility to see if the name is valid in the current domain as well as trusted domains. |
| 1398 | There is a time or date difference between the client and server. Verify that your client, server and the KDC are within five minutes of each other. Make sure time zones and date are correct as well. Check the event log to see if you are having any W32time errors. Kerberos defaults to five minutes of allowable time skew. Any time past five minutes causes errors. |

| | |
|---|---|
| 2 | The system cannot find the file specified. This message is very similar to the error 3 message and you troubleshoot them the same way. This message indicates that the computer cannot locate either the file or the path that you have told it to look for. This occurs if you try to connect to a UNC path that is not valid (for example, reading policy from SYSVOL). Find the path or file in question and attempt to navigate to that path under the proper credentials. If you want to navigate to the path under the computer context. use **/interactive** at a command line. If you perform a network trace during the event, you can get some good information about the problem. |
| 3 | The system cannot find the path specified. For more information, see Error 2. |
| 53 | The network path was not found. This usually is the sign of a name resolution problem and most of the time it is referencing a NetBIOS name resolution problem. Verify that the NetBIOS cache using the tool nbtstat is not causing this error. Also, make sure there is not a LMHOSTS file on the computer, and, if there is, make sure it is correct. Finally, verify that WINS is configured correctly. When you have exhausted NetBIOS resolution do not forget to look at DNS resolution because DNS can effect lookup names in WINS and vice versa. You can also identify the path and attempt to connect to it via IP address. Also, IPSec can generate this error message too. When a computer has IPSec turned off and it is attempting to communicate with a machine that has IPSec turned on, this error can occur. |
| 32 | The process cannot access the file because it is being used by another process. Another process has a lock on the file and is preventing access to the file. Restarting usually fixes this issue. If you discover that you need to restart often, you need to identify the process that is not releasing the file. You can use a tool like OH.exe to locate the offending process. After you identify the process, you will need assistance from the vendor to determine why the process is not releasing the file. Verify that you do not have any user mode applications that are holding the process open. To fix this problem, try stopping services, such as Realtime virus scanners, Open File Agents. |

**Return Codes that cannot be converted with the net command**

| Return Code | Description |
|---|---|
| 0xc000015b | Requested Logon type was not granted. Windows provides for multiple Logon types. Interactive Logon is logging on at the console of the machine. Network Logon is connecting to a resource on a remote machine. The function has requested to log on with certain credentials that do not permit the requested log on type. |
| 0xc000023c | The network is unreachable. There is an underlying networking problem. Start by troubleshooting basic connectivity. When using TCP/IP, check to make sure that the address, subnet mask and gateway are correct. To determine the location of the problem, use the ping utility to ping the loop back adapter first, then the IP address of the current machine that is having a problem, then the default gateway and then an address that is past the gateway. |
| 0xc000026d | DFS is unavailable. Troubleshoot basic connective to DFS shares. If possible, use a network sniffer, such as Network Monitor, to assess whether the problem is on the server or the client. |
| 0xc00000be | Bad Network Path. The path that was used to connect to another machine was invalid or there could be name resolution issues. Verify the path does not have invalid characters. Look for additional errors in the log that could support a name resolution problem. These errors can include 53 and 1355. Trying doing a directory of the path from a command prompt using the fully qualified name, the NetBIOS name and lastly the IP address itself. Also rule out the possibility of IPSec. |
| 0xc0000064 | No such user. A request specified a user account that is not on the system the function targeted. Verify that the user does exist. If the user requested resides in Active Directory make sure Active Directory replication is working, and that the user is present on the authenticating domain controller. |
| 0xc000006a | Wrong Password. The function attempted to use an invalid password for the credentials it was impersonating or in the context of. Check and make sure Active Directory replication is working. A large number of these events could be a service account that has had a password change but the entry in the machine has not been updated with the new password. Additionally, multiple versions of these events can trigger account lockouts. |

| | |
|---|---|
| 0x80090303 | Target Unknown. Typically, this error occurs when a security function fails. Start your troubleshooting by looking at the calling function to try to determine what Group Policy was attempting to do when it called the security related function. Look at the event logs to see if the computer is having other security related problems such as policy not applying, logon problems, or prompting for credentials when the credentials are good. Another area where you can see this is with Kerberos when the Service Principal Name that is being requested is not listed in the database. Use a network trace to determine if this is the cause. Also, turn on Kerberos logging at the targeted computer. |
| 0x8009030c | Logon Denied. Group Policy impersonates both the computer and the user when it determines the scope of policy, meaning it acts on behalf the user. This could be a machine account needing its password reset or a user account that has some problem. Check Active Directory to make sure the user account is set up correctly. Use a utility, like netdom or nltest, to test the computer account's password. |
| 0x80090324 | Time Skew. The majority of the time this is a Kerberos issue. Make sure your client, server and the KDC are within five minutes of each other. Verify that time zones and date are correct too. Check the event log to see if you are having any W32time errors. Kerberos defaults to five minutes of allowable time skew. Any time past that causes errors. |
| 0x80090311 | No authenticating authority. Normally, this implies that the machine cannot contact a Domain Controller or from a Kerberos perspective (that is, a KDC). Troubleshoot this by turning on Kerberos Logging and getting a network monitor trace to determine where the computer is trying to contact for authentication. If you are traversing a trust then Netlogon log on Domain Controllers in both domains and if possible run a simultaneous trace to help pinpoint the failure. |

**LDAP ERROR CODES**

Some of the messages return LDAP Error codes. You can recognize these error codes because the function usually starts with **ldap_**.

| Error Code | Description |
|---|---|
| 85 | Timeout. Verify network connectivity. Ping the target computer and check the response time. Verify that the targeted server is responding on port 389 and port 3268 if it is a Global Catalog server. |
| 82 | Local Error. Check the Directory Services event log and make sure Active Directory is healthy on this machine. Check for memory consumption as well as CPU time for LSASS. If these checks do not turn up any other errors, turn on Directory Services logging. |
| 53 | Unwilling to Perform. The specific command was refused by the LDAP Server. The most direct way to find the problem is to turn on Directory Service Logging on the targeted server that returned the error message. By getting a network monitor trace while this happening, you can identify the request that the client is sending to the server. The additional logging helps you determine why the server is returning the message. This is not always an error message. It is uncommon, but occasionally Group Policy requests that the computer do something that it cannot, which causes this error. |
| 7 | Authentication Method not supported. Look at the type of Authentication you are using to connect to the targeted LDAP server. In most cases, you should begin Kerberos Troubleshooting. Turn on Kerberos Logging on the client and the KDC. Also, a network trace might help you locate the problem. |
| 81 | Server Down. The targeted LDAP Server is not responding. Check the network, the proper LDAP ports for the given server, and any router connections. Start investigating the health of the targeted server. |

# Interpreting Security Settings log files

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic explains how to get started in interpreting Security Settings log files.

**Using Winlogon.log**
When you enable logging for the Security Settings client-side extension, a file named Winlogon.log is created in the %SYSTEMROOT%\Security\Logs folder. Winlogon.log provides debug logging information for security events and errors (Scecli.dll).

You can use Winlogon.log to look for security configuration-related errors such as the following:

- Mapping problems between account names and security IDs.
- Incorrect permissions for performing tasks (Access denied errors).
- Problems with trust relationship between the primary domain and trusted domains.
- Extended errors. These refer to generic errors that may be caused by a number of different problems.

For information about enabling Scecli logging, see the "Enabling Logging for Security Settings" section in Configuring Administrative Workstation Settings for Group Policy Troubleshooting.

To reproduce the policy failure, refresh the policy settings by typing the following at the command prompt: **gpupdate /force**.

**Example**
After you enable logging for Scecli, you can examine the Winlogon.log file to look for errors such as an account causing a Scecli failure or the wrong policy applying to a computer. For example, you can determine if an account was removed from policy or deleted, the account name contains typographical errors, or an account name cannot be resolved in Restricted Groups policy and as a result the security policy settings do not apply.

To identify the problem account, problem setting, and the problem GPO, perform the following tasks:

1. To find the problem account, type the following at the command prompt, and then press ENTER:

   **find /I "cannot find" %SYSTEMROOT%\security\logs\winlogon.log**

   The **Find** output identifies the problem account names, for example, "Cannot find MichaelAlexander." In this example, the user account MichaelAlexander does not exist in the domain, or it has a different spelling, for example, MichelleAlexander.

   Determine why this account cannot be resolved. Look for issues such as typographical errors, a deleted account, the wrong policy applied to this computer, or a trust problem. You may decide that the account has to be removed from the policy.

2. After you find the account problem, you can find the affected policy setting and GPO. To do this, type the following at the command prompt on the computer that is producing the Scecli error:

   "**c:\>find /i ?<account name>? %SYSTEMROOT%\security\templates\policies\gpt*.***" (without the quotation marks)

   The following example shows the **Find** command syntax and the results:

   **c:\>find /i "MichaelAlexander" %SYSTEMROOT%\security\templates\policies\gpt*.***

   ```
   D:\WINNT\SECURITY\TEMPLATES\POLICIES\GPT00000.DOM

   D:\WINNT\SECURITY\TEMPLATES\POLICIES\GPT00001.INF
   ```

```
D:\WINNT\SECURITY\TEMPLATES\POLICIES\GPT00002.INF SeInteractiveLogonRight = TsInternetUser,*S-1-5-32-549,*S-1-5-32-550,MichaelAlexander,*S-1-5-32-551,*S-1-5-32-544,*S-1-5-32-548

D:\WINNT\SECURITY\TEMPLATES\POLICIES\GPT00003.DOM
```

In this case, GPT00002.INF is the cached security template from the problem GPO that contains the problem security setting. It also identifies the problem setting as **SeInteractiveLogonRight**. The display name for **SeInteractiveLogonRight** is **Logon locally**.

For more information about the security settings constants (for example, SeInteractiveLogonRight) and their associated display names (for example, Logon locally), see the Microsoft Windows 2000 Server Resource Kit, "Distributed Systems Guide." The map of security settings constants is in the "User Rights" section of the Appendix.

3. After you identify the cached security template, you can determine which GPO contains the problem setting by searching the cached security template for the text "GPOPath=".

   In this example, you would see the following: **GPOPath={6AC1786C-016F-11D2-945F-00C04FB984F9}\MACHINE**. The GPO's GUID is {6AC1786C-016F-11D2-945F-00C04FB984F9}.

4. You can use the Resource Kit utility **Gpotool.exe** to find the friendly name for the GPO GUID. Type the following at the command prompt, and then press ENTER: **gpotool /verbose**.

   Search the output for the GUID that you identified in step 3. The four lines that follow the GUID contain the friendly name of the policy. For example:

```
Policy {6AC1786C-016F-11D2-945F-00C04FB984F9}

Policy OK

Details:

----------------------------------------------------------

DC: domcntlr1.wingtiptoys.com

Friendly name: Default Domain Controllers Policy
```

**GPOTool.exe** ships with the Microsoft Windows® 2000 Server Resource Kit and is also available as a free download from Gpotool.exe: Group Policy Verification Tool on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=17911). For more information see the Windows 2000 Server Resource Kit.

# Additional Resources for Group Policy

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For information about identifying and fixing Group Policy problems, see the following resources:

- Managing Windows XP Service Pack 2 Features Using Group Policy on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=29269)
    - This white paper explains how to deploy Group Policy settings that are delivered with Microsoft Windows XP Service Pack 2 (SP2) as well as how to assess and use new policy settings.
- Group Policy Collection on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=35202)
    - This page introduces Group Policy management concepts and architecture, summarizes the areas included in the Group Policy collection, and describes Group Policy scenarios.
- Tools and Settings Collection on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=4972)
    - These resources provide technical information about tools and settings related to Microsoft Windows Server 2003.
- Group Policy Deployment Kit: Designing a Managed Environment on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=15311)
    - This book provides the foundation you need to understand and deploy Group Policy within Windows Server 2003. It will help you to manage a diverse set of organizational and computing needs and to centrally manage domain-wide and forest-wide security, servers, computers, and users.
- Group Policy Infrastructure White Paper (Windows Server 2003) on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=60547)
    - This white paper explains Group Policy infrastructure and shows how Group Policy Management Console fits into this infrastructure. The white paper also includes detailed information about Group Policy processing as well as many best practices useful to Group Policy administrators.
- Administering Group Policy with GPMC on the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=14320)
    - This white paper provides an overview of Group Policy Management Console (GPMC) and outlines functionality necessary to administer Group Policy with the GPMC. Highlights of the white paper include new functionality such as backup, restore, import, copy, and reporting of Group Policy objects (GPOs).

# Security Technologies Operations Guide

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This Operations Guide provides troubleshooting information for security technologies in Microsoft® Windows Server™ 2003 and Windows® 2000 operating systems.

**In this guide**


- Troubleshooting Kerberos
- Troubleshooting Access Control

# Troubleshooting Kerberos

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides troubleshooting information for Kerberos authentication in Windows 2000 and Windows Server 2003 domains. It is designed to help you identify and resolve problems that are related to the Kerberos version 5 authentication protocol.

**In this guide**


- Introduction to Troubleshooting Kerberos
- Verifying Computer Settings for Troubleshooting Kerberos
- Configuring a Computer for Troubleshooting Kerberos
- Troubleshooting Kerberos Problems
- Additional Resources for Troubleshooting Kerberos



**Acknowledgments**

Produced by: Microsoft Windows Server User Assistance team

Project Writer: Brian Reisman

Project Editor: Carolyn Eller

Contributing Writers: Michiko Short, Mike Danseglio

Contributing Editors: David Longmuir, Vincent Abella

Technical Reviewers: David Christiansen, Xin Fan, Mohamed Fathalla, JK Jaganathan, Steve Light, Soumitra Sengupta, Tim Springston, Todd Stecher, Jonathan Stephens, Darol Timberlake, Joseph Vasil, Liqiang (Larry) Zhu, Leon Arber

# Introduction to Troubleshooting Kerberos

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are using this guide for the first time, review the following sections of this introduction.

**When to Use This Guide**
You should use this guide when:

- You have an issue related to authentication such as console logon, network logon, access to network resources, or remote access and Kerberos is the default authentication protocol.
- The System event log shows errors from any services that provide authentication, such as Kerberos Key Distribution Center (kdc), Local Security Authority Server (LsaSrv), or Net Logon (Netlogon).
- Failure audits in the Security event log show that the Kerberos protocol was being used when a logon failure occurred.

Do not use this guide to find out how to perform a task, such as enabling the Kerberos protocol or adding a program or port to the Kerberos exceptions list.

This guide assumes a basic understanding of Windows authentication using the Kerberos protocol, how it works, and why your organization uses it to authenticate users and services. You should also have a thorough understanding of how Kerberos authentication is deployed and managed in your organization. This includes an understanding of the mechanism your organization uses to configure and manage authentication settings generally and Kerberos settings specifically.

**How to Use This Guide**
This guide is divided into three sections. Each section addresses a type of problem, ranging from less complex to more complex.

Verifying Computer Settings for Troubleshooting Kerberos

This section provides a list of prerequisites and settings that must be verified before you troubleshoot. Read this section first.

Configuring a Computer for Troubleshooting Kerberos

This section describes how to configure your computer for troubleshooting. Read this section if you cannot find a solution in the Quick Fixes section.

Troubleshooting Kerberos Problems

This section provides step-by-step diagnostic procedures and possible solutions that help you identify and fix Kerberos-related problems.

# Verifying Computer Settings for Troubleshooting Kerberos

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin troubleshooting, verify that your computer is configured properly and that the Kerberos authentication infrastructure is set up and running properly.

**Settings to Verify Before You Troubleshoot**
Verify all of the following items:

Make sure you have administrative rights on the computer you are troubleshooting.

Install all critical updates and security updates for Windows Server 2003.

Update all of your software, including non-Microsoft software.

Restart the computer if you are running a server operating system.

Make sure that required services and server are available.

Make sure that the clocks are synchronized across the domain.

**Make sure you have administrative rights on the computer you are troubleshooting.**
You cannot modify Kerberos-related settings unless you are a member of the Administrators group on the computer that you are administering.

### To verify that you are a member of the Administrators group

1. Open the Computer Management snap-in.
2. In the console tree, double-click **Local Users and Groups**, and then click **Groups**.
3. In the details pane, double-click **Administrators** and verify that your account name or a group to which your account is a member appears in the **Members** list.

**Install all critical updates and security updates for Windows Server 2003.**
Some updates might be required for the Kerberos protocol to function properly.

### To verify that you have all critical updates and security patches

- Click **Start**, click **Windows Update**, and then follow the instructions that appear on your screen.

**Update all of your software, including non-Microsoft software.**
The Kerberos protocol might not function properly with some programs unless you update the programs with the most recent service pack or software update.

**Restart the computer if you are running a server operating system.**
If you are running any Windows Server 2003 operating system, and you have made configuration changes that affect Kerberos authentication, you need to restart your computer.

**Make sure that required services and servers are available.**
The Kerberos authentication protocol requires a functioning domain controller, Domain Name System (DNS) infrastructure, and network in order to work properly. Verify that you can access these resources before you begin troubleshooting the Kerberos protocol.

**Make sure that the clocks are synchronized across the domain.**
Many network services, including Kerberos authentication are dependant on time synchronization throughout the domain. You can manually synchronize a computer with the time on the domain.

### To synchronize the computer's time with the current time on the domain

1. Click **Start**, and then click **Run**.
2. Type **net time /domain /set**, and then click **OK**.

# Configuring a Computer for Troubleshooting Kerberos

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can use advanced troubleshooting techniques to identify and resolve Kerberos related problems, you need to configure your computer for troubleshooting. In addition, you need a basic understanding of troubleshooting concepts, procedures, and tools.

**Configuration Tasks for Troubleshooting**
To configure your computer for troubleshooting, perform the following tasks:

Install the Windows Support Tools

Configure Tools for Troubleshooting

Enable Account Logon Failure Auditing

**Install the Windows Support Tools**
The Windows Support Tools assist support personnel and network administrators in managing their networks and troubleshooting problems. They are not installed with the Windows operating system; you must install them separately from the \Support\Tools folder of the Windows operating system CD.

**Configure Tools for Troubleshooting**
Configure the following tools as needed:

- Ldifde

  Ldifde.exe is present on domain controllers but can be copied and used on client computers running Windows XP and Windows Server 2003. Ldifde provides a method to quickly extract and display certain Service Principal Names (SPNs) in a forest or domain.

- LDP

  LDP.exe is included with the Windows Support Tools.

- Setspn

  Setspn.exe is included with the Windows Support Tools.

- Tokensz

  For information about downloading Tokensz, see Tokensz.exe on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42933).

**Enable Account Logon Failure Auditing**
You can configure the Windows Server 2003 audit log so that it records information about failed logon attempts. To record this information in the audit log, you must enable the **Audit account logon events** policy.

**To enable Account Logon Failure auditing**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage the Audit Policy settings in your organization.
2. In the console tree, open **Computer Configuration**, open **Windows Settings**, open **Security Settings**, open **Local Policies**, and then click **Audit Policy**.
3. In the details pane, double-click **Audit account logon events**.
4. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.

5. To audit unsuccessful attempts, select the **Failure** check box.

# Troubleshooting Kerberos Problems

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Any authentication problem with console logon, network logon, access to network resources, or remote access might indicate some sort of Kerberos error because the Kerberos protocol is the default authentication protocol.

To determine whether a problem is occurring with Kerberos authentication, check the System event log for errors from any services such as Kerberos, kdc, LsaSrv, or Netlogon that provide authentication. If any such errors exist, there might be errors associated with the Kerberos protocol as well. In addition, failure audits in the Security event log might show that the Kerberos protocol was being used when a logon failure occurred.

**Where to start**

In most cases, a Kerberos-related error is a symptom of another service failing. The Kerberos protocol relies on many services that must be available and functioning properly in order for any authentication to take place. The following is a list of the services or conditions that you will want to make sure are functioning properly before you scrutinize the Kerberos protocol.

- Make sure that the network infrastructure is functioning properly and that all computers and services can communicate.
- Make sure that a domain controller is accessible.
- Make sure that DNS is configured properly and resolving host names and services appropriately.
- Make sure that the clocks are synchronized across the domain.

**To synchronize the computer's time with the current time on the domain**

1. Click **Start**, and then click **Run**.
2. Type **net time /domain /set**, and then click **OK**.

| ⚠Caution |
| --- |
| You should not use the Net time command to configure or set time when the Windows Time service is running. Also, running the command **Net time /querysntp** will display the name of an NTP server with which a computer is configured to synchronize, but that NTP server is used only when the computer's time client is configured as NTP or AllSync. Most domain member computers have a time client type of NT5DS, which means they synchronize time from the domain hierarchy. The only typical exception to this is the domain controller that functions as the primary domain controller (PDC) emulator operations master of the forest root domain, which is typically configured to synchronize time with an external time source. To view the time client configuration of a computer, run the **W32tm /query /configuration** command (from an elevated Command Prompt in Windows Vista and Windows Server 2008) and read the Type line in the command output. For more information, see How Windows Time Service Works (http://go.microsoft.com/fwlink/?LinkId=117753) and Windows Time Service Tools and Settings (http://go.microsoft.com/fwlink/?LinkID=42984). |

If you have examined all these conditions and are still having authentication problems or Kerberos errors, you need to look further for the solution to your problems. Problems can be caused by the way the Kerberos protocol is configured or they can be caused by the way other technologies that work with the Kerberos protocol are configured.

**In this section**


Service Logons Fail Due to Incorrectly Set SPNs

Authentication Uses NTLM instead of Kerberos

Authentication Fails in a Mixed Windows and UNIX Environment

Authentication Fails Due to User PAC

Authentication Using UDP Causes Errors

Authentication Errors are Caused by Unsynchronized Clocks

# Service Logons Fail Due to Incorrectly Set SPNs

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Kerberos authentication is not possible for services without properly set Service Principal Names (SPNs). SPNs are unique identifiers for services running on servers. Each service that uses Kerberos authentication needs to have an SPN set for it so that clients can identify the service on the network. It is registered in Active Directory under a user account as an attribute called **Service-Principal-Name**. The SPN is assigned to the account under which the service the SPN identifies is running. Any service can look up the SPN for another service. When a service wants to authenticate to another service, it uses that service's SPN to differentiate it from other services running on that computer.

In general, only one SPN should be set for each service. Multiple SPNs can cause clients to connect to the wrong system or the ticket may be encrypted with the wrong key.

**Cause**

No SPN has been set for a specific service or it has been set incorrectly. If an SPN is not set for a service, then clients have no way to locate that service. Thus, common results of not setting an SPN are the following errors:

- "0x6 - KDC_ERR_C_PRINCIPAL_UNKNOWN: Client not found in Kerberos database"
- "0x7 - KDC_ERR_S_PRINCIPAL_UNKNOWN: Server not found in Kerberos database"

These two errors usually indicate that an SPN has not been set correctly. There are many other errors for which the cause might be a missing or incorrectly set SPN including:

- "0x29 - KRB_AP_ERR_MODIFIED: Message stream modified"
- "0x1B - KDC_ERR_MUST_USE_USER2USER: Server principal valid for user2user only"
- "0x8 - KDC_ERR_PRINCIPAL_NOT_UNIQUE: Multiple principal entries in database"
- "0x1F - KRB_AP_ERR_BAD_INTEGRITY: Integrity check on decrypted field failed"

**Solution**

To enable a service to authenticate properly, you need to set an SPN for that service.

Because multiple services can run simultaneously under the same account, setting an SPN requires four pieces of information that will make the SPN unique. These four pieces of information uniquely identify any service running on a network and can be used to mutually authenticate to any service.

An SPN itself consists of *ServiceClass***/***Host***:***Port*, where *ServiceClass* is the service class of the SPN.

You can use the following procedures to configure an SPN.

**To add an SPN**

- Type the following at a command prompt:

  **setspn -A** *ServiceClass* **/** *Host* **:** *Port AccountName*

**To delete an SPN**

- Type the following at a command prompt:

  **setspn -D** *ServiceClass* **/** *Host* **:** *Port AccountName*

**To view the SPNs that are registered for an account**

- Type the following at a command prompt:

  **setspn -L** *AccountName*

**To reset the default SPN registrations for the host names for an account**

- Type the following at a command prompt:

  **setspn -R** *AccountName*

**Setspn Parameters**
The following section describes the parameters for the Setspn.exe tool:

**-A**
Adds the specified SPN for the computer.

**-D**
Deletes the specified SPN for the computer.

**-L**
Lists the currently registered SPN for the computer. This parameter is the default: setspn run without this parameter also lists the registered SPNs.

**-R**
Resets the default SPN registrations for the host names for Computer.

***ServiceClass***
SPN service class assigned to the service.

***Host***
The computer to which the SPN belongs is all the names by which a computer on which the service is running can be referenced. This usually includes a NetBIOS name, a fully qualified domain name (FQDN), and any aliases that might have been assigned to this computer. A separate SPN will need to be set for each name by which the computer can be referenced, with the Host parameter changing respectively.

***Port***
The port that the service is running on. It is recommended to include the port even if it is a default port for the service.

***AccountName***
The name of the domain account under which the service runs. If the service runs as Local System or the network service, you usually do not need to set an SPN explicitly for the service because most common SPN service classes will automatically be mapped to the HOST/ SPN which is in turn automatically generated for each computer account.

For more information about using Setspn.exe to manipulate service principal names for accounts, see Setspn.exe: Manipulate Service Principal Names for Accounts on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43030).

**Cause**
Duplicate SPNs are set for the service, and the Kerberos Distribution Center (KDC) does not know which one is correct. Thus, common results of duplicate SPNs are the following errors:

- 0x8 - KDC_ERR_PRINCIPAL_NOT_UNIQUE: Multiple principal entries in database

- 0x29 - KRB_AP_ERR_MODIFIED: Message stream modified errors

These two errors usually indicate that there are duplicate SPNs. Furthermore, there are many other errors for which the cause might be duplicate SPNs. Kerberos authentication is not possible without properly set SPNs.

**Solution**
To enable a service to authenticate properly, you need to make sure that the service has only one SPN. First, find the duplicate SPNs and then delete one of them.

**To find the duplicate SPNs by using the LDP tool**

1. Click **Start**, then click **Run**, type **LDP**, and then click **OK**.

2. On the **Connection** menu, click **Connect**.

3. If you are on the domain controller, leave the default settings, and then click **OK**. If you are not on the domain controller, type the domain controller name in the Server text box, and then click **OK**.

4. On the **Connection** menu, click **Bind**.
5. Type **User**, **Password**, and **Domain** in the corresponding text boxes, and then click **OK**.
6. On the **View** menu, click **Tree**.
7. In the **Tree View** dialog box, type the base distinguished name in the **BaseDN** text box or select it from the pull-down menu.
8. On the **Browse** menu, click **Search**.
9. In the **Search** dialog box, type the base distinguished name in the **BaseDN** box or select from the pull-down menu.
10. In the **Search** dialog box, type the following in the **Filter** box:

    serviceprincipalname=*SPN/FQDN*

    For SPN, type the Service Principal Name that the error refers to - for example, HOST for computer accounts, HTTP for Web services.
11. Under **Scope**, click the **Subtree** option.
12. Click **Run**.

For more information about using Ldp.exe to search Active Directory, see Using Ldp.exe to Find Data in the Active Directory on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=23064).

If there are duplicate SPNs listed, delete the appropriate SPN by using the following command:

**To delete an SPN**

- Type the following at a command prompt:

  **setspn -D** *ServiceClass* **/** *Host* **:** *Port AccountName*

For information about the parameters for setspn.exe, see Setspn parameters.

For more information about how to use the Windows Support Tools, see Introduction to Support Tools on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=38906).

# Authentication Uses NTLM instead of Kerberos

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You might find that the Security log recorded an event in which logon occurred by using NTLM when it should have occurred by using Kerberos authentication.

**Cause**
The system attempts authentication by using the Kerberos protocol but it fails. As a result, the system attempts to authenticate by using NTLM. Windows Server 2003, Windows XP, and Windows 2000 use an algorithm called Negotiate (SPNEGO) to negotiate which authentication protocol is used. Although the Kerberos protocol is the default, if the default fails, Negotiate will try NTLM.

**Solution**
Investigate by examining errors or warnings in the event logs.

**Cause**
A call to Negotiate returns NTLM as the only protocol available. This situation is much more difficult to diagnose. There are two common causes of this situation — when Internet Explorer is being used and the Kerberos protocol is not being attempted:

- **Enable Integrated Windows Authentication (requires restart)** setting is not selected in Internet Explorer 6.0.
- Internet Explorer is accessing a site in the Internet zone instead of the intranet zone.

**Solution**
Internet Explorer 6.0 will, by default, not attempt to use the Kerberos protocol to authenticate to any site.

**To enable Internet Explorer 6.0 to respond to a negotiate challenge and perform Kerberos authentication:**

1. In Internet Explorer, on the **Tools** menu click **Internet Options**.
2. Click the **Advanced** tab, in the **Security** group select the **Enable Integrated Windows Authentication (requires restart)** check box, and then click **OK**.
3. Restart Internet Explorer.

For more information, see Unable to Negotiate Kerberos Authentication After Upgrading to Internet Explorer 6 in the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=23045).

The second common cause is that Internet Explorer 6.0 is attempting to access a site located in the Internet zone. Internet zone sites are prevented from using Integrated Windows authentication because these protocols do not typically work through Web proxies, among other reasons. If a site is located in the Internet zone, Internet Explorer 6.0 does not attempt to use Kerberos authentication, and automatically tries NTLM. In all versions of Internet Explorer, when accessing a Web site to which you want to use Kerberos authentication, you must verify that the Web site appears as being in the local intranet zone. An icon in the lower right corner of the Internet Explorer window indicates what zone a Web site is in. It displays “Internet” for the Internet zone and “Local Intranet” for the intranet zone. If the Web site appears as being in the Internet zone, you must manually add the site to the local intranet sites list.

**To add an Internet site to the local intranet sites list**

1. On the **Tools** menu, click **Internet Options**.
2. Click the **Security** tab, then click **Local Intranet**, then click **Sites**, and then click **Advanced**.
3. In the box under **Add this Web site to the zone**, type the name of the Web site that you want to authenticate with Kerberos authentication, and then click **Add**.
4. Click **Close**, and then click **OK**.

After you perform this procedure, if you find that NTLM authentication is still being used, or that Kerberos is not even being attempted in a situation where Kerberos authentication should be used, contact Product Support Services for help in diagnosing the problem.

# Authentication Fails in a Mixed Windows and UNIX Environment

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The logon from a computer running Windows NT 4.0 in a mixed Windows and UNIX environment fails and the Active Directory domain controller returns an Access denied error.

In an environment where a trust exists between a Kerberos realm and an Active Directory domain, authentication data can come from one of two sources, either a UNIX Key Distribution Center (KDC) or the Active Directory domain controller.

If the authentication data is coming from a UNIX KDC, Windows users have account mappings set up for them to map their UNIX user account to a Windows user account. Normally, the password on the Windows account does not matter, because all authentications are done by the UNIX KDC.

**Cause**
The size of a ticket is too large to be transmitted reliably by means of the User Datagram Protocol (UDP). In a Windows environment, this message is purely informational. A computer running a Windows operating system will automatically try TCP if UDP fails.

In a mixed UNIX and Windows authentication environment, many errors can be attributed to the UNIX KDC not running the latest MIT distribution of the Kerberos protocol. One of the errors that can result is "0x34 - KRB_ERR_RESPONSE_TOO_BIG: Response too big for UDP, retry with TCP."

**Solution**
If this error occurs in a mixed operating system environment, consult the Vendor specific documentation to upgrade the UNIX KDCs to the latest MIT distribution of the Kerberos protocol, which supports TCP connections if UDP fails.

**Cause**
The "0x19 - KDC_ERR_PREAUTH_REQUIRED: Additional pre-authentication required" error often occurs in UNIX interoperability scenarios. MIT Kerberos clients do not request pre-authentication when they send a KRB_AS_REQ message. If pre-authentication is required (the default), Windows operating systems send this error. Most MIT-Kerberos clients respond to this error by giving the pre-authentication, in which case the error can be ignored, but some clients might not respond in this way.

**Solution**
Set the **Do not require Kerberos pre-authentication** flag on the user's account. Alternatively, consider upgrading to the most recent MIT reference distribution of Kerberos authentication.

**To set the Do not require Kerberos pre-authentication flag on the user's account using Active Directory Users and Computers:**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**. (Or, click the folder that contains the user account.)
3. Right-click the user account, and then click **Properties**.
4. On the **Account** tab, scroll through the **Account options** and select the **Do not require Kerberos pre-authentication** checkbox, and then click **OK**.

| **Note** |
| --- |
| To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated an assignment of administrative responsibility to a user, computer, group, or organization. |

**Cause**
When the KDC, server, or client receives a packet for which it does not it have a key of the appropriate encryption type, this can cause the "0xF - KDC_ERR_SUMTYPE_NOSUPP: KDC has no support for checksum type" error. The result is that the computer is unable to decrypt the ticket. This error is common in UNIX interoperability scenarios when a new account is created. The new

account might have an incompatible key associated with it.

**Solution**
Change the password on the account to re-create the key, which should eliminate the error.

**Cause**
The UNIX KDC attempts to use 3DES to encrypt its tickets which can cause the "0xE - KDC_ERR_ETYPE_NOTSUPP: KDC has no support for encryption type" error.

**Solution**
Consult the vendor-specific documentation to configure the UNIX KDC to use another encryption type, such as DES or RC4. Windows operating systems do not support 3DES.

**Cause**
The target server does not support the encryption type used by the KDC in some UNIX and Windows interoperability scenarios.

**Solution**
Consult the vendor-specific documentation to configure the target server to support the RFC standard encryption type RC4 or contact the vendor for assistance.

**Cause**
The Windows NT 4.0 operating system does not support Kerberos authentication. Therefore, if there are Windows NT–based computers on the network running services, any authentications involving these computers occur by using NTLM and these authentications are conducted by the domain controller. In this scenario, the password on the domain controller must match the password stored on the UNIX KDC. The passwords must match because the fallback to NTLM occurs transparently. If the passwords do not match, the domain controller returns an Access denied error because the user has provided a password that does not match the one stored on the user's Active Directory account.

**Solution**
Consult the vendor-specific documentation to reset the password on the account that the user’s UNIX principal is mapped to in order to match the password stored on the UNIX KDC.

# Authentication Fails Due to User PAC

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In certain scenarios when a user is a member of more than 70 groups, Kerberos authentication might fail. In order to provide the proper authorization information to a user after the user logs on with Kerberos authentication, the KDC will transmit Privilege Attribute Certificate (PAC) data in the ticket-granting ticket (TGT). The PAC contains various types of authorization data including groups that the user is a member of, rights the user has, and what policies apply to the user. When the client receives a ticket, the information contained in the PAC is used to generate the user's access token.

**Cause**

In order to optimize performance, the buffer size for the PAC is pre-allocated. The pre-allocated buffer size is usually adequate to hold all the required authorization data. However, if a user has a very high group membership — from over 70 to over 120, depending on what groups — the size of the PAC might exceed the pre-allocated buffer size. In such a case, the system generates a memory allocation error, PAC creation fails, and the Kerberos ticket-granting service either fails to generate a valid ticket or generates a ticket with an empty PAC. This sort of error usually manifests itself in the form of a memory allocation error, which gets reported as "0x3C - KRB_ERR_GENERIC: Generic error." This also can result in the failure of clients to apply Group Policy settings.

**Solution**

You can use the following procedure on a Windows Server 2003-based computer to calculate the number of groups that a user is a member of by using the Kerberos Token Size (Tokensz.exe) tool.

**To calculate the groups that a user belongs to by using the Tokensz.exe tool**

- Type the following at a command prompt:

  **tokensz /calc_groups** *ClientName* **/user:** *UserName* **/domain:** *UserDomain* **/password:** *ClientPassword* **/system**

The following describes the parameters used with this command:

**/calc_groups**
This option is only available if you are using Windows Server 2003 KDCs and running the tool on a computer running Windows Server 2003. The parameter passed to this option is a user and (optionally) a domain that this user is a member of. Tokensz will list all the groups that this user is a member of.

***ClientName***
The user that you are listing the groups for

**/user: *UserName***
The user name of an account with permission to query the group membership of the user.

**/domain: *UserDomain***
The domain that the above account exists in.

**/password: *ClientPassword***
The password that corresponds to the user account specified in the **/user** parameter.

**/system**
You must have administrator rights to use this switch, which runs the test under the Local System context of the computer that the test is being run on. In these cases, you must use the [**/user:***UserName*] [**/domain:***UserDomain*] [**/password:***ClientPassword*] switches, because the tool will be running under the Local System context. This switch is useful to see what groups are present in the Local System logon context.

| **Note** |
| --- |
| This command is available only if you are using Windows Server 2003 KDCs and running the tool on a computer running Windows Server 2003. The parameter passed to this option is a user and (optionally) a domain that this user is a member of. Tokensz.exe lists all the groups that this user is a member of. |

You can solve this problem in two ways:

- Reduce the number of groups that the user is a member of. Because nested groups are expanded out within the PAC, the actual number of groups that the user is a member of might be greater than you suspect.

- Alternatively, install a hotfix that will enable you to set the maximum size of a Kerberos token by using the registry. For information about installing this hotfix, see New Resolution for Problems That Occur When Users Belong to Many Groups on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=23044).

# Authentication Using UDP Causes Errors

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Packets sent along the network by using the User Datagram Protocol (UDP) sometimes get lost or fragmented on the way to their destination.

**Cause**
The UDP protocol can also sometimes cause the "KRB_ERR_RESPONSE_TOO_BIG: Response too big for UDP, retry with TCP" error.

The size of a ticket is too large to be transmitted reliably by means of UDP. In a Windows environment, this message is purely informational. A computer running a Windows operating system will automatically try TCP if UDP fails.

**Solution**
If this error occurs in a mixed operating systems environment, consult your vendor-specific documentation to upgrade the UNIX KDCs to the latest MIT distribution of the Kerberos protocol, which supports TCP connections if UDP fails.

For information about forcing Kerberos to use TCP, see How to Force Kerberos to Use TCP Instead of UDP on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=23043).

**Cause**
By default, Kerberos authentication uses UDP to transmit its data. UDP provides no guarantee that a packet sent along the network will reach its destination intact. Thus, in environments with a high amount of network congestion it is common for packets to get lost or fragmented on the way to their destination.

If you have discovered the "0x7 - KDC_ERR_S_PRINCIPAL_UNKNOWN: Server not found in Kerberos database" error message and the Service Principal Name (SPN) is set, or if the request failed for an initial TGT (requesting a TGT does not require any SPNs to be set manually), UDP fragmentation might be causing the failure.

If you have discovered the "0x28 - KRB_AP_ERR_MSG_TYPE: Invalid msg type" error message, then UDP is being attempted with User-to-User protocol. User-to-User is an extension of Kerberos authentication that enables secure servers to be run on personal computers.

**Solution**
Because the only way to decrease the likelihood of UDP fragmentation occurring is to reduce network traffic — a usually impractical solution — it is almost always better to configure the Kerberos authentication service to use TCP instead of UDP. TCP provides a guarantee that a packet that is sent will reach its destination intact and can therefore be used in any network environment.

In order to support the User-to-user extension to the Kerberos protocol you must force Kerberos authentication to use TCP.

For information about forcing Kerberos to use TCP, see How to Force Kerberos to Use TCP Instead of UDP on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=23043).

# Authentication Errors are Caused by Unsynchronized Clocks

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Clients can be prevented from authenticating by the mechanisms that Kerberos authentication uses to prevent "replay" attacks. In a replay attack, a malicious user captures the network traffic and replays it to trick the authenticating server into accepting the attacker as a legitimate user who is providing credentials.

Kerberos authentication prevents a replay attack by using two mechanisms:

- The Kerberos client on the local computer encrypts a timestamp inside the authenticator and then sends it to the Key Distribution Center (KDC). If the KDC verifies that the time it decrypts from the authenticator is within a specified amount of the local time on the KDC (the default is five minutes), the system can assume that the credentials presented are genuine.

- All tickets issued by the KDC have an expiration time. Thus, if a ticket is compromised, it cannot be used outside of a specified time range — usually short enough to make the risk of a replay attack minimal.

Because of these mechanisms, Kerberos authentication relies on the date and time that are set on the KDC and the client. If there is too great a time difference between the KDC and a client requesting tickets, the KDC cannot determine whether the request is legitimate or a replay. Therefore, it is vital that the time on all of the computers on a network be synchronized in order for Kerberos authentication to function properly.

Clock skew can be easily diagnosed by reviewing the information in the System log.

**Causes**
If the ticket presented to the server is not yet valid (in relationship to the server time), the "0x21 - KRB_AP_ERR_TKT_NYV: Ticket not yet valid" error will be recorded. The most probable cause is that the clocks on the KDC and the client are not synchronized.

In addition, the errors "0xB - KDC_ERR_NEVER_VALID: Requested start time is later than end time," "KDC_ERR_PREAUTH_FAILED: Pre-authentication information was invalid," and "KDC_ERR_NEVER_VALID: Requested start time is later than end time" are caused because the times on the KDC and the client do not match.

**Solution**
In this case no action is needed. In a Windows environment, the domain controller will supply the correct time on the domain controller to the Kerberos client. The Kerberos client then uses the correct domain controller time to attempt the authentication request a second time. Presuming that the user's credentials are valid, the user will be authenticated on the second try.

**Cause**
Typically, when the value for the **Maximum lifetime for user ticket Kerberos policy** setting is set too low, the "KRB_AP_ERR_TKT_EXPIRED: Ticket expired" error occurs.

**Solution**
Because ticket renewal is automatic, you should not have to do anything if you get this message.

**Cause**
If Kerberos authentication fails when it is being attempted between realms, the clocks in the KDC in the target realm and the KDC in the client realm might not be synchronized.

**Solution**
For non-Windows Kerberos realms, consult the operating system documentation to synchronize clocks on the KDCs in the two realms.

**Cause**
If a client requests the postdating of a Kerberos ticket, the "0xA - KDC_ERR_CANNOT_POSTDATE: Ticket not eligible for postdating" error can occur. Postdating is the act of requesting that a ticket's start time be set into the future. Postdating is not allowed therefore the client will not be authenticated.

**Solution**
No action is necessary. This behavior is by design.

**Cause**

Typically, when the difference between client timestamp in the authenticator or KRB_AS_REQ and the server is greater than **the Maximum tolerance for computer clock synchronization** setting in the domain policy, the "0x25 - KRB_AP_ERR_SKEW: Clock skew too great" error occurs. This error is logged if a client computer sends a timestamp whose value differs from that of the server's timestamp by more than the number of minutes found in the **Maximum tolerance for computer clock synchronization** setting in Kerberos policy.

**Solution**

Although the "0x25 - KRB_AP_ERR_SKEW: Clock skew too great" error might show up in the logs, it will not prevent a user from being authenticated. When this error is returned, the domain controller also supplies the correct time on the domain controller. The Kerberos client uses the correct domain controller time to attempt the authentication request a second time. Presuming that the user's credentials are valid, the user will be authenticated on the second try.

This error can more commonly occur as the number of portable, or frequently disconnected, computers in your network increases.

To avoid getting this error, you can change the value of the **Maximum tolerance for computer clock synchronization** setting to a higher value.

| ⚠**Caution** |
| --- |
| Beware that the higher you set this value, the more susceptible the network becomes to replay attacks. |

**To set Maximum tolerance for computer clock synchronization Kerberos policy by using Active Directory Users and Computers**

1. Click **Start**, click **All Programs**, click **Administrative Tools**, and then click **Active Directory Users and Computers**.
2. Right-click the domain whose policy you wish to change, and then click **Properties**.
3. Click the **Group Policy** tab, select the **Default Domain Policy**, and then click the **Edit** button to launch the Group Policy Object Editor.
4. Open **Computer Configuration**, open **Windows Settings**, and then open **Security Settings**.
5. Click **Account Policies**, double-click **Kerberos Policy**, and then double-click **Maximum tolerance for computer clock synchronization**.
6. Increase the value in the **Maximum tolerance** box, and then click **OK**.
7. You can wait for the policy change to propagate or you can run **gpupdate /force** on the client computers to force propagation immediately.

# Additional Resources for Troubleshooting Kerberos

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For additional information about troubleshooting Kerberos problems, see the following resources:

General resources

- Introduction to Support Tools on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=38906)
- Using Ldp.exe to Find Data in the Active Directory on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=23064)
- Ldifde on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=42656)
- Setspn.exe: Manipulate Service Principal Names for Accounts on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43030)
- Troubleshooting Kerberos Errors (http://www.microsoft.com/downloads/details.aspx?familyid=7dfeb015-6043-47db-8238-dc7af89c93f1&displaylang=en)
- Tokensz.exe on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42933)
- Kerberos protocol registry entries and KDC configuration keys in Windows Server 2003 in Microsoft Knowledge Base article 837361 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;837361)
- Http.sys registry settings for IIS in Microsoft Knowledge Base article 82019 (http://support.microsoft.com/kb/820129)
- Kerberos protocol registry entries that are defined under the Parameters subkey in Windows Server 2003, Windows XP, and Windows 2000 in Microsoft Knowledge Base article 244556 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;244556)

Kerberos Authentication and Ticket Sizing Issues

- Addressing Problems Due to Access Token Limitation (http://www.microsoft.com/downloads/details.aspx?familyid=22DD9251-0781-42E6-9346-89D577A3E74A&displaylang=en)
- New resolution for problems with Kerberos authentication when users belong to many groups in Microsoft Knowledge Base article 327825 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;327825)
- How to use Group Policy to add the MaxTokenSize registry entry to multiple computers on a domain controller that is running Windows Server 2003 or that is running Windows 2000 in Microsoft Knowledge Base article 938118 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;938118).
- Event ID 1053 is logged when you use the "Gpupdate /force" command, or you restart a Windows Server 2003-based domain controller in Microsoft Knowledge Base article 937535 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;937535)
- Error message when you use a Windows Server 2003-based domain controller to join a Windows XP-based client computer to a domain: "Not enough storage is available to complete this operation" in Microsoft Knowledge Base article 935744 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;935744)
- Domain join during an unattended setup fails with an unexpected error message in computers that are running Windows 2000, Windows XP, or Windows Server 2003 in Microsoft Knowledge Base article 920599 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;920599)

- You may receive an error message when you try to connect to an instance of SQL Server or SQL Server Desktop Engine by using Windows Authentication in Microsoft Knowledge Base article 867581 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;867581)

- "The information store could not be opened" error message when you try to retrieve the client permissions of a public folder in Exchange 2003 in Microsoft Knowledge Base article 842019 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;842019)

- The Mobile Information Server Message Processor crashes when a user tries to use Outlook Mobile Access in Microsoft Knowledge Base article 819529 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;819529)

- How to troubleshoot server ActiveSync HTTP error codes in Microsoft Knowledge Base article 330463 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;330463)

- How to troubleshoot inter-forest sIDHistory migration with ADMTv2 in Microsoft Knowledge Base article 322970 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;322970)

- DCOM Client May Put Memory on the Wire in Microsoft Knowledge Base article 300367 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;300367)

- SMS administrator issues after you modify the Kerberos MaxTokenSize registry value in Microsoft Knowledge Base article 297869 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;297869)

- Kerberos authentication may not work if user is a member of many groups in Microsoft Knowledge Base article 280830 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;280830)

- Internet Explorer logon fails due to an insufficient buffer for Kerberos in Microsoft Knowledge Base article 277741 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;277741)

- Internet Explorer Kerberos authentication does not work because of an insufficient buffer connecting to IIS in Microsoft Knowledge Base article 269643 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;269643)

- Group Policy may not be applied to users belonging to many groups in Microsoft Knowledge Base article 263693 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;263693)

- How to use SPNs when you configure Web applications that are hosted on IIS 6.0 in Microsoft Knowledge Base article 929650 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;929650)

- FIX: Error Message: "Timeout expired" Occurs When You Connect to SQL Server Over TCP/IP and the Kerberos MaxTokenSize is Greater Than 0xFFFF in Microsoft Knowledge Base article 313661 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;313661)

- FIX: You receive an "Error 413" error message when you try to log on to Site Builder by clicking to select the "Login as <CurrentUser>" check box in Content Management Server 2001 in Microsoft Knowledge Base article 867431 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;867431)

- INF: List of Bugs Fixed by SQL Server 7.0 Service Packs in Microsoft Knowledge Base article 313980 (http://support.microsoft.com/default.aspx?scid=kb;EN-US;313980)

Web applications and services, such as Internet Explorer, Internet Information Server, and Outlook Web Access may require additional Kerberos maximum settings adjusted, which are discussed in the following articles:

- MaxBufferedSendBytes – Used to resolve slow performance when Internet Server Application Programming Interface (ISAPI) applications or Common Gateway Interface (CGI) applications are hosted on Microsoft Internet Information Services (IIS) 6.0, as described in Microsoft Knowledge Base article 906977 (http://support.microsoft.com/kb/906977/)

- MaxClientRequestBuffer, MaxFieldLength, and MaxRequestBytes – Used to resolve HTTP 400 Bad Request (Request Header Too Long) errors, as described in Microsoft Knowledge Base article 920862 (http://support.microsoft.com/kb/920862).

- MaxPacketSize – Used to resolve Kerberos failures that results in "There are currently no logon servers available to service the logon request" error messages, as described in Microsoft Knowledge Base article (http://support.microsoft.com/kb/244474).

# Troubleshooting Access Control

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this guide**


- Introduction to Troubleshooting Access Control
- Verifying Computer Settings for Troubleshooting Access Control
- Access Control Quick Fixes
- Configuring a Computer for Access Control Troubleshooting
- Troubleshooting Access Control Problems
- Additional Resources for Troubleshooting Access Control


This guide provides troubleshooting information for access control in Microsoft® Windows® XP and Microsoft® Windows Server™ 2003. It is designed to help you identify and resolve problems that are related to access control.


**Acknowledgments**

Produced by: Microsoft Windows Server User Assistance team

Project Writer: Jennifer Allen

Project Editor: Carolyn Eller

Contributing Writers: Chris Gregory

Technical Reviewers: Satyajit Nath and Dave McPherson

# Introduction to Troubleshooting Access Control

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide explains how to troubleshoot access control. If you are using this guide for the first time, review the following sections of this introduction:

When to Use this Guide

How to Use this Guide

**When to Use This Guide**
You should use this guide when:

- You believe that a user is accessing a resource he or she should not have access to.
- A user has been changing discretionary access control lists (DACLs) or system access control lists (SACLs) and you need to discover who is performing the changes.
- A user who should have permission cannot access a resource.

Do not use this guide to find out how to perform a task, such as designing an access control strategy. Most access control problems arise due to the way the access control strategy was designed. While this guide can give you some guidance about how to best assign permissions and perform auditing, you should review the design documents listed in Additional Resources for Troubleshooting Access Control for more definitive advice.

This guide assumes that you have a basic understanding of access control, how it works, and why your organization uses it to control access to resources. You should also have a thorough understanding of how access control is deployed and managed in your organization. This includes an understanding of the mechanism your organization uses to configure and manage access control settings.

**How to Use This Guide**
This guide is divided into four sections. Each section addresses a type of problem, ranging from less complex to more complex.

Verifying Computer Settings for Troubleshooting Access Control

This section provides a list of prerequisites and settings that must be verified before you troubleshoot. Read this section first.

Access Control Quick Fixes

This section provides quick solutions to the most common access control problems. Read this section if you have verified your computer settings, but your problem persists.

Configuring a Computer for Access Control Troubleshooting

This section describes how to configure your computer for troubleshooting. Read this section if you cannot find a solution in the Quick Fixes section.

Troubleshooting Access Control Problems

This section provides step-by-step diagnostic procedures and possible solutions that help you identify and fix access control problems.

Additional Resources for Troubleshooting Access Control

# Verifying Computer Settings for Troubleshooting Access Control

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin troubleshooting, verify that your computer is configured properly and that access control is set up properly.

**Settings to Verify Before You Troubleshoot**
Verify all of the following items:

Determine what type of file system you are using

Make sure you have administrative rights on the computer you are troubleshooting

Install all critical updates and security updates for your system (Windows Server 2003 or Windows 2000 Server)

Update all of your software, including non-Microsoft software

**Determine what type of file system you are using.**
Windows Server 2003 and Windows 2000 Server support the FAT16, FAT32, and NTFS file systems. Because NTFS is required for many advanced, specific security options, it is strongly recommended that you use NTFS. Perform the following procedure to determine what file system is present on your computer.

### To determine your computer's file system

1. Open the Computer Management snap-in. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
2. In the console tree, expand **Storage**, and then click **Disk Management**.
3. In the details pane, note the file system listed under the **File System** column.

**Make sure you have administrative rights on the computer you are troubleshooting.**
You cannot modify access control settings unless you are a member of the Administrators group on the computer that you are administering.

### To verify that you are a member of the Administrators group

1. Open the Computer Management snap-in. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
2. In the console tree, double-click **Local Users and Groups**, and then click **Groups**.
3. In the details pane, double-click **Administrators** and verify that your account name or a group to which your account is a member appears in the **Members** list.

You can also tell if you have the appropriate administrative rights to configure access control by opening **Local Security Policy** in Control Panel under Administrative Tools. If you do not have administrative rights, all of the controls in the **Local Security Policy** user interface (UI) will appear dimmed and a warning will appear telling you that you need to be a computer administrator to change any settings.

**Install all critical updates and security updates for your system (Windows Server 2003 or Windows 2000 Server).**
Some updates might be required for the Windows access control system to function properly.

### To verify that you have all critical updates and security patches

- Click **Start**, click **Help and Support**, click **Windows Update**, and then follow the instructions that appear on your screen.

If you are using Windows Update Services, check that all updates have been approved and applied to your computer.

**Update all of your software, including non-Microsoft software.**
The Windows access control system might not function properly with some programs unless you update the programs with the most recent service pack or software update. Newer versions of many applications, such as antivirus programs, will

automatically configure access control and might resolve any problems you have.

You should also work with any vendors to ensure that all custom applications are up to date.

# Access Control Quick Fixes

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides the most common access control problems and their solutions. You can use the information in this section to resolve problems in the same way you would use a FAQ to find answers to common questions. Read this section before you begin any advanced troubleshooting.

**Common Access Control Problems**
Listed below are common access control problems and quick fixes:

You are unable to access a file or folder as a result of a Deny ACE applied to Everyone

Security policy is not being enforced

**You are unable to access a file or folder as a result of a Deny ACE applied to Everyone**
If a Deny access control entry (ACE) is applied to the Everyone security group, no one will be able to access the resource and only the owner will be able to change the permissions. Anyone, including the owner, who attempts to access the file or folder will receive an "Access is denied" error.

Perform the following procedure to determine whether a Deny ACE is applied to a user or group. You must be an administrator or the owner of the file or folder to perform this task.

**To determine whether a Deny ACE is applied to a user or group**

1. Right-click the file or folder that the user or group is unable to access and then click **Properties**.
2. Click the **Security** tab, and in **Group or user names**, click the user or group you want to examine ACEs for.
3. In **Permission for <user or group>**, examine the permissions and determine whether a Deny ACE is applied to a user or group.

You can recover access to the file or folder, though, by modifying the resource's permissions. Before you can modify the permissions, you must be the owner of the file or folder. Perform the following procedure if you are not the resource's owner. You must be an administrator or have been granted the **Take ownership of files or other objects** user right.

**To take ownership of a file or folder**

1. Open **Windows Explorer**, and then locate the file or folder you want to take ownership of.
2. Right-click the file or folder, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, and then click the **Owner** tab.
4. In the **Change owner to** box, do one of the following:
   - To change the owner to a user or group that is not listed, double-click **Other users and groups**. In the **Enter the object name to select (examples)** box, type the name of the user or group, and then click **OK**.
   - To change the owner to a user or group that is listed, click the new owner.
5. (Optional) To change the owner of all subcontainers and objects within the tree, select the **Replace owner on subcontainers and objects** check box.

| **Note** |
| --- |
| You can transfer ownership in two ways: 1) the current owner can grant the **Take ownership** permission to others, allowing those users to take ownership at any time. A user granted the **Take ownership** permission can take ownership of the object or assign ownership to any group that the user is a member of. 2) A user who has the **Restore files and directories** privilege can double-click **Other users and groups** and choose any user or group to assign ownership to. |

| **Note** |
| --- |

| |
|---|
| An administrator can take ownership of any file on the computer. |

Once you have taken ownership of the file or folder, perform the following procedure to remove the Deny ACE from the Everyone group's ACL.

**To remove the Deny ACE applied to the Everyone group**

1. Open **Windows Explorer**, and then locate the file or folder you want to modify permissions for.
2. Right-click the file or folder, click **Properties**, and then click the **Security** tab.
3. In the **Group or user names** box, select **Everyone**.
4. In the **Permissions for Everyone** box, select **Full Control**.

**Security policy is not being enforced**

After the local security policy or a policy administered through Group Policy has been modified, you might need to refresh the policy on a local computer if it is not enforcing the policy.

Perform the following procedure to refresh the security policy on a Windows 2000 Server–based computer.

**To force a Windows 2000 Server–based computer to refresh its security policy settings and Group Policy settings**

1. Open a command prompt.
2. Type **secedit /refreshpolicy**.

The **/refresh** option ignores all processing optimizations and reapplies all settings.

The Gpupdate command-line tool refreshes local Group Policy settings and Group Policy settings that are stored in Active Directory, including security settings. This command supersedes the now obsolete **/refreshpolicy** option for the **secedit** command for Windows Server 2003.

Perform the following procedure to refresh the security policy on a Windows Server 2003–based computer.

**To force a Windows Server 2003–based computer to refresh its security settings and Group Policy settings**

1. Open a command prompt.
2. Type **gpupdate /force**.

The **/force** option ignores all processing optimizations and reapplies all settings.

# Configuring a Computer for Access Control Troubleshooting

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can use advanced troubleshooting techniques to identify and fix access control problems, you need to configure your computer for troubleshooting. In addition, you need a basic understanding of troubleshooting concepts, procedures, and tools.

To configure your computer for troubleshooting access control, you can install Xcacls.exe.

**Install Xcacls.exe**

The Xcacls command-line tool sets all file-system security options that are accessible in Windows Explorer. Xcacls does this by displaying and modifying the access control lists (ACLs) of files.

The tool is part of the Windows 2000 Resource Kit Tools. To download the Xcacls tool, see Free Tool Downloads on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=18748).

# Troubleshooting Access Control Problems

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The problems that typically occur with the Windows access control system are problems with user access to files and folders.

Share permissions are set on folders only and apply to folders that have been shared on a file server or computer. NTFS permissions apply to discs that are formatted with the NTFS file system. NTFS permissions can be applied on local files and folders and shared folders.

**In this section**


- Anonymous user cannot access a shared folder
- Unknown user is attempting to access a file or folder he or she should not have access to
- Unknown user is changing ACLs on files or folders
- User cannot access a file or folder but should have permission
- Users cannot write to a shared folder after migration to Windows Server 2003

# Anonymous user cannot access a shared folder

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem typically occurs when someone logging on with a Guest account attempts to access a shared folder. Users or services that attempt to access an object anonymously are not granted access if the access control list (ACL) on the object includes the Everyone group.

**Cause**

While Windows NT 4.0 and Windows 2000 allowed anonymous access to the server, the inclusion of the Everyone security group in the anonymous user access token is disabled by default in Windows Server 2003 and Windows XP. In Windows XP and in Windows Server 2003, the Everyone group does not contain the security identifier (SID) "Anonymous." Anonymous access is only granted for objects whose ACL explicitly contains the anonymous SID.

**Solution**

Add the Anonymous SID to the Everyone access token.

| ⚠Caution |
| --- |
| For security reasons, it is strongly recommended that you do not allow anonymous access to your server. |

However, if anonymous access is needed, perform one of the following procedures.

Enable anonymous access on a local computer

Enable anonymous access on a domain controller

**Enable anonymous access on a local workstation or server computer**

Perform the following procedure to enable anonymous access on a local workstation or server computer.

**To enable anonymous access on a local workstation or server computer**

1. Open Local Security Settings. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.
2. In the console tree, double-click **Local Policies**, and then click **Security Options**.
3. In the details pane, right-click **Network access: Let Everyone permissions apply to anonymous users**, and then click Properties.
4. On the **Local Security Settings** tab, click **Enabled**, and then click **OK**.

**Enable anonymous access on a domain controller**

Perform the following procedure to enable anonymous access on a domain controller.

**To enable anonymous access on a domain controller**

1. Open Local Security Settings. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Domain Security Policy**.
2. In the console tree, double-click **Security Settings**, double-click **Local Policies**, and then click **Security Options**.
3. In the details pane, right-click **Network access: Let Everyone permissions apply to anonymous users**, and then click **Properties**.
4. Select the **Define this policy setting** check box, click **Enabled**, and then click **OK**.

For more information, see Article 278259, Everyone Group Does Not Include Anonymous Security Identifier, on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47167).

# Unknown user is attempting to access a file or folder he or she should not have access to

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem typically occurs when someone is attempting to access a file or folder he or she either has not been given permission to access or has been explicitly denied access. This problem can be caused unintentionally by a user mistakenly attempting to access a resource or it can be caused intentionally by a user trying to circumvent the access control policy on the file or folder.

**Cause**
When an audit policy is not in place, it is difficult to determine exactly who is attempting to access unauthorized resources.

**Solutions**
Because a Windows-based computer can be managed with Group Policy or be a standalone computer, there are two different solutions for this problem. Choose one of the following options and follow the corresponding procedures.

- Computer is managed with Group Policy
- Computer is not managed with Group Policy

Please check with your system administrator to determine whether the computer is managed with Group Policy.

**Computer is managed with Group Policy**
If the computer you would like to audit events for is managed by Group Policy, perform the following procedures to discover which user is attempting to access the file or folder.

1. Enable auditing of the object access event category
2. Apply auditing policy settings for an object using Group Policy
3. Analyze the Security events in Event Viewer
4. Adjust permissions or enforce a new security policy

| Note |
| --- |
| You can set up file and folder auditing only on NTFS drives. |

| Note |
| --- |
| To perform this set of procedures, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been granted the Manage auditing and security log right in Group Policy. As a security best practice, consider using Run as to perform this procedure. |

**Step One: Enable auditing of the object access event category**
Choose among the following object access auditing configuration options, depending on your environment:

- For your local computer
- For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed
- For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed

- For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain

**For your local computer**
Perform the following procedure:

**For your local computer**

1. Open **Local Security Policy**. Click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.
2. In the console tree, expand **Security Settings**, expand **Local Policies**, and click **Audit Policy**.
3. In the details pane, double-click **Audit object access**.
4. Select the **Failure** check box.

**For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed**
Perform the following procedure:

**For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed**

1. Open **Domain Controller Security Policy**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Domain Controller Security Policy**.
2. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
3. In the details pane, double-click **Audit object access**.
4. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
5. Select the **Failure** check box.

**For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed**
Perform the following procedure:

**For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed**

1. Open **Active Directory Users and Computers**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the Group Policy object (GPO) that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.
5. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
6. In the details pane, double-click **Audit object access**.
7. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
8. Select the **Failure** check box.

**For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain**
Perform the following procedure:

**For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a**

**domain**

1. Open Microsoft Management Console (MMC). Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. On the **Select Group Policy Object** page in the **Group Policy Wizard**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit—or create a new one—click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, and then click **Audit Policy**.
8. In the details pane, double-click **Audit object access**.
9. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
10. Select the **Failure** check box.

**Step Two: Apply auditing policy settings for an object by using Group Policy**

| **Important** |
| --- |
| Before setting up auditing for files and folders, you must enable object access auditing by defining auditing policy settings for the object access event category. If you do not enable object access auditing, you will receive an error message when you set up auditing for files and folders, and no files will be audited. |

| **Note** |
| --- |
| To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been granted the **Manage auditing and security log** right in Group Policy. As a security best practice, consider using Run as to perform this procedure. |

**Apply auditing policy settings for an object by using Group Policy**

1. Open Microsoft Management Console (MMC). Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. On the **Select Group Policy Object** page in the **Group Policy Wizard**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit—or create a new one, click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, and then click **File System**.
8. If you want to add a file or folder to this GPO to audit, right-click **File System**, and then click **Add File**. Browse to the file that you want, or make a new folder, and then click **OK**.
9. If you want to apply or modify auditing settings on a file or folder that has already been added to this GPO, in the details pane, right-click the file or folder, click **Properties**, and then click **Edit Security**.
10. Click **Advanced**, and then click the **Auditing** tab.
11. Do one of the following:

- To set up auditing for a new user or group, click **Add**. In **Name**, type the name of the user or group that you want, and then click **OK**.
- To view or change auditing for an existing group or user, click the name that you want, and then click **Edit**.
- To remove auditing for an existing group or user, click the name that you want, click **Remove**, click **OK**, and then skip the rest of this procedure.

12. Select the appropriate entry in the **Apply onto** list.
13. In the **Access** box, select the **Failed** check box.

    If you want to prevent files and subfolders in the tree from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.
14. Click **OK**.

For best system performance, minimize the number of entries in the SACL for an object. One entry in a SACL that contains 1000 users does not degrade system performance as much as 1000 separate entries.

Because the Security log is limited in size, select the files and folders to be audited carefully. Also, consider the amount of disk space that you want to devote to the Security log. The maximum size for the Security log is defined in **Event Viewer**.

If you are unsure about what user might be attempting to access the restricted file or folder, audit object access for groups. Make sure to change the location to the local computer if you want to include a local group, such as Users.

**Step Three: Analyze the Security events in Event Viewer**

After applying the audit policy to the desired file or folder, view the Security log in Event Viewer to determine what user has been attempting to access the restricted file or folder.

Perform the following procedure to view the Security log.

**To view the Security log**

1. Open **Event Viewer**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Event Viewer**.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. If you want to see more details about a specific event, in the details pane, double-click the event.

Once you have accessed Event Viewer, you should take advantage of the filter to eliminate irrelevant data in the Security log. Perform the following procedure to filter for Success and Failure events for the object access event category.

**To filter the Security log view**

1. Open **Event Viewer**.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. In the **File** menu, click **View**, and then click **Filter**.
4. In **Event types**, clear the **Information**, **Warning**, and **Error** check boxes and ensure that the **Success** and **Failure** check boxes are selected.
5. In the **Event Source** list, select **Security**.
6. In the **Category** list, select **Object Access**.
7. Click **OK**.

| **Note** |
| --- |
| You can use the other fields in the Filter page to further filter the Security events by event ID, user, and computer. You can also filter events by using a time-period range. Filtering by a range of time can be useful if you believe that someone tried to access a restricted file or folder during a specific time on a particular day. |

Once you have filtered the events, examine the Success and Failure audits to see if they apply to restricted files or folders that someone attempted to access. The user name of the person attempting to access the restricted file or folder will be listed within the security event, along with the action that the user was attempting to perform or successfully performed.

| **Important** |
| --- |
| The view filter that you set will remain on the event log unless you manually change the filter. Filtering the view of the event log can help you narrow the scope of the events presented and make it much easier to find potential security problems. |

**Step Four: Adjust permissions or enforce a new security policy**
After you have discovered which user is attempting to access a file or folder, you must then decide on a course of action to complete the solution. There are two possible solutions, depending upon the original intent of the person who attempted to access a file or folder.

If the user is trusted and attempted to access the file or folder believing that he or she was performing an action that complies with your security policy, you should either further refine your security policy or educate the user about the proper way to follow the policy.

If the user is not trusted and attempted to change the file or folder to intentionally contradict the established security policy, then you should lock down the user's permissions. Make sure to reevaluate your security policy and adhere to best practices. For guidance about applying permissions, see Users cannot write to a shared folder after migration to Windows Server 2003.

**Computer is not managed with Group Policy**
If the computer you would like to audit events for is not managed by Group Policy, perform the following procedures to discover which user is attempting to access the file or folder.

1. Enable auditing of the object access event category
2. Apply auditing policy to a specified object
3. Analyze the Security events in Event Viewer
4. Adjust permissions or enforce a new security policy

You must first decide what objects you would like to audit and whether you want to audit a specific computer or use Group Policy to manage auditing within an organization unit (OU) or domain.

| **Note** |
| --- |
| You can set up file and folder auditing only on NTFS drives. |

**Step One: Enable auditing of the object access event category**
Choose between the following object access auditing configuration options, depending on your environment:

- For your local computer
- For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed
- For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed
- For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain

**For your local computer**
Perform the following procedure:

**For your local computer**

1. Open **Local Security Policy**. Click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**,

and then double-click **Local Security Policy**.

2. In the console tree, expand **Security Settings**, expand **Local Policies**, and click **Audit Policy**.
3. In the details pane, double-click **Audit object access**.
4. Select the **Failure** check box.

**For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed**
Perform the following procedure:

### For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed

1. Open **Domain Controller Security Policy**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Domain Controller Security Policy**.
2. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
3. In the details pane, double-click **Audit object access**.
4. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
5. Select the **Failure** check box.

**For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed**
Perform the following procedure:

### For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed

1. Open **Active Directory Users and Computers**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the Group Policy object (GPO) that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.
5. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
6. In the details pane, double-click **Audit object access**.
7. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
8. Select the **Failure** check box.

**For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain**
Perform the following procedure:

### For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain

1. Open **Microsoft Management Console (MMC)**. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. In the File menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. On the **Select Group Policy Object** page in the **Group Policy Wizard**, click **Browse**.

5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit—or create a new one—click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, and then click **Audit Policy**.
8. In the details pane, double-click **Audit object access**.
9. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
10. Select the **Failure** check box.

**Step Two: Apply auditing policy to a specified object**

To perform the following procedure, you must be logged on as a member of the Administrators group or you must have been granted the **Manage auditing and security log** right in Group Policy.

| **Important** |
| --- |
| In addition to setting up auditing for files and folders, you must enable object access auditing by defining auditing policy settings for the object access event category. If you do not enable object access auditing, you will receive an error message when you set up auditing for files and folders, and no files or folders will be audited. |

**To apply or modify auditing policy settings for a local file or folder**

1. Open **Windows Explorer**. To open **Windows Explorer**, right click **Start** and click **Explore**.
2. Right-click the file or folder that you want to audit, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, and then click the **Auditing** tab.
4. Do one of the following:
   - To set up auditing for a new user or group, click **Add**. In **Enter the object name to select**, type the name of the user or group that you want, and then click **OK**.
   - To view or change auditing for an existing group or user, click its name, and then click **Edit**.
5. In the **Apply onto** box, click the location(s) where you want auditing to take place. This is only applicable for a folder.
6. In the **Access** box, indicate what actions you want to audit by selecting the appropriate check boxes, and then click **OK**:
   - To audit successful events, select the **Successful** check box.
   - To stop auditing successful events, clear the **Successful** check box.
   - To audit unsuccessful events, select the **Failed** check box.
   - To stop auditing unsuccessful events, clear the **Failed** check box.
   - To stop auditing all events, click **Clear All**. This option removes the auditing entry.
7. If you want to prevent subsequent files and subfolders of the original object from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.

If you see the one of the following, auditing has been inherited from the parent folder:

- In the **Auditing Entry for File or Folder** dialog box, in the **Access** box, the check boxes are unavailable.
- In the **Advanced Security Settings for File or Folder** dialog box, the **Remove** button is unavailable.

For best system performance, minimize the number of entries in the SACL for an object. One entry in a SACL that contains 1000 users does not degrade system performance as much as 1000 separate entries.

Because the Security log is limited in size, select the files and folders to be audited carefully. Also, consider the amount of disk space that you want to devote to the Security log. The maximum size for the Security log is defined in **Event Viewer**.

If you are unsure about what user might be attempting to access the restricted file or folder, audit object access for groups. Make sure to change the location to the local computer if you want to include a local group, such as Users.

**Step Three: Analyze the Security events in Event Viewer**
After applying the audit policy to the desired file or folder, view the Security log in Event Viewer to determine what user has been attempting to access the restricted file or folder.

Perform the following procedure to view the Security log.

**To view the Security log**

1. Open **Event Viewer**. To open **Event Viewer**, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Event Viewer**.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. If you want to see more details about a specific event, in the details pane, double-click the event.

Once you have accessed Event Viewer, you should take advantage of the filter to eliminate irrelevant data in the Security log. Perform the following procedure to filter for Success and Failure events for the object access event category.

**To filter the Security log view**

1. Open **Event Viewer**.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. In the **File** menu, click **View**, and then click **Filter**.
4. In **Event types**, clear the **Information**, **Warning**, and **Error** check boxes and ensure that the **Success** and **Failure** check boxes are selected.
5. In the **Event Source** list, select **Security**.
6. In the **Category** list, select **Object Access**.
7. Click **OK**.

| **Note** |
| --- |
| You can use the other fields in the Filter page to further filter the Security events by Event ID, user, and computer. You can also filter events by using a time-period range. Filtering by a range of time can be useful if you believe that someone tried to access a restricted file or folder during a specific time on a particular day. |

Once you have filtered the events, examine the Success and Failure audits to see if they apply to restricted files or folders that someone attempted to access. The user name of the person attempting to access the restricted file or folder will be listed within the security event, along with the action that the user was attempting to perform or successfully performed.

| **Important** |
| --- |
| The view filter that you set will remain on the event log unless you manually change the filter. Filtering the view of the event log can help you narrow the scope of the events presented and make it much easier to find potential security problems. |

**Step Four: Adjust permissions or enforce a new security policy**
After you have discovered which user is attempting to access a file or folder, you must then decide on a course of action to complete the solution. There are two possible solutions, depending upon the original intent of the person who attempted to access a file or folder.

If the user is trusted and attempted to access the file or folder believing that he or she was performing an action that complies

with your security policy, you should either further refine your security policy or educate the user about the proper way to follow the policy.

If the user is not trusted and attempted to change the file or folder to intentionally contradict the established security policy, you should lock down the user's permissions. Make sure to reevaluate your security policy and adhere to best practices. For guidance about applying permissions, see Users cannot write to a shared folder after migration to Windows Server 2003.

# Unknown user is changing ACLs on files or folders

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem typically occurs in any of the following circumstances:

- A user has excessive permissions.
- A user has the appropriate level of permissions and is unsure of how the Windows access control system works and mistakenly changes ACLs.
- A user has the appropriate level of permissions and intentionally changes ACLs to circumvent or alter an existing access control policy.

**Cause**
A user has been granted the NTFS Full Control permission on a file or folder or is an administrator.

**Solutions**
Because a Windows-based computer can be managed with Group Policy or be a standalone computer, there are two different solutions for this problem. Choose one of the following options and follow the corresponding procedures.

- Computer is managed with Group Policy
- Computer is not managed with Group Policy

Please check with your system administrator to determine whether the computer is managed with Group Policy.

**Computer is managed with Group Policy**
If the computer you would like to audit events for is managed by Group Policy, perform the following procedures to discover which user is attempting to access the file or folder.

1. Enable auditing of the object access event category
2. Apply auditing policy settings for an object by using Group Policy
3. Analyze the Security events in Event Viewer
4. Adjust permissions or enforce a new security policy

| **Note** |
| --- |
| You can set up file and folder auditing only on NTFS drives. |

| **Note** |
| --- |
| To perform this set of procedures, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been granted the Manage auditing and security log right in Group Policy. As a security best practice, consider using Run as to perform this procedure. |

**Step One: Enable auditing of the object access event category**
Choose among the following object access auditing configuration options, depending on your environment:

- For your local computer
- For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed

- For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed
- For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain

**For your local computer**
Perform the following procedure:

**For your local computer**

1. Open **Local Security Policy**. Click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.
2. In the console tree, expand **Security Settings**, expand **Local Policies**, and click **Audit Policy**.
3. In the details pane, double-click **Audit object access**.
4. Select the **Failure** check box.

**For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed**
Perform the following procedure:

**For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed**

1. Open **Domain Controller Security Policy**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Domain Controller Security Policy**.
2. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
3. In the details pane, double-click **Audit object access**.
4. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
5. Select the **Failure** check box.

**For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed**
Perform the following procedure:

**For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed**

1. Open **Active Directory Users and Computers**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the Group Policy object (GPO) that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.
5. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
6. In the details pane, double-click **Audit object access**.
7. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
8. Select the **Failure** check box.

**For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain**
Perform the following procedure:

### For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain

1. Open Microsoft Management Console (MMC). Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. In the File menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. On the **Select Group Policy Object** page in the **Group Policy Wizard**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit—or create a new one—click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, and then click **Audit Policy**.
8. In the details pane, double-click **Audit object access**.
9. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
10. Select the **Failure** check box.

**Step Two: Apply auditing policy settings for an object by using Group Policy**

| **Important** |
| --- |
| Before setting up auditing for files and folders, you must enable object access auditing by defining auditing policy settings for the object access event category. If you do not enable object access auditing, you will receive an error message when you set up auditing for files and folders, and no files will be audited. |

| **Note** |
| --- |
| To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been granted the **Manage auditing and security log** right in Group Policy. As a security best practice, consider using Run as to perform this procedure. |

### Apply auditing policy settings for an object by using Group Policy

1. Open Microsoft Management Console (MMC). Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. On the **Select Group Policy Object** page in the **Group Policy Wizard**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit—or create a new one, click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, and then click **File System**.
8. If you want to add a file or folder to this GPO to audit, right-click **File System**, and then click **Add File**. Browse to the file that you want, or make a new folder, and then click **OK**.
9. If you want to apply or modify auditing settings on a file or folder that has already been added to this GPO, in the details pane, right-click the file or folder, click **Properties**, and then click **Edit Security**.

10. Click **Advanced**, and then click the **Auditing** tab.
11. Do one of the following:
    - To set up auditing for a new user or group, click **Add**. In **Name**, type the name of the user or group that you want, and then click **OK**.
    - To view or change auditing for an existing group or user, click the name that you want, and then click **Edit**.
    - To remove auditing for an existing group or user, click the name that you want, click **Remove**, click **OK**, and then skip the rest of this procedure.
12. Select the appropriate entry in the **Apply onto** list.
13. In the **Access** box, select the **Failed** check box.

    If you want to prevent files and subfolders in the tree from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.
14. Click **OK**.

For best system performance, minimize the number of entries in the SACL for an object. One entry in a SACL that contains 1000 users does not degrade system performance as much as 1000 separate entries.

Because the Security log is limited in size, select the files and folders to be audited carefully. Also, consider the amount of disk space that you want to devote to the Security log. The maximum size for the Security log is defined in **Event Viewer**.

If you are unsure about what user might be attempting to access the restricted file or folder, audit object access for groups. Make sure to change the location to the local computer if you want to include a local group, such as Users.

**Step Three: Analyze the Security events in Event Viewer**

After applying the audit policy to the desired file or folder, view the Security log in Event Viewer to determine what user has been attempting to access the restricted file or folder.

Perform the following procedure to view the Security log.

**To view the Security log**

1. Open **Event Viewer**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Event Viewer**.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. If you want to see more details about a specific event, in the details pane, double-click the event.

Once you have accessed Event Viewer, you should take advantage of the filter to eliminate irrelevant data in the Security log. Perform the following procedure to filter for Success and Failure events for the object access event category.

**To filter the Security log view**

1. Open **Event Viewer**.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. In the **File** menu, click **View**, and then click **Filter**.
4. In **Event types**, clear the **Information**, **Warning**, and **Error** check boxes and ensure that the **Success** and **Failure** check boxes are selected.
5. In the **Event Source** list, select **Security**.
6. In the **Category** list, select **Object Access**.
7. Click **OK**.

**Note**

You can use the other fields in the Filter page to further filter the Security events by event ID, user, and computer. You can also filter events by using a time-period range. Filtering by a range of time can be useful if you believe that someone tried to access a restricted file or folder during a specific time on a particular day.

Once you have filtered the events, examine the Success and Failure audits to see if they apply to restricted files or folders that someone attempted to access. The user name of the person attempting to access the restricted file or folder will be listed within the security event, along with the action that the user was attempting to perform or successfully performed.

**Important**

The view filter that you set will remain on the event log unless you manually change the filter. Filtering the view of the event log can help you narrow the scope of the events presented and make it much easier to find potential security problems.

**Step Four: Adjust permissions or enforce a new security policy**

After you have discovered which user is attempting to access a file or folder, you must then decide on a course of action to complete the solution. There are two possible solutions, depending upon the original intent of the person who attempted to access a file or folder.

If the user is trusted and attempted to access the file or folder believing that he or she was performing an action that complies with your security policy, you should either further refine your security policy or educate the user about the proper way to follow the policy.

If the user is not trusted and attempted to change the file or folder to intentionally contradict the established security policy, then you should lock down the user's permissions. Make sure to reevaluate your security policy and adhere to best practices. For guidance about applying permissions, see Users cannot write to a shared folder after migration to Windows Server 2003.

**Computer is not managed with Group Policy**

If the computer you would like to audit events for is not managed by Group Policy, perform the following procedures to discover which user is attempting to access the file or folder.

1. Enable auditing of the object access event category
2. Apply auditing policy to a specified object
3. Analyze the Security events in Event Viewer
4. Adjust permissions of enforce a new security policy

You must first decide what objects you would like to audit and whether you want to audit a specific computer or use Group Policy to manage auditing within an organization unit (OU) or domain.

**Note**

You can set up file and folder auditing only on NTFS drives.

**Step One: Enable auditing of the object access event category**

Choose between the following object access auditing configuration options, depending on your environment:

- For your local computer
- For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed
- For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed
- For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain

**For your local computer**
Perform the following procedure:

**For your local computer**

1. Open **Local Security Policy**. Click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.
2. In the console tree, expand **Security Settings**, expand **Local Policies**, and click **Audit Policy**.
3. In the details pane, double-click **Audit object access**.
4. Select the **Failure** check box.

**For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed**
Perform the following procedure:

**For domain controllers only, when you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed**

1. Open **Domain Controller Security Policy**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Domain Controller Security Policy**.
2. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
3. In the details pane, double-click **Audit object access**.
4. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
5. Select the **Failure** check box.

**For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed**
Perform the following procedure:

**For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed**

1. Open **Active Directory Users and Computers**. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the Group Policy object (GPO) that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.
5. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Audit Policy**.
6. In the details pane, double-click **Audit object access**.
7. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
8. Select the **Failure** check box.

**For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain**
Perform the following procedure:

**For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain**

1. Open Microsoft Management Console (MMC). Click **Start**, click **Run**, type **mmc**, and then click **OK**.

2. In the File menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. On the **Select Group Policy Object** page in the **Group Policy Wizard**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit—or create a new one—click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, and then click **Audit Policy**.
8. In the details pane, double-click **Audit object access**.
9. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.
10. Select the **Failure** check box.

**Step Two: Apply auditing policy to a specified object**

To perform the following procedure, you must be logged on as a member of the Administrators group or you must have been granted the **Manage auditing and security log** right in Group Policy.

| **Important** |
| --- |
| In addition to setting up auditing for files and folders, you must enable object access auditing by defining auditing policy settings for the object access event category. If you do not enable object access auditing, you will receive an error message when you set up auditing for files and folders, and no files or folders will be audited. |

**To apply or modify auditing policy settings for a local file or folder**

1. Open **Windows Explorer**. To open **Windows Explorer**, right-click **Start** and click **Explore**.
2. Right-click the file or folder that you want to audit, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, and then click the **Auditing** tab.
4. Do one of the following:
   - To set up auditing for a new user or group, click **Add**. In **Enter the object name to select**, type the name of the user or group that you want, and then click **OK**.
   - To view or change auditing for an existing group or user, click its name, and then click **Edit**.
5. In the **Apply onto** box, click the location(s) where you want auditing to take place. This is only applicable for a folder.
6. In the **Access** box, indicate what actions you want to audit by selecting the appropriate check boxes, and then click **OK**:
   - To audit successful events, select the **Successful** check box.
   - To stop auditing successful events, clear the **Successful** check box.
   - To audit unsuccessful events, select the **Failed** check box.
   - To stop auditing unsuccessful events, clear the **Failed** check box.
   - To stop auditing all events, click **Clear All**. This option removes the auditing entry.
7. If you want to prevent subsequent files and subfolders of the original object from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.

If you see the one of the following, auditing has been inherited from the parent folder:

- In the **Auditing Entry for File or Folder** dialog box, in the **Access** box, the check boxes are unavailable.
- In the **Advanced Security Settings for File or Folder** dialog box, the **Remove** button is unavailable.

For best system performance, minimize the number of entries in the SACL for an object. One entry in a SACL that contains 1000 users does not degrade system performance as much as 1000 separate entries.

Because the Security log is limited in size, select the files and folders to be audited carefully. Also, consider the amount of disk space that you want to devote to the Security log. The maximum size for the Security log is defined in **Event Viewer**.

If you are unsure about what user might be attempting to access the restricted file or folder, audit object access for groups. Make sure to change the location to the local computer if you want to include a local group, such as Users.

**Step Three: Analyze the Security events in Event Viewer**
After applying the audit policy to the desired file or folder, view the Security log in Event Viewer to determine what user has been attempting to access the restricted file or folder.

Perform the following procedure to view the Security log.

**To view the Security log**

1. Open **Event Viewer**. To open **Event Viewer**, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Event Viewer**.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. If you want to see more details about a specific event, in the details pane, double-click the event.

Once you have accessed Event Viewer, you should take advantage of the filter to eliminate irrelevant data in the Security log. Perform the following procedure to filter for Success and Failure events for the object access event category.

**To filter the Security log view**

1. Open **Event Viewer**.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. In the **File** menu, click **View**, and then click **Filter**.
4. In **Event types**, clear the **Information**, **Warning**, and **Error** check boxes and ensure that the **Success** and **Failure** check boxes are selected.
5. In the **Event Source** list, select **Security**.
6. In the **Category** list, select **Object Access**.
7. Click **OK**.

| **Note** |
| --- |
| You can use the other fields in the Filter page to further filter the Security events by Event ID, user, and computer. You can also filter events by using a time-period range. Filtering by a range of time can be useful if you believe that someone tried to access a restricted file or folder during a specific time on a particular day. |

Once you have filtered the events, examine the Success and Failure audits to see if they apply to restricted files or folders that someone attempted to access. The user name of the person attempting to access the restricted file or folder will be listed within the security event, along with the action that the user was attempting to perform or successfully performed.

| **Important** |
| --- |
| The view filter that you set will remain on the event log unless you manually change the filter. Filtering the view of the event log can help you narrow the scope of the events presented and make it much easier to find potential security problems. |

**Step Four: Adjust permissions or enforce a new security policy**

After you have discovered which user is attempting to access a file or folder, you must then decide on a course of action to complete the solution. There are two possible solutions, depending upon the original intent of the person who attempted to access a file or folder.

If the user is trusted and attempted to access the file or folder believing that he or she was performing an action that complies with your security policy, you should either further refine your security policy or educate the user about the proper way to follow the policy.

If the user is not trusted and attempted to change the file or folder to intentionally contradict the established security policy, you should lock down the user's permissions. Make sure to reevaluate your security policy and adhere to best practices. For guidance about applying permissions, see Users cannot write to a shared folder after migration to Windows Server 2003.

# User cannot access a file or folder but should have permission

Updated: May 8, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem occurs when a user attempts to access a file or folder and receives an "Access Denied" error.

**Cause**
This problem results when an administrator, the owner of the file or folder, or a user with the Change Permissions permission for a file or Full Control permission for a folder, changes the access control list (ACL) on the file or folder. Some of the changes that could deny a user access to a file or folder they should be able to access include:

- The user or group was unintentionally removed from **Group or User names** on the file's or folder's **Security** tab.
- An explicit **Deny** has mistakenly been applied to the user or group.
- A conflict between share permissions and NTFS permissions was created when someone changes permissions without examining both types of permissions.

**Solutions**
There are two separate scenarios where a user may be denied access to a file or folder:

- A user is denied access to a local file or folder
- A user is denied access to a shared folder

**A user is denied access to a local file or folder**
There are three possible solutions for this problem:

Examine a user's permissions on a file or folder using Xcacls

Examine a user's or group's permissions on a file or folder

View a user's or group's effective permissions on files and folders

**Solution One: Examine a user's permissions on a file or folder using Xcacls**
The Xcacls command-line tool sets all file system security options that are accessible in Windows Explorer. Xcacls does this by displaying and modifying the ACLs of files. Perform the following procedure to investigate ACLs by using the Xcacls tool.

| **Important** |
| --- |
| Xcacls does not ship with Windows Server 2003. For information about how to install the Xcacls tool, see Configuring a Computer for Access Control Troubleshooting. |

### To examine a user's permissions on a file or folder using Xcacls

1. Open a command prompt. Click **Start**, click **Run**, and type **cmd**.
2. Type:

   **Xcacls** *FileOrFolderPath*

3. Examine the user's permissions and compare with the desired permissions.

The following is an example of the preceding Xcacls command:

**Xcacls C:\Tools**

**Solution Two: Examine a user's or group's permissions on a file or folder**

| **Note** |
| --- |

The use of access control lists (ACLs) to control users' or groups' permissions on a file or folder is supported on computers that use the NTFS file system but not on computers that use the FAT file system or that use Windows XP Simple File Sharing. If ACLs are not supported, the file or folder properties do not have Security properties available. Windows XP Home Edition is an example of an operating system that does not support ACLs.

**To examine a user's or group's permissions on a file or folder**

1. Right-click the file or folder you wish to administer permissions for, click **Properties**, and then click the **Security** tab.

   **Note**

   If you are logged in with an account that has not been granted rights to administer permissions on the file or folder, the **Security** tab is not available.

2. In **Group or User names**, select the user or group whose permissions you want to view.
3. In **Permissions for <group name or user name>**, examine the user's permissions and compare with the desired permissions.

**Solution Three: View a user's or group's effective permissions on files and folders**

You can also use the Effective Permissions tool to view effective permissions of a user or group. The Effective Permissions tool does not display share permissions; it displays only cumulative NTFS permissions for a particular user or group.

**To view effective permissions on files and folders**

1. Open **Windows Explorer**, and then locate the file or folder for which you would like to view effective permissions.
2. Right-click the file or folder, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, click the **Effective Permissions** tab, and then click **Select**.
4. In **Enter the object name to select (examples)**, enter the name of a user or group, and then click **OK**. The selected check boxes indicate the effective permissions of the user or group for that file or folder.

**A user is denied access to a shared folder**

Both share and NTFS permissions can be applied to a folder that is shared. However, share permissions and NTFS permissions behave differently. As a result, there are different processes for investigating permissions problems if share only or if a mixture of share and NTFS permissions have been applied. First you should determine what share permissions have been applied. Then determine which, if any, NTFS permissions have been applied.

**Important**

When share and NTFS permissions conflict, the most restrictive setting takes precedence.

Perform the following procedure to examine share permissions on a specific folder.

**To examine user and group share permissions on a shared folder**

1. Right-click the shared folder you want to examine, and then click **Properties**.
2. Click the **Sharing** tab, and then click **Permissions**.
3. Examine the permissions and change any settings that contradict the desired permissions.

After you examine the share permissions, use the Xcacls tool to determine what NTFS permissions apply to a specific user or group. Perform the following procedure to investigate ACLs by using the Xcacls tool.

**Important**

Xcacls does not ship with Windows Server 2003. For information about how to install the Xcacls tool, see Configuring a Computer for Access Control Troubleshooting.

**To examine user and group NTFS permissions on a shared folder**

1. Open a command prompt.
2. Type the following command:

   **Xcacls** *FileorFolderPath*

3. Examine the permissions and note any settings that contradict the desired permissions.
4. Open Windows Explorer, right-click the shared folder, and then click **Properties**.
5. Based on the outcome of the Xcacls data, adjust either the share permissions on the **Sharing** tab or the NTFS permissions on the **Security** tab. You can also use the Xcacls tool to adjust NTFS permissions using the command line.

Use the following guidelines when examining and adjusting share and NTFS permissions on a shared folder:

- Inherited Deny permissions do not prevent access to an object if the object has an explicit Allow permission entry.
- Explicit permissions take precedence over inherited permissions, even inherited Deny permissions.

NTFS permissions affect access to files and folders both locally and remotely and apply regardless of protocol. Share permissions, by contrast, apply only to shared folders (on a network). Share permissions do not restrict access to any local user, or to any terminal server user, on the computer on which you have set share permissions. Thus, share permissions do not provide privacy between users on a computer used by several users, nor on a terminal server accessed by several users.

If you want to manage folder access by using NTFS permissions exclusively, set share permissions to Full Control for Everyone. This frees you from having to think about share permissions, but because NTFS permissions are more complex than share permissions, using NTFS permissions correctly requires deeper understanding on your part. For more information about NTFS permissions, search for the term "NTFS permissions" on TechNet on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=38772).

| **Caution** |
| --- |
| Granting a user Full Control NTFS permission on a folder enables that user to take ownership of the folder unless the user is restricted in some other way. Be cautious in granting Full Control. |

# Users cannot write to a shared folder after migration to Windows Server 2003

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem typically occurs after a migration of computers from Windows NT 4.0 or Windows 2000 to Windows Server 2003. Certain default settings have changed in Windows Server 2003, and these changes might affect the expected behavior of your server.

**Cause**
The Everyone group has the Full Control share permission by default in Windows NT and Windows 2000. However, the Everyone group has the Read share permission by default in Windows Server 2003 and Windows XP. If the default permissions are applied to a Windows NT–based or Windows 2000–based computer before migration to Windows Server 2003, user accounts may no longer be able to write to the shared folders.

**Solutions**
To solve this problem, consider either creating new groups with the required permissions or adding the Full Control share permission for the Everyone group.

**Solution One: Create new groups and assign share permissions**
Create the groups that you need and assign the appropriate permissions on shared folders to these groups. As a best practice, assign permissions to groups rather than to users. Because it is inefficient to maintain user accounts directly, assigning permissions on a user basis should be the exception. The Everyone group is automatically assigned the Read permission, which is the most restrictive, to assist you in the administration of resources.

The following are also best practices for managing permissions and user rights:

- Assign permissions to an object as high on the tree as possible and then apply inheritance to propagate the security settings throughout the tree.

  You can quickly and effectively apply access control settings to all children or a subtree of a parent object. By doing this, you gain the greatest breadth of effect with the least effort. The permission settings you establish should be adequate for the majority of users, groups, and computers.

- Deny permissions should be used for certain special cases

  - Use Deny permissions to exclude a subset of a group which has Allowed permissions.

  - Use Deny to exclude one special permission when you have already granted Full Control to a user or group.

- If possible, avoid changing the default permission entries on file system objects, particularly on system folders and root folders. Changing default permissions can cause unexpected access problems or reduce security.

**Solution Two: Change share permissions for the Everyone group**
You can assign the following types of access permissions to shared folders or drives:

**Read**
Read is the default permission that is assigned to the Everyone group. Read allows:

- Viewing file names and subfolder names

- Viewing data in files

- Running program files

**Change**
Change is not a default permission for any group. The Change permission allows all Read permissions, plus:

- Adding files and subfolders
- Changing data in files
- Deleting subfolders and files

**Full Control**
Full Control is the default permission that is assigned to the Administrators group on the local computer. Full Control allows all Read and Change permissions, plus:

- Changing permissions (NTFS files and folders only).

Perform the following procedure to assign the Full Control share permission to the Everyone group.

**To assign the Full Control share permission to the Everyone group**

1. Open **Windows Explorer**.
2. Right-click the folder for which you want to administer share permissions, click **Properties**, and then click the **Sharing** tab.
3. Click **Permissions**.
4. Click **Everyone**, and then click **Full Control**.

# Additional Resources for Troubleshooting Access Control

Updated: June 8, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For specific information about troubleshooting access control problems, see the following resources:

- TechNet Newsgroup: Windows Server Security on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47169).
- Technet Newsgroup: Windows File System on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47172).
- TechNet Newsgroup: Windows NT File System on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47168).
- Article 834859, "The specified domain either does not exist or could not be contacted" error message in Windows Server 2003 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47170).

For general information about how access control works and how to manage and configure access control, see the following resource:

- Security Descriptors and Access Control Lists Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47171).

# Server Cluster Operations Guide

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Server Cluster Operations Guide provides troubleshooting information for the quorum resource in a server cluster running Microsoft® Windows Server™ 2003 or Windows Server 2003 with Service Pack 1 (SP1) operating systems.

**In this guide**


- Troubleshooting the Quorum Resource in a Server Cluster

# Troubleshooting the Quorum Resource in a Server Cluster

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide explains how to troubleshoot the quorum resource in a server cluster running Microsoft Windows Server 2003, Enterprise Edition, or Microsoft Windows Server 2003, Datacenter Edition, operating systems. The guide includes information for these operating systems as originally released and for these operating systems with Service Pack 1 (SP1). Some of the information will also apply to server clusters running the Windows 2000 Advanced Server or Windows 2000 Datacenter Server operating systems.

This guide covers only clusters using the single quorum device model, not the majority node set model.

To get started with this guide, review the following sections of this introduction.

When to Use This Guide

How This Guide Is Organized

Terminology Used in This Guide

**When to Use This Guide**

You should use this guide when, on a server cluster that uses a single quorum device (storage that is attached to all nodes of the cluster), the Cluster service will not start normally on any node, and the problem seems as if it might be related to the quorum resource. In some cases, you can use this guide when you can start the Cluster service, but messages in the event log or cluster log indicate that you have a problem with the quorum resource.

For information about distinguishing quorum problems from other problems that interfere with starting the Cluster service, see Troubleshooting Quorum Resource Problems.

Do not use this guide to find out how to perform a task, such as configuring the quorum resource for a new server cluster. You can find information about how to perform tasks and configure settings on the following pages on the Microsoft Web site:

- Designing and Deploying Server Clusters
- Checklists: Creating Server Clusters
- Server Cluster How To...

This guide assumes that you have a basic understanding of the quorum resource in a single quorum device server cluster, how the quorum works, and how the quorum helps ensure that all nodes in the cluster have the same view of the cluster membership and cluster configuration. You should also have a thorough understanding of how the quorum resource is configured in your organization.

**How This Guide Is Organized**

This guide is divided into several sections:

- Verifying Permissions, Hardware, and Software Before Troubleshooting the Quorum Resource

  This section provides a list of prerequisites to verify before you troubleshoot the quorum resource. Read this section first.

- Configuring a Computer for Troubleshooting the Quorum Resource in a Server Cluster

  This section describes how to configure your computer for troubleshooting the quorum resource. Read this section to prepare for troubleshooting.

- Troubleshooting Quorum Resource Problems

  This section provides step-by-step diagnostic procedures and possible solutions that help you identify and fix quorum resource problems in a single quorum device cluster.

- Additional Resources for Troubleshooting the Quorum Resource in a Server Cluster

This section provides additional resources that are related to server clusters and the quorum in a server cluster.

#### Terminology Used in This Guide

To use this guide, you should be familiar with basic terminology for a server cluster, particularly the following terms:

**single quorum device server cluster**

A cluster configuration that has two or more nodes and is configured so that every node is attached to one or more cluster storage devices. The cluster configuration data is stored on a single cluster storage device, called the quorum resource.

There are two other possible cluster configurations: a single node server cluster and a majority node set server cluster.

**cluster configuration database**

A database containing configuration data for all cluster objects. Copies of the cluster configuration database are located on the quorum resource and on each cluster node. Under normal conditions, the copies are synchronized and contain the same information.

**quorum resource**

A crucial cluster resource that maintains the configuration data that is necessary for recovery of the cluster. The quorum resource is generally accessible to other cluster resources so that any cluster node has access to the most recent changes to the cluster configuration database.

This document provides information for clusters where the quorum resource is located on a single quorum device, as described in "single quorum device server cluster," earlier in this list.

**quorum log**

A log maintained on the quorum resource to store cluster state and configuration changes during times when a node is down. The quorum log ensures that even if all nodes go down and the cluster must be formed again, those changes will not be lost. Also known as the *recovery log* or *change log*.

# Verifying Permissions, Hardware, and Software Before Troubleshooting the Quorum Resource

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin troubleshooting, you should verify permissions, hardware, and software updates, as described in the following list:

- Make sure that you have administrative rights on the computer you are troubleshooting.

  You cannot modify quorum resource settings unless you are a member of the Administrators group on the computer that you are administering.

  To verify that you are a member of the Administrators group on a computer, open the Computer Management snap-in. In the console tree, double-click **Local Users and Groups**, and then click **Groups**. In the details pane, double-click **Administrators** and verify that the account that you logged in with appears in the **Members** list.

- Confirm that your hardware is compatible with Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. For more information, see the following pages on the Microsoft Web site:

  - Windows Server Catalog

  - The Microsoft support policy for server clusters, the Hardware Compatibility List, and the Windows Server Catalog

- Confirm that cluster hardware is not disconnected or damaged. For example, check the cables to the cluster storage. If you use SCSI, check that the SCSI bus is properly terminated.

- Set aside time to review the software updates and service pack level on each node. Consider whether the latest service pack, critical updates, and security updates need to be tested and deployed on the cluster, and if so, choose a time to do so. Carry out this testing and deployment using standard methods for your organization.

  If you have not already planned for the testing and deployment of the latest service pack for Windows Server 2003, we strongly recommend that you plan for this. The service pack might enable the cluster to automatically correct future problems with the quorum resource — for example, corruption of a cluster configuration file on the quorum resource. For information about how to install service packs in a cluster, see article 174799 in the Microsoft Knowledge Base.

- Set aside time to review the drivers and software, including non-Microsoft drivers and software, on each node. Consider whether updated drivers and software need to be tested and deployed on the cluster, and if so, choose a time to do so. Carry out this testing and deployment using standard methods for your organization.

# Configuring a Computer for Troubleshooting the Quorum Resource in a Server Cluster

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can use advanced troubleshooting techniques to identify and fix quorum troubleshooting problems, we recommend that you obtain the following tool:

- **Clusterrecovery.exe**, available from Windows Server 2003 Resource Kit Tools on the Microsoft Web site.

  Clusterrecovery is especially useful for helping you to:

  - Replace a failed disk or recover from disk signature changes.
  - Move a disk resource (for example, the quorum resource) to a different disk in the cluster storage.
  - Restore registry checkpoints and cryptographic key checkpoints for cluster resources.

# Troubleshooting Quorum Resource Problems

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides troubleshooting information about a server cluster for which all of the following are true:

- The cluster uses a single quorum device (storage that is attached to all nodes of the cluster). In other words, the cluster is not a single node cluster and is not a majority node set cluster.
- The Cluster service will not start on any node.

  Note that in some situations, the Cluster service might start despite a problem with the quorum resource, and you can use this guide for troubleshooting in those situations. For example, a cluster in which the latest service pack has been applied might have repeated instances of Event ID 1147 related to the quorum, yet the Cluster service might be able to start on the nodes.
- The problem seems as if it might be related to the quorum resource. For more information, see Distinguishing Quorum Problems from Other Problems.

If you already have some understanding of the problem with the quorum resource, use the following topics to begin troubleshooting. If not, use the flowchart and other information in this topic to begin to determine which of the following topics apply to your situation.

- Files on the Cluster Quorum Might be Missing, Inaccessible, or Corrupt
- The Quorum Resource is on an Inaccessible or Nonfunctioning Disk
- The Cluster Configuration File on a Node is Corrupt

**Sections in This Topic**

This topic provides the following sections to help you begin troubleshooting:

- Distinguishing Quorum Problems from Other Problems
- Flowchart for Troubleshooting the Quorum Resource
- Understanding the Fixquorum Option for Starting the Cluster Service

**Distinguishing Quorum Problems from Other Problems**

It can be challenging to analyze a problem with starting the Cluster service and determine whether it is related in some way to the quorum resource. To try to distinguish different types of problems, review the following table:

| Area | Explanation |
|---|---|
| **Cluster service account**: On the nodes on which you are trying to start the Cluster service, are there problems with the Cluster service account?<br><br>Also, if you have Windows Server 2003 with the latest service pack, does Event ID 7041 appear in the system event log? This event was added in Windows Server 2003 Service Pack 1 and indicates that the Cluster service account does not have all the necessary user rights, such as the **Log on as a service** user right. | If no domain controller is available to authenticate the Cluster service account, the Cluster service cannot start. Other problems with the account can also prevent the Cluster service from starting, for example, if the password was allowed to expire or the account does not have the necessary rights, possibly as a result of a Group Policy setting.<br><br>To check for this type of problem, try to log on to the computer with the Cluster service account, or check the system event log for messages that indicate that the Cluster service account cannot log on. Also, review the Cluster service account to make sure that it has the permissions and rights described in Change the Account Under Which the Cluster Service Runs at the Microsoft Web site. |

| | |
|---|---|
| **Cluster log**: Is the cluster log read-only? Alternatively, is a policy preventing the cluster log from being modified? | If the cluster log cannot be modified, the Cluster service cannot start.<br><br>By default, the cluster log is called **Cluster.log** and is located in *systemroot***Cluster**. The default name of this log can be changed by changing the system environment variable called **ClusterLog**. |
| **Events in event log**: Do event log messages appear to be consistent with a quorum resource problem? | Specific events can indicate that a problem with the quorum is preventing cluster startup. However, an event can indicate a symptom rather than a root cause, so all events should be interpreted in context.<br><br>See the next section for information about some events that can indicate problems with the quorum resource. |

For more information about troubleshooting problems with starting the Cluster service, see article 266274, "How to Troubleshoot Cluster Service Startup Issues" in the Microsoft Knowledge Base.

If your startup issue appears to be related in some way to the quorum resource, see Flowchart for Troubleshooting the Quorum Resource.

**Event Messages That are Consistent with Quorum Resource Problems**
The following list describes some common event log messages that are consistent with quorum resource problems.

| **Important** |
|---|
| If your event log contains any of the messages in the following list, be sure to check that the cables for the storage are not damaged or disconnected. Also, follow the other recommendations in Verifying Permissions, Hardware, and Software Before Troubleshooting the Quorum Resource. |

### Event Messages That Are Consistent with Quorum Resource Problems

- **1034**: The disk associated with cluster disk resource *DriveLetter* could not be found. The expected signature of the disk was *DiskSignature*.

  This error can result when the disk signature of the quorum disk has been inadvertently changed. This can happen when you make changes through a disk or storage utility, or when you recreate the LUN containing the quorum resource. Note that the cluster identifies the quorum resource and other disk resources by disk signature, not just drive letter.

  Follow the flowchart in this topic, and see if The Quorum Resource is on an Inaccessible or Nonfunctioning Disk applies to your situation. Consider whether Clusterrecovery.exe, which is described in Configuring a Computer for Troubleshooting the Quorum Resource in a Server Cluster, would be useful for your situation.

- **1035**: Cluster disk resource *DriveLetter* could not be mounted.

  Follow the flowchart in this topic, and see if The Quorum Resource is on an Inaccessible or Nonfunctioning Disk applies to your situation.

- **1066**: Cluster disk resource *DriveLetter* is corrupt. Run **ChkDsk /F** to repair problems.

  This error might indicate that the volume used for the quorum resource has a file system corruption problem, possibly a transient problem. Other events in the logs might help with diagnosis.

  Follow the flowchart in this topic, and see if The Quorum Resource is on an Inaccessible or Nonfunctioning Disk applies to your situation.

- **1069**: Cluster disk resource *DriveLetter* failed.

  This error might indicate that the disk used for the quorum resource has a problem. Other events in the logs might help with diagnosis.

  Follow the flowchart in this topic, and see if The Quorum Resource is on an Inaccessible or Nonfunctioning Disk applies

to your situation.

- **1147** or **1148**:

  The Microsoft Clustering Service encountered a fatal error. The vital quorum log file 'q:\MSCS\quolog.log' could not be found. (Additional text is provided in this event.)

  If you have Windows Server 2003 with the latest service pack, you might see this event and yet find that the Cluster service started. When you apply the latest service pack, if the quorum log and cluster configuration files are missing from the quorum resource, the operating system can replace these files with information from a node.

  Follow the flowchart in this topic, and see if Files on the Cluster Quorum Might be Missing, Inaccessible, or Corrupt applies to your situation.

If you want to find out more details about a specific event that relates to the cluster, one source of information is the cluster log. The cluster log may be easiest to use when you already have a general sense of a problem and are looking for details. By default, the cluster log is called **Cluster.log** and is located in *systemroot***Cluster**. The log might have a different name, because the default name can be changed by changing the system environment variable called **ClusterLog**.

**Flowchart for Troubleshooting the Quorum Resource**

Use the following flowchart and the topics to which it refers to troubleshoot quorum resource problems. For information about the **/fixquorum** option mentioned in the flowchart, see Understanding the Fixquorum Option for Starting the Cluster Service, later in this topic.




**Understanding the Fixquorum Option for Starting the Cluster Service**

The troubleshooting techniques shown in the flowchart in this topic use the **/fixquorum** option for starting the Cluster service. It is useful to understand the following details about the **/fixquorum** option:

- When the **/fixquorum** option is used on a particular node, the Cluster service starts on that node, but the resources in the group that contains the quorum disk (usually the Cluster Group) remain offline. This allows you to attempt bringing the quorum resource online manually so that you can more easily diagnose problems with the quorum. Note that if this attempt to bring the quorum resource online fails, the Cluster service will stop again, even though it was started with the **/fixquorum** option.

- Only one node at a time can be started with the **/fixquorum** option. You cannot join any other nodes to the node that was started with this option. When you have corrected the problem that you were having with the quorum resource, stop the Cluster service and restart it without options.

- You can start the Cluster service with the **/fixquorum** option by opening a command prompt and typing:

  **net start clussvc /fixquorum**

  Or by typing:

  **net start clussvc /fq**

  You can also start the Cluster service with the **/fixquorum** option when using Services in Computer Management. To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**. On a cluster node, the Cluster service will appear in the list of services. Right-click the service, click **Properties**, type **/fixquorum** in **Start parameters**, and then click **Start**.

# Files on the Cluster Quorum Might be Missing, Inaccessible, or Corrupt

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To understand this topic in context, see the flowchart in Troubleshooting Quorum Resource Problems.

If you can start the Cluster service only when you use the **/fixquorum** option but then you succeed at bringing the quorum resource online (with Cluster Administrator), there is probably a problem with the configuration files on the quorum resource, as described in the "Cause" description in this topic. You can examine the files on the quorum resource and replace them if necessary.

**Cause**

There are a variety of possible causes. Like any files, the cluster configuration files on the quorum resource can become corrupted. Also, the files might be missing or they might be impossible to write to, for example, because of incorrect permissions or because the disk is full.

**Solution**

Use Windows Explorer to examine the files on the quorum resource. The files will be in a folder, almost always called MSCS, and in subfolders of that folder. Check the permissions on the files and the folder, and check whether the quorum disk is full. Then use the following tables to help you determine if any files are missing or corrupt. The procedures later in this topic can help you replace files as appropriate.

**Files That Should be Present on the Quorum Resource**

| File | Description |
|---|---|
| Chk*xxx*.tmp | Copy of the cluster configuration database (also known as a checkpoint file, but not to be confused with registry checkpoints for resources). If the file contains 0 bytes, it is definitely corrupt. Even if it contains more than 0 bytes, it could be corrupt.<br><br>If you have the original release of Windows Server 2003, the Cluster service will not start if the Chk*xxx*.tmp file is missing or corrupt. However, if you have Windows Server 2003 with SP1, in many situations the Cluster service can automatically re-create this file if it is missing or corrupt. |
| Quolog.log | The quorum log, which records changes to the cluster configuration database, but only those changes that occur while one or more nodes are down. The file exists even when all nodes are functioning, but information is added to it only when a node is offline. Information in the log is carefully marked according to sequence and timing so that it can be used correctly when nodes go down and come back up.<br><br>If you have the original release of Windows Server 2003, the Cluster service will not start if Quolog.log is missing or corrupt. However, if you have Windows Server 2003 with SP1, in many situations the Cluster service can automatically re-create this file if it is missing or corrupt. |

**Folders and Files That Might be Present on the Quorum Resource**

| File | Description |
|---|---|
| Folder called {*GUID*} | Folder for each resource that is configured with registry checkpoints or cryptographic key checkpoints. The resource GUID is the name of the folder. |
| {*GUID*}*.cpt | Checkpoint files for resources configured with registry checkpoints. |
| {*GUID*}*.cpr | Checkpoint files for resources configured with cryptographic key checkpoints. |

Use the following procedures to replace a missing or corrupted file.

**To replace a missing or corrupt Chkxxx.tmp file or quorum log or both on the quorum resource**

1. If the Cluster service is running (with the **/fixquorum** option), stop it by typing:

   **net stop clussvc**

2. On a node that was functioning correctly when problems with the quorum resource appeared, restart the Cluster service

with the **/resetquorumlog** option by typing:

**net start clussvc /rq**

This re-creates the Chk*xxx*.tmp file based on the cluster configuration information on that node.

You can accomplish the same action through Computer Management by opening Services and then, for the Cluster service, opening the **Properties** dialog box. From the **General** tab, stop the Cluster service if needed, specify **/resetquorumlog** in **Start parameters**, and then start the service with that parameter. (The parameter is used only once and does not persist.)

3. If the previous steps are not workable for you, and you have a backup containing quorum information (that is, a recent System State backup created on a node while the Cluster service was running), start Backup on a node. Before restoring the System State, in the **Confirm Restore** dialog box, click **Advanced**, and then click **Restore the Cluster Registry to the quorum disk and all other nodes**.

| **Important** |
| --- |
| Be sure to click this advanced option, or information will not be restored to the quorum resource, only to the node. |

4. If you have resources that are configured to use registry checkpoints or cryptographic key checkpoints and you think that these checkpoint files are also corrupt, see the next procedure.

**To replace corrupt registry checkpoint files or cryptographic key checkpoint files on the quorum resource**

1. Correct other problems with the quorum resource and start the Cluster service without options.
2. Run Clusterrecovery.exe, which is described in Configuring a Computer for Troubleshooting the Quorum Resource in a Server Cluster. Use Clusterrecovery to recover the resource checkpoint files, referring to Clusterrecovery Help as needed.
3. If the previous steps are not workable for you, and you have a recent backup of your registry checkpoint files or cryptographic key checkpoint files (that is, *.cpt or *.cpr files), restore from backup. To do this:
   a. Correct other problems with the quorum resource and start the Cluster service without options. Identify the resources that are configured to use registry checkpoints or cryptographic key checkpoints and take them offline.
   b. Identify the node that currently owns the quorum resource. On that node, restore the *.cpt or *.cpr files to the correct folder on the quorum resource. Refer to the table earlier in this topic for information about where the files should be restored.

# The Quorum Resource is on an Inaccessible or Nonfunctioning Disk

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To understand this topic in context, see the flowchart in Troubleshooting Quorum Resource Problems.

If you can start the Cluster service with the **/fixquorum** option and can use Cluster Administrator to bring nonquorum disk resources online, but cannot bring the quorum resource online, the disk used for the quorum resource might be inaccessible or might have failed. Note that when the quorum resource cannot be brought online in this way, the Cluster service will stop again, even though it was started with the **/fixquorum** option.

**Cause**
A disk might be inaccessible because of incorrectly configured hardware or software, for example, multipath I/O (MPIO) software. There are multiple causes for failure of a hard disk.

**Solution**
For further diagnosis, you can study events in the system log for indications of hardware errors. Then, use Device Manager to review the condition of the disk that the quorum resource is on. If there is a disk signature problem but the quorum disk can function otherwise, correct the disk signature problem. If the disk will not function, use either Clusterrecovery or Cluster Administrator to use a different disk for the quorum resource.

**To use Device Manager to review the condition of the disk that the quorum resource is on**

1. On the node on which you have started the Cluster service with the **/fixquorum** option, click **Start**, click **Programs** or **All Programs**, click **Administrative Tools**, and then click **Computer Management**.

2. Click **Device Manager**.

3. In the Details pane, click **Disk drives** and then try to view the disk that the quorum resource is on.

   If you cannot view the disk and want to further diagnose its condition, see diagnostic tools or information provided by the disk manufacturer, or contact the disk manufacturer for assistance.

   If you can view the disk, look in the event log for an instance of Event ID 1034 related to the quorum resource, and record the disk signature expected by the Cluster service for the quorum resource. Then use one of the following procedures to correct problems with the disk.

**To use Clusterrecovery to correct problems with the quorum being on a nonfunctioning disk**

- On the node on which you have started the Cluster service with the **/fixquorum** option, run Clusterrecovery.exe, which is described in Configuring a Computer for Troubleshooting the Quorum Resource in a Server Cluster. Use Clusterrecovery to substitute a new disk for the disk that has been having problems, referring to Clusterrecovery Help as needed.

  Clusterrecovery makes it relatively easy to configure a new disk to replace the disk that has developed a problem.

**To use a different disk for the quorum resource without using Clusterrecovery**

1. After starting the Cluster service with the **/fixquorum** option, from the disks available on the cluster storage, choose the disk that you will use for the quorum resource. If the disk has two or more NTFS partitions, ensure that all partitions on the disk are assigned drive letters.

2. Using Cluster Administrator, create a Physical Disk resource or other storage-class resource for the new disk (unless one has already been created).

3. In the console tree, click the **Resources** folder.

4. In the details pane, click the resource that you will use for the quorum resource. If the resource is online, on the **File** menu, click **Take Offline**.

5. In the console tree, click the cluster name and then on the **File** menu, click **Properties**.

6. On the **Quorum** tab, click **Quorum resource**, and then select the disk or storage-class resource you want to use as the

quorum resource for the cluster.

7. In **Partition**, if the disk has more than one partition, click the partition to be used for the quorum resource.
8. In **Root path**, type the path to the folder on the partition. For example:

   \MSCS\

9. In the console tree, click the **Resources** folder and then, in the details pane, click the new quorum resource.
10. On the **File** menu, click **Bring Online** to confirm that you can bring the new quorum resource online.
11. Stop the Cluster service by typing:

    **net stop clussvc**

12. On the same node on which you performed the preceding steps, restart the Cluster service by typing:

    **net start clussvc**

| **Important** |
| --- |
| You must perform this step, which forms the cluster, from the node on which you performed the preceding steps. This is because the new information that you configured on this node does not exist yet on the other nodes. After the other nodes have joined the cluster, they will receive the new information about the location of the quorum resource and will be able to form the cluster as needed. |

# The Cluster Configuration File on a Node is Corrupt

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To understand this topic in context, see the flowchart in Troubleshooting Quorum Resource Problems.

If you try to start the Cluster service with the **/fixquorum** option on one node at a time and discover that this fails on one node although it succeeds on another, the cluster configuration file on the node from which you cannot start the Cluster service might be corrupt. This is a separate problem from problems with the quorum resource, but it is covered here because you might discover it while troubleshooting a problem with the quorum resource.

**Cause**
There are a variety of possible causes. Like any files, the cluster configuration file (CLUSDB) on a node can become corrupted.

**Solution**
First, correct any problems with the quorum resource and get the cluster running again (without the problem node). Analyze the local disk on the problem node and take appropriate corrective steps, as described in the procedure that follows. If the local disk on the node appears to be free of problems or has been fixed, try to join the cluster from that node. In some cases, the node will be able to join, and the cluster configuration file on the node will be re-created automatically. However, if it appears that a corrupt cluster configuration file is preventing the node from joining the cluster, use the procedure in this section to get a good copy of the cluster configuration database from the quorum resource and copy it to the node.

**To analyze the state of the local disk on a node that cannot be started with the /fixquorum option**

1. On the node on which you cannot start the Cluster service with the **/fixquorum** option, try to load and unload the cluster registry hive (collection of keys, subkeys, and values) by carrying out the following steps:

   a. At a command prompt, type:

      **regedit**

   b. Click HKEY_LOCAL_MACHINE and make sure it is selected.

   c. On the **File** menu, click **Load Hive**.

   d. Browse to the *Systemroot***Cluster** folder, and then double-click **CLUSDB**. When you are prompted for a key name, type **cluster**.

   e. If you can successfully load the hive, under HKEY_LOCAL_MACHINE, make sure that **Cluster** is selected and then, on the **File** menu, click **Unload Hive**. When prompted for confirmation, click **Yes**. At this point, you might need to perform additional troubleshooting. One step to take is to compare the cluster hive on this node with the cluster hive on another node. Another step is to make sure the cluster hive is unloaded and then try starting the Cluster service again on the node. A third step is to check the system event log for errors that indicate that incorrect permissions or user rights are preventing the Cluster service account from accessing or loading the local registry hive.

   f. If you cannot load the hive, it is corrupt. Continue to the next step.

2. If **CLUSDB** is corrupt, or if you want to check the condition of the local disk, run a disk utility that provides preliminary information about whether the disk appears to have problems. For example, to find out whether the dirty bit is set on the disk (indicating the file system may be in an inconsistent state), you can type the following at a command prompt:

   **fsutil dirty query** *d* **:**

   where *d***:** is the disk you are interested in. Another option is to use the Vrfydsk.exe utility, which is available from Windows Server 2003 Resource Kit Tools at the Microsoft Web site. Using information from the previous step along with information about the age and previous history of the local disk, assess whether the disk probably contains errors and needs to be fixed. To fix the disk, at a command prompt, type:

   **chkdsk** *d* **: /f**

or

**chkdsk** *d* **: /f /r**

where *d***:** is the disk you want to fix. Note that on a large disk volume, this command will take a long time to run, especially if the **/r** option is used.

3. If information from the previous steps and other disk symptoms indicate that the local disk is failing, replace the disk, restore from backup, and assess whether further troubleshooting is needed.

If the local disk on the node appears to be free of problems or has been fixed, try to join the cluster from that node. In some cases, the node will be able to join, and the cluster configuration file on the node will be re-created automatically. However, if it appears that a corrupt cluster configuration file is preventing the node from joining the cluster, use the following procedure to get a good copy of the cluster configuration database from the quorum resource and copy it to the node.

**To replace a corrupt cluster configuration file on a cluster node**

1. In Cluster Administrator, view the functioning nodes in the cluster, and find the node that owns the quorum resource.
2. From the node that owns the quorum resource, view the files on the quorum resource, and locate the Chk*xxx*.tmp file.
3. On the problem node, which is not joined to the cluster at the moment, in the *systemroot***\cluster** folder, locate the CLUSDB file (which you have determined is corrupt) and then rename it.
4. Copy the Chk*xxx*.tmp file to the *systemroot***\cluster** folder on the problem node, and then rename that file **CLUSDB**.

   If the problem has been corrected on the node, you will be able to start the Cluster service with no start parameters.

# Additional Resources for Troubleshooting the Quorum Resource in a Server Cluster

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For specific information about troubleshooting server clusters, see the following resources on the Microsoft Web site:

- Windows Server 2003 Clustering Services Technology Center
- Change the Account Under Which the Cluster Service Runs
- How to Install Service Packs in a Cluster
- How to Troubleshoot Cluster Service Startup Issues

For information about best practices for server clusters, see the following resources:

- Server Cluster Best Practices
- Backing Up and Restoring Server Clusters
- Server Clusters: Network Configuration Best Practices for Windows 2000 and Windows Server 2003
- Server Clusters: Security Best Practices for Windows 2000 and Windows Server 2003
- Server Clusters: Storage Best Practices for Windows 2000 and Windows Server 2003

For general information about how server clusters work, see the following resources:

- Planning Your Server Cluster
- Using Server Clusters
- How a Server Cluster Works

# Terminal Server Licensing Operations Guide

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Terminal Server Licensing Operations Guide provides troubleshooting information for Terminal Server licensing processes in the Windows Server 2003, Windows Server 2003 with Service Pack 1 (SP1) and Windows 2000 Server operating systems.

**In this guide**


- Troubleshooting Terminal Server Licensing

# Troubleshooting Terminal Server Licensing

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides troubleshooting information for Terminal Server Licensing in the Microsoft Windows Server 2003 and Windows 2000 operating systems.

**In this guide**


- Introduction to Troubleshooting Terminal Server Licensing
- Verifying Computer Settings for Troubleshooting Terminal Server Licensing
- Quick Fixes for Terminal Server Licensing
- Configuring a Computer for Troubleshooting Terminal Server Licensing
- Troubleshooting Terminal Server Licensing Problems
- Additional Resources for Troubleshooting Terminal Server Licensing



**Acknowledgments**

Produced by: Microsoft Windows Server User Assistance team

Project Writer: Connie La Chasse

Project Editor: Stephanie Marr

Technical Reviewers: Craig Landis, Aimme Lirette, Russ Morgan, Aruna Somendra (KOKA), Tim Thompson, Dean Tsaltas, Joost van Haaren, Steve Wesa

Microsoft Most Valuable Professional (MVP) Reviewers: Christa Anderson

# Introduction to Troubleshooting Terminal Server Licensing

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are using this guide for the first time, review the following sections of this introduction:

When to Use This Guide

How to Use This Guide

**When to Use This Guide**
You should use this guide when:

- You have a problem that you believe is related to Terminal Server Licensing, but you do not know how to resolve it.
- You have configured Terminal Server Licensing settings, but Terminal Server Licensing does not behave the way you anticipated.
- A program is not working properly and you believe that Terminal Server Licensing is causing the problem.

This guide assumes a basic understanding of Terminal Server Licensing, how it works, and why your organization uses it to provide, catalog, and enforce license policy among Terminal Server clients. You should also have a thorough understanding of how Terminal Server Licensing is deployed and managed in your organization. This includes an understanding of the mechanism your organization uses to configure and manage Terminal Server Licensing settings.

**How to Use This Guide**
This guide is divided into five sections. Each section addresses a type of problem, ranging from less complex to more complex.

Verifying Computer Settings for Troubleshooting Terminal Server Licensing

This section provides a list of prerequisites and settings that must be verified before you troubleshoot. Read this section first.

Quick Fixes for Terminal Server Licensing

This section provides quick solutions to the most common Terminal Server Licensing problems. Read this section if you have verified your computer settings, but your problem persists.

Configuring a Computer for Troubleshooting Terminal Server Licensing

This section describes how to configure your computer for troubleshooting. Read this section if you cannot find a solution in the Quick Fixes section.

Troubleshooting Terminal Server Licensing Problems

This section provides step-by-step diagnostic procedures and possible solutions that help you identify and fix Terminal Server Licensing problems.

Additional Resources for Troubleshooting Terminal Server Licensing

See this section for a comprehensive list of resources available to help you troubleshoot Terminal Server Licensing.

# Verifying Computer Settings for Troubleshooting Terminal Server Licensing

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Settings to Verify Before You Troubleshoot**
Before you begin troubleshooting, verify that your computer is configured properly and that Terminal Server Licensing is set up and running correctly. Following are important general questions to consider before you begin more specific troubleshooting:

- Which operating system is the client running (the client that is attempting to connect to the terminal server)?
- Which operating system is the terminal server running (the terminal server to which the client is attempting to connect)?
- Do you have a computer that is running the Terminal Server Licensing service?
- Which operating system is the license server running?
- How many licenses are available on the license server?

After you have answered these questions, you can proceed with the following more specific troubleshooting steps:

Verify that you have administrative rights on the computer you are troubleshooting

Install all critical security updates for Windows 2000 or Windows Server 2003

Update all of your software, including non-Microsoft software

Verify that the operating system of the Terminal Server license server is compatible with the operating system of the terminal server

Review additional configuration requirements and considerations for implementing Terminal Server Licensing in different environments

Verify that Terminal Server Licensing is correctly installed on the license server

Verify that the Terminal Server Licensing service is started and that the startup type is set to Automatic

Verify that the license server is activated

Verify that the license server has a sufficient number of CALs for clients

Verify that the Terminal Server Licensing mode on the terminal server matches the type of CALs on the license server (applicable to Windows Server 2003 only)

Verify that the Authenticated Users group on the terminal server has the appropriate rights

Verify the value of the RestrictAnonymous registry key

Verify that the TS-License-Server object exists in Active Directory, if the license server is an Enterprise license server

**Verify that you have administrative rights on the computer you are troubleshooting**
You cannot modify Terminal Server Licensing settings unless you are a member of the Administrators group on the computer that you are administering.

**To verify that you are a member of the Administrators group**

1. Open the Computer Management snap-in.
2. In the console tree, double-click **Local Users and Groups**, and then click **Groups**.
3. In the details pane, double-click **Administrators** and verify that your account name or a group in which your account is a

member appears in the **Members** list.

You can also tell if you have the appropriate administrative rights to configure Terminal Server Licensing by opening the Terminal Server Licensing tool. If you do not have administrative rights, all of the controls in the Terminal Server Licensing user interface (UI) will appear dimmed and a warning will appear telling you that you must be a computer administrator to change any settings.

### To open Terminal Server Licensing

- Click **Start**, point to **Control Panel**, **Administrative Tools**, and then click **Terminal Server Licensing**.

**Install all critical updates and security updates for Windows 2000 or Windows Server 2003**
Some updates might be required for Terminal Server Licensing to function properly.

### To verify that you have all critical updates and security patches

- Click **Start**, click **Windows Update**, and then follow the instructions that appear on your screen.

**Update all of your software, including non-Microsoft software**
Terminal Server Licensing might not function properly unless you update your operating system and programs with the most recent service pack or software updates. For example, if your license servers or terminal servers are running Windows 2000 Service Pack 3 (SP3) or earlier, you should install Service Pack 4 (SP4) on those servers. Newer versions of many programs, such as antivirus programs, might resolve any problems you have.

**Verify that the operating system of the Terminal Server license server is compatible with the operating system of the terminal server**
Note the following operating system compatibility requirements for Terminal Server Licensing.

- A license server running Windows Server 2000 cannot issue client access licenses (CALs) to terminal servers running Windows Server 2003 operating systems.

- A license server running a Windows Server 2003 operating system can issue CALs to terminal servers running Windows Server 2000 or Windows Server 2003, or to terminal servers running Windows Server 2000 and terminal servers running Windows Server 2003 in a mixed environment.

Verify that both the license server and the terminal server are running either the release version of Windows Server 2003 or a pre-release (evaluation) version of Windows Server 2003. A license server that is running a release version of Windows Server 2003 cannot communicate with a terminal server running an evaluation version of Windows Server 2003, and vice-versa.

| **Note** |
| --- |
| It is recommended that you do not install Terminal Server Licensing on a compressed drive. This can cause the license server database to become corrupted. The default location for the license server database is systemroot\System32\LServer. If the license server database is corrupted, then you must remove and then reinstall Terminal Server Licensing. If you encounter problems removing and reinstalling Terminal Server Licensing, you might need to contact the Microsoft Clearinghouse. For information about how to contact the Clearinghouse, see Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkId=48885). |

**Review additional configuration requirements and considerations for implementing Terminal Server Licensing in different environments**
To ensure that license server discovery works as expected, review configuration requirements and considerations, and verify that the requirements have been met for your environment. For information, see Understanding Troubleshooting Considerations for Specific Terminal Server Licensing Environments.

**Verify that Terminal Server Licensing is correctly installed on the license server**
For step-by-step instructions, see Verify that Terminal Server Licensing is correctly installed on the license server.

**Verify that the Terminal Server Licensing service is started and that the startup type is set to Automatic**
For step-by-step instructions, see Verify that the Terminal Server Licensing service is started and the startup type is set to Automatic.

**Verify that the license server is activated**
Windows Server 2003 Terminal Services supports the following licensing modes: Per Device and Per User. A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL

gives one user the right to access a terminal server from an unlimited number of devices.

By default, when the Per Device licensing mode is used, after a client logs on and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. To issue clients permanent, Per Device licenses, a license server must be activated.

| **Note** |
| --- |
| If Per User licensing is used, after the terminal server discovers a license server, no further communication takes place between the servers, and client connections are allowed regardless of the number of licenses installed. Therefore, when Per User licensing is used, the license server does not need to be activated. |

To activate a license server, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886).

**To verify that a license server is activated**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, point to **Control Panel**, **Administrative Tools**, and then click **Terminal Server Licensing**.
2. In the console tree, click the license server for which you want to verify the activation status.
3. In the details pane, verify that activation status is Activated.
4. If the status is Not Activated and you have Internet connectivity, right-click the server name, click **Activate Server**, select **Automatic connection (recommended)**, and then click **Next**.

   If you do not have Internet connectivity, see Activate a Terminal Server License Server (http://go.microsoft.com/fwlink/?LinkId=48888) for alternate license server activation procedures.
5. Follow the instructions in the wizard.

If necessary, at a later time you can change any license server properties, such as activation method and company information, that you set during the activation process.

**Verify that the license server has a sufficient number of CALs for clients**

You can use Terminal Server Licensing to verify whether a terminal server has a sufficient number of permanent, Per Device CALs to support the number of clients who require remote connections. When Per Device licensing is used, after a temporary license expires, a client can only connect to a terminal server if the license server can issue a permanent CAL, or if the terminal server is still within its licensing grace period. For step-by-step instructions on how to verify the number of CALs a license server has, see Verify that the license server has a sufficient number of CALs.

**Important** Per User CALs are not enforced by Terminal Server. As a result, client connections can occur regardless of the number of licenses installed. This does not remove administrators from End User License Agreement (EULA) requirements to have a valid terminal server CAL for each user, however. Failure to have a Per User CAL for each user, if Per Device CALs are not being used, is a violation of the EULA. To ensure that you are in compliance with the EULA, make sure that you track the number of Per User licenses being used in your organization and that you have a sufficient number of licenses to provide a Per User CAL for each user.

The following text is from the EULA for Windows Server 2003:

"Two different TS CALs are available to you: 'Device' and 'User.' Each TS Device CAL permits one Device (used by any User) to conduct Windows Sessions on any of your Servers. Each TS User CAL permits one User (using any Device) to conduct Windows Sessions on any of your Servers. You may use a mix of TS Device CALs and TS User CALs simultaneously with the Server Software in your environment. You can have a Terminal Server request Per User licenses or Per Device (default) but not both simultaneously."

| **Note** |
| --- |
| On license servers running Windows Server 2003 Service Pack 1 (SP1), the numeric values for Per User licenses in the Total and Available columns in the Terminal Server Licensing tool are no longer displayed. Instead, the phrase "Not applicable" appears in these columns, to remind administrators that Per User licenses are not enforced. This phase does not indicate a problem with the installation or operation of Terminal Server Licensing. |

**Verify that the Terminal Server Licensing mode on the terminal server matches the type of CALs on the license server (applicable to Windows Server 2003 only)**

For Terminal Server Licensing to function properly, you must ensure that the Terminal Server Licensing mode that you configure on the terminal server matches the type of CALs you have purchased and installed on the license server. You can verify whether the Terminal Server Licensing mode and type of CALs purchased match by doing one of the following:

Using Terminal Services Configuration (the most direct method)

Using Group Policy (in Windows Server 2003 Service Pack 1)

| **Note** |
| --- |
| If you use Group Policy to verify the Terminal Server Licensing mode and if the Group Policy setting is not properly applied to the Group Policy object (GPO) to which the terminal server is assigned, then the licensing mode that is displayed in Group Policy might not be valid. Therefore, using Terminal Services Configuration to verify the Terminal Server Licensing mode is a more direct method. |

**Using Terminal Services Configuration**

**To verify that the Terminal Server Licensing mode on the terminal server matches the type of CALs on the license server using Terminal Services Configuration**

1. On the terminal server, open **Terminal Services Configuration**. To open Terminal Services Configuration, click **Start**, point to **Control Panel**, **Administrative Tools**, and then click **Terminal Services Configuration**.
2. In the console tree, click **Server Settings**.
3. In the details pane, right-click **Licensing Mode**, and then click **Properties**. Note the licensing mode (**Per User** or **Per Device**), and then click **OK**.

   | **Note** |
   | --- |
   | A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices. |

4. On the license server, open **Terminal Server Licensing**.
5. In the details pane, verify that the type of CALs installed on the license server match the Terminal Server Licensing mode that is configured on the terminal server.

**Using Group Policy**

**To verify that the Terminal Server Licensing mode on the terminal server matches the type of CALs on the license server using Group Policy**

1. On the terminal server, open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click **Set the Terminal Server licensing mode**.
3. If the setting is enabled, note the licensing mode (**Per Device** or **Per User**) that is selected in **Specify the licensing mode for the terminal server**.

   If the setting is disabled or not configured, the licensing mode that is specified during Setup or in Terminal Services Configuration is used.

   | **Note** |
   | --- |
   | A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices. |

4. Click **OK**.
5. On the license server, open Terminal Server Licensing.
6. In the details pane, verify that the type of CALs installed on the license server match the Terminal Server Licensing mode

that is configured on the terminal server.

**Verify that the Authenticated Users group on the terminal server has the appropriate rights**

For Terminal Server Licensing to function correctly, the Authenticated Users group on the terminal server must have the following two rights: **Access this computer from the network** and **Bypass Traverse Checking**. By default, the only Group Policy object that has the **Access this computer from the network** right defined is the Default Domain Controllers policy. Member servers and servers in a workgroup environment have the **Access this computer from the network** right defined in their local security policy.

**To verify that the Authenticated Users group on the terminal server has the appropriate rights**

1. On the terminal server, open the appropriate policy (local policy or Group Policy).
2. In the console tree, navigate to Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\.
3. In the details pane, right-click **Access this computer from the network**, and then click **Properties**.
4. If Authenticated Users is not in the list of users and groups that have been granted this right, click **Add User or Group**.
5. In **Select Users or Groups**, in **Enter the object name to select**, type Authenticated Users, click **OK**, and then click **OK** again to close the dialog box for this security setting.
6. In the console tree, right-click **Bypass Traverse Checking**, and then click **Properties**.
7. If **Authenticated Users** is not in the list of users and groups that have been granted this right, perform steps 5 and 6.

Additionally, users must also be members of the Remote Desktop Users group on the terminal server, as described in the following procedure, or you must assign the equivalent permissions manually. For more information, see Enabling users to connect remotely to the server (http://go.microsoft.com/fwlink/?LinkId=48887).

**To add users to the Remote Desktop Users group**

1. Open Computer Management. **T**o open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
2. In the console tree, click the **Local Users and Groups** node.
3. In the details pane, double-click the Groups folder.
4. Double-click **Remote Desktop Users**, and then click **Add**. By default, the Remote Desktop Users group is not populated. You must decide which users and groups should have permission to log on remotely, and then manually add them to the group.
5. On the **Select Users** dialog box, click **Locations...** to specify the search location.
6. Click **Object Types...** to specify the types of objects you want to search for.
7. Type the name you want to add in the **Enter the object names to select (examples):** box.
8. Click **Check Names**.
9. When the name is located, click **OK**.

**Verify the value of the RestrictAnonymous registry key**

If you are not using Windows Server 2003 terminal servers with Windows Server 2003-based license servers, verify that the **RestrictAnonymous** registry key value on the license server is set to **0** or **1**. In this scenario, if this registry key value is set to **2**, then the license server cannot issue CALS. When the **RestrictAnonymous** registry value is set to **2**, the access token built for non-authenticated users does not include the Everyone group, and because of this, the access token no longer has access to resources that grant permissions to the Everyone group.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**To verify the value of the RestrictAnonymous registry key**

1. On the license server, open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Locate, and then click, the following key in the registry:
   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\RestrictAnonymous**
3. Verify the value of the **RestrictAnonymous**registry key.
4. The three valid registry key values for **Restrict Anonymous** are equivalent to the following security settings in Local Security Policy:

   **0** = None (rely on default permissions)

   **1** = Do not allow enumeration of SAM accounts and shares

   **2** = No access without explicit anonymous permissions
5. If the value of the registry key is set to **2**, then perform the following steps.
6. Open Local Security Policy, and then navigate to Security Settings\Local Policies\Security Options. To open Local Security Policy, click **Start**, point to **Settings**, Control Panel, click **Administrative Tools**, and then click **Local Security Policy**.
7. Double-click **Additional restrictions for anonymous connections**, and then click **No access without explicit anonymous permissions**.
8. Restart the server.

   If the license server is a member of an Active Directory domain and a conflicting security setting is configured for the license server in Group Policy, the Group Policy setting overrides the local security setting. Therefore, to ensure that the required setting changes take effect (if you must change the settings) this procedure describes how to configure the setting change in Group Policy.

   In Windows Server 2003, you cannot set **RestrictAnonymous** to a value of **2** to prohibit anonymous connections. If you need to prohibit anonymous users from being granted the same access that is granted to members of the Everyone group, you must use the new **Everyone Network access: Let Everyone permissions apply to anonymous users** setting in Local Security Policies.

**Verify that the TS-License-Server object exists in Active Directory, if the license server is an Enterprise license server**
To determine the license server role (whether the license server is an Enterprise license server or a Domain license server), check the **Role** registry key value.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**To check the Role registry key value**

1. On the license server, open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Locate, and then click, the following key in the registry:
   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\TermServerLicensing\Parameters**

   If the registry key value is set to **0**, then the license server is a Domain license server.

   If the registry key value is set to **1**, then the license server is an Enterprise license server.

If the license server is an Enterprise License server, then the **TS-Enterprise-License-Server** object must exist in Active Directory.

**To verify that the TS-Enterprise-License-Server object exists in Active Directory**

1. Open Active Directory Sites and Services. To open Active Directory Sites and Services, click **Start**, click **Control Panel**, click **Performance and Maintenance**, click **Administrative Tools**, and then **double-click Active Directory Sites and Services**.

2. In the console tree, click **Sites**, and then click the site that contains the terminal server.
3. In the details pane, click **Licensing Site Settings**.
4. Right-click **TS-Enterprise-License-Server**, and then click **Properties**. The **SiteServer** attribute of the **TS-Enterprise-License-Server** object must specify the Distinguished Name (also known as DN) of the license server. For example: CN=TS-Enterprise-License-Server, CN=<**SiteName**>, CN=Configuration, DC=<**DomainName**>,DC=<**RootDomainName**>
5. On the **License Settings** tab, verify the name of the license server.

# Quick Fixes for Terminal Server Licensing

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides a summary of the most common Terminal Server Licensing problems and the solutions to those problems. You can use the information in this section to resolve problems in the same way you would use a FAQ to find answers to common questions. Read this section before you begin any advanced troubleshooting.

**Common Terminal Server Licensing Problems**
Listed below are common Terminal Server Licensing problems and quick fixes:

The Terminal Server licensing mode on the terminal server is set to Per Device, but the license server is set to issue Per User CALs

Windows-Based Terminal devices (thin clients) cannot connect to terminal servers

Clients connecting to a Windows 2000 terminal server cannot obtain a permanent Windows 2000 Terminal Server CAL

**The Terminal Server licensing mode on the terminal server is set to Per Device, but the license server is set to issue Per User CALs**
When there is a mismatch between the Terminal Server Licensing mode configuration on the terminal server and the configuration on the license server, remote sessions can be disconnected prematurely. You can resolve this problem by making sure that the Terminal Server licensing mode on the terminal server matches the type of licenses installed on the license server. For more information, see Change the Terminal Server Licensing mode.

| **Note** |
| --- |
| If you are running a Windows Server 2003 terminal server, after you use the Windows Components Wizard in Add or Remove Programs to add or to remove a Windows component, an invalid registry value might be set in the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\TerminalServer\Licensing Core\PolicyAcOn** subkey on the terminal server. This invalid value may unexpectedly change the Terminal Server Licensing mode and cause the terminal server to obtain only temporary licenses. To resolve this problem contact Microsoft Product Support Services to obtain the hotfix described in Knowledge Base article 834651, Terminal Services licensing mode changes from Per User or from Per Device to an invalid value after you add or remove a Windows component (http://go.microsoft.com/fwlink/?LinkId=48173) on the Microsoft Web site. This hotfix is applied to the terminal server. |

**Windows-based terminal devices (thin clients) cannot connect to terminal servers**
If you try to connect to connect to a terminal server from a Windows-based terminal device, the connection might be denied if:

- You update the electrically erasable programmable read-only memory (EEPROM) of the Windows-based terminal device
- You delete the following registry subkey: **HKEY_LOCAL_MACHINE\Software\Microsoft\MSLicensing\Store**
- The Windows-based terminal device has been issued a new temporary license
- You installed the hotfix described in Knowledge Base article 287687, Terminal Services Licensing Enhancements (http://go.microsoft.com/fwlink/?LinkId=48172) on the Microsoft Web site.

In such cases, the connection is denied because the license is treated as expired.

To resolve this problem, contact Microsoft Product Support Services to obtain the hotfix described in Knowledge Base article 287687 to apply to the terminal server, the license server, or both. For information about how to contact Microsoft Product Support Services, see Help and Support (http://go.microsoft.com/fwlink/?LinkId=48149) on the Microsoft Web site. If you are using a Citrix ICA client, you might also need to contact Citrix to obtain an appropriate fix.

**Clients connecting to a Windows 2000 terminal server cannot obtain a permanent Windows 2000 Terminal Server CAL**
In an environment with Windows 2000 terminal servers and Windows Server 2003 license servers, clients that connect to a Windows 2000 terminal server cannot obtain a permanent Windows 2000 client access license (CAL), even though Windows 2000 CALs are installed on the license server. Instead, the clients receive a Windows Server 2003 Terminal Server temporary license. This problem might occur because the CAL was not upgraded or renewed correctly.

To resolve this problem, do the following:

1. Install CALs automatically on the license server. Client computers can successfully upgrade from a temporary Windows Server 2003 Terminal Server license to a permanent Windows Server 2003 Terminal Server CAL.

2. Verify whether the clients can receive the new temporary licenses. If the clients cannot receive the new licenses, the terminal server might not be able to locate the license server. For troubleshooting procedures to resolve this issue, see The terminal server cannot locate the license server.

3. If the clients can connect to the terminal server, and if licenses are available on the license server and this problem still occurs, then contact Microsoft Product Support Services to obtain the hotfix described in Knowledge Base Article 837211, Client computers cannot obtain a permanent Windows 2000 Terminal Services Client Access License (CAL) from a Windows Server 2003 Terminal Services Licensing server (http://go.microsoft.com/fwlink/?LinkId=48146) on the Microsoft Web site, and then apply this hotfix to the license server. For information about how to contact Microsoft Product Support Services, see Help and Support (http://go.microsoft.com/fwlink/?LinkId=48149) on the Microsoft Web site.

# Configuring a Computer for Troubleshooting Terminal Server Licensing

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can use advanced troubleshooting techniques to identify and fix Terminal Server problems, it is helpful to configure your computer for troubleshooting. In addition, you need a basic understanding of troubleshooting concepts, procedures, and tools.

**Using Terminal Server Licensing Troubleshooting Tools**
Using the following tools can help you monitor Terminal Server Licensing status and diagnose the cause of Terminal Server licensing issues.

**Terminal Server Licensing**
You can use Terminal Server Licensing to do the following:

- View the list of discoverable license servers on your network.
- Connect to any discoverable license server on your network and view information about the licenses on that server.

  Using the Terminal Server Licensing tool, you can view the following information: the list of installed client license key packs, the total number of licenses in each client license key pack, and the number of available and issued licenses in each key pack; the date a CAL was issued to each computer; the name of the computer to which each CAL is issued; the issue and expiration dates of each temporary license; and the name of the computer to which the license was issued.

  | **Note** |
  | --- |
  | If you are running a Windows 2000 license server, expiration dates for temporary licenses are displayed only if you have Service Pack 3 (SP3) or later installed. |

- Manage Terminal Server licenses.
- Perform the following management tasks: activate a license server, install CALs, reactivate a license server, deactivate a license server, repeat the installation of CALs, and change the properties of a license server.

This tool is included with Windows Server 2003. You can access this tool on any Terminal Server license server. For more information, see Terminal Server Licensing administrative tool (http://go.microsoft.com/fwlink/?LinkId=48896).

**Terminal Services Licensing Reporter (Lsreport.exe)**
You can use Terminal Services Licensing Reporter (Lsreport.exe) to create a report with information about the CALs that are installed on license servers.

Lsreport.exe exports information from the license server's database into LSreport.txt, a tab-delimited text file. You can use this data to track license consumption and usage trends within your organization. This tool has been updated to include the client hardware ID, information that is useful for tracking the CALs that are issued to particular client devices.

To download Lsreport.exe, see Windows Server 2003 Resource Kit Tools (http://go.microsoft.com/fwlink/?LinkID=16721) on the Microsoft Download Center Web site.

**Terminal Services License Server Viewer (Lsview.exe)**
You can use Terminal Services License Server Viewer (Lsview.exe) to do the following:

- Display the name of each license server in the current domain and site, the type (enterprise or domain) of each license server in the domain, and the date and time that each license server is available.
- Create a log file with advanced diagnostic information about the discovery process.

To download Lsview.exe, see Windows Server 2003 Resource Kit Tools (http://go.microsoft.com/fwlink/?LinkID=16721) on the

Microsoft Download Center Web site.

**Terminal Server Client License Test tool (Tsctst.exe)**
You can use the Terminal Server Client License Test tool to display information about CALs that reside on client devices. By default, this tool displays information in the following fields: Issuer, Scope, Issued to computer, Issued to user, License ID, Type/Version, Valid From, and Expires On.

To download Tsctst.exe, see Windows Server 2003 Resource Kit Tools (http://go.microsoft.com/fwlink/?LinkID=16721) on the Microsoft Download Center Web site.

**Terminal Services Diagnostic Utility (Tstst.exe)**
You can use the Terminal Services Diagnostic Utility to display configuration information for terminal servers. This tool provides a grace period check to determine whether terminal servers are still within the licensing grace period, and it enumerates available license servers.

To download Tstst.exe, see Windows Server 2003 Resource Kit Tools (http://go.microsoft.com/fwlink/?LinkID=16721) on the Microsoft Download Center Web site. Note that there are eight specialty tools, each corresponding to a specific support scenario. The specialty tool that contains Tstst.exe is MPSRPT DirSrvc.exe.

# Troubleshooting Terminal Server Licensing Problems

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Terminal Server Licensing issues can result in client connection failures, license server activation failures, and license issuance and upgrade problems. The following topics provide troubleshooting steps to resolve these problems. Additionally, important updates for Terminal Server Licensing in Windows Server 2003 are described, including updates for Service Pack 1. Understanding these updates can help you prevent or minimize problems with Terminal Server Licensing.

**In this section**


- Understanding Windows Server 2003 updates for Terminal Server Licensing
- Understanding Troubleshooting Considerations for Specific Terminal Server Licensing Environments
- Fixing Client Connectivity Problems
- Fixing CAL Issuance and Upgrade Problems
- Fixing License Server Activation Problems and Other Licensing Problems

# Understanding Windows Server 2003 updates for Terminal Server Licensing

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Following are important changes for Terminal Server Licensing in Windows Server 2003, including updates for Service Pack 1 (SP1).

**Per Device Licenses Available in Windows Server 2003**
In Windows Server 2003, two types of client access licenses (CALs) are available: Per User and Per Device. A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices. In this case, only one CAL is needed instead of a CAL for each device.

By default, when the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs, however, are not enforced by Terminal Server. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.

**Important**

For Terminal Server Licensing to function correctly you must ensure that the Terminal Server Licensing mode that you configure on the terminal server matches the type of CALs you have purchased and installed on the license server. You can verify whether the Terminal Server Licensing mode and type of CALs purchased match by using Terminal Services Configuration, or in Windows Server 2003 SP1, Group Policy. For more information, see Verifying Computer Settings for Troubleshooting Terminal Server Licensing.

**Note**

On license servers running Windows Server 2003 SP1, the numeric values for Per User licenses in the Total and Available columns in the Terminal Server Licensing tool are no longer displayed. Instead, the phrase "Not applicable" appears in these columns, to remind administrators that Per User licenses are not enforced. This phrase does not indicate a problem with the installation or operation of Terminal Server Licensing.

**Required Licenses for Terminal Server in Windows Server 2003**
For Windows Server 2003 Terminal Server to function correctly there must be at least one Windows Server 2003 license server in your environment. Although Terminal Server provides a licensing grace period during which no license server is required, after the grace period ends, clients must obtain a valid Windows Server 2003 Terminal Server CAL to log onto a Windows Server 2003 terminal server. This includes Windows Server 2003 clients. Currently, all Terminal Server CALs that are available for purchase are Windows Server 2003 Terminal Server CALs.

If you are running Windows 2000 terminal servers with Windows Server 2003 license servers, the Windows 2000 terminal server can use can use Windows Server 2003 Per Device CALs or Windows 2000 CALs. You can downgrade these CALs by calling the Microsoft Clearinghouse and requesting that the CALs be downgraded to Windows 2000. For more information, see Windows Server 2003 Downgrade Rights (http://go.microsoft.com/fwlink/?LinkId=48210) on the Microsoft Web site and Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkID=48885).

Note that a Windows 2000 license server cannot install or issue Windows Server 2003 Terminal Server CALs.

**Terminal Server Licensing Grace Period Increased in Windows Server 2003**
In Windows 2000, the Terminal Server Licensing grace period was 90 days. In Windows Server 2003, the licensing grace period has increased to 120 days. The licensing grace period begins the first time a terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent CAL, or after 120 days, whichever comes first.

**New Group Policy Settings for Windows Server 2003 Service Pack 1**
Windows Server 2003 SP1 includes the following new Group Policy settings for Terminal Services Licensing:

- **Set the Terminal Server licensing mode**

  This new Group Policy setting determines the type of Terminal Server CAL a device or user requires to connect to a Terminal Server.

  When this setting is enabled, you can choose one of the following two licensing modes:

  - Per User. Each user connecting to the terminal server requires a Per User Terminal Server CAL.
  - Per Device. Each device connecting to the terminal server requires a Per Device Terminal Server CAL.

  If you enable this policy setting, the licensing mode that you specify overrides the licensing mode specified during Setup, or in Terminal Services Configuration (TSCC.msc).

  If you disable or do not configure this policy setting, Terminal Services uses the licensing mode specified during Setup or in Terminal Services Configuration.

- **Use the specified Terminal Server license servers**

  This new Group Policy setting determines whether terminal servers must first attempt to locate Terminal Server license servers that are specified in this policy setting before attempting to locate license servers elsewhere on the network.

  Specifying the name of a preferred licensing server in Group Policy saves time and helps eliminate roadblocks to successful configuration of your terminal servers. With the name of a specific licensing server added to Group Policy, Terminal Services does not need to search the network for a licensing server.

  During the automatic discovery process, terminal servers attempt to contact license servers in the following order:

1. Enterprise license servers or domain license servers that are specified in the**LicenseServers** registry key
2. Enterprise license servers that are specified in Active Directory
3. Domain license servers
4. Servers running Terminal Server Licensing in the same subnet
5. The local license server (if the terminal server is also running the Terminal Server Licensing service)

   If you enable this policy setting, terminal servers attempt to locate license servers that are specified in this setting, before following the automatic license server discovery process.

   If you disable or do not configure this policy setting, terminal servers follow the automatic license server discovery process.

| **Note** |
| --- |
| You are also prompted to specify a preferred license server during Setup, or you can specify the license server in Terminal Services Configuration, if the terminal server is running Windows Server 2003 SP1.<br><br>If you are upgrading a terminal server from Windows 2000 to Windows Server 2003, the compatibility report will now display a warning if a Windows Server 2003 license server is not located. |

**Show ToolTips for Licensing Problems on Terminal Server**
This new Group Policy setting allows you, after successfully logging on to a terminal server as an administrator, to display ToolTips that show any licensing problems with the terminal server, and also to display the expiration date of the terminal server's licensing grace period. If this Group Policy setting is not configured, ToolTips display is defined by registry settings.

Using ToolTips to view Terminal Server license statistics at a glance speeds administration tasks. By configuring Group Policy to show ToolTips for Terminal Server licenses, you do not need to open the Properties dialog box for specific licenses to view status and expiration information.

Allowing administrators to configure a global Terminal Server licensing mode makes it possible for them to implement unified license policies regardless of the configuration of Terminal Services client computers. With the new Group Policy setting, differences in configuration between terminal servers and clients can be resolved by defining a global policy that overrides other settings.

# Understanding Troubleshooting Considerations for Specific Terminal Server Licensing Environments

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To ensure that license server discovery works as expected, review the following configuration requirements and considerations, and verify that the requirements have been met for your environment.

Windows 2000 domain: Windows 2000 terminal servers and Windows 2000 license servers

Windows 2000 domain: Windows 2000 terminal servers and Windows Server 2003 license servers

Windows Server 2003 domain: Windows 2000 terminal servers and Windows 2000 license servers

Windows Server 2003 domain: Windows 2000 terminal servers and Windows Server 2003 license servers

Windows Server 2003 domain: Windows Server 2003 terminal servers and Windows Server 2003 license servers

Down level Windows clients, thin clients, and Macintosh clients

**Important**

The following two configurations will not function because a Windows 2000 license server cannot issue CALs to terminal servers running Windows Server 2003 operating systems:

- Windows 2000 domain: Windows Server 2003 terminal servers and Windows 2000 license servers
- Windows Server 2003 domain: Windows Server 2003 terminal servers and Windows 2000 license servers

Therefore, if you have an environment that includes Windows Server 2003 terminal servers, you must install a Windows Server 2003 license server. For more information, see Verify that Terminal Server Licensing is correctly installed on the license server.

**Windows 2000 domain: Windows 2000 terminal servers and Windows 2000 license servers**

- In this environment, you must install Terminal Server Licensing on a domain controller. Although you do not need to explicitly specify a license server, it is recommended that you do so. For more information, see Set a preferred license server.
- If the license server and the terminal server are not already running Windows 2000 Service Pack 4 (SP4), install this service pack on both computers. To download SP4, see Windows 2000 Service Packs (http://go.microsoft.com/fwlink/?LinkId=48213), on the Microsoft Web site.
- Also on the license server, apply the hotfix described in article 837321, You cannot connect to a Windows 2000 Terminal server if your Windows Server 2003 Terminal Services license is expired (http://go.microsoft.com/fwlink/?LinkId=48215), on the Microsoft Web site.
- On the terminal server, apply the hotfix described in article 829395, The Terminal server does not automatically shut down when the TSShutdn.exe command is invoked in Windows 2000 (http://go.microsoft.com/fwlink/?LinkId=48216), on the Microsoft Web site.
- Terminal Server licenses that are issued by license servers running Windows 2000 Service Pack 3 or later have an expiration date. Licenses are either reissued or returned to the Terminal Server CAL pool at a random interval between 52 and 89 days. Note that this setting cannot be configured, and that licenses are renewed seven days before they expire. Any Terminal Server licenses that are issued by license servers running earlier versions of Windows 2000 do not have an expiration date. To return these licenses to the CAL pool, you must remove the license certificate from the clients, and then call the Microsoft Clearinghouse to recover the license. To remove the license certificates from the client, delete the following registry key: **HKEY_LOCAL_MACHINE \Software\Microsoft\MSLicensing**. Licenses for clients that do not have a registry are stored in the registry of the terminal server.

| ⚠Caution |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**Windows 2000 domain: Windows 2000 terminal servers and Windows Server 2003 license servers**

- You can install Terminal Server Licensing on a Windows Server 2003 domain controller or a Windows Server 2003 member server. If you install Terminal Server Licensing on a domain controller, you do not need to explicitly specify a license server for the terminal server to discover. However, it is recommended that you do this to help ensure that license server discovery works as expected. For information about how to do this, see Set a preferred license server.

- If the terminal server is not already running Windows 2000 SP4, install this service pack. To download SP4, see Windows 2000 Service Packs (http://go.microsoft.com/fwlink/?LinkId=48213), on the Microsoft Web site.

- If you have thin clients in this environment, on the Windows Server 2003 license server, apply the hotfix described in article 88450, Event ID 1004 is logged when a thin client tries to obtain a license (http://go.microsoft.com/fwlink/?LinkId=48217) on the Microsoft Web site.

- If you have clients running Windows 2000 or Windows XP Professional in this environment, the license server issues a "free" Terminal Services CAL from its built-in pool to enable these clients to access the Windows 2000 terminal server. The use of this license is permitted under the Windows 2000 EULA. To issue these licenses, however, the license server must first be activated.

  **Note**

- For Windows Server 2003 terminal servers, there is no built-in pool of free licenses. You must install the appropriate number of CALs on the license server to enable all clients to connect to these terminal servers.

**Windows 2003 domain: Windows 2000 terminal servers and Windows 2000 license servers**

- If the license server and the terminal server are not already running Windows 2000 SP4, install this service pack on both computers. To download SP4, see Windows 2000 Service Packs (http://go.microsoft.com/fwlink/?LinkId=48213), on the Microsoft Web site.

- On the Windows 2000 license server, apply the hotfix described in article 837321, You cannot connect to a Windows 2000 Terminal server if your Windows Server 2003 Terminal Services license is expired (http://go.microsoft.com/fwlink/?LinkId=48215), on the Microsoft Web site.

- On the Windows 2000 terminal server, apply the hotfix described in article 829395, The Terminal server does not automatically shut down when the TSShutdn.exe command is invoked in Windows 2000 (http://go.microsoft.com/fwlink/?LinkId=48216), on the Microsoft Web site.

**Windows 2003 domain: Windows 2000 terminal servers and Windows Server 2003 license servers**

- If the terminal server is not already running Windows 2000 SP4, install this service pack. To download SP4, see Windows 2000 Service Packs (http://go.microsoft.com/fwlink/?LinkId=48213), on the Microsoft Web site.

- If you have thin clients in this environment, on the Windows Server 2003 license server, apply the hotfix described in article 88450, Event ID 1004 is logged when a thin client tries to obtain a license (http://go.microsoft.com/fwlink/?LinkId=48217) on the Microsoft Web site.

- If you have clients running Windows 2000 or Windows XP Professional in this environment, the license server issues a "free" Terminal Services CAL from its built-in pool to enable these clients to access the Windows 2000 terminal server. The use of this license is permitted under the Windows 2000 EULA. To issue these licenses, however, the license server must first be activated.

**Note**

- For Windows Server 2003 terminal servers, there is no built-in pool of free licenses. You must install the appropriate number of CALs on the license server to enable all clients to connect to these terminal servers.

**Windows 2003 domain: Windows Server 2003 terminal servers and Windows Server 2003 license servers**

- You can install Terminal Server Licensing on a Windows Server 2003 domain controller or a Windows Server 2003 member server. If you install Terminal Server Licensing on a domain controller, you do not need to explicitly specify a license server for the terminal server to discover. However, it is recommended that you do this to help ensure that license server discovery works as expected. For information about how to do this, see Set a preferred license server.

- If you have thin clients in this environment, on the Windows Server 2003 license server, apply the hotfix described in article 88450, Event ID 1004 is logged when a thin client tries to obtain a license (http://go.microsoft.com/fwlink/?LinkId=48217) on the Microsoft Web site.

**Down level Windows clients, thin clients, and Macintosh clients**
In addition to supporting clients that run Windows Server 2003, Windows 2000, and Windows XP Professional, Windows Server 2003 Terminal Server and Windows 2000 Terminal Server support clients running the following platforms:

- Windows XP Home Edition
- Windows XP Embedded Edition
- Windows Millennium Edition
- Windows 98
- Windows 95
- Windows NT 4.0
- Windows-based Terminal devices (embedded)
- Macintosh

To connect to a terminal server, these clients must have Terminal Server CALs available. Therefore, you should Verify that the license server has a sufficient number of CALs for clients.

# Fixing Client Connectivity Problems

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides a description of client connectivity problems that you might experience when Terminal Server is enabled.

When troubleshooting client connectivity problems, it is most useful to first determine whether there are server configuration issues or problems with network connectivity. Keep in mind that all communication during the Terminal Server Licensing process occurs between the client and the terminal server, and between the terminal server and the license server. The Terminal Server client never communicates directly with the license server. Although Terminal Server Licensing issues often result in error messages appearing on clients, these issues are most frequently caused by problems with the Terminal Server license server or the terminal server.

From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- The terminal server cannot locate the license server.
- The remote computer disconnected the session because of an error in the licensing protocol.
- The remote session was disconnected because there were network problems during the license protocol.
- A licensing error occurred while the client was attempting to connect. (Licensing timed out).
- The remote session was disconnected because there are no Terminal Server license servers available to provide a license.
- The remote session was disconnected because there are no Terminal Server client access licenses available for this computer.
- The remote session was disconnected because the local computer's client access license could not be upgraded or renewed.
- Because of a security error, the client could not connect to the terminal server.
- Event ID 26: Your terminal services temporary client license will expire in N days.
- Event ID 1011: The terminal services client has been disconnected because its temporary license has expired.
- Event ID 1026: The terminal server could not locate a license server in the domain.
- Event ID 1028: The terminal server client has been disconnected because its license could not be renewed.

# The terminal server cannot locate the license server

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process of locating a license server is called license server discovery. If you open Terminal Server Licensing and are prompted for the name of the license server, license server discovery is not working.

**Solution**

1. Review configuration requirements and considerations for implementing Terminal Server Licensing in different environments.
2. Verify that the client, the terminal server, and the license server can communicate.
3. Set a preferred license server.
4. Attempt again to verify whether the terminal server can discover the license server.

If the terminal server still cannot discover the license server, then perform the following steps.

1. Verify RDP-tcp connection settings.
2. Verify the license server role.
3. Verify that the Remote Registry service is started on the license server and that the startup type is set to Automatic.
4. Verify permissions on the Domain Controllers organizational unit.
5. Verify the value of the RestrictAnonymous registry key.
6. Verify that the administrative shares are shared on the license server.
7. Verify that File and Printer Sharing is bound to the network adapter and that the internal network adapter is at the top of the binding order.
8. Verify that the Terminal Server Licensing service is started on the license server and that the startup type is set to Automatic.
9. Verify that Terminal Server Licensing is correctly installed on the license server.
10. Verify that the license server is activated.
11. Remove the MSLicensing registry key on the client and verify permissions on the rebuilt key.

**Review Configuration Requirements and Considerations for Implementing Terminal Server Licensing in Different Environments**
To ensure that license server discovery works as expected, you must configure the terminal servers and license servers correctly. Keep in mind that configuration requirements vary by environment. For information about environment-specific configuration requirements and considerations, see Understanding Troubleshooting Considerations for Specific Terminal Server Licensing Environments.

**Verify that the Client, the Terminal Server, and the License Server Can Communicate**
If the terminal server and license server are installed on separate computers, ensure that there is no firewall on both servers or between them that blocks necessary ports. Terminal Server Licensing uses Remote Procedure Call (RPC) over port 135, and a dynamically assigned port above 1024. To enhance security, you can control which ports RPC is using so that your firewall router can be configured to forward traffic only to these Transmission Control Protocol (TCP) ports. For information, see How to configure RPC dynamic port allocation to work with firewalls (http://go.microsoft.com/fwlink/?LinkId=48218) on the Microsoft Web site.

Additionally, to ensure that clients can establish remote connections to terminal servers, verify that TCP port 3389 is open and listening on the terminal server and on the client. To do this, you can use commands such as Telnet (http://go.microsoft.com/fwlink/?LinkID=48891), Netstat (http://go.microsoft.com/fwlink/?LinkID=48892), or Portqry.

**To verify that the client, the terminal server, and the license server can communicate**

1. At the command prompt, run the **ping** command from each computer to each computer by using:
   - Fully qualified domain name (for example, Server1.contoso.com)
   - NetBIOS name (Server1)
   - The IP address (for example, 206.73.118.1)
2. If any of the **ping** commands fails, verify the DNS configuration on the network.

**Set a Preferred Licensing Server**
By default, terminal servers communicate with license servers in the same domain. This behavior might force them to use wide area network (WAN) connections to find a license server, even if there is a license server from another domain located within the same site. By explicitly specifying a license server, you enable the crossing of domains for license tracking and accounting purposes. For information about how to do this, see Set a preferred license server.

**Verify Whether the Terminal Server Can Discover the License Server**
To verify whether the terminal server can discover the license server, you can use Terminal Server Licensing, as described in the following procedures.

Using Terminal Server Licensing (on Windows Server 2003)

Using Terminal Server Licensing (on Windows 2000)

**Using Terminal Server Licensing (on Windows Server 2003)**
To connect to a license server, a terminal server must have the **Access this computer from the network** logon right for the target license server. Additionally, if the License Server Security Group setting in Group Policy is enabled, the terminal server computer account must be a member of this group. Enabling this Group Policy setting and applying it to your license server creates a local group called Terminal Server Computers. The license server issues licenses only to the terminal servers in this group. You must add the terminal servers for which you need to provide licenses to this group for each license server.

**To verify whether the terminal server can discover the license server (on Windows Server 2003)**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, click **Control Pane**l, double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**.

   When Terminal Server Licensing opens, it discovers all license servers in your workgroup or domain (and searches Active Directory for Enterprise License Servers) and it displays the names of those servers in the console tree. If no license servers are automatically displayed in the console tree, you are prompted to specify the name of a server as noted in the following steps.
2. On the **Action** menu, click **Connect**.
3. In **Server**, type the name of the license server to which you want to connect, and then click **OK**.

**Using Terminal Server Licensing (on Windows 2000)**
In addition to supporting clients that run Windows Server 2003, Windows 2000, and Windows XP Professional, Windows 2000 Terminal Server supports clients running the following platforms: Windows XP Home Edition, Windows XP Embedded Edition, Windows Millennium Edition, Windows 98, Windows 95, Windows NT 4.0, Windows-based Terminal devices, and Macintosh. These clients must have Terminal Server CALs available to connect to a Windows 2000 terminal server.

If you have clients running Windows 2000 or Windows XP Professional, the license server issues a "free" Terminal Services CAL from its built-in pool to enable these clients to access the Windows 2000 terminal server. The use of this license is permitted under the Windows 2000 Server EULA. To issue these licenses, however, the license server must first be activated.

For Windows Server 2003 terminal servers, there is no built-in pool of free licenses. You must install the appropriate number of CALs on the license server to enable all clients to connect to these terminal servers.

**To verify whether the terminal server can discover the license server (on Windows 2000)**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Terminal Services Licensing**.

2. Navigate to **All Servers**, and then click the name of the terminal server to which you want to connect.

   A list of installed CALs appears in the details pane. Verify that **Windows 2000 Terminal Services Client Access License** appears in the list.

3. If it does not appear, purchase and install the appropriate number of CALs to support your clients.

**Verify RDP-tcp Connection Settings**

Use Group Policy or Terminal Services Configuration to verify that RDP-tcp connection settings are configured as follows:

- The RDP-tcp connection is enabled.

- The **Maximum connections** setting on the **Network Adapter** tab of the **RDP-tcp Properties** dialog box is set to **Unlimited**.

- Verify whether the level encryption being used for the RDP-TCP connection is set to **High**. If the encryption level is set to **High**, keep in mind that this level encrypts data sent from client to server and from server to client by using strong 128-bit encryption. Clients that do not support this level of encryption cannot connect to a terminal server.

- Verify that the user attempting to log on to the terminal server over the RDP-tcp connection has the appropriate rights and permissions. Users must either be members of the Remote Desktop Users group, or you must assign the equivalent rights and permissions manually (users must have the Allow log on through Terminal Services (http://go.microsoft.com/fwlink/?LinkID=48894) right and User Access permissions). For more information, see Enabling users to connect remotely to the server (http://go.microsoft.com/fwlink/?LinkID=48887). If users are running the Citrix ICA client, they must have the appropriate permissions for the ICA-TCP connection.

For instructions on how to verify and modify RDP-tcp connection settings, see Configure Terminal Services Connections (http://go.microsoft.com/fwlink/?LinkId=48219).

**Verify the License Server Role**

**Caution** Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**To To verify the license server role**

1. On the license server, open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

2. Locate, and then click, the following key in the registry:

   **HKLM\System\CurrentControlSet\Services\TermServLicensing\Parameters\Role**

   Where

   **0** = Domain license server

   **1** = Enterprise license server

**Verify that the Remote Registry Service is Started on the License Server and that the Startup Type is Set to Automatic**

- You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

If you enable or disable a service and you encounter a problem starting the computer, you might be able to start the computer in Safe Mode. Then you can change the service configuration or restore the default configuration.

**Important** If you stop, start, or restart a service, any dependent services are also affected.

**To verify that the Remote Registry service is enabled on the license server and that the startup type is set to**

**Automatic**

1. On the license server, open **Services**. To open Services, click **Start**, click Control Panel, click **Performance and Maintenance**, click **Administrative Tools**, and then double-click **Services**.
2. In the list of services, right-click **Remote Registry**, and then click **Properties**.
3. On the **General** tab, verify that the service status is Started and that the startup type is Automatic.
4. If the status is not Started, click **Start** to start the service. If the startup type is not Automatic, click **Automatic** to ensure that the service will restart when the license server is restarted.

**Verify Permissions on the Domain Controllers OU**

For license server discovery to work in a Windows Server 2003 domain with Windows 2000 terminal servers and Windows 2000 license servers, the following permissions must exist on the Domain Controllers OU.

- System = Full Control
- Authenticated Users = Read
- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as (http://go.microsoft.com/fwlink/?LinkID=48886) to perform this procedure.

**To verify permissions on the Domain Controllers OU**

1. Open Active Directory Users and Computers. To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
2. On the **View** menu, select **Advanced Features**.
3. Right-click the **Domain Controllers** OU, and then click **Properties**.
4. On the **Security** tab, click **Advanced** to view the permission entries that exist for the object.
5. In the list, click **Authenticated Users**.
6. In the **Permissions** box, make sure that the **Allow** check box is selected for **Read.**
7. In the list, click **System**.
8. In the **Permissions** box, make sure that the **Allow** check box is selected for Full Control.

**Verify the Value of the RestrictAnonymous Registry Key**

If you are using Windows 2000 servers, verify that the **RestrictAnonymous** registry key value on the license server is set to**0**or **1**. In this scenario, if this registry key value is set to**2**, then the license server cannot issue CALS. When the**RestrictAnonymous** registry value is set to **2**, the access token built for non-authenticated users does not include the Everyone group, and because of this, the access token no longer has access to resources that grant permissions to the Everyone group.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**To verify the value of the RestrictAnonymous registry key**

1. On the license server, open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Locate, and then click, the following key in the registry:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\RestrictAnonymous**

3. Verify the value of the **RestrictAnonymous**registry key.

The three valid registry key values for **RestrictAnonymous** are equivalent to the following security settings in Local Security Policy:

**0** = None (rely on default permissions)

**1** = Do not allow enumeration of SAM accounts and shares

**2** = No access without explicit anonymous permissions

If the license server is a member of an Active Directory domain and a conflicting security setting is configured for the license server in Group Policy, the Group Policy setting overrides the local security setting.

4. If the value of the registry key is set to **2**, then set it to **1** or **0**.

   Also, verify that the equivalent value (**No access without explicit anonymous permissions**) is not enabled by Group Policy. To do this, open Local Security Policy, navigate to **Security Settings\Local Policies\Security Options\Additional restrictions for anonymous connections**, and then verify whether **No access without explicit anonymous permissions** is enabled. If it is, disable it. To open Local Security Policy, click **Start**, point to **Settings**, Control Panel, click **Administrative Tools**, and then click **Local Security Policy**.

5. Restart the server.

   In Windows Server 2003, if you need to prohibit anonymous users from being granted the same access that is granted to members of the Everyone group, use the new **Everyone Network access: Let Everyone permissions apply to anonymous users** setting in Local Security Policies.

#### Verify that the Administrative Shares are Shared on the License Server

**To verify that the administrative shares are shared on the license server**

1. At the command prompt, type **net share**, and then press ENTER.

2. Verify that the default shares, such as C$, Admin$, and IPC$, are returned.

   If you perform this procedure and the default administrative shares are not returned, verify that the Server service is started, that File and Print Sharing is enabled, and that the license server does not have a virus. Also, see HOW TO: Restore Administrative Shares That Have Been Deleted (http://go.microsoft.com/fwlink/?LinkId=48223) on the Microsoft Web site.

#### Verify that File and Printer Sharing is Bound to the Network Adapter and that the Internal Network Adapter is at the Top of the Binding Order

**To verify that File and Printer Sharing is bound to the internal network adapter and that the internal network adapter is at the top of the binding order**

1. Open Network and Dial-up Connections. To open Network and Dial-up Connections, click **Start**, point to **Settings**, and then click **Network and Dial-up Connections**.

2. Click the connection that you want to verify, and on the **Advanced**menu, click **Advanced Settings**.

3. On the **Adapters and Bindings** tab, in **Bindings for *AdapterName***, verify that the **File and Printer Sharing** check box is selected. If the check box is not selected, select it.

   To modify the protocol bindings order, you must be logged on as a member of the Administrators group.

4. Also in **Bindings for *AdapterName***, verify that the internal adapter is at the top of the list. If it is not at the top of the list, click the up and down arrows to move the internal adapter to the top of the list.

#### Verify that the Terminal Server Licensing Service is Started and that the Startup Type is Set to Automatic

For step-by-step instructions, see Verify that the Terminal Server Licensing service is started and the startup type is set to Automatic.

#### Verify that Terminal Server Licensing is Correctly Installed on the License Server

For step-by-step instructions, see Verify that Terminal Server Licensing is correctly installed on the license server.

#### Verify that the License Server is Activated

A license server that has been installed but not activated only issues temporary licenses. These temporary licenses allow clients to connect to the terminal server for 90 days.

**To verify that a license server is activated**

1. Open Terminal Server Licensing.
2. In the console tree, click the license server for which you want to verify the activation status.
3. In the details pane, verify that activation status is Activated.
4. If the status is Not Activated, right-click the server name, click **Activate Server**, and then follow the instructions in the wizard.

**Remove the MSLicensing Registry Key on the Client and Verify Permissions on the Rebuilt Key**

If all of the previous troubleshooting procedures fail, create a backup of the **MSLicensing**registry key and its subkeys on the client, and then remove the original key and subkeys.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**To remove the MSLicensing registry key on the client and verify permissions on the rebuilt key**

1. On the client, open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Locate, and then click, the following key in the registry:

   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSLicensing**
3. On the **Registry** menu, click **Export Registry File**.
4. In the **File name** box, type **mslicensingbackup**, and then click **Save**.
5. If you need to restore this registry key in the future, double-click **mslicensingbackup.reg**.
6. On the **Edit** menu, click **Delete**, and then click **Yes** to confirm the deletion of the **MSLicensing**registry subkey.
7. Close Registry Editor, and then restart the computer (when the client is restarted, the missing registry key is rebuilt).
8. On the client, open Registry Editor.
9. Locate, and then click, the following key in the registry:

   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSLicensing**
10. On the **Edit** menu, click **Permissions**. Users must have at least Read permissions.

# The remote computer disconnected the session because of an error in the licensing protocol

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**

The terminal server might not be able to locate the license server.

**Solution**

To resolve this issue, perform the troubleshooting procedures described in The terminal server cannot locate the license server.

# The remote session was disconnected because there were network problems during the license protocol

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
The terminal server might not be able to locate the license server.

**Solution**
Perform the troubleshooting procedures described in The terminal server cannot locate the license server.

# A licensing error occurred while the client was attempting to connect. (Licensing timed out)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
If you are using Internet Protocol security (IPsec) to help protect traffic over TCP between clients and terminal servers, then packet fragmentation might occur. As a result, some packets might not reach their destinations, and client connections to terminal servers might fail.

**Solution**
Configure IPsec to help protect traffic over UDP rather than over TCP, as described in Define IPSec Policies (http://go.microsoft.com/fwlink/?LinkID=48889).

**Cause**
If there is a firewall between the clients and terminal servers, packets might be fragmented due to the size of the maximum transmission unit (MTU) that can be transmitted over the network. As a result, the packet fragments might be dropped or received out of order.

**Solution**
Verify that the maximum transmission unit (MTU) that can be transmitted over the network is sufficient by running the **ping** command with the **-f** (don't fragment) parameter as follows:

**ping -l 1472 -f <Destination_IP_Address>**

If the MTU is not sufficient, you can change the MTU on the router or you can configure the computers to use a smaller MTU by modifying the registry. You can configure the computers to use a smaller MTU by creating a new MTU registry key entry and setting the value to 576 (or a larger value if appropriate, if the MTU is larger) or by creating a new **EnablePTMUDiscovery** registry key entry and setting the value to **0**. However, note that Windows Server 2003 And Windows 2000 TCP/IP use PMTU detection by default and query the NIC driver to determine the local MTU that is supported. Therefore, modifying the MTU parameter is generally not necessary and might result in reduced performance.

| ⚠Caution |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**Create a new MTU registry key and set it to 576 (or a larger value)**

1. Open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Locate, and then click, the following key in the registry:
   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\<Interface_GUID>**
3. Create a new DWORD registry entry called **MTU**.
4. Set the value to **576** (or a larger value, if appropriate).

**Create an EnablePMTUDiscovery registry key and set it to 0**

1. Open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Locate, and then click, the following key in the registry:
   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters**
3. Create a new DWORD registry entry called **EnablePMTUDiscovery**.
4. Set the value to **0**.

# The remote session was disconnected because there are no Terminal Server license servers available to provide a license

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
The terminal server might not be able to locate the license server.

**Solution**
Perform the troubleshooting procedures described in The terminal server cannot locate the license server.

**Cause**
The client might have exceeded its Terminal Server Licensing grace period (http://go.microsoft.com/fwlink/?LinkID=48895), and a license server has not yet been installed or activated to issue the client a client access license (CAL). After a temporary license expires, a client can only connect to a terminal server if the license server can issue a permanent CAL.

**Solution**
Install Terminal Server Licensing, activate the license server, then purchase and install a sufficient number of CALs to support the clients in your organization. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886) to perform this procedure.

| **Note** |
| --- |
| If you are installing an Enterprise License server and your environment includes multiple trusted Active Directory sites, each of which hosts a terminal server, you can use a script to publish a single license server across all of these sites. Doing this can greatly simplify deployment because it eliminates the need to deploy a separate license server in each site, and it eliminates the need to distribute CALs to license servers across multiple sites. For a sample script that you can use to accomplish this task, see the Microsoft Script Center (http://go.microsoft.com/fwlink/?LinkID=48230) on the Microsoft Web site. Note that this sample script is not supported under any Microsoft standard support program or service. |

**To install Terminal Server Licensing**

1. Open Add/Remove Programs in **Control Panel**.
2. Double-click **Add/Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. In Terminal Server Licensing Setup, do one of the following:
   - Click **Your entire enterprise** to choose an Enterprise license server. You can only install an Enterprise license server on a domain controller or a member server in a domain (if the member server is running Windows Server 2003), not on a stand-alone server. This role is published to Active Directory. If you are installing an Enterprise License Server, you must meet one of the following conditions: 1) You must be a member of the Enterprise Admins group in Active Directory or 2) You must have the right to create the TS-Enterprise-License-Server object in the same site in which the license server will reside in Active Directory.
   - Click **Your domain or workgroup** to choose a Domain license server. If you select this option, the computer on which you install the license server does not have to be a domain controller, if the computer is running Windows Server 2003. However, if the computer is running Windows 2000, the computer on which you install the license server must be a domain controller or a stand-alone server in a workgroup. If the computer on which you plan to install the license server is in a workgroup, it can only be automatically discovered by terminal servers on the same subnet. If you are installing a Domain license server, you must be a member of the Administrators group on the local computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
5. Provide the file system location where the license server database should be installed on the Terminal Server license server, click **Next**, and then click **Finish**.

   The default location for the license server database is systemroot\System32\LServer.

You can activate a Terminal Server license server by using one of three methods: Automatic, Web Browser, or Telephone. It is recommended that you activate a license server automatically, as described in the following procedure.

| **Note** |
| --- |
| To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886). |

**To activate a license server automatically.**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, point to **Control Panel**, **Administrative Tools**, and then click **Terminal Server Licensing**.
2. In the console tree, right-click the license server that you want to activate, click **Activate Server**, and then follow the instructions in the wizard.
3. In the **Activation** method list, select **Automatic connection (recommended)**, and then click **Next**.
4. Follow the instructions in the wizard.
5. If necessary, at a later time you can change any license server properties, such as activation method and company information, that you set during the activation process.

Then, Purchase CALs as required and Install CALs automatically.

# The remote session was disconnected because there are no Terminal Server client access licenses available for this computer

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
The license server might not have any remaining Per Device CALs to issue.

**Solution**

1. Use Terminal Server Licensing to Verify that the license server has a sufficient number of CALs for clients requiring remote connections.

2. If the license server does not have a sufficient number of permanent CALs, Purchase CALs as needed, and then Install CALs automatically, as needed.

**Cause**
In the following two cases, either no licenses, or temporary licenses that cannot be upgraded, are issued from the available pool of client access licenses (CALs) on a license server:

1. If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per User, then no licenses are issued from the available pool of licenses, because Per User licensing is not managed.

2. If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per Device but the license server is configured to issue only Per User CALs, then only temporary Per Device licenses are issued from the available pool of licenses. These licenses cannot be upgraded. When the temporary licenses are close to expiration, the following message appears on the terminal server: "Your terminal services temporary client license will expire in N days. Please contact your system administrator to get a permanent license." In this case, Event ID 26, Source: Application Pop-up appears in the application event log.

   Similarly, Terminal Server clients that run Windows XP or Windows 2000 might receive the following error message, if the client cannot connect to the terminal server for this reason: "The remote session was disconnected because the local computer client access license could not be upgraded or renewed." In this case, Event ID 1011, Source:TermService appears in the application event log on the terminal server, with the following message: "The terminal services client has been disconnected because its temporary license has expired."

**Solution**
Do one of the following, based on the requirements of your environment and license agreement:

- Call the Microsoft Clearinghouse and reactivate the license server in the correct mode.

- Change the Terminal Server Licensing mode.

**Cause**
The terminal server might not be able to locate the license server.

**Solution**
To resolve this issue, perform the troubleshooting procedures described in The terminal server cannot locate the license server.

# The remote session was disconnected because the local computer's client access license could not be upgraded or renewed

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After a temporary license expires, a client can only connect to a terminal server if the license server can issue a permanent CAL, or if the terminal server is still within its licensing grace period.

**Note**

If you are running Terminal Server Licensing in a Windows 2000 Service Pack 3 (SP3) or later environment, the client must connect to a terminal server and authenticate twice to obtain a permanent license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and is not yet issued, the terminal server issues the client a permanent, Per Device CAL.

**Cause**

If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per Device but the license server is configured to issue only Per User CALs, then only temporary Per Device licenses are issued from the available pool of licenses. These licenses cannot be upgraded. When the temporary licenses are close to expiration, the following message appears on the terminal server: "Your terminal services temporary client license will expire in N days. Please contact your system administrator to get a permanent license." In this case, Event ID 26, Source: Application Pop-up appears in the application event log.

Similarly, Terminal Server clients that run Windows XP or Windows 2000 might receive the following error message, if the client cannot connect to the terminal server for this reason: "The remote session was disconnected because the local computer client access license could not be upgraded or renewed." In this case, Event ID 1011, Source:TermService appears in the application event log on the terminal server, with the following message: "The terminal services client has been disconnected because its temporary license has expired."

**Solution**

Do one of the following, based on the requirements of your environment and license agreement:

- Call the Microsoft Clearinghouse and reactivate the license server in the correct mode. For more information, see Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkId=48885).

- Change the Terminal Server Licensing mode.

**Cause**

If you are running Terminal Server Licensing in a Windows 2000 (SP3) or earlier environment, expired temporary licenses might not be upgraded to permanent CALs.

**Solution**

1. On the terminal server and license server, install Service Pack 4 (SP4). To download SP4, see Windows 2000 Service Packs (http://go.microsoft.com/fwlink/?LinkID=48213) on the Microsoft Web site.

2. Use Terminal Server Licensing to Verify that the license server has a sufficient number of CALs for clients requiring remote connections.

3. If the license server does not have a sufficient number of permanent CALs, Purchase CALs, and then Install CALs automatically.

# Because of a security error, the client could not connect to the terminal server

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**

If you upgraded a Windows NT domain to Windows 2000 or Windows Server 2003, then the certificate on the terminal server might be corrupted. As a result, Windows 2000 Terminal Services clients might be repeatedly denied access to the terminal server.

**Solution**

On each terminal server, back up and then delete the **Certificate**, **X509 Certificate**, and**X509 Certificate ID** registry subkeys. Then, on each client, back up and then delete the **MSLicensing** registry key.

**Back Up and Then Delete the Certificate, X509 Certificate, and X509 Certificate ID Registry Subkeys**

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**To back up and then delete the Certificate, X509 Certificate, and X509 Certificate ID registry subkeys**

1. On each terminal server, create a backup of the registry.
2. Navigate to the following registry subkey: **HKLM\SYSTEM\CurrentControlSet\Services\TermServices\Parameters.**
3. On the **Registry** menu, click **Export Registry File**.
4. In the **File name** box, type **exported-parameters**, and then click **Save**.
5. If you need to restore this registry subkey in the future, double-click **exported-parameters.reg**.
6. Under the **Parameters r**egistry subkey, right-click each of the following values:
   - **Certificate**
   - **X509 Certificate**
   - **X509 Certificate ID**
7. Click **Delete**, and then click **Yes** to confirm the deletion.
8. Close Registry Editor, and then restart each terminal server.

**Back Up and Then Delete the MSLicensing Registry Key**

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**To back up and then delete the MSLicensing registry key**

1. On the client, navigate to the following registry subkey:
   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSLicensing.**
2. Click **MSLicensing**.
3. On the **Registry** menu, click **Export Registry File**.
4. In the **File name** box, type **mslicensingbackup**, and then click **Save**.

5. If you need to restore this registry key in the future, double-click **mslicensingbackup.reg**.
6. On the **Edit** menu, click **Delete**, and then click **Yes** to confirm the deletion of the **MSLicensing** registry subkey.
7. Close Registry Editor, and then restart the client.

**Solution**

If the client still cannot connect to the terminal server, then perform the following procedures.

1. On the client, back up and then delete the MSLicensing registry key and its subkeys.
2. On each terminal server, back up and then delete the Certificate, X509 Certificate, and X509 Certificate ID registry subkeys.
3. Deactivate the license server.
4. Reactivate the license server by using the Telephone connection method in the Terminal Server License Server Wizard.

**Deactivate the License Server**

You might need to deactivate a Terminal Server license server when the certificate of the server has expired or becomes corrupted, or when the server is being redeployed. You will be prompted to reactivate the Terminal Server license server when its registration has expired. When you deactivate a Terminal Server license server, you will not be able to license additional clients from this server until it is reactivated.

You can deactivate a server by using doing one of the following:

1. Using the Automatic method
2. Using the Telephone method

- You cannot deactivate a Terminal Server license server by using the Web browser connection method.

**Using the Automatic method**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886).

**To deactivate a license server using the Automatic method**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**.
2. In the console tree, right-click the Terminal Server license server you want to deactivate, point to **Advanced**, and then click **Deactivate Server**. The Terminal Server License Server Wizard starts.
3. In **Required Information**, confirm that your name, phone number, company, and country or region are correct, and then click **Next**.
4. Your request to deactivate the Terminal Server license server is sent to the Microsoft Clearinghouse for processing.
5. Click **Finish**.

**Using the telephone method**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886).

**To deactivate a license server using the Telephone method**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, click **Control Panel**, double-click

**Administrative Tools**, and then double-click **Terminal Server Licensing**.

2. In the console tree, right-click the Terminal Server license server you want to deactivate, point to **Advanced**, and then click **Deactivate Server**. The Terminal Server License Server Wizard starts.
3. Select **Telephone**, and then click **Next**.
4. Specify your location, and then click **Next**.
5. Call the telephone number displayed in the wizard, and give the Customer Support representative the Product ID that is displayed below the telephone number.
6. In the boxes in the wizard, type the 35-digit confirmation code provided by the Customer Support representative, then click **Next**.
7. Click **Finish**.

**Reactivate the License Server**

When you activate Terminal Server Licensing by using the Telephone option, Terminal Server Licensing uses a different certificate. You will receive a telephone call from the Microsoft Clearinghouse confirming that the license server has been reactivated. When a license server is reactivated, a record of your licenses is retained. Licenses that were already issued remain valid, and any un-issued licenses are also valid, but must be reissued through the Microsoft Clearinghouse.

You must reactivate a license server when its registration has expired.

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886).

**To reactivate a license server**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**.
2. In the console tree, right-click the Terminal Server license server you want to reactivate, point to **Advanced**, and then click **Reactivate Server**. The Terminal Server License Server Wizard starts.
3. In **Information Needed**, confirm that your name, phone number, and e-mail address are correct.
4. In **Reason**, select the appropriate reason for reactivating the license server, and click **Finish**.
5. Your request to reactivate the license server is sent to the Microsoft Clearinghouse for processing.

**Cause**

Windows XP-based clients might be attempting to connect to a Windows 2000 terminal server in a low-bandwidth network environment, in which client sessions are encrypted. In this case, IP packet fragmentation can cause encrypted frames that are sent by a client to be decrypted incorrectly.

**Solution**

- On the terminal server, install Windows 2000 Service Pack 4 (SP4). To download SP4, see Windows 2000 Service Packs (http://go.microsoft.com/fwlink/?LinkID=28213) on the Microsoft Web site .

**Cause**

The Remote Desktop Protocol (RDP) encryption settings on the terminal server and the client might not be compatible. For example, the terminal server might be running 128-bit encryption with an encryption level set to **High.** When this occurs, "Event ID 50, Source: TermDD" appears in the system event log on the terminal server.

**Solution**

Change the RDP encryption level on the terminal server to **Medium** or **Low** (if the terminal server is running Windows 2000) or to **Client Compatible** or **Low** (if the terminal server is running Windows Server 2003).

You can change the RDP encryption level on the terminal server by doing one of the following:

- Using Terminal Services Configuration (Windows 2000)

- Using Group Policy (Windows Server 2003)
- Using Terminal Services Configuration (Windows Server 2003)

  When you change the encryption level, the new encryption level takes effect the next time a user logs on. If you require multiple levels of encryption on one server, install multiple network adapters and configure each adapter separately.

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886).

**Using Terminal Services Configuration (Windows 2000)**

**To change the RDP encryption level Using Terminal Services Configuration (Windows 2000)**

1. Open Terminal Services Configuration. To open Terminal Services Configuration, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Terminal Services Configuration**.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection that you want to modify, and then click **Properties.**
4. On the **General** tab, in **Encryption level**, click **Low** or **Medium**:
   - The **Low** setting encrypts data sent from the client to the server by using either 40-bit or 56-bit encryption. A Windows 2000 terminal server uses 56-bit encryption when Windows 2000 clients connect to it and 40-bit encryption when earlier versions of the client connect to it.
   - The **Medium** level encrypts data sent from client to server and from server to client by using either 40-bit encryption or 56-bit encryption. A Windows 2000 terminal server uses 56-bit encryption when Windows 2000 clients connect to it and 40-bit encryption when earlier versions of the client connect.

**Using Group Policy (Windows Server 2003)**

Any encryption level settings that you configure in Group Policy override the configuration that you set by using the Terminal Services Configuration tool. Use this procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on as an administrator. Then, you must open Group Policy by using the Active Directory Users and Computers snap-in.

| **Important** |
| --- |
| You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. |

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886).

**To change the RDP encryption level using Group Policy (Windows Server 2003)**

1. Open Group Policy.
2. In Computer Configuration, Administrative Templates, Windows Components, Terminal Services, Encryption and Security, double-click the **Set client connection encryption level setting**, and verify that the setting is set to **Enabled**.
3. In the **Encryption Level** list, click **Client Compatible** or **Low Level**, and then click **OK**.
   - The **Client Compatible** setting encrypts data sent between the client and the server at the maximum key strength supported by the client.
   - The **Low** setting encrypts data sent from the client to the server using 56-bit encryption.

**Using Terminal Services Configuration (Windows Server 2003)**
Any encryption level settings that you configure in Group Policy override the configuration that you set by using the Terminal Services Configuration tool.

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886).

**To change the RDP encryption level using Terminal Services Configuration (Windows Server 2003)**

1. Open Terminal Services Configuration. To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection you want to modify, and then click **Properties**.
4. On the **General** tab, in **Encryption level**, click **Low** or **Client Compatible**.
   - The **Low** setting encrypts data sent from the client to the server using 56-bit encryption.
   - The **Client Compatible** setting encrypts data sent between the client and the server at the maximum key strength supported by the client.

**See Also**
**Other Resources**
Windows XP Clients Cannot Connect to a Windows 2000 Terminal Services Server

# Event ID 26: Your terminal services temporary client license will expire in N days

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per Device but the license server is configured to issue only Per User CALs, then only temporary Per Device licenses are issued from the available pool of licenses. These licenses cannot be upgraded. When the temporary licenses are close to expiration, the following message appears on the terminal server: "Your terminal services temporary client license will expire in N days. Please contact your system administrator to get a permanent license." In this case, Event ID 26, Source: Application Pop-up appears in the application event log.

Similarly, Terminal Server clients that run Windows XP or Windows 2000 might receive the following error message, if the client cannot connect to the terminal server for this reason: "The remote session was disconnected because the local computer client access license could not be upgraded or renewed." In this case, Event ID 1011, Source:TermService appears in the application event log on the terminal server, with the following message: "The terminal services client has been disconnected because its temporary license has expired."

**Solution**
Do one of the following, based on the requirements of your environment and license agreement:

- Call the Microsoft Clearinghouse and reactivate the license server in the correct mode. For more information, see Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkId=48885).

- Change the Terminal Server Licensing mode.

**Cause**
The terminal server might not be able to locate the license server.

**Solution**
To resolve this issue, perform the troubleshooting procedures described in The terminal server cannot locate the license server.

# Event ID 1011: The terminal services client has been disconnected because its temporary license has expired

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per Device but the license server is configured to issue only Per User CALs, then only temporary Per Device licenses are issued from the available pool of licenses. These licenses cannot be upgraded. When the temporary licenses are close to expiration, the following message appears on the terminal server: "Your terminal services temporary client license will expire in N days. Please contact your system administrator to get a permanent license." In this case, Event ID 26, Source: Application Pop-up appears in the application event log.

Similarly, Terminal Server clients that run Windows XP or Windows 2000 might receive the following error message, if the client cannot connect to the terminal server for this reason: "The remote session was disconnected because the local computer client access license could not be upgraded or renewed." In this case, Event ID 1011, Source:TermService appears in the application event log on the terminal server, with the following message: "The terminal services client has been disconnected because its temporary license has expired."

**Solution**
Do one of the following, based on the requirements of your environment and license agreement:

- Call the Microsoft Clearinghouse and reactivate the license server in the correct mode. For more information, see Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkId=48885).

- Change the Terminal Server Licensing mode.

**Cause**
The terminal server might not be able to locate the license server.

**Solution**
To resolve this issue, perform the troubleshooting procedures described in The terminal server cannot locate the license server.

# Event ID 1026: The terminal server could not locate a license server in the domain

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**

The process of locating a license server is called license server discovery. If license server discovery is not functioning correctly, Event ID 1026, Source:TermServ, might appear in the application event log, with the following message: "The terminal server could not locate a license server in the <domain> domain. Confirm that all license servers on the network are registered in WINS/DNS, accepting network requests, and the Terminal Server Licensing Service is running."

**Solution**

Set a preferred license server.

If your terminal server is running Windows Server 2003 with Service Pack 1 (SP1), you can do this by using Group Policy settings or Terminal Services Configuration, or you can use the registry or a Windows Management Instrumentation (WMI) script. If your terminal server is running Windows 2000 or Windows Server 2003 without SP1, you must use the registry or a WMI script. For three scripts that you can use to set preferred license servers, delete preferred license servers, or query preferred license servers, see Preferred License Server WMI Scripts (http://go.microsoft.com/fwlink/?LinkID=26220) in Windows Server 2003 Terminal Server Licensing on the Microsoft Web site.

**Note**

The license servers and the terminal servers must either be in the same domain or in domains that trust each other.

**Solution**

If explicitly specifying a license server does not resolve this issue, perform additional troubleshooting steps, as described in The terminal server cannot locate the license server.

# Event ID 1028: The terminal server client has been disconnected because its license could not be renewed

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are running a Windows Server 2003 terminal server, after you use the Windows Components Wizard in Add or Remove Programs to add or to remove a Windows component, an invalid registry value might be set in the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\TerminalServer\Licensing Core\PolicyAcOn** subkey on the terminal server. This invalid value might may unexpectedly change the Terminal Server Licensing mode and cause the terminal server to obtain only temporary licenses.

**Cause**

This problem might occur because the value of the **PolicyAcOn** registry entry changes from either **2** (the correct registry value for the Per Device licensing mode) or **4** (the correct registry value for the Per User licensing mode) to an invalid value of **0**. This behavior causes the terminal server to attempt to obtain temporary Per Device licenses. In this case, if you view Terminal Services Configuration to verify the licensing mode, the licensing mode appears as Per Device. Therefore, to verify whether this problem is causing the client disconnection, you must verify the value of the **PolicyAcOn** registry entry.

**Solution**

To resolve this problem contact Microsoft Product Support Services to obtain the hotfix described in Knowledge Base article 834651, Terminal Services licensing mode changes from Per User or from Per Device changes from Per User or from Per Device to an invalid value after you add or remove a Windows component, (http://go.microsoft.com/fwlink/?LinkID=48173) on the Microsoft Web site. This hotfix is applied to the terminal server. For information about how to contact Microsoft Product Support Services, see Help and Support (http://go.microsoft.com/fwlink/?LinkID=48149) on the Microsoft Web site.

# Fixing CAL Issuance and Upgrade Problems

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides a description of client access license (CAL) issuance and upgrade problems that you might experience when Terminal Server Licensing is enabled.

From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- The terminal server cannot locate the license server.
- The remote session was disconnected because there are no Terminal Server license servers available to provide a license.
- The remote session was disconnected because there are no Terminal Server client access licenses available for this computer.
- The remote session was disconnected because the local computer's client access license could not be upgraded or renewed.
- No licenses are being issued from the available pool of licenses.
- Event ID 26: Your terminal services temporary client license will expire in N days.
- Event ID 1011: The terminal services client has been disconnected because its temporary license has expired.
- Event ID 1026: The terminal server could not locate a license server in the domain.
- Event ID 1028: The terminal server client has been disconnected because its license could not be renewed.

# The terminal server cannot locate the license server

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The process of locating a license server is called license server discovery. If you open Terminal Server Licensing and are prompted for the name of the license server, license server discovery is not working.

**Solution**

1. Review configuration requirements and considerations for implementing Terminal Server Licensing in different environments.
2. Verify that the client, the terminal server, and the license server can communicate.
3. Set a preferred license server.
4. Attempt again to verify whether the terminal server can discover the license server.

If the terminal server still cannot discover the license server, then perform the following steps.

1. Verify RDP-tcp connection settings.
2. Verify the license server role.
3. Verify that the Remote Registry service is started on the license server and that the startup type is set to Automatic.
4. Verify permissions on the Domain Controllers organizational unit.
5. Verify the value of the RestrictAnonymous registry key.
6. Verify that the administrative shares are shared on the license server.
7. Verify that File and Printer Sharing is bound to the network adapter and that the internal network adapter is at the top of the binding order.
8. Verify that the Terminal Server Licensing service is started on the license server and that the startup type is set to Automatic.
9. Verify that Terminal Server Licensing is correctly installed on the license server.
10. Verify that the license server is activated.
11. Remove the MSLicensing registry key on the client and verify permissions on the rebuilt key.

**Review Configuration Requirements and Considerations for Implementing Terminal Server Licensing in Different Environments**
To ensure that license server discovery works as expected, you must configure the terminal servers and license servers correctly. Keep in mind that configuration requirements vary by environment. For information about environment-specific configuration requirements and considerations, see Understanding Troubleshooting Considerations for Specific Terminal Server Licensing Environments.

**Verify that the Client, the Terminal Server, and the License Server Can Communicate**
If the terminal server and license server are installed on separate computers, ensure that there is no firewall on both servers or between them that blocks necessary ports. Terminal Server Licensing uses Remote Procedure Call (RPC) over port 135, and a dynamically assigned port above 1024. To enhance security, you can control which ports RPC is using so that your firewall router can be configured to forward traffic only to these Transmission Control Protocol (TCP) ports. For information, see How to configure RPC dynamic port allocation to work with firewalls (http://go.microsoft.com/fwlink/?LinkId=48218) on the Microsoft Web site.

Additionally, to ensure that clients can establish remote connections to terminal servers, verify that TCP port 3389 is open and listening on the terminal server and on the client. To do this, you can use commands such as Telnet (http://go.microsoft.com/fwlink/?LinkID=48891), Netstat (http://go.microsoft.com/fwlink/?LinkID=48892), or Portqry.

**To verify that the client, the terminal server, and the license server can communicate**

1. At the command prompt, run the **ping** command from each computer to each computer by using:
   - Fully qualified domain name (for example, Server1.contoso.com)
   - NetBIOS name (Server1)
   - The IP address (for example, 206.73.118.1)
2. If any of the **ping** commands fails, verify the DNS configuration on the network.

**Set a Preferred Licensing Server**
By default, terminal servers communicate with license servers in the same domain. This behavior might force them to use wide area network (WAN) connections to find a license server, even if there is a license server from another domain located within the same site. By explicitly specifying a license server, you enable the crossing of domains for license tracking and accounting purposes. For information about how to do this, see Set a preferred license server.

**Verify Whether the Terminal Server Can Discover the License Server**
To verify whether the terminal server can discover the license server, you can use Terminal Server Licensing, as described in the following procedures.

Using Terminal Server Licensing (on Windows Server 2003)

Using Terminal Server Licensing (on Windows 2000)

**Using Terminal Server Licensing (on Windows Server 2003)**
To connect to a license server, a terminal server must have the **Access this computer from the network** logon right for the target license server. Additionally, if the License Server Security Group setting in Group Policy is enabled, the terminal server computer account must be a member of this group. Enabling this Group Policy setting and applying it to your license server creates a local group called Terminal Server Computers. The license server issues licenses only to the terminal servers in this group. You must add the terminal servers for which you need to provide licenses to this group for each license server.

**To verify whether the terminal server can discover the license server (on Windows Server 2003)**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, click **Control Pane**l, double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**.

   When Terminal Server Licensing opens, it discovers all license servers in your workgroup or domain (and searches Active Directory for Enterprise License Servers) and it displays the names of those servers in the console tree. If no license servers are automatically displayed in the console tree, you are prompted to specify the name of a server as noted in the following steps.
2. On the **Action** menu, click **Connect**.
3. In **Server**, type the name of the license server to which you want to connect, and then click **OK**.

**Using Terminal Server Licensing (on Windows 2000)**
In addition to supporting clients that run Windows Server 2003, Windows 2000, and Windows XP Professional, Windows 2000 Terminal Server supports clients running the following platforms: Windows XP Home Edition, Windows XP Embedded Edition, Windows Millennium Edition, Windows 98, Windows 95, Windows NT 4.0, Windows-based Terminal devices, and Macintosh. These clients must have Terminal Server CALs available to connect to a Windows 2000 terminal server.

If you have clients running Windows 2000 or Windows XP Professional, the license server issues a "free" Terminal Services CAL from its built-in pool to enable these clients to access the Windows 2000 terminal server. The use of this license is permitted under the Windows 2000 Server EULA. To issue these licenses, however, the license server must first be activated.

For Windows Server 2003 terminal servers, there is no built-in pool of free licenses. You must install the appropriate number of CALs on the license server to enable all clients to connect to these terminal servers.

**To verify whether the terminal server can discover the license server (on Windows 2000)**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Terminal Services Licensing**.

2. Navigate to **All Servers**, and then click the name of the terminal server to which you want to connect.

   A list of installed CALs appears in the details pane. Verify that **Windows 2000 Terminal Services Client Access License** appears in the list.

3. If it does not appear, purchase and install the appropriate number of CALs to support your clients.

**Verify RDP-tcp Connection Settings**

Use Group Policy or Terminal Services Configuration to verify that RDP-tcp connection settings are configured as follows:

- The RDP-tcp connection is enabled.

- The **Maximum connections** setting on the **Network Adapter** tab of the **RDP-tcp Properties** dialog box is set to **Unlimited**.

- Verify whether the level encryption being used for the RDP-TCP connection is set to **High**. If the encryption level is set to **High**, keep in mind that this level encrypts data sent from client to server and from server to client by using strong 128-bit encryption. Clients that do not support this level of encryption cannot connect to a terminal server.

- Verify that the user attempting to log on to the terminal server over the RDP-tcp connection has the appropriate rights and permissions. Users must either be members of the Remote Desktop Users group, or you must assign the equivalent rights and permissions manually (users must have the Allow log on through Terminal Services (http://go.microsoft.com/fwlink/?LinkID=48894) right and User Access permissions). For more information, see Enabling users to connect remotely to the server (http://go.microsoft.com/fwlink/?LinkID=48887). If users are running the Citrix ICA client, they must have the appropriate permissions for the ICA-TCP connection.

For instructions on how to verify and modify RDP-tcp connection settings, see Configure Terminal Services Connections (http://go.microsoft.com/fwlink/?LinkId=48219).

**Verify the License Server Role**

**Caution** Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**To To verify the license server role**

1. On the license server, open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

2. Locate, and then click, the following key in the registry:

   **HKLM\System\CurrentControlSet\Services\TermServLicensing\Parameters\Role**

   Where

   **0** = Domain license server

   **1** = Enterprise license server

**Verify that the Remote Registry Service is Started on the License Server and that the Startup Type is Set to Automatic**

- You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

If you enable or disable a service and you encounter a problem starting the computer, you might be able to start the computer in Safe Mode. Then you can change the service configuration or restore the default configuration.

**Important** If you stop, start, or restart a service, any dependent services are also affected.

**To verify that the Remote Registry service is enabled on the license server and that the startup type is set to**

**Automatic**

1. On the license server, open **Services**. To open Services, click **Start**, click Control Panel, click **Performance and Maintenance**, click **Administrative Tools**, and then double-click **Services**.
2. In the list of services, right-click **Remote Registry**, and then click **Properties**.
3. On the **General** tab, verify that the service status is Started and that the startup type is Automatic.
4. If the status is not Started, click **Start** to start the service. If the startup type is not Automatic, click **Automatic** to ensure that the service will restart when the license server is restarted.

**Verify Permissions on the Domain Controllers OU**

For license server discovery to work in a Windows Server 2003 domain with Windows 2000 terminal servers and Windows 2000 license servers, the following permissions must exist on the Domain Controllers OU.

- System = Full Control
- Authenticated Users = Read
- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as (http://go.microsoft.com/fwlink/?LinkID=48886) to perform this procedure.

**To verify permissions on the Domain Controllers OU**

1. Open Active Directory Users and Computers. To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
2. On the **View** menu, select **Advanced Features**.
3. Right-click the **Domain Controllers** OU, and then click **Properties**.
4. On the **Security** tab, click **Advanced** to view the permission entries that exist for the object.
5. In the list, click **Authenticated Users**.
6. In the **Permissions** box, make sure that the **Allow** check box is selected for **Read.**
7. In the list, click **System**.
8. In the **Permissions** box, make sure that the **Allow** check box is selected for Full Control.

**Verify the Value of the RestrictAnonymous Registry Key**

If you are using Windows 2000 servers, verify that the **RestrictAnonymous** registry key value on the license server is set to**0**or **1**. In this scenario, if this registry key value is set to**2**, then the license server cannot issue CALS. When the**RestrictAnonymous** registry value is set to **2**, the access token built for non-authenticated users does not include the Everyone group, and because of this, the access token no longer has access to resources that grant permissions to the Everyone group.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**To verify the value of the RestrictAnonymous registry key**

1. On the license server, open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Locate, and then click, the following key in the registry:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\RestrictAnonymous**

3. Verify the value of the **RestrictAnonymous**registry key.

The three valid registry key values for **RestrictAnonymous** are equivalent to the following security settings in Local Security Policy:

**0** = None (rely on default permissions)

**1** = Do not allow enumeration of SAM accounts and shares

**2** = No access without explicit anonymous permissions

If the license server is a member of an Active Directory domain and a conflicting security setting is configured for the license server in Group Policy, the Group Policy setting overrides the local security setting.

4. If the value of the registry key is set to **2**, then set it to **1** or **0**.

   Also, verify that the equivalent value (**No access without explicit anonymous permissions**) is not enabled by Group Policy. To do this, open Local Security Policy, navigate to **Security Settings\Local Policies\Security Options\Additional restrictions for anonymous connections**, and then verify whether **No access without explicit anonymous permissions** is enabled. If it is, disable it. To open Local Security Policy, click **Start**, point to **Settings**, Control Panel, click **Administrative Tools**, and then click **Local Security Policy**.

5. Restart the server.

   In Windows Server 2003, if you need to prohibit anonymous users from being granted the same access that is granted to members of the Everyone group, use the new **Everyone Network access: Let Everyone permissions apply to anonymous users** setting in Local Security Policies.

#### Verify that the Administrative Shares are Shared on the License Server

**To verify that the administrative shares are shared on the license server**

1. At the command prompt, type **net share**, and then press ENTER.

2. Verify that the default shares, such as C$, Admin$, and IPC$, are returned.

   If you perform this procedure and the default administrative shares are not returned, verify that the Server service is started, that File and Print Sharing is enabled, and that the license server does not have a virus. Also, see HOW TO: Restore Administrative Shares That Have Been Deleted (http://go.microsoft.com/fwlink/?LinkId=48223) on the Microsoft Web site.

#### Verify that File and Printer Sharing is Bound to the Network Adapter and that the Internal Network Adapter is at the Top of the Binding Order

**To verify that File and Printer Sharing is bound to the internal network adapter and that the internal network adapter is at the top of the binding order**

1. Open Network and Dial-up Connections. To open Network and Dial-up Connections, click **Start**, point to **Settings**, and then click **Network and Dial-up Connections**.

2. Click the connection that you want to verify, and on the **Advanced**menu, click **Advanced Settings**.

3. On the **Adapters and Bindings** tab, in **Bindings for *AdapterName***, verify that the **File and Printer Sharing** check box is selected. If the check box is not selected, select it.

   To modify the protocol bindings order, you must be logged on as a member of the Administrators group.

4. Also in **Bindings for *AdapterName***, verify that the internal adapter is at the top of the list. If it is not at the top of the list, click the up and down arrows to move the internal adapter to the top of the list.

#### Verify that the Terminal Server Licensing Service is Started and that the Startup Type is Set to Automatic

For step-by-step instructions, see Verify that the Terminal Server Licensing service is started and the startup type is set to Automatic.

#### Verify that Terminal Server Licensing is Correctly Installed on the License Server

For step-by-step instructions, see Verify that Terminal Server Licensing is correctly installed on the license server.

#### Verify that the License Server is Activated

A license server that has been installed but not activated only issues temporary licenses. These temporary licenses allow clients to connect to the terminal server for 90 days.

**To verify that a license server is activated**

1. Open Terminal Server Licensing.
2. In the console tree, click the license server for which you want to verify the activation status.
3. In the details pane, verify that activation status is Activated.
4. If the status is Not Activated, right-click the server name, click **Activate Server**, and then follow the instructions in the wizard.

**Remove the MSLicensing Registry Key on the Client and Verify Permissions on the Rebuilt Key**

If all of the previous troubleshooting procedures fail, create a backup of the **MSLicensing**registry key and its subkeys on the client, and then remove the original key and subkeys.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**To remove the MSLicensing registry key on the client and verify permissions on the rebuilt key**

1. On the client, open Registry Editor. To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Locate, and then click, the following key in the registry:

   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSLicensing**
3. On the **Registry** menu, click **Export Registry File**.
4. In the **File name** box, type **mslicensingbackup**, and then click **Save**.
5. If you need to restore this registry key in the future, double-click **mslicensingbackup.reg**.
6. On the **Edit** menu, click **Delete**, and then click **Yes** to confirm the deletion of the **MSLicensing**registry subkey.
7. Close Registry Editor, and then restart the computer (when the client is restarted, the missing registry key is rebuilt).
8. On the client, open Registry Editor.
9. Locate, and then click, the following key in the registry:

   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSLicensing**
10. On the **Edit** menu, click **Permissions**. Users must have at least Read permissions.

# The remote session was disconnected because there are no Terminal Server license servers available to provide a license

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
The terminal server might not be able to locate the license server.

**Solution**
Perform the troubleshooting procedures described in The terminal server cannot locate the license server.

**Cause**
The client might have exceeded its Terminal Server Licensing grace period (http://go.microsoft.com/fwlink/?LinkID=48895), and a license server has not yet been installed or activated to issue the client a client access license (CAL). After a temporary license expires, a client can only connect to a terminal server if the license server can issue a permanent CAL.

**Solution**
Install Terminal Server Licensing, activate the license server, then purchase and install a sufficient number of CALs to support the clients in your organization. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886) to perform this procedure.

| **Note** |
| --- |
| If you are installing an Enterprise License server and your environment includes multiple trusted Active Directory sites, each of which hosts a terminal server, you can use a script to publish a single license server across all of these sites. Doing this can greatly simplify deployment because it eliminates the need to deploy a separate license server in each site, and it eliminates the need to distribute CALs to license servers across multiple sites. For a sample script that you can use to accomplish this task, see the Microsoft Script Center (http://go.microsoft.com/fwlink/?LinkID=48230) on the Microsoft Web site. Note that this sample script is not supported under any Microsoft standard support program or service. |

**To install Terminal Server Licensing**

1. Open Add/Remove Programs in **Control Panel**.
2. Double-click **Add/Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. In Terminal Server Licensing Setup, do one of the following:
   - Click **Your entire enterprise** to choose an Enterprise license server. You can only install an Enterprise license server on a domain controller or a member server in a domain (if the member server is running Windows Server 2003), not on a stand-alone server. This role is published to Active Directory. If you are installing an Enterprise License Server, you must meet one of the following conditions: 1) You must be a member of the Enterprise Admins group in Active Directory or 2) You must have the right to create the TS-Enterprise-License-Server object in the same site in which the license server will reside in Active Directory.
   - Click **Your domain or workgroup** to choose a Domain license server. If you select this option, the computer on which you install the license server does not have to be a domain controller, if the computer is running Windows Server 2003. However, if the computer is running Windows 2000, the computer on which you install the license server must be a domain controller or a stand-alone server in a workgroup. If the computer on which you plan to install the license server is in a workgroup, it can only be automatically discovered by terminal servers on the same subnet. If you are installing a Domain license server, you must be a member of the Administrators group on the local computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
5. Provide the file system location where the license server database should be installed on the Terminal Server license server, click **Next**, and then click **Finish**.

   The default location for the license server database is systemroot\System32\LServer.

You can activate a Terminal Server license server by using one of three methods: Automatic, Web Browser, or Telephone. It is recommended that you activate a license server automatically, as described in the following procedure.

| **Note** |
| --- |
| To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider Using Run as (http://go.microsoft.com/fwlink/?LinkID=48886). |

**To activate a license server automatically.**

1. Open Terminal Server Licensing. To open Terminal Server Licensing, click **Start**, point to **Control Panel**, **Administrative Tools**, and then click **Terminal Server Licensing**.
2. In the console tree, right-click the license server that you want to activate, click **Activate Server**, and then follow the instructions in the wizard.
3. In the **Activation** method list, select **Automatic connection (recommended)**, and then click **Next**.
4. Follow the instructions in the wizard.
5. If necessary, at a later time you can change any license server properties, such as activation method and company information, that you set during the activation process.

Then, Purchase CALs as required and Install CALs automatically.

# The remote session was disconnected because there are no Terminal Server client access licenses available for this computer

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
The license server might not have any remaining Per Device CALs to issue.

**Solution**

1. Use Terminal Server Licensing to Verify that the license server has a sufficient number of CALs for clients requiring remote connections.

2. If the license server does not have a sufficient number of permanent CALs, Purchase CALs as needed, and then Install CALs automatically, as needed.

**Cause**
In the following two cases, either no licenses, or temporary licenses that cannot be upgraded, are issued from the available pool of client access licenses (CALs) on a license server:

1. If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per User, then no licenses are issued from the available pool of licenses, because Per User licensing is not managed.

2. If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per Device but the license server is configured to issue only Per User CALs, then only temporary Per Device licenses are issued from the available pool of licenses. These licenses cannot be upgraded. When the temporary licenses are close to expiration, the following message appears on the terminal server: "Your terminal services temporary client license will expire in N days. Please contact your system administrator to get a permanent license." In this case, Event ID 26, Source: Application Pop-up appears in the application event log.

   Similarly, Terminal Server clients that run Windows XP or Windows 2000 might receive the following error message, if the client cannot connect to the terminal server for this reason: "The remote session was disconnected because the local computer client access license could not be upgraded or renewed." In this case, Event ID 1011, Source:TermService appears in the application event log on the terminal server, with the following message: "The terminal services client has been disconnected because its temporary license has expired."

**Solution**
Do one of the following, based on the requirements of your environment and license agreement:

- Call the Microsoft Clearinghouse and reactivate the license server in the correct mode.

- Change the Terminal Server Licensing mode.

**Cause**
The terminal server might not be able to locate the license server.

**Solution**
To resolve this issue, perform the troubleshooting procedures described in The terminal server cannot locate the license server.

# The remote session was disconnected because the local computer's client access license could not be upgraded or renewed

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

After a temporary license expires, a client can only connect to a terminal server if the license server can issue a permanent CAL, or if the terminal server is still within its licensing grace period.

**Note**

If you are running Terminal Server Licensing in a Windows 2000 Service Pack 3 (SP3) or later environment, the client must connect to a terminal server and authenticate twice to obtain a permanent license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and is not yet issued, the terminal server issues the client a permanent, Per Device CAL.

**Cause**

If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per Device but the license server is configured to issue only Per User CALs, then only temporary Per Device licenses are issued from the available pool of licenses. These licenses cannot be upgraded. When the temporary licenses are close to expiration, the following message appears on the terminal server: "Your terminal services temporary client license will expire in N days. Please contact your system administrator to get a permanent license." In this case, Event ID 26, Source: Application Pop-up appears in the application event log.

Similarly, Terminal Server clients that run Windows XP or Windows 2000 might receive the following error message, if the client cannot connect to the terminal server for this reason: "The remote session was disconnected because the local computer client access license could not be upgraded or renewed." In this case, Event ID 1011, Source:TermService appears in the application event log on the terminal server, with the following message: "The terminal services client has been disconnected because its temporary license has expired."

**Solution**

Do one of the following, based on the requirements of your environment and license agreement:

- Call the Microsoft Clearinghouse and reactivate the license server in the correct mode. For more information, see Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkId=48885).

- Change the Terminal Server Licensing mode.

**Cause**

If you are running Terminal Server Licensing in a Windows 2000 (SP3) or earlier environment, expired temporary licenses might not be upgraded to permanent CALs.

**Solution**

1. On the terminal server and license server, install Service Pack 4 (SP4). To download SP4, see Windows 2000 Service Packs (http://go.microsoft.com/fwlink/?LinkID=48213) on the Microsoft Web site.

2. Use Terminal Server Licensing to Verify that the license server has a sufficient number of CALs for clients requiring remote connections.

3. If the license server does not have a sufficient number of permanent CALs, Purchase CALs, and then Install CALs automatically.

# No licenses are being issued from the available pool of licenses

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**

In the following two cases, either no licenses, or temporary licenses that cannot be upgraded, are issued from the available pool of client access licenses (CALs) on a license server:

1. If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per User, then no licenses are issued from the available pool of licenses, because Per User licensing is not managed.

2. If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per Device while the license server is configured to issue only Per User CALs, then only temporary Per Device CALs are issued from the available pool of licenses. These CALs cannot be upgraded. When the temporary CALs are within several days of expiring, "Event ID 26, Source: Application Pop-up" appears in the application event log on the client. The event message indicates the number of days remaining before the temporary license expires. Similarly, "Event ID 1011, Source:TermService" appears in the application event log on the terminal server.

**Solution**

Do one of the following, based on the requirements of your environment and license agreement:

- Call the Microsoft Clearinghouse and reactivate the license server in the correct mode. For more information, see "Locate the Microsoft Clearinghouse telephone number for your country or region" (http://go.microsoft.com/fwlink/?LinkId=48885).

- Change the Terminal Server Licensing mode.

# Event ID 26: Your terminal services temporary client license will expire in N days

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per Device but the license server is configured to issue only Per User CALs, then only temporary Per Device licenses are issued from the available pool of licenses. These licenses cannot be upgraded. When the temporary licenses are close to expiration, the following message appears on the terminal server: "Your terminal services temporary client license will expire in N days. Please contact your system administrator to get a permanent license." In this case, Event ID 26, Source: Application Pop-up appears in the application event log.

Similarly, Terminal Server clients that run Windows XP or Windows 2000 might receive the following error message, if the client cannot connect to the terminal server for this reason: "The remote session was disconnected because the local computer client access license could not be upgraded or renewed." In this case, Event ID 1011, Source:TermService appears in the application event log on the terminal server, with the following message: "The terminal services client has been disconnected because its temporary license has expired."

**Solution**
Do one of the following, based on the requirements of your environment and license agreement:

- Call the Microsoft Clearinghouse and reactivate the license server in the correct mode. For more information, see Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkId=48885).

- Change the Terminal Server Licensing mode.

**Cause**
The terminal server might not be able to locate the license server.

**Solution**
To resolve this issue, perform the troubleshooting procedures described in The terminal server cannot locate the license server.

# Event ID 1011: The terminal services client has been disconnected because its temporary license has expired

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**
If you are running Terminal Server Licensing in a Windows Server 2003 environment and the licensing mode on the terminal server is set to Per Device but the license server is configured to issue only Per User CALs, then only temporary Per Device licenses are issued from the available pool of licenses. These licenses cannot be upgraded. When the temporary licenses are close to expiration, the following message appears on the terminal server: "Your terminal services temporary client license will expire in N days. Please contact your system administrator to get a permanent license." In this case, Event ID 26, Source: Application Pop-up appears in the application event log.

Similarly, Terminal Server clients that run Windows XP or Windows 2000 might receive the following error message, if the client cannot connect to the terminal server for this reason: "The remote session was disconnected because the local computer client access license could not be upgraded or renewed." In this case, Event ID 1011, Source:TermService appears in the application event log on the terminal server, with the following message: "The terminal services client has been disconnected because its temporary license has expired."

**Solution**
Do one of the following, based on the requirements of your environment and license agreement:

- Call the Microsoft Clearinghouse and reactivate the license server in the correct mode. For more information, see Locate the Microsoft Clearinghouse telephone number for your country or region (http://go.microsoft.com/fwlink/?LinkId=48885).

- Change the Terminal Server Licensing mode.

**Cause**
The terminal server might not be able to locate the license server.

**Solution**
To resolve this issue, perform the troubleshooting procedures described in The terminal server cannot locate the license server.

# Event ID 1026: The terminal server could not locate a license server in the domain

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**

The process of locating a license server is called license server discovery. If license server discovery is not functioning correctly, Event ID 1026, Source:TermServ, might appear in the application event log, with the following message: "The terminal server could not locate a license server in the <domain> domain. Confirm that all license servers on the network are registered in WINS/DNS, accepting network requests, and the Terminal Server Licensing Service is running."

**Solution**

Set a preferred license server.

If your terminal server is running Windows Server 2003 with Service Pack 1 (SP1), you can do this by using Group Policy settings or Terminal Services Configuration, or you can use the registry or a Windows Management Instrumentation (WMI) script. If your terminal server is running Windows 2000 or Windows Server 2003 without SP1, you must use the registry or a WMI script. For three scripts that you can use to set preferred license servers, delete preferred license servers, or query preferred license servers, see Preferred License Server WMI Scripts (http://go.microsoft.com/fwlink/?LinkID=26220) in Windows Server 2003 Terminal Server Licensing on the Microsoft Web site.

**Note**

The license servers and the terminal servers must either be in the same domain or in domains that trust each other.

**Solution**

If explicitly specifying a license server does not resolve this issue, perform additional troubleshooting steps, as described in The terminal server cannot locate the license server.

# Event ID 1028: The terminal server client has been disconnected because its license could not be renewed

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are running a Windows Server 2003 terminal server, after you use the Windows Components Wizard in Add or Remove Programs to add or to remove a Windows component, an invalid registry value might be set in the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\TerminalServer\Licensing Core\PolicyAcOn** subkey on the terminal server. This invalid value might may unexpectedly change the Terminal Server Licensing mode and cause the terminal server to obtain only temporary licenses.

**Cause**

This problem might occur because the value of the **PolicyAcOn** registry entry changes from either **2** (the correct registry value for the Per Device licensing mode) or **4** (the correct registry value for the Per User licensing mode) to an invalid value of **0**. This behavior causes the terminal server to attempt to obtain temporary Per Device licenses. In this case, if you view Terminal Services Configuration to verify the licensing mode, the licensing mode appears as Per Device. Therefore, to verify whether this problem is causing the client disconnection, you must verify the value of the **PolicyAcOn** registry entry.

**Solution**

To resolve this problem contact Microsoft Product Support Services to obtain the hotfix described in Knowledge Base article 834651, Terminal Services licensing mode changes from Per User or from Per Device changes from Per User or from Per Device to an invalid value after you add or remove a Windows component, (http://go.microsoft.com/fwlink/?LinkID=48173) on the Microsoft Web site. This hotfix is applied to the terminal server. For information about how to contact Microsoft Product Support Services, see Help and Support (http://go.microsoft.com/fwlink/?LinkID=48149) on the Microsoft Web site.

# Fixing License Server Activation Problems and Other Licensing Problems

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides a description of terminal server license server activation problems, and other problems that you might experience when Terminal Server Licensing is enabled.

From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- The product ID generated by the Terminal Server License Server Activation Wizard doesn't work
- Event ID 43: Work Manager error can't startup work scheduler, Error code -1072167891

# The product ID generated by the Terminal Server License Server Activation Wizard doesn't work

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In a Windows Server 2003 environment, when you use the Terminal Server License Server Activation Wizard to activate a Terminal Services license server by telephone, you might not be able to use the product ID that appears on the License Server Activation page of the wizard to complete the activation of the license server.

**Cause**

This problem might occur when the Terminal Server License Server Activation Wizard incorrectly processes the product ID. When this problem occurs, Event ID 43 appears in the application event log on the license server, with the following message: "Work Manager error can't startup work scheduler, Error code -1072167891."

**Solution**

To resolve this problem, contact Microsoft Product Support Services to obtain a hotfix for this issue, as described in Knowledge Base article 832794, Terminal Server License Server Activation Wizard generates a product ID that does not work (http://go.microsoft.com/fwlink/?LinkId=48232) on the Microsoft Web site. For information about how to contact Microsoft Product Support Services, see Help and Support (http://go.microsoft.com/fwlink/?LinkId=48149) on the Microsoft Web site.

# Event ID 43: Work Manager error can't startup work scheduler, Error code -1072167891

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cause**

This behavior occurs because Microsoft Jet database engine transactions can use a significant amount of memory on a large database (a database that is larger than 80 megabytes [MB] in size). If a transaction hangs (runs for a long time before committing) or performs many modifications, Jet might consume the available version store.

**Solution**

In Windows Server 2003, you can use a new registry key to tune the default memory resources that are necessary for version store in Jet transactions. The version store maintains list of changes that are made to the database since the latest transaction began.

**To tune the memory resources that are necessary for version store in Jet**

1. On the license server, open Registry Editor. To open Registry Editor, click **Start**, click **Run,** type **regedit**, and then click **OK**.

2. Locate, and then click, the following key in the registry:

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\TermServLicensing\Parameters**

3. On the **Edit** menu, click **Add Value**, and then add the following registry value:

   Value name: **MaxVerPages**

   Data type: REG_DWORD

   Radix: Decimal

   Value data: 1024

   The valid range for Value is 256 to 2048

4. **Caution**

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

# Additional Resources for Troubleshooting Terminal Server Licensing

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For additional information about troubleshooting Terminal Server Licensing in Windows Server 2003 and Windows 2000, see the following resources.

**Known Issues**

**Windows Server 2003 and Windows 2000**

- Because of a security error, the client could not connect to the Terminal Server (http://go.microsoft.com/fwlink/?LinkId=46818) on the Microsoft Web site
- Client computers cannot obtain a permanent Windows 2000 Terminal Services Client Access License (CAL) from a Windows Server 2003 Terminal Services Licensing server (http://go.microsoft.com/fwlink/?LinkId=48146) on the Microsoft Web site
-
- **Windows Server 2003**
- Windows Server 2003 Terminal Server licensing issues and requirements for deployment (http://go.microsoft.com/fwlink/?LinkId=26381) on the Microsoft Web site
- License Server Security Group setting does not work in Windows Server 2003 Terminal Server (http://go.microsoft.com/fwlink/?LinkId=46819) on the Microsoft Web site
- "No License Server" Error Message Appears When You Try to Connect to Terminal Server After Mode Change (http://go.microsoft.com/fwlink/?LinkId=46820) on the Microsoft Web site
- Terminal Services Licensing mode changes from Per User to Per Device after you add or remove a Windows component (http://go.microsoft.com/fwlink/?LinkId=46821) on the Microsoft Web site
- Event ID 43: Work Manager error can't startup work scheduler, Error code -1072167891
- Terminal Server License Server Activation Wizard generates a product ID that does not work (http://go.microsoft.com/fwlink/?LinkId=48232) on the Microsoft Web site
-

**Windows 2000**

- Windows XP Clients Cannot Connect to a Windows 2000 Terminal Services Server (http://go.microsoft.com/fwlink/?LinkId=46822) on the Microsoft Web site

**How To**

The following resources provide additional information about common Terminal Server Licensing tasks in Windows Server 2003 and Windows 2000.

**Windows Server 2003**

- How to activate a License Server by using Terminal Server Licensing in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=46915) on the Microsoft Web site

- HOW TO: Deactivate or Reactivate a License Server Using Terminal Services Licensing - 2003 (http://go.microsoft.com/fwlink/?LinkId=46918) on the Microsoft Web site

- How to override the license server discovery process in Windows Server 2003 Terminal Services (http://go.microsoft.com/fwlink/?LinkId=36459) on the Microsoft Web site

- Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkId=38499) on the Microsoft Web site

**Windows 2000**

- HOW TO: Locate a Phone Number for the Microsoft Clearinghouse (http://go.microsoft.com/fwlink/?LinkId=46919) on the Microsoft Web site

- How to use the RestrictAnonymous Registry value in Windows 2000 (http://go.microsoft.com/fwlink/?LinkId=48242) on the Microsoft Web site

- Establishing Preferred Windows 2000 Terminal Services License Server (http://go.microsoft.com/fwlink/?LinkId=46953) on the Microsoft Web site

**Overview, FAQ, Tools**

The following resources provide overview information that can aid in troubleshooting Terminal Server Licensing in Windows Server 2003 and Windows 2000.

**Windows Server 2003 Overview**

- Terminal Services Licensing service discovery (http://go.microsoft.com/fwlink/?LinkId=46954) on the Microsoft Web site

- The Function of Terminal Server CALs in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=46956) on the Microsoft Web site

- Licensing Terminal Server in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=48245) on the Microsoft Web site

- Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkId=48247) on the Microsoft Web site

**Windows 2000 Overview**

- Terminal Services Licensing Enhancements (http://go.microsoft.com/fwlink/?LinkId=48172) on the Microsoft Web site

- Removing Terminal Server licenses from an RDP client (http://go.microsoft.com/fwlink/?LinkId=46958) on the Microsoft Web site

**Windows Server 2003 and Windows 2000 Overview**

- Description of the License Logging Service in Windows Server Operating Systems (http://go.microsoft.com/fwlink/?LinkId=46959) on the Microsoft Web site

**FAQs**

- "Terminal Server" section in Windows Server 2003 Pricing and Licensing: Frequently Asked Questions (http://go.microsoft.com/fwlink/?LinkId=46960) on the Microsoft Web site

- Terminal Server CAL Transition Plan: Frequently Asked Questions (http://go.microsoft.com/fwlink/?LinkId=46961) on the Microsoft Web site

- Frequently asked questions about Terminal Services (http://go.microsoft.com/fwlink/?LinkId=48251) on the Microsoft Web site

**Event and Error Messages**

- Microsoft TechNet Events and Errors Message Center (http://go.microsoft.com/fwlink/?LinkId=46963) on the Microsoft Web site

**Tools**

- Lsview.exe: Terminal Services License Server Viewer (http://go.microsoft.com/fwlink/?LinkId=46964) on the Microsoft Web site

- Tsctst.exe: Terminal Server Client License Dump Tool (http://go.microsoft.com/fwlink/?LinkId=48253) on the Microsoft Web site

- Lsreport.exe: Terminal Server Licensing Reporter (http://go.microsoft.com/fwlink/?LinkId=48253) on the Microsoft Web site

- Useful Terminal Services Resource Kit Utilities (http://go.microsoft.com/fwlink/?LinkId=46967) on the Microsoft Web site

-

- **Product Documentation**

- Windows Server 2003 Terminal Server Licensing Help (http://go.microsoft.com/fwlink/?LinkId=45409) on the Microsoft Web site

- Hosting Applications with Terminal Server (http://go.microsoft.com/fwlink/?LinkId=46968) in the Windows Server 2003 Deployment Kit, on the Microsoft Web site

# Step-by-Step Guide for Configuring Group Policy for Terminal Services

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Terminal Services provides configuration operations on a per-computer basis. Group Policy combined with Terminal Services provides central administration of terminal servers and remote desktop clients.

The basic scenario of this guide provides a way to assign the same user settings to all users logging on to Terminal Services. When not logged on to Terminal Services, users should see their normal configuration settings. Group Policy uses the loopback mechanism to apply the same consistent desktop to all Terminal Services users.

The basic scenario is then extended to assign the same user settings to all users logging on to Terminal Services while exempting domain administrators from these settings. This approach makes remote administration possible for domain administrators while providing a consistent logon experience for Terminal Services users.

# Loopback Technology Review

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Understanding Group Policy loopback technology is essential for configuring a consistent desktop for all Terminal Services users. This section introduces the Group Policy loopback feature with application to terminal servers.

**Group Policy loopback feature**

Group Policy is applied to the user or computer, based upon where the user or computer object is located in Active Directory® directory service. However, when applying a consistent logon experience to all users logging on to a terminal server, users will need policy applied to them, based on the location of the terminal server, not the location of the user object. The Group Policy loopback feature gives the administrator the ability to apply user Group Policy, based upon the terminal server Active Directory location.

To illustrate the loopback feature, we will use an example. In this example, you have full control over the computers and users in this domain because you have been granted domain administrator privileges.

The following illustration shows the Contoso.com domain, which is used for the purposes of this example. The gold lines connecting the Active Directory objects (domain and OU containers) are shown as solid lines. The GPO links between Active Directory objects (site, domain, and OU containers) and individual GPOs are shown as arrows.




**Figure 1 The Contoso.com domain**

The Terminal servers OU in the Contoso domain contains only the terminal servers. Users from the Marketing OU have GPOs A3, A1, A2, and A5 applied (in that order), regardless of the desktop computer to which they log on. The terminal servers located in the Terminal servers OU will have the computer settings from GPOs A3, A1, A2, A4, A6, and A7 applied (in that order) during computer Group Policy processing.

Loopback can be configured into two modes: replace and merge. For Terminal Services configuration purposes, the replace mode of the loopback policy setting is relevant. In this mode, the user's list of GPOs is not gathered. Only the list of GPOs based on the computer object location in the Active Directory is used. The loopback policy setting will be configured to replace mode in GPO A6. GPO A7 will be configured with the desired user-based policy settings for the Terminal Services desktop. Both GPO A6 and A7 are linked to the Terminal servers OU. This illustrates a best practice to separate the computer-based policy settings from the user-based policy settings to make troubleshooting Group Policy application problems easier.

For the Contoso.com domain example, users from the Marketing OU, when logging on to their normal desktops, will apply

user policy settings from GPOs A3, A1, A2, and A5 (in that order). When logging on to any terminal server, the users from the Marketing OU will evaluate a different set of user policy settings from GPOs A3, A1, A2, A4, A6, and A7 (in that order). You can easily see the difference in the set of user's policy settings based on the computer's location when seen in a table.

| Computer | GPOs evaluated to user when logging on |
|---|---|
| Normal desktop computer | A3, A1, A2, A5 |
| Terminal server | A3, A1, A2, A4, A6, A7 |

**Note**

The list of GPOs evaluated for a user logging on to a terminal server will include A3, A1, A2, A4, A6, and A7 (in that order). However, A6 will be found to only include computer settings and will not apply to a user.

In summary, administrators can use the loopback feature in replace mode to override the normal processing of user Group Policy settings. The user portion of the GPOs linked to the Terminal servers OU will be used for all terminal server logons, instead of evaluating the user GPOs in Group Policy processing normally applied when logging on to their desktop computer.

**Loopback processing and security filtering**

With security filtering, you can apply a GPO with user-based policy settings to a specific domain group with only user members. The computer policy will not apply, but the user policy settings will apply to any user belonging to the security group used in filtering.

Using the Contoso domain example, you would want to use security filtering to limit application of the user policy settings of GPO A7 to the Remote Desktop Users group. In this case, the Remote Desktop Users group will have the user policy settings of GPOs A3, A1, A2, A4, A6, and A7 applied to them and in that order. All other users will have only the user policy settings of GPOs A3, A1, A2, A4, and A6 applied to them and in that order. This method of limiting application of GPOs through security filtering will come in handy for Scenario 2.

**Loopback processing best practices**

When using loopback processing in replace mode to control the user policy processing settings, it is best to create separate GPOs for the user and computer configuration settings. Doing so will decrease the amount of confusion involved with using security filtering on a GPO—whether to filter for the user or computer side. Creating separate GPOs also will help in troubleshooting any issues surrounding the processing of Group Policy settings with loopback processing enabled.

**See Also**

**Other Resources**

Loopback Processing with Merge or Replace
Group Policy Loopback Support
Technical Overview of Windows Server 2003 Terminal Services

# Requirements for Configuring Group Policy for Terminal Services

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This step-by-step guide assumes you understand the basic concepts of Group Policy and using Group Policy Management Console.

Before running the scenarios presented in this document, you must follow these steps in order:

1. Prepare your environment.
2. Enable users to connect remotely to a terminal server.
3. Centrally enable Remote Desktop using Group Policy.
4. Grant access to your terminal servers with the Remote Desktop Users Group.

| **Note** |
| --- |
| It is assumed the firewall on the terminal server is off. If you have configured a firewall on your terminal server, you must allow inbound Remote Desktop exceptions through the firewall on the terminal server. |

**Prepare your environment**

You must have at least one Active Directory domain. You must have Group Policy Management Console (GPMC) installed on an administration computer. The administration computer should contain all the tools for administrating the domain-based operations.

In your domain, create a separate organizational unit (OU) for terminal servers. Move all terminal servers into that OU. Do not move any users or other computers into that OU.

It is assumed that you understand the basics of creating, managing, and administering Group Policy Active Directory domains. See Implementing Common Desktop Management Scenarios with the Group Policy Management Console for more information about creating, managing, and administering Group Policy Active Directory domains.

**Enable users to connect remotely to a terminal server**

To enable users to connect remotely to a terminal server, you must ensure that:

- Remote Desktop is enabled on the terminal server.
- Users have the appropriate rights and permissions to log on remotely to the server.

| **Note** |
| --- |
| To perform these two tasks, you must be logged on as a member of the Administrators group. These properties are configured on a per-server basis. You can use Group Policy to manage these properties centrally. |

**Centrally enable Remote Desktop using Group Policy**

It is recommended as a best practice to centrally enable Remote Desktop for all your terminal servers. Group Policy will allow you to centrally configure all your terminal servers instead of configuring the properties for each terminal server.

### To centrally enable Remote Desktop using Group Policy

1. To open Group Policy Management Console (GPMC), click **Start**, click **Run**, and then type **GPMC.msc**.
2. Create and link a GPO to the terminal server OU.
3. Right-click the GPO linked to the terminal server OU, and then click **Edit**.

4. In Computer Configuration\Administrative Templates\Windows Components\ Terminal Services, double-click the **Allow users to connect remotely using Terminal Services** policy setting.
5. Click **Enabled**.
6. Click **OK**.

| **Important** |
| --- |
| When you enable Remote Desktop on a computer, you enable the capability for other users and groups to log on remotely to the computer. However, you must also decide which users and groups should be able to log on remotely, and then manually add them to the Remote Desktop Users group. Domain administrators automatically have the ability to log on remotely to the terminal server. Do not add domain administrators to the Remote Desktop Users group.<br><br>You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. |

| **Note** |
| --- |
| If the **Allow users to connect remotely using Terminal Services** policy setting is set to "Not Configured," the **Enable Remote Desktop on this computer** policy setting (on the **Remote** tab of the **System Properties** dialog box) on the target computers takes precedence. Otherwise, the **Allow users to connect remotely using Terminal Services** policy setting takes precedence. |

**Grant access to your terminal servers with the Remote Desktop Users group**

The Remote Desktop Users group is one of the built-in users groups available when you install one of the Microsoft Windows Server 2003 operating systems. When you enable Remote Desktop, members of the Remote Desktop Users group and domain administrators are automatically able to log on remotely to a terminal server.

By default, the Remote Desktop Users group is not populated. Therefore, you must decide which users and groups should have access to log on remotely to a terminal server, and then add them to this group. Instead of populating the Remote Desktop Users group on a per-server basis, as a best practice it is recommended to centrally grant access to users via the Remote Desktop Users group for all your terminal servers using Group Policy. This approach will allow all domain administrators to be able to indirectly add their users to the Remote Desktop Users group for all terminal servers using a domain group.

| **Note** |
| --- |
| This procedure does not show you how to restrict access to the Remote Desktop Users group. Do not add domain administrators to the domain group configured as a member of the Remote Desktop Users group. By default, domain administrators have remote access rights to any terminal server. |

**To add a domain group to the Remote Desktop Users group via Group Policy**

1. To open Group Policy Management Console, click **Start**, click **Run**, and then type **GPMC.msc**.
2. Create and link a GPO named Restricted Groups to the terminal server OU.
3. Right-click the Restricted Groups GPO linked to the terminal server OU, and then click **Edit**.
4. You can configure the Restricted Groups setting in the following location in Group Policy Object Editor:

   **Computer Configuration\Windows Settings\Security Settings\Restricted Groups**
5. Right-click **Restricted Groups** and then click **Add Group**.
6. Click **Browse**, click **Locations**, select the locations you want to browse, and then click **OK**.
7. Type **Remote Desktop Users** in the **Enter the object names to select** text box and then click **Check Names**. Or, click **Advanced**, and then click **Find Now** to list all available groups.
8. Click the **Remote Desktop Users** group and then click **OK**.
9. Click **OK** in the **Add Groups** dialog box to close it. The **Remote Desktop Users Properties** dialog box is then displayed.
10. Click **Add** in the **Members of this group** section of the dialog box.

11. Click **Browse**.
12. Type the name of the domain group in the **Select Users or Groups** dialog box. Click **Check Names,** and then click **OK** to close this dialog box.
13. Click **OK** to close this dialog box to finish adding the domain group to the Remote Desktop Users group.

# Group Policy Terminal Services Scenarios

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following scenarios illustrate how Group Policy applies a consistent user desktop during logon through the Remote Desktop client while providing domain administrators full access for remote administration.

# Scenario 1: Administering Group Policy to Provide a Consistent Terminal Services Desktop

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Terminal Services client, called Remote Desktop Connection (RDC), provides a number of ways to configure the client. Some examples for which you may want to standardize the RDC clients in your organization for all user logons are:

- Client redirection features.
- Terminal Services profile path.
- Wallpaper during connection.
- Printing connection management.
- Other security related lockdown features.

In each of these examples, the desire is to configure all RDC clients to behave the same way on user logon.

**Administering Group Policy to Provide a Consistent Terminal Services Desktop**
You must have privileges delegated to you to create, link, and edit GPOs for a specific OU to perform the steps in this scenario. This scenario assumes that the requirements for these scenarios—as specified in Requirements for Configuring Group Policy for Terminal Services—have been fulfilled.

To configure Terminal Services desktop settings for a consistent user logon experience, you must perform the following tasks:

- Create a loopback GPO linked to the terminal server OU.
- Create user policy settings in a separate GPO linked to the terminal server OU.

**To create a loopback GPO linked to the terminal server OU**

1. To open Group Policy Management Console, click **Start**, click **Run**, and then type **GPMC.msc**.
2. Create and link a GPO to the terminal server OU. Use this GPO for setting the loopback policy setting in replace mode. Name this GPO appropriately for loopback capability.
3. Right-click the loopback GPO linked to the terminal server OU, and then click **Edit**.
4. In Computer Configuration\Administrative Templates\System\Group Policy, double-click the **User Group Policy loopback processing mode** setting.
5. In the **User Group Policy loopback processing mode** properties, click **Enabled**.
6. In the **User Group Policy loopback processing mode** properties, open the **Mode** drop-down list, and then click **Replace**.
7. Click **OK** to close the **User Group Policy loopback processing mode** properties.
8. Click **File,** and then click **Exit** to close Group Policy Object Editor.

**To create user policy settings in a separate GPO linked to the terminal server OU**

1. If Group Policy Management Console is not open, click **Start**, click **Run**, and then type **GPMC.msc**.
2. Create and link a GPO to the terminal server OU. Name this GPO, "User Settings for Terminal Services". Use this GPO for setting the user configuration settings for Terminal Services and other desktop configurations.
3. Right-click the User Settings for Terminal Services GPO linked to the terminal server OU, and then click **Edit**.

4. Configure only user-based policy settings in this GPO. Do not configure any machine-based policy settings in this GPO.
5. After configuring all the user-based policy settings, click **File** and then click **Exit** to close Group Policy Object Editor.

| **Note** |
| --- |
| This guide does not make recommendations about which Terminal Services configuration settings should be set for user logons. This guide makes recommendations only for setting the Terminal Services configuration settings to create a consistent user logon experience. |

# Scenario 2: Filtering Administrators from Terminal Services Group Policies

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Scenario 1 documented a procedure to administer Group Policy in order to configure all RDC clients to behave the same way on user logon. This creates a problem for administrators when the need arises to log on to a terminal server for remote administration. Domain administrators may not want Group Policy to configure their RDC client during remote administration, due to limitations that may be imposed by the Group Policy.

There are a number of ways to apply security filtering to exempt domain administrators from Group Policy application. This scenario presents one approach to applying security filtering to the User Settings for Terminal Services Group Policy object (GPO) from Scenario 1. The net effect of this scenario is to exempt domain administrators from having a specific desktop applied to them on logon through an RDC client.

**Filtering application of a Group Policy object (GPO)**
To perform the steps in this scenario, you must have privileges delegated to you to create, link, and edit GPOs for a specific OU.

This scenario assumes you have previously performed all the steps in Scenario 1: Administering Group Policy to Provide a Consistent Terminal Services Desktop, including creating the loopback processing and user policy settings for Terminal Services GPO.

This scenario will show you how to use security filtering to limit application of the User Settings for Terminal Services GPO to the user domain group when logging on to terminal server. The user domain group is the original group created to add Terminal Services users to the Remote Desktop Users group. The assumption is that no domain administrator has been added to the user domain group. By limiting the application of the user policy settings for the User Settings for Terminal Services GPO to the user domain group, you will have filtered out application to all domain administrators.

**To limit application of the user policy settings for Terminal Services GPO to the user domain group**

1. To open Group Policy Management Console, click **Start**, click **Run**, and then type **GPMC.msc**.
2. In the console tree, double-click **Group Policy Objects** to expand this folder displaying the GPOs for this domain.
3. Click the **User Settings for Terminal Services** GPO.
4. In the results pane, on the **Scope** tab, click **Add**.
5. In the **Enter the object name to select** box, type the name of the user domain group that you previously added to the Remote Desktop Users group. Click **OK**. GPMC will add the read permissions and "Apply Group Policy" permissions by default.
6. Click **Authenticated Users**.
7. Click **Remove** to *not* apply the User Settings for Terminal Services GPO to all other users, such as members of the Domain Administrators group.

| **Note** |
| --- |
| The user domain group is the original group created to add Terminal Services users to the Remote Desktop Users group. Do not add domain administrators to this group. |

# Additional References

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This list of links provides additional information.

- Implementing Common Desktop Management Scenarios with the Group Policy Management Console
- Deploying Remote Desktop for Administration
- Restricted Groups Policy Settings
- Add Users to the Remote Desktop Users Group
- Security Settings for Windows XP Clients
- Updates to Restricted Groups ("Member of") Behavior of User-Defined Local Groups

# Windows Firewall Operations Guide

Updated: March 2, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Firewall Operations Guide provides administering and troubleshooting information for Windows Firewall in Microsoft® Windows Server™ 2003 with Service Pack 1 (SP1).

**In this guide**


- Administering Windows Firewall
- Troubleshooting Windows Firewall

# Administering Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides information about administering Windows Firewall in Microsoft® Windows Server™ 2003 with Service Pack 1 (SP1). It is designed to help you perform scheduled and unscheduled administration tasks for managing and optimizing Windows Firewall. By following the best practices and procedures that are described in this guide, you can increase the effectiveness of Windows Firewall and reduce the need for troubleshooting.

**In this guide**


- Introduction to Administering Windows Firewall
- Managing Windows Firewall
- Optimizing Windows Firewall
- Monitoring Windows Firewall
- Securing Windows Firewall



**Acknowledgments**

Produced by: Microsoft Windows Server User Assistance team

Project Writer: Bill Gruber

Project Editor: Allyson Adley

Technical Reviewers: David Beder, Joseph Davies, Michael Surkan

# Introduction to Administering Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide explains how to administer Windows Firewall. These activities are part of the operating phase of the IT life cycle. If you are not familiar with the Windows Server 2003 administration guides, review the following sections of this introduction.

**When to Use This Guide**
You should use this guide when:

- You want to know how to perform a specific procedure.
- You want to learn about the scheduled and unscheduled maintenance tasks for administering Windows Firewall.
- You want to identify and understand the high-level operations objectives and tasks that are associated with Windows Firewall.

Do not use this guide to diagnose and solve Windows Firewall problems. Step-by-step troubleshooting information can be found in Troubleshooting Windows Firewall.

This guide can be used by organizations that have deployed Windows Server 2003 with Service Pack 1 (SP1). It includes information that is relevant to different roles within an IT organization, including IT administrators and managers who plan and implement the IT operations processes that are required to administer Windows Firewall in Windows Server 2003 SP1.

In addition, this guide provides more detailed procedures for operators who have varied levels of expertise and experience. Although the procedures provide operator guidance from start to finish, operators must have a basic proficiency with the Microsoft Management Console (MMC) and snap-ins, and know how to start administrative programs and access the command line.

**How to Use This Guide**
The operations areas are divided into the following types of content:

- **Objectives** are high-level goals for managing and securing Windows Firewall. Each objective consists of one or more high-level tasks that describe how to accomplish the objective. "Managing Windows Firewall" and "Optimizing Windows Firewall" are high-level objectives. Sometimes, to make the subject more clear, high-level objectives are broken down into lower-level objectives. For example, "Managing Windows Firewall" is an objective that contains several lower-level objectives, such as "Managing Resets, Startup, and Shutdown," "Managing Windows Firewall Profiles," and "Managing IPsec, Multicast, and ICMP Settings."
- **Tasks** are used to group related procedures and provide general guidance for achieving the goals of an objective. For example, "Enabling and Disabling Windows Firewall" and "Restoring Windows Firewall Default Settings" are tasks within an objective called "Managing Resets, Startup, and Shutdown."
- **Procedures** provide step-by-step instructions for completing the task. For example, "Turn Windows Firewall On or Off" and "Turn Windows Firewall On with No Exceptions" are procedures within a task called "Enabling and Disabling Windows Firewall."

If you are an IT manager or administrator who will be delegating tasks to operators within your organization, you will want to:

- Read through the objectives and tasks to determine how to delegate permissions and whether you need to install tools before operators perform the procedures for each task.
- Before assigning tasks to individual operators, ensure that you have all the tools installed where operators can use them.
- When necessary, create “tear sheets” for each task that operators perform within your organization. Cut and paste the task and its related procedures into a separate document and then either print these documents or store them online.

**Terminology Used in This Guide**

**Domain profile**

A set of Windows Firewall settings that are applied when a computer is connected to a network that contains the domain controllers for the domain in which its computer account resides. The domain profile is one of two profiles used by Windows Firewall to apply settings to a computer.

**Exception**

A port, program, or system service that is allowed to receive unsolicited traffic.

**Exceptions list**

A rules store that specifies which ports, programs, or system services are allowed to receive unsolicited traffic. To allow unsolicited incoming traffic through Windows Firewall, you add ports, programs, and system services to the exceptions list.

**Notification**

A **Windows Security Alert** dialog box that appears when a program that is not listed in the exceptions list attempts to listen for unsolicited traffic on a port. If you have administrative rights on a computer, the notification displays the option to add the program to the exceptions list, and thereby allow the program to receive unsolicited traffic. If you do not have administrative rights on a computer, the notification displays a warning that a program is attempting to listen on a port.

A **Windows Security Alert** dialog box does not appear when a system service attempts to listen for unsolicited incoming traffic; it appears only when a program attempts to listen for unsolicited incoming traffic.

**Port**

A logical communication endpoint representing a service or an application that listens for and receives IP packets. Ports are specified by a positive 16-bit decimal number and by the type of traffic (either UDP or TCP) that is expected to pass through the port. Most ports are predefined and considered "well-known" for specific services, such as DHCP and DNS. However, some ports are created dynamically by a server and assigned to the services and applications that need to listen for incoming traffic from clients.

**Program**

A software application that is usually started by a user and runs under the user's account. Programs run only while the user is logged on to a computer and usually consist of one or more executable (.exe) files and one or more dynamic-link library (.dll) files.

| **Note** |
| --- |
| Some system services run within their own .exe file and are started by a user (for example, the Telnet service, which runs in Tlntsvr.exe). Typically, these system services run under a privileged account, such as Local Service. They do not run under the user's account. In Windows Firewall, any system service that runs within its own .exe file is considered a program. |

**Scope**

A Windows Firewall setting that you configure for an exception. The scope setting controls from which addresses unsolicited traffic is allowed to originate. By default, the scope of an exception is any address, which includes any computer on the Internet. You can change the scope of an exception to locally reachable addresses or a list of individual Internet Protocol version 4 (IPv4) addresses or IPv4 address ranges.

**System service**

A software application that is usually not started by a user and runs under a privileged account, such as LocalSystem. System services run even when a user is not logged on to a computer and usually run as a separate process within Svchost.exe.

**Standard profile**

A set of Windows Firewall settings that are applied when a computer is not connected to a network that contains the domain controllers in which its computer account resides (for example, a public network, such as the Internet). The standard profile is one of two profiles used by Windows Firewall to apply settings to a computer.

# Managing Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To implement a successful management plan for Windows Firewall, you must address the following objectives:

Managing Resets, Startup, and Shutdown

Managing Windows Firewall Profiles

Managing Program, Port, and System Service Firewall Rules

Managing IPsec, Multicast, and ICMP Settings

Managing Windows Firewall Notifications

# Best Practices for Managing Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best Practices**

Use the following best practices when you administer Windows Firewall.

Use Security Configuration Wizard to configure Windows Firewall.

Use Group Policy to manage Windows Firewall.

Do not configure Windows Firewall settings on a computer-by-computer basis.

Do not configure per-connection settings.

Create program exceptions instead of port exceptions.

Open the minimum number of ports.

Use scope settings to limit the scope of an exception.

Turn off Windows Firewall if you are using a non-Microsoft host firewall.

**Use Security Configuration Wizard to configure Windows Firewall.**

Security Configuration Wizard (SCW) is the recommended method for configuring Windows Firewall settings in Windows Server 2003. Using SCW ensures that your computer is configured appropriately for the installed roles and enabled services and that no unnecessary ports are opened based on those roles and services. For more information, see Configuring Windows Firewall with SCW.

**Use Group Policy to manage Windows Firewall.**

If your organization uses Group Policy, use the Windows Firewall Group Policy settings to manage and configure Windows Firewall. This will ensure more consistency among your Windows Firewall configurations and will allow you to secure Windows Firewall so that administrators and users who are logged on locally cannot configure Windows Firewall settings. For more information, see Administering Windows Firewall with Group Policy on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43157).

**Do not configure Windows Firewall settings on a computer-by-computer basis.**

This can make your organization more accessible to attack because each computer has a different attack surface. When you configure Windows Firewall settings on a computer-by-computer basis, it becomes difficult to determine which computers are accessible to an attack by a worm or other malicious program that relies on unsolicited incoming traffic to spread.

**Do not configure per-connection settings.**

Per-connection settings (also called connection-specific settings or interface-specific settings) are applied without regard to the profile (standard or domain) used by the computer. This can make a computer more accessible to attack when it is connected to a public network, such as the Internet. You cannot restrict the scope of per-connection exceptions; these exceptions allow unsolicited traffic from any address. For more information, see Configuring Firewall Rules for Specific Connections.

**Create program exceptions instead of port exceptions.**

If you need to allow unsolicited incoming traffic through Windows Firewall, create a program exception instead of a port exception. When you create a program exception, Windows Firewall dynamically opens ports for the specified program and then closes those ports when the program is shut down. When you create a port exception, the specified port is open all of the time, which makes the computer more accessible to attack. For more information, see Configuring Program Firewall Rules.

**Open the minimum number of ports.**

If you have to open ports to allow unsolicited incoming traffic through Windows Firewall, be sure to open only those ports that you need. If a port handles intermittent or infrequent traffic, open and close the port as needed. If a computer configuration requires you to open numerous ports, you should evaluate the roles that are installed on the computer and the services that the computer is providing. This could indicate poor infrastructure design, and could make the computer a critical point of failure or attack.

**Use scope settings to limit the scope of an exception.**

If you create a program or port exception or enable a system service exception, be sure to configure the scope of the exception. The scope setting restricts an exception to computers that are directly reachable or to specific Internet Protocol version 4 (IPv4) addresses and address ranges. For more information about scope settings, see Configuring Scope Settings.

**Turn off Windows Firewall if you are using a non-Microsoft host firewall.**
Using Windows Firewall with another host firewall does not necessarily increase your security. In fact, it can cause your computer to behave erratically if the rules and policies of both firewalls are not complementary. You should turn off (disable) Windows Firewall if you are running a non-Microsoft host firewall.

**See Also**
**Concepts**
Best Practices for Optimizing Windows Firewall
Best Practices for Monitoring Windows Firewall
Best Practices for Securing Windows Firewall

# Managing Resets, Startup, and Shutdown

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because Windows Firewall and the Windows Firewall/Internet Connection Sharing service are disabled by default on Windows Server 2003, you must manage the initial startup of these components. In addition, you might have to manage the shutdown of Windows Firewall and the restoration of Windows Firewall settings.

To help you manage these functions, review Best Practices for Managing Windows Firewall and Known Issues for Managing Resets, Startup, and Shutdown

The following tasks are described in this objective:

- Configuring Windows Firewall with SCW
- Enabling and Disabling Windows Firewall
- Restoring Windows Firewall Default Settings

# Known Issues for Managing Resets, Startup, and Shutdown

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review the following known issues before you manage Windows Firewall resets, startup, and shutdown.

- Make sure the Windows Firewall/Internet Connection Sharing service is running before you start any programs or system services that are listed in the program exceptions list. If you start the Windows Firewall/Internet Connection Sharing service after you have started programs that are listed in the Windows Firewall exceptions list, restart your computer and then start your programs and system services. Windows Firewall cannot track the state of a program's traffic if the program is started before you start the Windows Firewall/Internet Connection Sharing service.

- Do not run Windows Firewall if you are running Routing and Remote Access. Routing and Remote Access uses a firewall that cannot be used with Windows Firewall. For more information, see Basic Firewall on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43159).

- Do not run Windows Firewall on a server that is running a perimeter firewall, such as Microsoft Internet Security and Acceleration (ISA) Server 2004. The protection that Windows Firewall provides in this situation is redundant and unnecessary. In addition, Windows Firewall can cause perimeter firewalls such as ISA Server to function improperly.

- Do not permanently run Windows Firewall on a domain controller. Instead, use Internet Protocol security (IPsec) to protect incoming traffic on a domain controller. You can use Windows Firewall on a temporary basis (for example, to protect a domain controller from attack while you install security updates, virus signatures, or security software).

- Do not use Windows Firewall on a computer that is running a non-Microsoft host firewall. Although it is possible to run two host firewalls on a single computer, it is not recommended that you do so. Host firewall implementations vary widely and there is no guarantee that a non-Microsoft host firewall and Windows Firewall will work well together.

- Do not change the Startup type or status of any service on which Windows Firewall relies. Changing the Startup type or status of a dependent service can prevent Windows Firewall from starting or running properly. For more information about these dependent services, see the section titled "How Windows Firewall Works" in the Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729).

- You might not be able to start Windows Firewall if Administrative Templates from the Windows XP Security Guide were applied to the computer before Windows XP with Service Pack 2 (SP2) was installed. For more information, see article ID 892199 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45337). The article describes why the Windows Firewall/Internet Connection Sharing service might not start and how you can configure the service so it behaves properly when you use Administrative Templates.

- When you reset Windows Firewall, the settings that you specified in Netfw.inf are reapplied to the computer; all existing Windows Firewall settings are deleted. Any Windows Firewall settings that you specified in an answer file, such as Unattend.txt, are deleted and are not restored. If you want to restore the Windows Firewall settings that you configured during installation with an answer file, you must configure the settings manually.

# Configuring Windows Firewall with SCW

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Security Configuration Wizard (SCW) is a tool used to reduce the attack surface of computers running a member of the Windows Server 2003 family with Service Pack 1 (SP1). SCW is the preferred tool for configuring Windows Firewall. You can use SCW to turn on Windows Firewall and configure Windows Firewall exceptions.

SCW does not configure Windows Firewall directly. It helps you create a security policy based on the roles performed by a server. After a policy is created, it can be edited or applied to one or more similarly configured servers. Applied policies can be rolled back to undo changes that have caused problems. To edit, apply, or roll back a security policy, the policy must have been created with SCW.

**When to perform this task**

You should use SCW to configure Windows Firewall when you install Windows Server 2003 with SP1. This includes:

- Installations of SP1 on servers that are running Windows Server 2003.
- Slipstream installations of Windows Server 2003 with SP1 on new servers.
- Upgrades from older operating systems to Windows Server 2003 with SP1.

Do not use SCW to configure Windows Firewall if you use Netfw.inf or an answer file to configure Windows Firewall during an automated installation or if you use domain-based Group Policy settings to configure Windows Firewall.

**Task requirements**

SCW is an optionally installed component. You must install SCW before you can use it. SCW can be installed only on computers running a member of the Windows Server 2003 family with SP1. The server that you are configuring must be running a member of the Windows Server 2003 family with SP1. For more information, see Install Security Configuration Wizard on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43158).

**Task procedures**

To complete this task, perform the following procedure:

Use SCW to Configure Windows Firewall

**See Also**
**Concepts**
Enabling and Disabling Windows Firewall
Restoring Windows Firewall Default Settings
Known Issues for Managing Resets, Startup, and Shutdown

# Use SCW to Configure Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Security Configuration Wizard (SCW) is an optional component that must be installed through Add or Remove Programs in Control Panel. SCW guides you through the process of creating a security policy, based on the roles performed by a given server. Once a policy is created, it can be edited or applied to one or more similarly configured servers.

SCW displays different pages, depending on the configuration settings you choose. If, during the Role-Based Service Configuration portion of SCW, you choose Remote access/VPN server, Internet Connection Sharing server, or Microsoft Internet Security and Acceleration Server 2004, you will not be able to select the Network Security portion of SCW, which is where Windows Firewall settings are configured.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**To use SCW to configure Windows Firewall**

You must use the graphical user interface to perform this procedure; you cannot use the command prompt or Group Policy to perform this procedure.

**To use SCW to configure Windows Firewall**

1. Open SCW, and follow the steps in the wizard until you reach the **Network Security** page.
2. Clear the **Skip this section** check box, and click **Next**.
3. Under **Select the ports to open**, review the ports that SCW will add to the Windows Firewall exceptions list. To prevent SCW from adding a port to the exceptions list, clear the check box next to the port.

   If you want to add additional ports or applications to the Windows Firewall exceptions list, click **Add**, and do the following:

   - To add a port to the exceptions list, in **Port number**, enter the port number, select **TCP** or **UDP** or both, and click **OK**.
   - To add an application to the exceptions list, click the **Approve Application** tab, and in **Application path**, enter the application path and name of the executable (.exe) file, and then click **OK**. Try to use environment variables for paths so that the administrative templates can be applied to any computer.
4. Click **Next**.
5. On the **Confirm Port Configuration** page, verify the port settings that SCW will use to configure Windows Firewall, and click **Next**.
6. On the **Registry Settings** page, select the **Skip this section** check box, and click **Next**.
7. On the **Audit Policy** page, select the **Skip this section** check box, and click **Next**.
8. On the **Internet Information Services** page, select the **Skip this section** check box, and click **Next**.
9. On the **Save Security Policy** page, click **Next**.
10. Enter a name and description for the security policy, and click **Next**.
11. Click **Apply now** to apply the security policy and configure Windows Firewall, and click **Next**.
12. After SCW has applied the security policy, click **Next**, and then click **Finish**.

**Note**

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Windows Firewall with SCW

# Enabling and Disabling Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To protect your computer with Windows Firewall, you must turn on (enable) Windows Firewall and start the Windows Firewall/Internet Connection Sharing service. If the Windows Firewall/Internet Connection Sharing service is not running and you attempt to start Windows Firewall, in the graphical user interface, a dialog box will be displayed to ask if you want to start the service; if you are using the command prompt, a warning will appear to inform you that the service has not been started.

| **Note** |
| --- |
| The Windows Firewall/Internet Connection Sharing service does not stop when you turn off Windows Firewall. Although unsolicited incoming traffic will not be blocked, as long as the Windows Firewall/Internet Connection Sharing service is still running, Windows Firewall events will be written to the security event log. |

**Turning Windows Firewall on and off**

If you turn on the Windows Firewall/Internet Connection Sharing service after you have started any programs or system services that are listed in the program exceptions list, you should do the following:

- Restart your computer.
- Start the programs that are listed in the exceptions list.

Windows Firewall cannot track the state of a program's traffic or a system service's traffic if the program or system service is started before the Windows Firewall/Internet Connection Sharing service is started. This problem typically occurs when you stop the Windows Firewall/Internet Connections Sharing service to install and configure software, and then turn on the Windows Firewall/Internet Connection Sharing service after you start the software program.

**Turning Windows Firewall on with no exceptions**

When you turn on Windows Firewall, you can also configure a special setting that blocks all traffic that has been added to the exceptions list. This setting appears in the graphical user interface as the **Don't allow exceptions** check box. When you use this setting, all of the exceptions in the Windows Firewall exceptions list are disabled and all unsolicited incoming TCP/IP traffic is blocked.

Using the **Don't allow exceptions** setting will prevent the display a **Windows Security Alert** dialog box when a program or system service attempts to listen for incoming traffic. However, Windows Firewall will continue to write notification events to the security log.

**Turning Windows Firewall on or off for a specific connection**

By default, when you turn Windows Firewall on or off, you are enabling or disabling Windows Firewall on every network connection on the computer. However, you can also turn Windows Firewall on or off on a connection-specific or interface-specific basis. When you do this, Windows Firewall filters traffic for each connection based on the connection-specific settings you configure.

When you turn on Windows Firewall globally, Windows Firewall is turned on for every connection on your computer unless you explicitly turn it off on a per-connection basis. However, if you turn off Windows Firewall globally, Windows Firewall is turned off for every connection on your computer; you cannot turn on Windows Firewall on a per-connection basis. In short, to turn Windows Firewall on or off on a per-connection basis, you must turn on Windows Firewall globally and then turn Windows Firewall on or off for each individual connection. The following table shows how global and per-connection settings are combined.

| Global Setting | Per-Connection Setting | Resultant Per-Connection Setting |
| --- | --- | --- |
| On | On | On |
| On | Off | Off |
| Off | Off | Off |

| | | |
|---|---|---|
| Off | On | Off |

**When to perform this task**

You should turn on Windows Firewall after you install Windows Server 2003 with Service Pack 1 (SP1). This includes:

- Installations of SP1 on servers that are running Windows Server 2003.
- Slipstream installations of Windows Server 2003 with SP1 on new servers.
- Upgrades from older operating systems to Windows Server 2003 with SP1.

| **Note** |
|---|
| You should not turn on Windows Firewall when you are running Routing and Remote Access, a perimeter firewall, such as Microsoft Internet Security and Acceleration (ISA) Server 2004, or a non-Microsoft host firewall. |

When your computer is connected to a public network, such as the Internet, or a nonsecure private network, you can turn on Windows Firewall and use the **Don't allow exceptions** setting to disable all exceptions in the exceptions list. You can also use the **Don't allow exceptions** setting when your computer or organization is being attacked by a malicious program that relies on unsolicited incoming traffic to spread and you need to protect your computer while you install security updates, virus signatures, or security software.

After you have started and configured it, you rarely need to turn off Windows Firewall except to install software or perform troubleshooting.

You usually do not need to configure Windows Firewall on a connection-specific basis, but you might want to if you have a multihomed computer and you do not want Windows Firewall to protect every connection on your computer.

**Task requirements**

No special tools are required to complete this task.

**Task procedures**

To complete this task, perform the following procedures:

Turn Windows Firewall On or Off

Turn Windows Firewall On or Off for a Specific Connection

Turn Windows Firewall On with No Exceptions

**See Also**
**Concepts**
Configuring Windows Firewall with SCW
Restoring Windows Firewall Default Settings
Known Issues for Managing Resets, Startup, and Shutdown

# Turn Windows Firewall On or Off

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to turn on (enable) or turn off (disable) Windows Firewall.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

**To turn Windows Firewall on or off**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To turn Windows Firewall on or off**

- Open Windows Firewall.

  To turn on Windows Firewall, click **On**, and then click **OK**.

  To turn off Windows Firewall, click **Off**, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To turn Windows Firewall on or off**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set opmode mode =** *mode*

Substitute values for the placeholder in italics. The following table lists possible values for the placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *mode* | enable, disable | Specifies whether to turn Windows Firewall on or off. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Enabling and Disabling Windows Firewall
Turn Windows Firewall On with No Exceptions
Turn Windows Firewall On or Off for a Specific Connection
Known Issues for Managing Resets, Startup, and Shutdown

# Turn Windows Firewall On or Off for a Specific Connection

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to turn on (enable) or turn off (disable) Windows Firewall on a per-connection basis. This is useful if your computer has multiple network connections (sometimes referred to as interfaces) and you do not want to protect every network connection with Windows Firewall.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

**To turn Windows Firewall on or off for a specific connection**

This procedure can be performed using the graphical user interface or the command prompt. You cannot use Group Policy to turn Windows Firewall on or off on a per-connection basis.

**Using the graphical user interface**

**To turn Windows Firewall on or off for a specific connection**

1. Open Windows Firewall.
2. Click the **Advanced** tab.

   To turn off Windows Firewall for a connection, in **Network Connection Settings**, clear the check box next to the connection that you do not want to protect with Windows Firewall, and then click **OK**.

   To turn on Windows Firewall for a connection, in **Network Connection Settings**, select the check box next to the connection that you want to protect with Windows Firewall, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To turn Windows Firewall on or off for a specific connection**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set opmode mode** = *mode* **interface = "** *interface* **"**

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *mode* | enable, disable | Specifies whether to turn Windows Firewall on or off. |
| *interface* | The name of any network connection that can exist on the computer, as displayed in the Network Connections folder. | Specifies the interface or connection on which you want to enable or disable Windows Firewall. You must enclose *interface* in quotation marks. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows

Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Enabling and Disabling Windows Firewall
Turn Windows Firewall On or Off
Turn Windows Firewall On with No Exceptions
Known Issues for Managing Resets, Startup, and Shutdown

# Turn Windows Firewall On with No Exceptions

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to turn on Windows Firewall and prevent all exceptions from taking effect. This procedure is useful when you are connected to a public network, such as the Internet, or a nonsecure private network. When you perform this procedure, all of the exceptions in the exceptions list are disabled and all unsolicited incoming TCP/IP traffic is blocked.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

**To turn on Windows Firewall with no exceptions**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To turn Windows Firewall on with no exceptions**

1. Open Windows Firewall.
2. Click **On**, select the **Don't allow exceptions** check box, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To turn Windows Firewall on with no exceptions**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set opmode mode = enable exceptions = disable**

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

**See Also**
**Concepts**

Enabling and Disabling Windows Firewall
Turn Windows Firewall On or Off
Turn Windows Firewall On or Off for a Specific Connection
Known Issues for Managing Resets, Startup, and Shutdown

# Restoring Windows Firewall Default Settings

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you restore Windows Firewall settings, all of the changes you have made to Windows Firewall are undone (deleted) and all Windows Firewall settings are reset to their original configuration. All of the exceptions in the Windows Firewall exceptions list are deleted. All optional settings, such as log file settings, Internet Control Message Protocol (ICMP) settings, and the notification setting, are reset to their default configuration. Restoring the default settings will turn on Windows Firewall for all network connections, which might cause some programs to stop working. This ensures that your computer is protected by Windows Firewall when you restore default settings. Restoring the default settings might also cause Internet Connection Sharing and Network Bridge to stop working correctly.

| **Note** |
| --- |
| If you used a Netfw.inf file to automate the configuration of Windows Firewall settings, the settings that you specified in Netfw.inf are reapplied to the computer when you restore Windows Firewall default settings. Any Windows Firewall settings that you specified in an answer file, such as Unattend.txt, are deleted and are not restored. If you want to restore the Windows Firewall settings that you configured during installation with an answer file, you must configure the settings manually. |

### When to perform this task

You might need to restore Windows Firewall default settings when you change a server's role or reconfigure a server. In this case, it might be easier to configure Windows Firewall from its default configuration instead of manually rolling back all of the configuration changes you have applied to the computer.

### Task requirements

No special tools are required to complete this task.

### Task procedures

To complete this task, perform the following procedure:

Restore Windows Firewall Default Settings

**See Also**
**Concepts**
Configuring Windows Firewall with SCW
Enabling and Disabling Windows Firewall
Known Issues for Managing Resets, Startup, and Shutdown

# Restore Windows Firewall Default Settings

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to reset Windows Firewall to its default configuration.

| **Note** |
| --- |
| Windows Firewall will be turned on when you perform this procedure. This ensures that your computer is protected by Windows Firewall when you restore default settings. |

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

### Special Considerations

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

#### To restore Windows Firewall default settings

This procedure can be performed using the graphical user interface or the command prompt. You cannot use Group Policy to restore Windows Firewall default settings.

### Using the graphical user interface

### To restore Windows Firewall default settings

1. Open Windows Firewall.
2. Click the **Advanced** tab, click **Restore Defaults**, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

### Using the command prompt

### To restore Windows Firewall default settings

- Type the following at the command prompt, and press ENTER:

  **netsh firewall reset**

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

### Notes

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Restoring Windows Firewall Default Settings
Known Issues for Managing Resets, Startup, and Shutdown

# Managing Windows Firewall Profiles

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you are configuring Windows Firewall settings, you must know which profile (standard or domain) Windows Firewall is using and how to specify a Windows Firewall profile. Both profiles contain the same set of Windows Firewall settings. However, you typically configure standard profile settings to be more restrictive than domain profile settings because the standard profile settings do not need to include applications and services that are used in a managed domain environment.

When you configure Windows Firewall in Control Panel, you are configuring settings in the profile that is currently in use. You cannot use Windows Firewall in Control Panel to configure settings in the profile that is not currently in use. When you configure Windows Firewall Group Policy settings, you can configure individual policy settings in either the domain profile or the standard profile. The Group Policy Object Editor shows you which profile you are configuring. Likewise, when you use the **netsh firewall** commands to configure Windows Firewall settings, you can choose to configure settings in either the domain profile or the standard profile; however, if you do not specify a profile, the settings are applied to the profile that is currently in use.

| **Note** |
| --- |
| If you do not configure standard profile settings, the default values for standard profile settings are applied whenever Windows Firewall uses the standard profile. These settings might not be appropriate when a computer is connected to a public network like the Internet. It is highly recommended that you configure both domain and standard profile settings and enable Windows Firewall for both profiles. |

For more information about the way Windows Firewall determines which profile to use, see the section titled "How Windows Firewall Works" in the Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729).

### When to perform this task

You should determine which profile Windows Firewall is using whenever you configure Windows Firewall settings through the graphical user interface or command prompt.

You should specify which profile you are configuring whenever you use the command prompt or the Group Policy Object Editor.

### Task requirements

No special tools are required to complete this task.

### Task procedures

To complete this task, perform the following procedure:

Determine Which Profile Windows Firewall Is Using

Specify Which Profile You Are Configuring

**See Also**
**Concepts**
Known Issues for Managing Windows Firewall Profiles

# Known Issues for Managing Windows Firewall Profiles

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review the following known issues before you manage Windows Firewall profiles:

- A Windows Firewall Setting Does Not Take Effect
- VPN and Remote Access Connections Fail

# Determine Which Profile Windows Firewall Is Using

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is useful when you are configuring settings with Windows Firewall in Control Panel and you need to know to which profile the settings will be applied. This procedure is also useful if you are configuring settings with the **netsh firewall** command and you have not specified to which profile the settings should be applied.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**To determine which profile Windows Firewall is using**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To determine which profile Windows Firewall is using**

- Open Windows Firewall.

  If Windows Firewall is using the domain profile, you will see **Windows Firewall is using your domain settings** displayed at the bottom of the **General** tab.

  If Windows Firewall is using the standard profile, you will see **Windows Firewall is using your standard settings** displayed at the bottom of the **General** tab.

**Using the command prompt**

**To determine which profile Windows Firewall is using**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall show currentprofile**

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**

**Concepts**

Managing Windows Firewall Profiles
Specify Which Profile You Are Configuring
Known Issues for Managing Windows Firewall Profiles

# Specify Which Profile You Are Configuring

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is useful when you are configuring Windows Firewall settings and you want to configure the settings in a specific profile.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**To specify which profile you are configuring**

This procedure can be performed using the graphical user interface, the command prompt, and Group Policy.

**Using the graphical user interface**

**To specify which profile you are configuring**

1. Open Windows Firewall.
2. If Windows Firewall is not currently using the profile that you want to configure, do one of the following:

   If Windows Firewall is using the domain profile and you want to configure settings in the standard profile, remove the computer from the domain or disconnect the computer from the network. This will force Windows Firewall to use the standard profile.

   If Windows Firewall is using the standard profile and you want to configure settings in the domain profile, join your computer to the domain in which its computer account resides.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To specify which profile you are configuring**

- Use the following **netsh firewall** commands:

  **netsh firewall set allowedprogram profile =** *profile*

  **netsh firewall set icmpsetting profile =** *profile*

  **netsh firewall set multicastbroadcastresponse profile =** *profile*

  **netsh firewall set notifications profile =** *profile*

  **netsh firewall set opmode profile =** *profile*

  **netsh firewall set portopening profile =** *profile*

  **netsh firewall set service profile =** *profile*

The *profile* parameter can have any of the following values:

| Value | Description |
|---|---|
| *current* | Specifies that you want to configure settings in the profile that Windows Firewall is currently using. |
| *domain* | Specifies that you want to configure settings in the domain profile. |

| | |
|---|---|
| *standard* | Specifies that you want to configure settings in the standard profile. |
| *all* | Specifies that you want to configure settings in the standard and domain profiles. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Using Group Policy**

**To specify which profile you are configuring**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, and then open **Windows Firewall**.
3. To configure settings in the domain profile, open **Domain Profile**. To configure settings in the standard profile, open **Standard Profile**.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Managing Windows Firewall Profiles
Determine Which Profile Windows Firewall Is Using
Known Issues for Managing Windows Firewall Profiles

# Managing Program, Port, and System Service Firewall Rules

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Because Windows Firewall blocks all incoming unsolicited TCP/IP traffic by default, you might need to configure program, port, and system service exceptions for programs or system services that are acting as servers, listeners, or peers. Program, port, and system service exceptions must be managed on an ongoing basis as your server roles or configurations change.

To help you manage Windows Firewall program, port, and system service exceptions, review Best Practices for Managing Windows Firewall and Known Issues for Managing Firewall Rules.

The following tasks are described in this objective:

- Configuring Program Firewall Rules
- Configuring Port Firewall Rules
- Configuring System Service Firewall Rules
- Configuring Firewall Rules for Specific Connections
- Configuring Scope Settings

# Known Issues for Managing Firewall Rules

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review the following known issues before you manage Windows Firewall exceptions.

- Do not use the **My network (subnet) only** scope setting unless you understand how this scope setting restricts the scope of an exception. For more information, see Configuring Scope Settings.

- Verify that you are configuring exceptions in the correct profile.

- You might have to enable Internet Control Message Protocol (ICMP) settings if you want to use the **ping** command. By default, Windows Firewall blocks all incoming ICMP traffic. For more information about enabling the **ping** command, see The Ping Command Times Out.

- Windows Firewall does not display a **Windows Security Alert** dialog box (notification) when a system service attempts to listen for unsolicited incoming traffic. You must use the security event log to identify system services that attempt to listen for unsolicited incoming traffic.

- If a computer is running Windows Server 2003 and the preconfigured Remote Desktop exception is enabled on a per-connection basis, the Remote Desktop exception will be enabled globally after you apply Service Pack 1 (SP1).

- You can use program exceptions to allow unsolicited incoming traffic through Windows Firewall only if the program uses Windows Sockets (Winsock) to create port assignments. If a program does not use Winsock to assign ports, you must determine which ports the program uses and add those ports to the exceptions list.

- You cannot use the **Custom list** scope setting to specify Internet Protocol version 6 (IPv6) addresses. You must use the **Any computer (including those on the Internet)** or **My network (subnet) only** settings to change the scope of an exception that receives IPv6 traffic.

- When you upgrade from Windows Server 2000 to Windows Server 2003, two shared fax folders, Faxclient$ and Fxssrvcp$, are created. The presence of these shared folders causes Windows Firewall to enable the File and Printer Sharing exception and the **Allow incoming echo request** setting.

- Some remote administration tools require you to configure exceptions on the server that is being managed and on the client that is being used to manage the server. For more information about remote administration tools and exceptions, see article ID 840634 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45342). This article describes why Windows Firewall might prevent you from remotely managing a computer and how you can configure Windows Firewall to overcome this problem.

- Some programs might stop working properly when you use Windows Firewall. For more information about the exceptions that are required for popular programs, see article ID 842242 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45343). The article describes why some program might not work properly with Windows Firewall and how you can configure Windows Firewall to overcome the problem.

- If you create an exception by modifying the registry, the exception might not show up in the Windows Firewall graphical user interface. In addition, you will need to stop and then start the Windows Firewall/Internet Connection Sharing service.

- If you modify the registry to create an exception and do not provide a name for the exception, the exception might not show up in the Windows Firewall graphical user interface. To see all exceptions, use the **netsh firewall show state verbose=enable** command. For more information, see article ID 897663 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=47734). The article describes why an exception might not show up in the graphical user interface and how you can configure the registry to overcome the problem.

# Configuring Program Firewall Rules

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you turn on Windows Firewall for the first time or restore Windows Firewall default settings, all unsolicited incoming TCP/IP traffic is blocked on all network connections. This means that any program or system service that attempts to listen for traffic on a TCP or UDP port will be unable to receive traffic. To allow programs and system services to receive unsolicited traffic through these ports, you must add the program or system service to the Windows Firewall exceptions list.

| **Note** |
| --- |
| In some cases, if you cannot add a program or system service to the exceptions list, you must determine which port or ports the program or system service uses and add the port or ports to the Windows Firewall exceptions list. |

To add a program or system service to the exceptions list, you must specify the executable (.exe) file used by the program or system service. A system service that runs within its own unique .exe file and is not hosted by a service container, such as Svchost.exe, is considered to be a program and can be added to the exceptions list. In the same way, a program that behaves like a system service and runs no matter if a user is logged on to the computer is also considered a program as long as it runs within its own unique .exe file. You cannot add a system service to the exceptions list unless the system service runs within its own unique .exe file. Also, you should not add service containers or programs that host services, such as Svchost.exe, Dllhost.exe, and Inetinfo.exe, to the exceptions list.

When you add a program to the exceptions list, Windows Firewall dynamically opens (unblocks) and closes (blocks) the ports required by the program. When the program is running and listening for incoming traffic, Windows Firewall opens the required ports; when the program is not running or is not listening for incoming traffic, Windows Firewall closes the ports. Because of this dynamic behavior, adding programs to the exceptions list is the recommended method for allowing unsolicited incoming traffic through Windows Firewall.

| **Note** |
| --- |
| You can use program exceptions to allow unsolicited incoming traffic through Windows Firewall only if the program uses Windows Sockets (Winsock) to create port assignments. If a program does not use Winsock to assign ports, you must determine which ports the program uses and add those ports to the exceptions list. |

In addition to adding program exceptions, you can also edit and delete program exceptions. Editing a program exception allows you to change the path or file name that is associated with the program and configure scope settings for the exception. Deleting a program exception removes the exception from the exceptions list and prevents the program from receiving unsolicited incoming traffic (unless a port exception or some other exception allows unsolicited incoming traffic to reach the program).

**Mitigating the Risks Associated with Exceptions**

Each time you add a program to the exceptions list, you make your computer more accessible to attack. A common form of network attack uses port scanning software to identify computers that have open and unprotected ports. By adding numerous programs to the exceptions list, you defeat the purpose of a firewall and increase the attack surface of your computer. This problem typically occurs when you configure a server for several different roles and you must open numerous ports to accommodate each of the server roles. You should closely evaluate the design of any server that requires you to open numerous ports. Servers that are configured for numerous roles or to provide numerous services can be a critical point of failure in your organization and might indicate poor infrastructure design.

To decrease your security risk, follow these guidelines when you configure program exceptions:

- **Create an exception only when you need it.** If you think a program might require a port for unsolicited incoming traffic, do not add the program to the exceptions list until you verify that the program attempted to listen for unsolicited traffic. By default, Windows Firewall displays a notification when a program attempts to listen for unsolicited traffic.

- **Never create an exception for a program that you do not recognize.** If Windows Firewall notifies you that a program has attempted to listen for unsolicited traffic, verify the name of the program and the .exe file before you add the program to the exceptions list.

- **Remove an exception when you no longer need it.** If you add a program to the exceptions list on a server and then change the server's role or reconfigure the services and applications on the server, be sure to update the exceptions list and remove exceptions that are no longer required.

### When to perform this task

You should configure a program exception when you can identify the .exe file for a program and you know that a program must receive unsolicited incoming traffic. You typically perform this task on an ongoing basis as your server roles and server configurations change.

### Task requirements

No special tools are required to complete this task.

### Task procedures

You can create most program exceptions by using the Windows Firewall notification feature. By default, Windows Firewall displays a **Windows Security Alert** dialog box (notification) whenever a program attempts to listen for incoming traffic and the incoming traffic is blocked. If you are a member of the Administrators group (or a member of a group that is a member of the Administrators group), the **Windows Security Alert** dialog box provides a **Keep Blocking** option, which adds the program to the exceptions list but does not enable the exception and an **Unblock** option, which adds the program to the exceptions list and enables the exception so that unsolicited incoming traffic can reach the program. If you are not a member of the Administrators group, a notification appears, but it does not display the options to unblock or continue to block the program. The notification is informational only.

| **Note** |
| --- |
| Many program exceptions are also created programmatically through the Windows Firewall application programming interfaces (APIs). Programs that use the Windows Firewall APIs can add themselves to the exceptions list without any user notification or input. |

The notification feature is useful if you do not know which programs are acting as servers, listeners, or peers (that is, are listening for unsolicited incoming traffic). However, you might have to use other methods to configure program exceptions if:

- You disable the notification feature.
- You are not logged on to a computer as an administrator when you see the **Windows Security Alert** dialog box.
- A program runs as a system service even though it runs within its own .exe file.
- You turn on Windows Firewall with the **Don't allow exceptions** option.

If you know that a program acts as a server, listener, or peer and you know the program's .exe file name and the path to the .exe file, use this procedure to add the program to the exceptions list:

Add a Program to the Firewall Rules List

If you know that a program acts as a server, listener, or peer, but you do not know the program's .exe file name, you might be able to find this information in the Windows Firewall Settings Technical Reference. For more information, see Windows Firewall Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43155).

If you do not know which programs act as servers, listeners, or peers, use the following procedure to identify programs that listen for unsolicited incoming traffic, and then add the program to the exceptions list:

Identify Blocked Servers, Listeners, and Peers

To edit or delete an existing program exception, use the following procedure:

Edit or Delete a Program Firewall Rule

**See Also**
**Concepts**
Known Issues for Managing Firewall Rules

Configuring Port Firewall Rules
Configuring System Service Firewall Rules
Configuring Firewall Rules for Specific Connections
Configuring Scope Settings

# Add a Program to the Firewall Rules List

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is useful when a program or system service runs within an executable (.exe) file and you want the program or system service to receive unsolicited incoming traffic.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should configure scope settings for any exceptions that you create or enable. For more information about scope settings, see Configuring Scope Settings.

**To add a program to the exceptions list**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To add a program to the exceptions list**

1. Open Windows Firewall.
2. Click the **Exceptions** tab, and then click **Add program**.

   If the program that you want to add to the exceptions list appears in **Programs**, click the program, and then click **OK**.

   If the program that you want to add to the exceptions list does not appear in **Programs**, click **Browse**, and then do the following:

   In the **Browse** dialog box, find the .exe file for the program that you want to add to the exceptions list, click **Open**, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To add a program to the exceptions list**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set allowedprogram program =** *program* **name =** *name* **mode = enable**

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *program* | Path and file name for the .exe file. | Specifies the program to add to the exceptions list. |

| | | |
|---|---|---|
| *name* | Any string less than 256 characters. | Specifies the friendly name for the exception, which is displayed in the graphical user interface. You must enclose *name* in quotation marks. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Program Firewall Rules
Known Issues for Managing Firewall Rules
Identify Blocked Servers, Listeners, and Peers
Identify Unblocked Servers, Listeners, and Peers
Edit or Delete a Program Firewall Rule

# Edit or Delete a Program Firewall Rule

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is useful when you change or delete the ports used by a program or system service and when you want to update the program settings in Windows Firewall.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should verify scope settings for any exceptions that you change. For more information about scope settings, see Configuring Scope Settings.

**To edit or delete a program exception**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To edit or delete a program exception**

1. Open Windows Firewall.
2. Click the **Exceptions** tab, and then click the program exception that you want to edit or delete.

   If you want to delete the program exception, click **Delete**, and then click **Yes**.

   If you want to edit the program exception, click **Edit**, make the changes you want to the program exception settings, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To edit or delete a program exception**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set allowedprogram program =** *program* **name =** *name* **mode = enable**

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *program* | Path and file name for the executable (.exe) file. | Specifies the program to add to the exceptions list. |
| *name* | Any string less than 256 characters. | Specifies the friendly name for the exception, which is displayed in the graphical user interface. You must enclose *name* in quotation marks. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows

Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Program Firewall Rules
Known Issues for Managing Firewall Rules
Add a Program to the Firewall Rules List
Identify Unblocked Servers, Listeners, and Peers
Identify Blocked Servers, Listeners, and Peers

# Identify Blocked Servers, Listeners, and Peers

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is useful when you do not know which programs act as servers, listeners, or peers, and you want to add those programs to the exceptions list.

**Administrative Credentials**

You do not need administrative credentials to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should verify scope settings for any exceptions that you change. For more information about scope settings, see Configuring Scope Settings.

**To identify programs that are listening for but cannot receive incoming traffic**

This procedure can be performed using the graphical user interface or the command prompt. The graphical user interface method will help you identify programs that have attempted to listen for incoming traffic. The command prompt method will help you identify programs that are currently listening for incoming traffic.

**Using the graphical user interface**

You can use the security log in Event Viewer to identify programs and system services that attempted to listen for unsolicited incoming traffic but whose incoming traffic was blocked by Windows Firewall. You must first configure auditing so that Windows Firewall events are written to the security log.

Use the following procedure to configure auditing:

Enable Auditing of Windows Firewall Events

After you configure auditing, use the following procedure to view Windows Firewall events in the security log:

View Windows Firewall Events in Event Viewer

Programs that attempted to listen for but were unable to receive incoming traffic have the following attributes:

Type: Failure Audit

Event ID: 861

**Using the command prompt**

You can use the **netsh firewall** command to identify programs that are currently listening for but cannot receive incoming traffic because Windows Firewall is blocking it.

**To identify programs that are listening for but cannot receive incoming traffic**

1. Type the following at the command prompt, and press ENTER:

   **netsh firewall show state verbose = enable**

2. Search the output text for **Ports on which programs want to receive incoming connections**.

   If this section does not appear in the output text, then there are no easily recognizable programs currently listening for unsolicited incoming traffic.

3. For each program listed in **Ports on which programs want to receive incoming connections**, determine whether the program's ports are already open. To do this, search the command line output text for **Ports currently open on all**

**network interfaces**. If the open ports and protocols match the ports and protocols used by the program, then the program is already able to receive unsolicited incoming traffic.

4. If the ports associated with a program are not already open, then the program is listening for but cannot receive unsolicited incoming traffic.

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Program Firewall Rules
Known Issues for Managing Firewall Rules
Identify Unblocked Servers, Listeners, and Peers

# Identify Unblocked Servers, Listeners, and Peers

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is useful when you want to determine which programs are able to receive unsolicited incoming traffic and which ports do not block unsolicited incoming traffic.

**Administrative Credentials**

You do not need administrative credentials to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should verify scope settings for any exceptions that you change. For more information about scope settings, see Configuring Scope Settings.

**To identify unblocked servers, listeners, and peers**
This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To identify unblocked servers, listeners, and peers**

1. Open Windows Firewall.
2. Click the **Exceptions** tab.
3. In **Programs and Services**, identify the program and port exceptions that are enabled. Enabled exceptions have a check mark next to the program or port exception name and indicate unblocked programs and ports.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

| **Note** |
| --- |
| If you create an exception by modifying the registry, the exception might not show up in the Windows Firewall graphical user interface. |

**Using the command prompt**

**To identify unblocked servers, listeners, and peers**

1. Type the following at the command prompt, and press ENTER:

   **netsh firewall show state verbose = enable**

2. Search the output text for **Ports currently open on all network interfaces**. The ports and programs listed in this section are unblocked and represent enabled program or port exceptions.

   If this section does not appear in the output text, then there are no easily recognizable program or port exceptions currently enabled.

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command**

**Prompt**.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Program Firewall Rules
Known Issues for Managing Firewall Rules
Identify Blocked Servers, Listeners, and Peers

# Configuring Port Firewall Rules

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you turn on Windows Firewall for the first time or restore Windows Firewall default settings, all unsolicited incoming TCP/IP traffic is blocked on all network connections. This means that any program or system service that attempts to listen for traffic on a TCP or UDP port will be unable to receive traffic. To allow programs and system services to receive unsolicited traffic through these ports, you must add the program or system service to the Windows Firewall exceptions list. In some cases, if you cannot add a program or system service to the exceptions list, you must determine which port or ports the program or system service uses and add the port or ports to the Windows Firewall exceptions list.

| **Note** |
| --- |
| A system service that runs within its own unique executable (.exe) file and is not hosted by a service container is considered to be a program. Such a system service can be added to the exceptions list. In the same way, a program that behaves like a system service and runs no matter if a user is logged on to the computer is also considered a program as long as it runs within its own unique .exe file. Do not add service containers or programs that host services, such as Svchost.exe, Dllhost.exe, and Inetinfo.exe, to the exceptions list. |

When you add a port to the exceptions list, you must specify the protocol (TCP or UDP) and port number. You cannot specify protocols other than TCP or UDP and you cannot add a port number without specifying either TCP or UDP. (For example, you cannot exclude traffic based on protocol alone.) When you add a TCP or UDP port to the exceptions list, the port is open (unblocked) whenever Windows Firewall is running and no matter if there is a program or system service listening for incoming traffic on the port. For this reason, if you need to allow unsolicited incoming traffic through Windows Firewall, you should create a program exception instead of a port exception. When you add a program to the exceptions list, Windows Firewall dynamically opens and closes the ports required by the program. When the program is running and listening for incoming traffic, Windows Firewall opens the required ports; when the program is not running or is not listening for incoming traffic, Windows Firewall closes the ports.

In addition to adding port exceptions to the exceptions list, you can edit and delete port exceptions. Editing a port exception allows you to change the port name, protocol, port number, and scope settings for the exception. Deleting a port exception closes (blocks) the port and prevents the port from receiving unsolicited traffic (unless another port exception or some other exception allows unsolicited incoming traffic to reach the program).

**Mitigating the Risks Associated with Exceptions**

Each time you add a program, system service, or port to the exceptions list, you make your computer more accessible to attack. A common form of network attack uses port scanning software to identify computers that have open and unprotected ports. By adding numerous programs, system services, and ports to the exceptions list, you defeat the purpose of a firewall and increase the attack surface of your computer. This problem typically occurs when you configure a server for several different roles, and you need to open numerous ports to accommodate each of the server roles. You should closely evaluate the design of any server that requires you to open numerous ports. Servers that are configured for numerous roles or configured to provide numerous services can be a critical point of failure in your organization and might indicate poor infrastructure design.

To decrease your security risk, follow these guidelines when you configure port exceptions:

- **Create an exception only when you need it.** If you think a program might require a port for unsolicited incoming traffic, do not add a port to the exceptions list until you verify that the program attempted to listen for unsolicited traffic. By default, Windows Firewall displays a notification when a program attempts to listen for unsolicited traffic. You can also use the security event log to determine whether a system service has attempted to listen for unsolicited incoming traffic.

- **Never create an exception for a program or system service that you do not recognize.** If Windows Firewall notifies you that a program has attempted to listen for unsolicited incoming traffic, verify the name of the program and the .exe file before you add a port to the exceptions list. Likewise, if you use the security event log to determine that a system service attempted to listen for unsolicited incoming traffic, verify that the system service is a valid component before you add a port to the exceptions list.

- **Remove an exception when you no longer need it.** If you add a port to the exceptions list on a server and then change the server's role or reconfigure the services and applications on the server, be sure to update the exceptions list

and remove port exceptions that are no longer required.

### When to perform this task

You should configure a port exception when you cannot configure a program exception and you know that a program or system service must receive unsolicited incoming traffic. You typically perform this task on an ongoing basis as your server roles and server configurations change.

### Task requirements

No special tools are required to complete this task.

### Task procedures

If you know that a program or system service acts as a server, listener, or peer and you know which ports are used by the program or system service, use this procedure to add the ports to the exceptions list:

Add a Port to the Firewall Rules List

If you know that a program or system service acts as a server, listener, or peer, but you do not know the ports that the program or system service uses, you might be able to find this information in the Windows Firewall Settings Technical Reference. For more information, see Windows Firewall Settings on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43155).

If you do not know which programs or system services act as servers, listeners, or peers, use the following procedure to identify the ports used by programs and system services to listen for unsolicited incoming traffic:

Identify Blocked Servers, Listeners, and Peers

To edit or delete an existing program exception, use the following procedure:

Edit or Delete a Port Firewall Rule

**See Also**
**Concepts**
Known Issues for Managing Firewall Rules
Configuring Program Firewall Rules
Configuring System Service Firewall Rules
Configuring Firewall Rules for Specific Connections
Configuring Scope Settings

# Add a Port to the Firewall Rules List

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to add a TCP or UDP port to the Windows Firewall exceptions list. When you add a port to the exceptions list, the port is always open; unsolicited incoming traffic is always allowed to pass through the port unless you select the **Don't allow exceptions** option when you turn on Windows Firewall. This procedure is useful when you know the port on which a program or service listens for incoming traffic and you want to run the program or service on your computer.

| **Note** |
| --- |
| Adding a port to the exceptions list can reduce the security of your computer because the port will be open any time the computer is running. You should add ports to the exceptions list only when it is not possible to add a program to the exceptions list. |

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

### Special Considerations

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should configure scope settings for any exceptions that you create or enable. For more information about scope settings, see Configuring Scope Settings.

#### To add a port to the exceptions list

This procedure can be performed using the graphical user interface or the command prompt.

### Using the graphical user interface

### To add a port to the exceptions list

1. Open Windows Firewall.
2. Click the **Exceptions** tab, and then click **Add Port**.
3. In **Name**, type a name for the port exception. This name will appear in the exceptions list.
4. In **Port number**, type the number of the port that you want to add to the exceptions list.
5. Click either **TCP** or **UDP**, depending on the type of traffic, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

### Using the command prompt

### To add a port to the exceptions list

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set portopening protocol** = *protocol* **port** = *port* **name** = *name* **mode** = **enable**

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *protocol* | TCP, UDP, All | Specifies the protocol for the port. Use All to specify both TCP and UDP. |
| *port* | Any number from 1 to 65,535 | Specifies the port number. |
| *name* | Any string less than 256 characters | Specifies the friendly name of the port exception, which is displayed in the graphical user interface. You must enclose *name* in quotation marks. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Port Firewall Rules
Known Issues for Managing Firewall Rules
Identify Unblocked Servers, Listeners, and Peers
Identify Blocked Servers, Listeners, and Peers
Edit or Delete a Port Firewall Rule

# Edit or Delete a Port Firewall Rule

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to change or delete a port exception in the Windows Firewall exceptions list. This procedure is useful when you change or delete the ports used by a program or system service and you want to update the port settings in Windows Firewall.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should verify scope settings for any exceptions that you change. For more information about scope settings, see Configuring Scope Settings.

**To edit or delete a port exception**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To edit or delete a port exception**

1. Open Windows Firewall.
2. Click the **Exceptions** tab, and then click the port exception that you want to edit or delete.

   If you want to delete the port exception, click **Delete**, and then click **OK**.

   If you want to edit the port exception, click **Edit**, make changes to the port exception settings, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To edit or delete a port exception**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set portopening protocol** = *protocol* **port** = *port* **name** = *name* **mode** = *mode*

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *protocol* | TCP, UDP, All | Specifies the protocol for the port. Use All to specify both TCP and UDP. |
| *port* | Any number from 1 to 65,535 | Specifies the port number. |
| *name* | Any string less than 256 characters | Specifies the friendly name of the port exception, which is displayed in the graphical user interface. You must enclose *name* in quotation marks. |

| | | |
|---|---|---|
| *mode* | enable, disable | Specifies whether to allow incoming traffic into the port. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Port Firewall Rules
Known Issues for Managing Firewall Rules
Add a Port to the Firewall Rules List
Add a Port to the Firewall Rules List for a Specific Connection

# Configuring System Service Firewall Rules

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you turn on Windows Firewall for the first time or restore Windows Firewall default settings, all unsolicited incoming TCP/IP traffic is blocked on all network connections. This means that any program or system service that attempts to listen for traffic on a TCP or UDP port will be unable to receive traffic. To allow programs and system services to receive unsolicited traffic through these ports, you must add the program or system service to the Windows Firewall exceptions list.

| **Note** |
| --- |
| If you cannot add a program or system service to the exceptions list, you must determine which port or ports the program or system service uses and add the port or ports to the Windows Firewall exceptions list. However, adding programs and system services to the exceptions list is the recommended way to control the traffic that is allowed through Windows Firewall. |

Windows Firewall provides four preconfigured system service exceptions that you can enable or disable. When you enable a preconfigured exception, Windows Firewall adds the appropriate programs and ports to the exceptions list so that the system service can receive unsolicited incoming traffic. When you disable a preconfigured system service, Windows Firewall deletes the programs and ports from the exceptions list. The following table lists the preconfigured exceptions you can configure in Windows Server 2003.

| Exception | Description |
| --- | --- |
| File and Printer Sharing | Adds TCP ports 139 and 445 and UDP ports 137 and 138 to the exceptions list. This setting allows a computer to receive unsolicited traffic to shared files, folders, and printers. |
| Remote Desktop | Adds TCP port 3389 to the exceptions list. This setting allows a computer to be managed remotely with the Remote Desktop Connection feature. |
| UPnP Framework | Adds TCP port 2869 and UDP port 1900 to the exceptions list. This setting allows a computer to receive UPnP discovery requests from other computers and devices. |
| Remote Administration | Adds TCP ports 135 and 445 to the exceptions list. Also adds Svchost.exe and Lsass.exe to the exceptions list to allow hosted services to open additional, dynamically-assigned ports, typically in the range of 1024 to 1034. This setting allows a computer to be remotely managed with administrative tools, such as the Microsoft Management Console (MMC) and Windows Management Instrumentation (WMI). It also allows a computer to receive unsolicited incoming Distributed Component Object Model (DCOM) and remote procedure call (RPC) traffic. |

| **Note** |
| --- |
| There is no predefined Remote Assistance exception in Windows Server 2003. If you want to use Remote Assistance on Windows Server 2003, you must enable the Remote Desktop exception or add TCP port 3389 to the Windows Firewall exceptions list. |

In most scenarios, all of the preconfigured exceptions are disabled by default. However, the File and Printer Sharing exception might be enabled by default after you perform an upgrade from an older operating system that has shared folders or printers. The Remote Desktop exception might be enabled by default after you perform an upgrade from an older operating system that has Remote Desktop Connection enabled.

In addition to enabling and disabling preconfigured exceptions, you can edit the File and Printer Sharing, Remote Desktop, and UPnP Framework exceptions. Editing a preconfigured exception allows you to enable and disable the programs and ports that are associated with the preconfigured exception. This is useful for troubleshooting and for those cases in which you want to modify a preconfigured exception to suit a specific server configuration but do not want to add additional exceptions to the exceptions list. You cannot delete any of the preconfigured exceptions; disabling a preconfigured exception does not remove it from the exceptions list.

| Note |
| --- |
| You cannot edit the Remote Administration exception nor can you enable or disable it in the graphical user interface. You must use the **netsh firewall** command or Group Policy to enable or disable the Remote Administration exception. |

**Mitigating the Risks Associated with Exceptions**

Each time you add a program, system service, or port to the exceptions list, you make your computer more accessible to attack. A common form of network attack uses port scanning software to identify computers that have open and unprotected ports. By adding numerous programs, system services, and ports to the exceptions list, you defeat the purpose of a firewall and increase the attack surface of your computer. This problem typically occurs when you configure a server for several different roles, and you need to open numerous ports to accommodate each of the server roles. You should closely evaluate the design of any server that requires you to open numerous ports. Servers that are configured for numerous roles or are configured to provide numerous services can be a critical point of failure in your organization and might indicate poor infrastructure design.

To decrease your security risk, follow these guidelines when you configure preconfigured exceptions:

- **Enable an exception only when you need it.** If you think a program might require a port for unsolicited incoming traffic, do not enable a preconfigured exception until you verify that the program attempted to listen for unsolicited traffic and that the ports it uses are those that are opened by the preconfigured exception.

- **Never enable an exception for a program or system service that you do not recognize.** If Windows Firewall notifies you that a program has attempted to listen for unsolicited incoming traffic, verify the name of the program and the executable (.exe) file before you enable a preconfigured exception. If you use the security event log to determine that a system service attempted to listen for unsolicited incoming traffic, verify that the system service is a valid component before you enable a preconfigured exception.

- **Disable preconfigured exceptions when you no longer need them.** If you enable a preconfigured exception on a server, and then change the server's role or reconfigure the services and applications on the server, be sure to update the exceptions list and disable the preconfigured exceptions that are no longer required.

### When to perform this task

You should configure the File and Printer Sharing, Desktop Connection, and UPnP Framework exceptions when your server provides or uses file and printer sharing, Remote Desktop Connection, or UPnP discovery. These exceptions are preconfigured for these services and features and are not meant to be used in any other way.

You should configure the Remote Administration exception on a server when you want to administer the server with a remote administration tool that uses RPC and DCOM. Malicious users often attempt to attack networks and computers using RPC and DCOM. You should contact the manufacturer of your remote administration tool to determine if it requires RPC and DCOM communication. If it does not, do not enable this exception.

### Task requirements

No special tools are required to complete this task.

### Task procedures

To complete this task, perform the following procedures:

Enable or Disable the File and Printer Sharing Firewall Rule

Enable or Disable the Remote Desktop Firewall Rule

Enable or Disable the UPnP Framework Firewall Rule

Enable or Disable the Remote Administration Firewall Rule

**See Also**

**Concepts**

Known Issues for Managing Firewall Rules
Configuring Program Firewall Rules
Configuring Port Firewall Rules
Configuring Firewall Rules for Specific Connections

# Enable or Disable the File and Printer Sharing Firewall Rule

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to open (enable) or close (disable) the TCP and UDP ports that are required for file and printer sharing. This procedure is useful when you share a folder or printer that you want other users to access.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should configure scope settings for any exceptions that you create or enable. For more information about scope settings, see Configuring Scope Settings.

**To enable or disable the File and Printer Sharing exception**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To enable or disable the File and Printer Sharing exception**

1. Open Windows Firewall.
2. Click the **Exceptions** tab.

   To enable the File and Printer Sharing exception, in **Programs and Services**, select the **File and Printer Sharing** check box, and click **OK**.

   To disable the File and Printer Sharing exception, in **Programs and Services**, clear the **File and Printer Sharing** check box, and click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To enable or disable the File and Printer Sharing exception**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set service type = fileandprint mode =** *mode*

Substitute values for the placeholder in italics. The following table lists possible values for the placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *mode* | enable, disable | Specifies whether to enable or disable the exception. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring System Service Firewall Rules
Known Issues for Managing Firewall Rules

# Enable or Disable the Remote Desktop Firewall Rule

Updated: July 9, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to open (enable) or close (disable) the TCP port required for Remote Desktop connections. This procedure is useful when you want to use Remote Desktop to connect to another computer.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should configure scope settings for any exceptions that you create or enable. For more information about scope settings, see Configuring Scope Settings.

**To enable or disable the Remote Desktop exception**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To enable or disable the Remote Desktop exception**

1. Open Windows Firewall.
2. Click the **Exceptions** tab.

   To enable the remote desktop exception, in **Programs and Services**, select the **Remote Desktop** check box, and click **OK**.

   To disable the remote desktop exception, in **Programs and Services**, clear the **Remote Desktop** check box, and click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To enable or disable the Remote Desktop exception**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set service type = remotedesktop mode =** *mode*

Substitute values for the placeholder in italics. The following table lists possible values for the placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *mode* | enable, disable | Specifies whether to enable or disable the exception. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring System Service Firewall Rules
Known Issues for Managing Firewall Rules

# Enable or Disable the UPnP Framework Firewall Rule

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to open (enable) or close (disable) the TCP and UDP ports required for UPnP discovery requests from other computers and devices. This procedure is useful when you want to facilitate the discovery and control of networked devices and services.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should configure scope settings for any exceptions that you create or enable. For more information about scope settings, see Configuring Scope Settings.

**To enable or disable the UPnP Framework exception**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To enable or disable the UPnP Framework exception**

1. Open Windows Firewall.
2. Click the **Exceptions** tab.

   To enable the File and Printer Sharing exception, in **Programs and Services**, select the **UPnP Framework** check box, and click **OK**.

   To disable the File and Printer Sharing exception, in **Programs and Services**, clear the **UPnP Framework** check box, and click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To enable or disable the UPnP Framework exception**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set service type = upnp mode =** *mode*

Substitute values for the placeholder in italics. The following table lists possible values for the placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *mode* | enable, disable | Specifies whether to enable or disable the exception. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring System Service Firewall Rules
Known Issues for Managing Firewall Rules

# Enable or Disable the Remote Administration Firewall Rule

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to open (enable) or close (disable) the ports that are used by many remote administration tools, including tools that use Windows Management Instrumentation (WMI), remote procedure calls (RPC), and Distributed Component Object Model (DCOM). This procedure is useful when you are remotely administering a computer and your remote administration tool uses WMI, RPC, or DCOM.

| **Important** |
| --- |
| You should enable the Remote administration exception only if your remote administration tools require RPC and DCOM. Malicious users often attempt to attack networks and computers using RPC and DCOM. It is recommended that you contact the manufacturer of your remote administration tool to determine if it requires RPC and DCOM communication. If it does not, do not enable this exception. |

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

### Special Considerations

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should configure scope settings for any exceptions that you create or enable. For more information about scope settings, see Configuring Scope Settings.

#### To enable or disable the Remote administration exception

This procedure can be performed using the command prompt or Group Policy. You cannot perform this procedure using the graphical user interface.

### Using the command prompt

### To enable or disable the Remote administration exception

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set service type = remoteadmin mode =** *mode*

Substitute values for the placeholder in italics. The following table lists possible values for the placeholder.

| Placeholder | Possible Values | Description |
| --- | --- | --- |
| *mode* | enable, disable | Specifies whether to enable or disable the exception. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

### Using Group Policy

### To enable or disable the Remote administration exception

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.

2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then open either **Domain Profile** or **Standard Profile**, depending on which profile you want to

configure.

3. In the details pane, double-click **Windows Firewall: Allow remote administration exception**.
4. In the **Windows Firewall: Allow remote administration exception properties** dialog box, on the **Settings** tab, click **Enabled** or **Disabled**.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring System Service Firewall Rules
Known Issues for Managing Firewall Rules

# Configuring Firewall Rules for Specific Connections

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you turn on Windows Firewall for the first time or restore Windows Firewall default settings, all unsolicited incoming TCP/IP traffic is blocked on all network connections. This means that any program or system service that attempts to listen for traffic on a TCP or UDP port will be unable to receive traffic. To allow programs and system services to receive unsolicited traffic through these ports, you must add the program or system service to the Windows Firewall exceptions list. In some cases, if you cannot add a program or system service to the exceptions list, you must determine which port or ports the program or system service uses and add the port or ports to the Windows Firewall exceptions list. However, adding programs and system services to the exceptions list is the recommended way to control the traffic that is allowed through Windows Firewall.

| **Note** |
| --- |
| A system service that runs within its own unique .exe file and is not hosted by a service container is considered to be a program and can be added to the exceptions list. In the same way, a program that behaves like a system service and runs no matter if a user is logged on to the computer is considered a program as long as it runs within its own unique .exe file. Do not add service containers or programs that host services, such as Svchost.exe, Dllhost.exe, and Inetinfo.exe, to the exceptions list. |

You usually configure Windows Firewall on a global basis. For example, when you turn on Windows Firewall, Windows Firewall is enabled on all of the network connections on your computer and all network connections that you create on your computer. Likewise, when you create an exception, the exception applies to all network connections on the computer and all network connections that you create on the computer.

You can also configure Windows Firewall on a connection-specific (per-connection) or interface-specific basis. When you configure Windows Firewall on a per-connection basis, you can configure Windows Firewall differently for every network connection on your computer. This is useful if your computer has multiple network adapters or multiple network connections and you do not want Windows Firewall enabled on all connections or you want to open different ports for each network connection.

The most commonly configured per-connection exception is a port exception. Per-connection port exceptions are cumulative (between the global and per-connection exceptions). Be sure you understand how global and per-connection exceptions work before you use them.

When you open a port on a global basis, the port is open on every network connection on your computer; you cannot close the port on a per-connection basis. To open a port on a specific connection only, you must close the port globally and then open the port on the specific connection. To close a port on a specific connection, you must close the port globally and make sure that the port is not open on the connection you are configuring. The following table shows how global and per-connection port exceptions are combined.

| Global Port Exception | Per-Connection Port Exception | Resultant Per-Connection Port Exception |
| --- | --- | --- |
| Enabled (port open) | Enabled (port open) | Enabled (port open) on every connection |
| Enabled (port open) | Disabled (port closed) | Enabled (port open) on every connection |
| Enabled (port open) | Not configured (no per-connection port exception) | Enabled (port open) on every connection |
| Disabled (port closed) | Enabled (port open) | Enabled (port open) |
| Disabled (port closed) | Disabled (port closed) | Disabled (port closed) |
| Disabled (port closed) | Not configured (no per-connection port exception) | Disabled (port closed) |

As with global exceptions, when you add a port to the exceptions list on a per-connection basis, you must specify the protocol (TCP or UDP) and the port number. You cannot specify protocols other than TCP or UDP, and you cannot add a port number without specifying either TCP or UDP. (For example, you cannot exclude traffic based on just a protocol.) When you add a TCP

or UDP port to the exceptions list, the port is open (unblocked) whenever Windows Firewall is running and no matter if there is a program or system service listening for incoming traffic on the port.

In addition, per-connection port exceptions are created and applied without regard to which Windows Firewall profile a computer is using. For example, if your computer is using the domain profile and you configure a port exception for a specific connection, the port exception will be applied when the computer is using the domain profile and when the computer is using the standard profile. Furthermore, you cannot configure scope settings for a per-connection port exception. The scope of a per-connection port exception includes any computer that has access to the specific network connection.

In addition to adding port exceptions to the exceptions list, you can edit and delete port exceptions on a per-connection basis. Editing a port exception allows you to change the exception name, protocol, and port number. Deleting a port exception closes (blocks) the port and prevents the port from receiving unsolicited incoming traffic (unless a port exception or some other exception exists to allow unsolicited incoming traffic to reach the program).

**Mitigating the Risks Associated with Exceptions**
Each time you add a program, system service, or port to the exceptions list, you make your computer more accessible to attack. A common form of network attack uses port scanning software to identify computers that have open and unprotected ports. By adding numerous programs, system services, and ports to the exceptions list, you defeat the purpose of a firewall and increase the attack surface of your computer. This problem typically occurs when you configure a server for several different roles, and you need to open numerous ports to accommodate each role. You should closely evaluate the design of any server that requires you to open numerous ports. Servers that are configured for numerous roles or to provide numerous services can be a critical point of failure in your organization and might indicate poor infrastructure design.

To decrease your security risk, follow these guidelines when you configure port exceptions:

- **Create an exception only when you need it.** If you think a program might require a port for unsolicited incoming traffic, do not add a port to the exceptions list until you verify that the program attempted to listen for unsolicited traffic. You can use the security event log to determine whether a program has attempted to listen for unsolicited incoming traffic.

- **Remove an exception when you no longer need it.** If you add a port to the exceptions list on a server and then change the server's role or reconfigure its services and applications, be sure to update the exceptions list and remove those port exceptions that are no longer required.

### When to perform this task

You should configure a per-connection port exception when you want to allow unsolicited incoming traffic through Windows Firewall on a specific connection. You should do this only if you know that a program or system service must receive unsolicited traffic. You typically perform this task on an ongoing basis as your server roles and server configurations change.

### Task requirements

No special tools are required to complete this task.

### Task procedures

To complete this task, perform the following procedure:

Add a Port to the Firewall Rules List for a Specific Connection

**See Also**
**Concepts**
Known Issues for Managing Firewall Rules
Configuring Program Firewall Rules
Configuring Port Firewall Rules
Configuring System Service Firewall Rules
Configuring Scope Settings

# Add a Port to the Firewall Rules List for a Specific Connection

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to add a TCP or UDP port to the Windows Firewall exceptions list on a per-connection basis. This is useful if your computer has multiple network connections (sometimes referred to as interfaces) and you want to open different ports on each network connection. When you add a port to the exceptions list, the port is always open; unsolicited incoming traffic is always allowed to pass through the port unless you select the **Don't allow exceptions** option when you turn on Windows Firewall.

| **Note** |
| --- |
| Adding a port to the exceptions list can reduce the security of your computer because the port will be open any time the computer is running. You should add ports to the exceptions list only when it is not possible to add a program (.exe file) to the exceptions list. |

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

#### To add a port to the exceptions list for a specific connection

This procedure can be performed using the graphical user interface or the command prompt. You cannot use Group Policy to configure per-connection port exceptions.

### Using the graphical user interface

#### To add a port to the exceptions list for a specific connection

1. Open Windows Firewall, and then click the **Advanced** tab.
2. In **Network Connection Settings**, click the connection that you want to configure, and click **Settings**.
3. Click **Add**.
4. In the **Service Settings** dialog box, in **Description of service**, type a display name for the port exception. The display name appears in the exceptions list.
5. In **Name or IP address of the computer hosting this service on your network**, type the name or IP address of the local computer.
6. In **External port number for this service**, type the port number to add to the exceptions list.
7. In **Internal port number for this service**, type the port number to add to the exceptions list.
8. Click either **TCP** or **UDP**, depending on the type of traffic, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

### Using the command prompt

#### To add a port to the exceptions list for a specific connection

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set portopening protocol** = *protocol* **port** = *port* **name** = *name* **mode** = *mode* **interface** = " *interface* "

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *protocol* | TCP, UDP, All | Specifies the protocol for the port. Use All to specify both TCP and UDP. |
| *port* | Any number between 1 and 65,535 | Specifies the port number for which you want to create an exception. |
| *name* | Any string less than 256 characters | Specifies the friendly name of the port exception. You must enclose *name* in quotation marks. |
| *mode* | enable, disable | Specifies whether to enable or disable the exception. |
| *interface* | The name of any network connection that can exist on the computer, as displayed in the Network Connections folder | Specifies the connection, or interface, on which to enable or disable the exception. You must enclose *interface* in quotation marks. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Firewall Rules for Specific Connections
Known Issues for Managing Firewall Rules
Configuring Port Firewall Rules

# Edit or Delete a Connection-Specific Port Firewall Rule

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to edit or delete a TCP or UDP port that is in the Windows Firewall exceptions list on a per-connection basis. This is useful if your computer has multiple network connections (sometimes referred to as interfaces) and you have opened different ports (created specific exceptions) on each network connection.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**To edit or delete a connection-specific port exception**

This procedure can be performed using the graphical user interface or the command prompt. You cannot use Group Policy to configure per-connection port exceptions.

**Using the graphical user interface**

**To edit or delete a connection-specific port exception**

1. Open Windows Firewall, and click the **Advanced** tab.
2. In **Network Connection Settings**, click the connection that the system service uses to communicate, and then click **Settings**.
3. Click the system service that you want to edit or delete.

   If you want to delete the system service exception, click **Delete**, and then click **OK**.

   If you want to edit the system service exception, click **Edit**, make changes to the system service exception settings, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To edit or delete a connection-specific port exception**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set portopening protocol** = *protocol* **port** = *port* **name** = *name* **mode** = *mode* **interface** = **"** *interface* **"**

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *protocol* | TCP, UDP, All | Specifies the protocol for the port. Use All to specify both TCP and UDP. |
| *port* | Any number between 1 and 65,535 | Specifies the port number for which you want to create an exception. |
| *name* | Any string less than 256 characters | Specifies the friendly name of the port exception. You must enclose *name* in quotation marks. |
| *mode* | enable, disable | Specifies whether to enable or disable the exception. |

| | | |
|---|---|---|
| *interface* | The name of any network connection that can exist on the computer, as displayed in the Network Connections folder | Specifies the connection, or interface, on which to enable or disable the exception. You must enclose *interface* in quotation marks. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- You cannot change the name of a per-connection port exception by clicking **Edit**. To rename a per-connection port exception, delete the exception and create a new exception with the new name.

**See Also**
**Concepts**
Configuring Firewall Rules for Specific Connections
Known Issues for Managing Firewall Rules
Configuring Port Firewall Rules

# Configuring Scope Settings

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

When you configure a program, port, or system service exception, you should also configure scope settings for the exception. Scope settings define from which addresses incoming traffic is allowed to originate, which defines the set of computers that are allowed to send traffic for an exception. There are three scope options:

**Any computer (including those on the Internet)**

This setting allows traffic from any Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) address through the exception. This setting might make your computer more accessible to malicious users or programs on the Internet.

**My network (subnet) only**

This setting allows traffic only from IPv4 or IPv6 addresses that can be reached directly by your computer. Windows Firewall determines whether the source IPv4 or IPv6 address of the incoming packet can be reached directly by querying the IPv4 and IPv6 routing tables. You can see all destinations that are considered directly reachable by typing the **route print** command at a command prompt. For the IPv4 routing table, all IPv4 addresses that match the routes in which the IPv4 address of the **Gateway** column equals the IPv4 address of the **Interface** column are considered directly reachable. For the IPv6 routing table, all IPv6 addresses that match routes in which the **Gateway** column is set to **On-link** are considered directly reachable. Therefore, the set of directly reachable addresses depends on your networking configuration, as specified by the IPv4 and IPv6 configuration of LAN-based connections (such as Ethernet and 802.11 wireless), dial-up connections, and broadband Internet connections. In some Internet configurations, all destinations are considered directly reachable.

For example, for a computer that is directly connected to a private home network only, the set of directly reachable unicast addresses is confined to those that match the IPv4 network ID of the private subnet. If the network connection is configured with an IPv4 address of 192.168.0.99 with a subnet mask of 255.255.0.0, only traffic from IPv4 addresses in the range 192.168.0.0 to 192.168.255.255 is allowed.

In another example, for a computer that is directly connected to both a private home network and the Internet through a cable modem, the set of directly reachable unicast addresses are those that match either the network ID of the private subnet or the cable modem provider subnet. For example, if the private network connection is configured with an IPv4 address of 192.168.0.1 and a subnet mask of 255.255.0.0 and the cable modem connection is configured with an IPv4 address of 131.107.17.211 and a subnet mask of 255.255.255.0, traffic received by either network connection is allowed from IPv4 addresses in ranges from 192.168.0.0 to 192.168.255.255 and from 131.107.17.0 to 131.107.17.255.

| **Important** |
| --- |
| The use of the **My network (subnet) only** option can make your computer more accessible than you expect. Make sure that you clearly understand how this option restricts scope before you use it. |

**Custom list**

This setting allows you to specify one or more IPv4 addresses or IPv4 address ranges separated by commas (with no spaces). IPv4 address ranges typically correspond to subnets. For IPv4 addresses, type the IPv4 address in dotted decimal notation. For IPv4 address ranges, you can specify the range using a dotted decimal subnet mask or a prefix length. When you use a dotted decimal subnet mask, you can specify the range as an IPv4 network ID (such as 10.47.81.0/255.255.255.0) or by using an IPv4 address within the range (such as 10.47.81.231/255.255.255.0). When you use a network prefix length, you can specify the range as an IPv4 network ID (such as 10.47.81.0/24) or by using an IPv4 address within the range (such as 10.47.81.231/24). The following is an example custom list:
10.91.12.56,10.7.14.9/255.255.255.0,10.116.45.0/255.255.255.0,172.16.31.11/24,172.16.111.0/24

You cannot specify a custom list for IPv6 traffic.

| **Important** |
| --- |

When you configure and enable an exception, you are instructing Windows Firewall to allow specific unsolicited incoming traffic sent from the specified scope (from any address, from an address that can be reached directly, or from a custom list). For any scope, enabling an exception makes the computer accessible to attacks based on incoming unsolicited traffic from computers that are assigned the allowed addresses and from malicious computers that spoof traffic. There is no way to prevent spoofed attacks from the Internet on connections assigned public IPv4 addresses except to disable the exception. Therefore, you should try to configure scope options so that the number of computers that are allowed to send unsolicited traffic through an exception is kept to a minimum. This will reduce, but not eliminate, the likelihood of a spoof attack.

**When to perform this task**

You should perform this task whenever you add, enable, or edit an exception. You typically perform this procedure on an ongoing basis as your server configurations and server roles change and your network architecture changes.

**Task requirements**

No special tools are required to complete this task.

**Task procedures**

To complete this task, perform the following procedures:

Change the Scope of a Firewall Rule

**See Also**
**Concepts**
Known Issues for Managing Firewall Rules
Configuring Program Firewall Rules
Configuring Port Firewall Rules
Configuring System Service Firewall Rules
Configuring Firewall Rules for Specific Connections

# Change the Scope of a Firewall Rule

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to configure the scope of an exception when you add or change a port, program, or system service exception. This procedure is useful when you want to create an exception but you want to make the exception less accessible, particularly to malicious users.

You cannot configure scope settings for port exceptions that are configured on a per-connection basis. The scope of a per-connection port exception includes any IP address. Also, you cannot configure scope settings for Internet Control Message Protocol (ICMP) exceptions. The scope of an ICMP exception includes any IP address.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

Using the **My network (subnet) only** option can make your computer more accessible than you expect. Make sure that you clearly understand how this option restricts scope before you use it.

**To change the scope of an exception**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To change the scope of an exception**

1. Open Windows Firewall, and click the **Exceptions** tab.
2. In **Programs and Services**, click the exception that you want to configure, and click **Edit**.
3. Click **Change scope**, and do one of the following:

   If you want to enable the exception for all computers, click **Any computer (including those on the Internet)**, and then click **OK**.

   If you want to enable the exception for computers that can be reached directly by your computer, click **My network (subnet) only**, and then click **OK**.

   If you want to enable the exception for specific IPv4 IP addresses or IPv4 IP address ranges, click **Custom list**, enter the specific IP addresses or IP address ranges, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To change the scope of an exception**

- If you are changing the scope of a program exception, type the following at the command prompt, and press ENTER:

  **netsh firewall set allowedprogram program =** *program* **name =** *name* **mode =** *mode* **scope =** *scope* **addresses =** *addresses*

- If you are changing the scope of a system service exception, type the following at the command prompt, and press ENTER:

  **netsh firewall set service type** = *type* **mode** = *mode* **scope** = *scope* **addresses** = *addresses*

- If you are changing the scope of a port exception, type the following at the command prompt, and press ENTER:

  **netsh firewall set portopening protocol** = *protocol* **port** = *port* **name** = *name* **mode** = *mode* **scope** = *scope* **addresses** = *addresses*

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *program* | Path and file name for the .exe file. | Specifies the program to add to the exceptions list. |
| *protocol* | TCP, UDP, All | Specifies the protocol for the port. Use All to specify both TCP and UDP. |
| *port* | Any number from 1 to 65,535 | Specifies the port number. |
| *type* | Fileandprint, Remoteadmin, Remotedesktop, Upnp, All | Specifies the service to configure. Use All to specify all services. |
| *name* | Any string less than 256 characters | Specifies the friendly name for the exception, which is displayed in the graphical user interface. You must enclose *name* in quotation marks. |
| *mode* | enable, disable | Specifies whether the exception is enabled or disabled. |
| *scope* | All, Subnet, Custom | Specifies the scope setting. |
| *addresses* | Comma-separated list of IPv4 addresses or ranges of IPv4 addresses | Specifies the IPv4 addresses or IPv4 address ranges if you use the **Custom scope** setting. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Scope Settings
Known Issues for Managing Firewall Rules

# Managing IPsec, Multicast, and ICMP Settings

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Firewall provides special settings for managing multicast and Internet Control Message Protocol (ICMP) traffic, as well as traffic protected by Internet Protocol security (IPsec).

To help you manage this type of traffic, review Best Practices for Managing Windows Firewall and Known Issues for Managing IPsec, Multicast, and ICMP Settings.

The following tasks are described in this objective:

Configuring ICMP Settings

Configuring IPsec Settings

Configuring Multicast Settings

# Known Issues for Managing IPsec, Multicast, and ICMP Settings

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review the following known issues before you manage Internet Protocol (IPsec), multicast, and Internet Control Message Protocol (ICMP) settings.

- Your computer might not be able to connect to multicast streams if you have Windows Firewall enabled. For more information about this issue, see article ID 883831 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45388). The article describes why Windows Firewall might prevent you from using Windows Media Player to connect to multicast and UDP streams and how you can configure Windows Firewall to overcome the problem.

- You might experience problems with ICMP traffic if you also use IPsec to secure traffic. For more information, see article ID 889527 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=45389). This article describes why you might have problems with ICMP traffic when you use IPsec to secure traffic and it provides information about how you can overcome the problem.

- The Ping Command Times Out.

# Configuring ICMP Settings

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure Windows Firewall so that ICMP version 4 (ICMPv4) and ICMP version 6 (ICMPv6) traffic is either blocked or allowed. The following table describes the ICMPv4 and ICMPv6 messages that you can control with Windows Firewall.

| ICMP message | Description |
|---|---|
| Allow incoming echo request | Corresponds to ICMPv4 Type 8 (Echo) and ICMPv6 Type 128 (Echo Request) messages. |
| Allow incoming timestamp request | Corresponds to ICMPv4 Type 13 (Timestamp) messages. |
| Allow incoming mask request | Corresponds to ICMP Type 17 (Address Mask Request) messages. |
| Allow incoming router request | Corresponds to ICMP Type 9 (Router Solicitation) messages. |
| Allow outgoing destination unreachable | Corresponds to ICMPv4 Type 3 (Destination Unreachable) and ICMPv6 Type 1 (Destination Unreachable) messages. |
| Allow outgoing source quench | Corresponds to ICMP Type 4 (Source Quench) messages. |
| Allow outgoing parameter problem | Corresponds to ICMP Type 12 (Parameter Problem) and ICMPv6 Type 4 (Parameter Problem) messages. |
| Allow outgoing time exceeded | Corresponds to ICMP Type 11 (Time Exceeded) and ICMPv6 Type 3 (Time Exceeded) messages. |
| Allow redirect | Corresponds to ICMP Type 5 (Redirect) and ICMPv6 Type 137 (Neight Discovery Redirect) messages. |
| Allow outgoing packet too big | Corresponds to ICMPv6 Type 2 (Packet Too Big) messages. |

If you do not enable the **Allow incoming echo requests** setting, commands that use the ICMP Echo message (also known as the ICMP Echo Request message), such as **ping** or **tracert**, will not work. If you are running network management software that uses ICMP Destination Unreachable messages, you need to enable the **Allow outbound destination unreachable** setting.

If you configure Windows Firewall so that traffic is allowed through TCP port 445, Windows Firewall will allow incoming ICMP Echo messages automatically. This is true even if you disable the **Allow incoming echo requests** setting, or you disable the **Windows Firewall: Allow ICMP exceptions** Group Policy setting, or you use the **netsh firewall set icmpsetting 8 disable** command. For example, there are two predefined service exceptions that allow traffic through TCP port 445: the File and Printer Sharing exception and the Remote Administration exception. If you enable either of these exceptions, and you allow unsolicited incoming traffic to pass through TCP port 445, other computers will be able to access your computer with the **ping** command.

**When to perform this task**

You should use these settings if your organization uses the **ping** or **tracert** commands for troubleshooting. Usually, you configure these settings only once or on an as-needed basis.

**Task requirements**

No special tools are required to perform this task.

**Task procedures**

To complete this task, perform the following procedure:

Block and Unblock ICMP Messages

**See Also**
**Concepts**
Known Issues for Managing IPsec, Multicast, and ICMP Settings

# Block and Unblock ICMP Messages

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to block or unblock Internet Control Message Protocol (ICMP) messages. This procedure is useful if you want to allow other computers to access your computer with the **ping** command or **tracert** command.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

**To block and unblock ICMP messages**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To block and unblock ICMP messages**

1. Open Windows Firewall.
2. Click the **Advanced** tab.
3. In **ICMP**, click **Settings**.
4. In the **ICMP Settings** dialog box, do one of the following:

   To unblock an ICMP message (allow the ICMP message to pass through Windows Firewall), select the check box next to the ICMP message and then click **OK**.

   To block an ICMP message (prevent the ICMP message from passing through Windows Firewall), clear the check box next to the ICMP message and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To block and unblock ICMP messages**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set icmpsetting type =** *type* **mode =** *mode*

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *type* | 2,3,4,5,8,9,11,12,13,17, All | Specifies the ICMP message type. All specifies that you want to configure all ICMP message types. |

| | | |
|---|---|---|
| *mode* | Enable, Disable | Specifies whether to enable (allow the ICMP message to pass through Windows Firewall) or disable (prevent the ICMP message from passing through Windows Firewall). |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

### Notes

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Known Issues for Managing IPsec, Multicast, and ICMP Settings
Configuring ICMP Settings

# Configuring IPsec Settings

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure Windows Firewall so that all incoming authenticated traffic protected by Internet Protocol security (IPsec) completely bypasses Windows Firewall processing on specific computers. To do this, you must enable the **Windows Firewall: Allow authenticated IPSec bypass** Group Policy setting. To specify which computers allow incoming IPsec-protected traffic to bypass Windows Firewall, you must provide a Security Descriptor Definition Language (SDDL) string that corresponds to the group accounts for the computers to which this policy applies.

| **Important** |
| --- |
| This information applies only to IPsec in Windows Server 2003 with Service Pack (SP) 1 and Windows XP with SP2. It does not apply to Microsoft Windows Vista® or Windows Server® 2008, or later versions of Windows. For information about using authenticated bypass on Windows Vista or Windows Server 2008, see How to Enable Authenticated Bypass or "Allow if Secure" at http://go.microsoft.com/fwlink/?linkid=111313 in the Windows Server 2008 Technical Library. |

The format of the SDDL string for a single group is:

**O:DAG:DAD:(A;;RCGW;;;** *SID* **)**

Where *SID* is the Security Identifier (SID) of a group account.

Use the Getsid.exe tool to obtain the SID of a group account. Getsid.exe is typically used to compare the SIDs of two accounts on different domain controllers, but you can also use it to obtain the SID of a specified user or group account.

To obtain a SID for a group account, use the following syntax:

**getsid ** *domain_controller group_account* **** *domain_controller group_account*

Where *domain_controller* is the computer name of a domain controller and *group_account* is the group account name.

The following example uses the Getsid.exe tool with a domain controller named EXAMPLE2 in the example.com domain and a group account named IPsecComputers:

```
C:\>getsid \\example2 IPsecComputers \\example2 IPsecComputers
The SID for account EXAMPLE\IPsecComputers matches account EXAMPLE\IPsecComputers
The SID for account EXAMPLE\IPsecComputers is
S-1-5-21-3575094098-3669797271-991787341-1127
The SID for account EXAMPLE\IPsecComputers is
S-1-5-21-3575094098-3669797271-991787341-1127
```

Although redundant, it provides the information you need. To allow computers in the IPsecComputers group to bypass Windows Firewall, use the following string to configure the **Windows Firewall: Allow authenticated IPSec bypass** Group Policy setting:

O:DAG:DAD:(A;;RCGW;;;S-1-5-21-3575094098-3669797271-991787341-1127)

If you have more than one group, then the syntax for the SDDL string is:

O:DAG:DAD:(A;;RCGW;;;*SID1*) (A;;RCGW;;;*SID2*) (A;;RCGW;;;*SID3*)...

If you enable the **Windows Firewall: Allow authenticated IPSec bypass** Group Policy setting, and a computer that is a member of one of the security groups on the SDDL list receives an IPsec-protected packet, Windows Firewall does not process it to determine whether it is allowed.

If you disable or do not configure this Group Policy setting, incoming IPsec-protected traffic is processed by Windows Firewall. If you enable this Group Policy setting and configure an SDDL list, upon disabling this setting, Windows Firewall deletes the list.

**When to perform this task**

You can use this setting if your organization uses both IPsec and Windows Firewall, and you assume that traffic protected by IPsec is safe. If you enable the **Windows Firewall: Allow authenticated IPSec bypass** Group Policy setting, you might need to modify the SDDL list to include additional security groups on an ongoing basis.

**Task requirements**

Getsid.exe. For more information, see Windows Support Tools at the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43235).

**Task procedures**

To complete this task, perform the following procedure:

Configure Authenticated IPsec Bypass

**See Also**
**Concepts**
Known Issues for Managing IPsec, Multicast, and ICMP Settings

# Configure Authenticated IPsec Bypass

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to allow authenticated traffic protected by Internet Protocol security (IPsec) to bypass Windows Firewall. When you perform this procedure, you must provide a Security Descriptor Definition Language (SDDL) string that contains a list of the computers or groups of computers that you want to bypass Windows Firewall. If you enable the **Windows Firewall: Allow authenticated IPSec bypass** Group Policy setting, and a computer that is a member of one of the security groups on the SDDL list receives an IPsec-protected packet, Windows Firewall does not process it to see whether it is allowed. This setting is useful if you use Windows Firewall and IPsec in your organization and assume that any computer using IPsec is protected and that its traffic should be allowed to pass through Windows Firewall.

**Important**

This information applies only to IPsec in Windows Server 2003 with Service Pack (SP) 1 and Windows XP with SP2. It does not apply to Microsoft Windows Vista® or Windows Server® 2008, or later versions of Windows. For information about using authenticated bypass on Windows Vista or Windows Server 2008, see How to Enable Authenticated Bypass or "Allow if Secure" at http://go.microsoft.com/fwlink/?linkid=111313 in the Windows Server 2008 Technical Library.

For information about creating SDDL strings, see Managing IPsec, Multicast, and ICMP Settings.

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

### Special Considerations

If you perform this procedure and enable the **Windows Firewall: Allow authenticated IPSec bypass** Group Policy setting, and then later disable the **Windows Firewall: Allow authenticated IPSec bypass** Group Policy setting, the SDDL strings that you entered will be deleted. Therefore, you should save the SDDL strings that you use to perform this procedure in case you need to perform this procedure again in the future.

#### To allow traffic protected by IPsec through Windows Firewall

This procedure can be performed through Group Policy only. You cannot use the graphical user interface or the command prompt to perform this procedure.

#### To allow traffic protected by IPsec through Windows Firewall

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, and then click **Windows Firewall**.
3. In the details pane, double-click **Windows Firewall: Allow authenticated IPSec bypass**.
4. In the **Windows Firewall: Allow authenticated IPSec bypass properties** dialog box, on the **Settings** tab, click **Enabled**.
5. In **Define IPSec peers to be exempted from firewall policy**, type the SDDL string that corresponds to the group accounts for the computers to which this policy applies, and then click **OK**.

**Note**

- The IPsec bypass only applies to computers that are the responder to an IPsec SA negotiation. The initiator of the session does not set the bypass flag for the connection, so the firewall on the initiator still inspects the packets. If the responder uses the same IPsec SA to connect back to the initiator, the packets are filtered out by the firewall on the initiator.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Group Policy settings must be refreshed before they take effect.

**See Also**
**Concepts**
Configuring IPsec Settings
Known Issues for Managing IPsec, Multicast, and ICMP Settings

# Configuring Multicast Settings

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure Windows Firewall so that Windows Firewall blocks all unicast responses to outgoing multicast or broadcast traffic. When you enable this setting, and your computer sends multicast or broadcast messages to other computers, Windows Firewall blocks the unicast responses sent by those other computers.

If you disable or do not configure this setting, and this computer sends a multicast or broadcast message to other computers, Windows Firewall waits as long as three seconds for unicast responses from the other computers and then blocks all later responses.

This setting has no effect if the unicast message is a response to a Dynamic Host Configuration Protocol (DHCP) broadcast message sent by this computer. Windows Firewall always permits those DHCP unicast responses. However, this setting can interfere with the NetBIOS messages that detect name conflicts.

**When to perform this task**

You should perform this task if your computer sends broadcast or multicast messages and you want Windows Firewall to block all unicast responses.

**Task requirements**

No special tools are required to perform this task.

**Task procedures**

To complete this task, perform the following procedure:

Block Unicast Responses to Multicast Traffic

**See Also**
**Concepts**
Known Issues for Managing IPsec, Multicast, and ICMP Settings

# Block Unicast Responses to Multicast Traffic

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to block unicast response traffic that is received in response to multicast or broadcast traffic.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

**To block unicast responses to multicast traffic**
This procedure can be performed using the command prompt or Group Policy. You cannot perform this procedure in the graphical user interface.

**Using the command prompt**

**To block unicast responses to multicast traffic**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set multicastbroadcastresponse mode = enable**

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Using Group Policy**

**To block unicast responses to multicast traffic**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, and then click **Windows Firewall**.
3. In the details pane, double-click **Windows Firewall: Prohibit unicast response to multicast or broadcast requests**.
4. In the **Windows Firewall: Prohibit unicast response to multicast or broadcast requests properties** dialog box, on the **Settings** tab, click **Enabled**.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Configuring Multicast Settings
Known Issues for Managing IPsec, Multicast, and ICMP Settings

# Managing Windows Firewall Notifications

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, Windows Firewall displays a **Windows Security Alert** dialog box (referred to as a notification) when a program attempts to listen for unsolicited incoming traffic and the incoming traffic is blocked. If you are a member of the Administrators group on the computer (or you are a member of a group that is a member of the Administrators group), the notification provides you with the following options:

**Unblock**. This adds the program to the exceptions list in the profile that Windows Firewall is currently using and enables the program exception. If you choose this option, the **Windows Security Alert** dialog box will not be displayed again for the program even if you disable the program exception. The **Windows Security Alert** dialog box might appear again if the program attempts to listen for incoming traffic when Windows Firewall is using the other profile.

**Keep Blocking**. This adds the program to the exceptions list in the profile that Windows Firewall is currently using, but does not enable the exception. If you choose this option, the **Windows Security Alert** dialog box will not be displayed again for the program because the program is in the exceptions list. The **Windows Security Alert** dialog box might appear again if the program attempts to listen for incoming traffic when Windows Firewall is using the other profile.

**Ask Me Later**. No action is taken (that is, the program is not added the exceptions list). If you choose this option, the **Windows Security Alert** dialog box will appear the next time the program attempts to listen for unsolicited incoming traffic.

If you are not a member of the Administrators group on the computer, you will still see a **Windows Security Alert** dialog box when a program attempts to listen for unsolicited incoming traffic and the incoming traffic is blocked; however, the notification is informational only. You are not given the ability to take any action; the **Unblock**, **Keep Blocking**, or **Ask Me Later** options are not displayed.

Windows Firewall does not display a notification in the user interface (UI) when a system service attempts to listen for incoming traffic on a port and the incoming traffic is blocked. This is also true for any program that runs like a system service (that is, a program that runs under an account that has higher privilege than a user account, for example, the Local System account, or a program that runs even when there is no user logged on to the computer). You can use the security event log to determine whether a system service, or a program that runs like a system service, attempts to listen for incoming traffic. To do this, you must enable **Audit process tracking** and **Audit policy change** settings in Group Policy. When you do this, Windows Firewall will write a Failure Audit with Event ID 861 to the security event log any time a program or system service attempts to listen for incoming traffic. For more information about Windows Firewall events, see Using the Security Log.

In addition, there are two other instances when Windows Firewall might not display notifications:

- Windows Firewall does not display notifications if you select the **Don't allow exceptions** check box in Windows Firewall in Control Panel, or if you enable the **Do not allow exceptions** setting in Group Policy, or if you use the **netsh firewall set opmode exceptions = disable** command. There are no Windows Firewall settings that allow you to override this behavior.

- Windows Firewall does not display notifications for programs that rely on the Winsock driver to dynamically bind to a UDP port. If a program uses this method (sometimes referred to as wildcard binds) to bind to a UDP port, you might be able to use the **netstat** command and other troubleshooting tools to determine which UDP port is being used, and then add that port to the exceptions list.

You can disable the notification feature, which prevents Windows Firewall from displaying a **Windows Security Alert** dialog box. This might be necessary on remotely managed servers.

**When to perform this task**

You should manage Windows Firewall notifications as the notifications are displayed. If you do not want notifications to appear, you only have to disable notifications once unless you enable notifications or reset Windows Firewall settings.

**Task requirements**

No special tools are required to complete this task.

**Task procedures**

To enable or disable the notification feature, use the following procedure:

Enable or Disable Windows Firewall Notifications

**See Also**
**Concepts**
Known Issues for Managing Windows Firewall Notifications

# Known Issues for Managing Windows Firewall Notifications

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review the following known issues before you manage Windows Firewall notifications.

- Windows Firewall does not display notifications if you select the **Don't allow exceptions** check box in Windows Firewall in Control Panel, or if you enable the **Do not allow exceptions** policy setting in Group Policy, or if you use the **netsh firewall set opmode exceptions = disable** command. There are no Windows Firewall settings that allow you to override this behavior.

- Windows Firewall does not display notifications for programs that rely on the Winsock driver to dynamically bind to a UDP port. If a program uses this method (sometimes referred to as wildcard binds) to bind to a UDP port, you might be able to use the **netstat** command and other troubleshooting tools to determine which UDP port is being used, and then add that port to the exceptions list.

- Windows Firewall does not display a notification when a system service attempts to listen for incoming traffic on a port and the incoming traffic is blocked. This is also true for any program that runs like a system service (that is, a program that runs under an account that has higher privilege than a user account, for example, the Local System account, or a program that runs even when there is no user logged on to the computer).

# Enable or Disable Windows Firewall Notifications

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure to enable or disable the notification feature in Windows Firewall. This is useful on remotely managed servers or on servers that are managed or monitored infrequently. The notification feature is enabled by default.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

**To enable or disable Windows Firewall notifications**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To enable or disable Windows Firewall notifications**

1. Open Windows Firewall.
2. Click the **Exceptions** tab, and do one of the following:

   To enable notifications, select the **Display a notification when Windows Firewall blocks a program** check box, and click **OK**.

   To disable notifications, clear the **Display a notification when Windows Firewall blocks a program** check box, and click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To add a port to the exceptions list**

- To enable notifications, type the following at the command prompt, and press ENTER:

  **netsh firewall set notifications mode = enable profile =** *profile*

- To disable notifications, type the following at the command prompt, and press ENTER:

  **netsh firewall set notifications mode = disable profile =** *profile*

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *profile* | Current, Domain, Standard, All | Specifies the profile that you want to configure. All indicates that you want to configure both the standard profile and the domain profile. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Windows Firewall does not display notifications if you select the **Don't allow exceptions** check box in Windows Firewall in Control Panel, or if you enable the **Do not allow exceptions** policy setting in Group Policy, or if you use the **netsh firewall set opmode exceptions = disable** command. There are no Windows Firewall settings that allow you to override this behavior.
- Windows Firewall does not display notifications for programs that rely on the Winsock driver to dynamically bind to a UDP port. If a program uses this method (sometimes referred to as wildcard binds) to bind to a UDP port, you might be able to use the **netstat** command and other troubleshooting tools to determine which UDP port is being used, and then add that port to the exceptions list.

**See Also**
**Concepts**
Known Issues for Managing Windows Firewall Notifications

# Optimizing Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To optimize Windows Firewall, see the following:

Best Practices for Optimizing Windows Firewall

Known Issues for Optimizing Windows Firewall

# Best Practices for Optimizing Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

As a best practice, you should upgrade your networking hardware on your server computers so you are using network adapters that support checksum offloading. Checksum offloading allows the network adapter to perform the checksum processing that is usually performed below the TCP/IP stack, which will reduce the load on your computer's processor.

| **Note** |
| --- |
| Checksum offloading will not usually impact the performance on a client computer. |

**See Also**
**Concepts**

Best Practices for Managing Windows Firewall
Best Practices for Securing Windows Firewall
Best Practices for Monitoring Windows Firewall

# Known Issues for Optimizing Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Running Windows Firewall can affect your computer's network performance by reducing the throughput, or bandwidth, that can be handled by your computer. The impact is usually minimal, unless your server handles a large amount of network traffic. In addition, if your computer primarily handles unsolicited incoming TCP traffic, the performance impact is usually minimal; however, if your computer primarily handles unsolicited incoming UDP traffic, the performance impact can be substantial. This is because short-lived connections like UDP connections require more processing than long-lived connections like TCP connections.

For example, the optimal type of connection would be a large file download that requires some early processing by Windows Firewall, but little or no processing as soon as the connection is established and the file transfer occurs. The impact of Windows Firewall on file servers and Web servers is minimal because these servers typically use TCP. In contrast, Windows Firewall can have a significant impact on a streaming media server that uses unicast broadcasts to deliver media. This is because streaming media servers typically use UDP to create a new connection to the media encoder each time a client connects to the live content stream and each of those connections must pass through Windows Firewall. Servers that are connected to the Internet and handle a high number of unsolicited incoming messages might also experience impact on network performance because although most of the messages will be blocked, Windows Firewall must still process each incoming message.

Even if your computer primarily handles UDP traffic, the impact on performance will not be that significant unless your computer's processor and network adapter are operating at close to maximum capacity. If your computer handles a substantial amount of incoming traffic when Windows Firewall is not turned on, and the network traffic is putting a high load on your processor, then Windows Firewall will impact your computer's network performance. However, this is often a symptom of poor network or infrastructure design; a problem like this one should be addressed no matter which firewall solution you use.

# Monitoring Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can monitor Windows Firewall activity by using the Windows Firewall log file and the security log. Both tools are useful for troubleshooting Windows Firewall problems and temporarily monitoring the behavior of Windows Firewall.

To help you monitor Windows Firewall activity, review Best Practices for Monitoring Windows Firewall, Known Issues for Using the Windows Firewall Log and Known Issues for Using the Security Log.

The following tasks are described in this objective:

- Using the Windows Firewall Log
- Using the Security Log

# Best Practices for Monitoring Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for monitoring Windows Firewall**
Use the following best practices when you monitor Windows Firewall.

Enable the Windows Firewall log file only when you need to.

Enable auditing of Windows Firewall events when you first turn on Windows Firewall.

**Enable the Windows Firewall log file only when you need to.**
You should enable the Windows Firewall log file only when you need to troubleshoot Windows Firewall problems or monitor Windows Firewall behavior over short periods of time. Permanently enabling the Windows Firewall log file makes your log files difficult to manage and uses disk space.

**Enable auditing of Windows Firewall events when you first turn on Windows Firewall.**
You should audit Windows Firewall events when you first turn on Windows Firewall. This will help you identify programs and services that are acting as servers, peers, or listeners. When you have identified servers, peers, or listeners and configured Windows Firewall exceptions, you should disable auditing of Windows Firewall events because auditing can impact the performance of servers that handle a large amount of network traffic.

**See Also**
**Concepts**
Best Practices for Managing Windows Firewall
Best Practices for Optimizing Windows Firewall
Best Practices for Securing Windows Firewall

# Using the Windows Firewall Log

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the Windows Firewall log file to monitor TCP and UDP connections and packets that are blocked by Windows Firewall. The log file provides source and destination IP addresses, port numbers, and protocols. For more information about log file structure, including descriptions of each log file parameter, see the section titled "Windows Firewall Tools and Settings" in the Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729).

By default, the log file is disabled, which means that no information is written to the log file. To use the log file, you must enable the logging of dropped packets, successful connections, or both. You can also configure other log file settings, including the name, location, and maximum size of the log file.

**When to perform this task**

Enable the Windows Firewall log file when you need to troubleshoot a Windows Firewall problem or you need to temporarily monitor Windows Firewall behavior.

**Task requirements**

No special tools are required to complete this task.

**Task procedures**

To complete this task, use the following procedures:

- Enable or Disable the Windows Firewall Log
- View the Windows Firewall Log File
- Change the Name and Location of the Windows Firewall Log File
- Change the Size of the Windows Firewall Log File

**See Also**
**Concepts**
Best Practices for Monitoring Windows Firewall
Using the Security Log

# Known Issues for Using the Windows Firewall Log

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review the following known issues before you use the Windows Firewall log file.

- Logging is available only on connections for which Windows Firewall is turned on. This applies if you configure per-connection settings because no information is written to the log file for connections that are not protected with Windows Firewall.
- Outgoing traffic that successfully reaches its destination is not logged. In addition, outbound traffic that is not blocked is not logged.
- The default size of the Windows Firewall log file is 4,096 kilobytes (KB); the maximum size is 32,767 KB.
- When the Windows Firewall log reaches the size that you set in **Size limit (KB)**, the log file is renamed and a new log file is created.

# Enable or Disable the Windows Firewall Log

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to turn on (enable) or turn off (disable) the Windows Firewall log file. This is useful if you are troubleshooting Windows Firewall problems or you are temporarily monitoring Windows Firewall behavior.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

**To enable or disable the Windows Firewall log**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To enable or disable the Windows Firewall log**

1. Open Windows Firewall, and click the **Advanced** tab.
2. In **Security Logging**, click **Settings**.

   To enable the log file, select the **Log dropped packets** check box or the **Log successful connections** check box or both check boxes, and click **OK**.

   To disable the log file, clear the **Log dropped packets** check box and the **Log successful connections** check box, and click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To enable or disable the Windows Firewall log**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set logging droppedpackets =** *droppedpackets* **connections =** *connections*

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *droppedpackets* | enable, disable | Specifies whether to turn on (enable) or turn off (disable) the logging of dropped packets. |
| *connections* | enable, disable | Specifies whether to turn on (enable) or turn off (disable) the logging of successful connections. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Known Issues for Using the Windows Firewall Log
View the Windows Firewall Log File
Change the Size of the Windows Firewall Log File
Change the Name and Location of the Windows Firewall Log File

# View the Windows Firewall Log File

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to view the data that has been written to the Windows Firewall log. This is useful if you are troubleshooting Windows Firewall problems or you are temporarily monitoring Windows Firewall behavior.

**Administrative Credentials**

No special administrative credentials are required to perform this task.

**Special Considerations**

No special considerations are required to perform this task.

**To view the Windows Firewall log**
This procedure can be performed using the graphical user interface.

**Using the graphical user interface**

**To view the Windows Firewall log**

1. Open Windows Firewall, and click the **Advanced** tab.
2. In **Security Logging**, click **Settings**.
3. In the **Log Settings** dialog box, click **Save As**.
4. Right-click the log file name, and then click **Open**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- By default, the Windows Firewall log file is named Pfirewall.log and is saved in the systemroot folder.
- You must enable the Windows Firewall log file for it to appear in the systemroot folder.

**See Also**
**Concepts**
Known Issues for Using the Windows Firewall Log
Enable or Disable the Windows Firewall Log
Change the Size of the Windows Firewall Log File
Change the Name and Location of the Windows Firewall Log File

# Change the Name and Location of the Windows Firewall Log File

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to change the name and location of the Windows Firewall log file. This is useful if you do not want users to easily find the log file or you want to have the log file written to a central location.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

**To change the name and location of the Windows Firewall log**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To change the name and location of the Windows Firewall log**

1. Open Windows Firewall, and click the **Advanced** tab.
2. In **Security Logging**, click **Settings**.
3. Click **Save As**, and navigate to the location where you want to save the log file.
4. In **File name**, type the new log file name, or leave it blank to accept the default name of pfirewall.log.
5. Click **Save**, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To change the name and location of the Windows Firewall log**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set logging filelocation = "** *filelocation* **"**

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *filelocation* | Any string that includes a valid path and file name for a text file. | You must enclose *filelocation* in quotation marks if there are spaces in the path or file name. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Known Issues for Using the Windows Firewall Log
Enable or Disable the Windows Firewall Log
View the Windows Firewall Log File
Change the Size of the Windows Firewall Log File

# Change the Size of the Windows Firewall Log File

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to change the size of the Windows Firewall log file. This is useful if you need to increase the size of the log file because you are logging substantial network activity or you want to log activity over a long period of time.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

**To change the size of the Windows Firewall log**

This procedure can be performed using the graphical user interface or the command prompt.

**Using the graphical user interface**

**To change the size of the Windows Firewall log**

1. Open Windows Firewall, and click the **Advanced** tab.
2. In **Security Logging**, click **Settings**.
3. In **Size limit (KB)**, enter the maximum size of the log file, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

**Using the command prompt**

**To change the size of the Windows Firewall log**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set logging maxfilesize =** *maxfilesize*

Substitute values for the placeholders in italics. The following table lists possible values for each placeholder.

| Placeholder | Possible Values | Description |
|---|---|---|
| *maxfilesize* | Any number between 1 and 32,767 kilobytes (KB). | Specifies the maximum size of the log file, in KB. |

If you get an "Access Denied" message when you run a command, you do not have administrative rights to configure Windows Firewall. If you get an "Ok" message but the command does not take effect, the setting might be managed by Group Policy.

**Notes**

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- The default size of the Windows Firewall log file is 4,096 KB.

**See Also**
**Concepts**
Known Issues for Using the Windows Firewall Log
Enable or Disable the Windows Firewall Log
View the Windows Firewall Log File
Change the Name and Location of the Windows Firewall Log File

# Interpreting the Windows Firewall Log

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Windows Firewall log file is useful for determining if Windows Firewall is the cause of program failures. It can also be used to help identify malicious activity, although it does not provide the information needed to track down the source of the activity. The Windows Firewall log file is not useful for analyzing the overall security of your network.

The Windows Firewall log is a plaintext file that can be viewed through any text editor. Notepad is the default text editor for the Windows Firewall log file. Depending on the size limit set for the file and the duration of the logging event, a single log file can contain thousands of text entries. To improve its readability in Notepad, disable word wrapping. This will preserve the column formatting. For more information, see View the Windows Firewall Log File.

The following is an example of a Windows Firewall log file:

```
#Version: 1.5
#Software: Microsoft Windows Firewall
#Time Format: Local
#Fields: date time action protocol src-ip dst-ip src-port dst-port size tcpflags tcpsyn
tcpack tcpwin icmptype icmpcode info path
2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - -
RECEIVE
2005-04-11 08:05:57 DROP UDP 123.45.678.90 255.255.255.255 1631 2234 37 - - - - - - -
RECEIVE
2005-04-11 08:05:58 OPEN UDP 123.45.678.90 123.456.78.90 500 500 - - - - - - - - -
2005-04-11 08:05:58 DROP UDP 123.45.678.90 123.456.78.255 138 138 299 - - - - - - -
RECEIVE
2005-04-11 08:06:02 CLOSE UDP 123.45.678.90 123.456.78.90 1027 53 - - - - - - - - -
2005-04-11 08:06:02 CLOSE UDP 123.45.678.90 123.456.78.90 137 137 - - - - - - - - -
2005-04-11 08:06:05 DROP UDP 0.0.0.0 255.255.255.255 68 67 328 - - - - - - - RECEIVE
2005-04-11 08:06:26 DROP TCP 123.45.678.90 123.456.78.90 80 1774 576 A 123456789
987654321 12345 - - - RECEIVE
2005-04-11 08:06:27 DROP TCP 123.45.678 90 123.456.78.90 80 1774 576 AP 123456789
987654321 12345 - - - RECEIVE
2005-04-11 08:08:58 DROP ICMP 123.45.678.90 123.456.78.90 7 7 78 - - - - 8 0 - RECEIVE
2005-04-11 08:09:29 OPEN TCP 123.45.678.90 123.456.78.90 1606 445 - - - - - - - - -
2005-04-11 08:09:30 CLOSE TCP 123.45.678.90 123.456.78.90 1607 139 - - - - - - - - -
2005-04-11 08:48:46 DROP TCP 123.45.678.90 123.456.78.90 80 1693 40 A 2351482979
694744025 64675 - - - RECEIVE
2005-04-11 08:48:46 DROP TCP 123.45.678.90 123.456.78.90 80 1693 40 FA 2351482979
694744025 64675 - - - RECEIVE
2005-04-11 08:52:26 INFO-EVENTS-LOST - - - - - - - - - - - - - 59 -
```

| **Note** |
| --- |
| The preceding log file example has been truncated. Many dropped broadcast packets, which are the most common type of entry, have been removed from this example. |

The log file has two sections: the header and the body. The header displays information about Windows Firewall and provides a column heading for the entries that appear in the body section.

The following table lists the information contained in the header of the log file:

| Item | Description |
| --- | --- |
| #Version: | Displays which version of the Windows Firewall security log is installed.<br><br>`#Version: 1.5` |
| #Software: | Displays the name of the software creating the log.<br><br>`#Software: Microsoft Windows Firewall` |
| #Time: | Indicates that all of the timestamps in the log are in local time.<br><br>`#Time Format: Local` |

| | |
|---|---|
| #Fields: | Displays a static list of fields that are available for security log entries, if data is available. These fields are listed in the following table.<br><br>**Note**<br>The hyphen (-) is used for fields for which no information is available. |

The body is the report of information gathered about traffic or attempts to pass through Windows Firewall. The body of the security log is a dynamic list; new entries appear at the bottom of the log.

The following table lists the information contained in the body of the log file:

| Item | Description |
|---|---|
| date | Displays the year, month, and day that the recorded transaction occurred. Dates are recorded in the following format:<br><br>YYYY-MM-DD<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |
| time | Displays the hour, minute, and seconds at which the recorded transaction occurred. Times are recorded in the following 24-hour format:<br><br>HH:MM:SS<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |
| action | Displays which operation was observed by Windows Firewall. The options available are OPEN, OPEN-INBOUND, CLOSE, DROP, and INFO-EVENTS-LOST. An INFO-EVENTS-LOST action indicates the number of events that occurred but were not recorded in the log.<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |
| protocol | Displays the protocol that was used for the communication. The options available are TCP, UDP, ICMP, and a protocol number for packets that are not TCP, UDP, or ICMP.<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |
| src-ip | Displays the source IP address (the IP address of the computer attempting to establish communication).<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |
| dst-ip | Displays the destination IP address of a communication attempt.<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |
| src-port | Displays the source port number of the sending computer. Only TCP and UDP display a valid src-port entry. All other protocols display a src-port entry of -.<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |
| dst-port | Displays the port number of the destination computer. Only TCP and UDP display a valid dst-port entry. All other protocols display a dst-port entry of -.<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |

| | |
|---|---|
| size | Displays the packet size, in bytes.<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |
| tcpflags | Displays the TCP control flags found in the TCP header of an IP packet:<br><br>• **A**ck Acknowledgment field significant<br>• **F**in No more data from sender<br>• **P**sh Push function<br>• **R**st Reset the connection<br>• **S**yn Synchronize sequence numbers<br>• **U**rg Urgent Pointer field significant<br><br>`2005-04-11 08:06:27 DROP TCP 123.45.678 90 123.456.78.90 80 1726 54 AP 123456789 987654321 12345 - - - RECEIVE` |
| tcpsyn | Displays the TCP sequence number in the packet.<br><br>`2005-04-11 08:06:278 DROP TCP 123.45.678 90 123.456.78.90 80 1726 54 AP 123456789 987654321 12345 - - - RECEIVE` |
| tcpack | Displays the TCP acknowledgement number in the packet.<br><br>`2005-04-11 08:06:27 DROP TCP 123.45.678 90 123.456.78.90 80 1726 54 AP 123456789 987654321 12345 - - - RECEIVE` |
| tcpwin | Displays the TCP window size, in bytes, in the packet.<br><br>`2005-04-11 08:06:27 DROP TCP 123.45.678 90 123.456.78.90 80 1726 54 AP 123456789 987654321 12345 - - - RECEIVE` |
| icmptype | Displays a number that represents the Type field of the ICMP message.<br><br>`2005-04-11 08:06:27 DROP ICMP 123.45.678.90 123.456.78.90 7 7 78 - - - - 8 0 - RECEIVE` |
| icmpcode | Displays a number that represents the Code field of the ICMP message.<br><br>`2005-04-11 08:06:27 DROP ICMP 123.45.678.90 123.456.78.90 7 7 78 - - - - 8 0 - RECEIVE` |
| info | Displays an entry that depends on the type of action that occurred. For example, an INFO-EVENTS-LOST action will result in an entry of the number of events that occurred but were not recorded in the log from the time of the last occurrence of this event type.<br><br>`2005-04-11 08:52:26 INFO-EVENTS-LOST - - - - - - - - - - - - 59 -` |
| path | Displays the direction of the communication. The options available are SEND, RECEIVE, FORWARD, and UNKNOWN.<br><br>`2005-04-11 08:05:57 DROP UDP 123.45.678.90 123.456.78.255 137 137 78 - - - - - - - RECEIVE` |

**When to perform this task**

- When attempting to determine if program failure is a result of your Windows Firewall configuration.
- When searching for malicious activity on your network.

**Task requirements**

No special tools are required to complete this task.

**Task procedures**

- Identify Program Failures
- Identify Malicious Activity

**See Also**
**Concepts**
Enable or Disable the Windows Firewall Log
Known Issues for Using the Windows Firewall Log
Change the Name and Location of the Windows Firewall Log File
Change the Size of the Windows Firewall Log File

# Identify Program Failures

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to determine if Windows Firewall is the source of your program failures. This procedure is useful if you are troubleshooting.

| **Note** |
| --- |
| If the failing program is a service and it is running as a driver, it is not possible to determine from the log file if the failure is related to Windows Firewall. For information about troubleshooting this type of issue, see A System Service Runs in Svchost.exe and Cannot Be Added to the Exceptions List. |

### Administrative Credentials

No special administrative credentials are required to perform this task.

### Special Considerations

No special considerations are required to perform this task.

#### To identify program failures

You can use this procedure to find any of the following, which can be useful for identifying the source of your problem:

- Disabled port openings.
- Dynamic port openings.
- Dropped packets with push and urgent flags.
- Dropped packets on the send path.

This method will verify if disabled ports remain closed.

#### To identify disabled port openings

1. With the Windows Firewall log file open in Notepad, press **Ctrl** + **F**, type the number of port that you know to be disabled, and then click **Find Next**.
2. Make sure each log entry corresponding to the port number in the **dst-port** field does not have OPEN in the **action** field.
3. If you find entries corresponding to the port number in which OPEN appears in the action field, write down the source IP address (**srcip**) and the destination IP address (**dstip**), and then see Fixing Program Problems for information about troubleshooting the issue.

This method will verify if Windows Firewall is allowing dynamic port openings for programs that use them.

#### To identify dynamic port openings

1. With the Windows Firewall log open in Notepad, scroll to the bottom of the file.
2. Launch the program and watch the log while it is running.
3. Look at the entries for the port used by the program. If those entries have OPEN-INBOUND in the action field, then Windows Firewall is not the source of the problem. If the port is not opening, see Fixing Program Problems for information about troubleshooting the issue.

The presence of dropped packets that are flagged with push or urgent might indicate that there is a problem with one of your programs.

#### To identify dropped packets with push and urgent flags

1. With the Windows Firewall log file open in Notepad, press **Ctrl** + **F**, type [space]U[space] or [space]PU[space], and then

click **Find Next**.

2. Make sure each log entry with U or UP in the **tcpflags** field does not have DROP in the **action** field.

3. If you find entries with U or PU and DROP together, write down the source IP address (**srcip**), the destination IP address (**dstip**), and the source port (**srcport**), and then see Fixing Program Problems for information about troubleshooting the issue.

The presence of dropped packets on the send path is rare, but might indicate that there is a problem with one of your programs.

**To identify dropped packets on the send path**

1. With the Windows Firewall log file open in Notepad, press **Ctrl** + **F**, type SEND, and then click **Find Next**.

2. Ensure each log entry with SEND in the **path** field does not have DROP in the **action** field.

3. If you find entries with SEND and DROP together, write down the source IP address (**srcip**) and source port (**srcport**), and then see Fixing Program Problems for information about troubleshooting the issue.

**See Also**
**Concepts**
Interpreting the Windows Firewall Log
View the Windows Firewall Log File

# Identify Malicious Activity

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to confirm that malicious activity is occurring within your network. This procedure is useful if you suspect such activity is taking place and you want to identify the targets of the attacks.

**Administrative Credentials**

No special administrative credentials are required to perform this task.

**Special Considerations**

No special considerations are required to perform this task.

**To identify malicious activity**

| **Note** |
| --- |
| Broadcast packets have destination IP addresses that end in 255. Disregard those packets when you search the Windows Firewall log file for malicious activity. |

**To identify malicious activity**

1. With the Windows Firewall log file open in Notepad, scroll through the file from beginning to end.
2. Look at each log entry with DROP in the **action** field and note whether the destination IP address (**dstip**) ends with a number other than 255.
3. If you find many such entries, take note of the destination IP addresses of the packets.
4. If those destination IP addresses are all the same, similar, or systematic, write down the source IP addresses (**srcip**) and the destination IP addresses. These dropped packets can be considered suspicious. Suspicious dropped packets often have systematic port hit entries as well.

**See Also**
**Concepts**
Interpreting the Windows Firewall Log
View the Windows Firewall Log File

# Using the Security Log

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the security log to track Windows Firewall events. The security log provides information about Windows Firewall activity, such as changes to Windows Firewall settings, Windows Firewall startup status, and notifications sent when programs and system services attempt to listen for incoming traffic and are blocked. For more information about the specific events that are written to the security log, see the section titled "Windows Firewall Tools and Settings" in the Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729).

By default, Windows Firewall events are not written to the security log. You must enable auditing in Group Policy to use the security log to track Windows Firewall events.

**When to perform this task**

You should enable the auditing of Windows Firewall events when you turn on Windows Firewall for the first time, when you need to troubleshoot Windows Firewall problems, or when you need to temporarily monitor Windows Firewall behavior.

**Task requirements**

No special tools are required to complete this task.

**Task procedures**

To complete this task, use the following procedures:

- Enable Auditing of Windows Firewall Events
- View Windows Firewall Events in Event Viewer

**See Also**
**Concepts**
Using the Windows Firewall Log
Best Practices for Monitoring Windows Firewall

# Known Issues for Using the Security Log

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review the following known issues before you use the Windows Firewall log file.

- You must enable the **Audit process tracking** and **Audit policy change** settings in Group Policy to have Windows Firewall events written to the security event log.

# Enable Auditing of Windows Firewall Events

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to enable auditing so that Windows Firewall events are written to the security log. This is useful for troubleshooting Windows Firewall problems or when you are configuring Windows Firewall for the first time and you want to monitor Windows Firewall behavior.

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

### Special Considerations

No special considerations are required to perform this procedure.

**To enable auditing of Windows Firewall events**
This procedure can be performed using Group Policy.

### Using Group Policy

### To enable auditing of Windows Firewall events

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Group Policy settings in your organization.
2. Open **Computer Configuration**, open **Windows Settings**, open **Security Settings**, open **Local Policies**, and then click **Audit Policy**.
3. Double-click **Audit process tracking**, select the **Success** and **Failure** check boxes, and then click **OK**.
4. Double-click **Audit policy change**, select the **Success** and **Failure** check boxes, and then click **OK**.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

### Notes

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Using the Security Log
Known Issues for Using the Security Log
View Windows Firewall Events in Event Viewer

# View Windows Firewall Events in Event Viewer

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use this procedure when you want to view Windows Firewall events that have been written to the security log. This is useful if you are troubleshooting Windows Firewall problems or you are configuring Windows Firewall for the first time and you want to monitor Windows Firewall behavior.

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

### Special Considerations

There are no special considerations to perform this procedure.

**To view Windows Firewall events in Event Viewer**
This procedure can be performed using the graphical user interface.

### Using the graphical user interface

### To view Windows Firewall events in Event Viewer

1. Open Event Viewer.
2. Click **Security**, and then look for event IDs 848 through 861.

If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

### Notes

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Using the Security Log
Known Issues for Using the Security Log
Enable Auditing of Windows Firewall Events

# Securing Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Firewall can be configured only by users who are members of the Administrators group or members of a group that is a member of the Administrators group. You can configure Windows Firewall so that local administrators are prevented from enabling or disabling Windows Firewall or configuring Windows Firewall settings. This is useful if you have a centrally managed environment and you want to strictly enforce your Windows Firewall configuration throughout your organization.

To help you secure Windows Firewall, review Best Practices for Securing Windows Firewall and Known Issues for Securing Windows Firewall.

The following tasks are described in this objective:

- Preventing Administrators from Turning Windows Firewall On or Off
- Preventing Administrators from Creating Exceptions

# Best Practices for Securing Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Use the following best practice when you secure Windows Firewall.

**Do not let users configure Windows Firewall settings.**

Do not grant users local administrative rights so they can configure Windows Firewall. By default, you must be a member of the Administrators group (or a member of a group that is a member of the Administrators group) to configure Windows Firewall. Giving users the ability to configure Windows Firewall is not recommended because it can result in individualized configurations that are difficult and costly to troubleshoot. It can also increase your organization's attack surface.

**See Also**
**Concepts**
Best Practices for Managing Windows Firewall
Best Practices for Optimizing Windows Firewall
Best Practices for Monitoring Windows Firewall

# Known Issues for Securing Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Review the following known issues before you secure Windows Firewall.

- If you do not configure the **Windows Firewall: Allow local program exceptions** Group Policy setting, the ability of administrators to define local program exceptions depends on the configuration of the **Windows Firewall: Define program exceptions** Group Policy setting. If the **Windows Firewall: Define program exceptions** Group Policy setting is not configured, administrators can define local program exceptions. If it is enabled or disabled, administrators cannot define local program exceptions.

- If you do not configure the **Windows Firewall: Allow local port exceptions** Group Policy setting, the ability of administrators to define local port exceptions depends on the configuration of the **Windows Firewall: Define port exceptions** Group Policy setting. If the **Windows Firewall: Define port exceptions** Group Policy setting is not configured, administrators can define local port exceptions. If it is enabled or disabled, administrators cannot define local port exceptions.

- If the **Prohibit use of Internet Connection Firewall on your DNS domain network** setting is enabled and the **Windows Firewall: Protect all network connections** Group Policy setting is enabled, then Windows Firewall is enabled when connected to the network from which the Group Policy object (GPO) was obtained with new Windows Firewall settings.

# Preventing Administrators from Turning Windows Firewall On or Off

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, you must be a member of the Administrators group (or a member of a group that is a member of the Administrators group) to enable or disable Windows Firewall. This prevents users from inadvertently turning Windows Firewall on or off, which can result in individualized configurations that are difficult to troubleshoot and can reduce your organization's overall security.

You can secure Windows Firewall even further by preventing local administrators from enabling or disabling Windows Firewall. This is useful if you rely on Windows Firewall and you always want it enabled, or you use a non-Microsoft host firewall and you always want Windows Firewall disabled. Preventing local administrators from turning Windows Firewall on or off is also useful in a centrally-managed environment, such as a Group Policy environment or an environment in which you want to strictly enforce Windows Firewall configuration and policy settings.

**When to perform this task**

You should perform this task when required by your organization's security plan or when you want to strictly enforce Windows Firewall configuration and policy settings.

**Task requirements**

No special tools are required to complete this task.

**Task procedures**

To complete this task, perform the following procedures:

Prevent Local Administrators from Turning Windows Firewall On or Off

**See Also**
**Concepts**
Best Practices for Securing Windows Firewall
Known Issues for Securing Windows Firewall
Preventing Administrators from Creating Exceptions

# Prevent Local Administrators from Turning Windows Firewall On or Off

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is useful if you want to prevent local administrators from turning Windows Firewall on or off.

**Administrative Credentials**

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Special Considerations**

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should verify scope settings for any exceptions that you change. For more information about scope settings, see Configuring Scope Settings.

**To prevent local administrators from turning Windows Firewall on or off**
This procedure can be performed using Group Policy. You cannot perform this procedure from the command prompt with the **netsh firewall** command or in the graphical user interface with Windows Firewall in Control Panel.

**Using Group Policy**

**To prevent local administrators from turning Windows Firewall on or off**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then open either **Domain Profile** or **Standard Profile**, depending on which profile you want to configure.
3. In the details pane, double-click **Windows Firewall: Protect all network connections**.
4. Do the following if you want to turn off Windows Firewall and prevent local administrators from turning on Windows Firewall:

   In the **Windows Firewall: Protect all network connections** properties dialog box, on the **Settings** tab, click **Disabled**, and then click **OK**.
5. Do the following if you want to turn on Windows Firewall and prevent local administrators from turning off Windows Firewall:

   In the **Windows Firewall: Protect all network connections** properties dialog box, on the **Settings** tab, click **Enabled**, and then click **OK**.

**Notes**

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Group Policy settings must be refreshed before they take effect.

**See Also**
**Concepts**
Preventing Administrators from Turning Windows Firewall On or Off
Known Issues for Securing Windows Firewall

# Preventing Administrators from Creating Exceptions

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, you must be a member of the Administrators group (or a member of a group that is a member of the Administrators group) to configure Windows Firewall exceptions. This prevents users from inadvertently configuring program and port exceptions, which can result in individualized configurations that are difficult to troubleshoot and can reduce your organization's overall security.

You can secure Windows Firewall even further by preventing local administrators from configuring Windows Firewall exceptions. This is useful if you have a centrally-managed environment, such as a Group Policy environment or an environment in which you want to strictly enforce Windows Firewall configuration and policy settings.

**When to perform this task**

You should perform this task when required by your organization's security plan or when you want to strictly enforce Windows Firewall configuration and policy settings.

**Task requirements**

No special tools are required to complete this task.

**Task procedures**

To complete this task, perform the following procedures:

Prevent Local Administrators from Creating Program Exceptions

Prevent Local Administrators from Creating Port Exceptions

**See Also**
**Concepts**
Best Practices for Securing Windows Firewall
Known Issues for Securing Windows Firewall
Preventing Administrators from Turning Windows Firewall On or Off

# Prevent Local Administrators from Creating Program Exceptions

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is useful if you want to prevent local administrators from using Windows Firewall in Control Panel or the **netsh firewall** command to configure program exceptions.

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

### Special Considerations

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should verify scope settings for any exceptions that you change. For more information about scope settings, see Configuring Scope Settings.

**To prevent local administrators from creating program exceptions**
This procedure can be performed using Group Policy. You cannot perform this procedure from the command prompt with the **netsh firewall** command or in the graphical user interface with Windows Firewall in Control Panel.

### Using Group Policy

### To prevent local administrators from creating program exceptions

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then open either **Domain Profile** or **Standard Profile**, depending on which profile you want to configure.
3. In the details pane, double-click **Windows Firewall: Allow local program exceptions**.
4. In the **Windows Firewall: Allow local program exceptions** properties dialog box, on the **Settings** tab, click **Disabled**, and then click **OK**.

### Notes

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Group Policy settings must be refreshed before they take effect.

**See Also**
**Concepts**
Preventing Administrators from Creating Exceptions
Known Issues for Securing Windows Firewall
Prevent Local Administrators from Creating Port Exceptions

# Prevent Local Administrators from Creating Port Exceptions

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure is useful if you want to prevent local administrators from using Windows Firewall in Control Panel or the **netsh firewall** command to configure port exceptions.

### Administrative Credentials

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

### Special Considerations

You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

For more information about Windows Firewall profiles, see Managing Windows Firewall Profiles.

You should verify scope settings for any exceptions that you change. For more information about scope settings, see Configuring Scope Settings.

**To prevent local administrators from creating port exceptions**
This procedure can be performed using Group Policy. You cannot perform this procedure from the command prompt with the **netsh firewall** command or in the graphical user interface with Windows Firewall in Control Panel.

### Using Group Policy

### To prevent local administrators from creating port exceptions

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then open either **Domain Profile** or **Standard Profile**, depending on which profile you want to configure.
3. In the details pane, double-click **Windows Firewall: Allow local port exceptions**.
4. In the **Windows Firewall: Allow local port exceptions** properties dialog box, on the **Settings** tab, click **Disabled**, and then click **OK**.

### Notes

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Group Policy settings must be refreshed before they take effect.

**See Also**
**Concepts**
Preventing Administrators from Creating Exceptions
Known Issues for Securing Windows Firewall
Prevent Local Administrators from Creating Program Exceptions

# Additional Resources for Administering Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For general information about how Windows Firewall works, see the following resources:

- Manually Configuring Windows Firewall in Windows XP Service Pack 2 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=33763).
- New Networking Features in Microsoft Windows XP Service Pack 2 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=33684).
- Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729).

For specific information about troubleshooting Windows Firewall problems, see the following resources:

- Troubleshooting Windows Firewall
- Troubleshooting Windows Firewall Settings in Windows XP Service Pack 2 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43102).
- Article ID 842242 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43103). This article describes why some programs might stop working after you install Windows XP with Service Pack 2 (SP2) and what you can do to overcome the problem.
- Article ID 884130 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43104). This article contains a list of programs that are known to experience loss of functionality when they run on Windows XP SP2.

# Troubleshooting Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This guide provides troubleshooting information for Windows Firewall in Microsoft® Windows Server™ 2003 with Service Pack 1 (SP1). It is designed to help you identify and resolve problems that are caused by Windows Firewall, including network access problems, problems running programs, and configuration and management problems.

**In this guide**


- Introduction to Troubleshooting Windows Firewall
- Verifying Computer Settings for Windows Firewall
- Windows Firewall Quick Fixes
- Configuring a Computer for Windows Firewall Troubleshooting
- Troubleshooting Windows Firewall Problems
- Additional Resources for Windows Firewall Troubleshooting



**Acknowledgments**

Produced by: Microsoft Windows Server User Assistance team

Project Writer: Bill Gruber

Project Editor: Allyson Adley

Technical Reviewers: David Beder, Joseph Davies, Erik Johnson, Michael Surkan

# Introduction to Troubleshooting Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

If you are using this guide for the first time, review the following sections of this introduction.

**When to Use This Guide**
You should use this guide when:

- You have a problem that you believe is related to Windows Firewall, but you do not know how to resolve it.
- You have configured Windows Firewall settings, but Windows Firewall does not behave the way you anticipated.
- A program is not working properly and you believe that Windows Firewall is causing the problem.

Do not use this guide to find out how to perform a task, such as enabling Windows Firewall or adding a program or port to the Windows Firewall exceptions list. Information about how to perform tasks and configure settings can be found in Administering Windows Firewall.

This guide assumes a basic understanding of what Windows Firewall is, how it works, and why your organization uses it to enhance network security. You should also have a thorough understanding of how Windows Firewall is deployed and managed in your organization. This includes an understanding of the mechanism your organization uses to configure and manage Windows Firewall settings.

**How to Use This Guide**
This guide is divided into four sections. Each section addresses a type of problem, ranging from less complex to more complex.

Verifying Computer Settings for Windows Firewall

This section provides a list of prerequisites and settings that must be verified before you troubleshoot. Read this section first.

Windows Firewall Quick Fixes

This section provides quick solutions to the most common Windows Firewall problems. Read this section if you have verified your computer settings, but your problem persists.

Configuring a Computer for Windows Firewall Troubleshooting

This section describes how to configure your computer for troubleshooting. Read this section if you cannot find a solution in the Quick Fixes section.

Troubleshooting Windows Firewall Problems

This section provides step-by-step diagnostic procedures and possible solutions that help you identify and fix Windows Firewall problems.

Additional Resources for Windows Firewall Troubleshooting

**Terminology Used in This Guide**
**Domain profile**
A set of Windows Firewall settings that are applied when a computer is connected to a network that contains the domain controllers for the domain in which its computer account resides. The domain profile is one of two profiles that Windows Firewall uses to apply settings to a computer.

**Exception**
A port, program, or system service that is allowed to receive unsolicited traffic.

**Exceptions list**
A special type of rules store that specifies which ports, programs, or system services are allowed to receive unsolicited traffic. To allow unsolicited incoming traffic through Windows Firewall, you add ports, programs, and system services to the exceptions list.

**Notification**

A **Windows Security Alert** dialog box that appears when a program that is not listed in the exceptions list attempts to listen for unsolicited traffic on a port. If you have administrative rights on a computer, the notification displays an option to add the program to the exceptions list, and thereby allow the program to receive unsolicited traffic. If you do not have administrative rights on a computer, the notification displays a warning that a program is attempting to listen on a port. A **Windows Security Alert** dialog box does not appear when a system service attempts to listen for unsolicted incoming traffic; it only appears when a program attempts to listen for unsolicited incoming traffic.

**Port**
A logical communication endpoint representing a service or an application that listens for and receives IP packets. Ports are specified by a positive 16-bit decimal number and the type of traffic that is expected to pass through the port--either UDP traffic or TCP traffic. Most ports have been predefined and are considered "well-known" for specific services, such as DHCP and DNS. However, some ports are created dynamically by a server and assigned to the services and applications that need to listen for incoming traffic from clients.

**Program**
A software application that is usually started by a user and runs under the user's account. Programs run only while the user is logged on to a computer and usually consist of one or more executable (.exe) files and one or more dynamic-link library (.dll) files.

| **Note** |
| --- |
| Some system services run within their own .exe file and are started by a user. However, these system services typically run under a priviledged account, such as Local Service, and do not run under the user's account. An example of this is the Telnet service, which runs in Tlntsvr.exe. |

**Scope**
A Windows Firewall setting that you configure for an exception. The scope setting controls which addresses unsolicited traffic is allowed to originate from. By default, the scope of an exception is any address, which includes any computer on the Internet. You can change the scope of an exception to locally reachable addresses or a list of individual Internet Protocol version 4 (IPv4) addresses or IPv4 address ranges.

**System service**
A software application that is usually not started by a user and runs under a privileged account, such as LocalSystem. System services run even when a user is not logged on to a computer and usually run as a separate process within Svchost.exe.

**Standard profile**
A set of Windows Firewall settings that are applied when a computer is not connected to a network that contains the domain controllers in which its computer account resides (for example, a public network, such as the Internet). The standard profile is one of two profiles that Windows Firewall uses to apply settings to a computer.

# Verifying Computer Settings for Windows Firewall

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you begin troubleshooting, you need to verify that your computer is configured properly and that Windows Firewall is set up and running properly.

**Settings to Verify Before You Troubleshoot**
Verify all of the following items:

Make sure you have administrative rights on the computer you are troubleshooting.

Install all critical updates and security updates for Windows Server 2003.

Update all of your software, including non-Microsoft software.

Make sure that Windows Firewall is enabled when you experience the problem.

Verify that the Group Policy setting, Prohibit use of Internet Connection Firewall on your domain network, is either disabled or not configured.

Make sure that Windows Firewall/Internet Connection Sharing service, Network Connections service, Remote Procedure Call service, and Windows Management Instrumentation service are started.

Restart the computer if you are running a server operating system.

**Make sure you have administrative rights on the computer you are troubleshooting.**
You cannot modify Windows Firewall settings unless you are a member of the Administrators group on the computer that you are administering.

### To verify that you are a member of the Administrators group

1. Open the Computer Management snap-in.
2. In the console tree, double-click **Local Users and Groups**, and then click **Groups**.
3. In the details pane, double-click **Administrators** and verify that your account name or a group to which your account is a member appears in the **Members** list.

You can also tell if you have the appropriate administrative rights to configure Windows Firewall by opening Windows Firewall in Control Panel. If you do not have administrative rights, all of the controls in the Windows Firewall user interface (UI) will appear dimmed and a warning will appear telling you that you need to be a computer administrator to change any settings.

**Install all critical updates and security updates for Windows Server 2003.**
Some updates might be required for Windows Firewall to function properly.

### To verify that you have all critical updates and security patches

- Click **Start**, click **Windows Update**, and then follow the instructions that appear on your screen.

**Update all of your software, including non-Microsoft software.**
Windows Firewall might not function properly with some programs unless you update the programs with the most recent service pack or software update. Newer versions of many applications, such as antivirus programs, will automatically configure Windows Firewall and might resolve any problems you have.

**Make sure that Windows Firewall is enabled when you experience the problem.**
If Windows Firewall is not enabled when you experience your problem, then Windows Firewall cannot be causing your problem.

### To verify that Windows Firewall is enabled

1. In Control Panel, open **Windows Firewall**.
2. In the **Windows Firewall** dialog box, verify that **On** is selected.

If **On** appears dimmed, you do not have administrative rights on the computer or your Windows Firewall settings are managed

through Group Policy.

**Verify that the Group Policy setting, Windows Firewall: Prohibit use of Internet Connection Firewall on your domain network, is either disabled or not configured.**
If enabled, this setting prevents anyone, including administrators, from enabling or configuring Windows Firewall. To change this policy setting, use the Group Policy Object Editor snap-in to edit the Group Policy objects (GPOs) that are used to manage Windows Firewall settings in your organization.

### To modify the Prohibit use of Internet Connection Firewall on your domain network setting

1. Open the Group Policy Object Editor snap-in to edit the GPO that is used to manage Windows Firewall settings in your organization.
2. Click **Computer Configuration**, click **Administrative Templates**, click **Network**, and then click **Network Connections**.
3. In the details pane, double-click the **Windows Firewall: Prohibit use of Internet Connection Firewall on your domain network** policy setting.
4. Select either the **Disabled** or **Not Configured** check box.

If this policy setting is enabled and you cannot configure it, you do not have administrative rights to configure Group Policy settings.

**Make sure that Windows Firewall/Internet Connection Sharing service, Network Connections service, Remote Procedure Call service, and Windows Management Instrumentation service are started.**

### To verify that a system service is started

1. Open the Services snap-in.
2. In the details pane, double-click the system service that you want to verify, and then check the **Service status**.

**Restart the computer if you are running a server operating system.**
If you turned on the Windows Firewall/Internet Connection Sharing service after you started a program or a system service that is listed in the program exceptions list, you should do the following:

- Restart your computer.
- Start the programs that are listed in the exceptions list.

Windows Firewall cannot track the state of a program's traffic or a system service's traffic if the program or system service is started before the Windows Firewall/Internet Connection Sharing service is started. This problem typically occurs when you stop the Windows Firewall/Internet Connections Sharing service to install and configure software, and then turn on the Windows Firewall/Internet Connection Sharing service after you start the software program.

# Windows Firewall Quick Fixes

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Most Windows Firewall problems do not require extensive troubleshooting and can be fixed by making a simple configuration change. This section provides the most common Windows Firewall problems and their solutions. You can use the information in this section to resolve problems in the same way you would use a FAQ to find answers to common questions. Read this section before you begin any advanced troubleshooting.

**Common Windows Firewall Problems**
Listed below are common Windows Firewall problems and quick fixes:


I turned on Windows Firewall, but the programs in the exceptions list are not working properly.

Windows Firewall displays a Windows Security Alert when a program attempts to listen on a port, but the alert tells me to contact my system administrator and does not prompt me to unblock the program.

Windows Firewall does not display a Windows Security Alert when a program or system service attempts to listen on a port.

No one can ping my computer with the ping command.

I disabled ping by blocking incoming echo traffic, but other computers can still ping my computer.

I cannot access shared folders or printers on a computer that has Windows Firewall enabled.

I cannot use Remote Desktop to access another computer.

I cannot use remote administration tools to manage my servers.

I cannot get my FTP program to work with Windows Firewall.

I added a program to the exceptions list in Windows Firewall in Control Panel, but the program does not show up in the exceptions list when I take the computer off the managed network.

Some or all of the settings in Windows Firewall appear dimmed.

I added a program to the exceptions list in Group Policy, but the program does not behave as though it's in the exceptions list.

I used the My network (subnet) only scope option, but traffic originating beyond my subnet is getting through Windows Firewall.

I used the My network (subnet) only option to change the scope of a program exception, but some computers in my subnet are being blocked by Windows Firewall.


**I turned on Windows Firewall, but the programs in the exceptions list are not working properly.**
This can happen on a server if you turn on Windows Firewall after you have started a program (for example, if you have started a Web server or an e-mail server and then start Windows Firewall). Windows Firewall might not be able to track the state of a program's network traffic if the program is started before you start Windows Firewall. This is true no matter how you start Windows Firewall (through Control Panel, Group Policy, or the command prompt).

**Windows Firewall displays a Windows Security Alert when a program attempts to listen on a port, but the alert tells me to contact my system administrator and does not prompt me to unblock the program.**
There are two conditions under which you might see this Windows Security Alert (referred to throughout this documentation as a notification).

Windows Firewall displays this notification when a program attempts to listen on a port and the user who started the program is not a member of the Administrators group. If you are not a member of the Administrators group, you cannot configure the exceptions list, so Windows Firewall does not prompt you to unblock the program by adding the program to the exceptions list. To change the notification so that you are prompted to unblock a program, you must log on to the computer using an account that is a member of the Administrators group. You can then either wait for the notification to appear and click **Unblock**, or you can use Windows Firewall to add the program to the exceptions list.

**To add a program to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.

2. On the **Exceptions** tab, click **Add Program**, and then follow the instructions that appear on your screen.

Windows Firewall also displays this notification if the **Windows Firewall: Allow local program exceptions** policy setting is disabled or if the **Windows Firewall: Allow local program exceptions** policy setting is not configured and the **Windows Firewall: Define program exceptions** policy setting is either disabled or enabled. Windows Firewall displays this notification under these conditions even if you are a member of the Administrators group. To change the notification so that you are prompted to unblock a program, you must either enable the **Windows Firewall: Allow local program exceptions** policy setting or, if this policy setting is not configured, change the **Windows Firewall: Define program exceptions** policy setting to **Not Configured**.

**To enable the Allow local program exceptions policy setting**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then click either **Domain Profile** or **Standard Profile**.
3. Double-click **Windows Firewall: Allow local program exceptions**, and on the **Settings** tab, click **Enabled**.

**To set the Define program exceptions policy setting to Not Configured**

1. Open the Group Policy Object Editor snap-in to edit the GPO that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then click either **Domain Profile** or **Standard Profile**.
3. Double-click **Windows Firewall: Define program exceptions**, and on the **Settings** tab, click **Not Configured**.

**Windows Firewall does not display a Windows Security Alert when a program or system service attempts to listen on a port.**
There are several reasons Windows Firewall might not display a Windows Security Alert (notification) when a program or system service attempts to listen on a port.

- Notifications are displayed only if you enable the display notification option and have disabled or not configured the **Windows Firewall: Prohibit notifications** policy setting.

**To enable notifications**

1. Open Windows Firewall.
2. On the **Exceptions** tab, select the **Display a notification when Windows Firewall blocks a program** check box.

**To disable the Prohibit notifications policy setting**

1. Open the Group Policy Object Editor snap-in to edit the GPO that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then click either **Domain Profile** or **Standard Profile**.
3. Double-click **Windows Firewall: Prohibit Notifications**, and on the **Settings** tab, select either **Not Configured** or **Disabled**.

- Windows Firewall will not display notifications for system services or programs that run when there is no user logged on to the computer. Most system services run this way, as do many server programs. However, Windows Firewall does write an event (Event ID 861) to the security event log for any system service or program that attempts to listen on a port, no matter if a notification is displayed. You must enable auditing for Windows Firewall events for Windows Firewall events to appear in the security event log.

**To enable auditing of Windows Firewall events**

1. Open the Group Policy Object Editor snap-in to edit the GPO that is used to manage Windows Firewall settings in your organization.

2. Open **Windows Settings**, open **Security Settings**, open **Local Policies**, and then click **Audit Policy**.
3. Double-click **Audit process tracking**, select the **Success** and **Failure** check boxes, and then click **OK**.
4. Double-click **Audit policy change**, select the **Success** and **Failure** check boxes, and then click **OK**.

| **Note** |
| --- |
| The only way to determine when a system service listens on a port is to enable auditing and view security event log entries. |

**To view Windows Firewall notification events in Event Viewer**

1. Open Event Viewer.
2. Click **Security**, and then look for Event ID 861.

- Windows Firewall cannot display notifications if you select the **Don't allow exceptions** check box in Windows Firewall in Control Panel or if you enable the **Windows Firewall: Do not allow exceptions** policy setting in the Group Policy Object Editor. There are no Windows Firewall settings that allow you to override this behavior. To display notifications, you must clear the **Don't allow exceptions** check box or set the **Windows Firewall: Do not allow exceptions** policy setting to **Disabled** or **Not Configured**; however, doing so lessens the level of security that Windows Firewall provides when you are connected to a public network, such as the Internet. Instead of allowing exceptions through Windows Firewall so notifications are displayed, it is recommended that you use Event Viewer to identify programs and system services that attempt to listen on a port.

**To configure Windows Firewall to allow exceptions**

1. Open Windows Firewall in Control Panel.
2. On the **General** tab, clear the **Don't allow exceptions** check box.

**To configure Group Policy so Windows Firewall allows exceptions**

1. Open the Group Policy Object Editor snap-in to edit the GPO that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then click either **Domain Profile** or **Standard Profile**.
3. Double-click **Windows Firewall:Do not allow exceptions**, and on the **Settings** tab, click either **Not Configured** or **Disabled**.

- Windows Firewall does not display notifications for programs that rely on the Winsock driver to dynamically bind to a UDP port (sometimes referred to as wildcard binds). If a program uses this method to bind to a UDP port, you might be able to use the **netstat** command and other troubleshooting tools to determine which UDP port is being used. For more information, see Fixing Program Problems.

**No one can ping my computer with the ping command.**

This happens because, by default, Windows Firewall blocks all Internet Control Message Protocol (ICMP) traffic. When someone pings your computer, an ICMP Echo message (also known as an Echo Request message) is sent to your computer; your computer sends back an ICMP Echo Reply message to the computer that sent the ICMP Echo message. By default, Windows Firewall discards the incoming ICMP Echo message. To configure your computer so that it does not discard ICMP Echo messages and can respond with ICMP Echo Reply messages, you must configure Windows Firewall to allow incoming echo traffic.

**To allow incoming echo request traffic**

1. Open Windows Firewall, and then click the **Advanced** tab.
2. In **ICMP**, click **Settings**.
3. In the **ICMP Settings** dialog box, select the **Allow incoming echo request** check box.

**I disabled ping by blocking incoming echo traffic, but other computers can still ping my computer.**

This happens when TCP port 445 is added to the exceptions list. By default, TCP port 445 is added to the exceptions list when you enable **File and Printer Sharing** in the Windows Firewall exceptions list or when you enable the **Windows Firewall: Allow remote administration exception** policy setting in Group Policy. To disable ping, you must remove TCP port 445 from the exceptions list. You can do this by disabling **File and Printer Sharing** in the Windows Firewall exceptions list or by editing **File and Printer Sharing** in the exceptions list settings so that TCP port 445 is disabled. You can also do this by disabling the **Windows Firewall: Allow remote administration exception** policy setting.

**To disable the File and Printer Sharing exception so ping is disabled**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Clear the **File and Printer Sharing** check box.

| **Note** |
| --- |
| Disabling the File and Printer Sharing exception will prevent users from accessing shared files, folders, and printers on your computer. |

**To edit the File and Printer Sharing exception so ping is disabled**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Click **File and Printer Sharing**, and then click **Edit**.
3. In the **Edit a Service** dialog box, clear the **TCP 445** check box.

| **Note** |
| --- |
| Editing the File and Printer Sharing exception can prevent other computers from accessing your shared folders, files, and printers. Be sure to test for the impact these changes might have on file and printer sharing functionality. |

**To disable the Allow remote administration exception so ping is disabled**

1. Open the Group Policy Object Editor snap-in to edit the GPO that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then click either **Domain Profile** or **Standard Profile**.
3. In the details pane, double-click **Windows Firewall: Allow remote administration exception**, and then click either **Not Configured** or **Disabled**.

| **Note** |
| --- |
| Disabling or not configuring the **Windows Firewall: Allow remote administration exception** policy setting can cause your remote administration tools and other programs to stop running properly. |

**I cannot access shared folders or printers on a computer that has Windows Firewall enabled.**
This happens when Windows Firewall blocks incoming traffic through UDP ports 137 and 138 and TCP ports 139 and 445. To allow access to shared folders and printers, you must enable the predefined exception for file and printer sharing, which will dynamically allow incoming traffic on UDP ports 137 and 138 and TCP ports 139 and 445.

**To enable the File and Printer Sharing exception**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Select the **File and Printer Sharing** check box.

**I cannot use Remote Desktop to access another computer.**
This happens when Windows Firewall blocks incoming traffic through TCP port 3389 on the computer you are trying to access. To use Remote Desktop to communicate with a remote computer, you must enable the predefined exception for Remote Desktop on the remote computer, which will dynamically allow incoming traffic on TCP port 3389.

**To enable the Remote Desktop exception**

1. On the remote computer, open Windows Firewall, and then click the **Exceptions** tab.
2. Select the **Remote Desktop** check box.

**I cannot use remote administration tools to manage my servers.**
This happens when Windows Firewall blocks incoming remote procedure call (RPC) or DCOM traffic on your servers. Remote administration tools, such the Microsoft Management Console (MMC) and Windows Management Instrumentation (WMI), use RPC and DCOM to communicate with servers. By default, Windows Firewall blocks communication through TCP ports 135 and 445. To fix this problem, you can enable the **Windows Firewall: Allow remote administration exception**policy setting on your servers. If you enable this policy setting, you must specify the Internet Protocol version 4 (IPv4) addresses or address ranges from which these incoming messages are allowed.

**To enable the Allow remote administration exception**

1. Open the Group Policy Object Editor snap-in to edit the GPO that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then click either **Domain Profile** or **Standard Profile**.
3. In the details pane, double-click **Windows Firewall: Allow remote administration exception**, click **Enabled**, and then, in **Allow unsolicited incoming messages from**, type the IPv4 addresses or address ranges for this exception.

| **Important** |
| --- |
| You should enable the **Windows Firewall: Allow remote administration exception** policy setting only if your remote administrative tools require RPC or DCOM. Malicious users often attempt to attack networks and computers using RPC and DCOM. It is recommended that you contact the manufacturer of your remote administration tool to determine if it requires RPC and DCOM communication. If it does not, do not enable the **Windows Firewall: Allow remote administration exception** policy setting. |

**I cannot get my FTP program to work with Windows Firewall.**
This happens when your FTP program uses active FTP instead of passive (PASV) FTP. Windows Firewall uses the Application Layer Gateway (ALG) Service to process active FTP traffic. If the ALG Service is not running, try starting it.

**To start the Application Layer Gateway Service**

1. Open the Services snap-in.
2. Double-click **Application Layer Gateway Service**.
3. On the **General** tab, if the **Service status** is not **Started**, click **Start**.

If the ALG Service is running and you still cannot get your FTP program to work with Windows Firewall, try adding the FTP program to the exceptions list.

**To add a program to the Windows Firewall exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Click **Add Program**, and then follow the directions that appear on your screen.

**I added a program to the exceptions list in Windows Firewall in Control Panel, but the program does not show up in the exceptions list when I take the computer off the managed network.**
This happens when you add a program to the exceptions list for one profile (for example, the domain profile), and the computer is using the exceptions list for the other profile (for example, the standard profile). An example of this occurs when a portable computer is running on your corporate domain and you add a program to the exceptions list by using Windows Firewall in Control Panel. In this case, the program is added to the exceptions list for the domain profile. If the portable computer is then used outside of the office and is not connected to the corporate domain, the exceptions list for the standard profile is used; the program will not show up in the exceptions list because it was added to the exceptions list for the domain profile only. To fix this problem, you must do the following:

- Determine to which exceptions list (domain profile or standard profile) the program was added.
- Add the program to the appropriate exceptions list.

**To determine to which exceptions list a program has been added**

- At the command prompt, type **netsh firewall show all**, and press ENTER.

**To add a program to the exceptions list for a specific profile**

1. To add a program to the domain profile, type the following:

   **netsh firewall set allowedprogram** *your_program_path\your_program your_program_name* **enable domain**

   Where *your_program_path* is the path to your program's executable (.exe) file, *your_program* is the name of your program's .exe file, and *your_program_name* is the friendly name that you want to show up in the exceptions list.

2. Press ENTER.

3. To add a program to the standard profile, type the following:

   **netsh firewall set allowedprogram** *your_program_path\your_program your_program_name* **enable standard**

   Where *your_program_path* is the path to your program's executable (.exe) file, *your_program* is the name of your program's .exe file, and *your_program_name* is the friendly name that you want to show up in the exceptions list.

4. Press ENTER.

**Some or all of the settings in Windows Firewall appear dimmed.**
This happens when you do not have the appropriate user rights on the computer you are trying to administer. It can also happen if your company uses Group Policy to manage Windows Firewall.

Windows Firewall dims all of the options in Windows Firewall if you are not a member of the Administrators group. To configure Windows Firewall options, have your domain administrator or accounts administrator add you to the Administrators group on the computer you want to administer.

Windows Firewall also dims any options that are configured through Group Policy. This includes options that are configured through Local Group Policy or through domain-based Group Policy. You can use the Resultant Set of Policy snap-in to determine which Windows Firewall policies are being applied to a computer as well as in which GPO the policies are stored.

**To determine which policies are being applied to the computer**

1. At the command prompt, type **rsop.msc**, and press ENTER.

2. In the console tree, open **Computer Configuration**, open **Administrative Templates**, open **Network**, and then click **Network Connections**.

3. In the console tree, click **Windows Firewall**, and double-click any policies that appear in the details pane to learn more about the policy.

4. In console tree, click **Domain Profile**, and double-click any policies that appear in the details pane to learn more about the policy.

5. In the console tree, click **Standard Profile**, and double-click any policies that appear in the details pane to learn more about the policy.

**I added a program to the exceptions list in Group Policy, but the program does not behave as though it's in the exceptions list.**
This can occur for several reasons. Attempt each of the following possible fixes in the order in which they appear.

Make sure that you used the correct profile when you added the program to the exceptions list. You can configure Windows Firewall policy settings in either a domain profile or a standard profile. Domain profile settings are applied when a computer is connected to a network that contains the domain controllers for the domain in which its computer account resides. Standard profile settings are applied when a computer is not connected to a network that contains the domain controllers for the domain in which its computer account resides, such as a public network.

**To verify Group Policy settings in the domain profile or standard profile**

1. Open the Group Policy Object Editor snap-in to edit the GPO that is used to manage Windows Firewall settings in your organization.

2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, and then open **Windows Firewall**.

3. In the console tree, click **Domain Profile** to see the Group Policy settings that are configured when the computer uses the domain profile.

4. In the console tree, click **Standard Profile** to see the Group Policy settings that are configured when the computer uses the standard profile.

Make sure that the Group Policy settings have been refreshed on the computer you are administering. Local policies take effect immediately, but domain-based Group Policies are refreshed every 90 minutes (unless the computer is a domain controller, in which case policies are refreshed every 5 minutes). If you are unsure whether any policy changes have been applied to the computer, you can force Group Policy to refresh the policies on a computer by running **gpupdate /force** at the command line.

Make sure that policies are being applied from the correct GPO. You can determine this by using the Resultant Set of Policy snap-in.

**To determine in which GPO a policy setting is stored**

1. At the command prompt, type **rsop.msc**, and press ENTER.

2. In the console tree, open **Computer Configuration**, open **Administrative Templates**, open **Network**, and then click **Network Connections**.

3. In the console tree, double-click any policies that appear in the details pane to determine in which GPO the policy is stored.

4. In the console tree, click **Windows Firewall**, and double-click any policies that appear in the details pane to determine in which GPO the policy is stored.

5. In console tree, click **Domain Profile**, and double-click any policies that appear in the details pane to determine in which GPO the policy setting is stored.

6. In the console tree, click **Standard Profile**, and double-click any policies that appear in the details pane to determine in which GPO the policy is stored.

Finally, make sure that you typed the path or program name correctly when you added it to the exceptions list. Windows Firewall does not check for errors, nor does it check whether the program exists. This allows you to add programs to the exceptions list before you install the program. Also, make sure that the path you used applies to the computer you are troubleshooting. If possible, use environment variables in the path so that the path is valid for multiple computer configurations.

**I used the My network (subnet) only scope option, but traffic originating beyond my subnet is getting through Windows Firewall.** This can happen because of the way in which Windows Firewall determines traffic that matches the **My network (subnet) only** scope option. When you use the **My network (subnet) only** scope option, Windows Firewall does not perform any IP configuration testing or subnet mask interpretation to determine whether the traffic is coming from a computer that is a neighbor on a locally attached subnet. Instead, Windows Firewall determines whether the traffic is originating from an address that is directly reachable based on routes in the IPv4 and Internet Protocol version 6 (IPv6) routing tables (which you can view with the **route print** command). If it is, this traffic is considered to match the **My network (subnet) only** scope. Because Windows Firewall defines traffic for the **My network (subnet) only** scope in this way, there are situations in which Windows Firewall allows unexpected or unanticipated traffic.

An example of this occurs when you have a multihomed server with two network interfaces: for example, one interface for the accounting department's subnet and one interface for the marketing department's subnet. If you use the **My network (subnet) only** scope option to restrict the scope of an exception on this computer, Windows Firewall will allow traffic originating from both subnets because the addresses for the address ranges of both subnets are directly reachable. To ensure that Windows Firewall allows only exceptions traffic from one subnet for IPv4 traffic, select the **Custom list** option and configure the IPv4 address range for the desired subnet. You cannot specify custom address ranges for IPv6 traffic.

Another example of this occurs when a computer is configured with a default route to the Internet for which all locations are directly reachable. This can occur when you use an ISP to directly connect a computer to the Internet through a high-speed connection. In this case, Windows Firewall sees all locations on the Internet as directly reachable when you use the **My network (subnet) only** scope option. You can restrict the scope of an exception by selecting the **Custom list** scope option and designating specific IPv4 addresses or IPv4 address ranges.

**To configure a scope with specific IPv4 addresses or IPv4 address ranges**

1. Open Windows Firewall, and then click the **Exceptions** tab.

2. If you are configuring the scope of a program or system service that is already in the exceptions list, click the program or system service, click **Edit**, and then do the following:

    a. Click the port that you want to configure, and then click **Change Scope**.

    b. Click **Custom list**, and then enter the comma-separated list of IPv4 addresses or IPv4 address ranges.

3. If you are adding a program or system service, click **Add Program**; if you are adding a port, click **Add Port**. Then do one of the following:

    a. In the **Add a Program** dialog box, click the program or system service that you want to configure, or click **Browse** and find the program or system service that you want to configure.

    b. In the **Add a Port** dialog box, type the exception name, type the port, and then select **TCP** or **UDP**.

4. Click **Change scope**.

5. In the **Change Scope** dialog box, click **Custom list**, and then enter the comma-separated list of IPv4 addresses and subnets.

| **Important** |
| --- |
| The **My network (subnet) only** or **Custom list** scope options do not necessarily increase your security. Malicious users can circumvent these scope restrictions by spoofing an IP address that appears to be directly reachable or part of the custom list. |

**I used the My network (subnet) only option to change the scope of a program exception, but some computers in my subnet are being blocked by Windows Firewall.**

This happens when a computer roams from one subnet to another. For example, if you create a shared folder on your computer, and then configure the **File and Printer Sharing** exception with the **My network (subnet) only** scope option, only directly reachable computers will be able to access your shared folder. But if one of the neighboring computers on your subnet is a portable computer, and a user moves the portable computer to a network segment outside your local subnet, the DHCP server will assign a new IP address configuration to the portable computer. When the portable computer attempts to access your shared folder, Windows Firewall determines that the computer is not directly reachable and discards the incoming connection requests.

To fix this problem, you can use the **Custom list** scope option to specify a list of IP addresses or IP address ranges from which traffic is allowed.

**To configure a scope with specific IPv4 addresses or IPv4 address ranges**

1. Open Windows Firewall, and then click the **Exceptions** tab.

2. If you are configuring the scope of a program or system service that is already in the exceptions list, click the program or system service, click **Edit**, and then do the following:

    a. Click the port that you want to configure, and then click **Change Scope**.

    b. Click **Custom list**, and then enter the comma-separated list of IPv4 addresses or IPv4 address ranges.

3. If you are adding a program or system service, click **Add Program**; if you are adding a port, click **Add Port**. Then do one of the following:

    a. In the **Add a Program** dialog box, click the program or system service that you want to configure, or click **Browse** and find the program or system service that you want to configure.

    b. In the **Add a Port** dialog box, type the exception name, type the port, and then select **TCP** or **UDP**.

4. Click **Change scope**.

5. In the **Change Scope** dialog box, click **Custom list**, and then enter the comma-separated list of IPv4 addresses and subnets.

**Important**

The **My network (subnet) only** or **Custom list** scope options do not necessarily increase your security. Malicious users can circumvent these scope restrictions by spoofing an IP address that appears to be directly reachable or part of the custom list.

**Note**

You cannot specify custom address ranges for IPv6 traffic.

# Configuring a Computer for Windows Firewall Troubleshooting

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can use advanced troubleshooting techniques to identify and fix Windows Firewall problems, you need to configure your computer for troubleshooting. In addition, you need a basic understanding of troubleshooting concepts, procedures, and tools.

**Configuration Tasks for Troubleshooting**
To configure your computer for troubleshooting, perform the following tasks:

Enable Windows Firewall logging

Enable Windows Firewall auditing

Enable network tracing

**Enable Windows Firewall logging**
The Windows Firewall log file contains information about the network traffic processed by Windows Firewall. To effectively troubleshoot problems with Windows Firewall, you must enable logging for both dropped packets and successful connections.

### To enable Windows Firewall logging

1. Open Windows Firewall in Control Panel.
2. Click the **Advanced** tab.
3. Under **Security Logging**, click **Settings**.
4. In the **Log Settings** dialog box, select the **Log dropped packets** and **Log successful connections** check boxes.

**Enable Windows Firewall auditing**
You can configure the Windows Server 2003 audit log so it records information about Windows Firewall. Specifically, the audit log can track changes that are made to Windows Firewall settings and it can track information about which programs and system services attempt to listen on a port. To record Windows Firewall information in the audit log, you must enable both the **Audit policy change** and **Audit process tracking** policies.

### To enable Windows Firewall auditing

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. In the console tree, open **Windows Settings**, open **Local Policies**, and then click **Audit Policy**.
3. In the details pane, double-click **Audit policy change**, select the **Success** and **Failure** check boxes, and then click **OK**.
4. In the details pane, double-click **Audit process tracking**, select the **Success** and **Failure** check boxes, and then click **OK**.

**Enable network tracing**
The IPNathlp trace log stored in the systemroot\tracing folder provides useful information about the network traffic passing to and from Windows Firewall. Network tracing is most commonly used to solve difficult-to-diagnose network access problems and problems with programs.

### To enable network tracing for Windows Firewall

- At the command line, type **netsh ras set tracing ipnathlp enable**, and press ENTER.

| **Important** |
| --- |
| Changing log file settings for Windows Firewall logging, auditing, and network tracing can substantially increase the size of your log files and use up disk space. After you are finished troubleshooting Windows Firewall issues, it is recommended that you configure log file settings, auditing settings, and network tracing settings as they were before you performed your troubleshooting tasks. |

# Troubleshooting Windows Firewall Problems

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Firewall typically causes three types of problems.

**In this section**


Fixing Network Access Problems

Fixing Program Problems

Fixing Configuration and Management Problems

# Fixing Network Access Problems

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides a description of some common network access problems that you might experience when Windows Firewall is enabled.

Typically, you encounter network access problems when you attempt to perform a task that uses the network or when you try to connect to a resource across a network, such as a printer, computer, or a shared folder or file. From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- File and Printer Sharing Does Not Work
- Computer Names Do Not Resolve When Used in a UNC Path
- VPN and Remote Access Connections Fail
- Remote Administration Tools Fail
- The Ping Command Times Out
- A Computer Receives Network Traffic Only from Its Local Subnet or Specific Computers

# File and Printer Sharing Does Not Work

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem typically occurs when someone attempts to access a resource on your computer, such as a shared folder or printer, and the person attempting to access the resource receives a "network path was not found" error.

**Cause**

File and printer sharing can fail when you turn on Windows Firewall because file and printer sharing requires your computer to respond to unsolicited incoming traffic on one or more of the following ports: TCP ports 139 and 445, and UDP ports 137 and 138. Windows Firewall blocks incoming traffic through these ports. To fix this problem, you need to unblock one or more TCP and UDP ports.

**Solution**

Make sure that you have enabled the File and Printer Sharing exception in the Windows Firewall exceptions list.

**To enable the File and Printer Sharing exception**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Select the **File and Printer Sharing** check box.

If file and printer sharing still does not work properly, verify that you are configuring **File and Printer Sharing** in the same profile that your computer is using. If you enable **File and Printer Sharing** in the domain profile, but your computer is using the standard profile, then **File and Printer Sharing** might not be enabled. You can use the Netsh tool to determine which profile your computer is currently using and which profile has **File and Printer Sharing** enabled.

**To determine which profile your computer is currently using**

- At the command prompt, type **netsh firewall show currentprofile**, and press ENTER.

**To determine which profile has File and Printer Sharing enabled**

- At the command prompt, type **netsh firewall show service**, and press ENTER.

If **File and Printer Sharing** is not enabled in the profile that your computer is using, you need to enable **File and Printer Sharing** in the appropriate profile.

**To enable File and Printer Sharing in a specific profile**

1. At the command prompt, type the following:

   **netsh firewall set service type = fileandprint mode = enable profile =** *your profile*

   Where *your profile* can be current, domain, standard, or all, depending on which profile must have **File and Printer Sharing** enabled.

2. Press ENTER.

If file and printer sharing still does not work properly, verify that the scope setting for the **File and Printer Sharing** exception is not preventing other computers from accessing your shared files and printers. The scope setting allows you to restrict an exception to a specific group of computers. Be sure that the scope setting for the **File and Printer Sharing** exception includes the computers that are trying to access the files and folders on your computer. There are three possible scope settings:

**Any computer (including those on the Internet)**, which does not restrict the exception.

**My Network (subnet) only**, which restricts the exception to computers that are directly reachable based on entries in the IP routing table.

**Custom list**, which restricts the exception to a list of Internet Protocol version 4 (IPv4) addresses or IPv4 address ranges. You cannot use the **Custom list** option for Internet Protocol version 6 (IPv6) traffic.

**To verify the scope of the File and Printer Sharing exception**

1. Open Windows Firewall, and then click the **Exceptions** tab.

2. Click **File and Printer Sharing**, and then click **Edit**.
3. Click **TCP 139**, and then click **Change Scope**.
4. Verify that the scope setting includes the computers that you want to access your computer. If it does not, change the scope so it is less restrictive.
5. Repeat steps 3 and 4 for **TCP 445**, **UDP 137**, and **UDP 138**.

| **Important** |
| --- |
| The **My network (subnet) only** or **Custom list** scope options do not necessarily increase your security. Malicious users can circumvent these scope restrictions by spoofing an IP address that appears to be directly reachable or is part of the custom list. |

If file and printer sharing still does not work properly and you are using a collaboration tool or a file-sharing tool, you might have to configure other exceptions to share files. If your vendor does not provide information about the settings that need to be configured for your collaboration or file sharing program, try adding the program to the exceptions list.

**To add a program or service to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. On the **Exceptions** tab, click **Add Program**, and then follow the instructions that appear on your screen.

# Computer Names Do Not Resolve When Used in a UNC Path

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem typically occurs when you use **Run** to access a shared folder and you specify a path in the form *server name**shared folder name*. If the UNC path does not resolve, you will see a "...network resource not found" error. This can also happen when you use the **net use** command to map a network drive.

**Cause**

Usually, UNC path names cannot resolve because Windows Firewall is preventing access to a naming service, such as WINS or DNS. To solve this problem, you need to add a port or a service to the Windows Firewall exceptions list on the server that is running the naming service.

This can also be caused if you're using NetBIOS and you have enabled the **Prohibit unicast response to multicast or broadcast requests** policy setting. If this policy setting is enabled, and a computer broadcasts a NetBIOS name query request to resolve a name, the unicast response from the target computer will be dropped by Windows Firewall. To fix this, you need to disable the **Windows Firewall: Prohibit unicast response to multicast or broadcast requests**policy setting.

**Solution**

If your organization uses WINS for name resolution, add UDP port 137 to the Windows Firewall exceptions list on your WINS server.

**To add UDP port 137 to the Windows Firewall exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. On the **Exceptions** tab, click **Add Port**.
3. In the **Add a Port** dialog box, in **Name**, type a friendly name for the port exception (for example, UDP port 137).
4. In **Port number**, type 137, and then click the **UDP** option button.

If your organization uses DNS for name resolution, add UDP and TCP port 53 to the Windows Firewall exceptions list on your DNS server.

**To add UDP port 53 and TCP port 53 to the Windows Firewall exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. On the **Exceptions** tab, click **Add Port**.
3. In the **Add a Port** dialog box, in **Name**, type a friendly name for the port exception (for example, TCP port 53).
4. In **Port number**, type 53.
5. Click the **TCP** option button, and then click **OK**.
6. On the **Exceptions** tab, click **Add Port**.
7. In the **Add a Port** dialog box, in **Name**, type a friendly name for the port exception (for example, UDP port 53).
8. In **Port number**, type 53.
9. Click the **UDP** option button, and then click **OK**.

If you use NetBIOS and UNC path names cannot resolve after you open the appropriate ports on your WINS server or DNS server, then try disabling the **Windows Firewall: Prohibit unicast response to multicast or broadcast requests**policy setting.

**Disable the Prohibit unicast response to multicast or broadcast requests policy setting**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open

**Windows Firewall**, and then click either **Domain Profile** or **Standard Profile**.

3. Double-click **Windows Firewall: Prohibit unicast response to multicast or broadcast requests**, and on the **Settings** tab, click **Disabled** or **Not Configured**.

# VPN and Remote Access Connections Fail

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, you encounter this problem when you try to connect to a corporate or private network using a virtual private network (VPN) or remote access connection and you are unable to connect or you receive an error message telling you that the connection failed.

**Cause**

This usually occurs when you configure VPN or remote access settings and exceptions in the Windows Firewall domain profile, but Windows Firewall uses the settings and exceptions that are in the standard profile. If you have not configured VPN or remote access settings and exceptions in the standard profile, the VPN and remote access connections will fail.

This happens because VPN and remote access connections are special connection types. A computer that connects through a VPN or remote access connection does not undergo any checks for domain membership when Windows Firewall determines the network on which a computer resides. For example, if your computer is connected to a public network, such as the Internet, and you attempt to create a VPN or remote access connection directly to a private network that contains the domain controllers of the domain in which your computer account resides, the VPN or remote access connection is ignored and Windows Firewall will determine that your computer is currently connected to a public network. Because your computer is connected to a public network, Windows Firewall uses the settings and exceptions in the standard profile instead of the settings in the domain profile.

**Solution**

To solve this problem, you need to configure VPN and remote access settings and exceptions in the standard profile. The following table lists some of the common ports that you might need to add to the exceptions list for your VPN configuration. It is not an exhaustive list. You might not need to add all of these ports to the exceptions list for your VPN configuration. For example, if you are allowing only Layer Two Tunneling Protocol (L2TP), you would not configure a filter for Point-to-Point Tunneling Protocol (PPTP).

| Service or Protocol | Port |
|---|---|
| PPTP | TCP 1723, UDP 1723 |
| L2TP | TCP 1701, UDP 1701 |
| Hypertext Transport Protocol (HTTP) | TCP 80 |
| Secure Hypertext Transport Protocol (HTTPS) | TCP 443 |
| Simple Mail Transfer Protocol (SMTP) | TCP 25, UDP 25 |
| Dynamic Host Control Protocol (DHCP). Required if the remote access server uses an external DHCP server. | TCP 67, UDP 67 |
| Post Office Protocol version 3 (POP3) | TCP 110 |
| Domain Name Service (DNS) | UDP 53 |
| IPsec | UDP 500 |
| IPsec with network address translation (NAT) | UDP 4500 |

# Remote Administration Tools Fail

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, you encounter this problem when you try to connect to a remote server with an administration tool, such as a Microsoft Management Console (MMC) snap-in.

**Cause**

Remote administration tools often fail because the server you are trying to manage blocks the unsolicited incoming traffic from the remote administration tool. To remotely manage a server with a remote administration tool, you must configure Windows Firewall on the remote server to allow unsolicited traffic from the remote administration tool.

**Solution**

You can solve this problem by configuring Windows Firewall settings on the remote server you are trying to administer. Usually, you need to add a port, program, or service to the Windows Firewall exceptions list on the server you are trying to manage. To determine which Windows Firewall settings you need to configure on the remote server, see "Windows Firewall Settings" in the Windows Server 2003 Tools and Settings documentation. Windows Firewall Settings provides information about the specific Windows Firewall settings you need to configure to use every remote administration tool that is available on Windows Server 2003.

If you cannot find your remote administration tool in Windows Firewall Settings, or the settings do not fix your problem, try enabling the **File and Printer Sharing** exception on the remote server.

**To enable the File and Printer Sharing exception**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Select the **File and Printer Sharing** check box, and click **Edit**.
3. In the **Edit a Service** dialog box, click **Change Scope**.
4. In the **Change Scope** dialog box, do one of the following:

   Click **Any computer (including those on the Internet)** if you want this exception to apply to all computers.

   Click **Custom list** if you want this exception to apply to specific Internet Protocol version 4 (IPv4) addresses or IPv4 address ranges.

If you still cannot use your remote administration tool to administer a remote server, try enabling the **Windows Firewall: Allow remote administration exception** policy setting on the remote server.

**To enable the Allow remote administration exception policy setting**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.
2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then click either **Domain Profile** or **Standard Profile**.
3. In the details pane, double-click **Windows Firewall: Allow remote administration exception**, and then click **Enabled**.

| **Important** |
| --- |
| You should enable the **Windows Firewall: Allow remote administration exception** policy setting only if your remote administrative tools require RPC or DCOM. Malicious users often attempt to attack networks and computers using RPC and DCOM. It is recommended that you contact the manufacturer of your remote administration tool to determine if it requires RPC and DCOM communication. If it does not, do not enable the **Windows Firewall: Allow remote administration exception** policy setting. |

# The Ping Command Times Out

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, you see this problem occur when you use the **ping** command to test network connectivity or to see if another computer is on the network.

**Cause**

The **ping** command times out because the computer you are trying to reach is blocking incoming Internet Control Message Protocol (ICMP) echo traffic.

**Solution**

To fix this problem, you need to configure Windows Firewall ICMP settings on the computer you are trying to reach so that incoming ICMP echo traffic is allowed.

**To allow incoming echo request traffic**

1. Open Windows Firewall, and then click the **Advanced** tab.
2. In **ICMP**, click **Settings**.
3. In the **ICMP Settings** dialog box, select the **Allow incoming echo request** check box.

# A Computer Receives Network Traffic Only from Its Local Subnet or Specific Computers

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, you see this problem occur when a computer beyond your local subnet tries to access a shared folder, printer, or file on your computer and the connection fails. You might also see this problem if you've added a port or program to the exceptions list and the port or program blocks traffic from beyond your local subnet, but does not block traffic from within your local subnet.

**Cause**

This usually occurs because the scope of one or more exceptions on your computer has been restricted to your local subnet or to a specific group of computers.

By default, exceptions are not restricted in scope (that is, an exception applies to any computer, including computers on the Internet). However, you can configure the scope of an exception to your local subnet or to specific Internet Protocol version 4 (IPv4) addresses.

**Solution**

To fix this problem, first determine whether any of the exceptions in your exceptions list have restricted scope. Next, if there any scope restrictions, you need to expand the scope restriction.

**To determine whether an exception has restricted scope**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. On the **Exceptions** tab, click a program or port, and click **Edit**.
3. Click **Change scope**.

   If an exception has restricted scope, **My network (subnet) only** or **Custom list** will be selected.

You can customize a scope restriction by using the **Custom list** option. The **Custom list** option allows you to enter specific IPv4 addresses or a range of IPv4 addresses. You cannot use the **Custom list** option for Internet Protocol version 6 (IPv6) traffic.

**To configure a scope with specific IP addresses or IP address ranges**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. If you are configuring the scope of a program or system service that is already in the exceptions list, click the program or system service, click **Edit**, and then do the following:
   a. Click the port that you want to configure, and then click **Change Scope**.
   b. Click **Custom list**, and then enter the comma-separated list of IP addresses or IPv4 address ranges.
3. If you are adding a program or system service, click **Add Program**; if you are adding a port, click **Add Port**. Then do one of the following:
   a. In the **Add a Program** dialog box, click the program or system service that you want to configure, or click **Browse** and find the program or system service that you want to configure.
   b. In the **Add a Port** dialog box, type the exception name, type the port, and then select **TCP** or **UDP**.
4. Click **Change scope**.
5. In the **Change Scope** dialog box, click **Custom list**, and then enter the comma-separated list of IPv4 addresses and IPv4 address ranges.

**Important**

The **My network (subnet) only** or **Custom list** scope options do not necessarily increase your security. Malicious users can circumvent these scope restrictions by spoofing an IP address that appears to be directly reachable or part of the custom list.

# Fixing Program Problems

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides a description of some common problems you might experience when you run a program with Windows Firewall enabled.

Typically, you encounter program problems when you try to run a program with Windows Firewall turned on or you turn Windows Firewall on while a program is running. From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- A Program in the Exceptions List Does Not Run Properly
- A Notification Says Windows Explorer Is Trying to Listen on a Port
- A User Cannot Determine What to Add to the Exceptions List
- A Program Removed from the Exceptions List Keeps Working
- There Is a Problem Running an Antivirus Program
- A System Service Runs in Svchost.exe and Cannot Be Added to the Exceptions List
- An Unexpected Program Appears in the Exceptions List

# A Program in the Exceptions List Does Not Run Properly

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, this problem occurs when you add a program's executable (.exe) file to the exceptions list and then try to run the program. Either the program will start, but then fail to run properly, or the program will not start at all.

**Cause**

Some programs depend on more than one .exe file. If this is the case, and the dependent .exe file is not in the exceptions list, the program will not run properly.

**Solution**

If you have already restarted your computer, but the problem persists, make sure that you added the correct .exe file to the exceptions list. If a program relies on more than one .exe file, you need to identify the .exe files that are associated with the program and add those .exe files to the exceptions list. You can use the **tasklist** command to determine whether a program relies on other .exe files.

**To determine whether a program relies on more than one .exe file**

1. Start the program that you want to evaluate.
2. At the command line, type **tasklist /m**, and press ENTER.
3. If a program relies on another .exe file, you will see it in the list of modules that are associated with the program.

If a program relies on other .exe files, add those files to the exceptions list.

**To add a program to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. On the **Exceptions** tab, click **Add Program**, and then follow the instructions that appear on your screen.

If the program still does not work properly, you might have to add ports to the exceptions list.

First, you need to identify the ports used by the program.

**To determine the ports used by a program**

1. Start the program that you want to evaluate.
2. At the command line, type **tasklist**, and then press ENTER.
3. Look up the process ID (PID) that is associated with the program you are evaluating. If the program relies on more than one .exe file, be sure to look up the PID for each .exe file that the program is using.
4. At the command line, type **netstat -a -o -n**, and then press ENTER.
5. Use the program's PID to determine the ports on which the program is listening.

Next, you need to determine which ports are enabled in the exceptions list.

**To identify ports that are enabled in the exceptions list**

- At the command line, type **netsh firewall show state**, and then press ENTER.

  Ports that are designated as open are enabled in the Windows Firewall exceptions list.

If a program uses a port, and the port is not enabled in the exceptions list, try adding the port to the exceptions list.

**To add a port to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Click **Add Port**.

3. In **Name**, type a friendly name for the port exception.
4. In **Port number**, type the port number used by the program.
5. Click either **TCP** or **UDP** to specify the type of port that corresponds to your port number.
6. Repeat steps 2 through 5 if your program uses multiple ports and the ports are not enabled in the exceptions list.

| **Important** |
| --- |
| Adding a port to the exceptions list can lessen the security of your computer because the port will be open any time the computer is running. You should add ports to the exceptions list only when it is not possible to add a program to the exceptions list. |

# A Notification Says Windows Explorer Is Trying to Listen on a Port

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, this problem occurs when you open Windows Explorer and see a **Windows Security Alert** dialog box (referred to throughout this documentation as a notification) that displays a message informing you that Explorer.exe is attempting to listen on a port. If you are a member of the Administrators group, the notification will ask if you want to unblock Explorer.exe. If you are not a member of the Administrators group, the notification will tell you to contact your Administrator if you want to unblock Explorer.exe.

**Cause**

When you see this notification, it usually means that Windows Explorer (Explorer.exe) is hosting a secondary program and the program it is hosting is trying to listen for incoming traffic on one or more ports. Because Explorer.exe is hosting the program, the program is running as a separate process within Explorer.exe. For this reason, Windows Firewall cannot determine the name of the program's executable (.exe) file. Instead, it can determine only that a process within Explorer.exe is attempting to listen for incoming traffic.

**Solution**

To fix this problem, trying using the **tasklist** command to identify the program that is attempting to listen for incoming traffic. If you can identify the program, try to add it to the exceptions list.

| **Important** |
| --- |
| Do not add Explorer.exe to the exceptions list. Doing so allows any program that is hosted by Explorer.exe to receive unsolicited incoming traffic, which can decrease the security of your computer. |

**To identify a program that is running within Windows Explorer**

1. Start Windows Explorer and wait for the Windows Firewall notification telling you that Explorer.exe is attempting to listen on a port.
2. At the command line, type **tasklist /m**, and press ENTER.
3. See if there is an .exe file listed with the modules that are running in Explorer.exe.

If there is an .exe file associated with Explorer.exe, try adding it to the exceptions list, and then shut down and restart Windows Explorer to see if you still get the Windows Firewall notification. If there is more than one .exe file, add the .exe files to the exceptions list one at a time and test Windows Explorer each time to determine which .exe files you need to add to the exceptions list. You might not need to add them all.

**To add a program to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. On the **Exceptions** tab, click **Add Program**, and then follow the instructions that appear on your screen.

If there are no .exe files running in Explorer.exe, you can try to determine which ports the program is using and then add those ports to the exceptions list.

**To determine the ports that are being used by Windows Explorer**

1. Start Windows Explorer and wait for the Windows Firewall notification telling you that Explorer.exe is attempting to listen on a port.
2. At the command line, type **tasklist /m**, and press ENTER.
3. Look up the process ID (PID) that is associated with Explorer.exe.
4. At the command line, type **netstat -a -o -n**, and then press ENTER.
5. Use the program’s PID to determine the ports on which the program is listening.

After you determine the ports that Explorer.exe is using, try adding the ports to the exceptions list.

**To add a port to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Click **Add Port**.
3. In **Name**, type a friendly name for the port exception.
4. In **Port number**, type the port number used by the program.
5. Click either **TCP** or **UDP** to specify the type of port that corresponds to your port number.
6. Repeat steps 2 through 5 if your program uses multiple ports and the ports are not enabled in the exceptions list.

| **Important** |
| --- |
| Adding a port to the exceptions list can lessen the security of your computer because the port will be open any time the computer is running. You should add ports to the exceptions list only when it is not possible to add a program to the exceptions list. |

# A User Cannot Determine What to Add to the Exceptions List

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, you have this problem when you try to add a program to the exceptions list, but you cannot find the program's executable (.exe) file or the program relies on secondary .exe files and dynamic-link library (.dll) files and you do not know which files to add to the exceptions list.

**Cause**

Some programs run as a system service and might not have an .exe file that you can add to the exceptions list. In addition, some programs rely on secondary .exe files or .dll files to process incoming traffic. In this case, adding the program's primary .exe file to the exceptions list will not open the port that the program needs because the program uses a secondary .exe file or .dll to process incoming traffic.

**Solution**

To fix this problem, ask the program vendor or read the program documentation to see which .exe file the program uses to listen for unsolicited traffic, and then add the .exe file to the exceptions list.

**To add a program to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. On the **Exceptions** tab, click **Add Program**, and then follow the instructions that appear on your screen.

If you cannot determine which .exe file to add to the exceptions list, you might be able to use Event Viewer to determine which .exe file the program uses to listen for unsolicited traffic.

| **Note** |
| --- |
| You must configure audit policies to see Windows Firewall events. For more information, see Configuring a Computer for Windows Firewall Troubleshooting. |

**To view Windows Firewall notification events in Event Viewer**

1. Open Event Viewer.
2. Click **Security**, and look for the most recent **Failure Audit** event type that has an **Event** identifier of 861.
3. Double-click the event.
4. The path and the name of the .exe file should appear next to **Path**.
5. If there is no information next to **Path**, write down the number that appears next to **Process Identifier**.

If you know the process identifier (PID) for the .exe file that attempted to listen for unsolicited traffic, you can use the **tasklist** command to determine the name of the .exe file.

**To determine which .exe file is associated with a process identifier**

1. At the command line type **tasklist**, and then press ENTER.
2. Use the PID to identify the .exe file.

If you still cannot determine which .exe file is listening for unsolicited traffic, try using the **tasklist** and **netstat** commands to determine which ports the program uses, and then add those ports to the exceptions list.

**To determine the ports used by a program**

1. Start the program that you want to evaluate.
2. At the command line, type **tasklist**, and then press ENTER.
3. Look up the PID that is associated with the program you are evaluating. If the program relies on more than one .exe file, be sure to look up the PID for each .exe file.

4. At the command line, type **netstat -a -o -n**, and then press ENTER.
5. Use the program's PID(s) to determine the ports on which the program is listening.

**To add a port to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Click **Add Port**.
3. In **Name**, type a friendly name for the port exception.
4. In **Port number**, type the port number used by the program.
5. Click either **TCP** or **UDP** to specify the type of port that corresponds to your port number.
6. Repeat steps 2 through 5 if your program uses multiple ports and the ports are not enabled in the exceptions list.

| **Important** |
| --- |
| Adding a port to the exceptions list can lessen the security of your computer because the port will be open any time the computer is running. You should add ports to the exceptions list only when it is not possible to add a program to the exceptions list. |

# A Program Removed from the Exceptions List Keeps Working

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, you see this behavior when you remove a program from the exceptions list and then run the program. Although the program is not on the exceptions list, it runs as though it is still listed in the exceptions list.

**Cause**

This can occur for several reasons.

- The program might not listen for unsolicited traffic all of the time. Some programs listen for unsolicited incoming traffic only when you enable certain options or use certain program features.
- There might be another exception on the exceptions list that opens a port used by the program. In this case, the program will still run normally even though it is not listed in the exceptions list.
- The program might use the Windows Firewall application programming interface (API) to add a required port to the exceptions list without notifying you.

**Solution**

You cannot prevent a program from using the Windows Firewall API to add a port to the exceptions list. If you need to prevent this, contact the program vendor or read the program documentation to see if there is a way to disable the feature that listens for incoming traffic. This might prevent the program from using the Windows Firewall APIs.

If the program consists of several components or features, and only one of those components or features is responsible for handling unsolicited incoming traffic, you can try to disable the feature or the component. The program vendor might be able to provide with a registry modification to do this.

If the program is using a port that is already listed in the exceptions list, contact the program vendor or read the program documentation to determine which ports the program uses. If you cannot determine this, run the program and then use the **tasklist** and **netstat** commands to determine which ports the program is using. To do this, you might need to contact the program vendor or read the program documentation to determine which program features you need to use or which program options you need to enable to force the program to listen on a port.

**To determine the ports used by a program**

1. Start the program that you want to evaluate.
2. At the command line, type **tasklist**, and then press ENTER.
3. Look up the process ID (PID) that is associated with the program you are evaluating. If the program relies on more than one .exe file, be sure to look up the PID for each .exe file that the program is using.
4. At the command line, type **netstat -a -o -n**, and then press ENTER.
5. Use the program's PID to determine the ports on which the program is listening.

After you identify the port that the program uses, see if the port is already listed on the exceptions list.

**To identify ports that are enabled in the exceptions list**

- At the command line, type **netsh firewall show state**, and then press ENTER.

  Ports that are designated as open are enabled in the Windows Firewall exceptions list.

Finally, if the port is in the exceptions list, you need to determine whether you can remove the port without causing other programs or system services to stop running. Removing a port from the exceptions list might prevent other programs and system services from running properly.

# There Is a Problem Running an Antivirus Program

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, you see this problem occur when an antivirus program tries to automatically download a file--such as a virus definition file, a signature file, or a program update--either from the Internet or from a server on your corporate network. When this happens, you might see a **Windows Security Alert** dialog box (notification) that displays a message informing you that the application is trying to use a port that is blocked, or the application might stop responding or it might behave unpredictably.

**Cause**

If the download is initiated by a central server and "pushed" down to your computer, then the download traffic will be considered unsolicited incoming traffic and Windows Firewall will block the traffic. Many antivirus programs use the file transfer protocol (FTP) to transmit program updates, virus signatures, and virus definitions to a computer, and because Windows Firewall blocks the FTP ports by default, the program cannot process the download request. Although many antivirus programs have been updated and use the Windows Firewall application programming interface (API) to open the FTP ports, not all programs do this. If a program does not automatically open the proper ports to allow an unsolicited download, the program might fail when Windows Firewall is enabled.

**Solution**

To solve this problem, contact the program vendor or read the program documentation to see if you can configure the antivirus program so files cannot be "pushed" down to your computer. In some cases, you can configure an antivirus program so it "pulls" files down at specified intervals or when you manually initiate an update or download.

If this is not acceptable behavior in your organization, contact the program vendor or read the program documentation to see if the antivirus program uses FTP to download program updates, virus signature files, or virus definition files. If the program use FTP, make sure that the Application Layer Gateway (ALG) Service is started. Windows Firewall uses the ALG Service to process active FTP traffic.

**To start the Application Gateway Service**

1. Open the Services snap-in.
2. Double-click **Application Layer Gateway Service**.
3. On the **General** tab, if **Service status** is not **Started**, click **Start**.

If the ALG Service is running and you still cannot get your antivirus program to work with Windows Firewall, try adding the antivirus program to the exceptions list.

**To add a program to the Windows Firewall exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Click **Add Program**, and then follow the directions that appear on your screen.

If this does not work, you need to determine which port(s) the antivirus program uses, and then add the port(s) to the exceptions list.

**To determine the ports used by a program**

1. Start the program that you want to evaluate.
2. At the command line, type **tasklist**, and then press ENTER.
3. Look up the process ID (PID) that is associated with the program you are evaluating. If the program relies on more than one .exe file, be sure to look up the PID for each .exe file that the program is using.
4. At the command line, type **netstat -a -o -n**, and then press ENTER.
5. Use the program's PID to determine the ports on which the program is listening.

After you identify the port that the antivirus program uses, try adding the port to the exceptions list.

**To add a port to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Click **Add Port**.
3. In **Name**, type a friendly name for the port exception.
4. In **Port number**, type the port number used by the program.
5. Click either **TCP** or **UDP** to specify the type of port that corresponds to your port number.
6. Repeat steps 2 through 5 if your program uses multiple ports and the ports are not enabled in the exceptions list.

| **Important** |
| --- |
| Adding a port to the exceptions list can reduce the security of your computer because the port will be open any time the computer is running. You should add ports to the exceptions list only when it is not possible to add a program to the exceptions list. |

# A System Service Runs in Svchost.exe and Cannot Be Added to the Exceptions List

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Typically, you see this behavior when you need to add a service to the exceptions list, but the service runs within Svchost.exe, not within an executable (.exe) file that you can add to the exceptions list.

**Cause**

Most system services run within Svchost.exe. It is not recommended that you add Svchost.exe to the exceptions list. Doing so will allow any service running within Svchost.exe to receive incoming traffic through Windows Firewall.

**Solution**

If the service uses one or more fixed ports, you might be able to add the ports to the exceptions list. You can use the **netstat** and **tasklist** commands to determine whether a service uses fixed ports. First, run **tasklist** to identify the process identifier (PID) that is associated with the Scvhost.exe process in which the service is running. Then, run **netstat** to determine whether there are any fixed ports being established by the Svchost.exe process in which the service is running. You need to restart your computer several times and perform these procedures each time you restart your computer to verify that the same ports are being assigned each time.

**To determine the process identifier for a service**

1. At the command line, type **tasklist /svc**, and press ENTER.
2. Look up the PID that is associated with the Svchost.exe in which the service is running.

**To determine which ports a process is using**

1. At the command line, type **netstat -a -o -n**, and press ENTER.
2. Look up the port numbers that are assigned to the PID you are evaluating.

If, after restarting your computer several times, you can determine that the same ports are being assigned to Svchost.exe, try adding those ports to the exceptions list.

**To add a port to the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Click **Add Port**.
3. In **Name**, type a friendly name for the port exception.
4. In **Port number**, type the port number used by the program.
5. Click either **TCP** or **UDP** to specify the type of port that corresponds to your port number.

| **Important** |
| --- |
| Adding a port to the exceptions list can lessen the security of your computer because the port will be open any time the computer is running. You should add ports to the exceptions list only when it is not possible to add a program to the exceptions list. |

If the service does not use static ports, but instead uses the RPC endpoint mapper (EPM) to allocate ports dynamically, you might be able to define the range of ports that the RPC EPM allocates and then add the defined ports to the exceptions list. For more information about defining a range ports that RPC assigns, see How to Configure RPC Dynamic Port Allocation to Work with Firewalls in the Microsoft Knowledge Base.

| **Important** |
| --- |
| Adding numerous ports to the exceptions list lessens the security of your computer and can increase the risk of attack. You should open numerous ports only as a last resort. |

# An Unexpected Program Appears in the Exceptions List

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can see this problem occur when you open Windows Firewall and click the **Exceptions** tab.

**Cause**
A program can use the Windows Firewall application programming interface (API) to automatically add itself and other programs, ports, and services to the Windows Firewall exceptions list without notifying you.

Group Policy might be applying policy settings that you are unaware of to your computer.

**Solution**
You cannot prevent a program from using the Windows Firewall API to add itself or other programs, ports, or system services to the exceptions list. You can, however, determine whether programs are being added to the exceptions list through Group Policy. To do this, you must run the Resultant Set of Policy snap-in.

**To determine which policies are being applied to a computer**

1. At the command prompt, type **rsop.msc**, and press ENTER.
2. In the console tree, open **Computer Configuration**, open **Administrative Templates**, open **Network**, and then click **Network Connections**.
3. In the console tree, click **Windows Firewall**, and double-click any policies that appear in the details pane to learn more about the policy.
4. In console tree, click **Domain Profile**, and double-click any policies that appear in the details pane to learn more about the policy.
5. In the console tree, click **Standard Profile**, and double-click any policies that appear in the details pane to learn more about the policy.

# Fixing Configuration and Management Problems

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section provides a description of some common problems that you might experience when you configure or manage Windows Firewall.

Typically, you encounter configuration and management problems when you are configuring Windows Firewall settings, either in the Windows Firewall user interface (UI) or through Group Policy. From the following list, choose the problem that best describes your situation, and then step through the suggested fix:

- A Program Named “Unspecified” Appears in the Exceptions List
- A Windows Firewall Setting Does Not Take Effect
- One or More Windows Firewall Settings Appear Dimmed

# A Program Named “Unspecified” Appears in the Exceptions List

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem most commonly occurs when you use the **netsh** command to add a program to the exceptions list.

**Cause**
When you use the **netsh** command to add programs to the exceptions list, you have the option of specifying the display name for the program. If you do not specify a display name, Windows Firewall uses “Unspecified” as the display name.

**Solution**
You can fix this problem by using the **netsh** command to add the application to the exceptions list with the proper display name. To do this, you need to know the name and path to the executable (.exe) file that is associated with the “unspecified” program.

**To identify the path and file name of a program in the exceptions list**

1. Open Windows Firewall, and then click the **Exceptions** tab.
2. Click **Unspecified**, and then click **Edit**.
3. Look in **Path** to see the path and file name of the .exe file.

After you identify the path and file name, you can change the display name of the program.

**To change the display name of a program in the exceptions list**

1. At the command prompt, type the following:

   **netsh firewall set allowedprogram program = *program path and file name* name = *display name***

   Where *program path and file name* is the path and file name of the executable file and *display name* is the display name that you want to use for the program.

2. Press ENTER.

# A Windows Firewall Setting Does Not Take Effect

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem most commonly occurs when you add a program or port to the exceptions list, but the program or port does not behave as though it's in the exceptions list. This problem can also occur when you configure other Windows Firewall settings, such as log file settings and notification settings, and the settings do not seem to take effect.

**Cause**

Usually this problem occurs because the exceptions and settings were configured in one profile and the computer is using another profile. Windows Firewall settings can be configured in two profiles: the domain profile, which is used when a computer is connected to a network in which the computer's domain account resides; and the standard profile, which is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network.

This problem can also occur if you add a program to the exceptions list, run the program, and then start Windows Firewall. In this case, Windows Firewall cannot track the ports that the program is using because they were assigned before Windows Firewall was started.

This problem can also occur when there is a Group Policy hierarchy problem. For example, if you configure Windows Firewall settings in a Group Policy Object (GPO) that inherits policy settings from another GPO, then your Windows Firewall settings might never get applied.

**Solution**

To see if you configured the Windows Firewall setting in the wrong profile, you must first determine which profile your computer is currently using.

**To determine which profile your computer is currently using**

- At the command prompt, type **netsh firewall show currentprofile**, and press ENTER.

Next, you need to verify your Windows Firewall settings in both the domain and standard profiles.

**To verify Windows Firewall settings in the domain and standard profiles**

- At the command line type **netsh firewall show config**, and then press ENTER.

  Windows Firewall configuration settings are displayed for both the standard and domain profiles.

If you configured a setting in a profile different from the one your computer is currently using, try configuring the setting in the appropriate profile.

If you still have a problem with the setting, try restarting your computer. Be sure to start Windows Firewall before you start any applications that are in the Windows Firewall exceptions list.

If you still have a problem with the setting, and your organization uses Group Policy to manage Windows Firewall settings, use the Resultant Set of Policy snap-in to determine how policy settings are being applied to the computer.

**To determine which policies are being applied to a computer**

1. At the command prompt, type **rsop.msc**, and press ENTER.
2. In the console tree, open **Computer Configuration**, open **Administrative Templates**, open **Network**, and then click **Network Connections**.
3. In the console tree, click **Windows Firewall**, and double-click any policies that appear in the details pane to learn more about the policy.
4. In console tree, click **Domain Profile**, and double-click any policies that appear in the details pane to learn more about the policy.
5. In the console tree, click **Standard Profile**, and double-click any policies that appear in the details pane to learn more about the policy.

# One or More Windows Firewall Settings Appear Dimmed

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This problem occurs when you open Windows Firewall in Control Panel and you cannot configure settings because they appear dimmed.

**Cause**

Windows Firewall settings appear dimmed when you do not have the appropriate administrative rights to configure Windows Firewall or Windows Firewall settings are being managed by Group Policy. You must be a member of the Administrators group to configure Windows Firewall settings. Domain-based policy settings override local Windows Firewall settings.

**Solution**

Make sure you are a member of the Administrators group on the computer you are trying to administer.

**To verify that you are a member of the Administrators group**

1. Open the Computer Management snap-in.
2. In the console tree, double-click **Local Users and Groups**, and then click **Groups**.
3. In the details pane, double-click **Administrators** and verify that your account name or a group to which your account is a member appears in the **Members** list.

You can also tell if you have the appropriate administrative rights to configure Windows Firewall by opening Windows Firewall in Control Panel. If you do not have administrative rights, all of the controls in the Windows Firewall user interface will appear dimmed and a warning will appear in the user interface (UI) telling you that you need to be a computer administrator to change any settings.

If you are a member of the Administrators group and some controls appear dimmed in the Windows Firewall UI, you can use the Resultant Set of Policy snap-in to determine which policy settings are being applied to the computer. Any controls that are managed by Group Policy will appear dimmed in the Windows Firewall UI.

**To determine which policies are being applied to a computer**

1. At the command prompt, type **rsop.msc**, and press ENTER.
2. In the console tree, open **Computer Configuration**, open **Administrative Templates**, open **Network**, and then click **Network Connections**.
3. In the console tree, click **Windows Firewall**, and double-click any policies that appear in the details pane to learn more about the policy.
4. In console tree, click **Domain Profile**, and double-click any policies that appear in the details pane to learn more about the policy.
5. In the console tree, click **Standard Profile**, and double-click any policies that appear in the details pane to learn more about the policy.

# Additional Resources for Windows Firewall Troubleshooting

Updated: March 28, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For specific information about troubleshooting Windows Firewall problems, see the following resources:

- Troubleshooting Windows Firewall Settings in Windows XP Service Pack 2 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43102).
- Article ID 842242 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43103). This article describes why some programs might stop working after you install Windows XP with Service Pack 2 (SP2) and what you can do to overcome the problem.
- Article ID 884130 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43104). This article contains a list of programs that are known to experience loss of functionality when they run on Windows XP SP2.

For information about Windows Firewall known issues and best practices, see the following resources:

- Best Practices for Managing Windows Firewall
- Known Issues for Managing Resets, Startup, and Shutdown
- Known Issues for Managing Windows Firewall Profiles
- Known Issues for Managing Firewall Rules
- Known Issues for Managing IPsec, Multicast, and ICMP Settings
- Known Issues for Managing Windows Firewall Notifications
- Known Issues for Optimizing Windows Firewall
- Known Issues for Securing Windows Firewall

For general information about how Windows Firewall works and how to manage and configure Windows Firewall, see the following resources:

- Administering Windows Firewall
- Manually Configuring Windows Firewall in Windows XP Service Pack 2 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=33763).
- New Networking Features in Microsoft Windows XP Service Pack 2 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=33684).
- Windows Firewall Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=42729).

# Windows Server 2003 PKI Operations Guide

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

By David B. Cross and Ayman AlRashed

Microsoft Corporation

This document provides a guide for administrators on how to configure and operate a Windows certification authority. Various operational scenarios, custom configuration information, sample commands, and best practices are provided.

Windows Server 2003 provides a flexible and low TCO solution for deploying a public key infrastructure. Due to the complexity of customer environments and various organization requirements, a Windows Server 2003 certification authority (CA) may require configuration changes. This white paper provides operational best practices and configuration walkthroughs for some of the more common scenarios. It is not intended to cover the entire scope of all operational scenarios and configuration parameters possible with the Windows Server 2003 CA.

**In This White Paper**

- Basic Administrative Tasks
- Migrating from a Stand-alone to an Enterprise CA
- Windows Server 2003 PKI and Role-Based Administration
- Role Separation
- CA Auditing
- Setting Up CA Auditing
- Auditing and Event Management
- CA Maintenance
- Custom CA Configuration
- Enrollment Processing
- Tuning CA Database Performance
- Viewing Extended Information
- Managing Subject Relative Distinguished Names in the Certificate Subject
- Enabling the Netscape Revocation Method
- Configuring the SMTP Exit Module
- Using SSL to Communicate with Exchange Server
- Related Links for Windows Server 2003 PKI Operations
- Appendix A for Windows Server 2003 PKI Operations

# Basic Administrative Tasks

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

For day to day tasks, it is usually preferable to create a standard procedure. A procedure is usually organization-dependent as the processes and people differ from organization to organization. There are usually common practices employed by most organizations when doing the common day to day administrative tasks.

**Adding Certificate Templates to a CA**

A certificate template profiles certificates based on their intended use. When requesting a certificate from a Microsoft certification authority (CA), depending on their access rights, the certificate requester will be able to select from a variety of certificate types that are based on certificate templates, such as User and Basic EFS. The certificate template saves users from low-level, technical decisions about the type of certificate that they need. Instead, they can rely on the judgment of their administrators and use the template name that indicates the purpose of the certificate. If none of the preset certificate templates meets your needs, you can create new certificate templates and customize them for a variety of different uses.

| **Note** |
| --- |
| In addition to assigning the correct permissions for enrollment on a certificate template in Active Directory; you also need to add the template to the list of certificate templates a CA can issue if you want your users to start enrolling for this template. |

| **Note** |
| --- |
| Only Windows Server 2003 and Windows 2000 Enterprise CAs can issue certificates based on certificate templates; stand-alone CAs cannot use certificate templates. |

| **Note** |
| --- |
| You need to be part of the Enterprise Admins or the Domain Admins, or you need to have enough permissions to write to the Certificate Templates container in Active Directory. |

**To change the permissions on a certificate template for user enrollment**

1. Right-click the Certificate Templates node in the Certification Authority snap-in and select Manage.
2. Double-click a certificate template.
3. On the Security tab, check the Allow boxes for the Read and Enroll permissions.

**To add a certificate template to a CA**

1. Right-click the Certificate Templates node in the Certification Authority snap-in, and on the New submenu, select Certificate Template to Issue.
2. Select the appropriate template and click OK.

| **Note** |
| --- |
| You need to be a CA Administrator to add templates to the CA. |

**Delegating Administration of Certificate Templates**

Although most of the CA-related tasks are achieved through the administration of the CA itself, certain tasks are controlled through Active Directory, such as administration of Certificate Templates.

**To delegate the administration of Certificate Templates**

1. Right-click the **Certificate Templates** node in the **Certification Authority** snap-in and select **Manage**.
2. Double-click a certificate template.

3. On the **Security** tab, check the **Allow** boxes for the **Read** and **Write** permissions.

**Issuing Certificates**

There are some questions you need to answer and document before issuing a certain certificate. These questions are more relevant to how the certificate is issued from an operational side rather than a technical side.

1. Does my organization currently employ a Certificate Practice Statement (CPS) for this CA? If it does, did the requester meet all the requirements for enrollment?

2. Are there special requirements that the person issuing the certificate (such as being an Officer) that I must fulfill as an administrator?

3. Are there any documented operational procedures for my organization that I must follow when issuing certificates (such as backup)?

4. Are there any special attributes that must be included in the certificate that are not included in the request (such as Certificate Policy)?

When these questions are answered, and all the requirements are fulfilled, issue the certificate by logging on as a user with Certificate Manager (CA Officer) permissions:

1. Left-click the **Pending Requests** node in the **Certification Authority** snap-in.

2. Right-click the request, then select **Issue** on the **All Tasks** submenu.

If one of the requirements is not met, you can either ensure that requirements are met (such as making the user supply more authentication information) then issuing the certificate, or you can deny the request.

**To deny the request**

1. Left-click the **Pending Requests** node in the **Certification Authority** snap-in.

2. Right-click the request, then select **Deny** on the **All Tasks** submenu.

In either case, make sure you document your actions and the answers to all four questions.

| **Important** |
| --- |
| The policy module will always re-process requests that are pended and if template, configuration, or user group information has changed after the request was originally submitted, the policy module will re-evaluate the request on the new information only. |

| **Note** |
| --- |
| To re-submit a failed request and issue the failed request, a user must have both the CA Officer and CA Admin permissions on the CA. Obviously, this capability will not be possible when role separation is enabled on the CA. |

**Revoking Certificates**

Although certificates are usually used to enhance the trust in an organization, removing the trust from a certain certificate is sometimes required. Before you revoke a certificate, make sure you answer and document the following questions:

1. What is the reason for revoking this certificate?

2. Who requested the revocation of this certificate?

3. Will I ever need this certificate again (such as verification of signatures or decryption of messages)? If yes, what is the need (that is, verification of signatures, decryption of messages, normal usage)?

4. Are there special requirements for the person revoking the certificate (such as being an Officer) that I must fulfill as an

administrator?

5. Are there any documented operational procedures for my organization that I must follow when revoking certificates (such as backup)?

When all of these questions are answered, and all the requirements are fulfilled, revoke the certificate.

**To revoke the certificate**

1. Left-click the Issued Certificates node in the Certification Authority snap-in.
2. Right-click the certificate and select **Revoke Certificate** on the **All Tasks** submenu.
3. Choose the correct reason for revocation and click **Yes**.

Make sure you document your actions and the answers to all five questions.

| **Note** |
| --- |
| If you answered yes to question 3, and the need will be full or normal usage, make sure you choose **Certificate Hold** as the reason. This is the only reason that could allow a revoked certificate to be unrevoked. |

If you revoke a certificate and the reason is Certificate Hold, and you decide later that you want to unrevoke the certificate, you need to answer and document the following questions:

1. Why am I revoking this certificate?
2. Who requested this task?
3. Are there special requirements from the person unrevoking the certificate (such as being an Officer) that I must fulfill as an administrator?
4. Are there any documented operational procedures for my organization that I must follow when revoking certificates (such as backup)?
5. Does my organization currently employ a Certificate Practice Statement (CPS) for this CA, and if it does, did the requester meet all the requirements for unrevoking the certificate?

When all of these questions are answered, and all the requirements are fulfilled, unrevoke the certificate.

**To unrevoke the certificate**

1. Left-click the Revoked Certificates node in the Certification Authority snap-in.
2. Right-click the revoked certificate, and select **Unrevoke Certificate** on the **All Tasks** submenu.

Make sure you document your actions and the answers to all four questions.

| **Note** |
| --- |
| Unrevoking a certificate is considered dangerous if misused. Ensure proper operations and documentation when you unrevoke a certificate. |

# Migrating from a Stand-alone to an Enterprise CA

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Despite the best planning intentions, it may be necessary to change the configuration of a Windows-based certification authority from a stand-alone mode to an enterprise mode. It may also be necessary to change the configuration of a CA that was first installed as an NT 4.0 certification authority that was included in the NT 4.0 Option Pack. For example, an NT 4.0 CA may be upgraded in place to a stand-alone CA and later converted to an enterprise CA to work with Exchange 2000. For NT 4.0 upgrade procedures, see the help files in Windows 2000 or Windows Server 2003. This section provides a walkthrough of the steps required to convert a Windows Server 2003 stand-alone CA to an enterprise CA.

| **Note** |
| --- |
| It is not possible to convert a root CA to a subordinate CA, or vice versa. |

The first step in migrating the CA is backing up the existing key pairs used by the CA and its database. To backup the CA keys and database, right-click the CA node in the MMC and choose **Back up CA** under **All Tasks**.

Next, you must back up the certificate database, the CA certificate, and the CA private key. Select **Private key and CA certificate** and **Certificate database and certificate database log**, then choose the appropriate path for the backup files.

| **Note** |
| --- |
| The backup path should not contain old backup files. Use the command-line tool Certutil.exe if you want to overwrite old backup files. |




Enter a strong password. This password is used to protect the CA's private key.

| **Important** |
| --- |
| Do not lose this password or you will not be able to restore the keys on the new CA. |

Review the summary, and click **Finish** to complete the backup.




You have now successfully backed up the CA keys and database. Next, you should remove the stand-alone CA from the server by uninstalling it. Uninstall the CA by removing the **Certificate Services** from the Windows Components.




Join the computer to a domain in the forest if it is not already joined to one.

**Best Practice** The recommended best practice is to install CAs as a member of the root domain in the forest to provide centralized administration and control of the PKI services. For additional best practices, see the Windows Server 2003 Resource Kit.

Reinstall the CA by adding the **Certificate Services** to the Windows Components.




Select **Enterprise root CA** as the CA Type, and select custom settings for the key generation.

| **Note** |
| --- |
| You must be an Enterprise Admin to install an Enterprise CA. |




Choose the CSP that has access to the old CA keys, and choose the same keys and certificate used with the old CA.

| **Note** |
| --- |
| If your CA has multiple keys, choose the original key used by the old CA. This can be determined by the number appended to the key where no number means oldest. |




Select **Preserve existing certificate database** to use the old database. This enables the new Enterprise CA to keep track of any pending requests to the old stand-alone CA, as well as any certificates issued or revoked by it.




When prompted for stopping the IIS service, click **Yes** to finish the installation of the CA.

# Windows Server 2003 PKI and Role-Based Administration

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Windows 2000 Server CA administration is changed significantly with Windows Server 2003 CA role-based administration. Windows 2000 Server administrators can perform any activity on a Windows 2000 Server CA, but once CA roles are assigned on a Windows Server 2003 CA, its administrators are subject to its roles. Administrators who could perform all tasks on a Windows 2000 Server CA will only be able to perform the tasks associated with their role on the Windows Server 2003 CA. After upgrading a Windows 2000 Server CA to a Windows Server 2003 CA, its administrators need to be assigned to the roles defined in the role-based administration for the Windows Server 2003 CA.

Windows Server 2003 was designed with the needs of organizations in mind to provide role-based administration of a public key infrastructure. Windows Server 2003 certification authorities were also designed to meet the role definitions defined in version 1.0 of Certificate Issuing and Management Components Family of Protection Profiles found at http://go.microsoft.com/fwlink/?LinkId=59668

Role-based administration can be used to organize CA Administrators into separate, predefined task-based roles, each with its own set of tasks. Roles are assigned using each user's security settings. Roles are assigned to a user by assigning that user the specific security settings that are associated with the role. A user who has one type of permission, such as Manage CA permission, may perform specific CA tasks that a user with another type of permission, such as Issue and Manage Certificates permission, may not perform. Role-based administration is supported by both Windows Server 2003 enterprise and stand-alone certification authorities.

Role-based administration involves CA roles users and groups. To assign a role to a user or group, you must assign the role's corresponding security permissions, group memberships, or user rights to the user or group. These security permissions, group memberships, and user rights are used to distinguish which users have which roles. The following table describes the CA roles of role-based administration and the groups relevant to role-based administration.

| Roles and Groups | Security Permission | Description |
|---|---|---|
| CA Administrator | **Manage CA** permission | Configure and maintain the CA. This is a CA role and includes the ability to assign all other CA roles and renew the CA certificate. This is a separate role from the local Administrator role. |
| Certificate Manager | Issue and Manage Certificates permission | Approve certificate enrollment and revocation requests. This is a CA Officer role. |
| Backup Operator | Back up files and directories and Restore files and directories permissions | Perform system backup and recovery. This is an operating system role. |
| Auditor | Manage auditing and security log permissions | Configure, view, and maintain audit logs. This is an operating system role. |
| Enrollees | Authenticated Users | Enrollees are clients who are authorized to request certificates from the CA. This is not a CA role for the purposes of administration. |
| Read | All (except Enrollees) | Allows an entity to read records from the database. |

With the default installation, all CA roles are assigned and modified by local Administrators on the computer, Enterprise Admins and Domain Admins (if joined to a domain). Local Administrators, Enterprise Admins and Domain Admins are CA Administrators by default on an Enterprise CA. Only local Administrators are CA Administrators by default on a stand-alone CA. If the stand-alone CA is joined to an Active Directory domain, Domain Admins are also CA Administrators.

The CA Administrator and Certificate Manager roles can be assigned to either Active Directory users or local users in the local Security Accounts Manager (SAM) database. As a best practice, it is recommended to assign roles to group accounts instead of individual user accounts. Only CA Administrator, Certificate Manager (as Officer in the CIMC), Auditor (as Auditor in the CIMC), and Backup Operator (as Operator in the CIMC) are CA roles. The other users described in the following table are relevant to

role-based administration and should be understood before assigning CA roles.

Only CA Administrators and Certificate Managers are assigned using the Certification Authority Microsoft Management Console (MMC) snap-in. Other roles, users, and groups are specified in their related consoles. To change the roles of a user, you must change the user's security permissions, group membership, or user rights.

When key archival is configured on an enterprise CA running Windows Server 2003 Enterprise Edition, the subject obtaining a certificate from a CA will provide their private key to the CA. The CA stores that private key in its database until key recovery is desired. Only a Certificate Manager can get the encrypted private key blob out of the CA database, which is then passed on to key recovery agents (KRAs). For more information, see the Key Archival and Management in Windows Server 2003 white paper.

**Roles and Activities**

Each CA role has a specific list of CA administration tasks associated with it. The following table lists all the CA administration tasks along with the roles in which they are performed. One of the most important distinctions is the local Administrator versus CA Administrator role. The local Administrator applies to the local operating system privilege which may be required to perform some tasks associated with the operations of the CA. The CA Administrator role applies only to specific tasks within the functionality of the CA. The local Administrator will always have full control of the system including the CA and cannot be blocked from taking control of the CA. Therefore, it is very important to keep this fact in mind when assigning operational and delegated roles to the CA for management purposes.

| **Activity** | **CA Administrator** | **Certificate Manager** | **Auditor** | **Backup Operator** | **Local Administrator** | **Notes** |
|---|---|---|---|---|---|---|
| Install CA | | | | | X | |
| Configure policy and exit module | X | | | | | |
| Stop and start the Certificate Services service | X | | | X (only stop) | | |
| Configure extensions | X | | | | | |
| Configure roles | X | | | | | |
| Renew CA keys and certificates | | | | | X | |
| Define key recovery agents | X | | | | | |
| Configure Certificate Managers restrictions | X | | | | | |
| Delete single row in database | X | | | | | |
| Delete multiple rows in database (bulk deletion) | | | | | X | |
| Enable role separation | | | | | X | |
| Issue and approve certificates | | X | | | | |
| Deny certificates | | X | | | | |

| | | | | | | |
|---|---|---|---|---|---|---|
| Revoke certificates | | X | | | | |
| Reactivate certificates placed on hold | | X | | | | |
| Enable, publish, or configure CRL schedule | X | | | | | |
| Recover archived key | | X | | | | Only a Certificate Manager can retrieve the encrypted key data structure from the database. The private key of a valid Key Recovery Agent is required to decrypt the key data structure and generate a PKCS#12 file. |
| Configure audit parameters | | | X | | | By default, the local Administrator holds the system audit privilege. |
| Audit logs | | | X | | | By default, the local Administrator holds the system audit privilege. |
| Back up system | | | | X | | By default, the local Administrator holds the system backup privilege. |
| Restore system | | | | X | | By default, the local Administrator holds the system restore (backup) privilege. |
| Read CA database | X | X | X | X | X | By default, the local Administrator holds the system audit and backup privileges. |
| Read CA configuration information | X | X | X | X | X | By default, the local Administrator holds the system audit and backup privileges. |

**Note**

By default, enrollees are allowed to read CA properties and certificate request lists (CRLs), and can request certificates. On an Enterprise CA, a user must also have Read and Enroll permissions on the certificate template to request a certificate. CA Administrators, Certificate Managers, the Auditor, and Backup Operators have implicit Read privileges on the CA. An Auditor is based on a user that holds the system audit privilege. The local Administrator by default on each machine always holds the system audit privilege. If role separation is enabled, a separate user must be configured to hold the system audit privilege.> A Backup Operator is based on a user that holds the system backup privilege. In addition, the Backup Operator has the ability to stop the Certificate Services service (but not start it). As many or as few roles may be configured and used. It is not necessary to define all roles if only one specific role is assigned to a security group. It is required to have both CA Administrator and CA Manager roles to do bulk deletion, so any person who has both the roles will be able to perform bulk deletion. If role separation is enabled, this feature is not available. For issuing failed requests, it is required to have both CA Administrator and CA Manager permissions. If role separation is enabled, this feature is not available.

**Assigning Roles**

The CA Administrator for a CA assigns users to the separate roles of role-based administration by giving each user the security settings required by a role. The CA Administrator can assign a user to more than one role, but the CA is more secure when each user belongs to one role only. When each CA role belongs to one user only, fewer CA tasks can be compromised if a user's account becomes compromised.

The default installation setting for a stand-alone CA is to have members of the local Administrators security group as CA Administrators. The default installation setting for an Enterprise CA is to have local Administrators, Enterprise Administrators, and Domain Administrators as CA Administrators. To limit the power of any of these accounts, they should be removed from the CA Administrator and Certificate Manager roles once all CA roles are assigned; and they should also be removed from the Administrators group on the CA computer if it is not a Domain Controller. To list the roles a current user holds with a given CA, see the sample script in Appendix A.

**Best Practice** As a best practice, group accounts that have been assigned CA Administrator or Certificate Manager roles should not be members of the local Admin security group. Also, CA roles should only be assigned to group accounts and not individual user accounts.

| **Note** |
| --- |
| Membership in the local Administrators group on the CA is required to renew the CA certificate. Members of this group are considered to be all powerful on the CA with administrative authority over all other CA roles. |

# Role Separation

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The separation of CA roles can be enforced using role separation. Once enforced, role separation only allows a user to be assigned a single role. If a user is assigned to more than one role and attempts to perform an operation on the CA, the operation is denied. For this reason, before role separation is enabled, a user should be assigned only one CA role. This feature is valuable for large enterprises where the separation of roles ensures that the compromise of a user's account does not compromise the entire CA administered by the user.

| **Important** |
|---|
| Before role separation is enabled, each user assigned a CA role on the CA must only be assigned a single CA role on that CA. If a user is assigned more than one CA role, when role separation is enabled, the Certificate Services service will detect that a user has more than one role and deny the user's attempts to operate the CA. |

Only members of the local Administrators security group on a CA can enable and disable role separation. Enabling role separation requires editing the registry of the Windows Server 2003 Enterprise Edition running the Certificate Services service. Once this registry setting is edited to enable role separation, any assigned roles are in effect until the local Administrator of the server disables role separation through the registry. CA roles can be assigned and changed by the CA Administrator while role separation is enabled or disabled. While role separation is enabled, the CA Administrator cannot assign a user to more than one CA role. If the CA Administrator attempts to assign a user to a second CA role, the operation is refused.

| **Important** |
|---|
| It is possible for a user assigned a role to become locked out of administering a CA when role separation is enabled if the user is also assigned to a second CA role. If the CA Administrator is assigned to a second role, or assigns another role holder to a second role, the CA Administrator violates the rules of role separation by allowing a user to have two roles. Once the user is assigned to two roles, role separation will not allow that user to perform any activity on the CA, including, in the case of the CA Administrator, the activity of removing himself from one of the roles. |

To correct this configuration, the local Administrator of the server must disable role separation, remove the CA Administrator from the second role, and then restart the Certificate Services service. Following these steps, role separation can be enabled again.

**Windows 2000 and Windows Server 2003 Role-Based Administration**

During the upgrade from a Windows 2000 CA to a Windows Server 2003 CA, Windows 2000 CA permissions are upgraded to Windows Server 2003 CA roles according to the rules in the following table.

| Windows 2000 Permission | Windows Server 2003 Role or Permission |
|---|---|
| **Manage CA** permission | CA Administrator **and** Certificate Manager |
| Revoke Certificate permission | Certificate Manager |
| Approve/Issue Certificate permission | Certificate Manager |
| **Enroll** permission | **Enroll** permission |
| **Read** permission | **Read** permission |
| All other permissions listed in the Windows 2000 CA advanced security settings | **Read** permission |

| **Note** |
|---|
| You can assign certification authority roles for role-based administration on servers running any version of the Windows Server 2003 family, but you can only *enable* role separation on servers running Windows Server 2003 Enterprise Edition and Windows Server 2003 Datacenter, including the 64-bit version of Windows Server 2003 Enterprise Edition and 64-bit version of Windows Server 2003 Datacenter. |

To enable role separation, open a command prompt window and type

```
certutil -setreg ca\RoleSeparationEnabled 1
```

The Certificate Services service must now be stopped and started.

To stop and start the Certificate Services service, at the command prompt, type

```
net stop certsvc
net start certsvc
```

To disable role separation, open a command prompt and type

```
certutil -delreg ca\RoleSeparationEnabled
```

Again, the Certificate Services service must be stopped and started.

To display the role separation setting, at the command prompt, type

```
certutil -getreg ca\RoleSeparationEnabled
```

The following command will display all CA information including CA role separation status:

```
Certutil.exe –cainfo role
```

**Role Separation Validation**
All role operations are performed through the ICertAdminD DCOM interface once the CA is configured into role separation mode. Role separation may not be enabled nor enforced when role assignments are made. Role separation validation is only enforced when a person (administrator, operator, and so on) performs an action. Role separation enforcement rules are stored and read by the CA in the registry as binary blobs. Each role is defined as a bit (allow/deny). For more information about this interface, see the Platform SDK in MSDN.

**Certificate Managers**
The capability of Certificate Managers is supported to prevent a CA Officer from issuing certificates to everyone. This is implemented through an authorization callback and is stored in a virtual Security Descriptor in the registry of the CA. A second access check is performed on the Officer role to validate what user(s) and group(s) they manage (approve and revoke certificates). The GUI for each Certificate Manager has a list of users and groups that the Officer can approve, revoke, and so on.

If an Officer attempts to approve a request for a user to which the Officer is not authorized to manage, an access denied error will occur. This does not reject the request or remove it from the pending approval queue. In addition, a new user or group may be added to the Certificate Managers authorized list AFTER the certificate request has been made.

**Backup/Restore and Auditing under Role Separation**
When you enable Role Separation, members of the local Administrators group, including the local Administrator account, will not be able to back up or restore the CA, nor will they be able to enable auditing on the CA. Because Administrators have the permission to back up and restore the CA as well as enable auditing on the CA, the CA will not allow them to do any tasks because they are assigned multiple roles.

To assign the backup permission to another user, either add the user to the local Backup Operators group, or assign the user the Backup files and directories privilege in the Local Security Policy snap-in under User Rights Assignment. To assign the restore permission to another user, either add the user to the local Backup Operators group, or assign the user the Restore files and directories privilege in the Local Security Policy snap-in under User Rights Assignment. To give the permission needed to enable auditing, assign the user the Manage auditing and security log in the Local Security Policy snap-in under User Rights Assignment.

| **Note** |
| --- |
| You may need to refresh the local security policy by typing **gpupdate.exe** at the command prompt. To open the Local Security Policy, type **secpol.msc** at the command prompt. |

# CA Auditing

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

Auditing certification authority (CA) operations is supported for Windows Server 2003 Enterprise Edition. The audit events will be logged in the Security log and can be viewed using the Event Viewer utility. CA auditing is dependent on system object access auditing, and therefore, it is necessary for the system administrator to first enable object access auditing on the target system.

CA auditing is enabled by selecting which group of CA operations to audit in the MMC snap-in. The following sections describe each group of CA operations that can be audited.

**CA Audit Groups**
The following group of events can be configured to be audited:

- Back Up and Restore the CA Database
- Change CA Configuration
- Change CA Security Settings
- Issue and Manage Certificate Requests
- Revoke Certificates and Publish CRLs
- Store and Retrieve Archived Keys
- Start and Stop Certificate Services

**Back Up and Restore the CA Database**
By enabling auditing on this group, successful or failed attempts to back up the CA database will be logged to the system Security log. In addition, the CA service will detect on restart that the CA database has been restored. The restore events are logged to the system Security log.

**Change CA Configuration**
By enabling auditing on this group, successful or failed attempts to change CA configuration will be logged to the system Security log. This includes the following operations:

- Add/Remove Templates to the CA
- Configure the CRL Publication Schedule
- Modify Request Disposition for the Policy Module
- Modify Publish Cert Flags for the Exit Module
- Configure CRL Distribution Points (CDP)
- Configure Authority Information Access (AIA)
- Change the Policy Module
- Change the Exit Module
- Configure Key Archival and Recovery (KAR)

**Change CA Security Settings**
By enabling auditing for this group, successful or failed attempts to change CA security settings will be logged to the system Security log. This includes the following operations:

- Configure CA Roles for Role-Based Administration of the CA
- Configure Restrictions on Certificate Managers
- Configure CA Auditing

**Issue and Manage Certificate Requests**
By enabling auditing for this group, successful or failed attempts to issue and manage certificate requests will be logged to the system Security log. This includes the following operations:

- Incoming Certificate Requests
- Certificate Issuance
- Certificate Import
- Deletion of Rows in the CA Database

**Revoke Certificates and Publish CRLs**
By enabling auditing for this group, successful or failed attempts to revoke certificates and publish CRLs will be logged to the system Security log. This includes the following operations:

- Certificate Revocation
- CRL Publication

**Store and Retrieve Archived Keys**
By enabling auditing for this group, successful or failed attempts to store and retrieve archived keys will be logged to the system Security log. This includes the following operations:

- Archival of Subject Keys
- Retrieval of Subject Keys

**Start and Stop Certificate Services**
By enabling auditing for this group, successful or failed attempts to start and stop Certificate Services will be logged to the system Security log. This includes the following operations:

- Starting Certificate Services
- Stopping Certificate Services

# Setting Up CA Auditing

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

CA auditing depends on system object access auditing to be enabled. Therefore, to set up CA auditing for a system, a system administrator will have to

1. Enable Object Access Auditing on the system.

2. Enable auditing for the CA by selecting which group of events to audit in the MMC snap-in.

The following sections describe these steps in detail.

**Enabling Object Access Auditing**
**When the CA Is on a Domain Controller**
**To enable object access auditing when the CA is on a Domain Controller (DC)**

1. Select Start > Programs > Administrative Tools > Domain Controller Security Policy.
2. Expand Default Domain Controllers Security.
3. Expand Computer Configuration.
4. Expand Windows Settings.
5. Expand Security Settings.
6. Expand Local Policies.
7. Select Audit Policy.
8. Right-click Audit object access and select Properties.
9. Check Define these policy settings.
10. Check Success and Failure under Audit these attempts.
11. Click OK.

**When the CA Is on a Member or a Workgroup Server**
**To enable object access auditing when the CA is on a member or a workgroup server**

1. Select Start > Programs > Administrative Tools > Local Security Policy.
2. Expand Local Policies.
3. Select Audit Policy.
4. Right-click Audit object access and select Properties.
5. Check Success and Failure under Audit these attempts.
6. Click OK.

**Enabling Auditing on the CA**
**To enable auditing of the CA**

1. Open the MMC snap-in.
2. Right-click the CA and select Properties.
3. Click the Audit tab.
4. Check which groups of CA operations to audit.

5. Click OK.

# Auditing and Event Management

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

**Event IDs Used by Certificate Services**

The following event IDs are currently used by Certificate Services:

772. The Certificate Manager denied a pending certificate request.
773. Certificate Services received a resubmitted certificate request.
774. Certificate Services revoked a certificate.
775. Certificate Services received a request to publish the certificate revocation list (CRL).
776. Certificate Services published the certificate revocation list (CRL).
777. A certificate request extension changed.
778. One or more certificate request attributes changed.
779. Certificate Services received a request to shut down.
780. Certificate Services backup started.
781. Certificate Services backup completed.
782. Certificate Services restore started.
783. Certificate Services restore completed.
784. Certificate Services started.
785. Certificate Services stopped.
786. The security permissions for Certificate Services changed.
787. Certificate Services retrieved an archived key.
788. Certificate Services imported a certificate into its database.
789. The audit filter for Certificate Services changed.
790. Certificate Services received a certificate request.
791. Certificate Services approved a certificate request and issued a certificate.
792. Certificate Services denied a certificate request.
793. Certificate Services set the status of a certificate request to pending.
794. The Certificate Manager settings for Certificate Services changed.
795. A configuration entry changed in Certificate Services.
796. A property of Certificate Services changed.
797. Certificate Services archived a key.
798. Certificate Services imported and archived a key.
799. Certificate Services published the CA certificate to Active Directory.
800. One or more rows has been deleted from the certificate database.
801. Role separation enabled.

**Breakdown of Shared Event IDs**

**796:**
Property: 29

Index: 0

Type: 4

Adding/removing template to/from CA. Value is list of resulting

templates by name and object identifier.

Property: 26

Index: <KRA cert index>

Type: 3

Adding KRA cert to CA. Value is Base64 representation of the

certificate.

Property: 25

Index: 0

Type: 1

Removing KRA certificate from CA. Value is the total KRA certificate count. For example, you can add 7 KRA cert to CA but configure it to use 3 only. In this case, the property 25 (CR_PROP_KRACERTCOUNT) will be 7 and the property 24 (CR_PROP_KRACERTUSEDCOUNT) will be 3.

Property: 24

Index: 0

Type: 1

Adding/removing number of KRA certificates to use for key archival. Value is resulting number of certificates to use. A value of 0 indicates that KAR is disabled. For example, you can add 7 KRA cert to CA but configure it to use 3 only. In this case, the property 25 (CR_PROP_KRACERTCOUNT) will be 7 and the property 24 (CR_PROP_KRACERTUSEDCOUNT) will be 3.

**795:**
Node:

Entry: CRLPeriod or CRLPeriodUnits or CRLDeltaPeriod or

CRLDeltaPeriodUnits

Describe change in CRL publication schedule. Value of 0 for

CRLDeltaPeriodUnits means Delta CRL publishing disabled.

Node: PolicyModules\CertificateAuthority_MicrosoftDefault.Policy

Entry: RequestDisposition

Value: 1

Set CA to issue incoming requests unless specified otherwise.

Node: PolicyModules\CertificateAuthority_MicrosoftDefault.Policy

Entry: RequestDisposition

Value: 257

Set CA to keep incoming requests pending.

Node: ExitModules\CertificateAuthority_MicrosoftDefault.Exit

Entry: PublishCertFlags

Value: 1

Allow certificates to be published to the file system.

Node: ExitModules\CertificateAuthority_MicrosoftDefault.Exit

Entry: PublishCertFlags

Value: 0

Disallow certificates to be published to the file system.

Node: ExitModules

Entry: Active

Change in active Exit module. Value specifies name of new module. Blank

means none.

Node: PolicyModules

Entry: Active

Change in active Policy module. Value specifies name of new module.

Node:

Entry: CRLPublicationURLs

Change in CDPs or AIAs. Value specifies resultant set of CDPs.

Node:

Entry: CACertPublicationURLs

Change in AIAs or CDPs. Value specifies resultant set of AIAs.

**CA Audit Specification**

The following tables provide more information about the data contained in the CA audit events.

**Certificate Request Events**

| Audit Event | Audit Data |
|---|---|
| Certificate Request Submission | Request ID<br>UPN of Requestor |
| Certificate Request Processing | Request ID<br>Distinguished Name of Subject<br>Result of Processing (Issue, Pend, or Deny) |
| Certificate Issuance | Request ID<br>Certificate Serial No<br>Hash of Certificate<br>Certificate Template and Version<br>Sequence Number from CSP |
| Certificate Publication | Request ID<br>Distinguished Name of Object Updated<br>DC Name<br>Certificate Serial No |

| | |
|---|---|
| Certificate Revocation | Certificate Serial No<br>Time for Revocation<br>Reason for Revocation |
| Key Archived | Request ID<br>UPN of Requestor<br>List of Hashes of Recovery Agents Certificate(s) |

**Certificate Management Audit Events**

| **Audit Event** | **Audit Data** |
|---|---|
| Certificate Revocation Request | Issuer Name and Serial No of Signing Certificate (if signed)<br>Revocation Reason<br>UPN of Certificate Manager |
| Request Resubmission | Request ID<br>UPN of Certificate Manager |
| Denied Request | Request ID<br>UPN of Certificate Manager |
| Certificate Import | Request ID<br>UPN of Certificate Manager |
| Archived Key Retrieval | Request ID<br>Certificate Serial Number<br>Hash of Encrypted Blob<br>UPN of Certificate Manager |

**CA Administration Audit Events**

| **Audit Event** | **Audit Data** |
|---|---|
| Service Start or Stop | Hash of the Certificate Services Database Directories<br>Hash of the Database Log Directories<br>List of All Hashes of the Certificate Services Certificates<br>Sequence Number from CSP |
| CA Certificate Renewal Request | URN of Requestor<br>SKI |
| CA Certificate Installation | UPN of Installer<br>Hash of Certificate<br>Issuer Name<br>AKI<br>SKI |

| | |
|---|---|
| CRL Creation and Publication | CRL Type<br>AKI<br>Hash of CRL<br>Base and/or Delta CRL<br>CRL This Update<br>CRL Next update<br>URL Used to Publish<br>SKI (Identifier of CA)<br>UPN of Service Manager |
| Configure CRL Publication Policy | List of All CRL Policies Entries<br>UPN of Service Manager |
| Selecting Policy Module | Name of Active Policy Module (relative registry path)<br>UPN of Service Manager |
| Selecting Exit Module | Name of Active Exit Module (relative registry path)<br>UPN of Service Manager |
| Configure Policy Module | Name of Policy Module<br>Configuration Entry Name<br>New Configuration Entry Value<br>UPN of Service Manager |
| Configure Exit Module | Name of Exit Module<br>Configuration Entry Name<br>New Configuration Entry Value<br>UPN of Service Manager |
| Certificate Template Update | Template Name<br>Template Major and Minor Version Nos<br>List of Template Attributes<br>UPN of Service Manager |
| Key Archive Policy Change | Subject Name of Key Recovery Agent Certificate<br>Hash of Key Recovery Agent Certificate<br>Number of Key Recovery Agent Certificates Used<br>UPN of Service Manager |
| Data Base Row Deletion | Table<br>Row<br>UPN of Service Manager |

| | |
|---|---|
| Configure Certificate Managers Restrictions | Enable/Disable Restrictions<br><br>UPN of Each Certificate Manager, List of Users to Manage, Type of ACE(Allow/Deny)<br><br>UPN of Service Manager |
| Configure CA Security | UPN of Each User, Control Access Type, Type of ACE(Allow/Deny)<br><br>UPN of Service Manager |
| Configure CDP | List of All CDPs<br><br>UPN of Service Manager |
| Configure AIA | List of All AIAs<br><br>UPN of Service Manager |

**Backup/Restore Events**

| | |
|---|---|
| Start Service Backup | UPN of Operator<br><br>Backup Type<br><br>Backup Set ID<br><br>Data Integrity Check On\Off |
| Finish Service Backup | |
| Cancel Service Backup | |
| Start Service Restore | UPN of Operator<br><br>Restore Type<br><br>Backup Set ID<br><br>Data Integrity Check On\Off |
| Finish Service Restore | Integrity Check OK (if integrity checking on) |
| Cancel Service Restore | UPN of Operator |

**Audit Events**

| | |
|---|---|
| Audit Filter Change | Value of New Audit Filter<br><br>UPN of Auditor |

# CA Maintenance

Updated: March 4, 2013

Applies To: Windows Server 2003 with SP1

Setting up the PKI is a small step compared to the management and operational tasks that are associated with the public key infrastructure.

**CA Remote Maintenance**

A CA that is connected to the network can be maintained locally or through a remote connection; however, the CA maintenance and administration tools are best designed for local operation. This is because the CA administration is a sensitive operation and should be kept as secure as possible.

If the Certificate Services MMC Snap-In should be used for remote administration, see *Users Allowed to Manage the CA Cannot Access It Remotely* [271470] to take the appropriate steps to make the CA remotely accessible.

Even if it is technically possible, a CA should not be maintained through a Terminal Server session because it increases the attack spectrum, and because some administration tools—like certutil.exe—do not work properly if used in a Terminal Server session.

| **Note** |
| --- |
| A Windows 2000 CA may not be managed using the Windows Server 2003 version of the Certification Authority MMC snap-in, or vice-versa. |

**Publish the CRL of an Offline CA**

The actual offline CRL publication should be performed at a minimum of several days before the actual expiration of the previously issued CRL. This should be performed to provide a safety factor in case the offline root CA has a hardware or publication failure. Adequate time should be allotted to ensure that any errors or failures can be corrected and to enable actual publication and replication of the CRL to all CDP locations.

Once the CDP extensions have been updated on the CA, new CRL(s) should be published so that all clients who download the CRL have the latest download information (for example, Delta CRL URL).

**To manually publish a CRL on an offline CA**

1. Select the **Revoked Certificates** node of the certification authority MMC snap-in.
2. Right-click, select **All Tasks**, and then click **Publish**.

   A new base and delta (if configured) CRL will be published.

   
   

   A prompt will be displayed asking to confirm which type(s) of CRL(s) should be published with this request. Since only base CRLs are being published from the offline root CA, only the **New CRL** option will be available.

3. Click OK.




4. Root CA certificates may be published manually to Active Directory using the Windows Server 2003 version of the **certutil.exe –dspublish** command from the command line if logged in as an account that is a member of the Enterprise Admin Group or a domain admin from the root (first) domain in the forest.

   The equivalent command in Windows 2000 is called **dsstore.exe**. For more information about pushing CA certificates manually into Active Directory, see the following Knowledge Base article:

   - *HOW TO: Use the Directory Services Store Tool to Add a Non-Windows 2000 Certification Authority (CA) to the PKI in Windows 2000* [313197]

**CRL Re-Sign**

In some scenarios, it may not be possible to publish a CRL from an offline CA. In this case, with Windows Server 2003, the old CRL may be re-signed without using the certification authority. This process assumes the availability of the CA private key(s) outside of the CA to actually sign the CRL. To update an expiring CRL, the old CRL file will need to be retrieved first. It will be available in Active Directory if the CA is an Enterprise CA or if Active Directory was accessible when the CA was installed, or in the **%windir%\System32\CertSrv\CertEnroll** directory on the CA computer itself.

The simple syntax for re-signing a CRL is

```
certutil -sign <existing CRL file name> <resigned CRL file name>
```

You can also add or remove serial numbers, or remove extensions, or change the length of time the CRL will be valid through the certutil.exe –sign command.

The default is to re-sign the CRL to be valid starting 10 minutes prior to the signature (to allow for clock skew), and a lifetime (NextUpdate) equal to the old CRL. Use the following command to publish the CRL to Active Directory. Certutil will state whether the object in Active Directory was updated or if it was already up-to-date.

```
certutil -dspublish <resigned CRL file name>
```

Publishing to a file://, ftp:// or an http:// location can be done by copying the CRL file manually to that location. The following command executed on the CA computer should display the next time the CA expects to wake up and publish the next CRL:

```
certutil -getreg ca\CRLNextPublish
```

Dumping the CRL with the certutil command will display the 1.3.6.1.4.1.311.21.4 (Next CRL Publish) extension, which should be equivalent to the CRLNextPublish registry value (but the syntax of the two displays is different). Certutil -sign strips this extension out of re-signed CRLs, because the next issue date is misleading at best after re-signing. Dumping a certificate issued by the CA with the certutil command will display the ldap:///, http://, and file:// URLs where the CRL is stored.

**Administrative Process for Offline CRL Publication**

The following is a sample process outline that may be followed for publishing an offline CA's CRL:

Several days before the current CRL is due to expire, the offline Root CA system is removed from its protected location (usually in secured storage such as double-locked closet, combination safe, or other physically well-protected locations); normally two or more staff are present (for example, one IT administrator as well as a supervisory manager).

The root CA computer is powered up and logged on with an account with appropriate permissions.

The MMC is started and the CA's CRL is published to the local drive.

The CRL is copied to a disk or other removable media.

The offline Root CA server is logged out and shut down, and placed into a secured storage.

The removable media is taken to the publishing or staging servers, and the CRL is copied to the appropriate location(s), according to the current CDP locations published in non-expired certificates from the CA.

| **Important** |
| --- |
| It is extremely important for administrators to perform regular (test) restores of their offline CAs to see if the procedure for backup/restore performs as expected. Disaster recovery procedures and testing are paramount in operating a public key infrastructure. |

**Certification Authority Renewal**
**Best Practices for Renewing CAs**
Certification authorities are renewed or replaced for a multitude of reasons. The following are the most common reasons for CA renewal:

- Increase the lifetime of the CA
- Change the key used by the CA
- Increase the key size of the CA
- Add certificate policies to the CA (qualified subordination)
- CRL partitioning

The first three reasons may be performed by using a capolicy.inf file when renewing the CA. If renewing a root CA, it is important to understand that the root CA certificate will need to be re-distributed to all clients that will trust the root CA. Otherwise, existing clients will only be aware of the exiting root CA certificate and will have no mechanism for discovery of the renewal event.

When a CA is renewed, various objects and attributes will be updated or changed. If the CA is an Enterprise root or subordinate CA, the following objects in Active Directory will be updated:

The updated CA certificate (cACertificate) and cross-certificate (CrossCertificatePair if renewal with a new key pair is performed) will be published to the AIA container.

A new CRL will be published for every key pair the CA has used.

The new certificate will be published to the NTAUTH object.

The new certificate will replace the existing certificate in the enrollment services container.

| **Note** |
| --- |
| If the enrollment services container had been previously deleted, it will be replaced on renewal and the default templates will be re-installed as well, if they have been deleted. |




**CRL Partitioning**

CRL partitioning is another main reason why administrators often renew an issuing CA. When a CA is renewed with a new key, a new key and certificate are generated for that CA. When a new key and certificate are generated, the CA will use the new key as well as any unexpired previous keys corresponding to previous certificates when generating revocation information. Therefore, a CA may be using multiple keys at the same time and will publish multiple CRLs corresponding to those keys. This may be seen in the Certification Authority MMC snap-in by selecting the CA properties.




The renewal status of the CA may also be determined by examining the CA certificate itself. The CA version extension will identify how many times a CA has been renewed and how many times with a new key. In this case, the CA has been renewed three times and with each instance, a new key, hence the 3.3 version number as displayed in the following screen.




Once a CA has been renewed with a new key, only the new key will be used in signing new certificates. The unexpired previous key(s) will continue to be used to sign CRLs for certificates that were signed using the previous keys. Therefore, a CA may publish multiple CRLs at the same time, each using a different key. This method of CA renewal may be an ideal method for CRL size control and effective CRL partitioning using the Microsoft CA.

**Automatic RootCA Cross-Certificate Generation**

Windows Server 2003 has introduced the capability for Microsoft root certification authorities with access to Active Directory to automatically issue and publish a cross-certificate for a root CA that has been renewed. For example, when a Windows Server 2003 root CA is renewed with a new key, the root will cross-certify the renewed root CA certificate as being a qualified subordinate to the old root CA certificate. For more information about qualified subordination, see the Planning and Implementing Qualified Subordination for Using Windows Server 2003 Enterprise Server white paper.

This functionality is especially important to customers who have had an existing root CA trusted by other organizations, bridge CAs, or cross-certified by other organizations. To configure or disable this functionality, the following commands may be performed on the root CA.

- To force the root CA to use the CrossCA certificate template, the following command should be run. Otherwise, without this flag, the CA will never use the CrossCA certificate template (even if it is available), and will fall back to generating a certificate without the template, using predefined extensions:

  ```
  certutil -setreg ca\CRLFlags +CRLF_USE_CROSS_CERT_TEMPLATE
  ```

- To disable automatic CrossCA certificate generation, run the following command:

  ```
  certutil -setreg ca\CRLFlags +CRLF_DISABLE_ROOT_CROSS_CERTS
  ```

- To enable automatic CrossCA certificate generation again, run the following command:

  ```
  certutil -setreg ca\CRLFlags -CRLF_DISABLE_ROOT_CROSS_CERTS
  ```

- To force the root CA to use the CAExchange certificate template when generating CA encryption certificates on demand, run the following command. Without this flag, the CA will use the CAExchange cert template when available, and fall back to generating a certificate without the template, using predefined extensions.

  ```
  certutil -setreg ca\CRLFlags +CRLF_USE_XCHG_CERT_TEMPLATE
  ```

**Key Backup**

If you are using a smartcard or other hardware key and the computer fails, you will need to move the smartcard or key device to another computer, install the CA certificate, and possibly add the KeyProvInfo property to the CA certificate, so the private key's CSP and container name, and so on, will be available. This may be accomplished using the **certutil.exe –repairstore** command (see below). This is usually done automatically when a smart card is inserted into a reader.

If you are using a software-based CSP and the computer fails, it is necessary to use the **certutil -backupkey** command prior to the hardware failure to save the CA's keys and certificates in a PFX file (PKCS #12) encrypted with a password, and use **certutil -restorekey** on the second computer.

To add the KeyProvInfo property, use the following command. Include the **-user** option if the certificate may have been imported into the HKEY_CURRENT_USER personal store.

```
certutil -repairstore my CACertSHA-1Hash
```

or

```
certutil -repairstore -user my CACertSHA-1Hash
```

If necessary, dump the certificate with **certutil.exe –dump <file name>** to display the SHA-1 certificate hash.

**Backup/Recovery**
**Best practice** Backing up the certification authority database, the CA certificate, and the CA keys is essential to protect against the loss of critical data. The CA should be backed up on a regular basis (daily, weekly, monthly) based on the number of certificates issued over the same interval. The more certificates issued, the more frequently the CA should be backed up.

For additional information, see the Microsoft Knowledge Base:

Certificate Server Does Not Create Backups of Installed Keys [216922] *(This article applies to Windows 2000 only.)*

**Removing a CA from Active Directory**
Removing a CA from any public key infrastructure or Active Directory environment may have a significant impact on applications and services. Therefore, careful planning is always recommended before removing a CA. Always perform a complete backup and maintain that backup for a period of time should a restore be required at a later date.

**Uninstalling an Enterprise Certification Authority**
To decommission a root certification authority, all outstanding certificates issued by that CA should be revoked. After revocation, the Certificate Revocation List (CRL) should be published.

1. Revoke all issued certificates.

    - Start the Certification Authority MMC snap-in.
    - Click the Issued Certificates folder.
    - Highlight one of the issued certificates, and then press CTRL+A to select all issued certificates.
    - Right-click the highlighted certificates, select All Tasks, and then click Revoke Certificate.
    - In the Certificate Revocation dialog box, select Cease of Operation as the reason for revocation.
    - Click OK.

2. Increase the CRL publication interval.

    - In the Certification Authority MMC snap-in, right-click the Revoked Certificates folder.
    - Select Properties.
    - Increase the Publication Interval to a suitably long value (5 years).

      The lifetime of the CRL should be longer than the remaining lifetime of the certificates that have been revoked.

- Clear the Publish Delta CRLs check box if it is selected.
- Click OK.

3. Publish the new CRL.

   - In the Certification Authority MMC snap-in, right-click the Revoked Certificates folder.
   - Select All Tasks, and then select Publish.
   - Choose New CRL as the type of CRL to publish.
   - Click OK.

4. Stop Certificate Services.

   - At the command prompt, type

     ```
     certutil -shutdown
     ```

5. List all the key stores for the local computer.

   - At the command prompt, type

     ```
     certutil -key
     ```

   This will display the names of all the installed Cryptographic Service Providers (CSP) and the key stores associated with each provider. Among the listed key stores, you will see the name of your CA listed several times. The following is the sample output.

```
Microsoft Strong Cryptographic Provider:
  Enterprise Root CA
    AT_SIGNATURE
  Enterprise Root CA(11)
    AT_SIGNATURE
  Enterprise Root CA(13)
    AT_SIGNATURE
  Enterprise Root CA(4)
    AT_SIGNATURE
  Enterprise Root CA(14)
    AT_SIGNATURE
  Enterprise Root CA(9)
    AT_SIGNATURE
  Enterprise Root CA(7)
    AT_SIGNATURE
  Enterprise Root CA(6)
    AT_SIGNATURE
  MS IIS DCOM Server
    AT_SIGNATURE, AT_KEYEXCHANGE
  Enterprise Root CA(2)
    AT_SIGNATURE
  Enterprise Root CA(12)
```

```
    AT_SIGNATURE
  Enterprise Root CA(16)
    AT_SIGNATURE
  Enterprise Root CA(1)
    AT_SIGNATURE
  Microsoft Internet Information Server
    AT_SIGNATURE, AT_KEYEXCHANGE
  Enterprise Root CA-Xchg(7)
    AT_KEYEXCHANGE
  Enterprise Root CA(5)
    AT_SIGNATURE
  Enterprise Root CA(8)
    AT_SIGNATURE
  Enterprise Root CA(15)
    AT_SIGNATURE
  Enterprise Root CA(3)
    AT_SIGNATURE
```

6. Delete the private keys associated with the CA.

   - At the command prompt, type

     ```
     certutil -delkey <CA Name>
     ```

   If your CA name contains spaces, enclose it in quotation marks.

   | **Note** |
   | --- |
   | Repeat this step for all the key containers of your CA. You will need to do this if your CA has multiple certificates. |

7. List the key stores again to verify that the private key for your CA has been deleted.

8. Use Add or Remove Programs to uninstall Certificate Services.

**Active Directory Objects**

When Microsoft Certificate Services is installed on a server that is a member of a domain, several objects are created in the Configuration container in Active Directory. These objects are

- certificateAuthority object
  - Located in CN=AIA,CN=Public Key Services,CN=Services,CN=Configuration,DC=
  - Contains the CA certificate for the Certification Authority
  - Published Authority Information Access location
- crlDistributionPoint object
  - Located in CN=<servername>,CN=CDP,CN=Public Key Service,CN=Services,CN=Configuration,DC=
  - Contains the CRL periodically published by the certification authority
  - Published CRL Distribution Point location
- certificationAuthority object

    - Located in CN=Certification Authorities,CN=Public Key Services,CN=Services,CN=Configuration,DC=
    - Contains the CA certificate for the certification authority
- pKIEnrollmentService object
    - Located in CN=Enrollment Services,CN=Public Key Services,CN=Services,CN=Configuration,DC=
    - Created by Enterprise CAs. Contains information about the types of certificates the CA has been configured to issue. Permissions on this object can control which users are allowed to enroll against this CA.
- msPKI-PrivateKeyRecoveryAgent object
    - Located in CN=KRA,CN=Public Key Services,CN=Services,CN=Configuration,DC=
    - Contains the Key Recovery Agent certificate(s)

When the CA is uninstalled, only the pKIEnrollmentService object is removed. The other objects are left in place because there are likely still outstanding certificates issued by the CA. In order for clients to successfully process these outstanding certificates, they need to locate the AIA and CDP paths in Active Directory. Good practice is to revoke all outstanding certificates (Reason: Cease of Operation), extend the lifetime of the CRL, and publish it in Active Directory. When those outstanding certificates are processed by the various clients, validation should fail and those certificates will not be used.

If maintaining the CDP and AIA in Active Directory is not a priority, these objects can be safely removed.

**To remove all Certification Services objects from Active Directory**

1. Start Active Directory Sites and Services.
2. Click the View menu option, and select Show Services Node.
3. Expand the Services, and then expand Public Key Services.
4. Select the AIA node.
5. In the right-hand pane, locate the certificateAuthority object for your Certification Authority. Delete the object.
6. Select the CDP node.
7. In the right-hand pane, locate the Container object for the server where Certification Services is installed. Delete the container and the objects it contains.
8. Select the Certification Authorities node.
9. In the right-hand pane, locate the certificateAuthority object for your Certification Authority. Delete the object.
10. Select the Enrollment Services node.
11. In the right-hand pane, verify that the pKIEnrollmentService object for your Certification Authority was removed when Certificate Services was uninstalled. If not, delete it.
12. Select the Certificate Templates node.
13. In the right-hand pane, delete all the Certificate Templates.

    | ⚠**Caution** |
    | --- |
    | Delete all the Certificate Templates only if no other Enterprise CAs are installed in the forest. If the templates are inadvertently deleted, restore the templates from backup. |

14. Click the Public key Services node and locate the NTAuthCertificates object.
15. If there are no other Enterprise or Stand-alone CAs installed in the forest, delete the object, otherwise leave it alone.

**CA Database**
When Certification Services is uninstalled, the CA database is left intact in case the CA is to be recreated on another server.

**To remove the CA database**

- Delete the %systemroot%\system32\certlog folder.

**Domain Controller Cleanup**
Once the CA has been taken down, the certificates that have been issued to all the domain controllers need to be removed. This can be done quite easily using DSSTORE.EXE from the Resource Kit.

### To remove old domain controller certificates

1. At the command prompt on a domain controller, type

```
certutil -dcinfo deleteBad
```

2. Certutil.exe will attempt to validate all the DC certificates issued to the domain controllers. Certificates that fail to validate will be removed.

At this point, you can reinstall Certificate Services. After the installation is finished, the new root certificate will be published to Active Directory. When the domain clients refresh their security policy, they will automatically download the new root certificate into their trusted root stores.

### To force application of the security policy

- At the command prompt, type

```
gpupdate /target:computer
```

**Best Practices for CRLs**
**Certificate Revocation List Defaults**
Certificates that are revoked prior to expiration will remain in a published base CRL for one full base CRL period (defined by CA) after they expire. Certificates that expire will no longer be included in published CRLs after one additional base CRL is expired.

| **Note** |
| --- |
| The CRL table for the CA can be exported and converted to a tab-separated file for Microsoft Excel or other programs using the following commands: |

```
certutil -view <name of CRL file> > crl.txt
```

The following registry setting may be enabled on a CA to ensure that revoked certificates that have now expired are not removed from the CRL. Although most applications do not check CRLs for certificates that have expired, it is sometimes desirable in specific scenarios to maintain a public list of signing certificates that have been revoked.

To enable this option on the CA, use the following command:

```
certutil -setreg ca\CRLFlags +CRLF_PUBLISH_EXPIRED_CERT_CRLS
```

**Application Reliability**
Many applications rely on the availability of the certificate revocation list (CRL) and will fail completely should the CRL be inaccessible or out of date. One such example is the smartcard logon process. During smartcard logon, the client validates the domain controller certificate and the domain controller validates the user (client) certificate. If either validation fails, the smartcard logon process will fail. Therefore, it is highly desirable to keep the following best practices in mind when publishing CRLs:

1. The CRL should be valid for a reasonable period of time that would allow recovery of the CA should hardware or software failures occur. For example, a one-hour CRL publication period is most likely not adequate time to perform a hardware or software restoration in case of failure.

2. Set a reasonable CRL overlap period to dampen and overcome CRL publication or replication issues. For CRL overlap settings, see the next section. The CRL schedule for frequently publishing issuing CAs must be able to survive network and server downtime as well as take the replication latency of Active Directory into consideration. The CRL publication schedule must be longer than the maximum replication latency. In addition, the validity of a certificate must allow enough time to repair a broken network connection or to restore a down system. To enable this, the publication schedule must be shorter than the validity period of the CRL.

3. The private key of the CA and a copy of the CRL may be kept securely offline to manually sign and publish a valid CRL through certutil.exe when a catastrophic failure occurs.

4. For immediate denial of logon certificates, the account in Active Directory should be disabled. It is more efficient to delete or disable user accounts when a certificate should be revoked when a user should be denied access immediately.

5. CRLs should be published using Active Directory methods whenever possible for highest availability and best network performance. Always consider the expected propagation delay with a minimum replication time of 10 minutes.

6. CRLs should not be published to Active Directory when the CRL publication period is shorter than the replication convergence time for the Active Directory forest.

**Revoking Large Numbers of Certificates**

When a large number of certificates are revoked, such as during an employee layoff, the Delta CRL size may increase significantly due to the large number of entries, and almost all clients will refer to the older Base CRL. This situation will happen even if a new base CRL is published right after the revocation of the certificate until the new base is fully propagated.

To overcome this particular scenario where the Delta CRL is very large, perform the following steps on the CA:

1. Modify the registry values under the following registry key:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration\<Name of CA>**

   - Set the **CRLOverLapPeriod** to minutes. The default is hours.

   - Set the **ClockSkewMinutes** to 1minute. The default is 10.

2. Restart the CA.

3. Publish a new Base CRL. The Base CRL will have a CRLPropagationComplete time that will be just two minutes and any subsequent Delta CRLs will refer to this base CRL.

Once this has been completed, you can then restore the **CRLOverLapPeriod** and **ClockSkew** to the default values.

**Controlling CRL Size**

Often, it may be necessary to perform CRL partitioning to control the size of a Base CRL that is published from the CA. This is especially important for controlling both the data size that is replicated to Active Directory as well as the size of the data object downloaded by clients when performing revocation checks on certificates. CRL partitioning may be performed through CA key renewal which effectively creates a partitioned CRL for all subsequently issued certificates. For more information about this process, see the section on CA renewal.

Depending on the revocation reason chosen, the CRL will grow linearly in size by 29 bytes for every certificate that is revoked and added to the CRL. Accordingly, revoked certificates will be removed from the CRL when the certificate reaches its original expiration date. CAs may consider renewing the CA with a new key every 100-125K certificates to maintain a reasonable CRL size. This issuance number is based on the assumption that approximately 10 percent of the issued certificates will be revoked prior to their natural expiration date. If the actual or planned revocation rate is higher or lower for your organization, adjust the key renewal strategy accordingly.

| **Important** |
| --- |
| The more keys and certificates used by a CA will affect its performance in service restarts as each certificate and key must be validated before the CA will be operational. |

**Note**

Windows 2000 and Windows Server 2003 do not support the issuance of the IDP extension for partitioned CRLs; however, Windows XP and Windows Server 2003 clients may use a partitioned CRL (using IDP extension). This is technically different from the method identified previously.

**Remove Expired CRLs**

By default, a CA will maintain an expired CRL in the database and will keep this CRL also in the directory at the last known CDP publication point for historical purposes. Once the key of a CA expires, the CRL is published one final time and no additional changes are made to that CRL. A best practice is to maintain this CRL in the CA database for long-term validation and audit purposes. However, it may be removed by using the following command:

```
certutil -setreg ca\CRLFlags + CRLF_DELETE_EXPIRED_CRLS
```

For more information about how CRL checking and revocation status works see Windows XP: Certificate Status and Revocation Checking and How Certificate Revocation Works

**See Also**

**Other Resources**

Large CRLs: What is added to a Certificate Revocation List (CRL)?

# Custom CA Configuration

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

This section highlights various custom configuration scenarios that may be employed with a Windows Server 2003 certification authority to achieve various operational environment requirements. Various scenarios may apply differently to stand-alone or enterprise CAs.

**Ignore Offline CRL Errors on the CA**
Normally, a Windows Server 2003 CA will always check revocation on all certificates in the PKI hierarchy (except the root CA certificate) before issuing an end-entity certificate. To disable this feature, use the following command on the CA, and then restart the CA service:

```
certutil -setreg ca\CRLFlags +CRLF_REVCHECK_IGNORE_OFFLINE
```

**Configure Serial Number Generation**
In a Windows 2000 CA, two types of fixed-length serial numbers are generated. The registry can be modified to generate one or the other type. The default serial number is (from high to low): a DWORD from GetTickCount() + a USHORT CA cert index (0 to start) + a DWORD RequestId (10 bytes/20 hexadecimal digits). The alternate form is: one byte derived from the registry + a DWORD RequestId + 8 bytes of CryptGenRandom output + a USHORT CA cert index + a DWORD RequestId (19 bytes/38 hexadecimal digits).

To enable the alternate form and set the byte derived from the registry, use the following command:

```
certutil -setreg ca\HighSerial 0x33
```

The byte value specified will be modified to clear the sign bit and to set a bit in the high nibble to work around serial number encoding ambiguity bugs in certain non-Microsoft PKI applications.

In a Windows Server 2003 CA, three types of fixed-length serial numbers are generated. The default and alternate forms are the same as in Windows 2000. The Windows 2000 alternate form uses a new random 8 bytes generated by CryptGenRandom for each serial number. The new alternate form for Windows Server 2003 uses a fixed random 8 bytes from CryptGenRandom, generated during the first attempt to issue a certificate, and saved in the registry as 8 bytes of fixed CryptGenRandom output + a USHORT CA cert index + a DWORD RequestId (14 bytes/28 hexadecimal digits).

To enable the new alternate form in the registry, use the following command:

```
certutil -setreg ca\HighSerial 0xffffffff
```

Since the fixed random 8 bytes from CryptGenRandom are encoded as a string and saved in the registry, you could set them directly and cause them to be used for new serial numbers. In fact, any length hexadecimal string could be set in the registry (but there must be an even number of digits). The number of bytes used from the registry will be reduced if it would overflow a total of 19 bytes for the serial number. The high byte is manipulated as described previously to avoid problems with certain non-Microsoft applications. The IETF standards specify a maximum of 20 byte serial numbers.

**CA Key Usage**
Normally, a stand-alone CA certificate will contain digital signature, Certificate Signing, and CRL Signing as key usage values. For a stand-alone CA to issue a subordinate CA certificate without the digital signature key usage value, the following command must be performed on the stand-alone CA and the CA service restarted before the subordinate CA request is issued:

```
certutil -setreg policy\EditFlags -EDITF_ADDOLDKEYUSAGE
```

| **Note** |
| --- |
| Enterprise CAs enforce the key usage based on the subordinate CA template settings. |

**Disable DN Length Enforcement**
The original CCITT spec for the OU field says it should be limited to 64 characters. Normally, the CA enforces x.500 name length standards on the subject extension of certificates for all requests. It is possible that deep OU paths may exceed normal length restrictions.

To disable name length enforcement, run the following command on the CA, and then restart the CA service:

```
certutil -setreg ca\EnforceX500NameLengths 0
```

To restore the default setting, run the following command on the CA, and then restart the CA service:

```
certutil -setreg ca\EnforceX500NameLengths 1
```

**Maintaining the CA Database**

Windows Server 2003 allows certificate records to be purged manually through certutil.exe should the need arise to remove expired certificate records. The **certutil.exe –deleterow** command may be used for this task. Only a local administrator may delete multiple rows from a CA database. Standard JET database tools such as **ESEUTIL.EXE** may also be used with the certificate server database to perform defragmentation maintenance, and so on. For additional information, see the Windows Server 2003 help files.

**Securing DCOM Interfaces**

The Windows Server 2003 CA does not enforce encryption on the ICertRequest or ICertAdmin DCOM interfaces by default. Normally, this setting is not required except in special operational scenarios and should not be enabled. Only Windows Server 2003 machines by default support DCOM encryption on these interfaces. For example, Windows XP clients will not by default enforce encryption on certificate request to a Windows Server 2003 CA.

To enable encryption to be forced, perform the following procedures from the command line:

```
certutil -setreg ca\InterfaceFlags [+|-]IF_ENFORCEENCRYPTICERTREQUEST
certutil -setreg ca\InterfaceFlags [+|-]IF_ENFORCEENCRYPTICERTADMIN
```

# Enrollment Processing

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

**Enabling Referrals Across Forests**
Windows Server 2003 native Active Directory domains now support Kerberos trust and referrals across forests. By default, a Windows Server 2003 CA will not chase a referral for user or computer information in a trusted forest. When referrals are not chased and the user information is not available, the request will be denied if the user is enrolling from another forest. Referral chasing is not enabled by default as unintended template enumeration and enrollment may occur in some scenarios.

To enable referral chasing, use the following command on the certification authority and stop and start the service:

```
certutil -setreg policy\EditFlags +EDITF_ENABLELDAPREFERRALS
```

| **Note** |
| --- |
| For this to work, you need to use Kerberos Forest Trust as opposed to normal External Domain Trust. For more information, see the Windows Server 2003 help files. |

**Enabling Netscape Browser Enrollment**
The following configuration change must be made to a Windows Server 2003 CA to permit Netscape 6.2.2 and later browsers to perform enrollment through the Web enrollment pages.

To enable the parsing of request attributes for subject information, which is required for Netscape browser enrollment, use the following command:

```
certutil -setreg ca\CRLFlags +CRLF_ALLOW_REQUEST_ATTRIBUTE_SUBJECT
```

The certification authority must be stopped and re-started for this change to take effect. If this is not enabled, Netscape clients will receive the following error in the event log when the enrollment fails: The request subject name is invalid or too long.

**Adding CRL Information in an Authority Revocation List**
Neither Windows clients nor Windows certification authorities use or process authority revocation lists. However, some manual steps may be followed to artificially create an ARL.

To place a CRL in the ARL attribute in the directory, use the following command where the CAName, MachineName, Domain information, and name of the CRL are specific to the PKI environment and must be specified by the administrator:

```
certutil -addstore "ldap:///CN=CAName(KeyIndex),CN=MachineName,CN=CDP,CN=Public Key
Services,CN=Services,CN=Configuration,DC=Domain2,DC=Domain1,DC=com?auth
orityRevocationList" <name of CRL file.crl>
```

To delete an old CRL from the same attribute, use the same command, but use the –delstore instead of –addstore flag and specify the hash of the CRL instead of the CRL file as the last parameter. To display all CRLs with that attribute, use the –store flag and no ending parameter.

**Allowing and Blocking Extensions in Certificate Requests**
The template extension processing takes place *__before__* the three object identifier lists in the registry are processed.

The following configuration change must be made to a Windows Server 2003 CA to configure the CA to allow custom extensions to be added or blocked in certificates issued by the CA. Due to the fact that custom extensions comprise custom ASN.1 data, the CA is unable to parse and validate the information contained in the extension. If the certificate request contains the custom extension information properly encoded, the CA can be configured to pass the extension into issued certificates without validating the information. The extension must exist in the request as the CA will not generate this information. The request should be evaluated by a Registration Authority process before processing such requests.

To enable custom extension passing by using the object identifier of the custom extension defined by the organization, use the following command:

**certutil -setreg** policy**EnableRequestExtensionList** +< object identifier of extension to be added>

The certification authority must be stopped and re-started for this change to take effect.

**Example** The Netscape certtype extension is not enabled by default on a Microsoft CA and must be enabled using the previous mechanism.

To accept the Netscape certtype revocation extension to be included in issued certificates, use the following command:

**certutil -setreg** policy**EnableRequestExtensionList** +2.16.840.1.113730.1.1

To add a custom extension during CA installation through a capolicy.inf file or to submit a request using certreq.exe and a policy.inf file, the following information can be added to the extensions section of the *.inf file as an example. The object identifier of the extension must be specified followed by the Base64 representation of the value to be included in the extension. For example, AwIBBg== is the Base64 representation of the ASN.1 value: 03 02 01 06.

[Extensions]

1.3.6.1.5.5.7.1.3= AwIBBg==

To disable a certificate extension from being added to a certificate that is included by default in certificates issued by an enterprise CA, such as the S/MIME capabilities extension, use the following command, and then re-start the CA:

```
certutil -setreg policy\DisableExtensionList +< object identifier of extension to be added>
```

| **Note** |
| --- |
| When encountering template extensions with conflicting object identifiers, an Enterprise CA will override the EnableRequestExtensionList and EnableEnrolleeRequestExtensionList registry value behavior. If the certificate request includes extensions whose object identifiers conflict with the template extensions, the template extensions will override the request extensions. The EnableRequestExtensionList and EnableEnrolleeRequestExtensionList registry values merely cause the disabled bit to be cleared for any extensions in the two registry object identifier lists that exist. Extensions supplied by the template will already have the disabled bit clear. The object identifier list in EnableEnrolleeRequestExtensionList is only processed for templates that are configured to allow the request supply subject information. The object identifier list in EnableRequestExtensionList is processed for all templates. The object identifier list in DisableExtensionList causes the disabled bit to be *__set__* for any extensions in the list that exist—without regard to the origin of the extension (request or template). This prevents the listed extensions from appearing in issued certificates. This object identifier list is also processed for all templates. |

**Adding an E-Mail Address to the SubjectAltName Extension**
The following configuration option applies to stand-alone CAs only. Enterprise CAs may add e-mail address information to certificates automatically (if specified in the template) based on the user account information in Active Directory. One of the few configuration options that may be set for subject naming on a stand-alone CA is the ability for the CA to place the e-mail address of the requestor (authenticated user making the request) in the SubjAltName extension of an issued certificate.

To set the policy module to allow this option, perform the following steps using regedit.exe on the stand-alone CA:

Set the following REG_SZ value to Email in the registry, and then restart the CA:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration\<*CAName*>\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\SubjectAltName2

The <*CAName*> is the name of the local CA.

Set the e-mail request attribute to the users e-mail name by specifying the following request attribute string in the Web enrollment pages: email:joe@northwindtraders.com. As an alternative, set the e-mail name as part of the subject name E= email:joe@northwindtraders.com as part of the full subject distinguished name, or just joe@northwindtraders.com if there is a separate e-mail text box on the Web page.

The following is the syntax for specifying a Subject Alt Name 2 extension as a request attribute; it will only work on a Windows Server 2003 CA server and requires enabling the flag to process this attribute in the registry:

SAN:1.2.3.4=
{asn}Base64String&email=sample@bar.com&dns=sample.bar.com&dn="CN=xxx,OU=xxx,DC=xxx"&url="http://sample.com/default.htlm"&ipaddress=172.134.10.134&oid=1.2.3.4&upn=sample@bar.com&
b8ce-4fb4-aa58-3d1dc0e36b39

To set the Subject Alt Name 2 extension to a specific UPN value:

SAN:upn=sample@bar.com

To set it to a specific DNS name:

SAN:dns=sample.bar.com

To set both:

SAN:upn=sample@bar.com&dns=sample.bar.com

To enable the registry flag for a CA:

certutil –setreg policy\EditFlags +EDITF_ATTRIBUTESUBJECTALTNAME2

or certutil –setreg policy\EditFlags +0x40000

And restart the CA.

**How to Change the Validity Period for Certificates Issued from a CA**
Enterprise CAs set the validity of issued certificates through the settings on Active Directory-based templates. Stand-alone CAs enforce the validity period of issued certificates based on registry values. To change the default validity period intervals for all certificates that are issued by a stand-alone CA, set the following registry values:

HKLM\system\currentcontrolset\services\certsvc\configuration\<ca name>\validityperiod (=days/months/years)

HKLM\system\currentcontrolset\services\certsvc\configuration\<ca name>\validityperiodunits (=number of above)

# Tuning CA Database Performance

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

**Increasing CA Database Session Limit**

By default, the Windows Server 2003 certification authority allows only 20 concurrent sessions to the JET BLUE database, which is used for storing certificate information, and so on. The CA itself can use several connections as well as client enrollment requests or management tools that view the database. In general, the 20 concurrent session limit should be sufficient for most operations. To increase the max session limit, use the following command to increase the maximum number of sessions to 30, which is the highest limit tested with the Windows Server 2003 certification authority:

```
certutil -setreg DBSessionCount 30
```

The certification authority must be stopped and re-started for this change to take place.

# Viewing Extended Information

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

**Viewing Request Attributes**

Request Attributes are name-value string pairs that are passed to the certificate server and stored in the database for possible use by the policy module or exit module. They are intended to be used for customer-specific purposes to control the behavior of a custom policy or exit module. They do not directly affect the certificate content. They may be used by the policy module to affect certificate content, but that is determined by the custom policy module, not the default policy module. To review a Request Attribute for a specific request, use the following command, replacing nnnn with the Request ID of the request being examined:

```
certutil -view -restrict requested=nnnn -out attrib:all
```

**Removing CA Information from the Directory**

Enterprise CA information is stored in the configuration container of Active Directory, most specifically in the Enrollment Services container of the Public Key Services node. Various pieces of information stored in the Public Key Services node of the configuration partition in Active Directory can be viewed or removed with the PKI Health Tool.

# Managing Subject Relative Distinguished Names in the Certificate Subject

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

For the subject relative distinguished names in certificates issued by a Windows Server 2003 certification authority, the following list contains the object identifiers that are supported in the platform:

COUNTRY_NAME "2.5.4.6" "Country"

ORGANIZATION_NAME "2.5.4.10" "OrganizationalUnit"

ORGANIZATIONAL_UNIT_NAME "2.5.4.11" "Organization"

COMMON_NAME "2.5.4.3" "CommonName"

LOCALITY_NAME "2.5.4.7" "Locality"

STATE_OR_PROVINCE_NAME "2.5.4.8" "State"

TITLE "2.5.4.12" "Title"

GIVEN_NAME "2.5.4.42" "GivenName"

INITIALS "2.5.4.43" "Initials"

SUR_NAME "2.5.4.4" "SurName"

DOMAIN_COMPONENT "0.9.2342.19200300.100.1.25" "DomainComponent"

RSA_emailAddr "1.2.840.113549.1.9.1" "EMail"

STREET_ADDRESS "2.5.4.9" "StreetAddress"

RSA_unstructName "1.2.840.113549.1.9.2" "UnstructuredName"

RSA_unstructAddr "1.2.840.113549.1.9.8" "UnstructuredAddress"

DEVICE_SERIAL_NUMBER "2.5.4.5" "DeviceSerialNumber"

By default, the following relative distinguished names elements are allowed in the subject of certificates and in the following order when also specified in a version 2 template:

- 0: EMail
- 1: CommonName
- 2: OrganizationalUnit
- 3: Organization
- 4: Locality
- 5: State
- 6: DomainComponent
- 7: Country

The default list can be displayed on a CA by running the following command:

```
certutil -getreg ca\SubjectTemplate
```

A relative distinguished names component can be added to the allowable list by running the following command. In this

example, the title is added to the end of the list in the registry of the CA.

```
certutil -setreg ca\SubjectTemplate +title
```

Because the strings are ordered, using the "+title" parameter to add a string to the REG_MULTI_SZ registry value may not produce a certificate with the Subject RDNs in the desired order. You can also use regedit.exe to editing the registry directly may be easiest way to specify the desired order.

You can set the value to the Windows 2003 defaults via the following command:

```
certutil -setreg ca\SubjectTemplate
"EMail\nCommonName\nOrganizationalUnit\nOrganization\nLocality\nState\nDomainComponent\
nCountry"
```

Variations on this command can be used to specify any subset of the 16 supported RDNs in any desired order.

To remove the DC= component from the subject of a certificate issued by a subordinate stand-alone CA, run the following command:

```
certutil -setreg ca\SubjectTemplate -DomainComponent
```

# Enabling the Netscape Revocation Method

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

To enable a legacy Netscape (iPlanet) application certificate revocation service with a Windows Server 2003 CA, runs the following command on the CA:

```
certutil -SetReg Policy\RevocationType +AspEnable
```

If the IIS (ASP) pages are to be hosted on a separate computer or if the default URL to be used by the Netscape application server is different from the default, it may be reviewed by using the following command-line example:

certutil -getreg Policy\RevocationURL

Where the value is stored in the following registry key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration\ <*CAName*> \PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\RevocationURL:

RevocationURL REG_SZ = https://%1/CertEnroll/nsrev_%3.asp

The URL may be changed in the registry. Restart the CA after making a change. The following replacement variables may be used in the revocation URL:

SERVERDNSNAME "%1"

SERVERSHORTNAME %2"

SANITIZEDCANAME "%3"

CERTFILENAMESUFFIX "%4"

DOMAINDN "%5"

CONFIGDN "%6"

SANITIZEDCANAMEHASH "%7"

CRLFILENAMESUFFIX "%8"

CRLDELTAFILENAMESUFFIX "%9"

DSCRLATTRIBUTE "%10"

DSCACERTATTRIBUTE "%11"

DSUSERCERTATTRIBUTE "%12"

DSKRACERTATTRIBUTE "%13"

DSCROSSCERTPAIRATTRIBUTE "%14"

Note that for this revocation service to work, the application, service, or account connecting to this URL must have READ permissions in the certification authority MMC snap-in. If IIS is using a local account, follow the steps for enabling anonymous access in IIS and allowing Anonymous READ access to the CA.

| **Important** |
| --- |
| Allowing anonymous access to the CA may expose privacy or security concerns. |

# Configuring the SMTP Exit Module

Updated: December 13, 2010

Applies To: Windows Server 2003 with SP1

The following information can be applied to Windows Server 2003 with SP1 Enterprise and Datacenter versions only.

A Windows Server 2003 certification authority (CA) has the capability to send SMTP (e-mail) messages to users, administrators, and Certificate Managers regarding various operations on the CA. This function is enabled through the default exit module of the CA and is configured in the registry. By default, no messages are enabled for the CA.

**Sample Configuration Batch File**

The following batch file can be used as a sample to configure the SMTP exit module capabilities of the CA without editing the registry directly. The SMTP exit module can use various values stored in the CA database. BodyArg is a list of database columns to be defined and later called by variable names such as %%1, %%2, and so on. Variables must be called in sequential order as they are defined. The following text can be used in a batch file to configure the SMTP exit module options on a CA:

**Important**

Before using the batch file below, ensure that you replace:

- <ExchangeServerNameOrIP> with the actual name or IP address of the Exchange Server you want to use.
- <Account> with the administrative mailbox account you want to use.
- <EmailAddress> with the administrative email address you want to use.

```
@echo off
cd\
%systemdrive%
:Setup_SMTP_Server // Section for setting the name of the exchange
server to be used and type of authentication to be used. 1 means to use
Basic authentication, 2 means to use NTLM authentication, 0 is for Anonymous
authentication
certutil -setreg exit\smtp\SMTPServer "<ExchangeServerNameOrIP>"
certutil -setreg exit\smtp\SMTPAuthenticate 1
:Setup_CA_For_Exit_Module // Section for turning events on or off. In
this case, on.
certutil -setsmtpinfo -p "<Account>" Administrator
certutil -setreg exit\smtp\eventfilter +EXITEVENT_CRLISSUED
certutil -setreg exit\smtp\eventfilter +EXITEVENT_CERTDENIED
certutil -setreg exit\smtp\eventfilter +EXITEVENT_CERTISSUED
certutil -setreg exit\smtp\eventfilter +EXITEVENT_CERTPENDING
certutil -setreg exit\smtp\eventfilter +EXITEVENT_CERTREVOKED
certutil -setreg exit\smtp\eventfilter +EXITEVENT_SHUTDOWN
certutil -setreg exit\smtp\eventfilter +EXITEVENT_STARTUP
:CrlIssued // Section for setting CRLIssued parameters.
certutil -setreg exit\smtp\CRLissued\To "<EmailAddress>"
certutil -setreg exit\smtp\CRLissued\From "<EmailAddress>"
certutil -setreg exit\smtp\CRLissued\CC "<EmailAddress>"
certutil -setreg exit\smtp\CRLissued\bodyformat "A new CRL has been
issued"
certutil -setreg exit\smtp\CRLissued\titleformat "A new CRL was issued by %%1"
certutil -setreg exit\smtp\CRLissued\BodyArg ""
certutil -setreg exit\smtp\CRLissued\TitleArg +"SanitizedCAName"
:Denied // Section for setting Denied parameters
certutil -setreg exit\smtp\Denied\From "<EmailAddress>"
certutil -setreg exit\smtp\Denied\CC "<EmailAddress>"
certutil -setreg exit\smtp\Denied\titleformat "Your certificate request
was denied by %%1"
certutil -setreg exit\smtp\Denied\BodyArg ""
Certutil -setreg exit\smtp\Denied\BodyFormat ""
call Stop_Start_CA
certutil -setreg exit\smtp\Denied\BodyArg +"Request.RequestID"
```

```
certutil -setreg exit\smtp\Denied\BodyArg +"Request.RequesterName"
certutil -setreg exit\smtp\Denied\BodyArg +"Request.SubmittedWhen"
certutil -setreg exit\smtp\Denied\BodyArg +"Request.DistinguishedName"
certutil -setreg exit\smtp\Denied\BodyArg +"Request.DispositionMessage"
certutil -setreg exit\smtp\Denied\BodyArg +"Request.StatusCode"
Certutil -setreg exit\smtp\Denied\BodyFormat +"Your Request ID is: %%1"
Certutil -setreg exit\smtp\Denied\BodyFormat +"The Requester Name is: %%2"
Certutil -setreg exit\smtp\Denied\BodyFormat +"The Request Submission Date was: %%3"
Certutil -setreg exit\smtp\Denied\BodyFormat +"Subject Name: %%4"
Certutil -setreg exit\smtp\Denied\BodyFormat +"Request Disposition Message: %%5"
Certutil -setreg exit\smtp\Denied\BodyFormat +"Request StatusCode: %%6"
certutil -setreg exit\smtp\Denied\TitleArg +"SanitizedCAName"
:Certificate_Issued // Section for setting Issued parameters.
certutil -setreg exit\smtp\Issued\From "<EmailAddress>"
certutil -setreg exit\smtp\Issued\CC "<EmailAddress>"
certutil -setreg exit\smtp\Issued\titleformat "Your certificate has been issued by %%1"
certutil -setreg exit\smtp\Issued\BodyArg +"RawCertificate"
Certutil -setreg exit\smtp\Issued\BodyFormat ""
net stop certsvc
call Stop_Start_CA
Certutil -setreg exit\smtp\Issued\BodyFormat +"Request ID: %%1"
Certutil -setreg exit\smtp\Issued\BodyFormat +"UPN: %%2"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Requester Name: %%3"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Serial Number: %%4"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Valid not before: %%5"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Valid not after: %%6"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Distinguished Name: %%7"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Certificate Template: %%8"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Certificate Hash: %%9"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Request Disposition Message: %%10"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Copy and paste the
following in Notepad, save and install"
Certutil -setreg exit\smtp\Issued\BodyFormat +"Binary Certificate: %%11"
:Certificate_Pending // Section for setting Pending parameters.
certutil -setreg exit\smtp\Pending\From "<EmailAddress>"
certutil -setreg exit\smtp\Pending\CC "<EmailAddress>"
certutil -setreg exit\smtp\Pending\titleformat "Your certificate is pending on %%1"
Certutil -setreg exit\smtp\Pending\BodyFormat ""
call Stop_Start_CA
Certutil -setreg exit\smtp\Pending\BodyFormat +"Request ID: %%1"
Certutil -setreg exit\smtp\Pending\BodyFormat +"UPN: %%2"
Certutil -setreg exit\smtp\Pending\BodyFormat +"Requester Name: %%3"
Certutil -setreg exit\smtp\Pending\BodyFormat +"Time submitted: %%4"
Certutil -setreg exit\smtp\Pending\BodyFormat +"Distinguished Name: %%5"
Certutil -setreg exit\smtp\Pending\BodyFormat +"Certificate Template used: %%6"
Certutil -setreg exit\smtp\Pending\BodyFormat +"Request Disposition Message: %%7"
:Certificate_Revoked // Section for setting Revoked parameters.
certutil -setreg exit\smtp\Revoked\From "<EmailAddress>"
certutil -setreg exit\smtp\Revoked\CC "<EmailAddress>"
certutil -setreg exit\smtp\Revoked\titleformat "Your certificate was revoked by %%1"
Certutil -setreg exit\smtp\Revoked\BodyFormat ""
call Stop_Start_CA
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Request ID: %%1"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Revoked when: %%2"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Effective: %%3"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Reason for being revoked: %%4"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"UPN: %%5"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Requester Name: %%6"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Serial Number: %%7"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Was not valid until: %%8"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Was not valid after: %%9"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Distinguished Name: %%10"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Certificate Template: %%11"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Certificate Hash: %%12"
Certutil -setreg exit\smtp\Revoked\BodyFormat +"Request Status: %%13"
:Certificate_Authority_Shutdown // Section for setting Shutdown parameters.
certutil -setreg exit\smtp\Shutdown\To "<EmailAddress>"
certutil -setreg exit\smtp\Shutdown\From "<EmailAddress>"
certutil -setreg exit\smtp\Shutdown\CC "<EmailAddress>"
:Certificate_Authority_Startup // Section for setting Startup parameters.
certutil -setreg exit\smtp\Startup\To "<EmailAddress>"
certutil -setreg exit\smtp\Startup\From "<EmailAddress>"
```

```
certutil -setreg exit\smtp\Startup\CC "<EmailAddress>"
:Stop_Start_CA // This is just a sub-routine for stopping and starting the CA.
net stop certsvc & net start certsvc
:Exit
echo Certificate Services SMTP Exit module has now been configured.
echo .
pause
exit
```

**Revocation Status**

The SMTP exit module allows for e-mail notification when a certificate has been revoked. However, one of the database columns, Request.RevokedReason, can only be used to return a revocation value to the user in the e-mail body, instead of the string. For example, "Key Compromise" would be displayed as a value of 1. The revocation values and their corresponding definitions as defined in RFC 2459 are as follows:

0 = Unspecified

1 = Key Compromise

2 = CA Compromise

3 = Change of Affiliation

4 = Superseded

5 = Cease of Operation

6 = Certificate Hold

It may be necessary to list the values in the body of the e-mail message to easily interpret the revocation reason.

# Using SSL to Communicate with Exchange Server

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

A best practice in using the SMTP exit module of a certification authority with a Microsoft Exchange 2000 Server is to configure the CA to communicate with the Exchange Server over SSL (SMTP Protocol). The following steps should be implemented to enable this enhanced security configuration once the SMTP exit module is already configured on the CA for operation:

Make sure the Root CA that issued the server certificate for the Exchange Server is trusted by the CA.

Configure the Exchange Server SMTP Protocol to use SSL by enrolling for a Server Certificate using the Wizard in the Exchange System Manager console.

Open the Exchange System Manager on the Exchange server by selecting Start >

Programs > Microsoft Exchange > System Manager.

In the left-pane view, expand Servers > SERVER_NAME > Protocols > SMTP > Default SMTP Virtual Server.

Right-click Default SMTP Virtual Server, and then select Properties.

Click the Access tab, and then select the Certificate button under Secure communication.

Follow the directions in the Wizard and when completed, go back to the Default SMTP Virtual Server Properties and select Communication.

Select the Require secure channel checkbox. It is recommended to set Require 128-bit encryption.

At a CMD prompt on the CA, type the following command:

```
certutil -setreg exit\smtp\http://schemas.microsoft.com/cdo/configuration/smtpusessl 1
```

| **Note** |
| --- |
| This will add a REG_SZ value: http://schemas.microsoft.com/cdo/configuration/smtpusessl with a value of 1 to the SMTP key in the registry. |

# Related Links for Windows Server 2003 PKI Operations

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

For the latest information about Windows Server 2003, see the Windows Server 2003 Web site at http://www.microsoft.com/windowsserver2003/default.mspx

Whats New in Windows XP and Windows Server 2003 PKI?

http://www.microsoft.com/technet/prodtechnol/winxppro/plan/default.mspx

Windows XP Technical Articles

http://www.microsoft.com/technet/prodtechnol/winxppro/default.mspx

CMC - Certificate Management Messages over CMS

http://www.ietf.org/rfc/rfc2797.txt?number=2797

Troubleshooting Certificate Status and Revocation white paper

http://www.microsoft.com/technet/prodtechnol/winxppro/support/tshtcrl.mspx

Certificate Auto-enrollment in Windows XP

http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/default.mspx

# Appendix A for Windows Server 2003 PKI Operations

Updated: August 13, 2009

Applies To: Windows Server 2003 with SP1

The following basic VBScript can be run by an end user or administrator to enumerate all the enterprise CAs in the forest, list the individual CAs, and display the role(s) held by the current user for each of the CA(s) as appropriate:

```
on error resume next
dim certadm
dim configobj
dim config
dim index
dim roles(9)
roles(0) = "CAADMIN"
roles(1) = "OFFICER"
roles(2) = "AUDITOR"
roles(3) = "OPERATOR"
roles(8) = "READ"
roles(9) = "ENROLL"
set certadm = CreateObject("CertificateAuthority.Admin")
Call ExitOnError("Createobject CertificateAuthority.Admin")
set configobj = CreateObject("CertificateAuthority.Config")
Call ExitOnError("Createobject CertificateAuthority.Config")
index = 0
do while index <> -1
    config = ""
    config = configobj.GetField("config")
    if(config <> "") then
        wscript.echo "CA: ", config
        PrintRoles(config)
    end if
    index = configobj.Next()
    Call ExitOnError("ICertConfig::Next")
loop
sub PrintRoles(config)
    value = certadm.GetMyRoles(config)
    Call ExitOnError("GetMyRoles")
    wscript.echo "Roles: 0x" & hex(value)
    for count=0 to 9
        if((count<4 or count>7) and value mod 2) then
            wscript.echo "      ", roles(count)
        end if
        value = value \ 2
    next
End Sub
Sub ExitOnError(failedCall)
'    if err.number <> 0 then
 '   wscript.echo "Error calling " + failedCall + " : " + err.value
  '  end if
End Sub
```

This is a preliminary document and may be changed substantially prior to final commercial release of the software described herein.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

2003. Microsoft Corporation. All rights reserved. Microsoft, Active Directory, Visual Studio, Windows, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

# How to Use Windows Firewall with a Server Cluster

Updated: March 2, 2005

Applies To: Windows Server 2003 R2

If you enable Windows Firewall for a server cluster running Windows Server™ 2003 Service Pack 1, you must configure the firewall so that the cluster nodes can communicate as needed. To do this, run the Security Configuration Wizard, which is an optional component in Windows Server 2003 Service Pack 1 that can be installed through Control Panel.

| **Note** |
| --- |
| To perform the following procedures, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. |

**Using the Security Configuration Wizard to Configure Windows Firewall**

To use the Security Configuration Wizard to configure Windows Firewall on a server cluster node, you must do the following:

- Install the Security Configuration Wizard (an optional component).
- Use the Security Configuration Wizard to create a new security policy for a server cluster node.
- Turn on Windows Firewall (if you have not already turned it on) and apply the security policy you created to the appropriate server cluster nodes.

After turning on Windows Firewall and applying the security policy, confirm that the recovery action is set correctly for the Cluster service.

| **Important** |
| --- |
| In Windows Server 2003, Windows Firewall resets its TCP session mapping table whenever a significant IP addressing change occurs. This means that when failover occurs in a cluster, and the virtual IP address is moved to another node, any existing TCP session mappings are lost. Because of this, for the firewall on the new node, inbound traffic appears to be unsolicited and will be dropped. TCP eventually times out and renegotiates the sessions through the firewall on the new node, but the timeout can cause the affected programs to experience delays of up to two minutes. |

**To install the Security Configuration Wizard**

1. In Control Panel, open **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. Select the **Security Configuration Wizard** check box, click **Next**, and then click **Finish**.

**To use the Security Configuration Wizard to create a new security policy for a server cluster node**

1. Make sure you have installed the server cluster nodes and have configured the clustered resources and resource groups.
2. Make sure you have installed the Security Configuration Wizard as described in the previous procedure.
3. In **Administrative Tools**, open **Security Configuration Wizard** and read the information on the first page of the wizard.
4. Click **Next** and then select **Create a new security policy**.
5. Click **Next** and then select a cluster node to use as a baseline for this security policy. Later, you can apply this policy to the selected server or to any other server with a similar configuration.

You must be logged on as an administrator of the selected server.

6. Click **Next** several times until you see the **Select Server Roles** page.

7. On the **Select Server Roles** page, for **View**, select **All roles**.

8. Select **Cluster server** and any server roles that correspond to the clustered resources on the server cluster nodes. For example, for a cluster with a File Share resource, select **Cluster server** and **File server**. Clear any roles that do not correspond to the clustered resources on server cluster nodes.

9. Click **Next**, select the client features that are necessary for the node to function, and clear any client features that are not necessary.

   Examples of client features are **Domain member**, **Microsoft networking client**, **DHCP client**, **DNS client**, and **WINS client**.

10. Click **Next** and, on the **Select Administration and Other Options** page, select **Join a cluster**.

11. Select or clear check boxes for other options as appropriate.

    Examples of options on this page include **Remote Windows administration** and **Time synchronization**.

12. Click **Next** and select additional services required for the node.

    The services that are available are listed. For example, if antivirus software is installed, the services it provides might be listed.

13. Click **Next** and choose an option for the handling of unspecified services.

    To learn more about unspecified services, click the link on this wizard page.

14. Click **Next** and review the service changes resulting from the selections you have made. If necessary, go back to previous pages and change your selections.

15. Click **Next**.

16. If needed, specify additional configuration details in the next three sections of the wizard, which cover Network Security, Registry Settings, and Audit Policy. You can probably skip those three sections if the following is true for your configuration:

    - You have already specified all the roles the servers will have, and there are no additional ports (beyond those needed for those server roles) that need to be opened.

    - The clients and servers that connect to your cluster nodes are running Windows Server 2003, Microsoft Windows® 2000, or Windows XP, and are not running operating systems earlier than these.

    - The cluster nodes do not require specialized auditing.

17. Use the remaining pages of the wizard to save and view the security policy and to specify whether to apply the policy now or later.

**To turn on Windows Firewall and apply the security policy to server cluster nodes**

As with any configuration change to a cluster, it is best to observe the cluster carefully as you make the following change. Apply the security policy to one node, then observe that node and make adjustments if necessary before applying the security

policy to the next node.

1. If you have not already turned on Windows Firewall, click **Start**, click **Control Panel**, and click **Windows Firewall**. If you are prompted to start the Windows Firewall/Internet Connection Sharing (ICS) service, click **Yes**. Click **On**, and then click **OK**.

2. In **Administrative Tools**, open **Security Configuration Wizard**, and then click **Next**.

3. Select **Apply an existing security policy**.

4. Specify the security policy that you created for cluster nodes, or click **Browse** and then select the security policy.

5. Click **Next** and select the server to which you want to apply the security policy.

6. Click **Next** twice.

   The wizard applies the security policy.

7. Click **Next**, and then click **Finish**.

8. As appropriate, test the cluster node to ensure that it functions correctly with Windows Firewall enabled and the security policy in effect.

   If the security policy does not work correctly, you can run the wizard again to roll it back, and then turn off Windows Firewall. At that point, you can run the wizard to edit the security policy as needed, and then repeat this procedure to apply the modified security policy.

9. Repeat steps 1 through 8 for each of the cluster nodes.

**Confirming the Recovery Action Setting for the Cluster Service**

After turning on Windows Firewall and applying a security policy to server cluster nodes, confirm that the recovery action is set correctly for the Cluster service.

**To confirm that the recovery action is set correctly for the Cluster service**

1. On a cluster node, in Control Panel, open **Administrative Tools**.

2. Open **Services**.

3. Double-click **Cluster service**.

4. On the **Recovery** tab, for **First failure**, make sure **Restart the Service** is chosen.

5. For **Second failure** and **Subsequent failures**, make sure **Restart the Service** is chosen.

Note that the preceding procedure will set the recovery action for the Cluster service, but will not immediately restart the Cluster service. If necessary, after completing the procedure, start the Cluster service, either in the Services snap-in or by opening a command prompt and typing **net start clussvc**.

# Using Windows Server 2003 with Service Pack 1 in a Managed Environment: Controlling Communication with the Internet

Applies To: Windows Server 2003 with SP1

This white paper provides information about the communication that flows between components in Windows Server 2003 with Service Pack 1 (SP1) and sites on the Internet, and it describes steps to take to limit, control, or prevent that communication in an organization with many users.

This white paper includes the following sections:


Introduction to Controlling Communication with the Internet for Windows Server 2003 with SP1

Activation, Registration, and Resulting Internet Communication in Windows Server 2003 with Service Pack 1

Certificate Support and Internet Communication

Device Manager, Hardware Wizards, and Internet Communication

Dynamic Update and Internet Communication

Event Viewer and Internet Communication

File Association Web Service and Internet Communication

Help and Support Center: Features That Communicate with the Internet

HyperTerminal and Internet Communication

Internet Explorer 6.0 and Resulting Internet Communication in Windows Server 2003 with Service Pack 1

Internet Information Services and Internet Communication

Internet Protocol Version 6 (IPv6) and Internet Communication

NetMeeting and Internet Communication

Online Device Help and Internet Communication

Outlook Express 6.0 and Resulting Internet Communication in Windows Server 2003 with Service Pack 1

Plug and Play and Internet Communication

Program Compatibility Wizard and Internet Communication

Remote Assistance and Internet Communication

Search Companion and Internet Communication

Terminal Server Licensing and Internet Communication

Windows Error Reporting and Internet Communication

Windows Media Player and Internet Communication

Windows Media Services and Internet Communication

Windows Time Service and Internet Communication

Windows Update, Automatic Updates, and Internet Communication

Appendix A: Resources for Learning About Automated Installation and Deployment

Appendix B: Resources for Learning About Group Policy

Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key

Appendix D: Differences in Windows Server 2003 Before and After Service Pack 1

Appendix E: Windows Firewall and Security Configuration Wizard

Appendix F: Internet Connection Sharing and Network Bridge

Appendix G: Message Queuing and Internet Communication

Appendix H: Connection Manager and Internet Communication

Appendix I: Passport Manager Administration

Appendix J: Online Ordering Wizards, Add Network Place Wizard, and New Connection Wizard

# Introduction to Controlling Communication with the Internet for Windows Server 2003 with SP1

Applies To: Windows Server 2003 with SP1

Products in the Microsoft® Windows Server™ 2003 family include a variety of technologies that communicate with the Internet to provide increased ease of use. Browser and e-mail technologies are obvious examples, but there are also technologies such as Automatic Updates that help you obtain the latest software and product information, including bug fixes and security patches. These technologies provide many benefits, but they also involve communication with Internet sites, which administrators might want to control.

Control of this communication can be achieved through a variety of options built into individual components, into the operating system as a whole, and into server components designed for managing configurations across your organization. For example, as an administrator, you can use Group Policy to control the way some components communicate. For some components, you can direct all communication to the organization's own internal Web site instead of to an external site on the Internet. In addition, in Windows Server 2003 with Service Pack 1 (SP1), you can use Windows Firewall and the Security Configuration Wizard to help you control aspects of your configuration such as which services are running and which ports are open.

This white paper provides information about the communication that flows between components in Windows Server 2003 with SP1 and sites on the Internet, and it describes steps to take to limit, control, or prevent that communication in an organization with many users. The white paper is designed to assist you, the administrator, in planning strategies for deploying and maintaining products in the Windows Server 2003 family with SP1 in a way that helps to provide an appropriate level of security and privacy for your organization's networked assets.

This white paper provides guidelines for controlling components in the following operating systems:

- Windows Server 2003 with SP1, Web Edition
- Windows Server 2003 with SP1, Standard Edition
- Windows Server 2003 with SP1, Enterprise Edition
- Windows Server 2003 with SP1, Datacenter Edition
- Windows Server 2003, Standard x64 Edition
- Windows Server 2003, Enterprise x64 Edition
- Windows Server 2003, Datacenter x64 Edition

For more information about the services and software available in each of the editions of Windows Server 2003, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=46002

Note that when selecting an appropriate edition of the operating system, you can use the basic security principle of running only the services and software that a particular server requires. For example, if you have a server that requires only the operating-system services and software in Windows Server 2003 SP1, Web Edition, you might choose this edition instead of another edition for that server.

The white paper is organized around individual components found in the Windows Server 2003 family, so that you can easily find detailed information for any component you are interested in.

**What This White Paper Covers and What It Does Not Cover**

This section describes the following:

- Types of components covered in this white paper
- Types of components not covered in this white paper

- Security basics that are beyond the scope of this white paper, with listings of some other sources of information about these security basics

**Types of Components Covered in This White Paper**
This white paper provides:

- Information about components that in the normal course of operation send information to or receive information from one or more sites on the Internet. An example of this type of component is Windows Error Reporting. If you choose to use this component, it sends information to a site on the Internet.

- Information about components that routinely display buttons or links that make it easy for you to initiate communication with one or more sites on the Internet. An example of this type of component is Event Viewer; if you open an event in Event Viewer and click a link, you are prompted with a message box that says, "**Event Viewer will send the following information across the Internet. Is this OK?**" If you click **OK**, information about the event is sent to a Web site, which replies with any additional information that might be available about that event.

- Brief descriptions of Windows Firewall and the Security Configuration Wizard, which are designed to help you control certain settings related to network communication (and therefore are related to communication across the Internet). These components can help you control aspects of your configuration such as which services are running and which ports are open.

- Brief descriptions of components like Microsoft Internet Explorer and Internet Information Services (IIS) that are designed to communicate with the Internet. It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running servers that communicate across the Internet. This white paper does, however, provide basic information about how components such as Internet Information Services work, and it provides suggestions for other sources of information about balancing your organization's requirements for communication across the Internet with your organization's requirements for protection of networked assets.

**Types of Components Not Covered in This White Paper**
This white paper does not provide:

- Information about managing or working with applications, scripts, utilities, Web interfaces, Microsoft ActiveX® controls, extensible user interfaces, the .NET Framework, and application programming interfaces (APIs). These are either applications, or are layers that support applications, and as such provide extensions that go beyond the operating system itself.

  Windows Installer is not covered in this white paper, although Windows Installer includes some technology that (if you choose) you can use for installing drivers or other software from the Internet. Such Windows Installer packages are not described here because they are like a script or utility that is created specifically for communication across the Internet.

  Note that among the applications not covered in this white paper are Web-based and server-based applications, for example, server-based applications for databases, e-mail, or instant messaging. You must work with your application software provider to learn what you can do to mitigate any risks that are part of using particular applications (including Web-based or server-based applications), scripts, utilities, and other types of software that run on products in the Windows Server 2003 family.

- Information about components that store local logs that could potentially be sent to someone or could potentially be made available to support personnel. This information is similar to any other type of information that can be sent through e-mail or across the Internet in other ways. You must work with your support staff to provide guidelines about the handling of logs and any other similar information you might want to protect.

**Security Basics That are Beyond the Scope of This White Paper**
This white paper is designed to assist you, the administrator, in planning strategies for deploying and maintaining Windows Server 2003 with SP1 in a way that helps provide an appropriate level of security and privacy for your organization's networked assets. The white paper does not describe security basics, that is, strategies and risk-management methods that provide a foundation for security across your organization. It is assumed you are actively evaluating and studying these security basics as a standard part of network administration.

Some of the security basics that are a standard part of network administration include:

- Monitoring. This includes using a variety of software tools, including tools to assess which ports are open on servers and clients.
- Virus-protection software.
- The principle of least privilege (for example, not logging on as an administrator if logging on as a user is just as effective).
- The principle of running only the services and software that are necessary—that is, stopping unnecessary services and keeping computers (especially servers) free of unnecessary software.
- Strong passwords -- that is, requiring all users and administrators to choose passwords that are not easily cracked.
- Risk assessment as a basic element in creating and implementing security plans.
- Software deployment and maintenance routines to help ensure that your organization's software is running with the latest security updates and patches.
- Defense-in-depth. In this context, defense-in-depth (also referred to as in-depth defense) means redundancy in security systems. An example is using firewall settings together with Group Policy to control a particular type of communication with the Internet.

**Other Sources of Information about Security Basics**

The following books and Web sites are a few of the many sources of information about the security basics described previously:

- Howard, Michael, et al. *Designing Secure Web-Based Applications for Microsoft Windows 2000*. Redmond, WA: Microsoft Press, 2000.
- Howard, Michael, and David LeBlanc. *Writing Secure Code*. Redmond, WA: Microsoft Press, 2002.
- Kaufman, C., R. Perlman, and M. Speciner. *Network Security: Private Communication in a Public World*. Upper Saddle River, New Jersey: Prentice-Hall Inc., 2002.
- Smith, B., B. Komar, and the Microsoft Security Team. *Microsoft® Windows® Security Resource Kit*. Redmond, WA: Microsoft Press, 2003.

  For more information, see the Microsoft Press Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29168
- The Security Guidance Center on the Microsoft Web site at:

  http://www.microsoft.com/security/guidance/
- The Prescriptive Architecture Guides on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=29413
- The Web page focused on security for Windows Server 2003 on the Microsoft Windows Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46003
- The Web page focused on security on the Microsoft Developer Network (MSDN®) Web site at:

  http://msdn.microsoft.com/security/

- The Web pages focused on security on the Microsoft TechNet Web site at:

  http://www.microsoft.com/technet/security/

  http://go.microsoft.com/fwlink/?LinkId=46004

# Activation, Registration, and Resulting Internet Communication in Windows Server 2003 with Service Pack 1

Applies To: Windows Server 2003 with SP1

This white paper provides information about the communication that flows between components in Windows Server 2003 with SP1 and sites on the Internet, and it describes steps to take to limit, control, or prevent that communication in an organization with many users.

This section of the white paper provides information about:

- The purposes of activation and registration associated with a new installation or an upgrade
- How a computer running Windows Server 2003 with SP1 communicates with sites on the Internet during activation and registration
- Choosing volume licensing so that product activation need not take place (to limit the flow of information to and from Internet sites)
- Using Group Policy to prevent a connection between a server running Windows Server 2003 with SP1 and the Microsoft Web site for online registration

**Purposes of Activation and Registration Associated with a New Installation or an Upgrade**
This subsection briefly describes the differences between product activation and registration, and then describes the purpose of each.

Product registration involves the provision of personally identifiable information, such as an e-mail address, to Microsoft for the purpose of receiving information about product updates and special offers. Registration is usually done on a per-product basis and is not required for Windows Server 2003. If registration is completed, all registration information is stored using a variety of security technologies. When you register, you can specify what kinds of communication you wish to receive and whether your information can be used in particular ways. For example, if you provide a mailing address during registration, you can specify whether your mailing address can be shared with other companies. Other than the uses you specify, your registration information is never loaned or sold outside Microsoft.

Product activation involves the authentication with Microsoft of nonpersonal information, including the product identifier for Microsoft Windows Server 2003 with Service Pack 1 (SP1) and a hardware hash representing the computer, for the purpose of reducing software piracy. (A hardware hash is a non-unique number generated from the computer's hardware configuration.) Activation of Windows Server 2003 is required in situations where the product is not purchased through a volume licensing program such as Microsoft Select License, Microsoft Enterprise Agreement, or Microsoft Open License. Many computer manufacturers can bypass activation on software preinstalled on a new computer by binding the software to the computer’s basic input/output system (BIOS). In this situation, no activation of that software is required. Detailed information about product activation can be found on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=29484

For more information about volume licensing, see "Choosing Volume Licensing So That Individual Product Activation Need Not Take Place," later in this section.

Activation is aimed at reducing software piracy as well as ensuring that Microsoft customers are receiving the product quality that they expect. Activation means that a specific product key becomes associated with the computer (the hardware) it is installed on. After that happens, that product key cannot be used for activation on other computers (unless you are enrolled in a special program that permits additional activations, for example, a program through the Microsoft Developer Network [MSDN]).

**Overview: Activation and Registration in the Context of a Managed Environment**
Product activation is an anti-piracy technology designed to verify that software products have been legitimately licensed. If you have software re-imaging rights granted under a Microsoft volume license agreement, and if you obtained Windows Server 2003 through a retail channel or preinstalled by the computer manufacturer, you can re-image it with the product that you licensed through one of the Microsoft volume licensing programs. With volume licensing, there is no need to perform product activation.

Product registration involves the provision of personally identifiable information, such as an e-mail address, to Microsoft for the purpose of receiving information about product updates and special offers. In a managed environment, you might want to prevent administrators from registering the product with Microsoft. You can use Group Policy to do this.

**How a Computer Communicates with Sites on the Internet During Activation and Registration**
Windows Server 2003 with SP1 can be activated through the Internet or by phone. When it is activated through the Internet, the operating system communicates with Web sites as follows:

- **Specific information sent or received**: During activation of Windows Server 2003 with SP1, the following information is sent to the activation server at Microsoft:

    - Request information, that is, protocol information necessary for successfully establishing communication with the activation server.

    - Product key information in the form of the product ID, plus the product key itself.

    - A hardware hash (a non-unique number generated from the computer's hardware configuration). The hardware hash does not represent anything about the software or any personal information (such as user name or e-mail address). It is based on the MD5 message-digest hash algorithm, and consists of a combination of partial MD5 hash values of various computer components. The hardware hash cannot be used to determine the make or model of the computer, nor can it be backward-calculated to determine the raw computer information.

    - Date and time.

    - The language being used on the system (so that any error message that is sent back can be in the correct language).

    - The operating system being activated (and the version number of the activation software).

  Depending on your preference, the preceding information is either sent over the Internet to the activation system at Microsoft, or the product key information and hardware hash (combined into one number) are called in by phone.

- **Default setting and ability to disable**: Product activation can only be disabled by installing the operating system with software acquired through one of the Microsoft volume licensing programs. Product activation can be bypassed by many computer manufacturers if they bind the product to the computer’s BIOS instead. In all other cases, product activation cannot be disabled.

- **Trigger and notification for activation**: When activation is required, the operating system provides a reminder each time you log on and at common intervals until the end of the activation grace period stated in the End-User License Agreement (thirty days is the typical grace period). With software acquired through one of the Microsoft volume licensing programs, there is no need for activation, and therefore there are no reminders that appear about activation.

- **Trigger and notification for registration**: Registration is optional for Windows Server 2003. You can register at activation time by choosing appropriate options on the Windows Product Activation interface. As an alternative, you can type **regwiz /r** to start the Registration Wizard for Windows Server 2003 with SP1. Before the wizard starts and in the first page of the wizard, brief explanations notify you that completing the wizard will cause the product to be registered.

- **Logging**: Entries that track the progress of activation and registration (for example, return codes and error codes) are logged into a text file, *systemroot***setuplog.txt**. This file can be used for troubleshooting if activation (or any part of setup) fails. If you choose to register the product, two entries are made in this text file. One entry records the country or region that was chosen for the operating system. A second entry records whether you chose to have Microsoft (or the computer manufacturer) send information about product updates and special offers. No other registration data is logged.

- **Privacy, encryption, and storage for activation data**: Customer privacy was a paramount design goal in building the product activation technology. No personal information, such as user name or e-mail address, is collected as part of activation. The data is encrypted (using HTTPS) during transmission and is stored on servers located in controlled facilities at Microsoft. The data is accessible to a restricted number of server and program support personnel who oversee and maintain the activation servers and the product activation program.

To review the Microsoft online privacy statement on activation, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=29923

- **Privacy, encryption, and storage for registration data**: When you register at activation time (through the Windows Product Activation interface), registration data is encrypted (using HTTPS) during transmission. When you register by using the Registration Wizard (which you start by typing **regwiz /r**), registration data is encrypted (using HTTPS) during transmission unless the wizard is unable to establish an HTTPS connection through port 443 with the Microsoft registration server. In this situation, registration data will be sent unencrypted, using HTTP through port 80.

  Registration data, which contains information that you choose to send to Microsoft, is stored on servers with restricted access that are located in controlled facilities. The data can be seen by customer service representatives and marketing personnel. To review the Microsoft online privacy statement on product registration, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29508

- **Transmission protocol and port**:

  - **For Windows Product Activation**: When the operating system is activated through the Internet and a modem is not used, the first transmission uses HTTP through port 80 and goes to wpa.one.microsoft.com/ to check the HTTP response code. A response code of less than 500 indicates that a product activation server is available. (With a modem, there is only a check to see whether the modem can currently be used to make a connection to the Internet.) If the product activation server can be reached (or for a modem, if a connection to the Internet can be made), any activation or registration data that is sent by Windows Product Activation uses HTTPS through port 443.

  - **For the Registration Wizard**: When you register by using the Registration Wizard (which you start by typing **regwiz /r**), HTTPS is used through port 443 unless the wizard is unable to establish an HTTPS connection through port 443 with the Microsoft registration server. In this situation, HTTP is used through port 80.

**Activation Improvements in Windows Server 2003 Compared to Windows XP**

After the original release of Windows XP, Microsoft introduced improvements to activation that are aimed at ensuring that customers receive the full benefits of purchasing valid licenses. Windows Server 2003 and Windows Server 2003 SP1 contain these improvements, as do Windows XP SP1 and Windows XP SP2. Some of the changes include:

- The addition of a three-day grace period, where re-activation is required because of a significant hardware change. (When Windows XP was originally released, there was no grace period for reactivation after a significant hardware change.)

- The ability for volume license customers to encrypt their volume license product keys in unattended installations.

For more information about the changes to activation that were made after the original release of Windows XP, see the Microsoft Web site at:

- http://go.microsoft.com/fwlink/?linkid=29225

- http://go.microsoft.com/fwlink/?linkid=29226

**Choosing Volume Licensing So That Individual Product Activation Need Not Take Place**

If you use the rights granted under a volume licensing agreement to purchase or re-image software, you cannot and need not perform activation on the individual computers that are installed under the volume license. Qualifying as a volume licensing customer is not difficult. Customers can qualify for a Microsoft volume license by purchasing as few as five licenses. For more information, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=29878

**Procedure for Preventing Online Product Registration with Microsoft**

Use the following procedure to configure Group Policy to prevent online registration with Microsoft.

1. See Appendix B: Resources for Learning About Group Policy for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.

3. In the details pane, double-click **Turn off Registration if URL connection is referring to Microsoft.com**, and then click **Enabled**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the **Restrict Internet communication** policy setting, which is located in Computer Configuration/Administrative Templates/System/Internet Communication Management. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

# Certificate Support and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of the certificate functionality built into Microsoft Windows Server 2003 with Service Pack 1 (SP1), including the benefits of Update Root Certificates
- How the Update Root Certificates component in Windows Server 2003 with SP1 communicates with sites on the Internet
- How to control Update Root Certificates to limit the flow of information to and from the Internet

**Benefits and Purposes of Certificate Functionality**

Certificates, and the public key infrastructure of which they are a part, support authentication and encrypted exchange of information on open networks, such as the Internet, extranets, and intranets. A certificate securely binds a public key to the entity that holds the corresponding private key. With certificates, host computers on the Internet no longer have to maintain a set of passwords for individual subjects who need to be authenticated as a prerequisite to access. Instead, the host merely establishes trust in a certification authority that certifies individuals and resources that hold private keys. The host can establish this trust through a certificate hierarchy that is ultimately based on a root certificate, that is, a certificate from a certification authority that establishes a well-defined level of integrity and security for the hierarchy.

Examples of times that a certificate is used are when you:

- Use a browser to engage in a Secure Sockets Layer (SSL) session
- Accept a certificate as part of installing software
- Accept a certificate when receiving an encrypted or digitally signed e-mail message

When learning about public key infrastructure, it is important to learn not only about how certificates are issued, but about how certificates are revoked and how information about those revocations is made available to clients. This is because certificate revocation information is crucial for an application that is seeking to verify that a particular certificate is currently (not just formerly) considered trustworthy. Certificate revocation information is often stored in the form of a certificate revocation list, although this is not the only form it can take. Applications that have been presented with a certificate might contact a site on an intranet or the Internet not only for information about certification authorities, but also for certificate revocation information.

In an organization where servers run Windows Server 2003 with SP1, you have a variety of options in the way certificates and certification revocation lists (or other forms of certificate revocation information) are handled. For more information about these options, see the references listed in the next subsection, "Overview: Using Certificate Components in a Managed Environment."

The Update Root Certificates component in Windows Server 2003 with SP1 is designed to automatically check the list of trusted authorities on the Microsoft Windows Update Web site when this check is needed by an application. Specifically, if the application is presented with a certificate issued by a certification authority that is not directly trusted, the Update Root Certificates component (if present) will contact the Microsoft Windows Update Web site to see if Microsoft has added the certification authority to its list of trusted authorities. If the certification authority has been added to the Microsoft list of trusted authorities, its certificate will automatically be added to the trusted certificate store on the computer. Note that the Update Root Certificates component is optional, that is, it can be removed or excluded from installation on a computer running Windows Server 2003 with SP1.

**Overview: Using Certificate Components in a Managed Environment**

In an organization where servers run Windows Server 2003 with SP1, you have a variety of options in the way certificates are handled. For example, you can establish a trusted root authority, also known as a root certification authority, inside your organization. The first step in establishing a trusted root authority is to install the Certificate Services component. Another step that might be appropriate is to configure the publication of certificate revocation information to the Active Directory® directory service. When implementing public key infrastructure, we recommend that you also learn about Group Policy as it applies to certificates. Procedures for these steps are provided in the resources listed at the end of this subsection.

When you configure a certification authority inside your organization, the certificates it issues can specify a location of your choosing for retrieval of additional evidence for validation. That location can be a Web server or a directory within your organization. Because it is beyond the scope of this white paper to provide full details about working with certification authorities, root certificates, certificate revocation, and other aspects of public key infrastructure, this section provides a list of conceptual information and a list of resources to help you learn about certificates.

Some of the concepts to study when learning about certificates include:

- Certificates and the X.509 V3 standard (the most widely used standard for defining digital certificates) as well as the public key infrastructure for X.509 (PKIX). PKIX is described in RFC 3280, which you can search for on the Internet Engineering Task Force (IETF) Web site at:

  http://www.ietf.org/rfc.html

  You can also learn about PKIX on the Internet Engineering Task Force (IETF) Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29924

- Standard protocols that relate to certificates, for example, Transport Layer Security (TLS), Secure Sockets Layer (SSL), and Secure Multipurpose Internet Mail Extensions (S/MIME).

- Encryption keys and how they are generated.

- Certification authorities, including the concept of a certification authority hierarchy and the concept of an offline root certification authority.

- Certificate revocation.

- Ways that Active Directory and Group Policy can work with certificates.

The following list of resources can help you as you plan or modify your implementation of certificates and public key infrastructure:

- Help for products in the Windows Server 2003 family.

  You can view Help for products in the Windows Server 2003 family on the Web at:

  http://go.microsoft.com/fwlink/?linkid=29881

- The *Microsoft Windows Server 2003 Deployment Kit* and the *Microsoft Windows Server 2003 Technical Reference*.

  You can view links on the Windows Deployment and Resource Kits Web site at:

  http://www.microsoft.com/windows/reskits/

- "Troubleshooting Certificate Status and Revocation," a white paper on the Microsoft TechNet Web site at:

  http://go.microsoft.com/fwlink/?LinkId=27081

- Links to information about public key infrastructure for Windows Server 2003 on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=29886

In a medium-size to large organization, for the greatest control of communication with the Internet, it is recommended that you manage the list of certification authorities yourself.

**How Update Root Certificates Communicates with Sites on the Internet**

This subsection focuses on how the Update Root Certificates component communicates with sites on the Internet. The previous

subsection, "Overview: Using Certificate Components in a Managed Environment" provides references for the configuration choices that control the way other certificate components communicate with sites on the Internet.

If the Update Root Certificates component is installed on a server, and has not been disabled through Group Policy, and an application is presented with a certificate issued by a root certification authority that is not directly trusted, the Update Root Certificates component communicates across the Internet as follows:

- **Specific information sent or received**: Update Root Certificates sends a request to the Windows Update Web site, asking for the current list of root certification authorities in the Microsoft Root Certificate Program. If the untrusted certificate is named in the list, Update Root Certificates obtains that certificate from Windows Update and places it in the trusted certificate store on the server. No user authentication or unique user identification is used in this exchange.

  The Windows Update Web site is located at:

  http://windowsupdate.microsoft.com/

- **Default setting and ability to disable**: Update Root Certificates is installed by default in Windows Server 2003 with SP1. You can disable this component with Group Policy, or you can remove it or exclude it from installation on a server.

- **Trigger and user notification**: Update Root Certificates is triggered when the user is presented with a certificate issued by a root certification authority that is not directly trusted. There is no user notification.

- **Logging**: Events containing information such as the following will be logged:

  **For Event ID 7**:

  Description: Successful auto update retrieval of third-party root list sequence number from: *URL_for_Windows_Update_Web_Site*

  **For Event ID 8**:

  Description: Failed auto update retrieval of third-party root list sequence number from: *URL_for_Windows_Update_Web_Site* with error: *hexadecimal_error_value*

- **Encryption, privacy, and storage**: When requests or certificates are sent to or from Update Root Certificates, no encryption is used. Microsoft does not track access to the list of trusted authorities that it maintains on the Microsoft Windows Update Web site.

- **Transmission protocol and port**: The transmission protocol is HTTP and the port is 80.

**Controlling the Update Root Certificates Component to Prevent the Flow of Information to and from the Internet**
If you want to prevent the Update Root Certificates component in Windows Server 2003 with SP1 from communicating automatically with the Microsoft Windows Update Web site, you can disable this component with Group Policy, or you can remove it or exclude it from installation on servers. You can exclude the component during deployment by using standard methods for unattended installation or remote installation, as described in Appendix A, "Resources for Learning About Automated Installation and Deployment." If you are using an answer file, the entry is as follows:

```
[Components]
Rootautoupdate = Off
```

For information about how to disable Update Root Certificates through Group Policy, see "To Disable the Update Root Certificates Component by Using Group Policy," later in this section.

**How Disabling, Removing, or Excluding Update Root Certificates from Servers Can Affect Applications**
If a server is presented with a certificate issued by a root certification authority that is not directly trusted, and the Update Root Certificates component is not installed on that server, you (or your application) will be prevented from completing the action that required authentication. For example, you might be prevented from installing software, viewing an encrypted or digitally signed e-mail message, or using a browser to engage in an SSL session.

**Procedures for Preventing Root Certificates from Being Updated on an Individual Computer**
The following procedures describe:

- How to use Group Policy to disable the Update Root Certificates component.
- How to use Control Panel to remove the Update Root Certificates component from an individual computer running Windows Server 2003 with SP1.
- How to exclude the Update Root Certificates component during unattended installation of Windows Server 2003 with SP1 by using an answer file.

**To Disable the Update Root Certificates Component by Using Group Policy**

1. See Appendix B: Resources for Learning About Group Policy for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.
2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.
3. In the details pane, double-click **Turn off Automatic Root Certificates Update**, and then click **Enabled**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the **Restrict Internet communication** policy setting, which is located in Computer Configuration/Administrative Templates/System/Internet Communication Management. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

**To Remove the Update Root Certificates Component from an Individual Computer Running Windows Server 2003 with SP1**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components** (on the left).
4. Scroll down the list of components to **Update Root Certificates**, and make sure the check box for that component is cleared.
5. Follow the instructions to complete the Windows Components Wizard.

**To Exclude the Update Root Certificates Component During Unattended Installation by Using an Answer File**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file. For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment.
2. In the [Components] section of the answer file, include the following entry:

   **Rootautoupdate = Off**

# Device Manager, Hardware Wizards, and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Device Manager.
- How hardware wizards communicate with sites on the Internet.

  Device Manager provides an interface for viewing the configuration of hardware devices, and the wizards help with installing and configuring the correct driver for a device. Therefore, this section describes how the wizards communicate with the Internet, while providing background about Device Manager.

- How to control hardware wizards to limit the flow of information to and from the Internet.

For procedures related to disabling Windows Update, see Windows Update, Automatic Updates, and Internet Communication in this white paper.

**Benefits and Purposes of Device Manager**
Device Manager provides users and administrators with information about how the hardware on their computers is installed and configured, and how the hardware interacts with the computer's applications. With Device Manager, administrators can update the drivers (or software) for hardware devices, modify hardware settings, and troubleshoot problems.

| **Note** |
| --- |
| Only administrators or users with administrative credentials can install or update device drivers. |

**Overview: Using Device Manager in a Managed Environment**
In Microsoft Windows Server 2003, you access Device Manager through Administrative Tools\Computer Management. Administrators or users with administrative credentials will typically use Device Manager to check the status of hardware and to update device drivers. Administrators who have a thorough understanding of computer hardware might also use Device Manager's diagnostic features to resolve device conflicts and change resource settings.

Device Manager can work in conjunction with hardware wizards and Windows Update to deliver updated drivers for installed hardware. On a computer running Windows Server 2003 with SP1, when a driver is being upgraded, the person at the computer is prompted to find out whether to search the Windows Update Web site for an appropriate device driver. The same prompt is displayed if the person has attached a new Plug and Play device for which the driver is not included as part of Windows Server 2003 with SP1. As an administrator, you can prevent this, if you choose, by using Control Panel or Group Policy. The following subsection provides details about how hardware wizards interact with the Internet.

**How Hardware Wizards Communicate with Sites on the Internet**
Hardware wizards search for device drivers in slightly different ways, depending on the type of device for which a driver is needed:

- An existing device for which a new device driver is available.

  To update a device driver, a person logged on to Windows Server 2003 as an administrator opens Device Manager, selects a hardware device, and clicks **Update Driver** on the **Action** menu. Alternatively, the person can open Device Manager, right-click a hardware device, and click **Update Driver**. This activates the Hardware Update Wizard. If Windows Update device driver searching has not been disabled, by default, the Hardware Update Wizard prompts the person to find out whether to search the Windows Update Web site for an updated device driver.

- A new Plug and Play device for which the driver is not included as part of Windows Server 2003 with SP1.

  To install a new Plug and Play device, a person logged on to Windows Server 2003 as an administrator first attaches the device. In some instances, Windows Server 2003 then finds the appropriate driver locally and installs it without input

from the user. In other instances, Windows Server 2003 cannot find an appropriate driver locally, and the Found New Hardware Wizard starts. If Windows Update device driver searching has not been disabled, by default, the Found New Hardware Wizard prompts the person to find out whether to search the Windows Update Web site for an updated device driver.

- A device that is not Plug and Play.

  After a person logged on to Windows Server 2003 as an administrator attaches a device that is not Plug and Play, he or she can use Add Hardware in Control Panel to start the Add Hardware Wizard. This wizard does not connect to the Internet for device drivers.

When the person at the computer responds to a prompt by agreeing to the Internet search for a device driver, the Hardware Update Wizard or Found New Hardware Wizard communicate with the Windows Update Web site. Therefore, much of the information in this subsection is the same as for Windows Update. Additional details are as follows:

- **Specific information sent or received**: See the section Windows Update, Automatic Updates, and Internet Communication in this white paper. None of the communication between the computer and the Internet uniquely identifies the user.

- **Default settings**: By default, the hardware wizards prompt to find out whether to search Windows Update for a device driver.

- **Triggers**: A person logged on as an administrator starts the Hardware Update Wizard, or adds a new Plug and Play device for which Windows Server 2003 with SP1 cannot find an appropriate driver locally.

- **User notification**: By default, when the Hardware Update Wizard is started, it asks whether to search Windows Update for the device driver, and offers the following choices:

  - Yes, this time only

  - Yes, now and every time I connect a device

  - No, not this time

  Choosing the second choice, **Yes, now and every time I connect a device**, causes the Hardware Update Wizard to stop offering the preceding three prompts, but these prompts can be restored through System Properties (in Control Panel). To restore the three prompts, in System Properties, click the **Hardware** tab, click the **Windows Update** button, and choose **Ask me to search Windows Update every time I connect a new device.**

  To use System Properties to turn off Windows Update Device Driver searching, see "Procedures for Controlling Communication Between Hardware Wizards and Windows Update," later in this section.

- **Logging**: Errors that result from problems installing hardware devices without drivers are logged to the event log.

- **Encryption, access, privacy, transmission protocol, and port**: See the section Windows Update, Automatic Updates, and Internet Communication in this white paper. To view the privacy statement for Windows Update, go to the Web site and click **Windows Update Privacy Statement**:

  http://windowsupdate.microsoft.com/

- **Ability to disable**: You cannot disable Device Manager, but you can use Control Panel or Group Policy to prevent hardware wizards from searching for drivers on Windows Update. Alternatively, you can turn off all access to Windows Update.

**Controlling Hardware Wizards to Limit the Flow of Information to and from the Internet**

To prevent hardware wizards from searching Windows Update for device drivers, you can use either Control Panel or Group Policy. You can also turn off all access to Windows Update by using Group Policy.

If you turn off all access to Windows Update, users will still be able to use Device Manager to view information about their hardware devices. For administrators to be able to update drivers, there is the option for manually downloading driver updates from the Windows Update Catalog, or from an intranet server, whereby they can be distributed on your managed network as needed.

For more information about the Windows Update Catalog, see the Windows Update Web site at:

http://windowsupdate.microsoft.com/

**Procedures for Controlling Communication Between Hardware Wizards and Windows Update**
The procedures for turning off Windows Update device driver searching by using either Control Panel or Group Policy are included here. When you turn off Windows Update device driver searching, Windows Server 2003 does not search the Windows Update Web site for device drivers, regardless of any action taken through Device Manager or the hardware wizards.

For the procedure to turn off all access to Windows Update or configure automatic updating, see the section Windows Update, Automatic Updates, and Internet Communication in this white paper.

**To Turn Off Windows Update Device Driver Searching by Using Control Panel**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **System**.
3. In **System Properties**, click the **Hardware** tab.
4. Click **Windows Update**.
5. Click **Never search Windows Update for drivers**.

**To Turn Off Windows Update Device Driver Searching by Using Group Policy**

1. See Appendix B: Resources for Learning About Group Policy for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.
2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.
3. In the details pane, double-click **Turn off Windows Update device driver searching,** and then click **Enabled**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the **Restrict Internet communication** policy setting, which is located in Computer Configuration/Administrative Templates/System/Internet Communication Management. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

# Dynamic Update and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Dynamic Update
- How Dynamic Update communicates with sites on the Internet
- How to control Dynamic Update to limit the flow of information to and from the Internet

**Benefits and Purposes of Dynamic Update**

With Dynamic Update, Setup for Microsoft Windows Server 2003 can check the Windows Update Web site for new Setup files, including drivers and other files, while the server operating system is being installed. In an interactive installation (in contrast to an unattended installation), the person installing Windows Server 2003 chooses whether to allow Dynamic Update to occur.

In a managed environment, if you are using Setup (Winnt32.exe) for unattended installation, you can create a shared folder on a server and deliver Dynamic Update files to destination computers from that shared folder (instead of downloading the files directly from the Windows Update Web site to the computer being installed). For additional information about how to do this, see "How Dynamic Update Communicates with Sites on the Internet" and "Controlling Dynamic Update to Limit the Flow of Information to and from the Internet," later in this section.

Whenever an important update is made to any crucial Setup file, that update is made available through Dynamic Update functionality built into the Windows Update Web site. Some of the updated files will be replacements (for example, an updated Setup file) and some will be additions (for example, a driver not available at the time that the Setup CD was created). All files on the Dynamic Update section of the Windows Update Web site are carefully tested, and only files that are important in ensuring that Setup runs well are made available.

Using Dynamic Update reduces the need to apply patches to recently installed systems, and makes it easier to run Setup with hardware that requires a driver that was recently added or updated on Windows Update. For example, if a new video adapter requires a driver that was recently added to Windows Update, with Dynamic Update, the video adapter is recognized and supported during Setup.

Dynamic Update downloads only the files that are required for a particular computer, which means that the Dynamic Update software briefly examines the computer hardware. No personal information, such as user name or e-mail address, is collected. The only purpose for examining the hardware is to select appropriate drivers for it. This keeps the download time as short as possible and ensures that only necessary drivers are downloaded to the hard disk.

**Overview: Using Dynamic Update in a Managed Environment**

If you do not want Dynamic Update to connect to the Windows Update Web site during the installation of Windows Server 2003, you have several options:

- **Creating a shared folder on a server and delivering Dynamic Update files to destination computers from that shared folder**: You can create a shared folder on a server in your organization, download Dynamic Update files to that server, and by using Winnt32.exe for unattended installations, ensure that when Setup for Windows Server 2003 is run in your organization, Dynamic Update uses the files you placed on the server and does not connect to the Internet.
- **Avoiding Dynamic Update**: You can avoid using Dynamic Update, which means Setup will use only the files and drivers provided on the CD for Windows Server 2003. For more information, see "Avoiding Dynamic Update," later in this section.

The subsections that follow provide more information about these options.

For additional sources of information about performing unattended installations, see Appendix A: Resources for Learning About Automated Installation and Deployment.

**How Dynamic Update Communicates with Sites on the Internet**

This subsection focuses on the communication that occurs between Dynamic Update and the Windows Update Web site during an interactive installation (or a pre-installation compatibility check) when the computer has access to the Internet. This

subsection also provides some description of the default behavior of Dynamic Update with unattended Setup.

For information about how you can control the behavior of Dynamic Update during unattended installations, see "Controlling Dynamic Update to Limit the Flow of Information to and from the Internet," later in this section.

- **Specific information sent or received**: When Dynamic Update contacts the Windows Update Web site, it sends only the information necessary for appropriate drivers to be selected. In other words, it collects only necessary information about the hardware (devices) on that particular computer. No personal information, such as user name or e-mail address, is collected.

  The Setup files and drivers downloaded by Dynamic Update consist only of files that are important in ensuring that Setup runs successfully. Files with minor updates that will not significantly affect setup are not made available through the Dynamic Update section of the Windows Update Web site. Some of the updated files will be replacements (for example, an updated Setup file) and some will be additions (for example, a driver not available at the time that the setup CD was created).

- **Default behavior and triggers**: Dynamic Update may connect to the Internet, depending on how Setup is run. The following two tables provide details. The first table provides contrasting scenarios to show the broad outlines of choices among command-line options and answer file entries. The second table provides details about additional scenarios.

**Three Contrasting Scenarios for Running or Preventing Dynamic Update**

| Scenario | Steps to Take and Effect on Dynamic Update | Does Dynamic Update Connect to the Internet? |
|---|---|---|
| Interactive installation in which you permit Dynamic Update to run | Start Setup from the CD or a network and run it interactively. You are offered the choice of running Dynamic Update or skipping Dynamic Update. | Yes, if you choose to run Dynamic Update. |
| Unattended installation in which you prevent Dynamic Update from running | Run Winnt32.exe with the **/unattend** and **/DUdisable** command-line options. If the **/DUdisable** option is used, Dynamic Update is not triggered, regardless of whether an answer file is used. | No. |
| Unattended installation in which you create a shared folder on a server and deliver Dynamic Update files to destination computers from that shared folder | Prepare a shared folder as outlined in "Creating a Shared Folder on a Server and Delivering Dynamic Update Files to Destination Computers from that Shared Folder," later in this section. Then choose one of two methods for handling installations:<br><br>One method is to run Winnt32.exe with **/DUShare** = *path_to_downloaded_files*. Dynamic Update uses the folder specified in the /DUShare option and does not connect to the Internet.<br><br>Another method is to create an answer file that includes an **[Unattended]** section with an entry that specifies **DUShare** = *path_to_downloaded_files*. Run the Winnt32.exe command with the **/unattend**:*answer_file* option. Dynamic Update uses the folder specified in the DUShare entry and does not connect to the Internet. | No, Dynamic Update uses the files in the shared folder that you created. |

**Additional Scenarios for Running or Preventing Dynamic Update**

| Scenario | Steps to Take and Effect on Dynamic Update | Does Dynamic Update Connect to the Internet? |
|---|---|---|

| | | |
|---|---|---|
| Pre-installation compatibility check | Insert the Setup CD and choose the appropriate options for checking system compatibility. You are offered the choice of running Dynamic Update or skipping Dynamic Update. | Yes, if you choose to run Dynamic Update. |
| Unattended Setup in which you do not use an answer file and you allow Dynamic Update to run | Run the Winnt32.exe command with the **/unattend** option, but do not provide the name of an answer file and do not specify **/DUdisable** or any other options that affect Dynamic Update. Dynamic Update is triggered under these conditions for both unattended installation and unattended upgrade. | Yes. |
| Unattended Setup in which you use an answer file and you allow Dynamic Update to run | Create an answer file that includes an **[Unattended]** section with an entry that specifies **DUDisable = No**. Run the Winnt32.exe command with the **/unattend**:*answer_file* option. Dynamic Update is triggered. (However, note that if you run Winnt32.exe from the command line with the **/DUdisable** option, Dynamic Update is always prevented from running.) | Yes. |
| Unattended Setup in which you prevent Dynamic Update by creating an answer file that does not specify any options that affect Dynamic Update | Run the Winnt32.exe command with the **/unattend**:*answer_file* command-line option. By default, if the answer file does not specify any options that affect Dynamic Update, Dynamic Update is disabled. | No. |

- **User notification**: During an interactive installation, the user is notified when the choice of whether to run Dynamic Update is offered. During an unattended installation, there is no notification (unattended installation by definition means that no user interaction is required).

- **Logging**: By default, the progress of Setup is logged in *systemroot*\Winnt32.log. By using command options for the Winnt32.exe command, you can control the name of the log and the level of detail it contains.

- **Encryption**: The data is transferred from Microsoft using HTTPS.

- **Access**: No information about the hardware (devices) on a particular computer is saved or stored, so no one can access this information. The information is used only to select appropriate drivers.

- **Privacy**: Dynamic Update is covered by the same privacy statement that covers Windows Update. To view the privacy statement for Windows Update, go to the Web site and click **Windows Update Privacy Statement**:

  http://windowsupdate.microsoft.com/

- **Transmission protocol and port**: The transmission protocol is HTTPS and the port is 443.

- **Ability to disable**: You can control the behavior of Dynamic Update by running Setup in specific ways, as shown in the previous table. (Of course, you can also disable Dynamic Update by preventing access to the Internet or by blocking HTTPS over port 443.)

If you do not want to disable Dynamic Update but only want to prevent it from communicating with an Internet site, as noted earlier you can create a shared folder on a server and deliver Dynamic Update files to destination computers from that shared folder.

**Controlling Dynamic Update to Limit the Flow of Information to and from the Internet**
As summarized in "Overview: Using Dynamic Update in a Managed Environment" earlier in this section, if you do not want Dynamic Update to connect to the Windows Update Web site during the installation of Windows Server 2003, you have several options. With the appropriate methods for unattended installation, you can create a shared folder on a server and deliver Dynamic Update files to destination computers from that shared folder. Another alternative is to avoid using Dynamic Update at all.

**Creating a Shared Folder on a Server and Delivering Dynamic Update Files to Destination Computers from that Shared Folder**
This subsection briefly describes the steps for creating a shared folder on a server and delivering Dynamic Update files to destination computers from that shared folder. The subsection also provides links to more detailed information. The steps can

be summarized as follows:

- Step 1: Determine what packages you need to download from the Windows Update Web site.
- Step 2: Download the packages and prepare them and the folder they are in for use with Dynamic Update. This step includes extracting files and placing them in folders, as well as running Winnt32.exe with the **/duprepare** option, which creates subfolders and copies appropriate files to those subfolders. This step also requires other actions, such as sharing the folder and setting permissions.
- Step 3: Configure the answer file and Winnt32.exe settings for Dynamic Update (and for any other configuration options you want).
- Step 4: Run the unattended installations.

For more detailed information about performing the preceding steps, see the *Microsoft Windows Server 2003 Deployment Kit*, specifically the book titled, *Automating and Customizing Installations*. To view the *Microsoft Windows Server 2003 Deployment Kit*, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?linkid=29887

For additional sources of information about performing unattended installations, see Appendix A: Resources for Learning About Automated Installation and Deployment.

**Avoiding Dynamic Update**

You can avoid using Dynamic Update, which means that Setup will use only the files and drivers provided on the CD for Windows Server 2003. The method by which you avoid using Dynamic Update depends on how you are performing the installation:

- **Interactive installation**: During interactive installation (not unattended installation), you can select **No** when offered the option to use Dynamic Update. As an alternative, you can ensure that the computer does not have Internet access.
- **Unattended Setup**: Dynamic Update is disabled when you run Setup in specific ways, as shown in the tables in "How Dynamic Update Communicates with Sites on the Internet," earlier in this section. As the tables show, the simplest way to ensure that Dynamic Update does not run during unattended Setup is to use the **/DUdisable** option in the command line. This ensures that Dynamic Update will not occur during the installation.

**How Avoiding Dynamic Update or Directing Dynamic Update to a Server on Your Network Can Affect Users and Applications**

Regardless of whether you use Dynamic Update, you can obtain updated system and driver files after installations are complete (for example, through Windows Update or a service pack). Allowing Dynamic Update to run during Setup, however, helps ensure Setup success.

If you create a shared folder on a server and deliver Dynamic Update files to destination computers from that shared folder (instead of downloading the files directly from Windows Update to the computers), you can control the exact set of updated files to be installed. By contrast, when you download the current set of Dynamic Update files directly from the Windows Update Web site to users' computers, you might introduce inconsistencies among your destination computers because the Windows Web Site is periodically updated, and you cannot control when these updates occur.

**Procedures for Controlling Dynamic Update**

For detailed descriptions of Dynamic Update and procedures for controlling it, see the *Microsoft Windows Server 2003 Deployment Kit*, specifically the book titled, *Automating and Customizing Installations*. To view the *Microsoft Windows Server 2003 Deployment Kit*, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?linkid=29887

# Event Viewer and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Event Viewer
- How Event Viewer communicates with sites on the Internet
- How to control Event Viewer to prevent the flow of information to and from the Internet

**Benefits and Purposes of Event Viewer**
Administrators can use Event Viewer to view and manage event logs. Event logs contain information about hardware and software problems and about security events on your computer. A computer running Microsoft Windows Server 2003 records events in at least three kinds of logs: application, system, and security. A computer running Windows Server 2003 which is configured as a domain controller records events in two additional logs, the Directory service log and the File Replication service log. A computer running Windows Server 2003 which is configured as a Domain Name System (DNS) server records events related to DNS in an additional log.

Other types of events and event logs might be available on a computer, depending on what services are installed.

**Overview: Using Event Viewer in a Managed Environment**
The Event Log service starts automatically when you start the operating system. Administrators access event logs for a server through **Administrative Tools\Event Viewer**. They can obtain detailed information about a particular event by either double-clicking the event or selecting the event and clicking **Properties** on the **Action** menu. The dialog box gives a description of the event, which can contain one or more links to Help.

Links can either be to Microsoft servers or to servers managed by the software vendor for the component that generated the event. On Windows Server 2003, most events that originate from Microsoft products will have standard text containing a URL at the end of the description ("For more information, see Help and Support Center at http://go.microsoft.com/fwlink/events.asp").

When you click the link, you are asked to confirm that the information presented can be sent over the Internet. If you click **Yes**, the information listed will be sent to the Web site named in the link. The parameters in the original URL will be replaced by a standard list of parameters whose contents are detailed in the confirmation dialog box. This list is provided in the next subsection under "Specific Information Sent or Received."

You may want to prevent users from sending this information over the Internet through this link and accessing a Web site. Alternatively, you may want to redirect the requests that result from users clicking links in Event Viewer so that the requests go to a Web server in your organization. In Windows Server 2003 with SP1, you can do either of these things through Group Policy.

**How Event Viewer Communicates with Sites on the Internet**
In order to access the relevant Help information provided by the link in the Event Properties dialog box, you must send the information listed about the event. The collected data is confined to what is needed to retrieve more information about the event from the Microsoft Knowledge Base. User names and e-mail addresses, names of files unrelated to the logged event, computer addresses, and other forms of personal information are not collected.

The exchange of information that takes place over the Internet is as follows:

- **Specific information sent or received**: Information about the event sent over the Internet includes the following:
    - Company name (software vendor)
    - Date and time
    - Event ID (for example, 1704)
    - File name and version (for example, userenv.dll, 5.2.3790.1830)

- Product name and version (for example, Microsoft Windows Operating System, 5. 2.3790.1830)
- Registry source (for example, Userenv)
- Type of event message (for example, Error)

The information the user receives is from the Web site named in the link.

- **Default settings**: Access to Event Viewer is enabled by default.
- **Triggers**: The user chooses to send information about the event over the Internet in order to obtain more information about the event.
- **User notification**: When a user clicks the link, a dialog box listing the information that will be sent is provided.
- **Logging**: This is a feature of Event Viewer.
- **Encryption**: The information may or may not be encrypted, depending on whether the link uses HTTP or HTTPS.
- **Access**: No information is stored.
- **Privacy**: For information about privacy, in Help and Support, type **Linking to Microsoft for Help and Support**.)
- **Transmission protocol and port**: Communication occurs over the standard port for the protocol in the URL, using either HTTP or HTTPS.
- **Ability to disable**: The ability to send information over the Internet or to be linked to a Web site can be prevented through a Group Policy setting.

**Controlling Event Viewer to Prevent the Flow of Information to and from the Internet**

You can prevent administrators from sending information across the Internet and accessing Internet sites through Event Viewer by configuring Group Policy. Alternatively, you can redirect the requests that result from a person clicking a link in Event Viewer so that the requests go to a Web server in your organization. You can do these things by configuring Group Policy.

These Group Policy settings affect only the flow of information to and from an intranet or the Internet through Event Viewer, not the other functions of Event Viewer.

**Procedures for Preventing the Flow of Information to and from the Internet Through Event Viewer**

The following procedures tell how to use Group Policy to prevent the flow of information to and from the Internet through Event Viewer.

**To Use Group Policy to Prevent the Flow of Information to and from the Internet Through Event Viewer**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.
2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.
3. In the details pane, double-click **Turn off Event Viewer "Events.asp" links**, and then click **Enabled**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the **Restrict Internet communication** policy setting, which is located in Computer Configuration/Administrative Templates/System/Internet Communication Management. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

The following procedure tells how to use Group Policy to redirect the requests that result from users clicking links in Event Viewer so that the requests go to a Web server in your organization.

**To Use Group Policy to Redirect Links in Event Viewer to a Web Server in Your Organization**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Event Viewer**.

3. In the details pane, double-click **Events.asp URL**, click **Enabled**, and then type in the URL for the Web page that you want Event Viewer links to go to. Click **OK**.

4. In the details pane, double-click **Events.asp program**, click **Enabled**, and then type the path for the program that should be used for displaying the URL that you typed in the previous step. If you want the page to be displayed in the Web browser and the Web browser is in the system path, you can type the name of the Web browser executable alone, for example, **iexplore.exe**.

5. In the details pane, double-click **Events.asp program command line parameters**, click **Enabled**, and then type any command line parameters required for the program you typed in the previous step. If the program you typed in the previous step does not use parameters, clear the text box.

| **Note** |
| --- |
| Even after the preceding settings go into effect, when a person clicks a link in Event Viewer, the notification still appears, stating that Event Viewer will send information across the Internet and asking for confirmation. Regardless of the notification, if you carry out the preceding procedure and redirect events to a Web server in your organization, the information goes to that server, not across the Internet. |

# File Association Web Service and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of the file association Web service
- How the file association Web service communicates with sites on the Internet
- How to control the file association Web service to limit the flow of information to and from the Internet

**Benefits and Purposes of the File Association Web Service**

The file association Web service in Microsoft Windows Server 2003 extends the scope of information stored locally by the operating system about file name extensions, file types, and the applications or components to use when opening a particular file type. Both the locally stored information and the file association Web service are intended to provide you with the ability to open (double-click) a file without having to specify which application or component to open it with. The operating system associates the file name extension (for example, .txt or .jpg) with a file type, and it opens each file type with the application or component specified for that file type. For example, file name extensions .htm and .html are both "HTML Document" file types.

The operating system first checks for the file association information locally. If no local information is available about the file name extension and its associated file type, the operating system offers you the option of looking for more information on a Microsoft Web site. For details about the URL for this Web site, see "How the File Association Web Service Communicates with Sites on the Internet," later in this section.

**Overview: Using the File Association Web Service in a Managed Environment**

To limit the flow of information from the file association Web service to the Internet, you have several options. You can use firewall settings, you can disable the file association Web service using Group Policy, you can configure automatic server-based software installation through Group Policy, and you can train those who work on servers so that they understand how to specify an association between a file type and the application or component that is used to open that file type. You can also use scripts to limit the file types that can be stored, viewed, or used on computers in your organization, which will limit the likelihood that anyone will need to obtain information about those file types.

**How the File Association Web Service Communicates with Sites on the Internet**

The file association Web service communicates with sites on the Internet as follows:

- **Specific information sent or received**: If the operating system does not find local information about a file name extension, it offers you the option of sending a query to look for more information on a Microsoft Web site. The site is language-specific. The file name extension that you double-click is appended to the query. The query takes the following form:

  **http://shell.windows.com/fileassoc/***nnnn***/xml/redir.asp?Ext=***AAA*

  where *nnnn* is a hexadecimal value used in Windows Server 2003 to map to a language identifier (that is, to an RFC1766 identifier), and *AAA* is the file name extension for which information is needed. An example of a hexadecimal value and its corresponding language identifier is 0409 for en-us, English (United States).

  For more information about these hexadecimal values, see information about the multiple language (MLang) registry settings on the MSDN Web site at:

  http://go.microsoft.com/fwlink/?linkid=29165

  To search for information about MLang registry settings or the Microsoft Internet Explorer Multiple Language application programming interface (MLang API), use the Search tool on the MSDN Web site at:

  http://msdn.microsoft.com/

- **Default setting and ability to disable**: The service is enabled by default. It can be disabled by using Group Policy, as described in "Disabling the File Association Web Service," later in this section.

There are ways of reducing the likelihood that a person will trigger the file association Web service. One basic way is to configure automatic, server-based software installation based on Group Policy settings. For more information, see article 816102, "HOW TO: Use Group Policy to Remotely Install Software in Windows Server 2003," in the Microsoft Knowledge Base at:

http://go.microsoft.com/fwlink/?linkid=29166

- **Trigger and user notification**: When you try to open a file (for example, by double-clicking the file), and there is no local information about the correct application or component to use when opening the file, the operating system offers the option either to "Use the Web service to find the appropriate program" or to "Select the program from a list."

- **Logging**: No events are logged by the file association Web service.

- **Encryption, storage, and privacy**: The file name extension sent in a query to the Internet is not encrypted. Nothing in the query identifies the person who triggered the query. If the local computer's browser is configured to store information about recently visited Internet sites, the browser will store the query containing the file name extension. Otherwise, the query containing the file name extension is not stored anywhere.

- **Transmission protocol and port**: The transmission protocol is HTTP and the port is 80.

**Controlling the File Association Web Service to Limit the Flow of Information to and from the Internet**

If you want to limit the flow of information from the file association Web service to the Internet, you can use one or more of the following methods:

- Use your firewall to block access to any Web site that contains the following string:

  http://shell.windows.com/fileassoc/

- Disable the file association Web service by using Group Policy, as described in "Disabling the File Association Web Service," later in this section.

- Configure automatic, server-based software installation. To do this, configure one or more servers with the Software Installation extension of Group Policy in Windows Server 2003. When you do this, if someone tries to open a file for which the corresponding application is not installed locally, a copy of the application (stored on a server) is installed automatically. In this situation, the file association Web service will not be triggered. For more information, see article 816102, "HOW TO: Use Group Policy to Remotely Install Software in Windows Server 2003," in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?linkid=29166

- Train those who work on servers to work with file associations as follows:

  - Instruct them that an association exists (stored by the local operating system) between a file name extension, a file type, and the application or component that is used to open that file type.

  - Provide those who work on servers with information about the file name extensions for the files they need to work with most often, the file type for each extension, and the application that should be used to open each file type. For example, file name extensions .htm and .html are both "HTML Document" file types.

  - Show those who work on servers how to use Control Panel, **Folder Options**, and the **File Type** tab in **Folder Options** to associate a file name extension with a file type, and a file type with an application. Explain to them that the operating system stores this information on the local computer.

  - Instruct those who work on servers to always click **Select the program from a list** if they see a message box offering the two options, **Use the Web service to find the appropriate program** or **Select the program from a list**.

- Use scripts to scan your organization's computers for file types that you do not want to be stored, viewed, or used. Take actions to ensure that these files do not remain on individual computers' hard disks. If unwanted file types do not exist on the hard disks, it decreases the need for anyone to obtain information about the file name extension used for that file type.

**Procedures that Limit Internet Communication Initiated by the File Association Web Service**

This section contains the following information:

- A procedure for disabling the file association Web service by using Group Policy.
- Procedures that can be used as a basis for training those who work on servers about file name extensions, file types, and the application or component that the operating system uses when opening a specific file type.

**Disabling the File Association Web Service**

The following procedure explains how to disable the file association Web service by using Group Policy.

**To Disable the File Association Web Service by Using Group Policy**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.
2. If you want the policy setting to apply to all users of a computer and to come into effect when the computer starts or when Group Policy is refreshed, click **Computer Configuration**. If you want the policy setting to apply to users and to come into effect when users log on or when Group Policy is refreshed, click **User Configuration**.
3. Click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.
4. In the details pane, double-click **Turn off Internet File Association service**, and then click **Enabled**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the **Restrict Internet communication** policy setting, which is located in Computer Configuration/Administrative Templates/System/Internet Communication Management or in User Configuration/Administrative Templates/System/Internet Communication Management. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

**Specifying Associations Between File Name Extensions, File Types, and Applications or Components**

You can use the following procedures as a basis for training those who work on servers about file name extensions, file types, and the application or component that the operating system uses when opening a specific file type.

**To Associate a File Name Extension with a File Type**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Folder Options**, and then click the **File Types** tab.
3. Click **New**.
4. Type a new or existing file name extension, and then click **Advanced**.
5. In **Associated File Type**, do one of the following:
    - Type or select New to create a file type to associate with the file name extension.

- Type or select a different file type to associate with the extension.

| **Note** |
| --- |
| When you type a file name extension in the **Create New Extension** dialog box, the Associated File Type list displays the file type that is associated with that extension. To select **New**, scroll to the top of the list. |

**To Associate a File Name Extension with an Application**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Folder Options**, and then click the **File Types** tab.
3. Under **Registered file types**, click a file type.
4. Click **Change**, and then choose the application you want to use to open this file.

**Related Links**

For more information about automatic server-based software installation based on Group Policy settings, see article 816102, "HOW TO: Use Group Policy to Remotely Install Software in Windows Server 2003," in the Microsoft Knowledge Base at:

http://go.microsoft.com/fwlink/?linkid=29166

# Help and Support Center: Features That Communicate with the Internet

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of the Headlines and Online Search features in Help and Support Center
- How the Headlines and Online Search features communicate with sites on the Internet
- How to control the Headlines and Online Search features to limit the flow of information to and from the Internet

**Benefits and Purposes of Headlines and Online Search**
Help and Support Center is a self-help portal that was first included in Microsoft Windows Millennium Edition. It is also included in all versions of Windows Server 2003, including Windows Server 2003 with Service Pack 1 (SP1). You can access Help and Support Center in a number of ways, including:

- Selecting **Help and Support** from the Start menu.
- Selecting **Help and Support Center** from the Help menus for Control Panel, Windows Explorer, My Network Places, My Computer, or My Documents.

**Headlines**
A useful feature of Help and Support Center is the Headlines area. This area is typically titled "Top Issues" and is usually located in the lower-right corner of the main window, unless the window has been customized by the OEM or modified for certain languages. The Headlines area might also contain a hyperlink to additional content that can be viewed in Help and Support Center. Headlines provides a dynamic source of content that you can visit to find help and support on current issues as well as those that were known at the time the operating system was released. For example, it may display links to topics about new security bulletins, software updates, or new Help content.

**Online Search**
Online Search, another useful feature of Help and Support Center, enables users to query online Web sites automatically when performing a search. By default, the Microsoft Knowledge Base is designated as one of the Web sites for online searches. OEMs often customize the Online Search feature by, for example, adding a check box to the search window to enable the search engine to query their OEM-specific Web sites for results. To produce the most informative results when querying the Microsoft Knowledge Base, certain information such as the version of the product installed is collected from the user's computer and uploaded to the servers hosting the Microsoft Knowledge Base.

**Overview: Using Headlines and Online Search in a Managed Environment**
By using Group Policy or using the Windows user interface, you can control the extent to which the Headlines and Online Search features access the Internet. More details on the methods and procedures for controlling these features are described in the following subsections.

**How Headlines and Online Search Communicate with Sites on the Internet**
**Headlines**
The Headlines ("Top Issues") area is updated only when there is Internet connectivity. You are not required or prompted to connect to the Internet. Help and Support Center uses information contained in the NewsSet.xml file (stored in the *systemroot*\PCHealth\Helpctr\Config folder) to determine:

- Whether or not to update the Headlines area
- How frequently to update the Headlines area
- Where on the Internet to obtain the Headlines updates

This subsection summarizes the communication process:

- **Specific information sent or received**: If there is Internet connectivity, when you start Help and Support Center, the Help and Support service (helpsvc) compares the current date to the date specified by the TIMESTAMP attribute in the NewsSet.xml file and calculates the total number of days that have elapsed since the last time Headlines was successfully updated.

  Then, if the number of elapsed days is greater than the number of days specified by the FREQUENCY attribute in NewsSet.xml, the Help and Support service connects to the Web site specified by the URL attribute and downloads an updated version of the file NewsVer.xml to the *systemroot*\PCHealth\Helpctr\Config\News folder. The administrator (or other person at the server) is not uniquely identified.

  | **Note** |
  | --- |
  | For Headlines supplied by Microsoft, the URL attribute in NewsSet.xml is: **http://go.microsoft.com/fwlink/?LinkID=11**. This redirects to a different site such as: http://windows.microsoft.com/windowsxp/newsver.xml |

- The downloaded NewsVer.xml file contains links to the news content files (news blocks) for Windows Server 2003 operating systems and the installed language. These news blocks contain the information used to update the Headlines area—that is, links to and descriptions of the latest information from Help and Support Center, Windows, or support-related articles posted on Microsoft Web sites.

  | **Note** |
  | --- |
  | If the OEM has customized the Headlines feature, then the OEM-supplied Headlines may have links to the OEM's Web site. |

- If there is no Internet connectivity, Help and Support Center displays an offline message in the Headlines area similar to the following:

  "When you are connected to the Internet, this area will display links to timely help and support information. If you want to connect to the Internet now, start the New Connection Wizard and see how to establish a Web connection through an Internet service provider."

- **Default and recommended settings**: The Headlines feature is enabled by default. Recommended settings are described in the next subsection, "Controlling Headlines and Online Search to Limit the Flow of Information to and from the Internet."

- **Triggers**: The Headlines feature is automatically triggered if there is Internet connectivity when you start Help and Support Center.

- **User notification**: You are not given the choice to select whether or not to update the Headlines area before an update is performed. An "Updating ..." status indicator is displayed in the Headlines area, however, to indicate when an update is being performed. Once Help and Support Center has completed checking for new headlines, the Headlines area is labeled "Updated: date," where date is the current date.

- **Logging**: No information related to Headlines is entered into the event log.

- **Encryption**: The data transferred to Microsoft is not encrypted.

- **Access**: The only data generated on servers at Microsoft from the process of updates to the Headlines area is a single number telling how many times a connection has been made, by any computer, to the link that supports Headlines updates. No computer is identified in the process of a Headlines update. The data can be viewed by the Microsoft group that provides support for the link through which Headlines is updated.

- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is 80.

- **Ability to disable**: You can disable Headlines by using Group Policy. For more information, see "Procedure for Disabling Headlines and Online Search," later in this section.

**Online Search**

Online Search can only query online Web sites like the Microsoft Knowledge Base when there is Internet connectivity; you are neither required nor prompted to connect to the Internet. When you perform a search in Help and Support Center, if you have set the search options to search the Microsoft Knowledge Base or an OEM-designated Web site, the search engine automatically searches the specified site.

This subsection summarizes the communication process:

- **Specific information sent or received**: To produce relevant results when querying the Microsoft Knowledge Base, certain information is collected from the computer and uploaded to a server at Microsoft that hosts the Microsoft Knowledge Base. The administrator (or other person at the server) is not uniquely identified. Following is a list of the information collected:

    - The search text string that was entered
    - The language code of the operating system
    - The product Knowledge Base to be searched (for example, the Windows Server 2003 family or Outlook)
    - The version of the operating system installed (for example, Standard Edition, Enterprise Edition, Datacenter Edition, or Web Edition)
    - The number of results specified for the result set
    - Titles field status (indicates whether or not to search the article title only)
    - Type field status (indicates whether to search using "all" or "any" of the search string)

- **Default and recommended settings**: Online Search is enabled by default. Recommended settings are described in the next subsection, "Controlling Headlines and Online Search to Limit the Flow of Information to and from the Internet."

- **Triggers**: Online Search is automatically triggered if you use the default search options or if you have set the search options to encompass searches on the Internet. (Online Search is also dependent on having Internet connectivity when the search is performed.)

- **User notification**: You are not notified before Help and Support Center performs Online Search. A permanent headline is provided in the Headlines area that tells how to set Online Search options, including how to turn the feature off.

- **Logging**: No information related to Online Search is entered into the event log.

- **Encryption**: The data transferred to Microsoft is not encrypted.

- **Access**: The data uploaded to the server at Microsoft is aggregated and clustered. Information about the most common queries is later made available to the Windows Product Support Services and Windows User Assistance teams to help in developing new content or in revising existing content.

- **Privacy**: Microsoft does not retrieve personal information, such as user name or e-mail address, from the computer during an online search. A permanent headline is provided in the Headlines area that tells how to set Online Search options, including how to turn the feature off.

- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is 80.

- **Ability to disable**: You can disable Online Search by using Group Policy or through the Help and Support Center user interface.

**Controlling Headlines and Online Search to Limit the Flow of Information to and from the Internet**

The following tables describe Group Policy settings and other configuration options for both Headlines and Online Search. For more information, see "Procedures for Disabling Headlines and Online Search," later in this section.

| **Headlines: Configuration Tool** | **Setting** | **Result** |
|---|---|---|
| Group Policy | In Computer Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings, enable **Turn off Help and Support Center "Did you know?"** content. | Help and Support Center no longer retrieves text for the Headlines area and instead replaces it with white space. |

| **Online Search: Configuration Tool** | **Setting** | **Result** |
|---|---|---|
| Help and Support Center user interface (Set Search Options pane) | Clear any check boxes for querying the Microsoft Knowledge Base or OEM Web sites for results. | Disables online searches. The search results window neither displays an area for the online Web sites (such as the Microsoft Knowledge Base) nor returns any search results from online Web sites. |
| Group Policy | In Computer Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings, enable **Turn off Help and Support Center Microsoft Knowledge Base search**. | Removes the Knowledge Base section from the Help and Support Center Set search options page. Searches for information from Help topics are performed locally only. |

| **Headlines and Online Search: Configuration Tool** | **Setting** | **Result** |
|---|---|---|
| Firewall | Block HTTP port 80. | Displays offline message text in Headlines area. The search results window displays an area for the online Web sites (such as the Microsoft Knowledge Base), but it does not return any search results from online Web sites. |

**How Controlling Headlines and Online Search Can Affect Users and Applications**

The Headlines area provides a good way to obtain up-to-date solutions to common problems, updated self-help content, and information about software and driver updates. If you decide to disable the Headlines feature, the Headlines ("Top Issues") area in the Help and Support Center user interface will be blank and links to new content or software update notifications will never be presented.

The Online Search feature enables you to obtain help from online Web sites. If you decide to disable the Online Search feature, you will only be able to query local Help content.

Disabling Headlines and Online Search will not affect any other applications.

**Procedures for Disabling Headlines and Online Search**

The following procedures describe how to:

- Disable the Headlines feature by using Group Policy
- Disable the Online Search feature on individual computers
- Disable the Online Search feature by using Group Policy

**To Disable the Headlines Feature by Using Group Policy**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click Computer Configuration, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.

3. In the details pane, double-click **Turn off Help and Support Center "Did you know?" content**, and then click **Enabled**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the Restrict Internet communication policy setting, which is located in Computer Configuration/Administrative Templates/System/Internet Communication Management. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

**To Disable the Online Search Feature on Individual Computers**

1. Click **Start**, and then click **Help** and **Support**.

2. Below the Search box, click **Set search options**.

3. Clear the **Microsoft Knowledge Base** check box and any other check boxes below it. For example, the manufacturer of your computer may have added a check box to enable your search to query their Web site for results.

4. Close Help and Support Center.

**To Disable the Online Search Feature by Using Group Policy**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.

3. In the details pane, double-click **Turn off Help and Support Center Microsoft Knowledge Base search**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the Restrict Internet communication policy setting, which is located in **Computer Configuration/Administrative Templates/System/Internet Communication Management**. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

# HyperTerminal and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of HyperTerminal
- How HyperTerminal communicates with sites on the Internet
- How to control HyperTerminal to prevent the flow of information to and from the Internet

**Benefits and Purposes of HyperTerminal**

HyperTerminal is a program that you can use to connect to other computers, Telnet sites, bulletin board systems (BBSs), online services, and host computers. HyperTerminal connections are made using a modem, a null modem cable (used to emulate modem communication), or an Ethernet connection.

HyperTerminal is not installed by default when a product in the Windows Server 2003 family is installed manually. HyperTerminal is installed by default, however, if the Windows Server 2003 product is installed using an answer file during automated installation.

HyperTerminal has capabilities beyond making connections to other computers. It can, for example, transfer large files from a computer onto your portable computer using a serial port rather than requiring you to set up your portable computer on a network. It can help debug source code from a remote terminal. It can also communicate with many older, character-based computers.

HyperTerminal records the messages passed to and from the computer or service on the other end of your connection. It can therefore serve as a valuable troubleshooting tool when setting up and using your modem. To make sure that your modem is connected properly or to view your modem's settings, you can send commands through HyperTerminal and check the results. HyperTerminal also has scroll functionality that enables you to view received text that has scrolled off the screen.

| **Note** |
| --- |
| HyperTerminal is designed to be an easy-to-use tool yet it is not meant to replace other full-featured tools. You can use HyperTerminal as described in this subsection, but you should not attempt to use HyperTerminal for more complex communication. For more information about what HyperTerminal does and does not support, see the Help topics at http://go.microsoft.com/fwlink/?linkid=29890. |

**Overview: Using HyperTerminal in a Managed Environment**

HyperTerminal is not installed by default when a product in the Windows Server 2003 family is installed manually. HyperTerminal is installed by default, however, if the Windows Server 2003 product is installed using an answer file during automated installation. In a managed environment, you might choose to prevent Hyperterminal from being installed on servers running Windows Server 2003. If you choose to allow Hyperterminal to be installed on servers, when deciding how to configure it, consider the following security issues:

- **Viruses**: Incoming files might contain viruses or malicious programs that could compromise or destroy data on your computer. To reduce this risk, use virus-scanning software and ensure that incoming files are from a reliable and trusted source.
- **ID and password**: HyperTerminal cannot automatically provide your login ID and password when you make a connection. If you provide a password when using HyperTerminal for a Telnet session, be aware that this password will be sent to the remote computer using plaintext (as with all Telnet connections).
- **Automatic download**: The automatic download feature of the Zmodem protocol can pose a security risk by allowing remote users to send files to your computer without your explicit permission. To avoid this risk, you should select a protocol other than Zmodem in the **Receive File** dialog box or you should clear the **Allow remote host initiated file transfers** check box on the **Settings** tab of Connection Properties.

Complete information about concepts and procedures associated with using or configuring HyperTerminal is beyond the scope of this white paper. For more information, access the HyperTerminal Help documentation in Help and Support Center on any computer running Windows Server 2003.

**How HyperTerminal Communicates with Sites on the Internet**

The exchange of information that takes place during the HyperTerminal connection is as follows:

- **Specific information sent or received**: When using HyperTerminal for Telnet connectivity, the user ID and password are sent in plaintext format (as with all Telnet connections). If files are being transmitted, only the protocol, terminal emulation data, and file-specific binaries are sent. The computer running HyperTerminal is identified by its IP address when the connection type is TCP/IP. The computer is not uniquely identified when the connection type is not TCP/IP.

- **Default settings**: HyperTerminal is not installed by default when a product in the Windows Server 2003 family is installed manually. HyperTerminal is installed by default, however, if the Windows Server 2003 product is installed using an answer file during automated installation. To remove or uninstall HyperTerminal, see "Controlling HyperTerminal to Prevent the Flow of Information to and from the Internet," later in this section.

- **Triggers**: When HyperTerminal is set to automatically answer incoming connections, a file transfer can be initiated if the Zmodem transfer protocol is used. Otherwise, communication through HyperTerminal is only triggered when the user deliberately initiates it.

- **User notification**: After a user starts a HyperTerminal connection session, the status of the connection that is currently open within HyperTerminal is displayed in the HyperTerminal title bar. The status of the file and text transfer is displayed in the HyperTerminal window during the transfer process. HyperTerminal does not display connection or transfer status information when the automatic download feature of the Zmodem protocol is used. For more information about the HyperTerminal automatic download feature, see "Overview: Using HyperTerminal in a Managed Environment," earlier in this section.

- **Encryption**: Information sent or received by HyperTerminal is not encrypted.

- **Transmission protocol and port**: The protocols used are Kermit, Xmodem, Xmodem-1K, Ymodem, Ymodem-G, and Zmodem transmissions protocols on port 23.

- **Ability to disable**: You can prevent Hyperterminal from being installed, remove it, or prevent it from being run. For more information, see "Controlling HyperTerminal to Prevent the Flow of Information to and from the Internet," later in this section.

**Controlling HyperTerminal to Prevent the Flow of Information to and from the Internet**

You can ensure that HyperTerminal cannot be used on a server running Windows Server 2003 by using the following methods, described in the procedures in this section:

- For unattended installations, specify an entry in an answer file that will prevent HyperTerminal from being installed.

- After installation, use Add or Remove Programs in Control Panel to remove HyperTerminal or to confirm that it is not installed.

- Use Group Policy to ensure that HyperTerminal cannot be run.

**To Ensure that HyperTerminal is Not Installed During Unattended Installation**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file. For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment.

2. In the [Components] section of the answer file, include the following entry:

   **hypertrm = Off**

**To Remove HyperTerminal from an Individual Computer Running Windows Server 2003**

1. Click **Start**, and then either point to **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components** (on the left).
4. Double-click **Accessories and Utilities**, and then double-click **Communications**.
5. Make sure the check box for the **HyperTerminal** component is cleared.
6. Follow the instructions to complete the Windows Components Wizard.

| **Note** |
| --- |
| You must have administrative credentials to complete this procedure. |

**To Use Group Policy to Prevent HyperTerminal from Being Run**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.
2. In Group Policy, click **User Configuration**, click **Administrative Templates**, and then click **System**.
3. In the details pane, double-click **Don't run specified Windows applications**.
4. Click **Enabled**, click **Show**, click **Add**, and then type the executable name:

   **Hypertrm.exe**

# Internet Explorer 6.0 and Resulting Internet Communication in Windows Server 2003 with Service Pack 1

Applies To: Windows Server 2003 with SP1

This white paper provides information about the communication that flows between components in Windows Server 2003 with SP1 and sites on the Internet, and it describes steps to take to limit, control, or prevent that communication in an organization with many users.

This section of the white paper provides information about:

- The benefits of Microsoft Internet Explorer 6 in Windows Server 2003.
- A description of Internet Explorer Enhanced Security Configuration, which is enabled by default when you install a product in the Windows Server 2003 family.
- Examples of the security-related features provided in Internet Explorer 6 (as compared to Internet Explorer 5).
- Procedures for working with security-related settings in Internet Explorer.
- Resources for learning about topics related to security in Internet Explorer 6. This includes resources that help you learn about:
    - The Internet Explorer Enhanced Security Configuration
    - Security and privacy settings in Internet Explorer 6.
    - Mitigating the risks inherent in Web-based applications and scripts.
    - Methods for deploying specific configurations of Internet Explorer 6 across your organization using Group Policy, the Internet Explorer Administration Kit (IEAK), or both.

| **Note** |
| --- |
| This section of the white paper describes Internet Explorer 6, but does not describe the related components Outlook Express 6 (the e-mail component in Windows Server 2003), the New Connection Wizard, or the error reporting tool in Internet Explorer. For information about these components, see the respective sections of this white paper (the error reporting tool in Internet Explorer is described in the Windows Error Reporting and Internet Communication section of this white paper). |

| **Note** |
| --- |
| The New Connection Wizard replaces the Network Connection Wizard and the Internet Connection Wizard in Windows 2000. |

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization where users connect to Web sites, run software from the Internet, download items from the Internet, and perform similar actions. This section, however, provides overview information as well as suggestions for other sources of information about how to balance users' requirements for Internet access with your organization's requirements for protection of networked assets.

The following are central locations in which to find information about Internet Explorer 6:

- Help for Internet Explorer (with Internet Explorer open, click the Help menu and select an appropriate option)
- The Internet Explorer page on the Microsoft Web site at:

http://www.microsoft.com/windows/ie/

- The Resource Kit for Internet Explorer. To learn about this and other Resource Kits, see the Microsoft TechNet Web site at:

  http://go.microsoft.com/fwlink/?linkid=29894

**Benefits and Purposes of Internet Explorer 6**

Internet Explorer 6 in Windows Server 2003 with SP1 is designed to make it easy to browse and interact with sites on an intranet or on the Internet. It differs from most of the other components described in this white paper in that its main function is to communicate with sites on the Internet or an intranet (which contrasts with components that communicate with the Internet in the process of supporting some other activity).

Internet Explorer 6 is also designed to be highly configurable, with security and privacy settings that can help protect your organization's networked assets while at the same time providing access to useful information and tools. In addition, Internet Explorer Enhanced Security Configuration, which is enabled by default when you install a product in the Windows Server 2003 family, helps make your computer more secure by limiting its exposure to malicious Web sites.

With this enhanced level of security, however, you might find that some Web sites are not displayed correctly in Internet Explorer when you are browsing from a server. Also, you might be prompted to enter your credentials when accessing network resources, such as files in shared folders with Universal Naming Convention (UNC) names. You can easily change the enhanced security settings.

If you want to establish a specific configuration on servers (instead of using Internet Explorer Enhanced Security Configuration), Internet Explorer 6 offers more security-related options and settings than were available in Internet Explorer 5. The subsections that follow provide more information about Internet Explorer Enhanced Security Configuration and about the security-related options and settings in Internet Explorer 6.

**Internet Explorer Enhanced Security Configuration**

Internet Explorer Enhanced Security Configuration is enabled by default when you install Windows Server 2003. With this configuration, each zone uses a higher security setting than was used by default in Windows 2000. You can easily change the enhanced security settings.

The following table outlines some of the differences that Internet Explorer Enhanced Security Configuration makes in security settings on a server. (For a description of zones, see "Examples of Security-Related Features Provided in Internet Explorer 6," later in this section.)

**Security settings with and without Internet Explorer Enhanced Security Configuration**

| Zone | With Internet Explorer Enhanced Security Configuration | Without Internet Explorer Enhanced Security Configuration (the same security levels as Windows 2000) |
|---|---|---|
| Internet zone | High security settings | Medium security settings |
| Trusted sites zone | Medium security settings | Low security settings |
| Local intranet zone | Medium-low security settings (intranet sites are not automatically detected) | Medium-low security settings (intranet sites are automatically detected) |

Also, with Internet Explorer Enhanced Security Configuration, several sites are added automatically to specific zones:

- The Windows Update Web site is added to the Trusted sites zone. This allows you to continue to get important updates for your operating system. For more information about Windows Update, see the Windows Update, Automatic Updates, and Internet Communication section of this white paper.

- The Windows Error Reporting site is added to the Trusted sites zone. This allows you to report problems you encounter with your operating system and search for fixes. For more information about Windows Error Reporting, see the Windows Error Reporting and Internet Communication section of this white paper.

- Several local computer sites (for example, http://localhost, https://localhost, hcp://system) are added to the Local intranet zone. This allows applications and code to work locally so that you can complete common administrative tasks.

You can enable or disable the Internet Explorer Enhanced Security Configuration for administrators, all other user groups, or both. For more information, see "To Remove Internet Explorer Enhanced Security Configuration and Restore the Default Internet Explorer 6 Security Settings," later in this section.

For more information about Internet Explorer Enhanced Security Configuration, see the resources listed in "Learning About Internet Explorer Enhanced Security Configuration," later in this section.

**Examples of the Security-Related Features Provided in Internet Explorer 6**
This subsection describes enhancements in some of the security-related features in Internet Explorer 6, as compared to Internet Explorer 5. These features include:

- A **Privacy** tab that provides greater flexibility in specifying whether cookies will be blocked from specific sites or types of sites. An example of a type of site that could be blocked is one that does not have a compact policy—that is, a condensed computer-readable privacy statement. (The **Privacy** tab was not available in Internet Explorer 5.)

- Security settings that specify how Internet Explorer 6 handles such higher-risk items as ActiveX controls, downloads, and scripts. These settings can be customized as needed, or they can be set to these predefined levels: high, medium, medium-low, or low. You can specify different settings for a number of zones, the most basic being the four preconfigured zones:

  - Local intranet zone: Contains addresses inside the boundary defined by your proxy server or firewall.

  - Trusted sites: Includes sites you designate as "trusted."

  - Restricted sites: Includes sites you designate as "restricted."

  - Internet zone: Includes everything that is not in another zone and is not on the local computer.

  You can also specify different settings for the customized zones that you add programmatically using the URL security zones application programming interface (API). For more information, search for "URL security zones" on the MSDN Web site at:

  http://msdn.microsoft.com/

- Support for content-restricted IFrames (inline floating frames). This type of support enables developers to implement these frames in a way that makes it more difficult for malicious authors to start e-mail-based or content-based attacks.

- Improvements in Windows Server 2003 Service Pack 1 (SP1) that increase the overall security and reliability of Internet Explorer 6. These improvements include a configurable pop-up blocker, an interface from which you can manage add-ons (programs that extend the capabilities of the browser), and enhancements to other security features.

For more information about features available in Internet Explorer, see "Resources for Learning About Topics Related to Security in Internet Explorer 6," later in this section, as well as the Internet Explorer page on the Microsoft Web site at:

http://www.microsoft.com/windows/ie/

**Procedures for Working with Security-Related Settings in Internet Explorer**
This subsection describes how to carry out the following:

- View security settings for zones in Internet Explorer

- Locate Group Policy settings that affect Internet Explorer, and view related Help

- Determine whether Internet Explorer Enhanced Security Configuration is enabled on a specific server

- Remove Internet Explorer Enhanced Security Configuration and restore the default Internet Explorer 6.0 security settings

**To View Security Settings for Zones in Internet Explorer**

1. On the server on which you want to view settings, start Internet Explorer by your preferred method, for example, by clicking the Internet Explorer icon on the taskbar.

2. On the **Tools** menu, click **Internet Options**.

3. Click the **Security** tab.

4. Select the zone for which you want to view security settings:

   - Internet
   - Local intranet
   - Trusted sites
   - Restricted sites

**To Locate Group Policy Settings that Affect Internet Explorer**

1. See Appendix B: Resources for Learning About Group Policy for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Internet Explorer**.

   View the available settings.

3. Click **User Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Internet Explorer**.

   View the available settings.

**To Determine Whether Internet Explorer Enhanced Security Configuration is Enabled on a Specific Server**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Add or Remove Programs**.

3. Click **Add/Remove Windows Components** (on the left).

4. Scroll down to **Internet Explorer Enhanced Security Configuration**. If the check box is selected, it is enabled. If the check box is cleared, it is not enabled.

5. If you want to see whether Internet Explorer Enhanced Security Configuration is enabled for administrator groups, all other user groups, or both, select **Internet Explorer Enhanced Security Configuration**, and then click **Details**.

**To Remove Internet Explorer Enhanced Security Configuration and Restore the Default Internet Explorer 6 Security Settings**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Add or Remove Programs**.

3. Click **Add/Remove Windows Components** (on the left).

4. Click **Internet Explorer Enhanced Security Configuration**, and then do one of the following:

   - To remove Internet Explorer Enhanced Security Configuration for both administrators and all other users, clear the **Internet Explorer Enhanced Security Configuration** check box, and then click **Next**.

   - To remove Internet Explorer Enhanced Security Configuration for administrators only or for users who are not in an administrator group, click **Details**, clear either the **For administrator groups** check box or the **For all other user groups** check box, and then click **Next**.

5. Follow the instructions to complete the Windows Components Wizard.

**Procedures for Setting the Security Level to High for Specific Web Sites**

The procedures that follow provide information about how to set the security level for a particular Web site to **High**, which prevents actions such as the running of scripts and the downloading of files from the site. For related information, see "Examples of the Security-Related Features Provided in Internet Explorer 6," earlier in this section, and "Learning About Security and Privacy Settings in Internet Explorer 6," later in this section.

**To Configure a Specific Computer with a Security Level of High for Specific Sites**

1. On the computer on which you want to configure a security level of High for specific sites, in Internet Explorer, click **Tools**, click **Internet Options**, and then click the **Security** tab.

2. Select **Restricted sites**.

3. Under **Security level for this zone**, make sure the slider for the security level is set to **High**. If the security level for the zone is **Custom** and you prefer to set it to **High**, click **Default Level** and make sure the slider for the security level is set to **High**.

   You can view the individual settings that make up **High** security by clicking **Custom Level**. For example, you can click **Custom Level** and then scroll down to confirm that for **High** security, the settings for active scripting and for file download are both **Disable**. After viewing the settings, click **Cancel**.

4. With **Restricted sites** still selected, click **Sites**.

5. In **Add this Web site to the zone**, type the Web site address. You can use an asterisk for a wildcard. For example, for Web sites at Example.Example.com and www.Example.com, you could type:

   **http://*.Example.com**

6. Click the **Add** button.

**To Use Group Policy to Set the Security Level to High for Specific Sites That Users or Administrators in Your Organization Might Connect To**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.

2. In Group Policy, click **User Configuration**, click **Windows Settings**, click **Internet Explorer Maintenance**, and then click **Security**.

3. In the details pane, double-click **Security Zones and Content Ratings**.

4. Under **Security Zones and Privacy**, click **Import the current security zones and privacy settings**, and then click **Modify Settings**.

5. Select **Restricted sites**.

6. Under **Security level for this zone**, make sure the slider for the security level is set to **High**. If the security level for the zone is **Custom** and you prefer to set it to **High**, click **Default Level** and make sure the slider for the security level is set to **High**.

   You can view the individual settings that make up **High** security by clicking **Custom Level**. For example, you can click **Custom Level** and then scroll down to confirm that for **High** security, the settings for file download and for active scripting are both **Disable**. After viewing the settings, click **Cancel**.

7. With **Restricted sites** still selected, click **Sites**.

8. In **Add this Web site to the zone**, type the Web site address. You can use an asterisk for a wildcard. For example, for Web sites at Example.Example.com and www.Example.com, you could type:

   **http://*.Example.com**

9. Click the **Add** button.

#### Resources for Learning About Topics Related to Security in Internet Explorer 6

This subsection lists resources that can help you learn about the following topics related to security in Internet Explorer 6:

- Internet Explorer Enhanced Security Configuration
- Security and privacy settings available in Internet Explorer 6
- Methods for mitigating the risks inherent in Web-based programs and scripts
- Ways to use Group Policy objects that control configuration settings for Internet Explorer 6
- The Internet Explorer Administration Kit

In addition, for information about unattended installation, see the resources listed in Appendix A: Resources for Learning About Automated Installation and Deployment.

| **Note** |
| --- |
| For information about Internet Explorer on clients running Windows XP Professional with Service Pack 2, that is, for information similar to what is provided in this white paper but focused on clients instead of servers, see "Using Windows XP Professional with Service Pack 2 in a Managed Environment: Controlling Communication with the Internet" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=29133. |

#### Learning About Internet Explorer Enhanced Security Configuration

For more information about Internet Explorer Enhanced Security Configuration, see one of the following:

- The informational pages displayed in Internet Explorer after you install a product in the Windows Server 2003 family. To view these pages, start Internet Explorer after completing the installation.
- Help topics in Internet Explorer. To view these topics, with Internet Explorer Enhanced Security Configuration enabled, start Internet Explorer, click **Help**, and then click **Enhanced Security Configuration**.
- Help topics in Help and Support Center. To view these topics, click **Start**, click **Help and Support**, and search for "enhanced security configuration."

#### Learning About Security and Privacy Settings in Internet Explorer 6

Two important sources of detailed information about the security and privacy settings in the version of Internet Explorer 6 in Windows Server 2003 with SP1 are as follows:

- "Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1," which contains information about the pop-

up blocker, the interface from which you can manage add-ons (programs that extend the capabilities of the browser), and enhancements to other security features. This document is available on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=46278

- Microsoft Internet Explorer 6 Resource Kit

  To learn about this and other resource kits, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=29894

  The Microsoft Internet Explorer 6 Resource Kit consists of a number of parts that include these titles:

  - "Privacy and Security Features"
  - "Preparation for Deployment"
  - "Customization and Installation"
  - "Maintenance and Support," including information about keeping programs updated
  - Appendices, including an appendix titled "Setting System Policies and Restrictions"

**Learning About Mitigating the Risks Inherent in Web-based Applications and Scripts**
In a network-based and Internet-based environment, code can take a variety of forms including scripts within documents, scripts within e-mail messages, or applications or other code objects running within Web pages. This code can move across the Internet and is sometimes referred to as "mobile code." Configuration settings provide ways for you to control the way Internet Explorer 6 responds when someone tries to run mobile code. Two examples of the ways you can customize the Internet Explorer configuration deployed in your organization are as follows:

- You can control the code (in ActiveX controls or in scripts, for instance) that administrators or operators can run. You can do this by customizing Authenticode® settings, which can, for example, prevent administrators or operators from running any unsigned code or enable them to only run code signed by specific authors.

- If you want to permit the use of ActiveX controls, but you do not want administrators or operators to download code directly from the Internet, you can specify that when Internet Explorer 6 looks for a requested executable, it goes to your own internal Web site instead of the Internet. For more information, see the white paper titled "Managing Mobile Code with Microsoft Technologies" at the end of this list, and search for "CodeBase".

You can use the following sources to learn more about mitigating the risks inherent in Web-based applications and scripts:

- To understand more about how a particular Microsoft programming or scripting language works, see the MSDN Web site at:

  http://msdn.microsoft.com/

- To learn about approaches to mitigating the risks presented by mobile code, see "Managing Mobile Code with Microsoft Technologies," a white paper on the TechNet Web site at:

  http://go.microsoft.com/fwlink/?linkid=29170

**Learning About Group Policy Objects That Control Configuration Settings for Internet Explorer 6**
You can control configuration settings for Internet Explorer 6 by using Group Policy objects (GPOs) on servers running Windows Server 2003. (You can also control the configuration of Internet Explorer by using the Internet Explorer Administration Kit. For more information, see "Learning about the Internet Explorer Administration Kit," later in this section.) For sources of information about Group Policy, see Appendix B: Resources for Learning About Group Policy.

To learn about specific Group Policy settings that can be applied to computers running Windows Server 2003 with SP1, see the following two sources of information:

- "Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1" on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46278

- The Group Policy Settings Reference on the Microsoft Download Center Web site at:

  http://go.microsoft.com/fwlink/?linkid=29911

**Learning About the Internet Explorer Administration Kit**

With the deployment technologies available in the Internet Explorer Administration Kit (IEAK), you can efficiently deploy Internet Explorer and control the configuration of Internet Explorer across your organization. (You can also control the configuration of Internet Explorer by using Group Policy. For more information, see "Learning about Group Policy objects that control configuration settings for Internet Explorer 6," earlier in this section.)

A few of the features and resources in the IEAK include:

- **Internet Explorer Customization Wizard**. Step-by-step screens guide you through the process of creating customized browser packages that can be installed on client desktops.

- **IEAK Profile Manager**. After you deploy Internet Explorer, you can use the IEAK Profile Manager to change browser settings and restrictions automatically.

- **IEAK Toolkit**. The IEAK Toolkit contains a variety of helpful tools, programs, and sample files.

- **IEAK Help**. IEAK Help includes many conceptual and procedural topics that you can view by using the **Index**, **Contents**, and **Search** tabs. You can also print topics from IEAK Help.

For more information about the IEAK, see the Windows Web site at:

http://go.microsoft.com/fwlink/?linkid=29479

# Internet Information Services and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Internet Information Services (IIS) in Microsoft Windows Server 2003
- For servers from which you want to offer content on an intranet or the Internet, descriptions of some of the security-related features offered in IIS 6.0, and suggestions for other sources of information about security and IIS 6.0

| **Note** |
| --- |
| For servers from which you do not want to offer content on an intranet or the Internet, you do not need to remove IIS, since by default it is not installed with most products in the Windows Server 2003 family. The exception is Windows Server 2003, Web Edition, on which IIS is installed by default. If you use a server as a Web server and then deploy it for some other purpose, remove IIS from that server. |

- Controlling Internet printing
- Subcomponents that are part of IIS, with instructions for finding out which subcomponents are installed on a given server
- Viewing Help for IIS
- Other sources of information about IIS

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running servers that communicate across the Internet. This section, however, provides overview information as well as suggestions for other sources of information about balancing your organization's requirements for communication across the Internet with your organization's requirements for protection of networked assets.

**Benefits and Purposes of IIS**

IIS 6.0 is one of the optional components in products in the Windows Server 2003 family. IIS is a component that provides an easy way to publish information on the Internet or an intranet. In a managed environment, IIS is usually installed on selected servers only. IIS includes innovative security features and a broad range of administrative features for managing Web sites. By using programmatic features like Active Server Pages (ASP and ASP.NET), you can more easily create and deploy scalable, flexible Web applications.

IIS is not installed by default with products in the Windows Server 2003 family other than Windows Server 2003, Web Edition. IIS and related components can be added by using either **Add or Remove Programs** in Control Panel or Manage Your Server. After IIS 6.0 is installed, it is configured by default in a "locked down" state. The locked down state means that IIS 6.0 accepts requests for static files only, until it is configured to serve dynamic content. It also means that all time-outs and settings are set to restrictive defaults. You can enable or disable IIS 6.0 functionality based on the needs of your organization by using IIS Manager. You can also enable IIS 6.0 functionality through programmatic and command-line interfaces.

For more information about IIS features, including features related to security, see the following Web sites:

- The product information page for IIS 6.0 at:

  http://go.microsoft.com/fwlink/?LinkId=46005

- The technical overview document for IIS 6.0 at:

  http://go.microsoft.com/fwlink/?linkid=22805

- The IIS page on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=31134

If you have a Web site on which you want to use Microsoft .NET Passport for authentication and you also want to use Passport Manager Administration, a component available for Windows Server 2003, see Appendix I: Passport Manager Administration.

**Examples of Security-Related Improvements in IIS 6.0**

IIS 6.0 includes a variety of settings and features related to security, some of which are listed in the following table. For additional information about security-related improvements in the version of IIS 6.0 in Windows Server 2003 with SP1, see the links in the previous section.

**Examples of security-related settings and features in IIS 6.0**

| Setting or feature | Description |
|---|---|
| Disabling through Group Policy | With Windows Server 2003, domain administrators can prevent users from installing IIS 6.0 on their computers. |
| Running as an account with limited privileges | IIS 6.0 worker processes run in a user context with limited privileges by default. This drastically reduces the attack surface of the Web server. |
| Secure ASP | All functions built into ASP pages always run as an account with limited privileges (anonymous user). |
| Recognized file extensions | IIS 6.0 serves requests only to files that have recognized file extensions and rejects requests to file extensions it doesn't recognize. |
| Command-line tools not accessible to Web users | Attackers often take advantage of command-line tools that are executable through the Web server. In IIS 6.0, the command-line tools cannot be executed by the default Web server identity. |
| Write protection for content | Once attackers get access to a server, they try to deface Web sites. By preventing anonymous Web users from overwriting Web content, these attacks can be mitigated. |
| Time-outs and limits | Product settings are set to aggressive and secure defaults. |
| Upload data limitations | Administrators can limit the size of data that can be uploaded to a server. |
| Buffer overflow protection | The Windows Administration Service in IIS will detect if a worker process had a buffer overflow and will exit that process. |
| File verification | The core server verifies that the requested content exists before it gives the request to a request handler (Internet Server Application Programming Interface [ISAPI] extension). |

For more information about creating Web sites with IIS 6.0 and maintaining appropriate levels of awareness and control over the communication to and from those sites, see the IIS Help. For information about viewing the Help, see "To View Help After Installing IIS," later in this section.

**Controlling Internet Printing**

Internet printing makes it possible for clients to use printers located anywhere in the world by sending print jobs using Hypertext Transfer Protocol (HTTP). Additionally, a computer running Windows Server 2003 can use IIS to create a Web page that provides information about printers and provides the transport for printing over the Internet.

For Internet printing, it is important to consider both the server and the client:

- **Server**: Internet printing is an optional component (not installed by default) of IIS 6.0. A server running Windows Server 2003 can be configured to act as a print server allowing Internet printing. In a managed environment, you might want to ensure that the Internet printing subcomponent of IIS is not installed. For information about how to do this, see "Procedures for Checking or Controlling the Installation of IIS Subcomponents," later in this section.

- **Client**: Clients (typically, running Windows XP, not Windows Server 2003) can install an Internet printer using a Web browser, the Add Printer Wizard, or the Run dialog box. To control whether clients can support Internet printing, see the section about Internet printing in the white paper titled "Using Windows XP Professional with Service Pack 2 in a Managed Environment: Controlling Communication with the Internet." You can view this white paper on the TechNet Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29133

**Answer File Entries and Registry Keys for IIS Subcomponents**

For reference purposes, the following table shows the syntax for answer file entries associated with IIS in the Windows Server 2003 family as well as the corresponding registry keys. Do not change the registry keys. They are shown for use in a script that could check whether a particular component is installed on a particular server. A registry key value of 0x00000000 means the component is not installed, and a value of 0x00000001 means the component is installed.

| **Note** |
| --- |
| For more details about answer file entries related to IIS components, follow the steps in "To View Help After Installing IIS," later in this section, and then search for the topic called "Installing IIS." In that topic, look for a table showing the answer file entries. |

Answer file entries and registry keys associated with IIS subcomponents for the Windows Server 2003 family

| **IIS subcomponent** | **Syntax for answer file entry (in the [Components] section)** | **Registry key (for use in a script that checks whether a component is installed): 0x00000000 means it is not installed; 0x00000001 means it is installed** |
| --- | --- | --- |
| IIS common files | iis_common = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\ Setup\OC Manager\Subcomponents\ iis_common |
| Active Server Pages (ASP) for IIS | iis_asp = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\ Setup\OC Manager\Subcomponents\ iis_asp |
| File Transfer Protocol (FTP) service | iis_ftp = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\ Setup\OC Manager\Subcomponents\ iis_ftp |
| IIS Manager (Microsoft Management Console [MMC] snap-in) | iis_inetmgr = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\ Setup\OC Manager\Subcomponents\ iis_inetmgr |
| Internet Data Connector | iis_internetdataconnector = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\ Setup\OC Manager\Subcomponents\ iis_internetdataconnector |
| Network News Transfer Protocol (NNTP) service | iis_nntp = On \| Off | HKEY_LOCAL_MACHINE\Software\ Microsoft\Windows\CurrentVersion\ Setup\OC Manager\Subcomponents\ iis_nntp |

| | | |
|---|---|---|
| Server-Side Includes | iis_serversideincludes = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\iis_serversideincludes |
| Simple Mail Transfer Protocol (SMTP) service | iis_smtp = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\iis_smtp |
| Web Distributed Authoring and Versioning (WebDAV) publishing | iis_webdav = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\iis_webdav |
| World Wide Web (WWW) service | iis_www = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\iis_www |
| Remote administration (HTML) | sakit_web = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\sakit_web |
| Internet Server Application Programming Interface (ISAPI) for Background Intelligent Transfer Service (BITS) server extensions | BitsServerExtensionsISAPI = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\bitsserverextensionsisapi |
| Background Intelligent Transfer Service (BITS) server extensions snap-in | BitsServerExtensionsManager = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\bitsserverextensionsmanager |
| FrontPage server extensions | fp_extensions = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\fp_extensions |
| Internet printing | inetprint = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\inetprint |

| | | |
|---|---|---|
| ActiveX control and sample pages for hosting Terminal Services client connections over the Web | TSWebClient = On \| Off | HKEY_LOCAL_MACHINE\Software\<br>Microsoft\Windows\CurrentVersion\<br>Setup\OC Manager\Subcomponents\<br>TSWebClient |

**Note**

For several of the subcomponents in the previous table, the software for the subcomponent is installed regardless of the answer-file entry, but the subcomponent cannot be used unless the answer-file entry is set to On (or the procedure is followed for installing the subcomponent through **Add or Remove Programs** in Control Panel). These subcomponents are Internet Data Connector, Server-Side Includes, and WebDAV publishing.

**Procedures for Checking or Controlling the Installation of IIS Subcomponents and the IIS Lockdown Tool**

The following procedures explain how to:

- View the registry keys listed in the table in the previous subsection
- View or change the IIS components currently installed on a computer running a product in the Windows Server 2003 family
- Specify answer file entries that control whether IIS subcomponents are included during unattended installation

**To View Registry Keys Related to IIS Subcomponents**

1. Open Registry Editor by clicking **Start**, clicking **Run**, and then typing **regedit**.

   **Caution**

   Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. Navigate to:

   HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Setup\OC Manager\Subcomponents\.

3. View the registry keys listed in the table in the previous subsection, and find the value associated with each key. A value of 0x00000000 means the component is not installed. A value of 0x00000001 means the component is installed.

4. Close **Registry Editor**.

**To View or Change the IIS Components Currently Installed on a Computer Running Windows Server 2003**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components** (on the left).
4. Select **Application Server** and then click **Details**.
5. Find **Internet Information Services (IIS)** in the list, and perform one of the following steps:

   - If IIS is installed and you want to remove it, clear the check box for IIS and complete the wizard.

- If IIS is not installed and you want to add the default set of IIS subcomponents, select the check box for IIS and complete the wizard.

- If you want to view or select from the list of IIS subcomponents, after selecting IIS, click **Details**.

| **Note** |
| --- |
| The Internet Printing component is in the list of subcomponents that appears when you click **Details**. |

6. Follow the instructions to complete the Windows Components Wizard.

**To Specify Answer File Entries That Control Whether IIS Subcomponents are Included During Unattended Installation**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file.

2. In the [Components] section of the answer file, add the appropriate entries listed in the table in "Answer File Entries and Registry Keys for IIS Subcomponents," earlier in this section. Ensure that the entries specify **Off** for components you do not want to install and **On** for components you want to install.

   If no IIS subcomponents are listed in an answer file for unattended installation of a product in the Windows Server 2003 family other than Windows Server 2003, Web Edition, the IIS subcomponents are not installed by default.

| **Note** |
| --- |
| For more details about answer file entries related to IIS components, follow the steps in the next procedure, "To View Help After Installing IIS," and then search for the Help topic called "Installing IIS." In that topic, look for a table showing the answer file entries. |

**To View Help After Installing IIS**

1. After installing IIS (including the IIS Manager subcomponent, which is included in default installations of IIS), click **Start**.

2. Either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

3. Double-click **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.

4. Click the **Help** menu and then click **Help Topics**.

**To Obtain the IIS Lockdown Tool**

- The IIS Lockdown Tool, designed for use on computers on which IIS is installed, is available from the Microsoft TechNet Web site at:

  http://go.microsoft.com/fwlink/?linkid=29896

**Related Links**

For more information about IIS, see the following Web sites:

- The product information page for IIS 6.0 at:

  http://go.microsoft.com/fwlink/?LinkId=46005

- The technical overview document for IIS 6.0 at:

  http://go.microsoft.com/fwlink/?linkid=22805

- The IIS page on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=31134

- The page for the IIS Lockdown Tool on the TechNet Web site at:

  http://go.microsoft.com/fwlink/?linkid=29896

# Internet Protocol Version 6 (IPv6) and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- An introduction to the IPv6 protocol
- The benefits of the IPv6 protocol
- How the IPv6 protocol can be used across the Internet
- How to control the IPv6 protocol to limit the flow of information to and from the Internet
- How to monitor and troubleshoot the IPv6 protocol after configuration is complete

**An Introduction to the IPv6 Protocol**

The current version of the Internet Protocol (known as IP version 4 or IPv4) has not been substantially changed since 1981, when the Internet Engineering Task Force (IETF) published the definitive specification of IP (RFC 791). IPv4 has proven to be robust, easily implemented, and interoperable. It has stood the test of scaling an internetwork to a global utility the size of today's Internet, which is a tribute to the protocol's initial design.

The initial design, however, did not anticipate the exponential growth of the Internet and the exhaustion of the IPv4 address space, or the effort required to maintain routing information. Because of the way in which IPv4 network IDs are allocated, there are routinely over 70,000 routes in the routing tables of Internet backbone routers. Most current IPv4 implementations are configured either manually or through a stateful address configuration protocol such as the Dynamic Host Configuration Protocol (DHCP). With more computers and devices using IP, there is a need for a simpler and more automatic configuration of addresses and other configuration settings that do not rely on the administration of a DHCP infrastructure.

Another factor driving the development of IPv6 is the need for improved security. Private communication over a public medium like the Internet requires encryption services that protect the data sent from being viewed or modified in transit. There is a standard for providing security for IPv4 packets (known as Internet Protocol security or IPSec). In IPv4, however, this standard is optional and proprietary solutions are prevalent. In IPv6, IPSec support is required.

To address these concerns, the IETF has developed a suite of protocols and standards known as IP version 6 (IPv6). This new version incorporates the concepts of many proposed methods for updating the IPv4 protocol.

For the latest set of RFCs and Internet drafts describing IPv6 and IPv4 coexistence and migration technologies, see the Internet Engineering Task Force (IETF) Web site at:

http://go.microsoft.com/fwlink/?LinkId=29136

(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

**Benefits and Purposes of the IPv6 Protocol**

An IPv6 address is four times as large as an IPv4 address. The global addresses used on the IPv6 portion of the Internet are designed to create an efficient, hierarchical, and summarized routing infrastructure that addresses the common occurrence of multiple levels of Internet service providers. On the IPv6 Internet, the backbone routers have an efficient and hierarchical addressing and routing infrastructure that uses smaller routing tables.

IPv6 supports both stateful address configuration (such as address configuration in the presence of a DHCP server) and stateless address configuration (address configuration in the absence of a DHCP server). The support for IPSec is an IPv6 protocol suite requirement. This requirement provides a standards-based solution for network security needs and promotes interoperability between different IPv6 implementations.

The new format of the IPv6 header is designed to minimize header validation and processing. In addition, a new field in the IPv6 header helps to define how traffic is handled and identified for quality of service delivery.

IPv6 can be extended for new features by adding extension headers after the IPv6 header. Unlike the IPv4 header, which can only support 40 bytes of options, the size of IPv6 extension headers is only constrained by the size of the IPv6 packet.

The new Neighbor Discovery protocol in IPv6 is a series of Internet Control Message Protocol for IPv6 (ICMPv6) messages that

manage the interaction of neighboring nodes. Neighbor Discovery replaces Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and ICMPv4 Redirect messages with efficient multicast and unicast messages.

The following table compares the key features of the IPv4 and IPv6 protocols.

**Comparison of Features in IPv4 and IPv6**

| IPv4 | IPv6 |
|---|---|
| Source and destination addresses are 32 bits (4 bytes) in length. | Source and destination addresses are 128 bits (16 bytes) in length. |
| IPSec support is optional. | IPSec support is required. |
| No identification of packet flow for Quality of Service (QoS) handling by routers is present within the IPv4 header. | Packet flow identification for QoS handling by routers is included in the IPv6 header using the Flow Label field. |
| Fragmentation is done by both routers and the sending host. | Fragmentation is not done by routers, only by the sending host. |
| Header includes a checksum. | Header does not include a checksum. |
| Header includes options. | All optional data is moved to IPv6 extension headers. |
| Address Resolution Protocol (ARP) uses broadcast ARP Request frames to resolve an IPv4 address to a link layer address. | ARP Request frames are replaced with multicast Neighbor Solicitation messages. |
| Internet Group Management Protocol (IGMP) is used to manage local subnet group membership. | IGMP is replaced with Multicast Listener Discovery (MLD) messages. |
| ICMP Router Discovery is used to determine the IPv4 address of the best default gateway and is optional | ICMP Router Discovery is replaced with ICMPv6 Router Solicitation and Router Advertisement messages and is required. |
| Broadcast addresses are used to send traffic to all nodes on a subnet. | There are no IPv6 broadcast addresses. Instead, a link-local scope all-nodes multicast address is used. |
| Must be configured either manually or through DHCP. | Does not require manual configuration or DHCP. |
| Uses host address (A) resource records in the Domain Name System (DNS) to map host names to IPv4 addresses. | Uses host address (AAAA) resource records in the Domain Name System (DNS) to map host names to IPv6 addresses. |
| Uses pointer (PTR) resource records in the INADDR.ARPA DNS domain to map IPv4 addresses to host names. | Uses pointer (PTR) resource records in the IP6.ARPA DNS domain to map IPv6 addresses to host names. |

For more information about Internet Protocol version 6, see the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=29519

**Using the IPv6 Protocol Across the IPv4 Internet**

On networks that do not have native support for IPv6 traffic, the IPv6 traffic is transmitted on the network by encapsulating the IPv6 packets with an IPv4 header, a technique known as IPv6 tunneling. One IPv6 tunneling technology that provides IPv6 connectivity across the IPv4 Internet is 6to4.

For more information about 6to4, see "Connection of IPv6 Domains via IPv4 Clouds," in RFC 3056 on the IETF Web site at:

http://go.microsoft.com/fwlink/?LinkId=29898

(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

If native IPv6 connectivity does not exist, a 6to4-capable computer makes a Domain Name System (DNS) query for 6to4 relay routers. By default, this DNS query is presently set to "6to4.ipv6.microsoft.com" and the response contains a well-known IPv4 anycast address. (An *anycast* address is one that identifies multiple nodes and interfaces.) As more 6to4 relay routers are added in the future, this address will be assigned to more computers that are owned by various Internet service providers (ISPs).

If the DNS query provides multiple addresses, the host selects an appropriate 6to4 relay router by sending an IPv6 packet to

each one and choosing the one that responds first.

| Note |
| --- |
| 6to4 tunneling is enabled when IPv6 services are not native to your network and there is an assigned public IPv4 Internet address. |

The use of IPv6 in Microsoft Windows Server 2003 is currently supported only when IPv4 is also installed.

**Controlling IPv6 Protocol Traffic**

You can stop the ingress or egress of IPv6 traffic on your network by configuring your network firewalls to block all IPv6 packets. For 6to4 traffic, you can configure your firewalls to block all IPv4 packets that include the IPv6 protocol designation of 41 in the Protocol field of the IPv4 packet header.

The default settings for a member of the Users group do not permit those users to install networking protocols. You should limit who is allowed to install the IPv6 protocol stack on network computers by carefully limiting the number of users that have administrative logon credentials.

You can use the Active Directory directory service and Group Policy to filter and control the ability to add new networking protocols or modify existing networking configurations. For more information about these configuration methods, see Appendix B: Resources for Learning About Group Policy. For information about installing and uninstalling IPv6, see the list of procedures in the next subsection.

**Procedures for Configuration of the IPv6 Protocol**

Installing and uninstalling the IPv6 protocol stack can be done by using the Network Connections folder or the command prompt.

The following procedures tell how to find information about security in relation to IPv6 and describe installing and uninstalling the IPv6 protocol stack by using the Network Connections folder.

**To Find Information About Security in Relation to IPv6**

1. View updated product documentation for Windows Server 2003 products on the Web at:

   http://go.microsoft.com/fwlink/?linkid=29881

2. In the table of contents, navigate as follows:

   Network Services\Managing Core Network Services\IP Version 6\IPv6 Concepts\IPv6 Overview

3. View topics below “IPv6 Overview,” especially “Security information for IPv6.”

**To Install IPv6 Using the Network Connections Folder**

1. Click **Start**.

2. Either click **Control Panel** and then double-click **Network Connections**, or point to **Settings**, click **Control Panel**, and then double-click **Network Connections**.

3. Right-click any local area connection, and then click **Properties**.

4. Click **Install**.

5. In the **Select Network Component Type** dialog box, click **Protocol**, and then click **Add**.

6. In the **Select Network Protocol** dialog box, click **Microsoft TCP/IP version 6**.

**To Uninstall IPv6 Using the Network Connections Folder**

1. Click **Start**.

2. Either click **Control Panel** and then double-click **Network Connections**, or point to **Settings**, click **Control Panel**, and then double-click **Network Connections**.

3. Right-click any local area connection, and then click **Properties**.

4. Click **Microsoft TCP/IP version 6** in the list of installed components, and then click **Uninstall**.

5. Restart the computer.

The following two procedures describe installing and uninstalling the IPv6 protocol stack from the command prompt.

**To Install IPv6 on a Computer from the Command Prompt**

1. To open a Command Prompt window, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. At the command prompt, type **netsh interface ipv6 install**, and then press ENTER.

**To Uninstall IPv6 from a Computer from the Command Prompt**

1. To open a Command Prompt window, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. At the command prompt, type **netsh interface ipv6 uninstall**, and then press ENTER.

3. Restart the computer.

| **Note** |
| --- |
| The IPv6 configuration procedures require that you have administrative credentials on the computer. |

**Monitoring and Troubleshooting the IPv6 Protocol**
The following procedures describe ways to view TCP/IP configurations.

**To Display the Complete List of TCP/IP Interface Configurations for a Computer from the Command Prompt**

1. To open a Command Prompt window, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. At the command prompt, type **ipconfig /all**, and then press ENTER.

**To Display the TCP/IP Routing Table from the Command Prompt**

1. To open a Command Prompt window, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. At the command prompt, type **route print**, and then press ENTER.

**Troubleshooting a Command-line Installation Error**
The installation of the IPv6 protocol stack requires that you have administrative credentials. If you attempt to install the IPv6 protocol from the command line without having the required account credentials, the result is an "Access is denied" error (0x800700005).

**Related Links**
**Web Resources**

- For more information about 6to4, see "Connection of IPv6 Domains via IPv4 Clouds," in RFC 3056 on the IETF Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29898

- For more information about IP version 6, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29519

- For information about security in relation to IPv6, see “To Find Information About Security in Relation to IPv6,” earlier in this white paper.

- For more information about IPv6 addressing, see "Internet Protocol Version 6 (IPv6) Addressing Architecture" in RFC 3513 on the IETF Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29135

- For the latest set of RFCs and Internet drafts describing IPv6 standards, see “IP Version 6 Working Group (ipv6)” at the IETF Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29136

- For the latest set of RFCs and Internet drafts describing IPv6 transition technologies, see “Next Generation Transition (ngtrans)” at the IETF Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29215

  (Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

**Printed References**

For more information about the IPv6 protocol suite, you can consult the following references.


- Davies, J. *Understanding IPv6*. Redmond, WA: Microsoft Press, 2002.

- Hagen, S. *IPv6 Essentials*. Sebastopol, CA: O'Reilly and Associates, Inc., 2002.

# NetMeeting and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of NetMeeting® conferencing software
- Using NetMeeting in a managed environment
- How NetMeeting communicates with sites on the Internet
- How to control NetMeeting to limit the flow of information to and from the Internet

**Benefits and Purposes of NetMeeting**

NetMeeting conferencing software is a feature of Windows Server 2003 with Service Pack 1 (SP1) that enables real-time communication and collaboration over the Internet or an intranet. From a computer running the Windows 95, Windows 98, Windows NT® 4.0, Windows 2000, or Windows XP operating system, users can communicate over a network with real-time voice and video technology. Users can work together on virtually any Windows application, exchange or mark up graphics on an electronic whiteboard, transfer files, or use the text-based chat program.

NetMeeting helps small and large organizations take full advantage of their corporate intranet for real-time communication and collaboration. On the Internet, connecting to other NetMeeting users is made easy with Internet Locator Service (ILS), enabling participants to call each other from a dynamic directory within NetMeeting or from a Web page. Features include remote desktop sharing, virtual conferencing using Microsoft Outlook, security features, and the ability to embed the NetMeeting user interface in an organization's intranet Web pages.

To learn more about the NetMeeting features, see "Microsoft NetMeeting 3 Features" on the Microsoft TechNet Web site at:

http://go.microsoft.com/fwlink/?LinkId=29175

**Overview: Using NetMeeting in a Managed Environment**

NetMeeting supports communication standards for audio, video, and data conferencing. NetMeeting users can communicate and collaborate with users of other standards-based, compatible products. They can connect by modem, Integrated Services Digital Network (ISDN), or local area network (LAN) using Transmission Control Protocol/Internet Protocol (TCP/IP). In addition, support for Group Policy in NetMeeting makes it easy for administrators to centrally control and manage the NetMeeting work environment.

You can use Active Directory directory service and Group Policy to configure NetMeeting to help meet your security requirements. You can also control the configuration of NetMeeting by using the NetMeeting Resource Kit. For more information about the NetMeeting Resource Kit, see "Alternate Methods for Controlling NetMeeting," later in this section.

NetMeeting components and features require that several ports be open from the firewall. For more information, see "NetMeeting and Firewalls," later in this section.

**How NetMeeting Communicates with Sites on the Internet**

NetMeeting provides an infrastructure for communication between network applications and services. In this infrastructure, NetMeeting is both an application and a platform for other applications or services. The components and services in NetMeeting provide real-time communication and collaboration over the Internet or an organization's intranet.

NetMeeting audio and video conferencing features are based on the H.323 standard infrastructure, which enables NetMeeting to interoperate with other H.323 standards-based products. (H.323 is a standard approved by the International Telecommunication Union [ITU] that defines how audiovisual conferencing data is transmitted across networks.) NetMeeting data conferencing features are based on the T.120 infrastructure, enabling NetMeeting to interoperate with other T.120 standards-based products. (The T.120 standard is a suite of communication and application protocols developed for real-time, multipoint data connections and conferencing.)

Detailed information about the H.323 and T.120 standards is beyond the scope of this white paper. Further information can be found on the following sites.

- For more information about the H.323 standard and NetMeeting, see Part 3, Chapter 11, "Understanding the H.323 Standard," in the *NetMeeting 3.0 Resource Kit* at:

http://go.microsoft.com/fwlink/?LinkId=36134

- For more information about the H.323 specification, see the Web sites of the International Telecommunication Union (ITU) and the International Multimedia Telecommunications Consortium (IMTC) at:

  http://go.microsoft.com/fwlink/?LinkId=29510

  http://www.imtc.org/

- To learn more about the T.120 standard and NetMeeting, see Part 3, Chapter 10, "Understanding the T.120 Standard," in the *NetMeeting 3.0 Resource Kit* at:

  http://go.microsoft.com/fwlink/?LinkId=36134

- For more information about the T.120 architecture, see the International Multimedia Teleconferencing Consortium (IMTC) Web site at:

  http://www.imtc.org/

(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

**NetMeeting Port Assignments**
When you use NetMeeting to call other users over the Internet, several IP ports are required to establish the outbound connection. The following table describes the port numbers, their functions, and the resulting connection.

**Port Assignments for NetMeeting**

| Port | Function | Outbound Connection |
|---|---|---|
| 389 | Internet Locator Service (ILS) | TCP |
| 522 | User Location Service (ULS) | TCP |
| 1503 | T.120 | TCP |
| 1720 | H.323 call setup | TCP |
| 1731 | Audio call control | TCP |
| 1024 through 65535 (dynamic) | H.323 call control | TCP |
| 1024 through 65535 (dynamic) | H.323 streaming | Real-Time Transfer Protocol (RTP) over User Datagram Protocol (UDP) |

For more information about NetMeeting communication ports and firewall configuration topics, see Part 2, Chapter 4, "Firewall Configuration" in the *NetMeeting 3.0 Resource Kit* at:

http://go.microsoft.com/fwlink/?LinkId=36134

**Controlling NetMeeting to Limit the Flow of Information to and from the Internet**
You can configure NetMeeting by using Group Policy objects (GPOs) on servers running Windows Server 2003. (You can also control the configuration of NetMeeting by using the NetMeeting Resource Kit. For more information, see "Alternate Methods for Controlling NetMeeting," later in this section.)

This subsection includes information about the following topics:

- NetMeeting and Group Policy
- NetMeeting security
- NetMeeting and firewalls

- Establishing a NetMeeting connection with a firewall
- Firewall limitations for NetMeeting

**NetMeeting and Group Policy**

Group Policy can be used to define the default NetMeeting configuration settings that will be automatically applied to users and computers. These settings determine which NetMeeting features and capabilities are available to a particular computer or to a particular group of users. Through the use of Group Policy you can enable, disable, or set configuration options for NetMeeting features or capabilities.

For additional information about Group Policy, see Appendix B: Resources for Learning About Group Policy.

You can use Group Policy to manage the following NetMeeting configuration options for users in your organization:

- NetMeeting Group Policy setting for computers
- NetMeeting Group Policy settings for users

**Configuring NetMeeting Setting for Computers Through Group Policy**

In Group Policy, the NetMeeting policy setting for computers is located in Computer Configuration\Administrative Templates\Windows Components\NetMeeting:

- **Disable remote Desktop Sharing**

For more information about how to use Group Policy to manage the NetMeeting computer setting, see "To Disable the NetMeeting Remote Desktop Sharing Feature Through Group Policy," later in this section.

| **Note** |
| --- |
| Group Policy settings for computers are applied when the operating system starts and during the periodic refresh cycle. |

**Configuring NetMeeting Settings for Users Through Group Policy**

The Group Policy configuration options that apply to users include general NetMeeting settings, and settings for application sharing, audio and video, and the options page.

**Configuring General NetMeeting Settings for Users Through Group Policy**

In Group Policy, general NetMeeting policy settings for users are located in User Configuration\Administrative Templates\Windows Components\NetMeeting:

- **Enable Automatic Configuration**: Configures NetMeeting to download settings for users each time it starts.
- **Disable Directory services**: Disables the directory feature—users will not log on to a directory server when NetMeeting starts. Users will not be able to view or make calls using the NetMeeting directory.
- **Prevent adding Directory servers**: Prevents the user from adding directory servers to the list of available directory servers they can use for placing calls.
- **Prevent viewing Web directory**: Prevents the user from viewing directories as Web pages in a browser.
- **Set the intranet support Web page**: Sets the Web address that NetMeeting will display when users choose the Online Support command from the NetMeeting **Help** menu.
- **Set Call Security options**: Sets the level of security for outgoing and incoming NetMeeting calls.
- **Prevent changing Call placement method**: Prevents the user from changing the way calls are placed, either directly or by means of a gatekeeper server.
- **Prevent automatic acceptance of Calls**: Prevents the user from turning on automatic acceptance of incoming calls.

- **Allow persisting automatic acceptance of Calls**: Sets automatic acceptance of incoming calls to be persistent.
- **Prevent sending files**: Prevents users from sending files to others in a conference.
- **Prevent receiving files**: Prevents users from receiving files from others in a conference.
- **Limit the size of sent files**: Sets the maximum file size that can be sent to others in a conference.
- **Disable Chat**: Disables the chat feature of NetMeeting.
- **Disable NetMeeting 2.x Whiteboard**: Disables the NetMeeting 2.x Whiteboard feature. (The 2.x feature provides compatibility with earlier versions of NetMeeting only.)
- **Disable Whiteboard**: Disables the whiteboard feature of NetMeeting.

| **Note** |
| --- |
| Group Policy settings for users are applied when a user logs on to the computer and during the periodic refresh cycle. |

**Configuring NetMeeting Application Sharing Settings Through Group Policy**

Group Policy settings for the NetMeeting Application Sharing feature are located in User Configuration\Administrative Templates\Windows Components\NetMeeting\Application Sharing:

- **Disable application Sharing**: Disables the NetMeeting application sharing feature completely. Users will not be able to host or view shared applications.
- **Prevent Sharing**: Prevents users from sharing anything themselves. They will still be able to view shared applications or desktops from others.
- **Prevent Desktop Sharing**: Prevents users from sharing their Windows desktop. They will still be able to share individual applications.
- **Prevent Sharing Command Prompts**: Prevents the user from sharing command prompts. Enabling this prevents the user from inadvertently sharing applications, since command prompts can be used to start other applications.
- **Prevent Sharing Explorer windows**: Prevents the user from sharing Windows Explorer windows. Enabling this prevents the user from inadvertently sharing applications, since Windows Explorer windows can be used to start other applications.
- **Prevent Control**: Prevents users from allowing others in a conference to control what they have shared. Enabling this enforces a read-only mode whereby the other participants cannot change the data in the shared application.
- **Prevent Application Sharing in true color**: Prevents users from sharing applications in true color, which uses more bandwidth.

| **Note** |
| --- |
| Group Policy settings for users are applied when a user logs on to the computer and during the periodic refresh cycle. |

**Configuring NetMeeting Audio and Video Settings Through Group Policy**

Group Policy settings for NetMeeting audio and video are located in User Configuration\Administrative Templates\Windows Components\NetMeeting\Audio & Video:

- **Limit the bandwidth of Audio and Video**: Configures the maximum bandwidth, specified in kilobytes per second, to be used for audio and video.

- **Disable Audio:** Disables the audio feature of NetMeeting. Users will not be able to send or receive audio.

- **Disable full duplex Audio**: Disables the full duplex audio mode. Users will not be able to listen to incoming audio while speaking into the microphone. Older audio hardware may not perform well when full duplex audio is enabled.

- **Prevent changing DirectSound Audio setting**: Prevents the user from changing the DirectSound® audio setting. DirectSound has better audio quality, but earlier audio hardware may not support DirectSound.

- **Prevent sending Video**: Prevents the user from sending video. Setting this option does not prevent the user from receiving video.

- **Prevent receiving Video**: Prevents the user from receiving video. Setting this option does not prevent the user from sending video.

| **Note** |
| --- |
| Group Policy settings for users are applied when a user logs on to the computer and during the periodic refresh cycle. |

**Configuring NetMeeting Options Settings Through Group Policy**

Group Policy settings for NetMeeting options are located in User Configuration\Administrative Templates\Windows Components\NetMeeting\Options Page. These settings correspond to elements of the NetMeeting Options page. To view this page, in NetMeeting, click **Tools** and then click **Options**. These Group Policy settings are as follows:

- **Hide the General page**: Removes the **General** tab on the **NetMeeting Options** page.

- **Disable the Advanced Calling button**: Disables the **Advanced Calling** button from the **General** page.

- **Hide the Security page**: Removes the **Security** tab on the NetMeeting **Options** page.

- **Hide the Audio page**: Removes the **Audio** tab on the NetMeeting **Options** page.

- **Hide the Video page**: Removes the **Video** tab on the NetMeeting **Options** page.

| **Note** |
| --- |
| Group Policy settings for users are applied when a user logs on to the computer and during the periodic refresh cycle. |

To learn about specific Group Policy settings that can be applied to computers running Windows Server 2003 SP1, see the Group Policy Settings Reference on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=29911

**NetMeeting Security**

The NetMeeting security architecture for data conferencing takes advantage of the existing, standards-compliant security features of Windows Server 2003 SP1 and Microsoft Internet Explorer. The NetMeeting security architecture utilizes a 40-bit encryption technology and has the following security features.

- **Password protection**: This feature enables the user to create or participate in a meeting that requires a password to join. Password protection helps to ensure that only authorized users participate in a password-protected meeting. A password is also required to use the remote desktop sharing feature.

- **User authentication**: This feature provides a way to verify the identity of a caller or meeting participant using a certificate.

- **Data encryption**: This feature helps to protect data exchanged during a meeting so that it is not easily read by any unauthorized parties that may intercept the data. The 40-bit data encryption applies to the whiteboard and chat features, shared applications, and transferred files. Audio and video communications are not encrypted.

NetMeeting security features integrate with security in Windows Server 2003 SP1 and Internet Explorer in a variety of ways, including the following:

- NetMeeting uses the NetMeeting private certificate store to provide personal certificates for user authentication and data encryption.
- NetMeeting uses the Windows certificate store to maintain NetMeeting certificates.
- NetMeeting uses Security Support Provider Interface (SSPI) functions to generate and process security tokens.

These security features can be implemented by an administrator or a NetMeeting user. Using the NetMeeting Resource Kit Wizard or Group Policy in NetMeeting, the administrator can enforce security settings that apply to all users. If allowed by the administrator, NetMeeting users can also select their own security settings in the NetMeeting user interface (UI) and change security settings for individual calls.

You can use the following sources to learn more about NetMeeting configuration and security topics.

- For more information about the NetMeeting Resource Kit Wizard, see Part 2, Chapter 2, "Resource Kit Wizard" in the *NetMeeting 3.0 Resource Kit* at:

  http://go.microsoft.com/fwlink/?LinkId=36134

- For more information about the security features available in NetMeeting, see Part 2, Chapter 5, "NetMeeting Security" in the *NetMeeting 3.0 Resource Kit* at:

  http://go.microsoft.com/fwlink/?LinkId=36134

**NetMeeting and Firewalls**

You can configure firewall components in a variety of ways, depending on your organization's specific security policies and overall operations. While most firewalls are capable of allowing primary (initial) and secondary (subsequent) Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) connections, it is possible that they are configured to support only specific connections based on security considerations. For example, some firewalls support only primary TCP connections, which some professionals view as the most reliable.

For NetMeeting multipoint data conferencing—program sharing, whiteboard, chat, file transfer, and directory access—your firewall only needs to pass through primary TCP connections on assigned ports. NetMeeting audio and video features require secondary TCP and UDP connections on dynamically assigned ports.

| **Note** |
| --- |
| NetMeeting audio and video features require secondary TCP and UDP connections. Therefore, when you establish connections through firewalls that accept only primary TCP connections, you are not able to use the audio or video features of NetMeeting. |

Detailed firewall configuration procedures for NetMeeting are beyond the scope of this white paper. For more information about NetMeeting firewall connections, see Part 2, Chapter 4, "Firewall Configuration" in the *NetMeeting 3.0 Resource Kit*, particularly the section titled, "Establishing a NetMeeting Connection with a Firewall" at:

http://go.microsoft.com/fwlink/?LinkId=36134

Microsoft NetMeeting can be configured to work with an organization's existing firewall security. Because of limitations in most firewall technology, however, few products are available that enable you to securely transport inbound and outbound NetMeeting calls containing audio, video, and data across a firewall. You should consider carefully the relative security risks of enabling different parts of a NetMeeting call in your firewall product. You must especially consider the security risks involved when modifying your firewall configuration to enable any component of an inbound NetMeeting call.

Some organizations have security or policy concerns that require them to limit how fully they support NetMeeting in their firewall configuration. These concerns are based on network capacity planning or weaknesses in the firewall technology being used. For example, security concerns might prohibit an organization from accepting any inbound or outbound flow of UDP data through the firewall. Because these UDP connections are required for NetMeeting audio and video features, disabling this function excludes audio and video features in NetMeeting for calls through the firewall. The organization can still use

NetMeeting data conferencing features such as program sharing, file transfer, whiteboard, and chat for calls through the firewall by allowing only TCP connections on ports 522 and 1503.

For more information about NetMeeting firewall security, see the section titled "Security and Policy Concerns," in the chapter of the NetMeeting Resource Kit from the previous link (scroll through the chapter until you find the section).

**Establishing a NetMeeting Connection with a Firewall**
When you use NetMeeting to call other users over the Internet, several IP ports are required to establish the outbound connection.

If you use a firewall to connect to the Internet, it must be configured so that the following IP ports are not blocked:

- TCP ports 389, 522, 1503, 1720, and 1731
- TCP and UDP ports (1024 through 65535)

To establish outbound NetMeeting connections through a firewall, the firewall must be configured to do the following:

- Pass through primary TCP connections on ports 389, 522, 1503, 1720, and 1731.
- Pass through secondary TCP and UDP connections on dynamically assigned ports (1024 through 65535).

The H.323 call setup protocol dynamically negotiates a TCP port for use by the H.323 call control protocol. Also, both the audio call control protocol and the H.323 call setup protocol dynamically negotiate UDP ports for use by the H.323 streaming protocol, the Real-Time Transfer Protocol (RTP). In NetMeeting, two UDP ports are designated on each side of the firewall for audio and video streaming, for a total of four ports for inbound and outbound audio and video. These dynamically negotiated ports are selected arbitrarily from all ports that can be assigned dynamically.

NetMeeting directory services require either port 389 or port 522, depending on the type of server you are using. The Microsoft Internet Locator Service (ILS), which supports LDAP for NetMeeting, requires port 389. The Microsoft User Location Service (ULS), developed for NetMeeting 1.0, requires port 522.

**Firewall Limitations for NetMeeting**
Some firewalls cannot support an arbitrary number of virtual internal IP addresses, or cannot do so dynamically. With these firewalls, you can establish outbound NetMeeting connections from computers inside the firewall to computers outside the firewall, and you can use the audio and video features of NetMeeting. Users outside the organization cannot, however, establish inbound connections from outside the firewall to computers inside the firewall. Typically, this restriction is due to limitations in the network implementation of the firewall.

| **Note** |
| --- |
| Some firewalls are capable of accepting only certain protocols and cannot handle TCP connections. For example, if your firewall is a Web proxy server with no generic connection-handling mechanism, you will not be able to use NetMeeting through the firewall. |

You can use the following sources to learn more about NetMeeting configuration and firewall topics.

- For more information about NetMeeting firewall connections, see Part 2, Chapter 4, "Firewall Configuration" in the *NetMeeting 3.0 Resource Kit*, particularly the section titled "Establishing a NetMeeting Connection with a Firewall" at:

  http://go.microsoft.com/fwlink/?LinkId=36134

- For more information about using NetMeeting and your firewall, see article 158623, "How to Establish NetMeeting Connections through a Firewall" in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?LinkId=29206

**Alternate Methods for Controlling NetMeeting**
You can create customized installation options for specific users or groups within your organization by using the NetMeeting Resource Kit Wizard. Additionally, you can use the NetMeeting Resource Kit Wizard to control user and computer access rights by creating custom configurations of client settings and specific features that you have selected to restrict or allow. For

example, you can control audio and video access, set data throughput limits and network speeds, and choose to display online support. The Resource Kit Wizard can also help you set up various configurations of NetMeeting for different types of users and different levels of security. It can help you save network bandwidth by restricting specific features. You can also use the Resource Kit Wizard both to change registry settings for all NetMeeting users and to implement such changes globally.

| **Note** |
| --- |
| By selecting certain options in the Resource Kit Wizard, be aware that you may be changing the NetMeeting user interface. For example, if you click **Restrict the Use of Video**, the **Video** tab does not appear in the NetMeeting user's **Options** dialog box. |

Part 4 of the Resource Kit for NetMeeting has an appendix that provides detailed information about responding to NetMeeting problems, including problem descriptions, causes, and resolutions.

For more information about the *NetMeeting 3.0 Resource Kit*, see the Microsoft TechNet Web site at:

http://go.microsoft.com/fwlink/?LinkId=36134

**Procedures for Configuring NetMeeting**

NetMeeting is designed to enhance the enterprise environment and enable users to communicate internally and externally with other NetMeeting users. You can use Group Policy to develop a NetMeeting feature management policy to support the specific business rules or communication policies that exist within your organization. For example, your organization may not want users to be able to access or use the NetMeeting chat feature from their computers. By using Active Directory and Group Policy, you can disable the chat feature from any or all computers that are affected by the application of the Group Policy configuration settings.

For lists of Group Policy settings that you can use to manage NetMeeting configuration options, see "NetMeeting and Group Policy," earlier in this section.

**Procedures for Managing NetMeeting Features Through Group Policy**

This subsection provides procedures for the following configuration methods:

- Locating the Group Policy objects (GPOs) for NetMeeting configuration settings. These are the settings listed in "NetMeeting and Group Policy," earlier in this section.
- Disabling the NetMeeting remote desktop sharing feature. This prevents users from using this feature.
- Disabling the NetMeeting advanced calling feature.
- Disabling the NetMeeting chat feature.

**To Locate the Group Policy Settings for NetMeeting User Configuration**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.
2. Click **User Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **NetMeeting**.
3. View the **Group Policy** objects that are available. For more information about these objects, see "NetMeeting and Group Policy," earlier in this section.

**To Disable the NetMeeting Remote Desktop Sharing Feature Through Group Policy**

Use the following steps to configure the Group Policy setting to prevent users from using the NetMeeting remote desktop sharing feature.

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.
2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **NetMeeting**.
3. In the **details** pane, double-click **Disable remote Desktop Sharing**.

4. Click **Enabled**.

| **Note** |
| --- |
| Group Policy settings for computers are applied when the operating system is initialized and during the periodic refresh cycle. |

**To Disable the NetMeeting Advanced Calling Feature Through Group Policy**
Use the following steps to configure the Group Policy setting to disable the advanced calling feature on the NetMeeting options page.

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.

2. Click **User Configuration**, click **Administrative Templates**, click **Windows Components**, click **NetMeeting**, and then click **Options Page**.

3. In the details pane, double-click **Disable the Advanced Calling** button, and then click **Enabled**.

**To Disable the NetMeeting Chat Feature Through Group Policy**
Use the following steps to configure the Group Policy setting to prevent the use of the NetMeeting Chat feature.

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.

2. Click **User Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **NetMeeting**.

3. In the **details** pane, double-click **Disable Chat**, and then click **Enabled**.

| **Note** |
| --- |
| Group Policy settings for users are applied when a user logs on to the computer and during the periodic refresh cycle. |

**Related Links**
**Web Resources**

- For more information about using NetMeeting and your firewall, see article 158623, "How to Establish NetMeeting Connections through a Firewall," in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?LinkId=29206

- For more information about NetMeeting, see Windows NetMeeting on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29514

- For more information about configuring NetMeeting, see the Windows NetMeeting Resource Kit on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=36134

- To learn more about NetMeeting features, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29175

- To view articles that explain how to use some of the features in NetMeeting, see the Microsoft NetMeeting How-to Guide on Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29207

- For more information about the H.323 specification, search for "H.323" on the ITUT Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29510

- For more information about the T.120 architecture, see the International Multimedia Teleconferencing Consortium (IMTC) Web site at:

  http://www.imtc.org/

(Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

**Printed References**

For more information about firewall design, policy, and security considerations for firewall design in general, you can consult the following reference.


- Chapman, D. Brent and Elizabeth D. Zwicky. *Building Internet Firewalls*. O'Reilly & Associates, Inc., 1995.

# Online Device Help and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Online Device Help
- How Online Device Help communicates with sites on the Internet
- How to control Online Device Help to limit the flow of information to and from the Internet

**Benefits and Purposes of Online Device Help**
Online Device Help (also known as the "Get help for my hardware device" wizard) delivers targeted content on problems with hardware and peripheral devices installed on the system. This mitigates the need for calls to support professionals to resolve hardware issues. You interact with Online Device Help in products in the Microsoft Windows Server 2003 family when installing new hardware (through the Found New Hardware Wizard).

At the conclusion of the Found New Hardware Wizard, when the system uses a device driver that is not found on the operating system installation CD or is not available through Windows Update, Online Device Help collects anonymous data on the problem device (including a unique hardware identifier for that device) and sends that information over the Internet to a server at Microsoft. If a match for the device is found, content on that device is then downloaded to the system and is displayed in the Help and Support Center user interface. This content may include:

- Information from the independent hardware vendor (IHV) about upcoming and planned device support.
- Links to the product compatibility area in Help and Support Center to enable users to search or browse the Windows Server Catalog Web site for compatible devices. The Windows Server Catalog Web site is located at:

  http://go.microsoft.com/fwlink/?LinkId=4303

- A link to the Web site of the IHV for that device.

The data provided by Online Device Help enables Microsoft to identify the number and system locale of users experiencing hardware problems due to missing drivers and to identify the most common problem devices. Microsoft works with these hardware vendors to provide targeted troubleshooting content on the most common problem hardware devices.

This section of the white paper explains how to control Online Device Help in a managed environment.

**Overview: Using Online Device Help in a Managed Environment**
You have control over whether to upload the data required by Online Device Help. You can also block Online Device Help at the firewall or through the Services snap-in. The configuration options and procedures for controlling Online Device Help are described later in this section.

**How Online Device Help Communicates with Sites on the Internet**
If no information for a particular hardware device is found on either the installation CD for the operating system or through Windows Update, you are prompted to release anonymous information about the hardware profile through Online Device Help. This subsection summarizes the communication process:

- **Specific information sent or received**: The following information is collected from the computer running Windows Server 2003 and is uploaded to a server at Microsoft. The user or administrator is not uniquely identified.
  - The hardware ID, also known as the PnPID (code that indicates the device manufacturer, device name, and version)
  - The time and date that the data was sent
  - Language code of the operating system, and platform and build information

- **Default and recommended settings**: Online Device Help is enabled by default. Recommended settings are described in the next subsection, "Controlling Online Device Help to Limit the Flow of Information to and from the Internet."

- **Triggers**: Online Device Help is triggered if no information for a particular hardware device is found after you complete the Found New Hardware Wizard.

- **User notification**: You are prompted to send anonymous hardware profile data to a server at Microsoft. If you opt to send this information, the privacy statement is displayed. You can view the contents of the hardware.xml file being uploaded by clicking a link on the privacy statement page.

- **Logging**: Errors that result from problems installing hardware devices without drivers are logged to the event log.

- **Encryption**: The data transferred to Microsoft is not encrypted.

- **Access**: The raw data uploaded to the server is accessible to operations engineers at Microsoft.com and is used in the Windows Hardware Quality Labs (WHQL) to improve Windows-compatible hardware and drivers.

- **Privacy statement**: Online Device Help is covered by its own privacy statement. The privacy statement (located in a file on your computer at systemroot\pchealth\helpctr\system\dfs\privacy.htm) is displayed when you opt to send the anonymous hardware profile data to Microsoft.

- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is 80.

- **Ability to disable**: You cannot disable Online Device Help directly. Disabling Internet access or HTTP port 80 will, however, block Online Device Help, as will disabling the Upload Manager service.

**Controlling Online Device Help to Limit the Flow of Information to and from the Internet**

You have control over whether to upload anonymous hardware profile information about your system. You cannot, however, disable Online Device Help directly. To block Online Device Help, you can restrict Internet access or configure the Services snap-in to disable the Upload Manager service (uploadmgr).

**How Controlling Online Device Help Can Affect Users and Applications**

If you decide to disable Online Device Help, users or administrators will not be prompted to upload anonymous hardware profile information and they will not receive up-to-date, targeted self-help content on hardware issues relating to missing or problem drivers.

| **Note** |
| --- |
| If you restrict Internet access to block Online Device Help, the feature will queue the data and periodically retry to upload the hardware profile information for some period of time. If an Internet connection becomes available during that period, Online Device Help will upload the queued data. If an Internet connection does not become available, no data will be uploaded. Users will not be impacted in either case. |

**Procedure for Controlling Online Device Help**

You cannot disable Online Device Help directly but can do so indirectly by disabling the Upload Manager service in Windows Server 2003.

**To Disable Online Device Help by Disabling the Upload Manager Service**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then double-click **Services**.

3. In the details pane, right-click **Upload Manager**, and then click **Properties**.

   The Upload Manager service is the underlying service used by Online Device Help.

4. Click the **Log On** tab, then click the hardware profile that you want to configure, and then click **Disable**.

# Outlook Express 6.0 and Resulting Internet Communication in Windows Server 2003 with Service Pack 1

Applies To: Windows Server 2003 with SP1

This white paper provides information about the communication that flows between components in Windows Server 2003 with SP1 and sites on the Internet, and it describes steps to take to limit, control, or prevent that communication in an organization with many users.

This section of the white paper provides:

- A description of Microsoft Outlook Express 6 in Windows Server 2003 with Service Pack 1 (SP1). This section also provides a comparison of Outlook and Outlook Express.
- Descriptions of new security-related features in Outlook Express 6 in Windows Server 2003 with SP1 (as compared to Outlook Express 5), with information about how these new features are configured at the desktop.
- Information about controlling Outlook Express 6 through Group Policy to limit the risk associated with e-mail attachments. The Group Policy setting that you use for this is **Block attachments that could contain a virus**.

| **Note** |
| --- |
| This section of the white paper describes Outlook Express 6 in Windows Server 2003 with SP1, but does not describe related components such as Internet Explorer 6, the New Connection Wizard, or the tool that can report errors that occur in Outlook Express. For information about these components, see the respective sections of this white paper (the error reporting tool is described in Windows Error Reporting and Internet Communication). |

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization where users send and receive e-mail messages, open attachments in e-mail messages, and perform similar actions. This section, however, provides information about features and configuration methods in Outlook Express 6 that can reduce the inherent risks associated with sending and receiving e-mail messages.

For more information about Outlook Express, see the following resources:

- Help for Outlook Express (which can be accessed in Outlook Express by clicking the **Help** menu and then selecting an appropriate option).
- The section about Internet Explorer 6 in this white paper, which describes security zones in Internet Explorer 6. These security zones are also used in Outlook Express 6.
- The Internet Explorer page on the Microsoft Web site at:

  http://www.microsoft.com/windows/ie/

- The Resource Kit for Internet Explorer (specifically, the chapter describing what's new in Internet Explorer 6). To learn about this and other Resource Kits, see the Microsoft TechNet Web site at:

  http://go.microsoft.com/fwlink/?linkid=29894

**Benefits and Purposes of Outlook Express 6**

Outlook Express 6 is designed to make it easy to send or receive e-mail messages and to browse or participate in newsgroups. It differs from many of the other components described in this white paper in that its main function is to communicate through the Internet or an intranet (in contrast to components that communicate with the Internet in the process of supporting some other activity).

Outlook Express is part of Windows Server 2003 with SP1, in contrast to Microsoft Outlook, which is an application included in Microsoft Office. Outlook provides comprehensive e-mail capabilities, including information management and collaboration

capabilities, useful to a wide spectrum of users from home to small business to large enterprise. Outlook Express, part of Windows Server 2003 with SP1, offers standard Internet e-mail and news access, useful to many home and small-business users. Outlook Express supports Post Office Protocol 3 (POP3), Internet Message Access Protocol (IMAP), Simple Mail Transfer Protocol (SMTP), and Network News Transfer Protocol (NNTP).

Outlook Express 6 offers more security-related options and settings than were available in Outlook Express 5. The following subsections describe the new options and ways of configuring them.

**New Security-Related Features in Outlook Express 6**

The version of Outlook Express in Windows Server 2003 with SP1 includes additional security-related features as compared to earlier versions of Outlook Express including Outlook Express 5. The following list describes these features. The table that follows this list shows how each option is configured in Outlook Express.

- **Warning about harmful e-mail**. To prevent e-mail messages from being sent without your knowledge, Outlook Express warns you when other programs, such as viruses or harmful attachments, attempt to send messages from your computer. This warning appears only if Outlook Express is configured as the default simple MAPI client, and another program attempts to use simple MAPI to programmatically send e-mail messages without presenting a visible user interface on the computer.

- **Blocking of potentially harmful attachments**. If this option is enabled, Outlook Express 6 blocks the opening or saving of specific e-mail attachments that are considered "unsafe." To determine whether an attachment is unsafe, Outlook Express 6 uses an updated API in Windows Server 2003 with SP1, the Attachment Manager API. The Attachment Manager API provides functionality to give each attachment a risk rating based on the extension, content type, registered handlers, and other heuristics. By using Group Policy, you can customize some aspects of Attachment Manager, such as the lists of high, medium, and low risk files.

  In addition, the prompts that are used for mail attachments, file downloads, shell process execution, and program installation have been modified to be both more consistent and clearer than they were in Windows Server 2003 without SP1.

  Blocking of potentially harmful attachments can be enabled or disabled through Group Policy as well as at the local computer. For more information about using this setting, see the table that follows and "To Locate the Group Policy Setting for Blocking E-mail Attachments in Outlook Express 6," later in this section.

  For more information about Attachment Manager and other changes that make the version of Outlook Express in Windows Server 2003 with SP1 more resistant than previous versions, see "Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1" on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46278

  To learn about Group Policy settings with which you can adjust Attachment Manager, in Group Policy, go to User Configuration\Administrative Templates\Windows Components\Attachment Manager. For a detailed explanation of a setting, select the setting and click the **Extended** tab, or open the setting and click the **Explain** tab.

- **Plain text format option for reading of e-mail**. Starting with Outlook Express 6.0 in Windows Server 2003 with SP1, Outlook Express can be configured to read all e-mail messages in plain text format. Some HTML e-mail messages may not appear correctly in plain text, but no active content in the e-mail message is run when this setting is enabled.

- **Blocking of downloads of external content (to help limit spam)**. If this option is enabled, Outlook Express 6 will not contact an external Web server when an e-mail contains a reference to an image that resides on that external Web server. Businesses that use spam sometimes incorporate such external references for the purpose of validating e-mail addresses that they use, after which they send repeated e-mails to the validated addresses. The image involved might be a single pixel image that is not visible to the e-mail recipient, who is unaware that his or her e-mail address has been validated. This option can be enabled or disabled at the local computer. For more information about using this setting, see the table that follows and "To Start Outlook Express 6 and View or Configure Security Settings," later in this section.

  This option is new in the version of Outlook Express in Windows Server 2003 with SP1. For more details about other changes that make this version of Outlook Express more resistant than previous versions, see "Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1" on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=46278

The following table shows how each option is configured in Outlook Express 6.

**Options for Configuring Outlook Express 6**

| Option to Configure in Outlook Express 6 | Menu to Click | Menu Item to Click | Tab to Click |
|---|---|---|---|
| Warning about harmful e-mail | **Tools** | **Options** | **Security** |
| Blocking of potentially harmful attachments (also configurable through Group Policy) | **Tools** | **Options** | **Security** |
| Blocking of the downloading of images and other external content in HTML e-mail (this helps limit spam) | **Tools** | **Options** | **Security** |
| Plain text format option for reading of all e-mail | **Tools** | **Options** | **Read** |

**Overview: Using Outlook Express 6 in a Managed Environment**

Although there are inherent risks associated with sending and receiving e-mail (and e-mail attachments), with Windows Server 2003 SP1, you can use several different features and configuration methods in Outlook Express 6 to reduce the risks:

- You can use the graphical user interface to configure the security-related features in Outlook Express 6. For more information, see "New Security-Related Features in Outlook Express 6," earlier in this section and "To Start Outlook Express 6 and View or Configure Security Settings," later in this section.

- You can use a Group Policy setting, **Block attachments that could contain a virus**, to limit the risk associated with e-mail attachments in Outlook Express 6. For more information, see "To Locate the Group Policy Setting for Blocking E-mail Attachments in Outlook Express 6," later in this section.

**Procedures for Working with Outlook Express 6**

This subsection provides procedures for the following:

- Opening the dialog box from which you can configure security settings for Outlook Express 6.

- Locating the Group Policy setting, **Block attachments that could contain a virus**.

  You can use this Group Policy setting in situations where you want Outlook Express 6 to be available but where you want to limit the risk associated with e-mail attachments. For more information about this policy setting, see "New Security-Related Features in Outlook Express 6," earlier in this section.

**To Start Outlook Express 6 and View or Configure Security Settings**

1. Click **Start**, point to **All Programs** or **Programs**, and then click **Outlook Express**.

2. On the **Tools** menu, click **Options**.

3. Click the **Security** tab and view or configure the settings, including the check boxes for the following options:

   - **Warn me when other applications try to send mail as me**.

   - **Do not allow attachments to be saved or opened that could potentially be a virus**.

   - **Block images and other external content in HTML e-mail**.

   You can also view or configure the security zones setting. Outlook Express 6 uses two of the same security zones that you configure in Internet Explorer 6. For more information about security zones, see Internet Explorer 6.0 and Resulting Internet Communication in Windows Server 2003 with Service Pack 1 in this white paper.

4. Click the **Read** tab, and view or configure the settings, including the check box for **Read all messages in plain text**.

**To Locate the Group Policy Setting for Blocking E-mail Attachments in Outlook Express 6**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **User Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Internet Explorer**.

   In the details pane, double-click **Configure Outlook Express**.

3. If you enable this policy, you can select or clear the check box for **Block attachments that could contain a virus**.

**Related Links**

- For more details about changes in the version of Outlook Express in Windows Server 2003 with SP1, see "Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1" on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46278

- For more information about security zones in Internet Explorer 6 (zones also used in Outlook Express 6), see Internet Explorer 6.0 and Resulting Internet Communication in Windows Server 2003 with Service Pack 1 in this white paper.

# Plug and Play and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Plug and Play
- How Plug and Play communicates with sites on the Internet
- How to control Plug and Play to prevent the flow of information to and from the Internet

**Benefits and Purposes of Plug and Play**

Windows Plug and Play provides ease of support for installing devices on computers in your network. You can simply plug in a Plug and Play device and Windows does the rest by searching locally for the driver and prompting to find out whether to also search Windows Update, installing the driver, updating the system, and allocating resources. After you install a Plug and Play device, the driver is configured and loaded dynamically, typically without requiring user input.

Plug and Play in Microsoft Windows Server 2003 with Service Pack 1 (SP1) provides the following functionality:

- Detects a Plug and Play device and determines its hardware resource requirements and device identification number (Plug and Play ID).
- Locates an appropriate device driver for newly installed devices.
- Allocates hardware resources.
- Dynamically loads, initializes, and unloads drivers.
- Notifies other drivers and applications when a new device is available.
- In conjunction with power management, handles stop and start processes for devices during hibernation, standby, and startup and shutdown operations.
- Supports a wide range of device types.

In order to install devices using the hardware wizards, you must be logged on as an administrator or a member of the Administrators group. You can then use the hardware wizards, such as the Hardware Update Wizard, to search the Windows Update site for device drivers. All drivers obtained through Windows Update are signed by Windows Hardware Quality Labs (WHQL). The WHQL provides compatibility testing services to test hardware and drivers for Windows operating systems.

| **Note** |
| --- |
| Some buses, such as Peripheral Component Interconnect (PCI) and universal serial bus (USB), take full advantage of Plug and Play. Older buses, such as Industry Standard Architecture (ISA), do not take full advantage of Plug and Play, and require more user interaction to ensure that devices are correctly installed. |

The Windows Update site is located at:

http://windowsupdate.microsoft.com/

**Overview: Using Plug and Play in a Managed Environment**

The Plug and Play feature is built into Windows Server 2003 and is always available. When a person who is logged on as an administrator installs a Plug and Play device, Windows Server 2003 first searches locally for an appropriate device driver. If the computer is connected to the Internet, Windows Server 2003 with SP1 prompts the person to find out whether to also search Windows Update for the latest device driver.

As an IT administrator in a highly managed network environment, you may want to control whether Windows Server 2003

with SP1 will search the Windows Update Web site for the latest device driver, and if so, whether the person installing or updating the device will be prompted before the Internet search begins. You can control these things through Group Policy.

There are also policy settings you can use to disable any access to Windows Update. If you do prevent access to Windows Update, there is the option for manually downloading the updates from the Windows Update Catalog, whereby they can be distributed on your organization's network as needed.

Using Group Policy to disable access to Windows Update, and to configure driver search locations, is described in "Controlling Automatic Device Updating to Prevent the Flow of Information to and from the Internet," later in this section.

**How Plug and Play Communicates with Sites on the Internet**

When a person logged on as an administrator installs new hardware, or updates the driver for existing hardware, Windows Server 2003 with SP1 will by default display a prompt to find out whether to search Windows Update for the latest device driver. If the person installing or updating the device consents to the Internet search, the interaction takes place as follows:

- **Specific information sent or received**: The Code Download Manager (CDM) calls Windows Update to find and download device drivers. The CDM also calls Help and Support Center, which logs Plug and Play IDs for devices that Microsoft does not have drivers for. Neither of these communications is under the direct control of Plug and Play. The CDM handles all of the communication between the computer and Windows Update. None of the communication between the computer and the Internet uniquely identifies the user.

- **Default setting**: Plug and Play is always available, and by default will prompt the person installing a device to find out whether to search the Windows Update Web site for the latest device driver.

- **Trigger and user notification**: When an administrator adds hardware or updates a driver on a computer, and the computer is connected to the Internet, by default, Windows Server 2003 with SP1 prompts to find out whether to search Windows Update for driver updates. The search is conducted if the administrator consents.

- **Logging**: If you use a Plug and Play driver with a device that is not Plug and Play, any associated issues or problems are recorded in the event log.

- **Encryption**: Data transfer is based on interaction with Windows Update. The data is transferred using HTTPS.

- **Transmission protocol and ports**: The transmission protocols and ports are HTTP 80 and HTTPS 443.

- **Ability to disable**: Plug and Play cannot be disabled, because system instability would result. You can disable access to Windows Update using Group Policy.

**Controlling Automatic Device Updating to Prevent the Flow of Information to and from the Internet**

Windows will automatically update device drivers using Plug and Play, and it will even search for compatible drivers for devices that are not Plug and Play. You therefore may want to exercise various levels of control over the ability of someone who logs on as an administrator to install new hardware and to update hardware devices and drivers.

You can use Group Policy to:

- Control whether Windows Update is included when Plug and Play searches for a device driver.

  This procedure is presented in the next subsection.

- Suppress the prompt that by default is displayed before Plug and Play begins searching the Windows Update Web site for a device driver. This setting only has an effect if you also use a setting to specify that Plug and Play will search the Windows Update Web site for device drivers.

  This procedure is presented in the next subsection.

- Turn off all access to Windows Update.

  If you turn off all access to Windows Update, it also means Plug and Play cannot search Windows Update. For more information about controlling access to Windows Update and for alternative approaches to updating such as Software Update Services, see the Windows Update, Automatic Updates, and Internet Communication section in this white paper.

**Procedure for Controlling Where Plug and Play Searches for Drivers**

When you install new hardware, Windows Server 2003 can potentially search four different locations for drivers in the following order: the hard drive, the floppy drive, the CD drive, and Windows Update. The default approach for Windows Server 2003 with SP1 is to search the first three locations, and then prompt you to find out whether to also search Windows Update. However, you can configure the driver search locations to remove selected locations.

This subsection includes procedures for configuring the following Group Policy settings:

- A setting that controls where Plug and Play searches for device drivers.
- A setting that specifically controls whether Plug and Play searches Windows Update for drivers. This setting is one in a collection of settings that control how various components communicate with the Internet,
- A setting that suppresses the prompt that by default is displayed before Plug and Play begins searching the Windows Update Web site for a device driver. This setting only has an effect if you also use a setting to specify that Plug and Play will search the Windows Update Web site for device drivers.

For additional procedures to configure policy settings for Windows Update, see the section Windows Update, Automatic Updates, and Internet Communication in this white paper.

**To Specify Driver Search Locations for Plug and Play Devices**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.
2. Click **User Configuration**, click **Administrative Templates**, and then click **System**.
3. In the details pane, double-click **Configure driver search locations**, and then click **Enabled**.
4. Select or clear check boxes to prevent or allow searching of floppy disk drives, CD-ROM drives, or Windows Update.

**To Disable Windows Update as a Driver Search Location**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.
2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.
3. In the **details** pane, double-click **Turn off Windows Update device driver searching**, and then click **Enabled**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the **Restrict Internet communication** policy setting, which is located in **Computer Configuration/Administrative Templates/System/Internet Communication Management**. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

**To Suppress the Prompt That is Displayed Before Windows Update is Searched for a Device Driver**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.
2. If you want the policy setting to apply to all users of a computer and to come into effect when the computer starts or when Group Policy is refreshed, click **Computer Configuration**. If you want the policy setting to apply to users and to come into effect when users log on or when Group Policy is refreshed, click **User Configuration**.

3. Click **Administrative Templates** and then click **System**.

4. In the details pane, double-click **Turn off Windows Update device driver search prompt**, and then click **Enabled**.

This setting only has an effect if you also use a setting to specify that Plug and Play will search the Windows Update Web site for device drivers.

**Related Links**
For more information about Windows Update, see the Windows Update Web site at:

http://windowsupdate.microsoft.com/

# Program Compatibility Wizard and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of the Program Compatibility Wizard.
- How the Program Compatibility Wizard communicates with sites on the Internet.
- How to control the Program Compatibility Wizard to prevent the flow of information to the Internet.
- How to obtain the Application Compatibility Toolkit. You can use the toolkit to improve the user experience with incompatible applications.

**Benefits and Purposes of the Program Compatibility Wizard**

There are some applications that work on earlier versions of Windows that might fail to function properly on Microsoft Windows Server 2003 with Service Pack 1 (SP1). This can happen for several reasons—an application may expect older formats of Windows data, may expect user information to be in specific locations or formats, or may not be designed for the tighter security implemented in Windows Server 2003 SP1. In addition, applications written exclusively for Windows 95, Windows 98, or Windows Millennium Edition may use direct hardware access, which can greatly reduce operating system stability. Because of its Windows NT heritage, Windows Server 2003 SP1 requires that hardware access be handled through the correct channels.

To enable a better user experience, Microsoft has integrated technologies for application compatibility into Windows Server 2003. These technologies are applied whenever an application is installed on the operating system, whether in the course of a system upgrade or during regular operations. Some of these technologies work automatically to apply compatibility fixes, while others can be selected by users or administrators. This section describes the Program Compatibility Wizard, which can be used to make setting adjustments for an incompatible application and run the application successfully.

This section also provides a description of the Application Compatibility Toolkit, which you can use to locate and address compatibility problems. For example, with one of the tools, you can create custom messages that notify users of the problems with an incompatible application and redirect users to your intranet site.

**Overview: Using the Program Compatibility Wizard in a Managed Environment**

IT administrators who want to get an application to work quickly, without addressing compatibility for the application throughout the organization, may choose to use the Program Compatibility Wizard. You can use the wizard in situations where you want to determine quickly whether the prepackaged compatibility fixes can resolve problems you encounter, particularly if you are working on a computer where the Application Compatibility Toolkit is not installed. For more information about the Application Compatibility Toolkit, see "Using the Application Compatibility Toolkit to Improve the User Experience with Incompatible Applications," later in this section.

In Windows Server 2003 you can access the Program Compatibility Wizard by default through **Programs\Accessories** or **All Programs\Accessories**. At the completion of the wizard you are presented with "Program Compatibility Data" that states, "Microsoft has created temporary files that contain information about the settings you selected and whether the problems were fixed. Sending this information to Microsoft will help us improve program compatibility." You can then choose to send this information to Microsoft or not. You can also choose to click a link to a Web site to view the data collection policy and you can choose to view the temporary files that will be sent. Allowing users or administrators to do this, however, may present a privacy problem for highly managed organizations.

| **Note** |
| --- |
| As an alternative to running the Program Compatibility Wizard, you can set the compatibility properties for an application manually through the **Compatibility** tab of a program's **Properties** sheet. To do this, right-click the program icon, click **Properties**, click **Compatibility**, and then change the compatibility settings for your application. |

You can use Group Policy to control where data collected by the Program Compatibility Wizard is sent. You can prevent data transfer to the Internet by using Group Policy settings related to error reporting and you can have data from the wizard sent to a server on your intranet instead of to Microsoft. For more information about these procedures, see the section of this white paper titled Windows Error Reporting and Internet Communication.

#### How the Program Compatibility Wizard Communicates with Sites on the Internet

Although you can control information sent by the Program Compatibility Wizard, it is designed to communicate over the Internet to expedite problem solving. This subsection lists details of the communication process:

- **Specific information sent or received**: The results of the Program Compatibility Wizard data, including settings and problems that were encountered with the application being installed, are sent to Microsoft. The user is not uniquely identified.

- **Default and recommended settings**: Use of the Program Compatibility Wizard is enabled by default. Recommended settings are discussed in the next subsection, "Controlling Program Compatibility Wizard Data to Prevent the Flow of Information to the Internet."

- **Trigger and notification**: In the last dialog box of the wizard, users are asked if they want to send information to Microsoft. Data is not sent automatically.

- **Logging**: There is no information related to the Program Compatibility Wizard entered into the event log.

- **Encryption**: HTTPS is used to perform the data transfer to Microsoft.

- **Access**: The Microsoft product group has access to the raw data only.

- **Privacy**: The privacy statement is the same as that associated with Windows Error Reporting (WER) data. A link to the privacy statement on the Web is provided in the wizard. This privacy statement is available at:

  http://go.microsoft.com/fwlink/?LinkId=825

- **Transmission protocol and port**: The transmission protocol used is HTTPS and the port is 443.

- **Ability to disable**: You cannot disable the Program Compatibility Wizard. Using Group Policy, you can prevent data from being sent to the Internet.

For more information about the type of information that is sent to Microsoft, how the data is used, encryption, and the privacy statement, see the section of this white paper titled Windows Error Reporting and Internet Communication.

#### Controlling Program Compatibility Wizard Data to Prevent the Flow of Information to the Internet

Using Group Policy, you can configure **Configure Error Reporting** policy setting to prevent data collected by the Program Compatibility Wizard from being sent to Microsoft. By using configuration options within error reporting you can have the data sent to a server on your intranet instead of to Microsoft. When you configure error reporting this way, you activate Corporate Error Reporting (CER).

The **Configure Error Reporting** policy setting is located in **Computer Configuration\Administrative Templates\System\Error Reporting**. For more information and procedures for configuring error reporting, see the section of this white paper titled Windows Error Reporting and Internet Communication.

If you use this approach for reporting errors, the user experience with the Program Compatibility Wizard does not change. The dialog box that presents the option of sending data to Microsoft is the same. If the user selects **Yes**, the data is sent to the designated server on your intranet.

#### Using the Application Compatibility Toolkit to Improve the User Experience with Incompatible Applications

When a user tries to run a low-level application—such as an antivirus or disk-access utility—that is known to be incompatible and compromise system integrity, Windows XP blocks the application and informs the user about it. To do this, Windows XP uses information in databases stored locally on the computer. Compatibility fixes are contained in a database file named SYSMAIN.SDB. The warning information used when an application cannot be run successfully is contained in a related database file, APPHELP.SDB.

Before Windows XP Service Pack 2 (SP2), the operating system would also contact a Microsoft Web site for the latest information about incompatible applications. As of SP2, Windows XP only uses the information stored locally on the computer.

You can customize the way Windows XP responds to programs that are known to compromise system integrity by using the Application Compatibility Toolkit. For example, you can use one of the tools in the toolkit, the Compatibility Administrator tool, to create custom messages that notify users of the problems with an incompatible application and redirect users to your

intranet site. To do this, you need to first download the Application Compatibility Toolkit, and then use tools such as the Compatibility Administrator tool.

**To Read About and Download the Application Compatibility Toolkit**

1. Read about the toolkit and find links for downloading it from the Windows Web site at:

   http://go.microsoft.com/fwlink/?LinkId=29880

2. Follow the installation instructions. Once you have installed the toolkit, you can view the Windows Application Compatibility 3.0 Reference and you can run the Compatibility Administrator tool to make the changes you need.

**To Create Custom Messages for Applications with Incompatibility Problems**

1. Make sure that the Application Compatibility Toolkit is installed by using the previous procedure.

2. Click **Start**, point to **Programs** or **All Programs**, point to **Microsoft Windows Application Compatibility Toolkit**, and then click **Compatibility Administrator Tool 3.0**.

3. In the console tree, click **Custom Databases**, and then click **New Database**.

4. On the toolbar, click **AppHelp**. The **Create a custom AppHelp message** dialog box appears.

5. Enter information as prompted in the dialog box.

6. Save the new database file.

| **Note** |
| --- |
| When you have completed your entries and saved the file, you can deploy your changes to multiple computers running Windows XP. See "Deploying Compatibility Fixes," in Compatibility Administrator Help. |

**Event Logging for Events Related to the Blocking of Incompatible Applications**

Events related to the blocking of incompatible applications are not logged by default. You can enable event logging in the Application log (**Control Panel\Administrative Tools\Event Viewer**) by using Group Policy.

**To Enable Event Logging for Events Related to the Blocking of Incompatible Applications**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Application Compatibility**.

3. In the details pane, double-click **Turn On Application Help Log Events**, and then click **Enabled**.

**Related Links**

- For more information about application compatibility resources, see “Windows Application Compatibility” at:

  http://go.microsoft.com/fwlink/?LinkId=29880

# Remote Assistance and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Remote Assistance
- How Remote Assistance communicates with sites on the Internet
- How to control Remote Assistance to prevent the flow of information to and from the Internet

**Benefits and Purposes of Remote Assistance**

With Microsoft Windows Server 2003, users and administrators in your organization can use Remote Assistance to get help from a member of your support staff. Users or administrators can also collaborate in other ways through screen sharing. Remote Assistance is a convenient way for support professionals to connect to a computer from another computer running a compatible operating system, such as Windows Server 2003 or Windows XP, and to show the users or administrators a solution to the problem.

Using Windows Messenger Service or an e-mail program, such as Microsoft Outlook or Outlook Express, you can provide support to a user by connecting to the user's computer. After you are connected, you can view the user's computer screen, communicate in real time about what you both see on the user's computer, send files, use voice communication, and use your mouse and keyboard to work on the user's computer.

**Overview: Using Remote Assistance in a Managed Environment**

On Windows Server 2003, with or without Service Pack 1 (SP1), Remote Assistance is disabled by default. By opening Control Panel\System and clicking the Remote tab, you can enable Remote Assistance (unless it is disabled through Group Policy). Then you can open Help and Support Center, click **Remote Assistance**, and then click **Invite someone to help you**. When you do this, you can have someone inside or outside your network connect to your server.

Although a firewall on your organization's network will likely prevent outsiders from connecting directly to a computer on your intranet, some firewall settings might permit users or administrators to connect remotely to someone within your intranet or outside your network through Remote Assistance. As an administrator in a highly managed environment, you might want to prevent the use of this feature. You can do this during your deployment of Windows Server 2003 with SP1, or post-deployment using Group Policy.

In a domain environment there is also the option of a support person or IT administrator offering unsolicited assistance. From Help and Support Center using **Tools\Help and Support Center Tools\Offer Remote Assistance**, an administrator in the domain may offer assistance to users in the same domain or trusted domains without being asked. However, users can decline the invitation. This capability can be strictly controlled with Group Policy. Controlling the use of unsolicited as well as solicited Remote Assistance is described further in the subsection "Controlling Remote Assistance to Prevent the Flow of Information to and from the Internet."

**How Remote Assistance Communicates with Sites on the Internet**

When a user (referred to as the "novice") initiates a request for assistance through either the e-mail option or the **Save invitation as a file** option in Remote Assistance, the operating system starts Help and Support Center. Help and Support Center then passes the information to Remote Assistance.

When the person who is being contacted (the "expert") accepts the invitation from the novice, Remote Assistance calls Help and Support Center application programming interfaces (APIs) to initiate the session. Help and Support Center relies on Terminal Services to negotiate the session. Help and Support Center passes the Remote Assistance invitation (the "ticket") file to Terminal Services. The Remote Assistance session is established using RDP (Remote Desktop Protocol), and uses port 3389 through Terminal Services on the novice computer.

There are safeguards built into the Remote Assistance feature. All sessions are encrypted and can be password-protected. The novice (user soliciting the assistance) sets the maximum time for the duration of the ticket. Also, firewalls on your organization's network can be configured to prevent communication associated with Remote Assistance, for example, Remote Assistance connections that are inbound to computers behind the firewall.

The following information presents additional details on how information transfer over the Internet takes place when a connection is made:

- **Specific information sent or received**: Information that is transmitted in a Remote Assistance ticket includes user name, IP address, and computer name. Information necessary to provide functionality for Remote Assistance (for example, screen sharing, file transfer, and voice) is sent in real time using point-to-point connections.

- **Default and recommended settings**: Anyone with access to Help and Support Center can access the Remote Assistance feature. Users can prevent someone from connecting to their computer by declining an invitation. You can also prevent someone from remotely controlling a server running Windows Server 2003 SP1 by using Control Panel settings or Group Policy.

- **Triggers**: A user or administrator establishes contact with the expert by sending an invitation through e-mail, instant messaging, or by saving an invitation as a file and transferring it manually, such as on a floppy disk, to the expert. Or, an expert offers unsolicited assistance to a user.

- **User notification**: Whether assistance is solicited or unsolicited, the novice is notified of the offer of assistance from the expert. The novice must accept the connection before Remote Assistance begins.

- **Logging**: Events such as a person initiating a connection or a user or administrator accepting or rejecting an invitation are recorded in the event logs. Windows Server 2003 Service Pack 1 (SP1) records more details than were recorded previously, including events such as taking and releasing control. sending and accepting files, and ticket creation and deletion. SP1 also records details such as whether assistance is solicited or unsolicited as well as more detailed user name and IP address information.

- **Encryption**: The RDP (Remote Desktop Protocol) encryption algorithm for the main Remote Assistance communication and the RTC (Real-Time Communication) encryption algorithm for voice are used. The RDP encryption algorithm is RC4 128-bit.

- **Access**: No information is stored at Microsoft.

- **Transmission protocol and port**: The port is 3389 and the transmission protocols are RDP and RTC. For Offer Remote Assistance, Distributed Component Object Model (DCOM) is also used.

- **Ability to disable**: This component can be disabled by using Group Policy or locally through Control Panel.

- **Firewall protection**: Any firewall that blocks port 3389 will not allow a Remote Assistance connection to users outside the firewall. This does not prevent users from within the network protected by the firewall from connecting to each other. If you close port 3389, you will block all Remote Desktop and Terminal Services events through it as well. If you want to allow these services but want to limit Remote Assistance requests, use Group Policy. If the port is opened only for outbound traffic, a user can request Remote Assistance by using Windows Messenger.

| **Note** |
| --- |
| Windows Server 2003 with SP1 includes enhancements to the firewall component, now called Windows Firewall. To use Remote Assistance with Windows Firewall enabled, in Windows Firewall, on the **Exceptions** tab, select **Remote Desktop**. Alternatively, you can configure Windows Firewall through Group Policy with a setting at Computer Configuration\Administrative Templates\Network\Network Connections\Windows Firewall, either in Domain Profile or in Standard Profile. The setting to configure is **Windows Firewall: Allow Remote Desktop exception**. |

For more information about the Remote Assistance connection process, see article 300692, "Description of the Remote Assistance Connection Process" in the Microsoft Knowledge Base at:

http://go.microsoft.com/fwlink/?LinkId=29212

**Controlling Remote Assistance to Prevent the Flow of Information to and from the Internet**
Administrators can control the use of Remote Assistance in the following ways:

- Group Policy to prevent Remote Assistance from being solicited from this computer

- Group Policy to prevent unsolicited Remote Assistance from being offered to this computer

- Local control of Remote Assistance through Control Panel

Group Policy settings are described in detail in this subsection. Procedures for disabling Remote Assistance are presented in the next subsection.

**Using Group Policy**

There are two Group Policy settings you can configure to control the use of Remote Assistance:

- **Solicited Remote Assistance**. Use this policy setting to determine whether Remote Assistance can be solicited from a given computer. In Solicited Remote Assistance the user of a computer explicitly requests help from another party.

- **Offer Remote Assistance**. Use this policy setting to determine whether a support person or IT administrator (expert) can offer remote assistance to a computer without a user explicitly requesting it first through e-mail, a file, or instant messaging.

These policy settings are located in Computer Configuration\Administrative Templates\System\Remote Assistance. Configuration options for these policy settings are described in the following table.

**Group Policy Settings for Controlling Remote Assistance**

| Policy Setting | Description |
|---|---|
| Solicited Remote Assistance (enabled) | When this policy setting is enabled, a user can create a Remote Assistance invitation that a person ("expert") can use at another computer to connect to the user's computer. If given permission, the expert can view the user's screen, mouse, and keyboard activity in real time.<br><br>Additional configuration options are available when you enable this policy setting. |
| Solicited Remote Assistance (disabled) | If the status is set to **Disabled**, users cannot request Remote Assistance and this computer cannot be controlled from another computer. |
| Solicited Remote Assistance (not configured) | If the status is set to **Not Configured**, the configuration of solicited Remote Assistance is determined by the Control Panel settings. |
| Offer Remote Assistance (enabled) | When this policy setting is enabled, a remote user or administrator can offer Remote Assistance to the computer. When you configure this policy setting, you have two choices: you can select either **Allow helpers to only view the computer** or **Allow helpers to remotely control the computer**. In addition to making this selection, when you configure this policy setting, you also specify the list of users or user groups that will be allowed to offer remote assistance. Administrators of this computer can offer remote assistance to it by default. They do not need to be added to the list. |
| Offer Remote Assistance (disabled or not configured) | If you disable or do not configure this policy setting, users or groups cannot offer unsolicited remote assistance to this computer. |

For additional configuration options, see the Remote Assistance policy settings in Group Policy. To find more information about editing Group Policy, see Appendix B: Resources for Learning About Group Policy.

**Procedures for Disabling Remote Assistance**

This section presents procedures administrators can use for disabling Remote Assistance through Group Policy or Control Panel.

**To Disable the Use of Remote Assistance Using Group Policy**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, and then click **Remote Assistance**.

3. In the details pane, double-click **Solicited Remote Assistance**, click **Disabled**, and then click **OK**.

4. In the details pane, double-click **Offer Remote Assistance**, click **Disabled**, and then click **OK**.

**To Disable the Use of Remote Assistance Through Control Panel**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **System**.

3. In **System Properties**, click the **Remote** tab.

4. Under **Remote Assistance**, clear the check box labeled **Turn on Remote Assistance and allow invitations to be sent from this computer**.

# Search Companion and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Search Companion
- How Search Companion communicates with sites on the Internet
- How to control Search Companion to prevent the flow of information to and from the Internet

**Benefits and Purposes of Search Companion**
The Microsoft Search Companion Web service enables you to search for files and folders on your computer, to search for files, people, and other computers on your internal network, and to search for information on the Internet. Search Companion uses Indexing Service to maintain an index of all the files on the computer, making searches faster.

When employing Search Companion, you can specify several search criteria. For example, you can search for files and folders by name, type, or size. You can find files based on when they were last modified, or search for files containing specific text. When searching for information on the Internet, Search Companion enables you to enter search queries in natural language (meaning informal, or conversational language). It then suggests the best way to conduct the search, and sends the query to Internet services that are most likely to yield positive results.

**Overview: Using Search Companion in a Managed Environment**
When you search the Internet using Search Companion, the following information is collected:

- The text of the Internet search query
- Grammatical information about the query
- The list of tasks (suggestions) that the Search Companion Web service recommends to refine the search
- Any tasks you selected from the recommendation list

Search Companion does not collect:

- Personal information, such as user name or e-mail address
- Demographic information

Microsoft does not use the information it collects to identify you individually or associate such information with other data sources that may contain personal data. Microsoft does not collect information when you search on the local system, LAN, or intranet.

The Search Companion Web service is designed to upgrade automatically as new features become available. It therefore uses the Internet connection periodically to check for and replace necessary files. You can use Group Policy to prevent the Search Companion Web service from upgrading automatically from the Internet.

If you want to disable the Search Companion Web service, you can do so by changing to Classic Search for the Internet. Microsoft Windows does not collect any query information when Classic Search is used. You can also disable Search Companion (change to Classic Search) by modifying the registry settings. The procedures for both of these methods are described later in this section of the white paper.

**How Search Companion Communicates with Sites on the Internet**
Search Companion in Microsoft Windows Server 2003 improves the search process by consolidating search tasks, optimizing searches for the most common scenarios, and offering suggestions for refining the search.

The form that collects search criteria will post information to an ASP page that displays the search results. The search pages use a combination of XML and Microsoft Visual Basic® development system, Scripting Edition (VBScript) and Microsoft

JScript® development software for accessing the search objects. Because the script is run on the server, you can view the search pages from any browser.

Search Companion uses XML files to define both the user interface and some functional parameters of its tasks (for example, what list of file extensions constitutes the "Music" category of files). The first time in each Search Companion session that an XML file is referenced, by default, Search Companion checks to see if a later version of that XML file is available from sa.windows.com. The "check" is really a file download request, conditioned on the modified date of the file. If there is a later version of the XML file, Search Companion downloads it and replaces the earlier version. The XML files are located in a language-specific subfolder of *systemroot*\Srchasst\, and if the current user does not have administrative credentials, the old XML file cannot be overwritten. The downloading of these XML files can be prevented through Group Policy.

This subsection describes various aspects of the data that is sent to and from the Internet through Search Companion, and how the exchange of information takes place:

- **Specific information sent or received**: When you search the Internet using Search Companion, the following information is collected regarding your use of the service: the text of your Internet search query, grammatical information about the query, the list of tasks that the Search Companion Web service recommends, and any tasks you select from the recommendation list.

- **Default and recommended settings**: Search Companion is enabled by default.

- **Triggers**: The user selects **Start\Search** and uses search options to search the Internet. The user can also start Internet Explorer and use a toolbar button, shortcut key, or menu to search the Internet.

- **User notification**: There is no provision in Search Companion for user review or notification of data sent.

- **Uniquely identify users**: The user is not uniquely identified. Session-based cookies are used to maintain state information, but these randomly assigned GUIDs do not persist across browser sessions.

- **Logging**: No information is collected when you search your local system, LAN, or intranet. The only "storage" is the Internet Information Services (IIS) log of the file request on the server at Microsoft that provides the Search Companion Web service. Search Companion does not record your choice of Internet search engines, and it does not collect or request any demographic information or personal information (such as user name or e-mail address).

- **Encryption**: There is no encryption of data.

- **Access**: No user information is collected. The IIS logs (described in the "Logging" item, earlier in this list) are cycled annually, that is, logs are retained for twelve months, and discarded in the thirteenth month following collection.

- **Privacy**: The privacy statement is located at the following Web site:

  http://sa.windows.com/privacy/

- **Transmission protocol and port**: The transmission protocol is HTTP and the port is 80.

- **Ability to disable**: The feature can be disabled by changing to Classic Search.

**Controlling Search Companion to Prevent the Flow of Information to and from the Internet**

You can disable the Search Companion Web service by changing preferences to Classic Search for the Internet. You can also disable Search Companion by changing the registry settings manually. In addition, you can prevent Search Companion from checking for and downloading updated versions of the XML files that it uses. Procedures for all of these approaches are provided in the following subsection.

**Procedures for Configuration of Search Companion**

Search Companion can be configured in several ways as described previously.

**To Change to Classic Search for the Internet through the User Interface**

1. Click **Start**, and then either click **Search**, or point to **Search** and click **On the Internet**.

2. Click **Change preferences**.

3. Click **Change Internet search behavior**.

4. Click **With classic Internet search** and click **OK**.

5. On the File menu, click **Close**.

The next Internet search you perform will use the preference you specified.

**To Change to Classic Search for the Internet Through the Registry Key**

1. Close Internet Explorer (all instances).

2. Open Registry Editor by clicking **Start**, clicking **Run**, and then typing **regedit**.

| **Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied. |

3. Navigate to **HKEY_CURRENT_USER\Software\Microsoft\Internet Explorer\Main**.

4. Look for an entry called **Use Search Asst**. If it exists, skip to step 7.

5. On the **Edit** menu, point to **New**, and then click **String value**.

6. Type **Use Search Asst** as the name for the new value (the type is REG_SZ), and then press **ENTER**.

7. Click **Use Search Asst**, and then on the **Edit** menu, click **Modify**.

8. For **Value data**, type:

   **no**

| **Note** |
| --- |
| Type the entry and value exactly as shown, including spaces and capitalization. |

**To Prevent Search Companion from Downloading Updated Versions of the XML Files That It Uses**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.

3. In the details pane, double-click **Turn off Search Companion content file updates**, and then click **Enabled**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the Restrict Internet communication policy setting, which is located in **Computer Configuration/Administrative Templates/System/Internet Communication Management**. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

# Terminal Server Licensing and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The purposes of Terminal Server Licensing
- How Terminal Server Licensing communicates with sites on the Internet
- How to control Terminal Server Licensing to limit the flow of information to and from the Internet

| Note |
| --- |
| In Microsoft Windows Server 2003, several components that use technology related to Terminal Server have different names than they had in Windows 2000. Terminal Server was previously known as Terminal Services in Application Server mode. Remote Desktop for Administration was previously known as Terminal Services in Remote Administration mode. |

**Purposes of Terminal Server Licensing**

Terminal Server Licensing is an administrative tool that helps manage your Terminal Server licenses. Terminal Server Licensing can be installed on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. If you install Terminal Server, you must also install Terminal Server Licensing (although not necessarily on the same computer). However, there is a grace period of 120 days during which you can use Terminal Server even though you have not set up license servers yet and have not installed client license key packs.

| Note |
| --- |
| Terminal Server and Terminal Server Licensing are not available in Windows Server 2003, Web Edition. Remote Desktop for Administration, however, is available in Windows Server 2003, Web Edition. |

When you set up Terminal Server on a server, users can access applications running on that server, which means they can run the applications you provide even if they are working from hardware that might not otherwise support the applications.

Microsoft has introduced new licensing options to address business customer needs and to complement the technical capabilities of Windows Server 2003. For details about these changes, see the Windows Server 2003 licensing overview at:

http://go.microsoft.com/fwlink/?LinkId=46007

**Overview: Using Terminal Server Licensing in a Managed Environment**

You can control the communication that occurs between the Terminal Server Licensing component and sites on the Internet by choosing the server or servers on which to install the Terminal Server Licensing component, and by choosing among three methods for activation. (You also choose among the same three activation methods for obtaining client license key packs, which are digital representations of a group of client access licenses.) The three methods for activation are as follows:

- **Automatic**. This method requires the server running Terminal Server Licensing to have an Internet connection.
- **Web Browser**. This method requires an Internet connection, but the connection can be made from any computer, not just the computer that is running Terminal Server Licensing.
- **Telephone**. This method is used to contact the nearest Customer Support Center and receive an ID number to activate the Terminal Server License Server by phone.

For more details on these methods, see the subsections that follow.

**How Terminal Server Licensing Communicates with Sites on the Internet**

A server running the Terminal Server Licensing component communicates with the Microsoft Clearinghouse (a database for managing licensing) on the Internet only when you activate Terminal Server Licensing or when you initiate subsequent transactions with Microsoft to obtain client license key packs. The following list describes the communication that occurs when

you activate or obtain client license key packs directly over the Internet or when you connect to the Microsoft Clearinghouse from a computer other than the one where Terminal Server Licensing is installed.

| **Note** |
| --- |
| The information in the following list applies only to activation or obtaining client license key packs over the Internet. It does not apply when you activate by phone.<br><br>Activation by phone is done as follows: When you activate by phone, you start the licensing wizard, choose a country or region from the list that is displayed, and call the number shown. |

The rest of this subsection describes various aspects of the data that is sent to and from the Internet from Terminal Server Licensing and how the exchange of information takes place.

- **Specific information sent or received**: The information sent to the Microsoft Clearinghouse includes company name, first and last name of the user, license server name, and license server ID. Client license key packs are returned to the Terminal Server License Server.

- **Default and recommended settings**: Terminal Server Licensing is not installed by default.

- **User notification and triggers**: The administrator triggers activating, obtaining client license packs, or deactivating Terminal Server Licensing by performing the steps described in the subsection "Procedures for Configuring Terminal Server Licensing." When the Terminal Server Licensing component starts, the administrator is notified that activating, obtaining client license packs, or deactivating will initiate communication with Microsoft.

- **Logging**: Terminal Server Licensing logs events in the system log. The events can be viewed through Event Viewer.

- **Encryption**: Terminal Server Licensing uses the HTTP protocol over SSL (Secure Sockets Layer) to communicate on the Internet and the Web.

- **Access**: The Microsoft Clearinghouse is the database Microsoft maintains to activate license servers and to issue client license key packs. Microsoft customer service representatives have access to the licensing information and are able to successfully re-create the information on your Terminal Server License Server if technical problems occur. The information you provide might also be used internally at Microsoft to perform aggregate quality testing of the Terminal Server Licensing program.

- **Privacy statement**: To see the privacy statement for Terminal Server Licensing, click Microsoft's Terminal Server Licensing privacy policy at the following Terminal Services Web site:

  http://go.microsoft.com/fwlink/?LinkId=46192

- **Transmission protocol and port**: HTTPS over port 443, and remote procedure call (RPC) over port 135.

- **Ability to disable**: Terminal Server Licensing is not installed by default. Once installed, however, it can be disabled by the procedures described later in this section.

**External Connector license**

In place of individual Terminal Server client access licenses (TS CALs), you have the option of purchasing the Terminal Server External Connector license. This license enables external users to access a company's terminal servers without the need to purchase individual TS CALS for them or their devices. Generally speaking, an external user is a person who is not an employee or similar personnel of a company or its affiliates, and who is not a customer of hosted services (for specific information, see the External Connector license).

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running servers that communicate with the Internet. For more information about security and the Internet, see the introduction to this white paper, or see the Microsoft TechNet Web site at:

http://go.microsoft.com/fwlink/?LinkId=29482

**Controlling Terminal Server Licensing to Limit the Flow of Information to and from the Internet**

You can control the communication that occurs between the Terminal Server Licensing component and sites on the Internet in the following ways:

- **Install the Terminal Server Licensing component on selected servers only**. This follows the basic principle of stopping unnecessary services and keeping computers (especially servers) free of unnecessary software. For information about choosing which computer or computers on which to install Terminal Server Licensing, see the *Microsoft Windows Server 2003 Deployment Kit* on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46078

  When searching in documentation about Terminal Server Licensing, use the phrases "domain license server" and "enterprise license server" to discover information about these two options.

- **Review the method you want to use for activating Terminal Server Licensing before starting the licensing wizard**. Automatic activation is the fastest method, but if you prefer, you can activate by the other methods mentioned previously (connecting to a Web site from a computer other than the one where Terminal Server Licensing is installed, or activating by phone).

  You are required to activate a license server before it can issue licenses to Terminal Server clients. You are required to activate a license server only once. When you activate the license server, Microsoft provides the server with a limited-use digital certificate that validates server ownership and identity. Microsoft uses the X.509 industry standard certificate for this purpose. Using this certificate, a license server can make subsequent transactions with Microsoft and receive client license key packs.

**Procedures for Configuring Terminal Server Licensing**

Terminal Server Licensing servers can be configured in several ways as described previously. This subsection provides procedures for:

- Installing and uninstalling Terminal Server Licensing
- Activating Terminal Server Licensing
- Deactivating Terminal Server Licensing
- Viewing Help for Terminal Server and Terminal Server Licensing

**To Install or Uninstall Terminal Server Licensing**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components** (on the left).
4. Select **Terminal Server Licensing**, click **Next**, and then do one of the following:
   - If Terminal Server Licensing is installed and you want to remove it, clear the check box for **Terminal Server Licensing** and complete the licensing wizard.
   - If Terminal Server Licensing is not installed and you want to add it, select the check box for **Terminal Server Licensing** and then click **Next**.
5. In Terminal Server Licensing setup, do one of the following:
   - If your network includes several domains, click **Your entire enterprise**, and then provide the database location. An enterprise license server can serve terminal servers on any Windows Server 2003 domain.

- If you want to maintain a separate license server for each domain, or if your network includes workgroups or Windows NT 4.0 domains, click **Your domain or workgroup**, and then provide the database location.

6. Follow the instructions to complete the Windows Components Wizard.

**To Activate a License Server**

As mentioned previously in this section, you must activate a Terminal Server License Server before it can issue licenses to Terminal Server clients. Use Terminal Server Licensing to activate a Terminal Server License Server through the Microsoft Clearinghouse.

You can find procedure checklists along with complete instructions for configuring and activating Terminal Server License Servers in Help and Support Center. For more detailed instructions, see the procedure below titled, "To view Help for Terminal Server and Terminal Server Licensing" and "To view Help topics specific to Terminal Server Licensing."

**Deactivating a license server**

If Terminal Server and Terminal Server Licensing are already installed and you want to deactivate the license server, use one of the following procedures. You might need to deactivate a license server when the certificate of the server has expired, when the server becomes corrupted, or when the server is being redeployed. Note that when a license server's registration has expired, you are prompted to reactive the license server (not deactivate it). When you deactivate a license server, you will not be able to license additional clients from this server until the license server is reactivated.

**To Deactivate a License Using the Automatic Method**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**.

3. In the console tree, right-click the license server you want to deactivate, point to **Advanced**, and then click **Deactivate Server**. The licensing wizard starts.

4. In **Information Needed**, confirm that your name, phone number (optional), and e-mail address (required if you are using the Internet method) are correct, and then click **Next**.

5. Your request to deactivate the license server is sent to Microsoft where it is processed.

| **Note** |
| --- |
| The information sent to the Microsoft Clearinghouse during deactivation is the same information sent during activation, which includes company name, first and last name of the user, license server name, and license server ID. |

6. Click **Finish**.

**To Deactivate a License Using the Telephone Method**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**.

3. In the console tree, right-click the Terminal Server License Server you want to deactivate, point to **Advanced**, and then click **Deactivate Server**. The Terminal Server License Server Wizard starts.

4. Select **Telephone** and then click **Next**.

5. Specify your location and then click **Next**.

6. Call the telephone number displayed in the wizard, and give the customer support representative the product ID that is displayed below the telephone number.

7. Type the 35-digit confirmation code provided by the customer support representative in the boxes in the wizard, and then click **Next**.

8. Click **Finish**.

| **Note** |
| --- |
| Deactivating a license server does not remove the component. As mentioned previously, follow the basic principle of stopping unnecessary services and keeping computers (especially servers) free from unnecessary software by removing the Terminal Server Licensing component if you no longer plan to use it as a license server. |

| **Note** |
| --- |
| You cannot deactivate a license server using either the fax or Internet connection methods. |

**To View Help for Terminal Server and Terminal Server Licensing**

1. Click **Start** and then click **Help and Support**.

2. Click **Software Deployment** and then expand **Terminal Services**.

**To View Help Topics Specific to Terminal Server Licensing**

1. Click **Start** and then click **Help and Support**.

2. Click **Software Deployment** and then expand **Terminal Services**.

3. Expand **Terminal Server** and then expand **Checklists: Setting up Terminal Server**.

   When you are finished viewing appropriate checklists, collapse the items you were viewing.

4. Expand **Terminal Server Licensing**.

5. To get information about the following procedures, expand **How To**:

   - Install Terminal Server Licensing
   - Activate Terminal Server Licensing server
   - Install client license key packs
   - Deactivate a Terminal Server Licensing server
   - Reactivate a Terminal Server Licensing server
   - Repeat the installation of a client license key pack
   - Connect to a Terminal Server Licensing server
   - Change Terminal Server Licensing server properties

**Related Links**

- For details about new licensing options for Windows Server 2003, as well as information about the licenses necessary

with Terminal Server, see the Windows Server 2003 licensing overview at:

http://go.microsoft.com/fwlink/?LinkId=46007

- For descriptions of other ways you can use Terminal Server, see the overviews on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=24427

- For information about controlling the issuance of Terminal Server licenses, on a server running Windows Server 2003, open Help and Support Center and search for the topic "To control the issuance of Terminal Server licenses."

# Windows Error Reporting and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Windows Error Reporting
- How Windows Error Reporting communicates with sites on the Internet
- How to control Windows Error Reporting to prevent the flow of information to and from the Internet

**Benefits and Purposes of Windows Error Reporting**

The Windows Error Reporting feature in Microsoft Windows Server 2003 provides a service that allows Microsoft to track and address errors relating to the operating system, Windows components, and applications. This service, called the Error Reporting service, gives administrators and users the opportunity to send data about errors to Microsoft and to receive information about them. Microsoft developers can use the Error Reporting service as a problem-solving tool to address customer problems in a timely manner and to improve the quality of Microsoft products.

When users or administrators send information to Microsoft, in some cases Microsoft may provide information in return, such as a way to work around a problem or a link to a Web site for updated drivers, patches, or Microsoft Knowledge Base articles.

**Overview: Using Windows Error Reporting in a Managed Environment**

In Windows Server 2003 with Service Pack 1 (SP1), error reporting is enabled by default and you can choose to report errors to Microsoft. When an error occurs, a dialog box is displayed allowing you to report the problem. If you choose to report the problem, technical information about the problem is collected and then sent to Microsoft over the Internet. No information is sent unless you confirm that the error report is to be sent to Microsoft. When you are logged on as an administrator, you can choose to report system and application errors. When you are not logged on as an administrator you can choose to report application errors.

On Windows Server 2003 with SP1, you can configure or disable error reporting through **Control Panel\System\Advanced**. You can configure error reporting to send specified information such as system errors (Stop errors) only, or errors for Windows components, such as Windows Explorer or Microsoft Internet Explorer. You can also send information for applications, such as Microsoft Word.

Since error reporting is a valuable service, we recommend that IT administrators not disable it, but that they control what information is reported and where it is sent. For an organization where privacy is a concern, we recommend that the IT department review and filter error reports before they are sent to Microsoft. Though it is not recommended, you can also completely disable error reporting by using Group Policy.

The best method to use to prevent the automatic flow of error reporting information to and from the Internet is to redirect error reports to a server on your intranet by using Group Policy and to set up Corporate Error Reporting (CER).If you have Software Assurance with your volume license, you can use the Corporate Error Reporting tool to manage error reports that have been redirected to a network server. You use the tool to review the redirected error reports and then filter the reports that are sent to Microsoft based on your policies and the data contained within the error report. The tool is also useful for determining the types of problems users are experiencing most often.

If you have not yet deployed Windows Server 2003 with SP1, you can use unattended installation files to configure error reporting in the same way as in Group Policy. If it is necessary in your organization to completely disable Windows Error Reporting, you can do so with the unattended installation file or with Group Policy. For more information about these methods, see "Controlling Error Reporting to Prevent the Flow of Information to and from the Internet," later in this section.

**How Windows Error Reporting Communicates with Sites on the Internet**

The data that Microsoft collects is used strictly for the purpose of tracking down and solving problems that users are experiencing. The information is stored on servers at Microsoft with limited access that are located in controlled facilities. This subsection describes various aspects of the data that is sent to and from the Internet during error reporting, and how the exchange of information takes place.

- **Specific information sent or received**: For Windows Server 2003 with SP1, Microsoft collects various types of information related to two types of errors, user mode or application errors, and kernel mode or operating system failures. Some information that uniquely identifies the user might unintentionally be collected as part of the crash report. This information, if present, is never used to identify a user. The specific data collected is described later in this

subsection. Also, Microsoft may send information about a problem, including links to Web sites.

- **Default and recommended settings**: Error reporting for application and system errors is enabled by default on servers running Windows Server 2003 with SP1. For more information about recommended settings, see "Controlling Error Reporting to Prevent the Flow of Information to and from the Internet," later in this section.

- **Triggers**: The opportunity to send an error report is triggered by application or system errors.

- **User notification**: A dialog box appears notifying users that an error has occurred and asks if they want to send an error report to Microsoft. Users can review the data that will be sent.

- **Logging**: Descriptions of system and application errors are recorded in the event log.

- **Encryption**: All data that could include personally identifiable information is encrypted (HTTPS) during transmission. The "crash signature," which includes such information as the application name and version, module name and version, and offset (location) is not encrypted.

- **Access**: Microsoft employees and contingent staff may access the error reports to maintain the Error Reporting service or improve Microsoft products, and may not use the reports for other purposes.

  If the error report indicates that one or more non-Microsoft products were involved in causing the problem, Microsoft may send the report to the respective companies. Qualified software or hardware developers (employed by Microsoft or one of its partners) will analyze the fault data and try to identify and correct the problem.

- **Privacy**: The privacy statement for Microsoft Error Reporting is located at the following Web site:

  http://go.microsoft.com/fwlink/?LinkId=825

  Details related to privacy of data are presented in "Types of Data Collected," later in this section.

- **Transmission protocol and port**: The transmission protocol is HTTP and the ports are HTTP 80 and HTTPS 443.

- **Ability to disable**: The feature can be disabled through Group Policy or by administrators on individual servers.

**Types of Errors Reported**

In Windows Server 2003 there are two types of errors that are reported, user mode and kernel mode.

**User Mode Reporting**

When a user mode error occurs, such as an application error, the Error Reporting service does the following:

- Displays an alert stating that Windows Server 2003 detected a problem.

  Users can choose to report the problem or not. If they do report it, they will see that the information is being sent to Microsoft.

- Sends a problem report to Microsoft.

  Users may then be queried for additional computer information (to complete the error report) and again may choose to send it or not.

- When more information is available, offers it to the user or administrator.

  Users or administrators might be offered the option of selecting **More Information**, which directs them to updated drivers, patches, or Microsoft Knowledge Base articles.

If the error report indicates that one or more non-Microsoft products were involved in causing the problem, Microsoft may send the report to the respective companies. Qualified software or hardware developers (employed by Microsoft or one of its

partners) will analyze the fault data and try to identify and correct the problem.

**Kernel Mode Reporting**
When a kernel-mode (system) error occurs, Windows Server 2003 displays a Stop message and writes diagnostic information to a memory dump file. When you restart the computer by using normal mode or Safe Mode (with networking) and log on to Windows Server 2003 as an administrator, the Error Reporting service gathers information about the problem and displays a dialog box that gives you the option of sending a report to Microsoft.

**Types of Data Collected**
The Error Reporting service collects information about the computer configuration, what the software was doing when the problem occurred, and other information directly related to the problem. The Error Reporting service does not intentionally collect anyone's name, address, e-mail address, computer name, or any other form of personally identifiable information. It is possible that such information may be captured in memory or in the data collected from open files, but Microsoft does not use it to identify users. The Error Reporting service collects Internet Protocol (IP) addresses, but the addresses are not used to identify users, and in many cases are the address of a Network Address Translation (NAT) computer or proxy server, not a specific client behind that NAT computer or proxy server. IP address information is used in aggregate by the operators who maintain the servers that receive error reports. The other use for IP address information is to locate error reports that come from computers inside Microsoft—errors on those computers can be more thoroughly investigated as needed.

In rare cases, such as problems that are especially difficult to solve, Microsoft may request additional data, including sections of memory (which may include memory shared by any or all applications running at the time the problem occurred), some registry settings, and one or more files from the user's computer. When additional data is requested, the user can review the data and choose to send the information or not.

In Windows Server 2003 with SP1 the specific types of data that are collected when application errors or kernel failures occur are as follows.

**Application Errors**
If an application error occurs for which Error Reporting is available and you choose to send the report, the information included is as follows:

- The Digital Product ID, which can be used to identify your license.
- Information regarding the condition of the computer and the application at the time when the error occurred. This includes data stored in memory and stacks, information about files in the application's directory, as well as the operating system version and the computer hardware in use. This information is packaged into a *minidump*—a small memory dump. The minidump contains the following:
- Exception information: This is information regarding the problem that occurred. It tells Microsoft what kind of instruction the application received that caused it to generate an error.
- System information: This is data about the kind of CPU (processor) you have and what operating system you are running.
- A list of all the modules that are currently loaded and their version information.
- A list of all the threads that are currently running. For each thread, the current context and the whole stack are collected.
- Global data.

The minidump data is shown as a hexadecimal representation that the user cannot read.

| **Note** |
| --- |
| For the exact specification of the minidump format, see the Microsoft Platform SDK, which is available on the MSDN Web site. |

**Windows Kernel Failures**
Windows kernel fault reports contain information about what your operating system was doing when the problem occurred. These event reports contain the minimum information that can help to identify why the operating system stopped unexpectedly. If you choose to send the report, it includes:

- The operating system name (for example, Microsoft Windows Server 2003).

- The operating system version (for example, 5.1.2600 0.0).
- The operating system language as represented by the locale identifier (LCID) (for example, 1033 for United States English). This is a standard international numeric abbreviation.
- The loaded and recently unloaded drivers. These identify the modules used by the kernel when the Stop error occurred, and the modules that were used recently.
- The list of drivers in the Drivers folder on your hard disk (systemroot\System32\Drivers).
- The file size, date created, version, manufacturer, and full product name for each driver.
- The number of available processors.
- The amount of random access memory (RAM).
- The time stamp that indicates when the Stop error occurred.
- The messages and parameters that describe the Stop error.
- The processor context for the process that stopped. This includes the processor, hardware state, performance counters, multiprocessor packet information, deferred procedure call information, and interrupts (requests from software or devices for processor attention).
- The process information and kernel context for the halted process. This includes the offset (location) of the directory table and the database that maintains the information about every physical page (block of memory) in the operating system.
- The process information and kernel context for the thread that stopped. This information identifies registers (data-storage blocks of memory in the processor) and interrupt request levels, and includes pointers to data structures for operating system data.
- The kernel-mode call stack for the interrupted thread. This is a data structure that consists of a series of memory locations and one or more pointers.

**Controlling Error Reporting to Prevent the Flow of Information to and from the Internet**

To prevent the automatic flow of information to and from the Internet when users and administrators report errors, you can configure error reporting in two ways: while deploying Windows Server 2003 using answer files with unattended or remote installation, or after deployment using Group Policy. There may be some aspects of error reporting that you want to configure using answer files, and others you may want to configure using Group Policy. Review the tables in this subsection to determine the configuration options that will work best for your organization.

**Using Unattended Installation**

You can configure error reporting by using standard methods for unattended or remote installation. You use the **[PCHealth]** section of an answer file to make entries for this feature. The following table describes those entries.

**Entries for Configuring Error Reporting in an Answer File (for Unattended Installation)**

| Entry | Description |
|---|---|
| ER_Display_UI | Specifies whether Setup notifies the user that an error has occurred and shows details about the error. When the entry is **ER_Display_UI = 0**, Setup does not notify the user that an error has occurred. |

| | |
|---|---|
| ER_Enable_Applications<br><br>ER_Include_EXE(n)<br><br>and<br><br>ER_Exclude_EXE(n) | **ER_Enable_Applications = All**<br><br>Reports errors for all applications except for those listed in ER_Exclude_EXE(*n*).<br><br>**ER_Enable_Applications = Listed**<br><br>Reports errors only for those applications listed in ER_Include_EXE(n). You can automatically include Microsoft applications by using ER_Include_MSApps.<br><br>**ER_Enable_Applications = None**<br><br>Reports no application errors.<br><br>Examples of entries that list included applications are:<br><br>**ER_Include_EXE1 = iexplore.exe**<br><br>**ER_Include_EXE2 = explorer.exe**<br><br>Examples of entries that list excluded applications are:<br><br>**ER_Exclude_EXE1 = calc.exe**<br><br>**ER_Exclude_EXE2 = notepad.exe** |
| ER_Enable_Kernel Errors | Specifies whether Windows reports errors in the Windows kernel. When the entry is **ER_Enable_Kernel Errors = 0**, Windows does not report errors in the Windows kernel. |
| ER_Enable_Reporting | Specifies whether Windows automatically reports errors. When the entry is **ER_Enable_Reporting = 0**, Windows does not report errors. |
| ER_Enable_Windows_ Components | Specifies whether to report errors in Windows components. When the entry is **ER_Enable_Windows_Components = 0**, Windows does not report errors in Windows components. To exclude individual Windows components, use ER_Exclude_EXE(*n*), as described earlier in this table. |
| ER_Force_Queue_Mode | Specifies whether to send all reports in queue mode. When the entry is **ER_Force_Queue_Mode = 0**, Windows does not send reports in queue mode. |
| ER_Include_MSApps | Specifies whether to track and report errors in Microsoft applications. When the entry is **ER_Include_MSApps = 0**, errors in Microsoft applications are not tracked and are not reported. |
| ER_Include_Shutdown_ Errs | Specifies whether to report shutdown errors. When the entry is **ER_Include_Shutdown_Errs = 0**, shutdown errors are not reported. |

For complete details about the entries for error reporting, see the resources listed in Appendix A: Resources for Learning About Automated Installation and Deployment. Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix).

**Using Group Policy**

You can use Group Policy to disable Windows Error Reporting, or you can use it to redirect error reports to a server on your intranet, after which you can enable Corporate Error Reporting. By controlling error reporting through Group Policy, you override actions users may take.

This subsection includes a list of the settings in both the **Error Reporting** policy and the **Advanced Error Reporting** policy.

**Error Reporting Policy Settings**

The error reporting policy settings that you configure for Corporate Error Reporting are located at Computer Configuration\Administrative Templates\System\Error Reporting. The following list describes the policy settings.

- **Configure Error Reporting, enabled**: Errors are reported to Microsoft through the Internet or to a server on your intranet. Enabling **Configure Error Reporting** will override any error reporting settings made using Control Panel. Default values will be used for any error reporting settings that are not configured, even if settings were adjusted through Control Panel.

  In **Configure Error Reporting**, you can select the following:

  - Do not display links to any Microsoft provided "more information" web sites

    - Do not collect additional files
    - Do not collect additional machine data
    - Force queue mode for application errors

  In **Configure Error Reporting**, you can enter:

    - Corporate upload file path
    - Text with which to replace instances of the word "Microsoft"

- **Configure Error Reporting, disabled**: Users will not be given the option to report errors. If **Display Error Notification** is enabled, users will still get a message indicating that a problem occurred, but they will not have the option to report it.
- **Configure Error Reporting, not configured**: A person logged in as an administrator will be able to adjust the setting using Control Panel, which is set to "Enable error reporting" by default on Windows Server 2003 SP1.
- **Display Error Notification, enabled**: This setting controls whether a user is given the choice to report an error. When enabled, the user will be notified that an error has occurred and will be given access to details about the error.
- **Display Error Notification, disabled**: The user is not given the choice of whether to report the error. If Configure Error Reporting is enabled, the error will be automatically reported, but the user will not be notified that an error has occurred.
- **Display Error Notification, not configured**: A person logged in as an administrator will be able to adjust the setting through Control Panel, which is set to enable notification by default.

**Advanced Error Reporting Policy Settings**

When you enable error reporting, you can choose to specify the types of errors that are reported. In a highly managed environment, administrators might want to do this based on the kinds of information included in the error report (see "Types of data collected," in the previous subsection).

With **Advanced Error Reporting settings** you can configure the following policy settings:

- **Default application reporting settings**
- **List of applications to always report errors for**
- **List of applications to never report errors for**
- **Report operating system errors**
- **Report unplanned shutdown events**

These policy settings are located in Computer Configuration\Administrative Templates\System\Error Reporting. When you configure these policy settings, they will override any adjustments to error reporting that administrators might make through Control Panel.

To find more information about editing Group Policy, see Appendix B: Resources for Learning About Group Policy.

**How Controlling Error Reporting Can Affect Administrators**

What administrators will see on a server when an error occurs depends on how you have configured the **Error Reporting** policy settings. Depending on which policy settings you have enabled and which options you have configured, you can have administrators input varying amounts of information during error reporting, or none at all. You can choose not to have any user interface when a fault occurs, or, you can have a notification that an error has occurred, but not allow for the opportunity to send a report.

Another factor in how the user interface is affected is how you have configured the following policy settings: **List of applications to always report errors for** and **List of applications to never report errors for**. For more information about these policy settings, see "Procedures for Configuring Error Reporting" and "Related Links," later in this section.

The following table presents two examples of what the user will see when an error occurs if you have enabled the **Error Reporting** policy settings and if you have entered a path to a server. The first option presents the recommended policy settings.

**How Sending Error Reports to an Intranet Server Affects the User Interface**

| Configuration Options | User Interface |
|---|---|
| **Configure Error Reporting** enabled; **Corporate file path** entered; **Display Error Notification** enabled | • User is notified that an error occurred<br>• User might be asked for additional data<br>• Reports go to an intranet server |
| **Configure Error Reporting** enabled; **Corporate file path** entered; **Display Error Notification** not enabled | • No user interface<br>• Reports automatically go to an intranet server |

**Procedures for Configuring Error Reporting**

The following procedures explain how to:

- Use the **Configure Error Reporting** policy setting so error reports are sent to a server on your intranet instead of to Microsoft, so you can then use the Corporate Error Reporting tool to filter reports.
- Locate the Group Policy settings for configuring error reporting.
- Locate the Group Policy setting for disabling error reporting.
- Prevent error reporting by using an answer file for unattended installation.

**To Enable Corporate Error Reporting**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.
2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, and then click **Error Reporting**.
3. In the details pane, double-click **Display Error Notification**, and then click **Enabled**.
4. Click **Next Setting**, and then under **Configure Error Reporting**, click **Enabled**.
5. In the **Corporate upload file path** box, enter a UNC (Universal Naming Convention) path (*servername**sharename*).

| Note |
|---|
| Administrators can then filter the error reports using the CER tool described in the previous subsection, "Controlling Error Reporting to Prevent the Flow of Information to and from the Internet." |

Use the following two procedures to locate the Group Policy settings described in "Using Group Policy," earlier in this section.

**To Locate Group Policy Settings for Configuring Error Reporting**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, and then click **Error Reporting**.

3. View the Group Policy settings that are available. For more information about these settings, see "Error Reporting Policy Settings," earlier in this section.

4. Click **Advanced Error Reporting settings**.

5. View the advanced settings that are available. For more information about these settings, see the list in "Advanced Error Reporting Policy Settings," earlier in this section.

**To Disable Windows Error Reporting by Using Group Policy**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.

3. In the details pane, double-click **Turn off Windows Error Reporting**, and then click **Enabled**.

| **Important** |
| --- |
| You can also restrict Internet access for this and a number of other components by applying the **Restrict Internet communication** policy setting, which is located in Computer Configuration/Administrative Templates/System/Internet Communication Management. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

**To Prevent Error Reporting by Using an Answer File for Unattended Installation**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file. For information about unattended installation, and for complete details about the entries for error reporting, see the resources listed in Appendix A: Resources for Learning About Automated Installation and Deployment. Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix).

2. In the [PCHealth] section of the answer file, create entries according to the table in "Using Unattended Installation," earlier in this section. For example, to disable error reporting, the entry is:

```
[PCHealth]
ER_Enable_Reporting = 0
```

**Related Links**

- For more information about Windows Error Reporting, see the article on the MSDN Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29903

- To obtain the Corporate Error Reporting tool, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29517

- To read the Microsoft privacy statement for error reporting, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=825

# Windows Media Player and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Windows Media Player
- How Windows Media Player communicates with sites on the Internet
- Procedures for configuration of Windows Media Player

**Resources Describing Windows Media Player Configuration Options**
This section of the white paper describes Microsoft Windows Media Player (also called the Player) in the context of use on a server. The Player is not commonly used in a server environment, unless it is being used to test a Windows Media server, so this section of the white paper does not provide extensive information about implementing a specific configuration for the Player. For example, this section does not provide information about allowing or preventing downloads of codecs or other software for the Player. (A codec, short for compressor/decompressor, is software that compresses or decompresses audio or video data.) This section of the white paper also does not describe Windows Media Services, which is described in a separate section of the white paper.

For information about implementing a specific desktop configuration for Windows Media Player 9 Series in your organization, see the white paper titled "Using Windows XP Professional with Service Pack 2 in a Managed Environment: Controlling Communication with the Internet." The white paper includes information about allowing or preventing downloads for the Windows Media Player 9 Series (for example, downloads of codecs or other software for the Player). This white paper can be found on the TechNet Web site at:

http://go.microsoft.com/fwlink/?LinkId=29133

Note that there are differences between Windows Media Player 9 Series in Windows XP SP2 and Windows Media Player 10 in Windows Server 2003 with SP1. However, the applicable Group Policy settings affect these two versions of the Player in the same way. Specifically, Group Policy settings in the following locations have the same effect on both Windows Media Player 9 Series in Windows XP SP2 and Windows Media Player 10 in Windows Server 2003 with SP1:

User Configuration\Administrative Templates\Windows Components\Windows Media Player

Computer Configuration\Administrative Templates\Windows Components\Windows Media Digital Rights Management

For a list of additional resources, see "Related Links," later in this section.

**Benefits and Purposes of Windows Media Player**
Windows Media Player 10 is the media player included with Microsoft Windows Server 2003 Service Pack 1 (SP1). The Player enables playing and organizing digital media files on computers and on the Internet. If you choose to use the Player on a server, you can listen to radio stations, search for and organize digital media files, and (with the right hardware) play CDs and DVDs, create custom CDs, and copy files to a portable device.

Windows Media Player 10 is described in this white paper because it can access media files on the Internet. In a highly managed network environment, you might want to control access to the Internet, including access gained through Windows Media Player. The rest of this section briefly describes how the Player communicates with the Internet, and how to control this communication.

**Overview: Using Windows Media Player in a Managed Environment**
Windows Media Player 10 in Windows Server 2003 SP1 is an integral component of the operating system and is installed by default. The Windows Media Player interface is accessed by navigating to Programs\Accessories\Entertainment or All Programs\Accessories\Entertainment from the Start menu. The Player is not commonly used in a server environment, unless it is being used to test a Windows Media server. If you do not want the Player accessible, you can use Group Policy to block access to the Player. The procedure for this configuration method is described later in this section.

**How Windows Media Player Communicates with Sites on the Internet**
The Windows Media Player has a number of features that connect to sites on the Internet. When a person at a server selects a feature such as the media **Guide** or **Skin Chooser\More Skins**, Windows Media Player connects to www.WindowsMedia.com through either a local area network (LAN) or a modem connection. When a person at a server selects a feature such as **Music**,

**Radio**, or **Movies & TV**, Windows Media Player connects to a related site, by default, a site at MSN.com such as radio.msn.com.

WindowsMedia.com is a Web site operated by Microsoft and is tightly integrated into Windows Media Player. Media Guide is a Web page provided by WindowsMedia.com. Other services provided by WindowsMedia.com include the Player updates and download support for codecs, skins, and visualizations. (A codec, short for compressor/decompressor, is software that compresses or decompresses audio or video data.)

To support the playback of secure content, Windows Media Player will also contact:

- Non-Microsoft digital rights management (DRM) license servers
- Microsoft DRM upgrade service

The other common Internet connections that Windows Media Player makes are to Windows Media servers run by content providers.

For details about how Windows Media Player 10 communicates with the Internet, including details about what information is sent, see "Related Links," later in this section, and see the privacy statement for Windows Media Player 10 on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=46465

**Procedure for Controlling Windows Media Player**

Windows Media Player can be configured as described previously. This subsection provides a procedure for using Group Policy to prevent users or administrators from starting Windows Media Player.

| **Note** |
| --- |
| For information about implementing a specific configuration for Windows Media Player in your organization, see the sources listed in "Related Links," later in this section. |

**To use Group Policy to Prevent Users and Administrators from Starting Windows Media Player**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.
2. Click **User Configuration**, click **Administrative Templates**, and then click **System**.
3. In the details pane, double-click **Don't run specified Windows applications**.
4. Select **Enabled**, click **Show**, click **Add**, and then enter the application executable name, Wmplayer.exe.

**Related Links**

- For more information about deploying and managing Windows Media Player in an enterprise environment, see the following pages on the Windows Media Web site at:

  http://www.microsoft.com/Windows/WindowsMedia/

  http://go.microsoft.com/fwlink/?LinkId=46006

- For the privacy statement for Windows Media Player 10, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46465

- For information about implementing a specific desktop configuration for Windows Media Player 9 Series in your organization, see the white paper titled "Using Windows XP Professional with Service Pack 2 in a Managed Environment: Controlling Communication with the Internet." This white paper can be found on the TechNet Web site at:

http://go.microsoft.com/fwlink/?LinkId=29133

- For information about Windows Media Services, see the corresponding section of this white paper.

# Windows Media Services and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Windows Media® Services on servers running Microsoft Windows Server 2003 with Service Pack 1 (SP1).

  | Note |
  | --- |
  | Windows Media Services 9 Series is included in most editions of Windows Server 2003. However, Windows Media Services is not included in the Web Edition or in the Itanium-based versions of Windows Server 2003. |

  Windows Media Services delivers advanced streaming functionality such as multicasting, wireless network support, Internet authentication, server plug-ins, and Cache/Proxy APIs.

- For servers from which you want to offer content that will be streamed to an intranet or the Internet, the following types of information are provided:

  - Examples of features in Windows Media Services 9 Series that help you control communication to and from a Windows Media server. Windows Media Services 9 Series is the version of Windows Media Services included with Windows Server 2003 SP1.

  - References to more detailed information about Windows Media Services, including information about ports and security-related topics.

  - Information about installing the Windows Media Services and the Windows Media Services subcomponents, along with instructions for viewing the Help that is included with Windows Media Services.

- For servers from which you do not want to offer content on an intranet or the Internet, information about excluding or removing Windows Media Services.

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running servers that communicate across the Internet. This section, however, provides overview information as well as suggestions for other sources of information about balancing your organization's requirements for communication across the Internet with your organization's requirements for protection of networked assets.

| Note |
| --- |
| This section of the white paper describes Windows Media Services (the server component), but it does not describe Windows Media Player (the client component) or Internet Information Services (IIS), both of which are involved in carrying out communication of multimedia content across the Internet. For information about these components, see the respective sections of this white paper. |

**Benefits and Purposes of Windows Media Services**

Windows Media Services is an optional component in Windows Server 2003. With Windows Media Services, you can manage and deliver Windows Media content over an intranet or the Internet. The clients receiving the content can render it as it is being received, that is, without downloading the content first. Streaming greatly reduces the wait time and storage requirements on the client. It also permits presentations of unlimited length, as well as live broadcasts.

For more information about features in Windows Media Services, see the sources in "Related Documentation and Links," later in this section.

**Examples of Features that Help You Control Communication to and from a Windows Media Server**

This subsection provides brief descriptions of some features in Windows Media Services 9 Series that help you control communication to and from a Windows Media server. These features are integrated with two aspects of basic functionality built into the Windows Server 2003 operating system:

- Authentication
- Authorization

**Authentication**
Authentication is a fundamental aspect of security for a server running Windows Media Services. It confirms the identity of any unicast client trying to access resources on your Windows Media server. Windows Media Services includes authentication plug-ins that you can enable in order to validate user credentials for unicast clients. Authentication plug-ins work together with authorization plug-ins: after users are authenticated, authorization plug-ins control access to unicast content.

Windows Media Services authentication plug-ins fall into the following categories:

- **Anonymous authentication**. These are plug-ins that do not exchange challenge and response information between the server and a player, such as the WMS Anonymous User Authentication plug-in.
- **Network authentication**. These are plug-ins that validate unicast clients based on user logon credentials, such as the WMS Negotiate Authentication plug-in.

When you make decisions about how authentication might affect users, consider the following points:

- For multicast streaming with Windows Media Services 9 Series, clients do not establish a connection, and therefore authentication and authorization do not apply for multicasting. (Multicast streaming is only available if you have Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, or the x64-based versions of these operating systems.)
- If a player is connected through HTTP, the player disconnects from the server each time the user stops, pauses, fast-forwards, or rewinds the content. If the user tries to continue receiving the content, the authentication and authorization process occurs again.

For more information about authentication and about the specific authentication plug-ins that you can enable for Windows Media Services, see the list in "Related Documentation and Links," later in this section.

**Authorization**
In order to control access to unicast content on your Windows Media server, unless you identify users only by IP address, you must enable one or more authentication plug-ins and also one or more authorization plug-ins. Authentication plug-ins verify the credentials of unicast clients attempting to connect to the server. Authorization plug-ins verify that the unicast client is allowed to connect to the server. Authorization occurs after authentication is successful.

You can enable authorization plug-ins to control the access to content by authenticated users. If you enable an authorization plug-in, with one exception, you must also enable an authentication plug-in for unicast clients to be able to access your publishing points. The exception is the WMS IP Address Authorization plug-in, which does not require an authentication plug-in to authenticate a unicast client.

Note that for multicast streaming with Windows Media Services 9 Series, clients do not establish a connection, and therefore authentication and authorization do not apply for multicasting. (Multicast streaming is only available if you have Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, or the x64-based versions of these operating systems.)

During the authorization process, the server checks the user against the set of access permissions for the resource to which the user is trying to connect.

For more information about authorization, see the list in "Related Documentation and Links," later in this section.

**Firewall Information for Windows Media Services**
This subsection provides information about configuring firewalls (or proxy servers or both) for use with Windows Media Services. For more information about firewalls, see the sources in "Related Documentation and Links," later in this section, as well as the Windows Media Web site at:

http://go.microsoft.com/fwlink/?linkid=29862

You can configure each control protocol plug-in (Microsoft Media Server [MMS] protocol, Real Time Streaming Protocol [RTSP], and HTTP) to use a specific port to make firewall configuration easier. If opening ports on your firewall is not possible,

Windows Media Services can stream content by using the HTTP protocol over port 80.

| Note |
| --- |
| Using HTTP to stream content is disabled by default. |

Windows Media Services was formerly known as Microsoft NetShow Services; some firewalls have a preconfigured NetShow setting, which may work for Windows Media Services.

**Configuring Firewalls for Unicast Streaming**

To configure a firewall for unicast streaming, you must open the ports on the firewall that are required for the connection protocols enabled on your server. If you are streaming content by using either the Microsoft Media Server (MMS) protocol or the Real Time Streaming Protocol (RTSP), you need to support both the User Datagram Protocol (UDP) and Transmission Control Protocol (TCP).

To enable Windows Media Player and other clients to use the HTTP, RTSP, or MMS protocols to connect to a Windows Media server that is behind a firewall, open the ports described in the following table.

**Ports to Open When Clients are Connecting Using HTTP, RTSP, or MMS Protocols**

| Ports | Description |
| --- | --- |
| In: TCP on ports 80, 554, and 1755 | The Windows Media server uses the TCP In ports to accept an incoming HTTP connection (port 80), RTSP connection (port 554), or MMS connection (port 1755) from Windows Media Player and other clients. |
| In: UDP on ports 1755 and 5005 | The Windows Media server uses UDP In port 1755 to receive resend requests from clients streaming by using MMSU (MMS used with UDP), and UDP In port 5005 to receive resend requests from clients streaming by using RTSPU (RTSP used with UDP). |
| Out: UDP on ports 1024 through 5000. | The Windows Media server uses UDP Out ports 1024 through 5000 to send data to Windows Media Player and other clients. |

To enable a distribution server that is behind a firewall to use the HTTP or RTSP protocols to stream content that originates from a server outside the firewall, open the ports described in the following table.

**Ports to Open When a Distribution Server is Behind a Firewall and Uses HTTP or RTSP to Stream Content that Originates from a Server Outside the Firewall**

| Ports | Description |
| --- | --- |
| In: UDP on ports 1024 through 5000 | The Windows Media server uses UDP In ports 1024 through 5000 to receive data from another server. |
| Out: TCP on ports 80 and 554 | The Windows Media server uses the TCP Out ports to establish an HTTP connection (port 80) or RTSP connection (port 554) to another server or encoder. |
| Out: UDP on port 5005 | When RTSPU distribution is used, the Windows Media server uses UDP Out port 5005 to send resend requests to another server. |

| Note |
| --- |
| If it is not possible to open all the UDP Out ports on a firewall, UDP packets sent by a Windows Media server may be blocked by the firewall and may not be able to reach the clients on the other side of the firewall. If this is the case, clients may still be able to receive a stream by automatically rolling over to a TCP-based protocol, such as HTTP or RTSPT (RTSP used with TCP). However, the rollover will cause a delay for the client receiving the stream. |

| Note |
| --- |
| If you know you will not be able to support UDP streaming through a firewall, you can decrease the rollover delay by clearing the UDP check box in the Unicast Data Writer plug-in Properties dialog box. For more information, see the Help for Windows Media Services. A procedure for viewing Help is included in "Procedures for Installing, Removing, or Excluding Windows Media Services and Its Subcomponents," later in this section. |

For more information, see the sources in "Related Documentation and Links," later in this section.

**Configuring Firewalls for Multicast Streaming**

If you distribute content using multicast streaming, network traffic is directed through the standard Class D IP addresses (224.0.0.1 through 239.255.255.255). For multicast streaming, you must enable multicast-forwarding on your network. The Internet Group Management Protocol (IGMP), supported by Windows Media Services, ensures that multicast traffic passes through your network only when a player requests a multicast connection, so that enabling multicasting on your routers does not flood your network.

The following firewall configuration enables multicast packets to traverse your firewall:

**IP multicast address range**: 224.0.0.1 through 239.255.255.255

To enable IP multicasting, you must allow packets sent to the standard IP multicast address range to come through your firewall. This IP multicast address range must be enabled on both the player and server sides, as well as on every router in between. For more information about multicasting, see the section called "Delivering a multicast stream" in the white paper on the Windows Media Web site at:

http://go.microsoft.com/fwlink/?linkid=29862

**Enabling Access to an Encoder Outside a Firewall**

Encoders use HTTP to connect to a server running Windows Media Services. By default, Windows Media Encoder uses port 80 for HTTP connections; however, the encoder administrator can specify a different port. If a different port is used, you must specify the same port when you identify the encoder connection URL for the Windows Media server and when opening the port on your firewall.

The following example of a firewall configuration allows a computer running Windows Media Encoder outside a firewall to access a Windows Media server behind a firewall by using HTTP:

**In/Out**: Transmission Control Protocol (TCP) on port 80.

(The In port is the port through which the server accepts connections. The Out port is the port through which the server sends data to clients.)

**Answer File Entries and Registry Keys for Windows Media Services Subcomponents**

For reference purposes, the following table shows the syntax for answer file entries associated with Windows Media Services in Windows Server 2003. The table also shows the corresponding registry keys. Do not change the registry keys. They are shown for use in a script that could check whether a particular subcomponent is installed on a particular server. A registry key value of 0x00000000 means the subcomponent is not installed, and a value of 0x00000001 means the subcomponent is installed.

| **Note** |
|---|
| For more details about answer-file entries related to Windows Media Services subcomponents, including information about dependencies between the entries, see the references listed in Appendix A: Resources for Learning About Automated Installation and Deployment. Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix). |

**Answer File Entries and Registry Keys for Windows Media Services Subcomponents for Windows Server 2003**

| **Windows Media Services subcomponent** | **Answer file entry (in the [Components] section)** | **Registry key (for use in a script that checks whether a subcomponent is installed): 0x00000000 means it is not installed; 0x00000001 mean it is installed** |
|---|---|---|
| Core Windows Media server components | wms = On \| Off | No key available (check for subcomponents by using other keys) |
| Windows Media Services Administrator for the Web | wms_admin_asp = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\<br>Windows\CurrentVersion\Setup\OC Manager\<br>Subcomponents\wms_admin_asp |
| Windows Media Services MMC snap-in | wms_admin_mmc = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\<br>Windows\CurrentVersion\Setup\OC Manager\<br>Subcomponents\wms_admin_mmc |

| | | |
|---|---|---|
| Multicast and Advertisement Logging Agent components | wms_isapi = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\<br>Windows\CurrentVersion\Setup\OC Manager\<br>Subcomponents\wms_isapi |
| Windows Media Services server components | wms_server = On \| Off | HKEY_LOCAL_MACHINE\Software\Microsoft\<br>Windows\CurrentVersion\Setup\OC Manager\<br>Subcomponents\wms_server |

**Procedures for Installing, Removing, or Excluding Windows Media Services and Its Subcomponents**
The following procedures explain how to:

- Add or remove Windows Media Services on a computer after setup is complete for Windows Server 2003.
- View the Help that is included with Windows Media Services.
- Prevent the installation of Windows Media Services during unattended installation by using an answer file.
- Specify answer file entries that control whether Windows Media Services subcomponents are included during unattended installation.

| **Note** |
|---|
| Windows Media Services 9 Series is included in most editions of Windows Server 2003. However, Windows Media Services is not included in the Web Edition or in the Itanium-based versions of Windows Server 2003. |

**To Add or Remove Windows Media Services on an Individual Computer after Setup is Complete**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components** (on the left).
4. Select **Windows Media Services**.
5. Perform one of the following steps:
   - If Windows Media Services is installed and you want to remove it, clear the check box for **Windows Media Services** and complete the wizard.
   - If Windows Media Services is not installed and you want to add it, select the check box for **Windows Media Services** and complete the wizard.
   - If you want to view the list of Windows Media Services subcomponents, after selecting **Windows Media Services**, click **Details**.

**To View the Help That Is Included with Windows Media Services**

1. Make sure that Windows Media Services is installed by using the previous procedure.
2. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
3. Double-click **Administrative Tools** and then click **Windows Media Services**.

4. Click the **Help** menu and then click **Help Topics**.

**To Specify Answer File Entries that Control Whether Windows Media Services Subcomponents are Included During Unattended Installation**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file.

2. In the [Components] section of the answer file, add the appropriate entries listed in the table in "Answer File Entries and Registry Keys for Windows Media Services Subcomponents," earlier in this section. Ensure that the entries specify **Off** for components you do not want to install and **On** for components you want to install.

   If no Windows Media Services subcomponents are listed in an answer file for unattended installation of Windows Server 2003, these components are not installed by default.

| **Note** |
| --- |
| For more information about unattended installation, and for details about dependencies between answer-file entries related to Windows Media Services subcomponents, see the references listed in Appendix A: Resources for Learning About Automated Installation and Deployment. Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix). |

**To Prevent the Installation of Windows Media Services During Unattended Installation by Using an Answer File**

1. Using the methods you prefer for unattended installation or remote installation, create an answer file. For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment.

2. In the [Components] section of the answer file, ensure that there are no entries for the subcomponents listed in the table in "Answer File Entries and Registry Keys for Windows Media Services Subcomponents," earlier in this section. If you want to list any of these subcomponents, ensure that the entries specify **Off**.

If no Windows Media Services subcomponents are listed in an answer file for unattended installation of Windows Server 2003, these subcomponents are not installed by default.

**Related Documentation and Links**
The following list of resources can help you as you plan or modify your implementation of Windows Media Services and Windows Media Player in your organization:

- For technical information about Windows Media, see “Windows Media in the Enterprise,” on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46006

  A variety of technical resources are available on the preceding Web site, including:

  - The Windows Media 9 Series Deployment Guide.

  - The Enterprise Deployment Pack (EDP) for Windows Media Player 9 Series, a downloadable packaging tool that simplifies the configuration, deployment, and management of Windows Media Player 9 Series.

  - Other technical articles.

- For information about deploying over an intranet, see “Deploying Windows Media 9 Series over an Intranet,” on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46008

- For information about using Windows Media Services 9 Series with a firewall, see "Firewall Information for Windows

Media Services 9 Series" on the Windows Media Web site at:

http://go.microsoft.com/fwlink/?linkid=29862

- For conceptual and how-to information about using Windows Media Services, including information about authentication, authorization, ports, and firewall settings, see the Help that is included with Windows Media Services. For information about installing Windows Media Services and viewing Help, see "Procedures for Installing, Removing, or Excluding Windows Media Services and Its Subcomponents," earlier in this section.

- For general information about features, and information about ports and firewall or proxy settings, search for the latest information on the Windows Media Web site at:

  http://www.microsoft.com/Windows/WindowsMedia/

- For information about Windows Media Services Software Development Kits (SDKs), see the Microsoft Developer Network Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46019

**Printed Reference**


- Birney, B., Tricia Gill, and members of the Microsoft Windows Media Team. Microsoft Windows Media Resource Kit. Redmond, WA: Microsoft Press, 2003.

  You can read a sample chapter and view information about the Microsoft Windows Media Resource Kit on the MS Press Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46020

# Windows Time Service and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about the following:

- The benefits of Windows Time Service
- How Windows Time Service communicates with sites on the Internet
- How to control Windows Time Service to limit the flow of information to and from the Internet
- How to monitor and troubleshoot Windows Time Service after configuration is complete

**Benefits and Purposes of Windows Time Service**

Many components of Microsoft Windows Server 2003 rely on accurate and synchronized time to function correctly. For example, without clocks that are synchronized to the correct time on all computers, Windows Server 2003 authentication might falsely interpret logon requests as intrusion attempts and consequently deny access to users.

With time synchronization, you can correlate events on different computers in an enterprise. With synchronized clocks on all of your computers, you ensure that you can correctly analyze events that happen in sequence on multiple computers. Windows Time Service automatically synchronizes a local computer’s time with other computers on a network to improve security and performance in your organization.

**Overview: Using Windows Time Service in a Managed Environment**

Computers keep the time on their internal clocks, which allows them to perform any function that requires the date or time. For scheduling purposes, however, the clocks must be set to the correct date and time, and they must be synchronized with the other clocks in the network. Without some other method in place, these clocks must be set manually.

With time synchronization, computers set their clocks automatically to match another computer's clock. One computer maintains very accurate time, and then all other computers set their clocks to match that computer. In this way, you can set accurate time on all computers.

Windows Time Service is installed by default on all computers running Windows Server 2003 and Windows XP. Windows Time Service uses Coordinated Universal Time (UTC), which is independent of time zone. Time zone information is stored in the computer's registry and is added to the system time just before it is displayed to the user.

By default, Windows Time Service starts automatically on computers running Windows XP. In a domain, time synchronization takes place when Windows Time Service turns on during system startup and periodically while the system is running. In the default configuration, the Net Logon service looks for a domain controller that can authenticate and synchronize time with the client. When a domain controller is found, the client sends a request for time and waits for a reply from the domain controller. This communication is an exchange of Network Time Protocol (NTP) packets intended to calculate the time offset and round-trip delay between the two computers.

Note that computers running Windows Server 2003 use the Network Time Protocol (NTP), while computers running Windows 2000 use the Simple Network Time Protocol (SNTP).

**How Windows Time Service Communicates with Sites on the Internet**

In Windows Server 2003, Windows Time Service automatically synchronizes the local computer's time with other computers on the network. The time source for this synchronization varies, depending on whether the computer is joined to a domain in the Active Directory directory service or to a workgroup.

**When a Server Running Windows Server 2003 is Part of a Workgroup**

In this scenario, the default setting for the time synchronization frequency is set to "once per week," and this default setting uses the time.windows.com site as the trusted time synchronization source. This setting will remain until you manually set it otherwise. One or more computers might be identified as a locally reliable time source by configuring Windows Time Service on those computers to use a known accurate time source, either by using special hardware or a time source available on the Internet. All other workgroup computers can be configured manually to synchronize their time with these local time sources.

**When a Server Running Windows Server 2003 is a Member of a Domain**

In this scenario, Windows Time Service configures itself automatically, using the Windows Time Service that is available on the domain controllers.

Windows Time Service on a domain controller can be configured as either a reliable or an unreliable time source. Windows Time Service running on a client will attempt to synchronize its time source with servers that are indicated as reliable. Windows Time Service can configure a domain controller within its domain as a reliable time source, and it synchronizes itself periodically with this source. These settings can be modified or overwritten, depending on specific needs.

**Communication Between Windows Time Service and the Internet**

The following list describes various aspects of Windows Time Service data that is sent to and from the Internet and how the exchange of information takes place.

- **Specific information sent or received**: The service sends information in the form of a Network Time Protocol (NTP) packet. For more information about Windows Time Service and NTP packets, see the references listed in "Related Links," later in this section.

- **Default and recommended settings**: Computers that are members of an Active Directory domain synchronize time with domain controllers by default. Domain controllers synchronize time with their parent domain controller. By default, the root parent domain controller will not synchronize to a time source. The root parent domain controller can be set to either synchronize to a known and trusted Internet-based time source, or a hardware time device that provides an NTP or SNTP interface. Its time accuracy can also be maintained manually.

  We recommend that you configure the authoritative time server to synchronize to a hardware source, not an Internet time source. For more information, see article 884776, "Configuring the Windows Time service against a large time offset" in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?LinkId=46021

- **Triggers and user notification**: Windows Time Service is started when the computer starts. Additionally, the service will continue to synchronize time with the designated network time source and adjust the computer time of the local computer when necessary. Notification is not sent to the user.

- **Logging**: Information related to the service is stored in the Windows System event log. The time and network address of the time synchronization source is contained in the Windows event log entries. Additionally, warning or error condition information related to the service is stored in the Windows System event log.

- **Encryption**: Encryption is not used in the network time synchronization for domain peers. (Authentication, however, is used.)

- **Information storage**: The service does not store information, as all information that results from the time synchronization process is lost when the time synchronization service request is completed.

- **Port**: NTP and SNTP use User Datagram Protocol (UDP) port 123 on time servers. If this port is not open to the Internet, you cannot synchronize your server to Internet SNTP or NTP servers.

- **Protocol**: The service on Windows Server 2003 implements NTP to communicate with other computers on the network.

- **Ability to disable**: Disabling the service might have indirect effects on applications or other services. Applications and services that depend on time synchronization, such as Kerberos V5 authentication protocol, may fail, or they may yield undesirable results if there is a significant time discrepancy among computers. Because most computers' hardware-based clocks are imprecise, the difference between computer clocks on the network usually increases over time.

**Controlling Windows Time Service to Limit the Flow of Information to and from the Internet**

Group Policy can be used to control Windows Time Service for computers that are running Windows Server 2003 to limit the flow of information to and from the Internet.

The synchronization type and NTP time server information can be managed and controlled through Group Policy. The Windows Time Service Group Policy object (GPO) contains configuration settings that specify the synchronization type. When the synchronization type is set to NT5DS, Windows Time Service synchronizes its time resource with the network domain controller. Alternatively, setting the type attribute to NTP configures Windows Time Service to synchronize with a specified NTP time server. The NTP server is specified by either its Domain Name System (DNS) name or its IP address when you select NTP as the synchronization type.

For more information about configuring Windows Time Service during deployment of products in the Windows Server 2003 family, see *Designing and Deploying Directory and Security Services* and *Designing a Managed Environment* in the *Microsoft Windows Server 2003 Deployment Kit* at:

http://go.microsoft.com/fwlink/?linkid=44319

Clients on a managed network can be configured to synchronize computer clock settings to an NTP server on the network to minimize traffic out to the Internet and to ensure that the clients synchronize to a single reliable time source. If you choose to do so, you can disable time synchronization for both non-domain and domain computers running Windows Server 2003 by using Group Policy. The procedures for configuring Windows Time Service are given at the end of this section of the white paper.

**How Windows Time Service Can Affect Users and Applications**

Windows components and services depend on time synchronization. For example, the Kerberos V5 authentication protocol on a Windows Server 2003 family domain has a default time synchronization threshold of five minutes. Computers that are more than five minutes out of synchronization on the domain will fail to authenticate using the Kerberos protocol. This time value is also configurable, allowing for greater or lesser thresholds. Failure to authenticate using the Kerberos protocol can prevent logons and access to Web sites, file shares, printers, and other resources or services within a domain.

When the local clock offset has been determined, the following adjustments are made to the time:

- If the local clock time of the client differs from the time on the server by more than the threshold amount, Windows Time Service will change the local clock time immediately. The threshold is five minutes if the computer is part of a domain. For more information about Windows Time Service settings in a domain, see "Related Links" later in this section.

  The threshold is one second if the computer is part of a workgroup. However, if a computer is part of a workgroup and the time differs from the time source by more than 15 hours, the time is not synchronized, as described later in this list.

- If the local clock time of the client differs from the server by less than the threshold amount, the service will gradually synchronize the client with the correct time.

- In a workgroup, if the local clock time of the client differs from the time on a time source by more than 15 hours, a workstation running Windows Time Service and using default settings will not synchronize with the time source. Such occurrences are rare, and are often caused by configuration setting errors. For example, if a user sets the date on the computer incorrectly, the time does not synchronize. Under these circumstances, most often the time is off by a day or more. Be sure to check the computer's calendar and ensure that the correct date has been set.

**Configuration Settings for Windows Time Service**

You can set the global configuration settings for Windows Time Service by using Group Policy. The settings that might be relevant to communication between Windows Time Service and the Internet are described in this subsection.

In Computer Configuration\Administrative Templates\System\Windows Time Service\Global Configuration Settings, there is only one setting that might, in certain scenarios, affect the way that Windows Time Service communicates when the computer is in a domain. This setting is **AnnounceFlags**, which controls whether this computer is marked as a reliable time server. A computer is not marked as reliable unless it is also marked as a time server. The settings are as follows:

- 0 Not a time server

- 1 Always a time server

- 2 Automatic time server, meaning the role is decided by Windows Time Service

- 4 Always a reliable time server

- Automatic reliable time server, meaning the role is decided by Windows Time Service

The default is 10, meaning that Windows Time Service decides the role.

In the Group Policy settings located in Computer Configuration\Administrative Templates\System\Windows Time Service\Time Providers, there are a number of settings that might affect the way that Windows Time Service communicates across the Internet. The following table describes some of these policy settings.

| Note |
| --- |
| The table lists the settings that most directly affect the way Windows Time Service communicates with time sources, but the table does not list all settings. For example, it does not list the setting that specifies the location of the Windows Time Service DLL or the setting that controls the logging of events for Windows Time Service. |

**Selected Group Policy Settings for Configuring the Windows Time Service NTP Client for Computers Running Windows Server 2003**

| Policy Setting | Effect of Setting | Default Setting |
| --- | --- | --- |
| NtpServer | Establishes a space-delimited list of peers from which a computer obtains time stamps, consisting of one or more DNS names or IP addresses per line. Computers connected to a domain must synchronize with a more reliable time source, such as the official U.S. time clock. This setting is used only when Type is set to NTP or AllSync.<br><br>0x01 SpecialInterval<br><br>0x02 UseAsFallbackOnly<br><br>0x04 SymmetricActive<br><br>0x08 NTP request in Client mode | time.windows.com, 0x1 |
| Type | Indicates which peers to accept synchronization from:<br><br>**NoSync**. The time service does not synchronize with other sources.<br><br>**NTP**. The time service synchronizes from the servers specified in the NtpServer registry entry.<br><br>**NT5DS**. The time service synchronizes from the domain hierarchy.<br><br>**AllSync**. The time service uses all the available synchronization mechanisms. | Default options<br><br>**NTP**. Use on computers that are not joined to a domain.<br><br>**NT5DS**. Use on computers that are joined to a domain. |
| CrossSiteSyncFlags | Determines whether the service chooses synchronization partners outside the domain of the computer.<br><br>None 0<br><br>PdcOnly 1<br><br>All 2<br><br>This value is ignored if the NT5DS value is not set. | 2 |
| ResolvePeerBackoffMinutes | Specifies the initial interval to wait, in minutes, before attempting to locate a peer to synchronize with. If the Windows Time Service cannot successfully synchronize with a time source, it will keep retrying, using the settings specified in ResolvePeerBackOffMinutes and ResolvePeerBackoffMaxTimes. | 15 |
| ResolvePeerBackoffMaxTimes | Specifies the maximum number of times to double the wait interval when repeated attempts fail to locate a peer to synchronize with. A value of zero means that the wait interval is always the initial interval in ResolvePeerBackoffMinutes. | 7 |
| SpecialPollInterval | Specifies the special poll interval in seconds for peers that have been configured manually. When a special poll is enabled, Windows Time Service will use this poll interval instead of a dynamic one that is determined by synchronization algorithms built into Windows Time Service. | 604800 (workgroup)<br><br>3600 (domain) |

For other sources of information about Group Policy, see Appendix B: Resources for Learning About Group Policy.

For information about configuring the authoritative time server in a domain, see article 884776, "Configuring the Windows Time service against a large time offset" in the Microsoft Knowledge Base at:

http://go.microsoft.com/fwlink/?LinkId=46021

**Procedures for Configuring Windows Time Service**

The following procedures explain how to set some of the Windows Time Service configuration settings available in Group Policy. For details about other Group Policy settings for Windows Time Service, see the table earlier in this section.

**To Set Group Policy for Windows Time Service Global Configuration Settings**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, and then click **Windows Time Service**.

3. In the details pane, double-click **Global Configuration Settings**, and then click **Enabled**. Configure settings as appropriate for your environment.

**To Configure the Group Policy Setting to Prevent Your Computer from Servicing Time Synchronization Requests from Other Computers on the Network**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Windows Time Service**, and then click **Time Providers**.

3. In the details pane, double-click **Enable Windows NTP Server**, and then select **Disabled**.

**Starting and Stopping Windows Time Service**

By default, Windows Time Service starts automatically at system startup. You can, however, start or stop the service manually by accessing services in Administrative Tools or by using the **net** command.

**To Manually Start Windows Time Service Using the Graphical Interface**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then double-click **Services**.

3. Select **Windows Time** from the list of services.

4. On the **Action** menu, click **Start** to begin the service.

**To Manually Stop Windows Time Service Using the Graphical Interface**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Administrative Tools**, and then double-click **Services**.

3. Select **Windows Time** from the list of services.

4. On the **Action** menu, click **Stop** to discontinue the service.

**To Manually Start Windows Time Service Using the Net Command**

1. To open a Command Prompt window, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. At the command prompt, type **net start w32time**, and then press ENTER.

**To Manually Stop Windows Time Service Using the Net Command**

1. To open a Command Prompt window, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. At the command prompt, type **net stop w32time**, and then press ENTER.

**Synchronizing Computers with Time Sources**

Use the following procedures to synchronize the internal time server with an external time source, and to synchronize the client time with a time server.

**To Synchronize an Internal Time Server with an External Time Source**

1. To open a Command Prompt window, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. At the command prompt, type the following, where *PeerList* is a comma-separated list of Domain Name System (DNS) names or Internet protocol (IP) addresses of the desired time sources:

   **w32tm /config /syncfromflags:manual /manualpeerlist:***PeerList*

   and then press ENTER.

3. Type **w32tm /config /update** and then press ENTER.

| **Note** |
| --- |
| The most common use of this procedure is to synchronize the internal network's authoritative time source with precise external time source. This procedure can be run on any computer running Windows 2000, Windows XP, or Windows Server 2003. |

| **Note** |
| --- |
| If the computer cannot reach the servers, the procedure fails and an entry is written to the Windows System event log. |

**To Synchronize the Client Time with a Time Server**

1. To open a Command Prompt window, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. At the command prompt, type **w32tm /resync**, and then press ENTER.

| **Note** |
| --- |
| The W32tm command-line tool is used for diagnosing problems that can occur with Windows Time Service. |

**Monitoring and Troubleshooting Windows Time Service**

In many cases, problems with Windows Time Service can be attributed to network configuration. If the network is not configured correctly, computers might not be able to communicate to send time samples back and forth. Viewing the contents of NTP packets can help you to identify exactly where a packet is blocked on a network. An error associated with Windows Time Service might occur when a computer is unable to synchronize with an authoritative source. You can use the W32tm command-line tool to assist you in troubleshooting this and other types of errors associated with Windows Time Service.

The W32tm command-line tool is the preferred command-line tool for configuring, monitoring, and troubleshooting Windows Time Service. For more information, search for "W32tm" in Help and Support Center.

**Procedure to Follow When a Computer Is Unable to Synchronize**

By default, a computer running Windows Time Service will not synchronize with a time source if the computer's time is more than 15 hours off. For information about scenarios in which this can occur, see "How Windows Time Service Can Affect Users and Applications," earlier in this section.

**To Resynchronize the Client Time with a Time Server**

1. To open a Command Prompt window, click **Start**, click **Run**, type **cmd**, and then click **OK**.

2. At the command prompt, type **w32tm /resync /rediscover** and then press ENTER.

| **Note** |
| --- |
| When you run the preceding command, it redetects the network configuration and rediscovers network resources, causing resynchronization. You can then view the event log for more information about why the time service does not synchronize. |

**Related Links**

For more information about configuring Windows Time Service during deployment of products in the Windows Server 2003 family, see *Designing and Deploying Directory and Security Services* and *Designing a Managed Environment* in the *Microsoft Windows Server 2003 Deployment Kit* at:

http://go.microsoft.com/fwlink/?LinkId=44319

For information about configuring the authoritative time server in a domain, see article 884776, "Configuring the Windows Time service against a large time offset" in the Microsoft Knowledge Base at:

http://go.microsoft.com/fwlink/?LinkId=46021

# Windows Update, Automatic Updates, and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about the following:

- The benefits of Windows Update and Automatic Updates
- How Windows Update and Automatic Updates communicate with sites on the Internet
- How to control Windows Update and Automatic Updates to limit the flow of information to and from the Internet

| **Important** |
| --- |
| This section describes methods for controlling the way the Automatic Updates component interacts with the Windows Update Web site. To control the way Automatic Updates interacts with Windows Update, also control the type of accounts that people log on with. If an account does not allow software to be installed (for example, if the account is a user account), only one option for Automatic Updates will function while that person is logged on. That option is the automatic download and installation of updates, which means that updates are installed on the computer at a regularly scheduled time, regardless of what type of account the person who uses the computer has, or whether that person is logged on at the time. |

**Benefits and Purposes of Windows Update and Automatic Updates**

**Windows Update**

Windows Update is an online catalog that can be used to support computers running Microsoft Windows operating systems, including Windows Server 2003 with Service Pack 1 (SP1). The catalog contains items such as drivers, critical updates, Help files, and Internet products. Windows Update scans the computer and provides a tailored selection of updates that apply only to the software and hardware on that specific computer. Windows Update then enables you choose updates for your computer's operating system and hardware. New content is added to the Windows Update Web site regularly, so you can always get the most recent and secure updates and solutions.

Windows Update contains two key components:

- **Web site control**: The Windows Update Web site includes an ActiveX Web control program that downloads and installs updates. The Windows Update team receives feedback from their customers on how to improve their Web site, and responds by periodically updating the Web control. The newest version of the Web control program is downloaded automatically when you visit the Windows Update site or when any of the other Windows features calls on the Windows Update control. Just like downloading an ActiveX control, you may receive a security dialog box that a Web control is attempting to be installed. You may not receive the dialog box if you have selected to always trust Microsoft as a content provider (using security settings in Microsoft Internet Explorer). If you do not click Yes on the security dialog box, the control will not be updated and you will not be able to access the Windows Update site.

- **Updates**: As needed, you can access the Windows Update Web site and select component updates to download and install. You control the downloading and installation of the updates. The Windows Update Web site is located at:

  http://windowsupdate.microsoft.com/

**Automatic Updates**

Automatic Updates is not enabled by default. The person who installs the operating system is prompted to enable this option following setup. When Automatic Updates is configured so that updates automatically download and install, the person who uses the computer does not need to visit special Web pages or remember to periodically check for new updates. Automatic Updates can be configured to use one of the following options:

- **Automatic download and installation of updates**: Windows Server 2003 downloads and installs updates automatically on a schedule specified by an administrator of the computer. Updates are installed regardless of what type of account the person who uses the computer has, or whether that person is logged on at the time.

- **Automatic download only**: Windows Server 2003 automatically starts the download whenever it finds updates available for the computer. The updates are downloaded in the background, enabling you to continue working uninterrupted. After the download is complete, if you are logged on as an administrator, an icon in the notification area will prompt you that the updates are ready to be installed.

- **Notification only**: Windows Server 2003 sends a notification after which an administrator of the computer can respond by downloading and installing any updates.

- **Turn off Automatic Updates**: It is left to you to go to the Windows Update Web site and download updates from time to time.

An administrator can decline a specific update that has been downloaded. The administrator can download those declined files again by opening the Performance and Maintenance category in Control Panel, clicking the **System** tool, clicking the **Automatic Updates** tab, and then clicking **Offer updates again that I've previously hidden**. If any of the previously declined updates can still be applied to the computer, those updates will appear the next time that Windows Server 2003 notifies the administrator of available updates.

For more information about using Control Panel to configure Automatic Updates, see "Procedures for Controlling Windows Update and Automatic Updates," later in this section.

**Alternatives to Windows Update and Automatic Updates**
For managed environments, there are several alternatives to Windows Update:

- Windows Update Catalog Web site

- Microsoft Software Update Services (SUS)

- Distribution software, such as Microsoft Systems Management Server, that can be used to distribute software updates

  For more information, see the documentation for your distribution software, and see Appendix A: Resources for Learning About Automated Installation and Deployment, especially the "Related Documentation and Links" subsection in that appendix.

**Windows Update Catalog Web Site**
By using the Windows Update Catalog site, you can use your own software distribution tools to deploy updates to Windows in a managed environment without requiring users to connect to the Windows Update Web site. The Windows Update Catalog site provides a comprehensive catalog of updates that can be distributed over a managed network. It provides a single location for Windows Update content and drivers that display the Designed for Windows logo. Administrators can search the site using keywords or predefined search criteria to select the relevant downloads and then download the updates to a location on their internal network.

An enhancement in products in the Windows Server 2003 family enables you to select updates that you plan to deploy later, which means that you can control how and when the updates are deployed. For additional information, see information about Windows Update on the Microsoft Web site at:

http://windowsupdate.microsoft.com/

**Microsoft Software Update Services (SUS)**
Microsoft Software Update Services (SUS) is a version of Windows Update designed for installation inside the boundary defined by an organization's firewall. This feature is very useful for organizations that

- Do not want their systems or users connecting to an external Web site

- Want to first test these updates before deploying them throughout their organizations

Microsoft Software Update Services enables administrators to quickly and reliably deploy critical updates to servers running Windows Server 2003 and Windows 2000 Server as well as desktop computers running Windows XP Professional and Windows 2000 Professional.

For more information about Software Update Services and updated versions of Software Update Services, see the Microsoft

Web site at:

http://go.microsoft.com/fwlink/?LinkId=29906

**Overview: Using Windows Update and Automatic Updates in a Managed Environment**
As an administrator, you can use Group Policy to block the use of Windows Update or to specify an internal server for Automatic Updates to use when searching for updates. You can also disable Automatic Updates using Control Panel or using the Group Policy Administrative template, Wuau.adm. Details on the methods and procedures for controlling these features are described in the following subsections.

**How Windows Update and Automatic Updates Communicate with Sites on the Internet**
This subsection summarizes the communication process.

- **Specific information sent or received**: Drivers and replacement files (critical updates, Help files, and Internet products) may be downloaded to the computer. The computer is uniquely identified and is logged in the download and installation success report, but the person using the computer is not uniquely identified.

- **Data storage and access**: Windows Update tracks the total number of unique computers that visit the Windows Update Web site. The success or failure of downloading and installing updates is also recorded but no personal information, such as user name or e-mail address, is recorded as part of this. The information that is recorded is stored on servers at Microsoft with limited access that are located in controlled facilities.

  For more information, see "Privacy," later in this list.

| **Note** |
| --- |
| If you want to block the use of the Windows Update Web site, you can apply Group Policy settings to specify an internal server for updates and for storing upload statistics. For more information, see "Procedures for Controlling Windows Update and Automatic Updates." |

- **Default and recommended settings**: By default, Windows Server 2003 allows access to the Windows Update Web site. Recommended settings are described in the next subsection, "Controlling Windows Update and Automatic Updates to Limit the Flow of Information to and from the Internet."

- **Triggers**: You control whether to run Windows Update. If Automatic Updates is enabled following setup, it is triggered about once per day when there is an Internet connection.

- **User notification**:

  - **Windows Update Web site**: You control whether to go to the Windows Update Web site to download files to your computers.

  - **Automatic Updates**: The way that Automatic Updates notifies you depends on how Automatic Updates is configured. For more information, see "Automatic Updates," earlier in this section.

| **Note** |
| --- |
| For information about configuring Automatic Updates, see "To Configure or Disable Automatic Updates Using Control Panel on a Computer Running Windows Server 2003 with SP1," later in this section. |

- **Logging**: Automatic Updates logs events to the event log.

- **Encryption**: Initial data is transferred using HTTPS, and updates are transferred using HTTP. The data packages downloaded to the system by Microsoft are digitally signed.

- **Privacy**: To view the privacy statement for Windows Update, see the Windows Update Web site, and click **Windows Update Privacy Statement**. The Windows Update Web site is located at:

http://windowsupdate.microsoft.com/

Automatic Updates is covered by the same privacy statement that covers Windows Update.

- **Transmission protocols and ports**: The transmission protocols and ports used are HTTP 80 and HTTPS 443.

- **Ability to disable**: You can use Group Policy to prevent the operating system from being updated through Windows Update, to prevent access to Windows Update commands, or both. You can use Group Policy to specify an internal server to use for Automatic Updates. You can disable Automatic Updates using Control Panel tools or Group Policy. Procedures for these methods are given at the end of this section.

**Controlling Windows Update and Automatic Updates to Limit the Flow of Information to and from the Internet**
The recommended methods for controlling Windows Update and Automatic Updates or both are as follows.

| **Important** |
| --- |
| When using these methods, also control the type of accounts that people log on with. If an account does not allow software to be installed (for example, if the account is a user account), only one option for Automatic Updates will function while that person is logged on. That option is the automatic download and installation of updates, which means that updates are installed on the computer at a regularly scheduled time, regardless of what type of account the person who uses the computer has, or whether that person is logged on at the time. |

- You can use Group Policy settings to disable both Windows Update and Automatic Updates.

    - To disable Windows Update and Automatic Updates by preventing the operating system from being updated through Windows Update, configure **Turn off access to all Windows Update features** in Computer Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings.

    - To disable Windows Update and Automatic Updates by preventing access to Windows Update commands, configure **Remove links and access to Windows Update** in User Configuration\Administrative Templates\Start Menu and Taskbar.

- You can use Group Policy to configure Automatic Updates so that instead of searching the Windows Update Web site, Automatic Updates searches your internal server for updates.

  To do this, configure **Specify intranet Microsoft update service location** in Computer Configuration\Administrative Templates\Windows Components\Windows Update. The server you specify in this setting must be one on which you are running Software Update Services or an updated version of Software Update Services.

- You can use Group Policy settings in the Administrative template Wuau.adm to selectively disable Automatic Updates.

  To do this, disable **Configure Automatic Updates** in Computer Configuration\Administrative Templates\Windows Components\Windows Update.

You can also configure Automatic Updates on individual computers by using Control Panel. For a description of the options available through Control Panel, see "Automatic Updates," earlier in this section.

**How Disabling Windows Update and Automatic Updates Can Affect Users or Administrators and Applications**
The following list shows the effects of two Group Policy settings, both of which prevent the use of Windows Update and Automatic Updates.

- Turn off access to all Windows Update features: This Group Policy setting is located in Computer Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings.

When you enable this setting, the operating system cannot be updated through Windows Update, and Automatic Updates is disabled. Users or administrators can still perform actions such as clicking the **Windows Update** option on the **Start** menu, and the Windows Update Web site will appear in the browser. However, it will not be possible to update the operating system through Windows Update, regardless of the type of account being used to log on.

- Remove links and access to Windows Update This Group Policy setting is located in User Configuration\Administrative Templates\Start Menu and Taskbar. When you enable this setting, you will not be able to access the Windows Update Web site from any of the following locations:

  - The **Windows Update** option on the **Start** menu

  - The **Tools** menu in Microsoft Internet Explorer

  - The **Windows Update** button in **Add New Programs** (**Add New Programs** is in Control Panel under **Add or Remove Programs**)

  Enabling this setting also disables Automatic Updates notifications—that is, the user for which this policy setting is enabled will neither be notified about nor receive critical updates from Windows Update.

Removing access to Windows Update also prevents Device Manager from automatically installing driver updates from the Windows Update Web site. For more information about controlling Device Manager, see the section of this white paper titled Device Manager, Hardware Wizards, and Internet Communication.

Blocking Windows Update and Automatic Updates will not block applications from running.

The Windows Update site is located at:

http://windowsupdate.microsoft.com/

**Procedures for Controlling Windows Update and Automatic Updates**
This subsection provides procedures for the following:

- Configuring or disabling Automatic Updates by using Group Policy.

- Preventing the operating system from being updated through Windows Update by using Group Policy. With this policy, commands for accessing Windows Update are visible and the Windows Update site can be viewed through the browser, although Windows Update cannot be used.

- Turning off access to Windows Update commands and to Automatic Updates by using Group Policy.

- Specifying an internal server for Windows Update by using Group Policy.

- Configuring or disabling Automatic Updates using Control Panel on a computer running Windows Server 2003 SP1.

**To Configure or Disable Automatic Updates Using Group Policy**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated (specifically, Wuau.adm), and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Windows Update**.

3. In the details pane, double-click **Configure Automatic Updates**.

4. Select **Not Configured**, **Enabled**, or **Disabled**. If you choose **Enabled**, choose from the available settings, which are equivalent to the Control Panel settings described in "Automatic Updates," earlier in this section.

| Note |
| --- |
| Disabling this setting disables Automatic Updates but does not block access to Windows Update. |

**To Prevent the Operating System from Being Updated Through Windows Update by Using Group Policy**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **System**, click **Internet Communication Management**, and then click **Internet Communication settings**.

3. In the details pane, double-click **Turn off access to all Windows Update features**.

| Important |
| --- |
| This policy also disables Automatic Updates.<br><br>You can also restrict Internet access for this and a number of other components by applying the **Restrict Internet communication** policy setting, which is located in Computer Configuration/Administrative Templates/System/Internet Communication Management. For more information about this Group Policy and the policies that it controls, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key. |

**To Turn Off Access to Windows Update Commands by Using Group Policy**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.

2. Click **User Configuration**, click **Administrative Templates**, and then click **Start Menu and Taskbar**.

3. In the details pane, double-click **Remove links and access to Windows Update**.

| Important |
| --- |
| This policy also disables Automatic Updates. |

**To Specify an Internal Server for Windows Update Using Group Policy**

1. As needed, see Appendix B: Resources for Learning About Group Policy, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Windows Components**, and then click **Windows Update**.

3. In the details pane, double-click **Specify intranet Microsoft update service location** and then click **Enabled**.

4. Specify the name of the internal server to function as the update server, and specify the name of the server to store upload statistics.

| Important |
| --- |
| You must specify an upgrade server and a server to store upload statistics, but they can be the same server. The server you specify as the upgrade server must be one on which you are running Software Update Services or an updated version of Software Update Services. |

**To Configure or Disable Automatic Updates Using Control Panel on a Computer Running Windows Server 2003 with SP1**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Automatic Updates**.

3. Choose from the available options, which are described in "Automatic Updates," earlier in this section.

# Appendix A: Resources for Learning About Automated Installation and Deployment

Applies To: Windows Server 2003 with SP1

This appendix provides the following:

- An overview of automated installation and deployment
- Procedures and resources for obtaining more information about automated installation and deployment

**Overview: Automated Installation and Deployment**

In the enterprise environment, it is often not cost-effective to install Microsoft Windows Server 2003 or Microsoft Windows XP Professional using the standard interactive setup on each computer. To greatly lower the total cost of ownership (TCO) and ensure configuration uniformity, you can perform an automated installation of Windows Server 2003 and Windows XP on multiple computers. By using an automated installation method, you can ensure that certain components and applications are not available on your organization's computers, or that certain components and applications are preconfigured in such a way that helps prevent unwanted communication over the Internet.

**Methods for Automating the Setup Process**

There are several options for automating the setup process. Any or all of the following tools can help ensure that all of your servers and clients are configured to appropriately limit communication over the Internet:

- Unattended setup using Setup (Winnt32.exe)

  Unattended setup enables you to simplify the process of setting up the operating system on multiple computers. To run an unattended setup, you can create and use an answer file, which is a customized script that answers Setup questions automatically. Then you can run Setup (Winnt32.exe) from the command line with the appropriate options for invoking unattended setup.

  Using Winnt32.exe, you can upgrade your previous version of the operating system using all user settings from the previous installation, or you can perform a fresh installation using the answer file that provides Setup with your custom specifications. The latter method is most likely the best option to limit component communication over the Internet, provided you use an appropriate answer file. Details about specific answer file entries are included in the respective component sections of this white paper.

- Remote Installation Services (RIS)

  With RIS, you can install the operating system by itself or a complete computer configuration, including desktop settings and applications. RIS installations can be either CD-based (through the use of Risetup.exe) or image-based (through the use of Riprep.exe). You can also specify which RIS server will provide installations to a given client computer, or you can allow any RIS server to provide the installation.

- Image-based installation using the System Preparation (Sysprep) tool.

  Image-based installation is also a good choice if you need to install an identical configuration on multiple computers. You typically use the Sysprep tool in conjunction with a non-Microsoft disk imaging tool or Microsoft Windows Server 2003 Automated Deployment Services (ADS) to perform imagebased installations.

  On a master computer, you install the operating system and any applications that you want installed on all of the target computers. Then you run Sysprep and a disk imaging utility. Sysprep prepares the hard disk on the master computer so that the disk imaging utility can transfer an image of the hard disk to the other computers. This method decreases deployment time dramatically compared to standard or unattended installations. You can customize the images so that only the files required for a specific configuration appear on the image, such as additional Plug and Play drivers that might be needed on various systems. The image can also be copied to a CD and distributed to remote sites that have slow links.

- System management software, such as Microsoft Systems Management Server (SMS)

  This type of software assists with the many tasks that are involved when you apply automated procedures to multiple servers and client computers throughout your organization. These tasks include the following:

  - Selecting computers that are equipped for the operating system and that you are ready to support
  - Distributing the operating system source files to all sites, including remote sites and sites without technical support staff
  - Monitoring the distribution to all sites
  - Providing the appropriate user rights to do the upgrade
  - Automatically initiating the installation of the software package with the possibility of having the user control the timing
  - Resolving problems related to the distributions or installations
  - Reporting on the rate and success of deployment

Using system management software helps to further ensure that all computers within your organization have received the standardized operating system configuration that helps prevent unwanted communication over the Internet.

**Using Scripts for Configuring Computers**
In addition to the automated installation methods described here, another common method of controlling the configuration of computers in a domain is to use scripts. For more information about scripts, see "Related Documentation and Links," at the end of this section.

**Procedures for Accessing Additional Information About Automated Setup Tools**
Windows Server 2003 includes extensive documentation describing unattended installation, RIS, and image-based installation with Sysprep. You can view this documentation from any computer that has Internet access (regardless of the operating system running on that computer). You can also view Help for Windows Server 2003 from any computer running Windows Server 2003. The following procedure gives the details.

**To Access Help for a Computer Running Windows Server 2003**

1. Open Help and Support by doing one of the following:

   - On any computer running a product in the Windows Server 2003 family, click **Start**, and then click **Help and Support**.
   - View product documentation on the Web at:

     http://go.microsoft.com/fwlink/?LinkId=29881

     On this site, click the link for the appropriate product.

2. Locate the specific topics as follows:

   - For unattended installation: Navigate to Getting Started\Installing and upgrading the operating system\Concepts\Planning for unattended Setup
   - For RIS: Navigate to Software Deployment\Remote Installation Services
   - For Winnt32.exe: Navigate to Administration and scripting tools\Command-line reference\Command-line reference A-Z\Winnt32

| Note |
| --- |
| Detailed information about unattended installation, RIS, and image-based installation with Sysprep is also available in the sources listed in "Related Documentation and Links." |

**Related Documentation and Links**

You can also find additional information about all of the topics described earlier in this appendix in a variety of other locations:

- On the CD for Windows Server 2003, you can find additional information about unattended installations in the *Microsoft Windows Corporate Deployment Tools User's Guide*, found in \Support\Tools\Deploy.cab. To view the guide and the associated reference information, make sure Deploy.chm and Ref.chm are in the same folder, and open Deploy.chm. For the most up-to-date information, be sure to view the version of the guide that is included in Windows Server 2003 with SP1, rather than an older version from the original release of Windows Server 2003.

- For additional information about unattended installation, RIS, and image-based installation with Sysprep, see the following sources.

    - *The Microsoft Windows Server 2003 Deployment Kit*, especially the book titled *Automating and Customizing Installations*. The *Microsoft Windows Server 2003 Deployment Kit* is available on the Web at:

      http://go.microsoft.com/fwlink/?LinkId=29887

    - "Installation and Setup Technologies" in the Core Operating System Collection of the *Windows Server 2003 Technical Reference* on the Microsoft Web site at:

      http://go.microsoft.com/fwlink/?LinkId=46281

- For general information about Group Policy, see Appendix B: Resources for Learning About Group Policy.

- To learn about specific Group Policy settings that can be applied to computers running Windows Server 2003 with SP1, see the Group Policy Settings Reference on the Microsoft Download Center Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29911

- For more information about scripting, see the Script Center on the TechNet Web site at:

  http://go.microsoft.com/fwlink/?LinkId=24771

- The Help documentation for Windows Server 2003, included in the product and on the Web, contains information about Windows Script Host. You can find the documentation on the Web at:

  http://go.microsoft.com/fwlink/?LinkId=29221

# Appendix B: Resources for Learning About Group Policy

Applies To: Windows Server 2003 with SP1

This appendix provides the following:

- A list of resources for learning about Group Policy and about the Group Policy settings that can be used with Windows Server 2003 with SP1.
- A procedure for updating Administrative templates (.adm files) to the set of templates included in Windows Server 2003 with SP1.

**Resources for Learning about Group Policy**

To learn about using Group Policy and the Group Policy Management Console (GPMC), see the following resources.

- Windows Server 2003 Group Policy page on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46022

- The Group Policy Technology Center Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29866

- *Designing a Managed Environment* in the *Microsoft Windows Server 2003 Deployment Kit* at:

  http://go.microsoft.com/fwlink/?linkid=44319

- The Group Policy Collection of the *Windows Server 2003 Technical Reference* at:

  http://go.microsoft.com/fwlink/?LinkId=29907

- Article 816662, "Recommendations for Managing Group Policy Administrative Template (.adm) Files" in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?LinkId=29128

To learn about Group Policy settings that can be used with Windows Server 2003 with SP1, see the following resources.

- The Group Policy Settings Reference on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29911

- "Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1" on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46278

**Installing Group Policy Management Console (GPMC)**

Group Policy Management Console (GPMC) unifies the management of all aspects of Group Policy. Using GPMC, you can manage Group Policy objects (GPOs), Windows Management Instrumentation (WMI) filters, and permissions on your network.

GPMC is available for download from the Microsoft Web site. For instructions for downloading and installing GPMC, see the Microsoft Windows Server 2003 Web site at:

http://go.microsoft.com/fwlink/?linkid=29909

To run GPMC, click **Start**, point to **Administrative Tools**, and then click **Group Policy Management**.

**Updating Administrative Templates**

You can use Group Policy to control registry keys using the settings supplied with Administrative templates. An updated set of Administrative templates is included with each new release of Windows and includes policy settings for earlier versions of Windows and earlier service packs as well as new policy settings. Use the following procedure to update the Administrative templates in your domain with the templates provided with Windows Server 2003 with SP1.

**To Update Administrative Templates to the Version Provided in Windows Server 2003 with SP1**

1. In the domain where you want to configure Group Policy, on a computer running Windows Server 2003 with SP1, ensure that you have installed GPMC. On that computer, log on as a member of the Domain Admins group, or with an account from which you can manage GPOs in the domain.

   For information about installing GPMC, see "Installing Group Policy Management Console (GPMC)," earlier in this section.

2. Click **Administrative Templates** under either **User Configuration** or **Computer Configuration** (it does not matter which one). Be sure that you click it so that it is selected (highlighted). When you select **Administrative Templates**, the Administrative templates are updated throughout the domain.

   For more information about managing .adm files, see article 816662, "Recommendations for Managing Group Policy Administrative Template (.adm) Files" in the Microsoft Knowledge Base at:

   http://go.microsoft.com/fwlink/?LinkId=29128

| **Caution** |
| --- |
| You should treat the .adm files that are provided with the operating system as read-only. These files will be updated from time to time in service packs and in future releases. To use an existing .adm file as a template for a custom .adm file, first copy the original .adm file into a new filename and then edit the new file. |

# Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key

Applies To: Windows Server 2003 with SP1

**The Internet Communication Management Key in Group Policy**

Windows Server 2003 Service Pack 1 (SP1) contains Group Policy settings that were not available in the original release of Windows Server 2003. Some of these settings can be found in Group Policy under a new key called **Internet Communication Management**. This new key is located within GPMC in Computer Configuration\Administrative Templates\System and in User Configuration\Administrative Templates\System. This appendix describes the settings under the **Internet Communication Management** key, and the way in which one setting, **Restrict Internet communication**, controls multiple other settings.

The settings under the **Internet Communication Management** key are designed to help you control the way components in Windows Server 2003 with SP1 communicate with the Internet. Other settings can also help with controlling the way components communicate with the Internet.

Individual settings under **Internet Communication Management** for specific components are described in the appropriate sections of this white paper and in the Explain text in the settings, as well as being described in this appendix. For example, the setting called **Turn off Event Viewer "Events.asp" links**, located under **Internet Communication Management**, is described in the section about Event Viewer, in the Explain text in the setting itself, and in this appendix. Many other Group Policy settings (beyond the ones described in this appendix) are also described in appropriate sections of this white paper.

**Controlling Multiple Settings Through the Restrict Internet Communication Setting**

There are multiple ways to configure the new Group Policy settings under **Internet Communication Management** in Windows Server 2003 with SP1. You can configure settings individually, which means you could configure, for example, **Turn off Event Viewer "Events.asp" links** differently from **Turn off Windows Error Reporting**. Alternatively, you can enable or disable the entire collection of settings at one time by configuring a setting that controls them all. This setting is called **Restrict Internet communication**.

If you want to enable or disable **Restrict Internet communication** and then create exceptions to this master setting by configuring individual settings in the **Internet Communication Management** key, you must use two GPOs. To do this, ensure that you understand how processing and precedence works for multiple GPOs (for example, see the Help for GPMC). Choose or create a GPO that has lower precedence than another GPO. In the GPO with lower precedence, enable or disable **Restrict Internet communication**. Then, in the GPO that has precedence over it, apply the individual settings that are exceptions to the master setting. If you do not use two GPOs when setting both **Restrict Internet communication** and individual settings that are exceptions to the master setting, the settings might not work as expected.

To check the effect of multiple Group Policy settings, you can use **Group Policy Results** in GPMC. For more information, see the Help in GPMC.

**Settings That Affect Computer Configuration**

This subsection of the appendix describes settings under **Internet Communication Management** in the Administrative template for computer configuration (which contrasts with the template for user configuration, described later in this appendix). For all these settings, the following is true:

- The settings apply to all users of a computer running Windows Server 2003 with SP1 and come into effect when the computer starts or when Group Policy is refreshed.
- Within GPMC, these settings are located in Computer Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings.
- All of the settings can be enabled or disabled in one step by enabling or disabling the master setting that controls them, **Restrict Internet communication**. This setting is located in Computer Configuration\Administrative Templates\System\Internet Communication Management. When you change this setting, the change affects all of the settings in the following list, even if you previously configured individual settings in different ways.

**Note**

This appendix does not describe all Group Policy settings that are available in Windows Server 2003 with SP1. It describes only the settings available under **Internet Communication Management**. For sources of information about new settings available in Windows Server 2003 with SP1, see Appendix B: Resources for Learning About Group Policy.

**Individual Settings That Affect Computer Configuration**

The following list describes the computer configuration settings under **Internet Communication Management** that are outlined in the previous subsection.

**Note**

More details about each setting are available in the Explain text for the setting. To view the text, select the setting in Group Policy and click the **Extended** tab, or open the setting and click the **Explain** tab.

- **Turn off the "Publish to Web" task for files and folders**: Specifies whether the tasks for publishing items to the Web are available from File and Folder Tasks in Windows folders. These tasks are **Publish this file to the Web**, **Publish this folder to the Web**, and **Publish the selected items to the Web**.

  This setting is also described in Appendix J: Online Ordering Wizards, Add Network Place Wizard, and New Connection Wizard in this white paper.

- **Turn off Internet download for Web publishing and online ordering wizards**: Specifies whether Windows should download a list of providers for the Web Publishing Wizard, the Add Network Place Wizard, and the Online Print Wizard. By default, Windows displays providers downloaded from a Windows Web site in addition to providers specified in the registry. If you enable this setting, Windows will not download providers and only the service providers that are stored in the local registry will be displayed. When Windows Server 2003 with SP1 has been installed but the Web publishing and online ordering wizards have not yet been used, no service providers are stored in the local registry. If this Group Policy setting is applied at that time, the wizards will not display links to service providers.

  For more information about the registry keys in which providers can be specified, see "Registering a Service" at the MSDN Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29134

  This setting is also described in Appendix J: Online Ordering Wizards, Add Network Place Wizard, and New Connection Wizard in this white paper.

- **Turn off the "Order Prints" picture task**: Specifies whether the "Order Prints Online" task is available from Picture Tasks in Windows folders.

  This setting is also described in Appendix J: Online Ordering Wizards, Add Network Place Wizard, and New Connection Wizard in this white paper.

- **Turn off the Windows Messenger Customer Experience Improvement Program**: Affects Windows XP with SP2, not Windows Server 2003. This setting specifies whether Windows Messenger collects anonymous information about how Windows Messenger software and service is used. With the Customer Experience Improvement program, users can allow Microsoft to collect anonymous information about how the product is used. This information is used to improve the product in future releases.

- **Turn off Help and Support Center "Did you know?" content**: Specifies whether to show the Headines area of Help and Support Center. By default, the Headlines area in Windows XP is labeled **Did you know?** and in Windows Server 2003 is labeled **Top Issues**. By default, this content is dynamically updated when users are connected to the Internet and open Help and Support Center.

  This setting is also described in the section of this white paper titled Help and Support Center: Features That Communicate with the Internet.

- **Turn off Help and Support Center Microsoft Knowledge Base search**: Specifies whether users can perform a

Microsoft Knowledge Base search from Help and Support Center. The Knowledge Base is an online source of technical support information and self-help tools for Microsoft products. By default, if a user performs a search through Help and Support Center, the Knowledge Base is included in the search. It can be excluded if a user opens the Search Options page in Help and Support Center and clears the Microsoft Knowledge Base check box.

This setting is also described in the section of this white paper titled Help and Support Center: Features That Communicate with the Internet.

- **Turn off Windows Error Reporting**: Specifies whether error reports from a system or application that has stopped responding are sent to Microsoft. Error reports are used to improve the quality of the product. This setting overrides any user setting made from the Control Panel for error reporting.

  This setting and other ways of controlling error reporting through Group Policy are described in Windows Error Reporting and Internet Communication in this white paper.

- **Turn off Internet File Association service**: Specifies whether to use the Web-based File Association service, or whether to use only locally stored information about file name extensions, file types, and the applications or components to use when opening a particular file type. The file association Web service is used only when a user tries to open a file and there is no locally stored information about the file name extension.

  This setting is also described in File Association Web Service and Internet Communication in this white paper.

- **Turn off Internet Connection Wizard if URL connection is referring to Microsoft.com**: Affects Windows XP with SP2, not Windows Server 2003. This setting specifies whether the Internet Connection Wizard can connect to Microsoft to download a list of Internet Service Providers (ISPs). If you enable this setting, and a user running the New Connection Wizard clicks **Choose from a list of Internet Service Providers (ISPs)**, finishes the wizard, and clicks **Refer me to more Internet Service Providers**, a message appears, saying that the user cannot complete the Internet Connection Wizard. This prevents users from retrieving the list of ISPs, which resides on Microsoft servers.

- **Turn off Event Viewer "Events.asp" links**: Specifies whether Internet links shown within events in Event Viewer are activated. When such a link is activated and the user clicks it, information that identifies the event is sent to a Microsoft Web site so that explanatory text, if available, can be sent back to the user.

  This setting and the information sent and received when an Event Viewer link is clicked are described in Event Viewer and Internet Communication in this white paper.

- **Turn off Automatic Root Certificates Update**: Specifies whether to automatically update root certificates using the list of trusted certification authorities that Microsoft maintains on the Windows Update Web site. If you enable this setting, when a user is presented with a certificate issued by an untrusted root authority, the user's computer will not contact the Windows Update web site.

  This setting and the way that root certificate updates work are described in Certificate Support and Internet Communication in this white paper.

- **Turn off Registration if URL connection is referring to Microsoft.com**: Specifies whether the Windows Registration Wizard connects to Microsoft.com for online registration.

  This setting and the differences between registration and activation are described in the section of this white paper titled Activation, Registration, and Resulting Internet Communication in Windows Server 2003 with Service Pack 1.

- **Turn off Search Companion content file updates**: Specifies whether Search Companion should automatically download content updates during local and Internet searches.

  This setting and the way that Search Companion works for local and Internet searches are described in Search Companion and Internet Communication in this white paper.

- **Turn off printing over HTTP**: Specifies whether to allow printing over HTTP from this computer. Note that this setting does not control whether the computer can act as an Internet print server.

Internet printing is described briefly in the subsection titled "Internet Printing" in Internet Information Services and Internet Communication in this white paper.

- **Turn off downloading of print drivers over HTTP**: Specifies whether to allow this computer to download print drivers over HTTP when needed.

  Internet printing is described briefly in the subsection titled "Internet Printing" in Internet Information Services and Internet Communication in this white paper.

- **Turn off Windows Update device driver searching**: Specifies whether Windows searches Windows Update for device drivers when no local drivers for a device are present.

  This setting is described in Plug and Play and Internet Communication in this white paper. Related details about the searching of Windows Update for device drivers are described in Device Manager, Hardware Wizards, and Internet Communication in this white paper.

- **Turn off access to all Windows Update features**: Specifies whether Windows Update can be used to update the operating system on this computer.

  This setting is described in Windows Update, Automatic Updates, and Internet Communication in this white paper.

- **Turn off Windows Movie Maker automatic codec downloads**: Affects Windows XP with SP2, not Windows Server 2003. This setting specifies whether Windows Movie Maker automatically downloads missing codecs for audio and video files. A codec—compressor/decompressor—is software that compresses or decompresses audio or video data.

- **Turn off Windows Movie Maker online Web links**: Affects Windows XP with SP2, not Windows Server 2003. This setting specifies whether links to Web sites are available in Windows Movie Maker.

- **Turn off Windows Movie Maker saving to online video hosting provider**: Affects Windows XP with SP2, not Windows Server 2003. This setting specifies whether users can send a final movie to a video hosting provider on the Web.

#### Settings That Affect User Configuration

This subsection of the appendix describes settings under **Internet Communication Management** in the Administrative template for user configuration (which contrasts with the template for computer configuration). For all these settings, the following is true:

- The settings apply to the individual user and come into effect when the user logs on or when Group Policy is refreshed.

- Within GPMC, these settings are located in User Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings.

- All of the settings can be enabled or disabled in one step by enabling or disabling the master setting that controls them, **Restrict Internet communication**. This setting is located in User Configuration\Administrative Templates\System\Internet Communication Management. When you change this setting, the change affects all of the settings in the following list, even if you previously configured individual settings in different ways.

| **Note** |
| --- |
| This appendix does not describe all Group Policy settings available in Windows Server 2003 with SP1. It describes only the settings available under **Internet Communication Management**. For sources of information about new settings available in Windows Server 2003 with SP1, see Appendix B: Resources for Learning About Group Policy. |

#### Individual Settings That Affect User Configuration

The user configuration settings under **Internet Communication Management** are as follows. For a description of a particular setting, find the setting under "Individual Settings That Affect Computer Configuration," earlier in this section.

Alternatively, you can select the setting in Group Policy and click the **Extended** tab, or open the setting and click the **Explain** tab.

- Turn off the "Publish to Web" task for files and folders
- Turn off Internet download for Web publishing and online ordering wizards
- Turn off the "Order Prints" picture task
- Turn off the Windows Messenger Customer Experience Improvement Program
- Turn off Internet File Association service
- Turn off printing over HTTP
- Turn off downloading of print drivers over HTTP
- Turn off Windows Movie Maker automatic codec downloads
- Turn off Windows Movie Maker online Web links
- Turn off Windows Movie Maker saving to online video hosting provider

# Appendix D: Differences in Windows Server 2003 Before and After Service Pack 1

Applies To: Windows Server 2003 with SP1

In Windows Server 2003 Service Pack 1 (SP1), Microsoft introduces a set of security technologies that help to improve the ability of computers running Windows Server 2003 to withstand malicious attacks from viruses and worms. The technologies include additional network protection, memory protection, and improved computer maintenance. Also included are technologies for safer e-mail handling and enhanced browsing security. In addition, SP1 includes enhanced Administrative templates for Group Policy settings. Together, these security technologies will help to make it more difficult to attack Windows Server 2003 with SP1, even if the latest updates are not applied. These security technologies together are particularly useful in mitigation against worms and viruses.

This white paper provides information about some of the improvements in Windows Server 2003 SP1 as compared to the original release of Windows Server 2003. In particular, Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key provides information about some of the Group Policy settings available only in SP1 and later. These settings are designed to provide specific ways to control the way components in Windows Server 2003 with SP1 communicate with the Internet.

For additional information about the differences between the original release of Windows Server 2003 and Windows Server 2003 with SP1, see the following resources:

- The release notes for Windows Server 2003 SP1 on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=44311

- "Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1" on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46278

- Other information about Windows Server 2003 SP1 on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=36940

  http://go.microsoft.com/fwlink/?LinkId=46225

# Appendix E: Windows Firewall and Security Configuration Wizard

Applies To: Windows Server 2003 with SP1

Windows Firewall is a host firewall technology that replaces Internet Connection Firewall in Windows Server 2003 with Service Pack 1 (SP1) and Windows XP with Service Pack 2 (SP2). Windows Firewall is off by default on Windows Server 2003 with SP1, but is easily turned on.

Security Configuration Wizard (SCW) is an attack surface reduction tool for Windows Server 2003 with SP1. When you run it, you can create, edit, apply, or roll back a security policy (stored as an XML file). Any security policy that you create takes effect only when you apply it to a server.

**Overview: Windows Firewall and Security Configuration Wizard**
Windows Firewall provides stateful inspection of incoming Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic and is designed to help protect you from network attacks that pass through your perimeter network or originate inside your organization, such as Trojan horse attacks, port scanning attacks, and worms. By running Windows Firewall on each of your clients and servers, you can extend your defense-in-depth strategy to the innermost layer of your security architecture.

SCW determines the minimum functionality required for a server's role or roles, and disables functionality that is not required. For example, it can disable unneeded services and block unused ports. SCW guides you through the process of creating, editing, applying, or rolling back a security policy based on the selected roles of the server. The security policies that are created with SCW are XML files that, when applied, configure services, network security, specific registry values, audit policy, and if applicable, Internet Information Services (IIS).

**Group Policy Settings that Affect Windows Firewall**
Windows Server 2003 with SP1 includes Group Policy settings through which you can configure Windows Firewall. (Windows XP with SP2 also includes these settings.)

**To Locate the Group Policy Settings for Windows Firewall**

1. See Appendix B: Resources for Learning About Group Policy, for information about using Group Policy. Ensure that your Administrative templates have been updated, and then edit an appropriate GPO.

2. Click **Computer Configuration**, click **Administrative Templates**, click **Network**, click **Network Connections**, and then click **Windows Firewall**.

3. Click either **Domain Profile** or **Standard Profile**.

4. View the **Group Policy** objects that are available.

For more details about any of the Group Policy settings, use a Group Policy interface to navigate to the setting and then click the **Extended** tab, or open the setting and then click the **Explain** tab.

There is a setting through which you can disable Windows Firewall. This setting is located in Computer Configuration\Administrative Templates\Network\Network Connections\Windows Firewall\Domain Profile and is called **Windows Firewall: Protect all network connections**. If you disable this policy setting, Windows Firewall does not run and cannot be started.

| **Note** |
| --- |
| In Computer Configuration\Administrative Templates\Network\Network Connections, the setting called **Prohibit use of Internet Connection Firewall on your DNS domain network** still exists. This setting has no effect if **Windows Firewall: Protect all network connections** is enabled or disabled. However, if **Windows Firewall: Protect all network connections** is set to **Not Configured**, your servers are in a DNS domain, and you enable **Prohibit use of Internet Connection Firewall on your DNS domain network**, it will prevent Windows Firewall from running (Internet Connection Firewall is the former name for Windows Firewall). |

**Resources for Learning About Windows Firewall and Security Configuration Wizard**

The following resources provide detailed information about using Windows Firewall and Security Configuration Wizard:

- For how-to and conceptual topics about Windows Firewall, along with information about how Windows Firewall differs from Internet Connection Firewall (the earlier version of Windows Firewall), see the Help for Windows Firewall on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=33577

- By using an answer file for unattended or remote installation and including entries in the [WindowsFirewall] section, you can control Windows Firewall settings at deployment time. For more information, see the references listed in Appendix A: Resources for Learning About Automated Installation and Deployment. Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix).

- For the Windows Firewall Operations Guide, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=43262

- For a collection of overviews and guides for using SCW, see "Security Configuration Wizard for Windows Server 2003" at the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?linkid=45503

- For instructions for using Security Configuration Wizard, see the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46018

  http://go.microsoft.com/fwlink/?LinkId=46017

# Appendix F: Internet Connection Sharing and Network Bridge

Applies To: Windows Server 2003 with SP1

Internet Connection Sharing and Network Bridge are features designed for home and small office networks, and can be used with Windows Firewall (formerly called Internet Connection Firewall or ICF). These features are included in some of the Microsoft Windows Server 2003 family operating systems. Information about these features is presented here so you as an IT administrator can be aware of these potential capabilities within your organization's network.

The following list describes the editions and versions of Windows Server 2003 that do and do not include Internet Connection Sharing, Windows Firewall (formerly called Internet Connection Firewall), and Network Bridge:

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

For more information about Windows Firewall, see Appendix E: Windows Firewall and Security Configuration Wizard.

This appendix includes the following:

- An overview of Internet Connection Sharing and Network Bridge.

- How Internet Connection Sharing and Network Bridge can be used in a managed environment.

- How to control or prevent the use of Internet Connection Sharing and Network Bridge.

**Overview: Internet Connection Sharing and Network Bridge**

Internet Connection Sharing and Network Bridge provide the following functionality:

- Internet Connection Sharing (ICS)

  ICS provides Internet access for a home or small office network by using one common connection as the Internet gateway. In an ICS network, a single computer is chosen to be the ICS host. The ICS host has at least two network adapters: one connected to the Internet, one or more connected to the private network. All Internet-destined traffic flows through the ICS host. ICS uses DHCP to assign private IP addresses on the network, and Network Address Translation (NAT) to allow multiple computers on the private network to connect to the public network through the ICS host.

  In a home or small office environment, there can be security benefits in using ICS. Only the ICS host is visible from the Internet. The private network is "hidden." Also, NAT blocks any network traffic that did not originate from the private network or is a response to traffic originating from the private network.

  In addition, ICS provides name resolution to the home network through a DNS proxy.

  | **Note** |
  | --- |
  | You should not use Internet Connection Sharing in an existing network with Windows Server 2003 domain controllers, Domain Name System (DNS) servers, gateways, Dynamic Host Configuration Protocol (DHCP) servers, or systems configured for static IP addresses. |

- Network Bridge

  Network Bridge removes the need for routing and bridging hardware in a home or small office network that consists of multiple LAN segments. Network Bridge forwards traffic among the multiple LAN segments, making them appear to be a

single IP subnet.

| ⚠Caution |
| --- |
| If neither Windows Firewall nor ICS is enabled on your network, do not set up Network Bridge between the public Internet connection and the private network connection. Setting up Network Bridge between the public Internet connection and the private network connection creates an unprotected link between your network and the Internet, leaving your network vulnerable to external attacks. When either Windows Firewall or ICS is enabled, this risk is mitigated. |

**Using Internet Connection Sharing and Network Bridge in a Managed Environment**

On Windows Server 2003 with SP1, Internet Connection Sharing and Network Bridge are not enabled by default (and Windows Firewall is also not enabled by default). Also, Internet Connection Sharing (ICS) is available only on computers that have two or more network connections. An administrator or user with administrative credentials can enable ICS by clicking the **Advanced** tab on network connections (**Control Panel\Network Connections**). Also, when running the New Connection Wizard, administrators can choose to enable ICS. ICS lets administrators configure a computer as an Internet gateway for a small network, and it provides network services such as name resolution through Domain Name System (DNS). It also provides addressing through Dynamic Host Configuration Protocol (DHCP) to the local private network.

The **Network Bridge** menu command **Bridge Connections** is available only when two or more network adapters are present. By default, Network Bridge is disabled, but administrators can use **Bridge Connections** to enable Network Bridge.

In a domain environment, you should not allow these features to be enabled or configured. See the following subsection for information about how to disable them.

It is important to be aware of all the methods users and administrators have for connecting to your networked assets, and to review whether your security measures provide in-depth defense (as contrasted with a single layer of defense, more easily breached).

**Controlling the Use of Internet Connection Sharing and Network Bridge**

You can block administrators and users from accessing ICS and Network Bridge by using answer files during initial installation and by using Group Policy post-deployment.

**Using Answer Files for Unattended or Remote Installation**

Using standard methods for preparing an unattended or remote installation, you can make entries in the [Homenet] section of the answer file. If you make entries in the [Homenet] section, also investigate making entries in the [WindowsFirewall] section. Thes sections include entries for installing home and small office networking settings for network adapters, Internet Connection Sharing, Windows Firewall, and Network Bridge. For example, to prevent users and administrators from enabling Internet Connection Sharing by using an answer file, the entry is as follows:

```
[Homenet]
EnableICS = No
```

For additional configuration options for [Homenet] and [WindowsFirewall] entries for the answer file, and for more information about unattended installation, see the references listed in Appendix A: Resources for Learning About Automated Installation and Deployment. Be sure to review the information in the Deploy.chm file (whose location is provided in that appendix).

**Using Group Policy to Disable Internet Connection Sharing and Network Bridge**

Group Policy settings for disabling Internet Connection Sharing and Network in your domain environment are as follows. For information about Group Policy settings for Windows Firewall, see Appendix E: Windows Firewall and Security Configuration Wizard.

| Note |
| --- |
| For more details about any of the Group Policy settings, use a Group Policy interface to navigate to the setting and then click the **Extended** tab, or open the setting and then click the **Explain** tab. For other sources of information about Group Policy, see Appendix B: Resources for Learning About Group Policy. |

- **Prohibit use of Internet Connection Sharing on your DNS domain network**, located in Computer Configuration\Administrative Templates\Network\Network Connections.

If you enable this policy setting, ICS cannot be enabled or configured by administrators, and the ICS service cannot run on the computer. In the **Advanced** tab in the Properties dialog box for a local area network (LAN) or remote access connection, under **Internet Connection Sharing**, it says "Internet Connection Sharing has been disabled by the Network Administrator."

Also, if you enable this policy setting, the Internet Connection Sharing page is removed from the New Connection Wizard.

- **Prohibit installation and configuration of Network Bridge on your DNS domain network**, located in Computer Configuration\Administrative Templates\Network\Network Connections.

  When you enable this policy setting, administrators cannot create a Network Bridge. Enabling this policy setting does not remove an existing Network Bridge from a computer.

| **Important** |
| --- |
| The preceding policy settings (both with have "DNS" in the name of the setting) are dependent on the network context that the computer is in. They apply only when a computer is connected to the same DNS domain network it was connected to when the policy setting was refreshed on that computer. If a computer is connected to a DNS domain network other than the one it was connected to when the policy setting was refreshed, the policy setting does not apply. |

For more information about home and small office networking features, see Help for Windows Server 2003 with SP1 on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=46009

# Appendix G: Message Queuing and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The purposes of Message Queuing
- How Message Queuing handles messages
- Finding additional sources of information about Message Queuing

**Purposes of Message Queuing 3.0**

Message Queuing (MSMQ) is one of the optional components in Microsoft Windows Server 2003. Information about Message Queuing is presented here so you as an IT administrator can be aware of the potential capabilities for use within your organization's network, if you determine that you will be installing Message Queuing.

Message Queuing 3.0, the version in products in the Microsoft Windows Server 2003 family, is a message queuing service that enables applications on different systems to communicate with each other. Message Queuing implements asynchronous communication by enabling applications to send messages to and receive messages from other applications that may be running on the same computer or on different computers connected by a network. With Message Queuing, these applications can communicate across the Internet and other heterogeneous networks, and with computers that might be temporarily offline.

**Overview: Using Message Queuing in a Managed Environment**

Applications can use Message Queuing to send messages and to continue running regardless of whether the receiving application is running or reachable over the network. To create a message, an application specifies the message fields or properties and supplies the field values. The application then issues an MSMQ application programming interface (API) call to send the message. The MSMQ queue manager (the designated server) transmits the message to the destination message queue. If the destination location is not connected to the network when the message is sent, the message queuing system stores the message at an interim location. The system forwards the message automatically when a connection is established.

To receive a message, an application issues an API call that reads the message from the queue. When an application receives a request message, it processes the request by reading the contents of the message and acting accordingly. If required, the receiving application can send a response message back to the original requester. While in transit between senders and receivers, MSMQ keeps messages in holding areas called queues. MSMQ queues protect messages from being lost in transit and provide a place to retrieve messages when the receivers are ready to receive them.

Message Queuing 3.0 provides support for sending messages over the Internet. This support includes referencing queues by URLs, sending and receiving XML-based messages that are formatted using Simple Object Access Protocol (SOAP), and providing security for HTTP messaging. MSMQ also provides support for Web farms and firewalls. For more information about sending messages over the Internet, refer to "Related Links," later in this section.

It is beyond the scope of this white paper to describe all aspects of maintaining appropriate levels of security in an organization running applications that communicate across the Internet. This section, however, provides overview information as well as suggestions for other sources of information about balancing your organization's requirements for running Internet applications with your organization's requirements for protection of networked assets.

**Controlling Message Queuing to Prevent the Flow of Information to and from the Internet**

**Installing and Uninstalling Message Queuing**

Message Queuing is not installed by default. However, if your organization has determined that MSMQ is an essential component of the business process, it can be installed as described in this subsection. Although MSMQ is highly configurable, it is beyond the scope of this white paper to describe all the configuration options available to you. If you will be using MSMQ, you can refer to the "Related Links" subsection following these procedures for more detailed information about deploying Message Queuing.

**To Install or Uninstall Message Queuing**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.
2. Double-click **Add or Remove Programs**.

3. Click **Add/Remove Windows Components** (on the left).

4. In Windows Components, select **Application Server**, and then click **Details**.

5. Click **Message Queuing** and do one of the following:

    - If Message Queuing is installed and you want to remove it, clear the check box for **Message Queuing** and complete the Message Queuing wizard.

    - If Message Queuing is not installed and you want to add it, select the check box for **Message Queuing**.

6. If you are installing Message Queuing, select **Details**, select one or more of the options listed, and complete the Windows Components wizard.

**Accessing the Operating System Help Documentation for Message Queuing**

The operating system has Help documentation describing the use of Message Queuing. You can view this documentation from any computer that has Internet access (regardless of the operating system running on that computer), or from any server running a product in the Windows Server 2003 family.

**To Access the Help Documentation for Message Queuing**

Open Help for Windows Server 2003 by doing one of the following:

- Click **Start**, and then click **Help and Support**. Click **Application and Development Support**, and then click **Using Message Queuing**.

- View Help on the Web at:

  http://go.microsoft.com/fwlink/?LinkId=29881

**Related Links**

For more information about Message Queuing, see the following references:

- The Message Queuing Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46016

- The Message Queuing and Queued Components (MSMQ) SDK documentation, as described in the following list:

    - "Using Message Queuing"

      http://go.microsoft.com/fwlink/?LinkId=46015

    - "Internet Messaging"

      http://go.microsoft.com/fwlink/?LinkId=46014

    - "HTTP Messages"

      http://go.microsoft.com/fwlink/?LinkId=46013

    - "New Features for Windows XP and the Windows Server 2003 family"

      http://go.microsoft.com/fwlink/?LinkId=46012

# Appendix H: Connection Manager and Internet Communication

Applies To: Windows Server 2003 with SP1

This section provides information about:

- The benefits of Connection Manager 1.3
- How Connection Manager communicates with sites on the Internet
- Controlling Connection Manager to limit the flow of information to and from the Internet

**Benefits and Purposes of Connection Manager 1.3**

Connection Manager 1.3 is a client dialer and connection software program included with products in Microsoft Windows Server 2003. It provides support for local and remote connections to your network service using a network of access points, such as those available worldwide through Internet service providers (ISPs). Connection Manager 1.3 includes features that both simplify and enhance the implementation of connection support for you and your users.

Connection Manager is implemented using the Connection Manager Administration Kit (CMAK). CMAK is designed for network administrators, information officers, and other team members who are responsible for the design, development, testing, distribution, and support of connection software for users who connect to your Internet or network service.

You can customize Connection Manager so that the set of files you distribute to your users (called a service profile) can be easily installed and run. Using Connection Point Services you install Phone Book Service, which enables the distribution of phone books for a Connection Manager service profile.

In addition to the support for basic dial-up connections, you can also use the CMAK wizard to incorporate support for virtual private network (VPN) connections using tunneling protocols to tunnel through a public network (such as when dialing into an ISP to access an organization's server). Connection Manager can create a VPN connection through a preexisting dial-up session, local area network (LAN), or digital subscriber line (DSL).

For more information about Connection Manager, see the instructions in "Accessing Help Documentation for Connection Manager," later in this appendix.

**How Connection Manager Communicates with Sites on the Internet**

Administrators control whether to enable the phonebook download feature and which server the client computer sends the request to. The client computer initiates the connection to the phone book server in order to update its data. The phone book server sends one of the following elements to the Connection Manager client while there is a connection:

- Phone book update file: The server sends a phone book update file when the version of the phone book of the connecting client is earlier than the version of the phone book server.
- Replacement phone book: This is sent when there is a difference of more than five versions between the client and the server phone book file versions.
- Standard HTTP error message

The following list describes various aspects of the Connection Manager information that is sent to and from the Internet and how the exchange of information takes place:

- **Specific information sent or received**: The information sent by Connection Manager in the phone book server query string contains the following elements:
    - **osarch**: Specifies the type of processor that the requesting computer is using.
    - **ostype**: Specifies the version of the Windows operating system of the computer making the request.
    - **cmver**: Specifies the version of Connection Manager of the computer making the request.

    - **lcid**: Specifies the locale identifier (LCID) describing the system default locale information of the computer making the request.
    - **pbver**: Specifies the current version of the phone book being used by the computer making the request.
    - **pb**: Specifies the name of the phone book being used by the computer making the request.

- **Default settings**: By default, Connection Point Services and the Connection Manager Administration Kit are not installed with Windows Server 2003. (You can add these components by using Add or Remove Programs in Control Panel.)
- **Triggers**: Connection Manager is triggered when users start it to connect to the Internet or their organization's network.
- **Logging**: Information is logged in the Connection Manager log file but not in the operating system event log. The logging feature for Connection Manager is used for troubleshooting and can be disabled.
- **Encryption**: The information contained in the query string is sent as plaintext and is not encrypted.
- **Storage**: The information that is sent is used by the phone book server to process the request being made, although the information is not stored.
- **Transmission protocol and port**: The transmission protocol used is HTTP and the port is 80.
- **Ability to disable**: On the server, you can prevent Connection Point Services and the Connection Manager Administration Kit from being installed, which prevents Connection Manager from being used. For clients, you can disable Connection Manager by not providing the user with a Connection Manager service profile. Also for clients, the Connection Manager phone book download and synchronization feature can be disabled using a Connection Manager service profile.

For more information about the Connection Manager Administration Kit (CMAK), see "Before you start: Understanding Connection Manager and the Administration Kit," on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=46011

You can also view the same information in Help, as described in "Accessing Help Documentation for Connection Manager," later in this appendix.

**Controlling Connection Manager to Limit the Flow of Information to and from the Internet**
Connection Manager client software is installed by default. Administrators determine when to create and distribute a Connection Manager service profile, without which Connection Manager cannot run.

You can limit the use of Connection Manager by creating and distributing a service profile only to users and administrators who need to use the Connection Manager client to initiate local or remote network access point connections. You can use the Connection Manager Administration Kit (CMAK) wizard to create a customized service profile to configure Connection Manager not to update the client phone book information.

For a list of items to consider when designing a service profile for Connection Manager, see "Security information for the Connection Manager Administration Kit" on the Microsoft Web site at:

http://go.microsoft.com/fwlink/?LinkId=46010

**Procedures for Installing the Connection Manager Administration Kit and for Viewing Help**
This subsection lists procedures for the following tasks:

- Installing the Connection Manager Administration Kit on your server
- Accessing Help documentation for Connection Manager

**Installing Connection Manager Administration Kit on Your Server**
The Connection Manager Administration Kit (CMAK) and Connection Point Services (used to install Phone Book Service) are not installed by default on Windows Server 2003. Use the following procedure to add these components.

**To Add the Connection Manager Administration Kit and Connection Point Services After Server Installation**

1. Click **Start**, and then either click **Control Panel**, or point to **Settings** and then click **Control Panel**.

2. Double-click **Add or Remove Programs**.

3. Click **Add/Remove Windows Components** (on the left).

4. In the Windows Components Wizard, scroll down and double-click **Management and Monitoring Tools**.

5. Select **Connection Manager Administration Kit** and **Connection Point Services** and any other subcomponents you want to install.

6. Follow the instructions to complete the Windows Components Wizard.

**Accessing Help Documentation for Connection Manager**

Windows Server 2003 has Help documentation describing the Connection Manager Administration Kit. You can view this documentation from any computer that has Internet access (regardless of the operating system running on that computer), or from any server running Windows Server 2003. The following procedure provides the details.

**To Access Help Documentation for a Server Running Windows Server 2003**

1. Open Help for Windows Server 2003 by doing one of the following:

   - On any computer running Windows Server 2003, click **Start**, and then click **Help and Support**.

   - View Help on the Web at:

     http://go.microsoft.com/fwlink/?LinkId=29881

     As appropriate, navigate from this Web site to the documentation for the server product you are using.

2. To view information about Connection Manager Administration Kit, navigate to Network Services\Managing Remote Connections\Connection Manager Administration Kit.

# Appendix I: Passport Manager Administration

Applies To: Windows Server 2003 with SP1

The following appendix provides information about:

- The benefits and purposes of Passport Manager Administration
- Using Passport Manager Administration in a managed environment
- Resources for learning about Passport Manager Administration

**Benefits and Purposes of Passport Manager Administration**
Passport Manager Administration is a utility designed for use on Web sites on which you want to use Microsoft .NET Passport for authentication. In Microsoft Windows Server 2003 family operating systems, Passport Manager Administration provides a graphical-interface alternative to editing the registry to change Passport Manager object settings. The most common use of the Passport Manager Administration utility is to change the default object settings that Passport Manager methods use if certain optional parameters are omitted when the methods are called.

Although some parameters of Microsoft .NET Passport methods are listed as optional, these methods still generate URLs or results that reflect inherent default values when called. These default parameter values can be used to provide consistent values for your Web site, such as the required time window within which all users must be authenticated.

Primary default object settings that affect Passport Manager implementation include: Time Window, Force Login, and Language ID.

**Using Passport Manager Administration in a Managed Environment**
By configuring options in Passport Manager Administration to work with software on your Web site, you can control how Web browsers are authenticated to your Web site. With Passport Manager Administration, you can configure the settings for cookies, such as choosing to disable the use of cookies. You can also choose to have cookies deleted once users sign out of your Web site. (Cookies are a form of data passed both ways between Web browsers and servers. Cookies sent to a user's browser have some implications for privacy.)

To ensure that cookies are deleted when users sign out of a Web site, make sure the following conditions are met:

- The Site ID value in Passport Manager Administration should be the value given to you when you received your Passport encryption key.
- The path and file name of the page to which cookies are written (specified in Passport Manager Administration) should match the ExpireCookieURL provided when you registered your site. Also, if cookies are written to a page in a subdomain of your own domain, you should precede the domain path with a dot (".").
- The path to the page where cookies are written should be the root directory of the Web site. This ensures that cookies written to all possible paths are deleted. If this is not done, not all the cookies may be deleted.
- If your site uses multiple servers with Passport Manager installed on them, you should verify that all of the servers are writing the cookies to the same domain and path.

**Related Links**
You can find additional information about Passport Manager Administration at these sites:

- For more information about .NET Passport and how to download it, see the SDK documentation on the MSDN Web site at:

  http://go.microsoft.com/fwlink/?linkid=7739

- The Help file for Passport Manager Administration utility provides detailed information about configuring settings for this utility. You can find this Help file at:

http://go.microsoft.com/fwlink/?linkid=7738

# Appendix J: Online Ordering Wizards, Add Network Place Wizard, and New Connection Wizard

Applies To: Windows Server 2003 with SP1

This appendix provides information about the following:

- Online ordering wizards and tasks
- Add Network Place wizard
- New Connection wizard

While these features are not primarily designed for use on servers, information about them is presented here so IT administrators can be aware of these potential capabilities within the organization's network.

**Online Ordering Wizards and Tasks**
In Microsoft Windows Server 2003, you can click tasks in Windows Explorer to run several different wizards to order picture prints online, sign up for a service that offers online storage space, or publish files that can be viewed in a browser. These wizards present you with one or more service providers from which to request the online service. The task or wizard obtains the names and URLs of these service providers from two sources: a list stored locally (in the registry) and a list stored on a Microsoft Web site. By default, Windows displays providers from a list on the Microsoft Web site in addition to providers listed in the registry.

As an administrator, you can use several Group Policy settings in Windows Server 2003 with SP1 to control the way the wizards and tasks work.

**Wizards and Tasks That Can Be Used for Online Ordering**
In Microsoft Windows Server 2003, you can use the following tasks in Windows Explorer and the following wizards for online ordering.

- **Online Print Ordering Wizard and related tasks**: In Windows Explorer, when common tasks are shown and you select either a picture or a folder that is configured as a picture folder, you can choose picture tasks including **Order prints online**. This starts the Online Print Ordering Wizard.

  A folder can be configured as a picture folder by right-clicking the folder, clicking **Properties**, clicking the **Customize** tab and, under **Use this folder type as a template**, selecting **Pictures or Photo Album**.

- **Add Network Place Wizard and related tasks**: This wizard and the related tasks are described in "Add Network Place Wizard and Web Publishing Wizard," later in this appendix. You can use the wizard to sign up for a service that offers online storage space.

- **Web Publishing Wizard and related tasks**: This wizard and the related tasks are described in "Add Network Place Wizard and Web Publishing Wizard," later in this appendix. You can use the wizard to publish files over the Internet or a local network so that those files can be viewed in a Web browser.

**Controlling the Way Online Ordering Wizards and Tasks Communicate with the Internet**
You can control the way online wizards and tasks communicate with the Internet by using Group Policy settings available in Windows Server 2003 with SP1. Because the following Group Policy settings are part of Windows Server 2003 with SP1 but were not part of the original release of Windows Server 2003, in order to use them, you must ensure that your Administrative templates have been updated before you can use these settings. For more information, see Appendix B: Resources for Learning About Group Policy.

- To prevent the preceding wizards and tasks from accessing the Microsoft Web site for the list of providers to offer to users, configure the following Group Policy setting in Computer Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings: **Turn off Internet download for Web publishing and online ordering wizards**.

By default, the Web publishing and online ordering wizards display providers downloaded from a Windows Web site in addition to providers specified in the registry. If you enable this setting, Windows will not download providers and only the service providers that are stored in the local registry will be displayed. When Windows Server 2003 with SP1 has been installed but the Web publishing and online ordering wizards have not yet been used, no service providers are stored in the local registry. If this Group Policy setting is applied at that time, the wizards will not display links to service providers.

For more information about the registry keys in which providers can be specified, see "Registering a Service" on the MSDN Web site at:

http://go.microsoft.com/fwlink/?LinkId=29134

- To disable the Web Publishing Wizard and related File and Folder Tasks, see "Add Network Place Wizard and Web Publishing Wizard," later in this section.

- To disable the Online Print Ordering Wizard, configure the following Group Policy setting in Computer Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings: **Turn off the "Order Prints" picture task**.

| **Note** |
| --- |
| If you want the policy settings to apply to all users of a computer and come into effect when the computer starts or when Group Policy is refreshed, use the settings listed in the preceding paragraphs. If you want the policy settings to apply to users and come into effect when users log on or when Group Policy is refreshed, configure the same settings in User Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings. |

For information about other settings that control the way wizards and other components communicate with the Internet, see Appendix C: Group Policy Settings Listed Under the Internet Communication Management Key.

**Add Network Place Wizard and Web Publishing Wizard**

In the Microsoft Windows Server 2003 operating systems, you can use the Add Network Place Wizard to create shortcuts to shared folders and resources on the network or on Web or File Transfer Protocol (FTP) servers. You can use the Web Publishing Wizard to publish files over the Internet or a local network so that those files can be viewed in a Web browser. If you do not have folders on a Web server already or do not have an account with a Web service provider, the Add Network Place Wizard or the Web Publishing Wizard can help you set these up.

This appendix includes the following:

- An overview of the Add Network Place Wizard and the Web Publishing Wizard

- How to control the use of the Add Network Place Wizard and the Web Publishing Wizard

**Overview: Add Network Place Wizard and Web Publishing Wizard**

The Add Network Place Wizard and the Web Publishing Wizard are enabled by default. You can use these wizards as follows:

- **Add Network Place Wizard**: You can access this wizard through **My Network Places\Add Network Place** or in Windows Explorer through **Tools\Map Network Drive** (through the link labeled "Sign up for online storage or connect to a network server").

  The wizard can be used to sign up for a service that offers online storage space, or to create a shortcut to a Web site, an FTP site, or other local network connection. To add a shortcut in My Network Places to a folder on a Web server, the Web server must support network places. Supporting network places requires the Web Extender Client (WEC) protocol and Microsoft FrontPage® Server Extensions, or the Web Distributed Authoring and Versioning (WebDAV) protocol and Internet Information Services (IIS). You must also have read and write access to the Web server.

- **Web Publishing Wizard**: In Windows Explorer, when common tasks are shown, you can access this wizard through several different tasks. These tasks are **Publish this file to the Web**, **Publish this folder to the Web**, and **Publish the selected items to the Web**.

You can use the Web Publishing Wizard to publish files over the Internet or a local network so that those files can be viewed in a Web browser.

You can remove access to the Add Network Place Wizard and the Web Publishing Wizard using Group Policy.

For more information about the WEC and WebDAV protocols, see "About Web Folder Behaviors" on the MSDN Web site at:

http://go.microsoft.com/fwlink/?LinkId=29223

**Controlling the Use of the Add Network Place Wizard and the Web Publishing Wizard**
To control the Add Network Place Wizard and the Web Publishing Wizard, use Group Policy.

**Controlling Whether the Wizards Download Information from a Microsoft Web site**
To control whether the Add Network Place Wizard and the Web Publishing Wizard display a list of service providers downloaded from a Windows Web site, use the Group Policy setting **Turn off Internet download for Web publishing and online ordering wizards**. This setting is described in "Online Ordering Wizards and Tasks," earlier in this appendix.

**Controlling the Use of the Add Network Place Wizard**
To block users and administrators from accessing the Add Network Place Wizard from the **Map Network Drive** command, configure the following Group Policy setting in User Configuration\Administrative Templates\Windows Components\Windows Explorer: **Remove "Map Network Drive" and "Disconnect Network Drive"**. This policy setting does the following:

- If you enable this policy setting, the system removes the **Map Network Drive** and **Disconnect Network Drive** commands from the toolbar and **Tools** menus in Windows Explorer and My Network Places. The system also removes these commands from menus that appear when you right-click the Windows Explorer or My Network Places icons.

- This policy setting does not prevent you from connecting to another computer by typing the name of a shared folder in the **Run** dialog box.

- This policy setting does not remove the **Add Network Place** task from My Network Places.

**Controlling the Use of the Web Publishing Wizard**
To control access to the Web Publishing Wizard and related tasks, configure a Group Policy setting. The File and Folder Tasks affected by this setting are **Publish this file to the Web**, **Publish this folder to the Web**, and **Publish the selected items to the Web**. The setting is located in Computer Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings, and is called **Turn off the "Publish to Web" task for files and folders**.

If you want the policy setting to apply to all users of a computer and come into effect when the computer starts or when Group Policy is refreshed, use the setting listed in the preceding paragraph. If you want the policy setting to apply to a particular set of users and come into effect when users log on or when Group Policy is refreshed, use the same setting in User Configuration\Administrative Templates\System\Internet Communication Management\Internet Communication settings.

**New Connection Wizard**
In Microsoft Windows Server 2003 operating systems, you use the New Connection Wizard to create Internet and other types of network connections for home and small office networks. While this feature is designed for home and small office use, information about this feature is presented here so IT administrators can be aware of these potential capabilities within the organization's network.

This appendix includes the following:

- An overview of the New Connection Wizard

- How to control the use of the New Connection Wizard

**Overview: New Connection Wizard**
The New Connection Wizard in Windows Server 2003 operating systems replaces the Windows 2000 Network Connection Wizard and Internet Connection Wizard. Administrators for a home or small office network can use the New Connection Wizard to create any type of network connection including Internet, incoming, dial-up, virtual private network (VPN), and direct connection.

You can start the New Connection Wizard in a variety of ways:

- By starting any program that requires an Internet connection when no Internet connection has yet been configured. An example of such a program is Internet Explorer.

- By opening Network Connections and clicking **Create a new connection**. Network Connections can be opened in a variety of ways, including through **Control Panel\Network Connections**.

- By opening Internet Options and, on the **Connections** tab, clicking the **Setup** button. Internet Options can be opened in a variety of ways, including through **Control Panel\Internet Options**.

- By clicking **Start**, clicking **Programs** or **All Programs**, clicking **Accessories**, clicking **Communications**, and then clicking **New Connection Wizard**.

**Controlling the Use of the New Connection Wizard**

You can control how administrators and users can use the New Connection Wizard by configuring Group Policy.

| **Note** |
| --- |
| Some of the Group Policy settings described in this section do not affect administrators. The descriptions provide details. |

- Group Policy setting: **Disable Internet Connection Wizard**

  This Group Policy setting is in User Configuration\Administrative Templates\Windows Components\Internet Explorer. It affects administrators as well as users.

  If you enable this policy, the **Setup** button on the **Connections** tab in the **Internet Options** dialog box appears dimmed.

  | **Note** |
  | --- |
  | Note that this policy overlaps with the **Disable the Connections page** policy (located in \User Configuration\Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel), which removes the **Connections** tab from the interface. |

- Group Policy setting: **Prohibit Access to the New Connection Wizard**

  This Group Policy setting is in User Configuration\Administrative Templates\Network\Network Connections. It affects users but not administrators.

  This policy setting determines whether users can use the New Connection Wizard, which creates new Internet or intranet connections.

- Group Policy setting: **Enable Windows 2000 Network Connections settings for Administrators**

  This Group Policy setting is also in User Configuration\Administrative Templates\Network\Network Connections.

  | **Note** |
  | --- |
  | This policy setting is intended to be used in a situation in which the Group Policy object (GPO) contains computers running both Windows 2000 and Windows Server 2003, and identical Network Connections policy setting behavior is required between those computers. |

With this policy setting enabled, policy settings that exist in both Windows 2000 and Windows Server 2003 behave the same for administrators.

The set of Network Connections policy settings that exists in Windows 2000 also exists in Windows Server 2003. In Windows 2000, all of these policy settings have the ability to prohibit the use of certain features by administrators. By default, Network Connections policy settings in Windows Server 2003 do not prohibit the use of features from administrators.

For information about using Group Policy, see Appendix B: Resources for Learning About Group Policy.

# Related Links for Controlling Communication with the Internet

Applies To: Windows Server 2003 with SP1

This section contains a list of the Web sites that are found in other sections of this white paper, with some additional Web sites that provide general information about Microsoft products.

**Links to Product Information, Support Information, TechNet, Microsoft Developer Network, and Information in Resource Kits**
The following sites provide information about Windows Server 2003 operating systems and other Microsoft products. The list includes sites containing product documentation as well as other basic sites that provide information about Microsoft operating systems and other Microsoft products:

- Windows Server 2003 operating systems:
    - http://go.microsoft.com/fwlink/?LinkId=34401
    - http://go.microsoft.com/fwlink/?LinkId=46002
    - http://go.microsoft.com/fwlink/?linkid=44311
    - http://go.microsoft.com/fwlink/?linkid=36940
    - http://go.microsoft.com/fwlink/?LinkId=46225
    - http://go.microsoft.com/fwlink/?LinkId=46278
- Windows Server 2003 Help on the Web:

  http://go.microsoft.com/fwlink/?linkid=29881
- Windows XP with Service Pack 2 (SP2), specifically, controlling the way that components in Windows XP with SP2 communicate with the Internet:

  http://go.microsoft.com/fwlink/?LinkId=29133
- Windows Server Catalog:

  http://go.microsoft.com/fwlink/?LinkId=4303
- Microsoft product support:

  http://support.microsoft.com/
- Microsoft TechNet:

  http://www.microsoft.com/technet/
- Microsoft Developer Network:

  http://msdn.microsoft.com/
- Prescriptive Architecture Guides:

  http://go.microsoft.com/fwlink/?LinkId=29413
- *Windows Deployment and Resource Kits* (includes links to the *Windows Server 2003 Deployment Kit* and the *Windows Server 2003 Technical Reference*):

http://www.microsoft.com/windows/reskits/

- Microsoft Resource Kits:

  http://go.microsoft.com/fwlink/?linkid=29894

- Internet Explorer Administration Kit:

  http://go.microsoft.com/fwlink/?linkid=29479

- NetMeeting 3 Resource Kit:

  http://go.microsoft.com/fwlink/?LinkId=36134

**Links to Information About Security, Management, and Deployment**
The following sites provide information about security, management, and deployment topics:

- Managing mobile code:

  http://go.microsoft.com/fwlink/?linkid=29170

- *Automating and Customizing Installations* in the *Microsoft Windows Server 2003 Deployment Kit*:

  http://go.microsoft.com/fwlink/?LinkId=29887

- "Installation and Setup Technologies" in the Core Operating System Collection of the *Windows Server 2003 Technical Reference* on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46281

- Using scripts:

  - http://go.microsoft.com/fwlink/?LinkId=24771
  - http://go.microsoft.com/fwlink/?LinkId=29221

- Security:

  - http://www.microsoft.com/security/guidance/
  - http://go.microsoft.com/fwlink/?LinkId=29168
  - http://go.microsoft.com/fwlink/?LinkId=46003
  - http://msdn.microsoft.com/security/
  - http://www.microsoft.com/technet/security/
  - http://go.microsoft.com/fwlink/?LinkId=46004

**Links to Information About Components in Windows Server 2003 SP1**
The following sites provide information about some of the components in Windows Server 2003 SP1:

- Add Network Place Wizard (specifically, related protocols called the Web Extender Client protocol and the Distributed Authoring and Versioning protocol):

http://go.microsoft.com/fwlink/?LinkId=29223

- Certificates, certificate status, certificate revocation, and Public Key Infrastructure (PKI):

  - http://go.microsoft.com/fwlink/?linkid=29886
  - http://go.microsoft.com/fwlink/?LinkId=27081

  For more information, also see "Links to Sites Maintained by Task Forces and Other Organizations," later in this section.

- Connection Manager:

  - http://go.microsoft.com/fwlink/?LinkId=46011
  - http://go.microsoft.com/fwlink/?LinkId=46010

- Dynamic Update:

  http://go.microsoft.com/fwlink/?linkid=29887

- File association Web service (specifically, the related process of using Group Policy to remotely install software):

  http://go.microsoft.com/fwlink/?linkid=29166

  For more information about the file association Web service, also see "Language codes" in this list.

- HyperTerminal:

  http://go.microsoft.com/fwlink/?linkid=29890

- Internet Connection Sharing and Network Bridge:

  http://go.microsoft.com/fwlink/?LinkId=46009

- Internet Explorer:

  http://www.microsoft.com/windows/ie/

  For more information about Internet Explorer, also see the following:

  - "Microsoft Resource Kits" and "Internet Explorer Administration Kit" in "Links to Product Information, Support Information, TechNet, Microsoft Developer Network, and Information in Resource Kits," earlier in this section.
  - "Language codes" in this list.

- Internet Information Services:

  - http://go.microsoft.com/fwlink/?LinkId=46005
  - http://go.microsoft.com/fwlink/?linkid=22805
  - http://go.microsoft.com/fwlink/?linkid=31134
  - http://go.microsoft.com/fwlink/?linkid=29896

- Internet Protocol version 6 (IPv6):

  http://go.microsoft.com/fwlink/?LinkId=29519

  For more information about IPv6, also see “Links to Sites Maintained by Task Forces and Other Organizations,” later in this section.

- Language codes (used in Internet Explorer and the file association Web service when a language is being specified):

  http://go.microsoft.com/fwlink/?linkid=29165

- Message Queuing and the Message Queuing SDK:

  - http://go.microsoft.com/fwlink/?LinkId=46016
  - http://go.microsoft.com/fwlink/?LinkId=46015
  - http://go.microsoft.com/fwlink/?LinkId=46014
  - http://go.microsoft.com/fwlink/?LinkId=46013
  - http://go.microsoft.com/fwlink/?LinkId=46012

- NetMeeting:

  - http://go.microsoft.com/fwlink/?LinkId=29514
  - http://go.microsoft.com/fwlink/?LinkId=29175
  - http://go.microsoft.com/fwlink/?LinkId=29206
  - http://go.microsoft.com/fwlink/?LinkId=29207

  For more information about NetMeeting, also see the following:

  - "Links to Sites Maintained by Task Forces and Other Organizations," later in this section.
  - "NetMeeting 3 Resource Kit" in "Links to Product Information, Support Information, TechNet, Microsoft Developer Network, and Information in Resource Kits," earlier in this section.

- Online Print Ordering Wizard:

  http://go.microsoft.com/fwlink/?LinkId=29134

- Passport Manager Administration:

  - http://go.microsoft.com/fwlink/?linkid=7739
  - http://go.microsoft.com/fwlink/?linkid=7738

- Program Compatibility Wizard (specifically, the Application Compatibility Toolkit, which can be used with the Program Compatibility Wizard):

  http://go.microsoft.com/fwlink/?LinkId=29880

- Remote Assistance:

http://go.microsoft.com/fwlink/?LinkId=29212

- Search Companion:

  http://sa.windows.com/privacy/

- Security Configuration Wizard (see also "Windows Firewall" in this list):
  - http://go.microsoft.com/fwlink/?linkid=45503
  - http://go.microsoft.com/fwlink/?LinkId=46018
  - http://go.microsoft.com/fwlink/?LinkId=46017
- Terminal Server Licensing (and Terminal Server):

  http://go.microsoft.com/fwlink/?LinkId=46192

  http://go.microsoft.com/fwlink/?LinkId=46078

  http://go.microsoft.com/fwlink/?LinkId=46007

  http://go.microsoft.com/fwlink/?linkid=24427

- Windows Error Reporting:
  - http://go.microsoft.com/fwlink/?LinkId=825
  - http://go.microsoft.com/fwlink/?LinkId=29903
  - http://go.microsoft.com/fwlink/?LinkId=29517
- Windows Firewall (see also "Security Configuration Wizard" in this list):
  - http://go.microsoft.com/fwlink/?linkid=33577
  - http://go.microsoft.com/fwlink/?linkid=43262
- Windows Media Player:
  - http://www.microsoft.com/Windows/WindowsMedia/
  - http://go.microsoft.com/fwlink/?LinkId=46006
  - http://go.microsoft.com/fwlink/?LinkId=46008
  - http://go.microsoft.com/fwlink/?LinkId=46465
- Windows Media Services:
  - http://go.microsoft.com/fwlink/?LinkId=29862
  - http://go.microsoft.com/fwlink/?LinkId=46019
  - http://go.microsoft.com/fwlink/?LinkId=46020

- Windows Time Service:
    - http://go.microsoft.com/fwlink/?LinkId=46021
- Windows Update and Automatic Updates:
    - http://windowsupdate.microsoft.com/
    - http://go.microsoft.com/fwlink/?LinkId=29906

**Links to Information About Licensing, Product Activation, and Registration**
The following sites provide information about licensing, product activation, and registration:

- http://go.microsoft.com/fwlink/?LinkId=29878
- http://go.microsoft.com/fwlink/?linkid=29225
- http://go.microsoft.com/fwlink/?LinkId=29484
- http://go.microsoft.com/fwlink/?LinkId=29923
- http://go.microsoft.com/fwlink/?linkid=29226
- http://go.microsoft.com/fwlink/?LinkId=29508

**Links to Sites Maintained by Task Forces and Other Organizations**
The following sites are maintained by the Internet Engineering Task Force:

- http://go.microsoft.com/fwlink/?LinkId=29898
- http://go.microsoft.com/fwlink/?LinkId=29135
- http://go.microsoft.com/fwlink/?LinkId=29136
- http://go.microsoft.com/fwlink/?LinkId=29215
- http://go.microsoft.com/fwlink/?LinkId=29924

The following site is maintained by the International Multimedia Telecommunications Consortium:

http://www.imtc.org/

The following site is maintained by the International Telecommunication Union:

http://go.microsoft.com/fwlink/?LinkId=29510

(Web addresses can change, so you might be unable to connect to the Web sites mentioned here.)

**Links to Information About Group Policy**
The following sites provide information about topics related to Group Policy:

- Windows Server 2003 Group Policy page on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=46022

- The Group Policy Technology Center Web site at:

http://www.microsoft.com/grouppolicy/

- *Designing a Managed Environment* in the *Microsoft Windows Server 2003 Deployment Kit* at:

  http://go.microsoft.com/fwlink/?linkid=44319

- The Group Policy Collection of the *Windows Server 2003 Technical Reference* at:

  http://go.microsoft.com/fwlink/?LinkId=29907

- Article 816662, "Recommendations for Managing Group Policy Administrative Template (.adm) Files" in the Microsoft Knowledge Base at:

  http://go.microsoft.com/fwlink/?LinkId=29128

- Information about the Group Policy Management Console at the Microsoft Windows Server 2003 Web site at:

  http://go.microsoft.com/fwlink/?linkid=29909

- The Group Policy Settings Reference on the Microsoft Web site at:

  http://go.microsoft.com/fwlink/?LinkId=29911

# User Data and Settings Management

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Published: March 1, 2002

By Craig Marl

This white paper explains the IntelliMirror user data and settings management features in Microsoft Windows Server 2003. These key components of change and configuration management can help administrators reduce total cost of ownership (TCO).

**In This White Paper**

User Profiles Overview in User Data and Settings Management

User Profile Structure

Enhancements to User Profiles in Windows Server 2003 and Windows XP

How to Configure a Roaming User Profile

Security Considerations when Configuring Roaming User Profiles

Best Practices for User Profiles

Folder Redirection Overview

How to Configure Folder Redirection

Security Considerations when Configuring Folder Redirection

Best Practices for Folder Redirection in User Data and Settings Management

Related Technologies: Offline Files and Synchronization Manager

Common Scenarios for IntelliMirror User Data and Settings Features

Appendix: Group Policy Settings for Roaming User Profiles

Related Links for User Data and Settings Management

User data includes the documents, images, spreadsheets, presentations and e-mail messages on a user's computer. User settings include application configurations, preferences, window sizes, toolbar settings and so forth on a users computer.

With Microsoft IntelliMirror management technologies, administrators can manage user data and settings in ways that reduce the total cost of ownership (TCO) for the computing systems.

By using IntelliMirror on both the server and client, administrators can protect and manage user data and settings. Non-recoverable data from local workstations can be copied to servers, where it can be easily backed up and centrally managed. Personalized data, applications, and settings can follow each user to different computers throughout the network. Administrators can easily replace faulty computers and restore all user data and settings on a new computer.

When fully deployed, IntelliMirror uses the Active Directory service and Group Policy for policy-based management of user desktops. A Windows 2000, Windows XP Professional or Windows Server 2003 computer can be automatically configured to meet specific requirements of a users business roles, group memberships, and location. Group Policy and the Active Directory are not necessary for every IntelliMirror feature. Some of the features can be set on the local level or through local polices. An organization can tailor use of IntelliMirror to its needs.

This paper discusses two of the key components that provide user data and settings management in IntelliMirrorUser Profiles and Folder Redirection. It also provides an architectural overview of these features, and presents sample scenarios showing how IntelliMirror is used throughout a computers lifecycle.

# User Profiles Overview in User Data and Settings Management

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A user profile describes the desktop computing configuration for a specific user, including the users environment and preference settings.

A profile is created the first time that a user logs on to a computer running Windows Server 2003, Windows XP, Windows 2000, or Windows NT Workstation. A user profile is a group of settings and files that defines the environment that the system loads when a user logs on. It includes all the user-specific configuration settings, such as program items, screen colors, network connections, printer connections, mouse settings, and window size and position. Profiles are not user policies and the user has a profile even if you don't use Group Policy.

A user's data can be stored on the local hard disk drive, or IntelliMirror can be set so that the data roams with the user wherever he or she logs on. User data can include shortcuts to executable files, personal files, and user settings, such as a custom dictionary.

Depending on how you manage your network, you or a user can define the desktop settings.

The following user profiles are available in Windows Server 2003, Windows XP Professional, and Windows 2000 Professional:

- **Local User Profile.** Created the first time that a user logs on to a computer, the local user profile is stored on a computer's local hard disk. Any changes made to the local user profile are specific to the computer on which the changes are made.

- **Roaming User Profile.** A copy of the local profile is copied to, and stored on a server share. This profile is downloaded every time that a user logs on to any computer on the network, and any changes made to a roaming user profile are synchronized with the server copy upon logoff.

- **Mandatory User Profile.** A type of profile that administrators can use to specify particular settings for users. Only system administrators can make changes to mandatory user profiles. Changes made by the user to desktop settings are lost when the user logs off.

- **Temporary User Profile.** A temporary profile is issued any time that an error condition prevents the users profile from being loaded. Temporary profiles are deleted at the end of each session - changes made by the user to their desktop settings and files are lost when the user logs off.

| **Note** |
| --- |
| If you need to provide managed desktop configurations for groups of users or computers, consider using Group Policy instead of mandatory profiles. |

**Advantages of User Profiles**

A primary goal of user profiles is to separate each users settings and data from that of other users and the local computer. Separating each users state provides several advantages:

- It allows for stateless computers. An organization can configure computers to store all the key user settings and data away from the local computer. This allows for much easier computer replacement and backup. When a computer needs replacing, it can simply be swapped outall of the users state information is safely maintained separately on the network and is independent of a particular computer. When the user logs onto the new computer for the first time, the server copy of the users state is copied to the new computer.

- It allows a users system and desktop customizations to travel with the user from computer to computer, without requiring the user to reconfigure any settings. When a user logs on to any computer on the network that supports the roaming profile, the users desktop appears just as that user left it before logging off. With roaming user support, users can share computers, but each user has his or her personal desktop (both roaming and mandatory profiles support this functionality).

# User Profile Structure

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A user profile consists of:

- A registry hive. The registry is a database used to store computer- and user-specific settings. Portions of the registry can be saved as files, called hives. These hives can then be reloaded for use as necessary. User profiles take advantage of the hive feature to provide roaming profile functionality. The user profile registry hive is the NTuser.dat in file form, and is mapped to the HKEY_CURRENT_USER portion of the registry when the user logs on. The NTuser.dat hive maintains the users environment preferences when the user is logged on. It stores those settings that maintain network connections, Control Panel configurations unique to the user (such as the desktop color and mouse), and application-specific settings. The majority of the settings stored in the registry are opaque to user profiles settings are owned and maintained by individual applications and operating system components.

- A set of profile folders stored in the file system. User profile files are stored in the filesystem in the Documents and Settings directory, in a per user folder. The user profile folder is a container for applications and other operating system components to populate with subfolders and per-user data, such as shortcut links, desktop icons, startup applications, documents, configuration files and so forth. Windows Explorer uses the user profile folders extensively for special folders such as the users desktop, start menu and my documents folder.

Together, these two components record user-configurable settings that can migrate from computer to computer.

The default location of user profiles was changed from the Windows NT 4.0 operating system to allow administrators to secure the operating system folders without adversely affecting user data. On a clean installed computer running Windows Server 2003, Windows XP or Windows 2000, profiles are stored in the %Systemdrive%\Documents and Settings folder. In contrast, on computers running Windows NT 4.0, profiles are stored inside the system directory, at %Systemroot%\profiles folder (typically WINNT\profiles).

| **Note** |
| --- |
| If you upgrade a computer from Windows NT 4.0, the profile location remains %Systemroot%\profiles. |

Table 1 below shows the location of user profiles for each of the possible installation scenarios:

**Table 1 User Profile Locations**

| Operating system | Location of user profile |
| --- | --- |
| Windows Server 2003 clean installation (no previous operating system) | %SYSTEMDRIVE%\Documents and Settings; for example, C:\Documents and Settings |
| Windows Server 2003 upgrade of Windows 2000 | SYSTEMDRIVE%\Documents and Settings; for example, C:\Documents and Settings |
| Windows Server 2003 upgrade of Windows NT 4.0 | %SYSTEMROOT%\Profiles; for example, C:\WinNT\Profiles |

**Configuration Preferences Stored in the Registry Hive**

The NTuser.dat file contains the following configuration settings:

- **Windows Explorer settings.** All user-definable settings for Windows Explorer, as well as persistent network connections.

- **Taskbar settings.**

- **Printer settings.** All network printer connections.

- **Control Panel.** All user-defined settings made in the Control Panel.

- **Accessories.** All user-specific application settings affecting the Windows environment, including: Calculator, Clock, Notepad, Paint, and HyperTerminal, among others.

- **Application Settings.** Many applications store some per user settings in the users registry hive (HKEY_CURRENT_USER). An example of these types of settings would be Microsoft Word 2000s toolbar settings.

**Configuration Preferences Stored in Profile Directories**

Figure 1 below shows the structure of the user profile.

**Figure 1: User Profile.**




Each users profile contains the following folders:

- **Application data***. Application-specific data, such as a custom dictionary for a word processing program. Application vendors decide what data to store in this directory.

- **Cookies**. Internet Explorer cookies.

- **Desktop**. Desktop items, including files and shortcuts.

- **Favorites**. Internet Explorer favorites

- **Local Settings***. Application settings and data that **do not roam** with the profile. Usually either machine specific, or too large to roam effectively.

    - **Application data**. Computer specific application data.

    - **History**. Internet Explorer history.

    - **Temp**. Temporary files.

    - **Temporary Internet Files**. Internet Explorer offline cache.

- **My Documents**. The new default location for any documents that the user creates. Applications should be written to save files here by default.

    - **My Pictures**. Default location for users pictures.

    - **My Music**. Default location for users music.

- **NetHood***. Shortcuts to My Network Places items.

- **PrintHood***. Shortcuts to printer folder items.
- **Recent**. Shortcuts to the most recently used documents.
- **SendTo**. Shortcuts to document storage locations and applications.
- **Start Menu**. Shortcuts to program items.
- **Templates***. Shortcuts to template items.

* These directories are hidden by default. To see these directories, change the View Options.

The Folder Redirection feature of IntelliMirror allows an administrator to redirect the location of certain folders in the user profile to a network location. When these redirected folders are accessed either by the operating system or by applications, the operating system automatically redirects to the location on a network share specified by the administrator. From a user perspective, this is similar to the roaming scenario because users have the same settings regardless of which computers they use. However unlike roaming, these settings actually remain on the network share. Folder redirection can be used with all types of user profiles: local, roaming, or mandatory.

Using Folder Redirection with local profiles can provide some of the benefits of roaming profiles (such as having a users data available at any computer or maintaining data on the server) without the need to implement roaming profiles. Remember though, using Folder Redirection with a local profile would only result in the users documents and files being available from all computers. To have settings and configurations move with the user, you would need to use roaming profiles.

Combining Folder Redirection with roaming profiles gives the benefit of roaming profiles, while minimizing network traffic caused by synchronization of the profile.

Folder redirection is accomplished using Group Policy. The use of Folder Redirection with roaming profiles is discussed later in this article.

Table 2 below lists the folders that roam with the profile by default and indicates whether they can be redirected using Group Policy.

**Table 2 Folders that Roam with the Profile**

| Folder Name | Description | Roams with profile by default | Redirect with Group Policy |
|---|---|---|---|
| Application Data | Per-user roaming application data. | Yes | Yes |
| Cookies | Users Internet Explorer cookies. | Yes | No |
| Desktop | Desktop items, including files and shortcuts. | Yes | Yes |
| Favorites | Users Internet Explorer favorites. | Yes | No |
| Local Settings | Temporary files and per-user non-roaming application data. | No | No |
| My Documents | Users documents. | Yes | Yes |
| NetHood | Shortcuts to My Network Places items. | Yes | No |
| PrintHood | Shortcuts to printer folder items. | Yes | No |
| Recent | Shortcuts to recently used documents | Yes | No |
| Send To | Shortcuts to document storage locations and applications. | Yes | No |

| | | | |
|---|---|---|---|
| Start Menu | Users personal start menu. | Yes | Yes |
| Templates | Per-user customized templates. | Yes | No |

**Non-Roaming Folders**
The default behavior of roaming user profiles in Windows NT 4.0 is to include all the folders in the user profile directory. Thus when a user first logs on, all folders within the profile folder are copied from the server to the client at logon and copied back at logoff,

Windows 2000 introduced a per-user local settings folder into the user profile that is not copied during log on or logoff. This folder is intended for the storage of operating system components and other applications can store non-roaming per-user data. A typical example of the usage of this folder is for Microsoft Internet Explorer to store a users Favorites in the roaming portion of the user profile but store the Temporary Internet Files in the local (non-roaming) portion of the user profile. This will allow a user to retain access to their favorite URLs, but will save copying of temporary cache files at logon and logoff.

On computers running Windows Server 2003, Windows XP or Windows 2000, the History, Local Settings, Temp and Temporary Internet Files folders do not roam by default. Other Non-Roaming Folders are configured using the Group Policy Object Editor. The path for this setting in the Group Policy name space is:

User Configuration\Administrative Templates\System\User Profiles\Exclude directories in roaming profile

Once enabled this allows multiple folder names to be defined, all relative to the root of the users profile. Once included in the policy these folders will not be copied to the local machine at logon, nor copied back to the server at logoff. This setting is likely to result in decreased time taken for a user to logon, by restricting the amount of data within a user profile that really does roam with the user.

**How Do Users Get Their Profile?**
The way in which users get their profiles depends on the type of profile they're configured to use. This section describes this process.

**Local Profile - New User**

1. The user logs on.

2. The operating system checks the list of user profiles located in HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\ProfileList to determine if a local profile exists for the user.

3. Because this is a new user, no local profile is found. If the computer is part of a domain, the operating system checks if a domain wide default profile exists in a folder named **Default User** on the domain controllers **NETLOGON** share.

    - If a domain wide profile exists, it is copied to a subfolder on the local computer with the user name under %SYSTEMDRIVE%\Documents and Settings\. For example, a new user with the user name JDoe would have a profile created in %SYSTEMDRIVE%\Documents and Settings\JDoe.

    - If a default domain profile does not exist, then the local default profile is copied from the %Systemdrive%\Documents and Settings\Default User folder to a subfolder on the local computer with a user name under %Systemdrive%\Documents and Settings\.

4. If the computer is not part of a domain, the local default profile is copied from the %Systemdrive%\Documents and Settings\Default User folder to a subfolder on the local computer with a user name under %Systemdrive%\Documents and Settings\.

5. The users registry hive (NTUSER.DAT) is mapped to the HKEY_CURRENT_USER portion of the registry.

6. The users %userprofile% environment variable is updated with the value of the local profile folder

7. When the user logs off, a profile is saved to the local hard disk of the computer.

**Local Profile - Existing User**

1. The user logs on.

2. Windows checks the list of user profiles located in HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\ProfileList to get the path to the users profile.

3. The users registry hive (NTUSER.DAT) is mapped to the HKEY_CURRENT_USER portion of the registry.

4. The users %userprofile% environment variable is updated with the value of the local profile folder.

5. When the user logs off, the profile is saved to the local hard disk of the computer.

**Roaming Profile - New User**

1. The user logs on.

2. The path to the users roaming profile is retrieved from the user object on the Domain Controller.

3. Windows checks to see if a profile exists in the roaming path, if no profile exists a folder is created.

4. Windows checks the list of user profiles located in HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\ProfileList to determine if a cached copy of the profile exists. If a local copy of the profile is not found, and the computer is part of a domain, Windows checks to determine if a domain wide default profile exists in the Default User folder on the domain controllers NETLOGON share.

    - If a domain wide profile exists, it is copied to a subfolder on the local computer with their user name under %Systemdrive%\Documents and Settings\.

    - If a default domain profile does not exist, then the local default profile is copied from the %Systemdrive%\Documents and Settings\Default User folder to a subfolder on the local computer with their user name under %Systemdrive%\Documents and Settings\.

5. The users registry hive (NTUSER.DAT) is mapped to the HKEY_CURRENT_USER portion of the registry.

6. The users %userprofile% environment variable is updated with the value of the local profile folder

7. The user can then run applications and edit documents as normal. When the user logs off, their local profile is copied to the path configured by the administrator. If a profile already exists on the server, the local profile is merged with the server copy (see merge algorithm later in this paper for more details).

**Roaming Profile - Existing User**

1. The user logs on.

2. The path to the users roaming profile is retrieved from the user object on the Domain Controller.

3. Windows checks to see if a profile exists in the roaming path, if no profile exists a folder is created.

4. Windows checks the list of user profiles located in HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\ProfileList to get the path to the users cached profile if it exists.

5. The contents of the local cached profile are compared with the copy of the profile on the server, and the two profiles are merged. (See the new merge algorithm later in this paper for more details).

6. The users registry hive (NTUSER.DAT) is mapped to the HKEY_CURRENT_USER portion of the registry.

7. The users %userprofile% environment variable is updated with the value of the local profile folder

8. The user can then run applications and edit documents as normal. When the user logs off, the local profile is copied to the path configured by the administrator. If a profile already exists on the server, the local profile is merged with the server copy.

# Enhancements to User Profiles in Windows Server 2003 and Windows XP

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

User data management and user settings management features provide several enhancements that increase the usability, resilience, and performance of user profiles:

- New Group Policy settings
- Support for Windows XP fast network logon
- New System Service profiles
- More robust roaming

**New Group Policy Settings**

User profiles policies have moved to their own node in the Group Policy Object Editor: Administrative Templates\ System\User Profiles. In addition there are three new policies under computer policies:

- **Prevent Roaming Profile Changes from Propagating to the server.** This policy determines if the changes a user makes to their roaming profile are merged with the server copy of their profile. If this policy is set at login, the user will receive their Roaming Profile, but any changes a user makes to their profile will not be merged to their roaming profile at logoff.
- **Add the Administrator security group to the roaming user profile share.** In Windows 2000, the default file permissions for newly generated roaming profiles are full control, or read and write access for the user, and no file access for the administrators group. This policy allows an administrator to choose the same behavior as Windows NT 4.0, and have the administrators group given full control of the user's profile directories.
- **Only allow local user profiles.** This setting determines if roaming user profiles are available on a particular computer. By default, when roaming profile users log on to a computer, their roaming profile is copied down to the local computer. Using this setting, an administrator can prevent users configured to use roaming profiles from receiving their roaming profile on a specific computer.

**Support for fast network logon**

**Fast network logon overview**

Windows XP includes new features that provide faster start up of a computer by not waiting for the network during boot and logon.

By default, Windows XP (but not Windows Server 2003) does not wait for the network to be fully initialized at startup and logon. Any existing users logging on are logged on using cached credentials, which results in shorter logon times. Because the computer doesn't wait for the network to be fully started, Group Policy is applied in the background once the network becomes available. This has a number of effects on the logon process:

- **Changes to some Group Policy extensions can take up to three logons to become effective.** Because background refresh becomes the default behavior, some policy extensions such as Software Installation and Folder Redirection may require as many as three logons to apply changes. This is because to be able to operate safely, these extensions require that no users be logged on, so they must be processed in the foreground before users are actively using the computer. In the case of Advanced folder redirection, because policy is evaluated based on security group membership three logons will be required: the first logon to update the cached user object (and security group membership), the second logon for policy to detect the change in security group membership and require a foreground policy application, and the third logon to actually apply folder redirection policy in the foreground.
- **Changes to some user object properties may take two logons to become effective.** Because users are logged on

using cached credentials, changes that are made to the user object, such as adding a roaming profile path, home directory, or user object logon script, may take up to two logons to be detected. The user object is updated in the background after the user has logged on, so any changes to its properties will not take effect until the next logon at least. Certain key fields in the user object are also checked at logoff, so if the user is logged on at the time of the change, the change may be detected at logoff.

**Changes in roaming user profiles to support fast network logon**
When a user logs onto a computer they have not logged onto for some time, the timestamps of the local user registry hives are compared to the server registry hive and the newest wins. Typically this is the server copy of the registry, as the timestamps are checked before the hive is loaded.

However, when Windows XP is operating in fast network logon mode, the user is always logged on with a cached profile. The effect of this is that when the system detects that the user is now a roaming user, the local registry hive has already been loaded and therefore the hive timestamp is always changed. This introduces the possibility that if a user logs onto multiple computers, an older profile can overwrite a newer server profile because the users roaming status is one step behind (due to the cached logon). To avoid this situation, the roaming profile algorithm treats the transition from local to roaming on a given computer as a special case:

At logon, checks are made to see which of the following conditions are true:

- Is this is the first roaming logon for a user that previously had a local profile?
- Is there a server copy of the users profile?

If all of these conditions are met, then the algorithm is subtly changed, the new algorithm does the following:

1. The contents of the local profile are merged with the server copy of the profile as normal, with the exception of the user registry hive (ntuser.dat).
2. The server copy of the user registry is always copied from the server profile to the local profile, regardless of the time stamps on the hives.

In all other cases, the algorithm remains unchanged. Its only the first time that a user becomes roaming on a specific computer that this check is made. Once the user has become a roaming user, the computer always waits for the network (so that the profile can be downloaded) and the behavior is exactly the same as Windows 2000.

With this change, it is recommended that if an administrator removes the profile path from a users user object, that they either rename or delete the corresponding profile folder. If the profile folder is not removed, and the administrator re-adds the same profile path, the user will receive the older server copy of their registry.

**New System Service profiles**
Windows XP and Windows Server 2003 introduce three new user profiles for the System:

- LocalService
- NetworkService
- System

**LocalService and NetworkService profiles**
The LocalService and NetworkService profiles are automatically created for two new built in user accounts that are used by the Service Control Manager to host services that do not need to run as the local system account. While these two new profiles are normal user profiles, because they are required by the system to run, they are treated slightly differently:

- LocalService and NetworkService profiles are not shown in the profiles list on the system properties dialog.
- Both LocalService and NetworkService profiles are located in %systemroot%\Documents and Settings by default, but are marked as super hidden so that they are not ordinarily visible.

**System Profile**
In Windows 2000, when an application or service used the LoadUserProfile API to load a user profile for a process running as the local system, Windows created a profile named %computername%$ , where %computername% is the name of the local computer. This could cause problems for some applications and services, because depending on whether the system profile was loaded, HKEY_CURRENT_USER could in fact resolve to different registries either HKEY_USERS \S-1-5-18 or HKEY_USERS\.DEFAULT depending on whether another component has loaded the SYSTEM profile.

To avoid this, Windows Server 2003 and Windows XP create a new profile for the system, located in %systemroot%\System32\Config\SystemProfile. This profile is always loaded, and is a link to HKEY_USERS\.DEFAULT. This ensures that system components always have a consistent profile and registry.

**More robust roaming**
In Windows 2000, an issue exists whereby poorly written applications and services that keep registry keys open during logoff prevent Windows from unloading the user's registry hive. When this occurs, changes that a user has made to their profile are not saved to the server. This has three symptoms:

- The user experience is impacted, as users may wonder why changes have not been saved when they log onto another computer.
- Since "locked" profiles never get unloaded, they end up using a lot of memory on a terminal server that has many users logging in.
- If a profile is marked for deletion at logoff (to cleanup the machine or for temporary profiles), profiles do not get deleted.

The three symptoms are solved as follows:

- In Windows 2000 when a user logs off, if the profile is "locked", Windows polls the profile for 60-seconds before giving up. Windows Server 2003 and Windows XP now save the users registry hive at the end of the 60-second delay, and roam the profile correctly.
- When the application or service closes the registry key "unlocking" the profile, Windows Server 2003 or Windows XP will unload the users registry, thus freeing the memory consumed by the profile.
- If a profile is marked for deletion at logoff, when the reference count drops to zero, it will be unloaded and deleted. In the event that the application never releases the registry key, Windows Server 2003 or Windows XP will delete all profiles marked for deletion at the next machine boot.

Note that if a roaming user logs onto back onto a machine that failed to unload her profile previously, and the profile has still not unloaded, the user will receive the currently loaded profile. If the user has logged onto another computer in the meantime, and made changes to their profile, the changes will not be available on the machine with the locked profile.

**Improved Merge Algorithm**
This section describes how Windows Server 2003 or Windows XP reconciles local and server copies of a users profile. To improve the experience of users, roaming profiles have a new algorithm to synchronize copies of a profile. This prevents problems from occurring when a user logs into two different computers simultaneously. The Windows NT 4.0 algorithm worked well in the most common cases where users logged on to only a single computer. However, when a user logged onto multiple computers at the same time, the user sometimes experienced unexpected behavior caused by the assumption that each computer had the master copy of the profile.

For Windows Server 2003, Windows XP, and Windows 2000, the algorithm was changed to support the merging of user profiles at the file level and to provide support for last writer wins.

To illustrate the behavior of the new algorithm, several examples are presented that compare the behavior of Windows NT 4.0 to the current model.

**Overview of the Windows NT 4.0 Algorithm**
In Windows NT 4.0, the algorithm is an Xcopy with full synchronization support, meaning it has the ability to mirror a profile from one location to another, and any extra files or directories in the destination location are removed. The algorithm is based on the concept that there is only one master profile at any one time. When the user is logged on, the master profile is on the local computer. When the user is not logged on, the master profile is on the server.

1. The user logs on to computer A, the primary computer.
2. The roaming profile is Xcopied from the server location to the local profile location.
3. The user creates documents, changes colors, and so on. All of these changes are stored in the local profile location.
4. As the user logs off the computer, the profile is Xcopied from the local profile location back to the server location.

This is an exact mirroring process. If there are any extra files in the server location, they are deleted to ensure that the server location is a duplicate of the local profile. As mentioned previously, this works well in the majority of cases, where a user logs on to only a single computer; but a user who logs on to multiple computers at the same time might experience unexpected behavior.

**Examples of Windows NT 4.0 Merge Algorithm Issues**
When using Windows NT 4.0, a problem arises if the user logs on at two or more computers. Building on the preceding example:

1. The user logs on to computer A.
2. The user logs on to computer B.
3. The user creates a document on computer A and stores it in the user profile.
4. The user logs off of computer A.
5. The user logs off of computer B.

The document that the user created in step 3 is deleted because, from the perspective of computer B, the master profile is stored locally. The extra files on the server must be deleted so that the local profile is currently the master server profile.

The Windows XP algorithm preserves the document because it is able to compare the time the document was created with the time the profile was loaded. If the document was created or modified after the profile load time, the file must be preserved because it came from a different source.

A similar problem can occur when files are modified. For example, suppose that the user has a document called Document.doc in his or her My Documents folder in the server copy of the profile:

1. The user logs on to computer A.
2. The user logs on to computer B.
3. The user modifies the document on computer A.
4. The user logs off computer A.
5. The user logs off computer B.

The changes made to the document on computer A are lost because when the user logged off computer B, the computer overwrote the new version of the document with the old one; the computer is programmed to recognize that it had the master version of the profile.

The current algorithm preserves the changes to the document because it compares the time the document was modified with the time the profile was loaded. This results in a much better experience for the user.

**Overview of Windows Server 2003 Merge Algorithm**
Windows Server 2003 and Windows XP merge user profiles at the file level. The merged profile contains the superset of files that are in the local computer and server copies of the users profile. In the case where the same file is in both the local and server copy of the profile, the file that was modified most recently is used. This means that new files are not deleted, and updated versions of existing files are not overwritten.

When a document or file is updated, the new algorithm compares the timestamp of the destination file with the timestamp of the source file. If the destination file is newer, it is not overwritten.

When a user logs on to a computer, the current time is saved; when the user logs off, this timestamp is used to determine which files are new in the server profile and which files have been deleted in the local profile. For example, if the server profile has a document in the **My Documents** folder called Review.doc and this file does not exist in the local profile, either it is a new file from a different computer, or it was originally in the local profile and the user deleted it. Knowing the time when this new profile was loaded, it is possible to compare it against the timestamp of Review.doc. If Review.doc was created or written to after the profile load time, the file must be preserved because it came from a different source. If the Review.doc timestamp is older than the profile load time, Review.doc must be deleted because it would have been copied to the local computer at load time.

In addition, some files might need to be removed from the local cache so that items that were deleted between sessions remain deleted. For example:

1. The user logs on to computer A.
2. The user creates or modifies a document on computer A.
3. The user logs on to computer B.
4. The user logs off computer B; computer B has a copy of the document.
5. The user deletes the document and logs off computer A.

To ensure that the files are deleted, the cached version of the profile is synchronized with the profile on the server when a user logs on. All files in the local cache that are not present in the server and that were not modified since the last logoff time are removed. By using these changes, Windows XP can merge user profiles.

There are several files that will **always** be copied because the timestamp on those files is changed as part of the profile unloading process. The ntuser.dat and ntuser.ini files will always be copied to the server.

**Quotas on Profile Size**

Often, over time a roaming user profile may potentially grow to a large size, as the user stores more data on their desktop, installs more applications and increases the complexity of their roaming environment. While much of this enhances the user experience, it may result in greater logon and logoff times for the user, as the data is copied down from or back up to the server.

While excluding some of the profile folders from roaming can reduce profile size, there may be situations in which it is more useful to actively restrict the overall quota size, to prevent it growing to an excessive size. The Proquota.exe program is a tool that you can set to monitor the size of a user's profile. If an individuals user profile exceeds the predetermined file limit, the user cannot log off from the computer until the user reduces the size of files.

Profile quota size is managed using the Group Policy snap-in. You can use the **Limit Profile Size** policy, available in the **User Configuration\Administrative Templates\System\User Profiles** node of the Group Policy snap-in to set the maximum size of the roaming user profile and to determine the systems response when the limit is reached. Click the Explain tab of this policy setting for more details.

**Deleting cached roaming profiles**

If you are concerned with disk size on a multi-user computer for example, a public computer where thousands of users can log on, you can also use the Group Policy setting that removes cached versions of the profile on logoff. The policy is called **Delete cached copies of roaming profiles**, and it is accessed under the **Computer Configuration\Administrative Templates\System\User Profiles** node of the Group Policy snap-in.

**Changes to the way Guest profiles are handled**

Windows 2000 and Windows NT 4.0 always delete the user profile of users belonging to the local Guests security group when users log off. Windows XP and Windows Server 2003, continue to delete the profile of guest users, only when the computer is joined to a domain. When the computer is part of a workgroup, the user profile of users belonging to the local Guests group is not deleted at logoff.

The exception is when the user is a member of the local Guests group AND a member of local Administrators, in this case the profile is NOT deleted when in a domain.

**Group Policy Settings for Roaming User Profiles**

See the Appendix for a list of policy settings related to roaming user profiles. For details about these policy settings, click the policys **Explain** tab.

# How to Configure a Roaming User Profile

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can configure a roaming profile by using the following procedure.

To configure a roaming profile for users:

1. Create a folder on the server where user profiles will be stored. This will be the top-level folder that contains all the individual user profiles.

2. Configure the folder as a shared folder, and give all users **Full Control** permissions.

3. Open the **Active Directory Users and Computers** snap-in and navigate to the individuals **User** object.

4. Right-click the **user's name** and click **Properties** on the shortcut menu.

5. Click the **Profile** tab.

6. For the **Profile Path**, type the path to the network share where the user profile is to be stored. For example, for a user whose network name is JDoe, the following path, \\NetworkShare\Profiles\%username%, would create a directory called JDoe in the Profiles share on the server used to store user profiles.

You can also populate the profile path by using an Active Directory Scripting Interface (ADSI) script. ADSI provides a single set of interfaces for managing network resources. Administrators can combine ADSI with Visual Basic Scripting Edition (VbScript) or Jscript scripts to manage resources in the directory service such as users, services, and so on.

To learn about ADSI and ADSI scripts, see the Microsoft Platform SDK.

**User Profiles Troubleshooting**

The first troubleshooting step should be to examine the Application event log on the client computer, and determine the error.

If this is a roaming profile, be sure to check for the correct permissions (see Security Considerations when Configuring Roaming User Profiles for the required permissions) - one of the most common causes of roaming user profile errors is incorrect permissions on the profile share.

In addition to logging events in the Application Event log, User Profiles can provide a detailed log to aid troubleshooting. To create a detailed log file for user profiles:

1. Start regedit and locate the following path:
   HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Winlogon

2. Create a new value called UserEnvDebugLevel as a REG_DWORD, and set the value to 30002 in hexadecimal format.

The log file can be found at: %windir%\debug\usermode\userenv.log

# Security Considerations when Configuring Roaming User Profiles

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**When creating the roaming profile share, limit access to the share to only users that need access.**

Because a users roaming profile contains personal information, such as documents and EFS certificates care should be taken to protect it as well as possible. In general:

- Restrict the share to only users that need access. Create a security group for users that have profiles on a particular share, and limit access to only those users.
- When creating the share, hide the share by putting a $ after the share name. This will hide the share from casual browsers; the share will not be visible in My Network Places.
- Unless you need special permissions on the profile folder, don't pre-create profile folders for the user, allow the system to create them.
- Only give users the minimum amount of permissions needed. The permissions needed are shown in the tables below:

**Table 3 NTFS Permissions for Roaming Profile Parent Folder**

| User Account | Minimum permissions required |
|---|---|
| Creator/Owner | Full Control, Subfolders And Files Only |
| Administrator | None |
| Security group of users needing to put data on share. | List Folder/Read Data, Create Folders/Append Data - This Folder Only |
| Everyone | No Permissions |
| Local System | Full Control, This Folder, Subfolders And Files |

**Table 4 Share level (SMB) Permissions for Roaming Profile Share**

| User Account | Default Permissions | Minimum permissions required |
|---|---|---|
| Everyone | Full Control | No Permissions |
| Security group of users needing to put data on share. | N/A | Full Control, |

**Table 5 NTFS Permissions for Each Users Roaming Profile Folder**

| User Account | Default Permissions | Minimum permissions required |
|---|---|---|
| %Username% | Full Control, Owner Of Folder | Full Control, Owner Of Folder |
| Local System | Full Control | Full Control |
| Administrators | No Permissions* | No Permissions |

| | | |
|---|---|---|
| Everyone | No Permissions | No Permissions |

**Unless the Add the Administrator security group to the roaming user profile share policy is set, in which case the Administrators group has Full Control. (requires Windows 2000 Service pack 2 or later)*

**Use at least Windows 2000 servers to host profile shares.**

Because a users roaming profile contains personal information which is copied to and from the client computer, and the server hosting the roaming profile, it is important to ensure that data is protected as it travels over the network.

The biggest potential threats to the privacy and integrity of a users data come from intercepting the data as it passes over the network, tampering with the data as it passes over the network, and spoofing the server hosting the users data.

Several features of Windows 2000 and Windows Server 2003 can help to secure a users data:

- **Kerberos** - Kerberos is standard on all versions of Windows 2000 and Windows Server 2003s, and ensures the highest level of security to network resources. While NTLM authenticates the client only, Kerberos authenticates the server and the client. When NTLM is used, the client doesn't know whether the server is valid this is particularly important if the client is exchanging personal files with the server, as is the case with Roaming Profiles. Kerberos provides better security than NTLM and is not available on Windows NT version 4.0 or earlier operating systems.

- **IPSec**- The IP Security Protocol (IPSec) provides network-level authentication, data integrity, and encryption ensuring that roamed data is:

    - Safe from data modification while enroute.

    - Safe from interception, viewing, or copying.

    - Safe from being accessed by unauthenticated parties.

- **SMB Signing**- The Server Message Block (SMB) authentication protocol supports message authentication, which prevents active message and "man-in-the-middle" attacks. SMB signing provides this authentication by placing a digital signature into each SMB, which is then verified by both the client and the server. In order to use SMB signing, you must either enable it or require it on both the SMB client and the SMB server. Note: SMB signing imposes a performance penalty; although it doesn't consume any more network bandwidth, it does use more CPU cycles on the client and server side.

**Always use the NTFS Filesystem for volumes holding users data.**

For the most secure configuration, configure servers hosting roaming profiles to use the NTFS File System. Unlike FAT, NTFS supports Discretionary access control lists (DACLs) and system access control lists (SACLs), which control who can perform operations on a file and what events will trigger logging of actions performed on a file.

# Best Practices for User Profiles

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To get the best experience possible from roaming user profiles, it is important to read all the documentation and plan your implementation thoroughly. This section presents best practices for using roaming user profiles.

**Turn off the fast logon enhancement**

With the fast logon enhancement in Windows XP when users change from a local to a roaming profile, it will take two logons on each machine for profile changes to be registered. This is because the user always logs on with cached credentials; therefore it takes one logon for the network to notice that the user has become roaming and the second logon to apply these settings.

To ensure the best possible experience, enable the setting **Always wait for the network at computer startup and logon**, located at Computer Configuration\Administrative Templates\System\Logon.

**Redirect the location of the My Documents Folder outside of the users Roaming Profile.**

To decrease initial logon time to a new computer, it is recommended that you redirect the location of the My Documents folder outside of the users roaming profile. The best way to do this is with Folder Redirection. If you don't have Active Directory enabled, you can do this with a logon script or instruct the user to do so manually.

**Let the system create profile folders for each user.**

To ensure that Roaming user profiles work optimally, create only the root profile share on the server, and let the system create the folders for each user. If you must create folders for the users, ensure that you have the correct permissions set. For details on the required permissions see **Security Considerations when Configuring Roaming User Profiles**.

**Don't use Offline Folders on Roaming Profile Shares.**

Make sure that you turn off Offline Folders for shares where roaming user profiles are stored. If you do not turn off Offline Folders for a users profile, you may experience synchronization problems as both Offline Folders and Roaming Profiles try to synchronize the files in a users profile.

| **Note** |
| --- |
| This does not affect using Offline Folders with redirected folders such as My Documents. |

**Do not use Encrypted File System (EFS) on files in a Roaming User Profile.**

The Encrypted File System is not compatible with files within Roaming User Profiles. If you encrypt profile folders or files using EFS the users profile will not roam.

| **Note** |
| --- |
| This does not affect encrypting files on remote shares. |

**Do not Set Disk Quotas too low for users with Roaming Profiles.**

If a users disk quotas are set too low, roaming profile synchronization may fail. Make sure enough disk space is allocated to allow the system to create a temporary duplicate copy of a users profile. Because the temporary profile is created in the users context as part of the synchronization process, it debits his or her quota.

**Use Group Policy loopback policy processing sparingly if you use roaming profiles.**

Group Policy loopback processing enables a different set of user type Group Policies to be applied based on the computer being logged onto. This policy is useful when you need to have user type policies applied to users of specific computers. There are two methods for doing this. One allows for the policies applied to the user to be processed, but to also apply user policies based on the computer that the user has logged onto. The second method does not apply the users settings based on where the user object is, but only processes the policies based on the computers list of GPOs.

Use caution when using loopback policy processing and roaming profilesespecially when users may roam between Windows 2000 or Windows XP-based computers and Windows NT 4.0-based computers. You may see some tattooing applications can store policy settings in HKCU\Software\Policies regardless of operating system version. Windows NT 4 also stored some explorer policy settings in HKCU\Software\Microsoft\windows\currentversion\explorer\policies. Windows 2000

and Windows XP clears these keys each time before re-applying current policy, but because Windows NT 4 does not clear them, you will get settings left if you roam from a Windows 2000-based machine.

**Roaming between the same operating system versions**

Because the contents of the users registry are opaque to roaming user profiles, you should minimize as much as possible significant differences between operating system installations on computers the users will roam between. When using roaming profiles, try to ensure that:

- The same application versions are installed.
- Applications are installed to the same path and drive.
- The operating system is installed on the same %systemdrive% and in the same %windir%.
- The operating system language and system locale are the same.

**Roaming between different operating system versions**

Although roaming between Windows 2000 and Windows XP should be a smooth process, there are some precautions you can take to minimize possible issues:

- If you can avoid roaming between versions of the operating system, then do so. There's nothing inherent in roaming that will cause problems, but the data that applications put in the profile may have unintended side effects on other versions of the operating system.
- Make sure that you have the same application versions installed.
- Make sure that applications are installed to the same path and drive.
- Make sure that the different versions of the operating system are installed on the same %systemdrive% and in the same %windir%.
- If Users roam between Windows NT 4.0-based clients and Windows XP- Windows 2000-based clients, consider setting the Profile Path during install on Windows XP or Windows 2000. Differences in the default profile path (%windir%\Profiles vs. %systemdrive%\Documents and Settings) may cause problems for users roaming between Windows NT 4.0-based clients and Windows XP- or Windows 2000-based clients. To minimize the chance of problems, make sure the path to the profile is the same on both clients.

# Folder Redirection Overview

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Folder redirection is a feature of IntelliMirror that allows users and administrators to redirect the path of a folder to a new location. The new location can be a folder on the local computer or a directory on a network share. Users have the ability to work with documents on a server as if the documents were based on the local drive. For example, you can redirect the My Documents folder, which is usually stored on the computer's local hard disk, to a network location. The documents in the folder are available to the user from any computer on the network. The My Documents folder is the location on the Windows Server 2003, Windows XP or Windows 2000 desktop where the user can save documents and graphic files.

Previously, administrators who wanted to redirect folders to the network had to use logon scripts to change registry values. In Windows Server 2003 and Windows XP, the same task can be accomplished by using Group Policy.

**Advantages of Using Folder Redirection**

Folder redirection provides a number of advantages. Some of the following benefits relate to redirecting any folder, but redirecting My Documents can be particularly advantageous.

- Even if a user logs on to various computers on the network, the users documents are always available.
- The system administrator can use Group Policy to set disk quotas, limiting the amount of space taken up by users' special folders.
- Data specific to a user can be redirected to a different hard disk on the user's local computer from the hard disk holding the operating system files. This protects the user's data if the operating system needs to be reinstalled.
- Data stored on a shared network server can be backed up as part of routine system administration. This is safer and it requires no action on the part of the user.

You can also combine Folder Redirection and roaming user profiles to decrease logon and logoff times for roaming and mobile users. Besides the improved availability and backup benefits of having the data on the network, users also have performance gains with low-speed network connections and subsequent logon sessions. Because only some of their documents are copied, performance is improved when the users profiles are copied from the server. Not all of the data in the user profile is transferred to the desktop each time the user logs on only the data that user requires.

When you combine the use of Folder Redirection and roaming user profiles, you can provide fast computer replacement. If a user's computer needs to be replaced, the data that a user requires can quickly be re-established on a replacement computer. By using Folder Redirection to redirect the My Documents and Application Data folders, in conjunction with roaming user profiles and Group Policy-based deployment of applications, an organization can move the key user state to a network location. This means the users documents, settings, and applications follow them, regardless of which Windows XP computer the user logs on to.

For increased availability, Offline File technology gives users access to My Documents even when they are not connected to the network. For more information, see the Complementary Technologies section later in this article. This is particularly useful for those who use laptop computers.

**Folders that Can Be Redirected**

My Documents, Application Data, Desktop, and Start Menu are the top level folders that can be redirected. These were identified as the key folders that an organization would need to redirect to preserve important user data and settings. There are several advantages to redirecting each of these folders. The usefulness of each will vary according to your organizations needs.

- **My Documents.** The place in the shell for users to save their documents and pictures. Because common dialog boxes in Windows Server 2003 and Windows XP point to the My Documents folder by default, there is a greater tendency for users to save files there. Data stored on a shared network server can be backed up as part of routine system administration, and is safer because it requires no action from the user.
- **Application Data.** Often applications place large data, such as dictionaries, in the Application Data portion of the user's profile, which roams with the user. To improve performance, Application Data was added to the list of folders that can be redirected. This means that users can still have access to Application Data (such as the custom dictionary), but without the

need to download the (possibly large) files at every logon.

- **Desktop.** Some organizations want to configure computers to look the same. By redirecting the desktop for a group of users, you can ensure that all users share the same desktop, with the same desktop items.

- **Start Menu.** For compatibility with Windows NT 4.0, Windows Server 2003 and Windows XP allow you to use Folder Redirection to redirect the **Start** menu folder. **Start** menu redirection is treated differently from other redirected folders in that the contents of the users **Start** menu are not copied to the redirected location. It is assumed that a redirected **Start** menu has been pre-created by an administrator and that all users **Start** menus will be the same. As a best practice for computers running Windows Server 2003 or Windows XP, do not use Folder Redirection to redirect the **Start** menu folder, use Group Policy to control what appears on the **Start** menu.

**Folder Redirection Improvements for Windows XP and Windows Server 2003**

This section provides information about the differences between Windows 2000 and Windows Server 2003.

**User Interface changes**

The Folder Redirection user interface has been simplified for Windows Server 2003. The main goals of these changes were to simplify the use of Folder Redirection by removing the requirement that administrators be familiar with environment variables such as %USERNAME%.

In addition to the simplified UI, several new redirection options have been added:

- **Create a folder for each user under the root path**

  Rather than having to enter a UNC path such as \\server\share\%username%\MyDocuments, the administrator can simply type in the path to the share such as \\server\share, and Folder Redirection will automatically append the user name and the folder name when the policy is applied. This removes the need for administrators to be familiar with environment variables, and minimizes the chances of errors and spelling mistakes.

- **Redirect to home directory (My Documents Only)**

  Windows Server 2003 and Windows XP allow you to redirect a users My Documents folder to their home directory. This option is intended only for organizations that have a legacy deployment of home directories and wish to transition users to the My Documents metaphor while maintaining compatibility with their existing home directory environment. **You should only select this option if you have already deployed home directories in your organization.**

  Folder redirection treats redirection to the home directory as a special case and certain checks are skipped:

  - Redirection to the home folder is only supported on Windows XP and Windows Server 2003 computers. **Redirection to the home directory policy will fail to apply on Windows 2000 computers.**

  - Security is not checked it is assumed that the administrator has set directory security correctly, no ACL check is made and no reACLing is done. The **Grant the user exclusive rights to My Documents** checkbox on the settings page of the property sheet is disabled.

  - No ownership checks are made - normally folder redirection will fail if a user is not the owner of the folder they are being redirected to. Because redirection to the home directory is intended for use in a legacy environment, this ownership check is skipped.

  - Users must have the home folder property correctly set on their user object - The folder redirection client side extension retrieves the actual path for the user's home directory from the user object at logon time. Users affected by Folder Redirection Policy must have this path correctly set or folder redirection will not apply.

- **Redirect to a specific path**

  This option is intended to allow an administrator to redirect folders to an alternate local drive/partition, or to enter unusual configurations not anticipated by the new user interface. Functionally it works in exactly the same way as the Windows 2000 folder redirection user interface.

- **Redirect to the local user profile**

  This option is intended to allow an administrator to redirect the selected folder to the default location in the local user profile, for example: %userprofile%\<Folder Name>. This setting can be used to remove redirection for a particular folder, without removing the GPO. Note: Setting the redirection option to Not Configured (or No Administrative policy specified in the Windows 2000 UI) **does not** redirect the folder to the local profile, this option means that Folder Redirection is not configured *if a folder was previously redirected it will continue to be redirected to the previous location*. If an administrator wished to return the folder to the local user profile they can use this redirection setting.

**My Pictures no longer shown in the Folder Redirection Node**
To simplify the user interface and to help support the best practice that the My Pictures folder should always follow the My Documents folder, the My Pictures folder is not shown in the Folder Redirection node for new GPOs. If you have previously redirected the My Pictures folder separately, the My Pictures node will still appear.

**Redirected Folders automatically made available offline**
By default in Windows XP and Windows Server 2003, any redirected shell folders such as My Documents, Desktop, Start Menu, and Application Data are automatically made available offline. This is in contrast to Windows 2000, which required administrators to configure the Administratively assigned offline files policy setting to ensure all files in the redirected folders were always available offline. This setting was difficult to use with advanced folder redirection, and involved extra administrative overhead.

The default behavior can be overridden by enabling the Do not automatically make redirected folders available offline policy. This setting can be found in the Group Policy Object Editor in the User Configuration\Administrative Templates\Network\Offline Files section.

Note that on Windows Server 2003 Offline files are disabled by default.

# How to Configure Folder Redirection

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Administrators manage Folder Redirection settings by using the Group Policy snap-in.

**To configure Folder Redirection:**

1. To start the Group Policy snap-in from the Active Directory Users and Computers snap-in, click Start, point to Programs, click Administrative Tools, and then click Active Directory Users and Computers.

2. In the MMC console tree, right-click the domain or the OU for which to access Group Policy, click Properties, and click Group Policy.

3. To create a new Group Policy object (GPO), right-click the domain or OU you want to associate with the GPO, select Properties from the context menu, and then in the domain or OU containers Properties page, click the Group Policy tab.

4. Click New, and type the name to use for the GPO. For example, type Redirect MyDocuments GPO.

5. Click Edit to open the Group Policy snap-in and edit the new GPO.

6. In the Group Policy console, expand the User Configuration, Windows Settings, and Folder Redirection nodes. Icons for the personal folders that can be redirected will be displayed.

7. To redirect any of these folders, right-click the folder name, click Properties, and then select one of the following options from the Setting drop-down box:

   - Basic - Redirect everyone's folder to the same location. All folders affected by this Group Policy object will be stored on the same network share.

   - Advanced Specify locations for various user groups. Folders are redirected to different network shares based on security group membership. For example, folders belonging to users in the Accounting group can be redirected to the Finance server, while folders belonging to users in the Sales group are redirected to the Marketing server.

8. On the My Documents Properties page, in the **Target folder location** drop down box select **Create a folder for each user under the root path**. In the Root Path text box, type the name of the shared network folder to use, or click **Browse** to locate it. Note: Unlike Windows 2000, you do not need to type in the %username% variable. The folder redirection code will automatically create a My Documents folder for each user, inside a folder based on their user name. For example, type **\\FolderServer\MyDocumentsFolders** rather than **\\FolderServer\MyDocumentsFolders\%username%** as you would on Windows 2000.

9. In the folders **Properties** dialog box, select the **Settings** tab, configure the options you want to use, and then click **Finish** to complete the Folder Redirection. The available options for settings are:

   - **Grant the user exclusive rights to My Documents.** If selected, this sets the NTFS security descriptor for the %username% folder to Full Control for the user and local system only; this means that administrators and other users do not have access rights to the folder. This option is enabled by default. Note: Changing this option after the policy has been applied to some users will only effect new users receiving the policy.

   - **Move the contents of My Documents to the new location.** Moves any document the user has in the local My Documents folder to the server share. This option is enabled by default.

   - **Leave the folder in the new location when policy is removed.** Specifies that files remain in the new location when the Group Policy object no longer applies. This option is enabled by default.

   - **Redirect the folder back to the local user profile location when policy is removed.** If enabled, specifies that the folder be copied back to the local profile location if the Group Policy object no longer applies.

The My Documents Properties page provides two additional options for the My Pictures folder:

**Make My Pictures a subfolder of My Documents.** If selected, when the My Documents folder is redirected, My Pictures remains a subfolder of My Documents. By default, My Pictures automatically follows the My Documents folder.

**Do not specify administrative policy for My Pictures.** If selected, Group Policy does not control the location of My Pictures; this is determined by the user profile.

An important point to note is that you should not pre-create the directory defined by user name. Folder Redirection will handle setting the appropriate ACLs on the folder. If you choose to pre-create folders for each user, be sure to set the permissions correctly (see the permissions tables in the Best Practices section later in this paper).

For more information about using the Group Policy snap-in and the Folder Redirection extension, refer to the Windows Server 2003 online Help and the Step-by-Step Guide to User Data and User Settings at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/directory/activedirectory/stepbystep/usrdata.mspx.

**Changing settings after Folder Redirection policy has been applied.**
It is possible to change the Folder Redirection options on the Settings tab after the policy has been applied, you should note that changing the value of the **Grant the user exclusive rights to <folder name>** setting will only apply to new users effected by the policy. Any existing users that received the policy will use the original **Grant the user exclusive rights to <folder name>** setting.

**Folder Redirection and environment variables**
The folder redirection client side extension is only able to process two environment variables: %username% and %userprofile%. Other environment variables such as %logonserver%, %homedrive% and %homepath% will not work with folder redirection.

**Folder Redirection and mapped drives**
Because folder redirection is processed early in the logon process, drives mapped via logon scripts (including the homedrive for folders other than My Documents), the folder redirection client side extension is not able to redirect to these locations. At the time that redirection takes place, the drives do not exist hence redirection fails.

**Folder Redirection Troubleshooting.**
Folder redirection processing contains 5 steps:

1. Determine which folders to redirect based on changes to policy at logon time.
2. Determine desired redirected location and verify access.
3. If folder does not exist: create folders, set ACLs.
4. If folder exists, check ACLs and ownership.
5. If desired, move contents.

Folder redirection failures only affect the folder redirection extension on a per folder basis. If you're pre-creating folders rather than letting the folder redirection extension automatically create the folder, typical errors include:

- Redirecting to a folder that is incorrectly ACLd.
- User is not the owner of the folder.
- Destination does not exist.

**Enabling logging**
In addition to logging events in the Application Event log, Folder Redirection can provide a detailed log to aid troubleshooting. To create a detailed log file for folder redirection, use the following registry key:

- HKLM\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics
- Set: FdeployDebugLevel = Reg_DWORD 0x0f

| Note |
| --- |
| The log file can be found at: %windir%\debug\usermode\fdeploy.log |

**Using Logon Scripts to Redirect Folders**

Although using Group Policy to redirect users folders is the recommended method, there are alternate ways to achieve similar results. You can use logon scripts to set the values of the User Shell Folders key in the registry, which will give you basic functionality similar to Folder Redirection.

Alternatively, you could use Windows NT 4.0 system policies to set the appropriate values. However if you choose to do this, you loose the advantages of using Group Policy to set folder paths, such as automatic moving of files when the path changes, and the registry settings will persist.

# Security Considerations when Configuring Folder Redirection

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**When creating the redirection share, limit access to the share to only users that need access.**

Because redirected folders contain personal information, such as documents and EFS certificates care should be taken to protect them as well as possible. In general:

- Restrict the share to only users that need access. Create a security group for users that have redirected folders on a particular share, and limit access to only those users.
- When creating the share, hide the share by putting a $ after the share name. This will hide the share from casual browsers; the share will not be visible in My Network Places.
- Only give users the minimum amount of permissions needed. The permissions needed are shown in the tables below:

**Table 12 NTFS Permissions for Folder Redirection Root Folder**

| User Account | Minimum permissions required |
|---|---|
| Creator/Owner | Full Control, Subfolders And Files Only |
| Administrator | None |
| Security group of users needing to put data on share. | List Folder/Read Data, Create Folders/Append Data - This Folder Only |
| Everyone | No Permissions |
| Local System | Full Control, This Folder, Subfolders And Files |

**Table 13 Share level (SMB) Permissions for Folder Redirection Share**

| User Account | Default Permissions | Minimum permissions required |
|---|---|---|
| Everyone | Full Control | No Permissions |
| Security group of users needing to put data on share. | N/A | Full Control, |

**Table 14 NTFS Permissions for Each Users Redirected Folder**

| User Account | Default Permissions | Minimum permissions required |
|---|---|---|
| %Username% | Full Control, Owner Of Folder | Full Control, Owner Of Folder |
| Local System | Full Control | Full Control |
| Administrators | No Permissions | No Permissions |
| Everyone | No Permissions | No Permissions |

**Use at least Windows 2000 servers to host redirected file shares.**

Because a users redirected files contain personal information which is copied to and from the client computer, and the server hosting the redirected folders, it is important to ensure that data is protected as it travels over the network.

The biggest potential threats to the privacy and integrity of a users data come from intercepting the data as it passes over the network, tampering with the data as it passes over the network, and spoofing the server hosting the users data.

Several features of Windows 2000 and Windows Server 2003 can help to secure a users data:

- **Kerberos** - Kerberos is standard on all versions of Windows 2000 and Windows Server 2003, and ensures the highest level of security to network resources. While NTLM authenticates the client only, Kerberos authenticates the server and the client. When NTLM is used, the client doesn't know whether the server is valid this is particularly important if the client is exchanging personal files with the server, as is the case with Roaming Profiles. Kerberos provides better security than NTLM and is not available on Windows NT version 4.0 or earlier operating systems.

- **IPSec**- The IP Security Protocol (IPSec) provides network-level authentication, data integrity, and encryption ensuring that roamed data is:

    - Safe from data modification while enroute.

    - Safe from interception, viewing, or copying.

    - Safe from being accessed by unauthenticated parties.

- **SMB Signing**- The Server Message Block (SMB) authentication protocol supports message authentication, which prevents active message and "man-in-the-middle" attacks. SMB signing provides this authentication by placing a digital signature into each SMB, which is then verified by both the client and the server. In order to use SMB signing, you must either enable it or require it on both the SMB client and the SMB server. Note: SMB signing imposes a performance penalty; although it doesn't consume any more network bandwidth, it does use more CPU cycles on the client and server side.

**Always use the NTFS Filesystem for volumes holding users data.**

For the most secure configuration, configure servers hosting redirected files to use the NTFS File System. Unlike FAT, NTFS supports Discretionary access control lists (DACLs) and system access control lists (SACLs), which control who can perform operations on a file and what events will trigger logging of actions performed on a file.

**Do not rely on EFS to encrypt users files when transmitted over the network**

When using the Encrypting File System (EFS) to encrypt files on a remote server, encrypted data is not encrypted when in transit over the network, but only when stored on disk.

The exceptions to this are when your system includes Internet Protocol security (IPSec) or Web Distributed Authoring and Versioning (WebDAV). IPSec encrypts data while it is transported over a TCP/IP network. If the file is encrypted before being copied or moved to a WebDAV folder on a server, it will remain encrypted during the transmission and while it is stored on the server.

**Encrypt the Offline Files cache**

While the Offline Files cache is protected on NTFS partitions by ACLs by default, encrypting the cache enhances security on a local computer. By default, the cache on the local computer is not encrypted, so any encrypted files cached from the network will not be encrypted on the local computer. This may pose a security risk in some environments.

When encryption is enabled, all files in the Offline Files cache are encrypted. This includes existing files as well as files added later. The cached copy on the local computer is affected, but the associated network copy is not.

The cache can be encrypted in one of two ways:

1. Via Group Policy, by enabling the **Encrypt the Offline Files Cache** setting, located at Computer Configuration\Administrative Templates\Network\Offline Files, in the Group Policy editor.

2. Manually, by selecting Tools and then Folder Options in the command menu of Windows Explorer. Select the Offline Files tab, and check the **Encrypt offline files to secure data** checkbox.

**Note**

Encryption of the Offline File cache is only available on Windows XP and above, it is not possible to encrypt the cache on Windows 2000 computers.

For more information about encrypting the Offline files cache, see How to Encrypt Offline Files at: http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/encryptoffline.mspx

**Let the system create folders for each user.**

To ensure that Folder Redirection works optimally, create only the root share on the server, and let the system create the folders for each user. Folder Redirection will create a folder for the user with appropriate security.

If you must create folders for the users, ensure that you have the correct permissions set, also note that if pre-creating folders you must clear the "grant the user exclusive rights to XXX checkbox on the settings tab of the Folder Redirection page. If you don't clear this checkbox, then Folder Redirection will first check a pre-existing folder to ensure the user is the owner. If the folder is pre-created by the administrator, this check will fail and redirection will be aborted. Folder Redirection will then log an event in the Application event log:

Error: Folder Redirection

Event ID: 101

Event Message:

Failed to perform redirection of folder XXXX. The new directories for the redirected folder could not be created. The folder is configured to be redirected to \\server\share, the final expanded path was \\server\share\XXX .

The following error occurred:

This security ID may not be assigned as the owner of this object.

It is strongly recommended that you do not pre-create folders, and allow Folder Redirection to create the folder for the user.

**Ensure correct permissions are set if redirecting to a users home directory.**

Windows Server 2003 and Windows XP allow you to redirect a users My Documents folder to their home directory. When redirecting to the home directory, the default security checks are not made - ownership and the existing directory security are not checked and any existing permissions are not changed - it is assumed that the permissions on the users home directory are set appropriately.

If you are redirecting to a users home directory, be sure that the permissions on the users home directory are set appropriately for your organization.

# Best Practices for Folder Redirection in User Data and Settings Management

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To get the best results from using Folder Redirection, it is important that you read the Group Policy documentation and plan your implementation thoroughly, particularly with respect to Group Policy.

To learn more about how Group Policy and Folder Redirection were developed in Windows 2000, see the Windows 2000 Management Services page at http://www.microsoft.com/windows2000/technologies/management/default.asp.

**Do not use redirect to home directory option unless you have already deployed home directories in your organization.**

Windows Server 2003 allows you to redirect a users My Documents folder to the users home directory. This option is intended only for organizations that have a legacy deployment of home directories and wish to transition users to the My Documents metaphor while maintaining compatibility with their existing home directory environment. You should only select this option if you have already deployed home directories in your organization.

**Let the system create folders for each user.**

To ensure that Folder Redirection works optimally, create only the root share on the server, and let the system create the folders for each user. If you must create folders for the users, ensure that you have the correct permissions set. See Security Considerations when Configuring Folder Redirection for the permissions required.

**Considerations for Group Policy Removal**

It is important to consider the behavior of your Folder Redirection policy settings when Group Policy is removed.

You specify policy removal options in the selected folders Properties page, shown in Figure 2 below. This is accessed under the User Configuration\Windows Settings\Folder Redirection node of the Group Policy snap-in by right-clicking a folder, and clicking Properties. See Configuring Folder Redirection.




Table 10 below summarizes what happens to Redirected Folders and their contents when the Group Policy object no longer applies.

**Table 10 Summary of Redirected Folders**

| Folder Properties page settings | When policy is removed: | |
|---|---|---|

| State of the Move the contents of special folder to the new location option | Setting selected for the Policy Removal option | |
|---|---|---|
| Enabled | **Redirect the folder back to the user profile location when policy is removed** | The special folder returns to its user profile location.<br>The folder contents are copied back to the user profile location.<br>The contents are not deleted from the redirected location.<br>The user can continue to access the contents, but only on the local computer. |
| Disabled | **Redirect the folder back to the user profile location when policy is removed** | The special folder returns to its user profile location.<br>**Note**<br>In this case, the folder contents are not copied or moved to the user profile location. As a result, the user can no longer see the contents. |
| Either Enabled or Disabled | **Leave the folder in the new location when policy is removed** | The special folder remains at its redirected location.<br>The contents remain at the redirected location.<br>The user continues to have access to the contents at the redirected folder. |

**Important**

Changing the redirection option to Not Configured **does not** redirect the folder to the local profile, this option means that Folder Redirection is not configured *if a folder was previously redirected it will continue to be redirected to the previous location*. If an administrator wished to return the folder to the local user profile they should use the **Redirect to the local user profile** setting.

**Use Offline Folder settings on a server share where the users data is stored.**

This is especially useful for users with laptops. In particular, redirected folders of any type should be coupled with Offline Files.

Table 11 below details the recommended configuration to use for Offline Files.

**Table 11 Recommended Configuration for Offline Files**

| Redirected Folder | Recommended Offline Folder Settings |
|---|---|
| My Documents | Auto-caching for documents or manual caching for documents (if you want users to have to manually make files and folders available offline). |
| My Pictures | Auto-caching for documents or manual caching for documents (if you want users to have to manually make files and folders available offline). |
| Application Data | Auto-caching for programs. |
| Desktop | Auto-caching for programs if the desktop is read only. |

**Always enable the Synchronize all offline files before logging off Group Policy setting.**

Enabling this setting ensures that offline files are fully synchronized and that all the files in the users redirected folder are available when they are working offline. If this setting is not enabled, the system only performs a quick synchronization, and as

a result only recently used files will be cached.

**Have My Pictures follow My Documents.**

It is recommended that you configure the My Pictures folder to follow the My Documents folder, unless you have a compelling reason not to.

**In general, accept the Default Settings for Folder Redirection.**

If you are storing roaming profiles on the same server as Offline Files is enabled, Redirected Folders ensure that Offline Files are set to synchronize at logon and logoff.

When a share is unavailable, Offline Files considers the whole server to be unavailable until the offline cache is manually synchronized. Roaming profiles will not be synchronized with the server while Offline Folders considers the server to be unavailable. If you are using Offline Files in conjunction with Folder Redirection and roaming user profiles, for the best performance you should ensure that you leave the default setting of synchronizing Offline Files at logoff enabled.

**Do not open multiple instances of Outlook 2000 when redirecting the Application Data folder.**

If a user has redirected Application Data and has multiple instances of the Outlook 2000 messaging and collaboration client running on different machines, and one of those instances is Outlook 2000, the user will experience poor performance opening e-mails. Outlook 2000 continually holds the file outcmd.dat open. This file stores information about toolbar customizations (such as position) made in Outlook. When another instance of Outlook attempts to access this file, it cannot do so since it is locked by the running instance of Outlook 2000. The second copy of Outlook keeps trying to access outcmd.dat many times, which causes a long delay when opening messages, replying to messages, and so on.

| **Note** |
| --- |
| Outlook version 2002 (included in Office XP) doesn't hold outcmd.dat openthis is only a problem when running Outlook 2000 on one of the machines. |

# Related Technologies: Offline Files and Synchronization Manager

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section highlights technologies that complement the Folder Redirection feature.

**Offline Files**

Offline Files is a feature introduced in Windows 2000 that complements Folder Redirection. Offline Files let users disconnect from the network and work as if they were still connected. When the computer is offline, the files and folders appear in the same directory as they did online as if they still resided in the same location on the network. This allows the user to edit files when they are disconnected. The next time they connect to the network, the offline changes are synchronized with the network share.

By using Offline Files, users can continue to work with a copy of network files even when they are not connected to a network. If your organization has mobile users with portable computers, Offline Files gives them access to their files when they are not connected to the network, and ensures that they are always working with the current version of network files. By using a cached version of the files, users can open and update files even when they are not connected to the network. Offline Files stores the data in the computer's cache to make network files available offline. The cache is a portion of disk space that a computer accesses when it is not connected to the network. The view of shared network items that you have made available offline remains as it is when connected, even if users lose a connection to the network or they remove a portable computer from the docking station. Users can continue to work with the Offline Files as they normally do when online. Users have the same access permissions to those files and folders as when they are connected to the network. When users dock a portable computer and the network connection is restored, any changes they made while working offline are updated to the network.

If two users on the network make changes to the same file, they can save their version of the file to the network, or keep the other user's version, or save both.

Shared files or folders on a Windows Server 2003 or Windows XP-based network can be available offline. You can also make files available for offline use from any computer that is sharing files using server message blockbased file and printer sharing, including Windows 2000, Windows 95, Windows 98, and Windows NT 4.0.

| **Note** |
| --- |
| The Offline Files feature is not available on Novell NetWare networks. |

When configuring a shared folder, you have the option of choosing whether all the files in the folder are automatically available offline, or whether a user must explicitly mark a file to be available offline.

Offline Files is a completely stand-alone technology, which means that you don't need to pair it with Folder Redirection and set up and configure network shares, but it works well if you do pair the two technologies. For example, if a shortcut to a file is available offline, that file is made available offline, but if a shortcut to a folder is available offline, the contents of that folder are not available offline. If you pair the two technologies, Offline Files and Folder Redirection, both the shortcut and the folder are available offline.

By using the manual caching for documents, users manually specify any files that they want available when they are working offline. Automatic caching for documents is recommended for folders that contain user documents. Opened files are automatically downloaded and made available when users work offline. Older copies of the files are automatically deleted to make room for newer and more recently accessed files. The automatic caching of programs is used for folders with read-only data or run-from-the-network applications. To ensure proper file sharing, the server version of the file is always opened.

**Synchronization Manager**

When using Offline Files and folders, users can synchronize all network resources by using the Synchronization Manager. The Synchronization Manager can be set to automatically synchronize some or all resources. For example, users can set certain files and folders to be synchronized every time they log on or off the network. The Synchronization Manager quickly scans the system for any changes, and if it detects changes, the resources are automatically updated. Only resources that have changed are updated vastly speeding up the synchronization process.

# Common Scenarios for IntelliMirror User Data and Settings Features

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This section presents sample scenarios that illustrate some of the practical uses of the user data and settings management features in IntelliMirror.

The scenarios present a snapshot of a users computer in various uses and stages throughout a typical lifecycle. Each of the scenarios fits into an entire picture or can be seen as a separate event and shows how IntelliMirror benefits the entire organization by reducing the time and effort associated with maintaining the computing environment.

The following scenarios are explained:

- The New Hire
- The Laptop User
- Computer Replacement
- A Shared Computer Environment

**The New Hire**

One of the most critical and time consuming IT tasks is setting up the new hire with a computer. In an organization that is using IntelliMirror, the new hire logs on to a new computer and finds documents and shortcuts already on the desktop. There are shortcuts to common files, URLs, and folders that are useful to all employees (for example, the employee handbook, a shortcut to the departmental shared documents store, and a shortcut to the users departmental guidelines and procedures). Desktop options, application configurations, Internet settings, and so on, are all configured to the corporate standard, ensuring that if the user needs to call the help desk, the support staff knows the configuration the user started with.

In this example, the user gets a pre-configured user profile that was set up for all new users, and was configured before the new hire logged on to the network. The administrator configured a computer to look and behave according to the corporate standard, and then using the User Profile[1] utility built into the System Control Panel application, copied the user profile to a Default User folder on the domain controllers Netlogon share. When the new hire logged onto the network for the first time, Windows copied this default profile to the local computer and used this profile as the basis for the new hires profile. In addition to configuring the default profile the user received, the administrator also used Group Policy to redirect the users My Documents folder to a network location, so that the users documents are safely stored on a network server and can be backed up regularly.

**The Laptop User**

A laptop user working at the office creates several documents and saves them to his or her My Documents folder. After saving documents, the user logs off, unplugs the laptop computer from the network and takes it home. While at home and off the network, the user continues to edit the documents saved earlier in My Documents.

The user returns to the office and logs on to the network. Since the user has done some offline work, a dialog box appears advising the user that data in My Documents has changed and is being synchronized with the network copy.

In this scenario, the users My documents folder has been redirected to a network server, the documents are transparently saved to the network location and also saved in the local computers cache (because the network folder is setup to be available offline), so that they are available when the computer is disconnected from the network.

The whole process can be transparent to the user; the experience is no different than saving documents to the local hard disk.

As soon as the user reconnects to the network, IntelliMirror attempts to reconnect to the network location of the redirected folders. When IntelliMirror reconnects, it determines if there are differences in the data between the local copy of the folder and the network copy. In this scenario, the user has made modifications to a document on the local computer. IntelliMirror identifies this change and prompts the user to update the version stored on the network.

**Computer Replacement**

The computer that the user is working on suddenly stops working because of a complete hardware failure. The user calls the support line, and about 20 minutes later a new computer, loaded only with the Windows XP Professional operating system arrives for the user. Without waiting for technical assistance, the user plugs in the new computer, connects it to the network, and boots it. The computer allows the user to log on to the corporate network, and the user finds that the desktop has the same appearance as the original computer that it replaced. It has the same color scheme, the users preferred background picture is on the screensaver, and all the application icons, shortcuts, and favorites are present. More importantly, all the users data files have been restored.

In a disaster recovery scenario, IntelliMirror assists in getting the users computer replaced and running quickly and with the minimum of support. In this example, because the user was configured to use roaming user profiles, a copy of the users working environment was safely stored on a network server. When the new computer arrived, the user was able to log on and the server copy of the users profile was downloaded to the new computer. An administrator could also have used Folder Redirection to redirect the users key folders such as My Documents and Application Data, to ensure that the users documents were safely stored on the server.

**A Shared Computer Environment**

In this scenario, a user works in a department where the computer he or she uses may change from day to day a call center or IT support environment, for example. The user is working on an important document late one night when the shift ends. The user saves the document and logs off the computer. When the user returns to work the next day, he or she logs onto the first available computer a different computer from the one used the previous night. The user logs onto the network, and sees that the desktop has the same look and feel as the original computer. The user opens the My Documents folder on the desktop and finds the document exactly where he or she saved it and continues the work started the previous night.

In this example, the user was configured to use roaming user profiles, so that a copy of the users working environment was stored on a network server. When the user logged onto the computer, the users existing preferences, shortcuts and documents were copied to the local computer, so that the user could continue working as if using the original computer. A variation on this scenario is using roaming profiles in conjunction with Folder Redirection. Users can have the same work environment and access to the same documents on any computer. Changes made on one computer are synchronized with the other computer the next time the user logs on.

[1] To access the User Profile utility, click Start, point to Settings, select Control Panel, select System, and then select the User Profiles tab in the System Properties dialog box.

# Appendix: Group Policy Settings for Roaming User Profiles

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Settings for Roaming User Profiles**

| Policy | Location in Group Policy Snap-in | Description |
|---|---|---|
| **Limit profile size** | User Configuration\Administrative Templates\System\User Profiles | Sets the maximum size of a roaming user profile and determines the system's response when a roaming user profile reaches the limit. |
| **Connect home directory to root of the share** | User Configuration\Administrative Templates\System\User Profiles | Restores the definitions of the %HOMESHARE% and %HOMEPATH% environment variables to those used in Windows NT 4.0 and earlier. |
| **Exclude directories in roaming profile** | User Configuration\Administrative Templates\System\User Profiles | Lets you add to the list of folders excluded from the user's roaming profile. |
| **Delete cached copies of roaming profiles** | Computer Configuration\Administrative Templates\System\ User Profiles | Determines whether the system saves a copy of a user's roaming profile on the local computer's hard disk drive when the user logs off. This policy and the related policies in this folder define a strategy for managing user profiles that reside on remote servers. Specifically, they indicate to the system how to respond when a remote profile is slow to load. |
| **Slow network connection timeout for user profiles** | Computer Configuration\Administrative Templates\System\ User Profiles | Defines a slow connection for roaming user profiles. If the server on which the user's roaming user profile resides takes longer to respond than the thresholds set by this policy permit, the system considers the connection to the profile to be slow. This policy and related policies in this folder together define the system's response when roaming user profiles are slow to load. |
| **Add the Administrator security group to the roaming user profile share**<br>*New for Windows XP and Windows Server 2003* | Computer Configuration\Administrative Templates\System\ User Profiles | This policy allows an administrator to choose the same behavior as Windows NT 4.0, and have the administrators group given full control of the user's profile directories. |
| **Prevent Roaming Profile Changes From Propagating to the server**<br>*New for Windows XP and Windows Server 2003* | Computer Configuration\Administrative Templates\System\ User Profiles | This policy determines if the changes a user makes to their roaming profile are merged with the server copy of their profile. If this policy is set, at login the user will receive their Roaming Profile, but any changes a user makes to their profile will not be merged to their roaming profile at logoff. |

| | | |
|---|---|---|
| **Only allow local profiles**<br>*New for Windows XP and Windows Server 2003* | Computer Configuration\Administrative Templates\System\ User Profiles | This setting determines if roaming user profiles are available on a particular computer. By default, when roaming profile users log on to a computer, their roaming profile is copied down to the local computer. Using this setting, an administrator can prevent users configured to use roaming profiles from receiving their roaming profile on a specific computer. |
| **Prompt user when slow link is detected** | Computer Configuration\Administrative Templates\System\ User Profiles | Notifies users when their roaming profiles are slow to load, letting a user decide whether to use a local copy or to wait for the roaming user profile.<br>If you disable this policy or do not configure it, when a roaming user profile is slow to load, the system does not notify the user. It loads the local copy of the profile. If you have enabled the Wait for remote user profile policy, the system loads the remote copy without prompting the user. |
| **Maximum retries to unload and update user profile** | Computer Configuration\Administrative Templates\System\ User Profiles | Windows 2000 Only. Determines how many times the system tries to unload and update the registry portion of a user profile. When the number of trials specified by this policy is exhausted, the system stops trying. As a result, the user profile might not be current, and local and roaming user profiles might not match. |
| **Do not detect slow network connections** | Computer Configuration\Administrative Templates\System\ User Profiles | Disables the slow link detection feature.<br>Slow link detection measures the speed of the connection between a user's computer and the remote server that stores the roaming user profile. When the system detects a slow link, the related policies in this folder tell the system how to respond. |
| **Wait for remote user profile** | Computer Configuration\Administrative Templates\System\ User Profiles | Directs the system to wait for the remote copy of the roaming user profile to load, even when loading is slow. The system waits for the remote copy when the user is notified about a slow connection, but does not respond in the time allowed. |
| **Timeout for dialog boxes** | Computer Configuration\Administrative Templates\System\ User Profiles | Determines how long the system waits for a user response before it uses a default value. The default value is used when the user does not respond to messages explaining that any of the following events has occurred:<br>• The system detects a slow connection between users' computers and the server that stores the users' roaming user profiles.<br>• The system cannot access users' server-based profiles when users log on or off.<br>• Users' local profiles are newer than their server-based profiles. |

# Related Links for User Data and Settings Management

Updated: March 1, 2002

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To learn more about how change and configuration management features were introduced in Windows 2000 Server, see the following Windows 2000 resources:

- Windows 2000 Management Services at: http://www.microsoft.com/windows2000/technologies/management/default.asp.
- Step-by-Step Guide to User Data and User Settings at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/directory/activedirectory/stepbystep/usrdata.mspx
- For the latest information about Windows XP, check out our Web site at http://www.microsoft.com/windowsxp/default.asp.

# SSL/TLS in Windows Server 2003

Updated: July 31, 2003

Applies To: Windows Server 2003 with SP1

Chris Crall, Mike Danseglio, and David Mowers

Microsoft Corporation

Published: July 2003

The Microsoft® Windows Server™ 2003 operating system supports Schannel, an implementation of three protocols (TLS 1.0, SSL 3.0 & SSL 2.0) that provide network security for applications and services. This paper focuses on the Secure Sockets Layer (SSL 3.0) protocol and Transport Layer Security (TLS 1.0) protocol. These protocols offer certificate-based authentication and secure data transfers using symmetric encryption keys.

**In This White Paper**

Introduction (SSL/TLS in Windows Server 2003)

Overview of SSL/TLS Encryption

SSL/TLS in Detail

SSL/TLS Scenarios

SSL and Firewalls

Performance Considerations

# Introduction (SSL/TLS in Windows Server 2003)

Updated: July 31, 2003

Applies To: Windows Server 2003 with SP1

How can you secure data being sent between applications across an untrusted network? The Security Support Provider Interface (SSPI) in the Microsoft® Windows Server™ 2003 operating system provides one answer to this very common security question. SSPI is an application interface that provides the security services for Windows Server 2003. SSPI supports Schannel, an interface that implements three industry standard protocols:

- Transport Layer Security (TLS) version 1.0
- Secure Sockets Layer (SSL) version 3.0
- Secure Sockets Layer (SSL) version 2.0

SSL/TLS is most widely recognized as the protocol that provides secure HTTP (HTTPS) for internet transactions between Web browsers and Web servers. It can also be used for other application level protocols such as FTP, LDAP, and SMTP. The SSL/TLS protocol enables server authentication, client authentication, data encryption and data integrity over networks such as the World Wide Web.

SSPI is documented in the Microsoft Platform Software Development Kit (SDK), which also includes sample applications.

**History and Standards for SSL and TLS**

SSL was developed by Netscape Communications Corporation in 1994 to secure transactions over the World Wide Web. Soon after, the Internet Engineering Task Force (IETF) began work to develop a standard protocol to provide the same functionality. SSL 3.0 was used as the basis for that work, which is known as the Transport Layer Security protocol (TLS). The implementation of the SSL/TLS protocol in Windows Server 2003 closely follows the specification defined in RFC 2246, "The TLS Protocol Version 1.0."

**Differences Between SSL and TLS**

Although there are some slight differences between SSL 3.0 and TLS 1.0, this paper will refer to the protocol as SSL/TLS. One important difference is that TLS 1.0 applies a Keyed-Hashing for Message Authentication Code (HMAC) algorithm, whereas SSL 3.0 applies the Message Authentication Code (MAC) algorithm. The HMAC produces an integrity check value as the MAC does, but with a hash function construction that makes the hash much harder to break. For more information about the HMAC, see "The Handshake Protocol" later in this paper.

| **Note** |
| --- |
| Although their differences are minor, TLS 1.0 and SSL 3.0 do not interoperate. If the same protocol is not supported by both parties, the parties must negotiate a common protocol to communicate successfully. |

**Benefits of SSL/TLS**

SSL/TLS provides numerous benefits to clients and servers, including:

**Strong authentication, message privacy, and integrity.** The primary feature of SSL/TLS is the ability to secure transmitted data using encryption. SSL/TLS also offers server authentication and, optionally, client authentication to prove the identities of parties engaged in secure communication. It also provides data integrity through an integrity check value. In addition to protecting against data disclosure through encryption, the SSL/TLS security protocol can be used to protect against masquerade attacks, man-in-the-middle or bucket brigade attacks, rollback attacks, and replay attacks.

**Interoperability.** SSL/TLS works with most Web browsers, including Microsoft Internet Explorer and Netscape Navigator, and on most operating systems and Web servers including the Microsoft® Windows operating system, UNIX, Novell, Apache (version 1.3 and later), Netscape Enterprise Server, and Sun Solaris. In addition, it is often integrated in news readers, LDAP servers, and a variety of other applications.

**Algorithm flexibility.** SSL/TLS provides options for the authentication mechanisms, encryption algorithms, and hashing algorithms that will be used during the secure session.

**Ease of deployment.** SSL/TLS is used transparently by many applications on Windows Server 2003. Using SSL for secure

browsing when using Internet Explorer and Internet Information Services (IIS) is as easy as selecting a check box.

**Ease of use.** Because SSL/TLS is implemented beneath the application layer, most of its operations are completely invisible to the client. This allows the client to have little or no knowledge of secure communications and still be protected from attackers.

#### Drawbacks of SSL/TLS

There are a few drawbacks to using SSL/TLS, including:

**Increased processor load.** This is the most significant drawback to implementing SSL/TLS. Cryptography, specifically public key operations, are CPU intensive. As a result, there is a performance penalty when using SSL. Unfortunately, there is no single answer to the frequently asked question: how much of performance penalty? The penalty varies widely depending on how often connections are established and how long they last. The greatest overhead occurs while connections are being set up.

**Administrative overhead.** An SSL/TLS environment is complex and requires maintenance; the system administrator needs to configure the system and manage certificates.

# Overview of SSL/TLS Encryption

Updated: July 31, 2003

Applies To: Windows Server 2003 with SP1

SSL/TLS is primarily used to encrypt confidential data sent over an insecure network such as the Internet. In the HTTPS protocol, the types of data encrypted include the URL, the HTTP header, cookies, and data submitted through forms. A Web page secured with SSL/TLS has a URL that begins with "https://".

The SSL/TLS security protocol is layered between the application protocol layer and TCP/IP layer, where it can secure and then send application data to the transport layer. Because it works between the application layer and the TCP/IP layer, SSL/TLS can support multiple application layer protocols.

The SSL/TLS protocol can be divided into two layers. The first layer is the Handshake Protocol Layer, which consists of three sub-protocols: the Handshake Protocol, the Change Cipher Spec Protocol, and the Alert protocol. The second layer is the Record Protocol Layer. Figure 1 illustrates the various layers and their components.

**Figure 1. SSL/TLS Protocol Layers**




**The Handshake Layer**

The Handshake Layer consists of three sub-protocols:

- **Handshake.** This sub-protocol is used to negotiate session information between the client and the server. The session information consists of a session ID, peer certificates, the cipher spec to be used, the compression algorithm to be used, and a shared secret that is used to generate keys.

- **Change Cipher Spec.** The Change Cipher Spec sub-protocol is used to change the *keying material* used for encryption between the client and server. Keying material is raw data that is used to create keys for cryptographic use. The Change Cipher Spec sub-protocol consists of a single message to tell other party in the SSL/TLS session, who is also known is the peer, that the sender wants to change to a new set of keys. The key is computed from the information exchanged by the Handshake sub-protocol.

- **Alert.** Alert messages are used to indicate a change in status or an error condition to the peer. There are a wide variety of alerts to notify the peer of both normal and error conditions. A full list can be found in RFC 2246, "The TLS Protocol Version 1.0.". Alerts are commonly sent when the connection is closed, an invalid message is received, a message cannot be decrypted, or the user cancels the operation.

**Handshake Sub-Protocol Functions**

The Handshake sub-protocol provides a number of very important security functions. It performs a set of exchanges that starts authentication and negotiates the encryption, hash, and compression algorithms.

**Authentication**

For authentication purposes, the Handshake Protocol uses an X.509 certificate to provide strong evidence to a second party that helps prove the identity of the party that holds the certificate and the corresponding private key. A certificate is a digital form of identification that is usually issued by a certification authority (CA) and contains identification information, a validity period, a public key, a serial number, and the digital signature of the issuer.

A CA is a mutually trusted third party that confirms the identity of a certificate requester (usually a user or computer), and then issues the requester a certificate. The certificate binds the requester's identity to a public key. CAs also renew and revoke certificates as necessary. For example, if a client is presented with a server's certificate, the client computer might try to match the server's CA against the client's list of trusted CAs. If the issuing CA is trusted, the client will verify that the certificate is authentic and has not been tampered with. Finally, the client will accept the certificate as proof of identity of the server.

For more information about certificates and how they are used, see http://go.microsoft.com/fwlink/?LinkId=18015.

| **Note** |
| --- |
| Windows Server 2003 includes software that allows it to function as a certification authority. For more information, see Help and Support Center in Windows Server 2003. |

**Encryption**

There are two main types of encryption: symmetric key (also known as shared secret key) and asymmetric key (also known as public key or public-private key). SSL/TLS uses both symmetric key and asymmetric key encryption.

- **Symmetric Key**. In symmetric key encryption, the same key is used to encrypt and decrypt the message. If two parties want to exchange encrypted messages securely, they must both possess a copy of the same symmetric key. Symmetric key cryptography is often used for encrypting large amounts of data because it is computationally faster than asymmetric cryptography. Typical algorithms include DES (Data Encryption Standard), 3-DES (Triple DES), RC2, RC4, and AES (Advanced Encryption Standard).

- **Asymmetric Key**. Asymmetric or public key encryption uses a pair of keys that have been derived together through a complex mathematical process. One of the keys is made public, typically by asking a CA to publish the public key in a certificate for the certificate-holder (also called the subject). The private key is kept secret by the subject and never revealed to anyone. The keys work together, with one being used to perform the inverse operation of the other: If the public key is used to encrypt data, only the private key of the pair can decrypt it; if the private key is used to encrypt, the public key must be used to decrypt. This relationship allows a public key encryption scheme to do two important things. First, anyone can obtain the public key for a subject and use it to encrypt data that only the user with the private key can decrypt. Second, if a subject encrypts data using its private key, anyone can decrypt the data by using the corresponding public key. This is the foundation for digital signatures. The most common algorithm is RSA (Rivest, Shamir & Adleman).

SSL/TLS uses public key encryption to authenticate the server to the client, and optionally the client to the server. Public key cryptography is also used to establish a *session key*. The session key is used in symmetric algorithms to encrypt the bulk of the data. This combines the benefit of asymmetric encryption for authentication with the faster, less processor-intensive symmetric key encryption for the bulk data.

**Hash Algorithms**

During the Handshake process the hash algorithm is also agreed upon. A hash is a one-way mapping of values to a smaller set of representative values, so that the size of the resulting hash is smaller than the original message and the hash is unique to the original data. A hash is similar to a fingerprint: a fingerprint is unique to the individual and is much smaller than the original person. Hashing is used to establish data integrity during transport. Two common hash algorithms are Message Digest 5 (MD5) and Standard Hash Algorithm 1 (SHA-1). MD5 produces a 128 bit hash value and SHA-1 produces a 160 bit value.

The hash algorithm includes a value used to check the integrity of the transmitted data. This value is established using either a MAC or an HMAC. The MAC uses a mapping function to represent the message data as a fixed-length, preferably smaller, value and then hashes the message. The MAC ensures that the data has not been modified during transmission. The difference between a MAC and a digital signature is that a digital signature is also an authentication method. SSL uses a MAC.

The HMAC is similar to the MAC but uses a hash algorithm in combination with a shared secret key. The shared secret key is appended to the data to be hashed. This makes the hash more secure because both parties must have the same shared secret key to prove the data is authentic. TLS uses an HMAC. For more information about HMAC see RFC 1024, "Keyed-Hashing for Message Authentication."

For more information about public key cryptography, see "Designing a Public Key Infrastructure" in *Designing and Deploying Directory and Security Services of the Microsoft® Windows® Server 2003 Deployment Kit*, (or see "Designing a Public Key Infrastructure" on the Web at http://go.microsoft.com/fwlink/?LinkID=4735).

**The Record Layer**

The protocol at the record layer receives and encrypts data from the application layer and delivers it to the Transport Layer. The Record Protocol takes the data, fragments it to a size appropriate to the cryptographic algorithm, optionally compresses it (or,

for data received, decompresses it), applies a MAC or HMAC (HMAC is supported only by TLS) and then encrypts (or decrypts) the data using the information negotiated during the Handshake Protocol.

# SSL/TLS in Detail

Updated: July 31, 2003

Applies To: Windows Server 2003 with SP1

This section provides a detailed explanation of the SSL/TLS protocol, specifically the handshake protocol, its associated messages and alerts, and the record protocol.

**The Handshake Protocol**

The handshake protocol is a series of sequenced messages that negotiate the security parameters of a data transfer session. Figure 2 illustrates the message sequence in the handshake protocol.

**Figure 2. Handshake Protocol Messages**




**Initial Client Message to Server**

**Client Hello.** The client initiates a session by sending a Client Hello message to the server. The Client Hello message contains:

- **Version Number.** The client sends the version number corresponding to the highest version it supports. Version 2 is used for SSL 2.0, version 3 for SSL 3.0, and version 3.1 for TLS. Although the IETF RFC for TLS is TLS version 1.0, the protocol uses 3.1 in the version field to indicate that it is a higher level (newer and with more functionality) than SSL 3.0.

- **Randomly Generated Data.** ClientRandom[32], the random value, is a 4-byte number that consists of the client's date and time plus a 28-byte randomly generated number that will ultimately be used with the server random value to generate a master secret from which the encryption keys will be derived.

- **Session Identification (if any).** The sessionID is included to enable the client to resume a previous session. Resuming a previous session can be useful, because creating a new session requires processor-intensive public key operations that can be avoided by resuming an existing session with its established session keys. Previous session information, identified by the sessionID, is stored in the respective client and server session caches.

- **Cipher Suite.** The A list of cipher suites available on the client. An example of a cipher suite is TLS_RSA_WITH_DES_CBC_SHA, where TLS is the protocol version, RSA is the algorithm that will be used for the key exchange, DES_CBC is the encryption algorithm (using a 56-bit key in CBC mode), and SHA is the hash function.

- **Compression Algorithm.** The requested compression algorithm (none currently supported).

The following is an example of a Client Hello message:

```
ClientVersion 3,1
ClientRandom[32]
SessionID: None (new session)
Suggested Cipher Suites:
   TLS_RSA_WITH_3DES_EDE_CBC_SHA
   TLS_RSA_WITH_DES_CBC_SHA
Suggested Compression Algorithm: NONE
```

#### Server Response to Client

**Server Hello.** The server responds with a Server Hello message. The Server Hello message includes:

- **Version Number.** The server sends the highest version number supported by both sides. This is the lower of: the highest version number the server supports and the version sent in the Client Hello message.

- **Randomly Generated Data.** ServerRandom[32], the Random Value, is a 4-byte number of the server's date and time plus a 28-byte randomly generated number that will be ultimately used with the client random value to generate a master secret from which the encryption keys will be derived

- **Session Identification (if any).** This can be one of three choices.

    - New session ID – The client did not indicate a session to resume so a new ID is generated. A new session ID is also generated when the client indicates a session to resume but the server can't or won't resume that session. This latter case also results in a new session ID.

    - Resumed Session ID– The id is the same as indicated in the client hello. The client indicated a session ID to resume and the server is willing to resume that session.

    - Null – this is a new session, but the server is not willing to resume it at a later time so no ID is returned.

- **Cipher Suite.** The server will choose the strongest cipher that both the client and server support. If there are no cipher suites that both parties support, the session is ended with a "handshake failure" alert.

- **Compression Algorithm.** Specifies the compression algorithm to use (none currently supported).

The following is an example of a Server Hello Message:

```
Version 3,1
ServerRandom[32]
SessionID: bd608869f0c629767ea7e3ebf7a63bdcffb0ef58b1b941e6b0c044acb6820a77
Use Cipher Suite:
TLS_RSA_WITH_3DES_EDE_CBC_SHA
Compression Algorithm: NONE
```

**Server Certificate.** The server sends its certificate to the client. The server certificate contains the server's public key. The client will use this key to authenticate the server and to encrypt the premaster secret.

The client also checks the name of the server in the certificate to verify that it matches the name the client used to connect. If the user typed www.contoso.com as the URL in the browser, the certificate should contain a subject name of www.contoso.com or *.contoso.com. Internet Explorer will warn the user if these names do not match.

**Server Key Exchange.** This is an optional step in which the server creates and sends a temporary key to the client. This key can be used by the client to encrypt the Client Key Exchange message later in the process. The step is only required when the public key algorithm does not provide the key material necessary to encrypt the Client Key Exchange message, such as when the server's certificate does not contain a public key.

**Client Certificate Request.** This is an optional step in which the server requests authentication of the client. This step might be used for Web sites (such as a banking Web site) where the server must confirm the identity of the client before providing sensitive information.

**Server Hello Done.** This message indicates that the server is finished and awaiting a response from the client.

#### Client Response to Server

**Client Certificate.** If the server sent a Client Certificate Request, the client sends its certificate to the server for client authentication. The client's certificate contains the client's public key.

**Client Key Exchange.** The client sends a Client Key Exchange message after computing the premaster secret using both random values. The premaster secret is encrypted by the public key from the server's certificate before being transmitted to the server. Both parties will compute the master secret locally and derive the session key from it.

If the server can decrypt this data and complete the protocol, the client is assured that the server has the correct private key. This step is crucial to prove the authenticity of the server. Only the server with the private key that matches the public key in

the certificate can decrypt this data and continue the protocol negotiation.

This message will also include the protocol version. The server will verify that it matches the original value sent in the client hello message. This measure guards against rollback attacks. Rollback attacks work by manipulating the message in order to cause the server and the client to use a less secure, earlier version of the protocol.

**Certificate Verify.** This message is sent only if the client previously sent a Client Certificate message. The client is authenticated by using its private key to sign a hash of all the messages up to this point. The recipient verifies the signature using the public key of the signer, thus ensuring it was signed with the client's private key.

**Change Cipher Spec.** This message notifies the server that all messages that follow the Client Finished message will be encrypted using the keys and algorithms just negotiated.

**Client Finished.** This message is a hash of the entire conversation to provide further authentication of the client. This message is the first message that the record layer encrypts and hashes.

#### Server Final Response to Client

**Change Cipher Spec Message.** This message notifies the client that the server will begin encrypting messages with the keys just negotiated.

**Server Finished Message.** This message is a hash of the entire exchange to this point using the session key and the MAC secret. If the client is able to successfully decrypt this message and validate the contained hashes, it is assured that the SSL/TLS handshake was successful, and the keys computed on the client machine match those computed on the server.

#### The Alert Sub-Protocol

The alert sub-protocol is a component of the Handshake protocol that includes event-driven alert messages that can be sent from either party. Following an alert message the session is either ended or the recipient is given the choice of whether or not to end the session. The alerts are defined in the TLS specification in RFC 2246. Table 1 lists the alert messages and their descriptions.

**Table 1. Event-Driven Alert Messages from the Alert Sub-Protocol**

| Alert Message | Fatal? | Description |
|---|---|---|
| unexpected_message | Yes | Inappropriate message |
| bad_record_mac | Yes | Incorrect MAC |
| decryption_failed | Yes | Unable to decrypt TLSCiphertext correctly |
| record_overflow | Yes | Record is more than 2^14+1024 bytes |
| handshake_failure | Yes | Unacceptable security parameters |
| bad_certificate | Yes | There is a problem with the Certificate |
| unsupported_certificate | No | Certificate is unsupported |
| certificate_revoked | No | Certificate has been revoked |
| certificate_expired | No | Certificate has expired |
| certificate_unknown | No | Certificate is unknown |
| illegal_parameter | Yes | Security parameters violated |
| unknown_ca | Yes | CA unknown |
| access_denied | Yes | Sender does not want to negotiate |
| decode_error | Yes | Unable to decode message |
| decrypt_error | Yes | Handshake cryptographic operation failed |
| export_restriction | Yes | Not in compliance with export regulations |

| | | |
|---|---|---|
| protocol_version | Yes | Protocol version not supported by both parties |
| insufficient_security | Yes | Security requirements no met |
| internal_error | Yes | Error not related to protocol |
| user_canceled | Yes | Not related to protocol failure |
| no_renegotiation | No | Negotiation request rejected. |

**The Record Protocol**

The record protocol receives the data from the application layer and:

- Fragments the data into blocks or reassembles fragmented data into its original structure.
- Numbers the sequence of data blocks in the message to protect against attacks that attempt to reorder data.
- Compresses or decompresses the data using the compression algorithm negotiated in the handshake protocol.
- Encrypts or decrypts the data using the encryption keys and cryptographic algorithm negotiated during the handshake protocol.
- Applies an HMAC (or a MAC for SSL 3.0) to outgoing data. Computes the HMAC and verifies that it is identical to the value transmitted in order to check data integrity when a message is received.

Once the record protocol has completed its operations on the data, it sends the data to the TCP/IP transport layer for transmission. If the data is incoming, it is sent to the appropriate process for reception.

# SSL/TLS Scenarios

Updated: July 31, 2003

Applies To: Windows Server 2003 with SP1

Many people think of SSL and TLS as protocols used with Web browsers for securely browsing the Internet. However, these are also general purpose protocols that can be used whenever authentication and data protection are necessary. The following examples depict a few uses of SSL/TLS today. This is not an exhaustive list. In fact the ability to access these protocols through the SSPI interface means that anyone take advantage of them for just about any application. Many applications are being modified to take advantage of the features of SSL/TLS.

**Secure transaction with an e-commerce Web site.** This is a typical use of SSL between a browser and a Web server. An example is an e-commerce shopping site where clients need to furnish their credit card numbers. The protocol would first confirm that the Web site's certificate was valid and then send the client's credit card information as cipher text. For this type of transaction, where the server's certificate is from a trusted source, only server-side authentication occurs. SSL/TLS would need to be enabled for the Web page, such as an order form, where the data transactions occur.

**Authenticated client access to a secure Web site.** Both the client and server need certificates from a mutually trusted CA. With Schannel, client certificates can be mapped on a one-to-one or many-to-one basis to their Windows Server 2003 user or computer accounts and can be managed by Active Directory Users and Computers. This is invisible to the users, who can be authenticated to a Web site without needing to supply a password.

If you want to give several users access to confidential material, you can create a group, map the users' certificates to the group, and give the group permissions to the material.

In one-to-one mapping, the server has a copy of the client's certificate; whenever the client logs in, the server verifies that they are identical. This one-to-one mapping is typically used for private material, such as a banking site where only one individual has the right to view a personal account.

**Remote Access.** Schannel is used to provide authentication and data protection when users remotely log in to Windows-based systems or networks. Telecommuting is a common use for this technology. Users can more securely access their e-mail or enterprise applications from home or while traveling, reducing the risk of exposure of the information to anyone on the Internet.

**SQL Access.** Microsoft® SQL Server™ provides the ability for administrators to require authentication of the client when connecting to the server running SQL Server. In addition, either the client or server can be configured to require encryption of the data transferred between them. Very sensitive information, such as financial or medical databases, can be protected to prevent unauthorized access and disclosure of information on the network.

**E-mail.** Exchange servers can use Schannel to protect data as it moves from server to server on the Intranet or Internet. Full end-to-end security might require the use of Secure/Multipurpose Internet Mail Extensions (S/MIME); however, the protection of data in a server-to-server exchange allows companies to use the Internet to securely transfer e-mail among divisions within the same company, subsidiaries and partners. This can be done regardless of whether S/MIME is used.

# SSL and Firewalls

Updated: July 31, 2003

Applies To: Windows Server 2003 with SP1

You must make some additional decisions if you need to conduct SSL/TLS transactions through a firewall. A firewall is a program that can exist in many different forms, but essentially functions as a barrier between your local area network (LAN) and the outside world. The SSL/TLS protocol interprets a computer on which a firewall is running as presenting a man-in-the-middle attack, which prevents the transaction from happening.

You can use one of two approaches to facilitate SSL/TLS transactions through a firewall:

- **Open the firewall to allow all traffic through a designated port.** The typical port for HTTP over SSL is 443. This port can be opened to allow traffic through to the destination Web server. Unfortunately, this means that the firewall can make security decisions based only on the apparent origin of the packet and its destination. The firewall cannot examine the encrypted data in the requests.

- **Configure the firewall or boundary system as a proxy server.** In this case, the boundary system is the destination for the SSL traffic from the client. The client will authenticate to the boundary system, which will then forward, or *proxy*, the requests to the internal system. The connection from the boundary system to the internal system might or might not be protected by using SSL. This presents an authentication problem because the proxy needs to transmit the authenticated identity of the original user to the internal system. It is not possible to use the certificate mapping features of Windows Server 2003 at the application server, because the authentication process that relies on the user's certificate takes place at the proxy.

# Performance Considerations

Updated: July 31, 2003

Applies To: Windows Server 2003 with SP1

Using the SSL protocol for encryption protects data traveling over the Internet, but it also imposes a performance penalty. SSL can slow down an application or Web site considerably. The most obvious performance degradation when using SSL/TLS results from the time required to establish a session and then encrypt and decrypt the data, all of which heavily use processor cycles.

There are ways to minimize this impact. This section discusses the performance issues related to the cryptography involved and ways to accelerate the performance of the public key operations in SSL/TLS by using hardware solutions. It also covers considerations for scaling Web sites that use SSL/TLS, and performance tuning options for Schannel in Windows Server 2003.

**Asymmetric Encryption**

The process of SSL/TLS key exchange can exact a performance penalty. Nearly all commercially available Web servers and clients implement and support RSA as the asymmetric key exchange algorithm. Although there are other key exchange algorithms supported by SSL/TLS, RSA is the most widely used algorithm.

**Offload Hardware Advantages**

The high computational cost of the RSA public-key encryption algorithm makes it attractive to use encryption offload hardware to accelerate processing. Using offload hardware offers two advantages over performing the RSA encryption in software:

- **Optimizes processing.** The general-purpose processors normally used in computers do most things very well. However, special processors optimized for RSA are able to perform specific mathematical operations much more efficiently.

- **Decreases system-wide load.** Offloading the large number of clock cycles required for RSA decryption to a specialized processor frees the computer's general purpose processor to do other things, such as serving Web pages to other clients, while the decryption completes.

**Offload Devices**

There are three types of offload devices: RSA offload; Hardware Security Modules (HSMs), and Internet Hardware Devices. The first two, RSA offload and HSM devices, are typically cards (such as PCI cards) that plug into the bus on the computer. They interact with the Windows Server 2003 operating system through the Crypto API (CAPI). Internet Hardware Devices are network devices.

- **RSA-offload only devices.** This type of device is designed to provide offloading of RSA operations from the host CPU. These devices are typically cards that are installed with a minimum of effort and little to no configuration. They perform the RSA operation only, and offer no additional features, such as improved key security.

- **Hardware Security Modules (HSM).** These devices are primarily targeted to deployments where security of the private key is a priority. By default, private keys are stored in software. With HSMs, not only are the RSA operations offloaded, but private keys never leave the HSM. These devices are normally tamper-proof and meet varying levels of security certification.

- **Internet Hardware Devices.** This loosely-defined class of devices is generally targeted at solving various problems in deploying e-commerce and internet solutions. These devices all offload SSL operations from the Web Servers by terminating the SSL connections at the device and then proxy the request as standard HTTP. Many of these devices have additional features such as content caching and load-balancing.

**Symmetric Encryption**

Symmetric encryption is so efficient that offloading it to specialized hardware is unnecessary, and might even be counterproductive. For most symmetric algorithms, fewer clock cycles are required to process the data to be encrypted or decrypted on the main processor than to move it from main memory, through the I/O bus, onto the offload hardware, back through the bus and back into memory. Add in the necessary trip through the bus to the network adapter, and it becomes clear why this is commonly referred to as the triple-trip problem. Because all these trips would themselves require resources, and because most symmetric algorithms are optimized for computer processing, it is far more efficient to simply perform the operation on the processor.

**Considerations for Scaling When Using SSL**

Scaling out a Web site means adding servers and then sending client requests to these additional computers, thereby increasing the performance of the site. This is often done by using technologies such as software load balancing, hardware load balancing, and DNS round robin. The traditional methods for scaling out a Web site can still be utilized with SSL/TLS, but you must also deal with a specific requirement imposed by SSL/TLS.

Once an SSL/TLS session is established between a specific client and a specific server, only that server can encrypt and decrypt the requests and responses for that client. Consequently, during an SSL/TLS session a client cannot send an SSL response to a server different from the one with which the client has already negotiated. This would result in an error, and a new SSL/TLS session would have to be negotiated between the client and the new server.

Therefore, when scaling out an SSL/TLS site you must use a method that allows a client to maintain its session with the same SSL/TLS server. This requirement can be met when using Network Load Balancing in Windows Server 2003 by configuring affinity as Single or Class C. The Single option specifies that Network Load Balancing direct multiple requests from the same client IP address to the same cluster host. Class C affinity specifies that Network Load Balancing direct multiple requests from the same TCP/IP Class C address range to the same cluster host.

Most load balancing solutions provide similar functionality. For more information about Network Load Balancing, see The Network Load Balancing Technical Overview link on the Web Resources page at http://go.microsoft.com/fwlink/?LinkID=291.

For example, you might have an SSL/TLS site that runs on a multiprocessor server and can process only 100 transactions per second (Tx/Sec), but you require a processing rate of 300 Tx/sec. To meet this requirement, you could create a Network Load Balancing cluster using two additional servers of the same configuration, and configure the cluster for Single affinity or Class C affinity. In theory you would have increased the performance of your site to 300 Tx/sec. In fact, you would also need to take into account some slight differences in the configuration and possible overhead from using Network Load Balancing. However, it is clearly possible to scale out an SSL/TLS site so long as you observe the requirements imposed by SSL/TLS.

Another technique for increasing performance of an SSL/TLS site is *scaling up*. Scaling up refers to adding larger servers, adding additional processors, or adding special SSL/TLS accelerators to your environment. Scaling up will allow more transactions to be performed by the increase hardware capacity. There are no special considerations for SSL/TLS when scaling up a server.

**Performance Tuning Parameters for SSL/TLS**

The administrator of a Windows Server 2003–based system can tune several Schannel parameters to control session reconnects. These parameters can be used to increase performance or save memory in some application scenarios.

**The Schannel Cache**

SSL/TLS supports a reconnect operation that can be used to enable a client and server to resume a previously negotiated session. This is desirable in many cases because the SSL/TLS reconnect does not require the time needed for a full RSA handshake. But there are cases in which reconnects can be troublesome, such as during performance testing or when the connection pattern of the site is that the same client never reconnects to the same Web server.

Reconnects are enabled by the Schannel cache, which keeps a list of previous SSL/TLS sessions that are established with the current credential handle. The sessions are referenced by the SSL/TLS session ID. Schannel currently maintains up to ten thousand cached sessions for a maximum of ten hours.

Schannel cache values are stored in the following registry subkey:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL

You can add or modify these values by using the registry editor Regedit.exe. Table 2 describes the entries and their values.

| ⚠Caution |
|---|
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Resource Kit Registry Reference for Windows Server 2003 at http://go.microsoft.com/fwlink/?LinkID=4543. |

**Table 2. Schannel Registry Settings**

| Entry | Datatype | Description |
|---|---|---|
| MaximumCacheSize | REG_DWORD | The maximum number of SSL/TLS sessions to maintain in the cache. The default value is 10,000 |

| | | |
|---|---|---|
| ClientCacheTime | REG_DWORD | The time, in milliseconds, to expire each client side cache element. The default is 10 hours. |
| ServerCacheTime | REG_DWORD | The time, in milliseconds, to expire each server side cache element. The default is 10 hours. |

Setting either MaximumCacheSize or ServerCacheTime to zero disables the server-side session cache and prevents reconnects. Increasing MaximumCacheSize or ServerCacheTime above the default values causes LSASS.EXE to consume additional memory. Each session cache element typically requires 2-4k bytes of memory.

**Additional Reading**

See the following resources for further information:

RFC 2246: *The TLS Protocol Version 1.0*.

ITU-T Recommendation X.509, available on the International Telecommunications Union Web site at http://go.microsoft.com/fwlink/?LinkId=3799.

# Choosing the Right Printer Drivers

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

**In This White Paper**


Introduction (Choosing the Right Printer Drivers)

User-mode vs. Kernel-mode Drivers

Unidrv and PostScript-based Drivers

Designed for Windows Drivers

Alternate Options

Summary (Choosing the Right Printer Drivers)

Related Links (Choosing the Right Printer Drivers)

# Introduction (Choosing the Right Printer Drivers)

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

The Microsoft Windows printing architecture consists of a print spooler and a set of printer drivers. By calling device-independent Win32 printing and Graphics Device Interface (GDI) functions, applications can create print jobs and send them to any of a variety of devices, including laser printers, vector plotters, raster printers, and fax machines. Printer drivers include a user interface (UI) component that allows users to control a printer's selectable options.

An application's calls to Win32 GDI functions are passed to the GDI graphics engine, which either spools the drawing instructions as an enhanced metafile (EMF) or, in conjunction with a printer driver, renders a printable image that can be sent to the spooler. Spooler components interpret EMF files, and they can insert page layout information and job control instructions into the data stream. The spooler then sends the data stream to the serial, parallel, USB, or network port driver associated with the target printer's I/O port. The device interprets the job, formatted in Printer Description Language (PDL) for the device, and outputs the printed document.

The printer drivers that you select to run on your Windows 2003 print server can have a great impact on its performance, reliability, and scalability.

Printer performance and reliability can be improved by selecting robust, well-tested drivers. These drivers are less likely to crash and therefore require less maintenance by the print administrator.

Scalability is affected by the quantity of unique drivers on a specific server. By reusing the same driver for multiple printers, you lower the chances that drivers will interact in a negative way and reduce the amount of spooler process and registry space used.

The following types of drivers, explained in detail later in this document, are recommended for achieving the optimal performance, reliability, and scalability goals:

- User-mode drivers
- Unidrv-based drivers
- PostScript drivers
- Drivers with the Designed for Windows logo

The following driver types should be avoided unless there are no other options:

- Kernel-mode or Microsoft Windows NT 4.0 drivers
- Monolithic drivers (that is, neither a Unidrv nor PostScript-based driver)
- Drivers without the Designed for Windows logo (unsigned drivers)

# User-mode vs. Kernel-mode Drivers

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

Drivers can be written in either user mode (also called version 3 drivers) or kernel mode (also called version 2 drivers). In Windows NT 4.0, drivers were moved into kernel mode to improve performance. However, when a kernel-mode driver fails, it can crash an entire system, whereas the failure of a user-mode driver causes only the current process to crash. Because of this difference, and because performance enhancements were made in Windows 2000 and Windows Server 2003, native drivers on Windows 2000 and later run in user mode. Windows Server 2003 can still run kernel-mode drivers, although this is not recommended for the stability reasons mentioned previously. It is also possible to install kernel-mode drivers using the Additional Drivers option on Windows 2003 print servers to support Windows NT 4.0 clients.

All drivers on the Windows Server 2003 CD are user mode, except those in the printers\nt4\i386 directory, which are kernel mode to support Windows NT 4.0 clients.

To check whether a driver that you have installed is user mode or kernel mode, do the following:

1. Click **Start**, and then choose **Printers Folder**.
2. Click **File**, and then click **Server Properties**.
3. Click the **Drivers** tab.
4. Look at the **Version** column for a specific driver. If the version indicates Windows NT 4.0 you have a kernel-mode driver. If the version is Windows 2000, Windows XP and Windows Server 2003, you have a user-mode driver.

| **Note** |
| --- |
| The Server Properties dialog box contains information and settings about the entire print server. |

# Unidrv and PostScript-based Drivers

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

Unidrv and Postscript are the standard, well-tested core rendering engines included in Windows Server 2003. In addition, there is a core plotter driver in Windows Server 2003 which a few drivers use. All drivers included on the Windows Server 2003 CD are based on the PostScript, plotter or Unidrv-based drivers, the majority of which were written in conjunction with Independent Hardware Vendors (IHVs) and are also available on the IHVs independent Web sites. These full-featured drivers are supported by both Microsoft and the IHVs.

Because approximately 3,800 drivers based on Unidrv or PostScript are included on the Windows Server 2003 CD and have been thoroughly tested by Microsoft, the IHVs, and multiple customers, these core engines have been proven to be reliable and stable. Drivers based on these engines are less likely to have problems than those based on monolithic engines.

To determine whether a driver you have installed is using the Unidrv or Postscript rendering engines, do the following:

1. Click **Start**, and then choose **Printers Folder**.
2. Click **File**, and then click **Server Properties**.
3. Click the **Drivers** tab.
4. Select a driver, and then click **Properties**.
5. Look at the **Driver File** list box of in the Driver Properties dialog box. If UNIDRV.DLL or PSCRIPT5.DLL appears, the driver is using one of these rendering engines.

| **Note** |
| --- |
| The Postscript rendering engine was co-written by Microsoft and Adobe Systems. You can find the same PostScript engine on the Adobe Web site (http://www.adobe.com) as is on the Windows Server 2003 CD. PostScript 5.1 was included in Windows 2000. PostScript 5.2 is included in Windows 2000 SP3 and Windows Server 2003. |

# Designed for Windows Drivers

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

The "Designed for Windows" logo identifies hardware and software products that have been designed and tested to meet Microsoft standards for compatibility with the Windows Server 2003 operating system. Products displaying the logo have passed rigorous testing that assures they are easy to use, stable, and take advantage of the new features of Windows Server 2003. Hardware is tested to specifications set by the Windows Hardware Quality Labs (WHQL).

Businesses that use products meeting the Designed for Windows logo criteria stand to gain the following benefits:

- **Lower support costs:** Drivers with the logo follow standard Windows look-and-feel guidelines. A consistent UI helps users get up to speed without phoning the help desk.
- **Help manage "DLL Hell":** Drivers with the logo do not overwrite other drivers or uninstall key components.
- **Proper use of operating system:** Products with the logo make proper use of the Windows registry and other key operating system files.
- **Compliance with the Americans with Disabilities Act and other equal rights legislation:** Products with the logo meet usability standards for a wide range of people.

A list of all drivers that currently have the Designed for Windows logo is available at http://www.microsoft.com/whdc/hcl/default.mspx.

| **Note** |
| --- |
| Most but not all drivers with the Designed for Windows logo are user-mode drivers. All drivers recently granted the Designed for Windows Server 2003 logo are user-mode drivers. However, some kernel-mode drivers with proven reliability were grandfathered in during the initial stages of the logo process during the release of Windows 2000. |

# Alternate Options

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

If you are unable to find user-mode drivers with the Designed for Windows logo that are based on the Unidriver or PostScript engines, try using a compatible driver for your printer that supports a majority of the important features. For more information, see the Print Server Upgrade, Migration, and Interoperability white paper at http://www.microsoft.com/windows2000/technologies/fileandprint/print/default.asp.

# Summary (Choosing the Right Printer Drivers)

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

As described in the preceding sections, the drivers chosen to run on your print server can affect the stability, reliability, and performance of the server. For best performance, look for user-mode drivers with the Designed for Windows logo that are Unidrv or Postscript based drivers.

# Related Links (Choosing the Right Printer Drivers)

Updated: March 1, 2003

Applies To: Windows Server 2003 with SP1

See the following resources for further information:

- For more information about deployment and management of Windows print servers, see the Windows Server 2003 Deployment Kit.
- For information about the point-and-print architecture, see the Windows Point and Print Technical Overview white paper.
- For more information about the Windows .NET architecture, see the Display and Print Devices sections in the Windows Driver Development Kit.
- For more information about Windows Print Services in general, visit the Print Services Technology Center.
- For the latest information about Windows Server 2003, see the Windows Server 2003 Web site.

# GPMC Scripting: Automate GPO management tasks

Updated: August 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By Darren Mar-Elia

Learn how to automate many common GPO management tasks with scripting interfaces included in the Group Policy Management Console (GPMC). This article is from the August 2003 issue of Windows & .NET Magazine.

**In This White Paper**

Getting Started with GPMC Scripting

Retrieving Permissions for a GPO

Obtaining RsoP Reports

GPMC Opens New Possibilities

In April, Microsoft released the Group Policy Management Console (GPMC), which provides a Microsoft Management Console (MMC)—based UI for easy management of Windows Server 2003 and Windows 2000 Group Policy. The GPMC represents a big step forward in Group Policy object (GPO) management capabilities as compared with Windows 2000's native tools. With the native tools, scripting GPO management is difficult. However, GPMC includes a set of scripting interfaces for automating many common GPO management tasks. Using these scripting interfaces, you can manage the Group Policy environment, including generating reports of GPO settings, creating and copying GPOs, and finding unlinked GPOs. Microsoft provides several GPMC scripts that cover many common scripting tasks. You can also create your own scripts to perform custom GPO management tasks.

Although you can manage Windows 2000 domain-based Group Policies, GPMC runs only on Windows 2003 and Windows XP Professional computers. (For more information about GPMC's requirements and features, see "Windows Server 2003's Group Policy Management Console," July 2003, http://www.windowsitpro.com, InstantDoc ID 39190.) You can download the GPMC from the Microsoft Download Center (http://www.microsoft.com/downloads/details.aspx?FamilyId=0A6D4C24-8CBD-4B35-9272-DD3CBFC81887&displaylang=en). When you install the GPMC, the system creates a folder called Scripts, which contains all the prewritten GPMC scripts. On a Windows 2003 or XP client, this folder is in the %programfiles%\gpmc directory. The main administrative scripts have a .wsf extension, which is one of the file formats associated with Windows Script Host (WSH). Scripts with the .wsf extension are XML-formatted files that can call other scripts written in VBScript or JScript, which means that one script can take advantage of both the VBScript and JScript scripting engines. For the scripts in this article, I use VBScript without relying on .wsf files.

The GPMC interfaces are implemented in gpmgmt.dll, which resides in the %programfiles%\gpmc directory. Microsoft geared these interfaces toward automating the GPMC functions as well as managing GPOs. Thus, you can use the interfaces not only to script GPMC operations such as creating mapping tables for GPO migrations but also to query and modify GPOs. However, the GPMC interfaces don't let you read or configure policy settings within a GPO. For example, you can't create a script that enables the *Remove Run from Start Menu* Administrative Template policy within a GPO. This limitation is unfortunate; nonetheless, the GPMC interfaces still provide a level of automation that surpasses what has been available to date. Let's take a look at how to get started with GPMC scripting and how you can use the GPMC objects to perform various administrative tasks, such as retrieving permissions for a GPO and obtaining Resultant Set of Policies (RSoP) reports.

# Getting Started with GPMC Scripting

Updated: August 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Learning to write GPMC scripts is fairly straightforward. All GPMC scripts that you write will follow the same basic steps. As with most new objects that you use in the WSH environment, you first need to create instances of, or instantiate, the objects you want to use. In all GPMC scripts, the first object you need to instantiate is the GPM object. This object is the root object in the GPMC object model. You need the GPM object to access other GPMC interfaces, which then provide access to further capabilities. For example, you need the GPM object to access the IGPMDomain interface, which lets you create a reference to an Active Directory (AD) domain. After you have the reference to the AD domain, you can call IGPMDomain's GetGPO method to access the IGPMGPO interface and create a reference to a particular GPO that you want to manage. From here, the IGPMGPO interface contains methods and properties for managing that GPO. You can learn more about the GPMC object model in the Help file called gpmc.chm, which is in the Scripts folder. You can also learn about the object model in the Microsoft Developer Network's (MSDN's) Group Policy Management Console Reference at http://msdn.microsoft.com/library/en-us/gpmc/gpmc/group_policy_management_console_reference.asp.

Another interface that you're likely to run across is IGPMConstants, which is a special interface in GPMC scripting. This interface provides a set of properties that represent GPO-related constants that you'll often need in your GPMC scripts. For example, suppose you need to set the permissions to control who can edit a GPO. You could use a complex set of file-system and AD ACLs to represent the Edit permission, but that approach would take a lot of code work. So, Microsoft provided the IGPMConstants interface to do the work for you. You can simply call the IGPMConstants interface's PermGPOEdit property to represent the appropriate permission. To access the IGPMConstants interface, you use the GPM object's GetConstants method. After using the GetConstants method to obtain a reference to the GPMConstants object, you can then use any of the GPMConstants properties within your scripts.

Listing 1 shows the code that you use to create the GPM and GPMConstants objects. Let's look at two sample scripts—GetGPOPerms.vbs and RSoPLogging.vbs—to see how you build on this code. Admittedly, these scripts aren't the most basic. I didn't want to duplicate the GPMC scripts that Microsoft provides.

**Listing 1 Code That Creates the GPM and GPMConstants Objects**

```
BEGIN COMMENT
' Code that creates the GPM object.
END COMMENT
Set GPMC = CreateObject("GPMgmt.GPM")
BEGIN COMMENT
' Code that creates the GPMConstants object.
END COMMENT
Set Constants = GPMC.GetConstants()
```

# Retrieving Permissions for a GPO

Updated: August 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The script in Listing 2, GetGPOPerms.vbs, demonstrates how to use several GPMC objects to list the permissions for a GPO in a test domain. GetGPOPerms.vbs begins by instantiating the GPM and GPMConstants objects. Next, the script accesses IGPMDomain, a useful interface that lets you retrieve information about a domain and manage GPOs within it. To access IGPMDomain, you use the GPM object's GetDomain method, which returns a GPMDomain object. As callout A in Listing 2 shows, the GetDomain method takes three arguments. The first argument is the name of the domain that stores the GPOs you want to manage. The domain name must be the domain's DNS name (e.g., mycompany.net). As callout A shows, you can hard-code this argument's value in the script. Another approach is to have the scripts' users provide the domain name at the command line when they launch the script.

The second argument lets you specify which domain controller (DC) you want to use to connect to the domain. A null string ("") signifies that you don't have a DC preference, which means that the GetDomain method will use the PDC emulator. For the third argument, you specify the option you want to use to find a DC with which to connect. You have three options: GPM_USE_ANYDC, (use any available DC), GPM_USE_PDC (use the PDC emulator DC), or GPM_DONOTUSE_W2KDC (use a DC running Windows 2003). As callout A shows, GetGPOPerms.vbs uses the GPMConstants object's UseAnyDC property to specify the GPM_USE_ANYDC option.

After you connect to the domain, the real fun starts. As the code at callout B in Listing 2 shows, you use the GPMDomain object's GetGPO method to retrieve the GPMGPO object that represents the GPO for which you want to list the permissions. To use GetGPOPerms.vbs, you need to replace the domain name mycompany.net with the DNS name of your AD domain. Notice that the GetGPO method's argument is the GPO's globally unique identifier (GUID) and not the GPO's friendly name. In this script, I included the GUID for the Default Domain Policy that's present in every AD domain. This GUID is the same for all AD domains.

If the scripts' users will be providing the necessary GPO information at the command line, having them input the hard-to-type GUID might not be an option. An alternative is to have the users provide the GPOs' friendly names; you can then obtain the corresponding GUID by using the GetGPObyName function, which Microsoft provides in the lib_commongpmcfunctions.js file in the Scripts folder. GetGPObyName uses the IGPMSearchCriteria interface to search all the GPOs in a domain and, on matching the friendly name entered, returns a GUID that the script can pass to GetGPO. However, GetGPObyName is a JScript function. If you'd rather use VBScript, you can write a VBScript version of GetGPObyName or use the GetGPOs, GPOName, and GPOGuid methods in the IADsTools COM object, which is part of the Windows 2000 Support Tools. For more information about these methods, see the article "Scripting with IADsTools," April 2003, http://www.windowsitpro.com, InstantDoc ID 38286.

Next, the script uses the GPMGPO object's GetSecurityInfo method to retrieve the permissions for the GPO. The GetSecurityInfo method returns a reference to a GPMSecurityInfo collection object, which the script assigns to the GPOSec variable. The GPMSecurityInfo object contains the set of permissions assigned to the GPO. The script then iterates through the collection and uses the GPMSecurityInfo object's Count property to count and return the number of permission entries in the collection.

To retrieve each permission entry, the script uses the GPMSecurityInfo object's Item property, which returns a reference to a GPMPermission object. After the script assigns this reference to the Ace variable, the script uses the GPMPermission object's Trustee property to access the GPMTrustee object. By calling the GPMTrustee object's TrusteeName property, the script determines the name of the user or group assigned to the current permission, then assigns the name to the PrincipalName variable.

The code at callout C in Listing 2 uses a Select Case statement to determine the security right assigned to that user or group. A GPO can have five different security rights, as defined in the IGPMConstants interface. The Select Case statement contains these five rights.

The first line in the Select Case statement tells the VBScript runtime engine to compare the Ace.Permission value (i.e., the GPMPermission object's Permission property value) to each case. When the Permission property value matches one of the five security rights, the script assigns a user-friendly description of that permission to the Perm variable. Finally, the script uses WSH's WScript.Echo command to output the user's or group's name and permission to the console screen.

**Listing 2 GetGPOPerms.vbs**

```
BEGIN COMMENT
' Create the GPM and GPMConstants objects,
then connect to the domain.
```

```
END COMMENT
Set GPMC = CreateObject("GPMgmt.GPM")
Set Constants = GPMC.GetConstants()
' BEGIN CALLOUT A
Set GPMCDomain = GPMC.GetDomain("mycompany.net", "",
Constants.UseAnyDC)
' END CALLOUT A
' BEGIN CALLOUT B
BEGIN COMMENT
' Create an object for the GPO for which you want to list the
permissions.
END COMMENT
Set MyGPO = GPMCDomain.GetGPO
("{31B2F340-016D-11D2-945F-00C04FB984F9}")
BEGIN COMMENT
' Get the permissions.
END COMMENT
Set GPOSec = MyGPO.GetSecurityInfo()
' END CALLOUT B
For indx=1 to GPOSec.Count
BEGIN COMMENT
   ' Set the ACE to the Ace variable.
END COMMENT
   Set Ace = GPOSec.Item(indx)
BEGIN COMMENT
   ' Find out the username or group name for the ACE.
END COMMENT
   Set UsrorGrp= Ace.Trustee
   PrincipalName=UsrorGrp.TrusteeName
' BEGIN CALLOUT C
BEGIN COMMENT
   ' Find out which permission the user or group has.
END COMMENT
   Select Case Ace.Permission
      Case Constants.permGPOApply
         Perm="Read and Apply Group Policy"
      Case Constants.permGPOEdit
         Perm="Edit Group Policy"
      Case Constants.permGPOEditSecurityAndDelete
         Perm="Edit Group Policy, Modify Security and Delete
Group Policy"
      Case Constants.permGPORead
         Perm="Read Group Policy"
      Case Constants.permGPOCustom
         Perm="Custom Permission"
   End Select
' END CALLOUT C
   WScript.Echo "The User or Group: " & PrincipalName & _
      " has the following permission: " & Perm
Next
```

# Obtaining RsoP Reports

Updated: August 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A useful feature of the GPMC is its ability to perform Group Policy logging and Group Policy planning. Using the GPMC interfaces, you can programmatically obtain the results from Group Policy—logging and Group Policy—planning sessions. For example, to obtain the results from a Group Policy—logging session, you need to use RSoP Windows Management Instrumentation (WMI) providers.

RSoPLogging.vbs, which Listing 3 shows, demonstrates how you can use the RSoP interfaces to execute a logging query and create a logging report in HTML format. The first few lines in the script create the GPM and GPMConstants objects. Next, the script uses the GPM object's GetRSOP method to create an instance of the GPMRSOP object. This method takes three parameters, the first of which specifies the RSoP mode. As the code at callout A in Listing 3 shows, one way you can provide this mode is to use the GPMConstants object's RSOPModeLogging property. If you were performing an RSoP planning session, you would use the RSOPModePlanning property instead. The second parameter specifies the path to the WMI namespace in which previous RSoP data reside. In this case, the parameter is a null string because no previous data exists. The last parameter is always 0.

After creating an instance of the GPMRSOP object, the script sets two properties—GPMRSOP object's LoggingComputer and LoggingUser properties—for the RSoP logging query. The LoggingComputer property specifies the name of the target machine (in this case, myworkstation), whereas the LoggingUser property specifies the name of the target user (in this case, Darren). Next, the script executes the logging query by calling the GPMRSOP object's CreateQueryResults method, which has no parameters.

Finally, the script calls the GPMRSOP object's GenerateReportToFile method, which takes two parameters. The first parameter specifies the type of report to generate (HTML or XML). The script uses the Constants object's ReportHTML property to specify an HTML report. If you prefer to receive an XML report, you can use the ReportXML property instead of the ReportHTML property. The second parameter specifies the pathname for the report.

The GenerateReportToFile method can return a reference to the GPMResult object. The GPMResult object has two properties—Result and Status—that you can use to determine when the report has finished running or failed to run successfully. However, in RSoPLogging.vbs, the generation of the report is the last task, so you don't need to know when the report is done. (You'll know the report is done when the script finishes executing.) Thus, the script doesn't store the reference to GPMResult.

**Listing 3 RSoPLogging.vbs**

```
BEGIN COMMENT
' Create the GPM and GPMConstants objects.
END COMMENT
Set GPMC = CreateObject("GPMgmt.GPM")
Set Constants = GPMC.GetConstants()
` BEGIN CALLOUT A
BEGIN COMMENT
' Create a reference to an RSoP object.
END COMMENT
Set RSOP= GPMC.GetRSOP(Constants.RSOPModeLogging,"",0)
BEGIN COMMENT
` END CALLOUT A
' Set the RSoP logging session’s properties.
END COMMENT
RSOP.LoggingComputer="myworkstation"
RSOP.LoggingUser="darren"
BEGIN COMMENT
' Execute the RSoP logging query and send
the results to an HTML file.
END COMMENT
RSOP.CreateQueryResults()
RSOP.GenerateReportToFile Constants.ReportHTML,
"c:\reports\myrsop.html"
```

# GPMC Opens New Possibilities

Updated: August 1, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The new GPMC interfaces are extremely flexible, powerful, and fairly well documented by Microsoft. They provide a lot more control over the Group Policy infrastructure than the Windows 2000 native tools. If you create custom scripts and use them in conjunction with the scripts that Microsoft provides, you can automate most GPO management tasks.

© 2003 *Windows & .NET Magazine*. All rights reserved.

We at Microsoft Corporation hope that the information in this work is valuable to you. Your use of the information contained in this work, however, is at your sole risk. All information in this work is provided "as -is", without any warranty, whether express or implied, of its accuracy, completeness, fitness for a particular purpose, title or non-infringement, and none of the third-party products or information mentioned in the work are authored, recommended, supported or guaranteed by Microsoft Corporation. Microsoft Corporation shall not be liable for any damages you may sustain by using this information, whether direct, indirect, special, incidental or consequential, even if it has been advised of the possibility of such damages. All prices for products mentioned in this document are subject to change without notice.

# Advanced Certificate Enrollment and Management

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

Complex infrastructure environments and branch-office deployment environments often dictate unique and advanced management techniques for managing a public key infrastructure (PKI) or certificate deployment to remote servers. This white paper explains several remote deployment scenarios along with the step-by-step procedures to perform X.509 certificate enrollment to implement a secure infrastructure.

**In This White Paper**

Introduction (Advanced Certificate Enrollment and Management)

Requesting Offline Domain Controller Certificates (Advanced Certificate Enrollment and Management)

Processing Domain Controller Certificates

Domain Controller Certificate Installation

Removing Domain Controller Certificates

Troubleshooting (Advanced Certificate Enrollment and Management)

Appendix 1: Identifying a Domain Controller GUID

Appendix 2: Sample Scripts

Appendix 3: Certreq.exe Syntax

Appendix 4: Certutil -setextension

Appendix 5: ASN.1 File Structure

Appendix 6: Encoding and Decoding with Hexadecimal, Binary, and Base64

Summary (Advanced Certificate Enrollment and Management)

Related Links (Advanced Certificate Enrollment and Management)

# Introduction (Advanced Certificate Enrollment and Management)

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

Complex infrastructure environments and branch-office deployment environments often dictate unique and advanced management techniques to manage a PKI or certificate deployment to remote servers. Network and infrastructure services, such as domain controllers, Internet Authentication Servers (IAS), Internet Information Services (IIS), and other stand-alone applications, often require X.509 certificate enrollment or provisioning to provide or enable secure protocols, messaging, or application services. In many of these deployment scenarios, automatic system provisioning or even traditional certificate enrollment may not be possible due to one or more reasons, such as:

- Stand-alone servers with no relationship to an Active Directory® domain
- Firewalls blocking required communication ports
- Complete lack of connectivity without a virtual private network (VPN) or Internet Protocol Security (IPSec) certificate credential to authenticate to the master network

For example, a branch office domain controller may be connected to the central site only through a firewall, and only port 25 is open for Simple Mail Transfer Protocol (SMTP) replication. The domain controller cannot automatically or manually enroll a domain controller certificate over Remote Procedure Call/Distributed Component Object Model (RPC/DCOM) and, therefore, SMTP replication will fail. Because the certification authority is located in the central site and the firewall is blocking RPC traffic, the branch office domain controller cannot contact the certification authority to enroll its certificate. In this situation, a domain controller certificate must be requested, processed, and installed in an asynchronous or offline process. This white paper explains several remote deployment scenarios along with the step-by-step procedures to perform X.509 certificate enrollment to implement a secure infrastructure.

**Supported Scenarios**
Although it is preferable to describe and support all environments, it is not always possible due to technology limitations or documentation requirements. This white paper documents and focuses on domain controller certificate enrollment for Windows 2000 and Windows Server™ 2003 domain controllers from a Windows 2000 or Windows Server 2003 stand-alone certification authority (CA) as well as from a Windows Server 2003 enterprise CA. Because of technical constraints, manual certificate enrollment from a Windows 2000 enterprise CA is not covered in detail.

**Assumptions**
In this white paper, it is assumed that all domain controllers have been configured and function properly. It is also assumed that the CA has been implemented according to the recommended best practices by Microsoft® as documented in the white paper at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws3pkibp.mspx

**Process Overview**
To provide a high-level overview of an advanced enrollment scenario, the following diagrams illustrate step-by-step references of the various procedures and processes that are described in detail in later sections.


Domain Controller Offline Processing with a Windows 2000 Stand-alone CA
Domain Controller
Certification Authority
Start
Is the Domain Controller Windows 2000?
Yes
Copy certutil.exe and certreq.exe to a new folder on the Domain Controller.
No
Determine the Domain Controller's GUID and FQDN.
Prepare the ASN.1 string with certificate extensions.
Prepare the INF file.
Run certreq -new to generate keys and build the request.
Verify the certificate request with certutil -dump.
Copy the request to a removable storage.
Is this the first offline certificate request?
No
2
Yes
1
Accept and install the certificate on the Domain Controller.
3
Is this an offline Directory E-mail Replication certificate?
Yes
Publish the certificate in Active Directory.
Verify the published certificate.
No
Finish
Set the registry flag to allow the CA to accept subject alternate names and restart the CA service.
1
Make removable storage available to the CA.
2
Are extensions part of the request?
No
Confirm that the CA policy module sets the request status to pending.
Yes
Submit the certificate request to the CA.
Are extensions part of the request?
No
Set extensions with certutil -setextensions.
Verify certificate extensions in the pending certificate request.
Issue the certificate.
Yes
Retrieve the certificate and store it on a removable storage.
3


Domain Controller Offline Processing with a Windows 2003 Stand-alone CA
Domain Controller
Certification Authority
Start
Is the Domain Controller Windows 2000?
Yes
No
Copy certutil.exe and certreq.exe to a new folder on the Domain Controller.
Determine the Domain Controller's GUID and FQDN.
Prepare the INF file.
Run certreq -new to generate keys and build the certificate request.
Verify the certificate request with certutil -dump.
Copy the request to a removable storage.
Is this the first offline certificate request?
No
Yes
2
1
Accept and install the certificate on the Domain Controller.
3
Is this an offline Directory E-mail Replication certificate?
Yes
No
Publish the certificate in Active Directory.
Verify the published certificate.
Finish
Set the registry flag to allow the CA to accept subject alternate names and restart the CA service.
1
Make removable storage available to the CA.
2
Are extensions part of the request?
No
Yes
Confirm that the CA policy module sets request status to pending.
Submit the certificate request to the CA.
Are extensions part of the request?
No
Yes
Set extensions with certutil -setextensions.
Verify certificate extensions in the pending certificate request.
Issue the certificate.
Retrieve the certificate and store the certificate on a removable storage.
3


Domain Controller Offline Processing with a Windows Server 2003 Enterprise CA
Domain Controller
Certification Authority
Start
Is the Domain Controller Windows 2000?
Yes
Copy certutil.exe and certreq.exe to a new folder on the Domain Controller.
No
Determine the Domain Controller's GUID and FQDN.
Prepare the INF file.
Run certreq -new to generate keys and build the request.
Verify the request with certutil -dump.
Copy the request to a removable storage.
Is this the first offline certificate request?
No
2
Yes
1
Install the certificate on the Domain Controller.
3
Is this an offline Directory E-mail Replication certificate?
Yes
Publish the certificate in Active Directory.
Verify the published certificate.
No
Finish
Set the registry flag to allow the CA to accept subject alternate names and restart the CA service.
1
Duplicate default certificate templates, Offline Directory E-mail Replication, and Domain Controller Authentication.
Configure templates in which the certificate subject and subject alternate name are supplied in the request.
Is the CA on Windows Server 2003 SP1 or later?
no
Run fixdctemplate.vbs.
yes
Attach certificate templates to the CA.
Make removable storage available to the CA.
2
Are extensions part of the request?
No
Confirm that the template sets the request status to pending.
Yes
Submit the certificate request to the CA.
Are extensions part of the request?
No
Set extensions with certutil -setextensions.
Verify certificate extensions in the pending certificate request.
Yes
Issue the certificate.
Retrieve the certificate and store the certificate on a removable storage.
3

# Requesting Offline Domain Controller Certificates (Advanced Certificate Enrollment and Management)

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

In branch-office scenarios with limited connectivity or those requiring a VPN or IPSec connection to the primary data center will often require an asynchronous enrollment process to enroll or provision the domain controller with an X.509 certificate. Asynchronous or (offline) certificate requests for domain controllers must be created with command-line tools because the operating system does not natively support an offline enrollment Wizard such as the one found in IIS. Since any type of certificate request can be created with the command-line tools, you can also build other certificate request types with these tools once you understand the general procedures. The following sections provide the steps required to enroll a domain controller for a certificate through an offline process.

| **Note** |
| --- |
| Some differences and capabilities exist between a Windows 2000 system and a Windows Server 2003 family operating system. The differences and required procedures specific to an operating system family are noted accordingly. |

**Preparing a Windows 2000 Domain Controller**

For a Windows 2000 family operating system, the first step is to prepare the Windows 2000 domain controller. Since Windows 2000 has no natively installed tools to create certificate requests at a command line, a Windows Server 2003 tool must be used. The version of the certreq.exe and certutil.exe command-line tools available on Windows 2000 have only limited capabilities and do not satisfy the requirements for offline certificate request processing, which is required for this scenario. For example, the Windows 2000 version of certreq.exe does not support the –new option which is required to create new certificate requests. Certutil.exe has many more options to verify and process certificates.

The Windows Server 2003 Administration Tools Pack includes the latest version of certreq.exe and certutil.exe and is required as part of this process. To install the tools on a Windows 2000 computer, you must first install the Windows Server 2003 Administration Tools Pack on a Windows XP or Windows Server 2003 system because it cannot be installed directly on a Windows 2000 computer.

The Windows Server 2003 Administration Tools Pack can be downloaded from the Web site at http://www.microsoft.com/downloads/details.aspx?FamilyID=c16ae515-c8f4-47ef-a1e4-a8dcbacff8e3&DisplayLang=en

To install the tools on a Windows 2000 system, perform the following steps.

| **Important** |
| --- |
| Once you have copied the new files to the Windows 2000 domain controller, two versions of these files will reside on the Windows 2000 computer. Do not remove the natively installed files on the Windows 2000 system because other applications, like the Certificates MMC Snap-in, depend on these native versions. Also, do not register certcli.dll and certadm.dll on the Windows 2000 system through the regsvr32.exe command. |

1. Log on to a computer that runs Windows XP Service Pack 1 or Windows Server 2003.
2. Install the Windows Server 2003 Administration Tools Pack.
3. Copy the following files to a removable storage medium such as a diskette.
    - certreq.exe
    - certutil.exe
    - certcli.dll
    - certadm.dll

4. Log off the computer.

5. Make the removable storage medium available to the domain controller that requires an offline domain controller certificate.

6. Log on to the domain controller as a domain administrator and create a new directory to store the files, for example, *%HOMEDRIVE%\W2K3AdmPak*.

| **Note** |
| --- |
| Do not include this path in your system search path to avoid conflicts with the tool versions that already exist on your computer. |

7. From the removable storage medium, copy the four files to the newly created directory.

**Generating an Offline Certificate Request**

This section applies to both Windows 2000 and Windows Server 2003 domain controllers. The next step is to generate an offline certificate request for the domain controller. As mentioned previously, you must create an offline certificate request with a command-line tool called certreq.exe. This tool supports a rich set of command-line options, but only a few options are required in this procedure. For more information about the certreq.exe tool and its syntax, see Appendix 3: Certreq.exe Syntax.

Certreq.exe requires a text (instruction) file to generate an appropriate X.509 certificate request for a domain controller. You can create the file with your preferred (ASCII) text editor and save the file with an *.inf extension to any directory on your hard drive.

In general, a domain controller certificate that is to be used for SMTP replication must meet the following requirements.

- The certificate must contain the "Certificate Template Name" (referenced by the object identifier 1.3.6.1.4.1.311.20.2) extension.

- The certificate must contain a "Subject Alternative Name" extension that includes the Globally Unique Identifier (GUID) of the domain controller account and the fully qualified domain name (FQDN) of the domain controller Domain Name System (DNS) host name. The subject alternative name is uniquely identified with object identifier 2.5.29.17.

Typically, both extensions are automatically inserted in the certificate by an enterprise CA when the certificate is enrolled automatically or manually while connected to a CA. However, when you request the domain controller certificate offline, you must provide these extensions explicitly in the offline request.

There are several different approaches to adding extensions in certificate requests and, ultimately, issued X.509 certificates. For a Windows 2000 CA, you can either include the extension(s) in the INF instruction file, or you can add the extension(s) to a pending certificate request. If you add the extension(s) to a pending certificate request, it is not necessary to also use the INF instruction file to perform the same task.

For a Windows Server 2003 CA, you can either use the same procedures as a Windows 2000 CA, or you can specify certificate extensions using certreq.exe from a command-line when the certificate request is submitted.

| **Important** |
| --- |
| When you submit extensions from the certreq.exe command-line tool, you cannot set the critical flag for the submitted extensions. Therefore, if an extension is required to be marked critical, you must use an INF instruction file as the submission method. |

**Identifying the Domain Controller GUID**

Active Directory replication via SMTP requires the domain controller GUID as an attribute in the domain controller certificate. Thus, you must identify the domain controller GUID if you are going to issue certificates for SMTP replication from a Windows Server 2003 enterprise or stand-alone CA. Since the GUID of the domain controller is stored in Active Directory, there are several ways to read this attribute from the domain controller's computer object. You do not need to be logged on to a domain

controller or the CA to read the GUID from Active Directory; any computer that has read permissions to objects in the domain is sufficient to complete this task. The following is a partial list of tools that can be used to determine the domain controller GUID from Active Directory.

- Adsiedit.exe (available as part of the Windows® support tools)
- Ldp.exe (available as part of the Windows support tools)
- Dsquery.exe (available as part of the Windows support tools)
- Netdiag.exe (available as part of the Windows support tools)
- Replmon.exe (available as part of the Windows support tools)
- Windows Management Instrumentation (WMI)

For more information, see the Microsoft Knowledge Base article "Determining the Server GUID of a Domain Controller" at http://support.microsoft.com/default.aspx?scid=kb;en-us;224544

Since it may be complicated in some scenarios to easily identify a domain controller GUID, a script is provided in Appendix 2 to look up the GUID from Active Directory automatically when the certificate request is generated. This script simplifies the process for administrators. For more information, see Reqdccert.vbs – Generates Domain Controller Certificate Requests.

**Creating an Offline Certificate Request**

Every X.509 certificate that is issued by a CA requires a unique certificate request to initiate the issuance process. If required, you can create several certificate requests as a batch process before submitting them to a CA.

However, to avoid potential errors and confusion, ensure that you are using unique names for your request files. It is necessary to create the certificate request as a member of the local Administrators group because only this group has interactive logon permissions to a domain controller by default. The first step is to define the input information in an INF file.

To create the INF file that will supply input information for the request, perform the following steps.

1. Log on as a member of the local Administrators group on the domain controller that requires a certificate.
2. Make the reqdccert.vbs script, found in Appendix 2, available locally to the domain controller.
3. If you plan to submit the certificate request to a Windows 2000 or Windows Server 2003 stand-alone CA, run the script from a command-line prompt without any additional parameters. In this case, the *Domain Controller* certificate template is assumed and the domain controller's GUID and DNS name are included as a subject alternative name in the issued certificate.

   If you plan to submit the certificate request to a Windows Server 2003 enterprise CA, you must specify the specific name of the certificate template that is used for enrollment. The name depends on the certificate template name that was previously chosen in **Certificate Template Creation**. In addition, you must specify if the certificate is enrolled for either domain controller authentication or for e-mail replication. This is required because the authentication and e-mail replication templates differ in the subject alternative name construction.

   According to your requirement, run one of the following commands at the command-line prompt.

   To create a certificate request for a domain controller authentication certificate with a custom template:

   ```
   cscript reqdccert.vbs <Templatename> A
   ```

   To create a certificate request for a directory e-mail replication certificate with a custom template:

   ```
   cscript reqdccert.vbs <Templatename> E
   ```

   To create a certificate request with the default domain controller template:

```
cscript reqdccert.vbs
```

The script creates an INF file that is used as an input file to build the certificate request. It also creates some additional syntax that contains the appropriate command syntax to submit the request to the CA and verify the certificate.

For a detailed description of the tasks that are performed by the script, see Reqdccert.vbs – Generates Domain Controller Certificate Requests in Appendix 2.

4. Open a command-prompt window. On a Windows 2000 domain controller, type

```
%HOMEDRIVE%\W2K3AdmPak\CERTREQ -new <dcname>.inf <dcname>.req
```

On a Windows Server 2003 domain controller, type

```
CERTREQ -new <dcname>.inf <dcname>.req
```

Replace the <dcname> variable with the hostname of the domain controller. If you are unsure of the name, look up the name of the INF file saved in your current directory. Note that it may take a few seconds or minutes before the command prompt returns. Depending on the size and available cycles of the system CPU as well as the key length that was set in the INF file, it may take several minutes for the key material to be generated.

| **Important** |
| --- |
| The certreq.exe tool performs several tasks that are often not apparent. First, it generates the public and private key material, and then, it generates the actual certificate request. The key material generation process is the reason why certificate requests must not be generated for several domain controllers on a single computer. The key material must be unique per domain controller to ensure a secure solution. |

**Viewing the Certificate Request**

At this stage, you have successfully generated the public and private key material and created the certificate request. The certificate request is stored in a file and is kept in the *Certificate Enrollment Requests* store of the local computer. To find the enrollment request that was generated, perform the following steps.

1. While still logged on as a member of the local Administrators group, start the Microsoft Management console.

2. Add the Certificates MMC Snap-In.

3. Select Computer Account when the Add window asks for the certificate store that should be managed.

4. In the Certificates MMC Snap-In, navigate to **Certificate Enrollment Requests** in the left pane.

   The certificate request appears in the right pane. It is expected that there is no more than one certificate request pending in this container. Otherwise, this would imply that other certificate request attempts have been submitted but have not yet been approved or accepted.

5. When you double-click the certificate request, the message “You have a private key that corresponds to this certificate” is displayed in the General tab at the bottom. As mentioned previously, certreq.exe has generated the certificate request and the key material. You can safely ignore the warning at the top of the window that says the certificate might be altered. This is an expected warning since the certificate has not yet been issued.

**Verifying a Certificate Request**

It is recommended that the offline domain controller certificate request be validated before it is sent to a CA for processing. Verification will ensure that all request fields are set properly and that the certificate, when processed and issued, will provide the expected functionality. To examine the previously generated certificate request, perform the following step.

- While still logged on as a member of the local Administrators group to the domain controller, type certutil –dump <dcname>.req at a command-line prompt, and then press Enter.

Replace the <dcname> option with the name used in the previous section to generate the offline request. If you have properly created a request that has used the DomainController certificate template, the command will display an output similar to the following:

```
PKCS10 Certificate Request:
Version: 1
Subject:
    CN=W2K3-BO-DC.contoso2.com
Public Key Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.1 RSA
    Algorithm Parameters:
    05 00
Public Key Length: 1024 bits
Public Key: UnusedBits = 0
    0000  30 81 89 02 81 81 00 95  b8 e9 18 84 df 95 ce 37
    0010  ce f6 af 32 40 43 14 d5  0f 7b 3b 76 36 8c dd 8b
    0020  7c 03 29 33 26 d3 84 c3  7e ae 25 34 ea 1e db 3a
    0030  b9 01 e3 a7 02 3c 6b 8f  66 99 c8 ac 51 70 03 bc
    0040  47 06 ef 2f 62 3e c3 8d  e1 51 bd 9d c8 7d 95 8c
    0050  08 0a bf 54 a6 f3 1d 2f  cd b8 7d 17 fc 4c 7d a5
    0060  a6 ce 90 d0 a3 21 c5 b0  c1 f0 de ae 00 43 16 cb
    0070  eb 73 01 e7 71 79 ed dd  72 d0 cc 4a 55 26 a2 99
    0080  03 21 dc d1 5b b4 b9 02  03 01 00 01
Request Attributes: 4
  4 attributes:
  Attribute[0]: 1.3.6.1.4.1.311.13.2.3 (Operating System Version)
    Value[0][0]:
        5.2.3790.2
  Attribute[1]: 1.3.6.1.4.1.311.21.20 (Client Information)
    Value[1][0]:
    Unknown Attribute type
    Client Id: = 1
    XECI_XENROLL -- 1
    User: CONTOSO2\administrator
    Machine: W2K3-BO-DC.contoso.com
    Process: certreq
  Attribute[2]: 1.2.840.113549.1.9.14 (Certificate Extensions)
    Value[2][0]:
    Unknown Attribute type
Certificate Extensions: 4
    2.5.29.14: Flags = 0, Length = 16
    Subject Key Identifier
        55 93 fd 45 5b 22 87 33 95 96 4a 77 e3 ff 08 08 f6 83 de fc
    2.5.29.17: Flags = 1(Critical), Length = 3c
    Subject Alternative Name
        Other Name:
             1.3.6.1.4.1.311.25.1= 0410 6661 6135 3636 3234 3831 6263 3866 6662
        DNS Name=W2K3-BO-DC.contoso.com
    2.5.29.37: Flags = 0, Length = 16
    Enhanced Key Usage
        Server Authentication (1.3.6.1.5.5.7.3.1)
        Client Authentication (1.3.6.1.5.5.7.3.2)
    2.5.29.15: Flags = 0, Length = 4
    Key Usage
        Digital Signature, Key Encipherment (a0)
  Attribute[3]: 1.3.6.1.4.1.311.13.2.2 (Enrollment CSP)
    Value[3][0]:
    Unknown Attribute type
    CSP Provider Info
    KeySpec = 1
    Provider = Microsoft RSA SChannel Cryptographic Provider
    Signature: UnusedBits=0
    0000  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
    Remaining 78 bytes are zero
Signature Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.5 sha1RSA
    Algorithm Parameters:
```

```
    05 00
Signature: UnusedBits=0
    0000  45 4c 4f ca 00 52 36 88  a0 2d 2e 45 f3 87 be ac
    0010  89 4f 4e d4 19 42 a7 e5  cb 7a 15 d5 eb e2 a0 96
    0020  6c 39 94 c5 71 d6 e7 03  10 9c 45 a0 ad c7 34 e6
    0030  f9 f2 31 da ce e2 2e 6f  b7 23 6f 12 53 6a 40 89
    0040  ba a9 2e 2f bc cf d4 00  72 18 82 05 33 ea 0e 20
    0050  0b e2 5c 4c 5a 57 b5 71  08 e0 3e d8 8b 0d 18 05
    0060  7b 11 4d b1 e9 db 16 e5  78 e8 c2 b2 ff bb c2 9d
    0070  e6 2b 17 83 dc 1d 43 fd  4e 0e 37 58 f0 ac a9 95
Signature matches Public Key
Key Id Hash(sha1): 55 93 fd 45 5b 22 87 33 95 96 4a 77 e3 ff 08 08 f6 83 de fc
CertUtil: -dump command completed successfully.
```

It should be expected from the output of the certificate request that the majority of attributes from the generated INF file appear in the actual encoded certificate request.

The reqdccert.vbs script creates an INF instruction file that contains the extension information for DNS name and GUID of the domain controller in an encoded format. This enables the certificate requested to be constructed with certreq.exe –new, whereby the extension information will be automatically included in the certificate request from the INF file.

**Deleting a Certificate Request**

Sometimes it may be necessary to either delete an unneeded or unused certificate request or even delete erroneous certificate requests due to mistakes in the INF input file. If you are logged on as a normal user, you can remove certificate requests only from your own certificate store. However, if you are logged on as Administrator, you can also remove certificate requests from the computer's certificate store. To delete a certificate request, perform the following steps.

1. Open the Certificates MMC Snap-In for the computer account (as described previously).
2. Navigate to Certificate Enrollment Requests in the left pane.
3. Remove the pending certificate request in the right pane by pressing Delete on your keyboard.
4. Close the Certificates MMC Snap-In.
5. If the request has already been submitted to a CA, it may be necessary to also deny the pending request on the CA.

The key material that was generated for the certificate request will remain on the system in the local system (computer) profile unless explicitly deleted by an administrator. Orphaned key material may only be removed manually at the command-line prompt using certutil.exe. However, the administrator will need to know the key container.

| **Important** |
| --- |
| If you delete keys manually from your computer, you will invalidate all data that was encrypted with an encryption key. If you have not implemented a key recovery mechanism, it is recommended that you leave unused keys on the system instead of deleting them. |

Use the following command to display the list of available key containers for the machine context.

```
certutil –store my
```

Use the following command for the current user's context.

```
certutil –store –user my
```

The output will display the key container name for each certificate. Key container names typically are created with random GUID strings as the name.

If you have determined the *keycontainername* for a specific certificate, you can delete the key container with the following command.

```
certutil.exe -delkey <KeyContainerName>
```

The -delkey option is supported only with the Windows Server 2003 version of certutil. On Windows 2000, you must add a prefix to the commands. The prefix is the path you have copied the Windows Server 2003 version of certutil to. In this white

paper, the *%HOMEDRIVE%\W2K3AdmPak* path is used.

**Transferring the Request to a CA**

Once the request file has been generated and validated on the domain controller, it can be transferred to the CA for processing and issuance. Transfer the following files to the CA.

- <dcname>-req  This file contains the certificate request that was generated on the domain controller.

- <dcname>-req.bat  This batch file was created by reqdccert.vbs and contains the correct command-line parameters to submit the request to the CA.

At this point, you can log off of your domain controller.

# Processing Domain Controller Certificates

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

Due to differences in the Certificate Services components in Windows 2000 and Windows Server 2003, the issuing process for domain controller certificates depends on the certificate template version and the specific certificate template. Since the processes and steps required are different, both types are outlined separately in the following sections.

**Certificate Templates**
The Windows 2000 and 2003 Server enterprise certification authority (CA) supports the concept of certificate templates. Certificate templates define how an enterprise CA should process a certificate request and generate a specific certificate type when issued. A difference exists between the *Domain Controller* certificate templates in a Windows 2000 and Windows Server 2003 Active Directory environment. It is necessary to understand the differences before you submit the certificate request to a CA. For a broader discussion on certificate templates, see the following references.

- Implementing and Administering Certificate Templates in Windows Server 2003 at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws03crtm.mspx

- Selecting Certificate Templates at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/library/depkit/C71D2CD3-82EF-4E3C-8746-1340D0EF4E9A.mspx

**Windows 2000 Certificate Templates**
A Windows 2000 CA provides only one certificate template for domain controller certificates: *Domain Controller*. It may be manually or automatically enrolled through Automatic Certificate Request Service (ACRS) settings in Group Policy. The certificate template name is hard-coded in the operating system and it must be used to enroll Windows 2000 domain controllers. In Windows 2000, you cannot modify certificate templates and thus, only the *Domain Controller* template may be used.

With a domain controller certificate that was issued by a Windows 2000 CA, a domain controller can use the certificate for any of the following purposes.

- Provide mutual authentication for Public Key Cryptography for Initial Authentication in Kerberos (PKINIT) or smart card logon.

- Encrypt Active Directory replication traffic if SMTP replication is enabled.

- Enable Lightweight Directory Access Protocol (LDAP) over Secure Sockets Layer (SSL) connections to Active Directory.

- Authenticate using SSL client authentication for an application that requires this functionality.

If Windows 2000 style auto-enrollment (ACRS in Group Policy) is enabled in a connected environment, a Windows 2000 domain controller will auto-enroll certificates based on the *Domain Controller* template.

**Windows Server 2003 Certificate Templates**
In contrast to a Windows 2000 enterprise CA, a Windows Server 2003 enterprise CA provides three templates for domain controller certificates. In Windows Server 2003, the *Domain Controller* certificate template is known as a V1 certificate template and exists to support auto-enrollment for Windows 2000 domain controllers. The two other certificate templates are *Directory E-mail Replication* and *Domain Controller Authentication*. Both certificate templates are V2 templates, which implies they can be modified, duplicated, and used for the new style certificate auto-enrollment.

Windows XP and Windows Server 2003 systems that are joined to an Active Directory domain support a more advanced auto-enrollment mechanism than previously available in Windows 2000. To support a Windows 2000 computer in a mixed environment with Windows XP and Windows Server 2003 systems, both methods are supported by a Windows Server 2003 CA. For more information about the new style auto-enrollment, see the white paper at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/plan/autoenro.mspx

The capabilities of the Windows 2000 V1 *Domain Controller* certificate template have been divided into two V2 certificate

templates to provide more specific functionality. Instead of having one multi-purpose domain controller certificate that can be used for almost everything, the *Domain Controller Authentication* certificate template is tailored for smart card logon support; the *Directory Email Replication* certificate is made for supporting SMTP e-mail replication. Thus, if you do not require Active Directory replication via SMTP, you do not have to deploy *Directory Email Replication* certificates.

| **Note** |
| --- |
| Both certificate templates are configured by default to supersede the former *Domain Controller* certificate template, which will be performed when both a Windows 2000 enterprise CA and domain controller are upgraded to Windows Server 2003. For more information about superseding templates, see the Implementing and Administering Certificate Templates in Windows Server 2003 white paper at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws03crtm.mspx |

To summarize the previous points, the following table provides an overview of the certificate templates that are used by default in Windows 2000 and Windows Server 2003.

| **Domain Controller Operating System** | **Windows 2000 Stand-alone CA** | **Windows 2000 Enterprise CA** | **Windows Server 2003 Stand-alone CA** | **Windows Server 2003 Enterprise CA** |
| --- | --- | --- | --- | --- |
| **Windows 2000** | Domain Controller | Domain Controller | Domain Controller | Domain Controller |
| **Windows Server 2003** | Domain Controller | Domain Controller | Domain Controller | • Domain Controller Authentication<br>• Directory Email Replication |

These same templates can also be used for domain controller offline certificate enrollment. In the case of a Windows Server 2003 enterprise CA, it will also be required to customize the *Domain Controller* certificate template. For more information about this process and requirements, see **Certificate Template Configuration**.

**Using Version 2 Templates with Windows 2000 Computers**

In a mixed environment where you may support or have multiple operating system versions installed, your choice of template version may be limited by the client operating system. Only Windows XP client systems are able to enroll for version 2 templates through the Certificates MMC Snap-In. Windows 2000 clients are not able to display or enroll for version 2 templates through the Certificates MMC Snap-In or automatic certificate request settings in Group Policy. Nevertheless, a Windows 2000 computer can be used to enroll for user certificates that are based on version 2 templates through the Windows Server 2003 Web enrollment pages. As an alternative, you could install certreq.exe on a Windows 2000 computer and request certificates based on version 2 templates manually using that method.

**Domain Controller Certificate Details**

Since many of the offline domain controller enrollment processes involve complex, manual procedures, the following section is provided to assist in understanding the most important characteristics of a domain controller certificate as a reference.

| **Certificate Purpose** | **Domain Controller** | **Domain Controller Authentication** | **Directory E-mail Replication** |
| --- | --- | --- | --- |
| Domain controller authentication–A domain controller can prove its identity to another party, such as a client computer, during smart card logon. | Yes | Yes | No |
| SSL–If a Web server is installed on the domain controller, the Web server can leverage the domain controller certificate to establish SSL connections with clients. This also supports LDAP over SSL connections to the domain controller. | Yes | Yes | No |
| Client authentication (SSL)–This would be used if the computer acts as an SSL client to another application on a separate server. | Yes | Yes | No |
| E-mail encryption–SMTP e-mail can be encrypted and signed with this certificate. | Yes | No | Yes |

A certificate that was issued by a CA based on the *Domain Controller* certificate template has the following characteristics.

- The subject contains the domain controllers FQDN prefix with the "CN=" relative distinguished name element.
- The certificate purposes (also known as extended key usage) are set to "Client Authentication (1.3.6.1.5.5.7.3.2)" and "Server Authentication (1.3.6.1.5.5.7.3.1)". The numbers in parentheses are the corresponding standard object identifier for each certificate purpose.
- The common name of the template is set to "DomainController".
- The Subject Alternative Name contains the domain controller's GUID in object identifier 1.3.6.1.4.1.311.25.1 and the FQDN of the domain controller.

A certificate that was issued based on the *Domain Controller Authentication* certificate template has the following characteristics.

- The subject of the certificate is empty.
- The certificate purposes (also known as extended key-usage) are set to "Client Authentication (1.3.6.1.5.5.7.3.2)", "Server Authentication (1.3.6.1.5.5.7.3.1)", and "Smart Card Logon (1.3.6.1.4.1.311.20.2.2)". The numbers in parentheses are the corresponding object identifier for each certificate purpose.
- The common name of the template is set to "Domain Controller Authentication" or the name of the template that was specified in the certificate request for this certificate type.
- The Subject Alternative Name extension contains the domain controller's fully qualified DNS name of the domain controller.

A certificate that was issued based on the *Directory Email Replication* certificate template has the following characteristics.

- The subject of the certificate is empty.
- The certificate purpose (also known as extended key-usage) is set to "Directory Service Email Replication (1.3.6.1.4.1.311.21.19)". The numbers in parentheses are the corresponding object identifier for each certificate purpose.
- The common name of the template is set to "DirectoryEmailReplication" or the name of the template that was specified in the certificate request for this certificate type.
- The Subject Alternative Name extension contains the domain controller's GUID in object identifier 1.3.6.1.4.1.311.25.1 and the FQDN of the domain controller.

**Windows Server 2003 Certificates and Publishing in Active Directory**

Certificates, which are enrolled from an enterprise CA with the *Domain Controller* or the *Directory Email Replication* certificate template, are published by default into the requester's object in Active Directory. When a certificate is auto-enrolled, the requestor is naturally a domain controller. These certificate templates publish certificates in Active Directory primarily to facilitate encrypted replication of Active Directory content with SMTP; both replication partners must have access to the public key (certificate) of their replication partner.

Only enterprise certification authorities publish certificates in Active Directory as an automatic process. If you enroll a *Domain Controller* certificate manually from a stand-alone CA, you must publish the certificate manually in Active Directory. Remember that stand-alone CAs do not offer the *Domain Controller Authentication* and *Directory Email Replication* certificate template format(s), so regardless of the CA version, you can only issue *Domain Controller* certificates from a stand-alone CA.

By default, certificates that have been built with the *Domain Controller Authentication* certificate template are not published in Active Directory because smart card logon requires the domain controller's certificate in Active Directory. Therefore, it is not recommended that *Domain Controller Authentication* certificates be published in Active Directory by manipulating the default *Domain Controller Authentication* certificate template.

**Windows 2000 Server CA Configuration**

By default, a Windows 2000 Server CA does not permit subject alternative names that are specified in a certificate request to

be accepted and inserted in the issued certificate. This applies for both stand-alone and enterprise CAs. This functionality is required to submit and process offline domain controller certificate requests. To permit this functionality, the CA configuration must be modified.

| **Important** |
| --- |
| Changing the CA configuration of the CA to permit subject alternative names in certificate requests will be a global setting and is not limited to a single template. Once this setting is enabled, the CA will accept attributes for subject alternative names for all certificate requests. |

To change the CA configuration to accept subject alternative names in requests, perform the following steps.

1. Log on to the CA as Administrator.

2. At a command-line prompt, type

   ```
   CERTUTIL -setreg policy\EditFlags +EDITF_ATTRIBUTESUBJECTALTNAME2
   ```

   This enables the CA to accept the "Subject Alternative Name" to be submitted in the request. If you submit more offline domain controller requests to this CA, you do not have to set the parameter again. It persists until it is manually reset to the default setting by an administrator. To reset the CA to the default setting, type the following at a command-line prompt.

   ```
   CERTUTIL -setreg policy\EditFlags -EDITF_ATTRIBUTESUBJECTALTNAME2
   ```

   | **Note** |
   | --- |
   | This setting will affect all certificate requests that are submitted to your CA. Any certificate request that provides a SubjectAltName2 will be recognized and processed by the CA. |

3. The following command will display the configuration parameters (EditFlags) that have been set before and after the change. You can verify the parameters explicitly with the following command.

   ```
   CERTUTIL -getreg policy\EditFlags
   ```

4. To enable the changes in all subsequently issued certificates, restart the certification authority service. To restart the CA, type the following at a command prompt, and then press Enter.

   ```
   NET STOP certsvc & NET START certsvc
   ```

**Issuing Domain Controller Certificates with a Windows 2000 CA**
A CA that is installed on a Windows 2000 computer can issue domain controller certificate requests for both Windows 2000 and Windows Server 2003 domain controllers.

With a Windows 2000 CA, no option is available to customize the *Domain Controller* certificate template. However, as a limited option, you may set the validity time of a domain controller certificate on a specific CA. For information about how to change the validity time, see the Microsoft Knowledge Base article *HOW TO: Change the Expiration Date of Certificates That Are Issued by a Windows Server 2003 or a Windows 2000 Server Certification Authority* at http://support.microsoft.com/default.aspx?id=254632

**Windows 2000 Stand-alone CA**
To issue domain controller certificates, it does not matter if the stand-alone CA is a domain member or a member of a workgroup. With both configurations, domain controller certificates can be issued. A stand-alone CA does not support certificate templates, therefore, no choice is required in this regard. A stand-alone CA will process a certificate request according to the information supplied in the request.

To issue a certificate for domain controllers from a Windows 2000 CA, you can use one of two approaches.

- Prepare a certificate request with the Windows Server 2003 version of certreq.exe that includes the domain controllers

GUID and its DNS name in the subject alternative name.

- Use a generic certificate request and add the certificate extensions to the certificate request while it is pending at the CA.

In the first case where you have included the subject alternative name in the certificate request, it is not required to set the certificate request status to pending before issuing the certificate.

Since this white paper illustrates how to include the subject alternative name in the certificate request in **Certificate Template Configuration**, the following section explains how to manipulate a pending request. You may use the following procedure with a different set of command-line parameters for other certificate types as well.

| **Note** |
| --- |
| From a technical viewpoint, both methods result in the same certificate being issued. However, it is more convenient to include extensions in the certificate request because no additional manipulation of a pending request is required. This method, however, does require the Windows Server 2003 version of certreq.exe to be used. |

**Windows 2000 CA Configuration**

Once you have set the *EditFlag* (see Windows 2000 Server CA Configuration), you should verify that the policy module will put all certificates into a pending mode before they are issued. To confirm the request handling of the CA, perform the following steps.

1. Log on to the computer where the Windows Server 2000 stand-alone CA is installed.

2. Open the Certification Authority MMC Snap-In.

3. In the left pane, select the CA object and select Properties on the Action menu.

4. In the Policy Module tab, click Properties.

5. Verify that the policy module of the stand-alone CA is set to the following value.

   *Set the certificate request status to pending*.

   | **Note** |
   | --- |
   | A stand-alone CA sets the certificate request status to pending by default. |

6. Click **OK** twice to confirm the setting and close the Properties page.

7. Close the Certification Authority MMC Snap-In.

**Issuing a Domain Controller Certificate**

Compared to a Windows 2003 CA, it is more complicated to issue certificates that require a given set of subject alternative names in the certificate. Unfortunately, unlike a Windows Server 2003 CA, you cannot specify certificate extensions such as the subject alternative name when you submit the certificate. The Windows 2000 CA does not support such functionality. To work around this limitation, you must include the certificate extensions in the certificate request or perform the certificate enrollment process in three distinct steps, with an optional fourth step.

1. The certificate request is submitted.

2. Certificate extensions are set in the pending certificate request.

3. The certificate is approved and issued.

4. The certificate can be manually verified for accuracy.

The following steps show how to submit the request, process the request, verify the request, and issue the certificate. In this case, the certificate request does not include the certificate extensions.

**Submit the Certificate Request**

1. Log on to the computer where the Windows Server 2000 stand-alone CA is installed.

2. Copy the certificate request (<dcname>.req) and the batch script (<dcname>-req.bat) that were previously created on the domain controller (see Creating an Offline Certificate Request) into a working file folder on the CA.

3. At a command-line prompt, use the <dcname>-req.bat script to run the required certreq command. The batch file will simply submit the certificate request to the CA with the following command.

   ```
   certreq -attrib "CertificateTemplate:DomainController" <dcname.req>
   ```

4. A window will appear where you can select the CA that will issue the certificate. Select the Windows 2000 issuing CA and click OK.

5. Note the RequestID that is shown after the command has finished.

**Process the Certificate Request**

The next step is to set the certificate extensions in the pending certificate request using either the Windows 2000 or the Windows Server 2003 version of certutil.

| **Important** |
| --- |
| The following command will override any existing information in the certificate request that matches the given object identifier. If you have specified, for example, a subject alternative name as part of the certificate request, the subject alternative name will be overridden with this command. |

1. Copy the <dcname>.asn file from the domain controller to the CA computer.

2. At a command-line prompt, run the following command.

   ```
   certutil -setextension <RequestID> 2.5.29.17 1 @<dcname>.asn
   ```

   The subject alternative name, which is identified by the object identifier 2.5.29.17, is set with the attributes that are defined in the <dcname>.asn file. The fourth parameter that is set to "1" marks the extension as critical. For a description of the file structure of the ASN file, see Appendix 5: ASN.1 File Structure.

**Verify the Pending Certificate Request**

You may verify the pending certificate request to validate that the extensions are properly inserted in the certificate request. If something is not correct or if the certificate request does not meet your requirements, you can validate the information at this point to ensure accuracy. Unfortunately, the Windows 2000 CA does not support viewing the properties of a pending certificate request in the Certification Authority MMC Snap-In. To verify that the certificate will be issued with the correct certificate template, you can use either the Windows 2000 or Windows Server 2003 version of certutil.

1. From a command-line prompt, run the following command.

   ```
   certutil -view -restrict RequestID=<RequestID> -out RequestAttributes
   ```

   Replace the <RequestID> with the RequestID that was recorded previously.

   The command output should look similar to the following:

   ```
   Row 1:
   ```

```
Request Attributes: "CertificateTemplate:DomainController"
```

2. To view the request attributes of all pending certificate requests, run the following command.

```
certutil -view -restrict disposition=9 -out RequestID,RequestAttributes
```

3. Verify that the certificate extensions have been set correctly. Type the following command at a command-line prompt.

```
certutil -view -restrict RequestID=<RequestID> -out ext:2.5.29.17
```

   Replace the <RequestID> with the RequestID that was recorded previously. The “2.5.29.17” value represents the object identifier of the certificate extension that you are referencing. If you do not know the exact object identifier of your extension, use “all” instead of the specific object identifier. For example:

```
certutil -view -restrict RequestID=<RequestID> -out ext:all
```

   Ensure that the extension you have configured in the previous section appears in the pending certificate request.

| **Note** |
| --- |
| All of the certutil -view commands may be used to query the certification authority database on a Windows 2000 or Windows Server 2003 CA. Once you understand the schema, you can query the database providing the appropriate field names. You can also use certutil -view -? for more samples to query the database. |

**Issue and Retrieve the Certificate**

Next, you must issue the pending certificate. For this section, you can use either the Windows 2000 or Windows Server 2003 version of certutil.

1. To issue the pending request, open the MMC snap-in. Run certsrv.msc at a command-line prompt, and then press Enter.

2. In the left pane, click the Pending Certificates container.

3. In the right pane, select the certificate that corresponds to the RequestID recorded previously.

4. On the Action menu, select All Tasks – Issue.

   Alternatively, you can issue certificates from the command line with the following command.

```
certutil -resubmit <RequestID>
```

   Replace <RequestID> with the RequestID value that was recorded previously by the certreq command.

5. Once the certificate has been issued, you must store the certificate as a file to transfer it to the domain controller. The following command will create two certificate files. The CER file contains only the domain controller certificate; the P7B file contains the domain controller certificate and all of its parent certificates. At a command-line prompt, run the following command.

```
CERTREQ -retrieve <RequestID> <dcname>.cer <dcname>.p7b
```

   Replace <RequestID> with the RequestID that was used in the previous commands.

   A window appears where you can select the CA that has issued the certificate.

6. Select the issuing CA and click OK.

7. Store the retrieved certificates on a diskette or other medium to transfer to the domain controller.

8. Log off the CA.

To install the certificate on the domain controller, see Domain Controller Certificate Installation.

**Windows 2000 Enterprise CA**
Unfortunately, domain controller certificates cannot be issued manually from a Windows 2000 enterprise CA for the following reasons.

- A Windows 2000 CA supports only hard-coded certificate templates so that you cannot duplicate and customize the default certificate template. If the original domain controller certificate was manipulated, it would affect all other domain controllers that are able to connect to the CA and use auto-enrollment to request certificates.

- A Windows Server 2000 CA issues certificates immediately without pending. Since the submission interface on the CA has limited functionality in Windows 2000 and, therefore, certificate manipulation is required in a pending state, support for offline domain controllers is particularly problematic.

Therefore, if support for offline certificate request processing is required, it is recommended that you install a Windows 2000 stand-alone CA or a Windows Server 2003 CA.

**Issuing Domain Controller Certificates with a Windows Server 2003 CA**
A Windows Server 2003 enterprise CA may support V2 templates and pending requests on a per template basis, as well as support for submitting request attributes in enrollment request from both enterprise and stand-alone CAs. Therefore, the following sections document the unique procedures for each CA type. V2 templates are only available if the CA was installed on Windows Server 2003, Enterprise Edition.

**Windows Server 2003 Stand-alone CA**
A stand-alone CA issues certificates based on fixed rules similar to those in templates. As described previously, a stand-alone CA issues certificates similar to the *Domain Controller* certificate template; thus, you cannot enroll *Directory Email Replication* or *Domain Controller Authentication* certificates with this CA type.

When issuing certificates manually, it does not matter if the stand-alone CA is a member of a domain or a workgroup. Both configurations are supported, and since the subject alternative name information was included in the certificate request when reqdccert.vbs was run, it is not required to place the certificate requests in a pending state to process them correctly.

To issue a domain controller certificate from a stand-alone Windows Server 2003 CA, perform the following steps.

1. Log on to the CA computer.

2. Copy the certificate request and the batch script (<dcname>-req.bat) that was created on the domain controller previously into a working folder on the CA.

3. From a command-line prompt, run the <dcname>-req.bat script to perform the certreq command. The script will request a certificate based on the *DomainController* certificate template because no other template exists for this purpose.

   The bat-file contains the following command.

   ```
   CERTREQ -attrib "CertificateTemplate:DomainController" <requestfile>
   ```

4. A window will appear where you can select the CA that will issue the certificate. Select the issuing CA and click OK.

5. Make a note of the RequestID that is shown after the command has finished. This value will be needed later.

6. If the enrollment handling is set to the default configuration on a stand-alone CA, you will have to approve and issue the pending certificate. If the CA is configured to issue certificates automatically, continue to step 15.

7. To verify and issue the pending request, open the MMC snap-in.

8. In the left pane, click the *Pending Certificates* container.

9. In the right pane, select the certificate that corresponds to the RequestID noted in step 5.

10. On the Action menu, select All Tasks - View Attributes/Extensions.

11. A window opens that displays the request properties. Click the Extensions tab.

12. Verify the value of the Subject Alternative Name tag. It should display the hexadecimal encoded GUID and the FQDN of the domain controller.

13. Click OK to close the window.

14. With the certificate request still selected in the right pane, choose All Tasks - Issue on the Action menu.

    Alternatively, you can issue certificates from the command-line prompt with the following command.

    ```
    certutil -resubmit <RequestID>
    ```

    Replace the <RequestID> variable with the RequestID that was noted in step 5.

15. Once the certificate has been issued, you will have to store the certificate as a file to transfer it to the domain controller. The following command will create two certificate files. The CER file contains only the domain controller certificate; the P7B file contains the domain controller certificate and all of its parent certificates. At a command-line prompt, run the following command.

    ```
    CERTREQ -retrieve <RequestID> <dcname>.cer <dcname>.p7b
    ```

    Replace the <RequestID> variable with the RequestID that was used in the previous commands.

16. A window will appear where you can select the CA that has issued the certificate. Select the issuing CA and click OK.

17. Store the retrieved certificates on a diskette or other medium to transfer to the domain controller.

18. Log off the CA.

Next, continue to Domain Controller Certificate Installation.

**Windows Server 2003 Enterprise CA**
An enterprise CA maintains and uses certificate templates in Active Directory for all certificate request processing and issuance. Therefore, domain controller certificates are formatted and issued based on the templates that are available and assigned to a CA. In the case of offline domain controller certificates, modification of the default certificate templates in Active Directory is necessary to support the enrollment processing.

Since the following section requires V2 templates, it is assumed that the enterprise CA was installed on a server with Windows Server 2003, Enterprise Edition.

**Certificate Template Creation**
In the following steps, duplicates of the existing *Domain Controller* certificate templates are created to support offline request processing. The following are the two primary reasons for modifications.

- By default, the offline *Directory Email Replication* certificate template publishes certificates in Active Directory. Normally, this would be satisfactory; however, the CA publishes an issued certificate into the certificate requestors Active Directory object and not into the certificate subject's Active Directory object. When a certificate request is submitted manually, the requestor is always the account of the user who created the certificate request. Thus, the CA would publish a domain controller certificate in the administrator's Active Directory user object. To work around this issue, a new certificate template is created that does not require the CA to publish the domain controller certificate in Active Directory. For information about manual certificate publication, see Publishing Domain Controller Certificates.

- Both *Directory Email Replication* and *Domain Controller Authentication* certificate templates require a special template flag to enable offline request processing. To leave the existing templates intact for all normal domain controller certificates processing, duplicates are made of both templates with the special settings.

| **Note** |
| --- |
| It is recommended that you do not manipulate the *Directory Email Replication* and *Domain Controller Authentication* default templates. They are used for auto-enrollment by Windows Server 2003 domain controllers that can connect to the enterprise CA. Modification of these templates may cause an interruption of service for auto-enrollment of connected domain controllers. |

Perform the following steps to duplicate the default *Domain Controller* certificate templates.

1. Log on with Enterprise Admin permissions to a domain computer that has Windows Server 2003 Administration Tools (AdminPak.msi) installed and is a member of the forest where the enterprise CA is installed. Any computer joined to the Active Directory forest can be used to maintain certificate templates since they are stored in Active Directory. Connectivity to a CA is not required at this point. By default, only members of the Enterprise Administrators group have permission to change certificate templates.

2. Click the Start button, and then point to Run. Type *certtmpl.msc* and press Enter.

3. The certificate Templates MMC Snap-In opens. In the right pane, select the Directory Email Replication template.

4. On the Action menu, select Duplicate template.

5. The Properties window of the new template appears. On the General tab, type *Offline Directory Email Replication* as Template display name. Click to clear the Publish certificate in Active Directory check box.

   Technically, any template name may be used; however, a meaningful label is recommended.

6. Click the Subject Name tab. Click to select the Supply in the request option and click OK.

7. In the right pane, select the Domain Controller Authentication template.

8. On the Action menu, select Duplicate template.

9. The Properties window of the new template appears. In the General tab, type *Offline Domain Controller Authentication* as Template display name. The Publish certificate in Active Directory check box is already clear because, by default, authentication certificates are not published in Active Directory.

10. Click the Subject Name tab. Click to select the Supply in the request option and click OK.

11. Close the Certificate Templates MMC Snap-In.

**Certificate Template Configuration**

Note that the template change is not required for CAs running Windows Server 2003 Service Pack 1 or later. If you have installed Windows Server 2003 Service Pack 1 or a later version on your CA, continue to **Issuing a Domain Controller Certificate**.

Typically, an enterprise CA will read information that is stored in several certificate attributes from the requestor's object in Active Directory. For example, the domain controller's common name or its GUID are derived from the domain controller's computer object and inserted into a certificate when a new domain controller certificate is manually or automatically enrolled through normal processes. If a domain controller certificate is requested manually through an offline or asynchronous process, all information that is required in the certificate must be explicitly specified in the certificate request.

When the certificate request is submitted to the CA, the request contains a field with the certificate template's common name, which is subsequently used by the CA to determine the enrollment and issuance policies. One such policy is whether the request may define the subject and the subject alternative name. By default, these attributes must not be set in a certificate

request, and the CA will ignore such information in a request. It is important to note that auto-enrollment requires that the request not contain the subject information to process requests properly.

To configure the CA to allow the subject alternative name to be specified in certificate requests, the certificate templates that have been created in a previous section must be modified. Fortunately, this is a one-time operation, and it is not necessary to undo the change after a domain controller certificate has been enrolled. The template change will replicate normally such as any other attribute value change in your Active Directory environment and will not generate any schema changes.

To configure and apply the template modification before issuing an offline domain controller certificate, perform the following steps.

1. Log on as the Enterprise Administrator to a computer that is a member of the forest where the enterprise CA was installed.

2. Make the script available to your local computer as shown in FixDCtemplate.vbs.

3. Run the following script with the specified parameters.

   ```
   fixdctemplate.vbs <Templatename>
   ```

   Replace <Templatename> with the template's common name *OfflineDirectoryEmailReplication*. (Do not use blank spaces.)

   For more information about all the steps performed by the script, see FixDCtemplate.vbs.

4. Run the following script again with the specified parameters.

   ```
   fixdctemplate.vbs <Templatename>
   ```

   In this case, replace <Templatename> with the template's common name *OfflineDomainControllerAuthentication*. (Do not use blank spaces.)

5. Log off the computer.

**Windows Server 2003 CA Configuration**

In the previous section, the certificate templates have been created and configured. Once the changes in Active Directory have been replicated throughout the forest, the CA will be able to access and use these templates. However, in order for the CA to use and issue certificates based on the templates, the templates must be manually published on each CA before they will be available.

Perform the following steps to publish a template in an enterprise CA.

1. Log on to the computer where the CA was installed as a local computer administrator.

2. Open the MMC snap-in.

3. Expand the CA object in the right pane and select Certificate Templates.

4. On the Action menu, select New - Certificate Template to Issue.

5. The template selection window appears. Press and hold the <CTRL> key and select the templates Offline Directory Email Replication and Offline Domain Controller Authentication.

6. Click OK to add the templates to the CA.

   Both templates will appear in the right pane of the MMC snap-in.

7. Close the MMC snap-in.

**Issuing a Domain Controller Certificate with a Windows Server 2003 CA**

The steps to issue a domain controller certificate from a Windows Server 2003 enterprise CA are similar to the Windows Server 2003 stand-alone CA procedures. However, since the Windows Server 2003 enterprise CA supports the new *Directory Email Replication* and *Domain Controller Authentication* certificate templates, these templates are used instead of the *Domain Controller* template.

Depending on the specific certificate template that is used, different attributes are specified as the subject alternative name. The *Directory Email Replication* requires the domain controllers GUID and its fully qualified DNS name (FQDN) in this extension. The *Domain Controller Authentication* template requires the FQDN only to be included in the subject alternative name extension. Therefore, you will find that the INF file and the certificate request look different for both certificate types.

| **Important** |
| --- |
| For offline domain controller certificate requests, never use the *Domain Controller* certificate template because a Windows Server 2003 CA supersedes that template with the *Directory Email Replication* and *Domain Controller Authentication* certificate templates. You will receive error 0x80094803 if you use the *Domain Controller* certificate template. |

To issue a domain controller certificate from a Windows Server 2003 enterprise CA, perform the following steps.

1. Log on to the CA computer.

2. Copy the certificate request and the batch script (<dcname>-req.bat) that was created on the domain controller in a previous step into a working folder on the CA.

3. From a command-line prompt, use the <dcname>-req.bat script to run the certreq command. The script will request a certificate based on the given certificate template.

   The bat file contains the following command if the *Directory Email Replication* certificate template is used.

   ```
   CERTREQ -attrib "CertificateTemplate:<TemplateName>" <requestfile>
   ```

   A window will appear where you can select the CA that will issue the certificate.

4. Select the issuing CA and click OK.

5. If you have changed the default enrollment handling in one of the *Domain Controller* certificate templates or in the CA policy module configuration, make a note of the RequestID that is shown after the previous command has finished to issue the pended certificate. Otherwise, a Save window appears where you can set the name of the certificate file. Type the name of the certificate and click OK. Continue to step 9.

6. If the certificate request was pended, run the following command at a command-line prompt. Otherwise, continue to step 7.

   ```
   certutil -resubmit <RequestID>
   ```

   Replace <RequestID> with the RequestID that was recorded previously.

7. Once the certificate is issued, you have to store it as a file to transfer it to the domain controller. The following command will create two certificate files. The CER-file contains only the domain controller certificate; the P7B file contains the domain controller certificate and all of its parent certificates. At a command-line prompt, run the following command.

   ```
   CERTREQ -retrieve <RequestID> <dcname>.cer <dcname>.p7b
   ```

   Replace <RequestID> with the RequestID that was used in the previous commands.

8. A window will appear where you can select the CA that has issued the certificate. Select the issuing CA and click OK.

9. Store the retrieved certificates on a diskette or other medium to transfer to the domain controller.

10. Log off the CA.

# Domain Controller Certificate Installation

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

Installing a certificate on a domain controller implies that the certificate and associated private key are available to the local system (computer) account. Since the key material has been previously generated and the certificate request is still pending at the domain controller, the certificate must be accepted (installed). Acceptance refers to the operation whereby the certificate and the key material are linked, and the certificate request is deleted from the *Certificate Enrollment Requests* container.

When an Administrator accepts a certificate, certreq will look into the *Certificate Enrollment Requests* container of the local machine store first, and if no corresponding request is found, it looks into the current user's (Administrator) *Certificate Enrollment Requests* container. Thus, you can accept machine certificate requests as Administrator.

**Installing Domain Controller Certificates**

The following procedure will install the domain controller certificate locally in the local system profile but not in Active Directory. To install the certificate on the target domain controller, perform the following steps with the Windows Server 2003 version of certreq and certutil. On a Windows 2000 domain controller, you must add a prefix to the commands. The prefix is the path you have copied the commands to. In this white paper, the *%HOMEDRIVE%\W2K3AdmPak* path is used.

1. Log on to the target domain controller.

2. Make the CER- and P7B-file from the previous section available to the domain controller.

3. From a command-line prompt, run the following command.

   ```
   CERTREQ -ACCEPT <dcname>.p7b
   ```

   Replace <dcname> with the name of the target domain controller. The command will not report any confirmation of success.

4. Verify that the certificate has been installed in the local system personal certificate store by running the following command at a command-line prompt:

   ```
   certutil -viewstore My
   ```

   A window will appear that displays all certificates that are available in the local computer personal store.

5. Log off your domain controller.

**Publishing Domain Controller Certificates**

As mentioned previously, a stand-alone CA is not capable of publishing certificates in Active Directory. However, SMTP replication requires the use of domain controller certificates published in Active Directory. For certificates that are enrolled asynchronously through an offline process, you must manually publish these certificates in Active Directory.

The following steps instruct you to examine the certificates that reside in the domain controller's local machine certificate store before the new certificate is published in Active Directory. The new certificate is published and you will be able to verify that the certificate was published properly.

Publishing certificates into computer objects in Active Directory requires write permissions for the userCertificate attribute that is part of any computer object. Administrators and members of the built-in domain group "Cert Publishers" have this permission by default.

The reqdccert.bat script creates a file called <dcname>-vfy.bat that contains the correct certutil command to verify the domain controller's certificate in Active Directory. You can use this batch file instead of typing the certutil –viewstore commands manually as described in the following steps.

Perform the following steps to manually view and publish a domain controller certificate in Active Directory.

1. Log on as domain administrator or a member of the Cert Publishers global group for the target domain controller. Technically, the publication can be performed at any computer that is a domain member, but for convenience, the domain controller is used in this scenario.

2. Verify that there are no certificates already published on the domain controller's Active Directory object.

   **Note**

   The following steps work only with the Windows Server 2003 version of certutil.exe. Thus, if you perform the following steps from a Windows 2000 domain controller, you must add a prefix to the certutil command. The prefix is the path you have copied the certutil command to. In this white paper, *%HOMEDRIVE%\W2K3AdmPak* is used.

   Run the following command from a command-line prompt.

   ```
   certutil -viewstore "ldap:///cn=<dcname>,ou=domain controllers,dc=<domainname>,dc=
   <com>?usercertificate"
   ```

   Replace the <dcname> variable with the name of the target domain controller and <domainname> and <com> variable names with the appropriate domain suffix.

   A window should appear with no certificates displayed. This is expected since no certificates have been published yet.

3. Click Cancel to close the window.

4. The certificate is published in Active Directory using the *userCertificate* attribute on the machine account object for the domain controller. Run the following command to write the certificate to the domain controller's Active Directory object.

   ```
   certutil –f –dspublish <dcname>.cer machine
   ```

   Replace the <dcname> variable with the name of the target domain controller.

   The command determines the proper Active Directory object by the subject information in the certificate. The publication will fail if no object can be found based on the subject information.

   **Note**

   The use of the "machine" parameter is a mandatory requirement in the previous command example.

5. To verify that the certificate was published successfully, perform the following steps from a command-line prompt.

   ```
   certutil -viewstore "ldap:///cn=<dcname>,dc=<domainname>,dc=<com>?usercertificate"
   ```

   If the domain controller's computer object has no certificates in the *userCertificate* attribute, the certutil output will display an empty list in the window. If "?userCertificate" was omitted from the command line parameters or an invalid object class was specified, an error message will appear such as the following:

   ```
   CertUtil: -viewstore command FAILED: 0x80092009 (-2146885623)
   CertUtil: Cannot find the requested object.
   ```

**Note**

It is always a good practice to verify the certificates in the requestor's Active Directory object as well. If the certificate templates have not been configured correctly and the certificate template was configured to publish certificates in Active Directory, these certificates may be published to the user account that created the certificate request. To examine certificates in the domain administrator's Active Directory user object, run the following command.

```
certutil -viewstore "ldap:///cn=Administrator,cn=Users,dc=<domainname>,dc=<com>"
```

If computer certificates are found in this object, follow the instructions to remove them in Removing Certificates from an Active Directory Computer Object.

# Removing Domain Controller Certificates

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

Occasionally, it may be necessary to remove or delete unwanted certificates locally in a certificate store, or remotely in Active Directory, to either correct mistakes or to perform periodic system cleanup. The following section describes the procedures to perform these activities. However, removing a certificate from a certificate store or an Active Directory object is not specific to domain controller certificates. The procedures are the same for any user or computer certificate type.

**Removing Certificates from a Local Certificate Store**

It may be necessary to remove certificates locally from a domain controller to ensure that only valid certificates are used by the domain controller or other applications. Expired or revoked certificates may be removed from a domain controller since the purpose of domain controller certificates is to encrypt replication traffic. Once replication has been performed, the replication data is discarded and there are no requirements to decrypt this information again. Therefore, such certificates may be safely deleted from the local machine profile.

To remove certificates from a dedicated domain controller, perform the following steps.

1. While logged on as a member of the local Administrators group, start the Microsoft Management console.
2. Add the Certificates MMC Snap-In.
3. Select Computer Account when prompted to select an account to manage.
4. In the Certificates MMC Snap-In, navigate to Personal in the left pane.
5. In the right pane, determine the domain controller certificate(s) by the template name as shown in the Certificate Templates column or select the certificate(s) by their intended purpose.
6. Delete the certificate(s) by selecting Delete on the Action menu.
7. Close the MMC Snap-In and log off.

**Removing Certificates from an Active Directory Computer Object**

In some cases, it may be necessary to remove certificates that are stored in an Active Directory object explicitly. Usually, this is the case if certificates have been enrolled manually. The auto-enrollment functionality in Windows XP and Windows Server 2003 can remove certificates from Active Directory objects when it determines that certificates in the Active Directory object have expired or are revoked. Also, the CA removes expired certificates when it publishes a new certificate in an object in Active Directory. For more information about the functionality performed by auto-enrollment, see the Certificate Autoenrollment in Windows Server 2003 white paper at

http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/autoenro.mspx

To manually remove a certificate from a domain controller object in Active Directory, perform the following steps.

1. Log on as a domain or enterprise administrator to a computer in your domain where the Windows Server 2003 version of certutil is installed and available.

   | **Note** |
   | --- |
   | The following steps work only with the Windows Server 2003 version of certutil.exe. Thus, if you perform the following steps from a Windows 2000 domain controller, you must add a prefix to the certutil command. The prefix is the path you have copied the certutil command to. In this white paper, the *%HOMEDRIVE%\W2K3AdmPak* path is used. |

2. At a command-line prompt, run the following command and press Enter.

```
certutil -viewdelstore "ldap:///cn=<dcname>,ou=domain controllers,<domainname>?
```

```
usercertificate"
```

The command contains the distinguished name of the object and specifies the LDAP attribute explicitly. If the LDAP attribute is missing or an invalid attribute was specified, the command will fail with an “Access denied” error message.

In the case of domain controller certificates, certificates are always stored in the *userCertificate* attribute. However, user objects may also have certificates in the *usercert* or *userSMIMEcertificate* attribute from other applications such as Outlook, which use these attributes as the preferred attribute for storing certificates.

A window will appear displaying the array of certificates on the user object. If the specified object contains certificates in the given attribute, the certificates are shown in the window. If no certificates are available, the list in the window will be empty.

3. To remove a specific certificate, select the certificate and click OK.

   The window will close and a status message will be displayed at the command-line prompt similar to the following: “Deleted certificate <certificate subject name>” where the <certificate subject name> displays the subject name of the certificate.

4. To close the window without deleting a certificate, click Cancel.

# Troubleshooting (Advanced Certificate Enrollment and Management)

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

Despite the level of detail and documented procedures in this white paper, some scenarios and environments may encounter problems that require troubleshooting. This section contains some of the common errors and troubleshooting tips for the procedures in this white paper to assist in resolving potential problems in your environment.

The first troubleshooting tool shows the full descriptive text for an error code displayed by an application or tool. If an unexpected error appears, you can display the corresponding error message text with the following command.

```
certutil -error <Hexadecimal_Error_ID>
```

For example, to see the error text for error 0x80094800, run the following command at a command-line prompt.

```
certutil -error 0x80094800
```

**Domain Controller Certificates Appear in User Objects**
If a certificate template is configured for the CA to publish a certificate in Active Directory during the enrollment and issuance process, the CA will choose the requestor's Active Directory object instead of the domain controller's computer object. To correct this issue, follow the steps in Removing Certificates from an Active Directory Computer Object.

**Certreq –new fails with error 0x80092023**
This error may occur if the subject name was not specified in an X.500 format. You cannot specify just the raw string as the common name. You must at least add the prefix "CN=" to the string.

**Certreq –submit fails with error 0x80094800**
This common error occurs when the template is not available on the CA where the certificate request was submitted. To correct this issue, make sure that the template name is spelled correctly in the certificate request or as an –attrib parameter. Also, it is important to verify that you have performed the steps in **Windows Server 2003 CA Configuration**.

**Certreq –submit fails with error 0x80094001**
This error occurs if you have performed the steps described in Issuing Domain Controller Certificates with a Windows Server 2003 CA where the request uses the V1 *Domain Controller* certificate template and includes a subject or subject alternative name. The V1 domain controller template instructs an enterprise CA to read the subject name of the requestor from Active Directory, but this fails with manual certificate requests. An administrator's user object in Active Directory will obviously never have the appropriate subject name for a domain controller.

You must not enroll manual certificate requests with V1 certificate templates on an enterprise CA.

**Certreq –submit fails with error 0x80094803**
This error typically occurs when a certificate template does not allow the subject or subject alternative name to be explicitly specified in the certificate request. By default, a template only allows information to be retrieved from Active Directory when building the subject or subject alternative name in a certificate. To correct this error for offline certificate enrollment, modify the specific template by running the fixdctemplate.vbs script in Appendix 2.

To verify the mandatory attributes of a subject alternative name in a template, run the following command at a command-line prompt with the Windows Server 2003 versions of certutil. On a Windows 2000 computer, you must add a prefix to the commands. The prefix is the path you have copied these commands to. In this white paper, the *%HOMEDRIVE%\W2K3AdmPak* path is used. You must be logged on as a member of the Authenticated User group in the Active Directory forest.

```
Certutil -v -dstemplate {CertificateTemplate_commonname}
```

For example, type

```
certutil -v -dstemplate OfflineDomaincontrollerauthentication
```

The command will display the properties of the specified template. Examine the parameters following *msPKI-Certificate-Name-Flag* in the output. Parameters that are indented in the output are disabled. Ensure that you have specified all attributes in the certificate request that are not indented in the output of the template properties.

**Certreq –submit fails with error 0x8009480e**
This error most commonly occurs when the template is not available on the CA to which you have submitted the request. This

situation may also be corrected by running the fixtemplate.vbs script in Appendix 2. See also Certreq –submit fails with error 0x80094803.

**Certutil –viewstore displays an empty dialog with no certificates**

This error most commonly occurs when an invalid object class was specified in the command-line parameter(s). You can use ADSIedit from the Windows Server 2003 Support Tools on the Windows Server 2003 CD-ROM to identify the correct object class and its distinguished name.

# Appendix 1: Identifying a Domain Controller GUID

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

Identifying the correct domain controller GUID for an SMTP replication certificate may pose a challenge for some administrators who are unfamiliar with the nuances of Active Directory and domain controller objects in the directory. To determine a specific domain controller GUID from a Windows XP or Windows Server 2003 computer joined to the Active Directory forest, perform the following steps.

| **Note** |
| --- |
| The dsquery utility is part of the Windows Server 2003 Administration Tools Pack and is not available on Windows 2000 computers. |

1. Log on to the computer with a domain account.

2. From a command-line prompt, run the following command.

   ```
   dsquery * "CN=<hostname>,OU=Domain Controllers,DC=<yourdomain>,DC=<yourdomain>" -
   scope base -attr objectguid
   ```

   You must replace the <hostname> variable with the name of the specific domain controller you want and the <yourdomain> variable with the domain name of your specific domain. For example:

   ```
   dsquery * "CN=DC01,OU=Domain Controllers,DC=contoso,DC=com" -scope base -attr
   objectguid
   ```

   The command will result in output similar to the following:

   ```
   Objectguid
   {57A8AAF4-686E-4128-8712-B6CA89FBF5BC}
   ```

3. Log off the computer.

# Appendix 2: Sample Scripts

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

This section contains two sample Visual Basic scripts that may be used to simplify or customize the various processes described in previous sections.

**Reqdccert.vbs – Generates Domain Controller Certificate Requests**
The reqdccert.vbs script makes it easier to create the correct INF file that is required to submit a certificate request to a Windows 2000 or Windows Server 2003 CA. Since it is fairly complicated to determine the GUID of a computer and create an ASN.1 file that contains both the GUID and the DNS name, the script generates the ASN.1 file automatically. In addition, a batch file is created that contains the certreq –submit command(s), which are required to submit the certificate request to the CA. Finally, a validation script is generated that allows you to verify the certificate(s) on a domain controller's computer object after you have published the domain controller certificate(s) in Active Directory.

If you plan to submit a certificate request to a Windows 2000 or a Windows Server 2003 stand-alone CA, run the script without any command-line parameters. If you plan to submit the certificate request to a Windows Server 2003 enterprise CA, you must specify the name of the certificate template that will be used to enroll the domain controller certificate.

| **Note** |
| --- |
| Since the script requires the Windows Server 2003 version of certutil for an internal encoding operation, it is highly recommended, on a Windows 2000 domain controller, to put this script in the same directory as the Windows Server 2003 version of certutil and certreq.exe. |

```
Set oArgs = WScript.Arguments
Set oShell = WScript.CreateObject("WScript.Shell")
'
' Parse command line
'
if oArgs.Count < 1 then
    sTemplateName = "DomainController"
    sType = "E"
else
    if ((oArgs(0) = "-?") or (oARgs.Count < 2)) then
        Wscript.Echo "Usage: reqdccert.vbs [Templatename] [Type]"
        Wscript.Echo "[Templatename] is the name of a V2 template"
        Wscript.Echo "[Type]         can be E for Email and A for Authentication
certificate"
        Wscript.Echo "If no option is specified, the DomainController certificate
template is used."
        Wscript.Quit 1
    else
        sTemplateName = oArgs(0)
                sType = oArgs(1)
    end if
end if
Set oFilesystem = CreateObject("Scripting.FileSystemObject")
Set objSysInfo = CreateObject("ADSystemInfo")
Set objDC = GetObject("LDAP://" & objSysInfo.ComputerName)
sGUID = objDC.GUID
sDNShostname = objDC.DNShostname
sHostname = objDC.cn
'##############################################################################
'
' Create the ASN.1 file
'
'##############################################################################
Dim aASNsubstring(2, 5)
Const HEX_DATA_LENGTH = 1
Const ASCIIDATA = 2
Const HEXDATA = 3
Const HEX_BLOB_LENGTH = 4
Const HEX_TYPE = 5
```

```
aASNsubstring(0, ASCIIDATA) = sDNShostname
aASNsubstring(0, HEX_TYPE) = "82"
'
' Convert DNS name into Hexadecimal
'
For i = 1 to Len(aASNsubstring(0, ASCIIDATA))
    aASNsubstring(0, HEXDATA) = aASNsubstring(0, HEXDATA) & _
                                    Hex(Asc(Mid(aASNsubstring(0, ASCIIDATA), i, 1)))
Next
aASNsubstring(0, HEX_DATA_LENGTH) = ComputeASN1 (Len(aASNsubstring(0, HEXDATA)) / 2)
'
' Build the ASN.1 blob for DNS name
'
sASN = aASNsubstring(0, HEX_TYPE) & _
       aASNsubstring(0, HEX_DATA_LENGTH) & _
       aASNsubstring(0, HEXDATA)
'
' Append the GUID as other name
'
if (sType = "E") then
    aASNsubstring(1, HEXDATA) = sGUID
    aASNsubstring(1, HEX_TYPE) = "A0"
    aASNsubstring(1, HEX_DATA_LENGTH) = ComputeASN1 (Len(aASNsubstring(1, HEXDATA)) /
2)
    sASN = sASN & _
           "A01F06092B0601040182371901" & _
           aASNsubstring(1, HEX_TYPE) & _
           "120410" & _
           aASNsubstring(1, HEXDATA)
end if
'
' Write the ASN.1 blob into a file
'
Set oFile = oFilesystem.CreateTextFile(sHostname & ".asn")
'
' Put sequence, total length and ASN1 blob into the file
'
oFile.WriteLine "30" & ComputeASN1 (Len(sASN) / 2) & sASN
oFile.Close
'
' Use certutil to convert the hexadecimal string into bin
'
oShell.Run "certutil -f -decodehex " & sHostname & ".asn " & _
                                        sHostname & ".bin", 0, True
'
' Use certutil to convert the bin into base64
'
oShell.Run "certutil -f -encode " & sHostname & ".bin " & _
                                    sHostname & ".b64", 0, True
'##############################################################################
'
' Create the INF file
'
'##############################################################################
Set iFile = oFilesystem.OpenTextFile(sHostname & ".b64")
Set oFile = oFilesystem.CreateTextFile(sHostname & ".inf")
oFile.WriteLine "[Version]"
oFile.WriteLine "Signature= " & Chr(34) & "$Windows NT$" & Chr(34)
oFile.WriteLine ""
oFile.WriteLine "[NewRequest]"
oFile.WriteLine "KeySpec = 1"
oFile.WriteLine "KeyLength = 1024"
oFile.WriteLine "Exportable = TRUE"
oFile.WriteLine "MachineKeySet = TRUE"
oFile.WriteLine "SMIME = FALSE"
oFile.WriteLine "PrivateKeyArchive = FALSE"
oFile.WriteLine "UserProtected = FALSE"
oFile.WriteLine "UseExistingKeySet = FALSE"
oFile.WriteLine "ProviderName = " & Chr(34) & _
                "Microsoft RSA SChannel Cryptographic Provider" & Chr(34)
oFile.WriteLine "ProviderType = 12"
oFile.WriteLine "RequestType = PKCS10"
```

```
oFile.WriteLine "KeyUsage = 0xa0"
oFile.WriteLine ""
oFile.WriteLine "[EnhancedKeyUsageExtension]"
oFile.WriteLine "OID=1.3.6.1.5.5.7.3.1"
oFile.WriteLine "OID=1.3.6.1.5.5.7.3.2"
oFile.WriteLine ";"
oFile.WriteLine "; The subject alternative name (SAN) can be included in the INF-file"
oFile.WriteLine "; for a Windows 2003 CA."
oFile.WriteLine "; You don't have to specify the SAN when submitting the request."
oFile.WriteLine ";"
oFile.WriteLine "[Extensions]"
iLine = 0
Do While iFile.AtEndOfStream <> True
    sLine = iFile.Readline
    If sLine = "-----END CERTIFICATE-----" then
        Exit Do
    end if
    if sLine <> "-----BEGIN CERTIFICATE-----" then
        if iLine = 0 then
            oFile.WriteLine "2.5.29.17=" & sLine
        else
            oFile.WriteLine "_continue_=" & sLine
        end if
        iLine = iLine + 1
    end if
Loop
oFile.WriteLine "Critical=2.5.29.17"
oFile.WriteLine ";"
oFile.WriteLine "; The template name can be included in the INF-file for any CA."
oFile.WriteLine "; You don't have to specify the template when submitting the request."
oFile.WriteLine ";"
oFile.WriteLine ";[RequestAttributes]"
oFile.WriteLine ";CertificateTemplate=" & sTemplateName
oFile.Close
iFile.Close
'##########################################################################
'
' Create the certreq.exe command-line to submit the certificate request
'
'##########################################################################
Set oFile = oFilesystem.CreateTextFile(sHostname & "-req.bat")
oFile.WriteLine "CERTREQ -attrib " _
                & Chr(34) & "CertificateTemplate:" & sTemplateName _
                & Chr(34) & " " & sHostname & ".req"
'
' The GUID structure needs to be reconstructed. The GUID is read
' as a string like f4aaa8576e6828418712b6ca89fbf5bc however the
' format that is required for the certreq command looks like
' 57a8aaf4-686e-4128-8712-b6ca89fbf5bc. The bytes are reordered
' in the following way:
'
'                           1111111111222222222333
'             Position 123456789012345678901234567890 12
'                      |------|--|--|--|--------------|
' Original GUID:       f4aaa8576e6828418712b6ca89fbf5bc
'
'                           11 1 1111 1112 222222222333
'             Position 78563412 1290 5634 7890 123456789012
'                      |------- |--- |--- |--- |----------|
' Reformatted GUID:    57a8aaf4-686e-4128-8712-b6ca89fbf5bc
'
oFile.WriteLine "REM "
oFile.WriteLine "REM !!! Only valid for Windows 2003 or later versions !!!"
oFile.WriteLine "REM If you do not specify certificate extensions in the *.INF file"
oFile.WriteLine "REM they can be specified here like the following example"
oFile.WriteLine "REM "
oFile.WriteLine "REM CERTREQ -submit -attrib " _
                & Chr(34) & "CertificateTemplate:" & sTemplateName _
                & "\n" _
                & "SAN:guid=" _
                & Mid(sGUID, 7, 2) _
                & Mid(sGUID, 5, 2) _
```

```
                 & Mid(sGUID, 3, 2) _
                 & Mid(sGUID, 1, 2) & "-" _
                 & Mid(sGUID, 11, 2) _
                 & Mid(sGUID, 9, 2) & "-" _
                 & Mid(sGUID, 15, 2) _
                 & Mid(sGUID, 13, 2) & "-" _
                 & Mid(sGUID, 17, 4) & "-" _
                 & Mid(sGUID, 21, 12) _
                 & "&DNS=" & sDNShostname & Chr(34) & " " & sHostname & ".req"
oFile.Close
'##############################################################################
'
' Create the certificate verification script
'
'##############################################################################
Set oFile = oFilesystem.CreateTextFile(sHostname & "-vfy.bat")
oFile.WriteLine "certutil -viewstore " & Chr(34) & objDC.distinguishedname & _
                "?usercertificate" & chr(34)
oFile.Close
'##############################################################################
'
' Compute the ASN1 string
'
'##############################################################################
Function ComputeASN1 (iStrLen)
    If Len(Hex(iStrLen)) Mod 2 = 0 then
        sLength = Hex(iStrLen)
    else
        sLength = "0" & Hex(iStrLen)
    end if
    if iStrLen > 127 then
        ComputeASN1 = Hex (128 + (Len(sLength) / 2)) & sLength
    else
        ComputeASN1 = sLength
    End If
End Function
```

**FixDCtemplate.vbs**

The fixDCtemplate.vbs script simplifies a change that is required for any certificate template that needs to accept the subject and subject alternative name as part of a certificate request. To change the *msPKI-Certificate-Name-Flag* in the certificate templates object, Active Directory Service Interfaces (ADSI) is used.

The *msPKI-Certificate-Name-Flag* attribute is actually a bit field, where the first bit determines whether the template accepts the subject and subject alternative name. Before changing the bit, the script determines if the bit is already set.

```
Set oArgs = WScript.Arguments
'
' Parse command line
'
if oArgs.Count <> 1 then
    Wscript.Echo "fixdctemplate {templatename}"
    WScript.Quit 1
end if
Set objRoot = GetObject("LDAP://rootDSE")
Set objTemplate = GetObject("LDAP://CN=" & oArgs(0) & _
                              ",CN=Certificate Templates," & _
                              "CN=Public Key Services,CN=Services," & _
                              objRoot.Get("configurationNamingContext"))
iNameFlag = objTemplate.Get("msPKI-Certificate-Name-Flag")
if iNameFlag = 1 then
    Wscript.Echo "Flag is already set"
else
    objTemplate.Put("msPKI-Certificate-Name-Flag"), 1
    objTemplate.SetInfo
    Wscript.Echo "Flag was set"
end if
```

# Appendix 3: Certreq.exe Syntax

Updated: May 20, 2013

Applies To: Windows Server 2003 with SP1

Certreq.exe is a command-line tool that is included in Windows 2000 Server, Windows XP, and Windows Server 2003 as well as in the Windows Server 2003 Administration Tools Pack. The version available with Windows XP and Windows Server 2003 is different from the version that is included in Windows 2000 Server. When using certreq.exe, you must distinguish between these versions.

| **Important** |
| --- |
| This document was written for Windows XP and Windows Server 2003. Some, but not all of the settings discussed in this article will apply to newer operating system. See Certreq as the reference for newer operating system versions. |

| **Tip** |
| --- |
| Several readers have asked for an actual example and mentioned that certreq is more difficult than using a graphical user interface. Microsoft KB article 931351 provides an example of how to perform a certificate request using a web interface as well as using Certreq. |

You can perform the following tasks with certreq.exe.

- Create new certificate requests. This includes the generation of key material that corresponds to the certificate request.
- Apply policies to a certificate request. This applies to CA certificates where various constraints are required. Since you cannot include certificate policies in the certificate request that is created by a CA, you have to apply the policy to the CA certificate request explicitly.
- Submit certificate requests to a certification authority. This feature also works with remote certification authorities, so that you can submit a request via the network to a remote CA.
- Retrieve certificates from a CA. Based on the certificate request ID, you can retrieve any certificate from a CA and save it to a file.
- Accept (and install) certificates or responses. This command is intended to accept a certificate or a response that contains a certificate. Once you have created a certificate request with certreq.exe, it remains in a pending state on your computer until you install the certificate that corresponds to the request.
- Sign certificate requests. In the case where the certificate template administrator has set the number of certificate enrollment agent signatures to a value greater than zero, you can sign certificate requests with this option.

The Windows 2000 version of certreq.exe can only submit certificate requests and retrieve issued certificates. All other options require the certreq.exe version that comes with Windows XP or Windows Server 2003. For more information about the command-line syntax, run certreq,exe with the following parameters to display all available syntax and parameter information.

```
Certreq.exe -v -?
```

**Creating Certificate Requests**

The –new parameter with certreq.exe is used to construct new certificate requests. Certreq.exe uses an INF file as an input option that defines the certificate request parameters. When the certificate request is constructed based on the INF file, key material is also generated. A certificate request process can be divided into the following elements.

- Read the INF file.
- Create the private and public key based on the information in the INF file, and store the private key in the local key store of the user or computer profile as appropriate.

- Create the certificate request based on the information in the INF file and store the request in a Base64-encoded file or, optionally, in binary form if the –binary option is selected.

**Certreq.exe INF File Structure**
Since the INF file allows for a rich set of parameters and options to be specified, it is difficult to define a default template that administrators should use for all purposes. Therefore, this section describes all the options to enable you to create an INF file tailored to your specific needs.

The following key words are used to describe the INF file structure.

- A *section* is an area in the INF file that covers a logical group of keys. A section always appears in brackets in the INF file.
- A *key* is the parameter that is to the left of the equal sign.
- A *value* is the parameter that is to the right of the equal sign.

For example, a minimal INF file would look similar to the following:

```
[NewRequest]
; At least one value must be set in this section
Subject = "CN=W2K3-BO-DC.contoso2.com"
```

| **Note** |
|---|
| You can remark lines by putting a semi-colon (;) in front of the line. |

The following are some of the possible sections that may be added to the INF file.

**[Version]**
This section is optional and is not required in an INF file. If present, it must contain the Signature key.

**Signature**
The signature key must be equal to the fixed string "$Windows NT$".

**[NewRequest]**
This section is mandatory for an INF file that acts as a template for a new certificate request. If this section is missing, the following error message will be displayed.

"INF file line not found 0xe0000102 (INF: -536870654)"

This section requires at least one key with a value. If this section is empty and has no keys, the following error message will be displayed.

"Incorrect function. 0x1 (WIN32: 1)"

**EncipherOnly**

| | |
|---|---|
| Syntax | EncipherOnly={Boolean} |
| Values | TRUE or FALSE |
| Default value | FALSE |
| Sample | `EncipherOnly = TRUE` |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | Yes (through the key usage extension) |

This parameter has an impact on the functional capabilities of the public-private key pair. If the value is set to TRUE, the key can exclusively be used for "encipherment" (encryption). If it is set to FALSE, the key can be used for "encipherment" and other purposes. This parameter refers only to the key type in CryptoAPI and has an indirect relationship only to the key usage

extension that may be included in an issued X.509 certificate. The default value of the certificate enrollment control is set to FALSE. This setting only affects AT_KEYEXCHANGE key types and then correspondingly limits the key usage to key Encipherment, data Encipherment, or both. The digital signature and non-repudiation key usage(s) will be disallowed.

**Exportable**

| | |
|---|---|
| Syntax | Exportable = {Boolean} |
| Values | TRUE or FALSE |
| Default value | FALSE |
| Sample | `Exportable = TRUE` |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

This parameter is ignored when the "UseExistingKeySet" key is set to TRUE because you can set the exportable flag only when a new key is created. You cannot change this flag for an existing key. If this attribute is set to TRUE, the private key can be exported with the certificate. To ensure a high level of security, private keys should not be exportable; however, in some cases, it might be required to make the private key exportable if several computers or users must share the same private key. For example, a Web server that runs SSL and is published with Internet Security and Acceleration (ISA) to the Internet requires a certificate with an exportable key because the certificate and key must be installed on the Web server and the ISA server as well. Also, if it is required that the keys be managed and backed up, they must be exportable to provide for this functionality.

This setting correlates with the template configuration if a Windows Server 2003 enterprise CA is used. A certificate template administrator can explicitly define if private keys should be exportable.

| **Note** |
|---|
| The template setting for exportable has no affect on the exportability properties of a key generated through certreq.exe -new. |

**KeyContainer**

| | |
|---|---|
| Syntax | KeyContainer = {Key_containerName} |
| Values | Random string value (GUID) |
| Default value | None |
| Sample | `KeyContainer = {C347BD28-7F69-4090-AA16-BC58CF4D749C}` |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

It is not recommended to set this parameter for new requests where new key material is generated. The key container is automatically generated and maintained by the system.

For requests where the existing key material should be used, this value can be set to the key-container name of the existing key.

Use the certutil –key command to display the list of available key containers for the computer context.

Use the certutil –key –user command for the current user's context.

If you need to find the key container for a specific certificate, use the certutil –store my command for machine certificates.

To find the key container for a specific certificate in the current user's certificate store, use the certutil –store –user my command.

If the "UseExistingKeySet" key is set to TRUE, the key container must be set explicitly in the INF file or in the renewal certificate.

**KeyLength**

| | |
|---|---|
| Syntax | KeyLength = {integer} |
| Values | Any valid key length that is supported by the cryptographic service provider |
| Default value | 1024 [depends on the certificate service provider (CSP)] |
| Sample | `KeyLength=2048` |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

The key length defines the length of the public and private key. The key length has an impact on the security level of the certificate. Greater key length usually provides a higher security level; however, some applications may have limitations regarding the key length.

**KeySpec**

| | |
|---|---|
| Syntax | KeySpec = {integer} |
| Values | AT_KEYEXCHANGE: 1<br>AT_SIGNATURE: 2 |
| Default value | 2 |
| Sample | `KeySpec=1` |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

| **Warning** |
|---|
| The KeySpec determines if the key can be used for signatures, for Key Exchange (encryption), or for both. If the KeySpec is set to a value of 2, the EncipherOnly key is ignored. Starting with Windows Vista and Windows Server 2003 R2, specifying EncipherOnly= in the INF will cause an error. |

**KeyUsage**

| | |
|---|---|
| Syntax | KeyUsage = {hexadecimal_value} |
| Values | CERT_DIGITAL_SIGNATURE_KEY_USAGE 0x80<br>CERT_NON_REPUDIATION_KEY_USAGE 0x40<br>CERT_KEY_ENCIPHERMENT_KEY_USAGE 0x20<br>CERT_DATA_ENCIPHERMENT_KEY_USAGE 0x10<br>CERT_KEY_AGREEMENT_KEY_USAGE 0x08<br>CERT_KEY_CERT_SIGN_KEY_USAGE 0x04<br>CERT_OFFLINE_CRL_SIGN_KEY_USAGE 0x02<br>CERT_CRL_SIGN_KEY_USAGE 0x02<br>CERT_ENCIPHER_ONLY_KEY_USAGE 0x01 |

| | |
|---|---|
| Default value | 0xC0 for AT_SIGNATURE (default)<br><br>For AT_KEYEXCHANGE, EncipherOnly = TRUE: 0x30<br><br>For AT_KEYEXCHANGE, EncipherOnly = FALSE: 0xf0 |
| Sample | `KeyUsage=0xa0` |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | Yes |

The key usage defines what the certificate key should be used for. The value is a bit field that is composed of the key usage flags as they are defined in the Windows Platform Software Development Kit (SDK). (See the include file wincrypt.h.) Caution should be used in defining the correct key usage for the certificate request based on the intended application usage.

| **Note** |
|---|
| The key usage extension has an indirect dependency with the enhanced key usage extension. For example, if you specify "Encrypting File System" as the enhanced key usage, you should also specify "Key Encipherment" as the key usage. |

To combine multiple key usages, summarize the values of the individual key usages by performing hexadecimal arithmetic. The default key usage is set to Digital Signature (0x80) and Non-Repudiation (0x40). The hexadecimal sum of 0x80 and 0x40 is in 0xC0.

The key usage parameter is primarily useful if certificate requests are submitted to stand-alone CAs. Only enterprise CAs have certificate key usages predefined through certificate templates. The key usage extension is set to critical by default.

**MachineKeySet**

| | |
|---|---|
| Syntax | MachineKeySet = {Boolean} |
| Values | TRUE or FALSE |
| Default value | FALSE |
| Sample | `MachineKeySet = TRUE` |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

This key is important when you need to create certificates that are owned by the computer and not a user. The key material that is generated is maintained in the security context of the security principal (user or computer account) that has created the request. When an administrator creates a certificate request on behalf of a computer, the key material must be created in the computer's security context and not the administrator's security context. Otherwise, the computer could not access its private key because it would be in the administrator's security context.

You must set this key to TRUE if you are creating requests for domain controllers, a Web server, or any other service that runs in the computer's security context.

**PrivateKeyArchive**

| | |
|---|---|
| Syntax | PrivateKeyArchive = {Boolean} |
| Values | TRUE or FALSE |
| Default value | FALSE |
| Sample | `PrivateKeyArchive = TRUE` |

| | |
|---|---|
| Supported with CA type | Windows Server 2003 |
| Can be manipulated in a pending request | No |

The *PrivateKeyArchive* setting works only if the corresponding RequestType is set to "CMC" because only the Certificate Management Messages over CMS (CMC) request format allows for securely transferring the requester's private key to the CA for key archival. Only the Windows Server 2003 enterprise CA supports private key archival. The Windows Server 2003 enterprise CA must be online and accessible to retrieve the CA encryption certificate directly during this process.

**ProviderName**

| | |
|---|---|
| Syntax | ProviderName = {CSP_stringname} |
| Values | The descriptive name of the certificate service provider |
| Default value | Microsoft Strong Cryptographic Provider |
| Sample | `ProviderName="Microsoft RSA SChannel Cryptographic Provider"` |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

The provider name is the display name of the CSP. If you do not know the provider name of the CSP you are using, run certutil –csplist from a command line. The command will display the names of all CSPs that are available on the local system.

**ProviderType**

| | |
|---|---|
| Syntax | ProviderType = {integer} |
| Values | The number that describes the provider type |
| Default value | 12 |
| Sample | `ProviderType=13` |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

The provider type is used to select specific providers based on specific algorithm capability such as "RSA Full". If you do not know the provider type of the CSP you are using, run certutil –csplist from a command-line prompt. The command will display the provider type of all CSPs that are available on the local system.

**RenewalCert**

| | |
|---|---|
| Syntax | RenewalCert={CertificateHash} |
| Values | The certificate hash of any certificate that is available at the computer where the certificate request is created |
| Default value | n/a |
| Sample | `RenewalCert=4EDF274BD2919C6E9EC6A522F0F3B153E9B1582D (must use double quotation marks if white space is used between pairs of hexadecimal values)` |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |

| | |
|---|---|
| Can be manipulated in a pending request | No |

If you need to renew a certificate that exists on the system where the certificate request is generated, you must specify its certificate hash as the value for this key. If you do not know the certificate hash, use the Certificates MMC Snap-In and look at the certificate that should be renewed. Open the certificate properties and see the "Thumbprint" attribute of the certificate. Certificate renewal requires either a PKCS#7 or a CMC request format.

| **Note** |
|---|
| You can only renew certificates that are time valid. Expired certificates cannot be renewed and must be replaced with a new certificate. |

**Requestername**

| | |
|---|---|
| Syntax | Requestername=samAccountName in Active Directory |
| Values | |
| Default value | n/a |
| Sample | `Requestername = "DOMAINNAME\username"` |
| Supported with CA type | Windows 2000 (with a PKCS#7 request type only)<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

The requester name can be specified for certificate requests if the *RequestType* is set to PKCS7 or CMC. If the *RequestType* is set to PKCS10, this key will be ignored. The *Requestername* can only be set as part of the request. You cannot manipulate the *Requestername* in a pending request.

**RequestType**

| | |
|---|---|
| Syntax | RequestType={string value} |
| Values | CMC<br>PKCS10<br>PKCS10-<br>PKCS7 |
| Default value | PKCS10 |
| Sample | `RequestType = CMC` |
| Supported with CA type | Windows 2000 (except CMC request type)<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

The RequestType determines the standard that is used to generate and send the certificate request.

**Silent**

| | |
|---|---|
| Syntax | Silent={Boolean} |
| Values | TRUE or FALSE |
| Default value | FALSE |
| Sample | `Silent = TRUE` |

| | |
|---|---|
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

By default, this option allows the CSP access to the interactive user desktop and request information such as a smart card PIN from the user. If this key is set to TRUE, the CSP must not interact with the desktop and will be blocked from displaying any user interface to the user.

**SMIME**

| | |
|---|---|
| Syntax | SMIME = {Boolean} |
| Values | TRUE \| FALSE |
| Default value | False (when KeySpec = AT_SIGNATURE) |
| Sample | `SMIME = TRUE` |
| Supported with CA type | Windows 2000 (if the policy module is configured in the registry to ignore the extension in the request)<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | Yes |

If this parameter is set to TRUE, an extension with the object identifier value 1.2.840.113549.1.9.15 is added to the request. The extension contains up to four object identifiers, depending on the CSP capability, which refer to symmetric encryption algorithms that may be used by Secure Multipurpose Internet Mail Extensions (S/MIME) applications such as Outlook.

When an S/MIME encrypted e-mail is sent, the sender's computer does not know exactly what encryption algorithms are supported by the receiver. Thus, the weakest encryption algorithm was chosen by default to ensure a high level of interoperability. Today, the majority of computers support encryption algorithms that can handle stronger keys than 40-bit. To give the sender of an S/MIME message a hint as to which encryption algorithms might be supported by the receiver, identifiers of supported symmetric algorithms are inserted in the certificate. Alternatively, Outlook can be configured through the registry to always use the default encryption algorithm. For more information, see the Knowledge Base article at http://support.microsoft.com/?id=307472

This functionality is supported with Outlook XP SP2 or a later version.

If the CA administrator has added the object identifier for S/MIME capabilities to the CA configuration parameter "*DisableExtensionList*", the CA will ignore the S/MIME capabilities in the request. For more information about how to add entries to the "*DisabledExtensionList*", see the Windows Server 2003 PKI Operations Guide at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws03pkog.mspx

**Subject**

| | |
|---|---|
| Syntax | Subject={String} |
| Values | Relative Distinguished Name string values |
| Default value | n/a |
| Sample | `Subject="CN=computer1.contoso.com"`<br>`Subject="CN=John Smith,CN=Users,DC=Contoso,DC=com" x` |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | Yes |

Several applications rely on the subject information in a certificate. Thus, it is recommended that a value for this key be specified. If the subject is not set here, it is recommended that a subject name be included as part of the subject alternative name certificate extension.

**UseExistingKeySet**

| | |
|---|---|
| Syntax | UseExistingKeySet={Boolean} |
| Values | TRUE or FALSE |
| Default value | FALSE |
| Sample | `UseExistingKeySet=TRUE` |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

This parameter is used to specify that an existing key pair should be used in building a certificate request. If this key is set to TRUE, you must also specify a value for the *RenewalCert* key or the KeyContainer name. You must not set the *Exportable* key because you cannot change the properties of an existing key. In this case, no key material is generated when the certificate request is built.

**UserProtected**

| | |
|---|---|
| Syntax | UserProtected={Boolean} |
| Values | TRUE or FALSE |
| Default value | FALSE |
| Sample | `UserProtected=TRUE` |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

If this key is set to TRUE, a CryptoAPI password window is displayed when the key is generated during the certificate request build process. You can optionally protect the key with a password in the window or choose to display only a window when the key is used within an application. Once the key is protected with a password, you must enter this password every time the key is accessed. For more information about strong private key protection, see http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/autoenro.mspx

**[EnhancedKeyUsageExtension]**

The enhanced key usage extension determines for what purposes a certificate can be used. To see a list of object identifiers that are supported by your CA, run the following command at a command-line prompt.

```
certutil -oid [OID string value]
```

| **Note** |
|---|
| You can add custom object identifiers to the CA configuration. Thus, the default list of object identifiers that is supplied with a Windows Server 2003 CA can be extended according to your requirements. |

For a list of object identifiers that are supported with the Microsoft PKI, see the Knowledge Base article "Object IDs Associated with Microsoft Cryptography" at http://support.microsoft.com/?id=287547.

**Object Identifier**

| | |
|---|---|
| Syntax | OID={OID_String} |
| Values | Any valid object identifier |
| Default value | n/a |

| | |
|---|---|
| Sample | `OID = 1.3.6.1.5.5.7.3.1` |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | Yes |

The object identifier is specified as a value that will be added to the enhanced key usage of the certificate. To specify multiple object identifiers, add a new line for every object identifier required. For example:

```
OID = 1.3.6.1.5.5.7.3.1
OID = 1.3.6.1.5.5.7.3.2
```

**[Extensions]**
Instead of specifying extensions in this section, you can set extensions at the time you submit the certificate request to a Windows Server 2003 CA or a Windows 2000 CA if the request is initially made pending.

**{OID}**

| | |
|---|---|
| Syntax | {OID} = {Base64Text} |
| Values | any valid object identifier that appears in the certificate |
| Default value | n/a |
| Sample | `2.5.29.17=MDmCFlcySy1CTy1EQy5jb25`<br>`_continue_=0b3NvMi5jb22gHwYJKwYBB`<br>`_continue_=AGCNxkBoBIEENKgYHzc`<br>`_continue_=F4dJmg+HRCPkkQ0=` |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | Yes |

Certificate extensions can be added to a request through this parameter. The value must be a Base64-encoded string. To split the string into several lines, use the _continue_ key word.

**Critical**

| | |
|---|---|
| Syntax | Critical = {coma-separated OIDList} |
| Values | Any valid object identifier that appears in the certificate [Extensions] section |
| Default value | n/a |
| Sample value | `Critical = 2.5.29.17,2.5.29.15` |
| Can be manipulated in a pending request | Yes |

This parameter carries the list of object identifiers of certificate extensions where the extension is set to critical. A certificate extension can also be set to critical with the certutil –setextension command during the time a certificate request is pending. For more information, see the section "certutil -setextension".

**[RequestAttributes]**
All keys from this section can be specified as part of the certreq -submit -attrib command. For example, instead of specifying the CertificateTemplate in the INF instruction file, you could use the following command at a command-line prompt.

```
certreq -submit -attrib "CertificateTemplate:Domaincontroller" dcname.req
```

**CertificateTemplate**

| | |
|---|---|
| Syntax | CertificateTemplate = {Templatename} |
| Values | Any valid certificate template name |

| | |
|---|---|
| Default value | n/a |
| Sample value | Critical = "DomainController" |
| Can be manipulated in a pending request | Yes |

The value of this key instructs the CA as to which certificate template should be used in processing the certificate request. The name must be specified as the common name of the certificate template or the object identifier of the template, not the template display name. A list of valid certificate template common names can be determined with the following command.

```
certutil -template
```

**SAN**

| | |
|---|---|
| Syntax | SAN = ["{tag}={value}&{tag}={value}]..."<br><br>To see an example of how to compose a request for a SAN see the "How to use the Certreq.exe utility to create and submit a certificate request that includes a SAN" section in Microsoft KB article 931351. |
| Values | A list of valid subjectalternatename2 extension tags.<br><br>*email, upn, dns, guid, url, ipaddress, oid*<br><br>Alternatively, any valid object identifier can be specified. |
| Default value | n/a |
| Sample value | See below; should always include double quotation marks around value(s) |
| Can be manipulated in a pending request | Yes |

The subject alternative name can be specified in various ways. A service principal name (SPN) is specified as a user principal name (UPN) value. The value is a string type by default but can have a prefix of a tag in section brackets "{}" that determines the data type of the value explicitly as exemplified in the following table.

| **SAN string** | |
|---|---|
| `2.5.29.17={octet}MBmCF1cySzMtQk8tREMuY29udG9zbzIuY29t` | An ASN.1 string is inserted in the subject alternative name with one or several encoded extensions.<br><br>The octet is a wrapper around a Base64 string encoded in binary format. |
| `1.2.3.4={utf8}SomeReadableText` | An ASN.1 string is inserted in the subject alternative name with one or several encoded extensions.<br><br>The value is a string encoded with a UTF8 wrapper. |
| `2.5.29.17={asn}=3019821757...` | An ASN.1 string is inserted in the subject alternative name with one or several encoded extensions. |
| `email=John.Smith@contoso.com` | The user's e-mail address is inserted in the subject alternative name. |
| `dns="w2k3bodc.contoso.com"` | The DNS name is set to the computer's FQDN. |
| `dn="CN=W2K3BODC,OU=Domain Controllers,DC=contoso,DC=com"` | The distinguished name is set to the domain controller's distinguished name. |

| | |
|---|---|
| `url="http://www.contoso.com/default.html"` | The URL is set to an http URL. |
| `ipaddress=172.134.10.134` | The IP is inserted in the subject alternative name. |
| `oid=2.5.29.17` | A given object identifier is inserted in the subject alternative name. |
| `upn=John.Smith@contoso.com` | The UPN name represents the user. |
| `guid=f7c3ac41-b8ce-4fb4-aa58-3d1dc0e36b39` | A GUID is inserted in the subject alternative name. |

Multiple tags and values must be combined with an ampersand (&). For example, to set the UPN and the GUID, you would have to set the subject alternative name as follows at a command-line prompt.

```
-attrib "SAN:upn=John.Smith@contoso.com&guid=f7c3ac41-b8ce-4fb4-aa58-3d1dc0e36b39"
```

If you insert the subject alternative name in the INF file, the syntax would look like this.

```
SAN="upn=John.Smith@contoso.com&guid=f7c3ac41-b8ce-4fb4-aa58-3d1dc0e36b39"
```

**-policy**
For a detailed description and examples of this option, see the "Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003" white paper at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws03qswp.mspx.

**-submit**
This is the default option for certreq.exe. If no option is specified explicitly at the command-line prompt, certreq.exe attempts to submit a certificate request to a CA. You must specify a certificate request file when using the –submit option. If this parameter is omitted, a common File Open window is displayed where you can select the appropriate certificate request file. For more information about the available command-line parameters for this option, see Certreq –submit fails with error 0x80094800 or type certreq -submit -? from a command line.

Certificate attributes from the RequestAttributes section in the INF instruction file can be alternatively specified as a command-line parameter for Windows 2000 Server and Windows Server 2003. For a list of valid subject alternative tags that may be used with the –submit parameter, see **SAN**. It is important to separate attribute types with "\n". Use an ampersand (&) to separate each specific subject alternative name attribute value.

The following are samples of certreq.exe commands.

```
certreq -attrib "SAN:email=John.Smith@contoso.com" -submit myreq.req
certreq -submit -
attrib "CertificateTemplate:Domaincontroller\nSAN:DNS=w2k3bodc.contoso.com" dcname.req
dcname.cer dcname.p7b
certreq -submit -attrib "SAN:UPN=admin@contoso.com&email=sample@bar.com&guid=f7c3ac41-
b8ce-4fb4-aa58-3d1dc0e36b39"
```

**-retrieve**
This option allows you to retrieve an issued certificate from the CA. Certificate owners or submitters can retrieve their certificate(s) from a Windows 2000 or 2003 CA. The certificate manager can retrieve all certificates from a Windows Server 2003 CA. This is not a security issue because certificates that are made available through the CA database contain only public information and only the original requestor may retrieve an issued certificate.

**-accept**
The final and most important step of the certificate enrollment process is to accept a certificate or response after it has been issued and retrieved from the CA. Certificates that are auto-enrolled or requested through the MMC Snap-In are accepted using the automation built into these applications. However, manual certificate requests, such as those submitted using certreq.exe, must be accepted and installed manually. The -accept parameter of certreq.exe links the previously generated private key with the issued certificate and removes the pending certificate request from the system where the certificate is requested. Accepting a CMC full response is required when archiving a private key on the CA.

**-sign**
Once the template administrator has configured the template to require one or more signatures on a certificate request, you must add these signatures to the request in order for the CA to process it correctly. The default Certificate Services Web enrollment pages can add only one signature to a certificate request, specifically when an Enrollment Agent certificate is used

to request a smart card certificate on behalf of another user. If more than one signature is required, you must add these signatures with certreq.exe –sign from the command line to the raw certificate request, one at a time.

# Appendix 4: Certutil -setextension

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

This option allows an administrator to manipulate pending requests on a CA. In some cases, it is not possible to specify every possible certificate extension in the certificate request. For example, if a request for a user certificate is submitted to the CA with the MMC snap-in, additional information is not allowed in the enrollment process.

Unfortunately, to apply an extension to a pending certificate, you have to create the ASN.1-encoded extension manually. Once you have constructed the ASN.1 binary large object (BLOB) information you want, you can add the extension to the certificate request that is pending. Encoding data into ASN.1 is not straightforward, and you have to understand the basic concept of this macro language used by certutil.exe. It is usually easier to find an existing request or certificate with the extensions you want, and then use certutil -v -dump to display the extensions as a hexadecimal dump.

For more information about ASN.1 encoding, see Appendix 5: ASN.1 File Structure.

The -setextension option is similar to the *[Extensions]* section in the INF file used by certreq.exe. This implies that the critical flag on certificate extensions can be set for distinct object identifiers and that specific value(s) can be set for these object identifiers using the command-line options.

| **Note** |
| --- |
| With an enterprise CA, template settings will always override certificate extension conflicts that have been set manually. But this is only true for the extensions that are explicitly set through the template. |

The syntax of the command line looks similar to the following example.

```
certutil.exe -setextension <RequestID> <ExtensionName> Flags {Long|Date|String|@InFile}
```

| Option | Description |
| --- | --- |
| Request-ID | The ID of the pending request. It can be determined with the MMC snap-in. |
| ExtensionName | The object identifier of the extension that should be modified. If you are unsure which object identifier is appropriate, dump an existing certificate and find the object identifier from there. |
| Flags | Used to set a certificate extension to critical. flag = 0 means the extension is non-critical; flag = 1 means the extension is critical. |
| {Long \| Date \| String \| @InFile} | Only the @InFile option applies. All other parameters should not be used. The @InFile can be Base64-encoded or a file in hexadecimal format. If the file contents could be interpreted as either Base64 or hexadecimal, hexadecimal is used. To force Base64 interpretation, use the "----- BEGIN -----" and "----- END -----" header. The content of the file depends on the extensions that should be manipulated. |

The following examples illustrate how to use the certutil -setextension command to set the key usage of a pending certificate on the CA to critical and constrain the key usage to certificate and certificate revocation list (CRL) signing.

```
echo 03 02 01 06>ku.txt
certutil.exe -setextension <Request-ID> 2.5.29.15 1 @ku.txt
```

# Appendix 5: ASN.1 File Structure

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

Compiling an ASN.1 BLOB requires advanced knowledge about ASN.1 encoding. When developing your own applications, it is highly recommended that you use an encoding function such as CCertEncodeAltName to create the ASN.1 BLOB. See ICertEncodeAltName (http://go.microsoft.com/fwlink/?LinkID=116167).

Manually developing and parsing such ASN.1 data structure(s) is time-consuming and prone to error due to the overall complexity of the data encoding rules.

The following is a sample ASN.1 BLOB that was used to add a server's GUID and its DNS name to the subject alternative name extension of a certificate. The first line in both boxes is the actual ASN.1 BLOB; the lines following explain per column what the field in the BLOB actually represents. Note that the full ASN.1 BLOB is the following sequence.

30468223636B696E64657230312E6575726F70652E636F72702E6D6963726F736F66742E636F6D
A01F06092B0601040182371901A0120410633035303536343461613133364326338

The following boxes explain in greater detail the breakdown of the ASN.1.

```
30468223636B696E64657230312E6575726F70652E636F72702E6D6963726F736F66742E636F6D
Sequence
  Total length of ASN.1 string
    ContextSpecificPrimitiveType
     Generalname see http://www.ietf.org/rfc/rfc2459.txt for a list of General
      Lenth of DNS name -------------------------------------------------------|
        DNS name ----------------------------------------------------------------|
A01F06092B0601040182371901A0120410633035303536343461613133364326338
ContextSpecificConstructedType
 Generalname see http://www.ietf.org/rfc/rfc2459.txt
  Length of othername -------------------------------------------|
    Other name OID-------|
                          ContextSpecificConstructedType
                           Generalname see http://www.ietf.org/rfc/rfc2459.txt
                            Length of datatype and othername
                              Other name datatype see
                              http://asn1.elibel.tm.fr/en/resources/tags.htm
                                Length of othername
                                  GUID -------------------------|
```

For more information about ASN.1 encoding, see the MSDN article, Example C Program: Converting Names from Certificates to ASN.1 and Back (http://go.microsoft.com/fwlink/?LinkID=116166).

# Appendix 6: Encoding and Decoding with Hexadecimal, Binary, and Base64

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

In many cases, data that must be inserted in a certificate or certificate request must be converted from a human-readable format into a computer-readable format. The most important formats to understand when working with X.509 certificates are hexadecimal, binary, and Base64. Since it is time-consuming to perform data conversions manually, this capability is natively provided with certutil.exe.

As mentioned previously, certutil.exe comes in two versions. In the following table, V1 represents the Windows 2000 version and V2 represents the Windows Server 2003 version.

| From | To Hexadecimal | To Binary | To Base64 |
|---|---|---|---|
| Hexadecimal | n/a | V1: -decodehex<br>V2: -decodehex | n/a |
| Binary | V1: n/a<br>V2: -encodehex | n/a | V1: -encode<br>V2: -encode |
| Base64 | n/a | V1: -decode<br>V2: -decode | n/a |

The table illustrates that you cannot convert data directly from Base64 into hexidecimal, and vice versa. However, you can use binary as an intermediate format to perform this kind of conversion. The reqdccert.vbs script in Appendix 2: Sample Scripts leverages the conversion capabilities of certutil.exe to perform the work on behalf of the administrator, without requiring programming knowledge. For example, to convert an ASN.1 BLOB into a Base64 format (which is required for the certreq.exe INF instruction file), the script uses certutil.exe to automate the following commands.

```
certutil –decodehex <dcname>.asn <dcname>.bin
certutil –encode <dcname>.bin <dcname>.b64
```

# Summary (Advanced Certificate Enrollment and Management)

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

As the white paper thoroughly demonstrates, Windows Server 2003 provides a robust and flexible platform for supporting even the most complex or demanding deployment environments. Although the primary focus of the white paper is to demonstrate the procedures for supporting offline domain controller X.509 certificate enrollment in a disconnected or branch-office deployment, the procedures directly apply to all complex application environments. Administrators may apply the concepts presented to support any type of network and infrastructure services. These include directory servers (ADAM), radius servers (IAS), Web servers (IIS), and other stand-alone applications that require X.509 certificate enrollment or provisioning to provide or enable secure protocols, messaging, or application services.

# Related Links (Advanced Certificate Enrollment and Management)

Updated: June 30, 2010

Applies To: Windows Server 2003 with SP1

See the following resources for further information:

- Requirements for Domain Controller Certificates from a Third-Party CA (Q291010) (http://go.microsoft.com/fwlink/?LinkId=160506)
- How to Enable LDAP over SSL with a Third-Party Certification Authority (http://go.microsoft.com/fwlink/?LinkID=15129)
- Certificate Autoenrollment in Windows Server 2003
- Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure
- Implementing and Administering Certificate Templates in Windows Server 2003
- Key Archival and Management in Windows Server 2003
- Windows Server 2003 PKI Operations Guide

For the latest information about Windows Server 2003, see the Windows Server 2003 Web site (http://go.microsoft.com/fwlink/?LinkID=4302).

# Administering Group Policy with Group Policy Management Console Abstract

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In conjunction with Windows Server 2003, Microsoft has released a new Group Policy management solution that unifies management of Group Policy. Microsoft Group Policy Management Console (GPMC) provides a single solution for managing all Group Policy-related tasks. GPMC lets administrators manage Group Policy for multiple domains and sites within one or more forests, in a simplified user interface (UI) with drag-and-drop support. Highlights include new functionality such as backup, restore, import, copy, and reporting of Group Policy objects (GPOs). These operations are fully scriptable, which lets administrators customize and automate management. Together these advantages make Group Policy much easier to use and help you manage your enterprise more cost-effectively.

For more information about this topic, see Administering Group Policy with Group Policy Management Console.

**In This White Paper**

Introduction (Administering Group Policy with Group Policy Management Console)

Group Policy Management Console Overview (Administering Group Policy with Group Policy Management Console)

Creating and Editing GPOs

Scoping GPOs

Group Policy Inheritance

Delegating Group Policy (Administering Group Policy with Group Policy Management Console)

Reporting on GPO Settings

GPO Details

GPO Operations

WMI Filters

Searching for GPOs

Group Policy Modeling

Group Policy Results (Administering Group Policy with Group Policy Management Console)

Platform Dependencies

GPMC Options

Internet Explorer Enhanced Security Configuration Considerations

Scripting Group Policy-related Tasks

Related Links (Administering Group Policy with Group Policy Management Console)

# Introduction (Administering Group Policy with Group Policy Management Console)

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft Group Policy Management Console (GPMC) is the new tool for Group Policy management that helps administrators manage an enterprise more cost-effectively by improving manageability and increasing productivity. It consists of a new Microsoft Management Console (MMC) snap-in and a set of scriptable interfaces for managing Group Policy.

GPMC simplifies the management of Group Policy by providing a single place for managing core aspects of Group Policy. It addresses the top Group Policy deployment requirements, as requested by customers, by providing the following functionality:

- A user interface (UI) that makes Group Policy much easier to use.
- Backup/restore of Group Policy objects (GPOs).
- Import/export and copy/paste of GPOs and Windows Management Instrumentation (WMI) filters.
- Simplified management of Group Policy-related security.
- HTML reporting of GPO settings and Resultant Set of Policy (RSoP) data.
- Scripting of policy related tasks that are exposed within this tool (not scripting of settings within a GPO).

**GPMC System Requirements**

GPMC helps you manage both Windows 2000 and Windows Server 2003-based domains with the Active Directory® service. In either case, the computer on which GPMC runs must be running one of the following operating systems:

- Windows Server 2003.
- Windows XP Professional with Service Pack 1 (SP1) and the Microsoft .NET Framework. In addition, a post-SP1 hotfix (Q326469) is required. This hotfix updates your version of gpedit.dll to version 5.1.2600.1186, which is required by GPMC. This hotfix is included with GPMC, and GPMC setup will prompt you to install it. However, if the language of GPMC does not match the language of your operating system, GPMC will not install the hotfix, and you will need to separately obtain and install this hotfix. This hotfix will be included in Windows XP Service Pack 2.

**Installing GPMC**

Installing GPMC is a simple process that involves running a Windows Installer (.MSI) package. All necessary files are installed to the \Program Files\GPMC folder.

1. Double-click the **gpmc.msi** package, and click **Next**.
2. Accept the End User License Agreement (EULA), and click **Next**.
3. Click **Close** to complete the installation.

Upon completion of the installation, the Group Policy tab that appeared on the Property pages of sites, domains, and organizational units (OUs) in the Active Directory snap-ins is updated to provide a direct link to GPMC. The functionality that previously existed on the original Group Policy tab is no longer available since all functionality for managing Group Policy is available through GPMC.

To open the GPMC snap-in directly, use any of the following methods:

- Click **Start**, click **Run**, type **GPMC.msc**, and then click **OK**.

- Click the **Group Policy Management** shortcut in the **Administrative Tools** folder on the Start Menu or in the Control Panel.

- Create a custom MMC console - Click **Start**, click **Run**, type **MMC**, and then click **OK**. Point to **File**, click **Add/Remove Snap-in**, click **Add**, highlight **Group Policy Management**, click **Add**, click **Close**, and then click **OK**.

To repair or remove GPMC, use **Add or Remove Programs** in Control Panel. Alternatively, run the gpmc.msi package, select the appropriate option, and click **Finish**.

# Group Policy Management Console Overview (Administering Group Policy with Group Policy Management Console)

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In the past, administrators have been required to use several Microsoft tools to manage Group Policy, such as the Active Directory Users and Computers, Active Directory Sites and Services, and Resultant Set of Policy snap-ins. GPMC integrates the existing Group Policy functionality exposed in these tools into a single, unified console, along with several new capabilities.

Built-in to GPMC is support for managing multiple domains and forests, making it possible for administrators to easily manage Group Policy across an enterprise. Administrators have complete control of which forests and domains are listed in GPMC, making it possible to display only pertinent parts of an environment.

By default, the first time GPMC is started it loads the forest and domain containing the user object logged on to the computer. Administrators can then specify which forests and domains to display. When the console is closed, GPMC automatically saves the last view and will return to that view the next time the user opens that console.

The console tree on the left side of the snap-in contains GPMC's root node **Group Policy Management**. Each forest appears as a sub node of GPMC's root node, and is named after the forest root domain for that forest, pre-pended with the word "Forest." Each forest has either three or four sub nodes of its own: Domains, Sites, Group Policy Modeling, and Group Policy Results. The Group Policy Modeling node is only shown in a forest that has the Windows Server 2003 schema for Active Directory. To perform a Group Policy Modeling analysis, you must also have at least one domain controller that is running Windows Server 2003.

Figure 1 shows GPMC with two forests, Contoso.com and Tailspintoys.com, added to the console. Tailspintoys.com is a Windows 2000 forest so the Group Policy Modeling node is not available.




**Figure 1**

Below is a description of each of the four major sub node types within each forest:

- **Domains:** This node contains sub nodes corresponding to domains within the forest. The domain nodes are named after the DNS names of the domain. Users can choose which domains to display within the console, by right clicking this node and selecting the **Show Domains** context menu option. Domain nodes are always shown as peers of one another, regardless of the actual DNS relationship between them. This is because Group Policy is not inherited across domains. The list in the details pane of Figure 2 shows the name of the domain and the domain controller in that domain used during GPMC operations.



**Figure 2**

- **Sites:** This node contains sub nodes corresponding to sites within the forest. Users can choose which sites to display within the console, by right clicking this node and selecting the **Show Sites** context menu option. All sites are shown as peers of one another. As with forests and domains, the list of sites displayed is preserved for future use when the console is closed. No sites are displayed by default. This is done to speed up console performance by not enumerating a potentially large number of sites in the forest, unless explicitly requested by the user.

- **Group Policy Modeling:** This node allows you to access the Resultant Set of Policy (RSoP) – Planning Mode capabilities of Windows Server 2003. This is a powerful new Group Policy management feature that allows the user to simulate policy settings applied to users and computers via Group Policy before actually applying the policies. You can simulate the policy deployment for any user and computer in the forest. This feature, known as Resultant Set of Policies (RSoP) – Planning Mode in Windows Server 2003, is integrated into GPMC as Group Policy Modeling. This feature requires at least one domain controller in the forest running Windows Server 2003, since the simulation is performed by a service that is only present on domain controllers running Windows Server 2003. Each Group Policy Modeling simulation is displayed as an individual sub node.

- **Group Policy Results:** This node allows you to access the Resultant Set of Policy (RSoP) – Logging Mode capabilities. In contrast to Group Policy Modeling, which is a simulation, Group Policy Results represents the actual resultant set of policy that was applied to a given user and computer. This information is obtained by directly querying the target user/computer. Each sub node represents a different RSoP query for a given user/computer combination. You can only obtain Group Policy Results data from computers that are running Windows XP or Windows Server 2003 and later.

**Managing Multiple Forests**

Multiple forests can be easily added to the console tree:

1. Right-click the root node **Group Policy Management**, and select **Add Forest...**



**Figure 3**

2. Specify the DNS or NetBIOS name of the desired domain in a forest that is not already loaded in GPMC, and click **OK**. If you specify a NetBIOS name, GPMC will attempt to determine the corresponding DNS name and will prompt you with a dialog box to confirm the entry.

The specified forest will appear as a sub node in the console tree and the forest is loaded into the console with the domain that was entered in the **Add Forest** dialog box.

To remove a forest node, simply right-click the node, and then select **Remove**.

By default you can only add a forest to GPMC if there is a 2-way trust with the forest of the user running GPMC. You can optionally enable GPMC to work with only 1 way trust or even no trust. To enable this functionality, clear the **Enable Trust Detection** check box on the **General** tab in the **Options** dialog box of GPMC.

If you need to add a forest to which you have no trust, you must also use the **Stored User Names and Passwords** tool to add credentials for the forest you want to connect to using the procedure below. For this procedure, assume you have no trust to a domain called "mydomain.myforest.contoso.com" and you want to manage that domain.

1. Start the Stored User Names and Passwords tool:

   - In Windows XP, click **Start**, click **Control Panel**, double-click **User Accounts**, click **Advanced**, and then click **Manage Passwords.**

   - In Windows Server 2003, click **Start**, click Control Panel, and then double-click **Stored User Names and Passwords**.

2. Add an entry for the forest containing the domain you want to manage. In this example, add the following entry:

   *.myforest.contoso.com

   For the user name and password, enter the name and password of an account that has the rights to access the domain you want to manage.

3. Start GPMC and disable trust detection by deselecting the **Enable Trust Detection** check box: on the **General** tab in the **Options** dialog box of GPMC.

4. Add the forest to the GPMC console by right-clicking **Group Policy Management** and then clicking **Add Forest**. Enter the name of the domain in the forest that you want to manage.

| **Note** |
| --- |
| When adding forests to which you have no trust, some functionality will not be available. For example, Group Policy Modeling is not available and it is not possible to open the Group Policy Object Editor on GPOs in the untrusted forest. The untrusted forest scenario is primarily intended to enable copying GPOs across forests. Microsoft will support the untrusted forest scenario on a limited basis, and will not be providing escalation support for issues arising from this scenario. |

**Domain Controllers in GPMC**

In each domain, GPMC uses the same domain controller for all operations in that domain. This includes all operations on the GPOs, OUs, security principals, and WMI filters that reside in that domain. In addition, when the Group Policy Object Editor is opened from GPMC, it always uses the same domain controller that is targeted in GPMC for the domain where that GPO is located.

In addition, GPMC uses the same domain controller for all operations on sites. Note that this domain controller is used to read and write information about what links to GPOs exist on any given site, but information regarding the GPO itself is obtained from the domain controller of the domain hosting the GPO.

Group Policy Management Console allows you to choose which domain controller to use for each domain, as well for all sites in a forest in Group Policy Management Console. You can choose from among these four options:

- Use the primary domain controller (PDC) emulator (default choice).

- Use any available domain controller.

- Use any available domain controller that is running a Windows Server 2003 family operating system. This option is useful is you are restoring a deleted GPO that contains Group Policy software installation settings. See *Restore* for more details.

- Use a specific domain controller that you specify.

Right-click the desired domain node and select **Change Domain Controller** to specify a particular domain controller to use for domain operations.

To specify a domain controller to use for operations on sites, right click the Sites node and click **Change Domain Controller**.

In either case, the **Change Domain Controller** dialog box appears. This dialog box provides four options for specifying a domain controller as shown in Figure 4. Selecting the **This domain controller:** radio button activates the list of domain controllers allowing GPMC to target any desired domain controller in a given domain.

| **Note** |
| --- |
| When choosing a domain controller for sites, you have the additional option of choosing which domain to use. Using the domain dropdown list, you can efficiently filter the list of domain controllers, and select a domain controller accordingly. |




**Figure 4**

**Selection of Domain Controllers**

By default, when you add a new domain to the console, GPMC uses the PDC emulator in that domain. For managing sites, GPMC uses the PDC emulator in the user's domain by default.

It is important to consider the choice of domain controller in order to avoid replication conflicts. This is especially important to consider since GPO data resides in both Active Directory and on SYSVOL, and two independent replication mechanisms must be used to replicate GPO data to the various domain controllers in the domain. If two administrators are simultaneously editing the same GPO on different domain controllers, it is possible for the changes written by one administrator to be overwritten by another administrator, depending on replication latency.

To avoid this situation, GPMC uses the PDC emulator in each domain as the default to help ensure that all administrators are using the same domain controller. However, it may not always be desirable to use the PDC. For example, if the administrator resides in a remote site, or if the majority of the users or computers targeted by the GPO are in a remote site, then the administrator may want to choose to target a domain controller at the remote location.

| **Important** |
| --- |
| If multiple administrators manage a common GPO, it is recommended that all administrators use the same domain controller when editing a particular GPO, to avoid collisions in File Replication Services (FRS). |

**Domain Contents Overview**

Within each domain, GPMC provides a policy-based view of Active Directory and the components associated with Group Policy, such as GPOs, WMI filters, and GPO links. The view in GPMC is similar to the view in Active Directory Users and Computers MMC snap-in in that it shows the OU hierarchy. However, GPMC differs from this snap-in because instead of showing users, computers, and groups in the OUs, it displays the GPOs that are linked to each container, as well as the GPOs

themselves.

Each domain node in GPMC displays the following items (see Figure 5):

- All GPOs linked to the domain.
- All top-level OUs and a tree view of nested OUs and GPOs linked to each of the OUs.
- The **Group Policy Objects** container showing all GPOs in the domain.
- The **WMI Filters** container showing all WMI Filters in the domain.




**Figure 5**

The **Group Policy Objects** container shows all of the GPOs for the domain. Each node in this container represents the actual GPO components from Active Directory and SYSVOL that collectively define that GPO. Figure 6 shows an expanded **Group Policy Objects** container for the Contoso.com domain containing 12 GPOs.




**Figure 6**

GPOs are not useful until they are linked to a site, domain, or OU (Scope of Management, or SOM). The settings defined in a GPO can only be applied when the GPO is linked to one or more of these SOMs. The link is not a component of the GPO; it is a component of the SOM to which it is linked. Therefore, the ability to manage links for a given SOM must be delegated on that SOM, not the GPO. In the GPMC tree view, GPO-links on a given SOM are shown as child nodes of that container.

This distinction between GPOs and GPO links was not readily apparent in Windows 2000, and a key goal for GPMC was to

make this distinction clearer. The GPMC user interface distinguishes between GPO-links and GPOs as follows:

- Location in the tree view. Actual GPOs are always shown under the **Group Policy Objects** node for a given domain, whereas links appear as child nodes of a site, domain, or OU. Note that the contents of the result panes for GPOs and GPO-links are identical.

- When you click a GPO-link, a confirmation dialog box is shown by default to indicate that you are viewing a GPO-link, not a GPO.

- The icons for GPO-links have a shortcut icon, to indicate that they are pointers to another object. For example, in Figure 5, note the Default Domain Policy GPO link at the domain level. The icon for this link has a shortcut icon to differentiate it from the icon for the actual Default Domain Policy GPO under the **Group Policy Objects** node (see Figure 6).

- The context menu that appears when you right click in the tree view is different depending on whether you are managing a GPO-link or a GPO. Right clicking a GPO exposes options that are primarily relevant for the actual GPO (such as backup and restore), whereas right clicking a GPO-link exposes options that are relevant to managing the link (such as "Enforced"). Note that some options, such as "Edit" are available on both context menus.

# Creating and Editing GPOs

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A Group Policy object (GPO) is a collection of policy settings that can be applied to a given set of users and/or computers. The process of applying a GPO to a set of users and/or computers is known as "scoping the GPO", which is described later in this white paper.

With GPMC, you can create GPOs using any of the following methods:

- Right-click any domain or OU and choose **Create and Link**. This option simultaneously creates a new GPO and links that GPO to that domain or OU.
- Right-click the **Group Policy Objects** node in any domain in GPMC and click **New**. This will create a new unlinked GPO.
- Using a script. For example, GPMC includes a sample script called CreateGPO.wsf that can be run from the command line to create a new GPO.
- Using any of the methods for copying GPOs. Copy operations are described later in this white paper in the section on GPO operations.

When you create a new GPO, it initially has no settings defined. To edit the settings in a GPO, use Group Policy Object Editor (known previously as the Group Policy snap-in or GPedit). You can open Group Policy Object Editor by right clicking a GPO in GPMC and selecting **Edit**.

# Scoping GPOs

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy is a powerful tool for managing the Windows 2000 (and later) environment. The value of Group Policy can only be realized through properly applying the GPOs to the Active Directory containers you want to manage. Determining which users and computers will receive the settings in a GPO is referred to as "scoping the GPO". Scoping a GPO is based on three factors:

- The site(s), domain(s), or organization unit(s) where the GPO is linked.
- The security filtering on the GPO.
- The WMI filter on the GPO.

This section will discuss how to utilize GPMC to properly manage the scope of a GPO and the Active Directory components you want to manage. Figure 7 shows the GPO scope tab.




**Figure 7**

**Linking GPOs**

The primary mechanism by which the settings in a GPO are applied to users and computers is by linking the GPO to a site, domain, or OU in Active Directory. The location where a GPO is linked is referred to as the Scope of Management, or SOM (also sometime abbreviated as SDOU in previous white papers). There are three types of SOMs: sites, domains, and OUs. A GPO can be linked to multiple SOMs, and a SOM can have multiple GPOs linked to it. A GPO must be linked to a SOM for it to be applied.

| **Note** |
| --- |
| GPOs are not stored on a per-SOM basis, but rather per domain. Thus, if you link a GPO to an OU, the GPO does not actually reside in that OU. A GPO is a per domain object that can be linked anywhere in the forest (although performance issues may exist where cross-domain links are used). One of the goals for GPMC is to make the distinction between links and actual GPOs clearer. |

With GPMC, you can link a GPO to SOMs using any of the following methods:

- Right click a domain or OU node, and choose **Create and Link a GPO here**. This option is analogous to choosing **New** in the old Group Policy user interface prior to GPMC. (This old Group Policy interface was the Group Policy tab on the properties dialog box of a site, domain, or OU in Active Directory Users and Computers or Active Directory Sites and

Services snap-ins.). Although this operation is presented as one action to the user, there are actually two operations taking place. First, a GPO is created in the domain, and second, a link to that GPO is created on the domain or OU from where you started the shortcut menu.

- Right click a site, domain, or OU node, and choose **Link an existing GPO here**. This option is analogous to choosing **Add** in the old Group Policy user interface prior to GPMC. This requires that the GPO already exists in the domain.

- Using drag and drop, drag a GPO from under the **Group Policy objects** node to the OU. This drag and drop functionality can only be performed within the same domain.

| **Note** |
| --- |
| By default, new user and computer accounts are created in the CN=Users and CN=Computers containers. These containers are not actually OUs, so GPOs cannot be directly applied to these containers, however, objects in these containers do inherit GPOs linked to the domain and sites. It is possible to specify a different container in which to place either new user accounts, new computer accounts, or both. This allows you to have greater control for applying GPOs to newly created user and computer objects, before they are moved to their final locations in Active Directory. To specify new locations for user or computer accounts, use Redirusr.exe (for users) or Redircomp.exe (for computers) in the domain you want to manage. The domain administrator can use these tools to specify an OU in which to place all new user or computer accounts when they are created. Redirusr.exe (for user accounts) and Redircomp.exe (for computer accounts) are two new tools included with Windows Server 2003 and are located in the %windir%\system32 directory. For more information about redirecting containers for users and computers, see article Q324949, "Redirecting the Users and Computers Containers in Windows Server 2003 Domains," in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Base link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources. |

**Security Filtering**

By default all Authenticated Users that are located in the SOM (and its children) where a GPO is linked will apply the settings in the GPO.

You can further refine which users and computers will receive the settings in a GPO by managing permissions on the GPO. This is known as security filtering. In order for GPO to apply to a given user or computer, that user or computer must have both **Read** and **Apply Group Policy** permissions on the GPO. By default, GPOs have permissions that allow the **Authenticated Users** group both of these permissions. This is how all authenticated users receive the settings of a new GPO when it is linked to a SOM (OU, domain or site). These permissions can be changed however to limit the scope of the GPO to a specific set of users, groups, and/or computers within the SOM(s) where it is linked.

GPMC simplifies how administrators manage security filtering for a GPO. Without GPMC, administrators were required to use the ACL editor, to manually set the **Read** and **Apply Group Policy** permissions for various security principals (users, computers, and groups) to modify the scope of the GPO. In GPMC, this process is no longer necessary. The administrator can just add or remove security principals in the security filtering section in the **Scope** tab for the GPO or the GPO link. This will automatically set or remove the **Read** and **Apply Group Policy** permissions for that security principal on that GPO. In the example in Figure 7, the security filtering on the “Common Managed Settings” GPO has been modified so that only members of the “Managed Users” group can receive the settings. Note that members of this group that are not located in either the “Engineering – Offsite” or “User Accounts” OUs will not receive the settings in this GPO.

If the administrator needs to access the detailed permissions, they are still available using the ACL editor when you can access using the **Advanced** button on the **Delegation** tab for the GPO. For example, if you need to set permissions to **Deny**, you can do this using the ACL editor. Groups with Deny permissions will appear as having **Custom** permissions on the **Delegation** tab. In general, it is recommended that you avoid the use of Deny permissions for managing Group Policy.

**Linking WMI Filters**

WMI Filters allow an administrator to dynamically determine the scope of GPOs based on attributes (available through WMI) of the target computer. A WMI filter consists of one or more queries that are evaluated to be either true or false against the WMI repository of the target computer. The WMI filter is a separate object from the GPO in the directory. To apply a WMI filter to a GPO, you link the filter to the GPO. This is shown in the WMI filtering section on the Scope tab of a GPO. Each GPO can have only one WMI filter; however the same WMI filter can be linked to multiple GPOs.

When a GPO that is linked to a WMI filter is applied on the target computer, the filter is evaluated on the target computer. If the WMI filter evaluates to false, the GPO is not applied. If the WMI filter evaluates to true, the GPO is applied. Note that client support for WMI filters exists only on Windows XP and later operating systems. Windows 2000 clients will ignore any WMI filter and the GPO is always applied, regardless of the WMI filter.

# Group Policy Inheritance

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy can be applied to users and computers at a site, domain, or OU. GPOs from parent containers are inherited by default. When multiple GPOs apply to these users and computers, the settings in the GPOs are aggregated. For most policy settings, the final value of a given policy setting is set only by the highest precedent GPO that contains that setting. (However, the final value for a few settings will actually be the combination of values across GPOs.) Precedence of GPOs determined by the order of processing for the GPOs. GPOs processed last have highest precedence. GPOs follow the *SDOU* rule for processing; site first, then domain and followed by OU including nested OUs. A nested OU is one that has another OU as its parent. In the case of nested OUs, GPOs associated with parent OUs are processed prior to GPOs associated with child OUs. In this processing order, sites are applied first but have the least precedence. OUs are processed last and have the highest precedence.

There are several Group Policy options that can alter this default inheritance behavior. These options include:

- Link Order – the precedence order for GPOs linked to a given container. The GPO link with Link Order of 1 has highest precedence on that container.
- Block Inheritance – the ability to prevent an OU or domain from inheriting GPOs from any of its parent container. Note that Enforced GPO links will always be inherited.
- Enforcement – (previously known as "No Override") the ability to specify that a GPO should take precedence over any GPOs that are linked to child containers. Enforcing a GPO link works by moving that GPO to the end of the processing order.
- Link Status – determines if a given GPO link is processed or not for the container to which it is linked.

These items are described in more detail below.

If multiple GPOs are linked to the same container and have settings in common, there must be a mechanism for reconciling the settings. This behavior is controlled by the link order. The lower the link order number, the higher the precedence. Information about the links for a given container are shown on the **Linked Group Policy Objects** tab of a given container, as in Figure 8. This pane shows if the link is enforced, if the link is enabled, the status of the GPO, if a WMI Filter is applied, when it was modified, and the domain container where it is stored. An administrator or users that have been delegated permissions to link GPOs to the container can change the link order by highlighting a GPO link and using the up and down arrows to move the link higher or lower in the link order list.




**Figure 8**

GPOs are inherited from parent containers. For example, a GPO linked to an OU will be inherited by child OUs. The **Group Policy Inheritance** tab for a given container shows all GPOs (except for GPOs linked to sites) that would be inherited from parent containers, as shown in Figure 9. The precedence column on this tab shows the overall precedence for all the links that would be applied to objects in that container, taking into account both Link Order and the Enforcement attribute of each link, as

well as Block Inheritance on any SOMs. Note that the Group Policy Inheritance tab does NOT show the impact of GPOs linked to sites, because it is not possible to determine which site would apply, unless a particular target computer is identified.




**Figure 9**

It is possible to prevent containers from inheriting GPOs linked to parent containers by blocking the inheritance on the OU or domain. Blocking inheritance on a container will prevent all GPOs from parent containers from applying to the blocked container, except for GPOs that are marked as Enforced. Administrators can right-click the domain or OU and select **Block Inheritance** from the context menu to set GPO blocking on the container. If inheritance is blocked for an OU or domain, its icon will appear with a blue exclamation mark in the console tree.

An administrator can prevent the settings in a GPO linked to a container from being overwritten by settings linked to GPOs in child containers (which normally would have higher precedence) by setting the GPO link to **Enforced** (formerly known as No Override). This also will prevent the GPO link from being blocked at containers that have been set to **Block Inheritance**. GPO-links that are enforced appear with a gray padlock icon in both the console tree and in the details pane.

A GPO link can be set to **Enabled** to allow it to be processed. If the link is not set to Enabled, processing of the linked GPO is disabled. The GPO link can be either enabled or disabled by right-clicking the link and selecting the **Link Enabled** option. A check beside this option indicates that the link is enabled and will be processed.

# Delegating Group Policy (Administering Group Policy with Group Policy Management Console)

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Group Policy is an essential component of management in Windows 2000 and Windows Server 2003 environments. A portion of Group Policy's power lies in the ability to delegate certain Group Policy tasks to other administrators. For example, the creation, linking, and editing of GPOs are independent permissions that can be delegated separately.

Group Policy delegation covers these areas:

- The ability to create GPOs in a domain.
- Permissions on an individual GPO (for example, read access, edit access).
- Permissions on an individual SOM. There are three permissions related to policy:
    - The ability to link GPOs to a SOM.
    - The ability to perform Group Policy Modeling analyses for objects in that SOM.
    - The ability to collect Group Policy Results data for objects in that SOM.
- The ability to create WMI filters in a domain.
- Permissions on an individual WMI filter (for example, edit access)

One of the key goals for GPMC was to simplify the management of Group Policy related permissions. GPMC does this by abstracting the low-level permissions on the object and managing them as a single unit, based on the task that an administrator wants to perform. Using GPMC, delegation tasks for a given object can be performed by navigating to the Delegation tab for that object.

Using this new model in GPMC, an administrator may never need the traditional ACL Editor. However, for administrators that prefer the flexibility of that ACL editor, it is still available through GPMC by clicking the Advanced button on the Delegation tab.

**Delegating Creation of GPOs**

Creation of GPOs is a right of the Group Policy Creator Owners (GPCO) group by default but can be delegated to any group or user. There are two methods to grant a group or user this right:

- Add the group or user to the Group Policy Creator Owners group. This was the only method available prior to GPMC.
- Explicitly grant the group or user permission to create GPOs. This method is newly available with GPMC.

You can manage this permission using the delegation tab on the **Group Policy Objects** container for a given domain in GPMC. This tab shows the groups that have permission to create GPOs in the domain, including the GPCO group. From this tab, you can modify the membership of existing groups with this permission, or add new groups.

The ability to grant users or groups permissions to create GPOs without using GPCO was added to facilitate the delegation of GPO creation to users outside the domain. Because the Group Policy Creator Owners group is a domain global group, it cannot contain members from outside the domain. Thus, prior to GPMC, this task could not be delegated to members outside the domain.

It is recommended that for users and groups within the domain that you continue to use the GPCO group to grant them GPO creation rights. If you require that users outside the domain have the ability to create GPOs, consider using the following steps.

- Create a new domain local group in the domain ("GPCO – External" for example)

- Grant that group GPO creation rights in the domain
- Create a global group in the external domain ("GPCO in Domain X" for example)
- Add external domain groups or users to that group
- Add the global group ("GPCO in Domain X") to the local group ("GPCO – External")

Adding a user to the membership of GPCO, or granting the user GPO creation permissions directly using the new method available in GPMC, is identical in terms of permissions. Granting a user the ability to create GPOs in the domain does not give the user the ability to edit or delete existing GPOs or the ability to link the GPO to a SOM. Users have full control of GPOs that they created.




**Figure 10**

Figure 10 shows the results pane for the delegation tab of the Group Policy objects container in the Contoso.com domain. Domain Admins can add new delegations or remove existing delegations through the Add and Remove buttons on the delegation tab. The Properties button will display the object properties for the selected user or group.

**Delegating an individual GPO**

GPMC has simplified the ability to delegate rights on an individual GPO. Prior to GPMC, administrators needed to rely on the ACL editor and required an intimate knowledge of the low-level permissions on the GPO. GPMC abstracts this so the administrator can manage the permissions on the GPO at the task level.

There are five categories of Allowed Permissions on a GPO.

- Read
- Edit settings
- Edit, delete, modify security
- Read (from Security Filtering)
- Custom

These permissions are managed using the delegation tab of a GPO, as shown in Figure 11.




**Figure 11**

The options listed above represent predefined combinations of ACLs. The corresponding underlying permissions for each are shown in Table 1 below:

**Table 1**

| Option | Underlying Permissions |
|---|---|
| Read | Allow Read Access on the GPO |
| Edit settings | Allow Read, Write, Create Child Objects, Delete Child Objects |
| Edit, delete, and modify security | Allow Read, Write, Create Child Objects, Delete Child Objects, Delete, Modify Permissions, and Modify Owner. This essentially grants full control on the GPO, except that the "Apply Group Policy" right is not set. |
| Read (from Security Filtering) | This setting cannot be set directly, but appears if the user has Read and Apply Group Policy rights to the GPO (on the **Scope** tab of the GPO). |
| Custom | Any other combinations of rights, such as the use of deny permissions, will show up as **Custom** in the display. Custom rights cannot be set using the Add button, but they can be removed using the Remove button. They can be set only if using the ACL editor directly (which can be started with the Advanced button). |

You can grant users permissions on a GPO using the Add button. This opens the object picker so you can find the desired user or group to set the permission level. You can then set the permission level by selecting either "Read", "Edit", or "Edit, Delete, Modify Security" permissions.

Note that the Apply Group Policy permission, which is used for Security Filtering, cannot be set using the Delegation tab. Setting Apply Group Policy is used for scoping the GPO, so is managed on the Scope tab. Note that when you grant a user Security Filtering on the Scope tab, you are actually setting both the Read and Apply Group Policy rights.

**Delegating Policy-Related Permissions on SOMs**

GPMC enables you to easily set the following three permissions on SOMs that are related to Group Policy:

- The ability to link GPOs to a SOM.
- The ability to perform Group Policy Modeling analysis for objects in that SOM.
- The ability to collect Group Policy Results data for objects in that SOM.

These permissions are viewed and managed using the **Delegation** tab on a given SOM, and selecting the desired option from the permissions dropdown. You can add new delegations or remove existing delegations through the Add and Remove buttons on this tab. The Properties button will display the object properties for the selected user or group. The Advanced button will display the Security properties dialog box (the ACL Editor) for the SOM, from which you can view the actual

permissions for the individual attributes of the SOM.

When you grant permissions using the **Add** button, you will be prompted to specify whether the permission should apply only to the current container, or whether it should be inherited to child containers. Using the **Add** button only grants "Allow" permissions. It is strongly recommended that you avoid using "Deny". However, if you must set a "Deny" permission, you can do so by pressing **Advanced** button.

| **Note** |
| --- |
| The "Applies to" column is not present on sites as they can not have child containers.<br><br>Group Policy Modeling or Group Policy Results permissions cannot be delegated on sites so these options are not available in the permissions dropdown on site nodes. |

**Delegating Linking of GPOs**

The settings in a GPO are applied to users and computers by linking the GPO to a SOM (site, domain, or OU) that contains the user or computer objects, either as a direct child or indirectly through inheritance. The ability to link GPOs to a SOM is a permission that is specific to that SOM. At the lowest level, the permission equates to having read and write access to the gPLink and gPOptions attributes on the SOM. However, with GPMC, there should be no need to manage these attributes individually. GPMC abstracts this permission as a single permission called "Link GPOs." This permission also grants the ability to manage link order, block inheritance, and set the enforced attribute on GPO-links to this SOM.

This permission can be managed using the delegation tab on the SOM in GPMC when you select the **Link GPOs** option in the permission dropdown.

Figure 12 shows the **Link GPOs** permissions on the SOM delegation tab for the Child.Contoso.com domain.




**Figure 12**

**Delegating Group Policy Modeling**

Group Policy Modeling allows the user to simulate the resultant set of policy for objects in a domain or OU. This feature is described in more detail later in the white paper. This section discusses delegation aspects only for Group Policy Modeling.

This feature is only available to Domain Administrators by default but can be delegated to other users or groups. At its lowest level, this delegation equates to granting the user or group the "**Generate Resultant Set of Policy (Planning)**" permission on an OU or domain. This permission is available in any forest that has the Windows Server 2003 schema. Without GPMC, this attribute was available only on the advanced page of the ACL editor for a given domain or OU. GPMC simplifies the management of this permission by exposing this permission directly on the **Delegation** tab for any domain or OU. The administrator selects the "**Perform Group Policy Modeling Analyses**" setting from the Permissions drop-down box. This will display the Name, Applies To, Setting, and Inherited properties for the delegations. Figure 13 shows the **Perform Group Policy Modeling Analyses** permissions on Delegation tab for the Contoso.com domain.




**Figure 13**

**Delegating Group Policy Results**

Group Policy Results allows the user to read RSoP logging data for objects in the domain or OU. Group Policy Results is described in more detail later in the white paper. This section discusses delegation aspects only for Group Policy Results.

By default, only users with local administrator rights on the target computer can remotely access Group Policy results data; however this right can be delegated to other users or groups. Delegation is performed on either a domain or OU. Users with this permission can read Group Policy Results data for any object in that container, and in child containers if the permission is specified to be inherited.

At its lowest level, this delegation equates to granting the user or group the "**Generate Resultant Set of Policy (Logging)**" permission on an OU or domain. This permission is available in any forest with the Windows Server 2003 schema. Without GPMC, this attribute was available only on the advanced page of the ACL editor for a given container. GPMC simplifies the management of this permission by exposing this permission directly on the **Delegation** tab for the domain or OU. The administrator selects the "**Read Group Policy Results Data**" setting from the Permissions drop-down box. This will display the users and groups that have this permission for the domain or OU, if the setting applies to only this container or to the container and all child objects, if the setting is allowed, and if the permission is inherited from a parent container. Figure 14 shows the **Read Group Policy Results data** permissions on the Delegation tab for the Contoso.com domain.




**Figure 14**

**Delegating Creation of WMI Filters**

WMI filters are a new feature in Windows Server 2003. When a new WMI filter is created it is stored in the WMIPolicy container in the domain's System container in Active Directory. It is the permissions applied on this container that govern the rights a user has to create, edit, and delete WMI Filters. This section is limited to delegation aspects of WMI filters. For more information about WMI filters, see the WMI Filters section in this white paper.

There are two levels of permission for creating WMI filters, which can only be delegated by Domain Admins:

- Creator Owner: Allows the user to create new WMI Filters in the domain, but does not grant them permissions on WMI filters created by other users. This permission is analogous to the GPO creation permission.

- Full Control: Allows the user to create WMI filters, and grants them full control on all WMI Filters in the domain, whether they created them or not. There is no corresponding permission level for GPOs because permissions in the sysvol are not inherited, so it's not technically possible to guarantee full control to all GPOs

By default, Domain Admins and Enterprise Admins have Full Control and Group Policy Creator Owners have Creator Owner permissions.

To delegate either permission to a user or group, use the Add button on the delegation tab of the WMI Filters pane. Figure 15 shows the delegation tab for the **WMI Filters** container.




**Figure 15**

An administrator can **Add**, **Remove**, and view **Properties** for WMI Filter delegations from the **Delegation** tab. Selecting **Add** will prompt for a user or group before selecting the permission level (Creator Owner or Full Control) to assign to the user or group. Selecting **Remove** will prompt for confirmation that the delegation should be removed. Selecting **Properties** will display the user or group properties for that object.

**Delegating an individual WMI Filter**

GPMC has the ability to delegate rights on an individual WMI Filter. There are two levels of permissions that can be granted to a user or group on an individual WMI Filter.

- Edit – allows the user or group to edit the WMI Filter.

- Full Control – allows the user or group to edit, delete, and modify security on the WMI Filter.

These permissions are managed using the **Delegation** tab of an individual WMI Filter, as shown in Figure 16.




**Figure 16**

The **Delegation** tab shows the users and groups that have permissions on the WMI Filter, the permission level, and if the permission is inherited from a parent container. Buttons on this tab allow the user to Add or Remove users and groups to the delegation list for the WMI Filter.

Note that all users have **Read** access to all WMI filters. GPMC does not allow this permission to be removed. If the **Read** permission were removed, this could cause policy processing on the client to fail. Therefore, GPMC deliberately doesn't offer the capability to remove this permission.

# Reporting on GPO Settings

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Settings tab of the GPO or GPO link pane in GPMC shows an HTML report that displays all the defined settings in the GPO. Clicking this tab will generate a report of the settings in the GPO. Figure 17 shows a typical report. This report can be generated by any user with read access to the GPO.

Without GPMC, users that did not have write access to a GPO could not read and view the settings in that GPO. This is because the Group Policy Object Editor requires the user to have read and write permissions to the GPO to open it. Some examples of users that might need to read and view but not edit a GPO include security audit teams that need to read but not edit GPO settings, helpdesk personnel that are troubleshooting a Group Policy issue, and OU administrators that may need to read and view the settings from inherited GPOs. With GPMC, these users now have read access to the settings.

The HTML reports also make it easy for the administrator to view all settings that are contained in a GPO at a glance. By clicking the **Show All** option at the top of the report, the report is fully expanded and all settings are shown. Alternatively, administrators can expand and contract individual sections within the report by clicking the heading for each section.

For settings under the Administrative Templates section of the report, you can view a description of the setting by clicking the setting name in the report. This opens a new window with the Explain text for that policy setting.

GPMC also solves some common reporting requirements including the ability to document all the settings in a GPO to a file for printing or viewing. Using a context menu, users can either print the reports, or save them to a file as either HTML or XML. Note that saved reports include the contents of the Settings tab, as well as additional information that that is shown on the Scope, Details, and Delegation pages in the UI.

To view a saved report directly in a Web browser, you must use Internet Explorer 6 or Netscape 7. Netscape 7 does not support functionality that enables you to show or hide data in reports.




**Figure 17**

In addition, GPMC provides similar reports for Group Policy Modeling (Resultant Set of Policy – planning) and Group Policy Results (Resultant Set of Policy –logging). This is described later in the paper in those respective sections.

The reports generated by GPMC display all settings that are contained in a GPO, except:

- Within the IE Maintenance section of reports:
    - The reports indicate only whether Content Ratings and Connectoids are deployed, and do not report the details of those settings.
    - If Preference Mode is specified, this will be indicated, however the new settings that are only available in Preference mode will not be displayed.

- Settings for the following cookie settings which were available in IE 5.5 but not in IE 6 are not displayed:

  Allow per-session cookies (not stored)

  Allow cookies that are stored on your computer

- Within Security Zones and Privacy, the details of customized Java settings, if specified, are not shown. Customized Java settings will appear as "Custom."

- The core information for Wireless and IPSec settings is displayed, however some details for these settings are not displayed.

**Administrative Templates Background**
Administrative templates, (or .adm files), enable administrators to control registry settings using Group Policy. Windows comes with a predefined set of Administrative template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a Group Policy object (GPO). These .adm files are stored in two locations by default: inside the GPO's folder on Sysvol and in the %windir%\inf folder on the local computer.

As new versions of Windows are released, new policy settings are added. In addition to supporting these new settings, each successive version of Windows supports all registry policy settings that were available in earlier versions of Windows. For example, the Windows Server 2003 family supports all registry policy settings available in Windows 2000 and Windows XP.

It is important to understand that .adm files are not the actual settings that are deployed to client operating systems. The .adm file is simply a template file that provides the friendly name for the setting and an explanation. This template file is used to populate the user interface. The settings that are deployed to clients are contained in the registry.pol file inside the GPO. On Windows XP and Windows Server 2003, each registry setting contains a "Supported on" tag that indicates which operating system versions support that policy setting. If a setting is specified and deployed to a client operating system that does not support that setting, the settings are ignored.

Because all successive iterations of .adm files include settings from earlier versions, and because there is no harm if a new setting is applied inadvertently to a computer running an earlier operating system that does not support that setting, it is recommended to always create and edit GPOs from a computer that has the latest .adm files available.

**Administrative Templates and GPMC**
GPMC uses administrative templates (.adm files) to display the friendly names of policy settings in the Administrative Templates section when generating HTML reports for GPOs, Group Policy Modeling, and Group Policy Results. The reports generated by GPMC can display settings based on custom .adm files as well.

GPMC handles administrative templates for GPOs differently than the Group Policy Object Editor. The change in behavior was a deliberate design decision to simplify the behavior with .adm files, to avoid complications with .adm file version conflicts, and to improve performance in GPMC.

To generate a report, GPMC looks by default in the following locations for .adm files:

- %Windir%\inf on the local computer where GPMC is running. If an .adm file is found, it is used, regardless of its timestamp.

- If the .adm file is not found, GPMC looks in the GPO's folder on SYSVOL

The user can specify an alternate path for where to find .adm files using the custom search location option on the Reporting tab of GPMC options. If specified, this takes precedence over the previous locations.

When searching for a given .adm file, GPMC will only use the first .adm file it finds in the listed search order. If there are policy settings in the GPO for which no .adm file can be found, these settings will be displayed in the report in a section called "Extra Registry Settings" which displays the registry keys and values for those settings.

As noted above, GPMC looks by default on the local computer for .adm files first, since the user-specified location is not specified by default. When running GPMC on Windows XP, this means that settings that were not available in Windows XP or Windows 2000 may be displayed in the report as "Extra Registry Settings." This situation only occurs if both of the following are true:

- You have set one of the settings in the Administrative Templates section of the GPO that is new for Windows Server 2003.

- You are generating a report of that GPO on a computer running Windows XP.

  The workaround here is to store the Windows Server 2003 versions of the .adm files somewhere, and specify that location in the custom search location mentioned above.

Unlike the Group Policy Object Editor, GPMC itself never transfers newer versions of .adm files to the SYSVOL. It simply reads the .adm files found using the algorithm described above in order to generate the report. However, when the Group Policy Object Editor is opened (either from GPMC or using other means) .adm files may be transferred to the sysvol, as described below.

**Administrative Templates and Group Policy Object Editor**

- The Group Policy Object Editor uses .adm files to display available policy settings in the Administrative Templates section of a GPO.

- By default it attempts to read .adm files from the GPO (from the Sysvol on the domain controller). Alternatively, the .adm file can be read from the local workstation computer. This behavior can be controlled by a policy setting.

- By default, if the version of the .adm file found on the local computer is newer (based on the time stamp of the file) than the version on the Sysvol, the local version is copied to the Sysvol and is then used to display the settings. This behavior can be controlled by a policy setting.

- If the GPO contains registry settings for which there is no corresponding .adm file, these settings cannot be seen in the Group Policy Object Editor. However, the policy settings are still active and will be applied to users or computers targeted by the GPO.

# GPO Details

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are a variety of details about GPOs that are useful for troubleshooting and other purposes. For example, since the data for a given GPO is stored in both Active Directory and Sysvol, it's important to have details about of these components. As shown in Figure 18, the **Details** tab of any GPO displays this type of information and other useful attributes.




**Figure 18**

The **Details** tab of a GPO shows the following information:

- The domain where the GPO is defined.
- The owner of the GPO. In most cases this is the user that created the GPO. However, if a member of the Domain Admins group creates the GPO, then Domain Admins will be the owner.
- The date and time when the GPO was created.
- The date and time when the GPO was last modified.
- The version number for the user configuration component of the GPO. This is broken out into a version number for the Active Directory component and the Sysvol component. In a healthy GPO, these numbers should match. However, the user version numbers need not match the computer version numbers.
- The version number for the computer configuration component of the GPO. This is broken out into a version number for the Active Directory component and one for the Sysvol component. In a healthy GPO, these numbers should match. However, the computer version numbers need not match the user version numbers.
- The unique ID of the GPO, also known as the GPO GUID.
- The GPO Status. This indicates whether either the user configuration or computer configuration of the GPO is enabled or disabled. It can have four possible values:
    - Enabled.
    - User configuration settings disabled.
    - Computer configuration settings disabled.

- All settings disabled.

**Ensuring consistency of permissions on a GPO**

Each Group Policy object (GPO) is stored partly in the Sysvol on the domain controller and partly in Active Directory. GPMC, Group Policy Object Editor, and the old Group Policy user interface provided in the Active Directory snap-ins present and manage a GPO as a single unit. For example, when you set permissions on a GPO in GPMC, GPMC is actually setting permissions on objects in both Active Directory and the Sysvol.

It is essential that the permissions in the Active Directory component are consistent with the Sysvol component for a given GPO. It is not recommended that you manipulate these separate objects independently outside of GPMC and Group Policy Object Editor. Doing so can potentially cause Group Policy processing on the client to fail, or certain users that should normally have access may no longer be able to edit a GPO. Furthermore, file system objects and directory service objects don't have the same available permissions since they are different types of objects. So in the event of a permissions mismatch, it might not be immediately apparent how to make them consistent.

To help you ensure that the security for the Active Directory and Sysvol component of a given GPO is consistent, GPMC will automatically check the consistency of the permissions of any GPO when you navigate to the GPO using GPMC. If it detects a problem with that GPO, you will be presented with one of the following dialog boxes, depending on whether you have permission to modify security on that GPO:

- "The permissions for this GPO in the SYSVOL folder are inconsistent with those in Active Directory. It is recommended that these permissions be consistent. To change the permissions in SYSVOL to those in Active Directory, click OK."

- "The permissions for this GPO in the SYSVOL folder are inconsistent with those in Active Directory. It is recommended that these permissions be consistent. Contact an administrator who has rights to modify security on this GPO."

If you have permission to modify security on the GPO, you should click OK in this dialog box. This action will reset the permission on the Sysvol component of the GPO to be consistent with the existing permissions on the Active Directory component of the GPO. Note that the information presented on the Delegation tab (and Security Filtering section) for a GPO is based on the Active Directory component of the GPO, so this will simply correctly ensure that the permissions you see in the GPMC UI are being applied to that GPO.

| **Note** |
| --- |
| When running GPMC in a Windows 2000 domain, clicking on either the Default Domain Policy or the Default Domain Controllers Policy opens the dialog box described above. This is the result of a bug in Windows 2000 that is expected to be fixed in Windows 2000 Service Pack 4. The issue occurs because the Access Control List (ACL) on the Sysvol portion of the GPO is mistakenly set to inherit permissions from the parent folder. If you have permissions to modify security on the default GPOs, you should click OK in this dialog box. This will correct the problem by modifying the ACLs on the Sysvol portion to make them consistent with the ACLs on the Active Directory component. In this case, it will remove the inheritance attribute in the Sysvol. |

# GPO Operations

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

GPO operations refer to the ability to **backup** (export), **restore**, **import**, and **copy** GPOs.

Backing up a GPO consists of making of copy of GPO data to the file system. Note that the **Backup** function also serves as the export function for GPOs. Backed up GPOs can be used either in conjunction with Restore or Import operations.

Restoring a GPO takes an existing GPO backup and re-instantiates it back in the domain. The purpose of a restore is to reset a specific GPO back to the identical state it was in when it was backed up. Since a restore operation is specific to a particular GPO, it is based on the GUID and domain of the GPO. Therefore, a restore operation cannot be used to transfer GPOs across domains.

Importing a GPO allows you to transfer settings from a backed up GPO to an existing GPO. You can perform this operation within the same domain, across domains, or across forests. This allows for many interesting capabilities such as staging of a test GPO environment in a lab before importing into a production environment.

Restoring and Importing a GPO will remove any existing settings already in the GPO. Only the settings in the backup will be in the GPO when these operations are complete.

Copying a GPO is similar to an export/import operation only the GPO is not saved to a file system location first. In addition, a copy operation creates a new GPO as part of the operation, whereas an import uses an existing GPO as its destination. These operations are discussed in additional detail in the following sections.

**Backup**

Backing up a GPO places a copy of all relevant GPO data into a specified file system location. The relevant data includes:

- The GPO GUID and domain.
- GPO Settings.
- The Discretionary Access Control List (DACL) on the GPO.
- WMI filter link (but not the filter itself).

A backup operation only backs up components of a GPO that are in the GPO in Active Directory and in the GPO file structure in SYSVOL. The operation does not capture items stored outside the GPO, such as WMI filters, links on a SOM to that GPO, and IP Security policies. These are separate objects with their own set of permissions, and it is possible that an administrator performing either a backup or restore operation may not have the required permissions on those other objects. In addition, in the case of WMI filters and IP Security policies, these can be used by multiple GPOs, so it is not necessarily desirable to restore these objects when restoring a single GPO. To avoid these complications, the backup and restore operations only handle the contents of the GPO itself.

Note that the link to a WMI filter is actually an attribute of the GPO, so the link is included as part of the GPO backup and restore operation; even though the WMI filter itself is not. Customers who want to backup and restore WMI filters should use the import and export capabilities that are available on the WMI filter node in GPMC. Similarly, a backup operation only includes the link to any IP Security policy set in a GPO, because IPSec policies are stored outside the GPO. You must use the **Export Policies** and **Import Policies** commands of the IP Security Policy Management snap-in to back up and restore the IPSec policies themselves.

The backup also contains an XML report of the GPO settings, a date and time stamp, and the user supplied description. The report can be viewed as HTML from within GPMC. Each backup is given a unique ID. This allows for multiple GPOs and/or multiple versions of the same GPO to be backed up to the same file system location. With GPMC, users can identify backups stored in a given file system location based on GPO GUID, the GPO friendly name, the date and time stamp of the backup, the domain name, and the description of the backup. You can manage backups from within GPMC using the Manage Backups dialog box, described later in the "Managing Backups" section.

Administrators can backup one or more GPOs using the following methods:

- Right click a GPO under the **Group Policy objects** node and choose **Back up...** from the context menu.

- Right click one or more GPOs in the **Contents** tab of the **Group Policy objects** node and choose **Back up...** from the context menu. This will backup the selected GPO(s).

- On the **Group Policy Objects** node, right click and choose the **Back Up All...** option. This will backup all GPOs in the domain. This is shown in the figure below.

- Use GPO backup scripts. You can either write your own scripts, or you can use the sample scripts included with GPMC in the GPMC\scripts folder. There are two scripts **BackupGPO.wsf** and **BackupAllGPOs.wsf** that are included with GPMC that you can use to back up GPOs.

When you backup a GPO, you must supply the file system location, and optionally, a description for the backup.

| **Important** |
| --- |
| The administrator should take precautions to secure this location to trusted administrators only. This should be done to protect the GPO backups from malicious or accidental tampering that could compromise the environment if they were used as the basis for a restore or import operation. |




**Figure 19**

Backing up a GPO requires read access to the GPO and write access to the file system location where the backup will be stored.

**Managing Backups**

GPMC provides a way to manage the GPO backups. This feature can be accessed in 2 ways:

- Right-clicking on the **Domains** container and selecting **Manage Backups** from the context menu. This will show backed up GPOs for all domains and forests in the specified file system location.

- Right-clicking on the **Group Policy Objects** container in a domain in GPMC and selecting **Manage Backups**. This will only show GPOs backed up for that domain in the specified file system location.

Either method of starting **Manage Backups** will provide a dialog box for selecting the location of your stored GPO backups. Select a file system location containing valid GPO backups. Figure 20 shows the **Manage Backups** dialog box.




**Figure 20**

In the Manage Backups dialog box, you can sort, delete, restore, and view the backup settings for GPOs. Note the check box that gives the user an option to only show the latest version for each backed up GPO. The dialog box displays the originating domain for the GPO, the display name of the GPO, the timestamp when the GPO was backed up, a user supplied description when the backup was taken, and the GPO GUID.

Upon selecting either the "Restore" or "Delete" buttons, the user is presented with a confirmation dialog box before the operation is performed. Selecting the "View Settings..." button opens a new instance of the user's default web browser which displays a report of the settings in the selected GPO. Selecting "Browse" allows the user to choose the backup location. It is possible to select multiple GPOs for deleting but not for viewing or restoring.

GPO backups can also be managed from script. The sample script QueryBackupLocation.wsf, in the GPMC\Scripts folder, will display information about backups stored in a particular file system location.

**Restore**

Restoring a GPO restores the GPO to a previous state. A restore operation can be used in both of the following cases: the GPO was backed up but has since been deleted, or the GPO is live and you want to roll back to a known previous state. A restore operation retains the original GPO GUID even if the restore is recreating a deleted GPO. This is a key difference between the restore operation and the import or copy operations discussed in later sections of this white paper.

A restore operation replaces the following components of a GPO:

- GPO Settings.
- ACLs on the GPO.
- WMI filter links (but not the filters themselves).

The restore operation does **not** restore links to a SOM (Scope of Management). Any existing links will continue to be used, for example, when restoring an existing GPO to a previous state. However, if the user has deleted a GPO and all links to the GPO, the user must add these links back after restoring the GPO. To facilitate recreating these links, you can view the report in the backup to identify all links in the domain of the GPO.

You can restore GPOs using any of the following methods:

- To restore an existing GPO, right-click the GPO in the **Group Policy objects** container and select **Restore from Backup...** This opens the **Restore Group Policy Object Wizard.**

- To restore a GPO that has been deleted since it was backed up, use the **Manage Backups** dialog box (refer to Figure 20 in the Managing Backups section).

- Use GPO restore scripts. You can either write your own scripts, or you can use the sample scripts included with GPMC in the GPMC\scripts folder. There are two scripts **RestoreGPO.wsf** and **RestoreAllGPOs.wsf** that are included with GPMC that you can use to restore GPOs.

There is also an option to view the settings in the backup before restoring.

The permissions necessary to perform a restore of a GPO are different for restoring an existing GPO or if you are restoring a GPO that has been deleted since it was backed up.

**Table 2**

| GPO State | Permissions needed |
|---|---|
| Existing GPO | The user must have **Edit settings, delete, and modify security** permissions on the GPO in order to restore it, as well as read access to the file system location where the backup is stored. Notice that this does NOT require GPO creation rights. |
| Deleted GPO | The user must have the right to create GPOs in the domain, as well as read access to the file system location where the backup is stored. The person that performs the restore will become the new creator owner for the GPO. |

When restoring a GPO, GPMC will handle the GPO version number differently depending on the type of restore.

- Restore of existing GPO: The GPO version number is incremented by one over the existing version number of the live GPO. This is done to guarantee that clients will apply the GPO.

- Restore of a previously deleted GPO: The GPO version number in the GPO backup is retained in the restored GPO.

**Special Considerations for Software Installation GPOs**

When restoring a deleted GPO that contains Software Installation settings, some side effects are possible depending on the circumstances under which the GPO is restored. This section describes those side effects and how you can avoid them.

When restoring a GPO that has been deleted, it is possible that:

- Cross-GPO upgrade relationships that upgrade applications in the GPO being restored, if any, are not preserved after restore. A cross-GPO upgrade is one where the administrator has specified that an application should upgrade another application, and the two applications are not deployed in the same GPO. Note that cross-GPO upgrade relationships are preserved when applications—in the GPO being restored—upgrade applications in other GPOs.

- If the client has not yet seen that the GPO has been deleted (either because the user has not re-logged on or rebooted since the deletion of the GPO), and the application was deployed with the option to “Uninstall this application when it fall out of scope of management,” then the next time the client logs on:

    - Published apps that the user has previously installed will be removed.

    - Assigned applications will be uninstalled before re-installation.

This behavior occurs because when the GPO is restored, the object in Active Directory that represents the application (the "deployment object") is assigned a new GUID. Because the GUID is different than the GUID of the original deployment object, Windows interprets this as a different application.

The solution is to restore the GPO using the original GUID for the deployment object. However, because the GUID is controlled by Active Directory, the only way to re-use the original GUID is to re-animate the tombstone of the deleted deployment object. Tombstone re-animation is a new feature of Windows Server 2003.

GPMC automatically attempts to re-animate these objects when performing a restore. Re-animation will succeed if each of the following is true:

- GPMC is targeting a domain controller that is running Windows Server 2003.

- The user performing the restoration has permissions to re-animate objects in the domain. By default, only Domain Admins and Enterprise Admins have this permission. You can delegate this right to additional users at the domain level using the ACL editor.

- The time elapsed between deletion and restoration of the GPO does not exceed the tombstone lifetime specified in Active Directory.

If re-animation fails, a new deployment object will be created and the previously mentioned side effects will occur.

Therefore, as a general rule, if you deploy software using Group Policy, it is recommended that you perform the restoration of GPOs that contain application deployments using a domain controller running Windows Server 2003 and that you grant the tombstone re-animation right to the users who will be performing restoration of those GPOs.

**Considerations for Domain Rename**
GPO Backups taken prior to a domain rename operation cannot be restored after a renaming a domain. If you are renaming a domain, be sure to backup your GPOs as soon as you complete the rename procedure. For more details on domain rename, see the documentation on Domain Rename at
http://www.microsoft.com/windowsserver2003/downloads/domainrename.mspx.

**Import**
The Import operation transfers settings into an existing GPO in Active Directory using a backed up GPO in the file system location as its source. Import operations can be used to transfer settings across GPOs within the same domain, cross-domain in the same forest, or cross-domain in a separate forest.

Import operations are ideally suited for migrating Group Policy across environments where there is no trust. For example, if you have separate test and production forests with no trust, you can use import operations to migrate GPOs that you develop and verify in the test forest to the production environment. You need to be able to access a file system location from the production environment where the backed up GPOs are stored. Depending on the settings in the GPO, you may also want to use a migration table (described later in this paper) to effectively transfer the settings.

The target of an import operation is an existing GPO. The target GPO's ACLs, links to any SOMs, and link to a WMI filter are not modified during an Import operation. Importing a GPO requires edit permissions on the target GPO.

Import operations can be performed using either of the following methods:

- Right click a GPO under the Group Policy Objects node and click Import Settings. This starts a wizard that guides you through the process of selecting a backup and optionally specifying a migration table if appropriate.

- Using either the **ImportGPO.wsf** or **ImportAllGPOs.wsf** scripts that are included with GPMC.

**Copy**
A copy operation transfers settings using an existing GPO in Active Directory as the source and creates a new GPO as its destination. A copy operation can be used to transfer settings to a new GPO either in the same domain, cross-domain in the same forest, or cross-domain in a separate forest. Since a copy operation uses an existing GPO in Active Directory as its source, trust is required between the source and destination domains, or you must use the **Stored User Names and Passwords** tool as described earlier in the section "Managing Multiple Forests," to gain access to the untrusted forest. Copy operations are ideally suited for migrating GPOs between production environments, as well as for migrating Group Policy that has been staged and tested in a test domain or forest to a production environment.

Copying a GPO requires GPO creation rights in the target domain and read access to the source GPO.

When copying a GPO, the administrator has two options for specifying the DACL to use on the destination GPO:

- Use the default permissions that are used when creating new GPOs.

- Preserve the DACL from the source GPO. Note that if a migration table (described in the next section) is specified during the copy operation, the migration table will apply to any security principals in the DACL.

There are some minor differences between same-domain and cross-domain copy operations.

- When copying a GPO within the same domain, any link to a WMI filter is preserved. However, when copying a GPO to a new domain, the link is dropped because WMI filters can only be linked to GPOs within the same domain.

- When copying a GPO within the same domain, any link to an IPSec policy is preserved. However, when copying a GPO to a new domain, the link is dropped because it is unknown whether a corresponding IPSec policy exists in the new domain, and it is not always desirable to link to the original IPSec policy in the original domain.

- When copying a GPO within the same domain, the user is presented with a simple choice of choosing the DACL. However, for cross-domain copy operations, GPMC presents a wizard to facilitate the operation. In the cross-domain case, the user has the ability to specify a migration table (described below).

Copy operations can be performed using any of the following methods:

- Right-click the source GPO and choose copy, and then right-click the **Group Policy Objects** container in the desired destination domain and choose the paste option.

- Using drag and drop to drag the source GPO to the **Group Policy Objects** container in the destination domain.

- Using the **CopyGPO.wsf** command-line script that is included with GPMC.

**Using migration tables to facilitate cross-domain import and copy operations**
This section discusses the use of migration tables to facilitate cross-domain import and copy operations.

The key challenge when migrating GPOs from one domain to another is that some information in the GPO is actually specific to the current domain where the GPO is defined. When transferring the GPO to a new domain, it may not always be desirable, or even possible, to use the exact same settings. Items that can be specific to a domain include references to security principals (users, groups, and computers) and UNC paths. The solution is to modify these references in the GPO that are domain-specific, during the import or copy operation, so that the settings in the destination GPO are written with the appropriate information for the destination domain. GPMC supports this capability using migration tables.

A migration table is a file that maps references to users, groups, computers, and UNC paths in the source GPO to new values in the destination GPO. The migration table consists of one or more mapping entries. Each mapping entry consists of a source type, source reference, and destination reference. If you specify a migration table when performing an import or copy, each reference to the source entry will be replaced with the destination entry when writing the settings into the destination GPO. Note to use a migration table, the security principals specified in the destination entries of the migration table must already exist when you perform the import or copy operation.

The migration table will apply to any references in the settings within a GPO, whether you are performing an import or copy operation. In addition, during a copy operation, if you choose the option to preserve the ACL on the GPO, the migration table will also apply to both the ACL on the GPO and the ACLs on any software installation settings in the GPO.

Migration tables have their own file name extension, .migtable. The mapping information is stored as XML. You can create and edit migration tables using the Migration Table Editor (MTE) described later in this section. A sample migration table file is included with GPMC installation at "%ProgramFiles%\GPMC\Scripts" as SampleMigrationTable.migtable.

**Which settings are impacted by Migration Tables:**
There are two types of settings that may not transfer well across domains:

1. Users, groups, and computers (security principals) that are referenced in the settings of the GPO, and optionally for copy operations, in the DACL for the GPO itself and the DACL for any software installation settings in the GPO. Security principals don't transfer well for several reasons:

    - Domain local groups are never valid in external domains.

    - If there is no trust between source and destination domains, then none of the security principals in the source domain will be available in the destination domain.

    - Even in situations with trust between source and destination domains, it may not always be appropriate to use the exact same group in the new domain.

2. UNC paths. When you are migrating GPOs from test to production, users in the production environment may not have access to paths in the test environment, and vice versa.

The following items can contain security principals and can be modified using a migration table:

- Security policy settings of the following types:
    - User rights assignment
    - Restricted groups
    - System services
    - File system
    - Registry
- Advanced folder redirection policy settings.
- The GPO DACL, if you choose to preserve it during a copy operation.
- The DACL on software installation objects – only preserved if the option to copy the GPO DACL is specified.

The following items can contain UNC paths, which can be updated to new values as part of the migration process:

- Folder redirection policy settings.
- Software installation policy settings (for software distribution points).
- Pointers to scripts deployed through Group Policy (startup, shutdown, logon, and logoff) that are stored outside the GPO. Note that the script itself is not copied as part of the GPO copy operation, unless the script is stored inside the source GPO.

Note that security principals and UNC paths referenced in other types of settings not mentioned above, such as Administrative Templates and Software Restriction Policies, cannot be mapped using migration tables.

**Options for specifying migration tables**

Migration tables are specified when performing import and copy operations. There are three options for using migration tables with import and copy:

- Do not use a migration table. This option will copy all references in the source GPO exactly as is.
- Use a migration table. This option will map any references in the GPO that are in the migration table. References in the source GPO that are not contained in the migration table will be copied as is.
- Use a migration table exclusively. This option requires that all references to security principals and UNC paths be specified in the migration table. If a security principal or UNC path is referenced in the source GPO and is not included in the migration table, the import or copy operation will not proceed. This option is useful to ensure you have accounted for all security principals and UNC paths in your migration table.

In addition, cross-domain copy operations will apply the migration table to the DACL on the GPO (and any software installation settings) if you choose the option to “Preserve or migrate the existing permissions” and you specify a migration table.

When performing a copy or import, the wizard will scan the source GPO or backup to determine if there are any references to

security principals or UNC paths in the GPO. If there are, you will have the opportunity to specify an existing migration table or to create a new migration table. Note that during cross-domain copy operations, if you specify the option to "Preserve or migrate the permissions on the GPO", the wizard will always present the opportunity to specify a migration table, since a DACL by definition contains security principals.

**Creating Migration Tables**

The Migration Table Editor (MTE) in GPMC allows the user to create and edit migration tables .You can open the MTE using any of the following methods:

- From either the Import Settings or Cross Domain Copying wizards
- By right clicking the **Group Policy Objects** node in any domain and choosing **Open Migration Table Editor**.
- From the **Domains** node in any forest and choosing **Open Migration Table Editor**.
- By double clicking an existing .migtable file in Explorer.
- By manually running the program at %programfiles%\gpmc\mtedit.exe.

A migration table consists of one or more mapping entries, as shown in Figure 21 below. Each mapping entry has three pieces of information: Source type, source name, and destination name.

- Source Type: The type of domain-specific information in the domain for the source GPO. Migration tables support the following types:
    - User
    - Computer
    - Domain Local Group
    - Domain Global Group
    - Universal Group
    - UNC Path
    - Free Text or security identifier (SID). Note: This category is only for use with security principals that are specified as free text and raw SIDs.
- Source Name: The exact name of the User, Computer, Group or UNC Path referenced in the source GPO. The type of the source name in the source GPO or backup must match the source type specified in the migration table. Security principals can be specified using any of the following formats:
    - UPN. For example, "someone@example.com".
    - SAM. For example, "example\someone".
    - DNS. For example, "example.com\someone".
    - Free text. For example, "someone". You must specify the type as Free Text or SID in this case.
    - SID. For example, S-1-11-111111111-111111111-1111111111-1112. You must specify the type as Free Text or SID in this case.
- Destination Name: The destination name specifies how the name of the User, Computer, Group or UNC Path in the

source GPO should be treated upon transfer to the destination GPO. There are four choices:

- *Same as source*. Copy without changes. Equivalent to not putting the source value in the migration table at all. If specified, the destination entry will appear as <Same As Source> in the MTE.
- *None*. Removes the User, Computer or Group from the GPO. If specified, the destination entry will appear as <None> in the MTE. This option cannot be used with UNC paths.
- *Map by relative name*. For example, map SourceDomain\Group1 to TargetDomain\Group1. If specified, the destination entry will appear as <Map by Relative name> in the MTE. This option cannot be used with UNC paths.
- *Explicitly specify value*. In the destination GPO, replace the source value with the exact literal value specified in the migration table. The same format for destination entries is supported as for source names, except that raw SIDs cannot be specified.




| Source Name | Source Type | Destination Name |
|---|---|---|
| testdomain3.com\Computer3$ | Computer | testdomain4.com\Computer4$ |
| Computer5$@testdomain5.com | Computer | Computer6$@testdomain6.com |
| testdomain1\Computer1$ | Computer | testdomain2\Computer2$ |
| testdomain1\Group02 | Domain Global Group | <Map by Relative name> |
| testdomain1.com\Group1 | Domain Global Group | testdomain2.com\Group2 |
| Group01@testdomain1.com | Domain Global Group | <None> |
| testdomain3\Group3 | Domain Local Group | Group4@testdomain4.com |
| S-1-1-11-111111111-111111111-1111111111-11 | Free Text or SID | testdomain1\Group02 |
| Group05 | Free Text or SID | Group06 |
| EVERYONE | Free Text or SID | <None> |
| S-1-1-11-111111111-111111111-1111111111-11 | Free Text or SID | <Same As Source> |
| \\domain01\Share1\Script1.bat | UNC Path | \\domain02\share2\Script2.bat |
| \\domain01.test.com\share1\FolderRedirect1 | UNC Path | \\domain02.test.com\share2\FolderRedirect2 |
| \\domain01\deploy1 | UNC Path | \\domain02\deploy2 |
| Group03@testdomain3.com | Universal Group | <Same As Source> |
| Group5@testdomain5.com | Universal Group | Group6@testdomain6.com |
| testdomain1\User1 | User | testdomain2\User2 |
| User5@testdomain5.com | User | User6@testdomain6.com |
| testdomain3.com\User3 | User | testdomain4.com\User4 |

**Figure 21**

To add entries in the migration table using MTE, you can either manually type in the appropriate information or you can scan an existing GPO or GPO backup to auto-populate existing security and UNC paths in the source column. When using the auto-population methods, the destination will be set to <Same as Source> by default. The user must then adjust the destination entries in the MTE.

GPMC users can also create migration table files using the **CreateMigrationTable.wsf** script included with GPMC. This script performs an auto-populate operation that is available in the MTE. The user must then edit the file manually to set the target ACLs and paths. Alternatively, you can create a migration table manually using any XML editor.

**Scenarios for Copy and Import**

There are two key scenarios GPMC enables with the copy and import operations .

- Test-to-production migration.
- Production-to-production migration.

### Test-to-production migration

The test-to-production scenario is usually a multi-forest environment with separate forests for test and production which are often physically isolated. In this scenario, all groups, users, and computers in all domains will likely need to be mapped.

Sometimes there is a trust between the forests while other times there is not. GPMC handles either scenario, and typically a

migration table will be required. In cases with no trust, all security principals and UNC paths must be mapped using a migration table. Also, if there is no trust, you must either use import, or if using copy, you must supply alternate credentials using the **Stored User Names and Passwords** tool as described earlier in the section "Managing Multiple Forests," to gain access to the untrusted forest.




Assume for this scenario that the administrator wants to copy or import GPO X in domain B in the TEST forest to Domain E in the PRODUCTION forest. GPO X contains security settings such as the Log on locally setting under User Rights Assignment that reference security principals throughout the TEST forest. These security principals may reside in the same domain as the GPO or potentially in different domains in the TEST forest. In the drawing above, these security principals are represented as A\user, A\group, etc, where "A" is the domain, and "user" is the relative name of the security principal.

In this scenario, when copying GPO X to the new domain, the administrator will want to map all security principals that are referenced in the GPO. At a minimum, the domain prefix will need to change and potentially the relative names as well.

This scenario can be handled by using a migration table.

**Production-to-production migration**

The production-to-production scenario has all domains in production environments and there is a trust between the source and target domains. The diagram below assumes A, B, and C are all production domains in the same forest. Note that you could also have this scenario across forests provided there was a trust in place.




Assume for this scenario that the administrator wants to copy GPO X from production domain B to production domain C. There are potentially two alternative behaviors that are desired in this scenario:

- Retain all references to security principals in the copied GPO. Functionally, this would be equivalent to cross-linking GPO X to domain C (in terms of the settings that are applied from the GPO). This scenario is handled when you choose no migration table.

- Map some or all of the security principals to new values. For example for security principals defined in domain B, it may make sense to map them to domain C when copying the GPO from B to C. In particular, and domain local groups must be mapped. However, it probably makes sense to leave references defined in domain A as is. This scenario can also be handled using a migration table.

**Creating a Staging Environment**

Most administrators will want to deploy Group Policy to a test environment before introducing it to their production environment. Deploying Group Policy ideally requires the following steps:

- Gather management requirements for the target platform.
- Determine what management tasks can be done through Group Policy.
- Test Group Policy in staging environment.
- Use GPMC to migrate Group Policy to production environment.

GPMC can be used for planning, creating, testing, and migrating the Group Policy environment.

For best results, you should have a test environment that closely resembles the production environment, in terms of domains, OUs, GPOs, and groups. Using this test environment, you can test and validate changes to your policy deployment in a controlled, isolated manner that does not impact production users. Once the change has been validated, you can use the import and/or copy functions to migrate the GPOs to production.

Creating a staging environment that closely resembles the production environment is essential for this task. GPMC facilitates this by providing two sample scripts that help you create a new test environment, based on an existing environment.

- **CreateXMLFromEnvironment.wsf**: This script uses the information in a live domain to generate an XML file and a set of GPO backups that represent the policy information for that domain. The XML file captures information such as OU structure, groups and users, GPOs and the settings contained in them, links to GPOs, security on GPOs, and WMI filters. By running this script against a production domain, you can capture the essential policy information for that domain for later re-use.

- **CreateEnvironmentFromXML.wsf**: This script populates a domain with policy information such as OU structure, groups, and users, GPOs and the settings contained in them, links to GPOs, security on GPOs, and WMI filters using an XML file and a set of GPO backups referenced in the XML. The XML file required as the input for this script can be generated using the previous script. By using this second script in conjunction with the XML file previously generated, you can replicate the contents of one domain to another.

For more details on using these two scripts, see the chapter, “Staging Group Policy Deployments” in the Windows Server 2003 Deployment Kit.

# WMI Filters

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

WMI Filters are a new feature in Windows Server 2003 and Windows XP. WMI Filters allow an administrator to dynamically determine the scope of GPOs based on attributes of the target computer. This provides the administrator with the potential to dramatically extend the filtering capabilities for GPOs well beyond the previously available security filtering mechanism.

A WMI filter is a separate object that can be linked to a GPO. When the GPO is applied on the target computer, the filter is evaluated on the target computer. A WMI filter consists of one or more queries that are evaluated against the WMI repository of the target computer. If the total set of queries evaluates to false, the GPO is not applied. If all queries evaluate to true, the GPO is applied. Each query is written using the WMI Query Language (WQL), which is a SQL-like language for querying the WMI repository.

Note that client support for WMI filters exists only on Windows XP and later operating systems. Windows 2000 clients will ignore any WMI filter and the GPO is always applied, regardless of the WMI filter.

Each GPO can have only one WMI filter. However, the same WMI filter can be linked to multiple GPOs. Like GPOs, WMI filters are per domain objects.

Figure 7 in the Scoping GPOs section shows a GPO scope pane with a link to the GPO for the "XP Systems" WMI filter. WMI filters are only available in domains that have the Windows Server 2003 configuration. Although none of the domain controllers need to be running Windows Server 2003, you must have run **ADPrep /DomainPrep** in this domain. ADPrep is a utility included on the Windows Server 2003 CD and must be run before upgrading an existing Windows 2000 domain to Windows Server 2003. If ADPrep /DomainPrep has not been run in a Windows 2000 domain, the **WMI Filters** node will not be present, and the **GPO scope** tab will not have a WMI filters section.

The user can create new WMI filters from the **WMI Filters** container in the GPMC console. Right-clicking either the **WMI Filters** container or the Contents pane for this node allows the user to create a new WMI filter or to import a filter that was previously exported. Selecting **New** will present the user with the dialog box in Figure 22.




**Figure 22**

The user enters a name for the WMI filter, an optional description, and then one or more WQL queries using the Add button. Note that for each query, you must specify the WMI namespace where the query is to be executed. The default namespace is root\CIMv2, which should be appropriate for most scenarios. When the user presses the **Save** button, the query syntax is checked before the WMI filter can be saved.

There are three methods for linking a WMI filter to a GPO:

1. On the **Scope** tab of the GPO, use the **WMI filtering** dropdown to select a WMI filter to link to the GPO.

2. On the **General** tab of a WMI filter, right-click the **GPOs that use this WMI Filter** section and select **Add** (shown below in Figure 23). Selecting **Add** will present the user with a list of GPOs in the domain to which the user can link the WMI filter. You can only link to one GPO at a time from this pane.

3. Using drag and drop, drag a WMI filter onto a GPO.




**Figure 23**

There is no option to link to a GPO in another domain because WMI filters can not be linked to GPOs in a different domain. Note that the ability to link a WMI filter to a GPO requires edit access to the GPO, as the link to the WMI filter is an attribute of the GPO.

The General tab on the WMI filter pane allows the user to edit a WMI filter. If the user does not have write access to the WMI filter, the **Edit Filter** button is grayed out. The user can modify the WMI Filter name, description, and the WMI query from this dialog box.

# Searching for GPOs

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

GPMC provides extensive capabilities to search for GPOs within a domain or across all domains shown in a forest. This search feature allows you to search for GPOs based on:

- Display name of the GPO.
- Whether or not a specific domain contains links to the GPO.
- The permissions set on the GPO.
- The WMI filter that is linked to the GPO.
- The *type* of policy settings that have been set in the User Configuration or Computer Configuration in the GPO, such as folder redirection or security settings. Note that you cannot search based on the individual settings configured in a GPO.
- GUID of the GPO.

Figure 24 shows the GPO Search dialog box.




**Figure 24**

The GPO search by name and by GUID functions allow the user to search for a GPO by the GPO display name or by GUID associated with the GPO. The search by User Configuration or Computer Configuration allows the user to find GPOs that contain certain types of Group Policy settings. Refer to Table 3 for a list of the Group Policy settings allowed in this search feature.

The search by security group feature allows the user to find GPOs that have certain permissions applied to them. You can search for GPOs that either explicitly have these permissions (or explicitly do not have these permissions) or that have these permissions effectively applied to them (or not). An explicit permission on a GPO means the security principal is directly referenced in the ACL on the GPO. An effective permission means the security principal has permissions on the GPO either as

the result of an explicit ACE, or because of its group membership. These factors combine to give a security principal the merged or effective set of permissions they have on the GPO.

Searching by GPO-Link allows the user to find unlinked or cross-domain linked GPOs. For example, if you are searching for GPOs in a given domain and you perform a search where GPO links do NOT exist in that domain, this search type will return the list of unlinked GPOs.

Searching by GPOs that link to a WMI filter allows the user to find all GPOs that link to a specified WMI filter.

Table 3 summarizes the GPO search actions and how they can be used.

**Table 3**

| Search Item | Search Condition | Value |
|---|---|---|
| GPO name | Contains<br>Does not Contain<br>Is Exactly | GPO Display name |
| GPO-Link | Exist in<br>Does not Exist in | Domain name(s)<br>[All Sites] |
| Security Group | Has this explicit permission<br>Does not have this explicit permission<br>Has this effective permission<br>Does not have this effective permission | Apply Settings<br>Edit Settings<br>Edit Settings, Delete, Modify Security<br>Read Settings |
| Linked WMI filter | Is<br>Is not | WMI Filter name |
| User Configuration | Contains<br>Does not Contain | Folder Redirection<br>Internet Explorer Branding<br>Registry<br>Scripts<br>Software Installation |
| Computer Configuration | Contains<br>Does not Contain | EFS Recovery<br>IP Security<br>Microsoft Disk Quota<br>QoS Packet Scheduler<br>Registry<br>Scripts<br>Security<br>Software Installation<br>Wireless Group Policy |
| GPO GUID | Equals | GUID |

**Note**

When searching based on user or computer configuration, if a setting is enabled, and then all the settings in that extension are removed, there can be false-positive search for certain types of settings. This happens because the GPO has the extension listed as active. The extensions with this behavior are Security Settings, Software Installation, Folder Redirection, Internet Explorer Maintenance, and Encrypting File System (EFS).

# Group Policy Modeling

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 has a powerful new Group Policy management feature that allows the user to simulate a policy deployment that would be applied to users and computers before actually applying the policies. This feature, known as Resultant Set of Policy (RSoP) – Planning Mode in Windows Server 2003, is integrated into GPMC as Group Policy Modeling. This feature requires a domain controller that is running Windows Server 2003 in the forest, because the simulation is performed by a service that is only present on Windows Server 2003 domain controllers. However, with this feature, you can simulate the resultant set of policy for any computer in the forest, including those running Windows 2000.

In Figure 25, note that the Contoso.com Windows Server 2003 forest has a **Group Policy Modeling** container in the tree pane and the Tailspintoys.com Windows 2000 forest does not have this container.




**Figure 25**

The contents tab on the **Group Policy Modeling** node displays a summary of all Group Policy Modeling queries that the user has performed. This is shown in Figure 25. For each query, GPMC shows the following data:

- Name – This is the user-supplied name of the modeling results.
- User – This is the user object (or the OU where the user object is located) that forms the basis of the modeling query.
- Computer – This is the computer object (or the OU where the computer object is located) that forms the basis of the modeling query.
- Last refresh time – This is the last time the planning query was refreshed.

The Group Policy Modeling Wizard can be opened from the Group Policy Modeling container, the domain node, or from any OU. When the Group Policy Modeling Wizard is started from one of the SOM containers, the wizard automatically passes the SOM data to the wizard and pre-populates the User and Computer Selection page of the wizard.

| **Note** |
| --- |
| For users that are familiar with the RSoP MMC snap-in in Windows Server 2003, the Group Policy Modeling Wizard is a newer version of the RSoP wizard, running in Planning mode. Because all RSoP functionality provided by the RSoP MMC snap-in is included in GPMC, along with new functionality such as HTML reporting of RSoP data, it is recommended that users access all RSoP functionality primarily through GPMC, rather than the standalone RSoP MMC snap-in. |

Figure 26 shows the Group Policy Modeling Wizard's Summary of Selections dialog box prior to running the modeling analysis. The settings in the summary pane are the answers supplied by the user while running the modeling wizard.




**Figure 26**

Once the user completes the Group Policy Modeling Wizard, a new node in the console is created to display the results. These nodes are persistent across GPMC console sessions. The user must manually remove any Group Policy Modeling nodes that are no longer desired.

For a given Group Policy Modeling query, the node contains three tabs as shown in Figure 27.

- Summary - this contains an HTML report of the summary information including the list of GPOs, security group membership, and WMI filters.
- Settings - this contains an HTML report of the simulated policy settings that would be applied in this simulation.
- Query - this lists the parameters that were used to generate the query.

Using the context menu on this new node, the user is able to:

- Save the results report to the file system. This saves the contents of both the Summary and Settings tabs as a single file (either HTML or XML).
- Re-run the query. Choosing the option to re-run the query will re-run the simulation and re-generate the data displayed in the report.
- Create a new query using the original as a template.
- Start the RSoP MMC snap-in, by choosing the "Advanced View" option. The RSoP snap-in includes the same data that is shown in the HTML report, but also shows precedence information. For example, if three GPOs set the same setting, only 1 GPO will actually set that setting. The HTML will tell you the final value and which GPO actually set it, whereas the traditional RSoP snap-in will also identify all GPOs that attempted to set that setting and the corresponding value of that setting. This is shown on the precedence tab when you double click a setting in the RSoP MMC snap-in.

The summary tab (shown in Figure 27) shows a summary of the RSoP data for the user configuration and computer configuration. For both sections, the following information is shown:

- General information including the name of the user, computer, and/or SOM for which the RSOP data was gathered.
- A list of GPOs that are in scope for the given user, computer, computer, and/or container, and the SOM to which each GPO was linked. This includes a list of GPOs that would be applied, as well as GPOs that were in scope, but would not be

applied on the target.

- Simulated security group membership of the targeted user and/or computer.
- List of WMI filters that are linked to the GPOs and whether they were assumed in the query to be true or not.




**Figure 27**

The settings tab (shown in Figure 28) displays a report of the final value of all policy settings that would be applied, and the GPO (for example, "Winning GPO") that would be responsible for setting each value.




**Figure 28**

The query tab displays the parameters of the query entered by the user that were used to generate the data. This tab includes data that the user entered in the wizard to generate the query such as:

- Last time the query was refreshed.
- Domain controller on which the simulation was run.
- User name or SOM name for user settings.
- Computer name or SOM name for computer settings.

- If slow link processing was simulated.
- The site that the computer is assumed to be in, for this simulation.
- Whether Loopback processing was assumed, and if so, the mode (“none”, “merge mode”, or “replace mode”).
- Simulated alternate user location.
- Simulated alternate computer location.
- Simulated security group membership of user object.
- Simulated security group membership of computer object.
- WMI filters that were assumed to be true for the computer object
- WMI filters that were assumed to be true for the user object

# Group Policy Results (Administering Group Policy with Group Policy Management Console)

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This feature allows administrators to determine the resultant set of policy that was applied to a given computer and (optionally) user that logged on to that computer. The data that is presented is similar to Group Policy Modeling data, however, unlike Group Policy Modeling, this data is not a simulation. It is the actual resultant set of policy data obtained from the target computer. Unlike Group Policy Modeling, the data from Group Policy Results is obtained from the client, and is not simulated on the DC. The client must be running Windows XP, Windows Server 2003 or later. It is not possible to get Group Policy Results data for a Windows 2000 computer. (However, with Group Policy Modeling, you can simulate the RSoP data).

| **Note** |
| --- |
| Technically, a Windows Server 2003 DC is not required to access Group Policy Results. However, by default, only users with local admin privileges on the target computer can remotely access Group Policy Results data. This can be delegated to additional users (as previously described), however, the ability to delegate RSoP data is only available in Active Directory forests that have the Windows Server 2003 schema (for example, you have run ADPrep /ForestPrep) in that forest. |

Each Group Policy Results query is represented by a node in the tree view under the Group Policy Results container. Each node has three tabs:

- Summary – this is analogous to the information shown for the corresponding tab on a Group Policy Modeling node. In particular, this page shows the component status for the various Group Policy extensions. This information tells you whether there were any issues with a particular extension and is a good place to begin troubleshooting.

- Settings – this is analogous to the information shown for the corresponding tab on a Group Policy Modeling node.

- Events – this tab shows all policy-related events from the target computer (see Figure 29). Note that to gather this data, the user performing the query must have access to remotely view the event log. By default, this access is granted to all users on Windows XP, but not on Windows Server 2003. This data is useful for troubleshooting Group Policy issues. For example if the summary report indicates that a particular Group Policy component failed to process, you may be able to determine why by looking for errors and warnings in the event log.




**Figure 29**

# Platform Dependencies

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

GPMC exposes features that are available in the underlying operating system. Because new features have been added to Group Policy since Windows 2000, certain features will only be available in GPMC depending on the operating system that has been deployed on the domain controllers and clients. This section describes these dependencies. In general, there are four key issues that determine whether a feature is available in GPMC:

- Windows Server 2003 Active Directory schema must be available to delegate Group Policy Modeling or Group Policy Results
- Windows Server 2003 domain controller must be available to run Group Policy Modeling
- Windows Server 2003 domain configuration (ADPrep /DomainPrep) must be available to use WMI Filters
- Clients must be running Windows XP or Windows Server 2003 in order to generate Group Policy Results data.

Windows and Active Directory platform dependencies are summarized below in Table 4.

**Table 4 - Group Policy Version Dependencies**

| Dependency | Feature | Reason |
|---|---|---|
| Windows Server 2003 Active Directory Schema | Delegation of Group Policy Modeling and Group Policy Results | The Generate Resultant Set of Policy (Logging) and Generate Resultant Set of Policy (Planning) permissions needed for this operation are only available with the Windows Server 2003 Active Directory schema |
| Windows Server 2003 Domain Controller in the forest | Group Policy Modeling | The simulation is performed by the Resultant Set of Policy Service which is only available on domain controllers running Windows Server 2003 |
| Windows Server 2003 domain configuration (DomainPrep) | WMI Filters | ADPREP /DomainPrep configures the domain for Windows 2003 Active Directory including configuration for WMI Filters |
| Clients must be running Windows XP or Windows Server 2003 | Group Policy Results | Clients must be instrumented to log Group Policy Results data when policy is processed. This capability is only available on the listed systems |

There is no dependency from the Group Policy perspective on whether a domain is in native mode or mixed mode.

# GPMC Options

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The options dialog box allows the user to customize how certain features work in GPMC. The options dialog box is available from the **View** menu in GPMC. The options menu includes tabs for Columns, Reporting, and General options which allow the user to control GPMC as described below:

- Columns

  This tab allows you to control the columns displayed and their order for the following list views in GPMC: **Group Policy Inheritance**, **Group Policy Objects**, **Linked Group Policy Objects** and **WMI Filters**. The columns to display can be selected by checking the appropriate column names in the **Columns** tab of the GPMC options. The order can be set by moving column names up or down in the list.

- Reporting

  This tab allows you to customize the location where GPMC searches for .adm files used by reporting functions in GPMC. The reporting functions in GPMC require the .adm file for the GPO to display properly. The search location option allows the user to specify a custom search folder for the .adm files, which takes precedence over the default option of first searching the Windows folder then the SYSVOL folder of the GPO. This option is discussed further in the "Reporting on GPO Settings" section earlier in this white paper.

- General

  This tab allows you to:

  - Enable or disable trust detection. By default you can only add a forest to GPMC for which there is a 2-way trust with the forest of the user running GPMC. You can optionally enable GPMC to work with only 1-way trust. By using the **Stored User Names and Passwords** in Windows XP and Windows Server 2003, you can also enable access to un-trusted forests, as described earlier in the “Managing Multiple Forests” section,

  - Enable or disable confirmation dialog boxes to distinguish between GPOs and GPO Links. As mentioned previously, GPMC distinguishes between GPOs and GPO-links in several ways. One method is a confirmation dialog box which appears by default when a user clicks on a GPO-link. This option controls whether that confirmation is displayed.

  - Display the domain controller name beside the domain name. By default, GPMC displays the domain controller that GPMC is using for each domain in the result pane when you go to the **Domains** node in the console. Using this option you can specify that GPMC should always display the domain controller in parentheses as part of the name of the domain in the console tree.

# Internet Explorer Enhanced Security Configuration Considerations

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 includes a new default security configuration for Internet Explorer, called Internet Explorer Enhanced Security Configuration (ESC). ESC impacts the Security Zones and Privacy settings within the Internet Explorer Maintenance settings of a GPO. The Security Zones and Privacy settings can either be ESC enabled or not.

- When you edit settings for Security Zones and Privacy settings in a GPO from a computer where ESC is enabled, that GPO will contain ESC-enabled settings. When you look at the HTML report for that GPO, the Security Zones and Privacy header will be appended with the text **(Enhanced Security Configuration enabled)**.

- When you edit settings for Security Zones and Privacy settings in a GPO from a computer where ESC is not enabled, that GPO will contain ESC-disabled settings. ESC is not enabled on any computer running Windows 2000 or Windows XP, nor on computers running Windows Server 2003 where ESC has been explicitly disabled.

ESC settings deployed through Group Policy will only be processed on and applied by computers where ESC is enabled. ESC settings will be ignored on computers where ESC is not enabled (all computers running Windows 2000 and Windows XP, and Windows Server 2003 computers where ESC has been explicitly disabled). The converse is also true: A GPO that contains non-ESC settings will only be processed on and applied by computers where ESC is not enabled.

Furthermore, ESC impacts the functionality contained in the HTML reports produced by GPMC as follows.

- On computers with ESC enabled a prompt appears when you attempt to view reports in GPMC. This happens because the reports contain a script that allows you to expand and collapse sections of the reports using **Show** and **Hide**. To use the **Show** and **Hide** functionality in these reports, you must add the *About:security_MMC.exe* site to the Trusted Sites zone in Internet Explorer. This site represents all Web pages that are hosted inside MMC. You can do this by clicking **Add** on the **Internet Explorer** prompt. This opens the **Trusted sites** dialog box with the correct entry (*About:security_mmc.exe*) for the page being called by GPMC. Click **Add,** and then click **Close** in the **Trusted sites** dialog box to add this site to the Trusted sites zone. If the *About:security_mmc.exe site* is not added to the Trusted sites zone, the reports appear fully expanded and cannot be collapsed.

- In addition, the Explain text for a given setting in the HTML report is available by clicking on any setting in the Administrative Templates section of the GPO, assuming you have already added *About:security_MMC.exe* to the list of trusted sites. A prompt appears when you attempt to view the Explain text for a given administrative template setting in the HTML report. This prompt will ask if you want to add the *about:blank* site to the list of trusted sites. This is not recommended because it could significantly compromise the security of Internet Explorer on that computer. If you do not add the *about:blank* site to the list of trusted sites, you will not lose significant functionality in GPMC. You will still be able to view the Explain text, however, the Print and Close buttons in the Explain text dialog box will not be functional. To close the dialog, use the close box in the upper right corner.

# Scripting Group Policy-related Tasks

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The GPMC user interface is based on a set of COM interfaces that accomplish most of the operations performed by GPMC. These interfaces are available to Windows scripting technologies like JScript and VBScript as well as programming languages such as Visual Basic and VC++. For example, the following capabilities are scriptable using these interfaces:

- Creating/deleting/renaming GPOs.
- Linking/unlinking GPOs and WMI filters.
- Delegation:
    - Security on GPOs and WMI filters.
    - Group Policy-related security on sites, domains, OUs.
    - Creation rights for GPOs and WMI filters.
- Generating reports of GPO settings.
- Generating reports of RSOP data.
- Backup/Restore of GPOs.
- Import/Export, Copy/Paste of GPOs.
- Search for GPOs, WMI filters, SOMs, and backups.

These interfaces are discussed in detail in the GPMC software development kit (SDK) located in the %programfiles%\gpmc\scripts\gpmc.chm help file on systems where GPMC has been installed. The contents of the GPMC SDK are also available in the Platform SDK.

GPMC comes with a number of sample scripts (written mostly in VBScript but some JScript) that form a toolkit of scripts that administrators can use to directly administer a Group Policy environment or as examples to build more elaborate management tools. The scripts are installed in the %programfiles%\gpmc\scripts directory. Table 5 shows a list of scripts that are provided to do the associated types of Group Policy administrative tasks:

**Table 5**

| Administrative task | Script name | Description |
|---|---|---|
| Back up a GPO | BackupGPO.wsf | Backs up all GPOs in a domain to the specified backup directory. |
| Back up all GPOs in a domain | BackupAllGPOs.wsf | Given a GPO name or a GUID, backs up the GPO to a specified backup directory. |
| Create a GPO with default options | CreateGPO.wsf | Creates a GPO with the specified name, in the current domain, using the default options. |
| Create a migration table | CreateMigrationTable.wsf | Populates the entries of a migration table with security principals and UNC paths that are referenced in a GPO or backup. |

| | | |
|---|---|---|
| Copy a GPO | CopyGPO.wsf | Creates a new GPO and copies the settings from the source GPO into the new destination GPO, given a source GPO name or GUID and a new destination GPO name. |
| Create a policy environment using an XML representation | CreateEnvironmentFromXML.wsf | Reads an XML file that specifies a policy environment; for example, OUs, GPOs, links, and security groups. The script can either create the environment in a domain by creating the objects, or delete the environment by deleting objects specified in the XML file. |
| Create an XML representation of a policy environment | CreateXMLFromEnvironment.wsf | Reads an existing policy environment and creates an XML file representing that environment. The XML file captures information about OUs, GPOs, and GPO links, and security on GPOs. You can use this script in conjunction with the CreateEnvironmentFromXML.wsf script to create a replica of domain for staging purposes. |
| Delete a GPO | DeleteGPO.wsf | Deletes the specified GPO when given a GPO name or GUID. By default the script deletes links to that GPO within the same domain. |
| Grant Permissions for all GPOs in a Domain | GrantPermissionOnAllGPOs.wsf | Grants a user or group the specified level of permission for all GPOs in the specified domain. |
| Generate a report for a GPO | GetReportsForGPO.wsf | Creates an HTML and XML report for a given GPO at a given location in the file system. |
| Generate a report for all GPOs in the domain | GetReportsForAllGPOs.wsf | Creates HTML and XML reports for all GPOs in the domain, at a given location in the file system. |
| Import settings into a GPO | ImportGPO.wsf | Imports the settings from the specified backup to the existing destination GPO in the domain |
| Import multiple GPOs into a domain | ImportAllGPOs.wsf | Creates a new GPO and imports settings into that GPO for each backed-up GPO stored at a specific file system location. |
| Restore a GPO | RestoreGPO.wsf | Restores a backed-up GPO. |
| Restore all GPOs | RestoreAllGPOs.wsf | Restores all GPOs that are stored at a given file system location |
| Grant permissions for GPOs linked to a domain, OU, or site | SetGPOSecurityBySOM.wsf | Grants a user or group the specified permission type for all GPOs that are linked to a specified domain, OU, or site. You can specify **Read, Apply, Edit, FullEdit,** or **None** for the permission type. |
| Set GPO permissions | SetGPOPermissions.wsf | Sets the permission level for a security principal on a given GPO. You can specify **Read, Apply, Edit, FullEdit,** or **None** for the permission type. |
| Set permissions to create GPOs | SetGPOCreationPermissions.wsf | Grants or removes the ability to create GPOs in a domain for a given security principal. |

| | | |
|---|---|---|
| Set policy-related permissions on a given site, domain, or OU | SetSOMPermissions.wsf | Sets policy-related permissions on a given scope of management (SOM). A SOM is any site, domain, or OU. |
| List all GPOs in a domain | ListAllGPOs.wsf | Prints all GPOs in the specified domain. |
| List disabled GPOs | FindDisabledGPOs.wsf | Prints all GPOs in the specified domain that are disabled or partially disabled. |
| List GPO information | DumpGPOInfo.wsf | Prints the information for a specific GPO, including creation time, modification time, owner, status, version number, links, security groups that filter the GPO, and security groups that have full control, edit, read, or custom permissions. |
| List scope of management information | DumpSOMInfo.wsf | Prints all information for a specific Scope of Management (SOM), including GPO links and policy related permissions on the SOM. A SOM is any site, domain, or OU. |
| List GPO by policy extension | FindGPOsByPolicyExtension.wsf | Prints all GPOs in the specified domain for which a specific policy extension is configured; for example, find all GPOs that contain the Software Installation or Folder Redirection policy settings. |
| List GPOs by security group | FindGPOsBySecurityGroup.wsf | Prints all GPOs that for which a given security principal has the specified permission on that GPO. You can specify **Read, Apply, Edit,** or **Fulledit** for the permission type. |
| List GPOs with duplicate names | FindDuplicateNamedGPOs.wsf | Prints all GPOs in the specified domain that have duplicate names. |
| List GPOs without Apply permission | GPOsWithNoSecurityFiltering.wsf | Prints all GPOs in the specified domain that do not apply to anyone because **Apply** permission is not set on the GPO. |
| List GPOs Orphaned in SYSVOL | FindOrphanedGPOsInSYSVOL.wsf | Finds and prints all GPOs in SYSVOL with no corresponding component in Active Directory. |
| List domains and OUs with external GPO links | FindSOMsWithExternalGPOLinks.wsf | Prints all domains and OUs in the specified domain that link to a GPO in a different domain. |
| List unlinked GPOs in a domain | FindUnlinkedGPOs.wsf | Prints all GPOs in the specified domain that have no links. Links outside the domain, including site links, are not checked. |
| Print the scope of management policy tree | ListSOMPolicyTree.wsf | Prints all SOMs in the specified domain with the list of GPOs that are linked to the domain and each OU. |
| List GPO backups in a given file system location | QueryBackupLocation.wsf | Prints information about all backed up GPOs at the file system location specified by the user. |

All of the scripts are intended for command line operation. The user can execute a script using the following command (using the CreateGPO script as an example) from a command shell:

*Cscript CreateGPO.wsf*

Alternatively, you can set Cscript to be the default scripting engine by using this command:

*cscript //H:cscript*

The user will need to run the scripts from the %programfiles%\gpmc\scripts directory or by specifying the path to the scripts directory. To see usage details for any script, use the "/?" command line option.

# Related Links (Administering Group Policy with Group Policy Management Console)

Updated: April 7, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

See the following resources for further information:

- Enterprise Management with the Group Policy Management Console at http://www.microsoft.com/windowsserver2003/gpmc/default.mspx.
- Migrating GPOs Across Domains with GPMC at http://www.microsoft.com/windowsserver2003/gpmc/migrgpo.mspx.

For the latest information about Windows Server 2003, see the Windows Server 2003 Web site at http://www.microsoft.com/windowsserver2003.

For information about Group Policy in general see http://www.microsoft.com/grouppolicy. TechNet also features a Group Policy page at http://www.microsoft.com/technet/grouppolicy.

# Maintaining the MIIS 2003 Database

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic provides basic guidelines for monitoring and backing up your Microsoft® Identity Integration Server 2003 database. Because MIIS 2003 uses Microsoft® SQL Server™ 2000 for database services and storage, this topic contains information about SQL Server 2000 in addition to information about MIIS 2003.

**In This White Paper**


Overview of Maintaining MIIS 2003

Overview of the SQL Server 2000 Implementation for MIIS 2003

Configuring SQL Server 2000 to Support MIIS 2003

SQL Server 2000 Databases Used by MIIS 2003

SQL Server 2000 Backup Types and Recovery Models

Backup Strategies for Your MIIS Database

Creating a New Backup Device

Scheduling Full Database Backups by Using the Database Maintenance Plan Wizard

Scheduling Differential Database Backups by Using SQL Server Enterprise Manager

Scheduling Backups of the System Databases

Backing Up the MIIS 2003 Encryption Key

Restoring the MIIS Database

Recommendations for Backing Up and Restoring the MIIS Database

Ensuring Failover and High Availability

Monitoring the MIIS Database

Establishing a Performance Baseline

Setting Alerts to Monitor System Events

Additional Ways to Monitor the MIIS Database

Optimizing the MIIS Database

Additional Resources (Maintaining the MIIS 2003 Database)

# Overview of Maintaining MIIS 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

MIIS 2003 uses SQL Server 2000 as its data store. Because your MIIS 2003 data is a critical asset to your organization, it is important to back up your MIIS database regularly. It is also important to monitor your installation of SQL Server 2000 on a regular basis to assure optimum performance and reliability.

This topic assumes that you already have a thorough working knowledge of MIIS 2003 and some familiarity with SQL Server 2000. For more information about SQL Server 2000, see Microsoft SQL Server 2000 Books Online, which you can install during SQL Server 2000 installation and then access from the Start menu through All Programs. Optionally, see "Getting Started with SQL Server Books Online" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=7107.

# Overview of the SQL Server 2000 Implementation for MIIS 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to understand the relationship between SQL Server 2000 editions and the editions of MIIS 2003. It is also important to understand some general recommendations for running SQL Server 2000 and the MIIS database. For example, backup types and recovery models, which are a part of SQL Server database maintenance, are critical to the safety concerns of your MIIS 2003 data.

**SQL Server 2000 and Windows Server 2003 Supported Editions**
Each edition of MIIS 2003 is supported by specific editions of Microsoft® Windows Server™ 2003 and SQL Server 2000. Table 1 lists the Windows Server 2003 and SQL Server 2000 editions that support each edition of MIIS 2003.

**Table 1 Windows Server 2003 and SQL Server 2000 Editions**

| This MIIS 2003 Edition | Is Supported by These Windows Server 2003 Editions | And by These SQL Server 2000 Editions |
|---|---|---|
| Standard Edition | Standard Edition<br><br>Enterprise Edition | Enterprise Edition (with SP3*)<br><br>Standard Edition (with SP3) |
| Enterprise Edition | Enterprise Edition | Enterprise Edition (with SP3) |

* Service Pack 3.

**MIIS 2003 Requirements for SQL Server 2000**
The following are MIIS 2003 requirements for SQL Server 2000:

- SQL Server 2000 and MIIS 2003 can be installed on the same server or different servers, but they must be located in the same domain.
- If you configure SQL Server 2000 to use a communications port other than the default SQL Server port 1433, you must use the SQL Server Client Network Utility on the computer on which you run MIIS 2003 to set the communications port to be used by SQL Server clients on that computer.

For more information about setting the communications port, see “Configuring SQL Server 2000 to Support MIIS 2003” later in this topic.

**Recommendations for Running SQL Server 2000 to Support MIIS 2003**
The following recommendations for running SQL Server 2000 will help you to improve the performance and efficiency of your MIIS 2003 installation:

- Do not attempt to manually modify data that is stored in the MIIS database. Modify the database only by using Identity Manager, which can determine if and what changes are necessary.
- Maintain your transaction log files for the MIIS database on a separate drive from the drive that contains the MIIS database.
- To ensure availability and provide redundancy in large-scale environments, use SQL Server clustering with MIIS 2003, Enterprise Edition.
- In addition to using the performance improvement methods described elsewhere in this document, use performance tuning to optimize SQL Server 2000 performance.

For information about SQL Server 2000 performance tuning, see “Chapter 20 - RDBMS Performance Tuning Guide for Data Warehousing” on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkID=17516.

# Configuring SQL Server 2000 to Support MIIS 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To meet the requirements and recommendations for MIIS 2003, be sure that your SQL Server 2000 implementation is configured correctly. First, verify that the correct version of SQL Server 2000 is running in your environment. Then, check whether SQL Server 2000 is configured to use the correct communications port, and reconfigure it, if necessary.

**To verify that the correct version of SQL Server 2000 is running**

1. In Microsoft SQL Server, click **Query Analyzer**.
2. In the **Connect to SQL Server** dialog box, type your server name, click a connection option, and then click **OK**.
3. In the **Query** box, type or select the appropriate version.
4. On the **Query** menu, click **Execute**.

The query returns the version of SQL Server 2000 that is running. Ensure that you are running version number **8.00.760**, which indicates SQL Server 2000 SP3 or later.

By default, SQL Server 2000 uses communications port 1433. If SQL Server 2000 is configured to use a port other than 1433, you must use the SQL Server Client Network Utility on the computer running MIIS 2003 in order to change the communications port. You must do this even if MIIS 2003 is running on the same computer that SQL Server 2000 is running on.

**To determine which communications port that SQL Server 2000 is using**

1. In **Microsoft SQL Server**, click **Server Network Utility**.
2. On the **General** tab, under **Enabled Properties**, click **TCP/IP**, and then click **Properties**. The configured port is listed in **Default port**.

**To change the communications port that SQL Server 2000 is using**

1. On the computer that is running MIIS 2003, open **Microsoft SQL Server**, and then click **Client Network Utility**. (If this item does not appear, you must first install the SQL Server client tools from the SQL Server 2000 installation media.)
2. On the **General** tab, under **Enabled protocols by order**, click **TCP/IP**, and then click **Properties**.
3. In the **Default port** box, type the communications port in use by SQL Server 2000, and then click **OK**.
4. Click **OK** to close the Client Network Utility.

| **Important** |
| --- |
| The server and client port numbers must match. |

# SQL Server 2000 Databases Used by MIIS 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The implementation of SQL Server 2000 that is used by MIIS 2003 contains both an MIIS database for use by MIIS 2003 and several SQL Server system databases for use by SQL Server 2000. In addition, a SQL Server 2000 implementation can include several sample databases.

**To view the SQL Server system databases and the MIIS database**

1. In **Microsoft SQL Server**, click **Enterprise Manager**.

2. Expand **Microsoft SQL Servers**, expand **SQL Server Group**, expand *servername* (where servername represents the computer on which SQL Server 2000 is running), and then expand **Databases**.

3. Click a database to view.

   To view more information about a database, right-click the database, and then click **Properties**.

**MicrosoftIdentityIntegrationServer Database in SQL Server 2000**

When MIIS 2003 is installed, a SQL Server 2000 database that is named **MicrosoftIdentityIntegrationServer** is created within SQL Server 2000. The **MicrosoftIdentityIntegrationServer** database contains all the metadirectory data used by MIIS 2003, including the following:

- Metaverse data
- Management agent data
- Connector space data
- Rules extension data
- Object link data
- Management agent run profiles
- Management agent run history
- Joiner log

The **MicrosoftIdentityIntegrationServer** database does not contain any of the following:

- Event log files
- Contents of the directory at C:\ProgramFiles\Microsoft Identity Integration Server\MaData
- Administrative scripts

**SQL Server 2000 System Databases**

In addition to the MIIS database, SQL Server 2000 contains several system databases, which are created when you install SQL Server 2000.

**Master** The database that controls the operation of each instance of SQL Server. It is installed automatically with each instance of SQL Server and keeps track of user accounts, remote user accounts, and remote servers with which each instance can interact. It also tracks ongoing processes, configurable environment variables, system error messages, tapes and disks available on the system, and active locks.

**Msdb** The database used to store scheduling information, and backup and restore history information.

**Distribution** A database on the Distributor that stores data for replication, including transactions, snapshot jobs, synchronization status, and replication history information.

**Model** A database installed with SQL Server that provides the template for new user databases. SQL Server 2000 creates a new database by copying the contents of the model database and then expanding it to the size requested.

| ⚠**Caution** |
| --- |
| Do not delete the SQL Server 2000 system databases. These databases are critical to the core functionality of SQL Server 2000. |

**Sample Databases**

In addition to the MIIS database and SQL Server system databases, your implementation of SQL Server 2000 might include two sample databases: **Northwind** and **pubs**. To save disk space in your production environment, you can delete these two databases from SQL Server 2000.

**To delete the Northwind and pubs sample databases**

1. In **Microsoft SQL Server**, click **Enterprise Manager**.

2. Expand **Microsoft SQL Servers**, expand **SQL Server Group**, expand *servername* (where *servername* represents the computer on which SQL Server 2000 is running), and then expand **Databases**.

3. Right-click the sample database that you want to delete, click **Delete**, and then click **Yes** to confirm the deletion.

# SQL Server 2000 Backup Types and Recovery Models

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

SQL Server 2000 provides three different backup types and three different recovery models. You can combine these backup types and recovery models in different ways, depending on your data recovery requirements.

**SQL Server 2000 Backup Types**

You can perform three different types of backups on the MIIS database: database (or full), differential, and transaction log.

**Database (or Full)** Full database backups are the default and the starting point for all other types of backups. A full database backup captures the entire database, including all entries in the transaction log—excluding any unallocated extents in the files. Pages are read directly from disk to increase the speed of the operation.

**Differential** Differential backups capture *all* the data that has changed since the *last* full database backup. Differential backups will increase the speed of the backup operation as well as the restore. Because only the changed or newly allocated extents (bitmap tracking) are captured, differential backups are faster and smaller than full database backups. During a recovery, the last differential backup can be restored without applying a series of the individual transaction log backups or differential backups. All changed data is captured in the last differential backup.

**Transaction log** Transaction log backups serially capture modifications to the database. Backups of the transaction log provide a history of the transactions that have taken place within the database. The backups of the log are then used in the recovery process to restore the database fully, to a point in time. Transaction log backups are applied to recover a database by rolling forward (redoing) any committed changes not reflected in the database and rolling back (undoing) uncommitted transactions. Log backups are smaller and are taken more frequently than full or differential backups.

For more information about these backup types, see "Chapter 3 - Backing Up and Restoring Databases" on the Microsoft TechNet Web site at http://go.microsoft.com/fwlink/?LinkID=15117.

**SQL Server 2000 Recovery Models**

SQL Server 2000 provides three different recovery models: Full Recovery, Bulk-Logged Recovery, and Simple Recovery. The recovery model determines whether and how transaction logs can be backed up. The recovery model that you choose has a significant impact on system performance (especially during heavy data loads) and also on your ability to recover data.

By default, the MIIS database is configured to use the Simple Recovery model. Administrators can choose to enhance their data recoverability options by using the Full Recovery or Bulk-Logged Recovery model. These enhanced methods can dramatically improve disk performance by reducing the amount of disk space used by the log. However, using the Full Recovery or Bulk-Logged Recovery model also makes backup and restoration procedures more complex.

Be sure that you thoroughly understand the differences between the recovery models before you apply the Full Recovery or Bulk-Logged Recovery model. Table 2 lists and describes the SQL Server 2000 recovery models.

**Table 2 SQL Server 2000 Recovery Models**

| Model | Benefits | Work-Loss Exposure | Recovery Point |
|---|---|---|---|
| Full Recovery | No work is lost due to a lost or damaged data file.<br><br>Can recover to a specified point in time (for example, prior to an application or user error). | None. However, if the log is damaged, changes made since the most recent log backup are lost and must be redone. | Can recover to any point in time. |
| Bulk-Logged Recovery | Permits high-performance bulk copy operations.<br><br>Minimal log space is used by bulk operations. | If the log is damaged, or bulk operations occurred after the most recent log backup, changes made since the last backup must be redone. | Can recover to the end of any backup. |

| | | | |
|---|---|---|---|
| Simple Recovery | Permits high-performance bulk copy operations.<br><br>Reclaims log space to keep space requirements small. | Changes made after the most recent database or differential backup are lost and must be redone. | Can recover to the end of any backup. |

| **Note** |
|---|
| For more information about recovery models, see "Backing Up and Restoring Databases" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=10575. |

You can change your recovery model to accommodate changes in your recovery requirements.

**To change SQL Server 2000 recovery models**

1. In **Microsoft SQL Server**, click **Enterprise Manager**.

2. Expand **Microsoft SQL Servers**, expand **SQL Server Group**, expand *servername* (where *servername* represents the computer on which SQL Server 2000 is running), and then expand **Databases**.

3. Right-click the MIIS database, and then click **Properties**.

4. On the **Options** tab, in the **Model** list, click a recovery model.

After you change the recovery model, it is recommended that you perform a full database backup.

# Backup Strategies for Your MIIS Database

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to establish and document a backup and restore plan for your MIIS 2003 database, being sure to thoroughly test your plan in a simulation of your production environment.

| **Note** |
| --- |
| For more information about developing backup and restore plans for SQL Server databases, see "Chapter 3 – Backing Up and Restoring Databases" on the Microsoft TechNet Web site at http://go.microsoft.com/fwlink/?LinkID=15117. |

Unless specified otherwise, this section provides instructions for using the Simple Recovery model, which is the default recovery model for MIIS 2003.

For information about backing up the SQL Server 2000 databases when you are using the Full Recovery or Bulk-Logged Recovery models, see "Backing Up and Restoring Databases" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=10575.

Table 3 lists the recommended backup activities that are necessary to ensure the recoverability of your MIIS database, including the back up frequency for each.

**Table 3 Recommended Backup Activities**

| Backup Activity | When You Install | Intraday | Daily | Monthly | Quarterly |
| --- | --- | --- | --- | --- | --- |
| Create a backup device. | X | | | | |
| Determine an offsite storage location for your backups. | X | | | | |
| Determine and document your backup procedures. | X | | | | |
| Perform differential backups. | | X | | | |
| Perform transaction log backups. | | X | | | |
| Automatically shrink unused database file space. | | | X | | |
| Perform full database backups. | | | X | | |
| Back up the MIIS 2003 encryption keys. | | | | X | |
| Back up the **master** and **msdb** system databases in SQL Server. | | | | X | |
| Back up the entire MIIS 2003 installation. | | | | X | |
| Back up the entire SQL Server 2000 installation. | | | | X | |
| Perform a test recovery of your data in a test environment by using your production backups. | | | | | X |
| Review your backup procedures and amend as necessary | | | | | X |

**Note**

If your SQL Server implementation also contains other databases that are used by MIIS 2003 or other applications, be certain to develop a backup and restore plan that includes those databases. For a more detailed administrative schedule for SQL Server 2000, see the SQL Server 2000 Operations Guide on the Microsoft TechNet Web site at http://www.microsoft.com/technet/prodtechnol/sql/2000/maintain/sqlops0.mspx.

Backing up your MIIS database involves first creating a backup device within SQL Server 2000 and then implementing your database backup plan. Finally, you must back up the MIIS 2003 encryption keys. You can back up SQL Server 2000 databases by using the Database Maintenance Plan Wizard or SQL Server Enterprise Manager.

# Creating a New Backup Device

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Before you can create backups of your MIIS database, you must create a new backup device in SQL Server 2000 by using SQL Server Enterprise Manager. The physical backup device (a tape drive, for example) must already be installed and configured before you perform this procedure.

**To create a new backup device by using SQL Server Enterprise Manager**

1. In **Microsoft SQL Server**, click **Enterprise Manager**.
2. Expand **SQL Server**, and then expand **Management**.
3. Right-click **Backup**, and then click **New Backup Device**.
4. In the **Name** box, enter the name of the device, and then click **OK**.

| **Note** |
| --- |
| For more information about SQL Server 2000 backup devices, see "Backing Up and Restoring Databases" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=10575. |

# Scheduling Full Database Backups by Using the Database Maintenance Plan Wizard

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The easiest way to back up your MIIS database is to use the Database Maintenance Plan Wizard, which is available in SQL Server Enterprise Manager. By using the Database Maintenance Plan Wizard, you can create a job that runs full database backups and transaction log backups on a scheduled basis.

The procedures for backing up the MIIS database are the same as those for backing up SQL Server databases.

| **Important** |
| --- |
| It is recommended that you schedule transaction log backups to occur more frequently than the database backups (for example, every four hours). Backing up the transaction log truncates (removes) logged transactions and can improve performance. |

**To schedule backups by using the Database Maintenance Plan Wizard**

1. In **Microsoft SQL Server**, click **Enterprise Manager**.
2. Expand the server group where the **MicrosoftIdentityIntegrationServer** database is located, and then expand the server.
3. On the **Tools** menu, click **Database Maintenance Planner**.
4. On the **Start** page, click **Next**.
5. On the **Select databases** page, click **These databases**, click **MicrosoftIdentityIntegrationServer**, and then click **Next**.
6. On the **Update Data Optimization** page, if you want the database to be shrunk automatically, freeing up unused space, do the following:

   | **Note** |
   | --- |
   | For an optional method of shrinking the database, see "Using the Auto Shrink Option to Maintain a Minimum of 25% Unused Space" later in this section.<br><br>It is recommended that you choose one of these methods to automatically (and regularly) shrink database size. |

   - Select **Remove unused space from database files**, and then in the **When it grows beyond box**, keep the default value of 50 megabytes and in the **Amount of free space to remain after shrink** box, keep the default value of 10%,
   - Click **Change**, and then specify a schedule that shrinks the database, setting a time that is at least one hour before your scheduled full backups.

     Allowing one hour usually gives sufficient time for the database shrink to complete before the scheduled backup begins.
7. Click **Next** until you reach the **Specify the Database Backup Plan** page.
8. On the **Specify the Database Backup Plan** page, do the following:

   - Select the **Back up the database as part of the maintenance plan** and **Verify the integrity of the backup**

**when complete** check boxes.

- Under **Location to store the backup file**, click the tape or disk on which you want to store the backup, and then click **Change**.
- In the **Edit Recurring Job Schedule** dialog box, specify the frequency, time, and duration of your backup job.
- Make sure that the **Enable schedule** check box is selected, and then click **OK**.

9. (Optional) If using the Full Recovery or Bulk-Logged Recovery model, do the following:

   - On the **Specify the Transaction Log Backup Plan** page, select **Back up the transaction log as part of the maintenance plan**, and **Verify the integrity of the backup when complete**.
   - Under **Location to store the backup file**, click the tape or disk on which you would like to store the backup, and then click **Change**.
   - In the **Edit Recurring Job Schedule** dialog box, specify the frequency, time, and duration of your backup job.
   - Make sure that the **Enable schedule** check box is selected, and then click **OK**.

10. On the **Specify Transaction Log Backup Disk Directory** page, select a location for the transaction log backup, and then specify if you want older backup files automatically removed.

11. Accept the default selections on the remaining wizard pages until you reach the **Completing the Database Maintenance Plan Wizard** page, and then in the **Plan Name** textbox, type a name for your backup plan.

**Using the Auto Shrink Option to Maintain a Minimum of 25% Unused Space**

In addition to shrinking the database automatically (see step 6 of the preceding procedure), you can also use the Auto Shrink option to automatically shrink the database each time the database file contains more than 25% unused space. It is recommended that you choose one of these methods to automatically shrink database size on a regular basis.

For more information about shrinking the database, see Microsoft SQL Server Books Online.

**To enable Auto Shrink**

- In **Databases**, right click the **MicrosoftIdentityIntegrationServer** database, click **Properties**, click **Options**, and then in **Settings**, select **Auto Shrink**.

# Scheduling Differential Database Backups by Using SQL Server Enterprise Manager

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use SQL Server Enterprise Manager to schedule differential database backups of the MIIS database.

**To back up the database by using SQL Server Enterprise Manager**

1. In **Microsoft SQL Server**, click **Enterprise Manager**.
2. Expand the server group where the **MicrosoftIdentityIntegrationServer** database is located, and then expand the server.
3. On the **Tools** menu, click **Backup Database**.
4. In the **Database** list, click **MicrosoftIdentityIntegrationServer**.
5. Under **Backup**, click **Database - differential**.
6. Under **Destination**, select the backup device on which you want the backup to be performed.
7. Select the **Schedule** check box, and then click the ellipsis (...) next to the **Schedule** check box.
8. In the **Edit Schedule** dialog box, click **Change**, and then set the frequency of the backup schedule in the **Edit Recurring Schedule** dialog box.
9. Click **OK** three times to return to Enterprise Manager.

| **Note** |
| --- |
| For more information about differential database backups, see "Chapter 3 - Backing Up and Restoring Databases" on the Microsoft TechNet Web site at http://go.microsoft.com/fwlink/?LinkID=15117. |

# Scheduling Backups of the System Databases

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

It is important to back up the **master** and **msdb** databases any time they change. At a minimum, it is recommended that you perform monthly backups of the system databases.

**Master** database changes occur as a result of:

- Adding logins or other login security-related operations in SQL Server 2000.
- Changing any server-wide or database configurations.
- Creating or removing logical backup devices.
- Configuring the server for distributed queries and remote procedure calls, such as adding linked servers or SQL Server 2000 remote logins.

**Msdb** database changes occur as a result of:

- Scheduling tasks.
- Storing Data Transformation Services (DTS) packages created by using the DTS Import/Export Wizard to an instance of SQL Server.
- Maintaining online backup and restore history.
- Replicating.

You can schedule backups of the **master** and **msdb** databases by using the Database Maintenance Plan Wizard procedure for scheduling full database backups and substituting the appropriate database on the **Select databases** page. For more information about this procedure, see “Scheduling Full Database Backups by Using the Database Maintenance Plan Wizard” earlier in this topic.

# Backing Up the MIIS 2003 Encryption Key

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You need to regularly back up the encryption key that MIIS 2003 uses to encrypt data and credentials. Failure to back up encryption keys can lead to loss of encrypted data and credentials. To back up the encryption key, use the MIIS Key Management Utility.

The MIIS Key Management Utility has two different user interfaces: command line and a Windows wizard. For most backup and restore operations of the encryption key, the command line interface is sufficient. For more advanced backup operations of the encryption key, for example, to add a new encryption key to the key set or to abandon an existing key set and create a new set, you need to use the MIIS Key Management wizard.

To back up the encryption key, you must be logged on as a member of an MIIS security group and have administrative credentials on the local computer.

**To back up the MIIS 2003 encryption key by using the command line**

1. On the server where MIIS 2003 is running, change the directory by typing the following at the command line:

   **cd %Programfiles%\Microsoft Identity Integration Server\bin**

2. Type the following at the command line to save the encryption key to a floppy disk:

   **miiskmu /e a:keyback.bin /u miisadmin ***

   **Table 4 Parameters for the Miiskmu Command**

| Parameter | Description |
|---|---|
| /e | Exports the key set to a file. |
| **a:***keyback.bin* | For the a: drive, specifies the file name in which you are saving the encryption key. |
| /u | Specifies the MIIS 2003 service account credentials. |
| *miisadmin* | Specifies the MIIS service account. |

**To back up the MIIS 2003 encryption key by using the MIISkmu wizard**

1. Click **Start**, point to **All Programs** and **Microsoft Identity Integration Server**, and then click **Key Management Utility**.

2. Follow the instructions in the wizard.

| **Note** |
|---|
| When using the MIISkmu wizard, you might need to stop the Microsoft Identity Integration Server service in Windows Services in order to enable advanced options. |

# Restoring the MIIS Database

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The steps required to recover from an MIIS database failure vary according to the severity of the failure. Based on the severity of the failure, you might need to complete one or several steps, including recovering your MIIS 2003 server, recovering your SQL Server 2000 server, and restoring the MIIS database and the SQL Server 2000 system databases.

If you need to recover your server before you restore the MIIS database (for example, because of a server hardware failure), you will need to complete some or all of the following steps:

- **Reinstall the operating system on the server.** Ensure that the server role is configured in the way that it was previously, either as a standalone server or as a member of an Active Directory® directory service domain. If the server is a standalone server, use the Windows Server 2003 Backup tool to restore the local Security Accounts Manager (SAM) database. Note that you can use the same server name or a different name.

- **Install SQL Server 2000.** Use the same installation options that you used for the previous installation.

- **Restore the MIIS database and the SQL Server 2000 system databases.** If you are not using the simple recovery model or if there are differential backups since the last full backup was created, you must also restore those differential backups to recover incomplete transactions.

- **Install MIIS 2003 on your server.** Use the same configuration for the options that you used for the previous installation. When you are prompted for the encryption key during Setup, insert the disk that includes the MIIS 2003 encryption key.

The procedures for restoring the MIIS database are the same as those for restoring other SQL Server databases. For general information about restoring SQL Server databases, see "Chapter 3 – Backing Up and Restoring Databases" on the Microsoft TechNet Web site at http://go.microsoft.com/fwlink/?LinkID=15117.

# Recommendations for Backing Up and Restoring the MIIS Database

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following recommendations will help you to optimize your process for maintaining the MIIS database:

- In addition to making regular SQL Server database backups, use the Windows Server 2003 Backup tool to create an Automated System Recovery (ASR) of the servers running MIIS 2003 and SQL Server 2000 on a regular (at least monthly) basis. You can use an ASR set to restore your servers if a disaster occurs.

  For more information about creating an ASR set, see "Safeguarding your system" in Help and Support Center for Windows Server 2003.

- Test backup and restore procedures on a regular basis and before a disaster occurs. As a guide, test your backup and restore procedures at least once every three months.

- Always keep your backups offsite in a secure location, to lessen any losses that could occur from fire, disaster, or accident.

- Make a full backup of your MIIS database everyday.

- On SQL Server 2000, back up the **master** and **msdb** system databases at least monthly.

- After you complete a backup, verify all your backup media.

# Ensuring Failover and High Availability

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can create a standby server that is running MIIS 2003 before you set up your primary server running MIIS 2003. If the SQL Server 2000 installation that contains your MIIS database resides on a separate server from the primary and standby MIIS 2003 installations, you can activate the standby server to ensure that MIIS 2003 is available even if your primary system fails. To activate a standby server running MIIS 2003, you need to run MIISactivate on the standby server by using the Microsoft Identity Integration Server service account credentials and the most recent backup copy of the MIIS 2003 encryption key.

For more information about MIISactivate, see “MIISactivate: Server activation tool” in MIIS 2003 Help. For more information about the MIIS 2003 encryption keys, see “MIISkmu: Encryption key management tool” in MIIS 2003 Help. For more information about backing up your encryption key, see “Backing Up the MIIS 2003 Encryption Key” earlier in this topic.

# Monitoring the MIIS Database

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Monitoring the MIIS database requires that you first establish a performance baseline for your implementation. The goal of baseline monitoring is to document how the SQL Server 2000 installation that hosts the MIIS database operates under normal conditions. Establishing a baseline helps you identify problems more quickly.

After you document your performance baseline, you need to continue monitoring on an ongoing basis the performance of the SQL Server 2000 installation that hosts the MIIS database. A number of tools are available to assist you with ongoing performance monitoring.

# Establishing a Performance Baseline

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Baseline performance monitoring involves establishing a performance baseline for your system. A performance baseline includes a single performance chart accompanied by an interpretation of the results, based on your environment. Many elements of the chart, such as the timeframe, vary according to the environment. You can use System Monitor in Windows Server 2003 to establish your performance baseline. The System Monitor chart can be created in real time or based on a performance log file. It is recommended that you base your chart on a log file, because this allows you to record statistics for an extended period.

When establishing your performance baseline, you must select the sampling interval and the server from which to monitor. The sampling interval determines the amount of disk space that your baseline requires. (Using performance counters on the server does not add much overhead to the server running SQL Server.) If you find that your baseline data is taking up too much disk space, you can use a larger interval between data samples. The larger the sampling interval you use, however, the less accurate your graph will be.

You can monitor remotely; however, use of the counters across a network connection for an extended period of time can congest traffic on your network. If you have disk space on your server for the performance log files, it is recommended that you record performance log information locally.

Because of performance impacts related to the number of counters being sampled and the frequency with which sampling occurs, it is important to test the number of counters and the frequency of data collection so you can determine the balance that meets the needs of your environment while still providing useful performance information. For the initial performance baseline, however, it is recommended that you use the highest number of counters possible and the highest frequency available.

Table 5 lists the recommended performance counters to include in your performance baseline.

**Table 5 Performance Counters for Baseline Performance Monitoring**

| Object\Counter | Definition | Recommendations |
|---|---|---|
| Memory\Pages/sec | The number of pages read from or written to disk to resolve hard page faults. This counter serves as a primary indicator of the types of faults that cause system-wide delays. | Although it is normal to have some spikes, this counter generally remains at or close to zero. |
| Network Interface\Bytes total/sec | Number of bytes traveling over the network interface per second. This counter only reflects the local network connection. | If this value stays below 50 percent of your available network bandwidth, the network adapter on the server running SQL Server 2000 is not likely to cause any performance bottlenecks. |
| PhysicalDisk\Disk Transfers/sec | The rates of read and write operations on the disk. Define a counter for each physical disk on the server. | |
| PhysicalDisk\Avg. Disk Queue Length | The average numbers of both read and write requests that were queued for the selected disk during the sample interval. | The number of requests should not exceed two times the number of spindles constituting the physical disk. If the number of requests is too high, you can add additional disks or replace the existing disks with faster disks. |
| SQLServer:Memory Manager\Total Server Memory | The total memory in use by SQL Server 2000. | Add memory to the server if this value is generally higher than the amount of physical memory in the server. |

| | | |
|---|---|---|
| SQLServer:Access Methods\Full Scans/sec | The number of unrestricted full scans. These can either be base table or full index scans. | |
| SQLServer:Buffer Manager\Buffer Cache Hit Ratio | The percentage of pages that were found in the buffer pool without having to incur a read from disk. | When this percentage is high, your server is operating at optimal disk I/O efficiency. If this value decreases over time, you might consider adding physical memory to your server. |
| SQLServer:Databases\Log Growths | The total number of log growths for the selected database. | Run against your application database instance. |
| SQLServer:Databases Application Database\Percent Log Used | The percent of space in the log that is in use. | Run against your application database instance. |
| SQLServer:Databases Application Database\Transactions/sec | The number of transactions started for the database. | Run against your application database instance. |
| SQLServer:General Statistics\User Connections | The number of users connected to the system. | Research any dramatic shifts in this value. |
| SQLServer:Latches\Average Latch Wait Time | The average latch wait time, in milliseconds, for latch requests that had to wait. | If this number is high, your server might have resource limitations. |
| SQLServer:Locks\Average Wait Time | The average amount of wait time, in milliseconds, for each lock request that resulted in a wait. | |
| SQLServer:Locks\Lock Waits/sec | The number of lock requests that could not be satisfied immediately and required the caller to wait before the lock was granted. | |
| SQLServer:Locks\Number of Deadlocks/sec | The number of lock requests that resulted in a deadlock. | |
| SQLServer:Memory Manager\Memory Grants Pending | The current number of processes waiting for a workspace memory grant. | |

For information about how to use System Monitor to create performance charts and logs, see “Monitoring performance” in Help and Support Center for Windows Server 2003.

# Setting Alerts to Monitor System Events

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

An alert is a notification that is generated as a result of an event. You can use alerts for ongoing system monitoring; for example, you can set an alert to notify you if CPU usage reaches 100 percent.

To monitor your MIIS 2003 implementation, it is recommended that you use alerts for the following events:

- **Errors affecting service.** Whenever possible, ensure that each known critical error has an alert associated with it. If you cannot associate the error with an alert, consider writing a query to test for the error state, and then run the query on an ongoing basis, configuring an alert to be generated when the query finds the error.

- **Deadlocks.** Ensure that when the value of any SQLServer:Locks\Number of Deadlocks counter exceeds a given threshold, it generates an alert. If an application does not typically cause deadlocks, set this threshold to 1 or 2. If you run concurrent management agents, set this threshold from 10 to 20.

- **CPU usage.** If your application typically freezes or degrades at a certain level of CPU usage, set an alert to be generated when the CPU usage exceeds that level.

- **Disk usage.** If your application freezes or degrades at a certain disk usage level (or queue level), set an alert to be generated when the disk usage exceeds that level.

- **Table scans.** Ensure that your application database generates an alert if excessive table scans occur. First, establish the baseline value for excessive table scans in your environment. To determine your baseline, monitor the SQLServer:Access Methods\Full Scans/sec counter for a period of time; then set an alert. If you have many small tables that do not require indexes, a large number of acceptable table scans can occur before the number becomes excessive.

For information about setting and managing alerts, see “Performance Logs and Alerts” in Help and Support Center for Windows Server 2003.

# Additional Ways to Monitor the MIIS Database

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In addition to using performance counters and alerts, MIIS 2003 and SQL Server 2000 provide the following ways to monitor the MIIS database and SQL Server 2000 system databases:

- Use the error log in SQL Server Enterprise Manager.
- Use Event Viewer to monitor SQL Server events recorded in the Windows application log.
- Use Current Activity in SQL Server Enterprise Manager for ongoing monitoring.
- Use SQLDiag.exe to collect detailed information.
- Use a Query Analyzer session if SQL Server Enterprise Manager is unavailable.

**Use the Error Log in SQL Server Enterprise Manager**

You can use the error log in SQL Server Enterprise Manager for ongoing performance monitoring. The SQL Server error log enables you to determine how well the SQL Server 2000 installation that hosts the **MicrosoftIdentityIntegrationServer** database is functioning.

**To view the SQL Server error log**

1. In **SQL Server Enterprise Manager**, expand a server group, and then expand a server.
2. Expand **Management**, and then expand **SQL Server Logs**.
3. Click the SQL Server log that you want to view.

Error log details appear in the details pane.

**Use Event Viewer to Monitor SQL Server Events**

Event Viewer allows you to monitor SQL Server events recorded in the Windows application log. This log is separate from the SQL Server error log. SQL Server 2000 messages can be identified by their source, **MSSQLSERVER** or **SQLSERVERAGENT**.

**To view messages in the Windows application log**

1. Open **Event Viewer**.
2. In the console tree, click **Application Log**.
3. In the details pane, right-click an event, and then click **Properties**.

**Use Current Activity for Ongoing Monitoring**

You can use Current Activity in SQL Server Enterprise Manager for ongoing monitoring. Current Activity displays a snapshot of information regarding processes, user activity, and locks held by processes and on objects. Because using this tool requires no overhead, it can be helpful to view Current Activity when you encounter a problem state. Note that the information in Current Activityis only a snapshot; for ongoing performance monitoring, you must refresh this snapshot regularly.

You can use Current Activity to do the following:

- Monitor blocked and blocking transactions.
- Monitor users that are currently connected and their last executed statement.
- View all locks by database object.

- Terminate a selected process or send a message to the user who is executing a problematic transaction to terminate the transaction.

| **Important** |
| --- |
| Keep in mind that if you decide to terminate a process, the command might require a lengthy rollback. Hence, use this feature sparingly. |

**To view the Current Activity window**

1. In **SQL Server Enterprise Manager**, expand a server group, and then expand a server.
2. Expand Management, and then expand Current Activity.
3. Click **Process Info**.

The current server activity is displayed in the details pane.

**Use SQLDiag.exe to Collect Detailed Information**

You can gather detailed information about the current state of SQL Server 2000 — including current user information, DLL versions, configuration information, and database size information — by using the Sqldiag.exe tool, which is located in *drive*:\ProgramFiles\Microsoft SQL Server\Mssql\Binn\.

To run this tool, type the following at the command line:

**sqldiag -O filename -X**

**Table 6 Parameters for the Sqldiag Command**

| Parameter | Description |
| --- | --- |
| -O | Directs all Sqldiag.exe output to a single file. |
| *filename* | Specifies the name of the file where you want to direct the command output. |
| -X | Excludes error logs. |

For more information about Sqldiag.exe parameters, type **sqldiag /?** at the command line.

Sqldiag.exe can take up to a few minutes to complete. When Sqldiag.exe completes successfully, it writes a file named *filename*, which lists detailed SQL Server 2000 information, to the server.

**Use a Query Analyzer Session if SQL Server Enterprise Manager Is Unavailable**

There might be times when SQL Server Enterprise Manager cannot access SQL Server 2000 but a Query Analyzer session can. In these cases, you can use the **SP_WHO** SQL Server stored procedure to discover information about current connection count and connection details.

For more information, see "SP_WHO" in the Microsoft SQL Server 2000 Books Online, which you can install during SQL Server 2000 installation and then access from the Start menu through All Programs. Optionally, see "Getting Started with SQL Server Books Online" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=7107.

# Optimizing the MIIS Database

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The performance of a server running MIIS 2003 is significantly affected by the performance of the MIIS database. To maintain optimum performance, you might need to modify the default settings for SQL Server 2000 and the MIIS database.

**Configuring the Transaction Log to Optimize Performance**

The transaction log significantly affects performance. A transaction can be small or large depending on the type of data processing that is required. For example, renaming a forest can require the renaming of each object in the forest. Collectively, renaming all objects in the forest constitutes a single large transaction. A transaction can also be small, such as when a single user object is changed and that change must be synchronized among other data sources. In either case, whether the transaction is small or large, the performance of the transaction log can be significantly affected by disk performance, growth, and truncation (clearing).

To optimize your transaction log, you might need to do any of the following:

- Create the transaction log on a separate disk or redundant array of independent disks (RAID) device. The transaction log file is written serially; therefore, using a separate dedicated disk allows the disk heads to stay in place for the next write operation.

- When using a Full Recovery or Bulk-Logged Recovery model, be sure to clear the log periodically to prevent the log from filling up and diminishing performance. For more information about backing up and truncating the transaction log, see "Scheduling Full Database Backups by Using the Database Maintenance Plan Wizard" earlier in this topic.

- Never delete the transaction log file. Deleting the transaction log file can lead to corruption of the database.

- When running one or more large transactions, set the log file growth percentage to a number greater than 10%, which is the default setting. A single large transaction can generate a large amount of transaction log data. If the log file growth percentage is not set high enough to handle that growth, the number of transaction log growth requests can significantly slow or stall performance.

**To change the transaction log file growth**

1. In **Enterprise Manager**, right-click the **MicrosoftIdentityIntegrationServer** database, then click **Properties**, and then click the **Transaction Log** tab.

2. In **File Growth**, select **In megabytes**, and then enter a number equal to 1/50th (or up to four gigabytes, whichever is less) of space available on the drive where the transaction log resides.

**Clearing Management Agent Run History**

Another way to optimize the performance of a server running MIIS 2003 is to regularly clear the management agent run history. You can clear run histories from the MIIS 2003 Server up to a specified date by using MIISClearRunHistory in the Resource Tool Kit for Microsoft Identity Integration Server 2003, an unsupported set of tools that is available on the Microsoft Download Center at http://go.microsoft.com/fwlink/?LinkID=22788. Clearing run histories helps minimize the amount of space required for the MIIS database.

# Additional Resources (Maintaining the MIIS 2003 Database)

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

These resources contain additional information related to this topic.

**Related Information**

- "Getting Started with SQL Server Books Online" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=7107.
- The SQL Server 2000 Operations Guide on the Microsoft TechNet Web site at http://www.microsoft.com/technet/prodtechnol/sql/2000/maintain/sqlops0.mspx.
- "Features Supported by the Editions of SQL Server 2000" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=17515 for more information about the features of different editions of SQL Server 2000.
- "Backing Up and Restoring Databases" on the Microsoft MSDN Web site at http://go.microsoft.com/fwlink/?LinkID=10575.
- "Chapter 3 – Backing Up and Restoring Databases" in the Microsoft® SQL Server™ Administrator's Companion on the Microsoft TechNet Web site at http://go.microsoft.com/fwlink/?LinkID=15117 for more information about developing backup and restore plans for SQL Server databases.
- "Chapter 20 – RDBMS Performance Tuning Guide for Data Warehousing" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkID=17516.

**Related Help Topics**

- "MIISactivate: Server activation tool" in Microsoft Identity Integration Server 2003 Help.
- "MIISkmu: Encryption key management tool" in Microsoft Identity Integration Server 2003 Help.
- "Monitoring performance" in Help and Support Center for Windows Server 2003 for information about how to use System Monitor to create performance charts and logs.
- "Safeguarding your system" in Help and Support Center for Windows Server 2003 for information about creating an ASR set.
- "Performance Logs and Alerts" in Help and Support Center for Windows Server 2003 for information about setting and managing alerts.
- "The role of Microsoft SQL Server 2000" in Microsoft Identity Integration Server 2003 Help.
- "Backing up Microsoft Identity Integration Server 2003" in Microsoft Identity Integration Server 2003 Help.
- "Restoring Microsoft Identity Integration Server 2003" in Microsoft Identity Integration Server 2003 Help.
- "Activating a standby Microsoft Identity Integration Server 2003 server" in Microsoft Identity Integration Server 2003 Help.

**Related Microsoft Web Sites**

- "Microsoft SQL Server" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkID=1697.

- "Microsoft Identity Integration Server 2003" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkID=6999.

# Network Load Balancing Event Messages

Updated: July 1, 2003

Applies To: Windows Server 2003 with SP1

Published: July 1, 2003

**In This White Paper**


- Microsoft Windows 2000 Server and Windows Server 2003 Common NLB Events
- Microsoft Windows Server 2003-only NLB Events


This document contains a complete list of all Network Load Balancing (NLB) event messages, their causes, and detailed instructions for remedial actions, if any. The event severity and the format of each event are specified to aid the development of custom scripts.

The first section contains the events that are common for both Microsoft Windows Server 2000 and Microsoft Windows Server 2003 releases. The difference in the event formats and user actions between the two server versions is explained. The second section contains only those events that are new in Microsoft Windows Server 2003.

| **Note** |
| --- |
| The event source reported for all of the messages described in this document will be WLBS. |

**About Alert Severity**

Three levels of alert severity are indicated in Network Load Balancing event messages. Some events do not have an alert severity associated with them, and they are noted in this white paper.

- **E- Error.** A significant problem, such as loss of data or loss of functionality. Administrator action will likely be required to correct the error.
- **W - Warning.** An event that is not necessarily significant, but might indicate a possible future problem. Administrator action may or may not be required to correct the error.
- **I - For information only.** An event that typically describes the successful operation of an application, driver, or service. Typically no administrator action is required.

# Microsoft Windows 2000 Server and Windows Server 2003 Common NLB Events

Updated: July 1, 2003

Applies To: Windows Server 2003 with SP1

**Event ID: 1**

**Alert Severity:** Unused

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: %3. %4*

**Windows 2000 Server Event Text:** *%4 %1: %2. %3*

**Event ID: 2**

**Alert Severity:** Unused

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: %3. %4*

**Windows 2000 Server Event Text:** *%4 %1: %2. %3*

**Event ID: 3**

**Alert Severity:** Unused

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: %3. %4*

**Windows 2000 Server Event Text:** *%4 %1: %2. %3*

**Event ID: 4**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: %3 started. %4*

**Windows 2000 Server Event Text:** *%4 %1: %2 started. %3*

**Error Message Explanation:** This event is logged to inform you that cluster operations have started in response to a user-initiated action.

**User Action (Windows Server 2003):** No user action is required.

**User Action (Windows 2000 Server):** No user action is required.

**Event ID: 5**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Cluster mode started with host ID %2. %3*

**Windows 2000 Server Event Text:** *%4 %1 :cluster mode started with host ID %2. %3*

**Error Message Explanation:** This event is logged to inform you that cluster operations have started in response to a user-initiated action.

**User Action (Windows Server 2003):** No user action is required.

**User Action (Windows 2000 Server):** No user action is required.

**Event ID: 6**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Cluster mode stopped. %3. %4*

**Windows 2000 Server Event Text:** *%4 %1: cluster mode stopped. %2. %3*

**Error Message Explanation:** This event is logged to inform you that cluster operations have stopped in response to a user-initiated action.

**User Action (Windows Server 2003):** No user action is required.

**User Action (Windows 2000 Server):** No user action is required.

**Event ID: 7**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Internal error. Please contact technical support. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: Internal error. Please contact technical support. %2. %3*

**Error Message Explanation:** Network Load Balancing Clustering cannot complete the task because an error occurred.

**User Action (Windows Server 2003):** Contact your product support person. For more information about product support, see Help and Support. Alternatively, see Product Support Services at http://www.microsoft.com/homepage/ms.htm.Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**User Action (Windows 2000 Server):** Contact Microsoft Product Support Services. For more information about product support, see Help. Alternatively, see Product Support Services at http://www.microsoft.com/homepage/ms.htm. Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**Event ID: 8**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Error registering driver with NDIS. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: Error registering driver with NDIS. %2. %3*

**Error Message Explanation:** The Network Load Balancing driver cannot register with the Network Driver Interface Specification (NDIS). Therefore, the host will not respond to the TCP/IP traffic on this adapter.

**User Action (Windows Server 2003):** Restart your computer. If the problem continues, use Network Connections in Control Panel to disable and then re-enable the network adapter. If this does not resolve the problem, contact your product support person. For more information about product support, see Help and Support. Alternatively, see Product Support Services at http://www.microsoft.com/homepage/ms.htm. Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**User Action (Windows 2000 Server):** Restart your computer. If the problem continues, use Network Connections in Control Panel to disable and then re-enable the network adapter. If this does not resolve the problem, contact Microsoft Product Support Services. For more information about product support, see Help. Alternatively, see Product Support Services at http://www.microsoft.com/homepage/ms.htm. Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**Event ID: 9**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Driver does not support the media type that was requested by TCP/IP. NLB will not bind to this adapter. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: driver does not support media type that was requested by TCP/IP. %2. %3*

**Error Message Explanation:** In order for Network Load Balancing (NLB) to work, it must bind to an Ethernet network adapter. This event is logged when the NLB driver cannot bind to an Ethernet network adapter because it is not supported.

**User Action (Windows Server 2003):** Verify that a supported Ethernet network adapter is installed in this computer, and then bind NLB to it.

**User Action (Windows 2000 Server):** Verify that a supported Ethernet network or Fiber Distributed Data Interface (FDDI) network adapter is installed in this computer, and then bind NLB to it.

**Event ID: 10**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Driver could not allocate a required memory pool. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: driver could not allocate required memory pool. %2. %3*

**Error Message Explanation:** The Network Load Balancing (NLB) driver could not allocate memory.

**User Action (Windows Server 2003):** To increase the amount of available memory, try one of more of the following:

1. Add additional random access memory (RAM) to the host.

2. Close programs on the host that might be consuming memory.

3. Force the NLB driver to reload. Reloading will occur if you restart the host or disable all network adapters that are bound to NLB and then re-enable those adapters.

**User Action (Windows 2000 Server):** To increase the amount of available memory, try one or more of the following:

1. Add additional random access memory (RAM) to the host.

2. Close programs on the host that might be consuming memory.

3. Using **Network Connections** in Control Panel disable and then re-enable the network adapter that is bound to NLB.

**Event ID: 11**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Driver could not open device \DEVICE\ %3. %4*

**Windows 2000 Server Event Text:** *%4 %1: driver could not open device. %2. %3*

**Error Message Explanation:** Network Load Balancing (NLB) cannot bind to the network adapter. A possible reason is that the medium is not supported.

**User Action (Windows Server 2003):** Bind NLB to an Ethernet network adapter. If you are already using an Ethernet network adapter, examine the status code of the event message. The returned status code is included in the Data area at the bottom of the Event Properties dialog box.

The status code will correspond to a specific error condition as follows:

- 0xC0010019 = The adapter specified is not supported. In this case, you need to use a different adapter.

- 0xC000009A = The requested operation failed because NDIS could not allocate sufficient memory or initialize the state it uses to track an open binding. For this status, try adding additional memory.

- If the status code is not listed here or you cannot resolve the problem based on the information provided, contact your product support person. For more information about product support, see Help and Support. Alternatively, see Product Support Services at http://www.microsoft.com/homepage/ms.htm. Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**User Action (Windows 2000 Server):** Bind NLB to an Ethernet network adapter. If you are already using an Ethernet network adapter, examine the status code of the event message. The returned status code is included in the **Data** area at the bottom of the **Event Properties** dialog box.

The status code will correspond to a specific error condition as follows:

- 0xC0010019 = The adapter specified is not supported. In this case, you need to use a different adapter.

- 0xC000009A = The requested operation failed because NDIS could not allocate sufficient memory or initialize the state it uses to track an open binding. For this status, try adding additional memory.

- If the status code in the message is not listed here or you cannot resolve the problem based on the information provided, contact your product support person. For more information about product support, see Help. Alternatively, see Product Support Services at http://www.microsoft.com/homepage/ms.htm. Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**Event ID: 12**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Driver could not announce virtual NIC \DEVICE\ %3 to TCP/IP. %4*

**Windows 2000 Server Event Text:** *%4 %1: driver could not announce virtual NIC %2 to TCP/IP. %3*

**Error Message Explanation:** This event is logged to inform you that a call to Network Driver Interface Specification (NDIS) in order to bind the TCP/IP protocol to Network Load Balancing failed.

**User Action (Windows Server 2003):** No user action is required.

**User Action (Windows 2000 Server):** No user action is required.

**Event ID: 13**

**Alert Severity:** W

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Driver temporarily ran out of resources. Increase the %3 value under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\{GUID} key in the registry. %4*

**Windows 2000 Server Event Text:** *%4 %1: driver temporarily ran out of resources. Increase the %2 value under \SYSTEM\CurrentControlSet\Services\%4\Parameters key in the registry. %3*

**Error Message Explanation:** *When the Network Load Balancing (NLB) driver loads, NLB allocates blocks of memory for internal use. While this driver was loading, more memory was required than was initially allocated. Usually this is an indication that the load-balanced program is very heavily utilized.*

**User Action (Windows Server 2003):** *NLB does not automatically increase the amount of memory used for its internal processes. Instead, advanced users can manually increase the value in the registry. The path to the registry key specified in the logged event is incorrect. The correct path is: HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\WLBS\Parameters\Interface\{GUID}. The registry key to change is specified in the message. The GUID is the globally unique identifier of the network adapter that is experiencing the problem. This is the network adapter that is bound to NLB. If NLB is, or has been, bound to multiple network adapters, you can find the correct adapter by looking at the value for the ClusterIPAddress key for each GUID entry in the registry.*

*It is recommended that you double the existing value of this registry key. After increasing the value, you must re-bind NLB to the adapter that was experiencing the problem. To do this, do one of the following:*

- *Restart the host,*
- Disable and then re-enable the appropriate network adapter, or
- Remove and then re-add NLB to the appropriate network adapter

If increasing the value does not correct the problem, contact your product support person. For more information about product support, see Help and Support. Alternatively, see Product Support Services at http://www.microsoft.com/homepage/ms.htm.

Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

For more information about Windows Registry best practices, see Help and Support.

**User Action (Windows 2000 Server):** NLB does not automatically increase the amount of memory used for its internal processes. Instead, advanced users can manually increase the value of the specified registry key. It is recommended that you double the existing value. After increasing the value, reload NLB. To do this, do one of the following:

- *Restart the host,*
- Remove and then re-add NLB to the network adapter that is bound to NLB, or
- Disable and then re-enable the appropriate network adapter.

*If increasing the value does not correct the problem, contact Microsoft Product Support Services. For more information about product support, see Help. Alternatively, see Product Support Services at* http://www.microsoft.com/homepage/ms.htm.

Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

For more information about Windows Registry best practices, see Help.

**Event ID: 14**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Cluster network address %3 is invalid. Please check the NLB configuration and make sure that the cluster network address consists of six hexadecimal byte values separated by dashes. %4*

**Windows 2000 Server Event Text:** *%4 %1: cluster network address %2 is invalid. Please check %4 Setup page and make sure that it consists of six hexadecimal byte values separated by dashes. %3*

**Error Message Explanation:** *The cluster network MAC address is not in the following form xx-xx-xx-xx-xx-xx, where x is a hexadecimal value. Leading zeroes are optional.*

**User Action (Windows Server 2003):** *To change the MAC address to a valid form, open the***Properties** dialog box for this Network Load Balancing (NLB) cluster. If the **Network address** text box is available, NLB is set up for manual configuration of the network address and you should change the address to a correct MAC address.

*If the Network address text box is unavailable (it is dimmed), then NLB is configured to calculate the MAC address automatically and there is a problem doing so. Contact Microsoft Product Support Services for assistance. For more information about product support, see Help. Alternatively, see Product Support Services at* http://www.microsoft.com/homepage/ms.htm.

Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**User Action (Windows 2000 Server):** To change the MAC address to a valid form, open the Properties dialog box for this Network Load Balancing (NLB) cluster. If the **Network address** text box is available, NLB is set up for manual configuration of the network address and you should change the address to a correct MAC address.

*If the Network address text box is unavailable (it is dimmed), then NLB is configured to calculate the MAC address automatically and there is a problem doing so. Contact Microsoft Product Support Services for assistance. For more information about product support, see Help. Alternatively, see Product Support Services at* http://www.microsoft.com/homepage/ms.htm.

Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**Event ID: 15**

**Alert Severity:** W

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Dedicated IP address %3 is invalid. The driver will proceed assuming that there is no dedicated IP address. Please check the NLB configuration and make sure that the dedicated IP address consists of four decimal byte values separated by dots. %4*

**Windows 2000 Server Event Text:** *%4 %1: dedicated IP address %2 is invalid. Will proceed assuming that there is no dedicated IP address. Please check %4 Setup page and make sure that it consists of four decimal byte values separated by dots. %3*

**Error Message Explanation:** *The specified dedicated IP address is not valid. The cluster will continue to converge. However, you will not be able to use the dedicated IP address for such tasks as managing the host or for other non-cluster related communication.*

**User Action (Windows Server 2003):** *No user action is required. However, if you want to use a dedicated IP address, you must enter a valid address.*

**User Action (Windows 2000 Server):** *No user action is required. However, if you want to use a dedicated IP address, you must enter a valid address.*

**Event ID: 16**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Cluster IP address %3 is invalid. Please check the NLB configuration and make sure that the cluster IP address consists of four decimal byte values separated by dots. %4*

**Windows 2000 Server Event Text:** *%4 %1: cluster IP address %2 is invalid. Please check %4 Setup page and make sure that it consists of four decimal byte values separated by dots. %3*

**Error Message Explanation:** *This event is logged when the cluster IP address in the registry is not valid because it was manually changed to an incorrect address .*

**User Action (Windows Server 2003):** *Use Network Load Balancing Manager to change the cluster IP address to a valid IP address. For more information, see Help and Support.*

**User Action (Windows 2000 Server):** *Use Network Load Balancing to change the cluster IP address to a valid IP address. For*

*more information, see Help.*

**Event ID: 18**

**Alert Severity:** W

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Duplicate cluster subnets detected. The network may have been inadvertently partitioned. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: duplicate cluster subnets detected. The network may have been inadvertently partitioned. %2 %3*

**Error Message Explanation:** *The hosts in this cluster formed two or more separate clusters. This is usually caused by network partitioning that prevents the Network Load Balancing (NLB) heartbeats of one or more hosts from reaching all of the other hosts in the cluster. This event is logged when the network problem is corrected and the hosts converge back into a single cluster.*

**User Action (Windows Server 2003):** Investigate *the cause of the network partitioning. No investigation of NLB is necessary.*

**User Action (Windows 2000 Server):** *Investigate the cause of the network partitioning. No investigation of NLB is necessary.*

**Event ID: 19**

**Alert Severity:** W

**Windows Server 2003 Event Text:** NLB Cluster %2 %1: Could not allocate additional connection descriptors. Increase the %3 value under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\{GUID} key in the registry. %4

**Windows 2000 Server Event Text:** *%4 %1: could not allocate additional connection descriptors. Increase the %2 value under \SYSTEM\CurrentControlSet\Services\%4\Parameters key in the registry. %3*

**Error Message Explanation:** *Descriptors are used to track connections that are being load-balanced by Network Load Balancing (NLB). A fixed number of descriptors (controlled by a registry key value) are allocated when the NLB driver is loaded. This event is logged to warn you that the pool of descriptors has been exhausted. It might be necessary for NLB to reuse descriptors that are currently tracking existing active connections in order to track the state of new connections. This might have unexpected consequences and could result in legitimate connections being reset.*

**User Action (Windows Server 2003):** *Advanced users can increase the maximum number of descriptors that are allocated by changing the appropriate key in the registry. The registry key specified in the message is incorrect. Instead, you should increase the value of the following registry key:*
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\WLBS\Parameters

\Interface\{GUID}\MaxDescriptorAllocs.

*The GUID is the globally unique identifier of the network adapter that is experiencing the problem. This is the network adapter that is bound to NLB. If NLB is, or has been, bound to multiple network adapters, you can find the correct adapter by looking at the value for the ClusterIPAddress key for each GUID entry in the registry.*

*It is recommended that you double the existing value of this registry key. After increasing the value, you must re-bind NLB to the adapter that was experiencing the problem (as identified by the GUID in the message). To do this, do one of the following:*

- *Restart the host,*
- Disable and then re-enable the appropriate network adapter, or
- Remove and then re-add NLB to the appropriate network adapter.

*If increasing the value does not correct the problem, contact your product support person. For more information about product support, see Help and Support. Alternatively, see Product Support Services at* http://www.microsoft.com/homepage/ms.htm.

Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

For more information about Windows Registry best practices, see Help and Support.

**User Action (Windows 2000 Server):** Advanced users can double the number of descriptors that are allocated by changing the specified key in the registry. After increasing the value, reload NLB. To do this, do one of the following:

- *Restart the host,*

- Disable and then re-enable the appropriate network adapter, or
- Remove and then re-add NLB to the appropriate network adapter.

*If increasing the value does not correct the problem, contact Microsoft Product Support Services. For more information about product support, see Help. Alternatively, see Product Support Services at* http://www.microsoft.com/homepage/ms.htm.

Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

For more information about Windows Registry best practices, see Help.

**Event ID: 20**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Port rules with duplicate single host priority %3 detected on the network. Please check the NLB configuration on all hosts that belong to the cluster and make sure that they do not contain single host rules with duplicate priorities for the same port range. %4*

**Windows 2000 Server Event Text:** *%4 %1: port rules with duplicate single host priority %2 detected on the network. Please check %4 Setup dialog on all hosts that belong to the cluster and make sure that they do not contain single host rules with duplicate priorities for the same port range. %3*

**Error Message Explanation:** *When you configure a port rule as a single host, traffic to the port or ports governed by that port rule is handled exclusively by the host whose priority has the lowest numeric value. This event is logged when the local host's single host priority is identical to the single host priority of another host. The cluster will not converge until the problem is corrected.*

User Action (Windows Server 2003): *Using Network Load Balancing (NLB) Manager, edit the single host port rule on one of the hosts to give it a unique priority within the cluster. For more information about using NLB Manager to assign a unique priority to a host, see Help and Support.*

**User Action (Windows 2000 Server):** *Edit the single host port rule on one of the hosts to give it a unique priority within the cluster. For more information about using Network Load Balancing to assign a unique priority to a host, see Help.*

**Event ID: 21**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Host %3 does not have the same number or type of port rules as this host. Please check the NLB configuration on all hosts that belong to the cluster and make sure that they all contain the same number and type of port rules. %4*

**Windows 2000 Server Event Text:** *%4 %1: host %2 does not have the same number or type of port rules as this host. Please check %4 Setup dialog on all machines that belong to the cluster and make sure that they all contain the same number and the same type of port rules. %3*

**Error Message Explanation:** *This event is logged when another host in the cluster either contains a different number of port rules from the local host, or its configured port rules conflict with the port rules of the local host. The cluster will not converge until the problem is corrected.*

**User Action (Windows Server 2003):** *Using Network Load Balancing (NLB) Manager, edit the port rules as appropriate on each host so that the port rules are consistent on all hosts in the cluster. For more information about editing port rules in NLB, see Help and Support.*

**User Action (Windows 2000 Server):** *Edit the port rules as appropriate on each host so that the port rules are consistent on all hosts in the cluster. For more information about editing port rules in Network Load Balancing, see Help.*

**Event ID: 23**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Enabled traffic handling for rule containing port %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: enabled traffic handling for rule containing port %2. %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, this host will handle new load balanced traffic for the specified port rule.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 24**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Disabled ALL traffic handling for rule containing port %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: disabled ALL traffic handling for rule containing port %2. %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, network traffic arriving over ports governed by the specified port rule will no longer be handled by this host.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 25**

**Alert Severity:** W

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Port %3 not found in any of the valid port rules. %4*

**Windows 2000 Server Event Text:** *%4 %1: port %2 not found in any of the valid port rules. %3*

**Error Message Explanation:** *This event is logged when a user issues a command to control or query the state of a load-balanced port, but not applicable port rule is found. The command was ignored.*

**User Action (Windows Server 2003):** *No user action is required unless you expected this port to be governed by a port rule. In that case, you must configure a port rule to govern this port. For more information about configuring a port rule in Network Load Balancing, see Help and Support.*

**User Action (Windows 2000 Server):** *No user action is required unless you expected this port to be governed by a port rule. In that case, you must configure a port rule to govern this port. For more information about configuring a port rule in Network Load Balancing, see Help.*

**Event ID: 26**

**Alert Severity:** W

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Port rules cannot be enabled/disabled while cluster operations are stopped. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: port rules cannot be enabled/disabled while cluster operations are stopped. %2 %3*

**Error Message Explanation:** *This event is logged when a disable or enable command is issued to a host which has been either stopped or suspended. The command was ignored.*

**User Action (Windows Server 2003):** *Before issuing a disable or enable command, you must start (or resume and start) Network Load Balancing on the host.*

**User Action (Windows 2000 Server):** *Before issuing a disable or enable command, you must start (or resume and start) Network Load Balancing on the host.*

**Event ID: 28**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Host %3 converged with host(s) %4 as part of the cluster.*

**Windows 2000 Server Event Text:** *%4 %1: host %2 converged with host(s) %3 as part of the cluster.*

**Error Message Explanation:** *This event is logged to inform you that the specified host has joined the cluster.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 29**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Host %3 converged as DEFAULT host with host(s) %4 as part of the*

*cluster.*

**Windows 2000 Server Event Text:** *%4 %1: host %2 converged as DEFAULT host with host(s) %3 as part of the cluster.*

**Error Message Explanation:** *This event is logged to inform you that the specified host has joined the cluster. This host is now the default host and will handle all traffic that is not governed by any port rules.*

**User Action (Windows Server 2003):** *No user action is required. If you do not want this host to be the default host, reconfigure the host priority settings. For more information about reconfiguring host priorities in Network Load Balancing, see Help and Support.*

**User Action (Windows 2000 Server):** *No user action is required. If you do not want this host to be the default host, reconfigure the host priority settings. For more information about reconfiguring host priorities in Network Load Balancing, see Help.*

**Event ID: 30**

**Alert Severity:** W

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Dedicated network mask %3 is invalid. The driver will ignore the dedicated IP address and network mask. Please check the NLB configuration and make sure that the dedicated network mask consists of four decimal byte values separated by dots. %4*

**Windows 2000 Server Event Text:** *%4 %1: dedicated network mask %2 is invalid. Will ignore dedicated IP address and network mask. Please check %4 Setup page and make sure that it consists of four decimal byte values separated by dots. %3*

**Error Message Explanation:** *This event is logged when the subnet mask entered for the dedicated IP address is not valid. The cluster will continue to operate. However, you cannot use the dedicated IP address, so you cannot use it to manage the host or for other non-cluster-related communication.*

**User Action (Windows Server 2003):** *No user action is required. The cluster will continue to operate. However, if you want to use the dedicated IP address, enter a valid subnet mask. For more information about entering a subnet mask in Network Load Balancing, see Help and Support.*

**User Action (Windows 2000 Server):** *No user action is required. The cluster will continue to operate. However, if you want to use the dedicated IP address, enter a valid subnet mask. For more information about entering a subnet mask in Network Load Balancing, see Help.*

**Event ID: 31**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Cluster network mask %3 is invalid. Please check the NLB configuration and make sure that the cluster network mask consists of four decimal byte values separated by dots. %4*

**Windows 2000 Server Event Text:** *%4 %1: cluster network mask %2 is invalid. Please check %4 Setup page and make sure that it consists of four decimal byte values separated by dots. %3*

**Error Message Explanation:** *This event is logged when the subnet mask entered for the cluster IP address is not valid.*

**User Action (Windows Server 2003):** *Correct the subnet mask for the cluster IP address. For more information about correcting the subnet mask in Network Load Balancing, see Help and Support.*

**User Action (Windows 2000 Server):** *Correct the subnet mask for the cluster IP address. For more information about correcting the subnet mask in Network Load Balancing, see Help.*

**Event ID: 32**

**Alert Severity:** W

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Both the dedicated IP address and network mask must be specified. The driver will ignore the dedicated IP address and network mask. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: both dedicated IP address and network mask have to be entered. Will ignore dedicated IP address and network mask. %2 %3*

**Error Message Explanation:** *This event is logged when either the dedicated IP address or its subnet mask was not entered. The cluster will continue to operate. However, you cannot use the dedicated IP address, so you cannot use it to manage the host or for other non-cluster-related communication.*

**User Action (Windows Server 2003):** No user action is required. The cluster will continue to operate. However, if you want to use the dedicated IP address, enter a valid dedicated IP address and subnet mask. *For more information about entering a*

*dedicated IP address and subnet mask in Network Load Balancing, see Help and Support.*

**User Action (Windows 2000 Server):** No user action is required. The cluster will continue to operate. However, if you want to use the dedicated IP address, enter a valid dedicated IP address and subnet mask. *For more information about entering a dedicated IP address and subnet mask in Network Load Balancing, see Help.*

**Event ID: 33**

**Alert Severity:** W

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Remote control request rejected from %3 due to an incorrect password. %4*

**Windows 2000 Server Event Text:** *%4 %1: remote control request rejected from host %2 due to incorrect password. %3*

**Error Message Explanation:** *This event is logged when remote control is enabled on this server and a remote control request was received for the cluster. However, the remote control password received is not correct. A legitimate user might have accidentally entered the wrong password, or someone else might have tried to guess the password.*

**User Action (Windows Server 2003):** The source IP address of the remote control request is provided in the message. Use this information to determine the computer from which the remote control command was issued. This might help identify who issued the remote control command.

**Warning:** The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service, and information disclosure. We strongly recommend that you do not enable remote control, and that instead you use NLB Manager or other remote management tools such as Windows Management Instrumentation (WMI).

**User Action (Windows 2000 Server):** The source IP address of the remote control request is provided in the message. Use this information to determine the computer from which the remote control command was issued. This might help identify who issued the remote control command.

**Warning:** The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service, and information disclosure. We strongly recommend that you do not enable remote control, and that instead you use NLB Manager or other remote management tools such as Windows Management Instrumentation (WMI).

**Event ID: 34**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Error querying parameters from the registry key \HLKM\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\%3. Please fix the NLB configuration and run wlbs reload followed by wlbs start. %4*

**Windows 2000 Server Event Text:** *%4 %1: error querying parameters from the registry key %2. Please run %4 setup dialog to fix the problem and run %4 reload followed by %4 start. %3*

**Error Message Explanation:** *This event is logged when there is an error processing the Network Load Balancing settings from the registry. For example, the registry key could not be opened, or a value that is not valid was read from one or more registry parameters. The most likely cause of this error is manual configuration of the registry instead of making configuration changes through the NLB properties dialog box.*

**User Action (Windows Server 2003):** *The registry hive specified in the message is incorrect. The correct hive is: HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\WLBS\Parameters\Interface\{GUID}. The GUID identified in the message is the correct GUID. The GUID reference is the globally unique identifier of the network adapter that is experiencing the problem.*

*Advanced users should examine the registry keys in this hive to determine which registry parameter caused the event. To do this, consider comparing the registry entries of a host in the cluster that is working correctly to the registry entries of the host that is reporting the error. Also examine the event log for related events that might indicate more specific problems with the registry.*

*For more information about Windows Registry best practices, see Help and Support.*

**User Action (Windows 2000 Server):** *Refer to the message to determine the registry hive in which the error occurred. Advanced users should examine the registry keys in that hive to determine which registry parameter caused the event. To do this, consider comparing the registry entries of a host in the cluster that is working correctly to the registry entries of the host that is reporting the error. Also examine the event log for related events that might indicate more specific problems with the registry.*

*For more information about NLB registry settings, see the Windows 2000 Server Resource Kit . For more information about*

*Windows Registry best practices, see Help.*

**Event ID: 35**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Cluster mode cannot be enabled due to parameter errors. All traffic will be passed through to TCP/IP. Restart cluster operations after fixing the problem by running wlbs reload followed by wlbs start. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: cluster mode cannot be enabled due to parameter errors. All traffic will be passed to TCP/IP. Restart cluster operations after fixing the problem by running %4 reload followed by %4 start. %2 %3*

**Error Message Explanation:** *This event is probably logged because of incorrectly configured settings that are being read from the registry; in particular, IP addresses, IP subnet masks, or MAC addresses. Until this problem is corrected, this host will remain stopped and will not join the cluster.*

**User Action (Windows Server 2003):** *Look for additional Network Load Balancing (NLB) events that indicate errors reading from the registry or other configuration problems. You can also use NLB Manager to diagnose configuration problems. After you determine the specific error, correct the problem, and then run nlb reload' followed by 'nlb start'. If you cannot determine the specific error, contact your product support person.*

*For more information about product support, see Help and Support. Alternatively, see Product Support Services at* http://www.microsoft.com/homepage/ms.htm. *Under the* **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**User Action (Windows 2000 Server):** *Look for additional Network Load Balancing events that indicate errors reading from the registry or other configuration problems. After you determine the specific error, correct the problem, and then run wlbs reload' followed by wlbs start'. If you cannot determine the specific error, contact Microsoft Product Support Services. For more information about product support, see Help. Alternatively, see Product Support Services at* http://www.microsoft.com/homepage/ms.htm. *Under the* **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

**Event ID: 36**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Registry parameters successfully reloaded. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: registry parameters successfully reloaded. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that the registry parameters have been reloaded successfully.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 37**

**Alert Severity:** W

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Version mismatch between the driver and control programs. Please make sure that you are running the same version of all NLB components. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: version mismatch between the driver and control programs. Please make sure that you are running the same version. %2 %3*

**Error Message Explanation:** *This event is logged when incompatible Network Load Balancing binaries are installed, for example, if a Windows 2000 version of wlbs.exe is installed with a version of wlbs.sys created for a member of the Windows Server 2003 family of operating systems.*

**User Action (Windows Server 2003):** *Replace the conflicting binary.*

**User Action (Windows 2000 Server):** *Replace the conflicting binary.*

**Event ID: 38**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Adjusted traffic handling for rule containing port %3. %4*

**Windows 2000 Server Event Text:** *%4 %1: adjusted traffic handling for rule containing port %2. %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, the load weight of a port rule has changed.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 39**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Disabled NEW traffic handling for rule containing port %3. %4*

**Windows 2000 Server Event Text:** *%4 %1: disabled NEW traffic handling for rule containing port %2. %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, new network traffic is disabled for this specific port rule.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 40**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Adjusted traffic handling for all port rules. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: adjusted traffic handling for all port rules. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, the load weight of all port rules has changed.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 41**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Disabled NEW traffic handling for all port rules. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: disabled NEW traffic handling for all port rules. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, new network traffic is disabled for all port rules.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 41**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Disabled NEW traffic handling for all port rules. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: disabled NEW traffic handling for all port rules. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, new network traffic is disabled for all port rules.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 42**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Enabled traffic handling for all port rules. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: enabled traffic handling for all port rules. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, this host will handle new load-balanced traffic for all port rules.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 43**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Disabled ALL traffic handling for all port rules. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: disabled ALL traffic handling for all port rules. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, network traffic arriving over ports governed by any port rule will no longer be handled by this host.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 43**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Disabled ALL traffic handling for all port rules. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: disabled ALL traffic handling for all port rules. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, network traffic arriving over ports governed by any port rule will no longer be handled by this host.*

**User Action (Windows Server 2003):** *No user action is required.*

**User Action (Windows 2000 Server):** *No user action is required.*

**Event ID: 44**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Connection draining started. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: connection draining started. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, this host will stop accepting new network traffic. When there are no more active connections, this host will stop its cluster operations.*

**User Action (Windows Server 2003):** *No user action is required. For more information about the drain command, see Network Load Balancing in Help and Support.*

**User Action (Windows 2000 Server):** *No user action is required. For more information about the drain command, see Network Load Balancing in Help.*

**Event ID: 45**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Connection draining interrupted. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: connection draining interrupted. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that the host was draining when a stop or suspend command was issued by a user.*

**User Action (Windows Server 2003):** *No user action is required. For more information about the drain command, see Network Load Balancing in Help and Support.*

**User Action (Windows 2000 Server):** *No user action is required. For more information about the drain command, see Network Load Balancing in Help.*

**Event ID: 46**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: %3 remote control request received from %4.*

**Windows 2000 Server Event Text:** *%4 %1: %2 remote control request received from host %3.*

**Error Message Explanation:** *This event is logged when the cluster has received and is processing a remote control command.*

**User Action (Windows Server 2003):** *Verify that you know who is issuing the remote control commands and that this person is authorized to do so. The source IP address of the remote control request is provided in the message. Use this to determine the computer from which the remote control command was issued. This might help identify the person who issued the remote control command.*

**Warning:** The Network Load Balancing (NLB) remote control option presents many security risks, including the possibility of data tampering, denial of service, and information disclosure. It is strongly recommended that you do not enable remote control, and that you instead use NLB Manager or other remote management tools such as Windows Management Instrumentation (WMI).

**User Action (Windows 2000 Server):** *Verify that you know who is issuing the remote control commands and that this person is authorized to do so. The source IP address of the remote control request is provided in the message. Use this to determine the computer from which the remote control command was issued. This might help identify the person who issued the remote control command.*

**Warning:** The Network Load Balancing (NLB) remote control option presents many security risks, including the possibility of data tampering, denial of service, and information disclosure. It is strongly recommended that you do not enable remote control, and that you instead use NLB Manager or other remote management tools such as Windows Management Instrumentation (WMI).

**Event ID: 47**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Cluster mode control resumed. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: cluster mode control resumed. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, cluster operations have resumed.*

**User Action (Windows Server 2003):** *No user action is required. For more information about the resume command, see Network Load Balancing in Help and Support.*

**User Action (Windows 2000 Server):** *No user action is required. For more information about the resume command, see Network Load Balancing in Help.*

**Event ID: 48**

**Alert Severity:** I

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Cluster mode control suspended. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: cluster mode control suspended. %2 %3*

**Error Message Explanation:** *This event is logged to inform you that, in response to a user-initiated action, cluster operations have been suspended.*

**User Action (Windows Server 2003):** *No user action is required. For more information about the suspend command, see Network Load Balancing in Help and Support.*

**User Action (Windows 2000 Server):** *No user action is required. For more information about the suspend command, see Network Load Balancing in Help.*

**Event ID: 49**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *This error cannot occur on Windows Server 2003 systems.*

**Windows 2000 Server Event Text:** *%4 %1: Network Load Balancing can only be installed on systems running Windows 2000 Advanced Server or Datacenter Server. %2 %3*

**Error Message Explanation:** *To use Network Load Balancing, you must install Windows 2000 Advanced Server or Windows 2000 Datacenter Server on this computer. NLB is also available on all members of the Windows Server 2003 family.*

**User Action (Windows Server 2003):** *This error cannot occur on Windows Server 2003 systems.*

**User Action (Windows 2000 Server):** *To use NLB, install Windows 2000 Advanced Server or Datacenter Server. NLB is also available on all members of the Windows Server 2003 family.*

**Event ID: 50**

**Alert Severity:** E

**Windows Server 2003 Event Text:** *NLB Cluster %2 %1: Cannot add cluster MAC address. The maximum number of multicast MAC addresses are already assigned to the NIC. %3 %4*

**Windows 2000 Server Event Text:** *%4 %1: cannot add cluster MAC address. The maximum number of multicast MAC addresses are already assigned to the NIC. %2 %3*

**Error Message Explanation:** *The network adapter has the maximum number of multicast MAC addresses assigned to it already. Other programs might have added multicast addresses with the result that Network Load Balancing cannot add the MAC addresses needed to operate an NLB cluster in multicast mode.*

**User Action (Windows Server 2003):** *Perform one or more of the following to correct the problem.*

1. Use a network adapter that allows an increased number of multicast MAC addresses to be configured.

2. Operate the cluster in unicast mode.

*For more information about unicast and multicast mode, see Network Load Balancing in Help and Support.*

**User Action (Windows 2000 Server):** Perform one or more of the following to correct the problem.

1. Use a network adapter that allows an increased number of multicast MAC addresses to be configured.

2. Operate the cluster in unicast mode.

*For more information about unicast and multicast mode, see Network Load Balancing in Help.*

# Microsoft Windows Server 2003-only NLB Events

Updated: July 1, 2003

Applies To: Windows Server 2003 with SP1

*These events are new in Windows Server 2003 and are therefore only logged on computers running Windows Server 2003.*

**Event ID: 51**

**Alert Severity:** E

**Event Text:** *NLB Cluster %2 %1: Cannot bind NLB to this adapter due to memory allocation failure. %3 %4*

**Error Message Explanation:** *This event is logged when there is not enough memory to bind Network Load Balancing to this network adapter.*

**User Action:** To increase the amount of available memory, try one or more of the following:

1. Add additional random access memory (RAM) to the host.
2. Close programs on the host that might be consuming memory.
3. Force the NLB driver to reload. Reloading will occur if you restart the host or disable all network adapters that are bound to NLB and then re-enable those adapters.

**Event ID: 52**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1: Registry parameters successfully reloaded without resetting the driver. %3 %4*

**Error Message Explanation:** Certain changes, such as adding a new port rule, require Network Load Balancing to be stopped and restarted. Other changes do not require NLB to be stopped and restarted. This event is logged to inform you that changes were made to the NLB configuration which did not require NLB to be stopped and restarted.

**User Action:** No user action is required.

**Event ID: 53**

**Alert Severity:** E

**Event Text:** *NLB Cluster %2 %1: The adapter to which NLB is bound does not support dynamically changing the MAC address. NLB will not bind to this driver. %3 %4*

**Error Message Explanation:** To be operated in unicast mode, Network Load Balancing must alter the unicast MAC address on the network adapter that it will be using. This event is logged when the network adapter to which NLB is trying to bind does not support dynamic changing of its unicast MAC address.

User Action: Replace the network adapter with one that supports changing the unicast MAC address dynamically. Alternatively, operate the cluster in multicast mode.

**Event ID: 54**

**Alert Severity:** W

**Event Text:** *NLB Cluster %2 %1: Host %3* MAY be a Windows 2000 or NT 4.0 host participating in a cluster utilizing virtual cluster support. Virtual Clusters are only supported in a homogeneous Windows Server 2003 cluster. %4

**Error Message Explanation:** Virtual clusters, which are Network Load Balancing clusters where specific port rules are assigned to specific virtual IP addresses, are configured on this host, but at least one host in this cluster has one or more port rules that are not valid and appear to be missing the virtual IP address specification. Therefore, either:

- A Windows 2000 or Windows NT 4.0 host exists in the cluster. Virtual clusters require a cluster containing only hosts running products in the Windows Server 2003 family, or
- A Windows Server 2003 host with an incorrectly configured port rule exists in the cluster.

The cluster will not converge until the problem is corrected.

**User Action:** If there is a Windows 2000 or Windows NT 4.0 host in the cluster, either remove it or stop using virtual clusters on the Windows Server 2003 hosts.

If the cluster contains only Windows Server 2003 hosts, then use NLB Manager to check the port rule configuration of all hosts to ensure that all hosts have compatible port rule configurations.

For more information about virtual clusters and port rules, see Network Load Balancing in Help and Support.

**Event ID: 55**

**Alert Severity:** E

**Event Text:** *NLB Cluster %2 %1:* Inconsistent bi-directional affinity (BDA) teaming configuration detected on host %3. The team in which this cluster participates will be marked inactive and this cluster will remain in the converging state until consistent teaming configuration is detected. %4

**Error Message Explanation:** A bidirectional affinity (BDA) team is a set of adapters on a single host that act in concert. This event is logged when the local host and the remote host specified in the message have inconsistent BDA configurations. The BDA configuration must be identical on all hosts in a cluster. This is enforced through the Network Load Balancing heartbeat mechanism. The cluster will not converge until the problem is corrected.

**User Action:** If this event is followed by an NLB 58 event in Event Viewer, no user action is needed. If this event is not followed by an NLB 58 event, then you need to change your BDA configuration. BDA configuration is specified by the Microsoft Internet Security and Acceleration Server (ISA Server) configuration. For more information about configuring BDA, see the relevant ISA Server documentation.

**Event ID: 56**

**Alert Severity:** E

**Event Text:** *NLB Cluster %2 %1:* Invalid bi-directional affinity (BDA) teaming port rules detected. Each member of a BDA team may have no more than one port rule whose affinity must be single or class C if multiple host filtering is specified. To add this cluster to a team, correct the port rules and run 'wlbs reload' followed by 'wlbs start'. %3 %4

**Error Message Explanation:** A bidirectional affinity (BDA) team is a set of adapters on a single host that act in concert. A team ID identifies the different network adapters that belong to the same BDA team. The team ID is a {GUID} that uniquely identifies the team to which an adapter belongs. This event is logged when this adapter has been configured as a member of a BDA team, but the specified team ID is either missing or not valid.

**User Action:** Correct the BDA team ID, and then run nlb reload' followed by 'nlb start'. BDA configuration is specified by the Microsoft Internet Security and Acceleration Server (ISA Server) configuration. For more information about configuring BDA, see the relevant ISA Server documentation.

**Event ID: 57**

**Alert Severity:** E

**Event Text:** *NLB Cluster %2 %1:* Invalid bi-directional affinity (BDA) team ID detected. Team IDs must be a GUID of the form {xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx}. To add this cluster to a team, correct the team ID and run 'wlbs reload' followed by 'wlbs start'. %3 %4

**Error Message Explanation:** A bidirectional affinity (BDA) team is a set of adapters on a single host that act in concert. This event is logged when this adapter has been configured as a member of a BDA team, but the configured port rules do not comply with the requirements for using BDA. The use of BDA requires a single port rule that covers all ports (0 65535), configured as either Single or Class C affinity.

**User Action:** Correct the BDA teaming port rules, and then run nlb reload' followed by 'nlb start'. BDA configuration is specified by the Microsoft Internet Security and Acceleration Server (ISA Server) configuration. For more information about configuring BDA, see the relevant ISA Server documentation.

**Event ID: 58**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Consistent bi-directional affinity (BDA) teaming configuration detected again. The team in which this cluster participates has been re-activated. %3 %4

**Error Message Explanation:** A bidirectional affinity (BDA) team is a set of adapters on a single host that act in concert. This event is logged to inform you that a valid and consistent BDA configuration has been detected and therefore the team has

been activated.

**User Action:** No user action is required.

**Event ID: 59**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* The bi-directional affinity (BDA) team which this cluster has attempted to join already has a designated master. This cluster will not join the team and will remain in the stopped state until the parameter errors are corrected and the cluster is restarted by running 'wlbs reload' followed by 'wlbs start'. %3 %4

**Error Message Explanation:** A bidirectional affinity (BDA) team is a set of adapters on a single host that act in concert. A single adapter from each team must be assigned the role of master for that team. This event is logged when this adapter attempts to join a BDA team as the master but the team already has an active master.

**User Action:** Correct the BDA team parameters, and then run nlb reload' followed by 'nlb start'. BDA configuration is specified by the Microsoft Internet Security and Acceleration Server (ISA Server) configuration. For more information about configuring BDA, see the relevant ISA Server documentation.

**Event ID: 60**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* The bi-directional affinity (BDA) team in which this cluster participates has no designated master. The team is inactive and will remain inactive until a master for the team is designated and consistent BDA teaming configuration is detected. %3 %4

**Error Message Explanation:** A bidirectional affinity (BDA) team is a set of adapters on a single host that act in concert. One adapter on each cluster host must be the master for that team. This event is logged when the team that this adapter has just joined does not have an active master yet. This might be a transient condition due to the binding order of adapters in the system.

**User Action:** If this event is followed by an NLB 61 event in Event Viewer, no user action is required.

If this event is not followed by an NLB 61 event, then you need to change your BDA configuration to make it consistent across all hosts. BDA configuration is specified by the Microsoft Internet Security and Acceleration Server (ISA Server) setup. For more information about configuring BDA, see the relevant ISA Server documentation.

**Event ID: 61**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* This cluster has joined a bi-directional affinity (BDA) team as the designated master. If the rest of the team has been consistently configured, the team will be activated. %3 %4

**Error Message Explanation:** A bidirectional affinity (BDA) team is a set of adapters on a single host that act in concert. A single adapter from each team must be assigned the role of master for that team. This event is logged to inform you that this adapter has joined a BDA team as the master.

**User Action:** No user action is required.

**Event ID: 62**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* This cluster has left a bi-directional affinity (BDA) team in which it was the designated master. If other adapters are still participating in the team, the team will be marked inactive and will remain inactive until a master for the team is designated and consistent BDA teaming configuration is detected. %3 %4

**Error Message Explanation:** A bidirectional affinity (BDA) team is a set of adapters on a single host that act in concert. One adapter on each cluster host must be the master for that team. This event is logged when the master has just left the BDA team and a new master must be configured for the team to operate. Note that this can be a transient condition caused by a re-configuration of BDA teaming.

**User Action:** If this event is followed by an NLB 61 event in Event Viewer, no user action is required.

If this event is not followed by an NLB 61 event, then you need to change your BDA configuration. BDA configuration is specified by the Microsoft Internet Security and Acceleration Server (ISA Server) setup. For more information about configuring BDA, see the relevant ISA Server documentation.

**Event ID: 63**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Initiating convergence on host %3. Reason: Host %4 is joining the cluster.

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** If this event is followed by an NLB 28 event, NLB 29 event, NLB 71 event, or NLB 72 event in Event Viewer, no user action is required.

If this event is not soon followed by one of these other events, use Network Load Balancing (NLB) Manager to determine the specific configuration problem and then correct as appropriate. For more information about Network Load Balancing, see Help and Support.

**Event ID: 64**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Initiating convergence on host %3. Reason: Host %4 has conflicting configuration.

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** If this event is followed by an NLB 28 event, NLB 29 event, NLB 71 event, or NLB 72 event in Event Viewer, no user action is required.

If this event is not soon followed by one of these other events, use Network Load Balancing (NLB) Manager to determine the specific configuration problem and then correct as appropriate. For more information about Network Load Balancing, see Help and Support.

**Event ID: 65**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Initiating convergence on host %3. Reason: Host %4 is converging for an unknown reason.

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** If this event is followed by an NLB 28 event, NLB 29 event, NLB 71 event, or NLB 72 event in Event Viewer, no user action is required.

If this event is not soon followed by one of these other events, use Network Load Balancing (NLB) Manager to determine the specific configuration problem and then correct as appropriate. For more information about Network Load Balancing, see Help and Support.

**Event ID: 66**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Initiating convergence on host %3. Reason: Host %4 is configured with the same host ID.

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** If this event is followed by an NLB 28 event, NLB 29 event, NLB 71 event, or NLB 72 event in Event Viewer, no user action is required.

If this event is not soon followed by one of these other events, use Network Load Balancing (NLB) Manager to determine the specific configuration problem and then correct as appropriate. For more information about Network Load Balancing, see Help and Support.

**Event ID: 67**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Initiating convergence on host %3. Reason: Host %4 is configured with a conflicting number of port rules.

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** If this event is followed by an NLB 28 event, NLB 29 event, NLB 71 event, or NLB 72 event in Event Viewer, no user action is required.

If this event is not soon followed by one of these other events, use Network Load Balancing (NLB) Manager to determine the specific configuration problem and then correct as appropriate. For more information about Network Load Balancing, see Help and Support.

**Event ID: 68**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Initiating convergence on host %3. Reason: Host %4 has changed its port rule configuration.

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** If this event is followed by an NLB 28 event, NLB 29 event, NLB 71 event, or NLB 72 event in Event Viewer, no user action is required.

If this event is not soon followed by one of these other events, use Network Load Balancing (NLB) Manager to determine the specific configuration problem and then correct as appropriate. For more information about Network Load Balancing, see Help and Support.

**Event ID: 69**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Initiating convergence on host %3. Reason: Host %4 is leaving the cluster.

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** If this event is followed by an NLB 28 event, NLB 29 event, NLB 71 event, or NLB 72 event in Event Viewer, no user action is required.

If this event is not soon followed by one of these other events, use Network Load Balancing (NLB) Manager to determine the specific configuration problem and then correct as appropriate. For more information about Network Load Balancing, see Help and Support.

**Event ID: 71**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Host %3 converged with %4 host(s) as part of the cluster. Note: The list of cluster members was too large to fit in this event log - see the first data dump entry for a binary map of the cluster membership or run 'wlbs query' to see a list of cluster members.

**Error Message Explanation:** This event is logged to inform you that the specified host has joined the cluster.

**User Action:** No user action is required. If you do not want this host to be the default host, reconfigure the host priority settings. For more information about Network Load Balancing, see Help and Support.

**Event ID: 72**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Host %3 converged as DEFAULT host with %4 host(s) as part of the cluster. Note: The list of cluster members was too large to fit in this event log - see the first data dump entry for a binary map of the cluster membership or run 'wlbs query' to see a list of cluster members.

**Error Message Explanation:** This event is logged to inform you that the specified host has joined the cluster. This host will be the default host and will handle all traffic that is not governed by any port rules.

**User Action:** No user action is required. If you do not want this host to be the default host, reconfigure the host priority settings. For more information about Network Load Balancing, see Help and Support.

**Event ID: 73**

**Alert Severity:** E

**Event Text:** *NLB Cluster %2 %1:* Cluster IGMP multicast IP address %3 is invalid. Please check the NLB configuration and make sure that the cluster IGMP multicast IP address consists of four decimal byte values separated by dots. %4.

**Error Message Explanation:** The Internet Group Management Protocol (IGMP) multicast IP address in the registry is not valid.

**User Action:**

1. Open the following registry key: HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\WLBS\Parameters\Interface\{GUID}, where the GUID is the globally unique identifier of the network adapter that is experiencing the problem. This is the network adapter that is bound to Network Load Balancing. If NLB is, or has been, bound to multiple network adapters, you can find the correct adapter by looking at the value for the ClusterIPAddress key for each GUID.

2. Examine the IPtoMcastIP key and verify that it is a non-zero value. If it is a value of zero, change it to a value of one.

3. In Network Connections, open NLB Properties (do not use NLB Manager).

4. Click OK. This action forces NLB to recalculate the IGMP multicast address automatically and subsequently reload the NLB driver.

*If these actions do not resolve the problem, contact Microsoft Product Support Services. For more information about product support, see Help and Support. Alternatively, see Product Support Services at* http://www.microsoft.com/homepage/ms.htm. Under the **Resources** area in the navigation pane, point to **Support**, and then click **Support Home**.

For more information about Windows Registry best practices, see Help and Support.

**Event ID: 74**

**Alert Severity:** W

**Event Text:** *NLB Cluster %2 %1:* Unable to update the NLB host state in the registry. This will affect the ability of NLB to persist the current state across a reboot if NLB has been configured to do so. To predict the state of NLB at the next reboot, please delete the HKLM\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\{GUID}\%3 registry value, which will cause the state at the next reboot to revert to the user-specified preferred initial host state. %4.

**Error Message Explanation:** A Network Load Balancing (NLB) host can be in one of three states: started, stopped or suspended. The preferred state of NLB on startup can be configured to be any one of the three available states. Each time the server restarts, the host assumes that initial state. However, NLB can be configured to remember a suspended state. This can be used to override the preferred initial state if a host is restarted while suspended and this option has been selected. The host will remain suspended upon completion of the restart, regardless of the specified initial state. NLB therefore must track the last known state of the host to correctly configure the host during a restart. If NLB cannot save that information to the registry, the startup state of the cluster at this restart and subsequent restarts might not be what is expected.

**User Action:** The registry key specified in the message is incorrect. To be able to predict the state of the NLB host at the next restart, advanced users should delete the following registry key, if it exists:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\WLBS\Parameters\Interface\{GUID}\HostState. The GUID is the globally unique identifier of the network adapter that is experiencing the problem. This will be the network adapter that is bound to NLB. If NLB is, or has been, bound to multiple network adapters, you can find the correct adapter by looking at the value for the ClusterIPAddress key for each GUID entry in the registry.

Deleting this key ensures that on the next restart, NLB will use the preferred initial host state.

For more information about Windows Registry best practices, see Help and Support.

**Event ID: 75**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* Current NLB host state successfully updated in the registry.

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** No user action required.

**Event ID: 76**

**Alert Severity:** W

**Event Text:** *NLB Cluster %2 %1:* Last known NLB host state not found in the registry. The state of this host will revert to the user-specified preferred initial host state. %3 %4

**Error Message Explanation:** A Network Load Balancing (NLB) host can be in one of three states: started, stopped or suspended. The preferred state of NLB on startup can be configured to be any one of the three available states. Each time the server restarts, the host assumes that initial state. However, NLB can be configured to remember a suspended state. This can be used to override the preferred initial state if a host is restarted while suspended and this option has been selected. The host will remain suspended upon completion of the restart, regardless of the specified initial state. NLB therefore must track the last known state of the host to correctly configure the host during a restart. If NLB cannot read that information from the registry, the startup state of the cluster at this restart might not be what is expected.

**User Action:** No user action is required. NLB will attempt to correct this condition over time.

**Event ID: 77**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* This host is configured to persist 'Started' states across reboots. Because the last known state of this host was started, the host will be re-started regardless of the user-specified preferred initial host state. To stop the host, use 'wlbs stop' or 'wlbs suspend'. %3 %4

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** No user action is required. If you want to change the default state, use Network Load Balancing (NLB) Manager to make the change. For more information about Network Load Balancing, see Help and Support.

**Event ID: 78**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* This host is configured to persist 'Stopped' states across reboots. Because the last known state of this host was stopped, the host will remain stopped regardless of the user-specified preferred initial host state. To start the host, use 'wlbs start'. %3 %4

**Error Message Explanation:** This event is logged for informational purposes only.

User Action: No user action is required. If you want to change the default state, use Network Load Balancing (NLB) Manager to make the change. For more information about Network Load Balancing, see Help and Support.

**Event ID: 79**

**Alert Severity:** I

**Event Text:** *NLB Cluster %2 %1:* This host is configured to persist 'Suspended' states across reboots. Because the last known state of this host was suspended, the host will remain suspended regardless of the user-specified preferred initial host state. To start the host, use 'wlbs resume' followed by 'wlbs start'. %3 %4

**Error Message Explanation:** This event is logged for informational purposes only.

**User Action:** No user action is required. If you want to change the default state, use Network Load Balancing (NLB) Manager to make the change. For more information about Network Load Balancing, see Help and Support.

**Event ID: 80**

**Alert Severity:** W

**Event Text:** *NLB Cluster %2 %1:* Windows 2000 or NT 4.0 host(s) are participating in this cluster. Consult the relevant NLB documentation concerning the potential limitations of operating a mixed cluster. %3 %4

**Error Message Explanation:** Certain features of Network Load Balancing are not designed to work in a mixed cluster, that is, in a cluster with hosts running a product in the Windows Server 2003 family along with hosts running Windows 2000 or Windows NT 4.0. These features include virtual clusters, bidirectional affinity, and IPSec load balancing (such as VPN).

**User Action:** These features can be used if the cluster contains only hosts running a product in the Windows Server 2003 family. If you require any of these features, remove all hosts running Windows 2000 and Windows NT 4.0 from the cluster. For more information about Network Load Balancing, see Help and Support.

**Event ID: 81**

**Alert Severity:** W

**Event Text:** *NLB Cluster %2 %1:* NLB was unable to open the TCP connection callback object and will therefore be unable to track TCP connections. %3 %4

**Error Message Explanation:** Network Load Balancing tracks TCP connections to ensure affinity to the host servicing the connection. If NLB cannot open a communication channel with the TCP/IP stack, it cannot guarantee, under certain conditions, that a given TCP connection is serviced by the same host for its duration.

**User Action:** Try restarting the host. If restarting fails to establish communication with the TCP/IP stack, advanced users can force NLB to track TCP connections itself by setting the following REG_DWORD value to 0 (zero):
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\WLBS\Parameters\Global\EnableTCPNotification

**Note**

The key will not exist by default. You must create it.

You must restart the computer after making this change. This latter method of tracking TCP connections is less robust than using the TCP/IP stack, but it ensures TCP connection affinity in most cases. This method is not recommended if you are using multiple network adapters with asymmetric traffic flow (that is, using different network adapters to handle outbound vs. inbound NLB traffic).

For more information about Windows Registry best practices, see Help and Support.

**Event ID: 82**

**Alert Severity:** E

**Event Text:** *NLB Cluster %2 %1:* The network adapter to which NLB is bound does not support NDIS packet indications, which is commonly caused by outdated network adapter hardware and/or drivers. In order to properly load-balance incoming load, NLB requires that adapters support NDIS packet indications. %3 %4

**Error Message Explanation:** Your computer is using a network adapter that is either older or non-server-class. Network Load Balancing (NLB) requires that network adapters support Network Driver Interface Specification (NDIS) packet indications. NLB no longer accepts the older NDIS Packet Receive indications from network adapters. Incoming packets that arrive via this older mechanism will be dropped by NLB.

**User Action:** Verify that the drivers for the network adapter are up-to-date and that your network adapter is listed in the Windows Catalog as a supported adapter for the Windows Server 2003 family of products. If necessary, install a different (newer) network adapter that supports NDIS packet receives. For more information, see the Windows Catalog at http://go.microsoft.com/fwlink/?LinkID=37604.

**Event ID: 83**

**Alert Severity:** E

**Event Text:** *NLB Cluster %2 %1:* Duplicate dedicated IP address %3 was discovered on the network. Please check the NLB configuration on all hosts that belong to the cluster and make sure that they all contain unique dedicated IP addresses. %4

**Error Message Explanation:** Another host in the cluster has an identical dedicated IP address.

**User Action:** Check the Network Load Balancing configuration on all hosts that are members of the cluster, and verify that they are all configured with unique dedicated IP addresses. For more information about Network Load Balancing, see Help and Support.

**Event ID: 84**

**Alert Severity:** W

**Event Text:** *NLB Cluster %2 %1:* NLB was unable to open the NLB public connection callback object and will therefore be unable to track connections. %3 %4

**Error Message Explanation:** Network Load Balancing (NLB) tracks TCP connections to ensure affinity to the host servicing the connection. In a Microsoft Internet Security and Acceleration Server (ISA Server)/NLB deployment, if NLB cannot open a communication channel with the firewall engine, it cannot guarantee, under certain conditions, that a given TCP connection is serviced by the same host for its duration.

**User Action:** This communication medium with the firewall engine is required for ISA Server/NLB deployments. Check the ISA Server installation to verify that it is correctly configured, and then try restarting the host.

**Event ID: 4096**

**Alert Severity:** W

**Event Text:** Configuration completed. Microsoft Cluster service is also installed on this machine, and this may prevent the proper operation of both.

**Error Message Explanation:** This computer is a member of a server cluster. Network Load Balancing can interfere with a server cluster's use of network adapters: therefore, Microsoft does not support this configuration.

**User Action:** Either do not install NLB on this computer, or remove this computer from the server cluster.

**Event ID: 12288**

**Alert Severity:** I

**Event Text:** Configuration update %2 started by %1 on adapter %3

**Error Message Explanation:** This event is logged to inform you that a user action was initiated using Network Load Balancing Manager.

**User Action:** No user action is required.

**Event ID: 12289**

**Alert Severity:** I

**Event Text:** Configuration update %2 completed with status %1 on adapter %3. Information log contains: %n%4

**Error Message Explanation:** This event is logged to inform you that a user action was initiated using Network Load Balancing Manager.

**User Action:** No user action is required.

# Implementing Common Desktop Management Scenarios with the Group Policy Management Console

Updated: May 3, 2004

Applies To: Windows Server 2003 with SP1

This white paper describes six scenarios for using Group Policy to manage users' desktop environments. This document is a companion to a set of preconfigured Group Policy object (GPO) scenarios, available for download from this page. These GPO scenarios are intended to be starting points from which you can develop settings tailored to your environment. These scenarios illustrate of the range of features manageable by using Group Policy. This white paper includes coverage of Windows Server 2003 and the Group Policy Management Console (GPMC).

**In This White Paper**

Introduction (Implementing Common Desktop Management Scenarios with the Group Policy Management Console)

Understanding and Using the Scenarios

Deploying the Scenarios

Configuring Specific Features

Switching Between Scenarios

Extending the Scenarios

Appendix A: GPO Scenario Policy Settings

Appendix B: Running CommonScenarios.msi

# Introduction (Implementing Common Desktop Management Scenarios with the Group Policy Management Console)

Updated: March 2, 2005

Applies To: Windows Server 2003 with SP1

**Group Policy Overview**

The IntelliMirror® management technologies provide Active Directory-based change and configuration management of user and computer settings on computers running a member of the Microsoft® Windows® Server 2003 or Microsoft Windows® 2000 families of operating systems, or the Microsoft® Windows® XP Professional operating system. Group Policy provides the infrastructure for IntelliMirror by allowing you to specify settings for registry-based policy, security, software installation, scripts, folder redirection, Remote Installation Services, and Internet Explorer maintenance.

The Group Policy settings that you create are contained in a Group Policy object (GPO). By linking a GPO to selected Active Directory Service containers – sites, domains, and organizational units (OUs) – you can apply these settings to the users and computers in those Active Directory containers. Group Policy inheritance and precedence determine where and how you link GPOs. By default, options set in GPOs linked to higher levels of Active Directory containers – sites, domains and OUs – are inherited by all containers at lower levels, though inheritance does not occur across domains. Because lower-level GPOs apply last, they override higher-level GPOs and can provide lower-level OUs with a different set of Group Policy settings.

To manage GPOs, you use the Group Policy Management Console (GPMC) or its scripting interfaces. GPMC is a new tool, released at the launch of Windows Server 2003. It is important to note, however, that GPMC is a very effective tool for managing Group Policy in Windows 2000 domains and, with the exception of a few new features, does not require Windows Server 2003 to be running in your environment.

It is assumed that the reader of this white paper has an understanding of basic Group Policy principles.

**About the Common Desktop Scenarios**

Group Policy is a rich and flexible technology providing the capability to efficiently manage a large number of computer and user accounts through a centralized, "one-to-many" model. With this flexibility comes the potential for complexity. For example, Group Policy in Windows Server 2003 exposes almost 1,000 configurable settings. This can initially appear a daunting task – how does the administrator assess the relative importance of these settings, and what features enabled by Group Policy might be considered when deploying a solution?

This Common Desktop Scenarios white paper provides a structured, tested, and consistent set of pre-packaged GPOs and associated documentation, with the goal of lowering the "barrier of entry" for those assessing Group Policy.

The scenarios described in this white paper provide a good starting point from which you can begin evaluating and understanding Group Policy. By implementing these scenarios in a test environment, you should be able to:

- Quickly understand the scope of Group Policy and consider how to use a wide number of settings.
- List some of the important solutions enabled by Group Policy.
- Gain familiarity with some of the new functionality introduced with GPMC, particularly the backup/import of GPOs and Group Policy reports.
- Reach conclusions about how to implement Group Policy in your production environment.

| **Note** |
| --- |
| The approach taken in this white paper has significant differences to previously released versions. These differences primarily reflect the new features available in GPMC and are described fully later in this document. |

**How to Use the Common Desktop Scenarios**

The GPOs provided with this white paper were created using the Backup capability of GPMC. This new tool provides a single entry point for management of Group Policy in your environment and can manage both Windows 2000 and Windows Server 2003 domains. You can download the Group Policy Management Console (GPMC) from the Microsoft Web site

(http://go.microsoft.com/fwlink/?linkid=12946). You can import these GPOs into your environment as a first step toward implementing the common scenarios. The details of each scenario are described in this document and further documented in a spreadsheet, as well as in HTML-based reports generated using the Group Policy Reports capability of GPMC for each of the GPOs.

In many cases, a scenario might deliver a specific computer/user configuration that is close to the required environment for your production environment and might not need significant changes. In other cases, it might be necessary to substantially modify the GPOs provided to ensure alignment with your business goals.

The GPOs can be implemented and validated in a test environment, modified (where necessary) to map to your specific needs, and – after appropriate testing – copied or imported into a production environment.

The mere act of importing the scenario GPOs into your domain has no immediate impact on computers or users. To affect target accounts, the GPOs must be linked to an appropriate Scope of Management (SOM) – a site, domain, or OU which you want to manage using these GPOs. This means that the reader with limited test resources might choose to import the GPOs into a production environment and – before applying to a set of computer or user accounts – link to a SOM that has a more limited number of accounts, effectively forming a controlled pilot program. After completing the pilot program and testing any required changes, the GPOs can be linked to broader SOMs. This approach is most effective when computers and users are targeted through OUs, and the majority of this white paper assumes that this is the targeting mechanism used (unless noted otherwise).

| **Note** |
| --- |
| You can download the CommonScenarios.msi installation package, which includes this document and supporting files, from the Microsoft Download Center. |

**Overview of the Scenarios**
The following is a list of the scenarios along with typical examples.

**Lightly Managed**
Use this scenario for power users or developers who require considerable control over their computer. You can also use this scenario in an organization where tightly managed desktops are not acceptable to users or where desktop management is highly delegated. Along with the other scenarios, the Lightly Managed scenario supports increased security and promotes consistency of user experience, both of which can be beneficial even where a tightly managed desktop is not appropriate.

The Lightly Managed scenario has the following characteristics:

- Is the least managed of all of the scenarios.
- Allows users to customize most settings that affect them but prevents them from making harmful system changes.
- Includes settings that reduce help desk costs and user downtime.
- Supports *free-seating*, which means users can sit down at any computer and access all their resources, applications, and data as if they were sitting at their own computer. This also simplifies your file-backup scenarios, because users' files are all stored on designated file servers.
- Typically has a core set of applications assigned to either the user or the computer, which are always available. Users can also install applications that have been published for them, which they can choose to install.

**Mobile**
The Mobile scenario is relevant to mobile/laptop computers and their users. This scenario pays particular attention to the disconnected user who frequently needs to work offline and occasionally resynchronize with the corporate network.

The Mobile scenario has the following characteristics:

- Can be used by users who are away from the office most of the time, who log on using low-speed, dial-up links, but who also occasionally log on using high-speed network links.
- Can also be used by users who are away from the office only occasionally and who log on by using remote access or remote network links.

- Allows users continuous access to their data and configuration settings whether the computer is connected to or disconnected from the network.

- Partially supports free-seating (can optionally support full free-seating) to facilitate centralized data backup and to enable users to access important data and settings from additional computers.

- Allows users to disconnect from the network without logging off or shutting down.

**Multi-User**
Use this scenario in a university computer laboratory or library where users can save some customizations, such as desktop wallpaper and color scheme preferences, but are not allowed to change hardware or connection settings.

The Multi-User scenario has the following characteristics:

- Allows basic customization of the desktop environment. Users can save desktop configurations, but they cannot customize network, hardware, and system settings.

- Supports free-seating; users can log onto any computer and get their data and settings. No cached state is maintained on the computer when they leave.

- Users have restricted write access to the local computer and can only write data to their user profile and to redirected folders.

- Has a set of applications that are always available (assigned), as well as applications that can be installed and removed as necessary (published).

- Is highly secure.

**AppStation**
The AppStation scenario is used when you require highly restricted configurations with only a few applications. Use this scenario in vertical applications such as marketing, claims and loan processing, and customer-service scenarios.

The AppStation scenario has the following characteristics:

- Allows minimal customization by the user.

- Allows users to access a small number of applications appropriate to their job role.

- Does not allow users to add or remove applications.

- Supports free-seating.

- Provides a simplified desktop and Start menu.

- Users have restricted write access to the local computer and can only write data to their user profile and to redirected folders.

- Is highly secure.

**TaskStation**
Use the TaskStation scenario when you need the computer dedicated to running a single application, such as on a manufacturing floor, as an entry terminal for orders, or in a call center.

The TaskStation scenario is similar to the AppStation scenario, with the following changes:

- It has only one application installed, which automatically starts when the user logs on.

- No desktop or Start menu is present.

**Kiosk**
Use this scenario in a public area, such as in an airport where passengers check in and view their flight information. Because the computer is normally unattended, it needs to be highly secure.

The Kiosk scenario has the following characteristics:

- Is a public workstation.
- Runs only one application.
- Uses only one user account and automatically logs on. The system automatically resets to a default state at the start of each session.
- Runs unattended.
- Is highly secure.
- Is simple to operate, with no logon procedure.
- Does not allow users to make changes to the default user or system settings.
- Does not save data to the disk.
- Is always on (the user cannot log off or shut down the computer).

A workstation that uses the Kiosk scenario is similar to a TaskStation, but users are anonymous in that they all share a single user account that automatically logs on at computer startup. This is achieved by modifying the Kiosk machine in a manner described later in this document. Users cannot customize their environment and user states are not preserved.

Although user sessions are usually anonymous, a user can log on to an application-specific account, such as to a Web-based application through Internet Explorer (assuming Internet Explorer is the "kiosk application" launched at startup).

The dedicated application could be a Line of Business (LOB) application, an application hosted in Internet Explorer, or another application, such as one available in Microsoft Office. The default application should not be Windows Explorer or any other shell-like application. Windows Explorer allows more access to the computer than is appropriate for a Kiosk computer. Be sure the command prompt is disabled and Windows Explorer cannot be accessed from any application you use for this purpose.

Applications used for kiosk scenarios should be carefully checked to ensure they do not contain "back doors" that allow users to circumvent system policies. For example, they should not allow users access to applications that access the file system. Ideally, you should only use applications that comply with The Application Specification for Windows 2000, are Certified for Windows, and that check for Group Policy settings before giving users access to prohibited features. For more information, see the "The Application Specification for Windows 2000" at the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=1636). Older applications will not normally be Group Policy-aware, so try to disable any features that allow users to bypass administrative policy.

The registry entries **Run** and **RunOnce** are disabled in the Kiosk scenario through associated policy settings.

| **Important** |
| --- |
| Applications that use the **RunOnce** entry to finish an installation or upgrade will fail when the **Do not process the Run Once list** Group Policy setting is enabled. |

Please see the "Scenario Comparison" table in Appendix A for a comparison of features across scenarios.

# Understanding and Using the Scenarios

Updated: May 3, 2004

Applies To: Windows Server 2003 with SP1

The scenarios are implemented through the use of GPOs provided in the .MSI package with this white paper. By linking appropriate combinations of GPOs to OUs for computer and user settings, you can be implement each of the scenarios in your environment. This section describes the general characteristics of the GPOs and how they can be linked to OUs.

#### How the Scenarios Are Designed

### Separate GPOs for Computer and User Policy Settings

Most scenarios are directly associated with two GPOs – one for computer configuration and another for user configuration. For example, the **Lightly Managed** scenario is directly associated with the **Lightly Managed (Machine)** GPO for computer configuration and the **Lightly Managed (User)** GPO for user configuration. This approach simplifies troubleshooting, makes the application of GPOs somewhat more intuitive, and is a Group Policy best practice.

### "Base GPOs" for Common Settings

Many of the scenarios share a considerable number of policy settings and associated values. This version of the Common Scenarios white paper introduces the concept of the "base GPO," which is more relevant in a real-world, production environment than the monolithic GPOs packaged with previous versions.

The base GPOs (two each for **Lightly Managed** and **Highly Managed**) contain Group Policy settings common to other "child" scenarios (see the next section). The GPOs for these child scenarios extend these base GPOs by enabling, disabling, or changing the values of a relatively small number of settings tailored to their specific requirements.

### Scenario/GPO Relationships

To implement any of the scenarios, the GPOs must be linked to Scope of Management (SOM) containers - a site, domain or OU. Most of the GPOs packaged with the white paper are designed to be linked to OUs, rather than sites or domains.

| **Note** |
| --- |
| The GPOs provided are intended to assemble common policy settings and their relationships are technically independent of any OU hierarchy you might implement. For scenarios that require more than one GPO you might either link separate GPOs in a chain of OUs (Group Policy inheritance) or link multiple GPOs to a single OU (Group Policy precedence) – these options are covered in more detail later in this white paper. |

The following list illustrates the relationships between GPOs.

#### Machine Policy GPOs

### Lightly Managed

Mobile – No differences from Lightly Managed GPO, so no machine GPO exists for Mobile

### Highly Managed

AppStation - No differences from Highly Managed GPO, so no machine GPO exists for AppStation

Multi-User – GPO provided

TaskStation – GPO provided

Kiosk – GPO provided

#### User Policy GPOs

### Lightly Managed

Mobile – GPO provided

### Highly Managed

AppStation - GPO provided

Multi-User – GPO provided

TaskStation – GPO provided

Kiosk – GPO provided

For example, the AppStation scenario shares many of its Group Policy settings with the Highly Managed GPOs, as follows:

- On the machine side of policy for the AppStation scenario, the settings are identical in all regards to that of the Highly Managed scenario and, as such, there is no AppStation-specific GPO for machine policy settings. To implement machine settings appropriate to the AppStation scenario, it is only necessary to ensure that the Highly Managed (Machine) GPO is linked to a SOM containing the target machines.

- By comparison, some differences exist for user policy settings between Highly Managed and AppStation. For this reason, a GPO is provided for the user policy settings in the AppStation scenario, which contains differences from the Highly Managed user GPO. To implement user settings for the AppStation scenario, it is necessary that both the Highly Managed (User) and AppStation (User) GPOs are linked – directly or otherwise - to a SOM containing the target users. This can be achieved through either Group Policy inheritance or precedence when GPOs are linked to the same container.

The AppStation scenario is an exception - most scenarios are implemented through two GPOs – one for computer configuration and another for user configuration (the AppStation scenario uses only user configuration). In addition, most are associated with the base scenarios, **Lightly Managed** and **Highly Managed**, through Group Policy inheritance and/or precedence. With these issues in mind, the GPOs listed in Table 1 are effective for each scenario.

**Table 1. Scenario GPOs**

| Scenario Name | Base GPO Name | Scenario-Specific GPO Name |
|---|---|---|
| **Lightly Managed** | | |
| **Computer** | Lightly Managed (Machine) | N/A (no changes from base) |
| **User** | Lightly Managed (User) | N/A (no changes from base) |
| **Mobile** | | |
| **Computer** | Lightly Managed (Machine) | N/A (no changes from base) |
| **User** | Lightly Managed (User) | Mobile (User) |
| **AppStation** | | |
| **Computer** | Highly Managed (Machine) | N/A (no changes from base) |
| **User** | Highly Managed (User) | AppStation (User) |
| **Multi-User** | | |
| **Computer** | Highly Managed (Machine) | Multi-User (Machine) |
| **User** | Highly Managed (User) | Multi-User (User) |
| **TaskStation** | | |
| **Computer** | Highly Managed (Machine) | TaskStation (Machine) |
| **User** | Highly Managed (User) | TaskStation (User) |
| **Kiosk** | | |
| **Computer** | Highly Managed (Machine) | Kiosk (Machine) |
| **User** | Highly Managed (User) | Kiosk (User) |

**Purpose of the Highly Managed GPOs**

The Highly Managed scenario is a virtual scenario – its purpose is to provide a set of settings common to other scenarios

(specifically, the Multi-User, AppStation, TaskStation, and Kiosk scenarios). As such, the Highly Managed GPOs are not intended (or tested) to operate in their own context. Instead, the other scenarios are enabled by ensuring that the Highly Managed GPOs affect accounts through Group Policy inheritance or precedence. In short, no computer or user account should exist in OUs designed specifically for a Highly Managed scenario.

**Creating a Test Environment**

In previous versions of this white paper, a considerable number of scripts were necessary to create and configure the GPOs provided with the white paper. With the advent of GPMC – and specifically its Import functionality – these scripts are no longer necessary. The GPOs have been backed up and provided with this white paper, which allows you to import them directly into your environment.

To assist in creation of this environment, a script included with GPMC, CreateEnvirommentFromXML.wsf, is used to create a hierarchy of OUs, create GPOs, and link the GPOs to the OUs appropriately. A script called CreateCommonScenarios.cmd is provided with this white paper to call this script with the appropriate parameters.

After successfully running the script, you will have OUs and linked GPOs configured similar to those shown in Figure 1.




**Figure 1. GPMC Screenshot of Scenario GPOs**

**Linking GPOs to OUs to Create Scenarios**

There are two broad approaches available for using the GPOs to create scenarios in your environment. To illustrate these options, the AppStation scenario is used, which is an extension of the Highly Managed scenario (many of the AppStation settings are defined by the Highly Managed GPOs).

**Linking GPOs to a Hierarchy of OUs (Group Policy Inheritance)**

With this option, create an OU hierarchy where the AppStation OUs (one each for Computer and User accounts) are children to the Highly Managed OUs. Figure 2 illustrates how the GPOs would be linked.




**Figure 2. Linking Scenario GPOs Using Inheritance**

Note that because the machine policy for the AppStation scenario does not differ from that of the Highly Managed scenario, no AppStation (Machine) GPO exists. Technically, it is not necessary to place AppStation affected computers in the AppStation OU

because they might reside in the Highly Managed OU and be affected by the same settings. However, for clarity and to easily accommodate the potential for future differences between Highly Managed and AppStation GPOs (on the machine side), there might be value in retaining the AppStation OU. After the appropriate GPOs are linked to these OUs, the computer and user accounts would be moved into the appropriate AppStation OUs.

**Linking GPOs Directly to OUs (Group Policy Precedence)**
In this approach, you only create two OUs (one each for Computer and User accounts). For each, both the Highly Managed and AppStation GPOs are linked directly to the OU. This is illustrated by Figure 3.




**Figure 3. Linking Scenario GPOs Using Precedence**

After linking the GPOs, it is important to ensure that the precedence of the GPOs is such that the AppStation GPOs take priority over Highly Managed GPOs. GPMC provides a simple and intuitive interface for enforcing this precedence. This is achieved by selecting a container, such as an OU, in the GPMC tree view and using the **Linked Group Policy Objects** tab in the right pane to adjust ordering of GPOs linked to that container.

After creating the OUs and GPO links, move the computer and user accounts into the appropriate AppStation OUs.

**Differences with Earlier Versions of the Common Scenarios**
An earlier version of this white paper focused on Windows 2000 domains. Since that time, a number of Group Policy enhancements have become available, most notably GPMC. The following areas differ between the Windows 2000 and Windows Server 2003 common scenarios.

**Base GPOs**
As described earlier, the base GPOs (Lightly Managed and Highly Managed) are a new concept in this version of the white paper.

**GPMC-Based GPO Deployment**
GPMC provides a significant number of features – Backup, Import, Copy and so on – that significantly simplify managing GPOs. This white paper takes full advantage of these new capabilities by providing backed-up copies of GPOs which can easily be imported into your test environment.

**Environment Creation Script (CreateCommonScenarios.cmd)**
GPMC includes a series of example scripts which illustrate the automation of common Group Policy operations. One such script is CreateEnvironmentFromXML.wsf which uses an XML file to recreate an environment (typically OUs, GPOs, and GPO links) in your domain. An XML file (CommonScenarios.xml) is included with this white paper to facilitate the use of the script to largely automate the creation of your environment.

A command file – CreateCommonScenarios.cmd – is packaged with this white paper to streamline the installation of a sample environment.

It is important to note that while the script must be run from either a Windows XP or Windows Server 2003 computer, the domain in which it creates the sample environment can be either a Windows 2000 or Windows Server 2003 domain.

**Group Policy Reports Documentation**
GPMC provides a new way to report the contents of GPOs – Group Policy Reports. These are HTML reports that document all aspects of each GPO. Reports for each of the GPOs are provided with this CommonScenarios.msi as a documentation tool. When you make changes to the GPOs, GPMC can help you document the changes by using Group Policy Reports.

**New Group Policy Settings**
Windows Server 2003 delivers a number of new settings not just for that operating system but also for Windows XP Professional. Many of these settings are used in the updated scenario GPOs.

**Consolidated Spreadsheet of Group Policy Settings**

The Windows 2000 and Windows XP versions of the scenarios installed one spreadsheet for each scenario. This revised version includes a consolidated spreadsheet – CommonScenarios.xls - and each scenario is listed in a separate column. Using Microsoft Excel column filtering, this can be a helpful way to quickly compare GPO settings.

# Deploying the Scenarios

Updated: March 2, 2005

Applies To: Windows Server 2003 with SP1

This section describes the steps necessary to implement the scenarios. As a prerequisite, it is assumed that you have created a fully operational Active Directory domain and have validated the integrity and operation of the underlying DNS infrastructure. Group Policy is dependent on a well-configured, reliable Active Directory environment.

| **Note** |
| --- |
| You can download the CommonScenarios.msi installation package, which includes this document and supporting files, from the Microsoft Download Center. |

The following steps are necessary to implement the scenarios:

1. Run the CommonScenarios.msi package to copy the GPOs, scripts, and associated documentation (such as this white paper) to your administrative workstation.

2. Create an appropriate OU environment.

3. Use GPMC to import the scenario GPOs into your environment.

4. Link the GPOs to OUs.

   Two options are available to cover steps 2 - 4:

   a. Automatically create an OU hierarchy, GPOs, and GPO links using a script provided with this white paper (CreateCommonScenarios.cmd).

   b. Manually implement each step (create OUs, import GPOs, and so on). This is a more flexible approach but will take a little longer to implement.

On completion of these steps – and regardless of the approach taken (scripted or manual) - you will perform the following steps to complete configuration and testing of the scenarios:

- Configure specific scenario features (see the "Configuring Specific Features" section)

- Create Computer and User accounts for use within the scenarios.

- Test scenarios.

**Test and Production Environment Considerations**

Before installing the scenarios into your environment (described in the next section), you should understand the options available for testing. There are two primary ways in which you can incorporate the GPOs into your environment for testing; by linking the GPOs to test OUs in your production domain or by using a separate test domain (either in the same or a different forest).

**Using Test OUs in Your Production Environment**

Where a separate test domain (the preferred approach) is not feasible, the scenario GPOs can be imported to a production environment and linked to OUs created in that domain specifically for test purposes. These OUs should be well-segmented from those used for regular production purposes. The benefit of this approach is that you need not create infrastructure to support a separate domain (domain controllers, networking, and so on). The disadvantage is that misconfiguration of accounts, GPOs, or GPO links can directly affect the production environment.

**Using a Separate Test Domain**

With this approach, you use a domain that is separate from that of your production environment (trusts may or may not exist

between the test and production environments). This can allow you to run a more realistic test by modeling the test domain on your existing domain, so you can test domain-wide settings and other interactions as needed. A significant benefit of using this approach is that the impact of errors through inappropriately linked or filtered GPOs is significantly less than is evident in a production environment. The primary drawback of this approach is that you need to set up domain controllers and other supporting infrastructure (networking, DNS, and so on) that implements the test domain. If possible, use this approach.

**Considerations Across Forests and Domains**
The two domains used for test and production purposes might or might not be in the same forest. Further, if they are in separate forests, a forest trust might or might not exist, because a forest trust can only be established between Windows Server 2003 forests.

Where a trust exists between the domains (for example, where the test and production domains are in the same forest and the default transitive trusts exist), an administrator with the appropriate rights in each domain can use GPMC to copy GPOs between domains. This can be as simple as dragging and dropping GPOs between domains using the GPMC graphical user interface. Note that if the GPOs in your test environment include security principals (for example, group names) or UNC paths (for example, when specifying your Folder Redirection parameters), migration tables will help you handle any changes you might need to make across domain boundaries. The Migration Table Editor – part of GPMC – provides a simple interface for editing migration tables. For more information, see the "Migrating GPOs Across Domains with GPMC" white paper at the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=14321).

Where a trust does not exist, the follow options are available:

- Create your OU structure and use the GPMC **Import** function to create new GPOs in the production environment. The backup from which you import GPOs might be that provided with this white paper or – if you have already tailored the GPOs to your needs – from a backup set that you created.

- Use the CreateEnvironmentFromXML.wsf script (a sample script installed with GPMC – see GPMC Help for more information) to create the default OU environment in your production domain and, if necessary, rename OUs and GPOs to better reflect naming/structural conventions in your production environment.

- Use the Stored User Names and Password feature of Windows XP to store credentials for a user in a non-trusted domain.

**Test Environment Recommendation**
Where adequate resources are available (domain controllers, networking infrastructure, and so on), it is strongly recommended that the GPOs are imported into a test domain. Using a separate domain provides you with a relatively safe environment in which to assess the GPOs and also allows more flexibility in regards to testing domain-wide Group Policy settings.

**Installing the Common Scenario Scripts and GPOs**
Packaged with this white paper is a series of scripts and GPO backups that form the basis of your deployment of the common scenarios. The CommonScenarios.msi file installs these components – including the white paper itself – into the "%programfiles%\Microsoft\Group Policy Common Scenarios" folder and creates menu options on the Start Menu for some of these components (for example, a link to the white paper and its associated spreadsheet). After running this package, the following directory structure is created:

**%programfiles%\Microsoft\Group Policy Common Scenarios**

**\Documentation (Common Scenarios white paper and spreadsheet)**

**\GPO-Backups (GPMC-originated backups of the scenario GPOs)**

**\GPO-Reports (GPMC-originated reports for each GPO in HTML format)**

**\Environment (XML representation of example OU structure for scenarios)**

**\Scripts (scripts used to create example environment)**

It is assumed that the computer on which this directory is created already has GPMC installed (the supplied scripts include calls to GPMC scripting interfaces). GPMC requires either Windows XP Professional or Windows Server 2003 (see GPMC documentation for further details about requirements). The directory in which the common scenarios files are installed is hereafter referred to as the <**installdir**>, (typically %programfiles%\Microsoft\Group Policy Common Scenarios).

In addition to installing the files, a **Group Policy Common Scenarios** menu on the user's **Start/Programs** menu is created. This menu provides access to the white paper, the spreadsheet, a shortcut to the Group Policy Reports directory, and a

command line configured to start in the %programfiles%\Microsoft\Group Policy Common Scenarios\Scripts directory.

After you run the installer, you must import the GPOs associated with the scenarios into your domain. There are two options to achieve this: a quick method to create an overall common scenarios test environment (a sample OU structure, GPOs, and GPO links), and a more manual process where creation of OUs, GPOs and GPO links are separate steps.

**Deployment Option 1: Quick Setup Using CreateCommonScenarios.cmd**
A script is included to assist with installing a representative OU structure, GPOs, and GPO links into your environment. This script is named **CreateCommonScenarios.cmd** and is installed into the <**installdir**>**\Scripts** directory. The script calls a sample script installed by GPMC –(CreateEnvironmentFromXML.wsf) which is installed in the **%programfiles%\GPMC\Scripts** directory.

CreateCommonScenarios.cmd initiates the following tasks:

- Creates an example OU hierarchy, representing locations for computer and user accounts
- Creates the scenario GPOs, including default permissions
- Links the scenario GPOs to the appropriate OUs

**Running CreateCommonScenarios.cmd**
To run the script, do the following:

1. Click **Start**, click **Programs**, click **Group Policy Common Scenarios**, and then click **Common Scenarios Scripts Command Prompt**.
2. From the command prompt, type **CreateCommonScenarios**, and then press **Enter**.

**Results of Running the Script**
By running the script against your domain, you create a self-contained test environment with the appropriate OUs and GPOs linked to those OUs. Note that might still need to create or edit GPOs linked at the domain level, especially where account policy settings are modified but this is specific to your domain and outside the scope of the scenario GPOs. To avoid potential conflict with your existing environment, no domain-level policy settings are implemented through the scenarios.

**Modifying the Behavior of CreateCommonScenarios.cmd**
The CreateCommonScenarios.cmd script is a wrapper for the CreateEnvironmentFromXML.wsf script, which is installed with GPMC. In some cases you might need to change the default behavior of this script. For example, if you want to modify the domain controller against which the script is run. Any parameters passed directly to the CreateCommonScenarios.cmd script are passed on to the CreateEnvironmentFromXML.wsf script. For more information about the parameters supported by the CreateEnvironmentFromXML.wsf script, see GPMC Help.

**Deployment Option 2: Manual Deployment Steps Using the GPMC GUI**
This method is more time-consuming but allows you greater control over your OU design and GPO links.

**Create an Appropriate OU Test Environment**
You need to create OUs for the users and computers you want to manage within your domain. If your environment requires a more complex OU structure than that described in this document, refer to the "Designing the Active Directory Logical Structure" chapter of the *Designing and Deploying Directory and Security Services* book in the *Windows Server 2003 Deployment Kit*. For more information, see the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=18341).

**Use GPMC to Import GPOs**
The GPOs provided with this white paper can also be imported into your environment manually, using a GPMC sample script, as follows:

```
cscript %programfiles%\gpmc\scripts\ImportAllGPOs.wsf <installdir>\GPO-Backups
```

**Note**

"Cscript" can be omitted if you have previously configured cscript as the default scripting environment for your machine. More specifically, if you have previously run:

**cscript //hh:cscript**

this will set cscript as the default environment, in which case the following command would work:

**%programfiles%\gpmc\scripts\ImportAllGPOs.wsf <installdir>\GPO-Backups**

The ImportAllGPOs.wsf script will result in all the GPOs provided in the <installdir>GPO-Backups directory being created in your environment and their settings imported. On completion of this step you will then need to link the GPOs to the OUs you have created.

The same result can be achieved through the GPMC MMC-snap-in, do the following:

1. Create empty GPOs (no policy settings) in your environment by right-clicking the Group Policy object node and selecting **New** from the context menu. Create one GPO for each of the Common Scenario GPOs, using the names documented in the "How the Scenarios Are Designed" section of this white paper.

2. In GPMC, right-click each GPO and then click **Import**. Go to the <installdir>\GPO-Backups directory and import policy settings from the appropriate backed-up GPO.

**Link GPOs to OUs**

To apply the settings of a GPO to the users and computers of a domain, site, or OU, you need to add a link to that GPO. You can add one or more GPO links to each domain, site, or OU by using GPMC. Keep in mind that creating and linking GPOs is a sensitive privilege that should be delegated only to administrators who are trusted and understand Group Policy.

To link an existing GPO

1. In GPMC, right-click a domain or OU, and then click **Link an Existing GPO here**.

2. In the **Select GPO** dialog box, click the GPO which you want to link, and then click **OK**. In the left pane, the GPO is displayed beneath the domain or OU to which it is linked.

You can simultaneously link multiple GPOs to an Active Directory object by holding down the **CTRL** key while selecting GPOs.

**Post-Installation Configuration**

After you have created an environment in which you wish to host your scenario GPOs – either automatically by using the CreateCommonScenarios.cmd script or manually – you must carry out a number of additional steps to finalize the configuration of certain aspects of your environment. This section describes each of these steps.

**Configure Scenario Features**

Configuration for Roaming User Profiles and Redirected Folders require you to enter UNC paths specific to your environment (into the user objects within the Active Directory and to the appropriate GPOs, respectively). See the "Configuring Specific Features" section for more details.

**Create Computer and User Accounts**

Using the Active Directory Users and Computers snap-in, create a sufficient number of user accounts to allow you to test each of the scenarios.

It is important to note that any one scenario is implemented through the combination of user and computer accounts, each of which should be affected by GPOs associated with the same scenario. For example, to achieve the TaskStation scenario, a user in the CommonScenarios/Users/TaskStation OU might log on to a computer in the CommonScenarios/Computer/TaskStation OU. Any "cross matching" of user and computer scenarios is untested and might cause unforeseen results.

**Migrating the Scenarios to a Production Environment**

After testing and potentially customizing the Common Scenario environment, you might wish to move into a production environment. The following approaches can be used:

- **Manual creation in production environment.** With this approach, you create each of the OUs from either GPMC or Active Directory Users and Computers. The GPOs are either copied (if a trust exists between the source and target environments) or created/imported (where file-based GPO backups are used). Then link the GPOs to the OUs and move

the users/computers into the OUs as appropriate.

- **Use the CreateXMLFromEnvironment and CreateEnvironmentFromXML scripts.** The CreateXMLFromEnvironment script – included with GPMC as a sample script – allows the administrator to create an XML file representing all the policy-related objects (OUs, GPOs, and GPO links). Optionally, you can specify a starting OU rather than the entire domain. The script creates a representation of the OU and its linked GPOs in the specified XML file. This file can then be used by the CreateEnvironmentFromXML script to create a mirrored OU structure and associated GPOs and links in the target environment. When planned carefully, these two scripts can help create an initial framework in your production environment.

For more information about managing GPOs across domains, see the "Migrating GPOs Across Domains with GPMC" white paper at the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=14321).

**Removing the Scenarios from Your Environment**

Removing the scenarios from your environment is a relatively simple task. Use the following steps to remove them:

- Move computer and user accounts that you need to retain into alternative OUs, as necessary. Reconfigure user objects, as appropriate, to modify Roaming User Profile paths.

- Validate that each scenario GPO is linked to the expected SOMs (using the **Scope** tab within GPMC). This step is important to ensure that no unexpected cross-domain or other links exist.

- Delete each scenario GPO.

- Validate that the scenario OUs are empty, and then use Active Directory Users and Computers to remove those OUs.

- If you added the domain GPO, validate, unlink, and delete that GPO as appropriate.

If you have used the CreateCommonScenarios.cmd script to create your sample environment, then the CreateEnvironmentFromXML.wft sample script was called implicitly. This script has an **/undo** switch that removes the environment described by the XML file passed as an argument. See GPMC online Help for further details on this script.

In addition to these issues, please refer to the section "Switching Between Scenarios" for more factors that might also be relevant when removing scenarios (for example, "tattooing" the registry with security settings).

# Configuring Specific Features

Updated: May 3, 2004

Applies To: Windows Server 2003 with SP1

This section describes the various features configured in the common scenarios GPOs.

**Roaming User Profiles**
A user profile is a group of settings and files that define the user's environment. A profile includes program items, screen colors, network connections, window sizes and positions, and so on. Roaming User Profiles (RUP) enable the server-based storage of user profiles, which means that users can move between computers and see an identical environment. The RUP is downloaded to the user when s/he logs on and, by default, is stored back on the server when the user logs off. This feature is one component of the concept of "free-seating" – the capability for users to roam between computers yet maintain an identical environment.

**Scenarios in Which Roaming User Profiles are Used**
Due to the order in which logon, profile creation/loading, and the application of GPOs occurs, it is not possible to specify the location of a user profile using Group Policy. For this reason, specifying the user profile location is a distinct and separate step from the application of GPOs. Therefore, to ensure that a user is configured correctly, it is important to move that user into the appropriate OU and, additionally, configure the user object as described below.

The common scenarios leverage Roaming User Profiles as follows:

- **Used:** Lightly Managed, Highly Managed, Multi-User, AppStation, and TaskStation
- **Optionally Used:** Mobile User
- **Not Used:** Kiosk

**Roaming User Profiles Configuration Steps**
For the purposes of illustration, this document assumes a server named **CommonServer** is available to store user profiles. The following steps are necessary to create a share for user profiles and configure user accounts appropriately.

1. On the CommonServer computer, create a folder called **profiles**.
2. Share this folder as **profiles$**.
3. Set share permissions to **Full Control** for the **Everybody** group (security will be enforced by NTFS permissions when the user profile folders are created).
4. Ensure that caching for this folder is disabled (only use caching for folders where you do not store profiles).
5. For each account for which Roaming User Profiles is required, set the profile path in the user object to \\CommonServer\profiles$\%user name% (the Active Directory Users and Computers MMC snap-in is the most common way to edit this parameter). At the time the user profile is created, the %user name% environment variable will be resolved to the name of the user.

**Roaming User Profile Notes**
It is a best practice to allow the user-specific accounts (in the **Profiles** folder) to be automatically created at the time the user profile is first established. This ensures that the appropriate set of NTFS permissions and ownership are set on the folder.

| **Important** |
| --- |
| Roaming User Profiles have their own caching mechanism, which can interfere with Offline Files synchronization and can lead to unexpected behavior and loss of data. Therefore, make sure that caching is disabled for the shares where you store profiles. |

**Redirected Folders**

Folder Redirection allows the contents of a folder in the user's profile to be redirected to a location on the network. For example, you can move My Documents, which is typically part of the user's profile and cached on the local drive, to a folder in the user's home directory on the network. Enabling this feature is useful because the My Documents and Application Data folders often contain large amounts of data. Enabling Folder Redirection helps speed up logon and logoff time because the contents of the redirected folders are not copied along with the rest of the user profile.

In most cases when you use Folder Redirection, you will combine it with Offline Files so that users can access cached copies of the redirected folders when disconnected from the network. In fact, with Windows XP, all redirected folders are also cached locally by default (this is not the case with Windows 2000).

If you have a high volume of users that use a single computer, consider enabling the **At logoff, delete local copy of user's offline files** Group Policy setting to prevent the local hard drive from filling up with cached files from multiple users. This setting is located under Computer Configuration\Administrative Templates\Network\Offline Files.

| **Caution** |
| --- |
| You should use this policy setting with caution. If a user has been working offline and has not synchronized their changes, their changes will be lost at logoff when their offline files are deleted. |

When you implement Folder Redirection, the destination folders are automatically created, and security is configured on these folders automatically. You can change the security of redirected folders by checking or clearing **Grant the user exclusive rights to foldername**,which is located in this path: User Configuration\Windows Settings\Folder Redirection*foldername* - **Settings** tab.

For example, if you redirect the My Documents folder to \\CommonServer\Docs*%User name%*, do not create a **User Name** folder in advance. If the folder exists in advance and the current user is not the owner of the folder and its contents, the redirection process fails. If the folder does not exist prior to folder redirection, however, it is created and the user is made the owner of that folder. If you use the **Create a folder for each user under the root path** option in the **Target Folder Location** dropdown box, Folder Redirection automatically appends %user name% to the root path you specify.

| **Important** |
| --- |
| If you must create the folder in advance, ensure that the user is the owner of the folder and its contents. To change the ownership of a file or folder, use Windows Explorer or the Subinacl.exe utility. For more information, see the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=18341). |

| **Note** |
| --- |
| When Group Policy Folder Redirection settings go out of scope, by default, they leave the redirection in place. The **Settings** tab has an option that allows you to redirect the folders to the local computer; however, avoid using this option. The default setting is the safest because no data is moved when a Folder Redirection policy falls out of scope. Any new Folder Redirection policy then redirects as appropriate. |

#### Scenarios in Which Folder Redirection is Used

The common scenarios leverage Folder Redirection as follows:

- **My Documents and Application Data Redirection:** Lightly Managed, Highly Managed, Mobile, Multi-User, AppStation and TaskStation

- **None:** Kiosk

#### Redirected Folders Configuration Steps

Because a server name is specified when configuring redirected folders, the GPOs provided with this CommonScenarios.msi do not specify Folder Redirection properties. To fully implement the scenarios in your environment, you must carry out the following steps.

1. On the CommonServer computer, create a folder called **redirected**.

2. Share this folder as **redirected$** (the full share name would therefore be *CommonServer**redirected$*).

3. Set share permissions for the *redirected*$ share to **Full Control** for the **Everyone** group (security will be enforced by

NTFS permissions when the user profile folders are created).

4. For both the **Lightly Managed (User)** and **Highly Managed (User)** GPOs, do the following:

    a. In GPMC, right-click the GPO, and then click **Edit**.

    b. Within the Group Policy Object Editor, go to *User Configuration/Windows Settings/Folder Redirection*.

    c. Right-click the **My Documents** node and click **Properties**.

    d. In the **Properties** dialog box, change the **Setting** list to **Basic – Redirect everyone's folder to the same location**.

    e. Leave the **Target Folder Location** list set to **Create a folder for each user under the root path**.

    f. Set the **Root Path** field to \\CommonServer\redirected$, and then click **OK**. Folder Redirection automatically appends %user name% to the path specified.

    g. Right-click the **Desktop** node and click **Properties**.

    h. In the **Properties** dialog box, change the **Setting** dropdown box to **Basic – Redirect everyone's folder to the same location**.

    i. Leave the **Target Folder Location** list set to **Create a folder for each user under the root path**.

    j. Set the **Root Path** field to \\CommonServer\redirected$\, and then click **OK**. Folder Redirection automatically appends %user name% to the path specified.

    k. Close the GPO.

**Redirected Folders Notes**

The steps above ensure that all scenarios – except the Kiosk scenario – implement redirected folders. Because the Kiosk scenario is a child of the Highly Managed scenario, it is necessary to disable the Folder Redirection setting specified in the Highly Managed (User) GPO. This is achieved – through the supplied Kiosk GPOs – by redirecting the appropriate folders back to the local user profile location; this is, in essence, the same as disabling Folder Redirection.

**Internet Explorer Configuration**

There are two types of Internet Explorer Group Policy settings: the standard Windows registry settings and a collection of registry settings and files used to configure Internet Explorer. The standard settings are located in the Group Policy Object Editor under Administrative Templates\Windows Components\Internet Explorer, and the configuration settings are located in User Configuration\Windows Settings\Internet Explorer Maintenance.

The Internet Explorer Maintenance settings can be set in two modes: **policy mode** or **preference mode**. Policy mode is enforced and automatically resets any settings users might change (if they have the appropriate permissions) when Group Policy is applied.

| **Note** |
| --- |
| Like all registry-based Group Policy settings, Internet Explorer Maintenance policy mode settings reapply only when the GPO changes. Unlike most registry-based Group Policy settings, configuring an Internet Explorer Maintenance setting does not prevent the user from changing that setting. |

Preference mode sets initial default values that the user can change, subject to other Group Policy settings that affect the user. The preference settings are only applied again when the administrator makes changes to the preference settings. To allow users to change these preferences, they must be able to access Internet Options on the Tools menu. You allow users access by configuring settings in the Group Policy Object Editor under Computer (or User) Configuration\Administrative Templates\Windows Components\Internet Explorer.

You cannot use policy and preference modes together in a single GPO. If you need both modes, you must use two separate

GPOs. The included scenarios are configured using policy mode, so you might need to create an additional GPO using preference mode for your environment.

**Kiosk Scenario Configuration**
The Kiosk scenario uses a single account that is a member of the Domain Users group (all new accounts in a domain are added to this group) with no special privileges. To enable the use of this account with no user intervention, the AutoLogon feature is used. At startup, the operating system logs on automatically using the account and password specified in the registry key below.

**Kiosk User Account Configuration**
The account used by the Kiosk scenario should be configured as follows:

- Password never expires
- User cannot change password

To configure the account for auto logon, a tool that is part of the Windows XP PowerToys can be used to specify the user name and password. For more information about the PowerToys, see http://go.microsoft.com/fwlink/?LinkId=101706.

**Specifying the Kiosk Application**
The packaged GPOs specify Internet Explorer as the auto launch application for both the TaskStation and Kiosk scenarios. This is handled by enabling the \User Configuration\Administrative Templates\System\Custom User Interface Group Policy setting. The value used for this policy (for both scenarios) is:

%programfiles%\internet explorer\iexplore.exe –k

This has the result of automatically launching Internet Explorer (in full screen mode because of the –k switch) when a user logs on. In effect, the policy is used to replace Windows Explorer as the shell. By updating the Custom User Interface Group Policy setting, you can dictate the application to be launched in the TaskStation and Kiosk scenarios in your own environment.

**Resetting Kiosk Settings to a Default State**
In some cases, you might want to reset the Kiosk environment to a known state. Although Kiosk users are prevented from making most changes to a Kiosk-based computer, a few application-specific settings cannot easily be managed by using Group Policy settings. For example, if a user resizes the Internet Explorer window, the window starts up in that size for all subsequent users.

You can use a mandatory user profile to reset the Kiosk settings. To do this, use the following procedure:

1. Log on to the user account and configure the appropriate settings.
2. Log off of the user account and then log on as an Administrator.
3. Copy the profile to a local or network directory and then rename the profile root to OldName.man.
4. Modify the user object and specify this directory as the profile path.

After you carry out this procedure, if a Kiosk user logs on, the computer settings reflect the settings defined in the mandatory profile.

**Disk Maintenance**
Applications used on a Kiosk-based computer can write data to the disk drive. To prevent the disk drive from filling up, set a disk quota that leaves at least 100 megabytes (MB) free on the system disk.

You should consider using the disk quota in conjunction with a scheduled script that removes older temporary files each night. If the Kiosk application creates a lot of temporary files, the script might also need to run disk defragmentation to maintain system performance.

**Disabling Logoff Capability for Kiosk Users**
Kiosk users cannot log off the Kiosk account by pressing **Ctrl+Alt+Del** or by using the **Start** menu. When the computer starts up, it automatically logs on to the Kiosk account.

**Important**

In the Kiosk scenario, Kiosk users are prevented from logging off; however, in the Kiosk GPOs provided with this white paper, logoff capability is enabled because it makes testing easier. Therefore, before deploying the scenario, you need to disable the logoff feature.

A disadvantage of disabling the logoff feature is that an administrator cannot easily log on to the computer. The administrator cannot restart the computer from the console because shutdown and restart are disabled for the kiosk user. To resolve this, you can use one of the following solutions:

**Solution 1**

1. Restart the computer by doing one of the following:

   - Execute a remote restart using **shutdown.exe**.

     Perform a hardware reset or power off of the computer (this solution is not recommended because it could lead to data loss or system malfunction).

2. When the computer restarts, hold down the **Shift** key to bypass the automatic logon feature and access the standard logon dialog box.

**Solution 2**

In this option, the logon feature remains enabled. The administrator logs off from the Kiosk account and then holds down the Shift key to bypass the automatic logon feature. The disadvantage to this method is that a user could do the first part of this operation. Although they would be unable to log on, this would leave the logon dialog box displayed and render the Kiosk computer unusable until it is restarted.

# Switching Between Scenarios

Updated: May 3, 2004

Applies To: Windows Server 2003 with SP1

Joining a particular scenario is relatively simple – move the appropriate computer and user accounts into OUs to which the appropriate scenario GPOs have been linked. For example, if a user is to experience the Multi-User scenario, his or her account must exist in an OU to which the Multi-User (User) GPO has been linked and must log on to a computer in an OU to which the Multi-User (Computer) GPO has been linked.

In general, moving between scenarios is simply a matter of moving computers and users into the appropriate OUs. However, you might need to do some additional work to complete the transition between scenarios related to security settings and Roaming User Profiles.

Note that moving users between OUs requires a logon to ensure that this is reflected and that the appropriate GPOs are applied. Similarly, moving a computer to a different OU requires that the computer is restarted for the new OU to be recognized.

**Security Settings**
A number of security settings "tattoo" the registry – once applied through a GPO, they remain active even when the GPO is moved out of scope for the target computer or user (by unlinking the GPO, moving the account from an affected OU, or disabling the GPO). Therefore, simply switching a computer or user account to an OU associated with a different scenario will result in some settings from the original scenario remaining. This is particularly evident when transitioning to a scenario with fewer security settings than the original.

There are a number of options to address this issue:

- Ensure that the GPOs associated with the new scenario explicitly disable/enable security settings that have been set in the GPOs associated with the original scenario. For example, if a security setting is configured in a GPO for the original scenario, the GPO for the new scenario must configure that security setting, even if the security setting isn't directly relevant to the new scenario. While this makes for a more deterministic application of these settings (the eventual result of the security setting is independent of the configuration in the original scenario), this does result in all GPOs needing to include the superset of all possible security settings, which can be cumbersome to manage and track.

- Temporarily move the computer and user accounts to a "clean up" OU that explicitly sets all security settings to a default state. After these have been implemented (allowing Group Policy to refresh by waiting at least 120 minutes or by using the **gpupdate** command), the accounts are moved to the appropriate OUs for the new scenario. One drawback with this approach is the potential for a delay to ensure that the temporary policy settings are applied to the computer and user accounts. Additionally, it is necessary for the user account be used to log on at least once to ensure that the "clean up" GPOs associated with user-based security settings are applied.

- As an administrator on the target machine, manually remove the security settings that no longer apply. This is clearly the most labor-intensive approach, requires significant knowledge of how each security setting is implemented, and is generally not an option with a significant number of computers.

**Roaming User Profiles**
The Roaming User Profile attribute for a user (the location of the roaming profile, if set) is stored in the user account object within Active Directory. Due to the order in which profiles are applied, as related to the application of Group Policy, this parameter cannot be set through Group Policy. As such, if a user moves between scenarios where the use of Roaming User Profiles changes (for example, where the initial scenario uses a Roaming User Profile but the target scenario does not), then it is necessary to update the user object to reflect this change.

# Extending the Scenarios

Updated: May 3, 2004

Applies To: Windows Server 2003 with SP1

You can use the common scenarios GPOs as a starting point for your own custom scenarios. For more information about the features and capabilities of Group Policy, see the *Designing a Managed Environment* book in the *Windows Server 2003 Deployment Kit*. For more information, see the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=18341).

**Software Distribution through Group Policy**
The scenarios do not implement any form of software distribution, although this is a feature of Group Policy. By creating .MSI packages and associating them with targeted GPOs, it is possible to manage applications using Group Policy, including support for initial deployment, transforms, and "self-healing" installations (detection and automatic repair of missing components such as DLLs). Group Policy enables you to automatically install and maintain software installation on target computers or make it available for user installation.

**Software Restriction Policies**
Software restriction policies, new with Windows XP and Windows Server 2003, provide a policy-driven mechanism that enables you to identify the programs that are running on computers in your domain, and to control their ability to run. By using software restriction policies, you can:

- Control what programs run on your system. For example, you can apply a rule that does not allow certain file types to run in the mail attachment directory of your e-mail program if you are concerned about users receiving viruses through e-mail.
- Run only digitally signed scripts.
- Allow users to run only specific files on multi-user computers. For example, if you have multiple users who use a single computer, you can set up software restriction policies and Access Control List settings so that users cannot make changes to the computer.
- Decide who can add trusted publishers to a computer.
- Control whether software restriction policies affect all users or only certain users who use a computer.
- Prevent any files from running on a local computer. For example, if you are aware of a known virus, you can disallow a hash of that virus so that the computers in your domain cannot run that program.

**Creating Default OUs for New Machine and User Accounts**
By default, all new computer or user accounts are created in the Computer or User containers, respectively. Because these are not OUs, it is not possible to link GPOs to them. However, using two new tools provided with Windows Server 2003, you can specify that all new accounts will be created in specific OUs. You do this by first creating OUs for new user and computer accounts and then running Redirusr.exe (for user accounts) and/or Redircmp.exe (for computer accounts) once for each domain. From this point, all new user and computer accounts will be placed in the targeted OUs. These tools are included on the Windows Server 2003 CD. You can run either of these tools or both of them.

For more details, see article 324949, "Redirecting the Users and Computers Containers in Windows Server 2003 Domains," in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=4441).

# Appendix A: GPO Scenario Policy Settings

Updated: May 3, 2004

Applies To: Windows Server 2003 with SP1

The Group Policy settings for each scenario (GPOs for both computer and user policy settings) are documented in the accompanying spreadsheet, CommonScenarios.xls. Using Excel's column-filtering capability, the spreadsheet allows you to easily browse through the settings associated with each scenario.

In addition, HTML-based GPO reports are provided in the **<installdir>\GPO-Reports** directory. One report exists for each GPO provided with CommonScenarios.msi and provides a good deal of information about each GPO.

**Scenario Comparison Table**

Table 2 lists the feature characteristics of each scenario.

**Table 2. Scenario Features**

| | **Lightly Managed** | **Mobile** | **Multi-User** | **AppStation** | **TaskStation** | **Kiosk** |
|---|---|---|---|---|---|---|
| Number of users | Multiple | 1 | Multiple | Multiple | Multiple | 1 (anonymous) |
| User profile type | Roaming | Roaming | Roaming | Roaming | Roaming | Local |
| Profile persistence at logoff | Cached | Cached | Removed at logoff | Cached | Removed at logoff | N/A |
| Folder Redirection | My Documents and AppData | My Documents and AppData | My Documents and AppData | My Documents and AppData | My Documents and AppData | No |
| User can customize | Almost all settings | Some or most settings | Some settings | Few settings | None | None |
| Task bar and Start Menu | Yes | Yes | Yes | Yes | No | No |
| Assigned Applications | Multiple | Multiple | Multiple | Few | 1 (usually computer assigned) | 1 (computer assigned) |
| Published applications | Yes | Yes | Yes | No | No | No |
| Security context | User or Power User | User or Power User | User | User | User | User |
| Based on security template | Secure Workstation | Secure Workstation | Highly Secure Workstation | Highly Secure Workstation | Highly Secure Workstation | Highly Secure Workstation |

**Notes:**

- The scenarios are based on the security templates listed; however, in each scenario, the templates have been modified.
- The following significant modifications are made for compatibility reasons:
- Mandatory digital signing of SMB traffic is disabled.
    - Mandatory encryption of secure channel communications is disabled.
    - LAN Manager Authentication Level is not specified.

**Permissions Needed for Folder Redirection**

When setting up folder redirection, it is recommended that you create the root share only on the server, and let the system create the folders for each user. For the best experience, set the share permissions to **Full Control** for the security groups you're redirecting, and set the NTFS permissions to **Full Control** on this folder, subfolders, and files.

If you must create folders for the users, ensure that you set the correct permissions. Tables 3, 4, and 5 below show the default and minimum permissions required for folder redirection.

**Table 3. NTFS Permissions Needed for Root Folder**

| User account | Folder redirection defaults | Minimum permissions needed |
|---|---|---|
| Creator/owner | Full Control, this folder, subfolders, and files | Full Control, this folder, subfolders, and files |
| Local Administrator | Full Control, this folder, subfolders, and files | Full Control, this folder, subfolders, and files |
| Everyone | Full Control, this folder, subfolders, and files | List Folder/Read data, Create Files/Write Data, Create Folders/Append Data - This Folder only |
| Local System | Full Control, this folder, subfolders, and files | Full Control, this folder, subfolders, and files |

**Table 4. Share level (SMB) Permissions Required for Root Folder**

| User account | Folder redirection defaults | Minimum permissions needed |
|---|---|---|
| Everyone | Full Control | Everyone - no permissions.<br><br>Use security group that matches the users who will need to put data on share. |

**Table 5. NTFS Permissions Required for each User's Redirected Folder**

| User account | Folder redirection defaults | Minimum permissions needed |
|---|---|---|
| *%User Name%* | Full Control, owner of folder | Full Control, owner of folder |
| Local System | Full Control | Full Control |
| Everyone | Traverse Folder, Read Attributes, Read Extended Attributes, and Read Permissions | Everyone - no permissions |

# Appendix B: Running CommonScenarios.msi

Updated: March 2, 2005

Applies To: Windows Server 2003 with SP1

This white paper comes with a number of supporting files – spreadsheets, GPO backups, and HTML-based Group Policy reports. These are installed by the installation package CommonScenarios.msi.

| **Note** |
| --- |
| You can download the CommonScenarios.msi installation package, which includes this document and supporting files, from the Microsoft Download Center. |

When you install this package, the files listed in Table 6 are copied into the **%programfiles%\Microsoft\CommonScenarios** directory and a new menu – **Group Policy Common Scenarios** – is created on **Start/Programs**.

**Table 6. Files Installed by CommonScenarios.msi Package**

| Directory | FileName | Description |
| --- | --- | --- |
| \Documentation | CommonScenariosWS2003.doc | This white paper. |
| \Documentation | CommonScenariosWS2003.xls | Spreadsheet listing each setting specified for all scenarios. |
| \Environment | CommonScenarios.xml | File representing sample environment for common scenarios and for use by CreateEnvironmentFromXML script. |
| \GPO-Backups | [GUID] | One subdirectory for each GPO included in the common scenarios. |
| \GPO-Reports | [GPO Name].htm | HTML-based report for each GPO (generated from GPMC). |
| \Scripts | CreateCommonScenarios.cmd | Command file to install a sample OU structure, create the common scenarios GPOs, and link GPOs. |

# Key Archival and Management in Windows Server 2003

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Updated: December 6, 2004

| Tip |
| --- |
| An updated version of this information is on the TechNet Wiki as Active Directory Certificate Services PKI - Key Archival and Management. |

By David B. Cross and Avi Ben-Menahem

Windows Server 2003, Enterprise Edition introduces several new features in the area of Public Key Infrastructure (PKI) technologies and certification authorities (CAs). One area of new functionality is private key archival, recovery, and management. This white paper covers best practices and procedural steps in a key recovery strategy as well as migration procedures for moving from a Microsoft Exchange Key Management Server (KMS) environment to a Windows Server 2003 certification authority.

For information about upcoming enhancements to key archival and recovery, see Key Archival and Recovery in Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkID=92523).

**On This Page**


- Introduction
- Understanding Manual Key Archival
- Understanding Automatic Key Archival
- Understanding User Key Recovery
- Implementing Key Archival Walkthrough
- Migrating Exchange KMS to Windows Server 2003 CA
- Troubleshooting - Key Archival and Management in Windows Server 2003
- Appendix A: Certificate Request Structure
- Appendix B: Additional Information
- Appendix C: Useful Commands


**Introduction**

Windows Server 2003, Enterprise Edition introduces significant advancements in the area of data protection and private key recovery. Windows 2000 introduced the capability for data recovery with the implementation of Encrypting File System (EFS). EFS in both Windows 2000 and Windows Server 2003 supports the use of Data Recovery Agents (DRAs) to decrypt files that have been encrypted by other users. With the expanded functionality of the Windows Server 2003 Certificate Services to offer key archival and recovery services, an enterprise may choose to implement different strategies based on the needs for data protection and data recovery.

Microsoft first offered key archival and recovery features in the Exchange Server 4.0 product line through the KMS component of the Exchange Server. The KMS can act as a Registration Authority (RA) to Microsoft Certification Services to provide user registration, key archival, key recovery, and certificate publishing capabilities to an Exchange Server e-mail and collaboration system. The KMS allows an administrator to recover the lost or corrupted encryption private keys of a Microsoft Outlook® user and generate a new signing key. The Outlook client as well as many other Secure/Multipurpose Internet Mail Extensions (S/MIME)–enabled mail clients support the use of separate signing and encryption key pairs. This method helps to isolate an

organization from non-repudiation issues. For more information about Key Management Server and Windows PKI interoperability, see Appendix B: Additional Information.

The Windows Server 2003 key archival solution also provides migration flexibility to organizations that desire to migrate their existing Exchange 2000 KMS solution to a more integrated archival solution in the Windows Server 2003 CA or for those customers who have public key certificates from third-party CAs. The Windows Server 2003 PKI solution allows import and migration of third-party public keys and certificates into the Windows Server 2003 CA as an escrow solution as well.

Windows Server 2003 key archival can be performed in two different ways: manual key archival and automatic key archival. Manual key archival refers to a process where users manually export private keys from their keys store and send them to a CA Administrator who in turn imports them to the protected CA database. Automatic key archival is done as part of the certificate enrollment process. It is possible to mark certificate templates to require key archival. In such a case, the private key will be sent to the CA as part of the certificate request and will be archived automatically by the CA. Manual and automatic key archival are discussed in depth later in this document.

**Understanding Key Archival and Recovery**

Key recovery implies that the private key portion of a public-private key pair may be archived and recovered. Private key recovery does not recover any data or messages. It merely enables a user to retrieve lost or damaged keys, or for an administrator to assume the role of a user for data access or data recovery purposes. In many applications, data recovery cannot occur without first performing key recovery. Figure 1 demonstrates a typical key archival and recovery scenario.




**Figure 1: Key Archival and Recovery**

1. The user requests a certificate from a CA and provides a copy of the private key as part of the request. The CA which is processing the request archives the encrypted private key in the CA database and issues a certificate to the requesting user.

2. The issued certificate can be used by an application such as EFS to encrypt sensitive files.

3. If at some point the private key is lost or damaged, the user can contact the company's Certificate Manager to recover the private key. The Certificate Manager, with the help of the Key Recovery Agent (KRA), recovers the private key, stores it in a protected file format, and sends it back to the user.

4. Once the user stores the recovered private key in the user's local keys store, it can be used by an application such as EFS to decrypt previously encrypted files or to encrypt new ones.

**Understanding EFS Data Recovery**

Encrypting a file always includes a risk that it cannot be read again. The owner of the private key, without which a file cannot be decrypted, might leave the organization without decrypting all of his or her files. Worse yet, he or she might intentionally or accidentally encrypt critical shared files so that no one else can use them. A user's profile might be damaged or deleted, meaning that the user no longer has the private key needed to decrypt encrypted files. Data recovery is one of the methods used for recovering files or data when encrypted files cannot be decrypted. By implementing data recovery and DRAs, every encrypted file's encryption key (FEK) is encrypted by using the DRA's public key in addition to the user's public key. By using the associated private key, any designated DRA can decrypt any encrypted file within the scope of the EFS recovery policy. Figure 2 demonstrates encrypted file format when data recovery is implemented.

**Figure 2: File Encryption Using Data Recovery**

For more information about Data Recovery, see Appendix B: Additional Information.

**Understanding When to Use Data Recovery versus Key Recovery**

There is no definitive answer of whether key recovery or data recovery should be used over the other option. Both solutions have technical advantages and disadvantages and the decision to use one or both is subjective. Both are viable solutions separately and in summation. This section is written solely to identify some of the key advantages and disadvantages so that an organization may make an informed decision.

Key recovery should be used when the policy of a given company or enterprise permits the private keys and certificates of a given user(s) to be retrieved. Key archival and recovery imply that a person other than the original user may gain access to the private keys of another user. If a company or enterprise does not permit a person other than the original requester to ever have access to private keys of other users, key archival and recovery should not be used.

Data recovery should be used when a company or enterprise requires the ability to recover data, but not access to the individual private keys of a user. For example, when private key recovery is provided, new operations can be performed when using the key, not just data decryption.

Data recovery and key recovery should not be used when a company or enterprise wants to protect data from all parties except the original user. If data should not be accessed or recovered by any person other than the author/owner, neither data recovery nor key recovery processes should be implemented.

**Table 1: Key Recovery versus Data Recovery**

| | Key Recovery | Data Recovery |
|---|---|---|
| Certificates re-enrollment is not required. | X | |
| Existing certificates revocation is not required. | X | |
| Administrators do not have access to the user's private key. | | X |
| Non-repudiation assurance. | | X |
| Re-Encryption of all previously encrypted data is NOT required. | X | |

**Key Recovery Best Practice**

Some key best practices should be followed when implementing key archival and recovery in an organization. First, key recovery by an administrator implies impersonation and that administrators who perform key archival should be highly trusted. Implementation of operations of a key archival and recovery solution should be carefully planned and monitored for security purposes.

**Other Best Practices**

The following additional best practices are also recommended:

- If a key is known to be compromised, it should be immediately revoked and never used again.
- Keys or certificates that are used to secure high-value transactions, or are identified as high-value certificates should not be recovered except under extreme circumstances.
- Private keys that are used for signing should never be archived or recovered as this introduces non-repudiation issues.

- When a key has been compromised or lost, it should be revoked before allowing key recovery. Despite key compromise, it may be necessary to recover a key for the purpose of decrypting old data and encrypting with a new key.
- Issue the least number of certificates and key pairs possible to ease key management and reduce user confusion.
- Issue encryption certificates with longer lifetimes than signing certificates.
- Issue only one valid key pair for a given application purpose at one time.
- Develop a recovery process that combines role separation of Certificate Managers and KRAs.
- Minimize the number of CAs that archive keys for a certificate purpose, that is, if possible, do not archive keys for users across a large number of CAs as recovery operations become confusing.
- If performing key archival over the Certification Services Web enrollment pages, it is important to use Secure Sockets Layer (SSL) on the enrollment Web site to protect the enrollment traffic from tampering or malicious interception.
- Use roaming user profiles if possible to reduce user confusion, accidental key loss, and the need for manual import and export operations to enable key roaming.

**Requirements**

Key archival and recovery using a Windows Server 2003 CA has several technical dependencies.

- Enrollment requires the Certificate Management Protocol using CMS (CMC) protocol, which is only available in Windows XP client, Windows Server 2003 clients, and through the xenroll.dll ActiveX control in the CA Web enrollment interface. Through the Web enrollment interface, Windows 2000 and Windows Millennium Edition may enroll for certificates with key archival.

  | **Note** |
  | --- |
  | Version 2 templates and key archival are not available to Windows 2000 Microsoft Management Console (MMC) enrollment. |

  | **Note** |
  | --- |
  | Windows 2000 and Windows Millennium Edition require that the key be marked for export during key archival enrollment. |

- Active Directory® must have the Windows Server 2003 schema extensions.
- Mandating key archival requires a version 2 certificate template.
- Key archival is only available on a CA running Windows Server 2003, Enterprise Edition.
- Only a Microsoft Enterprise Certification Authority may be used.

| **Note** |
| --- |
| Clients running down-level operating systems, such as Windows NT® 4.0 or Windows 98, will not be able to view certificate templates that allow key archival in the Web enrollment interface. |

# Understanding Manual Key Archival

Updated: December 6, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Manual key archival is supported on the Windows Server 2003 CA as a separate operation from enrollment while still offering centralized key archival. Users may export their private keys into *.pfx files [Public Key Certificate Standard (PKCS) #12 format] or through Outlook into the *.epf format. The following section describes the procedure to export private keys manually from a Windows client so that they may be manually archived on the CA. This is especially useful for users who have enrolled with third-party CAs that do not support key archival.

**Exporting Keys and Certificates**

Keys and certificates may be exported on Windows clients by one of two methods.

- PKCS #12 (*.pfx file) export from the Certificates MMC snap-in on Windows 2000 or Windows Server 2003
- *.epf file format from the Outlook 2000 or Outlook 2002 client

If a user has enrolled for Exchange Advanced Security with version 1 certificates (first offered with Exchange 5.0 Key Management Server), direct export from Outlook into the *.epf file format will be necessary. X.509 version 1 certificates and keys may not be exported into PKCS #12 format on the Windows client.

If only X.509 version 3 certificates have been used, the PKCS #12 format may be used.

**Exporting Keys from the MMC**

**To export the certificate and private key while logged in as the user**

1. Click the **Start** button, and then click **Run**.
2. Type **mmc.exe** and press **Enter**.

   An empty MMC shell appears.
3. Select the **Console** menu, and then select **Add/Remove Snap-in**.

   A dialog box appears with a list of all the snap-ins that have been added to this MMC shell.
4. Click **Add**.

   A list appears with all the registered snap-ins on the current machine.
5. Click the **Certificates** snap-in and click **Add**, choose **My User Account**, and then click **Finish**.
6. In the **Add Standalone Snap-in** dialog box, click **Close**. In the **Add/Remove Snap-in** dialog box, click **OK**.

   The MMC now contains the personal certificate store for the currently logged-on user.
7. Expand the tree view of the certificate store. Click through **Certificates - Current User, Personal**, and then **Certificates**. When you click the **Certificates** folder on the left, the right-hand pane will display a list of all the certificates for the currently logged on user.
8. Right-click the certificate intended for export.
9. Choose **All Tasks**, and then **Export** on the **Context** menu.

   A wizard will guide you through the export process.
10. Click **Yes**, **export the private key**, and then click **Next**.

When exporting a private key, the *.pfx file format is used. The *.pfx file format is based on the PKCS #12, which is used to specify a portable format for storing or transporting a user's private keys, certificates, and miscellaneous secrets. For more information about the PKCS #12, see Appendix B: Additional Information.

11. Select the appropriate check boxes, and then click **Next**.

    As a best practice, strong private key protection should also be used as an extra level of security on the private key when exporting. The private key should be deleted only if you are performing archival and will no longer use the key on that machine.

12. The *.pfx file format (PKCS #12) allows a password to protect the private key stored in the file. Choose a strong password, and then click **Next**.

13. The last step is to save the actual *.pfx file. The certificate and private key can be exported to any writeable device, including a network drive or disk. After typing or browsing for a file name and path, click **Next**.

Once the *.pfx file and private key have been exported, the file should be secured on a stable media and transferred in a secure manner to the CA on which the key will be imported in accordance with the organization's security guidelines and practices.

**Exporting Keys from Outlook**

**To export a key from Outlook**

1. In Outlook, click the **Tools** menu, and then select **Options**. Click the **Security** tab, and then click **Import/Export** (Figure 3).

**Figure 3: Outlook Security Options**

2. Click **Export your Digital ID** to a file, and then complete the **Filename**, **Password**, and **Confirm** (password confirmation) text boxes (Figure 4).



**Figure 4: Outlook Import/Export Digital ID Options**

3. Copy this file to a location accessible by the server or manually transport it using a disk.

#### Importing a Key Manually on a CA

### To import a key manually on a CA

- On the certificate server, open a command prompt window, and run the following command.

  **C:\CertUtil.exe –f –importKMS <name of file>**

| **Note** |
| --- |
| The –f flag is required when the key and certificate pair have not been issued from the CA in question. |

The file may be in one of three formats.

- KMS export file
- PKCS #12 format (*.pfx file)
- Outlook export format (*.epf)

| **Important** |
| --- |
| The previous command will work only after the CA was configured for key archival. For more information about the actions required for enabling key archival, see Implementing Key Archival Walkthrough. |

# Understanding Automatic Key Archival

Updated: December 6, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following sections document in detail how the key archival and key recovery processes work in Windows Server 2003 as well as step-by-step guides for implementing the features in production.

**Understanding the Archival Process**

The general steps in the archival of a user's private key are described in the following steps and in Figure 5.

1. The client finds CAs in Active Directory (enrollment services container in the configuration partition).
2. The client makes an authenticated Distributed Component Object Model (DCOM) connection to the CA and requests the CA's Exchange certificate (encryption certificate).
3. The CA sends the exchange certificate to the client.
4. The client validates that the CA's exchange certificate has been signed by the same key as the CA signing certificate and performs a revocation status check. This ensures that only the intended CA may decrypt the certificate request containing the private key.
5. The client encrypts the private key corresponding to the request with the CA exchange certificate public key, builds a CMC request, and sends a CMC full PKI request to the CA.
6. The CA validates that the encrypted private key cryptographically pairs with the public key in the certificate request.
7. The CA validates the signature on the request with the public key in the request.
8. The CA encrypts the private key from the user request with a random Triple Data Encryption Standard (3DES) symmetric key and then encrypts the symmetric key with one or more KRA public keys.
9. The CA saves the encrypted key binary large object (BLOB) containing the encrypted private key and the symmetric key encrypted with one or more KRA public keys to the CA database.
10. The CA processes the certificate request normally.
11. The CA responds to the client with a CMC full PKI response.

| **Note** |
| --- |
| Only RSA Security encryption keys may be archived in the CA database. Signature only keys as well as non-RSA key pairs will not be archived. Denied and resubmitted requests will also not archive private keys. |




**Figure 5: The Protocol for Key Archival**

**Certificate Request**
A certificate request can be driven from at least three different sources in Windows Server 2003. The sources built into the operating system are the Web browser (Internet Explorer), the certificate enrollment wizard, and auto-enrollment. All three methods call xenroll.dll (Microsoft Enrollment Control), which provides various methods to support certificate request generation and certificate installation. One of the formats the certificate request uses is the CMC format for certificate requests, which supports an optional encrypted data payload. Technically, any client that supports the CMC protocol may enroll with an Enterprise CA and request that the private key be archived by the CA. The CA enforces the archival option through a template flag. For more information about the CMC request format, see Appendix A: Certificate Request Structure.

**CA Exchange Certificate Retrieval**
Before the private key can be encrypted and delivered to the CA server, the client must first retrieve the CA's exchange certificate. The CA exchange certificate is an encryption certificate for the Certification Authority that can be used by clients to encrypt their private keys as part of their certificate request. The CA exchange certificate is issued by the CA itself where the subject and issuer are almost the same. However, the subject contains a suffix to distinguish the certificate from a root CA. Getting the CA exchange certificate can be accomplished using the existing ICertRequest::GetCACertificate method. The caller of the ActiveX xenroll.dll control will need to retrieve this certificate and verify it prior to calling xenroll.dll to create the request and to encrypt the private key. Xenroll will also perform this action. The client will verify that the exchange certificate is signed by the same key that issued the CA certificate; this step is performed to provide additional security and to prevent man-in-the-middle attacks. For that reason, a CA may not use a certificate issued by another CA for an exchange certificate.

**CA Exchange Certificate Generation**
A Windows Server 2003 CA will automatically generate a short-lived exchange certificate for use of the key archival mechanisms. By default, if the CA Exchange template is available, it will be used for extensions and the validity and overlap periods. If the validity and overlap periods in the registry differ from the template, the registry values are rewritten, so there will be consistent behavior if the template is unavailable in the future. If the template is not available, hard-coded extensions will be used along with the registry validity and overlap period settings.

If the following registry flag is configured using certutil.exe, the template must be available or the attempt to generate a CA Exchange certificate will fail. In addition, the machine object of the CA (LOCAL SYSTEM) must have Enroll access to the template. This flag is not currently set by setup; it must be manually applied.

**certutil –setreg ca\CRLFlags +CRLF_USE_XCHG_CERT_TEMPLATE**

The CA Exchange certificate template has a default expiration of one week and a template overlap period of one day. Note that the previous settings apply to both enterprise and stand-alone CAs.

**Restricting Key Archival**
Windows Server 2003 may be effectively prevented from archiving private keys through the use of qualified subordination and policy constraints. When submitting and processing a request to a parent or root CA for a subordinate CA certificate, the parent CA may exclude the key archival policy in the subordinate CA. For more information about Qualified Subordination, see Appendix B: Additional Information.

**Private Key Payload**
For Windows Server 2003 clients and servers, the encrypted private key is stored in an unauthenticated attribute as a Cryptographic Message Syntax (CMS) EnvelopedData object. This poses no concern from a security perspective as the private key is later validated to cryptographically match the public key in the certificate request, which is also signed by the same private key. For more information about the format and object structure, see Appendix A: Certificate Request Structure.

**Key Verification**
The CA must check that the encrypted private key is cryptographically associated with the public key in the certificate request. Because the client encrypted the private key using the public key from the CA's exchange certificate, the CA can decrypt the private key payload contained in the request. For encryption keys, the CA encrypts randomly generated data with the public key in the request, and then decrypts the data with the private key passed in the unauthenticated attributes of the CMC request. The decrypted data is verified against the original random data. If any of the validation steps fail, or if the data does not match, the request is rejected. For signing keys, the CA will reject the request and will not archive a key that is marked for signature only.

To verify the cryptographic association between the public and private keys, the CA loads the key pair into a cryptographic service provider (CSP) that supports the specified algorithms. The CA will use the default-enhanced RSA CSP on the CA by default or a hardware CSP, if available. Both the encrypted private key material payload and the private key loaded into the server-based CSP are securely discarded since the private key will be archived using a different key controlled by a recovery agent and is not available for the CA after verification of the payload. The decrypted private key memory is zeroed out prior to freeing the memory.

**Key Encryption**
Once the private key has been successfully verified, the CA must then encrypt the private key and archive it for future recovery. The CA will not encrypt the private key using its own key(s) to provide separation between the recovery agent role and the CA administrator role. The certificate request process and certificate issuance are completed after key archival has occurred. For more information about role separation, see Configuring the CA to Allow Key Archival.

The private key will be encrypted by the CA using a dynamically generated symmetric key that is itself encrypted using a public key from a KRA certificate (or certificates) configured on the CA. Each certificate request and private key archival operation will generate a new random 3DES symmetric key using the same CSP that was used to generate the CA signing key. This ensures that a hardware security module (HSM) can be used to generate the long-term keys that protect the individual private keys. The symmetric key algorithm is 3DES by default and cannot be changed.

If multiple recovery agents are associated with the CA, the symmetric key will be encrypted individually with each recovery agent's public key to allow any one recovery agent to perform a recovery operation. The actual KRA(s) used may randomly vary if a round-robin selection of KRAs is chosen. For example, if a CA Administrator configures four KRA certificates on a CA for use, but requires a minimum of two KRAs to be used, when the CA service starts, a random start point in the KRA list is chosen. Two of the four KRAs will be chosen from that start list start point and then for each subsequent archival operation, the list order will be incremented by two. Therefore, a round robin affect is achieved for all key archival operations. These encrypted keys along with the issued certificate will be collectively referred to as the recovery BLOB and are stored as a PKCS #7–encrypted BLOB in the database.

**Key Archival**
The private key database is the same as the database used to store the certificate requests. The Windows Server 2003 Certification Authority database has been extended to support storing the encrypted private key along with the associated encrypted symmetric key and issued certificate. The recovery BLOB will be stored in the same row as the signed certificate request and any other information the CA persists in its database for each request transaction. The actual encrypted BLOB is stored as an encrypted PKCS #7 BLOB. For more information about the format and object structure, see Appendix A: Certificate Request Structure.

The Microsoft Certification Authority uses the Microsoft Jet database engine upon which various Jet utilities may be used for maintenance purposes.

# Understanding User Key Recovery

Updated: December 6, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A recovery operation is initiated by an end-entity that has lost access to a private key. The key recovery process is explained in Figure 6.




**Figure 6: Key Recovery Process**

The following are the key recovery steps.

1. The Certificate Manager (CA Officer) locates and retrieves the user's encrypted private key from the CA database. The encrypted BLOB(s) is protected by access control lists (ACLs) to ensure that only a valid Certificate Manager is able to copy the BLOB from the database. Since it is encrypted with the KRA's public key, the Certificate Manager cannot extract the user's private key but it can retrieve it from the CA. The encrypted PKCS #7 BLOBs in the database contain the Issuer name and serial number of each KRA certificate for KRA identification purposes. Once extracted, the Certificate Manager sends it to the appropriate KRA for decryption. For more information about the format and structure of the recovery BLOB, see Appendix A: Certificate Request Structure.

2. The KRA decrypts the user's private key and stores it in a password-protected file format (PKCS #12) and sends it to the user.

3. The user imports the key to the user's local, protected, keys store.

It is at the discretion of organizational policy whether a KRA and a Certificate Manager may be combined into one role or separate roles. For better security, it is recommended that the Certificate Manager and the KRA roles be separated. Because the Windows Server 2003 Certification Authority supports role separation, an organization could easily implement a two-step process to recover the private key(s) of a user.

**Understanding Role Separation**

Role-based administration is used to organize CA administrators into separate, predefined task-based roles. It is recommended that the management roles are distributed among several individuals in your organization to ensure that a single individual cannot compromise PKI services. Role separation enables one person to audit the actions of another person.

The Common Criteria PKI management roles in Windows Server 2003 include the following:

1. **CA Administrator** Configures and maintains the CA, designates other CA administrators and certificate managers, and renews CA certificates.

2. **Certificate Manager** Approves or denies certificate enrollment requests and revokes issued certificates.

3. **Backup Operator** Performs backups of the CA database, the CA configuration, and the CA's private and public key pair (also known as a key pair).

4. **Auditor** Defines what events are audited for Certificate Services and reviews the security log in Windows Server 2003 for success and failure audit events that are related to Certificate Services.

| **Note** |
| --- |
| Role-based administration is supported by both Windows Server 2003 Enterprise CAs and stand-alone CAs running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. |

To help determine role separation, you can use the Common Criteria specification, which defines security standards for all forms of network security and includes specifications for managing PKIs.

For more information about Role-Separation and Common Criteria, see Appendix B: Additional Information.

**Understanding the Key Recovery Agent Role**
KRAs are Information Technology (IT) administrators who can decrypt users' archived private keys. An organization can assign KRAs by issuing KRA certificates to designated administrators and configure them on the CA. The KRA role is not one of the default roles defined by the Common Criteria specifications but a virtual role that can provide separation between Certificate Managers and the KRAs. This allows the separation between the Certificate Manager, who can retrieve the encrypted key from the CA database but not decrypt it, and the KRA, who can decrypt private keys but not retrieve them from the CA database. For more information about how to implement Key Recovery Agents, see Implementing Key Archival Walkthrough.

**Key Recovery Agent Certificate**
A new certificate template exists in the Windows Server 2003 schema to support the KRA role. This certificate will use a unique Extended Key Usage extension to identify a KRA certificate. A Windows Server 2003 Certification Authority will only use certificates that have been properly formatted as per the following information, although it is not necessary for the certificate to contain the Microsoft-specific extensions.

KRA certificates, when issued by an Enterprise CA, are automatically published in the configuration container of Active Directory. The actual certificates are published to the userCertificate attribute of the KRA object when issued to an IT administrator. The location in the configuration container is as follows:

**(CN=KRA,CN=Public Key Services,CN=Services,CN=Configuration,DC=<domainname>)**.

The KRA certificate contains the following X.509v1 fields.

- Version
- Serial Number
- Signature Algorithm
- Valid From
- Valid To
- Subject
- Issuer
- Public Key

The KRA certificate contains the following X.509v3 extensions identified in RFC 3280.

- Authority Key Identifier
- Subject Key Identifier
- Authority Information Access
- Key Usage
- Subject Alternative Name
- CDP (CRL Distribution Point)
- Extended Key Usage (Key Recovery Object Identifier = 1.3.6.1.4.1.311.21.6)
- Application Policies (Policy Identifier = Key Recovery Agent)

The KRA certificate will contain the following X.509v3 extensions specific to Microsoft.

- Certificate Template Name
- Certificate Template Information

| **Note** |
| --- |
| None of the extensions is marked as critical; however, the Key Recovery Object Identifier must exist in the Extended Key Usage extension for the certificate to be used. |

**Understanding the Key Recovery Process**

Users may require their private keys to be recovered for a multitude of reasons. Key recovery from a Windows Server 2003 CA may be accomplished using the command-line tool, certutil.exe.

The recovery of a private key is a manual process that requires the user(s) to contact an administrative authority to perform the necessary processes. It should be a best practice of any organization to

1. Validate the identity of a user requesting key recovery.
2. Separate the roles of CA Officer and KRA as a minimum of two physical persons.
3. Develop a mechanism to securely deliver the private keys and passwords to end users. Examples could include e-mail of the *.pfx file and leaving a voice mail with the password for the *.pfx file for the user. Discussion of these best practices is beyond the scope of this white paper.

**Finding Recovery Candidates**

A CA Officer (Certificate Manager) can perform recovery of private key(s) using the CN (CommonName), UPN (UserPrincipalName), down-level name (domainname\username), certificate serial number of the certificate, or an SHA1 (Secure Hash algorithm) hash (thumbprint) of the certificate. The process is known as finding candidate certificates for recovery.

Examples:

- User1@nwtraders.com (denotes UPN)
- nwtraders\user1 (denotes the down-level name)
- Users\nuser1 (denotes a user in the default users container of Active Directory)
- User1 (denotes the CN)
- <serial number of the certificate>
- <SHA1 hash (thumbprint) of certificate>

**Finding Recovery Candidates by Using Command-Line Tools**

To recover a user using their CN, type the following command.

**Certutil -getkey <cn of user> outputblob**

If only one certificate has been issued to that user, that certificate and private key will be generated into the specified <outputblob> file. If more than one certificate has been issued to that CN, certutil.exe will list the serial number(s) of all the certificates issued to that CN on a specific Certification Authority.

Example:

```
"user1.nwtraders.com\CA1"

Serial Number: 1464be10000000000007
Subject: CN=user1, CN=Users, DC=nwtraders, DC=com
NotBefore: 1/13/2001 11:51 AM
NotAfter: 1/13/2002 11:51 AM
Template: KeyA, KeyArchival
Cert Hash(sha1): a2 7f 77 bc 2f 7b eb 26 bd 3e ed 43 b8 2a 24 04
2e 55 b8 64


Serial Number: 1464fcbc000000000008

Subject: CN=user1, CN=Users, DC=nwtraders, DC=com
NotBefore: 1/13/2001 11:51 AM
NotAfter: 1/13/2002 11:51 AM
Template: KeyA, KeyArchival
Cert Hash(sha1): 21 bd 88 2c 2a 84 0c e5 57 7b 2a bf f0 43 4b b3
ed bf 02 5a
```

Once the actual serial number of the certificate is known, the command may be run again with the serial number instead of the CN to retrieve the actual recovery BLOB.

| **Important** |
| --- |
| If a CA Administrator has configured the server to use, for example, three of a possible five available KRAs, certutil –recoverkey will list the three KRAs that were used to encrypt the selected subject's private key. At the time of archival, the three KRAs were randomly selected (out of five possible) by the CA to encrypt the subject's private key. To perform the recovery operation, it will be necessary to have one of any of the KRAs listed by certutil -recoverkey. |

#### Finding Recovery Candidates by Using the Certification Authority MMC Snap-In

A CA Officer may determine the serial number of a certificate that has been archived.

**To determine the serial number of an archived certificate**

1. Log on to a CA or a Windows XP Professional machine that has the **Administration Tools** pack installed as a CA Officer who has **Certificate Management** authority of the user(s) in question.
2. On the **Administrative Tools** menu, open **Certification Authority**.
3. In the console tree, expand **Certification Authority**, and then click **Issued Certificates**.
4. On the **View** menu, click **Add/Remove Columns**.
5. In the **Add/Remove Columns** dialog box, in the **Available Column list**, select **Archived Key**, and then click **Add**.

   Archived Key should now appear in the Displayed Columns listing.
6. Click **OK**.
7. In the details pane, scroll to the right and ensure that the last issued key has a value in the Archived Key column.
8. Double-click the user certificate.
9. Click the **Details** tab
10. Write down the serial number of the certificate (do not include spacing between digit pairs). This is required for recovery.
11. Click **OK**.
12. Close **Certification Authority**.

**Command-Line Key Recovery**
Command line key recovery is done by using the certutil.exe command-line tool. This tool is installed by default on the certificate server and it can be installed separately on a Windows XP Professional or Windows Server 2003 machine as part of the Windows Server 2003 Administration Tool Pack. The certutil.exe command structure was designed to perform batch and scripted processes.

**Recovery Using the Certificate Serial Number**
Once the serial number(s) is known, key recovery can occur through the certutil.exe command- line tool.

### To recover private key(s)

1. Open a command-prompt window.

2. At the command prompt, type **Certutil -getkey ################### outputblob** where **###################** is the serial number of the certificate that should be recovered and outputblob is the file name of the encrypted BLOB that is extracted from the CA database.

3. If a specific CA is to be targeted instead of all Enterprise CAs in the forest, use the following syntax.

   **Certutil -config <machine name\CA name> -getkey ################### outputblob**

4. The outputblob file is a PKCS #7 file containing the KRA certificate(s) and the user certificate with its entire certificate chain. The inner content is an encrypted PKCS #7 structure containing the private key [encrypted to the KRA certificate(s)].

5. To decrypt the PKCS #7 encrypted BLOB, the logged on user must have the private key of one or more of the KRAs for the target encrypted BLOB. If the CA Officer is not a Key Recovery Agent, the CA Officer must transfer the encrypted BLOB file to a user that holds a KRA private key for further processing.

6. At the command prompt, type the following command:

   **Certutil -recoverkey outputblob user.pfx –p password** where:

   - outputblob is the encrypted PKCS #7 file to be decrypted.
   - user.pfx is the output file name for the user private keys in PKCS #12 format.
   - password is the password for the *.pfx file.

   The PKCS #12 format allows for the private key file to be protected with a password. Certutil.exe will prompt the KRA for a password when generating the PKCS #12 file.

   The system will search for a valid private key in the store of the logged in user that corresponds to a valid KRA certificate that was used to encrypt the recovery BLOB. If the user does not have the proper private key in their local store, they will receive an error.

7. Enter and confirm a password for the file that cannot be randomly guessed. The user should be given the password through a secure out-of-band mechanism that is separate from the file itself to ensure strong security in the recovery process.

| **Note** |
| --- |
| An event log message will be generated when a private key is recovered from the database. The event log message will be similar to the following: |

```
Event Type: Success Audit
Event Source: Security
```

```
Event Category: Object Access
Event ID: 787
Date:  2/19/2001
Time:  5:23:45 PM

User:  NWTRADERS\User1

Computer: SERVER1

Description: Certificate Services retrieved an archived key.
Request ID 12
```

**Recovery Batch File**

The certutil.exe –getkey command is an advanced tool that can be used with explicit commands and also build batch file syntax to perform a complete recovery operation. The following is the command-line syntax for certutil.exe –getkey.

Usage:

**CertUtil -GetKey [Options] SearchToken [RecoveryBlobOutFile]**

**Retrieve archived private key recovery blob**

**SearchToken**—Used to select the keys and certificates to be recovered.

**RecoveryBlobOutFile**—Output file containing a certificate chain and an associated private key, still encrypted to one or more KRA certificates.

If the following command is performed using any of the previous SearchToken criteria, the certutil.exe tool will output the complete syntax that can be used.

Example:

**certutil -v -getkey user1@northwindtraders.com >myBatchfile.bat**

Outputfile:

```
@goto start
Querying northwind5.nwtraders.com\TestCA9....................


"northwind5.nwtraders.com\TestCA9"
  Serial Number: 611e23c200030000000e
  Subject: E=user1@nwtraders.com, CN=user1, CN=Users, DC=nwtraders, DC=com
  UPN:user1@nwtraders.com
  NotBefore: 1/12/2002 1:24 PM
  NotAfter: 1/12/2003 1:24 PM
  Template: EFS2
  Version: 3
  Cert Hash(sha1): d6 41 99 e7 e7 24 73 34 02 41 53 2d 29 11 a8 3b e6 aa 12 2f

  Serial Number: 6131f9c300030000000f
  Subject: E=user1@nwtraders.com, CN=user1, CN=Users, DC=nwtraders, DC=com
  UPN:user1@nwtraders.nttest.microsoft.com
  NotBefore: 1/12/2002 1:45 PM
  NotAfter: 1/12/2003 1:45 PM
  Template: EFS2
  Version: 3
  Cert Hash(sha1): 1a cc 8d 26 7f 9f 70 6c 65 05 a0 84 8c 4c e9 b7 b4 6c 66 a3

:start

@echo Version 3 certificates and keys:
CertUtil -config "northwind5.nwtraders.com\TestCA9" -getkey
611e23c200030000000e "user1-611e23c200030000000e.rec"

CertUtil -config "northwind5.nwtraders.com\TestCA9" -getkey
6131f9c300030000000f "user1-6131f9c300030000000f.rec"

CertUtil -p "UQcYLsJ(57s]FuBl" -recoverkey -user "user1-611e23c200030000000e.rec" "
user1-611e23c200030000000e.p12"

CertUtil -p "UQcYLsJ(57s]FuBl" -recoverkey -user "user1-6131f9c300030000000f.rec" "
```

```
user1-6131f9c300030000000f.p12"

CertUtil -p "UQcYLsJ(57s]FuBl,iG1-bOt&tqdvJiu1" -MergePFX -user "user1-611e23c20003
0000000e.p12,user1-6131f9c300030000000f.p12" "user1.p12"

@del user1-611e23c200030000000e.rec
@del user1-611e23c200030000000e.p12
@del user1-6131f9c300030000000f.rec
@del user1-6131f9c300030000000f.p12
@echo PASSWORD: "iG1-bOt&tqdvJiu1"

@goto exit
CertUtil: -GetKey command FAILED: 0x8002802c (-2147319764)
CertUtil: Ambiguous name.
```

The following is an explanation of the previous syntax, which can be run by Certificate Managers that hold a valid KRA private key.

1. The tool finds all possible CAs that are registered in Active Directory. Then, it will query each CA for keys that have been archived for the specified user.

   ```
   @goto start
   Querying northwind5.nwtraders.com\TestCA9.....................
   ```

2. The tool will show the archived keys found for the user on each CA that is queried. It will list the most important details on each archived key found that will be useful for Certificate Managers and administrators.

   ```
   "northwind5.nwtraders.com\TestCA9"
     Serial Number: 611e23c200030000000e
     Subject: E=user1@nwtraders.com, CN=user1, CN=Users, DC=nwtraders, DC=com
     UPN:user1@nwtraders.com
     NotBefore: 1/12/2002 1:24 PM
     NotAfter: 1/12/2003 1:24 PM
     Template: EFS2
     Version: 3
     Cert Hash(sha1): d6 41 99 e7 e7 24 73 34 02 41 53 2d 29 11 a8 3b e6 aa 12 2f

     Serial Number: 6131f9c300030000000f
     Subject: E=user1@nwtraders.com, CN=user1, CN=Users, DC=nwtraders, DC=com
     UPN:user1@nwtraders.nttest.microsoft.com
     NotBefore: 1/12/2002 1:45 PM
     NotAfter: 1/12/2003 1:45 PM
     Template: EFS2
     Version: 3
     Cert Hash(sha1): 1a cc 8d 26 7f 9f 70 6c 65 05 a0 84 8c 4c e9 b7 b4 6c 66 a3
   ```

3. The tool will retrieve the encrypted recovery BLOB(s) from the CA. Note that the administrator must be a Certificate Manager for the recovered user(s) on the CA.

   ```
   @echo Version 3 certificates and keys:
   CertUtil -config "northwind5.nwtraders.com\TestCA9" -getkey
   611e23c200030000000e "user1-611e23c200030000000e.rec"

   CertUtil -config "northwind5.nwtraders.com\TestCA9" -getkey
   6131f9c300030000000f "user1-6131f9c300030000000f.rec"
   ```

4. The tool will decrypt the recovery BLOB(s) and generate *.pfx files (PKCS #12 format) for each recovered key and set a random password as denoted after the –p parameter. Note that the administrator must have a valid KRA private key to perform this step.

   ```
   CertUtil -p "UQcYLsJ(57s]FuBl" -recoverkey -user "user1-611e23c200030000000e.rec" "
   ```

```
CertUtil -p "UQcYLsJ(57s]FuBl" -recoverkey -user "user1-611e23c200030000000e.rec"
user1-611e23c200030000000e.p12"

CertUtil -p "UQcYLsJ(57s]FuBl" -recoverkey -user "user1-6131f9c300030000000f.rec" "
user1-6131f9c300030000000f.p12"
```

5. The tool will merge the multiple *.pfx files (if applicable) into a single *.pfx file to simplify the process for the user installing recovered keys. The certutil.exe –MergePFX command is automatically used to perform this process.

```
CertUtil -p "UQcYLsJ(57s]FuBl,iG1-bOt&tqdvJiu1" -MergePFX -user "user1-
611e23c200030000000e.p12,user1-6131f9c300030000000f.p12" "user1.p12"
```

6. The tool will clean up any temporary files used during the process.

```
@del user1-611e23c200030000000e.rec
@del user1-611e23c200030000000e.p12
@del user1-6131f9c300030000000f.rec
@del user1-6131f9c300030000000f.p12
@echo PASSWORD: "iG1-bOt&tqdvJiu1"

@goto exit
```

Note that the following output is expected when the user has multiple archived keys.

```
CertUtil: -GetKey command FAILED: 0x8002802c (-2147319764)
CertUtil: Ambiguous name.
```

This output implies that a full key recovery could not be performed because the command line was not specific enough to retrieve a specific recovery BLOB.

**CA Officer Does Not Have a KRA Certificate**

As mentioned previously, a CA Officer need not also hold a KRA certificate. In the case where the CA Officer does not hold a valid KRA certificate and only has permissions to retrieve a recovery BLOB, running certutil.exe –getkey <serial number> with no output file name will list the certificate hashes (thumbprints) of all of the KRA(s) that may be used to decrypt (recover) the recovery BLOB in question. The CA Officer can then determine which KRA(s) may be used and send the recovery BLOB to one of the valid KRA(s).

The previous procedure is extremely useful in an organization that has a round-robin selection of KRA(s) when archiving private keys. In a round-robin archival, it may be difficult to always predict which KRA(s) were used to encrypt any specific private key of a user. It is probably a good best practice to keep the list of all KRA certificates hashes accessible so that they may be readily identified in the previous scenario.

The following certutil command may be used to list all the information from Active Directory, including KRA certificate hashes.

**Certutil.exe –DS – v > c:\output.txt**

| **Note** |
| --- |
| This command will dump all PKI information and certificates from Active Directory and may be more usable when placing the output in a text file as shown previously. |

**Setting the Timeout Option**

When keys are recovered from a CA using the command line or the key recovery tool, the recovery tool(s) will build a complete chain for the end-entity certificate, if possible. Sometimes a chain may not be able to be built after a long time due to infrastructure changes, CA certificate availability, and so on; the tool may also timeout when trying to build these chains. The certutil –recoverkey command and other commands that verify chains or construct *.pfx files accept a –t MilliscondCount option. This allows key recovery to avoid a 15-second timeout for each user key when building a chain. Fetching the certificate specified in the AIA extension Uniform Resource Locators (URLs) can time out when the recovered keys are associated with expired encryption certificates and the CAs have been decommissioned. For more information about certificate chains and status, see Appendix B: Additional Information.

**Using the Key Recovery Tool**

The key recovery tool is a Windows Server 2003 Resource Kit tool for recovering keys from a Windows Server 2003 CA

instead of using the certutil.exe command-line tool.

The tool may be launched from the command line by running krecover.exe. The tool is dependent on the existence of the certutil.exe and other binaries, which are available by default in Windows Server 2003 and in the Administration tool pack of Windows XP Professional. The tool is also dependent on one dynamically linked library (crutredir.dll), which is installed by the resource kit.

When launched, the tool will look similar to the following:




**Figure 7: Key Recovery Tool Interface**

The tool will automatically enumerate the forest for the currently logged on user and identify Enterprise CAs that may have archived private keys stored. The tool will filter out the following from the list of available CAs.

- CAs that are not Enterprise CAs
- CAs that are not running Windows Server 2003, Enterprise Edition
- CAs that are running Windows 2000
- CAs that the current user does not have CA Administrator or CA Manager permissions for

The following are several options enabled by the tool.

- **Search**—Allows a Certificate Manager to search for archived keys for a user based on several search criteria options.
- **Recover**—Allows a Certificate Manager who holds a KRA private key to recover a user from a CA database.
- **Show KRA**—Displays the valid KRA certificates that may decrypt the user private key.
- **Retrieve BLOB**—Allows a Certificate Manager to retrieve the encrypted user's private key from the database for the purpose of transferring to a KRA.
- **Decrypt BLOB**—Allows a KRA to decrypt and create a *.pfx file for an encrypted BLOB retrieved from the database by a Certificate Manager.

If the user does not have adequate permissions on the CA to recover a user or is not a Certificate Manager, an error similar to the following will be displayed.

**Figure 8: Failed to enumerate the KRA certificates Error Message**

**Recovering a Specific User's Keys**

To recover the private key(s) of a specific user, enter a user name and click **List** to list all the certificates for that user with an archived key. In the **User Name** field, enter the user name (as described previously), or domain\user_name, or user_name@northwindtraders.com—any input that is valid for the certutil -getkey command.

| **Important** |
| --- |
| For the key recovery tool to recover a key for a user, the same permissions apply as in the certutil.exe command-line as described in the previous section. |

If no key(s) is found for the selected user on the selected CA(s), the following dialog box will appear.

**Figure 9: No archived keys were found Message**

– or –

**Figure 10: No archived keys were found Message**

Here, you can select one or more certificates to be recovered, and then click **Recover**. If you select more than one certificate, you will be prompted if you want all the keys to be recovered in a single *.pfx file, or if you want separate *.pfx files for every certificate.




**Figure 11: Single or Separate *.pfx Prompt Message**

In either case, you will be prompted for a destination for the *.pfx file. The dialog box looks like the following:




**Figure 12: Output File Dialog Box**

When saving a *.pfx file, the tool will prompt for a location to save the *.pfx file and require a password to be set on the PKCS #12 file. When recovering a single private key, the default name for the *.pfx file will be the certificate serial number that corresponds to the private key. The output file name may be changed.

If the tool cannot recover a key for a user, the following dialog box will be displayed.




**Figure 13: Recovery Failure Message**

If you do not have a valid certificate and private key for a KRA that was used to encrypt and archive the private key of the user on the CA, the following error will be displayed.

**Figure 14: Invalid Certificate and Private Key Message**

**Recovering Version 1 and Version 3 Certificates**

A Windows Server 2003 CA may be used to archive both X.509, and version 1 and version 3 certificates as detailed previously. The key recovery tool allows recovery of both types, although x.509 version 1 certificates must be exported separately as they may only be exported in the *.epf format for Outlook. The Windows Server 2003 CA supports the following version 1 certificates and keys to be imported and exported from the database.

- Profile Version=2, CAST=3, Protection=40, EPFALG_CAST_MD5
- Profile Version=2, CAST=3, Protection=64, EPFALG_CAST_MD5
- Profile Version=3, HashCount=0, Protection=128, EPFALG_RC2_SHA
- Profile Version=3, Protection=128, EPFALG_3DES

If no version 1 certificates are in the recovery selection, the user will not be prompted any differently. If one x.509 version 1 certificate is in the selection, the user will be prompted for an *.epf file name for export. If more than one x.509 version1 certificate is selected, the user will be prompted whether the user would like one *.epf file or separate *.epf files for each key recovered. X.509 version 1 certificates will be recovered first and then the x.509 version 3 certificates will be recovered. If both version 1 and version 3 certificates have been selected, the user will be prompted twice whether the certificates should be

recovered in a single file or separate files. X.509 version 1 and version 3 certificates may not be recovered in a single file together.

**Importing Recovered Keys**

The following steps would be performed by an end user who has received the recovered key file (*.pfx file) and associated password from an administrator. It is assumed that the password and associated file have been transmitted to the user in a method consistent with the organization's security policy.

**To import a recovered key for an individual user**

1. Log on as the user who needs to recover their private key(s).
2. Click the **Start** button, and then click **Run**.
3. Type **mmc.exe**, and then press **Enter**. An empty MMC shell starts up.
4. Select the **Console** menu, and then click **Add/Remove Snap-in**. A dialog box appears with a list of all the snap-ins that have been added to this MMC shell.
5. Click **Add**.

   A list appears with all the registered snap-ins on the current machine.
6. Double-click the **Certificates** snap-in, choose **My User Account**, and then click **Finish**.
7. In the **Add Standalone Snap-in** dialog box, click **Close**, and then click **OK** in the **Add/Remove Snap-in** dialog box.

   The MMC now contains the personal certificate store for the user.
8. Expand the tree view of the certificate store. Click **Certificates**, **Current User**, **Personal**, and then **Certificates**.
9. In the console tree, right-click **Personal**, click **All Tasks**, and then click **Import**.
10. In the **Certificate Import Wizard** dialog box, click **Next**.
11. On the **Files to Import** page, in the **File name** box, type the name and path of the recovered key file (*.pfx), and then click **Next**.
12. On the **Password** page, in the **Password** box, type the password for the key file, and then click **Next**.
13. On the **Certificate Store** page, click **Automatically select the certificate store based on the type of certificate**, and then click **Next**.
14. On the **Completing the Certificate Import Wizard** page, click **Finish**. The wizard will report that the import was successful.

| **Note** |
| --- |
| A *.pfx file can also be selected (double-click) to invoke the certificate import wizard. |

# Implementing Key Archival Walkthrough

Updated: December 6, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The first step in enabling key archival on a CA is enrolling for one or more KRA certificates.

**Enrolling a Key Recovery Agent**
The first step is to enroll KRAs for KRA Certificates. The following section explains the steps for enrolling a KRA.

**Configuring the Certificate Templates**
A certificate template suitable for creating KRA certificates is installed in Active Directory. By default, only an Enterprise Administrator or a Domain Administrator may request a KRA certificate as defined by the default ACLs on the KRA certificate template. The certificate template ACLs can be viewed in the Certificate Templates MMC snap-in; in addition, using the Certificate Templates MMC snap-In, certificate templates can be cloned or edited.

| **Note** |
| --- |
| Only a domain with the Windows Server 2003 schema will support version 2 templates and only a Windows Server 2003, Enterprise Edition may issue a version 2 template certificate. |

**To configure a certificate template**

1. Log on as an Enterprise or Domain Administrator to the CA machine.

2. Click the **Start** button, click **Run**, and then type **certtmpl.msc**.

3. Click **OK**.




**Figure 15: Certificate Templates MMC Snap-In**

4. In the details pane, double-click **Key Recovery Agent**.

5. In the **Key Recovery Agent Properties** dialog box, click the **Security** tab.



**Figure 16: Key Recovery Agent Template Properties**

6. Add the appropriate user(s) or group(s) with both Read and Enroll permission.

7. Click **OK** to close the dialog box.

Next, the Certification Authority must be configured to issue this type of certificate.

**Certificate Template Permissions**
For a user or a computer to enroll for a certificate template, it must have appropriate permissions [access control entries (ACEs)] set on the template in Active Directory. A user or computer must have both Enroll and Read permissions to enroll for a selected certificate template. The Read permission allows the template to be discovered by the user and the Enroll permission is enforced by the enterprise CA when a user requests a certificate for a selected template. The enterprise CA must also have Read permissions on a template to enumerate the template in the directory and issue certificates based on that template. Normally, the enterprise CA is included in the Authenticated Users group, which has Read permissions by default on a template.

The Full Control permission is given to Enterprise Administrators by default on installation of a fresh Windows Server 2003 domain. If a domain has been upgraded from Windows 2000, Enterprise Administrators will not have this permission by default and the Full Control permission allows a user to set or modify the permissions on a selected template.

The Autoenroll permission is set on a template when a user or computer wants to automatically enroll for a selected certificate template. The Autoenroll permission is needed in addition to the Enroll permission for a user to enroll for a given certificate template.

The Write permission allows a user to modify the contents of a certificate template. Note that only a version 2 certificate with a Windows Server 2003 schema may be modified and version 1 certificate templates may only have the ACLs modified.

**Smart Card Support**
Smart cards are supported for use in conjunction with KRA certificates. It may be necessary to use a smart card and CSP that supports at least an 8-KB smart card to enroll for a KRA certificate on a smart card. If a smart card does not have adequate memory to support a KRA certificate, the following error will be generated on enrollment.

**Error: 0x80100028**

**An attempt was made to write more data than would fit in the target object**

All recovery operations are supported using a Smartcard. The system will prompt the recovery agent to insert an appropriate Smartcard when the key is needed to decrypt the recovery BLOB.

**Configuring an Enterprise CA to Issue KRA Certificates**
An Enterprise CA must be configured to issue a KRA certificate.

**To configure an EnterpriseCA to issue a KRA certificate**

1. On the **Administrative Tools** menu, open the **Certification Authority** snap-in.

2. In the console tree, expand **Certification Authority**, and then click **Certificate Templates**.

3. Right-click the **Certificate Templates** node, click **New**, and then click **Certificate Template to Issue**.

4. In the **Select Certificate Template** dialog box, click **Key Recovery Agent**, and then click **OK**.

5. Close the **Certification Authority** MMC snap-in.

6. The last manual step is to add the Certification Authority machine account to the Pre-Win2K Compatible Access group in every domain in which users will be using key archival. If this mandatory step is not performed, a CA Officer may not be able to manage groups of users. The Pre-Win2K Compatible Access group allows the CA to enumerate a user account and determine the group membership for CA Manager's capability.

**Enrolling a User with a KRA Certificate**

A user may enroll for a certificate with a CA by using the Certificates MMC snap-in or through the CA Web pages. In the case of the KRA template, the template is marked to be "pended" by the CA, which requires that the certificate request be approved first by a CA Administrator or a Certificate Manager before it is issued. Pended certificate requests may only be retrieved through the Web enrollment interface or through the auto-enrollment process. For more information, see Appendix B: Additional Information.

| **Important** |
| --- |
| It is strongly recommended not to automatically enroll KRA certificates as this may cause confusion for CA Administrators when automatic enrollment is unintentionally initiated resulting in additional KRA certificates in Active Directory. |

**KRA Web Enrollment**

**To enroll through a Web page**

1. Connect to the CA using a Web URL, for example:

   http://<CA Machine Name>/Certsrv

   | **Note** |
   | --- |
   | The Microsoft Certification Authority uses an ActiveX control known as xenroll.dll, which can be downloaded to client browsers that support ActiveX controls. Windows XP and Windows Server 2003 clients both ship with the ActiveX control pre-installed. |

   A Web site will open allowing you to request a certificate.

2. Select **Request a Certificate**.

3. On the next Web page, select **Advanced Certificate Request**.

4. Select **Create and submit a request to this CA**.

5. The next page will allow the user to select various configuration options, including the type of certificate to request. Choose **Key Recovery Agent** as the Certificate Template.

| ⚠Caution |
| --- |
| A key size of 8192 or larger may take several hours to generate on the client. (Key pairs are always generated by the client CSP, not the CA.) This may slow public key operations on the CA when keys are archived. Key sizes of 2048 are much more reasonable for standard security needs; a key size of 8192 is used only as an example. |

The keys should be marked as exportable. This will enable an administrator (KRA) to export the private keys from the local store of the workstation to a disk or other medium for safe storage. Strong private key protection is also recommended, as this will require a password to be used when accessing the private key.

6. When finished, click **Submit**.

   The Web page will show that the request is being generated. When the key generation and request is complete, if the request had included Enable strong private key protection, a dialog box will appear showing that a protected item is being created (KRA private key).

**Figure 17: Private Key Security Level Dialog Box**

The dialog box will present the opportunity to set the security level on the private key. If the workstation will be used for KRA operations and the private key is to be kept in the protected store, it is recommended that the security level be set to High. This will ensure that access to the private key will require a separate password.

**Figure 18: Private Key Security Level Dialog Box**

7. Click **Next**.

**Figure 19: Private Key Security Level Dialog Box**

8. Choose a password, confirm it, and then click **Finish**.

**Figure 20: Private Key Security Level Dialog Box**

9. Click **OK** to confirm. The Web page will appear and offer a link to install the certificate. If the Certification Authority is configured to set all certificate requests to pending, the user will have to return to the Web link to install the certificate after the CA administrator or a CA Officer has approved the request.

| **Important** |
| --- |
| The user may have to return to the Web link using the same machine that was used to generate the request. This is due to the fact that certificate pending information is stored as a browser cookie. If auto-enrollment is enabled in Active Directory for the user, the user will not be required to actually return to the Web page. The auto-enrollment process will automatically retrieve the certificate for the user when available. For more information about the auto-enrollment process, see the Certificate Auto-Enrollment in Windows XP white paper in Appendix B: Additional Information. |

| **Important** |
| --- |
| If the issuing CA is in a hierarchy, a second option will appear on the Web page to download the entire path (certificate chain). You should choose to download the entire chain. |

The default KRA certificate template requires that the certificate request be pended and not issued automatically. When a certificate request is pended, a CA Officer (Certificate Manager) must manually issue the certificate, assuming that the request was valid. Pending certificate requests can be issued through the Certification Authority MMC and by selecting the pending request node.

The Certificate Manager (or administrator of the server, if role separation is not enabled) can right-click the certificate request

and choose to issue or deny the request. For more information about Certificate Managers, see Configuring Certificate Managers.

After the certificate has been issued, the user (KRA) can return to the Web pages to retrieve the pending request.

The user should select **View the status of a pending certificate request**.

The Web page will display all the pending requests for that user that have been requested from that machine. As mentioned previously, this is managed through Web browser cookies. The user should select the appropriate certificate.

The last page allows the user to install the selected certificate and private key into the local MY store of the user.

#### Configuring the CA to Allow Key Archival

This section details the steps required to configure the CA to allow key archival.

Note: If the Certification Authority is enabled to enforce role separation, only a CA Administrator may configure KRAs on a CA. Role separation is enabled through the registry and only a local server administrator may configure the registry. The easiest way to enable the CA for role separation is to use the certutil.exe command-line tool:

**certutil -setreg ca\RoleSeparationEnabled 1**

It is necessary to stop and start certificate services for the setting to take affect.

| **Note** |
| --- |
| Certutil.exe and other tools may be installed on a Windows XP Professional machine by installing the Administrative Tools (adminpak.msi) that are found in the \i386 directory on all Windows Server 2003 CD-ROM media. |

#### Enabling a Key Recovery Agent

**To enable a KRA**

1. Log on as Administrator of the server or CA Administrator, if role separation is enabled.
2. On the **Administrative Tools** menu, open **Certification Authority**.
3. In the console tree, select the CA.
4. Right-click the CA name, and then click **Properties**.
5. Click the **Recovery Agents** tab.




**Figure 21: Recovery Agents Tab**

6. To enable key archival, click **Archive the key**.

7. By default, the CA will only use one KRA. However, a KRA certificate must first be selected for the CA to begin archival. To select a KRA certificate, click **Add**.

   The system will find valid KRA certificates and display the available KRA certificates.

   KRA certificates are normally published to Active Directory by an Enterprise CA when enrollment occurs. KRA certificates are stored under the KRA container in the Public Key Services branch of the configuration partition in Active Directory. Since a CA may issue multiple KRA certificates, each KRA certificate will be added to the multi-valued userAttribute attribute of the CA object.

8. Select one certificate and click **OK**. You may view the highlighted certificate to ensure that you have selected the intended certificate.




**Figure 22: Recovery Agents Tab**

9. After one or more KRA certificates have been added, click OK to enable key archival on the CA. However, Certificate Services must be stopped and started to enable the use of the selected KRAs. KRA certificates are only processed at service start.

**Configuring Certificate Managers**

To recover the private keys of a user, the CA enforces that a person be a Certificate Manager (if defined) and also holds a private key for a valid KRA certificate. As a best practice, most organizations separate these two roles. By default, the CA Administrator is a Certificate Manager for all users unless otherwise explicitly defined. A KRA is not necessarily a CA Officer (Certificate Manager). They may be segmented as separate roles. A KRA is defined as a person who holds a private key for a valid KRA certificate.

A CA Officer is defined as a Certificate Manager who has the security permission on the CA to Issue and Manage Certificates. The security permissions are configured on a CA using the Security tab on the CA Properties in the Certification Authority MMC snap-in.



**Figure 23: Security Tab**

A CA Administrator can define more specific CA Officer Roles on a Certification Authority by using the Certificate Managers Restrictions tab. By default, any user or group that has the permission to Issue and Manage Certificates is also a CA Officer and can issue, revoke, or export a recovery BLOB for any other user who has been issued a certificate from that CA.




**Figure 24: Certificate Managers Restrictions Tab**

A CA Administrator can define that individual CA Officers have the ability to issue, revoke, and recover certificates for defined sets of user(s) and group(s) by selecting the Restrict certificate managers option. Once the option has been selected, a CA Administrator may define CA Officers' roles as required.

| **Important** |
| --- |
| Once this is complete, it is also necessary to add the machine account of the CA to the Pre W2K Compatible Access Group of every domain in which users will be archived and recovered. This is necessary to provide proper group membership enumeration for role restriction enforcement when performing recovery functions. Certificate Managers will not be able to recover users unless this task is performed first in every user domain. |

**Configuring User Templates**

To enable a user template for key archival, select the template that should enforce key archival and select the Archive subject's encryption private key check box. Once the check box has been selected, the CA will always enforce key archival for that template.




**Figure 25: Request Handling Tab**

| **Note** |
| --- |
| A Windows Server 2003 CA will not allow archival of a key marked for signature only and will reject the request, even if sent programmatically to the CA. |

# Migrating Exchange KMS to Windows Server 2003 CA

Updated: December 6, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following are the summary steps for migrating Exchange 2000 Server KMS to a Windows Server 2003 certification authority.

1. If running Exchange 5.5 KMS, upgrade to Exchange 2000.
2. Configure Windows Server 2003 CA for key archival.
3. Ensure that the certificate is available for database migration.
4. Enable the foreign certificate import option on the CA, if necessary.
5. Export the Exchange KMS database.
6. Import the Exchange database into the CA.

**Important**

Before migrating KMS to a Windows Server 2003 CA, it is important to consider the version 1 CRL that is published by the Exchange KMS for Outlook clients in the Exchange Global Address List (GAL). If KMS is migrated to a Windows Server 2003 CA, the v1 certificates can no longer be revoked and it is recommended that a KMS migration is only performed when all V1 certificates are expired and/or are no longer being issued by KMS. If x.509 version 3 certificates are being issued by KMS with a Windows 2000 CA, the existing CA will need to be maintained to publish CRL(s) until all the original certificates issued by KMS have expired.

**Creating an Export Certificate**

When a KMS migration to a Windows Server 2003 CA is performed, the export file from the KMS must be encrypted with a public key certificate and then subsequently decrypted by the Windows Server 2003 CA. The CA may or may not have an encryption certificate available to be used for this process. It is absolutely critical that an encryption certificate and private key be installed in the machine store (local machine) of the CA to facilitate KMS migration. Since the process runs as SYSTEM, any encryption certificate and private key available in the machine store may be used.

To view the certificates installed in the local machine store, open the Certificates MMC console for the local machine and view the certificates under the Personal store. A Secure Sockets Layer (SSL) or machine authentication certificate will suffice for use in this scenario. The certificate corresponding to the private key that will be used should be manually exported and made available during the KMS migration process. For more information about certificate enrollment and exporting certificates, see the Windows Server 2003 help files. If importing a certificate and key to be used by the CA (*.pfx file), ensure that the certificate is marked for export when importing on the CA. Otherwise, the CA may not be able to use the key and certificate for key import purposes.

**Important**

The export certificate used by the KMS should not have a key size greater than 1024 bits as this may cause errors on import to the Windows Server 2003 CA.

**Important**

A Windows Server 2003 always has an Exchange certificate (encryption certificate) available for the purpose of key archival. Do not attempt to use this certificate for the purpose of migrating the KMS database as it will not be usable by the CA for this purpose.

**Enabling Foreign Certificates Import**

If the KMS contains x.509 version 1 certificates and private keys, and/or if the KMS was not configured to use the same CA with

Windows 2000, the foreign certificate import option must be enabled on the Windows Server 2003 CA.

**Foreign Certificate Import**

By default, a CA does not allow certificates (or keys) to be imported on the CA that were issued by another CA. A CA must be enabled to accept certificates and keys into the database that were issued by a foreign CA. (An Exchange 5.5 KMS issuing version 1 certificates is also considered a foreign CA.)

### To import a foreign CA

1. Run the following command in a command-prompt window on the CA.

   **certutil –setreg ca\KRAFlags +KRAF_ENABLEFOREIGN**

2. Once that has competed, restart the service.

| **Important** |
| --- |
| When foreign certificates are being imported on a CA, the –f switch must be used with certutil to inform the CA that the keys and certificates will be foreign. The command line would be as follows |

**Certutil.exe –f –importKMS [name of import file]**

**Exporting Users' Keys from Exchange 2000 KMS**

| **Important** |
| --- |
| Before an export of data from the KMS occurs, a full backup of the KMS should be performed and validated before continuing. An export of data from a KMS is destructive and will remove the keys from the KMS database. |

| **Important** |
| --- |
| If the KMS or the CA is online when the export occurs, the KMS will attempt to revoke all version 3 certificates that are exported. If this occurs, it is important to re-enroll all users immediately with the Windows Server 2003 CA to allow continued S/MIME encryption operations. Otherwise, take the CA offline, so the KMS export operation will not revoke the existing certificates. |

### To perform the export operation on the KMS

1. Start the Exchange System Manager.




2. Point to the **Advanced Security** node, right-click **Key Manager**, click **All Tasks**, and then click **Export Users**.




3. In the **Key Management Service password** box, type the password (the default password for KMS is "password"), and then click **OK**.




The Exchange KMS Key Export Wizard will start.

4. Click **Next**.




5. Click **Browse** to select the Certificate that will be used to encrypt the export file. This is the certificate file created in the previous section.




6. Browse for the certificate that will be used to encrypt the export file to the CA. This is the certificate created in the previous section. Click Open.




7. On the Encryption Certificate screen, click Next.




8. When this screen appears, use **Windows Explorer** to find and open the certificate that you chose from the screen in step 6. You will need to validate this certificate with the **Exchange KMS Key Export Wizard**.

9. Copy the first eight characters from the Certificate thumbprint field in the certificate chosen to encrypt the KMS export file (Figure 33).




10. Type the first eight characters of the certificate thumbprint in the **Thumbprint** field (Figure 34), and then click **Next**.




11. Type the name of the export file (Figure 35). Do not type in a path, only the file name. It will be saved in the following location by default. This is based on the default installation for Exchange.

    **C:\program files\exchsrvr\KMSDATA**

    This file will not have an extension.

12. Click **Next**.




13. You may select an alphabetic list of users or select by mailbox store, server, or administrative group.




14. In this case, select all of the administrative groups, and then click **Next**.




15. To start the Export process after selecting the users or administrative group(s), click **Next**.




The records will be exported. On average, approximately 100 records will be exported a minute. The actual performance will vary depending on the hardware configuration.




16. When complete, click **Next**.



17. The results will be displayed. Click **Finish**.

> **Note**
> If large numbers of users are exported, KMS may generate multiple export files and split the exported keys across the multiple files. In this case, all export files should be re-imported to the new CA.




The export file will be located in the following folder.

**C:\program files\exchsrvr\KMSDATA**

18. Copy the KMS export file to the server that will accept the import file.

**Importing Users' Keys**

The Windows Server 2003 CA allows not only key archival, but also certificate and key importation to the CA database. Certificate and key importation is important in providing migration services for Exchange KMS as well as for providing migration and escrow operations for certificates that were enrolled using a third-party CA. The Windows Server 2003 CA supports both certificate import as well as key import. Certificate import does not require that key archival be enabled on the CA, but key import does.

**To import users' keys**

1. Log in as the CA Administrator.

2. Open a command prompt window.

3. Change to a directory containing the KMS import file.

4. Run the following command.

   **CertUtil.exe –f –importkms <name of export file>**

   The output will indicate the status of the import process and the number of user keys imported and archived to the CA. The number of imported user keys should match the output from the KMS. The following is a sample successful output.

```
Processing KMS exports from:
    O=microsoft,C=US

KMS export file signature verifies
```

```
Lock box opened, symmetric key successfully decrypted
.................................

Users: 6

Ignored signature certificates: 25
Certificates with keys: 17
Foreign certificates imported: 17
Certificates imported: 17

Keys: 17
Keys archived: 17
CertUtil: -ImportKMS command completed successfully.
```

# Troubleshooting - Key Archival and Management in Windows Server 2003

Updated: December 6, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting KRA Configuration**

This section identifies a number of common mistakes in configuring the KRAs on a CA. The most common error in archiving user private keys on a CA is that the CA is either not configured for key archival or does not have any valid KRA certificate(s) added.

1. The number of recovery agents required by the CA must be less than or equal to the number of available KRA certificates. If an invalid number of recovery agents is entered, the following error message will appear.




**Figure 42: Invalid Number of Recovery Agents Error Message**

2. If an error occurs in trying to validate the KRA certificate when Certificate Services is started, the Recovery Agents tab on the Certification Authority will show that the selected KRA certificate is invalid. This can occur due to a corrupted certificate, corrupted registry entry, deleted certificate, revoked certificate, and so on. Figure 43 shows an example of a corrupted certificate or registry entry on the Recovery Agents tab as shown in the Status column of the selected KRA.




**Figure 43: Invalid KRA Certificate on the Recovery Agents Tab**

3. Similarly, a revoked KRA certificate will also show an error on the Recovery Agents tab when Certificate Services is stopped and started. The error will be displayed in the status column of the KRAs certificates listing.




**Figure 44: Revoked KRA Certificate on the Recovery Agents Tab**

**Loading KRA Certificates**

When certificate services starts on a Certification Authority, the CA attempts to load the KRA(s) defined by the CA Administrator. If the CA is unable to load one or more KRA(s), event log messages will be generated; however, certificate services will continue to start. If the CA is unable to load a KRA(s) successfully as defined by a CA Administrator, the CA will deny all requests for key archival and not issue any certificates that have key archival defined in the certificate template. The following event log messages may appear in the Certification Authority's Application Log when an error occurs in loading KRA certificates. The event log messages indicate that action is required by a CA Administrator to properly configure or reconfigure KRAs.

Event Type: Error

Event Source: CertSvc

Event Category: None

Event ID: 83

Date: 12/20/2000

Time: 8:24:24 AM

User: N/A

Computer: SERVER1

Description:

Certificate Services encountered an error loading key recovery certificates. Requests to archive private keys will not be accepted. The system cannot find the file specified.
0x80070002 (WIN32: 2)

This is a global error that can be caused by one of several conditions.

- The Certification Authority cannot open the KRA store on the local machine.
- The Certification Authority cannot find a corresponding certificate in the KRA store on the local machine that matches the hash of a certificate set in the registry as a KRA.
- The registry has been edited incorrectly or is corrupted.

- The count of KRA certificate hashes in the registry equals zero.
- A certificate hash in the registry corresponds to a certificate in the KRA store that is not a KRA certificate type.
- The KRA certificates are revoked, expired, or invalid.

Event Type: Error

Event Source: CertSvc

Event Category: None

Event ID: 82

Date: 12/27/2000

Time: 9:05:25 AM

User: N/A

Computer: SERVER1

Description:

Certificate Services could not load any valid key recovery certificates. Requests to archive private keys will not be accepted.

This error is usually caused when none of the certificates specified in the user interface (UI) (which corresponds to the registry) is a valid KRA certificate. This event log message is usually accompanied by the previous global event log message.

Event Type: Error

Event Source: CertSvc

Event Category: None

Event ID: 84

Date: 1/24/2003

Time: 08:49:27

User: N/A

Computer: SERVER1

Description:

Certificate Services will not use key recovery certificate 6 because it could not be verified for use as a Key Recovery Agent. CN=User1, OU=Users, DC=nwtraders, DC=com The revocation function was unable to check revocation because the revocation server was offline. 0x80092013 (-2146885613)

This error usually occurs when the CA receives an error when retrieving the CRL to check the status of the KRA certificate.

Event Type: Error

Event Source: CertSvc

Event Category: None

Event ID: 98

Date: 1/24/2003

Time: 08:49:28

User: N/A

Computer: SERVER1

Description:

Certificate Services encountered errors validating configured key recovery certificates. Requests to archive private keys will no longer be accepted.

Event Type: Error

Event Source: CertSvc

Event Category: None

Event ID: 85

Date: 12/27/2000

Time: 9:05:25 AM

User: N/A

Computer: SERVER1

Description:

Certificate Services ignored key recovery certificate 0 because it could not be loaded. Cannot find object or property. 0x80092004 (-2146885628)

This error usually occurs when a specific KRA certificate cannot be found in the KRA store on the local machine of the Certification Authority. Specifically, a KRA certificate has been specified in the UI or registry, and the certificate has been deleted or corrupted in the KRA store. This event log message is usually accompanied by a more global event log message.

**KRA Certificate Status**
When certificate services starts on a Certification Authority, the CA attempts to load the configured KRA(s). The CA must validate the status of each KRA certificate. If the CA is unable to retrieve a current CRL for each KRA certificate, the CA will not be able to load and use that KRA.

The following event log message will be logged in the application event log of the CA.

Event Type: Error

Event Source: CertSvc

Event Category: None

Event ID: 84

Date: 1/12/2001

Time: 11:47:23 AM

User: N/A

Computer: SERVER1

Description:

Certificate Services ignored key recovery certificate 1 because it could not be verified for use as a Key Recovery Agent. CN=User1, OU=Users, DC=nwtraders, DC=com The revocation function was unable to check revocation because the revocation server was offline.

0x80092013 (-2146885613)

**Importing Exchange KMS Export File**
The Windows Server 2003 CA may fail during the importation of the KMS data file if the key size used for the export certificate is greater than 1024 bits in size. The Windows Server 2003 CA may fail with the following message when a key size of greater than 1024 bits is used.

Processing KMS exports from:

CN=Certification Authority, OU=Test, O=Contoso, C=US

KMS export file signature verifies

CertUtil: -ImportKMS command FAILED: 0x80070057 (WIN32: 87)

CertUtil: The parameter is incorrect.

**User Enrollment Errors**
A user certificate request for a template that requires key archival will be denied if one of the following conditions exists.

- No KRA has been defined on the CA.
- No KRA can be successfully loaded. (KRA certificates are revoked, expired, and so on.)
- The minimum number of KRA certificates defined by the CA Administrator cannot be loaded.

If the user enrolls through a Web page, the following text will display on the Web page.

Your request failed. An error occurred while the server was processing your request.

Contact your administrator for further assistance.

**Request Mode**:

newreq - New Request **Disposition**:

(never set) **Disposition message**:

(none) **Result**:

Cannot archive private key. The certification authority is not configured for key archival. 0x8009400a (-2146877430) **COM Error Info**:

CCertRequest::Submit Cannot archive private key. The certification authority is not configured for key archival. 0x8009400a (-2146877430) **LastStatus**:

Cannot archive private key. The certification authority is not configured for key archival. 0x8009400a (-2146877430) **Suggested Cause**:

No suggestions.

If enrolling through the MMC, the following error will be displayed.

**Figure 45: Incorrect Certificate Request Error Message**

The CA will also log the following error to the application event log of the CA.

Event Type: Error

Event Source: CertSvc

Event Category: None

Event ID: 21

Date: 1/12/2001

Time: 4:23:39 PM

User: N/A

Computer: SERVER1

Description:

Certificate Services could not process request 16094 due to an error: Cannot archive private key. No valid key recovery agent is defined for this certification authority. 0x8009400b (-2146877429). The request was for NWTRADERS\User1.

If the CA is unable to retrieve a current CRL for the CA itself or any of its parent CA(s), it will be unable to issue a certificate when a user submits a request. If the CA does not have a valid CRL for itself, the following error message will be displayed in the application event log of the CA.

Event Type: Warning

Event Source: CertSvc

Event Category: None

Event ID: 53

Date: 1/6/2001

Time: 11:24:05 AM

User: N/A

Computer: SERVER1

Description:

Certificate Services denied request 1471 because the revocation function was unable to check revocation because the revocation server was offline. 0x80092013 (-2146885613). The request was for CN=user1, OU="Test", O="NWTraders", L=Redmond, S=WA, C=US, E=user1@nwtraders.com. Additional information: Denied by Policy Module

**Certificate Template Not Supported by the CA**
If a user tries to enroll with a CA for a template that is not supported by that CA, the following event log message will be entered in the CA application event log.

Event Type: Warning

Event Source: CertSvc

Event Category: None

Event ID: 53

Date: 1/16/2001

Time: 2:07:02 PM

User: N/A

Computer: SERVER1

Description:

Certificate Services denied request 8 because the requested certificate template is not supported by this CA. 0x80094800 (-2146875392). The request was for NWTRADERS\Administrator. Additional information: Denied by Policy Module The request was for certificate template (1.3.6.1.4.1.311.21.8.4144336743.1329436349.2065260953.3989445610.1.27) that is not supported by the Certificate Services policy.

**Client CSP Does Not Permit Key Export**
For the client enrollment process to generate and send a private key to the CA, the key must be marked as exportable when the key is generated. If the certificate template is not set to allow key exportable or if the third-party CSP (if applicable) does not support exportable keys, enrollment will fail and the enrollment wizard will return an error that the key is not exportable. Third-party CSPs may report varying errors, such as “catastrophic failure”, when this occurs. If a Windows 2000 or Windows Millennium Edition client performs enrollment with key archival, the following error may appear if the key is not marked for export.

0x80090009 - NTE_BAD_FLAGS

**Note**

If the CSP supports the one-time flag for key archival, known as (CRYPT_ARCHIVABLE), the key export flag is not required. The Microsoft default software CSPs support this flag. However, Windows 2000 and Windows Millennium Edition clients do not support this flag and must allow the key to be exported for enrollment to work with key archival.

**Certification Authority CSP Not Supported for Key Archival Functions**

If a software or hardware CSP is not capable of performing the symmetric and public operations for encrypting the private key(s) of users in the CA database, the CA will log an event in the application event log:

Event Type: Warning

Event Source: CertSvc

Event Category: None

Event ID: 86

Date: 12/27/2001

Time: 8:13:54 AM

User: N/A

Computer: NORTHWIND5

Description:

Certificate Services could not use the provider specified in the registry for encryption keys. The keyset is not defined. 0x80090019 (-2146893799)

For more information, see the Help and Support Center at http://go.microsoft.com/fwlink/events.asp

Event Type: Warning

Event Source: CertSvc

Event Category: None

Event ID: 88

Date: 12/27/2001

Time: 8:13:54 AM

User: N/A

Computer: NORTHWIND5

Description:

Certificate Services switched to the default provider for encryption keys.

For more information, see the Help and Support Center at http://go.microsoft.com/fwlink/events.asp

To verify which CSP the CA is using for encryption operations associated with key archival, run the following command from the CA.

**Certutil –getreg ca\EncryptionCSP\Provider**

**Certificate and Key Import Issues**

If the CA has not been configured for key archival or if the CA has not been configured for foreign key import, the following error will occur when attempting to import a KMS export file or a foreign key import. Note that the keys were not archived in the following message.

```
Processing KMS exports from:
    O=microsoft, C=US
KMS export file signature verifies
Lock box opened, symmetric key successfully decrypted
CertUtil: Invalid Signature.
```

```
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.
CertUtil: Invalid Signature.


Users: 6

Ignored signature certificates: 25
Certificates with keys: 17
Certificates not imported: 17

Keys: 17
Keys not archived: 17
CertUtil: -ImportKMS command completed successfully.
```

**Troubleshooting Key Recovery Issues**

**Unable to Recover User**

If a CA performing key archival is also enabled for role separation with specific Certificate Manager restrictions, a Certificate Manager may not be able to recover a user certificate until the machine account of the CA has been added to the Pre W2K Compatible Access Group of the domain in which the recover user belongs. This is a necessary requirement for the CA to enumerate the group memberships of Certificate Managers and recovered users to ensure that proper restrictions are enforced.

**Missing KRA Certificate in the CA Registry**

If one of the recipient KRA certificates from the HKEY_LOCAL_MACHINE KRA certificate store on the Certification Authority is deleted, key recovery tools, such as certutil –getkey, will fail because the server cannot find the KRA certificate to include in the recovery BLOB. The following error message will be displayed when this error occurs.

```
certutil -getkey "1b 4a b7 1e 00 00 00 00 00 1d"
Querying server1.nwtraders.com\CA1............

"server1.nwtraders.com\CA1"
   1b4ab71e00000000001d  CN="Users

Administrator"
CertUtil: -GetKey command FAILED: 0x80092004 (-2146885628)
CertUtil: Cannot find object or property
```

Note that the KRA certificate must be available in the registry on the CA, not the machine where the recovery tool(s) are used.

# Appendix A: Certificate Request Structure

Updated: December 6, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This appendix provides additional detailed information about the key archival process regarding the certificate request structure.

**ASN.1 Structure**
A certificate request for key archival to the CA is a CMC Full PKIRequest message as specified in RFC 2797. The ASN.1 structure used by the Windows Server 2003 CA is demonstrated in Figure 46.




**Figure 46: CMC Request Message**

**Understanding the PKCS #7 Message Content Structure**
The first section of the CMC message contains a PKCS #7 message that has the relevant elements for generating a certificate request.

**Understanding the controlSequence TaggedAttribute Element**
The TaggedAttribute element in the message contains the following information.

- Extensions—The Extensions section of the TaggedAttribute element contains the following extensions.
    - Application Policies
    - Template Information
    - Key Usage
    - Enhanced Key Usage
- Attributes—The Attributes section of the TaggedAttribute element contains the following data.
    - Common Name
    - Template Name to be used
    - Hash of the encrypted private key BLOB
- Other request information

**Understanding the reqSequence TaggedRequest Element**
The reqSequence TaggedRequest element contains a nested PKCS #10 message. This message contains the user's public key in addition to other information relevant for generating the certificate.

**Understanding the cmsSequence TaggedContentInfo Element**
The cmsSequence TaggedContentInfo element can contain nested PKCS #7 and CMC messages. In a standard archival request, this element is not used.

**Understanding the otherMsgSequence OtherMsg Element**
Not Used

**Understanding the Signatures Structure**
The signatures section of the CMC message contains one or more signatures used to sign the request. The following is an example of the signatures section.

```
Signer Count: 1

Signer Info[0]:
Signature matches request Public Key
CMSG_SIGNER_INFO_CMS_VERSION(3)
CERT_ID_KEY_IDENTIFIER(2)
    0000  81 92 56 3a c4 31 f8 82  0c 54 c9 d0 98 4f d8 c5
    0010  34 63 9e cc
Hash Algorithm:
    Algorithm ObjectId: 1.3.14.3.2.26 sha1
    Algorithm Parameters: NULL
Encrypted Hash Algorithm:
    Algorithm ObjectId: 1.2.840.113549.1.1.1 RSA
    Algorithm Parameters: NULL
Encrypted Hash:
    0000  c1 ae 90 a7 a3 0b 52 66  ea c4 d0 04 17 2e 94 95
    0010  14 20 06 ...
```

**Understanding the Authenticated Attributes Structure**
The authenticated attributes section contains additional authenticated attributes, such as Content Type, Message Digest, and Client Information. The following is an example of the authenticated attributes section.

```
Authenticated Attributes[0]:
  3 attributes:
  Attribute[0]: 1.2.840.113549.1.9.3 (Content Type)
    Value[0][0]:
    Unknown Attribute type
    1.3.6.1.5.5.7.12.2 CMC Data
  Attribute[1]: 1.2.840.113549.1.9.4 (Message Digest)
    Value[1][0]:
    Unknown Attribute type
    Message Digest:
        5e 1f 0f f0 28 a4 fe 91 0d c2 2f 1a 18 78 7e 2e 10 7f 17 39
  Attribute[2]: 1.3.6.1.4.1.311.21.20 (Client Information)
    Value[2][0]:
    Unknown Attribute type
    Client Id: = 1
    XECI_XENROLL -- 1
    User: CONTOSO0\avibm
    Machine: dcross-stress.contoso.com
    Process: certreq
```

**Understanding the Unauthenticated Attributes Structure**
The unauthenticated attributes section contains the encrypted private key. The private key is contained in an enveloped PKCS #7 message that is encrypted to the CA's exchange key. Since this is an unauthenticated attribute, the SHA1 hash of the PKCS #7 message is included as one of the attributes of the controlSequence TaggedAttribute attributes.

The following is an example of the unauthenticated attributes section.

```
Unauthenticated Attributes[0]:
  1 attributes:

  Attribute[0]: 1.3.6.1.4.1.311.21.13 (Encrypted Private Key)
    Value[0][0]:
    Unknown Attribute type
================ Begin Nesting Level 1 ================
```

```
PKCS7 Message:
  CMSG_ENVELOPED(3)
  CMSG_ENVELOPED_DATA_PKCS_1_5_VERSION(0)
  Content Type: 1.2.840.113549.1.7.1 PKCS 7 Data

PKCS7 Message Content:
0000    d4 a6 31 b6 5a ee 62 90  cc 17 b1 7a 6a 0d 40 9a
..1.Z.b....zj.@.
0010    33 fd 11 14 0b ae 12 bd  3b 32 b8 73 af cc 1b 76
3.......;2.s...v ...
```

**Performing Binary Export for a Request**

To view and decode a CMS key archival request from a Windows Server 2003 CA, it is necessary to do a binary export directly from the CA database. A binary export can be easily achieved through the Certification Authority MMC snap-in or by using the certutil.exe command-line tool.

**Binary Request Export Using the Certification Authority MMC Snap-In Walkthrough**

**To export a binary request using the Certification Authority MMC Snap-in**

1. Log on to the CA machine using a CA Administrator account.

2. Open the **Certification Authority** MMC snap-in.

3. Click the **Issued Certificates** folder.

4. If the binary request column has not been previously added to the database view, it must be added to support a binary request export. To add a column to the view, click **View** on the menu bar, and then select the **Add/Remove Columns** menu option.

5. In the **Add/Remove Columns** dialog box, select the **Binary Request** field in the **Available Columns** list box on the left.

6. Click **Add**, and then click **OK**.

Next, a binary request can be exported.

1. Select a request from the issued certificates view, and then click the **Action** menu.

2. Select **Export Binary Data** on the **All Tasks** menu.

3. In the **Export Binary Data** dialog box, choose **Binary Request** as the column you want to export.

4. Click **OK**.

The data will be exported into ASCII format that can be opened in Notepad using notepad.exe.

| **Note** |
| --- |
| Following the previous steps will generate a dump of the certificate archival request only; it does not include the private key material. To dump a full certificate archival request including the private key material, follow the command-line option. |

**Binary Request Export Using the CertUtil.exe Command-Line Tool Walkthrough**

To use the certutil.exe to view the certificate request including private key material, a request file has to be generated first.

**To generate a request file**

1. Run Notepad.exe.

2. Paste the following certificate request information into Notepad.

   [Version]

Signature= "$Windows NT$"

[NewRequest]

Subject = "CN=Test Subject"

KeySpec = 1

Exportable = FALSE

PrivateKeyArchive = TRUE

[RequestAttributes]

CertificateTemplate = EFS

| **Note** |
| --- |
| Make sure that the CA is configured for key archival before starting this process. In this example, the EFS template is used; this should be changed to an existing certificate template that allows private key archival. |

3. Save the file as **CertificateRequest.inf**, and then close **Notepad**.

4. Open the command-line window.

5. Type the following command.

   **Certreq –new CertificateRequest.inf CertificateRequest.req**

   **Notes**:

   - This command will prompt you to select the CA to fetch the CA exchange certificate from, and to encrypt the private key to.

   - This command will write the request to a file named by the last argument on the command line: CertificateRequest.req.

   - To avoid using the CA selection dialog, you can specify the CA via -config **CAMachineDNSName\CACertCommonName** before or after the –new option.

6. Type the following command.

   **certreq -submit CertificateRequest.req KeyArchival.cer KeyArchival. p7b KeyArchival.rsp**

   This command will prompt you to select the CA to submit the request to.

   **Notes**:

   - This command will write the newly issued certificate, a PKCS7 containing only the issued certificate and chain, and the full CMC response to files named by the last three arguments on the command line: KeyArchival.cer, KeyArchival.p7b, and KeyArchival.rsp, respectively.

   - To avoid the U/I, you can specify the CA via -config **CAMachineDNSName\CACertCommonName** before or after the **–submit**.

7. Type the following command.

   **certreq -accept KeyArchival.rsp**

   This command verifies the response, installs the certificate, and associates it with the private key.

8. Type the following command.

   **Certutil –privatekey –dump CertificateRequest.req >CertificateRequest.txt**

   This command will generate a dump of the certificate archival request into the CertificateRequest.txt file.

9. Type the following command.

   **Certutil –privatekey –dump KeyArchival.rsp >CertificateResponse.txt**

   This command will generate a dump of the certificate archival response into the CertificateResponse.txt file.

For non-Windows Server 2003 clients or servers enrolling to a Windows Server 2003 CA, the format of the request may be different. The reason is that non-Windows Server 2003 platforms may not support CMC data structures and, therefore, may not be able to encode the request information inside a PKIData object. Instead, the request information may be inside the Data body but not encoded as a PKIData object.

| **Note** |
| --- |
| certreq.exe and other tools may be installed on a Windows Server 2003 Professional machine by installing the Administrative Tools (adminpak.msi) that are located in the \i386 directory on all Windows Server 2003 CD-ROM media. |

**CMC Request and Response Examples**

Request:

```
SEQUENCE :
 OBJECT IDENTIFIER :  signedData [1.2.840.113549.1.7.2]
 CONTEXT SPECIFIC (0) :
    SEQUENCE :
       INTEGER : 3
       SET :
          SEQUENCE :
             OBJECT IDENTIFIER :  sha1 [1.3.14.3.2.26]
             NULL :
       SEQUENCE :
          OBJECT IDENTIFIER :  [1.3.6.1.5.5.7.12.2]
          CONTEXT SPECIFIC (0) :
             OCTET STRING :
                SEQUENCE :
                   SEQUENCE :
                      SEQUENCE :
                         INTEGER : 2
                         OBJECT IDENTIFIER :  [1.3.6.1.5.5.7.7.8]
                            SET :
                               SEQUENCE :
                                  INTEGER : 0
                               SEQUENCE :
                                  INTEGER : 1
                               SEQUENCE :
                                  SEQUENCE :
                                     OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.10]
                                     OCTET STRING :
                                        SEQUENCE :
                                           SEQUENCE :
                                              OBJECT IDENTIFIER :  encryptedFileSystem
[1.3.6.1.4.1.311.10.3.4]
                                  SEQUENCE :
                                     OBJECT IDENTIFIER :  keyUsage [2.5.29.15]
                                     OCTET STRING :
```

```
                                    BIT STRING UnusedBits:5 :
                                          20
                                 SEQUENCE :
                                    OBJECT IDENTIFIER :  extKeyUsage [2.5.29.37]
                                       OCTET STRING :
                                          SEQUENCE :
                                             OBJECT IDENTIFIER :  encryptedFileSystem
[1.3.6.1.4.1.311.10.3.4]
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.7]
                                       OCTET STRING :
                                          SEQUENCE :
                                             OBJECT IDENTIFIER :
 [1.3.6.1.4.1.311.21.8.4014942.3497959.5914804.3829722.12246394.103.3066650.1537810]
                                             INTEGER : 100
                                             INTEGER : 2
                           SEQUENCE :
                              INTEGER : 3
                              OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.10.10.1]
                              SET :
                                 SEQUENCE :
                                    INTEGER : 0
                                    SEQUENCE :
                                       INTEGER : 1
                                    SET :
                                       SEQUENCE :
                                          OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.21]
                                          SET :
                                             OCTET STRING :
                                                9231E6C0B87445190EA2CA934B2807FF799
                                                3C59F
                           SEQUENCE :
                              INTEGER : 4
                              OBJECT IDENTIFIER :  [1.3.6.1.5.5.7.7.18]
                              SET :
                                 OCTET STRING :
                                    436572746966696361746554656D706C6174653D4172636
                                    869766554657374426173696345465326
                        SEQUENCE :
                           CONTEXT SPECIFIC (0) :
                              INTEGER : 1
                              SEQUENCE :
                                 SEQUENCE :
                                    INTEGER : 0
                                    SEQUENCE :
                                       SET :
                                          SEQUENCE :
                                             OBJECT IDENTIFIER :  commonName [2.5.4.3]
                                             PRINTABLE STRING :
'Test Subject'
                                    SEQUENCE :
                                       SEQUENCE :
                                          OBJECT IDENTIFIER :  rsaEncryption
[1.2.840.113549.1.1.1]
                                          NULL :
                                       BIT STRING UnusedBits:0 :
                                          SEQUENCE :
                                             INTEGER :
                                                00DAFF7C6859557C698CDA4598222E8E90E
                                                EB481889531E9F67F10C081F2545B060BE7
                                                714E755325AC710774764DCA8120C6BEB7B
                                                6EF74B0260EDD56DD299B242A94EE83C420
                                                AC7FF0E694122E26EF67670782223C4E8D8
                                                12C98047F24E10CF6A26FEBEEB826638924
                                                F36B697CEA02EFC4CA0D108CB85047266AD
                                                27DE582D181A1
                                             INTEGER : 65537
                                    CONTEXT SPECIFIC (0) :
                                       SEQUENCE :
                                          OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.13.2.3]
                                          SET :
                                             IA5 STRING :
```

```
                                        '5.2.3790.2'
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.20]
                                       SET :
                                          SEQUENCE :
                                             INTEGER : 1
                                             UTF8 STRING :
                                                'dcross-stress.contoso.com'
                                             UTF8 STRING :
                                                'CONTOSO0\avibm'
                                             UTF8 STRING :
                                                'certreq'
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.13.2.2]
                                       SET :
                                          SEQUENCE :
                                             INTEGER : 1
                                             BMP STRING :
                                                'Microsoft Strong Cryptographic P'
                                                'rovider'
                                             BIT STRING UnusedBits:0 :
                                                0000000000000000000000000000000
                                                0000000000000000000000000000000
                                                0000000000000000000000000000000
                                                0000000000000000000000000000000
                                                0000000000000000000000000000000
                                                0000000000000000000000000000000
                                                0000000000000000000000000000000
                                                0000000000000000000000000000000
                                                0000000000000000
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  extensionReq
[1.2.840.113549.1.9.14]
                                       SET :
                                          SEQUENCE :
                                             SEQUENCE :
                                                OBJECT IDENTIFIER :  sMIMECapabilities
[1.2.840.113549.1.9.15]
                                                OCTET STRING :
                                                   SEQUENCE :
                                                      SEQUENCE :
                                                         OBJECT IDENTIFIER :  rc2CBC
[1.2.840.113549.3.2]
                                                         INTEGER : 128
                                                      SEQUENCE :
                                                         OBJECT IDENTIFIER :  rc4
[1.2.840.113549.3.4]
                                                         INTEGER : 128
                                                      SEQUENCE :
                                                         OBJECT IDENTIFIER :  desCBC
[1.3.14.3.2.7]
                                                      SEQUENCE :
                                                         OBJECT IDENTIFIER :  DES-
EDE3-CBC [1.2.840.113549.3.7]
                                             SEQUENCE :
                                                OBJECT IDENTIFIER
:  subjectKeyIdentifier [2.5.29.14]
                                                OCTET STRING :
                                                   OCTET STRING :
                                                      8192563AC431F8820C54C9D098
                                                      4FD8C534639ECC
                                             SEQUENCE :
                                                OBJECT IDENTIFIER
:  [1.3.6.1.4.1.311.21.10]
                                                OCTET STRING :
                                                   SEQUENCE :
                                                      SEQUENCE :
                                                         OBJECT IDENTIFIER
:  encryptedFileSystem [1.3.6.1.4.1.311.10.3.4]
                                             SEQUENCE :
                                                OBJECT IDENTIFIER :  keyUsage
[2.5.29.15]
```

```
                                                   OCTET STRING :
                                                      BIT STRING UnusedBits:5 :
                                                         20
                                                SEQUENCE :
                                                   OBJECT IDENTIFIER :  extKeyUsage
[2.5.29.37]
                                                   OCTET STRING :
                                                      SEQUENCE :
                                                         OBJECT IDENTIFIER
:  encryptedFileSystem [1.3.6.1.4.1.311.10.3.4]
                                                SEQUENCE :
                                                   OBJECT IDENTIFIER
:  [1.3.6.1.4.1.311.21.7]
                                                   OCTET STRING :
                                                      SEQUENCE :
                                                         OBJECT IDENTIFIER :
 [1.3.6.1.4.1.311.21.8.4014942.3497959.5914804.3829722.12246394.103.3066650.1537810]
                                                         INTEGER : 100
                                                         INTEGER : 2
                              SEQUENCE :
                                 OBJECT IDENTIFIER :  sha1withRSAEncryption
[1.2.840.113549.1.1.5]
                                 NULL :
                              BIT STRING UnusedBits:0 :
                                 31E945A575155D8F91E972DB26A52C8FAE16D7F5074365D
                                 C2E585C8718AB09A4FBB67D8A78A63C76B14482A1DEDCAA
                                 5B234035F3CFFABCAF3DEC24C5944ACE46A1BAFE857F310
                                 7C21105C817FA88C0CCB23B88D2684327B40CB99E9A059F
                                 3B95BAC6423740CA1B46B4DC58664863325004DCA2857C2
                                 2B4117942CC7D39E86900
                    SEQUENCE :
                    SEQUENCE :
          SET :
             SEQUENCE :
                INTEGER : 3
                CONTEXT SPECIFIC (0) :
                   8192563AC431F8820C54C9D0984FD8C534639ECC
                SEQUENCE :
                   OBJECT IDENTIFIER :  sha1 [1.3.14.3.2.26]
                   NULL :
                CONTEXT SPECIFIC (0) :
                   SEQUENCE :
                      OBJECT IDENTIFIER :  contentType [1.2.840.113549.1.9.3]
                      SET :
                         OBJECT IDENTIFIER :  [1.3.6.1.5.5.7.12.2]
                   SEQUENCE :
                      OBJECT IDENTIFIER :  messageDigest [1.2.840.113549.1.9.4]
                      SET :
                         OCTET STRING :
                            5E1F0FF028A4FE910DC22F1A18787E2E107F1739
                   SEQUENCE :
                      OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.20]
                      SET :
                         SEQUENCE :
                            INTEGER : 1
                            UTF8 STRING :
                               'dcross-stress.contoso.com'
                            UTF8 STRING : 'CONTOSO0\avibm'
                            UTF8 STRING : 'certreq'
                SEQUENCE :
                   OBJECT IDENTIFIER :  rsaEncryption [1.2.840.113549.1.1.1]
                   NULL :
                OCTET STRING :
                   C1AE90A7A30B5266EAC4D004172E949514200653AA5EA3C2BF17C7731DA8EB
                   1A635CE1DC4F5AD9FB44EF2D9E8C9F961800DBEBC1ADE14E0459A8B46880DF
                   01A177FC9B02B89113638F3A6A3B3ED0765BD16B905D6BCB404F65E79AAB12
                   97F2F9F52D68D13373D41D510D97A954800368F8DEDEE13D8635EBF4364512
                   17407F1A
                CONTEXT SPECIFIC (1) :
                   SEQUENCE :
                      OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.13]
                      SET :
```

```
                        SEQUENCE :
                           OBJECT IDENTIFIER :  envelopedData [1.2.840.113549.1.7.3]
                           CONTEXT SPECIFIC (0) :
                              SEQUENCE :
                                 INTEGER : 0
                                 SET :
                                    SEQUENCE :
                                       INTEGER : 0
                                       SEQUENCE :
                                          SEQUENCE :
                                             SET :
                                                SEQUENCE :
                                                   OBJECT IDENTIFIER
:  domainComponent [0.9.2342.19200300.100.1.25]
                                                   IA5 STRING :
                                                      'com'
                                             SET :
                                                SEQUENCE :
                                                   OBJECT IDENTIFIER
:  domainComponent [0.9.2342.19200300.100.1.25]
                                                   IA5 STRING :
                                                      'contoso'
                                             SET :
                                                SEQUENCE :
                                                   OBJECT IDENTIFIER :  commonName
[2.5.4.3]
                                                   PRINTABLE STRING :
                                                      'TestEnrollment'
                                          INTEGER :
                                             18D0100D00000000005B
                                       SEQUENCE :
                                          OBJECT IDENTIFIER :  rsaEncryption
[1.2.840.113549.1.1.1]
                                          NULL :
                                       OCTET STRING :
                                          A41AAE9CDA66F283D6D4BC829D2F58BCECFD3F
                                          5A57EC8AE14021179AE5F93F03AE90747FD300
                                          4573ED78F802E02AB3C6ADEDEAA367069DA399
                                          8E1D2D34ABEEFF0F8DE2CB76078C56D883BD94
                                          D7CE9C5CD75F5E3F442A467F74E07C5A434E4A
                                          F1BDD6EC493F3A870764B6CC6446FA5D674255
                                          D93F248DE23E0D96902C79901800
                                 SEQUENCE :
                                    OBJECT IDENTIFIER :  data [1.2.840.113549.1.7.1]
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  DES-EDE3-CBC
[1.2.840.113549.3.7]
                                       OCTET STRING :
                                          06003B8D3EB4B44C
                                    CONTEXT SPECIFIC (0) :
                                       D4A631B65AEE6290CC17B17A6A0D409A33FD11140
                                       BAE12BD3B32B873AFCC1B76A4022D0FB2B50E431A
                                       1E48C8D45865EC5B730D7357D61C9495235143381
                                       19CDBF34C5455B73C9FF38AEFBC4E32DD8145647B
                                       46B0B4A60D29D062051F116C6BA49253D4590944A
                                       7CCB70F43E7E850B34DE55074B3C5FF5AE1C5A18C
                                       6BC271D1F2BC3FBBE19558252C894110CC801292D
                                       63DA1485BDB957270E6C1A38FE33D672EA3E8D031
                                       CD7BCFEF5C738818DCC43A6F76F3EC81701C561DF
                                       9FA6032C47236D9A16973BDF6A033F4925CC5B491
                                       C00C635C65F744C8FEBE19B1EDD2172AE3A7CFB70
                                       87A6BCAE7BB52BCEEF412889C4A45ACAE0ACC0E43
                                       A14C7AA34FB4B4C49360ADD0C65D1494B792E04D7
                                       8D43C2EDB79974B5C08C87E0C72767C26A2EBF6F0
                                       E273269D139F2D6F451301944B76218D9BD4C5931
                                       50C79FA5DA1AF1383E5342EC2F5318E2404774345
                                       B82A0CB4EE26FC0D59A1D18EDBBEFF6135675D014
                                       293470B301CC59387C4E627E1F6B038A158A927B9
                                       160387104BFC5466B7FB4107DF02D136E076F2CAD
                                       94718ADD9F93C0D376A80A6E3796C6236888E6517
                                       1D36A0F3BFAA8B8E44FC8DA426F3F19128A910D83
                                       71A7D68CDFEFCA0BBF32888D8AC679975AE43BB6D
```

```
                                        209D61F82EEA2463616E905177E929CFD3D85C8ED
                                        8ED1EDCECA01CA1580960E87D57591817C863FE33
                                        757F527DC7C6457ED5CEDC3BE1597A05BFB10A145
                                        522C98AF266A992CC607434D3421D57A80195D052
                                        557AE89652193B840FC27CB343C2C242445453E78
                                        9E6E397DFC84363B4EAE801DF1BE2993D1AF13256
                                        A1390C4B7D51127CC55FF0B1184D4E87967961E86
                                        B722E1048C0
```


Response:


```
SEQUENCE :
    OBJECT IDENTIFIER :  signedData [1.2.840.113549.1.7.2]
    CONTEXT SPECIFIC (0) :
       SEQUENCE :
          INTEGER : 3
          SET :
             SEQUENCE :
                OBJECT IDENTIFIER :  sha1 [1.3.14.3.2.26]
                NULL :
          SEQUENCE :
             OBJECT IDENTIFIER :  [1.3.6.1.5.5.7.12.3]
             CONTEXT SPECIFIC (0) :
                OCTET STRING :
                   SEQUENCE :
                      SEQUENCE :
                         SEQUENCE :
                            INTEGER : 1
                            OBJECT IDENTIFIER :  [1.3.6.1.5.5.7.7.1]
                            SET :
                               SEQUENCE :
                                  INTEGER : 0
                                  SEQUENCE :
                                     INTEGER : 1
                                  UTF8 STRING : 'Issued'
                         SEQUENCE :
                            INTEGER : 2
                            OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.10.10.1]
                            SET :
                               SEQUENCE :
                                  INTEGER : 0
                                  SEQUENCE :
                                     INTEGER : 1
                                  SET :
                                     SEQUENCE :
                                        OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.17]
                                        SET :
                                           OCTET STRING :
                                              DE73D68A50323310A01EEDDF66188213DC9
                                              CD490
                                     SEQUENCE :
                                        OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.21]
                                        SET :
                                           OCTET STRING :
                                              9231E6C0B87445190EA2CA934B2807FF799
                                              3C59F
                      SEQUENCE :
                      SEQUENCE :
          CONTEXT SPECIFIC (0) :
             SEQUENCE :
                SEQUENCE :
                   CONTEXT SPECIFIC (0) :
                      INTEGER : 2
                   INTEGER :
                      172B1FB96BBBF2BA49A64EBEA41833EF
                   SEQUENCE :
                      OBJECT IDENTIFIER :  sha1withRSAEncryption [1.2.840.113549.1.1.5]
                      NULL :
                   SEQUENCE :
                      SET :
                         SEQUENCE :
                            OBJECT IDENTIFIER :  domainComponent
```

```
[0.9.2342.19200300.100.1.25]
                            IA5 STRING : 'com'
                      SET :
                         SEQUENCE :
                            OBJECT IDENTIFIER :  domainComponent
[0.9.2342.19200300.100.1.25]
                            IA5 STRING : 'contoso'
                      SET :
                         SEQUENCE :
                            OBJECT IDENTIFIER :  commonName [2.5.4.3]
                            PRINTABLE STRING :
                               'TestEnrollment'
                   SEQUENCE :
                      UTC TIME : '040210162354Z'
                      UTC TIME : '090210162738Z'
                   SEQUENCE :
                      SET :
                         SEQUENCE :
                            OBJECT IDENTIFIER :  domainComponent
[0.9.2342.19200300.100.1.25]
                            IA5 STRING : 'com'
                      SET :
                         SEQUENCE :
                            OBJECT IDENTIFIER :  domainComponent
[0.9.2342.19200300.100.1.25]
                            IA5 STRING : 'contoso'
                      SET :
                         SEQUENCE :
                            OBJECT IDENTIFIER :  commonName [2.5.4.3]
                            PRINTABLE STRING :
                               'TestEnrollment'
                   SEQUENCE :
                      SEQUENCE :
                         OBJECT IDENTIFIER :  rsaEncryption [1.2.840.113549.1.1.1]
                         NULL :
                      BIT STRING UnusedBits:0 :
                         SEQUENCE :
                            INTEGER :
                               00E23136361B94412ABD67C376C6AC882B50F45D9AD28719C1
                               5B0F3125CB352E19F5A381A33FF2971CC4702747BD94C3EE93
                               75493C1A48F5174BE1F8135CCFB641F3EE6042C4771E8E176A
                               7B65E49E407903072C28E2CC92153454664630FDA3CC70A805
                               086B586592AF45BFFE5CC82DCF1ED622DD9BE4ECF64D635600
                               9338C96F7D2EF77447F3ACD2AFC9C76EBC7A77DAAA9245A0EE
                               0398D041B37DD78BD77C46D84A808AECDB88EC4319B1E6ADB9
                               19053A84D3403163003EE696F65E0A55F5EA7A4955870D451E
                               E4A0AB684EE6ED503437A3F4388DC96A00A9F7D26E3527B3D0
                               F657EFB8E431B24A97ADBD1475DAF545B9754856200E640E42
                               CA8BF78614A953
                            INTEGER : 65537
                   CONTEXT SPECIFIC (3) :
                      SEQUENCE :
                         SEQUENCE :
                            OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.20.2]
                            OCTET STRING :
                               BMP STRING : 'CA'
                         SEQUENCE :
                            OBJECT IDENTIFIER :  keyUsage [2.5.29.15]
                            OCTET STRING :
                               BIT STRING UnusedBits:1 :
                                  86
                         SEQUENCE :
                            OBJECT IDENTIFIER :  basicConstraints [2.5.29.19]
                            BOOLEAN : 'FF'
                            OCTET STRING :
                               SEQUENCE :
                                  BOOLEAN : 'FF'
                         SEQUENCE :
                            OBJECT IDENTIFIER :  subjectKeyIdentifier [2.5.29.14]
                            OCTET STRING :
                               OCTET STRING :
                                  10C8E49879236E65350924C24EFB074EFB5F4AA0
```

```
                        SEQUENCE :
                           OBJECT IDENTIFIER :  cRLDistributionPoints [2.5.29.31]
                           OCTET STRING :
                              SEQUENCE :
                                 SEQUENCE :
                                    CONTEXT SPECIFIC (0) :
                                       CONTEXT SPECIFIC (0) :
                                          CONTEXT SPECIFIC (6) :
                                             'ldap:///CN=TestEnrollment,CN=dcross'
                                             '-stress,CN=CDP,CN=Public%20Key%20Se'
                                             'rvices,CN=Services,CN=Configuration'
                                             ',DC=contoso,DC=com?certificateRevoc'
                                             'ationList?base?objectClass=cRLDistr'
                                             'ibutionPoint'
                                          CONTEXT SPECIFIC (6) :
                                             'http://dcross-stress.contoso.com/Ce'
                                             'rtEnroll/TestEnrollment.crl'
                        SEQUENCE :
                           OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.1]
                           OCTET STRING :
                              INTEGER : 0
                 SEQUENCE :
                    OBJECT IDENTIFIER :  sha1withRSAEncryption [1.2.840.113549.1.1.5]
                    NULL :
                 BIT STRING UnusedBits:0 :
                    CA9E6760A8DFB0D213E90D7450B5C7A7C5C920760D01EB45E4F46A23780841
                    40EDE1A37BA123934C06A39F9638F86C9A50258E43E71DE44239A20DFD6EAE
                    C636F6B50C964EF23A72B349F35530A96CC99AF8937F22F684AF5E39E64C90
                    F49C0D87621BBB13DE9FAF84609C26C5ECEB37F479CAEF826D36C19FD5C80D
                    B865D0C6FF287DE8FF0CD3FE0476E514ED82D9A23DCB684D28E3B93A229A7B
                    D4DAF89E9A2F2D62599B91E8746830BCF88947611A82E9893137ABBA74B489
                    6C9C1492DCA2A7FA75F46451C7838EC0E9FB5D9222D3895C116C2C13E3995F
                    6D56ACB5F62FD7B764FAAB5AF0B5EA73AF3211B40AE44697DCB6E0D28E88E9
                    00037A506832C0BA
              SEQUENCE :
                 SEQUENCE :
                    CONTEXT SPECIFIC (0) :
                       INTEGER : 2
                    INTEGER : '18E922D0000000000060'
                    SEQUENCE :
                       OBJECT IDENTIFIER :  sha1withRSAEncryption [1.2.840.113549.1.1.5]
                       NULL :
                    SEQUENCE :
                       SET :
                          SEQUENCE :
                             OBJECT IDENTIFIER :  domainComponent
[0.9.2342.19200300.100.1.25]
                             IA5 STRING : 'com'
                       SET :
                          SEQUENCE :
                             OBJECT IDENTIFIER :  domainComponent
[0.9.2342.19200300.100.1.25]
                             IA5 STRING : 'contoso'
                       SET :
                          SEQUENCE :
                             OBJECT IDENTIFIER :  commonName [2.5.4.3]
                             PRINTABLE STRING :
                                'TestEnrollment'
                    SEQUENCE :
                       UTC TIME : '040812185455Z'
                       UTC TIME : '050812185455Z'
                    SEQUENCE :
                       SET :
                          SEQUENCE :
                             OBJECT IDENTIFIER :  domainComponent
[0.9.2342.19200300.100.1.25]
                             IA5 STRING : 'com'
                       SET :
                          SEQUENCE :
                               OBJECT IDENTIFIER :  domainComponent
[0.9.2342.19200300.100.1.25]
                               IA5 STRING : 'contoso'
```

```
                        SET :
                           SEQUENCE :
                              OBJECT IDENTIFIER :  commonName [2.5.4.3]
                              PRINTABLE STRING : 'Users'
                        SET :
                           SEQUENCE :
                              OBJECT IDENTIFIER :  commonName [2.5.4.3]
                              PRINTABLE STRING :
                                 'Avi Ben-Menahem'
                     SEQUENCE :
                        SEQUENCE :
                           OBJECT IDENTIFIER :  rsaEncryption [1.2.840.113549.1.1.1]
                           NULL :
                        BIT STRING UnusedBits:0 :
                           SEQUENCE :
                              INTEGER :
                                 00DAFF7C6859557C698CDA4598222E8E90EEB481889531E9F6
                                 7F10C081F2545B060BE7714E755325AC710774764DCA8120C6
                                 BEB7B6EF74B0260EDD56DD299B242A94EE83C420AC7FF0E694
                                 122E26EF67670782223C4E8D812C98047F24E10CF6A26FEBEE
                                 B826638924F36B697CEA02EFC4CA0D108CB85047266AD27DE5
                                 82D181A1
                              INTEGER : 65537
                     CONTEXT SPECIFIC (3) :
                        SEQUENCE :
                           SEQUENCE :
                              OBJECT IDENTIFIER :  sMIMECapabilities
[1.2.840.113549.1.9.15]
                              OCTET STRING :
                                 SEQUENCE :
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  rc2CBC [1.2.840.113549.3.2]
                                       INTEGER : 128
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  rc4 [1.2.840.113549.3.4]
                                       INTEGER : 128
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  desCBC [1.3.14.3.2.7]
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  DES-EDE3-CBC
[1.2.840.113549.3.7]
                           SEQUENCE :
                              OBJECT IDENTIFIER :  subjectKeyIdentifier [2.5.29.14]
                              OCTET STRING :
                                 OCTET STRING :
                                    8192563AC431F8820C54C9D0984FD8C534639ECC
                           SEQUENCE :
                              OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.10]
                              OCTET STRING :
                                 SEQUENCE :
                                    SEQUENCE :
                                       OBJECT IDENTIFIER :  encryptedFileSystem
[1.3.6.1.4.1.311.10.3.4]
                           SEQUENCE :
                              OBJECT IDENTIFIER :  keyUsage [2.5.29.15]
                              OCTET STRING :
                                 BIT STRING UnusedBits:5 :
                                    20
                           SEQUENCE :
                              OBJECT IDENTIFIER :  extKeyUsage [2.5.29.37]
                              OCTET STRING :
                                 SEQUENCE :
                                    OBJECT IDENTIFIER :  encryptedFileSystem
[1.3.6.1.4.1.311.10.3.4]
                           SEQUENCE :
                              OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.21.7]
                              OCTET STRING :
                                 SEQUENCE :
                                    OBJECT IDENTIFIER :

[1.3.6.1.4.1.311.21.8.4014942.3497959.5914804.3829722.12246394.103.3066650.1537810]
                                     INTEGER : 100
```

```
                        INTEGER : 2
                     SEQUENCE :
                        OBJECT IDENTIFIER :  authorityKeyIdentifier [2.5.29.35]
                        OCTET STRING :
                           SEQUENCE :
                              CONTEXT SPECIFIC (0) :
                                 10C8E49879236E65350924C24EFB074EFB5F4AA0
                     SEQUENCE :
                        OBJECT IDENTIFIER :  cRLDistributionPoints [2.5.29.31]
                        OCTET STRING :
                           SEQUENCE :
                              SEQUENCE :
                                 CONTEXT SPECIFIC (0) :
                                    CONTEXT SPECIFIC (0) :
                                       CONTEXT SPECIFIC (6) :
                                          'ldap:///CN=TestEnrollment,CN=dcross'
                                          '-stress,CN=CDP,CN=Public%20Key%20Se'
                                          'rvices,CN=Services,CN=Configuration'
                                          ',DC=contoso,DC=com?certificateRevoc'
                                          'ationList?base?objectClass=cRLDistr'
                                          'ibutionPoint'
                                       CONTEXT SPECIFIC (6) :
                                          'http://dcross-stress.contoso.com/Ce'
                                          'rtEnroll/TestEnrollment.crl'
                     SEQUENCE :
                        OBJECT IDENTIFIER :  authorityInfoAccess
[1.3.6.1.5.5.7.1.1]
                        OCTET STRING :
                           SEQUENCE :
                              SEQUENCE :
                                 OBJECT IDENTIFIER :  caIssuers
[1.3.6.1.5.5.7.48.2]
                                 CONTEXT SPECIFIC (6) :
                                    'ldap:///CN=TestEnrollment,CN=AIA,CN=Publi'
                                    'c%20Key%20Services,CN=Services,CN=Configu'
                                    'ration,DC=contoso,DC=com?cACertificate?ba'
                                    'se?objectClass=certificationAuthority'
                              SEQUENCE :
                                 OBJECT IDENTIFIER :  caIssuers
[1.3.6.1.5.5.7.48.2]
                                 CONTEXT SPECIFIC (6) :
                                    'http://dcross-stress.contoso.com/CertEnro'
                                    'll/dcross-stress.contoso.com_TestEnrollme'
                                    'nt.crt'
                     SEQUENCE :
                        OBJECT IDENTIFIER :  subjectAltName [2.5.29.17]
                        OCTET STRING :
                           SEQUENCE :
                              CONTEXT SPECIFIC (0) :
                                 OBJECT IDENTIFIER :  [1.3.6.1.4.1.311.20.2.3]
                                 CONTEXT SPECIFIC (0) :
                                    UTF8 STRING :
                                       'avibm@contoso.com'
            SEQUENCE :
               OBJECT IDENTIFIER :  sha1withRSAEncryption [1.2.840.113549.1.1.5]
               NULL :
            BIT STRING UnusedBits:0 :
               9D0000D2CC5668BEE443EBDE5EE4CADA5D61C17C00B262A3F231726FD2E7A8
               500603B89BE123D577FA2AE592567FB96743A6AE9B57AE089B1C205D6552F5
               5D60DD825D94D27301527FDB275035473DFC16A4F0C4886036A50CA1D320E3
               D284744CC0E552D1FFB24CD6110E6B17C86F830B5CC7A7E1791930320373CA
               C4E667BC372983597713CF8608389A6C82F9079FF8666C867BF2243DE5A22C
               20DBDBAD788A77758B68D9260EA5040A2F5C97C1AD80144F06F714D20BF671
               96BE5774D16080A9EAA5933C3C7EA34AE3F41DC001E0C2F83EA7AFAADA4812
               D0F27C48E288A20C44F085F328CCE6F478D6E4E89131D8EF43DA7B23DA39C9
               8CB15DE2EBA2BC8F
         SET :
            SEQUENCE :
               INTEGER : 1
               SEQUENCE :
                  SEQUENCE :
                     SET :
```

```
                        SEQUENCE :
                           OBJECT IDENTIFIER :  domainComponent
[0.9.2342.19200300.100.1.25]
                           IA5 STRING : 'com'
                     SET :
                        SEQUENCE :
                           OBJECT IDENTIFIER :  domainComponent
[0.9.2342.19200300.100.1.25]
                           IA5 STRING : 'contoso'
                     SET :
                        SEQUENCE :
                           OBJECT IDENTIFIER :  commonName [2.5.4.3]
                           PRINTABLE STRING :
                              'TestEnrollment'
                  INTEGER :
                     172B1FB96BBBF2BA49A64EBEA41833EF
               SEQUENCE :
                  OBJECT IDENTIFIER :  sha1 [1.3.14.3.2.26]
                  NULL :
               CONTEXT SPECIFIC (0) :
                  SEQUENCE :
                     OBJECT IDENTIFIER :  contentType [1.2.840.113549.1.9.3]
                     SET :
                        OBJECT IDENTIFIER :  [1.3.6.1.5.5.7.12.3]
                  SEQUENCE :
                     OBJECT IDENTIFIER :  messageDigest [1.2.840.113549.1.9.4]
                     SET :
                        OCTET STRING :
                           17CEEAA968CDD0A92DFC7E9AA174F87755AD8A87
               SEQUENCE :
                  OBJECT IDENTIFIER :  rsaEncryption [1.2.840.113549.1.1.1]
                  NULL :
               OCTET STRING :
                  C5D3AC35D5AAE5766640A2EF87D8ED005BB9BD63D51B10D803EEFEA1261161
                  3031241F695A2EDFF0240EE624D22FECB5AB6B74FD97A5DB12B3B873558AD5
                  0BC6DB59E438A7150A27749F53CBA447CD0751D7D49EEE3EBD1BBB20234887
                  5DD11DE26764DDEB2EBFA1E0023DD8CECF9C2530E2D0886FF26EAB747635A7
                  A57B7CA154BD0083A1DA891A35C3CD7EF5BA735FBCD2FD811FABE68C988C4D
                  172572BE63AE0575CF646756D4E66B2B127A699119368AAFB8B54661D317AF
                  2DF2622A0FFF01F18D5EF261E830107BD7F58848813CA6C0F8BF681A214E37
                  13618340D6DE9594829FB2B2DB1CFF973DC01F22D982846E474DDB9767D1BF
                  51E8C66F934593B5
```

**Recovery BLOB Structure**

When stored in the CA database, the private key is stored as a PKCS #7 message, encrypted with a 3DES symmetric key that is encrypted to the KRA(s) public key as a column in the CA database. When the recovery BLOB is retrieved by the certutil –getkey command, the encrypted PKCS #7 and the KRA certificate hashes are retrieved from the database. Also, the encrypted PKCS #7 is wrapped inside a signed PKCS #7 to allow collecting the previous certificates and attaching them to the signed PKCS #7. The PKCS #7 is not protected with a password since it is already protected by the public key of the recovery agent(s). The outer PKCS #7 wrapper can contain the certificate chains for the recovery agent(s) and the end-entity to facilitate the recovery operations and construction of the end-entity PKCS #12 file. Figure 47 illustrates the recovery BLOB structure.

The recovery BLOB consists of wrapping the encrypted PKCS #7 in the database in another (signed) PKCS #7 to allow a number of certificates to be included in the recovery BLOB. The returned certificates include the full chain of the user certificate being recovered, the chain of the signing CA certificate (which may differ from the CA certificate under which the user certificate was issued), and the KRA certificates to which the key was encrypted. The szOID_ARCHIVED_KEY_CERT_HASH(1.3.6.1.4.1.311.21.16) is an attribute containing the SHA-1 hash of the certificate for the key being recovered, attached as an authenticated attribute to the CA signature of the recovery BLOB. This allows certutil -recoverkey recoveryblobfile to also display the Subject name of the KRA certificate(s) used to protect the private key BLOB.




**Figure 47: Recovery BLOB**

**ASN.1 Structure**

The following is the ASN.1 structure of the PKCS #7 EnvelopedData object.

```
EnvelopedData ::= SEQUENCE {

version                    Version,

recipientInfos             RecipientInfos,

encryptedContentInfo       EncryptedContentInfo

}
```

Storing the recovery BLOB as an enveloped PKCS #7 enables a recovery agent to retrieve the recovery BLOB from the CA database. The recovery agent's private key is used to decrypt the EncryptedContentInfo to extract the PKCS #12 data. The following is the ASN.1 structure of the EncryptedContentInfo body.

```
EncryptedContentInfo ::= SEQUENCE {

contentType                    ContentType,

contentEncryptionAlgorithm     ContentEncryptionAlgorithmIdentifier,

encryptedContent[0]            IMPLICIT EncryptedContent OPTIONAL

}
```

By definition, there can be multiple recovery agent certificates specified by RecipientInfo, where IssuerAndSerialNumber is used to disambiguate between multiple recovery agent certificates. Only the recovery agent certificates included in the RecipientInfo body of the enveloped PKCS #7 object can be used to recover the archived key material. The following is the ASN.1 structure of the RecipientInfo body.

```
RecipientInfo ::= SEQUENCE {

version                    Version,

issuerAndSerialNumber      IssuerAndSerialNumber,

keyEncryptionAlgorithm     KeyEncryptionAlgorithmIdentifier,

encryptedKey               EncryptedKey

}
```

# Appendix B: Additional Information

Updated: December 6, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**General**
Windows Server 2003 documentation (http://go.microsoft.com/fwlink/?LinkID=17343)

**White Papers**
Implementing and Administering Certificate Templates in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkID=84928)

Certificate Autoenrollment in Windows XP (http://go.microsoft.com/fwlink/?LinkId=103241)

Windows Server 2003 PKI Operations Guide (http://go.microsoft.com/fwlink/?LinkID=59548)

PKI Enhancements in Windows XP Professional and Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkID=19644)

Encrypting File System in Windows XP and Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkID=22413)

Key Archival and Management in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=103251)

Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=103243)

**Standards**
Internet X.509 Public Key Infrastructure – Certificate and Certificate Revocation List (CRL) Profile (http://go.microsoft.com/fwlink/?LinkID=71066)

Certificate Management Messages over CMS (CMC) (http://go.microsoft.com/fwlink/?LinkId=103244)

PKCS #12 - Personal Information Exchange Syntax Standard (http://go.microsoft.com/fwlink/?LinkId=103245)

Common Criteria (http://go.microsoft.com/fwlink/?LinkId=103246)

CIMC Protection Profile (http://go.microsoft.com/fwlink/?LinkId=103247)

# Appendix C: Useful Commands

Updated: December 6, 2004

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following commands are available with certutil.exe in the Windows Server 2003 Administration Tools pack in performing key archival and recovery functions.

- **certutil – ImportKMS** accepts a *.pfx, *.epf, or a KMS export file, and archives the contents in the CA database.
- **certutil –ConvertEPF** converts a *.pfx (PKCS #12) file to a *.epf-formatted file for import into Outlook.
    - *certutil –ConvertEPF* and *–MergePFX* are similar in that they both accept a comma-separated list of input files that are merged and placed in the output file. Input files for both commands may be *.pfx files, *.epf files, or a mixture of the two.
    - The –cast parameter should be used to specify the CAST encryption algorithm.
    - The –silent option may be specified to suppress the UI.
- **certutil –ImportPFX** accepts a *.pfx or a *.epf file, and installs the certificates and keys into the HKLM (local machine) MY store. If the –User parameter is specified, the key and certificate will be imported into the HKCU (user profile) MY store. The CSP to be used may also be overridden by specifying the name with the –csp parameter.
- **certutil –getkey** retrieves the archived private key recovery BLOB from the CA database.
- **certutil –recoverkey** recovers the archived private key.
- **certutil –verifykeys** verifies the private/public key pair.

# ADFS Operations Guide

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

This Operations Guide provides administering and troubleshooting information for Active Directory Federation Services (ADFS) in the Microsoft® Windows Server™ 2003 R2, Enterprise Edition, and Microsoft Windows Server 2003 R2, Datacenter Edition, operating systems (for Federation Service, Federation Service Proxy, and ADFS Web Agent components); and the Microsoft Windows Server 2003 R2, Standard Edition, operating system (for ADFS Web Agent components).

**In this guide**


- Administering Active Directory Federation Services
- Troubleshooting Active Directory Federation Services


**See Also**
**Other Resources**
Active Directory Federation Services (ADFS)
Active Directory Federation Services Deployment Guide
Active Directory Federation Services Design Guide
ADFS Step-by-Step Guide

# Administering Active Directory Federation Services

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

This guide provides administering information for Active Directory® Federation Services (ADFS) in the Windows Server® 2003 R2, Enterprise Edition, and Windows Server 2003 R2, Datacenter Edition, operating systems (for Federation Service, Federation Service Proxy, and ADFS Web Agent components); and the Windows Server 2003 R2, Standard Edition, operating system (for ADFS Web Agent components).

**In this guide**


- Introduction to Administering ADFS
- Managing ADFS Components
- Managing ADFS Partnerships
- Managing Accounts and Account Stores
- Managing Claims and Claim Mapping



**Acknowledgments**

Produced by: Microsoft Windows Server User Assistance team

Project Writer: Mary Hillman

Project Editor: Femila Anilkumar

Contributing Writer: Nick Pierson

Technical Reviewers: Derek Del Conte, Vijay Gajjala, Dan Hartop, Ed Johns, Ryan D. Johnson, Jagadeesh Kalki, Carol Li, Vani Nori, Harini Raghavan, Rahul Shelar

# Introduction to Administering ADFS

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) is a component in Microsoft® Windows Server™ 2003 R2 that provides Web single-sign-on (SSO) technologies that allow the authentication of a user to multiple Web applications over the life of a single online session. ADFS accomplishes SSO by securely sharing digital identity and entitlement rights, or claims, across security and enterprise boundaries.

ADFS provides a robust environment that requires few frequent maintenance tasks. However, in operating a federation environment, you might have to perform certain tasks on a regular basis and others only as needed. This guide provides information and instructions for performing such tasks.

**Planning for ADFS Operations**

Operating Active Directory Federation Services (ADFS) consists of tasks and procedures for updating configurations for ADFS components as well as the installed applications and Windows components, including Windows Certificate Services, Internet Information Services (IIS), Active Directory directory service, and Active Directory Application Mode (ADAM).

When managing ADFS operations, you will need to update configurations for the ADFS components, including the ADFS servers (including federation servers and federation server proxies), ADFS Trust Policy, ADFS Web Agents, ADFS partnerships, ADFS account stores, and ADFS claims.

Before you begin, prepare a plan that establishes a baseline operating environment and addresses operational needs and actions.

To plan your ADFS operations environment, perform the following tasks:

- Assess your IT environment and establish a baseline.
- Determine your operational needs.

**Assess Your IT Environment and Establish a Baseline**

- Understand the details of the federated Web sites and partners that the ADFS deployment must support to effectively and securely operate ADFS servers. For information about planning for specific ADFS scenarios, see the Active Directory Federation Services Design Guide (http://go.microsoft.com/fwlink/?LinkId=63486).
- Review any service specifications that were produced during the planning and deployment process, along with any service-level requirements defined in service level agreements between partner IT organizations.

You will need to understand the following environmental conditions and requirements when you establish your operations baseline and to accommodate growth and modifications to your IT environment:

- Supported partners: When using ADFS, you are usually working with partner organizations. When establishing identity federation, determine the organizations with which you want to form a partnership. After a baseline ADFS deployment is in place, operating with partners involves adding partners, deleting partners, and updating partner information. Changes to partnerships can be required for a variety of reasons. For example, your ADFS deployment might require partnership updates if your partner changes its business significantly, your organization becomes part of a larger organization or a federation of organizations, or your organization is acquired by a different company. In any scenario where you are federating identities from multiple domains, you will need to be aware of the domains (partners) that you are currently supporting and all additional domains that represent potential partners.
- Supported application types: Some ADFS applications require access to operating system resources, while others are "claims aware." It is important to understand the type of applications that ADFS will support so that administration requirements can be formulated.
- Logical and physical architectural diagrams or deployment topology: You will need to know whether ADFS is working in a set of farmed servers or a single server. You must understand where your network deploys firewalls and proxies. You

must also be aware of the location of resources and whether the users are accessing resources from within your organization or from outside the organization, or both.

- Certificate and trust information: You should understand how the certificates in the environment have been acquired and used. For example, it is important to understand whether the certificates follow a chain up to a root certification authority (CA) and how your certificates are obtained so that you can address certificate renewals. It is important to understand how certificate revocation works in your environment.

- User account management requirements: You must understand how users gain access to resources in your environment, whether external users have access to resources in your domain, and whether you have enabled fine-grained control or are leveraging groups to dictate access control.

Data for these conditions provide a starting point for establishing a baseline for the operations environment and for setting the proper level of service.

**Determine Operational Needs**
Performing operations require that tasks are assigned to the appropriate server administrators on the teams that support the ADFS deployment. The ADFS operations team must establish processes for managing the following ADFS components and their related configurations:

- Federation Service:

    - If you find that the federation server or server farm is not meeting scalability, performance, or reliability requirements, you can add an additional federation server.

    - If you want to monitor access or diagnose failures, you can modify certain logging settings at the federation server.

    - Many federation server tasks encompass establishing and managing partnerships, resources, and accounts, as explained later in this topic.

- Federation Service Proxy (optional):

    - If you find that the federation server proxy or server farm is not meeting scalability, performance, or reliability requirements, you can add an additional federation server proxy.

    - You might add a proxy to an existing deployment as part of enabling Internet access to your existing resources.

    - If the client authentication model changes, you can change the federation server proxy to handle this authentication model.

- Federated application(s):

    - When you add a new Web application that is protected by ADFS--for example, if your organization has a new Web site for purchase order management--you will need to add the application to the ADFS deployment.

    - When the type of user information that the application requires to make its decisions changes, you will need to update claims in your ADFS deployment.

    - When the URL for the application changes, you will need to update URLs in your ADFS deployment.

- Partnerships:

    - When establishing a new relationship due to acquisitions, mergers, business contracts, and so on, you will need to add and remove partnerships in your ADFS deployment.

    - When changing an existing business relationship, you might need to update properties in your ADFS partnership.

- Accounts:
    - In federated scenarios where the application is a Windows NT token-based application, you will need to manage resource groups in the resource partner.
    - When you establish a new relationship, you will need to add additional accounts to your account store.
    - When additional users and groups need access to an existing application, you will need to add these accounts to your account store.

**When to Use This Guide**
You should use this guide when:

- You have ADFS deployed in a test or production environment.
- You want to add or remove ADFS components.
- You want to make changes to the configuration of ADFS components.
- You want to add or remove Web applications that authenticate through ADFS.

This guide assumes a basic understanding of ADFS, how it works, and why your organization uses it to federate Web applications. You should also have a thorough understanding of how ADFS is deployed and managed in your organization, including an understanding of the mechanism your organization uses to configure and manage ADFS settings. To learn about ADFS concepts and scenarios, see the following ADFS documentation:

- Overview of ADFS (http://go.microsoft.com/fwlink/?LinkId=63491)
- Active Directory Federation Services (ADFS) (http://go.microsoft.com/fwlink/?LinkId=57765)
- Active Directory Federation Services Design Guide (http://go.microsoft.com/fwlink/?LinkId=63486)

This guide can be used by organizations that have deployed Microsoft Windows Server 2003 R2. It includes information that is relevant to different roles within an IT organization, including IT operations management and administrators. It contains high-level information that is required to plan an ADFS operations environment. This information provides management-level knowledge of ADFS and the IT processes required to operate it.

In addition, this guide contains more detailed procedures that are designed for operators who have varied levels of expertise and experience. Although the procedures provide operator guidance from start to finish, operators must have a basic proficiency with the Microsoft Management Console (MMC) and snap-ins and know how to start administrative programs and access the command line. If operators are not familiar with ADFS, it might be necessary for IT planners or IT managers to review the relevant operations in this guide and provide the operators with parameters or data that must be entered when the operation is performed.

**How to Use This Guide**
The operations areas are divided into the following types of content:

- Tasks pertain to group-related procedures and provide general guidance for achieving the goals of an objective. In this guide, "Managing ADFS Web Agents and Applications" is an example of a task.
- Procedures provide step-by-step instructions for completing tasks. In this guide, "Add an ADAM account store" is an example of a procedure topic.

If you are an IT manager who will be delegating tasks to operators within your organization, you will want to:

- Read through the tasks to determine whether you need to install tools before operators perform the procedures for each

task.

- Before assigning tasks to individual operators, ensure that you have all the tools installed where operators can use them.

- When necessary, create "tear sheets" for each task that operators perform in your organization. Cut and paste the task and its related procedures into a separate document, and then either print these documents or store them online, depending on the preference of your organization.

**Technologies and Terminology Used in This Guide**

Active Directory Federation Services (ADFS) uses terminology from several different technologies, including certificate services, Internet Information Services (IIS), Active Directory, Active Directory Application Mode (ADAM), and Web Services (WS*).

For more information about these technologies, see:

- Certificate services: Public Key Infrastructure for Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=19936) and Public Key Infrastructure (http://go.microsoft.com/fwlink/?LinkId=54917).

- IIS: Windows Server Internet Information Services (IIS) (http://go.microsoft.com/fwlink/?LinkId=63492).

- Active Directory: Active Directory Collection in the Microsoft Windows Server 2003 Technical Reference on (http://go.microsoft.com/fwlink/?LinkId=63494).

- ADAM: Active Directory Application Mode Technical Reference (http://go.microsoft.com/fwlink/?LinkId=63506).

- Web Services: Security Specifications (http://go.microsoft.com/fwlink/?LinkId=44191).

For a list of ADFS terms and definitions, see Terminology used in ADFS (http://go.microsoft.com/fwlink/?LinkId=63507).

**See Also**

**Other Resources**

Active Directory Federation Services (ADFS)
Overview of ADFS
Active Directory Federation Services Design Guide
ADFS Step-by-Step Guide
Public Key Infrastructure for Windows Server 2003
Public Key Infrastructure
Certificate Services Technical Reference
Windows Server Internet Information Services (IIS)

# Managing ADFS Components

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) has three components:

- Federation Service: Functions as a security token service and routes authentication requests from external user accounts in partner organizations and clients on the Internet.
- Federation Service Proxy: Functions as a proxy for the Federation Service in a perimeter network. This component is optional in an ADFS deployment.
- ADFS Web Agent: Provides authorization (consumes security tokens) and either allows or denies access to two types of Web applications, as follows:
    - Claims-aware applications: Makes authorization decisions based on claims.
    - Windows NT token-based applications: Uses Windows-based authorization mechanisms.

The following objectives are part of managing ADFS components:

- Managing the Federation Service
- Managing the Federation Service Proxy (Optional)
- Managing ADFS Web Agents and Applications
- Backing Up ADFS Components

**See Also**
**Other Resources**
Understanding ADFS Components

# Managing the Federation Service

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The Federation Service component of Active Directory Federation Services (ADFS) functions as a security token service. The act of installing the Federation Service component on a computer makes that computer a federation server.

The following objectives are part of managing the Federation Service:

- Managing a Federation Server Farm
- Managing Trust Policy on Federation Servers
- Managing Certificates Used by Federation Servers

**See Also**
**Concepts**
Managing the Federation Service Proxy (Optional)
Managing ADFS Web Agents and Applications
Backing Up ADFS Components

# Managing a Federation Server Farm

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

To provide load-balancing of security services in larger Active Directory Federation Services (ADFS) deployments, you can install additional federation servers. A Federation Service must be able to verify tokens issued by all federation servers in that farm. To accomplish this verification, the Federation Service uses verification certificates. Therefore, implementing a server farm requires you to configure verification certificates in the shared trust policy for every token-signing certificate that is in use by any server in the farm.

The following properties are shared by all federation servers in a server farm:

- TrustPolicy.xml file: This file contains the information for a Federation Service. This file must be accessible to all servers in the server farm, either by location in a shared directory or by using a file distribution method that ensures the replication of updates, such as Distributed File System (DFS).

- Certificate: Federation servers in the farm can each use a different token-signing certificate, or they can all use the same certificate. When using the same certificate, every federation server must have its own local copy of that certificate configured in the certificate store, with access to the private key.

- Account store: All servers recognize the same Active Directory domain or Active Directory Application Mode (ADAM) store for user authentication.

The following tasks for managing a Federation Service server farm are described in this objective.

- Changing the Location of the ADFS Directory

- Implementing a Server Farm of Federation Servers

- Adding a New Federation Server

- Removing a Federation Server

**See Also**
**Other Resources**
Distributed File System (DFS)

# Changing the Location of the ADFS Directory

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

To change the Active Directory Federation Services (ADFS) directory to another location, you can perform the following procedure on a federation server or federation server proxy. The default location for the directory is C:\ADFS.

| **Important** |
| --- |
| Microsoft does not support changing the location of the ADFS directory. If you alter its location, we cannot guarantee that you can upgrade normally to future versions of Windows operating systems or that ADFS will function after you upgrade. |

To complete this procedure, you must finish all of the following steps in the order that they are presented:

1. Copy the contents of the ADFS directory to the new location.
2. Change the virtual directory path in Internet Information Services (IIS) Manager.
3. Modify the registry to point to the new ADFS directory.

| **Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, back up any valued data on the computer. |

### Administrative credentials

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

#### Copy the contents of the ADFS directory to the new location

1. Start Windows Explorer, and then create a new directory named ADFS in the desired location on the computer.
2. Copy the contents of the C:\ADFS directory (including subfolders) to the new location.

| **Important** |
| --- |
| If you plan to upgrade the Windows Server operating system to a newer version in the future, leave the C:\ADFS directory intact. To do so, copy, rather than move, the directory to the new location. |

#### Change the virtual directory path in IIS Manager

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. Double-click *YourComputerName***(local computer)**, double-click **Web Sites**, and then double-click **Default Web Site**.
3. In the list of virtual directories, right-click **adfs**, and then click **Properties**.
4. On the **Virtual Directory** tab, in **Local path**, type the path to the new ADFS\STS directory, and then click **OK**.

   For example, if you want to move the directory to D:, type **d:\adfs\sts** in **Local path**.

#### Modify the registry to point at the new ADFS directory

1. Open Registry Editor (Regedit.exe).
2. Navigate to the **HKEY_LOCAL_MACHINE_SYSTEM\Microsoft\WebSSO** registry key.
3. Right-click **InitialDirectory**, click **Modify**, under **Value data**, type the path to the new ADFS directory, and then click

**OK**.

For example, if you want to move the directory to D:, then type **d:\adfs** in **Value data**.

4. Close Registry Editor.

# Implementing a Server Farm of Federation Servers

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can implement a server farm of servers that are running the Federation Service component of Active Directory Federation Services (ADFS) by appropriately installing and configuring token-signing certificates. Use the following options when installing the certificate and the Federation Service to implement a federation server farm:

- During installation of the token-signing certificate for a server, if you select the option to install the certificate into the local certificate store, the certificate becomes available during installation of the Federation Service.
- During installation of the Federation Service, the Windows Components Wizard provides options that affect how the token-signing certificate is configured:
    - Select token-signing certificate: You can use this option to select the installed token-signing certificate from the local certificate store.
    - Use an existing trust policy: If you select this option, ADFS automatically adds the public portion of the selected token-signing certificate to the shared trust policy of the Federation Service as the verification certificate.

Use the following methods for installing and sharing certificates to implement a federation server farm:

- Use a separate token-signing certificate for each server and generate the respective verification certificates during Federation Service installation:
    - Install a separate token-signing certificate on each server.
    - During Federation Service installation, select the installed certificate and the shared trust policy.
- Share both public and private portions of the same certificate by using an image of the server:
    - Install a token-signing certificate on one server.
    - During Federation Service installation, select the installed certificate and the shared trust policy.
    - Create an image of this server and use this image to create all additional servers in the server farm.
- Share both public and private portions of the same certificate by importing the certificate file that is provided by a public certification authority (CA) into the local certificate store.
    - Obtain a token-signing certificate from a public CA.
    - Use physical media to import the certificate into the local certificate store of each server.
    - During Federation Service installation, select the CA-provided token-signing certificate and the shared trust policy file.
- Share both public and private portions of the same certificate by exporting the private key:
    - Install a single token-signing certificate from an enterprise CA on a server and export the private key. This method requires that the token-signing certificate was generated and placed directly into the local certificate store by the enterprise CA and that private keys were marked as exportable.
    - Export the private key certificate to a file and protect it accordingly.

- Prior to Federation Service installation, import the exported private key certificate into the local certificate store on each additional server.

- During Federation Service installation, select the imported certificate and the shared trust policy. It is not necessary to export the public key because the trust policy is shared.

Completion of this task is accomplished during the course of Adding a New Federation Server.

**See Also**
**Concepts**
Managing Certificates Used by Federation Servers
Managing Token-signing Certificates
Rolling Over a Token-signing Certificate

# Adding a New Federation Server

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you want to add a new federation server to an existing Active Directory Federation Services (ADFS) deployment, you must configure the server as an application server, install and configure certificates, and install the Federation Server component of ADFS according to the method of implementing a server farm that you are using. You can also set event logging according to server needs.

**Task requirements**

You need the following to perform the procedures for this task:

| **Note** |
| --- |
| If an existing server image is being used to prepare the additional federation server, procedures in this task are not required. Use your imaging software to create a new federation server. |

- An installed Secure Sockets Layer (SSL) certificate. For information about how to acquire SSL certificates, see Obtaining Server Certificates (http://go.microsoft.com/fwlink/?LinkId=62479).
- If the token-signing certificate is shared among servers in the server farm, an existing token-signing certificate for the Federation Service.
- An existing Federation Service.
- The location of the shared Trustpolicy.xml file for the Federation Service.

To complete this task, perform the following procedures:

1. Install prerequisite applications.
2. If an existing token-signing certificate private key is shared among the servers in the server farm, go to step 3.

   If a new token-signing certificate is to be installed into the local certificate store on the new server, install a token-signing certificate on the new server, as follows:

   - If you are using Microsoft Certificate Services as an enterprise certification authority (CA), obtain a new client authentication certificate. For more information about obtaining a client authentication certificate, see Submit an advanced certificate request via the Web to a Windows Server 2003 CA (http://go.microsoft.com/fwlink/?LinkId=64020). Specify installing the certificate into the local certificate store.
   - If you are using a different enterprise CA or a public CA, follow the instructions provided by the CA.
   - Alternatively, Create a self-signed, token-signing certificate.
3. Configure the token-signing certificate according to the method of server farm implementation you are using:

   - If you installed a separate token-signing certificate into the local certificate store and are not sharing the private key, no other certificate configuration is required prior to Federation Service installation.
   - If you are sharing the public and private portions of the same certificate that has been provided by a public certification authority (CA), import the certificate to the local certificate store prior to Federation Service installation. For instructions to import the certificate, see Import a certificate (http://go.microsoft.com/fwlink/?linkid=20040).

- If you are sharing the public and private portions of the same exportable certificate that has been provided by an enterprise CA, import the private key into the local certificate store by performing the following procedures prior to Federation Service installation:

  Export the private key portion of a token-signing certificate.

  Import a certificate from the Microsoft Web site (http://go.microsoft.com/fwlink/?linkid=20040).

4. Install the Federation Service on an additional federation server

5. Configure event logging on a federation server

**See Also**
**Concepts**
Implementing a Server Farm of Federation Servers
Removing a Federation Server
Adding a New Federation Server Proxy

# Install prerequisite applications

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) requires that the following applications be installed on a computer running Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Standard Edition, (only for servers hosting the ADFS Web Agent) so that the computer can host the Federation Service, Federation Service Proxy, or ADFS Web Agent components of ADFS:

- Internet Information Services (IIS)
- Microsoft .NET Framework 2.0
- Microsoft ASP.NET (required for claims-aware applications only)

When you add the Application Server Windows component, these three applications are installed.

**Administrative credentials**

To complete this procedure, you must be a member of the local Administrators group.

**To add the Application Server Windows component**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard**, select the **Application Server** check box, and then click **Next**.
4. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server

# Create a self-signed, token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use the following procedure to create a self-signed, code-signing certificate that also creates and installs a private key. To perform this procedure, use the Makecert.exe utility. Makecert.exe is available in the Microsoft .NET Framework 2.0 Software Development Kit (SDK) (x86) (http://go.microsoft.com/fwlink/?LinkId=79548).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create a self-signed, token-signing certificate**

1. Open a command prompt.
2. Type the appropriate **makecert** syntax.
   - Example command:

**makecert -r -pe -n "CN=CertForADFS" -b 01/01/2006 -e 01/01/2007 -eku 1.3.6.1.5.5.7.3.3 -ss my -sr localMachine -sky exchange -sp "Microsoft RSA SChannel Cryptographic Provider" -sy 12 "CertForADFS.cer"**

| **Note** |
| --- |
| Track certificate expiration dates to make sure that certificates are replaced before they expire. You can do this using the Active Directory Federation Services snap-in for certificates of the current organization and also for partners that you configure in the trust policy. |

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Rolling Over a Client Authentication Certificate

# Export the private key portion of a token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Every federation server in an Active Directory Federation Services (ADFS) server farm must have access to the private key of the token-signing certificate. If you are implementing a server farm of federation servers that share a single, exportable private key certificate that is issued by an enterprise certification authority (CA), the private key portion of the existing token-signing certificate must be exported to make it available for importing into the certificate store on the new server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To export the private key of a token-signing certificate**

1. Click Start, point to Administrative Tools, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service**, and then click **Properties**.
3. On the **General** tab, click **View**.
4. In the **Certificate** dialog box, click the **Details** tab.
5. On the **Details** tab, click **Copy to File**.
6. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
7. On the **Export Private Key** page, select **Yes, export the private key**, and then click **Next**.
8. On the **Export File Format** page, select **Personal Information Exchange = PKCS #12 (.PFX)**, and then click **Next**.
9. On the **Password** page, type and confirm the password that is required to share the token-signing certificate. You will need this password when you select the exported token-signing certificate when installing the Federation Service.
10. On the **File to Export** page, specify the certificate file, and then click **Next**.
11. On the **Completing the Certificate Export Wizard** page, click **Finish**.
12. Validate the success of your export by confirming that the file you specified is created at the specified location.

| **Important** |
| --- |
| So that this certificate can be imported to the local certificate store on the new server, you must transfer the file to physical media and protect its security during transport to the new server. It is extremely important to guard the security of the private key. |

13. Import the exported certificate into the certificate store on the new server prior to installing the Federation Service. For information about how to import the certificate, see Import a certificate (http://go.microsoft.com/fwlink/?linkid=20040).

**See Also**
**Concepts**
Implementing a Server Farm of Federation Servers

# Install the Federation Service on an additional federation server

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you install an additional federation server in a server farm, you add the Federation Service component of Active Directory Federation Services (ADFS) and use the trust policy file that is used by other federation servers in the server farm. A token-signing certificate must have been installed or imported into the local certificate store. The trust policy file must also be available on the network.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To install the Federation Service component of ADFS on an additional server**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, select the **Active Directory Services** check box, and then click **Details**.
4. In the **Active Directory Services** dialog box, select the **Active Directory Federation Services (ADFS)** check box, and then click **Details**.
5. In the **Active Directory Federation Services (ADFS)** dialog box, select the **Federation Service** check box, and then click **OK**. If a message appears stating that ASP.NET 2.0 was not previously enabled, click **Yes** to enable it, and then click **OK**.
6. In the **Active Directory Services** dialog box, click **OK**.
7. In the Windows Components Wizard, click **Next**.
8. On the **Federation Service** page, click **Select token signing certificate**, and then click **Select** to select the appropriate token-signing certificate from the local certificate store.
9. In the **Select Certificate** dialog box, click the exported token-signing certificate, and then click **OK**.
10. Under **Trust policy**, click **Use an existing trust policy**, and then click **Browse**.
11. In the **Browse** dialog box, click the shared TrustPolicy.xml file, and then click **Open**.
12. If you are prompted for the location of the installation files, navigate to *R2 Installation Folder*\cmpnents\r2, and then click **OK**.
13. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Implementing a Server Farm of Federation Servers
Export the private key portion of a token-signing certificate

# Configure event logging on a federation server

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Servers that are running the Federation Service component of Active Directory Federation Services (ADFS) log ADFS Federation Service events in the Application event log. These events report information about the operation of the components of the local organization and the partner organizations that are covered by a trust policy.

| **Note** |
| --- |
| When it is manually configured, ADFS also can log debug information. Debug logs are located in *%systemdrive%*\ADFS\logs. For more information about how to configure debug logging, see Configuring ADFS Servers for Troubleshooting (http://go.microsoft.com/fwlink/?LinkId=74970). |

The following types of events are available and enabled by default in ADFS:

- **Error:** Information about a significant problem of which the user should be aware, usually involving a loss of functionality or data.
- **Warning:** Indicates a problem that is not immediately significant, but that may signify conditions that could cause future issues.
- **Info:** Information about a significant, successful operation.
- **Success audit:** Indicates an audited security event when an audited access attempt is successful; for example, a successful logon attempt.
- **Failure audit:** Indicates a security event that occurs when an audited access attempt fails; for example, an inbound token was not valid.
- **Detailed success:** A success audit event with detailed information about each token involved in the transaction, including claims information.
- **Detailed failure:** A failure audit event with detailed information about each token involved in the transaction, including claims information.

You can select the levels that you want to enable and disable.

| **Note** |
| --- |
| Audit object access must be turned on for success or failure to allow the Federation Service to log errors. For more information, see Audit object access (http://go.microsoft.com/fwlink/?LinkId=79749). |

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

#### To change the event types that are logged by ADFS

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click the **Trust Policy** node, and then click **Properties**.
3. Scroll to the **Event Log** tab.
4. Under **Event log level**, click to select and deselect event log types, and then click **OK**.

# Removing a Federation Server

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you no longer need a server that is running the Federation Service component of Active Directory Federation Services (ADFS), perform the following procedures:

- Remove the Federation Service ADFS component, as described in Remove an ADFS component.
- If you have a federation server farm and the trust policy is shared on the network, disconnect the mapped network drive for the connection to the shared trust policy folder.
- Delete the token-signing certificate from the certificate store. For information about how to delete a certificate from a certificate store, see Delete a certificate (http://go.microsoft.com/fwlink/?LinkId=62715.
- If this is the only instance of the Federation Service, do the following:
    - Notify the partner that this issuer is no longer valid and tell them to remove the partner node for this Federation Service, as described in Removing ADFS Partners.
    - Decommission the Web server and change the authentication for any Web applications.

**See Also**
**Concepts**
Removing a Federation Server Proxy
Removing ADFS Partners

# Remove an ADFS component

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you want to remove a server that is running the Federation Service or Federation Service Proxy component of Active Directory Federation Services (ADFS) from a Federation Service, use the following procedure to remove the respective ADFS component.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To remove an ADFS component**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, select the **Active Directory Services** check box, and then click **Details**.
4. In the **Active Directory Services** dialog box, select the **Active Directory Federation Services (ADFS)** check box, and then click **Details**.
5. In the **Active Directory Federation Services (ADFS)** dialog box, clear the **Federation Service** check box or the **Federation Service Proxy** check box, and then click **OK**.
6. In the **Active Directory Services** dialog box, click **OK**.
7. In the Windows Components Wizard, click **Next**.
8. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Remove the ADFS Web Agent component of ADFS

# Managing Trust Policy on Federation Servers

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The Active Directory Federation Services (ADFS) trust policy file defines the set of parameters that a Federation Service requires to identify partners, certificates, account stores, claims, and various properties of these entities that are associated with the Federation Service.

**Task requirements**

You need the following to perform the procedures for this task:

- Active Directory Federation Services MMC snap-in
- TrustPolicy.xml file located in %systemdrive%\ADFS

The following procedures for managing trust policy on federation servers are described in this task. Use these procedures on an as-needed basis.

- Change the Federation Service trust policy location
- Change the Federation Service URI
- Change the Federation Service endpoint URL
- Configure the Federation Service default Web pages
- Change the primary display name for a trust policy
- Change a claims transform module
- Change the token lifetime for a Federation Service
- Change the trust policy refresh period
- Change the Windows domain trust cache refresh period

**See Also**
**Concepts**
Managing ADFS Partnerships
Managing Accounts and Account Stores
Managing Claims and Claim Mapping
Managing Certificates Used by Federation Servers

# Change the Federation Service trust policy location

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The TrustPolicy.xml file contains data that defines a Federation Service in Active Directory Federation Services (ADFS). By sharing the folder that contains this file, you make the data available to other federation servers in a federation server farm that comprises the Federation Service.

If the location of the TrustPolicy.xml file changes, make sure that you update the location on each server in the Federation server farm.

Perform this procedure on a federation server that hosts the TrustPolicy.xml file you want to move. After copying the TrustPolicy.xml file to its new location, update the path on each server in the server farm, and then delete the old TrustPolicy.xml file.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the location of the trust policy file**

1. Copy the TrustPolicy.xml file from its current location to its new location and share the folder appropriately. (The default location for TrustPolicy.xml is *%systemdrive%*\ADFS).
2. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
3. Right-click **Federation Service**, and then click **Properties**.
4. On the **General** tab, replace the text in **Trust policy file** with the new path to the TrustPolicy.xml file, or click **Browse** to select it.
5. Confirm the new path in **Trust policy file**, and then click **OK**.
6. Repeat steps 2 through 5 on each federation server in the server farm.
7. After successfully reconfiguring each federation server with the new trust policy location, delete the original trust policy file.

# Change the Federation Service URI

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), the Federation Service Uniform Resource Identifier (URI) uniquely identifies a Federation Service. This URI also identifies the federation server farm membership of the federation server. URIs are case-sensitive.

In a federated scenario, the URI that is specified in the properties of the Trust Policy node of a Federation Service also identifies the Federation Service in the properties of the respective account or resource node in the partner organization. Therefore, if you change the URI of a partner Federation Service, you must change its value in the corresponding partner node in the trust policy of the other partner. If the partners are in separate organizations, you must communicate this change to the administrator who makes these changes.

URI values must match according to standard URI comparison rules, as described in Request for Comments (RFC) 3986, "Uniform Resource Identifier (URI): Generic Syntax" (http://go.microsoft.com/fwlink/?LinkId=65481).

| **Note** |
| --- |
| Changes to the Federation Service URI are received by federation servers in the Federation Service through the shared TrustPolicy.xml file. However, this change must be made manually in the corresponding Federation Service of the partner organization. During this time, users who have already signed on using the URI will have to authenticate again if they return to the same Federation Service, such as when they try to access another site. The contents of the cookie in the access token, which lasts 10 hours by default, will no longer match the Federation Service. For this reason, the user will be prompted again for credentials |

Perform the following procedure on a federation server in the Federation Service whose URI you want to change.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the Federation Service URI in the Trust Policy properties**

1. On the federation server whose Federation Service URI you want to change, click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click **Trust Policy**, and then click **Properties**.
3. On the **General** tab, in **Federation Service URI**, type the new URI, and then click **OK**.

The URI value must also be updated in the properties of the corresponding resource or account partner node in the partner Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the Federation Service URI in the account partner or resource partner properties**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners** if you are logged on to the resource federation server, or double-click **Resource Partners** if you are logged on to the account federation server.
3. Right-click the account or resource partner whose Federation Service URI has changed, and then click **Properties**.
4. On the **General** tab, in **Federation Service URI**, type the new URI, and then click **OK**.

| **Note** |
| --- |
| This URI must match the URI in the trust policy of the corresponding partner organization according to standard URI comparison rules. |

# Change the Federation Service endpoint URL

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), the Federation Service endpoint Uniform Resource Locator (URL) is the single location, or "public URL," that is used to contact all federation servers in a server farm. If the public name changes, you must change the Federation Service endpoint URL.

The Federation Service endpoint URL that is specified in the properties of the Trust Policy node of one Federation Service also identifies the Federation Service in the properties of the respective account or resource node in the partner organization. For example, for the account partner Federation Service in an ADFS deployment that has both a resource partner and an account partner, the Federation Service endpoint URL that is specified in the properties of the account partner node on the respective resource federation server must match the Federation Service endpoint URL in the Trust Policy properties on the account federation server. Consequently, if you change the endpoint URL of the Federation Service on the Trust Policy node, an administrator in the partner organization must also change the Federation Service endpoint URL on the corresponding account partner node.

If you install the Federation Service Proxy component of ADFS in a perimeter network, the Federation Service endpoint URL must point to the Federation Service Proxy that forwards requests to the Federation Service. In this case, the Federation Service endpoint URL contains the URL of the Federation Service Proxy rather than the partner Federation Service.

| **Note** |
| --- |
| Changes to the Federation Service endpoint URL are received by federation servers through the shared TrustPolicy.xml file. However, this change must be made manually in the trust policy of the corresponding partner organizations. |

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the Federation Service endpoint URL**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click **Trust Policy**, and then click **Properties**.
3. On the **General** tab, in **Federation Service endpoint URL**, select the domain portion of the URL, replace the selected text with the new URL, and then click **OK**.

Perform the following procedure on the federation server where you are updating the Federation Service URL on the respective account partner or resource partner after it has been changed in the trust policy for the home Federation Service.

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the Federation Service endpoint URL in the account partner or resource partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners** if you are logged on to the resource federation server, or double-click **Resource Partners** if you are logged on to the account federation server.
3. Right-click the account or resource partner whose Federation Service endpoint URL has changed, and then click **Properties**.
4. On the **General** tab, in **Federation Service endpoint URL**, select the domain portion of the URL, replace the selected text with the new endpoint URL, and then click **OK**.

**See Also**
**Concepts**
Set the application URL for an application
Set the return URL for a claims-aware application
Set the return URL for a Windows NT token-based application

# Configure the Federation Service default Web pages

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When a Web resource is protected by Active Directory Federation Services (ADFS), three .aspx files specify the default Web pages that are presented before and after access to the resource:

- clientlogon.aspx: This page is presented by the Federation Service or Federation Service Proxy for collecting credentials from the user. It may be in the form of the current Windows credentials using Windows Integrated authentication or Basic authentication, or it might provide for forms-based user name and password entry.

- discoverclientrealm.aspx: This page is presented by the resource Federation Service or Federation Service Proxy when the realm of the client is not known.

- signout.aspx: This page is presented by the resource Federation Service or Federation Service Proxy after the user signs out of the resource Web page.

The default location of these files on federation servers and federation server proxies is %systemdrive%\ADFS\sts\ls.

You can change the names of these files, but the location of the files must be %systemdrive%\ADFS\sts\ls or a subdirectory of this location. Use Windows Explorer to change the file names or locations, and then change their location in the Active Directory Federation Services snap-in.

Perform this procedure on a federation server or federation server proxy.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the Federation Service default Web pages**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service** on a federation server, or **Federation Service Proxy** on a federation server proxy, and then click **Properties**.
3. On the **Web Pages** tab, change the names of the files in **Client logon page**, **Account partner discovery page**, or **Client logoff page**, and then click **OK**.
4. Validate the effect of the change by performing an action on the Web application. Verify that you receive the new logon page.

| **Note** |
| --- |
| In certain cases (for example, when you are using existing facilities offered by ADFS and IIS for authentication), you might additionally need to copy IIS settings related to authentication, as well as the auth directory (%systemdrive%\ADFS\sts\ls\auth*) where ADFS stores the different files that leverage IIS authentication. For example, ...\auth\integrated stores files that are used for integrated authentication using Windows credentials. Similarly, ...\auth\sslclient stores files that are used for SSL authentication. |

# Change the primary display name for a trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Each Federation Service in an Active Directory Federation Services (ADFS) deployment has a primary display name. The discoverclientrealm.aspx page displays these names during logon so that the user can select the correct account realm for his organization.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the primary display name for a trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click the **Trust Policy** node, and then click **Properties**.
3. Click the **Display Name** tab.
4. In **Display name for this trust policy**, type the new primary display name, and then click **OK**.

# Change a claims transform module

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can configure servers that are running the Federation Service component of Active Directory Federation Services (ADFS) to use a claims transform module to manage claims. If you use a claims transform module and you make a change to the class name that the transform module uses, use the following procedure to make the change in the trust policy.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the claims transform module**

1. On the federation server whose transform module you want to change, click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click **Trust Policy**, and then click **Properties**.
3. In the **Trust Policy Properties** dialog box, click the right arrow to scroll to the **Transform Module** tab, and then click the **Transform Module** tab.
4. In **DLL file**, type the path to the dynamic-link library (DLL) file, or click **Browse** to select it.
5. In **Class name**, type the namespace-qualified class name that the transform module will use, and then click **OK**.

**See Also**
**Concepts**
Configure a claims transform module

# Change the token lifetime for a Federation Service

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Each server that is running the Federation Service component of Active Directory Federation Services (ADFS) issues Security Assertion Markup Language (SAML) tokens. These security tokens are transferred between clients and federation servers within authentication cookies. Authentication cookies can be issued by both the Federation Service and the ADFS Web Agent.

| **Note** |
| --- |
| At the Federation Service, the security token in an authentication cookie holds the organization claims for the client. For more information about SAML tokens, see WS-Federation: Passive Requestor Profile (http://go.microsoft.com/fwlink/?LinkId=64813). |

After the Federation Service validates the client once, the authentication cookie is written to the client. Further authentication takes place through use of the cookie rather than through repeated authentication of the client credentials. In this way, the authentication cookie facilitates single sign-on (SSO). The period of time that this authentication cookie can be used by the client is configurable in the trust policy of the Federation Service as the "token lifetime."

| **Note** |
| --- |
| Authentication cookies are not used by federation server proxies. For more information about federation server proxies, see Federation Service Proxy (http://go.microsoft.com/fwlink/?LinkId=62784). |

The default value for token lifetime is 600 minutes, or 10 hours. The minimum value is one minute.

| **Caution** |
| --- |
| If you set the token lifetime value too low, the SSO experience of extranet users is degraded. Users will be forced to authenticate again in a forms-based logon as soon as the cookie expires. Only internal corporate users who use integrated authentication will continue to experience single sign-on. |

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the token lifetime for a trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click the **Trust Policy** node, and then click **Properties**.
3. Scroll to the **Advanced** tab, and then click the **Advanced** tab.
4. In **Token lifetime (minutes)**, type or scroll to a new number of minutes, and then click **OK**.

**See Also**
**Concepts**
Change the trust policy refresh period
Change the Windows domain trust cache refresh period
Change the token lifetime for an application

# Change the trust policy refresh period

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Each Active Directory Federation Services (ADFS) component uses the trust policy refresh period to determine how often the component checks for changes to the trust policy, and loads the trust policy if it has changed. The default value for this setting is 60 minutes. The minimum value is five minutes.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the trust policy refresh period**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Trust Policy**, and then click **Properties**.
3. Scroll to the **Advanced** tab, and then click the **Advanced** tab.
4. In **Trust policy refresh period (minutes)**, type or scroll to a new number of minutes, and then click **OK**.

**See Also**
**Concepts**
Change the token lifetime for a Federation Service
Change the Windows domain trust cache refresh period

# Change the Windows domain trust cache refresh period

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

On a server that is running the Federation Service component of Active Directory Federation Services (ADFS), the refresh period of the Windows domain trust cache specifies how often the Federation Service refreshes Windows trust information. The default value is 60 minutes.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the refresh period for the Windows domain trust cache**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Trust Policy**, and then click **Properties**.
3. Scroll to the **Advanced** tab, and then click the **Advanced** tab.
4. In **Windows domain trust cache refresh period (minutes)**, type or scroll to a new number of minutes, and then click **OK**.

**See Also**
**Concepts**
Change the token lifetime for a Federation Service
Change the trust policy refresh period

# Managing Certificates Used by Federation Servers

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Servers that are running the Federation Service component of Active Directory Federation Services (ADFS) are required to use the following types of certificates:

- Secure Sockets Layer (SSL) server authentication certificates: Federation servers use SSL server authentication certificates to secure Web services traffic for communication with Web clients and federation server proxies. If you use a standalone certification authority (CA) specifically for your ADFS deployment, you need to manually request, download, and install SSL certificates. These certificates are requested and installed through the Internet Information Services (IIS) snap-in. For more information about using SSL certificates, see Configuring Secure Sockets Layer (http://go.microsoft.com/fwlink/?LinkId=62785) and Obtaining Server Certificates (http://go.microsoft.com/fwlink/?linkid=62479).

- Token-signing certificates: Each federation server uses a token-signing certificate to digitally sign all security tokens that it produces. A token-signing certificate can be any certificate that has a digital signature key usage (KU), such as a server authentication certificate or code-signing certificate. Enhanced key usage (EKU) is not required for token-signing certificates. For best results, use a certificate other than the SSL server authentication certificate that you installed on the federation server. Token-signing certificates are installed differently, depending on the server farm method, as follows:

    - For the first server in the Federation Service, a token-signing certificate must be requested and installed. You can install a token-signing certificate by connecting to an enterprise CA, a public CA (for example, Verisign), or by creating a self-signed certificate. For information about installing token-signing certificates when using Microsoft Certificate Services as your enterprise CA, see Submit an advanced certificate request via the Web to a Windows Server 2003 CA (http://go.microsoft.com/fwlink/?linkid=64020). For information about installing a token-signing certificate from a public CA, contact your public CA. For information about creating self-signed certificates, see Create a self-signed, token-signing certificate.

    - For additional federation servers within a single server farm, you can reuse the same token-signing certificate by sharing its private key, or you can create a unique certificate for each server. When multiple certificates are used, each server in that farm signs tokens with a unique private key. However, you must configure each server with the public keys from all servers in that farm by adding public keys for all certificates to the trust policy as verification certificates.

- Verification certificates: Verification certificates are the public key portion of the token-signing certificates of federation servers, and are used to ensure that the security token was issued by a trusted federation server and that it was not modified. Each federation server requires a verification certificate for every token-signing certificate that should be accepted when presented to that federation server, including its own token-signing certificate. In a federated scenario, verification certificates are also associated with each account partner to verify tokens that claim to be issued by that partner's federation server(s). In farmed scenarios where each federation server uses a different token-signing certificate, there must be a verification certificate that corresponds to each of those servers. If all token-signing certificates that are issued to a set of federation servers are issued by the same CA, you can use the exported public key of that CA certificate for the verification certificate in the trust policy and for use by partners.

The following tasks for managing certificates on federation servers are described in this objective.

- Managing Token-signing Certificates

- Managing Verification Certificates

- Rolling Over a Token-signing Certificate

**See Also**
**Concepts**
Managing Certificates Used by Federation Server Proxies

**Other Resources**

Understanding Certificates Used by ADFS
Public Key Infrastructure for Windows Server 2003
Public Key Infrastructure
Certificate Services Technical Reference

# Managing Token-signing Certificates

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Servers that are running the Federation Service component of Active Directory Federation Services (ADFS) in an account Federation Service require token-signing certificates to sign security tokens that the servers produce. You can view and change the current certificate as needed. You can also manage the certificate revocation list (CRL) to ensure that only valid certificates are in use in the Federation Service.

**Task requirements**

You need the following to perform the procedures for this task:

- A Federation Service in an account role
- A certification authority or the ability to create self-signed certificates
- Active Directory Federation Services snap-in

To complete this task, perform the following procedures on an as-needed basis:

- Create a self-signed, token-signing certificate
- View the current token-signing certificate
- Turn CRL checking on or off
- Export the public key portion of a token-signing certificate
- Export the private key portion of a token-signing certificate
- Change the token-signing certificate that a federation server uses

# Create a self-signed, token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use the following procedure to create a self-signed, code-signing certificate that also creates and installs a private key. To perform this procedure, use the Makecert.exe utility. Makecert.exe is available in the Microsoft .NET Framework 2.0 Software Development Kit (SDK) (x86) (http://go.microsoft.com/fwlink/?LinkId=79548).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create a self-signed, token-signing certificate**

1. Open a command prompt.
2. Type the appropriate **makecert** syntax.
   - Example command:

**makecert -r -pe -n "CN=CertForADFS" -b 01/01/2006 -e 01/01/2007 -eku 1.3.6.1.5.5.7.3.3 -ss my -sr localMachine -sky exchange -sp "Microsoft RSA SChannel Cryptographic Provider" -sy 12 "CertForADFS.cer"**

| **Note** |
| --- |
| Track certificate expiration dates to make sure that certificates are replaced before they expire. You can do this using the Active Directory Federation Services snap-in for certificates of the current organization and also for partners that you configure in the trust policy. |

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Rolling Over a Client Authentication Certificate

# View the current token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Each Active Directory Federation Services (ADFS) federation server uses a token-signing certificate to digitally sign all security tokens that it produces. You can view the certificate to check its expiration date, revocation status, and other details.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To view the current token-signing certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service**, and then click **Properties**.
3. On the **General** tab, under **Token-signing certificate**, click **View**.
4. On the **General** tab in the **Certificate** property sheet, assess the validity of the certificate by noting the **Valid from** and **to** dates, and whether the certificate has a private key associated with it, and then click **OK** twice.

# Turn CRL checking on or off

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Certificate revocation list (CRL) checking is the process of searching for revoked certificates on a server. Active Directory Federation Services (ADFS) uses the Cryptographic Application Programming Interface (CAPI) to validate token-signing certificates and to verify that they have been revoked. In ADFS, CRL checking is used only for token-signing certificates. It is turned on by default for federation servers and ADFS-enabled Web servers.

You can use the script that is provided at the end of this topic to turn CRL checking on or off on a per-organization basis to meet your specific security requirements. Depending on how the script is configured, it can turn CRL verification on or off for the organization itself, which applies to both the federation server and the Web server, or for an account partner.

| **Caution** |
| --- |
| Turning CRL checking off is not a security best practice, and it has the potential to compromise your ADFS infrastructure. However, some organizations may choose to disable CRL checking or configure it to behave in a certain way. |

Using this script will not affect CRL settings for Secure Sockets Layer (SSL) certificates that are defined in Internet Information Services (IIS) Manager. The IIS metabase uses the MD_CERT_NO_REVOC_CHECK setting to validate SSL client authentication certificates and server authentication certificates that are sent to a Web server. This IIS setting is not used by clients to validate ADFS token-signing certificates or by ADFS servers to validate partner certificates. For more information about how to disable CRL checking for SSL certificates using IIS, see SSL and Certificates (IIS 6.0) (http://go.microsoft.com/fwlink/?LinkId=73717) on the Microsoft Web Site.

| **Note** |
| --- |
| In scenarios in which ADFS servers do not have connectivity to a CRL distribution point on the Internet and the SSL certificates and token-signing certificates that are assigned to those servers were issued by a public certification authority (CA), you may have to disable CRL checking settings in both IIS (using the topic in the previous paragraph) and in ADFS (using this script) on each of the servers. Otherwise, disabling both of these settings simultaneously is not required. |

Before communication between an ADFS server and a Web browser client can occur, an SSL channel must be established between the ADFS server and the client. For more information about establishing an SSL channel, see SSL and Certificates (IIS 6.0) (http://go.microsoft.com/fwlink/?LinkId=73717) on the Microsoft Web Site.

For more information about certificate revocation, see Certificate Revocation and Status Checking (http://go.microsoft.com/fwlink/?linkid=27081) and RevocationFlags Enumeration (http://windowssdk.msdn.microsoft.com/en-us/library/system.web.security.singlesignon.revocationflags.aspx) on the Microsoft Web site.

The following table briefly describes the various arguments that are used in the script.

| | |
| --- | --- |
| Usage | TpCrlChk.vbs TrustPolicy.xml TrustRealmUri RevocationFlags |

| | |
|---|---|
| Arguments | • TrustPolicy.xml—Full path to the trust policy file<br>• TrustRealmUri—Uniform Resource Identifier (URI) of the trust realm whose setting must be changed<br>• RevocationFlags—One of the following:<br>  • None<br>  • CheckEndCert<br>  • CheckEndCertCacheOnly<br>  • CheckChain<br>  • CheckChainCacheOnly<br>  • CheckChainExcludeRoot<br>  • CheckChainExcludeRootCacheOnly |
| Examples | • Cscript TpCrlChk.vbs TrustPolicy.xml 5—Sets the revocation flags to CheckChainExcludeRoot, which is the recommended default.<br>• Cscript TpCrlChk.vbs TrustPolicy.xml 0—Sets the revocation flags to None, which means no revocation checking will be done. |

**Script Text**

```
'Option Explicit

Dim tpf ' Trust policy factory
Dim cf  ' Claim Factory

Dim tpFileName  ' Trust policy file name
Dim trUri       ' TrustRealm Uri
Dim revFlagsStr ' RevocationFlags enum in string form

Dim tp  ' TrustPolicy
Dim tr  ' TrustedRealm
Dim revFlags    ' RevocationFlags enum
Dim found       ' Did we find the realm in the trust policy?

'----------------------------------------
' Echo usage.
'----------------------------------------
Sub Usage()
    WScript.StdErr.WriteLine("Usage:")
    WScript.StdErr.WriteLine("TpCrlChk.vbs TrustPolicy.xml TrustRealmUri
RevocationFlags")
    WScript.StdErr.WriteLine()
    WScript.StdErr.WriteLine("Arguments:")
    WScript.StdErr.WriteLine("TrustPolicy.xml - Full path to the trust policy file")
    WScript.StdErr.WriteLine("TrustRealmUri   - Uri of the trust realm whose setting
must be changed")
    WScript.StdErr.WriteLine("RevocationFlags - One of the following:")
    WScript.StdErr.WriteLine("                      None")
    WScript.StdErr.WriteLine("                      CheckEndCert")
    WScript.StdErr.WriteLine("                      CheckEndCertCacheOnly")
    WScript.StdErr.WriteLine("                      CheckChain")
    WScript.StdErr.WriteLine("                      CheckChainCacheOnly")
    WScript.StdErr.WriteLine("                      CheckChainExcludeRoot")
```

```
    WScript.StdErr.WriteLine("                      CheckChainExcludeRootCacheOnly")
    WScript.Quit
End Sub


'----------------------------------------
' Fetch the RevocationFlags enum value.
'----------------------------------------
Function GetRevFlags(revFlagsStr)
    If (revFlagsStr = "None") Then
        GetRevFlags = 0
    ElseIf (revFlagsStr = "CheckEndCert") Then
        GetRevFlags = 1
    ElseIf (revFlagsStr = "CheckEndCertCacheOnly") Then
        GetRevFlags = 2
    ElseIf (revFlagsStr = "CheckChain") Then
        GetRevFlags = 3
    ElseIf (revFlagsStr = "CheckChainCacheOnly") Then
        GetRevFlags = 4
    ElseIf (revFlagsStr = "CheckChainExcludeRoot") Then
        GetRevFlags = 5
    ElseIf (revFlagsStr = "CheckChainExcludeRootCacheOnly") Then
        GetRevFlags = 6
    Else
        Call Usage()
    End If
End Function



'----------------------------------------
' Get the parameters.
'----------------------------------------

Dim ArgObj
Set ArgObj = WScript.Arguments

If (ArgObj.Count < 3) Then
    Call Usage()
End If

tpFileName = ArgObj.Item (0)
trUri      = ArgObj.Item(1)
revFlags   = GetRevFlags(ArgObj.Item(2))

'----------------------------------------
' Do the job.
'----------------------------------------

WScript.StdOut.WriteLine("Loading trust policy: " & tpFileName)

'
' Create factories
'
Set tpf = CreateObject("System.Web.Security.SingleSignOn.TrustPolicyFactory")
Set cf  = CreateObject("System.Web.Security.SingleSignOn.ClaimFactory")

'
' Load the TrustPolicy
'
Set tp  = tpf.Load(tpFileName, 0) ' initialize certs = false

'
' Find the realm and set the revocation flags
'
found = 0
If (tp.TrustPolicyEntryUri = trUri) Then
    '
    ' Hosted realm attributes
    '
    WScript.StdOut.WriteLine("Changing the setting for this Federation service: " &
trUri)
```

```
    found = 1
    tp.VerificationMethod.RevocationCheckFlags = revFlags
Else
    '
    ' Trusted Realms
    '
    For Each tr in tp.TrustedRealms
        If (tr.TrustPolicyEntryUri = trUri) Then
            WScript.StdOut.WriteLine("Changing the setting for this Account partner: "
& trUri)
            found = 1
            tr.VerificationMethod.RevocationCheckFlags = revFlags
            Exit For 'since the Uri is unique
        End If
    Next

    If (found = 0) Then
        WScript.StdOut.WriteLine("Error: " & trUri & " is neither this Federation
Service nor an Account partner.")
        WScript.Quit
    End If

End If

'----------------------------------------
' Save the TrustPolicy
'----------------------------------------
WScript.StdOut.Write("Saving changed trust policy...")

tp.Write(tpFileName)

WScript.StdOut.WriteLine("done.")
```

# Export the public key portion of a token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

A token-signing certificate is used by an Active Directory Federation Services (ADFS) federation server to digitally sign all security tokens that it produces. Verification certificates are used by the server that receives the token to validate that the security token was issued by a trusted federation server and that the token was not modified. To provide verification certificates to servers that will be processing tokens issued by the trusted federation servers, you can export the public key portion of the token-signing certificate of a federation server that issues the tokens.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To export the public key portion of a token-signing certificate**

1. Click Start, point to Administrative Tools, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service**, and then click **Properties**.
3. On the **General** tab, under **Token-signing certificate**, click **View**.
4. In the **Certificate** dialog box, click the **Details** tab.
5. On the **Details** tab, click **Copy to File**.
6. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
7. On the **Export Private Key** page, make sure that **No, do not export the private key** is selected, and then click **Next**.
8. On the **Export File Format** page, select **DER encoded binary X.509 (.CER)**, and then click **Next**.
9. On the **File to Export** page, specify the certificate file in **File name**, and then click **Next**.

   | **Note** |
   | --- |
   | So that this certificate can be imported to other federation servers as a verification certificate, you will need to securely transfer the file to administrators in your organization and in the partner organization. |

10. On the **Completing the Certificate Export Wizard** page, click **Finish**.
11. Validate success by checking to see that the file you specified was created at the specified location.

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Add a verification certificate to the trust policy
Add a verification certificate to an account partner

# Export the private key portion of a token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Every federation server in an Active Directory Federation Services (ADFS) server farm must have access to the private key of the token-signing certificate. If you are implementing a server farm of federation servers that share a single, exportable private key certificate that is issued by an enterprise certification authority (CA), the private key portion of the existing token-signing certificate must be exported to make it available for importing into the certificate store on the new server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To export the private key of a token-signing certificate**

1. Click Start, point to Administrative Tools, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service**, and then click **Properties**.
3. On the **General** tab, click **View**.
4. In the **Certificate** dialog box, click the **Details** tab.
5. On the **Details** tab, click **Copy to File**.
6. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
7. On the **Export Private Key** page, select **Yes, export the private key**, and then click **Next**.
8. On the **Export File Format** page, select **Personal Information Exchange = PKCS #12 (.PFX)**, and then click **Next**.
9. On the **Password** page, type and confirm the password that is required to share the token-signing certificate. You will need this password when you select the exported token-signing certificate when installing the Federation Service.
10. On the **File to Export** page, specify the certificate file, and then click **Next**.
11. On the **Completing the Certificate Export Wizard** page, click **Finish**.
12. Validate the success of your export by confirming that the file you specified is created at the specified location.

| **Important** |
| --- |
| So that this certificate can be imported to the local certificate store on the new server, you must transfer the file to physical media and protect its security during transport to the new server. It is extremely important to guard the security of the private key. |

13. Import the exported certificate into the certificate store on the new server prior to installing the Federation Service. For information about how to import the certificate, see Import a certificate (http://go.microsoft.com/fwlink/?linkid=20040).

**See Also**
**Concepts**
Implementing a Server Farm of Federation Servers

# Change the token-signing certificate that a federation server uses

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Each Active Directory Federation Services (ADFS) federation server uses a token-signing certificate to digitally sign all security tokens that it produces. Only one token-signing certificate can be in effect on a federation server. If you have installed a new token-signing certificate on a federation server and you want that certificate to be used, you will need to select that certificate in ADFS.

Perform this procedure on the account or resource federation server whose token-signing certificate you want to change.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the token-signing certificate on a federation server**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service**, and then click **Properties**.
3. On the **General** tab, under **Token-signing certificate**, click **Select**.
4. In the **Select Certificate** dialog box, click the token-signing certificate you want to use, and then click **OK**.
5. In the **Federation Service Properties** dialog box, click **OK**.

   In certain cases, the Federation Service might be running in a different account and you will be prompted to allow this account to have access to the private key. Based on your deployment, give access to the private key for this account.
6. In the **Federation Server Configuration** message box, click **Yes** to add the new certificate to the verification certificates in the trust policy.

   If you have already added this certificate to the trust policy as the verification certificate, you will not be prompted to add this new certificate.
7. On the **General** tab, under **Token-signing certificate**, click **View** to check that the selected certificate is being used.

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate

# Managing Verification Certificates

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Servers that are running the Federation Service component of Active Directory Federation Services (ADFS) in a resource Federation Service require verification certificates to validate the security tokens that are produced by any trusted federation servers, including the same federation server.

In a federated scenario, federation servers in a resource partner also re-sign certificates. Therefore, federation servers in an account resource partner must also have verification certificates for certificates that are signed by resource federation servers.

You can add, remove, and view these certificates as needed.

**Task requirements**

You need the following to perform the procedures for this task:

- A federated ADFS deployment
- Exported token-signing certificate
- Active Directory Federation Services snap-in

To complete this task, perform the following procedures on an as-needed basis:

- Add a verification certificate to the trust policy
- Add a verification certificate to an account partner
- View the current verification certificate
- Remove a verification certificate

**See Also**
**Concepts**
Managing Token-signing Certificates
Rolling Over a Token-signing Certificate

# Add a verification certificate to the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When a token-signing certificate is replaced on a server that is running the Federation Service component of Active Directory Federation Services (ADFS), the public key portion of the new token must be added as a verification certificate to federation servers that receive tokens from that Federation Service.

This procedure is not usually necessary because the verification certificate is added to the trust policy automatically when you select a token-signing certificate for use by the server. Use this procedure under the following conditions:

- You are restoring an old trust policy from a version that predates the **Valid from** date of the token-signing certificated.
- You answered "No" to the question of whether to add the verification certificate to the trust policy in the procedure to Change the token-signing certificate that a federation server uses.

### Administrative credentials

To complete the procedure in this topic, you must be a member of the Administrators group on the local computer.

**To add a verification certificate to the trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click **Trust Policy**, and then click **Properties**.
3. Click the **Verification Certificates** tab, and then click **Add**.
4. In the **Browse for Verification Certificate file** dialog box, locate the certificate file that you want to add.
5. Select the certificate file, and then click **Open**.
6. In the **Trust Policy Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Change the token-signing certificate that a federation server uses
Add a verification certificate to an account partner
Export the public key portion of a token-signing certificate
Remove a verification certificate

# Add a verification certificate to an account partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When a token-signing certificate is replaced on a server that is running the Federation Service component of Active Directory Federation Services (ADFS), you must export the new token to a file and then add the token to the trust policy that is used by federation servers that receive tokens from that Federation Service. In addition, if the Federation Service is acting in the account role, the verification certificate must be added to the account node in the respective resource partner.

After you export the public portion of the token-signing certificate, use the following procedure to add it as the verification certificate on the account node in the resource partner.

| **Note** |
| --- |
| To prevent downtime, do not remove the existing verification certificate until the new token-signing certificate has been added on the originating federation server. To prevent the need for re-authentication, wait 10 hours (the default lifetime for an access token) or until all access tokens have expired. |

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a verification certificate to an account partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click Trust Policy, double-click Partner Organizations, double-click Account Partners, right-click the account partner, and then click **Properties**.
3. Click the **Verification Certificates** tab, and then click **Add**.
4. In the **Browse for Verification Certificate file** dialog box, locate the certificate file that you want to add.
5. Select the certificate file, and then click **Open**.
6. In the **Trust Policy Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Add a verification certificate to the trust policy
Export the public key portion of a token-signing certificate
Remove a verification certificate

# View the current verification certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), verification certificates are used to verify that a security token was issued by a valid federation server and that it was not modified. A verification certificate is an exported copy of the token-signing certificate of any trusted federation server.

Use this procedure to view the certificate when you want to see the information about the verification certificate that is currently being used by a federation server. For example, use this procedure if you want to ensure that the certificate is good, view the server that issued the certificate, or view its date of expiration.

**Administrative credentials**

To complete this procedure , you must be a member of the Administrators group on the local computer.

**To view a verification certificate**

1. On a federation server that hosts the Federation Service for which you want to view verification certificates, click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click **Trust Policy**, and then click **Properties**.
3. Click the **Verification Certificates** tab.
4. Click the certificate that you want to view, and then click **View**.

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Add a verification certificate to the trust policy
Change the token-signing certificate that a federation server uses
Remove a verification certificate

# Remove a verification certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

After you add a new verification certificate on a server that is running the Federation Service component of Active Directory Federation Services (ADFS), you must remove the old verification certificate. Removing an old verification certificate is part of the prescribed certificate rollover process by which you replace token-signing and verification certificates in a manner that prevents downtime. Use this procedure in accordance with the instructions in "Rolling Over a Token-signing Certificate."

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To remove a verification certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Remove the certificate as follows:
   - If you are removing the verification certificate from the trust policy on an account or resource federation server, double-click **Federation Service**, right-click the **Trust Policy** node, and then click **Properties**.
   - If you are removing the verification certificate from the account partner on a trusting resource federation server, double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Account Partners**, right-click the account partner node, and then click **Properties**.
3. Click the **Verification Certificates** tab.
4. Click the certificate that you want to remove, click **Remove**, and then click **OK**.

# Rolling Over a Token-signing Certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you need to replace a token-signing certificate on a server that is running the Federation Service component of Active Directory Federation Services (ADFS), either due to expiration or suspicion of tampering, use the procedures in this task to roll over the certificate in a manner that eliminates any significant lapse of certificate validity on the servers.

Configure the public key of every certificate that you replace as a verification certificate on any other servers in the same federation server farm. This configuration occurs automatically when you select the token-signing certificate for use on the server. In addition, if the Federation Service is acting in the account role, the certificate must be configured as a verification certificate for the account partner by any trusting resource partners. To accomplish this configuration, you must export the public key of the certificate to make it available for adding to the account partner.

When the same token-signing certificate is shared among multiple servers in a server farm, the private key of the new token-signing certificate must be exported and then imported on all other servers in the server farm. If a different token-signing certificate is installed on each server in the federation server farm, this additional import step is not required and each server uses a unique private key. The public key is always shared as the verification certificate through the trust policy.

**Task requirements**

You need the following to perform the procedures for this task:

- If you are using self-signed certificates, you can use the makecert.exe utility, which is available for download on the Microsoft Web site.
- Active Directory Federation Services MMC snap-in.

To complete this task, perform the following procedures:

1. Install a new token-signing certificate, as follows:
    - If you are using Microsoft Certificate Services as an enterprise certification authority (CA), obtain a new code signing certificate according to the instructions in "Submit an advanced certificate request via the Web to a Windows Server 2003 CA" (http://go.microsoft.com/fwlink/?linkid=64020. Specify installing the certificate into the local certificate store.
    - If you are using a different enterprise CA or a public CA, follow the instructions provided by the CA.
    - Alternatively, Create a self-signed, token-signing certificate.
2. Configure the private key, if needed, on all servers in a server farm where federation servers use the same private key, as follows:
    - If you are implementing a server farm of federation servers that share a single, exportable private key certificate that is issued by an enterprise certification authority (CA) directly into the certificate store, first use the procedure Export the private key portion of a token-signing certificate to make it available for importing into the local certificate stores of the other servers in the farm. Then, on all other servers in the farm, import the exported certificate into the local store.
    - If you are implementing a server farm of federation servers that share a single, exportable private key certificate that is issued by a public certification authority (CA), import the certificate into the local certificate store.

    For information about how to import the certificate into the certificate store, see Import a certificate (http://go.microsoft.com/fwlink/?linkid=22763).

3. If the server is acting in the account role, do the following:

    a. On the federation server for which you obtained a new token-signing certificate in step 1, use the procedure Export the public key portion of a token-signing certificate to create a file that can be used as a verification certificate.

    b. Provide the exported certificate to the resource partner and instruct the resource partner to perform the procedure Add a verification certificate to an account partner to configure the account partner with the new verification certificate.

    c. Confirm that the partner organization has added the new verification certificate.

4. Select the new verification certificate to Change the token-signing certificate that a federation server uses. Newly issued tokens will now use this certificate. This procedure also adds the verification certificate to the trust policy. Perform this procedure as follows:

    - If you are sharing the same token-signing certificate among all servers in the server farm, perform this procedure on one federation server in the server farm.

    - If you are using separate token-signing certificates for each server in the server farm, perform this procedure on each federation server in the server farm.

5. Inform the partner organizations that it is safe to remove the original verification certificate.

6. On all servers on which you selected the new verification certificate in step 4, remove verification certificates, as follows:

    a. Use the procedure Remove a verification certificate to remove the old verification certificate from the trust policy.

    | **Note** |
    | --- |
    | To prevent the need for users who have already logged on to be authenticated again,, wait 10 hours (the default lifetime of an access token for the Federation Service) or until all access tokens have expired before you remove the existing verification certificate. |

    b. In the case of a Federation Service in the account role, instruct the resource partner to remove the old verification certificate from the account partner in the resource Federation Service. In this case, follow the same procedure but use the **Verification Certificates** tab in the account partner node properties instead of the **Trust Policy** node properties.

7. Delete the old token-signing certificate from the certificate store. For information about how to delete a certificate from a certificate store, see Delete a certificate (http://go.microsoft.com/fwlink/?linkid=62715).

**See Also**
**Concepts**
Rolling Over a Client Authentication Certificate

# Create a self-signed, token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use the following procedure to create a self-signed, code-signing certificate that also creates and installs a private key. To perform this procedure, use the Makecert.exe utility. Makecert.exe is available in the Microsoft .NET Framework 2.0 Software Development Kit (SDK) (x86) (http://go.microsoft.com/fwlink/?LinkId=79548).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create a self-signed, token-signing certificate**

1. Open a command prompt.
2. Type the appropriate **makecert** syntax.
   - Example command:

**makecert -r -pe -n "CN=CertForADFS" -b 01/01/2006 -e 01/01/2007 -eku 1.3.6.1.5.5.7.3.3 -ss my -sr localMachine -sky exchange -sp "Microsoft RSA SChannel Cryptographic Provider" -sy 12 "CertForADFS.cer"**

| **Note** |
| --- |
| Track certificate expiration dates to make sure that certificates are replaced before they expire. You can do this using the Active Directory Federation Services snap-in for certificates of the current organization and also for partners that you configure in the trust policy. |

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Rolling Over a Client Authentication Certificate

# Export the public key portion of a token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

A token-signing certificate is used by an Active Directory Federation Services (ADFS) federation server to digitally sign all security tokens that it produces. Verification certificates are used by the server that receives the token to validate that the security token was issued by a trusted federation server and that the token was not modified. To provide verification certificates to servers that will be processing tokens issued by the trusted federation servers, you can export the public key portion of the token-signing certificate of a federation server that issues the tokens.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To export the public key portion of a token-signing certificate**

1. Click Start, point to Administrative Tools, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service**, and then click **Properties**.
3. On the **General** tab, under **Token-signing certificate**, click **View**.
4. In the **Certificate** dialog box, click the **Details** tab.
5. On the **Details** tab, click **Copy to File**.
6. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
7. On the **Export Private Key** page, make sure that **No, do not export the private key** is selected, and then click **Next**.
8. On the **Export File Format** page, select **DER encoded binary X.509 (.CER)**, and then click **Next**.
9. On the **File to Export** page, specify the certificate file in **File name**, and then click **Next**.

   | **Note** |
   | --- |
   | So that this certificate can be imported to other federation servers as a verification certificate, you will need to securely transfer the file to administrators in your organization and in the partner organization. |

10. On the **Completing the Certificate Export Wizard** page, click **Finish**.
11. Validate success by checking to see that the file you specified was created at the specified location.

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Add a verification certificate to the trust policy
Add a verification certificate to an account partner

# Export the private key portion of a token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Every federation server in an Active Directory Federation Services (ADFS) server farm must have access to the private key of the token-signing certificate. If you are implementing a server farm of federation servers that share a single, exportable private key certificate that is issued by an enterprise certification authority (CA), the private key portion of the existing token-signing certificate must be exported to make it available for importing into the certificate store on the new server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To export the private key of a token-signing certificate**

1. Click Start, point to Administrative Tools, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service**, and then click **Properties**.
3. On the **General** tab, click **View**.
4. In the **Certificate** dialog box, click the **Details** tab.
5. On the **Details** tab, click **Copy to File**.
6. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
7. On the **Export Private Key** page, select **Yes, export the private key**, and then click **Next**.
8. On the **Export File Format** page, select **Personal Information Exchange = PKCS #12 (.PFX)**, and then click **Next**.
9. On the **Password** page, type and confirm the password that is required to share the token-signing certificate. You will need this password when you select the exported token-signing certificate when installing the Federation Service.
10. On the **File to Export** page, specify the certificate file, and then click **Next**.
11. On the **Completing the Certificate Export Wizard** page, click **Finish**.
12. Validate the success of your export by confirming that the file you specified is created at the specified location.

| **Important** |
| --- |
| So that this certificate can be imported to the local certificate store on the new server, you must transfer the file to physical media and protect its security during transport to the new server. It is extremely important to guard the security of the private key. |

13. Import the exported certificate into the certificate store on the new server prior to installing the Federation Service. For information about how to import the certificate, see Import a certificate (http://go.microsoft.com/fwlink/?linkid=20040).

**See Also**
**Concepts**
Implementing a Server Farm of Federation Servers

# Add a verification certificate to an account partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When a token-signing certificate is replaced on a server that is running the Federation Service component of Active Directory Federation Services (ADFS), you must export the new token to a file and then add the token to the trust policy that is used by federation servers that receive tokens from that Federation Service. In addition, if the Federation Service is acting in the account role, the verification certificate must be added to the account node in the respective resource partner.

After you export the public portion of the token-signing certificate, use the following procedure to add it as the verification certificate on the account node in the resource partner.

| **Note** |
| --- |
| To prevent downtime, do not remove the existing verification certificate until the new token-signing certificate has been added on the originating federation server. To prevent the need for re-authentication, wait 10 hours (the default lifetime for an access token) or until all access tokens have expired. |

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a verification certificate to an account partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click Trust Policy, double-click Partner Organizations, double-click Account Partners, right-click the account partner, and then click **Properties**.
3. Click the **Verification Certificates** tab, and then click **Add**.
4. In the **Browse for Verification Certificate file** dialog box, locate the certificate file that you want to add.
5. Select the certificate file, and then click **Open**.
6. In the **Trust Policy Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Add a verification certificate to the trust policy
Export the public key portion of a token-signing certificate
Remove a verification certificate

# Change the token-signing certificate that a federation server uses

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Each Active Directory Federation Services (ADFS) federation server uses a token-signing certificate to digitally sign all security tokens that it produces. Only one token-signing certificate can be in effect on a federation server. If you have installed a new token-signing certificate on a federation server and you want that certificate to be used, you will need to select that certificate in ADFS.

Perform this procedure on the account or resource federation server whose token-signing certificate you want to change.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the token-signing certificate on a federation server**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service**, and then click **Properties**.
3. On the **General** tab, under **Token-signing certificate**, click **Select**.
4. In the **Select Certificate** dialog box, click the token-signing certificate you want to use, and then click **OK**.
5. In the **Federation Service Properties** dialog box, click **OK**.

   In certain cases, the Federation Service might be running in a different account and you will be prompted to allow this account to have access to the private key. Based on your deployment, give access to the private key for this account.
6. In the **Federation Server Configuration** message box, click **Yes** to add the new certificate to the verification certificates in the trust policy.

   If you have already added this certificate to the trust policy as the verification certificate, you will not be prompted to add this new certificate.
7. On the **General** tab, under **Token-signing certificate**, click **View** to check that the selected certificate is being used.

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate

# Remove a verification certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

After you add a new verification certificate on a server that is running the Federation Service component of Active Directory Federation Services (ADFS), you must remove the old verification certificate. Removing an old verification certificate is part of the prescribed certificate rollover process by which you replace token-signing and verification certificates in a manner that prevents downtime. Use this procedure in accordance with the instructions in "Rolling Over a Token-signing Certificate."

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To remove a verification certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Remove the certificate as follows:
   - If you are removing the verification certificate from the trust policy on an account or resource federation server, double-click **Federation Service**, right-click the **Trust Policy** node, and then click **Properties**.
   - If you are removing the verification certificate from the account partner on a trusting resource federation server, double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Account Partners**, right-click the account partner node, and then click **Properties**.
3. Click the **Verification Certificates** tab.
4. Click the certificate that you want to remove, click **Remove**, and then click **OK**.

# Managing the Federation Service Proxy (Optional)

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The Federation Service Proxy component of Active Directory Federation Services (ADFS) functions as a proxy in a perimeter network (also known as demilitarized zone, extranet, or screened subnet) for the Federation Service. The Federation Service Proxy component is not a required component in an ADFS deployment.

The act of installing the Federation Service Proxy component on a computer makes that computer a federation server proxy. Federation server proxies forward requests to federation servers that are not accessible from the Internet.

The following tasks for managing the Federation Service Proxy are described in this objective.

- Managing a Federation Server Proxy Farm
- Managing Certificates Used by Federation Server Proxies

**See Also**
**Concepts**
Managing the Federation Service
Managing ADFS Web Agents and Applications
Backing Up ADFS Components

# Managing a Federation Server Proxy Farm

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

To provide load-balancing of proxy services in larger Active Directory Federation Services (ADFS) deployments, you can install additional federation server proxies to create server farms. When multiple federation server proxies are deployed in the same perimeter network or demilitarized zone, you must configure Federation Service Proxy (FSP) verification certificates in the trust policy for every Federation Service that is protected by the Federation Service Proxy. This configuration implements the federation service proxy farm.

The following tasks for managing a federation server proxy farm are described in this objective.

- Changing the Location of the ADFS Directory
- Adding a New Federation Server Proxy
- Removing a Federation Server Proxy

**See Also**
**Concepts**
Managing a Federation Server Farm

# Changing the Location of the ADFS Directory

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

To change the Active Directory Federation Services (ADFS) directory to another location, you can perform the following procedure on a federation server or federation server proxy. The default location for the directory is C:\ADFS.

| **Important** |
| --- |
| Microsoft does not support changing the location of the ADFS directory. If you alter its location, we cannot guarantee that you can upgrade normally to future versions of Windows operating systems or that ADFS will function after you upgrade. |

To complete this procedure, you must finish all of the following steps in the order that they are presented:

1. Copy the contents of the ADFS directory to the new location.
2. Change the virtual directory path in Internet Information Services (IIS) Manager.
3. Modify the registry to point to the new ADFS directory.

| **Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, back up any valued data on the computer. |

### Administrative credentials

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

#### Copy the contents of the ADFS directory to the new location

1. Start Windows Explorer, and then create a new directory named ADFS in the desired location on the computer.
2. Copy the contents of the C:\ADFS directory (including subfolders) to the new location.

| **Important** |
| --- |
| If you plan to upgrade the Windows Server operating system to a newer version in the future, leave the C:\ADFS directory intact. To do so, copy, rather than move, the directory to the new location. |

#### Change the virtual directory path in IIS Manager

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. Double-click *YourComputerName***(local computer)**, double-click **Web Sites**, and then double-click **Default Web Site**.
3. In the list of virtual directories, right-click **adfs**, and then click **Properties**.
4. On the **Virtual Directory** tab, in **Local path**, type the path to the new ADFS\STS directory, and then click **OK**.

   For example, if you want to move the directory to D:, type **d:\adfs\sts** in **Local path**.

#### Modify the registry to point at the new ADFS directory

1. Open Registry Editor (Regedit.exe).
2. Navigate to the **HKEY_LOCAL_MACHINE_SYSTEM\Microsoft\WebSSO** registry key.
3. Right-click **InitialDirectory**, click **Modify**, under **Value data**, type the path to the new ADFS directory, and then click

**OK**.

For example, if you want to move the directory to D:, then type **d:\adfs** in **Value data**.

4. Close Registry Editor.

# Adding a New Federation Server Proxy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you have an existing Active Directory Federation Services (ADFS) deployment and you want to add a new federation server proxy, you must configure the server as an application server, install and configure certificates, and then install the Federation Server Proxy component of ADFS. You can also set event logging according to server needs.

Unlike farmed federation servers, farmed sets of federation server proxies are not required to share certificate private keys, but must share the certificate public key with the Federation Service that they protect.

However, you must export the public key from each federation server proxy and add this key to the shared trust policy of the Federation Service as an FSP verification certificate.

During installation, you configure the federation server proxy with the DNS host name of the protected Federation Service. This name is added automatically to create the FS URL, which the federation service proxy uses to communicate with the Federation Service.

**Task requirements**

You need the following to perform the procedures for this task:

- An installed Secure Sockets Layer (SSL) certificate. For information about how to acquire SSL certificates, see Obtaining Server Certificates (http://go.microsoft.com/fwlink/?linkid=62479).
- An existing Federation Service.
- The Domain Name System (DNS) host name of the Federation Service that this federation server proxy will protect.
- An installed client authentication certificate for the Federation Service Proxy. For information about installing client authentication certificates when using Microsoft Certificate Services as your enterprise certification authority (CA), see Submit an advanced certificate request via the Web to a Windows Server 2003 CA (http://go.microsoft.com/fwlink/?linkid=64020).

To complete this task, perform the following procedures:

1. Install prerequisite applications
2. Install the Federation Service Proxy on an additional federation server proxy
3. Export the public key portion of a client authentication certificate
4. Add a Federation Service Proxy certificate to the trust policy
5. Configure event logging on a federation server proxy

**See Also**
**Concepts**
Rolling Over a Client Authentication Certificate

# Install prerequisite applications

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) requires that the following applications be installed on a computer running Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Standard Edition, (only for servers hosting the ADFS Web Agent) so that the computer can host the Federation Service, Federation Service Proxy, or ADFS Web Agent components of ADFS:

- Internet Information Services (IIS)
- Microsoft .NET Framework 2.0
- Microsoft ASP.NET (required for claims-aware applications only)

When you add the Application Server Windows component, these three applications are installed.

**Administrative credentials**

To complete this procedure, you must be a member of the local Administrators group.

**To add the Application Server Windows component**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard**, select the **Application Server** check box, and then click **Next**.
4. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Checklist: Installing a federation server
Checklist: Installing a federation server proxy
Checklist: Installing an ADFS-enabled Web server

# Install the Federation Service Proxy on an additional federation server proxy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you install the Federation Service Proxy component of Active Directory Federation Services (ADFS) on an additional server in a federation server proxy server farm, you identify the protected Federation Service by its DNS host name and select the installed client authentication certificate.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To install the Federation Service Proxy component of ADFS on an additional federation server proxy**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, select the **Active Directory Services** check box, and then click **Details**.
4. In the **Active Directory Services** dialog box, select the **Active Directory Federation Services (ADFS)** check box, and then click **Details**.
5. In the **Active Directory Federation Services (ADFS)** dialog box, select the **Federation Service Proxy** check box, and then click **OK**. If a message appears stating that ASP.NET 2.0 was not previously enabled, click **Yes** to enable it, and then click **OK**.
6. In the **Active Directory Services** dialog box, click **OK**.
7. In the Windows Components Wizard, click **Next**.
8. On the **Federation Service** page, click **Select client authentication certificate**, and then click **Select** to select the installed certificate.
9. In the **Select Certificate** dialog box, click the client authentication certificate, and then click **OK**.
10. If you are prompted for the location of the installation files, navigate to *R2 Installation Folder*\cmpnents\r2, and then click **OK**.
11. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Export the public key portion of a client authentication certificate
Add a Federation Service Proxy certificate to the trust policy

# Export the public key portion of a client authentication certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The public key portion of a client authentication certificate for the federation server proxy must be added to the trust policy on a federation server so that the Federation Service can authenticate the federation server proxy. By exporting the public key portion of the client authentication certificate, you create a file that can be imported into the trust policy. You can use the following procedure on the federation server proxy computer to export the public key portion of its client authentication certificate.

**To export the public key portion of a client authentication certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service Proxy**, and then click **Properties**.
3. On the **General** tab, under **FSP client authentication certificate**, click **View**.
4. In the **Certificate** dialog box, click the **Details** tab, and then click **Copy to File**.
5. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
6. On the **Export Private Key** page, ensure that **No, do not export the private key** is selected, and then click **Next**.
7. On the **Export File Format** page, ensure that **DER encoded binary X.509 (.CER)** is selected, and then click **Next**.
8. On the **File to Export** page, type or browse to the location and file name that you want to use for the exported certificate, and then click **Next**.
9. On the **Completing the Certificate Export Wizard** page, verify that the information that you provided is accurate, and then click **Finish**.
10. In the **Certificate Export Wizard** dialog box, click **OK**.
11. In the **Certificate** dialog box, click **OK**.
12. In the **Federation Service Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Certificate requirements for federation server proxies

# Add a Federation Service Proxy certificate to the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

So that the Federation Service can authenticate the federation server proxy, the public key portion of the client authentication certificate for the federation server proxy must be added to the trust policy on a federation server with which the federation server proxy communicates. You can use the following procedure to add the client authentication certificate for the federation server proxy from a file that you have exported.

| **Note** |
| --- |
| The Trust Policy user interface (UI) in the Active Directory Federation Services snap-in refers to client authentication certificates for federation service proxies as **Federation Service Proxy (FSP) certificates**. |

The Federation Service Proxy certificate should chain to a trusted root in the Federation Service. Perform this procedure on a federation server that hosts the trust policy to which you want to add a Federation Service Proxy certificate.

**To add a Federation Service Proxy certificate to the trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Trust Policy**, and then click **Properties**.
3. On the **FSP Certificates** tab, click **Add**.
4. In the **Browse for Federation Service Proxy Certificate file** dialog box, navigate to the certificate file that you want to add, select the certificate file, and then click **Open**.
5. In the **Trust Policy Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Certificate requirements for federation server proxies

# Configure event logging on a federation server proxy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Federation servers log Active Directory Federation Services (ADFS) Federation Service events in the Application and Security event logs. On a federation server proxy, events in the Application log contain additional information about errors regarding contact with the Federation Service. In addition, when a federation server proxy is in effect, the Federation Service events contain information about the proxy certificates that are used.

Use the following procedure to specify the level of events that you want to be logged on a server that is running the Federation Service Proxy. Event logging for a federation server proxy is set in the web.config file. By default, this file is located in %systemdrive%\ADFS\sts. You can apply the following logging types in the Web.config file:

- **Error:** Information about a significant problem of which the user should be aware, usually involving a loss of functionality or data.
- **Warning:** Indicates a problem that is not immediately significant, but that may signify conditions that could cause future issues.
- **Info:** Information about a significant, successful operation.
- **SuccessAudit:** Indicates an audited security event that occurs when an audited access attempt is successful, for example, a successful logon attempt.
- **FailureAudit:** Indicates a security event that occurs when an audited access attempt fails; for example, authentication failed.
- **DetailedSuccess:** A success audit event with detailed information about each token involved in the transaction, including claims information.
- **DetailedFailure:** A failure audit event with detailed information about each token involved in the transaction, including claims information.
- **Everything:** Enables all logging levels.

Use the following procedure to configure event logging levels on a federation server proxy. Perform this procedure on the federation server proxy.

| **Note** |
| --- |
| For security events to be registered, object access auditing must be enabled in the security policy. For more information, see Configuring ADFS Servers for Troubleshooting (http://go.microsoft.com/fwlink/?LinkId=79242). |

### Administrative credentials

To complete this procedure, you must have Read-Write access to the web.config file.

**To configure event logging for a federation server proxy using the web.config file**

1. In Notepad, open the web.config file in %systemdrive%\ADFS\sts.
2. Search for **<logonserver>**.
3. Add the **<auditlevel>** XML element under **<logonserver>**, as follows:

   **<auditlevel>** *Value* **</auditlevel>**

   Where *Value* is one of the following, or the combined values of two or more:

   - Error = 1

- Warning = 2
- Info = 4
- SuccessAudit = 16
- FailureAudit = 32
- DetailedSuccess = 64
- DetailedFailure = 128
- Everything = 247

4. Save and close the web.config file.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy

# Removing a Federation Server Proxy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you no longer need a federation server proxy in an Active Directory Federation Services (ADFS) deployment, remove ADFS and the client authentication certificate.

To complete this task, perform the following procedures:

- Remove the Federation Service Proxy ADFS component, as described in Remove an ADFS component.
- Delete the client authentication certificate from the certificate store. For information about how to delete a certificate from a certificate store, see Delete a certificate (http://go.microsoft.com/fwlink/?linkid=62715).
- Use the procedure Remove a Federation Service Proxy (FSP) certificate from the trust policy to remove the corresponding FSP certificate from any Federation Service trust policy that references it.

# Remove an ADFS component

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you want to remove a server that is running the Federation Service or Federation Service Proxy component of Active Directory Federation Services (ADFS) from a Federation Service, use the following procedure to remove the respective ADFS component.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To remove an ADFS component**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, select the **Active Directory Services** check box, and then click **Details**.
4. In the **Active Directory Services** dialog box, select the **Active Directory Federation Services (ADFS)** check box, and then click **Details**.
5. In the **Active Directory Federation Services (ADFS)** dialog box, clear the **Federation Service** check box or the **Federation Service Proxy** check box, and then click **OK**.
6. In the **Active Directory Services** dialog box, click **OK**.
7. In the Windows Components Wizard, click **Next**.
8. On the **Completing the Windows Components Wizard** page, click **Finish**.

**See Also**
**Concepts**
Remove the ADFS Web Agent component of ADFS

# Remove a Federation Service Proxy (FSP) certificate from the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When a new client authentication certificate has been added to a server that is running the Federation Service Proxy component of Active Directory Federation Services (ADFS), or when you remove a federation server proxy from a server farm, the public portion of the client authentication certificate (the FSP certificate) must be removed from the trust policy of the Federation Service.

Perform the following procedure on a federation server to remove an FSP certificate from the trust policy.

**Administrative credentials**

To complete the procedure in this topic, you must be a member of the Administrators group on the local computer.

**To remove a Federation Service Proxy authentication certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click **Trust Policy**, and then click **Properties**.
3. Click the **FSP Certificates** tab.
4. Click the certificate that you want to remove, and then click **Remove**, and then click **OK**.

# Managing Certificates Used by Federation Server Proxies

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Servers that are running the Federation Service Proxy component of Active Directory Federation Services (ADFS) are required to use the following types of certificates:

- Secure Sockets Layer (SSL) server authentication certificates: Federation server proxies use SSL server authentication certificates to secure Web services traffic for communication with Web clients. Federation server proxies are usually exposed to computers on the Internet that are not included in your enterprise public key infrastructure (PKI). For this reason, you should use a server authentication certificate that is issued by a public (third-party) certification authority (CA) (for example, Verisign). For more information about using SSL certificates, see Configuring Secure Sockets Layer (http://go.microsoft.com/fwlink/?linkid=62785) and Obtaining Server Certificates (http://go.microsoft.com/fwlink/?linkid=62479).

- SSL client authentication certificates: Each federation server proxy uses a client authentication certificate to authenticate to the Federation Service. You can use any certificate with client authentication enhanced key usage (EKU) and that chains to a trusted root CA on the federation server as a client authentication certificate for the federation server proxy. In addition, you must explicitly add the client authentication certificate to the trust policy. However, only the federation server proxy stores the private key that is associated with the federation server proxy client authentication certificate. You can install a client authentication certificate by connecting to an enterprise CA or by creating a self-signed certificate.

  | **Important** |
  | --- |
  | Do not use a certificate that was issued by your enterprise CA for client authentication of an Active Directory user (especially a domain administrator) because the private key is stored on the federation server proxy. Storing such a private key on the federation server proxy allows an administrator or other successful attacker to assume the identity that the certificate represents. |

  For information about installing client authentication certificates when using Microsoft Certificate Services as your enterprise CA, see "Submit an advanced certificate request via the Web to a Windows Server 2003 CA" (http://go.microsoft.com/fwlink/?linkid=64020). For information about creating self-signed certificates, see Create a self-signed, token-signing certificate.

### Task requirements

You need the following to perform the procedures for this task:

- Active Directory Federation Services MMC snap-in

The following procedures for managing certificates on federation server proxies are described in this task. Use these procedures on an as-needed basis.

- Managing Client Authentication Certificates

- Rolling Over a Client Authentication Certificate

**See Also**
**Concepts**
Managing Certificates Used by Federation Servers
**Other Resources**
Understanding Certificates Used by ADFS

# Managing Client Authentication Certificates

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Servers that are running the Federation Service Proxy component of Active Directory Federation Services (ADFS) require Secure Sockets Layer (SSL) client authentication certificates to authenticate to the Federation Service.

**Task requirements**

You need the following to perform the procedures for this task:

- A certification authority or the ability to create self-signed certificates
- A server that is running the Federation Service Proxy component of ADFS
- Active Directory Federation Services snap-in

To complete this task, perform the following procedures on an as-needed basis:

- Create a self-signed, token-signing certificate
- View the current client authentication certificate
- Export the public key portion of a client authentication certificate
- Add a Federation Service Proxy certificate to the trust policy
- Change the client authentication certificate that a federation server proxy uses

**See Also**
**Concepts**
Rolling Over a Client Authentication Certificate

# Create a self-signed, token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use the following procedure to create a self-signed, code-signing certificate that also creates and installs a private key. To perform this procedure, use the Makecert.exe utility. Makecert.exe is available in the Microsoft .NET Framework 2.0 Software Development Kit (SDK) (x86) (http://go.microsoft.com/fwlink/?LinkId=79548).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create a self-signed, token-signing certificate**

1. Open a command prompt.
2. Type the appropriate **makecert** syntax.
   - Example command:

**makecert -r -pe -n "CN=CertForADFS" -b 01/01/2006 -e 01/01/2007 -eku 1.3.6.1.5.5.7.3.3 -ss my -sr localMachine -sky exchange -sp "Microsoft RSA SChannel Cryptographic Provider" -sy 12 "CertForADFS.cer"**

| **Note** |
| --- |
| Track certificate expiration dates to make sure that certificates are replaced before they expire. You can do this using the Active Directory Federation Services snap-in for certificates of the current organization and also for partners that you configure in the trust policy. |

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Rolling Over a Client Authentication Certificate

# View the current client authentication certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

An Active Directory Federation Services (ADFS) federation server proxy requires you to install a client authentication certificate. You can view the certificate if you want to check its expiration date, revocation status, and other details.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To view the current Federation Service Proxy client authentication certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service Proxy**, and then click **Properties**.
3. On the **General** tab, under **FSP client authentication certificate**, click **View**.
4. After viewing the certificate, click **OK** twice.

**See Also**
**Concepts**
Change the client authentication certificate that a federation server proxy uses

# Export the public key portion of a client authentication certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The public key portion of a client authentication certificate for the federation server proxy must be added to the trust policy on a federation server so that the Federation Service can authenticate the federation server proxy. By exporting the public key portion of the client authentication certificate, you create a file that can be imported into the trust policy. You can use the following procedure on the federation server proxy computer to export the public key portion of its client authentication certificate.

**To export the public key portion of a client authentication certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service Proxy**, and then click **Properties**.
3. On the **General** tab, under **FSP client authentication certificate**, click **View**.
4. In the **Certificate** dialog box, click the **Details** tab, and then click **Copy to File**.
5. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
6. On the **Export Private Key** page, ensure that **No, do not export the private key** is selected, and then click **Next**.
7. On the **Export File Format** page, ensure that **DER encoded binary X.509 (.CER)** is selected, and then click **Next**.
8. On the **File to Export** page, type or browse to the location and file name that you want to use for the exported certificate, and then click **Next**.
9. On the **Completing the Certificate Export Wizard** page, verify that the information that you provided is accurate, and then click **Finish**.
10. In the **Certificate Export Wizard** dialog box, click **OK**.
11. In the **Certificate** dialog box, click **OK**.
12. In the **Federation Service Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Certificate requirements for federation server proxies

# Add a Federation Service Proxy certificate to the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

So that the Federation Service can authenticate the federation server proxy, the public key portion of the client authentication certificate for the federation server proxy must be added to the trust policy on a federation server with which the federation server proxy communicates. You can use the following procedure to add the client authentication certificate for the federation server proxy from a file that you have exported.

| **Note** |
| --- |
| The Trust Policy user interface (UI) in the Active Directory Federation Services snap-in refers to client authentication certificates for federation service proxies as **Federation Service Proxy (FSP) certificates**. |

The Federation Service Proxy certificate should chain to a trusted root in the Federation Service. Perform this procedure on a federation server that hosts the trust policy to which you want to add a Federation Service Proxy certificate.

**To add a Federation Service Proxy certificate to the trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Trust Policy**, and then click **Properties**.
3. On the **FSP Certificates** tab, click **Add**.
4. In the **Browse for Federation Service Proxy Certificate file** dialog box, navigate to the certificate file that you want to add, select the certificate file, and then click **Open**.
5. In the **Trust Policy Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Certificate requirements for federation server proxies

# Change the client authentication certificate that a federation server proxy uses

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), the federation server proxy uses a client authentication certificate to communicate securely with the federation server. If you have installed a new client authentication certificate on a federation server proxy and you want that certificate to be used, you will need to select that certificate in ADFS.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the client authentication certificate on a federation server proxy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service Proxy**, and then click **Properties**.
3. On the **General** tab, under **FSP client authentication certificate**, click **Select**.
4. In the **Select Certificate** dialog box, click the client authentication certificate that you want to use, and then click **OK**.
5. In the **Federation Service Proxy Properties** dialog box, click **OK**.
6. If the certificate is not trusted by any of the trusted toot certification authorities, after you click **OK**, a Federation Service message box appears, stating: "The Federation Service is running as an account that does not have access to the certificate private key. Do you want to grant permission to this account to read the private key?" Click **Yes** so that the certificate is trusted.
7. On the **General** tab, under **Token-signing certificate**, click **View** to check that the selected certificate is being used.

# Rolling Over a Client Authentication Certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When a client authentication certificate must be replaced on a server that is running the Federation Service Proxy component of Active Directory Federation Services (ADFS), use the procedures in this task to roll over the certificate in a manner that eliminates any significant lapse of certificate validity on the servers.

After a new client authentication certificate is installed on the first federation server proxy, the public portion of the client authentication certificate must be exported and added to the trust policy for the Federation Service.

**Task requirements**

You need the following to perform the procedures for this task:

- If you are using self-signed certificates, you can use the makecert.exe utility, which you can download from http://go.microsoft.com/fwlink/?linkid=63617.
- Active Directory Federation Services MMC snap-in.

To complete this task, perform the following procedures:

1. Install a new token-signing certificate, as follows:
   a. If you are using Microsoft Certificate Services as an enterprise certification authority (CA), obtain a new client authentication certificate according to the instructions in "Submit an advanced certificate request via the Web to a Windows Server 2003 CA" (http://go.microsoft.com/fwlink/?linkid=64020). Specify installing the certificate into the local certificate store.
   b. If you are using a different enterprise CA or a public CA, follow the instructions provided by the CA.
   c. Alternatively, use the procedure Create a self-signed, token-signing certificate.
2. On the federation server for which you obtained a new client authentication certificate in step 1, use the procedure Export the public key portion of a client authentication certificate to create a file that can be used as an FSP verification certificate.
3. Make the exported certificate file available to the Federation Service administrator who must add it to the trust policy.
4. Instruct the administrator to use the procedure Add a Federation Service Proxy certificate to the trust policy to the trust policy to add the exported certificate to the trust policy of the Federation Service that this federation server proxy is protecting.
5. On a federation server proxy, use the procedure Change the client authentication certificate that a federation server proxy uses to select the new certificate.
6. To remove the old FSP verification certificate from the Federation Service, inform the Federation Service administrator to use the procedure Remove a Federation Service Proxy (FSP) certificate from the trust policy.
7. Delete the old client authentication certificate from the certificate store. For information about how to delete a certificate from a certificate store, see Delete a certificate (http://go.microsoft.com/fwlink/?linkid=62715).

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate

# Create a self-signed, token-signing certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use the following procedure to create a self-signed, code-signing certificate that also creates and installs a private key. To perform this procedure, use the Makecert.exe utility. Makecert.exe is available in the Microsoft .NET Framework 2.0 Software Development Kit (SDK) (x86) (http://go.microsoft.com/fwlink/?LinkId=79548).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create a self-signed, token-signing certificate**

1. Open a command prompt.
2. Type the appropriate **makecert** syntax.
   - Example command:

**makecert -r -pe -n "CN=CertForADFS" -b 01/01/2006 -e 01/01/2007 -eku 1.3.6.1.5.5.7.3.3 -ss my -sr localMachine -sky exchange -sp "Microsoft RSA SChannel Cryptographic Provider" -sy 12 "CertForADFS.cer"**

| **Note** |
| --- |
| Track certificate expiration dates to make sure that certificates are replaced before they expire. You can do this using the Active Directory Federation Services snap-in for certificates of the current organization and also for partners that you configure in the trust policy. |

**See Also**
**Concepts**
Rolling Over a Token-signing Certificate
Rolling Over a Client Authentication Certificate

# Export the public key portion of a client authentication certificate

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The public key portion of a client authentication certificate for the federation server proxy must be added to the trust policy on a federation server so that the Federation Service can authenticate the federation server proxy. By exporting the public key portion of the client authentication certificate, you create a file that can be imported into the trust policy. You can use the following procedure on the federation server proxy computer to export the public key portion of its client authentication certificate.

**To export the public key portion of a client authentication certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service Proxy**, and then click **Properties**.
3. On the **General** tab, under **FSP client authentication certificate**, click **View**.
4. In the **Certificate** dialog box, click the **Details** tab, and then click **Copy to File**.
5. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
6. On the **Export Private Key** page, ensure that **No, do not export the private key** is selected, and then click **Next**.
7. On the **Export File Format** page, ensure that **DER encoded binary X.509 (.CER)** is selected, and then click **Next**.
8. On the **File to Export** page, type or browse to the location and file name that you want to use for the exported certificate, and then click **Next**.
9. On the **Completing the Certificate Export Wizard** page, verify that the information that you provided is accurate, and then click **Finish**.
10. In the **Certificate Export Wizard** dialog box, click **OK**.
11. In the **Certificate** dialog box, click **OK**.
12. In the **Federation Service Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Certificate requirements for federation server proxies

# Add a Federation Service Proxy certificate to the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

So that the Federation Service can authenticate the federation server proxy, the public key portion of the client authentication certificate for the federation server proxy must be added to the trust policy on a federation server with which the federation server proxy communicates. You can use the following procedure to add the client authentication certificate for the federation server proxy from a file that you have exported.

| **Note** |
| --- |
| The Trust Policy user interface (UI) in the Active Directory Federation Services snap-in refers to client authentication certificates for federation service proxies as **Federation Service Proxy (FSP) certificates**. |

The Federation Service Proxy certificate should chain to a trusted root in the Federation Service. Perform this procedure on a federation server that hosts the trust policy to which you want to add a Federation Service Proxy certificate.

**To add a Federation Service Proxy certificate to the trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Trust Policy**, and then click **Properties**.
3. On the **FSP Certificates** tab, click **Add**.
4. In the **Browse for Federation Service Proxy Certificate file** dialog box, navigate to the certificate file that you want to add, select the certificate file, and then click **Open**.
5. In the **Trust Policy Properties** dialog box, click **OK**.

**See Also**
**Concepts**
Checklist: Installing a federation server proxy
Certificate requirements for federation server proxies

# Change the client authentication certificate that a federation server proxy uses

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), the federation server proxy uses a client authentication certificate to communicate securely with the federation server. If you have installed a new client authentication certificate on a federation server proxy and you want that certificate to be used, you will need to select that certificate in ADFS.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the client authentication certificate on a federation server proxy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service Proxy**, and then click **Properties**.
3. On the **General** tab, under **FSP client authentication certificate**, click **Select**.
4. In the **Select Certificate** dialog box, click the client authentication certificate that you want to use, and then click **OK**.
5. In the **Federation Service Proxy Properties** dialog box, click **OK**.
6. If the certificate is not trusted by any of the trusted toot certification authorities, after you click **OK**, a Federation Service message box appears, stating: "The Federation Service is running as an account that does not have access to the certificate private key. Do you want to grant permission to this account to read the private key?" Click **Yes** so that the certificate is trusted.
7. On the **General** tab, under **Token-signing certificate**, click **View** to check that the selected certificate is being used.

# Remove a Federation Service Proxy (FSP) certificate from the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When a new client authentication certificate has been added to a server that is running the Federation Service Proxy component of Active Directory Federation Services (ADFS), or when you remove a federation server proxy from a server farm, the public portion of the client authentication certificate (the FSP certificate) must be removed from the trust policy of the Federation Service.

Perform the following procedure on a federation server to remove an FSP certificate from the trust policy.

**Administrative credentials**

To complete the procedure in this topic, you must be a member of the Administrators group on the local computer.

**To remove a Federation Service Proxy authentication certificate**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click **Trust Policy**, and then click **Properties**.
3. Click the **FSP Certificates** tab.
4. Click the certificate that you want to remove, and then click **Remove**, and then click **OK**.

# Managing ADFS Web Agents and Applications

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

ADFS Web Agents are Active Directory Federation Services (ADFS) components that consume security tokens and either allow or deny a user access to a Web application. ADFS Web Agents are available for two types of Web applications:

- ADFS Web Agent for Windows NT token-based applications: Internet Information Services (IIS) applications that use traditional Windows native authorization mechanisms. This type of application does not use ADFS claims.
- ADFS Web Agent for claims-aware applications: Applications that have been written to use ADFS claims for authorization.

The following tasks for managing ADFS Web Agents and applications are described in this objective.

- Installing and Removing ADFS Web Agents
- Adding and Configuring a Windows NT Token-based Application
- Adding and Configuring a Claims-aware Application
- Managing Security for Web Applications

**See Also**
**Other Resources**
Controlling Access to Web-based Applications

# Installing and Removing ADFS Web Agents

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can add and remove Active Directory Federation Services (ADFS) Web Agents on servers that are running Internet Information Services (IIS).

**Task requirements**

You need the following to perform the procedures of this task:

- To add the ADFS Web Agent for Windows NT token-based application, the Web server must have a Secure Sockets Layer (SSL) server authentication certificate installed.

The following procedures are provided in this task:

- Install the ADFS Web Agent component of ADFS
- Remove the ADFS Web Agent component of ADFS

**See Also**
**Concepts**
Adding and Configuring a Windows NT Token-based Application
Adding and Configuring a Claims-aware Application

# Install the ADFS Web Agent component of ADFS

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can add an Active Directory Federation Services (ADFS) Web Agent to a Web server in a resource partner organization. You use the same procedure to add either the ADFS Web Agent for Windows NT token–based applications or the ADFS Web Agent for claims-aware applications. This action makes that server an ADFS-enabled Web server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add an ADFS Web Agent to a Web server**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard**, select the **Active Directory Services** check box, and then click **Details**.
4. In the **Active Directory Services** dialog box, select the **Active Directory Federation Services (ADFS)** check box, and then click **Details**.
5. In the **Active Directory Federation Services (ADFS)** dialog box, click **ADFS Web Agents**, and then click **Details**.
6. In the **ADFS Web Agents** dialog box, select Web Agent subcomponents as follows, and then click **OK**:
   - To add ADFS Web Agent support for applications that can process ADFS claims directly for authorization, select the **Claims-aware applications** check box.

     | **Note** |
     | --- |
     | If your Web server will host only claims-aware applications, do not select the **Windows NT token–based applications** check box. |

   - To add ADFS Web Agent support for applications that require Windows native authorization mechanisms (that is, for applications that do not directly consume ADFS claims), select the **Windows NT token–based applications** check box.

     | **Note** |
     | --- |
     | If your Web server will host only Windows NT token–based applications, do not select the **Claims-aware applications** check box. |

7. In the **Active Directory Federation Services (ADFS)** dialog box, click **OK**.
8. In the **Active Directory Services** dialog box, click **OK**.
9. In the Windows Components Wizard, click **Next**.
10. If you are prompted for the location of installation files, navigate to *R2 installation files*\cmpnents\r2, and then click **OK**.
11. On the **Completing the Windows Components Wizard** page, click **Finish**.

# Remove the ADFS Web Agent component of ADFS

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you no longer want to use Active Directory Federation Services (ADFS) for authenticating and authorizing a Windows NT token-based application or claims-aware application, or both, you can remove one or both ADFS Web Agents from a Web server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To remove an ADFS Web Agent**

1. Click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In **Add or Remove Programs**, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard**, select the **Active Directory Services** check box, and then click **Details**.
4. In the **Active Directory Services** dialog box, select the **Active Directory Federation Services (ADFS)** check box, and then click **Details**.
5. In the **Active Directory Federation Services (ADFS)** dialog box, click **ADFS Web Agents**, and then click **Details**.
6. In the **ADFS Web Agents** dialog box, clear the Web Agent subcomponent check box for the Web agent subcomponent or subcomponents that you want to remove, and then click **OK**.
7. In the **Active Directory Federation Services (ADFS)** dialog box, click **OK**.
8. In the **Active Directory Services** dialog box, click **OK**, and then, in the **Windows Components Wizard**, click **Next**.
9. On the **Completing the Windows Components Wizard** page, click **Finish**.

# Adding and Configuring a Windows NT Token-based Application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have applications that use the Active Directory Federation Services (ADFS) Web Agent for Windows NT token-based applications for authorization, you can configure them in Internet Information Services (IIS) Manager. Use the ADFS Web Agent tab in the properties pages for the following IIS nodes:

- Web Sites node: Use this node to configure the Federation Service URL. Clients are directed to the Federation Service URL when they contact the associated Federation Service.
- *SiteName* nodes and virtual directories: Use this node to enable the ADFS Web agent for Windows NT token-based applications and to configure cookies and the Return URL. The Return URL directs clients to the application page after successful authentication.

**Task requirements**

You need the following to perform the procedures for this task:

- Active Directory Federation Services MMC snap-in
- Internet Information Services (IIS) Manager MMC snap-in
- Regedit.exe

To complete this task, perform the following procedures:

1. Add a new Windows NT token-based application to the Federation Service
2. Enable or disable a Web application
3. Set the cookie path for a Windows NT token-based application
4. Set the cookie domain for a Windows NT token-based application
5. Set the Federation Service URL for a Windows NT token-based application
6. Set the return URL for a Windows NT token-based application
7. Configure event logging for a Windows NT token-based application
8. Configure authentication methods for a federated application

**See Also**
**Concepts**
Adding and Configuring a Claims-aware Application

# Add a new Windows NT token-based application to the Federation Service

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you configure your Active Directory Federation Services (ADFS)–enabled Web server to host a Windows NT token–based application, you must add a reference to that application in the trust policy so that you can specify application properties for the Federation Service. You can use the following procedure to add a new Windows NT token–based application to the Federation Service. You perform this procedure on a federation server in the resource partner.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a Windows NT token–based application to the Federation Service**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Applications**, point to **New**, and then click **Application**.
3. On the **Welcome to the Add Application Wizard** page, click **Next**.
4. On the **Application Type** page, click **Windows NT token–based application**, and then click **Next**.
5. On the **Application Details** page, in **Application display name**, type the name of the application.
6. In **Application URL**, type the Uniform Resource Locator (URL) for your application—for example, https://www.treyresearch.net/*ApplicationName*/—and then click **Next**.

   | **Note** |
   | --- |
   | The value in **Application URL** on this property page must match precisely the **Return URL** value that you specify in the **ADFS Web Agent** tab for the Windows NT token-based application (on the ADFS-enabled Web server). |

7. On the **Accepted Identity Claim** page, click one of the following identity claim types, based on the needs of your application, and then click **Next**:
   - **User principal name (UPN)**
   - **E-mail**
8. On the **Enable this Application** page, ensure that the **Enable this application** check box is selected, and then click **Next**.
9. On the **Completing the Add Application Wizard** page, click **Finish**.

# Enable or disable a Web application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Web applications are enabled by default when you create them in the Active Directory Federation Services (ADFS) snap-in. You can disable an application if it is currently enabled. If you disable an application, access to the application is prevented.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable or disable a Web application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Applications**.
3. Right-click the application you want to enable or disable and do one of the following:
   - If the application is disabled and you want to enable it, click **Enable**.
   - If the application is enabled and you want to disable it, click **Disable**.

| **Note** |
| --- |
| You can also perform this procedure by right-clicking the application, clicking **Properties**, and then selecting or clearing the **Enabled** check box. |

**See Also**
**Concepts**
Configure authentication methods for a federated application

# Set the cookie path for a Windows NT token-based application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The cookie path specifies the location within an Internet Information Services (IIS) Web site virtual directory for which cookies will be sent in response to an application request by an Active Directory Federation Service (ADFS) client. If no cookie path is set, the default path points to a location that depends on whether a cookie domain is set, as follows:

- Cookie domain is set: The path continues from / below the domain that applies for the requested address on the basis of the cookie domain.
- No cookie domain is set: The path continues from / below the domain specified in the Federation Service Uniform Resource Identifier (URI).

If you specify a cookie path (for example, /test01), the cookie is sent for requests under that path (for example, /test01, /test01/index.html, /test01/test05/, and so on).

The cookie path setting is configured in the ADFS Web Agent for an application. For a Windows NT token-based application, configure the ADFS Web Agent in IIS.

Perform this procedure on a Web server in the resource Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To set the cookie path for a Windows NT token-based application**

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. Double-click *YourComputerName* **(local computer)**, and then double-click **Web Sites**.
3. Right-click your Web site, and then click **Properties**.
4. On the **ADFS Web Agent** tab, in **Cookie path**, change the path as needed, and then click **OK**.

**See Also**
**Concepts**
Set the cookie domain for a Windows NT token-based application

# Set the cookie domain for a Windows NT token-based application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use the cookie domain setting to share an application at a higher level than the domain level that is specified in the Federation Service Uniform Resource Identifier (URI). In this way, you can expand the scope of requests for which a cookie will be sent.

If you do not configure a cookie domain, cookies are sent only for requests where the domain that is specified matches the domain in the Federation Service Uniform Resource Identifier (URI). For example, if no cookie domain is set and the domain in the Federation Service URI is Sales.Adatum.com, cookies are sent for only requests where the request URL matches Sales.Adatum.com. However, if you set Adatum.com as the cookie domain, cookies are sent for Sales.Adatum.com plus requests for any other domain with the suffix Adatum.com. For example, cookies are also sent for Northwest.Adatum.com.

The cookie domain setting is configured in the ADFS Web Agent for an application. For a Windows NT token-based application, configure the ADFS Web Agent in Internet Information Services (IIS).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To set the cookie domain for a Windows NT token-based application**

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. Double-click *YourComputerName* **(local computer)**, and then double-click **Web Sites**.
3. Right-click your Web site, and then click **Properties**.
4. On the **ADFS Web Agent** tab, in **Cookie domain**, change the domain as needed, and then click **OK**.

**See Also**
**Concepts**
Set the cookie path for a Windows NT token-based application

# Set the Federation Service URL for a Windows NT token-based application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The Federation Service Uniform Resource Locator (URL) in Active Directory Federation Services (ADFS) defines the URL that is used for all Web sites and Web applications on a Web server where an ADFS Web Agent is enabled. If the Federation Service URL for a Web server that hosts a Windows NT token-based application changes, you must update the ADFS Web Agent in Internet Information Services (IIS).

Perform this procedure on a Web server in the resource Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To set the Federation Service URL for a Windows NT token-based application**

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. Double-click *YourComputerName* **(local computer)**, right-click **Web Sites**, and then click **Properties**.
3. On the **ADFS Web Agent** tab, in **Federation Service URL**, change the URL as needed, and then click **OK**.

# Set the return URL for a Windows NT token-based application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The return Uniform Resource Locator (URL) in Active Directory Federation Services (ADFS) is configured on the ADFS Web Agent on a Web server to reflect the URL that opens a requested Web application. If this URL changes for a Windows NT token-based application, you must change its value in Internet Information Services (IIS).

The return URL must match the application URL that you configure in the properties of the application in the resource Federation Service. Therefore, if you change one, you must change the other.

Perform this procedure on an ADFS Web server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To set the return URL for a Windows NT token-based application**

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, double-click *WebServerName* **(local computer)**, double-click **Web Sites**, right-click any application or site (typically, the **Default Web Site**), and then click **Properties**.
3. On the **ADFS Web Agent** tab, in **Return URL (Uniform Resource Locator)**, type the URL for the application, and then click **OK**.
4. If you have not already done so, change the application URL to match the new value of the return URL.

**See Also**
**Concepts**
Set the application URL for an application

# Set the application URL for an application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The application Uniform Resource Locator (URL) provides the path to the Web application in an Active Directory Federation Services (ADFS) deployment. If the location of the Web application changes, you must change this URL to reflect the new location.

The application URL must match the return URL that you configure in the ADFS Web Agent. Therefore, if you change one, you must change the other.

Perform the following procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**ToTo change the application URL**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Applications**.
3. Right-click the application whose URL you want to change, and then click **Properties**.
4. On the **General** tab, in **Application URL**, select the portion of the URL that you want to change, type the new value, and then click **OK**.
5. If you have not already done so, change the return URL to match the new value of the application URL, using the appropriate procedure for the application type.

**See Also**
**Concepts**
Set the return URL for a Windows NT token-based application
Set the return URL for a claims-aware application

# Configure event logging for a Windows NT token-based application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

On a Web server that is running the Active Directory Federation Services (ADFS) Web Agent for Windows NT token–based applications, you can configure the types of events that you want to be logged for Windows NT token–based applications.

| **Note** |
| --- |
| Event logging is enabled differently for Windows NT token–based applications and claims-aware applications. |

Use the following procedure to specify the types of events that you want to be logged for Windows NT token–based applications in the Application event log on the ADFS-enabled Web server. You set event logging for Windows NT token–based applications in the registry of the Web server.

| **Caution** |
| --- |
| Editing the registry incorrectly may damage your system severely. Before you make changes to the registry, back up any valued data on the computer. |

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

#### To configure event logging for Windows NT token–based applications

1. Open Regedit. Click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Navigate to:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ifssvc\ Parameters**
3. Right-click **Parameters**, click **New**, and then click **DWORD Value**.
4. In the new value file name box, type the following, and then press Enter:

   **ADFSEvent**
5. Double-click the new entry, and then, in **Value data**, provide a value for one of the following levels, or add values to configure multiple levels, and then click **OK**:
   - Warning: 0x01
   - Information: 0x02
   - Success: 0x04
   - Failure: 0x08

# Configure authentication methods for a federated application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

By default, a Web application that you configure in Active Directory Federation Services (ADFS) accepts any of the following authentication methods:

- Windows Integrated authentication (also referred to as Windows NT Challenge/Response authentication): A secure form of authentication because the user name and password are hashed before they are sent across the network.

- User name and password authentication (also called Basic authentication): A widely used, industry-standard method for transmitting user names and passwords across the network in an unencrypted form.

- Certificate or Secure Sockets Layer / Transport Layer Security (SSL/TLS) client authentication: A form of digital identification for your server and for clients that request information from your server. Its function is similar to that of a passport, or other official identity card, that identifies the person carrying it. Certificates are part of the SSL features in Internet Information Services (IIS) that establish a secure connection for sensitive information.

If you want to use a specific authentication method or methods, you can use the following procedure to select one or more methods that will apply to users of the application. You perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure authentication methods for a federated application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Applications**, right-click the application whose authentication methods you want to change, and then click **Properties**.
3. On the **Authentication Methods** tab, to select from the available authentication methods, do one of the following:
   - To allow all possible authentication methods in the list, select the **Any** check box.
   - To allow one or more of the available authentication methods, clear the **Any** check box, and then select one or more of the available methods.
4. When you finish selecting authentication methods, click **OK**.

# Adding and Configuring a Claims-aware Application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can configure the Active Directory Federation Services (ADFS) Web Agent for claims-aware applications by using the Web.config file in the Web application directory (typically \Inetpub\wwwroot*virtualDirectoryName*) on the Web server.

**Task requirements**

You need the following to perform the procedures for this task:

- You must have access to trust policy properties for the resource Federation Service.
- You must have access to the Web.config file for the Web application.
- Active Directory Federation Services MMC snap-in
- Internet Information Services (IIS) Manager MMC snap-in
- Notepad or other text editor

To complete this task, perform the following procedures:

- Add a new claims-aware application to the Federation Service
- Enable or disable a Web application
- Set the cookie path for a claims-aware application
- Set the cookie domain for a claims-aware application
- Set the Federation Service URL for a claims-aware application
- Set the return URL for a claims-aware application
- Set the application URL for an application
- Configure event logging for a claims-aware application
- Configure authentication methods for a federated application

# Add a new claims-aware application to the Federation Service

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you configure your Active Directory Federation Services (ADFS)–enabled Web server to host a claims-aware application, you must add a reference to that application in the trust policy so that you can specify application properties for the Federation Service. You can use the following procedure to add a new claims-aware application to the Federation Service. You perform this procedure on a federation server in the resource partner.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new claims-aware application to the Federation Service**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Applications**, point to **New**, and then click **Application**.
3. On the **Welcome to the Add Application Wizard** page, click **Next**.
4. On the **Application Type** page, click **Claims-aware application**, and then click **Next**.
5. On the **Application Details** page, in **Application display name**, type the name of the application.
6. In **Application URL**, type the Uniform Resource Locator (URL) for your application—for example, https://www.treyresearch.net/*ApplicationName*/—and then click **Next**.

   | **Note** |
   | --- |
   | The value in **Application URL** on this property page must match precisely the Return URL value (in **<returnurl> </returnurl>**) that you specify in the web.config file of the claims-aware application (on the ADFS-enabled Web server). |

7. On the **Accepted Identity Claims** page, select the check boxes for the claim types that will be recognized by your claims-aware application, and then click **Next**.
8. On the **Enable this Application** page, ensure that the **Enable this application** check box is selected, and then click **Next**.
9. On the **Completing the Add Application Wizard** page, click **Finish**.

# Enable or disable a Web application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Web applications are enabled by default when you create them in the Active Directory Federation Services (ADFS) snap-in. You can disable an application if it is currently enabled. If you disable an application, access to the application is prevented.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable or disable a Web application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Applications**.
3. Right-click the application you want to enable or disable and do one of the following:
   - If the application is disabled and you want to enable it, click **Enable**.
   - If the application is enabled and you want to disable it, click **Disable**.

| **Note** |
| --- |
| You can also perform this procedure by right-clicking the application, clicking **Properties**, and then selecting or clearing the **Enabled** check box. |

**See Also**
**Concepts**
Configure authentication methods for a federated application

# Set the Federation Service URL for a claims-aware application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The Federation Service Uniform Resource Locator (URL) in Active Directory Federation Services (ADFS) defines the URL that is used for all Web sites and Web applications on a federation Web server where the ADFS Web Agent is enabled. If the Federation Service URL for a Web server that hosts a claims-aware application changes, you must update the ADFS Web Agent by editing the Web.config file for the application.

**Administrative credentials**

To complete this procedure, you must have read-write access to the Web.config file.

**To set the Federation Service URL for a claims-aware application using the Web.config file**

1. In Notepad or other text editor, open the Web.config file that is in the Web application directory (typically \Inetpub\wwwroot*ApplicationName*) on the Web server.
2. Search for **<fs>**.
3. In the current entry (for example, https://adfsresource.trey.net/adfs/fs/federationserverservice.asmx), select the domain name portion of the URL, and then replace the name with a different fully qualified domain name.
4. Save and close the Web.config file.

# Set the return URL for a claims-aware application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The return Uniform Resource Locator (URL) in Active Directory Federation Services (ADFS) that you configure in the ADFS Web Agent on a Web server reflects the URL that opens a requested Web application. If this URL changes for a claims-aware application, you must change its value in the Web.config file.

The return URL must match the application URL that you configure in the properties of the application in the resource Federation Service. Therefore, if you change one, you must change the other.

### Administrative credentials

To complete this procedure, you must have read-write access to the Web.config file.

**To set the return URL for a claims-aware application using the Web.config file**

1. In Notepad or other text editor, open the Web.config file that is in the Web application directory (typically \Inetpub\wwwroot*ApplicationName*) on the Web server.
2. Search for **<returnurl>**.
3. In the current entry, select the portion of the URL that is changing, and then type the new value.
4. Save and close the Web.config file.
5. If you have not already done so, change the application URL to match the new value of the return URL.

**See Also**
**Concepts**
Set the application URL for an application

# Set the application URL for an application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The application Uniform Resource Locator (URL) provides the path to the Web application in an Active Directory Federation Services (ADFS) deployment. If the location of the Web application changes, you must change this URL to reflect the new location.

The application URL must match the return URL that you configure in the ADFS Web Agent. Therefore, if you change one, you must change the other.

Perform the following procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**ToTo change the application URL**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Applications**.
3. Right-click the application whose URL you want to change, and then click **Properties**.
4. On the **General** tab, in **Application URL**, select the portion of the URL that you want to change, type the new value, and then click **OK**.
5. If you have not already done so, change the return URL to match the new value of the application URL, using the appropriate procedure for the application type.

**See Also**
**Concepts**
Set the return URL for a Windows NT token-based application
Set the return URL for a claims-aware application

# Set the cookie path for a claims-aware application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The cookie path specifies the location within an Internet Information Services (IIS) Web site virtual directory for which cookies will be sent in response to an application request by an Active Directory Federation Service (ADFS) client. If no cookie path is set, the path defaults to a location that depends on whether a cookie domain is set, as follows:

- Cookie domain is set: / below the domain that applies for the requested address on the basis of the cookie domain.
- No cookie domain is set: / below the domain specified in the Federation Service Uniform Resource Identifier (URI).

If you specify a cookie path (for example, /test01), the cookie is sent for requests under that path (for example, /test01, /test01/index.html, /test01/test05/, and so on).

The cookie path setting is configured in the ADFS Web Agent for an application. For a claims-aware application, configure the ADFS Web Agent in the Web.config file.

**Administrative credentials**

To complete this procedure, you must have read-write access to the Web.config file.

**To set the cookie path for a claims-aware application using the Web.config file**

1. In Notepad or other text editor, open the Web.config file that is in the Web application directory (typically \Inetpub\wwwroot*ApplicationName*) on the Web server.
2. Search for the opening tag string **<cookies**.
3. In the current entry for **<path>**, select the existing path and replace it by typing the new path. If the path entry does not exist, type a new one (for example: **<path>https://**).
4. Save and close the Web.config file.

**See Also**
**Concepts**
Set the cookie domain for a claims-aware application

# Set the cookie domain for a claims-aware application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

On a Web server that is running the Active Directory Federation Services (ADFS) Web Agent for claims-aware applications, you can use the cookie domain setting to share an application at a higher level than the domain level that is specified in the Federation Service Uniform Resource Identifier (URI). In this way, you can expand the scope of requests for which a cookie will be sent.

If you do not configure a cookie domain, cookies are sent for only requests where the domain that is specified matches the domain in the Federation Service Uniform Resource Identifier (URI). For example, if no cookie domain is set and the domain in the Federation Service URI is Sales.Adatum.com, cookies are sent for only requests where the request URL matches Sales.Adatum.com. However, if you set Adatum.com as the cookie domain, cookies are sent for Sales.Adatum.com plus requests for any other domain with the suffix Adatum.com. For example, cookies are also sent for Northwest.Adatum.com.

The cookie domain setting is configured in the ADFS Web Agent for an application. For a claims-aware application, configure the ADFS Web Agent in the Web.config file.

**Administrative credentials**

To complete this procedure, you must have read-write access to the Web.config file.

**To set the cookie domain for a claims-aware application using the Web.config file**

1. In Notepad or other text editor, open the Web.config file that is in the Web application directory (typically \Inetpub\wwwroot*ApplicationName*) on the Web server.
2. Search for the opening tag string **<cookies**.
3. In the current entry for **<domain>**, select the existing domain and replace it by typing the new domain name. If the **<domain>** entry does not exist, add the entry within the **<cookies>** tag, as follows:

   **<domain>** *DNSDomainName* **</domain>**
4. Save and close the Web.config file.

**See Also**
**Concepts**
Set the cookie path for a claims-aware application

# Configure event logging for a claims-aware application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Use the following procedure to specify the level of events that you want to be logged for claims-aware applications in the Application event log on the Web server that is protected by Active Directory Federation Services (ADFS).You set event logging for claims-aware applications in the Web.config file for the application.

You can apply the following event logging settings in the web.config file:

- **Error:** An alert that provides information about a significant problem of which the user should be aware, usually involving a loss of functionality or data.
- **Warning:** An alert that indicates a problem that is not immediately significant but that may signify conditions that could cause future issues.
- **Info:** An alert that provides information about a significant, successful operation.
- **SuccessAudit:** A security event that occurs when an audited access attempt succeeds, for example, a successful logon attempt.
- **FailureAudit:** A security event that occurs when an audited access attempt fails, for example, a failed attempt to open a file.
- **DetailedSuccess:** A success audit event that provides detailed information about each token that is involved in the transaction, including claims information.
- **DetailedFailure:** A failure audit event that provides detailed information about each token that is involved in the transaction, including claims information.
- **Everything:** Enables all logging levels.

### Administrative credentials

To complete this procedure, you must have Read-Write access to the Web.config file.

#### To configure event logging for a claims-aware application

1. In Notepad, open the Web.config file in the directory that stores the claims-aware application.
2. Search the file for <websso>.
3. Add or change the **<auditlevel>** entry under **<websso>**, as follows:

   **<auditlevel>** *Value* **</auditlevel>**

   Where *Value* is one of the following or the combined values of two or more:

   - Error = 1
   - Warning = 2
   - Info = 4
   - SuccessAudit = 16
   - FailureAudit = 32
   - DetailedSuccess = 64

- DetailedFailure = 128
- Everything = 247

4. Save and close the web.config file.

# Configure authentication methods for a federated application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

By default, a Web application that you configure in Active Directory Federation Services (ADFS) accepts any of the following authentication methods:

- Windows Integrated authentication (also referred to as Windows NT Challenge/Response authentication): A secure form of authentication because the user name and password are hashed before they are sent across the network.

- User name and password authentication (also called Basic authentication): A widely used, industry-standard method for transmitting user names and passwords across the network in an unencrypted form.

- Certificate or Secure Sockets Layer / Transport Layer Security (SSL/TLS) client authentication: A form of digital identification for your server and for clients that request information from your server. Its function is similar to that of a passport, or other official identity card, that identifies the person carrying it. Certificates are part of the SSL features in Internet Information Services (IIS) that establish a secure connection for sensitive information.

If you want to use a specific authentication method or methods, you can use the following procedure to select one or more methods that will apply to users of the application. You perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure authentication methods for a federated application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Applications**, right-click the application whose authentication methods you want to change, and then click **Properties**.
3. On the **Authentication Methods** tab, to select from the available authentication methods, do one of the following:
   - To allow all possible authentication methods in the list, select the **Any** check box.
   - To allow one or more of the available authentication methods, clear the **Any** check box, and then select one or more of the available methods.
4. When you finish selecting authentication methods, click **OK**.

# Managing Security for Web Applications

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use a common set of procedures to manage security for Web applications that are protected by Active Directory Federation Services (ADFS) Web Agents for both Windows NT token-based applications and claims-aware applications.

**Task requirements**

You need the following to perform the procedures for this task:

- You must have the following information about how to control access to the application:
    - The appropriate authentication methods.
    - The security token protection methods.
    - The appropriate security token lifetime.
- Active Directory Federation Services (ADFS) MMC snap-in.
- The Web.config file located in %systemdrive%\ADFS\sts on a federation server.

To complete this task, perform the following procedures on an as-needed basis:

- Configure authentication methods for a federated application
- Configure the security token protection method for a federated application
- Change the token lifetime for an application
- Configure a policy page for a Web site

# Configure authentication methods for a federated application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

By default, a Web application that you configure in Active Directory Federation Services (ADFS) accepts any of the following authentication methods:

- Windows Integrated authentication (also referred to as Windows NT Challenge/Response authentication): A secure form of authentication because the user name and password are hashed before they are sent across the network.

- User name and password authentication (also called Basic authentication): A widely used, industry-standard method for transmitting user names and passwords across the network in an unencrypted form.

- Certificate or Secure Sockets Layer / Transport Layer Security (SSL/TLS) client authentication: A form of digital identification for your server and for clients that request information from your server. Its function is similar to that of a passport, or other official identity card, that identifies the person carrying it. Certificates are part of the SSL features in Internet Information Services (IIS) that establish a secure connection for sensitive information.

If you want to use a specific authentication method or methods, you can use the following procedure to select one or more methods that will apply to users of the application. You perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure authentication methods for a federated application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Applications**, right-click the application whose authentication methods you want to change, and then click **Properties**.
3. On the **Authentication Methods** tab, to select from the available authentication methods, do one of the following:
   - To allow all possible authentication methods in the list, select the **Any** check box.
   - To allow one or more of the available authentication methods, clear the **Any** check box, and then select one or more of the available methods.
4. When you finish selecting authentication methods, click **OK**.

# Configure the security token protection method for a federated application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Security tokens that are received by Active Directory Federation Services (ADFS) federation servers are protected during transit by one of two methods:

- Public Key Infrastructure (PKI): A PKI is implemented as a hierarchy of certification authorities that verify identities. When a PKI is in place, a signature is embedded into the token that protects the token from tampering.

- Domain service account: A domain service account, which is identified by a service principal name (SPN), runs under an account that is trusted for delegation and that can impersonate a client to gain access to resources. By default, this account is the Internet Information Services (IIS) application pool identity that hosts a claims-aware application and the identity of the ADFS Web Agent Authentication Service that hosts a Windows NT token–based application. When a token is transferred in a domain service account with this setting, the token contains a binary Kerberos V5 signature for the configured SPN. This signature protects the token from tampering.

You can use the following procedure to change the security token protection method for a federated application on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the security token protection method for a federated application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Trust Policy**, double-click **My Organization**, and then double-click **Applications**.
3. Right-click the application whose security token protection method you want to change, and then click **Properties**.
4. On the **General** tab, under **Security token protection method**, do one of the following, and then click **OK**:
   - If your deployment uses certificates that are issued by a certification authority (CA), select **Public Key Infrastucture (PKI)**.
   - If your deployment does not use certificates that are issued by a CA, select **Domain service account** and then, in **service principal name (SPN) of service account**, type the SPN of the account.

**See Also**
**Concepts**
Determine your security token protection method

# Change the token lifetime for an application

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The token lifetime for an Active Directory Federation Services (ADFS) Web application indicates the time during which the security token contained in an ADFS authentication cookie remains in effect for the application. You can change this value in the application properties.

Perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the token lifetime for an application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Trust Policy**, double-click **My Organization**, and then double-click **Applications**.
3. Right-click the application whose token lifetime you want to change, and then click **Properties**.
4. On the **Advanced** tab, in the **Token lifetime (minutes)** box, type or scroll to a new number of minutes, and then click **OK**.

**See Also**
**Concepts**
Change the token lifetime for a Federation Service

# Configure a policy page for a Web site

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In a Web site that is protected by Active Directory Federation Services (ADFS), you can configure a policy page that will be provided to users who request a token for some or all resource partners. The policy page notifies the user of special protections that are in place. To configure the policy page, use the sample policy.aspx file in %systemdrive%\ADFS\SampleFiles to create an .aspx page that contains the appropriate text, and then use the following procedure to configure the Web.config file to open the page during logon.

Perform this procedure on an account federation server or account federation server proxy.

**Administrative credentials**

To complete this procedure, you must have read-write access to the Web.config file.

**To configure a policy page using the Web.config file**

1. Use Notepad or other text editor to open the Web.config file that is located in %systemdrive%\ADFS\sts on the federation server.
2. Search for **<forms>** under **<logonserver>**.
3. Add or change the <policy> entry under <forms>, as follows:

   **<policy>** *PathToPolicyFileName* **.aspx</policy>**
4. Save and close the Web.config file.

| **Note** |
| --- |
| This setting takes effect immediately for all users that subsequently reach the account Federation Service. However, users who already have a token and do not return to the Federation Service until that token expires are not affected. |

**See Also**
**Concepts**
Enable enhanced identity privacy

# Backing Up ADFS Components

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Backing up Active Directory Federation Services (ADFS) components to preserve a snapshot of ADFS state is critical for ensuring a recovery option in the event of lost data or hardware failure.

ADFS state is maintained in the following places:

- TrustPolicy.xml file. The default location is *%systemdrive%*\adfs\sts.
- Web.config and other files under *%systemdrive%*\ADFS\..., especially any customized Web pages (*%systemdrive%*\ADFS\sts\ls).
- IIS metabase files (MetaBase.xml and MBSchema.xml) in *%systemroot%*\System32\Inetsrv (included in system state backup).
- Windows registry (included in system state backup).
- Custom transform module (.dll) or files related to the custom transform module.

Always back up system state on any federation server, federation server proxy, or Web server that you are backing up. System state contains the following components:

- System startup (boot) files. These files are required for Windows Server 2003 to start.
- System registry.
- IIS settings.
- Class registration database of component services. The Component Object Model (COM) is a binary standard for writing component software in a distributed systems environment.
- Windows Clustering or Certificate Services, if installed.

Use the following table to identify components to back up on servers that are running ADFS components.

| ADFS Component Running on Server | Files and Components to Back Up |
|---|---|
| Federation Service | • TrustPolicy.xml file<br>• Web.config and other files under *%systemdrive%E*\ADFS\...<br>• System state<br>• Custom transform module (.dll) and related files, if any |
| Federation Service Proxy | • Web.config and other files under *%systemdrive%*\ADFS\...<br>• System state |

| | |
|---|---|
| ADFS Web Agent | • *%systemdrive%*\ADFS\...<br>• System state |

| **Note** |
|---|
| For a claims-aware application, ADFS settings are contained in the Web.config file. For a Windows NT token-based application, ADFS settings are contained in the IIS metabase files. |

**Task requirements**

You need the following to perform the procedures for this task:

- Backup software

To complete this task, perform the following procedures:

- Back up ADFS components on a federation server, federation server proxy, or Web server

# Back up ADFS components on a federation server, federation server proxy, or Web server

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

This procedure provides instructions for using Ntbackup.exe to create a backup for Active Directory Federation Services (ADFS) components on federation servers, federation server proxies, and Web servers that are running the ADFS Web Agent. Ntbackup is available on all servers that are running Windows Server 2003. However, you can use any backup software to back up the components identified in Backing Up ADFS Components.

**Administrative credentials**

To complete this procedure, you must be a member of the Backup Operators group on the local computer.

**To back up ADFS components**

1. To start the Windows Server 2003 backup utility, click **Start**, click **Run**, type **ntbackup**, and then click **OK**.

| **Note** |
| --- |
| This procedure provides steps for backing up in Wizard Mode. By default, the **Always Start in Wizard Mode** check box is selected in the Backup or Restore Wizard. If the **Welcome to the Backup Utility Advanced Mode** page appears, click **Wizard Mode** to open the Backup or Restore Wizard. |

2. On the **Welcome to the Backup or Restore Wizard** page, click **Next**.
3. On the **Backup or Restore** page, click **Back up files and settings**, and then click **Next**.
4. On the **What to Back Up** page, click **Let me choose what to back up**, and then click **Next**.
5. On the **Items to Back Up** page, double-click **My Computer**.
6. In the expanded list below **My Computer**, expand **Local Disk (*DRIVE*:)** and make selections by placing check marks in the boxes for each directory that you want to back up, according to the ADFS component that is running on the server (see the table in Backing Up ADFS Components).
7. In the list below **My Computer**, place a check mark in the **System State** box, and then click **Next**.
8. On the **Backup Type, Destination, and Name** page, browse to select a location to save the backup, and type a name for the backup.
9. On the **Completing the Backup or Restore Wizard** page, click **Finish**.

**See Also**
**Concepts**
Backing Up ADFS Components

# Managing ADFS Partnerships

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When an Active Directory Federation Services (ADFS) deployment supports a Web resource that must be accessed by other organizations across the Internet, you must manage the partnership between your organization and a partner organization. You can add and remove partners and configure the properties that allow users in the account partner to access resources and that allow the resource Web applications to appropriately recognize, authenticate, and authorize those users.

The following tasks for managing ADFS partnerships are described in this objective:

- Adding a New Account Partner
- Adding a New Resource Partner
- Configuring Windows Trust for Account and Resource Partners
- Removing ADFS Partners

**See Also**
**Other Resources**
Partner organizations

# Adding a New Account Partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use the New Account Partner wizard to add an account partner to the resource Federation Service in Active Directory Federation Services (ADFS). This wizard allows you to create an account partner that requires manual configuration of the trust policy or imports an existing policy file that is provided (exported) by an existing federation server in the account partner, as follows:

- Manually configure the trust policy: If the corresponding account partner organization has not installed ADFS yet or does not plan to provide you with an exported policy file, create the account partner to use the values that you provide. To do so, you must have the following information about the account partner Federation Service:

  - Display name: The name of the account partner Federation Service. This name appears in the list of realms that is presented to clients that request access to a Web site that is protected by ADFS authentication and authorization.

  - Federation Service Uniform Resource Identifier (URI): Uniquely identifies the Federation Service, and identifies this server as a member of the account Federation Service, in the form **urn:federation:***OrganizationName*.

  - Federation Service endpoint Uniform Resource Locator (URL): The URL that will be used by clients to access a Web server in this Federation Service, in the form **https://***FullyQualifiedDomainName***/adfs/ls/**. If a federation server proxy is installed, the fully qualified domain name should be the host name of the federation server proxy. If no federation server proxy is installed, the fully qualified domain name should be the name you are using to represent federation servers in the account Federation Service.

- Import a trust policy file: When you import a policy file when creating an account partner, the properties of the new account partner are automatically entered in the trust policy by the New Account Partner wizard. This information is derived from the account Federation Service trust policy file, thereby eliminating configuration errors. To import a policy file, the account partner must have exported its generic policy file or partner policy file and provided you with the file or its shared location.

**Task requirements**

You need the following to perform the procedures for this task:

- Active Directory Federation Services snap-in running on a federation server.

To complete this task, perform the following procedures, as needed:

- Add a new account partner by manually configuring the trust policy

- Export an account or resource policy file to a partner organization

- Add a new account partner by importing an existing policy file

**See Also**
**Concepts**
Adding a New Resource Partner

# Add a new account partner by manually configuring the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

To add a new account partner in Active Directory Federation Services (ADFS) and manually configure the trust policy, perform the following procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new account partner by manually configuring the trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. Right-click **Account Partners**, point to **New**, and then click **Account Partner**.
4. On the **Welcome to the Add Account Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, click **No**, and then click **Next**.
6. On the **Account Partner Details** page, do the following, and then click **Next**
   - In **Display name**, type the display name of the account partner.
   - In **Federation Service URI**, type the Uniform Resource Identifier (URI) of the account partner Federation Service.
   - In **Federation Service endpoint URL**, type the endpoint Uniform Resource Locator (URL) as follows:

     **With a Federation Service Proxy:** If the Federation Service Proxy is installed on servers in a perimeter network of the account domain, use the fully qualified domain name (FQDN) of any federation server proxy.

     **Without a Federation Service Proxy:** If your deployment does not use federation server proxies, use the FQDN of the account partner Federation Service.
7. On the **Account Partner Verification Certificate** page, type the path to the verification certificate, or browse to it, and then click **Next**.
8. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**, and then go to step 10.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
9. On the **Federated Web SSO with Forest Trust** page, do one of the following, and then click **Next**:
   - To accept users in all domains that are trusted by the account partner, click **All Active Directory domains and forests**. Any user that can authenticate to the account partner will be accepted.
   - To accept user accounts that are located in some of the domains that are trusted by the account partner, click **The following Active Directory domains and forests**. Then, in **New, trusted Active Directory domain or forest**, type the name of a domain or forest, and then click **Add**. Only users from the specified domains will be accepted.
10. On the **Account Partner Identity Claims** page, select one or more identity claims to share with the resource partner,

and then click **Next**:

- If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

| **Important** |
| --- |
| When UPN claims or e-mail claims are used to make authorization decisions, it is essential that each account partner use a unique UPN suffix or e-mail suffix. If two account partners use the same UPN suffix or e-mail suffix, it might not be possible to identify users uniquely. This condition might allow a user from one account partner to receive the permissions that are intended for a user in another account partner. This condition might also introduce a significant security weakness because an administrator can intentionally create user accounts to impersonate users from one of your other account partners. |

| **Note** |
| --- |
| If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable because UPN claims are required for this scenario. |

- If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.

- If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.

11. If you selected **UPN Claim** as an identity claim, on the **Accepted UPN Suffixes** page, type the accepted suffix, click **Add**, and then click **Next**. You may click to accept **All UPN suffixes** if you selected the **Federation Web SSO with Forest Trust** option.
12. If you selected **E-mail Claim** as an identity claim, on the **Accepted E-mail Suffixes** page, type the accepted suffix, click **Add**, and then click **Next**. You may click to accept **All E-mail suffixes** if you selected the **Federation Web SSO with Forest Trust** option.

| **Note** |
| --- |
| Common name claims require no additional information. |

13. On the **Enable this Account Partner** page, if you do not want to enable the account partner now, clear the **Enable this account partner** check box, and then click **Next**.
14. To add the new account partner and close the wizard, click **Finish**.

**See Also**
**Concepts**
Change the Federation Service endpoint URL
View the current verification certificate

# Export an account or resource policy file to a partner organization

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have created an Active Directory Federation Services (ADFS) resource or account partner in your side of a federated partnership, you can export a trust policy file that has information about both your Federation Service and the Federation Service of the respective account or resource partner Federation Service. The policy file contains the following information that the prospective partner can use to configure its Federation Service trust policy:

- Resource Display Name
- Resource URI
- Resource Federation Server Proxy URL
- Account Display Name
- Account URI
- Account Federation Server Proxy URL
- Account Verification Certificate

For example, if you are a resource partner, you can export your partner policy file and provide it to the account partner organization. When the account partner adds a resource partner for your organization and selects the option to import your policy file, the Add Resource Partner wizard uses the imported file to automatically update the trust policy with the correct information for both organizations.

Perform this procedure on a federation server that hosts the account or resource partner that represents the Federation Service whose policy file you are exporting. The **Export Policy** command in the Active Directory Federation Services snap-in creates the file with the name and location that you provide.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To export an account or resource policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners** or **Resource Partners**, depending on the federation server role.
3. Right-click the account or resource partner, and then click **Export Policy**.
4. In the **Export Partner Policy** dialog box box, click **Browse** to browse to the location for the partner policy file (or type the path and file name in **Policy file Location**, and then click **OK**).
5. In the **Save As** dialog box, click **Save** and then click **OK**.
6. Notify the partner organization and make the exported file available to the partner organization.

**See Also**
**Concepts**
Export a generic policy file to a partner organization

# Add a new account partner by importing an existing policy file

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have received an exported Active Directory Federation Services (ADFS) trust policy file from the account partner organization, perform the following procedure on a resource federation server to automatically configure the new account partner by importing the policy file.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new account partner by importing an existing policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, double-click **Trust Policy**, and then double-click **Partner Organizations**.
3. Right-click **Account Partners**, point to **New**, and then click **Account Partner**.
4. On the **Welcome to the Add Account Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, click **Yes**. In **Partner interoperability policy file**, type or browse to the policy file that you want to import, and then click **Open**.
6. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**, and then go to step 8.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
7. On the **Federated Web SSO with Forest Trust** page, do one of the following, and then click **Next**:
   - To accept users in all domains that are trusted by the account partner, click **All Active Directory domains and forests**. Any user that can authenticate to the account partner will be accepted.
   - To accept user accounts that are located in some of the domains that are trusted by the account partner, click **The following Active Directory domains and forests**. Then, in **New, trusted Active Directory domain or forest**, type the name of a domain or forest, and then click **Add**. Only users from the specified domains will be accepted. If you need to remove a domain or forest, click **Remove**.
8. On the **Account Partner Identity Claims** page, select one or more identity claims to share with the resource partner, and then click **Next**:
   - If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

| **Important** |
| --- |
| When UPN claims or e-mail claims are used to make authorization decisions, it is essential that each account partner use a unique UPN suffix or e-mail suffix. If two account partners use the same UPN suffix or e-mail suffix, it might not be possible to uniquely identify users. This condition might allow a user from one account partner to receive the permissions that are intended for a user in another account partner. This condition might also introduce a significant security weakness because an administrator can intentionally create user accounts to impersonate users from one of your other account partners. |

| **Note** |
| --- |

| |
|---|
| If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable because UPN claims are required for this scenario. |

- If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.
- If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.

9. If you selected **UPN Claim** as an identity claim, on the **Accepted UPN Suffixes** page, type the accepted suffix, click **Add**, and then click **Next**. You may click to accept **All UPN suffixes** if you selected the **Federation Web SSO with Forest Trust** option.
10. If you selected **E-mail Claim** as an identity claim, on the **Accepted E-mail Suffixes** page, type the accepted suffix, click **Add**, and then click **Next**. You may click to accept **All E-mail suffixes** if you selected the **Federation Web SSO with Forest Trust** option.

| **Note** |
|---|
| Common name claims require no additional information. |

11. On the **Enable this Account Partner** page, if you do not want to enable the account partner now, clear the **Enable this account partner** check box, and then click **Next**.
12. To add the new account partner and close the wizard, click **Finish**.

**See Also**
**Concepts**
Export an account or resource policy file to a partner organization
Add a new account partner by manually configuring the trust policy

# Adding a New Resource Partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can use the New Resource Partner wizard to add a resource partner to the account Federation Service. This wizard allows you to create a resource partner that requires manual configuration of the trust policy or imports an existing policy file that is provided (exported) by an existing federation server, as follows:

- Manually configure the trust policy: If the corresponding resource partner organization does not plan to provide you with an exported policy file, create the resource partner to use the values that you provide. To do so, you must have the following information about the resource partner:

    - Display name: The name the resource partner uses to identify itself.

    - Federation Service Uniform Resource Identifier (URI): Uniquely identifies the Federation Service, and identifies this server as a member of the account Federation Service, in the form **urn:federation:***OrganizationName*.

    - Federation Service endpoint Uniform Resource Locator (URL): The URL that will be used by clients to access a server in this Federation Service, in the form **https://***FullyQualifiedDomainName***/adfs/ls/**

- Import a policy file: When you import a policy file during resource partner creation, the properties of the new resource partner are automatically entered in the trust policy by the New Resource Partner wizard according to the information in the resource policy file, thereby eliminating configuration errors. To import a policy file, the resource partner must have exported its generic policy file or partner policy file and provided you with the file or its shared location.

**Task requirements**

You need the following to perform the procedures for this task:

- Active Directory Federation Services snap-in running on a federation server.

To complete this task, perform one of the following procedures:

- Add a new resource partner by manually configuring the trust policy

- Add a new resource partner by importing an existing policy file

**See Also**
**Concepts**
Adding a New Account Partner

# Add a new resource partner by manually configuring the trust policy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

To add a new resource partner in Active Directory Federation Services (ADFS) and manually configure the trust policy, perform the following procedure on an account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new resource partner by manually configuring the trust policy**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. In the console tree, right-click **Resource Partners**, point to **New**, and then click **Resource Partner**.
4. On the **Welcome to the Add Resource Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, ensure that **No** is selected, and then click **Next**.
6. On the **Resource Partner Details** page, do the following, and then click **Next**:
   - In **Display name**, type the display name of the resource partner.
   - In **Federation Service URI**, type the Uniform Resource Identifier (URI) for the resource partner Federation Service.
   - In **Federation Service endpoint URL**, type the endpoint Uniform Resource Locator (URL) of the resource partner Federation Service.
7. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
8. On the **Resource Partner Identity Claims** page, select one or more identity claims to share with the resource partner, and then click **Next**:
   - If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

     | **Note** |
     | --- |
     | If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable because UPN claims are required for this scenario. |

   - If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.
   - If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.
9. If you selected **UPN Claim** as an identity claim, on the **Select UPN Suffix** page, do one of the following, and then click

**Next**:

- To pass all UPN suffixes through without replacing them, click **Pass all UPN suffixes through unchanged**.
- To replace all UPN suffixes with a different suffix, click **Replace all UPN suffixes with the following**, and then type the suffix that you want to use to replace all UPN suffixes.

10. If you selected **E-mail Claim** as an identity claim, on the **Select E-mail Suffix** page, do one of the following, and then click **Next**:

    - To pass all e-mail suffixes without replacing them, click **Pass all E-mail suffixes through unchanged**.
    - To replace all e-mail suffixes with a different suffix, click **Replace all E-mail suffixes with**, and then type the suffix that you want to use to replace all e-mail suffixes.

| **Note** |
| --- |
| Common name claims require no additional information. |

11. On the **Enable this Resource Partner** page, if you do not want to enable the resource partner now, clear the **Enable this resource partner** check box, and then click **Next**.
12. To add the new resource partner and close the wizard, click **Finish**.

**See Also**
**Concepts**
Add a new resource partner by importing an existing policy file

# Export an account or resource policy file to a partner organization

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have created an Active Directory Federation Services (ADFS) resource or account partner in your side of a federated partnership, you can export a trust policy file that has information about both your Federation Service and the Federation Service of the respective account or resource partner Federation Service. The policy file contains the following information that the prospective partner can use to configure its Federation Service trust policy:

- Resource Display Name
- Resource URI
- Resource Federation Server Proxy URL
- Account Display Name
- Account URI
- Account Federation Server Proxy URL
- Account Verification Certificate

For example, if you are a resource partner, you can export your partner policy file and provide it to the account partner organization. When the account partner adds a resource partner for your organization and selects the option to import your policy file, the Add Resource Partner wizard uses the imported file to automatically update the trust policy with the correct information for both organizations.

Perform this procedure on a federation server that hosts the account or resource partner that represents the Federation Service whose policy file you are exporting. The **Export Policy** command in the Active Directory Federation Services snap-in creates the file with the name and location that you provide.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To export an account or resource policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners** or **Resource Partners**, depending on the federation server role.
3. Right-click the account or resource partner, and then click **Export Policy**.
4. In the **Export Partner Policy** dialog box box, click **Browse** to browse to the location for the partner policy file (or type the path and file name in **Policy file Location**, and then click **OK**).
5. In the **Save As** dialog box, click **Save** and then click **OK**.
6. Notify the partner organization and make the exported file available to the partner organization.

**See Also**
**Concepts**
Export a generic policy file to a partner organization

# Add a new resource partner by importing an existing policy file

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have received a policy file from your resource partner organization that you can import, perform the following procedure on the account federation server on which you are creating the new resource partner.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add a new resource partner by importing an existing policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. Right-click **Resource Partners**, point to **New**, and then click **Resource Partner**.
4. On the **Welcome to the Add Resource Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, do the following, and then click **Next**:
   - Click **Yes**.
   - In **Partner interoperability policy file**, browse to or type the location of the resource partner policy file.
6. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
7. On the **Resource Partner Identity Claims** page, select one or more identity claims that the account partner will provide to the resource partner, and then click **Next**:
   - If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

     | **Note** |
     | --- |
     | If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable because UPN claims are required for this scenario. |

   - If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.
   - If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.
8. If you selected **UPN Claim** as an identity claim, on the **Select UPN Suffix** page, select one of the following, and then click **Next**.
   - To pass all UPN suffixes through without replacing them, click **Pass all UPN suffixes through unchanged**.
   - To replace all UPN suffixes with a different suffix, click **Replace all UPN domain suffixes with the following**, type the suffix that you want to use to replace all UPN suffixes, and then click **Add**.

9. If you selected **E-mail Claim** as an identity claim, on the **Select E-mail Suffix** page, do one of the following, and then click **Next**:

   - To pass all e-mail suffixes without replacing them, click **Pass all e-mail suffixes through unchanged**.

   - To replace all UPN suffixes with a different suffix, click **Replace all E-mail suffixes with**, and then type the suffix that you want to use to replace all e-mail suffixes.

| **Note** |
| --- |
| Common name claims require no additional information. |

10. On the **Enable this Resource Partner** page, if you do not want to enable the resource partner now, clear the **Enable this resource partner** check box, and then click **Next**.

11. To add the new resource partner and close the wizard, click **Finish**.

**See Also**
**Concepts**
Export an account or resource policy file to a partner organization
Add a new resource partner by manually configuring the trust policy

# Configuring Windows Trust for Account and Resource Partners

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

An Active Directory Federation Services (ADFS) deployment in a federated scenario can be configured with or without a Windows trust relationship. In a Federated Web single sign-on (SSO) scenario, a federation trust is created between the two partners, but a Windows trust relationship, if one exists between the two partners, is not used. In a Federated Web SSO with Forest Trust scenario, an organization that has a Windows trust relationship between the two Active Directory forests (the resource forest trusts the account forest) configures ADFS partners to use the Windows trust relationship.

The Windows trust between the forests of the two partners (either a forest trust between two Windows Server 2003 forests or an external trust between Windows Server 2003 or Windows 2000 Server domains in each forest) must be enabled for ADFS by both the resource partner and the account partner. The account partner must be configured to select the domains that are to be included in the trust relationship. You must also configure the resource partner in the account Federation Service.

**Task requirements**

You need the following to perform the procedures for this task:

- A Windows trust between the two partners (two Windows Server 2003 forests or an external trust between Windows Server 2003 or Windows 2000 Server domains in each forest), where the resource forest or domain has a trust relationship with the account forest or domain.
- A federation server in each partner organization.
- The Active Directory Federation Services snap-in.

This task provides procedures to begin using or discontinue Windows trust in both partners:

- Configure an account partner to use Windows trust
- Configure a resource partner to use Windows trust
- Discontinue Windows trust for an account partner
- Discontinue Windows trust for a resource partner

**See Also**
**Other Resources**
Administering Domain and Forest Trusts

# Configure an account partner to use Windows trust

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You must configure the account partner to select the domains that are to be included in the trust relationship. For example, if a Windows Server 2003 forest trust is in place from the resource partner forest to the account partner forest, the trust is transitive to all domains in the trusted account forest.

If you want to grant access to resources in the resource forest to users in only some domains in the account forest, you can specify only those domains. Otherwise, you can allow all trusted domains (all domains in the account forest and any forest that is trusted by the account forest) to be granted access.

Use the following procedure to enable Windows trust for the account partner that will participate in the Federated Web SSO with Forest Trust design scenario.

Perform this procedure on a federation server in the resource partner organization.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure an account partner to use Windows trust**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners**.
3. Right-click the account partner that you want to configure to use Windows trust, and then click **Properties**.
4. On the **Windows Trust** tab, click **Use Windows trust relationship**.
5. In **Trusted Active Directory domains and forests**, do one of the following, and then click **OK**:
   - Click **All Active Directory domains and forests** if you want to allow users in all trusted domains in the account partner forest and in any forest that is trusted by the account partner forest, and then click **OK**.
   - Click **Specified Active Directory domains and forests (press Enter to separate entries)** if you want to name only the domains where you want to allow federated access. Type a domain name, press ENTER, and then repeat this action to add each domain in the account partner forest and in any other trusted forests for the users that you want to grant access to resources.

**See Also**
**Concepts**
Configure a resource partner to use Windows trust
Discontinue Windows trust for an account partner

# Configure a resource partner to use Windows trust

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Use the following procedure to enable Windows trust for a resource partner in an Active Directory Federation Services (ADFS) Federated Web SSO with Forest Trust scenario.

| **Note** |
| --- |
| When you enable the Windows trust option in the account Federation Service, you are sending actual security identifiers (SIDs) to the resource partner organization over the Internet, which may be a security risk. These SIDs are packaged in the ADFS Security Assertion Markup Language (SAML) token. Therefore, enable this option only when you are using the Federated Web SSO with Forest Trust design. This design is meant to establish secure communication within the same organization. |

Perform this procedure on an account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure a resource partner to use Windows trust**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Resource Partners**.
3. Right-click the resource partner for which you want to configure Windows trust, and then click **Properties**.
4. On the **General** tab, click **Use Windows trust relationship for this partner**, and then click **OK**.

**See Also**
**Concepts**
Configure an account partner to use Windows trust
Discontinue Windows trust for a resource partner

# Discontinue Windows trust for an account partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you want to stop using Windows trust for an account partner in an Active Directory Federation Services (ADFS) scenario, you can configure the account partner to not use the Windows trust relationship.

Perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To discontinue using Windows trust for an account partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners**.
3. Right-click the account partner for which you want to stop using Windows trust, and then click **Properties**.
4. On the **Windows Trust** tab, clear the **Use Windows trust relationship** check box, and then click **OK**.

**See Also**
**Concepts**
Configure an account partner to use Windows trust

# Discontinue Windows trust for a resource partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you want to stop using Windows trust for a resource partner, you can configure the resource partner to not use the Windows trust relationship.

Perform this procedure on an account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To discontinue using Windows trust for a resource partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Resource Partners**.
3. Right-click the resource partner for which you want to stop using Windows trust, and then click **Properties**.
4. On the **General** tab, clear the **Use Windows trust relationship for this partner** check box, and then click **OK**.

# Removing ADFS Partners

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

A resource partner in an Active Directory Federation Services (ADFS) deployment can have more than one account partner, and vice versa. When a partner organization is to be removed from the federation, you must remove the partner node in the trust policy on the other side of the ADFS partnership. The partner that is being removed from the federation must remove its federation servers and federation server proxies, as well. For information about removing servers that are running ADFS components, see Managing ADFS Components.

**Task requirements**

You need the following tool to perform the procedures for this task:

- Active Directory Federation Services MMC snap-in

To complete this task, perform the following procedures:

- Delete an existing account partner
- Delete an existing resource partner
- Managing ADFS Components

# Delete an existing account partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Use this procedure to delete the account partner node for an account partner that is no longer a member of your Active Directory Federation Services (ADFS) deployment.

Perform this procedure on a federation server in the resource Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To delete an account partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners**.
3. Right-click the account partner you want to delete, and then click **Delete**.
4. In the **Delete Account Partner** dialog box, click **Yes** to confirm the deletion.

# Delete an existing resource partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Use this procedure to delete the resource partner node for a resource partner that is no longer a member of an Active Directory Federation Services (ADFS) deployment.

Perform this procedure on a federation server in the account Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To delete a resource partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Resource Partners**.
3. Right-click the resource partner you want to delete, and then click **Delete**.
4. In the **Delete Resource Partner** dialog box, click **Yes** to confirm the deletion.

# Managing Partner Relationships

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In an Active Directory Federation Services (ADFS) partnership, you can manage privacy aspects of the partner experience, export your trust policy for a partner, and manage resource (impersonation) accounts. You can also enable or disable a partner when you need to temporarily suspend the partnership.

**Task requirements**

You need the following to perform the procedures for this task:

- Active Directory Federation Services MMC snap-in.

To complete this task, perform the following procedures on an as-needed basis:

- Enable enhanced identity privacy
- Export a generic policy file to a partner organization
- Export an account or resource policy file to a partner organization
- Enable or disable a resource partner
- Enable or disable an account partner
- Configure resource account optionsr

**See Also**
**Other Resources**
Partner organizations

# Enable enhanced identity privacy

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Enhanced identity privacy is an optional setting that you can configure on a resource partner in the account Federation Service in an Active Directory Federation Services (ADFS) deployment. This setting hashes the user-name portion of outgoing user principal name (UPN) claims and e-mail claims. It substitutes the common name with a random value. If you select the **Enable enhanced identity privacy** option in ADFS, the resource partner will not be able to correlate identity claims to personally identifiable user information.

The enhanced identity privacy setting affects the information that is sent in identity claims, based on the claim type that is being used to transfer the user identity, as follows:

- UPN and e-mail claim types: The user component of the UPN and e-mail name is hashed, replacing the user component in the identity claim of the security token. In this way, each resource partner can uniquely identify each user without revealing their true identity.

- Common name claim types: The common name identity claim is populated with a randomly generated, globally unique identifier (GUID), which ensures that the identity claim is unique per session with the resource partner and that multiple sessions by the same user cannot be tracked.

Enable this setting if you want to:

- Prevent collusion between partners in correlating identity claims to personally identifiable user information.

- Prevent simple dictionary attacks against the user-name hash.

For more information about the effects of enhanced identity privacy, see Partner organizations (http://go.microsoft.com/fwlink/?LinkId=62227).

Perform this procedure on an account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable enhanced identity privacy on a resource partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Resource Partners**.
3. Right-click the resource partner that will begin using enhanced identity privacy, and then click **Properties**.
4. On the **General** tab, click **Enable enhanced identity privacy**, and then click **OK**.

**See Also**
**Concepts**
Configure a policy page for a Web site

# Export a generic policy file to a partner organization

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

To begin an Active Directory Federation Services (ADFS) federated partnership with another organization, prior to adding resource and account partners, you can export a generic policy file that the other organization can use to configure your organization as either its account or resource partner. For example, if you are administering a resource organization, you can export a generic policy file that the account organization can use to configure your organization as its resource partner in ADFS. You must make the file available to the partner organization, which can then import the file.

The exported generic policy file contains the following information:

- Your organization's display name
- Federation Service endpoint URL
- Federation Service URI
- Verification certificate (for an exported account partner only)

When the partner organization imports this policy file when adding a resource or account partner that represents your organization, the wizard (Add Resource Partner Wizard or Add Account Partner Wizard) automatically configures the new partner with the correct information for your organization, eliminating the possibility of mistyping the various required identifiers.

Perform this procedure on a federation server that hosts the policy file you want to share with another organization that will become a partner in a federated partnership. The **Export Policy** command in the Active Directory Federation Services snap-in creates the file with the name and location that you provide.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To export a generic policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Trust Policy**, and then click **Export Policy**.
3. In the **Export Generic Partner Policy** dialog box, click **Browse** to browse to the location for the generic policy file (or type the path and file name in **Policy file Location**, and then click **OK**).
4. In the **Save As** dialog box, in **File name**, type a name for the generic trust policy file, click **Save**, and then click **OK**.
5. Make the exported file available to the partner organization.

**See Also**
**Concepts**
Export an account or resource policy file to a partner organization

# Export an account or resource policy file to a partner organization

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have created an Active Directory Federation Services (ADFS) resource or account partner in your side of a federated partnership, you can export a trust policy file that has information about both your Federation Service and the Federation Service of the respective account or resource partner Federation Service. The policy file contains the following information that the prospective partner can use to configure its Federation Service trust policy:

- Resource Display Name
- Resource URI
- Resource Federation Server Proxy URL
- Account Display Name
- Account URI
- Account Federation Server Proxy URL
- Account Verification Certificate

For example, if you are a resource partner, you can export your partner policy file and provide it to the account partner organization. When the account partner adds a resource partner for your organization and selects the option to import your policy file, the Add Resource Partner wizard uses the imported file to automatically update the trust policy with the correct information for both organizations.

Perform this procedure on a federation server that hosts the account or resource partner that represents the Federation Service whose policy file you are exporting. The **Export Policy** command in the Active Directory Federation Services snap-in creates the file with the name and location that you provide.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To export an account or resource policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners** or **Resource Partners**, depending on the federation server role.
3. Right-click the account or resource partner, and then click **Export Policy**.
4. In the **Export Partner Policy** dialog box box, click **Browse** to browse to the location for the partner policy file (or type the path and file name in **Policy file Location**, and then click **OK**).
5. In the **Save As** dialog box, click **Save** and then click **OK**.
6. Notify the partner organization and make the exported file available to the partner organization.

**See Also**
**Concepts**
Export a generic policy file to a partner organization

# Enable or disable a resource partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) resource partners are enabled by default when you create them in the Active Directory Federation Services snap-in. If a resource partner is currently enabled, you can disable it. If a resource partner is currently disabled, you can enable it.

Perform this procedure on an account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**ToTo enable or disable a resource partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Resource Partners**.
3. Right-click the resource partner you want to enable or disable and do one of the following:
   - If the resource partner is disabled and you want to enable it, click **Enable**.
   - If the resource partner is enabled and you want to disable it, click **Disable**.

| **Note** |
| --- |
| You can also perform this procedure by right-clicking the resource partner, clicking **Properties**, and then selecting or clearing the **Enable this partner** check box. |

# Enable or disable an account partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) account partners are enabled by default when you create them in the Active Directory Federation Services snap-in. If an account partner is currently enabled, you can disable it. If an account partner is currently disabled, you can enable it.

Perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**ToTo enable or disable an account partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners**.
3. Right-click the account partner you want to enable or disable and do one of the following:
   - If the account partner is disabled and you want to enable it, click **Enable**.
   - If the account partner is enabled and you want to disable it, click **Disable**.

| **Note** |
| --- |
| You can also perform this procedure by right-clicking the account partner, clicking **Properties**, and then selecting or clearing the **Enable this partner** check box. |

# Configure resource account options

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

By adjusting resource account options in the resource Federation Service, resource partner administrators can refine how access control for federated users is administered and delegated on an account partner–by–account partner basis. Use the following procedure to adjust resource account options .

| **Note** |
| --- |
| We recommend that you first review information about each resource account option before you modify the default resource account setting. For more information about resource account options, see Select the optimal resource account option. |

Perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure resource account options**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners**.
3. Right-click the account partner whose resource account behavior you want to change, and then click **Properties**.
4. Click the **Resource Accounts** tab.
5. Select one of the following options for using resource accounts for this account partner, and then click **OK**:
   - **Resource accounts exist for all users** (Resource group claim mappings are not checked.)
   - **Resource accounts exist for some users (prefer resource account)** (First, check for resource accounts. If they do not exist, use the resource groups in the token.)
   - **Resource accounts exist for some users (prefer groups in token)** (First, process resource groups in tokens. If none exist, check for resource accounts.)
   - **No resource accounts exist for this account partner** (Check only for resource groups in tokens.)

**See Also**
**Concepts**
Map an organization group claim to a resource group

# Managing Accounts and Account Stores

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) uses account stores to log on users and extract security claims for those users. You can configure multiple account stores for a single Federation Service and define their priority. The Federation Service uses Lightweight Directory Access Protocol (LDAP) to communicate with account stores.

ADFS supports two types of account stores: Active Directory account stores and Active Directory Application Mode (ADAM) account stores. Only one Active Directory account store can be used for an ADFS deployment. However, you can use multiple ADAM account stores.

Enabling and disabling an account store is the same procedure for both types of account stores.

- Enable or disable an account store

The following tasks for managing account stores are described in this objective.

- Managing Active Directory Account Stores
- Managing ADAM Account Stores
- Using Multiple Account Stores

**See Also**
**Other Resources**
Account Stores

# Enable or disable an account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If an account store needs to be taken offline temporarily, you can disable the store in Active Directory Federation Service (ADFS) to stop authentication requests by federation servers. If you have disabled an account store in ADFS and you want to enable it, you can use the same procedure to enable the store.

Perform this procedure on an account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable or disable an account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Account Stores**.
3. Right-click the account store you want to enable or disable and do one of the following:
   - If the account store is disabled and you want to enable it, click **Enable**.
   - If the account store is enabled and you want to disable it, click **Disable**.

| **Note** |
| --- |
| You can also perform this procedure by right-clicking the account store, clicking **Properties**, and then selecting or clearing the **Enable this account store** check box. |

# Managing Active Directory Account Stores

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) is tightly integrated with Active Directory. ADFS retrieves user attributes and authenticates users against Active Directory. ADFS also uses Windows Integrated Authentication and security tokens that Active Directory creates.

Management of the Active Directory account store is limited to adding, removing, and enabling or disabling the account store.

**Task requirements**

You need the following tool to perform the procedures for this task:

- Active Directory Federation Services MMC snap-in.

To complete this task, perform the following procedures on an as-needed basis:

- Add an Active Directory account store
- Remove an Active Directory account store
- Enable or disable an account store

**See Also**
**Other Resources**
Account Stores

# Add an Active Directory account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If user and computer accounts that require access to a resource that is protected by Active Directory Federation Services (ADFS) are stored in Active Directory, you must add Active Directory as an account store on a federation server in the Federation Service that authenticates the accounts.

An Active Directory forest can have only one Active Directory instance. Therefore, you can add only one Active Directory store for the respective Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add an Active Directory account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Account Stores**, point to **New**, and then click **Account Store**.
3. On the **Welcome to the Add Account Store Wizard** page, click **Next**.
4. On the **Account Store Type** page, ensure that **Active Directory** is selected, and then click **Next**.
5. On the **Enable this Account Store** page, ensure that the **Enable this account store** check box is selected, and then click **Next**.
6. On the **Completing the Add Account Store Wizard** page, click **Finish**.

**See Also**
**Concepts**
Add an ADAM account store

# Remove an Active Directory account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Use this procedure to delete the Active Directory account store node for an Active Directory Federation Services (ADFS) Federation Service that no longer uses an Active Directory account store.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To delete the Active Directory account store node**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Account Stores**.
3. Right-click **Active Directory**, and then click **Delete**.
4. In the **Delete Account Store** dialog box, click **Yes** to confirm the deletion.

# Managing ADAM Account Stores

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Application Mode (ADAM) provides data storage and retrieval for directory-enabled applications, without the dependencies that are required for Active Directory. A Federation Service can use multiple ADAM account stores when they are needed to accommodate sets of users for different applications. To differentiate between ADAM stores, the Federation Service assigns a unique value to each ADAM store in the trust policy, in the form of the uniform resource identifier (URI).

**Task requirements**

You need the following to perform the procedures for this task:

- An ADAM instance
- ADAM ADSI Edit
- Ldp, a Windows Server 2003 Support Tool, available in the SUPPORT\TOOLS folder on the Windows Server 2003 operating system CD.
- Active Directory Federation Services MMC snap-in

To complete this task, perform the following procedures on an as-needed basis:

- Prepare an ADAM instance for use with ADFS
- Add an ADAM account store
- Change the server name or IP address for an ADAM account store
- Change the display name for an ADAM account store
- Change the port number for an ADAM account store
- Change the search base for an ADAM account store
- Change the user name attribute for an ADAM account store
- Enable or disable TLS and SSL for an ADAM account store

**See Also**
**Other Resources**
Account Stores

# Prepare an ADAM instance for use with ADFS

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Before you can use an Active Directory Application Mode (ADAM) instance as an account store in your Active Directory Federation Services (ADFS) deployment, you must perform two preliminary procedures:

- Set an attribute to enable user accounts
- Configure the member attribute with the federation server security identifier (SID) to enable federation servers to search the ADAM store

**Enable ADAM User Accounts**

On ADAM instances running on Windows Server 2003, where local or domain password policy restrictions are in effect, the ADAM user account is disabled by default. Before you can enable the user account, you must set a password that meets the password policy restrictions that are in effect. This rule does not apply to ADAM instances running on Windows XP Professional.

To enable user accounts, set the msDS-UserAccountDisabled attribute value to **False**. Be sure that the user account has been configured with a userPassword attribute value that meets policy requirements.

Use the following procedure to enable a user account in the ADAM account store.

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable an ADAM user account**

1. Open ADAM ADSI Edit and connect to the ADAM instance.
2. Right-click an ADAM user, and then click **Properties**.
3. In the **Attribute** column, click **msDS-UserAccountDisabled**, and then click **Edit**.
4. Click **False**, and then click **OK** twice.

   If an ADAM-ADSIEdit message appears stating that the password cannot be updated because the value does not meet requirements for the domain, right-click the user account and click **Reset Password**. Then repeat this procedure.

**Configure the Federation Server SID**

To enable federation servers to search the ADAM account store, you need to add the machine account SID of the account federation server to the member attribute in the Readers role of the ADAM instance.

Use the following procedures to prepare ADAM for searches by federation servers.

- Obtain the machine account SID of the federation server
- Add the SID to the member attribute in ADAM

### Administrative credentials

To complete this procedure, you must be a member of the Domain Users group in the Active Directory domain of the federation server.

**To obtain the SID of the federation server**

1. Open Ldp and connect and bind to the Active Directory domain to which the federation server is joined.
2. On the **View** menu, click **Tree**.
3. Expand the tree to locate the computer object of the federation server.

4. Double-click the computer object and view the properties in the results pane.
5. Make a note of the value in **1>objectSid**.

Perform the next procedure to add the SID you obtained to the member attribute in ADAM.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add the federation server SID to the member attribute for an ADAM instance**

1. In Ldp, connect and bind to the ADAM instance.
2. On the **View** menu, click **Tree**.
3. Double-click the ADAM instance and then double click the **CN=Roles** container.
4. Right-click the **CN=Readers** container, and then click **Modify**.
5. In **Attribute**, type **member**.
6. In **Values**, type the SID value as follows, and then click **Enter**:

   **<SID=** *objectSIDValue* **>**
7. Click **Run** to modify the attribute, and then click **Close**.

# Add an ADAM account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you use multiple Active Directory Application Mode (ADAM) stores for user accounts that require access to one or more Web applications that are protected by Active Directory Federation Services (ADFS), you can add the ADAM account stores to the Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To add an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Account Stores**, point to **New**, and then click **Account Store**.
3. On the **Welcome to the Add Account Store Wizard** page, click **Next**.
4. On the **Account Store Type** page, ensure that **Active Directory Application Mode (ADAM)** is selected, and then click **Next**.
5. On the **ADAM Store Details** page, in **Account store display name**, type the name of the ADAM account store as you want it to be displayed in the Active Directory Federation Services snap-in user interface (UI).
6. In **Account store URI**, type the Uniform Resource Identifier (URI) for the ADAM account store, and then click **Next**.

| **Note** |
| --- |
| The account store URI uniquely identifies the ADAM instance among multiple ADAM account stores. |

7. On the **ADAM Server Settings** page, do the following, and then click **Next**:
   a. In **ADAM server name or IP address**, type the name or IP address of the ADAM server.
   b. In **Port number**, type the TCP/IP port number for the account service. Accept the default of 389 unless Active Directory is installed on the same server, in which case you must use a different port.
   c. In **LDAP search base distinguished name**, type the distinguished name of the ADAM instance.
   d. In **User name LDAP attribute**, type the name of the user name attribute that users provide during logon (for example, **userPrincipalName** or **sAMAccountName**).
8. On the **Identity Claims** page, select one or more identity claims that will be provided by the account store, and then click **Next**:
   a. If the account store provides user principal name (UPN) identity claims, select the **User Principal Name (UPN)** check box, and then type the Lightweight Directory Access Protocol (LDAP) attribute name to which UPN identity claims map (the attribute whose value is the user's UPN, usually **userPrincipalName**).
   b. If the account store provides e-mail identity claims, select the **E-mail** check box, and then type the LDAP attribute name to which e-mail identity claims map (the attribute whose value is the user's a-mail name, usually **userPrincipalName**).
   c. If the account store provides a common name identity claim, select the **Common Name** check box, and then type the LDAP attribute name to which the common name identity claim maps (the attribute whose value is the user's common name, usually **displayName**).

9. On the **Enable this Account Store** page, ensure that the **Enable this account store** check box is selected, and then click **Next**.
10. On the **Completing the Add Account Store Wizard** page, click **Finish**.

**See Also**
**Concepts**
Add an Active Directory account store

# Change the server name or IP address for an ADAM account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If the name or IP address of the server that hosts an Active Directory Application Mode (ADAM) instance changes, you must change the corresponding server setting in the properties of the ADAM account store in Active Directory Federation Services (ADFS).

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the server name or IP address for an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Account Stores**.
3. Right-click the ADAM account store whose server name or IP address you want to change, and then click **Properties**.
4. On the **Settings** tab, in **Server name or IP address**, replace the existing server name or IP address with the new value, and then click **OK**.

**See Also**
**Concepts**
Change the port number for an ADAM account store
Change the search base for an ADAM account store
Change the user name attribute for an ADAM account store
Enable or disable TLS and SSL for an ADAM account store
Change the URI for an ADAM account store
Change the display name for an ADAM account store

# Change the display name for an ADAM account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

The ADAM display name identifies the ADAM store to the user during logon. You can change this name in Active Directory Federation Services (ADFS), if needed.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the URI for an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Account Stores**.
3. Right-click the ADAM account store whose display name you want to change, and then click **Properties**.
4. On the **General** tab, in **Display name**, type the new display name, and then click **OK**.

**See Also**
**Concepts**
Change the URI for an ADAM account store
Change the server name or IP address for an ADAM account store
Change the port number for an ADAM account store
Change the search base for an ADAM account store
Change the user name attribute for an ADAM account store
Enable or disable TLS and SSL for an ADAM account store

# Change the port number for an ADAM account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If the communication port that is used by an Active Directory Application Mode (ADAM) instance changes on the host server, you must change the corresponding server setting in the properties of the ADAM account store in Active Directory Federation Services (ADFS). The default port for Lightweight Directory Access Protocol (LDAP) traffic is 389, and ADAM uses this port if a different port is not specified. However, if Active Directory is subsequently installed on the server that hosts the ADAM store, the ADAM port conflicts with the Active Directory port and must be changed for ADAM.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the port number for an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Account Stores**.
3. Right-click the ADAM account store whose port number you want to change, and then click **Properties**.
4. On the **Settings** tab, in **Port number**, replace the existing port number with the new value, and then click **OK**.

**See Also**
**Concepts**
Change the server name or IP address for an ADAM account store
Change the search base for an ADAM account store
Change the user name attribute for an ADAM account store
Enable or disable TLS and SSL for an ADAM account store
Change the display name for an ADAM account store
Change the URI for an ADAM account store

# Change the search base for an ADAM account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

An Active Directory Application Mode (ADAM) store is discovered through a Lightweight Directory Access Protocol (LDAP) search. The search base for an ADAM account store specifies the top node from which the LDAP search for the ADAM instance is to be performed. Typically, the search base is the top node in the directory tree. However, if you want to reduce the search time, you can specify a more restricted path. For example, the search base distinguished name DC=Adatum,DC=com begins searching at Adatum and searches all domains in the path that contains Adatum.com. If the ADAM instance is located on a server in DomainB.DomainA.Adatum.com, you can avoid searching through Adatum and DomainA by specifying the search base of DC=DomainB,DC=DomainA,DC=Adatum.com.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the search base for an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Account Stores**.
3. Right-click the ADAM account store whose search base you want to change, and then click **Properties**.
4. On the **Settings** tab, in **Search base Distinguished Name (DN)**, replace the existing distinguished name with the new value, and then click **OK**.

**See Also**
**Concepts**
Change the server name or IP address for an ADAM account store
Change the port number for an ADAM account store
Change the user name attribute for an ADAM account store
Enable or disable TLS and SSL for an ADAM account store
Change the display name for an ADAM account store
Change the URI for an ADAM account store

# Change the user name attribute for an ADAM account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) can provide authentication for users whose accounts are stored in an Active Directory Application Mode (ADAM) partition. To process user credentials, ADFS must be configured to recognize ADAM users by the attribute that the user provides when logging on to the Web site. This attribute is typically **userPrincipalName** or **sAMAccountName**.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the search base for an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Account Stores**.
3. Right-click the ADAM account store whose user name attribute you want to change, and then click **Properties**.
4. On the **Settings** tab, in **User name attribute**, replace the existing attribute with the new attribute, and then click **OK**.

**See Also**
**Concepts**
Change the server name or IP address for an ADAM account store
Change the port number for an ADAM account store
Change the search base for an ADAM account store
Enable or disable TLS and SSL for an ADAM account store
Change the display name for an ADAM account store
Change the URI for an ADAM account store

# Enable or disable TLS and SSL for an ADAM account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

By default, Transport Layer Security (TLS) and Secure Sockets Layer (SSL) protocols are not enabled for Active Directory Application Mode (ADAM) account stores in Active Directory Federation Services (ADFS). You can enable or disable these protocols as needed.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable or disable TLS/SSL protocols for an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Account Stores**.
3. Right-click the ADAM account store whose search base you want to change, and then click **Properties**.
4. On the **Settings** tab, enable or disable TLS/SSL as follows, and then click **OK**:
   - If the **Enable Transport Layer Security (TLS) and Secure Sockets Layer (SSL) protocols** check box is cleared (disabled) and you want to enable it, select the check box.
   - If the **Enable Transport Layer Security (TLS) and Secure Sockets Layer (SSL) protocols** check box is selected (enabled) and you want to disable it, clear the check box.

**See Also**
**Concepts**
Change the server name or IP address for an ADAM account store
Change the port number for an ADAM account store
Change the search base for an ADAM account store
Change the user name attribute for an ADAM account store
Change the display name for an ADAM account store
Change the URI for an ADAM account store

# Using Multiple Account Stores

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have multiple Active Directory Application Mode (ADAM) account stores that use Web applications that are protected by Active Directory Federation Services (ADFS), you must configure each ADAM store in the Federation Service with a unique uniform resource identifier (URI). In addition, you can set a priority for connecting to each store.

**ADAM Store URI**
Account stores are identified in Active Directory Federation Services (ADFS) by a Uniform Resource Identifier (URI). The account store URI is different from the Federation Service URI.

The URI for an Active Directory account store is always the same (urn:federation:activedirectory) because a forest has only one Active Directory store; this URI is provided automatically by the system. However, you can have multiple instances of Active Directory Application Mode (ADAM) account stores in a Federation Service. To uniquely identify ADAM stores within the trust policy, each ADAM store requires a unique URI. The URI can use a Uniform Resource Locator (URL) format, such as ldap//:*ADAMInstanceName*, or a Uniform Resource Name (URN), such as urn:federation:*ADAMInstanceName*. ADFS does not impose or check the URI format; however, ADFS does check for uniqueness.

When you use the exact ADAM instance name in the URI, ADFS searches that store first. If the URI does not contain the ADAM instance name, ADFS searches all ADAM stores until it finds a match for the user. Therefore, although ADFS accepts the ADAM account store URI as long as it is unique, you can improve efficiency of the ADAM search by including the exact ADAM instance name in the URI.

**Task requirements**

You need the following to perform the procedures for this task:

- Active Directory Federation Services MMC snap-in

To complete this task, perform the following procedures on an as-needed basis:

- Change account store priority
- Change the URI for an ADAM account store

**See Also**
**Other Resources**
Account Stores

# Change account store priority

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If you have more than one account store in a Federation Service, you can determine the order in which Active Directory Federation Services (ADFS) contacts the stores when it authenticates users. If most users are stored in one of the stores, you can set that store to have a higher priority than other stores.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the priority of an account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, and then double-click **My Organization**.
3. Right-click **Account Stores**, and then click **Store Priority**.
4. In the **Account Store Prioritization** dialog box, click the store in the list that you want to move to the top of the list, and then click the **Up** button repeatedly until that store is at the top of the list.
5. If you want to reposition any other stores in the list, select each store and use the **Up** and **Down** buttons as needed.
6. When all stores are in the correct priority order, click **OK**.

# Change the URI for an ADAM account store

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) identifies account stores with a Uniform Resource Identifier (URI). You can change the URI in the account store properties.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the URI for an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then double-click **Account Stores**.
3. Right-click the ADAM account store whose URI you want to change, and then click **Properties**.
4. On the **General** tab, in **URI**, type the new URI, and then click **OK**.

**See Also**
**Concepts**
Change the display name for an ADAM account store
Change the server name or IP address for an ADAM account store
Change the port number for an ADAM account store
Change the search base for an ADAM account store
Change the user name attribute for an ADAM account store
Enable or disable TLS and SSL for an ADAM account store

# Managing Claims and Claim Mapping

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In an Active Directory Federation Services (ADFS) deployment, claims are statements about users that are carried within security tokens and are used by Web applications to make authorization decisions. Claims originate from either an account store or an account partner.

When administering an account store in a Federation Service, you perform the following claim management procedures:

- Create a claim that represents a user or group of users from your organization.
- Extract the claim to map it to a local security account (either a group or user) in your directory service.
- Map the claim to an outgoing identity that you name to appropriately represent the type of users in the claim. During a request for access to a resource, ADFS sends this claim to the resource Federation Service in a security token.

When administering an account partner in a resource Federation Service, you perform the following management procedures:

1. Create a claim for your organization.
2. Map an incoming claim (the name of which has been communicated to you by the account Federation Service) to your local organization claim. This claim is used by the Web server to make authorization decisions about the user or users represented by the claim.

The following tasks for managing claims are described in this objective.

- Exposing Account Store Attributes as Claims
- Mapping Claims as Part of Application Authorization
- Creating, Deleting, and Configuring Claims

**See Also**
**Other Resources**
Understanding Claims

# Exposing Account Store Attributes as Claims

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) uses claims to provide specific information about users to a Web application. Claims are populated with information (attributes) from the account store that hosts the user account. For example, a claim might extract the user's name, identity, key, group, privilege, or capability. The claim is passed in the security token, which the Web server uses to make authorization decisions for access to the requested application.

An organization group claim, when used with an Active Directory Application Mode store, requires mapping the claim to an ADAM attribute.

An organization custom claim maps the claim to an attribute when used with either an ADAM store or an Active Directory store.

**Task requirements**

You must meet the following conditions to perform the procedures for this task:

- ADFS must be installed to create at least one federation server in your forest or realm.
- The Active Directory Federation Services MMC snap-in must be running on the federation server.
- Active Directory or Active Directory Application Mode (ADAM) must be available in the ADFS forest or realm.
- You must have a plan for creating claims and mapping them to the appropriate attributes of a user account in ADAM if you are managing an account Federation Service, or to a set of organization claims if you are managing a resource Federation Service.

To complete this task, perform the following procedures on an as-needed basis:

- Map an organization group claim to an ADAM attribute and value (group claim extraction)
- Map an organization custom claim to an Active Directory or ADAM user attribute (custom claim extraction)

**See Also**
**Other Resources**
Understanding Claims

# Map an organization custom claim to an Active Directory or ADAM user attribute (custom claim extraction)

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Whether you use Active Directory or Active Directory Application Mode (ADAM) as the Active Directory Federation Services (ADFS) account store for an account Federation Service, an organization custom claim maps to an administratively assigned Lightweight Directory Access Protocol (LDAP) attribute for the user that the claim identifies. This mapping is called a custom claim extraction.

For example, if the user is to be identified by position, you might create the organization custom claim Position and use the Title attribute to identify the user's position. If the Title attribute is present in the Active Directory or ADAM store, the corresponding organization custom claim is generated with the value of the Title attribute. Suppose that the Title attribute of the user account has the value "Software Engineer." In this case, the organization custom claim Position is generated for this user with the value "Software Engineer." If the Title attribute is not found for the user account, the Position claim is not generated for the user.

Perform this procedure in the account Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To map an organization custom claim to an Active Directory or ADAM user attribute**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Account Stores**, right-click **ADAM** or **Active Directory**, point to **New**, and then click **Custom Claim Extraction**.
3. In the **Create a New Custom Claim Extraction** dialog box, in **Attribute**, type the LDAP attribute name for the user.
4. In **Map to this Organization Claim**, select the organization custom claim to map to the attribute, and then click **OK**.

# Map an organization group claim to an ADAM attribute and value (group claim extraction)

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you use Active Directory Application Mode (ADAM) as the Active Directory Federation Services (ADFS) account store for a Federation Service, an organization group claim maps to a Lightweight Directory Access Protocol (LDAP) attribute and value of the user account in ADAM. This mapping is called a group claim extraction. For example, suppose that the organization group claim Manager is mapped to the ADAM user account attribute **memberOf** and the value **CN=ADAMTestGroup,CN=Users,DC=adatum,DC=com**. In this case, if the ADAM store user account for the logged-on user has the **memberOf** attribute and that attribute has a value of **CN=ADAMTestGroup,CN=Users,DC=adatum,DC=com**, the organization group claim Manager is generated for the user. If both the **memberOf** attribute and the corresponding value that is specified in the group claim extraction are not present on the user account, the organization group claim is not generated.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To map an organization group claim to an ADAM attribute and value**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Account Stores**, right-click **ADAM**, point to **New**, and then click **Group Claim Extraction**.
3. In the **Create a New Group Claim Extraction** dialog box, in **Attribute** and **Value**, type the LDAP attribute and its value, respectively.
4. In **Map to this Organization Claim**, select the organization group claim to map to the ADAM attribute and value, and then click **OK**.

**See Also**
**Concepts**
Map an organization group claim to an Active Directory group (group claim extraction)
Map an organization custom claim to an Active Directory or ADAM user attribute (custom claim extraction)

# Mapping Claims as Part of Application Authorization

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Claims that are used by Active Directory Federation Services (ADFS) make it possible for an application in one organization to recognize and authorize users from a different organization or from the Internet. By creating claims to represent users in the account organization, and mapping them to a group that is recognized by the resource, you can enable authorization of users for applications that are protected by ADFS in a federated scenario.

**Task requirements**

You must meet the following conditions to perform the procedures for this task:

- ADFS must be installed to create at least one federation server in your forest or realm.
- The Active Directory Federation Services snap-in must be running on the federation server.
- Active Directory or Active Directory Application Mode (ADAM) must be available in the account forest or realm.
- You must have a plan for creating claims and mapping them to the appropriate users and groups if you are managing an account Federation Service, or to a set of local claims if you are managing a resource Federation Service.

To complete this task, perform the following procedures on an as-needed basis:

- Map an organization group claim to an Active Directory group (group claim extraction)
- Map an organization group claim to a resource group
- Create an incoming group claim mapping
- Create an incoming custom claim mapping
- Create an outgoing group or custom claim mapping
- Change the organization claim mapping of an outgoing group or custom claim
- Change the organization claim mapping of an incoming group or custom claim

**See Also**
**Other Resources**
Understanding Claims

# Map an organization group claim to an Active Directory group (group claim extraction)

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

When you use Active Directory as the Active Directory Federation Services (ADFS) account store for an account Federation Service, you map an organization group claim to a security group in Active Directory. This mapping is called a group claim extraction.

Perform this procedure in the account Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To map an organization group claim to an Active Directory group**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, double-click **Account Stores**, right-click **Active Directory**, point to **New**, and then click **Group Claim Extraction**.
3. In the **Create a New Group Claim Extraction** dialog box, click **Add**, and then select the Active Directory security group that you want to map to a group claim.
4. In **Map to this Organization Claim**, select the group claim to map to the Active Directory security group, and then click **OK**.

**See Also**
**Concepts**
Map an organization group claim to an ADAM attribute and value (group claim extraction)
Map an organization custom claim to an Active Directory or ADAM user attribute (custom claim extraction)

# Map an organization group claim to a resource group

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Organization group claims in the resource Federation Service can be mapped to resource groups, which reside in Active Directory, so that federated users who have the organization group claim are treated as members of that resource group. Therefore, you must create a security group (which is the resource group) in the resource partner forest that represents the users who have the organization group claim but do not have domain user accounts in the resource partner forest. After you create the organization group claim, map the claim to the resource group that you created in the resource partner forest.

| **Note** |
| --- |
| This procedure is not required if the resource partner forest trusts the account partner forest and the Windows trust option is selected in the Active Directory Federation Services snap-in in both the account Federation Service and the resource Federation Service. |

Perform this procedure on a resource federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To map an organization group claim to a resource group**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then click **Organization Claims**.
3. In the details pane, right-click the organization group claim that requires mapping to a local group, click **Properties**, and then click the **Resource Group** tab.
4. Select the **Map this claim to the following local resource group** check box, and then click the **...** button.
5. In **Enter the object name to select**, type the name of the resource group that you want to map to the group claim, and then click **OK**.
6. In the **Group Claim Properties** dialog box, click **OK**.

**See Also**
**Concepts**
When to use resource groups

# Create an incoming group claim mapping

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), incoming group claim mappings are used in the resource Federation Service to transform group claims that are sent by an account partner into organization claims that can be used by the resource partner to make authorization decisions.

For example, an account partner might send a security token for a user that contains the group claim SalesReps. Because the resource partner cannot make authorization decisions based on the account user's membership in the SalesReps group, an incoming group claim mapping is used to map the incoming group claim that is named SalesReps in the account Federation Service to the organization group claim that is named Purchasers in the resource Federation Service. The resource itself provides access to the local security group to which the Purchasers claim is mapped.

Perform this procedure on a resource federation server. To perform this procedure, you must have created an organization group claim to which you can map the incoming claim.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create an incoming group claim mapping**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Account Partners**, right-click your account partner, point to **New**, and then click **Incoming Group Claim Mapping**.
3. In the **Create a New Incoming Group Claim Mapping** dialog box, in **Incoming group claim name**, type the name of the group claim that your account partner sends to you.
4. In **Organization group claim**, select the group claim that you will use in your organization to map the incoming group claim to, and then click **OK**.

**See Also**
**Concepts**
Create an organization group or custom claim
Change the organization claim mapping of an incoming group or custom claim

# Create an incoming custom claim mapping

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), an organization custom claim maps to a user attribute. Incoming custom claim mappings are used in the resource Federation Service to map custom claims that are sent by an account partner to claims that can be used by the resource partner to make authorization decisions.

For example, an account partner might send a security token for a user that contains the custom claim EmployeeID, which maps to a user attribute in the account partner directory database. Because the resource partner cannot make authorization decisions based on the account user's EmployeeID attribute value, an incoming custom claim mapping is used to map the custom organization claim EmployeeID, which is recognized in the account partner, to the organization claim CustomerID, which is recognized in the resource partner.

Perform this procedure on a federation server in the resource Federation Service. To complete this procedure, you must have created an organization custom claim to which you can map the incoming claim.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create an incoming custom claim mapping**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Account Partners**, right-click your account partner, click **New**, and then click **Incoming Custom Claim Mapping**.
3. In the **Create a New Incoming Custom Claim Mapping** dialog box, in **Incoming custom claim name**, type the name of the custom claim that your account partner sends to you.
4. In **Organization custom claim**, select the custom claim that you will use in your organization to map the incoming custom claim to, and then click **OK**.

**See Also**
**Concepts**
Create an organization group or custom claim
Change the organization claim mapping of an incoming group or custom claim

# Create an outgoing group or custom claim mapping

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), an organization claim (group or custom) in the account Federation Service must be mapped to an outgoing claim, which the resource Federation Service will receive when an account organization user is requesting access to a resource. On the resource federation server, this claim is received as an incoming claim, which is likewise configured to map to a local organization claim, which the resource Federation Service uses to make authorization decisions.

Perform this procedure on an account federation server. To perform this procedure, you must have created an organization group or custom claim to which you can map the outgoing claim.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create an outgoing group or custom claim mapping**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Resource Partners**, right-click the resource partner, point to **New**, and then click one of the following:

   If you are mapping a group claim, click **Outgoing Group Claim Mapping**, and then create the mapping as follows:

   a. In the **Create a New Outgoing Group Claim Mapping** dialog box, in **Organization group claims**, select the group claim in the account organization that you want to map to the outgoing claim.

   b. In **Outgoing group claim name**, type the name of the outgoing group claim that you want to send to the resource partner, and then click **OK**.

   If you are mapping a custom claim, click **Outgoing Custom Claim Mapping**, and then create the mapping as follows:

   a. In the **Create a New Outgoing Custom Claim Mapping** dialog box, in **Organization custom claims**, select the custom claim in the account organization that you want to map to the outgoing claim.

   b. In **Outgoing custom claim name**, type the name of the outgoing custom claim that you want to send to the resource partner, and then click **OK**.

**See Also**
**Concepts**
Create an organization group or custom claim
Change the organization claim mapping of an outgoing group or custom claim

# Change the organization claim mapping of an outgoing group or custom claim

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), when you create an organization group claim or organization custom claim in an account Federation Service, you map it to an outgoing claim that is sent to the resource Federation Service. If you subsequently replace the organization claim with a different claim or want to use a different organization claim for the group or user identified in the claim, you can change the mapping of the existing outgoing claim to the new organization custom or group claim. You can also rename the outgoing custom or group claim.

Perform this procedure in the account Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the organization claim mapping of an outgoing group or custom claim**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Resource Partners**, and then click the resource partner whose organization claim mapping you want to change.
3. In the details pane, right-click the organization group or custom claim whose outgoing claim mapping you want to change, and then click **Properties**.
4. On the **General** tab, in **Organization group claims** or **Organization custom claims**, select the new organization group or custom claim that you want to map to the outgoing group or custom claim.
5. In **Outgoing group claim name** or **Outgoing custom claim name**, change the name, if needed, and then click **OK**.

# Change the organization claim mapping of an incoming group or custom claim

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), when you create an organization group claim or organization custom claim in a resource Federation Service, you map it to an incoming claim that is sent from the account Federation Service. If you subsequently replace the organization claim with a different claim or want to use a different organization claim for the group or user identified in the claim, you can change the mapping of the existing incoming claim to the new organization custom or group claim. You can also rename the incoming custom or group claim.

Perform this procedure in the resource Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the organization claim mapping of an incoming group or custom claim**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Account Partners**, and then click the account partner whose organization claim mapping you want to change.
3. In the details pane, right-click the incoming claim whose claim mapping you want to change, and then click **Properties**.
4. On the **General** tab, in **Organization group claim** or **Organization custom claim**, select the new organization group or custom claim to which the incoming group or custom claim will map.
5. In **Incoming group claim name** or **Incoming custom claim name**, change the name, if needed, and then click **OK**.

# Creating, Deleting, and Configuring Claims

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Use the procedures in this task to create, delete, and configure auditing for Active Directory Federation Services (ADFS) organization group and custom claims. You can also configure a claims transform module and manage identity claims.

**Task requirements**

You must meet the following conditions to perform the procedures for this task:

- ADFS must be installed to create at least one federation server in your forest or realm.
- The Active Directory Federation Services snap-in must be running on the federation server.
- Active Directory or Active Directory Application Mode (ADAM) must be available in the account forest or realm.
- You must have a plan for creating claims and mapping them to the appropriate users and groups if you are managing an account Federation Service, or to a set of local claims if you are managing a resource Federation Service.

To complete this task, perform the following procedures on an as-needed basis:

- Create an organization group or custom claim
- Delete an organization group or custom claim
- Change the auditing limitation for an organization group or custom claim
- Configure a claims transform module
- Change the domain suffix for an incoming or outgoing e-mail claim
- Enable or disable an organization identity claim for an account or resource partner

**See Also**
**Other Resources**
Understanding Claims

# Create an organization group or custom claim

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), an organization group claim is used by the Federation Service to represent a user's membership in a group or role. An organization custom claim is used by the Federation Service to provide custom information, such as an employee identification number, about a user. Group claims and custom claims are created in the same way. You can use the following procedure to create an organization group claim or a custom claim.

Administrators in the account Federation Service create organization group claims and custom claims to represent account users. Administrators in the resource Federation Service create corresponding organization group claims and custom claims to represent groups and users that can be recognized as resource users in the resource Federation Service. Because outgoing claims in the account Federation Service map to incoming claims in the resource Federation Service, the resource Federation Service is able to accept the credentials that are provided by the account Federation Service.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To create an organization group or custom claim**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, right-click **Organization Claims**, point to **New**, and then click **Organization Claim**.
3. In the **Create a New Organization Claim** dialog box, in **Claim name**, type a unique name that identifies the group or role for a group claim, or type the name of a custom claim.

   | **Note** |
   | --- |
   | You can create a group claim and a custom claim with the same name. |

4. If you want to limit the auditing of the claim to the claim name and omit the value, select the **Limit the auditing of this claim** check box.
5. For **Specify the claim type**, click **Group claim** or **Custom claim**, and then click **OK**.

**See Also**

**Concepts**

Create an incoming group claim mapping
Create an incoming custom claim mapping
Create an outgoing group or custom claim mapping

# Delete an organization group or custom claim

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

If a group or user no longer requires access to a resource that is protected by Active Directory Federation Services (ADFS), you can delete the organization group claim or custom claim that represents the group or user to reduce management overhead.

For example, an account partner might have provided a group claim or custom claim to represent a capability of a group or role. If this condition changes and the group claim or custom claim is no longer valid in the resource domain, you can delete the group claim.

At the account partner, a group that existed before might no longer be valid. In this case, the corresponding group claim can be removed.

If a partner has been deleted and any claims are specific to the deleted partner, you can delete these claims as well.

Perform this procedure on a federation server in your organization.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To delete an organization group or custom claim**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then click **Organization Claims**.
3. In the details pane, right-click the claim to be deleted, and then click **Delete**.
4. In the **Delete group claim** message box, click **Yes** to confirm the deletion.

# Change the auditing limitation for an organization group or custom claim

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) uses claims to transfer user information that is used for authorization in the resource Federation Service. By default, auditing of an organization group or custom claim is not limited, which means that both the claim name and its value are audited, or shared, when the claim is produced or mapped. If you limit the auditing of an organization group or custom claim, the ADFS audit logs indicate only the name of the claim; the value of the claim is omitted. Omitting the value protects the privacy of sensitive data, such as the groups to which a user belongs.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To change the auditing limitation on an organization group or custom claim**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **My Organization**, and then click **Organization Claims**.
3. In the details pane, right-click the organization group or custom claim whose auditing you want to change, and then click **Properties**.
4. Click to select or clear the **Limit the auditing of this claim** check box to change the existing setting.

# Configure a claims transform module

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

A claims transform module is custom code that manipulates organization, input (incoming), and output (outgoing) claims. Typically, transform modules use the corporate and input claims to produce additional output claims. However, the claim transform module can enumerate, add, delete, and modify claims in any of the claim sets.

Store the claims transform module in %systemdrive%\adfs\sts\bin. This location provides the following advantages:

- ASP.NET keeps a shadow copy of the dynamic-link library (DLL), which allows the DLL to be replaced without stopping the Federation Service, thereby preventing downtime.
- File security is inherited from the \adfs\sts directory.
- The module can be backed up along with all other Active Directory Federation Services (ADFS) files.

After you deploy the transform module to the federation server, perform the following procedure on the account federation server or resource federation server that is configured with the trust policy whose claims transform module you are configuring.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure a claims transform module**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click the **Trust Policy** node, and then click **Properties**.
3. On the **Transform Module** tab, configure the DLL file and class name for the module as follows:
   a. **DLL file:** Click **Browse** to navigate to the DLL that implements the claim transform module, and then click **Open**.

      | **Note** |
      | --- |
      | This DLL must be a managed-code assembly. |

   b. **Class name:** Type the namespace-qualified class name that implements the claim transform interface (**IClaimTransform**, which is defined in System.Web.Security.SingleSignOn.ClaimTransforms.dll): The namespace qualified name should be of the format namespace.classname.
4. Click **OK** to save the configuration.

# Change the domain suffix for an incoming or outgoing e-mail claim

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), an e-mail identity claim is an organization identity claim that is created when you create an account or resource partner. The e-mail identity claim is automatically mapped to the outgoing (from the account Federation Service) and incoming (to the resource Federation Service) e-mail claim. In the account Federation Service, you specify whether all domain suffixes can be sent to the resource Federation Service, or whether you will replace all domain suffixes with a specific domain suffix. In the resource Federation Service, you can specify the domain suffixes that the resource Federation Service will accept or, if your partnership does not extend across the Internet, you can accept all domain suffixes.

| **Note** |
| --- |
| To manage the properties of an identity claim, the claim must first be enabled. |

### Administrative credentials

To complete these procedures, you must be a member of the Administrators group on the local computer.

**To change the domain suffix for an outgoing e-mail identity claim in the account Federation Service**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Resource Partners,** and then click the resource partner whose e-mail claim you want to manage.
3. In the details pane, right-click the e-mail identity claim, and then click **Properties**.
4. Under **Settings**, select **Replace all domain suffixes with** if it is not selected.
5. In **Specify the new suffix**, type the domain name suffix, and then click **OK**.

**To change the domain suffixes for an incoming e-mail identity claim in the resource Federation Service**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Account Partners,** and then click the account partner whose e-mail claim you want to manage.
3. In the details pane, right-click the e-mail identity claim, and then click **Properties**.
4. Under **Settings**, select **Accept some domain suffixes** if it is not selected.
5. In **Accepted domains (press Enter to separate entries)**, type the domain name suffix or suffixes, and then click **OK**.

**See Also**
**Concepts**
Enable or disable an organization identity claim for an account or resource partner

# Enable or disable an organization identity claim for an account or resource partner

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

In Active Directory Federation Services (ADFS), organization identity claims are created when you create an account or resource partner. They are incoming claims on the resource partner and outgoing claims on the account partner. Identity claims are not enabled unless you specify the identity type (UPN, e-mail, or common name) when you create the partner. You can enable or disable organization identity claims after the partner is created.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable or disable an organization identity claim**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, double-click **Account Partners** or **Resource Partners**, and then click the partner whose organization identity claim you want to manage.
3. In the details pane, right-click the organization identity claim that you want to enable or disable, and then click **Properties**.
4. On the **General** tab, enable or disable the identity claim as follows, and then click **OK**:
   - To enable the claim when it is disabled, select the **Enabled** check box.
   - To disable the claim when it is enabled, clear the **Enabled** check box.

# Troubleshooting Active Directory Federation Services

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

This guide provides troubleshooting information for Active Directory Federation Services (ADFS) in the Microsoft Windows Server 2003 R2, Enterprise Edition, and Microsoft Windows Server 2003 R2, Datacenter Edition, operating systems (for Federation Service, Federation Service Proxy, and ADFS Web Agent components); and the Microsoft Windows Server 2003 R2, Standard Edition, operating system (for ADFS Web Agent components).

**In this guide**


- Verifying ADFS Computer Settings and Connectivity
- Configuring ADFS Servers for Troubleshooting
- Troubleshooting ADFS Problems



**Acknowledgments**

Produced by: Microsoft Windows Server User Assistance team

Project Writer: Mary Hillman

Project Editor: Femila Anilkumar

Contributing Writers: Nick Pierson, Edward Gomes

Technical Reviewers: Derek Del Conte, Vijay Gajjala, Dan Hartop, Ryan D. Johnson, Jagadeesh Kalki, Carol Li, Harini Raghavan, Vani Nori

# Verifying ADFS Computer Settings and Connectivity

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Before you begin troubleshooting, you must isolate the location of the problem by performing preliminary tests to verify that your federation servers, Web servers, and client computers are communicating and that Active Directory Federation Services (ADFS) is set up and running properly.

**Verifying Settings to Locate the Point of Failure**
ADFS components are deployed among several servers, each of which requires settings that affect various server-server and client-server interactions. The best way to effectively troubleshoot and solve a problem is to find the exact location of failure between ADFS components.

To help you better understand how to locate the point of failure, the following sample scenario is provided.

**Sample Scenario**
Using this sample diagram of a client requesting access to an ADFS-protected application in a Federated Web with Single Sign-on (SSO) scenario, you can easily distinguish how ADFS servers enumerate connections. This information is meant to assist you in determining which connection point might be causing the failure. If you are already familiar with the location of the failure in your environment, skip to the section where you suspect the failure is occurring. By testing for success at each step, you can better determine where the problem lies.

For more details and instructions for implementing this sample Federated Web SSO scenario, see the ADFS Step-by-Step Guide (http://go.microsoft.com/fwlink/?linkid=63445).




| Note |
| --- |
| To improve user experience, always obtain Secure Sockets Layer (SSL) certificates for any external (Internet-facing) computers from a well-known third-party certification authority (CA). Because well-known third-party CAs are already trusted by the client browser, the user does not receive prompts that the browser does not trust the root CA certificate. |

**Verification Steps to Locate the Point of Failure**
Each topic below corresponds to a number in the preceding diagram and provides information and step-by-step procedures to help eliminate the error that is causing the failure.

1. Verify Connectivity and Initial Request from the Client

2. Verify Web Server Redirection to the Resource Federation Server

3. Verify Home Realm Discovery

4. Verify Client Authentication in the Account Domain

5. Verify Account Server Redirection to the Resource Federation Server

6. Verify Resource Server Redirection to the Web Server

**1. Verify Connectivity and Initial Request from the Client**
In this step, you verify that the ADFS client can reach the Web site on the Web server before the client has received any

previous token.

When you enter the uniform resource locator (URL) of the Web application in the client browser for the first time, the client sends an initial GET request to the ADFS Web server. The ADFS Web server does not allow the client access to the requested page at this point because the client does not have an authentication token to present to the Web server at that point. If the ADFS client fails to receive access to the page, you get an error and the URL in the browser indicates the point of failure. Common errors might be Domain Name System (DNS) failures and 401 access denied errors.

Make sure you have good connectivity and name resolution between the account and resource domains. Clients coming from the Internet should be able to locate domain controllers.

Domain controllers, federation servers, Web servers, and clients need to be able to locate each other by a fully qualified domain name (FQDN). Ping by IP address and FQDN, and use Nslookup.exe to test DNS connectivity. In addition, to make sure that all the necessary ports are open, you can download PortqryUI.exe from the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=64729

Use the following tests to verify connectivity from the client:

- Test a TCP/IP configuration by using the ping command (http://go.microsoft.com/fwlink/?LinkId=64765)
- Verify DNS server responsiveness using the **nslookup** command (http://go.microsoft.com/fwlink/?LinkId=64766)

If DNS resolution appears to be the connectivity problem, see Troubleshooting Domain Name System (http://go.microsoft.com/fwlink/?LinkId=62177). You should also check the Web server to verify that the **ASP.NET v2.0.50727** extension is set to **Allowed** using the Internet Information Services (IIS) Manager snap-in. The typical symptom that occurs when this extension is set to **Prohibited** is a ***Page cannot be displayed error*** in the browser.

**2. Verify Web Server Redirection to the Resource Federation Server**

In this step, verify that the redirect of the ADFS client to the resource federation server is occurring properly and that the client can reach the Federation Service endpoint URL.

| **Note** |
| --- |
| The Federation Service URL is used by clients to access a server in the resource Federation Service. This URL has the form **https://***FullyQualifiedDomainName***/adfs/ls/** |

The ADFS Web server issues a standard HTTP 302 REDIRECT to the client, which directs the client to the resource federation server. This redirect occurs to the resource federation server because the ADFS Web server knows about only its own federation server and requires that incoming ADFS tokens are signed by its federation server. The ADFS client is able to communicate with the resource federation server because it trusts the CA that issues the SSL server certificate for the resource federation server. As in Step 1 (Verify Connectivity and Initial Request from the Client), if the client fails to contact the resource federation server, an error appears and the URL on the browser indicates the point of failure. Common errors are DNS failures and 401 access denied. In this case, use the tests that are described in step 1.

| **Important** |
| --- |
| If the Federation Service endpoint URL value has recently been modified in the Windows NT token–based Web agent, clients may experience logon problems to Windows NT token–based applications that are located on the Web server within one hour from the time when the URL value was changed. After the one hour has passed, clients should be able to log on to the Windows NT token–based applications without problems. To avoid this one-hour delay, you can run **iisreset** from a command line immediately after the URL value changes. |

**3. Verify Home Realm Discovery**

In this step, verify that the default page is displayed in the browser and the proper home realm of the client is available in the list.

When the ADFS client first requests a resource, the resource federation server has no information about the realm of the client. The resource federation server responds to the ADFS client with an HTTP 200, which provides the Client Realm Discovery page, where the user selects the home realm from a list. The list values are populated from the display name property in the account Federation Service trust policy. This value is known to the resource Federation Service through the account partner node properties.

**The client realm discovery page fails to appear**

If the client realm discovery page fails to appear, do the following on the client computer:

1. Delete all cookies from the client browser to remove the persistent home realm cookie, if present.

2. Try again to connect to the Web site. The Web server attempts to redirect to the resource federation server.

### Administrative credentials

To complete this procedure, you must be a member of the Users group on the local computer.

**To delete all cookies in Internet Explorer**

1. In Internet Explorer, on the **Tools** menu, click **Internet Options**.
2. On the **General** tab, click **Delete Cookies**.
3. In the **Delete Cookies** dialog box, click **OK**.
4. In the **Internet Options** dialog box, click **OK**.

**Deleting cookies does not solve the problem**
If the ADFS client still fails to receive access to the realm discovery page, the problem is with the resource Federation Service because this redirection is the final step before the client is sent to the account Federation Service.

In this case, do the following:

1. Look at the URL in the Web browser of the client to determine whether the failure is occurring on the Web server or on the resource federation server.

2. If the URL still shows the Web server address and nothing further, you can assume that the issue is on the Web server. In this case, go to the next step.

3. Look at the Application log on the Web server to see what errors are logged. Review the User Action section of the event and take proper troubleshooting steps according to any errors. If these steps do not solve the problem, go to the next step.

4. Use the following procedure to confirm that the Web server can communicate with the Federation Service on the resource federation server.

   **Administrative credentials**

   To complete this procedure, you must be a member of the Administrators group on the local computer.

   **To check the Federation Service URL**

   a. On the ADFS Web server, open Internet Information Services (IIS) Manager, right-click **Web Sites**, and then click **Properties**.

   b. On the **ADFS Web Agent** tab, copy the **Federation Service URL**.

   c. Paste the URL into the browser address line, and then press Enter.

If the URL resolves properly, a Federation Service Web page appears that lists supported operations for the service.

| **Note** |
| --- |
| You might be prompted for a client authentication certificate, which is standard behavior. |

**The Federation Service URL does not resolve**
If the URL does not resolve, assuming you have verified name resolution and connectivity, the error is most likely a typing error in the Federation Service URL on the Web server. Use the following procedure to check for potential typing errors for the Windows NT token-based Web agent.

**To check the Federation Service URL of the Windows NT token-based Web agent**

1. On the ADFS Web server, open Internet Information Services (IIS) Manager, right-click **Web Sites**, and then click **Properties**.

2. On the **ADFS Web Agent** tab, check that the URL specified in the **Federation Service URL** field is typed correctly.

Use the following procedure to check for potential typing errors in the web.config file that is associated with the claims-aware application.

**To check the Federation Service URL in the web.config file**

1. Using Notepad, edit the web.config file associated with your claims-aware application.

2. Find the Federation Service URL value in the web.config file by locating the **<fs>FederationServiceURLValue</fs>** entry. Check that the value and format specified in **FederationServiceURLValue** is typed correctly. An example of the Federation Service URL located in the web.config file used in the ADFS Step-by-Step guide is **<fs>https://adfsresource.treyresearch.net/adfs/fs/federationserverservice.asmx</fs>**.

**Confirm that Web service extensions are configured properly**
Confirm that **ASP.NET v2.0.50727** (for claims-aware applications) and/or the **ADFS Web Component Extension** (for Windows NT token-based applications) are set to **Allowed** in the Internet Information Services (IIS) Manager snap-in on the ADFS Web server.

When the ADFS client successfully reaches the home realm discovery page and selects its domain, the ADFS client submits a POST back to the resource federation server that includes the home realm. In response, the resource federation server issues a 302 redirect to the ADFS client so that the client can be authenticated for its domain and receive the appropriate claims from the account federation server.

**4. Verify Client Authentication in the Account Domain**
If the address in the ADFS client browser indicates failure of client redirection from the resource federation server back to the account federation server, the client cannot be authenticated by a domain controller in the account forest. In this case, check the following:

- The value of the Federation Service endpoint URL that is configured on the account partner node in the resource Federation Service trust policy is the same value as the Federation Service endpoint URL that is configured on the Trust Policy node in the account Federation Service. These values must match.

- Whether or not the user account is configured to use the **Smart card is required for interactive logon** option in the account properties. If this option is enabled for a user account that is attempting to access an ADFS-enabled Web application, authentication to that application will not be successful.

**Disable Java Scripting**
JavaScript is used to automatically redirect the client to various points, including posting tokens. When JavaScript is disabled, the automatic redirect will be prevented and a submit button will be displayed instead. This button allows the client to walk through each step more easily as part of troubleshooting the configuration.

### Administrative credentials

To complete this procedure, you must be a member of the Users group on the local computer.

**To disable JavaScript in Internet Explorer**

1. Open Internet Explorer.

2. On the **Tools** menu, click **Internet Options**.

3. On the **Security** tab, click **Custom Level**.

4. Scroll to the **Scripting** category.

5. Under **Active scripting**, click **Disable**, and then click **OK** twice.

**Check the Federation Service endpoint URL**
Use the following procedures to check the value of the Federation Service endpoint URL on the Trust Policy node in the account Federation Service against the value on the account partner node in the resource Federation Service.

### Administrative credentials

To complete these procedures, you must be a member of the Administrators group on the local computer.

**To check the Federation Service endpoint URL in the account Federation Service**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, right-click the **Trust Policy** node, and then click **Properties**.
3. On the **General** tab, check the value in the **Federation Service endpoint URL** box. If you need to change the value, select the domain portion of the URL, replace the selected text with the new URL, and then click **OK**.

Use the following procedure to check the Federation Service endpoint URL on the account partner node in the resource Federation Service.

**To check the Federation Service endpoint URL in the resource Federation Service**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Double-click **Federation Service**, double-click **Trust Policy**, double-click **Partner Organizations**, and then double-click **Account Partners** if you are logged on to the resource federation server, or **Resource Partners** if you are logged on to the account federation server.
3. Right-click the account or resource partner whose Federation Service endpoint URL has changed, and then click **Properties**.
4. On the **General** tab, check the value in the **Federation Service endpoint URL** box. If you need to change the value, select the domain portion of the URL, replace the selected text with the new URL, and then click **OK**.

**5. Verify Account Server Redirection to the Resource Federation Server**
After the client is authenticated and the account federation server receives authorization information from the domain controller, the account federation server builds the SAML token with the claims for the user, writes the cookie to the ADFS client. The account federation server then issues a POST REDIRECT with the SAML token and the redirect URL of the resource federation server. The account federation server signs the SAML token with its token-signing certificate. The resource federation server must have been configured with a verification certificate from the account federation server. The verification certificate contains the public key portion of the token-signing certificate of the account federation server. The resource federation server requires this certificate to verify that the account federation server sent the token.

If the verification certificate from the account federation server is not installed on the resource federation server, token validation fails. Use the following procedure to check for the existence of the verification certificate on the resource federation server.

View the current verification certificate

**The verification certificate is invalid**
If the verification certificate is present but has expired, follow the steps to roll over the verification certificate.

Rolling Over a Token-signing Certificate

**The verification certificate is missing**
If the verification certificate is not present, use the following procedure to first export the public portion of the token-signing certificate on the account federation server, and then you will need to import it into the certificate store on the resource federation server.

Export the public key portion of a token-signing certificate

**6. Verify Resource Server Redirection to the Web Server**
After receiving the token from the account federation server, the following events occur:

1. Resource federation server maps the incoming claims in the token to the appropriate claims for the ADFS Web server, and then issues a token for the ADFS Web server.
2. The resource federation server writes the cookie to the ADFS client and provides the token for the ADFS Web server to the client, along with the URL for the ADFS Web server. This information is provided to the client in the form of a POST REDIRECT.
3. The resource federation server signs the token with its own token-signing certificate. The ADFS Web server uses the

public key to verify that the token is signed by the resource federation server. The ADFS Web obtains the public key by making a call to the resource federation server. By using this public key, the ADFS Web server can verify that the token is signed by the resource federation server. The ADFS Web server then allows the appropriate access to the client.

If redirection to the ADFS Web server fails, the page fails to appear. Causes for failure can include:

- The application URL in the trust policy does not match the return URL specified in the ADFS Web Agent tab on the Web server hosting the application.
- The endpoint URL is incorrect for a federation server or federation server proxy

**The application URL does not match the return URL**
If the page fails to appear with an error such as "Server Error in '/adfs' Application," use the following procedures to check that the application URL in the properties of the application node on the resource federation server matches the return URL in the ADFS Web Agent tab on the application Web site.

1. Set the return URL for a Windows NT token-based application
2. Set the return URL for a claims-aware application
3. Set the application URL for an application

You can also check the security and application logs in Event Viewer on the resource federation server and ADFS Web server for warnings and errors that provide problem-specific information.

**The endpoint URL is incorrect for a federation server proxy**
If a federation server proxy is in place in a perimeter network in the adatum.com domain, both the resource federation server in Treyresearch.net and the account federation server in Adatum.com must be configured with the location of the federation server proxy.

If you use a federation server proxy in the account domain, check the endpoint URL on both the account partner node and resource partner node on the respective federation servers. If the URL does not point to the federation server proxy, change its value. Use the procedure "To change the Federation Service endpoint URL" in the following procedure:

Change the Federation Service endpoint URL

**The client authentication certificate for the federation server proxy is not configured**
You must ensure that an SSL certificate and a client authentication certificate is installed on each server that is running the federation server proxy component. The public key portion of the client authentication certificate must be present on the account federation server and the resource federation server. If you have more than one federation server proxy, the same SSL and client authentication certificate can be shared among all federation server proxies. Alternately, you can obtain a separate SSL and client authentication certificate for each federation server proxy, export each certificate to a file, and import all of them to the Federation Service trust policy.

| **Note** |
| --- |
| When certificates are shared among servers in the same role, you export both the public and private portions of the certificates. |

In the account Federation Service, if the client is internal, the client goes directly to the account federation server and the steps, as illustrated in the sample Federated Web SSO diagram, occur. If the client requests access remotely from the Internet, the client is redirected to the federation server proxy.

Use the following information and procedures for configuring the client authentication certificate and installing its public key on the federation servers.

Managing Client Authentication Certificates

**See Also**
**Concepts**
Configuring ADFS Servers for Troubleshooting

# Configuring ADFS Servers for Troubleshooting

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

This topic provides the following sections to help you configure and understand settings and logs that can aid problem-solving on Active Directory Federation Services (ADFS) servers.

- Configuring Cache Settings on ADFS servers
- Configuring ADFS Servers to Record Auditing of ADFS Events to the Security Log
- Configuration Tasks for Troubleshooting
- Understanding Default Event Logs

**Configuring Cache Settings on ADFS servers**

The following table describes all of the various ADFS-related cache settings, where they are located, the values associated with them and the ADFS components they affect.

| Cache Setting Description | Cache Setting Location | ADFS Component / Sub-Component | Configuration Parameter | Units | Default Value | Minimum Value | Maximum Value |
|---|---|---|---|---|---|---|---|
| Lifetime of an entry in the Security Assertion Markup Language (SAML) token cache | Registry(DWORD Value) | Windows NT Token-based Web Agent / ADFS Authentication Service | System\\CurrentControlSet\\Services\\IFSSVC\\Parameters**TokenCacheEntryLifetime** | Minutes | 60 minutes | 5 minutes | Maximum DWORD |
| How often the SAML token cache is scavenged | Registry (DWORD Value) | Windows NT Token-based Web Agent / ADFS Authentication Service | System\\CurrentControlSet\\Services\\IFSSVC\\Parameters**TokenCacheScavengePeriod** | Minutes | 15 minutes | 5 minutes | Maximum DWORD |
| Total number of entries in the SAML token cache | Registry (DWORD Value) | Windows NT Token-based Web Agent / ADFS Authentication Service | System\\CurrentControlSet\\Services\\IFSSVC\\Parameters**TokenCacheSize** | N/A | 10000 | 0 | Maximum DWORD |
| How long cookie cache entries are guaranteed to live | Registry (DWORD Value) | Windows NT Token-based Web Agent / ADFS Web Agent ISAPI Extension | Software\\Microsoft\\ADFS\\WebServerAgent**CookieCacheEntryLifetime** | Seconds | 5 minutes | 1 minute | 8 hours |
| How often the cookie cache is scavenged | Registry (DWORD Value) | Windows NT Token-based Web Agent / ADFS Web Agent ISAPI Extension | Software\\Microsoft\\ADFS\\WebServerAgent**CookieCacheScavengeInterval** | Seconds | 15 minutes | 1 minute | 1 hour |
| How long token-information cache entries are guaranteed to live | Registry (DWORD Value) | Federation Service and Windows NT Token-based Web Agent | System\\CurrentControlSet\\Control\\Lsa\\WebSso\\Parameters**CacheEntryLifetime** | Seconds | 1 hour | 1 minute | 10 hours |
| How often the token information cache is scavenged | Registry (DWORD Value) | Federation Service and Windows NT Token-based Web Agent | System\\CurrentControlSet\\Control\\Lsa\\WebSso\\Parameters**CacheScavengeInterval** | Seconds | 15 minutes | 1 minute | 1 hour |
| Total number of entries in the SAML token cache | Web.config (Integer) | Claims-aware Web Agent | /websso/**tokenCacheSize** | N/A | 10000 | 0 | Integer Maximum Value |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| Lifetime of an entry in the SAML token cache | Web.config (Integer) | Claims-aware Web Agent | /websso/**tokenCacheEntryLifetime** | Minutes | 60 minutes | 5 minutes | Integer Maximum Value |
| How often the SAML token cache is scavenged | Web.config (Integer) | Claims-aware Web Agent | /websso/**tokenCacheScavengePeriod** | Minutes | 15 minutes | 5 minutes | Integer Maximum Value |
| How often the Windows trust information is updated | Trust policy (Integer) | Federation Service | **WindowsTrustCacheUpdatePeriodInMinutes** | Minutes | 60 minutes | 5 minutes | Integer Maximum Value |
| Total number of entries in the SAML token cache | Trust policy (Integer) | Federation Service | **TokenCacheCapacity** | N/A | 10000 | 0 | Integer Maximum Value |
| Lifetime of an entry in the SAML token cache | Trust policy (Integer) | Federation Service | **TokenCacheEntryLifetimeInMinutes** | Minutes | 60 minutes | 5 minutes | Integer Maximum Value |
| How often the SAML token cache is scavenged | Trust policy (Integer) | Federation Service | **TokenCacheScavengePeriodInMinutes** | Minutes | 15 minutes | 5 minutes | Integer Maximum Value |
| How often the trust policy is checked for updates (this is in addition to change notifications in the trust policy file) | Trust policy (Integer) | Federation Service | **TrustPolicyUpdatePeriodInMinutes** | Minutes | 60 minutes | 5 minutes | Integer Maximum Value |

**Configuring ADFS Servers to Record Auditing of ADFS Events to the Security Log**
All ADFS-related audits that are made specifically to the security log are considered by the system to be object access–type audits, which by default are ignored by the system. For this reason, to ensure that ADFS-related audits (specifically **Success Audits** and **Failure Audits**) appear in the Security log, you need to manually configure the Local Security Policy, using the procedure below. You must apply the steps in this procedure to all of the ADFS servers (federation servers, federation server proxies, and Web servers hosting the ADFS Web Agent) before enabling success or failure auditing in the Trust Policy properties of the ADFS snap-in. This will allow the Federation Service to log either success or failure errors.

This procedure has no effect on the events that ADFS writes to the application log.

**To configure the Windows Security Log to support auditing of ADFS events**

1. Click **Start**, point to **Administrative Tools**, and then click **Local Security Policy**.
2. Double-click **Local Policies**, and then click **Audit Policy**.
3. In the details pane, double-click **Audit object access**.
4. On the **Audit object access Properties** page, select either **Success** or **Failure**, or both, and then click **OK**.
5. Close the **Local Security Settings** snap-in.
6. At a command prompt, type **gpupdate /force** and then press ENTER to immediately refresh the local policy.

**Configure event logging for a federation server**
Use the following procedure to specify the types of events that you want to be logged on a server that is running the Active Directory Federation Service:

- Configure event logging on a federation server

**Configure event logging for a federation server proxy**
Use the following procedure to specify the types of events that you want to be logged on a server that is running the ADFS Federation Service Proxy:

- Configure event logging on a federation server proxy

**Configure event logging on a Web server**
Events logged on Web servers that are running an ADFS Web Agent are configured according to the application type that the agent supports. Event logging is configured differently for Windows NT token-based applications and claims-aware applications:

- Windows NT token-based applications: On Web servers that are running the ADFS Web Agent for Windows NT token-based applications, event logging for these applications is set in the registry on the Web server. Use the following procedure to specify the types of events that you want to be logged for Windows NT token-based applications on the Web server:

  Configure event logging for a Windows NT token-based application

- Claims-aware applications: On Web servers that are running the ADFS Web Agent for claims-aware applications, event logging for these applications is set in the Web.config file for the application. Use the following procedure to specify the types of events that you want to be logged for claims-aware applications on the Web server:

  Configure event logging for a claims-aware application

**Configuration Tasks for Troubleshooting**
Before you use advanced troubleshooting techniques to identify and solve ADFS problems, configure your computer for troubleshooting.

To configure your computer for troubleshooting, perform the following tasks:

- Configure ADFS Event Logging
- Configure ADFS Debug Logging
- Disable JavaScript
- Enable ASP.NET Debug Output
- Configure an ASP.NET Error Page

**Configure ADFS Event Logging**
You can configure event logging on federation servers, federation server proxies, and Web servers. ADFS events are logged in the Application event log and the Security event log.

| **Important** |
| --- |
| You must turn on audit object access at each of the federation servers, for ADFS-related audits to appear in the Security log. This will allow the Federation Service to log either success or failure errors. For more information about how to turn on audit object access, see Audit object access (http://go.microsoft.com/fwlink/?LinkId=62686). |

**Configure ADFS Debug Logging**
Event logs are generally descriptive, intended to help you understand what is happening. However, the default events do not always provide the level of detail that is needed for effective troubleshooting. In this case, configure ADFS debug logging.

ADFS provides several levels of debug information that are available for troubleshooting ADFS problems. Use the procedures in this section for enabling debug logging and setting appropriate debug logging levels on federation servers, federation proxy servers, and Web servers that are running ADFS Web Agents.

| **Note** |
| --- |
| Debug logging consumes space and resources on the computer. Enabling debug logging is recommended only if you are troubleshooting a problem and need more information than is provided in events. Otherwise, do not enable debug logging. |

The debug log file is located in %systemdrive%\ADFS\logs.

**Debug log filename format**
If debug logging is enabled on a federation server, the log filename in the C:\ADFS\logs directory has the following format:

adfs*yyyymmdd-hhmmss*.log

In the name of the file, the number following "adfs" represents the date of the log and the number following the dash (-) represents the beginning time of the log.

**Debug log tags**
Depending on the level of debug logging you enable, you will see the following tags in debug logs:

[INFO] - Displays information about events, such as redirects with protocol Uniform Resource Locators (URLs), token validations, or claim mappings.

[VERBOSE] - Displays information about events, such as sign-in requests, responses, token contents, Web method calls, and security identifier (SID) information.

[ERROR] - Displays events for significant problems in the debug log.

[WARNING] - Displays events, which are not necessarily significant but that may cause future problems.

[EVENTLOG] - Displays all ADFS events.

Although all information in the log file could be useful, you can look at the lines that are tagged [ERROR] and [WARNING] first to quickly assess the problem.

For example, the following section of a debug log file shows that certificate chain validation is failing.

```
----------------

2005-11-09T19:46:47 [INFO] Requesting token for https://adfsweb.treyresearch.net/ from FS using inbound token.

2005-11-09T19:46:47 [VERBOSE] Parse: Token NOT found in cache

2005-11-09T19:46:47 [VERBOSE] SAML: effectivetime = 11/09/2005 19:46:53

expirationtime = 11/09/2005 20:46:53

2005-11-09T19:46:50 [WARNING] VerifyCertChain: Cert chain did not verify - error code was 0x80092013

2005-11-09T19:46:50 [ERROR] KeyInfo processing failed because the trusted certificate does not have a a valid certificate chain. Thumbprint =
BAF02C45AF23389CC7FEC615615056021E107C3E

2005-11-09T19:46:50 [WARNING] Failing signature verification because the KeyInfo section failed to produce a key.

2005-11-09T19:46:50 [WARNING] SAML token signature was not valid: AssertionID = _cbe6e3ca-fb90-4a93-a789-b925856163d0

2005-11-09T19:46:50 [VERBOSE] Processing FS response: policy version is a9d515c1-6965-4aa7-a78e-3cfc77f0dd2a - 16

2005-11-09T19:46:50 [INFO] Token issuance request to FS failed: ValidationFailure

2005-11-09T19:53:14 [VERBOSE] Processing HTTP GET: https://adfsresource.treyresearch.net/adfs/ls/?
wa=wsignin1.0&wreply=https://ADFSWeb.TREYRESEARCH.NET/&wct=2005-11-09T19:53:13Z&wctx=https://adfsweb.treyresearch.net/default.aspx

2005-11-09T19:53:14 [VERBOSE] Received SignIn Request.

2005-11-09T19:53:14 [VERBOSE] HOMEREALM: Realm could not be determined.
```

```
2005-11-09T19:53:14 [INFO] Received signin request via query string.

2005-11-09T19:53:14 [VERBOSE] Sign In Request Dump

-------------------
```

As you can see in the log text, even the thumbprint of the certificate is provided.

You can run the following command against the .cer file of the certificate to get more information about the failure.

**certutil –v –urlfetch -verify** *CertFileName* **.cer**

#### Set ADFS debug levels on federation servers

On federation account, resource, and proxy servers, you can use the Windows UI to enable debug logging and set levels to increase the detail of feedback in the logs.

Perform the following procedure on an account or resource federation server or federation proxy server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To set ADFS debug levels on federation servers and federation proxy servers**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Federation Service** or **Federation Service Proxy**, and then click **Properties**.
3. On the **Troubleshooting** tab, select debug levels as appropriate, and then click **OK**.

| **Note** |
| --- |
| To see descriptions for each debug level, click **Help** on the **Troubleshooting** tab. |

#### Enable ADFS authentication package debug logging on ADFS account federation servers

The account federation server uses the ADFS authentication package (ifsAp.dll) for mapping client certificates.

Perform the following procedures on an ADFS account federation server.

**Administrative credentials**

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable debug logging for the ADFS authentication package on an account federation server**

1. Open Regedit.
2. Navigate to:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\lsa\WebSSO\Parameters**
3. Right-click **Parameters**, click **New**, and then click **DWORD Value**.
4. In the new value file name box, type the following, and then press Enter:

   **DebugLevel**
5. Double-click the new entry and then, in **Value data**, type the following, and then click **OK**:

   **FFFFFFFF**

#### Enable ADFS debug logging on Web servers

To enable debug logging on Web servers, you need to edit the registry on the servers that you are troubleshooting.

You can enable debug logging for the following components:

- The ADFS Web Agent running on ADFS Web servers has two components:
  - ADFS Token Authentication service (ifssvc.exe), which validates incoming tokens and cookies. Debug logging creates ifssvc.log.
  - ADFS Web Agent Internet Server Application Programming Interface (ISAPI) extension (ifsext.dll), which handles the protocols that are used by ADFS to authenticate requests; and the ADFS Web Agent ISAPI filter (ifsfilt.dll), which assists the extension and enables user name logging in the Internet Information Services (IIS) log files. Debug logging creates the ifsext_StsAppPool1.log and ifsfilt_StsAppPool1.log, respectively in the %systemdrive%\ADFS\Logs directory.
- In addition, the ADFS Web Agent authentication package (ifsAp.dll) is used by Windows NT token-based applications for generating tokens when Service-for-User (S4U) is not available. Debug logging creates ifsap.log.

You can enable debug logging for each of these components in the registry on ADFS Web servers.

**Administrative credentials**

To complete these procedures, you must be a member of the Administrators group on the local computer.

**To enable debug logging for the ADFS Token Authentication service**

1. Open Regedit.
2. Navigate to:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ifssvc\Parameters**
3. Right-click **Parameters**, click **New**, and then click **DWORD Value**.
4. In the new value file name box, type the following, and then press Enter:

**DebugPrintLevel**

5. Double-click the new entry and then, in **Value data**, type the following, and then click **OK**:

   **FFFFFFFF**

**To enable debug logging for the ADFS ISAPI extension and filter**

1. Open Regedit.
2. Navigate to:

   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\ADFS\WebServerAgent**

3. Right-click **WebServerAgent**, click **New**, and then click **DWORD Value**.
4. In the new value file name box, type the following, and then press Enter:

   **DebugPrintLevel**

5. Double-click the new entry and then, in **Value data**, type the following, and then click **OK**:

   **FFFFFFFF**

**To enable debug logging for the ADFS Web Agent authentication package (for Windows NT token-based applications)**

1. Open Regedit.
2. Navigate to:

   **HKEY_LOCAL_MACHINE_SYSTEM\CurrentControlSet\Lsa\WebSso\ Parameters**

3. Right-click **Parameters**, click **New**, and then click **DWORD Value**.
4. In the new value file name box, type the following, and then press Enter:

   **DebugLevel**

5. Double-click the new entry and then, in **Value data**, type the following, and then click **OK**:

   **FFFFFFFF**

**Disable JavaScript**
JavaScript is used to automatically redirect the client to various points, including posting tokens. When JavaScript is disabled, the automatic redirect is prevented and a submit button is displayed instead. This button allows the client to walk through each step more easily as part of troubleshooting the configuration.

### Administrative credentials

To complete this procedure, you must be a member of the Users group on the local computer.

**Disable JavaScript in Internet Explorer**

1. Open Internet Explorer.
2. On the **Tools** menu, click **Internet Options**.
3. On the **Security** tab, click **Custom Level**.
4. Scroll to the **Scripting** category.
5. Under **Active scripting**, click **Disable**, and then click **OK** twice.

**Enable ASP.NET Debug Output**
You can generate debug output for ASP.NET files by configuring the Web.config file on the computer.

| **Note** |
| --- |
| Generating debug output for ASP.Net files enable clients to view server state information that could be misused by malicious users. |

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To enable ASP.NET debug output**

1. In Notepad, open the Web.config file in *SYSTEMDRIVE*\ADFS\sts.
2. Search for **<system.web>**.
3. Add the **<compilation debug>** entry under **<system.web>**, as follows:

   **< compilation debug="true" />**

4. Save and close the Web.config file.

**Configure an ASP.NET Error Page**
You can configure an error page to display when ASP.NET errors prohibit the opening of a page. Use the Web.config file to configure this error page.

### Administrative credentials

To complete this procedure, you must be a member of the Administrators group on the local computer.

**To configure an ASP.NET error page**

1. In Notepad, open the Web.config file in *SYSTEMDRIVE*\ADFS\sts.
2. Search for **<system.web>**.

3. Add the following entries under **<system.web>**:

   **<customErrors mode="Off" defaultRedirect="Errors.aspx">**

   **<error statusCode="404" redirect="PageNotFound.aspx" />**

   **</customErrors>**

4. Save and close the Web.config file.

**Understanding Default Event Logs**
By default, some events are logged by ADFS components in Event Viewer. The following categories of default events provide information about successful operations that are logged in the event logs and for which you will see events in Event Viewer.

- Default Events for Token-based Applications on a Web Server
- Default Events for Claims-aware Applications on a Web Server
- Default Auditing Events for Token-based Applications on a Federation Server
- Default Auditing Events for Claims-aware Applications on a Federation Server

**Default Events for Token-based Applications on a Web Server**
The following types of events appear when an ADFS client successfully receives access to the requested application Web page. These events occur in an environment similar to the sample scenario presented in Verifying ADFS Computer Settings and Connectivity.

If a Web server that is running an ADFS Web Agent is able to retrieve trust information successfully from the Federation Service, the following event is generated in the Application log on the Web server:

```
Event Type:Information
Event Source:ADFS ISAPI Extension
Event Category:None
Event ID:101
Date:11/10/2005
Time:10:45:23 AM
User:N/A
Computer:ADFSWEB
Description:
The ADFS Web Agent for Windows NT token-based applications successfully retrieved trust information from the Federation
Service.
```

The following event appears in the Security log on the ADFS Web server:

```
Event Type:Success Audit
Event Source:Security
Event Category:Logon/Logoff
Event ID:540
Date:11/10/2005
Time:10:46:17 AM
User:urn:federation:adatum\adamcar@adatum.com
Computer:ADFSWEB
Description:
Successful Network Logon:
 User Name:adamcar@adatum.com
 Domain:urn:federation:adatum
 Logon ID:(0x0,0x25E31)
 Logon Type:3
 Logon Process:
 Authentication Package:IfsAp
 Workstation Name:-
 Logon GUID:-
 Caller User Name:ADFSWEB$
 Caller Domain:TREYRESEARCH0
 Caller Logon ID:(0x0,0x3E7)
 Caller Process ID: 1568
 Transited Services: -
 Source Network Address:-
 Source Port:-
```

In the preceding event text, Caller Process ID: 1568 is the process ifssvc.exe on the ADFS Web server.

**Default Events for Claims-aware Applications on a Web Server**
You will notice the following event if the ADFS Web server is able to retrieve ADFS trust information successfully from the Federation Service.

```
Event Type:Information
Event Source:ADFS
Event Category:None
Event ID:621
Date:11/10/2005
Time:4:09:26 PM
User:N/A
Computer:ADFSWEB
Description:
The ADFS Web Agent for claims-aware applications successfully retrieved trust information from the Federation Service.
GUID: d977fee6-175b-4532-bc24-5ac54d137d57
Version: 17
Federation Service Uniform Resource Locator (URL):
https://adfsresource.treyresearch.net/adfs/fs/federationserverservice.asmx
Federation Service Uniform Resource Identifier (URI): urn:federation:treyresearch
Federation Service Endpoint URL: https://adfsresource.treyresearch.net/adfs/ls/
Federation Service Domain Account: TREYRESEARCH0\ADFSRESOURCE$
```

You will also see the following event below in the Security log.

```
Event Type:Success Audit
Event Source:ADFS ASP.NET Module Auditor
Event Category:Object Access
Event ID:560
Date:11/10/2005
Time:4:10:11 PM
User:NT AUTHORITY\NETWORK SERVICE
Computer:ADFSWEB
Description:
The client presented a valid inbound token as evidence.
Token ID: _ad5a3694-860d-4063-95a3-3b0163fad3ca
Identity: adamcar@adatum.com
```

**Default Auditing Events for Token-based Applications on a Federation Server**

Federation servers in the account and the resource domains log similar events in the security log if object access auditing is turned on. These events provide details about claims, including the time that the claim was presented and the requested user.

**Default auditing events for a token-based application on an account federation server**

The following default auditing events are logged in the Security log on the account federation server:

- Event ID 500: Indicates that a token request was received by adfsaccount.
- Event ID 510: Contains details of the resource token that was issued by adfsaccount.
- Event ID 520: Contains details of the logon accelerator token that was issued by adfsaccount.
- Event ID 550: Contains the list of claims that were retrieved from the account store.

The following sample events are generated on the federation server named adfsaccount.adatum.com, in an ADFS deployment similar to that described in Verifying ADFS Computer Settings and Connectivity.

**Event ID 500:**

```
Event Type:Success Audit
Event Source:ADFS Federation Service Auditor
Event Category:Object Access
Event ID:500
Date:11/10/2005
Time:4:10:13 PM
User:NT AUTHORITY\SYSTEM
Computer:ADFSACCOUNT
Description:
Transaction ID: {c727f983-de98-4c88-947c-06d5f914659a}

A token request was received directly by the Federation Service. The request for target 'urn:federation:treyresearch' was
approved, and one or more tokens were issued.
Target URI: urn:federation:treyresearch

A resource token was issued. Depending on the audit policy, further details of this token may be written to a 510 event
with the same transaction ID.
Token ID: _952de51d-e827-41db-8332-1afa84aa51d2
Identity: adamcar@adatum.com

A logon accelerator was issued. Depending on the audit policy, further details of this token may be written to a 520 event
with the same transaction ID.
Token ID: _b17a22ce-96af-4a59-a885-c1252872c6ea
Identity:

The client did not present a logon accelerator token as evidence.

The client presented valid credentials. Depending on the audit policy, the list of generated claims may be written to a 550
event with the same transaction ID.
Authentication method: Windows integrated authentication
Username: ADATUM0\adamcar
```

**Event ID 510:**

```
Event Type:Success Audit
Event Source:ADFS Federation Service Auditor
Event Category:Object Access
Event ID:510
Date:11/10/2005
Time:4:10:13 PM
User:NT AUTHORITY\SYSTEM
Computer:ADFSACCOUNT
Description:
Transaction ID: {c727f983-de98-4c88-947c-06d5f914659a}

This event contains the details of the resource token that was issued as part of the referenced transaction.

Token ID: _952de51d-e827-41db-8332-1afa84aa51d2
Issuer: urn:federation:adatum
Audience: urn:federation:treyresearch
Effective time: 11/10/2005 4:10:13 PM
Expiration time: 11/10/2005 5:10:13 PM
Claim source:
Authentication methods:
MethodTime
urn:federation:authentication:windows2005-11-11T00:10:13Z
UPN: adamcar@adatum.com
Email: [Claim not present]
Common name: [Claim not present]
Groups: (0 sensitive values omitted)
```

```
SharePointMapping
Custom claims:
NameValue
[Custom claims not present]
```

**Event ID 520:**

```
Event Type:Success Audit
Event Source:ADFS Federation Service Auditor
Event Category:Object Access
Event ID:520
Date:11/10/2005
Time:4:10:13 PM
User:NT AUTHORITY\SYSTEM
Computer:ADFSACCOUNT
Description:
Transaction ID: {c727f983-de98-4c88-947c-06d5f914659a}

This event contains the details of the logon accelerator token that was issued as part of the referenced transaction.

Token ID: _b17a22ce-96af-4a59-a885-c1252872c6ea
Issuer: urn:federation:adatum
Audience: urn:federation:adatum
Effective time: 11/10/2005 4:10:13 PM
Expiration time: 11/11/2005 2:10:13 AM
Claim source: urn:federation:activedirectory
Authentication methods:
MethodTime
urn:federation:authentication:windows2005-11-11T00:10:13Z
UPN: adamcar@adatum.com
Email: [Claim not present]
Common name: [Claim not present]
Groups: (0 sensitive values omitted)
Trey SharePoint Claim
Custom claims:
NameValue
[Custom claims not present]
```

**Event ID 550:**

```
Event Type:Success Audit
Event Source:ADFS Federation Service Auditor
Event Category:Object Access
Event ID:550
Date:11/10/2005
Time:4:10:13 PM
User:NT AUTHORITY\SYSTEM
Computer:ADFSACCOUNT
Description:
Transaction ID: {c727f983-de98-4c88-947c-06d5f914659a}

This event contains the list of claims that were retrieved from the account store as part of the referenced transaction.

UPN: adamcar@adatum.com
Email: [Claim not present]
Common name: [Claim not present]
Groups: (0 sensitive values omitted)
Trey SharePoint Claim
Custom claims:
NameValue
[Custom claims not present]
```

**Default auditing events for a token-based application on a resource federation server**
The following default auditing events are logged on the resource federation server:

- Event ID 500: Indicates that a token request was received by adfsresource.
- Event ID 510: Contains the details of the resource token that was issued by adfsresource.
- Event ID 520: Contains the details of the logon accelerator token that was issued by adfsresource.
- Event ID 540: Contains the details of the token that was presented by the client.

The following sample events are generated on the federation server named adfsresource.treyresearch.net, in an ADFS deployment similar to that described in Verifying ADFS Computer Settings and Connectivity.

**Event ID 500:**

```
Event Type:Success Audit
Event Source:ADFS Federation Service Auditor
Event Category:Object Access
Event ID:500
Date:11/10/2005
Time:4:10:12 PM
User:NT AUTHORITY\SYSTEM
Computer:ADFSRESOURCE
Description:
Transaction ID: {fb4e7da3-ca0b-43ca-a0c0-42a333ce0d80}

A token request was received directly by the Federation Service. The request for target
'https://adfsweb.treyresearch.net:8081/sampleapp/' was approved, and one or more tokens were issued.
Target URI: https://adfsweb.treyresearch.net:8081/sampleapp/

A resource token was issued. Depending on the audit policy, further details of this token may be written to a 510 event
with the same transaction ID.
Token ID: _ad5a3694-860d-4063-95a3-3b0163fad3ca
```

Identity: adamcar@adatum.com

A logon accelerator was issued. Depending on the audit policy, further details of this token may be written to a 520 event with the same transaction ID.
Token ID: _ea256664-ac66-4560-9a4f-97a26bf21fa0
Identity:

The client did not present a logon accelerator token as evidence.

The client presented a valid inbound token as evidence. Depending on the audit policy, further details of this token may be written to a 540 event with the same transaction ID.
Token issuer: urn:federation:adatum
Token ID: _952de51d-e827-41db-8332-1afa84aa51d2
Identity: adamcar@adatum.com

For more information, see Help and Support Center at http://go.microsoft.com/fwlink/events.asp.

**Event ID 510:**

Event Type:Success Audit
Event Source:ADFS Federation Service Auditor
Event Category:Object Access
Event ID:510
Date:11/10/2005
Time:4:10:12 PM
User:NT AUTHORITY\SYSTEM
Computer:ADFSRESOURCE
Description:
Transaction ID: {fb4e7da3-ca0b-43ca-a0c0-42a333ce0d80}

This event contains the details of the resource token that was issued as part of the referenced transaction.

Token ID: _ad5a3694-860d-4063-95a3-3b0163fad3ca
Issuer: urn:federation:treyresearch
Audience: https://adfsweb.treyresearch.net:8081/sampleapp/
Effective time: 11/10/2005 4:10:12 PM
Expiration time: 11/10/2005 5:10:12 PM
Claim source: urn:federation:adatum
Authentication methods:
MethodTime
urn:federation:authentication:windows2005-11-11T00:10:13Z
UPN: adamcar@adatum.com
Email: [Claim not present]
Common name: [Claim not present]
Groups: (0 sensitive values omitted)
[Groups not present]
Custom claims:
NameValue
[Custom claims not present]

For more information, see Help and Support Center at http://go.microsoft.com/fwlink/events.asp.

**Event ID 520:**

Event Type:Success Audit
Event Source:ADFS Federation Service Auditor
Event Category:Object Access
Event ID:520
Date:11/10/2005
Time:4:10:12 PM
User:NT AUTHORITY\SYSTEM
Computer:ADFSRESOURCE
Description:
Transaction ID: {fb4e7da3-ca0b-43ca-a0c0-42a333ce0d80}

This event contains the details of the logon accelerator token that was issued as part of the referenced transaction.

Token ID: _ea256664-ac66-4560-9a4f-97a26bf21fa0
Issuer: urn:federation:treyresearch
Audience: urn:federation:treyresearch
Effective time: 11/10/2005 4:10:12 PM
Expiration time: 11/11/2005 2:10:12 AM
Claim source: urn:federation:adatum
Authentication methods:
MethodTime
urn:federation:authentication:windows2005-11-11T00:10:13Z
UPN: adamcar@adatum.com
Email: [Claim not present]
Common name: [Claim not present]
Groups: (0 sensitive values omitted)
Adatum Share Point Claim
Custom claims:
NameValue
[Custom claims not present]

For more information, see Help and Support Center at http://go.microsoft.com/fwlink/events.asp.

**Event ID 540:**

Event Type:Success Audit
Event Source:ADFS Federation Service Auditor
Event Category:Object Access
Event ID:540
Date:11/10/2005
Time:4:10:12 PM
User:NT AUTHORITY\SYSTEM
Computer:ADFSRESOURCE
Description:

```
Transaction ID: {fb4e7da3-ca0b-43ca-a0c0-42a333ce0d80}

This event contains the details of the token that was presented by the client as part of the referenced transaction.

Token ID: _952de51d-e827-41db-8332-1afa84aa51d2
Issuer: urn:federation:adatum
Audience: urn:federation:treyresearch
Effective time: 11/10/2005 4:10:13 PM
Expiration time: 11/10/2005 5:10:13 PM
Claim source:
Authentication methods:
MethodTime
urn:federation:authentication:windows2005-11-11T00:10:13Z
UPN: adamcar@adatum.com
Email: [Claim not present]
Common name: [Claim not present]
Groups: (0 sensitive values omitted)
SharePointMapping
Custom claims:
NameValue
[Custom claims not present]
```

**Default Auditing Events for Claims-aware Applications on a Federation Server**

A claims-aware application generates events that are similar to the events that are generated by a token-based application.

# Troubleshooting ADFS Problems

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) scenarios rely on the interactions among several components, all of which must be configured correctly to achieve single sign-on (SSO) for ADFS clients. The components that make up an ADFS deployment include the following:

- Multiple ADFS components, including:
    - Federation Services
    - ADFS Web Agent
    - ADFS Proxy Server (optional)
- Internet Information Services (IIS) Web sites
- Active Directory or Active Directory Application Mode (ADAM) directory store
- Certificate services
- Domain Name System (DNS) infrastructure

Troubleshooting content that is included in this guide provides the information that you need to identify and solve problems that occur within and between components in an ADFS deployment. The identified problems and solutions apply to computers that are running the following software versions:

- IIS servers that are running ADFS servers, including account and resource federation servers and ADFS proxy servers: These IIS servers must be running Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Datacenter Edition.
- Domain controllers: Domain controllers must be running Windows Server 2003 R2, Windows Server 2003, or Windows 2000 Server.
- IIS servers running ADFS Web Agent: These IIS servers must be running Windows Server 2003 R2, Enterprise Edition; Windows Server 2003 R2, Datacenter Edition; or Windows Server 2003 R2, Standard Edition.

**In this section**


- A certificate has been issued to the wrong site

# A certificate has been issued to the wrong site

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

This problem typically occurs when you are attempting to connect to a Web site that is configured to use Active Directory Federation Services (ADFS) for single-sign-on (SSO) authentication and authorization. For this problem, the user might see error messages, as follows:

- If the resource Federation Service certificate name does not match the resource Federation Service host name, "The page cannot be displayed" error message appears.
- If either the account Federation Service certificate name or the Web server certificate name does not match the appropriate Federation Service host name, the following Internet Explorer warning appears: "The name on the security certificate is invalid or does not match the name of the site."

**Cause**

A Secure Sockets Layer (SSL) server authentication certificate Uniform Resource Locator (URL) does not match the appropriate URLs, as follows:

- The Federation Service endpoint URL that is configured in any of the following:
  - The account Federation Service
  - The resource Federation Service
  - The Web server
- The Federation Service URL in the ADFS Web Agent

**Solution**

Perform the following checks to determine where an incorrect URL exists.

- Verify the DNS Name in the Server Authentication Certificate
- Verify the Federation Service URL in the ADFS Web Agent

**Verify the DNS Name in the Server Authentication Certificate**

Use the following procedure to verify that the Domain Name System (DNS) names in the SSL server authentication certificates for account and resource federation servers match the DNS names in the Federation Service endpoint URL for the respective account and resource Federation Services.

Perform the following procedure on a federation server in both the account and resource Federation Services.

**To verify the DNS name in the server authentication certificate for a federation server**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. Right-click **Trust Policy**, and then click **Properties**.
3. Under **Federation Service endpoint URL**, note the DNS name in the first portion of the URL, and then click **Cancel**.
4. Open the server authentication certificate (*ServerName*.cer) for the federation server.
5. Make sure the DNS name in **Issued to** is the same as the DNS name in the Federation Service endpoint URL in step 3. If the names do not match, do one of the following:
   - If the **Federation Service endpoint URL** is incorrect, change the value.

- If the **Issued to** value in the certificate is incorrect, remove the current certificate, and then request and install a new SSL server authentication certificate using the correct DNS name.

Perform the following procedure on a Web server that is running an ADFS Web Agent.

**To verify the DNS name in the server authentication certificate for a Web server**

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services**.
2. Right-click **Web Site**, and then click **Properties**.
3. On the **Directory Security** tab, under **Secure Communications**, click **View Certificate**.
4. Make sure that the DNS name in **Issued to** is the same as the DNS name of the Web server.
5. If the names do not match, remove the current certificate, and then request and install a new SSL server authentication certificate using the correct DNS name.

For information about how to request and install SSL server authentication certificates, see Public Key Infrastructure for Windows Server 2003 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=19936).

**Verify the Federation Service URL in the ADFS Web Agent**
Use the following procedure to make sure that the DNS name in the Federation Service URL on the Web server that is running the ADFS Web Agent matches the DNS name in the Federation Service endpoint URL for the resource partner Federation Service.

Check the Federation Service endpoint URL value for the resource Federation Service that you noted in the preceding procedure against the Federation Service URL that is configured in the ADFS Web Agent on the Web server. The DNS names in these two URLs must match exactly.

**To verify the Federation Service URL in the ADFS Web Agent**

1. To check the ADFS Web Agent in IIS (for Windows NT token-based applications):
   a. On the Web server, click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
   b. Double-click *YourServerName* **(local computer)**, right-click **Web Sites**, and then click **Properties**.
   c. On the **ADFS Web Agent** tab, under **Federation Service URL**, verify that the URL path specifies **https://***DNSNameOfResourcePartner***/adfs/fs/FederationServerService.asmx**.
   d. Verify that the DNS name in the first portion of the URL matches the DNS name in **Federation Service endpoint URL** in step c.
2. To check the ADFS Web Agent in the Web.config file (for claims-aware applications):
   a. In Notepad or another text editor, open the Web.config file that is located in the application folder on the Web server.
   b. Search for **<fs>**.
   c. Verify that the URL path specifies **https://***DNSNameOfResourcePartner***/adfs/fs/FederationServerService.asmx**.
   d. Verify that the DNS name in the first portion of the URL matches the DNS name in **Federation Service endpoint URL** in step c.
3. If the path is incorrect or if the DNS name does not match the name in **Federation Service endpoint URL**, either change the Federation Service URL as needed in the ADFS Web Agent, or change the Federation Service endpoint URL in the resource Federation Service.

**See Also**

**Concepts**

Change the Federation Service endpoint URL
Set the Federation Service URL for a Windows NT token-based application
Set the Federation Service URL for a claims-aware application

# IIS 6.0 Operations Guide

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0 with the Microsoft® Windows Server 2003™ operating system provides integrated, reliable, scalable, secure, and manageable Web server capabilities over an intranet, the Internet, or an extranet. IIS is a tool for creating a strong communications platform of dynamic network applications. Organizations of all sizes use IIS to host and manage Web pages on the Internet or on their intranet, to host and manage FTP sites, and to route news or mail using the Network News Transfer Protocol (NNTP) and the Simple Mail Transfer Protocol (SMTP). IIS 6.0 leverages the latest Web standards like Microsoft ASP.NET, XML, and Simple Object Access Protocol (SOAP) for the development, implementation, and management of Web applications. IIS 6.0 includes new features designed to help organizations, IT professionals, and server administrators achieve their goals of performance, reliability, scalability, and security for potentially thousands of Web sites either on a single IIS server or on multiple servers.

| **Important** |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down — which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Enabling and Disabling Dynamic Content in IIS 6.0 node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0. |

This version of IIS was designed to provide the following benefits and features.

| IIS 6.0 Feature | Description |
| --- | --- |
| Reliability | IIS 6.0 uses a new request-processing architecture and application isolation environment that enables individual Web applications to function within a self-contained worker process. This environment prevents one application or Web site from stopping another, and reduces the amount of time administrators spend restarting services to correct problems related to applications. The new environment also includes proactive Application Pool Health. |
| Scalability | IIS 6.0 introduces a new kernel-mode driver for HTTP parsing and caching, specifically tuned to increase Web server throughput and scalability of multiprocessor computers, thereby significantly increasing the following:<br>• The number of sites a single IIS 6.0 server can host<br>• The number of concurrently-active worker processes<br>Also, by Configuring Startup and Shutdown Time Limits, IIS allocates resources to active sites, as opposed to wasting resources on idle requests. |

| | |
|---|---|
| Security | IIS 6.0 provides significantly improved security over earlier versions of IIS. To reduce the attack surface of systems, IIS is not installed by default on the operating systems in the Windows Server 2003 family. Administrators must explicitly select and install IIS. IIS installs by default in a locked-down state, capable of serving only static content. Using the Web Service Extensions node, Web site administrators can Configuring IIS For Dynamic Content IIS functionality based on the individual needs of their organization. IIS 6.0 includes a variety of Security in IIS 6.0 features and technologies to help ensure the integrity of your Web and FTP site content, as well as the data transmitted through your sites. IIS security features include the following security-related tasks:<br>• Authentication in IIS 6.0<br>• Access Control with IIS 6.0<br>• IIS 6.0 Encryption<br>• Certificates_IIS_SP1_Ops<br>• Auditing in IIS 6.0 |
| Manageability | To meet the needs of a diverse set of organizations, IIS provides a variety of manageability and administration tools. Administrators can configure an IIS 6.0 server using IIS Manager, Using Command-Line Administration Scripts, or by directly Enabling Edit-While-Running in IIS 6.0. Administrators can also Administering Servers Remotely in IIS 6.0 IIS servers and sites. |
| Enhanced Development | Compared to earlier operating systems, the Windows Server 2003 family offers an improved developer experience with About ASP.NET and IIS integration. ASP.NET recognizes most ASP code while providing greater functionality for building enterprise-class Web applications that can work as a part of the Microsoft .NET Framework. Using ASP.NET allows you to take full advantage of the features of the common language runtime, such as type safety, inheritance, language interoperability, and versioning. IIS 6.0 also offers support for the latest Web standards, including XML, SOAP, and Internet Protocol Version 6 Features. |
| Application Compatibility | IIS 6.0 is compatible with most existing applications, based on feedback from thousands of customers and independent software vendors (ISVs). Also, to ensure maximum compatibility, IIS 6.0 can be configured to run in IIS 5.0 Isolation Mode in IIS 6.0. |

The IIS 6.0 Operations Guide includes the following information:

- Getting Started with IIS 6.0: Describes key features in IIS 6.0, outlines the differences between this and earlier versions of IIS, and introduces basic concepts and tools, such as how to install IIS and how to open IIS Manager.
- Common Administrative Tasks in IIS 6.0: Includes a series of short procedural topics describing how to complete routine or common IIS 6.0 tasks.
- Site Setup: Describes how to set up a new Web, FTP, SMTP, or NNTP site.
- Server Administration Guide: Includes conceptual and procedural information regarding administrative tasks performed on IIS servers and sites.
- Security in IIS 6.0: Describes how to configure your Web server to strengthen the security of your Web sites and to carry out other essential security functions.
- Performance Tuning: Describes how to maximize the performance of your IIS 6.0 solution.
- Logging Site Activity: Describes how to use operating system and IIS logs to collect information about server events, server performance, and user activity.

- Error Messages in IIS 6.0: Describes how to configure IIS to send default and custom error messages.
- Web Application Guide: Describes how to develop and manage Web applications by using ASP and ASP.NET.
- IIS 6.0 Administration Reference: Describes interfaces, objects, and data structures that are available for programming and administering IIS, as well as metabase properties.
- Troubleshooting in IIS 6.0: Describes processes for resolving issues with your IIS 6.0 installation.

# Getting Started with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section includes general information about key features and changes in this version of Internet Information Services (IIS). It also includes information about how to install IIS and access IIS Manager to administer IIS sites, as well as links to resources on the Web.

| **Important** |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Enabling and Disabling Dynamic Content in IIS 6.0 node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0. |

**In This Section**

This section includes the following information:

- IIS 6.0 Features: Describes key features in this version of IIS.
- What's Changed: Describes important differences between IIS versions 4.0, 5.0, 5.1 and 6.0.
- IIS Manager: Describes how to access the IIS graphical user interface.
- Frequently Asked Questions about IIS 6.0: Answers common questions about IIS.
- Additional IIS 6.0 Resources: Provides links to information on the Web.
- Release Notes_IIS_SP1_Ops: Explains how to access late-breaking technical information, known problems, and information about how to get technical support for IIS.
- Before Configuring IIS: Discusses server configurations that should be addressed before attempting to configure IIS.
- Installing IIS: Describes various methods for installing IIS and how to add or remove components.
- Uninstalling IIS: Describes two methods for removing IIS from your server.

# IIS 6.0 Features

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS 6.0 includes new features designed to help businesses, IT professionals, and Web administrators achieve their goals of performance, reliability, scalability, and security for potentially thousands of Web sites either on a single IIS server or on multiple servers. This section describes new and notable features in IIS 6.0.

**In This Section**

- Reliability Features: Describes new modes of operation and World Wide Web Publishing Service (WWW service) processes.
- Security Features_IIS_SP1_Ops: Describes new features designed to secure your Web sites, data, and your IIS server.
- Performance Features: Describes performance and scalability features.
- Web Application Technologies: Describes the IIS 6.0 implementation of Microsoft® ASP.NET, ASP, and the ADSI provider.
- Administrative Tools and Features: Describes metabase changes and new administrative tools.
- Internet Standards: Describes support for Web and Internet standard protocols.
- Internet Protocol Version 6 Features: Describes the IIS 6.0 functionality that supports Internet Protocol Version 6 (IPv6).

# Reliability Features

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS 6.0 is more reliable than previous versions due to a new request processing architecture that provides an application isolation environment which enables individual Web applications to function within their own, self-contained worker process. The following are just a few of the improvements that make IIS 6.0 a more reliable Web administration environment.

**Modes of Operation — New**
You can configure IIS 6.0 to run in either worker process isolation mode, which runs all processes in an isolated environment, or IIS 5.0 isolation mode, in which you can run Web applications that are not compatible with worker process isolation mode. In worker process isolation mode, you can isolate anything — from an individual Web application to multiple sites — in its own self-contained World Wide Web Publishing Service (WWW service) worker process, preventing one application or site from stopping another. Separating applications or sites into separate processes simplifies a number of management tasks.

**Worker Processes — New**
In worker process isolation mode, worker processes handle all user code and are completely isolated from the core WWW service, Inetinfo.exe. Because these ISAPI applications run separately from the WWW service, an application failure prevents all services hosted by the WWW service from failing. Only the worker process that hosts the ISAPI application is affected. Worker processes can be configured to run on specific CPUs, which allows you greater control of balancing system resources.

**Multiple Application Pools — New**
Worker process isolation mode allows customers to create multiple application pools, where each application pool can have a unique configuration. Performance and reliability are enhanced because these application pools receive their requests directly from the kernel instead of from the WWW service.

**Application Pool Assignments — New**
An application pool can be configured in worker process isolation mode to serve anything from one Web application, to multiple applications, to multiple sites. Assigning an application to an application pool allows you to further isolate applications, and is as easy as configuring what pool that application should be routed to in the metabase. Sites, by default, are considered to be a simple application, where the root namespace "/" is configured as an application.

**Application Pool Queue Length Limits — New**
When running in worker process isolation mode, you can choose to limit the number of requests that should be queued for each application pool in the HTTP protocol stack (HTTP.sys). This allows you to fine-tune the environment in which your applications run by managing application pool requests and system resources.

**Web Garden Configuration on Multi-Processor Computers — New**
An application pool is associated with a specified worker process. For improved reliability, you can configure a Web garden, which is an application pool that has multiple worker processes serving the requests routed to that pool. If one worker process fails, another worker process can continue to process requests. On a computer with multiple processors, you can obtain even greater reliability by establishing affinity between worker processes and specific processors. Using processor affinity, a lock on one process does not block other processes running on the multiprocessor computer.

**Kernel-Mode Queuing — New**
The IIS 6.0 kernel-mode driver, HTTP.sys, places requests destined for application pools in a queue if a worker process in that pool fails. Because the WWW service manages the processes, the WWW service can initialize another process without causing an interruption in service to users. Requests are sent from the queue until the queue is empty or until the WWW service is stopped.

**Health Monitoring — New**
In worker process isolation mode, the WWW service monitors the health of worker processes and, if the process is unresponsive, terminates the process and replaces it with another. Health monitoring helps keep processes running and provides users with more reliable service.

**Idle Process Timeout — New**
You can control system resources by configuring application pools in worker process isolation mode to have their worker processes request a shutdown if the worker processes are idle for a configurable amount of time, such as minutes, hours, or days. You can then request a forced application startup when IIS receives a request for that application pool.

**Rapid-Fail Protection — New**

You can configure IIS in worker process isolation mode to use rapid-fail protection to automatically disable an application pool when a particular application pool suffers consecutive multiple failures.

**Orphan Worker Process — New**

A worker process that fails to respond to queries from the WWW Service Administration and Monitoring component can be configured in worker process isolation mode to "separate" from the application pool instead of being terminated, as a new worker process is started in its place. Because the process is still running, you can perform diagnostics on it. For more information, see Application Pool Health.

**Recycling Worker Processes — New**

You can configure IIS, in worker process isolation mode, to periodically restart worker processes in an application pool. When a new process is initialized, it takes over any requests to that application pool while the former worker process finishes processing its requests from the queue. Recycling worker processes allows you to manage troublesome applications, such as those with memory leaks, without interrupting service to your users.

# Security Features_IIS_SP1_Ops

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

IIS 6.0 includes a variety of security features and technologies to ensure the integrity of your Web and FTP site content, as well as data transmitted through your sites. IIS security features cover the following security-related tasks: authentication, access control, encryption, certificates, and auditing. This topic describes new and notable security features.

| **Important** |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0. |

**Digest Authentication**

Digest authentication allows robust authentication of users across proxy servers and firewalls. In addition, Anonymous authentication, Basic authentication, and Integrated Windows authentication are still available.

**Advanced Digest Authentication — New**

Advanced Digest authentication makes improvements over Basic authentication in that credentials are sent over the network as an MD5 hash and are stored as such in the Active Directory® of the domain controller. This mechanism makes it extremely difficult for intruders to discover users' passwords and does not require you to modify your applications.

**Secured Communications**

Secure Sockets Layer (SSL) 3.0 and Transport Layer Security (TLS) provide a secure way to exchange information between clients and servers. In addition, SSL 3.0 and TLS provide a way for the server to verify who the client is before the user logs on to the server. In IIS, client certificates are exposed to both ISAPI and Active Server Pages (ASP), so that programmers can track users through their sites. Also, IIS can map the client certificate to a Windows user account so that administrators can control access to system resources based on the client certificate.

**Server-Gated Cryptography**

Server-Gated Cryptography (SGC) is an extension of SSL that allows financial institutions with export versions of IIS to use strong 128-bit encryption. Although SGC capabilities are built into IIS, a special SGC certificate is required to use SGC. For more information about SGC, see Server-Gated Cryptography.

**Selectable Cryptographic Service Provider — New**

Secure Sockets Layer provides a secure way to exchange information between clients and servers. However, the CPU has to perform intensive cryptography, which degrades performance. IIS offers the Selectable Cryptographic Service Provider (CSP), which allows you to select a cryptographic provider that suits your needs. Each provider can create a public and private key for encrypting data sent to and from the Web server. The private key is stored at the server on hardware, a PCI card, a SmartCard, or in the registry — as it is for the two default providers that Microsoft installs. Storing the private key on hardware allows you to plug into hardware-based accelerator cards, that perform cryptographic computations, instead of the server. It is easy to select providers from IIS Manager to use Microsoft CryptoAPI providers or installed CryptoAPI providers from other companies. All CryptoAPIs implement the same methods so that you can switch between providers without having to change your code.

**Configurable Worker Process Identity — New**

To thwart system attackers and malicious users, you can configure application pools and therefore the worker process executing within, to run under an account with a lower level of permissions than LocalSystem. If you provide services to Internet users, you can allow your customers to upload static content and executable code. Erroneous code will not cause the World Wide Web Publishing Service (WWW service) or computer to fail. Only the application will fail.

**Security Wizards**

The following security wizards simplify server administration tasks:

- The **Web Server Certificate Wizard** simplifies certificate administration tasks, such as creating certificate requests and managing the certificate life cycle.

- The **CTL Wizard** helps you configure your certificate trust lists (CTLs). A CTL is a list of trusted certification authorities (CAs) for a particular directory. CTLs are especially useful for Internet service providers (ISPs) who have several Web sites on their server and need to have a different list of approved certification authorities for each site.

**IP and Internet Domain Restrictions**
You can assign or deny Web access to individual computers, groups of computers, or entire domains.

**Kerberos V5 Authentication Protocol Compliance**
IIS is fully integrated with the Kerberos V5 authentication protocol implemented in members of the Windows Server 2003 family, allowing you to pass authentication credentials among connected computers running Windows.

**Certificate Storage**
IIS certificate storage is now integrated with the Windows CryptoAPI storage. The Windows Certificate Manager provides a single point of entry that allows you to store, back up, and configure server certificates.

# Performance Features

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A new generation of applications puts a greater demand on performance and scalability attributes of Web servers. Increasing the speed at which HTTP requests can be processed and allowing more applications and sites to run on one server translates directly into fewer servers needed to host a site. The following are a few of the performance improvements included in IIS 6.0.

**HTTP.sys — New**
The Hyper Text Transfer Protocol (HTTP) stack in members of the Windows Server 2003 family is implemented as a kernel-mode device driver called HTTP.sys. HTTP.sys is a new feature specific to the Windows Server 2003 family, and is a part of the networking subsystem. IIS is a user-mode process that runs in a completely different execution mode than kernel processing. HTTP.sys is a single point of contact for all incoming (server-side) HTTP requests that provides high performance connectivity for HTTP server applications. HTTP.sys is also responsible for overall connection management, bandwidth throttling, and text-based logging. HTTP.sys implements a flexible Universal Resource Identifier (URI) response cache, called flexible caching, which enables HTTP applications to cache static data and service cached HTTP requests completely in kernel mode with no transition to user mode. HTTP.sys implements a URI namespace mapping mechanism called application pools. With application pools, HTTP applications can claim portions of the URI namespace, which allows HTTP.sys to route HTTP requests directly to the appropriate HTTP applications. As a result of these changes, IIS has increased throughput and scalability of multiprocessor computers, which significantly increases:

- The number of sites a single IIS 6.0 server can host.
- The number of concurrently-active worker processes.

**ASP Disk Caching — New**
Active Server Pages (ASP) processes templates that contain ASP scripts, stores the processed templates in a cache, and serves the cached templates to clients. By default, 250 templates are cached in-memory. If a site uses ASP heavily, this in-memory template cache cannot store all the necessary templates. IIS includes a persistent cache, so templates are cached to disk if the in-memory cache does not have enough space. If the ASP page is requested again, Asp.dll doesn't need to recompile the source code and can load the compiled template from disk. Caching ASP templates enhances performance because cached ASP templates are not processed each time they are called. You can further improve performance by changing the number of cached templates for all applications with low (in process) application protection, medium (pooled) application protection, or individually for applications with high (isolated) application protection.

For more information about ASP caching, see Configuring ASP Applications.

**Asynchronous CGI — New**
IIS handles the Common Gateway Interface (CGI) implementation asynchronously, thereby removing the potential for blocked CGI processing responses.

**Capacity Planning Tracing — New**
The IIS tracing feature for capacity planning traces application workloads against ASP script resource usage, such as HTTP response/request times. You can collect and analyze event trace data under different workloads and different hardware configurations and then use this data to help you with capacity planning.

**Centralized Binary Logging — New**
Centralized Binary Logging allows multiple sites to write to a single log file in a binary, unformatted manner. This logging format improves performance, as the data does not have to be formatted in any specific manner like W3C Extended, IIS, or NCSA formats. With Microsoft SQL Server-TM- support, multiple machines can deposit logs in one centralized database.

**Quality of Service —New**
Quality of Service (QoS) encompasses the set of methods or processes that a service-based organization implements to maintain a specific level of quality. In the context of Windows, QoS is a set of service requirements that the network must meet to assure an adequate service level for data transmission. The goal of QoS is to ensure that particular sites or applications do not monopolize server resources, like memory or CPU cycles. As an administrator, you can control the level and quality of service for your users by using the IIS implementation of connection limits and timeouts, application pool queue length limits, bandwidth throttling, and process accounting.

**Bandwidth Throttling**

If the network or Internet connection used by your Web server is also used by other services such as e-mail or news, you may want to limit the bandwidth used by your Web server so that bandwidth is available for those other services. If your Web server hosts more than one Web site, you can individually throttle the bandwidth used by each site.

**CPU Monitoring**

CPU Monitoring is a tool that monitors and automatically shuts down worker processes that consume large amounts of CPU resources. CPU monitoring is enabled for individual application pools. CPU monitoring can also help you determine which sites have malfunctioning scripts or CGI processes.

**Large Memory Support for x86.0**

IIS 6.0 (HTTP.sys) can be configured to cache up to 64 GB on an x86 system if booted with /PAE and without /3GB. This is a new feature as IIS 5.0 cached in virtual memory which was capped at 3GB.

# Web Application Technologies

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The Windows Server 2003 family offers an improved developer experience with ASP.NET and IIS integration. Microsoft® ASP.NET recognizes most ASP code while providing greater functionality for building enterprise-class Web applications that can work as a part of the Microsoft .NET Framework. Using ASP.NET allows you to take full advantage of the features of the common language runtime, such as type safety, inheritance, language interoperability, and versioning. IIS 6.0 also offers support for the latest Web standards, including XML, Simple Object Access Protocol (SOAP) and Internet Protocol Version 6 (IPv6.0).

**ASP.NET -- New**
ASP.NET is a unified Web application platform that provides the services necessary to build and deploy enterprise-class Web applications and XML Web services. ASP.NET is available and is installed, by default, on the operating systems in the Windows Server 2003 family.

| **Important** |
| --- |
| This feature is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. For more information, see "Features unavailable on 64-bit versions of the Windows Server 2003 family" in Help and Support Center for Windows Server 2003. |

ASP.NET offers a new programming model and infrastructure for more secure, scalable, and stable applications that can target any browser or device. ASP.NET is part of Microsoft .NET Framework, a computing environment that simplifies application development in the highly distributed environment of the Internet.

.NET Framework includes the common language runtime, which provides core services such as memory management, thread management, code security, and the .NET Framework class library, a comprehensive and object-oriented collection of types that developers can use to create applications. Note that ASP.NET is not supported by 64-bit systems.

**Active Server Pages**
You can create dynamic content by using server-side scripting and components to create browser-independent dynamic content. Microsoft Active Server Pages (ASP) provides an easy-to-use alternative to Common Gateway Interface (CGI) and ISAPI by allowing content developers to embed any scripting language or server component into their HTML pages. ASP provides access to all of the HTTP request and response streams, as well as standards-based database connectivity and the ability to customize content for different browsers.

**ADSI 2.0**
ADSI 2.0 provides administrators and application developers the ability to add custom objects, properties, and methods to the existing ADSI provider, giving administrators even more flexibility in configuring their sites.

**Wildcard Application Mappings**
Wildcard application mappings provide administrators the ability to configure a Web site or virtual directory to run an ISAPI application at the beginning of every request to that Web site or virtual directory, regardless of the extension of the requested file.

**Related Information**

- For more information about ASP.NET, see About ASP.NET and Creating ASP.NET Web Applications.
- For more information about Active Server Pages, see About ASP.
- For more information about ADSI 2.0, see IIS Administration Technologies in the IIS Software Development Kit (SDK) on MSDN.
- For more information about wildcard application mappings, see Installing Wildcard Application Mappings.

# Administrative Tools and Features

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To meet the needs of a diverse set of customers, IIS provides a variety of manageability and administrative tools. Administrators can configure a server running IIS 6.0 with IIS Manager, administration scripts, or by directly editing the IIS plaintext configuration file. Administrators can also remotely administer IIS servers and sites.

**Plaintext Metabase — New**
The metabase is a repository for most IIS configuration values. Metabase reengineering has lead to dramatic improvements in server startup and shutdown times, in addition to enhancing overall metabase performance and usability. This version of IIS includes an XML metabase configuration file, which can be edited manually or programmatically. The metabase is stored in two plaintext XML files:

- MetaBase.xml contains the configuration values for IIS.
- MBSchema.xml stores the XML metabase schema and enforces correct metabase configuration.

Because both the MetaBase.xml and MBSchema.xml files are stored as plain text, they can be read using any plaintext editor.

**Storing Metabase History — New**
You can quickly revert to a previous version of the metabase, thereby preventing long service interruptions. Automatic tracking of changes to the metabase are written to disk by copying each version of the metabase file into a history folder. Each history file is identified with a unique version number, which is then available for the metabase restore processes. A new set of Admin Base Object (ABO), WMI, and ADSI methods are exposed, which allows users to enumerate and restore the metabase configuration programmatically.

**Edit Metabase File While IIS Is Running — New**
You can edit the metabase file directly while IIS is running and the changes are immediately applied to the in-memory metabase. This capability, known as *edit-while-running*, is possible because MetaBase.xml is a text file. Direct editing can be accomplished using common text tools, like Microsoft Notepad, when the edit-while-running feature is enabled.

**Import and Export Site and Application Configuration — New**
In order to propagate site and application configuration settings across multiple servers for any node level, IIS supports two Admin Base Object (ABO) methods, Import() and Export(). Using these methods, you can export one node, or an entire tree, to an XML file from any level of the metabase, along with the inherited configuration, and then import one node, or an entire tree, from an XML file, as well as the inherited configurations. In addition you can password-protect secure data with these ABO methods.

**Create Server-Independent Metabase Backups — New**
Using the BackupWithPassword() and RestoreWithPassword() methods of the ADSI IIsComputer object, you can programmatically back up or restore the metabase with a password. The session key is encrypted with an optional user-supplied password during backup and is not based on the machine key.

**Metabase Reliability — New**
Using the configuration backup/restore feature, you can create a secure or unsecured backup of your metabase. Administrators can also restore copies of the metabase to other computers through IIS Manager. These backup methods provide a way to restore only your metabase settings, not your content files.

**IIS WMI Provider — New**
The Windows Server 2003 family of products features Windows Management Instrumentation (WMI), a scalable management infrastructure, as a means to improve management of Windows servers in your environment. The IIS WMI provider provides a level of manageability that is functionally equivalent to the IIS ADSI provider, while supporting an extensible schema. Furthermore, schema extensions developed using ADSI are imported into the IIS WMI provider, making the benefits of WMI even greater. For more information about WMI, see "IIS Administration Technologies" in the IIS Software Development Kit (SDK) on MSDN.

**FTP User Isolation — New**
FTP user isolation is a solution for Internet service providers (ISP) and application service providers that need to offer their customers individual FTP directories to upload files and Web content. FTP user isolation prevents users from viewing or

overwriting other users' Web content by restricting users to their own directory. A user cannot navigate higher up the directory tree because the user's top-level directory appears as the root of the FTP service. Within the user's specific site, the user still has the ability to create, modify, or delete files and folders.

**FrontPage 2002 Server Extensions — New**
FrontPage Server Extensions provide Web-based and command-line administration for extending virtual servers. Extending virtual servers with FrontPage Server Extensions enables the site owner to author the site in FrontPage and delegate site ownership and administration credentials.

FrontPage® 2002 Server Extensions from Microsoft is an update of FrontPage 2000 Server Extensions, and includes new security features, such as roles and rights, and new features for monitoring server health and Web site usage.

**Command-Line Administration Scripts — New**
Command-line administration makes it easier to perform management tasks efficiently. IIS provides scripts for the following tasks:

- Create, delete, start, stop, and list Web sites
- Create, delete, start, stop, and list FTP sites
- Create or delete Web virtual directories
- Create or delete applications
- Export or import IIS configuration
- Back up or restore IIS configuration

**UTF-8 Logging Support — New**
IIS supports log file recording in UTF-8 that allows administrators to log their URL hits in their native language instead of English. This setting, configurable on the Web-service level, instructs the HTTP protocol stack (HTTP.sys) in which format to write the log files: UTF-8 or the local code page.

**Remote Administration**
IIS includes a variety of tools to help you remotely administer your IIS Web servers. These tools include the Remote Administration (HTML) tool, Terminal Services, and programmatic features like ADSI, WMI, and a new set of command-line administration scripts.

The IIS Remote Administration (HTML) tool allows you to remotely administer IIS across the Internet or your intranet through your Web browser.

Terminal Services is a feature of the Windows Server 2003 family of products that allows you to run 32-bit Windows applications on terminals and terminal emulators running on personal computers and other computer desktops. This enables you to remotely administer Windows services, such as IIS, as if you were at the server console. You can administrate from older legacy PCs, or even from devices such as UNIX workstations with compatible client software. However, non-Windows-based client devices require third-party add-on software.

**Centralized Administration**
Administration tools for IIS use Microsoft Management Console (MMC). MMC hosts programs, such as IIS Manager, that administrators can use to manage their servers. You can use IIS Manager from a computer running a member of the Windows Server 2003 family to administer a computer on your intranet running IIS on a Windows server product.

**Remotable Certification Object —New**
Using scripts, you can process Secure Sockets Layer (SSL) certificates on your local or a remote computer. A Component Object Model (COM) object allows you to perform the add, remove, back up, and restore functions, which circumnavigates the constraints imposed by MMC.

For more information, see Managing Server Certificates Programmatically in IIS 6.0.

# Internet Standards

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS 6.0 offers support for the latest Web standards, including HTTP 1.1, TCP/IP, FTP, SMTP, NNTP, and PICS ratings.

**Standards Based**
IIS 6.0 complies with the HTTP 1.1 standard, which includes features such as PUT and DELETE, the ability to customize HTTP error messages, and support for custom HTTP headers.

**Multiple Sites, One IP Address**
With support for host headers, you can host multiple web sites on a single server running a member of the Windows Server 2003 family with only one IP address. This is useful for Internet service providers and corporate intranets hosting multiple sites.

**IIS Extensions to FTP — New**
IIS extends FTP in two significant areas. The IIS FTP User Isolation feature allows you to contain users to their own FTP directory, thus preventing them from viewing or overwriting other users' Web content. In addition, IIS supports multiple character sets for FTP, which allows your users access to their native-language worldwide.

**UTF-8 Support — New**
Support for Unicode and UTF-8 now extends to file names and URLs. Active Server Pages (ASP) can open any file name using the Unicode file name string. UTF-8 URLs are converted to a Unicode representation and presented to ASP.

**Web Distributed Authoring and Versioning**
Web Distributed Authoring and Versioning (WebDAV) enables remote authors to create, move, or delete files, file properties, directories, and directory properties on your server over an HTTP connection.

**News and Mail**
You can use News Network Transport Protocol (NNTP) and Simple Mail Transfer Protocol (SMTP) services to set up intranet news and mail services that work in conjunction with IIS.

**PICS Ratings**
You can apply Platform for Internet Content Selection (PICS) ratings to sites that contain content for mature audiences.

**FTP Restart**
FTP Restart enables File Transfer Protocol (FTP) file downloads to be resumed without having to download the entire file again if an interruption occurs during data transfer.

**HTTP Compression**
HTTP Compression provides faster transmission of pages between the Web server and compression-enabled clients. It compresses and caches static files, and performs on-demand compression of dynamically generated files.

# Internet Protocol Version 6 Features

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This version of IIS provides Internet services to clients connecting over the next generation of Internet Protocol (IP) known as IP version 6, or IPv6. IPv6 is included with the Microsoft Windows XP networking platform and the Microsoft Windows Server 2003 family.

The Internet Server API (ISAPI) framework provides the appropriate local- and remote-host server variables for IPv6 network addresses: LOCAL_ADDR and REMOTE_ADDR. When clients connect over IPv6, these variables store the IPv6 address.

When writing applications that use the ISAPI server variables for network addressing over IPv6, remember that all buffers allocated for network address storage can be a maximum of 128 bits in size (56 character string), whereas buffers allocated for network address storage over IPv4 are limited to 32 bits (16 character string).

After the protocol stack is installed, IIS automatically begins supporting IPv6 on your Web server.

Web sites that are already running must be restarted before they begin listening for IPv6 requests. Sites created after IPv6 support is enabled automatically listen for IPv6.

**Functionality**

The IIS core functionality has not changed as a result of IPv6 support; however, only a subset of all IIS 6.0 functionality is available for IPv6.

- IIS Manager does not display IPv6 addresses as it does for IPv4 addresses.
- Bandwidth Throttling is not supported for IPv6 Web sites.
- The ServerBindings metabase property does not support storing IPv6 literal addresses (as defined in RFC 2732). This behavior limits IIS to support host-header routing only.
- The **IP Address Restrictions** feature in IIS does not support IPv6 addresses or IPv6 prefixes.
- The EnableReverseDnsLookup metabase property is not supported. Setting this property to true does not cause the REMOTE_HOST server variable to return the DNS name of the client, as it does for IPv4. REMOTE_HOST will always contain the IPv6 address regardless of the **EnableReverseDNSLookup** setting.
- Site routing based on IP addresses is not supported for IPv6. Server address with sites that are configured to route based on IPv4 addresses will not respond to IPv6 requests.
- IPv6 is enabled for all sites. You cannot configure IPv6 support at the machine level. In other words, you cannot configure individual sites to respond to IPv6 traffic while other sites on the same server respond to IPv4 traffic. It is possible, however, to limit IPv6 traffic to a specific site if that site is configured to use IP-based routing on an IPv4 address.
- Logging. IIS writes IPv6 addresses to the log file when IPv6 is enabled and client computers connect to the server using IPv6 addresses. Log parsing tools must support IPv6 address formats if they are to be used with log files for IPv6 sites.
- SSL. Due to the IP routing restriction for IPv6, IIS deployments designed for IPv6 addresses are limited to one Secure Sockets Layer (SSL) site per computer.

**Related Information**

- For information about how to install the IPv6 network protocol stack, see IPv6 and IIS 6.0.

# What's Changed

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are some notable and important differences in default behavior and settings between IIS 4.0, IIS 5.0, IIS 5.1, and IIS 6.0.

This topic includes the following information:

- Core Functionality and Services
- Metabase Configuration
- Administration
- Programmatic Administration
- Active Server Pages (ASP)
- ASP Hang Detection
- Security
- Performance
- IIS Utility Components
- IIS on 64-bit Versions of the Windows Server 2003 Family

The following table summarizes the important differences between versions of IIS.

| | **IIS 4.0** | **IIS 5.0** | **IIS 5.1** | **IIS 6.0** |
|---|---|---|---|---|
| **Platform** | Windows NT 4.0 | Windows 2000 | Windows XP Professional | Windows Server 2003 family |
| **Architecture** | 32-bit | 32-bit | 32-bit and 64-bit | 32-bit and 64-bit |
| **Application process model** | TCP/IP kernel<br>MTX.exe | TCP/IP kernel<br>DLLhost.exe (multiple DLL hosts in medium- or high-application isolation) | TCP/IP kernel<br>DLLhost.exe (multiple DLL hosts in medium- or high-application isolation) | HTTP.sys kernel<br>When IIS is running in IIS 5.0 isolation mode: Inetinfo.exe (for in-process applications) orDLLhost.exe (for out-of-process applications)<br>When IIS is running in Worker Process Isolation Mode: W3wp.exe (multiple worker processes) |
| **Metabase configuration** | Binary | Binary | Binary | XML |

| | | | | |
|---|---|---|---|---|
| **Security** | Windows authentication<br>SSL | Windows authentication<br>SSL<br>Kerberos | Windows authentication<br>SSL<br>Kerberos<br>Security wizard | Windows authentication<br>SSL<br>Kerberos<br>Security wizard<br>Passport support |
| **Remote administration** | HTMLA | HTMLA | No HTMLA<br>Terminal Services | Remote Administration Tool (HTML)<br>Terminal Services |
| **Cluster support** | In Windows NT 4.0 | IIS clustering | Windows support | Windows support |
| **WWW services** | IIS on Windows NT 4.0 | Personal Web Manager on Windows 9x<br>IIS on Windows 2000 | IIS optionally on Windows XP Professional | IIS on a member of the Windows Server 2003 family |

**Core Functionality and Services**

IIS 6.0 has been redesigned to take advantage of the base Windows kernel, HTTP.sys. This allows for built-in response and request caching and queuing, as well as the ability to route application process requests directly to the worker processes, which improves reliability and performance.

IIS 6.0 introduces two modes of operation in order to configure your application environments: worker process isolation mode and IIS 5.0 isolation mode. The default isolation mode upon installing IIS 6.0 depends on whether you perform a clean installation or an upgrade.

- After a clean install of IIS 6.0, IIS runs in worker process isolation mode.
- After an upgrade from an earlier version of IIS 6.0, the isolation mode is the same as configured on the previously-installed version of IIS 6.0.
- After an upgrade from IIS 5.0 or IIS 4.0, IIS 6.0 runs in IIS 5.0 isolation mode by default to maintain compatibility with your existing applications.

For information about switching from one isolation mode to the other, see Configuring Isolation Modes.

**IIS 5.0 Isolation Mode**

IIS 5.0 isolation mode manages application processes in a similar fashion to the process management in IIS 5.0: all in-process applications run inside Inetinfo.exe, and out-of-process applications run in separate DLL hosts. Some existing applications may not have been written to run concurrently, or to store session state separately from the application. Therefore, running processes in IIS 5.0 isolation mode ensures compatibility for most existing applications. The following illustration shows how application processes are handled in IIS 5.0 isolation mode.




**Worker Process Isolation Mode**

When configured to execute in worker process isolation mode, all application code runs in an isolated environment. This design removes some of the existing bottlenecks. Worker process isolation mode allows the administrator to isolate anything from an individual Web application to multiple sites in their own self-contained worker process. This prevents one application or site from stopping another. In addition, separating applications or sites into their own process space simplifies a number of management tasks, such as restarts (independent of all other sites or applications running on the system), changing a component used by the application, debugging, monitoring counters, throttling resources, and so forth. The following illustration shows how applications are managed by IIS in worker process isolation mode.




HTTP requests are routed to the correct application pool queue, which means that user mode worker processes serving an application pool pull the requests directly from the kernel and eliminate the unnecessary process hops encountered when

sending a request to an out-of-process DLL host. In IIS 6.0, there is no longer the notion of in-process applications; all necessary HTTP application run-time services, such as ISAPI extension support, are equally available in any application pool. This design prevents a malfunctioning HTTP application or Web site from disrupting other HTTP applications (or other Web sites) served from other processes on that computer. Unloading components becomes easier because with isolated application processes, the process can, if necessary, be terminated to unload all resources, with no effect on other content or applications being served from other processes. It is also beneficial to be able to leverage other operating system services available at the process level (for example CPU throttling), per application pool.

Furthermore, critical portions of worker process isolation mode that maintain the overall functioning of World Wide Web Publishing Service (WWW service) run entirely outside of the worker processes. The IIS 6.0 kernel-mode driver, HTTP.sys, which is the universal HTTP processor for Windows, and the WWW Service Administration and Monitoring component isolate the critical portions of the core Web server. Both of these components are protected and do not allow third-party code to be loaded into them. This design prevents a malfunctioning HTTP application from disrupting WWW services on the server.

For more information about isolation modes, see Application Isolation Modes.

**Metabase Configuration**
The metabase for IIS 6.0 is stored in an XML file instead of in binary format as it was in earlier versions of IIS. The location remains the same, but the ways it can be manipulated -- updated, rolled back, restored, and extended -- have changed. There are two significant files instead of one: MetaBase.xml and MBSchema.xml.

**Administration**

- In IIS 4.0, applications could either be run in the same process as the Internet service or in a separate process. In IIS 5.0 and 5.1, applications can be grouped into pooled processes for increased performance and improved scalability. For more information, see About Configuring Applications. In IIS 6.0 worker process isolation mode, applications can grouped in any number of application pools.

- The Application Mappings property sheet contains a list of Hypertext Transport Protocol (HTTP) verbs that are processed by an application that is mapped to certain file types. This list of verbs represents a change from IIS 4.0. In IIS 4.0, the list contained verbs that were *excluded*, or not processed. This change was made to accommodate new HTTP verbs as they are added to the protocol. For more information about application mappings, see Setting Application Mappings in IIS 6.0.

- Clustering is not a feature of IIS 6.0 (IISsynche.exe is not supported). Clustering is a feature of the Windows Server 2003 family. For information about Windows Clustering (MSCS), see Windows Server 2003 family Help.

- The location of custom error files has changed from IIS 4.0 to IIS 5.0.

- The Web-based Internet Services Manager (HTML) has been replaced by a Web application called the Remote Administration (HTML) tool. To remotely administer IIS by using the Remote Administration (HTML) tool, see Administering Servers Remotely in IIS 6.0.

**Programmatic Administration**
In previous versions of IIS, programmatic administration of IIS was possible with Admin Base Objects (ABO) from compiled C++ applications, or with Active Directory Services Interfaces (ADSI) from C++ or script files. IIS 6.0 includes a provider for Windows Management Instrumentation (WMI), a technology that allows administrators to control all services and applications programmatically. For more information about WMI, see "IIS Administration Technologies" in the IIS Software Development Kit (SDK) on MSDN.

**Active Server Pages**
Beginning with IIS 6.0, Microsoft Active Server Pages (ASP) can be used along with Microsoft ASP.NET. For information about configuring IIS to run ASP.NET applications, see About ASP.NET. For news on changes to ASP functionality in IIS 6.0, see Important Changes in ASP.

Because the worker process, W3wp.exe, runs as the Network Service account in IIS 6.0 worker process isolation mode, you must configure Launch and Access permissions to enable ASP debugging for Script Debugger and Visual InterDev. For more information, see Enabling ASP Debugging.

**ASP Hang Detection**
When an IIS Web site is busy, there may be instances when the maximum number of ASP threads has been spawned and some of the ASP threads are hung, resulting in degraded performance. IIS 6.0 has the ability to solve the problem of hung threads by recycling the worker process that hosts that particular instance of the ASP ISAPI extension, ASP.dll. When ASP threads are hung

in IIS 6.0, ASP.dll calls the ISAPI server support function **HSE_REQ_REPORT_UNHEALTHY**, and the WWW service recycles the worker process that hosts ASP.dll, and make an entry in the event log.

For more information about ISAPI server support functions, see **ServerSupportFunction** in the ISAPI Extension Reference at MSDN® Online.

**Security**

One of the most important changes in IIS 6.0 addresses Web server security. In order to take a more proactive stance against malicious users and attackers, IIS is not installed by default on members of the Microsoft Windows Server 2003 family.

| **Important** |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0. |

With the Web Server Certificate Wizard and the CTL Wizard, you can synchronize Web and NTFS security settings, obtain and install server certificates, and create and modify certificate trust lists. You can also select a cryptographic service provider (CSP) for encrypting data with a certificate.

Other security changes in IIS 6.0 include the following:

- Disabled on upgrades: The World Wide Web Publishing Service (WWW service) is disabled on Windows Server 2003 family upgrades, unless one of the following is true:

    - You have already run the IIS Lockdown Wizard on your Windows 2000 Server before starting the upgrade process. The IIS Lockdown Wizard reduces surface attack by disabling unnecessary features, and it allows you to decide which features to enable for your site. The IIS Lockdown Wizard is available at IIS Lockdown Tool.

| **Important** |
| --- |
| If you use the WWW service, we strongly recommend that you run the IIS Lockdown Wizard on your Windows 2000 Server before upgrading to a product in the Windows Server 2003 family. The IIS lockdown Wizard will help secure your computer by disabling or removing unnecessary features that are present in your Windows 2000 Server installation. These features would otherwise have remained on your machine after upgrading, leaving your server vulnerable to attacks. |

- The registry key RetainW3SVCStatus has been added to the registry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC. Under RetainW3SVCStatus, you can add any value and then assign a DWORD value to it. For example, you can create the key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\RetainW3SVCStatus\do_not_disable with the DWORD value of 1.

- In the unattended install case, an entry "DisableWebServiceOnUpgrade = false" exists in the unattended install script.

- Disabling IIS through Group Policy: With members of the Windows Server 2003 family, domain administrators can prevent users from installing IIS on their machines.

- Running as an account with a low level of access rights: IIS worker processes run in a user context of few access rights. This drastically reduces the effect of potential attacks.

- Secure ASP: All ASP built-in functions always run as the account with very few access rights, IUSR_*computername*.

- Restriction on running executables: In order to run most executables in the system folder (such as cmd.exe) you must be a member of the Administrators group, the LocalSystem, Interactive, or Service account. This restriction limits remote access to Administrators, so anonymous users cannot run executables.

- Patch management: With patch management, administrators can get the latest security patches installed without interrupting service.

- Known extensions: IIS serves requests only to files with known file name extensions. The server rejects requests for content if the file name extension is not mapped to a known file type or to an application that services requests for that file type.

| **Note** |
| --- |
| The word *extensions* has two meanings: Either Web service extensions that enable pages to serve dynamic content, for example, .asp or .aspx; or file name extensions, which indicate the file type, such as .exe, .txt, or .inc. |

Use the **WebSvcExtRestrictionList** metabase key to lock down DLLs and executables that handle dynamic content. To lock down file name extensions, use the MimeMap feature in the IIS metabase.

- Write protection for content: Anonymous users (running as IUSR_*computername* account) are denied write access to Web content by default.

- Timeouts and limits: In IIS 6.0, settings are set to aggressive and secure defaults to minimize attacks due to timeouts and limits that were previously too generous.

- Upload data limitations: Administrators can limit the data that can be uploaded to a server.

- Buffer overflow protection: Worker processes detect and exit programs if a buffer overflow is detected.

- File verification: IIS verifies whether the requested content exists before it gives the request to a request handler (ISAPI extension).

- Index this resource: This permission is now enabled by default.

- Script source access: This permission, which allows access to the *source code* of scripts in ASP pages and other scripts, is new and is disabled by default. It is available if either the **Read** or **Write** permission is selected.

- Subauthentication: This is no longer enabled by default on a new installation of IIS 6.0. For more information, see Configuring Subauthentication.

- UNC authentication: In this version of IIS, the UNC authentication method checks for user credentials. For more information, see UNC Authentication in IIS 6.0.

- New policy: The "Prevent IIS from Installing" policy has been added to the Windows Server 2003 family of products. This policy allows a domain administrator to control which computers in the domain can install IIS. For more information, see "Group Policy" in Help and Support Center for Windows Server 2003.

- Fortezza: Support for this has been removed.

**Performance**

To limit the amount of memory allocated to ASP pages, IIS has set the default value of the AspScriptFileCacheSize Metabase Property to 250 ASP pages, and the default value of the AspScriptEngineCacheMax Metabase Property to 125 script engines. The **ASPScriptFileCacheSize** can be set higher on sites with a large set of frequently requested ASP pages. This improves performance because ASP page compilation is substantially slower than retrieving pages from cache. On a site with only a small number of frequently requested ASP pages, memory can be saved by setting this number to a smaller value.

**IIS Utility Components**

- Collaboration Data Objects for Windows NT Server (CDONTS): CDONTS has been removed from the Windows Server 2003 family. If your Web applications use CDONTS, you can convert them to Microsoft Collaboration Data Objects

(CDO). Most methods in CDONTS have matching methods in CDO, but might be named differently. For reference material for CDO in the Platform Software Developer Kit (PSDK), see Overview of CDO at MSDN Online.

- IIS utility components not installed: Ad Rotator, Browser Capabilities, Content Linker, Content Rotator, Counters, Logging Utility, My Info, Page Counter, Status, and tools are not installed with IIS 6.0. However, if you upgrade your Web server from a previous version of IIS, the utility components are not removed.

**IIS on 64-bit Versions of the Windows Server 2003 Family**

On the 64-bit versions of the Windows Server 2003 family of operating systems, IIS runs as a 64-bit application. This means that 32-bit applications cannot be called from IIS on the 64-bit versions of the Windows Server 2003 family of operating systems. For example, the Jet database engine will not convert to a 64-bit application, so you cannot use ActiveX® Data Objects (ADO) to open a Microsoft Access database from an ASP page. However, you can still use ADO to access other drivers, like SQL and Exchange.

# IIS Manager

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS Manager is a graphical interface for configuring your application pools or your Web, FTP, SMTP, or NNTP sites. With IIS Manager, you can configure IIS security, performance, and reliability features. You can add or delete sites; start, stop, and pause sites; back up and restore server configurations; and create virtual directories for better content management, to name only a few of the administrative capabilities. In previous releases of IIS, this tool was called the Internet Service Manager.

| **Important** |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down — which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the **Web Service Extensions** node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0. |

**Procedures**

**To start IIS Manager**

- From the **Start** menu, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.

**To start IIS Manager from the Run dialog box**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **inetmgr**, and click **OK**.

You can also access IIS from the Computer Management window. Accessing IIS in this way does not give you the range of administration options offered by IIS Manager; however, it does offer quick access and limited management options for your Web sites.

**To start IIS Manager from the Computer Management window**

1. From the **Start** menu, right-click **My Computer**, and click **Manage**.
2. In the console tree, expand the **Services and Applications** node.
3. Click **Internet Information Services**. The names and states of your Web sites appear in the details pane.
4. In the console tree, expand the **Internet Information Services** node and any subsequent Web site nodes to see a list of directories and virtual directories for that Web site.

**Related Information**

- For more information about using IIS Manager and administering IIS, see the Server Administration Guide.
- For information about administering IIS remotely, see Administering Servers Remotely in IIS 6.0.

# Frequently Asked Questions about IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1


1. Why do features like FrontPage Server Extensions, WebDAV publishing, ASP, and server-side includes not work on IIS 6.0 by default?
2. What is a Web service extension and how do I use the Web Service Extensions folder?
3. How do I publish documents or Web pages?
4. How do I create a virtual directory on a Web or FTP site?
5. How do I create a Web site?
6. How do I create multiple Web sites?
7. How do I replicate Web content?
8. How do I stop and restart Internet services?
9. Where can I find information about accessibility features for disabled users?
10. How do I get technical support for IIS?
11. Can I change the name of my Web site and also redirect requests for the old site name to the new one?
12. Can I use IIS Manager to remotely administer my site through a proxy server?
13. How do I determine whether my drive is configured to use NTFS?
14. How do I convert a drive or partition formatted for the file allocation table (FAT) file system to the NTFS file system?
15. Can I set up an FTP site with different levels of access for different users?
16. How can I confirm that a server certificate is attached to a Web site?
17. Can I attach more than one server certificate to a Web site?
18. Can I attach the same server certificate to more than one site?
19. Can I attach a server certificate to an FTP site?
20. Should I create a backup copy of my server certificate?
21. How do I determine how many users are connected to my FTP service?
22. Does ASP debugging work in IIS 6.0 worker process isolation mode?
23. Why do features like FrontPage Server Extensions, WebDAV publishing, ASP, and server-side includes not work on IIS 6.0 by default?

# Why do features like FrontPage Server Extensions, WebDAV publishing, ASP, and server-side includes not work on IIS 6.0 by default?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In order to take a more proactive stance against malicious users and attackers, IIS is not installed on members of the Microsoft Windows Server 2003 family by default. Furthermore, when you initially install IIS, the service is installed in a highly secure and locked mode. By default, IIS serves only static content — meaning features like ASP, ASP.NET, server-side includes, WebDAV publishing, and FrontPage Server Extensions do not work. You can serve dynamic content and unlock these features using IIS Manager. See Configuring IIS 6.0 Web Service Extensions.

**Related Sections**

For more information, see Security in IIS 6.0.

# What is a Web service extension and how do I use the Web Service Extensions folder?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The **Web Service Extensions** folder is the user interface for the new IIS 6.0 lockdown feature. This feature is a manifest of ISAPI extensions and CGIs with user-specified permissions, meaning, administrators must set the permissions to allow specific ISAPIs and CGIs to run on your server. Administrators can also specify the names of ISAPIs or CGIs that are forbidden to run on your server. Before loading an ISAPI extension .dll file or CGI .exe file, IIS checks this manifest for the permissions on the file. If the file is permitted to run, then the request proceeds normally. If the file is not permitted to run, then IIS returns a 404.2 error response to the client machine. The HTML page for a 404.2 error looks like a standard 404 error page, so the client machine processes the request as though the file did not exist. IIS logs the 404.2 error, which administrators can view to assess problems or potential threats against the server.

**Related Sections**

For more information, see Configuring IIS For Dynamic Content.

# How do I publish documents or Web pages?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

**To publish documents or Web pages**

1. Move your files to the \Inetpub\Wwwroot directory.
2. Type **http://***servername***/***filename* in the address bar of your Internet browser to see your published files.

For more information, see Site Setup.

# How do I create a virtual directory on a Web or FTP site?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use IIS Manager to create virtual directories on your Web site. For procedural information about how to accomplish this task, see Using Virtual Directories.

# How do I create a Web site?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you install IIS on a computer running a member of the Windows Server 2003 family, a default Web site is set up for you. You can publish your content here immediately.

# How do I create multiple Web sites?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To create multiple Web sites, you must first ensure that each site has unique identification. To accomplish this, you need to contact your network administrator to either obtain multiple IP addresses or to assign multiple host header names. For more information, see Domain Name Resolution.

# How do I replicate Web content?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the Content Deployment service (formerly known as the Content Replication System, or CRS) to replicate content and settings across servers in a Web cluster. Content Deployment is a feature of Microsoft Content Management Server. For more information about Content Deployment service, see the Content Management Server 2002 Product Documentation on the Windows Server System Web site.

# How do I stop and restart Internet services?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Internet services are most effectively shut down and restarted differently from other Windows services. For information about stopping and restarting IIS, see Restarting IIS.

# Where can I find information about accessibility features for disabled users?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

For general information about accessibility for Windows, see "Accessibility for people with disabilities" in Help and Support Center for Windows Server 2003.

# How do I get technical support for IIS?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If IIS was pre-installed as part of an original equipment manufacturer (OEM) package, or if it was otherwise included with your personal computer purchase, you must contact the manufacturer of your computer or server for technical support. Refer to the documentation that came with your computer or server for the telephone number of the manufacturer.

If you acquired this product separately at retail, and have a technical question that you cannot answer by using the IIS 6.0 Operations Guide, see the Microsoft Help and Support Center Web site.

# Can I change the name of my Web site and also redirect requests for the old site name to the new one?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure your Web site to respond to both the old name and the new name, provided your network correctly routes both requests to your computer. This way, visitors can still reach your site by using the old name, and will be informed of the new name. For instructions on how to assign names to Web sites, see Naming Web Sites.

# Can I use IIS Manager to remotely administer my site through a proxy server?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Yes. To learn about remote administration using IIS Manager and other tools, see Administering Servers Remotely in IIS 6.0.

# How do I determine whether my drive is configured to use NTFS?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

1. In Windows Explorer, right-click the drive, and click **Properties** to open its property sheet.
2. On the **General** tab, the file system type is listed.

# How do I convert a drive or partition formatted for the file allocation table (FAT) file system to the NTFS file system?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

At the command prompt, type **convert** *drive letter:* **/FS:NTFS**.

An NTFS drive cannot be converted back to FAT. For more information, see "File systems" in Windows Help.

# Can I set up an FTP site with different levels of access for different users?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

**To limit access to only certain computers**

1. In IIS Manager, right-click the FTP site, and click **Properties**.
2. Click the **Directory Security** tab, and click **Edit**.

   You can assign or deny access to all computers or subnets except those you specify.

On an NTFS file system, you could also use a single virtual directory and set the NTFS permissions to assign or deny different types of access to different Windows user accounts. For more information, see Isolating FTP Users.

# How can I confirm that a server certificate is attached to a Web site?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

1. In IIS Manager, right-click the Web site, and click **Properties**.

2. Click the **Directory Security** tab.

3. Under **Secure communications**, if the **View Certificate** button is activated, there is a certificate attached to the Web site.

   If the button is not activated, you must attach a server certificate to the site to use the Secure Sockets Layer (SSL) features.

# Can I attach more than one server certificate to a Web site?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

No. Each Web site can have only one server certificate attached to it.

# Can I attach the same server certificate to more than one site?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Yes. A server certificate can be attached to as many Web sites as needed.

# Can I attach a server certificate to an FTP site?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

No. FTP sites do not support Secure Sockets Layer (SSL) features.

# Should I create a backup copy of my server certificate?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Yes. Your server certificate is a valuable investment, and is the key to your server's Secure Sockets Layer features. To create a backup copy of your server certificate, copy the entire certificate on to a floppy disk and store it in a safe place.

# How do I determine how many users are connected to my FTP service?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

1. In IIS Manager, right-click an FTP site, and click **Properties**.
2. On the **FTP Site** tab, click **Current Sessions**.

   The **FTP User Sessions** dialog box appears indicating the number of connected users.

For more information, see Performance Tuning.

# Does ASP debugging work in IIS 6.0 worker process isolation mode?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Yes, when you configure Launch and Access permissions to enable ASP debugging in worker process isolation mode for Script Debugger and Visual InterDev. For more information, see Enabling ASP Debugging.

# Additional IIS 6.0 Resources

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The following is a list of Web sites that can provide helpful supplemental information that isn't included in this documentation set. The links, including their content, are subject to change without notice.

- Internet Information Services
- Internet Information Services Community Center
- Internet Information Services Software Development Kit
- Microsoft
- Microsoft Windows
- Microsoft TechNet
- Microsoft Service Providers
- Microsoft Windows Update
- MSDN Online
- Active Server Pages
- ASP.NET
- Microsoft .NET Framework Community
- Microsoft .NET Framework Software Development Kit
- Data Access and Storage Developer Center
- Microsoft Learning
- Microsoft FrontPage
- Microsoft Help and Support
- InterNIC List of Accredited Domain Registrars
- World Wide Web Consortium
- 15 Seconds
- Internet Engineering Task Force
- National Security Agency Central Security Service -- Information Assurance
- 25+ ASP Tips to Improve Performance and Style

# Release Notes_IIS_SP1_Ops

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

The release notes for Internet Information Services (IIS) 6.0 are incorporated into the release notes for the Microsoft Windows Server 2003 family of products. For late-breaking technical information, known problems, and upgrade information, please refer to the following: readme.htm, relnotes.htm, read1st.txt, and \docs*.txt.

# Before Configuring IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Before configuring IIS, it is important to understand issues users sometimes face. By addressing these issues now, you can enjoy a smoother configuration process. This topic outlines these issues and also links you to additional information that you might want to review before beginning your IIS configuration.

This topic includes the following information:

- Internet Connection Firewall
- Remote Administration
- IP Address and Domain Name Registration
- Permissions and Access Control

**Internet Connection Firewall**

The Windows Server 2003 family provides a software-based firewall to prevent unauthorized connections to your server from remote computers. The Internet Connection Firewall is disabled by default; however, if you have enabled the firewall in its default configuration after installing a member of the Windows Server 2003 family and before installing IIS, clients will not be able to connect to your server. The following procedure configures Internet Connection Firewall to allow clients to initiate Web and other IIS-related connections to your server.

**To configure Internet Connection Firewall for IIS**

1. From the **Start** menu, click **Control Panel**.
2. Double-click **Network Connections**.
3. Right-click the **Local Area Connection**, and click **Properties**.
4. Click the **Advanced** tab.
5. If you do not want to use the Internet Connection Firewall, make sure the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box is disabled, and click **OK**.
6. If you do want to use the Internet Connection Firewall, make sure the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box is enabled, and click **Settings**.
7. On the **Services** tab, enable a service to which you want to allow access to clients.
8. In the **Service Settings** dialog box that appears after enabling a service, do one of the following:
   - If you are enabling a service on the same computer on which you are working, the correct computer name is already filled in. Click **OK**.
   - If you are enabling a service on a different computer on your network, type the name or IP address of the computer hosting the service you are enabling, and click **OK**.
9. Repeat steps 7 and 8 until all the services you want accessible to clients are enabled.

**Remote Administration**

Many network administrators configure servers remotely, that is, from a computer other than the server being configured. If you are installing IIS on one server, and configuring and administering that IIS installation from a different computer, you should configure Remote Administration on the server running IIS. You can use the following tools for this purpose:

- IIS Manager: You can use IIS Manager on your server to remotely connect to and administer an intranet server running IIS 4.0, IIS 5.0, or IIS 5.1 (IIS 3.0 is not supported).

- Terminal Services: Terminal Services does not require you to install IIS Manager on the remote client computer because, once connected to the server running IIS, you use IIS Manager on the Web server as if logged on locally.

- Remote Administration (HTML) Tool: You can use the Remote Administration (HTML) tool to administer your IIS Web server from any Web browser on your intranet. This version of the Remote Administration tool runs only on servers running IIS 6.0.

For information about how to configure remote administration, see Administering Servers Remotely in IIS 6.0.

**IP Address and Domain Name Registration**
If you are configuring IIS on a server connected directly to the Internet, that is, not within a corporate network or behind a router, you must have on hand the IP address and domain name of each site you will create in this IIS configuration. For more information, see Domain Name Resolution and Hosting Multiple Web Sites on a Single Server.

If you are configuring IIS on a server within a corporate network or behind a router, you must ask the network administrator for an IP address, and if applicable, to configure the router and/or firewall to allow traffic to flow to and from the IIS server. Registering a domain name is usually not necessary in these cases because IIS servers within a corporate network are typically for Intranet use only.

When obtaining IP addresses for your Web sites, make sure they are static IP addresses. Dynamic IP addresses are not suitable for any kind of IIS configuration.

**Permissions and Access Control**
If you are configuring IIS so that it will be relying on other servers or network components, such as routers, to serve content, make sure that the other servers or network components have been configured to allow access to the IIS server, and that you have the appropriate authentication information on hand before you begin configuring IIS. For more information, see Authentication in IIS 6.0 and Access Control in IIS 6.0.

**Related Information**

- For a checklist of additional issues to consider before configuring IIS, see "Application server role: Configuring an application server" in Windows Help.

# Installing IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This topic describes 3 methods for installing IIS:

- Using the Configure Your Server Wizard.
- Using Add or Remove Programs from Control Panel.
- Using unattended setup.

**Note**

You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI.

This topic also lists the directories created on install, describes the IIS initial configuration backup, and briefly describes IIS optional components.

**Important**

To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down — which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Services Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0.

Microsoft strongly recommends installing IIS on an NTFS-formatted drive. NTFS is a more powerful and secure file system than FAT and FAT32. For more information, see Securing Files with NTFS Permissions.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the runas command to run IIS Manager as an administrator. At a command prompt, type **runas**

**/user:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**.

**Procedures**

**To install IIS using the Configure Your Server Wizard**

1. From the **Start** menu, click **Manage Your Server**.
2. Under **Managing Your Server Roles**, click **Add or remove a role**.
3. Read the preliminary steps in the **Configure Your Server Wizard** and click **Next**.
4. Under **Server Role**, click **Application server (IIS, ASP.NET)** and then click **Next**.

   By default, the wizard installs and enables IIS, COM+, and DTC.
5. If you want to serve either of the optional technologies (FrontPage Server Extensions or ASP.NET), on the **Application Server Options** page, select the appropriate check boxes, and then click **Next**.

6. Read the summary and click **Next**.
7. Complete the wizard, and then click **Finish**.

| **Note** |
| --- |
| The **Configure Your Server Wizard** enables ASP.NET by default, unlike the **Add/Remove Windows components** install method below. |

For more information about the **Configure Your Server Wizard**, see "Configuring Your Server Wizard" in Windows Help.

**To install IIS, add components, or remove components using Control Panel**

1. From the **Start** menu, click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. In the **Components** list box, click **Application Server**.
5. Click **Details**.
6. Click **Internet Information Services Manager**.
7. Click **Details** to view the list of IIS optional components. For a detailed description of IIS optional components, see "Optional Components" in this topic.
8. Select all optional components you wish to install.

   | **Note** |
   | --- |
   | The **World Wide Web Publishing Service** optional component includes important subcomponents like the **Active Server Pages** component and **Remote Administration (HTML)**. To view and select these subcomponents, click **World Wide Web Publishing Service** and then click **Details**. |

9. Click **OK** until you are returned to the **Windows Component Wizard**.
10. Click **Next** and complete the **Windows Component Wizard**.

**Unattended Setup**

To simplify the process of setting up IIS on multiple computers running a member of the Windows Server 2003 family, you can run setup unattended. To do this, create and use an answer file, which is a customized script that automatically answers the setup questions.

For information about how to create an answer file and to view a table of all IIS unattended setup parameters, see Appendix E: Unattended Setup.

**IIS Directories**

IIS installs the following directories:

- \InetPub
- **systemroot**\Help\IISHelp
- **systemroot**\System32\InetSrv
- **systemroot**\System32\InetSrv\MetaBack

These directories contain user content and cannot be moved. You can, however, select the location of your Wwwroot and Ftproot directories at installation by using a script during unattended setup. If you uninstall IIS, the IISHelp directory is removed. The InetPub and InetSrv directories remain on your computer.

**IIS Initial Configuration Backup**
When you first install IIS, a backup of the initial metabase configuration is automatically created in the *systemroot*\System32\InetSrv\MetaBack directory. This backup can be used to restore the IIS configuration to its state immediately following IIS installation. This is a useful tool for solving metabase corruption or configuration problems, and can help you recover a known good configuration without needing to reinstall IIS. This backup is not password protected, and can only be used to restore settings on the system on which it was created. See Backing Up and Restoring the Metabase in IIS 6.0 for information about restoring the initial IIS configuration backup.

It is strongly recommended that following IIS installation, and before any configuration changes are made, you create a password-protected backup of the IIS configuration. Unlike the automatic initial configuration backup, a password-protected backup is system independent, and can be used to restore settings on other systems running IIS 6.0. See Backing Up and Restoring the Metabase in IIS 6.0 for information about creating a password-protected backup of the IIS configuration.

**Optional Components**
IIS includes optional components that you can enable or disable at anytime through the **Add or Remove Programs** item in Control Panel, or using unattended setup. Descriptions of these components and the impact to your current IIS settings are described below.

Files for Active Server Pages (ASP), Internet Data Connector, server-side includes, and WebDav are installed, but disabled by default on a clean installation. All IIS functionality is enabled by default on an upgrade. When one of these components is selected or cleared from the **Windows Components Wizard** or by using an unattended file, the component is enabled or disabled in the WebSvcExtRestrictionList Metabase Property.

**BITS Server Extensions**
The Background Intelligent Transfer Service (BITS) is a background file transfer mechanism and queue manager, also known as a drizzle service. BITS throttles file requests to minimize bandwidth consumption and enhance the end-user experience. Enable BITS with IIS to maintain Web server quality of service. From the Windows Components Wizard, click **Details** to enable the following BITS Server components:

- **BITS Server Extensions ISAPI**: Enable this option to drizzle IIS requests with the help of the BITS server.
- **BITS Server Extensions Snap-in**: Enabling this option to access and view the BITS graphical user interface (GUI).

**Common Files**
For the sake of added security in your server environment, you may choose to deselect some of the common files. However, if you deselect the Common Files option, all of the common files are subsequently deselected and IIS will not be installed on your machine. Therefore, if you want to install IIS on your computer, leave this option selected. If you want to limit the services and components installed with IIS, deselect the individual components that are listed below the Common Files component.

**File Transfer Protocol (FTP) Server**
The File Transfer Protocol (FTP) is used to copy files to and from remote computer systems on a network that uses Transmission Control Protocol/Internet Protocol (TCP/IP).

**FrontPage 2002 Server Extensions**
FrontPage 2002 Server Extensions allow you to view and manage a Web site in a graphical user interface using FrontPage as the authoring environment. FrontPage allows you to quickly create Web sites on your server, as well as, create, edit, and post Web pages to IIS remotely. While you are creating your site, FrontPage keeps a connection open to IIS, saving and changing the Web files so you can view your site. If you choose not to install the FrontPage Server Extensions, you will have to manually copy all of your Web content, configure your settings, and in some cases, manually register applications already registered in FrontPage. Microsoft FrontPage Server Extensions are not supported for resources in Microsoft Clustering.

**Internet Information Services Manager**
IIS Manager is a graphical user interface to administer your Web site. In previous releases of IIS, this tool was called the Internet Service Manager. Without IIS Manager, you can still manage your server, but you must use coded scripts that call on the IIS APIs to create sites, applications, virtual directories, and security settings.

**NNTP Service**
Use Network News Transfer Protocol (NNTP) to distribute network news messages to NNTP servers and to NNTP clients (news readers) on the Internet. NNTP provides for the distribution, inquiry, retrieval, and posting of news articles by using a reliable stream-based transmission of news on the Internet. NNTP is designed so that news articles are stored on a server in a central database, thus users can select specific items to read. Indexing, cross-referencing, and expiration of aged messages are also provided.

If you have NNTP installed, you can view Microsoft News (NNTP) Service Help by typing **file:\\%systemroot%\help\news.chm** in your browser address bar and pressing ENTER.

**SMTP Service**

Use Simple Mail Transfer Protocol (SMTP) to set up intranet mail services that work in conjunction with IIS. SMTP is a TCP/IP protocol for sending messages from one computer to another on a network. This protocol is used on the Internet to route e-mail.

If you have SMTP installed, you can view Microsoft Mail (SMTP) Service Help by typing **file:\\%systemroot%\help\mail.chm** in your browser address bar and pressing ENTER.

**World Wide Web Publishing Service**

IIS serves pages to the Internet and the World Wide Web. This component must be installed for IIS to perform its primary service. If you do not select this option, you disable IIS.

The World Wide Web Publishing Service (WWW service) includes the following subcomponents:

- About ASP: Select to enable ASP on your server. If this option is not selected, all .asp requests return a 404 error.
- **Internet Data Connector**: Select to enable Internet Data Connector on your server. If this option is not selected, all .idc requests return a 404 error.
- Administering Servers Remotely in IIS 6.0: Select to enable remote Web administration of your IIS Web server from any Web browser on your intranet. Once you install IIS and view your Web sites through IIS Manager, IIS creates a site called **Administration**.
- **Remote Desktop Web Connection**: Select to enable connectivity to a computer's desktop from a remote location and run applications as if you were sitting at the console.
- Using Server-Side Include Directives: Select to enable server-side include files on your server. If this option is not selected, all .shtm, .shtml, and .stm requests return a 404 error.
- Web Authoring with WebDAV: Select to allow Web Distributed Authoring and Versioning (WebDAV) on your server. WebDAV is similar to File Transfer Protocol, with the exception being that WebDAV allows any WebDAV client to publish and change content in a WebDAV directory using HTTP.
- Web Site Administration: Select to install the World Wide Web Publishing Service. If this option is not selected, IIS does not run on your server.

**Related Information**

- For information about unattended setup, see "Planning for Unattended Setup" in Windows Help.

# Uninstalling IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This topic describes how to uninstall IIS using either of the following 2 methods:

- Using the Configure Your Server Wizard.
- Using Add or Remove a Windows Component in Control Panel.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To uninstall IIS using the Configure Your Server Wizard**

1. From the **Start** menu, click **Manage Your Server**.
2. Under **Managing Your Server Roles**, click **Add or remove a role**.
3. Read the preliminary steps in the **Configure Your Server Wizard** and click **Next**.
4. Under **Server Role**, click **Application server (IIS, ASP.NET)** and then click **Next**.
5. In the **Role Removal Configuration** dialog, check the **Remove the application server role** check box.
6. Complete the wizard, and then click **Finish**.

**To uninstall IIS using Add or Remove a Windows Component in Control Panel**

1. From the **Start** menu, click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. In the **Components** list click the **Application Server** check box.
5. Click **Next**.
6. Click **Close** when uninstallation of IIS is complete.

**Related Information**

For more information about the Configure Your Server Wizard, see "Configuring Your Server Wizard" in Windows Help.

# Common Administrative Tasks in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section provides procedures and conceptual information for completing common administrative tasks. Some of the tasks are new to this version of IIS, others are similar to previous versions, but the procedures have changed. Some tasks remain unchanged, but may be unfamiliar to new administrators.

**In This Section**

This section includes the following information:

- Enabling Web Service Extensions: Describes the state of the default installation and how to enable Web service extensions.
- Creating Web or FTP Sites: Describes how to create Web site configurations and how to install and use FTP services.
- Creating Virtual Directories in IIS 6.0: Describes how to create virtual directories within Web and FTP sites.
- Renaming Virtual Directories: Describes how to rename virtual directories by editing the MetaBase.xml file.
- Configuring Authentication: Describes how to set up and use user authentication methods.
- Creating Application Pools in IIS 6.0: Describes how to group Web applications into application pools.
- Creating and Isolating Applications in IIS 6.0: Describes how to create and manage applications.
- Obtaining and Backing Up SSL Certificates: Describes how to set up and use Secure Sockets Layer (SSL) certification on your sites.
- Backing Up and Restoring the Metabase in IIS 6.0: Describes how to back up your site configurations.
- Redirecting Web Sites in IIS 6.0: Describes how to automatically direct users to the correct page on your site.
- Hosting Multiple Web Sites: Describes how to create and host multiple Web sites.
- Assigning Resources to Applications in IIS 6.0: Describes how to control the amount of resources an application uses.
- Controlling Access to Applications in IIS 6.0: Describes how to secure your applications with permissions and restrictions.
- Enabling ASP.NET: Describes how to take advantage of dynamic content by enabling ASP.NET.
- Enabling ASP Pages in IIS 6.0: Describes how to allow IIS to process requests for Active Server Pages.
- Saving Configurations: Describes how to save metabase and application configurations.
- Starting and Stopping Services: Describes how to start and stop IIS services.
- Configuring Recycling in IIS 6.0: Describes how to periodically restart worker processes assigned to an application pool.
- Administering Servers from the Command Line in IIS 6.0: Describes how to use powerful scripting and programming tools to access and configure settings.
- Administering Servers Remotely in IIS 6.0: Describes the tools to use to administer your sites.
- Enabling Network File System Support: Describes how to allow IIS to access files on a UNIX or Linux-based system.

# Enabling Web Service Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down — which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

### To enable Web service extensions

1. In IIS Manager, click the Web Service Extensions folder.
2. In the details pane, select the Web service extension that you want to enable, and then click **Allow**.
3. To see the properties of a Web service extension, select an extension, and then click **Properties**.

| **Important** |
| --- |
| If you choose to allow all unknown ISAPI extensions or CGI extensions to run on your Web server, your Web server could be susceptible to computer viruses or worms that exploit these technologies. Instead, select only those specific ISAPI extensions or CGI extensions that you want to run on your Web server. |

#### Related Information

- For more information about Web service extensions, see Configuring IIS For Dynamic Content.

# Creating Web or FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS creates a default Web site configuration on your hard disk at the time of installation. You can use the \Inetpub\Wwwroot directory to publish your Web content, or create any directory or virtual directory you choose. The File Transfer Protocol (FTP) service must be installed and started in order to create an FTP site. It is not installed by default.

Creating a Web or FTP site using IIS Manager does not create content, but merely creates a directory structure and configuration files from which to publish the content.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To use the default Web site**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click **Default Web Site**, and then click **Properties**.
2. On the **Web Site** tab, under **Web site description**, type the name of your Web site in the **Description** box.
3. Click **OK**. The name of the new site appears in IIS Manager.

**To create a new Web site**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** folder, point to **New**, and then click **Web Site**. The **Web Site Creation Wizard** appears.
2. Click **Next**.
3. In the **Description** box, type the name of your site, and then click **Next**.
4. Type or click the IP address (the default is **All Unassigned**), TCP port, and host header (for example, *www.mysite.com*) for your site, and then click **Next**.

   | **Important** |
   | --- |
   | To help ensure that user requests reach the correct Web site, configure a unique identity for each site on the server by distinguishing each Web site with at least one of three unique identifiers: a host header name, an IP address, or a TCP port number. |

   | **Note** |
   | --- |
   | Using unique host header names is the preferred way to identify multiple Web sites on a single server. For more information about choosing unique identifiers, see Hosting Multiple Web Sites on a Single Server and Adding Web Sites to Your Server in the IIS 6.0 Operations Guide, which is accessible from IIS Manager. |

5. In the **Path** box, type or browse to the directory that contains, or will contain, the site content, and then click **Next**.
6. Select the check boxes for the Web site access permissions you want to assign to your users, and then click **Next**.
7. Click **Finish**.

To change these and other settings later, right-click the Web site, and click **Properties**.

**To install FTP services**

1. From the **Start** menu, click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. From the **Components** list box, click **Application Server**, and then click **Details**.
5. From the **Subcomponents of Application Server** list box, click **Internet Information Services (IIS)**, and then click **Details**.
6. From the **Subcomponents of Internet Information Services** list box, select the **File Transfer Protocol (FTP) Service** check box.
7. Click **OK** twice.
8. Click **Next**. You might be prompted for the Windows Server 2003 family CD or the network install path.
9. Click **Finish**, and then click **Close**.

| **Note** |
| --- |
| IIS Manager creates a default FTP site during the FTP services installation. You can use the \Inetpub\Ftproot directory to publish your content, or you can create another directory. |

**To create a new FTP site**

1. In IIS Manager, double-click the local computer, right-click the **FTP Sites** folder, point to **New**, and then click **FTP Site**.
2. Click **Next**.
3. In the **Description** box, type the name of your site, and then click **Next**.
4. Type or click the IP address (the default is **All Unassigned**) and TCP port for your site, and then click **Next**.
5. Click the user isolation option you want, and then click **Next**.
6. In the **Path** box, type or browse to the directory that contains or will contain shared content, and click **Next**.
7. Select the check boxes for the FTP site access permissions you want to assign to your users, and then click **Next**.
8. Click **Finish**.

To change these and other settings later, right-click the FTP site, and then click **Properties**.

**Related Information**

- For more information about configuring Web sites, see Web Site Setup.
- For more information about configuring FTP sites, see FTP Site Setup.

# Creating Virtual Directories in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In most cases, the content you publish to your Web or FTP site is located in a root or home directory on your computer, such as C:\Inetpub\Wwwroot\. However, there might be instances when the content is located somewhere else, or even on a remote computer.

To publish from any directory not contained within your home or root directory, you can create a virtual directory. A virtual directory is a directory that is not contained in the home directory but appears to client browsers as though it were.

You can create a virtual directory through IIS Manager or by using Windows Explorer. Additionally, you can create a Web virtual directory by using the Iisvdir.vbs script, and a FTP virtual directory by using the Iisftpdr.vbs script.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a virtual directory by using IIS Manager**

1. In IIS Manager, expand the local computer, expand the **Web Sites** or **FTP Sites** folder, right-click the site or folder within which you want to create the virtual directory, point to **New**, and then click **Virtual Directory**. The **Virtual Directory Creation Wizard** appears.
2. Click **Next**.
3. In the **Alias** box, type a name for the virtual directory. (Choose a short name that is easy to type because the user types this name.)
4. Click **Next**.
5. In the **Path** box, type or browse to the physical directory in which the virtual directory resides, and then click **Next**.
6. Under **Allow the following permissions**, select the check boxes for the access permissions you want to assign to your users, and then click **Next.**

   | **Important** |
   | --- |
   | For security reasons, when selecting access permissions, consider allowing only the default Read permission. By restricting permissions in this way, you can help avoid attacks against your Web site by malicious users. For more information about setting access permissions, see Securing Virtual Directories and Access Control in Help and Support Center for Windows Server 2003. |

7. Click **Finish**. The virtual directory is created below the currently selected folder level.

**To create a virtual directory by using Windows Explorer**

1. Open Windows Explorer.
2. Right-click the folder you want to be a virtual directory, and click **Sharing and Security**.
3. Click the **Web Sharing** tab.
4. Click **Share this folder**.
5. In the **Alias** box, type the name for the virtual directory.
6. Click **OK** twice.

**To create a Web virtual directory by using the Iisvdir.vbs script**

1. From the **Start** menu, click **Run.**
2. In the **Open** box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following:

   **cscript %SystemRoot%\system32\iisvdir.vbs /create** *SampleWebSite*[**/***Path*] *VirtualDirectorydrive***:***path*

   where *SampleWebSite*, *VirtualDirectory*, and *path* is the physical directory, as appropriate.

For more information about creating Web virtual directories by using the Iisvdir.vbs command and its parameters, type the following at the command prompt: **iisvdir /create /?**.

**To create an FTP virtual directory by using the Iisftpdr.vbs script**

1. From the **Start** menu, click **Run.**
2. In the **Open** box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following:

   **cscript %SystemRoot%\system32\iisftpdr.vbs /create** *FTPSite*[**/***Path*] *VirtualDirectory drive***:***path*

   where *FTPSite*, *VirtualDirectory*, and *path* is the physical directory, as appropriate.

For more information about creating FTP virtual directories by using the Iisftpdr.vbs command and its parameters, type the following at the command prompt: **iisftpdr /create /?**.

**Related Information**

- For more information about Iisvdir.vbs, see Creating Web Virtual Directories Using Iisvdir.vbs.
- For more information about Iisftpdr.vbs, see Creating FTP Virtual Directories Using Iisftpdr.vbs.

# Renaming Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In IIS 6.0, you cannot use IIS Manager to rename a virtual directory. Instead, you can rename a virtual directory by editing the MetaBase.xml file using a text editor, such as Microsoft Notepad.

**Note**

Always back up the metabase before editing MetaBase.xml. For information about backing up the metabase, see Backing Up and Restoring the Metabase in IIS 6.0.

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters).

**Procedures**

### To rename a virtual directory by using Notepad

1. Open MetaBase.xml in Notepad. The default location for MetaBase.xml is C:\Windows\System32\Inetsrv.
2. On the **Edit** menu, click **Find**, and then locate the virtual directory name that you want to change.
3. Replace the current virtual directory name with a new name.
4. Close the MetaBase.xml file and when prompted, save the changes.
5. In IIS Manager, right-click the local computer, and then select **All Tasks**.
6. On the **All Tasks** menu, click **Save Configuration to Disk**

**Note**

When you rename a newly created virtual directory, it takes at least two minutes for the change to appear in IIS Manager if no other changes are pending. To see the change immediately in IIS Manager, close and then reopen IIS Manager.

# Configuring Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can require users to provide a valid Microsoft® Windows® user account name and password before they access any information on your server. This identification process is called *authentication*. Authentication, like many of the features in IIS, can be set at the Web site, directory, or file level. IIS provides the following authentication methods to control access to the content on your server.

**Web Methods**

- Anonymous Authentication in IIS 6.0: Gives users access to the public areas of your Web site without prompting them for a user name or password.
- Basic Authentication in IIS 6.0: Requires a previously assigned Windows account user name and password, also known as credentials.
- Digest Authentication in IIS 6.0: Offers the same functionality as Basic authentication, while providing enhanced security in the way that a user's credentials are sent across the network.
- Advanced Digest Authentication in IIS 6.0: Collects user credentials and stores them on the domain controller as an MD5 hash, or message digest.
- Integrated Windows Authentication in IIS 6.0: Collects information through a secure form of authentication where the user name and password are hashed before being sent across the network.
- .NET Passport Authentication in IIS 6.0: Provides a single sign-in service through SSL, HTTP redirects, cookies, JavaScript, and strong symmetric key encryption.

**FTP Methods**

- Configuring Anonymous FTP Authentication: Gives users access to the public areas of your FTP site without prompting them for a user name or password.
- Configuring Basic FTP Authentication: Requires users to log on with a user name and password corresponding to a valid Windows user account.

# Creating Application Pools in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when running in worker process isolation mode. |

With IIS 6.0 running in worker process isolation mode, you can group Web applications into application pools. Application pools allow specific configuration settings to be applied to groups of applications, and the worker processes servicing those applications. Any Web directory or virtual directory can be assigned to an application pool.

By creating new application pools and assigning Web sites and applications to them, you can make your server more efficient and reliable, and your other applications always available, even when the applications in the new application pool terminate.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a new application pool**

1. In IIS Manager, double-click the local computer, right-click **Application Pools**, point to **New**, and then click **Application Pool**.
2. In the **Application pool ID** box, type the name of the new application pool.
3. Under **Application pool settings**, click either **Use default settings for new application pool** or **Use existing application pool as template**.
4. If you selected **Use existing application pool as template**, from the **Application pool name** list box, click the application pool to be used as a template.
5. Click **OK**.

**Related Information**

- For more information about creating application pools, see Configuring Application Pools in IIS 6.0.
- For information about configuring worker processes, see Configuring Application Pool Identity with IIS 6.0.

# Creating and Isolating Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To create an application, you designate a directory as the starting point (application root) for the application. You can then set properties for the application. Each application can have a friendly name; this name appears in IIS Manager and gives you a way to distinguish between applications. The application name is not used anywhere else.

Web sites are root-level applications by default. When you create a Web site, a default application is created at the same time. You can use this root-level application, remove it, or replace it with a new application by removing it and creating a new application.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create an application**

1. In IIS Manager, expand the local computer, right-click the directory that is the application starting point, and then click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab.
3. In the **Application settings** section, click **Create**. If you see the **Remove** button instead of the **Create** button, an application has already been created.
4. In the **Application name** box, type a name for your application.
5. Click **OK**.

Isolating applications means configuring them to run in a process (memory space) separate from the Web server and other applications. You can run IIS 6.0 in one of two isolation modes: worker process isolation mode or IIS 5.0 isolation mode. IIS cannot run in both modes simultaneously.

Worker process isolation mode is the default mode that the server runs in on a clean installation. In this mode, you can isolate an application by adding it to an application pool that includes isolation settings. For more information about creating and configuring application pools, see Configuring Application Pools in IIS 6.0.

| **Note** |
| --- |
| Before you isolate an application, verify in which application isolation mode the computer is running. You can determine the application isolation mode in which IIS is running by the presence (worker process isolation mode) or absence (IIS 5.0 isolation mode) of the Application Pools folder. |

**To isolate an application in worker process isolation mode**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click the application you want to isolate, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab (depending on the application).
3. In the **Application settings** section, in the **Application pool** list box, click an application pool.
4. Click **Apply**, and then click **OK**.

IIS 5.0 isolation mode allows you to run applications in IIS 6.0 that were developed for earlier versions of IIS. Use this mode only if an application has a compatibility issue when it runs in worker process isolation mode and you cannot resolve the problem.

**To isolate an application in IIS 5.0 isolation mode**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click the application you want to isolate, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab (depending on the application). If you are in the directory listed as the **Starting Point** directory, the **Application name** box is already filled in.
3. From the **Execute permissions** list box, click the appropriate process option.
4. Click **OK**. The Web server finishes processing any current requests for the application before creating a separate process. At the next request for the application, the application will run in the appropriate memory space.

| **Note** |
| --- |
| Server-side include (SSI) and Internet Database Connector (IDC) applications cannot be run in a memory space separate from the Web server. |

**Related Information**

- For more information about creating applications, see Creating Applications in IIS 6.0.

# Obtaining and Backing Up SSL Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Secure Sockets Layer (SSL) certificates contain information used in establishing identities over a network, a process called authentication. Similar to conventional forms of identification, certificates enable Web servers and users to authenticate each other before establishing a connection.

*Server certificates* contain information about the server that allows the client to positively identify the server before sharing sensitive information. *Client certificates* contain personal information about the clients requesting access to your site that allow you to positively identify them before allowing them access to the site.

This topic is limited to obtaining, installing and backing up server certificates. For information about obtaining client certificates, see Obtaining Client Certificates in IIS 6.0.

There are two ways to obtain a server certificate. You can issue your own certificate, or you can obtain a certificate from a certification authority.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the runas command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To issue your own server certificate**

1. Use Microsoft Certificate Services 2.0 to create a customizable service for issuing and managing certificates. You can create server certificates for the Internet or for corporate intranets, giving your organization complete control over certificate management policies. For more information, see Microsoft Certificate Services Help.

   -or-

   Use the Web Server Certificate Wizard to request and install your server certificate.

**To obtain a server certificate from a certification authority**

1. Find a certification authority that provides services that meet your business needs and then request a server certificate.

   -or-

   Use the Web Server Certificate Wizard to create a certificate request, which you can send to the certification authority.

2. After the certificate has been processed and returned to you, use the Web Server Certificate Wizard to install the certificate.

It is important to safeguard the certificate and key pair; always back them up to a disk and keep the disk in a secure place.

**To create a backup copy of your server certificate and private key**

1. Locate the correct certificate store. This is typically the Local Computer store in Certificate Manager.
2. If you do not have Certificate Manager installed in Microsoft Management Console (MMC), you need to install it.
3. Right-click the certificate in the Personal store, point to **All Tasks**, and click **Export**.
4. Select **Yes, export the private key**.
5. Follow the wizard default settings, and enter a password for the certificate backup file when prompted.
6. Do not select **Delete the private key if export is successful**, because this will disable your current server certificate.
7. Complete the wizard to export a backup copy of your server certificate.

If you already have Certificate Manager installed in MMC, it points to the correct Local Computer certificate store.

**To add Certificate Manager to MMC**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **mmc**, and then click **OK**. The **Microsoft Management Console** appears.
3. In the **File** menu, click **Add/Remove Snap-in**.
4. On the **Standalone** tab, click **Add**.
5. From the **Available Standalone Snap-ins** list box, click **Certificates**, and click **Add**.
6. Click the **Computer account** option, and click **Next**.
7. Click the **Local computer (the computer this console is running on)** option, and click **Finish**.
8. Click **Close**, and then click **OK**.

**Related Information**

- For more information about obtaining a certificate, see Obtaining and Installing Server Certificates.

# Backing Up and Restoring the Metabase in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS administrators can create backup files using IIS Manager or a programmatic administration script. The backup files are copies of the metabase configuration file (MetaBase.xml) and the matching metabase schema file (MBSchema.xml). Using the metabase configuration backup and restore feature, the metabase can be restored from the backup files.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To create a portable backup (password required)**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Backup/Restore Configuration**.
2. Click **Create Backup**.
3. In the **Configuration backup name** box, type a name for the backup file.
4. Select the **Encrypt backup using password** check box, type a password into the **Password** box, and then type the same password in the **Confirm password** box.
5. Click **OK**, and then click **Close**.

**To create a non-portable backup (password not required)**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Backup/Restore Configuration**.
2. Click **Create Backup**.
3. In the **Configuration backup name** box, type a name for the backup file.
4. Click **OK**, and click **Close**.

**To restore the metabase backup**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Backup/Restore Configuration**.
2. In the **Backups** list box, click the version of the **Automatic Backup** file that you want to restore, and click **Restore**. If prompted for a password, type the password you chose to secure the backup.
3. When a confirmation message appears, click **Yes**.
4. Click **OK**, and then click **Close**.

For more information about backing up and restoring the metabase, see Working with the Metabase in this reference, and Backing Up and Restoring the Metabase in the IIS 6.0 Operations Guide, which is accessible from IIS Manager.

**Related Information**

- Fore information about metabase reliability, see Backing Up and Restoring the Metabase.
- For more information about command-line administration scripts, see Using Command-Line Administration Scripts.

# Redirecting Web Sites in IIS 6.0

Updated: June 29, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When a browser requests a page or program on your Web site, the Web server locates the page identified by the URL and returns it to the browser. When you move a page on your Web site, you can't always correct all of the links that refer to the old URL of the page. To make sure that browsers can find the page at the new URL, you can instruct the Web server to *redirect* the browser to the new URL.

You can redirect requests for files in one directory to a different directory, to a different Web site, or to another file in a different directory. When the browser requests the file at the original URL, the Web server instructs the browser to request the page by using the new URL.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To redirect requests to another Web site or directory**

1. In IIS Manager, expand the local computer, right-click the Web site or directory you want to redirect, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab.
3. Under **The content for this source should come from**, click **A redirection to a URL**.
4. In the **Redirect to** box, type the URL of the destination directory or Web site. For example, to redirect all requests for files in the Catalog directory to the NewCatalog directory, type **/NewCatalog**.

**To redirect all requests to a single file**

1. In IIS Manager, expand the local computer, right-click the Web site or directory you want to redirect, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab.
3. Under **The content for this source should come from**, click **A redirection to a URL**.
4. In the **Redirect to** box, type the URL of the destination file.
5. Select the **The exact URL entered above** check box to prevent the Web server from appending the original file name to the destination URL.

You can use wildcards and redirect variables in the destination URL to precisely control how the original URL is translated into the destination URL.

You can also use the redirect method to redirect all requests for files in a particular directory to a program. Generally, you should pass any parameters from the original URL to the program, which you can do by using redirect variables.

**To redirect requests to a program**

1. In IIS Manager, expand the local computer, right-click the Web site or directory you want to redirect, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab.
3. Under **The content for this source should come from**, click **A redirection to a URL**.

   In the Redirect to box, type the URL of the program, including any redirect variables needed to pass parameters to the program. For example, to redirect all requests for scripts in a Scripts directory to a logging program that records the requested URL and any parameters passed with the URL, type **/Scripts/Logger.exe?URL=***$V***+PARAMS=***$P*. *$V* and *$P* are redirect variables.

4. Select the **The exact URL entered above** check box to prevent the Web server from appending the original file name to the destination URL.

**Related Information**

- For more information about redirect variables, see Redirect Reference.

- For information about how to use the URL Rewrite module to configure URL redirection in IIS 7, see Using the URL Rewrite Module.

# Hosting Multiple Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To create and host multiple Web sites, you must first ensure that each site has a unique identification. To accomplish this, you will need to either obtain multiple IP addresses or assign multiple host header names to a single IP address. Host header names are the friendly names for Web sites, such as www.microsoft.com.

Obtaining and maintaining multiple IP addresses is usually a task reserved for large corporations and Internet service providers (ISPs), while assigning multiple host header names is a fairly simple procedure accomplished through IIS Manager. Your computer or network must be using a name resolution system (typically DNS) in order to use multiple host header names.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To assign a host header name**

1. In IIS Manager, double-click the local computer, right-click the **Web Sites** folder, right-click the Web site, and click **Properties**.
2. On the **Web Site** tab, click **Advanced**.
3. Click **Add**.
4. In the appropriate boxes, type an IP address, TCP port, and host header value for your Web site.

   You can have multiple server certificates, multiple IP addresses, and multiple SSL ports per server.

**Note**

Although Windows Server 2003 added DNS support for UTF-8 characters, IIS 6.0 does not support Unicode characters in host headers. If you use the Web Site Creation Wizard to create a Web site and try to give the site a host header that contains Unicode characters, you will receive an error indicating that you must use only standard characters. If you use **adsutil.vbs** to configure the ServerBindings Metabase Property to use a host header containing Unicode characters, the change will be made in the metabase, but you will not be able to access the Web site.

**Related Information**

- For information about hosting multiple Web sites with multiple IP addresses, see Creating Multiple Sites Using Multiple IP Addresses.
- For more information about using host header names, see Hosting Multiple Web Sites on a Single Server.
- For more information about name resolution, see Domain Name Resolution.

# Assigning Resources to Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are several ways to assign, or limit, the amount of resources an application can use. You can use application property sheets to control performance, cache, and process options of individual applications, or you can use *bandwidth throttling* to limit resources for an entire site or application pool. A third alternative is to control access to your application through the NTFS file system and Web server security features.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To assign resources using property sheets**

1. In IIS Manager, expand the local computer, right-click the Web site or root directory of an application, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab.
3. In the **Application settings** section, click **Configuration**.
4. On the **Mappings** and **Options** tabs, select the appropriate check boxes to enable or disable caching, buffering, and session states.

**To globally assign resources using bandwidth throttling**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** folder, and then click **Properties**.
2. Click the **Performance** tab.
3. Under **Bandwidth** throttling, select **Limit the total network bandwidth available for all Web sites on this server** check box.
4. In the **Maximum bandwidth** box, click the up and down arrows or type the maximum number of kilobytes per second you want each site contained in the directory to use.
5. Under **Web site connections**, select either **Unlimited** or **Connections limited to**:. If you select **Connections limited to**:, click the up and down arrows in the box or type the maximum number of kilobytes per second you want each site contained in the directory to use.

**To assign resources using bandwidth throttling to an individual Web site**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click the Web site to which you want to assign resources, and click **Properties**.
2. Click the **Performance** tab.
3. Under **Bandwidth throttling**, select **Limit the network bandwidth available to this Web site** check box.
4. In the **Maximum bandwidth** box click the up and down arrows or type the maximum number of kilobytes per second you want each site contained in the directory to use.
5. Under **Web site connections**, select either **Unlimited** or **Connections limited to:**. If you select Connections limited to:, click the up and down arrows in the box or type the maximum number of Web site connections.

**Related Information**

- For more information about configuring servers for applications, see About Configuring Servers for Applications.
- For more information about bandwidth throttling, see Throttling Bandwidth.
- For more information about application configurations, see Configuring Applications in IIS 6.0.
- For information about using the NTFS file system and Web server security features, see Controlling Access to Applications in IIS 6.0.

# Controlling Access to Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can control which users and computers are allowed access to your Web server and its resources. You can use both the NTFS file system and Web server security features to assign users specific access rights to directories and files. You can also use IP address restrictions to limit access by specific computers.

- Best Practices for Securing Files with NTFS Permissions: With the NTFS file system, you can limit access to your Web server's files and directories. You can configure the file and directory permissions that set the access level assigned to a particular user account or user group. For example, you can configure your Web server to enable a specific user to view and execute a file, while excluding all other users from accessing that file.

- Securing Sites with Web Site Permissions: Web permissions, unlike NTFS permissions, apply to all users accessing your Web sites. NTFS permissions apply only to a specific user or group of users with a valid Windows account. NTFS controls access to physical directories on your server, whereas Web permissions control access to virtual directories on your Web site.

- Securing Sites with IP Address Restrictions: With IP address restrictions, you can configure your Web server to assign or deny specific computers, groups of computers, or domains access to Web sites, directories, or files. For example, if your intranet server is connected to the Internet, you can prevent Internet users from accessing your Web server by assigning access only to members of your intranet, and explicitly denying access to outside users.

# Enabling ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

**Important**

To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information, see Enabling Web Service Extensions. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0.

Microsoft® ASP.NET is a unified Web development platform that provides the services necessary for developers to build enterprise-class Web applications. While ASP.NET is largely syntax compatible with ASP, it also provides a new programming model and infrastructure for more secure, scalable, and stable applications. You can feel free to augment your existing ASP applications by incrementally adding ASP.NET functionality to them.

ASP.NET is a compiled, .NET Framework-based environment. You can write applications in any language that is compatible with the .NET Framework, including Microsoft Visual Basic® .NET, Microsoft Visual C#®, and Microsoft JScript® .NET. In addition, the entire Microsoft .NET Framework is available to any ASP.NET application. Developers can easily access the benefits of these technologies, such as the managed common language runtime (CLR) environment, type safety, and inheritance.

**Note**

If you use the Configure Your Server wizard to install IIS 6.0, ASP.NET is enabled when the wizard runs. If you use **Add or Remove Programs** in Control Panel to install IIS 6.0, you must ensure that you check **ASP.NET** in the **Subcomponents of Application Server** list to install ASP.NET with IIS 6.0. Then, you must enable ASP.NET Web service extensions so your ASP.NET applications will run correctly.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To enable ASP.NET by using IIS Manager**

1. In IIS Manager, expand the local computer, and then click **Web Service Extensions**.
2. In the details pane, click **ASP.NET**, and then click **Allow**.

**Note**

If your applications require ISAPI, CGI, ASP, WebDAV, or other extensions to operate, you will need to enable those extensions as well.

**Related Information**

- For information about enabling Web service extensions, see Enabling Web Service Extensions.
- For an introduction to ASP concepts, see About ASP.
- For information about developing ASP.NET applications, see the Web Application Guide.

# Enabling ASP Pages in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Microsoft Active Server Pages (ASP) is a server-side scripting environment that you can use to create dynamic and interactive Web pages, and build powerful Web applications. When the server receives a request for an ASP file, it processes server-side script code contained in the file to build the HTML Web page that is sent to the browser. In addition to server-side script code, ASP files can contain HTML (including related client-side scripts), as well as calls to Component Object Model (COM) components that perform a variety of tasks, such as connecting to a database or processing business logic.

**Important**

To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information, see Enabling Web Service Extensions. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0.

**Procedures**

**To enable ASP pages**

1. In IIS Manager, expand the local computer, and then click **Web Service Extensions**.
2. In the details pane, click **Active Server Pages**, and then click **Allow**.

**Note**

If your ASP pages require ISAPI, CGI, Microsoft ASP.NET, WebDAV, or other extensions to operate, you will need to enable those extensions as well.

**Related Information**

- For more information about ASP pages, see Enabling ASP Pages.
- For an introduction to ASP concepts, see About ASP.
- For information about creating ASP pages and developing Web applications, see the Web Application Guide.

# Saving Configurations

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

After you have your site and application up and running the way you want, you can save all or part of the configurations for a backup copy, or for import and export to other sites or computers.

IIS automatically makes a backup copy of the metabase configuration and schema files each time the metabase changes. Administrators can also create backup files on demand, or create backup copies of individual site or application configurations and then export and import them to and from other sites or computers.

Backup files contain only configuration data; they do not include your content (.asp files, .htm files, .dll files, etc.).

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To save a metabase configuration**

1. In IIS Manager, right-click the **local computer**, point to **All Tasks**, and then click **Backup/Restore Configuration**.
2. Click **Create Backup**.
3. In the **Configuration backup name** box, type a name for the backup file.
4. To restore your backup files to a different computer, provide a password that will be used as the encryption key by doing the following:
   - Select the **Encrypt backup using password** check box, type a "strong password" in the **Password** box, and then type the same password in the **Confirm password** box.
5. Click **OK**, and then click **Close**.

**To restore a metabase configuration**

1. In IIS Manager, right-click the **local computer**, point to **All Tasks**, and then click **Backup/Restore Configuration**.
2. From the **Backups** list box, click a previous backup version, and click **Restore**.
3. When a confirmation message appears, click **Yes**.
4. Click **OK**, and then click **Close**.

**To save a site or application configuration**

1. In IIS Manager, right-click the site or application you want to back up, point to **All Tasks**, and click **Save Configuration to a File**.
2. In the **File name** box, type a file name.
3. In the **Path** box, type or browse to the location where you want to save the file.
4. Click **OK**.

**Related Information**

- For information about saving configurations using command-line tools, see Using Command-Line Administration Scripts.

- For information about importing and exporting site and application configurations, see Metabase Import and Export in IIS 6.0.

- For information about strong passwords, see "Strong passwords" in Help and Support Center for Windows Server 2003.

# Starting and Stopping Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You might need to restart IIS before certain configuration changes take effect or when applications become unavailable. You can restart IIS using the **Restart IIS** command in IIS Manager or by using the IISReset command-line utility. Both methods allow you to stop, start, and restart IIS Internet services, as well as reboot your computer.

Restarting or stopping IIS, or rebooting your Web server, is a severe action. When you restart the Internet service, all sessions connected to your Web server (including Internet, FTP, SMTP, and NNTP) are dropped. Any data held in Web applications is lost. All Internet sites are unavailable until Internet services are restarted. For this reason, you should avoid restarting, stopping, or rebooting your server if at all possible. For a list of features designed to improve IIS reliability and remedy the need to restart IIS, see Restarting IIS.

If the **IISReset** command-line utility is disabled, then the command-line or IIS Manager calls that require IISReset.exe will fail and return an error stating that **IISReset** is disabled. Actions that fail include the **Restart IIS** command in IIS Manager and Service Control Manager (SCM) recovery configuration actions that use the **IISReset** command-line utility (for example, the default IIS Admin SCM recovery path). However, SCM recovery actions that do not use the **IISReset** command-line utility continue to function (for example, the default World Wide Web Publishing Service SCM recovery path that restarts the WWW service).

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To restart IIS services**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, then click **Restart IIS**.
2. In the **What do you want IIS to do** drop-down list, click **Restart Internet Services on** *computer name*. You can also choose to reboot the computer, stop the Internet service, or start the Internet service.
3. IIS attempts to stop all services before restarting.

**To start, stop, or pause individual sites**

- In IIS Manager, right-click the site you want to start, stop, or pause; and click **Start**, **Stop**, or **Pause**.

**To restart IIS by using the IISReset command-line utility**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following:

   **Iisreset /noforce** *ComputerName*

   If you are logged on locally, the *ComputerName* parameter is not required. If you are remotely administering a server running IIS, the *ComputerName* parameter is the NetBIOS name of the computer on which you want to restart IIS.

   | **Important** |
   | --- |
   | Use the */noforce* parameter to help prevent data loss in case the IIS services cannot be stopped within the one minute time-out period. If you are certain that it is safe to force IIS to restart, you can omit the */noforce* parameter. However, be aware that you could lose data if you do not include this parameter. |

4. IIS attempts to stop all services before restarting. The **IISReset** command-line utility waits up to one minute for all

services to stop if you do not include the */noforce* parameter. If the services cannot be stopped within one minute, all IIS services are terminated, and IIS restarts.

The **iisreset** command provides several additional parameters. For example, you can adjust the time-out value by using the **iisreset/timeout** command. For more information about the parameters that are available with the **IISReset** command-line utility, open a command-prompt window, and type **iisreset /?** at the command prompt.

For more information about starting and stopping services or using the **IISReset** command-line utility to restart IIS, see Restarting IIS in the IIS 6.0 Operations Guide, which is accessible from IIS Manager, and also see IIS 6.0 Administration Scripts, Tips, and Tricks in this reference.

**Related Information**

- For more information about starting and stopping services, see Restarting IIS.
- For information about using command-line administration, see Using Command-Line Administration Scripts.

# Configuring Recycling in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when running in worker process isolation mode. |

In worker process isolation mode, IIS can be configured to periodically restart worker processes assigned to an application pool, allowing you to recycle faulty Web applications. Recycling keeps problematic applications running smoothly, especially when it is not feasible to modify the application code. This ensures that those application pools remain healthy, and that system resources can be recovered.

You can configure worker processes to restart based on elapsed time, number of requests served, scheduled times, or memory usage; or you can configure them to start on demand.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure a worker process to be recycled after a set elapsed time**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and click **Properties**.
2. On the **Recycling** tab, select the **Recycle worker processes (in minutes)** check box.
3. Click the up or down arrow to set the number of minutes you want to elapse before the worker process is recycled.
4. Click **OK**.

**To configure a worker process to be recycled after a set number of processing requests**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and click **Properties**.
2. On the **Recycling** tab, select the **Recycle worker process (in requests)** check box.
3. Click the up or down arrow to set the number of requests to be processed before the worker process is recycled.
4. Click **OK**.

**To configure a worker process to be recycled at scheduled times**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and click **Properties**.
2. On the **Recycling** tab, select the **Recycle worker processes at the following times** check box.
3. Click **Add** to add a time to the list, **Remove** to delete a time from the list, or **Edit** to change an existing time at which the worker process will be recycled.
4. Click **OK**.

| **Note** |
| --- |
| When recycling is set to occur at scheduled times, it can occur off-schedule if the system time is altered. To avoid unintended changes in scheduled recycling times, recycle the scheduled worker processes right after the system time is changed. |

**To configure a worker process to be recycled after consuming a set amount of memory**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and click **Properties**.
2. On the **Recycling** tab, under **Memory recycling**, select the **Maximum virtual memory (in megabytes)** or **Maximum used memory (in megabytes)** check box.
3. Click the up or down arrow to set memory limits.
4. Click **OK**.

**Related Information**

- For more information about recycling worker processes, see Recycling Worker Processes with IIS 6.0.
- For information about worker process isolation mode, see Worker Process Isolation Mode in IIS 6.0.

# Administering Servers from the Command Line in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Each system on your computer comes with a graphical user interface (GUI), such as IIS Manager, to configure and manage your settings. However, if administrators manage systems frequently, or access large amounts of data on multiple servers, then using each GUI is time consuming. It may be faster to use command-line scripts that can manage single or multiple systems programmatically.

IIS provides powerful scripting and programming tools that you can use to access and configure settings from within a command-line script or compiled application. You can use these tools to create, delete, start, stop, pause, and list various sites and applications, and to copy, import, and export configurations.

You can use command-line scripts for the following:

- Managing Web Sites Using Scripts: Describes how to create, delete, start, stop, pause, and list Web Sites on a local or remote computer.
- Managing FTP Sites Using Scripts: Describes how to create, delete, start, stop, pause, and list FTP Sites on a local or remote computer.
- Managing Web Virtual Directories Using Scripts: Describes how to create, delete, and list Web virtual directories on a local or remote computer.
- Managing FTP Virtual Directories Using Scripts: Describes how to create, delete, and list FTP virtual directories on a local or remote computer.
- Managing IIS Configurations Using Scripts: Describes how to back up, restore, delete, list, copy, import, and export IIS configurations.
- Managing Applications and Web Service Extensions Using Scripts: Describes how to enable and list applications; add and remove application dependencies; enable, disable, and list Web service extensions; and add, remove, enable, disable, and list individual files.

# Administering Servers Remotely in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can administer your server remotely by running IIS on an intranet or the Internet. You can use the following tools for this purpose:

- IIS Manager: You can use IIS Manager on your server to remotely connect to and administer an intranet server running IIS 5.0, IIS 5.1, or IIS 6.0 (IIS 3.0 and IIS 4.0 are not supported).

- Terminal Services: Terminal Services does not require you to install IIS Manager on the remote client computer because, once connected to the server running IIS, you use IIS Manager on the Web server as if you are logged on to the server locally.

- Remote Administration (HTML) Tool: You can use the Remote Administration (HTML) tool to administer your IIS Web server from any Web browser on your intranet. This version of the Remote Administration (HTML) tool is supported only on servers running Windows Server 2003 with IIS 6.0.

**Note**

If you install the Remote Administration (HTML) Tool for IIS 6.0 on a server that has been upgraded to Windows Server 2003 from Microsoft Windows NT 4.0, you may receive an error when you try to view the administrative Web site. This error occurs because, by default, buffering is set to **False** in IIS 4.0. However, in IIS 6.0, buffering is set to **True** by default, and this setting is not changed during upgrade. To view the administrative Web site in IIS 6.0, you must enable buffering.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To administer your intranet server remotely by using IIS Manager**

1. Start IIS Manager on any computer on your network that is running a member of the Windows Server 2003 family.

2. To connect to a remote computer running IIS, right-click the local computer, and then click **Connect**.

3. In the **Connect to Computer**dialog box, in the **Computer name** box, type the name of the computer to which you want to connect, or click **Browse** to browse to the computer, and then click **OK**.

**Note**

If you do not have TCP/IP and a name resolution server, such as Windows Internet Name Service (WINS) installed, you might not be able to connect to an IIS computer using the computer name. As an alternative, you can also use the IP address of the IIS computer, or you can add the host name and IP address to the local hosts file on the computer, which is located at %systemdrive%\Windows\system32\drivers\etc\hosts. For more information, see "TCP/IP" in Help and Support Center for Windows Server 2003.

**To administer your server remotely using Terminal Services**

1. Install the Terminal Services client on the computer you are using to connect.

2. While the remote computer is running, start Terminal Services and identify the name of the remote computer.

3. From the Terminal Services window, administer IIS as you would locally. You can start IIS Manager on any computer on

your network that is running Windows, or you can open a Web-based server-management appliance. You can also run scripts from the Terminal Services window.

**To enable the Remote Administration (HTTP) tool through Control Panel**

1. In **Control Panel**, click **Add or Remove Programs**.
2. In the **Add or Remove Programs** dialog box, click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard** dialog box, in **Components**, click **Application Server**, and then click **Details**.
4. In the **Application Server** dialog box, in **Subcomponents of Application Server**, click **Internet Information Services (IIS)**, and then click **Details**.
5. In the **Internet Information Services (IIS)** dialog box, in**Subcomponents of Internet Information Services (IIS)**, click **World Wide Web Service**, and then click **Details**.
6. In the **World Wide Web Service** dialog box, in **Subcomponents of World Wide Web Service**, select the **Remote Administration (HTML)** check box, and then click **OK**.
7. Click **OK** two more times, and then, in the **Windows Components Wizard** dialog box, click **Next**.
8. After setup is complete, click **Finish** to close the wizard.

**To enable buffering on the administrative Web site (for upgrades from IIS 4.0 to IIS 6.0 only)**

1. In IIS Manager, double-click *servername* (local computer).
2. Double-click **Web Sites**.
3. Right-click the Administration Web site, and then click **Properties**.
4. Click **Home Directory**, click **Configuration**, and then click **Options**.
5. Check the **Enable buffering** check box, and then click **OK** twice.

**To view the Remote Administration (HTML) tool from IIS Manager**

- Expand the local computer, expand the **Web Sites** folder, right-click the **Administration** Web site, and click **Browse**.

**To administer an IIS Web server with the Remote Administration (HTML) tool**

- Open your intranet site from a Web browser and type the following in the address bar: **https://***hostname***:8098**

  Replace *hostname* with the name of the computer that you want to connect to and administer.

For more information about how to manage servers remotely by using the Remote Administration (HTML) tool, see HOW TO: Remotely Administer Internet Information Services in Windows Server 2003.

# Enabling Network File System Support

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If IIS must access files on a UNIX or Linux system, file name case sensitivity can be an issue unless Network File System (NFS) support is enabled in IIS.

Unix and Linux both support mixed-case file names, and IIS fully supports requesting static files in a case-sensitive manner. An issue arises, however, when IIS makes a subsequent request for a file from its static file cache. Because all file names are converted to uppercase letters in the IIS cache, any request after the first request from the IIS static file cache might fail or return the wrong file.

The solution to this issue is to disable the IIS static file cache so that all file requests are issued fresh, thereby retaining correct file name case. The static file cache can be disabled for individual virtual directories on Web sites, or globally for all of IIS.

| **Note** |
| --- |
| Changing this setting has no effect on how ASP files and templates are cached. |

**Procedures**

**To disable the static file cache for a particular Web site virtual directory**

- Edit the metabase and set the **MD_VR_NO_CACHE** property to 1. For example:

  `LM/W3SVC/`*n*`/ROOT/`*vdir*`MD_VR_NO_CACHE=1`

  If you subsequently want to enable the static file cache, set this property to 0.

**To disable the static file cache for all of IIS**

- Edit the registry and add the binary value **DisableStaticFileCache=1** to the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters key.

  If you subsequently want to enable the static file cache, set this value to 0.

# Site Setup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section describes how to complete a first-time basic setup of a Web server, FTP server, SMTP (e-mail) server, and NNTP (news) server. Additionally, publishing information to your Web and FTP sites is discussed.

Internet Information Services (IIS) makes no distinction whether the site you create is an internal-only site (intranet) or an Internet site. Setup, configuration, and administration procedures are the same for both kinds of sites. Differences exist only in the specific configuration information you supply to IIS.

When you install a member of the Microsoft Windows Server™ 2003 family, IIS creates a minimal configuration consisting of a default Web site in a highly secure mode limited to static content. Other IIS services and options, such as FTP, SMTP, and NNTP are not installed by default. This section also describes how to install these services.

**In This Section**

This section includes the following information:

- Web Site Setup: Describes the basics of managing a Web site's infrastructure, from setting a site home directory and default Web Page, to redirecting requests and dynamically altering Web pages.
- FTP Site Setup: Describes installation of the FTP service, and changing default FTP settings globally and for specific FTP sites.
- SMTP Server Setup: Describes installation of the SMTP service and the service's capabilities.
- NNTP Server Setup: Describes installation of the NNTP service and how hosting a news server can be an effective tool for managing discussions.
- Web Authoring with WebDAV: Describes the features of WebDAV and how to set up a WebDAV publishing directory on an IIS server.

# Web Site Setup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Whether your site is on an intranet or the Internet, the principles of providing content are the same. You place your Web files in directories on your server so that users can establish an HTTP connection and view your files with a Web browser. But beyond simply storing files on your server, you must manage how your site is deployed and how your site evolves. Today, an engaging Web site is seldom a static collection of pages. Most successful server administrators are kept busy accommodating ever changing Web content.

This section describes the basics of managing a Web site's infrastructure, from redirecting requests to dynamically altering Web pages.

This section includes the following information:

- Creating a Web Site from a Template: Describes how to create a new Web site from a template file.
- Setting Home Directories: Describes how to set or change the home directory of a Web site using IIS Manager or a command-line administration script.
- Setting Up Default Documents: Describes the purpose of a default document and where it should be placed, as well as file naming restrictions.
- Adding Footers to Web Pages: Describes how to automatically include HTML code in the bottom of Web pages.
- Using Virtual Directories: Describes the concept of virtual directories and how to use them to keep published data compartmentalized on the server while remaining transparent to the user.
- Domain Name Resolution: Describes the concept of name resolution, and how to register a domain name for your site.
- Using FrontPage Server Extensions to Extend Web Sites: Describes the benefits of using FrontPage® Server Extensions from Microsoft for Web site management, and explains how to set up the extensions.
- Hosting Multiple Web Sites on a Single Server: Describes how to effectively consolidate Web sites by hosting them on one server.

# Creating a Web Site from a Template

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you have a standard Web site configuration that you use or plan to use frequently, you can use IIS Manager to create a Web site template. To do this, export the configuration of the site to a file, and then use that file to create new sites. The exported configuration is stored as a metabase XML configuration file; it does not contain content. After you export the configuration file, you can import that file into a different directory on the same Web server or to a different Web server.

You can also export and import sites and application pools as configuration files to quickly migrate sites and application pools between servers. This feature allows you to import Web applications if the developer stored the application in a configuration file.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To create a Web site template file**

1. In IIS Manager, double-click the local computer, right-click the site or application that you want to back up, point to **All Tasks**, and then click **Save Configuration to a File**.
2. In the **Save Configuration to a File** dialog box, in the **File name** box, type a file name.
3. In the **Path** box, type the full path to the configuration file, or click **Browse** to browse to the location where you want to save the file, and then click **OK**.

**To create a new site from a template file**

1. In IIS Manager, double-click the local computer, right-click the **Web Sites** folder, point to **New**, and then click **Web Site (from file)**.
2. In the **Import Configuration** dialog box, in the **File** box, type the full path to the configuration file, or click **Browse** to browse for the file.
3. Click **Read File**.

   The **Select a configuration to import box** is populated.
4. In **Select a configuration to import**, click the configuration that you want to import, and then click **OK**.

# Setting Home Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Each Web site must have a home directory. The default Web site home directory is *LocalDrive*:\Inetpub\Wwwroot. You can change a Web site home directory using IIS Manager or by editing the MetaBase.xml file directly. Both methods are described in this topic.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To change the home directory of a Web site using IIS Manager**

1. In IIS Manager, expand the local computer, expand the **Web Sites** directory, right-click the Web site you wish to change, and click **Stop**.
2. Use Windows Explorer, to rename the *LocalDrive*:\Inetpub\Wwwroot directory to the name of your choice. Alternatively, you can copy the entire \Wwwroot directory tree to a new location.
3. In IIS Manager, right-click your Web site, and click **Properties**.
4. Click the **Home Directory** tab, and under **The content for this resource should come from**, click **A directory located on this computer**, **A share located on another computer**, or **A redirection to a URL**, depending on where your home directory is located.
5. In the **Local path** box, type the path name, share name, or URL of your directory.

   **Note**

   If you select a directory on a network share, you might need to enter a user name and password to access the resource. IUSR_*computername* is the default account used if another account is not specified. If you use an account with administration credentials on the server, clients can gain access to server operations. This seriously jeopardizes the security of your network. For more information about security user rights see, "Security Best Practices" in Help and Support Center for Windows Server 2003.

6. In IIS Manager, expand the **Web Sites** folder, right-click the Web site you just changed, and click **Start**.

**To change the home directory of a Web site by editing the MetaBase.xml file**

1. In IIS Manager, expand the **Web Sites** folder, right-click the Web site you wish to change, and click **Stop**.
2. Use Windows Explorer to rename the \Wwwroot directory to the name of your choice. Alternatively, you can copy the entire \Wwwroot directory tree to a new location and then rename the directory.
3. From the *systemroot*\System32\Inetsrv directory, open the MetaBase.xml file in a text editor.
4. To change the Web home directory, search for "/W3svc/*1*/root" and change the Path= value to the new location of the Web home directory. Substitute the value *1* with the value corresponding to the Web site to change.
5. In IIS Manager, expand the **Web Sites** folder, right-click the Web site you've just changed, and click **Start**.

**Related Information**

- For more information about the structure of the IIS metabase, see Metabase Structure in IIS 6.0.

# Setting Up Default Documents

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can establish a default document for client requests to a Web site that does not specify a document name. Default documents can be a directory's home page or an index page. Before enabling the default document in IIS Manager, you must create the default document and place it in either the root directory or a virtual directory on the appropriate Web site.

You cannot use commas in a default document name (for example, my,file.html). Commas are used only to delimit multiple default documents in the metabase. If you include a comma in the name of a default document, IIS uses the comma to notate two default documents.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To set up or change the default document**

1. In IIS Manager, double-click the local computer, right-click the **Web Sites** folder or an individual Web site folder, and then click **Properties**.

   **Note**

   Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element.

2. Click the **Documents** tab.
3. Select the **Enable default content page** check box.
4. Click **Add** to add a new default document to the list.
5. Click the document you want to remove from the list and click **Remove**.
6. Click a document from the list and click **Move Up** or **Move down** to change the order in which default documents are served to client requests.
7. Click **OK**.

# Adding Footers to Web Pages

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure the Web server to automatically insert an HTML-formatted file to the bottom of every Web page sent out by your Web server (appending footers or scripts to ASP pages is not supported). For example, your file can contain HTML formatting instructions to add a simple text message and logo image to your Web pages. Once enabled, IIS automatically appends the footer file to the designated Web page or file each time the Web page or file is accessed.

| **Note** |
| --- |
| Document footers can reduce Web server performance, especially if a Web page is frequently accessed. |

A footer file should be not be a complete HTML document. It should contain only those HTML tags necessary for formatting the appearance of footer content. For example, a footer file that adds your organization's name to the bottom of every page should consist of text and the HTML tags necessary for formatting the font and color.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

### To add a footer to every Web page

1. Create an HTML footer file and save it on your Web server's hard disk.
2. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click a Web site, directory, or file, and click **Properties**.
3. Click the **Documents** tab, and select the **Enable document footer** check box.
4. In the box below **Enable document footer**, type either the full local path to the footer file or click **Browse** to navigate to it.
5. Click **OK**.

# Using Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A virtual directory is a friendly name, or *alias*, either for a physical directory on your server hard drive that does not reside in the home directory, or for the home directory on another computer. Because an alias is usually shorter than the path of the physical directory, it is more convenient for users to type. The use of aliases is also secure because users do not know where your files are physically located on the server and therefore cannot use that information to modify your files. Aliases also make it easier for you to move directories in your site. Rather than changing the URL for the directory, you change the mapping between the alias and the physical location of the directory.

If your Web site contains files that are located in a directory other than the home directory, or on other computers, you must create virtual directories to include those files in your Web site. To use a directory on another computer, you must specify the directory's Universal Naming Convention (UNC) name, and provide a user name and password for access rights.

If you want to publish from any directory not contained within your home directory, you must create a virtual directory. For example, suppose you are setting up a Web site for your marketing team on your company's intranet. The following table shows the mapping between the physical location of the files and the URL that accesses the files.

| **Physical Location** | **Alias** | **URL** |
|---|---|---|
| C:\Inetpub\wwwroot | *home directory* (none) | http://SampleWebSite |
| \\Server2\SalesData | Customers | http://SampleWebSite/Customers |
| D:\Inetpub\wwwroot\Quotes | None | http://SampleWebSite/Quotes |
| D:\Inetpub\wwwroot\OrderStatus | None | http://SampleWebSite/OrderStatus |
| D:\Marketing\PublicRel | PR | http://SampleWebSite/PR |

Both virtual directories and physical directories (directories without an alias) appear in IIS Manager. A virtual directory is indicated by a gear icon. The following illustration shows the Web site example described above, where /Customers and /PR are virtual directories.




For a simple Web site, you might not need to add virtual directories. You can simply place all of your files in the site's home directory. If you have a complex site, or want to specify different URLs for different parts of your site, you can add virtual directories as needed. To make a virtual directory accessible from multiple sites, you must add the virtual directory to each site.

There are three ways to create or delete a virtual directory:

- Using IIS Manager; either using the Virtual Directory Creation Wizard or by importing a configuration file.

- Using Windows Explorer (this method requires that your hard drive is formatted with the NTFS file system).
- Using the lisvdir.vbs administration script.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a virtual directory using the Virtual Directory Creation Wizard**

1. In IIS Manager, double-click the local computer, double-click the Web site or FTP site to which you want to add a virtual directory, right-click the Web Site or folder within which you want to create the virtual directory, point to **New**, and then click **Virtual Directory**.
2. Click **Next**.
3. In the **Alias** box, type a name for the virtual directory, and then click **Next**. The alias is the name that the user types, and it should be short and easy to type.
4. In the **Directory** box, type or browse to the physical directory in which the virtual directory resides, and then click **Next**.
5. In the **Access Permissions** dialog box, set the access permissions appropriate to your needs.
6. Click **Next**, and then click **Finish**.

**To create a virtual directory by importing a configuration file**

1. In IIS Manager, double-click the Web site or FTP site to which you want to add a virtual directory, right-click the Web Site or folder within which you want to create the virtual directory, point to **New**, then click **Virtual Directory (from file)**.
2. Click **Next**.
3. In the **Import configuration** box, click **Browse** and navigate to the file containing the virtual directory configuration to import.
4. Click **Read File**, click the virtual directory configuration to import, and then click **OK**.

| **Note** |
| --- |
| The following method is available only if your server's hard disk is formatted with the NTFS file system. |

**To create a virtual directory using Windows Explorer**

1. In Windows Explorer, browse to the folder where to want to create a virtual directory for your site, right-click the directory, and click **Sharing and Security**.
2. Click the **Web Sharing** tab, and in the **Share on** list box, click the site for which you are creating a virtual directory, and click **Share this folder**.
3. In the **Alias** box, type an alias name for this virtual directory.
4. Set the **Access permissions** and **Applications permissions** by clicking the appropriate settings.
5. Click **OK**, and then click **OK** again. You can verify that the virtual directory was created by starting IIS Manager and double-clicking the Web site.

| **Important** |
| --- |

| |
|---|
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas/profile /User:***MyComputer\Administrator***cmd** to open a command window with administrator rights and then type **cscript.exe** ScriptName (include the script's full path and any parameters). |

**To create a virtual directory using the Iisvdir.vbs administration script**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cscript iisvdir.vbs /create** "*SampleWebSite*" *VirtualDirectoryName x:\path*. Substitute your Web site name, virtual directory name, and path to the physical directory as appropriate.
3. Click **OK**.

| **Note** |
|---|
| Deleting a virtual directory does not delete the corresponding physical directory or files. |

**To delete a virtual directory using IIS Manager**

1. In IIS Manager, expand the site containing the virtual directory you want to delete, right-click the virtual directory, click **Delete**, and then click **Yes**.

**To delete a virtual directory using Windows Explorer**

1. In Windows Explorer, browse to the folder containing the virtual directory you want to delete, right-click the directory, and click **Sharing and Security**.
2. Click the **Web Sharing** tab, click **Remove**, and then click **Yes**.
3. Click **OK**. You can verify that the virtual directory was deleted by starting IIS Manager and double-clicking the Web site.

| **Note** |
|---|
| The following method does not work for root virtual directories. |

**To delete a virtual directory using the Iisvdir.vbs administration script**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cscript iisvdir.vbs /delete** "*SampleWebSite*" *VirtualDirectoryName*. Substitute your Web site name and virtual directory name as appropriate.

**Related Information**

- For detailed information about the Iisvdir.vbs command-line administration script, search for "Iisvdir.vbs" in Help and Support Center for Windows Server 2003.
- For information about Security for virtual directories, see Securing Virtual Directories and "Access control" in Help and Support Center for Windows Server 2003.

# Domain Name Resolution

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In order for users to reach your site on an intranet or the Internet, you must have a unique IP address that identifies your computer on the network. This address takes the form of a string of four numbers separated by periods (for example, 192.168.1.42). Because a numeric address is difficult for people to remember, text names, or domain names, are used (for example, www.example.com). Before a user can type a domain name into a Web browser and access a Web page, the name has to be assigned to a corresponding IP address. The process of translating a domain name to an IP address is called name resolution.

**Procedures**

**To make your site available from the Internet**

1. Establish a connection to the Internet. Lease a line from an Internet Service Provider (ISP) or a telephone company. Your provider installs a cable at your site that is plugged into the network adapter card in your computer. Usually a router is also installed between your computer and the ISP. When shopping for a provider, be sure to ask what type of connection is provided and whether the provider guarantees a minimum bandwidth. Also, check on the type of security offered.

2. Obtain an IP address. Your ISP gives you an IP address, an Internet gateway (default gateway), and a subnet mask.

3. Verify your computer's network identity. To verify your IP address, at the command prompt, type **ipconfig**, and then press ENTER. Your current IP address, subnet mask, and default gateway are displayed.

   To verify your computer's name, right-click the **My Computer** icon on your desktop, and click **Properties**. The **Network Identification** tab shows your computer name.

4. Register a domain name. Many ISPs can register your domain name for you, or you can register it yourself. There are domain name registrars located around the world, available for the registration of domain names under many top-level domains (such as .com, .net, .org, as well as country/region code domains such as .us for the United States and .uk for the United Kingdom). These registrars charge a yearly fee to register and maintain your domain name. This fee varies from registrar to registrar, as do the services available. Your domain may be hosted by your ISP, the registrar you choose, or advanced users may choose to host it themselves.

**Related Information**

- For more information about name resolution, search for "dns" or "name resolution" in Windows Help.
- For information about assigning identification information to Web sites; see Naming Web Sites.
- For information about adding Web sites, see Hosting Multiple Web Sites on a Single Server.
- For information about obtaining a domain name, see the Accredited Registrar Directory at the InterNIC Web site.

# Using FrontPage Server Extensions to Extend Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

FrontPage Server Extensions from Microsoft provides Web-based and command-line administration for extending virtual servers. Extending virtual servers by using FrontPage Server Extensions enables the site owner to author the site in FrontPage and delegate site ownership and administration credentials.

FrontPage 2002 Server Extensions from Microsoft is an update of FrontPage 2000 Server Extensions, and includes new security features, such as roles and rights, and new features for monitoring server health and Web site usage.

If you build and maintain Web sites for customers and coordinate these efforts with multiple authors, and then the features included in FrontPage 2002 Server Extensions can help you accomplish the following tasks:

- Manage Web sites either on the local server or remotely, by using HTML Administration pages or a quick command-line interface.
- Secure Web sites and grant authoring, browsing, site management, or other user rights to authorized users.
- Analyze site usage to learn who is viewing the site and how often.
- Track errors on the server to help prevent site or server crashes.
- Upgrade FrontPage Web site management functionality with features such as forms, search tools, and usage analysis.

*Extending* a Web site means enabling FrontPage Server Extensions features which improve the management of content development and security of your site. Extending a Web site adds another level of security to the site, and after a Web site is extended, a FrontPage client can open the site and author it in FrontPage. A site that is not extended cannot be opened or authored in FrontPage.

| ⚠**Caution** |
| --- |
| Sites or virtual directories that are configured to use Active Directory isolation or FTP load balancing should not be mapped to physical directories that are used for Web sites that use FrontPage Server Extensions. Doing so could enable users to view any files in that folder structure over the network. |

**Extending Web Sites**

Before you can extend Web sites by using FrontPage Server Extensions, you must create a virtual server for every Web site that you plan to extend. A virtual server differs from a virtual directory. A virtual server, in technical terms, is a virtual machine that resides on an HTTP server but appears to the user as a separate HTTP server. Several virtual servers can reside on one computer, each running its own programs and each with individualized access to input and peripheral devices. Each virtual server has its own domain name and IP address and appears, to the user, to be an individual Web site. Some Internet service providers use virtual servers for clients who want to use their own domain names.

By default, IIS provides a working virtual server named **Default Web Site**. This virtual server points to the content directory *systemdrive*\Inetpub\Wwroot. To create a new virtual server, you must create a content directory.

FrontPage Server Extensions does not support Unicode content directory names or WebDAV. In addition, to remotely administer a Web site that uses SharePoint™ Team Services (STS) from Microsoft, you must assign the _vti_bin virtual directory of new FrontPage Server Extensions virtual servers to the MSSharePointPool application pool. Do not change the MSSharePointPool application pool configuration.

If you have an STS Web site, the following types of pages do not work without excluding the virtual directory that contains the pages from the STS ISAPI filter:

- ASP pages
- ASP.NET pages

- PHP pages
- PerlScript pages

**Procedures**

**To enable these types of pages to work on your STS Web site, do the following**

1. In the **Central Administration** Web site for STS, choose to **Manage** your virtual directory or Web site.
2. Under **Defined Managed Paths**, type the path to the virtual directory that contains your excluded pages (such as <code>\MyASPPages</code>) and select **Exclude**.

Virtual servers should not be nested because virtual servers nested within each other are not visible to FrontPage Server Extensions. For example, if you must create two virtual servers, the first content directory could be located in **systemroot**\Inetpub*Test*; and the second content directory could be located in **systemroot**\Inetpub*Test2*. If the *Test2* directory were located in the *Test* directory, all the permissions in *Test2* would be overwritten by *Test*.

| **Important** |
| --- |
| Never change permissions or Access Control Lists (ACLs) at the content directory level. To maintain a secure environment, always make permissions changes and other security-related changes by using the **FrontPage 2002 Server Administration** tool. |

After creating a virtual server, create the corresponding Web site in IIS.

After creating the Web site in IIS, you are ready to extend the site. Note that a site created in FrontPage and added to IIS is already extended.

Before starting the procedure of extending the Web site, check to make sure you have enabled the FrontPage Server Extensions .dll files. The following .dll files must be installed on your machine before you can extend Web pages.

- *systemdrive*\Program Files\Common Files\Microsoft Shared\Web Server Extensions\50\Isapi_vti_adm**Fpadmdll.dll**
- *systemdrive*\Program Files\Common Files\Microsoft Shared\Web Server Extensions\50\Isapi**Shtml.dll**
- *systemdrive*\Program Files\Common Files\Microsoft Shared\Web Server Extensions\50\Isapi_vti_adm**Admin.dll**
- *systemdrive*\Program Files\Common Files\Microsoft Shared\Web Server Extensions\50\Isapi_vti_aut**Author.dll**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To extend a site using the FrontPage 2002 Server Administration tool**

1. From the **Start** menu, point to **All Programs**, point to **Administrative Tools**, and click **Microsoft SharePoint Administrator**. If IIS returns a 404 error, "The page cannot be found," check to make sure that the FrontPage 2002 Server Extensions optional component is enabled in the **Windows Components Wizard** in the **Add or Remove Programs** item of the Control Panel.
2. On the **FrontPage 2002 Server Administration** site, check to see that the virtual server and corresponding Web site that you created earlier are listed in the **Virtual Servers** section.
3. To extend the Web site, click the **Extend** link next to the virtual server name.
4. Click **Submit**. The **FrontPage 2002 Server Administration** tool adds FrontPage Server Extensions template directories to the content directory of your virtual server, and other files that contain metadata.

**Related Information**

- For more information about FrontPage Server Extensions 2002, see the Microsoft FrontPage Server Extensions Overview on MSDN.
- For more information about how to extend and administering sites by using FrontPage Server Extensions 2002, see the SharePoint Team Services on TechNet.
- For more information about how to administer FrontPage Server Extensions programmatically, see the FrontPageWeb Metabase Property.
- For more information about programmatically extending and administering sites that use FrontPage Server Extensions 2002, see the SharePoint Team Services SDK on MSDN.

# Hosting Multiple Web Sites on a Single Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS supports multiple Web sites on a single server. For example, rather than using three different servers to host three different Web sites, you can install all three Web sites on the same server. Consolidating Web sites saves hardware resources, conserves space, and reduces energy costs.

To ensure that user requests reach the correct Web site, you must configure a unique identity for each site on the server. To do so, you must distinguish each Web site with at least one of three unique identifiers: a host header name, an IP address, or a unique TCP port number.

Web sites hosted on the same server can be distinguished using the following unique identifiers.

- Host header name
- IP address
- TCP port number

Using host header names to uniquely identify sites on a server is recommended for most situations.

This section includes the following information:

- Creating Multiple Sites Using Multiple IP Addresses: Describes how to use unique IP addresses to to distinguish multiple sites.
- Creating Multiple Sites Using Host Header Names: Describes how to use host headers to distinguish multiple sites.
- Creating Multiple Sites Using Ports: Describes how to use nonstandard TCP port numbers to distinguish multiple sites.

# Creating Multiple Sites Using Multiple IP Addresses

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Some server configurations require the use of unique IP addresses to distinguish multiple Web sites on the same server.

IIS enables you to configure Web sites to listen on all available IP addresses, or on an individual IP address. If you use multiple IP addresses to distinguish between different sites on the same server, you must configure IIS to assign each site a unique IP address.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To add a Web site using a unique IP address identifier using the Web Site Creation Wizard**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** directory, point to **New**, and then click **Web Site**.
2. Click **Next**.
3. In the **Description** box, type the name you have selected for the Web site, and click **Next**.
4. In the **Enter the IP address to use for this Web site** box, click the unique IP address reserved for this site.
5. In the **TCP port this Web site should use** box, type the port number used by all sites on the server. This should typically be left at the default of 80.
6. If SSL encryption is enabled on the server, type the SSL port number, and then click **Next**. If SSL encryption is not enabled on the server, the **SSL port** box does not appear.
7. In the **Path** box, type or browse to the path of your Web site home directory.
8. To create a secured or private Web site, clear the **Allow anonymous access to this Web site** check box, and then click **Next**. (Web sites are configured for anonymous access by default.)
9. In the **Web Site Access Permissions** box, set the permissions for the home directory.
10. Click **Next**, and then click **Finish**.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**To add a Web site using a unique IP address identifier with the Iisweb.vbs command-line script**

1. If the new Web site will be in a new directory, create the new directory.
2. From the **Start** menu, click **Run**.
3. In the **Open** box, type **cmd**, and then click **OK**.
4. At the command prompt, switch to the *systemroot*\system32 directory.
5. Type the following at the command prompt and then press ENTER:

**cscript iisweb.vbs /create** *Home Directory*

*Site Description* **" /i** *Unique IP Address* **/b** *Port*

| **Note** |
| --- |
| The Iisweb.vbs script supports many more options for creating a Web site. To see the options available, see Creating Web Sites Using Iisweb.vbs. |

**Related Information**

- For detailed information about the different methods of hosting multiple Web sites on a single server, see Hosting Multiple Web Sites on a Single Server.
- For information about Web site names, see Domain Name Resolution.

# Creating Multiple Sites Using Host Header Names

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Every Web site has a descriptive name, and can support one or more host header names. Organizations that host multiple Web sites on a single server often use host headers because this method enables them to create multiple Web site identities without using a unique IP address for each site.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To add a Web site using a host header identifier using the Web Site Creation Wizard**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** directory, point to **New**, and then click **Web Site**.
2. Click **Next**.
3. In the **Description** box, type the name you have selected for the Web site, and then click **Next**.
4. In the **Enter the IP address to use for this Web site** box, click the IP address used by all sites on the server.
5. In the **TCP port this Web site should use** box, type the port number used by all sites on the server.
6. In the **Host Header for this Web site (Default:None)** box, type the host header name to identify the Web site. The host header name must contain the full name of the site, for example, www.microsoft.com.
7. If SSL encryption is not enabled on the server, the **SSL port** box does not appear. If SSL encryption is enabled on the server, type the SSL port number, and then click **Next**. Note that you cannot use host headers with SSL encryption.
8. In the **Path** box, type or browse to the path of your Web site home directory.
9. To create a secured or private Web site, clear the **Allow anonymous access to this Web site** check box, and click **Next**. (Web sites are configured for anonymous access by default.)
10. In the **Web Site Access Permissions** box, set the permissions for the home directory.
11. Click **Next**, and then click **Finish**.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**To add a Web site using a host header identifier using the Iisweb.vbs command-line administration script**

1. If the new Web site will be in a new directory, create the new directory.
2. From the **Start** menu, click **Run**.
3. In the **Open** box, type **cmd**, and then click **OK**.
4. At the command prompt, switch to the *systemroot*\system32 directory.
5. Type the following at the command prompt and then press ENTER:

**cscript iisweb.vbs /create** *HomeDirectory SiteDescription* **" /d** *Host Header* **/i** *IP Address* **/b** *Port*

| **Note** |
| --- |
| The Iisweb.vbs script supports many more options for creating a Web site. To see the options available, see Creating Web Sites Using Iisweb.vbs. |

**Related Information**

- For detailed information about the different methods of hosting multiple Web sites on a single server, see Hosting Multiple Web Sites on a Single Server.
- For information about Web site names, see Domain Name Resolution.

# Creating Multiple Sites Using Ports

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Although not recommended for production servers, you can use nonstandard TCP port numbers to create unique Web site identifiers for site development and testing purposes. You can do this by creating a Web site identity that consists of a TCP/IP address or host header name, appended with a nonstandard TCP port number. Because TCP/IP endpoints are defined by a combination of IP address (or the corresponding name) and port number, configuring a Web site with a nonstandard port number creates a unique endpoint, or identity, for the site.

If you use a nonstandard TCP port number to identify a Web site, users must know the nonstandard TCP port number assigned to the Web site, and append the port number to the Web site's name or IP address in their Web browser address bar.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To add a Web site using a unique port identifier using the Web Site Creation Wizard**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** directory, point to **New**, and then click **Web Site**.
2. In the Web Site Creation Wizard, click **Next**.
3. In the **Description** box, type the name that you have selected for the Web site, and then click **Next**.
4. In the **Enter the IP address to use for this Web site** box, click the IP address used for all sites on this server.
5. In the **TCP port this Web site should use** box, type the unique port number reserved for this site.
6. Optional: In the **Host Header for this Web site (Default:None)** box, type the host header name to identify the Web site. The host header name must containing the full name of the site, for example, www.microsoft.com.
7. If SSL encryption is enabled on the server, type the SSL port number, and then click **Next**. If SSL encryption is not enabled on the server, the **SSL port** box does not appear.
8. In the **Path** box, type or browse to the path of your Web site home directory.
9. To create a secured or private Web site, clear the **Allow anonymous access to this Web site** check box, and click Next. (Web sites are configured for anonymous access by default.)
10. In the **Web Site Access Permissions** box, set the permissions for the home directory.
11. Click **Next**, and then click **Finish**.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**To add a Web site using a unique port identifier with the Iisweb.vbs command-line script**

1. If the new Web site will be in a new directory, create the new directory.
2. From the **Start** menu, click **Run**.

3. In the **Open** box, type **cmd**, and then click **OK**.
4. At the command prompt, switch to the *systemroot*\system32 directory.
5. Type the following at the command prompt and then press ENTER:

   **cscript iisweb.vbs /create** *Home Directory*

   *Site Description* **" /i** *IP Address* **/b** *Unique Port*

| **Note** |
| --- |
| The Iisweb.vbs script supports many more options for creating a Web site. To see the options available, see Creating Web Sites Using Iisweb.vbs. |

**Related Information**

- For detailed information about the different methods of hosting multiple Web sites on a single server, see Hosting Multiple Web Sites on a Single Server.
- For information about port filtering, see TCP/IP Port Filtering.

# FTP Site Setup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

IIS includes the File Transfer Protocol (FTP) service for publishing and managing files. This version of IIS includes Isolating FTP Users to help administrators (particularly Internet hosting providers) efficiently secure and commercialize FTP services for their customers.

The FTP service is not installed by default. To set up an FTP site, you must first install the FTP service through the Control Panel. Installing the FTP service creates a default FTP site, which you can then customize to your needs using IIS Manager.

**Procedures**
**To install the FTP service**

1. From the **Start** menu, click **Control Panel**, double-click **Add or Remove Programs**, and then click **Add/Remove Windows Components**.
2. From the **Components** list, click **Application Server**, and then click **Details**.
3. From the **Subcomponents of Application Server** list, click **Internet Information Services (IIS)**, and then click **Details**.
4. From the **Subcomponents of Internet Information Services (IIS)** list, select the **File Transfer Protocol (FTP) Service** check box, and then click **OK**.
5. Click **Next**. You might be prompted for the Windows Server 2003 product CD or the network install path.
6. Click **Finish**.

| **Caution** |
| --- |
| Caution FTP sites or virtual directories that are configured to use Active Directory isolation or FTP load balancing should not be mapped to physical directories that are used for Web sites that use FrontPage® Server Extensions from Microsoft. Doing so can allow users to view any files in that folder structure over the network. |

When you install the FTP service, a default FTP directory is created at *LocalDrive*:\Inetpub\Ftproot.

Setting up the FTP service for the first time involves first setting global FTP settings, then settings for the default FTP site, and finally adding the content to the FTP site. IIS uses an inheritance model, which means that settings on higher levels are automatically inherited by lower levels. Settings at lower levels can be edited individually to override inherited settings from the next level up.

If you change a setting at a lower level, then later change a setting at a higher level that conflicts with the lower-level setting, you will be prompted to choose whether you want to change the lower-level setting to match the new higher-level setting.

Before setting up the FTP service, see Securing FTP Sites with IIS 6.0 for clarification on security issues associated with the FTP service.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To globally change FTP settings**

1. In IIS Manager, expand the local computer, right-click the **FTP Sites** folder, and click **Properties**.
2. Change the default settings on the property pages as needed. The changes are inherited by all FTP sites on your computer. For information about individual settings, click **Help**.

**To change settings for the default FTP site**

1. In IIS Manager, expand the local computer, expand the **FTP Sites** folder, right-click **Default FTP site**, and click **Properties**.
2. Change the default settings on the property pages to replace any inherited settings as needed. For information about individual settings, click **Help**.

**To publish content on your FTP site**

- Copy or move your content files into the default FTP publishing directory.

| **Note** |
| --- |
| If your network has a name resolution system, for example, DNS, visitors can type **ftp://** followed by your computer name in the address box of their browsers to reach your site. If your network does not have a name resolution system, visitors must type **ftp://**, followed by the numerical IP address of your computer. |

**Related Information**

- For information about administering or adding an FTP site, see FTP Site Administration.
- For conceptual information describing hosting multiple sites on one computer, see Hosting Multiple Web Sites on a Single Server.
- To add virtual directories to an FTP site, see Using Virtual Directories with FTP Sites.
- To customize the appearance of your FTP site, see Setting FTP Messages and Directory Output Style.
- For information about isolating FTP users and the additional security it can provide, see Isolating FTP Users.

# SMTP Server Setup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The Simple Mail Transfer Protocol (SMTP) service provided by IIS is a simple component for delivering outgoing e-mail messages. Delivery of a message is initiated by transferring the message to a designated SMTP server. Based on the domain name of the recipient e-mail address, the SMTP server initiates communications with a Domain Name System (DNS) server, which looks up and then returns the host name of the destination SMTP server for that domain.

Next, the originating SMTP server communicates with the destination SMTP server directly through Transmission Control Protocol/Internet Protocol (TCP/IP) on port 25. If the user name of the recipient e-mail address matches one of the authorized user accounts on the destination server, the original e-mail message is transferred to that server, waiting for the recipient to pick up the message through a client program.

In the case where the originating SMTP server cannot communicate directly with the destination server, the SMTP service can transfer messages through one or more intermediate relay SMTP servers. A relay server receives the original message and then delivers it to the destination server, or redirects it to another relay server. This process is repeated until the message is delivered or a designated timeout period passes.

The SMTP service is not installed by default. You must install the SMTP service using the Control Panel. Installing the SMTP service creates a default SMTP configuration which you can then customize to your needs using IIS Manager.

Many articles about deploying and configuring the SMTP service are available by searching for "smtp" on the Microsoft Developer Network (MSDN) Web site.

**Procedures**

**To install the SMTP service**

1. From the **Start** menu, click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. From the left pane, click **Add/Remove Windows Components**.
4. From the **Components** list, click **Application Server**, and then click **Details**.
5. From the **Subcomponents of Application Server** list, click **Internet Information Services (IIS)**, and then click **Details**.
6. From the **Subcomponents of Internet Information Services (IIS)** list, select the **SMTP Service** check box.
7. Click **OK**.
8. Click **Next**. You might be prompted for the Windows Server 2003 family CD or the network install path.
9. Click **Finish**.

When you install the SMTP service, a default SMTP server configuration is created with a message store in *LocalDrive*:\Inetpub\Mailroot.

When setting up the SMTP service for the first time, you can configure global settings for a SMTP virtual server, as well as settings for individual components of the virtual server.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To configure global SMTP settings**

1. In IIS Manager, expand the local computer, right-click **Default SMTP Virtual Server**, and click **Properties**.

2. Change the default settings on the property pages as needed. For information about individual settings, click **Help**.

**To configure SMTP virtual server components settings**

1. In IIS Manager, expand the local computer, expand **Default SMTP Virtual Server**, right-click the component you want to configure, and click **Properties**.

2. Change the default settings on the property pages as needed. For information about individual settings, click **Help**.

**Related Information**

- SMTP Administration

# NNTP Server Setup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Using a mail-based mechanism for hosting discussions of interest to a variety of people can become inefficient as the number of discussion participants grows. Sending a separate copy of the message to each of the interested parties consumes large quantities of bandwidth, CPU resources, and disk space among the many recipients and mail servers hosting them.

Significant resources can be saved if popular messages are hosted in a single database. This single database is known as a news server and the messages on such a server are called news items or articles. Using the Network News Transfer Protocol (NNTP) service, you can create your own news server to host discussions. NNTP is the protocol for the distribution, inquiry, posting, and retrieval of news items between a news server and client.

Access to the news server can be open to everyone, or it can require user authentication. The NNTP service allows fine tuning of security policies and is capable of handling a large number of virtual domains and independent server instances on a single computer.

The NNTP service is not installed by default. You must install the NNTP service using the Control Panel. Installing the NNTP service creates a default NNTP configuration, which you can then customize to your needs using IIS Manager.

Many articles about deploying and configuring the NNTP service are available by searching for "nntp" on the Microsoft Developer Network (MSDN) Web site.

**Procedures**

**To install the NNTP service**

1. From the **Start** menu, click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. From the left pane, click **Add/Remove Windows Components**.
4. From the **Components** list, click **Application Server**, and then click **Details**.
5. From the **Subcomponents of Application Server** list, click **Internet Information Services (IIS)**, and then click **Details**.
6. From the **Subcomponents of Internet Information Services (IIS)** list, select the **NNTP Service** check box.
7. Click **OK**.
8. Click **Next**. You might be prompted for the Windows Server 2003 family CD or the network install path.
9. Click **Finish**.

When you install the NNTP service, a default NNTP virtual server configuration is created with a message store in *LocalDrive*:\Inetpub\Nntpfile\Root.

When setting up the NNTP service for the first time, you can configure global settings for an NNTP virtual server, as well as settings for individual components of the virtual server.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To configure global NNTP settings**

1. In IIS Manager, expand the local computer, right-click **Default NNTP Virtual Server**, and click **Properties**.
2. Change the default settings on the property pages as needed. For information about individual settings, click **Help**.

**To configure NNTP virtual server component settings**

1. In IIS Manager, expand the local computer, expand **Default NNTP Virtual Server**, right-click the component you want to configure, and click **Properties**.
2. Change the default settings on the property pages as needed. For information about individual settings, click **Help**.

**Related Information**

- NNTP Administration

# Web Authoring with WebDAV

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Web Distributed Authoring and Versioning (WebDAV) extends the HTTP/1.1 protocol to allow clients to publish, lock, and manage resources on the Web. This section describes how to set up a WebDAV publishing directory on an IIS server, and how clients can connect to the server to edit and manipulate files.

| **Note** |
| --- |
| WebDAV is an optional component of IIS 6.0 and is not installed by default. See Installing IIS for information about installing WebDAV. After WebDAV is installed, it can be enabled through the Web Services Extensions node in IIS Manager. |

This section includes the following information:

- About WebDAV: Describes how WebDAV can create directories from which remote users can publish and manipulate files.
- Creating Publishing Directories: Describes how to set up a WebDAV publishing directory.
- Managing WebDAV Security: Provides tips on how to optimize IIS and Windows Server 2003 family security to secure your WebDAV site.
- Searching WebDAV Directories: Describes how to configure a WebDAV directory so that users can search it for content and document properties.
- WebDAV Redirector_IIS_SP1_Ops: Describes how WebDAV Redirector treats files on the Web like they were a mapped network drive.

| **Note** |
| --- |
| WebDAV publishing is an implementation of the HTTP 1.1 protocol as specified in RFC 2518, and it is supported in IIS versions 5.0, 5.1, and 6.0. |

**Related Information**

- For more information about file and directory security, see "Encrypting File System" (EFS) in Help and Support Center for Windows Server 2003. EFS is a new feature in Windows Server 2003, Standard Edition. Also, see Securing Files with NTFS Permissions.
- For information about HTTP Extensions for WebDAV Distributed Authoring, see RFC 2518 at the Internet Engineering Task Force Web site.

# About WebDAV

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Web Distributed Authoring and Versioning (WebDAV) extends the HTTP/1.1 protocol to allow clients to publish, lock, and manage resources on the Web.

Integrated into IIS, WebDAV allows clients to do the following:

- Manipulate resources in a WebDAV publishing directory on your server. For example, users who have been assigned the correct rights can copy and move files around in a WebDAV directory.
- Modify properties associated with certain resources. For example, a user can write to and retrieve a file's property information.
- Lock and unlock resources so that multiple users can read a file concurrently. However, only one person can modify the file at a time.
- Search the content and properties of files in a WebDAV directory.

Setting up a WebDAV publishing directory on your server is as straightforward as setting up a virtual directory through IIS Manager. After you have set up your publishing directory, users who have been assigned the correct rights can publish documents to the server and manipulate files in the directory. Before you can set up a WebDAV directory, you must install Windows XP Professional or a member of the Windows Server 2003 family.

**Important**

To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0.

**WebDAV Clients**

You can access and publish to a WebDAV directory through one of the following Microsoft products or through any other client that supports the industry standard WebDAV protocol. For the specific procedure on how to access and publish through these Microsoft products, consult the specific product's Help.

- Windows clients (Windows 2000 and Windows XP): Connect to a WebDAV directory by adding the directory to the list of Network Places and display the contents as if it were part of the same file system on your local computer. Once connected, you can drag and drop files, retrieve and modify file properties, and complete many other file-system tasks. You can also connect using the command-line client (known as WebDAV Redirector). This client allows you to use existing applications across the Web and share files through firewalls and proxy servers.
- Internet Explorer (versions 5.0 and 6.0): Connect to a WebDAV directory by opening the target directory as a Web folder and complete the same file-system tasks as Windows clients.
- Microsoft Office products (Office 2000 and Office XP): Create, publish, edit, and save documents directly into a WebDAV directory through any application in Office 2000 or Office XP.

**Note**

Even if users connect from behind a firewall, they can still publish on a WebDAV directory if they have the correct permissions and if the firewall is configured to allow publishing.

When enabling WebDAV publishing on your intranet, ensure that all WebDAV clients are running the WebClient service.

**Procedures**

**To check the status, or to enable the WebClient service on a WebDAV client machine**

1. From the **Start** menu, point to **Administrative Tools**, and click **Computer Management**.
2. In the details pane, double-click **Services and Applications**.
3. Double-click **Services**.
4. Scroll down, right-click **WebClient**, and click **Properties**.
5. In the Startup type list box, click **Automatic**.
6. Click **Apply**.
7. In the **Service status** section, click **Start**.
8. Click **OK**.

**Searching in WebDAV**

Once connected to a WebDAV directory, you can quickly search the files on that directory for content as well as properties. For example, you can search for all files that contain the word *table* or for all files written by *Fred*.

**Integrated Security**

WebDAV is integrated with the Windows Server 2003 family and IIS, which means WebDAV takes advantage of the security features offered by the platform and the Web server, including permissions control and discretionary access control lists (DACLs) in the NTFS file system. For information about IIS security, see Security in IIS 6.0.

Clients with proper user rights can write to a WebDAV directory, so it is vital that you control who accesses your directory. IIS has reinforced Integrated Windows Authentication by building in support for the Kerberos V5 security protocol. (Note that Integrated Windows authentication and Kerberos V5 are not the same thing. Integrated Windows authentication now supports Kerberos V5.) By selecting Integrated Windows authentication, you can make sure that only clients with the correct user rights can access and write to the WebDAV directory on your intranet. For information about how the Kerberos V5 protocol works, see "Kerberos V5 protocol" in Windows Server 2003 family Help.

In addition, IIS supports Digest authentication and Advanced Digest authentication. Created for Windows domain servers, Digest and Advanced Digest authentication offer tighter security for passwords and for transmitting information across the Internet. For information about Digest authentication, see Digest Authentication in IIS 6.0. For information about Advanced Digest authentication, see Advanced Digest Authentication in IIS 6.0.

**Related Information**

- For more information about file and directory security, see "Encrypting File System" in Help and Support Center for Windows 2003. EFS is a new feature in Windows Server 2003 family.

# Creating Publishing Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Before setting up your WebDAV publishing directory, ensure that your publishing directory resides in an NTFS partition and be sure the WebDAV extension is enabled in IIS Manager.

WebDAV publishing and file management requires the following permissions on the *NTFS directory*:

- **Read**: Enables users to read the contents of a file.
- **Read and Execute**: Enables user to read a file and run scripts or Common Gateway Interfaces (CGIs).
- **List**: Enables users to view the contents of the directory.
- **Write**: Enables users to access and change the source of a script and to publish files.
- **Modify**: Enables users to rename or delete a directory or file.

| **Note** |
| --- |
| Assigning Write access does not give clients the ability to modify Active Server Pages (ASP) or any other script-mapped files. To allow these files to be modified, you must assign Write permission and Script source access after creating the virtual directory. For information about setting these permissions, see Securing Sites with Web Site Permissions. |

WebDAV publishing and file management requires the following permissions on the *IIS virtual directory*:

- **Read**: Enables users to read the contents of a file.
- **Directory Browsing**: Enables users to view the contents of the directory.
- **Write**: Enables users to access and change the source of a script and to publish files.
- **Indexing** (optional): Enables users to search a directory.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set up a publishing directory**

1. Create a directory on your Web server and set the desired NTFS permissions. You can create this directory anywhere on your Web server.
2. In IIS Manager, create a virtual directory and set the desired virtual directory permissions. For instructions, see Using Virtual Directories.
3. Type **WebDAV** as the alias for this virtual directory, and link it to the physical directory you created in Step 1.

After you finish setting up a WebDAV virtual directory, you can allow clients to publish to it. For information about how users can connect to the directory through any of the Microsoft WebDAV clients, see About WebDAV.

**Related Information**

- For more information about NTFS and IIS user rights, see Access Control with IIS 6.0.

# Managing WebDAV Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This topic provides a brief overview of recommended security practices for remote publishing. It briefly describes how to protect your server and content by authenticating client connections to your server and by controlling access to content on your server. Included in this topic are descriptions of the following:

- Authenticating Clients
- Controlling Access

| **Note** |
| --- |
| For security reasons and to to enable DAV custom properties, ensure that your publishing directory resides on an NTFS partition. To learn more about NTFS partitions, see Windows Server 2003 family Help. |

**Authenticating Clients**

The best way to configure a WebDAV directory depends on the kind of publishing that you want to do. When you create a virtual directory through IIS, Anonymous and Integrated Windows authentication are both turned on. Although this default configuration works well for clients connecting to your server, reading content on a Web page, and running scripts, it does not work well with clients publishing to a directory and manipulating files in that directory.

IIS offers the following authentication methods:

- Kerberos is the primary security protocol for authentication within a domain. Kerberos is the best option for WebDAV client authentication and file security.
- Anonymous authentication grants anyone access to the directory. You should turn off anonymous access to your WebDAV directory. Without controlling who has access, your directory could be attacked by unknown clients.
- Basic authentication sends passwords over the connection in clear text. Clear text passwords can be intercepted and read. Turn on Basic authentication only if you encrypt passwords by using Secure Sockets Layer.
- Digest authentication is a good choice for publishing information on a server that is accessed over the Internet and through firewalls because the passwords are sent over the network as an MD5 hash. However, passwords are stored as plain text in Active Directory.
- Advanced Digest authentication is an improvement over Digest authentication because in addition to sending passwords over the network as an MD5 hash, the passwords are also stored in Active Directory as an MD5 hash rather than plaintext. This makes Advanced Digest the best choice for publishing information on a server that is accessed over the Internet and through firewalls.
- Integrated Windows authentication works best when you are setting up a WebDAV directory on an intranet.
- .NET Passport authentication uses cookies to validate user credentials.

**Controlling Access**

This section describes how you can control access to your WebDAV directory by coordinating IIS and Windows Server 2003 permissions, and how you can protect your script files.

**Configuring Web Permissions**

The following are various ways to configure Web permissions based on the purpose of the material you are publishing:

- Read, Write, and directory browsing enabled: Turning on these permissions allows clients to see a list of resources, modify them (except for those resources without Write permission), publish their own resources, and manipulate files.

- Write enabled; and Read and directory browsing disabled: If you want clients to publish private information on the directory, but do not want others to see what has been published, set Write permission and do not set Read or directory browsing permission. This configuration works well if clients are submitting ballots or performance reviews.

- Read and Write enabled; and directory browsing disabled: Set this configuration if you want to rely on obscuring file names as a security method. However, be aware that security by obscurity is a low-level security precaution because an attacker could guess file names by trial and error.

- Index this resource enabled: Be sure to enable Indexing Service if you plan to let clients search directory resources.

**Controlling Access with DACLs**

WebDAV takes advantage of the security features offered by the platform and the Web server, including permissions control and discretionary access control lists (DACLs) in the NTFS file system. When setting up a WebDAV publishing directory on an NTFS file system drive, make sure the **Everyone** group has Read permission only. Then assign Write permission to specific individuals or groups.

**Protecting Script Code**

If you have script files in your publishing directory that you do not want to expose to clients, you can deny access to these files by verifying that **Script source access permission** is not assigned. Executable files are treated as static HTML files unless **Scripts and Executables** is enabled for the directory.

To prevent .exe files from being downloaded and viewed as HTML files, but to allow .exe files to run, on the **Virtual Directory** property sheet of the publishing directory, change the Execute Permissions to **Scripts and Executables**.

This level of permission makes all executable files subject to the **Script source access** setting. When **Script source access** is selected, clients with Read permission can see all executables; and clients with Write permission can edit them, as well as run them.

With the following permissions, clients can write to an executable file that does not appear in the Application Mapping:

- Write permission is assigned.

- Execute Permissions is set to **Scripts only**.

With the following permissions, clients can write to any executable file, regardless of whether it appears in the Application Mapping:

- **Script source access** is assigned.

- Execute Permissions is set to **Scripts and Executables**.

**Related Information**

- For more information about Web permissions, see Securing Sites with Web Site Permissions.

- For more information about Kerberos authentication, see Help and Support Center for Windows Server 2003.

# Searching WebDAV Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

After you create a WebDAV publishing directory, you may want to allow users to search for content and file properties.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set up your publishing directory for searching**

1. If you want to let clients search for resource properties, make sure the directory is on a drive formatted with the NTFS file system. If you put the directory on a drive formatted with the file allocation table (FAT) file system, clients can search for resource content but cannot search for resource properties.
2. Right-click the **My Computer** icon on your server's desktop, and then click **Manage**.
3. In the navigation pane of the Computer Management console, expand the Services and Applications node, and then select **Services**.
4. In the details pane, right-click **Indexing Service**, and then select **Properties**.
5. On the **General** tab of the Indexing Service Properties window, set the **Startup type** to **Automatic**, and then click the **Start** button.
6. After the Indexing Service starts, click **OK**, and then close the Computer Management Window.
7. In IIS Manager, expand the local machine, right-click your WebDAV directory, and click **Properties**.
8. Click the **Virtual Directory** tab and make sure **Index this resource** and **Read** access options are selected. If **Index this resource** is not selected, Indexing Service will not create a catalog for that directory, and no one will be able to search it. If **Read** access has not been selected, a client can search the directory, but will not be able to see the results of the search.

**Related Information**

- For information about how to enable the Indexing Service in the Windows Component Wizard, see Installing IIS.
- For details about setting IIS permissions for a virtual directory, see Securing Sites with Web Site Permissions.

# WebDAV Redirector_IIS_SP1_Ops

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

WebDAV Redirector is a remote file system over the WebDAV protocol that allows Windows XP client machines to connect to your WebDAV publishing directory through the command line. Once connected, all existing Win32 applications (without any modification) can create, modify, and save files on Web shares (IIS virtual directories). The WebDAV Redirector enables you to manipulate files on the Web as though the files exist on a mapped network drive.

Your WebDAV Web site (the site that contains the virtual directory mapped to your WebDAV publishing directory) must be set up exclusively as a WebDAV site. Therefore, you must delete all FrontPage custom headers. If your WebDAV Web site is set up as both a DAV and FrontPage site, IIS defaults to the FrontPage custom header and blocks DAV command-line connections.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable WebDAV Redirector on your Web server**

1. In IIS Manager, expand the local computer, right-click your WebDAV Web site (the site that contains the virtual directory mapped to your WebDAV publishing directory) and click **Properties**.
2. Click the **HTTP Headers** tab.
3. Remove any headers listed in the **Custom HTTP headers** list box.
4. Click **Apply**.
5. Click **OK**.

After you have set up your WebDAV Publishing Directory, you can connect to that directory through the command prompt using Windows XP clients. Once connected, all existing Win32 applications (without any modification) can create, modify, and save files on Web shares (IIS virtual directories).

# Server Administration Guide

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section describes how to administer your server and Web sites by using the many tools and features developed in Internet Information Services (IIS) 6.0 to improve server management, performance, reliability, and security. This section also includes conceptual information regarding some of the key changes to the IIS architecture and to the metabase.

| **Important** |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting. |

**In This Section**

This section includes the following information:

- IIS Metabase: Describes the IIS 6.0 Extensible Markup Language (XML) metabase, a hierarchical repository for most IIS configuration values.
- Web Site Administration: Describes how to manage Web site infrastructure, including how to manage multiple Web sites, redirect requests, and restart IIS.
- FTP Site Administration: Describes how to manage File Transfer Protocol (FTP) site infrastructure, including multiple FTP sites, data transmission, and user isolation.
- NNTP Administration: Describes how to manage Network News Transfer Protocol (NNTP) site infrastructure, including restricting access to newsgroups, configuring expiration policies, and NNTP security.
- SMTP Administration: Describes how to manage Simple Mail Transfer Protocol (SMTP) site infrastructure, including setting connections, protocol logging, and SMTP security.
- Configuring Applications: Describes how to configure IIS Manager so that various applications -- including Microsoft® Active Server Pages (ASP), Internet Service API (ISAPI), and Common Gateway Interfaces (CGI) -- run on your IIS server.
- Configuring Servers for Applications: Describes how to enable IIS to run in either IIS 5.0 application isolation mode or worker process isolation mode. Also describes how to enable dynamic content in IIS 6.0.
- Using Command-Line Administration Scripts: Describes how to use scripts to perform server administration tasks.
- Administering IIS Programmatically: Introduces the three programmatic interfaces that are supported by IIS 6.0.

# IIS Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The metabase is a repository for most Internet Information Services (IIS) configuration values. The metabase is a plaintext XML file and can be edited manually or programmatically. The metabase is also extensible in a highly efficient manner. As your IIS deployment grows, so does the metabase; however, by using an inheritance model, explicit declarations of duplicate values is avoided, reducing overhead when configuration values need to be read from the metabase.

This section includes the following information:

- Enabling Edit-While-Running in IIS 6.0: Explains how to enable metabase edit-while-running by using IIS Manager and the command line.
- Backing Up and Restoring the Metabase in IIS 6.0: Discusses backing up and restoring the metabase.
- Metabase Import and Export in IIS 6.0: Describes site cloning, application cloning, and creating a metabase template to import from other servers running IIS.
- Metabase Auditing: Describes enabling and disabling the feature that audits changes to IIS metabase settings.

**See Also**
**Concepts**
Metabase Property Reference
Working with the Metabase

# Enabling Edit-While-Running in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Enabling metabase edit-while-running requires that the metabase history feature is enabled and set to a value of at least **10**. The second procedure below requires IIS to be restarted. Restarting IIS stops all sites for a short time. To enable direct metabase edit without interruption of service, use the first procedure below.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable metabase edit-while-running using IIS Manager**

1. In IIS Manager, right-click the local computer, and then click **Properties**.
2. Select the **Enable Direct Metabase Edit** check box.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**To enable metabase edit-while-running from the command line**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. At the command prompt, type **iisreset /stop**, and press ENTER. Leave the command window open.

Open the MetaBase.xml file with a plain text editor. By default, the MetaBase.xml file is located in the systemroot\system32\inetsrv directory.

| **Important** |
| --- |
| All text within MetaBase.xml is case sensitive. |

1. In the IIsComputer node, change the value of the **EnableEditWhileRunning** property from **0** (**FALSE**) to **1** (**TRUE**). The change should look as follows:

```
<IIsComputer Location ="/LM"

EnableEditWhileRunning="1"

EnableHistory="1"

MaxBandwidth="4294967295"

MaxHistoryFiles="10">
```

2. Save changes to the MetaBase.xml file.

3. At the command prompt, type **iisreset /start**, and press ENTER.

**Related Information**

- For information about enabling metabase history, see Configuring the Metabase History Feature in IIS 6.0.
- For more information about the **EnableEditWhileRunning** property, see EnableEditWhileRunning Metabase Property.
- For information about restarting IIS, see Restarting IIS.

# Backing Up and Restoring the IIS Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS administrators can create backup files using IIS Manager or a programmatic administration script. The backup files are copies of the metabase configuration file (MetaBase.xml) and the matching metabase schema file (MBSchema.xml). Using the metabase configuration backup and restore feature, the metabase can be restored from the backup files.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a portable backup (password required)**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Backup/Restore Configuration**.
2. Click **Create Backup**.
3. In the **Configuration backup name** box, type a name for the backup file.
4. Select the **Encrypt backup using password** check box, type a password into the **Password** box, and then type the same password in the **Confirm password** box.
5. Click **OK**, and then click **Close**.

**To create a non-portable backup (password not required)**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Backup/Restore Configuration**.
2. Click **Create Backup**.
3. In the **Configuration backup name** box, type a name for the backup file.
4. Click **OK**, and click **Close**.

**To restore the metabase backup**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Backup/Restore Configuration**.
2. In the **Backups** list box, click the version of the **Automatic Backup** file that you want to restore, and click **Restore**. If prompted for a password, type the password you chose to secure the backup.

**Related Information**

- For information about metabase reliability, see Backing Up and Restoring the Metabase in IIS 6.0
- For more information about command-line administration scripts, see Using Command-Line Administration Scripts.

# Configuring the Metabase History Feature in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, the metabase history feature is enabled and the number of history file pairs (MetaBase.xml and MBSchema.xml) that are saved within the history folder is defined by the default value of the MaxHistoryFiles Metabase Property. If the EnableHistory Metabase Property is not written in the MetaBase.xml file, then IIS interprets the absence of the property to mean that the history feature is enabled and the MaxHistoryFiles Metabase Property is set to the default value. For information about the default values of properties, see the Metabase Property Reference.

**Procedures**

**The procedure title**

1. Open the MetaBase.xml file by using a text editor (MetaBase.xml is located in the systemroot\system32\inetsrv directory), and then navigate to the /LM level.
2. If the EnableHistory Metabase Property and MaxHistoryFiles Metabase Propertyare not in the metabase, add them at the /LM level.
3. Set the value of the **EnableHistory** property to **0** (zero).
4. Save the MetaBase.xml file.

**To enable the metabase history feature**

1. Open the MetaBase.xml file by using a text editor (MetaBase.xml is located in the systemroot\system32\inetsrv directory), and then navigate to the /LM level.
2. Set the value of the EnableHistory Metabase Property property to **1** or remove the **EnableHistory** property from MetaBase.xml.
3. Save the MetaBase.xml file.

**To change the number of history files that are stored in the history folder**

1. Open the MetaBase.xml file by using a text editor (MetaBase.xml is located in the systemroot\system32\inetsrv directory), and then navigate to the /LM level.
2. Change the value of the MaxHistoryFiles Metabase Property.
3. Save the MetaBase.xml file.

| ⚠Caution |
| --- |
| Disabling history files can have serious consequences. Without history files, you might have to reconfigure the IIS server from an empty state. Also, you should ensure that the history files are protected by access control lists (ACLs). By default, only members of the Administrators group can view the history files. For information about configuring ACLs, see "Access control" in Help and Support Center for Windows Server 2003. |

**Related Information**

- For more information about the metabase history feature and how it is used, see The Metabase History Feature.

# Restoring the Metabase with History Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS periodically creates history files consisting of copies of the metabase configuration file (MetaBase.xml) and the matching metabase schema file (MBSchema.xml). Restoring the metabase from history files is similar to restoring the metabase from backup files.

Use the first procedure below to restore the metabase from a history file using IIS Manager. If you want to restore the metabase without using IIS Manager, use the second procedure below.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To restore the metabase using IIS Manager**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Backup/Restore Configuration**.
2. In the **Backups** list, click the **Automatic Backup** you want to restore, and then click **Restore**.
3. Read the message that appears and click **Yes** if you want to continue.
4. Click **OK**.

**To manually restore the metabase**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and then click **OK**.
3. Stop the IIS services by typing **iisreset /stop** at the command prompt, and then press **ENTER**. The command prompt displays the following:

   ```
   Attempting stop...

   Internet services successfully stopped
   ```

   If the IIS services do not stop, and it is possible to restart the computer, type **iisreset /stop /rebootonerror**, and then press ENTER.
4. Make copies of your current metabase files. At the command prompt, type **cd %systemroot%\System32\Inetsrv**, and then press ENTER to navigate to the folder where the files are. Next, type **copy MBSchema.xml MBSchema.old**, and then press ENTER to copy the schema file. Finally, type **copy MetaBase.xml Metabase.old**, and then press ENTER to copy the metabase configuration file.
5. Locate the metabase history files. At the command prompt, type **cd %systemroot%\System32\Inetsrv\History**, and then press ENTER.
6. Select the version of history files that you want. At the command prompt, type **dir**, and then press ENTER. All of the files appear in the following format:

   ```
   <date> <time> <file size> MBSchema_<identifying number>_0000000000.xml

   <date> <time> <file size> MetaBase_<identifying number>_0000000000.xml
   ```

   When you restore the IIS metabase to a previous version, you will need matching metabase schema and configuration files. Select matching files by their identifying number.

7. Copy the chosen history schema file over the existing metabase schema. At the command prompt, type **copy /y MBSchema_<identifying number>_0000000000.xml %systemroot%\System32\Inetsrv\MBSchema.xml**, and then press ENTER. With this command, you copy the file and rename it at the same time.

8. Copy the matching history configuration file over the existing metabase configuration. At the command prompt, type **copy /y Metabase_<identifying number>_0000000000.xml %systemroot%\System32\Inetsrv\MetaBase.xml**, and then press ENTER. Again, this command copies the file and renames it at the same time.

| **Important** |
| --- |
| If you mix up the names or copy a schema and configuration file that do not match, the metabase restoration will fail when you start the IIS services. |

9. Restart IIS. At the command prompt, type **iisreset /start**, and then press ENTER. The command prompt displays the following:

```
Attempting start...

Internet services successfully started
```

If the IIS services do not restart, repeat these steps beginning at step 4, and copy a different set of history files over the current metabase files.

**Related Information**

- For more information about metabase history files, see Backing Up and Restoring the Metabase in IIS 6.0.
- For more information about metabase key names and paths, property inheritance, security, and reliability, see Metabase Structure in IIS 6.0.

# Backing Up the Metabase with Backup Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Using the configuration backup/restore feature, you can generate complete backup copies of the IIS metabase. You can create backup files using either IIS Manager or a command-line administration script. For information about backing up and restoring the metabase using scripts, see Backing Up IIS Configurations Using Iisback.vbs.

| **Note** |
| --- |
| The metabase is locked while the backup is in progress. |

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:** *Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

### To back up the IIS metabase

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Backup/Restore Configuration**.
2. Click **Create Backup** and type a name for the backup file in the **Configuration backup name** box. Note that the backup name is not case-sensitive.
3. Optional: If you want to create a secure backup (for portability between computers), select the **Encrypt backup using password** check box, type a password in the **Password** box, and then type the same password in the **Confirm password** box.
4. Click **OK**.

**Related Information**

- For information about backing up and restoring the IIS metabase using a command-line administration script, see Managing IIS Configurations Using Scripts
- For more information about backing up the metabase, see Backing Up and Restoring the Metabase in IIS 6.0.
- For more information about the initial metabase configuration backup, see Backing Up IIS Configurations Using Iisback.vbs.

# Restoring the Metabase with Backup Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS administrators can restore backup files using either IIS Manager or a command-line administration script. You can restore a metabase backup to the same computer on which the backup was made, or onto a different computer. For information about restoring the metabase using scripts, see Restoring IIS Configurations Using Iisback.vbs.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To restore an IIS metabase backup**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Backup/Restore Configuration**.
2. In the **Backups** list, click the backup you want to restore and click **Restore**.

   | **Note** |
   | --- |
   | If you want to restore IIS to its initial configuration, restore the backup named **Initial Backup** (created automatically by IIS setup). |

3. Read the message that appears and click **Yes** if you want to continue.
4. If you are restoring a secure backup, you are prompted for the password you typed when the backup was created. Type the password and click **OK**. Note that passwords are case sensitive.
5. Click **OK**.

| **Important** |
| --- |
| Before restoring a metabase from one computer to a different computer, you must first delete the computer-specific information from the metabase file. See "Related Information" below. |

**Related Information**

- For information about backing up and restoring the IIS metabase using a command-line administration script, see Managing IIS Configurations Using Scripts.
- For more information about restoring the metabase using backup files, see Backing Up and Restoring the Metabase in IIS 6.0.
- For information about removing computer-specific information from a metabase copy, see Machine-Specific and System-Specific Information.

# Metabase Import and Export in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can copy elements of a metabase in order to clone applications or Web sites, or to produce a metabase template. These copied elements may be moved to another computer or added back to the metabase of the computer from which the elements were copied.

This section includes the following operational information:

- Exporting Metabase Elements with IIS Manager: Explains how to create export files for metabase elements using IIS Manager. Export files must be created before the import feature can be used.
- Importing Metabase Elements with IIS Manager: Explains how to import metabase element files using IIS Manager.
- Moving Sites and Applications to Another Computer in IIS 6.0: Discusses the procedure for creating a template of an import file and special considerations when importing and exporting sites and applications.

# Exporting Metabase Elements with IIS Manager

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the metabase import/export feature to create an export file, containing specifically selected elements of a metabase configuration file. This export file can later be imported into the metabase configuration file on the same computer or on another computer running a member of the Microsoft Windows Server™ 2003 family. To create an export file using a command-line script, see Exporting IIS Configurations Using Iiscnfg.vbs.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To export metabase elements**

1. In IIS Manager, double-click the **Local Computer**, right-click the application pool, Web site, FTP site, or virtual directory that you want to export, point to **All Tasks**, and then click **Save Configuration to a File**.
2. In the **File name** box, type the name that you want to give to your export file.
3. In the **Path** box, type or browse to the path for where you want the export file to be saved. For security purposes, select a folder on which you have already restricted access by setting strict Access Control with IIS 6.0 permissions. Any files created in such a folder inherit the access control list (ACL).
4. If you want to encrypt your export file for increased security, select the **Encrypt configuration using password** check box, and type a strong password in the **Password** and **Confirm password** boxes.
5. Click **OK** and verify that your export file has been created.

**Related Information**

- For more information about exporting an IIS configuration, see Exporting IIS Configurations Using Iiscnfg.vbs.
- For more information about importing an IIS configuration, see Importing IIS Configurations Using Iiscnfg.vbs.
- For more information about metabase backup and restore, see Backing Up and Restoring the Metabase in IIS 6.0.

# Importing Metabase Elements with IIS Manager

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you have created a metabase export file, you can import that file into the metabase configuration file on the same computer or another computer running a member of the Microsoft Windows Server™ 2003 family. To import a file by using a command-line script, see Importing IIS Configurations Using Iiscnfg.vbs.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To import metabase elements**

1. Before importing metabase elements from a file, make sure you have removed or changed all machine- and system-specific information. For more information, see Machine-Specific and System-Specific Information.
2. In IIS Manager, double-click the **Local Computer**; right-click the Web site, FTP site, virtual directory, or application pool node to which you want to import configuration data; point to **New**; and then, depending on the type of node to which you want to import configuration data, click **Web Site (from file)**, **FTP Site (from file)**, **Virtual Directory (from file)**, or **Application Pool (from file)**.
3. In the **File** box, type the path to your import file or click **Browse** to locate it.
4. Click **Read File**, click the object whose configuration you want to import, and then click **OK**.
5. If the element already exists in the metabase, the following message appears: This <*element*> already exists. You can choose **Create a new** <*element*> or **Replace the existing** <*element*>, and then click **OK**. Note that if you choose to replace the existing element, you lose all previous configuration settings.

**Related Information**

- For more information about exporting an IIS configuration, see Exporting IIS Configurations Using Iiscnfg.vbs.
- For more information about importing an IIS configuration, see Importing IIS Configurations Using Iiscnfg.vbs.
- For more information about metabase backup and restore, see Back Up and Restore the IIS Metabase.

# Moving Sites and Applications to Another Computer in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the metabase export feature to configure multiple computers with the same basic IIS configuration settings. The procedures below illustrate how to create complete metabase templates for porting the World Wide Web Publishing Service (WWW service), or elements of a metabase, to another computer.

The first procedure below creates a complete, system-independent copy of the metabase configuration and saves it to a file. The file can then be imported to other IIS 6.0 servers without overwriting or conflicting with settings on the target system. The second procedure below creates a template for importing elements of a configuration to another computer.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**Procedures**

**To port an entire system-independent metabase configuration to another computer**

1. Open a command prompt, and switch to the systemroot\system32 directory.
2. Type **cscript iiscnfg.vbs /export /f** *path******filename***.xml /d** *password* **/sp**, and press ENTER.
3. Copy the .xml file you created to a temporary directory on the target system.
4. On the target system, open a command prompt and switch to the systemroot\system32 directory.
5. Type **cscript iiscnfg.vbs /import /f** *path******filename***.xml /d** *password* **/sp / /dp / /children /inherited /merge**, and press ENTER.
6. Before relying upon the imported elements in a production mode, test your computers after importing metabase elements.

**To port metabase elements to another computer**

1. Configure an IIS server to the state that you want ported to other computers.
2. Export a metabase element to a file, either using IIS Manager; the Exporting IIS Configurations Using Iiscnfg.vbs that is included with IIS; a custom script using the **IIsComputer.Export** method of the Windows Management Instrumentation (WMI) provider for IIS or the **Export** method of the Active Directory Services Interfaces (ADSI) provider for IIS; or a custom application using Admin Base Objects (ABO). For more information about these methods, see Using IIS Programmatic Administration in the IIS Software Development Kit (SDK) on MSDN.

   | **Note** |
   | --- |
   | IIS Manager includes inherited settings by default. If you are not using IIS Manager, you can specify the inherit flag if you want to include all inherited settings. |

3. Clean up the export file by removing or changing all the machine- and system-specific information, and removing all other settings that you do not want applied to other computers. The cleaned-up export file becomes your import file.
4. On the new computers, create any physical folders that are needed for site and virtual directory paths, copy any content files (.asp files, ISAPIs, and so on) that need to be ported, and register any components that were copied over.
5. Import your file onto each computer to which you want to apply the settings. Make sure to use the merge flag if you do not want settings in the import file to overwrite any existing settings on the computer. Choose the inherit flag if you used it to create the export file.

6. Before relying upon the imported elements in a production mode, test your computers after importing metabase elements.

**Related Information**

- For more information about importing an IIS configuration, see Importing IIS Configurations Using Iiscnfg.vbs.
- For more information about exporting an IIS configuration, see Exporting IIS Configurations Using Iiscnfg.vbs.
- For more information about behaviors of imported or exported metabases, see Metabase Import/Export Behavior.
- For more information about metabase import and export options, see Managing IIS Configurations Using Scripts.

# Auditing the Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

Beginning with the SP1 release, IIS 6.0 features metabase auditing to allow tracking of each change that is made to the metabase. Changes to the metabase are recorded under the IIS-metabase event source in the NT Security log as follows:

- What was changed (metabase node, property, and old and new values).
- When the change was made (date and time).
- Who made the change (domain and user name).
- Success or failure of the change attempt (HRESULT).

Metabase auditing is disabled by default. Enabling metabase auditing requires these tasks:

- Globally enable object auditing in Windows Server 2003.
- Enable metabase auditing for IIS 6.0 SP1.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName***"mmc systemroot\system32\inetsrv\iis.msc"**. |

**To globally enable object auditing in Windows Server 2003**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **gpedit.msc**, and then click **OK**.
3. In the Group Policy Editor, expand **Local Computer Policy**, and then click **Computer Configuration**/**Windows Settings/Security Settings/ Local Policies/Audit Policy**.
4. Double click **Audit object access**.
5. Check both **Success** and **Failure** check boxes.
6. Click **OK**.

To globally disable object auditing, repeat steps, and then clear the **Success** and **Failure** check boxes.

**To enable or disable metabase auditing**

1. From the **Start** menu, open the command-line window.
2. Enter one of the following commands to execute the script, iiscnfg.vbs, filling in the path to the metabase location:

   **iiscnfg /enableAudit** *path*

   **iiscnfg /disableAudit** *path*

   For example, to enable metabase auditing on a Web virtual directory, “MyVdir”, type:

   **c:\%systemroot%\system32\iiscnfg.vbs /enableAudit W3SVC/1/ROOT/MyVdir**

For example, to enable metabase auditing on the entire metabase, type:**c:\%systemroot%\system32\iiscnfg.vbs /enableAudit** */ /r*

3. Press **Enter**.

**To verify that metabase auditing is working**

1. Change a metabase setting.
2. In the NT Security event log, verify that the change to the metabase is recorded under the event-source name "IIS-metabase."

For more information about metabase auditing, see Metabase Auditing.

For more information about event messages produced by metabase auditing, see Metabase Auditing Event Messages.

# Web Site Administration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Suppose you just deployed an expensive, new Web site only to discover that a company logo was omitted from all of the site's pages. Or, imagine that you are suddenly asked to route thousands of users to alternate Web sites while technicians scramble to troubleshoot a malfunctioning computer. Such issues only begin to hint at the day-to-day challenges of successfully managing a Web site. And although effective Web site management ultimately depends on an administrator's skill, there are a number of basic tools and procedures that you can use to deal with the most common management tasks and emergencies.

**Getting Started**

To begin with, you should set up your Web sites by indicating which directories contain the documents that you want to publish. The Web server cannot publish documents that are not within these specified directories. So, the first step in deploying a Web site should be to first determine how you want your files organized. You then use IIS Manager to identify which directories are part of the site.

If you want to get started right away without having to create a special directory structure and your files are all located on the same hard disk of the computer running Internet Information Services (IIS), you can publish your documents immediately by copying your Web files into the default home directory, *LocalDrive*:\Inetpub\Wwwroot. (For an FTP site, copy your files into *LocalDrive*:\Inetpub\Ftproot.) Intranet users then can access these files with the following URL: http*://servername/filename*. To learn how to host multiple sites and organize disparate Web content, see the procedures in this section. They explain such things as how to change the default location of your files and how to add directories that are located on other hard disks or on computers across the network.

**Home Directories**

Each Web or FTP site must have one home directory. The *home directory* is the central location for your published pages. It contains a home page or index file that welcomes visitors and contains links to other pages in your site. The home directory is mapped to your site's domain name or to your server name. For example, if your site's Internet domain name is www.microsoft.com and your home directory is C:\Website\Microsoft, then browsers use the URL http://www.microsoft.com to access files in your home directory. On an intranet, if your server name is AcctServer, then browsers use the URL http://acctserver to access files in your home directory.

A default home directory is created when you install IIS, and when you create a new Web site. For instructions to change the home directory, see Setting Home Directories.

**Virtual Directories**

To publish from any directory not contained within your home directory, you must create a *virtual directory*. A virtual directory is a directory that is not contained in the home directory but appears to client browsers as though it were.

A virtual directory has an *alias*, or a name that Web browsers use to access that directory. Because an alias is usually shorter than the path name of the directory, it is more convenient for users to type. An alias is more secure because users do not know where your files are physically located on the server and cannot use that information to modify your files. Aliases make it easier for you to move directories in your site. Rather than changing the URL for the directory, you change the mapping between the alias and the physical location of the directory.

For example, suppose you are setting up a Web site for your marketing team on your company's intranet. The following table shows the mapping between the physical location of the files and the URL that accesses the files.

| Physical Location | Alias | URL |
|---|---|---|
| C:\Inetpub\Wwwroot | *home directory* (none) | http://SampleWebSite |
| \\Server2\SalesData | Customers | http://SampleWebSite/Customers |
| D:\Inetpub\Wwwroot\Quotes | None | http://SampleWebSite/Quotes |
| D:\Inetpub\Wwwroot\OrderStatus | None | http://SampleWebSite/OrderStatus |
| D:\Marketing\PublicRel | PR | http://SampleWebSite/PR |

Both virtual directories and physical directories (directories without an alias) appear in IIS Manager. A virtual directory is indicated by a wheel icon. The following illustration shows the Web site example described above, where /Customers and /PR

are virtual directories.

For a simple Web site, you might not need to add virtual directories. You can simply place all of your files in the site's home directory. If you have a complex site or want to specify different URLs for different parts of your site, you can add virtual directories as needed.

**Reroute Requests with Redirects**
When a browser requests a page on your Web site, the Web server locates the page identified by the URL and returns it to the browser. When you move a page on your Web site, you can't always correct all of the links that refer to the old URL of the page. To make sure that browsers can find the page at the new URL, you instruct the Web server to give the browser the new URL. The browser uses the new URL to request the page again. This process is called "redirecting a browser request" or "redirecting to another URL." Redirecting a request for a page is similar to using a forwarding address with a postal service. The forwarding address ensures that letters and packages addressed to your original residence are delivered to your new residence.

Redirecting a URL is useful when you are updating your Web site and want to make a portion of the site temporarily unavailable, or when you have changed the name of a virtual directory and want links to files in the original virtual directory to access the same files in the new virtual directory.

**Server-Side Includes**
Often, it can be useful to dynamically alter Web content after the content has been requested, but before it gets returned to the browser. IIS includes a feature called server-side includes that provides this functionality.

Using server-side includes (SSI), you can carry out a number of of Web site management activities from adding dynamic time stamps to running a special shell command each time a file is requested. SSI commands, called *directives*, are added to Web pages at design time. When a page is requested, the Web server parses out all the directives it finds in a Web page and then executes them. A commonly used SSI directive inserts, or *includes*, the contents of a file into a Web page. For example, if you are required to continually update a Web page advertisement, you could use SSI to include the advertisement's HTML source in the Web page. To update the advertisement, you need only modify the file containing the advertisement's HTML source. You do not have to know a scripting language to use SSI; simply follow the correct directive syntax.

**ASP and ASP .NET**
ASP is a server-side scripting environment that you can use to create dynamic and interactive Web pages, and build powerful Web applications. When the application server receives a request for an ASP file, it processes server-side script code contained in the file to build the HTML Web page that is sent to the browser. In addition to server-side script code, ASP files can contain HTML (including related client-side scripts) as well as calls to COM components that perform a variety of tasks, such as connecting to a database or processing business logic. ASP requires you to use a scripting language, such as VBScript or JScript.

ASP .NET is the next generation of the Microsoft server-side scripting environment. It provides a new programming model and infrastructure, enabling Web developers to build and deploy enterprise-class Web applications that are more secure, scalable, and stable than ever before.

**Related Information**

- For information about using virtual directories with Web sites, see Using Virtual Directories.
- For information about using virtual directories with an FTP site, see Using Virtual Directories with FTP Sites.
- For more information about redirecting URLs, see Redirecting Requests to Files, Directories, or Programs.
- For more detailed procedural and reference information about server-side includes, see Using Server-Side Include Directives.
- For information about using ASP to Manage Web Site Content, see Active Server Pages on the Microsoft Developer Network (MSDN) Web site.
- For information about ASP .NET to Manage Web Site Content, see ASP .NET on the Microsoft Developer Network (MSDN) Web site.

# About Web Site Administration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Suppose you just deployed an expensive, new Web site only to discover that a company logo was omitted from all of the site's pages. Or, imagine that you are suddenly asked to route thousands of users to alternate Web sites while technicians scramble to troubleshoot a malfunctioning computer. Such issues only begin to hint at the day-to-day challenges of successfully managing a Web site. And although effective Web site management ultimately depends on an administrator's skill, there are a number of basic tools and procedures that you can use to deal with the most common management tasks and emergencies.

**Getting Started**

To begin with, you should set up your Web sites by indicating which directories contain the documents that you want to publish. The Web server cannot publish documents that are not within these specified directories. So, the first step in deploying a Web site should be to first determine how you want your files organized. You then use IIS Manager to identify which directories are part of the site.

If you want to get started right away without having to create a special directory structure and your files are all located on the same hard disk of the computer running Internet Information Services (IIS), you can publish your documents immediately by copying your Web files into the default home directory, *LocalDrive*:\Inetpub\Wwwroot. (For an FTP site, copy your files into *LocalDrive*:\Inetpub\Ftproot.) Intranet users then can access these files with the following URL: http*://servername/filename*. To learn how to host multiple sites and organize disparate Web content, see the procedures in this section. They explain such things as how to change the default location of your files and how to add directories located on other hard disks or on computers across the network.

**Home Directories**

Each Web or FTP site must have one home directory. The *home directory* is the central location for your published pages. It contains a home page or index file that welcomes visitors and contains links to other pages in your site. The home directory is mapped to your site's domain name or to your server name. For example, if your site's Internet domain name is www.microsoft.com and your home directory is C:\Website\Microsoft, then browsers use the URL http://www.microsoft.com to access files in your home directory. On an intranet, if your server name is AcctServer, then browsers use the URL http://acctserver to access files in your home directory.

A default home directory is created when you install IIS, and when you create a new Web site. For instructions to change the home directory, see Setting Home Directories.

**Virtual Directories**

To publish from any directory not contained within your home directory, you must create a virtual directory. A *virtual directory* is a directory that is not contained in the home directory but appears to client browsers as though it were.

A virtual directory has an *alias*, or a name that Web browsers use to access that directory. Because an alias is usually shorter than the path name of the directory, it is more convenient for users to type. An alias is more secure because users do not know where your files are physically located on the server and cannot use that information to modify your files. Aliases make it easier for you to move directories in your site. Rather than changing the URL for the directory, you change the mapping between the alias and the physical location of the directory.

For example, suppose you are setting up a Web site for your marketing team on your company's intranet. The following table shows the mapping between the physical location of the files and the URL that accesses the files.

| Physical Location | Alias | URL |
|---|---|---|
| C:\Inetpub\Wwwroot | *home directory* (none)" | http://SampleWebSite |
| \\Server2\SalesData" | Customers | http://SampleWebSite/Customers |
| D:\Inetpub\Wwwroot\Quotes | *None* | http://SampleWebSite/Quotes |
| D:\Inetpub\Wwwroot\OrderStatus | *None* | http://SampleWebSite/OrderStatus |
| D:\Marketing\PublicRel | PR | http://SampleWebSite/PR |

Both virtual directories and physical directories (directories without an alias) appear in IIS Manager. A virtual directory is indicated by a wheel icon. The following illustration shows the Web site example described above, where /Customers and /PR

are virtual directories.




For a simple Web site, you might not need to add virtual directories. You can simply place all of your files in the site's home directory. If you have a complex site or want to specify different URLs for different parts of your site, you can add virtual directories as needed.

**Reroute Requests with Redirects**

When a browser requests a page on your Web site, the Web server locates the page identified by the URL and returns it to the browser. When you move a page on your Web site, you can't always correct all of the links that refer to the old URL of the page. To make sure that browsers can find the page at the new URL, you instruct the Web server to give the browser the new URL. The browser uses the new URL to request the page again. This process is called "redirecting a browser request" or "redirecting to another URL." Redirecting a request for a page is similar to using a forwarding address with a postal service. The forwarding address ensures that letters and packages addressed to your original residence are delivered to your new residence.

Redirecting a URL is useful when you are updating your Web site and want to make a portion of the site temporarily unavailable, or when you have changed the name of a virtual directory and want links to files in the original virtual directory to access the same files in the new virtual directory.

**Server-Side Includes**

Often, it can be useful to dynamically alter Web content after the content has been requested, but before it gets returned to the browser. IIS includes a feature called server-side includes that provides this functionality.

Using server-side includes (SSI), you can carry out a number of of Web site management activities from adding dynamic time stamps to running a special shell command each time a file is requested. SSI commands, called *directives*, are added to Web pages at design time. When a page is requested, the Web server parses out all the directives it finds in a Web page and then executes them. A commonly used SSI directive inserts, or *includes*, the contents of a file into a Web page. For example, if you are required to continually update a Web page advertisement, you could use SSI to include the advertisement's HTML source in the Web page. To update the advertisement, you need only modify the file containing the advertisement's HTML source. You do not have to know a scripting language to use SSI; simply follow the correct directive syntax.

**ASP and ASP .NET**

ASP is a server-side scripting environment that you can use to create dynamic and interactive Web pages, and build powerful Web applications. When the application server receives a request for an ASP file, it processes server-side script code contained in the file to build the HTML Web page that is sent to the browser. In addition to server-side script code, ASP files can contain HTML (including related client-side scripts) as well as calls to COM components that perform a variety of tasks, such as connecting to a database or processing business logic. ASP requires you to use a scripting language, such as VBScript or JScript.

ASP .NET is the next generation of the Microsoft server-side scripting environment. It provides a new programming model and infrastructure, enabling Web developers to build and deploy enterprise-class Web applications that are more secure, scalable, and stable than ever before.

**Related Information**

- For information about using virtual directories with Web sites, see Using Virtual Directories.

# Naming Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A server running IIS can host multiple Web sites. Each Web site is said to be running on a virtual server. Each virtual Web site has a descriptive name and can support one or more host header names. Host header names make it possible to host multiple domain names on one computer.

You can have multiple server certificates, multiple IP addresses, and multiple SSL ports per server.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To name a Web site**

1. In IIS Manager, expand the local computer, right-click the Web site, and click **Properties**.
2. On the **Web Site** tab, in the **Description** box, type a descriptive name for the site, and then click **OK**.

**To assign a host header name to a Web Site**

1. Ask your system administrator to add the new site name mapping for your static IP address to your name resolution system (typically DNS).
2. In IIS Manager, right-click the Web site to which you which you want to assign a host header name, and click **Properties**.
3. On the **Web Site** tab, click **Advanced**.
4. Click **Add** to assign a host header name, IP address, and port for the Web site.
5. Click **OK** twice.

**Note**

If you attempt to create a host header for a Web site containing a Unicode character, you will receive an error indicating that you must only use standard characters. This is because the Domain Name System (DNS) standard does not recognize Unicode characters.

**Related Information**

- For more information about using host header names, see Creating Multiple Sites Using Host Header Names.

# Starting and Stopping Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, sites start automatically when your computer starts. Stopping a site does not interfere with other Internet services that are running. Pausing a site prevents the site from accepting new connections but does not affect requests that are already being processed. Starting a site resumes Internet services for the site.

You can start, stop, or pause a Web site using either IIS Manager or a command-line administration script. Both methods are detailed below.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To start, stop, or pause a Web site using IIS Manager**

- In IIS Manager, right-click the site you want to start, stop, or pause; and click **Start**, **Stop**, or **Pause**.

**Note**

If a site stops unexpectedly, IIS Manager might not correctly indicate the state of the site. In IIS, right-click the **Web Sites** folder and click **Refresh** to see the current state of all Web sites.

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters).

**To start, stop, or pause a Web site using a command-line administration script**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. At the command prompt, type **cscript iisweb.vbs /***action* **"***WebSiteName***"** and press ENTER. Substitute start, stop, or pause for *action*. Substitute your actual Web site name (for example "Default Web Site") for *WebSiteName*.

**Related Information**

- For information about restarting all of your Internet services, see Restarting IIS.
- To quickly establish a site that uses the default settings, see Site Setup.
- To change properties on existing Web sites, see Changing Default Web Site Settings Web site.

# Working with MIME Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Multipurpose Internet Mail Extensions (MIME) types instruct a Web browser or mail application how to handle files received from a server. For example, when a Web browser requests an item on a server, it also requests the MIME type of the object. Some MIME types, like graphics, can be displayed inside the browser. Others, such as word processing documents, require an external helper application to be displayed.

When IIS delivers a mail message to a mail application, or a Web page to a client Web browser, it also sends the MIME type of the data it is sending. If there is an attached or embedded file in a specific format, IIS also tells the client application the MIME type of the embedded or attached file. The client application then knows how to process or display the data being received from IIS.

IIS serves only files with known extensions registered in the MIME types list or with the operating system. IIS allows you to configure additional MIME types and change or remove MIME types. Removing a MIME type in IIS does not block access to that MIME type by other applications if it is also registered with the operating system.

IIS is preconfigured to recognize a default set of global MIME types. These MIME types are recognized by all Web sites you create in IIS. MIME types can also be defined at the Web site and directory levels, independent of one another or the types defined globally. When you view MIME types at the Web site or directory level, only the types unique to that level are displayed, not all types inherited from the next level up.

| **Note** |
| --- |
| If you apply a MIME type at the global level after modifying the same MIME type at a lower level, the global-level MIME type overrides the modified MIME type at the lower level. |

| **Note** |
| --- |
| IIS returns a 404.3 error if a client request refers to a file name extension that is not defined in the MIME types. |

**In This Section**

This section contains the following information:

- Viewing Available MIME Types: Explains how to view available MIME types by using IIS Manager and Adsutil.vbs.
- Creating Global MIME Types: Explains how to add MIME types at the global level in IIS.
- Creating MIME Types for a Web site or directory: Explains how to add MIME types at the Web site or directory level in IIS.
- Editing Global MIME Types: Explains how to edit MIME types at the global level in IIS.
- Editing MIME Types for a Web site or directory: Explains how to edit MIME types at the Web site or directory level in IIS.
- Removing Global MIME Types: Explains how to remove MIME types at the global level in IIS.
- Removing MIME Types from a Web site or directory: Explains how to remove MIME types at the Web site or directory level in IIS.

# Viewing Available MIME Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To help prevent attackers from sending files that could compromise the security of your Web server, the IIS static file handler serves only the files that are explicitly listed in the list of Multipurpose Internet Mail Extensions (MIME) types. IIS is preconfigured to recognize a default set of global MIME types that are recognized by all of the Web sites that you create with IIS.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To see a list of all available MIME types by using IIS Manager**

1. In IIS Manager, right-click the local computer, and then click **Properties**.
2. In **MIME types**, click **MIME Types**.

   The registered MIME types are displayed in **Registered MIME types (file extensions)**. You can add, change, or remove MIME types from the list.

**To see a list of all available MIME types by using Adsutil.vbs**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and then click **OK**.
3. At the command prompt, switch to the *drive*\Inetpub\AdminScripts directory.
4. Type the following at the command prompt and then press ENTER:

   **cscript.exe adsutil.vbs enum /Mimemap**

**Related Sections**

- For more information about MIME types, see MIME Types.
- For more information about adding, editing, or removing MIME types, see Working with MIME Types.
- For more information about using Adsutil.vbs, see Using the Adsutil.vbs Administration Script.

# Creating Global MIME Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you have MIME types that you want all of your Web sites to recognize, you can add the new MIME types at the global level in IIS.

| **Important** |
| --- |
| You can also configure IIS to serve undefined file types by adding a wildcard character (*) MIME type. Do not use wildcard MIME-types on production servers. Doing so can result in IIS serving unrecognized files and displaying sensitive information to users. Wildcard MIME-types are intended for testing purposes or in scenarios where Internet Server API (ISAPI) filters have been developed specifically to handle these wildcard scenarios, for example, a custom authentication ISAPI. |

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

### To add a global MIME type

1. In IIS Manager, expand the local computer, right-click the computer on which you want to add a MIME type, and click **Properties**.
2. Click **MIME Types**.
3. Click **New**.
4. In the **Extension box**, type the file name extension.
5. In the **MIME type** box, type a valid MIME type.

   To create a MIME type for an undefined MIME type, type an asterisk (*****) in the **Extension** box, and type **application/octet-stream** in the **MIME type** box.

   To create a MIME type for a file without an extension, type a period (**.**) in the **Extension** box, and type your MIME type in the **MIME type** box.
6. Click **OK**.

# Creating MIME Types for a Web site or directory

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you have MIME types that you want a particular Web site or directory to recognize, you can add the new MIME types at the Web site or directory level in IIS.

| **Important** |
| --- |
| You can also configure IIS to serve undefined file types by adding a wildcard character (*****) MIME type. Do not use wildcard MIME-types on production servers. Doing so can result in IIS serving unrecognized files and displaying sensitive information to users. Wildcard MIME-types are intended for testing purposes or in scenarios where Internet Server API (ISAPI) filters have been developed specifically to handle these wild card scenarios, for example, a custom authentication ISAPI. |

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To add a MIME type to a Web site or directory**

1. In IIS Manager, right-click the Web site or Web site directory for which you want to add a MIME type, and click **Properties**.
2. Click the **HTTP Headers** tab.
3. Click **MIME Types**.
4. Click **New**.
5. In the **Extension** box, type the file name extension.
6. In the **MIME type** box, type a valid MIME type. If you define a MIME type that has already been defined at a higher level, you are prompted to select the level where the MIME type should reside.

   To create a MIME type for an undefined MIME type, type an asterisk (*****) in the **Extension** box, and type **application/octet-stream** in the **MIME type** box.

   To create a MIME type for a file without an extension, type a period (**.**) in the **Extension** box, and type your MIME type in the **MIME type** box.
7. Click **OK**.

The following illustration shows the relationship between a MIME type defined in IIS and a file type defined on a client computer.

New MIME type being defined in IIS.




Matching File Type and Application defined on client.



# Editing Global MIME Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you need to make changes to global MIME types, you can edit them at the global level in IIS.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To edit a global MIME type**

1. In IIS Manager, expand the local computer, right-click the computer on which you want to edit a MIME type, and click **Properties**.
2. Click **MIME Types**.
3. Select the item you want to edit from the **Registered MIME types (file extensions)** list.
4. Click **Edit**.
5. Edit the information in the **Extension** and **MIME type** boxes as necessary.
6. Click **OK**.

# Editing MIME Types for a Web site or directory

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you need to make changes to MIME types recognized by a particular Web site or directory, you can edit them at the Web site or directory level in IIS.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To edit a MIME type for a Web site or directory**

1. In IIS Manager, right-click the Web site or Web site directory for which you want to add a MIME type, and click **Properties**.
2. Click the **HTTP Headers** tab.
3. Click **MIME Types**.
4. Select the item you want to edit from the **Registered MIME types (file extensions)** list.
5. Click **Edit**.
6. Edit the information in the **Extension** and **MIME type** boxes as necessary.
7. Click **OK**.

# Removing Global MIME Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you want to globally restrict IIS from serving files with certain extensions, you can remove the related MIME types from the global list of registered MIME types in IIS Manager.

If you use IIS Manager to remove a MIME type mapping, you must also verify that the MIME type does not exist in the registry at HKEY_CLASSES_ROOT*extension*. If the MIME type also exists in the registry, you must remove it. If the MIME type continues to exist in the registry, files with that extension continue to be served. Be aware, however, that if you remove a MIME type mapping from the registry, other applications may be affected, because the registry database is shared by multiple applications.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

### To remove a Global MIME type

1. In IIS Manager, expand the local computer, right-click the computer on which you want to add a MIME type, and click **Properties**.
2. Click **MIME Types**.
3. From the **Registered MIME types** list, click the MIME type you want to remove, and then click **Remove**.
4. Click **OK**.

**Related Information**

For information about the inheritance and merge behavior of MIME types, see MIME Types.

# Removing MIME Types from a Web site or directory

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you do not want to serve files with certain extensions from a Web site or directory, you can remove the related MIME types from the list of MIME types registered for that particular Web site or directory.

If you use IIS Manager to remove a MIME type mapping, you must also verify that the MIME type does not exist in the registry at HKEY_CLASSES_ROOT*extension*. If the MIME type also exists in the registry, you must remove it. If the MIME type continues to exist in the registry, files with that extension continue to be served. Be aware, however, that if you remove a MIME type mapping from the registry, other applications may be affected, because the registry database is shared by multiple applications.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

### To remove a MIME type from a Web site or directory

1. In IIS Manager, right-click the Web site or Web site directory from which you want to remove a MIME type, and click **Properties**.
2. Click the **HTTP Headers** tab.
3. Click **MIME Types**.
4. From the **Registered MIME types** list, click the MIME type you want to remove, and then click **Remove**.
5. Click **OK**.

**Related Information**

For information about the inheritance and merge behavior of MIME types, see MIME Types.

# Changing Default Web Site Settings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

During IIS installation, default values are set for the various properties assigned to Web sites. Properties can be set at the global level, affecting all Web sites on a server, on the individual site level, at the directory level, or at the file level. IIS uses an inheritance model, which means that settings on higher levels are automatically inherited by lower levels. Settings at lower levels can be edited individually to override inherited settings from the next level up.

If you change a setting at a lower level, then later change a setting at a higher level that conflicts with the lower level setting, you will be prompted to choose whether you want to change the lower-level setting to match the new higher-level setting.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set master properties for all sites assigned to a computer**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** folder, and click **Properties**.
2. Click **Help** on any of the tabs for more information about those settings.

**To set properties for a single Web site**

1. In IIS Manager, right-click the Web site whose settings you want to change, and click **Properties**.
2. Click **Help** on any of the tabs for more information about those settings.

# Implementing Incremental Web Site Numbering

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In IIS 5.1 and earlier, IIS-generated Web site numbers incrementally. The first site was numbered 1, the second site numbered 2, and so on. Beginning with IIS 6.0, the default numbering scheme has changed. Upon installation of IIS 6.0, IIS generates a default Web site that is numbered 1. IIS then generates the Web site identification number, based on a hash of the site name, for any additional sites that you create.

You can retain the incremental numbering scheme of IIS 5.1 and earlier. This might be preferable if you have existing scripts, applications, or logging tools that rely on incremental site numbering.

See "Related Information" at the end of this topic for links to specific step-by-step procedures related to the general steps described.

| **Important** |
| --- |
| This topic contains information about modifying the registry. Using Registry Editor incorrectly can cause serious problems that require reinstalling the operating system. Because Registry Editor bypasses the standard safeguards that prevent you from entering settings that are conflicting or likely to degrade performance or damage your system, exercise caution when making changes to the registry. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. For information about how to edit the registry, see the "Changing Keys and Values" topic in Registry Editor Help. |

**Procedures**

### Implementing Incremental Web Site Numbering on an Existing IIS 6.0 Installation

1. Enable incremental Web site numbering on your production IIS 6.0 server and on a staging IIS 6.0 server.
2. Export the entire metabase from the production server and import this metabase to the staging server.
3. Enable the metabase edit-while-running feature on the staging server, and then stop the World Wide Web Publishing service (WWW service).
4. Open the MetaBase.xml file on the staging server in Notepad or another plain text editor.
5. Find the first Web site that you want to renumber in the MetaBase.xml file. Perform a search-and-replace through the entire metabase, replacing every occurrence of the original Web site number with the incremental number that you want to assign to the site. Repeat the search-and-replace pass for all Web sites, incrementing each successive Web site number by 1.
6. Save the MetaBase.xml file and restart the WWW service.
7. Test your sites and Web applications by requesting content from each of them. If any failures occur, identify the cause and correct it. Most failures will likely result from leftover references to the original Web site number.
8. Disable edit-while-running (once debugging is complete), test the sites again for functionality, and then export the entire metabase from the staging server and import it onto the production server.
9. Restart the WWW service on your production server and then validate functionality by making a request to the home page of your main site.

| **Important** |
| --- |
| If you want to use incremental site numbering in IIS 6.0, it is best to enable incremental site numbering immediately after installing IIS 6.0, and before creating any new Web sites. This avoids the possibility of accidentally creating new Web sites with numbers out of sequence. |

### Notes

- Enabling incremental Web site numbering does not change the identification number of existing sites.

- If your server already contains Web sites that were created by using the IIS 6.0 default numbering scheme, and you enable incremental Web site numbering, the next site created will be numbered with the first unused incremental number. For example, if you enable incremental site numbering on a server with existing sites numbered 1, 2, 3, and 10, the next site number will be 4.

- When remotely administering a server running IIS 6.0 with incremental site numbering enabled, make sure that the server from which you are performing the remote administration also has incremental site numbering enabled. This is necessary because the value of the IncrementalSiteIDCreation registry entry on the system from which you are performing the administration determines how site identification numbers are generated on the IIS server being administered.

**Related Information**

- For information about the IncrementalSiteIDCreation registry entry, see Global Registry Entries.

- For information about how to enable metabase edit-while-running, see Enabling Edit-While-Running in IIS 6.0.

- For information about exporting and importing the metabase, see Exporting IIS Configurations Using Iiscnfg.vbs and Importing IIS Configurations Using Iiscnfg.vbs.

# Using Content Expiration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you have time-sensitive information on your Web site, you can configure settings that will ensure that obsolete information isn't cached by proxies or Web browsers. You can configure Web site content to automatically expire at any time. When content expiration is enabled, the Web browser compares the current date to the expiration date to determine whether to display a cached page or request an updated page from the server. Server-side technology, such as Microsoft ASP.NET can be used to dynamically change the content served. Typically, time-sensitive information is limited to individual files, directories, or Web sites; however, you can also set content expiration for all Web sites on a computer.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set the expiration of Web site content**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** folder, an individual Web site folder, a virtual directory, or a file; and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element. |

2. Click the **HTTP Headers** tab.
3. Select the **Enable content expiration** check box.
4. Click **Expire immediately**, **Expire after**, or **Expire on**, and type the appropriate expiration information in the corresponding boxes.
5. Click **OK**.

To verify that expiration dates are working correctly, look for Not Modified messages in the HTTP response packet of your Network Monitor capture summary.

**Related Information**

- For more information about Web site settings, see Changing Default Web Site Settings.
- For information about additional ways to optimize performance, see Optimizing a Page-Load Performance Profile.

# Using Content Ratings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure your Web server's content rating features to embed descriptive labels in HTTP headers of your Web page. Some Web browsers, such as Microsoft Internet Explorer 3.0 or later, can detect these content labels and help users to identify potentially objectionable Web content. Your Web server's default Platform for Internet Content Selection (PICS)-based rating system uses a system developed by the Internet Content Rating Association (ICRA), which rates content according to levels of violence, nudity, sexual content, and offensive language. Before setting your Web content ratings, you should complete a ICRA questionnaire to obtain the recommended content ratings for your particular Web content.

If you are running a publicly viewable Web site, it is good practice to have the content rated. Doing so does not automatically make your site unavailable to users because the users must also enable content ratings and set content thresholds on their client computer for the rating information to have any effect on their browsing experience.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set content ratings**

1. In IIS Manager, expand the local computer; right-click a Web site, directory, or file; and click **Properties**.
2. Click the **HTTP Headers** tab.
3. In the **Content rating** section, click **Edit Ratings**.
4. Click the **Ratings** tab and select the **Enable ratings for this content** check box.
5. In the **Category** list, click a ratings category. Drag or click the rating slider to set the level of potentially objectionable material for the category. Each setting displays a description of the rating level.
6. Under **E-mail address of person rating this content**, type the e-mail address of the person who rated the content.
7. Click the arrow next to the **Expire on** box and click a ratings expiration date from the calendar.
8. Click **OK**.

# Using Server-Side Include Directives

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Server-side include (SSI) directives instruct the Web server to insert various types of content into a Web page. These types of content can include the value of a server variable (sometimes called a CGI variable), the output of a shell command, a modification date and/or size of a file, or the contents of another file (when the #include statement is used). Directives can also include a string of SSI errors.

You can use SSI directives to instruct the Web server to include text, graphics, or application output into a Web page just before sending that information to the browser. You can use SSI to include, for example, a time/date stamp, a copyright notice, or a form for a customer to fill out and return. Including a file is an easy way to incorporate text or links to graphics that are repeated in many files. SSI provides special preprocessing directives that help you access information about a file or display the value of a server variable. SSI also provides a directive that you can use to insert the output from an application or a shell command into a Web page. Rather than adding the content or output to every file, you can save the content into an include file. The include files are called through a simple statement that instructs the Web server to insert the content or output into the applicable Web pages.

IIS processes SSI directives while it processes the Web page. When the Web server reaches an SSI directive, it inserts the contents of the include file, or the output of the application directly into the Web page. If the include file contains an SSI directive, that file or output is also inserted. In addition to the basic directive used to include a file, you can use SSI directives to insert information about a file (such as its size) or to run an application or a shell command.

By default, IIS is set to process server-side include files with the file name extensions .stm, .shtm, or .shtml, although you can configure other file name extensions to be processed as well. IIS also looks for and processes the #include directive in files with the .asp file name extension. Ensure that all files containing SSI directives have a defined SSI file name extension. Furthermore, SSI directives, as well as other script maps, work in any file whose file name extension is mapped to the ssinc.dll. If, for some reason, the extensions were deleted from the applications mappings list, use the procedure below to add the SSI file name extensions to the list.

Also be sure to put SSI files into a directory with Scripts or Execute access permissions. You should limit user access to this directory to avoid attacks by malicious users. For more information, see Access Control with IIS 6.0 and Setting NTFS Permissions for Directories or Files.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable server-side includes**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** folder (to enable SSI on all Web sites) or right-click a specific Web site, and click **Properties**.
2. Click the **Home Directory** tab.
3. In the **Application settings** section, click **Configuration**.
4. On the **Mappings** tab, click **Add**.
5. In the **Add/Edit Application Extension Mapping** dialog box, enter all of the required information.
6. Click **OK**.

To disable server-side includes on your IIS Web server, you must remove SSI extensions from the application mappings list. If your Web server receives a request for a page and the SSI extension has been removed, IIS will not process the SSI directive within the page. You can disable the processing of server-side include directives on all Web sites or on a specific Web site.

**To disable server-side includes**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** folder (to disable SSI on all Web sites) or right-click a specific Web site, and click **Properties**.
2. Click the **Home Directory** tab.
3. In the **Application settings** section, click **Configuration**.
4. On the **Mappings** tab, scroll down in the **Application Mappings** list, click the SSI extension you wish to disable, and then click **Remove**.
5. Click **OK**.

**Related Information**

- For information about enabling the execution of shell commands, see the SSIEnableCmdDirective section in WWW Service Registry Entries.
- For specific programmatic information about implementing server-side include directives, search for "ssi" on the Microsoft Developer Network (MSDN) Web site.

# Redirecting Requests to Files, Directories, or Programs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can redirect requests for files in one directory to a different directory, to a different Web site, or to a file in a different directory. When the browser requests the file from the original URL, the Web server instructs the browser to request the page from the redirected URL.

An advanced capability of the redirect method allows you to redirect all requests for files in a particular directory to a program. Generally speaking, you should also pass any parameters from the original URL to the program, which you can do by using redirect variables.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To redirect requests to another directory or Web site**

1. In IIS Manager, expand the local computer, right-click the Web site, virtual directory, or directory within a Web site to which you want to add a redirect, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab.
3. Under **The content for this resource should come from**, click **A redirection to a URL**.
4. In the **Redirect to** box, type the path to the destination directory or the URL of the Web site. For example, to redirect all requests for files in the /Catalog directory to the /NewCatalog directory on another Web site, type **http://***sitename***/NewCatalog**. The redirect must be to a fully qualified URL.
5. Set the required flags as specified below, and then click **OK**. More than one flag is allowed.
   - Select **The exact URL entered above** to indicate that the value provided for the destination is an absolute target location.
   - Select **A directory below this one** to indicate that redirection should occur only once, because the destination is in a subdirectory of the original URL. This setting also instructs IIS to redirect requests only if they are to subfolders or files of the original URL.
   - Select **A permanent redirection for this resource** to indicate that this redirection is permanent for this resource.
   - Clear **A permanent redirection for this resource** to indicate that this redirection is temporary for this resource.

**To redirect all requests to a single file**

1. In IIS Manager, right-click the Web site or directory, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab.
3. Under **The content for this resource should come from**, click **A redirection to a URL**.
4. In the **Redirect to** box, type the URL of the destination file. The redirect must be to a fully qualified URL.
5. Under **The client will be sent to**, select The exact URL entered above check box to prevent the Web server from appending the original file name to the destination URL.
6. Click **OK**.

You can use wildcards and redirect variables in the destination URL to precisely control how the original URL is translated into the destination URL.

**To redirect requests to a program**

1. In IIS Manager, right-click the Web site or directory, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab.
3. Under **The content for this resource should come from**, click **A redirection to a URL**.
4. In the **Redirect to** box, type the URL of the program, including any redirect variables needed to pass parameters to the program.

   For example, to redirect all requests for scripts in a Scripts directory to a logging program that records the requested URL and any parameters passed with the URL, type **/Scripts/Logger.exe?URL=$V+PARAMS=$P**. $V and $P are redirect variables.
5. Under **The client will be sent to**, select **The exact URL entered above** check box to prevent the Web server from appending the original file name to the destination URL.
6. Click **OK**.

**Related Information**

- For more information about writing redirection strings with variables and flags, see the Redirect Reference.

# Restarting IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You might need to restart Internet Information Services (IIS) before certain configuration changes take effect or when applications become unavailable. Restarting IIS is the same as first stopping IIS, and then starting it again, except it is accomplished with a single command. There are two ways to restart IIS:

- Using IIS Manager.
- Using the **IISReset** command-line utility.

Both methods allow you to stop, start, and restart IIS Internet services. The **Restart IIS** command in IIS Manager and the **IISReset** command-line utility do not restart the HTTP protocol stack (HTTP.sys) or Internet services outside of IIS. In previous versions of IIS, restarting IIS using the **IISReset** command-line utility was a common practice for recovering from an unresponsive application.

Restarting or stopping IIS, or rebooting your Web server, is a severe action. When you restart the Internet service, all sessions connected to your Web server (including Internet, FTP, SMTP, and NNTP) are dropped. Any data held in Web applications is lost. All Internet sites are unavailable until Internet services are restarted. For this reason, you should avoid restarting, stopping, or rebooting your server if at all possible. IIS 6.0 includes application pool recycling and several other features that provide alternatives to restarting IIS. For a list of features designed to improve IIS reliability and remedy the need to restart IIS, see "Alternatives to Restarting IIS" in this topic.

| **Note** |
| --- |
| Changes to the metabase can be lost when restarting IIS. To avoid losing metabase changes and to trigger history files that back up the metabase, use the SaveData method. For more information about the SaveData method, see "SaveData" in the Platform SDK on MSDN. |

With IIS 6.0, the World Wide Web Publishing Service (WWW service) lives in the service host, Svchost.exe. The FTP, NNTP, and SMTP services and the IIS metabase, known as the IIS Admin service, lives in Inetinfo.exe. If the IIS Admin service terminates abnormally, IIS restarts automatically. This feature is known as Automatic Restart. Previously, in IIS 5.0, if the IIS Admin service terminated abnormally, both the WWW service and IIS Admin service had to be restarted because they shared the same application space. In IIS 6.0, if the IIS Admin service terminates abnormally, the WWW service does not go down because the IIS Admin service and the WWW service run in separate process spaces. In this case, the WWW service acknowledges that the metabase has terminated abnormally and checks to see if the **IISReset** command-line utility is configured on the IIS Admin service. If **IISReset** is configured on the IIS Admin service, IIS waits for the IIS Admin service to start again and reconnects the WWW service.

All of the Internet services listed below, if installed, are affected when you restart IIS. Not all of the services listed below are installed by default.

| **Service** | **Description** |
| --- | --- |
| IIS Admin service | This service manages all the services of IIS other than the WWW service (FTP, NMTP, and SMTP). |
| WWW service | This service provides Web connectivity between clients and Web sites. |
| HTTP SSL service | This service provides secure Web connectivity between clients and Web sites. |
| FTP service | This service provides FTP connectivity and administration through IIS Manager. |
| SMTP service | This service transports electronic mail across the network. |
| NNTP service | This service transports network news across the network. |

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To restart IIS using IIS Manager**

1. In IIS Manager, right click the local computer, point to **All Tasks**, then click **Restart IIS**.
2. In the **What do you want IIS to do** list, click **Restart Internet Services on** *computername*.
3. IIS attempts to stop all services before restarting. IIS waits up to five minutes for all services to stop. If the services cannot be stopped within five minutes, all IIS services are terminated, and IIS restarts. In addition, clicking **End now** forces all IIS services to stop immediately, and IIS is restarted.

**Caution**

If IIS is terminated before its services have stopped normally, data loss can result. If possible, back up your IIS configuration before restarting IIS. See Backing Up and Restoring the Metabase for information about backing up your IIS configuration.

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters).

**To restart IIS using the IISReset command-line utility**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. At the command prompt, type **iisreset /noforce** *computername*, and press ENTER.

   **Note**

   The **/noforce** parameter is recommended as a safeguard against data loss in case the IIS services cannot all be stopped within the one minute timeout period. If you are certain that it is safe to force IIS to restart, you can omit the **/noforce** parameter. In addition, if you are logged on locally, the *computername* parameter is not required. If you are administering an IIS server remotely, the *computername* parameter is the NetBIOS name of the computer on which you wish to restart IIS.

4. IIS attempts to stop all services before restarting. The **IISReset** command-line utility waits up to one minute for all services to stop. If the services cannot be stopped within one minute, all IIS services are terminated, and IIS restarts.

**Note**

You can adjust the timeout value using the **iisreset /timeout** command.

The IISReset command-line utility has several options available. For more information about the options available with the **IISReset** command-line utility, open a command-prompt window, and type **iisreset /?** You can also search for "iisreset" in Windows Help.

**Automatic Restart**

Automatic Restart is enabled by default for the IIS Admin service and WWW service. If a process within either service terminates abnormally, the service restarts automatically. The WWW service restarts using the **net start** command. The IIS

Admin service restarts using the **iisreset** command. These commands are issued internally without user intervention. You can disable Automatic Restart using the following procedures.

**To disable restart of IIS Admin service**

1. From the **Start** menu, right-click **My Computer**, and click **Manage**.
2. In the console tree, expand the **Services and Applications** node, and click **Services**.
3. Right-click **IIS Admin Service**, and click **Properties**.
4. Click the **Recovery** tab, and in the **First failure**, **Second failure**, and **Subsequent failures** list boxes, click **Take No Action**.
5. Click **OK**.

The following command can only be issued on the local machine.

**To disable restart of IIS Admin service from the command line**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. At a command prompt, type **iisreset /disable**, and press ENTER. This command disables IIS and prevents all **iisreset** calls from executing.

| **Note** |
| --- |
| To re-enable the IIS Admin service, at the command prompt, type **iisreset /enable**, and press ENTER. |

**To disable restart of the WWW service**

1. From the **Start** menu, right-click **My Computer**, and click **Manage**.
2. In the console tree, expand the **Services and Applications** node, and click **Services**.
3. Right-click **World Wide Web Publishing Service**, and click **Properties**.
4. Click the **Recovery** tab, and in the **First failure**, **Second failure**, and **Subsequent failures** list boxes, click **Take No Action**.
5. Click **OK**.

**Alternatives to Restarting IIS**

If you have to restart IIS often because of a faulty application, consider debugging or rewriting the application. To ensure that your applications do not interfere with other services and cause the need to restart IIS, it is recommended that you enable the following:

- Worker Process Isolation Mode in IIS 6.0: Worker process isolation mode is the default service mode in IIS. Worker process isolation mode delivers all the benefits of IIS 6.0 new architecture: robust application pooling, automated restarts, scalability, debugging, and precise performance tuning.
- Recycling Worker Processes with IIS 6.0: In worker process isolation mode, IIS can be configured to periodically restart worker processes in an application pool, allowing you to manage precisely those worker processes that are faulty.
- Configuring Rapid-Fail Protection in IIS 6.0: Worker process isolation mode can be configured so that if a particular application pool experiences consecutive multiple failures, the application can be automatically disabled, allowing for rapid-fail protection.
- Enabling CPU Monitoring in IIS 6.0: CPU monitoring is a tool for monitoring and automatically shutting down worker processes that consume large amounts of CPU time.

**Related Information**

- For more information about improving the performance of IIS network resources using related features like bandwidth throttling, HTTP Compression, and connection limits, see Administering Network Resources.

- For more information about configuring IIS to run Web applications and how to deal with reliability issues caused by Web applications, see Configuring Servers for Applications.

- For information about starting and stopping individual Web and FTP sites, see Starting and Stopping Web Sites.

# FTP Site Administration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

Whether your site is on an intranet or the Internet, the principles of providing a place to upload and download files using FTP are the same. You place your files in directories on your FTP server so that users can establish an FTP connection and transfer files with an FTP client or FTP-enabled Web browser. But beyond simply storing files on your server, you must manage how your site is deployed, and more importantly, how your site evolves. This section presents the basics of managing the infrastructure of an FTP site, from securing your site to hosting multiple sites.

This version of Internet Information Services (IIS) includes Isolating FTP Users to help administrators, and particularly Internet hosting providers, efficiently secure and commercialize the FTP services for their customers.

The FTP service is not installed by default on the operating systems in the Microsoft Windows Server™ 2003 family. To install the FTP service, see FTP Site Setup.

This section includes the following information:

- Changing FTP Site Home Directories: Describes the concept of a home directory and methods for changing the home directory of an FTP site.
- Naming FTP Sites: Describes assigning a descriptive name to an FTP site.
- Setting FTP Messages and Directory Output Style: Describes the message and directory output options.
- Using Virtual Directories with FTP Sites: Describes the concept of a virtual directory and methods for creating and removing a virtual directory.
- Stopping and Starting FTP Sites: Describes why you would need to stop and restart your FTP sites and how to perform these actions.
- Changing Default FTP Site Settings: Describes how to change default settings globally or on an individual site.
- Creating Multiple FTP Sites: Describes how to use IP addresses or port numbers to differentiate multiple FTP sites.
- Adding FTP Sites to Your Server: Describes the process of adding a new FTP site to a server running IIS.
- Redirecting FTP Requests to Directories or Network Shares: Describes forwarding file requests to a different directory or network share.
- Modes and Data Transmission: Describes the different FTP connection and data transfer modes supported by IIS.
- Server-to-Server FTP Transfer: Describes how to configure IIS for server-to-server FTP transfer by changing the EnablePortAttack registry key. The topic also suggests using Web Distributed Authoring and Versioning (WebDAV) publishing as a file transfer alternative.
- Restarting FTP Uploads and Downloads: Describes the FTP **rest** command, which addresses the problem of losing a network connection while downloading a file.
- Securing FTP Sites: Describes some of the misconceptions about FTP security and how to establish a secure FTP site.
- Isolating FTP Users: Describes the concept of FTP user isolation and which type of isolation to use to restrict users to their own directories.

# Adding FTP Sites to Your Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

IIS enables you to create multiple FTP sites on a single server. To add a site to a server, you must prepare the server and associated network services, and create a unique identity for the site.

**Preparing to Add an FTP Site**

Adding an FTP site to a server that is already hosting an FTP site requires some preparation.

- Multiple FTP sites can use the same IP address. If you use a unique IP address to identify a new FTP site, however, you must obtain a static IP address from your organization or Internet service provider (ISP) and configure the TCP/IP settings on the server.

- Standard FTP services use TCP port 21 by default, and it is generally not recommended to use any other port for FTP services. However, if you use a nonstandard TCP port number to identify a new FTP site for special situations (typically used on private Web sites for development and testing purposes), select a TCP port number above 1023 so it does not conflict with well-known port numbers assigned by the Internet Assigned Numbers Authority.

- Use Windows Explorer to create a home directory for the FTP site content. Create subdirectories to store other content as needed. To organize home directories for multiple FTP sites on the same server, you can create a top-level directory for storing all home directories, and then create subdirectories for each site. You can create a home directory on the local server, or as a uniform naming convention (UNC) path on a network share. You can also create virtual directories that map to the physical directories.

**Adding an FTP Site**

IIS provides the following two methods for adding a new FTP site:

- IIS Manager.

- The Iisftp.vbs command-line administration script.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To add an FTP site through IIS Manager**

1. In IIS Manager, double-click the local computer, right-click the **FTP Sites** folder or a specific FTP site folder, point to **New**, and then click **FTP Site**.

2. In the FTP Site Creation Wizard, click **Next**.

3. In the **Description** box, type the name that you have selected for the FTP site, and then click **Next**.

4. Under **Enter the IP address to use for this FTP site**, click **All Unassigned** to allow FTP to respond to all unassigned IP addresses on the server, or click a specific IP address for the site.

   **Note**

In the FTP Site Creation Wizard, **All Unassigned** refers to IP addresses that are assigned to a computer but not assigned to a specific site. The default FTP site uses all of the IP addresses that are not assigned to other sites. Only one site can be set to use unassigned IP addresses for a given port number. For more information about how sites are identified, see Creating Multiple FTP Sites.

5. The TCP port is assigned to port 21 by default. To use a nonstandard port number to create a unique identity for a private FTP site, under **Type the TCP port for this FTP site**, type a new port number above 1023, and then click **Next**.
6. Click the appropriate user isolation method for this site, and then click **Next**.
7. In the **Path** box, type or browse to the path of your FTP site home directory, and then click **Next**.
8. Select **Read** permissions, or select **Read** and **Write** permissions for the site.
9. Click **Next**, and then click **Finish** to close the FTP Site Creation Wizard.

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters).

**To add an FTP site by using the Iisftp.vbs command-line administration script**

1. If the new FTP site will be in a new directory, create the new directory.
2. From the **Start** menu, click **Run**.
3. In the **Open** box, type **cmd**, and then click **OK**.
4. At the command prompt, switch to the **systemroot**\System32 directory.
5. Type the following, and then press ENTER:

   **cscript iisftp.vbs /create** *HomeDirectory* **"** *SiteDescription* **"** *IPAddress* **/b** *PortNumber*

   For example, the following command creates and starts an FTP site called MyFTPSite on port 21, which is at IP address 123.456.789 with C:\Inetpub\Ftproot\NewDirectory as its home directory:

   **cscript iisftp.vbs /create c:\inetpub\ftproot\newdirectory "myftpsite" /i 123.456.789 /b 21**.

**Note**

The Iisftp.vbs script supports many more options for creating an FTP site. To see the options available, search for "iisftp.vbs" in Help and Support Center for Windows Server 2003.

**Related Information**

- For information about installing the FTP service, see FTP Site Setup.
- For more information about creating a home directory, see Changing FTP Site Home Directories.
- For information about how to change properties on existing FTP sites, see Changing Default FTP Site Settings.
- For information about how to add virtual directories, see Using Virtual Directories with FTP Sites.
- For information to help you decide which user isolation method is appropriate for your site, see Isolating FTP Users.
- For information about how to perform common tasks from command-line scripts, see Using Command-Line Administration Scripts.

# Naming FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

A server running IIS can host multiple FTP sites. Each site runs on a virtual server, and each FTP site has a descriptive name.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To name an FTP site**

1. In IIS Manager, expand the **FTP Sites** folder, right-click the FTP site you want to name, and click **Properties**.
2. On the **FTP Site** tab, in the **Description** box, type a descriptive name for the site.
3. Click **OK**.

**Related Information**

- For information about installing the FTP service, see FTP Site Setup.
- For instructions about how to add new FTP sites, see Adding FTP Sites to Your Server.

# Using Virtual Directories with FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

Overview and procedures for creating and deleting virtual directories for FTP sites using IIS Manager and command-line scripts.

A virtual directory is a friendly name, or alias, either for a physical directory on your server hard disk that does not usually reside under the FTP site's home directory. The use of aliases is secure because users do not know where your files are physically located on the server relative to the FTP site's home directory, and therefore cannot use that information to attempt to modify your files. Aliases also make it easier for you to move directories in your site. Rather than changing the URL for the directory, you need only change the mapping between the alias and the physical location of the directory.

Another advantage of aliases is that you can publish content in multiple directories, to be accessed by all users, and control read/write permissions for each virtual directory separately. Even with User Isolation enabled, you can share public content by creating a virtual directory with access permission for all users.

If your FTP site contains files that are located in a directory other than the home directory, or on other computers, you must create virtual directories to include those files in your FTP site. To create a virtual directory that points to a physical directory on another computer, you must specify the full UNC path to the directory, and provide a user name and password for user rights.

If you want to publish from any directory not contained within your home directory, you must create a virtual directory. For example, suppose you are setting up a FTP site for your marketing team on your company's intranet. The following table shows an example of how you might construct the mapping between the physical location of the files and the URL that accesses the files.

| URL | Physical Location | Alias |
|---|---|---|
| ftp://SampleFTPSite | C:\Inetpub\Ftproot | home directory |
| ftp://SampleFTPSite/PR | C:\Marketing\PublicRel | PR |
| ftp://SampleFTPSite/PR/OldPR | C:\Documents\Old | OldPR |
| ftp://SampleFTPSite/PRPublic | C:\Documents\Public | PRPublic |
| ftp://SampleWebSite/Customers | //Server2/SalesData | Customers |

Both virtual directories and physical directories (directories without an alias) appear in IIS Manager. A modified folder icon indicates a virtual directory. The following illustration shows the FTP site example described above, where Customers, PR, and PRPublic are virtual directories.




You can also define nested virtual directories. For example, within the PR virtual directory, you can create another virtual directory, for example, named Oldpr. The content in Oldpr is accessible using the URL ftp://SampleFTPSite/PR/Oldpr.

**Note**

The FTP server does not include the names of any virtual directories when the content of the current working directory is requested by a client application. Physical directories, however, are listed.

**Important**

Access to the content of a virtual directory depends on the current working directory. For example, if a user is connected to the root directory of SampleFTPSite, assuming that the user has the necessary permissions, the user can change the working directory to the following paths: /PRPublic, /PR, /PR/Oldpr, and /Customers only if the user knows the site structure. For example, if the current working directory is /PRPublic, the user will receive an error message when attempting to change the working directory to /PR by typing the command **cwd /PR**, but will successfully change the working directory to /PR by typing the command **cwd ./PR**.

Also, a user can use a combination of physical directory names and virtual directory names to change the working directory. For example, assume that under the physical path to which the PR virtual directory is mapped, there is a physical folder named Test, and the current working directory is the root of SampleFTPSite. The user can change the working directory to the physical Test folder using the command **cwd PR/Test**.

For a basic FTP site, you can place all of your files in the site's home directory, and most likely do not need to add virtual directories. If you have a complex site, or want to specify different URLs for different parts of your site, you can add virtual directories as needed. To make a virtual directory accessible from multiple sites, you must add the virtual directory to each site.

There are two ways to create or delete a virtual directory:

- Using IIS Manager.
- Using the iisftpdr.vbs administration script.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To create a virtual directory using IIS Manager**

1. In IIS Manager, expand the local computer, expand the **FTP Sites** folder, expand the FTP site to which you want to add a virtual directory, right-click the site or folder within which you wish to create the virtual directory, point to **New**, and then click **Virtual Directory**.
2. Click **Next**.
3. In the **Alias** box, type a name for the virtual directory, and click **Next**. This is the name the user types, and should be short and easy to type.
4. In the **Directory** box, type or browse to the physical directory in which the virtual directory resides, and click **Next**.
5. Under **Allow the following permissions**, select the check boxes next to the access permissions appropriate to your needs, and click **Next**.
6. Click **Finish**. The virtual directory is created below the currently selected folder level.

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe***ScriptName* (include the script's full path and any parameters).

**To create a virtual directory using the Iisftpdr.vbs administration script**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. At the command prompt, type **cscript iisftpdr.vbs /create** *SampleFTPSite VirtualDirectoryName x*:****_path_, and press ENTER. Substitute your FTP site name, virtual directory name, and path to the physical directory as appropriate.

| **Note** |
| --- |
| Deleting a virtual directory does not delete the corresponding physical directory or files. |

**To delete a virtual directory using IIS Manager**

- In IIS Manager, expand the FTP site containing the virtual directory you want to delete, right-click the virtual directory, click **Delete**, and then click **Yes**.

| **Note** |
| --- |
| The below method does not work for root virtual directories. |

**To delete a virtual directory using the Iisftpdr.vbs administration script**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. At the command prompt, type **cscript iisftpdr.vbs /delete** *SampleFTPSite VirtualDirectoryName*, and press ENTER. Substitute Your FTP site name and virtual directory name as appropriate.

**Related Information**

- For information about installing the FTP service, see FTP Site Setup.
- For information about Security for virtual directories, see Securing Virtual Directories, and see Access Control with IIS 6.0.

# Changing Default FTP Site Settings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

During installation of IIS, default values are set for the various properties assigned to FTP sites. Properties can be set at the global level, affecting all FTP sites on a server; on the individual site level; or at the directory or file level. IIS uses an inheritance model, which means that settings at higher levels are automatically inherited by lower levels. Settings at lower levels can be edited individually to override inherited settings from the next level up.

If you change a setting at a lower level, then later change a setting at a higher level that conflicts with the lower level setting, you are prompted to choose whether you want to change the lower-level setting to the new higher-level setting.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To globally set properties for all FTP sites**

1. In IIS Manager, expand the local computer, right-click the **FTP Sites** folder, and click **Properties**.
2. Click **Help** on any of the tabs for more information about those settings.

**To set properties for a single FTP site**

1. In IIS Manager, expand the **FTP Sites** folder, right-click the FTP site whose settings you want to change, and click **Properties**.
2. Click **Help** on any of the tabs for more information about those settings.

**Related Information**

- For information about installing the FTP service, see FTP Site Setup.

# Creating Multiple FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

You can create multiple FTP sites using multiple IP addresses and multiple ports. While creating multiple sites with multiple IP addresses is a common and recommended practice, it can be more complicated because, by default, clients call port 21 when using the FTP protocol. Therefore, if you create multiple FTP sites using multiple ports, you need to inform users of the new port number so their FTP clients can locate and connect to the port.

If you create a new site using the same port as an existing site with the same IP address, the new site will not start. The general rule is that you can have multiple sites using the same IP and port, but only one site from this group can run at a time. If you try to start another site from this group, you receive an error message.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create multiple FTP sites using multiple IP addresses**

1. In IIS Manager, expand the local computer, right-click the **FTP Sites** folder, point to **New**, and click **FTP Site**.
2. Click **Next**.
3. In the **Description** box, type a description of your FTP site, and then click **Next**.
4. Under **Enter the IP address to use this FTP site**, type a new IP address, and leave the TCP port setting at **21**.
5. Complete the rest of the **FTP Site Creation Wizard**.

**To create multiple FTP sites using multiple ports**

1. In IIS Manager, expand the local computer, right-click the **FTP Sites** folder, point to **New**, and click **FTP Site**.
2. Click **Next**.
3. In the **Description** box, type a description of your FTP site, and then click **Next**.
4. Under **Enter the IP address to use for this FTP site**, type your Web server's IP address.
5. Under **Type the TCP port for this FTP site**, change the TCP port from the default setting of 21 to an unallocated port number. If you are uncertain which port numbers are already allocated, choose a high number, for example a number between 5000 and 5999.
6. Complete the rest of the **FTP Site Creation Wizard**.

**Related Information**

- For information about installing the FTP service, see FTP Site Setup.

# Stopping and Starting FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

The need to stop and restart an FTP site typically indicates a problem with the site. If your FTP site is slow, unresponsive, or is not running properly in some other way, stopping and then restarting the site can resolve a temporary problem with the software.

By default, FTP sites start automatically when your computer restarts. Starting a site restarts or resumes a specific site, or the FTP service if it was not previously running, which might immediately fix problems. Stopping a site disables the specific FTP site, disconnecting all client connections to the site. Stopping the FTP service stops all FTP sites and unloads the FTP service. Pausing a site prevents the FTP service from accepting new connections but does not affect requests that are already being processed.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To start, stop, or pause an FTP site**

- In IIS Manager, expand the local computer, expand the **FTP Sites** folder, right-click the site you want to start, stop, or pause, and click **Start**, **Stop**, or **Pause**.

If a site stops unexpectedly, IIS Manager might not correctly indicate the state of the site. In IIS Manager, right-click the **FTP Sites** folder and click **Refresh** to see the current state of all FTP sites.

**Related Information**

- For information about installing the FTP service, see FTP Site Setup.
- For information about restarting all of your Internet services,see Restarting IIS.
- To change properties on existing FTP sites, see Changing Default FTP Site Settings.

# Setting FTP Messages and Directory Output Style

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

By default, the IIS FTP service sends a banner message that identifies the FTP service, and the fact that FTP runs in IIS, to the client. The message can be a greeting that welcomes users to the FTP site when they log on, an exit message (see note) that users see when they log off, a message that informs users when the maximum number of connections has been reached, or a banner message. Also, you can send the client informational messages about your FTP sites. The message can be a greeting, an exit message, or information about connection status. These messages are blank by default.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To suppress the default FTP banner for all FTP sites on a computer**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd** and click **OK**.
3. At the prompt, type the following command and then press ENTER:

   **Cscript.exe adsutil.vbs SET MSFTPSVC/SuppressDefaultFTPBanner true**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To globally set messages on all FTP sites**

1. In IIS Manager, expand the local computer, right-click the **FTP Sites** folder, and then click **Properties**.
2. Click the **Messages** tab.
3. Type messages in any of, or all of, the **Banner**, **Welcome**, **Exit**, and **Maximum connections** boxes. These messages will appear on all FTP sites.
4. Click **Apply**, and then click **OK**.

**To set messages on a specific FTP site**

1. In IIS Manager, right-click the FTP site for which you want to change messages, and then click **Properties**.
2. Click the **Messages** tab.
3. Type messages in any of, or all of, the **Banner**, **Welcome**, **Exit**, and **Maximum connections** boxes. Messages set at the individual FTP site level override messages set at the global level.
4. Click **Apply**, and then click **OK**.

**Setting the Directory Output Style**

You can set the directory output style of all FTP sites, or an individual FTP site to be in either MS-DOS® or UNIX format. The directory output style defines the format that is used when an FTP client sends a **/list** command. The default directory output

style is MS-DOS format.

| Note |
| --- |
| Although the exit message box can contain up to 30,000 characters, the FTP exit message size is limited to 1024 characters, including the reply code, blank, carriage return, and null. To prevent the exit message from being truncated, limit the length of the exit message to 1017 characters. |

The UNIX directory listing style displays the year in a four-digit format when the year in the date of the file is different from the year in the date of the FTP server. No year is returned in cases where the year in the date of the file is the same as the year in the date of the FTP server.

**To set the global FTP directory output style**

1. In IIS Manager, right-click the **FTP Sites** folder, and then click **Properties**.
2. Click the **Home Directory** tab.
3. Under **Directory listing** style, click **UNIX** or **MS-DOS®**.
4. Click **Apply**, and then click **OK**.

**To set the FTP directory output style for an individual FTP site**

1. In IIS Manager, expand the **FTP Sites** folder, right-click the FTP site for which you want to change the directory listing style, and then click **Properties**.
2. Click the **Home Directory** tab.
3. Under **Directory listing style**, click **UNIX or MS-DOS®**. The output style set at the individual FTP site level overrides the output style set at the global level.
4. Click **Apply**, and then click **OK**.

| Note |
| --- |
| The MS-DOS directory listing style displays the year in a two-digit format by default. You can change this setting to display the year in a four-digit format. |

**Related Information**

- For information about installing the FTP service, see FTP Site Setup.

# Changing FTP Site Home Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

Each FTP site on a computer must have its own home directory. The default home directory for the default FTP site is *LocalDrive*:\Inetpub\Ftproot.

There are two ways to change the home directory of an FTP site:

- Use IIS Manager
- Edit the MetaBase.xml file directly.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To change the home directory of an FTP site using IIS Manager**

1. In IIS Manager, expand the local computer, expand the **FTP Sites** folder, right-click the FTP site whose home directory you want to change, and click **Stop**.
2. Right-click the FTP site again, and click **Properties**.
3. Click the **Home Directory** tab.
4. Under **The content for this resource should come from**, click the location of your home directory.

   **Note**

   If you select a directory on a network share, you might need to enter a user name and password to access the resource. IUSR_computername is the default account used if another account is not specified. If you use an account with administrative credentials on the server, clients can gain access to server operations. This seriously jeopardizes the security of your network. For more information about security see, Security Best Practices in Windows Help.

5. Click **Apply**, and then click **OK**.
6. Right-click the FTP site you've just changed, and click **Start**.
7. Optional: Move or copy the content for the FTP site from the original home directory to the new home directory.

**To change the home directory of an FTP site by editing the MetaBase.xml file**

1. In IIS Manager, right-click the local computer, point to **All Tasks**, and click Select **Restart IIS...**.
2. In the **Stop/Start/Reset** dialog box, click **Stop Internet Services**, and click **OK**.
3. Expand the **FTP Sites** folder, right-click the FTP site you want to change, and click **Stop**.
4. In Windows Explorer, from the *localdrive*\System32\Inetsrv folder, open the MetaBase.xml file in a text editor.
5. To change the FTP home directory, search for "/msftpsvc/*1*/root" and change the **Path=** value to the new location of the FTP home directory. Substitute the value *1* with the value corrresponding to the FTP site you want to change.
6. In IIS Manager, right-click the local computer, point to **All Tasks**, and click **Select Restart IIS...**.

7. In the **Stop/Start/Reset** dialog box, click **Start Internet Services**, and click **OK**.
8. Optional: Move or copy the content for the FTP site from the original home directory to the new home directory.

**Related Information**

- For information about installing the FTP service, see FTP Site Setup.

# Securing FTP Sites with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

FTP is commonly misunderstood as a secure means for transferring data, because the FTP server can be configured to require a valid user name and password combination prior to granting access. Be aware that neither the credentials specified at logon nor the data itself is encrypted or encoded in any way. All credentials are sent across the network in plaintext. In other words, all FTP data can be easily intercepted and analyzed by any station on any network between the FTP client and FTP server. The risk of plaintext credentials is that someone other than the intended users could log on to FTP and download the files you have placed there. If you intend to place sensitive data on your FTP site or if secure communication between clients and your FTP server are important, consider using FTP over an encrypted channel such as a Virtual Private Network secured with Point-to-Point Tunneling Protocol or Secure Internet Protocol (IPSec). You should also consider using WebDAV, which utilizes Secure Sockets Layer (SSL).

| **Caution** |
| --- |
| FTP sites or virtual directories that are configured to use Active Directory isolation or FTP load balancing should not be mapped to physical directories that are used for Web sites that use FrontPage® Server Extensions from Microsoft. Doing so can allow users to view any files in that folder structure over the network. |

- For information about installing the FTP service, see FTP Site Setup.
- To learn how to set up a private network over the Internet using Point-to-Point Tunneling Protocol (PPTP) or how to set up secure IP communications with a Virtual Private Network (VPN) and IPSec, see "PPTP", "VPN", and "IPSec" in Help and Support Center for Windows Server 2003.

# Isolating FTP Users

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

FTP user isolation is a solution for Internet service providers (ISPs) and application service providers who want to offer their customers individual FTP directories for uploading files and Web content. FTP user isolation prevents users from viewing or overwriting other users' Web content by restricting users to their own directories. Users cannot navigate higher up the directory tree because the top-level directory appears as the root of the FTP service. Within their specific site, users have the ability to create, modify, and delete files and folders.

FTP user isolation is a site property, not a server property. It can be turned on or off for each FTP site.

FTP user isolation supports the following three isolation modes, each of which enables different levels of isolation and authentication:

- Do not isolate users: This mode does not enable FTP user isolation. It is designed to work similarly to earlier versions of IIS.
- Isolate users: This mode authenticates users against local or domain accounts before they can access the home directory that matches their user name.
- Isolate users using Active Directory: This mode authenticates user credentials against a corresponding Active Directory container, rather than searching the entire Active Directory, which requires large amounts of processing time.

| **Note** |
| --- |
| This mode requires an Active Directory server running on an operating system in the Windows Server 2003 family. A Windows 2000 Active Directory can also be used but requires manual extension of the User Object schema. To learn more about setting up an Active Directory server, see Help and Support Center for Windows Server 2003. |

**In this section:**

- Creating a New FTP Site with Isolate Users Mode: Describes how to create a new FTP site with **Isolate users** mode.
- Converting an Existing FTP Site to Isolate Users Mode: Describes how to convert an existing FTP site to **Isolate users** mode.
- Creating a New FTP Site with Isolate Users Using Active Directory Mode: Describes how to create a new FTP site with **Isolate users using Active Directory** mode.
- Enabling Anonymous Access for an FTP Site Configured with Isolate Users Using Active Directory Mode: Describes how to enable anonymous access for an FTP site configured with **Isolate users using Active Directory** mode.
- Converting an Existing FTP Site to Isolate Users Using Active Directory Mode: Describes how to prepare and convert an existing FTP site to **Isolate users using Active Directory** mode.

**Related Information**

- For more information about FTP user isolation modes, see Hosting Multiple FTP Sites with FTP User Isolation.
- For information about installing the FTP service, see FTP Site Setup.

# Creating a New FTP Site with Isolate Users Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

**Isolate users** mode authenticates users against local or domain accounts before they can access the home directory that matches their user name. All user home directories are in a directory structure under a single FTP root directory where each user is placed and restricted to their home directory. Users are not permitted to navigate out of their home directory. If users need access to dedicated shared folders, you can also establish a virtual root. This mode does not authenticate against Active Directory directory service.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To create a new FTP site with Isolate users mode**

1. In IIS Manager, click the local computer, right-click the **FTP Sites** folder, point to **New**, and then click **FTP Site**.
2. In **Welcome to the FTP Site Creation Wizard**, click **Next**.
3. In **FTP Site Description**, type a description for the FTP site, and then click **Next**.
4. In **IP Address and Port Settings**, type an IP address and port number, and then click **Next**.
5. In **FTP User Isolation**, click **Isolate users**, and then click **Next**.
6. In **FTP Site Home Directory**, type or browse to the FTP site home directory, and then click **Next**.
7. Enable the Read and Write permissions as appropriate, and then click **Next**, and click **Finish**.

| **Important** |
| --- |
| Adhere to the following convention when creating each user's FTP site directories: |

- If anonymous access is allowed, create the subdirectories **LocalUser** and **LocalUser\Public**under the FTP site home directory.
- If users of the local computer log on with their individual account user names (rather than as anonymous users), create the subdirectory **LocalUser**and a separate directory **LocalUser***UserName* under the FTP site home directory for each user allowed to connect to the FTP site.
- If users of different domains log on with their explicit *Domain**UserName* credentials, create a subdirectory for each domain (by using the name of the domain) under the FTP site root directory. Under each domain directory, create a directory for each user. For example, to support access by user Contoso\user1, create the Contoso and Contoso\user1 directories.

**Related Information**

# Converting an Existing FTP Site to Isolate Users Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

You can convert an existing FTP site to **Isolate users** mode after you upgrade to Windows Server 2003.

**To convert an existing FTP site to Isolate users mode**

1. In IIS Manager, click the local computer, double-click the **FTP Sites** folder, right-click the FTP site that you want to convert, and then click **Stop**.

| **Important** |
| --- |
| You must stop the FTP service because you need to move user directories to new locations. |

2. If users log on to the FTP site with local machine accounts or if you allow anonymous access to the FTP site, complete the following steps:
   a. Under the FTP home directory, create a subdirectory named **LocalUser**.
   b. If you allow anonymous access to the site, create a subdirectory under **LocalUser** named **Public**. This will be the home directory for anonymous users.
   c. Move the content for the anonymous user from the existing location to the new **Public** subdirectory.
   d. For each FTP user with a local computer account, create a subdirectory under **LocalUser** with the user account name. Move content from existing user directories to these new directories.
3. If domain users access the site, complete all of the following steps:
   a. Create a subdirectory under the FTP home directory for each domain for which users are accessing the site. The subdirectories must be named by the domain names.
   b. For each user, create a subdirectory under the corresponding domain-named subdirectory, and move existing content into the new subdirectories.
4. To finish the conversion, do one of the following:
   - Use the FTP Site Creation Wizard to create a new FTP site, and point the site path to the FTP home directory.
   - Edit the metabase, set the **UserIsolationMode** metabase property to **1** for the site and the path in the corresponding **/Root**virtual root section to the new site home directory.

| **Important** |
| --- |
| Before you edit the metabase, verify that you have a backup copy that you can restore if a problem occurs. |

5. In IIS Manager, click the local computer, double-click the **FTP Sites** folder, right-click the converted FTP site, and then click **Start**.

**Related Information**

# Creating a New FTP Site with Isolate Users Using Active Directory Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

**Isolate users using Active Directory** mode authenticates user credentials against a corresponding Active Directory container, rather than searching the entire Active Directory, which requires large amounts of processing time.

| **Note** |
| --- |
| This mode requires an Active Directory server running on an operating system in the Windows Server 2003 family. A Windows 2000 Active Directory can also be used but requires manual extension of the User Object schema. To learn more about setting up an Active Directory server, see Help and Support Center for Windows Server 2003. |

Specific FTP server instances can be dedicated to each customer to ensure data integrity and isolation. When a user's object is located within the Active Directory container, the **msIIS-FTPRoot** and **msIIS-FTPDir** properties are extracted to provide the full path to the user's home directory. If the FTP service can successfully access the path, the user is placed within the home directory, which represents the FTP root location. The user sees only their FTP root location and is, therefore, restricted from navigating higher up the directory tree. The user is denied access if either the **msIIS-FTPRoot** or **msIIS-FTPDir** property do not exist, or, if these two together do not form a valid and accessible path.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas/user:***Administrative_AccountName***"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To create FTP sites with Isolate users using Active Directory mode**

1. In IIS Manager, click the local computer, right-click the FTP Sites folder, point to New, and click FTP Site.
2. In the **Welcome to the FTP Site Creation Wizard**, click **Next**.
3. In **FTP Site Description**, type a description for the FTP site, and then click **Next**.
4. In **IP Address and Port Settings**, type an IP address and port, and then click **Next**.
5. In **FTP User Isolation**, click **Isolate users using Active Directory**, and then click **Next**.
6. In the **User name** text box, type the user name, using the *Domain\User* format, or browse to the user name. Choose a user with minimal domain privileges. This user name is used to access Active Directory and read the home directory properties.
7. In the **Password** text box, type the password of the user.
8. In the **Enter the default Active Directory domain** text box, type or browse to the default domain name.

   | **Note** |
   | --- |
   | This domain name is used for the users who do not specify their user domain when they log on. In other words, a user connecting with the user name Domain1\User1 is authenticated against Domain1, while a user connecting as User2 is authenticated against the default logon domain. If a default domain is not named and a user does not specify a domain name, access is denied for all but anonymous users. Type the base domain name only, not the fully qualified name. For example, type **MyDomain**, not **MyDomain.dept.microsoft.com**. |

9. Click **Next**. You are prompted to re-enter the password for the user entered in the previous steps.
10. Enable the Read and Write permissions as appropriate, and then click **Next**, and click **Finish**.

**Related Information**

# Enabling Anonymous Access for an FTP Site Configured with Isolate Users Using Active Directory Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

By default, anonymous access is disabled to sites created in **Isolate users using Active Directory** mode.

**To enable anonymous access for the FTP site configured with Isolate users using Active Directory mode**

- Configure the metabase properties as shown in the following example. You can do so by using the adsutil.vbs SET command-line tool.
    - **adsutil set /msftpsvc/6634/AllowAnonymous TRUE**
    - **adsutil set /msftpsvc/6634/AnonymousOnly FALSE**
    - **adsutil set /msftpsvc/6634/AnonymousUserName MyDomain\LowPrivUser**
    - **adsutil set /msftpsvc/6634/AnonymousUserPass PaSsWoRd**

| **Note** |
| --- |
| When a site is created with Isolate Users Using Active Directory mode, the **Path** property of the root FTP virtual directory (which, for the other isolation modes identifies the home directory) is set empty. Also, the **AccessFlags** property of the root FTP virtual directory contains the **AccessNoPhysicalDir** flag. Do not alter these two values. If you change or remove them, further access to the site is not allowed. |

**Related Information**

- For more information about FTP user isolation, see Hosting Multiple FTP Sites with FTP User Isolation.
- For more information about using adsutil.vbs, see Using the Adsutil.vbs Administration Script.

# Converting an Existing FTP Site to Isolate Users Using Active Directory Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

You can convert an existing FTP site to **Isolate users using Active Directory** mode after you upgrade to Windows Server 2003. Before you convert an existing FTP site to the **Isolate users using Active Directory** mode, complete the following procedure.

**To prepare to convert an existing FTP site to Isolate users using Active Directory mode**

1. Plan user distribution over your network file server resources.
2. Create server shares and user directories for all users that have access to FTP.
3. Make sure to create a directory for anonymous access (if you want to enable anonymous user connections).
4. For each user, set the **msIIS-FTPRoot** and **msIIS-FTPDir** properties in Active Directory to point to the new home directory using Iisftp.vbs, by typing the following at the command line:

   **Iisftp.vbs /SetADProp** *UserName* **FTPRoot** *Server* **** *Share*

   **Iisftp.vbs /SetADProp** *UserName* **FTPDir** *Directory*

   For complete Iisftp.vbs syntax, type **Iisftp.vbs /SetADProp /?**at the command prompt.

You can complete the conversion by using the FTP Site Creation Wizard (recommended), or by following the procedure below to configure the metabase.

**To convert an existing FTP site to Isolate users using Active Directory mode**

1. In IIS Manager, click the **local computer**, double-click the **FTP Sites** folder, right-click the FTP site that you want to convert, and then click **Stop**.
2. Move existing user content into the new directories.
3. Edit the metabase directory or use adsutil.vbs to configure the following metabase properties:

   - **UserIsolationMode**: Set to **2**.
   - **ADConnectionUserName**: Set to a user that has permission to read Active Directory properties. Use *Domain**UserName*format.
   - **ADConnectionPassword**: Set to the password for the user in **ADConnectionUserName**.
   - **DefaultLogonDomain**: Set to the default domain name.
   - For the site root virtual directory: Set the **Path** property to an empty string, and add the value **AccessNoPhysicalDir**to the **AccessFlags**property, using the **|** operator; for example: **AccessFlags=AccessRead|AccessNoPhysicalDir**.
   - **AllowAnonymous**, **AnonymousUserName**, and **AnonymousPassword**: See To enable anonymous access for the Isolate Users Using Active Directory FTP site earlier in this topic for information about how to set these properties.

   | **Important** |
   | --- |
   | Before you edit the metabase, verify that you have a backup copy that you can restore if a problem occurs. |

4. In IIS Manager, click the **local computer**, double-click the **FTP Sites** folder, right-click the converted FTP site, and then

click **Start**.

**Related Information**

# Modes and Data Transmission

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

When using FTP, there are two things to consider. The first is the connection mode, which can be either passive or active; and the second is the data transfer mode, which can be stream, block, or compressed.

The IIS-based FTP service supports both active and passive mode connections, depending on the method that is specified by the client. IIS does not support disabling either active or passive connection modes.

The default data transfer mode for IIS-based FTP is stream. IIS does not currently support block or compressed data transfer mode.

Unlike HTTP and other protocols used on the Internet, the FTP protocol uses a minimum of two connections during a session: a half-duplex connection for control, and a full-duplex connection for data transfer. By default, TCP port 21 is used on the server for the control connection, but the data connection is determined by the method that the client uses to connect to the server, as detailed below.

- *Active-mode FTP connections* are sometimes referred to as "client-managed" because the client sends a **port** command to the server, over the control connection. The command requests the server to establish a data connection from TCP port 20 on the server to the client, using the TCP port that is specified by the **port** command.

- *Passive-mode FTP connections* are sometimes referred to as "server-managed", because after the client issues a **pasv** command, the server responds with one of its transient ports used as the server-side port of the data connection. After a data connection command is issued by the client, the server connects to the client using the port immediately above the client-side port of the control connection.

**Data Transfer and Firewalls**

The most common problem encountered with FTP over the Internet involves data transfers through a proxy server, a firewall, or a Network Address Translation (NAT) device. In most cases these network security devices allow the control connection to be established over TCP port 21 (that is, the user can successfully log on to the FTP server), but when the user attempts a data transfer such as DIR, LS, GET, or PUT, the FTP client appears to stop responding because the network security device is blocking the data connection port that is specified by the client. If the network security device supports logging, you can verify port blocking by reviewing the deny/reject logs on the network security device.

The following table includes FTP clients provided by Microsoft, and the connection mode supported by each client.

| FTP Client | Transfer Mode |
|---|---|
| Command line | Active |
| Internet Explorer 5.1 and earlier | Passive |
| Internet Explorer 5.5 and later | Active and Passive |
| FrontPage 1.1 to FrontPage 2002 | Active |

- For information about installing the FTP service, see FTP Site Setup.

- For more information about FTP data transmission, refer to the FTP protocol specification at the World Wide Web Consortium Web site.

- For information about setting the port range used for passive-mode FTP connections, see the PassivePortRange Metabase Property.

# Restarting FTP Uploads and Downloads

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

If your connection closes unexpectedly while uploading or downloading a file using FTP, it might not be necessary to transfer the entire file again, because the IIS -based FTP service supports restarting failed FTP uploads and downloads. If your FTP client software supports FTP restart, you can reestablish the FTP connection and restart the file transfer; and the file transfer will automatically pick up where it left off.

The command-line FTP client included with Microsoft Windows does not support restarting downloads. You must use a third-party command-line FTP client for that functionality. If you are using a command-line FTP client, uploads can be resumed by using the **append** command to add data to an existing file. Downloads can be restarted by first issuing a **rest** command with the desired offset, and then issuing the **retr** command.

| **Note** |
| --- |
| The IIS implementation of FTP restart is limited to an offset of less than four gigabytes. |

- For information about installing the FTP service, see FTP Site Setup.
- For information about the commands supported by the Microsoft Windows command-line FTP client, see the Microsoft Developer Network (MSDN) Web site.
- For the official FTP specification with the entire FTP command set, see the World Wide Web Consortium Web site.

# Server-to-Server FTP Transfer

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

You can configure IIS to allow server-to-server FTP transfers by changing the EnableDataConnTo3rdIP and EnablePasvConnFrom3rdIP keys in the system registry. By default, the EnablePasvConnFrom3rdIP registry key is disabled to prevent malicious attacks against the FTP service. Although details for configuring server-to-server transfers are listed below, Microsoft does not recommend enabling the EnablePasvConnFrom3rdIP registry key on IIS-based FTP servers connected to the Internet.

| **Important** |
| --- |
| Using Registry Editor incorrectly can cause serious problems that require reinstalling the operating system. Because Registry Editor bypasses the standard safeguards that prevent you from entering settings that are conflicting or likely to degrade performance or damage your system, exercise caution when making changes to the registry. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. For information about how to edit the registry, see the "Changing Keys and Values" topic in Registry Editor Help. |

The procedure below lists the steps required to transfer a binary file (called MyFile.xls) from Server1 (IP Address: 192.168.0.10) to Server2 (IP Address: 192.168.1.20) using the command-line FTP client on Workstation1. You must have Write permissions granted on Server2 for this example to work. You must also have two computers running the IIS-based FTP service.

The server that initiates the data channel (Server1 in this

example) must have the following value set in the registry:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSFTPSVC\Parameters\EnableDataConnTo3rdIP=1

Server2 must have the following value set in the registry:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSFTPSVC\Parameters\EnablePasvConnFrom3rdIP=1

In normal client/server FTP transfers, the user is only responsible for entering "friendly" commands to the server; the client responds automatically. When a server-to-server FTP transfers is performed, another server is the client. Therefore, the user is responsible for sending all FTP commands to both servers (in RFC959 format, preceded by the "literal" keyword).

**Procedures**
**To transfer a binary file between FTP servers**

1. Connect to Server1 using a command-line FTP client. (This connection will be referred to as Client1).
2. Connect to Server2 using a command-line FTP client. (This connection will be referred to as Client2).
3. Issue the following command in Client2: **literal pasv**.

   The computer responds with: 227 Entering Passive Mode (192,168,1,20,15,61).

   Note that 15,61 represents the low and high order bytes of the port number (port 3901 in this case), which is randomly assigned with each FTP connection. As a result, the last two numbers of this string will be different.
4. Issue the following command in Client1: **literal port 192,168,1,20,15,61**.

   The computer responds with: 200 Port Command Successful.

   Use the IP and port string returned in step 3.
5. Issue the following command in both Client1 and Client2: **literal type i**.

   The computer responds with: 200.
6. Type **set to I**.
7. Issue the following command in Client1: **literal retr myfile.bin**.

The computer responds with: 125 Data Connection already open; Transfer Starting.

8. Issue the following command in Client2: **literal stor myfile.bin**.

   The computer responds with: Opening BINARY mode data connection for myfile.bin.

9. Continue issuing the **literal stat** command on Client2 until it returns a 211 response indicating "No data connection is listed".

   This 211 response message is several lines long and indicates that the transfer is no longer in progress.

10. To transfer files from Server2 to Server1, issue the **literal retr** command from Client1 and the **literal stor** command from Client2.

**File Transfer Alternatives**

If you are concerned about FTP security, or if a network security device like a firewall prevents server-to-server data transmission, consider using a file transfer alternative like Web authoring with WebDAV. WebDAV is the recognized standard for HTTP file transfers.

**Related Information**

- For information about **port**, **pasv**, **stor**, **retr**, **stat**, and other literal FTP commands, see the World Wide Web Consortium Web site.

# Redirecting FTP Requests to Directories or Network Shares

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

You can redirect requests for files in one directory to a different directory or to a network share. When an FTP client requests the file from the original URL, the FTP server instructs the client to request the file from the redirected resource. Redirection is useful when users have grown accustomed to using a particular URL to access certain files.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the runas command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To redirect requests to another directory or network share**

1. In IIS Manager, expand the **FTP Sites** folder, right-click the FTP site or directory within an FTP site to which you want to add a redirect, and click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab.
3. Under **The content for this resource should come from**, click the location from which the files should be served.
4. In the **Local path** box, type or browse to the path from which you want the files served.
5. Click **OK**.

**Making an FTP virtual directory visible to users**

You can configure your FTP site so that a user browsing the site sees a directory entry for a virtual directory. While the directory entry the user sees is a physical directory, when the user navigates into the physical directory, the contents listed are actually of a virtual directory you have specified. This is possible because IIS lists virtual paths before listing physical paths. This makes it appear to the user that they are seeing the contents of a physical directory, when in fact they are seeing the content in a virtual directory.

**To make an FTP virtual directory visible to users**

1. Create a physical directory in the user's home FTP directory. For example, if your FTP site is configured for user isolation, the user's home directory will either be LocalUser\Public (for anonymous users), or LocalUser*username* (for authenticated users).
2. In IIS Manager, expand the local computer, expand the **FTP Sites** folder, expand the FTP site to which you want to add a virtual directory, right-click the site or folder within which you wish to create the virtual directory, point to **New**, and then click **Virtual Directory**.
3. Click **Next**.
4. In the **Alias** box, type the name for the virtual directory and click **Next**.

   | **Important** |
   | --- |
   | The alias must exactly match the name of the directory you created in step 1. |

5. In the **Directory** box, type the path to, or browse to the physical directory where the content for this virtual directory resides, and click **Next**.
6. Under **Allow the following permissions**, select the check boxes next to the access permissions appropriate to your needs, and click **Next**.

7. Click **Finish**. The virtual directory is created within the FTP site.

   When a user browses the FTP site, the physical directory you created in step 1 is displayed in the folder listing. When the user then navigates to the physical directory, however, the content of the redirected virtual directory is actually displayed

| **Note** |
| --- |
| You can also use this procedure with nested virtual directories. |

**Related Information**

- For information about how to set up a user's home FTP directory using user isolation, see Isolating FTP Users.
- For information about installing the FTP service, see FTP Site Setup.

# NNTP Administration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Newsgroups make it easy for users in your organization to engage in discussions on various topics with each other and with users outside of the organization. Whether enabling a departmental newsgroup or providing a full feed of Internet newsgroups, the principals of setting up a news server are the same.

This section includes the following information:

- Creating NNTP Virtual Servers: Describes multiple-server strategies and how to create virtual Network News Transfer Protocol (NNTP) servers.
- Creating Virtual Directories in IIS 6.0: Describes how to create and modify virtual directories.
- Configuring NNTP Virtual Servers: Introduces the settings options for NNTP virtual servers.
- Starting, Stopping, or Pausing NNTP Virtual Servers: Describes how to start, stop, and pause NNTP virtual servers and the difference between managing servers and services.
- Starting, Stopping, or Pausing the NNTP Service: Describes how to start, stop, and pause NNTP services and how it differs from managing virtual servers.
- NNTP Security: Introduces NNTP security concepts.
- Managing Newsgroups: Introduces concepts in managing newsgroups including how to create, moderate, edit, delete, and configure them.
- Configuring Expiration Policies: Describes how to set and modify newsgroup expiration policies.

# Creating NNTP Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can create as many NNTP virtual servers as you need. Each virtual server should have a unique IP address/TCP port combination. Although it is possible to have multiple NNTP virtual servers on a single IP address by using a different TCP port for each one, this configuration is not recommended because the client software must then be specially configured to use the nonstandard TCP port.

When you create an NNTP virtual server, you are prompted to enter two directory paths: a path for internal files and a path for newsgroup files. If possible, put these two paths on different disk partitions because the internal files can be corrupted if the disk partition runs out of disk space. Keeping them on a separate disk partition reduces the chance that the disk partition can run out of disk space.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a new NNTP virtual server**

1. In IIS Manager, expand the local computer, and then right-click an existing NNTP virtual server, point to **New**, and click **Virtual Server**.
2. Follow the instructions of the New NNTP Virtual Server Wizard. Be sure to select an IP address that is not being used by another NNTP virtual server.
3. The new NNTP virtual server will start automatically. However, if you have assigned it an IP address/TCP port combination that is already in use, it won't start. You must reconfigure the IP address/TCP port combination to one that is unique.
4. Configure the new NNTP virtual server.

**Related Information**

- For information about creating virtual directories, see Creating Virtual Directories in IIS 6.0.

# Creating Virtual Directories with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can create as many virtual directories as you want. Create virtual directories that match the structure of your newsgroup hierarchy.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a virtual directory**

1. In IIS Manager, expand the local computer, and then expand the NNTP virtual server that you want.
2. Right-click **Virtual Directories**, point to **New**, and then click **Virtual Directory**.
3. Follow the New NNTP Virtual Directory Wizard. Create a newsgroup subtree for this virtual directory.
4. Select either **File System** (a local location) or **Remote** as a storage location for the news content. This location can be configured later if necessary by clicking **Contents** in the **Virtual Directory** box.
5. If you selected **File System**, type a path to the local folder, or click **Browse** to select a folder. If you selected **Remote**, type a path to the remote location, or click **Browse** to select a remote folder. For a remote location, you must also provide a user name and password that Microsoft NNTP Service can use to log on to that remote computer. Type a name or click **Browse** to search for one. Click **Check** to ensure that the user name/password combination you plan to use is valid.

| **Note** |
| --- |
| If you store content in a remote location, the remote computer must be configured with at least read/write permissions for the user name you provide. Make sure that the remote location is secure and that it is maintained by trusted administrators. |

**To modify a virtual directory**

1. In IIS Manager, expand the local computer, and then expand the NNTP virtual server that you want.
2. Right-click the virtual directory that you want to modify, and then click **Properties**.
3. Change the options in the virtual directory's **Properties** box as needed.

**Related Information**

- For information about creating an NNTP virtual server, see Creating NNTP Virtual Servers.

# Configuring NNTP Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

For each Network News Transfer Protocol (NNTP) virtual server you can configure security, create newsgroups, set newsgroup limits, and so on. Most configuration options are assigned default values when a new NNTP virtual server is created.

Each NNTP virtual server is listed separately in IIS Manager. Each one has its own properties, expiration policies, and virtual directories. Be sure to select the correct NNTP virtual server before changing properties. To open the properties dialog box for an NNTP virtual server, right-click the NNTP virtual server, and then click **Properties**.

Besides creating newsgroups for the new NNTP virtual server, you must specify the following virtual server settings:

| Tab | Settings |
|---|---|
| Access | Configure client access to the NNTP virtual server and to establish transmission security. |
| General | Specify IP address, connection limits, path header, and logging settings. |
| Security | Add Windows user accounts and groups to the list of NNTP virtual server operators. |
| Settings | Configure a maximum limit for the size of news articles and news feeds, and specify the SMTP server to which moderated messages are sent. |

This section includes the following information:

- NNTP Access Tab: Describes the options that are available on the **Access** tab.
- NNTP General Tab: Describes the options that are available on the **General** tab.
- NNTP Security Tab: Describes the options that are available on the **Security** tab.
- NNTP Settings Tab: Describes the options that are available on the **Settings** tab.

# NNTP Access Tab

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the **Access** tab to restrict access to individual newsgroups or sets of newsgroups on the basis of the user name, or to restrict access to all newsgroups on the basis of the client computer IP address.

The **Access** tab contains the following options:

| Option | Description |
|---|---|
| Access control | Authentication options. You can choose one or more options. To view the options, click **Authentication**. |
| Allow anonymous | Unrestricted user access to newsgroups that do not have permissions set on their directories. User names and passwords are not required. |
| Anonymous | Opens the **Anonymous Account** dialog box, which allows you to select a Windows account that will be used when someone logs on anonymously. |
| Username | The user name of the Windows user account used by all anonymous connections. This account defines the privileges of anonymous users. To select an account already defined on this computer, click **Browse**. |
| Password | The password for the anonymous user account. |
| Enable automatic password synchronization | The automatic update of the password in the NNTP service when the password in the computer running a member of the Windows Server 2003 family is changed. |
| Basic authentication | Authentication that requires the user to provide a valid Windows user name and password that are authorized for access to that directory. The user name and password are sent as unencrypted clear text across the network. Permissions are set on the directory. |
| Integrated Windows Authentication | Authentication that encrypts the user name and password and then sends them to the NNTP service. The user name and password must be authorized for access to the requested newsgroup directory. Microsoft Outlook Express supports this protocol. Permissions are set on the directory. |
| Enable SSL client authentication | Enables secure client authentication using Secure Sockets Layer (SSL) encryption and certificates. A server SSL certificate is required. |
| Require SSL client authentication | Requires clients to provide an encrypted user name and password using SSL. |
| Enable client certificate mapping to Windows user accounts | Uses the Windows account mapped to the SSL client certificate to authenticate the user. Click **Client Mappings** to enter certificates and Windows account names. |
| Secure communication | Click **Certificate** to start the Web Server Certificate Wizard. |
| Connection control | The specific IP addresses that are allowed or not allowed access to this Network News Transfer Protocol (NNTP) virtual server. By default, all IP addresses are granted access. To view the options, click **Connection**. |

| | |
|---|---|
| All except the list below | Allows access to this NNTP virtual server to all IP addresses except for those listed. |
| Only the list below | Denies access to this NNTP virtual server to all IP addresses except for those listed. |
| Add | Click this button to open the **Computer** window where you can add computers or domains to the list. |
| IP address | The IP address for a single computer that is being added. Click **DNS Lookup** to find the IP address associated with the DNS name of the computer. |
| Subnet address | The base IP address for a group of computers that is being added. A subnet mask is required. |
| Subnet mask | The subnet mask that specifies the range of IP addresses for a group of computers that is being added. |
| Name | The domain name of a group of IP addresses that is being added. |

**Related Information**

- For information about NNTP security, see NNTP Security.
- For additional information about the **Access** tab, see Restricting Access to Newsgroups.

# NNTP General Tab

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the **General** tab to specify how clients connect to the NNTP service.

The **General** tab contains the following options:

| "Option" | "Description" |
|---|---|
| IP address | The IP address of the Network News Transfer Protocol (NNTP) virtual server. |
| Advanced | Opens the **Advanced** dialog box where you can enable multiple IP addresses for the NNTP virtual server and configure TCP and Secure Sockets Layer (SSL) port numbers. |
| TCP port | The TCP port used by clients to connect to the NNTP service. You can enter any unique TCP port number. The default is port119. |
| SSL port | The TCP port used by clients to connect to the NNTP service using SSL. You can enter any unique TCP port number. The default is port 563. |
| Limit number of connections to | When the check box is selected, this option limits the number of simultaneous client connections to the NNTP virtual server. The default is 5,000 connections. |
| Connection time out (minutes) | The number of minutes before an inactive client is disconnected. The default is 10 minutes. |
| Path header | The string used for the Path line in each news posting. The **Path header** displays the path a news message takes to reach its destination. When a message is forwarded from a news server, a descriptive string is added to the message in the Path line. Any punctuation character except the period (.) can separate the Path header. For information about how NNTP Path statements are generated, see Request for Comments (RFC) 1036.<br><br>The default value is the Domain Name System (DNS) name of the NNTP virtual server. You do not need to change the default value. |
| Enable logging | Enables logging to track users of the NNTP virtual server. The default log file directory is **systemroot**\System32\LogFiles\Nntpsvc1. To enable logging, you must also select the **Log access** check box in the **Virtual Directory** dialog box.<br><br>**Note**<br>Auditing server activity with log files is a good way to detect if unauthorized external users are attempting to access your virtual server or if internal users are trying to access resources that they do not have permissions to access. For more information about auditing and using Event Viewer, see Windows Server 2003 family Help. |
| Active log format | The transaction log format. To specify log properties, click **Properties**. |

# NNTP Security Tab

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the **Security** tab to specify the operators who are allowed access to the NNTP virtual server.

The **Security** tab contains the following options:

| Option | Description |
|---|---|
| Operators | The operators or groups of operators that are authorized to access this NNTP virtual server. The default value is the Administrators group for the computer. |
| Add | Opens the **Select Users or Groups** dialog box where you can select new operators. |
| Remove | Removes the selected operator. |
| Object Types (in the **Select Users or Groups** dialog box) | Objects from which to select operators. |
| Locations (in the **Select Users or Groups** dialog box) | The available locations in the selected domain. |
| Enter the object names to select (in the **Select Users or Groups** dialog box) | Provides a space to type the object names that you want to find. |

**Related Information**

- For information about NNTP security, see NNTP Security.
- For information about restricting operator access, see Restricting Operator Access.

# NNTP Settings Tab

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the **Settings** tab to set limits on posting articles, to enable control messages, and to enable moderated newsgroups.

The **Settings** tab contains the following options:

| Option | Description |
|---|---|
| Allow client posting | When selected, this option allows clients to post articles to newsgroups on this NNTP virtual server. |
| Limit post size (KB) | Specifies, in kilobytes, the largest single article that a user can post. To have no limit, clear the check box. |
| Limit connection size (MB) | Specifies, in megabytes, the maximum amount of data that a user can post to a newsgroup during a connection. To have no limit, clear the check box. |
| Allow feed posting | When selected, this option allows clients to post articles to newsfeeds on this NNTP virtual server. |
| Limit post size (KB) | Specifies, in kilobytes, the largest single article that a user can post to a newsfeed. To have no limit, clear the check box. |
| Limit connection size (MB) | Specifies, in megabytes, the maximum amount of data that a user can post to a newsfeed during a connection. To have no limit, clear the check box. |
| Allow servers to pull news articles from this server | Allows other NNTP servers to use pull feeds to get more than one article at a time from this server. |
| Allow control messages | Allows clients and servers to send control messages that create and remove newsgroups and delete individual articles. |
| SMTP server for moderated groups | The DNS name of the SMTP mail server that sends articles to the moderators of newsgroups or specifies the directory path where these articles are placed. The directory must be a local directory. |
| Default moderator domain | The default domain for moderated newsgroups. The domain must be a qualified domain name. Articles posted to moderated news groups that have no specified moderator are sent to *news_group_name@default_moderator_domain*, where *news_group_name* is the name of the newsgroup to which the article is sent and *default_moderator_domain* is the value that you specify for this option. |
| Administrator e-mail account | An e-mail address for receiving the non-delivery reports (NDRs) of moderated newsgroup articles that cannot be delivered to the specified moderator. To enable sending NDRs, create a new DWORDvalue named MailFromHeader with a value of 1in the following registry key:<br>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet \Services\NntpSvc\Parameters\ |

# Starting, Stopping, or Pausing NNTP Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use IIS Manager in the Microsoft Management Console (MMC) to start, stop, and pause the default NNTP virtual server and any other virtual servers that you have created.

You can also start, stop, and pause the NNTP service. However, stopping the service stops all of the NNTP virtual servers that are running on your computer. When the NNTP service is stopped, you cannot use MMC to perform administrative functions on any NNTP virtual server.

It is important to make sure that only trusted and qualified administrators in your organization have the necessary permissions to start or stop an NNTP virtual server. For more information, see Restricting Operator Access.

When an NNTP virtual server is started, it accepts new connections from users. When you stop an NNTP virtual server, it disconnects current users and does not accept new connections. When you pause an NNTP virtual server, it does not accept new connections but continues to service existing connections. If you want to stop an NNTP virtual server but not disconnect users currently connected to the virtual server, you can also choose to shut down gracefully.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To start, stop, or pause an NNTP virtual server**

1. In IIS Manager, expand the local computer, and then click the NNTP virtual server that you want.
2. On the **Action** menu, click **Start**, **Stop**, or **Pause**.

**To shut down an NNTP virtual server gracefully**

1. Pause the virtual server to prevent new connections.
2. To check for current users, use IIS Manager. Click **Current Sessions**. All current users are listed.
3. If you want to block current users from posting new articles, click the NNTP Settings Tab, and then clear the **Allow client posting** check box.
4. Stop the NNTP virtual server after all of the current users have disconnected.

**Related Information**

- For information about starting, stopping, or pausing the NNTP service, see Starting, Stopping, or Pausing the NNTP Service.

# Starting, Stopping, or Pausing the NNTP Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The NNTP service runs as a service on computers running a member of the Windows Server 2003 family. There can be multiple NNTP virtual servers on a computer, but only one instance of the NNTP service.

It is important to make sure that only trusted and qualified administrators in your organization have the necessary permissions to start or stop the NNTP service. For more information, see Restricting Operator Access.

You can start, stop, or pause each virtual server independently of one another while the NNTP service is running. It is important to remember, however, that pausing or stopping the entire NNTP service will affect all of the NNTP virtual servers. When the service is stopped, you cannot use Microsoft Management Console (MMC) to perform administrative functions on any NNTP virtual server.

The one exception is that, in MMC only, you can start a virtual server while the NNTP service is stopped. Starting an NNTP virtual server will also restart the entire the NNTP service. This, in turn, will start all other NNTP virtual servers that had been running when the NNTP service was originally stopped.

When you start the NNTP service, it accepts new connections from users. When you stop the NNTP service, it disconnects current users and does not accept new connections. When you pause the NNTP service, every running NNTP virtual server stops accepting new connections but continues to service existing connections. If you want to stop the NNTP service but not disconnect users currently connected to each of the virtual servers, you can also choose to shut down gracefully.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To start, stop, or pause the NNTP service**

1. From the **Start** menu, point to **Administrative Tools**, and then click **Component Services**.
2. In the **Component Services** box, click **Services (Local)**.
3. In the details pane, right-click **Network News Transfer Protocol (NNTP)**, and then click **Start**, **Stop**, or **Pause**.

**Note**

When you restart the NNTP service, all NNTP virtual servers that were running when you shut down the service will automatically restart. However, an NNTP virtual server that was already stopped when the service was shut down will remain stopped, even after the NNTP service is restarted.

**To shut down the NNTP service gracefully**

1. Follow the previous procedure to pause the NNTP service. This option prevents new connections.
2. Check for current users on each of your NNTP virtual servers. In IIS Manager, expand each NNTP virtual server and click **Current Sessions**. All current users are listed.
3. If you want to block current users from posting new articles, click the NNTP Settings Tab, and then clear the **Allow client posting** check box.
4. Stop each NNTP virtual server after all of the current users have disconnected.
5. After all of the current users have disconnected from each virtual server, stop the NNTP service.

**Related Information**

- For information about starting, stopping, or pausing NNTP virtual servers, see Starting, Stopping, or Pausing NNTP

Virtual Servers.

# NNTP Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, any user who can connect to the Network News Transfer Protocol (NNTP) service can view all newsgroups. These users connect anonymously, which means that they do not have to supply a user name or password. The Setup program creates a different Windows user account that is used by all anonymous users.

You can configure security to control who has access to specific newsgroups and who can operate the NNTP service. You can also restrict access on the basis of the client computer IP address. If users connect to the NNTP service over a public network, you can encrypt that connection by using the Secure Sockets Layer (SSL) protocol so that others cannot intercept the articles being sent and received.

Security is based on Windows user accounts and directory permissions. User access to newsgroups is restricted on the basis of user name and password.

This section includes the following information:

- Restricting Access to Newsgroups: Describes how to restrict newsgroup access and set authentication methods.
- Limiting Access by IP Address: Describes how to restrict access to the NNTP service by the IP address of the client computer.
- Restricting Operator Access: Describes how to add and remove NNTP operator privileges.
- Securing Connections with SSL: Describes Secure Sockets Layer (SSL) concepts and how to implement them.

# Restricting Access to Newsgroups

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can control access to individual newsgroups or sets of newsgroups by setting permissions for the directories that contain those newsgroups. You can set permissions for an individual directory or for a set of directories. You can also limit access to all newsgroups according to the IP address of the client computer.

If you do not want to restrict access to a newsgroup, enable anonymous access and do not set any permissions for the directory that contains the newsgroup. To enable anonymous access, on the NNTP Access Tab, click the **Authentication** button under **Access control**, and then select the **Allow anonymous** check box.

**Procedures**

**To restrict access to a newsgroup**

1. Create Windows user accounts for users. It is recommended that you organize users into groups and then grant permissions by group to simplify administration. If possible, use accounts and groups that are already established in your organization.
2. Define permissions for the directory that contains the newsgroup. Be sure to give the LocalSystem account full access to all newsgroup directories so that the Network News Transfer Protocol (NNTP) service has access to newsgroup files.
3. Set the authentication method used by the NNTP service.

The NNTP service supports two methods for authentication of users:

- Basic authentication
- Integrated Windows authentication

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To set the authentication method**

1. In IIS Manager, expand the local computer, right-click the NNTP virtual server, and then click **Properties**.
2. On the **Access** tab, under **Access control**, click **Authentication**.
3. Select one or more of the following:
   - The **Allow anonymous** check box.
   - The **Basic authentication** check box.
   - The **Integrated Windows Authentication** check box.
   - The **Enable SSL client authentication** check box.

| **Important** |
| --- |
| If you are going to require only Basic authentication, it is highly recommended that you also enable SSL to prevent user names and passwords from being compromised during transmission. |

4. To require Secure Sockets Layer (SSL) authentication, select the **Require SSL client authentication** check box.
5. To associate client certificates with Windows user accounts, select the **Enable client certificate mapping to Windows user accounts** check box, and then click **Client Mappings**.

**Related Information**

- For information about limiting access by IP address, see Limiting Access by IP Address.
- For information about restricting operator access, see Restricting Operator Access.
- For information about securing connections with SSL, see Securing Connections with SSL.

# Limiting Access by IP Address

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can limit access to the NNTP service by the IP address of the client computer. By default, all IP addresses have access to the NNTP service.

You can either allow or deny access to a specific list of IP addresses. IP addresses can be specified individually or as a group by using a subnet mask. You can also specify IP addresses by using a domain name, but doing so adds the overhead of a DNS lookup for each connection.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To limit access by IP address**

1. In IIS Manager, expand the local computer, right-click the NNTP virtual server, and then click **Properties**.
2. On the **Access** tab, under **Connection control**, click **Connection**.
3. In the **Connection** box, click **All except the list below** to deny access to a list of IP addresses, or click **Only the list below** to allow access to a list of IP addresses.
4. In the **Connection** box, click **Add** for each IP address you want to grant or deny access to the NNTP virtual server.
5. In the **Computer** box, click one of the following options: **Single Computer**, **Group of computers**, or **Domain**. Then proceed as follows:
   - For the **Single Computer** option: In **IP address**, type the IP address of the computer. (If you don't know the IP address, click **DNS Lookup**.)
   - For the **Group of computers** option: In **Subnet address**, type the IP address. Then, in **Subnet mask**, type the subnet mask for the group of computers.
   - For the **Domain** option: In **Name**, type the domain name of the computer.
6. Click **OK**.

**Related Information**

- For information about restricting operator access, see Restricting Operator Access.
- For information about securing connections with SSL, see Securing Connections with SSL.

# Restricting Operator Access

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Operators of the NNTP service must be assigned operator user rights. By default, everyone in the Administrators group on computers running a member of the Windows Server 2003 family is assigned operator user rights.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

### To add an operator

1. In IIS Manager, expand the local computer, right-click the NNTP virtual server, and then click **Properties**.
2. On the **Security** tab, click **Add**.
3. In the **Select Users or Groups** box, click **Locations**, and then click the domain of the operator you want to add.
4. In the **Select Users or Groups** box, click **Object Types**, click the operator or group, and then click **OK**.
5. In the **Select Users or Groups** box, type the object name to select, or click **Advanced** to search for a particular object.
6. After you make all of your selections, click **OK**.

### To remove an operator

1. In IIS Manager, expand the local computer, right-click the NNTP virtual server, and then click **Properties**.
2. On the **Security** tab, in the **Operators** list, click the operator that you want to remove.
3. On the **Security** tab, click **Remove**, and then click **OK**.

**Related Information**

- For information about limiting access by IP address, see Limiting Access by IP Address.
- For information about securing connections with SSL, see Securing Connections with SSL.

# Securing Connections with SSL

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Secure Sockets Layer (SSL) provides a secure, encrypted connection between the NNTP service and the client computer. SSL can protect your private information when users connect across a public network such as the Internet.

SSL support requires an SSL certificate, and this certificate has to be installed on a computer running a member of the Windows Server 2003 family. SSL must also be supported by the client software. Microsoft Outlook Express supports SSL.

If a server certificate is installed, the NNTP service uses SSL whenever a client requests it. You can optionally require SSL for all newsgroups or for newsgroups located in a virtual directory. You can also use SSL authentication of clients.

**Procedures**

**To enable SSL**

- Obtain and install an SSL server certificate.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To require an SSL connection for access to newsgroups**

1. In IIS Manager, expand the local computer, expand the NNTP virtual server, and then click **Virtual Directories**.
2. In the details pane, right-click the virtual directory that you want to modify, and then click **Properties**.
3. On the **General** tab, click **Secure**.
4. In the **Security** box, select the **Require secure channel** check box.

**To enable SSL authentication for clients**

1. In IIS Manager, expand the local computer, right-click the NNTP virtual server, and then click **Properties**.
2. On the **Access** tab, click **Authentication**.
3. In the **Authentication Methods** box, select the **Enable SSL client authentication** check box.

**Related Information**

- For information about limiting access by IP address, see Limiting Access by IP Address.
- For information about restricting operator access, see Restricting Operator Access.

# Managing Newsgroups

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

As soon as you create your newsgroups, users can begin posting articles to them. If you want a moderator to review new articles before they are posted, define a moderated newsgroup.

You can change the properties of an existing newsgroup or delete the newsgroup if you no longer need it. You can also select options that enable the NNTP service to delete articles after a specified number of hours or when the amount of disk space used by newsgroups reaches specified limits. Finally, you can filter which newsgroups are visible to users and the number that will display whenever the **Newsgroups** node is selected under the Network News Transfer Protocol (NNTP) virtual server.

This section includes the following information:

- Creating Newsgroups in IIS 6.0: Describes how to use the New Newsgroup Wizard to create a newsgroup.
- Moderating Newsgroups in IIS 6.0: Introduces the concept of moderating newsgroups, how to enable moderating, and how to create a moderated newsgroup.
- Editing and Deleting Newsgroups: Describes how to edit the properties of an existing NNTP service newsgroup or delete a newsgroup that you no longer need.
- Defining Newsgroup Limits and Expirations: Describes how to limit the length of time articles are kept in an NNTP service newsgroup by defining expiration policies for newsgroups.
- Limiting Newsgroup Enumeration in IIS 6.0: Describes how to limit the number of newsgroups that will be displayed in IIS Manager.

# Creating Newsgroups in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the New Newsgroup Wizard to create a newsgroup. When you create a new newsgroup, the NNTP service creates the directory for the newsgroup automatically. If you want a moderator to review new articles before they are posted, you can create a moderated newsgroup.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a newsgroup**

1. In IIS Manager, expand the local computer, and then expand the NNTP virtual server that you want.
2. Right-click **Newsgroups**, point to **New**, and then click **Newsgroup**.
3. Follow the steps in the New Newsgroup Wizard:
   - In **Name**, type the name of the newsgroup you want to create.
   - In **Description**, type a brief description of the newsgroup function.
   - If you want, in **Pretty name** type another name for the newsgroup, which will be returned by the **List Prettynames NNTP** command. The pretty name is another alias for the newsgroup that can be displayed by clients.
   - Click **Finish**.

**Related Information**

- For information about moderated newsgroups, see Moderating Newsgroups in IIS 6.0.
- For information about editing and deleting newsgroups, see Editing and Deleting Newsgroups.
- For information about defining newsgroup limits and expirations, see Defining Newsgroup Limits and Expirations.
- For information about limiting newsgroup enumeration, see Limiting Newsgroup Enumeration in IIS 6.0.

# Moderating Newsgroups in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you want to control what articles are posted to a newsgroup, you should create a moderated newsgroup. Articles submitted to a moderated newsgroup are not posted until they are approved by the moderator you have chosen for that newsgroup.

| **Note** |
| --- |
| Microsoft recommends having a moderator for every newsgroup, especially if users outside your organization are allowed to make postings on your server. A moderator can filter inappropriate material and he or she can detect if someone is trying to compromise information on the server or compromise the server itself. |

When a user submits an article to a moderated newsgroup, the NNTP service sends the article to the newsgroup moderator. The moderator can either approve the article and send it back to the NNTP service to be posted or discard the article. For discarded articles, the moderator can send a message to the article's author to explain why the article was rejected.

Any NNTP client can be used to post messages to a moderated newsgroup. However, the moderator needs to use an NNTP client that supports moderating functions.

To use moderated newsgroups, you must specify either a Simple Mail Transfer Protocol (SMTP) mail server for sending articles directly to the moderators or a directory where the articles are stored for the moderators.

You can also specify a default moderator domain. If you choose to use the default moderator for a newsgroup, articles are sent to *newsgroup_name@default_domain*, where *newsgroup_name* is the name of the newsgroup (the periods have been replaced by underscores). For example, articles posted to the newsgroup rec.sports.tennis, with a default domain of contoso.com, are sent to rec_sports_tennis@contoso.com.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

### To enable moderated newsgroups

1. In IIS Manager, expand the local computer, right-click the NNTP virtual server, and then click **Properties**.
2. On the **Settings** tab, in **SMTP server for moderated groups**, type the DNS name or the IP address of the SMTP server, or type the path to the directory where articles are stored for the moderators. The directory path must be a local directory. You should normally set this value to the pickup directory of the SMTP server. By default, this directory is *localdrive*\Inetpub\Mailroot\Pickup.
3. On the **Settings** tab, in **Default moderator domain**, type the e-mail domain for default moderators.
4. Click **OK**.
5. If you specified an SMTP server, be sure that the user ID under which the NNTP service is running (by default, the System Account) is defined as an account on the SMTP server.

### To create a moderated newsgroup

1. If necessary, create the newsgroup. After you finish the New Newsgroup Wizard, the procedure is the same for both new and existing newsgroups.
2. In IIS Manager, expand the local computer, expand the NNTP virtual server, and then click **Newsgroups**.
3. In the details pane, right-click the newsgroup that you want, and then click **Properties**.
4. On the **General** tab, select the **Moderated** check box.

5. In **Moderator**, type the e-mail address of the moderator, or click **Set Default** to use the default moderator.
6. Click **OK**.

**Related Information**

- For information about creating newsgroups, see Creating Newsgroups in IIS 6.0.
- For information about editing and deleting newsgroups, see Editing and Deleting Newsgroups.
- For information about limiting newsgroup enumeration, see Limiting Newsgroup Enumeration in IIS 6.0.

# Editing and Deleting Newsgroups

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can edit the properties of an existing NNTP service newsgroup or delete a newsgroup that you no longer need.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

### To edit or delete a newsgroup

1. In IIS Manager, expand the local computer, expand the NNTP virtual server, and then click **Newsgroups**.
2. In the details pane, right-click the newsgroup that you want, and then click one of the following:
   - **Properties**, to edit the newsgroup properties.
   - **Rename**, to rename the newsgroup.
   - **Delete**, to remove the newsgroup.

**Related Information**

- For information about creating newsgroups, see Creating Newsgroups in IIS 6.0.

# Defining Newsgroup Limits and Expirations

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can limit how long articles are kept in an NNTP service newsgroup by defining expiration policies for newsgroups. An expiration policy can apply to a single newsgroup or to any number of newsgroups. You can define as many policies as you need. In each case, the oldest articles are deleted first.

If you do not specify an expiration policy for a newsgroup, you should delete articles manually when they are no longer needed.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

#### Procedures

### To create an expiration policy

1. In IIS Manager, expand the local computer, and then expand the NNTP virtual server that you want.
2. Right-click **Expiration Policies**, point to **New**, and then click **Expiration Policy**.
3. Follow the steps in the New NNTP Expiration Policy Wizard.

You can modify an existing expiration policy, including changing the newsgroups to which the policy applies, with the **Expiration Policy** dialog box.

### To modify an expiration policy

1. In IIS Manager, expand the local computer, expand the NNTP virtual server that you want, and then click **Expiration Policies**.
2. In the details pane, right-click the expiration policy that you want to modify, and then click **Properties**.
3. Change the options in the **Expiration Policy** box as needed.

**Note**

Use the **Move Up** and **Move Down** buttons to specify which newsgroups the expiration policy should apply to and in what order. For example, if you want to include all of the alt.tv newsgroups (alt.tv.*) in the policy, but exclude the alt.tv.cable newsgroup from the policy, meaning that you want to apply an expiration policy to all of the alt.tv newsgroups except the alt.tv.cable newsgroup, use the **Move Up** and **Move Down** buttons to include alt.tv.* first and exclude alt.tv.cable second.

#### Related Information

- For information about editing and deleting newsgroups, see Editing and Deleting Newsgroups.

# Limiting Newsgroup Enumeration in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the **Find Newsgroups** dialog box to limit the number of newsgroups that will be displayed when you select **Newsgroups**, which is beneath the NNTP virtual server node.

Potentially, a server can store thousands of newsgroups; in some cases, it might not be possible to enumerate all of them. Administrators may want to filter the newsgroups that will be enumerated. You can filter newsgroups by type, by number, or both.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To limit the number of enumerated newsgroups**

1. In IIS Manager, expand the local computer, and then expand the NNTP virtual server that you want.
2. Right-click **Newsgroups**, and then click **Limit groups enumeration**.
3. In the **Newsgroups** field, type the kind of the newsgroup you want to display when **Newsgroup** is selected. All other types will be filtered out. Use wildcards, such as **alt.***. In this example, only newsgroups under the **alt** hierarchy would be visible to users. By default, every kind of newsgroup is enumerated.

   | **Note** |
   | --- |
   | You can also type the name of a specific newsgroup in this space, but then only that newsgroup will be visible. Only one wildcard or one newsgroup can be entered at a time. |

4. If you do not want to have a limit on the number of enumerated newsgroups, clear the **Limit number of items on the result pane** check box. This check box is selected by default.
5. In **Maximum number of items**, type the maximum number of newsgroups that you want to enumerate on this virtual server. This limit will only apply when the **Limit number of items on the result pane** check box is selected. The default is 100 newsgroups.

**Related Information**

- For information about creating newsgroups, see Creating Newsgroups in IIS 6.0.
- For information about editing and deleting newsgroups, see Editing and Deleting Newsgroups.

# Configuring Expiration Policies

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Expiration policies can be set separately for each virtual server. Expiration policies specify the article age limit for newsgroups; they also specify the newsgroups to which the policy applies. If a newsgroup is assigned to more than one expiration policy, the more restrictive policy applies.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a new expiration policy**

1. In IIS Manager, expand the local computer, and then expand the NNTP virtual server that you want.
2. Right-click **Expiration Policies**, point to **New**, and then click **Expiration Policy**.
3. Follow the steps in the New NNTP Expiration Policy Wizard. The wizard will prompt you for the name of an expiration policy, the newsgroups it will affect, and the expiration period. These properties are described in the table below.

**To edit an existing expiration policy**

1. In IIS Manager, expand the local computer, expand the NNTP virtual server that you want, and then click **Expiration Policies**.
2. In the details pane, right-click the expiration policy that you want to edit, and then click **Properties**.
3. Change the options in the **Expiration Policy** box as needed.

| **Note** |
| --- |
| Use the **Move Up** and **Move Down** buttons to specify which newsgroups the expiration policy should apply to and in what order. For example, if you want to include all of the alt.tv newsgroups (alt.tv.*) in the policy, but exclude the alt.tv.cable newsgroup from the policy, meaning that you want to apply an expiration policy to all of the alt.tv newsgroups except the alt.tv.cable newsgroup, use the **Move Up** and **Move Down** buttons to include alt.tv.* first and exclude alt.tv.cable second. |

The following properties can be set for expiration policies:

| Option | Description |
| --- | --- |
| Articles older than (Hours) | Removes articles that are older than a specified number of hours. The default is 168 hours (7 days). To have no age limit, clear the check box. |
| Affected newsgroups | A list of the newsgroups to which the policy applies. The asterisk (*) wildcard character specifies multiple newsgroups. For example, rec.* denotes and includes all newsgroups starting with "rec.". |
| Add | Click this button to open the **Add Newsgroup** dialog box where you can type in the names of newsgroups to add to this expiration policy. Use the asterisk (*) wildcard character to specify multiple newsgroups. |
| Remove | To remove a newsgroup from an expiration policy, select the newsgroup in the list of **Affected newsgroups** and click **Remove**. |

| | |
|---|---|
| Move Up | To move a newsgroup up in an expiration policy, select the newsgroup in the list of **Affected newsgroups** and click **Move Up**. |
| Move Down | To move a newsgroup down in an expiration policy, select the newsgroup in the list of **Affected newsgroups** and click **Move Down**. |

**Related Information**

- For information about creating expiration policies, see Defining Newsgroup Limits and Expirations.

# Rebuilding the IIS NNTP Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can rebuild the IIS NNTP service by using IIS Manager.

You can choose one of two rebuild types:

- **Standard**. Rebuilds only group.lst files, based on the contents of the virtual directories. This is the faster of the two options.
- **Thorough**. Rebuilds all internal files, such as .hsh and .lst files.

First run the standard rebuild. If the problem is not fixed when you run the standard rebuild, run a thorough rebuild.

| **Note** |
| --- |
| You must stop the NNTP service news server before you can rebuild it. |

**To rebuild the NNTP service by using IIS Manager**

1. In IIS Manager, click the local computer, right-click the NNTP virtual server, and then click **Stop**.
2. Right-click the NNTP virtual server again, point to **All Tasks**, and then click **Rebuild Server**. If **Rebuild Server** is not available, then the server might still be running. Check to make sure that it is stopped.
3. In the **Rebuild NNTP Virtual Server** dialog box, do one of the following:
   - To run a standard rebuild, in **Rebuild type**, move the slider to **Standard**, and then click **Start**.
   - To run a thorough rebuild, in **Rebuild type**, move the slider to **Thorough**, and then click **Start**.

For more information about common problems that you can fix by rebuilding the NNTP service, see Rebuilding the NNTP Service.

# Enabling Newsfeeds

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Newsfeeds send articles between the IIS NNTP service and other NNTP servers. With newsfeeds, the articles that are posted to one server can be read by users of the other servers. Newsfeeds are not enabled by default.

**To enable newsfeeds**

1. In IIS Manager, click the local computer, right-click the NNTP virtual server, and then click **Properties**.
2. Click the **Settings** tab, select the **Allow feed posting** check box, and then do any of the following:
   - To limit the size of articles posted to this news site, select the **Limit post size (KB)**check box, and type or select the size (1 to 32,000 kilobytes) of the largest article that is accepted.
   - To limit the size of a newsfeed, select the **Limit connection size (MB)**check box, and type or select the size (1 to 4000 megabytes) of the largest newsfeed to accept.
   - To allow servers to pull articles from this news site, select the **Allow servers to pull news articles from this server**check box.

For information about the different kinds of newsfeeds, and adding, deleting, and modifying newsfeeds with Rfeed.vbs, see Managing Newsfeeds.

# SMTP Administration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The Internet-standard Simple Mail Transfer Protocol (SMTP) can be used to transport and deliver messages based on specifications in Request for Comments (RFC) 821 and RFC 822. The SMTP server also includes enhancements that build upon the basic delivery functions of the protocol. There are options available that give you control over the routing and delivery of messages, and that provide secure communications.

This section includes the following information:

- Configuring SMTP Virtual Servers: Describes how to set an SMTP virtual server display name and the IP address/TCP port combination.
- Starting, Stopping, or Pausing SMTP Virtual Servers: Introduces the concepts of starting, stopping, and pausing an SMTP virtual server and how it differs from the same actions on a service.
- Starting, Stopping, or Pausing the SMTP Service: Describes the concept of multiple virtual servers and how to start, stop, or pause them individually.
- Configuring Startup Settings: Describes how to configure the default state of the SMTP service at startup.
- Setting Connections: Describes how to set incoming and outgoing connection limits.
- Creating Additional SMTP Virtual Servers: Describes how to create new SMTP virtual servers.
- Setting Up Virtual Servers for Clustering: Introduces the concept of clustered servers and how to configure virtual servers for that environment.
- Enabling Protocol Logging: Describes how to enable logging in ASCII or ODBC formats.
- SMTP Security: Introduces security and access concepts.
- Delivering Messages: Describes the settings that determine how delivery occurs, including routing, transmission, and security options.
- Managing Messages_IIS_SP1_Ops: Describes how to set connection limits to maximize performance, how to set limits for message size and number of recipients, and how to manage undeliverable messages.
- LDAP Routing: Describes how to configure the SMTP service so that it will consult an LDAP server to resolve senders and recipients.
- Monitoring Message Transactions: Describes four protocol log formats and how to implement them.
- SMTP Domains: Introduces local and remote domain concepts.

# Configuring SMTP Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are two identifying features for a Simple Mail Transfer Protocol (SMTP) virtual server: the display name and the IP address/TCP port combination. Setup assigns a Microsoft Management Console display name for the default SMTP virtual server. You can keep the name (Default SMTP Virtual Server) or change it using the following steps. You can also select the IP address that will be associated with this SMTP virtual server.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**.

**Procedures**

**To identify an SMTP virtual server**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Rename**.
2. Type a new name for the virtual server name if you want to change it from the default. Do not use extended characters when naming your SMTP virtual server.
3. Right-click the SMTP virtual server, and click **Properties**.
4. On the **General** tab, in the **IP address** list, click the IP address for this virtual server. The SMTP virtual server can respond to connection requests for all IP addresses configured on the computer. To identify the TCP port for each IP address configured for the virtual server, click **Advanced**. Port 25 is the SMTP standard TCP port and is recommended. More than one virtual server can use Port 25, provided they are associated with different IP addresses.

**Note**

When you change an SMTP virtual server's configuration, you may want to use the **Pause** command. Pausing prevents new client connections but enables Microsoft SMTP Service to continue to process existing client connections and deliver queued messages.

# Starting, Stopping, or Pausing SMTP Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The Default SMTP Virtual Server starts upon installation of the SMTP service. You can pause, stop, and start it in IIS Manager.

You can also start, stop, and pause the entire SMTP service. However, if you have more than one virtual server, stopping the service affects all of the Simple Mail Transfer Protocol (SMTP) virtual servers running on your computer. When the SMTP service is stopped, you cannot use IIS in Microsoft Management Console (MMC) to perform administrative functions on any SMTP virtual server.

| **Important** |
| --- |
| Make sure only trusted administrators in your organization have the necessary permissions to start or stop an SMTP virtual server. For more information, see Setting Operator Permissions. |

You can stop and start the SMTP service manually. However, while it is operating, you must be careful when stopping, pausing, or restarting the service to minimize the impact on users.

If the default startup setting is **Manual**, you can use IIS in Microsoft Management Console (MMC) to start a Simple Mail Transfer Protocol (SMTP) virtual server.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To start an SMTP virtual server**

- In IIS Manager, expand the local computer, right-click the SMTP virtual server, and click **Start**.

You can stop a Simple Mail Transfer Protocol (SMTP) virtual server for configuration changes and maintenance.

**To stop an SMTP virtual server**

- In IIS Manager, expand the local computer, right-click the SMTP virtual server, and click **Stop**.

You can pause a Simple Mail Transfer Protocol (SMTP) virtual server for configuration changes and maintenance. Pausing prevents new client connections, but it enables the virtual server to continue processing existing client connections and delivering queued messages.

**To pause an SMTP virtual server**

- In IIS Manager, expand the local computer, right-click the SMTP virtual server, and click **Pause**.

**Related Information**

- For information about managing the SMTP service, see Starting, Stopping, or Pausing the SMTP Service.

# Starting, Stopping, or Pausing the SMTP Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The SMTP service runs as a service on Windows Server 2003, Standard Edition and Windows Server 2003, Enterprise Edition, and it starts upon installation. Although there is only one SMTP service on a computer, it is possible to have more than one Simple Mail Transfer Protocol (SMTP) virtual server. You can start, stop, or pause each virtual server independently of one another while the SMTP service is running.

If you have more than one SMTP virtual server, it is important to remember that pausing or stopping the entire SMTP service will affect all of the SMTP virtual servers. When the service is stopped, you cannot use IIS in Microsoft Management Console (MMC) to perform administrative functions on any SMTP virtual server.

| **Note** |
| --- |
| The one exception is that you can start a virtual server while the SMTP service is stopped. Starting an SMTP virtual server will also restart the entire SMTP service. This, in turn, will start all other SMTP virtual servers that had been running when the SMTP service was originally stopped. |

When you start the SMTP service, it accepts new connections from users. When you stop the SMTP service, it does not accept new connections. When you pause the SMTP service, every running SMTP virtual server will cease accepting new connections but will continue to service existing connections.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To start, stop, or pause the SMTP service**

1. From the **Start** menu, point to **Administrative Tools**, and then click **Component Services**.
2. In the console tree, click **Services (Local)**.
3. In the details pane, right-click **Simple Mail Transfer Protocol (SMTP)**, and then click **Start**, **Stop**, or **Pause**.

| **Note** |
| --- |
| When you restart the SMTP service, all SMTP virtual servers that had been running when you shut down the service will automatically restart. However, an SMTP virtual server that was already stopped when the service was shut down will remain stopped, even after the SMTP service restarts. |

**Related Information**

- For information about starting, stopping, or pausing an SMTP virtual server, see Starting, Stopping, or Pausing SMTP Virtual Servers.

# Configuring Startup Settings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use Administrative Tools to configure the default state of the SMTP service at startup.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in theAdministrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**To set SMTP Startup Type**

1. From the Start menu, point to Administrative Tools, and then click Component Services.
2. In the console tree, click Services (Local).
3. In the details pane, right-click Simple Mail Transfer Protocol (SMTP), and then click Properties.
4. On the General tab, in the Startup type list, click Automatic, Manual, or Disabled.

# Setting Connections

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A connection is initiated whenever a message is sent to or received from a remote server.

| **Note** |
| --- |
| Designating the TCP port that the SMTP service uses to receive incoming messages is done when you configure the Simple Mail Transfer Protocol (SMTP) virtual server. |

Setting connection limits and imposing connection time-outs can make it more difficult for someone to initiate a malicious attack (such as denial of service) against your virtual server.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure incoming connections**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. On the **General** tab, select the **Limit number of connections to** check box (the default is no limit), and set the following options.

| Option | Description |
| --- | --- |
| **Limit number of connections to** | When the check box is selected, this option specifies the number of concurrent connections for incoming message delivery. The minimum is 1 connection. When the check box is not selected, no limit is imposed. |
| **Connection time-out (minutes)** | Specifies the time allowed before an inactive connection is closed. The default is 10 minutes. |

**To configure outbound connections**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Delivery** tab, and click **Outbound connections**.
3. In the **Outbound Connections** dialog box, select a check box and set the following options.

| Option | Description |
| --- | --- |
| **Limit number of connections to** | When the check box is selected, this option specifies the total number of simultaneous outbound connections to all remote domains that can exist at one time. The default is 1,000 connections. The minimum is 1.<br><br>This setting can be used to improve system performance in conjunction with the **Limit number of messages per connection to** check box on the **Messages** tab. |
| **Time-out (minutes)** | Specifies the time allowed before an inactive connection is closed. The default is 10 minutes. |

| | |
|---|---|
| **Limit number of connections per domain to** | When the check box is selected, this option limits connections to any single remote domain. The default is 100 connections. This number should be less than or equal to the value for the **Limit number of connections to** option. |
| **TCP port** | Designates the TCP port used for outgoing transmissions. The default is port 25. The outgoing port setting can be the same as that for incoming transmissions or it can be different. |

**Related Information**

- For information about configuring an SMTP virtual server, see Configuring SMTP Virtual Servers.
- For information about limiting the number of messages per connection, see Limiting the Number of Messages per Connection.

# Creating Additional SMTP Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In most cases, you should need only one Simple Mail Transfer Protocol (SMTP) virtual server. However, if you are hosting multiple domains and want to have more than one default domain, for example, you can create multiple SMTP virtual servers. To an end user, each SMTP virtual server appears as a separate server with a unique IP address/TCP port combination.

When you create an SMTP virtual server, you are prompted to enter a path to your home directory. This directory must be local to the computer on which the SMTP service runs.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a new SMTP virtual server**

1. If the SMTP virtual server will be using a new IP address, configure an IP address. From the **Start** menu, click **Control Panel**, and then double-click **Network Connections**. See Windows Server 2003 family Help for more information about Network Connections.
2. After you configure the IP address, in IIS Manager, right-click an existing SMTP virtual server, point to **New**, and then click **Virtual Server**.
3. Follow the steps in the **New SMTP Virtual Server Wizard**. Be sure to select an IP address/TCP port combination that is not being used by another SMTP virtual server. The recommended TCP port is 25, which is the SMTP standard TCP port. More than one virtual server can use the same TCP port provided they are configured with different IP addresses.
4. If the default startup setting for SMTP Service is set to **Automatic**, the new SMTP virtual server will start automatically. If it doesn't start, it is because you selected an IP address/TCP port combination that is already in use.
5. Configure the new SMTP virtual server.

**Related Information**

- For information about configuring an SMTP virtual server, see Configuring SMTP Virtual Servers.

# Setting Up Virtual Servers for Clustering

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A server cluster is a group of independent computer systems, known as nodes, working together as a single system to ensure that mission-critical applications and resources remain available to clients. Server clusters provide high availability. High availability in a Simple Mail Transfer Protocol (SMTP) server cluster means that, if one virtual server fails, its work is dispersed to the remaining virtual servers in the cluster, ensuring mail service is not interrupted. The other benefit of server clusters is the failover of server resources, which is when a virtual server goes down and another one comes online to assumes its role.

In clustering, a node is a system that has a working installation of Windows Server 2003, Enterprise Edition and the Cluster service. Microsoft recommends creating a new virtual server on your node for clustering, rather than reconfiguring your default SMTP virtual server.

For more information about Windows Clustering and server clusters, see Windows Server 2003, Enterprise Edition Help. That documentation contains detailed information about setting up and administrating server clusters.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set up an SMTP virtual server for clustering**

1. If necessary, change the default startup setting of the SMTP Service to **Manual**.
2. Before creating a new virtual server, in IIS Manager, right-click **Default SMTP Virtual Server**, and then click **Properties**.
3. On the **General** tab, click **Advanced**.
4. Under **Address**, double-click **All Unassigned**.
5. In the **Identification** box, change the TCP port number from 25, which is the default setting, to another port number -- any unassigned port will work. Type in the new port number, making sure it is unique, and then click **OK**.
6. Create a new SMTP virtual server. After it is created, you will have to manually start the new virtual server. Use the default TCP port on the new virtual server, which is port 25.

   | **Important** |
   | --- |
   | When you get to the **Select Home Directory and Default Domain** screen of the New SMTP Virtual Server Wizard, be sure to type paths to the shared resource that will be used in clustering. Do not type paths to a local hard drive. |

7. Run the **iiscnfg /copy** command to copy the IIS configuration on your node to the other nodes in the server cluster. At a command prompt, navigate to the systemroot\System32\Inetsrv directory on your node, and then type **iiscnfg /copy /ts target server /tu userid /tp password** where the following is true:
   - **/ts** is another node in the same cluster.
   - **/tu** is the user ID to use when connecting to the target server.
   - **/tp** is the password associated with the specified user ID.

| **Important** |
| --- |

The **/copy** operation does not copy the server content, such as Web pages and FTP files, that is associated with the IIS configuration. This command changes the computer-specific and system-specific properties in the metabase so that they are valid on the target computer. However, it does not adjust the directory or file paths. As a result, you might need to configure valid paths on the target computer.

Repeat the **iiscnfg /copy** command for each node in the server cluster. For more information about the **iiscnfg /copy** command, see Copying IIS Configurations Using Iiscnfg.vbs.

1. Check the TCP port numbers on your newly created SMTP virtual server and on the default virtual server. Specifically, view your local node from other nodes in the server cluster, and make sure that the TCP port settings are identical to the settings on the local node. On all nodes, the default SMTP virtual server should no longer point to the default TCP port (25), but the new SMTP virtual server should be assigned port 25. Settings must be identical on every node for clustering to work.

2. Start **Cluster Administrator**. Follow the wizard and associate a resource type to the new SMTP virtual server.

**Note**

In the **Parameters** box, be sure to select your new SMTP virtual server from the list, and not the default SMTP virtual server. See Windows Server 2003 Enterprise Edition Help for more information about Cluster Administrator.

**Related Information**

- For information about configuring the default startup setting, see Configuring Startup Settings
- For information about creating additional SMTP virtual servers, see Creating Additional SMTP Virtual Servers.
- For information about starting, stopping, or pausing SMTP virtual servers, see Starting, Stopping, or Pausing SMTP Virtual Servers.

# Enabling Protocol Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the **General** tab to enable transaction logging and to select the format. When setting up the log file, keep in mind that the log formats and default file names are the same as those used by other IIS services. If you choose the default names, all transactions for all services are recorded in the same file. You can set up a separate file for each service if you want to maintain separate records.

If you set up the file in a location other than the default, make sure it is stored on a local drive and not on a network.

Auditing server activity with log files is a good way to detect if unauthorized external users are attempting to access your virtual server, or if internal users are trying to access resources they do not have permission to access. For more information about auditing and using Event Viewer, see Windows Server 2003 family Help.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable logging that uses one of the ASCII text formats**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. On the **General** tab, select the **Enable logging** check box.
3. In the **Active log format** list, click a transaction log format. If you chose the **W3C Extended Log File Format**, click the **Advanced** tab, and then click the items you want to track.
4. Click **Properties**, and then set the log file size and location.

**To enable logging that uses the ODBC format**

1. Set up an ODBC-compliant database.
2. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
3. On the **General** tab, select the **Enable logging** check box.
4. In the **Active log** format list, click **ODBC Logging**.
5. Click **Properties**, and then set the log file size and location.

**Related Information**

- For information about monitoring message transactions, see Monitoring Message Transactions.
- For information about setting up an ODBC-compliant database, see Setting Up an ODBC-Compliant Database.

# SMTP Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can specify security settings for a Simple Mail Transfer Protocol (SMTP) virtual server on the two security tabs. These settings apply to all domains on the virtual server.

| Tab | Settings |
|---|---|
| **Security** | Designates the user accounts that have operator permissions to run the SMTP service. Operators have the ability to configure SMTP virtual server properties. |
| **Access** | Contains settings to require authentication and the use of Transport Layer Security (TLS) for all incoming connections. You can also grant or deny access to specific computers or networks. |

This section includes the following information:

- Setting Operator Permissions: Describes how to assign and remove operator permission on SMTP virtual servers.
- Requiring Authentication for Incoming Connections: Describes how to disable authentication for incoming messages, set clear text authentication for incoming messages, and use integrated Windows authentication to authenticate incoming messages.
- Requiring Authentication for Outbound Messages: Describes how to disable authentication for outgoing messages, set Basic authentication for outgoing messages, and set Integrated Windows authentication for outgoing messages.
- Setting IP Access Restrictions to Servers: Describes how to set IP address access restrictions.
- Configuring SMTP Virtual Server Relay Restrictions: Describes how to add and remove relay restrictions from an SMTP virtual server.
- Requiring TLS Encryption: Describes how to create and manage key certificates and how to set TLS encryption levels for the server.

# Setting Operator Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can designate which user accounts can have operator permissions for the Simple Mail Transfer Protocol (SMTP) virtual server. After Windows user accounts are set up, you can grant permissions by selecting the account from a list. These permissions can be rescinded by removing the account from the list of virtual server operators.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To assign operator permissions**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Security** tab, and click **Add**.
3. Select a Windows user account, and then click **OK**. The selected account will now appear in Operators.

**To remove operator permissions**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Security** tab, and in the Operators list, click the operator you want to remove.
3. Click **Remove**.

| **Note** |
| --- |
| You can also remove operator permissions from the Windows user accounts that are listed in Operators. |

**Related Information**

- For information about requiring authentication for incoming connections, see Requiring Authentication for Incoming Connections.
- For information about requiring authentication for outbound messages, see Requiring Authentication for Outbound Messages.

# Requiring Authentication for Incoming Connections

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are three authentication methods available. You can choose one, two, or all three methods. All three are set by default.

| **Important** |
| --- |
| The default mail relay settings on the SMTP virtual server allow only mail from authenticated users to be relayed. Therefore, to allow users in your organization to send mail through the virtual server, users must first be authenticated. By choosing one of the methods in the Authentication dialog box, your users will be able to send mail and, at the same time, unauthorized users will not be able to use your virtual server to relay mail. |

| Authentication option | Description |
| --- | --- |
| **Anonymous access** | An account name or password is not required. You can use this option to disable authentication for the Simple Mail Transfer Protocol (SMTP) virtual server. |
| **Basic authentication** | An account name and a password are sent using plaintext. Specify a Windows domain that is appended to the account name for authentication. |
| **Integrated Windows authentication** | A Windows account name and password are authenticated using this option. |

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To disable authentication for incoming messages**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Access** tab, and under **Access control**, click **Authentication**.
3. Select the **Anonymous access** check box, and then clear the remaining check boxes for the other options.

**To set clear text authentication for incoming messages**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Access** tab, and under **Access control**, click **Authentication**.
3. Select the **Basic authentication** check box.
4. In the **Default domain** box, type a Windows domain name. This default domain differs from the SMTP virtual server default domain.

**To use Integrated Windows authentication to authenticate incoming messages**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Access** tab, and under **Access control**, click **Authentication**.
3. Select the **Integrated Windows Authentication** check box.

**Related Information**

- For information about requiring authentication for outbound messages, see Requiring Authentication for Outbound Messages.

# Requiring Authentication for Outbound Messages

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure the Simple Mail Transfer Protocol (SMTP) virtual server to provide the authentication credentials required by a receiving server. There are three types of authentication available: anonymous, Basic (plaintext), and Integrated Windows authentication. Anonymous requires no authentication. With the plaintext option, the account name and password of the server you're connecting to are transmitted in plaintext. The Integrated Windows authentication option requires a Windows account name and password.

The option set here can be overridden for a specific remote domain. Overriding the authentication settings for a remote domain enables you to set the virtual server authentication level to handle most of the transmissions, while allowing exceptions for individual addresses. The following table describes several configuration examples.

| SMTP transmissions | Authentication option |
|---|---|
| Messages are commonly sent to multiple addresses. | Disable authentication for the SMTP virtual server. If attempts to deliver messages to an address fail because of authentication requirements, add a remote domain for the address. Then enable authentication for the domain at the same level required by the server. |
| Messages are commonly sent to one address, which requires authentication. | Determine what level of authentication is required to connect. Then enable authentication for the SMTP virtual server using the same level. If you want to then send messages to other addresses, set up remote domains and set different authentication options. If you use this option, it is likely that the account name used is the one that identifies the computer set up as the smart host. |

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To disable authentication for outgoing messages**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Delivery** tab, and click **Outbound Security**.
3. Click **Anonymous access**.
4. Clear all other options.

**To set Basic authentication for outgoing messages**

1. In IIS Manager, select the SMTP virtual server, and then click **Properties** on the **Action** menu.
2. On the **Delivery** tab, click **Outbound Security** to open the **Outbound Security** dialog box.
3. Click **Basic authentication**.
4. Under **User name** and **Password**, type the account name and password that will grant you access to the computer you are connecting to.

**Important**

If Basic authentication is your only authentication method, it is strongly recommended that you also require TLS encryption to avoid unauthorized detection of user names and passwords.

**To set Integrated Windows authentication for outgoing messages**

1. In IIS Manager, select the SMTP virtual server, and then click **Properties** on the **Action** menu.
2. On the **Delivery** tab, click **Outbound Security** to open the **Outbound Security** dialog box.
3. Select the **Integrated Windows Authentication** check box.
4. Under **Account** and **Password**, type a Windows account name and password that will grant you access to the computer you're connecting to.

**Related Information**

- For information about requiring authentication for incoming connections, see Requiring Authentication for Incoming Connections.
- For information about encrypting outgoing messages, see Requiring TLS Encryption.

# Setting IP Access Restrictions to Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can grant or deny Simple Mail Transfer Protocol (SMTP) virtual server access to specific IP addresses. By default, the SMTP virtual server is accessible to all IP addresses. You can set restrictions by specifying a single IP address, a group of addresses using a subnet mask, or a domain name.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set IP address access restrictions**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Access** tab, and under **Connection control**, click **Connection**.
3. Click either **Only the list below** or **All except the list below**.
4. To add to the list of computers, click **Add**.
5. To delete from the list of computers, select a listing, and then click **Remove**.

**Related Information**

- For information about setting or removing relay restrictions from a virtual server, see Configuring SMTP Virtual Server Relay Restrictions.
- For information about requiring TLS encryption, see Requiring TLS Encryption.

# Configuring SMTP Virtual Server Relay Restrictions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS includes a full-featured SMTP virtual server that you can use to receive and relay e-mail messages to other SMTP servers on your network or to servers on the Internet. The relay function is useful for internal network clients that might have to forward mail to other SMTP servers, and it is useful for IIS programs that need access to an SMTP server to forward mail.

For a user or computer to relay e-mail messages through an SMTP virtual server, the following two conditions must be met:

- The user or computer can access the SMTP virtual server.
- The SMTP virtual server is configured to relay e-mail messages to other domains.

When an SMTP virtual server relays e-mail messages, it can forward mail that is addressed to any e-mail domain. With this feature, an SMTP virtual server can forward mail to any internal or external network SMTP server for which it can resolve an MX record. However, if the SMTP virtual server is accessible to Internet users, mail relay should not be enabled. With mail relay enabled, malicious users might forward e-mail to your SMTP virtual server, distributing unwanted messages to other computers and reducing the available bandwidth for your internal connection.

By default, the SMTP service blocks computers from relaying unwanted mail through the virtual server. To enable relay access through the SMTP virtual server, click **Relay** on the **Access** tab. By default, all computers are blocked except those that meet the authentication requirements that are designated in the **Authentication** box, which you can view by clicking **Authentication** on the **Access** the tab.

You can also allow messages to be relayed to a specific remote domain. The domain setting overrides the SMTP virtual server setting. For more information about relaying messages to a remote domain, see Configuring Remote Domains.

If you enable mail relay on your SMTP virtual server, then you can specify the relay restrictions that are described in the following table.

| Option | Description |
|---|---|
| **Only the list below** | This option allows only the computers specified in the list to relay messages through the SMTP virtual server. |
| **All except the list below** | This option allows all computers, except the computers that are specified in the list, to relay messages through the SMTP virtual server. This option is set by default, along with the **Allow all computers which successfully authenticate to relay, regardless of the list above** option. |
| **Add** and **Remove** | Clicking these buttons allows you to grant or deny relay access by adding to or removing from the list of computers. |
| **Allow all computers which successfully authenticate to relay, regardless of the list above** | This option allows computers that meet authentication requirements set in the **Authentication** box to relay messages to the SMTP virtual server. This option is set by default. |

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To add relay restrictions to an SMTP virtual server**

1. In IIS Manager, right-click the SMTP Virtual Server for which you want to add relay restrictions, and then click **Properties**.
2. Click the **Access** tab, and then click **Authentication**.
3. To enable the appropriate level of authentication for your server, select either (or both) the **Basic authentication** or the **Integrated Windows Authentication** check box, clear the **Anonymous access** check box, and then click **OK**.

| **Note** |
| --- |
| If you enable Anonymous access and do not enable Basic authentication and Integrated Windows authentication, then authentication is no longer enabled, which means that all users and computers can access the SMTP virtual server. |

4. On the **Access** tab, under **Relay restrictions**, click **Relay**.
5. In the **Relay Restrictions** box, click **Add**, and then do the following to add a single computer, a group of computers, or a domain:
   - To add a single computer, click **Single computer**, type the IP address of the computer that you want to add, and then click **OK**.
   - To add a group of computers, click **Group of computers**, type the subnet address and the subnet mask of the group into the corresponding boxes, and then click **OK**.
   - To add a domain, click **Domain**, type the domain name that you want to add, and then click **OK**.
6. To apply your configuration changes, click **OK** twice.

**To remove relay restrictions from an SMTP virtual server**

1. In IIS Manager, right-click the SMTP virtual server for which you want to remove relay restrictions, and then click **Properties**.
2. Click the **Access** tab, and then click **Relay**.
3. In the **Relay Restrictions** box, select either the **Only the list below** or the **All except the list below** check box.
4. If you want to add exceptions, click **Add** and then specify the computer, group of computers, or domain for which you want to retain relay restrictions.

**Related Information**

- For information about setting IP access restrictions to the server, see Setting IP Access Restrictions to Servers.
- For information about requiring TLS encryption, see Requiring TLS Encryption.

# Requiring TLS Encryption

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can require that all clients use Transport Layer Security (TLS) encryption, a generic security protocol similar to Secure Sockets Layer (SSL), to connect to the default Simple Mail Transfer Protocol (SMTP) virtual server. This option secures the connection, but it is not used for authentication.

When requiring Basic authentication on your virtual servers, it is strongly recommended that you also use TLS encryption. Without encryption, user names and passwords can be easily intercepted.

To use TLS encryption for the virtual server, you must create key pairs and configure key certificates. Clients can then use TLS to encrypt the session with the SMTP service, therefore, all messages sent. The SMTP service can also use TLS to encrypt sessions with remote servers.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**.

**Procedures**

**To create and manage key certificates**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Access** tab, and under **Secure communication**, click **Certificate** to set up new key certificates and manage installed key certificates for the SMTP virtual server.

Key pairs consist of a number of bits that indicate the key's security level. You can strengthen security by increasing the encryption level from 40 bits (the default) to 128 bits. The greater the number of bits, the more difficult the item is to decrypt. Users attempting to secure access must use the same encryption level that you set or messages will be returned with a non-delivery report (NDR).

**To set TLS encryption levels for the server**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Access** tab, and under **Access control**, click **Authentication**.
3. Click **Basic authentication**.
4. Select the **Require TLS encryption** check box.

**Note**

There are two additional TLS options available. To use TLS for all outgoing connections, click **Outbound Security** on the **Delivery** tab, and then click **TLS encryption**. Also, if a server you commonly connect to requires the use of TLS for all incoming connections, you can create a remote domain and click **TLS encryption** when creating the domain.

**Related Information**

- For information about setting IP access restrictions to the server, see Setting IP Access Restrictions to Servers.
- For information about setting or removing relay restrictions from a virtual server, see Configuring SMTP Virtual Server Relay Restrictions.

# Delivering Messages

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

After a connection has been opened and the receiving server has acknowledged that it is ready to receive data, messages can be transmitted for delivery. This section describes the settings that determine how delivery occurs.

This section includes the following information:

- Message Delivery Options: Provides an overview of routing, transmission, and security options.
- Configuring Retry Attempts and Intervals: Describes how to configure the server for mail that can't be delivered on the first attempt.
- Setting the Message Hop Count: Describes how to limit the number of servers a message is allowed to pass through before delivery.
- Setting the Masquerade Domain: Describes how to replace a local domain name used in the Mail From lines.
- Setting Fully Qualified Domain Names: Describes the concept of fully qualified domain names and how to change them.
- Configuring Smart Hosts: Describes how to route all outgoing messages for remote domains through a smart host instead of sending them directly to the domain.
- Enabling Reverse DNS Lookup: Describes how to enable reverse DNS lookup.

# Message Delivery Options

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the **Delivery** tab to set all delivery and routing options. Settings can be grouped into three categories, which are listed below.

**Routing Options**

- Smart host: Designates a server through which to route all outgoing messages.
- Message hop count: Determines the maximum number of servers a message is routed through before being considered undeliverable.
- Fully qualified domain name (FQDN): Clarifies the address to use in message exchanger (MX) records.

**Transmission Options**

- Retry attempts and retry interval: Determines how many times to resend a message, and at what intervals, before the message is considered undeliverable.
- Delivery using Pickup directory: Allows you to transmit messages composed as a text file.

**Security Options**

- Outbound security: Allows you to use authentication and Transport Layer Security (TLS) encryption for outgoing messages.
- Reverse DNS lookup: Verifies that the message actually originated from the computer and the domain listed in the From field.
- Masquerade domains: Replaces any local domain name used in any Mail From lines in the protocol with a different domain name. This is the name you want stamped on outgoing messages.

# Configuring Retry Attempts and Intervals

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If a message cannot be delivered on the first attempt, the SMTP service sends it again from the Queue directory after a specified time. You can set the interval between delivery attempts. You can also designate the number of times to attempt to deliver a message. After the limit is reached, the message is returned to the sender with a non-delivery report (NDR) and copies of the message and NDR are sent to the location you designate. The NDR is placed in the Queue directory and goes through the same delivery process as messages. When the NDR reaches the maximum number of retry attempts, the NDR and message are sent to the Badmail directory.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure retry attempt and interval settings**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Delivery** tab.
3. In the **First retry interval (minutes)** box, type a value for the amount of time to wait before retrying message delivery. The default is 15 minutes.
4. In the **Second retry interval (minutes)** box, type a value for the amount of time to wait before retrying message delivery. The default is 30 minutes, which is 30 minutes after the **First retry interval**.
5. In the **Third retry interval (minutes)** box, type a value for the amount of time to wait before retrying message delivery. The default is 60 minutes, which is 60 minutes after the **Second retry interval**.
6. In the **Subsequent retry interval (minutes)** box, type a value for the amount of time to attempt delivery before posting a notification. The default is 240 minutes.

**Delay Notification**

To allow for network delays, you can set a delay period to expire before sending the delivery notification. The minimum value is 1 minute, the default is 12 hours, and the maximum value is 9999 days. Use the drop-down menu beside the value field to use minutes, hours, or days.

**Expiration Timeout**

Type a value for messages that have not been delivered after all retries and delays have expired. The minimum value is 1 minute, the default is 2 days, and the maximum value is 9999 days. Use the drop-down menu beside the value field to use minutes, hours, or days.

# Setting the Message Hop Count

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When a message is delivered, it may be routed to a number of servers before reaching its final destination. You can designate how many servers the message is allowed to pass through. This is called the *hop count*.

After the hop count is set, the Simple Mail Transfer Protocol (SMTP) server counts the hops listed in the Received lines of the message header. When the number of Received fields exceeds the maximum hop count setting, the message is returned to the sender with a non-delivery report (NDR).

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set the message hop count**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Delivery** tab, and click **Advanced**.
3. In the **Maximum hop count** box, type a value for the number of hops a message can take between the source and destination servers. The default is 15 hops.

**Related Information**

- For information about designating a server for all outgoing messages, see Configuring Smart Hosts.
- For information about setting a fully qualified domain name, see Setting Fully Qualified Domain Names.

# Setting the Masquerade Domain

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The masquerade domain replaces any local domain name used in any Mail From lines in the protocol. The replacement occurs on the first hop only.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set the masquerade domain**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Delivery** tab, and click **Advanced**.
3. In the **Masquerade domain** box, type a domain name that you want to appear in message headers, instead of the actual name of the domain.

| **Note** |
| --- |
| All replies to such messages will be routed through the SMTP virtual server that uses the masquerade domain. |

**Related Information**

- For information about enabling reverse DNS lookup, see Enabling Reverse DNS Lookup.

# Setting Fully Qualified Domain Names

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

An e-mail domain must be able to be resolved through Domain Name System (DNS). There are two DNS records that are used to resolve an e-mail domain. In most cases, a mail exchanger (MX) record is set up to associate an e-mail domain with the fully qualified domain name (FQDN) of one or more Simple Mail Transfer Protocol (SMTP) virtual servers that serve that domain. Each SMTP server referenced in the MX record must have an address (A) record. The A record maps a given FQDN to its IP address.

It is possible to just have an A record set up for an e-mail domain. In this scenario, the A record maps the domain to the IP address or addresses of the SMTP virtual server or servers that serve that domain. Adding an MX record, however, is recommended over using an A record by itself, because an MX record allows an SMTP administrator to specify an ordered list of servers to use for clients sending mail to that e-mail domain. Microsoft SMTP Service always checks first for an MX record before falling back to an A record, so setting up MX records on your virtual server can improve performance. And, in some cases, the A record is used for other purposes, such as HTTP, although the MX record is generally used only for SMTP. The MX record allows one server to handle http://example.com (HTTP clients use the A record) and another server to handle someone@example.com (SMTP clients use the MX record).

On the SMTP service, there are two options for specifying an FQDN. You can use the name specified on the **Network Identification** tab of System Properties in Control Panel, or you can specify a unique FQDN for the SMTP virtual server you are configuring.

At startup, the name designated on the **Network Identification** tab of System Properties is automatically used for the FQDN. If you change the name (either manually or by joining a domain), the new name is automatically used for the FQDN the next time the computer is restarted. No action is required to update the FQDN for the virtual server.

To override the automatic use of the computer and domain names on the **Network Identification** tab, change the FQDN in the **Advanced Delivery** dialog box, accessed through the **Delivery** tab. The SMTP service can then use the designated name instead of the one specified on the **Network Identification** tab.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To change the FQDN**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Delivery** tab, and click **Advanced**.
3. In the **Fully-qualified domain name** box, type the FQDN.

**Related Information**

- For information about setting the message hop count, see Setting the Message Hop Count.
- For information about designating a server for all outgoing messages, see Configuring Smart Hosts.

# Configuring Smart Hosts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can route all outgoing messages for remote domains through a smart host instead of sending them directly to the domain. This enables you to route messages over a connection that may be more direct or less costly than other routes. The smart host is similar to the route domain option for remote domains. The difference is that, after a smart host is designated, all outgoing messages are routed to that server. With a route domain, only messages for the remote domain are routed to a specific server.

| **Important** |
| --- |
| Make sure your designated smart host is secure and administered by a trusted authority, especially when forwarding sensitive information. |

If you set up a smart host, you can still designate a different route for a remote domain. The route domain setting overrides the smart host setting.

You can identify the smart host by fully qualified domain name (FQDN) or an IP address (but if you change the IP address, you would have to change it on every virtual server as well). If you use an IP address, enclose it in brackets ([ ]) to increase system performance. The SMTP service checks first for a server name, and then an IP address. The brackets identify the value as an IP address, so the DNS lookup is bypassed.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set up a smart host**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Delivery** tab, and click **Advanced**.
3. In the **Smart host** box, type the name of the smart host server. You can type a string to represent a name or enter an IP address.
4. If you want the SMTP service to attempt to deliver remote messages directly before forwarding them to the smart host server, select the **Attempt direct delivery before sending to smart host** check box. The default is to send all remote messages to the smart host, not to attempt direct delivery.

**Related Information**

- For information about setting the message hop count, see Setting the Message Hop Count.
- For information about setting a fully qualified domain name, see Setting Fully Qualified Domain Names.

# Enabling Reverse DNS Lookup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you select this option, the SMTP service will attempt to verify that the client's IP address matches the host/domain submitted by the client in the EHLO/HELO command. If the reverse DNS lookup is successful, the RECEIVED header will remain intact. If the verification is unsuccessful, "unverified" appears after the IP address in the RECEIVED header of the message. If the reverse DNS lookup fails, "RDNS failed" will appear in the RECEIVED header of the message.

Because this feature verifies addresses for all incoming messages, its use could affect SMTP service performance. Clear the check box to disable the feature.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

### To enable reverse DNS lookup

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Delivery** tab, and click **Advanced**.
3. Select the **Perform reverse DNS lookup on incoming messages** check box.

**Related Information**

- For information about setting the masquerade domain, see Setting the Masquerade Domain.

# Managing Messages_IIS_SP1_Ops

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

This section describes how to set connection limits to maximize performance, how to set limits for message size and number of recipients, and how to manage undeliverable messages.

This section includes the following information:

- Setting Message Size Limits: Describes how to set message and session limits.
- Setting Recipient Limits: Describes how to set the maximum number of recipients for a single message sent in one connection.
- Limiting the Number of Messages per Connection: Describes how to limit the number of messages sent in a single connection.
- Storing Non-Delivery Reports: Describes how to set a storage location for non-delivery reports (NDRs).

# Setting Message Size Limits

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are two message size limit settings. The first, **Limit message size to**, is a preferred message limit for the virtual server. This is what the SMTP service will advertise as the maximum message size this Simple Mail Transfer Protocol (SMTP) virtual server will accept. If a mail client sends a message that exceeds the limit, they will get an error. If a remote server supports EHLO, it will detect the advertised maximum message size value when it connects to the SMTP virtual server and won't attempt to deliver a message that exceeds the limit. Instead it will simply send a non-delivery report (NDR) to the sender of the message. A remote server that does not support EHLO will try to send a message that exceeds the size limit, but will still end up sending an NDR to the sender when the message doesn't go through.

**Limit session size to** is the maximum amount of data accepted during the total connection. This is the sum of all messages sent during the connection (applying to the message body only).

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set message size limits**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Messages** tab, select the **Limit message size to (KB)** check box, and then type a value (in kilobytes) for the maximum size of a message. The default is 2048 KB. The minimum value is 1 KB.
3. Select the **Limit session size to (KB)** check box, and then type a value to indicate the maximum total size (in kilobytes) of all messages in a given connection. This number will always be larger than the maximum message size and should be set carefully because the connecting message transfer agent (MTA) is likely to resubmit the message repeatedly. The default size is 10240 KB. This value should be greater than or equal to the value entered for **Limit message size to (KB)**.

**Related Information**

- For information about setting recipient limits, see Setting Recipient Limits.
- For information about limiting the number of messages per connection, see Limiting the Number of Messages per Connection.
- For information about storing non-delivery reports, see Storing Non-Delivery Reports.

# Setting Recipient Limits

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can determine the maximum number of recipients for a single message sent in one connection. The default is 100, which is the minimum number specified in Request for Comment (RFC) 821. Many clients return messages with a non-delivery report (NDR) after an error message is received, indicating that the maximum number of recipients has been exceeded. The SMTP service does not return messages in this instance. It opens a new connection immediately and processes the remaining recipients. For example, if the recipient limit is set to 100 and a message with 105 recipients is received, the first 100 are delivered in one connection. Then, a new connection is opened, and the message is processed for the remaining five recipients.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set recipient limits**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Messages** tab, select the **Limit number of recipients per message to** check box, and then type a number to represent the recipient limit. To impose no limit, clear the check box.

**Related Information**

- For information about setting message size limits, see Setting Message Size Limits.
- For information about limiting the number of messages per connection, see Limiting the Number of Messages per Connection.
- For information about storing non-delivery reports, see Storing Non-Delivery Reports.

# Limiting the Number of Messages per Connection

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This option enables you to limit the number of messages sent in a single connection. It also provides a method to improve system performance by allowing the use of multiple connections to deliver messages to a remote domain. When the set limit is reached, a new connection is automatically opened and the transmission continues until all messages are delivered.

For example, if you commonly send a large number of messages to certain remote domains, you could set the **Limit number of messages per connection to** value to a relatively small number, such as 20. As a result, when sending 100 messages in one session, the SMTP service immediately opens a new connection after the first 20 are sent, another after the next 20 are sent, and so on. In this case, there could be up to five simultaneous connections delivering queued messages to one destination. Message delivery would be faster because fewer messages are delivered simultaneously, instead of in one long stream over one connection.

To determine a value for the limit, review the **Messages Sent/sec** performance counter for the SMTP Server object in System Monitor. The **Limit number of messages per connection to** value should be less than the value indicated by the performance counter. If the counter indicates a value of 30, and you set your maximum connections to 50, no simultaneous connection would be opened because the server would not exceed 30 messages per second. It would work as though the messages were sent in one long stream over one connection. This setting affects outgoing messages only. You can use it to increase your server output speed, but it has no effect on the rate that other servers process incoming messages.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To limit the number of messages sent in one connection**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Messages** tab, select the **Limit number of messages per connection to** check box, and then type a value in the box. The default is 20.

**Related Information**

- For information about setting message size limits, see Setting Message Size Limits.
- For information about setting recipient limits, see Setting Recipient Limits.
- For information about storing non-delivery reports, see Storing Non-Delivery Reports.

# Storing Non-Delivery Reports

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When a message is undeliverable, the SMTP service returns it to the sender with a non-delivery report (NDR). You can also designate that copies of the NDR be sent to a location of your choice. If the NDR cannot be delivered to the sender, a copy of the message is put in the Badmail directory.

All NDRs go through the same delivery process as other messages, including attempts to resend the message. If the NDR has reached the retry limit and cannot be delivered to the sender, a copy of the message is placed in the Badmail directory. Messages placed in the Badmail directory cannot be delivered or returned.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

### To set a storage location for NDRs

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **Messages** tab, in the **Send copy of Non-Delivery Report to** box, type the e-mail address of the mailbox you want to use to store copies of the NDRs. This address can be any valid SMTP e-mail address and is optional. To disable the feature, clear the text box.
3. In the **Badmail directory** box, type the directory you want to use to store undeliverable messages. The default location is *X*:\Inetpub\Mailroot\Badmail, where *X* is the drive on which IIS is installed. Click **Browse** to select another folder. You can designate a different directory, provided it is on the same computer as the SMTP service.

**Related Information**

- For information about setting message size limits, see Setting Message Size Limits.
- For information about setting recipient limits, see Setting Recipient Limits.
- For information about limiting the number of messages per connection, see Limiting the Number of Messages per Connection.

# LDAP Routing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Lightweight Directory Access Protocol (LDAP) is an Internet protocol that enables you to access directory information from an LDAP server. With permission to use LDAP, you can browse, read, and search directory listings on the LDAP server.

With the **LDAP Routing** tab, you can configure the SMTP service so that it will consult an LDAP server to resolve senders and recipients. For example, you can use Active Directory directory service as an LDAP server, and use Active Directory Users and Computers to create a group mailing list that is automatically expanded on the Simple Mail Transfer Protocol (SMTP) virtual server.

This section describes how to configure the SMTP service as an LDAP client so you can gain access to directory information on an LDAP server.

This section includes the following information:

- Determining Authentication Levels: Describes the four binding types that determine the level of authentication used for connections.
- Preliminary Steps to Configuring LDAP Connections: Describes how to plan before configuring Lightweight Directory Access Protocol (LDAP).
- Setting LDAP Routing Information: Describes how to enable Lightweight Directory Access Protocol (LDAP) routing, and set options for LDAP routing.
- Selecting Schema Types: Describes the schema types available from the **Schema** drop-down menu on the **LDAP Routing** tab.

# Determining Authentication Levels

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are four binding types that determine the level of authentication used for connections. The option selected affects other options.

| **Binding type** | **Description** |
|---|---|
| Anonymous | No authentication is required. Select this option when the server does not require account and password information. |
| Plaintext | This option allows the transmission of account information and passwords in plaintext. When this option is selected, account and password information is required. |
| Windows SSPI | The client and server negotiate to determine which level of authentication is acceptable by both, and use the strongest method available. When this option is selected, account and password information is required. |
| Service accounts | Uses the account information of the user under which the SMTP service is running. This is controlled by **Component Services**, under **Administrative Tools**. |

# Preliminary Steps to Configuring LDAP Connections

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Because some settings on the **LDAP Routing** tab affect others, it is recommended that you plan before configuring Lightweight Directory Access Protocol (LDAP).

1. Determine the schema type to use.

2. Determine the level of authentication that will be needed to bind to the LDAP directory. If the LDAP server you connect to does not require authentication, no account information is necessary.

3. Determine if you need to specify a naming context. A naming context may be necessary if the user accounts reside in a subcontainer within the LDAP directory. Specifying a naming context may allow the LDAP directory to return results more quickly and, in some cases, is necessary if there are multiple naming contexts on the LDAP server.

**Related Information**

- For information about determining the schema type to use, see Selecting Schema Types.

- For information about determining the level of authentication, see Determining Authentication Levels.

# Setting LDAP Routing Information

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This topic describes how to enable Lightweight Directory Access Protocol (LDAP) routing, and set options for LDAP routing.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set LDAP routing information**

1. In IIS Manager, right-click the SMTP virtual server, and then click **Properties**.
2. Click the **LDAP Routing** tab.
3. Select the **Enable LDAP routing** check box.
4. Set options as described in the following table.

| Option | Description |
| --- | --- |
| Server | Type the name of the computer running the LDAP directory. The field is not applicable when you use the Exchange LDAP Service schema type, because the service will find a nearby server and use it automatically. |
| Schema | Select a schema type. |
| Binding | Select the authentication method to use when binding to the LDAP server. **Plain text** transmits passwords in plaintext. To use the strongest authentication method available for the SMTP virtual server, select **Windows SSPI**. For no authentication, use **Anonymous**, which is the default. **Service accounts** will use the account information of whom ever is trying to bind to the LDAP server. |
| Domain | Type the domain of the account you want to use to bind to the LDAP directory. This option applies if you select the **Plaintext** or **Windows SSPI** binding types. |
| User name | Type the distinguished name of the account that you want to use to bind to the LDAP directory. For example, cn=user1, ou=users, o=company. This option applies if you select the **Plain text** or **Windows SSPI** binding types. |
| Password | Type the password used for logging on to the directory service. This option applies if you select the **Plain text** or **Windows SSPI** binding types. |
| Base | Type the distinguished name of a container in the directory service you're accessing. This setting specifies where you want the SMTP service to start searching in the LDAP directory. The search will begin in the container and will continue to subcontainers if necessary. |

**Related Information**

- For information about selecting a schema type, see Selecting Schema Types.
- For information about selecting the authentication method, see Determining Authentication Levels.

# Selecting Schema Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The following schema types are available from the **Schema** list box on the **LDAP Routing** tab.

| Schema type | Description |
|---|---|
| Active Directory | Select this option to use Active Directory directory service as the Lightweight Directory Access Protocol (LDAP) server. To manage mail boxes, use Active Directory Users and Computers. |
| Site Server Membership Directory | Use this option when you have installed Microsoft LDAP Service as part of Microsoft Commercial Internet System 2.0 Mail and you are using it to manage mailboxes. |
| Exchange LDAP Service | Select this option to use a Microsoft Site Server 3.0 or later LDAP server to manage mailboxes. |

**Related Information**

- For information, see Preliminary Steps to Configuring LDAP Connections.
- For information about setting LDAP routing information, see Setting LDAP Routing Information.

# Monitoring Message Transactions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Protocol logs track the commands a Simple Mail Transfer Protocol (SMTP) virtual server receives from SMTP clients.

Four formats are available for tracking. For the most part, they track the same information. The differences among them are depth of detail, type of display, and analytical uses. When choosing a format, decide what you want to track, how you want to review the data, and how you want to track transactions from other Internet Information Services (IIS) services.

The default format is World Wide Web Consortium (W3C) Extended Log File Format, but you can choose a different one. For three of the formats, transactions are written to a text file. You can determine the logging period and maximum size of the file. When the size limit is reached, a new file is created. Back up and remove older files so the system doesn't run out of disk space when writing to the current transaction log file. There is also a format that enables you to collect data in an ODBC-compliant database.

| **Note** |
| --- |
| The log formats and default file names are the same as those used by other IIS services. Plan what you want to track for all the services, how many files to use, and how to set file size before choosing a format for the SMTP service. |

This section includes the following information:

- Log Formats: Describes the four formats available for transaction logs.
- Setting Up an ODBC-Compliant Database: Describes how to set up an open database connectivity (ODBC)-compliant dabase for protocol logging.

**Related Information**

- For information about enabling protocol logging, see Enabling Protocol Logging.

# Log Formats

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Four formats are available for transaction logs. The log formats and default file names are the same as those used by other IIS services. Plan what you want to track for all the services, how many files to use, and how to set file size before choosing a format for the SMTP service.

**Microsoft IIS Log File Format**

- Information is written to a comma-delimited ASCII text file.
- Data is fixed; you cannot customize the log.
- There are multiple records for a single transmission.

**NCSA Common Log File Format**

- Information is written to an ASCII text file that uses the National Center for Supercomputing Applications (NCSA) format.
- Data is fixed; you cannot customize the log.
- There are multiple records for a single transmission.

**W3C Extended Log File Format**

- The World Wide Web Consortium (W3C) Extended Log File Format is the default logging format for the SMTP service, as well as other IIS services.
- Data is variable; you can choose what you want to track.
- This format is a good option for limiting log size.
- Information is written to an ASCII text file.
- This format includes some field options that apply to the Web and file transfer protocol (FTP) services only.
- There are multiple entries for a single transmission.

**ODBC Logging**

- You must set up an open database connectivity (ODBC)-compliant database before using this format.
- Information is written to a database.
- There are multiple records for a single transmission.

The following table describes the fields and data types for an ODBC transaction log file.

| Field name | Data type | Description |
|---|---|---|
| ClientHost | varchar(255) | Client IP address |
| Username | varchar(255) | Client domain name |

| | | |
|---|---|---|
| LogTime | datetime | Connection date and time |
| Service | varchar(255) | Internet Information Services (IIS) service |
| Machine | varchar(255) | Computer name |
| ServerIP | varchar(50) | Server IP address |
| ProcessingTime | Integer | Processing time in milliseconds |
| BytesRecvd | integer | Bytes received by server |
| BytesSent | Integer | Bytes sent by server |
| ServiceStatus | Integer | Simple Mail Transfer Protocol (SMTP) protocol reply code |
| Win32Status | integer | Windows Server 2003 status or error code (a 0 valueindicates success) |
| Operation | varchar(255) | SMTP protocol command |
| Target | varchar(255) | Recipient |
| Parameters | varchar(255) | |

**Related Information**

- For information about monitoring message transactions, see Monitoring Message Transactions.
- For information about enabling protocol logging, see Enabling Protocol Logging.

# Setting Up an ODBC-Compliant Database

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This topic describes how to set up an open database connectivity (ODBC)-compliant database for protocol logging.

1. Create an ODBC-compliant database by using a database program such as Microsoft Access or Microsoft SQL Server-TM-.

2. Create a table in the database that contains the fields listed in the following table.

| Field name | Data type |
|---|---|
| ClientHost | varchar(255) |
| Username | varchar(255) |
| LogTime | datetime |
| Service | varchar(255) |
| Machine | varchar(255) |
| ServerIP | varchar(50) |
| ProcessingTime | integer |
| BytesRecvd | integer |
| BytesSent | integer |
| ServiceStatus | integer |
| Win32Status | integer |
| Operation | varchar(255) |
| Target | varchar(255) |
| Parameters | varchar(255) |

# SMTP Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Domains in the SMTP service are different from DNS domains and Windows domains. The SMTP service domains are used for organizing messages for delivery.

The Simple Mail Transfer Protocol (SMTP) virtual server has at least one domain--the default local domain. You can add more domains and configure them to be local or remote. You can delete any domain except the default domain.

**Local Domains**
A local domain is a DNS domain that is serviced by the local SMTP server. Any message with a local domain name that arrives at an SMTP server must be delivered locally to a Drop directory or returned to the sender with a non-delivery report (NDR). Local domains are sometimes referred to as service domains or supported domains. E-mail addresses with local domain names are often referred to as local addresses.

If the domain is local, you can designate it as default or alias. There is one default domain. It is used to stamp message headers that lack a domain specification. An alias domain is an alias of the default domain. If you add a domain and assign it as the new default, the previous default changes to an alias domain.

**Remote Domains**
Domains that are not local are known as remote or nonlocal domains, and e-mail addresses with remote domain names are referred to as nonlocal addresses. The SMTP service looks up remote domains in DNS.

If you want to set unique delivery requirements for a specific remote domain, you can add a remote domain and configure it accordingly. For example, you can add a remote domain and require that the SMTP service always use Transport Layer Security (TLS) encryption in sessions with that domain. Or, you can change the routing so that messages sent to one remote domain are routed to another remote domain. Use the **Domain Properties** dialog box to configure domains.

This section includes the following information:

- Creating SMTP Domains: Describes how to create alias and remote domains.
- Deleting SMTP Domains: Describes how to delete an SMTP domain.
- Configuring Local SMTP Domains: Introduces the concept of configuring local domains.
- Configuring Remote Domains: Introduces the concept of configuring remote domains.

# Creating SMTP Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can create two types of domains in the SMTP service: alias and remote. Alias domains allow you to create secondary domains that point to the default domain and use its settings, including the Drop directory. Any message sent to an alias domain is stamped with the default domain name.

Remote domains can be set for domains to which you commonly send messages. For each remote domain, you can set a predetermined delivery route and require that Transport Layer Protocol (TLS) encryption be used in all sessions with that domain. You can also use a wildcard character in the name so that all inclusive domains for the domain you are creating use the same settings. Use an asterisk (*) as the first character, followed by a period (.). For example, you can use the asterisk as a wildcard in the following format:

```
*.example.com
```

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

### To create an alias domain

1. In IIS Manager, expand the SMTP virtual server, right-click **Domains**, point to **New**, and then click **Domain**.
2. Use the **New SMTP Domain Wizard** to set up a local (alias) domain.

### To create a remote domain

1. In IIS Manager, expand the SMTP virtual server, right-click **Domains**, point to **New**, and then click **Domain**.
2. In **Welcome to the New SMTP Domain Wizard**, ensure that the **Remote** option is selected, and click **Next**.
3. In **Domain Name**, in the **Name** text box, type a name for the new domain, and then click **Finish**.

   You can use a wildcard character. For example, type ***.contoso.com**if you want mail to be delivered to any of the contoso.com SMTP servers.
4. Right-click the domain you just created, and then click **Properties**.
5. Select the **Allow incoming mail to be relayed to this domain**check box.
6. Click **Outbound Security**, and then configure the authentication and provide the credentials required by the SMTP server to which the smart host will connect.

**Related Information**

- For more information about alias and remote domains, see Organizing Messages Using SMTP Domains.
- For information about deleting SMTP domains, see Deleting SMTP Domains.

# Deleting SMTP Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can delete domains that you have added. You cannot delete the default domain, unless you first add an alias domain and then make that domain the default domain.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

### To delete a domain

1. In IIS Manager, expand the SMTP virtual server, and click **Domains**.
2. In the details pane, right-click the domain you want to delete, and click **Delete**.

**Related Information**

- For information about creating SMTP domains, see Creating SMTP Domains.

# Configuring Local SMTP Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section describes how to set options for existing local domains.

This section includes the following information:

- Designating Default Domains: Describes how to rename the default domain.
- Configuring Alias Domains: Describes how to set up alias domains that use the same settings as the default domain.
- Configuring Drop Directories: Describes how to identify a directory in which to drop all incoming messages.

# Designating Default Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The default domain is used to stamp messages from addresses that do not have a domain. A Simple Mail Transfer Protocol (SMTP) virtual server can have one default domain that cannot be deleted.

To name a default domain, you can use the name specified on the **DNS** tab for the TCP/IP protocol in the Network application in Control Panel. This domain name is also used for all other services. Alternatively, you can specify a unique domain to serve as the default for the SMTP service only.

At startup, the name designated on the **DNS** tab for the TCP/IP protocol in the Network application in Control Panel is automatically used for the default domain. If you change the name on the **DNS** tab, the new name is used automatically for the default domain the next time the service is started. No action is required to update the default domain for the SMTP service.

To override the automatic use of the network domain, specify the default domain on the **General** tab. The SMTP service can then use the designated name instead of the network domain.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To rename the default domain**

1. In IIS Manager, expand the SMTP virtual server, and then click **Domains**.
2. In the details pane, right-click a domain name, and then click **Rename**.
3. Type a new name for the default domain.

**Related Information**

- For information about configuring an alias domain, see Configuring Alias Domains.
- For information about configuring a drop directory, see Configuring Drop Directories.

# Configuring Alias Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can set up alias domains that use the same settings as the default domain. Messages received by the SMTP service for the alias domain are placed in the Drop directory designated for the default domain.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

### To configure an alias domain

1. In IIS Manager, expand the SMTP virtual server, right-click **Domains**, point to **New**, and then click **Domain**. The **New SMTP Domain Wizard** starts.
2. Click **Alias**, and then click **Next**.
3. In the **Name** box, type a name for the alias domain, and then click **Finish**.

**Related Information**

- For information about designating a default domain, see Designating Default Domains.
- For information about configuring a drop directory, see Configuring Drop Directories.

# Configuring Drop Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

For the default domain, you must identify a directory in which to drop all incoming messages. The SMTP service places messages for alias domains in the Drop directory used for the default domain. You can designate any directory as the Drop directory, provided it is local to the computer for the SMTP service and is not already assigned as the Pickup directory. By default, this is a subdirectory of the SMTP Mailroot directory.

Be sure your Drop directory is in a secure location and can be accessed only by trusted administrators. Otherwise, unauthorized users may be able to access messages in the directory.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

### To designate a Drop directory

1. In IIS Manager, expand the SMTP virtual server, and then click **Domains**.
2. In the details pane, right-click the default domain, and then click **Properties**.
3. In the **Drop directory** box, type the directory used to drop all incoming messages. The default location is *LocalDrive*:\Inetpub\mailroot\Drop. Or click **Browse** and select the directory. Make sure that you select the Drop directory and not the Pickup directory.

# Configuring Remote Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can set unique delivery requirements for a specific remote domain by adding a domain and configuring it accordingly. For example, you can set a predetermined delivery route, set authentication, and require that Transport Layer Security (TLS) encryption be used on all connections to the domain. If the remote domain is not specifically configured, the SMTP service will not perform any of these operations when delivering to this domain. It will, however, complete a normal DNS lookup.

This section includes the following information:

- Configuring ATRN for Remote Domains: Describes how to create a remote domain for an address that periodically reviews messages and enable ATRN for the domain.
- Configuring Authentication for Remote Domains: Describes how to configure the Simple Mail Transfer Protocol (SMTP) virtual server to provide the authentication credentials required by the remote domain.
- Configuring SMTP Virtual Server Relay for Remote Domains: Describes how to configure an SMTP virtual server to relay mail to a remote domain.
- Enabling TLS Encryption for Remote Domains: Describes how to enable TLS encryption for a remote domain.

# Configuring ATRN for Remote Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The SMTP service can hold mail for clients that connect periodically to download messages. In this case, the client issues an ATRN command. The remote Simple Mail Transfer Protocol (SMTP) service then starts sending messages to the ATRN domain. You can create a remote domain for an address that periodically reviews messages and enable ATRN for the domain.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable ATRN transmissions for a remote domain**

1. In IIS Manager, expand the SMTP virtual server, and then click **Domains**.
2. In the details pane, right-click the remote domain you want to change, and click **Properties**.
3. Click the **Advanced** tab, and then select the **Queue messages for remote triggered delivery** check box.
4. Click **Add** or **Remove** as necessary to select Windows user accounts that will be allowed to use the ATRN command in this domain.

# Configuring Authentication for Remote Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure the Simple Mail Transfer Protocol (SMTP) virtual server to provide the authentication credentials required by the remote domain. There are two types of authentication available: basic (plaintext) and Integrated Windows authentication. With the basic option, the account name and password of the server you're connecting to is transmitted in clear text. The Integrated Windows Authentication option requires a Windows account name and password. You can also disable authentication, which is the default option.

The selected option overrides the SMTP virtual server authentication setting on the **Access** tab. The following table describes several configuration examples.

| **SMTP virtual server transmissions** | **Authentication option** |
|---|---|
| Messages are commonly sent to multiple addresses. | Disable authentication for the SMTP virtual server. If attempts to deliver messages to an address fail because of authentication requirements, add a remote domain for the address. Then enable authentication for the domain at the same level required by the server. |
| Messages are commonly sent to one address, which requires authentication. | Determine what level of authentication is required to connect. Then enable authentication for the SMTP virtual server using the same level. If you want to then send messages to other addresses, set up remote domains and set different authentication options. If you use this option, it is likely that the account name used is the one that identifies the computer set up as the smart host. |

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**.

**Procedures**

### To disable authentication for outgoing messages

1. In IIS Manager, expand the SMTP virtual server, and then click **Domains**.
2. In the details pane, right-click the remote domain you want to change, and click **Properties**.
3. On the **General** tab, click Outbound Security.
4. Click **Anonymous access**, and then click **OK**.

### To set Basic authentication for outgoing messages

1. In IIS Manager, expand the SMTP virtual server, and then click **Domains**.
2. In the details pane, right-click the remote domain you want to change, and click **Properties**.
3. On the **General** tab, click **Outbound Security**.
4. Click **Basic authentication**.
5. In the **User name** box, type the user name; and in the **Password** box, type the password of the computer to which you are connecting.
6. Click **OK**.

**Important**

If Basic authentication is your only authentication method, it is strongly recommended that you also require TLS encryption to avoid unauthorized detection of user names and passwords.

**To set Integrated Windows authentication for outgoing messages**

1. In IIS Manager, expand the SMTP virtual server, and then click **Domains**.
2. In the details pane, right-click the remote domain you want to change, and click **Properties**.
3. On the **General** tab, click **Outbound Security**.
4. Click **Integrated Windows Authentication**.
5. In the **Account** box, type the Windows account name; and in the **Password** box, type the password.
6. Click **OK**.

# Configuring SMTP Virtual Server Relay for Remote Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure an SMTP virtual server to relay incoming mail to your SMTP/POP3 server. The SMTP virtual server can also accept and relay mail to other domains within your organization.

Specifying a relay server overrides the smart host setting in the Advanced Delivery box of the SMTP virtual server.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure an SMTP virtual server to relay mail to a remote domain**

1. In IIS Manager, double-click the SMTP virtual server that you want to configure, right-click **Domains**, point to **New**, and then click **Domain**. The New SMTP Domain Wizard starts.
2. Click **Remote**, and then click **Next**.
3. In the **Name** box, type a name for the remote domain, and then click **Finish**.

   | **Note** |
   | --- |
   | You can specify a single domain, or you can specify more than one domain by using an asterisk (*) as a wildcard character. The asterisk must be at the beginning of the name and be separated from the rest of the name by a period. |

4. In IIS Manager, right-click the new remote domain, and then click **Properties**.
5. On the **General** tab, select the **Allow incoming mail to be relayed to this domain** check box to allow the SMTP server to act as a mail relay.
6. On the **General** tab under **Route domain**, click **Forward all mail to smart host**, and then type the fully qualified domain name or the IP address of the internal network corporate mail server through which you would like to route messages for this remote domain.

   | **Note** |
   | --- |
   | If you use an IP address, enclose the IP address in brackets []. The brackets identify the value as an IP address, bypassing DNS lookup and increasing performance. |

7. Click **OK**, and then stop and restart the SMTP virtual server.

After you configure the remote domain, all mail that is addressed to the remote domain is relayed to the smart host that you configured. Mail that is not deliverable is stored in the Inetpub\Mailroot\Badmail folder.

# Enabling TLS Encryption for Remote Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you choose not to use Transport Layer Security (TLS) encryption for the Simple Mail Transfer Protocol (SMTP) virtual server, but you select the **TLS encryption** check box for this remote domain, SMTP will use it when connecting to the remote domain. Domain configurations always override the virtual server configurations.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable TLS encryption for a remote domain**

1. In IIS Manager, expand the SMTP virtual server, and then click **Domains**.
2. In the details pane, right-click the remote domain you want to change, and click **Properties**.
3. On the **General** tab, click **Outbound Security**.
4. Click **Integrated Windows Authentication**.
5. In the **Account** box, type the Windows account; and in the **Password** box, type the password.
6. Select the **TLS** encryption check box, and then click **OK**.

| **Note** |
| --- |
| If the remote domain does not support TLS encryption, all messages will be returned with a non-delivery report (NDR). |

# Configuring Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section identifies some of the configuration options that are available to Web-based applications through IIS Manager, including performance, caching, and debugging options.

This section includes the following information:

- About Configuring Applications: Defines Web-based applications and explains how your Web site's directory structure is related to them.
- Creating Applications in IIS 6.0: Explains how to set up a directory so that it is the starting point, or root, of an application.
- Setting Application Mappings in IIS 6.0: Describes how to associate a file name extension with the program that processes those files.
- Configuring ASP.NET Applications in Worker Process Isolation Mode: Describes how to configure Microsoft ASP.NET applications when running in worker process isolation mode. In worker process isolation mode, attributes defined in the Machine.config file can now only be defined through Internet Information Services (IIS).
- Configuring ASP Applications: Outlines how to set properties for Microsoft Active Server Pages (ASP) applications.
- Enabling ASP Debugging: Explains how to check and correct ASP pages that are run in worker process isolation mode.
- Configuring CGI Applications: Provides information about installing and configuring Common Gateway Interface (CGI) applications on the Web server.
- Configuring ISAPI Extensions: Explains how to fine-tune server performance by caching Internet Server API (ISAPI) extensions.
- Installing ISAPI Filters: Describes ISAPI filters and explains how to associate them with a Web server event.
- Installing Wildcard Application Mappings: Describes how to configure a Web site to run an ISAPI application regardless of the extension of the requested file.
- Configuring Applications to Use COM+ Services: Describes how to configure ASP applications to use component services.

**Related Information**

- For information about setting application-related server configurations that are dependent on an isolation mode, see Configuring Servers for Applications.

# About Configuring Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you create an application, you use IIS Manager to designate the application's starting-point directory, which is also called the application root, in your Web site. You can have more than one application per Web site. The default Web site that is created when you install IIS is an application starting point. Every file and directory that is beneath the starting-point directory in your Web site is considered part of the application until another starting-point directory is found. Thus, you use directory boundaries to define the scope of an application.

| **Important** |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Enabling Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0. |

IIS 6.0 supports many types of applications and their configuration options. You can create and then configure your applications for optimal performance and debugging; define a default script language; enable caching, buffering, and parent paths; and associate a file name extension with a type of program. These options can be configured using IIS Manager or the metabase.

An important change in IIS 6.0 that may affect your current Active Server Pages (ASP) applications is that parent paths are not enabled for ASP applications by default. For information about how to enable parent paths, see Configuring ASP Applications.

**Application Modes**

IIS 6.0 runs in one of two distinct modes of operation: worker process isolation mode and IIS 5.0 isolation mode. Application code runs in an isolated environment, but without the performance penalty of previous versions of IIS. Application isolation is the separation of applications by process boundaries that prevent the applications from affecting one another. Application isolation is configured differently for each of the two IIS isolation modes.

By default, after a clean install of IIS 6.0, IIS runs in worker process isolation mode. In worker process isolation mode, you have the option of pooling application processes. Also, for ASP.NET applications, application configuration is defined through the Web.config or Machine.config files. If you are running in worker process isolation mode, then you must configure the attributes defined in the Machine.config file through IIS.

**Related Information**

- For more information about the configuration options that are available to Web-based applications through IIS Manager, see Configuring Applications in IIS 6.0.
- For information about configuring application isolation modes, see Configuring Isolation Modes.
- For information about the advantages of pooling application processes, and on how to create and configure application pools, see Configuring Application Pools in IIS 6.0.

# Creating Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To create an application, you must first designate a directory as the starting point, which is also called the application root, for the application. You can then set properties for the application. Each application can have a friendly name. This name appears in IIS Manager and gives you a way to distinguish between applications. The application name is not used anywhere else.

| **Note** |
| --- |
| Web sites are root-level applications by default. When you create a Web site, a default application is created at the same time. You can use this root-level application, remove it, or replace it with a new application. |

You can remove a directory from the application boundaries so that requests to files in that directory and its subdirectories cannot start the application. Removing a directory from an application boundary does not delete the directory from either your Web site or from your computer's hard disk.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To create an application**

1. In IIS Manager, expand the local computer, right-click the starting-point directory of the application you are creating, and then click **Properties**.
2. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
3. In the **Application settings** area, click **Create**.
4. In the **Application name** box, type a name for your application.
5. In the **Execute Permissions** list box, set permissions by doing one of the following:
   - Click **None** to prevent any programs or scripts from running.
   - Click **Scripts only** to enable applications mapped to a script engine to run in this directory without having permissions set for executables. Setting permissions to **Scripts only** is more secure than setting them to **Scripts and Executables** because you can limit the applications that can be run in the directory.
   - Click **Scripts and Executables** to allow any application to run in this directory, including applications mapped to script engines and Windows binaries (.dll and .exe files).
6. Click **OK**.

| **Note** |
| --- |
| If you see the **Remove** button instead of the **Create** button, an application has already been created. |

**To remove a directory from the application**

1. In IIS Manager, expand the local computer, right-click the directory that you want to remove, and then click **Properties**.
2. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
3. In the **Application settings** area, click **Remove**.

| **Note** |
| --- |

To stop an application and unload it from memory, click **Unload**. If the **Unload** button is dimmed, you are not in the application's starting-point directory.

**Related Information**

- For information about configuring application isolation modes, see Configuring Isolation Modes.
- For more information about pooling application processes, see Configuring Application Pools in IIS 6.0.

# Setting Application Mappings in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are many programming and scripting languages that you can use to develop Web applications. Thus, IIS uses the file name extension of a requested resource on your Web site to determine which Internet Server API (ISAPI) or Common Gateway Interface (CGI) program to run to process a request. For example, a request for a file ending with a .aspx extension causes the Web server to call the ASP.NET program (aspnet_isapi.dll) to process the request. The association of a file name extension with an ISAPI or CGI program is called *application mapping*.

IIS 6.0 is preconfigured to support common application mappings. You can add or remove mappings for all of the applications on a Web site or for an individual application.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To map an extension to one or more applications**

1. In IIS Manager, double-click the local computer, right-click the **Web Sites** folder or an individual Web site folder, and then click **Properties**.

   **Note**

   Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element.

2. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
3. In the Application settings area, click **Configuration**, and then click the **Mappings** tab.
4. On the **Mappings** tab, click **Add**.
5. In the **Executable** box, type the path to the ISAPI or CGI program that will process the file. You must specify a program in a local directory on the Web server.
6. In the **Extension** box, type the file name extension that you want to be associated with the ISAPI or CGI program. When the Web server receives a URL identifying a file with this extension, it calls the associated program to process the request.
7. To allow processing for files of this type in a directory with Script permission, select the **Script engine** check box. When a directory has Script permission set (instead of Execute permission), only files associated with applications that are designated script engines can be processed in the directory.
8. Select the **Verify that file exists** check box to allow the server to check that the requested script exists before mapping an extension to an application.

**Note**

On the **Mappings** tab there is a **Verbs** column, which was formerly called **Exclusions** in IIS 4.0. Changing to the word *verbs* addresses the issue of future compatibility, as more new HTTP verbs will be added to the Hypertext Transfer Protocol.

**To remove an application mapping**

1. In IIS Manager, double-click the local computer, and then click the starting-point directory of the application you want.
2. Right-click the **Web Sites** folder, and then click **Properties**.

3. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
4. In the **Application settings** area, click **Configuration**, and then click the **Mappings** tab.
5. On the **Mappings** tab, click **Remove**.

   Requests for files with this file name extension are no longer processed in this Web site or directory.

**Related Information**

- For information about how to configure a Web site to run an ISAPI application regardless of the extension of the requested file, see Installing Wildcard Application Mappings.

# Configuring ASP.NET Applications in Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you install ASP.NET on a computer running Windows Server 2003, there are two process model modes that you can choose. By default, ASP.NET uses the IIS 6.0 worker process isolation mode. This mode allows you to configure application pools, recycle worker processes, enable Web gardening, set idle time-outs and processor affinity, and use a variety of application health maintenance features. On computers running a member of the Windows Server 2003 family, your ASP.NET applications should use the worker process isolation mode unless you encounter specific compatibility issues.

The **<processModel>** element configures the ASP.NET process model settings on an IIS Web server. The **<processModel>** section can be set only within the Machine.config file, and it affects all ASP.NET applications running on the server.

**Configuring the <processModel> Element Using IIS Manager**

| **Important** |
| --- |
| These features of IIS 6.0 are available only when running in worker process isolation mode. |

**timeout Equivalent Setting in IIS 6.0 Worker Process Isolation Mode**

The **timeout** attribute specifies the number of minutes until ASP.NET launches a new worker process to take the place of the current one. The default setting is **Infinite**. The **Recycle worker processes** entry is set to 120 minutes, as shown below, and is on by default.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

### To configure the ASP.NET attribute timeout using IIS Manager

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and click **Properties**.
2. On the **Recycling** tab, select the **Recycle worker processes (in minutes)** check box.
3. Click the up or down arrow to set the number of minutes you want to elapse before the worker process is recycled.
4. Click **OK**.

**idleTimeout Equivalent Setting in IIS 6.0 Worker Process Isolation Mode**

The worker processes or application pools in IIS 6.0 can be shut down after remaining idle for a specific amount of time. The **Shutdown worker process** entry is set to 20 minutes and is on by default.

### To configure the ASP.NET attribute idleTimeout using IIS Manager

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. Click the **Performance** tab, and under **Idle timeout**, select the **Shutdown worker processes after being idle for** check box.
3. In the **minutes** box, type the number of idle (no processing requests) minutes that you want to elapse before shutting down the idle worker process.
4. Click **OK**.

**shutdownTimeout Equivalent Setting in IIS 6.0 Worker Process Isolation Mode**

The **shutdownTimeout** attribute specifies the amount of time the worker process is given to shut itself down gracefully. The

time is expressed in hr:min:sec string format. The worker process **Shutdown time limit** is set to 90 seconds, as shown below.

**To configure the ASP.NET attribute shutdownTimeout using IIS Manager**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and click **Properties**.
2. Click the **Health** tab, and under **Shutdown time limit**, in the **seconds** box next to **Worker process must shutdown within**, type the desired number of seconds.
3. Click **OK**.

**requestLimit Equivalent Setting in IIS 6.0 Worker Process Isolation Mode**

The worker processes or application pools in IIS 6.0 can be set to recycle after a specific number of requests. This feature is not on by default.

**To configure the ASP.NET attribute requestLimit using IIS Manager**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. Select **Recycle worker process (in requests)** and type the maximum amount of virtual memory allowed before the worker process is recycled.
3. Click **OK**.

**requestQueueLimit Equivalent Setting in IIS 6.0 Worker Process Isolation Mode**

The **requestQueueLimit** attribute specifies the maximum allowed number of requests that will be queued in the Internet Server API (ISAPI) while waiting for the worker process to start after an abnormal termination. Note that this setting does not apply in the case of a clean shut down or normal cycle. The queue limit in IIS 6.0 is set in the kernel and is set to 1000 by default.

**To configure the ASP.NET attribute requestQueueLimit using IIS Manager**

1. In IIS Manager, expand the local computer, expand the **Application Pools** folder, right-click the application, and then click **Properties**.
2. Click the **Performance** tab.
3. In the **Request queue limit** section, select the **Limit the kernel request queue to** check box, and click the up and down arrows in the **requests** box to set the maximum number of queued requests.
4. Click **Apply**, and then click **OK**.

**memoryLimit Equivalent Setting in IIS 6.0 Worker Process Isolation Mode**

In IIS 6.0 worker process isolation mode, worker processes can be recycled based on both virtual memory and physical memory. Unlike the Machine.config files, the memory limit is not a percentage value, it is a finite value measured in megabytes.

Recycling based on physical or virtual memory is not on by default, but if selected, the default value is 500 megabytes (MB) for the virtual memory and 192 MB for the physical memory.

**To configure the ASP.NET attribute memoryLimit using IIS Manager**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and click **Properties**.
2. On the **Recycling** tab, under Memory recycling, select the **Maximum virtual memory (in megabytes)** or **Maximum used memory (in megabytes)** check box.
3. Click the up or down arrow to set memory limits.
4. Click **OK**.

**webGarden and cpuMask Equivalent Settings in IIS 6.0 Worker Process Isolation Mode**

The **webGarden** attribute controls Web garden mode, or using multiple affinitized worker processes. The **cpuMask** controls the number of processes and the way the Web garden works.

Since **webGarden** enables the use of multiple processes, it means that all processes will have their own copy of things like

application state, in-process session state, caches, and static data.

When using Web gardens, keep in mind the following:

- If your application uses session state, it must choose an out-of-process provider (NT Service or SQL).
- Application state and application statics are per process, not per computer.
- Caching is per process, not per computer.

Web gardens can be also set up in IIS 6.0 worker process isolation mode by increasing the maximum number of worker processes per application pool from 1 to n (as required).

**To configure the ASP.NET attribute webGarden and cpuMask**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. Click the **Performance** tab, and under **Web garden**, in the **Maximum number of worker processes** box, type the number of worker processes that you want to assign to the application pool.
3. Click **OK**.

**userName and password Equivalent Settings in IIS 6.0 Worker Process Isolation Mode**

The **userName** attribute allows the worker process to run with a different Windows identity than the default process identity. By default, this attribute is empty and the process runs as the IIS identity. If valid credentials are presented in the **userName** and **password** attributes, the process runs as the given account. When present, these attributes cause the worker process to run as the configured Windows identity. The default is blank attributes. In worker process isolation mode, the default is to run as Network Service. However, they can be set to a configurable account as follows.

**To configure the ASP.NET attributes userName and password**

1. In IIS Manager, expand the local computer, expand the **Application Pools** folder, right-click the application pool you would like to configure, and click **Properties**.
2. Click the **Identity** tab.
3. Click **Configurable**.
4. Click **Browse**, and under **Enter the object name to select**, type the account name in which you want your worker process to run, and click **OK**.
5. In the **Password** box, type the password associated with this account. If there is no password associated with the account, leave the **Password** box blank.
6. Click **OK**.

| **Important** |
| --- |
| When you configure a worker process identity for an application pool using a configurable account, you must add the account you created to the IIS_WPG group in order for the worker process to run. |

**pingFrequency and pingTimeout Equivalent Settings in IIS 6.0 Worker Process Isolation Mode**

The **pingFrequency** attribute is the time interval in standard process model format (hr:min:sec). The default is 30 seconds. The **pingTimeout** attribute is the time interval in standard process model format (hr:min:sc). The default is 5 seconds.

In IIS 6.0 worker process isolation mode, pinging is enabled by default. The default value of **pingFrequency** is 30 seconds after the response comes back from the worker process. The equivalent setting for **pingTimeout** is PingResponseTime Metabase Property, which is 90 seconds and can be changed directly in the metabase.

**To configure the ASP.NET attribute pingFrequency and pingTimeout**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.

2. Click the **Health** tab, and select the **Enable pinging** check box.
3. In the **seconds** box next to **Ping worker process**, type the number of seconds that you want to elapse between pings.
4. Click **OK**

**Related Information**

- For information about configuring the ASP.NET worker process for a hosted environment so that users cannot obtain privileged access to information on the server, see Configuring ASP.NET Process Identity.
- For more information about the ASP.NET configuration system, see ASP.NET Configuration.

# Configuring ASP Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can set properties that control performance and other factors of the Active Server Pages (ASP) applications that you have installed on your Web server. For example, you can turn on the use of session state in the application or set the default scripting language. Application properties apply to all ASP pages in the application unless you directly override the property on an individual page. You can also set these properties in the metabase. If the property in the metabase conflicts with the property you set in the property sheet then the last property setting overrides.

**Important**

To help minimize the attack surface of IIS, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0.

**In This Section**

This section includes the following information:

- Enabling Session State Persistence: Describes how to enable session state persistence for ASP applications.
- Enabling Buffering: Describes how to enable buffering for an ASP application.
- Enabling Parent Paths: Describes how to enable parent paths.
- Changing the Default ASP Language: Describes how to change the default ASP language.
- Setting the ASP Script Timeout: Describes how to change the ASP script timeout.
- Configuring ASP Template Caching: Describes how to change the settings for caching ASP script files and how to create a disk cache directory.

**Related Information**

- For information about side-by-side assemblies, see What's New in ASP.
- For information about NTFS and Web site permissions, see Securing Files with NTFS Permissions and Securing Sites with Web Site Permissions.
- For information about changing ASP configuration options programmatically, see Metabase Property Reference.
- For more information about pooling applications, see Configuring Application Pools in IIS 6.0.

# Enabling Session State Persistence

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can enable session state persistence for ASP applications. If you enable session state persistence, the server will create a new Session object for each connection and session state will be accessible, session storage will be allowed, **Session_OnStart** and **Session_OnEnd** will occur, and **ASPSessionID** cookie will be sent to the client. If you do not enable session state persistence, state access and storage are not allowed, events are not processed, and no cookie is sent. By default, session state persistence is enabled with the session timeout at 20 minutes. The corresponding metabase property is AspAllowSessionState Metabase Property.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To enable session state persistence for an ASP application**

1. In IIS Manager, double-click the local computer, right-click the **Web Sites** folder or an individual Web site folder, and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element. |

2. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
3. In the **Application settings** area, click **Configuration**, and then click the **Options** tab.
4. On the **Options** tab, in the **Application configuration** section, select the **Enable session state** check box.
5. In the **Session timeout** box, click the up or down arrow to set the number of minutes until the session expires.
6. Click **OK**.

# Enabling Buffering

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Enabling buffering on an ASP application specifies whether output from an ASP application will be buffered. All output from the application is collected in the buffer before the buffer is flushed to the client browser. If you do not enable buffering, output from an ASP script is written to the client browser as it becomes available. By default, buffering is enabled. The corresponding metabase property is AspBufferingOn Metabase Property.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To enable buffering for an ASP application**

1. In IIS Manager, double-click the local computer, right-click the starting-point directory of the application you want to configure, and then click **Properties**.
2. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
3. In the **Application settings** area, click **Configuration**, and then click the **Options** tab.
4. On the **Options** tab, in the **Application configuration** section, select the **Enable buffering** check box.
5. Click **OK**.

# Enabling Parent Paths

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Enabling parent paths specifies whether an ASP page allows paths relative to the current directory (using the ..\ notation). If set to true, this property constitutes a potential security risk, because an include path could access critical or confidential files outside the root directory of the application.

In IIS 6.0, parent paths are no longer enabled by default. This affects your application if it has a Web page that contains the #include server-side include directive and uses ".." notation to refer to a parent directory. Enabling parent paths corresponds to the metabase setting, AspEnableParentPaths Metabase Property.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To enable parent paths**

1. In IIS Manager, double-click the local computer, right-click the starting-point directory of the application you want to configure, and then click **Properties**.
2. Click the **Directory** tab, and then click **Configuration**.
3. Click the **Options** tab.
4. In the **Application configuration** section, select the **Enable parent paths** check box.
5. Click **OK**.

# Changing the Default ASP Language

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The default ASP language affects ASP files that do not use the <%@Language%> directive. The default language is Microsoft Visual Basic® Scripting Edition (VBscript). Changing the default ASP language corresponds to the metabase setting, AspScriptLanguage Metabase Property.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To change the default ASP language**

1. In IIS Manager, double-click the local computer, right-click the starting-point directory of the application you want to configure, and then click **Properties**.
2. Click the **Directory** tab, and then click **Configuration**.
3. Click the **Options** tab.
4. In the **Application configuration** section, in the **Default ASP language** box, type the language you want to use.
5. Click **OK**.

# Setting the ASP Script Timeout

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you set the ASP script timeout, you specify, in seconds, the default length of time that ASP pages allow a script to run before terminating the script and writing an event to the Microsoft Windows® Event Log. The minimum value for this property is one second and the default value is 90 seconds. The corresponding metabase property is AspScriptTimeout Metabase Property.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To change ASP script timeout**

1. In IIS Manager, double-click the local computer, right-click the starting-point directory of the application you want to configure, and then click **Properties**.
2. Click the **Directory** tab, and then click **Configuration**.
3. Click the **Options** tab.
4. In the **Application configuration** section, in the **ASP script timeout** box, click the up or down arrow to set the number of seconds before an error message is returned.
5. Click **OK**.

# Configuring ASP Template Caching

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP processes the files that contain ASP scripts, stores these processed files in a cache, and then serves the cached files to clients. Caching ASP script files enhances performance because cached ASP scripts are not compiled each time they are called. The corresponding metabase property is AspScriptFileCacheSize Metabase Property. AspScriptFileCacheSize Metabase Property specifies the number of precompiled script files to cache.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**To change the settings for caching ASP script files**

1. In IIS Manager, double-click the local computer, double-click the **Web Sites** folder, and then click **Properties**.
2. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
3. In the **Application Settings** area, click **Configuration**, and then click the **Cache Options** tab.

   **Note**

   If the **Configuration** button is not active, the directory or starting point does not contain an isolated application. You can create a new application by clicking **Create**.

4. In the **ASP File Cache** section, configure the ASP cache setting by doing one of the following:
   - To cache all requested .asp files, select the **Cache all ASP files in memory** check box. The maximum amount of files that can be cached is 2,000,000,000.
   - To set the maximum number of files cached in memory, select the **Cache limited ASP files in memory** check box, and then click the up or down arrow to set the maximum number of files. Files in excess of this number are persisted to the disk cache if you select the **Cache remaining ASP files on disk** check box. Or you can limit the total number of ASP files cached by selecting the **Cache limited ASP files on disk** check box, and then clicking the up or down arrow to set the maximum number of files.
   - To turn off the cache, select the **Do not cache ASP files** check box.
5. Click **OK**.

The ASP disk cache begins caching files into a directory when the ASP memory cache is full. Furthermore, a file has to be requested at least twice before it becomes a candidate for the ASP disk cache. The corresponding metabase property, AspMaxDiskTemplateCacheFiles Metabase Property, specifies the maximum number of compiled ASP templates that can be stored. AspDiskTemplateCacheDirectory Metabase Property contains the name of the directory that ASP uses to store compiled ASP templates to disk after overflow of the in-memory cache.

**Important**

In order for disk caching to work, ASP needs a properly configured disk cache directory in which to store the compiled ASP files. This means that the Administrators and IIS_WPG groups have full control on the disk cache directory and that the disk cache directory is a local directory on the Web server.

**To create a disk cache directory**

1. In IIS Manager, double-click the local computer, double-click the **Web Sites** folder, right-click the Web site that you want, and then click **Properties**.

2. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
3. In the **Application Settings** area, click **Configuration**.
4. Click the **Cache Options** tab, and in the **Disk cache directory** box browse to the directory used for disk caching.
5. Click **OK**.

| **Note** |
| --- |
| ASP requests might fail unless the IIS_WPG group has Read/Write and Delete permission for the ASP template cache directory.<br><br>For more information about ASP template caching and the ASP disk cache, see ASP Template Caching. |

# Configuring ASP Script Engine Caching

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP code is compiled and executed by Microsoft® Visual Basic® Scripting Edition (VBScript) by default. To improve performance and scalability of ASP code, ASP creates script engines. ASP script engines contain the object and marshalling data that VBScript.dll needs to execute the request. Each entry in the ASP script file cache can point to one or more entries in the ASP script engine cache.

**To change the settings for caching ASP script engines**

1. In IIS Manager, right-click the **Web Sites** folder, and then click **Properties**.
2. Click the **Home Directory** tab.
3. In the **Application Settings** area, click **Configuration**, and then click the **Cache Options** tab.

| **Note** |
| --- |
| If the **Configuration** button is not active, the directory or starting point does not contain an isolated application. You can create a new application by clicking **Create**. |

4. Set the value for **Number of script engines cached**.

To programmatically change the maximum number of script engines ASP pages will keep cached in memory, use the AspScriptEngineCacheMax Metabase Property.

# Enabling ASP Debugging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| For security and performance reasons, do not enable debugging in a production environment. |

You can set debugging options for ASP pages and control which types of error messages are sent. In order to enable debugging for ASP pages in worker process isolation mode, you must also configure the Access permissions and the Launch permissions for the debugger.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

### To enable ASP debugging

1. In IIS Manager, double-click the local computer, right-click the **Web Sites** folder or an individual Web site folder, and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element. |

2. Click the **Home Directory** tab, and then click **Configuration**.
3. Click the **Debugging** tab, and then select the **Enable ASP server-side script debugging** check box.
4. Click **Send detailed ASP error messages to client** if you want to send the client very detailed debugging information, or click **Send the following text error message to client** and type the text you want to send to the client.
5. Click **OK**.

Because the worker process identity runs as the Network Service account in worker process isolation mode, you must follow these directions to debug ASP applications.

| **Note** |
| --- |
| It is recommended that you do not change the worker process identity to run as LocalSystem. |

### To configure Launch and Access permissions to enable ASP debugging in worker process isolation mode for Script Debugger and Visual InterDev

1. If you haven't done so already, install Microsoft® Script Debugger or Microsoft Visual InterDev® debugger.
2. From the **Start** menu, click **Run**, and then type the following: **dcomcnfg**
3. In the **Console Root** file, double-click **Component Services**, double-click **Computers**, double-click **My Computer**, and double-click **DCOM Config**.
4. Right-click **Machine Debug Manager** and then click **Properties**.
5. Click the **Security** tab.
6. Under **Launch Permissions**, click **Edit**, click **Add**, type **Network Service**, and then click **OK**. If you are joined to a

domain, you might need to click **Locations** and change the location to the local computer.

7. Under **Access Permissions**, click **Edit**, click **Add**, type **Network Service**, and then click **OK**. If you are joined to a domain, you might need to click **Locations** and change the location to the local computer.

8. Click **OK** to apply the changes.

| **Note** |
| --- |
| To enable ASP debugging for Microsoft Visual Studio® .NET, you must add the Network Service account to the Debugger Users group. |

**Related Information**

- For more information about the Microsoft Script Debugger tool, see Microsoft Script Debugger on MSDN® Online.

- For more information about debugging in Visual Studio, see Adding Debugger Users and Configuring DCOM.

# Configuring CGI Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS 6.0 supports Common Gateway Interface (CGI) applications. This topic provides administrative information about installing and configuring CGI applications on the Web server, and setting the timeout value for a CGI script.

**Important**

To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Serve 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Enabling Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0.

CGI programs are executed when the Web server receives a URL that contains the CGI program name and any parameters that it requires. If your CGI program is compiled into an executable (.exe) file, you must give the directory that contains the program Execute permissions so that users can run the program. If your CGI program is written as a script, for example a Perl script, then you can give the directory either Execute permissions or Script permissions. To use Script permissions, the script interpreter must be marked as a script engine.

**Note**

The IIS_WPG group does not have the right to start CGI processes by default. If you create a new account and add it to the IIS_WPG group to run a worker process identity, you still must grant this new account two user rights to start CGI processes. These user rights are **Adjust memory quotas for a process** and **Replace a process level token**. For more information about how to modify user rights, see "Edit local security settings" or "Edit a security setting on a Group Policy object" in Help and Support Center for Windows Server 2003.

If you want to develop new CGI applications, Microsoft provides several alternatives. For faster execution, consider developing an Internet Server API (ISAPI) extension instead. For ease of development, consider developing an ASP.NET application.

This section contains the following information:

- Installing CGI Applications in IIS 6.0: Explains how to install CGI applications on your Web server.
- Configuring the Maximum Number of Concurrent CGI Requests: Explains how to configure the maximum number of concurrent CGI requests.
- Configuring CGI Application Timeouts: Explains how to configure timeout values for CGI applications.

# Installing CGI Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To install and configure CGI applications**

1. Set up a directory for your CGI programs. For extra security, separate your CGI programs from your content files. You do not need to name the directory Cgi-bin, although you can do so if you want. For more information, see Using Virtual Directories.

2. If your CGI programs are scripts, obtain and install the appropriate script interpreter. For example, to run Perl scripts, obtain a Perl interpreter. Windows operating systems do not provide versions of Perl, SED, or AWK. Instead, you must obtain these interpreters from third-party developers.

3. If your CGI programs are .exe files, give the directory Execute permissions. If your CGI programs are scripts, you can give the directory either Execute permissions or Script permissions. For more information, see Access Control in IIS 6.0.

   If you choose to give the directory Script permissions, you must mark the script interpreter as a script engine in the directory properties. Only interpreters that are marked as script engines are allowed to execute in the directory. Executable files (.dll and .exe files) cannot be directly executed; that is, a browser request cannot launch an executable file on the Web server by including the program name in the URL. By enabling **Scripts only** permissions and selecting the **Script engine** check box, you can safely put content files (such as .htm or .gif files) in the same directory as your CGI scripts. Content files are displayed in the browser and scripts are executed, but no one can run an unauthorized program and script commands are not displayed in the browser.

4. For CGI scripts, create an application mapping between the file name extension of your script and the script interpreter.

   IIS maps file name extensions to an interpreter. For example, if you are using Perl scripts stored in files with a .pl extension, map the .pl extension to the program that runs Perl scripts. Map .bat and .cmd files to the command interpreter (Cmd.exe). For more information, see Setting Application Mappings in IIS 6.0.

5. Set NTFS permissions. If you are using NTFS access permissions, verify that all users who need to run the program have Execute permissions for the directory. If your Web site accepts anonymous users, make sure that the anonymous user (the IUSR_*computername* account) has Execute permissions.

IIS supports most of the standard server environment variables.

For greater security, the Web server does not pass the command interpreter's special characters to a mapped CGI application. By default, the following characters cannot be included in the URL that calls the CGI application:

```
| ( , ; % < >
```

**Related Information**

- For general information about programming CGI applications, refer to the MSDN Online Library.

# Configuring the Maximum Number of Concurrent CGI Requests

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To prevent performance degradation and to decrease the possibility of security exposures on your Web server, you can configure the maximum number of simultaneous CGI requests allowed to run at one time.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

| **Important** |
| --- |
| Using Registry Editor incorrectly can cause serious problems that require reinstalling the operating system. Because Registry Editor bypasses the standard safeguards that prevent you from entering settings that are conflicting or likely to degrade performance or damage your system, exercise caution when making changes to the registry. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. For information about how to edit the registry, see "Changing Keys and Values" in Registry Editor Help. |

**Procedures**

**To configure the maximum number of concurrent CGI requests**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **Regedit.exe**, and click **OK**.
3. Navigate to the following location in the registry:
   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\W3SVC\Parameters.
4. From the **Edit** menu, point to **Add**, click **DWORD Value** and then add the following registry value:

   Name: **MaxConcurrentCgisExecuting**

   Type: **REG_DWORD**

   Data: IIS uses a default value of 256 if this registry key does not exist. You can set the value between 0 and 4294967295 (unlimited) depending on your needs. Using a low value means that fewer CGI requests can run concurrently on the Web server; however, using a high value can cause performance degradation.
5. Quit Registry Editor.

**Related Information**

- For more information about the registry key **MaxConcurrentCgisExecuting**, see WWW Service Registry Entries.
- For general information about programming CGI applications, refer to the MSDN Online Library.

# Configuring CGI Application Timeouts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Setting the timeout value for a CGI application involves manually setting the CGITimeout metabase property.

For information about this metabase property, see the CGITimeout Metabase Property.

The CGITimout property can be set at a global level and at the individual site and site element levels. This gives you the flexibility to configure an overall CGITimeout level, and then customize the setting for individual sites and site elements.

If you will be editing this setting while IIS is running, make sure to enable the Metabase Edit-While Running setting. See Enabling Edit-While-Running in IIS 6.0 for more information.

# Configuring User Rights for Nondefault Identities to Run CGI Processes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS 6.0 worker processes use the CreateProcessAsUser API to start CGI processes. The CreateProcessAsUser API must have the SE_ASSIGNPRIMARYTOKEN_NAME and SE_INCREASE_QUOTA_NAME user rights to succeed. The Network Service, Local Service, and LocalSystem user accounts have these user rights. If you change the identity of a worker process and want to run CGI processes, ensure that the new identity has these two user rights.

**To assign user rights to an account on the local computer**

1. From the **Start** menu, point to **Administrative Tools**, and then click **Local Security Policy**.
2. In the **Local Security Settings** dialog box, double-click **Local Policies**, and then double-click **User Rights Assignment**.
3. In the details pane, double-click **Adjust memory quotas for a process**. This is the SE_INCREASE_QUOTA_NAME user right.
4. Click **Add User or Group**, and, in the **Enter the object names to select** box, type the user or group name to which you want to assign the user right, and then click **OK**.
5. Click **OK** again, and then, in the details pane, double-click **Replace a process level token**. This is the SE_ASSIGNPRIMARYTOKEN_NAME user right.
6. Click **Add User or Group**, and, in the **Enter the object names to select** box, type the user or group name to which you want to assign the user right, and then click **OK**.

# Configuring ISAPI Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By adjusting several settings that are associated with the ISAPI extensions (ISAPI DLLs) that run on your server, you can fine-tune server performance. Logging usage data, such as the number of requests and the time they were made, can help you determine how to make the adjustments.

| **Important** |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Enabling Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting IIS 6.0. |

#### Caching ISAPI extensions

ISAPI DLLs can be loaded and cached so that additional requests can be processed without calling them again. Otherwise, the DLLs are unloaded after they are run. Caching reduces the time a user waits for a DLL to load, and most ISAPI extensions benefit from caching. You should clear this option only for special circumstances, such as debugging ISAPI extensions.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

### To enable caching

1. In IIS Manager, expand the local computer, and then click the starting-point directory of the application you want.
2. Right-click the directory, and then click **Properties**.
3. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
4. In the **Application settings** area, click **Configuration**, and then click the **Mappings** tab.
5. On the **Mappings** tab, select the **Cache ISAPI extensions** check box.

#### Related Information

- For more information about logging server activity, see Logging Site Activity.

# Working with ISAPI Filters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Like Internet Server API (ISAPI) extensions, ISAPI filters are programs that respond when the application server receives an HTTP request. However, unlike ISAPI extensions, ISAPI filters are always run in the server process. ISAPI filters are different from applications because they are driven by Web server events instead of a client request. You can associate an ISAPI filter with a particular Web server event; the filter is then notified every time its associated event occurs. For example, a filter could be notified when a Read or Write event occurs and then encrypt the raw data to be returned to the client.

| **Note** |
| --- |
| If an HTTP request triggers an event for which the filter is registered, the filter will receive the data contained in the request regardless of whether the request is for a file, a CGI application, or an ISAPI extension. |

This section contains the following information:

- Installing ISAPI Filters: Explains how to install ISAPI filters on a server and how to change the load order of the filters.
- Checking ISAPI Filter Status: Explains how to view the status for ISAPI filters.

# Installing ISAPI Filters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can install filters for all of the sites on an application server (global filters), and you can install filters for individual Web sites (site filters). If you install both global filters and site filters, the two filter lists are merged for the site.

When several filters have registered for the same event, they are called sequentially. Filters with a higher priority are run before filters with a lower priority. If several filters have the same priority, global filters that are set in the master properties are run before the filters that are set at the site level. Filters with the same priority at the same inheritance level are run in the order in which they were loaded. You can change the filter load order by adjusting the properties for the Web server or the Web site.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To add a filter to a Web server or Web site**

1. In IIS Manager, expand the local computer, right-click the Web server or Web site to which you want to add a filter, and then click **Properties**.
2. Click the **ISAPI Filters** tab, and then click **Add**.

| **Note** |
| --- |
| If you are adding filters to a Web site, you will not see any global filters inherited from the Web server's master properties. You will see only the filters installed for the Web site, even though both sets of filters are run. |

3. In the **Filter name** text box, type the name of the filter.
4. In the **Executable** text box, type the name of the DLL file or click **Browse** to navigate to it, and then click **OK**.
5. Click the up or down arrow to change the load order of a filter.

| **Note** |
| --- |
| You can only change the load order of filters that have the same priority. |

6. If you have added or changed a global filter, you must stop and restart the Web server to load the new filters into memory.

**Related Information**

- For information about checking the status of ISAPI filters, see Checking ISAPI Filter Status.
- For information about how to fine-tune server performance by caching ISAPI extensions, see Configuring ISAPI Extensions.
- For more information about ISAPI filters, see ISAPI Reference in the IIS Software Development Kit (SDK).

# Checking ISAPI Filter Status

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you add an ISAPI filter to a site on IIS 6.0, the Status column in the list of filters is empty. Until a request is made for the site and an attempt is made to load the filter, the status of the filter is not known. After the request is made, the status of the filter is displayed either as a red arrow if the filter did not load or as a green arrow if the filter loaded successfully.

| **Note** |
| --- |
| This behavior is different from the behavior in IIS 5.0, where the filter was loaded dynamically as soon as it was added to the site, regardless of whether a request was made to the site. |

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

### To check the ISAPI filter status

1. In IIS Manager, right-click the Web site or virtual server on which the ISAPI filter is installed, and then click **Properties**.
2. Click the **ISAPI Filters** tab, and then, in the **Status** column, check the ISAPI filter status.

   A green arrow indicates that the last ISAPI filter installed correctly. A red arrow indicates that a filter failed the last time the worker process started.

**Related Information**

- For information about adding ISAPI filters and changing the load order of filters, see Installing ISAPI Filters.
- For information about how to fine-tune server performance by caching ISAPI extensions, see Configuring ISAPI Extensions.
- For more information about ISAPI filters, see ISAPI Reference in the IIS Software Development Kit (SDK).

# Installing Wildcard Application Mappings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In IIS 6.0, administrators can configure a Web site or virtual directory to run an Internet Server API (ISAPI) application at the beginning of every request to that Web site or virtual directory, regardless of the extension of the requested file.

When an IIS Web server receives a request from a client, the Web server looks at the extension of the file that is named in the request to determine which ISAPI or Common Gateway Interface (CGI) application handles that file. For example, if a request comes in for a Web page called Default.asp, IIS knows from the application mapping list that Default.asp is handled by the Asp.dll ISAPI application.

If you want certain tasks to be executed for every client request before the requested page is sent to its mapped application, you can write an ISAPI Filter which intercepts requests as they arrive. However, ISAPI filters have the following limitations:

- They cannot access the body of a request, only the headers.
- They can only be configured per site, not per virtual directory.
- They should not contain long-running operations because ISAPI filters are synchronous.
- They are less flexible because the ISAPI filter API has less functionality than the ISAPI extension API, which is used to write ISAPI applications.
- They are difficult to manage over time.
- They are event-driven.

By using wildcard application maps, you can intercept every request before the requested page is sent to its mapped application. The effect is like having an Installing Wildcard Application Mappings that handles every file name extension. This is why the term *wildcard* is used to name this feature. Applications using wildcard application mapping can only be ISAPI applications.

A client request takes the following route:

- An incoming request is handled by any existing ISAPI filters that are installed on the root Web site. The ISAPI filters are executed in order.
- The request is sent to any existing wildcard application map applications, which are also executed in order. If the incoming request is for a Web file that is in a virtual directory, and there are wildcard application maps installed on the virtual directory and the root Web site, then only the wildcard application maps that are installed on the virtual directory are executed. If there are no wildcard application maps installed on the virtual directory, then the wildcard application maps that are installed on the Web site are executed. In other words, wildcard application maps are only inherited if none exist in the application where the requested Web file resides.
- The IIS server looks at the application mappings table to determine which ISAPI application or CGI script specifically processes the requested file.
- The response can be appended to at each stage and then sent back to the client.

**Using Wildcard Application Mappings**

Wildcard application maps can be used in the following scenarios:

- Highly secure Web applications, such as banking applications, that require that every user who makes a request is authenticated against a database of user names, passwords, and account numbers. For example, once the user is authenticated, certain rules are created that specify what information the user can access. If the wildcard application map determines that the user has no rights, it can prevent the request from going any further, or it can send the request to a

Web page other than the one that was requested. This custom authentication and authorization scheme can be implemented in an ISAPI application and installed as a wildcard application map. This way, all users must go through authentication, regardless of the file name extension they request.

- E-commerce Web sites that display advertisements targeted to the interests of individual users. For example, when returning users make a request for a page, a wildcard application map application can examine their identity and then look up items that they previously viewed or purchased. Once the wildcard script map transfers control to the page that was requested, an advertisement that is appropriate to the user can be displayed.

**Inserting Wildcard Application Mappings**

By inserting wildcard application mapping, you can instruct IIS to allow ISAPI applications known as interceptors to intercept and examine every request before the requested page is sent to its mapped application. The effect is an application mapping that handles every file name extension.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**To add a wildcard application mapping to a Web server or Web site**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click the Web site or virtual directory that you want, and then click **Properties**.
2. Click the appropriate tab: **Home Directory**, **Virtual Directory**, or **Directory**.
3. In the **Application settings** area, click **Configuration**, and then click the **Mappings** tab.
4. To install a wildcard application map, do the following:
   - On the **Mappings** tab, click **Insert**.
   - Type the path to the DLL in the **Executable** text box or click **Browse** to navigate to it, select the **Application Engine** check box if the DLL is a script engine, and then click **OK**.
5. To edit or delete a wildcard application map, do the following:
6. In the **Application extensions** list, click the script map that you want to change.
7. Click **Edit** to alter the script map, or click **Remove** to remove the script map.
8. To change the order in which the wildcard application maps execute, rearrange the wildcard application maps in the **Wildcard application maps (order of implementation)** box by clicking the **Move Up** or **Move Down** button until they match the execution order you want.

   | **Note** |
   | --- |
   | Do not list the same application more than once in your list of wildcard application maps. |

**Related Information**

- For information about developing wildcard application mappings, see Developing ISAPI Extensions and Filters and ISAPI Reference in the IIS Software Development Kit (SDK).
- For information about regular application mappings, see Setting Application Mappings in IIS 6.0.

# Configuring Applications to Use COM+ Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

COM+ provides a number of services that can be used by Active Server Pages (ASP) applications, and it provides one service, transaction processing, that can be used by ASP pages. For more information about COM+ Services in ASP, see Improvements to the ASP Programming Environment. For information about how to configure these services in COM+, either programmatically or by using Component Services Manager, see the related information in the COM+ Software Development Kit (SDK) or online at COM+ (Component Services).

| **Note** |
| --- |
| Batch Logon is required for COM+ server-activated application identities. COM+ will add this automatically when you enter a new account for the identity of a COM+ application using the UI. Batch Logon gives your application the ability to execute code across servers, but does not allow the account to logon interactively to home server. |

This section contains the following information:

- Enabling Side-by-Side Assemblies using IIS Manager: Explains how to enable side-by-side assemblies using IIS Manager.
- Configuring Launch Permissions for COM+ Server Objects: Explains how to configure launch permissions for distributed Web applications.

# Enabling Side-by-Side Assemblies using IIS Manager

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Side-by-side (SxS) assemblies allow ASP applications to specify which version of a SxS-supported system DLL to use, such as MSVCRT, MSXML, COMCTL, GDIPLUS, and so on. For example, if your ASP application relies on MSVCRT version 2.0, you can ensure that your application still uses MSVCRT version 2.0 even after service packs are applied to the server. Any new version of MSVCRT is still installed on the computer, but version 2.0 remains and is used by your application. SxS-supported DLLs are stored in %WINDIR%\WinSxS.

Configuring SxS assemblies requires that you know the path to the DLL, and that the COM+ manifest file exists in every virtual directory that needs to use the DLL. The COM+ manifest is an XML file that has information about where a DLL is installed. IIS does not verify that the manifest exists.

You can enable side-by-side assemblies on the IIS side programmatically or by using IIS Manager.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To enable side-by-side assemblies on the IIS side using IIS Manager**

1. In IIS Manager, right-click a Web site or virtual directory and click **Properties**.
2. Click the **Virtual Directory** tab, and click **Configuration**.
3. If the configuration button is not available, it is because you have not created an application for this virtual directory. Click **Create** to create an application.
4. In the **Application Configuration** dialog box, click the **Options** tab.
5. Select the **Enable Side by Side assemblies** check box.
6. In the **Manifest file name** box, type the name of the COM+ manifest file.
7. Click **OK** twice.

**Related Information**

For information about enabling side-by-side assemblies on the IIS side programmatically, see Improvements to the ASP Programming Environment.

# Configuring Launch Permissions for COM+ Server Objects

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

Changes in Windows Server 2003 Component Object Model (COM) services and IIS 6.0 have security implications for COM+ objects in ASP pages, ISAPI extensions, or any other application that runs on IIS. To minimize security risks, the new process model for IIS 6.0 requires all worker processes to run under an identity other than LocalSystem. If the identity of the worker process does not have launch permissions for that object, the object will not be successfully created. You must grant sufficient user rights to allow the IIS worker process to successfully create the object without granting so many rights that the identity poses a security risk.

**Important**

For Windows Server 2003, the IUSR_*computername* and IWAM_*computername* accounts are created and granted launch permissions to DCOM components by default when installing IIS. For Windows Server 2003 Service Pack 1 (SP1), the accounts are still created, but are no longer granted launch permissions to DCOM components by default when installing IIS. If you had IIS installed prior to upgrading to SP1, the permissions will not be changed. However, if you install IIS after upgrading to SP1, COM objects will fail to launch under the IUSR_*computername* and IWAM_*computername* accounts if your application uses default launch permissions for those accounts. In this situation you will receive an access denied error. Use the procedure below to grant explicit launch permissions to specific user(s) or group(s) to launch COM objects instead of using default launch permissions.

The following recommendations can provide sufficient security while reducing administrative overhead:

- Instead of adding individual users directly to the launch permissions for an object, create a group on the computer on which the COM+ object runs, and then add the individual users to this group. Then, add this group to the launch permissions for the object.

- If a COM+ object was created from the **GetExtensionVersion** method of an ISAPI extension, from an ISAPI filter, or from an ISAPI extension that is running as the identity of the current process, add the IIS_WPG group to the launch permissions for the object.

  **Note**

  You cannot use this approach with COM+ objects that are running in-process.

**Procedures**

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**To configure launch permissions for distributed Web applications**

1. Click **Start**, and then click **Control Panel**.
2. Double-click **Administrative Tools**, and then click **Component Service**.
3. In the left pane of Component Services, double-click **Component Services**, double-click **Computers**, and then click **My Computer**.
4. Click the **DCOM Config** folder.
5. In the details pane, find the object for which you want to modify the default launch permissions, right-click the icon for the object, and then click **Properties**.
6. Click the **Security** tab.

7. Under **Launch Permissions**, click **Customize**, and then click **Edit**.

| **Important** |
| --- |
| You should not modify the default launch permissions for an application. |

# Configuring IIS to Run 32-bit Applications on 64-bit Windows

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Windows Server 2003™, Service Pack 1 enables IIS 6.0 to run 32-bit Web applications on 64-bit Windows using the Windows-32-on-Windows-64 (WOW64) compatibility layer. IIS 6.0 using WOW64 is intended to run 32-bit personal productivity applications needed by software developers and administrators, including 32-bit Internet Information Services (IIS) Web applications. For more information about running 32-bit applications on 64-bit Windows, see Running 32-bit Applications on 64-bit Windows.

**To enable IIS 6.0 to run 32-bit applications on 64-bit Windows**

1. Open a command prompt and navigate to the %systemdrive%\Inetpub\AdminScripts directory.
2. Type the following command:

   **cscript.exe adsutil.vbs set W3SVC/AppPools/Enable32BitAppOnWin64 "true"**

3. Press ENTER.

**See Also**
**Concepts**
Configuring IIS to Run 32-bit FrontPage Server Extensions on 64-bit Windows
Configuring IIS to Run 32-bit ASP.NET Applications on 64-bit Windows

# Configuring IIS to Run 32-bit FrontPage Server Extensions on 64-bit Windows

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you intend to run 32-bit applications that use FrontPage Server Extensions on 64-bit Windows, you must configure IIS to create 32-bit worker processes. For more information about running 32-bit applications on 64-bit Windows, see Running 32-bit Applications on 64-bit Windows.

IIS cannot run 32-bit and 64-bit applications concurrently on the same server.

**To enable IIS 6.0 to run 32-bit FrontPage Server Extensions on 64-bit Windows**

1. Open a command prompt and navigate to the %systemdrive%\Inetpub\AdminScripts directory.
2. Type the following command:

   **cscript.exe adsutil.vbs set W3SVC/AppPools/Enable32BitAppOnWin64 "true"**
3. Press ENTER.
4. Download and install the FrontPage Server Extensions 2002 Web Download.

**See Also**
**Concepts**
Configuring IIS to Run 32-bit Applications on 64-bit Windows

# Configuring IIS to Run 32-bit ASP.NET Applications on 64-bit Windows

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you intend to run 32-bit ASP.NET applications on 64-bit Windows, you must configure IIS to create 32-bit worker processes. For more information about running 32-bit applications on 64-bit Windows, see Running 32-bit Applications on 64-bit Windows.

IIS cannot run 32-bit and 64-bit applications concurrently on the same server.

**To enable IIS 6.0 to run 32-bit ASP.NET applications on 64-bit Windows**

1. Open a command prompt and navigate to the %systemdrive%\Inetpub\AdminScripts directory.
2. Type the following command:

   **cscript.exe adsutil.vbs set W3SVC/AppPools/Enable32BitAppOnWin64 true**

3. Press ENTER.
4. Download and install the Microsoft .NET Framework Version 1.1 Redistributable Package.

**See Also**
**Concepts**
Configuring IIS to Run 32-bit Applications on 64-bit Windows

# Synchronizing the Application Path of a Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you use **SqlServer** or **StateServer** session state mode, session state might be lost when you run an ASP.NET Web application in a Network Load Balancing server cluster. To maintain session state across different Web servers in the cluster, the application path of the Web site (for example, \LM\W3SVC\2) in the IIS metabase must be the same for all of the Web servers in the cluster.

You can change the Web site instance ID by manually editing the metabase on each Web server or by using the script Moveinstance.vbs as described in this section. For more information about manually editing the IIS metabase, see Working with the Metabase.

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters).

**To synchronize the application path of a Web site**

1. Create a text file, and then name the file Moveinstance.vbs.
2. Add the following script code to Moveinstance.vbs. This code modifies the instance IDs of the Web sites so that they are the same:

```
Dim WebService
Dim oldstr
Dim newstr
Dim args
Set args = WScript.Arguments
If args.Count < 1 Then
    Wscript.Echo "Must have original instance id and new instance id" &     chr(10)
& chr(13) & _
    "usage:  moveinstance.vbs 1 5"  & chr(10) & chr(13) & _
"Moves instance 1 to instance 5"
    WScript.Quit()
End If
Set WebService = GetObject("IIS://LocalHost/W3SVC")
oldstr = args(0) 'old instance
newstr = args(1) 'new instance
WebService.MoveHere oldstr,newstr
WebService.SetInfo
Set WebService = nothing
Set args=nothing
WScript.echo "DONE"
```

3. Save Moveinstance.vbs.
4. From the same directory in which you saved Moveinstance.vbs, run the script from a command prompt. For example, at the command prompt, type the following:

   **cscript moveinstance.vbs 1 5**

   and then press ENTER. This changes the instance ID of a Web site from 1 to 5 in the metabase.

   **Note**

When you use **StateServer** mode, make sure that the [machineKey] section in the Machine.config file has the same keys on every server in the Web farm.

For more information about session state and ASP.NET Web applications, see Preserving Session State with ASP.NET Web Applications in Network Load Balancing.

# Configuring FastCGI Applications

Updated: February 16, 2009

Applies To: Windows Server 2003

**In This Section**

This topic includes links to the following information:

- Configuring the FastCGI Extension for IIS 6.0: Describes how to configure the FastCGI module for IIS 6.0.
- Using FastCGI to Host PHP Applications on IIS 6.0: Describes how to use FastCGI to host PHP applications on IIS 6.0.

# Configuring the FastCGI Extension for IIS 6.0

Updated: February 16, 2009

Applies To: Windows Server 2003

Since version 1.0, the Microsoft® Internet Information Services (IIS) Web server has supported the Common Gateway Interface (CGI) protocol specification. CGI, in its basic form, is a standards-based protocol that allows information servers, such as IIS, to interface with external applications. FastCGI was introduced to the market to address scalability shortcomings of CGI.

This article describes the shortcomings of earlier versions of CGI and defines the need for the FastCGI extension for Internet Information Services 6.0. In addition, this article describes configurations settings of the FastCGI extension.

**CGI, ISAPI, and FastCGI on IIS 6.0**

**CGI**

As stated in the introduction, CGI is a protocol that allows information servers to interface with external applications. Because HTTP is stateless, it makes sense that any requests made over HTTP create a new instance of the external application in a new operating system process.

Within the new process, the stdin handle is remapped so that it receives request data from the client, the stdout handle is remapped so that it writes response data to the client, and the command-line and operating system environment variables are set to provide other server and request information to the CGI process.

The disadvantage with CGI on IIS is the relatively expensive process creation on Windows® operating systems. Every HTTP request creates a new process, performs the work inside the CGI application, and shuts down the process. On operating systems with very lightweight process creation, the performance is bound by the work done inside the CGI application. On operating systems such as Windows where process creation is expensive, the performance of the CGI application is bound by spinning up the new process. That is why CGI has performed well in UNIX-based systems but has not been recommended for IIS.

**ISAPI**

Despite the disadvantage of CGI on Windows, IIS can keep up with, and often surpass, the performance of other Web servers. Use of Internet Server Application Programming Interface (ISAPI) enables this performance. Unlike CGI, ISAPI is internal to the Web server process. When a new request is made for an ISAPI application, a new process is not created. Instead, the Web server calls an entry point in a dynamic-link library (DLL) that is loaded into the Web server process. If the ISAPI application is written with an understanding of how the operating system threading model works, the performance is very fast.

For many years, PHP has run on IIS, both through ISAPI and CGI implementations. However, both implementations have disadvantages when they run on IIS. As with all CGI applications, the CGI implementation of PHP has a disadvantage due to the performance characteristics of process creation on Windows. The ISAPI implementation has a disadvantage caused by threading issues.

When PHP runs as an ISAPI, it runs inside the Web server process in a highly multithreaded environment. While the PHP implementation is thread-safe, many popular extensions to PHP are not thread-safe. If you use a non-thread-safe extension to PHP with ISAPI, the server could become unstable. Hence, many applications cannot run in the ISAPI PHP implementations whereas others run very well in this environment.

**FastCGI**

FastCGI offers a compromise solution that delivers both performance and stability. FastCGI allows the host CGI process to remain alive after one request finishes so that the process may be reused for another request. Since the process can be reused many times over, the cost of process creation on Windows drops out of the equation.

The technical difference between ordinary CGI and FastCGI is that FastCGI has a layer in the process that maps the FastCGI protocol into the stdin, stdout, and the other standard streams that CGI uses. Many third-party libraries can be linked into existing CGI source code with very minor modifications that will enable them to work with FastCGI.

FastCGI on IIS runs on top of ISAPI. It is comprised of applications, the Application Manager, and FastCGI protocol support code.

Because Web servers handle multiple, concurrent requests, they must have a pool of processes available that are ready to handle incoming requests. In the FastCGI handler, this pool of processes is called an application (to avoid confusion with IIS applications, this article uses the term *process pool*). There are several properties of a process pool that you may want to

manage. For example, you probably want to specify the number of processes in the pool, or the number of requests that a process can accept before it is shut down and recycled.

The FastCGI handler supports multiple process pools. This capability lets you run more than one kind of FastCGI on one server—for example, both PHP and Ruby. You may also have multiple sites on your server, and you may not want requests for those sites to share the same processes, especially if you want the site processes to run as different users. The application manager is the part of FastCGI that handles multiple process pools.

**Configuring FastCGI**

## Script Maps

To route requests to the FastCGI handler, you must associate the FastCGI handler with IIS. In IIS 6.0, you do this by using an IIS configuration setting called a script map. Script maps associate file name extensions with the ISAPI handler that executes when that file type is requested. In addition, the script map has an optional setting that verifies that the physical file associated with the request exists before allowing the request to be processed. For security reasons, this is the default setting.

However, sometimes you need to allow a request to be processed that is not associated with a physical file. The ISAPI extension that contains the FastCGI handler is called fcgiext.dll. To configure the FastCGI handler to accept PHP requests, create a new script map that associates the ".php" file name extension with fcgiext.dll. For Ruby, map the ".rb" extension to fcgiext.dll and so on. In IIS 6.0 it is also possible to use a wildcard in the script map. In this situation, all requests are routed to a single ISAPI extension. If you create a wildcard script map for fcgiext.dll, all requests will use FastCGI, regardless of the file name extension requested.

**To create a script map for FastCGI handler on IIS 6.0**

1. Start inetmgr.exe.
2. Double-click the icon for the local computer.
3. Right-click **Web Sites** and select **Properties**.
4. Click **Home Directory** tab.
5. Click **Configuration**.
6. Click **Add**.
7. Browse to %*windir*%\system32\inetsrv\ to select fcgiext.dll as the Executable.
8. Enter **.php** in the **Extension** field.
9. Select **Limit to** and enter **GET,HEAD,POST**.
10. Select **Script Engine** and **Verify that file exists**.
11. Click **OK**.

**FCGIEXT.ini**

Once you have created a script map, you must configure the FastCGI extension settings. Use the Fcgiext.ini file found in %*windir*%\system32\inetsrv\. This file contains various configuration sections and settings that change the behavior of the FastCGI handler.

Regardless of the FastCGI application type that you want to run, you must include a Types section in your configuration file. The purpose of the Types section is to associate a file name extension with a specific process pool and process. If you configure IIS script maps to point both .php and .rb files to the FastCGI handler, you must configure the Types section to tell FastCGI how these extensions differ from each other. Following is a sample configuration file that maps:

[Types]

php=c:\php\php-cgi.exe

php:123154536=PHP Site 1

php:123154537=PHP Site 2

*=Wildcard Mapping

The first line of this file shows that PHP file name extensions are associated with c:\php\php-cgi.exe, which is the PHP CGI application handler. The second and third lines of this sample use site instance identifiers to limit which sites have access to the PHP handler.

You can create configuration on a process-name basis, a file name extension (which is optionally site-specific), and a section name. The section name supports declaring arguments for the process. Since different pools running in the same process may require different arguments, you cannot depend on the process name alone to define all the settings for a given process pool. However, to maintain compatibility with previous technical preview releases of the FastCGI handler, the path of the executable is used to create the starting configuration, as in the following sample:

[c:\php\php-cgi.exe]

QueueLength=999

MaxInstances=20

InstanceMaxRequests=500

Since this section does not contain an ExePath setting, the FastCGI handler uses the section name for the path of the executable. The section below is for PHP Site 1, which only applies to the site that uses 123154536 as its ID:

[PHP Site 1]

ExePath=c:\php\php-cgi.exe

Arguments=site1

QueueLength=1001

MaxInstances=20

IdleTimeout=200

ActivityTimeout=20

RequestTimeout=60

InstanceMaxRequests=1000

In this section, .php requests are associated on that site with c:\php\php-cgi.exe. Additionally, when a process starts, "site1" is passed as an argument. This differs from the procedures followed on PHP Site 2, which looks like the following:

[PHP Site 2]

ExePath=c:\php\php-cgi.exe

Arguments=site2

QueueLength=1001

MaxInstances=20

IdleTimeout=400

ActivityTimeout=40

RequestTimeout=120

InstanceMaxRequests=1000

In this case, you pass "site2" as the argument when you start c:\php\php-cgi.exe.

**Note:** The Arguments setting is optional, and you can pass multiple arguments using a space delimiter.

If you need a space within a single argument, use quotation marks around the complete argument, just as you would do when using the command line.

**Note:** Every time that you change the FastCGI extension configuration in the Fcgiext.ini file, you must restart the application pools associated with Web sites that use FastCGI extension before the configuration changes will take effect.

**Application Pool Settings**

FastCGI extension has a set of configuration settings that control application pool behavior. This section lists all the settings supported by FastCGI and their formats:

- ExePath – This setting specifies the physical path of the process executable that will be used in the pool.

- Arguments – This setting specifies the arguments to pass to each process in the pool at start time. This setting is optional.

- EnvironmentVars – This setting specifies the environment variables that are set for the process executable associated with this pool. This setting uses the following format:

  EnvironmentVars=Name:Value,Name:Value,...,Name:Value

  Example:

  EnvironmentVars=VARIABLE1:10000,VARIABLE2:20000

If the environment variable value contains a space character, then enclose the value in quotation marks. If the environment variable value contains a comma character (","), then this character is escaped with "/". Similarly, if the environment variable value contains "/" then it should be escaped also.

Example:

EnvironmentVars=VARIABLE1:"C://PHP",VARIABLE2="1/,2/,3"

- Protocol – This setting specifies the protocol that must be used to communicate with the FastCGI process. The allowed values are **NamedPipe** and **Tcp**. If not specified, the default value is **NamedPipe**.

- QueueLength – This setting specifies the maximum number of requests to an application's process pool that are queued before the FastCGI handler starts returning errors to clients. This indicates that the application is too busy. If not specified, the default value is 1000.

- MaxInstances – This is the highest number of process instances allowed in the process pool. The default is 10. Note that the FastCGI handler will not create this number of processes unless they are needed. If your application never receives more than two concurrent requests, your application only creates two processes.

- InstanceMaxRequests – This is the number of requests sent to a process in the pool before it is shut down and recycled. The default value is 1000.

- IdleTimeout – This is the number of seconds that a process can remain idle without working on a request before it shuts down. The default is 300 seconds.

- ActivityTimeout – This is the number of seconds that the FastCGI handler waits for I/O activity from a process before it is terminated. The default is 30 seconds.

- RequestTimeout – This is the maximum amount of time that a FastCGI process is allowed to handle a request before it is terminated. The default is 90 seconds.

- ResponseBufferLimit – Data from FastCGI processes is buffered before it is returned to the client as a response. This property specifies the amount of response data, in bytes, that is buffered for requests to an application. This buffer is flushed to the client as soon as it is full or when the response is complete, whichever occurs first. If not specified, the default value is 4194304 (4 MB).

- FlushNamedPipe – There are some cases where a FastCGI application might not read all of the data from the named pipe that communicates with the Web server. If this occurs, the Web server waits for a read that does not come. This, in turn, causes a deadlock on that member of the process pool. This most often occurs when a FastCGI process exits abnormally. For instance, the maximum number of requests a process can handle may be set internally, but that number may be lower than the InstanceMaxRequests setting. Setting FlushNamedPipe to 1 causes FastCGI to flush data that could lead to

this condition. The default is 0.

- UnhealthyOnQueueFull – If the value of this variable is 1, the worker process hosting is flagged to IIS as unhealthy any time that the application's request queue is filled. IIS checks health whenever it pings the worker process. If that worker process has been flagged as unhealthy, it and everything it is hosting will be recycled. If not specified, the default value is 0.

**Wildcard Mapping**

The asterisk, "*", association is used for wildcard mapping. This applies to IIS 6.0, which allows a special kind of script map that applies to all requests. If you configure fcgiext.dll as a wildcard handler, this association catches any requests that do not have a more specific configuration associated with the request. For example, the Types section described earlier in this article does not have a ".py" extension configured. If you use the sample section earlier in this article to model a request for a file that is named "File.py," the "*" association is used and the request is routed to the process pool called "Wildcard Mapping."

Consider this section:

[Wildcard Mapping]

ExePath=c:\echo\echo.exe

QueueLength=1000

MaxInstances=20

InstanceMaxRequests=100

IgnoreExistingFiles=1

;IgnoreExistingDirectories=1

Most of this example resembles the other process pool sections, but it has two additional settings, IgnoreExistingFiles and IgnoreDirectories, that are only used for wildcard script maps.

Consider the situation where you have configured an IIS script map that does not verify that a physical file exists before executing the handler. In this scenario, it is convenient to let the FastCGI handler handle requests that are not associated with physical files, but still to let IIS handle other files in the directory. This scenario explains the reason and need for these two settings.

- IgnoreExistingFiles – When set to 1, the FastCGI handler checks if the requested URL is associated with a physical file. If the request is associated with a physical file, the FastCGI handler defers processing the request so that it will be processed externally from the FastCGI handler. The default value is 0.

- IgnoreExistingDirectories – When set to 1, the FastCGI handler checks whether the requested URL is associated with a physical directory. If the request is associated with a physical directory, the FastCGI handler defers processing the request so that it is processed externally from the FastCGI handler. This setting is useful for allowing default documents and directory listings to be served. The default value is 0.

**Site-Specific Wildcard Mappings**

It is also possible to set a site-specific wildcard mapping using the same site ID syntax that file name extensions use (that is, "*:12345678=Section Name"). Note, however, that the search order for file name extensions is significant. When a request reaches the FastCGI handler, the associations in the Types section are checked in the following order:

1. A specific file name extension association specific to the site ID.

2. A '*' association specific to the site ID.

3. A specific file name extension association that has no specified site ID.

4. A '*' association that has no specified site ID.

**Using FastCGI Configuration Script**
To simplify and automate the configuration steps described in previous sections, the configuration script (fcgiconfig.js) is provided with the installation of FastCGI extension. The script is found in %*windir*%\system32\inetsrv. This script adds and removes script maps and modifies application pool settings in the Fcgiext.ini file.

### Add New FastCGI Mapping

To add new FastCGI mapping, run fcgiconfig.js with the –add switch. When using this switch, you must provide the following parameters:

- -section:<Section name>. This parameter specifies the name of the section that will be added to the fcgiext.ini file in [Types] block.
- -extension:<file name extension>. This parameter specifies what file name extensions are associated with this section.
- -path:<file path of the CGI executable>. This specifies the absolute file path of the CGI executable that processes requests for files that have extension specified in –extension parameter.
- -site:<site id>. This optional parameter specifies the site to which the section should be added. If not specified, then the section is added to all sites on your Web server.
- -norecycle. By default, the configuration script recycles all applications pools on IIS 6.0 so that configuration changes can take effect. This optional parameter prevents this recycling.

Examples:

>cscript fcgiconfig.js -add -section:"PHP" -extension:php -path:"C:\PHP\php-cgi.exe"

The example given earlier in this section adds script map for .php extension and updates the Fcgiext.ini file.

As an option, you can specify the site to which the script map is applied:

>cscript fcgiconfig.js -add -section:"PHP" -extension:php -path:"C:\PHP\php-cgi.exe" –site:1

This example adds script map to "Default Web Site" only.

### Remove Existing FastCGI Mapping

To remove existing FastCGI mapping, run fcgiconfig.js with –remove switch. When using this switch, you must provide the following parameters:

- -section:<Section name>. This parameter specifies the name of the section that is removed to the Fcgiext.ini file in [Types] block.
- -norecycle. By default, configuration script recycles all applications pools on IIS 6.0 so that configuration changes take effect. This optional parameter can be used to prevent this recycling.

This example removes the FastCGI mapping for PHP applications.

>cscript fcgiconfig.js -remove -section:"PHP"

### Configuring Existing FastCGI Mapping

You can use the configuration script for all FastCGI configuration properties that were described earlier in this article. To set the configuration properties, use the –set switch. When using this switch, you must provide the following parameters:

- -section:<Section name>. This parameter specifies the section to which the configuration change is applied.
- -<Parameter Name>:<Parameter Value>. This specifies the name of the parameter to change and its new value. The complete list of possible parameters is found in the Application Pool Settings section.
- -norecycle. Use this parameter if you do not want configuration script to recycle application pools after applying

configuration change.

This example sets the FastCGI process pool configuration property InstanceMaxRequests for the PHP section:

```
>cscript fcgiconfig.js -set -section:"PHP" –InstanceMaxRequests:10000
```

# Using FastCGI to Host PHP Applications on IIS 6.0

Updated: February 16, 2009

Applies To: Windows Server 2003

This article describes how to use the Microsoft® Internet Information Services (IIS) 6.0 FastCGI extension to set up and run PHP applications on Windows XP and Windows Server 2003 operating systems.

**What is FastCGI?**
FastCGI is a standard protocol that allows the Common Gateway Interface (CGI) executables in application frameworks to interface with the Web server. It differs from the standard CGI protocol in that FastCGI reuses CGI processes for multiple requests, which provides a significant performance boost compared to CGI. IIS FastCGI support enables IIS to host common CGI programs such as PHP or Ruby on Rails by using the FastCGI protocol and, by doing so, offer high performance and stability for production deployment of such application frameworks. IIS FastCGI support involves the following items:

- The IIS Web server (FastCGI Extension Go Live release supports IIS 6.0)
- The IIS FastCGI extension
- The CGI program (such as Php-cgi.exe)

The Web server dispatches HTTP requests to your application to the FastCGI component, which in turn starts the CGI program executable and forwards the request for processing. Once the request is finished and the response is returned back to the server and sent to the client, the CGI process is reused for a subsequent request. This helps avoid the performance penalty of starting a new process for each request, which gives you much better performance and scalability in a production environment.

To learn more about the FastCGI protocol, see the FastCGI Specification.

**Installing the FastCGI Extension for IIS 6.0**
The FastCGI extension for 32–bit and 64–bit Windows XP and Windows Server 2003 can be downloaded here: http://www.iis.net/downloads/default.aspx?tabid=34&g=6&i=1521.

When you run the FastCGI installer, it copies FastCGI specific files to the "%*windir*%\system32\inetsrv" folder and then registers and enables the FastCGI Web Server Extension. Of the files that the installer copies, the following files are of particular significance:

- fcgiext.dll — This is the actual FastCGI handler that communicates with FastCGI-enabled processes for processing requests.
- fcgiext.ini — This is the configuration file that contains mapping for file name extensions to FastCGI processes. It also contains the configuration of FastCGI process pools.
- fcgiconfig.js — This configures the FastCGI extension. The script updates the Fcgiext.ini file, modifies the IIS metabase, and recycles the Web service application pools if there have been changes.

**Download and Install PHP**
The FastCGI extension is fully compatible with the current official PHP 4.4.x and PHP 5.x distribution for Windows®. It is available from www.php.net/downloads. Since the FastCGI extension hosts and executes PHP processes in a single-threaded manner, you should download and install the non-thread-safe PHP build available for Windows. The non-thread-safe PHP build provides a major performance boost for a single-concurrency environment like FastCGI where a lack of thread safety is acceptable. After downloading the package, extract it to c:\php\. Navigate to c:\php\, and rename php.ini-recommended to php.ini.

You can also use the FastCGI extension with existing PHP 4.4.x or PHP 5.x installations.

**Note:** PHP distributions starting from PHP 5.2.1 contain the performance enhancements developed by Zend that will improve the performance of the PHP engine on Windows. These latest versions provide the optimal PHP performance on the IIS/FastCGI platform.

Zend Core provides a commercial PHP distribution that contains the PHP performance improvements for Windows.

**Configuring the FastCGI Extension to Work with PHP**

Once PHP is installed on Windows, you have two options for configuring FastCGI: you can use the configuration script provided with the installation of the FastCGI extension, or you can perform all the configuration tasks manually by modifying IIS configuration settings and the Fcgiext.ini file. Depending on the option that you select, see the corresponding sections in the remainder of this topic.

**Option 1: Using Script for Configuring the FastCGI Extension to Work with PHP**

The configuration script with the name fcgconfig.js is provided with the installation of the FastCGI extension and is located in %*windir*%\system32\inetsrv.

### To configure the FastCGI extension to work with PHP

1. Open a command-line window, change the current directory to %*windir*%\system32\inetsrv, and run the configuration script to register the PHP CGI program as the one that will be processing .php extensions:

   cscript fcgiconfig.js -add -section:"PHP" -extension:php -path:"C:\PHP\php-cgi.exe"

   (Replace C:\PHP\php-cgi.exe with the path of php-cgi.exe if you installed to a directory other than C:\PHP.)

   **Note:** If you do not want to register the PHP extension to be processed by the FastCGI component on your entire server, and you only want to register it for a specific Web site, then add a "–site:[siteId]" argument.

   Example:

   cscript fcgiconfig.js -add -section:"PHP" -extension:php -path:"C:\PHP\php-cgi.exe" –site:1

   **Note:** For a complete list of configuration parameters supported by the FastCGI extension, see the FastCGI extension documentation.

**Option 2: Manually Configuring the FastCGI Extension to Work with PHP**

To configure the IIS and FastCGI extension manually, you must create script mappings for the PHP extension in the IIS metabase and modify Fcgiext.ini file located in %*windir*%\system32\inetsrv.

### To create script mapping

1. Start Inetmgr.exe.
2. Double-click the icon for the local computer.
3. Right-click **Web Sites** and select **Properties**.
4. Click the **Home Directory** tab.
5. Click **Configuration**.
6. Click **Add**.
7. Browse to the fcgiext.dll located in %*windir*%\system32\inetsrv.

   **Note:** If you must use FastCGI in 32–bit mode on a 64–bit computer, then browse to %*windir*%\SysWOW64\inetsrv.
8. Enter **.php** in the **Extension** field.
9. Select **Limit to** and enter **GET,HEAD,POST**.
10. Select **Script Engine** and **Verify that file exists**.
11. Click **OK**.

### To modify the fcigext.ini file

1. Once the script mapping is added, modify the %*windir*%\System32\inetsrv\fcgiext.ini file.
2. Add an extension to application mapping (php=PHP) to the [Types] section.
3. Add [PHP] section with ExePath=c:\php\php-cgi.exe (assuming you installed PHP files to C:\PHP folder).

[Types]

...

php=PHP

...

[PHP]

ExePath=c:\php\php-cgi.exe

After all the changes to the Fcgiext.ini file are saved, restart the application pool associated with the Web site that hosts your PHP applications.

**Configuring the Default Document in IIS**

After configuring the FastCGI Extension to work with PHP by either of the two methods detailed earlier in this section, configure a default application document in IIS. Most PHP applications use a file that is named Index.php as the default application document. Configure IIS to treat this file as the default content page.

### To configure the default document in IIS

1. Start Inetmgr.exe.
2. Double-click the icon for the local computer.
3. Right-click **Web Sites**, and select **Properties**.
4. Click the **Documents** tab.
5. Click **Add**, and then enter **index.php** in the Default content page field.
6. Click **OK**.
7. If prompted for Inheritance Overrides, select the child nodes, and then click **OK**.

**Setting FastCGI Configuration for Optimal Functionality, Security, and Performance with PHP**

Full PHP functionality support and performance of PHP applications that are running on the FastCGI extension were the focus and main goal for the Microsoft IIS team. During the IIS team's internal testing and collaboration with Zend Technologies, they identified the set of configuration settings for the FastCGI extension and PHP that provides optimal functionality and performance for PHP applications when they are running on Windows and are using the FastCGI extension.

### To configure your server

1. Modify the php.ini file as follows:

   **Note:** Remove the ";", if it is present at the beginning of the following configuration lines.

   - Set **fastcgi.impersonate = 1**. FastCGI under IIS supports the ability to impersonate the security tokens of the calling client. This lets IIS define the security context under which the request runs.

   - Set **cgi.fix_pathinfo=1**. cgi.fix_pathinfo provides *real* PATH_INFO/PATH_TRANSLATED support for CGI. PHP's previous behavior was to set PATH_TRANSLATED to SCRIPT_FILENAME and not care what PATH_INFO is. For more information about PATH_INFO, see the CGI specifications. Setting this to 1 causes PHP CGI to fix its paths to comply with the specification.

   - Set **cgi.force_redirect = 0**.

2. Set the FastCGI configuration settings for the PHP section by running fcgiconfig.js as follows:

   **Note:** These commands should be run from the %*windir*%\system32\inetsrv\ folder.

   - Set the FastCGI process pool property InstanceMaxRequests to 10000. This setting specifies that the FastCGI extension will recycle Php-cgi.exe after it has processed 10000 requests successfully.

```
> cscript fcgiconfig.js -set -section:"PHP" -InstanceMaxRequests:10000
```

3. Configure the FastCGI extension to set PHP_FCGI_MAX_REQUESTS environment variables for the PHP process to 10000. This setting instructs Php-cgi.exe to recycle itself after it has processed 10000 requests successfully.

   ```
   > cscript fcgiconfig.js -set -section:"PHP" -EnvironmentVars:PHP_FCGI_MAX_REQUESTS:10000
   ```

   **Note:** You can configure **InstanceMaxRequests** and **PHP_FCGI_MAX_REQUESTS** to use numbers other than 10000. As a rule, always make sure that the value of **InstanceMaxRequests** is less than or equal to the value of **PHP_FCGI_MAX_REQUESTS**.

**Testing PHP CGI**

After installing the FastCGI extension and registering and configuring PHP CGI, you can request and use your PHP application.

To test, create and request a phpinfo.php page in your C:\inetpub\wwwroot folder that contains the following:

```
<?php

phpinfo();

?>
```

Note that the Server API indicates that PHP is hosted in FastCGI mode.

**Installing phpBB Application for Ase with the FastCGI Extension**

The phpBB Internet forum solution is a popular open-source application that provides a large and very customizable set of features. These include powerful permissions systems, private messaging, search functions, a customizable template and language system, and support for multiple databases. For more information about phpBB, see http://www.phpbb.com. For phpBB community–specific information, see http://www.phpbb.com/community/.

This section describes the steps that are required to install phpBB to work with the FastCGI extension on IIS 6.0.

The following steps assume that you have completed the setup and configuration of the FastCGI extension and PHP libraries as described in the previous sections.

**Download and Unpack the Application**

First, download the latest stable release of phpBB. As of January, 2009, the current release of phpBB (phpBB 3.0.2) is available at http://www.phpbb.com/downloads/?sid=ee349e0ae5950a30df04461023f7e34f. Once you download the package, uncompress it and copy all the files and folders to C:\Inetpub\wwwroot\phpBB2.

**Set Up the Database**

In order to use phpBB, you must have access to a database. For this walkthrough, we use Microsoft® SQL Server® 2005 Express Edition with mixed authentication. You can also use Microsoft® SQL Server® 2005, Microsoft® Office Access®, MySQL, or other databases.

Before starting the installation, create a database on the database server. Also create a database user and grant this user database ownership permission to the database.

Open **Microsoft SQL Server Management Studio**, and then click the **New Query** button. Enter the following script into the query window, and then click **Execute**. This script creates the database and user necessary for the phpBB application and should return "Command(s) completed successfully." in the messages pane.

```
USE [master]

GO

CREATE DATABASE phpbb3

GO

CREATE LOGIN [phpbb3]

WITH PASSWORD=N'phpbb3',

DEFAULT_DATABASE=[phpbb3],

CHECK_EXPIRATION=OFF,
```

```
CHECK_POLICY=OFF

GO

USE [phpbb3]

GO

CREATE USER [phpbb3] FOR LOGIN [phpbb3] WITH DEFAULT_SCHEMA=[dbo]

USE [phpbb3]

GO

sp_addrolemember 'db_owner','phpbb3'
```

**Set Up ODBC Connection to the phpBB3 Database**
To form the connection between the phpBB3 application and the database, create an ODBC connection. In the example used in this article, we use SQL Server 2005 Express Edition with an ODBC connection.

**To create an ODBC connection**

1. Click **Start**, **Administrative Tools**, **Data Source (ODBC)**.
2. Click the **System DSN** tab.
3. Click **Add**.
4. Select **SQL Server**, and then click **Finish**.
    a. Name: **phpbb3**
    b. Description:
    c. Server: **<HOSTNAME>\SQLEXPRESS**
    d. Click **Next**.
5. Select **With SQL Server authentication using a login ID and password entered by the user**.
    - Check **Connect to SQL Server to obtain default settings for the additional configuration options**.
    - Login ID: **sa**
    - Password: **<sa password>**
    - Click **Next**.
6. Select **Change the default database to:**, and select **phpbb3** from the drop-down menu.
7. Click **Next**.
8. Click **Finish**.
9. Click **Test Data Source**. The results should be "TESTS COMPLETED SUCCESSFULLY!"
10. Click **OK**.
11. Click **OK**.
12. Click **OK**.

**Set Up and Configure the Application**
Open your Web browser and browse to http://localhost/phpBB3/index.php. This redirects you to the phpBB setup page. Enter the required installation information and replace the Database Server Hostname with the name of your database server.

**To set up and configure the application**

1. Click **Install** tab.
2. Click **Proceed to next step**.
3. Requirements:
    a. **Files and Directories** may appear as "Unwritable."
    b. Navigate to c:\inetpub\wwwroot\phpbb3\ folder.
    c. Change NTFS security for the following folders so that **Internet Guest Account** has **write** access to the following folders:
        i. Cache/
        ii. Files/
        iii. Store/
    d. Click **Test again**.
    e. All compatibility requirements should pass.
    f. Click **Start Install**.
4. Database settings: (SQL Server 2005 Express, mixed authentication)
    a. Database type: **MSSQL Server 2000+ via ODBC**
    b. Database server hostname or DSN: **phpbb3** (created earlier in ODBC configuration section)
    c. Database name: **phpbb3**
    d. Database username: **phpbb3**
    e. Database password: **phpbb3**
    f. Click **Proceed to next step**.
    g. "Successful connection" message should be displayed. Click **Proceed to next step**.
5. Administrator details
    a. Administrator username: **admin**
    b. Administrator password: <create admin password>
    c. Confirm administrator password: <create admin password>
    d. Contact e-mail address: <enter contact e-mail>
    e. Confirm contact e-mail address: <enter contact e-mail>
    f. Click **Proceed to next step**.
    g. Results should be "Tests passed." Click **Proceed to next step**.

6. Configuration file
    a. To download configuration file, click **Download**.
    b. Save config.php to c:\inetpub\wwwroot\phpbb3\ folder. Overwrite, if you are prompted,.
    c. Click **Done**.
    d. Click **Proceed to next step**.
7. Advanced settings
    a. Leave default settings for this installation.
    b. Click **Proceed to next step**.
8. Click **Proceed to next step**.
9. Click **Login**.
10. You will be redirected to the Welcome page. Click **Logout** in the left column.

**Note:** If your Web site is configured to allow anonymous users, then you must grant access permissions to "Internet Guest Account."

#### Test the phpBB Application

### To test that phpBB has been installed successfully and create a forum

1. Navigate to http://localhost/phpBB3, and then click **Login**.
2. Enter the administrator username and password (this walkthrough used username: admin).
3. Once logged in, click on the **Administration Control Panel** link at the bottom of the page to view the administration panel.
4. This causes phpBB to ask you to re-authenticate. It then takes you to the administration panel.
5. On this panel, click **Manage Forums** under the **Quick Access** group.
6. Enter the desired forum name, and then click **Create new forum**.
7. Leave the default settings on the Create new forum page, and then click **Submit**.
8. You will then be redirected to the Forum permissions page. You can assign permissions later. Click the **Manage forums** link in the left column.
9. You can now see your new forum listed on the Forum administration page.

#### Getting More Information

To discuss the FastCGI extension for IIS 6.0 or file bug reports, use the IIS.NET FastCGI forum.

To see a list of known issues with the IIS 6.0 FastCGI, see this forum: http://forums.iis.net/1103/ShowForum.aspx.

To get more information about running various PHP applications on IIS, see:

PHP support in IIS 6.0 and IIS 7.0

PHP Community forum

#### Related Content

### Articles

- Setting up FastCGI for PHP

- Installing PHP on Windows Vista with FastCGI
- PHP Questions
- Installing FastCGI Support On Server Core

# Configuring Servers for Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In Web development, Web site and application code is complex and ever changing. Web sites and applications change as frequently as every week or month. Consequently, you require hosting services that actively manage your run-time environment, while safeguarding applications and managing performance.

This section explains how to use IIS Manager to configure your server to manage and execute Web-based applications.

This section includes the following information:

- About Configuring Servers for Applications: Discusses things to consider when you configure your server to run Web-based applications.
- Configuring IIS For Dynamic Content: Describes the necessary tasks to enable and disable the Web service extensions that process dynamic information.
- Configuring Isolation Modes: Describes configuring a server's modes of operation, and explains the ramifications of each mode.
- Worker Process Isolation Mode: Describes the WWW Service Administration and Monitoring features of Internet Information Services (IIS) 6.0 running in worker process isolation mode, and provides procedures for using them.
- IIS 5.0 Isolation Mode: Explains how to use the features of IIS 6.0 running in IIS 5.0 isolation mode.

**Related Information**

- For information about creating and configuring applications, see Configuring Applications in IIS 6.0.

# About Configuring Servers for Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0 delivers Web hosting services through an adjustable architecture that you can use to manage server resources with improved stability, efficiency, and performance. IIS separates applications into isolated pools and automatically detects memory leaks, defective processes, and over-utilized resources. When problems occur, IIS manages them by shutting down and redeploying faulty resources and connecting faulty processes to analytical tools.

IIS can run in either of two mutually exclusive modes of operation:

- Worker process isolation mode. This is the default mode of IIS 6.0, isolates key components of the World Wide Web Publishing Service (WWW service) from the effects of errant applications, and it protects applications from each other by using the worker process component. Use worker process isolation mode unless you have a specific compatibility issue that makes the use of IIS 5.0 isolation mode necessary. Web sites that serve static content or simple ASP applications should be able to move to IIS 6.0 running in worker process isolation mode with little or no modification.

- IIS 5.0 isolation mode. With this mode, you can run applications that are incompatible with worker process isolation mode because they were developed for earlier versions of IIS. Applications that run correctly on IIS 5.0 should run correctly on IIS 6.0 in IIS 5.0 isolation mode.

Worker process isolation mode provides better default security for running Web applications than IIS 5.0 isolation mode. By default, worker processes run with the Network Service identity. The Network Service account has lower access rights than the default account for IIS 5.0 isolation mode. Web applications that run in-process in IIS 5.0 application mode run as LocalSystem. The LocalSystem account can read, execute, and change most of the resources on the computer.

The default isolation mode upon installing IIS 6.0 depends on whether you perform a clean installation or an upgrade.

- After a clean install of IIS 6.0, IIS runs in worker process isolation mode.

- After an upgrade from an earlier version of IIS 6.0, the isolation mode is the same as configured on the previously-installed version of IIS 6.0.

- After an upgrade from IIS 5.0 or IIS 4.0, IIS 6.0 runs in IIS 5.0 isolation mode by default to maintain compatibility with your existing applications.

For information about switching from one isolation mode to the other, see Configuring Isolation Modes.

**Worker Process Isolation Mode**

IIS 6.0 introduces worker process isolation mode, which runs all Web applications in an isolated environment. When you run IIS in worker process isolation mode, applications can be configured to run in separate application pools. Each application pool is a logical representation of a configurable worker process and links to the applications in the pool. Worker processes operate independently of each other; they can fail without affecting other worker processes. The pooling of applications protects applications from the effects of worker processes that support other application pools. In this way, applications are protected from each other.

In worker process isolation mode, Hypertext Transfer Protocol (HTTP) requests are routed directly to an in-kernel application pool queue serving the configured application. Worker processes that serve an application pool pull the requests directly from the queue, avoiding process-switching overhead.

To further protect your WWW service, IIS 6.0 isolates critical World Wide Web Publishing Service (WWW service) components, such as the HTTP protocol stack (HTTP.sys) and WWW Service Administration and Monitoring, from the effects of third-party code running in worker processes. HTTP.sys receives and queues requests for WWW services. When a worker process enters an unhealthy state, and thus stops processing requests, HTTP.sys continues to process requests. Meanwhile, the WWW service detects that the worker process is unhealthy and shuts it down. If there is demand for a new worker process to serve requests (HTTP.sys has requests queued), the WWW service starts a new worker process to pick up the queued requests from HTTP.sys. Even though a worker process has failed, the WWW service continues to process requests and shields the user from experiencing a loss of service.

IIS 6.0 worker process isolation mode delivers the following specific improvements over earlier versions of IIS:

- **Robust performance** Isolation prevents Web applications and Web sites from affecting each other or the WWW service. Reboots of the operating system and restarting of the WWW service are avoided.

- **Self-healing** Automated management provides auto-restart of failed worker processes and periodic restart of deteriorating worker processes.

- **Scalability** Web gardens allow more than one worker process to serve the same application pool.

- **Process affinity** enables the connection of worker processes to specific processors on multi-CPU servers.

- **Automated debugging** The debugging feature enables the automatic assignment of failing worker processes to debugging tools.

- **CPU limiting** This monitoring feature enables controlling the amount of CPU resources that an application pool consumes in a configured amount of time.

**IIS 5.0 Isolation Mode**

If your server must run applications that are incompatible with worker process isolation mode, you can configure IIS to run in IIS 5.0 isolation mode, which supports the process model and features of IIS 5.0. If you run IIS in IIS 5.0 isolation mode, some features of worker process isolation mode are not available.

**Related Information**

- For information about creating application pools, see Configuring Application Pools in IIS 6.0.

- For more information about when to use IIS 5.0 isolation mode and how to configure IIS for IIS 5.0 isolation mode, see Configuring Isolation Modes.

# Configuring Isolation Modes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can run IIS 6.0 in one of two modes: *worker process isolation mode* or *IIS 5.0 isolation mode*. Worker process isolation mode is the default mode, that is, the mode that the server runs in on a clean installation. You can use IIS 5.0 isolation mode to run applications developed for older versions of IIS that are determined to be incompatible with worker process isolation mode.

This topic includes the following information:

- Configuring IIS for worker process isolation mode.
- Configuring IIS for IIS 5.0 isolation mode.

**Configuring IIS for Worker Process Isolation Mode**

Worker process isolation mode is the default service mode in IIS. Worker process isolation mode delivers all the benefits of IIS 6.0 new architecture: robust application pooling; automated restarts, scalability, debugging; and finely-tuned performance tuning. Web applications run with the Network Service identity, which provides a security advantage: the Network Service account has lower access privileges than LocalSystem. Use worker process isolation mode unless you must run applications that might conflict with it, as stated below in "Configuring IIS for IIS 5.0 Isolation Mode."

To complete the following procedure, you must restart IIS, which will temporarily interrupt the World Wide Web Publishing Service (WWW service).

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To configure IIS for worker process isolation mode**

1. In IIS Manager, expand the local computer, right-click **Web Sites**, and then click **Properties**.
2. Click the **Service** tab, clear the **Run WWW service in IIS 5.0 isolation mode** check box, and then click **OK**.
3. To start the WWW service, click **Yes**.

If the switch to worker process isolation mode is successful, a folder named **Application Pools** appears in the **IISManager** listing for your local computer. You can always determine which isolation mode IIS is running by the presence (worker process isolation mode) or absence (IIS 5.0 isolation mode) of the **Application Pools** folder.

**Configuring IIS for IIS 5.0 Isolation Mode**

You must use IIS 5.0 isolation mode for applications that conflict with worker process isolation mode until the applications are modified.

The following application characteristics conflict with worker process isolation mode:

- Dependency on Inetinfo.exe: If the application must run in the Inetinfo.exe process, it must be run in IIS 5.0 isolation mode because applications do not run in Inetinfo.exe in worker process isolation mode.
- Requires Read Raw Data Filters: Read Raw Data Filters are available in IIS 5.0 isolation mode only.
- Requires Dllhost.exe: Applications that must be run in a Dllhost.exe environment can be run only in IIS 5.0 isolation mode because Dllhost.exe is not available in worker process isolation mode.

If the IIS 6.0 service is running in worker process isolation mode (the IIS 6.0 default mode), and you must run applications that

do not meet the requirements for worker process isolation mode, switch to IIS 5.0 isolation mode. This means you will not be able to take advantage of worker process isolation and the other features of worker process isolation mode.

After you complete the following procedure, you must restart the WWW service, which will temporarily interrupt the service.

**To configure IIS for IIS 5.0 isolation mode**

1. In IIS Manager, expand the local computer, right-click **Web Sites**, and then click **Properties**.
2. Click the **Service** tab, select the **Run WWW service in IIS 5.0 isolation mode** check box, and then click **OK**.
3. To start the WWW service, click **Yes**.

**Configuring IIS for an application mode using Adsutil.vbs**

You can also configure IIS for an application isolation mode by setting the metabase property **IIs5IsolationModeEnabled** to **true** or **false**. Use the following procedure to configure this metabase property from a command line.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**To configure IIS for an application isolation mode by using Adsutil.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
2. To set the application isolation mode, type one of the following commands at the command prompt:
   - To configure the server to run in worker process isolation mode, type:

     **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/IIs5IsolationModeEnabled FALSE**

   - To configure the server to run in IIS 5.0 isolation mode, type:

     **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/IIs5IsolationModeEnabled TRUE**

   | **Note** |
   | --- |
   | If you configured the server to run in IIS 5.0 isolation mode, you will need to set the level of application protection, or isolation, for each application running on the server. For instructions, see Isolating Applications in IIS 6.0. |

3. To restart IIS, at the command prompt, type:

   **iisreset /noforce**

   For more information about restarting IIS from the command line, see Restarting IIS.

To verify that you successfully changed to worker process isolation mode, check whether an **Application Pools** folder appears when you open IIS Manager on your local computer. You can always determine which isolation mode that IIS is running by the presence (in worker process isolation mode) or absence (in IIS 5.0 isolation mode) of the Application Pools folder.

| **Important** |
| --- |
| Do not confuse the **DefaultAppPool** application pool, which appears by default when you configure a computer to run in worker process isolation mode, with the **MSSharePointAppPool** application pool, which appears by default when you install FrontPage2002 Server Extensions or SharePoint™ Team Services from Microsoft. |

| **Note** |
| --- |

Do not use, delete, or make changes to the **MSSharePointAppPool** application pool. It holds files that are used internally by FrontPage Server Extensions and SharePoint Team Services. For more information, see Using FrontPage 2002 Server Extensions with IIS 6.0.

# Worker Process Isolation Mode in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section explains how use IIS Manager to configure Internet Information Services (IIS) to execute Web-based applications with IIS in worker process isolation mode.

This section includes the following information:

- Configuring Application Pools in IIS 6.0: Explains the advantages of pooling application processes, and describes how to create and configure application pools.
- Recycling Worker Processes with IIS 6.0: Describes how to configure IIS to periodically restart worker processes in an application pool.
- Configuring Worker Processes for Idle Timeout: Describes how to configure idle process timeout, which is a resource management feature.
- Processor Affinity: Explains the concept of process/processor affinity.
- Configuring Web Gardens with IIS 6.0: Describes how to configure multiple application pools running under one worker process.
- Application Pool Health: Explains the concepts of application pool health and health monitoring.
- Configuring Health Monitoring: Describes how to enable IIS to monitor the health of each worker process.
- Configuring Rapid-Fail Protection in IIS 6.0: Describes how to protect resources when the requests to a worker process fail rapidly.
- Configuring Startup and Shutdown Time Limits: Explains how to configure a worker process to startup on demand when the first request for a Uniform Resource Locator (URL) is received.
- Configuring Application Pool Identity with IIS 6.0: Describes the security considerations that are related to configuring the identity of an application pool.

**Related Information**

- For an overview of IIS worker process isolation mode, see About Configuring Servers for Applications.
- For information about configuring applications, see Configuring Applications in IIS 6.0.

# Configuring Application Pools in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when running in worker process isolation mode. |

An application pool is a configuration that links one or more applications to a set of one or more worker processes. Because applications in an application pool are separated from other applications by worker process boundaries, an application in one application pool is not affected by problems caused by applications in other application pools.

By creating new application pools and assigning Web sites and applications to them, you can make your server more efficient and reliable, as well as making your other applications always available, even when the worker process serving the new application pool has problems.

**Guidelines for Creating Application Pools**

Consider the following guidelines when configuring application pools:

- To isolate Web applications on a Web site from Web applications on other sites running on the same computer, create an individual application pool for each Web site.
- For enhanced security, configure a unique user account (process identity) for each application pool. Use an account with the least user rights possible, such as Network Service in the IIS_WPG group.
- If there is a test version of an application on the same server with the production version of the application, separate the two versions into different application pools. This isolates the test version of the application.
- As a design consideration, if you want to configure an application to run with its own unique set of properties, create a unique application pool for that application.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To create a new application pool**

1. In IIS Manager, expand the local computer, right-click **Application Pools**, point to **New**, and then click **Application Pool**.
2. In the **Application pool name** box, type the name of the new application pool.
3. If the ID that appears in **Application pool ID** box is not the ID that you want, type a new ID.
4. Under **Application pool settings**, click the appropriate setting. If you click **Use existing application pool as template**, in **Application pool name** box, right-click the application pool that you want to use as a template.
5. Click **OK**.

Application pools allow you to apply configuration settings to groups of applications and the worker processes that service those applications. Any Web site, Web directory, or virtual directory can be assigned to an application pool.

**To assign an application to an application pool**

1. In IIS Manager, right-click the application that you want to assign to an application pool, and then click **Properties**.
2. Click the **Virtual Directory**, **Directory**, or **Home Directory tab**.

3. If you are assigning a directory or virtual directory, verify that Application name is filled in. If the **Application name** box is not filled in, click **Create**, and then type a name.

4. In the **Application pool** list box, click the name of the application pool to which you want to assign the Web site.

# Configuring Application Pool Identity with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. |

The identity of an application pool is the name of the account under which the application pool's worker process runs. By default, application pools operate under the Network Service account, which has low-level user access rights. That is, this account provides better security against attackers or malicious users who might attempt to take over the computer on which the World Wide Web Publishing Service (WWW service) is running. The LocalService account also has low-level user access rights, which is helpful for situations that do not require access to resources on remote computers. You can configure application pools to run as LocalSystem, which is an account with more user rights than the Network Service or LocalService account. However, be mindful that running an application pool under an account with increased user rights presents a high security risk.

For example, suppose that an Internet Service Provider (ISP) wants to allow customers to upload Common Gateway Interface (CGI) applications and then add them to an application pool. Running CGI-enabled applications in a separate application pool under the Network Service account, with its lower user rights, reduces the risk that these applications will be used to attack the server.

For more information about application pool identities, see Configuring Application Pool Identity in IIS 6.0.

For more information about issues that can arise when you use a configurable account as the application pool identity, see Knowledge Base article 871179, "You receive an 'HTTP Error 401.1 - Unauthorized: Access is denied due to invalid credentials' error message when you try to access a Web site that is part of an IIS 6.0 application pool."

For more information about using a configurable account with Kerberos authentication, see the following resources:

- Knowledge Base article 326985, "HOW TO: Troubleshoot Kerberos-Related Issues in IIS"
- Troubleshooting Kerberos Delegation
- Kerberos authentication for load-balanced Web sites

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To change the account under which an application pool runs using IIS Manager**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. Click the **Identity** tab, and click either **Predefined** or **Configurable**. **Predefined** refers to standard service names, such as Network Service (the default), LocalSystem, or Local Service. **Configurable** refers to registered user names.
   - If you click **Predefined**, click a predefined account in the list box.
   - If you click **Configurable**, in the **User name** and **Password** boxes, type the user name and password of the account under which you want the worker process to operate. Then, add the account to the IIS_WPG group.
3. Click **OK**.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**To change the account under which an application pool runs using the Adsutil.vbs administration script**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. Switch to the %systemroot%\system32\inetsrv directory.
4. At the command prompt type **cscript adsutil.vbs SET W3SVC/AppPools/***AppPoolName***/WAMUserName** *Username*
5. At the command prompt type **cscript adsutil.vbs SET W3SVC/AppPools/***AppPoolName***/WAMUserPass** *Password*

**To programmatically change the account under which an application pool runs**

- In order to change the identity of a worker process in IIS 6.0, set three properties in the metabase under the application pool key that you want to change. Using "TestPool" as an example, and assuming the application pool named "TestPool" already exists, the following JScript snippet would change the application pool identity.

```
var appPool = GetObject("IIS://localhost/w3svc/AppPools/TestPool" );
// Tell IIS to use configurable identity for worker process
appPool.AppPoolIdentityType = 3;
// Set user name
appPool.WAMUserName = "myuser";
// Set password
appPool.WAMUserPass = "myuserpass";
// Save changes to the IIS metabase
appPool.SetInfo();
```

# Configuring Worker Processes for Idle Timeout

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. |

When it is necessary to conserve system resources by terminating unused worker processes, you can configure a worker process to gracefully close after a specified period of time. You can use this feature to better manage the resources when the processing load is heavy, when identified applications consistently fall into an idle state, or when new processing space is not available. You can also start additional worker processes to replace a worker process that is finished.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure application pool idle timeout**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. On the **Performance** tab, under **Idle timeout**, select the **Shutdown worker processes after being idle for** check box.
3. In the **minutes** box, type the number of idle minutes (with no processing requests) that you want to elapse before shutting down the idle worker process. The default value is 20 minutes.
4. Click **OK**.

Optionally, you can configure idle timeout by setting the **IdleTimeout** metabase property.

**To configure application pool idle timeout by using Adsutil.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
2. At the command prompt, type:

   ```
   cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set
   W3SVC/AppPools/ApplicationPoolName/IdleTimeout n
   ```

   Replace *n* with the number of minutes of idle time to allow before shutting down the idle worker process; the default time-out interval is 20 minutes.

For more information about setting application pool idle timeout, see the IdleTimeout Metabase Property.

**Related Information**

- For information about application pool settings affecting starting and stopping of a worker process, see Configuring Startup and Shutdown Time Limits.

# Processor Affinity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. |

On multi-CPU servers, application pools can be configured to establish affinity between worker processes and multiple CPUs to more efficiently use CPU caches. Processor affinity is used in conjunction with the processor affinity mask setting to specify CPUs.

You can set processor affinity from the command line by using the **set** command with Adsutil.vbs. To set processor affinity, configure the **SMPAffinitized** and **SMPProcessorAffinityMask** metabase properties. Assigning processor affinity involves these steps:

- Enable processor affinity. Set the **SMPAffinitized** metabase property to **true**, which indicates that a particular worker process assigned to an application pool can also be assigned to a specified CPU.

- Assign processor affinity. Set the **SMPProcessorAffinityMask** metabase property, which configures a hexadecimal processor mask that indicates which CPU the worker processes in an application pool are bound to.

**Procedures**

**To enable processor affinity for an application pool**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.

2. At the command prompt, type:

   **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/SMPAffinitized TRUE**

**To assign processor affinity by binding an application pool to a specified processor**

- At a command prompt, type:

  **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/SMPProcessorAffinityMask** *MaskValue*

  Replace *MaskValue* with the hexadecimal value that binds the application pool to the appropriate processor or processors. See Table 3.8 for sample hexadecimal mask values.

For example, if you want to bind the worker processes that serve your FinanceAppPool to processors 0, 1, 2, and 3 of an eight-processor computer, type the following command:

**cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/FinanceAppPool/SMPProcessorAffinityMask 0xF**

**Related Information**

- For more information about processor affinity, see Assigning Processor Affinity.

- For information about processor affinity mask settings, see SMPProcessorAffinityMask Metabase Property and SMPAffinitized Metabase Property.

# Configuring Web Gardens with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. |

In IIS 6.0 worker process isolation mode, application pools enhance Web site or application reliability by isolating applications and the worker processes that service those applications. For even greater reliability, you can configure an application pool to be supported by multiple worker processes. An application pool that uses more than one worker process is called a Web garden. The worker processes in a Web garden share the requests that arrive for that particular application pool. If a worker process fails, another worker process can continue to process requests.

| **Note** |
| --- |
| Web gardens are different from Web farms. A Web garden is configured on a single server by specifying multiple worker processes for an application pool. Web farms use multiple servers for a Web site. |

Creating a Web garden for an application pool can also enhance performance in the following situations:

- Robust processing of requests: When a worker process in an application pool is tied up (for example, when a script engine stops responding), other worker processes can accept and process requests for the application pool.
- Reduced contention for resources: When a Web garden reaches a steady state, each new TCP/IP connection is assigned, according to a round-robin scheme, to a worker process in the Web garden. This helps smooth out workloads and reduce contention for resources that are bound to a worker process.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

### To configure a Web garden by using IIS Manager

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. Click the **Performance** tab, and under **Web garden**, in the **Maximum number of worker processes** box, type the number of worker processes that you want to assign to the application pool. (You must type a number greater than 1 for the application pool to become a Web garden.
3. Click **OK**.

Optionally, you can configure a Web garden by setting the metabase property **MaxProcesses**. The **MaxProcesses** property determines the maximum number of worker processes that an application pool allows to service its requests. A value of zero indicates an unmanaged application pool that is not served by a worker process.

The default value for the **MaxProcesses** property is 1, which is the default number of worker processes that service an application pool. To configure an application pool so that it is a Web garden, set the **MaxProcesses** property to a value greater than 1.

| **Important** |
| --- |

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters).

**To configure a Web garden by using Adsutil.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
2. At the command prompt, type:

   **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/MaxProcesses** *n*

   Replace *n* with the number of worker processes that you want to service the application pool.

For more information about configuring Web gardens, see MaxProcesses Metabase Property.

**Related Information**

- For information about resource allocation, see Processor Affinity.

# Application Pool Health

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when running in worker process isolation mode. |

In worker process isolation mode, application pools can be configured to monitor the health of their worker processes as well as the health of the entire pool. Monitoring the health of a worker process includes detecting that the worker process is not able to serve requests and taking appropriate action. For example, if a worker process fails to respond to a ping request by the World Wide Web Publishing Service (WWW service,) the worker process probably does not have threads available for processing incoming requests. When this happens, the WWW service will either terminate the worker process, or release the worker process and leave it running, and start a new worker process to replace it. An administrator can pre-configure an action to take when an unhealthy worker process is released, such as attaching the worker process to a debugger.

In addition to monitoring the health of worker processes, the WWW service can also detect ongoing problems with the entire application pool. For example, if worker processes are terminating abnormally every few seconds, the WWW service can determine that the worker process is unhealthy and stop it, preventing the unhealthy applications from affecting applications in other pools.

The following conditions cause an application pool to start or stop:

- Initiation of rapid-fail protection causes an application pool to stop.
- A job object hitting it's time limit causes an application pool to stop, followed by a start when the time window expires.
- A configuration error caused by trying to use a nonexistent identity causes an application pool to stop.
- A Windows administrator performs a demand stop or start on an application pool.

**Monitoring Application Pool Health**

Worker process pinging enables the WWW service to detect that a worker process is unable to respond to requests (In other words, the worker process is unhealthy). The ping is a message sent between the WWW service and the worker process. If the ping succeeds, then the WWW service assumes the application pool is healthy. If the ping fails (no response from the worker is received), then the WWW service assumes that a problem exists with the worker process. If a problem exists, the WWW service will either terminate the worker process or release it, and then will start a new worker process when it is needed to serve new requests. If the worker process is released but still running, then the WWW service will run the action configured by the administrator.

The setting of the WWW service pinging feature and the frequency of the ping are controlled by the PingingEnabled Metabase Property and the PingInterval Metabase Property.

**ISAPI Extensions Can Declare Themselves Unhealthy**

An ISAPI extension application can be built to programmatically signal IIS that it needs to be recycled. This can be accomplished through the new Server Support function: HSE_REQ_REPORT_UNHEALTHY. See ServerSupportFunction Extension Function at MSDN Online.

To use this function effectively, the IIS server that is running your application must have worker process pinging enabled (see Configuring Worker Process Pinging), because it is during the ping operation that the WWW service detects that the ISAPI is unhealthy and the worker process should be recycled. The ISAPI will need an internal mechanism for determining its unhealthy state, such as monitoring the status of its internal thread pool. You should consider that this type of programming shuts down the worker process the ISAPI extension is running in. Therefore, all applications running in that worker process will be restarted.

The ASP ISAPI extension implements logic that takes advantage of this feature as it monitors the status of its internal thread pool. If too many of its threads enter a blocking state, it will signal a recycle.

**Limitations of Health Detection**

Health detection cannot be used to determine application failures that do not cause the worker process to crash nor block the

available threads in the worker process. As an example, an application that is returning invalid response codes such as HTTP 500 errors, but otherwise functioning normally, will still respond to a ping from the WWW service, unless the application was a custom ISAPI extension that implemented specific code to indicate its unhealthy state.

**Enabling Debugging Action**

When a worker process is determined to be unhealthy, instead of terminating it, you might want the WWW service to keep it running for the purpose of debugging it, while bringing up a new worker process to serve requests. By enabling debugging on an application pool, you signal the WWW service not to terminate the worker process, but to release it from serving an application pool, and leave it running.

In addition to running an unhealthy worker process in a released state, you can configure the WWW service to launch an executable application or script (for example, an application that sends e-mail to administrators notifying them that the failure could be configured as an enable-debugging feature). The *process id* of the process deemed unhealthy is the first argument sent to the executable application or script.

| **Note** |
| --- |
| A worker process that has been released and left running may still terminate. The WWW service will have left it running. If the worker process recovers from its unhealthy state, it will detect that it has no relationship to the WWW service, and self-terminate. It would then be possible to find a log entry stating that a worker process was released, but find no evidence of the worker process running. |

If you enable running unhealthy worker processes to be released, consider that you will need to deal with blocked worker processes, as they will not be removed from memory by IIS. There could be large numbers of failed worker processes running on your computer if administrators are not properly handling worker processes that are kept alive for debugging purposes. Also, consider that these worker processes may be tying up resources needed by other processes. You may need to terminate them quickly in order to free up those resources.

**Rapid-Fail Protection**

Rapid-fail protection stops application pools when too many worker processes assigned to it are found to be unhealthy in a specified period of time.

When an application pool is stopped, HTTP.sys will either return an out-of-service message (503: Service Unavailable) or connection resets based on the configuration of the **LoadBalancerCapabilities** property of the application pool. Also, when an application pool is stopped automatically, you can configure an action (a debugging action, for example) to notify the administrator that the application pool has stopped.

| **Note** |
| --- |
| If you host multiple applications on a single computer, you should be careful to configure load balancers or switching hardware to reroute only the traffic intended for a failed application pool. Do not route requests away from healthy application pools; they are still able to receive and process requests. |

Rapid-fail protection reduces processing overhead for problematic applications as the requests do not enter user-mode processing. Thus, other application pools are protected from the unhealthy application pool.

You can set rapid-fail protection two ways:

1. Configure IIS to place the application in a rapid-fail state based on the number of the worker process failures in a given time period, measured in minutes.

2. Place an application in a rapid-fail state manually.

**Scenario for Monitoring Application Pool Health**

Monitoring application pool health and taking corrective actions requires, at a minimum, that you take the following steps:

1. Allow the WWW service to detect unhealthy applications by enabling the WWW service to ping worker processes. See Configuring Worker Process Pinging.

2. Configure rapid-fail protection to allow the disabling of application pools when the worker processes assigned to them crash a set number of times in a specified period. See Configuring Rapid-Fail Protection in IIS 6.0.

**Related Information**

- For more information about configuring IIS for application pool health monitoring, see Configuring Worker Process Pinging.

- For information about setting rapid-fail protection, see Configuring Rapid-Fail Protection in IIS 6.0.

# Configuring Health Monitoring

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. |

The World Wide Web Publishing Service (WWW service) monitors the health of worker processes by periodically pinging them to determine their responsiveness. If a worker process fails to respond to a ping, the WWW service terminates the worker process and creates another worker process as a replacement. Additionally, the WWW service maintains a communication channel to each worker process, and it can detect a drop in the communication channel, which indicates a worker process failure. When a worker process fails on its own accord, the WWW service starts another worker process in its place.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To enable application pool health monitoring by using IIS Manager**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. Click the **Health** tab, and select the **Enable pinging** check box.
3. In the box next to **Ping worker process every (frequency in seconds):**, type the number of seconds that you want to elapse between pings.
4. Click **OK**.

You can also configure the WWW service pinging feature and the frequency of the ping by setting the metabase properties **PingingEnabled** and **PingInterval**. Use the following procedures to configure these metabase properties from the command line.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**To enable application pool health monitoring by using Adsutil.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
2. At the command prompt, type:
   - **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/PingingEnabled TRUE**

**To change the ping interval by using Adsutil.vbs**

- At the command prompt, type:
  - **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName*

**/PingInterval** *n*

In the preceding syntax, replace *n* with the number of seconds that you want to elapse between pings.

For more information about setting these metabase properties see PingingEnabled Metabase Property, PingInterval Metabase Property, and Using the Adsutil.vbs Administration Script.

# Configuring Rapid-Fail Protection in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. |

When an application pool detects that too many worker processes assigned to it have become unhealthy in a given period of time, rapid-fail protection is initiated. The communication channel with the World Wide Web Publishing Service (WWW service) is dropped. The WWW service detects the loss of communication and initiates the appropriate actions, typically sending an error or warning to the event log and then restarting the worker process. Worker process isolation mode can be configured so that if a particular application pool experiences consecutive multiple failures, the application pool can be automatically disabled, allowing for rapid-fail protection to start. The application pool is taken out of service and placed in a state such that the kernel-mode driver immediately returns an out-of-service message (503: Service Unavailable) to any requests to that application pool.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To configure rapid-fail protection by using IIS Manager**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. Click the **Health** tab, and select the **Enable rapid-fail protection** check box.
3. In the **Failures** box, type the number of worker process failures to be detected before disabling the worker process.
4. In the **Time period** box, type the number of minutes during which failure totals are accumulated.
5. Click **OK**.

**Configuring Rapid-Fail Protection from the Command Line**

Use the following procedures to configure rapid-fail protection from the command line, and to change the default settings for rapid-fail protection in the **RapidFailProtectionInterval** metabase property and the **RapidFailProtectionMaxCrashes** metabase property.

**To configure rapid-fail protection for an application pool by using Adsutil.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
2. At the command prompt, type:

   **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/RapidFailProtection TRUE**

By default, rapid-fail protection is activated if five process failures occur within a five-minute interval. You can reset the accumulation interval by updating the **RapidFailProtectionInterval** metabase property, or you can reset the maximum number of crashes to allow during the interval by updating the **RapidFailProtectionMaxCrashes** metabase property.

**To specify the maximum number of minutes during which worker process failures are accumulated**

- At the command prompt, type:

  **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/RapidFailProtectionInterval** *n*

In the preceding syntax, replace *n* with the number of minutes that you want to set for the accumulation interval.

The **RapidFailProtectionMaxCrashes** property specifies the maximum number of worker process failures that are allowed within the minutes specified by the **RapidFailProtectionInterval** property. The default value for the **RapidFailProtectionMaxCrashes** property is 5.

**To specify the maximum number of worker process failures before rapid-fail protection occurs**

- At the command prompt, type:

  **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/RapidFailProtectionMaxCrashes** *n*

  In the preceding syntax, replace *n* with the maximum number of worker process failures you want to allow before rapid-fail protection occurs.

If you set the **RapidFailProtection** metabase property to **true** but do not set either the **RapidFailProtectionInterval** property or the **RapidFailProtectionMaxCrashes**property, IIS applies the default settings for these properties.

For more information about the rapid-fail metabase properties, see RapidFailProtection Metabase Property, RapidFailProtectionInterval Metabase Property, and RapidFailProtectionMaxCrashes Metabase Property.

**Related Information**

- For information about Web gardens, see Configuring Web Gardens with IIS 6.0.

# Configuring Startup and Shutdown Time Limits

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When the World Wide Web Publishing Service (WWW service) detects that a worker process is unhealthy, it marks the worker process for termination, and then, within configured time limits, it terminates the worker process and starts a new worker process to replace it. Worker processes shut down gracefully (finish request processing) when the shutdown time period is long enough to allow for this. It is important to tune the shutdown time period to fit the load characteristics of the applications that are being processed.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To configure a new worker process to start up within a specified time by using IIS Manager**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and click **Properties**.
2. Click the **Health** tab, and under **Startup time limit**, in the box next to **Worker process must startup within (time in seconds)**, type the desired number of seconds.
3. Click **OK**.

**To configure a worker process to shut down within a specified time by using IIS Manager**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and click **Properties**.
2. Click the **Health** tab, and under **Shutdown time limit**, in the box next to **Worker process must shutdown within (time in seconds)**, type the desired number of seconds.
3. Click **OK**.

Optionally, you can configure shutdown and startup time limits by setting the **StartupTimeLimit** or **ShutdownTimeLimit** metabase properties.

**To set a worker process to shut down or start up within a specified time by using Adsutil.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
   - To set a shutdown time limit, at the command prompt, type:

     **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/***ApplicationPoolName***/ShutdownTimeLimit** *n*

   - To set a startup time limit, at the command prompt, type:

     **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/***ApplicationPoolName***/StartupTimeLimit** *n*

     Replace *n* with the number of seconds that you want to elapse before shutting down or starting the worker process.

For more information about setting startup and shutdown times for the worker processes, see StartupTimeLimit Metabase Property and ShutdownTimeLimit Metabase Property.

# Identifying the Worker Process Process ID

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The nature of some problems, especially those in which an application is hanging, leaking memory, or generally misbehaving, make it necessary to know which application pool goes with which worker process (or processes, if Web gardens are enabled). With the advent of the new worker process isolation mode in IIS 6.0, a Web server potentially can have thousands of individual instances of W3wp.exe running at any given time. Looking for the process ID in Task Manager is pointless in such a case because it is impossible to differentiate between the different worker processes servicing the application pools.

**Procedures**

In worker process isolation mode, you can find the process ID by using the command-line script Iisapp.vbs, which is stored in *systemroot*\System32, and is designed specifically to obtain a list of running worker processes, their process IDs, and the application pool names.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**To find a worker process ID by using Iisapp.vbs**

- Run Iisapp.vbs with no parameters.

For more information about using Iisapp.vbs, see Listing Running Web Applications Using Iisapp.vbs.

In earlier versions of IIS and when IIS 6.0 is running in IIS 5.0 isolation mode, identifying the process ID is relatively easy. If a component installed in a COM+ library application has been activated from a worker process, the procedure for identifying the process ID is the same.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To find the process ID in versions prior to IIS 6.0 and in IIS 5.0 isolation mode**

- From the **Start** menu, point to **Administrative Tools**, and then click **Component Services**.
- In the console tree of the **Component Services** window, double-click **Component Services**, double-click **Computers**, double-click **My Computer**, and then double-click **Running Processes**.
- Under this node, you see the various COM+ applications and their associated process IDs.

# Recycling Worker Processes with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. |

In worker process isolation mode, IIS can be configured to periodically restart worker processes, so that you can recycle faulty Web applications. This helps ensure that applications in those pools remain healthy and that system resources are recovered.

You can configure worker processes to restart based on elapsed time, number of requests served, scheduled times, and two types of memory usage.

This section includes the following information:

- How Worker Process Recycling Works in IIS 6.0
- Determining When to Use Worker Process Recycling
- Configuring Worker Processes for Recycling
- Logging Worker Process Recycling Events in IIS 6.0

# How Worker Process Recycling Works in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Depending on how an application pool is configured for recycling, the World Wide Web Publishing Service (WWW service) in IIS 6.0 can recycle the assigned worker process in two ways:

1. By default, the WWW service establishes an *overlapped recycle*, in which the worker process that is to be terminated is kept running until after a new worker process is started.

2. Alternatively, the WWW service can terminate a worker process and then start a new worker process if the workload justifies this action.

| **Note** |
| --- |
| When the WWW service recycles a worker process, it does not disconnect the existing TCP/IP connection. The HTTP protocol stack (HTTP.sys) establishes and maintains TCP/IP connections. |

In an overlapped recycle scenario, the process that is targeted for a recycle continues to process requests while the WWW service simultaneously creates a replacement worker process. The new worker process is started before the old worker process stops, and requests are then directed to the new process. This design prevents service interruptions, because the old process remains in communication with HTTP.sys to handle requests until it shuts down. Because the shutdown timeout value of an overlapped shutdown or startup is configurable, the worker process can be terminated while it is still serving requests, if it does not finish draining requests within the time limit.

When you configure an application pool to recycle worker processes based on elapsed time, all the worker processes are recycled during the elapsed time set, but they are not recycled all at once. Application recycling is spread over the time set to reduce interruptions of service to client requests.

Similarly, when you configure an application pool to recycle applications based on the number of requests processed, recycling takes place at intervals to spread the overhead that is associated with the worker-process recycling.

# Determining When to Use Worker Process Recycling

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Consider the following general guidelines when you decide whether to enable worker process recycling. The optimal solution is to fix the application that is causing the problem. However, recoding is not always an option, especially in cases when other application code is running that cannot be modified.

Consider using worker process recycling under the following circumstances:

- Your Web server is hosting problematic applications that you cannot fix.
- An application is failing unexpectedly or experiencing undetermined or intermittent problems.
- An application is leaking memory as a result of performance monitoring.
- You have previously implemented a temporary reset solution, such as scheduled execution of the **IISReset** command-line utility.

You might not need to use worker process recycling under the following circumstances:

- Your Web server is hosting a Web site that contains only static content and no custom ISAPI applications.
- Your Web server is hosting an application that has already been thoroughly tested and does not suffer from memory or resource allocation problems.

To effectively use worker process recycling, review the criteria on which recycling can be based, as listed in the following table.

| **Recycle Method** | **Description** | **When to Use** |
|---|---|---|
| ISAPI request | Recycles the worker process based on a request from an ISAPI in the application pool. | Use this method when ISAPI extensions can declare themselves unhealthy. |
| Elapsed time | Recycles the worker process based on a time frame that is specified by the user in minutes. | Use this method when applications have problems running for extended time periods. |
| Number of requests | Recycles the worker process when the number of Hypertext Transfer Protocol (HTTP) requests exceeds a certain threshold. | Use this method when applications are failing based on the number of requests they receive. |
| Scheduled times | Recycling can occur at specified times within a 24-hour period. | Use this method when conditions are similar to the conditions for elapsed time. |
| Virtual memory (reserved plus used) | Recycles the worker process when the worker process virtual memory reaches a certain threshold. | Use this method when the memory heap becomes highly fragmented, which is caused by applications reserving memory multiple times. The symptom is a steady increase in virtual memory. |
| Used memory | Recycles the worker process when the memory that is used by the W3wp.exe process reaches a certain threshold. | Use this method when some applications have memory leaks. |

| | | |
|---|---|---|
| Demand | Recycling happens when an IIS administrator can control the recycling of an entire application pool by using Microsoft® Management Console (MMC) or a script. | Use this method when there is an application pool that is causing problems while the other sites are up and running. Consider recycling the application, as opposed to resetting the entire WWW service. |

# Configuring Worker Processes for Recycling

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure worker processes to restart based on one of several options, including elapsed time, number of requests served, scheduled times, and memory usage; or you can configure worker processes to start on demand. In IIS 6.0, worker process recycling is available only when IIS is running in worker process isolation mode.

For more information about recycling worker processes, see Running IIS 6.0 as an Application Server.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To recycle a worker process immediately on demand**

- In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool for which you want to recycle the worker process, and then click **Recycle**.

**To configure a worker process to be recycled after a set elapsed time**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. On the **Recycling** tab, select the **Recycle worker processes (in minutes)** check box.
3. In the box to the right of **Recycle worker processes (in minutes)**, type the number of minutes that you want to elapse before the worker process is recycled, and then click **OK**.

**To configure a worker process to be recycled after a set number of processing requests**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. On the **Recycling** tab, select the **Recycle worker process (number of requests)** check box.
3. In the box to the right of **Recycle worker process (number of requests)**, type the number of requests to be processed before the worker process is recycled, and then click **OK**.

**To configure a worker process to be recycled at scheduled times**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. On the **Recycling** tab, select the **Recycle worker processes at the following times** check box.
3. To schedule or adjust the time at which the worker process will be recycled, do one of the following:
   - Click **Add** to add a time to the list, and then click **OK**.
   - Click **Remove** to delete a scheduled time, and then click **OK**.
   - Click **Edit** to change an existing time, and then click **OK**.

**Note**

When recycling is set to occur at scheduled times, it might occur off-schedule if the system time is altered. To avoid unintended changes in the time that recycling occurs, recycle the scheduled worker processes soon after the system time is changed.

**To configure a worker process to be recycled after consuming a set amount of memory**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. On the **Recycling** tab, under **Memory recycling**, select the **Maximum virtual memory (in megabytes)** check box.
3. In the box to the right of **Maximum virtual memory (in megabytes)**, type the maximum amount of virtual memory allowed before the worker process is recycled.
4. Select the **Maximum used memory (in megabytes)** check box.
5. In the box to the right of **Maximum used memory (in megabytes)**, type the maximum amount of memory allowed before the worker process is recycled, and then click **OK**.

# Logging Worker Process Recycling Events in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, IIS does not log worker process recycling events. You can, however, enable logging of the worker process recycling events that you specify. Unlike other IIS activity, which is logged in a location and file that you specify, worker process recycling events are written to the system event log.

Logging worker process recycling events can help you with troubleshooting IIS. For example, when troubleshooting an ISAPI extension that implements the new server support function REPORT_UNHEALTHY, you can specify that IIS logs an event for each instance that the ISAPI reports itself unhealthy.

The following table describes the worker process recycling events for which you can enable logging. To enable logging for an event, set the corresponding metabase property to **true**.

| Metabase Property | Description |
|---|---|
| **AppPoolRecycleTime** | Logs event on time-based recycles. |
| **AppPoolRecycleRequests** | Logs event on number of request-based recycles. |
| **AppPoolRecycleSchedule** | Logs event on schedule-based recycless. |
| **AppPoolRecycleMemory** | Logs event on memory-based recycle. |
| **AppPoolRecycleIsapiUnhealthy** | Logs event when worker processes request recycles because an ISAPI reported unhealthy. |
| **AppPoolRecycleOnDemand** | Logs event when an administrator requests a recycle of all processes in the application pool. |
| **AppPoolRecycleConfigChange** | Logs an event if an application pool is recycled because one of the application pool properties that requires a recycle to take effect has changed. |
| **AppPoolRecyclePrivateMemory** | Logs an event if an application pool is recycled based on private memory. |

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters).

**Procedures**

**To enable logging of worker process recycling**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following and then press **ENTER**: **cd %systemroot%\inetpub\adminscripts**
4. At the command prompt, type the following and then press **ENTER**: **cscript adsutil.vbs set w3svc/AppPools/***AppPoolName/EventName* **true**

   Replace *AppPoolName* with the name of the application pool and *EventName* with the name of the event for which you want to enable logging. For example, the following command enables logging of memory-based recycling events for the default AppPool: cscript adsutil.vbs set w3svc/AppPools/DefaultAppPool/AppPoolRecycleMemory true

| Note |
| --- |
| To disable logging of worker process recycling events, end the command-line with **false**. |

The following table shows an example system event log entry along with the event ID for each type of recycling event.

| Recycling Event | Log Entry Generated | Event ID |
| --- | --- | --- |
| **AppPoolRecycleTime** | A worker process with process ID *IDNumber* serving application pool *AppPoolName* has requested a recycle because the worker process reached its allowable processing limit. | 1074 |
| **AppPoolRecycleRequests** | A worker process with process ID *IDNumber* serving application pool *AppPoolName* has requested a recycle because it reached its request limit. | 1075 |
| **AppPoolRecycleSchedule** | A worker process with process ID *IDNumber* serving application pool *AppPoolName* has requested a recycle because it reached its scheduled recycle time. | 1076 |
| **AppPoolRecycleMemory** | A worker process with process ID *IDNumber* serving application pool *AppPoolName* has requested a recycle because it reached its virtual memory limit. | 1077 |
| **AppPoolRecycleIsapiUnhealthy** | A worker process with process ID *IDNumber* serving application pool *AppPoolName* has requested a recycle because an ISAPI it loaded reported unhealthy and requested a recycle. | 1078 |
| **AppPoolRecycleOnDemand** | An administrator has requested a recycle of all worker processes in application pool *AppPoolName*. | 1079 |
| **AppPoolRecycleConfigChange** | The worker processes serving application pool *AppPoolName* are being recycled due to one or more configuration changes in the application pool properties which necessitate a restart of the processes. | 1080 |
| **AppPoolRecyclePrivateMemory** | A worker process with process ID *IDNumber* serving application pool *AppPoolName* has requested a recycle because it reached its private bytes memory limit. | 1177 |

**Related Information**

- For information about setting the LogEventOnRecycle metabase property, see LogEventOnRecycle Metabase Property.
- For information about viewing the system event log, see "Event Viewer" in Help and Support Center for Windows Server 2003.

# IIS 5.0 Isolation Mode in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section explains how to use IIS Manager to configure Internet Information Services (IIS) to execute Web-based applications in IIS 5.0 isolation mode.

This section includes the following information:

- Isolating Applications in IIS 6.0: Describes how to run an application in a memory space that is separate from the Web server memory space.
- Stopping Isolated Applications: Describes how to stop execution of an isolated application.

**Related Information**

- For information about creating or configuring applications, see Configuring Applications in IIS 6.0.
- For an overview of IIS 5.0 isolation mode, see About Configuring Servers for Applications.

# Isolating Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in IIS 5.0 isolation mode. |

*Isolating* applications means configuring them to run in a process (memory space) that is separate from the Web server and other applications. You can configure applications to have one of three levels of application protection:

- Low (in-process) application protection.
- Medium (pooled) application protection.
- High (isolated) application protection.

Note that server-side include (SSI) and Internet Database Connector (IDC) applications cannot be run in a memory space that is separate from the Web server.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set or change the level of application protection**

1. In IIS Manager, expand the local computer, right-click the Web site or the starting-point directory for an application, and then click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab. If you are in the directory that is listed as the **Starting Point** directory, the **Application name** box is filled in.
3. In the **Application protection** box, click the appropriate level of protection.
4. Click **OK**.

The Web server finishes processing any current requests for the application before it creates a separate process. At the next request for the application, the application will run in the appropriate memory space.

You can also configure isolation by updating the **AppIsolated** metabase property setting for each application. For information about how to do this, see the AppIsolated Metabase Property, and Code Examples to Configure Metabase Properties.

| **Note** |
| --- |
| If you need to stop an application that has isolation set to Medium (pooled) or High, see Stopping Isolated Applications. |

# Stopping Isolated Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in IIS 5.0 isolation mode. |

You can stop and unload isolated applications from memory without stopping the Web server, which makes it easier to test and debug applications.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User:Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To stop and unload an isolated application**

1. In IIS Manager, expand the local computer, right-click the Web site or the starting-point directory (the directory on which the application is set) for an application, and then click **Properties**.
2. Click the **Home Directory**, **Virtual Directory**, or **Directory** tab. If you are in the directory that is listed as the **Starting Point** directory, the **Application name** box is filled in.
3. Click **Remove**. If the **Remove** button is unavailable, you are not in the application's starting-point directory.

# Configuring IIS For Dynamic Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

When you install IIS 6.0, it is locked down—only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. By default, features, such as ASP, ASP.NET, Server-Side Includes Events, WebDAV publishing, and FrontPage® Server Extensions do not work unless they are enabled. If you do not enable this functionality after installing IIS, IIS returns a 404 error.

| **Note** |
| --- |
| If you install the Microsoft Visual Studio .NET 2003 development system with IIS 6.0 to host Web applications or XML Web services, you must enable ASP and FrontPage Server Extensions. |

You can configure these request handlers, which are called Web service extensions, by manipulating the Web Service Extensions node in IIS Manager.

This section includes the following information:

- Enabling and Disabling Dynamic Content in IIS 6.0: Explains how to enable or disable an individual Web service extension that is registered in the Web Service Extensions node in IIS Manager. Also discusses how to disable all Web Service extensions at once.
- Adding a Web Service Extension: Describes how to register a Web service extension so that it appears in the Web Service Extensions node in IIS Manager.
- Renaming a Web Service Extension: Explains how to rename a Web service extension that you have added to the list of Web service extensions in IIS.
- Configuring an Application to Call a Web Service Extension: Describes how to specify which Web service extensions an application can call.

# Enabling and Disabling Dynamic Content in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

You can enable or disable Web service extensions individually if they are registered in the Web Service Extensions node in IIS Manager. When enabled features are no longer required, you can disable them through the Web Service Extensions node in IIS Manager.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**
**To enable or disable dynamic content**

1. In IIS Manager, expand the local computer, and then click **Web Service Extensions**.
2. In the details pane, click the Web service extension that you want to enable or disable.
3. Do one of the following:
   - To enable a disabled Web service extension, click **Allow**.
   - To disable an enabled Web service extension, click **Prohibit**.
   - To view the properties of a Web service extension, click **Properties**.
4. Click **OK**.

| **Important** |
| --- |
| If you choose to allow all unknown ISAPI extensions or CGI extensions to run on your Web server, your Web server could be susceptible to computer viruses or worms that exploit these technologies. Instead, select only those specific ISAPI extensions or CGI extensions that you need to run on your Web server. |

You can disable all Web service extensions that are registered on the local computer with one setting in the Web Service Extensions node of IIS Manager.

**To disable all Web service extensions**

1. In IIS Manager, expand the local computer, and then click **Web Service Extensions**.
2. Click **Prohibit all extensions**. A warning message appears.
3. To disable all extensions, click **Yes**. To cancel the action, click **No**. To see more information, click **Help**.
4. If you click **Yes**, the status of each Web service extension is disabled.

# Adding a Web Service Extension

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

Before you can use, or to deny the use of, an HTTP request handler that is not in the list of Web service extensions, you must first register it by adding the HTTP request handler to the list of Web service extensions.

**Procedures**
**To add a new Web service extension**

1. In IIS Manager, expand the computer name, and then click **Web Service Extensions**.
2. In the **details** pane, click **Add a new Web service extension**.
3. In the **Extension name** box, type the name of the new Web service extension.
4. Click **Add**.
5. In the **Path to file** box, type the path or click **Browse** to navigate to any files that the new Web service extension requires, and click **OK**.
6. Optionally, select the **Set extension status to Allowed** check box to automatically set the status of the new Web service extension to **Allowed**.
7. Click **OK**.

# Renaming a Web Service Extension

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

You can rename a Web service extension you have added to IIS. Web service extensions included with IIS, however, cannot be renamed.

You can rename a Web service extension either in IIS Manager, or by directly editing the MetaBase.xml file. Using IIS Manager, the Web service extension you want to rename must first be removed, and then added using the new name. Renaming a Web service extension by editing the MetaBase.xml file is a more direct approach. Both methods are explained here.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To rename a Web service extension using IIS Manager**

1. In IIS Manager, expand the local computer, and then click **Web Service Extensions**.
2. In the details pane, right-click the Web service extension that you want to rename, and then click **Delete**.
3. At the confirmation prompt, click **Yes**.
4. Add the Web service extension, using the new name for the extension.

**To rename a Web service extension by editing the MetaBase.xml file**

1. Make sure that the metabase edit-while-running feature enabled.
2. Open the \Windows\System32\Inetsrv\MetaBase.xml file in a text editor.
3. Search for the name of the Web service extension that you want to rename, and then edit the name.
4. Save and close the MetaBase.xml file.
5. To verify that the Web service extension was renamed, in IIS Manager, right-click the **Web Service Extensions** folder, and click **Refresh**.

   The Web service extension is displayed with the new name in the details pane.

**Related Information**

- For more information about adding a Web service extension, see Adding a Web Service Extension.

# Configuring an Application to Call a Web Service Extension

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

Before an application can call a Web service extension, you must configure the application to make calls to that extension. A Web service extension can accept calls from more than one application.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**
**To allow an application to call a Web service extension**

1. In IIS Manager, expand the local computer, and then click **Web Service Extensions**.
2. Click **Add and allow extensions for a specific application**.
3. From the **Application list** box, click the name of the application. The Web service extension that the application is allowed to call appears in the **Extensions to be allowed** box.
4. Click **OK**.

# Using Command-Line Administration Scripts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can extend administrative control by using scripts to perform server administration tasks. You can use scripts to automate tasks, remotely administer sites and resources, and take advantage of batch files to create and manage objects. For additional information about using command lines, see the "Command-Line Reference" in Help and Support Center for Windows Server 2003.

Internet Information Services (IIS) contains eight supported command-line scripts that use the IIS Windows Management Instrumentation (WMI) provider to configure and manage IIS metabase configurations. Microsoft supports the command-line scripts that are included in IIS, as long as the scripts are not modified. If you need to modify a supported script, you must save it under a new file name, leaving the original script unmodified.

This section includes the following information:

- Managing Web Sites Using Scripts: Describes how to create and manage Web sites by using a command-line script.
- Managing FTP Sites Using Scripts: Describes how to create and manage FTP sites by using a command-line script.
- Managing Web Virtual Directories Using Scripts: Describes how to create and manage Web virtual directories by using a command-line script.
- Managing FTP Virtual Directories Using Scripts: Describes how to create and manage FTP virtual directories by using a command-line script.
- Managing IIS Configurations Using Scripts: Describes how to export and import partial or complete IIS configurations and how to back up and restore IIS configurations by using a command-line script.
- Managing Applications and Web Service Extensions Using Scripts: Describes how to manage Web applications and Web service extensions by using a command-line script.

**Related Information**

- For information about the Active Directory Service Interfaces (ADSI) utility, Adsutil.vbs, see "IIS Administration Technologies" in the IIS Software Development Kit (SDK) on MSDN.
- For information about the **IISReset** command-line utility, see Restarting IIS.

# Using the Adsutil.vbs Administration Script

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Adsutil.vbs is an IIS administration utility that uses Microsoft Visual Basic Scripting Edition (VBScript) with Active Directory Service Interfaces (ADSI) to manipulate the IIS configuration. This script should be run using CScript, which is installed with Windows Script Host.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. From the command prompt, type **runas /profile /***user:mymachine******administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Usage**

Cscript.exe adsutil.vbs COMMAND <path> [<param>...]

Cscript.exe adsutil.vbs COMMAND [<path> [<parameters>...]]

| **Command** | **Description** |
| --- | --- |
| GET Path | Display chosen parameter. |
| SET Path Value | Assign a new value. |
| ENUM Path "/P" | Enumerate all parameters for the path. /P Enumerate the paths only (no data). |
| ENUM_ALL "/P" | Enumerate all parameters. /P Enumerate the paths only (no data). |
| DELETE Path | Delete the path or parameter. |
| CREATE Path [KeyType] | Create the path and assign it the **KeyType**. |
| APPCREATEINPROC Path | Create an in process application. |
| APPCREATEOUTPROC Path | Create an out-of-process application. |
| APPDELETE Path | Delete the application (if present). |
| APPUNLOAD Path | Unload an out-of-process application. |
| APPGETSTATUS Path | Get status of the application. |
| FIND Path | Find the paths where a parameter is set. |
| START_SERVER Path | Start the Web site. |
| STOP_SERVER Path | Stop the Web site. |
| PAUSE_SERVER Path | Pause the Web site. |
| CONTINUE_SERVER Path | Unpause the Web site. |
| HELP | Print all available commands. |

**Notes**

- <Path> is the path of the node for which you are setting the property, combined with the name of the property that you are setting. For example, if you want to set the **ServerComment** for Web server number 1, the command would be as

follows:

```
adsutil SET w3svc/1/ServerComment "Web Server Number 1"
```

The next time IIS Manager is opened, it will reflect the name of the Web server as "Web Server Number 1."

- The switch "-s:*server name*" can be used after any command to execute the command on a remote computer. (See first example below.)

**Examples**

- Cscript.exe adsutil.vbs GET W3SVC/1/ServerBindings -s:remotecomputer1
- Cscript.exe adsutil.vbs SET W3SVC/1/ServerBindings ":81:"
- Cscript.exe adsutil.vbs CREATE W3SVC/1/Root/MyVdir "IIsWebVirtualDir"
- Cscript.exe adsutil.vbs START_SERVER W3SVC/1
- Cscript.exe adsutil.vbs ENUM /P W3SVC

# Managing Web Sites Using Scripts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisweb.vbs, which is stored in *systemroot*\System32, to create and manage Web sites on computers running a member of the Microsoft Windows Server™ 2003 family with Internet Information Services (IIS) 6.0.

Iisweb.vbs performs some of the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

This section contains the following information:

- Creating Web Sites: Describes how to create Web site configurations on local or remote computers by using iisweb.vbs.
- Deleting Web Sites: Describes how to delete Web site configurations on local or remote computers by using iisweb.vbs.
- Starting Web Sites: Describes how to start Web sites on local or remote computers by using iisweb.vbs.
- Stopping Web Sites: Describes how to stop Web sites on local or remote computers by using iisweb.vbs.
- Pausing Web Sites: Describes how to pause Web sites on local or remote computers by using iisweb.vbs.
- Listing Web Sites: Describes how to list Web sites on local or remote computers by using iisweb.vbs.

# Creating Web Sites Using Iisweb.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisweb.vbs, which is stored in *systemroot*\System32, to create Web site configurations on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. This command does not create or destroy content, but it sets up the directory structure and some IIS configuration files.

When you use iisweb.vbs to create a new Web site, you specify only the basic properties that are needed to create the site and identify its contents. Iisweb.vbs uses the same default properties that IIS Manager uses when establishing new Web sites, and it adheres to the same rules for inheriting properties. To configure the more advanced properties of the Web site, use IIS Manager.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisweb /create** *Path SiteName* [**/b** *Port*] [**/i** *IPAddress*] [**/d** *HostHeader*] [**/dontstart**] [**/s** *Computer* [**/u** [*Domain*\]*User***/p** *Password*]]

**Parameters**

*Path*

Specifies the location of content files for the Web site. The path must be local, such as C:\Projects\HTML. If the specified path does not exist, then iisweb.vbs creates the path.

The **Path** parameter must immediately precede the **SiteName** parameter in the command. Otherwise, iisweb.vbs does not interpret the site information correctly.

*SiteName*

Required. Specifies a name for the Web site.

**/b** *Port*

Specifies a TCP port number for the Web site. The default port is 80.

**/i** *IPAddress*

Specifies an IP address for the Web site. The default is **All Unassigned**, which assigns to the site all of the IP addresses on the computer that are not assigned to other sites. Only one site on each IIS server can be set to **All Unassigned**.

**/d** *HostHeader*

Specifies the host header name for the Web site, such as www.microsoft.com. By default, the site does not have a host header name and must be identified by its IP address or port number.

Your DNS server must be configured to associate the host header name with the Internet server. Otherwise, the site will not start.

**/dontstart**

Indicates that the Web site should not be started automatically as soon as it is created, which is contrary to the IIS default.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

| **Note** |
| --- |
| Iisweb.vbs does not verify the port number, IP address, or host header of a Web site, and it does not verify that these bindings are unique within the server. If you create a site with invalid or conflicting bindings, the site will not start. |

**Examples**
**Example 1:**
The following example creates a single Web site configuration on a local computer. The source files for the "My Vacations" Web site are located in the C:\Rome directory. The command uses the **/d** parameter to specify a host header for the Web site. It also uses the **dontstart** parameter so that the Web site is not started automatically, providing time to add the host header name to DNS before starting the site.

**iisweb /create C:\Rome "My Vacations" /d www.reskit.com /dontstart**

In response, **iisweb** displays the following message, along with the basic properties of the new Web site. In this example, "Server" represents the machine name where drive C: is found, "Site Name" is the name you assigned to your Web site, "Metabase Path" represents the metabase key assigned by IIS and is analogous to a registry key in the system registry, "IP" is unassigned by default, "Host" is the header name you assigned, "Port" is set to 80 by default, "Root" is the directory where the Web files reside, and "Status" is set to "STOPPED" as directed.

```
Connecting to server ...Done.

Server = IIS-TEST

Site Name = My Vacations

Metabase Path = W3SVC/1211348328

IP = ALL UNASSIGNED

Host = www.reskit.com

Port = 80

Root = C:\Rome

Status = STOPPED
```

**Example 2:**
The following example creates the Marketing Web site configuration on a remote server (SVR01). It associates the site with files stored in the D:\New Initiatives\Marketing\HTMFiles directory on SVR01. The command uses the **/i** parameter to specify an IP address for the site. It uses the **/s** parameter to specify the remote computer, and the **/u** and **/p** parameters to run the script with the permissions of the user's Administrator account.

**iisweb /create C:\New Initiatives\Marketing\HTMFiles "Marketing" /i 172.30.163.244 /s SVR01 /u Admin6 /p A76QVJ32#**

In response, **iisweb** displays the basic properties of the new Web site. In this example, "Server" represents the server machine name, "Site Name" is the name you assigned to your Web site, "Metabase Path" represents the metabase key assigned by IIS and is analogous to a registry key in the system registry, "IP" represents the assigned DNS address, "Host" is unspecified, "Port" is set to 80 by default, "Root" is the directory where the Web files reside, and "Status" is set to "STARTED" by default.

```
Server = SVR01

Site Name = Marketing

Metabase Path = W3SVC/1907510956

IP = 172.30.163.244

Host = NOT SPECIFIED

Port = 80

Root = C:\New Initiatives\Marketing\HTMFiles

Status = STARTED
```

**Related Information**

- For information about deleting Web site configurations, see Deleting Web Sites Using Iisweb.vbs.
- For information about listing Web site configurations, see Listing Web Sites Using Iisweb.vbs.
- For information about starting Web sites, see Starting Web Sites Using Iisweb.vbs.
- For information about stopping Web sites, see Stopping Web Sites Using Iisweb.vbs.
- For information about pausing Web sites, see Pausing Web Sites Using Iisweb.vbs.

# Deleting Web Sites Using Iisweb.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisweb.vbs, which is stored in *systemroot*\System32, to delete IIS Web site configurations on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0.

This command is useful if you plan to move the site to a new URL or server, or remove it entirely. The content within the site is unaffected by this command, but the site will be inaccessible to users.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type cscript.exe ScriptName (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisweb /delete** *WebSite* [*WebSite*...][**/s** *Computer* [**/u**[*Domain*\]*User***/p** *Password*]]

**Parameters**

*WebSite*

Required. Specifies the site name or metabase path of the Web site. Web sites must be uniquely identified. If more than one Web site has the same descriptive name, you must use the metabase path to identify the Web site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** [*Domain*\]*User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**

**Example 1:**

The following example deletes a single Web site configuration on a local computer. The **/s** parameter is omitted because the local computer is the default setting. The **/u** and **/p** parameters are omitted because the current user of the local computer must be logged on as an administrator to run the command.

**iisweb /delete "My First Novel"**

In response, **iisweb** displays the following message indicating that the site configuration has been deleted, and identifying the Web site by its metabase path:

```
Connecting to server ...Done.
```

```
Server W3SVC/409413479 has been DELETED.
```

**Example 2:**
The following example deletes multiple Web site configurations from a remote computer. The "Finance", "Work Group", and "Logo" Web sites reside on SRV01. The command identifies the first two sites by name and the "Logo" Web site by its metabase path, "W3SVC/79116006". In addition, the command uses the **/s** parameter to specify the remote computer and the **/u** and **/p** parameters to run the commands with the permissions of the user's administrator account.

**iisweb /delete Finance "Work Group" W3SVC/79116006 /s SRV01 /u Alice /p p@##word**

In response, **iisweb** displays the following message indicating that the site configurations have been deleted, and identifying each Web site by its metabase path:

```
Connecting to server ...Done.

Server W3SVC/1977437537 has been DELETED.

Server W3SVC/1509060625 has been DELETED.

Server W3SVC/79116006 has been DELETED.
```

**Related Information**

- For information about creating Web site configurations, see Creating Web Sites in IIS 6.0.
- For information about listing Web site configurations, see Listing Web Sites Using Iisweb.vbs.
- For information about starting Web sites, see Starting Web Sites Using Iisweb.vbs.
- For information about stopping Web sites, see Stopping Web Sites Using Iisweb.vbs.
- For information about pausing Web sites, see Pausing Web Sites Using Iisweb.vbs.

# Starting Web Sites Using Iisweb.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS starts Web sites automatically when they are created, unless the administrator specifically prohibits it with the **/dontstart** parameter. An administrator might prohibit an automatic start, for example, to allow time for the site to be added to DNS.

An administrator might also pause or stop a Web site for troubleshooting purposes. In cases like these, you can use the command-line script iisweb.vbs, which is stored in *systemroot*\System32, to start or restart a Web site.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisweb/start** *WebSite [WebSite...] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

Parameters

*WebSite*

Required. Specifies the site name or metabase path of the Web site. Web sites must be uniquely identified. If more than one Web site has the same descriptive name, you must use the metabase path to identify the Web site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following example starts the "Products" Web site on the local computer. The command identifies the site by its name. In this example, the **/s** parameter is omitted because the server is local. The **/u** and **/p** parameters are omitted because the current user of the local computer must be logged on as an administrator in order to run the command.

**iisweb /start Products**

In response, **iisweb** displays the following message confirming that the site has been started, and identifying it by its metabase path:

```
Connecting to server ...Done.
```

```
Server W3SVC/1977437537 has been STARTED.
```

**Example 2:**
The following example starts the "Products" and "Marketing" Web sites on the SVR01 remote computer. The command identifies the "Products" site by its metabase path, "W3SVC/1509060225", and the "Marketing" site by name. In addition, the command uses the **/s** parameter to specify the remote computer, and the **/u** and **/p** parameters to run the commands with the permissions of the user's administrator account.

**iisweb /start w3svc/1509060225 Marketing /s SVR01 /u Alice /p p@##word**

In response, **iisweb** displays the following message confirming that the sites have been started, and identifying each by its metabase path:

```
Connecting to server ...Done.

Server W3SVC/1509060225 has been STARTED.

Server W3SVC/1977437537 has been STARTED.
```

**Related Information**

- For information about creating Web site configurations, see Creating Web Sites in IIS 6.0.
- For information about pausing Web site configurations, see Pausing Web Sites Using Iisweb.vbs.
- For information about deleting Web site configurations, see Deleting Web Sites Using Iisweb.vbs.
- For information about listing Web site configurations, see Listing Web Sites Using Iisweb.vbs.
- For information about stopping Web sites, see Stopping Web Sites Using Iisweb.vbs.

# Stopping Web Sites Using Iisweb.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If your Web site is not behaving as expected, or if the site is slow or unresponsive, stopping and restarting the site might clear a temporary problem with the software—for example, a faulty worker process.

In cases like these, you can use the command-line script iisweb.vbs, which is stored in *systemroot*\System32, to stop a Web site from responding. The content of the site is unchanged, but it will be inaccessible to users until the site is restarted.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisftp/stop** *WebSite [WebSite...] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**
*WebSite*

Required. Specifies the site name or metabase path of the Web site. Web sites must be uniquely identified. If more than one Web site has the same descriptive name, you must use the metabase path to identify the Web site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following example stops the "Products" Web site on the local computer. The command identifies the site by its name. In this example, the **/s** parameter is unnecessary because the server is local. The **/u** and **/p** parameters are omitted because the current user of the local computer must be logged on as an administrator in order to run the command.

**iisweb /stop Products**

In response, **iisweb** displays the following message confirming that the site has been stopped, and identifying it by its metabase path:

```
Connecting to server ...Done.
```

```
Server W3SVC/1977437537 has been STOPPED.
```

**Example 2:**
The following example stops the "Products" and "Marketing" Web sites on the SVR01 remote computer. The command identifies the "Products" site by its metabase path, "W3SVC/1509060225", and the "Marketing" site by name. In addition, the command uses the **/s** parameter to specify the remote computer, and the **/u** and **/p** parameters to run the commands with the permissions of the user's administrator account.

**iisweb /stop w3svc/1509060225 Marketing /s SVR01 /u Alice /p p@##word**

In response, **iisweb** displays the following message confirming that the sites have been stopped, and identifying each by its metabase path:

```
Connecting to server ...Done.

Server W3SVC/1509060225 has been STOPPED.

Server W3SVC/1977437537 has been STOPPED.
```

**Related Information**
For information about creating Web site configurations, see Creating Web Sites in IIS 6.0.

For information about deleting Web site configurations, see Deleting Web Sites Using Iisweb.vbs.

For information about listing Web site configurations, see Listing Web Sites Using Iisweb.vbs.

For information about starting Web sites, see Starting Web Sites Using Iisweb.vbs.

For information about pausing Web sites, see Pausing Web Sites Using Iisweb.vbs.

# Pausing Web Sites Using Iisweb.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Pausing a Web site is similar to stopping it, except that pausing allows existing processes to continue to their conclusion. Administrators might want to use this function to troubleshoot situations where stopping the Web site abruptly would result in severe issues for users.

You can use the command-line script iisweb.vbs, which is stored in *systemroot*\System32, to pause a Web site indefinitely. Pausing prevents Internet services from accepting new connections, but does not affect requests that are already being processed.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisweb/pause** *WebSite [WebSite...] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**
*WebSite*

Required. Specifies the site name or metabase path of the Web site. Web sites must be uniquely identified. If more than one Web site has the same descriptive name, you must use the metabase path to identify the Web site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following example pauses the "Products" Web site on the local computer. The command identifies the site by its name. In this example, the **/s** parameter is omitted because the server is local. The **/u** and **/p** parameters are omitted because the current user of the local computer must be logged on as an administrator in order to run the command.

**iisweb /pause Products**

In response, **iisweb** displays the following message confirming that the site has been paused, and identifying it by its metabase path:

```
Connecting to server ...Done.

Server W3SVC/1977437537 has been PAUSED.
```

**Example 2:**
The following example pauses the "Products" and "Marketing" Web sites on the SVR01 remote computer. The command identifies the "Products" site by its metabase path, "W3SVC/1509060225", and the "Marketing" site by name. In addition, the command uses the **/s** parameter to specify the remote computer, and the **/u** and **/p** parameters to run the commands with the permissions of the user's administrator account.

**iisweb /pause w3svc/1509060225 Marketing /s SVR01 /u Alice /p p@##word**

In response, **iisweb** displays the following message confirming that the sites have been paused, and identifying each by its metabase path:

```
Connecting to server ...Done.

Server W3SVC/1509060225 has been PAUSED.

Server W3SVC/1977437537 has been PAUSED.
```

**Related Information**

- For information about creating Web site configurations, see Creating Web Sites in IIS 6.0.
- For information about deleting Web site configurations, see Deleting Web Sites Using Iisweb.vbs.
- For information about listing Web site configurations, see Listing Web Sites Using Iisweb.vbs.
- For information about starting Web sites, see Starting Web Sites Using Iisweb.vbs.
- For information about stopping Web sites, see Stopping Web Sites Using Iisweb.vbs.

# Listing Web Sites Using Iisweb.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Administrators frequently need to check the status and/or properties of their Web sites. To check the status of your Web site, you can use the command-line script iisweb.vbs, which is stored in *systemroot*\System32, to query or list Web sites on remote or local computers.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisweb/query** *[WebSite [WebSite...]]*

*[* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

*WebSite*

Limits the query to the specified Web site. Type the name or metabase path of the Web site. The default displays all of the Web sites on the IIS server. Web sites must be uniquely identified. If more than one Web site has the same descriptive name, you must use the metabase path to identify the Web site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

| **Note** |
| --- |
| The query output includes the site name, metabase path, and the IP address, or **All** to indicate all unassigned addresses. |

**Examples**

**Example 1:**

The following example displays the Web sites on the local computer. Omitting the **WebSite** parameter limits the search to the local computer, eliminating the need for the **/s** parameter. The **/u** and **/p** parameters are omitted because the user running the command prompt must be logged on as an administrator for the local machine.

**iisweb /query**

In response, **iisweb** displays all of the IIS Web sites that are on the local computer, including their site names, metabase paths, status, IP addresses, ports, and host names.

```
Site Name (Metabase Path) Status IP Port Host

Default Web Site (W3SVC/1) STARTED ALL 90 N/A

Products (W3SVC/1948497947) STARTED 172.30.163.244 80 N/A

Finance (W3SVC/1211348328) STARTED 172.30.152.249 80 N/A

Marketing (W3SVC/1907510956) STARTED ALL 95 www.contoso.com
```

**Example 2:**
The following example searches for the "Marketing" Web site on the remote computer SVR01. It uses the **/s** parameter to specify the remote computer, and the **/u** and **/p** parameters to run the script with the permissions of the user's administrator account.

**iisweb /query Marketing /s SVR01 /u Admin6 /p A76QVJ32#**

In response, **iisweb** displays the "Marketing" Web site on SVR01.

```
Site Name (Metabase Path) Status IP Port Host

Marketing (W3SVC/1907510956) STOPPED ALL 80 www.contoso.com
```

**Example 3:**
The following example displays the status of the "Marketing" and "Finance" Web sites on the local computer. It uses site names to identify the Web sites, although metabase paths are also valid.

**iisweb /query Marketing Finance**

In response, **iisweb** displays the status of the "Marketing" and "Finance" Web sites.

```
Site Name (Metabase Path) Status IP Port Host

Marketing (W3SVC/1907510956) STARTED ALL 80 www.contoso.com

Finance (W3SVC/1509060625) STARTED 192.168.125.225 80 www.contosofinance.com
```

**Related Information**

- For information about creating Web site configurations, see Creating Web Sites in IIS 6.0.
- For information about deleting Web site configurations, see Deleting Web Sites Using Iisweb.vbs.
- For information about starting Web sites, see Starting Web Sites Using Iisweb.vbs.
- For information about stopping Web sites, see Stopping Web Sites Using Iisweb.vbs.
- For information about pausing Web sites, see Pausing Web Sites Using Iisweb.vbs.

# Managing FTP Sites Using Scripts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisftp.vbs, which is stored in *systemroot*\System32, to create and manage File Transfer Protocol (FTP) sites on computers running a member of the Microsoft Windows Server™ 2003 family with Internet Information Services (IIS) 6.0.

Iisftp.vbs performs some of the same operations that are available in IIS Manager. You can use either tool to administer IIS FTP sites.

This section includes the following information:

- Creating FTP Sites Using Iisftp.vbs: Describes how to create and isolate FTP sites on local or remote computers by using iisftp.vbs.
- Deleting FTP Sites Using Iisftp.vbs: Describes how to delete FTP sites on local or remote computers by using iisftp.vbs.
- Starting FTP Sites Using Iisftp.vbs: Describes how to start FTP sites on local or remote computers by using iisftp.vbs.
- Stopping FTP Sites Using Iisftp.vbs: Describes how to stop FTP sites on local or remote computers by using iisftp.vbs.
- Pausing FTP Sites Using Iisftp.vbs: Describes how to pause FTP sites on local or remote computers by using iisftp.vbs.
- Listing FTP Sites Using Iisftp.vbs: Describes how to list FTP sites on local or remote computers by using iisftp.vbs.
- Setting Active Directory User Isolation Using Iisftp.vbs: Describes how to get and set user isolation on local or remote computers by using iisftp.vbs.

# Creating FTP Sites Using Iisftp.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisftp.vbs, which is stored in *systemroot*\System32, to create File Transfer Protocol (FTP) site configurations on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. This command does not create or destroy content, but it sets up the directory structure and IIS configuration files.

When you use iisftp.vbs to create a new FTP site, you specify only the basic properties that are needed to create the site and identify its contents. Iisftp.vbs uses the same default properties that IIS Manager uses when establishing new FTP sites, and it adheres to the same rules for inheriting properties. To configure the more advanced properties of the FTP site, use IIS Manager.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisftp /create** *Path SiteName* **iisftp /create** *Path SiteName*[**/b** *Port*] [**/i** *IPAddress*] [**/dontstart**] [**/isolation** *{AD|Local}* [**/ADdomain DomainName /ADadmin** *[Domain\]User* **/ADpass** *Password]]* [**/s** *Computer* [**/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

*Path*

Specifies the location of content files for the FTP site. The path must be local, such as C:\Projects\HTML. If the specified path does not exist, then iisftp.vbs creates the path.

The **Path** parameter must immediately precede the **SiteName** parameter in the command. Otherwise, iisftp.vbs does not interpret the site information correctly.

*SiteName*

Required. Specifies a name for the FTP site.

**/b** *Port*

Specifies a TCP port number for the FTP site. Port 21 is the default.

**/i** *IPAddress*

Specifies an IP address for the FTP site. The default is **All Unassigned**, which assigns to the site all of the IP addresses on the computer that are not assigned to other sites. Only one site on each IIS server can be set to **All Unassigned**.

**/dontstart**

Indicates that the FTP site should not be started automatically as soon as it is created, which is contrary to the IIS default.

**/isolation** *{AD|Local}*

Provides for one of two user isolation modes, AD (Active Directory) and Local. If the **/isolation** switch is not present, the site will not use user isolation.

**/ADdomain** *DomainName*

Only valid when the **/isolation** switch is set to AD. This value is the name of the Active Directory domain.

**/ADadmin** *[Domain\]User*

Only valid when the **/isolation** switch is set to AD. This value is the name of the administrator in the *user*, *domain\user*, or *user@domain* format.

**/ADpass** *Password*

Only valid when the **/isolation** switch is set to AD. This value is the administrator's password.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

| **Note** |
| --- |
| Iisftp.vbs does not verify the port number or IP address of an FTP site, and it does not verify that these bindings are unique within the server. If you create a site with invalid or conflicting bindings, the site will not start. |

**Examples**

**Example 1**

The following example creates the "Archive" FTP site on the local computer, and associates it with content in the D:\Public\Archives directory. The command omits all optional parameters and accepts the default values. The "Archives" subdirectory does not exist, so iisftp.vbs will create it in the "Public" directory.

**iisftp /create D:\Public\Archives Archive**

In response, **iisftp** displays the following message along with the basic properties of the new FTP site. In this example, "Server" represents the machine name where drive D: is found, "Site Name" is the name you assigned to your FTP site, "Metabase Path" represents the metabase key assigned by IIS and is analogous to a registry key in the system registry, "IP" is unassigned by default, "Port" is set to 21 by default, "Root" is the directory where the FTP files reside, and "Status" is set to "STARTED" by default.

```
Connecting to server ...Done.

Server = RESKIT

Site Name = Archive

Metabase Path = MSFTPSVC/1452008083

IP = ALL UNASSIGNED

Port = 21

Root = D:\Public\Archives

Status = STARTED
```

**Example 2**

The following example creates the "Drivers" FTP site configuration on a remote server. It associates the site with files stored in the C:\Public\Download directory on SVR16. The command uses the **/i** parameter to specify an IP address for the site and the **/dontstart** parameter to prevent the site from starting automatically. It uses the **/s** parameter to specify the remote computer, and the **/u** and **/p** parameters to run the script with the permissions of the user's administrator account.

**iisftp /create C:\Public\Download Drivers /i 172.31.69.150 /dontstart /s SVR16 /u Admin6 /p p@ssWor#**

In response, **iisftp** displays the following message along with the basic properties of the new FTP site. In this example, "Server" represents the server machine, "Site Name" is the name you assigned to your FTP site name, "Metabase Path" represents the metabase key assigned by IIS and is analogous to a registry key in the system registry, "IP" represents the assigned DNS address, "Port" is set to 21 by default, "Root" is the directory where the FTP files reside, and "Status" is set to "STOPPED" as directed.`Connecting to server ...Done.`

```
Server = SVR16

Site Name = Drivers

Metabase Path = MSFTPSVC/1932955329

IP = 172.31.69.150

Port = 21

Root = C:\Public\Download

Status = STOPPED
```

Example 3

The following example creates the "ChangePages" FTP site configuration on a remote server using Active Directory isolation. It associates the site with files stored in the C:\Public\Update directory on SVR16. The command uses the **/dontstart** parameter to prevent the site from starting automatically. The command also uses the **/isolation** parameter with a value of AD to specify Active Directory isolation, and uses the **ADdomain**, **ADadmin**, and **ADpass** parameters to provide the name of the Active Directory domain, as well as the domain administrator name and password.The **/s** parameter specifies the server name, and the **/u** and **/p** parameters specify the user account and password for the administrator of the remote server, under whose permissions the iisftp.vbs script runs.

**iisftp /create C:\Public\Update ChangePages /dontstart /isolation AD /ADdomain DomainA /ADadmin Admin5 /ADpass p@ssWor# /s SVR16 /u Admin6 /p p@SSwor#**

In response, iisftp displays the following message along with the basic properties of the new FTP site. In this example, "Server" represents the server machine; "Site Name" is the name assigned to the FTP site, in this example, ChangePages; "Metabase Path" represents the metabase key assigned by IIS and is analogous to a registry key in the system registry; "IP" is unassigned by default; "Port" is set to 21 by default; "Root" is the directory where the FTP files reside; "IsoMode" is the isolation mode, which is Active Directory; and "Status" is set to STOPPED as specified in the command.

```
Connecting to server....Done.

Server = SVR16

Site Name = ChangePages

Metabase Path = MSFTPSCV/945765436

IP = ALL UNASSIGNED

Port = 21

Root = C:\Public\Update

IsoMode = AD

Status = STOPPED
```

**Related Information**

- For information about deleting FTP configurations, see Deleting FTP Sites Using Iisftp.vbs.
- For information about starting FTP sites, see Starting FTP Sites Using Iisftp.vbs.
- For information about stopping FTP sites, see Stopping FTP Sites Using Iisftp.vbs.
- For information about pausing FTP sites, see Pausing FTP Sites Using Iisftp.vbs.
- For information about listing FTP sites, see Listing FTP Sites Using Iisftp.vbs.
- For information about setting user isolation, see Setting Active Directory User Isolation Using Iisftp.vbs.

# Deleting FTP Sites Using Iisftp.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisftp.vbs, which is stored in *systemroot*\System32\, to delete IIS File Transfer Protocol (FTP) site configurations on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0.

This command is useful if you plan to move the site to a new URL or server, or if you intend to remove it entirely. The content within the site is unaffected by this command, but the site will be inaccessible to users.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisftp /delete** *FTPSite [FTPSite...] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

*FTPSite*

Required. Specifies the site name or metabase path of the FTP site. FTP sites must be uniquely identified. If more than one FTP site has the same descriptive name, you must use the metabase path to identify the FTP site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**

**Example 1:**

The following example deletes a single FTP site configuration on a local computer. The content of the "Downloads" FTP site remains unchanged, but will be unavailable to users. The **/s** parameter is omitted because the local computer is the default setting. The **/u** and **/p** parameters are omitted because the current user of the local computer must be logged on as an administrator to run the command.

**iisftp /delete Downloads**

In response, **iisftp** displays the following message indicating that the site configuration has been deleted, and identifying the FTP site by its metabase path. The metabase path is the IIS equivalent of an entry in the system registry.

```
Connecting to server ...Done.
```

```
Server MSFTPSVC/2109607139 has been DELETED.
```

**Example 2**

The following example deletes multiple FTP site configurations from a remote computer. The "Downloads", "Drivers", and "Library" FTP sites reside on SVR01. The command identifies the first two sites by name and the "Logo" FTP site by its metabase path, "MSFTPSVC/1932955329". In addition, the command uses the **/s** parameter to specify the remote computer and the **/u** and **/p** parameters to run the commands with the permissions of the user's administrator account.

**iisftp /delete Downloads Drivers MSFTPSVC/2109607139 /s Svr01 /u Tom /p p@##word**

In response, **iisftp** displays the following message indicating that the site configurations have been deleted, and identifying each FTP site by its metabase path.

```
Connecting to server ...Done.

Server MSFTPSVC/2109607139 has been DELETED.

Server MSFTPSVC/1879143292 has been DELETED.

Server MSFTPSVC/1932955329 has been DELETED.
```

**Related Information**

- For information about creating FTP configurations, see Creating FTP Sites Using Iisftp.vbs.
- For information about starting FTP sites, see Starting FTP Sites Using Iisftp.vbs.
- For information about stopping FTP sites, see Stopping FTP Sites Using Iisftp.vbs.
- For information about pausing FTP sites, see Pausing FTP Sites Using Iisftp.vbs.
- For information about listing FTP sites, see Listing FTP Sites Using Iisftp.vbs.
- For information about setting user isolation, see Setting Active Directory User Isolation Using Iisftp.vbs.

# Starting FTP Sites Using Iisftp.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS starts File Transfer Protocol (FTP) sites automatically when they are created, unless the administrator specifically prohibits it with the **/dontstart** parameter. An administrator might prohibit an automatic start, for example, to allow time for the site to be added to DNS.

An administrator might also pause or stop an FTP site for troubleshooting purposes. In cases like these, you can use the command-line script iisftp.vbs, which is stored in *systemroot*\System32, to start or restart an FTP site.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisftp/start** *FTPSite [FTPSite...] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**
*FTPSite*

Required. Specifies the site name or metabase path of the FTP site. FTP sites must be uniquely identified. If more than one FTP site has the same descriptive name, you must use the metabase path to identify the FTP site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following example starts the "Products" FTP site on the local computer. The command identifies the site by its name. In this example, the **/s** parameter is omitted because the server is local. The **/u** and **/p** parameters are omitted because the current user of the local computer must be logged on as an administrator in order to run the command.

**iisftp /start Products**

In response, **iisftp** displays the following message confirming that the site has been started, and identifying it by its metabase path. The metabase path is the IIS equivalent of an entry in the system registry.

```
Connecting to server ...Done.
```

```
Server MSFTPSVC/1452008085 has been STARTED.
```

**Example 2:**
The following example starts the "Products" and "Marketing" FTP sites on the SVR01 remote computer. The command identifies the "Products" site by its metabase path, "MSFTPSVC/1452008085", and the "Marketing" site by name. In addition, the command uses the **/s** parameter to specify the remote computer, and the **/u** and **/p** parameters to run the commands with the permissions of the user's administrator account.

**iisftp /start MSFTPSVC/1452008085 Marketing /s SVR01 /u Jim /p p@##word**

In response, **iisftp** displays the following message, confirming that the sites have been started, and identifying each site by its metabase path:

```
Connecting to server ...Done.

Server MSFTPSVC/1452008085 has been STARTED.

Server MSFTPSVC/1879143292 has been STARTED.
```

**Related Information**

- For information about creating FTP configurations, see Creating FTP Sites Using Iisftp.vbs.
- For information about deleting FTP configurations, see Deleting FTP Sites Using Iisftp.vbs.
- For information about stopping FTP sites, see Stopping FTP Sites Using Iisftp.vbs.
- For information about listing FTP sites, see Listing FTP Sites Using Iisftp.vbs.
- For information about setting user isolation, see Setting Active Directory User Isolation Using Iisftp.vbs.

# Stopping FTP Sites Using Iisftp.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If your File Transfer Protocol (FTP) site is not behaving correctly, or if the site is slow or unresponsive, stopping and restarting the site might clear a temporary problem with the software—for example, a faulty worker process.

In cases like these, you can use the command-line script iisftp.vbs, which is stored in *systemroot*\System32, to stop an FTP site from responding. The content of the site is unchanged, but it will be inaccessible to users until the site is restarted.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisftp/stop** *FTPSite [FTPSite...] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**
*FTPSite*

Required. Specifies the site name or metabase path of the FTP site. FTP sites must be uniquely identified. If more than one FTP site has the same descriptive name, you must use the metabase path to identify the FTP site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following example stops the "Drivers" FTP site on the local computer. The command identifies the site by its name. In this example, the **/s** parameter is omitted because the server is local. The **/u** and **/p** parameters are omitted because the current user of the local computer must be logged on as an administrator in order to run the command.

**iisftp /stop Drivers**

In response, **iisftp** displays the following message confirming that the site has been stopped, and identifying the site by its metabase path. The metabase path is the IIS equivalent of an entry in the system registry.

```
Connecting to server ...Done.
```

```
Server MSFTPSVC/1879143292 has been STOPPED.
```

**Example 2:**
The following example stops the "Downloads" and "Drivers" FTP sites on the SVR01 remote computer. The command identifies the "Downloads" site by its metabase path, "MSFTPSVC/2109607139", and the "Drivers" site by name. In addition, the command uses the **/s** parameter to specify the remote computer, and the **/u** and **/p** parameters to run the commands with the permissions of the user's administrator account.

**iisftp /stop MSFTPSVC/2109607139 Drivers /s SVR01 /u Tom /p p###word**

In response, iisftp displays the following message confirming that the sites have been stopped, and identifying each site by its metabase path:

```
Connecting to server ...Done.

Server MSFTPSVC/2109607139 has been STOPPED.

Server MSFTPSVC/1879143292 has been STOPPED.
```

**Related Information**

- For information about creating FTP configurations, see Creating FTP Sites Using Iisftp.vbs.
- For information about deleting FTP configurations, see Deleting FTP Sites Using Iisftp.vbs.
- For information about starting FTP sites, see Starting FTP Sites Using Iisftp.vbs.
- For information about pausing FTP sites, see Pausing FTP Sites Using Iisftp.vbs.
- For information about listing FTP sites, see Listing FTP Sites Using Iisftp.vbs.
- For information about setting user isolation, see Setting Active Directory User Isolation Using Iisftp.vbs.

# Pausing FTP Sites Using Iisftp.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Pausing a File Transfer Protocol (FTP) site is similar to stopping it; except that pausing the FTP site lets existing processes continue to their conclusion. Administrators might want to use this function to troubleshoot situations where stopping the FTP site abruptly would result in severe issues for users.

You can use the command-line script iisftp.vbs, which is stored in *systemroot*\System32, to pause a Web site indefinitely. Pausing prevents FTP services from accepting new connections, but it does not affect requests that are already being processed.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisftp/pause** *FTPSite [FTPSite...] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**
*FTPSite*

Required. Specifies the site name or metabase path of the FTP site. FTP sites must be uniquely identified. If more than one FTP site has the same descriptive name, you must use the metabase path to identify the FTP site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following example pauses the "Drivers" FTP site on the local computer. The command identifies the site by its name. In this example, the **/s** parameter is omitted because the server is local. The **/u** and **/p** parameters are omitted because the current user of the local computer must be logged on as an administrator in order to run the command.

**iisftp /pause Drivers**

In response, **iisftp** displays the following message confirming that the site has been paused, and identifying the site by its metabase path. The metabase path is the IIS equivalent of an entry in the system registry.

```
Connecting to server ...Done.

Server MSFTPSVC/1879143292 has been PAUSED.
```

**Example 2:**
The following example pauses the "Downloads" and "Drivers" FTP sites on the SVR01 remote computer. The command identifies the "Downloads" site by its metabase path, "MSFTPSVC/2109607139", and the "Drivers" site by name. In addition, the command uses the **/s** parameter to specify the remote computer, and it uses the **/u** and **/p** parameters to run the commands with the permissions of the user's administrator account.

**iisftp /pause MSFTPSVC/2109607139 Drivers /s SVR01 /u Tom /p pa##word**

In response, **iisftp** displays the following message confirming that the sites have been paused, and identifying each site by its metabase path:

```
Connecting to server ...Done.

Server MSFTPSVC/2109607139 has been PAUSED.

Server MSFTPSVC/1879143292 has been PAUSED.
```

**Related Information**

- For information about creating FTP configurations, see Creating FTP Sites Using Iisftp.vbs.
- For information about deleting FTP configurations, see Deleting FTP Sites Using Iisftp.vbs.
- For information about starting FTP sites, see Starting FTP Sites Using Iisftp.vbs.
- For information about stopping FTP sites, see Stopping FTP Sites Using Iisftp.vbs.
- For information about listing FTP sites, see Listing FTP Sites Using Iisftp.vbs.
- For information about setting user isolation, see Setting Active Directory User Isolation Using Iisftp.vbs.

# Listing FTP Sites Using Iisftp.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Administrators frequently need to check the status and/or properties of their File Transfer Protocol (FTP) sites. To check the status of your site, you can use the command-line script iisftp.vbs, which is stored in *systemroot*\System32, to query or list FTP sites on remote or local computers.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisftp/query** *[FTPSite [FTPSite...]]*

*[* **/s** *Computer [* **/u** *[Domain* **** *]User* **/p** *password]]*

**Parameters**

*FTPSite*

Required. Specifies the site name or metabase path of the FTP site. FTP sites must be uniquely identified. If more than one FTP site has the same descriptive name, you must use the metabase path to identify the FTP site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain* **** *]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

| **Note** |
| --- |
| The query output includes the site name, metabase path, port, status, and the IP address, or *All* to indicate all unassigned addresses. |

**Examples**

**Example 1:**

The following example displays the FTP sites that are on the local computer. Omitting the *FTPSite* parameter limits the search to the local computer, eliminating the need for the **/s** parameter. The **/u** and **/p** parameters are omitted because the user running the command prompt must be logged on as an administrator for the local machine.

**iisftp /query**

In response, **iisftp** displays all of the IIS FTP sites that are on the local computer, including site names, metabase paths, status, IP addresses, and ports.

```
Connecting to server ...Done.

Site Name (Metabase Path) Status IP Port

Default FTP Site (MSFTPSVC/1) STARTED ALL 21

My Site (MSFTPSVC/858812021) STOPPED ALL 21

Drivers (MSFTPSVC/1932955329) STARTED 172.30.163.249 21

Downloads (MSFTPSVC/1879143292) STARTED ALL 25

Library (MSFTPSVC/2109607139) STARTED 172.30.163.126 21
```

**Example 2:**
The following example searches for the "Download" FTP site on the remote computer, SVR01. It uses the **/s** parameter to specify the remote computer, as well as the **/u** and **/p** parameters to run the script with the permissions of the user's administrator account.

**iisftp /query Download /s SVR01 /u Admin6 /p A76QVJ32#**

In response, **iisftp** displays the "Download" FTP site on SVR01.

```
Site Name (Metabase Path) Status IP Port

Downloads (MSFTPSVC/1879143292) STARTED ALL 25
```

**Example 3:**
The following example displays the status of the "Downloads" and "Drivers" FTP sites on the local computer. It uses site names to identify the FTP sites, although metabase paths are also valid.

**iisftp /query Downloads Drivers**

In response, **iisftp** displays the status of the "Downloads" and "Drivers" FTP sites.

```
Site Name (Metabase Path) Status IP Port

Downloads (MSFTPSVC/1879143292) STARTED ALL 25

Drivers (MSFTPSVC/1932955329) STARTED 172.30.163.249 21
```

**Related Information**

- For information about creating FTP configurations, see Creating FTP Sites Using Iisftp.vbs.
- For information about deleting FTP configurations, see Deleting FTP Sites Using Iisftp.vbs.
- For information about starting FTP sites, see Starting FTP Sites Using Iisftp.vbs.
- For information about stopping FTP sites, see Stopping FTP Sites Using Iisftp.vbs.
- For information about pausing FTP sites, see Pausing FTP Sites Using Iisftp.vbs.
- For information about setting user isolation, see Setting Active Directory User Isolation Using Iisftp.vbs.

# Setting Active Directory User Isolation Using Iisftp.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisftp.vbs, which is stored in *systemroot*\System32, to query and set Active Directory directory service properties for a user's File Transfer Protocol (FTP) home directory. For more information about user isolation, see Isolating FTP Users.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisftp /GetADProp** *UserID [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**iisftp /SetADProp** *UserID {msIIS-FTPDir|msIIS-FTPRoot} PropertyValue [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

**/GetADProp**

Returns property values for a specific Active Directory user.

**/SetADProp**

Sets the property values for a specific Active Directory user.

*UserID*

Required. Specifies the logon ID of the Active Directory user.

*msIIS-FTPDir|msIIS-FTPRoot*

Required to set the home directory properties. Specifies isolation at either the directory or root level.

*PropertyValue*

Required to set the home directory properties. Specifies the value of the home directory and relative path.

**/s** *Computer*

Specifies the name or IP address of a remote computer (with no backslashes). The default is the local computer.

**/u** *[Domain\]User*

Connects to the Active Directory with the permissions of the specified user account in the user or *domain\user* format. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following example returns the Active Directory properties of user "BSmith" on the local computer:

**iisftp /GetADProp BSmith**

In response, **iisftp** displays the following message, which is followed by the Active Directory properties for BSmith in the default domain. In this example, "Server" represents the machine name where drive D: is found, "Site Name" is the name you assigned to the FTP site, "Metabase Path" represents the metabase key assigned by IIS, "IP" is unassigned by default, "Port" is set to 21 by default, "Root" is the directory where the FTP files reside, "IsoMode" is set to "FTPRoot", and "Status" is set to "STARTED" by default.

```
Connecting to server ...Done.

Server = FTPServer

Site Name = Smittys

Metabase Path = MSFTPSVC/1452008083

IP = ALL UNASSIGNED

Port = 21

Root = D:\Public\Archives

IsoMode = FTPRoot

Status = STARTED
```

**Example 2:**
The following example sets the home directory properties in the Active Directory User object properties for user "Sales\MJones" on a remote computer:

**iisftp /SetADProp Sales\MJones FTPDir MikesDir /s SVR16 /u Sales\Admin6 /p p@ssWor#**

In response, **iisftp** displays the following message, which includes the Active Directory properties for MJones. In this example, "Server" represents the machine name where drive D: is found, "Site Name" is the name you assigned to the FTP site, "Metabase Path" represents the metabase key assigned by IIS, "IP" is unassigned by default, "Port" is set to 21 by default, "Root" is the directory where the FTP files reside, "IsoMode" is set to "FTPDir", and "Status" is set to "STARTED" by default.

```
Connecting to server ...Done.

Server = SVR16

Site Name = MikesDir

Metabase Path = MSFTPSVC/1452004045

IP = ALL UNASSIGNED

Port = 21

Root = D:\Public\Archives

IsoMode = FTPDir

Status = STARTED
```

**Related Information**

- For information about creating FTP configurations, see Creating FTP Sites Using Iisftp.vbs.
- For information about deleting FTP configurations, see Deleting FTP Sites Using Iisftp.vbs.
- For information about starting FTP sites, see Starting FTP Sites Using Iisftp.vbs.
- For information about stopping FTP sites, see Stopping FTP Sites Using Iisftp.vbs.
- For information about pausing FTP sites, see Pausing FTP Sites Using Iisftp.vbs.
- For information about listing FTP sites, see Listing FTP Sites Using Iisftp.vbs.

# Managing Web Virtual Directories Using Scripts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisvdir.vbs, which is stored in *systemroot*\System32 to create, delete, and list Web virtual directories on computers running a member of the Microsoft Windows Server™ 2003 family with Internet Information Services (IIS) 6.0.

Iisvdir.vbs performs some of the same operations that are available in IIS Manager. You can use either tool to administer IIS Web virtual directories.

This section includes the following information:

- Creating Web Virtual Directories Using Iisvdir.vbs: Describes how to create Web virtual directories on local or remote computers by using iisvdir.vbs.
- Deleting Web Virtual Directories Using Iisvdir.vbs: Describes how to delete Web virtual directories on local or remote computers by using iisvdir.vbs.
- Listing Web Virtual Directories Using Iisvdir.vbs: Describes how to list Web virtual directories on local or remote computers by using iisvdir.vbs.

# Creating Web Virtual Directories Using Iisvdir.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script Iisvdir.vbs, which is stored in *systemroot*\System32, to create a new Web virtual directory on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. This command does not create or destroy content, but it sets up the virtual directory structure and IIS configuration files.

When you use Iisvdir.vbs to create a new Web virtual directory, you specify only the basic properties that are needed to create the site and identify its contents. Iisvdir.vbs uses the same default properties that IIS Manager uses when establishing new virtual directories, and it adheres to the same rules for inheriting properties. To configure the more advanced properties of the virtual directory, use IIS Manager.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /**User:MyComputer****Administrator **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisvdir /create** *WebSite[/VirtualPath] Name PhysicalPath [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

*WebSite*

Required. Specifies the descriptive name, or the metabase path, of the Web site.

*VirtualPath*

Specifies a path, if necessary, to the virtual directory within the Web site. This parameter is required when the virtual directory is not located at the root of the Web site.

*Name*

Required. Specifies a name for the virtual directory. Virtual directory names do not have to be unique. However, when a Web site includes a virtual directory and a physical directory with the same name, the physical directory will not be visible on the Internet.

*PhysicalPath*

Specifies a physical directory where the content for the virtual directory resides. You must specify a path on the local computer, such as C:\Project\HTML. If the specified directory does not exist, **iisvdir** will create it.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

| **Note** |
| --- |
| The **WebSite** (or **WebSite/VirtualPath**), **Name**, and **PhysicalPath** parameters must always appear in the specified order on the command line. Otherwise, Iisvdir.vbs does not interpret the information correctly. |

**Examples**

**Example 1:**

The following example creates the "Insurance" virtual directory at the root of the "Finance" Web site on the local computer. It associates the directory with content that is currently stored in the C:\Projects\Insurance\HTML directory.

**iisvdir /create Finance Insurance c:\projects\insurance\html**

In response, **iisvdir** displays the following success message along with the basic properties of the new virtual directory. In this example, "Virtual Path" reflects the virtual directory structure, "ROOT" represents the physical directory where the content resides, and "Metabase Path" represents the metabase key that is assigned by IIS.

```
Connecting to server ...Done.

Virtual Path = Finance/Insurance

ROOT = c:\projects\insurance\html

Metabase Path = W3SVC/1509060625/ROOT/Insurance
```

**Example 2:**

The following example creates the "Updates" virtual subdirectory on the "Finance" Web site on the remote computer. The command uses the metabase path of the "Finance" Web site, "W3SVC/1509060625", to identify the Web site, and it indicates the virtual path, "Finance/Insurance", by appending it to the Web site name. In addition, the command associates the "Updates" directory with content that is stored in C:\Newstuff\Web on the remote computer.

The example also uses the **/s** parameter to identify the remote computer, as well as the **/u** and **/p** parameters to run Iisvdir.vbs with the permissions of the user's administrator account.

**iisvdir /create W3SVC/1509060625/Insurance Updates C:\Newstuff\Web /s SVR01 /u Admin01 /p p@SSw#rD2**

In response, **iisvdir** displays the basic properties of the new Web site.

```
Connecting to server ...

Virtual Path = Finance/Insurance/Updates

ROOT = C:\Newstuff\Web

Metabase Path = W3SVC/1509060625/ROOT/Insurance/Updates
```

**Related Information**

- For information about deleting Web virtual directories, see Deleting Web Virtual Directories Using Iisvdir.vbs.
- For information about Listing Web virtual directories, see Listing Web Virtual Directories Using Iisvdir.vbs.

# Deleting Web Virtual Directories Using Iisvdir.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisvdir.vbs, which is stored in *systemroot*\System32, to delete a virtual directory on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. This command does not create or destroy content, but simply deletes the virtual directory structure.

You should not use this command to delete Web sites, as the sites may still be listed, but partially disabled. To delete a Web site, use Deleting Web Sites Using Iisweb.vbs

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisvdir /delete** *WebSite[/VirtualPath] Name [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**
*WebSite*

Required. Specifies the descriptive name, or the metabase path, of the Web site.

*VirtualPath*

Specifies a path, if necessary, to the virtual directory within the Web site. This parameter is required when the virtual directory is not located at the root of the Web site.

*Name*

Required. Specifies a name for the virtual directory. Virtual directory names do not have to be unique. However, when a Web site includes a virtual directory and a physical directory with the same name, the physical directory will not be visible on the Internet.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**

The following example deletes the "Insurance" virtual directory from the "Finance" Web site on the local computer. As a result, not only is the "Insurance" directory deleted, but all of the physical and virtual subdirectories are deleted as well.

**iisvdir /delete Finance/Insurance**

The **/s**, **/u**, and **/p** parameters are unnecessary in this instance because the default computer is local and the user running the command must be logged on as an administrator.

In response, **iisvdir** displays the following message:

```
Connecting to server ...Done.

Web directory Finance/Insurance has been DELETED.
```

**Example 2:**
The following example deletes the "Updates" virtual subdirectory called on the "Finance" Web site of the remote computer. The command uses the metabase path of the "Finance" Web site, "W3SVC/1509060625", to identify the Web site. It indicates the virtual path, "Finance/Insurance", by appending it to the Web site name.

The example also uses the **/s** parameter to identify the remote computer, as well as the **/u** and **/p** parameters to run iisvdir.vbs with the permissions of the user's administrator account.

**iisvdir /delete W3SVC/1509060625/Insurance/Updates /s SVR01 /u Admin01 /p p@SSw#rD2**

In response, **iisvdir** displays the following message:

```
Connecting to server ...Done.

Web directory Finance/Insurance/Updates has been DELETED.
```

**Related Information**

- For information about creating Web virtual directories, see Creating Web Virtual Directories Using Iisvdir.vbs.
- For information about Listing Web virtual directories, see Listing Web Virtual Directories Using Iisvdir.vbs.

# Listing Web Virtual Directories Using Iisvdir.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisvdir.vbs, which is stored in *systemroot*\System32, to list the virtual directory on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0.

The **/query** operation only displays virtual directories at the root of the Web site, or path, and does not search recursively. Also, it does not display physical directories.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisvdir /query** *WebSite[/VirtualPath][* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**
*WebSite*

Required. Specifies the descriptive name, or the metabase path, of the Web site.

*VirtualPath*

Specifies a path, if necessary, to the virtual directory within the Web site. This parameter is required when the virtual directory is not located at the root of the Web site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following example displays the virtual directories at the root of the "Finance" Web site on the local computer:

**iisvdir /query Finance**

The **/s**, **/u**, and **/p** parameters are unnecessary in this instance because the default computer is local and the user running the command must be logged on as an administrator.

In response, **iisvdir** displays the two virtual directories at the root of "Finance" Web site. These directories appear at the "Finance" root even though their physical locations are unrelated. In this example, "Alias" represents the names of the virtual directories, while "Physical Root" is their actual location.

| **Note** |
| --- |
| This display does not include virtual directories that are subdirectories of the site. The procedure for finding subdirectories is demonstrated below in Example 2. |

```
Alias Physical Root

/Personnel D:\Corpdir\FinanceWeb\People

/Insurance C:\Marketing\Insurance\HTMFiles
```

**Example 2:**
The following example displays the virtual directories that are subdirectories of the "Insurance" virtual directory on the "Finance" Web site. The command specifies the "Insurance" directory by using its virtual path.

The example also uses the **/s** parameter to identify the remote computer, as well as the **/u** and **/p** parameters to run iisvdir.vbs with the permissions of the user's administrator account.

**iisvdir /query Finance\Insurance**

In response, iisvdir.vbs reveals the "Current" subdirectory of the "Insurance" virtual directory, and its physical location.

```
Alias Physical Root

/Current C:\Insurance\Monthly\200204
```

**Related Information**

- For information about creating Web virtual directories, see Creating Web Virtual Directories Using Iisvdir.vbs.
- For information about deleting Web virtual directories, see Deleting Web Virtual Directories Using Iisvdir.vbs.

# Managing FTP Virtual Directories Using Scripts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisftpdr.vbs, which is stored in *systemroot*\System32, to create, delete, and list File Transfer Protocol (FTP) virtual directories on computers running a member of the Microsoft Windows Server™ 2003 family with Internet Information Services (IIS) 6.0.

Iisftpdr.vbs performs some of the same operations that are available in IIS Manager. You can use either tool to administer IIS FTP virtual directories.

This section includes the following information:

- Creating FTP Virtual Directories Using Iisftpdr.vbs: Describes how to create FTP virtual directories on local or remote computers by using iisftpdr.vbs.
- Deleting FTP Virtual Directories Using Iisftpdr.vbs: Describes how to delete FTP virtual directories that are on local or remote computers by using iisftpdr.vbs.
- Listing FTP Virtual Directories Using Iisftpdr.vbs: Describes how to list FTP virtual directories that are on local or remote computers by using iisftpdr.vbs.

# Creating FTP Virtual Directories Using Iisftpdr.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisftpdr.vbs, which is stored in *systemroot*\System32\, to create new FTP virtual directories on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. This command does not create or destroy content, but simply sets up the virtual directory structure and IIS configuration files.

When you use iisftpdr.vbs to create a new FTP virtual directory, you specify only the basic properties needed to create the site and identify its contents. Iisftpdr.vbs uses the same default properties that IIS Manager uses when establishing new virtual directories, and it adheres to the same rules for inheriting properties. To configure the more advanced properties of the directory, use IIS Manager.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /**User:MyComputer****Administrator **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisftpdr /create** *FTPSite[VirtualPath] Name PhysicalPath [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

*FTPSite*

Required. Specifies the descriptive name or the metabase path of the FTP site.

*VirtualPath*

Specifies a path, if necessary, to the virtual directory within the Web site. This parameter is required when the virtual directory is not located at the root of the Web site.

*Name*

Required. Specifies a name for the virtual directory. Virtual directory names do not have to be unique. However, when an FTP site includes a virtual directory and a physical directory with the same name, the physical directory will not be visible on the Internet.

*PhysicalPath*

Specifies a physical directory where the content for the virtual directory resides. You must specify a path on the local computer, such as C:\Project\FTP. If the specified directory does not exist, iisftpdr will create it.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

| **Note** |
| --- |
| The **FTPSite** (or **FTPSite/VirtualPath**), **Name**, and **PhysicalPath** parameters must always appear in the specified order on the command line. Otherwise, iisftpdr.vbs does not interpret the information correctly. |

**Examples**

**Example 1:**

The following example creates the virtual directory "Scripts" at the root of the "Test Downloads" FTP site on the local computer. It associates the directory with content that is currently stored in the C:\Libraries\Scripts directory.

**iisftpdr /create "Test Downloads" Scripts C:\Libraries\Scripts**

In response, **iisftpdr** displays the following success message along with the basic properties of the new virtual directory. In this example, "Virtual Path" reflects the virtual directory structure, "ROOT" represents the physical directory where the content resides, and "Metabase Path" represents the metabase key assigned by IIS and is analogous to a registry key in the system registry.

```
Connecting to server ...Done.

Virtual Path = Test Downloads/Scripts

ROOT = C:\Libraries\Scripts

Metabase Path = MSFTPSVC/577648624/ROOT/Scripts
```

**Example 2:**

The following example adds the virtual directory "Updates" to the FTP site of a remote server. The command creates the new virtual directory on the "Finance" FTP site as a subdirectory of the existing virtual directory, "Documents".

The first parameter in the command, **Finance\Documents**, identifies the FTP site and virtual path. The second parameter, **Updates**, specifies the name of the virtual directory, and the third parameter specifies the physical directory on the server that stores the content of the virtual directory.

The command uses the **/s** parameter to identify the server computer, as well as the **/u** and **/p** parameters to run iisftpdr.vbs with the permissions of the user's administrator account.

**iisftpdr /create Finance/Documents Updates C:\inetpub\ftproot\documents\updates /s SVR01 /u Admin01 /p p@SSw#rD2**

In response, **iisftpdr** displays the basic properties of the new FTP site. In this example, the "Finance" FTP site and its "Documents" subdirectory existed on the SVR01 IIS server before the command was issued. If the FTP site or the subdirectory did not exist, the command would have failed.

```
Connecting to server ...Done.

Virtual Path = Finance/Documents/Updates

ROOT = C:\inetpub\ftproot\documents\updates

Metabase Path = MSFTPSVC/2109607139/ROOT/Documents/Updates
```

**Related Information**

- For information about deleting FTP virtual directories, see Deleting FTP Virtual Directories Using Iisftpdr.vbs.
- For information about listing FTP virtual directories, see Listing FTP Virtual Directories Using Iisftpdr.vbs.

# Deleting FTP Virtual Directories Using Iisftpdr.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisftpdr.vbs, which is stored in *systemroot*\System32, to delete an FTP virtual directory on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. This command does not create or destroy content, but simply deletes the virtual directory structure.

You should not use this command to delete FTP sites, as the sites may still be listed, but partially disabled. To delete an FTP site, use Deleting FTP Sites Using Iisftp.vbs.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisftpdr /delete** *FTPSite[VirtualPath]/Name [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**
*FTPSite*

Required. Specifies the descriptive name or the metabase path of the FTP site.

*VirtualPath*

Specifies a path to the virtual directory within the FTP site. This parameter is required when the virtual directory is not located at the root of the FTP site.

*Name*

Required. Specifies a name for the virtual directory. Virtual directory names do not have to be unique. However, when an FTP site includes a virtual directory and a physical directory with the same name, the physical directory will not be visible on the Internet.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**

The following example deletes the "Insurance" virtual directory from the "Finance" FTP site on the local computer. As a result, not only is the directory deleted, but all of the physical and virtual subdirectories are deleted as well.

**iisftpdr /delete Finance/Insurance**

The **/s**, **/u**, and **/p** parameters are unnecessary in this instance because the default computer is local and the user running the command must be logged on as an administrator.

In response, **iisftpdr** displays the following message:

```
Connecting to server ...Done.

FTP directory Finance/Insurance has been DELETED.
```

**Example 2:**
The following example deletes the "Updates" virtual subdirectory on the "Finance" FTP site on the remote computer. The command uses the "Finance" FTP site metabase path, "MSFTPSVC/2109607139", to identify the FTP site. It indicates the virtual path, "Finance/Insurance", by appending it to the FTP site name.

The example also uses the **/s** parameter to identify the remote computer, as well as the **/u** and **/p** parameters to run iisftpdr.vbs with the permissions of the user's administrator account.

**iisftpdr /delete MSFTPSVC/2109607139/Insurance/Updates /s SVR01 /u Admin01 /p p@SSw#rD2**

In response, **iisftpdr** displays the following message:

```
Connecting to server ...Done.

FTP directory Finance/Insurance/Updates has been DELETED.
```

**Related Information**

- For information about creating FTP virtual directories, see Creating FTP Virtual Directories Using Iisftpdr.vbs.
- For information about listing FTP virtual directories, see Listing FTP Virtual Directories Using Iisftpdr.vbs.

# Listing FTP Virtual Directories Using Iisftpdr.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisftpdr.vbs, which is stored in *systemroot*\system32, to list FTP virtual directories on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0.

The **/query** operation displays only virtual directories and not physical directories. Further, it displays only virtual directories at the root of the FTP site or path and does not search recursively.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisftpdr /query** *FTPSite[/VirtualPath][* **/s** *Computer [* **/u** *[Domain \ ]User* **/p** *Password]]*

**Parameters**

*FTPSite*

Required. Specifies the descriptive name or the metabase path of the Web site.

*VirtualPath*

Specifies a path, if necessary, to the virtual directory within the Web site. This parameter is required when the virtual directory is not located at the root of the Web site.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**

**Example 1:**

The following example displays the virtual directories at the root of the "Test Downloads" FTP site on the local computer:

**iisftpdr /query "Test Downloads"**

The **/s**, **/u**, and **/p** parameters are unnecessary in this instance because the default computer is local and the user who is running the command must be logged on as an administrator.

In response, **iisftpdr** displays the "Scripts" virtual directory at the root of the "Test Downloads" FTP site. Note that the directory appears at the root even though its physical location is unrelated.

| **Note** |
| --- |
| This display does not include virtual directories that are subdirectories of the site. The procedure for finding subdirectories is demonstrated below in Example 2. |

```
Alias Physical Root

/Scripts C:\Libraries\Scripts
```

"Alias" represents the names of the virtual directories while "Physical Root" describes their actual locations.

**Example 2**
The following example displays the virtual directories that are subdirectories of the "Scripts" virtual directory on the "Test Downloads" FTP site. The command specifies the "Scripts" virtual directory by using its virtual path. Note that the quotation marks enclose only the path component that requires them.

### iisftpdr /query "Test Download"\Scripts

In response, iisftpdr displays the VBScripts subdirectory of the "Scripts" virtual directory.

```
Alias Physical Root

/VBScripts D:\Scripts\VBScripts
```

**Related Information**

- For information about creating FTP virtual directories, see Creating FTP Virtual Directories Using Iisftpdr.vbs.
- For information about deleting FTP virtual directories, see Deleting FTP Virtual Directories Using Iisftpdr.vbs.

# Managing IIS Configurations Using Scripts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use two command-line scripts, which are stored in *systemroot*\System32, to create and manage Internet Information Services (IIS) configuration files on computers running a member of the Microsoft Windows Server™ 2003 family with IIS 6.0.

Iisback.vbs creates and manages backup copies of the IIS configuration (metabase and schema) of a remote or local computer.

Iiscnfg.vbs imports and exports all, or selected, elements of an IIS metabase on a local or remote computer, or it copies the entire IIS configuration (metabase and schema) to another computer.

This section includes the following information:

- Backing Up IIS Configurations Using Iisback.vbs: Describes how to back up configurations on local or remote computers by using iisback.vbs.
- Restoring IIS Configurations Using Iisback.vbs: Describes how to restore configurations on local or remote computers by using iisback.vbs.
- Deleting IIS Backup Configurations Using Iisback.vbs: Describes how to delete configurations on local or remote computers by using iisback.vbs.
- Listing IIS Backup Configurations Using Iisback.vbs: Describes how to list configurations on local or remote computers by using iisback.vbs.
- Copying IIS Configurations Using Iiscnfg.vbs: Describes how to copy configurations on local or remote computers by using iiscnfg.vbs.
- Importing IIS Configurations Using Iiscnfg.vbs: Describes how to import configurations on local or remote computers by using iiscnfg.vbs.
- Exporting IIS Configurations Using Iiscnfg.vbs: Describes how to export configurations on local or remote computers by using iiscnfg.vbs.
- Saving Configurations to Disk Using Iiscnfg.vbs: Describes how to save configurations to disk on local or remote computers by using iiscnfg.vbs.

# Backing Up IIS Configurations Using Iisback.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisback.vbs, which is stored in *systemroot*\System32, to create and manage backup copies of the IIS configuration (metabase and schema) of a remote or local computer. Administrators can use this script tool to create a backup copy of their IIS configuration, to restore an IIS configuration from a backup copy, and to list and delete backup copies.

Iisback.vbs performs the same backup and restore operations that are available in IIS Manager. You can use either tool to view and manage backup copies.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /**User:MyComputer**\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Overview: Key concepts in understanding the backup and restore procedures.
- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Overview**

Each **/backup** operation creates two files, an .MDx file to store the metabase and an .SCx file to store the schema, where the x is the version number of the backup copy. IIS and iisback.vbs store backup copy files in the **systemroot**\System32\inetsrv\MetaBack directory.

The metabase and schema of an IIS configuration include system-specific and session-specific properties. Do not copy or import the metabase or schema of one IIS server to another IIS server without modification. To copy all or part of a metabase configuration from one system to another, use Copying IIS Configurations Using Iiscnfg.vbs.

To prevent unauthorized use of backup copies, you can use IIS 5.1 or IIS 6.0, and iisback.vbs to encrypt the backup copy with a password. The password encrypts the session key which, in turn, encrypts all of the properties for which the secure attribute is set.

You can read an encrypted backup copy (only the session key and secure properties are encrypted) and you can delete the files. However, you cannot use the backup copy in a restore operation unless you provide the encrypting password. Also, you cannot remove the password encryption from a backup copy or change the encrypting password.

If you do not use password encryption, the session key and secure properties are encrypted with a blank password, which prevents you from reading these values but allows any member of the Administrators group to restore the metabase from the backup copy. IIS 5.1 and IIS 6.0 do not support machine key encryption of backup copies.

**Syntax**

**iisback /backup** *[* **/b** *BackupName] [* **/v** *{Integer* | **HIGHEST_VERSION | NEXT_VERSION** *}] [* **/overwrite** *] [* **/e** *EncryptingPassword] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

**/b** *BackupName*

Specifies the name of the backup copy. The default is *SampleBackup*.

**/v** {*Integer* | **NEXT_VERSION | HIGHEST_VERSION**}

Specifies the version number of the backup copy. NEXT_VERSION is the default.

| ***Integer*** | **Specifies the version number of the backup copy.** |
|---|---|
| NEXT_VERSION | Specifies the highest version number plus 1. NEXT_VERSION isthe default value. |
| HIGHEST_VERSION | Reuses the highest version number. Because this option replacesa backup copy, the **/overwrite** parameter is required whenusing this value. This parameter is only valid when previousversions of the file are stored on the computer. |

**/overwrite**

Permits iisback.vbs to replace an exiting backup copy with the newly created backup copy. Without this parameter, commands used to create a backup copy with the same name and version number as an existing backup copy fail.

**/e** *EncryptingPassword*

Encrypts the backup copy with the specified backup password. You can use any string as the password.

The **/e** parameter creates a secure backup by encrypting the session key and secure properties with the specified password. Administrators must provide the password to use the backup in a restore operation.

If you omit this parameter or omit the password argument, then the backup is encrypted with a blank password. As a result, only the secure properties are encrypted and any administrator can use the backup copy to restore the metabase.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following command creates a new backup copy of the IIS configuration on the local computer. It uses the **/b** parameter to assign the name ReskitBkp to the file.

**iisback /backup /b ReskitBkp**

In response, **iisback** displays the following success message indicating that the backup copy has been created, and that the default version value, NEXT_VERSION, is assigned to the file. Because this is the first backup copy named ReskitBkp, **iisback** assigns the version number 0 to the backup copy.

```
Backup ReskitBkp version NEXT_VERSION has been CREATED.
```

The new backup copy is displayed in the following list of backup copies on the server:

```
Backup Name Version # Date and Time

ReskitBkp 0 1/9/2001 5:00:04 AM
```

The ReskitBkp backup copy consists of two files, ReskitBkp.MD0, which contains the metabase and ReskitBkp.SC0, which contains the schema. Both files are stored in the **systemroot**\System32\inetsrv\MetaBack directory.

**Example 2:**
The following command creates a backup copy of the IIS configuration on the \\SVR01 remote server. It uses the **/s** parameter to provide the name of the remote computer, and the **/u** and **/p** parameters to run the script with the permissions of the user's Administrator account.

It also includes the **/b** parameter to assign the file name Svr01Bkp to the backup copy and the **/v** parameter to specify the

version number (15) of the new backup copy. The command uses the **/e** parameter to encrypt the backup copy with the backup password 7W*48Hv6#.

**iisback /backup /s svr01 /u admin21 /p Rrr3Qv7s /b Svr01Bkp /v 15 /e 7W*48Hv6#**

In response, iisback displays the following success message:

```
Backup Svr01Bkp version 15 has been CREATED.
```

This command created the Svr01Bkp, version 15 backup copy of the \\SVR01 IIS configuration and stored the data in the Svr01Bkp.MD15 and Svr01Bkp.SC15 files in the **systemroot**\System32\inetsvr\MetaBack directory of the remote computer.

Listing IIS Backup Configurations Using Iisback.vbs shows the backup copies that are stored on the remote server. The **backup** command created the final entry in the list.

```
Backup Name Version # Date and Time

AsiaSvr 1 1/8/2001 9:18:25 PM

Svr01Bkp 11 1/3/2001 3:00:42 AM

Svr01Bkp 12 1/4/2001 3:00:12 AM

Svr01Bkp 13 1/5/2001 3:01:00 AM

Svr01Bkp 14 1/6/2001 3:00:00 AM

Svr01Bkp 15 1/9/2001 5:18:04 AM
```

Although it is not evident from the list, the new backup copy is encrypted and you must provide the backup password, 7W*48Hv6#, in order to use the backup in a restore operation.

**Example 3:**
The following command replaces the most recent backup copy of the local IIS configuration with a newer backup. You can use this command format to update the backup copy that you maintain for the server.

The command uses the **/b** parameter to specify the name of the backup copy. It uses the **/v HIGHEST_VERSION** parameter to direct iisback.vbs to replace the latest version of the backup copy named CurrentBackup, instead of creating a new version. It also uses the **/overwrite** parameter, which is required with **/v HIGHEST_VERSION** to permit iisback.vbs to replace an existing backup.

**iisback /backup /b CurrentBackup /v HIGHEST_VERSION /overwrite**

In response, **iisback** displays the following success message:

```
Backup CurrentBackup version 0 has been CREATED.
```

The command created new versions of the CurrentBackup.md0 and CurrentBackup.sc0 files, which replaced files with the same names.

**Related Information**

- For information about restoring IIS configurations, see Restoring IIS Configurations Using Iisback.vbs.
- For information about deleting IIS configurations, see Deleting IIS Backup Configurations Using Iisback.vbs.
- For information about listing IIS configurations, see Listing IIS Backup Configurations Using Iisback.vbs.

# Restoring IIS Configurations Using Iisback.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisback.vbs, which is stored in *systemroot*\System32, to restore the IIS configuration of a server.

Iisback.vbs performs the same backup and restore operations that are available in IIS Manager. You can use either tool to view and manage backup copies.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Overview: Key concepts in understanding the backup and restore procedures.
- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Overview**

Each backup operation creates two files, an .MDx file to store the metabase and an .SCx file to store the schema, where the x is the version number of the backup copy. IIS and iisback.vbs store backup copy files in the *systemroot*\System32\inetsrv\MetaBack directory.

The metabase and schema of an IIS configuration include system-specific and session-specific properties. Do not copy or import the metabase or schema of one IIS server to another IIS server without modification. To copy all or part of a metabase configuration from one system to another, use Copying IIS Configurations Using Iiscnfg.vbs.

To prevent unauthorized use of backup copies, you can use IIS 5.1 or IIS 6.0, and iisback.vbs to encrypt the backup copy with a password. The password encrypts the session key which, in turn, encrypts all of the properties for which the secure attribute is set.

You can read an encrypted backup copy (only the session key and secure properties are encrypted), and you can delete the files. However, you cannot use the backup copy in a restore operation unless you provide the encrypting password. Also, you cannot remove the password encryption from a backup copy nor can you change the encrypting password.

If you do not use password encryption, the session key and secure properties are encrypted with a blank password, which prevents you from reading these values but allows any member of the Administrators group to restore the metabase from the backup copy. IIS 5.1 and IIS 6.0. do not support machine key encryption of backup copies.

| **Note** |
| --- |
| The **/restore** operation is a complex operation that can take several minutes to complete. A brief delay is normal for this operation. During **/restore**, Web pages, FTP sites, and SMTP and NNTP services are temporarily stopped. Service is restored when the operation completes. |

**Syntax**

**iisback /restore/b** *BackupName [* **/v** *{Integer* | **HIGHEST_VERSION** *}] [* **/e** *EncryptionPassword] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**
**/b** *BackupName*

Specifies the name of the backup copy.

**/v** *{Integer* | **HIGHEST_VERSION** *}*

Specifies the version number of the backup copy. HIGHEST_VERSION is the default.

| **Integer** | **Specifies the version number of the backup copy.** |
|---|---|
| HIGHEST_VERSION | Specifies the backup copy with the highest version number. |

**/e** *Encryption Password*

Provides the password of an encrypted backup copy. When the session key in a backup copy is encrypted, you must provide its password to use the copy in the **/restore** operation.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following command restores the IIS configuration from the TestBkp, version 3 backup copy. It uses the **/b** parameter to specify the name of the backup copy and the **/v** parameter to specify the version.

**iisback /restore /b TestBkp /v 3**

In response, **iisback** displays the following success message, indicating that the IIS configuration is restored:

```
Backup TestBkp version 3 has been RESTORED.
```

To verify that the configuration is restored, compare the current configuration, which is stored in *systemroot*\System32\inetsrv\MetaBase.xml and MBSchema.xml, with the backup copy files, TestBkp.MD3 and TestBkp.SC3, in the *systemroot*\System32\inetsrv\MetaBack directory.

**Example 2:**
The following command restores the IIS configuration from the Svr01Bkp, version 15 backup copy. The **/b** parameter is required to specify the backup copy name and the **/e** parameter to specify the backup password, 7W*48Hv6#. This command omits the **/v** parameter and uses the default value, HIGHEST_VERSION.

**iisback /restore /b Svr01Bkp /e 7W*48Hv6#**

In response, **iisback** displays the following success message, indicating that the IIS configuration is restored:

```
Backup Svr01Bkp version HIGHEST_VERSION has been RESTORED.
```

**Related Information**

- For information about backing up IIS configurations, see Backing Up IIS Configurations Using Iisback.vbs.
- For information about deleting IIS configurations, see Deleting IIS Backup Configurations Using Iisback.vbs.
- For information about listing IIS configurations, see Listing IIS Backup Configurations Using Iisback.vbs.

# Deleting IIS Backup Configurations Using Iisback.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisback.vbs, which is stored in *systemroot*\System32, to delete the backup IIS configuration of a local or remote computer.

Deleting backup files moves the files from the *systemroot*\System32\inetsrv\MetaBack directory to the Recycle Bin. It does not affect the files that are in the *systemroot*\System32\inetsrv\MetaBack\History directory.

Iisback.vbs performs the same delete operations that are available in IIS Manager. You can use either tool to view and manage backup copies.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Overview: Key concepts in understanding the backup and restore procedures.
- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Overview**

Each backup operation creates two files, an .MDx file to store the metabase and an .SCx file to store the schema, where the x is the version number of the backup copy. IIS and iisback.vbs store backup copy files in the *systemroot*\System32\inetsrv\MetaBack directory.

The metabase and schema of an IIS configuration include system-specific and session-specific properties. Do not copy or import the metabase or schema of one IIS server to another IIS server without modification. To copy all or part of a metabase configuration from one system to another, use Copying IIS Configurations Using Iiscnfg.vbs.

To prevent unauthorized use of backup copies, you can use IIS 5.1 or IIS 6.0, and iisback.vbs to encrypt the backup copy with a password. The password encrypts the session key which, in turn, encrypts all of the properties for which the secure attribute is set.

You can read an encrypted backup copy (only the session key and secure properties are encrypted), and you can delete the files. However, you cannot use the backup copy in a restore operation unless you provide the encrypting password. Also, you cannot remove the password encryption from a backup copy or change the encrypting password.

If you do not use password encryption, the session key and secure properties are encrypted with a blank password, which prevents you from reading these values but allows any member of the Administrators group to restore the metabase from the backup copy. IIS 5.1 and IIS 6.0. do not support machine key encryption of backup copies.

**Syntax**

**iisback /delete** *[* **/b** *BackupName] [* **/v** *{Integer* | **HIGHEST_VERSION}** *] [* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

**/b** *BackupName*

Required. Specifies the name of the backup copy.

**/v** *{Integer* | **HIGHEST_VERSION** *}*

Required. Specifies the version number of the backup copy.

| Integer | Specifies the version number of the backup copy. |
|---|---|
| HIGHEST_VERSION | Indicates the backup copy with the highest version number. |

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**
**Example 1:**
The following command deletes the Svr01Bkp, version 0 backup copy. It uses the **/b** parameter to specify the Svr01Bkp series and the **/v** parameter to specify the version number. Both parameters are required to identify the backup copy.

**iisback /delete /b Svr01Bkp /v 0**

In response, **iisback** displays the following message, indicating that the backup copy is deleted:

```
Backup Svr01Bkp version 0 has been DELETED.
```

**Example 2**
This example demonstrates the use of a simple batch file command to delete a series of backup copies of the IIS configuration. You can use this command to delete old files from your archive. You can type the command at the command line or save it in a batch file.

```
For /L %%i in (0, 1, 15) do iisback /delete /b Svr01Bkp /v %%i
```

This batch command deletes Svr01Bkp versions 1 through 15 by issuing the **iisback /delete** command fifteen times, once for each version.

For a more practical application, make the backup name and version numbers variable. Or, use a **dir** or Listing IIS Backup Configurations Using Iisback.vbs command to detect the backup with the lowest version number, and then append a command that deletes the oldest backup copy (lowest version number) to each command to create a new backup copy.

**Related Information**

- For information about backing up IIS configurations, see Backing Up IIS Configurations Using Iisback.vbs.
- For information about restoring IIS configurations, see Restoring IIS Configurations Using Iisback.vbs.
- For information about listing IIS configurations, see Listing IIS Backup Configurations Using Iisback.vbs.

# Listing IIS Backup Configurations Using Iisback.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisback.vbs, which is stored in *systemroot*\System32, to list backup copies of the IIS configuration (metabase and schema) of a remote or local computer.

The **/list** operation lists only the backup copies that are stored in the default location, *systemroot*\System32\Inetsrv\MetaBack. If you move the backup copy files to a different directory, the backup copy does not appear in the list.

Iisback.vbs performs the same backup and restore operations that are available in IIS Manager. You can use either tool to view and manage backup copies.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Example: Sample code and an explanation of the results.

**Syntax**

**iisback /list** *[* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Example**

The following command lists the backup copies that are stored on the remote server, Svr16. It uses the **/s** parameter to specify the server name, the **/u** parameter to specify the user's Administrator account, and the **/p** parameter to specify the password for the Administrator account.

**iisback /list /s Svr16 /u Domain06\User1 /p R34*9W@4b**

In response, **iisback** displays all of the backup copies that are on the remote server, Svr16:

```
Backup Name Version # Date and Time

AsiaSvr 0 12/8/2001 9:18:25 PM

AsiaSvr 1 1/14/2001 9:02:33 PM

Svr01Bkp 11 1/3/2001 3:00:42 AM
```

```
Svr01Bkp 12 1/4/2001 3:00:12 AM

Svr01Bkp 13 1/5/2001 3:01:00 AM

Svr01Bkp 14 1/6/2001 3:00:00 AM

Svr01Bkp 15 1/9/2001 5:18:04 AM
```

| **Note** |
| --- |
| **iisback /list** uses Coordinated Universal Time (UTC) for date and time stamps. Also, without parameters, the **iisback /list** command lists all of the backup copies that are stored on the local computer. |

**Related Information**

- For information about backing up IIS configurations, see Backing Up IIS Configurations Using Iisback.vbs.
- For information about restoring IIS configurations, see Restoring IIS Configurations Using Iisback.vbs.
- For information about deleting IIS configurations, see Deleting IIS Backup Configurations Using Iisback.vbs.

# Copying IIS Configurations Using Iiscnfg.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iiscnfg.vbs, which is stored in *systemroot*\System32, to copy the IIS metabase and schema from one computer to another. Both computers must be running a member of the Windows Server 2003 family with IIS 6.0.

By default, iiscnfg.vbs copies the metabase of the local computer to the target computer that is specified by the **/ts** parameter. However, you can use the **/s** parameter to copy the metabase of another computer to the target computer.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer\Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Overview: Key concepts in understanding the copy procedures.
- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Overview**

The **/copy** operation uses iisback.vbs to make a backup copy of the source metabase and the schema. Afterward, it copies the backup copy files (.MDx and .SCx) to the target computer and uses iisback.vbs to replace the metabase and schema of the target computer with the backup copy.

Although you can do these operations manually, **/copy** provides a convenient one-step method to replicate an IIS configuration. The **/copy** operation replaces the features of iissync.exe, a tool formerly included in Windows.

The **/copy** operation does not copy the server content, such as Web pages and FTP files, that is associated with the IIS configuration. Instead, **/copy** changes the computer-specific and system-specific properties in the metabase so that they are valid on the target computer. However, it does not adjust the directory or file paths. As a result, paths in the configuration might not be valid on the target computer.

**Syntax**

**iiscnfg /copy /ts** *TargetComputer* **/tu** *TargetUser* **/tp** *TargetPassword* **[/s** *Computer* **[/u** *[Domain\]User* **[/p** *Password]]]*

**Parameters**

**/ts** *TargetComputer*

Required. Copies the IIS metabase to the specified computer. Type the IP address or computer name without backslashes.

**/tu** *TargetUser*

Required (even when the current user of the local computer is an administrator of the target computer). Specifies the account used to connect to the target computer. Type a user account in *[Domain]\User* format. This account must be a member of the Administrators group on the target computer.

**/tp** *TargetPassword*

Required. Specifies the password for the account that is specified by the **/tu** parameter.

**/s** *Computer*

Copies the IIS metabase of the specified remote computer to the target computer. This parameter runs the script on the remote computer. Type the name or IP address of a remote computer without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Required only when the current user is not an administrator of the remote computer. Runs the script with the permissions of the specified user account. The default is the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**
**Example 1:**
The following example copies the IIS configuration of the local computer to the Svr01 server. The command uses the **/ts** parameter to identify the target server, and the **/tu** and **/tp** parameters to provide the user's administrator account and password on the Svr01 server.

**iiscnfg /copy /ts SVR01 /tu SVR01\Admin06 /tp P@ssWord6**

In response, **iiscnfg** displays a list of the steps it takes to copy the IIS configuration.

```
Microsoft (R) Windows Script Host Version 5.6

Copyright (C) Microsoft Corporation 1996-2001. All rights reserved.

Backing up server 127.0.0.1

cmd /c iisback /backup /b iisreplback /overwrite

Backup complete.

Mapping local drive D: to admin share on server 127.0.0.1

Mapping local drive E: to admin share on server Svr01

Copying backup files...

cmd /c copy /Y D:\System32\inetsrv\metaback\iisreplback.* E:\System32\inetsrv\metaback

Unmapping local drive D:

Restoring on server Svr01

cmd /c iisback /restore /s svr01 /u SVR01\Admin06 /p P@ssWord6 /b iisreplback

Restore complete.

Unmapping local drive E:
```

This display shows that **iiscnfg** issues commands to Backing Up IIS Configurations Using Iisback.vbs, first to create a backup copy of the metabase and schema and then to restore (or paste) it to the target computer. To do so, it must copy the backup copy files (iisreplback.*) to the target computer.

The IP address, 127.0.0.1, is the loopback address and refers generically to the local computer. This address is used only when copying the configuration of the local computer.

**Example 2:**
The following example copies the IIS configuration of a remote computer, Svr06, to the MainBkp server. The command uses the **/s** parameter to identify the remote computer and the **/u** parameter to provide the user's administrator password on the Svr06 computer. It also uses the **/ts** parameter to identify the target server, and the **/tu** and **/tp** parameters to provide the user's administrator account and password on the MainBkp server.

Please note that although the **/p** parameter is optional, the **/tp** parameter is required. If you omit it, the command fails.

**iiscnfg /copy /s Svr06 /u Domain01\Admin01 /ts MainBkp /tu MainBkp\Admin01 /tp P@ssWord6**

In response, **iiscnfg** prompts you for the password for the account specified by the **/u** parameter, which specifies the permissions under which the script runs, and then it displays a list of the steps it takes to copy the IIS configuration. The phrase "Restore complete" indicates that the **/copy** operation was successful.

```
Microsoft (R) Windows Script Host Version 5.6
```

```
Copyright (C) Microsoft Corporation 1996-2001. All rights reserved.

Enter the Password:

Processing...

Backing up server Svr06

cmd /c iisback /backup /s Svr06 /u Domain01\Administrator /p B##jdt! /b iisreplback /overwrite

Backup complete.

Mapping local drive G: to admin share on server Svr06

Mapping local drive I: to admin share on server MainBkp

Copying backup files...

cmd /c copy /Y G:\System32\inetsrv\metaback\iisreplback.* I:\System32\inetsrv\metaback

Unmapping local drive G:

Restoring on server MainBkp

cmd /c iisback /restore /s MainBkp /u MainBkp\Admin01 /p P@ssWord6 /b iisreplback

Restore complete.

Unmapping local drive I:
```

**Related Information**

- For information about importing IIS configurations, see Importing IIS Configurations Using Iiscnfg.vbs.
- For information about exporting IIS configurations, see Exporting IIS Configurations Using Iiscnfg.vbs.
- For information about saving IIS configurations, see Saving Configurations to Disk Using Iiscnfg.vbs.

# Importing IIS Configurations Using Iiscnfg.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iiscnfg.vbs, which is stored in *systemroot*\System32, to add configuration settings from an Extensible Markup Language (XML) metabase export file to the metabase of an IIS server.

The computer issuing the command must be running Microsoft Windows XP or a member of the Windows Server 2003 family. The computer that the command affects must be running a member of the Windows Server 2003 family with IIS 6.0.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Overview: Key concepts in understanding the copy procedures.
- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Overview**

The metabase of an IIS server includes system-specific and computer-specific properties. Do not export the configuration of one IIS server to another IIS server without modification. Instead, use Backing Up IIS Configurations Using Iisback.vbs or the Importing IIS Configurations Using Iiscnfg.vbs**/copy** operation, which replaces system-specific and computer-specific properties so that the configuration is valid on the destination computer.

Changes made by using iiscnfg.vbs are effective immediately, but they aren't reflected in the MetaBase.xml file until the file is updated (every five minutes or every 50 changes). Use **iiscnfg /save** to review your changes immediately.

**Syntax**

**iiscnfg /import /f** *[Path\]FileName.xml* **/sp** *SourcePath* **/dp** *DestiniationPath[* **/d** *EncryptingPassword] [* **/inherited** *] [* **/children** *] [* **/merge** *][* **/s** *Computer [* **/u** *[Domain\]User [* **/p** *Password]]]*

**Parameters**

**/f** *[Path\]FileName.* **xml**

Required. Specifies the path and file name of the XML source file.

**/sp** *SourcePath*

Required. Specifies the location in the XML file of the keys being imported. Enter the metabase path as it appears in the XML file. To import all of the keys in the source file, enter the metabase path of the root key.

**/dp** *DestinationPath*

Required. Specifies the metabase path into which the keys are placed. Enter the metabase path as it appears in the XML file.

| **Note** |
| --- |

The imported key must be of the type expected in its new position in the metabase path. If the key is not of the expected type, the configuration will not operate properly. For example, do not import the event logging configuration (/LM/Logging) to a Web site configuration path, such as /LM/W3SVC/3. If you specify a path that already exists in the metabase, the imported keys overwrite and replace the keys currently in that metabase path, without warning.

**/d** *EncryptingPassword*

Required only when the configuration file is encrypted. Specifies the password that was used to encrypt the configuration data.

**/inherited**

Imports the inherited properties of the keys. This parameter is only effective when the source file includes inherited properties.

**/children**

Recursively imports the subkeys of the specified key.

**/merge**

Combines keys in the XML file with the existing metabase keys. Without this parameter, the keys in the XML file replace existing keys in the same metabase path. The **/merge** parameter adds keys to the metabase that appear only in the XML file, retains keys in the metabase that are not in the XML file, and applies the values in the XML file when a key appears in both the XML file and the metabase.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**
**Example 1:**
The following command imports the NewSite Web site configuration from C:\Config.xml file to the metabase. The command uses the **/f** parameter to specify the name and location of the source file. Next, it uses the **/sp** parameter to indicate that the configuration is in the /lm/w3svc/5/Root/NewSite path in the Config.xml file. Finally, it uses the **/dp** parameter to specify that the keys in that path should be placed in the /lm/w3svc/1/Root/401KApp location in the metabase.

**iiscnfg.vbs /import /f c:\config.xml /sp /lm/w3svc/5/Root/NewSite /dp /lm/w3svc/1/Root/401KApp**

In response, **iiscnfg** displays the following success message:

```
Configuration imported from /lm/w3svc/5/Root/NewSite in c:\config.xml to /lm/w3svc/1/Root/401KApp in the
Metabase.
```

As a result of this command, the NewSite configuration now appears in IIS as the first Web site, with metabase path W3SVC/1.

**Example 2:**
The following two-step example imports a custom logging configuration from C:\IISLogging.xml to the IIS metabase of a remote computer. You can use this method, or an extension of this method, to propagate a configuration to a group of servers.

The first command copies the IIsLogging.xml file to the remote server, Svr126.

**copy IIsLogging.xml \\Svr126\c$**

The second command imports IIsLogging.xml to the IIS configuration of Svr126. It uses the **/f** parameter to identify the source file, the **/sp** parameter to specify the /LM/Logging/Custom Logging metabase path to import, and the **/dp** parameter to place the imported keys in the same location in the target configuration, /LM/Logging/Custom Logging. (Use quotation marks to enclose path elements that include spaces.) Next, it uses the **/children** parameter to import all of the subkeys of the Logging

key. Finally, to act on Svr126, it uses the **/s** parameter to specify the server, and the **/u** and **/p** parameters to run the script with the permissions of the user's administrator account on the remote computer.

**iiscnfg /import /f C:\IISLogging.xml /sp "/lm/logging/custom logging" /dp "/lm/logging/custom logging" /children /s Svr126 /u Svr126\Administrator /p p@ssword1#**

In response, **iiscnfg** displays the following success message:

```
Connecting to server ... Done.

Configuration imported from /lm/logging/custom logging in file C:\IISLogging.xml to

/lm/logging/custom logging in the Metabase.
```

As a result, the /LM/Logging/Custom Logging key and its subkeys have been copied from the local metabase to the metabase of Svr126.

**Example 3:**
This example demonstrates how the **/merge** parameter works by comparing a standard **import** operation with an **import** operation that uses the **/merge** parameter. Both of the commands in this example update the properties of the Investments virtual directory of a Web site by using the same import file, Invest_vdir.xml.

The following sample shows the Investments metabase key before any import operations are run. (This sample is taken from a backup copy of the metabase created by using iisback.vbs).

**Metabase before import**

```
<IIsWebVirtualDir Location="/LM/W3SVC/1509060625/root/Investments"

AccessFlags="AccessRead | AccessScript"

AppFriendlyName="Investments"

AppIsolated="2"

AppRoot="/LM/W3SVC/1509060625/Root/Investments"

DirBrowseFlags="DirBrowseShowDate | DirBrowseShowTime | DirBrowseShowSize | DirBrowseShowExtension |
DirBrowseShowLongDate | EnableDefaultDoc"

Path="E:\public\Download"

>

</IIsWebVirtualDir>
```

The following sample shows the content of the Invest_vdir.xml file that will be used in the **/import** operations. (This sample was created by exporting the Investments metabase key and then editing the export file.) Please note that the following properties of the Investments key are changed in the XML file:

- The value of the **AccessFlags** property is changed.
- The **AuthFlags** property is added.
- The **DirBrowseFlags** property is omitted.

**Import File Content**

```
<IIsWebVirtualDir Location="/LM/W3SVC/1509060625/root/Investments"

AccessFlags="AccessRead"

AuthFlags="AuthAnonymous | AuthNTLM"

AppFriendlyName="Investments"

AppIsolated="2"

AppRoot="/LM/W3SVC/1509060625/Root/Investments"

Path="E:\public\Download"

>

</IIsWebVirtualDir>
```

Command 1: Standard Import

The following command imports the Invest_vdir.xml file to the metabase without using the **/merge** parameter. The command directs IIS to replace the Investments key in the metabase with the Investments key in the XML file. The command uses the **/f** parameter to specify the import file, and it uses the same metabase path in the source path (**/sp**) and destination path (**/dp**) parameters.

**iiscnfg /import /f e:\iistest\Invest_vdir.xml /sp /LM/W3SVC/1509060625/root/Investments /dp /LM/W3SVC/1509060625/root/Investments**

As a result of this command, the Investments key in the metabase is now identical to the Investments key in the Invest_vdir.xml import file.

Command 2: Import and merge

The following command merges the Investments key in the Invest_vdir.xml file to the metabase. The **/merge** parameter directs IIS to add new properties from the import file to the metabase, to leave the metabase properties that are not in the import file unchanged, and to apply the import file value when a property appears in both the import file and in the metabase. This command is identical to Command 1, except that the **/merge** parameter is added.

**iiscnfg /import /f e:\iistest\Invest_vdir.xml /sp /LM/W3SVC/1509060625/root/Investments /dp /LM/W3SVC/1509060625/root/Investments /merge**

As a result of this command, the content of the Investments key in the metabase is a combination of its sources, as shown in the following sample. Note that the following properties of the Investments key that are in the metabase are changed.

- The value of the **AccessFlags** property, which appeared in the metabase and the XML file, is changed. The value in the XML file now appears in the metabase.
- The **AuthFlags** property, which appeared only in the XML file, is added.
- The **DirBrowseFlags** property, which appeared only in the metabase, is unchanged.

```
<IIsWebVirtualDir Location="/LM/W3SVC/1509060625/root/Investments"

AccessFlags="AccessRead"

AppFriendlyName="Investments"

AppIsolated="2"

AppRoot="/LM/W3SVC/1509060625/Root/Investments"

AuthFlags="AuthAnonymous | AuthNTLM"

DirBrowseFlags="DirBrowseShowDate | DirBrowseShowTime | DirBrowseShowSize | DirBrowseShowExtension |
DirBrowseShowLongDate | EnableDefaultDoc"

Path="E:\public\Download"

>

</IIsWebVirtualDir>
```

**Related Information**

- For information about copying IIS configurations, see Copying IIS Configurations Using Iiscnfg.vbs.
- For information about exporting IIS configurations, see Exporting IIS Configurations Using Iiscnfg.vbs.
- For information about saving IIS configurations, see Saving Configurations to Disk Using Iiscnfg.vbs.

# Exporting IIS Configurations Using Iiscnfg.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iiscnfg.vbs, which is stored in *systemroot*\System32, to copy all or part of an IIS metabase to an XML file, in an encrypted or unencrypted format. The XML file can be used in an **/import** operation to copy all or part of the metabase to another IIS configuration.

The computer issuing the command must be running Windows XP or a member of the Windows Server 2003 family. The computer that the command affects must be running a member of the Windows Server 2003 family with IIS 6.0.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Overview: Key concepts in understanding the export procedures.
- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Overview**

The metabase of an IIS server includes system-specific and computer-specific properties. Do not export the configuration of one IIS server to another IIS server without modification. Instead, use Backing Up IIS Configurations Using Iisback.vbs or the Exporting IIS Configurations Using Iiscnfg.vbs**/copy** operation, both of which replace system-specific and computer-specific properties so that the configuration is valid on the destination computer.

Iiscnfg.vbs provides an option to include inherited properties in the export file.

To prevent unauthorized use of export files, iiscnfg.vbs allows you to encrypt the export file with a password. The password encrypts the session key which, in turn, encrypts all of the properties for which the secure attribute is set.

You can read an encrypted export file (only the session key and secure properties are encrypted) and you can delete the file. However, you cannot use the export file in an **/import** operation unless you provide the encrypting password. Also, you cannot remove the password encryption from an export file or change the encrypting password.

The **/copy** operation changes the computer-specific and system-specific properties in the metabase so that they are valid on the target computer. However, it does not adjust the directory or file paths. As a result, paths in the configuration might not be valid on the target computer.

**Syntax**

**iiscnfg /export/f** *[Path\]FileName.***xml/sp** *SourcePath [***/d** *EncryptingPassword] [***/inherited***] [***/children***] [***/s** *Computer [***/u** *[Domain\]User [***/p** *Password]]]*

**Parameters**

**/f** *[Path\]FileName.* **xml**

Required. Specifies a unique name and location for the export file. If you specify an existing file or if the directories in the path you specify do not exist, the command fails. Also, the file path must be local. When exporting a file from a remote computer, the **/f** parameter refers to a path on the remote computer.

**/sp** *SourcePath*

Required. Specifies the metabase keys to export. Enter a metabase path, including /lm/ *root*, to indicate the selected keys.

**/d** *EncryptingPassword*

Encrypts the export file with the specified password. You can use any string as the password. The **/d** parameter creates a secure export file by encrypting the session key and secure properties with the specified password. You must provide the password to use the file in an **/import** operation.

If you omit this parameter or omit the password argument, then the export file is encrypted with a blank password. As a result, only the secure properties are encrypted and any administrator can use the export file to import a metabase configuration.

**/inherited**

Adds the inherited properties of the exported keys to the export file. Without this parameter, the export file does not include properties inherited from higher-level keys. If you select this option, iiscnfg.vbs adds the inherited properties of the exported keys to an IIsInheritedProperties key in the export file. When you import keys from this file, you can elect to import the inherited properties with the keys.

| **Note** |
| --- |
| When inherited properties are imported, the imported keys retain the properties they inherited in the exporting configuration. If inherited properties are not imported, the keys inherit the properties of the importing configuration. |

**/children**

Recursively adds the subkeys of the specified key to the export file.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**
**Example 1:**
The following example exports the configuration of the Finance Web site to the C:\Finance\Finance.xml file. The command uses the **/sp** parameter to specify the metabase path of the Web site, which is /lm/W3SVC/2, and the **/children** parameter to export all of the subkeys of the configuration. It also uses the **/inherited** parameter to include inherited properties in the export file. This ensures that the Web site configuration retains these properties even if it is imported into a different system.

**iiscnfg /export /f c:\finance\finance.xml /sp /lm/w3svc/2 /children /inherited**

In response, **iiscnfg** displays the following success message:

```
Configuration exported from /lm/w3svc/2 to file c:\finance\finance.xml
```

The resulting file includes all of the keys and subkeys within the /lm/w3svc/2 path. It also includes an IisInheritedProperties key that contains the values of all of the inherited properties.

**Example 2:**
The following example exports the IIS configuration of the local computer to the D:\IIStest\Config.xml file. It uses the **/f** parameter to specify the file name and location and the **/sp** parameter to specify the root key (**/**) of the metabase. The **/children** parameter adds all of the subkeys of the root key recursively to the export file. (Without the /children parameter, only the root key is exported.) Finally, the command uses the **/d** parameter and a password to encrypt the export file.

**iiscnfg /export /f d:\iistest\config.xml /sp / /children /d p@sswoRd61**

In response, **iiscnfg** displays the following success message:

```
Configuration / has been exported to d:\iistest\config.xml
```

Even a complete export of the configuration is not identical to the MetaBase.xml file for the system. The session key, access control lists (ACLs), and passwords differ.

You cannot tell by looking at the export file that it is encrypted. However, if you compare an encrypted and unencrypted version of the same file, you can see that the values of the secure properties are changed in the encrypted version. Therefore, the file cannot be used in an import command without the password.

**Related Information**

- For information about copying IIS configurations, see Copying IIS Configurations Using Iiscnfg.vbs.
- For information about importing IIS configurations, see Importing IIS Configurations Using Iiscnfg.vbs.
- For information about saving IIS configurations, see Saving Configurations to Disk Using Iiscnfg.vbs.

# Saving Configurations to Disk Using Iiscnfg.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iiscnfg.vbs, which is stored in *systemroot*\System32, to immediately save configuration changes to disk. Administrators can use this script tool to make changes immediately visible in the metabase.xml.

Iiscnfg.vbs performs the same backup and restore operations that are available in IIS Manager. You can use either tool to view and manage backup copies.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Overview: Key concepts in understanding the backup and restore procedures.
- Syntax: The order in which you type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Overview**

Any changes that are made to the metabase, whether by using scripts or IIS Manager, are written to memory first and to disk at the next scheduled update, which might be several minutes later. As a result, an administrator might make metabase changes and then open metabase.xml only to find the changes aren't reflected. The **/save** command in iiscnfg.vbs resolves the problem by saving changes to the disk immediately.

**Syntax**

**iiscnfg /save** *[* **/s** *Computer [* **/u** *[Domain\]User* **/p** *Password]]*

**Parameters**

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. The default is the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter.

**Examples**

**Example 1:**

The following command saves the IIS configuration to disk on the local computer:

**iiscnfg /save**

In response, **iiscnfg** displays the following message:

```
Configuration saved
```

**Example 2:**
The following command saves the IIS configuration to disk on the local computer:

**iiscnfg /save /s SourceServer /u Administrator /p Kj30W**

In response, **iiscnfg** displays the following message:

```
Configuration saved
```

**Related Information**

- For information about copying IIS configurations, see Copying IIS Configurations Using Iiscnfg.vbs.
- For information about importing IIS configurations, see Importing IIS Configurations Using Iiscnfg.vbs.
- For information about exporting IIS configurations, see Exporting IIS Configurations Using Iiscnfg.vbs.

# Managing Applications and Web Service Extensions Using Scripts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use two command-line scripts, which are stored in *systemroot*\system32, to configure and manage applications, Web service extensions, and individual files on computers running a member of the Microsoft Windows Server™ 2003 family with Internet Information Services (IIS) 6.0.

Iisext.vbs can enable and list applications; add and remove application dependencies; enable, disable, and list Web service extensions; and add, remove, enable, disable, and list individual files.

Iisapp.vbs can list applications that are running.

This section includes the following information:

- Enabling Applications Using Iisext.vbs: Describes how to enable applications on local or remote computers using iisext.vbs.
- Listing Applications Using Iisext.vbs: Describes how to list applications on local or remote computers using iisext.vbs.
- Adding Application Dependencies Using Iisext.vbs: Describes how to add application dependencies on local or remote computers using iisext.vbs.
- Removing Application Dependencies Using Iisext.vbs: Describes how to remove application dependencies on local or remote computers using iisext.vbs.
- Enabling Web Service Extensions Using Iisext.vbs: Describes how to enable Web service extensions on local or remote computers using iisext.vbs.
- Disabling Web Service Extensions Using Iisext.vbs: Describes how to disable Web service extensions on local or remote computers using iisext.vbs.
- Listing Web Service Extensions Using Iisext.vbs: Describes how to list Web service extensions on local or remote computers using iisext.vbs.
- Adding Web Service Extension Files Using Iisext.vbs: Describes how to add Web service extension files on local or remote computers using iisext.vbs.
- Enabling Web Service Extension Files Using Iisext.vbs: Describes how to enable Web service extension files on local or remote computers using iisext.vbs.
- Removing Web Service Extension Files Using Iisext.vbs: Describes how to remove Web service extension files on local or remote computers using iisext.vbs.
- Listing Web Service Extension Files Using Iisext.vbs: Describes how to list Web service extension files on local or remote computers using iisext.vbs.
- Listing Running Web Applications Using Iisapp.vbs: Describes how to list applications currently running on local or remote computers using iisapp.vbs.

# Enabling Applications Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to enable applications on computers running a member of the Windows Server 2003 family with IIS 6.0. The application must have at least one preexisting dependency in the ApplicationDependencies Metabase Property.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisext /EnApp** *Application [* **/s** *Computer [* **/u** *[Domain\]User [* **/p** *Password]]]*

**Parameters**
*Application*

Required. Specifies the name of the application to be enabled.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**
**Example 1:**
The following command enables the application "Payroll" on the local computer:

**iisext /EnApp Payroll**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Enabling application complete.
```

**Example 2:**
The following command enables the application "TravelPlans" on a remote computer:

**iisext /EnApp TravelPlans /s RemoteComputer /u Manager /p Pa$$Werd**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Enabling application complete.
```

**Related Information**

- For information about listing applications, see Listing Applications Using Iisext.vbs.

# Listing Applications Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to list applications that can be enabled on computers running a member of the Windows Server 2003 family with IIS 6.0. These applications are listed in the ApplicationDependencies Metabase Property.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Example: Sample code and an explanation of the results.

**Syntax**

**iisext /ListApp** *[* **/s** *Computer [* **/u** *[Domain* **** *]User [* **/p** *Password]]]*

**Parameters**

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain* **** *]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Example**

The following command lists all of the applications that are in ApplicationDependencies Metabase Property on the local computer:

**iisext /ListApp**

In response, **iisext** displays all of the applications that are on the local computer.

```
Connecting to server ...Done

Active Server Pages

Internet Data Connector

Server Side Includes

WebDAV

Payroll
```

TravelPlans

# Adding Application Dependencies Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to add dependencies between applications and one or more Web service extensions on computers running a member of the Windows Server 2003 family with IIS 6.0. If the application does not exist in the ApplicationDependencies Metabase Property, **iisext** will create it.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisext /AddDep** *Application ID [ID] [* **/s** *Computer [* **/u** *[Domain\]User [* **/p** *Password]]]*

**Parameters**

*Application*

Required. Specifies the name of the application that is dependent on a Web service extension.

*ID*

Required. Specifies the ID of the Web service extension upon which the application is dependent. Additional dependencies are optional.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**

**Example 1:**

The following command makes the application "Payroll" dependent on the Web service extension "timesheet" on the local computer. In this example, the files that are required for the Web service extension "timesheet" were previously added to and enabled on the local computer.

**iisext /AddDep Payroll timesheet**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Adding dependency complete.
```

**Example 2:**
The following command makes the application "TravelPlans" dependent on the Web service extensions "FlightSchedules" and "HotelRooms" on a remote computer. In this example, the files that are required for the Web service extensions "FlightSchedules" and "HotelRooms" were previously added to and enabled on the remote computer.

**iisext /AddDep TravelPlans FlightSchedules HotelRooms /s RemoteComputer /u Manager /p Pa$$Werd**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Adding dependency complete.
```

**Related Information**

- For information about removing application dependencies, see Removing Application Dependencies Using Iisext.vbs.

# Removing Application Dependencies Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to remove dependencies between applications and one or more Web service extensions on computers running a member of the Windows Server 2003 family with IIS 6.0.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisext /RmDep** *Application ID [ID] [* **/s** *Computer [* **/u** *[Domain* **** *]User [* **/p** *Password]]]*

**Parameters**

*Application*

Required. Specifies the name of the application that is dependent on a Web service extension.

*ID*

Required. Specifies the ID of the Web service extension upon which the application is dependent. Additional dependencies are optional.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain* **** *]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**

**Example 1:**

The following command removes the dependency of the application "Payroll" on the Web service extension "timesheet" on the local computer:

**iisext /RmDep Payroll timesheet**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Removing dependency complete.
```

**Example 2:**
The following command removes the dependency of the application "TravelPlans" on the Web service extensions "FlightSchedules" and "HotelRooms" on a remote computer:

**iisext /RmDep TravelPlans FlightSchedules HotelRooms /s RemoteComputer /u Manager /p Pa$$Werd**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Removing dependency complete.
```

**Related Information**

- For information about adding application dependencies, see Adding Application Dependencies Using Iisext.vbs.

# Enabling Web Service Extensions Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to enable a Web service extension. The Web service extension must have at least one file associated with it in the WebSvcExtRestrictionList Metabase Property. The command can be executed on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisext /EnExt** *ID [* **/s** *Computer [* **/u** *[Domain\]User [* **/p** *Password]]]*

**Parameters**

*ID*

Required. Specifies the ID of the Web service extension that is going to be enabled.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**

**Example 1:**

The following command enables the Web service extension "Payroll" on the local computer:

**iisext /EnExt Payroll**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Enabling extension complete.
```

**Example 2:**

The following command enables the Web service extension "Inventory" on a remote computer:

**iisext /EnExt Inventory /s RemoteComputer /u Manager /p Pa$$Werd**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Enabling extension complete.
```

**Related Information**

- For information about disabling Web service extensions, see Disabling Web Service Extensions Using Iisext.vbs.
- For information about listing Web service extensions, see Listing Web Service Extensions Using Iisext.vbs.

# Disabling Web Service Extensions Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to disable a Web service extension by disabling all of the files with a specific Web service extension ID. The command can be executed on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisext /DisExt** *ID [* **/s** *Computer [* **/u** *[Domain\]User [* **/p** *Password]]]*

**Parameters**
*ID*

Required. Specifies the ID of the Web service extension that is going to be disabled.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**
**Example 1:**
The following command disables the Web service extension "FlightTimes" on the local computer:

**iisext /DisExt FlightTimes**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Disabling extension complete.
```

**Example 2:**
The following command disables the Web service extension "PassengerList" on a remote computer:

**iisext /DisExt PassengerList /s RemoteComputer /u Manager /p Pa$$Werd**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Disabling extension complete.
```

**Related Information**

- For information about enabling Web service extensions, see Enabling Web Service Extensions Using Iisext.vbs.
- For information about listing Web service extensions, see Listing Web Service Extensions Using Iisext.vbs.

# Listing Web Service Extensions Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to list Web service extension IDs for all of the defined Web service extensions. The extensions must reside on a computer running a member of the Windows Server 2003 family with IIS 6.0.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Example: Sample code and an explanation of the results.

**Syntax**

**iisext /ListExt** *[* **/s** *Computer [* **/u** *[Domain* **** *]User [* **/p** *Password]]]*

**Parameters**

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain* **** *]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Example**

The following command lists all of the Web service extension IDs that are in the **WebSvcExtRestrictionList** property of the metabase on the local computer:

**iisext /ListExt**

In response, **iisext** displays all of the Web service extension IDs that are on the local computer.

```
Connecting to server ...Done

SSINC

Payroll

ASP

Inventory

ASP.NET
```

```
WEBDAV

HTTPODBC
```

**Related Information**

- For information about enabling Web service extensions, see Enabling Web Service Extensions Using Iisext.vbs.
- For information about disabling Web service extensions, see Disabling Web Service Extensions Using Iisext.vbs.

# Adding Web Service Extension Files Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to add a single file to the WebSvcExtRestrictionList Metabase Property. The command can be executed on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. Web service extension files can be either ISAPI extensions or CGI files.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisext /AddFile** *Drive:\Path\FileName Access {0|1} ID Deletable {0|1} ShortDesc*[**/s** *computer* [**/u** *[domain\]user***/p** *password]]*

**Parameters**

*Drive:\Path\FileName*

Required. Specifies the file name and the path of the Web service extension file that is going to be added.

*Access*

Required. Numerically designates whether the file is disabled (set to 0) or enabled (set to 1) after it is added.

*ID*

Required. Specifies the ID of the Web service extension that is associated with the file.

*Deletable*

Required. Numerically designates whether the file can be deleted from the Web service extension restriction list by using **iisext** or IIS Manager. Set to 0 to prevent the file from being deleted, or set to 1 to allow the file to be deleted.

*ShortDesc*

Required. A short description of the Web service extension that is associated with the file. This text can be localized, and it should be the same for all of the files that are associated with the same Web service extension.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**/?**

Displays Help at the command prompt.

**Examples**
**Example 1:**
The following command adds the file "Timesheet.exe" to the **WebSvcExtRestrictionList** property of the metabase on the local computer. The *Access* flag is set to "0" to disable the file on startup, the *ID* flag identifies the associated Web service extension as "Payroll", and the *Deletable* flag is set to "1" to allow the file to be deleted from the list of Web service extensions.

**iisext /AddFile c:\Source\Timesheet.exe 0 Payroll 1 Payroll**

In response, **iisext** displays the following:

```
Connecting to server …Done.

Adding extension file complete.
```

**Example 2:**
The following command adds the file "Payscale.dll" to the **WebSvcExtRestrictionList** property of the metabase on a remote computer. The *Access* flag is set to "1" to enable the file on startup, the *ID* flag identifies the associated Web service extension as "Payroll", and the *Deletable* flag is set to "0" to prevent the file from being deleted from the list of Web service extensions.

**iisext /AddFile c:\Source\Payscale.dll 1 Payroll 0 Payroll /s RemoteComputer /u Manager /p Pa$$Werd**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Adding extension file complete.
```

**Related Information**

- For information about enabling Web service extension files, see Enabling Web Service Extension Files Using Iisext.vbs.
- For information about disabling Web service extension files, see Disabling Web Service Extensions Using Iisext.vbs.
- For information about removing Web service extension files, see Removing Web Service Extension Files Using Iisext.vbs.
- For information about listing Web service extension files, see Listing Web Service Extension Files Using Iisext.vbs.

# Enabling Web Service Extension Files Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to enable a single file that is in the WebSvcExtRestrictionList Metabase Property property of the metabase. The command can be executed on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. Web service extension files can be either ISAPI extensions or CGI files.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisext /EnFile** *Drive:\Path\FileName [* **/s** *computer [* **/u** *[domain\]user* **/p** *password]]*

**Parameters**
*Drive:\Path\FileName*

Required. Specifies the file name and the path of the Web service extension file that is going to be enabled.

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**
**Example 1:**
The following command enables the file "Timesheet.exe" in the WebSvcExtRestrictionList Metabase Property on the local computer:

**iisext /EnFile c:\Source\Timesheet.exe**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Enabling extension file complete.
```

**Example 2:**
The following command enables the file "Taxes.dll" in the WebSvcExtRestrictionList Metabase Property on a remote computer:

**iisext /EnFile c:\Source\Taxes.dll /s RemoteComputer /u Manager /p Pa$$Werd**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Enabling extension file complete.
```

**Related Information**

- For information about adding Web service extension files, see Adding Web Service Extension Files Using Iisext.vbs.
- For information about disabling Web service extension files, see Disabling Web Service Extensions Using Iisext.vbs.
- For information about removing Web service extension files, see Removing Web Service Extension Files Using Iisext.vbs.
- For information about listing Web service extension files, see Listing Web Service Extension Files Using Iisext.vbs.

# Removing Web Service Extension Files Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to remove a single deleteable file from the WebSvcExtRestrictionList Metabase Property. The command can be executed on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. Web service extension files can be either ISAPI extensions or CGI files.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters).

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**

**iisext /RmFile** *Drive:\Path\FileName [* **/s** *Computer [* **/u** *[Domain\]User [* **/p** *Password]]]*

**Parameters**

*Drive\Path\FileName*

Required. Specifies the file name and the path of the Web service extension file that is going to be removed from the Web service extension restriction list (**WebSvcExtRestrictionList**).

**Note**

If the file has been flagged as undeletable, **iisext** returns the error message, "Error trying to delete extension. Make sure it exists and is deletable."

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain\]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is specified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Examples**

**Example 1:**

The following command removes the file "Accounting.dll" from the **WebSvcExtRestrictionList** property of the metabase on the local computer:

**iisext /RmFile c:\Source\Accounting.dll**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Removing extension file complete.
```

**Example 2:**
The following command removes the file "Taxes.exe" from the **WebSvcExtRestrictionList** property of the metabase on a remote computer:

**iisext /RmFile c:\Source\Taxes.exe /s RemoteComputer /u Manager /p Pa$$Werd**

In response, **iisext** displays the following:

```
Connecting to server ...Done.

Removing extension file complete.
```

**Related Information**

- For information about adding Web service extension files, see Adding Web Service Extension Files Using Iisext.vbs.
- For information about enabling Web service extension files, see Enabling Web Service Extension Files Using Iisext.vbs.
- For information about disabling Web service extension files, see Disabling Web Service Extensions Using Iisext.vbs.
- For information about listing Web service extension files, see Listing Web Service Extension Files Using Iisext.vbs.

# Listing Web Service Extension Files Using Iisext.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisext.vbs, which is stored in *systemroot*\system32, to list the Web service extension files that are in the WebSvcExtRestrictionList Metabase Property. The command can be executed on local or remote computers running a member of the Windows Server 2003 family with IIS 6.0. Web service extension files can be either ISAPI extensions or CGI files.

Iisext.vbs performs the same operations that are available in IIS Manager. You can use either tool to administer IIS Web sites.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:MyComputer\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Example: Sample code and an explanation of the results.

**Syntax**

**iisext /ListFile** *[* **/s** *Computer [* **/u** *[Domain\]User [* **/p** *Password]]]*

**Parameters**

**/s** *Computer*

Runs the script on the specified remote computer. Type the computer name or IP address without backslashes. By default, the script runs on the local computer.

**/u** *[Domain* **** *]User*

Runs the script with the permissions of the specified user account. This account must be a member of the Administrators group on the remote computer. By default, the script runs with the permissions of the current user of the local computer.

**/p** *Password*

Specifies the password of the user account that is pecified in the **/u** parameter. If you omit this parameter, the script prompts you for the password and obscures the text you type.

**Example**

The following command lists all of the Web service extension files that are in the **WebSvcExtRestrictionList** property of the metabase on the local computer:

**iisext /ListFile**

In response, **iisext** displays all of the Web service extension files that are on the local computer.

```
Connecting to server ...Done

E:\Windows\system32\inetsrv\asp.dll

E:\Windows\system32\inetsrv\httpodbc.dll

E:\Windows\system32\inetsrv\ssinc.dll

E:\Windows\system32\inetsrv\httpext.dll
```

```
E:\Windows\system32\Microsoft.NET\Framework\v1.0.3705\aspnet_isapi.dll

c:\Source\Taxes.dll

c:\Source\Accounting.dll

c:\Source\Taxes.exe

c:\Source\Timesheet.exe

c:\Source\Payscale.dll
```

**Related Information**

- For information about adding Web service extension files, see Adding Web Service Extension Files Using Iisext.vbs.
- For information about enabling Web service extension files, see Enabling Web Service Extension Files Using Iisext.vbs.
- For information about disabling Web service extension files, see Disabling Web Service Extensions Using Iisext.vbs.
- For information about removing Web service extension files, see Removing Web Service Extension Files Using Iisext.vbs.

# Listing Running Web Applications Using Iisapp.vbs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the command-line script iisapp.vbs, which is stored in *systemroot*\system32, to view any worker processes that are currently running. This command is often used for troubleshooting.

The computer issuing the command must be running Windows XP or a member of the Windows Server 2003 family. The computer that the command affects must be running a member of the Windows Server 2003 family with IIS 6.0.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

This topic includes the following information:

- Syntax: The order in which you must type a command and any arguments and options that follow it.
- Parameters: The values that are given to variables in the command.
- Examples: Sample code and an explanation of the results.

**Syntax**
**iisapp** *[* **/a** *AppPoolName* | **/p** *PID]*

**Parameters**
**/a** *AppPoolName*

Optional. Specifies the name of a particular application pool.

**/p** *PID*

Optional. Specifies an application pool by its ID number.

**Examples**
**Example 1:**
The following command displays all of the application pools that are running on the local computer:

**iisapp**

In response, **iisapp** displays all of the currently running applications, identifying each application pool by its process ID (PID) and application pool ID (AppPoolID). **Iisapp** omits all of the optional parameters and accepts the default values.

```
W3wp.exe PID: 2232 AppPoolID: DefaultAppPool

W3wp.exe PID: 2608 AppPoolID: MyAppPool
```

**Example 2:**
Use the following command to view a specific application by entering its PID:

**iisapp /p 2608**

In response, **iisapp** displays the specified application.

**W3wp.exe PID: 2608 AppPoolID: MyAppPool**

# Administering IIS Programmatically

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Graphical user interfaces (GUIs) on your computer—such as Internet Information Services (IIS) Manager, Computer Management Console, Control Panel, or Windows Explorer—help you manage and configure your settings. However, if you manage systems frequently or access large amounts of data on numerous servers, then using a collection of GUIs can be time consuming. In these cases, it is faster to use command-line scripts that can change single or multiple systems programmatically.

For example, using IIS Manager to add a script mapping to 20 Web servers will take many mouse clicks and key strokes. However, using a command-line script that takes an input file containing the names of your Web servers and the name of the property you want to change, you can make all of these changes from one server with a single command.

IIS 6.0 comes with three programmatic interfaces (also called providers), Windows Management Instrumentation (WMI), Active Directory Service Interface (ADSI), and Administrative Base Objects (ABO). Because the ABO interface can only be accessed from a compiled application, ABO is only discussed in the IIS SDK.

By using the WMI and ADSI providers, you can efficiently monitor the IIS metabase, and access and configure settings from within a command-line script or compiled application.

**The IIS WMI Provider**

Windows Management Instrumentation (WMI) is a technology that allows administrators to programmatically manage and configure the Windows operating system and Windows applications. Beginning with IIS 6.0, IIS includes a WMI provider that exposes programming interfaces that can be used to query and configure the IIS metabase.

**The IIS ADSI Provider**

The IIS ADSI provider exposes COM Automation objects that you can use within command-line scripts, ASP pages, or custom applications to change IIS configuration values stored in the IIS metabase. IIS ADSI objects can be accessed and manipulated by any language that supports automation, such as Microsoft® Visual Basic® Scripting Edition (VBScript), Microsoft JScript®, Perl, Microsoft Active Server Pages (ASP), Visual Basic, Java, or Microsoft C++.

**Comparing the ADSI and the IIS WMI Providers**

The following table compares the architecture and features of the IIS ADSI provider with those of the IIS WMI provider.

| Issue | IIS ADSI provider | IIS WMI provider |
|---|---|---|
| Query capabilities | No. ADSI has no provision for queries. | Yes. By querying on metabase key types, the IIS WMI provider returns only the data you need. |
| Object model and access routes | COM: Scripts and programs. | COM: Scripts, programs, and UI tools; for example, the WMI Object Browser. |
| Extensible schema | Yes. ADSI provider supports metabase schema extensions. | No. The IIS WMI provider can return existing schema extensions, but cannot extend the metabase schema. |
| Association or containment of related data | Properties are related to IIS metabase keys by containment.<br><br>You can use the ADSI container object methods of the IIS Admin Objects to manipulate keys in the IIS metabase. You can create, delete, and move keys by creating, deleting, and moving IIS Admin Objects within container objects. You can also enumerate contained objects, such as virtual directories or servers with container object methods. ADSI supports property inheritance. | Properties are related to IIS metabase keys by containment or by association. Also, in conjunction with other providers, WMI supports associations with managed objects not in the metabase.<br><br>An association in WMI describes a relationship between classes and is in itself defined by a class. This powerful concept allows management information about an entire system of associated components to be viewed and traversed for tasks such as troubleshooting. Navigating associations to other classes is not limited by containment. |

**Related Information**

- For more information about the IIS ADSI and IIS WMI providers, see IIS 6.0 Administration Scripts, Tips, and Tricks.
- For a more detailed comparison of the IIS ADSI and IIS WMI providers, see IIS Administration Technologies on MSDN.
- For detailed information about developing administrative scripts and applications using the IIS ADSI and IIS WMI providers, see Using IIS Programmatic Administration on MSDN.

# Security in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section describes how to configure your Web server to strengthen the security of your Web sites and to carry out other essential security functions.

**In This Section**

This section includes the following information:

- IIS 6.0 Security Best Practices: Provides a list of best practices for securing your Web site and server.
- Authentication in IIS 6.0: Provides information about how to confirm the identity of users that are attempting to establish connections to restricted content.
- Access Control with IIS 6.0: Provides information about how users access and manipulate your resources.
- IIS 6.0 Encryption: Provides information about how to protect the privacy of your information with Secure Sockets Layer (SSL) encryption features.
- Certificates_IIS_SP1_Ops: Provides information about how to establish secure connections by using certificates and SSL features.
- Auditing in IIS 6.0: Provides information about how to monitor security activities to prevent tampering and unauthorized access.

Related Information

- For information about isolating server applications so that they will be run in a process that is separate from the Web server process, see Isolating Applications in IIS 6.0.
- For information about verifying the network name and address of your computer, see Domain Name Resolution.

# IIS 6.0 Security Best Practices

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

**General Best Practices**

- Log on with the least credentials. Log on to your computer using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator.
- Reduce the attack surface. Disable all services you do not need, including IIS services such as FTP, NNTP or SMTP. If a feature or service is not enabled, then there is no need to secure it.
- Do not download or run programs from untrusted sources. Programs can contain instructions to violate security in a number of ways including data theft, denial of service, and data destruction.
- Keep virus scanners up to date. Virus scanners frequently identify infected files by scanning for a signature that is a known component of a previously identified virus. The scanners keep these virus signatures in a signature file, which is usually stored on the local hard disk. Because new viruses are discovered frequently, this file should also be updated frequently for the virus scanner to easily identify all current viruses.
- Keep current with all software updates. Software updates provide solutions to known security issues. Check software provider Web sites periodically to see if there are new updates available for software used in your organization.
    - The new process model in IIS 6.0 includes process recycling, which means an administrator can easily install most IIS updates and most new worker process DLLs without any interruption of service.
    - Auto Update version 1.0 provides three options to customers: notify update availability the moment it is available; download the update, and notify that it has been downloaded; and scheduled install. For more information, see "Windows Automatic Updates" in Help and Support Center for Windows Server 2003.
- Use NTFS. The NTFS file system is more secure than the FAT or FAT32 file system.
- Assign strong NTFS permissions for your resources.
- Exercise caution with domain controllers. If you use a domain controller as an application server, be aware that if security is compromised on the domain controller, then security is compromised on the entire domain.

**IIS Best Practices**

- Restrict write access permissions for the IUSR_*computername* account. This will help limit the access anonymous users have to your computer.
- Store executable files in a separate directory. This makes it easier to assign access permissions and audit for administrators.
- Create a group for all anonymous user accounts. You can deny access permissions to resources based on this group membership.
- Deny execute permissions for anonymous users to all executables in Windows directories and subdirectories.
- Use IP address restriction if administering IIS remotely. For more information, see Securing Sites with IP Address Restrictions.
- Assign the most restrictive permissions possible. For example, if your Web site is used only for viewing information,

assign only **Read** permissions. If a directory or site contains applications, assign **Scripts Only** permissions instead of **Scripts and Executables** permissions. For more information, see Securing Sites with Web Site Permissions.

- Do not assign **Write** and **Script source access** permissions or **Scripts and Executables** permissions. Use this combination with extreme caution. It can allow a user to upload potentially harmful executable files to your server and run them. For more information, see Securing Sites with Web Site Permissions.

- Enable data encryption in all WMI-based remote administration scripts. For more information, see Encrypting Data When Running WMI–Based Remote Administration Scripts.

- Ensure that the VeriSign Intermediate Root CA on your Web server is up to date. Verify the expiration date, and update the Intermediate Root CA if necessary. The new VeriSign Intermediate Root CA has the following properties:

  ```
  Issued to: www.verisign.com/CPS Incorp.by Ref. LIABILITY LTD.(c)97 VeriSign
  Issued by: Class 3 Public Primary Certification Authority
  Valid from: 4/16/97 to 10/24/11
  ```

  For information about verifying the expiration date, see Determining the Intermediate Root CA Version on a Web Server. For information about updating the Intermediate Root CA, see the VeriSign Support Web site.

**Related Information**

For information about security best practices for the Windows Server 2003 family, see "Security Best Practices" in Help and Support Center for Windows Server 2003.

# Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can require users to provide a valid Microsoft® Windows® user account name and password before they access any information about your server. This identification process is called authentication. Authentication can be set at the Web or FTP site, directory, or file level. IIS provides authentication methods to control access to your Web sites and FTP sites.

This section includes the following information:

- Anonymous Authentication in IIS 6.0: Describes how to configure anonymous authentication for IIS 6.0 and explains how to use the default anonymous user account, IUSR_*ComputerName*.
- Basic Authentication in IIS 6.0: Describes how to configure Basic authentication for IIS 6.0 and explains security considerations for the token cache.
- Digest Authentication in IIS 6.0: Describes how to configure both Digest authentication and sub-authentication in IIS 6.0.
- Advanced Digest Authentication in IIS 6.0: Describes how to configure Advanced Digest authentication in IIS 6.0.
- Integrated Windows Authentication in IIS 6.0: Describes how to configure Integrated Windows authentication in IIS 6.0. Integrated Windows authentication is the default authentication method for Windows Server™ 2003 operating systems.
- UNC Authentication in IIS 6.0: Describes how to configure UNC, or passthrough, authentication in IIS 6.0.
- .NET Passport Authentication in IIS 6.0: Describes how to configure .NET Passport authentication to work with IIS 6.0.
- FTP Site Authentication in IIS 6.0: Describes how to configure both anonymous and basic FTP site authentication in IIS 6.0.

# Anonymous Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Anonymous authentication gives users access to the public areas of your Web or FTP site without prompting them for a user name or password. By default, the IUSR_*computername* account is used to allow anonymous access.

In IIS 6.0, anonymous authentication no longer requires the Allow log on locally user right, because NETWORK_CLEARTEXT is now the default logon type for anonymous and basic authentication.

During setup, the IUSR_*computername* account is added to the Guests group on the computer running IIS. Guests have the same access as members of the Users group by default, except for the Guest account, which is further restricted.

This section includes the following information:

- Enabling Anonymous Authentication: Describes how to enable anonymous authentication for a Web site.
- Changing the Account Used for Anonymous Authentication: Describes how to configure an account for anonymous authentication.
- Configuring Subauthentication: Describes how to configure subauthentication to manage passwords for anonymous accounts.

# Enabling Anonymous Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the following procedure to enable anonymous authentication for a Web site.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable Anonymous authentication**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** folder, an individual Web site folder, a virtual directory, or a file; and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element. |

2. Click the **Directory Security** or **File Security** tab, and then, in the **Authentication and access control** section, click **Edit**.
3. Select the **Enable anonymous access** check box.
4. Click **OK** twice.

**Related Information**

- For information about changing the account for anonymous authentication, see Changing the Account Used for Anonymous Authentication.
- For information about other types of authentication, see Authentication in IIS 6.0.

# Changing the Account Used for Anonymous Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can change the account that is used for Anonymous authentication in IIS Manager, either at the Web server service level or for individual virtual directories and files. You can also change the security settings for the IUSR_*computername* account in Windows by using the Group Policy Manager snap-in of the Microsoft Management Console (MMC). However, if the anonymous user account does not have permission to access a specific file or resource, your Web server will refuse to establish an anonymous connection for that resource. For more information, see Securing Sites with Web Site Permissions.

**Important**

When you change the IUSR_*computername* account, the changes affect every anonymous HTTP request that a Web server processes. Use caution if you modify this account.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To change the account used for anonymous authentication**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** folder, an individual Web site folder, a virtual directory, or a file; and then click **Properties**.

   **Note**

   Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element.

2. Click the **Directory Security** or **File Security** tab, and then, in the **Authentication and access control** section, click **Edit**.
3. Select the **Enable anonymous access** check box.
4. Click **Browse** and type or browse to the valid Windows user account that you want to use for anonymous access.
5. Click **OK** three times.

**Related Information**

- For information about enabling anonymous authentication, see Enabling Anonymous Authentication.
- For information about other types of authentication, see Authentication in IIS 6.0.

# Configuring Subauthentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Subauthentication enables IIS to manage passwords on anonymous accounts. In earlier versions of IIS, the subauthentication component, Iissuba.dll, was enabled by default.

Because using Iissuba.dll can create a security risk, IIS 6.0 does not enable subauthentication by default. However, you can use subauthentication to manage passwords for anonymous accounts by meeting the following requirements:

- For applications to which you grant anonymous access, the worker process runs as LocalSystem. For more information about configuring worker process identities, see Configuring Worker Process Identities.
- The subauthentication component, Iissuba.dll, is registered.
- The **AnonymousPasswordSync** metabase property on the IISWebService node is enabled (set to **true)**.

The process for configuring subauthentication differs depending on whether you are configuring it on a new installation of IIS 6.0 or after upgrading to IIS 6.0 from an earlier installation of IIS with subauthentication enabled.

This section includes the following information:

- Configuring Subauthentication on a New Installation of IIS 6.0: Describes how to configure subauthentication on a new installation of IIS 6.0.
- Configuring Subauthentication in IIS 5.0 Isolation Mode: Describes how to configure subauthentication on a new installation of IIS 6.0 that is configured to run in IIS 5.0 isolation mode.
- Configuring Subauthentication After Upgrading to IIS 6.0: Describes how to configure subauthentication after upgrading your server to IIS 6.0 from an earlier version of IIS that uses subauthentication.

# Configuring Subauthentication on a New Installation of IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, after installing Windows Server 2003 and IIS 6.0, IIS runs in worker process isolation mode and subauthentication is disabled. Subauthentication is disabled when the AnonymousPasswordSync Metabase Property is set to **false**.

To enable anonymous password synchronization, ensure that your system meets the following requirements:

- The subauthentication component, Iissuba.dll, must be registered.
- The application pool of the Web site runs as the LocalSystem user account. However, running as the LocalSystem user account might be a security risk because it allows the worker process full access to the entire system.
- The AnonymousPasswordSync metabase property must be enabled (set to **true**).

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters). |

**To configure subauthentication and enable anonymous password synchronization**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. At the command prompt, type the following and then press ENTER:

   **rundll32 %systemroot%\system32\iissuba.dll,RegisterIISSUBA**

4. Set the identity of the application pool to the LocalSystem user account. For step-by-step information about setting the application pool identity, see Configuring Application Pool Identity in IIS 6.0.

   | **Important** |
   | --- |
   | Running as the LocalSystem user account might pose a security risk because it allows the worker process full access to the entire system. |

5. At the IisWebService node, set the AnonymousPasswordSync metabase property to **true**.

When you no longer want to use subauthentication, revoke the subauthentication components registration.

**To revoke the subauthentication components registration**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. At the command prompt, type the following and then press ENTER:

   **rundll32 %systemroot%\system32\iissuba.dll,UnregisterIISSUBA**

**Related Information**

- For information about configuring subauthentication on a new installation of IIS 6.0 that has been configured to run in IIS 5.0 isolation mode, see Configuring Subauthentication in IIS 5.0 Isolation Mode.

- For information about configuring subauthentication after upgrading your server to IIS 6.0 from an earlier version of IIS that uses subauthentication, see Configuring Subauthentication After Upgrading to IIS 6.0.

# Configuring Subauthentication in IIS 5.0 Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you perform a new installation of IIS 6.0 and then configure the server to run in IIS 5.0 isolation mode, the worker process that is assigned to in-process applications will run as LocalSystem by default. However, for subauthentication to work properly, you must still register Iissuba.dll and configure the AnonymousPasswordSync Metabase Property.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To configure subauthentication after a new installation and after switching to IIS 5.0 isolation mode**

1. Register Iissuba.dll by opening a command prompt and then typing the following:

   **rundll32** *%windir%* **\system32\iissuba.dll,RegisterIISSUBA**

2. At the IisWebService node, set the AnonymousPasswordSync metabase property to **true**.

**Related Information**

- For information about configuring subauthentication on a new installation of IIS 6.0, see Configuring Subauthentication on a New Installation of IIS 6.0.

- For information about configuring subauthentication after upgrading your server to IIS 6.0 from an earlier version of IIS that uses subauthentication, see Configuring Subauthentication After Upgrading to IIS 6.0.

# Configuring Subauthentication After Upgrading to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you upgrade a server to IIS 6.0 from an earlier version of IIS that uses subauthentication to manage passwords on anonymous accounts, subauthentication is enabled by default. The AnonymousPasswordSync Metabase Property is set to **true**. However, two additional configuration tasks must be completed before subauthentication will work correctly in IIS 6.0. Your event log should have entries about this.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To configure subauthentication after upgrading to IIS 6.0 from an earlier version of IIS configured to use subauthentication**

1. Register Iisuba.dll by opening a command prompt and then typing the following:

   **rundll32** *%windir%* **\system32\iissuba.dll,RegisterIISSUBA**

2. Run all worker processes that use Anonymous authentication as LocalSystem. For more information about configuring worker process identities, see Configuring Worker Process Identities.

**Related Information**

- For information about configuring subauthentication on a clean installation of IIS 6.0, see Configuring Subauthentication on a New Installation of IIS 6.0.

- For information about configuring subauthentication on a new installation of IIS 6.0 that is configured to run in IIS 5.0 isolation mode, see Configuring Subauthentication in IIS 5.0 Isolation Mode.

# Basic Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The Basic authentication method is a widely used, industry-standard method for collecting user name and password information. Basic authentication transmits user names and passwords across the network in an unencrypted form. You can use your Web server's encryption features, in combination with Basic authentication, to secure user account information transmitted across the network. For easier administration, add each user to a group that has access to the necessary files.

| **Important** |
| --- |
| Enabling Basic authentication does not automatically configure your Web server to authenticate users. You must also create Windows user accounts and set appropriate NTFS permissions. For more information about setting NTFS permissions, see Setting NTFS Permissions for Directories or Files. |

This section includes the following information:

- Enabling Basic Authentication and Configuring the Realm Name: Describes how to enable Basic authentication and configure the realm name.
- Setting the Default Logon Domain: Describes how to set the default logon domain for use with Basic authentication.
- Configuring Token Cache for Basic Authentication: Describes how to configure token caching so user tokens are not stored in the cache too long.

# Enabling Basic Authentication and Configuring the Realm Name

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Users that are authenticated with Basic authentication must log on with a valid Windows user name and password.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable Basic authentication and configure the realm name**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** folder, an individual Web site folder, a virtual directory, or a file; and then click **Properties**.

   Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element.

2. Click the **Directory Security** or **File Security** tab, and then, in the **Authentication and access control** section, click **Edit**.

3. In the **Authenticated access** section, select the **Basic authentication** check box.

4. Because Basic authentication sends passwords over the network unencrypted, a dialog box appears, asking if you want to proceed. Click **Yes** to proceed.

5. In the **Default domain** box, either type the domain name you want to use, or click **Select** to browse to a new default logon domain. If the **Default domain** box is filled in, the name is used as the default domain. If the **Default domain** box is left empty, IIS uses the domain of the computer that is running IIS as the default domain. IIS configures the value of the DefaultLogonDomain Metabase Property, which determines the default domain used to authenticate clients accessing your IIS server using Basic authentication. However, the domain specified by DefaultLogonDomain is used only when a client does not specify a domain in the logon dialog box that appears on the client computer.

6. Optionally, you can enter a value in the **Realm** box, which configures the value of the Realm Metabase Property. If the Realm property is set, its value appears on the client's logon dialog box, when Basic authentication is used. The value of Realm is sent to the client for informational purposes only, and is not used to authenticate clients using Basic authentication.

7. Click **OK** twice.

**Related Information**

- For information about setting the default logon domain, see Setting the Default Logon Domain.

- For more information about permissions, see "Access control" in Help and Support Center for Windows Server 2003.

# Setting the Default Logon Domain

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Users authenticated with Basic authentication must log on with a valid Windows user name and password. The user name usually includes:

- A Windows *domain* name.
- An account user name.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set the default logon domain**

1. In IIS Manager, right-click the **Web Sites** folder, Web site, directory, virtual directory, or file, and click **Properties**.

| **Note** |
| --- |
| Configuration settings made at the Web Sites folder level can be inherited by all Web sites. |

1. Click the **Directory Security** or **File Security** tab, depending on the level at which you want to configure security settings.
2. In the **Anonymous access and authentication** control section, click **Edit**.
3. In the **Authenticated access** section, select the **Basic authentication** check box.
4. Because Basic authentication sends passwords over the network unencrypted, a dialog box appears, asking if you want to proceed. Click **Yes** to proceed.
5. In the **Default domain** box, either type the domain name you want to use, or click **Select** to browse to a new default logon domain. If the **Default domain** box is filled in, the name is used as the default domain. If the Default domain box is left empty, IIS uses the domain of the computer that is running IIS as the default domain. IIS configures the value of the DefaultLogonDomain Metabase Property, which determines the default domain used to authenticate clients accessing your IIS server using Basic authentication. However, the domain specified by DefaultLogonDomain is used only when a client does not specify a domain in the logon dialog box that appears on the client computer.
6. Click **OK** twice.

**Related Information**

- For information about Basic authentication, see Basic Authentication.

# Configuring Token Cache for Basic Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Basic authentication stores user tokens in a token cache. If you log on using Basic authentication with an account that has a high level of user logon rights, a successful attacker could use the account to gain access to the resources on your computer. There are several ways to minimize this threat:

- Do not log on, or allow anyone to log on, using Basic authentication with an account that has a high level of user logon rights.

- Configure the token cache to flush all tokens by setting the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters FlushTokenCache registry key to any non-zero value, or set a short time span for cached user tokens by setting the Global Registry entry, UserTokenTTL, to less than the default of 15 minutes. Setting UserTokenTTL to 0 disables TTL-based flushing of tokens. When TTL-based flushing is disabled, user tokens remain cached until either IIS is restarted or the worker process is recycled. See UserTokenTTL in Global Registry Entries.

Before you edit the registry, make sure you understand how to restore it if a problem occurs. For information about how to restore the registry, see the "Restoring the Registry" and the "Restoring a Registry Key" topics in Registry Editor Help.

| **Important** |
| --- |
| Using Registry Editor incorrectly can cause serious problems that require reinstalling the operating system. Because Registry Editor bypasses the standard safeguards that prevent you from entering settings that are conflicting or likely to degrade performance or damage your system, exercise caution when making changes to the registry. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. For information about how to edit the registry, see "Changing Keys and Values" in Registry Editor Help. |

Note that you should back up the registry before you edit it. If you are running Microsoft® Windows NT® Server or newer server operating system from Microsoft, you should also update your Emergency Repair Disk (ERD).

For information about how to edit the registry, see the "Changing Keys and Values", "Add and Delete Information in the Registry", and "Edit Registry Data" topics in Registry Editor Help.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To configure token caching**

1. From the **Start** menu, click **Run**.

2. In the **Open** box, type Regedit.exe, and click **OK**.

3. Navigate to and double-click the following key in the registry:

4. HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\InetInfo\Parameters.

5. From the **Edit** menu, point to **Add**, click DWORD Value and then add the following registry value:

   Name: **UserTokenTTL**

   Type: **REG_DWORD**

Data: *number of seconds* to cache user tokens, or **0** to disable token caching

6. Quit Registry Editor.
7. Restart IIS.

**Related Information**

- For information about setting the default logon domain, see Setting the Default Logon Domain.
- For information about the **DefaultLogonDomain** property, see DefaultLogonDomain Metabase Property.
- For more information about permissions, see "Access control" in Help and Support Center for Windows Server 2003.

# Digest Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Digest authentication offers the same functionality as Basic authentication; however, Digest authentication provides a security improvement in the way that a user's credentials are sent across the network. Digest authentication transmits credentials across the network as an MD5 hash, or message digest, where the original user name and password cannot be deciphered from the hash. Digest authentication is available to Web Distributed Authoring and Versioning (WebDAV) directories. Digest authentication requires HTTP 1.1 compliance.

This section includes the following information:

- Configuring Digest Authentication in IIS 6.0: Describes how to configure Digest authentication.
- Configuring Subauthentication in a Windows 2000 Domain: Describes how to configure subauthentication so that Digest authentication can be used in a Windows 2000 domain.

# Configuring Digest Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Configuring Digest authentication on the server running IIS requires the following two tasks:

- Enable Digest authentication for Windows domain servers.
- Configure the realm name.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable Digest authentication for Windows domain servers and configure the realm name**

1. In IIS Manager, right-click the **Web Sites** folder, Web site, directory, virtual directory, or file, and click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web Sites folder level can be inherited by all Web sites. |

2. Click the **Directory Security** or **File Security** tab, depending on the level at which you want to configure security settings.
3. In the **Anonymous access and authentication control** section, click **Edit**.
4. In the **Authenticated access** section, select the **Digest authentication for Windows domain servers** check box.
5. In the **Realm** box, type the realm name, or click **Select** to browse for a domain.
6. Click **OK** twice.
7. Restart the W3SVC service.

| **Note** |
| --- |
| If Basic authentication is enabled for the site, virtual directory, or folder you are configuring, the **Default domain** box will also be available. However, only **Realm** is meaningful to Digest authentication. |

**Related Information**

- For information about the **Realm** property, see Realm Metabase Property.
- For information about Advanced Digest authentication, see Advanced Digest Authentication in IIS 6.0.

# Configuring Subauthentication in a Windows 2000 Domain

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To use Digest authentication in IIS 6.0 when the domain controller is running Windows 2000 Server, you must enable subauthentication, which is not installed by default on IIS 6.0. There are three steps required to enable subauthentication:

- Register the subauthentication component, Iissuba.dll.
- Set the UseDigestSSP Metabase Property to **false**. See Advanced Digest Authentication in IIS 6.0 for more information.
- Set the identity of the application pool to LocalSystem. See Configuring Application Pool Identity with IIS 6.0 for more information.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To register the subauthentication component**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following and then press ENTER:

   **rundll32 %systemroot%\system32\iissuba.dll,RegisterIISSUBA**.
4. For any application pools that use Digest authentication, set the identity as LocalSystem.

When you no longer want to use subauthentication, unregister the subauthentication component.

**To unregister the subauthentication component**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following and then press ENTER:

   **rundll32 %systemroot%\system32\iissuba.dll,UnregisterIISSUBA**.

**Related Information**

- For more information about configuring subauthentication in IIS 6.0, see Configuring Subauthentication.

# Advanced Digest Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Under Advanced Digest authentication, user credentials are stored on the domain controller as an MD5 hash. Because credentials are stored in Active Directory as an MD5 hash, user passwords cannot be feasibly discovered by anyone with access to the domain controller, not even by the domain administrator. Advanced Digest authentication is available to Web Distributed Authoring and Versioning (WebDAV) directories. In IIS 6.0, Advanced Digest authentication is preferred over Digest authentication, but Digest authentication is still available. Advanced Digest authentication relies on the HTTP 1.1 protocol.

Advanced Digest authentication uses the UseDigestSSP Metabase Property. This metabase key is a switch between Digest and Advanced Digest Security Support Provider Interface (SSPI) code. After the key has been set, the only valid property values are 1 (**true**), 0 (**false**), or empty. If the property is set to **true**, the new SSPI code for Advanced Digest authentication is used. In all other cases (**false**, empty, or not set), IIS uses the Digest authentication code.

| **Note** |
| --- |
| The World Wide Web Publishing Service (WWW service) must be restarted before changes to UseDigestSSP take effect. |

Configuring Advanced Digest authentication on the server running IIS requires the following three tasks:

- Enable Digest authentication for Windows domain servers.
- Configure the realm name.
- Set the UseDigestSSP metabase property to **true**. You can configure the UseDigestSSP metabase property at the W3SVC level of the metabase. A child key inherits its configuration from the level above it.

| **Important** |
| --- |
| If you follow the first two procedures, but do not configure the UseDigestSSP metabase property, you will be using Digest authentication, not Advanced Digest authentication. |

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc** *%systemroot%***\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable Advanced Digest authentication and configure the realm name for Windows domain servers**

1. In IIS Manager, right-click the **Web Sites** folder, Web site, directory, virtual directory, or file, and click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web Sites folder level can be inherited by all Web sites. |

2. Click the **Directory Security** or **File Security** tab, depending on the level at which you want to configure security settings.
3. In the **Anonymous access and authentication control** section, click **Edit**.
4. In the **Authenticated access** section, select the **Digest authentication for Windows domain servers** check box.
5. In the **Realm** box, type the realm name, or click **Select** to browse for a domain. If Basic authentication is enabled for the site, virtual directory, or folder you are configuring, the **Default domain** box will also be available. However, only the

realm is meaningful for Advanced Digest authentication.

6. Click **OK** twice.
7. Restart the W3SVC service.

**Related Information**

- For more information about Digest authentication, see Digest Authentication in IIS 6.0, and the RFC 2617 specification as defined at the World Wide Web Consortium Web site.
- For more information about the **Realm** metabase property, see Realm Metabase Property.
- For more information about Advanced Digest Authentication, see Advanced Digest Authentication.
- For more information about Active Directory or domain controllers, see "Domains" in Help and Support Center for Windows Server 2003.

# Integrated Windows Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

Integrated Windows authentication (formerly called NTLM, and also referred to as Windows NT Challenge/Response authentication) is a secure form of authentication because the user name and password are hashed before being sent across the network. When you enable Integrated Windows authentication, the user's browser proves its knowledge of the password through a cryptographic exchange with your Web server, involving hashing.

| **Important** |
| --- |
| Integrated Windows authentication is disabled by default if you install Windows Server 2003 Service Pack 1 (SP1) as part of a slipstream installation of a Windows Server 2003 operating system. If you install Windows Server 2003 SP1 by itself as an upgrade to a Windows Server 2003 operating system, the setting for Integrated Windows authentication is unchanged from its Windows Server 2003 setting. Integrated Windows authentication is enabled by default for Windows Server 2003 operating systems. |

This section includes the following information:

- Configuring Integrated Windows Authentication in IIS 6.0: Describes how to configure Integrated Windows authentication.
- Configuring Constrained Delegation for Kerberos: Describes how to configure constrained delegation when using Kerberos as the authentication method.
- Forcing NTLM Authentication: Describes how to force Integrated Windows authentication to use NTLM instead of Kerberos.

# Configuring Integrated Windows Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

Integrated Windows authentication was formerly called NTLM. This authentication method sends a hash of the user name and password across the network.

Integrated Windows authentication is disabled by default if you install Windows Server 2003 Service Pack 1 (SP1) as part of a slipstream installation of a Windows Server 2003 operating system. If you install Windows Server 2003 SP1 by itself as an upgrade to a Windows Server 2003 operating system, the setting for Integrated Windows authentication is unchanged from its Windows Server 2003 setting. Integrated Windows authentication is enabled by default for Windows Server 2003 operating systems.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure Integrated Windows authentication**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** folder, an individual Web site folder, a virtual directory, or a file; and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the **Web Sites** level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element. |

2. Click the **Directory Security** or **File Security** tab, and then, in the **Authentication and access control** section, click **Edit**.
3. In the **Authenticated access** section, select the **Windows Integrated Authentication** check box.
4. Click **OK** twice.

**Related Information**

- For more information about Integrated Windows authentication, see Integrated Windows Authentication.
- For information about configuring Integrated Windows authentication to force NTLM authentication, see Forcing NTLM Authentication.
- For information about configuring worker process identity when using Kerberos authentication, see Configuring Constrained Delegation for Kerberos.

# Configuring Constrained Delegation for Kerberos

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

*Constrained delegation*, which is new in Windows Server 2003, is intended to be used by service accounts, which should have registered Service Principal Names (SPNs), instead of by a regular user account, which typically does not have an SPN.

The Setspn.exe command-line utility allows you to read, modify, and delete SPNs for an Active Directory property. Setspn.exe is available in the Support Tools pack located on your Windows Server 2003 CD-ROM.

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer\Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters).

**Important**

You must be a domain administrator to set an SPN.

**To configure constrained delegation**

1. Install the Support Tools pack on your Windows Server 2003 CD-ROM.
2. Click **Start**, point to **Programs** or **All Programs**, point to **Windows Support Tools**, and then click **Command Prompt**.
3. At the command prompt, type the following command:

   ```
   setspn -a http/SiteName.DomainName.com Domain \User
   ```

   where *SiteName.DomainName.com* is the Web site whose application pool runs with the custom identity you created.

   For example, the following command assigns "http/contosohr.contoso.com" to the custom identity CONTOSO\cd1hr. The **http/** prefix is a service class that identifies this as an SPN for a Web site.

   ```
   setspn -a http/contosohr.contoso.com CONTOSO\cd1hr
   ```

   Here is another way to specify the same SPN, using just the SiteName:

   ```
   setspn -a http/contosohr CONTOSO\cd1hr
   ```

If you register duplicate SPNs accidentally, you can use Setspn.exe to delete the duplicate SPN. For more information about Setspn.exe syntax, see Setspn Syntax.

**Note**

If you are configuring servers running IIS 6.0 in a Windows 2000 Server domain, you can use either the version of Setspn.exe that comes with Windows Server 2003 or the version that comes with Windows 2000 Server.

**Related Information**

- For more information about configuring application pool identity, see Configuring Application Pool Identity with IIS 6.0.
- For more information about Kerberos and service registration, see Integrated Windows Authentication.
- For more information about NTLM authentication, see "Authentication" in Help and Support Center for Windows Server 2003.
- For more information about Kerberos, see "Kerberos" in Help and Support Center for Windows Server 2003.

- For more information about constrained delegation, see "Delegating authentication" in Help and Support Center for Windows Server 2003.

# Forcing NTLM Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In the following situations, Kerberos authentication fails and you must force IIS to use NTLM authentication by setting the NTAuthenticationProviders Metabase Property to **NTLM**.

- Kerberos fails when you isolate Web sites on a virtual directory level by configuring worker process identities as different domain accounts.
- Kerberos authentication fails if you are using Integrated Windows authentication, are not using a WINS or DNS name for the server running IIS, and you want to use a local user account or the LocalService account as a worker process identity. It fails because Active Directory does not trust these accounts.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To force NTLM authentication**

1. In IIS Manager, right-click the local computer, and then click **Properties**.
2. Select the **Enable Direct Metabase Edit** check box, and then click **OK**.
3. Click **Start**, click **Run**, type **cmd**, and then click **OK**.
4. At the command prompt, type the following command to change to the directory where the MetaBase.xml file is located:

   **cd %systemroot%\system32\inetsrv**

5. To open the file with Notepad, at the command prompt, type the following:

   **notepad MetaBase.xml**

6. In the **<IISWebServer>** section, locate the **NTAuthenticationProviders** metabase property and modify its setting to read "NTLM".
7. Save and close the MetaBase.xml file.

**Related Information**

- For information about metabase edit-while-running, see Enabling Edit-While-Running in IIS 6.0.
- For more information about NTLM authentication, see "Authentication" in Help and Support Center for Windows Server 2003.
- For more information about Kerberos, see "Kerberos" in Help and Support Center for Windows Server 2003.

# UNC Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In IIS 6.0, you can use delegation of authentication, a feature of IIS 6.0 and the Windows Server 2003 family, to pass the authenticated credentials of the request user to a remote file server on which a Universal Naming Convention (UNC) share is located. Delegation of authentication, also called UNC authentication, allows virtual directories mapped to remote UNC shares to use local user authentication for access to the shares. This is preferable to the alternative of using a static user name and password, because you can limit access to specific portions of the UNC share for each user or group. Delegation of authentication works only with authentication methods that can perform delegation, for example, Basic authentication or Kerberos authentication.

IIS can be configured to secure a virtual directory as you create it or to configure an existing directory.

This section includes the following information:

- Creating a Virtual Directory with a UNC Path Describes how to create a secure virtual directory that uses a UNC path.
- Securing an Existing Virtual Directory with a UNC Path Describes how to configure a UNC path to secure an existing virtual directory.
- Configuring Constrained Delegation for IIS: Describes how to configure constrained delegation.

# Creating a Virtual Directory with a UNC Path

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the following procedure to create a secure virtual directory with a UNC path.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a secure virtual directory with a UNC path**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click a Web site, Web application, directory, or virtual directory, point to **New**, and then click **Virtual Directory**.
2. Click **Next**.
3. In the **Alias** box, type the alias for the new virtual directory, and click **Next**.
4. In the **Path** box, type the path to the directory on the UNC share, and click **Next**.
5. The **Virtual Directory Creation Wizard** recognizes the path format as UNC and the **User Name and Password** dialog box appears.
6. To impersonate the requested user credentials to the remote computer, select the **Always use the authenticated user's credentials when validating access to the network resource** check box and leave the **User name** and **Password** boxes blank. Click **Next** and proceed to step 8.
7. To secure the virtual directory and share with a static user name and password, specify the user name as Domain\username in the **User name** box, specify a password in the **Password** box, and clear the **Always use the authenticated user's credentials when validating access to the network resource** check box.
8. Select the appropriate check boxes for the access permissions you want to allow, and clear the check boxes for the permissions you want to deny.
9. Click **Next** to complete the wizard.

**Related Information**

- For information about delegation, see "Delegation of Authentication" in Help and Support Center for Windows Server 2003.

# Securing an Existing Virtual Directory with a UNC Path

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the following procedure to further secure an existing virtual directory with a UNC path.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To secure an existing virtual directory with a UNC path**

1. In IIS Manager, right-click the virtual directory you want to secure, and click **Properties**.
2. Click **Connect As...**.
3. To use delegation to pass the local user's credentials to the remote computer, select the **Always use the authenticated user's credentials when validating access to the network resource** check box and leave the **Username** and **Password** boxes blank. Proceed to step 5.
4. To secure the virtual directory and share with a static user name and password, specify the user name as Domain\username in the **User name** box, specify a password in the **Password** box, clear the **Always use the authenticated user's credentials when validating access to the network resource** check box, and click **Next**.
5. Click **OK**.

**Related Information**

- For information about delegation, see "Delegation of Authentication" in Help and Support Center for Windows Server 2003.

# Configuring Constrained Delegation for IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure your server to allow clients to authenticate using any authentication protocol, and IIS will access content on the file server on behalf of the authenticated clients. For more information about constrained delegation, see UNC Authentication.

| **Important** |
| --- |
| You must be a domain administrator to perform the following procedure. |

**To configure constrained delegation for IIS**

1. From the **Start** menu, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
2. In the **Active Directory Users and Computers** dialog box, ensure that your domain is in the **Windows Sever 2003** mode by doing the following:
   - In the left pane, click your domain name.
   - From the **Action** menu, select **Raise Domain Functional Level**.
   - In the **Raise Domain Functional Level** dialog box, if your domain is in Windows 2000 native or Windows 2000 mixed mode, click **Windows Server 2003** from the **Select an available domain functional level** list box, and then click **Raise**.

| **Important** |
| --- |
| After you change your domain level, you cannot reverse the action. The change can take up to 15 minutes to propagate. For more information about domain functional levels, see "Domain and forest functionality" in Help and Support Center for Windows Server 2003. |

3. In each of the two **Raise Domain Functional Level** dialog boxes that appear, click **OK**.
4. Double-click the domain, and then click the **Computers** folder. A list of computers in the domain appears in the right pane.
5. Right-click the Web server computer name, and then select **Properties**.
6. Click the **Delegation** tab, and then select the **Trust this computer for delegation to specified services only** check box.
7. Click the authentication type that you want to use, and then click **Add**.
8. In the **Add Services** dialog box, click **Users or Computers**.
9. In the **Select Users or Computers** dialog box, in the **Enter the object names to select** text box, search for or type the name of the file server that you want to use, and then click **OK**.
10. In the **Properties** dialog box, click **Add** to add the HOST and CIFS services for the target file server. Add only the services that you are sure must present delegated credentials, and then click **OK**.

# .NET Passport Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Microsoft® .NET Passport is a user-authentication service and a component of the Microsoft .NET framework. The .NET Passport single sign-in service and express purchase service enable your business to deliver a fast, convenient, and secure way for consumers to sign in and make transactions on your site. With the .NET Passport single sign-in service, you can map sign-in names to information in your databases, which enables you to offer .NET Passport members a personal Web experience through targeted ads, promotions, and content. Using .NET Passport in this way potentially helps you increase customer loyalty, sales, and advertising revenues.

.NET Passport uses standard Web technologies and techniques, such as Secure Sockets Layer (SSL), HTTP redirects, cookies, Microsoft JScript®, and strong symmetric key encryption to deliver the single sign-in service. The .NET Passport service is compatible with both Microsoft Internet Explorer version 4.0 or later, and Netscape Navigator version 4.0 or later. In addition, .NET Passport is compatible with some versions of UNIX.

This section includes the following information:

- Setting Up .NET Passport in IIS 6.0: Describes how to configure and test your .NET Passport configuration in IIS 6.0 and against .NET Passport production servers.

- Enabling .NET Passport Authentication in IIS 6.0: Describes how to implement .NET Passport authentication on your Web sites through IIS Manager.

# Setting Up .NET Passport in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Before setting up Microsoft .NET Passport authentication on your Web sites in a production environment, you are required to test IIS against .NET Passport preproduction servers. By working through this process, you confirm that your IIS server and the .NET Passport server are communicating correctly, your site(s) is registered with .NET Passport (which might involve signing forms and agreements), and your Web sites have the required site IDs. You must complete each process for every Web site you want to enable with .NET Passport authentication.

This topic includes the following information:

- .NET Passport Environments
- Testing and Preproduction (PREP)
- Mapping to Active Directory
- Configuring for Preproduction (PREP)
- Configuring for Production
- Passport Manager Administration Utility

With .NET Passport authentication on members of the Windows Server 2003 family, the default .NET Passport **SecureLevel** setting is 10. This means that new sites using .NET Passport authentication (and default settings) require an Secure Sockets Layer (SSL) server certificate. You can change the **SecureLevel** setting for a site by changing a registry value.

| **Important** |
| --- |
| Using Registry Editor incorrectly can cause serious problems that require reinstalling the operating system. Because Registry Editor bypasses the standard safeguards that prevent you from entering settings that are conflicting or likely to degrade performance or damage your system, exercise caution when making changes to the registry. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. For information about how to edit the registry, see "Changing Keys and Values" in Registry Editor Help. |

**Procedures**

**To change the SecureLevel setting for the default Web site**

- Enter a new value in the registry for the following key:

  HKEY_LOCAL_MACHINE\Software\Microsoft\Passport\SecureLevel

**To change the SecureLevel setting for any Web site other than the default Web site**

- Enter a new value in the registry for the following key:

  HKEY_LOCAL_MACHINE\Software\Microsoft\Passport\Sites\<*Site Name*>\SecureLevel

| **Note** |
| --- |
| This documentation is designed to be used in conjunction with the .NET Passport Software Development Kit (SDK) documentation (starting with version 2.1). If you have not already done so, review the SDK documentation, specifically the section titled "How to Implement .NET Passport." You may find it helpful to browse a sample .NET Passport Web site to better understand what a site running .NET Passport looks like. You can also review the code samples or simply use the code in the sample site as a template for your own site. For more information about the .NET Passport SDK and the sample Web site, see Microsoft .NET Passport for Developers. |

**.NET Passport Environments**
Setting up .NET Passport on your Web sites involves testing and configuring IIS against the following three .NET Passport environments:

- Default installation environment: By default, the Microsoft .NET Passport SDK and Internet Information Services (IIS) (with the Passport Manager object) are configured for testing. The test or default installation environment uses a site ID of 1 (one), and has a default encryption key instead of a private, site-specific key.

- Preproduction: The Preproduction (PREP) environment enables sites to verify their development efforts against .NET Passport servers without access to real-world .NET Passport user identifications and profiles. While in PREP mode, prospective .NET Passport participating sites can manipulate data, create new users, and run other tests against the user base without affecting existing deployed .NET Passport sites.

- Production: The .NET Passport Production environment is shared by all working and approved .NET Passport participating sites after they are deployed to the public.

**Testing and Preproduction (PREP)**
You can determine whether your site is in the default installation environment by checking the **Site ID** field in the Passport Manager Administration utility. If the value in the **Site ID** field is 1, your site is in the default installation environment. While your site is in this environment, you can perform an initial evaluation of the single sign-in service (SSI) and test any dynamic content on your site against the Passport Manager object. You can develop your Web site and any applications on your site before registering your site and requesting a site ID, or while you are waiting for the site ID and key after registration. After you have a site ID, there should be no reason to use the default installation environment again. The restrictions imposed by this mode are intended only to prevent developers from accessing certain .NET Passport features before having signed the necessary agreements or contracts.

| **Note** |
| --- |
| The preproduction environment requires a .NET Passport preproduction logon account. To set up a preproduction logon account, see Microsoft .NET Services Manager. |

**Limitations in the Default Installation Environment**
The following operational limitations exist in the default installation environment:

- Cannot read full core profile data or profile cookies: A Web site running in the default installation environment indicates that the administrator has not yet signed the necessary agreements that specify a site's requirements for privacy. For this reason, a site running in the default installation environment cannot access all .NET Passport user information. .NET Passport user information displayed on a test site contains several fields of default placeholder values, as generated by the logon server.

- Cannot use co-branding: Co-branding support is dependent on several URLs that you must provide when registering as a .NET Passport participating site. Until you register and provide these URLs, the co-branding logo and text exist as placeholders. For details about how to implement co-branding after your site has registered as a .NET Passport participating site, see ".NET Passport Cobranding Overview" in .NET Passport SDK Help.

- Cannot sign users out from sign-out page: On a registered .NET Passport site, when members choose to sign out by clicking the .NET Passport sign-out link, they are redirected to a central page that enables deletion of all .NET Passport cookies from all of the sites the member visited during the session. This is not the case in the default installation environment, so .NET Passport cookies remain on the user's browser until the user closes the browser entirely. If the .NET Passport user chose to save their password (thus making all session cookies persistent cookies), .NET Passport cookies written in the test site's domain are still not deleted. For this reason, when testing browser behavior in the default installation environment, you may occasionally need to quit the browser to reproduce a clean .NET Passport sign-in.

**Testing Passport with Active Server Pages**
If your Web site uses Active Server Pages (ASP) to query information from the Passport authentication ticket, you must create a Passport Manager object. The Passport Manager object is a server-side object for .NET Passport single sign-in and profile services. The Passport Manager object uses cookies and query string data as intermediaries for querying a central user store.

To create a Passport Manager object in an ASP page, use the ASP **Server.CreateObject** method to create a local object

reference as follows:

```
set oMgr = Server.CreateObject("Passport.Manager")
```

The Passport Manager object should only be created once on each ASP page. Multiple object instantiations on a single page can lead to unexpected results when the page loads or unloads. Also, multiple instantiations will slow performance with extra memory overhead and multiple encryption and decryption of the same cookies. One way to assure a single instantiation is to create the object only once in a header file that is included on each page.

Do not create the Passport Manager object in an ASP application in Session or Application scope, unless you are specifically using that instance for utility methods such as **GetDomainAttribute** or **DomainFromMemberName** (with input parameter).

The following code sample instantiates the Passport Manager object in an ASP page to obtain user information from a .NET Passport cookie. Web sites running in the default installation environment return generic values. Web sites running in PREP or Production environments return valid user information.

```
%
'
' Instantiate the Passport Manager.
'
Set oMgr = Server.CreateObject("Passport.Manager.1")
'
' Use HTTPS if SecureLevel > 1.
'
thisURL = Server.URLEncode("https://" & Request.ServerVariables("SERVER_NAME") &
Request.ServerVariables("SCRIPT_NAME") ) signoutURL = Server.URLEncode("https://" &
Request.ServerVariables("SERVER_NAME") & "/signout.asp")
If oMgr.IsAuthenticated(3600,FALSE) Then
'
' Convert the binary PUID to a hexadecimal format for displaying.
'
strMemberIDHigh = oMgr.Profile("MemberIDHigh")
strMemberIDLow = oMgr.Profile("MemberIDLow")
strHex = Hex(strMemberIDHigh)
strPUID = String(8 - Len(strHex), "0") & strHex
strHex = Hex(strMemberIDLow)
strPUID = strPUID & String(8 - Len(strHex), "0") & strHex
```

For more information, see "Passport Manager Methods" or "Passport Manager Properties" in .NET Passport SDK (starting with version 2.1) Help.

#### Testing ASP Against the .NET Passport Default Installation Environment

### To test ASP against the .NET Passport default installation environment

1. Enable .NET Passport authentication on a virtual directory. You can use the code sample above in the Default.asp page and include the Default.asp page in the virtual directory. To learn more about enabling .NET Passport authentication on a Web site or a virtual directory, see Enabling .NET Passport Authentication in IIS 6.0.

   | **Note** |
   | --- |
   | Microsoft Windows XP client computers used to test .NET Passport must be configured for .NET Passport default installation and PREP environments. To configure your Windows XP client computers to connect to these environments, see ".NET Passport SDK: System Requirements" in .NET Passport SDK (starting with version 2.1) Help. More specifically, read the section titled "Client Computer Requirements." |

2. On a remote client computer, open a Web browser, type the virtual directory URL in the address bar, and press ENTER. The client connects to the Web site and is then redirected to the .NET Passport server. Passport reads the site ID as 1, which tells the Passport service that your server is still in the default installation environment. You are then redirected to the .NET Passport test environment for credential authentication. Windows 2000 or earlier clients open a logon Web page. Windows XP clients (configured according to "Client Computer Requirements" in .NET Passport SDK Help) open the **Stored User Names and Passwords** dialog box.

3. Type valid .NET Passport credentials and click **OK**.

4. .NET Passport validates the credentials and displays hard-coded, generic user information. The server running IIS collects a cookie with the generic information. The display does not contain any confidential material. The .NET Passport User ID (PUID) is not real because, at this point, your IIS server has not established a contract with .NET Passport.

**Mapping to Active Directory**

In certain scenarios (for example Business-to-Business [B2B] or Business-to-Consumer [B2C]) you might want to set up Access Control Lists (ACL) on specific Web pages or content and then have selective members or groups access the Web pages using .NET Passport authentication. As an example, a large company with multiple vendors might establish a timecard reporting site for all contingent staff. By setting up the timecard reporting online, vendors can access the Web page from any Web browser, and the company can control security using Passport authentication.

To configure.NET Passport authentication that uses ACLs, the .NET Passport user must have a PUID mapped to an account in Active Directory directory service. A Passport ticket contains the PUID. The initial mapping of PUIDs to accounts in Active Directory is called provisioning. Provisioning is described later in this topic.

The Active Directory attribute **altSecurityIdentity** is used to map the Passport user to an Active Directory account. The value of **altSecurityIdentity** looks like this: **kerberos:***0124ABCD78787878***@domain.com**. This essentially is a User Principal Name (UPN) where the PUID is the UPN prefix and everything after the "@" is the UPN suffix. The UPN suffix determines where the mapping lives. The UPN suffix should be the domain where the mapping is created or exists. In other words, in the sample mapping **kerberos:***0124ABCD78787878***@domain.com**, *domain.com* is the UPN suffix. The UPN suffix can be configured in the IIS authentication dialog. The UPN suffix value must be the same as the value configured in the IIS authentication dialog or the same as the domain name of the server.

The IIS metabase property AuthFlags contains the available settings for Windows authentication schemes and the file access authentication flags. By default, IIS 6.0 with .NET Passport authentication enabled tries to map the .NET Passport user to an account in Active Directory. This default behavior can create performance overhead for sites that do not want to use .NET Passport Active Directory mapping. The mapping behavior can be controlled by the **PassportRequireADMapping** flag. The flag values are:

- PassportRequireADMapping set to 0: IIS does not attempt Active Directory mapping. The request is handled as Anonymous User.
- PassportRequireADMapping set to 1 (default setting): In this situation the worker process must have Trusted Computing Base (TCB) privileges, which means that the worker process acts as part of the operating system. IIS then attempts to map to an Active Directory account (called LsaLogonUser). If this attempt fails, the request is handled as Anonymous User.
- PassportRequireADMapping set to 2: IIS enforces a mapping from the .NET Passport account to an Active Directory account before returning the requested Web page. If the mapping fails, IIS returns a 401 error.

For .NET Passport mapping to work, the account used for the IIS worker process must meet the following criteria:

- The IIS worker process must be running as LocalSystem if the computer is a domain member. To change the account on the worker process, see Configuring Worker Process Identities.
- If using a service account, the domain account must have TCB privilege on the local box (meaning, the account acts as part of the operating system).
- The account must be a member of the pre-Windows 2000 compatible access group, or have Read access to list group memberships.

The following WMI script sets the **PassportRequireADMapping** flag in the IIS metabase:

```
dim g_PPADMapping : g_PPADMapping = 0
call ValidateArgs(WScript.Arguments)
Set g_Locator = CreateObject("WbemScripting.SWbemLocator")
Set g_IIS = g_Locator.ConnectServer ("","root\MicrosoftIISv2","", "")
Set g_NodeObj = g_IIS.Get("IIsWebVirtualDirSetting='" & WScript.Arguments(0) & "'")
g_NodeObj.PassportRequireADMapping.Value = g_PPADMapping
g_NodeObj.Put_
WScript.Echo WScript.Arguments(0) & "/PassportRequireADMapping changed to: " &
g_PPADMapping
sub ValidateArgs(args)
      if args.Count <> 2 then
            call Usage
      end if
      if args(1) = "no" then
            g_PPADMapping = 0
```

```
      elseif args(1) = "try" then
            g_PPADMapping = 1
      elseif args(1) = "enforce" then
            g_PPADMapping = 2
      else
            call Usage
      end if
end sub
sub Usage
      WScript.echo "Usage: " & WScript.ScriptName & " metabase_path no|try|enforce"
      WScript.Echo "Example: " & WScript.ScriptName & """w3svc/1/root/passportvdir
enforce"""
      WScript.Echo " ------- Enforces mapping for virtual directory passportvdir"
      WScript.Quit
end sub
```

**Provisioning**

To establish ACLs on a Web page for .NET Passport authentication, PUIDs need to be mapped to a corresponding account in Active Directory. This process is called provisioning. The sample code in this section provisions PUIDs to Active Directory accounts. To provision the Active Directory, you must have administrator credentials on the Active Directory server. Normally this provisioning is accomplished by a .NET Passport-enabled Web page that collects information about the .NET Passport user and then saves the data so an administrator can perform the provisioning against Active Directory.

The following sample code is taken from Subscribe.asp. Use this code on your .NET Passport-enabled ASP page to collect .NET Passport user IDs and create Active Directory user accounts with the .NET Passport mapping.

```
%
 '
 ' Instantiate Passport Manager.
  '
 Set objPassportManager = Server.CreateObject("Passport.Manager.1")
 '
 ' Get any available information about the user from Passport.
 '
 strCurrentUserEmail = objPassportManager.Profile("PreferredEmail")
 strMemberIDHigh = objPassportManager.Profile("MemberIDHigh")
 strMemberIDLow = objPassportManager.Profile("MemberIDLow")
 strHex = Hex(strMemberIDHigh)
  '
 ' Convert the PUID into the correct form.
 '
 strPUID = String(8 - Len(strHex), "0") & strHex
 strHex = Hex(strMemberIDLow)
 strPUID = strPUID & String(8 - Len(strHex), "0") & strHex
 strLastName = objPassportManager.Profile("LastName")
 strFirstName = objPassportManager.Profile("FirstName")
 '
 ' Connect to Active Directory and find the default naming context for the default
domain.
 '
 set rootDSE = GetObject("LDAP://RootDSE")
 if Err.Number <> 0 then
  Response.Write "GetObject Error = " & Hex(Err.Number) & " (Source: " & Err.Source &
")"
 end if
 Err.Clear
 strDomainDN = rootDSE.Get("defaultNamingContext")
 if Err.Number <> 0 then
  Response.Write "Get Domain DN Error = " & Hex(Err.Number) & " (Source: " & Err.Source
& ")"
 end if
 Err.Clear
 strdnsName = rootDSE.Get("dnsHostName")
      FirstDot = Instr( strdnsName, "." )
      strDomainName = Right( strdnsName, Len( strdnsName ) - FirstDot )
 strContainerDN = "cn=users, " & strDomainDN

 ' Bind to the user's container.
 set cont = GetObject("LDAP://" & strContainerDN )
 if Err.Number <> 0 then
```

```
  Response.Write "GetObject Users Container Error = " & Hex(Err.Number) & " (Source: "
& Err.Source & ")"
 end if
 Err.Clear
 ' Create the user account.
 Set user = cont.Create("user", "cn=" & strCurrentUserEmail)
  if Err.Number <> 0 then
  Response.Write "Create User Error = " & Hex(Err.Number) & " (Source: " & Err.Source &
")"
 end if
 user.put "samAccountName", strPUID
 strFullAccountName = strPUID & "@" & strDomainName
  user.put "altSecurityIdentities", "Kerberos:" & strFullAccountName
      ' Set the UPN of the user to equal the PP e-mail address.
      user.put "userPrincipalName", strCurrentUserEmail
      user.put "givenName", Request.form("fname")(1)
      user.put "sn", Request.form("lname")(1)
 '
      '   IMPORTANT!!!
      ' Before going live with a site based on this sample code,
      ' the following must be changed so that a random password
      ' is assigned to each account. A recommended way to do this
      ' is to use the CAPICOM Utilities.GetRandom function.
 '
      user.put "userPassword", "89sd53ET!23kL"
 ' Write the property cache.
 user.SetInfo
 ' Refresh the property cache.
 user.GetInfo
 ' Clear the account disabled flag and set neccessary flags.
 flags = user.get("userAccountControl")
 flags = flags OR "H00010000
 flags = flags AND "HFFFFFFFD
 user.put "userAccountControl", flags
 user.SetInfo
%
```

**Configuring for Preproduction**

Real-world .NET Passport users probably do not have existing e-mail names within the PREP environment. Part of your site's development and testing effort might require that you first create a store of users within the PREP environment and then use these accounts for testing against the .NET Passport authentication and profile-access portions of your site's code. The .NET Passport server code that runs the services in the PREP environment is essentially identical to the code that runs *live* in the Production environment; Production and PREP are released in tandem. This means that *going live* should be a matter of simply stopping the servers on a site currently running against the PREP environment, reconfiguring Passport Manager on those servers to run against the Production environment instead, and then restarting the servers.

For more information about getting a PREP site ID, see "Registering Your .NET Passport Site" in the .NET Passport SDK (starting with version 2.1) Help. For more information about PREP Passports, see "Get a PREP Passport" and "Sign in to PREP Passport" in .NET Passport SDK (starting with version 2.1) Help. Also, see "Going Live: Deploying Passport Manager and Site Code."

Upon completion of the registration process, you will see a confirmation page that indicates your new .NET Passport sign-in name. This is the name you will use while in the development phase.

**Configuring for Production**

All Microsoft .NET Passport users already have an account in this environment. Users of Microsoft MSN Internet Access and MSN Explorer, as well as users who registered directly at the Passport.com Web site, also have passports in this environment. Within the Production environment, sites can expect to handle many real-world .NET Passport users who might already have been authenticated at various other Passport participating sites.

To deploy the .NET Passport code of your site, the Passport Manager server-side object must be reconfigured to work against the Production environment. To deploy against the Production environment, your site must also register for a site ID in the Production environment, and must supply some configuration information about your site that is specific to the Production deployment. Approval of a site ID in the Production environment might require some basic remote checking of your site's .NET Passport implementation by the .NET Passport team, and might also require some additional contracts or agreements.

**Important**

Do not place co-branded URLs in a virtual directory on which .NET Passport authentication is enabled. If a co-branded URL resides in a virtual directory on which .NET Passport authentication is enabled, client requests are challenged to authenticate before the co-branded information is passed to the client. This can result in a situation where a client might not be able to authenticate and view the desired page.

To deploy against the .NET Passport Production environment, see "Registering Your Passport Web Site" in the .NET Passport SDK (starting with version 2.1) Help. To disable the default installation environment and enable normal operation, enter a different site ID in the **Site ID** field and commit changes. Note that if you change the site ID to something other than 1, a matching encryption key must also be installed. The key, as well as the unique site ID assigned to any particular participating site, can be obtained only by registering as a .NET Passport participating site.

**Note**

Be sure to reconfigure your Windows XP clients for .NET Passport production mode as well. For more information, see ".NET Passport SDK: System Requirements" in the .NET Passport SDK (starting with version 2.1) Help. More specifically, read the section titled "Client Computer Requirements."

**Passport Manager Administration Utility**

The Microsoft Passport Manager Administration utility is a graphical user interface (GUI) alternative to editing the registry when changing Passport Manager object settings. The most common use of the Passport Manager Administration utility is to change the default settings. For details about how defaults can be set and affect the behavior of Passport Manager methods, see "Setting Passport Manager Defaults" in .NET Passport SDK Help.

**To access the Passport Manager Administration utility**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **msppcnfg.exe**, and click **OK**.

**Related Information**

- For helpful whitepapers that include information about everything from the benefits of enabling .NET Passport on your Web sites, to the .NET Passport Software Development Kit (SDK), see the Microsoft .NET Passport Web site.
- For more information about IIS and Web site security, see Security in IIS 6.0.

# Enabling .NET Passport Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When .NET Passport is enabled, requests coming into IIS must contain .NET Passport credentials either on the query string or within a cookie. The credentials also have to be valid, meaning the ticket has not expired. If IIS does not detect .NET Passport credentials, requests are redirected to the .NET Passport sign-in page.

.NET Passport uses cookies, which contain information that can be compromised. However, .NET Passport authentication can be used over a Secure Sockets Layer (SSL) connection, which reduces the potential of replay attacks.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To enable .NET Passport authentication on a Web site**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click the Web site on which you want to enable .NET Passport authentication, and click **Properties**.
2. Click the **Directory Security** tab.
3. In the **Anonymous access and authentication control** section, click **Edit**.
4. Select the **.NET Passport Authentication** check box. There are fundamental differences in the way .NET Passport validates user credentials, so .NET Passport cannot be used with other authentication methods. When .NET Passport authentication is selected, all other authentication methods are unavailable.
5. Click **OK**.

**Related Information**

- For comprehensive developer and technical support resources to help businesses implement .NET Passport services, including information about licensing and implementing .NET Passport, see the Microsoft .NET Passport for Developers Web site.
- For information about configuring .NET Passport authentication programmatically, see AuthFlags Metabase Property in the Metabase Property Reference.
- For more information about SSL, see Configuring SSL on a Web Server or Web Site.
- For more information about secure protocols, see Authentication in IIS 6.0.

# FTP Site Authentication in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure an FTP server to allow anonymous access to FTP resources or to require Basic authentication in the form of a user name and password corresponding to a valid Windows user account.

| **Note** |
| --- |
| Digest and Integrated Windows authentication cannot be used with FTP sites. Available authentication settings must be set at the site level for FTP sites. |

This section includes the following information:

- Configuring Anonymous FTP Authentication: Describes how to enable anonymous FTP authentication for an FTP site.
- Configuring Basic FTP Authentication: Describes how to enable basic FTP authentication for an FTP site.

# Configuring Anonymous FTP Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure your FTP server to allow anonymous access to FTP resources. If you select Anonymous FTP authentication for a resource, all requests for that resource are accepted without prompting the user for a user name or password. This is possible because IIS automatically creates a Windows user account called IUSR_*computername*, where *computername* is the name of the server on which IIS is running.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**.

**Procedures**

**To enable Anonymous FTP authentication**

1. If the IUSR_*computername* account is not used for Anonymous FTP authentication, you must create a Windows user account appropriate for the authentication method and add the account to a Windows user group.
2. Configure NTFS permissions for the directory or files for which you want to control access, using the user account you selected in step 1.
3. In IIS Manager, double-click the local computer; right-click the **FTP Sites** folder, an individual FTP site folder, a virtual directory, or a file; and then click **Properties**.

   **Note**

   Configuration settings made at the FTP Sites level are inherited by all of the FTP sites on the server. You can override inheritance by configuring the individual site or site element.

4. Click the **Security Accounts** tab.
5. Select the **Allow anonymous connections** check box.
6. To allow your users to gain access by Anonymous authentication only, select the **Allow only anonymous connections** check box.
7. In the **User name** and **Password** boxes, type the anonymous logon user name and password that you want to use, and then click **OK**. The user name is the name of the anonymous user account, which is typically designated as IUSR_*computername*.
8. Set the appropriate NTFS permissions for the anonymous account.

**Important**

If you change the security settings for your Web site or virtual directory, your Web server prompts you for permission to reset the security settings for the child keys of that site or directory. If you choose to accept these settings, the child keys inherit the security settings from the parent site or directory.

**Related Information**

- For information about creating Windows user accounts, see Best Practices for Securing Files with NTFS Permissions.
- For information about setting NTFS permissions for an anonymous account, see Setting NTFS Permissions for Directories or Files.

- For information about setting metabase properties, see Configuring the Metabase.

# Configuring Basic FTP Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To establish an FTP connection with your Web server by using the Basic FTP authentication method, users must log on with a user name and password corresponding to a valid Windows user account.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable the Basic FTP authentication method**

1. Create a Windows user account appropriate for the authentication method. If appropriate, add the account to a Windows user group.
2. Configure NTFS permissions for the directory or file for which you want to control access.
3. In IIS Manager, right-click the FTP site, directory, virtual directory, or file, and then click **Properties**.
4. Click the **Security Accounts** tab.
5. Clear the **Allow anonymous connections** check box.
6. Click **OK**.
7. Set the appropriate NTFS permissions for the account.

| **Important** |
| --- |
| If you change the security settings for your Web site or virtual directory, your Web server prompts you for permission to reset the security settings for the child nodes of that site or directory. If you choose to accept these settings, the child nodes inherit the security settings from the parent site or directory. |

**Related Information**

- For information about creating Windows user accounts, see Best Practices for Securing Files with NTFS Permissions.
- For information about setting NTFS permissions for an anonymous account, see Setting NTFS Permissions for Directories or Files.
- For information about setting metabase properties, see Configuring the Metabase.

# Access Control with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can control which users and computers are allowed to access your Web server and its resources. You can use both NTFS and Internet Information Services (IIS) security features, such as Web permissions and IP address restrictions, to specific access rights to Web sites, directories, and files.

This section includes the following information:

- Securing Files with NTFS Permissions: Describes how to use NTFS permissions to secure the files and directories that comprise your Web site.
- Securing Sites with Web Site Permissions: Describes how to set Web permissions to protect your Web content.
- Allowing Anonymous Access to Web Sites: Describes how to set up an account for anonymous access to Web sites.
- Configuring Worker Process Identities: Describes how to configure the worker process identity so an application can run with the appropriate access rights for your Web server.
- Securing Sites with IP Address Restrictions: Describes how to prevent specific computers, groups of computers, or entire networks from accessing your Web content.
- Securing Virtual Directories: Describes how to secure virtual directories.

**Related Information**

- For information about FTP user isolation, which prevents users from viewing or overwriting other users' Web content, see Isolating FTP Users.
- For information about configuring content rating features to warn users about objectionable Web content, see Using Content Ratings.
- For information about configuring your Web server's Secure Sockets Layer (SSL) security features, see Configuring SSL on a Web Server or Web Site.

# Securing Files with NTFS Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

It is strongly recommended that you use the NTFS file system, instead of the FAT or FAT32 file system, for your application server. With NTFS, you can limit access to your Web server's files and directories. You can also configure the access level granted to a particular user or group for the files and directories on your server.

This section includes the following information:

- Best Practices for Securing Files with NTFS Permissions: Provides best practices for configuring the files and folders that comprise your Web site to use NTFS security features.

- Setting NTFS Permissions for Directories or Files: Provides procedural information about configuring NTFS permissions, which define the level of file or directory access assigned to a Windows user account.

# Best Practices for Securing Files with NTFS Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Observe the following best practices when you set NTFS permissions:

- Assign permissions to groups rather than to users. Because it is inefficient to maintain user accounts directly, assigning permissions on a user basis should be the exception.
- If possible, avoid changing the default permission entries on file system objects, particularly on system folders and root folders. Changing default permissions can cause unexpected access problems or reduce security.
- Never deny the Everyone group access to an object. If you deny Everyone permission to an object, it includes administrators. A better solution is to remove the Everyone group, as long as you give other users, groups, or computers permissions to that object. You may also need to assign Full Control to the Administrators group and LocalSystem.
- Inherited Deny permissions do not prevent access to an object if the object has an explicit Allow permission entry. Explicit permissions take precedence over inherited permissions, including inherited Deny permissions.
- Deny permissions should be used only for the following special cases:
    - To exclude a subset of a group which has Allowed permissions.
    - To exclude one special permission when you have already assigned Full Control to a user or group.
- Use care when configuring NTFS permissions for your Web site. Inappropriately set permissions can deny valid users access to required files and directories. For example, even though a user has the correct user rights to view and execute a program, the user might not have permission to access a particular dynamic-link library (DLL) that is required to run that program. To guarantee users secure and uninterrupted file access, place related files in the same directory, and then assign the appropriate NTFS permissions to the directory.

**Related Information**

- For more information about NTFS permissions, see "Access Control" in Help and Support Center in Windows Server 2003.
- For more information about converting a volume from the FAT or FAT32 file system to NTFS, see "Convert" in Help and Support Center in Windows Server 2003.

# Setting NTFS Permissions for Directories or Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can control access to your Web site's directories and files by setting NTFS access permissions. You can use NTFS permissions to define the level of access that you want to grant to specific users, and groups of users. The appropriate configuration of file and directory permissions is crucial for preventing unauthorized access to your resources.

The procedures in this section explain how to secure Web sites using IIS Manager, which allows setting permissions on Web sites, directories, and virtual directories. Files inherit permissions from directories. To set permissions on individual files, use a programmatic solution, such as Adsutil.vbs IIS Administration Utility.

This section includes the following information:

- Securing a Web Site Using NTFS Permissions: Describes how to set NTFS permissions on a Web site.
- Securing a Web Site Using NTFS Special Permissions: Describes how to set NTFS special permissions on a Web site.

# Securing a Web Site Using NTFS Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can strengthen the security of a Web site by configuring NTFS permissions for directories, virtual directories, or the Web site itself.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To secure a Web site by using NTFS permissions**

1. In IIS Manager, expand the local computer, right-click a Web site or file, and then click **Permissions**.

   Do one of the following:

| Task | Procedure |
| --- | --- |
| Add a group or user that does not appear in the **Group or user names** list box. | a. Click **Add**.<br>b. In the **Enter the object name to select** box, type the name of the user or group, and then click **OK**. |
| Change or remove permissions from an existing group or user. | In the **Group or user names** list box, click the name of the group or user. |

2. To allow or deny a specific permission, in the **Permissions for** *User or Group* list box, select the **Allow** or **Deny** check box.

| **Important** |
| --- |
| Inherited Deny permissions do not prevent access to an object if the object has an explicit Allow permission entry. Explicit permissions take precedence over inherited permissions, including inherited Deny permissions. |

**Related Information**

- For more information about access control, see "Access Control" in Help and Support Center for Windows Server 2003.

# Securing a Web Site Using NTFS Special Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

With NTFS permissions, you can assign special permissions to groups or users. Special permissions are permissions on a more detailed level. For better management, you should assign broad-level permissions to users or groups, where it is applicable. For descriptions of permissions, see "Permissions for Files or Folders" in Help and Support Center for Windows Server 2003.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To secure a Web site by using NTFS special permissions**

1. In IIS Manager, expand the local computer, right-click a Web site, directory, or virtual directory, and then click **Permissions**.
2. Click the **Advanced** tab, and then do one of the following:

| To | Do this |
| --- | --- |
| Set special permissions for an additional group or user | Click **Add**, and in the **Enter the object name to select** box, type the name of the user or group, and then click **OK**. |
| View or change special permissions for an existing group or user | Click the name of the group or user and then click **Edit**. |
| Remove an existing group or user and its special permissions | Click the name of the group or user and then click **Remove**. If the **Remove** button is unavailable, clear the **Allow inheritable permissions from the parent to propagate to this object and all child objects. Include these with entries exclusively defined here.** check box, and then click **Remove**. Click **OK** and skip steps 3-6 below. |

3. In the **Permissions** box, select or clear the appropriate **Allow** or **Deny** check boxes.
4. In the **Apply onto** list box, click the folders or subfolders you want these permissions to be applied to.
5. To prevent the subfolders and files from inheriting these permissions, clear the **Apply these permissions to objects and/or containers within this container only** check box.
6. Click **OK** three times.

| **Important** |
| --- |
| You should assign permissions to the highest-level folders possible and then apply inheritance to propagate the settings to lower-level subfolders and files. For more information about inheritance, see "How Inheritance Affects File and Folder Permissions" in Help and Support Center for Windows Server 2003. |

**Related Information**

- For more information about access control, see "Access Control" in Help and Support Center for Windows Server 2003.

# Securing Sites with Web Site Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Web site permissions are not meant to be used in place of NTFS permissions. Instead, they are used with NTFS permissions to strengthen the security of your Web site content. You can configure your Web site's access permissions for specific sites, directories, and files. Unlike NTFS permissions, Web site permissions affect everyone who tries to access your Web site.

The following conditions apply to setting permissions:

- If Web site permissions conflict with NTFS permissions for a directory or file, the more restrictive settings are applied.
- Disabling permissions restricts all users. For example, disabling the Read permission restricts all users from viewing a file, regardless of the NTFS permissions applied to those users' accounts. However, enabling the Read permission can allow all users to view that file, unless NTFS permissions that restrict access have also been applied.
- If both IIS and NTFS permissions are set, the permissions that explicitly deny access take precedence over permissions that grant access.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set permissions for Web content (including WebDAV)**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** folder, an individual Web site folder, a virtual directory, or a file; and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element. |

2. On the **Home Directory**, **Virtual Directory**, or **File** property sheet, select or clear any of the following check boxes (if available):

   - **Read** (selected by default) Users can view directory or file content and properties.
   - **Write** Users can change directory or file content and properties.
   - **Script Source Access** Users can access source files. If Read is selected, then source can be read, if Write is selected, then source can be written to. Script Source Access includes the source code for scripts. This option is not available if neither Read nor Write is selected.

     | **Important** |
     | --- |
     | When you select **Script source access**, users might be able to view sensitive information, such as a user name and password. They might also be able to change source code that runs on your server, and thereby significantly affect your server's configuration and performance. |

   - **Directory browsing** Users can view file lists and collections.

- **Log visits** A log entry is created for each visit to the Web site.
- **Index this resource** Allows Indexing Service to index this resource. This allows searches to be performed on the resource.

3. Under **Execute Permissions** select the appropriate level of script execution:
   - **None** Do not run scripts or executables on the server.
   - **Scripts only** Run only scripts on the server.
   - **Scripts and Executables** Run both scripts and executables on the server.
4. Click **OK**.

**Related Information**

- For information about securing WebDAV, see Managing WebDAV Security.
- For information about setting NTFS permissions, see Securing Files with NTFS Permissions.

# Allowing Anonymous Access to Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Anonymous access, the most common Web site access control method, allows anyone to visit the public areas of your Web sites. In IIS 6.0, anonymous users are assigned by default to the IUSR_*computername* account, which is a valid Windows account that is a member of the **Guests** group. The IUSR_*computername* account can be defined on a computer or on a domain.

**Procedures**

**To create a new "anonymous logon" user account**

1. Change the account used for anonymous authentication. For more information, see Changing the Account Used for Anonymous Authentication.
2. Assign the following NTFS permissions on the Web site directory for the anonymous account: **Read & Execute**, **List Folder Contents**, and **Read**. When setting permissions on a Web site, be aware that all subfolders and files inherit these permissions by default. For more information about changing NTFS permissions on a Web site directory, see Setting NTFS Permissions for Directories or Files and "Access Control" in Help and Support Center for Windows Server 2003.

For security reasons, in order to run most executables (such as Cmd.exe) in the system folder, you must be a member of the Administrators group, the LocalSystem, Interactive, or Service account. Because this limits remote access to administrators, an anonymous user will not be able to run these executables (such as a CGI program).

| **Note** |
| --- |
| In IIS 6.0, the IUSR_*COMPUTERNAME* account has been denied write access to Web content by default. |

**Related Information**

- For information about creating a user account, see "Creating user and group accounts" in Help and Support Center for Windows Server 2003.
- For information about modifying user rights, see "Edit Local Security Settings" or "Edit a Security Setting on a Group Policy Object" in Help and Support Center for Windows Server 2003.

# Configuring Worker Process Identities

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In previous versions of IIS, the worker processes ran as the LocalSystem account. Because the LocalSystem account has access to almost all resources on the operating system, this had serious security implications. IIS 6.0 offers better security in that the worker process runs under the new built-in Network Service account by default. IIS 6.0 also allows you to configure which account worker processes run as. You have the option of using one of the three predefined accounts, or creating your own account.

This section includes the following information:

- Configuring a Worker Process Identity Using a Predefined Account: Describes how to specify a predefined account as the worker process identity of an application pool.
- Configuring a Worker Process Identity Using a Configurable Account: Describes how to create a new account as the worker process identity of an application pool.
- Configuring Applications to Run Independently of Each Other: Describes how to isolate applications by configuring them in separate application pools that run with different worker process identities.

# Configuring a Worker Process Identity Using a Predefined Account

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure a predefined account, such as Network Service, Local Service, or Local System, as the worker process identity for an application pool.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure a worker process identity for an application pool using a predefined account**

1. In IIS Manager, expand the local computer, expand the **Application Pools** folder, right-click the application pool you want to configure, and then click **Properties**.
2. Click the **Identity** tab.
3. Click **Predefined**, and in the list box beside it, click **Network Service**, **Local Service**, or **Local System**.

   | **Note** |
   | --- |
   | Network Service, Local Service and LocalSystem are members of the IIS_WPG group. |

4. Click **OK**.

**Related Information**

- For information about adding user accounts to groups, see "Add a Member to a Local Group" in Help and Support Center for Windows Server 2003.
- For information about application pools, see Configuring Application Pools in IIS 6.0.
- For information about configuring NTFS permissions, see Best Practices for Securing Files with NTFS Permissions.

# Configuring a Worker Process Identity Using a Configurable Account

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure a custom account to use as the worker process identity for an application pool.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure a worker process identity for an application pool using a configurable account**

1. In IIS Manager, expand the local computer, expand the **Application Pools** folder, right-click the application pool you want to configure, and then click **Properties**.
2. Click the **Identity** tab.
3. Click **Configurable**.
4. Click **Browse**, and under **Enter the object name to select**, type the account name in which you want your worker process to run, and then click **OK**.
5. In the **Password** box, type the password associated with this account. If there is no password associated with the account, leave the **Password** box blank.
6. Click **OK**.
7. Add the account you just created to the IIS_WPG group. For information about adding user accounts to groups, see "Add a Member to a Local Group" in Help and Support Center for Windows Server 2003.
8. If the account you just created must be able to start CGI processes, assign the following user rights to this account: **Adjust memory quotas for a process** and **Replace a process level token**. For information about configuring user rights for CGI applications, see Configuring CGI Applications.

**Related Information**

- For information about application pools, see Configuring Application Pools in IIS 6.0.
- For information about configuring NTFS permissions, see Best Practices for Securing Files with NTFS Permissions.

# Configuring Applications to Run Independently of Each Other

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

*Application isolation* is the separation of applications through process boundaries to prevent them from affecting one another.

For worker process isolation mode, you can configure isolation using application pools. You can either place single applications into their own application pool or group multiple applications into a single application pool.

**Procedures**

**To configure applications to run independently of each other**

1. Create a new account and configure a worker process identity for an application pool.
2. Add the account to the IIS_WPG group. The IIS_WPG group has **Read & Execute**, **List Folder Contents**, and **Read** permissions on every Web site directory by default. In a case where the account is not in the IIS_WPG group and does not have the appropriate permissions, the worker process will fail to start.

**Related Information**

- For information about adding user accounts to groups, see "Add a Member to a Local Group" in Help and Support Center for Windows Server 2003.
- For information about application pools, see Configuring Application Pools in IIS 6.0.
- For information about configuring NTFS permissions, see Best Practices for Securing Files with NTFS Permissions.

# Securing Sites with IP Address Restrictions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure your Web sites and FTP sites to grant or deny specific computers, groups of computers, or domains access to Web sites, FTP sites, directories, or files. For example, if your intranet server is connected to the Internet, you can prevent Internet users from accessing your server by granting access only to members of your intranet, and explicitly denying access to outside users.

| **Note** |
| --- |
| IP address restrictions apply only to IPv4 addresses. |

| **Note** |
| --- |
| In Windows 2003 Server Service Pack 1 (SP1), HTTP.sys will not cache content from a Web site with restricted IP addresses. In this case, you might consider using Internet Protocol Security (IPSec) instead of IP address restriction. For more information about IPSec, see Creating an IPsec Policy to Restrict Ports. |

This section includes the following information:

- Granting or Denying Access to a Computer: Describes how to configure IP address restrictions to grant or deny access to a computer.
- Granting or Denying Access to a Domain: Describes how to configure IP address restrictions to grant or deny access to a domain.
- Granting or Denying Access to a Group of Computers: Describes how to configure IP address restrictions to grant or deny access for a group of computers.

# Granting or Denying Access to a Computer

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the following procedure to grant or deny access to a specific computer.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To grant or deny access to a computer**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** or **FTP Sites** folder, an individual Web or FTP site, a virtual directory, or a file; and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web or FTP Sites level are inherited by all of the Web or FTP sites on the server. You can override inheritance by configuring the individual site or site element. |

2. Click the **Directory Security** or **File Security** tab, and then do one of the following:

   - For Web sites, in the **IP address and domain name restrictions** section, click **Edit**.
   - For FTP sites, continue to the next step.

3. Click **Granted access** or **Denied access**. When you select **Denied access**, you deny access to all computers and domains, except to those that you specifically grant access. When you select **Granted access**, you grant access to all computers and domains, except to those that you specifically deny access.
4. Click **Add**, and then click **Single computer**.
5. Click **DNS Lookup** to search for computers or domains by name, rather than by IP address.
6. Type the DNS name for the computer. IIS searches on the current domain for the computer, and if found, enters its IP address in the **IP address box**.

   The following information is important to remember when using the **DNS Lookup** feature:

   - It causes a performance decrease on your server while it is looking up DNS addresses.
   - A user accessing your Web server through a proxy server appears to have the IP address of the proxy server.
   - Some user server access problems can be corrected by entering the "*.*domainname*.com" syntax, rather than the "*domainname*.com" syntax.

7. Click **OK** three times.

**Related Information**

- For more information about IP address restrictions, see Setting IP Address and Domain Name Restrictions.
- For more information about subnet masks, see "Subnet masks" in the Windows Server 2003 Help and Support Center.

# Granting or Denying Access to a Domain

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the following procedure to grant or deny access to a specific domain.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To grant or deny access to a domain**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** or **FTP Sites** folder, an individual Web or FTP site, a virtual directory, or a file; and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web or FTP Sites level are inherited by all of the Web or FTP sites on the server. You can override inheritance by configuring the individual site or site element. |

2. Click the **Directory Security** or **File Security** tab, and then do one of the following:
   - For Web sites, in the **IP address and domain name restrictions** section, click **Edit**.
   - For FTP sites, continue to the next step.
3. Click **Granted access** or **Denied access**. When you select **Denied access**, you deny access to all computers and domains, except to those that you specifically grant access. When you select **Granted access**, you grant access to all computers and domains, except to those that you specifically deny access.
4. Click **Add**, and then click **Domain name**.
5. In the **Domain name** box, type the domain name, and then click **OK** three times.

**Related Information**

- For more information about IP address restrictions, see Setting IP Address and Domain Name Restrictions.
- For more information about subnet masks, see "Subnet masks" in the Windows Server 2003 Help and Support Center.

# Granting or Denying Access to a Group of Computers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A group of computers can be either denied or granted access based upon their network ID and a subnet mask.

For example, if the host computer has an IP address of 172.16.16.1 and a subnet mask of 255.255.0.0, all of the computers in that subnet would have IP addresses that began with 172.16. To select all of the computers in the subnet, enter 172.16.16.1 in the **Network ID** box and 255.255.0.0 in the **Subnet Mask** box.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the runas command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To grant or deny access to a group of computers**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** or **FTP Sites** folder, an individual Web or FTP site, a virtual directory, or a file; and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web or FTP Sites level are inherited by all of the Web or FTP sites on the server. You can override inheritance by configuring the individual site or site element. |

2. Click the **Directory Security** or **File Security** tab, and then do one of the following:
   - For Web sites, in the **IP address and domain name restrictions** section, click **Edit**.
   - For FTP sites, continue to the next step.
3. Click **Granted access** or **Denied access**. When you select **Denied access**, you deny access to all computers and domains, except to those that you specifically grant access. When you select **Granted access**, you grant access to all computers and domains, except to those that you specifically deny access.
4. Click **Add**, and then click **Group of computers**.
5. In the **Network ID** box, type the IP address of the host computer.
6. In the **Subnet mask** box, type the subnet ID for the computer you want grant or deny access to, and then click **OK** three times.

**Related Information**

- For more information about IP address restrictions, see Setting IP Address and Domain Name Restrictions.
- For more information about subnet masks, see "Subnet masks" in the Windows Server 2003 Help and Support Center.

# Securing Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can secure virtual directories by using local user authentication. IIS supports controlling access to directories and files on a local computer using all Windows authentication methods.

This section includes the following information:

- Securing a Virtual Directory: Describes how to create a secure virtual directory on the local computer.
- Changing Security on an Existing Virtual Directory: Describes how to change the security setting for an existing virtual directory.

# Securing a Virtual Directory

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure access permissions when you create a virtual directory.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a secure virtual directory mapped to resources on the local computer**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click the Web Site or folder within which you want to create the virtual directory, point to **New**, and then click **Virtual Directory**.
2. Click **Next**.
3. In the **Alias** box, type the alias for the new virtual directory, and then click **Next**.
4. In the **Path** box, type the path to the local directory, and then click **Next**.
5. Select the appropriate check boxes for the access permissions you want to allow, and clear the check boxes for the permissions you want to deny.
6. Click **Next** to complete the wizard.

**Related Information**

- For information about changing the security setting on an existing virtual directory, see Changing Security on an Existing Virtual Directory.
- For information about using virtual directories, see Using Virtual Directories.

# Changing Security on an Existing Virtual Directory

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To protect the contents of an existing virtual directory, you can change the access permissions on the virtual directory.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To change security on an existing virtual directory mapped to resources on the local computer**

1. In IIS Manager, right-click the virtual directory you want to secure, and click **Properties**.
2. On the **Virtual Directory** tab, under **Local Path**, select the check boxes for the access permissions you want to allow, and clear the check boxes for the permissions you want to deny.
3. Click **OK**.

**Related Information**

- For information about creating a secure virtual directory, see Securing a Virtual Directory.
- For information about using virtual directories, see Using Virtual Directories.

# URL Authorization in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

URL authorization helps simplify access management by authorizing user access to the URLs that comprise a Web application. IIS 6.0 works with Authorization Manager, a management tool available with the Windows Server 2003 operating systems, to implement IIS URL authorization. For information about Authorization Manager, see "Authorization Manager" in Help and Support Center for Windows Server 2003.

To use URL authorization in IIS 6.0 you must perform the following configuration tasks, in the order specified below:

- Disabling Anonymous Access for an Application, Virtual Directory, or Web Site Describes how to disable anonymous access for applications, virtual directories, or Web sites that use URL authorization.
- Enabling Wildcard Scriptmapping for an Application, Virtual Directory, or Web Site Describes how to enable the Urlauth.dll ISAPI interceptor.
- Adding Urlauth.dll as a New Web Service Extension Describes how to add Urlauth.dll to the list of Web service extensions.
- Creating an Authorization Policy Store in Authorization Manager Describes how to create a policy store in Authorization Manager for the URL.
- Setting the Scope for the Application Describes how to set the scope for URL authorization.
- Adding the IIS Worker Process to the Readers Role Describes how to configure the Readers role in Authorization Manager to accept the IIS worker process identity.
- Setting Metabase Properties for URL Authorization Describes how to set IIS metabase properties to ensure that URL authorization works correctly.

# Disabling Anonymous Access for an Application, Virtual Directory, or Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Integrated Windows authentication is the default authentication method for Windows Server 2003 operating systems. It is required for applications, virtual directories, or Web sites that use URL authorization. Anonymous access is not permitted.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc** *%systemroot%***\system32\inetsrv\iis.msc"**. |

**Procedures**

**To disable anonymous access for an application, virtual directory, or Web site**

1. In IIS Manager, expand the local computer and then expand the **Web Sites** folder.
2. Do one of the following:

   Expand the Web site that contains the application for which you want to configure URL authorization, right-click the application directory, and then click **Properties**.

   Or

   Expand the Web site that contains the virtual directory for which you want to configure URL authorization, right-click the virtual directory, and then click **Properties**.

   Or

   Right-click the Web site where you want to configure URL authorization, and then click **Properties**.
3. Click the **Directory Security** tab, and in the **Authentication and access control** section, click **Edit**.
4. Clear **Enable anonymous access**, ensure that **Integrated Windows authentication** is checked in the **Authenticated access** section, and then click **OK**.

**Related Information**

- For more information about Integrated Windows authentication, see Integrated Windows Authentication in IIS 6.0.

# Enabling Wildcard Scriptmapping for an Application, Virtual Directory, or Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

URL authorization uses *wildcard scriptmapping* as implemented by an ISAPI interceptor, Urlauth.dll. To use Urlauth.dll, you must first enable it for each Web site, application, or virtual directory where you want to configure URL authorization.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc** *%systemroot%***\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure wildcard scriptmapping for an application, virtual directory, or Web site**

1. In IIS Manager, expand the local computer and then expand the **Web Sites** folder.
2. Do one of the following:

   Expand the Web site that contains the application for which you want to configure URL authorization, right-click the application directory, and then click **Properties**.

   Or

   Expand the Web site that contains the virtual directory for which you want to configure URL authorization, right-click the virtual directory, and then click **Properties**.

   Or

   Right-click the Web site where you want to configure URL authorization, and then click **Properties**.
3. Click the **Directory** tab (for an application), the **Virtual Directory** tab (for a virtual directory), or the **Home Directory** tab (for a Web site), and then in the **Application settings** section, click **Configuration**.
4. Click the **Mappings** tab, and then in the **Wildcard application maps** section, click **Insert**.
5. In the **Add/Edit Application Extension Mapping** box, click **Browse** and browse to the **Windows\system32\inetsrv** directory.
6. Click **urlauth.dll**, click **Open**, and then click **OK** three times.

**Related Information**

- For more information about application mappings, see Installing Wildcard Application Mappings.

# Adding Urlauth.dll as a New Web Service Extension

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you want to use a Web service extension that is not in the default list of extensions, you must add the file to the list, and then enable the Web service extension.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc** *%systemroot%***\system32\inetsrv\iis.msc"**. |

**Procedures**

**To add Urlauth.dll as a new Web service extension**

1. In IIS Manager, expand the local computer, right-click the **Web Service Extensions** node, and then click **Add a new Web service extension**.
2. In **Extension Name**, type the following text exactly as it appears below:

   ```
   URL Authorization
   ```

   and then click **Add**. If you do not type this text exactly as it appears, IIS URL authorization fails without sending a warning message.
3. Click **Browse** on the **Add file** dialog box, browse to *systemroot***\system32\inetsrv**, select urlauth.dll, click **Open**, and then click **OK** twice.
4. Check **Set extension status to Allowed**, and then click **OK**.

**Related Information**

- For more information about application mappings, see Installing Wildcard Application Mappings.
- For more information about web service extensions, see Enabling and Disabling Dynamic Content in IIS 6.0.

# Creating an Authorization Policy Store in Authorization Manager

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

An authorization store is a database that stores Authorization Manager policy.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc** *%systemroot%***\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create an authorization policy store in Authorization Manager**

1. Click **Start**, click **Run**, type

   ```
   Azman.msc
   ```

   and then click **OK**.

2. In the console tree, right-click **Authorization Manager** and then click **Options**.
3. Confirm that **Developer mode** is selected, and then click **OK**.
4. In the console tree, right-click **Authorization Manager** and then click **New Authorization Store**.
5. Click **XML file** and, in **Store name**, type the name of the store, for example, C:\MyStore.xml.
6. Optionally, in **Description**, type information about the new store, and then click **OK**.
7. In the console tree, right-click the store file you just created and then click **New Application**.
8. In **New Application**, in **Name**, type the following text exactly as it appears below:

   ```
   IIS 6.0 URL Authorization
   ```

   If you do not type this text exactly as it appears, IIS URL authorization fails without sending a warning message.

9. Optionally, in **Description**, type information about this usage of IIS URL Authorization, and then click **OK**.
10. In the console tree, double-click **Authorization Manager**, double-click the store name, double-click the application name (IIS 6.0 URL Authorization), double-click **Definitions**, right-click **Operation Definitions**, and then click **New Operation Definition**.
11. In **New Operation Definition**, in **Name**, type the following text exactly as it appears below:

    ```
    AccessURL
    ```

    In **Operation number**, type the numeral **1**, and then click **OK**. If you do not type this text exactly as it appears, IIS URL authorization fails without sending a warning message.

**Related Information**

- For more information about Authorization Manager, see Help and Support Center in Windows Server 2003.

# Setting the Scope for the Application

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A *scope* is a subdivision within an application that separates selected resources from other resources that are used by the application. Scopes can help to prevent unintended resource sharing, as well as to support auditing and delegation.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc** *%systemroot%***\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set the scope for the application**

1. If Authorization Manager is not already open, click **Start**, click **Run**, type

   ```
   Azman.msc
   ```

   and then click **OK**.

2. In the console tree, right-click **IIS 6.0 URL Authorization** and then click **New Scope**.

3. In **Name**, type the following text exactly as it appears below:

   ```
   WebApp
   ```

   and then click **OK**. If you do not type this text exactly as it appears, IIS URL authorization fails without sending a warning message.

4. In the console tree, double-click **IIS 6.0 URL Authorization**, double-click **Definitions**, right-click **Role Definitions**, and then click **New Role Definition**.

5. In **New Role Definition**, in **Name**, type the following text exactly as it appears below:

   ```
   Viewer
   ```

   and then click **Properties**.

   If you do not type this text exactly as it appears, IIS URL authorization fails without sending a warning message.

6. Click the **Definition** tab, click **Add**, and then click the **Operations** tab.

7. Select the **AccessURL** checkbox and then click **OK** twice.

8. In the console tree, under **IIS 6.0 URL Authorization**, under **WebApp**, right-click **Role Assignments**, and then click **Assign Roles**.

9. Select the **Viewer** checkbox and then click **OK**.

10. In the details pane, right-click the role assignment you created above, and then click **Assign Windows Users and Groups**.

11. In **Enter the object names to select (examples)**, type your user name, and then click **OK**.

**Related Information**

- For more information about Authorization Manager, see Help and Support Center in Windows Server 2003.

# Adding the IIS Worker Process to the Readers Role

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, IIS runs in the Network Service account. You can, however, configure an IIS worker process to run in a different account.

If you use a remote authorization store, such as Active Directory or a remote XML file-based store, and you run IIS in the default Network Service context, you must add the Active Directory account of the Web server running IIS to the store's Readers role.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc** *%systemroot%***\system32\inetsrv\iis.msc"**. |

**Procedures**

### To add the IIS worker process to the Readers role

1. If Authorization Manager is not already open, click **Start**, click **Run**, type

   ```
   Azman.msc
   ```

   and then click **OK**.

2. In the console tree, right-click **Authorization Manager**, click **Open Authorization Store**, click **Browse**, click the name of the authorization store file you want to update, click **Open**, and then click **OK**.

3. In the console tree, right-click the name of the store, and then click **Properties**.

4. Click the **Security** tab, and in the **Authorization manager user role** list, click **Reader**, and then click **Add**.

5. In the **Enter the object names to select (examples)** box, type the name of the IIS worker process, and then click **OK** twice.

**Related Information**

- For more information about application mappings, see Installing Wildcard Application Mappings.

- For more information about configuring a worker process account, see Configuring a Worker Process Identity Using a Configurable Account.

# Setting Metabase Properties for URL Authorization

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To use URL authorization, you must set the required metabase properties on the application, virtual directory, or Web site where you want to authorize user access. You can either edit the metabase directly or run a script to set metabase properties.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile/User:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To set metabase properties for URL authorization by editing the metabase directly**

- Edit the metabase directly to set the following metabase properties:
  - AzEnable Metabase Property: Enables URL authorization for the virtual directory, application, or URL that corresponds to the entry in the metabase.
  - AzStoreName Metabase Property: Associates an Authorization Manager store with the virtual directory, application, or URL.
  - AzScopeName Metabase Property: Associates the virtual directory, application, or URL with a scope. This scope will be the name of a scope in the IIS 6.0 URL authorization application in the Authorization Manager policy store referred to in the **AzStoreName** attribute. If no scope or an empty string is specified, the default scope of the IIS 6.0 URL authorization will be used.
  - AzImpersonationLevel Metabase Property: Determines the impersonation behavior for the application. This allows you to configure the Web application to impersonate the client user, the IIS worker process, or the IUSR_*computername* account for the worker process. Each setting significantly changes the environment and implied design of the Web application.

The sample script, written in Microsoft Visual Basic® Scripting Edition (VBScript), marks the root of the first site as a URL in MyAZScope, which is defined in the MyAZStore.xml file. Users with URLAccess rights in this scope will be able to access the site.

**To set metabase properties for URL authorization by using a script**

- Use the following sample script as a guide to setting the required metabase properties:

```
var objVDir = GetObject("IIS://localhost/w3svc/1/root")

objVDir.AzEnable = true

objVDir.AZStoreName = "MSXML://d:\MyAZStore.xml"

objVDir.AzScopeName = "MyAZScope"

objVDir.AZImpersonationLevel = 0

objVDir.SetInfo()
```

**Related Information**

- For more information about the IIS Metabase, see Working with the Metabase.

# IIS 6.0 Encryption

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use encryption to help protect the private information that clients exchange with your Web server, such as credit card numbers or phone numbers. You can also encrypt data when running WMI-based remote administration scripts. In IIS, the foundation for encryption is the Secure Sockets Layer (SSL) protocol, which establishes an encrypted communication link with users.

This section contains the following information:

- Configuring Secure Sockets Layer: Describes how to configure a Web server or Web site to use SSL for secure connections.
- Setting Encryption Strength: Describes how to configure encryption strength in IIS Manager.
- Encrypting Data When Running WMI–Based Remote Administration Scripts: Describes how to modify WMI-based administration scripts to establish an encrypted remote connection.

# Configuring Secure Sockets Layer

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section contains the following information:

- Configuring SSL on a Web Server or Web Site: Describes how to configure a Web server or Web site to use SSL for secure connections.
- Redirecting Bookmarked Requests to SSL-Enabled Content: Describes how to prevent 404 errors by redirecting requests from http:// to https:// URLs.
- Enabling Kernel-Mode SSL: Describes how to enable kernel-mode SSL for a Web Server.

# Configuring SSL on a Web Server or Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure Secure Sockets Layer (SSL) security features on a Web server or a Web site to verify the integrity of your content, verify the identity of users, and encrypt network transmissions.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To configure SSL on a Web server or a Web site**

1. In IIS Manager, double-click the local computer, and then double-click the **Web Sites** folder.
2. Right-click the Web site or file that you want to protect with SSL, and then click **Properties**.
3. Under **Web site identification** click **Advanced**.
4. In the **Advanced Web site identification** box, under **Multiple identities for this Web site**, verify that the Web site IP address is assigned to port 443, the default port for secure communications, and then click **OK**. Optionally, to configure more SSL ports for this Web site, click **Add** under **Multiple identities of this Web site**, and then click **OK**.
5. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Edit**.
6. In the **Secure Communications** box, select the **Require secure channel (SSL)** check box.
7. To enable SSL client certificate authentication and mapping features, select the **Enable client certificate mapping** check box, click **Edit**, add the 1-to-1 or many-to-1 mappings you need, and then click **OK** three times.

**Note**

If you set your Web site to require SSL, as in step 6 above, and you have not completed setting up SSL for the site, then users browsing your site will receive this error: "HTTP Error 403.4 - Forbidden: SSL is required to view this resource." To avoid this condition, either complete all the steps in the list above, or go back and clear the **Require Secure Channel (SSL)** check box (see step 6).

**Related Information**

- For information about enabling client certificates, see Enabling Client Certificates in IIS 6.0.
- For information about client certificate mapping, see Mapping Client Certificates to User Accounts in IIS 6.0.

# Redirecting Bookmarked Requests to SSL-Enabled Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you update a Web site configuration to require SSL, users who have previously bookmarked pages on the site will receive a 403.4 "Forbidden: SSL required" error message when they request those pages after SSL is configured. You can prevent this situation by creating a custom error page that automatically redirects the http: call to an https: call.

The example in the following procedure uses ASP, but you can use any dynamic content handler. Because dynamic content handlers are not enabled by default when IIS is installed, you should ensure that the appropriate Web service extension is enabled.

**Procedures**

**To redirect bookmarked requests to SSL-enabled content**

1. Create an ASP page called RedirectToSSL.asp and save it in the root directory of your Web site.
2. Include redirection code similar to the following code example. Note that this code example does not include any error-handling code.

```
<%@Language=VBScript %>
<%
strQueryString = Request.QueryString
strAddress = Right(strQueryString, Len(strQueryString) - Instr(1,strQueryString,
";"))
strSecure = Replace(strAddress, "http:", "https:", 1, 1)
Response.Redirect strSecure
%>
```

3. In IIS Manager, right-click the SSL-enabled Web site and then click **Properties**.
4. In the **Secure communications** section on the **Directory Security** tab, click **Edit**.
5. Click **Require secure channel (SSL)** and then click **OK**.
6. On the **Custom Errors** tab, select the 403;4 error message and then click **Edit**.
7. In the **Message Type** field on the **Edit Custom Error Properties** dialog box, choose **File**, click **Browse** and browse to the location of the RedirectToSSL.asp page, and click **OK** twice.
8. Navigate to the location of the RedirectToSSL.asp page, right-click the page, and then click **Properties**.
9. In the **Secure communications** section on the **File Security** tab, click **Edit**.
10. Clear **Require secure channel (SSL)** and then click **OK** twice.

**Related Information**

- For information about enabling web service extensions, see Enabling Web Service Extensions.
- For more information about SSL, see About Secure Sockets Layer.

# Enabling Kernel-Mode SSL

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Enabling kernel-mode Secure Sockets Layer (SSL) improves SSL performance because encryption and decryption operations take place in the faster processing space of the kernel. However, enabling kernel-mode SSL disables user-mode SSL. For more information, see Kernel-Mode SSL.

| **Important** |
| --- |
| Using Registry Editor incorrectly can cause serious problems that require reinstalling the operating system. Because Registry Editor bypasses the standard safeguards that prevent you from entering settings that are conflicting or likely to degrade performance or damage your system, exercise caution when making changes to the registry. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. |

Note that you should back up the registry before you edit it. If you are running Microsoft® Windows NT® Server or a newer server operating system from Microsoft, you should also update your Emergency Repair Disk (ERD).

For information about how to edit the registry, see the "Changing Keys and Values", "Add and Delete Information in the Registry", and "Edit Registry Data" topics in Registry Editor Help.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer***\Administrator cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To configure SSL to run in kernel mode**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type Regedit.exe, and click **OK**.
3. Navigate to and double-click the following key in the registry:

   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\HTTP\Parameters
4. From the **Edit** menu, point to **Add**, click **DWORD Value** and then add the following registry value:

   Name: **EnableKernelSSL**

   Type: **REG_DWORD**

   Data: Set this to 1 to use kernel-mode SSL instead of user-mode SSL.
5. Quit Registry Editor.
6. Restart IIS.

**Related Information**

- For information about restarting an IIS service, see Restarting IIS.
- For information about configuring SSL, see Configuring SSL on a Web Server or Web Site.

# Configuring SSL Host Headers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Organizations that host multiple Web sites on a single server often use host headers to create multiple Web sites without requiring a unique IP address for each site. For more information about hosting multiple Web sites on a single server, see Hosting Multiple Web Sites.

You can configure Web sites that use host headers to serve protected content over a Secure Sockets Layer (SSL) connection, that is, a connection that uses https:// instead of http://. To use SSL with host headers, you must obtain and install a wildcard server certificate. After you configure SSL host headers for a Web site, protected content is served only over an https:// connection.

| **Important** |
| --- |
| The following conditions apply to the use of SSL host headers: |

- SSL host headers cannot be configured by using the IIS Manager UI.
- Using SSL host headers requires that the wildcard certificate be installed on each Web site from which you want to serve protected content. This adds overhead to site management, because you must manually ensure that multiple sites are kept in sync with each other.
- You must configure secure bindings for each Web site that uses the wildcard server certificate to prevent unauthorized use of that certificate.

This section includes the following information:

- Obtaining and Installing a Wildcard Server Certificate: Describes how to request a wildcard server certificate and install it on a Web site.
- Configuring Server Bindings for SSL Host Headers: Describes how to configure the **SecureBindings** metabase property to create SSL host headers for each Web site that you want to enable to use the wildcard server certificate.
- Ensuring That Secure Content Is Served Over HTTPS Only: Describes how to prevent unauthorized Web sites from using the wildcard certificate.

# Obtaining and Installing a Wildcard Server Certificate

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A wildcard server certificate works with many SSL sites, because a wildcard (*) is used to stand for the host header name for each of the sites. All SSL-enabled Web sites that use the same IP/port binding and are distinguished only by their host header names must use the same wildcard server certificate.

| **Important** |
| --- |
| Not all Certification Authorities (CAs) issue wildcard certificates. |

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To obtain a wildcard server certificate**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.
2. Right-click the Web site for which you want to obtain a wildcard server certificate, and then click **Properties**.
3. On the **Directory Security** tab, under **Secure communications**, click **Server Certificate**.
4. In the Web Server Certificate Wizard, click **Create a new certificate**.
5. Follow the Web Server Certificate Wizard, which will guide you through the process of requesting a new server certificate. On the **Your Site's Common Name** page, type a name in the **Common name** box, using the following format:

   *.< *sitename* >

   for example, *.contoso.com.

   By default, the certificate request file is saved as C:\Certreq.txt, but the wizard allows you to specify a different location.
6. Click **Finish** to complete the wizard.

After you receive the wildcard server certificate from the Certification Authority (CA), assign the certificate on all Web sites that have the same IP/port binding and are distinguished only by host header name.

**Related Topics**

- For information about installing a server certificate, see Installing Server Certificates.
- For general information about certificates, see "SSL and Certificates" in Managing a Secure IIS 6.0 Solution.

# Configuring Server Bindings for SSL Host Headers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In addition to installing a wildcard server certificate on SSL-enabled Web sites, you must also configure the **SecureBindings** metabase property on each site so it contains the host header name of the site. All SSL-enabled sites that use the same IP/port binding and are distinguished only by host header name must use the same wildcard server certificate.

| **Important** |
| --- |
| You must configure secure bindings for all SSL-enabled Web sites that use the wildcard server certificate to prevent unauthorized use of the certificate. |

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To configure the SecureBindings metabase property for SSL host headers**

1. Click **Start**, click **Run**, type

   ```
   cmd
   ```

   in the **Open** box, and then click **OK**.

2. Type the following command at the command prompt:

   `cscript.exe adsutil.vbs set /w3svc/` *<site identifier>* `/SecureBindings ":443:`*<host header>* `"`

   where *host header* is the host header for the Web site, for example, site2.contoso.com or site4.contoso.com.

**Related Topics**

- For information about preventing unauthorized use of a wildcard server certificate, see Ensuring That Secure Content Is Served Over HTTPS Only.
- For information about the **SecureBindings** metabase property, see SecureBindings Metabase Property.
- For information about Adsutil.vbs, see Using the Adsutil.vbs Administration Script.

# Ensuring That Secure Content Is Served Over HTTPS Only

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

After a certificate is installed on a site, it can be served over HTTP or HTTPS. To ensure that only SSL requests are served, you must configure the **AccessSSL** metabase property to force SSL content requests only.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To force SSL content requests only**

1. Click **Start**, click **Run**, type

   ```
   cmd
   ```

   and then click **OK**.

2. Type the following command at the command prompt:

   ```
   cscript.exe adsutil.vbs set /w3svc/< site identifier >/AccessSSL TRUE
   ```

   where <*site identifier*> is the unique number that identifies the site.

**Related Topics**

- For information about configuring server bindings for SSL host headers, see Configuring Server Bindings for SSL Host Headers.
- For information about Adsutil.vbs, see Using the Adsutil.vbs Administration Script.
- For information about **AccessSSL**, see AccessSSLFlags Metabase Property.

# Setting Encryption Strength

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure your Web server to require a 128-bit minimum session-key strength for all Secure Sockets Layer (SSL) communications. This is the default session-key strength for Microsoft Windows Server 2003.

If you set a minimum 128-bit key strength, users attempting to establish a secure communications channel with your server must use a browser capable of communicating with a 128-bit session *key*. The session key is not the same as an SSL *key pair*, which is used to negotiate and establish a secure communication link. For information about upgrading browsers to 128-bit encryption capability, see How to Upgrade Internet Explorer to 128-Bit Encryption on the Windows Support Web site.

To establish encrypted communications, you must have a valid server certificate installed.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**.

**Procedures**

**To set encryption strength**

1. In IIS Manager, double-click the local computer, and then right-click the Web site, directory, or file that you want and click **Properties**.
2. On the **Directory Security** or **File Security** tab, under **Secure Communications**, click **Edit**.
3. In the **Secure Communications** box, select the **Require secure channel (SSL)** check box.
4. If 128-bit encryption is required, select the **Require 128-bit Encryption** check box.
5. Click **OK**.

**Note**

If you open a Server Gated Cryptography (SGC) certificate, you might receive a notice on the **General** tab that reads as follows: "The certificate has failed to verify for all of its intended purposes." This notice is issued because of the way SGC certificates interact with Windows, and does not necessarily indicate that the certificate does not work correctly.

**Related Information**

- For information about server and client certificates, see Certificates_IIS_SP1_Ops.
- For information about Server-Gated Cryptography (SGC), an extension of SSL that allows financial institutions with export versions of IIS to use strong encryption, see Server-Gated Cryptography.
- For information about SSL encryption, see About Secure Sockets Layer.

# Encrypting Data When Running WMI–Based Remote Administration Scripts

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

WMI-based administration scripts operating on an IIS server over a remote connection send data over the network in clear text. You can enhance security by modifying WMI-based administration scripts to establish an encrypted remote connection.

| **Important** |
| --- |
| For Windows Server 2003 with Service Pack 1 (SP1), WMI-based remote administration scripts must be modified to encrypt authentication credentials and data as outlined in this topic. If you do not modify the scripts, they will fail. |

Use one of the following two methods to establish an encrypted remote connection:

- Use the **IWbemLocator::ConnectServer** method when an administrator user name and password must be passed to the remote server.
- Use the **GetObject** function when an administrator user name and password do not need to be passed to the remote server. This is the case when the administrator user name and password on the remote server and on the server running the script are the same.

To establish an encrypted connection to a remote server when an administrator user name and password must be provided, use the **ConnectServer** method to connect to the server, and then set the **AuthenticationLevel** property to encrypt the connection. The following code example illustrates this.

| **Note** |
| --- |
| The **ConnectServer** call encrypts the administrator user name and password. |

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

```
var locatorObj=new ActiveXObject("WbemScripting.SWbemLocator");
provObj=locatorObj.ConnectServer(Server_Name,"root/MicrosoftIISv2",
AdministratorUserName, AdministratorPassword);
var wbemAuthenticationLevelPktPrivacy=6;
provObj.Security_.AuthenticationLevel=wbemAuthenticationLevelPktPrivacy;
```

To establish an encrypted connection to a remote server when an administrator user name and password do not need to be provided, use the **GetObject** function, and then set the **AuthenticationLevel** property to **pktPrivacy**. The following code example illustrates this.

```
provObj=GetObject( "winmgmts:
{AuthenticationLevel=pktPrivacy}!//Server_Name/root/MicrosoftIISv2");
```

**Related Information**

- For more information about security settings in WMI scripts, see the Scripting API for WMI on MSDN Online.

# Creating an IPsec Policy to Restrict Ports

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Internet Protocol security (IPsec) is designed to encrypt data as it travels between two computers, protecting it from modification and interpretation. To use IPsec, you must define how computers that attempt to connect will trust each other and how the computers will secure their traffic. To implement these specifications, create and apply an IPsec policy. IPsec policies support policy-based stateful packet filtering rules that can be used with IPsec authentication and encryption to provide robust end-to-end security.

**Example: Creating an IPsec Policy**
This section provides step-by-step instructions for creating an example IPsec policy that allows inbound requests to connect only to ports 80 and 443. Port 80 is the default port for HTTP requests, and port 443 is the default port for HTTPS requests. The policy blocks requests for all other ports on the server.

In contrast with typical firewall or packet filtering rules, there is no way to order the list of rules in an IPsec policy. The rule engine matches traffic with rules according to specificity. If a packet matches more than one rule, the engine will apply the most specific rule to the packet. In the following example, packets that match the filter list that allows connections to port 80 also match the filter list that blocks all incoming traffic. Because the former list is more specific, the rule engine uses that list to make its decision. Therefore, traffic that connects to port 80 or port 443 is passed to the server, while everything else is blocked. You need to create multiple filter lists to make the policy effective.

There are several steps to creating an IPsec policy:

1. Create IPsec filter lists. Filter lists include ports, protocols, and directions and trigger a decision when traffic matches an item in the list. This example requires you to create three filter lists. The first two filter lists apply to inbound traffic that attempts to connect to port 80 or 443, and the third filter list applies to all ports.

2. Set filter actions. Filter actions are the required response when traffic matches a filter list. This example uses only permit and block actions for its IIS IPsec policy.

3. Create the IPsec policy. The IPsec policy is a collection of rules, which are the correlation of a filter list with a filter action. You can have only one active, or assigned, policy at any particular time.

The following steps are for creating IPsec policies using the Local Security Settings snap-in in Windows Server 2003. You can also create IPsec policies using the **Netsh** commands for IPsec. For more information about **Netsh** commands for IPsec, see IPsec.

To create IPsec filter lists, begin by changing the security settings on the local computer.

**To create the IPsec filter lists**

1. From the **Start** menu, point to **All Programs**, point to **Administrative Tools**, and then select **Local Security Policy**.

2. In the **Local Security Settings** dialog box, click **IP Security Policies on Local Computer**. The right pane displays the default Windows Server 2003 policies.

3. Right-click the right pane, and then click **Manage IP filter lists and filter actions**.

4. In the **Manage IP filter lists and filter actions** dialog box, on the **Manage IP Filter Lists** tab, click **Add**.

5. In the **IP Filter List** dialog box, in the **Name** box, type a name for your filter list, such as Inbound HTTP, and type a description if you want one. This is the filter list that applies to all inbound HTTP connections.

6. Click **Add**. The IP Filter Wizard appears. Create a filter list with the following specifications:

    - Description: Optional.

    - Source address: **Any IP Address**.

- Destination address: **My IP Address**. Or click **A specific IP address**, and then type the IP address of the interface that is connected to the Internet.
- Protocol type: **TCP**.
- Protocol Port: **From any port**.
- To this port: **80**.

7. In the **Completing the IP Filter Wizard** screen, clear the **Edit properties** check box, and then click **Finish**.
8. In the **IP Filter List** dialog box, click **OK**.
9. Repeat steps 1 through 8, this time to create a filter list that applies to destination port 443 for HTTPS connections. Name the filter list Inbound HTTPS, or something similar.
10. Repeat steps 1 through 8, this time to create a filter list that applies to all destination ports. This filter list will apply to the policy that blocks all inbound traffic. Name the filter list All Inbound, or something similar.

After you create the IPsec filter list, you must create a filter action that will occur when incoming traffic matches the criteria in your filter lists. For this example, two actions are necessary. The first is a permit action for allowing requests to connect to ports 80 and 443. The action to permit traffic is an existing default, so you do not need to create it. You must create the second action, which is to block traffic to all other ports.

**To set the filter actions**

1. After you create the filter lists, in the **Local Security Settings** dialog box, right-click the right pane, and then click **Manage IP filter lists and filter actions**.
2. In the **Manage IP filter lists and filter actions** dialog box, click the **Manage Filter Actions** tab, and then click **Add**. The Filter Action Wizard appears.
3. Create a filter action with the following specifications:
   - Name: **Block**
   - Description: Optional
   - Filter Action General Options: **Block**
4. In the **Completing the IP Security Filter Action Wizard** screen, clear the **Edit properties** check box, and then click **Finish**.
5. In the **Manage IP filter lists and filter actions** dialog box, click **Close**.

After you create the IPsec filter lists and filter actions, you must create the IPsec policies and define the rules that link the lists to the actions.

**To create the IPsec policies**

1. After you set the filter actions, in the **Local Security Settings** dialog box, right-click the right pane, and then select **Create IP Security Policy**. The IP Security Policy Wizard appears.
2. Create a policy with the following specifications:
   - Name: **Packet Filter**
   - Description: Optional
3. In the **Requests for Secure Communication** screen, clear the **Activate the default response rule** check box, and then click **Next**.

4. In the **Completing the IP Security Policy Wizard** screen, make sure that the **Edit properties** check box is selected, and then click **Finish**. The **New IP Security Policy Properties** dialog box appears. Do not close this dialog box.

After you create the IPsec filter lists, filter actions, and IPsec policies, you must add rules to the policy and associate the IPsec filter lists that you created with the rules.

**To add rules to the policy**

1. After you create the IPsec policies, in the **New IP Security Policy Properties** dialog box, click **Add**. The Security Rule Wizard appears.
2. Create a rule with the following specifications:
   - Tunnel endpoint: **This rule does not specify a tunnel**.
   - Network type: **All network connections**.
   - IP filter lists: **All Inbound** (or the name of your filter list that applies to inbound traffic on any port).
   - Filter Action: **Block**.
3. In the **Completing the Security Rule Wizard** screen, clear the **Edit properties** check box, and then click **Finish**.
4. In the **New IP Security Policy Properties** dialog box, click **Close**.
5. Repeat steps 1 through 4 to create rules that apply to the Inbound HTTP and Inbound HTTPS filter lists. For these rules, choose the **Permit** filter action.

After you create the IPsec filter lists, filter actions, and policies and add rules to the policy and associate the IPsec filter lists with rules, you must assign the IPsec policy to the server. Only one policy can be assigned at any time.

**To assign the IPsec policy**

- In the **Local Security Settings** dialog box, right-click the policy you created, and click **Assign**.

The policy is applied immediately, and IPsec starts processing packets according to the rules in the policy. You do not need to reboot the server. To stop the policy, right-click the policy in the **Local Security Settings** dialog box, and click **Un-assign**.

# Certificates_IIS_SP1_Ops

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

Certificates are a form of digital identification for your server, and for clients that are requesting information from your server. Their function is similar to that of a passport, or other official identity card, which identifies the person carrying it. Certificates are part of the Secure Sockets Layer (SSL) features of Internet Information Services (IIS) that establish a secure connection over which sensitive information can be sent.

**This section includes the following information:**

- Server Certificates
    - Obtaining Server Certificates: Describes how to obtain a server certificate.
    - Installing Server Certificates: Describes how to install a server certificate.
    - Backing Up Server Certificates: Describes how to back up a server certificate.
    - Determining the Intermediate Root CA Version on a Web Server: Describes how to determine the intermediate Root CA version on a Web server.
    - Configuring a List of Trusted Certification Authorities: Explains how to add new certification authorities to your server's list of trusted issuing authorities.
    - Managing Server Certificates Programmatically in IIS 6.0: Describes samples that use the certificate component to import, copy, save, and back up server certificates.
- Client Certificates
    - Obtaining Client Certificates in IIS 6.0: Describes how to obtain user digital identifications, called client certificates.
    - Enabling Client Certificates in IIS 6.0: Describes how to enable SSL authentication of client certificates.
    - Backing Up Client Certificates in IIS 6.0: Describes how to create a copy of a client certificate.
    - Checking the Status of Client Certificates in IIS 6.0: Explains how to verify the validity of your users' client certificates.
    - Mapping Client Certificates to User Accounts in IIS 6.0: Describes how to map client certificates in one of three ways.

**Related Information**

- For information about regulating how users access your Web server files and directories, see Access Control with IIS 6.0.
- For information about setting up or changing encryption settings, see IIS 6.0 Encryption.

# Obtaining Server Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can obtain server certificates from a third-party certification authority (CA), which may require you to provide proof of identity before issuing a certificate. You can also issue your own server certificates by using an online CA, such as Microsoft Certificate Services, which installs the certificate at the time of your request. When you use the Server Certificate Wizard to obtain a server certificate, the process is referred to as creating a server certificate.

This section includes the following information:

- Requesting a New Server Certificate from an Online CA: Describes how to request a server certificate by using the Web Server Certificate Wizard.
- Obtaining a Server Certificate from a Third-party CA: Describes how use a server certificate request to obtain a server certificate from a certification authority.
- Issuing Your Own Server Certificates: Describes how to act as your own certification authority to issue server certificates for the Internet or for corporate intranets.

# Requesting a New Server Certificate from an Online CA

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the Web Server Certificate Wizard, available through IIS Manager, to request a new server certificate that will be installed immediately by an online CA.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**. |

**Procedures**

**To request a new server certificate from an online CA**

1. In IIS Manager, double-click the local computer, and then double-click the **Web Sites** folder.
2. Right-click the Web site or file for which you want to request a certificate, and then click **Properties**.
3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.
4. In the Web Server Certificate Wizard, click **Create a new certificate**.
5. On the **Delayed or Immediate Request** wizard page, select **Send the request immediately to an online certification authority.**
6. Complete the rest of the steps in the Web Server Certificate Wizard and then click **Finish**.

**Related Information**

- For information about other ways to obtain certificates, see Obtaining Server Certificates.
- For general information about certificates, see SSL and Certificates.
- For information about online certification authorities, see "Certificates Services" in Help and Support Center for Windows Server 2003.

# Obtaining a Server Certificate from a Third-party CA

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the Web Server Certificate wizard to obtain a server certificate from a third-party certification authority (CA). The wizard generates a certificate request, which you then send to the CA. Some certification authorities require you to prove your identity before they process your request or issue a certificate.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To obtain a server certificate from a third-party certification authority**

1. In IIS Manager, double-click the local computer, and then double-click the **Web Sites** folder.
2. Right-click the Web site or file for which you want to request a certificate, and then click **Properties**.
3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.
4. In the Web Server Certificate Wizard, on the **Delayed or Immediate Request** page, click **Prepare the request now, but send it later**. By default, the certificate request file is saved as C:\Certreq.txt, but the wizard allows you to specify a different location.
5. Complete the rest of the steps in the Web Server Certificate Wizard and then click **Finish**.
6. Send the request to the certification authority. The CA will process the request and then send you the certificate.

**Related Information**

- For information about other ways to obtain server certificates, see Obtaining Server Certificates.
- For information about installing server certificates, see Installing Server Certificates.
- For general information about certificates, see SSL and Certificates.

# Issuing Your Own Server Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use Microsoft Certificate Services to create a customizable service for issuing and managing certificates. You can create server certificates for the Internet or for corporate intranets, giving your organization complete control over certificate management policies. For more information, see "Certificate services" in Help and Support Center for Windows Server 2003.

Online requests for server certificates can be made only to local and remote Enterprise Certificate Services and remote standalone Certificate Services. The IIS Web Server Certificate Wizard does not recognize a standalone installation of Certificate Services on the same computer when requesting a certificate. To get around this, use the offline certificate request to save the request to a file and then process as an offline request. For information about standalone certification authorities, see "Standalone certification authorities" in Help and Support Center for Windows Server 2003.

If you are not using an online certification authority, you will need to save the request file generated by the Web Server Certificate Wizard to disk and send it to the CA. When the response is received, you can start the wizard and it will begin where it left off. If you are replacing a certificate, IIS will continue to use the old certificate until the new request is completed. For information about using the Web Server Certificate Wizard to save a request file, see Obtaining a Server Certificate from a Third-party CA.

**Related Information**

- For information about other ways to obtain certificates, see Obtaining Server Certificates.
- For general information about certificates, see SSL and Certificates.

# Installing Server Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

After you have obtained a server certificate, you can install it. When you use the Server Certificate Wizard to install a server certificate, the process is referred to as *assigning* a server certificate.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To install a server certificate using the Web Server Certificate Wizard**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.
2. Right-click the Web site or file that you want, and then click **Properties**.
3. On the **Directory Security** or **File Security** tab, under **Secure communications**, click **Server Certificate**.
4. In the Web Server Certificate Wizard, click **Assign an existing certificate**.
5. Follow the Web Server Certificate Wizard, which will guide you through the process of installing a server certificate.

| **Note** |
| --- |
| When you use the Web Server Certificate Wizard to assign a certificate, you must specify a password before the certificate can be assigned to your Web server. |

**Related Information**

- For information about obtaining server certificates, see Obtaining Server Certificates.
- For general information about certificates, see SSL and Certificates.

# Backing Up Server Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the Certificate Manager MMC snap-in to back up server certificates.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

#### Procedures

### To back up a server certificate by using Certificate Manager

1. If you do not have **Certificate Manager** installed in the MMC, install it.
2. Locate the correct certificate store. This is typically the **Local Computer** store in Certificate Manager.
3. In the **Personal** store, click the certificate that you want to back up.
4. On the **Action** menu, point to **All tasks**, and then click **Export**.
5. In the Certificate Manager Export Wizard, click **Yes, export the private key**.
6. Follow the wizard default settings, and enter a password for the certificate backup file when prompted.

   | **Important** |
   | --- |
   | Do not select **Delete the private key if export is successful**, because this will disable your current server certificate. |

7. Complete the wizard to export a backup copy of your server certificate.

### To add Certificate Manager to the MMC

1. To open an MMC console:
   - From the **Start** menu, click **Run**.
   - In the **Open** box, type the following: **mmc**.
   - Click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. In the **Add/Remove Snap-in** box, click **Add**.
4. In the **Available Standalone Snap-ins** list, click **Certificates**, and then click **Add**.
5. In the **Certificates** snap-in box, click **Computer Account**, and then click **Next**.
6. In the **Select Computer** box, click **Local Computer**, and then click **Finish**.

#### Related Information

- For general information about certificates, see Server Certificates.

# Determining the Intermediate Root CA Version on a Web Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The VeriSign Global Server ID Intermediate Root certification authority (128-bit SSL) expired on January 7, 2004. Servers running IIS that have not been updated with the new Global Server ID Intermediate Root certification authority may encounter problems when clients try to establish SSL sessions by using the Secure Hypertext Transfer Protocol after January 7, 2004.

| **Note** |
| --- |
| Servers that use VeriSign 40-bit Secure Server ID certificates are not affected by this expiration. |

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To determine the Intermediate Root CA version on a Web server**

1. From the **Start** menu, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add/Remove Snap-in** dialog box, click **Add**, select **Certificates** from the snap-in list, and then click **Add**.
4. Choose the **Computer account** option and then click **Next**.
5. Select **Local Computer (the computer this console is running on)**, and then click **OK**.
6. Click **Close**, and then click **OK**.
7. In the console window, expand **Certificates**, and then expand **Intermediate Certificate Authorities**.
8. Click the **Certificates** folder.
9. In the **Issued To** column, locate the certificate with "www.verisign.com/CPS Incorp.by Ref. LIABILITY LTD. ©97 Verisign" as the certificate name.
10. Right-click the certificate, and then click **Open**.

    If the certificate has expired, the following message is displayed on the **General** tab:

    "This certificate has expired or is not yet valid."

    If the certificate is current, the following message is displayed on the **General** tab:

    "This certificate is intended for the following purposes:"

    and lists the purposes for which the certificate is to be used.
11. Note the **Valid from date to date** range, and then click **OK**.
12. If you need to update the Intermediate Root CA on your Web server, follow the instructions on the Verisign Web site.

# Configuring a List of Trusted Certification Authorities

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure computers running a member of the Windows Server 2003 family with IIS 6.0 to accept certificates from a predefined list of certification authorities (CAs). You can then automatically verify client certificates against your CTL. CTLs are especially useful for Internet Service Providers (ISPs) who have several Web sites on one server and need a different list of approved certification authorities for authenticating clients at each site.

Each of your Web sites can be configured to accept certificates from a different list by using different certificate trust lists (CTLs). For example, an intranet administrator creates a different list of trusted certification authorities (CAs) for each department's Web site on an intranet, in the form of a CTL. As a result, IIS only accepts certificates from certification authorities that are on the department's CTL. When members of the department log on with a client certificate from a CA on the department's CTL, they are automatically authenticated.

This section includes the following information:

- Creating a Certificate Trust List on Your Server: Describes how to create a certificate trust list by using the Certificate Trust List Wizard.

- Modifying a Certificate Trust List on Your Server: Describes how to modify a certificate trust list.

# Creating a Certificate Trust List on Your Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the Certificate Trust List Wizard to create CTLs and to add new root certificates to your CTLs. By default, IIS installs the most commonly used CA root certificates.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To create a certificate trust list on your server**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.

   | **Note** |
   | --- |
   | CTLs can be applied only at the Web site level, not at the virtual directory or file level. CTLs are not available on FTP sites. |

2. Right-click the Web site that you want, and then click **Properties**.
3. On the **Directory Security** tab, under **Secure Communications**, click **Edit**.

   | **Note** |
   | --- |
   | The **Edit** button is enabled only if a server certificate has been installed on the site. For information about obtaining and installing a server certificate, see Obtaining Server Certificates and Installing Server Certificates. |

4. In the **Secure Communications** box, select the **Enable certificate trust list** check box, and then click **New**.
5. Follow the steps in the Certificate Trust List Wizard, which will guide you through the process of creating a CTL.

**Related Information**

- For general information about certificates, see SSL and Certificates.

# Modifying a Certificate Trust List on Your Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the Certificate Trust List Wizard to edit existing CTLs.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To modify a certificate trust list on your server**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.
2. Right-click the Web site that you want, and then click **Properties**.
3. On the **Directory Security** tab, under **Secure Communications**, click **Edit**.
4. In the **Secure Communications** box, scroll through the **Current CTL** list, click the CTL that you want to modify, and then click **Edit**.
5. Follow the steps in the Certificate Trust List Wizard, which will guide you through the process of modifying a CTL.

**Related Information**

- For general information about certificates, see SSL and Certificates.

# Managing Server Certificates Programmatically in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use scripts in conjunction with **IISCertObj**, a Component Object Model (COM) object, to manage certificates remotely and programmatically. You can use **IISCertObj** to perform the following tasks:

- Import copies of a certificate from a .pfx file (a file containing a PFX-encrypted certificate and a private key) from a server onto multiple servers.
- Save backup copies of certificates on a central site.
- Copy a certificate from one server onto multiple servers.

To learn about obtaining SSL server certificates, see Obtaining Server Certificates.

**Sample Scripts**

After acquiring an SSL server certificate from a certification authority, distributing copies to all of the servers requires one or more of the following steps:

- Import copies of master certificates onto multiple servers.
- Save backup copies of certificates in a central archive.
- Export certificates across multiple servers.

In IIS, if you use Windows Script Host (WSH) or Active Server Pages (ASP), you can execute scripts that communicate with **IISCertObj** programmatically to import, archive, or export certificates.

**Import Certificates to Multiple Servers (CertImport.vbs)**

Large sites often need to use one SSL certificate for secure user logon to multiple servers. It would be very time-consuming to use the MMC to add copies of the certificate to each server. What you need is a scripted solution that installs copies of the same certificate to all of the targeted servers.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

Save the following script as CertImport.vbs. Modify the command-line statement arguments to match your network resources. Then use the command-line statement to import a certificate from a certificate store on one server to other servers.

```
Option Explicit
Dim iiscertobj, pfxfile, pfxfilepassword, InstanceName, WebFarmServers, IISServer
Set iiscertobj = WScript.CreateObject("IIS.CertObj")
pfxfile = WScript.Arguments(0)
pfxfilepassword = WScript.Arguments(1)
InstanceName = WScript.Arguments(2)
WebFarmServers = split(WScript.Arguments(3), ",")
iiscertobj.UserName = WScript.Arguments(4)
iiscertobj.UserPassword = WScript.Arguments(5)
For Each IISServer in WebFarmServers
  iiscertobj.ServerName = IISServer
  iiscertobj.InstanceName = InstanceName
  iiscertobj.Import pfxfile, pfxfilepassword, true, true
Next
```

Command-line statement:

```
Certimport.vbs cert.pfx pfxpassword w3svc/1 iisserver1,iisserver2,iisserver3 Administrator aal34290
```

**Save Certificates in a Central Archive (Save_all_certs.vbs)**
Using the **IISCertObj** export method, you can archive a backup of each certificate on your server farm to a central site.

Save the following script as Save_all_certs.vbs in your text editor. Modify the command-line statement arguments to match your network resources. Then use the command-line statement to export copies of certificates to the central site.

```
Option Explicit
Dim iiscertobj, targetServer, targetServers, pfxbasename, pfxpassword, InstanceName
Set iiscertobj = WScript.CreateObject("IIS.CertObj")
pfxbasename = WScript.Arguments(0)
pfxpassword = WScript.Arguments(1)
InstanceName = WScript.Arguments(2)
targetServers = split(WScript.Arguments(3), ",")
iiscertobj.UserName = WScript.Arguments(4)
iiscertobj.UserPassword = WScript.Arguments(5)
iiscertobj.InstanceName = InstanceName
For Each targetServer in targetServers
  iiscertobj.ServerName = targetServer
  iiscertobj.Export pfxbasename + targetServer + ".pfx", pfxpassword, true, false,
false
Next
```

Command-line statement:

```
Save_all_certs.vbs C:\certbackup\ adsf-0324 w3svc/1 iisserver2,iisserver3,iisserver4 Administrator aal34290
```

**Copy a Certificate from an Existing Server to a New Server (CertCopy.vbs)**
With the script and command-line statement below you can use the **copy** method of **IISCertObj** to copy a certificate to a new server after you add the server to a server farm.

Save the script below as certcopy.vbs. Modify the command-line statement arguments to match your network resources. Then run the command-line statement.

```
Option Explicit
Dim iiscertobj, targetServer, targetServers, targetInstance
Set iiscertobj = WScript.CreateObject("IIS.CertObj")
iiscertobj.ServerName = WScript.Arguments(0)
iiscertobj.Instancename = WScript.Arguments(1)
targetServers = split(WScript.Arguments(2), ",")
targetInstance = WScript.Arguments(3)
iiscertobj.UserName = WScript.Arguments(4)
iiscertobj.UserPassword = WScript.Arguments(5)
For Each targetServer in targetServers
  iiscertobj.Copy true, true, targetServer, targetInstance
Next
```

Command-line statement:

```
Certcopy.vbs iisServer1 w3svc/1 iisServer2 w3svc/1 Administrator asdf-0324
```

# Obtaining Client Certificates in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can obtain a client certificate from a *certification authority* (CA), which is a trustworthy third-party organization. Before the CA issues a certificate, you may need to verify your identity. The identification information required by the CA varies, depending on the type of certificate that you want to obtain. How these certificates are obtained and implemented depends upon the browser that is being used. The procedures in this topic are for Microsoft Internet Explorer version 3.0 and later only. For client certificate procedures for other browsers, see the browser documentation.

**Procedures**

### To obtain a client certificate

1. Choose a certification authority. For information about choosing a certification authority, see Obtaining and Installing Server Certificates.
2. Refer to the specific ordering instructions that are provided by the certification authority.
3. Order your certificate. Some client certificates are valid only for the computer and Web browser that are used to generate the certificate request. When this is the case, be sure to make the online request from the computer where you intend to install the certificate.
4. When you receive your client certificate, use your Web browser's security features to install the certificate. (Some Web browsers may refer to client certificates as "browser certificates" or "personal certificates.")

**Related Information**

- For information about enabling client certificates, see Enabling Client Certificates in IIS 6.0.
- For information about backing up client certificates, see Backing Up Client Certificates in IIS 6.0.

# Enabling Client Certificates in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can require users attempting to access your Web site to log on with a client certificate. Requiring a client certificate is just one aspect of protecting your server against unauthorized access. Any user with a valid and trusted client certificate can establish a secure connection and access your resource. To protect your Web content from unauthorized access you must do one of the following:

- Use Basic, Digest, or Integrated Windows authentication, in addition to requiring a client certificate.
- Create a Windows account mapping for client certificates. For more information, see Mapping Client Certificates to User Accounts in IIS 6.0.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures, or you must have been delegated the appropriate authority. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc** *systemroot***\system32\inetsrv\iis.msc"**. For information about delegating administrative authority, see "Delegating administration" in Help and Support Center for Windows Server 2003. |

**Procedures**

**To enable client certificates**

1. In IIS Manager, double-click the local computer, and then right-click the Web site, directory, or file that you want and click **Properties**.
2. If you have not previously obtained a server certificate, click the **Directory Security** tab, and then under **Secure Communications**, click **Server Certificate**. For more information, see Obtaining Server Certificates.
3. If you have previously obtained a server certificate, click the **Directory Security** or **File Security** tab, and then under **Secure Communications**, click **Edit**.
4. In the **Secure Communications** box, select the **Require secure channel (SSL)** check box. Requiring a secure channel means that users cannot connect to this site without using a secure link (that is, the link's URL must begin with https://).
5. Under **Client certificates** select one of the following to enable client Certificate authentication:
   - **Accept client certificates** Users can access the resource with a client certificate, but the certificate is not required.
   - **Require client certificates** The server will request a client certificate before connecting the user to the resource. Users without a valid client certificate will be denied access.
   - **Ignore client certificates** Users with or without a client certificate will be granted access.

**Related Information**

- For information about obtaining client certificates, see Obtaining Client Certificates in IIS 6.0.
- For information about backing up client certificates, see Backing Up Client Certificates in IIS 6.0.

# Backing Up Client Certificates in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use the Certificate Export Wizard to create a backup copy of a certificate.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

### To back up a client certificate

1. On the **Tools** menu in Internet Explorer, version 5.0 or later, click **Internet Options**, and then click the **Content** tab.
2. On the **Content** tab, click **Certificates**, and then click the **Personal** tab.
3. On the **Personal** tab, in the list of certificates, click the certificate that you want to back up, and then click **Export**.
4. In the Certificate Export Wizard, click **Next**.
5. On the next page in the wizard, click **No, do not include any private keys in the export**, and then click **Next**.
6. On the next page in the wizard, click **Base64 Encoded X.509 (*.CER)**, and then click **Next**.
7. Follow the remaining steps in the Certificate Export Wizard to complete the procedure.

**Related Information**

- For information about obtaining client certificates, see Obtaining Client Certificates in IIS 6.0.
- For information about enabling client certificates, see Enabling Client Certificates in IIS 6.0.
- For information about checking the status of client certificates, see Checking the Status of Client Certificates in IIS 6.0.

# Checking the Status of Client Certificates in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Certification authorities (CAs) cannot physically revoke your users' client certificates. However, they may publish Certificate Revocation Lists (CRLs) that are copied onto your computer, where you can search them for client certificates that are in revoked status. For information about retrieving CRLs from a CA and storing them on your computer, see "Retrieve a certificate revocation list" in Help and Support Center for Windows Server 2003.

The metabase properties that control CRL checking can be set or viewed using a COM object, or WMI scripts, or ADSI scripts. For information about configuring the metabase, see Configuring the Metabase.

To use a certificate revocation list to check the status of client certificates, take the following actions:

- Enable CRL checking.
- Optionally, configure the CRL on your computer to refresh at a fixed interval, even when the CRL on your computer is still valid.
- Optionally, change the default time interval for refreshing the CRL on your computer at a fixed interval.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To enable and disable CRL checking**

- Set the CertCheckMode Metabase Property. CRL Checking is enabled by default. The CRL will be refreshed by the CA when a new CRL is issued, unless you intervene by setting a CRL refresh interval.

**To set the CRL refresh interval**

- Change the setting of the RevocationFreshnessTime Metabase Property. You can refresh the CRL on your computer with the CRL on the CA server even when the cached CRL on your computer is valid.

**To change the default interval to a custom time**

- Set the RevocationURLRetrievalTimeout Metabase Property to a time in milliseconds.

**Related Information**

- For information about obtaining client certificates, see Obtaining Client Certificates in IIS 6.0.
- For information about enabling client certificates, see Enabling Client Certificates in IIS 6.0.
- For information about backing up client certificates, see Backing Up Client Certificates in IIS 6.0.
- For more information about WMI and ADSI, see IIS Administration Technologies on MSDN.

# Mapping Client Certificates to User Accounts in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Mapping client certificates allows you to automatically authenticate users who log on with client certificates, without requiring the use of other supported authentication methods such as Basic, Digest, or Integrated Windows authentication.

This section includes the following information:

- Mapping Client Certificates One-to-One: Describes how to map individual client certificates to Windows user accounts.
- Mapping Client Certificates Many-to-One: Describes how to use wildcard-matching rules to map more than one client certificate to a Windows user account.
- Mapping Client Certificates with Directory Service Mapping: Describes using native Windows Active Directory features to authenticate users with client certificates.

| **Note** |
| --- |
| The client certificate mapping programs for both IIS 6.0 and Windows Active Directory Services accept certificates formatted as either binary or Base-64 encoded files. |

# Mapping Client Certificates One-to-One

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

One-to-one mapping maps individual client certificates to Windows user accounts. The server compares the copy of the client certificate it holds with the client certificate that is sent by the browser. The two must be identical for the mapping to proceed. If a client gets another certificate containing all of the same user information, it must be mapped again.

This section includes the following information:

- Exporting a Client Certificate for One-to-One Mapping: Describes how to export a client certificate so it can be used for IIS one-to-one mapping.

- Mapping a Specific Client Certificate to a User Account: Describes how to map a specific client certificate one-to-one to a Windows user account.

# Exporting a Client Certificate for One-to-One Mapping

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Some client certificates need to be *exported* for use in IIS one-to-one mapping. For more information about whether you must export client certificates, contact your certification authority.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To export a client certificate for one-to-one mapping**

1. On the **Tools** menu in Internet Explorer, version 5.0 or later, click **Internet Options**, and then click the **Content** tab.
2. On the **Content** tab, click **Certificates**, and then click the **Personal** tab.
3. On the **Personal** tab, in the list of certificates, click the certificate that you want to export, and then click **Export**.
4. In the Certificate Export Wizard, click **Next**.
5. On the next page in the wizard, click **No, do not include any private keys in the export**, and then click **Next**.
6. On the next page in the wizard, click **Base64 Encoded X.509 (*.CER)**, and then click **Next**.

   | **Note** |
   | --- |
   | The client certificate mapping programs for both IIS 6.0 and Windows Active Directory Services accept certificates formatted as either binary or Base-64 encoded files. |

7. Follow the remaining steps in the Certificate Export Wizard to complete the procedure.
8. Map the exported certificate one-to-one with a Windows user account. For information about mapping a certificate one-to-one with a Windows user account, see Mapping a Specific Client Certificate to a User Account.

**Related Information**

- For information about how to add a client certificate mapping using wildcard rules, see Mapping Client Certificates Many-to-One.
- For information about when to use each type of client certificate mapping, see Mapping Strategies.
- For general information about certificates, see SSL and Certificates.

# Mapping a Specific Client Certificate to a User Account

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the following procedure to map a specific client certificate to a user account.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To map a specific client certificate to a user account**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.
2. Right-click the Web site for which you want to configure authentication, and then click **Properties**.
3. Click the **Directory Security** tab, and then, in the **Secure Communications** section, click **Edit**.
4. In the **Secure Communications** box, select the **Enable client certificate mapping** check box, and then click **Edit**.
5. In the **Account Mappings** box, click the **1-to-1** tab.
6. On the **1-to-1** tab, either add a new certificate by clicking **Add**, or edit an existing mapping by selecting the mapping and clicking **Edit Map**.
7. If you are adding a new certificate, browse to the certificate file and open it.

   | **Note** |
   | --- |
   | If you cannot find the certificate file, it might first need to be exported. For information about exporting a certificate for use in one-to-one mapping, see Exporting a Client Certificate for One-to-One Mapping. |

8. In the **Map to Account** box, enter a map name for the mapping. This is the name that will be displayed in the selection list on the **Account Mappings** box.
9. Either type or browse to a Windows user account. Type the password of the account to which the certificate is being mapped.
10. Click **OK**.
11. Repeat these steps to map other certificates, or to map this certificate to other accounts.

**Related Information**

- For information about how to map a client certificate by using wildcard rules, see Mapping Client Certificates Many-to-One.
- For information about when to use each type of client certificate mapping, see Mapping Strategies.
- For general information about certificates, see SSL and Certificates.

# Mapping Client Certificates Many-to-One

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Many-to-one mapping uses *wildcard* matching rules that verify whether a client certificate contains specific information, such as issuer or subject. This mapping does not compare the actual client certificate, but rather accepts all of the client certificates that fulfill the specific criteria. If a client gets another certificate containing all of the same user information, the existing mapping will work.

When using many-to-one mapping, keep the following information in mind:

- Specific client certificate mappings always take precedence over wildcard mappings.
- Some client certificates offer more identifying information and may contain additional custom subfields. For information about certificate formats, contact your certification authority.

This section includes the following information:

- Adding a Client Certificate Mapping Using Wildcard Rules: Describes how to use wildcard rules to map more than one client certificate to a Windows user account.
- Editing Existing Wildcard Rules for Client Certificate Mapping: Describes how to edit existing wildcard rules.

# Adding a Client Certificate Mapping Using Wildcard Rules

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you create wildcard rules to map client certificates, you should use matching rules that are as specific as possible. A good wildcard rule matches information from several different fields and subfields. For example, the names "Accounting", "Shipping", and "Sales" can appear in the organization unit subfield of more than one company's client certificate. A matching rule that maps certificates based exclusively on this subfield would probably result in unintended mappings.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To add a client certificate mapping using wildcard rules**

1. In IIS Manager, double-click the local computer, and then double-click the **Web Sites** folder.
2. Right-click the Web site for which you want to configure wildcard certificate mapping, and then click **Properties**.
3. Click the **Directory Security** tab, and then in the **Secure Communications** section, click **Edit**.
4. In the **Secure Communications** box, select the **Enable client certificate mapping** check box, and then click **Edit**.
5. In the **Account Mappings** box, click the **Many-to-1** tab, and then click **Add**.
6. In the **General** box, type a name for the rule. This is the name that will be displayed in the selection list on the **Account Mappings** box. You can create rules for future use or disable rules without deleting them by selecting or clearing the **Enable this wildcard rule** check box. Click **Next**.
7. In the **Rules** box, click **New**.
8. In the **Edit Rule Element** box, select the appropriate criteria and click **OK**.

   **Note**

   Steps 6 and 7 can be repeated to define the rule more stringently.

9. When you have finished editing the rule, click **Next**.
10. In the **Mapping** box, either type or browse to a Windows user account. Type the password of the account to which the rule is being mapped.

    **Note**

    If the account you are mapping to is on a computer that is a member of a workgroup, you must specify the computer name and the account name. For example, if you are mapping to the "RegionalSales" account on the computer called "Sales1", the mapping account name would be "Sales1\RegionalSales".

11. Click **OK**.
12. Repeat these steps to create other mapping rules.
13. Use the **Move Up** and **Move Down** buttons to establish the precedence given to the rules. Rules that are higher in the list take precedence.

**Related Information**

- For information about how to map a specific client certificate to a user account, see Mapping Client Certificates One-to-One.
- For information about when to use each type of client certificate mapping, see Mapping Strategies.
- For general information about certificates, see SSL and Certificates.

# Editing Existing Wildcard Rules for Client Certificate Mapping

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the following procedure to edit existing wildcard rules for client certificate mapping.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To edit an existing wildcard rule for client certificate mapping**

1. In IIS Manager, expand the local computer, and then expand the **Web Sites** folder.
2. Right-click the Web site that you want to configure authentication for, and then click **Properties**.
3. Click the **Directory Security** tab, and then in the **Secure Communications** section, click **Edit**.
4. In the **Secure Communications** box, select the **Enable client certificate mapping** check box, and then click **Edit**.
5. In the **Account Mappings** box, click the **Many-to-1** tab.
6. On the **Many-to-1** tab, in the list of rules, click the rule that you want to modify, and then click **Edit Rule**.
7. Make the necessary changes, and when you are finished editing the rule, click **OK**.

**Related Information**

- For information about how to map a specific client certificate to a user account, see Mapping Client Certificates One-to-One.
- For information about mapping a client certificate to an Active Directory user account, see Mapping Client Certificates with Directory Service Mapping.
- For information about when to use each type of client certificate mapping, see Mapping Strategies.
- For general information about certificates, see SSL and Certificates.

# Mapping Client Certificates with Directory Service Mapping

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| If you use directory service mapping, you cannot use one-to-one or many-to-one mapping for the entire World Wide Web Publishing Service (WWW service). |

Directory service (DS) certificate mapping uses native Windows Active Directory features to authenticate users with client certificates. For information about mapping a certificate to an Active Directory user account, see "Map a certificate to a user account" in Help and Support Center for Windows Server 2003.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

### To enable Directory Service mapping

1. In IIS Manager, right-click the **Web Sites** node, and then click **Properties**.
2. On the **Directory Security** tab, select the **Enable the Windows directory service mapper** check box, and then click **OK**.

**Related Information**

- For information about how to map a specific client certificate to a user account, see Mapping Client Certificates One-to-One.
- For information about how to map many user accounts to a single client certificate, see Mapping Client Certificates Many-to-One.
- For information about when to use each type of client certificate mapping, see Mapping Strategies.
- For general information about certificates, see SSL and Certificates.

# Auditing in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can use security auditing techniques to track the activities of users and to detect unauthorized attempts to access your NTFS file system directories and files. Activities that you can audit include the following:

- User attempts to log on, both successful and unsuccessful.
- User attempts to access restricted accounts.
- User attempts to execute restricted commands.

| **Note** |
| --- |
| In IIS 6.0, successful account logons, including every impersonation of the anonymous account, are audited by default. If your server is hosting busy Web sites, this can cause the security event log to fill quickly with entries. To disable auditing of successful logons in the security event log, see "Auditing security events" in Help and Support Center for Windows Server 2003. |

**Related Information**

- For information about how to log user Web site and FTP site activity, see Logging Site Activity.
- For information about how to control access to your Web site and FTP site, see Access Control with IIS 6.0.

# This Topic Is No Longer Available

We're sorry—the topic that you requested is no longer available. Use the search box below to find related information.

# Performance Tuning

Updated: June 30, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Web server administrators often monitor their Web servers to create a performance baseline. A performance baseline is a collection of data that indicates how the servers are performing when everything is running smoothly. Before making changes to their servers in a production environment, administrators tune their servers in a test environment to achieve the established performance baseline. By tuning their servers, Web server administrators maximize the Web server's ability to handle HTTP requests and minimize Web application response times, which create a better experience for clients accessing their Web servers.

**In This Section**

This section includes the following information:

- Using HTTP Compression: Describes how to improve performance between client computers and servers by reducing file sizes before the files are sent to the client computer. Reducing the file size as described here is referred to as compression.
- Enabling HTTP Keep-Alives: Describes how to use HTTP Keep-Alive requests to maintain an open connection.
- Limiting Connections: Describes how to set limits on the number of connections to a Web server.
- Setting Connection Timeouts: Describes how to set connection time-out values on a Web server.
- Throttling Bandwidth: Describes how to change the bandwidth used by a Web server and individual Web sites.
- Setting the IIS Object Cache Time Period: Describes how to set the IIS object cache time period.
- Maximizing Data Throughput for Network Applications: Describes how to maximizing data throughput for network applications.
- Setting the SSL Cache Time-out Interval: Describes how to set the SSL cache time-out interval.
- Measuring Load Time for a Web Page: Describes how to measure load time for a Web page.
- Enabling CPU Monitoring in IIS 6.0: Describes how to monitor and stop problematic applications.
- Configuring Application Pool Queue-Length Limits: Describes how to limit the number of requests IIS stores in any application pool queue so the requests do not grow to such a large size as to exhaust server resources.
- Scalability: Provides information and links to detailed resources about scalability.
- Configuring Last-Modified Caching: Describes how to configure last-modified caching.
- Configuring ASP Change-Notification Caching: Describes how to configure ASP change-notification caching.

**Related Sections**

- For more information about each of the performance features described in this section and for a more detailed introduction into what it means to monitor and tune an IIS Web server, see Optimizing IIS 6.0 Performance.
- For more information about performance tuning, how to work with counters, how to create and configure counters, trace logs, or alerts, and a list of best practices; see "Monitoring Performance" in Help and Support Center for Windows Server 2003.

- For information about monitoring the performance of your ASP.NET applications, see Overview of ASP.NET Performance Monitoring.

- For more information about the Web Capacity Analysis Tool for performance testing and monitoring, download the IIS 6.0 Resource Kit Tools and their accompanying documentation.

# Viewing Counter Data in the Performance Console

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In its default view, System Monitor provides a list of active counters, including their object and instance, directly below the graphical display of counter data. Remember that many counters display the last observed value or a cumulative count, not an average value or rate.

If you are new to Windows Server 2003 performance monitoring, open the Performance console and take a look at the data provided by a few counters.

**To select performance objects and counters to view in System Monitor**

1. In Control Panel, double-click **Administrative Tools**, and then double-click **Performance**.
2. On the System Monitor toolbar, click the **Add** button (**+**) to open the **Add Counters** dialog box.
3. Click the **Performance object** list box to view a list of available performance objects, and then select an object to monitor.

   When you select an object, the list box on the left shows the available counters for that object and the list box on the right displays the instances for that object.
4. To select specific counters for the performance object, click **Select counters from list** (the default option), and then select the counters that you want to view.

   After you select a counter, you can click **Explain** to view a description.
5. If you want to monitor one or more specific instances of the counter (when available), select from the list of instances in the box on the right. Windows provides a default selection if you do not select a specific instance.
6. Click **Add** to begin monitoring.

# Using HTTP Compression

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To more efficiently use available bandwidth, enable IIS HTTP compression. HTTP compression provides faster transmission time between compression-enabled browsers and IIS, regardless of whether your content is served from local storage or a UNC resource. You can compress static files and application response files. Compressing application response files is usually called *dynamic compression*.

This section includes the following information:

- Enabling HTTP Compression: Describes how to enable compression.
- Customizing the File Types IIS Compresses: Describes how to add or remove file types from IIS compression.

**Related Information**

- For a detailed discussion of how HTTP compression works, how to test compression, and detailed compression configuration options, see Using HTTP Compression for Faster Downloads.
- For more information about improving the performance of IIS network resources using related features like bandwidth throttling and connection time-outs, see Administering Network Resources.

# Enabling HTTP Compression

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Introduction to and procedures for enabling HTTP compression using IIS Manager and command-line scripts.

You can enable HTTP compression server-wide or on a specific directory. HTTP compression improves bandwidth utilization and speeds up Web site performance. In addition, you can compress static or dynamic responses.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName***"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable global HTTP compression by using IIS Manager**

1. In IIS Manager, double-click the local computer, right-click the **Web Sites** folder, and then click **Properties**.
2. Click the**Service** tab, and in the **HTTP compression** section, select the **Compress application files** check box to enable compression for dynamic files.
3. Select the **Compress static files** check box to enable compression for static files.
4. In the **Temporary directory** box, type the path to a local directory or click **Browse** to locate a directory. Once a static file is compressed, it is cached in this temporary directory until it expires, or the content changes. The directory must be on the local drive of an NTFS–formatted partition. The directory cannot be compressed or shared, and the access control lists (ACLs) for the directory must include Full Control access to the identity of the application pool or to the IIS_WPG group.
5. Under **Maximum temporary directory size**, click a folder size option. If you specify a maximum size under **Limited to (in megabytes)** (the default setting is 95 MB), then when the limit is reached, IIS automatically cleans up the temporary directory by applying the "least recently used" rule.
6. Click **Apply**, and then click **OK**.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator***cmd** to open a command window with administrator rights and then type **cscript.exe***ScriptName* (include the script's full path and any parameters). |

**To enable global HTTP compression by using Adsutil.vbs**

1. Open a command prompt.
2. To enable dynamic compression, type the following at the command prompt and then press ENTER:

   **cscript adsutil.vbs set w3svc/filters/compression/parameters/HcDoDynamicCompression true**

3. To enable static compression, type the following at the command prompt and then press ENTER:

   **cscript adsutil.vbs set w3svc/filters/compression/parameters/HcDoStaticCompression true**

**To enable HTTP Compression for Individual Sites and Site Elements**

1. To enable static compression for only a single directory, first disable global static compression (if it is enabled) and then enable static compression at that directory. For example, to enable static compression for a directory at http://www.contoso.com/Home/StyleSheets, perform the following steps:

a. Disable global static compression by executing the following command at a command prompt:

**adsutil set w3svc/filters/compression/parameters/HcDoStaticCompression false**

a. Enable static compression at this directory by executing the following command at a command prompt:

**adsutil set w3svc/1/root/Home/StyleSheets/DoStaticCompression true**

2. To disable static compression for only a single directory, first enable global static compression (if it is disabled) and then disable static compression at that directory. For example, to enable static compression for a directory at http://www.contoso.com/Home/StyleSheets, perform the following steps:

   a. Disable global static compression by executing the following command at a command prompt:

   **adsutil set w3svc/filters/compression/parameters/HcDoStaticCompression true**

   a. Enable static compression at this directory by executing the following command at a command prompt:

   **adsutil set w3svc/1/root/Home/StyleSheets/DoStaticCompression false**

**Related Information**

- For a detailed discussion of how HTTP compression works, how to test compression, and detailed compression configuration options, see Using HTTP Compression for Faster Downloads.

- For more information about improving the performance of IIS network resources using related features like bandwidth throttling and connection timeouts, e Administering Network Resources.

- For information about configuring the types of files IIS compresses server-wide, see Customizing the File Types IIS Compresses.

# Customizing the File Types IIS Compresses

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can customize the file types IIS compresses. Enabling static file compression compresses .htm, .html, and .txt files. Enabling dynamic compression compresses .asp, .dll, and .exe files.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To add one or more file types to the server-wide static compression configuration**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. Type both of the following commands:
   - **cscript adsutil.vbs SET W3SVC/Filters/Compression/Deflate/HcFileExtensions "htm" "html" "txt" "newext"** where *newext* is a file type you want to compress (for example, Microsoft Word or Excel documents), then press **ENTER**. You can add multiple file types separated by spaces.
   - **cscript adsutil.vbs SET W3SVC/Filters/Compression/gzip/HcFileExtensions "htm" "html" "txt" "newext"** where *newext* is a file type you want to compress, then press **ENTER**. You can add multiple file types separated by spaces.

| **Note** |
| --- |
| To remove one or more file types from the server-wide static compression configuration, repeat the previous two commands, leaving out the file type you want to remove. |

**To add one or more file types to the server-wide dynamic compression configuration**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and click **OK**.
3. Type both of the following commands:
   - **cscript adsutil.vbs SET W3SVC/Filters/Compression/Deflate/HcScriptFileExtensions "asp" "dll" "exe" "newext"** where newext is a file type you want to compress (for example, aspx, a commonly used ASP extension), then press **ENTER**. You can add multiple file types separated by spaces.
   - **cscript adsutil.vbs SET W3SVC/Filters/Compression/gzip/HcScriptFileExtensions "asp" "dll" "exe" "newext"** where newext is a file type you want to compress, then press **ENTER**. You can add multiple file types separated by spaces.

| **Note** |
| --- |
| To remove one or more file types from the server-wide dynamic compression configuration, repeat the previous two commands, leaving out the file type you want to remove. |

**Related Information**

For a detailed discussion of how HTTP compression works, how to test compression, and detailed compression configuration options, see Using HTTP Compression for Faster Downloads.

For more information about improving the performance of IIS network resources using related features like bandwidth throttling and connection timeouts, see Administering Network Resources.

# Enabling HTTP Keep-Alives

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Most Web browsers request that the server keep the client connection open while the server sends multiple elements (.htm files and .gif or .jpeg files) to the client. Keeping the client connection open in this way is referred to as an HTTP Keep-Alive. Keep-Alive is an HTTP specification that improves server performance. HTTP Keep-Alives are enabled by default.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable HTTP Keep-Alives**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click the Web site, and click **Properties**.
2. On the **Web Site** tab, in the **Connections** section, click the **Enable HTTP Keep-Alives** check box.
3. Click **Apply**, and then click **OK**.

**Related Information**

- For information about HTTP Keep-Alives, including security requirements and how disabling keep-alives can impact the server, see Enabling HTTP Keep-Alives to Keep Connections Open.
- For more information about improving the performance of IIS network resources using related features like HTTP Compression, bandwidth throttling, and connection limits, see Administering Network Resources.

# Limiting Connections

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Connection limits restrict the number of simultaneous client connections to your Web sites and your Web server. Limiting connections conserves memory and protects against malicious attacks designed to overload your Web server with thousands of client requests.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set a global WWW or FTP service connection limit**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** or **FTP Sites** folder, and then click **Properties**.
2. Click the **FTP Site** tab to limit connections on the FTP service, or click the **Performance** tab to limit connections on the WWW service.
3. Click the **Connections limited to** option, and in the box next to it, type the maximum number of simultaneous connections you want to allow on your Web or FTP service.
4. Click **Apply**, and then click **OK**.

**To set a Web or FTP site connection limit**

1. In IIS Manager, expand the local computer, expand the **Web Sites** or **FTP Sites** folder, right-click the Web or FTP site on which you want to set the connections limit, and click **Properties**.
2. Click the **FTP Site** tab to limit connections on the FTP site, or click the **Performance** tab to limit connections on the Web site.
3. Click the **Connections limited to** option, and in the box next to it, type the maximum number of simultaneous connections you want to allow on your Web or FTP site.
4. Click **Apply**, and then click **OK**.

**Related Information**

- For more detailed information about IIS connection limits, including the security ramifications of selecting the **Unlimited** option, see Limiting Connections to Manage Resources.
- For more information about improving the performance of IIS network resources using related features like HTTP Compression and bandwidth throttling, see Administering Network Resources.

# Setting Connection Timeouts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Connection time-outs help reduce the loss of processing resources consumed by idle connections. When you enable connection time-outs, IIS enforces the time-outs at the connection level.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To set a global WWW or FTP service connection time-out value**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** or **FTP Sites** folder, and click Properties.
2. On the **Web Site** or **FTP Site** tab, in the **Connection timeout** box, type the maximum number of seconds that IIS should maintain an idle connection before resetting the connection.
3. For the WWW service, verify that the **Enable HTTP Keep-Alives** box is selected. For more information, see Enabling HTTP Keep-Alives.
4. Click **Apply**, and then click **OK**.

**To set a connection time-out value for a specific Web or FTP site**

1. In IIS Manager, expand the local computer, expand the **Web Sites** or **FTP Sites** folder, right-click a Web or FTP site, and click **Properties**.
2. On the **Web Site** or **FTP Site** tab, in the **Connection timeout** box, type the maximum number of seconds that IIS should maintain an idle connection before resetting the connection.
3. For the WWW service, verify that the **Enable HTTP Keep-Alives** box is selected. For more information, see Enabling HTTP Keep-Alives.
4. Click **Apply**, and then click **OK**.

**Related Information**

- For more information about connection timeouts, including which counters to use to monitor connection timeouts, see Setting Connection Timeouts to Save Resources.
- For more information about improving the performance of IIS network resources using related features like bandwidth throttling, HTTP Compression, and connection limits, see Administering Network Resources.
- For information about setting CGI application timeouts, see Configuring CGI Application Timeouts.

# Throttling Bandwidth

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If the network or Internet connection used by your Web server is also used by other services such as e-mail or news, you may want to limit the bandwidth used by your Web server so it is available for those other services. If your Web server hosts more than one Web site, you can individually throttle the bandwidth used by each site.

This section includes the following information:

- Enabling Packet Scheduler: Describes how to enable Packet Scheduler when configuring a site to use bandwidth throttling without using IIS Manager.
- Enabling Bandwidth Throttling: Describes how to enable bandwidth throttling globally for the WWW service and for individual Web sites.
- Disabling Bandwidth Throttling: Describes how to disable bandwidth throttling for individual Web sites.

**Related Information**

- For more information about when and where to use bandwidth throttling, see Throttling Bandwidth to Manage Service Availability.
- For more information about improving the performance of IIS network resources using related features like HTTP Compression and connection limits, see Administering Network Resources.

# Enabling Packet Scheduler

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Bandwidth throttling uses Packet Scheduler to manage when data packets are sent. When you configure a site to use bandwidth throttling by using IIS Manager, Packet Scheduler is automatically installed, and IIS automatically sets bandwidth throttling to a minimum of 1024 kilobytes per second (KBps). However, if you use another method, such as Active Directory® Service Interfaces (ADSI) or WMI, you must install Packet Scheduler in order for bandwidth throttling to work correctly.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To enable Packet Scheduler**

1. From the **Start** menu, click **Control Panel**, click **Network Connections**, right-click any of the listed connections, and then click **Properties**.
2. On the **General** tab under **This connection uses the following items**, click any of the listed items, and then click **Install**.
3. In the **Select Network Component Type** dialog box, click **Service**, and then click **Add**.
4. In the **Select Network Service** dialog box, click **QoS Packet Scheduler**, and then click **OK** to close all open dialog boxes.

When using ADSI or WMI to configure bandwidth throttling, you must set bandwidth throttling to 1024 KBps or more, because Packet Scheduler cannot enforce bandwidth throttling settings below that level. Also, you need to uninstall Packet Scheduler if it is not in use by individual sites or the WWW service as a whole.

**Related Information**

- For more information about administering IIS with ADSI or with WMI, see ADSI and WMI Providers in IIS in the IIS Platform SDK on MSDN.

# Enabling Bandwidth Throttling

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The following procedures tell how to globally enable bandwidth throttling for the WWW service or to enable bandwidth throttling for a specific Web site.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To globally throttle WWW service bandwidth**

1. In IIS Manager, expand the local computer, right-click the **Web Sites** folder, and click **Properties**.
2. Click the **Performance** tab, and in the **Bandwidth throttling** section, select the **Limit the total network bandwidth available for all Web sites on this server** check box.
3. In the **Maximum bandwidth (in kilobytes per second)** box, type the maximum number of kilobytes per second that you want each site contained in the directory to use. This value must be at least 1024 Kbps.
4. Click the **Connections limited to** option, and in the box next to it, type the maximum number of simultaneous connections that you want to allow on your Web or FTP service.
5. Click **Apply**, and then click **OK**.

**To throttle the bandwidth used by an individual Web site**

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click the Web site, and click **Properties**.
2. Click the **Performance** tab, and in the **Bandwidth throttling** section, select the **Limit the network bandwidth available to this Web site** check box.
3. In the **Maximum bandwidth (in kilobytes per second)** box, type the maximum number of kilobytes per second that you want this site to use. This value must be at least 1024 Kbps.
4. Click **Apply**, and then click **OK**.

# Disabling Bandwidth Throttling

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The following procedure tells how to disable bandwidth throttling for an individual site. The procedure is slightly different if you are overriding global bandwidth throttling. Either task can also be performed by using ADSI or WMI to update the **MaxBandwidth** metabase property for the site.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To disable bandwidth throttling on an individual site**

1. If global bandwidth throttling is not in effect, on the **Performance** tab, under **Bandwidth throttling**, clear the **Limit the network bandwidth available to this Web site** check box. Alternatively, you can set the **MaxBandwidth** metabase property to **-1** by using ADSI or WMI.
2. If global bandwidth throttling is in effect, but you want an individual site to have virtually unlimited bandwidth, on the Performance tab, under Bandwidth throttling, do not clear the **Limit the network bandwidth available to this Web site** check box (the global setting automatically applies to the individual site when the check box is cleared). Instead, enter a very high number in the **Maximum bandwidth (in kilobytes per second)** box. Alternatively, you can set the **MaxBandwidth** metabase property to **-2** by using ADSI or WMI.

| **Note** |
| --- |
| Bandwidth throttling is not supported for requests that come to a site by means of an IPv6 address. The **MaxBandwidth** metabase property does not affect IPv6 network traffic. |

**Related Information**

- For more information about the **MaxBandwidth** property, see MaxBandwidth Metabase Property.
- For more information about administering IIS with ADSI or with WMI, see ADSI and WMI Providers in IIS in the IIS Platform SDK on MSDN.

# Setting the IIS Object Cache Time Period

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The value of ObjectCacheTTL controls the static file cache by specifying the Time To Live (TTL), which sets the length of time that objects are held in cached memory. If an object in the memory cache has not been referenced for the defined period, that object is phased out of the cache.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see IIS 6.0 Registry Reference. |

Use the following procedure to reset the time that an unused object remains in the object cache by adding the ObjectCacheTLL entry to the registry.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to edit the registry. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to open a command window as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **regedit.exe** to open the registry editor. |

**To reset the period that unused objects remain in the cache**

1. From the **Start** menu, click **Run**, type **regedit.exe**, and then click **OK**.
2. In the registry editor, navigate to the following subkey:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters
3. Right-click the Parameters subkey, point to **New**, and then click **DWORD Value**.
4. In the **New Value** box, type **ObjectCacheTTL**.
5. Right-click **ObjectCacheTTL**, and then click **Modify**.
6. Under **Base**, click **Decimal**.
7. In the **Value Data** box, type the number of seconds that you want an unused object to remain in the cache, and then click **OK**.

   The default value is 30 (seconds). You can enter any value from zero, to disable caching, through 4,294,967,295 (unlimited), to disable the object-cache scavenger and allow cached objects to remain in the cache until the cached object changes.

# Maximizing Data Throughput for Network Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, servers running Windows Server 2003 are configured to give preference to the file system cache over the working sets of processes when allocating memory (through the server property **Maximize data throughput for file sharing**). Although IIS 6.0–based servers benefit from a large file system cache, giving the file system cache preference often causes the IIS 6.0 pageable code to be written to disk, which results in lengthy processing delays. To avoid these processing delays, set server properties to maximize data throughput for network applications.

**To maximize data throughput on the server for network applications**

1. In Control Panel, double-click **Network Connections**, right-click **Local Area Connection**, and then click **Properties**.
2. Select **File and Printer Sharing for Microsoft Networks**, and then click **Properties**.
3. Under **Optimization**, select **Maximize data throughput for network applications**.

# Setting the SSL Cache Time-out Interval

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you expect thousands of users to connect to your site by using SSL, estimate how long you expect SSL sessions to last, and then set the value of the ServerCacheTime entry to a number slightly higher than your estimate. Do not set the value much higher than your estimate, because the resulting time-out interval might cause your server to retain stale data in the cache.

| **Caution** |
| --- |
| The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see IIS 6.0 Registry Reference. |

Before changing the SSL cache time-out interval, make sure that HTTP Keep-Alives are enabled (HTTP Keep-Alives are enabled by default). SSL sessions do not expire when you use them with HTTP Keep-Alives except when the browser closes the connection.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to edit the registry. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to open a command window as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **regedit.exe** to open the registry editor. |

**To configure the ServerCacheTime registry entry**

1. In the registry editor, navigate to the following subkey:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders \SCHANNEL.

2. Right-click the SCHANNEL subkey, point to **New**, and then click **DWORD Value**.
3. In the **New Value** box, type the following: **ServerCacheTime**
4. Right-click the **ServerCacheTime** entry and then click **Modify**.
5. Under **Base**, click **Decimal**.
6. In the **Value Data** box, type the value (in milliseconds) that you want to assign for the cache time (1 minute = 60,000 milliseconds), and then click **OK**. See Table 6.28 for commonly used cache times converted to milliseconds.

**Table 6.28 Calculating ServerCacheTime Values for Secure Session Caching**

| **Desired Cache Time (1 minute = 60,000 milliseconds)** | **ServerCacheTime Value (in Milliseconds)** |
| --- | --- |
| No secure session caching | 0 (turns off session caching) |
| 2 minutes (default setting for the Microsoft® Windows NT® version 4.0 operating system) | 120000 |
| 5 minutes (default setting for the Microsoft® Windows® 2000 operating system) | 300000 |
| 10 hours (default setting for Windows Server 2003, Windows 2000 with Service Pack 2 [SP2] or later, and Windows XP) | 36000000 |

For more information about configuring the SSL cache time-out interval, see Configuring ServerCacheTime for SSL Sessions.

# Measuring Load Time for a Web Page

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Measure your first-access and second-access page-load times by using the procedures that follow.

**To measure a Web page's first-access load time**

1. Set up your WAN simulator to represent your 50th percentile domestic user. For example, your settings might be 60 milliseconds roundtrip delay, 50 kilobits per second, and no packet loss.
2. Set the site that you want to test as the default home page for your browser.
3. Browse to a site other than the one you are testing.

   If you are still on the test site when you clear your cache, the site refreshes after you clear the cache, reloading the files that you want to clear.
4. Clear the cache of all files by doing the following:
   a. On the **Tools** menu of Internet Explorer, click **Internet Options**.
   b. On the **General** tab, under **History**, click **Clear History**.
5. Close your browser.

   You must close TCP connections between tests, or your tests will be affected by TCPs slow-start algorithm, which retains the window size from the previous page load, producing inaccurate results.
6. Start your Network Monitor capture.
7. Open your browser to the test site (your new home page).
8. Stop the capture when the page is fully loaded and run Network Monitor. Determine how much total time elapsed while the page was loading.
9. Copy your data into a Microsoft® Excel worksheet for further analysis.

| **Important** |
| --- |
| Do not use CTRL+F5 (refresh the Web page — regardless of time stamp) as a substitute for this procedure. Although using CTRL+F5 is similar to this process in that it forces a full reload, it does not disconnect the TCP connections before the new page load. |

When you test a Web page for its second-access load time, you can use files that are saved in the client's cache.

**To measure a Web page's second-access load time**

1. Use the same WAN parameters and browser default settings that you used to determine first-access load time.
2. Close your browser without clearing the cache.
3. Restart your Network Monitor capture.
4. Open your browser to the test page.
5. Stop the capture when the page is fully loaded, and measure the time from the first TCP SYN from the client to the server, until the last byte is received.
6. Copy the results to a second Excel worksheet for further analysis.

| **Important** |
| --- |

Do not use F5 (refresh the Web page) as a substitute for this procedure because it does not provide the same results as a second page-load. Using F5 forces Internet Explorer to check the date on all cached files regardless of whether the dates will expire in the future.

After obtaining a page-load performance profile, analyze the data and remove obstacles to the pages efficient performance by using the information in Optimizing a Page-Load Performance Profile. For more information about first-access and second-access page-load times, see Obtaining a Page-Load Performance Profile.

# Enabling CPU Monitoring in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The CPU monitoring feature monitors and automatically shuts down worker processes that consume large amounts of CPU time. CPU monitoring is enabled for individual application pools.

Before enabling CPU monitoring, take note of the following:

- IIS must be operating in Worker Worker Process Isolation Mode.
- CPU monitoring is available only for application pools.
- CPU monitoring does not apply to CGI applications.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Enabling CPU Monitoring by Using IIS Manager**

When you enable CPU monitoring by using IIS Manager, IIS allows, by default, a maximum CPU use of 50 percent. To change the default value, follow this procedure.

**To enable CPU monitoring**

1. In IIS Manager, expand the local computer, expand the **Application Pools** folder, right-click the application pool for which you want to enable CPU accounting, and click **Properties**.
2. Click the **Performance** tab, and select the **Enable CPU monitoring** check box.
3. In the **Maximum CPU use** box, click the up and down arrows to set the maximum percentage of the CPU the application pool should use. If the application pool uses more than the designated maximum, IIS generates an error in the Windows Events log.
4. In the **Refresh CPU usage numbers (in minutes)** box, click the up and down arrows to set the refresh rate.
5. In the **Action performed when CPU usage exceeds maximum CPU use** list box, click the appropriate action for the designated application pool.
   - To have IIS write an event to the system log when the application pool reaches the maximum CPU usage, without shutting the application pool down, click **No Action**.
   - To shut down the application pool in addition to logging an event in the system log, click **Shutdown**.

     By default, **No Action** is selected.

     | **Important** |
     | --- |
     | Shutting down the application pool shuts down all worker processes that serve that application pool. |

6. Click **Apply**, and then click **OK**.

**Enabling CPU Monitoring from the Command Line**

Enabling CPU monitoring from the command line can take as many as three steps:

- Set the **CPULimit** metabase property, which limits the worker processes in a specified application pool to a percentage of CPU time.

- Set the **CPUResetInterval** metabase property, which specifies the time interval for CPU monitoring.

- Set the **CPUAction** metabase property, which specifies the type of action you want IIS to take, such as write to the event log or shut down worker processes that exceed the CPU limit.

#### Set the CPU Limit

The**CPULimit**property configures the maximum percentage of CPU time that the worker processes in an application pool are allowed to consume over the period of time set by the **CPUResetInterval** property. If the limit set by the **CPULimit** property is exceeded, an event is written to the event log and an optional set of events can be triggered as determined by the **CPUAction** property. Setting the value of the **CPULimit** property to zero disables CPU monitoring.

To start CPU monitoring for an application pool, set the **CPULimit** property to a value greater than zero. To calculate the value to set, determine the maximum percentage of CPU time that you want the worker processes in an application pool to use — for example, 50 percent (which is the default value). Next multiply that percentage by 1,000 (50  1,000 = 50,000). Hence, to limit CPU use to 50 percent, set a value of 50000 for the CPU limit.

**To set the CPU limit for CPU monitoring by using Adsutil.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.

2. At the command prompt, type:

   **cscript %Systemroot%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/CPULimit** *n*

   Replace *n* with the maximum percentage of CPU use to allow for the application pool in 1/1,000ths of a percentage.

#### Set the CPU Reset Interval

After turning on CPU monitoring for the application, use Adsutil.vbs to set a value for the **CPUResetInterval** property, which specifies the reset period (in minutes) for CPU monitoring and throttling limits on the application pool. When the number of minutes elapsed since the last CPU monitoring reset equals the number specified by this property, IIS resets the CPU timers for both the logging and limit intervals. The default value is five minutes. Setting the value of this property to zero disables CPU monitoring.

| **Important** |
| --- |
| The value of the **CpuResetInterval** property must be greater than the time between logging operations; otherwise, IIS resets counters before logging has occurred, and CPU monitoring does not occur. |

**To set the CPU monitoring reset interval by using Adsutil.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.

2. At the command prompt, type:

   **cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/CPUResetInterval** *n*

   Replace *n* with the number of minutes for the reset interval. To disable CPU monitoring, type **0**.

#### Set the CPU Action

Finally, set the **CPUAction** property, which configures the actions that IIS takes when Microsoft® Windows NT® job objects run. Only one Windows NT job object exists per application pool; therefore, you must configure the **CPUAction** property separately for each application pool. If you do not set the **CPUAction** metabase property, IIS assigns the default value, which is zero (event logging only).

**To set the CPUAction value by using Adsutil.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.

2. At a command prompt, type:

**cscript %SystemDrive%\Inetpub\AdminScripts\adsutil.vbs set W3SVC/AppPools/** *ApplicationPoolName* **/CPUAction** *ActionValue*

Replace *ActionValue* with the appropriate value: **0** logs an error without shutting down the process; **1** logs an error and shuts down the process.

For more information about setting the CPU action, see the CPULimit Metabase Property, the CPUResetInterval Metabase Property, and theCPUAction Metabase Property.

**Related Information**

- For more information about CPU Monitoring, including explanations of the corresponding metabase properties and options for advanced CPU monitoring with Windows System Resource Manager, see Enabling CPU Monitoring.

- For information about related features that help manage and conserve resources when IIS is running in Worker Process Isolation Mode, see Managing Resources in Worker Process Isolation Mode.

# Configuring Application Pool Queue-Length Limits

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

*Application pool queue-length limits* prevent large numbers of requests from queuing up and overloading your server. When application pool queue-length limits are enabled, IIS monitors the number of requests for a designated application pool queue before queuing a new request. If adding the new request to the queue exceeds the queue size, the server rejects the request and sends an uncustomizable 503-error response to the client.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To change an application pool queue-length limit using IIS Manager**

1. In IIS Manager, expand the local computer, expand the **Application Pools** folder, right-click the application, and then click **Properties**.
2. Click the **Performance** tab.
3. In the **Request queue limit** section, select the **Limit the kernel request queue to** check box, and click the up and down arrows in the **requests** box to set the maximum number of queued requests.

   **Note**

   If you clear the **Limit the kernel request queue to** check box, or if you set the **AppPoolQueueLength** metabase property to zero, IIS does not enforce a kernel request limit. Without a kernel request limit, IIS could queue an unlimited number of requests and your server could potentially run out of memory. For this reason, you should set kernel request-queue limits on all production servers.

4. Click **Apply**, and then click **OK**.

**Related Information**

- For more information about application pool queue-length limits, as well as other features and strategies for managing memory resources on your server, see Optimizing Memory Usage.
- For information about application pools, including guidelines for creating application pools, see Creating Application Pools.
- For information about features and strategies to help you monitor the health and performance of your application pools, see Ensuring Application Health in Worker Process Isolation Mode.

# Scalability

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

**Scalability** is the capability of a piece of hardware or software to easily expand to meet future computing needs. Each version of IIS is designed for better performance and scalability than the previous version. However, scalability can be impaired by hardware configurations and bottlenecks in poorly designed or poorly tested Web applications. Administrators can dramatically improve the scalability of their IIS server environment with careful planning and assessment of their infrastructure, which includes hardware, distributed applications, and custom components.

**Related Information**

- For information about scaling IIS, including information about adjusting IIS queue and cache settings and how to scale IIS with UNC centralized content storage, see Web Server Scalability.

- For information about monitoring IIS performance and scalability, including which counters to use, the tools available, and how to collect useful performance data, see Optimizing IIS 6.0 Performance.

# Configuring Last-Modified Caching

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To allow for greater scalability of UNC-based storage systems, IIS 6.0 implements a new last-modified cache algorithm for static and ASP files. The new caching algorithm does not request change notification for each directory structure; instead, it requests only the last-modified date of the cached file. If the file is new, the cache entry is updated with the new content and served. If the file has not changed, IIS serves the cached version of the file. By default, IIS checks the last-modified date for both ASP and static files if more than five seconds have elapsed since the file was last checked. Otherwise, IIS assumes the file has not changed, and it serves the existing content.

| **Note** |
| --- |
| The server side includes (SSI) file handler (Ssinc.dll) makes use of the static file cache, so this behavior occurs for .stm, .shtml, and any other files that are mapped to this DLL. |

IIS 6.0 uses the last-modified caching algorithm by default. To specify change-notification caching or to change the TTL value of the cached files, you must edit the registry.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the IIS 6.0 Registry Reference. |

**To set change-notification tracking or TTL for static files**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **Regedit.exe**, and then click **OK**.
3. Navigate to the following subkey: HKLM\System\CurrentControlSet\Services\Inetinfo\Parameters
4. Do one or both of the following:
   - To enable change-notification caching, modify the value of the registry entry **DoDirMonitoringForUnc** by setting it to **1**.
   - To change the TTL setting from the default of five seconds, change the value of the registry entry **FileAttributeCheckThreshold** to the number of seconds you want to use.
5. Close the registry editor.

When stale content is being served by the Web Server even though the content has been updated on the remote file server, you can disable the Web Server cache to attempt to isolate the problem.

**To disable the static file cache**

1. Navigate to the following subkey in the registry:

   HKLM\System\CurrentControlSet\Services\Inetinfo\Parameters
2. Add DisableMemoryCache, a DWORD value, and set it to **1**.

Then disable the ASP file cache to complete disabling the Web Server cache. For information about disabling ASP file caching, see Configuring ASP Template Caching.

# Configuring ASP Change-Notification Caching

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS 6.0 provides a scalable change-notification mechanism for ASP pages. Because change notification occurs for unique directory structures, instead of for every directory, the number of work items required for change notification is reduced.

To specify change-notification caching or to change the TTL value of the cached files, you must edit the registry.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the IIS 6.0 Registry Reference. |

**To set change-notification tracking or TTL for ASP files**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **Regedit.exe**, and then click **OK**.
3. Navigate to the following subkey: HKLM\System\CurrentControlSet\Services\ASP\Parameters
4. Do one or both of the following:
   - To enable change-notification caching, change the value of the registry entry **EnableChangeNotificationForUNC** by setting it to **1**.
   - To change the TTL setting from the default of five seconds, change the value of the registry entry **FileMonitoringTimeoutSeconds** to the number of seconds you want to use.
5. Close the registry editor.

# Logging Site Activity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can collect information about user activity by enabling logging for your Web sites. Logging information in Internet Information Services (IIS) goes beyond the scope of the event logging or performance monitoring features provided by Microsoft® Windows®. The logs can include information such as who has visited your site, what the visitor viewed, and when the information was last viewed. You can use the logs to assess content popularity or to identify information bottlenecks.

**In This Section**

This section includes the following information:

- Enabling Logging: Describes how to enable or disable logging for a Web or FTP site, or for directories.
- Enabling UTF-8 Format for Non-English Languages and Security: Describes how to log data in UTF-8 format.
- Customizing W3C Extended Logging: Describes how to change which fields are logged when using W3C Extended logging.
- Enabling W3C Centralized Logging: Describes how to save all log data for all sites on the server in a single log file.
- Saving Log Files: Describes how to specify the log-file directory and the options for creating new log files.
- Centralized Binary Logging in IIS 6.0: Describes how to save all log files for all Web sites on a specific IIS server in a centralized container, in binary format.
- Remote Logging: Describes how to store log data on a remote share using a UNC path.
- Substatus Error Codes in Log Files: Describes the IIS 6.0 implementation of substatus error codes, how to log these codes in IIS log files, and how to use these codes to debug IIS problems.

Related Information

- For more information about logging, including examples and descriptions of various IIS log file formats and information about HTTP.sys error logging, see Analyzing Log Files section.
- For information about auditing techniques to track the activities of users and to detect unauthorized attempts to access your NTFS file system directories and files, see Auditing in IIS 6.0.
- For an explanation of W3C Extended logging fields, including a detailed technical breakdown of the **time-taken** field, see W3C Extended Log File Format.
- For more information about HTTP status codes (100 - 500) see HTTP Status Codes in IIS 6.0.
- For information about client error codes (400 - 404) and how to troubleshoot common request problems, see HTTP 40x - - Client Error Codes.

# Enabling Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can choose a format and enable logging for individual Web sites and FTP sites. After you enable logging on a Web site or FTP site, all traffic to the site (including virtual directories) is written to the corresponding file for each site.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable logging on a Web or FTP site**

1. In IIS Manager, double-click the local computer, double-click the **Web Sites** or **FTP Sites** folder, right-click the Web site or FTP site for which you want to enable logging, and then click **Properties**.
2. On the Web Site or FTP Site tab, select the **Enable logging** check box.
3. In the **Active log format** box, click a format. By default, the format is **W3C Extended Log File Format**.

   | **Note** |
   | --- |
   | If you select ODBC logging, click **Properties** and type the ODBC Data Source Name (DSN) and the name of the table within the database in the text boxes. If a user name and password are required to access the database, type the necessary credentials, and then click **OK**. |

   | **Note** |
   | --- |
   | Logging to a UNC share is not supported by IIS FTP. You must configure the FTP log files location to a path on the local machine. |

4. Click **Apply**, and then click **OK**.

**Related Information**

- For more information about logging, including examples and descriptions of various IIS log file formats and information about HTTP.sys error logging, see the Analyzing Log Files section.
- For an explanation of W3C Extended logging fields, including a detailed technical breakdown of the **time-taken** field, see W3C Extended Log File Format.

# Enabling UTF-8 Format for Non-English Languages and Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Note** |
| --- |
| IIS does not support the UTF-8 format for FTP site log files. |

*UTF-8* is a method of character encoding that allows for both single and multibyte characters in one string. If you want to read text-based (W3C Extended, IIS, and NCSA Common) log files in a language other than English, you can enable the UTF-8 format. You also need the UTF-8 format if you serve URLs in a language other than the native server code page, for example, if you serve Chinese URLs on a Japanese system. In addition, because IIS 6.0 tries to serve URLs in code pages other than the server default code page by default, you might want to enable the UTF-8 format for security reasons. It is a good security practice to enable the UTF-8 format in case of an attack based on URLs in UTF-8 that might not translate correctly to the default code page.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

### To enable UTF-8 format

1. In IIS Manager, right-click the local computer, and then click **Properties**.
2. In **UTF-8 Logging**, select the **Encode Web logs in UTF-8** check box, and then click **OK**.

#### Related Information

- For more information about formats of log files in IIS, see Log File Formats in IIS.

# Customizing W3C Extended Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you selected the W3C Extended log file format for a Web or FTP site, you can customize the logged data. You can also limit unnecessary data and reduce the size of the log file by omitting unneeded fields.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To customize W3C Extended logging**

1. In IIS Manager, expand the local computer, expand the **Web** or **FTP Sites** folder, right-click the Web or FTP site, and click **Properties**.
2. On the **Web Site** or **FTP Site** tab, select the **Enable logging** check box (if it is not already selected).
3. In the **Active log format** list box, click **W3C Extended Log File Format**.
4. Click **Properties**.
5. On the **Advanced** tab, select the properties you want to log.
6. Click **OK**.

**Related Information**

- For an explanation of W3C Extended logging fields, including a detailed technical breakdown of the **time-taken** field, see W3C Extended Log File Format.
- For information about the **Protocol Substatus** field, see Substatus Error Codes in Log Files.

# Enabling W3C Centralized Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

W3C centralized logging is a global configuration on the server where all Web sites write data to a single log file. This single log file uses the W3C Extended log file format and can be viewed in a text editor. W3C centralized logging is available in Windows Server 2003 Service Pack 1 or later. For more information about this logging option, see W3C Centralized Logging.

By default, log files are stored in the Windows/system32/logfiles/W3SVC directory. After you enable, W3C centralized logging, IIS creates a new log file and all Web sites begin writing data to that single source.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To enable W3C Centralized Logging**

1. From a command prompt, type
   **Adsutil.vbs set w3svc/CentralW3CLoggingEnabled true**
2. The World Wide Web Publishing Service (WWW service) must be stopped and restarted for changes to take effect. From the command prompt, type **net stop W3SVC**, and press **ENTER**. The WWW service stops.
3. Type **Net start w3svc**

**See Also**
**Concepts**
CentralW3CLoggingEnabled Metabase Property

# Saving Log Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, IIS creates a new log file for each Web site in the *systemroot*\System32\LogFiles directory. However, you can specify the directory into which log files are saved, and you can determine when new log files are started. To protect logged data, set appropriate Access Control with IIS 6.0 on the log file directory. Consider allowing only the administrator and the IIS_WPG group to have access to the log file directory.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**Procedures**

**To set options for saving log files**

1. In IIS Manager, expand the local computer, expand the Web or FTP Sites directory, right-click the Web or FTP site, and click **Properties**.
2. On the **Web Site** or **FTP Site** tab, click **Properties** next to the **Active log format** list box.
3. Select the log schedule to use when starting a new log file.

   **Note**

   "Midnight" is midnight local time for all log file formats except the W3C Extended format. For W3C Extended log file format, "midnight" is midnight Greenwich Mean Time (GMT), by default, but can be changed to midnight local time. To open new W3C Extended logs using local time, select the **Use local time for file naming and rollover** check box. The new log starts at midnight local time, but the time recorded in the log files is still GMT.

4. Under **Log file directory**, type the directory where log files should be saved. For information about saving log files on a remote share, see Remote Logging.
5. Click **Apply**, and then click **OK** twice.

**Related Information**

- For information about saving log data in raw, binary format to conserve hard-disk space and to reduce memory consumption, see Centralized Binary Logging.
- For more information about IIS and centralized-content storage, see Improving Scalability Through UNC-Based Centralized Content Storage.

# Centralized Binary Logging in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Centralized binary logging is the process by which multiple Web sites write binary, unformatted log data to a single log file. When enabled, all Web sites on your server write log data to a single log file.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Procedures**

**To enable centralized binary logging for all Web sites on your server running IIS**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cscript.exe adsutil.vbs SET W3SVC/CentralBinaryLoggingEnabled true**, and click **OK**.
3. The World Wide Web Publishing Service (WWW service) must be stopped and restarted for changes to take effect. From the command prompt, type **net stop W3SVC**, and press ENTER. The WWW service stops.

From the command prompt, type **net start W3SVC**, and press ENTER. The WWW service starts and centralized binary logging is enabled.

The binary log file is created in the W3SVC folder located at *systemroot*\System32\LogFiles\.

**Related Information**

- For more detailed information about centralized binary logging, including information about how to extract data from the raw log file, see Centralized Binary Logging.
- For more information about W3C centralized logging, see W3C Centralized Logging.
- For more information about logging, including examples and descriptions of various IIS log file formats and information about HTTP.sys error logging, see Analyzing Log Files.

# Remote Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can write log data to a remote share over a network using a full Universal Naming Convention (UNC) path for centralized log file storage and backup. Be aware that remote logging can negatively affect performance because IIS writes the log file and log file data over the network.

| **Note** |
| --- |
| Logging to a UNC share is not supported by IIS FTP. You must configure FTP log files location to a path on the local machine. |

| **Important** |
| --- |
| It is highly recommended that you enable Internet Protocol security (IPSec) between your Web server running IIS and the remote server before configuring remote logging. If IPSec is not enabled between the Web server running IIS and the remote server, data packets containing log data are potentially at risk of being intercepted by malicious individuals and wire-sniffing applications while the data packet travels through the network. |

**Procedures**

**Setting up remote logging requires the following actions:**

1. Configuring IIS to Log Data on a Remote Share
2. Configuring Permissions for Remote Logging
3. Setting Up a Null Session for Cross-Domain Logging (if applicable)

**Related Information**

- For more information about IIS and centralized-content storage, see Improving Scalability Through UNC-Based Centralized Content Storage.

# Configuring IIS to Log Data on a Remote Share

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure IIS to write log file data to a remote share. In the remote share, IIS creates a unique directory for each Web site, for example **W3SVC***X*, where *X* is a random number generated by IIS to represent the specific Web site. IIS also creates the log file with exclusive write access, so that multiple machines cannot write to the same log file.

**Important**

You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc systemroot\system32\inetsrv\iis.msc"**.

**Procedures**

**To log Web site data on a remote share**

1. Create a log file directory on a remote server in the same domain as your Web server running IIS.
2. Change the directory properties so the directory is a share and assign the **Everyone** group **Full Control** permissions. If you want to write log files to a remote server in a different domain, see Setting Up a Null Session for Cross-Domain Logging.
3. Ensure that your server running IIS has **Full Control** access permission on the remote share and read and write permissions on the remote log file directory. For more information, see Configuring Permissions for Remote Logging.
4. In IIS Manager, expand the local computer, right-click the **Web Sites** folder, and click **Properties**.
5. On the **Web Site** tab, ensure that the **Enable logging** check box is selected.
6. In the **Active log format** list box, click a log file format.
7. Click **Properties**.
8. Click the **General** tab, and in the **Log file directory** box, type the full UNC path. For example, type *servername**LogFiles*, where *servername* represents the name of the remote server, and *LogFiles* represents the name of the share where the log files are stored.
9. Click **Apply**, and then click **OK**. All Web sites within the directory begin logging data to the remote share.

**Note**

Logging to a UNC share is not supported by IIS FTP. You must configure the FTP log files location to a path on the local machine.

**Related Information**

- For information about saving log data in raw, binary format to conserve hard-disk space and to reduce memory consumption, see Centralized Binary Logging in IIS 6.0.
- For more information about IIS and centralized-content storage, see Improving Scalability Through UNC-Based Centralized Content Storage.

# Configuring Permissions for Remote Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS can store log files on a remote share as long as the remote computer allows IIS to create log files and write data to the remote share. Use the following procedure to configure permissions on the remote machine so IIS can create log files and write data to the remote share.

**Procedures**

**To configure permissions for remote logging**

1. On the remote computer, navigate to *systemroot**System32*, right-click the **LogFiles** folder, and then click **Sharing and Security**.
2. On the **Sharing** tab, click **Share this folder**, and then click **Permissions**.
3. Click **Add**.
4. Click **Object Types**.
5. Select the **Computers** check box, and click **OK**. If you like, you can deselect all other options.
6. In the **Enter the object name to select** box, type the object name in the form *Domain**WebServer*, and then click **OK**.
7. In the **Group or user names** list, select the *Domain**WebServer* object, and in the **Permissions** section, select the **Allow** check box next to **Full Control**.
8. In the **Group or user names** list, select **Everyone**.
9. In the **Permissions** section, clear all permissions and click **OK**. The remote computer now has the appropriate access permissions.
10. To set the appropriate file permissions, click the **Security** tab.
11. Select the *Domain**WebServer* object and in the **Permissions** section, select the **Allow** check box next to **Full Control**.
12. Click **Apply**, and then click **OK**.

**Related Information**

- For more information about permissions in IIS see, Setting NTFS Permissions for Directories or Files.
- For more information about IIS and centralized-content storage, see Improving Scalability Through UNC-Based Centralized Content Storage.

# Setting Up a Null Session for Cross-Domain Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Remotely logging files to a UNC share in a different domain requires configuring the remote share as a null session share. A null session connection is an unauthenticated connection. When IIS attempts to access a remote Microsoft Windows server resource, such as a file share, using a null session, the operation might fail if the file share is not configured as a null session share, or if there are any registry, group, or policy restrictions set on the server hosting the file share.

| **Important** |
| --- |
| Using Registry Editor incorrectly can cause serious problems that require reinstalling the operating system. Because Registry Editor bypasses the standard safeguards that prevent you from entering settings that are conflicting or likely to degrade performance or damage your system, exercise caution when making changes to the registry. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. For information about how to edit the registry, see "Changing Keys and Values" in Registry Editor Help. |

**Procedures**

**To enable null session access**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **Regedit.exe**, and click **OK**.
3. Navigate to and double-click the following key in the registry:

   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\LanmanServer\Parameters\NullSessionShares.

   | **Note** |
   | --- |
   | **NullSessionShares** is a REG_MULTI_SZ value. |

4. On a new line in the **NullSessionShares** key, type the name of the share that you want to access with a null session, for example, **public**.
5. Navigate to and click the following key in the registry:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA.

6. From the **Edit** menu, point to **Add**, click **DWORD Value** and then add the following registry value:

   Name: **RestrictAnonymous**

   Type: **REG_DWORD**

   Data: **0**

7. Quit Registry Editor.
8. Restart the server.

If the remote share does not allow your IIS server to write the log file to the designated share after restarting the server, you might need to adjust the Windows security groups and security policies on the remote machine to enable anonymous access.

**Related Information**

- For more information about security groups and security policies, see "Security" in Help and Support Center for Windows Server 2003.
- For more information about IIS and centralized-content storage, see Improving Scalability Through UNC-Based Centralized Content Storage.

# Substatus Error Codes in Log Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When enabled, substatus error codes are written to IIS log files whenever a client request fails.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

### To log substatus error codes

1. In IIS Manager, expand the local computer, expand the **Web Sites** folder, right-click the Web site, and click **Properties**.
2. On the **Web Site** tab, ensure that the **Enable logging** check box is selected.
3. In the **Active log format** list box, click **W3C Extended Log File Format**.
4. Click **Properties**.
5. Click the **Advanced** tab, and under **Extended properties**, select the **Protocol Substatus** check box.
6. Click **OK** twice.

**Related Information**

- For an explanation of W3C Extended logging fields, including a detailed technical breakdown of the **time-taken** field, see W3C Extended Log File Format.
- For more information about HTTP status codes (100 - 500) see HTTP Status Codes in IIS 6.0.
- For information about client error codes (400 - 404) and how to troubleshoot common request problems, see HTTP 40x - - Client Error Codes.

# Error Messages in IIS 6.0

Updated: June 30, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

You can configure Internet Information Services (IIS) to send default HTTP 1.1 error messages or custom error messages. Custom error messages can be mapped to a file name or to a URL.

This section includes the following information:

- Enabling Standard HTTP 1.1 Error Messages: Describes how to use the detailed (also known as friendly) HTTP 1.1 errors that are included with IIS.

- Configuring Custom Error Messages: Provides information about customizing Microsoft® Active Server Pages (ASP) error files.

- Enabling ASP Error Processing: Describes how to use the 500-100.asp custom error file to process ASP errors.

- Disabling HTTP Friendly Error Messages in Internet Explorer: Describes how to show literal HTTP error codes used with to troubleshooting.

  | **Note** |
  | --- |
  | HTTP.sys, the kernel-mode driver, does not serve custom errors. If a request does not reach user mode, HTTP.sys returns standard HTTP 1.1 error messages to the browser. HTTP.sys might return standard error messages for the following errors, even if you have specified a custom error message: |

  - 400: Cannot resolve the request.

  - 414: Request URL is too large and therefore unacceptable on the Web server.

  - 501: Header values specify a configuration that is not implemented.

  For details about why HTTP.sys returned one of these standard error messages, view the HTTP error logs. By default, HTTP error logs are created in the following location: *systemroot*\System32\LogFiles\HTTPERR.

- Disabling HTTP Friendly Error Messages in Internet Explorer: Describes how to disable HTTP friendly error messages in Internet Explorer for troubleshooting purposes.

**Related Sections**

- For more information about HTTP status codes (100–500) see HTTP Status Codes in IIS 6.0.

- For information about client error codes (400–404) and how to troubleshoot common request problems, see HTTP 40x -- Client Error Codes.

- For information about how to log substatus error codes and what substatus error codes mean, see Substatus Error Codes in Log Files.

# Enabling Standard HTTP 1.1 Error Messages

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, IIS sends custom error messages because these messages contain more information to help users understand why they have received an error. You can, however, configure IIS to send standard, or default, HTTP 1.1 error messages.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable standard HTTP 1.1 error messages**

1. In IIS Manager, double-click the local computer; right-click the Web site, virtual directory, directory, or file for which you want to enable standard HTTP 1.1 error messages; and then click **Properties**.
2. Click the **Custom Errors** tab.
3. In the **Error messages for HTTP errors** list, click the HTTP error(s) that you want to change to standard HTTP 1.1 error(s), click **Set to Default**, and then click **OK**.

**Related Information**

- For information about how to configure custom error messages, see Configuring Custom Error Messages.
- For more information about HTTP status codes (100 - 500) see HTTP Status Codes in IIS 6.0.
- For information about client error codes (400 - 404) and how to troubleshoot common request problems, see HTTP 40x - - Client Error Codes.
- For information about logging substatus error codes and what substatus error codes mean, see Substatus Error Codes in Log Files.

# Configuring Custom Error Messages

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS can send standard HTTP 1.1 error messages or custom error messages. Custom error messages provide more informative, or friendly, feedback than standard error messages. By default, IIS is configured to send custom error messages from a file stored in the *systemroot*\Help\IisHelp\Common folder. You can also configure IIS to return a URL to a custom error message.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure custom error messages**

1. In IIS Manager, double-click the local computer; right-click the **Web Sites** folder, an individual Web site folder, a virtual directory, or a file; and then click **Properties**.

   | **Note** |
   | --- |
   | Configuration settings made at the Web Sites level are inherited by all of the Web sites on the server. You can override inheritance by configuring the individual site or site element. |

2. Click the **Custom Errors** tab.

3. In the **Error messages for HTTP errors** list, click the HTTP error that you want to change, and then click **Edit**.

   | **Note** |
   | --- |
   | The following errors are not customizable: 400, 403.9, 411, 414, 500, 500.11, 500.14, 500.15, 501, 503, and 505. |

4. In the **Message Type** list box, click either **File** to return a custom error file or **URL** to direct the request to a custom error URL on the local machine.

   | **Important** |
   | --- |
   | If your custom error is an .asp page, you must select **URL**. If you do not select **URL**, you risk returning .asp source code to the client. |

5. If you selected **File**, type the path to the file or click **Browse** to navigate to the file. Custom error messages are installed by default to the *systemroot*\Help\IisHelp\Common folder. The file names are numbers that correspond to the specific HTTP errors; for example, 400.htm, 401-1.htm, and so on.

   or

   If you selected **URL**, type the path to the Web site or virtual directory. The URL must be a Web site or a virtual directory on the local machine. In addition, the custom error URL must exist in the application pool that directs the request to the custom error URL. If you store custom error pages in a virtual directory, that virtual directory must run in the same application pool as the rest of your Web site. Otherwise, the worker process cannot serve the custom error page when it is requested.

6. Click **OK**, and then click **OK** again.

If you upgraded from IIS 4.0, and you previously created your own custom error files or modified the custom error files that IIS 4.0 installed in the *systemroot*\Help\Common folder, your custom error files have been moved to the *systemroot*\Help\IisHelp\Common folder and given a .bak file name extension. You can retrieve your custom errors and reuse

them. For example, the custom error file 400.htm that IIS 4.0 installed is now named 400.bak and is located in the *systemroot*\Help\IisHelp\Common folder (in addition to the new IIS custom error files).

**Related Information**

- For more information about HTTP status codes (100 - 500) see HTTP Status Codes in IIS 6.0.
- For information about client error codes (400 - 404) and how to troubleshoot common request problems, see HTTP 40x -- Client Error Codes.
- For information about logging substatus error codes and what substatus error codes mean, see Substatus Error Codes in Log Files.

# Enabling ASP Error Processing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The default Web site transfers the processing of Active Server Pages (ASP) errors to the 500-100.asp file (by default). If you create a new Web site, ASP errors return standard HTTP Status Codes (known as **Default** in IIS Manager. If you want to customize ASP errors, you can either map the 500-100 error to the 500-100.asp file or create an error-processing ASP file of your own. If you create an ASP error-processing file of your own, you must enable the custom error file through IIS Manager.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To enable ASP error processing by mapping to a URL**

1. Create a file that contains your custom error message, and save it to a virtual directory.
2. In IIS Manager, double-click the local computer; right-click the Web site, virtual directory, directory, or file for which you want to enable ASP error processing; and then click **Properties**.
3. Click the **Custom Errors** tab.
4. In the **Error messages for HTTP errors** list, click the HTTP error that you want to change, and then click **Edit**.
5. In the **Message Type** list box, click **URL**.
6. In the **URL** box, type the URL that points to your customized error message by typing the path to the URL, beginning with the virtual directory name.
7. Click **OK**, and then click **OK** again.

**Related Information**

- For more information about HTTP status codes (100 - 500) see HTTP Status Codes in IIS 6.0.
- For information about client error codes (400 - 404) and how to troubleshoot common request problems, see HTTP 40x - - Client Error Codes.

# Disabling HTTP Friendly Error Messages in Internet Explorer

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Internet Explorer displays friendly error messages when it encounters a problem during communication with an HTTP server. Friendly error messages are appropriate for the end user. However, they can prevent administrators from seeing the literal HTTP error codes and explanations, which makes troubleshooting difficult, if not impossible. During the discovery phase of troubleshooting, you can gather and record the HTTP errors that have occurred if you have disabled friendly messages.

**To disable friendly HTTP error messages in Internet Explorer**

1. On the **Tools** menu in Internet Explorer, click **Internet Options**.
2. Click the **Advanced** tab.
3. In the list of options under **Browsing**, clear the **Show friendly HTTP error messages** check box.
4. Click **Apply**, and then click **OK**.

# Web Application Guide

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section includes information about configuring Web applications using Microsoft® ASP.NET. In this section you will find information about configuration, security, and deployment.

**In This Section**

This section includes the following information:

- About the .NET Framework: Provides an introduction to the Microsoft .NET Framework.
- About ASP.NET: Provides information about ASP.NET configuration, deployment, and security.
- About ASP: Provides an overview of what is new in Microsoft Active Server Pages (ASP).

Related Information

- For information about writing server-side scripts with ASP and creating ASP applications, see the Active Server Pages Guide on MSDN® Online or in the Internet Information Services (IIS) Software Development Kit (SDK).
- Redirect Reference: This section explains how to format a server redirection correctly. For reference topics related to ASP, see the IIS Web Application Technology Reference section of the IIS Software Developer Kit (SDK).

# About the .NET Framework

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server™ 2003 family. For more information, see "Features unavailable on 64-bit versions of the Windows Server family" in Help and Support Center for Windows 2003. |

The Microsoft .NET Framework is a new computing platform that simplifies application development in the distributed environment of the Internet. The .NET Framework is designed to provide an object-oriented programming environment that guarantees the safe execution of code, and to eliminate the performance problems of scripted environments.

The .NET Framework has two main components: the common language runtime and the .NET Framework class library. Microsoft ASP.NET hosts the runtime to provide core services such as memory management, thread management, and remoting while enforcing strict type safety and other forms of code accuracy that ensure security and robustness. Code management is a fundamental principle of the runtime; code that targets the runtime is called *managed code*, while code that does not target the runtime is called *unmanaged code*.

The .NET Framework class library is a comprehensive, object-oriented collection of reusable types that developers can use to create ASP.NET applications.

**Common Language Runtime Features**

The common language runtime (CLR) manages memory, thread execution, code execution, code safety verification, compilation, and other system services. These features are intrinsic to the managed code that runs on the common language runtime hosted by ASP.NET. Managed components are granted varying degrees of trust, depending on a number of factors including their origin (such as the Internet, enterprise network, or local computer). This means that a managed component might or might not be able to perform sensitive functions, even if it is being used in the same active application.

The CLR enforces code access security. Applications can be granted access to specific resources but cannot access any resources outside of those limits.

The CLR also enforces code robustness by implementing a strict type- and code-verification structure called the common type system (CTS). The CTS ensures that all managed code is self-describing, which means that managed code can consume other types and instances regardless of which language compiler was used to develop the software.

In addition, the managed environment of the CLR eliminates many common software issues. For example, the CLR automatically handles object layout and manages references to objects, releasing them when they are no longer in use. This automatic memory management resolves the two most common application errors: memory leaks and invalid memory references. The CLR is designed to enhance performance. Although the CLR provides many standard runtime services, managed code is never interpreted. Just-in-time (JIT) compiling enables all managed code to run in the native machine language of the system on which it is executing. Meanwhile, the memory manager further improves performance.

**.NET Framework Class Library Features**

The .NET Framework class library is a collection of reusable types that tightly integrate with the CLR. The class library is object-oriented, providing types from which applications can derive additional functionality. This object orientation makes the .NET Framework types easy to use and it reduces the time associated with learning new features of the .NET Framework. In addition, third-party components can integrate seamlessly with classes in the .NET Framework.

The .NET Framework types enable you to accomplish a range of common programming tasks such as string management, data collection, database connectivity, and file access. In addition, the class library includes types that support a variety of specialized development scenarios.

**Related Information**

- For information about writing, building, testing, and deploying .NET Framework applications, see the Microsoft .NET Framework Software Development Kit.

# About ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server™ 2003 family. For more information, see "Features unavailable on 64-bit versions of the Windows Server family" in Help and Support Center for Windows 2003. |

Microsoft ASP.NET is more than just the next generation of Active Server Pages (ASP). It provides an entirely new programming model for creating network applications that take advantage of the Internet. This section includes information about configuration, security, and deployment of ASP.NET applications.

FrontPage Server Extensions does not support unicode content directory names or Web Distributed Authoring and Versioning (WebDAV). New FrontPage Server Extensions virtual servers should be assigned to the MSSharePointPool application pool. Also, do not make changes to the MSSharePointPool application pool configuration.

| **Note** |
| --- |
| If you have a Web site that uses Sharepoint™ Team Services (STS) from Microsoft, the following types of pages do not work without excluding the virtual directory containing the pages from the STS ISAPI filter: |

- ASP pages
- ASP.NET pages
- PHP pages
- PerlScript pages

**Procedures**

**To enable these types of pages to work on your STS Web site**

1. In the **Central Administration** Web site for STS, choose to **Manage** your virtual directory or Web site.
2. Under **Defined Managed Paths**, type the path to the virtual directory that contains your excluded pages (such as \MyASPPages) and select **Exclude**.

This section includes the following information:

- Overview of ASP.NET: Describes the ASP.NET programming model and infrastructure.
- Why Use ASP.NET?: Describes the differences between ASP.NET and ASP.
- Introduction to XML Web Services Created with ASP.NET: Defines XML Web services and describes how to use them.
- ASP.NET Administration: Describes installing, configuring, and deploying ASP.NET applications on a computer running Microsoft Windows Server™ 2003.

**Related Information**

- For more information about creating ASP.NET Web applications, see Creating ASP.NET Web Applications
- For more information about security considerations for ASP.NET Web applications, seeASP.NET Web Application Security

# Overview of ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP.NET is a unified Web application platform that provides the services necessary to build and deploy enterprise-class Web applications. ASP.NET offers a new programming model and infrastructure for more secure, scalable, and stable applications that can target any browser or device.

ASP.NET is part of the Microsoft .NET Framework, a computing environment that simplifies application development in the highly distributed environment of the Internet. The .NET Framework includes the common language runtime (CLR), which provides core services such as memory management, thread management, and code security. It also includes the .NET Framework class library, which is a comprehensive, object-oriented collection of types that developers can use to create applications.

ASP.NET offers the following benefits:

- Manageability: ASP.NET uses a text-based, hierarchical configuration system that simplifies applying settings to your server environment and Web applications. Because configuration information is stored as plain text, new settings can be applied without the aid of local administration tools. Any changes to the configuration files are automatically detected and applied to the application. For more information, see Overview of ASP.NET Configuration.

- Security: ASP.NET provides default authorization and authentication schemes for Web applications. Developers can easily add to, remove, or replace these schemes depending on the needs of the application. For more information, see Overview of ASP.NET Security.

- Ease of Deployment: An ASP.NET application is deployed to a server by simply copying the necessary files to the server. No server restart is required — not even to deploy or replace running compiled code. For more information, see Overview of ASP.NET Deployment.

- Enhanced Performance: ASP.NET is compiled code running on the server. Unlike traditional Active Server Pages (ASP), ASP.NET can take advantage of early binding, just-in-time (JIT) compilation, native optimization, and caching services right out of the box for improved performance. For more information, see Overview of ASP.NET Performance Monitoring.

- Flexible Output Caching: ASP.NET can cache page data, portions of a page, or whole pages, depending on the needs of the application. Cached items can be dependent on files or other items in the cache, or they can be refreshed based on an expiration policy.

- Internationalization: ASP.NET uses Unicode internally to represent request and response data. Internationalization settings can be configured for each computer, each directory, and each page.

- Mobile Device Support: ASP.NET supports any browser on any device. Developers use the same programming techniques to target new mobile devices that they use for traditional desktop browsers.

- Scalability and Availability: ASP.NET was designed to be scalable, with features specifically tailored to improve performance in clustered and multiprocessor environments. Moreover, processes are closely monitored and managed by Internet Information Services (IIS) and the ASP.NET CLR so that if one misbehaves, a new process can be created in its place, which helps keep your application available to handle requests. For more information, see Overview of ASP.NET Process Isolation.

- Tracing and Debugging: ASP.NET provides tracing services that can be enabled during debugging at both the application level and the page level. You can choose to view the information either in a page or using the application-level trace viewing tool. ASP.NET supports local and remote debugging with .NET Framework debugging tools, both during development and when the application is in production. When an application is put into production, tracing statements can be left in the production code with no impact on performance.

- Integration with the .NET Framework: Because ASP.NET is part of the .NET Framework, the power and flexibility of the

entire platform is available to Web applications. The .NET class library, as well as solutions for messaging and data access, are all seamlessly accessible from the Web. ASP.NET is language-independent, so developers can choose the language that best applies to the application. In addition, CLR interoperability preserves existing investments in COM-based development.

- Compatibility with Existing ASP Applications: ASP and ASP.NET can run side by side on an IIS Web server without interference; there is no chance of corrupting an existing ASP application by installing ASP.NET. Only files with a .aspx file name extension are processed by ASP.NET. Files with an .asp file name extension will continue to be processed by the ASP engine. You should note, however, that session state and application state are not shared between ASP and ASP.NET pages.

ASP.NET enables two features for distributed applications: Web Forms and XML Web services. These two features are supported by the same configuration and debugging infrastructure.

- Web Forms technology enables you to build powerful form-based Web pages. Web Forms pages use reusable built-in or custom components to simplify the code of a page.

- XML Web services that are created using ASP.NET enable you to access servers remotely. Using XML Web services, businesses can provide programmatic interfaces to their data or business logic, which in turn can be obtained and manipulated by client and server applications. XML Web services enable the exchange of data using standards such as XML messaging and HTTP across firewalls in client/server and server/server scenarios. Programs written in any language and running on any operating system can call XML Web services.

**Related Information**

- For more information about how ASP.NET works with IIS 6.0, see How ASP.NET Works with IIS 6.0.

- For more information about deploying ASP.NET applications on IIS 6.0, see Deploying ASP.NET Applications in IIS 6.0.

# Why Use ASP.NET?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Microsoft ASP.NET is more than just the next generation of Active Server Pages (ASP). It provides an entirely new programming model for creating network applications that take advantage of the Internet.

**Improved Performance and Scalability**

- Compiled Execution: ASP.NET is much faster than classic ASP, while preserving the "just hit save" update model of ASP. No explicit compile step is required. ASP.NET automatically detects any change, dynamically compiles files if needed, and stores the compiled results to reuse for subsequent requests. Dynamic compilation ensures that your application is always up to date, and compiled execution makes it fast. Most applications migrated from classic ASP to ASP.NET see a 3x to 5x increase in pages served.

- Rich Output Caching: ASP.NET output caching can dramatically improve the performance and scalability of your application. When output caching is enabled on a page, ASP.NET executes the page once and saves the result in memory before sending it to the user. When another user requests the same page, ASP.NET serves the cached result from memory without re-executing the page. Output caching is configurable, and it can be used to cache individual regions or an entire page.

- Web Farm Session State: ASP.NET session state lets you share session data across all machines in a Web farm. Now a user can hit different servers in the Web farm over multiple requests and still have full access to session data.

**Enhanced Reliability**

- Memory Leak, Dead Lock, and Crash Protection: ASP.NET automatically detects and recovers from errors such as dead locks and memory leaks to ensure that your application is always available. For example, when a memory leak is detected, ASP.NET automatically starts up a new copy of the ASP.NET worker process and directs all new requests to the new process. After the old process has finished processing pending requests, it is gracefully disposed of and the leaked memory is released.

**Easy Deployment**

- "No Touch" Application Deployment: With ASP.NET you can deploy an entire application by copying it to the server. Configuration settings are stored in an XML file within the application.

- Dynamic Update of Running Application: ASP.NET lets you update compiled components without restarting the Web server. Unlike classic COM components that required the Web server to be manually restarted when an update was deployed, ASP.NET automatically detects the change and starts using the new code.

- Easy Migration Path: ASP.NET runs side by side on IIS with classic ASP applications on Microsoft Windows 2000 and Windows XP, and on members of the Windows Server 2003 family. You can migrate one application at a time, or even single pages. ASP.NET even lets you continue to use your existing classic COM business components.

**New Application Models**

- XML Web Services: XML Web services allow applications to communicate and share data over the Internet, regardless of operating system or programming language. ASP.NET makes exposing and calling XML Web services simple.

- Mobile Web Device Support: ASP.NET mobile controls let you target over 80 mobile Web devices using ASP.NET. You write the application once, and the mobile controls automatically generate pages for the requesting device.

**Developer Productivity**

- Easy Programming Model: ASP.NET makes building real-world Web applications dramatically easier with server controls that let you build great pages with far less code than classic ASP.

- Flexible Language Options. ASP.NET supports not only Microsoft Visual Basic Scripting Edition (VBScript) and Microsoft JScript but also more than 25 .NET languages, including built-in support for Visual Basic .NET, Microsoft C#, and JScript .NET.

- Rich Class Framework: The .NET Framework class library offers over 4,500 classes that encapsulate rich functionality such as XML, data access, file upload, regular expressions, image generation, performance monitoring and logging, transactions, message queuing, and SMTP mail.

# Introduction to XML Web Services Created with ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

An XML Web service is a programmable entity that provides a particular function, such as application logic, and is accessible to systems using Internet standards such as XML, HTTP, and SOAP. Applications written in any language and running on any operating system can call XML Web services.

An XML Web service can be used internally by a single application or exposed externally over the Internet for use by any number of applications. Because it is accessible through a standard interface, an XML Web service allows heterogeneous systems to work together as a single web of computation. One of the core characteristics of an XML Web service is the high degree of abstraction that exists between the implementation and consumption of a service. Because an XML Web service is created and accessed using XML-based messaging, the XML Web service provider and the client need no knowledge of each other beyond inputs, outputs, and location.

The XML Web services infrastructure provides a discovery mechanism to locate XML Web services, a service description for defining how to use those services, and standard wire formats for communication.

*Discovery* is the process of locating one or more documents that describe a particular XML Web service using the Web Services Description Language (WSDL). Once a service is found, the WSDL document provides a description of the interactions that the service supports. XML Web services communicate using open wire formats, which are protocols that can be understood by any system capable of supporting the most common Web standards. SOAP is the key protocol for XML Web service communication.

ASP.NET enables you to build and publish XML Web services using the same programming constructs that are used to build and publish ASP.NET Web pages. XML Web services created with ASP.NET use the industry standards shown in the following table.

| **Industry standard** | **Use in XML Web services created with ASP.NET** |
|---|---|
| XML | The text format used when communicating with XML Web services using SOAP. When communicating with XML Web services using the HTTP-GET and HTTP-POST protocols, XML is used to encode the response. |
| SOAP | An XML-based message exchange protocol used for communication between XML Web services and their clients. |
| Web Services Description Language (WSDL) | Describes the contract of messages that an XML Web service can interpret when communicating with an XML Web service client. |
| XML Schema Definition language (XSD) | Provides a universal type system, allowing data types to be defined and passed across platforms. For an XML Web service, XSD defines the structure and data types for the XML encapsulated within a SOAP message sent to and from an XML Web service. |
| application/x-www-form-url encoded | A MIME type used for encoding parameters on a URL. This encoding is used for encoding request parameters to XML Web services using the HTTP-GET and HTTP-POST protocols. |

# ASP.NET Administration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

| **Important** |
| --- |
| This feature is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server™ 2003 family. For more information, see "Features unavailable on 64-bit versions of the Windows Server family" in Help and Support Center for Windows 2003. |

This document is for administrators responsible for establishing and maintaining Microsoft® ASP.NET applications on production servers. It covers common concepts and tasks that administrators need to be familiar with, from setting up ASP.NET on server computers to monitoring ASP.NET applications once they are deployed.

This section includes the following information:

- Overview of ASP.NET Setup: Describes how to install and uninstall ASP.NET on the platforms that support it and how to identify and solve common problems with a new ASP.NET installation.
- Overview of ASP.NET Web Applications: Defines ASP.NET applications and the file types and directories they contain.
- Overview of ASP.NET Configuration: Describes how to configure ASP.NET applications and includes reference information about each of the ASP.NET configuration elements.
- Overview of ASP.NET Security: Describes how to work with authorization and authentication in ASP.NET applications.
- Overview of ASP.NET Deployment: Describes how to distribute ASP.NET applications and components from development and test servers to production servers.
- Overview of ASP.NET Process Isolation: Describes the process models that Internet Information Services (IIS) 6.0 and ASP.NET use.
- Overview of ASP.NET Performance Monitoring: Describes the tools available to monitor the performance of your ASP.NET applications and includes a list of performance counters that are exposed by ASP.NET.
- Additional Resources for ASP.NET: Contains a list of other sources of information about administering ASP.NET applications.

# Overview of ASP.NET Setup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Microsoft ASP.NET is installed with the Microsoft .NET Framework through Web downloads, Microsoft Visual Studio .NET installations, and as a service on all products in the Microsoft Windows Server™ 2003 family.

| **Note** |
| --- |
| ASP.NET is a part of the .NET Framework. You must install the .NET Framework to install ASP.NET. |

This section describes the necessary steps to install and enable ASP.NET on all platforms that support it. It also explains how to uninstall ASP.NET and includes a troubleshooting section that addresses common problems encountered once ASP.NET is installed.

This section includes the following information:

- Installing ASP.NET: Describes the software requirements for ASP.NET and how to install it on servers running a product in the Windows Server 2003 family.
- Uninstalling ASP.NET: Describes how to remove an ASP.NET installation from a server.
- Troubleshooting an ASP.NET Installation: Explains how to solve common problems encountered once ASP.NET is installed on a server.
- ASP.NET IIS Registration Tool: Explains how to use the ASP.NET IIS Registration tool to create script maps that point from your application to a particular version of ASP.NET and the .NET Framework.

Related Information

- ASP.NET Administration

# Installing ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP.NET is supported on the Windows Server™ 2003 family, Windows 2000 (Professional, Server, and Advanced Server), and Windows XP Professional for both client and server applications.

A server running a member of the Microsoft Windows Server 2003 family can be configured as an application server, with ASP.NET as an option that you can enable while configuring the application server role. To deploy ASP.NET Web applications to a production server, you must be sure to enable the ASP.NET and IIS roles on the production server before you distribute the application.

| **Note** |
| --- |
| If you want to install ASP.NET on a domain controller, there are special steps you must take to make the installation work correctly. For more information, see Knowledge Base article Q315158, ASP.NET Does Not Work with the Default ASPNET Account on a Domain Controller. |

ASP.NET, along with the .NET Framework version 1.1, is installed as a part of Windows Server 2003. You simply need to add it as a new program from Control Panel or enable it by using the Configure Your Server wizard.

| **Note** |
| --- |
| When you use the Configure Your Server wizard or the **Add or Remove Programs** dialog box to install ASP.NET on a server running Windows Server 2003, ASP.NET is automatically enabled in IIS Manager. However, if you install ASP.NET from a Web download or as part of an application such as Visual Studio .NET, you must enable ASP.NET manually. For instructions, see the "To enable ASP.NET in IIS Manager on a server running Windows Server 2003" procedure, later in this topic. |

**Procedures**

**To install ASP.NET on a server running Windows Server 2003 using the Configure Your Server wizard**

1. From the **Start** menu, click **Manage Your Server**; in the **Manage Your Server** window, click **Add or remove a role**.
2. In the **Configure Your Server Wizard**, click **Next**, and in the **Server Role** dialog box, check **Application Server (IIS, ASP.NET)** and then click **Next**.
3. In the **Application Server Options** dialog box, select the **Enable ASP.NET** check box, click **Next**, and then click **Next** again.
4. If necessary, insert your Windows Server 2003 installation CD in the CD-ROM drive and then click **Next**.
5. When the installation is complete, click **Finish**.

**To install ASP.NET on a server running Windows Server 2003 using the Add or Remove Programs dialog box**

1. From the **Start** menu, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In the **Add or Remove Programs** dialog box, click **Add/Remove Windows Components**.
3. In the **Components** box in the Windows Components wizard, select the **Application Server** check box, and then click **Next**.
4. When the Windows Components wizard has finished configuring Windows Server 2003, click **Finish**.

**To enable ASP.NET in IIS Manager on a server running Windows Server 2003**

1. From the **Start** menu, click **Run**.
2. In the **Open** box in the **Run** dialog box, type **inetmgr**, and then click **OK**.
3. In IIS Manager, expand the local computer and then click **Web Service Extensions**.

4. In the details pane, right-click **ASP.NET** and then click **Allow**. The status of ASP.NET changes to **Allowed**.

| **Note** |
| --- |
| If your applications require ISAPI, CGI, ASP, WebDAV, or other extensions to operate, you must enable those extensions also. |

**Related Information**

- For general information about ASP.NET setup, see Overview of ASP.NET Setup.
- For information about uninstalling ASP.NET, see Uninstalling ASP.NET.
- For information about troubleshooting an installation of ASP.NET, see Troubleshooting an ASP.NET Installation.

# Uninstalling ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

On Windows Server™ 2003, you can use the Configure Your Server wizard or use the Add or Remove Programs dialog box in Control Panel to disable or uninstall ASP.NET.

| **Note** |
| --- |
| There are specific issues that you need to understand if you are uninstalling a version of ASP.NET from a server with multiple versions. For more information, see Managing Multiple Versions of ASP.NET on MSDN. |

The following procedures describe how to remove ASP.NET and the .NET Framework from all server platforms that support them. You might also need to uninstall to troubleshoot an installation of ASP.NET.

**Procedures**

**To uninstall ASP.NET from a computer running Windows Server 2003 using the Configure Your Server wizard**

1. From the **Start** menu, click **Manage Your Server**; in the **Manage Your Server** window, click **Add or remove a role**.
2. In the Configure Your Server wizard, click **Next**, and in the **Server Role** dialog box, clear **Application Server (IIS, ASP.NET)**, and then click **Next**.
3. When the setup program is complete, click **Finish**.

**To uninstall ASP.NET and the .NET Framework from a computer running Windows Server 2003**

1. From the **Start** menu, point to **Control Panel**, and then click **Add or Remove Programs**.
2. In the **Currently installed programs** box, click **Microsoft .NET Framework (***language***)** *version*, and then click the **Change/Remove** button.

| **Note** |
| --- |
| In the program name, *language* is a placeholder that represents the language of the version of the .NET Framework that you want to remove from your server, and *version* represents the version number of the installation you want to remove from your server. |

1. In the **Add or Remove Programs** dialog box, click **Yes**.
2. Answer any questions from the setup program.

**Related Information**

- For general information about ASP.NET setup, see Overview of ASP.NET Setup.
- For information about installing ASP.NET, see Installing ASP.NET.
- For information about troubleshooting an installation of ASP.NET, see Troubleshooting an ASP.NET Installation.

# Troubleshooting an ASP.NET Installation

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are a few problems that commonly occur with new ASP.NET installations:

- You have installed ASP.NET and the .NET Framework on the server, but when you request a page from your application, nothing happens and the request does not resolve.

  You need to install and enable IIS on the server before you install the .NET Framework. Either you can uninstall the .NET Framework from the server, install IIS, then reinstall the .NET Framework, or you can install IIS, then use the ASP.NET IIS Registration tool (Aspnet_regiis.exe) to configure the script maps that IIS uses for ASP.NET file name extensions. For more information, see Uninstalling ASP.NET, Setting Application Mappings in IIS 6.0, and the "To repair IIS mappings for ASP.NET using Aspnet_regiis.exe" procedure later in this topic.

- You have installed ASP.NET and the .NET Framework with IIS installed and enabled, then uninstalled and reinstalled IIS. When you request a page from your application, nothing happens and the request does not resolve.

  When you uninstalled and reinstalled IIS, the script maps that IIS uses for ASP.NET were broken. Use Aspnet_regiis.exe to configure the script maps that IIS uses for ASP.NET file name extensions. For more information, see the "To repair IIS mappings for ASP.NET using Aspnet_regiis.exe" procedure later in this topic.

- You have installed and enabled IIS, installed ASP.NET and the .NET Framework, deployed your application, and requested a page, but you receive one of the following error messages:

  - Access denied to 'C:\Inetpub\Wwwroot' directory. Failed to start monitoring directory changes.

  - Server cannot access application directory 'C:\Inetpub\Wwwroot\ *Virtual Directory Name* \'. The directory does not exist or is not accessible because of security settings.

  The proper permissions have not been set for the Web site or virtual directory. ASP.NET requires read, execute, and list access for the ASPNET account (the Aspnet_wp.exe process account) on the root Web site or on any virtual directory. These settings are necessary so that ASP.NET can access the content files and can monitor file changes. To correct this problem, follow the steps in the following procedure.

**Procedures**

**To add read, execute, and list permissions for the ASPNET account on the root Web site or on a virtual directory**

1. In Windows Explorer, browse to the folder that contains the root Web site (which is C:\Inetpub\Wwwroot by default) or the virtual directory.
2. Right-click the folder and then click **Properties**.
3. On the **Security** tab, click **Add**.
4. Type *ComputerName***\ASPNET** (for example, on a computer named Webdev, type **Webdev\ASPNET**) and then click **OK**.
5. Allow the following permissions for the ASPNET account: Read & Execute, List Folder Contents, and Read.
6. Click **OK** to close the **Properties** dialog box and to save the changes.

| **Note** |
| --- |
| You do not need to perform these steps if the Everyone group or the Users group has read access to the root Web site or virtual directory. |

**To repair IIS mappings for ASP.NET using Aspnet_regiis.exe**

1. From the **Start** menu, click **Run**.
2. In the **Open** box in the **Run** dialog box, type **cmd**, and then click **OK**.
3. At the command prompt in the new window, type the following line:

   **"%systemroot%\Microsoft.NET\Framework** *version* **\aspnet_regiis.exe" -i**.

4. In this path, *version* represents the version number of the .NET Framework installed on your server. You must replace this placeholder with the actual version number when you type the command.

   | **Note** |
   | --- |
   | You must include the quotation marks in the command. |

   For more information about Aspnet_regiis.exe, use repeat step 3 with **-?** as an argument, or see ASP.NET IIS Registration Tool (Aspnet_regiis.exe).

Also, on Windows Server 2003, if you have installed the .NET Framework and ASP.NET from a Web download or through Visual Studio .NET, you must enable ASP.NET manually from IIS Manager. For more information, see Installing ASP.NET.

| **Note** |
| --- |
| If you want to install ASP.NET on a domain controller, there are special steps you must take to make the installation work properly. For more information, see Knowledge Base article Q315158, ASP.NET Does Not Work with the Default ASPNET Account on a Domain Controller. |

**Related Information**

- For general information about ASP.NET setup, see Overview of ASP.NET Setup.
- For information about installing ASP.NET, see Installing ASP.NET.
- For information about uninstalling ASP.NET, see Uninstalling ASP.NET.

# ASP.NET IIS Registration Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When multiple versions of the .NET Framework are executing side by side on a single computer, the ASP.NET ISAPI version that is scriptmapped to an ASP.NET application determines which version of the common language runtime is used for the application. The ASP.NET IIS Registration tool (Aspnet_regiis.exe) allows an administrator or installation program to easily update the scriptmaps for an ASP.NET application to point to the ASP.NET ISAPI version associated with the tool. The tool can also be used to display the status of all installed versions of ASP. NET, register the ASP.NET version coupled with the tool, create client-script directories, and perform other configuration operations.

```
Aspnet_regiis [options]
```

You can specify one or more of the following options.

| Option | Description |
|---|---|
| -c | Installs the client-side scripts for ASP.NET, such as client-side validation scripts, to the aspnet_client subdirectory of each IIS site directory.<br>**Note**<br>Only the client-side scripts for the ASP.NET version associated with Aspnet_regiis.exe are installed. |
| -e | Removes the client-side scripts for ASP.NET from the aspnet_client subdirectory of each IIS site directory.<br>**Note**<br>Only the client-side scripts for the ASP.NET version associated with Aspnet_regiis.exe are removed. |
| -ea | Removes the client-side scripts for all versions of ASP.NET from the Aspnet_client subdirectory of each IIS site directory. |
| -i | Installs the version of ASP.NET associated with Aspnet_regiis.exe and updates the scriptmaps at the IIS metabase root and below.<br>**Note**<br>Only the scriptmaps for applications that use an earlier version of the ASP.NET are updated. Applications that use a later version are not affected. |
| -ir | Installs the version of ASP.NET associated with Aspnet_regiis.exe and only registers ASP.NET in IIS.<br>**Note**<br>This option does not update the scriptmaps. To install ASP.NET and update the scriptmaps, use the **-i** option. |
| -k *path* | Removes the scriptmaps to all versions of ASP.NET from all ASP.NET applications at the specified application root path and its subdirectories. |
| -kn *path* | Removes the scriptmaps to all versions of ASP.NET from the ASP.NET application only at the specified application root path.<br>**Note**<br>This option does not affect applications in subdirectories of path. |

| | |
|---|---|
| -lk | Lists the path and version of all IIS metabase keys where ASP.NET is scriptmapped.<br>**Note**<br>The keys that inherit ASP.NET scriptmaps from a parent key are not displayed. |
| -lv | Lists the status and installation path of all versions of ASP.NET installed on the computer. |
| -r | Updates all scriptmaps in the IIS metabase and below to point to the ASP.NET ISAPI version associated with Aspnet_regiis.exe.<br>**Note**<br>All existing scriptmaps are updated to point to the ASP.NET ISAPI version associated with Aspnet_regiis.exe, regardless of the current version. |
| -s *path* | Installs the scriptmap pointing to the ASP.NET ISAPI version associated with Aspnet_regiis.exe to all ASP.NET application at the specified application root path and its subdirectories. All existing scriptmaps in the specified path and below that use an earlier version of the ASP.NET ISAPI are updated. |
| -sn *path* | Installs the scriptmap pointing to the ASP.NET ISAPI version associated with Aspnet_regiis.exe to the ASP.NET application at the specified application root path. All existing scriptmaps in the specified path that use an earlier version of the ASP.NET ISAPI are updated. [[LEFT]] Note<br>**Note**<br>This option does not affect applications in subdirectories of *path*. |
| -u | Uninstalls the version of ASP.NET associated with Aspnet_regiis.exe from the computer. Existing scriptmaps to this version of the ASP.NET ISAPI are automatically remapped to the highest remaining ASP.NET ISAPI version installed. |
| -ua | Uninstalls all versions of ASP.NET from the computer. |
| -? | Displays the command syntax and options for the tool. |

**Remarks**

When multiple versions of ASP.NET are installed on a computer, ASP.NET is said to be running *side-by-side*. In this setup, Internet Information Services (IIS) needs to know which version of the ASP.NET ISAPI (aspnet_isapi.dll) should process a page in an ASP.NET application. The ASP.NET ISAPI version associated with an ASP.NET application determines which version of the common language runtime is used for the application. An ASP.NET application is associated with an ASP.NET ISAPI version through a scriptmap in IIS. To simplify the configuration process for an ASP.NET application, each version of ASP.NET comes with a linked version of Aspnet_regiis.exe.

**Note**

A unique version of Aspnet_regiis.exe is included with each version of the .NET Framework. Because each version of the tool is applicable only to its associated version of the .NET Framework, be sure to use the appropriate version of the tool to configure an ASP.NET application.

Aspnet_regiis.exe is commonly used with the **-s** or **-sn** option to remap an ASP.NET application to the version of the .NET Framework that is associated with the tool. Use the **-s** option to update the applications at the specified root path and in all its subdirectories. If you don't want applications in subdirectories to be updated, use the **-sn** option. To update the scriptmaps of all existing ASP.NET applications on the computer at once, use the **-r** option.

**Note**

The *path* parameter refers to the root path of the application, not the physical path. For example, `W3SVC/1/ROOT/SampleApp1`.

Conversely, you can use the tool to remove the scriptmaps to any version of ASP.NET from an application using the **-k** or **-kn** option and specifying the root path of the application.

| **Note** |
| --- |
| If the root path specified inherits its scriptmap from a parent root path, the **-k** and **-kn** options will have no effect. |

The tool can also be used to install and uninstall the linked version of ASP.NET. Use the **-i** option to install ASP.NET and to update the scriptmaps of all existing ASP.NET applications. Use the **-ir** option to install ASP.NET without updating the scriptmaps. To uninstall the ASP.NET version associated with the tool, use the **-u** option. If you want to uninstall all versions of ASP.NET from the computer, use the **-ua** option.

You can use Aspnet_regiis.exe to view information about ASP.NET. To list the status and installation path of all installed versions of ASP.NET, use the **-lv** option. If you would like to see the paths to all IIS metabase keys where ASP.NET is scriptmapped, use the **-lk** option.

Client-side scripts, such as client-side validation, can be installed and removed using Aspnet_regiis.exe. To install the client-side script for the version of ASP.NET associated with the tool to the aspnet_client subdirectory of each IIS site directory, use the **-c** option. To remove the client-side script for just the ASP.NET version associated with the tool, use the **-e** option. To remove the client-side script for all installed versions of ASP.NET, use the **-ea** option.

For more information about side-by-side execution in ASP.NET, see Side-by-Side Support in ASP.NET on MSDN. For more information about script maps and application root path, see Setting Application Mappings in IIS 6.0.

**Examples**

The following command installs the scriptmaps pointing to the ASP.NET version associated with Aspnet_regiis.exe to the SampleApp1 application and all its sub-applications.

```
Aspnet_regiis -s W3SVC/1/ROOT/SampleApp1
```

The following command only updates the scriptmaps for the SampleApp1 application, without affecting applications in subdirectories.

```
Aspnet_regiis -sn W3SVC/1/ROOT/SampleApp1
```

The following command installs the ASP.NET version associated with the tool and updates the scriptmaps of all existing ASP.NET applications. Note that only applications that are currently scriptmapped to an earlier version of ASP.NET are affected.

```
Aspnet_regiis -i
```

The following command installs the ASP.NET version associated with the tool, but does not update the scriptmaps of existing ASP.NET applications

```
Aspnet_regiis -ir
```

The following command displays the status and installation path of all versions of ASP.NET installed on the computer.

```
Aspnet_regiis -lv
```

**Related Information**

- For general information about ASP.NET setup, see Overview of ASP.NET Setup.
- For information about installing ASP.NET, see Installing ASP.NET.
- For information about uninstalling ASP.NET, see Uninstalling ASP.NET.
- For more information about side-by-side operation of multiple versions of ASP.NET, see "Side-by-Side Support in ASP.NET" on MSDN.
- For information about troubleshooting an installation of ASP.NET, see Troubleshooting an ASP.NET Installation.

# Overview of ASP.NET Web Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A Microsoft ASP.NET application is defined as all the files and executable code that can be invoked from a virtual directory and its subdirectories on a Web server. These can include Web pages (.html files), Web Forms pages (.aspx files), Web Forms user controls (.ascx files), XML Web Services (.asmx files), HTTP handlers, HTTP modules, and other files (such as image and configuration files). Any script maps that are associated with the version of the Microsoft .NET Framework that you are using are also a part of an ASP.NET application.

If an ASP.NET application has been developed using Microsoft Visual Studio .NET, a project DLL is stored in the \Bin directory under the application's root directory. Any other DLLs that provide functionality that the application uses, whether it has been developed using Visual Studio or not, must reside in the \Bin directory as well. These can include compiled ASP.NET custom server controls, business logic, and database access logic.

The relationship between an ASP.NET application and Internet Information Services (IIS) is this: IIS serves all requests for publishable ASP.NET files through aspnet_isapi.dll (the process model for ASP.NET). IIS does not process the server-side code; the ASP.NET engine does that and then returns the output to IIS (or another Web server, possibly) to be served from that software.

This section includes the following information

- Layout of ASP.NET Web Applications: Describes the guidelines for creating a file structure for an ASP.NET application.
- Viewing the ASP.NET Version: Describes how to determine which version of the .NET Framework, and ASP.NET, an application is using.

**Related Information**

- ASP.NET Administration
- Overview of ASP.NET Deployment
- Overview of ASP.NET Configuration

# Layout of ASP.NET Web Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

An ASP.NET application must reside in an IIS virtual directory, known as the application root. ASP.NET applications can contain compiled assemblies (usually DLL files that contain business logic); a well-known directory, always named \Bin, where precompiled code is stored; configuration settings stored in a text-based, human-readable Web.config file; pages; server controls; and XML Web services.

Any precompiled code that is not shared with other applications on the server must be stored in the application's \Bin directory. This is the local assembly cache for the application.

| **Note** |
| --- |
| If you intend to use the same assembly in multiple applications on your Web server, you should deploy that assembly to the global assembly cache. For more information, see Deploying Shared Components to the Global Assembly Cache. |

Web.config files store application-wide configuration settings in an XML-based text file. This means that you can use any standard text editor or XML parser to create them, and that they are readable. If you do not include a Web.config file in the application root, configuration settings are determined by those set in the Machine.config file, the configuration file for the entire server. A version of the Machine.config file is installed when you install the .NET Framework. For more information about configuration files, see Overview of ASP.NET Configuration.

The following illustration shows a sample file system layout for an ASP.NET application.




The application in this illustration includes two DLL files in the \Bin directory, a Default.aspx page, a user control named Menu.ascx, an XML Web service named MyWebService.asmx, and a Global.asax file. In addition, the application is configured using three configuration files: the machine-level Machine.config file in the system root directory, the site-level Web.config file in the C:\Inetpub\Wwwroot directory, and the application-level Web.config file in the application root directory. Configuration settings stored in the application root override the settings in the site's Web.config file and the Machine.config file, while the site's Web.config file overrides settings in the Machine.config file.

| **Note** |
| --- |
| ASP.NET applications are composed of virtual directory structures, not physical directory structures. Configuration settings are applied to virtual paths. |

**Related Information**

- Overview of ASP.NET Web Applications
- Overview of ASP.NET Deployment
- Overview of ASP.NET Configuration
- Viewing the ASP.NET Version

# Viewing the ASP.NET Version

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The version of the Framework used by an ASP.NET application is determined by the version of the aspnet_isapi.dll file that is registered under the IIS script maps for that application. A script map is a three-way association of a file name extension, such as .aspx; an HTTP verb associated with a request, such as GET; and the DLL that contains instructions for what the Web server should do when it receives an HTTP request for a file with the specified file name extension. You can set these script maps to the application root directory itself, which is an IIS virtual directory, or allow them to be inherited from a parent application or site.

**Procedures**

**To view script maps for file name extensions in your application**

1. From the **Start** menu, click **Run**, type **inetmgr** in the **Open** box, and then click **OK**.
2. In IIS Manager, expand the **local computer**, expand **Web Sites**, expand **Default Web Site**, and then click the name of your application.
3. Right-click the application name, and then click **Properties**.
4. On the **Directory** tab, click **Configuration**. All the script maps associated with your application appear alphabetically by file name extension in the **Application Mappings** box.

**Configuring Script Maps**

ASP.NET provides a command-line tool, called the ASP.NET IIS Registration tool (Aspnet_regiis.exe), to make it easy to configure script maps for your applications. If you have multiple versions of ASP.NET installed on your Web server, a version of this tool is provided with each. You can use the tool associated with a particular version of ASP.NET to register your application with that version.

The **-s** and **-sn** commands included in Aspnet_regiis.exe install the current version to the specified path, included as an argument on the command line. The **-s** command installs the specified version of ASP.NET to the specified application root and all its virtual subdirectories. This command updates all existing script maps in the specified path and below that use an earlier version of the ASP.NET ISAPI. The **-sn** command installs the specified version of ASP.NET to the directory specified in the path argument only; it does not install the version to subdirectories. For example, the following entry on the command line creates a script map between the version of ASP.NET associated with the tool version and the `SampleApp1` directory, along with all its subdirectories:

```
aspnet_regiis.exe -s W3SVC/1/ROOT/SampleApp1
```

Aspnet_regiis.exe has other utility purposes as well, such as registering a version of ASP.NET (which includes registering the Aspnet_isapi.dll, setting up the state server, setting up the worker process account, and so on) and creating client script directories for that version of ASP.NET. It can also list all installed versions of ASP.NET and which applications use each version. For more information about this tool, see ASP.NET IIS Registration Tool.

For more information about running multiple ASP.NET applications on a server using different versions of the .NET Framework, see Side-by-Side Support in ASP.NET on MSDN.

**Related Information**

- Overview of ASP.NET Web Applications
- Overview of ASP.NET Deployment
- Layout of ASP.NET Web Applications

# Overview of ASP.NET Configuration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The Microsoft® ASP.NET configuration system features an extensible infrastructure that enables you to define configuration settings when your ASP.NET applications are first deployed so that you can add or revise configuration settings at any time with minimal impact on operational Web applications and servers.

The ASP.NET configuration system provides the following benefits:

- Configuration information is stored in XML-based text files. You can use any standard text editor or XML parser to create and edit ASP.NET configuration files.

- Multiple configuration files, all named Web.config, can appear in multiple directories on an ASP.NET Application server. Each Web.config file applies configuration settings to its own directory and all child directories below it. Configuration files in child directories can supply configuration information in addition to that inherited from parent directories, and the child directory configuration settings can override or modify settings defined in parent directories. The root configuration file named *systemroot*\Microsoft.NET\Framework*versionNumber*\CONFIG\Machine.config provides ASP.NET configuration settings for the entire Web server.

- At run time, ASP.NET uses the configuration information provided by the Web.config files in a hierarchical virtual directory structure to compute a collection of configuration settings for each unique URL resource. The resulting configuration settings are then cached for all subsequent requests to a resource. Note that inheritance is defined by the incoming request path (the URL), not the file system paths to the resources on disk (the physical paths).

- ASP.NET detects changes to configuration files and automatically applies new configuration settings to Web resources affected by the changes. The server does not have to be rebooted for the changes to take effect. Hierarchical configuration settings are automatically recalculated and recached whenever a configuration file in the hierarchy is changed. The <processModel> section is an exception.

- The ASP.NET configuration system is extensible. You can define new configuration parameters and write configuration section handlers to process them.

- ASP.NET helps protect configuration files from outside access by configuring Internet Information Services (IIS) to help prevent direct browser access to configuration files. HTTP access error 403 (forbidden) is returned to any browser attempting to request a configuration file directly.

This section includes the following information:

- ASP.NET Configuration Files Format: Explains how ASP.NET configuration files are organized.

- Hierarchical Configuration Architecture in ASP.NET: Explains configuration setting inheritance.

**Related Information**

- ASP.NET Administration

# ASP.NET Configuration Files Format

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Configuration information for ASP.NET resources is contained in a collection of configuration files, each named Web.config. Each configuration file contains a nested hierarchy of XML tags and subtags with attributes that specify the configuration settings. Because the tags must be well-formed XML, the tags, subtags, and attributes are case-sensitive. Tag names and attribute names are camel-cased, which means that the first character of a tag name is lowercase and the first letter of any subsequent concatenated words is uppercase. Attribute values are Pascal-case, which means that the first character is uppercase and the first letter of any subsequent concatenated words is uppercase. Exceptions are **true** and **false**, which are always lowercase.

All configuration information resides between the <configuration> and </configuration> root XML tags. Configuration information between the tags is grouped into two main areas: the configuration section handler declaration area and the configuration section settings area.

Configuration section handler declarations appear at the top of the configuration file between <configSections> and </configSections> tags. Each declaration contained in a <section> tag specifies the name of a section that provides a specific set of configuration data and the name of the .NET Framework class that processes configuration data in that section.

The configuration section settings area follows the <configSections> area and contains the actual configuration settings. There is one configuration section for each declaration in the <configSections> area. Each configuration section contains subtags with attributes that contain the settings for that section.

The following Web.config file example declares two configuration <section> handlers. One manages application settings and the other manages session state.

```
<configuration>
<configSections>
<section name="appSettings"
type="System.Configuration.NameValueFileSectionHandler,
System, Version=1.0.3300.0,
Culture=neutral, PublicKeyToken=b77a5c561934e089"/>
<section name="sessionState"
type="System.Web.SessionState.SessionStateSectionHandler,
System.Web, Version=1.0.3300.0, Culture=neutral,
PublicKeyToken=b03f5f7f11d50a3a"
allowDefinition="MachineToApplication"/>
</configSections>
<appSettings>
<add key="dsn" value="localhost;uid=MyUserName;pwd=;"/>
<add key="msmqserver" value="server\myqueue"/>
</appSettings>
<sessionState cookieless="true" timeout="10"/>
</configuration>
```

You need to declare a configuration section handler only once. You can place it in the server's root Machine.config file or in a Web.config file in the virtual directory containing the Web application files. Configuration files in subdirectories automatically inherit configuration handlers declared in parent directories. For more information, see Configuration Inheritance in ASP.NET.

Configuration settings are often nested together under section grouping tags. These top-level section tags typically represent the namespace to which the configuration settings apply. For example, the top-level <system.net> tag represents the settings for network classes, and the <system.web> tag represents the settings for the ASP.NET classes.

The following example shows tag nesting.

```
<configuration>
```

```
<configSections>
<sectionGroup name="system.net">
<section name="authenticationModules"
type="System.Net.Configuration.NetAuthenticationModuleHandler,
System, Version=1.0.3300.0, Culture=neutral,
PublicKeyToken=b77a5c561934e089"/>
<section name="webRequestModules"
type="System.Net.Configuration.WebRequestModuleHandler,
System, Version=1.0.3300.0, Culture=neutral,
PublicKeyToken=b77a5c561934e089"/>
</sectionGroup>
<sectionGroup name="system.web">
<section name="authorization"
type="System.Web.Configuration.AuthorizationConfigHandler,
System.Web, Version=1.0.3300.0, Culture=neutral,
PublicKeyToken=b03f5f7f11d50a3a"/>
<section name="sessionState"
type="System.Web.SessionState.SessionStateSectionHandler,
System.Web, Version=1.0.3300.0, Culture=neutral,
PublicKeyToken=b03f5f7f11d50a3a"
allowDefinition="MachineToApplication"/>
</sectionGroup>
</configSections>
<system.net>
<! -- Net Class Settings would go here. -->
</system.net>
<system.web>
<authorization>
<allow users="*"/> <!-- Allow all users -->
<!-- Allow or deny specific users.
allow users="[comma separated list of users]"
roles="[comma separated list of roles]"/>
<deny users="[comma separated list of users]"
roles="[comma separated list of roles]"/>
-->
</authorization>
<sessionState
sqlConnectionString="data source=localhost;
Integrated Security=SSPI;
Initial Catalog=northwind"
cookieless="false"
timeout="10"/>
</system.web>
</configuration>
```

The ASP.NET configuration infrastructure makes no assumptions about the types of configuration data that the infrastructure supports. Configuration section handler classes process all Web.config data. You can use the predefined configuration section handlers that are supplied with the .NET Framework or you can create your own handlers to process your own custom

configuration data.

# Hierarchical Configuration Architecture in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP.NET applies configuration settings to resources in a hierarchical manner. Site-wide Web.config configuration files supply configuration information to the directories in which they are located and to all child directories. The configuration settings for a Web resource are supplied by the configuration file located in the same directory as the resource and by all configuration files in all parent directories.

Configurations settings are inherited from file to file in a virtual path, with the Machine.config file, which helps control configuration settings for the entire server, at the root. Web.config files that are stored in the virtual path help control their sites and virtual roots.

**In This Section**

Configuration Inheritance in ASP.NET: Describes how to design virtual paths so that configuration problems will not arise in your ASP.NET applications.

Configuration <location> Settings in ASP.NET: Describes how to customize configurations settings to apply to a specific directory or resource in that directory.

Locking Configuration Settings in ASP.NET: Describes how to make configuration settings on a site unchangeable, particularly for administrators who host Web applications.

**Related Sections**

- For general information about ASP.NET configuration, see Overview of ASP.NET Configuration.
- For guidelines on creating a file structure for an ASP.NET application, see Layout of ASP.NET Web Applications.

# Configuration Inheritance in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When the server receives a request for a particular Web resource, ASP.NET computes the configuration settings for that resource hierarchically, using all the configuration files located in the virtual directory path for the requested URL. For example, you could have a Web site with the following file structure, where the application root directory is the application virtual directory (Vdir).

Normally, the last configuration setting overwrites settings for the same section provided in parent directories, but the section handler might implement a different inheritance scheme.

For example, as the administrator, you could configure the application settings to allow all users access to the ASP.NET resources in the application root directory (the physical directory mapped by the application virtual directory) but allow only selected users access to the ASP.NET resources in the two subdirectories.

Assume that there is a Web.config file in the SubDir1 directory and none in the application root or SubDir2 directories. In this case, ASP.NET is using two configuration files. The highest-level file is the one located in the *systemroot*\Microsoft .NET \Framework*versionNumber*\CONFIG directory. This file, named Machine.config, is at the machine level, and all ASP.NET directories and subdirectories inherit its settings. Machine.config is shipped with the .NET Framework and contains many default ASP.NET settings. The default setting for the security configuration section of this file allows all users to access all URL resources. There is no configuration file in the example's application root directory that modifies security, so all users have access to the ASP.NET resources in it (because that directory inherits from the machine-level configuration file). If the Web.config file in the SubDir1 directory contains a security configuration section that allows access only to certain users, then SubDir2 inherits that setting. So, all users have access to the ASP.NET resources in the application root directory, but only certain users have access to the ASP.NET resources in SubDir1 and SubDir2.

Configuration settings for virtual directories are independent of physical directory structure, and virtual directories must be organized carefully to avoid configuration problems. For example, you could have an application named MyResource.aspx with the following physical directory structure: C:\Subdir1\Subdir2\MyResource.aspx.

Assume that a configuration file is located in Subdir1 and a virtual directory named Vdir1 is mapped to c:\Subdir1 and a virtual directory named Vdir2 is mapped to c:\Subdir1\Subdir2. If a client accesses the resource physically located at c:\Subdir1\Subdir2\MyResource.aspx using the URL http://localhost/vdir1/subdir2/MyResource.aspx, the resource inherits configuration settings from Vdir1. However, if the client accesses the same resource using the URL http://localhost/vdir2/MyResource.aspx, it does not inherit settings from Vdir1. Therefore, creating virtual directories in this manner could cause unexpected results or even an application failure and is not recommended.

| **Note** |
| --- |
| The ASP.NET configuration system applies only to ASP.NET resources (those registered to be handled by ASP.NET using the Aspnet_isapi.dll). The configuration system does not provide authorization for non-ASP.NET resources by default. ASP, HTML, TXT, GIF, and JPEG files, for example, are accessible by all users. In the preceding example, if directory browsing is enabled and no other restrictions are in place, all users can view non-ASP.NET files located in the application root directory, SubDir1, and SubDir2. For more information about ASP.NET security, see Overview of ASP.NET Security. |

**Related Information**

- Hierarchical Configuration Architecture in ASP.NET

# Configuration <location> Settings in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Configuration settings can be applied to specific resources by using a **<location>** tag with an appropriate **path** attribute. The **path** attribute can be used to identify a specific file or child directory to which unique configuration settings apply.

For example, the following configuration file specifies settings at three levels:

- Settings that apply to the current directory and all child directories (everything contained within the top <configuration> tag).
- Settings that apply to the Sub1 child directory (everything contained within the <location> tag with a path attribute set to Sub1).
- Settings that apply to the Sub2 child directory (everything contained within the <location> tag with a path attribute set to Sub2).

```
<configuration>
<system.web>
<sessionState cookieless="true" timeout="10"/>
</system.web>
<!-- Configuration for the "Sub1" subdirectory. -->
<location path="sub1">
<system.web>
<httpHandlers>
<add verb="*" path="Sub1.Scott" type="Sub1.Scott"/>
<add verb="*" path="Sub1.David" type="Sub1.David"/>
</httpHandlers>
</system.web>
</location>
<!-- Configuration for the "Sub2" subdirectory. -->
<location path="sub2">
<system.web>
<httpHandlers>
<add verb="*" path="Sub2.Scott" type="Sub2.Scott"/>
<add verb="*" path="Sub2.David" type="Sub2.David"/>
</httpHandlers>
</system.web>
</location>
</configuration>
```

**Related Information**

Hierarchical Configuration Architecture in ASP.NET

# Locking Configuration Settings in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, configuration files located in subdirectories override and extend all configuration settings defined in parent configuration files. In application hosting scenarios, administrators often want to lock or make some settings on a site unchangeable to prevent modification. For example, an administrator might want to lock the sandbox security settings for hosted applications to help prevent Web users from attacking the system.

Administrators can lock configuration settings by adding an **allowOverride="false"** attribute to a **<location>** directive. This tells the configuration system to throw an error if a lower-level configuration file attempts to override any configuration section defined within this locked **<location>** directive.

The following configuration file example (which could be stored at either the main system level or at the site level) locks the trust level of two different ASP.NET applications (`application1` and `application2`).

```
<configuration>
<location path="application1" allowOverride="false">
<system.web>
<trust level="High"/>
</system.web>
</location>
<location path="application2" allowOverride="false">
<system.web>
<trust level="Medium"/>
</system.web>
</location>
</configuration>
```

Any attempt to override the configuration settings in the preceding example by the configuration settings in the following example would generate a configuration system error.

```
<configuration>
<system.web>
<trust level="Full"/>
</system.web>
</configuration>
```

**Related Information**

- Hierarchical Configuration Architecture in ASP.NET

# ASP.NET Settings Schema Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The ASP.NET configuration section schema contains elements that help control how ASP.NET Web applications behave. Where a default value is specified for an attribute, the default value is set in the Machine.config file that is located at *systemroot*\Microsoft.NET\Framework*versionNumber*\Config\Machine.config.

For detailed information about the ASP.NET settings schema elements, see ASP.NET Settings Schema.

**Related Information**

For more information about ASP.NET settings, see Overview of ASP.NET Configuration.

# Overview of ASP.NET Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Most Web sites need to selectively restrict access to some portions of the site. You can think of a Web site as somewhat analogous to an art gallery. The gallery is open for the public to come in and browse, but there are certain parts of the facility, such as the business offices, that are accessible only to people with certain credentials, such as employees. When a Web site stores its customers' credit card information in a database, for example, ASP.NET helps protect the database from public access. ASP.NET security addresses this and many other security issues.

ASP.NET, in conjunction with Internet Information Services (IIS), can authenticate user credentials such as names and passwords using any of the following authentication methods:

- Windows: Basic, digest, or Integrated Windows Authentication (NTLM or Kerberos).
- Microsoft Passport authentication
- Forms authentication
- Client Certificate authentication

ASP.NET helps control access to site information by comparing authenticated credentials, or representations of them, to NTFS file system permissions or to an XML file that lists authorized users, authorized roles (groups), or authorized HTTP verbs.

The topics in this section describe the specifics of ASP.NET security.

This section includes the following information:

- How Security Works in ASP.NET: Provides an overview of ASP.NET security.
- Architecture in ASP.NET: Provides an overview of ASP.NET infrastructure and subsystem relationships, as related to security.
- Data Flow in ASP.NET: Describes the security data flow for two common scenarios.
- Authentication in ASP.NET: Describes ASP.NET authentication providers.
- Authorization in ASP.NET: Describes two fundamental ways to authorize access to a resource.
- Impersonation in ASP.NET: Describes how and when to use ASP.NET Impersonation.
- Authentication and ASP.NET Impersonation: Describes how to create ASP.NET application with incorporated security.
- Security of ASP.NET Applications in Hosted Environments: Describes ASP.NET security features for multi-application Web servers.

**Related Information**

- ASP.NET Administration

# How Security Works in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Helping protect Web sites against unauthorized access is a critical, complex issue for Web developers. A successful system requires careful planning, and Web site administrators and programmers must have a clear understanding of the options for securing their site.

ASP.NET works in concert with the Microsoft .NET Framework and Internet Information Services (IIS) to help provide Web application protection. To help protect an ASP.NET application, you must perform the two fundamental functions described in the following table.

| Security function | Description |
|---|---|
| Authentication in ASP.NET | Assures that the user is, in fact, who the user claims to be. The application obtains credentials (various forms of identification, such as name and password) from a user and validates those credentials against some authority. If the credentials are valid, the entity that submitted the credentials is considered an authenticated identity. |
| Authorization in ASP.NET | Limits access rights by granting or denying specific permissions to an authenticated identity. |

IIS can also grant or deny access based on a user's host name or IP address. Any further access authorization is performed by NTFS file access permission's URL authorization.

It is helpful to understand how all the various security subsystems interact. Since ASP.NET is built on the Microsoft .NET Framework, the ASP.NET application developer also has access to all the built-in security features of the .NET Framework, such as code access security and role-based user-access security. For details about the security capabilities of ASP.NET, see ASP.NET Code Access Security.

**Related Information**

- Overview of ASP.NET Security

# Architecture in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section provides an overview of the ASP.NET infrastructure and subsystem relationships, as they relate to the subject of security. The following illustration shows the relationships among the security systems in ASP.NET.




As the illustration shows, all Web clients communicate with ASP.NET applications through Internet Information Services (IIS). IIS deciphers and optionally authenticates the request. If Allow Anonymous is set to **true**, no authentication occurs. IIS also finds the requested resource (such as an ASP.NET application), and, if the client is authorized, returns the appropriate resource.

In addition to the built-in ASP.NET features, an ASP.NET application can use the low-level security features of the .NET Framework. For more information, see the "Key Security Concepts" topic in .NET Framework Help.

**Integrating with IIS**

When considering ASP.NET authentication, you should understand the interaction with IIS authentication services.

IIS always assumes that a set of credentials maps to a Microsoft Windows NT account and uses them to authenticate a user. There are three different kinds of authentication available in IIS 5.0 through IIS 6.0: basic, digest, and Integrated Windows Authentication (NTLM or Kerberos). You can select the type of authentication to use in IIS administrative services. For more information about IIS authentication, see Authentication in IIS 6.0.

If you request a URL containing an ASP.NET application, the request and authentication information are handed off to the application. ASP.NET provides the two additional types of authentication described in the following table.

| ASP.NET authentication provider | Description |
|---|---|
| Forms authentication | A system by which unauthenticated requests are redirected to an HTML form using HTTP client side redirection. The user provides credentials and submits the form. If the application authenticates the request, the system issues an authentication ticket in a cookie that contains the credentials or a key for reacquiring the identity. Subsequent requests are issued with the cookie in the request headers; they are authenticated and authorized by an ASP.NET handler using whatever validation method the application developer specifies. |
| Passport authentication | Centralized authentication service provided by Microsoft that offers a single logon and core profile services for member sites. |

**Using ASP.NET Configuration Files**

ASP.NET configuration, of which security is a part, has a hierarchical architecture. All configuration information for ASP.NET is contained in files named Web.config and Machine.config. Web.config can be placed in the same directories as the application files. The Machine.config file is in the Config directory of the install root. Subdirectories inherit a directory's settings unless overridden by a Web.config file in the subdirectory. In a Web.config file, there are sections for each major category of ASP.NET functionality. To see an example of the way in which the hierarchical configuration system works for security, see Hierarchical Configuration Architecture in ASP.NET.

The security section of a Web.config file is organized as follows.

```
<authentication mode="[Windows|Forms|Passport|None]">

<forms name="[name]"
```

```
loginUrl="[url]"
protection="[All|None|Encryption|Validation]"
path="[path]" timeout="[minutes]"
requireSSL="[true|false]"
slidingExpiration="[true|false]">
<credentials passwordFormat="[Clear|MD5|SHA1]">
<user name="[UserName]"
password="[password]"/>
</credentials>
</forms>
<passport redirectUrl="internal"/>
</authentication>
<authorization>
<allow users="[comma separated list of users]"
roles="[comma separated list of roles]"/>
<deny users="[comma separated list of users]"
roles="[comma separated list of roles]"/>
</authorization>
<identity impersonate ="[true|false]"
userName="[domain\user_name]"
password="[user_password]"/>
<trust level="[Full|High|Medium|Low|Minimal]"
originUrl=""/>
<securityPolicy>
<trustLevel name="Full" policyFile="internal"/>
<trustLevel name="High" policyFile="web_hightrust.config"/>
<trustLevel name="Medium" policyFile="web_mediumtrust.config"/>
<trustLevel name="Low" policyFile="web_lowtrust.config"/>
<trustLevel name="Minimal" policyFile="web_minimaltrust.config"/>
</securityPolicy>
```

The default settings for these elements are shown in the following table.

| **Default value** | **Description** |
|---|---|
| <allow roles= > | No default value. |
| <allow users="*"> | All |
| <authentication mode="Windows"> | The authentication mode cannot be set at a level below the application root directory. |
| <credentials passwordFormat="SHA1"> | The hashing algorithm to be used on passwords. |
| <deny roles=""> | Empty |
| <deny users=""> | Empty |

| | |
|---|---|
| <forms loginUrl="login.aspx"> | If you set the mode to **Forms**, and if the request does not have a valid ticket (cookie), this is the URL to which the request is directed for a forms-based logon. |
| <forms name=".ASPXAUTH"> | Default cookie name. |
| <forms path="/"> | Path |
| <forms protection="All"> | Type=[All\|None\|Encryption\|Validation]. |
| <forms requireSSL="false"> | Specifies that an encrypted connection is not required to transmit the authentication cookie. |
| <forms slidingExpiration="true"> | Specifies that sliding expiration is enabled. |
| <forms timeout="30"> | Time-out in minutes. 30 minutes is the default. |
| <identity impersonate="false"> | Impersonation is disabled by default. |
| <identity userName=""> | Empty |
| <identity password=""> | Empty |
| <passport redirectUrl="internal"> | If you set the mode to **Passport**, and if the requested page requires authentication (anonymous users are denied access) but the user has not logged on with Passport, then the user will be redirected to this URL. |
| <trustLevel name="Full" policyFile="internal"/> | Default policy file for Full trust level. |
| <trustLevel name="High" policyFile="web_hightrust.config"/> | Default policy file for High trust level. |
| <trustLevel name="Medium" policyFile="web_mediumtrust.config"/> | Default policy file for Medium trust level. |
| <trustLevel name="Low" policyFile="web_lowtrust.config"/> | Default policy file for Low trust level. |
| <trustLevel name="Minimal" policyFile="web_minimaltrust.config"/> | Default policy file for Minimal trust level. |
| <user name=""> | Empty |
| <user password=""> | Empty |

There are three major subsections to a Web.config file: authentication, authorization, and identity. The values for each of the security elements are usually set by overriding a section of the computer-level configuration file with a similar section in an application configuration file placed in the application root directory. All subdirectories automatically inherit those settings. However, subdirectories can have their own configuration files that override other settings.

| **Note** |
|---|
| ASP.NET configuration applies only to ASP.NET resources (those registered to be handled by Aspnet_isapi.dll). Unfortunately, ASP.NET configuration cannot provide authorization for non-Aspnet_isapi.dll resources, so TXT, HTML, GIF, JPEG, ASP, and other types of files are still accessible by all users, subject to IIS permissions. For example, although the ASP.NET resources in a directory might be restricted by a Web.config file, all users can still view the files located in that directory if directory browsing is turned on and no other restrictions are in place. |

You can avoid this situation by explicitly mapping such files, but not directories, to Aspnet_isapi.dll using the IIS administration tool. However, there could be a performance impact if you do this.

You can use the <**location**></**location**> tags to specify a particular file or directory to which settings should apply. For more information about how to use the <**location**> tag, see Configuration <location> Settings in ASP.NET and <location> Element.

**Related Information**

- Overview of ASP.NET Security

# Data Flow in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are a number of different ways to design security into ASP.NET applications. This section describes the data flow for two common scenarios: impersonation and forms authentication using cookies.

**Scenario 1: Impersonation**
This scenario relies on Internet Information Services (IIS) authentication and Microsoft Windows NT® file access security to minimize security programming in the ASP.NET application itself. The data flow is shown in the following illustration.

The illustration shows the following sequence of events:

1. A request for access comes to IIS from a network client.

2. IIS authenticates the client using basic, digest, or Integrated Windows Authentication (NTLM or Kerberos).

3. If the client is authenticated, IIS hands the authenticated request over to ASP.NET.

4. The ASP.NET application impersonates the requesting client using the access token passed from IIS, and relies on NTFS file permissions for granting access. The ASP.NET application needs only to verify that in the ASP.NET configuration file, the impersonation-enable directive is set to **true**; no ASP.NET security code needs to be written.

   Notice that if impersonation is not enabled, the application runs with the IIS process identity. For Microsoft Windows® 2000 Server and Windows XP, the default identity is a User account named ASPNET that is created automatically when ASP.NET is installed. For products in the Microsoft Windows Server 2003 family, the default identity is the Network Service account. If you want to restrict access, you must use some other means of authorization, such as URL authorization.

   For more details about using impersonation in ASP.NET applications, see Impersonation in ASP.NET and Using IIS Authentication with ASP.NET Impersonation.

5. If access is granted, the ASP.NET application returns the requested page through IIS.

**Scenario 2 - Forms Authentication**
In this scenario an application uses ASP.NET forms authentication, a process that enables the application to collect credentials such as name and password directly from the client requestor and make its own determination about their authenticity. IIS authentication is not used by the application, but IIS authentication settings are important to the ASP.NET forms authentication process. Unless you decide to reject all requests that do not meet the criteria for the enabled method of IIS authentication, you must enable the IIS **Anonymous Access** setting.

| **Note** |
| --- |
| If you do not enable anonymous access in IIS, requests not meeting the criteria for IIS authentication will be rejected and never reach the ASP.NET application. |

The data flow in this scenario is shown in the following illustration.




This illustration shows the following sequence of events:

1. A client generates a request for a protected resource.

2. IIS receives the request, and if the requestor is authenticated by IIS, or if IIS anonymous access is enabled, the request gets passed on to the ASP.NET application. Because the authentication mode in the ASP.NET application is set to forms in this case, IIS authentication is not used.

3. If there is no cookie attached to the request, ASP.NET redirects the request to a logon page, the path of which resides in the application's configuration file. On the logon page, the client user enters the required credentials (usually a name and password).

4. The application code checks the credentials to confirm their authenticity, usually in an event handler. If the credentials are authenticated, the application code attaches a ticket (as a cookie) containing the user name, but not the password. If authentication fails, the request is usually returned with an Access Denied message or the logon form is presented again.

5. After a ticket is issued by the application, ASP.NET just checks the ticket for validity using a message authentication check. Applications do not need the credentials in the *.config files. In fact, ASP.NET does not check them after the cookie is issued, even if they are present.

6. If the user is authenticated, ASP.NET checks authorization and can either allow access to the originally requested, protected resource or redirect the request to some other page, depending on the design of the application. It can also direct the request to a custom authorization module where the credentials are tested for authorization to access the protected resource. If authorization fails, ASP.NET always redirects to the logon page.

7. If the user is authorized, access is granted to the protected resource; or the application might require an additional test of the credentials before authorizing access to the protected resource, depending on the design of the application.

**Related Information**

- Overview of ASP.NET Security

# Authentication in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

Authentication is the process of obtaining identification credentials such as name and password from a user and validating those credentials against some authority. If the credentials are valid, the entity that submitted the credentials is considered an authenticated identity. Once an identity has been authenticated, the authorization process determines whether that identity has access to a given resource.

ASP.NET implements authentication through authentication providers, the code modules that contain the code necessary to authenticate the requestor's credentials. ASP.NET supports the authentication providers described in the following table.

| **ASP.NET authentication provider** | **Description** |
|---|---|
| Forms authentication | A system by which unauthenticated requests are redirected to an HTML form using HTTP client-side redirection. The user provides credentials and submits the form. If the application authenticates the request, the system issues a cookie that contains the credentials or a key for reacquiring the identity. Subsequent requests are issued with the cookie in the request headers; they are authenticated and authorized by an ASP.NET event handler using whatever validation method the application developer specifies. |
| Passport authentication | Centralized authentication service provided by Microsoft that offers a single logon and core profile services for member sites. |
| Windows authentication | ASP.NET uses Windows authentication in conjunction with Internet Information Services (IIS) authentication. Authentication is performed by IIS in one of three ways: basic, digest, or Integrated Windows Authentication. When IIS authentication is complete, ASP.NET uses the authenticated identity to authorize access. |

**Important**

Integrated Windows authentication is disabled by default if you install Windows Server 2003 Service Pack 1 (SP1) as part of a slipstream installation of a Windows Server 2003 operating system. If you install Windows Server 2003 SP1 by itself as an upgrade to a Windows Server 2003 operating system, the setting for Integrated Windows authentication is unchanged from its Windows Server 2003 setting. Integrated Windows authentication is enabled by default for Windows Server 2003 operating systems.

To enable an authentication provider for an ASP.NET application, create an entry for the application configuration file as follows.

```
// Web.config file

<authentication mode= "[Windows|Forms|Passport|None]"/>
```

The mode is set to one of the authentication modes: **Windows**, **Forms**, **Passport**, or **None**. The default is Windows. If the mode is **None**, ASP.NET does not apply any additional authentication to the request - this can be useful when you want to implement a custom authentication scheme, or if you are solely using anonymous authentication and want the highest possible level of performance.

The authentication mode cannot be set at a level below the application root directory. As is the case with other ASP.NET modules, subdirectories in the URL space inherit authentication modules unless explicitly overridden.

**Related Information**

- Windows Authentication Provider in ASP.NET
- Passport Authentication Provider in ASP.NET
- Forms Authentication Provider in ASP.NET

# Windows Authentication Provider in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

The **WindowsAuthenticationModule** provider relies on Internet Information Services (IIS) to provide authenticated users, using any of the mechanisms that IIS supports. If you want to help protect site security with a minimum of ASP.NET coding, this is the provider configuration you should use. The provider module constructs a **WindowsIdentity** object. The default implementation constructs a **WindowsPrincipal** object and attaches it to the application context. The **WindowsPrincipal** object maps identities to Windows groups.

If you use IIS authentication, the provider module uses the authenticated identity passed in from IIS. IIS authenticates the identity using basic, digest, or Integrated Windows authentication, or some combination of them. You can use impersonation and NTFS ACL permissions to restrict or allow access to protected resources.

| **Important** |
| --- |
| Integrated Windows authentication is disabled by default if you install Windows Server 2003 Service Pack 1 (SP1) as part of a slipstream installation of a Windows Server 2003 operating system. If you install Windows Server 2003 SP1 by itself as an upgrade to a Windows Server 2003 operating system, the setting for Integrated Windows authentication is unchanged from its Windows Server 2003 setting. Integrated Windows authentication is enabled by default for Windows Server 2003 operating systems. |

An important reason to use the **WindowsAuthenticationModule** provider is to implement an impersonation scheme that can use any of the authentication methods that might have already been performed by IIS before passing the request to the ASP.NET application. To do this, set the authentication mode to Windows, and confirm that the **impersonate** element is set to **true**, as shown in the following example:

```
<authentication mode="Windows"/>
<identity impersonate="true"/>
```

Note that configuring an ASP.NET application has no effect on the IIS Directory Security settings. The systems are completely independent and are applied in sequence. In addition to selecting an authentication mode for an ASP.NET application, it is also important to configure IIS authentication appropriately.

Next, you must set the NTFS ACLs to allow access only to the proper identities.

**Related Information**

- Overview of ASP.NET Security

# Passport Authentication Provider in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Passport authentication is a centralized authentication service provided by Microsoft that offers a single logon and core profile services for member sites. This benefits the user because it is no longer necessary to log on to access new encrypted resources or sites. If you want your site to be compatible with Passport authentication and authorization, this is the provider you should use. This topic provides some introductory material about Microsoft .NET Passport and the ASP.NET support for it. For more information, see the Passport documentation located at http://www.passport.com. In order to access the documentation, you must get a Passport and register.

Passport is a cookies-based authentication service. A sample transaction conversation using Passport authentication might look similar to the following:

1. A client issues an **HTTP GET** request for a protected resource, such as http://www.contoso.com/default.aspx.
2. The client's cookies are examined for an existing Passport authentication ticket. If the site finds valid credentials, the site authenticates the client. If the request does not include a valid authentication ticket, the server returns status code 302 and redirects the client to the Passport Logon Service. The response includes a URL in the query string that is sent to the Passport logon service to direct the client back to the original site.
3. The client follows the redirect and requests the original resource again, this time with the Passport cookie.
4. The Passport logon server presents the client with a logon form.
5. The client fills out the form and does a **POST** back to the logon server, using Secure Sockets Layer (SSL).
6. The logon server authenticates the user and redirects the client back to the original URL (http://www.contoso.com/default.aspx). The response contains an encrypted Passport cookie in the query string.
7. The client follows the redirect and requests the original protected resource again, this time with the Passport cookie.
8. Back on the originating server, the **PassportAuthenticationModule** detects the presence of the Passport cookie and tests for authentication. If successful, the request is then authenticated.

Subsequent requests for protected resources at the site are authenticated at the originating server using the supplied ticket. Passport also makes provisions for ticket expiration and reusing tickets on other member sites.

Passport uses the Triple DES encryption scheme. When member sites register with Passport, they are granted a site-specific key. The Passport logon server uses this key to encrypt and decrypt the query strings passed between sites.

| **Note** |
| --- |
| To use Passport authentication, you must register your site with the Passport service, accept the license agreement, and install the .NET Passport SDK prior to use. |

The ASP.NET **PassportAuthenticationModule** provides a wrapper around the .NET Passport SDK for ASP.NET applications, and provides Passport authentication services and profile information from an **IIdentity**-derived class called **PassportIdentity**.

As is the case with **WindowsIdentity**, the primary purpose of handling the **PassportAuthentication_OnAuthenticate** event is to attach a custom **IPrincipal** object to the context. A special **IIdentity**-derived class called **PassportIdentity** provides an interface to the Passport profile information and methods to encrypt and decrypt Passport authentication tickets.

**Procedures**

**To implement Passport authentication in an ASP.NET application**

1. Download, install, and configure the .NET Passport SDK from http://www.passport.com/business. You must complete a

registration form to obtain the SDK. Windows Server 2003 includes the .NET Passport SDK and does not require this step.

2. Set up Passport as the authentication mode in the application configuration file as follows.

```
<authentication mode= "Passport"/>
```

3. Using the Passport documentation and the .NET Passport SDK functionality, implement Passport authentication and authorization.

**Related Information**

- Overview of ASP.NET Security

# Forms Authentication Provider in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Forms authentication generally refers to a system in which unauthenticated requests are redirected to an HTML form, using HTTP client-side redirection. Forms authentication is a good choice if your application needs to collect its own user credentials at logon time through HTML forms. The user provides credentials and submits the form. If the application authenticates the request, the system issues a cookie that contains the credentials or a key for reacquiring the identity. Subsequent requests are issued with the cookie in the request headers. The requests are authenticated and authorized by an ASP.NET event handler using whatever validation method the application specifies.

Note that forms authentication is often used for personalization, where content is customized for a known user. In some of these cases, identification is the issue rather than authentication, so it is enough merely to store the user name in a durable cookie and use that cookie to access the user's personalization information.

For detailed information about how to develop ASP.NET applications that use forms authentication, see documentation on the forms authentication provider in the MSDN Library.

**Forms Authentication Across Applications**

ASP.NET supports forms authentication in a distributed environment, either across applications on a single server or in a Web farm. Several configuration attributes in the <forms> element configuration section must match for all applications participating in shared forms authentication. The directives **name**, **protection**, **path**, **validationKey**, **descriptionKey**, and **validation** in the following example must be identical (unless otherwise noted) across the applications.

```
<configuration>

<system.web>

<authentication>

<forms name=".ASPXAUTH"

loginUrl="logon.aspx"

protection="all" <!-- Protection must be identical.-->

timeout="30"

path="/" > <!-- Path must have a compatible scope.-->

</authentication>

<!-- Validation and decryption keys must exactly match and cannot

be set to "AutoGenerate". The validation algorithm must also

be the same. -->

<machineKey>

validationKey=
"C50B3C89CB21F4F1422FF158A5B42D0E8DB8CB5CDA1742572A487D9401E3400267682B202B746511891C1BAF47F8D25C07F6C39A104696DB51F17C529AD3CABE"

decryptionKey= "8A9BE8FD67AF6979E7D20198CFEA50DD3D3799C77AF2B72F"

validation="SHA1"

isolateApplications="false"

</machineKey>

</system.web>

</configuration>
```

The encryption and validation keys and the encryption scheme used for cookie data must be exactly the same. If the settings do not match, cookies cannot be shared.

After a cookie has been issued, expiration of the cookie is tracked based on the **Expires** property value in the cookie itself. This means that if two applications have different **timeout** attributes, the expiration date and time that was set when each cookie was originally issued is retained throughout the lifetime of the cookie. When a cookie is updated, the cookie's original expiration is used to compute the new expiration. The only time the configuration **timeout** value is used is when the cookie is initially created.

**Related Information**

- <forms> Element
- Overview of ASP.NET Security

# Authorization in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The purpose of authorization is to determine whether an identity should be granted the requested type of access to a given resource. There are two fundamental ways to authorize access to a given resource:

- File authorization

  File authorization is performed by the **FileAuthorizationModule**, and is active when you use Windows authentication. It does an access control list (ACL) check of the .aspx or .asmx handler file to determine if a user should have access. Applications can further use impersonation to get resource checks on resources that they are accessing. For more information about impersonation, see Authentication and ASP.NET Impersonation.

- URL authorization

  URL authorization is performed by the **URLAuthorizationModule** class, which maps users and roles to pieces of the URL namespace. This module implements both positive and negative authorization assertions. That is, the module can be used to selectively allow or deny access to arbitrary parts of the URL namespace for certain sets, users, or roles.

The **URLAuthorizationModule** is available for use at any time. You only need to place a list of users and/or roles in the <**allow**> or <**deny**> elements of the <**authorization**> section of a configuration file.

To establish the conditions for access to a particular directory, you must place a configuration file that contains an <**authorization**> section in that directory. The conditions set for that directory also apply to its subdirectories, unless configuration files in a subdirectory override them. The general syntax for this section is as follows.

```
<[element] [users] [roles] [verbs]/>
```

An element is required. Either the *users* or the *roles* attribute must be included. Both can be included, but both are not required. The verbs attribute is optional.

The permissible elements are <**allow**> and <**deny**>, which grant and revoke access, respectively. Each element supports three attributes, which are defined in the following table.

| Attribute | Description |
|---|---|
| roles | Identifies a targeted role for this element. The associated **IPrincipal** object for the request determines the role membership. You can attach arbitrary **IPrincipal** objects to the context for a given request and they can determine role membership in whatever way you like. For example, the default **WindowsPrincipal** class uses Microsoft Windows NT groups to determine role membership. |
| users | Identifies the targeted identities for this element. |
| verbs | Defines the HTTP verbs to which the action applies, such as **GET**, **HEAD**, and **POST**. |

Anonymous users are also denied.

The following example grants access to Kim and members of the Admins role, while denying it to John and all anonymous users:

```
<authorization>
<allow users="Kim"/>
<allow roles="Admins"/>
<deny users="John"/>
<deny users="?"/>
</authorization>
```

Both users and roles can refer to multiple entities by using a comma-separated list, as shown in the following example.

```
<allow users="John, Kim, contoso\Jane"/>
```

Notice that the domain account, `contoso\Jane`, must include both the domain and user name combination.

In addition to identity names, there are two special identities, as shown in the following table.

| Identity | Description |
|---|---|
| * | Refers to all identities |
| ? | Refers to the anonymous identity |

To allow John and deny everyone else, one might construct the following configuration section.

```
<authorization>
<allow users="John"/>
<deny users="*"/>
</authorization>
```

The following example lets everyone do a **GET**, but only Kim can use **POST**.

```
<authorization>
<allow verb="GET" users="*"/>
<allow verb="POST" users="Kim"/>
<deny verb="POST" users="*"/>
</authorization>
```

Rules are applied using the following heuristics:

- Rules contained in configuration files at lower directory levels take precedence over rules at higher directory levels. The system determines which rule takes precedence by constructing a merged list of all rules for a URL, with the most recent (nearest in the hierarchy) rules at the head of the list.

- Given a set of merged rules for a URL, the system starts at the head of the list and checks rules until the first match is found. Note that the default configuration for ASP.NET contains an <**allow users="*"**> element, which authorizes all users. If no rules match, the request is allowed unless otherwise denied. If a match is found and the match is a <**deny**> element, it returns the 401 status code. Applications or sites can easily configure a <**deny users="*"**> element at the top level of their site or application to prevent this behavior.

  If an <**allow**> matches, the module does nothing and lets the request be processed further.

There is also a <**location**> tag that you can use to specify a particular file or directory to which settings wrapped by that tag (between <**location**> and <**/location**> tags) should apply.

**Related Information**

- Overview of ASP.NET Security

- Hierarchical Configuration Architecture in ASP.NET

- Configuration <location> Settings in ASP.NET

- Authentication and ASP.NET Impersonation

# Impersonation in ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When using impersonation, ASP.NET applications can optionally execute with the identity of the client on whose behalf they are operating. The usual reason for doing this is to avoid dealing with authentication and authorization issues in the ASP.NET application code. Instead, you rely on Internet Information Services (IIS) to authenticate the user and either pass an authenticated token to the ASP.NET application or, if unable to authenticate the user, pass an unauthenticated token. In either case, the ASP.NET application impersonates whichever token is received if impersonation is enabled. The ASP.NET application, now impersonating the client, then relies on the settings in the NTFS directories and files to allow it to gain access, or not. Be sure to format the server file space as NTFS, so that access permissions can be set.

Impersonation is disabled by default. For ASP compatibility, the user must explicitly enable impersonation. If impersonation is enabled for a given application, ASP.NET always impersonates the access token that IIS provides to ISAPI extensions. That token can be either an authenticated user token, or the token for the anonymous user (such as IUSR_*ComputerName*). The impersonation occurs regardless of the type of authentication being used in the application.

Only application code is impersonated; compilation and configuration are read as the process token. The result of the compilation is put in the Temporary ASP.NET files directory. The account that is being impersonated needs to have read/write access to this directory. If an application is on a universal naming convention (UNC) share, ASP.NET will always impersonate the token provided to IIS to access that share unless a configured account is used. If an explicit configured account is provided, ASP.NET will use that account in preference to the IIS UNC token. Applications that require per-request impersonation can simply be configured to impersonate the user making the request.

Impersonation is disabled at the computer level by default and, unless overridden, all the application domains inherit this setting. You can enable impersonation by putting a configuration file in the application root directory. For more information about the ASP.NET configuration system, see Overview of ASP.NET Configuration.

As is the case with other configuration directives, this directive applies hierarchically. It is respected by nested applications in the hierarchy, unless explicitly overridden. The default value for this setting is as follows.

```
<impersonation enable="false"/>
```

A minimal configuration file to enable impersonation for an application might look similar to the following example.

```
<!-- Web.config file. -->
<identity impersonate="true"/>
```

There is also name support for running an application as a configurable identity. For example:

```
<identity impersonate="true" userName="contoso\Jane" password="pass"/>
```

This enables the entire application to run as `contoso\Jane`, regardless of the identity of the request, as long as the password is correct. This type of impersonation can be delegated to another computer.

You can programmatically read the identity of the impersonated user, as shown in the following example.

```
[Visual Basic]
Dim username As String = System.Security.Principal.WindowsIdentity.GetCurrent().Name
[C#]
String username = System.Security.Principal.WindowsIdentity.GetCurrent().Name;
```

In the preceding example, **userName** and **password** are stored in clear text in the configuration file. Although IIS will not transmit .config files in response to a user agent request, configuration files can be read by other means, for instance by an authenticated user with proper credentials on the domain that contains the server. To help maintain security, the identity section supports storage of encrypted **userName** and **password** attributes in the registry as shown in the following example.

```
userName="registry:HKLM\Software\AspNetIdentity,Name"
password="registry:HKLM\Software\AspNetIdentity,Password"
```

The portion of the string after the keyword **registry** and up to the first comma indicates the **registry** key that ASP.NET must

open. The portion after the comma contains a single string value name from which ASP.NET will read the credentials. The comma is required and the credentials must be stored in the HKLM hive. If the configuration format is incorrect, ASP.NET will not launch the worker process and the current account creation failure code path will be followed.

The portion of the string after the keyword registry and before the comma indicates the name of the registry key that ASP.NET opens. The portion after the comma contains a single string value name from which ASP.NET will read the credentials. The comma is required, and the credentials must be stored in the HKLM hive. If the configuration format is incorrect, ASP.NET will not launch the worker process and the current account creation failure code path will be followed.

The credentials must be in REG_BINARY format, containing the output of a call to the Windows API function **CryptProtectData**. You can create the encrypted credentials and store them in the registry with the ASP.NET Set Registry console application (Aspnet_setreg.exe), which uses CryptProtectData to accomplish the encryption. To download Aspnet_setreg.exe, along with the Visual C++ source code and Help, visit the Web site www.asp.net and search for **aspnet_setreg**.

You should configure access to the key storing the encrypted credentials so that access is provided only to Administrators and SYSTEM. Because the key will be read by the ASP.NET process running as SYSTEM, you should set the following permissions:

Administrators:F

SYSTEM:F

CREATOR OWNER:F

ProcessAccount:R

This provides two lines of defense to protect the data.

The ACL permissions require the identity accessing the data to be an Administrator. An attacker must run code on the server (**CryptUnprotectData**) to recover the credentials for the account.

**Related Information**

- Overview of ASP.NET Security

# Authentication and ASP.NET Impersonation

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The examples in this section show how to create ASP.NET applications that incorporate various kinds of industry-standard security technologies. For examples of how to create applications that use forms authentication, see information about Forms Authentication Provider on MSDN.

**Using IIS Authentication with ASP.NET Impersonation**

In this scenario, you are setting up an application on an intranet Web site for posting employee information. However, some of the information is for managers only. The manager information can be posted to a subdirectory of the general employee information, so that access to it can be limited. The scenario also assumes that:

- You are using a Microsoft Windows NT or Microsoft Windows 2000 server.
- The hard disk is formatted for NTFS.
- Internet Information Services (IIS) 6.0 is the Web server.
- All employees needing access are using Windows platforms.

You, as administrator:

1. Create the files and directories shown in the following illustration.




2. Create a Windows group called Managers that contains all users who should have access to the ManagerInfo.aspx file.
3. Set up Windows authentication using the IIS administration tool.
4. Set the **impersonate** attribute in the <identity> element in the ASP.NET configuration file to **true**.
5. Set the NTFS access control list (ACL) for the ManagerInformation directory to allow access to only those identities that are in the Windows Manager group. Note that the local system still needs access as well so that the ASP.NET process itself can read the files. Settings in this directory typically resemble the following:
    a. Remove access from the Everyone group, if it has access.
    b. Deny anonymous users.
    c. Add accounts that are to have access privileges.
    d. Give the system account access.

This helps provide the necessary security without the necessity of writing any code.

**Related Information**

- Secure ASP.NET Applications
- Overview of ASP.NET Security

# Security of ASP.NET Applications in Hosted Environments

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP.NET offers an excellent environment for hosting multiple applications on a single Web server. Application Service Providers (ASPs) and Internet Service Providers (ISPs) are able to host multiple applications from different outside sources on a shared server and keep the applications isolated from each other and from critical resources on the server. ASP.NET's configuration system is designed to make it easy to isolate Web applications and to grant code access security permissions to individual applications.

**In This Section**

Configuring Process Identity for ASP.NET: Describes the trust level and policy settings that you can use to help protect the configuration of the server that hosts your ASP.NET applications.

Required Access Control Lists (ACLs) for ASP.NET: Lists the permissions that ASP.NET needs to access specific server resources.

Code Access Security for ASP.NET: Describes the importance of configuring your Web applications so that they will accept code only from trusted sources.

Trust Levels and Policy Files for ASP.NET: Describes the trust level and policy settings that you can use to securely configure the server that hosts your ASP.NET applications.

**Related Sections**

Overview of ASP.NET Security:Describes how to work with authorization and authentication in ASP.NET applications.

Overview of ASP.NET Configuration: Describes how to configure ASP.NET applications and includes reference information about each of the ASP.NET configuration elements.

# Configuring Process Identity for ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The ASP.NET worker process (Aspnet_wp.exe) should run with a different and less privileged identity than the Windows default System identity. This allows the process to run as a relatively untrusted user so that hosted Web applications do not have administrative privileges and cannot compromise the integrity of the server. If the security of the ASP.NET worker process were breached, only a relatively weak account would be compromised and the intruder would not have administrative access to the entire server.

To run the ASP.NET worker process with its own account under Windows 2000 and Windows Server, you can apply the following two attributes to the **<processModel>** configuration element:

- **userName**

  The name of the account under which the process will run.

- **password**

  The clear-text password for the account. There are risks associated with storing clear-text passwords in a configuration file. For applications published on the Internet, you should use alternative means of running the application with a specific identity, or you should store encrypted credentials in the registry. For more information about storing credentials in the registry, see "<processModel>" on MSDN.

The following example shows how to set these attributes in the <**processModel**> section of a configuration file to run the worker process under a local user account.

```
<system.web>
    <processModel enable="true"
                  userName="LOCALMACHINE\IUSR_SomeUser"
                  password="password"/>
</system.web>
```

System and Machine are two predefined **userName** identity settings. The System account runs the worker process with the same identity as Internet Information Services (IIS) itself (typically System). The Machine account runs the worker process with a special unprivileged account named ASPNET. With either identity, the password is **AutoGenerate**, which means that the process does not have to supply credentials to the operating system.

To enable running the worker process with the special ASPNET account, add the following lines to the Machine.config file. If you change the process account in Machine.config, you must reset IIS or stop the Aspnet_wp.exe process.

```
<system.web>
    <processModel enable="true"
                  userName="Machine"
                  password="AutoGenerate"/>
</system.web>
```

With the new process model enabled on Internet Information Services version 6, ASP.NET does not use its own process model. Its identity is inherited from the IIS worker process.

The following example is supported only if you explicitly add the appropriate account rights to the account under which the worker process runs. There are risks associated with storing clear-text passwords in a configuration file. For applications published on the Internet, you should use alternative means of running the application with a specific identity, or you should store encrypted credentials in the registry. For more information, see "<processModel>"on MSDN.

```
<system.web>
    <identity impersonate="true"
              userName="LOCAL\SomeUser"
              password="password"/>
</system.web>
```

The technique described in the following paragraphs is commonly used by administrators hosting ASP sites and it will work in an ASP.NET environment with a weaker process identity.

The application grants access to an anonymous user (this assumes that the authentication scheme allows anonymous users), and the anonymous account is configured to be the desired account for the application. Also, the application uses per-request impersonation and the application code executes as the impersonated account. The process account still compiles and reads configuration because impersonation is used only with the request handler.

If you alternate between running the worker process with an account that has administrator privileges and a weaker account, the weaker account will not be able to read the Local Security Authority (LSA) secret used for Forms authentication, so all Forms authentication users must be reauthenticated. This does not apply when going from an account without administrative privileges to an account with administrative privileges. It does apply to applications with the following configuration.

```
<system.web>
    <authentication mode="Forms"/>
    <machineKey validationKey="AutoGenerate"
                decryptionKey="AutoGenerate"/>
</system.web>
```

**Related Information**

- Overview of ASP.NET Security

# Required Access Control Lists (ACLs) for ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The following table shows which type of access permissions ASP.NET must have in order to access the listed locations.

| Location | Access type | Account | Comments |
|---|---|---|---|
| *installroot*\ASP.NET Temporary Files | Read/Write/Execute | Process or configured impersonation. | This is the location for dynamically compiled files. Beneath this location, application code generation takes place in a discrete directory for each application. The root location itself can be configured using the **tempDir** attribute of the **<compilation>** section. |
| **systemroot**\assembly | Read/Execute | Process or configured impersonation. | This is the location of the global assembly cache. |
| **systemroot**\system32 | Read/Execute | Process | Contains system DLLs loaded by the .NET Framework. |
| **systemroot**\temp | Read/write/delete | Process | Used for XML Web services support. |
| *userprofile* directory | Read/write | Process | Used by the global assembly cache lock files and the common language runtime security configuration caching mechanism. If the user profile directory for the account does not exist, ASP.NET will use the default *userprofile* directory. |
| *application* directory | Read | Process or configured impersonation. | This is the location for application files. |
| *installroot* hierarchy | Read/Execute | Process or configured impersonation | ASP.NET must be able to access the system assemblies referenced in the Machine.config file in the config subdirectory under *installroot*. |
| Default Web Site monitoring | Read | Process | Used for file change notification monitoring. |

**Related Information**

- Overview of ASP.NET Security

# Code Access Security for ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

One of the major benefits of using ASP.NET to host multiple Web sites is the common language runtime's support for code access security to help provide an execution platform for server applications. Code is assigned to a security zone classification based on evidence about the code's origin (evidence such as a strongly named assembly or the URL of origin, for example). Without the ability to configure security for individual, unrelated applications installed on a common server, code in an ASP.NET page belonging to one application would be able to read the files, for example, the Web.config file, in another application. Applications that run with full trust are constrained only by the operating system account under which they are executing.

Code access security can be easily handled for an individual assembly by strongly naming it and adding policy for that assembly. However, many ASP.NET assemblies are dynamically generated as a result of dynamic page compilation and therefore are not strongly named, so an additional facility is required.

ASP.NET lets you assign to an application a configurable level of trust that corresponds to a predefined set of permissions. By default (unless you explicitly alter the configuration), applications receive a level of trust commensurate to the evidence they present. For example, local applications must run in the MyComputer zone with the **Full** trust permission set, and applications located on a Universal Naming Convention (UNC) share must run in the Intranet zone with the **LocalIntranet** restricted permission set. If you want to run a Web application with less than **Full** trust, you must use one of the predefined trust levels defined in ASP.NET Trust Levels and Policy Files to enforce a partial trust policy.

You can use the following configuration settings to override the default behavior and associate an application with a given security policy.

```
<location path="MyApp" allowOverride="false">
<trust level="High"
originUrl="http://www.contoso.com"/>
</location>
```

The <**trust**> configuration tag can apply to the machine level or to any application root directory in the hierarchy.

If you want to set policy for the entire site, you can do so by specifying the root of the site as the path location, as follows.

```
<location path="ContosoSite" allowOverride="false">
<trust level="High"
originUrl="http://www.contoso.com"/>
</location>
```

You should take care to specify the <location> directive using the **allowOverride** attribute if you do not want applications to be able to specify their own trust level (this is the case in most shared server installations).

The following table lists the default supported attributes for the <**trust**> element.

| Attribute | Description | Default supported values |
|---|---|---|
| level | Specifies the security zone under which the application will run. | Full, High, Medium, Low, and Minimal. |
| originUrl | Allows an administrator to set the application's URL of origin.<br>This allows the permissions that rely on the notion of host to function correctly | Well-formed HTTP URLs. |

By default, the policies associated with the various security levels are granted the permissions shown in the following table.

| | **Level** | | | | |
|---|---|---|---|---|---|
| Permission | Full | High | Medium | Low | Minimal |

| AspNetHostingPermission | Full | High | Medium | Low | Minimal |
|---|---|---|---|---|---|
| Environment | Un | Un | Read: TEMP, TMP, OS, USERNAME, COMPUTERNAME | | |
| FileIO | Un | Un | Read, Write, Append, PathDiscovery:Application Directory | Read, Path Discovery: Application Directory | |
| IsolatedStorage | Un | Un | AssemblyIsolationByUser, Unrestricted UserQuota | 1 MB quota (can be varied for individual sites), AssemblyIsolationByUser" | |
| Reflection | Un | ReflectionEmit | | | |
| Registry | Un | | | | |
| Security | Un | Execution, Assertion, ControlPrincipal, ControlThread, RemotingConfiguration | Execution, Assertion, ControlPrincipal, ControlThread, RemotingConfiguration | Execution | Execution |
| Socket | Un | Un | | | |
| WebPermission | Un | Un | Connect to origin host (if configured) | | |
| DNS | Un | Un | Un | | |
| Printing | Un | Default Printing | Default Printing | | |
| OleDBPermission | Un | | | | |
| SqlClientPermission | Un | Un | AllowBlankPassword=false | | |
| EventLog | Un | | | | |
| Message Queue | Un | | | | |
| Service Controller | Un | | | | |
| Performance Counters | Un | | | | |
| Directory Service | Un | | | | |

Un = Unrestricted

Blank = No Permission

When a permission level is available but is not explicitly mentioned in security policy, applications running with the **Full** level of trust can use it (since they have the **Full** trust permission set). Applications running with lower trust levels will not be able to use resources protected by nonexplicit permissions unless you specifically alter policy to allow it.

As the table shows, **High** trust applications have read/write permission for files in their application directories, and **Low** trust applications have read-only permission for files in their application directories. Because **FileIO** permission relies on a physical path (for example, `c:\MyAppPath`), ASP.NET uses a tokenized statement in the policy files that is replaced at run time with

relevant path information for the application.

Note that **WebPermission** allows the application to connect to the origin host. This mechanism works from ASP.NET by providing an optional **OriginUrl** attribute on the <**trust**> section for a given application. The **OriginUrl** attribute is used to replace the **$OriginHost$**variable in policy files, as demonstrated in the following code from Web_hightrust.config.

```
<IPermission class="WebPermission" version="1">
<ConnectAccess>
<URI uri="$OriginHost$"/>
</ConnectAccess>
</IPermission>
```

**SocketPermission** takes a host name or a dotted IP string (possibly with wildcard characters) and **WebPermission** takes a regular expression that includes the protocol (for example, `http://backendmachine/.*`). If you want to change this, you can alter the policy files (or create new ones by copying the defaults) with the desired permissions. For example, you can alter the **SocketPermission** section from the ASP.NET named permission set to grant TCP socket connectivity to "backend1" and "backend2" on port 8080, as follows.

```
<IPermission class="SocketPermission" version="1">
<ConnectAccess>
<ENDPOINT host="backend1" transport="Tcp" port="8080"/>
<ENDPOINT host="backend2" transport="Tcp" port="8080"/>
</ConnectAccess>
</IPermission>
```

**Related Information**

- Overview of ASP.NET Security

# Trust Levels and Policy Files for ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Trust levels are associated with policy files using the <securityPolicy> configuration element, which is only valid in a site level (Web.config) configuration file.

```
<system.web>
<securityPolicy>
<trustLevel name="Full" policyFile="internal"/>
<trustLevel name="High" policyFile="web_hightrust.config"/>
<trustLevel name="Medium" policyFile="web_mediumtrust.config"/>
<trustLevel name="Low" policyFile="web_lowtrust.config"/>
<trustLevel name="Minimal" policyFile="web_minimaltrust.config"/>
</securityPolicy>
</system.web>
```

You should take care to specify a <location> directive using the **allowOverride=False** attribute if you do not want applications to be able to specify their own trust level (as is the case in most shared server scenarios). You can remove or add trust levels by adding a new entry to the configuration section specifying the trust level name and the policy file to be used. The default trust files will be installed in the config directory under the location where the aspnet_isapi.dll file is stored. This is the same location as the Machine.config configuration file and the runtime security policy file.

The **Full** trust Level is a special case. Because it is exactly equivalent to having full trust in the local machine zone, the ASP.NET host does not apply any additional policy to these applications. Therefore, it is mapped to an internal handler and the host does not add additional policy to the application domain for **Full** trust applications.

In order to preserve the default settings, ASP.NET will ship two copies of each default file. One copy will be as described above in the configuration section, the second will have a .default.config extension and will not be consumed directly by the system.

**ASP.NET Policy Mechanics**

The application domain policy for the various trust levels consists of a policy level (for example, an instance of the **System.Security.Policy.PolicyLevel** class that is loaded and applied to the application domain as an application domain policy. The policy tree itself is the same for **High**, **Medium**, **Low**, and **Minimal** trust levels (no policy is applied for **Full**), except for the set of granted permissions in the ASP.NET permission set. The policy tree with the permissions granted to the **High** trust level is shown in the table below.

To find the location of the global assembly cache, check for "`CacheLocation`" in the following registry key: HKLM\Software\Microsoft\Fusion. If this key does not exist, the location defaults to `%windir%\assembly`.

The three code groups that use a **UrlMembershipCondition** object rely on paths that are dynamically computed at run time; ASP.NET loads the policy files and replaces the specially delimited arguments (for example, it replaces `$AppDirUrl$` with the proper path before performing the policy resolution).

After computing the application domain policy for a given application (**Full** trust applications are exempt from this step), the policy is applied to the application domain. Application domain policy is intersected with the various other policy levels (Enterprise, Machine, and User) so that it can only further restrict, not expand, the existing policy.

To see how application domain policy is used to restrict granted permissions, consider the fictitious permission sets and policy levels shown in the following table.

| Level | Name | Permissions |
|---|---|---|
| Enterprise | P1 | {A, B, C, D} |
| Machine | P2 | {A, B, C} |
| User | P3 | {A, B, C} |
| Application Domain | P4 | {A, B} |

The simplified application permission (ignoring permission requests) of these policy sets would be the intersection of the permissions granted at the various levels. Thus, the intersection of P1, P2, P3, and P4 is {A, B}.

**Managing Policy Files**

Policy files can be hand-edited. You can use the .NET Framework Configuration Tool or the Code Access Security Policy Tool to administer security policy for the enterprise, machine, or user levels. For more information about using these tools, see the "Administering Security Policy" topic in .NET Framework Help.

**Related Information**

- Overview of ASP.NET Security

# Overview of ASP.NET Deployment

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Deployment is the process of distributing a completed application or component to be installed on other computers. Microsoft ASP.NET is designed to make Web application deployment easy. To that purpose, it includes the following features:

- When a newer version of a Web.config, .aspx, .asmx, or any other ASP.NET file is copied to an existing Web application, ASP.NET detects that the file has been updated. It then loads a new version of the Web application to handle all new requests, while allowing the original instance of the Web application to finish responding to any current requests. When all requests have been satisfied for the original application, it is automatically removed This feature eliminates downtime while Web applications are being updated.

- All ASP.NET application configuration settings are stored in a Web.config file, an XML-based text file stored in the application root directory, for easy deployment of application settings. If you modify these settings when the application is running, ASP.NET automatically detects the changes, reapplies all the settings, and recompiles the application for you. For more information, see Overview of ASP.NET Configuration.

| **Note** |
| --- |
| This does not apply to Internet Information Services (IIS) or COM configuration settings. ASP.NET application settings also can be made in IIS, such as security or process model configuration. Configuration changes made in IIS are not automatically detected and updated by ASP.NET. For more information, see IIS Metabase. |

- No registration is required for local application components. To include any precompiled managed code in your application, you simply need to place it in the local assembly cache. The local assembly cache is the \Bin directory, which is a subdirectory of the application root directory. In addition, any assembly stored in the \Bin directory can be replaced while the application is running.

- All ASP.NET pages are dynamically compiled when they are first requested. The application then uses the compiled class to satisfy requests until a change is made to the page. This allows you to make changes to a page and replace it in a running application.

To deploy application files and local component files, you can simply copy them to the directory structure on the production server. To do this on a network you can use either the XCOPY command-line tool or Microsoft Windows Explorer. To deploy files over the Internet, you can use the FTP tool of your choice. If you are developing your ASP.NET application using Microsoft Visual Studio .NET, use Visual Studio to deploy any Microsoft FrontPage Server Extensions associated with your application. For more information, see Visual Studio .NET Help.

This section describes how to deploy ASP.NET applications and individual application files to a production server. In addition, it explains how to support multiple versions of ASP.NET on a single server and how to ensure that individual applications on that server reference the correct version of ASP.NET. Also, this section describes how to send a redistributable version of the Microsoft .NET Framework to a production server in an application setup program.

This section includes the following information:

- Overview of Deploying ASP.NET Applications: Describes the options for distributing ASP.NET applications and components from a development or test server to a production environment.

- Side-By-Side Support for ASP.NET: Describes how you can run multiple applications that use different versions of the .NET Framework.

- .NET Framework Redistribution: Describes how you can use the redistributable version of the .NET Framework to support applications on a production server.

**Related Information**

- ASP.NET Administration

# Overview of Deploying ASP.NET Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you are ready to deploy your ASP.NET Web application to a production server, you have many options. Your choice depends on the complexity of your application, whether it uses local-only or shared assemblies, and what tools you used to develop the application.

For ASP.NET applications that do not require modifications to IIS settings or registration of COM objects, you can simply copy the application's files to the directory structure on the production server, using one of the following tools:

- The XCOPY command-line tool can copy files from one computer to another on an intranet or internal network.
- Windows Explorer can copy files from one computer to another on an intranet or internal network.
- You can use the FTP tool of your choice to copy application files from one computer to another over the Internet.
- If your application has been developed using Visual Studio .NET, you can use the **Copy Project** command to copy files from one computer to another on an intranet or internal network. For more information, see Visual Studio .NET Help .

| **Note** |
| --- |
| Front Page Server Extensions must be installed on the target server to use the **Copy Project** command. |

These choices have limitations, however. None performs any IIS configuration that your application might require, and if your application uses COM objects, they will not be registered on the target server. Also, if your application uses shared assemblies stored in the server's global assembly cache, you cannot use any of these methods. For more information, see Deploying Shared Components to the Global Assembly Cache and the "Global Assembly Cache" topic in .NET Framework Help.

| **Note** |
| --- |
| You should deploy components that are used only by your application to the \Bin directory. |

If your application is complex -- for example, if it includes shared assemblies or relies on specific IIS settings -- deploying your application to production servers using Windows Installer 2.0 is recommended. For more information, see the "Deploying Applications and Components" topic in Visual Studio .NET Help.

**In This Section**

Creating an ASP.NET Application Root Directory with IIS

Describes how to use IIS to create a virtual directory or to mark a directory in the C:\Inetpub\Wwwroot directory as an application.

Creating a Directory Structure for an ASP.NET Application

Describes how to create a directory structure for your ASP.NET applications.

Deploying ASP.NET Application Files and Local Components

Describes how to distribute all files associated with your ASP.NET application to a production server.

Deploying Local Application Components

Describes how to compile and distribute component assemblies to your ASP.NET application's \Bin directory.

Deploying Shared Components to the Global Assembly Cache

Describes the global assembly cache and how to distribute shared component assemblies to it.

Deploying HTTP Handlers and HTTP Modules

Describes how to distribute compiled assemblies that contain classes that implement the **IHttpHandler** and **IHttpModule** interfaces. This includes required configuration settings for including new HTTP handlers and HTTP modules in your application.

# Creating an ASP.NET Application Root Directory with IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

All ASP.NET applications use IIS virtual directories as their application root directories. The procedures in this topic show you how to create a virtual directory in IIS.

**Procedures**

**To prepare for this procedure**

1. Create a new directory. This walkthrough uses the directory C:\myWebApp.
2. Next, open IIS Manager and create a virtual directory using one of the following procedures.

**To create a virtual directory on a computer running Windows Server 2003**

1. From the **Start** menu, point to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. Expand the **local computer**, and then expand **Web Sites**.
3. Right click **Default Web Site**, point to **New**, and then click **Virtual Directory**.
4. In the Virtual Directory Creation wizard, click **Next**.
5. Enter the name of your application in the **Alias** box, and then click **Next**.
6. Enter the path to the physical directory that you created to prepare for this walkthrough, C:\myWebApp, and then click **Next**.
7. In the **Access Permissions** page, select the access permissions that you want for your application, and then click **Next**.
8. Click **Finish**.

**Related Information**

- Overview of Deploying ASP.NET Applications
- Creating a Directory Structure for an ASP.NET Application
- Deploying Local Application Components
- Deploying ASP.NET Application Files and Local Components
- Deploying Shared Components to the Global Assembly Cache
- Deploying HTTP Handlers and HTTP Modules

# Creating a Directory Structure for an ASP.NET Application

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

How you organize the directories that your applications contain affects any virtual or relative paths that are included in the files that your application contains. In addition, if your application uses local compiled code, you must include a \Bin directory directly beneath your application root. ASP.NET automatically resolves any references to components stored there.

The procedures in this topic explain which directories to include in the root directory of your ASP.NET Web application. A short discussion of possible directory structures follows, and a sample directory structure with potential file names is included at the end of the topic.

**Procedures**

**To prepare for this walkthrough**

- Create a root directory for your application and mark it as an application, as described in Creating an ASP.NET Application Root Directory with IIS.

**To create a directory structure for your ASP.NET Web application**

1. At the command prompt or in Windows Explorer, open the root directory you created.
2. Create a new directory and name it **Bin**.
3. Create any other directories that your application requires. For example, you could create a directory to contain all the images that your application contains.

The \Bin directory is the default assembly cache for your Web application. Unless your Web application shares DLLs with another application, place all DLLs that your application uses in the \Bin directory to ensure that references to the objects contained in them resolve correctly. These assemblies can include business logic or compiled custom server controls. If your application shares assemblies with other applications, place the shared DLLs in the global assembly cache. For more information, see Deploying Shared Components to the Global Assembly Cache.

| **Note** |
| --- |
| These instructions do not apply to assemblies compiled by Visual Studio .NET when you create an ASP.NET Web application in that environment. For information about how to deploy Web applications developed using Visual Studio .NET, see Visual Studio Help. |

If you want to modify the default configuration for your Web application, create a Web.config file, make the modifications you want, and save it in the root directory. For more information, see Overview of ASP.NET Configuration.

In addition, if you write any application-wide functions, including handlers for Application_Error events, place those in a Global.asax file and save the file in the application root directory. For more information, see the "ASP.NET Applications" topic in .NET Framework Help .

In the following sample directory structure, C:\myWebApp is the application root directory. For the application to function properly, the myWebApp directory must be an IIS virtual directory.

| **Note** |
| --- |
| This simple application includes the \Bin directory and an |Img directory, which stores the image associated with the application. Your directory structure might be more complex than this. |

```
c:\myWebApp

Default.aspx

Menu.ascx

Service.asmx

Web.config
```

```
Global.asax

\Bin

myBizLogic.dll

\Img

myImage.gif
```

**Related Information**

- Deploying ASP.NET Applications in IIS 6.0
- Creating a Directory Structure for an ASP.NET Application
- Deploying Local Application Components
- Deploying ASP.NET Application Files and Local Components
- Deploying Shared Components to the Global Assembly Cache
- Deploying HTTP Handlers and HTTP Modules

# Deploying Local Application Components

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP.NET application components are all of the DLL files that the application uses to perform a variety of tasks. They can include business logic and database access logic for three-tier applications, as well as custom ASP.NET server controls that you have developed. To work for an application, each of the DLLs associated with these components must be deployed to your application's \Bin directory.

**Procedures**

**To deploy ASP.NET application components from the command line**

1. From the **Start** menu, click **Run**.
2. In the **Open** box in the **Run** dialog box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following command:

   **xcopy** < *source path* > < *destination path* >

| **Note** |
| --- |
| In this command, <*source path*> is a placeholder that represents the full path to the compiled component files that you want to deploy, including the drive and directory names. You can omit the file names if you want all files in the directory to be copied. The <*destination path*> placeholder represents the full path to the directory where the files should be placed. |

The following example command copies all compiled components from the `C:\devApp\bin` directory to the `D:\liveApp\bin` directory.

```
xcopy c:\devApp\bin d:\liveApp\bin
```

1. Answer any questions about the files or directory you are copying.

The following example copies a single DLL from a \Bin directory on one drive to a \Bin directory on another.

```
xcopy c:\devApp\bin\myAssembly.dll d:\liveApp\bin\
```

**To compile and deploy an assembly to an application's \Bin directory from the command line**

1. From the **Start** menu, click **Run**.
2. In the **Open** box in the **Run** dialog box, type **cmd**, and then click **OK**.
3. Use the **cd** command to navigate to the directory that you want. This command uses relative paths to navigate through the directories and drives on a computer or network.

   The root location of any path that this command uses is the local drive, often C:\. The examples shown in this procedure assume that you have navigated to the directory where the source files are located. The following use of the **cd** command moves the command prompt from a directory two directories beneath the local drive root to a myApp directory directly beneath the root.

   ```
   cd..\..\myApp
   ```

4. At the command prompt, type the following command:

   `csc /t:library /out:` < *destination path* > < *source path* >

   `vbc /t:library /out:` < *destination path* > < *source path* >

| **Note** |
| --- |

In this command, <*destination path*> is a placeholder that represents the path to the .dll file that is compiled by this command, and <*source path*> represents the path to the source file. Each <*destination path*> should include \Bin and the full name you assign to the .dll file. Both <*destination path*> and <*source path*> are relative paths.

The following example compiles `HelloObj.dll` from the command line from a `HelloObj.cs` or `HelloObj.vb` source file. The destination path is a relative path from the directory where the command line currently is to the \Bin directory of the ASP.NET application where the .dll is used. The source file is in the same directory as the command line.

```
csc /t:library /out:..\..\..\..\bin\HelloObj.dll HelloObj.cs
```

```
vbc /t:library /out:..\..\..\..\bin\HelloObjVB.dll HelloObj.vb
```

**Related Information**

- Overview of Deploying ASP.NET Applications
- Creating a Directory Structure for an ASP.NET Application
- Deploying Local Application Components
- Deploying ASP.NET Application Files and Local Components
- Deploying Shared Components to the Global Assembly Cache
- Deploying HTTP Handlers and HTTP Modules

# Deploying ASP.NET Application Files and Local Components

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To deploy an ASP.NET application, simply copy the application files you have created from your development computer to the production Web server that will host your application. You can use the XCOPY command-line tool or any FTP application to copy the files from one location to another.

To deploy assemblies that you want to share across Web applications, such as assemblies that contain custom ASP.NET server controls, you must deploy them to the global assembly cache. For more information, see the "Global Assembly Cache" topic in .NET Framework Help.

**Procedures**

**To deploy ASP.NET application files from the command line**

1. From the **Start** menu, click **Run**.
2. In the **Open** box in the **Run** dialog box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following command:

   **xcopy** < *source path* > < *destination path* >

| **Note** |
| --- |
| In this command, <*source path*> is a placeholder that represents the full path to the source files you want to copy, including the drive, the directory, and the file names. You can omit the file names if you want all files in the directory to be copied. The <*destination path*> placeholder represents the full path to the directory where the copied files should be placed. |

The following example command copies all files from the `C:\myWebApp` directory to the `D:\liveApp` directory.

```
xcopy c:\myWebApp d:\liveApp
```

1. Answer any questions about the files or directory you are copying.

You can use this procedure when you want to update a DLL stored in your application's \Bin directory, or if you want to update any of your other application files. The following example copies a single DLL from a \Bin directory on one drive to a \Bin directory on another.

```
xcopy c:\devApp\bin\myAssembly.dll d:\liveApp\bin\
```

You can also use this command to update files in your application once you have deployed it. Although you can copy an entire directory at once, individual files must be copied one at a time. You can exclude subdirectories, files with a specific file name extension, or specific file names from being copied by using the XCOPY **/exclude** option. For more information about how to use the XCOPY tool, open Help for your operating system and search for XCOPY.

| **Note** |
| --- |
| You must use physical directory names with the XCOPY tool. You cannot use virtual directory names. |

**Related Information**

- Overview of Deploying ASP.NET Applications
- Creating a Directory Structure for an ASP.NET Application
- Deploying Local Application Components
- Deploying Shared Components to the Global Assembly Cache
- Deploying HTTP Handlers and HTTP Modules

# Deploying Shared Components to the Global Assembly Cache

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you plan to share a component between two or more applications on your Web server, you must deploy the component to the server's computer-wide code cache, called the global assembly cache. Each computer on which the common language runtime is installed has a global assembly cache. The global assembly cache stores assemblies that are shared by several applications on the computer. For more information, see the "Global Assembly Cache" topic in .NET Framework Help.

You should share assemblies by installing them into the global assembly cache only when you need to. As a general guideline, keep assembly dependencies private and locate assemblies in the application directory unless sharing an assembly is explicitly required. Assemblies do not have to be in the global assembly cache to be accessible to COM interop or unmanaged code.

| **Note** |
| --- |
| There are scenarios in which you should not install an assembly into the global assembly cache. If you place one of the assemblies that make up an application in the global assembly cache, you can no longer replicate or install the application by using the XCOPY tool to copy the application directory. You must move the assembly in the global assembly cache as well. |

There are two ways to install an assembly into the global assembly cache:

- Using Microsoft Windows Installer 2.0.

  This is the recommended and most common way to add assemblies to the global assembly cache. The installer provides reference counting of assemblies in the global assembly cache, plus other benefits. The only time you should not use the installer is when you are installing to a development or test server.

- Using the Global Assembly Cache tool (Gacutil.exe).

  You can use Gacutil.exe to add strong-named assemblies to the global assembly cache and to view the contents of the global assembly cache.

| **Note** |
| --- |
| Gacutil.exe is only for development purposes and should not be used to install production assemblies into the global assembly cache. |

For more information, see the "Global Assembly Cache Tool (Gacutil.exe)" topic in .NET Framework Help.

**Procedures**

**To install a strong-named assembly into the global assembly cache**

1. From the **Start** menu, click **Run**, type **cmd**, and then click **OK**.
2. At the command prompt, type the following command:

   **gacutil -i** < *assemblyname* >

| **Note** |
| --- |
| In this command, *assembly name* is the name of the assembly to install in the global assembly cache. |

The following example installs an assembly with the file name Hello.dll into the global assembly cache.

```
gacutil -i Hello.dll
```

The .NET Framework SDK also provides a Windows shell extension called the Assembly Cache Viewer (Shfusion.dll), which you can use to drag and drop assemblies into the global assembly cache.

**Related Information**

- Overview of Deploying ASP.NET Applications
- Creating a Directory Structure for an ASP.NET Application
- Deploying Local Application Components
- Deploying ASP.NET Application Files and Local Components
- Deploying Shared Components to the Global Assembly Cache
- Deploying HTTP Handlers and HTTP Modules

# Deploying HTTP Handlers and HTTP Modules

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP.NET provides **IHttpHandler** and **IHttpModule** interfaces that allow you to use APIs that are as powerful as the Internet Server API (ISAPI) programming interfaces available with IIS but with a simpler programming model. HTTP handler objects are functionally similar to IIS ISAPI extensions, and HTTP module objects are functionally similar to IIS ISAPI filters.

ASP.NET maps HTTP requests to HTTP handlers. Each HTTP handler enables processing of individual HTTP URLs or groups of URL extensions within an application. HTTP handlers have the same functionality as ISAPI extensions with a much simpler programming model.

An HTTP module is an assembly that handles events. ASP.NET includes a set of HTTP modules that can be used by your application. For example, the **SessionStateModule** is provided by ASP.NET to supply session-state services to an application. You can also create custom HTTP modules to respond to either ASP.NET events or user events.

**Procedures**

**To compile, deploy, and register HTTP handlers**

1. Compile and deploy the .NET class for the HTTP handler in the \Bin directory under the application's virtual root. For more information, see Deploying Local Application Components.

2. Register the HTTP handler, synchronous or asynchronous, in the application's Web.config configuration file. The following example maps all HTTP requests to the classes **MyHandler.New** and **MyHandler.Fin** in the assembly **MyHandler** that is in the file **MyHandler.dll**.

```
<configuration>

<system.web>

<httpHandlers>

<add verb="*" path="MyHandler.New"

type="MyHandler.New, MyHandlerAssembly" />

<add verb="*" path="*.myNewFileExtension"

type="MyHandler.Fin, MyHandlerAssembly" />

</httpHandlers>

<system.web>

</configuration>
```

3. The <**add**> elements within the <**httpHandlers**> configuration section have three attributes: **path**, **type**, and **verb**. For more information, see <add> Element for <httpHandlers>.

4. Register the HTTP handler extension as a script map in IIS. For more information, see Setting Application Mappings in IIS 6.0.

**To compile, register, and deploy HTTP modules**

1. Compile and deploy the .NET class for the HTTP module in the \Bin directory under the application's virtual root. For more information, see Deploying Local Application Components.

2. Register the HTTP module in the application's Web.config configuration file. The following example registers an **OutputCache** module to an application.

```
<configuration>

<system.web>

<httpModules>
```

```
<add name="OutputCache"

type="System.Web.Caching.OutputCacheModule, System.Web,

Version=1.0.2800.0, Culture=neutral,

PublicKeyToken=b03f5f7f11d50a3a"/>

</httpModules>

<system.web>

</configuration>
```

For more information about working with and registering HTTP handlers and HTTP modules, see information about HTTP handlers or HTTP modules in the MSDN Library.

**Related Information**

- Overview of Deploying ASP.NET Applications
- Creating a Directory Structure for an ASP.NET Application
- Deploying Local Application Components
- Deploying ASP.NET Application Files and Local Components
- Deploying Shared Components to the Global Assembly Cache
- Deploying HTTP Handlers and HTTP Modules

# Side-By-Side Support for ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Traditionally, when a component or application is updated on a computer, the older version is removed and replaced with the newer version. If the new version is not compatible with the previous version, this usually breaks other applications that use the component or application. The .NET Framework provides support for side-by-side execution, which allows multiple versions of an assembly or application to be installed on the same computer, at the same time. Since multiple versions can be installed simultaneously, managed applications can select which version to use without affecting other applications that use a different version.

In ASP.NET, applications are said to be running side by side when they are installed on the same computer but use different versions of the .NET Framework. The following topics describe side-by-side support in ASP.NET and how to configure ASP.NET applications for side-by-side execution.

**In This Section**

Overview of ASP.NET Side-By-Side Support

Describes side-by-side support in ASP.NET and how it relates to side-by-side support in the .NET Framework.

Managing Multiple ASP.NET Versions

Describes how to upgrade, uninstall, and redistribute multiple versions of ASP.NET.

Configuring an ASP.NET Application for an ASP.NET Version

Describes how to configure an ASP.NET application to use a specific version of ASP.NET when multiple versions of the .NET Framework are installed on a computer.

Assembly Unification in the .NET Framework

Describes how an ASP.NET application determines which version of the .NET Framework to use if the application and the components it uses were built on different versions of the runtime.

**Related Sections**

Overview of ASP.NET Deployment

Describes how to distribute ASP.NET applications and components from development and test servers to production servers.

# Overview of ASP.NET Side-By-Side Support

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The .NET Framework allows you to install multiple versions of the runtime on the same computer. By default, when the .NET Framework is installed on a computer with an existing installation, all ASP.NET applications are automatically updated to use this version of the .NET Framework. The only exceptions are applications that are bound to an incompatible version of the runtime or to a later version of the runtime. Although later versions of the .NET Framework are designed to be backwards compatible, you might want to configure an ASP.NET application to use an earlier version.

**Script Maps for ASP.NET Applications**

When multiple versions of the .NET Framework are installed on the same computer, each installation contains an associated version of the ASP.NET ISAPI. An ASP.NET application uses the ASP.NET ISAPI to determine which version of the .NET Framework to use for the application. An ASP.NET application can be configured to use any of the installed ASP.NET ISAPI versions. To specify the ASP.NET ISAPI version to use for an ASP.NET application, a script map is registered in Internet Information Services (IIS) for the application.

A script map associates a file name extension and HTTP verb with the appropriate ISAPI for script handling. For example, when IIS receives a request for an .aspx file, the script map for the application directs IIS to forward the requested file to the appropriate version of the ASP.NET ISAPI for processing. The script map for each ASP.NET application is normally set in the Internet Information Services management console and can be applied directly to an application, or inherited from a parent application. By default, when the .NET Framework is installed, the script maps for all existing ASP.NET applications on the computer are automatically updated to use the ASP.NET ISAPI version associated with the installation, unless the application uses a later or an incompatible version.

To make it easier to reconfigure the script map for an ASP.NET application, each installation of the .NET Framework comes with an associated version of the ASP.NET IIS Registration tool (Aspnet_regiis.exe). By default, this tool is installed in the following directory:

**%systemroot%**\Microsoft.NET\Framework*versionNumber*

Administrators can use this tool to remap an ASP.NET application to the ASP.NET ISAPI version associated with the tool.

| **Note** |
| --- |
| Because Aspnet_regiis.exe is associated with a specific version of the .NET Framework, administrators must use the appropriate version of Aspnet_regiis.exe to reconfigure the script map for an ASP.NET application. Aspnet_regiis.exe only reconfigures the script map of an ASP.NET application to the ASP.NET ISAPI version associated with the tool. |

The tool can also be used to display the status of all installed versions of ASP.NET, register the associated version of ASP.NET, create client-script directories, and perform other configuration operations.

For more information about script maps and IIS configuration, see IIS 6.0 Operations Guide. For more information about updating script maps for an ASP.NET application, see ASP.NET IIS Registration Tool (Aspnet_regiis.exe).

**ASP.NET Performance Counters**

Performance counters allow you to monitor the operation of ASP.NET from the Performance console. Each installation of ASP.NET installs performance counters that apply to each individual version of ASP.NET, as well as to all installed versions.

Each version of ASP.NET installs a separate set of performance counter objects to monitor the ASP.NET system and ASP.NET applications. These objects use the naming conventions "ASP.NET v.X.X.XXXX.X" and "ASP.NET Apps v.X.X.XXXX.X", respectively, where v.X.X.XXXX.X denotes the ASP.NET version and build numbers.

In addition, ASP.NET also installs ASP.NET system and ASP.NET applications performance counters that apply to all installed versions of ASP.NET. Since these counters are shared between all versions of ASP.NET, they always correspond to the latest installed version on the machine. If the latest version of ASP.NET is uninstalled, these counters are automatically re-associated with the latest remaining version of ASP.NET.

**ASP.NET State Service**

ASP.NET state service is used to manage session state on a computer. All compatible versions of ASP.NET that are installed on the same computer share the same state service. The state service version that is used is always the one that is installed with the latest version of ASP.NET. This version is registered under aspnet_state. Therefore, when this version of ASP.NET is

uninstalled, the latest remaining version on the computer is then registered and used in its place.

**ASP.NET SQL State Server**
ASP.NET SQL state server is used to manage session state in SQL on a computer. All compatible versions of ASP.NET that are installed on the same computer share the same SQL state server. The SQL state server version that is used is always the one that is installed with the latest version of ASP.NET. When this version of ASP.NET is uninstalled, the latest remaining version on the computer is then registered and used in its place.

**ASPNET User Account and Filter**
All installed versions of ASP.NET share the same ASPNET user account (for the worker process) and Aspnet_filter.dll (for storing and removing the session ID from cookie-less session-state requests). The ASPNET user account and Aspnet_filter.dll always corresponds to the latest installed version of ASP.NET. When this version of ASP.NET is uninstalled from a computer, the latest remaining version on the computer is then registered and used in its place.

**Related Information**

- Side-by-Side Support in ASP.NET
- Managing Multiple ASP.NET Versions
- Configuring an ASP.NET Application for an ASP.NET Version

# Managing Multiple ASP.NET Versions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The .NET Framework allows you to install multiple versions of the runtime on the same computer. This allows an ASP.NET application to use the version it needs to run without affecting other applications. The following sections describe how to install, uninstall, and redistribute multiple versions of ASP.NET on the same computer.

**Installing Multiple Versions of ASP.NET**

To add a version of ASP.NET to a computer, run the .NET Framework redistributable package setup program (Dotnetfx.exe) for the version you want to install. By default, all existing ASP.NET applications are automatically reconfigured during installation to use the installing version of the runtime if the following conditions are met:

- The installing .NET Framework version is later than the version that is currently mapped to the application.
- The installing .NET Framework version is compatible with the version mapped to the application. In general, the version number determines compatibility. Minor revision and build number changes are usually compatible, whereas major revision numbers are not.

| **Caution** |
| --- |
| Although ASP.NET applications are automatically updated to use the installing version of ASP.NET if the preceding conditions are met, custom configuration settings in the current Machine.config file are not transferred to the installing Machine.config file. If your application uses custom configuration settings, be sure to either manually update the new Machine.config file or use the ASP.NET IIS Registration tool (Aspnet_regiis.exe) to remap the application to the previous version of ASP.NET. For more information about remapping an application, see Configuring an ASP.NET Application for an ASP.NET Version. |

To prevent the automatic remapping of all existing applications to the installing version of the .NET Framework, use the **/noaspupgrade** command-line option with the Dotnetfx.exe setup program. The following shows the entire command line:

**Dotnetfx.exe /c:"install /noaspupgrade"**

For more information about the .NET Framework redistributable package, see Redistributing the .NET Framework 1.1 on MSDN.

**Web Service Extensions in IIS 6.0**

If you're using Internet Information Services 6.0, you might need to enable ASP.NET in the Internet Information Services management console after installing ASP.NET. IIS 6.0 introduces a new Web Service Extensions folder, which allows an administrator to selectively enable and disable functionality in IIS. When ASP.NET version 1.1 or later is installed on a computer that uses IIS 6.0, the installation script creates a metabase key for that version of ASP.NET. This metabase key allows ASP.NET to appear in the console as a configurable component. A separate key is created for each installed version of ASP.NET, allowing each version to be separately configurable in the console. When a version of ASP.NET is uninstalled from a computer, the respective key is also removed.

| **Note** |
| --- |
| When ASP.NET is installed through the Add or Remove Programs item in Control Panel, or the Configure Your Server wizard, ASP.NET is enabled in the Web Service Extensions by default. However, if it is installed through Microsoft Visual Studio .NET, Windows Update, or the .NET Framework redistributable package (Dotnetfx.exe), ASP.NET is disabled by default. |

**To manage items in the Web Service Extensions folder**

1. Open IIS Manager, expand **local computer**, and then click the **Web Service Extensions** folder.
2. Click the **Extended** tab, and then click the Web Service Extension item you want to configure.
3. Click the **Allow** or **Prohibit** button to enable or disable the item.
4. Click the **Add a new Web service extension** link to add a new item to the console.

**Uninstalling a Version of ASP.NET**

If you have multiple versions of ASP.NET installed on a machine, you can uninstall a version on an individual basis. Uninstalling ASP.NET is usually accomplished by uninstalling the .NET Framework through the **Add or Remove Programs** Control Panel item. During the uninstall process, the setup program calls the ASP.NET IIS Registration tool (Aspnet_regiis.exe) with the **-u** option (uninstall flag).

| **Note** |
| --- |
| If you just want to uninstall a version of ASP.NET without uninstalling the associated .NET Framework, you can use Aspnet_regiis.exe directly. Each installation of the .NET Framework contains an associated version of the tool. To uninstall only ASP.NET, use the **-u** option of the tool version associated with the ASP.NET version you want to uninstall. For more information, see the ASP.NET IIS Registration Tool. |

The **-u** option of Aspnet_regiis.exe performs the following actions:

- All ASP.NET applications currently mapped to the ASP.NET version being uninstalled are remapped to the next-latest compatible version on the computer. If no compatible version exists, the script map for the application is removed completely.

  | **Caution** |
  | --- |
  | Removing the script maps will cause IIS to serve ASP.NET pages as literal text. This could potentially expose source code to the client. |

- The ASP.NET state service is shared across all compatible versions of ASP.NET and is always mapped to the latest installed version of ASP.NET. If the ASP.NET version you are uninstalling is currently the latest version, the next-highest compatible version on the computer is then registered to replace it. If no compatible version exists, the ASP.NET State Service is uninstalled.

- Performance objects and their associated performance counters specific to the version being uninstalled are removed. The generic performance objects and counters that apply to all installed versions of ASP.NET, ASP.NET and ASP.NET Applications, are directed to the latest installed version remaining on the computer.

- The ASPNET user account is shared across all compatible versions of ASP.NET. If no additional ASP.NET installations remain on the computer, the ASPNET user account and the associated access control lists are removed.

For information about other effects of installing multiple versions of ASP.NET, see "Overview of ASP.NET Side-By-Side Support" on MSDN.

**Third-Party Redistribution of the .NET Framework**

Third-party applications can package and install the .NET Framework as part of their setup under a special redistribution license. This allows an application to install the specific version of the .NET Framework it needs to run. The .NET Framework redistributable package (Dotnetfx.exe) is provided for an application to deploy with the runtime. The default installation script of Dotnetfx.exe might not suit the needs of your application. It is up to the third-party application developer to use the options described below when installing the .NET Framework as part of an application's setup.

For more information about the .NET Framework redistributable package, see "Redistributing the .NET Framework 1.1" on MSDN. For more information about the EULA of the redistributable package, see the Where to Obtain the **.NET Framework Redistributable Package** topic in .NET Framework Help.

**Disabling Automatic Script Map Updates**

By default, when the .NET Framework is installed on a computer with an existing installation, all ASP.NET applications are automatically updated to use this version of the .NET Framework, unless the applications are bound to an incompatible version of the runtime or to a later version of the runtime. Usually, only a system administrator performs this type of installation. When the .NET Framework is installed as part of an application's setup, it should not force other applications on the computer to use that version of the runtime.

As part of the setup script, Dotnetfx.exe normally calls the ASP.NET IIS Registration tool (Aspnet_regiis.exe) with the **-i** option to install ASP.NET. This option attempts to update the script maps of all existing ASP.NET applications. To prevent the installation from automatically updating other applications on the computer, the **/noaspnetupgrade** option is provided as part of Dotnetfx.exe. This option causes the setup program to use the **-ir** option of Aspnet_regiis.exe instead, which skips the automatic

script map updates, default documentation, and MIME swapping steps. The following shows the entire command line.

**Dotnetfx.exe /c:"install /noaspupgrade"**

**Manually Updating Script Maps**

Sometimes, it is necessary to manually set up the script maps for an application. If the application needs its own virtual directory that targets the installing version of the .NET Framework, make sure the application's setup program also explicitly calls Aspnet_regiis.exe with the **-s** or **-sn** option to set up the script maps.

The following shows a sample command line that updates the script maps for an application called `SampleApp1`.

```
Aspnet_regiis.exe -s W3SVC/1/ROOT/SampleApp1.
```

For more information, see the ASP.NET IIS Registration Tool.

# Configuring an ASP.NET Application for an ASP.NET Version

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, when the .NET Framework is installed on a computer with an existing installation, all ASP.NET applications are automatically updated to use this version of the .NET Framework. The only exceptions are applications that are bound to an incompatible version of the runtime or to a later version of the runtime. Although later versions of the .NET Framework are designed to be backward compatible, you might want to configure an ASP.NET application to use an earlier version. The following sections describe the process to configure an ASP.NET application to target a specific version of the .NET Framework.

**Viewing the Script Map for an ASP.NET Application**

When managing a computer with multiple versions of the .NET Framework installed, it is often useful to view the script map for an ASP.NET application to determine which version is used by the application. You can view the script map for an ASP.NET application using the Internet Information Services management console.

### To view the script map for an ASP.NET application

1. Open IIS Manager, expand the local computer by clicking the plus sign, and navigate to the folder that contains the ASP.NET application.
2. Right-click the folder and click **Properties**. The application's **Properties** dialog box appears.
3. On the **Directory** tab, click the **Configuration** button. The **Application Configuration** dialog box appears.
4. On the **Mappings** tab, select an ASP.NET application extension, such as .asmx or .aspx.

The Executable Path column of the dialog box will list the path to the ASP.NET ISAPI version used by the application. By default, the ASP.NET ISAPI is installed in the following location:

*systemroot*\Microsoft.NET\Framework*versionNumber*

The version number shown in the path indicates the version number of the ASP.NET ISAPI used by the application. The ASP.NET ISAPI version determines which version of the runtime is used by the application.

**Updating Script Maps Using Aspnet_regiis.exe**

To make it easier to reconfigure the script map for an ASP.NET application, each installation of the .NET Framework comes with an associated version of the ASP.NET IIS Registration tool (Aspnet_regiis.exe). Administrators can use this tool to remap an ASP.NET application to the ASP.NET ISAPI version associated with the tool.

| **Note** |
| --- |
| Because Aspnet_regiis.exe is linked to a specific version of the .NET Framework, administrators must use the appropriate version of Aspnet_regiis.exe to reconfigure the script map for an ASP.NET application. Aspnet_regiis.exe only reconfigures the script map of an ASP.NET application to the ASP.NET ISAPI version associated with the tool. |

The tool can also be used to display the status of all installed versions of ASP.NET, register the associated version of ASP.NET, create client-script directories, and perform other configuration operations.

### To use Aspnet_regiis.exe to update a script map for an ASP.NET application

1. Open a command window. (Click **Start**, click **Run**, type **cmd**, and then click **OK**.)
2. Navigate to the directory of the Aspnet_regiis.exe version you want to use. Remember that each version of the .NET Framework comes with its own version. The file is usually located in the following directory:

   *systemroot*\Microsoft.NET\Framework*versionNumber*

3. Use the **-s** or **-sn** option of Aspnet_regiis.exe along with the path to the application to setup the script maps. The following shows a sample command line that updates the script maps for an application called `SampleApp1`.

   ```
   Aspnet_regiis.exe -s W3SVC/1/ROOT/SampleApp1
   ```

**Application Pooling in IIS 6.0**

IIS versions earlier than 6.0 use the ASP.NET process model (Aspnet_wp.exe). Under the ASP.NET process model, each unique application version automatically runs in a separate process at run time. All applications that target the same version of the runtime share the same process (or processes in Web garden mode). However, IIS 6.0 uses the IIS 6.0 process model (W3wp.exe) and introduces a new isolation feature called application pooling. Application pooling allows applications to run together in one or more processes, as long as they share the same pool designation. Applications that are assigned different application pools never run in the same process.

When multiple versions of the .NET Framework are installed on a computer that uses IIS 6.0, you might encounter the following error message in the Application Event log.

It is not possible to run different versions of ASP.NET in the same IIS process. Please use the IIS Administration Tool to reconfigure your server to run the application in a separate process.

This error occurs when more than one version of ASP.NET is configured to run in the same process. Different versions of the .NET Framework and runtime cannot coexist side-by-side in the same process. Therefore, an ASP.NET application that uses a particular version of the runtime must not share a process with an application that uses a different version. This error commonly occurs when two or more applications are mapped to different versions of ASP.NET, but share the same application pool.

To use application pooling, first make sure IIS 5.0 isolation mode is disabled. Next, navigate to the desired application in the Internet Information Services management console and assign it an application pool. If you want to isolate your application from the existing application pools, create a new pool designation.

| **Note** |
| --- |
| The process isolation for each application is configured manually. |

**To disable IIS 5.0 isolation mode in IIS 6.0**

1. Open IIS Manager and expand **local computer**.
2. Right-click the **Web Sites** folder and click **Properties**.
3. On the **Service** tab, clear the **Run Web service in IIS 5.0 isolation mode** check box.

**To assign a pool designation to an ASP.NET application in IIS 6.0**

1. Open IIS Manager and expand **local computer**, and then navigate to the folder that contains the ASP.NET application.
2. Right-click the folder and click **Properties**. The application's **Properties** dialog box appears.
3. On the **Directory** tab, select the desired pool designation from the **Application Pool** drop down list.

**To create a pool designation in IIS 6.0**

1. Open IIS Manager and expand **local computer**.
2. Right-click the **Application Pools** folder, point to **New**, and then click **Application Pool**. The **Add New Application Pool** dialog box appears.
3. Enter the new pool designation in the **Application** pool text box, and then click **OK**.

**Related Information**

- Side-By-Side Support for ASP.NET
- Overview of ASP.NET Side-By-Side Support

# Assembly Unification in the .NET Framework

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In side-by-side execution, it is possible for an application to be composed of components that are built using different versions of the .NET Framework. This can create conflicts between the dependencies of the components. For example, suppose component A is built using the .NET Framework version 1.0 and needs to use version 1.0 of the System.Data assembly. Component B is built using the .NET Framework version 1.1 and needs to use version 1.1 of the System.Data assembly. If the assembly versions are incompatible, but are loaded at the same time, this could lead to unintended type casting exceptions, or even worse problems.

**How an Application Determines Which .NET Framework Version to Use**

To solve this problem, when an application uses components that are built on different versions of the .NET Framework, the runtime version that is associated with an application determines which version of the .NET Framework assemblies are used by the application and all its components. In the previous example, if the application is associated with version 1.1 of the .NET Framework, version 1.1 of the System.Data assembly is then loaded and shared across all components used by the application. The reference from component A to version 1.0 of the System.Data assembly is promoted at run time to reference version 1.1.

You can override this behavior by adding <**bindingRedirect**> elements to the Machine.config or Web.config files. This allows an application to use an updated assembly version that was designed to replace an existing assembly for certain types of applications. For example, if an updated version of System.Web.Service.dll is released in the future that supports SOAP version 1.2, you might want your application to use that version, instead of the version that was installed with the runtime originally.

| **Note** |
| --- |
| For more information, see the <bindingRedirect> topic in .NET Framework Help. |

This behavior can also be overridden in the ASP.NET host configuration file (Aspnet.config). ASP.NET uses this file to ensure that the version of System.Web.dll and System.Web.RegularExpressions.dll always matches the runtime version associated with an application, regardless of any overrides in the Web.config file.

Sometimes you might want to use a component that was built using a later version of ASP.NET in an application built with an earlier version. Remember that the ASPNET ISAPI version associated with an application always determines the version of the runtime used for an application. If the application is configured to use the earlier version of the runtime, the component will be automatically redirected at runtime to use that version, as well. When using a component built with a later version of the .NET Framework in an application built on an earlier version, consider the following:

- Make sure the component does not use any features or depend on any behaviors specific to the later version of the .NET Framework. These features might not be available in the earlier version of the runtime.

- ASP.NET uses script maps in IIS to bind an application to a version of the runtime. <**supportedRuntime**> and <**requiredRuntime**> configuration elements are not applicable to ASP.NET applications.

Alternatively, if the runtime version used to build the component is installed on the computer, you could also reconfigure the application to use the later runtime version. Since later versions of the .NET Framework are designed to be backward compatible, the application should work without any modifications. For information about configuring an ASP.NET application to use a specific version of the runtime, see Configuring an ASP.NET Application for an ASP.NET Version.

**Related Information**

- Side-by-Side Support in ASP.NET

# .NET Framework Redistribution

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you deploy an ASP.NET application to a production server, you also need to install the .NET Framework redistributable package (Dotnetfx.exe) to support the ASP.NET server controls and other portions of the application that run on the server. However, your application can use one version of the redistributable package while other applications on the server uses another.

When you include the .NET Framework redistributable package with an application that you are installing to a production server, ASP.NET is automatically updated. If the server hosts other applications that use a different version of ASP.NET, you need to incorporate the following command in your application's setup program:

**Dotnetfx.exe /c:"install /NOASPUPGRADE"**

Once the application is installed, you must create a script map between the application's virtual directory to the version of the .NET Framework that you installed. For more information about how to create custom script maps, see Setting Application Mappings in IIS 6.0.

For more information about the Microsoft .NET Framework redistributable package, see information in Dotnetfx.exe in the MSDN Library.

**Related Information**

- Overview of ASP.NET Deployment
- Overview of Deploying ASP.NET Applications
- Side-By-Side Support for ASP.NET

# Overview of ASP.NET Process Isolation

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When Microsoft ASP.NET is enabled on a server running Microsoft Windows Server-TM- 2003, by default it runs in the Internet Information Services (IIS) 6.0 worker process isolation mode. Using worker process isolation mode enables application pooling, recycling, and various application health detection. However, if your application has compatibility problems in this mode, you can switch to run IIS 6.0 to IIS 5.0 isolation mode.

If you install ASP.NET and the Microsoft .NET Framework on a server running Windows XP Professional or Windows 2000 Server, ASP.NET automatically provides its own process model to run in.

This section outlines the process model, and discusses how to use the ASP.NET process model.

This section includes the following information:

- Process Model Modes: Describes the differences between the IIS 6.0 worker process isolation process model and the ASP.NET process model.
- Using the ASP.NET Process Model: Describes how to enable and configure the ASP.NET process model in IIS 6.0.

**Related Information**

- ASP.NET Administration

# Process Model Modes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you install ASP.NET on a server running Windows Server 2003, ASP.NET applications are executed by IIS 6.0. IIS 6.0 runs in one of two distinct modes of operation, called application isolation modes. These are worker process isolation mode, the default, and IIS 5.0 isolation mode. IIS functions differently in the two isolation modes.

When ASP.NET runs in the worker process isolation mode, it runs in the W3wp.exe worker process and picks up the IIS 6.0 application pool process properties for its operation.

When running on IIS 6.0 in worker process isolation mode, ASP.NET uses the IIS 6.0 process model settings. For information about how to migrate ASP.NET process model configuration settings to IIS 6.0 application pool settings, see "Configuring IIS 6.0 to Host ASP.NET Applications" in Migrating IIS Web Sites to IIS 6.0.

When you configure IIS 6.0 to run in IIS 5.0 isolation mode, ASP.NET runs in its own process model, Aspnet_wp.exe, and uses its own configuration settings. These settings are stored in the computer-wide Machine.config configuration file.

In almost all situations, ASP.NET applications running on Windows Server 2003 should be run in worker process isolation mode. If your application uses read raw data filters, you should use IIS 5.0 isolation mode instead.

When installed on servers running Windows XP or Windows 2000 Server, ASP.NET runs in IIS 5.1 or IIS 5.0, respectively, and runs in the Aspnet_wp.exe process by default. The ASP.NET process model provides process recycling, health detection, and support for affinities between worker processes and particular CPUs on a server with multiple CPUs. For more information, see Using the ASP.NET Process Model.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**Procedures**

**To configure IIS 6.0 for IIS 5.0 isolation mode**

1. In IIS Manager, double-click the local computer, right-click **Web Sites**, and then click **Properties**.
2. On the **Service** tab, select the **Run WWW service in IIS 5.0 isolation mode** check box, and then click **OK**. The **Internet Services Manager** dialog box appears.
3. To start the Web service, click **Yes.**

| **Note** |
| --- |
| When you have completed this procedure, you must restart the Web service, which will temporarily interrupt the service. |

**Related Information**

- Overview of ASP.NET Process Isolation
- Using the ASP.NET Process Model

# Using the ASP.NET Process Model

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

One of the most important requirements for .NET Framework applications is reliability. An architecture in which applications run inside the server process (in IIS, Inetinfo.exe) does not produce a solid foundation for building reliable applications that can continue to run over a long period of time. Too many resources are shared on the process level, and it is too easy for an error to bring down the entire server process.

To solve this problem, ASP.NET provides an out-of-process execution model, Aspnet_wp.exe, which helps protect the server process from user code. It also enables you to apply heuristics to the lifetime of the process to improve the availability of your Web applications. Using asynchronous interprocess communication helps provide a balance between performance, scalability, and reliability.

Process model settings are exposed in the root configuration file for the computer, Machine.config. The configuration section is named <processModel> Element and is shown in the following example. On computers running Windows 2000 and Windows XP, the process model is enabled by default (`enable="true"`).

**Note**

On computers running Windows Server 2003, ASP.NET runs in IIS 6.0 worker process isolation mode by default. If your application is designed with specific ASP.NET process model settings, to be able to use them on Windows Server 2003 you must configure the server to run in IIS 5.0 isolation mode. For more information, see Configuring Isolation Modes.

```
<processModel
enable="true"
timeout="Infinite"
idleTimeout="Infinite"
shutdownTimeout="0:00:05"
requestLimit="Infinite"
requestQueueLimit="5000"
restartQueueLimit="10"
memoryLimit="60"
webGarden="false"
cpuMask="0xffffffff"
userName="machine"
password="AutoGenerate"
logLevel="Errors"
clientConnectedCheck="0:00:05"
comAuthenticationLevel="Connect"
comImpersonationLevel="Impersonate"
responseRestartDeadlockInterval="00:09:00"
responseDeadlockInterval="00:03:00"
maxWorkerThreads="20"
maxIoThreads="20"
/>
```

Most of configuration settings in this example help control when a new worker process begins to serve requests (gracefully taking the place of an old worker process). The **userName** and **password** attributes define the account under which the ASP.NET worker process runs. These default to the values **machine** and **autogenerate**, respectively. These values tell ASP.NET to use the built-in ASPNET account and to use a cryptographically strong random password stored in the Local Security Authority (LSA) for that account.

**Note**

| |
|---|
| On computers running Windows Server 2003, by default ASP.NET runs as part of the IIS 6.0 (W3wp.exe) process using the default identity NETWORK SERVICE. If IIS5 isolation mode is configured on a computer running Windows Server 2003, the ASP.NET worker process runs in the ASPNET account. |

If you want a more privileged process, you can set the **userName** attribute to **System**, which causes the ASP.NET worker process to run with the same identity as the Inetinfo.exe process. This process runs by default as the System identity. When so configured, the ASP.NET worker process will have the right to access nearly all resources on the local server. In Windows 2000, Windows XP, and Windows Server 2003 family systems, the System account also has network credentials and can access network resources as the machine account.

| **Caution** |
|---|
| Caution |

Changing the **userName** attribute to **System** enables access to and the ability to alter nearly all the resources on the computer, creating a significant security risk to your server.

The following example changes the **userName** attribute in the <**processModel**> section of the Machine.config file to run the process as System.

```
<processModel userName="system" password="autogenerate" />
```

The process model supports two types of recycling: reactive and proactive.

**Reactive Process Recycling**
Reactive process recycling occurs when a process is not functioning properly or is unable to serve requests. The process typically displays detectable symptoms, such as deadlocks, access violations, memory leaks, and so on, in order to trigger a process recycle. You can help control the conditions that cause a process restart by using the **requestQueueLimit**, **memoryLimit**, and **shutdownTimeout** configuration attributes. For more information about these attributes, see the <processModel> Element configuration section.

**Proactive Process Recycling**
Proactive process recycling restarts the worker process periodically even if the process is healthy. This can be a useful way to prevent denials of service due to conditions the process model is unable to detect. A process can be restarted after a specific number of requests or after a time-out period has elapsed. You can enable proactive process recycling using the **timeout**, **idleTimeout**, and **requestLimit** configuration attributes. For more information, see the <processModel> Element configuration section.

**Logging Process Model Events**
The process model can write events to the Windows event log when process cycling occurs. This is controlled by the **logLevel** configuration attribute. For more information, see the <processModel> Element configuration section.

When a cycling event occurs, if logging is enabled for that event, the following entries are written to the application event log.

| **Shutdown reason** | **Event Log Type** | **Description** |
|---|---|---|
| Unexpected | Error | The ASP.NET worker process shut down unexpectedly. |
| RequestQueueLimit | Error | The ASP.NET worker process restarted because the request queue limit was exceeded. |
| RequestLimit | Information | The ASP.NET worker process restarted because the request limit was exceeded. |
| Timeout | Information | The ASP.NET worker process restarted because the time-out interval was met. |
| IdleTimeout | Information | The ASP.NET worker process shut down because the idle time-out interval was met. |
| MemoryLimitExceeded | Error | The ASP.NET worker process restarted because the process memory limit was exceeded. |

**Enabling Web Gardens on Multiprocessor Computers**
The process model helps enable scalability on multiprocessor computers by distributing the work to several processes, one per

CPU, each with the processor affinity set to its own CPU. This eliminates cross-processor lock contention and is ideal for large symmetric multiprocessing (SMP) systems. This technique is called Web gardening. To enable Web gardening, you need to set the **cpuMask** and **webGarden** configuration settings in the <**processModel**> element. The following example enables Web gardening on a server running ASP.NET.

```
<configuration>

<system.web>

<processModel

enable="true"

timeout="15"

webGarden="true"

maxWorkerThreads="25"

maxIoThreads="25"/>

</system.web>

</configuration>
```

Web gardening has some side effects that you should be aware of:

- If your application uses session state, it must choose an out-of-process provider (NT Service or SQL).
- Application state and application statistics are per process, not per computer.
- Caching is per process, not per computer.

**Related Information**

- Overview of ASP.NET Process Isolation
- Process Model Modes
- processModel Element

# Overview of ASP.NET Performance Monitoring

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are a number of tools available to help you test and monitor the performance of your Web application. Microsoft Visual Studio .NET provides Application Center Test 1.0 (ACT), which simulates a large group of users by opening multiple connections to a Web server and rapidly sending HTTP requests. ASP.NET includes a number of performance counters that you can use to track the execution of your application. You can also use the built-in ASP.NET tracing feature to track code execution for a page or an application.

ACT is designed to stress-test Web servers and to analyze performance and scalability problems with Web applications. For more information about ACT, either install Visual Studio .NET and open ACT Help, or see information located in the MSDN Library.

When you are doing steady-state running analysis, it is often best to ignore the first request and any one-time initialization costs for objects. For example, when requesting an ASP.NET page for the first time, there is always the cost of compiling an instance of then **Page** class. Once that has occurred, performance will improve for subsequent requests.

This section includes the following information:

- Using the System Monitor (PerfMon) with ASP.NET Applications: Describes the System Monitor tool and how to use it to monitor the performance of ASP.NET applications.
- System and Application Performance Counters for ASP.NET: Lists the performance counters included with ASP.NET.
- Related Performance Counters: Provides guidelines for monitoring an ASP.NET application's performance using the System Monitor tool (PerfMon).

**Related Information**

- ASP.NET Administration

# Using the System Monitor (PerfMon) with ASP.NET Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Unlike traditional Active Server Pages (ASP), which exposed performance counters globally for the entire server computer, most of the ASP.NET performance counters are exposed per application. In the System Monitor tool (PerfMon), per-application counters are available under the **ASP.NET Applications** performance object. If there are multiple applications on the server, you need to specify a particular application instance when selecting a counter to monitor. However, there is a special __Total__ application instance in PerfMon that aggregates the counter values for all applications on a server. The value of __Total__ is usually a sum of the counter values.

ASP.NET also exposes global-only counters that are not bound to a particular application instance. These counters are located under the ASP.NET performance object.

**Procedures**

**To view the available ASP.NET counters on a computer running Windows Server 2003**

1. From the **Start** menu, point to **Administrative Tools**, and then click **Performance**.
2. Right-click the System Monitor details pane and click **Add Counters**.
3. To monitor any computer on which the monitoring console is run, click **Use local computer counters**.

   Or, to monitor a specific computer, regardless of where the monitoring console is run, click **Select counters from computer** and specify a computer name or IP address.
4. In **Performance object**, click an object to monitor.
5. To monitor all counters, click **All counters**.

   Or, to monitor only selected counters, click **Select counters from list** and select the counters you want to monitor.
6. To monitor all instances of the selected counters, click **All instances**.

   Or, to monitor only selected instances, click **Select instances from list** and select the instances you want to monitor.
7. Click **Add**, and then click **OK**.

For information about each of the ASP.NET system and ASP.NET application performance counters, see System and Application Performance Counters for ASP.NET.

**Related Information**

- Overview of ASP.NET Performance Monitoring
- System and Application Performance Counters for ASP.NET
- Related Performance Counters

# System and Application Performance Counters for ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP.NET supports two groups of performance counters: system and application. The former is exposed in PerfMon in the ASP.NET performance counter object; the latter is exposed in the ASP.NET Applications performance object. There is a significant difference between the State Server Sessions counters found in the ASP.NET performance object, which apply only to the server computer on which the state server is running, and the Sessions counters found in the ASP.NET Applications performance object, which apply only to user sessions that occur in process.

| **Note** |
| --- |
| The value associated with each performance counter is updated every 400 milliseconds. |

When monitoring the performance of your ASP.NET Web applications, you should always track the performance counters listed in the following table.

| Performance object | Performance counter |
| --- | --- |
| ASP.NET | Application Restarts |
| ASP.NET | Requests Queued |
| ASP.NET | Worker Process Restarts |
| ASP.NET Applications | Errors Total |
| ASP.NET Applications | Requests/Sec |
| Processor | % CPU Utilization |

The % CPU Utilization counter monitors CPU utilization on the Web server computer. Low CPU utilization or the inability to maximize CPU utilization regardless of client load can signify contention for resource or a lock in your Web application.

In addition, the performance counters listed in the following table can be valuable in determining problems with your Web applications' performance.

| Performance object | Performance counter |
| --- | --- |
| ASP.NET Applications | Pipeline Instance Count |
| .NET CLR Exceptions | # of Exceps Thrown |
| System | Context Switches/sec |

The # of Exceps Thrown counter displays the number of exceptions thrown in an application, since these can have performance implications. Some code paths, however, rely on exceptions for their proper functioning. For example, the HttpResponse.Redirect method always throws an uncatchable exception, ThreadAbortException. Therefore, it can be more useful to track the number of exceptions thrown using the Errors Total counter to see if the exception generated an error on the application.

The Context Switches/sec counter measures the rate at which thread contexts are switched by all CPUs in the Web server computer. A high number for this counter usually indicates either high contention in locks or many switches between user and kernel mode by the thread. Further investigation with sampling profilers and other tools might be warranted.

The counters in the ASP.NET and ASP.NET Applications performance objects are detailed in the following lists.

**ASP.NET System Performance Counters**

ASP.NET supports the following ASP.NET system performance counters. These aggregate information for all ASP.NET applications on a Web server computer, or they apply generally to a system of ASP.NET servers running the same applications. These can include Web farms and Web gardens.

| Performance Counter | Description |
|---|---|
| Application Restarts | The number of times that an application has been restarted during the Web server's lifetime. Application restarts are incremented with each Application_OnEnd event. An application restart can occur because of changes to the Web.config file, changes to assemblies stored in the application's \Bin directory, or too many changes in Web Forms pages. Unexpected increases in this counter can mean that unforeseen problems are causing your Web application to shut down. In such circumstances you should investigate as soon as possible.<br>**Note**<br>This value is reset every time the Internet Information Services (IIS) host is restarted. |
| Application Running | The number of applications running on the server computer. |
| Requests Disconnected | The number of requests that have been disconnected due to a communication failure. |
| Requests Queued | The number of requests waiting for service from the queue. When this number starts to increment linearly with respect to client load, the Web server computer has reached the limit of concurrent requests that it can process. The default maximum for this counter is 5,000. You can change this setting in the computer's Machine.config file. |
| Requests Rejected | The total number of requests not executed because of insufficient server resources to process them. This counter represents the number of requests that return a 503 HTTP status code, indicating that the server is too busy. |
| Request Wait Time | The number of milliseconds that the most recent request waited for processing in the queue. |
| Session State Server Connections Total | The total number of session-state connections to the computer on which out-of-process session-state data is stored. |
| Session SQL Server Connections Total | The total number of session-state connections to the Microsoft SQL Server-TM- database in which session-state data is stored. |
| State Server Sessions Abandoned | The number of user sessions that have been explicitly abandoned. These are sessions that are ended by specific user actions, such as closing the browser or navigating to another site. This counter is available only on the computer where the state server service (aspnet_state) is running. |
| State Server Sessions Active | The number of currently active user sessions. This counter is available only on the computer where the state server service (aspnet_state) is running. |
| State Server Sessions Timed Out | The number of user sessions that have become inactive through user inaction. This counter is available only on the computer where the state server service (aspnet_state) is running. |
| State Server Sessions Total | The number of sessions created during the lifetime of the process. This counter is the cumulative value of State Server Sessions Active, State Server Sessions Abandoned, and State Server Sessions Timed Out. This counter is available only on the computer where the state server service (aspnet_state) is running. |

| | |
|---|---|
| Worker Process Restarts | The number of times a worker process has been restarted on the server computer. A worker process can be restarted if it fails unexpectedly or when it is intentionally recycled. In the case of unexpected increases in this counter you should investigate as soon as possible. |
| Worker Process Running | The number of worker processes running on the server computer. |

**ASP.NET Application Performance Counters**

ASP.NET supports the following application performance counters, which you can use to monitor the performance of a single instance of an ASP.NET application. A unique instance appears for these counters, named __Total__, which aggregates counters for all applications on a Web server (similar to the global counters described in the first section of this topic). The __Total__ instance is always available. The counters will display zero when no applications are present on the server.

| **Performance Counter** | **Description** |
|---|---|
| Anonymous Requests | The number of requests that are using anonymous authentication. |
| Anonymous Requests/Sec | The number of requests per second that are using anonymous authentication. |
| Cache Total Entries | The total number of entries in the cache. This counter includes both internal use of the cache by the ASP.NET page framework and external use of the cache through exposed APIs. |
| Cache Total Hits | The total number of hits from the cache. This counter includes both internal use of the cache by the ASP.NET page framework and external use of the cache through exposed APIs. |
| Cache Total Misses | The number of failed cache requests per application. This counter includes both internal use of the cache by ASP.NET and external use of the cache through exposed APIs. |
| Cache Total Hit Ratio | The ratio of hits to misses for the cache. This counter includes both internal use of the cache by ASP.NET and external use of the cache through exposed APIs. |
| Cache Total Turnover Rate | The number of additions and removals to the total cache per second. It is useful in helping determine how effectively the cache is being used. If the turnover is large, the cache is not being used efficiently. |
| Cache API Entries | The total number of entries in the application cache. |
| Cache API Hits | The total number of hits from the cache when it is accessed through only the external cache APIs. This counter does not track any use of the cache internally by ASP.NET. |
| Cache API Misses | The total number of failed requests to the cache when accessed through the external cache APIs. This counter does not track any use of the cache internally by ASP.NET. |
| Cache API Hit Ratio | The cache hit-to-miss ratio when accessed through the external cache APIs. This counter does not track any use of the cache internally by ASP.NET. |
| Cache API Turnover Rate | The number of additions and removals to the cache per second, when used through the external APIs (excluding internal use by the ASP.NET page framework). It is useful in helping determine how effectively the cache is being used. If the turnover is large, then the cache is not being used effectively. |

| | |
|---|---|
| Compilations Total | The total number of compilations that have taken place during the lifetime of the current Web server process. This occurs when a file with an .aspx, .asmx, .ascx, or .ashx extension, or a code-behind source file, is dynamically compiled on the server.<br><br>**Note**<br>This number will initially climb to a peak value as requests are made to all parts of an application. Once a compilation occurs, however, the resulting binary is saved to disk, where it is reused until its source file changes. This means that, even in the event of a process restart, the counter can remain at zero (inactive) until the application is modified or redeployed. |
| Debugging Requests | The number of requests that occur while debugging is enabled. |
| Errors During Preprocessing | The number of errors that occurred during parsing. Excludes compilation and run-time errors. |
| Errors During Compilation | The number of errors that occur during dynamic compilation. Excludes parser and run-time errors. |
| Errors During Execution | The total number of errors that occur during the execution of an HTTP request. Excludes parser and compilation errors. |
| Errors Unhandled During Execution | The total number of unhandled errors that occur during the execution of HTTP requests.<br><br>**Note**<br>An unhandled error is any uncaught run-time exception that escapes user code on the page and enters the ASP.NET internal error-handling logic. Exceptions to this rule occur when:<br><br>• Custom errors are enabled, an error page is defined, or both.<br>• The Page_Error event is defined in user code and either the error is cleared (using the HttpServerUtility.ClearError method) or a redirect is performed. |
| Errors Unhandled During Execution/Sec | The number of unhandled exceptions per second that occur during the execution of HTTP requests. |
| Errors Total | The total number of errors that occur during the execution of HTTP requests. Includes any parser, compilation, or run-time errors. This counter is the sum of the Errors During Compilation, Errors During Preprocessing, and Errors During Execution counters. A well-functioning Web server should not generate errors. If errors occur in your ASP.NET Web application, they may skew any throughput results because of very different code paths for error recovery. Investigate and fix any bugs in your application before performance testing. |
| Errors Total/Sec | The number of errors per second that occur during the execution of HTTP requests. Includes any parser, compilation, or run-time errors. |
| Output Cache Entries | The total number of entries in the output cache. |
| Output Cache Hits | The total number of requests serviced from the output cache. |

| | |
|---|---|
| Output Cache Misses | The number of failed output-cache requests per application. |
| Output Cache Hit Ratio | The percentage of total requests serviced from the output cache. |
| Output Cache Turnover Rate | The number of additions and removals to the output cache per second. If the turnover is large, the cache is not being used effectively. |
| Pipeline Instance Count | The number of active request pipeline instances for the specified ASP.NET application. Since only one execution thread can run within a pipeline instance, this number gives the maximum number of concurrent requests that are being processed for a given application. In most circumstances it is better for this number be low when under load, which signifies that the CPU is well utilized. |
| Request Bytes In Total | The total size, in bytes, of all requests. |
| Request Bytes Out Total | The total size, in bytes, of responses sent to a client. This does not include standard HTTP response headers. |
| Requests Executing | The number of requests currently executing. |
| Requests Failed | The total number of failed requests. Any and all status codes greater than or equal to 400 will increment this counter.<br>**Note**<br>Requests that cause a 401 status code will increment this counter and the Requests Not Authorized counter. Requests that cause a 404 or 414 status code will increment this counter and the Requests Not Found counter. Requests that cause a 500 status code will increment this counter and the Requests Timed Out counter.<br>**Note**<br>The equivalent ASP counter would also increment on requests rejected, which cannot be done (because the rejection is done by IIS and not the process model). |
| Requests Not Found | The number of requests that failed because resources were not found (status code 404, 414). |
| Requests Not Authorized | The number of requests that failed due to no authorization (status code 401). |
| Requests Succeeded | The number of requests that executed successfully (status code 200). |
| Requests Timed Out | The number of requests that timed out (status code 500). |
| Requests Total | The total number of requests since the service was started. |
| Requests/Sec | The number of requests executed per second. This represents the current throughput of the application. Under constant load, this number should remain within a certain range, barring other server work (such as garbage collection, cache cleanup thread, external server tools, and so on). |
| Sessions Active | The number of sessions currently active. This counter is supported only with in-memory session state. |
| Sessions Abandoned | The number of sessions that have been explicitly abandoned. This counter is supported only with in-memory session state. |

| | |
|---|---|
| Sessions Timed Out | The number of sessions that timed out. This counter is supported only with in-memory session state. |
| Sessions Total | The number of sessions timed out. This counter is supported only with in-memory session state. |
| Transactions Aborted | The number of transactions aborted. |
| Transactions Committed | The number of transactions committed. |
| Transactions Pending | The number of transactions in progress. |
| Transactions Total | The total number of transactions since the service was started. |
| Transactions/Sec | The number of transactions started per second. |

**Related Information**

- Overview of ASP.NET Performance Monitoring
- Using the System Monitor (PerfMon) with ASP.NET Applications
- Related Performance Counters

# Related Performance Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The following sections list other performance objects and counters that can be useful when monitoring your ASP.NET applications. The list is alphabetical by performance object, then by counter.

For a full list of performance counters on a computer running a version of the .NET Framework, see the "Performance Counters" topic in .NET Framework Help . For objects and counters associated specifically with ASP.NET, see System and Application Performance Counters for ASP.NET.

**.NET CLR Exceptions**
### # of Exceptions Thrown/Sec

The total number of managed exceptions thrown per second. As this number increases, performance degrades.

**.NET CLR Loading**
### Bytes in loader Heap

The number of bytes committed by the class loader across all application domains.

### Current App Domains

The current number of application domains loaded in the process. The value of this counter should be the same as the number of virtual directories, plus 1 (to account for the default application domain).

### Current Assemblies

The current number of assemblies loaded in the process. Note that assemblies cannot be unloaded from an application domain; the entire application domain must be unloaded. This is an issue for sites with a large number of pages.

**.NET CLR Memory**
### # Bytes in all heaps

Displays the sum of the **Gen 0 Heap Size**, **Gen 1 Heap Size**, **Gen 2 Heap Size**, and the **Large Object Heap Size** counters. This counter indicates, in bytes, the current memory allocated on the garbage collection heaps.

### # Gen 0 Collections

Displays the number of times that generation 0 objects (that is, the youngest, most recently allocated objects) have been reclaimed by garbage collection since the application started.

Generation 0 garbage collection occurs when the available memory in generation 0 is not sufficient to satisfy an allocation request. This counter is incremented at the end of a generation 0 garbage collection. Higher-generation garbage collections include all lower-generation collections. This counter is explicitly incremented when a higher-generation (generation 1 or 2) garbage collection occurs.

This counter displays the last observed value. The **_Global_** counter value is not accurate and should be ignored.

### # Gen 1 Collections

Displays the number of times that generation 1 objects have been reclaimed by garbage collection since the application started.

The counter is incremented at the end of a generation 1 garbage collection. Higher-generation garbage collections include all lower generation collections. This counter is explicitly incremented when a higher-generation (generation 2) garbage collection occurs.

This counter displays the last observed value. The **_Global_** counter value is not accurate and should be ignored.

### # Gen 2 Collections

Displays the number of times that generation 2 objects have been reclaimed by garbage collection since the application started. The counter is incremented at the end of a generation 2 garbage collection (also called a full garbage collection).

This counter displays the last observed value. The _**Global**_ counter value is not accurate and should be ignored.

**Memory**
### Available Mbytes

The amount of physical memory available to processes running on the computer, in megabytes. It is calculated by adding the amount of space on the zeroed, free, and standby memory lists. Free memory is ready for use; zeroed memory is filled with zeros to prevent later processes from seeing data used by a previous process; standby memory is removed from a process's working set (its physical memory) to be stored on the server's hard drive but is still available to be recalled. This counter displays the last observed value only; it is not an average.

**Process**
### % Processor Time

The percentage of the elapsed time on the processor that it took all threads associated with a single process to perform execution instructions. An instruction is the basic unit of execution in a computer; a thread is the object that executes instructions; and a process is the object created when a program is run. Code executed to handle some hardware interrupts and trap conditions are included in this count.

### Handle Count

The total number of handles currently open that are associated with this process. This number is equal to the sum of the handles currently open for each thread in this process.

### Private Bytes

The current amount of memory, in bytes, that this process has allocated that cannot be shared with other processes.

### Thread Count

The number of threads currently active in this process. An instruction is the basic unit of execution in a processor, and a thread is the object that executes instructions. Every running process has at least one thread.

**Processor**
### % Processor Time

The percentage of sampled time that the processor spends to execute a non-Idle thread. It is calculated by measuring how long the idle thread is active in the sample interval and subtracting that time from the interval duration. (Each processor has an idle thread that consumes cycles when no other threads are ready to run). This counter is the primary indicator of processor activity, and displays the average percentage of busy time observed during the sample interval.

**Web Service**
### Bytes Received/Sec

The number of bytes per second received by the Web service.

### Bytes Sent/Sec

The number of bytes per second sent by the Web service.

### Current Connections

The current number of connections established with the Web service.

### ISAPI Extension Requests/Sec

The number of ISAPI Extension requests per second that are being processed by the Web service.

### Total Method Requests/Sec

The number of HTTP requests per second received by the Web service.

**Related Information**

- Overview of ASP.NET Performance Monitoring
- System and Application Performance Counters for ASP.NET

# Additional Resources for ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

There are a number of other useful sources of information about ASP.NET, including Microsoft ASP.NET Help found in the **Creating ASP.NET Applications** section of Microsoft .NET Framework Help.

In addition, the following resources contain a great deal of information about ASP.NET:

- The official ASP.NET Web site contains links to tools, forums, articles, sample applications, a gallery of ASP.NET server controls, links to third-party Web sites, and more.

- The official .NET Framework Community Website Web site contains links to tools, white papers, sample applications and third-party Web sites. This site can help you understand how to access other portions of the .NET Framework from your ASP.NET applications.

- The MSDN Library contains technical articles and columns about ASP.NET, as well as links to other resources. Go to the ASP .Net Developer Center, expand the **.NET Development** node in the Table of Contents, and then click **ASP.NET**.

- The ASP.NET QuickStart, included when you install the .NET Framework, contains overview information and samples of how to develop and administer ASP.NET applications.

**Related Information**

- ASP.NET Administration

# About ASP

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Microsoft Active Server Pages (ASP) is a server-side scripting environment that you can use to create dynamic and interactive Web pages, and to build powerful Web applications. When the server receives a request for an ASP file, it processes server-side script code contained in the file to build the HTML Web page that is sent to the browser. In addition to server-side script code, ASP files can contain HTML (including related client-side scripts) as well as calls to COM components that perform a variety of tasks, such as connecting to a database or processing business logic.

FrontPage Server Extensions does not support unicode content directory names or WebDAV. New FrontPage Server Extensions virtual servers should be assigned to the MSSharePointPool application pool. Also, do not make changes to the MSSharePointPool application pool configuration.

| **Note** |
| --- |
| If you have a Web site that uses Sharepoint™ Team Services (STS) from Microsoft, the following types of pages do not work without excluding the virtual directory containing the pages from the STS ISAPI filter: |

- ASP pages
- ASP.NET pages
- PHP pages
- PerlScript pages

**Procedures**

**To enable these types of pages to work on your STS Web site**

1. In the **Central Administration** Web site for STS, choose to **Manage** your virtual directory or Web site.
2. Under **Defined Managed Paths**, type the path to the virtual directory that contains your excluded pages (such as <code>\MyASPPages</code>) and select **Exclude**.

The topics in this section introduce ASP, explain the basic concepts of scripting with ASP, and discuss more complex application issues such as how to maintain state.

| **Important** |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0. |

This section includes the following information:

- Checklist: ASP Security: Provides a checklist of issues to monitor when developing ASP applications that include user input or database access.
- What's New in ASP: Describes new features added in this release.
- Important Changes in ASP: Describes functionality and behavior that has changed for this release.

**Related Information**

- For information about writing server-side scripts with ASP and creating ASP applications, see the Active Server Pages Guide on MSDN Online or in the Internet Information Services (IIS) Software Development Kit (SDK).

- Redirect Reference: This section explains how to format a server redirection correctly. For reference topics related to ASP, see the IIS Web Application Technology Reference section of the IIS Software Developer Kit (SDK).

# Checklist: ASP Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The following is a list of issues to keep in mind when you are developing Web pages using ASP.

| | Step | Reference |
|---|---|---|
| | Never trust user input to be of an appropriate size or contain appropriate characters. Always verify user input before using it to make decisions. The best option is to create a COM+ component which you can call from an ASP page to verify user input. You can also use the Server. HTMLEncode method, the Server. URLEncode method, or one of the code examples at the bottom of this page. | Writing Secure Code |
| | Do not create database connection strings in an ASP page by concatenating strings of user input together. Malicious users can inject code in their input to gain access to your database. If you are using an SQL database, use stored procedures for creating database connection strings. | Writing Secure Code |
| | Do not use the default SQL administrator account name, *sa*. Everyone who uses SQL knows that the sa account exists. Create a different SQL administrative account with a strong password and delete the sa account. | "Passwords" in Help and Support Center for Windows 2003. |
| | Before you store client user passwords, hash them, base64encode them, or use **Server.HTMLEncode** or **Server.URLEncode** to encode them. You can also use one of the code examples at the bottom of this page to verify the characters in the client's password. | "Encrypting File System" in Help and Support Center for Windows 2003. |
| | Do not put administrative account names or passwords in administration scripts or ASP pages. | Writing Secure Code |
| | Do not make decisions in your code based on request headers because header data can be fabricated by a malicious user. Before using request data, always encode it or use one of the code examples below to verify the characters it contains. | Writing Secure Code |
| | Do not store secure data in cookies or hidden input fields in your Web pages. | Writing Secure Code |
| | When writing ISAPI applications, filters, or COM+ objects, watch out for buffer over-runs caused by assuming sizes of variables and data. Also watch out for canonicalization issues that can be caused by interpreting things, like absolute path names, as relative path names or URLs. | Writing Secure Code |
| | When you switch an ASP application from running in Single Threaded Apartment (STA) to Multi-Threaded Apartment (MTA), the impersonation token becomes obsolete. This can cause the application to run with no impersonation, effectively letting it run with the identity of the process which might allow access to other resources. If you must switch threading models, disable the application and unload it before you make the change. | Configuring Applications to Use COM+ Services; "COM+ Services" in the IIS Software Development Kit (SDK) |

**Code Example**

This code example includes a function that removes potentially harmful characters from a string that is sent to the function. In both examples above, the code page is specified to ensure proper encoding. The example below is in Microsoft Visual Basic Scripting Edition (VBScript):

```
<%@ LANGUAGE="VBScript" %>

<%

Response.CodePage = 1252

Response.Write("Hello, " & RemoveBadCharacters(Request.Form("UserName")))

Response.Write("<BR>This is why you received an error:")
```

```
Function RemoveBadCharacters(strTemp)

Dim regEx

Set regEx = New RegExp

regEx.Pattern = "[^\s\w]"

regEx.Global = True

RemoveBadCharacters = regEx.Replace(strTemp, "")

End Function

%>
```

The example below is in Microsoft JScript:

```
<%@ LANGUAGE="JScript" %>

<%

Response.CodePage = 1252;

Response.Write("Hello, " + RemoveBadCharacters(Request.Form("UserName")));

Response.Write("<BR>This is why you received an error:");

function RemoveBadCharacters(strTemp) {

strTemp = strTemp.replace(/[^\s\w]/g,"");

return strTemp;

}

%>
```

**Related Information**

- IIS 6.0 Security Best Practices

# What's New in ASP

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In addition to internal performance enhancements and optimizations, this version of Active Server Pages (ASP) has the following new features:

- Better International and UTF-8 Support
- Improved POST Support
- ASP Hang Detection
- Caching of Popular Files
- UNC Enhancements
- COM+ Services in ASP
    - **Apartment Model Selection**
    - **Side-by-side Assemblies**
    - **COM+ Partitions**
    - **Tracker**
    - **Transactions**
- New Metabase Properties for ASP
- ASP Debugging

**Better International and UTF-8 Support**
UTF-8 support has been expanded to all ASP built-in object properties and methods. The last version of ASP supported UTF-8 only for Response.Write.

**Improved POST Support**
ASP can now read chunked-encoded POST data from a client.

**ASP Hang Detection**
When an IIS Web site is busy there may be instances when the maximum number of ASP threads has been spawned and some of the ASP threads are hung, resulting in degraded performance. HSE_REQ_REPORT_UNHEALTHY, and the World Wide Web Publishing Service (WWW service) will recycle the worker process hosting ASP.dll and make an entry in the event log.

| **Note** |
| --- |
| Threads are considered *hung* if they do not respond to a timeout. |

**Caching of Popular Files**
ASP now caches some of the most-recently used files on disk, in addition to keeping a memory cache. The default location of the disk cache is **%systemroot%\System32\inetsrv\ASP** Compiled Templates. You can change the location of the disk cache by setting the AspDiskTemplateCacheDirectory Metabase Property.

**UNC Enhancements**

ASP can now handle UNC paths in ASP scripts.

**COM+ Services in ASP**

If you wanted to use COM+ services from an ASP application in previous versions of IIS, you had to create a COM component that would call the methods of those services. Now, you can configure your ASP application to use the following COM+ services without having to create a COM component:

- Apartment Model Selection:

  ASP is now capable of running all of its threads in a multi-threaded apartment (MTA). If your COM components are primarily free-threaded or both-threaded, running the ASP threads as MTA can improve performance significantly.

  To enable an ASP application to run in an MTA, you can use the metabase setting, AspExecuteInMTA Metabase Property, at the application level. This means that you can have one application running on ASP MTA threads and a second application running on ASP STA (single-threaded apartment) threads. The default for ASP threads continues to be STA.

  | **Important** |
  | --- |
  | When you switch an ASP application from running in STA to MTA (or from MTA to STA), the impersonation token becomes obsolete. This can cause the application to run with no impersonation, effectively letting it run with the identity of the process which might allow access to other resources. If you must switch threading models, disable the application and unload it before you make the change. |

  The following example sets the Default Web Site application (*W3SVC/1/ROOT*) to execute in MTA:

```
    On Error Resume Next
    set providerObj = GetObject("winmgmts://MyMachine/root/MicrosoftIISv2")
    ' Get a reference to the ASP application called Default Web Site
    set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
    WScript.Echo "Before: AspExecuteInMTA = " &
IIsWebVirtualDirSettingObj.AspExecuteInMTA
    ' Set the ASP application to execute in MTA
    IIsWebVirtualDirSettingObj.AspExecuteInMTA = 1
    IIsWebVirtualDirSettingObj.Put_()
    WScript.Echo "After: AspExecuteInMTA = " &
IIsWebVirtualDirSettingObj.AspExecuteInMTA
```

- Side-by-side Assemblies:

  Side-by-side (SxS) assemblies allow ASP applications to specify which version of a system DLL or classic COM component to use, such as WinHTTP 5.1, Shell Common Controls version 6.0 (Comctl32.dll), GDI Plus version 1.0 (GDIplus.dll), and Visual C++ Run-time Libraries version 6.0. For example, if your ASP application relies on MSXML version 2.0, you can ensure that your application still uses MSXML version 2.0 even after service packs are applied to the server. Any new version of MSXML is still installed on the computer, but version 2.0 remains and is used by your application. Configuring SxS assemblies requires that you know the path to the DLL, and that the COM+ manifest file exists in every virtual directory that needs to use the DLL. The COM+ manifest is an XML file that has information about where a DLL is installed. IIS does not verify that the manifest exists. A manifest looks like the following file excerpt:

```
    <?xml version="1.0" encoding="UTF-8" standalone="yes"?>
    <assembly xmlns="urn:schemas-microsoft-com:asm.v1" manifestVersion="1.0">
    <assemblyIdentity publicKeyToken="75e377300ab7b886" type="win32"
name="Test4Dir" version="1.0.0.0" processorArchitecture="x86"/>
    <file name="DirComp.dll" hash="35ca6f27b11ed948ac6e50b75566355f0991d5d9"
hashalg="SHA1">
    <comClass clsid="{6C6CC20E-0F85-49C0-A14D-D09102BD7CDC}"
progid="DirComp.PathInfo" threadingModel="apartment"/>
    <typelib tlbid="{AA56D6B8-9ADB-415D-9E10-16DD68447319}" version="1.0"
```

```
    helpdir=""/>
        </file>
        </assembly>
```

You can enable side-by-side assemblies on the IIS side programmatically or by using IIS Manager.

**To enable side-by-side assemblies on the IIS side using IIS Manager**

1. In IIS Manager, right-click a Web site or virtual directory and click **Properties**.
2. Click the **Virtual Directory** tab, and click **Configuration**.
3. If the configuration button is not available, it is because you have not created an application for this virtual directory. Click **Create** to create an application.
4. In the **Application Configuration** dialog box, click the **Options** tab.
5. Select the **Enable Side by Side assemblies** check box.
6. In the **Manifest file name** box, type the name of the COM+ manifest file.
7. Click **OK** twice.

To enable side-by-side assemblies programmatically, set the **AspEnableSxs** flag of the AspAppServiceFlags Metabase Property. Also set the AspSxsName Metabase Property to the name of the COM+ manifest name. Set both metabase properties at the application level.

| **Important** |
| --- |
| Only one version of a system DLL can be used in any application pool, even though this feature is configurable at the application level. For example, if application App1 uses MDAC, version 2.5 and application App2 uses MDAC, version 2.4, then App1 and App2 should not be in the same application pool. If they are, the application that is loaded first has its version of MDAC loaded, and the other application is forced to use it until the applications are unloaded. |

The following example enables side-by-side assemblies on the Default Web Site application (*W3SVC/1/ROOT*). Notice that after setting only the **AspEnableTracker** property, the **AspAppServiceFlags** property changes as well:

```
On Error Resume Next
    set providerObj = GetObject("winmgmts://MyMachine/root/MicrosoftIISv2")
    ' Get a reference to the ASP application called Default Web Site
    set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
    WScript.Echo "Before: AspEnableSxs = " & IIsWebVirtualDirSettingObj.AspEnableSxs
    WScript.Echo "        AspSxsName = " & IIsWebVirtualDirSettingObj.AspSxsName
    WScript.Echo "        AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
    ' Set the ASP application to enable COM+ side-by-side assemblies
    IIsWebVirtualDirSettingObj.AspEnableSxs = 1
    ' Set the AspSxsName property
    IIsWebVirtualDirSettingObj.AspSxsName = "VersionInfo"
    ' Save the values to the IIS metabase
    IIsWebVirtualDirSettingObj.Put_()
    ' Get the reference again in order to refresh the AspAppServiceFlags property.
    set IIsWebVirtualDirSettingObj = Nothing
    set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
    WScript.Echo "After: AspEnableSxs = " & IIsWebVirtualDirSettingObj.AspEnableSxs
    WScript.Echo "       AspSxsName = " & IIsWebVirtualDirSettingObj.AspSxsName
    WScript.Echo "       AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
```

- COM+ Partitions:

  COM+ Partitions can be used to isolate Web applications into their own COM+ partitions. This is useful to prevent one Web application from accessing the private COM+ applications, configuration information, and data of another Web

application. COM+ partitions can hold different versions of your own custom COM components. For example, if you host Web sites for two competing companies that both use COM+ in their Web applications, you can use COM+ partitions to ensure that one company's Web application cannot access the COM+ components in the other company's Web applications. If one of those companies asks you to change certain features in a COM+ application that they both use, you can isolate the new version of that COM+ application in the partition that is linked to their Web application.

To enable COM+ partitions on the IIS side, set the **AspUsePartition** flag of the AspAppServiceFlags Metabase Property at the application level. The partition is identified by a GUID (created using the Component Services Manager snap-in), which can be set at the AspPartitionID Metabase Property. If no partition is specified, the default system partition is used. For more information, please see "Creating and Configuring COM+ Partitions" in the COM+ SDK or online at COM+ (Component Services).

| **Important** |
| --- |
| Only one version of a COM+ component can be used in any application pool, even though this feature is configurable at the application level. For example, if application App1 uses version 1.0 of a custom COM+ application called Shop.dll, and application App2 uses version 2.0 of Shop.dll, then App1 and App2 should not be in the same application pool. If they are, the application that is loaded first has its version of Shop.dll loaded, and the other application is forced to use it until the applications are unloaded. |

The following example enables partitions on the Default Web Site application (*W3SVC/1/ROOT*). Notice that after setting only the **AspEnableTracker** property, the **AspAppServiceFlags** property changes as well.

```
    On Error Resume Next
    set providerObj = GetObject("winmgmts://MyMachine/root/MicrosoftIISv2")
    ' Get a reference to the ASP application called Default Web Site
    set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
    WScript.Echo "Before: AspUsePartition = " &
IIsWebVirtualDirSettingObj.AspUsePartition
    WScript.Echo "        AspPartitionID = " &
IIsWebVirtualDirSettingObj.AspPartitionID
    WScript.Echo "        AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
    ' Set the ASP application to enable COM+ partitioning
    IIsWebVirtualDirSettingObj.AspUsePartition = 1
    ' Set the AspPartitionID property to the GUID configured in Component Services
Manager
    ' when you created the COM+ partition
    IIsWebVirtualDirSettingObj.AspPartitionID = "{00000000-0000-0000-0000-
000000000000}"
    ' Save the values to the IIS metabase
    IIsWebVirtualDirSettingObj.Put_()
    ' Get the reference again in order to refresh the AspAppServiceFlags property.
    set IIsWebVirtualDirSettingObj = Nothing
    set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
    WScript.Echo "After: AspUsePartition = " &
IIsWebVirtualDirSettingObj.AspUsePartition
    WScript.Echo "       AspPartitionID = " & IIsWebVirtualDirSettingObj.AspPartitionID
    WScript.Echo "       AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
```

- Tracker:

  COM+ tracker allows administrators or developers to debug ASP applications. For example, if a Web application is causing problems on your server, you can enable COM+ tracker to determine when ASP pages are being loaded, when COM components are loaded, and when threads leave a page. After you have debugged your application, you can disable COM+ tracker to return your application to normal performance speed.

  To enable COM+ tracker on the IIS side, set the **AspEnableTracker** flag of the AspAppServiceFlags Metabase Property at the application level.

The following example enables tracking on the Default Web Site application (*W3SVC/1/ROOT*). Notice that after setting only the **AspEnableTracker** property, the **AspAppServiceFlags** property changes as well:

```
    On Error Resume Next
    set providerObj = GetObject("winmgmts://MyMachine/root/MicrosoftIISv2")
    ' Get a reference to the ASP application called Default Web Site
    set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
    WScript.Echo "Before: AspEnableTracker = " &
IIsWebVirtualDirSettingObj.AspEnableTracker
    WScript.Echo "        AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
    ' Set the ASP application to enable COM+ tracking
    IIsWebVirtualDirSettingObj.AspEnableTracker = 1
    IIsWebVirtualDirSettingObj.Put_()
    ' Get the reference again in order to refresh the AspAppServiceFlags property.
    set IIsWebVirtualDirSettingObj = Nothing
    set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
    WScript.Echo "After: AspEnableTracker = " &
IIsWebVirtualDirSettingObj.AspEnableTracker
    WScript.Echo "       AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
```

- Transactions: Transaction support in ASP pages has been available since IIS version 4.0. However, IIS has changed its internal implementation of transactions to take advantage of COM+ services. IIS no longer creates a COM component to house a transaction but accesses the COM+ transaction service directly, thus greatly improving the speed of transaction processing. As always, transactions can only be configured at the page level using the @TRANSACTION directive and the events of the ASP built-in object called ObjectContext.

**New Metabase Properties for ASP**

The following metabase settings have been added for this release of ASP.

| Metabase Property Name | Purpose |
|---|---|
| AspAppServiceFlags Metabase Property | Enables features like COM+ side-by-side assembly (formerly known as Fusion), partitioning, and so on. |
| AspDiskTemplateCacheDirectory Metabase Property | Specifies the location of the ASP disk cache. |
| AspExecuteInMTA Metabase Property | Enables ASP threads to execute in a multi-threaded apartment. |
| AspKeepSessionIDSecure Metabase Property | Sends the ASP session cookie to a browser securely. |
| AspMaxDiskTemplateCacheFiles Metabase Property | Specifies the ASP Disk cache maximum. |
| AspPartitionID Metabase Property | Specifies the COM+ partition to use for the application. |
| AspRunOnEndAnonymously Metabase Property | Enables ASP to run the global.asa Application_OnEnd and Session_OnEnd events anonymously. |
| AspBufferingLimit Metabase Property | Specifies the limit of the buffer size. |
| AspMaxRequestEntityAllowed Metabase Property | Specifies the maximum number of bytes allowed in the entity body of an ASP request. |

| | |
|---|---|
| AspSxsName Metabase Property | Allows ASP applications to specify which version of a system DLL or classic COM component to use. |

**ASP Debugging**

Because the worker process, W3wp.exe, runs as the Network Service account in IIS 6.0 worker process isolation mode, you must configure Launch and Access permissions to enable ASP debugging for Script Debugger and Visual InterDev. For more information, see Enabling ASP Debugging.

# Important Changes in ASP

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

ASP has undergone several important changes and enhancements to improve security and performance:

| Note |
| --- |
| To learn about new ASP features, see What's New in ASP. |

- ASP Pages are Disabled by Default
- Using Parent Paths
- The "Templates Cached" Performance Counter
- Global.asa Events
- ASP Metabase Properties Removed
- ASP Debugging

**ASP Pages are Disabled by Default**

| Important |
| --- |
| To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down -- which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable these features, IIS returns a 404 error. You can enable these features through the Web Service Extensions node in IIS Manager. For more information about how to troubleshoot 404 errors and other issues, see Troubleshooting in IIS 6.0. |

**Using Parent Paths**

As a security precaution, the AspEnableParentPaths Metabase Property setting is now set to 0 by default. This affects your application if it has a Web page that contains the #include server-side include directive and uses ".." notation to refer to a parent directory.

**The "Templates Cached" Performance Counter**

The **AspTemplatesCached**performance counter now counts the number of cached ASP files, both on disk and in memory. In previous versions, because there was no disk-based caching, this counter referred only to the in-memory cache. There is also a new counter called **AspInMemoryTemplatesCached** that reflects the number of ASP pages in the memory cache; it's the number of templates cached in total memory.

**Global.asa Events**

Earlier versions of ASP executed events in the security context (or user identity) of the host process because during these events there is no user context. This caused problems such as access denied errors when writing to a file in the Session_OnEnd event. ASP, by default, now runs the global.asa events, Application_OnEnd, and Session_OnEnd, anonymously. This can be controlled by the metabase setting, **AspRunOnEndAnonymously**.

**ASP Metabase Properties Removed**

The following metabase settings are no longer used by ASP, even though they are still in the metabase for compatibility with existing administration scripts

| Metabase Property Name |
| --- |
| AspAllowOutOfProcComponents Metabase Property |

| |
|---|
| AspTrackThreadingModel Metabase Property |
| AspExceptionCatchEnable Metabase Property |

**ASP Debugging**

Because the worker process, W3WP, runs as the Network Service account in IIS 6.0 worker process isolation mode, you must configure Launch and Access permissions to enable ASP debugging for Script Debugger and Visual InterDev. For more information, see Enabling ASP Debugging.

# Troubleshooting in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Many of the design changes in Internet Information Services (IIS) 6.0 directly address the need to secure the World Wide Web Publishing Service (WWW service) as a whole, and Web and FTP sites in particular. You might experience errors if you have not enabled certain features or services that are locked, or disabled, by default. This topic describes some of the symptoms of these errors and the processes to remedy them (or a link to a topic that describes how to remedy the error). For a comprehensive list of changes in IIS 6.0, including a list of security changes and a brief explanation of the new process model, see What's Changed in IIS 6.0.

**In This Section**


- Dynamic or Static Content Errors
- Connection Errors
- Miscellaneous Errors

# Dynamic or Static Content Errors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section includes troubleshooting information about issues with dynamic content and static content.

**In This Section**


- Applications are denied access to resources
- Requests for dynamic content return 404 error
- All requests with /bin in the URL are rejected and return a 404 error
- Requests for static files return 404 error
- Worker process recycling drops application session state
- Server-side include directives (#include) return 404 error (for .stm files) or 0131 error (for .asp files)
- ASP generates Permission Denied errors in event log for global.asa
- CGI processes do not start
- ASP.NET pages are returned as static files
- Collaboration Data Objects for Windows NT Server fail

# Applications are denied access to resources

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

After a clean install, IIS 6.0 runs in Worker Process Isolation Mode in IIS 6.0. Applications running in this mode use the Network Service identity, by default. Network Service is an account with few user rights and therefore provides better security by restricting access to resources on the Web server. If you migrate applications to IIS 6.0 while the server is in worker process isolation mode, and if your applications previously ran in-process (in Inetinfo.exe) as LocalSystem, the applications may fail to access resources because of the restrictions set forth by the Network Service identity. The LocalSystem account has access to almost all resources on the operating system, and therefore creates serious security implications. You should avoid using the LocalSystem account when possible. If it is absolutely necessary to use the LocalSystem account on an application, run that application in a new application pool in its own virtual directory so you can reduce the attack surface by isolating the application. As an alternative, and if your application needs permission to use the Trusted Computing Base (TCB), run the application as a configurable identity and assign the TCB permission to the configurable identity. This alternative, however, still presents a security risk because the TCB permission is very powerful.

| **Note** |
| --- |
| By default IIS 6.0 worker processes run under the low-privileged Network Service user account. This account does not have the appropriate permissions to write to the System event log, so application events, including CGI events, are written to the Application event log. |

For more information, see Configuring Worker Process Identities and IIS and Built-in Accounts.

# Requests for dynamic content return 404 error

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

To help minimize the attack surface of the server, IIS 6.0 is not installed on Windows Server 2003 by default. When you first install IIS 6.0, it is locked down — which means that only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and WebDAV publishing. If you do not enable this functionality after installing IIS, by default on this denial, IIS returns a generic 404 custom error page to prevent disclosure of configuration information. IIS also writes the 404 error with the substatus code of 2 (404.2) in the W3C Extended log files, by default.

**Procedures**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /User:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To enable or disable a Web service extension**

1. In IIS Manager, expand the local computer, and then click **Web Service Extensions**.
2. In the details pane, click the Web service extension that you want to enable or disable, and then do one of the following:
   - To enable a disabled Web service extension, click **Allow**, and then click **OK**.
   - To disable an enabled Web service extension, click **Prohibit**, and then click **OK**.

To enable or disable Web service extensions programmatically, see WebSvcExtRestrictionList Metabase Property.

# All requests with /bin in the URL are rejected and return a 404 error

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This occurs when IIS 6.0 and ASP.NET are both installed. In order to take a more proactive stance against malicious users and attackers, the ASP.NET ISAPI filter, aspnet_filter.dll, blocks incoming request containing /bin in the URL. This behavior occurs server-wide, regardless whether the request is for static or dynamic content.

The preferred solution to this issue is to modify the path to content on the server so that /bin is not necessary in any request.

If the content URL cannot be modified, an alternative solution is to set a registry key that stops the ASP .NET ISAPI filter from filtering requests containing /bin in the URL. This is a server-wide setting.

**Procedures**

| **Important** |
| --- |
| Setting the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\ASP.NET\ registry key can allow a malicious user access to programs and content in the /bin directory. |

**To disable /bin filtering**

1. Start Registry Editor and navigate to the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\ASP.NET\ key.
2. In the details pane, right-click, point to **New**, and click **DWORD Value**.
3. In the **Name** box, type the following: **StopBinFiltering**.
4. Double-click the **StopBinFiltering** value, and in the Value data box type **1**.
5. Click **OK**, and then close Registry Editor.
6. To reenable /bin filtering, set the **StopBinFiltering** value to 0.

# Requests for static files return 404 error

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

For requests to static content, this version of IIS serves requests for files with known file name extensions only, a feature called Known Extensions. If a request is made for a resource whose file name extension is not mapped to a known extension in the MimeMap Metabase Property, IIS denies the request and logs a 404 error with the substatus code of 3 (404.3) in the W3C Extended log files (by default). To prevent disclosure of configuration information, IIS is configured to return the generic 404 custom error page by default on this denial. To resolve this error, use IIS manager to add the file name extension to the Multipurpose Internet Mail Exchange (MIME) map on a specific directory or Web site.

If adding the specific file name extension to the MIME map does not remedy the problem, turn off the Known Extensions feature on a specific directory or a Web site. To turn off the Known Extensions feature and thereby allow IIS to serve files with any extension for a specific site or directory, add the ***,application/octet-stream** value to the list of MIME maps (enter the asterisk in the **Extension** text box and **application/octet-stream** in the **MIME Type** text box on the **MIME Type** dialog box). For more information about how to add this value to the MIME map, see Working with MIME Types.

If you are still experiencing problems, you can turn off the Known Extensions feature on the Web server and allow IIS to serve files with any extension for all sites and directories. Turn off the Known Extensions feature by adding the ***,application/octet-stream** to the global MIME map. You must wait until the worker process has recycled or restart the World Wide Web Publishing Service (WWW service) before changes take effect.

| **Important** |
| --- |
| Turning off Known Extensions on the Web server means IIS will serve files with any extension. Malicious users could potentially exploit this configuration to download sensitive information from your Web server. An IIS server that serves files with any extension is not secure. |

| **Note** |
| --- |
| Tools like UrlScan can be configured to block the processing of certain file name extensions. Check your UrlScan settings. |

For more information about adding or editing the MIME map, see Working with MIME Types.

# Worker process recycling drops application session state

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, Worker Processes recycle after a preconfigured amount of time. If your ASP applications are not designed to store session state while a worker process is recycled, then session state in that ASP application can be lost. To remedy this problem, you can either store session state in a database or disable worker process recycling.

**Procedures**

**To disable worker process recycling**

1. In IIS Manager, expand the local computer, expand **Application Pools**, right-click the application pool, and then click **Properties**.
2. On the **Recycling** tab, clear the **Recycle worker processes (in minutes)** check box.
3. Click **OK**.

# Server-side include directives (#include) return 404 error (for .stm files) or 0131 error (for .asp files)

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If your ASP page uses the **#include** server-side include directive and the ".." notation to refer to a parent directory, the directive will return an error unless you have reconfigured the AspEnableParentPaths Metabase Property. This property is set to **false** by default.

| **Important** |
| --- |
| If you set the **AspEnableParentPaths** property to **true**, a potential security risk is introduced because an include path might access critical or confidential files outside of the application root directory. |

**Procedures**

**To enable parent paths through IIS Manager**

1. In IIS Manager, expand the local computer, right-click the starting-point directory of the application you want to configure, and then click **Properties**.
2. Click the **Directory** tab, and then click **Configuration**.
3. Click the **Options** tab.
4. In the **Application configuration** section, select the **Enable parent paths** check box.
5. Click **OK**.

# ASP generates Permission Denied errors in event log for global.asa

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Earlier versions of ASP executed events in the security context (or user identity) of the host process because there is no user context during these events. This caused problems, such as access denied errors when writing to a file in the Session_OnEnd event. ASP, by default, now runs the global.asa events, Application_OnEnd and Session_OnEnd, anonymously (the default value is **true**).

To change this setting programmatically, see AspRunOnEndAnonymously Metabase Property.

# CGI processes do not start

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If your CGI processes do not run, ensure that the CGI Web service extension has been enabled. See Requests for dynamic content return 404 error. Also, CGIs will not start unless the account on which the CGI processes run are assigned certain user rights. You can add the account as a member of the IIS_WPG group and assign the account the following two user rights:

- Adjust memory quotas for a process.
- Replace a process level token.

If you use anonymous authentication and you call external programs, including PHP programs, from CGI code, the operation may fail. The failure occurs because in IIS 6.0 the default anonymous user account, IUSR_*Computername*, is explicitly denied access to all executable programs in the %windir%\System32 folder for security reasons. To resolve this problem, remove the IUSR denial access control list (ACL) on the appropriate command-line tools that reside in the %windir%\System32 folder, and grant the IUSR_*Computername* account **Read & Execute** access.

**Procedures**

**To assign user rights to an account on the local computer**

1. From the **Start** menu, point to **Administrative Tools**, and then click **Local Security Policy**.
2. Expand **Security Settings**, double-click **Local Policies**, and then double-click **User Rights Assignment**.
3. In the details pane, double-click the policy you want to change.
4. Click **Add User or Group**.
5. In the **Enter the object names to select** box, type the user or group name.
6. Click **OK**.

# ASP.NET pages are returned as static files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you installed IIS 6.0 without installing ASP.NET, ASP.NET files can be returned as static files. This error can also occur if you reinstalled IIS 6.0 without reregistering ASP.NET. To learn how to remedy this error, see ASP.NET IIS Registration Tool and use the **i** option.

# Collaboration Data Objects for Windows NT Server fail

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Collaboration Data Objects for Microsoft Windows NT Server (CDONTS) has been removed from the Microsoft Windows Server-TM- 2003 family. If your Web applications use CDONTS, you can convert them to Microsoft Collaboration Data Objects (CDO). Most methods in CDONTS have matching methods in CDO, but might be named differently.

For reference material for CDO in the Platform Software Developer Kit (PSDK), see Overview of CDO at MSDN Online.

# Connection Errors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section includes troubleshooting information about connection errors.

**In This Section**


- Client requests receive a "503: Service Unavailable" error
- Anonymous accounts (IUSR_computername) attempting subauthentication logon receive 401 error
- Clients cannot connect to server
- UNC connections are denied access
- Access denied to console applications in System32 directory
- Client requests time out or error out
- Client cannot renegotiate request and returns an HTTP 413 error
- Clients receive 403.16 when logging on to a Web server over an SSL connection
- Clients browsing to a Web site receive HTTP 403.4 - Forbidden: SSL is required to view this resource error, but the Web site is not configured to use SSL
- UPN logons not accepted by IIS 6.0 server in Windows 2000 domain

# Client requests receive a "503: Service Unavailable" error

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The 503: Service Unavailable error is a noncustomizable error that means the server is currently unable to process the request. There are many reasons why the server might not be able to process a request. For example, the administrator might have shut down an application pool to perform maintenance. Another possible cause of the 503 error is that the application pool queue was already full when the request arrived.

If the error was generated because the application pool queue was full at the time of the request, you will see error code 503 along with the message "QueueFull" when viewing the Httperr.log file. The Httperr.log file is located in the **%Systemroot%**\System32\LogFiles directory.

To remedy this problem, increase the application pool queue length limit as detailed in Configuring Application Pool Queue-Length Limits.

# Anonymous accounts (IUSR_computername) attempting subauthentication logon receive 401 error

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The subauthentication component, Iissuba.dll, is not enabled by default in IIS 6.0. In earlier versions, Iissuba.dll allowed IIS to manage passwords on anonymous accounts, which created a potential security risk. In IIS 6.0, you can use subauthentication to manage passwords for anonymous accounts by meeting the following requirements:

- For applications which you assign anonymous access, the worker process runs as LocalSystem.
- The subauthentication component, Iissuba.dll, is registered.
- The AnonymousPasswordSync Metabase Property is enabled (set to **true**).

The actions taken to meet the above requirements differ between clean installs of IIS 6.0 and upgrades to IIS 6.0 from installations of IIS with subauthentication configured.

For information about the procedures to configure sub-authentication, see Anonymous Authentication in IIS 6.0.

# Clients cannot connect to server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The Windows Server 2003 family provides a software-based firewall to prevent unauthorized connections to your server from remote computers. The Internet Connection Firewall (ICF) is disabled by default. However, if you have enabled the firewall in its default configuration after installing a member of the Windows Server 2003 family and before installing IIS, clients will not be able to connect to your server. The following procedure configures ICF to allow clients to initiate Web and other IIS-related connections to your server.

**Procedures**

**To configure Internet Connection Firewall for IIS**

1. From the **Start** menu, click **Control Panel**.
2. Double-click **Network Connections**.
3. Right-click **Local Area Connection**, and click **Properties**.
4. Click the **Advanced** tab.
5. If you do not want to use the ICF, make sure the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box isn't available, and click **OK**.
6. If you do want to use the ICF, make sure the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box is enabled, and click **Settings**.
7. On the **Services** tab, enable a service to which you want to allow access to clients.
8. In the **Service Settings** dialog box that appears after enabling a service, do one of the following:
   - If you are enabling a service on the same computer you are working, the correct computer name is already filled in. Click **OK**.
   - If you are enabling a service on a different computer on your network, type the name or IP address of the computer hosting the service you are enabling, and click **OK**.
9. Repeat steps 7 and 8 until all the services you want accessible to clients are enabled.

# UNC connections are denied access

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The **Universal Naming Convention (UNC) authentication** method, also called **UNC Passthrough authentication**, determines the credentials to be used for gaining access to a UNC share on a remote computer. Beginning with IIS 6.0, the UNC authentication looks at the request user and the credentials stored in the UNCUserName Metabase Property and UNCPassword Metabase Property properties of the metabase to determine the credentials to pass through to the computer with the UNC share, in the following way:

1. **If UNCUserName** is specified (not empty) and **UNCPassword** is valid, the metabase user credentials are sent as the user identity for access to the remote share. If **UNCUserName** is specified (not empty) and **UNCPassword** is not valid, a 500 Internal Server Error: Invalid Username or Password message is sent to the client.

2. If **UNCUserName** is empty, then the credentials of the request user (either an authenticated set of credentials for authenticated requests or IUSR_*computername* credentials for anonymous requests) are sent as the user identity for access to the remote share.

| **Note** |
| --- |
| The **UNCAuthenticationPassthrough** metabase key is no longer used for UNC authentication. |

# Access denied to console applications in System32 directory

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Requests that use console applications in the Windows System32 directory, such as Cmd.exe, are denied access unless the remote user making the request is an authenticated member of the Administrators group. The denial is the result of special access control lists (ACLs) on all console application programs in the Windows System32 directory to restrict access to only administrators, LocalSystem, interactive users, and services. The ACL restriction does not affect local logon users that should have access, nor does the restriction affect your own custom CGI executable programs.

# Client requests time out or error out

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In IIS 6.0, settings are set to aggressive and secure defaults to minimize attacks due to time-outs and limits that were previously too generous. IIS enforces the following time-outs at the connection level:

- Limits on Response Buffering: The default value for the AspBufferingLimit Metabase Property is 4 MB. If ASP scripts buffer more than this, they error-out. There was no limit to buffering prior to IIS 6.0.

- Limits on posts: The AspMaxRequestEntityAllowed Metabase Property enforces a maximum ASP post size of 204,800 bytes, with each individual field limited to 100 KB. There was no limit to posts prior to IIS 6.0.

- The **ServerListenTimeout** metabase property is not used used by the WWW service in IIS 6.0, but remains available for the FTP, SMTP, and NNTP services. For the WWW service, **ServerListenTimeout** has been replaced by the following metabase properties:

    - ConnectionTimeout Metabase Property: This property specifies the amount of time, in seconds, the server waits before disconnecting an inactive connection.

    - MinFileBytesPerSec Metabase Property: When IIS responds to a client request, the **MinFileBytesPerSec** property determines the length of time the client has to receive the entire response. If the client machine takes too long to receive the entire response, the kernel-mode driver, HTTP.sys, terminates the connection according to the time-out value.

    - HeaderWaitTimeout Metabase Property: When a client connects to the Web server, the client computer is given a time limit to send in all headers for the request (demarked by a final double \r\n). If the complete header set for the request is not received within the time period indicated by **HeaderWaitTimeout**, HTTP.sys resets the connection. You can configure the value of **HeaderWaitTimeout**.

- Header size limitation: By default, HTTP.sys only accepts requests where the request header is less than 16 KB. This means that if HTTP.sys does not receive the terminating <CRLF><CRLF> sequence within 16 KB, HTTP.sys considers the request malicious and terminates the connection. You can change the header size limitation by adjusting the value in the MaxRequestBytes registry key.

# Client cannot renegotiate request and returns an HTTP 413 error

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If client certificates are enabled on a Web server, Web site, or on individual directories or files on the site, clients might see an HTTP 413 error when uploading large files.

If a client sends a long HTTP request, for example, a POST request, to a Web server running IIS 6.0, the IIS worker process might receive enough data to parse request headers, but not receive the entire request entity body. When the IIS worker process detects that client certificates are required to return data to the client, IIS attempts to renegotiate the client connection. However, the client cannot renegotiate the connection because it is waiting to send the remaining request data to IIS.

If client renegotiation is requested, the request entity body must be preloaded using SSL preload. SSL preload will use the value of the **UploadReadAheadSize** metabase property, which is used for ISAPI extensions. However, if **UploadReadAheadSize** is smaller than the content length, an HTTP 413 error is returned, and the connection is closed to prevent deadlock. (Deadlock occurs because a client is waiting to complete sending a request entity, while the server is waiting for renegotiation to complete, but renegotiation requires that the client to be able to send data, which it cannot do).

The solution is to ensure that client is not blocked from sending the entire entity body. To do so, change the value of **UploadReadAheadSize** to a value larger than the content length.

The following example shows how to set the value for **UploadReadAheadSize** to 64 KB on the Web server.

```
cscript adsutil.vbs set w3svc/1/uploadreadaheadsize 65536
```

**Note:** Setting the UploadReadAheadSize property above 64KB is not recommended for web sites that do not require client certificate authentication, otherwise you run the risk of allowing an anonymous Denial of Service (DOS) attack. When UploadReadAheadSize needs to be set at a value higher than 64KB, it is recommended that you set the SSLAlwaysNegoClientCert property to True, which will require that the client has a valid client certificate installed.

For more information about **UploadReadAheadSize**, see UploadReadAheadSize Metabase Property.

For more information see SSLAlwaysNegoClientCert Metabase Property.

# Clients receive 403.16 when logging on to a Web server over an SSL connection

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The VeriSign Global Server ID Intermediate Root certification authority (128-bit SSL) expired on January 7, 2004. Servers running IIS that have not been updated with the new Global Server ID Intermediate Root certification authority may encounter problems when clients try to establish SSL sessions by using the Secure Hypertext Transfer Protocol after January 7, 2004. Clients may see the following error message:

```
403.16 -- Client certificate is untrusted or invalid.

The client certificate has expired or is not yet valid.
```

| **Note** |
| --- |
| Servers that use VeriSign 40-bit Secure Server ID certificates are not affected by this expiration. |

**Procedures**

**To resolve this issue, take the following steps:**

1. Determine the version of the Intermediate Root CA that is currently active on your Web server. For more information, see Determining the Intermediate Root CA Version on a Web Server.
2. If necessary, obtain the current Intermediate CA certificate from the VeriSign Web site. On the VeriSign Web site, search for the Microsoft IIS 5.0 heading, and then click **Get Intermediate CA Here (If Required)** in Step 1.
3. Follow the instructions at the VeriSign Web site to perform the following tasks:
   a. Remove the expired Intermediate Root CA from your Web server.
   b. Install the current Intermediate Root CA on your Web server.
   c. Reboot your Web server and test the SSL connection.

# Client receives 403.16 error when IIS cannot process a complete certificate chain

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When IIS receives an incoming Secure Sockets Layer (SSL) request from a client, it attempts to build its certificate chain before sending its certificate information back to the client. Beginning with IIS 6.0, IIS does not automatically download intermediate certificates by default, which was the behavior in previous versions of IIS. Instead, IIS 6.0 fails the Certificate Revocation List (CRL) check for the certificate because it cannot find the HTTP location for the CRL Distribution Point (CDP). The user then sees an HTTP 403.16 error message.

To rectify this error, add the certificate to the Intermediate Certification Authorities folder on the local computer.

**To add an intermediate certificate to Intermediate Certification Authorities**

1. Download the intermediate certificate to the desktop of the computer where you want to install the certificate.
2. Right-click the certificate, click **Install Certificate** to open the Certificate Import wizard, and then click **Next**.
3. On the **File to Import** page, click **Next**.
4. On the **Password** page, type a password, if available, in the **Password** field, and then click **Next**.
5. On the **Certificate Store** page, click **Place all certificates in the following store** and then click **Browse**.
6. On the **Select Certificate Store** dialog box, check **Show physical stores**.
7. Expand the **Intermediate Certification Authorities** folder, select the **Local Computer** folder, and then click **OK**.
8. Click **Next** and then click **Finish**.

# Clients browsing to a Web site receive HTTP 403.4 - Forbidden: SSL is required to view this resource error, but the Web site is not configured to use SSL

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

It is possible that the Web site configuration has been altered to require SSL, although no certificate has been associated with the Web site. This might happen because the **Edit** button on the **Directory Security** tab is available although no certificate is associated with the Web site. The **Edit** button allows changes to be made to the site's SSL configuration even if a certificate has been removed.

To correct this problem, change the Web site configuration to not require SSL.

**Procedures**

**To remove SSL from a Web site by using IIS Manager**

1. In IIS Manager, right-click the Web site whose SSL configuration you want to change, and then click **Properties**.
2. On the **Directory Security** tab, in the **Secure Communications** area, click **Edit**.
3. Clear the **Require Secure Channel (SSL)** check box, and then click **OK**.

**To remove SSL from a Web site by using adsutil.vbs**

1. Click **Start**, click **Run**, type **cmd**, and then click **OK**.
2. Change to the inetpub\adminscripts directory.
3. At the command prompt, type the following command for the Web site whose SSL configuration you want to change, substituting # with the Web site identifier:

   **cscript adsutil.vbs set w3svc/** # **/AccessSSL False**

# UPN logons not accepted by IIS 6.0 server in Windows 2000 domain

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Although your IIS 6.0 server is a member of a Windows 2000 domain and is correctly configured for Digest authentication as well as subauthentication, logging on with a user principal name (UPN) in the form *username@domain.com* fails. This failure occurs because Digest authentication accepts domain logons in two formats only. If a user logs on to a server in the same domain where the user's account resides, *username* alone is sufficient for authentication. For logons with a user account from another domain, use the *domain**username* format.

# FTP server terminates user connections and returns MSFTPSVC event message 8

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you host a high-load FTP site with the FTP anonymous home directory on an UNC share, you may encounter the following event message in the Event Log: "FTP Server could not create a client worker thread for user at host *ip_address*. The connection to this user is terminated. The data is the error."

This happens when a large number of clients (approximately 1000) attempt to access the same UNC path at the same time and all of the redirector connections are in use. To remedy this problem, enable caching of the anonymous home directory with the registry key CacheAnonymousHomeDirInfo, and set the related timeout values in the registry keys AnonymousHomeDirInfoInterval and AnonymousHomeDirInfoGrace.

**Related Information**

- For more information about the registry keys CacheAnonymousHomeDirInfo, AnonymousHomeDirInfoInterval, and AnonymousHomeDirInfoGrace, see FTP Service Registry Entries.

- For more information about MSFTPSVC event messages, see FTP Service-Related Events.

# Miscellaneous Errors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This section includes troubleshooting information about miscellaneous errors.

**In This Section**


- File requests to UNIX or Linux server return wrong file or error
- Cannot locate /Scripts or /Msadc directory
- ISAPI filter does not show up as "loaded" in UI
- During an upgrade to Windows Advanced Server, Limited Edition or to Windows Server 2003, Datacenter Edition, clustered IIS resources fail
- When attempting to download a file from an FTP site that is hosted on IIS 6.0, a "File not found" or similar error is returned, and the file name in the error message contains a caret (^) character
- Cannot make immediate request for a Web server certificate
- When adding a new Web site, the Web site is not added, and the error message "No more memory is available for security information updates" is displayed (also, error 1131 might appear in the system Event Log)
- Remote Desktop Connection shortcut does not appear on START menu and Remote tab does not appear on System Properties after upgrade from Windows 2000 with IIS
- Warnings received on upgrade
- Client certificate warnings on upgrade
- Install certificate after deleting the pending certificate request
- User receives ERROR_BAD_EXE_FORMAT on 64-bit Windows

# File requests to UNIX or Linux server return wrong file or error

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If IIS must access files on a UNIX or Linux system, file name case sensitivity can be an issue unless Network File System (NFS) support is enabled in IIS.

UNIX and Linux both support mixed-case file names, and IIS fully supports requesting static files in a case-sensitive manner. An issue arises, however, when IIS makes a subsequent request for a file from its static file cache. Because all file names are converted to uppercase letters in the IIS cache, any request after the first request from the IIS static file cache might fail or return the wrong file.

To remedy this problem, disable the IIS static file cache so that all file requests are issued fresh, thereby retaining correct file name case. The static file cache can be disabled for individual virtual directories on Web sites, or globally for all sites.

| **Note** |
| --- |
| Changing this setting has no effect on how ASP files and templates are cached. |

**Procedures**

**To disable the static file cache for a particular Web site virtual directory**

- Edit the metabase and set the **MD_VR_NO_CACHE** property to **1**.

**To disable the static file cache for all sites**

- Edit the registry and add the binary value **DisableStaticFileCache**=**1** to the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters key.

# Cannot locate /Scripts or /Msadc directory

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

By default, the **/scripts** and **/msadc** directories in IIS 5.0 allowed scripts and executables to run. These directories were removed from IIS 6.0 because, if a malicious user gained access to one of these directories, the user could run a script or executable and potentially gain control of the Web server. If your server configuration requires a directory of this nature, you will need to create the directory and establish the appropriate NTFS Permissions.

# ISAPI filter does not show up as "loaded" in UI

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In an effort to optimize resources in IIS 6.0, an ISAPI filter is not loaded until a request is made to a Web site that requires the ISAPI filter. Until this request is made, IIS Manager does not display the status of the ISAPI filter. Also, if the ISAPI filter requires the SF_NOTIFY_READ_RAW_DATA filter notification, the filter will not load while IIS is running in worker process isolation mode. Review the Application Event Log for events from W3SVC-WP to verify that the filter did not load. To remedy this problem, run IIS in IIS 5.0 isolation mode or contact the vendor of the ISAPI filter for compatibility updates.

| **Important** |
| --- |
| If the ISAPI filter is restricted by an access control list (ACL) in such a way that the IIS worker process identity cannot load it, requests receive a 503 error. To remedy this, set the ACLs on the ISAPI filter DLL to allow access for the IIS_WPG group. |

**Procedures**

**To configure IIS for IIS 5.0 isolation mode**

1. In IIS Manager, expand the local computer, right-click **Web Sites**, and then click **Properties**.
2. Click the **Service** tab, select the **Run WWW service in IIS 5.0 isolation mode** check box, and then click **OK**.
3. To start the WWW service, click **Yes**.

# During an upgrade to Windows Advanced Server, Limited Edition or to Windows Server 2003, Datacenter Edition, clustered IIS resources fail

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS 6.0 is not supported during rolling upgrades of server clusters. To upgrade a clustered IIS resource, you must replace the existing IIS resource with a Generic Script Resource. For instructions on how to do this, search for "generic script resource" in Help and Support Center for Windows Server 2003.

# When attempting to download a file from an FTP site that is hosted on IIS 6.0, a "File not found" or similar error is returned, and the file name in the error message contains a caret (^) character

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This error occurs when the requested file contains one or more Unicode characters in its file name. For security reasons, certain characters are disallowed in file names. When Unicode characters are used, they are converted to the closest ANSI equivalent that is available in the code page of the server operating system. Unicode characters that do not have exact ANSI equivalents might be replaced with a disallowed ANSI character. To prevent a disallowed character from appearing in the file name, it is converted to a caret (^).

To avoid this issue, do not use Unicode characters in file names that are used on an FTP site. If Unicode characters must be permitted, only use those that have a direct equivalent (one-to-one mapping) to the server code page.

# Cannot make immediate request for a Web server certificate

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

After installing Certificate Services on a server and configuring it to be either an Enterprise Root certification authority (CA) or a standalone certification authority on a server, you cannot immediately request a Web server certificate by using the Web Server Certificate Wizard on that same server. The **Send the request immediately to an online certification authority** option is grayed and cannot be selected.

It takes a few hours for IIS to recognize that a new CA has come onine. After IIS recognizes the new CA, re-run the Web Server Certificate Wizard to generate the certificate request. The **Send the request immediately to an online certification authority** option will then be available.

If you must have a certificate immediately, run the Web Server Certificate Wizard and select **Prepare the request now, but send it later**. The wizard will generate a certificate request file that can be processed immediately by the new CA.

# When adding a new Web site, the Web site is not added, and the error message "No more memory is available for security information updates" is displayed (also, error 1131 might appear in the system Event Log)

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This problem occurs because insufficient non-paged pool memory has not been allocated for HTTP.sys. By default, IIS automatically determines the number of Web sites that can be created based on the amount of physical memory available. For example, on a computer that has 64 megabytes (MB), the non-paged pool memory value is set to 16. On a computer that has 256 MB, the non-paged pool memory value is set to 64.

To resolve this issue, create a MaxEndPoints registry value, and set it to a number that is at least 1 greater than the number of Web sites you want to host with IIS. To do so, follow these steps:

1. Start Registry Editor and navigate to the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\HTTP\Parameters key.

2. In the details pane, right-click, point to **New**, and and select **DWORD Value**.

3. In the **Name** box, type the following: **MaxEndPoints**

4. Double-click the MaxEndPoints value, and in the **Value data** box type a number that is at least 1 greater than the number of Web sites that you want to host, but no larger than 1024, then click **OK**, and close Registry Editor.

# Remote Desktop Connection shortcut does not appear on START menu and Remote tab does not appear on System Properties after upgrade from Windows 2000 with IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Although the Start menu shortcut and the Remote tab have not been installed, the Remote Desktop Connection application itself is installed. The user interface elements may be hidden if the permissions on the \Inetpub directory or the \Inetpub\Adminscripts directory deny full control to the Local System account.

There are several ways to address this issue.

- Work around the issue by running Remote Desktop Connection from the command line.
- Restore the Remote tab on the System Properties page.
- Restore the Start menu shortcut.

**Procedures**

**To run Remote Desktop Connection from the command line**

1. Click **Start** and then click **Run**.
2. In the **Open** box, type

```
mstsc.exe
```

and then click **OK**.

**To restore the Remote tab on the System Properties page**

1. In Windows Explorer, right-click %windir%\System32\remotepg.dll, and then click **Open With**.
2. Click **Open With**, click **Select the program from a list**, and then click **OK**.
3. Click **Browse**, navigate to %windir%\System32\regsvr32.exe, and select this file.
4. Click **Open**, click **OK**, and then click **OK** again.

**To restore the Start menu shortcut for Remote Desktop Connection**

1. In Windows Explorer, right-click \Inetpub\Adminscripts and then click **Sharing and Security**.
2. On the **Security** tab, click **Add**, type

```
System
```

and then click **OK**.

1. Grant the account full control, and then click **OK**.
2. Insert the Windows Server 2003 CD, click **Perform Additional Tasks**, and then click **Set up Remote Desktop Connection**.
3. Accept the license agreement, click **Next**, click **Next** again, and then click **Install**.
4. Click **Finish** when installation is complete, and then click to exit the CD.

# Warnings received on upgrade

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you upgrade a server with IIS installed to Windows Server 2003 with IIS 6.0, you may receive the following warning messages:

The following items are not compatible with Windows. If you continue, you may not be able to use these items, even after Setup is complete.

- IIS World Wide Web Publishing Service (W3SVC) will be disabled during upgrade.
- IIS is being installed on a FAT volume. This will disable IIS security features.

The first message occurs when you upgrade a server without previously running the IIS Lockdown Tool or configuring the RetainW3SVCStatus registry entry. By default, because of increased security measures, the World Wide Web Publishing Service (WWW service) is not enabled after you upgrade from the Windows 2000 Server family with IIS 5.0 unless you take one of the following actions on a manual or unattended upgrade:

- Run the IIS Lockdown Tool on the computer that is running Windows 2000 Server before you start the upgrade process. The IIS Lockdown Tool reduces the attack surface of your server by disabling unnecessary features, and the tool provides the option to enable and customize features for your site. For more information about the IIS Lockdown Tool, see IIS Lockdown Tool 2.1.

  | **Important** |
  | --- |
  | If you do not run the IIS Lockdown Tool, after the upgrade is complete, you must ensure that all unnecessary IIS features have been removed or disabled. You must also ensure that the enabled features are configured with the highest security settings that your organization can support. |

- Add the RetainW3SVCStatus registry entry to the registry prior to upgrading.
- Add the ?DisableWebServiceOnUpgrade = False? entry to the [InternetServer] section of an unattended installation script that you plan to run for the upgrade.
- Enable the WWW service after the upgrade is completed.

  | **Caution** |
  | --- |
  | Incorrectly editing the registry might severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**Procedures**

**To use regedit.exe to add the RetainW3SVCStatus registry entry to the registry**

1. Click **Start**, click **Run**, type

```
regedit
```

in the **Run** box, and then click **OK**.

1. Navigate to the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC registry entry.
2. Click to select the registry entry, point to **New** on the **Edit** menu, and then click **Key**.
3. Type

```
RetainW3SVCStatus
```

and then click **Enter**.

1. On the **Edit** menu, point to **New**, and then click **DWORD Value**.

2. Type

```
do_not_disable
```

and then click **Enter**.

1. On the **Edit** menu, click **Modify**.

2. Type

```
1
```

and then click **Enter**.

**To enable the WWW service after upgrading**

1. Click **Start**, click **Administrative Tools**, and then click **Services**.
2. Right-click **World Wide Web Publishing** and then click **Properties**.
3. Click the **General** tab, set **Startup type** to Automatic, and then click **Apply**.
4. Click **Start** to start the service, and then click **OK** after the service has started.

The second message occurs when the partition on which you are installing IIS 6.0 is formatted as FAT. Installing IIS on a FAT volume disables Windows Server 2003 security features that require the NTFS file system. For example, NTFS supports Active Directory, which is necessary for domains, user accounts, and other important security features. You should convert the partition from FAT to NTFS before continuing the upgrade.

# Client certificate warnings on upgrade

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you have a Web site on Internet Information Services (IIS) 5.0 that requires client certificates, when you upgrade the server to Microsoft Windows Server 2003 with IIS 6.0, clients that connect to the site may receive one of the following error messages even if the client certificates are not controlled by a certificate trust list (CTL):

```
HTTP 403.16 Forbidden: Client certificate untrusted or invalid.

-or-

HTTP 403.16 Forbidden: Client certificate is ill-formed or is not

trusted by the web server.

-or-

HTTP Error 403.7: Forbidden: SSL client certificate is required.
```

In addition, when the client accesses the Web site, one of two things may happen:

- The client may not receive the Client Authentication dialog box in the browser, and therefore cannot choose a client certificate that allows access to the Web site.
- If the client receives the Client Authentication dialog box, the certificate list in the **Client Authentication** dialog box may not display the required client certificate.

These issues may occur if the client certificate was created by a certification authority that the IIS computer does not trust. In IIS 5.0, you could specify a CTL that contained certification authorities whose root certification authority (CA) certificates were installed in the personal certificate store on the local computer. However, in IIS 6.0, the root CA certificates must be installed in the Trusted Root Certification Authorities certificate store on the local computer. This is because IIS 6.0 verifies certificates based on the rules specified in the CryptoAPI. The CryptoAPI rejects certificates if the root CA certificates are not installed in the Trusted Root Certification Authorities certificate store on the local computer.

To resolve the error and display the certificate in the browser, you must install the root CA certificate in the Trusted Root Certification Authorities certificate store on the local computer.

**Procedures**

**To install the root CA certificate in the Trusted Root Certification Authorities certificate store on the local computer**

1. Add a Certificates snap-in for the local computer. For more information about adding a Certificates snap-in, see "To manage certificates for a computer" in Windows Server 2003 Help and Support Center.
2. Export the certificate from the Personal Certificate store on the local computer. For more information about exporting a certificate, see "To export a certificate" in Windows Server 2003 Help and Support Center.
3. Import the certificate to the Trusted Root Certification Authorities certificate store on the local computer. For more information about importing a certificate, see "To import a certificate" in Windows Server 2003 Help and Support Center.

# Install certificate after deleting the pending certificate request

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS stores the private key for a certificate as the pending request. Deleting the pending request deletes the association of the private key with IIS, but the private key still exists in the certificate store. To install the certificate without having the pending request available, you can use version 5.2.3718.0 of the Certutil.exe command-line tool that is available through the Certificate Services MMC snap-in in Windows Server 2003. For more information about Certutil.exe, see Certutil.exe.

**Procedures**

**To install a Web server certificate that lacks a pending certificate request**

1. Click **Start**, point to **Run**, type **cmd**, and then click **OK**.
2. Navigate to the directory where Certutil.exe is stored; by default, this is %windir%\system32.
3. Type the following command at the command prompt: **certutil -addstore my** *certnew.cer*

   where *certnew.cer* is the name of the certificate you received from the certification authority (CA). You should see the following message: **CertUtil: -addstore command completed successfully**.
4. Navigate to the directory where you stored the certificate you received from the CA. Right-click the certificate and then point to **Properties**.
5. Click the **Details** tab and select <All> in the **Show** drop-down list.
6. In the **Field** list, select **Thumbprint** to display its value in the view pane.
7. Select the **Thumbprint** value in the view pane and then click CTRL+C.
8. Return to the command prompt window and type the following command: **certutil -repairstore my "***thumbprint***"**

   where *thumbprint* is the value of the **Thumbprint** field. Be sure to type the double quotes as part of the command. If the command is successful, the following message is displayed: "Encryption test passed CertUtil: = repairstore command completed successfully."
9. Install the server certificate on your Web server.

**Important**

If the certutil command does not complete successfully, the following error message is displayed: "Certutil: -repairstore command FAILED: 0x80090011 (-2146893807) Certutil: Object was not found." This message indicates that the private key for the certificate does not exist in the certificate store. You cannot install the certificate you obtained from the CA. Instead, you must generate a new certificate request, obtain the new certificate, and install that new certificate on your Web server.

**See Also**

**Concepts**

Installing Server Certificates

# Cannot launch COM objects with IUSR_computername or IWAM_computername accounts

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

For Windows Server 2003, the IUSR_*computername* and IWAM_*computername* accounts are created and granted launch permissions to DCOM components by default when installing IIS. For Windows Server 2003 Service Pack 1 (SP1), the accounts are still created, but are no longer granted launch permissions to DCOM components by default when installing IIS.

If you had installed IIS prior to upgrading to SP1, the permissions will not be changed. However, if you install IIS after upgrading to SP1, Component Object Model (COM) objects will fail to launch under the IUSR_*computername* and IWAM_*computername* accounts if your application uses default launch permissions for those accounts. In this situation you will receive an "Access denied" error message, because COM+ is passing credentials that are not valid to Windows Server 2003. To remedy this problem, grant explicit launch permissions to specific user(s) or group(s) to launch COM objects instead of using default launch permissions. You should also ensure that IIS, COM+, and Windows all have the same password for the IWAM_*computername* account. For more information about configuring launch permissions, see Configuring Launch Permissions for COM+ Server Objects.

# Computer runs slowly when scripts using ADSI scripts to query IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

When you run a program that uses multiple threads and Active Directory Service Interfaces (ADSI) scripts to query Internet Information Services (IIS), memory usage on your computer might greatly increase, causing the program to stop responding.

This problem occurs on computers that are running one of the following operating systems:

- Microsoft Windows Server 2003
- Microsoft Windows XP Professional
- Microsoft Windows 2000

When you use multiple threads to query IIS using ADSI, memory is used by each of these threads. This problem occurs because the information in the IIS ADSI provider per-thread cache is not cleared when a thread is terminated. Because this memory is not released, the program might run out of memory.

There is a hotfix available to correct this problem. See Knowledge Base article 823700, FIX: Your computer may run slowly when you use Active Directory Service Interfaces scripts in a program to administer IIS on a computer that is running Windows Server 2003, Windows 2000, or Windows XP, for information about how to obtain the hotfix.

# User receives "The system cannot find the path specified" message when clicking on a Web site in IIS Manager

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

This occurs when a user deletes a physical directory for a Web site with virtual directories from a location other than IIS Manager, such as in Windows Explorer or through the command line. When physical directories are deleted this way, the virtual directories continue to exist in the IIS metabase and to serve content. When the user expands or clicks on the Web site in IIS Manager, they receive the message "The system cannot find the path specified" because the folder containing the content for that Web site was deleted; however, IIS Manager will continue to display the Web site with the deleted physical directory as long as virtual directories still exist and serve content under that Web site. To fix this issue, you must correct the path to the home directory for the Web site.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**. |

**To set the location of the home directory for a Web site by using IIS Manager**

1. In IIS Manager, expand the local computer, expand **Web Sites**, right-click the Web site with the deleted physical directory, and then click **Properties**.
2. On the **Home Directory** tab of the Web site **Properties** dialog, click **Browse**.
3. In the **Browse For Folder** dialog box, locate a folder with content for the Web site or click **Make New Folder** to create a folder where you will put the content, and click **OK**.
4. Click **Apply** and then **OK**.

# User receives ERROR_BAD_EXE_FORMAT on 64-bit Windows

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The most common problems with 32-bit applications on 64-bit Windows involve image mismatch. An image mismatch occurs when a 64-bit process tries to load a 32-bit DLL or a 32-bit process tries to load a 64-bit DLL. Image mismatch errors that cause a worker process to fail during startup are logged in the Windows Events log as 0xC1: ERROR_BAD_EXE_FORMAT. Image mismatch errors that do not cause the worker process to fail during startup are logged in the IIS log files as Win32 error 0xC1: ERROR_BAD_EXE_FORMAT. To remedy these errors, configure IIS to launch worker processes that are the same bitness as your applications. For information about how to configure IIS to run 32-bit applications on 64-bit Windows, see Configuring IIS to Run 32-bit Applications on 64-Bit Windows. For more information about running 32-bit applications on 64-bit Windows, see Running 32-bit Applications on 64-bit Windows.

# Advanced Troubleshooting with IIS Tracing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Event Tracing for Windows (ETW) is a feature of the Windows operating system that allows you to troubleshoot issues with HTTP requests in IIS components like Active Server Pages (ASP), Internet Server Application Programming Interface (ISAPI) extensions, and the Secure Sockets Layer (SSL) Filter service. If an HTTP request in IIS fails or becomes unresponsive while ETW is enabled, you can view ETW trace data, called events, to determine which component caused the failure.

This section includes the following information:

Tracing Requests: Describes the various procedures for tracing different components in IIS.

Tracing Startup and Shutdown Issues in the IIS Admin Service: Describes how to trace the IIS Admin Service when IIS is having trouble starting up or shutting down.

**See Also**
**Other Resources**
Troubleshooting IIS 6.0 with Tracing

# Tracing Requests

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

In Windows Server 2003 (Service Pack 1 or later) IIS includes the following tracing features. These features leverage the Event Tracing for Windows (ETW) infrastructure.

- IIS Currently-Executing Requests Tracing: This tracing feature provides general statistics and details about all requests executing on the server at the moment tracing was started. If the CPU on your server is spiking or if requests become unresponsive, currently-executing requests tracing can help you understand which URL is being requested, which application pool the URL resides in, and other similar details. This feature does not include the option to specify which components or URLs to trace. For more detailed information, see Currently Executing Requests Tracing.

- IIS Request-Based Tracing: This tracing feature tracks HTTP requests in IIS components. Request-based tracing can help you target the IIS component processing a request when the request failed or became unresponsive. Request-based tracing allows you to define either a specific URL or a group of URLs to trace. Request-based tracing events can be viewed in real time, but the feature can return a substantial amount of data (depending on server traffic), so it is much more practical to view the data in a trace log. For more detailed information about request-based tracing, see IIS Request-Based Tracing.

This section includes the following information:

- How to Create a Provider File for Request-Based Tracing: Describes how to create a simple text file that defines which IIS components to trace and how much data tracing should return.

- How to Trace Requests in IIS: Describes how to trace various IIS components, called providers, using IIS request-based tracing.

- How to Trace Requests for a Specific URL or Set of URLs: Describes how to trace requests for a specific URL or set of URLs using request-based tracing.

- How to Trace All Requests Currently Executing in IIS Worker Processes: Describes how to trace requests that are executing inside an IIS worker process using the currently-executing requests tracing command-line feature.

- How to Trace Requests Using the Performance Logs and Alerts MMC: Describes how to use the Performance Logs and Alerts MMC to trace IIS providers on demand or according to a schedule.

- How to Process and View Trace Log Files: Describes how to process and view trace log data using Log Parser 2.2 and the DumpTraceReqs.js sample script.

# How to Create a Provider File for Request-Based Tracing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS components are called **providers** in the context of Event Tracing for Windows (ETW) because they provide trace data. When you start request-based tracing in IIS, you need to tell ETW which provider(s) to trace. You also need to tell ETW how much data each provider should report. The verbosity level is the name of the setting that indicates the amount of data requested. You define provider(s) and verbosity levels in a simple text file called a provider file. When you execute the command to start tracing, you specify the name of the provider file (and path, if necessary). ETW reads the provider file and begins tracing the provider(s) at the verbosity levels defined in that file.

In Windows Server 2003 (Service Pack 1 or later) the following IIS providers report trace data to ETW:

- IIS: WWW Server
- IIS: SSL Filter
- IIS: ISAPI Extension
- IIS: ASP
- IIS: IISAdmin

The IIS: WWW Server provider includes the following flags to enable various functional areas or subcomponents in the WWW server provider. You can define one or more functional areas in your provider file by specifying the flag number or the friendly flag name.

| Functional Area | Flag | Friendly Flag Name |
|---|---|---|
| UrlFilter | 1 | UseURLfilter |
| Authentication | 2 | IISAuthentication |
| Security | 4 | IISSecurity |
| ISAPI Filter | 8 | IISFilter |
| Static File | 10 | IISStaticFile |
| CGI | 20 | IISCGI |
| Compression | 40 | IISCompression |
| Cache | 80 | IISCache |
| WWW Server All | 0xFFFFFFFE | IISAll |

Verbosity levels include the following:

| Level | Name | Description |
|---|---|---|
| 0 | General | IIS providers log only those few events required to provide some context to the request activity. These events include a **General_Request_Start** event that logs the URL and verb for the request and a **General_Request_End** event. |

| | | |
|---|---|---|
| 1 | Fatal Error | IIS providers report events for actions that can cause a process to exit or actions that are about to cause a process to exit. |
| 2 | Error | IIS providers report events when components cannot continue processing a request. The request is immediately rejected.<br>**Note**<br>These events typically indicate a problem on the server side rather than a client-induced problem. |
| 3 | Warning | IIS providers report events when components experience an unexpected result/error but can continue processing the request. |
| 4 | Information | IIS providers report general information about requests. |
| 5 | Verbose | IIS providers report detailed information about requests. |

**To create a provider file**

1. Create a new file in a text editor like Notepad.
2. Enter a provider name in double quotes, flags (if any) or zero (0), and the verbosity level.

   For example:

   "IIS: WWW Server" 0xFFFFFFFE 5

   "IIS: Active Server Pages (ASP)" 0 5

   "IIS: WWW Isapi Extension" 0 5
3. Start a new line for each provider, if you plan to trace more than one provider.
4. Save and close the provider file.

**See Also**
**Concepts**
Tracing Startup and Shutdown Issues in the IIS Admin Service
**Other Resources**
Configuring the IIS Request-Based Tracing Session

# How to Trace Requests in IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you are actively troubleshooting a problem in IIS and you believe the problem exists in a specific component like ASP or the SSL Filter service, you can trace requests in that component to locate the source of the problem. If you are uncertain where the problem exists, you can trace requests in multiple components and view the resulting trace data to locate the source of the problem.

Request-based tracing uses the following commands or switches:

| | |
|---|---|
| **logman start** | Starts the trace session using Logman.exe as the controller. |
| **logman stop** | Stops the trace session. |
| *session name* | Any name you want to give this trace session. |
| **-pf** *provider file name* | Identifies the name of the provider file. This can also be a path to the file. For more information, see How to Create a Provider File for Request-Based Tracing. |
| **-ets** | Tells Logman.exe to start the trace session immediately. |

**To trace requests in IIS components**

1. From a command prompt, type **logman start** *session name* **-pf** *provider file name* **-ets** and press ENTER.

   For example: **logman start ASPtrace -pf ASPproviders -ets**

2. Event Tracing for Windows prints to the screen details about the trace session you just started, including the name of the session, the file name where the trace data will be collected (*session name*.etl by default), the providers listed in the provider file, and whether or not the command was successful.

3. Allow the trace session to run until you have reproduced the problem or until your sites have processed enough requests to produce a manageable data set.

4. From the command prompt, type **logman stop** *session name* **-ets** and press ENTER.

**See Also**
**Tasks**
How to Process and View Trace Log Files
**Other Resources**
Troubleshooting IIS 6.0 with Tracing

# How to Trace Requests for a Specific URL or Set of URLs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

If you are actively troubleshooting a problem and you know the URL where the problem exists, you can trace that specific URL. If you do not know which URL is causing the problem, but you suspect it may be an Active Server Page (ASP) or Microsoft ASP.NET page, you can trace a set of URLs.

Request-based tracing for a specific URL or set of URLs uses the following commands or switches:

| | |
|---|---|
| **logman start** | Starts the trace session using Logman.exe as the controller. |
| **logman stop** | Stops the trace session. |
| *session name* | Any name you want to give this trace session. |
| **-pf** *provider file name* | Identifies the name of the provider file. This can also be a path to the file. For more information, see How to Create a Provider File for Request-Based Tracing. |
| **-ets** | Tells Logman.exe to start the trace session immediately. |

**To trace requests for a specific URL**

1. Configure the TraceUriPrefix Metabase Property for the URL you want to trace.

   a. Open a command prompt and navigate to the %windir%\Inetpub\AdminScripts directory.

   b. Type **adsutil.vbs set w3svc/***n***/TraceUriPrefix=***path to URL* and press ENTER. *N* is the number of the site as it is listed in the metabase, e.g. w3svc/1 is the default Web site (by default). *Path to URL* is the physical path to the file you want to trace, e.g. %windir%\Inetpub\wwwroot\asp\default.asp

2. Set the **UseUrlFilter** flag.

   a. Create a new provider file or open an existing provider file. See How to Create a Provider File for Request-Based Tracing.

   b. Enter one of the following items in your provider file:

      **"IIS: WWW Server" UseUrlFilter 5**

      The **UseUrlFilter** flag enables tracing for the URL specified in the **TraceUriPrefix** metabase property.

      -or-

      **"IIS: WWW Server" 0xFFFFFFFE 5**

      The 0xFFFFFFFE flag enables all the WWW server tracing flags, including the **UseUrlFilter** flag.

3. Save and close the provider file.

4. From a command prompt, type **logman start** *session name* **-pf** *provider file name* **-ets** and press ENTER.

   For example: **logman start AllAspTrace -pf AllAspUrls -ets**

5. Event Tracing for Windows prints to the screen details about the trace session you just started, including the name of the session, the file name where the trace data will be collected (*session name*.etl by default), the providers listed in the

provider file, and whether or not the command was successful.

6. Allow the trace session to run until you have reproduced the problem or until your sites have processed enough requests to produce a manageable data set.

7. From the command prompt, type **logman stop** *session name* **-ets** and press ENTER.

   To view the trace data, see How to Process and View Trace Log Files.

**To trace requests for a set of URLs**

1. Configure the TraceUriPrefix Metabase Property for the URLs you want to trace.

   a. Open a command prompt and navigate to the %windir%\Inetpub\AdminScripts directory.

   b. Type **adsutil.vbs set w3svc/***n***/TraceUriPrefix=***path to URL*extension* and press ENTER. *N* is the number of the site as it is listed in the metabase, e.g. w3svc/1 is the default Web site (by default). *Path to URL*extension* is the physical path to the group of files you want to trace, e.g. %windir%\Inetpub\wwwroot\asp*.asp or %windir%\Inetpub\wwroot\aspnet*.aspx

2. Set the **UseUrlFilter** flag.

   a. Create a new provider file or open an existing provider file. See How to Create a Provider File for Request-Based Tracing.

   b. Enter the ASP provider, if you are going to trace a group of .asp files, for example: "IIS: Active Server Pages (ASP)" 0 5

   c. Type one of the following items in your provider file:

      **"IIS: WWW Server" UseUrlFilter 5**

      The **UseUrlFilter** flag enables tracing for the URL specified in the **TraceUriPrefix** metabase property.

      **-or-**

      **"IIS: WWW Server" 0xFFFFFFFE 5**

      The 0xFFFFFFFE flag enables all the WWW server tracing flags, including the **UseUrlFilter** flag.

3. Save and close the provider file.

4. From a command prompt, type **logman start** *session name* **-pf** *provider file name* **-ets** and press ENTER.

   For example: **logman start AllAspTrace -pf AllAspUrls -ets**

5. Event Tracing for Windows prints to the screen details about the trace session you just started, including the name of the session, the file name where the trace data will be collected (*session name*.etl by default), the providers listed in the provider file, and whether or not the command was successful.

6. Allow the trace session to run until you have reproduced the problem or until your sites have processed enough requests to produce a manageable data set.

7. From the command prompt, type **logman stop** *session name* **-ets** and press ENTER.

   To view the trace data, see How to Process and View Trace Log Files.

# How to Trace All Requests Currently Executing in IIS Worker Processes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use currently-executing requests tracing to find out which requests are executing in a worker process when the worker process has become slow or unresponsive. When enabled, currently-executing requests tracing prompts worker processes to report statistics and details about each request in each worker process.

Currently-executing requests tracing does not require a provider file. Instead, currently-executing requests tracing uses the following commands or switches:

| | |
|---|---|
| **logman start** | Starts the trace session using Logman.exe as the controller. |
| **logman stop** | Stops the trace session. |
| *session name* | Any name you want to give this trace session. |
| **-p "IIS: Request Monitor"** | Identifies which provider (-p) should be traced. |
| **-ets** | Tells Logman.exe to start the trace session immediately. |

**To trace all requests currently executing in IIS worker processes**

1. Open a command window and type **logman start** *session name* **–p "IIS: Request Monitor" -ets** and press ENTER.
2. Event Tracing for Windows prints to the screen details about the trace session you just started, including the name of the session, the file name where the trace data will be collected (*session name*.etl by default), and whether or not the command was successful.
3. Allow the trace session to run until you have reproduced the problem or until your sites have processed enough requests to produce a manageable data set.
4. From the command prompt, type **logman stop** *session name* **-ets** and press ENTER.

   To view the trace data, see How to Process and View Trace Log Files.

# How to Trace Requests Using the Performance Logs and Alerts MMC

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The Performance Logs and Alerts snap-in to the Microsoft Management Console (MMC) allows you to trace IIS providers. You can manually start a trace or you can configure Event Tracing for Windows to trace IIS providers according to a schedule. You cannot, however, configure verbosity levels (the amount of data returned by IIS providers) when using the Performance Logs and Alerts MMC. For more information about verbosity levels, see Verbosity levels.

**To trace IIS providers using the Performance Logs and Alerts MMC**

1. From the Start menu, right-click **My Computer**, and click **Manage**.
2. In the Computer Management MMC, expand the **Performance Logs and Alerts** node, right-click **Trace Logs**, and click **New Log Settings**.
3. Type a name for the new trace session.
4. In the trace session dialog box, click **Nonsystem Providers**.
5. Click the **Add** button.
6. In the **Add Nonsystem Providers** dialog box, select the providers you want to trace. Hold the CTRL key to select multiple providers.
7. Click **OK**.
8. If you want to start the trace immediately, click **OK**. If you want to configure other tracing options, including an automatic tracing schedule or the location and naming of log files, click the other tabs on this dialog box.
9. To stop tracing, right-click the trace session you just created and click **Stop** in the drop-down menu.

   To view the trace data, see How to Process and View Trace Log Files.

# How to Process and View Trace Log Files

Updated: August 22, 2005

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1

Trace data is written in binary format to a trace log (*filename*.etl). To view the trace data, the log file must be processed by a tool like LogParser. LogParser version 2.2 or later includes a sample script called DumpTraceReqs.js. DumpTraceReqs.js was specifically designed to format trace log files for easy viewing. You do not need to be familiar with LogParser to process trace logs.

| **Note** |
| --- |
| The following procedure does not apply to IIS Admin trace logs. To process and view IIS Admin trace logs, see How to Process IIS Admin Trace Log Files Using Trace Report. |

**To process and view trace log files**

1. Download Logparser, version 2.2 or later from the Microsoft Download Center.
2. Register Logparser.dll on your copy of Windows Server 2003 (version SP1 or later).
3. From a command prompt, navigate to the directory where you installed LogParser.
4. Type **regsvr32 logparser.dll** and press ENTER.
5. Open a command prompt and navigate to the directory where DumpTraceReqs.js is stored (%systemroom%\LogParser\Samples\Scripts, by default).
6. Type **dumptracereqs** *tracelogname* **>***newfilename*.txt and press ENTER.
7. *Tracelogname* is the name of the trace log .etl file and *newfilename*.txt is the name of a new file where the post-processed data will be saved. LogParser and DumpTraceReqs.js process the trace log file and store the post-processed data in the *newfilename*.txt file display the results to the screen.
8. To view the data, type **notepad** *newfilename*.txt and press ENTER.

**See Also**
**Concepts**
Tracing Events Reference

# Tracing Startup and Shutdown Issues in the IIS Admin Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

The IIS Admin service reads the IIS metabase.xml file and the related in-memory configuration services. Use this provider if you notice that the IIS Admin service has trouble starting up or appears to be unresponsive at any time.

This section includes the following information:

- How to Trace the IIS Admin Service
- How to Process IIS Admin Trace Log Files Using Trace Report

**See Also**
**Other Resources**
Tracing Startup and Shutdown Problems in the IIS Admin Service

# How to Trace the IIS Admin Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

Use the following procedure to trace the IIS Admin service if the service fails to startup or shutdown. IIS Admin tracing uses the following parameters:

| | |
|---|---|
| **logman start** | Starts the trace session using Logman.exe as the controller. |
| **logman stop** | Stops the trace session. |
| *session name* | Any name you want to give this trace session. |
| **-p "IIS: IISADMIN Global"** | Identifies which provider (-p) should be traced. |
| **(startup,shutdown)** | Tells Logman.exe to trace the IIS Admin service for startup and shutdown events. You can specify one or both flags, as needed. |
| **-ets** | Tells Logman.exe to start the trace session immediately. |

**To trace the IIS Admin service**

1. From a command prompt, type **logman start** *session name* **-p "IIS: IISADMIN Global"(startup,shutdown) -ets** and press ENTER.

   For example: **logman start IISAdminTrace -p "IIS: IISADMIN Global" (startup,shutdown) -ets**

2. Event Tracing for Windows prints to the screen details about the trace session you just started, including the name of the session, the file name where the trace data will be collected (*session name*.etl by default), and whether or not the command was successful.

3. If necessary, start or stop the IIS Admin service (depending on the nature of the problem you are trying to troubleshoot). Type **net stop IIS Admin** or **net start IIS Admin** and press ENTER.

4. Allow the trace session to run until you have reproduced the problem.

5. From the command prompt, type **logman stop** *session name* **-ets** and press ENTER.

   To view the trace data, see How to Process IIS Admin Trace Log Files Using Trace Report.

**See Also**
**Other Resources**
Tracing Startup and Shutdown Problems in the IIS Admin Service

# How to Process IIS Admin Trace Log Files Using Trace Report

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

IIS Admin trace data is written in binary format to a trace log (*filename*.etl). To view the trace data, the log file must be processed using the Trace Report tool that ships with Windows Server 2003.

**To process and view an IIS Admin trace log**

1. Open a command prompt and navigate to the directory where your IIS Admin trace log .etl file is stored.
2. Type **tracerpt** *filename*.etl and press ENTER. Trace Report processes the .etl file and creates two new files in the directory where you executed this command: summary.txt and dump.csv.
3. To view the file, type **notepad dump.csv** and press ENTER.

**See Also**
**Concepts**
Tracing Events Reference

# Windows Server 2003 Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003 Technical Reference**
Windows Server 2003 Technical Reference documentation provides comprehensive information about the technologies included in the Microsoft Windows Server 2003 operating systems, including Active Directory, Group Policy, and core operating system, high availability and scalability, networking, storage, and security technologies. This set of Technical Reference documentation is designed to assist IT planners and administrators by providing the foundational information about the technology components of the operating system. Windows Server 2003 Technical Reference documentation is not designed to provide prescriptive guidance about how to deploy, operate, or maintain Windows Server 2003 technologies.

**Structure**
Windows Server 2003 Technical Reference includes a number of collections covering each of the core technology areas for Windows Server 2003. Each collection is organized into subjects that provide detailed information about components of the technology, including a basic description of the component, a discussion of how the technology works, and information about tools and settings related to the technology.

In addition, an A-Z list of the technologies documented allows you to quickly find information about a specific feature or technology.

**Coverage**
Windows Server 2003 Technical Reference is a dynamic documentation set. Technical Reference documentation will continue to expand to cover additional technologies and tools. In addition, the content will be updated to reflect product upgrades, service packs, and critical updates. Please let us know if you do not find what you are looking for.

# A to Z Technologies List

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Symbols | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**Symbols**
4GT

802.11 Wireless

**A**
Access

Access Tokens

Account Policies

Active Directory

Active Directory Application Mode

Active Directory Data Store

Active Directory Domains and Forests

Active Directory DNS Support

Active Directory Functional Levels

Active Directory Installation

Active Directory Installation, Upgrade, and Migration

Active Directory Replication Model

Active Directory Replication Topology

Active Directory Schema

Active Directory Searches

Service Publication and Service Principal Names

Administrative Templates Extension

Asynchronous Transfer Mode(ATM)

Authentication

Authorization

**B**
Backup

Basic Disks and Volumes

**C**
CA Certificates

Certificates

Certificate Services

Clustering

Computer Browser Service

Connection Manager

Core Group Policy

**D**

Data Defragmentation

Data Security Technologies

Data Streaming and E-mail

Demand Dial Routing

Device Drivers

Device Management

Dial-up Remote Access

Digest

Disk Defragmenter

Disk Quotas

Disks and Volumes

Distributed File System (DFS)

DNS

DNS Support for Active Directory

Domain and Forest Trusts

Domain Controller Roles

Domain Rename

Dynamic Disks and Volumes

Dynamic Host Configuration Protocol (DHCP)

**E**

Emergency Management Services

Encrypting File System

Event Log

**F**

FAT File System

File Replication Service FRS

File Services

File System Policy Settings

File Systems

Folder Redirection Extension

**G**

Global Catalog

Group Policy

Group Policy Components

Group Policy Management Console

Group Policy Object Editor

**H**
High Availability

Installation and Setup

**I**
Interactive Logon

Internet Authentication Service

Internet Explorer Maintenance Extension

Internet Protocol Security

Internet Protocol, Version 6 (IPv6)

IPSec

IPSec Policy Extension

IP Security Policy Settings

IPv4 Multicast Routing

**J**
No entries

**K**
Kerberos

**L**
Local Policy Settings

Logon

**M**
No entries

**N**
Name Resolution

Network Access Technologies

Network Address Translation

Network Configuration

Network Interoperability Technologies

Network Load Balancing

Network Monitoring

Network Printing

Networking

NTFS File System

**O**
Operations Masters

**P**
Permissions

Physical Address Extension (PAE) X86

Plug and Play

Post Office Protocol, Version 3 (POP3)

Printing

Public Key Infrastructure (PKI)

Public Key Policy Settings

**Q**
Quality of Service

**R**
Recovery and Reliability

Recovery Console

Registry Policy Settings

Remote Installation Services (RIS)

Remote Installation Services Extension

Remote Procedure Call (RPC)

Remote Storage

Remote Access Technologies

Removable Storage

Resource and Memory Management

Restricted Groups Policy Settings

Resultant Set of Policy

Routing

**S**
Scripts Extension

Security

Security Identifiers

Security Descriptors and Access Control Lists

Security Policy Settings

Security Principals

Security Settings Extension

Server Clusters

Services for Netware

Setup

Shadow Copies for Shared Folders

Shutdown Event Tracker

Simple Network Management Protocol (SNMP)

Software Installation Extension

Software Restriction Policies

Software Restriction Policy Settings

Storage

Sysprep

System Services

**T**
TCP/IP

Telnet

Terminal Services

TLS/SSL

Trust Technologies

**U**
Unattended Installation

Unicast Routing

**V**
Virtual Disk Service

VPN

Volume Shadow Copy Service

**W**
Windows Firewall

Windows Time Service

WINS

Wireless Network Policies Extension

Wireless Network (IEEE 802.11) Policy Settings

**X**
No entries

**Y**
No entries

**Z**
No entries

# New and Updated Collections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 Technical Reference documentation is periodically expanded with new Collections, Subcollections, Subjects, and sections. Existing sections are also updated as new or changed information becomes available.

**Technologies Collections**
**Active Directory Collection**
**Active Directory Application Mode**

- Active Directory Application Mode Technical Reference **NEW**
    - What Is Active Directory Application Mode? **NEW**
    - How Active Directory Application Mode Works **NEW**
    - Active Directory Application Mode Tools and Settings **NEW**

**Active Directory Structure and Storage Technologies**

- Domains and Forests Technical Reference **NEW**
    - What Are Domains and Forests? **NEW**
    - How Domains and Forests Work **NEW**
    - Domains and Forests Tools and Settings **NEW**
- Active Directory Schema Technical Reference
    - What Is the Active Directory Schema?
    - How the Active Directory Schema Works
    - Active Directory Schema Tools and Settings **NEW**
- Data Store Technical Reference
    - What Is the Data Store?
    - How the Data Store Works
    - Data Store Tools and Settings
- DNS Support for Active Directory Technical Reference
    - What Is DNS Support for Active Directory?
    - How DNS Support for Active Directory Works
    - DNS Support for Active Directory Tools and Settings

**Domain Controller Roles**

- Global Catalog Technical Reference **NEW**
    - What Is the Global Catalog? **NEW**
    - How the Global Catalog Works **NEW**
    - Global Catalog Tools and Settings **NEW**
- Operations Masters Technical Reference **NEW**
    - What are Operations Masters? **NEW**
    - How Operations Masters Work **NEW**
    - Operations Masters Tools and Settings **NEW**

**Active Directory Replication Technologies**

- Active Directory Replication Model Technical Reference
    - What Is the Active Directory Replication Model?
    - How the Active Directory Replication Model Works
    - Active Directory Replication Tools and Settings
- Active Directory Replication Topology Technical Reference
    - What Is Active Directory Replication Topology?
    - How Active Directory Replication Topology Works
    - Active Directory Replication Tools and Settings

**Active Directory Search and Publication Technologies**

- Active Directory Searches Technical Reference **NEW**
    - What Are Active Directory Searches? **NEW**
    - How Active Directory Searches Work **NEW**
    - Active Directory Searches Tools and Settings **NEW**
- Service Publication and Service Principal Names Technical Reference
    - What Are Service Publication and Service Principal Names?
    - How Service Publication and Service Principal Names Work
    - Service Publication and Service Principal Names Tools and Settings

**Active Directory Installation, Upgrade, and Migration Technologies**

- Active Directory Functional Levels Technical Reference
    - What Are Active Directory Functional Levels?
    - How Active Directory Functional Levels Work
    - Active Directory Functional Levels Tools and Settings
- Active Directory Installation Wizard Technical Reference **NEW**
    - What Is the Active Directory Installation Wizard? **NEW**
    - How the Active Directory Installation Wizard Works **NEW**
- Domain Rename Technical Reference **NEW**
    - What Is Domain Rename? **NEW**
    - How Domain Rename Works **NEW**

**Core Operating System Collection**
**Printing Technologies**

- Network Printing Technical Reference
    - What Is Network Printing?
    - How Network Printing Works
    - Network Printing Tools and Settings

**Recovery and Reliability Technologies**

- Recovery Console Technical Reference
    - What Is Recovery Console?
    - Recovery Console Tools and Settings
- Emergency Management Services Technical Reference
    - What Is Emergency Management Services?
    - How Emergency Management Services Works
    - Emergency Management Services Tools and Settings
- Shutdown Event Tracker Technical Reference
    - What Is Shutdown Event Tracker?
    - How Shutdown Event Tracker Works
    - Shutdown Event Tracker Tools and Settings

- Backup Technical Reference
    - What Is Backup?
    - How Backup Works
    - Backup Tools and Settings
- Shadow Copies for Shared Folders Technical Reference
    - What Is Shadow Copies for Shared Folders?
    - How Shadow Copies for Shared Folders Work
    - Shadow Copies for Shared Folders Tools and Settings
- Volume Shadow Copy Service Technical Reference
    - What Is Volume Shadow Copy Service?
    - How Volume Shadow Copy Service Works
    - Volume Shadow Copy Service Tools and Settings

**Device Management Technologies**

- Device Driver Technical Reference
    - What Is a Device Driver?
    - How Device Drivers Work
    - Device Drivers Tools and Settings
- Plug and Play Technical Reference
    - What Is Plug and Play?
    - How Plug and Play Works
    - Plug and Play Tools and Settings

**Installation and Setup Technologies**

- Setup Technical Reference
    - What Is Setup?
    - How Setup Works
    - Setup Tools and Settings
- Unattended Installation Technical Reference
    - What Is Unattended Installation?

    - How Unattended Installation Works
    - Unattended Installation Tools and Settings
- Sysprep Technical Reference
    - What Is Sysprep?
    - How Sysprep Works
    - Sysprep Tools and Settings
- Remote Installation Services Technical Reference
    - What Is Remote Installation Services?
    - How Remote Installation Services Work
    - Remote Installation Services Tools and Settings

**Group Policy Collection**
**Core Group Policy Technical Reference**

- What Is Core Group Policy?
- How Core Group Policy Works
- Core Group Policy Tools and Settings

**Group Policy Components**

- Administrative Templates Extension Technical Reference
    - What Is Administrative Templates Extension?
    - How Administrative Templates Extension Works
    - Administrative Templates Extension Tools and Settings
- Group Policy Software Installation Extension Technical Reference
    - What Is Group Policy Software Installation Extension?
    - How Group Policy Software Installation Extension Works
    - Group Policy Software Installation Extension Tools and Settings
- Security Settings Extension Technical Reference
    - What Is Security Settings Extension?
    - How Security Settings Extension Works
    - Security Settings Extension Tools and Settings

- IPSec Policy Extension Technical Reference
    - What Is IPSec Policy Extension?
    - How IPSec Policy Extension Works
    - IPSec Policy Extension Tools and Settings
- Software Restriction Policies Technical Reference
    - What Are Software Restriction Policies?
    - How Software Restriction Policies Work
    - Software Restriction Policies Tools and Settings
- Scripts Extension Technical Reference
    - What Is Scripts Extension?
    - How Scripts Extension Works
    - Scripts Extension Tools and Settings
- Wireless Network Policies Extension Technical Reference
    - What Is Wireless Network Policies Extension?
    - How Wireless Network Policies Extension Works
    - Wireless Network Policies Extension Tools and Settings
- Folder Redirection Extension Technical Reference
    - What Is Folder Redirection Extension?
    - How Folder Redirection Extension Works
    - Folder Redirection Extension Tools and Settings
- Internet Explorer Maintenance Extension Technical Reference
    - What Is Internet Explorer Maintenance Extension?
    - How Internet Explorer Maintenance Extension Works
    - Internet Explorer Maintenance Extension Tools and Settings
- Remote Installation Services Extension Technical Reference
    - What Is Remote Installation Services Extension?
    - How Remote Installation Services Extension Works

- Remote Installation Services Extension Tools and Settings

**Group Policy Administrative Tools**

- Group Policy Management Console Technical Reference
  - What Is Group Policy Management Console?
  - How Group Policy Management Console Works
  - Group Policy Management Console Tools and Settings
- Group Policy Object Editor Technical Reference
  - What Is Group Policy Object Editor?
  - How Group Policy Object Editor Works
  - Group Policy Object Editor Tools and Settings
- Resultant Set of Policy Technical Reference
  - What Is Resultant Set of Policy?
  - How Resultant Set of Policy Works
  - Resultant Set of Policy Tools and Settings

**High Availability and Scalability Collection**
**Clustering Technologies**

- Server Clusters Technical Reference
  - What Is a Server Cluster?
  - How a Server Cluster Works
  - Server Cluster Tools and Settings **NEW**
- Network Load Balancing Technical Reference
  - What Is Network Load Balancing?
  - How Network Load Balancing Works
  - Network Load Balancing Tools and Settings

**Resource and Memory Management Technologies**

- 4GT Technical Reference
  - What Is 4GT?
  - How 4GT Works

    - 4GT Tools and Settings
- PAE X86 Technical Reference
    - What Is PAE X86?
    - How PAE X86 Works
    - PAE X86 Tools and Settings
- Windows System Resource Manager Technical Reference
    - What is Windows System Resource Manager?
    - How Windows System Resource Manager Works

**Networking Collection**
**Core Networking Technologies**

- TCP/IP Technical Reference
    - What Is TCP/IP?
    - How TCP/IP Works
    - TCP/IP Tools and Settings
- IPv6 Technical Reference
    - What Is IPv6?
    - How IPv6 Works
    - IPv6 Tools and Settings

**Routing Technologies**

- Demand Dial Routing Technical Reference **NEW**
    - What Is Demand Dial Routing? **NEW**
    - How Demand Dial Routing Works **NEW**
    - Demand Dial Tools and Settings **NEW**
- IPv4 Multicasting Technical Reference **NEW**
    - What Is IPv4 Multicasting? **NEW**
    - How IPv4 Multicasting Works **NEW**
    - IPv4 Multicasting Tools and Settings **NEW**
- NAT Technical Reference

    - What Is NAT?
    - How NAT Works
    - NAT Tools and Settings
- Unicast IPv4 Routing Technical Reference **NEW**
    - What Is Unicast IPv4 Routing? **NEW**
    - How Unicast IPv4 Routing Protocols and Services Work **NEW**

**Network Monitoring Technologies**

- SNMP Technical Reference
    - What Is SNMP?
    - How SNMP Works
- Computer Browser Service Technical Reference
    - What Is Computer Browser Service?
    - How Computer Browser Service Works

**Name Resolution Technologies**

- DNS Technical Reference
    - What Is DNS?
    - How DNS Works
    - DNS Tools and Settings
- WINS Technical Reference
    - What Is WINS?
    - How WINS Works
    - WINS Tools and Settings

**Network Configuration Technologies**

- DHCP Technical Reference
    - What Is DHCP?
    - How DHCP Works
    - DHCP Tools and Settings
- Windows Time Service Technical Reference

- What Is Windows Time Service?
- How Windows Time Service Works
- Windows Time Service Tools and Settings **NEW**

**Network Access Technologies**

- Connection Manager Technical Reference
  - What Is Connection Manager?
  - How Connection Manager Works
  - Connection Manager Tools and Settings
- IAS Technical Reference
  - What Is IAS?
  - How IAS Works
  - IAS Tools and Settings
- VPN Technical Reference
  - What Is VPN?
  - How VPN Works
  - VPN Tools and Settings
- 802.11 Wireless Technical Reference
  - What Is 802.11 Wireless?
  - How 802.11 Wireless Works
  - 802.11 Wireless Tools and Settings

Network Interoperability Technologies **NEW**

- RPC Technical Reference **NEW**
  - What Is RPC? **NEW**
  - How RPC Works **NEW**
- Services for Netware Technical Reference **NEW**
  - What Is Services for Netware **NEW**
  - How Services for Netware Works **NEW**

- Services for Netware Tools and Settings **NEW**

**Remote Access Technologies**

- Dial-up Remote Access Technical Reference **NEW**
  - What Is Dial-up Remote Access? **NEW**
  - How Dial-up Remote Access Works **NEW**
  - Dial-up Remote Access Tools and Settings **NEW**
- Telnet Technical Reference
  - What Is Telnet?
  - How Telnet Works
  - Telnet Tools and Settings
- Terminal Services Technical Reference
  - What Is Terminal Services?
  - How Terminal Services Works
  - Terminal Services Tools and Settings

**Data Streaming and E-mail Technologies**

- ATM Technical Reference
  - What Is ATM?
  - How ATM Works
- POP3 Service Technical Reference
  - What Is POP3 Service?
  - How POP3 Service Works
  - POP3 Service Tools and Settings
- QoS Technical Reference **NEW**
  - What Is QoS? **NEW**
  - How QoS Works **NEW**
  - QoS Tools and Settings **NEW**

**Storage Technologies Collection**
**File Services Technologies**

- Disk Quotas Technical Reference

Disk Quotas Technical Reference

    - What Are Disk Quotas?
    - How Disk Quotas Work
    - Disk Quotas Tools and Settings
- DFS Technical Reference
    - What Is DFS?
    - How DFS Works
    - DFS Tools and Settings
- FRS Technical Reference
    - What Is FRS?
    - How FRS Works
    - FRS Tools and Settings

**File Systems Technologies**

- NTFS Technical Reference
    - What Is NTFS?
    - How NTFS Works
    - NTFS Tools and Settings
- FAT Technical Reference
    - What Is FAT?
    - How FAT Works
    - FAT Tools and Settings
- Disk Defragmenter Technical Reference
    - What Is Disk Defragmenter?
    - How Disk Defragmenter Works
    - Disk Defragmenter Tools and Settings

**Disks and Volumes**

- Basic Disks and Volumes Technical Reference
    - What Are Basic Disks and Volumes?

    - How Basic Disks and Volumes Work
    - Basic Disks and Volumes Tools and Settings
- Dynamic Disks and Volumes Technical Reference
    - What Are Dynamic Disks and Volumes?
    - How Dynamic Disks and Volumes Work
    - Dynamic Disks and Volumes Tools and Settings

**Storage Services Technologies**

- Virtual Disk Service Technical Reference
    - What Is Virtual Disk Service?
    - How Virtual Disk Service Works
    - Virtual Disk Service Tools and Settings
- Remote Storage Technical Reference
    - What Is Remote Storage?
    - How Remote Storage Works
    - Remote Storage Tools and Settings
- Removable Storage Technical Reference
    - What Is Removable Storage?
    - How Removable Storage Works
    - Removable Storage Tools and Settings

**Windows Security Collection**
**Logon and Authentication Technologies**

- Digest Authentication Technical Reference **NEW**
    - What is Digest Authentication? **NEW**
    - How Digest Authentication Works **NEW**
    - Digest Authentication Tools and Settings **NEW**
- Interactive Logon Technical Reference **NEW**
    - What is Interactive Logon? **NEW**
    - How Interactive Logon Works **NEW**

    - Interactive Logon Tools and Settings **NEW**
- Kerberos Authentication Technical Reference
    - What Is Kerberos Authentication?
    - How the Kerberos Version 5 Authentication Protocol Works
    - Kerberos Authentication Tools and Settings
- TLS/SSL Technical Reference **NEW**
    - What is TLS/SSL? **NEW**
    - How TLS/SSL Works **NEW**
    - TLS/SSL Tools and Settings **NEW**

**Authorization and Access Control Technologies**

- Security Descriptors and Access Control Lists Technical Reference **NEW**
    - What Are Security Descriptors and Access Control Lists? **NEW**
    - How Security Descriptors and Access Control Lists Work **NEW**
    - Security Descriptors and Access Control Lists Tools and Settings **NEW**
- Access Tokens Technical Reference
    - What Are Access Tokens?
    - How Access Tokens Work
    - Access Tokens Tools and Settings
- Permissions Technical Reference **NEW**
    - What Are Permissions? **NEW**
    - How Permissions Work **NEW**
    - Permissions Tools and Settings **NEW**
- Security Principals Technical Reference **NEW**
    - What Are Security Principals? **NEW**
    - How Security Principals Work **NEW**
    - Security Principals Tools and Settings **NEW**
- Security Identifiers Technical Reference **NEW**
    - What Are Security Identifiers? **NEW**

    - How Security Identifiers Work **NEW**
    - Security Identifiers Tools **NEW**

**Data Security Technologies**

- Encrypting File System Technical Reference **NEW**
    - What is Encrypting File System? **NEW**
    - How Encrypting File System Works **NEW**
    - Encrypting File System Tools and Settings **NEW**
- IPSec Technical Reference
    - What Is IPSec?
    - How IPSec Works
    - IPSec Tools and Settings **NEW**

**PKI Technologies**

- CA Certificates Technical Reference **NEW**
    - What Are CA Certificates? **NEW**
    - How CA Certificates Work **NEW**
    - CA Certificates Tools and Settings **NEW**
- Certificates Technical Reference
    - What Are Certificates?
    - How Certificates Work
    - Certificates Tools and Settings
- Certificate Services Technical Reference
    - What Is Certificate Services?
    - How Certificate Services Works
    - Certificate Services Tools and Settings

**Trust Technologies**

- Domain and Forest Trusts Technical Reference
    - What Are Domain and Forest Trusts?

- How Domain and Forest Trusts Work
- Domain and Forest Trust Tools and Settings
- Security Considerations for Trusts **NEW**

**Tools and Settings Collection**
**Security Policy Settings**

- Account Policies
- Local Policies
- Event Log
- Restricted Groups
- System Services
- Registry
- File System
- Wireless Network (IEEE 802.11) Policies **NEW**
- Public Key Policy Settings **NEW**
- Software Restriction Policy Settings **NEW**
- IP Security Policies on Active Directory **NEW**

**Resource Kit Tools Help**
**Windows Support Tools**
**Windows Firewall Settings**

- Windows Firewall Settings: Optional Components **NEW**
- Windows Firewall Settings: Remote Administration Tools **NEW**
- Windows Firewall Settings: Server Roles **NEW**
- Windows Firewall Settings: Services **NEW**

# Technologies Collections

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Technical Reference documentation for the Microsoft Windows Server 2003 family of operating systems includes the following technology collections:

### Active Directory Collection

This collection provides in-depth technical reference information about the Windows Server 2003 Active Directory directory service and its components. These components include logical structure, physical structure, data storage, replication, and service publication. In addition, this collection includes technical reference information about the tools that are used to install, upgrade, and migrate an Active Directory environment.

### Core Operating System Collection

This collection describes the core Windows Server 2003 operating system technologies. These technologies provide basic functionality such as printing services, data recovery, and troubleshooting.

### Group Policy Collection

This collection provides in-depth technical information about Windows Server 2003 Group Policy and its supporting technologies, including core Group Policy, Group Policy components, and Group Policy administrative tools. These technologies enable you to deliver, apply, and centrally manage one or more configurations, including Administrative Templates (registry-based policy), folder redirection, scripts, security, and software installation settings, to a set of targeted users and computers within an Active Directory environment.

### High Availability and Scalability Collection

This collection describes the Windows Server 2003 high availability technologies that are designed to reduce unscheduled outages and mitigate the impact of scheduled downtime for servers. This collection also describes the Windows Server 2003 scalability technologies used to increase or decrease computing capacity in order to address the capacity needs of an organization.

### Networking Collection

This collection provides in-depth technical reference information about Windows Server 2003 technologies that make up a networking infrastructure based on the TCP/IP suite of protocols, services, and other industry-standard networking protocols and technologies.

### Storage Technologies Collection

This collection describes the Windows Server 2003 storage technologies that are designed to help you store, access, and manage data on servers.

### Windows Security Collection

This collection describes Windows Server 2003 security technologies, including authentication, authorization, and public key infrastructure (PKI), and provides detailed information about security policy settings.

# Active Directory Collection

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Active Directory Domain Services Collection**
Active Directory Domain Services (AD DS) directory service is the distributed directory service that is included with Microsoft Windows Server operating systems. AD DS enables centralized, secure management of an entire network, which might span a building, a city, or multiple locations throughout the world.

AD DS includes the following:

- AD DS on a Windows Server Network
- Active Directory Lightweight Directory Services (AD LDS)
- Structure and Storage Technologies
- Domain Controller Roles
- Replication Technologies
- Search and Publication Technologies
- Installation, Upgrade, and Migration Technologies

In distributed computing environments, networked computers and other devices communicate over remote connections to accomplish tasks through client/server applications. Distributed environments require a central repository of information and integrated services that provide the means to manage network users, services, devices, and additional information that administrators want to store.

Organizations operating a distributed environment need to have a way to manage network resources and services. As the organization grows, the need for a secure and centralized management system becomes more critical.

A directory service provides a centralized location to store information in a distributed environment about networked devices and services and the people who use them. A directory service also implements the services that make this information available to users, computers, and applications. A directory service is both a database storage system (directory store) and a set of services that provide the means to securely add, modify, delete, and locate data in the directory store.

AD DS is typically used for one of three purposes:

- Internal directory. Used within the corporate network for publishing information about users and resources within the enterprise. A company's internal directory may be accessible to employees when they are outside the company network using a secure connection such as a virtual private network (VPN) connection, but it is not accessible to non-employees.

- External directory. These are directories typically located on servers in the perimeter network or demilitarized zone (DMZ) at the boundary between the corporate local area network (LAN) and the public Internet. External directories are typically used to store information about customers, clients, and business partners who access external applications or services. They are also made available to customers, clients, and business partners to provide them with selected business information such as catalogs and so on.

- Application directory. Application directories store "private" directory data that is relevant only to the application in a local directory, perhaps on the same server as the application, without requiring any additional configuration to Active Directory. The personalization data, which is only interesting to the portal application and does not need to be widely replicated, can be stored solely in the directory associated with the application. This solution reduces replication traffic on the network between domain controllers.

**AD DS on a Windows Server Network**

AD DS is the information hub of the operating system. The following figure shows AD DS as the focal point of the Windows Server network used to manage identities and broker relationships between distributed resources so they can work together.

**Active Directory on a Windows Server Network**




Active Directory provides:

- A central location for network administration and delegation of administrative authority. You have access to objects representing all network users, devices, and resources and the ability to group objects for ease of management and application of security and Group Policy.

- Information security and single sign-on for user access to network resources. Tight integration with security eliminates costly tracking of accounts for authentication and authorization between systems. A single user name and password combination can identify each network user, and this identity follows the user throughout the network.

- Scalability. AD DS includes one or more domains, each with one or more domain controllers, enabling you to scale the directory to meet any network requirements.

- Flexible and global searching. Users and administrators can use desktop tools to search AD DS. By default, searches are directed to the global catalog, which provides forest-wide search capabilities.

- Storage for application data. AD DS provides a central location to store data that is shared between applications and with applications that need to distribute their data across entire Windows networks.

- Systematic synchronization of directory updates. Updates are distributed throughout the network through secure and cost-efficient replication between domain controllers.

- Remote administration. You can connect to any domain controller remotely from any Windows-based computer that has administrative tools installed.

- Single, modifiable, and extensible schema. The schema is a set of objects and rules that provide the structure requirements for AD DS objects. You can modify the schema to implement new types of objects or object properties.

- Integration of object names with Domain Name System (DNS), the Internet-standard computer location system. AD DS uses DNS to implement an IP-based naming system so that AD DS services and domain controllers are locatable over standard IP both on intranets and the Internet.

- Lightweight Directory Access Protocol (LDAP) support. LDAP is the industry standard directory access protocol, making AD DS widely accessible to management and query applications. AD DS supports LDAPv3 and LDAPv2.

**Active Directory Lightweight Directory Services (AD LDS)**
Active Directory Lightweight Directory Services (AD LDS) is a directory service designed to meet the needs of organizations that cannot rely solely on AD DS to provide directory services for directory-enabled applications. While AD DS offers many benefits for managing network infrastructure, organizations often need a more flexible directory service to support directory-enabled applications. AD LDS is a Lightweight Directory Access Protocol (LDAP) directory service designed specifically for directory-enabled applications.

**Structure and Storage Technologies**
AD DS uses domains and forests to represent the logical structure of the directory hierarchy. Domains are used to manage the various populations of users, computers, and network resources in your enterprise. The forest represents the security boundary for AD DS. Within domains you can create organizational units to subdivide the various divisions of administration.

The logical structure of AD DS includes a two-dimensional definition that can be viewed as a hierarchy, even though the objects themselves are stored in a flat database file. In addition to its own name, each object stores the name of the container directly above it in the hierarchy. That container object stores the name of its superior container, and so on, up to the root container. In this way, a logical structure is imposed that can be viewed by using AD DS tools as a tree of containers. By virtue of a hierarchical naming system, the objects in the tree appear to be nested inside (contained by) other objects.

The AD DS schema defines the types of objects that are available to the directory service. The schema is stored in the schema partition, which is also defined as an object in the directory. The attributes and classes in AD DS are stored in the schema partition as directory objects called schema objects. It is possible for Administrators to add their own classes or attributes to an existing object type. However, the default schema provides all of the classes and attributes that AD DS needs to function.

AD DS uses objects to store and reference data in the directory. The AD DS database file (Ntds.dit) provides the physical storage of all AD DS objects for a single forest. Although there is a single directory, some directory data is stored within domains while other data is distributed throughout the forest, without regard for domain boundaries. Beginning with Windows Server 2003, data can also be distributed to domain controllers according to applications that use the data, where the scope of distribution can be set according to the needs of the application.

Any updates made to data in the directory are automatically distributed to the appropriate domain controllers by means of AD DS replication. By replicating data according to directory partitions, AD DS provides a data repository that is logically centralized (maintains a single point of administration) but physically distributed (is synchronized on multiple domain controllers throughout the network).

**Replication Technologies**
Objects in the directory are distributed among the domain controllers in a forest, and all domain controllers can be updated directly. AD DS replication is the process by which the changes that are made on one domain controller are automatically synchronized with other domain controllers. Data integrity is maintained by tracking changes on each domain controller and updating other domain controllers in a systematic way. By default, AD DS replication uses a connection topology that is created automatically. This replication topology makes optimal use of physical network connections and frees administrators from having to determine which domain controllers replicate with one another. The replication topology can also be created manually. AD DS replication is designed to maximize directory consistency and minimize the impact to network traffic.

**Domain Controller Roles**
A domain controller is a server that has the AD DS server role installed.

### Note

- Implementations of Microsoft Windows NT 3.51 and Microsoft Windows NT 4.0 operating systems also have domain controllers, but they do not support AD DS.

When you install Windows Server on a computer, you can choose to configure a server role for that computer. When you want to create a new forest, a new domain, or an additional domain controller in an existing domain, you configure the server as a domain controller by installing AD DS.

By default, a domain controller stores one domain directory partition consisting of information about the domain in which it is located, plus the schema and configuration directory partitions for the entire forest. A domain controller can also store one or more application directory partitions.

Whereas every domain controller stores the objects for only one domain, a domain controller that is designated as a global

catalog server stores the objects from all domains in the forest. For each object that is not in the domain for which the global catalog server is authoritative as a domain controller, a limited set of attributes is stored in a partial replica of a corresponding domain. The partial replicas on a global catalog server are not writable — you cannot update an object in a partial replica on a global catalog server, but only on a domain controller that stores a full replica. Thus a global catalog server stores its own full, writable domain replica (all objects and all attributes) plus a partial, read-only replica of every other domain in the forest. The attributes that are replicated to the global catalog servers are the attributes that are most likely to be used to search for the object in AD DS. These attributes are identified by default in the schema as being included in the partial attribute set of the global catalog.

The global catalog makes it possible for clients to search AD DS without having to be referred from server to server until the domain controller that has the domain that stores the requested object is found. By default, AD DS searches are directed to global catalog servers. The first domain controller in a forest is automatically created as a global catalog server. Thereafter, you can designate other domain controllers to be global catalog servers if they are needed.

All domain controllers can receive updates to any writable object that they store (with the exception of schema updates, which can be made only on the one domain controller in the forest that has the role of schema master). The day-to-day operations that are associated with managing users, groups, and computers are typically multimaster operations — that is, changes to these objects can be made on any domain controller. When a client application updates an object on a domain controller, the domain controller automatically replicates the change to all other domain controllers in the same domain if the change is a domain change or to all other domain controllers in the forest if the change is a configuration or schema change.

There are some operations, however, that are not performed as multimaster operations because they must occur at only one place and time. For these operations, there are specially designated domain controllers that manage the operations singly. Some master operations, required at the forest level, include the schema master and the domain naming master. Others, required at the domain level, include the PDC emulator, RID master and infrastructure master. Domain controllers that hold these special roles are called operations masters.

**Search and Publication Technologies**

Successful operation of an AD DS forest depends on clients and services being able to locate domain controllers. The success of domain controller location depends on the registration of information in DNS and the availability of that information. AD DS uses DNS to locate networked computers by resolving computer names to IP addresses. The Net Logon service on domain clients and domain controllers interacts with Windows server application programming interfaces (APIs) and DNS to provide a domain controller locator service (Locator). Locator finds requested service-specific and site-specific domain controllers.

After a domain controller has been located, LDAP is used to retrieve information from the directory. AD DS stores objects that provide information about the real objects that exist in an organization's network and that are associated with one or more domains, such as users, specific groups of users, computers, applications, services, files, and distribution lists. AD DS makes this information available to administrators, network users, and applications throughout the organization through LDAP. LDAP enables clients to query, create, update, and delete information stored in a directory service. The LDAP protocol is the AD DS core protocol, and is the preferred and most common way of interacting with AD DS.

The creation, storage, and maintenance of information in AD DS is called service publication. Directory-enabled services and applications can publish globally useful information, such as service availability and properties, in AD DS. This allows client processes to find and connect to any directory-enabled service as needed, and network clients and administrators to find, connect to, and manage services.

**Installation, Upgrade, and Migration Technologies**

The installation or removal of AD DS is performed by the Active Directory Installation Wizard. Before installing AD DS on a server, the wizard will verify that the server is eligible to run AD DS. After the prerequisites have been met, a user interface is used to gather information specific to the environment in which AD DS will be installed. Finally, the wizard configures the directory service, making the server a domain controller.

Part of the directory configuration process includes configuring the AD DS schema. The schema contains a master list of all classes (object types) and attributes that can be used in the directory. The Active Directory Preparation Tool (ADPrep) is used to prepare an AD DS forest and domain for a newer version of the directory service. One of several tasks accomplished by ADPrep is updating the AD DS schema. If you do not prepare your AD DS infrastructure, the upgrade will fail.

After installing or upgrading AD DS, you can enable the appropriate domain or forest functional level based on an assessment of your current environment. The functional level of a domain or forest defines the set of advanced AD DS features that are available in that domain or forest. The functional level of a domain or forest also defines the set of Windows operating systems that can run on the domain controllers in that domain or forest. Functional levels provide configuration support for the AD DS features and ensure compatibility with domain controllers running earlier operating systems.

Depending on the design of your environment, you might opt to restructure it instead of upgrading. For example, if your Windows NT 4.0 environment consists of multiple domains, rather than upgrading each domain it might be more productive

to restructure the environment by consolidating some of those domains. Or if your Windows 2000 environment was poorly designed and you are upgrading your environment to Windows Server 2003, it might benefit you to restructure your existing environment before or after the upgrade takes place. You can perform either of these tasks by using the Active Directory Migration Tool (ADMT). ADMT includes wizards that automate migration tasks such as copying users, groups, and service accounts; moving computers; migrating trusts; and performing security translation. When you use ADMT to restructure Windows NT 4.0 domains, ADMT copies the accounts that are migrated, so that when the accounts are created in the target domain, they continue to exist in the source domain. The primary security identifiers (SIDs) for the accounts can be migrated to the SID history in the target domain. SID history maintains resource permissions when you migrate accounts, thus enabling access to resources in the source domain.

Another method for restructuring an AD DS environment is to rename a domain. You can use the domain rename process to change the names of your domains, and you can also use it to change the structure of the domain trees in your forest. This process involves updating the Domain Name System (DNS) and trust infrastructures as well as Group Policy and service principal names (SPNs).

The ability to rename domains provides you with the flexibility to make important name changes and forest structural changes as the needs of your organization change. Using domain rename, you can not only change the name of a domain, but you can change the structure of the domain hierarchy and change the parent of a domain or move a domain located in one domain tree to another domain tree.

# Active Directory Application Mode

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2012, Windows Server 2012 R2

Active Directory Application Mode (ADAM) is a new mode of the Active Directory directory service that is designed to meet the specific needs of organizations that use directory-enabled applications. While Active Directory supports directory-enabled applications, as well as the server operating system, some directory-enabled applications have requirements that Active Directory does not meet. For example, some directory-enabled applications require schema changes that administrators may not want to make to Active Directory.

In addition, organizations may want to:

- Support directory-enabled applications but not implement Active Directory domains and forests.
- Support directory-enabled applications outside their existing domains and forests.
- Use X.500 naming conventions for top-level directory partitions.
- Run multiple directory service instances on a single server.

ADAM is designed to support these and other directory service scenarios. ADAM runs completely independently from Active Directory, and ADAM has its own schema. You can make changes to the ADAM schema with no impact to the Active Directory schema. ADAM does not require the presence of Active Directory domain controllers, domains, or forests. Therefore, organizations that have not implemented Active Directory can install and run ADAM. ADAM supports X.500 naming conventions for top-level directory partitions, and you can run multiple instances of ADAM on a single server.

ADAM runs on the following:

- Domain controllers running operating systems in the Microsoft Windows Server 2003 family (Note that Microsoft Windows Server 2003, Web Edition cannot be a domain controller.)
- Member servers running operating systems in the Windows Server 2003 family (except for Windows Server 2003, Web Edition)
- Client computers running Microsoft Windows XP Professional

# Active Directory Application Mode Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2012, Windows Server 2012 R2

Active Directory Application Mode (ADAM) is a new mode of the Active Directory directory service that is designed specifically for directory-enabled applications. ADAM is a Lightweight Directory Access Protocol (LDAP) directory service that runs as a user service, rather than as a system service. You can run ADAM on servers and domain controllers running operating systems in the Microsoft Windows Server 2003 family (except for Windows Server 2003, Web Edition) and also on client computers running Microsoft Windows XP Professional. ADAM does not require the deployment of domains or domain controllers. You can run multiple instances of ADAM concurrently on a single computer, with an independently managed schema for each ADAM instance.

**In this subject**


- What Is Active Directory Application Mode?
- How Active Directory Application Mode Works
- Active Directory Application Mode Tools and Settings

# What Is Active Directory Application Mode?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- The Business Need
- The ADAM Solution
- LDAP Directories Compared to Relational Databases
- ADAM Compared to Active Directory
- ADAM Scenarios
- Related Information


Active Directory Application Mode (ADAM) is a new mode of Active Directory that is designed specifically for directory-enabled applications. ADAM is a Lightweight Directory Access Protocol (LDAP) directory service that runs as a user service, rather than as a system service. You can run ADAM on servers and domain controllers running operating systems in the Windows Server 2003 family (except for Windows Server 2003, Web Edition) and also on client computers running Windows XP Professional.

ADAM contains many Active Directory components and features, but ADAM differs from Active Directory in several key ways, as described in "ADAM Compared to Active Directory" later in this section. Some of the scenarios for which ADAM is designed are described in "ADAM Scenarios" later in this section.

**The Business Need**

ADAM is a directory service that is designed to meet the needs of organizations that cannot rely solely on Active Directory for providing directory services for directory-enabled applications. While Active Directory offers many benefits for managing network infrastructure, organizations often need a more flexible directory service to support directory-enabled applications.

For example, most directory-enabled applications require the directory service to be extended with application-specific schema extensions. Such schema extensions are supported in Active Directory. However, because Active Directory is mission critical for overall network operations, many administrators consider Active Directory to be "off limits" with respect to application-specific schema changes. In addition, many LDAP directory usage scenarios, including development, portal, and legacy application scenarios, are hindered by the forest and domain requirements of Active Directory.

**The ADAM Solution**

ADAM is a new mode of Active Directory that is designed for organizations that require flexible support for directory-enabled applications. ADAM provides data storage and retrieval for directory-enabled applications, without the dependencies that are required for Active Directory. ADAM provides much of the same functionality as Active Directory, including multimaster replication, but ADAM does not require the deployment of domains or domain controllers. You can run multiple instances of ADAM concurrently on a single computer, with an independently managed schema for each ADAM instance. You can run ADAM on servers and domain controllers running operating systems in the Windows Server 2003 family (except for Windows Server 2003, Web Edition) and on computers running Windows XP Professional.

The following figure illustrates the ADAM solution. In the figure, the computer running ADAM is hosting three separate ADAM instances, each with independently managed schema and data. Each ADAM instance is supporting a different directory-enabled application. The computers in this figure can reside within or outside an Active Directory forest.

**Active Directory Application Mode Solution**




**LDAP Directories Compared to Relational Databases**

Both LDAP directories (such as ADAM and Active Directory) and relational databases (such as Microsoft SQL Server 2000) can be used to support applications. However, LDAP directories and relational databases are designed for different purposes. Understanding the basic differences between these two technologies can help you choose the right technology for your application.

The following table summarizes the basic design differences between LDAP directories and relational databases.

**Design Differences Between LDAP Directories and Relational Databases**

| **Directories** | **Databases** |
|---|---|
| Optimized for search and read operations. | Optimized for write operations. |
| Object-oriented, hierarchical data design. Data objects in the directory represent entities such as users, computers, and shared resources. These data objects can be organized hierarchically in containers. | Relational data design. Data is organized in tables of rows and columns. Data from one table can be linked to data in another table. |
| Uses standardized, extensible schemas. | Does not use schemas. |
| Designed for replication and distributed management. | Designed for central storage and administration of data. |
| Precise security, down to the object and attribute level. | Less precise security, only down to the row and column level. |
| Loose data consistency between replication partners. | Transactional — guaranteed data consistency. Referential integrity across relational tables and concurrency control with file and record locking. |

Given these design differences, LDAP directories and relational databases are each particularly suited to certain types of applications.

LDAP directories are well suited for:

- Applications that manage data that is needed by clients throughout the network. Because LDAP directories can replicate data throughout a network easily and efficiently, LDAP directories are well suited to applications that require wide data distribution.

- Applications with a high ratio of search and read operations to write operations. Because LDAP directories are optimized for search and read operations, they are best suited to applications in which data changes infrequently.

- Applications that benefit from hierarchical storage and security models. LDAP directory services are particularly useful

when the data that they store benefits from container-style storage and inheritance-based security. For example, data that is related to enterprise "white pages" usually maps well to a hierarchical model. In other cases in which data corresponds to large business-to-consumer (B2C) e-commerce sites — where users are stored as huge, flat lists, LDAP directories may not be as well suited.

Good candidate applications for LDAP directories include the following:

- White pages applications that store contact and location information about people
- User, security, and resource management applications
- Applications that store network configuration and service policies

Relational databases are well suited for:

- Applications that require complex transactional updates. Directory services do not support complex transactions across multiple object types or tables. Therefore, applications that must observe atomicity, consistency, isolation, and durability (ACID), such as inventory databases, are best suited to relational databases.
- Applications that require change history information. With LDAP directories, it is not practical to create a centralized log that reflects all the updates that have been applied on all replicas throughout the network. Therefore, applications that require a log of changes that are made to data (for example, for auditing reasons) are best suited to relational databases.
- Applications that do not tolerate replication latency and inconsistency. Because of the store-and-forward style of replication that is used by directory services and the fact that directory service replicas are frequently connected by unreliable network connections, directory-enabled applications cannot be designed to rely on updates that propagate between replicas on any particular schedule or on updates that arrive in a particular order.

Good candidate applications for relational databases include the following:

- Airline reservation systems
- Banking systems
- Accounting applications that require transactional update protections and updates that typically occur quickly
- Centralized data collection
- Hardware inventory applications. Most data about the hardware configuration of client and server computers is not relevant to applications throughout the network.
- Process control. Any application that depends on the timing of data propagation between replicas is not well suited for a directory service.

**ADAM Compared to Active Directory**

ADAM is a special mode of Active Directory that is based on the same source code as Active Directory. To a large extent, ADAM operates just like Active Directory. However, because ADAM is designed specifically to support directory-enabled applications, rather than server operating system operations, there are some key differences between ADAM and Active Directory. These differences are related to:

- Supported operating systems
- Schema
- Partition naming

- Directory server consolidation
- Replication
- Dependencies
- Performance and scalability

**Supported Operating Systems**
**The following table lists the operating systems that are supported by Active Directory and ADAM.**
**Operating Systems Supported by Active Directory and ADAM**

| Directory Service | Platforms Supported |
|---|---|
| Active Directory | • Domain controllers running Windows 2000 Server<br>• Domain controllers running Windows 2000 Advanced Server<br>• Domain controllers running Windows 2000 Datacenter Server<br>• Domain controllers running Windows Server 2003, Standard Edition<br>• Domain controllers running Windows Server 2003, Enterprise Edition<br>• Domain controllers running Windows Server 2003, Datacenter Edition |
| ADAM | • Computers running Windows XP Professional<br>• Domain controllers and servers running Windows Server 2003, Standard Edition<br>• Domain controllers and servers running Windows Server 2003, Enterprise Edition<br>• Domain controllers and servers running Windows Server 2003, Datacenter Edition |

When you run ADAM on Windows XP Professional, there are certain considerations to keep in mind. For more information, see "ADAM Servers" in "How Active Directory Application Mode Works."

**Schema**
The schema defines what kind of information the directory can hold. With Active Directory, a single schema applies to the entire forest; that is, all domain controllers must use the same schema. With ADAM, each stand-alone ADAM instance or ADAM configuration set can have its own schema. In addition, the schema of each ADAM instance is fully independent from the Active Directory schema.

Another difference between ADAM and Active Directory is the size of the base schema. The base schema of Active Directory is designed to support overall server operating system operations, and it is therefore very large. The base schema of ADAM is very small — just large enough to bootstrap the ADAM instance. The ADAM schema is designed to be extended with the schema elements that are required for the applications that the ADAM instance supports.

Finally, while only certain object classes in Active Directory can be used to bind to (or authenticate against) Active Directory, you can give any object class in ADAM the ability to bind to ADAM by adding the **msDS-bindableobject** auxiliary class and the **unicodePwd** attribute to the schema definition of an object class. For more information about bindable object classes, see "ADAM Security Principals" in "How Active Directory Application Mode Works."

**Partition Naming**
A directory partition is a logical portion of the directory data store. Each ADAM directory partition has its own unique distinguished name. ADAM partition naming is more flexible than partition naming in Active Directory. Active Directory supports only Domain Name System (DNS)–style (DC=) names for top-level directory partitions. Such names can contain only DC= naming components. For example, Active Directory can have a partition named DC=Microsoft,DC=COM, but it cannot have a partition named DC=Microsoft,C=US. In contrast to Active Directory, ADAM can have a directory partition named

DC=Microsoft,C=US.

ADAM also supports both DNS-style and X.500-style names for top-level directory partitions. ADAM supports all the distinguished name components that are listed in the following table.

**ADAM Directory Partition Distinguished Name Components**

| Distinguished Name Component | Description |
|---|---|
| C= | Country/region |
| CN= | Common name |
| DC= | Domain component |
| L= | Location |
| O= | Organization |
| OU= | Organizational unit |

For more information about directory partitions, see "Directory Partitions" in "How Active Directory Application Mode Works."

**Directory Server Consolidation**
Because of the limitations that are placed on Active Directory by forests and domains, each domain controller can run only one instance of Active Directory. However, with ADAM you can run multiple directory service instances on each computer. Running multiple instances of ADAM on a single computer helps you lower your total cost of ownership (TCO) by consolidating directory operations onto a single server or just a few servers.

**Replication**
Active Directory domain controllers replicate with each other based on membership in a domain or forest. ADAM instances replicate with each other based on membership in a configuration set. A configuration set is a group of ADAM instances that share and replicate a common configuration partition and schema partition. ADAM instances cannot replicate with Active Directory. Furthermore, ADAM instances replicate on a schedule that is completely independent of the Active Directory replication schedule, even when ADAM is running in an Active Directory forest.

**Dependencies**
One key difference between ADAM and Active Directory is that ADAM is free of some of the underlying dependencies that are required by Active Directory. For example, ADAM does not require or rely on forests or domains. In addition, ADAM does not rely on the File Replication service (FRS) or DNS, both of which are required for Active Directory to function properly.

**Performance and Scalability**
Because ADAM is based on the same code as Active Directory, ADAM has the same scalability and performance characteristics as Active Directory. In other words, ADAM supports a comparable number of users, sites, and objects as Active Directory. In addition, ADAM LDAP and replication performance is comparable to that of Active Directory.

**ADAM Scenarios**
ADAM is designed to support a number of usage scenarios, including scenarios that are related to the following:

- Application-specific directories
- Application development
- Support for legacy applications
- Programmatic access to Microsoft Identity Integration Server (MIIS) 3.0
- Extranet access management

**Application-specific Directories**
In the application-specific directory scenario, a directory-enabled application (such as a portal application) requires an LDAP data store for application-specific data that is associated with application users that reside in Active Directory. Because this application-specific data requires schema extensions that the organization does not want to make to Active Directory, this application-specific data store must reside outside Active Directory. In this scenario, the directory-enabled application can use

Active Directory for authentication of users and for service publication while it uses ADAM as a store for the application-specific data that is associated with each user. If the application requires other application-specific data, such as policy and management information, this data can also be stored in ADAM. By using the user principals in Active Directory for authentication and for controlling access to objects in ADAM, the organization can achieve single sign-on (SSO). This design helps avoid the need for each ADAM directory to have its own user database, which prevents a proliferation of user IDs and passwords for end users every time that a new directory-enabled application is introduced on the network. The following figure illustrates an ADAM application-specific directory scenario.

**ADAM Application-specific Directory Scenario**




The ADAM features that are described in the following sections play an important role in application-specific directory scenarios.

**Storage of Application-specific Data**

An application can use ADAM to store "private" directory data, which is relevant only to the application, in an application-specific directory service — perhaps on the same server as the application — without requiring any additional configuration to Active Directory. The application-specific data, which does not need to be widely replicated, is stored solely in the ADAM directory that is associated with the application. This design has an added benefit of reducing replication traffic on the network between domain controllers. In addition, because ADAM is based on the same database engine as Active Directory — the Extensible Storage Engine (ESE), the ADAM data store scales to millions of objects and gigabytes of storage.

**Tailored Schema**

Business data is often organized uniquely for a particular application. ADAM is useful when an organization must use a directory-enabled application that is based on a unique business logic or schema. For example, an application developer may want to expose application data through LDAP in a way that is specific to the clients of the application. The application requires a different schema than the existing Active Directory installation for the server operating system. Without altering the configuration of the Active Directory, the organization is able to provide a directory-enabled application with a tailored schema that meets the organization's business needs, data requirements, and workflow processes. The workflow and associated data that is represented by the application directory can then be stored in the most intuitive way possible, independent of the existing directory structure. In addition, ADAM helps avoid schema conflicts by providing organizations with a separate and independent directory service for each directory-enabled application.

**Local Management in the Organization**

In many organizations, departments implement applications that are relevant to their operations. Because these applications are specific to a department, the information that is stored by each application may not be relevant to the rest of the organization. Each department may need to manage its own local directory service independently. For example, a department may require replication schedules that differ from the replication schedules that are offered by the enterprise directory. ADAM supports this requirement by enabling each department to manage its own directory service independently.

**Optional Centralized Management**

An ADAM instance can be deployed on a locally administered, departmental server or on a centrally managed server, offering a choice of distributed or centralized management. A centrally administered server can host multiple, independently configured instances of ADAM on the same computer, which provides greater efficiency through server consolidation and ease of management. Each separate instance of ADAM can be used by a different directory-enabled application.

**Windows NT 4.0 Domains**

You can install ADAM on a computer running Windows Server 2003 or Windows XP Professional that is joined to a Windows NT 4.0 domain. Any directory-enabled application that can run on ADAM can also run in a Windows NT 4.0 domain. Because ADAM works with any Windows-integrated security infrastructure, the Windows NT 4.0 domain can be used to authenticate domain users for ADAM.

**Application Development**

Developers can use ADAM to prototype an application for Active Directory. ADAM uses the same programming model as Active Directory and provides a similar administration experience. Developers can work with a local instance of ADAM on a developer workstation and then move the application to Active Directory.

**Active Directory Application Mode Setup Wizard**

One feature that enhances the use of ADAM for application development scenarios is the Active Directory Application Mode

Setup Wizard. ADAM provides a simple setup wizard that requires little manual input. The setup can easily be scripted, which is beneficial for unattended and silent installations as part of application installation. Developers can install and uninstall ADAM or use multiple instances of ADAM during the development phase. No reboot is required during setup.

**Locally Installed, Locally Managed**

Consider a scenario in which an application developer is developing a directory-enabled application. Having the directory on a server or domain controller requires the organization to provide a server-class development server. This server must be managed, and installing or uninstalling the directory can affect other users. ADAM does not have to be installed on a server or domain controller; it runs effectively on client computers. Therefore, you can deploy an application without using the resources of a dedicated information technology (IT) professional. ADAM runs on the following:

- Computers running Windows XP Professional
- Domain controllers and servers running:
    - Windows Server 2003, Standard Edition
    - Windows Server 2003, Enterprise Edition
    - Windows Server 2003, Datacenter Edition

**Support for Legacy Applications**

In a legacy application scenario, suppose that your organization has an established directory using the X.500-style attribute value assertions of o=<organization> and c=<country>, and the organization relies on this directory to serve legacy applications. In addition, the organization wants to migrate to Active Directory. In this scenario, to enable migration to Active Directory, you can deploy ADAM to service the applications that rely on X.500-style naming, which Active Directory does not support. You can deploy Active Directory in the organization to provide a shared security infrastructure, while ADAM provides direct support for the legacy applications. As applications are rewritten to avoid the use of o= and c=, the applications can be migrated to Active Directory. You can use a metadirectory application, such as MIIS 3.0, to synchronize data from ADAM to Active Directory automatically for a seamless migration experience. The following figure illustrates an ADAM legacy application scenario.

**ADAM Legacy Application Scenario**




**Programmatic Access to MIIS 3.0**

In a scenario of programmatic access to MIIS 3.0, ADAM provides a read/write interface to MIIS 3.0. MIIS 3.0 synchronizes data between multiple data stores, including LDAP directory services, SQL databases, and other data stores. MIIS 3.0 stores its data in SQL Server 2000. By design, the data that is held by MIIS 3.0 is not available for modification; the data can only be updated when synchronization with one of the connected data sources occurs. You can use ADAM to enable LDAP programmatic access to MIIS 3.0 data. First, the data from an installation of MIIS 3.0 is synchronized to an existing ADAM instance. The data in ADAM can then be modified programmatically through one of the LDAP programmatic interfaces to ADAM. After the data has been updated in ADAM, the data can be resynchronized with MIIS 3.0, updating the data that is held by MIIS 3.0. The following figure illustrates an ADAM scenario of programmatic access to MIIS 3.0.

**ADAM Scenario Featuring Programmatic Access to MIIS 3.0**




**Extranet Access Management**

In an extranet access management scenario, an organization uses a Web portal application (such as OpenNetwork DirectorySmart or Netegrity SiteMinder) to handle extranet access management. The Web portal application stores the authorization information as data in an LDAP directory and uses the directory for authentication purposes only. In such a scenario, ADAM can be used as the LDAP directory store to provide simple authentication support for the Web portal application. Because ADAM can host user objects that are not Windows security principals but are authenticated with LDAP simple binds, all the user information, as well as authorization data for these applications, can be stored in ADAM. This configuration works equally well in heterogeneous environments, where Active Directory is not present as the infrastructure provider. In this case, the Web clients are serviced by the portal application, which can be running on any platform, while ADAM is used as a simple LDAP store. The following figure illustrates an ADAM extranet access management scenario.

**ADAM Extranet Access Management Scenario**




**Related Information**

The following resources contain additional information that is relevant to this section:

- How Active Directory Application Mode Works

# How Active Directory Application Mode Works

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- ADAM Architecture
- ADAM Protocols
- ADAM Interfaces
- ADAM Physical Structure
- ADAM Logical Structure
- ADAM Security
- ADAM Processes and Interactions
- Network Ports Used by ADAM
- Related Information


Active Directory Application Mode (ADAM) is a new mode of Active Directory that is designed specifically for directory-enabled applications. ADAM is a Lightweight Directory Access Protocol (LDAP) directory service that runs as a user service, rather than as a system service. You can run ADAM on servers and domain controllers running operating systems in the Windows Server 2003 family (except for Windows Server 2003, Web Edition) and also on client computers running Windows XP Professional.

ADAM does not require the deployment of domains or domain controllers. You can run multiple instances of ADAM concurrently on a single computer, with an independently managed schema and independently managed data for each ADAM instance.

The ideal environment for ADAM includes the following:

- A computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows XP Professional and containing hardware that meets the minimum hardware requirements of each respective operating system
- For environments consisting of multiple, replicating ADAM instances, the presence of a fully functioning replication topology
- A regular backup schedule

This section describes the fundamental elements of ADAM, including its architecture, protocols, interfaces, physical structure, logical structure, security concepts, processes and interactions, and network ports.

**ADAM Architecture**

The ADAM architecture consists of several components that work together to provide directory services. These components include the following:

- Interfaces: LDAP and replication (REPL)and management interface
- Directory System Agent (DSA)

- Database layer
- Extensible Storage Engine (ESE)
- Directory database

The ADAM architecture is illustrated in the following figure.

**ADAM Architecture**




The following table describes the ADAM components.

**ADAM Components**

| Component | Description |
|---|---|
| Interfaces (LDAP and REPL) | The data store interfaces provide a way for directory clients and other directory servers to communicate with the data store. |
| DSA | The DSA, which runs as Adamdsa.dll on each ADAM instance, provides the interfaces through which directory clients and other ADAM instances gain access to the directory database. In addition, the DSA enforces directory semantics, maintains the schema, guarantees object identity, and enforces data types on attributes. |
| Database layer | The database layer is an application programming interface (API) that resides in Adamdsa.dll and provides an interface between applications and the directory database to protect the database from direct interaction with applications. Calls from applications are never made directly to the database; they go through the database layer. In addition, because the directory database is flat — with no hierarchical namespace — the database layer provides the database with the abstraction of an object hierarchy. |
| ESE | The ESE, which runs as Esent.dll, manages the tables of records — each with one or more columns — that make up the directory database. |
| Directory database | The data store stores directory information in a single database file called Adamntds.dit. |

**DSA**

The DSA runs on an ADAM instance as Adamdsa.dll, and it provides access to the directory database. On computers running multiple instances of ADAM, a single, shared copy of Adamdsa.dll supports all ADAM instances.

**Note**

- Multiple instances of ADAM running on a single computer share all ADAM-related binaries. Therefore, a hotfix that is applied to one ADAM instance affects all ADAM instances running on that computer.

The DSA runs as a user service, in the security context of the service account under which the ADAM instance is running. Clients can use one of the supported interfaces to connect (bind) to the DSA and then search for, read, and write to ADAM objects and their attributes.

An object in the configuration partition of the directory, the NTDS Settings object (with a distinguished name of CN=NTDS Settings,CN=*computername*$*instancename*,CN=Servers,CN=Default-First-Site-Name,CN=Sites,CN=Configuration,CN={*GUID*}), represents the DSA. This object contains configuration information about the DSA.

In addition to providing the interfaces through which directory clients gain access to directory data, the DSA provides the following functionality.

**Object identification**
Every object in ADAM has a permanent globally unique identifier (GUID) and an object name. The object name can be changed; all permanent references to the object are kept in terms of the GUID. The object name is used for hierarchy navigation and display purposes. The DSA maintains the GUID association with an object when the object's name changes, for example, when the object is moved to a different folder.

**Schema enforcement**
The DSA ensures that data in the directory adheres to the data definitions that are provided by the directory schema. The schema is the set of rules that determines what kind of data the directory can hold.

### Note

- If an update does not produce a conflict with the schema at the originating replica, the update is considered acceptable at all replicas. Therefore, replicated updates do not perform schema checks, and you do not have to wait until the schema replicates before creating instances of a new object or attribute.

**Access control enforcement**
The DSA reads security IDs (SIDs) on the access token to identify the current user and the user's group memberships, and it enforces access control based on the identity and group memberships of the user.

**Support for replication**
The API that is called to initiate replication is implemented in the DSA.

**Referrals**
The DSA manages directory hierarchy information (referred to as "knowledge"), which it receives from the database layer. The DSA is responsible for cross-references of partition objects.

**DSA GUID and Invocation ID**
Both the DSA and the ADAM directory database are represented uniquely and have their own respective GUIDS. The DSA GUID is the GUID of the NTDS Settings object (class **nTDSDSA**). The value of the DSA GUID is stored in the **objectGUID** attribute of the NTDS Settings object of the ADAM server object, which resides in the Sites container in the configuration directory partition (with a distinguished name of CN=NTDS Settings,CN=*computername*$*instancename*,CN=Servers,CN=Default-First-Site-Name,CN=Sites,CN=Configuration,CN={*GUID*}).

The DSA GUID is created when the ADAM instance is initially installed. The DSA GUID is destroyed only if the ADAM instance is uninstalled. The DSA GUID ensures that the DSA remains recognizable when the computer on which ADAM is running is renamed. The DSA GUID is not affected by ADAM backup and restore processes.

The ADAM directory database has its own GUID, which the DSA uses to identify the database instance (the version of the database). The database GUID is stored in the **invocationId** attribute on the nTDSDSA object. Unlike the DSA GUID, which never changes for the lifetime of the ADAM instance, **invocationId** changes in two scenarios:

- During the ADAM restore process, to ensure the consistency of the replication process. The change in **invocationId** enables ADAM to distinguish changes that are made before and after a database restore.

- When an application directory partition is removed from and added back again to the ADAM instance.

**Database Layer**
The database layer is an API that resides in Adamdsa.dll and provides an object view of the directory database, making the data accessible to the DSA as a set of hierarchical containers. By applying schema semantics to database records, the database layer isolates the upper components of the directory service from the underlying database system. The database layer is an internal interface. No database access calls are made directly to the ESE; instead, all database access goes through the database layer.

In the directory database, each object is identified by its relative distinguished name, which is unique in the object's parent container. The relative distinguished name and the chain of successive parent object names make up the object's distinguished name. The database stores the relative distinguished name for each object, as well as a reference to the parent object. The database layer follows these parent references and concatenates the successive relative distinguished names to form distinguished names, thereby defining the object hierarchy.

The database layer is also responsible for the creation, retrieval, and deletion of individual records (objects), attributes within records, and values within attributes. To carry out these functions, the database layer uses the schema cache (an in-memory structure in the DSA) to get information about the attributes that it needs.

**ESE**

The ESE is a Windows component that is used by ADAM, as well as by several other Windows components, as an interface to the data that is stored in an indexed and sequential access method (ISAM) database. (The ADAM directory database is an ISAM database.) The ESE is responsible for indexing the data in the database file and for transferring the data in and out of the database. The purpose of the ESE is to enable applications to store and retrieve data by using the ISAM. The ESE provides applications with a consistent data state by means of transacted data update and retrieval. A crash recovery mechanism maintains data consistency, even in the event of a system crash. Transactions in the ESE are highly concurrent, making the ESE suitable for server applications. The ESE caches data intelligently to ensure high-performance access to data. In addition, the ESE is resource efficient, which makes it suitable for applications that perform auxiliary roles.

The version of the ESE that is used with ADAM is implemented in Esent.dll. Each ADAM instance running on a computer shares a single a copy of Esent.dll, which is located in *windir*\ADAM.

The following characteristics of the ESE make it well suited to the storage needs of ADAM. The ESE:

- Supports databases of up to 16 terabytes (TB) in size, and it can hold many millions of objects per instance.
- Supports indexing.
- Supports multivalue attributes.
- Supports update operations that are transacted for stability and integrity across system failures.
- Can be backed up while the ADAM instance is online.
- Handles sparsely populated objects well; that is, space in the database is not reserved for attributes that do not have values.

The ADAM schema defines all the attributes that are required and allowed for a given object. However, the ESE reserves storage only for the space that is used — that is, only for the attributes that have values, not for all possible attributes. For example, if a user object has 50 attributes defined in the schema and you create a user with values for only 4 attributes, storage space is allocated only for those 4 attributes. If more attributes are added later, more storage is allocated for them.

The ESE implements the search and retrieval functionality of the underlying database. Also, the ESE is able to store attributes that can have multiple values. For example, the database can store multiple phone numbers for a single user without requiring a different phone number attribute for each phone number.

The ESE provides transactional views of the database. The cost of providing these views is that any object that is modified in a transaction has to be temporarily copied so that two views of the object can be provided: one to the thread inside that transaction and one to threads in other transactions. This copy must remain as long as any two transactions in the process have different views of the object. The repository that holds these temporary copies is called the version store. Because the version store requires contiguous virtual address space, it has a size limit. If a transaction is open for a long time while changes are being made (either in that transaction or in others), eventually the version store can be exhausted. At this point, no further database updates are possible.

**ADAM Protocols**

The primary protocol that ADAM uses is LDAP, which runs on top of TCP/IP. In addition, the data store uses remote procedure call (RPC) for replication and management functions.

The following table describes the protocols that are used by ADAM.

### ADAM Protocols

| Protocol | Description |
|---|---|
| LDAP | LDAP is a directory service protocol that specifies directory communications. It runs directly over TCP/IP, and it can also run over User Datagram Protocol (UDP) connectionless transports. Clients can use LDAP to query, create, update, and delete information that is stored in a directory service over a TCP connection through the TCP default port 389. Active Directory supports LDAP v2 (Request for Comments (RFC) 1777, "Lightweight Directory Access Protocol") and LDAP v3 (RFC 2251, "Lightweight Directory Access Protocol v3"). LDAP v3 is an industry standard that can be used with any directory service that implements the LDAP protocol. LDAP is the preferred and most common way of interacting with Active Directory.<br><br>Historically, LDAP is a simplified ("lightweight") version of Directory Access Protocol (DAP), which is the original protocol that was used to interact with X.500 directories. X.500 defines an earlier set of standards that was developed by the International Organization for Standardization (ISO). LDAP is simpler than DAP in two key ways:<br><br>• Rather than using its own protocol stack according to the Open Systems Interconnection (OSI) networking model, LDAP communicates over Internet Protocol (IP) by using either UDP or TCP.<br><br>• LDAP syntax is easier to use than DAP syntax.<br><br>For these reasons, LDAP is widely used and accepted as the standard protocol for directory service access.<br><br>The following key aspects of LDAP characterize it:<br><br>• The protocol is carried directly over TCP for connection-oriented transport (receipt of data is acknowledged) and over UDP for connectionless transport (sent or received data is not acknowledged).<br><br>• Most protocol data elements can be encoded as ordinary strings, for example, as distinguished names.<br><br>• Referrals to other servers can be returned to the client.<br><br>• Simple Authentication and Security Layer (SASL) mechanisms can be used with LDAP to provide associated security services.<br><br>• Attribute values and distinguished names can be internationalized through the use of the ISO 10646 character set.<br><br>• The protocol can be extended to support new operations, and controls can be used to extend existing operations.<br><br>• The schema is published through an attribute on the directory root object (rootDSE) for use by clients. |
| RPC | The REPL interface uses RPC for management tasks, including replication and management functions, that are not well suited to the LDAP protocol. RPC is a powerful, robust, efficient, and secure interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality that resides in a different process. That different process can be on the same computer, on the local area network (LAN), or across the Internet. |

**ADAM Interfaces**

As shown in the "ADAM Architecture" figure earlier in this section, network clients and other directory servers obtain access to ADAM through one of the interfaces that are described in the following table.

### ADAM Interfaces

| Interface | Description |
|---|---|
| LDAP | The primary interface for ADAM. Directory clients use LDAP v3 or LDAP v2 to connect to the DSA through the LDAP interface. The LDAP interface is part of Wldap32.dll. LDAP v3 is backward compatible with LDAP v2. |

| | |
|---|---|
| REPL | Provides functionality for finding data about ADAM instances, converting the names of network objects between different formats, manipulating service principal names (SPN) and DSAs, and managing replication of ADAM instances. |
| System.DirectoryServices<br><br>(ADSI for .NET Framework) | A namespace in the Microsoft .NET Framework that provides simple programming access to LDAP directories, such as ADAM. This API requires the .NET Framework to be installed. |
| Active Directory Service Interfaces (ADSI) (Ads*.dll) | A set of Component Object Model (COM) interfaces that abstracts the capabilities of directory services from different network providers (such as LDAP) in a distributed computing environment to present a single set of directory service interfaces for managing network resources. |
| Directory Services Markup Language (DSML) | Provides a means of representing directory structural information and directory operations as an XML document. |

**ADAM Physical Structure**

The ADAM physical structure consists of the following:

- Computers running ADAM (ADAM servers)
- ADAM instances
- Configuration sets
- Sites

**Note**

- Because ADAM is based on the same code as Active Directory, the service and scalability limits of ADAM are comparable to Active Directory. Therefore, ADAM supports a comparable number of sites as Active Directory.

**ADAM Servers**

A computer running ADAM is called an ADAM server. ADAM servers can run on any of the following operating systems:

- Windows Server 2003, Standard Edition
- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition
- Windows XP Professional

When you run ADAM on a computer running Windows XP Professional, certain considerations apply. The following table lists these considerations.

**Running ADAM on Windows Server 2003 and on Windows XP Professional**

| ADAM Feature | Windows Server 2003 | Windows XP Professional |
|---|---|---|
| Support for multiple processors | Supported. | Not supported. On multiple-processor computers, ADAM uses only one process at a time to prevent resource monopolization. |
| Maximum cache value | No limit. | Limited to one half of physical RAM. |

| | | |
|---|---|---|
| Data and log file locations | Data and log files can reside on different drives. | Data and log files must reside on the same drive. |
| Password policies | Supported. | Not supported. |
| Auditing | Supported. For more information, see “Auditing” later in this section. | Not supported. |
| ADAM service account requirements | See “Service Account” later in this section. | See “Service Account” later in this section. |
| Stopping and restarting an ADAM instance | You can stop and then restart an ADAM instance immediately. | If an LDAP client is connected when you shut down the ADAM instance, you may need to wait as long as two minutes before restarting the ADAM instance. This time is required to allow any existing connections to clear. |
| Backing up ADAM | Multiple ADAM instances running on the same computer can be backed up simultaneously. | For computers running multiple ADAM instances, only one ADAM instance can be backed up at a time. |
| Full ADSI functionality | No hotfixes are required. | For complete ADSI functionality with Windows XP Professional, you must install hotfix 817583. For more information about this hotfix, see article 817583, “Active Directory Services Does Not Request Secure Authorization Over an SSL Connection,” in the Microsoft Knowledge Base. |

**Using ADAM on Windows XP Professional as a Personal Data Store**
On computers running Windows XP Professional, ADAM has been designed primarily as a personal data store. It is recommended that you run ADAM on Windows Server 2003 if you need scalability to support a growing number of simultaneously connected clients or if you intend to use ADAM on a centralized directory server.

**Running ADAM on Windows XP Professional in a Workgroup**
When you install ADAM on a computer running Windows XP Professional that is joined to a workgroup, you must set the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\forceguest** to **0**. (The default is **1**.) Otherwise, users connecting to ADAM over the network are forced to a security context of Guest, and binds to ADAM fail.

**ADAM Instances**
Each time that you install ADAM, you create a unique ADAM instance. Each ADAM instance consists of the following:

- Program files
- Data files
- Registry keys
- Service name
- Service account
- Event log
- Name:port

**Program Files**
After you install ADAM, the program files and tools of an ADAM instance reside in the *windir*\ADAM directory. If you install multiple instances of ADAM on a single computer, the program files and tools in the *windir*\ADAM directory are shared by all ADAM instances running on that computer.

**Note**

- If you apply a hotfix to one instance of ADAM, all ADAM instances on the same computer are affected.

**Data Files**

The data files for a given ADAM instance are installed in the directory Program Files\Microsoft ADAM*instancename*\Data, where *instancename* is the name that you provide for the ADAM instance during installation. If you install multiple instances of ADAM on a single computer, each instance that you install has a separate data directory. For example, if you install two ADAM instances named, respectively, instance1 and instance2 on a computer, the computer has the following ADAM data directories:

Program Files\Microsoft ADAM\instance1\Data

Program Files\Microsoft ADAM\instance2\Data

**Registry Keys**

Registry keys for ADAM are stored in the following location:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_instancename**

For a detailed explanation of the ADAM registry keys, see "Active Directory Application Mode Tools and Settings."

**Service Name**

During setup, you assign a name to the ADAM instance, and that name is used in the creation of the file directory structure and registry keys for ADAM. In addition, the name that you assign is used to create the service name, service display name, and service description, as shown in the example in the following table.

**ADAM Service Name**

| Name Supplied During Setup | Service Name | Service Display Name | Service Description |
|---|---|---|---|
| Instance1 | ADAM_instance1 | Instance1 | Blank (no default description provided) |

The name that you specify for an ADAM instance during setup must meet the following requirements:

- It must be unique with respect to other ADAM instances running on the same computer.
- It must be no longer than 44 characters.
- It must use characters only from the ranges of a through z, A through Z, or 0 through 9.
- The name "ntds" cannot be used.

The service display name appears in Add or Remove Programs and in the Services snap-in in Microsoft Management Console (MMC). The service description appears in the Services snap-in. You can modify the service display name and the service description anytime after installation by using the **sc** command-line tool.

**Service Account**

Each ADAM instance runs as a user service in the security context of the service account that is specified for that instance. The service account for an ADAM instance is specified at installation, and it should be modified only with the **dsmgmt** command-line tool.

**Note**

- Change the ADAM service account only with the **dsmgmt** command-line tool. Changing the ADAM service account with the Services snap-in will cause errors.

The type of ADAM service account that you can choose for an ADAM instance depends on the Windows workgroup or domain environment in which you install the ADAM instance and on whether the ADAM instance participates in a configuration set.

The Active Directory Application Mode Setup Wizard helps you choose a service account for ADAM by only allowing you to choose a service account that is valid for a particular configuration. The following table shows ADAM service accounts that are valid for a given environment.

**ADAM Service Accounts**

| Environment | First ADAM Instance Service Account | Replica ADAM Instances Service Account |
|---|---|---|
| Workgroup (that is, no domain or forest) | Network service | Replica ADAM instances not allowed |
| Workgroup (that is, no domain or forest) | Workstation user | Workstation user |
| Windows 2000 Server or Windows Server 2003 domain or forest | Network service | Network service or domain user |
| Windows 2000 Server or Windows Server 2003 domain or forest | Workstation user | Workstation user |
| Windows 2000 Server or Windows Server 2003 domain or forest | Domain user | Domain user or network service |
| Windows NT 4.0 domain | Workstation user | Workstation user |
| Windows NT 4.0 domain | Domain user | Domain user |

**Note**

- When a workstation (local) user account is used on the first ADAM instance in a configuration set, all subsequent ADAM instances in the same configuration set must use an identical local workstation account name and password as the ADAM service account.

**Event Log**

Each ADAM instance that is installed on a computer writes events to the event log. These events appear in Event Viewer under **ADAM (***instancename***)**, where *instancename* is the name that you assigned to the ADAM instance.

**Name:Port**

For a directory-enabled application to communicate with an ADAM instance, the application must specify the network basic input/output system (NetBIOS) name, Domain Name System (DNS) name, or IP address of the computer on which ADAM is running. In addition, the application must specify the LDAP or Secure Sockets Layer (SSL) communication port that is being used by the ADAM instance. For example, a directory-enabled application that references an ADAM instance running on a computer with a DNS name of ADAMSrv1.microsoft.com — with a NetBIOS name of ADAMSrv and an IP address of 192.168.0.1 and using communication port 389 for LDAP connections — can address the ADAM instance by using one of the following:

- ADAMSrv1.microsoft.com:389
- ADAMSrv:389
- 192.168.0.1:389

**Configuration Sets**

ADAM instances replicate data based on participation in a configuration set. A configuration set is a group of ADAM instances that replicate a common schema partition and configuration partition. (For more information about partitions, see "Directory Partitions" later in this document.) ADAM instances in a configuration set can also replicate any number of application directory partitions. ADAM instances in a configuration set are not required to replicate all application directory partitions in the configuration set; a single ADAM instance can replicate all — or any subset of — the application directory partitions in its configuration set. An ADAM instance cannot, however, replicate an application directory partition from a different configuration set.

The following figure shows an example of two ADAM configuration sets, each with two ADAM instances. As shown in the figure, a single computer can run multiple ADAM instances, each in a different configuration set.

**ADAM Configuration Sets**




**Creating Configuration Sets**

When you install an ADAM instance, you can install a unique instance or you can install a replica of an existing instance, as shown in the following figure, which illustrates the **Setup Options** page in the Active Directory Application Mode Setup Wizard.

**Active Directory Application Mode Setup Wizard Setup Options**




The following table describes each of these options as they relate to configuration sets.

**Setup Options**

| Setup Option | Meaning |
|---|---|
| **A unique instance** | Choosing this setup option starts a new configuration set. |
| **A replica of an existing instance** | Choosing this option causes the ADAM instance being installed to join to an existing configuration set. |

**Note**

- An ADAM instance can be joined to a configuration set only during installation of the instance. After an ADAM instance is created, it cannot be added to or removed from a configuration set.

**Sites**

Like Active Directory, ADAM uses topology information, which is stored as site objects and site link objects in the configuration directory partition, to build the most efficient replication topology for a configuration set. You can use sites in ADAM to represent the physical structure, or topology, of your network. You can use ADAM ADSI Edit to define site objects and site link objects.

### Note

- For specific instructions about creating site objects, site link objects, and setting site link costs, see the ADAM Administrator's Guide. To view the ADAM Administrator's Guide, after you install ADAM click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Help**.

**Site Objects**
You can create a site object in ADAM to represent each area of your network (typically, a building or a group of buildings) in which all computers are connected by high-speed bandwidth. You can then move the directory objects of ADAM instances that are located in each area of the network into the corresponding site object. By default, the directory objects of each ADAM instance that you create belong to a single default site (CN=Default-First-Site-Name,CN=Sites,CN=Configuration,CN={*GUID*}).

**Site Link Objects**
You can use site link objects in the directory to represent low-speed bandwidth connections between the sites on your network. Site link objects help you tune replication on your network; you can use them to determine the relative cost of replication across each site link. By default, the default site (CN=Default-First-Site-Name) belongs to the default site link object (CN=DEFAULTIPSITELINK,CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,CN={*GUID*}). You can add new sites that you create to the default site link object. You can also create additional site link objects to represent each of the actual site links on your network.

**Site Link Costs**
Each site link object in the directory has an associated "cost" that represents the relative preference that ADAM gives to one site link compared to another. You can edit the cost attribute on a site link object to specify the relative cost of replication across that site link. ADAM takes into consideration site link costs when it builds the replication topology. ADAM considers a site link with a higher cost to be less desirable than a site link with a lower cost. When multiple paths for replication are available, ADAM uses the site link with the lower cost.

**ADAM Logical Structure**
The ADAM logical structure consists of the following:

- Directory database
- rootDSE
- Namespaces
- Directory partitions
- Schema

**Directory Database**
All data that is stored by an ADAM instance is held in the directory database. The ADAM directory database is organized into two or more partitions. ADAM stores directory data in a hierarchical directory store. By default, the directory database for a given ADAM instance is called Adamntds.dit, and it is located in the following directory:

ProgramFiles\Microsoft ADAM*instancename*\Data\Adamntds.dit

**rootDSE**
At the root of the ADAM directory tree is a DSA-specific entry (DSE), which is not part of any directory partition. The rootDSE represents the top of the logical namespace for one ADAM instance. The rootDSE attributes contain information about the ADAM instance, including its capabilities and configuration. The rootDSE is visible, by design, to anonymous users. (The rootDSE is the *only* object in ADAM that is visible by default to anonymous users.)

### Note

- Some non-Microsoft security scanners report an issue regarding the ability of anonymous users to read the rootDSE. However, the rootDSE is visible to anonymous users by design, and it does not represent a security issue.

There is only one root for a given directory instance, but the information that is stored in the root is specific to the ADAM instance to which you connect. Among other things, the attributes of rootDSE identify the directory partitions (the application, schema, and configuration directory partitions) that are specific to one ADAM instance. In this way, the rootDSE provides a

"table of contents" for a given ADAM instance.

The rootDSE is always accessible for queries through an anonymous connection.

#### rootDSE Operational Attributes

The rootDSE publishes information through attributes on the rootDSE object. RFC 2251 and RFC 2252, "Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions," define a set of seven rootDSE operational attributes that LDAP v3 servers are expected to publish, and these attributes are described in the following table. All LDAP servers recognize these attribute names, but when an attribute corresponds to a feature that the server does not implement, the attribute is absent.

**rootDSE Operational Attributes**

| Attribute | Description |
|---|---|
| **supportedLDAPVersion** | LDAP versions that the server implements. Domain controllers running Windows Server 2003 and Windows 2000 Server support LDAP v2 and LDAP v3. |
| **namingContexts** | Naming contexts (directory partitions) that this server masters (stores as a writable replica) or shadows (stores as a read-only replica). A client uses this attribute to choose suitable base objects for searching when the client contacts a server. |
| **subschemaSubentry** | The name of a subschema entry, which is used to administer information about the schema, in particular, the object classes and attribute types that are supported. |
| **altServer** | Uniform Resource Locators (URLs) of other servers that can be contacted when this server becomes unavailable. If the server cannot locate any other servers, this attribute is absent. Clients can cache this information in case their preferred LDAP server becomes unavailable. This attribute is absent by default for Active Directory servers. |
| **supportedExtension** | Object identifiers that identify the extended LDAP operations that the server supports. LDAP extensions include such operations as different query types, paging operations, and sorting methods, which each have a different object identifier. This attribute is absent by default for Active Directory servers. |
| **supportedControl** | Object identifiers that identify the LDAP controls that the server supports. If the server does not support any controls, this attribute is absent. Server controls extend LDAP functionality; examples include a control to move objects across domains and a control to delete an entire subtree of a container object. |
| **supportedSASLMechanisms** | The names of the SASL mechanisms that the server supports. SASL is a standard for negotiating an authentication mechanism and, optionally, an encryption mechanism to provide client authorization for accessing Active Directory. By default, Generic Security Service API (GSSAPI) is supported. |

**Note**

- A new extended LDAP operation on ADAM instances and domain controllers running Windows Server 2003 enables client refresh of a dynamic entry in the directory. Object identifier = 1.3.6.1.4.1.1466.101.119.1 is defined and published in the **supportedExtension** attribute of the rootDSE object.

#### rootDSE Informational Attributes

In addition to the previous operational attributes, ADAM also supports the informational attributes in the following table.

**rootDSE Informational Attributes**

| Attribute | Description |
|---|---|
| **configurationNamingContext *** | The naming context (directory partition) for the forest Configuration container |
| **currentTime *** | The current time in the generalized time format |

| | |
|---|---|
| **defaultNamingContext *** | The default naming context (directory partition) for a particular ADAM instance as set in the **msDS-defaultNamingContext** attribute on the NTDSA object at:<br><br>(CN=NTDS Settings,CN=*computername*$*instancename*,CN=Servers,CN=Default-First-Site-Name,CN=Sites,CN=Configuration,CN={*GUID*})<br><br>Setting a value for **msDS-defaultNamingContext** is optional. |
| **dnsHostName *** | The DNS name of this ADAM instance |
| **domainControllerFunctionality *** | The numeric level of Active Directory functionality for the ADAM instance |
| **dsaVersionString** | The version stamp of Adamdsa.dll. This is visible only to administrators. |
| **dsSchemaAttrCount** | The total number of attributes that are defined in the schema |
| **dsSchemaClassCount** | The total number of classes that are defined in the schema |
| **dsSchemaPrefixCount** | The total number of prefixes in the schema |
| **dsServiceName *** | The distinguished name of the NTDS settings object that represents this ADAM instance in the configuration directory partition |
| **forestFunctionality *** | The Windows forest functional level. With Active Directory:<br><br>• 0 equals Windows 2000 functional level<br><br>• 1 equals Windows Server 2003 interim functional level<br><br>• 2 equals Windows Server 2003 functional level<br><br>With ADAM, **forestFunctionality** is set to 2 by default. |
| **highestCommittedUSN *** | The highest update sequence number (USN) that is committed to the database on this ADAM instance |
| **isSynchronized *** | A Boolean indicator for whether the ADAM instance has completed its initial synchronization with replica partners |
| **msDS-PortLDAP** | The LDAP port used by the ADAM instance |
| **msDS-PortSSL** | The SSL port used by the ADAM instance |
| **msDS-PrincipalName** | The name of the current user. This is visible only to authenticated users. |
| **msDS-ReplAllInboundNeighbors**<br><br>**msDS-ReplAllOutboundNeighbors**<br><br>**msDS-ReplConnectionFailures**<br><br>**msDS-ReplLinkFailures**<br><br>**msDS-ReplPendingOps**<br><br>**msDS-ReplQueueStatistics** | Attributes that make replication information available through LDAP. By using these attributes, you can obtain information with LDAP that otherwise must be obtained with RPC and the DsReplicaGetInfo API. For more information about DsReplicaGetInfo, see DsReplicaGetInfo in the Microsoft Platform SDK on MSDN. |
| **msDS-TopQuotaUsage** | A sorted list (highest to lowest) of the top 10 quotas that are used across all directory partitions that are held by the ADAM instance |

| | |
|---|---|
| **pendingPropagations** | The number of work items in the security descriptor propagator queue |
| **schemaNamingContext *** | The naming context (directory partition) for the forest schema |
| **serverName *** | The fully qualified distinguished name for this ADAM instance |
| **serviceAccountInfo** | Contains details about the service account that is used by this ADAM instance. This service account information is used during replica promotions to detect invalid configurations (configurations for which replication will not work). This is visible only to ADAM administrators. |
| **spnRegistrationResult** | The most recent error that relates to service principal name (SPN) registration. This is visible only to ADAM administrators. |
| **supportedCapabilities *** | The object identifier values (including 1.2.840.113556.1.4.1851 and 1.2.840.113556.1791) that indicate the additional capabilities of an ADAM instance, such as support for LDAP policies |
| **supportedConfigurableSettings** | The list of configurable settings that are supported by the ADAM instance |
| **supportedLDAPPolicies *** | Supported LDAP administrative query policies, such as the maximum size of a result set (**MaxResultSetSize**) and maximum page size (**MaxPageSize**) |
| **tokenGroups** | The current user's individual and group SIDs |
| **validFSMOs** | Operations master roles (also known as flexible single master operations (FSMO) roles) that are held by the ADAM instance. Valid operations master roles for ADAM are the schema operations master role and the domain naming operations master role. |

*Default attributes (All other attributes must be explicitly requested.)

**Note**

- You can retrieve an active users individual and group SIDs by explicitly querying the **tokenGroups** attribute on the rootDSE.

**Namespaces**

Each ADAM directory partition has its own unique distinguished name. Unlike Active Directory, which supports only DNS-style (DC=) names for top-level directory partitions, ADAM supports both DNS-style and X.500-style names for top-level directory partitions. However, if you develop an application for ADAM but later intend to migrate the application to Active Directory, use only DC= naming components in directory partition names. Using only DC= naming components helps to prevent compatibility problems when you migrate the application from ADAM to Active Directory.

For directory partition naming, ADAM supports the distinguished name components that are listed in the following table.

**ADAM Directory Partition Distinguished Name Components**

| Distinguished Name Component | Description |
|---|---|
| C= | Country/region |
| CN= | Common name |
| DC= | Domain component |
| L= | Location |
| O= | Organization |
| OU= | Organizational unit |

**OUs in a Namespace**

Organizational Unit containers can be created only in certain types of parent objects. Therefore, if you want to create organizational units (OUs) in a given ADAM directory partition, consider this limitation when you name the directory partition.

OUs can be created only under the following objects:

- OU (OU=)
- Country/region (C=)
- Organization (O=)
- Domain (DC=)

For example, you can create an OU in an application partition named O=Microsoft,C=US, but you cannot create an OU in an application partition named L=Microsoft,C=US.

**Directory Partitions**

The ADAM directory store is organized into logical directory partitions, or naming contexts. There are three different types of directory partitions: configuration, schema, and application. Each ADAM directory store must contain a single configuration directory partition and a single schema directory partition. The directory store can contain zero or more application directory partitions. The following table describes each directory partition type, its distinguished name, whether it is required or optional, and its contents.

**Directory Partition Types**

| Partition Type | Distinguished Name | Required/Optional | Contents |
|---|---|---|---|
| Configuration | CN=Configuration,CN={*GUID*} | Required | ADAM configuration information |
| Schema | CN=Schema,CN=Configuration,CN={*GUID*} | Required | Class and attribute definitions |
| Application | (Assigned during partition creation) | Optional | Application data |

**Note**

- The distinguished names of the configuration partition and the schema partition contain the GUID of the ADAM instance as a part of their distinguished names.

**Configuration Directory Partition**

The configuration directory partition holds information about ADAM replication scheduling and replica sets, information that defines the other partitions in the replication set, information about the users and groups in the replica set, and other information.

The following table describes the default containers in the configuration directory partition.

**Default Containers in the Configuration Directory Partition**

| Container | Purpose |
|---|---|
| CN=DirectoryUpdates,CN=Configuration,CN={*GUID*} | Not currently used. |
| CN=Extended-Rights,CN=Configuration,CN={*GUID*} | Stores objects of the class **controlAccessRight** that applications can use to extend standard access control. |
| CN=ForeignSecurityPrincipals,CN=Configuration,CN={*GUID*} | Stores proxy objects for security principals that are from Active Directory domains. |
| CN=LostAndFoundConfig,CN=Configuration,CN={*GUID*} | Stores configuration directory partition objects that are being created in containers that are being deleted simultaneously on other ADAM instances in the same configuration set. If, before replication, an object is created in or moved to a location that no longer exists after replication, the “lost” object is added to the LostAndFoundConfig container. |

| | |
|---|---|
| CN=NTDS Quotas,CN=Configuration,CN={*GUID*} | Stores objects of the class **msDS-QuotaControl** that contain object ownership quota assignments for the configuration directory partition. Quotas limit the number of objects that a user (including inetOrgPerson), group, computer, or service can own in a configuration directory partition or application directory partition. |
| CN=Partitions,CN=Configuration,CN={*GUID*} | Stores the cross-references to every directory partition in the configuration set, including the configuration partition, the schema partition, and all application partitions. During LDAP searches, these cross-references to directory partitions make referrals to other domains possible. |
| CN=Roles,CN=Configuration,CN={*GUID*} | Stores the default groups for a given partition. |
| CN=Services,CN=Configuration,CN={*GUID*} | Stores data for various networking services and applications. |
| CN=Sites,CN=Configuration,CN={*GUID*} | Identifies all the sites in the network, the ADAM instances in those sites, and the replication topology. The contents of this container take the form of replication transports, subnets, and the first site and site link that are created, which are called Default-First-Site-Name and Default-First-Site-Link, respectively. |

**Schema Directory Partition**

The schema directory partition contains the definitions for the types of data that the directory database can hold. The definitions in the schema partition maintain data consistency for the ADAM directory service. In addition, applications can refer to the schema partition to determine the types of data that an ADAM instance allows. You can also extend the schema so that ADAM can hold data that is specific to a particular application. For more information about schemas, see the "Active Directory Schema Technical Reference."

**Application Directory Partitions**

Application directory partitions hold the data that your applications use. You can create an application partition during ADAM setup or anytime after installation. Depending on your application, you might extend the schema manually or your application might automatically extend the schema for you. Typically, you manage data in a given application directory partition through your application. After the application directory partition is created, ADAM holds the application partition reference objects in CN=Configuration,CN=Partitions.

### Note

- With ADAM you can create users in an application directory partition. ADAM differs from Active Directory in this respect, because Active Directory does not allow the creation of users in an application directory partition.

**Schema**

The ADAM schema uses object classes and attributes to define the kinds of objects and data that can be created and stored in an ADAM directory. Each ADAM configuration set has its own independently manageable schema, which is stored in the schema directory partition. In keeping with the ADAM design concepts of simplicity and flexibility, the base (or default) ADAM schema contains only the classes and attributes that are needed to start an ADAM instance. The schema can be extended with new classes and attributes, either by administrators or by the applications themselves. In addition, unneeded schema classes and attributes can be deactivated. As with all objects in the directory, access control lists (ACLs) protect schema objects so that only authorized users can alter the schema. Every object in an ADAM directory is an instance of an object class that is defined in a schema.

**Object Classes**

An object class represents a category of objects, such as users, printers, or applications, that share a set of common characteristics. The definition for each object class contains a list of the attributes that can be used to describe instances of the class. For example, the **User** class has attributes such as **givenName**, **surname**, and **streetAddress**. Attributes for a class are divided into those attributes that an object of that class must contain and additional attributes that the object may contain. The definition of each class also lists the classes whose objects can be parents of the objects of that class.

**Attributes**

The schema defines each attribute. The definition for each attribute includes unique identifiers for the attribute, the syntax for the attribute, optional range limits for the attribute values, whether the attribute can have only one value or multiple values, and whether the attribute is indexed. The directory schema defines each attribute exactly once. Each attribute can then be referenced by multiple object classes. For example, the **description** attribute is defined once, and then it is referenced by many

object classes.

**Single-value and multivalue attributes**
Attributes can be single valued or multivalued. An instance of a single-value attribute can contain only a single value. An instance of a multivalue attribute can contain multiple values, which must all use the same syntax. Each value of a multivalue attribute must be unique.

### Note

- A multivalue attribute stores its values in random order. Therefore, you should not make any directory decisions — programmatic or manual — based on the order of the values in a multivalue attribute.

**Indexed attributes**
Indexing attributes helps to improve the performance of queries that are based on the indexed attributes. Both single-value and multivalue attributes can be indexed; however, classes cannot be indexed. Attributes can be marked for indexing through their schema definition. Indexing an attribute makes it possible for users to use wildcard symbols (*) as prefixes and suffixes when they specify a search string.

When you mark an attribute as indexed, all instances of the attribute are added to the index, not just the instances that are members of a particular class. Indexing attributes, particularly multivalue attributes, can negatively affect replication and object creation time, as well as directory database size. Therefore, you should only index commonly used attributes.

For general information about schemas, see "Active Directory Schema Technical Reference."

**Schema Extensions**
Like the Active Directory schema, the ADAM schema can be extended through the importing of .ldf files into the schema. The ADAM package contains four .ldf files that can be imported as an option on the **Import LDIF Files** page of the Active Directory Application Mode Setup Wizard. Together, these files contain several **user** class schema definitions, along with objects that can be used with Windows Authorization Manager.

The following table describes each of the optional ADAM .ldf files.

### ADAM .ldf Files

| .ldf file | User Classes | Description |
|---|---|---|
| ms-user.ldf | **Person**<br>**Organizational-Person**<br>**User** | Used to create user objects in the ADAM directory. |
| ms-inetorgperson.ldf | **Person**<br>**Organizational-Person**<br>**User**<br>**inetOrgPerson** | Used to create user objects in the ADAM directory of the **inetOrgPerson** class (as defined in RFC 2798, "Definition of the inetOrgPerson LDAP Object Class"). |
| ms-userproxy.ldf | **User-Proxy** | Used to create proxy objects in ADAM for use in bind redirection. For more information about proxy objects and bind redirection, see "Bind Redirection for ADAM Proxy Objects" later in this section. |
| ms-azman.ldf | Not applicable | Used to prepare ADAM for use with Windows Authorization Manager. |

These .ldf files are installed to the *windir*\ADAM directory, and they can be imported anytime after ADAM installation.

### Note

- Each of these .ldf files contains within it the command-line instruction for importing the file into the ADAM schema. To import one of these files after ADAM installation:

1. Open the file in a text editor, such as Notepad.

2. Copy the command-line instruction for importing the file from the beginning of the file.

3. Open an ADAM tools command prompt. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.

4. Paste the copied command into the command prompt.

5. Run the command.

**ADAM Security**

Because ADAM is based on the same code as Active Directory, security in ADAM works very similarly to security in Active Directory. ADAM uses the same basic set of security concepts as Active Directory: authentication, authorization, users, groups, ACLs, tokens, and so on. This section describes how these security concepts are implemented in ADAM.

**Default Groups in ADAM**

ADAM uses groups to provide access to directory data. Like Active Directory, ADAM contains its own set of default groups that are created at installation time. After ADAM is installed, a set of default groups exist in each directory partition. In addition, you can create your own custom groups.

ADAM provides three default groups in each configuration and application directory partition. The groups reside in the CN=Roles container of each partition. These default groups include the following:

- Administrators (CN=Administrators,CN=Roles)

- Readers (CN=Readers,CN=Roles)

- Users (CN=Users,CN=Roles)

**Note**

- The schema partition does not contain any groups.

The following figure illustrates the default ADAM groups for each directory partition.

**Default ADAM Groups by Partition**




The following group resides only in the configuration directory partition: Instances (CN=Instances,CN=Roles).

**Note**

- ADAM uses the Instances group for replication. You should not add users to the Instances group.

You can use ADAM ADSI Edit to view the ADAM groups in the CN=Roles container. The following figure shows the groups that are contained in the CN=Roles container in the configuration partition of an ADAM instance.

**Groups in the ADAM Configuration Partition**




The following table describes the purpose and default members of each of the default ADAM groups.

| Partition | Group | Purpose | Default Members |
|---|---|---|---|
| Configuration | Administrators | This group administers the ADAM instance. | The ADAM administrator that is assigned during ADAM setup. The ADAM administrator is always a Windows security principal. |
| Configuration | Readers | This group has Read access to the configuration partition (including the schema). | None |
| Configuration | Users | This group is a computed group. | All users that are defined in any application directory partition |
| Application (0 to *n*) | Administrators | This group administers the application partition in which the group resides. | Members of the Administrator's group from the configuration partition |
| Application (0 to *n*) | Readers | This group has Read access to the application partition in which the group resides. | None |
| Application (0 to *n*) | Users | This is a computed group. | All users from the respective application directory partition |

You can create additional groups in any application partition. You cannot create additional groups in the configuration partition.

**Partition Membership and ACL Restrictions**

This section describes the partition membership and ACL restrictions in ADAM.

**ADAM users**

Cross-partition memberships are not allowed for ADAM users. In other words, an ADAM user cannot reside in groups in more than one directory partition. In addition, ADAM users can be assigned permissions only on objects in their own directory partition; they cannot be assigned permissions on objects outside their own directory partition.

By default, ADAM users cannot be created in the configuration directory partition, and they cannot be members of a group in the configuration directory partition. However, there is a hotfix for ADAM that you can use to create ADAM users in the configuration directory partition, which allows ADAM users to be administrators of an ADAM instance. For more information about this hotfix, contact Microsoft Product Support Services.

**Windows users**

Cross-partition memberships are allowed in ADAM for Windows users (both local and domain Windows users). In other words, a Windows user can reside in groups in more than one directory partition, including the configuration directory partition. In addition, Windows users can be assigned permissions on objects outside their own directory partition.

**Group Memberships in ADAM**

ADAM groups can contain ADAM users (also called ADAM security principals), as well as Windows users (also called Windows security principals), as members. An ADAM user is a user that is created in an ADAM directory partition. A Windows user is a user that exists as a Windows user on the local computer or as a domain member of an Active Directory domain.

An ADAM user or Windows user becomes a member of an ADAM group by being added to the multivalue **member** attribute on the respective group object. You can use ADAM ADSI Edit to add users to ADAM groups.

### Note

- Although ADAM provides a number of default groups, no default users exist in ADAM, except for the Windows security principal that you specify as the ADAM administrator during ADAM setup. The user that you specify as the ADAM administrator becomes a member of the Administrators group in the configuration directory partition.

#### Restrictions on the Administrators Group

By default, only Windows security principals can be made members of the Administrators group on the configuration directory partition. Therefore, by default only Windows security principals can administer an ADAM instance. However, there is a hotfix for ADAM that you can use to create ADAM users in the configuration directory partition that allows ADAM users to be administrators of an ADAM instance. For more information about this hotfix, contact Microsoft Product Support Services.

#### Group Scope

The scope of an ADAM group (or of an ADAM user) is restricted to the partition in which the group (or user) is created. In other words, a group that is created in partition A cannot be used to assign permissions on resources in partition B. The exception to this rule is the Administrators group in the configuration directory partition, whose members have full control of all objects in all partitions. In addition, a user in partition B cannot be added to a group in partition A. This includes groups in the configuration directory partition.

#### Security Principals in ADAM

The term security principal refers to any object (typically, a user) that has a unique SID and that can be assigned permissions to directory objects. In ADAM, security principals can be any of the following:

- ADAM security principals that are created in an ADAM directory partition
- Windows users that are defined on a local computer
- Windows users that are defined in an Active Directory domain

#### ADAM Security Principals

An ADAM security principal is a user or other bindable object that is created in ADAM. A newly installed ADAM instance does not include any default ADAM users. In fact, the default ADAM schema does not contain any user object classes. However, you can easily import the schema extensions that are provided with ADAM to add several different user object classes to the ADAM schema. These object classes include the **User** object class (which is very similar to the **User** object class in Active Directory) and the **inetOrgPerson** object class (which is derived from the **User** object class). You can import these object classes through the Active Directory Application Mode Setup Wizard.

In addition, you can make any object class in the ADAM schema a security principal by adding the **msDS-bindableobject** auxiliary class and the **unicodePwd** attribute to the schema definition of the object class in the ADAM schema.

### Note

- When you create an ADAM security principal, ADAM generates a SID for the security principal automatically.

#### Windows Security Principals

In addition to supporting ADAM security principals, ADAM supports the use of Windows security principals. In other words, Windows security principals can be assigned permissions on objects in ADAM, and Windows security principals can become members of ADAM groups. ADAM does not authenticate Windows security principals. Instead, Windows security principals are authenticated by Active Directory (for domain-based Windows security principals) or by the Local Security Authority (LSA) on the local computer (for local Windows security principals).

#### Default ACLs

Just like Active Directory, ADAM uses ACLs to control access to objects in ADAM. However, unlike Active Directory, the majority of ACLs in ADAM are on the partition head objects (the top-level containers of partitions). These ACLs are inherited throughout the entire partition on which they reside.

By default, the ACLs in the following table exist on the partition head objects in ADAM.

**ADAM Default Partition Head ACLs**

| Partition Head | Group: Default ACLs |
|---|---|
| Configuration | Administrators: Full Control<br><br>Readers: Read |
| Schema | Administrators (from the configuration partition): Full Control<br><br>Readers: Read |
| Application (0 to *n*) | Administrators: Full Control<br><br>Readers: Read<br><br>Instances: five replication-related access control entries (ACE) |

Other default ACLs for ADAM include the ACLs in the following table.

**Other Default ADAM ACLs**

| Security Principal | ACLs |
|---|---|
| Everyone | • "Query self quota" control on quota objects |
| Administrators group | • Full Control on the Roles container<br><br>• Full Control on the Quotas container<br><br>• Full Control on the Lost and Found container |
| Readers group | • Read control on the Roles container (inheritable)<br><br>• Read control on the Lost and Found container |
| **User** and **inetOrgPerson** classes | • "Change password" control for "self" in the default security descriptor in the schema |

**Adding SIDs to ACLs**

By adding the SIDs of particular users to the ACL of a particular object, you can provide access to that object to the security principals that the SIDs represent. To add a SID to an ACL in ADAM, you can use the Dsacls.exe command-line tool.

The general syntax to add a SID to an ACL using **dsacls** for granting permissions is as follows:

```
dsacls "\\
servername:portnumber
\Object DN" /G "User DN":
security descriptor
```

For example, to grant Delete permissions to a user named Mary Baker (CN=Mary Baker,OU=ADAM users,O=Microsoft) on the object ADAM testers (CN=ADAM testers,OU=ADAM users,O=Microsoft,C=US), you use the following syntax:

```
dsacls "\\
servername
:
portnumber
\CN=ADAM testers,OU=ADAM users,O=Microsoft,C=US" /G "CN=Mary Baker,OU=ADAM
users,O=Microsoft,C=US":SD;;
```

The general syntax to add a SID to an ACL using **dsacls** for denying permissions is as follows:

```
dsacls "\\servername:portnumber\Object DN" /D "User DN":security descriptor
```

For example, to deny the Delete, Delete Child, and Delete Tree permissions to Mary Baker (CN=Mary Baker,OU=ADAM users,O=Microsoft) on the ADAM users OU (OU=ADAM users,O=Microsoft,C=US), you use the following syntax:

```
dsacls "\\
servername
:
portnumber
\OU=ADAM users,O=microsoft,C=US" /D "CN=Mary Baker,OU=ADAM
users,O=Microsoft,C=US":SDDCDT;;
```

For additional information about **dsacls** syntax, type **dsacls /?** at an ADAM command prompt. For more information about using **dsacls**, see article 281146, "How to Use Dsacls.exe in Windows 2000," in the Microsoft Knowledge Base.

**Matching Tokens Against an ACL**

In ADAM, when a user is authenticated, ADAM (or the LSA, depending on the type of security principal) creates a security access token for that user. An access token contains the users name, the groups to which that user belongs, a SID for the user, and all of the SIDs for the groups to which the user belongs. The information in the access token is used to determine a user's level of access to objects whenever the user attempts to access them. The SIDs in the access token are compared with the list of SIDs that make up the discretionary access control list (DACL) for the object to ensure that the user has sufficient permission to access the object.

**Authentication in ADAM**

Authentication is the act of ensuring the identity of a user before granting that user access to resources. ADAM implements authentication to ensure the identity of users.

ADAM supports several types of authentication, or binding, depending on the type of user being authenticated:

- Anonymous binding, for anonymous users (users that do not supply a password)
- Simple LDAP binding, for ADAM security principals (security principals that reside in ADAM)
- SASL binding, for Windows security principals (Windows security principals that reside on a local computer or in Active Directory)
- Bind redirection, for ADAM proxy users (security principals that reside in ADAM but that contain a reference to a security principal that resides in Active Directory)

### Note

- Proxy users are designed for special circumstances and should only be used as a last resort, when Windows principals cannot be used directly. For more information, see "Bind Redirection for ADAM Proxy Objects" later in this section.

When a user attempts to authenticate against, or bind to, ADAM, the user might be authenticated by ADAM, the LSA, or Active Directory, depending on the type of user attempting the bind. The following table lists the types of authentication that ADAM supports, the types of users that are appropriate to each authentication method, and the authenticating authority that handles the authentication.

### Authentication Methods in ADAM

| Type of Authentication | Type of User | Authenticating Authority |
|---|---|---|
| Anonymous | Anonymous | Not applicable. The user is not authenticated. |

| | | |
|---|---|---|
| Simple LDAP bind | ADAM security principal | ADAM |
| SASL bind (using Kerberos, NTLM, or negotiated) | Local Windows security principal | LSA on the local computer |
| SASL bind (using Kerberos, NTLM, or negotiated) | Domain Windows security principal | Active Directory |
| Bind redirection (Simple LDAP bind to ADAM, then SASL bind to Active Directory) | ADAM proxy object | LSA and Active Directory |

The types of authentication that are used with ADAM are described in the following sections.

**Anonymous Authentication**

Anonymous authentication occurs when a user attempts to bind to ADAM without providing a password. Instead, the user presents ADAM with an empty (null) password in a simple LDAP bind request. In return, ADAM provides the user with anonymous access to the directory. Anonymous authentication is useful for any applications that require unrestricted, public access to data and for situations in which you do not want to maintain individual accounts for users.

**Note**

- By default, ADAM does not allow anonymous binding. However, you can enable anonymous binding in ADAM by modifying the **dsHeuristics** attribute. For more information, see "To allow anonymous LDAP binding to an ADAM instance" in the ADAM Administrators Guide. To view the ADAM Administrators Guide, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Help**.

By default, the only object in ADAM to which an anonymous user has access is the rootDSE object. (The rootDSE object is visible to anonymous users whether or not anonymous binding is enabled.) If you enable anonymous binding, you can grant anonymous access to any object or attribute in the directory by assigning permissions on the object or attribute to NT Authority\Anonymous Logon (whose well-known SID is S-1-5-7). In addition to assigning permissions on the object or attribute, to allow anonymous users to read an object or attribute in ADAM, you must also grant the **ADS_RIGHT_DS_LIST_OBJECT** permission on the parent of the object or attribute. You can use **dsacls** to provide access to an object in the directory.

**Note**

- To grant anonymous access to an entire partition, grant anonymous access permissions on the partition head.

When a user binds to ADAM anonymously, ADAM generates a token that contains the following: S-1-5-7 (NT Authority\Anonymous Logon).

**Note**

- Tokens that ADAM generates for a user are exposed externally through an attribute on the rootDSE object called **tokenGroups**. The **tokenGroups** attribute is new with ADAM, and it does not appear in Active Directory.

**Simple LDAP Bind for ADAM Security Principals**

Simple LDAP binding occurs when a user binds to ADAM as a security principal that exists as an object in the ADAM directory. This authentication method uses a simple LDAP bind; no SASL options are available with basic authentication. This simple LDAP bind can be accomplished using the distinguished name of the security principal or the user principal name (UPN) of the security principal. In contrast to UPNs in Active Directory, the UPN for an ADAM security principal is not required to contain the @ symbol.

In a simple LDAP bind, the user password is sent in plaintext format to ADAM. For this reason ADAM requires an encrypted channel — using SSL between the client that connects to ADAM and the ADAM server — for basic authentication to occur securely. To enable SSL connections to ADAM, you must have certificates installed on the computer running ADAM. In addition, the clients that connect to ADAM should trust the certification authority (CA) that issues the certificate to the ADAM server. You can obtain certificates from a trusted CA, or, for internal use, you can issue certificates from your own internal CA by installing and using certificate services on Windows 2000 Server or Windows Server 2003.

**Notes**

- Before you attempt to use a certificate with ADAM, you must ensure that the service account under which ADAM is

running has Read access to the certificate that you install. The certificate is located in the following directory: C:\Documents and Settings\All Users\Application Data\Microsoft\Crypto\RSA\MachineKeys

- Although it is not recommended for production environments, you can disable the requirement for SSL in basic authentication by using the **ds-behavior** option of the Dsmgmt.exe command-line tool.

Any ADAM object can be used to bind to ADAM, as long as the object class from which the object is derived contains the **msDS-bindableobject** auxiliary class and the **unicodePwd** attribute. Adding **the msDS-bindableobject** auxiliary class to an object class causes ADAM to generate a SID automatically for any object that is created from that class.

**Note**

- The default ADAM schema contains no bindable object classes. However, several optional user class object definitions are provided with ADAM, including **User** and **inetOrgPerson**. You can import these user classes during ADAM installation.

For a user that is authenticated with basic authentication, the ADAM token contains the following SIDs:

- User SID
- Everyone [S-1-1-0]
- NT AUTHORITY\Authenticated Users [S-1-5-11]
- SID of the user's group from the user's partition
- SIDs of any additional groups that are determined through transitive group expansion

**Note**

- To determine the contents of a user's token, you can view the **tokenGroups** attribute on the rootDSE object.

**Password policy settings and account lockouts**

When basic authentication is used and ADAM is running on a computer running Windows Server 2003, ADAM supports and enforces the password policy settings and account lockout settings that are provided by Windows Server 2003, including the following:

- **Minimum age**
- **Maximum age**
- **Complexity**
- **History**
- **Too many failed logon attempts**
- **Disabling and enabling of accounts**

If the server on which ADAM is running belongs to a workgroup, the server's local password policy settings and account lockout settings are implemented. If the server on which ADAM is running belongs to a domain, the password policy settings and account lockout settings from Active Directory are implemented. Therefore, it is possible for two ADAM instances in the same configuration set to enforce different password policy settings and account lockouts.

**Note**

- You can disable the enforcement of password policy settings in ADAM by setting **ADAMDisablePasswordPolicies**, a value in the attribute **msDS-Other-Settings** on CN=Directory

Service,CN=Windows NT,CN=Services,CN=Configuration,CN=GUID, to **1**.

**SASL Bind for Windows Security Principals**

Integrated authentication occurs when a Windows security principal attempts to bind to an ADAM instance. ADAM does not authenticate Windows users itself. Instead, ADAM relies on the Windows Security API to authenticate the user. Integrated authentication uses SASL binds. All SASL mechanisms that are supported in Active Directory are supported in ADAM. These mechanisms include Kerberos, NTLM, and Negotiated (default). Integrated authentication supports the use of domain\user, UPN, or distinguished name as authentication credentials.

When ADAM receives a SASL bind request, ADAM forwards the request to Active Directory. If the bind request is authenticated by Active Directory, Active Directory returns a token to ADAM. ADAM uses that token in its internal context, augments it with ADAM groups that the Windows user is a member of across all the application partitions that are hosted by the ADAM instance, and then performs transitive group expansion across all the application partitions. That is, ADAM determines group memberships for the Windows user that result from the user being a member of a group that itself is a member of a different group.

For a user that is authenticated using integrated authentication, the ADAM token contains the following SIDs:

- SIDs of all groups from the Windows token that is obtained from Active Directory
- SIDs of all ADAM groups of which the Windows security principal is a member
- SIDs of any additional groups that are determined through transitive group expansion

**Bind Redirection for ADAM Proxy Objects**

ADAM bind redirection should be used only in special cases where an application can perform a simple LDAP bind to ADAM but the application still needs to associate the user with a security principal in Active Directory.

ADAM bind redirection occurs when a bind to ADAM is attempted using a special object called a proxy object. A proxy object is an object in ADAM that represents a security principal in Active Directory. Each proxy object in ADAM contains the SID of a user in Active Directory. When a user attempts to bind to a proxy object, ADAM takes the SID that is stored in the proxy object, together with the password that is supplied at bind time, and presents the SID and the password to Active Directory for authentication. A proxy object in ADAM does not store a password, and users cannot change their Active Directory passwords through ADAM proxy objects.

Because the initial bind request is a simple LDAP bind request, the password is presented in plaintext to ADAM. For this reason an SSL connection is required by default between the directory client and ADAM. ADAM uses Windows Security APIs to present the password to Active Directory.

### Note

- Because the password that is presented to ADAM during ADAM bind redirection is an Active Directory password that is presented in plaintext, you should not disable the requirement for SSL when using bind redirection.

**Proxy objects**

Proxy objects (and proxy object classes) do not exist by default in ADAM. However, you can import a proxy object class into the ADAM schema during ADAM installation. A proxy object can be created from any object class that contains the **msDS-bindProxy** auxiliary class. The **msds-BindProxy** class possesses a single “must contain” attribute, **ObjectSid**, which holds the SID of the associated Active Directory security principal. You can set the value of **ObjectSid** only at the time that the object is created. After a proxy object is created, the value of its **ObjectSid** attribute cannot be modified. You can set the **ObjectSid** of a proxy object to the SID of any local Windows user or to any user who is a member of a domain or forest that is trusted by the computer on which ADAM is running.

### Note

- You cannot create a proxy object for a domain user in an ADAM directory partition that already contains a foreign principal object (FPO) for that same domain user.

For a user that is authenticated through bind redirection, the token that is generated for the user is identical to a token that is generated for a Windows security principal that is binding to ADAM through integrated authentication.

**ADAM Processes and Interactions**
This section describes the following ADAM processes and interactions:

- Locating an ADAM instance
- Replication
- LDAP referrals
- Service publication
- Auditing
- ADAM and Microsoft Authorization Manager
- Backup and restore

**Locating an ADAM Instance**
For a directory-enabled application to be able to communicate with an ADAM instance, the application must specify the NetBIOS name, DNS name, or IP address of the computer on which ADAM is running. In addition, the application must specify the LDAP or SSL communication port that is used by the ADAM instance. For example, a directory-enabled application referencing an ADAM instance running on a computer with a DNS name of ADAMSrv1.microsoft.com — with a NetBIOS name of ADAMSrv and an IP address of 192.168.0.1 and using communication port 389 for LDAP connections — can address the ADAM instance by using one of the following:

- ADAMSrv1.microsoft.com:389
- ADAMSrv:389
- 192.168.0.1:389

**Locating ADAM Instances in Active Directory Environments**
When you run ADAM in an Active Directory environment, each ADAM instance creates and maintains service connection point (SCP) objects about itself in Active Directory. For more information, see "Service Publication" later in this document.

**Server Object Names**
By default, the relative distinguished name of a server object in ADAM is *servername*$*instancename*, where *servername* represents the computer on which the ADAM instance is running and *instancename* represents the service name of the ADAM instance.

**Modifying Settings on a Computer Running ADAM**
You initially specify the computer name, DNS name, NetBIOS name, and IP address when you set up the computer. You initially specify the communication port numbers that are used by a given ADAM instance when you install the ADAM instance.

You can change these settings on a computer running ADAM anytime. However, when you update any of these settings, you may also need to reconfigure directory-enabled applications to properly address the ADAM instance. For example, if the host name of the computer on which ADAM is running changes, your directory-enabled application may need to be updated accordingly.

**Replication**
ADAM uses replication to provide fault tolerance and load balancing for directory services. ADAM uses a type of replication called multimaster replication. Through replication, ADAM copies directory data updates that are made to a directory partition on one ADAM instance to other ADAM instances that hold copies of the same directory partition. ADAM instances that hold copies of the same directory partition or partitions form a logical grouping called a configuration set. ADAM prevents replication conflicts through the use of change tracking information.

### Note

- An ADAM configuration set maintains its own replication topology, separate from any Active Directory replication topology that might also exist. Directory partitions cannot be replicated between ADAM instances and Active Directory domain controllers.

**Multimaster Replication**
With multimaster replication, you can make changes to directory data on any ADAM instance. ADAM replicates these changes to other members of the configuration set automatically. Multimaster replication is characterized by loose data consistency with convergence. That is, when you make changes to data on a given directory partition at one ADAM instance, replicas of that directory partition that are stored on other ADAM instances become inconsistent with the most up-to-date replica of the directory partition (the partition where the changes were made). However, as changes get replicated through the configuration set, all partition replicas once again become identical; that is, they converge to the most recent data.

**Preventing Replication Conflicts**
What if two different users make changes to the same data on replicas of the same directory partition on two different ADAM instances? In this case, each ADAM instance attempts to replicate the changes, which creates a conflict. To resolve this conflict, replication partners that receive these conflicting changes examine the attribute data that is contained in the changes, each of which holds a version and a time stamp. ADAM instances accept the change with the higher version and discard the other change. If the versions are identical, ADAM instances accept the change with the more recent time stamp.

If two or more values in a multivalue attribute on an object are updated simultaneously on two different ADAM instances, only one of the updated values will be replicated. In other words, simultaneous updates to a multivalue attribute that occur on two different ADAM instances are considered to be in conflict, even if the updates apply to different values in the multivalue attribute. The only exception to this rule is for linked-value attributes (such as group memberships), which do allow for simultaneous updates to different values in the linked-value attribute.

**Sites and Replication**
In ADAM, the replication topology is the set of physical connections that is used to replicate directory updates among ADAM instances in a configuration set. ADAM builds the most efficient replication topology for a configuration set based on the site or sites in which members of the configuration set reside. ADAM treats replication within a site differently than replication between sites.

**Building the intersite replication topology**
One ADAM server in each site, called the intersite topology generator (ISTG), is assigned to build the topology. Together, these ISTGs build a least-cost, spanning-tree algorithm that eliminates redundant replication paths between sites. The intersite replication topology is updated regularly to respond to any changes that occur in the network. You can control intersite replication through the information that you provide when you create your site links.

**Determining when intersite replication occurs**
ADAM preserves bandwidth between sites by minimizing the frequency of replication and by enabling you to schedule the availability of site links for replication. By default, intersite replication across each site link occurs every 180 minutes (3 hours). You can adjust this frequency to match your specific needs. (Remember that increasing this frequency increases the amount of bandwidth that is used by replication.) In addition, you can schedule the availability of site links for use by replication. By default, a site link is available to carry replication traffic 24 hours a day, 7 days a week. You can limit this schedule to specific days of the week and times of the day. You can, for example, schedule intersite replication so that it only occurs after normal business hours.

**Building the intrasite replication topology**
The Knowledge Consistency Checker (KCC) on each ADAM instance automatically builds the most efficient replication topology for intrasite replication using a bidirectional ring design. This bidirectional ring topology attempts to create at least two connections to each ADAM instance (for fault tolerance) and no more than three hops between any two ADAM instances (to reduce replication latency). To prevent connections of more than three hops, the topology can include shortcut connections across the ring. The KCC updates the replication topology regularly.

The KCC in ADAM operates identically to the KCC in Active Directory. For more information about how the KCC in Active Directory works, see "How Active Directory Replication Topology Works."

**Determining when intrasite replication occurs**
Intrasite replication is optimized for speed, rather than bandwidth, because bandwidth within a site is assumed to be high speed. Intrasite replication occurs automatically on the basis of change notification, and it begins when a directory update occurs. By default, the source ADAM instance waits 15 seconds and then sends an update notification to its closest replication partner. If the source ADAM instance has more than one replication partner, subsequent notifications go out by default at three-second intervals to each partner. After receiving notification of a change, a partner ADAM instance sends a directory update request to the source ADAM instance. The source ADAM instance responds to the request with a replication operation. The three-second notification interval prevents the source ADAM instance from being overwhelmed with simultaneous update requests from its replication partners.

If no directory updates occur in a given time period, intrasite replication still occurs, based on a scheduled interval. By default,

this scheduled interval is once per hour. For information about modifying this time period, see the ADAM Administrator's Guide. After you install ADAM, to view the ADAM Administrator's Guide, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Help**.

#### Ensuring Replication Security

To ensure replication security, ADAM authenticates replication partners before replication, and replication authentication always occurs over a secure channel. ADAM replication partners authenticate each other using the service account that is specified for each respective ADAM instance. ADAM uses the Security Support Provider Interface (SSPI) to establish the appropriate authentication security level between replication partners. The method that is used for replication authentication within a configuration set depends on the value of the **msDS-ReplAuthenticationMode** attribute on the configuration directory partition. After replication partners authenticate, all replication traffic between the two partners is encrypted.

The following table describes the security levels for ADAM replication and the corresponding authentication mode (**msDS-ReplAuthenticationMode)** attribute value for each security level. The default replication security level for a new, unique ADAM instance is 1, unless a local workstation user account is specified as the ADAM service account. If a local workstation account is specified as the ADAM service account, the replication security level is 0.

**ADAM Replication Security Levels**

| Replication Security Level | Authentication Mode | Description | Supported Environments |
|---|---|---|---|
| Mutual authentication with Kerberos | 2 | Kerberos authentication, using SPNs, is required. If Kerberos authentication fails, the ADAM instances will not replicate. | The configuration set must be fully contained in an Active Directory domain, forest, or forest trust. |
| Negotiated | 1 | Kerberos authentication (using SPNs) is attempted first. If Kerberos fails, NTLM authentication is attempted. If NTLM fails, the ADAM instances will not replicate. | The configuration set can contain Windows NT 4.0 member servers. This is useful for transitioning between negotiated pass-through mode and mutual authentication with Kerberos. |
| Negotiated pass-through | 0 | All ADAM instances in the configuration set use an account name and password that are identical to those of the ADAM service account. | The configuration set can include computers that are joined to one or more workgroups or that are joined to multiple domains or forests without trust relationships. |

To help maintain ADAM replication security, the following best practices are recommended:

- Use the highest level of replication security that your environment can support.
- In Active Directory environments, run ADAM on member servers — rather than on domain controllers — whenever possible.
- If you run ADAM on a domain controller in an Active Directory environment, do not use the Network Service account as the ADAM service account. Instead, use a domain user account that does not have administrative privileges.
- In workgroup and Windows NT 4.0 environments, do not use an account with administrative privileges as an ADAM service account.
- Use separate configuration sets for applications with strict isolation requirements.

#### Replication Monitoring

You can use the Repadmin.exe command-line tool to monitor ADAM replication.

**Repadmin /showutdvec** shows, from the perspective of the ADAM instance on which the command is run, the most recent change that the ADAM instance has received from each of its replication partners, both directly and transitively.

You can use **Repadmin /replicate** to force replication between ADAM instance replication partners.

For more information about the **repadmin** command, type **repadmin /?** at an ADAM tools command prompt.

**LDAP Referrals**

In accordance with RFC 2251, ADAM supports the use of LDAP referrals. With LDAP referrals, you can use ADAM to direct LDAP clients to a more appropriate location in response to a query from the client. Also in accordance with RFC 2251, ADAM can return referrals in response to any LDAP operation request, except "unbind" and "abandon" requests, which have no responses. ADAM maintains referral data in crossRef objects that are stored in the partitions container of the configuration directory partition.

ADAM creates and maintains crossRef objects for all directory partitions in the configuration set, including the configuration and schema directory partitions. These crossRef objects are used to generate referrals in response to queries that request data about objects that exist in the configuration set but that are not contained in the ADAM instance handling the request.

At a minimum, at least one URL must be present in the referral. When possible, ADAM returns at least three referrals to provide referral redundancy to clients in distributed environments. A referral is not returned for a **singleLevel** or **wholeSubtree** search in which the search scope spans multiple directory partitions and several different servers must be contacted to complete the operation. Instead, continuation references are returned.

If a crossRef object with a matching directory partition exists, ADAM uses it for referrals. However, if a crossRef object does not exist, but the queried namespace uses DNS-style (DC=) naming, ADAM constructs a DNS name to be used as a referral. If the queried namespace does not use DNS-style naming, a superior reference is used, if one exists. A superior reference can be configured by setting the **superiorDnsRoot** attribute on the configuration partition crossRef object. If a superior reference is not configured (it is not configured by default), an operational error is returned.

**Generating Referrals from a crossRef Object**

If the **dNSRoot** attribute has been set on the crossRef object (**dNSRoot** is not set by default), ADAM returns the DNS name as the referral for the queried namespace. (By default, the **dNSRoot** attribute is set to NULL.) Then, the client can chase the referral through registered DNS records for each host that contains the queried namespace.

If the **dNSRoot** attribute is NULL, the referral is generated based on other information on the crossRef object, as follows:

- ADAM refers to the **msDS-NC-Replica-Locations** attribute to find the NTDS Settingsobjects that represent the instances hosting this partition (for example, CN=NTDS Settings,CN=Machine1$instance1,CN=Servers,CN=Default-First-Site-Name,CN=Sites,CN=Configuration,CN={*GUID*}). If more than three ADAM instances host the directory partition, a randomly selected set of three instances is chosen.

- From the NTDS Settings object (the authoritative source in the directory for communications port information), ADAM retrieves the **msDS-PortLDAP** attribute for the communications port.

- From the parent object (Server Object), ADAM retrieves the **dNSHostName** attribute (instance1.machine1.domain.com). **dNSHostName** is the authoritative source for the server host name, and it can be either a DNS host name or a host IP address. If the server name changes, this object can be relied on to have the updated information.

- ADAM then constructs the referral based on the server and communications port information. ADAM returns multiple referrals for each request to ensure that the client can pursue the request, even if one or more hosts represented in the referrals cannot be reached. ADAM selects a random number of locations (three by default) and continues to use those defaults for five minutes, after which another three hosts are selected. Both the number of hosts and the number of minutes are configurable. If the number of specified hosts is greater than the actual number of instances holding the directory partition, all of the hosts are returned as referrals, and they will continue to be used indefinitely. This information is kept in the crossref cache, and it is refreshed anytime the crossref cache is rebuilt. The crossref cache is automatically rebuilt when **msds-nc-replica-locations** changes, that is, when an ADAM instance is added or removed.

- You can modify the values in the following table on the **msDS-Other-Settings** attribute of the directory service object — or on a crossRef object — to configure the referral behavior of an ADAM instance.

  **ADAM Referral Behavior Values**

| Value Name | Default | Meaning |
|---|---|---|
| **MaxReferrals** | 3 | Maximum number of referrals returned by ADAM |
| **ReferralRefreshInterval** | 5 minutes | Time interval at which referral information is updated |

| | | |
|---|---|---|
| **SelfReferralOnly** | 0 | Setting the value to 1 prevents ADAM from providing referral information. |

- Values that you set on a crossRef object take precedence over any values that you set on the directory service object.

**Generating Referrals Without crossRef Objects**
If a crossRef object does not exist for the search base, ADAM determines the DNS name if the distinguished name that is presented in the request uses DNS-style (DC=) syntax. Otherwise, ADAM uses the value of the **supref** attribute of the configuration directory crossRef object to generate a referral if **supref** contains a value. (This attribute is NULL by default, and it must be configured explicitly.) Otherwise, no referrals are generated.

**Service Publication**
In Active Directory environments, ADAM uses SCP objects to publish ADAM service information in the Active Directory directory store. Administering ADAM SCP objects includes the following operations:

- Manually creating SCPs when necessary
- Removing SCP objects when necessary
- Modifying permissions on SCP objects when necessary
- Managing SCPs efficiently
- Checking the ADAM event log for SCP errors

**Creating SCPs**
By default in Active Directory environments, the ADAM service creates SCPs in Active Directory when the ADAM service starts. To create SCPs successfully, the ADAM service account must have sufficient rights in Active Directory. If SCP creation fails, ADAM writes an event regarding the failure to the ADAM event log.

The default location for the ADAM SCP object is under the computer object that represents the computer on which ADAM is running. This default location can be altered by specifying a different location on the **SCPPublishingService** object.

When it is used as the ADAM service account on a computer that is joined to a domain, the Network Service account usually has sufficient rights to create SCPs in Active Directory. However, when it is used on an ADAM instance running on a domain controller, the Network Service account does not have sufficient rights to create SCPs in Active Directory.

**Removing SCPs**
An ADAM instance deletes its SCP from Active Directory when you remove the ADAM instance from the computer. Removing the SCP requires sufficient administrative credentials. If SCP removal fails, client applications may be directed to a nonexistent ADAM instance.

**Modifying SCPs**
An ADAM instance checks and updates, if necessary, its SCP each time that the ADAM instance starts. At startup, the ADAM instance searches the global catalog for its own GUID and retrieves the distinguished name of the SCP object. The ADAM instance then binds to that distinguished name and updates the SCP object as necessary. In addition, the ADAM instance reviews the SCP object on an hourly basis by default to confirm its validity, particularly regarding any directory partitions that have been added to or removed from the ADAM instance since the SCP object was last updated.

You can modify the default time interval at which ADAM reviews the SCP object by adding a **Server information update interval (mins)** value to the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\instancename\Parameters** and setting this value to the time interval (in minutes) that you want to use.

### Notes

- To prevent an ADAM instance from updating its SCP object, or to prevent an ADAM instance from recreating the SCP object after you have deleted the SCP object, set the registry key value **HKLM\System\CurrentControlSet\Services*instance name*\Parameters\Server information update interval (mins)** to **0**.
- To prevent the publication of SCPs for all ADAM instances in an ADAM configuration set, set the attribute **Enabled** on the object **CN=SCP Publication Service,CN=Directory**

**Service,CN=Windows NT,CN=Services,CN=Configuration,CN={*GUID*}** to **False**.

For more information about publishing and using SCPs, see the Microsoft Platform SDK on MSDN and search for "Service Publication."

**Managing ADAM SCPs**
One efficient way to manage SCPs for ADAM is to create an SCP container for your ADAM configuration set in Active Directory. In this container, place the computer objects of the computers on which ADAM instances are running. In addition, create a group called, for example, "ADAM instances," and then assign permissions on the SCP container to the group. Also, delegate control of the group to the assigned ADAM administrator. Then, each time the ADAM administrator installs a new ADAM instance under a new service account, the ADAM administrator can simply add the new service account to the ADAM instances group, and the ADAM instance creates and maintains its SCPs transparently.

**Checking the ADAM Event Log**
SCP-related errors do not prevent an ADAM instance from functioning properly. ADAM does, however, report SCP-related errors in the ADAM event log so that errors can be resolved. ADAM reports SCP-related errors as follows:

- When an SCP creation or update fails, the ADAM instance reports the error and points to an .ldf file that can be used to resolve the problem manually. Using this .ldf file to create or update the SCPs requires administrative credentials in Active Directory.

- When ADAM fails to remove an SCP object, ADAM reports this failure both in the audit log of the appropriate domain controller and in the ADAM event log.

**Auditing**
As with Active Directory, you can audit objects in ADAM for such activities as creation, modification, deletion, and reads. You can also audit ADAM objects at the attribute level. That is, you can specify specific attributes that you want to audit. However, in contrast to Active Directory, no user interface (UI) exists in ADAM for specifying which objects or attributes to audit. To enable auditing in ADAM, you must modify system access control lists (SACLs) for objects and attributes programmatically or by using Ldp.exe.

### Notes

- ADAM does not support auditing when it runs on Windows XP Professional.

- Before implementing auditing in ADAM, you should apply the hotfixes that are available for ADAM. For more information about ADAM hotfixes, contact Microsoft Product Support Services.

**ADAM and Microsoft Authorization Manager**
Microsoft Authorization Manager can use ADAM as a policy store for authorization policy. The ADAM installation process allows the installation of the Authorization Manager schema. The Authorization Manager schema must be installed and the container in which the store is to be created must exist before the Authorization Manager store is created. You can manage the policy store in ADAM either through the Authorization Manager scriptable COM interfaces or by using the Authorization Manager UI.

### Note

- Currently, you cannot use ADAM to store security principals for Authorization Manager.

**Backup and Restore**
By default, each instance of ADAM running on an ADAM server stores its database file, Adamntds.dit, and the associated log files in Program Files\Microsoft ADAM*Instancename*\Data. These files should be included as part of the regular backup plan of your organization. You can back up the directory stores using Ntbackup or any Windows Logo Program, non-Microsoft backup utility.

**Restoring ADAM**
When you restore a database to an existing ADAM instance, you must stop the ADAM instance before you run the restore operation. In addition, it is recommended that you move (or delete) the existing database and log files from the ADAM instance before beginning the restore operation.

**Authoritative Restore**

If objects in the directory are inadvertently deleted or modified, and if those objects are replicated in a configuration set, you must authoritatively restore those objects so that the correct version of the objects is replicated. To authoritatively restore directory data, run the **dsdbutil** utility after you have restored the data but before you restart the ADAM instance. With **dsdbutil**, you can mark directory objects for authoritative restore. When an object is marked for authoritative restore, its metadata version number is changed so that the number is higher than any other metadata version number in the configuration set. This ensures that any data you restore is properly replicated throughout the configuration set.

**Backup Latency Interval**

On computers running Windows Server 2003 with Service Pack 1 (SP1) and ADAM, a new NTDS Replication event provides a warning to administrators when an ADAM partition has not been backed up. Event ID 2089 provides the backup status of each directory partition that a domain controller stores, including application directory partitions. Specifically, event ID 2089 is logged in the Directory Service event log when partitions in the Active Directory forest are not backed up within a backup latency interval. The value for the backup latency interval is stored as a **REG_DWORD** value in the **Backup Latency Threshold (days)** entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

By default, the value of **Backup Latency Threshold (days)** is half the value of the tombstone lifetime of the forest. If halfway through the tombstone lifetime a directory partition has not been backed up, event ID 2089 is logged in the Directory Service event log and continues daily until the directory partition is backed up.

This event serves as a warning to administrators and monitoring applications to make sure that domain controllers are backed up before the tombstone lifetime expires. However, it is recommended that backups be taken at a much higher frequency than the default value of **Backup Latency Threshold (days)**.

**Network Ports Used by ADAM**

The following table lists the network ports that are used by ADAM.

### Default Port Assignments for ADAM

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | None | 389 |
| LDAP SSL | None | 636 |
| RPC Endpoint Mapper | 135 | 135 |

**Related Information**

The following resources contain additional information that is relevant to this section:

- Article 817583, “Active Directory Services Does Not Request Secure Authorization Over an SSL Connection,” in the Microsoft Knowledge Base
- Active Directory Application Mode Tools and Settings
- Microsoft Platform SDK on MSDN for more information about DsReplicaGetInfo (in “DsReplicaGetInfo”)
- Active Directory Schema Technical Reference
- Article 281146, “How to Use Dsacls.exe in Windows 2000,” in the Microsoft Knowledge Base
- How Active Directory Replication Topology Works
- Microsoft Platform SDK on MSDN for more information about publishing and using SCPs (search for “Service Publication”)
- “Windows Server 2003 Active Directory Application Mode” on the Microsoft Web Site

# Active Directory Application Mode Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2012, Windows Server 2012 R2

### In this section


- ADAM Tools
- ADAM Registry Entries
- ADAM WMI Classes
- Network Ports Used by ADAM


This section contains information about the tools, registry entries, Windows Management Instrumentation (WMI) classes, and network ports that are associated with Active Directory Application Mode (ADAM).

**ADAM Tools**
The following tools are associated with ADAM.

### Note

Some ADAM tools are updated versions of tools that are provided with Active Directory. When you administer ADAM with the tools that are described in this section, you should always use the version of the tool that ships with ADAM. By default, these tools can be found in the *windir*\ADAM directory on the computer on which ADAM is installed.

For more information about the tools in this section, see the ADAM Administrator's Guide. To view the ADAM Administrator's Guide, after you install ADAM click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Help**.

**ADAM-adsiedit.msc: ADAM ADSI Edit**

**Category**
This tool ships with ADAM.

### Version compatibility

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers and servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition |
| Computers running:<br>• Windows XP Professional | Computers running ADAM:<br>Windows XP Professional |

ADAM ADSI Edit is a Microsoft Management Console (MMC) snap-in tool that you can use to view and modify directory objects.

**ADAMsetup: ADAM Setup**

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers and servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running:<br>• Windows XP Professional | Runs locally. |

ADAM Setup performs an unattended installation of ADAM using a preconfigured setup script file.

**ADAMuninstall: ADAM Uninstall**

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| • Computers running ADAM | Runs locally. |

ADAM Uninstall uninstalls an ADAM instance.

**Csvde.exe: Csvde**

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers and servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running ADAM:<br>Windows XP Professional |

You can use Csvde to import and export data from Active Directory using files that store data in the comma-separated value (CSV) file format standard. Csvde also supports batch operations that are based on the CSV file format.

To find more information about Csvde, see "Command-Line References" in Tools and Settings Collection.

**Dsacls: Dsacls**

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers and servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running ADAM:<br>Windows XP Professional |

Dsacls displays and changes permissions (access control entries (ACEs)) in the access control list (ACL) of ADAM objects.

The ACEs that you add by using Dsacls must be object-specific permissions that override the default partition permissions that are defined in the ADAM schema. Do not add ACEs unless you are well informed about security for ADAM objects.

If you specify an object without additional parameters, Dsacls displays the ACEs in the ACL.

#### Dsdbutil: Dsdbutil

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| • Computers running ADAM | Runs locally. |

Dsdbutil provides management facilities for ADAM. This tool is intended for use by experienced administrators.

**Note**

- Some Dsdbutil commands can be used only against ADAM instances that are not currently running.

#### Dsdiag: Dsdiag

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers and servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running ADAM:<br>Windows XP Professional |

Dsdiag analyzes the state of the ADAM directory service and reports any problems to assist in troubleshooting. Dsdiag provides detailed information about how to identify abnormal behavior in the system.

Dsdiag consists of a framework for running tests and a series of tests to verify different functional areas of the system. This framework selects which directory services are tested, according to scope directives from the user.

**Dsmgmt: Dsmgmt**

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers and servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running ADAM:<br>Windows XP Professional |

Dsmgmt provides management facilities for ADAM, including partition management and Lightweight Directory Access Protocol (LDAP) policies.

**Ldifde.exe: Ldifde**

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers and servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running ADAM:<br>Windows XP Professional |

You can use Ldifde to create, modify, and delete directory objects. You can also use Ldifde to extend the schema, export ADAM user and group information to other applications or services, and populate ADAM with data from other directory services.

**Ldp.exe: Ldp**

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers and servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running:<br>• Windows XP Professional | Computers running an LDAP directory service, such as ADAM |

Ldp is an LDAP graphical user interface (GUI) tool that you can use to perform operations such as connect, bind, search, modify, add, and delete against any LDAP-compatible directory, such as ADAM. You can also use Ldp to view objects, along with their metadata, that are stored in ADAM (for example, security descriptors and replication metadata).

You can use the online **dbdump** feature in Ldp to view values that are stored in the database while the domain controller is running. You can trigger **dbdump** by modifying the **dumpDatabase** attribute on the rootDSE.

**Repadmin: Repadmin**

**Category**
This tool ships with ADAM.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers and servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition |
| Computers running:<br>• Windows XP Professional | Computers running ADAM:<br>Windows XP Professional |

Repadmin performs tasks that are related to replication, including managing and modifying replication topology, forcing replication events, and displaying replication metadata and up-to-dateness vectors.

**ADAM Registry Entries**
The following registry entries are associated with ADAM.

The registry entries under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics** control the logging level for the component or process that is specified in the entry name. The value for each entry is set to an integer from and including 0 (no logging) through 5 (most verbose logging).

The registry entries under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters** control or contain information about the configuration of an ADAM instance.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_instancename\Diagnostics**
The following registry entries are located under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics**.

**6 Garbage Collection**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of events that are generated when objects that are marked for deletion are actually deleted.

**7 Internal Configuration**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of internal operations.

**8 Directory Access**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of read and write operations to directory objects from all sources.

**9 Internal Processing**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of events that are related to internal directory service operations.

**11 Initialization/Termination**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of events that are generated by starting and stopping ADAM.

**12 Service Control**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of ADAM service events.

**13 Name Resolution**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of events that are generated by the resolution of addresses and ADAM names.

**14 Backup**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of events that are related to backing up ADAM. Specifically, controls the logging of events that occur when Extensible Storage Engine (ESE) database records are read or written during backup.

**16 LDAP Interface Events**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of events that are related to LDAP.

**22 DS RPC Client**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of events that are related to communication between ADAM instances. Examples of logged events include remote procedure call (RPC) errors, canceled calls, and service principal name (SPN)–related operations.

**23 DS RPC Server**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of events that are related to an ADAM instance acting as an RPC server. An ADAM instance acts as an RPC server, for example, during outbound replication and replication setup operations.

**24 DS Schema**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Diagnostics

**Version**
Computers running ADAM.

Controls the logging of events that are related to schema errors and operations. Such errors and operations include schema additions, deletions, modifications, look-up errors, look-up failures, and caching errors.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_instancename\Parameters**
The following registry entries are located under
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

**BinPath**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the path to the program files for the ADAM instance.

**CommonBinPath**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the path to the ADAM program files that are shared by all ADAM instances running on a given computer.

**Configuration NC**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the distinguished name of the configuration directory partition.

**Database Backup Path**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Determines the directory that is used as the target directory when online backups of the directory database are performed.

**Database Log Files Path**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Determines the directory path that is used to store ADAM log files.

**Database Logging/Recovery**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Controls a Microsoft Jet database engine parameter called *JET_paramRecovery* that determines whether database operations are logged.

**DS Drive Mappings**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Tracks local drive mapping names so that the database file can be located if drive mappings are modified.

**DSA Database File**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Determines the file that is used by the ADAM instance for storing objects.

**DSA Working Directory**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Specifies the working directory of the ADAM instance.

**Long Instance Name**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the instance name of the ADAM instance.

**Machine DN Name**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the distinguished name of the computer on which ADAM is running.

**Port LDAP**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the port number that is currently being used for LDAP.

**Port SSL**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the port number that is currently being used for Secure Sockets Layer (SSL).

**Schema Version**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the ADAM schema version for which a particular operating system is configured.

**Service Account SID**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the security ID (SID) of the account that is being used as the service account for the ADAM instance.

**System Schema Version**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_*instancename*\Parameters

**Version**
Computers running ADAM.

Contains the version of the ADAM schema at the time that a backup is taken. This value is used to prevent an incompatible schema version from being restored from backup.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ADAM_instancename\Diagnostics**
The following registry entry is located under
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

**Backup Latency Threshold (Days)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2003 with Service Pack 1 (SP1)

**Default value**
Half the value of the tombstone lifetime of the forest.

When the value is reached, logs event ID 2089 in the Directory Service event log, warning administrators and monitoring applications to make sure that domain controllers are backed up before the tombstone lifetime expires.

**ADAM WMI Classes**

The following table lists and describes the WMI classes that are associated with ADAM.

**WMI Classes Associated with ADAM**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| **rootDSE** | root\directory\LDAP | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running ADAM:<br>Windows XP Professional |
| **DS_LDAP_Class_Containment** | root\directory\LDAP | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running ADAM:<br>Windows XP Professional |
| **DS_LDAP_Instance_Containment** | root\directory\LDAP | Domain controllers and servers running ADAM:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Computers running ADAM:<br>Windows XP Professional |

For more information about these WMI classes, search for “Mapping Active Directory to WMI” in the WMI SDK documentation on MSDN.

**Network Ports Used by ADAM**

The network ports that are used by the data store are listed in the following table.

**Port Assignments for the Data Store**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | None | 389 |
| LDAP SSL | None | 636 |
| RPC Endpoint Mapper | 135 | 135 |

# Active Directory Structure and Storage Technologies

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

Administrators use Active Directory to store and organize objects on a network (such as users, computers, devices, and so on) into a secure hierarchical containment structure that is known as the logical structure. Although the logical structure of Active Directory is a hierarchical organization of all users, computers, and other physical resources, the forest and domain form the basis of the logical structure. Forests, which are the security boundaries of the logical structure, can be structured to provide data and service autonomy and isolation in an organization in ways that can both reflect site and group identities and remove dependencies on the physical topology.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

Domains can be structured in a forest to provide data and service autonomy (but not isolation) and to optimize replication with a given region. This separation of logical and physical structures improves manageability and reduces administrative costs because the logical structure is not affected by changes in the physical structure. The logical structure also makes it possible to control access to data. This means that you can use the logical structure to compartmentalize data so that you can control access to it by controlling access to the various compartments.

The data that is stored in Active Directory can come from many diverse sources. With so many different data sources and so many different types of data, Active Directory must employ some type of standardized storage mechanism so that it can maintain the integrity of the data that it stores. In Active Directory, objects are used to store information in the directory, and all objects are defined in the schema. The object definitions contain information, such as data type and syntax, that the directory uses to ensure that the stored data is valid. No data can be stored in the directory unless the objects that are used to store the data are first defined in the schema. The default schema contains all the object definitions that Active Directory needs to function; however, you can also add object definitions to the schema.

While the directory is exposed to you through a logical structure that consists of elements such as domains and forests, the directory itself is implemented through a physical structure that consists of a database that is stored on all domain controllers in a forest. The Active Directory data store handles all access to the database. The data store consists of both services and physical files. These services and physical files make the directory available, and they manage the processes of reading and writing the data inside the database that exists on the hard disk of each domain controller.

**Active Directory Structure and Storage Architecture**

The Active Directory structure and storage architecture consists of four parts:

- Active Directory domains and forests. Forests, domains, and organizational units (OUs) make up the core elements of the Active Directory logical structure. A forest defines a single directory and represents a security boundary. Forests contain domains.

- Domain Name System (DNS) support for Active Directory. DNS provides a name resolution service for domain controller location and a hierarchical design that Active Directory can use to provide a naming convention that can reflect organizational structure.

- Schema. The schema provides object definitions that are used to create the objects that are stored in the directory.

- Data store. The data store is the portion of the directory that manages the storage and retrieval of data on each domain controller.

  The following figure illustrates the Active Directory data structure and storage architecture.

**Active Directory Data Structure and Storage Architecture**




**Active Directory Domains and Forests**

Domains partition the directory into smaller sections within a single forest. This partitioning results in more control over how data is replicated so that an efficient replication topology can be established and network bandwidth is not wasted by replicating data where it is not required. OUs make it possible to group resources in a domain for management purposes, such as applying Group Policy or delegating control to administrators.

The following figure illustrates the relationships of OUs, domains, and forests in the logical structure architecture.

**Logical Structure Architecture**




**DNS Support for Active Directory**

Active Directory uses DNS as its domain controller location mechanism. When any of the principal Active Directory operations, such as authentication, updating, or searching, is performed, domain joined computers use DNS to locate Active Directory domain controllers, and these domain controllers use DNS to locate each other. For example, when a network user with an Active Directory user account logs on to an Active Directory domain, the user's computer uses DNS to locate a domain controller for the Active Directory domain to which the user wants to log on.

To log on to a network that consists of an Active Directory forest, a client workstation must first be able to locate a nearby domain controller. The domain controller is necessary for initial authentication of both the workstation and the user and for subsequent authorization of the user for the files and resources to which the user needs access. The support that is provided to Active Directory by DNS enables a client workstation to locate a domain controller.

**Active Directory Schema**

The Active Directory schema contains definitions for all the objects that are used to store information in the directory. There is one schema per forest. However, a copy of the schema exists on every domain controller in the forest. This way, every domain controller has quick access to any object definition that it might need, and every domain controller uses the same definition when it creates a given object. The data store relies on the schema to provide object definitions, and the data store uses those definitions to enforce data integrity. The result is that all objects are created uniformly, and it does not matter which domain controller creates or modifies an object because all domain controllers use the same object definition.

The following figure illustrates the relationship of the schema to the data store in the schema architecture.

**Schema Architecture**




**Active Directory Data Store**

The Active Directory data store is made up of several components that together provide directory services to directory clients. These components include the following:

- Four interfaces:
    - Lightweight Directory Access Protocol (LDAP)
    - Replication (REPL) and domain controller management interface
    - Messaging API (MAPI)
    - Security Accounts Manager (SAM)
- Three service components:
    - Directory System Agent (DSA)
    - The database layer
    - Extensible Storage Engine (ESE)
- The directory database where the data is actually stored

The following figure illustrates the relationships of these components in the data store architecture.

**Data Store Architecture**




**Active Directory Structure and Storage Components**

You can define some components for structure and storage in Active Directory, while others are defined by the system and cannot be modified.

- Forests, domains, and OUs are components that constitute the logical structure of Active Directory. You define them during the installation of Active Directory.
- DNS support for Active Directory includes components that are used to locate domain controllers and that use DNS naming schemes. Each domain in a forest must adhere to DNS naming schemes, and domains are organized in a root and subordinate domain hierarchy.
- The schema is a single component that exists inside the directory. The schema contains definitions of the objects that are

used to store information in the directory. These object definitions include two primary components: classSchema objects and attributeSchema objects.

- The data store consists of three layers of components. The first layer provides the interfaces that clients need to access the directory. The second layer provides the services that perform the operations that are associated with reading data from and writing data to the directory database. The third layer is the database itself, which exists as a single file on the hard disk of each domain controller.

**Active Directory Domains and Forests**

The logical structure of Active Directory is a hierarchical structure of Active Directory domains and OUs in a forest. The relationships of the components in the logical structure control access to stored data, and they control how information is replicated between the various domain controllers in the forest. The main components of the Active Directory logical structure are described in the following table.

**Domain and Forest Components**

| Component | Description |
|---|---|
| Forest | A forest is the highest level of the logical structure hierarchy. An Active Directory forest represents a single self-contained directory. A forest is a security boundary, which means that administrators in a forest have complete control over all access to information that is stored inside the forest and to the domain controllers that are used to implement the forest. |
| Domain | Domains partition the information that is stored inside the directory into smaller portions so that the information can be more easily stored on various domain controllers and so that administrators have a greater degree of control over replication. Data that is stored in the directory is replicated throughout the forest from one domain controller to another. Some data that is relevant to the entire forest is replicated to all domain controllers. Other data that is relevant only to a specific domain is replicated only to domain controllers in that particular domain. A good domain design makes it possible to implement an efficient replication topology. This is important because it enables administrators to manage the flow of data across the network, that is, to control how much data is replicated and where that replication traffic takes place. |
| OU | OUs provide a means for administrators to group resources, such as user accounts or computer accounts, so that the resources can be managed as one unit. This makes it much easier to apply Group Policy to multiple computers or to control the access of many users to a single resource. OUs also make it easier to delegate control over resources to various administrators. |

**DNS Support for Active Directory**

In Active Directory, DNS is the means by which directory clients locate, or discover, domain controllers. The primary components of the architecture for DNS support of Active Directory include the domain controller Locator, Active Directory domain names in DNS, and Active Directory DNS objects.

The following table describes the Active Directory components that help directory clients locate nearby domain controllers.

**Active Directory DNS Support Components**

| Active Directory/DNS Component | Description |
|---|---|
| Locator | Locator, which is implemented in the Net Logon service, enables a client to locate a domain controller. Locator contains Internet Protocol (IP)/DNS–compatible and Windows NT 4.0–compatible locators, which provide interoperability in a mixed Active Directory environment. |
| Active Directory domain names in DNS | Every Active Directory domain has a DNS domain name (for example, cohovineyard.com), and every domain joined computer has a DNS name (for example, server1.cohovineyard.com). Architecturally, domains and computers are represented both as objects in Active Directory and as nodes in DNS. |
| Active Directory DNS objects | When DNS data is stored in Active Directory, each DNS zone is an Active Directory container object (class **dnsZone**). The dnsZone object contains a DNS node object (class **dnsNode**) for every unique name in that zone. The dnsNode object has a **dnsRecord** multivalued attribute that contains a value for every resource record that is associated with that DNS name. |

For more information about DNS support for Active Directory, see "DNS Support for Active Directory Technical Reference."

**Active Directory Schema**

Everything that is stored in Active Directory is stored in an object. A definition for every type of object is stored in the schema. The definitions themselves consist of two types of objects: class objects and attribute objects. Classes define groups of attributes that are used to describe common objects. New object definitions are created by combining various class objects and attribute objects to make new combinations that contain the necessary attributes to meet the storage requirements of the new object type. The two main types of object definitions that are stored in the Active Directory schema are described in the following table.

**Schema Components**

| Component | Description |
|---|---|
| classSchema objects | classSchema objects are object definitions that are stored in the schema, and they are used to define classes. classSchema objects define groups of attributes that have something in common. For example, an object that is used to store a user account needs to store the user's logon name, first name, last name, and password. It is possible to create a **user** class that has a logon name attribute, a first name attribute, a last name attribute, and a password attribute. Anytime a new user account is created, the directory uses the **user** class as the definition, and every user account object that is created uses those attributes. classSchema objects can be nested to create more complex objects. |
| attributeSchema objects | attributeSchema objects define the individual attributes of a single object. For example, a user account object has a number of attributes that are used to store and define various pieces of data that are related to a user account, such as a logon name attribute and a password attribute. Each of these attributes also has its own attributes that specify the type of data that it stores, the syntax of the data that it stores, and whether or not the attribute is required or optional. The directory service uses attributeSchema objects to store data and verify that the stored data is valid. |

**Active Directory Data Store**

The Active Directory data store is implemented on every domain controller in the forest. The data store consists of components that store and retrieve data inside the directory. The components of the Active Directory data store are described in the following table.

**Data Store Components**

| Component | Description |
|---|---|
| Interfaces (LDAP, REPL, MAPI, SAM) | The data store interfaces provide a way for directory clients and other directory servers to communicate with the data store. |
| DSA (Ntdsa.dll) | The DSA (which runs as Ntdsa.dll on each domain controller) provides the interfaces through which directory clients and other directory servers gain access to the directory database. In addition, the DSA enforces directory semantics, maintains the schema, guarantees object identity, and enforces data types on attributes. |
| Database layer | The database layer is an application programming interface (API) that resides in Ntdsa.dll and provides an interface between applications and the directory database to protect the database from direct interaction with applications. Calls from applications are never made directly to the database; they go through the database layer. In addition, because the directory database is flat — with no hierarchical namespace — the database layer provides the database with an abstraction of an object hierarchy. |
| ESE (Esent.dll) | The ESE (which runs as Esent.dll) communicates directly with individual records in the directory database on the basis of an object's relative distinguished name attribute. |
| Database files | The data store stores directory information in a single database file. In addition, the data store also uses log files, to which it temporarily writes uncommitted transactions. |

# Domains and Forests Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

Domains and forests are logical representations of the directory hierarchy in the Active Directory directory service. Active Directory domains and forests allow administrators to organize elements of a network (such as users, computers, devices, resources, and application-specific data from Active Directory–enabled applications) into a non-physical hierarchical structure, known as the logical structure. You can use domains and forests to provide an information repository and services to make information available to users and applications.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

**In this subject**


- What Are Domains and Forests?
- How Domains and Forests Work
- Domains and Forests Tools and Settings

# What Are Domains and Forests?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**What Are Domains and Forests?**

**In this section**


- The Logical Structure of Active Directory
- The Physical Structure of Active Directory
- What Are Domains?
- What Are Forests?
- Related Information


The Active Directory directory service is a distributed database that stores and manages information about network resources, as well as application-specific data from directory-enabled applications. Active Directory allows administrators to organize objects of a network (such as users, computers, and devices) into a hierarchical collection of containers known as the logical structure. The top-level logical container in this hierarchy is the forest. Within a forest are domain containers, and within domains are organizational units.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

Understanding domains and forests requires understanding the possible relationships they might have in Active Directory. The relationships between these logical containers might be based on administrative requirements, such as delegation of authority, or they might be defined by operational requirements, such as the need to provide for data isolation. Service administrators can use domain and forest containers to meet a number of specific requirements, including:

- Implementing an authentication and authorization strategy for sharing resources across a network
- Providing a mechanism for centralizing management of multiple domains and forests
- Providing an information repository and services to make information available to users and applications
- Organizing objects of a network (such as users, computers, devices, resources, and application specific data from directory-enabled applications) into a non-physical hierarchical structure

To learn more about domains and forests, you must first understand the logical and physical structures of Active Directory. This section describes how those structures differ, and defines domains and forests in terms of the logical structure.

**The Logical Structure of Active Directory**

Active Directory stores network object information and implements the services that make this information available and usable to users. Active Directory presents this information through a standardized, logical structure that helps you establish and understand the organization of domains and domain resources in a useful way. This presentation of object information is referred to as the logical structure because it is independent of the physical aspects of the Active Directory infrastructure, such as the domain controllers required for each domain in the network.

**Benefits of the Logical Structure**

The logical structure provides a number of benefits for deploying, managing, and securing network services and resources. These benefits include:

- Increased network security. The logical structure can provide security measures such as autonomy for individual groups or complete isolation of specific resources.

- Simplified network management. The hierarchical nature of the logical structure simplifies configuration, control, and administration of the network, including managing user and group accounts and all network resources.

- Simplified resource sharing. The logical structure of domains and forests and the relationships established between them can simplify the sharing of resources across an organization.

- Low total cost of ownership. The reduced administration costs for network management and the reduced load on network resources that can be achieved with the Active Directory logical structure can significantly lower the total cost of ownership.

An efficient Active Directory logical structure also facilitates the system integration of features such as Group Policy, enabling desktop lockdown, software distribution, and administration of users, groups, workstations, and servers. In addition, the logical structure can facilitate the integration of services such as Exchange 2000, public key infrastructure (PKI), and domain-based distributed file system (DFS).

**Components of the Logical Structure**

The logical structure consists of leaf object and container object components that must conform to the hierarchical structure of an Active Directory forest. Leaf objects are objects that have no child objects, and are the most basic component of the logical structure. Container objects store other objects and occupy a specific level within the forest hierarchy.

The relationships among the components of the logical structure control access to stored data and determine how that data is managed across one or more domains within a single forest. The components that make up the Active Directory logical structure are described in the following table.

**Components of the Active Directory Logical Structure**

| Component | Description |
|---|---|
| Organizational Units | Organizational units are container objects. You use these container objects to arrange other objects in a manner that supports your administrative purposes. By arranging objects in organizational units, you make it easier to locate and manage them. You can also delegate the authority to manage an organizational unit. Organizational units can be nested in other organizational units.<br><br>You can arrange objects that have similar administrative and security requirements into organizational units. Organizational units provide multiple levels of administrative authority, so that you can apply Group Policy settings and delegate administrative control. This delegation simplifies the task of managing these objects and enables you to structure Active Directory to fit your organization's requirements. |
| Domains | Domains are container objects. Domains are a collection of administratively defined objects that share a common directory database, security policies, and trust relationships with other domains. In this way, each domain is an administrative boundary for objects. A single domain can span multiple physical locations or sites and can contain millions of objects. |
| Domain Trees | Domain trees are collections of domains that are grouped together in hierarchical structures. When you add a domain to a tree, it becomes a child of the tree root domain. The domain to which a child domain is attached is called the parent domain.<br><br>A child domain might in turn have its own child domain. The name of a child domain is combined with the name of its parent domain to form its own unique Domain Name System (DNS) name such as Corp.nwtraders.msft. In this manner, a tree has a contiguous namespace. |
| Forests | A forest is a complete instance of Active Directory. Each forest acts as a top-level container in that it houses all domain containers for that particular Active Directory instance. A forest can contain one or more domain container objects, all of which share a common logical structure, global catalog, directory schema, and directory configuration, as well as automatic two-way transitive trust relationships. The first domain in the forest is called the forest root domain. The name of that domain refers to the forest, such as Nwtraders.msft. By default, information in Active Directory is shared only within the forest. In this way, the forest is a security boundary for the information that is contained in that instance of Active Directory. |

| | |
|---|---|
| Site Objects | Sites are leaf and container objects. The sites container is the topmost object in the hierarchy of objects that are used to manage and implement Active Directory replication. The sites container stores the hierarchy of objects that are used by the Knowledge Consistency Checker (KCC) to effect the replication topology. Some of the objects located in the sites container include NTDS Site Settings objects, subnet objects, connection objects, server objects, and site objects (one site object for each site in the forest). The hierarchy is displayed as the contents of the Sites container, which is a child of the Configuration container. |

By understanding the purpose and hierarchical structure of these components, you can complete a variety of tasks, including installing, configuring, managing, and troubleshooting Active Directory. Although the logical structure of Active Directory is a hierarchical organization of all users, computers, and other physical resources, the forest and domain form the basis of the logical structure.

Forests, which are the security boundaries of the logical structure, can be structured to provide data and service autonomy and isolation in an organization in ways that can both reflect site and group identities and remove dependencies on the physical topology. Domains can be structured within a forest to provide data and service autonomy (but not isolation) and to optimize replication with a given region.

This separation of logical and physical structures improves manageability and reduces administrative costs because the logical structure is not impacted by changes in the physical structure, such as the addition, removal, or reorganization of users and groups.

**Note**

- You can view and manage components of the logical structure by using the Active Directory Users and Computers, Active Directory Domains and Trusts, and Active Directory Schema Microsoft Management Console (MMC) snap-ins, and other tools.

**The Physical Structure of Active Directory**

The physical structure of Active Directory is represented by a set of physical components which, when configured correctly, can help optimize the transmission of network replication and authentication in ways specifically tailored to fit your physical network. Physical network components, such as domain controllers and physical subnets, are used to facilitate network communication and to set physical boundaries around network resources. More specifically, the physical structure of Active Directory represents all of the physical subnets in your enterprise network, the domain controllers in those physical subnets (commonly referred to as Sites in Active Directory) and the replication connections between the domain controllers.

Sites and subnets are represented in Active Directory by site and subnet objects. Computers on TCP/IP networks are assigned to site objects based on their location in a physical subnet or a set of physical subnets. Physical subnets group computers in a way that identifies their physical proximity on the network. Each site object is associated with one or more subnet objects. Subnet objects are used during the process of domain controller location to find a domain controller in the same site as the computer that is logging on. This information also is used during Active Directory replication to determine the best routes between domain controllers. This enables you to use network bandwidth more efficiently. For example, it ensures that when users log on to the network they are authenticated by the authentication authority nearest to the user, thus reducing the amount of network traffic.

The physical domain controller contains the directory partitions that are replicated. Although the logical and physical structures are defined and configured independently, they have interdependencies that affect replication.

**Note**

- You can view the physical structure of Active Directory by using Active Directory Sites and Services.

For more information about the network components associated with the physical structure of Active Directory, see the "Active Directory Replication Topology Technical Reference."

**What Are Domains?**

Domains are logical directory components that you create to manage the administrative requirements of your organization. The logical structure is based on the administrative requirements of an organization, such as the delegation of administrative authority, and operational requirements, such as the need to control replication. In general, domains are used to control where in the forest replication of domain data occurs and organizational units are used to further organize network objects into a logical hierarchy and delegate control to appropriate administrative support personnel.

A domain is a partition in an Active Directory forest. Partitioning data enables organizations to replicate data only to where it is

needed. In this way, the directory can scale globally over a network that has limited available bandwidth. Domains can also be defined as:

- Containers within a forest
- Units of Policy
- Units of Replication
- Authentication and Authorization Boundaries
- Units of Trust

Each domain has a domain administrators group. Domain administrators have full control over every object in the domain. These administrative rights are valid within the domain only and do not propagate to other domains.

**Domains as Containers Within a Forest**
Within the scope of a forest, a domain is a container. Objects in that container inherently trust each other and the security services located in that same container. Each time you create a new domain container in a forest, a two-way, transitive trust relationship is automatically created between the new domain and its parent domain. Trusts are logical relationships established between domains to allow pass-through authentication in which a trusting domain honors the logon authentications of a trusted domain. Because all domain containers within a forest are joined together by two-way transitive trusts, objects within one domain container also inherently trust all other objects and security services located in every domain container located in that forest.

Domain containers are used to store and manage Active Directory objects, some of which have a physical representation. All of the objects within the domain container are stored in a central database that stores only objects created within that domain. Some objects represent people or groups of people, while others represent computers or network servers. Examples of Active Directory objects that have a physical representation are user, computer, and group objects.

While domains contain objects that represent physical things, they also contain objects that are used to help self-regulate domain operations such as trusted domain objects (TDOs) and site link objects. Domain containers can also hold subordinate containers such as organizational units. The following figure shows where objects are stored within the logical structure of a domain.

**Domain Containment Structure Within a Forest**




**Organizational Units and Objects**
Organizational units are used to group objects, including accounts and resources in a domain, for administrative purposes, such as the application of Group Policy or delegation of authority. Control over an organizational unit, including the objects within it, is determined by the access control lists (ACLs) on the organizational unit and on the objects in the organizational unit.

**Organizational Units**
The organizational unit is a particularly useful type of directory object contained within domains. Each organizational unit is an Active Directory container within a domain into which users, groups, computers, and other organizational units of the domain can be placed. An organizational unit cannot contain objects from other domains.

An organizational unit is the smallest scope or unit to which Group Policy settings can be assigned or to which administrative authority can be delegated. A hierarchy of organizational units can be extended as necessary to model the hierarchy of an

organization within a domain. The administrative model of the organizational unit can be scaled to any size. For more information about Group Policy, see "How Core Group Policy Works."

Administrative authority can be delegated for individual organizational units or for multiple organizational units. Organizational units can be nested to create a hierarchy within a domain and form logical administrative units for users, groups, and resource objects, such as printers, computers, applications, and file shares. The organizational unit hierarchy within a domain is independent of the structure of other domains: Each domain can implement its own hierarchy. Likewise, domains that are managed by the central authority can implement similar organizational unit hierarchies. The structure is flexible, which allows organizations to create an environment that mirrors the administrative model, whether it is centralized or decentralized.

**Objects**

Active Directory objects represent the physical entities that make up a network. An object stores an instance of a class. A class is defined in the Active Directory schema as a specific set of mandatory and optional attributes — that is, an attribute can be present in an object in Active Directory only when that attribute is permitted by the object's class. Classes also contain rules that determine which classes of objects can be superior to (parents of) a particular object of the class. Each attribute is also defined in the directory schema. The attribute definitions determine the syntax for the values the attribute can have.

Creation of an Active Directory object requires specification of values for the attributes of the object in its particular class, consistent with the rules of the directory schema. For example, creating a user object requires specification of alphanumeric values for the user's first and last names and the logon identifier, which are mandatory according to the directory schema. Other non-mandatory values can also be specified, such as telephone number and address.

Applications that create or modify objects in Active Directory use the directory schema to determine what attributes the object must and might have, and what those attributes can look like in terms of data structures and syntax constraints. The directory schema is maintained forest-wide so that all objects created in the directory conform to the same values.

Objects are either container objects or leaf objects. A container object stores other objects, and as such occupies a specific level in a subtree hierarchy. An object class is a container if at least one other class specifies it as a possible superior, so any object class defined in the schema can become a container. A leaf object does not store other objects, so it occupies the endpoint of a subtree. For more information about Active Directory Objects, see "How Domains and Forests Work" and "How the Active Directory Schema Works."

**Domains as Units of Policy**

A domain defines a scope or unit of policy within a forest. Some policy settings apply only to the scope of a domain, that is, the policy settings are domain-wide. Account policies, for example, apply uniformly to all user accounts in the domain. Although a domain is not the smallest unit of policy that can be assigned (policies can be assigned to organizational units) it is the most commonly used unit when splitting administrative duties between departments and subsidiaries located in different geographical locations. As shown in the following figure, you might need to create multiple domains to provide for policy variance among domains throughout a forest. A domain is also considered a unit of access control, in that it can be used for business groups requiring general autonomy.

**Delegation of Domains to Domain Admins that Require Different Policies or Autonomy**




You cannot define different account policies for different organizational units in the same domain. Policy settings are stored as Group Policy objects (GPOs) in Active Directory. A GPO establishes how domain resources can be accessed, configured, and used. The policies associated with a GPO are applied only within the domain and not across domains. A GPO can be associated with one or more Active Directory containers, such as a site, domain, or organizational unit.

Account policies and Public Key policies have domain-wide scope and are set at the domain GPO level. All other policies can be specified at the level of the organizational unit. Some policies that can be applied only at the domain container level include:

- Password policy. Determines the rules, such as password length, that must be met when a user sets a password.
- Account lockout policy. Defines rules for intruder detection and account deactivation.
- Kerberosticket policy. Determines the lifetime of a Kerberos ticket. A Kerberos ticket is obtained during the logon process and is used for network authentication. A particular ticket is only valid for the lifetime specified in the policy.

Because organizational units provide multiple levels of administrative authority, you can use them to systematically apply Group Policy settings and delegate administrative control. Delegation simplifies the task of managing these objects and enables you to structure Active Directory to fit the requirements of your organization. Using delegated authority in conjunction with GPOs and group memberships allows one or more administrators to be assigned rights and permissions to manage objects in an entire domain or in one or more organizational units within the domain.

For more information about Group Policy, see "How Core Group Policy Works."

**Domains as Units of Replication**

The physical significance of a domain is that it is a unit of replication. In fact, with the exception of application partitions, which replicate only non-security principal objects, the domain is the smallest unit of replication that can be administered within an Active Directory forest. This is because all objects that are located within a domain container, also referred to as domain data, are replicated to other domain controllers within that same domain, regardless of whether those domain controllers are located over a local area network (LAN) or wide area network (WAN) connection.

As shown in the following figure, Active Directory multi-master replication manages the transfer of domain objects to the appropriate domain controllers automatically, keeping domain data up-to-date among all domain controllers in the domain, regardless of location.

**Replication of Domain Data Within Each Domain in the Forest**




All domain controllers in the forest are also updated with changes to forest-wide data. For more information about replication at the Forest level, see "Forests as Units of Replication" later in this section Domain-wide replication relies on three components of the Active Directory physical structure to be in place for optimal performance, these include:

**Domain Controllers**

The physical domain controller contains the domain partition data that is replicated to other domain controllers in that same domain. A domain partition stores only the information about objects located in that domain. All domain controllers in a domain receive changes and replicate those changes to the domain partition stored on all other domain controllers in the domain. In this way, all domain controllers are peers in the domain and manage replication as a unit.

Domain controllers also have two non-domain directory partitions that store forest-wide data, which includes the directory schema and configuration data. Optionally, domain controllers, application partitions can be configured to be located on designated domain controllers anywhere in a forest. Unlike the schema and configuration partitions, application partitions are not located on every domain controller in a forest. For more information about the replication of forest-wide data, see "Forests as Units of Replication" later in this section.

Partitioning Active Directory data into three physical partitions on each domain controller helps to control replication so that data is replicated only to where it is needed. In this way, Active Directory can scale globally over a network that has limited available bandwidth.

**Sites**

Within the scope of a forest, sites are a representation of the physical network topology. This includes physical subnet and site definitions. Replication of updates to domain data occurs between multiple domain controllers to keep replicas synchronized. Multiple domains are common in large organizations, as are multiple sites in disparate locations. In addition, domain controllers for the same domain are commonly placed in more than one site.

Therefore, replication must often occur both within sites and between sites to keep domain and forest data consistent among domain controllers that store the same directory partitions. Replication within sites generally occurs at typical LAN speeds between domain controllers that are on the same network segment. Replication between sites usually occurs over WAN links that might be costly in terms of bandwidth. To accommodate the differences in distance and cost of replication within a site and between sites, the intrasite replication topology is used to optimize speed, and the intersite replication topology is used to minimize cost based on a configurable replication schedule.

The Knowledge Consistency Checker (KCC) is a distributed application that runs on every domain controller and is responsible for creating the connections between domain controllers that collectively form the replication topology. The KCC uses Active Directory data to determine where to create connections between source domain controllers and destination domain controllers.

Intersite replication is optimized for bandwidth efficiency, and directory updates between sites occur automatically based on a configurable schedule.

**Domain-Wide Operations Masters**

Active Directory supports multi-master replication of the directory data store between all domain controllers in the domain. However, some changes are impractical to perform in multi-master fashion, so only a single domain controller, called the operations master, accepts requests for such changes. Because these roles can be transferred to other domain controllers within the domain or forest, they are sometimes referred to as operations master roles.

There are three operations master roles per domain. These include the Relative Identifier (RID) Master, Primary Domain Controller (PDC) emulator, and Infrastructure Master. These three roles must be unique in each domain, so each domain can have only one RID master, one PDC emulator, and one Infrastructure Master. For information about forest-wide roles, see "Forest-Wide Operations Master Roles" later in this section. For more information about replication, see "How Active Directory Replication Topology Works."

**Domains as Authentication and Authorization Boundaries**

By default, a domain provides authentication and authorization services for all its accounts in order to facilitate logons and single sign-on access control to shared resources within the boundaries of the domain. The process of authentication ensures that users are who they claim to be. Once identified, the user can be authorized access to a specific set of network resources based on permissions.Authorization takes place through the mechanism of access control, using access control lists (ACLs) that define permissions on file systems, network file and print shares, and entries in Active Directory.

Authorization is determined not only by the identity of the user but also the membership of the user in one or more security groups. In fact, the preferred method of controlling domain-wide resources is to grant access to groups rather than to individuals, adjusting the level of authorization for a group according to the needs of its members. This method of controlling access makes it easier to keep ACLs up-to-date on domains that have thousands of users and objects. Group membership can be managed centrally by anyone with the appropriate administrative credentials. You can change the level of authorization a particular user has to many resources simply by adding or removing the user from a group. The following figure shows when authentication and authorization for a user in a given domain occur.

**Authentication and Authorization of a User Within the Same Domain**




Authentication is not limited to users. Computers and services are also authenticated when they make network connections to other servers. For example, domain joined computers connect Active Directory in their domain for policy information during startup. Computers and services also prove their identity to clients that request mutual authentication. Mutual authentication prevents an intruder from adding another computer as an imposter between the client and the real network server. The Kerberos version 5 authentication protocol is a technology for single sign-on to network resources. Kerberos V5 protocol is used to provide single sign-on to network services within a domain, and to services residing in trusted domains. The Kerberos V5 protocol verifies both the identity of the user and of the network services, providing mutual authentication.

Authentication and authorization services in one domain can be extended to accounts that are located in trusted domains. This can be done by using trusts. Trusts are logical relationships established between domains to allow pass-through authentication in which a trusting domain honors the logon authentications of a trusted domain. Consequently, trust relationships inherently

allow domain-specific authentication and authorization services to be extended, thereby enabling single sign-on access control capabilities throughout a forest. Because the domain trust relationships between all domains in the forest are bidirectional by default, authentication in one domain is sufficient for referral or pass-through authentication to resources in other domains in the forest.

**Domains as Units of Trust**

A domain is the smallest container within Active Directory that can be used in a trust relationship. All domains within a forest are connected by Kerberos-based trusts. Kerberos-based trusts are two-way and transitive in nature. Trusts act as authentication pipelines that must be present in order for users in one domain to access resources in another domain. All authentication requests made between domains located inside or outside of a forest must occur over trusts. Trust relationships within a forest are created as implicit two-way transitive trusts.

### Note

- "Access to resources" in any discussion of trust relationships always assumes the limitations of access control.

Within a forest, trust relationships are automatically created between the forest root domain and any tree root domains on the one hand, and all child domains that are subordinate to the forest root domain on the other. Because trust relationships are two-way and transitive by default, users and computers can be authenticated between any domain containers in the forest, as shown in the following figure.

**Domains as Units of Trust and the Authentication Paths they Provide**




In accordance with DNS naming standards, Active Directory domains within a single forest are created in an inverted tree structure, with the forest root domain at the top. This tree structure requires that trusts exist between domains to facilitate security of communications. Adding a domain tree to a forest requires specification of the forest root domain, which results in the establishment of a trust relationship between the root domain of the new tree and the forest root domain. For more information about trusts and root domains, see "Forests as Collections of Domain Containers that Trust Each Other" later in this section.

**What Domains Are Not**

Domains are not security boundaries within the scope of Active Directory and do not provide complete isolation in the face of possible attacks by malicious service administrators who might manage that forest. This is because a malicious service administrator, such as a member of the Domain Admins group, can use nonstandard tools and procedures to gain full access to any domain in the forest or to any computer in the forest. For example, service administrators in a non-root domain can make themselves members of the Enterprise Admins or Schema Admins group.

However, administrative rights do not flow across domain boundaries, nor do they flow down through a domain tree. For

example, if you have a domain tree with domains A, B, and C, where A is the parent domain of B and B is the parent domain of C, users with administrative rights in domain A do not have administrative rights in B, nor do users with administrative rights in domain B have administrative rights in domain C. For a user to obtain administrative rights in a given domain, a higher authority must grant them. This does not mean, however, that an administrator cannot have administrative rights in multiple domains; it simply means that all rights must be explicitly defined.

For more information about isolation, see "Forests as Security Boundaries" later in this section.

**What Are Forests?**
At its highest level, a forest is a single instance of Active Directory. Therefore, a forest is synonymous with Active Directory, meaning that the set of all directory partitions in a particular Active Directory instance (which includes all domain, configuration, schema and optional application information) makes up a forest. This means that when you have multiple forests in an enterprise they will, by default, act separately from each other as if they were the only directory service in your organization.

This behavior, however, is easily be modified so that multiple forests can share Active Directory responsibilities across an enterprise. This is done by creating external or forest trust relationships between the forests. In this way, each forest can be connected with every other forest to form a collaborative directory service solution for any enterprise with business needs that include multiple forest collaboration.

Forests can also be defined as:

- Collections of Domain Containers that Trust Each Other
- Units of Replication
- Security Boundaries
- Units of Delegation

**Forests as Collections of Domain Containers that Trust Each Other**
Within the scope of Active Directory, a forest is a collection of domain containers that inherently trust each other and other security services that are located in that same forest. All domain containers in a forest share a common global catalog, directory schema, and directory configuration, as well as automatic two-way transitive trust relationships. Because all of the domain containers are inherently joined through two-way transitive trusts, all authentication requests made from any domain in the forest to any other domain in the same forest will be granted. In this way, all security principals that are located in domain containers within a forest inherently trust each other.

Forests can be used to segregate domain containers into one or more unique DNS namespace hierarchies known as domain trees. Although each domain tree consists of a unique namespace the schema and configuration data for the forest are shared throughout all the domain containers in a forest irrespective of namespace. Each domain container in a forest must adhere to DNS naming schemes and all domains are organized in a root and subordinate domain hierarchy. Root domains, such as forest root and tree root domains, define the DNS namespace for their subordinate domains. Although a forest can consist of multiple domain trees, it represents one organization. However, an organization can have multiple forests. For more information about multiple forests, see "Forests as Units of Delegation" later in this section. As shown in the following figure, the namespace for each root domain must be unique.

**Domain Containers, Root Domains and DNS Namespaces Within a Forest**




**Forest Root Domain**

Although trees in a forest do not share the same DNS namespace, a forest does have a single root domain, called the forest root domain. The forest root domain is, by definition, the first domain created in the forest. The Enterprise Admins and Schema Admins groups are located in this domain. By default, members of these two groups have forest-wide administrative credentials. In Windows 2000 Active Directory, the forest root domain cannot be deleted, changed, or renamed. In Windows Server 2003 and later versions of Active Directory, the forest root domain cannot be deleted, but it can be restructured or renamed.

Objects are located within Active Directory domains according to a hierarchical path, which includes the labels of the Active Directory domain name and each level of container objects. The full path to the object is defined by the distinguished name (also known as a DN). The distinguished name is unambiguous (identifies one object only) and unique (no other object in the directory has this name). By using the full path to an object, including the object name and all parent objects to the root of the domain, the distinguished name uniquely and unambiguously identifies an object within a domain hierarchy.

The distinguished name of the forest root domain is used to locate the configuration and schema directory partitions in the namespace. The distinguished names for the Configuration and Schema containers in Active Directory always show these containers as child objects in the forest root domain. For example, in the child domain Noam.wingtiptoys.com (where Wingtiptoys.com is the name of the forest root domain), the distinguished name of the Configuration container is cn=configuration,dc=wingtiptoys,dc=com. The distinguished name of the Schema container is cn=schema,cn=configuration,dc=wingtiptoys,dc=com. However, this naming convention provides only a logical location for these containers.

These containers do not exist as child objects of the forest root domain, nor is the schema directory partition actually a part of the configuration directory partition: They are separate directory partitions. Every domain controller in a forest stores a copy of the configuration and schema directory partitions, and every copy of these partitions has the same distinguished name on every domain controller.

**Tree Root Domain**

A domain tree represents a unique DNS namespace: it has a single root domain, known as the tree root domain, and is built as a strict hierarchy: Each domain above the tree root domain has exactly one superior, or parent, domain (the forest root domain). The namespace created by this hierarchy, therefore, is contiguous — each level of the hierarchy is directly related to the level above it and to the level below it. In other words, the names of domains created beneath the tree root domain (child domains) are always contiguous with the name of the tree root domain. The DNS names of the child domains of the root domain of a tree reflect this organization; therefore, the children of a root domain called Somedomain are always children of that domain in the DNS namespace (for example, Child1.somedomain, Child2.somedomain, and so forth).

Multiple domain trees can belong to a single forest and do not form a contiguous namespace; that is, they have noncontiguous DNS domain names. Although the roots of the separate trees have names that are not contiguous with each other, the trees share a single overall namespace because names of objects can still be resolved by the same Active Directory instance. A forest exists as a set of cross-reference objects and trust relationships that are known to the member trees. Transitive trusts at the root domain of each namespace provide mutual access to resources.

The domain hierarchy in a forest determines the transitive trust links that connect each domain. Each domain has a direct trust link with its parent and each of its children. If there are multiple trees in a forest, the forest root domain is at the top of the trust tree and, from a trust perspective, all other tree roots are children. That means authentication traffic between any two domains in different trees must pass through the forest root.

**Forests as Units of Replication**

Unlike domains, forests are the largest unit of replication that can be administered in an Active Directory environment. To efficiently synchronize data between domain controllers throughout a forest, Active Directory replication transfers updates according to directory partition. Directory partitions are used to help organize how replication occurs within a forest. Active Directory uses four distinct directory partition types to store and copy four different types of data.

One directory partition contains domain data; the other three are forest-wide partitions, and contain configuration, schema, and application data. This storage and replication design provides directory information to users and administrators throughout the domain. Each domain controller receives directory updates to the data stored in its domain only, as well as updates to the data in the two directory partitions that store configuration and schema data for the forest. As shown in the following figure, schema and configuration data must be replicated to all domain controllers that exist in a forest, while optional Application data is replicated only to domain controllers within the same forest that are specified as replication partners.

**Forest-Wide Replication Scope**




The following table describes each of the three forest-wide partitions in more detail.

**Forest-Wide Directory Partitions**

| Partition | Description |
|---|---|

| | |
|---|---|
| Schema | The schema partition contains the forest-wide schema. Each forest has one schema so that the definition of each object class is consistent. The schema is the formal definition of all object and attribute data that can be stored in the directory. Active Directory domain controllers include a default schema that defines many object types, such as user and computer accounts, groups, domains, organization units, and security policies. Administrators and programmers can extend the schema by defining new object types and attributes or by adding new attributes for existing objects. Schema objects are protected by access control lists, ensuring that only authorized users can alter the schema. The schema partition is replicated to each domain controller in the forest. |
| Configuration | The configuration partition describes the topology of the forest as well as other forest, domain and domain controller settings. This configuration data includes a list of all domains, trees, and forests and the locations of the domain controllers and global catalogs. The configuration partition is replicated to each domain controller in the forest. |
| Application | Applications and services can use application directory partitions to store application-specific data. Data stored in the application directory partition is intended to satisfy cases where information needs to be replicated but not necessarily on a global scale. Application directory partitions can contain any type of object, except security principals. Application directory partitions are usually created by the applications that will use them to store and replicate data. One of the benefits of an application directory partition is that, for redundancy, availability, or fault tolerance, the data in it can be replicated to different domain controllers in a forest. The data can be replicated to a specific domain controller or any set of domain controllers anywhere in the forest. This differs from a domain directory partition in which data is replicated to all domain controllers in that domain.<br><br>Only domain controllers running Windows Server 2003 or higher can host a replica of an application directory partition. Application directory partitions are created, configured, and managed by the administrator. |

All forest replication is Multi-Master with the exception of the three domain-wide and two forest-wide operations master roles. Forest-wide replication requires domain controllers and three other components of the Active Directory physical structure to be in place for optimal performance. These components are forest-wide operations master roles, sites, and global catalogs.

**Forest-Wide Operations Master Roles**

There are two operations master roles assigned for the entire forest. These include the domain naming master and the schema master. Each forest must have one and only one schema master and one domain naming master, so these two roles must remain in the forest root domain. The other three operations master roles are assigned to each domain in the forest. These three roles must be unique in each domain, so each domain can have only one relative ID master, one PDC emulator, and one infrastructure master.

In an Active Directory forest with only one domain and one domain controller, that domain controller has all the operations master roles. For more information about operations master roles, see "How Operations Masters Work."

**Sites**

Sites in Active Directory represent the physical structure, or topology, of the network. Within the scope of a forest, sites are a representation of the physical network topology, including physical subnet and site definitions. When you establish sites, domain controllers within a single site communicate frequently. This communication minimizes the latency within the site; that is, the time required for a change that is made on one domain controller to be replicated to other domain controllers. You create sites to optimize use of the bandwidth between domain controllers in different locations. For more information about sites, see "How Active Directory Replication Topology Works."

**Global Catalogs**

The global catalog stores a full copy of all Active Directory objects in the directory for its host domain and a partial copy of all objects for all other domains in the forest. Users in a forest do not need to be aware of directory structure because all users see a single directory through the global catalog. Applications and clients can query the global catalog to locate any object in a forest. The global catalog is hosted on one or more domain controllers in the forest. It contains a partial replica of every domain directory partition in the forest. These partial replicas include replicas of every object in the forest, including the attributes most frequently used in search operations and the attributes required to locate a full replica of the object. A global catalog is created automatically on the first domain controller in the forest. Optionally, other domain controllers can be configured to serve as global catalogs.

A global catalog serves the following roles:

- Enables user searches for directory information about objects throughout all domains in the forest
- Resolves user principal names (UPNs) during authentication, when the authenticating domain controller does not have information about the account

- Supplies universal group membership information in a multiple domain environment
- Validates references to objects in other domains in the forest

For more information about global catalogs and their roles, see "How the Global Catalog Works."

**Forests as Security Boundaries**

Each forest is a single instance of the directory, the top-level Active Directory container, and a security boundary for all objects that are located in the forest. This security boundary defines the scope of authority of the administrators. In general, a security boundary is defined by the top-level container for which no administrator external to the container can take control away from administrators within the container. As shown in the following figure, no administrators from outside a forest can control access to information inside the forest unless first given permission to do so by the administrators within the forest.

**Enterprise Administrators Outside of a Forest Can Administer Only Within Their Own Forest**




A forest is the only component of the Active Directory logical structure that is a security boundary. By contrast, a domain is not a security boundary because it is not possible for administrators from one domain to prevent a malicious administrator from another domain within the forest from accessing data in their domain. A domain is, however, the administrative boundary for managing objects, such as users, groups, and computers. In addition, each domain has its own individual security policies and trust relationships with other domains.

**Forests as Units of Delegation**

The forest is the largest management unit of Active Directory as well as the ultimate unit of autonomy and isolation of authority. Members of the Enterprise Admins and forest root Domain Admins security groups are known as enterprise administrators. Enterprise administrators control the Active Directory objects in the configuration container that do not belong to any one domain, including Enterprise Certification Authority objects and other configuration data that affects all domains in the forest.

Because enterprise administrators have authority over all domains in the forest, the domain administrators in each domain must trust the enterprise administrators. You cannot truly restrict enterprise administrators from managing objects in any domain in the forest. Enterprise administrators can always regain control over objects. Some organizations with political challenges, such as those frequently encountered in mergers and acquisitions, might find the scope of this enterprise authority too great and require more than one forest. If your organization requires strict isolation of authority between domains, you must deploy multiple forests with manually created trusts between domains in the different forests.

Because each forest is administered separately, adding additional forests to your organization increases the management overhead of the organization. The number of forests that you need to deploy is based on the autonomy and isolation requirements of each group within your organization. Reasons to create multiple forests in your organization include:

- To secure data within each forest. Sensitive data can be protected so that only users within that forest can access it.
- To isolate directory replication within each forest. Schema changes, configuration changes, and the addition of new domains to a forest have forest-wide impact only within that forest, not on a trusting forest.

Because the forest is a security boundary, each forest does not trust or allow access from any other forest by default. However, in Windows Server 2003 and higher Active Directory, transitive trust relationships can be manually established between forests to establish cross-forest access to resources, so that users in one forest can access resources in another forest.

Forest trusts help you to manage a segmented Active Directory infrastructure within your organization by providing support for accessing resources and other objects across multiple forests. By using forest trusts, you can link two different Windows Server 2003 or higher forests to form a one-way or two-way transitive trust relationship. A forest trust allows administrators to federate two Active Directory forests with a single trust relationship to provide a seamless authentication and authorization experience across the forests.

When two forests are connected by a forest trust, authentication requests made using the Kerberos V5 or NTLM protocols can be routed between forests to provide access to resources in both forests. Trust security settings and firewalls can be configured between domains in different forests that are joined by trust relationships to limit cross-forest connectivity to specific users or applications. For more information about trust security settings, see "Security Considerations for Trusts."

A forest trust can be created only between a forest root domain in one Windows Server 2003 or higher forest and a forest root domain in another Windows Server 2003 or higher forest. Forest trusts can be created between two forests only and cannot be implicitly extended to a third forest. This means that if a forest trust is created between Forest 1 and Forest 2, and another forest trust is created between Forest 2 and Forest 3, Forest 1 does not have an implicit trust with Forest 3. The following figure shows two separate forest trust relationships between three forests in a single organization.

**Two Forest Trusts Between Three Windows Server 2003 Forests**




**Related Information**

The following resources contain additional information that is relevant to this section.

- How Active Directory Replication Topology Works
- How Core Group Policy Works
- How Domains and Forests Work
- How Operations Masters Work
- How the Active Directory Schema Works
- How the Global Catalog Works
- Security Considerations for Trusts

# How Domains and Forests Work

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**How Domains and Forests Work**

**In this section**


- Domain and Forest Structure
- Active Directory Objects
- Domain and Forest Protocols and Programming Interfaces
- Domain and Forest Processes and Interactions
- Network Ports Used by Domains and Forests
- Related Information


Active Directory stores data for an entire forest. A forest is a distributed database, which is made up of directory partitions spread across multiple computers. A domain is one partition of the database; each domain contains Active Directory objects, such as security principal objects (users, computers, and groups) to which you can grant or deny access to network resources. All domain data stored in the domain directory partition is replicated to domain controllers in that domain only.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

The directory partitions that store configuration and schema information are replicated to domain controllers in all domains. In this way, Active Directory provides a data repository that is logically centralized but physically distributed. Because all domain controllers store forest-wide configuration and schema information, a domain controller in one domain can reference a domain controller in any other domain if the information that it is requesting is not stored locally.

This section details how domains and forests work, discusses the various components of domains and forests, describes several common processes that depend on domains and forests, and lists the network ports related to domains and forests.

**Domain and Forest Structure**

A forest consists of a hierarchical structure of domain containers that are used to categorically store information about objects on the network. Domain containers are considered the core functional units in the forest structure. This is because each domain container in a forest is used primarily to store and manage Active Directory objects, most of which have a physical representation (such as people, printers, or computers). Forests also provide the structure by which domain containers can be segregated into one or more unique Domain Name System (DNS) namespace hierarchies known as domain trees.

In addition, the domain tree hierarchy is based on trust relationships — that is, the domain containers are linked by intra-forest trust relationships. When it is necessary for domain containers in the same organization to have different namespaces, you can create a separate tree for each namespace. In Active Directory, the roots of trees are linked automatically by two-way, transitive trust relationships. Trees linked by trust relationships form a forest. A single tree that is related to no other trees constitutes a forest of one tree. The domain and forest structure is made up of the following components:

- Cross-References
- Trust Relationships
- Forest Root

- Domain Trees and Child Domains
- Domain Names

For more information about Active Directory Domains and DNS, see "How DNS Support for Active Directory Works."

This section describes the structure and function of these components, and describes how this structure helps administrators manage the network so that users can accomplish business objectives.

**Cross-References**

Cross-references enable every domain controller to be aware not only of the partitions that it holds, but of all directory partitions in the forest. The information contained within cross-references form the glue that holds the pieces of the domain and forest structure together. Because Active Directory is logically partitioned, and directory partitions are the discrete components of the directory that replicate between domain controllers, either all objects in a directory partition are present on a particular domain controller or no objects in the directory partition are present on the domain controller. For this reason, cross references have the effect of linking the partitions together, which allows operations such as searches to span multiple partitions.

Cross-references are stored as directory objects of the class **crossRef** that identify the existence and location of all directory partitions, irrespective of location in the directory tree. In addition, these objects contain information that Active Directory uses to construct the directory tree hierarchy. Values for the following attributes are required for each cross-reference:

- **nCName .** The distinguished name of the directory partition that the **crossRef** object references. (The prefix **nC** stands for naming context, which is a synonym for directory partition.) The combination of all of the **nCName** properties in the forest defines the entire directory tree, including the subordinate and superior relationships between partitions.
- **dNSRoot.** The DNS name of the domain where servers that store the particular directory partition can be reached. This value can also be a DNS host name.

**How Cross-Reference Information is Propagated Throughout the Domain and Forest Structure**

For every directory partition in a forest, there is an internal cross-reference object stored in the Partitions container (cn=Partitions,cn=Configuration,dc=*ForestRootDomain*). Because cross-reference objects are located in the Configuration container, they are replicated to every domain controller in the forest, and thus every domain controller has information about the name of every partition in the forest. By virtue of this knowledge, any domain controller can generate referrals to any other domain in the forest, as well as to the schema and configuration directory partitions.

When you create a new forest, the Active Directory Installation Wizard creates three directory partitions: the first domain directory partition, the configuration directory partition, and the schema directory partition. For each of these partitions, a cross-reference object is created automatically. Thereafter, when a new domain is created in the forest, another directory partition is created and the respective cross-reference object is created. When the configuration directory partition is replicated to the new domain controller, a cross-reference object is created on the domain naming master and is then replicated throughout the forest.

**Note**

- The state of cross-reference information at any specific time is subject to the effects of replication latency.

For more information about cross-reference objects, see "How Active Directory Searches Work." Cross-reference objects can also be used to generate referrals to other directory partitions located in another forest through external cross-references.

**External Cross-References**

An external cross-reference is a cross-reference object that can be created manually to provide the location of an object that is not stored in the forest. If your Lightweight Directory Access Protocol (LDAP) clients submit operations for an external portion of the global LDAP namespace against servers in your forest, and you want servers in your forest to refer the client to the correct location, you can create a cross-reference object for that directory in the Partitions container. There are two ways that external cross-references are used:

- To reference external directories by their disjoint directory name (a name that is not contiguous with the name of this directory tree). In this case, when you create the cross-reference, you create a reference to a location that is not a child of any object in this directory.

- To reference external directories by a name that is within the Active Directory namespace (a name that is contiguous with the name of this directory tree). In this case, when you create the cross-reference, you create a referenceto a location that is a child of a real object in this directory.

Because the domain component (dc=) portion of the distinguished names of all Active Directory domains matches their DNS addresses, and because DNS is the worldwide namespace, all domain controllers can generate external referrals to each other automatically.

**Trust Relationships**

Active Directory provides security across multiple domains through intra-forest trust relationships. When there are trust relationships between domains in the same forest, the authentication mechanism for each domain trusts the authentication mechanism for all other trusted domains. If a user or application is authenticated by one domain, its authentication is accepted by all other domains that trust the authenticating domain. Users in a trusted domain have access to resources in the trusting domain, subject to the access controls that are applied in the trusting domain.

**Note**

- Default intra-forest trust relationships are created at the time the domains are created. The number of trust relationships required to connect *n* domains is *n*–1, whether the domains are linked in a single, contiguous parent-child hierarchy or constitute two or more separate contiguous parent-child hierarchies.

A trust relationship is a relationship established between two domains that allows users in one domain to be recognized by a domain controller in the other domain. Trusts let users access resources in the other domain, and also let administrators manage user rights for users in the other domain. Account authentication between domains is enabled by two-way, transitive trust relationships. All domain trusts in an Active Directory forest are two-way and transitive and are have the following attributes:

- **Two-way.** When you create a new child domain, the child domain automatically trusts the parent domain, and vice versa. At the practical level, this means that authentication requests can be passed between the two domains in both directions.

- **Transitive.** A transitive trust reaches beyond the two domains in the initial trust relationship. For example, if Domain A and Domain B (parent and child) trust each other, and if Domain B and Domain C (also parent and child) trust each other, Domain A and Domain C also trust each other (implicitly), even though no direct trust relationship between them exists.

At the level of the forest, a trust relationship is created automatically between the forest root domain and the root domain of each domain tree added to the forest, with the result that complete trust exists between all domains in an Active Directory forest. At the practical level, because trust relationships are transitive, a single logon process lets the system authenticate a user (or computer) in any domain in the forest. As shown in the following figure, this single logon process lets the account access resources on any domain in the forest.

**Transitive Trusts Facilitate Cross-Domain Access to Resources With a Single Logon**




**Note**

- Single logons enabled by trusts do not necessarily imply that the authenticated user has rights and permissions in all domains in the forest. This is because in any discussion of trust relationships, access to resources always assumes the limitations of access control.

**Forest Root Domain**

The first domain created in the forest is called the forest root domain. When you create a new tree, you specify the root domain of the initial tree, and a trust relationship is established between the root domain of the second tree and the forest root domain. If you create a third tree, a trust relationship is established between the root domain of the third tree and the forest root domain. Because all trust relationships created within a forest are transitive and two-way, the root domain of the third tree also has a two-way trust relationship with the root domain of the second tree. In Windows 2000 Active Directory, the forest root domain cannot be deleted, changed, or renamed. In Windows Server 2003 Active Directory or higher, the forest root domain cannot be deleted, but it can be restructured or renamed.

The distinguished name of the forest root domain is used to locate the configuration and schema directory partitions in the namespace. The distinguished names for the configuration and schema containers in Active Directory always show these containers as child objects in the forest root domain. For example, in the child domain Noam.wingtiptoys.com, the distinguished name of the Configuration container is cn=configuration,dc=wingtiptoys,dc=com. The distinguished name of the Schema container is cn=schema,cn=configuration,dc=wingtiptoys,dc=com. However, this naming convention provides only a logical location for these containers.

The containers do not exist as child objects of the forest root domain, nor is the schema directory partition actually a part of the configuration directory partition. They are separate directory partitions. Every domain controller in a forest stores a copy of the configuration and schema directory partitions, and every copy of these partitions has the same distinguished name on every domain controller. The following operations occur when you create the forest root domain:

- The Schema container and the Configuration container are created.

- The Active Directory Installation Wizard assigns the PDC emulator, RID master, domain naming master, schema master, and infrastructure master roles to the domain controller.

The Enterprise Admins and Schema Admins groups are located in this domain. By default, members of these two groups have forest-wide administrative credentials.

**Domain Trees and Child Domains**

A forest is a collection of one or more domain trees, organized as peers and connected by two-way, transitive trust

relationships. A single domain constitutes a tree of one domain, and a single tree constitutes a forest of one tree. Thus, a forest is synonymous with Active Directory — that is, the set of all directory partitions in a particular directory service instance (which includes all domains and all configuration and schema information) makes up a forest.

Trees in the same forest do not form a contiguous namespace. They form a noncontiguous namespace that is based on different DNS root domain names. However, trees in a forest share a common directory schema, configuration, and global catalog. (The global catalog is a domain controller that stores all objects of all domains in an Active Directory forest, which makes it possible to search for objects at the forest level rather than at the tree level.) This sharing of common schema and configuration data, in addition to trust relationships between their roots, distinguishes a forest from a set of unrelated trees. Although the roots of the separate trees have names that are not contiguous with each other, the trees share a single overall namespace because names of objects can still be resolved by the same Active Directory instance.

**Note**

- The directory schema and configuration data are shared because they are stored in separate logical directory partitions that are replicated to domain controllers in every domain in the forest. The data about a particular domain is replicated only to domain controllers in the same domain.

**Domain Trees**

A domain tree is a DNS namespace: it has a single root domain and is built as a strict hierarchy; each domain below the root domain has exactly one superior, or parent, domain. The namespace created by this hierarchy, therefore, is contiguous — each level of the hierarchy is directly related to the level above it and to the level below it, if any. In Active Directory, the following rules determine the way that trees function in the namespace:

- A tree has exactly one name. The name of the tree is the DNS name of the domain at the root of the tree.

- The names of domains created beneath the root domain (child domains) are always contiguous with the name of the tree root domain.

- The DNS names of the child domains of the tree root domain reflect this organization; therefore, the children of a tree root domain called Somedomain are always children of that domain in the DNS namespace (for example, Child1.somedomain, Child2.somedomain, and so forth).

**Note**

- Tree and forest hierarchies are specific to Active Directory domains. A Windows NT 4.0 domain that is configured to trust or to be trusted by a Active Directory domain is not part of the forest to which the Active Directory domain belongs.

The forest structure provides organizations with the option of constructing their enterprise from separate, distinct, noncontiguous namespaces. Having a separate namespace is desirable under conditions where, for example, the namespace of an acquired company should remain intact. If you have business units with distinct DNS names, you can create additional trees to accommodate the names.

**Note**

- Before creating a new domain tree when you want a different DNS namespace, consider creating another forest. Multiple forests provide administrative autonomy, isolation of the schema and configuration directory partitions, separate security boundaries, and the flexibility to use an independent namespace design for each forest.

When you create a new domain tree, you specify the root domain of the initial tree, and a trust relationship is established between the root domain of the new tree (the second tree) and the forest root domain. If you create a third tree, a trust relationship is established between the root domain of the third tree and the forest root domain. Because a trust relationship is transitive and two-way, the root domain of the third tree also has a two-way trust relationship with the root domain of the second tree. The following operations occur when you create a new tree root domain in an existing forest:

- Location of a source domain controller in the forest root domain and synchronization of domain system time with the system time of the source domain controller.

- Creation of a tree-root trust relationship between the tree root domain and the forest root domain, and creation of the trusted domain object (TDO) in both domains. The tree-root trust relationship is two-way and transitive.

**Child Domains**
Child domains can represent geographical entities (for example, the United States and Europe), administrative entities within the organization (for example, sales and marketing departments), or other organization-specific boundaries, according to the needs of the organization. Domains are created below the root domain to minimize Active Directory replication and to provide a means for creating domain names that do not change. Changes in the overall domain architecture, such as domain collapses and domain re-creation, create difficult and potentially IT-intensive support requirements. A good namespace design should be capable of withstanding reorganizations without the need to restructure the existing domain hierarchy.

Each time you create a new child domain, a two-way transitive trust relationship (known as the parent-child trust) is automatically created between the parent and new child domain. In this way, transitive trust relationships flow upward through the domain tree as it is formed, creating transitive trusts between all domains in the domain tree. The parent-child relationship is a naming and trust relationship only. Administrators in a parent domain are not automatically administrators of a child domain. Likewise, policies set in a parent domain do not automatically apply to child domains.

The following operations occur when you create a child domain in an existing tree:

- Verification of the name that you provide as a valid child domain name.
- Location of a source domain controller in the parent domain and synchronization of the system time of the child domain with the system time of the source domain controller.
- Creation of parent-child TDOs in the System folder on both the parent domain and the child domain. The TDO objects identify two-way transitive trust relationships between the child domain and the parent domain.
- Replication of the Active Directory Schema container and the Configuration container from a domain controller in the parent domain.

**Domain Names**
Active Directory uses DNS naming standards for hierarchical naming of Active Directory domains and computers. For this reason, domain and computer objects are part of both the DNS domain hierarchy and the Active Directory domain hierarchy. Although these domain hierarchies have identical names, they represent separate namespaces.

The domain hierarchy defines a namespace. A namespace is any bounded area in which standardized names can be used to symbolically represent some type of information (such as an object in a directory or an Internet Protocol [IP] address) and that can be resolved to the object itself. In each namespace, specific rules determine how names can be created and used. Some namespaces, such as the DNS namespace and the Active Directory namespace, are hierarchically structured and provide rules that allow the namespace to be partitioned. Other namespaces, such as the Network Basic Input/Output System (NetBIOS) namespace, are flat (unstructured) and cannot be partitioned.

The main function of DNS is to map user-readable computer names to computer-readable IP addresses. Thus, DNS defines a namespace for computer names that can be resolved to IP addresses, or vice versa. In Windows NT 4.0 and earlier, DNS names were not required; domains and computers used NetBIOS names, which were mapped to IP addresses by using the Windows Internet Name Service (WINS). Although DNS names are required for Active Directory domains and Windows 2000–, Windows XP–, and Windows Server 2003–based computers, NetBIOS names also are supported in Windows Server 2003 for interoperability with Windows NT 4.0 domains and with clients that are running Windows NT 4.0 or earlier, Windows for Workgroups, Windows 98, or Windows 95. There is no interoperability between Windows Server 2008 and higher domains and Windows NT 4.0 domains.

**Note**

- WINS and NetBIOS are not required in an environment where computers run only Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7 or Windows Server 2008 R2, but WINS is required for interoperability between Windows 2000 Server– and Windows Server 2003–based domain controllers, computers that are running earlier versions of Windows, and applications that depend on the NetBIOS namespace — for example, applications that call NetServerEnum and other so-called Net* application programming interfaces (APIs) that depend on NetBIOS. There is no interoperability between Windows Server 2008 based domains and Windows NT 4.0 domains.

Active Directory domains have both DNS names and NetBIOS names. In general, both names are visible to end users. The DNS

names of Active Directory domains include two parts, a prefix and a suffix. The DNS prefix is the first label in the DNS name of the domain. The suffix is the name of the Active Directory forest root domain. For example, the first label of the DNS name for the Child1.wingtiptoys.com domain is Child1 and is referred to as the DNS prefix. The name of the forest root domain in that same forest is wingtiptoys.com and is referred to as the DNS suffix.

**Active Directory Domain Names and the Internet**

Active Directory domain names can exist within the scope of the global Internet DNS namespace. When an Internet presence is required by an individual or organization, the Active Directory namespace is maintained as one or more hierarchical Windows 2000 domains beneath a root domain that is registered as a DNS namespace. Registration of individual and organizational root domain DNS names ensures the global uniqueness of all DNS names and provides for the assignment of network addresses that are recorded in the global DNS database. Registration of the DNS name for the root domain of the individual or organization also grants that individual or organization the authority to manage its own hierarchy of child domains, zones, and hosts within the root domain.

### Note

- An organization might or might not choose to be part of the global Internet DNS namespace. However, even if the root domain of the organization is not registered as an Internet DNS namespace, the DNS service is required to locate Windows 2000–, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7 and Windows Server 2008 R2–based computers in general and Windows 2000 Server,Windows Server 2003, Windows Server 2008 and Windows Server 2008 R2–based domain controllers in particular.

For more information about domain names, see "How DNS Support for Active Directory Works."

**Active Directory Objects**

Active Directory objects represent the entities that make up a network. An object is a distinct, named set of attributes that represents something concrete, such as a user, a printer, or an application. When you create an Active Directory object, Active Directory generates values for some of the attributes of the object; others you provide. For example, when you create a user object, Active Directory assigns a globally unique identifier (GUID) and a security identifier (SID), and you provide values for such attributes of the user as the given name, surname, logon identifier, and so on. This section describes the key identifiers assigned to objects by Active Directory and their associated naming schemes

**Object Uniqueness**

Each object in Active Directory is associated with at least one identifier that identifies that object as unique in an enterprise. This object identifier is referred to as a globally unique identifier (GUID). Another identifier, referred to as a security ID (SID), is used on security-enabled objects so that authentication and authorization services can determine its origin and validity within the domain or forest. Active Directory objects that are security-enabled are referred to as security principals.

A security principal is an object managed by Active Directory that is automatically assigned a SID for logon authentication and for access to resources. A security principal can be a user account, computer account, or a group, and is unique within a single domain. A security principal object must be authenticated by a domain controller in the domain in which the security principal object is located, and it can be granted or denied access to network resources.

**GUIDs**

Every object in Active Directory has a GUID, a 128-bit number assigned by the Directory System Agent when the object is created. The GUID, which cannot be altered or removed, is stored in an attribute that is required for every object, not just security principal objects. The GUID of each object is stored in its Object-GUID (objectGUID) property. When storing a reference to an Active Directory object in an external store (for example, a Microsoft SQL Server database), the objectGUID value should be used.

Active Directory uses GUIDs internally to identify objects. For example, the GUID is one of the properties of an object that is published in the global catalog. Searching the global catalog for the GUID of a User object will yield results if the user has an account somewhere in the enterprise. In fact, searching for any object by Object-GUID might be the most reliable way of finding the object you want to find. This is because the values of other object properties can change, but the Object-GUID never changes. When an object is assigned a GUID, it always keeps that value.

**Security IDs (SIDs)**

When a new security principal is created, Active Directory stores the SID of the security principal in the Object-SID (objectSID) property of the object and assigns the new object a GUID. Each security principal (as well as the domain itself) has a SID, which is the property that authoritatively identifies the object to the Windows security subsystem. The SID of a user, group, or computer is derived from the SID of the domain to which the object belongs; this SID is made up of the value of the domain SID and one additional 32-bit component called the relative identifier (RID).

In Windows 2000 and later operating systems, ACLs are used to identify users and groups by SID, not by GUID — even ACLs

on resources in Active Directory. A user gains access to, for example, a Group Policy object, based on one of the SIDs belonging to the user, not based on the GUID for the User object.

**SIDs and Migrations**

When an employee moves to a different location or to a new position, their user account might need to be moved or migrated to a different domain within the organization. Migrating a user account from one domain to another replaces the SID of the account with a new SID and new RID assigned by the new domain. For example, if Nicolette moves from North America to Europe, but stays in the same company, her account can be transferred with her. An administrator with the appropriate credentials can simply move her User object from, say, the Noam.wingtiptoys.com domain to the Euro.wingtiptoys.com domain. When the account is moved, the User object for her account needs a new SID. The domain identifier portion of a SID issued in Noam is unique to Noam, so the SID for her account in Euro has a different domain identifier. The RID portion of a SID is unique relative to the domain, so if the domain changes, the RID also changes.

Thus when a User object moves from one domain to another, a new SID must be generated for the user account and stored in the Object-SID property. Before the new value is written to the property, the previous value is copied to another property of a User object, SID History (SIDHistory). This property can hold multiple values. Each time a User object moves to another domain, a new SID is generated and stored in the Object-SID property and another value is added to the list of old SIDs in SIDHistory. When a user logs on and is successfully authenticated, the domain authentication service queries Active Directory for all of the SIDs associated with the user — the current SID of the user, any old SIDs of the user, and the SIDs for the groups to which the user belonged. All of these SIDs are returned to the authentication client and are included in the access token of the user. When the user tries to gain access to a resource, any one of the SIDs in the access token, including one of the SIDs in SIDHistory, can be used to authorize the user to the resource.

For more information about SIDHistory, see "How Security Identifiers Work" and "Security Considerations for Trusts."

**Object Names**

Object names are used to identify accounts in an Active Directory network. Objects have both LDAP-related names and logon names. Each object name represents a unique attribute for that object.

**LDAP-Related Names**

Active Directory is a Lightweight Directory Access Protocol (LDAP)-compliant directory service. In Windows Server 2003 and higher operating systems, all access to Active Directory objects occurs through LDAP. LDAP defines what operations can be performed in order to query and modify information in a directory, and how information in a directory can be accessed in compliance with established security requirements. Therefore, you can use LDAP to find or enumerate directory objects and to query or administer Active Directory.

It is possible to query by LDAP distinguished name (which is itself an attribute of the object), but because these names are difficult to remember, LDAP also supports querying by other attributes (for example, by using the attribute "color" to find "color" printers). This lets you find an object without having to know the distinguished name. If your organization has several domains, it is possible to use the same user name or computer name in different domains.

Like some other types of object names, LDAP-related names can change. The SID, globally unique ID, LDAP distinguished name, and canonical name generated by Active Directory uniquely identify each user, computer, or group in the forest. If the security principal object is renamed or moved to a different domain, the SID, LDAP relative distinguished name, LDAP distinguished name, and canonical name change, but the globally unique ID generated by Active Directory does not change.

**Distinguished Name**

Objects are located within Active Directory domains according to a hierarchical path, which includes the labels of the Active Directory domain name and each level of container objects. The full path to the object is defined by the distinguished name (also known as a DN). The name of the object itself, separate from the path to the object, is defined by the relative distinguished name.

The distinguished name is unambiguous (identifies one object only) and unique (no other object in the directory has this name). By using the full path to an object, including the object name and all parent objects to the root of the domain, the distinguished name uniquely and unambiguously identifies an object within a domain hierarchy. It contains sufficient information for an LDAP client to retrieve information the about the object from the directory.

For example, a user named Samantha Smith works in the marketing department of a company as a promotions coordinator. Therefore, her user account is created in an organizational unit that stores the accounts for marketing department employees who are engaged in promotional activities. The user identifier for Samantha Smith is Ssmith, and she works in the North American branch of the company. The root domain of the company is wingtiptoys.com, and the local domain is noam.wingtiptoys.com. The following distinguished name is of the user object Ssmith in the noam.wingtiptoys.com domain.

cn=Ssmith,ou=Promotions,ou=Marketing,dc=noam,dc=wingtiptoys,dc=com

### Note

- Active Directory tools do not display the LDAP abbreviations for the naming attributes domain component (dc=), organizational unit (ou=), common name (cn=), and so forth. These abbreviations are shown only to illustrate how LDAP recognizes the portions of the distinguished name. Most Active Directory tools display object names in canonical form, as described later in this section. Because distinguished names are difficult to remember, it is useful to have other means for retrieving objects. Active Directory supports querying by attribute (for example, the building number where you want to find a printer), so an object can be found without having to know the distinguished name.

**Relative Distinguished Name**

The relative distinguished name of an object is the part of the name that is an attribute of the object itself — the part of the object name that identifies this object as unique from its siblings at its current level in the naming hierarchy. Using the distinguished name mentioned earlier, the relative distinguished name of the object is SSmith. The relative distinguished name of the parent object is Promotions. The maximum length allowed for a relative distinguished name is 255 characters, but attributes have specific limits imposed by the directory schema. For example, in the case of the common name (cn), which is the attribute type often used for naming the relative distinguished name, the maximum number of characters allowed is 64.

Active Directory relative distinguished names are unique within a specific parent — that is, Active Directory does not permit two objects with the same relative distinguished name under the same parent container. However, two objects can have identical relative distinguished names but still be unique in the directory because within their respective parent containers, their distinguished names are not the same. (For example, the object cn=SSmith,dc=noam,dc=wingtiptoys,dc=com is recognized by LDAP as being different from cn=SSmith,dc=wingtiptoys,dc=com.)

The relative distinguished name for each object is stored in the Active Directory database. Each record contains a reference to the parent of the object. By following the references to the root, the entire distinguished name is constructed during an LDAP operation.

**Canonical Name**

By default, the user interface in Windows 2000 and higher operating systems, displays object names that use the canonical name, which lists the relative distinguished names from the root downward and without RFC 1779 naming attribute descriptors; it uses the DNS domain name (the form of the name where domain labels are separated by periods). For the LDAP distinguished name in the previous example, the respective canonical name would appear as follows: noam.wingtiptoys.com/marketing/promotions/ssmith

### Note

- If the name of an organizational unit contains a forward slash character (/), Active Directory requires an escape character in the form of a backslash (\) to distinguish between forward slashes that separate elements of the canonical name and the forward slash that is part of the organizational unit name. The canonical name that appears in Active Directory Users and Computers properties pages displays the escape character immediately preceding the forward slash in the name of the organizational unit. For example, if the name of an organizational unit is Promotions/Northeast and the name of the domain is Wingtiptoys.com, the canonical name is displayed as Wingtiptoys.com/Promotions\/Northeast.

**Logon Names**

A unique logon name is required for user security principals to gain access to a domain and its resources. Security principals are objects to which Windows security is applied in the form of authentication and authorization. Users are security principals, and they are authenticated (their identity is verified) at the time they log on to the domain or local computer. They are authorized (allowed or denied access) when they use resources.

In Windows 2000 and higher operating systems, user security principals require a unique logon name to gain access to a domain and its resources. Security principal objects might be renamed, moved, or contained within a nested domain hierarchy. The names of security principal objects must conform to the following guidelines:

- The name cannot be identical to any other user, computer, or group name in the domain. It can contain up to 20 uppercase or lowercase characters except for the following:" / \ [ ] : ; | = , + * ? < >

- A user name, computer name, or group name cannot consist solely of periods (.) or spaces.

The two types of logon names are user principal name and Security Account Manager account names.

**User Principal Name**

In Active Directory, each user account has a user principal name (UPN) in the format *user@DNS-domain-name*. A UPN is a friendly name assigned by an administrator that is shorter than the LDAP distinguished name used by the system and easier to remember. The UPN is independent of the DN of the user object, so a user object can be moved or renamed without affecting the user logon name. When logging on using a UPN, users do not have to choose a domain from a list on the logon dialog box.

The three parts of a UPN are the UPN prefix (user logon name), the @ character, and the UPN suffix (usually a domain name). The default UPN suffix for a user account is the DNS name of the Active Directory domain where the user account is located. For example, the UPN for user Frank Miller, who has a user account in the Wingtiptoys.com domain (if Wingtiptoys.com is the only domain in the tree), is FMiller@Wingtiptoys.com. The UPN is an attribute (userPrincipalName) of the security principal object. If the userPrincipalName attribute of a User object has no value, the User object has a default UPN of userName@DnsDomainName.

If your organization has many domains forming a deep domain tree organized by department and region, default UPN names can become unwieldy. For example, the default UPN for a user might be Sales.westcoast.microsoft.com. The logon name for a user in that domain is *user*@sales.westcoast.microsoft.com. Instead of accepting the default DNS domain name as the UPN suffix, you can simplify both administration and user logon processes by providing a single UPN suffix for all users. (The UPN suffix is used only within the Active Directory domain and is not required to be a valid DNS domain name.) You can choose to use your e-mail domain name as the UPN suffix — *userName*@microsoft.com. This gives the user in the example the UPN name of *user*@microsoft.com.

For a UPN–based logon, a global catalog might be necessary, depending on the user logging on and the domain membership of the computer where the user logs on. A global catalog is needed if the user logs on with a non-default UPN and the computer account of the user is in a different domain than the user account of the user. For example, if, instead of accepting the default DNS domain name as the UPN suffix (in the example *user*@sales.westcoast.microsoft.com), you provide a single UPN suffix for all users (so that the user then becomes simply *user*@microsoft.com), a global catalog is required for logon.

You use the Active Directory Domains and Trusts tool to manage UPN suffixes for a domain. UPNs are assigned at the time a user is created. If you have created additional suffixes for the domain, you can select from the list of available suffixes when you create the user or group account. The suffixes appear in the list in the following order:

- Alternate suffixes (if any; last one created appears first)
- Root domain
- The current domain

**SAM Account Name**

A Security Accounts Manager (SAM) account name is required for compatibility with Windows NT 3.*x* and Windows NT 4.0 domains. The user interface in Windows 2000 or higher operating systems, refers to the SAM account name as User logon name (pre-Windows 2000). There is no interoperability between Windows Server 2008 or higher domains and Winows NT 3 and Windows NT 4.0 domains.

SAM account names are sometimes referred to as flat names because — unlike DNS names — SAM account names do not use hierarchical naming. Because SAM names are flat, each one must be unique in the domain.

**Organizational Units**

Active Directory allows administrators to create a hierarchy within a domain that meets the needs of their organization. The object class of choice for building these hierarchies is the **organizationalUnit**, a general-purpose container that can be used to group most other object classes together for administrative purposes. An organizational unit in Active Directory is analogous to a directory in the file system; it is a container that can hold other objects. You can use organizational units for purposes such as creating an administrative hierarchy, applying Group Policy, and delegating control of administration.

**Administrative Hierarchy**

Organizational units can be nested to create a hierarchy within a domain and form logical administrative units for users, groups, and resource objects, such as printers, computers, applications, and file shares. The organizational unit hierarchy within a domain is independent of the structure of other domains; each domain can implement its own hierarchy. Likewise, domains that are managed by a central authority can implement similar organizational unit hierarchies. The structure is completely flexible, which allows organizations to create an environment that mirrors the administrative model, whether it is centralized or decentralized.

**Group Policy**

Group Policy can be applied to organizational units to define the abilities of groups of computers and users that are contained within the organizational units. Levels of control range from complete desktop lockdown to a relatively autonomous user

experience. Group Policy can affect functionality, such as what applications are available to a group of users, what features within an application are accessible on a particular computer, where documents are saved, and can affect access and user permissions. Group Policy also affects where, when, and how application and operating system updates or special scripts are applied.

Group Policy settings are stored as Group Policy objects in Active Directory. A Group Policy object can be associated with one or more Active Directory containers, such as a site, domain, or organizational unit.

**Delegation of Control**

The Active Directory object-based security model implements default access control that is propagated down a subtree of container objects. You can use this technology to determine the security for an entire group of objects according to the security that you set on the organizational unit that contains the objects. By default, most Active Directory objects inherit ACEs from the security descriptor on their parent container object. If necessary, you can change the inherited permissions. However, as a best practice, you should avoid changing the default permissions or inheritance settings on Active Directory objects unless you have additional security requirements.

### Note

- Because Active Directory is indexed, you can organize the tree to match your administrative model, instead of having to organize it for ease of browsing.

Inheritance enables the access control information defined at a container object in Active Directory to apply to the security descriptors of any subordinate objects, including other containers and their objects. One benefit this provides is that it eliminates the need to apply permissions each time a child object is created. You can apply or delegate administrative control over directory objects to organizational units by setting access control.

This inheritance of access effectively delegates administrative control to individuals in the organization. The best way to take full advantage of delegation and inherited control on directory objects is to organize the hierarchy to match the way that the directory is administered.

**User Accounts and Access Control**

Active Directory authenticates and authorizes security principals such as user, inetorgperson, and computer accounts to access shared resources on the network. The Local Security Authority (LSA) is the security subsystem responsible for all interactive user authentication and authorization services on a local computer. The LSA also processes authentication requests made through the Kerberos V5 protocol or the NTLM protocol in Active Directory.

Once the identity of a user is verified in Active Directory, the LSA on the authenticating domain controller creates a security access token for that user. An access token contains the name of the user, the groups to which that user belongs, a SID for the user, SIDs included in the SIDHistory property and all of the SIDs for the groups to which the user belongs.

The information in the access token is used to determine the level of access a user has to resources whenever the user attempts to access them. The SIDs in the access token are compared with the list of SIDs that make up the discretionary access control list (DACL) on the resource to ensure that the user has sufficient permission to access the resource. This is because the access control process identifies user accounts by SID rather than by name.

### Note

- When a domain controller provides an access token to a user, the access token only contains information about membership in domain local groups if the domain local groups are local to the domain where the domain controller is located.

**User Authorization**

In addition to securing network access through user authentication, Active Directory protects shared resources by facilitating user authorization. Once a user logon has been authenticated by Active Directory, the user rights assigned to the user through security groups and the permissions assigned on the shared resource determine if the user will be authorized to access that resource. This authorization process protects shared resources from unauthorized access and permits access to only authorized users or groups.

Administrators can use access control to manage user access to shared resources for security purposes. In Active Directory, access control is administered at the object level by setting different levels of access, or permissions, to objects, such as Full Control, Write, Read, Delete, or No Access. Access control in Active Directory defines how users can use Active Directory objects. By default, most permissions on objects in Active Directory are at the most secure setting.

The elements that define access control permissions on Active Directory objects include security descriptors and the concept of

object inheritance.

**Security Descriptors**
Access control permissions are assigned to securable objects and Active Directory objects to control how different users can use each object. A securable object, or shared resource, is an object that is intended to be used over a network by one or more users, and includes files, printers, folders, and services. All securable objects and Active Directory objects store access control permissions in security descriptors.

A security descriptor contains two access control lists (ACLs) used to assign and track security information for each object: these are the discretionary access control list (DACL) and the system access control list (SACL).

Discretionary access control lists (DACLs). DACLs identify the users and groups that are assigned or denied access permissions on an object. If a DACL does not explicitly allow access to a user, or to any group that a user is a member of, the user is implicitly denied access to that object. By default, a DACL is controlled by the owner of an object or the person who created the object (In Windows, the creator of an object is also the owner). The DACL contains access control entries (ACEs) that determine user access to the object.

System access control lists (SACLs). SACLs identify the users and groups that you want to audit when they successfully access or fail to access an object. Auditing is used to monitor events related to system or network security, to identify security breaches, and to determine the extent and location of any damage. By default, a SACL is controlled by the owner of an object or the person who created the object. A SACL contains ACEs that determine whether to record a successful or failed attempt by a user to access an object using a given permission, for example, Full Control or Read.

Active Directory allows you to apply access control permissions to objects at very low levels, including the ability to assign permissions on a per-attribute basis. To view DACLs and SACLs on Active Directory objects using Active Directory Users and Computers, on the View menu, click Advanced Features to access the Security tab for each object. You can also use the DSACLS support tool to manage access control lists in Active Directory.

By default, DACLs and SACLs are associated with every Active Directory object, which helps reduce attacks to your network by malicious users or any accidental mistakes made by domain users.

**Computer Accounts**
Each computer account created in Active Directory has a relative distinguished name, a pre-Windows 2000 computer name (Security Accounts Manager account name), a primary DNS suffix, a DNS host name, and a service principal name (SPN). The administrator enters the computer name when creating the computer account. This computer name is used as the LDAP relative distinguished name.

Active Directory suggests the pre-Windows 2000 name using the first 15 bytes of the relative distinguished name. The administrator can change the pre-Windows 2000 name at any time.

The DNS name for a host is called a full computer name and is a DNS fully qualified domain name (FQDN). The full computer name is a concatenation of the computer name (the first 15 bytes of the SAM account name of the computer account without the $ character) and the primary DNS suffix (the DNS domain name of the domain in which the computer account exists). It is listed on the **Computer Name** tab in System Properties in Control Panel.

By default, the primary DNS suffix portion of the FQDN for a computer must be the same as the name of the Active Directory domain where the computer is located. To allow different primary DNS suffixes, a domain administrator might create a restricted list of allowed suffixes by creating the msDS-AllowedDNSSuffixes attribute in the domain object container. This attribute is created and managed by the domain administrator using Active Directory Service Interfaces (ADSI) or LDAP.

The SPN is a multivalue attribute. It is usually built from the DNS name of the host. The SPN is used in the process of mutual authentication between the client and the server hosting a particular service. The client finds a computer account based on the SPN of the service to which it is trying to connect. The SPN can be modified by members of the Domain Admins group.

**Domain and Forest Protocols and Programming Interfaces**
The primary protocol that is used by domain controllers in every domain throughout the forest is LDAP, which runs on top of TCP/IP. LDAP is both a protocol and an API. In addition, the secured communications between domain controllers must use the remote procedure call (RPC) protocol for Messaging Application Programming Interface (MAPI), replication, domain controller management, and SAM-related operations.

**LDAP**
LDAP is a directory service protocol that specifies directory communications. It runs directly over TCP/IP, and it can also run over User Datagram Protocol (UDP) connectionless transports. Clients can use LDAP to query, create, update, and delete information that is stored in a directory service over a TCP connection through the TCP default port 389. Active Directory supports LDAP v2 (RFC 1777) and LDAP v3 (RFC 3377). LDAP v3 is an industry standard that can be used with any directory service that implements the LDAP protocol. LDAP is the preferred and most common way of interacting with Active Directory.

Historically, LDAP is a simplified (lightweight) version of Directory Access Protocol (DAP), which is the original protocol that was used to interact with X.500 directories. X.500 defines an earlier set of standards that was developed by the International Organization for Standardization (ISO). LDAP is simpler than DAP in two key ways:

- Rather than using its own protocol stack according to the Open Systems Interconnection (OSI) networking model, LDAP communicates over Internet Protocol (IP) by using either UDP or TCP.
- LDAP syntax is easier to use than DAP syntax.

For these reasons, LDAP is widely used and accepted as the standard protocol for directory service access. The following key aspects characterize LDAP:

- The protocol is carried directly over TCP for connection-oriented transport (receipt of data is acknowledged) and over UDP for connectionless transport (sent or received data is not acknowledged).
- Most protocol data elements can be encoded as ordinary strings, for example, as distinguished names.
- Referrals to other servers can be returned to the client.
- Simple Authentication and Security Layer (SASL) mechanisms can be used with LDAP to provide associated security services. SASL Digest is supported as the authentication standard for LDAP.
- Attribute values and distinguished names can be internationalized using the ISO 10646 character set.
- The protocol can be extended to support new operations, and controls can be used to extend existing operations.
- The schema is published through an attribute on the directory root object (rootDSE) for use by clients.

For more information about LDAP, see "How Active Directory Searches Work."

**RPC**
Active Directory uses remote procedure call (RPC) for replication (REPL) and domain controller management communications, MAPI communications, and SAM-related communications. RPC is a powerful, robust, efficient, and secure interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality located in a different process. That different process can be on the same computer, on the local area network (LAN), or across the Internet.

**Authentication Protocols**
Domain controllers authenticate users and applications by using one of two protocols: either the Kerberos version 5 authentication protocol or the NTLM authentication protocol. When two Active Directory domains or forests are connected by a trust, authentication requests made using these protocols can be routed to provide access to resources in both forests.

**NTLM**
The NTLM version 2 (NTLMv2) authentication protocol is a challenge/response authentication protocol. It is supported in Windows Server 2008, Windows Vista, Windows Server 2003, Windows 2000, and Windows XP, but it is not the default authentication protocol. Kerberos v5 is the default for these versions except when exchanging communications with a computer running Windows NT Server 4.0 or earlier. Networks with this configuration are referred to as mixed-mode. NTLM is also the authentication protocol for computers that are not participating in a domain, such as stand-alone servers and workgroups.

In Windows 7 and Windows Server 2008 R2, you can restrict the usage of NTLM. For more information see, Introducing the Restriction of NTLM Authentication (http://technet.microsoft.com/en-us/library/dd560653(WS.10).aspx)

When the NTLM protocol is used between a client and a server, the server must contact a domain authentication service on a domain controller to verify the client credentials. The server authenticates the client by forwarding the client credentials to a domain controller in the client account domain.

**Kerberos Version 5 Protocol**
The Kerberos version 5 protocol is the default authentication protocol used by computers running Windows 2000 or later operating system versions. This protocol is specified in RFC 1510 and is fully integrated with Active Directory, server message block (SMB), HTTP, and RPC, as well as the client and server applications that use these protocols. In Active Directory domains, the Kerberos protocol is used to authenticate logons when any of the following conditions is true:

- The user who is logging on uses a security account in an Active Directory domain.
- The computer that is being logged on to has an operating system that is Windows 2000 or later.
- The computer that is being logged on to is joined to an Active Directory domain.
- The computer account and the user account are in the same forest.
- The computer from which the user is trying to access resources is located in a non-Windows-based operating system Kerberos version 5 realm.

If any computer involved in a transaction does not support the Kerberos version 5 protocol, the NTLM protocol is used.

The authentication protocol of choice for Active Directory authentication requests is Kerberos V5. In contrast to NTLM, when the Kerberos protocol is used, the server does not have to contact the domain controller. Instead, the client gets a ticket for a server by requesting one from a domain controller in the server account domain; the server validates the ticket without consulting any other authority.

For more information about Kerberos, see "How the Kerberos Version 5 Authentication Protocol Works."

**Active Directory APIs**
You can use the following application programming interfaces (APIs) to access information in any LDAP directory including Active Directory:

- Active Directory Service Interfaces (ADSI)
- LDAP C API
- REPL
- MAPI
- SAM

**ADSI**
Active Directory Service Interfaces (ADSI) provides a simple, powerful, object-oriented interface to Active Directory. ADSI makes it easy for programmers and administrators to create directory programs by using high-level tools, such as Microsoft Visual Basic, without having to worry about the underlying differences between the different namespaces.

ADSI is supplied as a software development kit that enables you to build or buy programs that give you a single point of access to multiple directories in your network environment, whether those directories are based on LDAP or another protocol. ADSI is fully scriptable for ease of use by administrators.

ADSI also enables access to Active Directory by exposing objects stored in the directory as Component Object Model (COM) objects. A directory object is manipulated using the methods available on one or more COM interfaces. ADSI has a provider-based architecture that allows COM access to different types of directories for which a provider exists.

**LDAP C**
The LDAP C API, defined in Internet standard RFC 1823, is a set of low-level C-language APIs to the LDAP protocol. Microsoft supports LDAP C APIs on all Windows platforms.

Developers have the choice of writing Active Directory-enabled applications using LDAP C APIs or ADSI. LDAP C APIs are most often used to ease portability of directory-enabled applications to the Windows platform. ADSI is a more powerful language and is more appropriate for developers writing directory-enabled code on the Windows platform.

LDAP is the primary interface for the data store. Directory clients use LDAP v3 to connect to the DSA through the LDAP interface. The LDAP interface is part of Wldap32.dll. LDAP v3 is backward compatible with LDAP v2.

**REPL**
The REPL management interface provides functionality for finding data about domain controllers, converting the names of network objects between different formats, manipulating SPNs and DSAs, and managing replication of servers.

**MAPI**
Messaging clients gain access to the Microsoft Exchange Server directory service by using MAPI address book providers. For compatibility with existing messaging clients, Active Directory supports the MAPI-RPC address book provider, which provides access to Active Directory, for example, to find the telephone number of a user.

**SAM**
Security Accounts Manager (SAM) is a proprietary interface for connecting to the DSA on behalf of clients that use Windows NT 4.0 or earlier. These clients use Windows NT 4.0 networking APIs to connect to the DSA through SAM. Replication with Windows NT 4.0 backup domain controllers (BDCs) goes through the SAM interface as well.

**Domain and Forest Processes and Interactions**
Many network related operations depend on domains and forests in order to complete various tasks. This section describes some of the processes and interactions that commonly occur within the boundaries of domains or forests.

**Logging on to a Domain**
When a user with an account in an Active Directory domain logs on at the keyboard of a computer running a Windows 2000 or later operating system, the logon request is processed in three stages:

1. The user requests admission to the ticket-granting service for the domain.

   This is accomplished through an Authentication Service (AS) Exchange between the Kerberos security support provider (SSP) on the computer and the Key Distribution Center (KDC) in the domain in which the user account exists. The result is a ticket-granting ticket (TGT) that the user can present in future transactions with this KDC.

2. The user requests a ticket for the computer.

   This is accomplished through a Ticket-Granting Service (TGS) Exchange between the Kerberos SSP on the computer and the KDC for the domain in which the computer account exists. The result is a session ticket that the user can present when he or she requests access to system services on the computer.

3. The user requests admission to Local System services on the computer.

   This is accomplished when the Kerberos SSP on the computer presents a session ticket to the LSA on the computer.

If the account domain of the computer is different from the account domain of the user, an extra step is involved. Before the Kerberos SSP can request a session ticket for the computer, it must ask the KDC in the domain where the user account exists for a TGT that is good for admission to the KDC in the domain where the computer account exists. The SSP can then present the TGT to the KDC in the domain of the computer and get a session ticket for the computer.

The precise steps in the logon process depend on how the computer is configured. With standard configurations of Windows, interactive users log on with a password. In another optional configuration of Windows, users log on with a smart card. Although the basic process is the same for both configurations, there are some differences. For more information about the domain logon process, see "How Interactive Logon Works."

**Processing Authentications Across Domains and Forests**
When a request for authentication is referred to a domain, before the domain controller in that domain authenticates the user to access resources in the domain, it must determine whether a trust relationship exists with the domain from which the request comes, as well as the direction of the trust and whether the trust is transitive or nontransitive. The authentication process between trusted domains varies according to the authentication protocol in use. The Kerberos version 5 and NTLM protocols in Windows 2000 Server, Windows Server 2003, Windows Server 2008, and Window Server R2, process referrals for authentication to a domain differently, as do other authentication protocols, such as Digest and SChannel, that Windows 2000 Server, Windows Server 2003, Windows Server 2008, and Window Server R2 support.

In an Active Directory environment the Kerberos-based authentication process is most commonly used. To access a shared resource in another domain by using Kerberos authentication, a computer where the user logs on first requests a ticket from a domain controller in its account domain to the server in the trusting domain that hosts the requested resource. This ticket is then issued by an intermediary trusted by both the requesting computer and the server. The computer then presents this trusted ticket to the server in the trusting domain for authentication. This process, however, becomes more complex when a workstation in one forest attempts to access data on a resource computer in another forest.

In this case, the Kerberos authentication process contacts the domain controller for a service ticket to the SPN of the resource computer. Once the domain controller queries the global catalog and determines that the SPN is not in the same forest as the domain controller, the domain controller sends a referral for its parent domain back to the workstation. At that point, the

workstation queries the parent domain for the service ticket and continues to follow the referral chain until it reaches the domain where the resource is located. For more detailed information about how authentication requests are processed across domains and forests, see "How Domain and Forest Trusts Work."

**Joining a Computer to a Domain**

Joining a computer to a domain creates the computer account object for the client in an Active Directory location where all computer accounts are created by default during a join operation. The default location is set to the **Computers** container in Active Directory. A computer account differs from a user account in that it identifies the computer to the domain, while a user account identifies a user to a computer.

The act of joining a computer to a domain creates an account for the computer on the domain if it does not already exist. When you join a client computer running Windows 2000 or later, the following events occur:

- The domain name is validated.
- A domain controller in the domain is located through a call to the DsGetDcName API.
- A session is established with the domain controller under the security context of the passed-in credentials that are supplied in the **Network Identification** tab under **System Properties** in **Control Panel**.
- The computer account is enabled. If the flags so specify (NETSETUP_ACCT_CREATE), the APIs create the computer account on the domain controller.
- The local password for this account is created in the Local Security Authority (LSA).
- The local primary domain information LSA policy is set to refer to the new domain. This includes the domain name and the domain SID.

  **Note**

  - For clients running Windows 2000 or later, the LSA policy consists of the domain name, domain SID, DNS domain name, DNS forest name, and domain GUID.
- The DNS name assigned to the local computer is updated.
- The local group membership is changed to add members of the Domain Admins group to the Local Accounts Administrators group.
- The Net Logon trusted domain cache is initialized to the trusted domains domain list.
- For clients running Windows 2000 or later, the Windows Time Service is enabled and started.
- The Net Logon service is started.

To join a workstation or member server to a domain, you can use the Netdom tool. For example, to join a workstation to the Wingtiptoys.com domain in the engineering organizational unit, type the following command at the workstation:

**Netdom join /d:wingtiptoys.com /OU:OU=engineering,DC=wingtiptoys,DC=com**

When a computer joins a domain, the following changes occur on domain controllers running Windows 2000 Server, Windows Server 2003, Windows Server 2008 and Window Server 2008 R@:

- A computer object is created. The name of this object is generated by appending a dollar sign ($) to the name (uppercase letters) of the client.
- On Windows 2000 Server and later versions of Microsoft Server operating systems–based domain controllers only, the Net Logon service creates SPNs on the computer object.

**Netsetup.log**

When joining a computer to an Active Directory domain, the Networking Setup (NetSetup) utility installs all the necessary Microsoft supported networking components. The Netsetup.log file provides information about attempts to join domains and records any errors that might prevent the join operation from succeeding.

**Registering DNS Names and Locating Domain Controllers**
When a Windows 2000 Server and later versions of Microsoft Server operating systems –based member server is promoted to a domain controller by installing Active Directory, the Net Logon service registers the DNS resource records necessary for network hosts and services to locate the domain controller on the network. When network hosts and services attempt an operation (such as joining a domain) that requires a domain controller, the locator mechanism attempts to locate the domain controller through DNS.

DNS is used because every Active Directory–based domain controller dynamically registers SRV records in DNS. The SRV records enable servers to be located by service type (for example, LDAP) and protocol (for example, TCP). Because domain controllers are LDAP servers that communicate over TCP, SRV records can be used to find the DNS computer names of domain controllers. In addition to registering LDAP-specific SRV records, Net Logon also registers Kerberos v5 authentication protocol–specific SRV records to enable locating servers that run the Kerberos Key Distribution Center (KDC) service.

Domain controllers not only register their DNS domain names on a DNS server, but also register their NetBIOS names by using a transport-specific mechanism (for example, WINS). Thus, a DNS client locates a domain controller by querying DNS, and a NetBIOS client locates a domain controller by querying the appropriate transport-specific name service. The domain controller locator service consists of two main parts:

- Locator finds which domain controllers are registered with a DNS server.
- Locator submits a DNS query to the DNS server to locate a domain controller in the specified domain.

After this query is resolved, an LDAP User Datagram Protocol (UDP) lookup is sent to one or more of the domain controllers listed in the response to the DNS query to ensure their availability. Finally, the Net Logon service caches the discovered domain controller to aid in resolving future requests.

For more information about the DC Locator process, see "How DNS Support for Active Directory Works."

**Raising Domain and Forest Functional Levels**
When Active Directory is installed on a server, a set of basic Active Directory features is enabled by default. In addition to the basic Active Directory features on individual domain controllers, there are new domain- and forest-wide Active Directory features available when all domain controllers in a domain or forest are running Windows Server 2003. This also holds true for Windows Server 2008 and Windows Server 2008 R2.

To enable the new domain-wide features, all domain controllers in the domain must be running Windows Server 2003, and the domain functional level must be raised to Windows Server 2003. To enable new forest-wide features, all domain controllers in the forest must be running Windows Server 2003, and the forest functional level must be raised to Windows Server 2003. The same is true for Windows Server 2008 and Windows Server 2008 R2.

Before raising the forest functional level to Windows Server 2003, verify that all domains in the forest are set to the domain functional level of Windows 2000 native or Windows Server 2003. Note that domains that are set to the domain functional level of Windows 2000 native will automatically be raised to Windows Server 2003 at the same time the forest functional level is raised to Windows Server 2003. For more detailed information about Windows Server 2008 and Windows Server 2008 R2 functional levels, see the "How Active Directory Functional Levels Work."

**Replicating Directory Partitions**
Active Directory uses RPC over IP to transfer replication data between domain controllers. RPC over IP is used for both intersite and intrasite replication. To keep data secure while in transit, RPC over IP replication uses both authentication (using the Kerberos V5 authentication protocol) and data encryption.

When a direct or reliable IP connection is not available, replication between sites can be configured to use the Simple Mail Transfer Protocol (SMTP). However, SMTP replication functionality is limited, and requires an enterprise certification authority (CA). SMTP can only be used to replicate the configuration, schema and application directory partitions, and does not support the replication of domain directory partitions. For more detailed information about the replication process, see "How Active Directory Replication Topology Works."

**Network Ports Used by Domains and Forests**
The following tables list the network ports that are associated with domains and forests.

### Port Assignments for Raising Active Directory Functional Levels

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP SSL | N/A | 636 |

**Port Assignments for Data Store**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP SSL | N/A | 636 |
| RPC Endpoint Mapper | 135 | 135 |
| Global Catalog LDAP | N/A | 3268 |
| Global Catalog LDAP SSL | N/A | 3269 |

**Port Assignments for Service Publication and SPNs**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP SSL | N/A | 636 |
| RPC Endpoint Mapper | 135 | 135 |
| Global Catalog LDAP | N/A | 3268 |
| Global Catalog LDAP SSL | N/A | 3269 |
| Kerberos | 88 | 88 |

**Port Assignments for Raising Active Directory Searches**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP SSL | N/A | 636 |
| Global Catalog LDAP | N/A | 3268 |
| Global Catalog LDAP SSL | N/A | 3269 |

**Port Assignments for Global Catalogs**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | N/A | 3268 |
| LDAP | N/A | 3269 (global catalog Secure Sockets Layer [SSL]) |
| LDAP | 389 | 389 |
| LDAP | N/A | 686 (SSL) |
| RPC/REPL | 135 | 135 (endpoint mapper) |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |

| | | |
|---|---|---|
| SMB over IP | 445 | 445 |

**Port Assignments for Replication**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | N/A | 686 (SSL) |
| RPC/REPL | N/A | 135 (endpoint mapper) |
| LDAP | N/A | 3268 |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

**Port Assignments for Operations Masters**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | N/A | 686 (SSL) |
| RPC/REPL | N/A | 135 (endpoint mapper) |
| Netlogon | N/A | 137 |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

**Port Assignments for Interactive Logon**

| Service Name | UDP | TCP |
|---|---|---|
| Kerberos | 88 | 88 |
| Local Security Authority (LSA) RPC | Dynmaic RPC | Dynamic RPC |
| NTLM | Dynamic | Dynamic |

**Port Assignments for Kerberos V5 Protocol**

| Service Name | UDP | TCP |
|---|---|---|
| DNS | 53 | 53 |
| Kerberos | 88 | 88 |

**Port Assignment for DC Locator**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |

The following table shows the list of ports that might need to be opened before you establish trusts.

**Ports Required for Trusts**

| Task | Outbound Ports | Inbound Ports | From–To |
|---|---|---|---|
| Set up trusts on both sides from the internal forest | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution — portmapper (135 TCP)<br>Net Logon fixed port | N/A | Internal domain domain controllers–External domain domain controllers (all ports) |
| Trust validation from the internal forest domain controller to the external forest domain controller (outgoing trust only) | LDAP (389 UDP)<br><br>Microsoft SMB (445 TCP)<br><br>Endpoint resolution — portmapper (135 TCP)<br>Net Logon fixed port | N/A | Internal domain domain controllers–External domain domain controllers (all ports) |
| Use Object picker on the external forest to add objects that are in an internal forest to groups and DACLs | N/A | LDAP (389 UDP and TCP)<br><br>Windows NT Server 4.0 directory service fixed port<br><br>Net Logon fixed port<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution portmapper (135 TCP) | External server–Internal domain PDCs (Kerberos)<br><br>External domain domain controllers–Internal domain domain controllers (Net Logon) |
| Set up trust on the external forest from the external forest | N/A | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP) | External domain domain controllers–Internal domain domain controllers (all ports) |
| Use Kerberos authentication (internal forest client to external forest) | Kerberos (88 UDP) | N/A | Internal client–External domain domain controllers (all ports) |
| Use NTLM authentication (internal forest client to external forest) | N/A | Endpoint resolution – portmapper (135 TCP)<br>Net Logon fixed port | External domain domain controllers–Internal domain domain controllers (all ports) |

| | | | |
|---|---|---|---|
| Join a domain from a computer in the internal network to an external domain | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution — portmapper (135 TCP) Net Logon fixed port<br><br>Windows NT Server 4.0 directory service fixed port | N/A | Internal client–External domain domain controllers (all ports) |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Active Directory Schema Technical Reference
- Data Store Technical Reference
- DNS Support for Active Directory Technical Reference
- Active Directory Replication Model Technical Reference
- Logon and Authentication Technologies
- Domain Controller Roles
- Active Directory Search and Publication Technologies
- Active Directory Installation, Upgrade, and Migration Technologies
- Windows Support Tools

# Domains and Forests Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Tools for Managing Domains and Forests
- Domains and Forests Registry Entries
- Domains and Forests Group Policy Settings
- Domains and Forests WMI Classes
- Network Ports Used by Domains and Forests
- Related Information


Administrators can use a number of methods to configure and manage Active Directory domain and forest environments. This section contains information about the tools, registry entries, Group Policy settings, Windows Management Instrumentation (WMI) classes, and network ports that are associated with Active Directory domains and forests.

**Note**

- When using Windows Server 2003 Active Directory administrative tools to connect to a domain controller running Windows 2000 you must first make sure that the Windows 2000–based domain controller to which you are connecting has Service Pack 3 or later installed. This is because Windows Server 2003 administrative tools sign and encrypt all LDAP traffic by default. If business reasons do not permit the installation of Service Pack 3 or later on domain controllers running Windows 2000 it is possible to disable this default behavior.

**Tools for Managing Domains and Forests**

The following tools are associated with domains and forests.

**Adsiedit.exe: ADSI Edit**

**Category**

ADSI Edit is included when you install Windows Server 2003 Support Tools.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

ADSI Edit is a Microsoft Management Console (MMC) tool that uses Active Directory Service Interfaces (ADSI), which ultimately uses the LDAP protocol. This tool can be used to view and modify directory objects in the Active Directory database.

To find more information about ADSI Edit, see "Adsiedit Overview."

**Csvde.exe: Csvde**

**Category**
Csvde is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Csvde to import and export data from Active Directory by using files that store data in the comma-separated value (CSV) file format. Csvde also supports batch operations that are based on CSV.

To find more information about Csvde, see Command-Line References in the "Tools and Settings Collection."

**Dcdiag.exe: Domain Controller Diagnostic Tool**

**Category**
The Domain Controller Diagnostic Tool command-line tool (Dcdiag) is included when you install the Windows Server 2003 Support Tools.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use the Domain Controller Diagnostic Tool to verify external trusts. This tool cannot be used to verify trust relationships based on the Kerberos version 5 authentication protocol; to verify Kerberos V5-based trust relationships, the recommended method is to use the Netdom tool. Using the Domain Controller Diagnostic Tool you can scope your external trust verification by site or by domain controller, check for trust establishment, check secured channel setup, and check for ticket validity between each pair of domain controllers. By default, errors are flagged. In verbose mode, successes are printed as well.

You can use the Domain Controller Diagnostic Tool to verify that there are sufficient resources for the DNS infrastructure when deploying the Windows 2000 Server or Windows Server 2003 Active Directory directory service. This tool analyzes the state of domain controllers in a forest or enterprise and reports any problems to assist in troubleshooting. As an end-user reporting program, the Domain Controller Diagnostic Tool queries the directory service infrastructure and uses the results to identify abnormal behavior in the system. The Domain Controller Diagnostic Tool provides a framework for executing tests to verify different functional areas of the system. This framework selects which domain controllers are tested according to scope directives from the user, such as enterprise, site, or single server.

**Dcpol.msc: Domain Controller Security Policy**

**Category**
Domain Controller Security Policy is a snap-in for MMC and is automatically installed when you install Active Directory. You can also use Domain Controller Security Policy on computers not running Active Directory by installing the Administration Tools Pack (Adminpak.msi).

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can set security settings for a domain controller in Domain Controller Security Policy.

For more information about Domain Controller Security Policy, see Help and Support Center in Windows Server 2003.

**Dcpromo.exe: Active Directory Installation Wizard**

**Category**
An Active Directory wizard that is included with Windows Server 2003 and is available from the command line or from the Configure Your Server Wizard on any computer running Windows Server 2003.

**Version compatibility**
This tool is compatible with computers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

The Active Directory Installation Wizard provides a graphical user interface for setting up a domain controller by installing Active Directory and, optionally, DNS. The Active Directory Installation Wizard can also be used on a Windows NT 4.0 primary domain controller (PDC) when upgrading it to Windows Server 2003 and forming a new forest, to raise the forest functional level to Windows Server 2003 interim, if appropriate.

**Domain.msc: Active Directory Domains and Trusts**

**Category**
An Active Directory Administrative Tools MMC snap-in that is automatically installed on all domain controllers running Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

**Note**

- You cannot run the Windows Server 2003 Administration Tools Pack (Adminpak.msi) on a computer that is running Windows XP Professional, Windows XP Home Edition, or Windows XP 64-Bit Edition Version 2003 without Windows XP Service Pack 1 (SP1).

Active Directory Domains and Trusts provides a graphical interface in which you can view all domains in the forest. Using this tool, an administrator can manage each of the domains in the forest, trust relationships between domains, configure the functional level for each domain or forest, and configure the alternative user principal name (UPN) suffixes for a forest.

Active Directory Domains and Trusts can be used to accomplish most trust related tasks. It can be used to target all Active Directory domain controllers and can verify all Active Directory trust types. Trust verification takes place between two domains by enumerating all of the domain controllers in each domain. If you choose to have Active Directory Domains and Trusts create both sides of the trust at once, the trust password is automatically generated.

For more information about Active Directory Domains and Trusts, see Help in Active Directory Domains and Trusts.

**Dompol.msc: Domain Security Policy**

**Category**
Domain Security Policy is a snap-in for MMC and is automatically installed when you install Active Directory. You can also use Domain Controller Security Policy on computers not running Active Directory by installing the Administration Tools Pack (Adminpak.msi).

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Security settings for a domain are set in Domain Security Policy. Group Policy settings can be applied to lock-down which users are allowed to log on to the server as well as who can access the server from the network.

For more information about Domain Security Policy, see Help and Support Center in Windows Server 2003.

**Dsa.msc: Active Directory Users and Computers**

**Category**
An Active Directory Administrative Tools MMC snap-in that is automatically installed on all Windows Server 2003 domain controllers running Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Active Directory Users and Computers provides a graphical user interface that can be used to manage users and computers in Active Directory domains.

Additionally, LDAP Query can be used in this tool for the following:

- To identify domain controllers running Windows NT 4.0
- To connect to a domain

**Dsadd.exe: Dsadd**

**Category**
Dsadd is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsadd to add specific types of objects to the directory.

To find more information about Dsadd, see Command-Line References in the "Tools and Settings Collection."

**Dsget.exe: Dsget**

**Category**
Dsget is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsget to display the selected properties of a specific object in the directory.

- To find more information about Dsget, see Command-Line References in the "Tools and Settings Collection."

**Dsmod.exe: Dsmod**

**Category**
Dsmod is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsmod to modify an existing object of a specific type in the directory.

To find more information about Dsmod, see Command-Line References in the "Tools and Settings Collection."

**Dsmove.exe: Dsmove**

**Category**
Dsmove is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsmove to move a single object in a domain from its current location in the directory to a new location. You can also use Dsmove to rename a single object without moving it in the directory tree.

To find more information about Dsmove, see Command-Line References in the "Tools and Settings Collection."

**Dsquery.exe: Dsquery**

**Category**
Dsquery is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsquery to perform searches against Active Directory according to specified criteria. To find more information about Dsquery, see Command-Line References in the "Tools and Settings Collection."

**Dsrm.exe: Dsrm**

**Category**
Dsrm is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsrm to delete an object of a specific type, or any general object, from the directory.

To find more information about Dsrm, see Command-Line References in the "Tools and Settings Collection."

**Dssite.msc: Active Directory Sites and Services**

**Category**
An Active Directory Administrative Tools MMC snap-in that is automatically installed on all Windows Server 2003 domain controllers running Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Active Directory Sites and Services to create, modify, and delete site configuration objects in Active Directory, including sites, subnets, connection objects, and site links. You can also use Active Directory Sites and Services to create the intersite topology, including mapping subnet addresses to sites, creating and configuring site links, creating manual connection objects, forcing replication over a connection, setting a domain controller to be a global catalog server, and selecting preferred bridgehead servers.

For more information about Active Directory Sites and Services, see Help and Support Center in Windows Server 2003.

**Ldifde.exe: Ldifde**

**Category**
Ldifde is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Ldifde to create, modify, and delete directory objects on domain controllers. You can also use Ldifde to extend the schema, export Active Directory user and group information to other applications or services, and populate Active Directory with data from other directory services.

To find more information about Ldifde, see Command-Line References in the "Tools and Settings Collection."

**Ldp.exe: Ldp**

**Category**
Ldp is included when you install Windows Server 2003 Support Tools.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Ldp is a Lightweight Directory Access Protocol (LDAP) graphical user interface (GUI) tool that you can use to perform operations such as connect, bind, search, modify, add, and delete against any LDAP-compatible directory, such as Active Directory. You can also use Ldp to view objects that are stored in Active Directory, along with their metadata, such as security descriptors and replication metadata.

To find more information about Ldp, see "Windows Support Tools."

**Netdiag.exe: Network Connectivity Tester**

**Category**
The Network Connectivity Tester command-line tool (Netdiag) is included when you install Windows Server 2003 Support Tools.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

The Netdiag command-line tool examines .dll files, output from other tools, and the system registry to find potential problems. You can use Netdiag to troubleshoot connectivity over the secured channel that exists between a workstation and a domain controller.

For the various trust related tasks that can be performed using this tool, see the table Trust Tools Comparison by Task earlier in this section. To find more information about Netdiag, see "Windows Support Tools."

**Netdom.exe: Windows Domain Manager**

**Category**
The Windows Domain Manager command-line tool (Netdom) is included when you install Windows Server 2003 Support Tools.

**Version compatibility**
This tool is compatible with computers running Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Web Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Netdom is a command-line tool that allows you to create and manage Active Directory trust relationships (except forest trusts) and can help reduce the number of steps needed to create a trust by using Active Directory Domains and Trusts. You can also use the Netdom command line tool to complete batch management of trusts, join computers to domains, verify trusts (including forest trusts) and secured channels, and obtain information about the status of trusts.

Netdom can be targeted at all Active Directory domain controllers and can verify all Active Directory trust types. Verification is accomplished between two domains by enumerating the domain controllers in each domain. If you choose to have Netdom create both sides of the trust at once the trust password is automatically generated.

To find more information about Netdom, see "Windows Support Tools."

**Nltest.exe: NLTest**

**Category**
The NLTest command-line tool is included when you install Windows Server 2003 Support Tools.

**Version compatibility**
This tool is compatible with computers running Windows XP Professional; Windows Server 2003, Standard Edition; Windows Server 2003, Web Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

You can use the NLTest command-line tool to perform trust-related network administrative tasks such as testing the trust relationship between a Windows–based computer that is a member of a domain and the domain controller on which its

computer account is located. In domains where an external trust is defined, NLTest can be used to test the trust relationship between all domain controllers in the trusting domain and a domain controller in the trusted domain. Nltest can also be used to verify any secured channel.

To find more information about NLTest, see "Windows Support Tools."

**Ntdsutil.exe: Ntdsutil**

**Category**
Ntdsutil is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition |

Ntdsutil.exe provides management capabilities for Active Directory. You can use Ntdsutil.exe to perform Active Directory database maintenance, manage and control single-master operations, and remove replication metadata left behind by domain controllers that are removed from the network without uninstalling Active Directory. You can also use Ntdsutil to create application directory partitions and perform authoritative restore operations. This tool is intended for use by experienced administrators.

To find more information about Ntdsutil, see Command-Line References in the "Tools and Settings Collection."

**Repadmin: Repadmin**

**Category**
Repadmin is included when you install Windows Server 2003 Support Tools.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Administrators can use Repadmin to monitor and manage replication between domain controllers. You can determine the last successful replication of all directory partitions, identify inbound and outbound replication partners, identify the current bridgehead servers, view object metadata, and generally manage Active Directory replication topology. You can use Repadmin to force replication of an entire directory partition or of a single object. You can also list domain controllers in a site.

To find more information about Repadmin, at a command prompt type **repadmin /?** or see Command-Line References in the Tools and Settings Collection.

**Schmmgmt.msc: Active Directory Schema**

**Category**
An Active Directory Administrative Tools MMCsnap-in that is automatically installed on all domain controllers running Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Active Directory Schema is a graphical user interface that can be used to manage Active Directory objects and their associated attributes. The Active Directory Schema snap-in allows members of the Schema Admins group to manage the schema through a graphical interface. You can create and modify classes and attributes and also specify what attributes are indexed and what attributes are replicated to the Global Catalog. The tool should only be used in a test environment because it does not permit the user to set some important values on new schema objects.

Before the snap-in can be used, it must be registered so that it appears as an available snap-in for MMC.

**Setspn.exe: Setspn**

**Category**
Setspn is included when you install Windows Server 2003 Support Tools.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Administrators can use this command-line tool to read, modify, and delete values in the **servicePrincipalNames** attribute on an Active Directory service account object.

To find more information about Setspn, see "Windows Support Tools."

**Domains and Forests Registry Entries**

The following registry entries are associated with domains and forests.

- For data store related registry entries, see "Data Store Tools and Settings."
- For global catalog related registry entries, see "Global Catalog Tools and Settings."
- For replication related registry entries, see "Active Directory Replication Tools and Settings."
- For logon related registry entries, see "Interactive Logon Tools and Settings."
- For Kerberos related registry entries, see "Kerberos Authentication Tools and Settings."
- For Access Token related registry entries, see "Access Tokens Tools and Settings."

**Domains and Forests Group Policy Settings**

The following tables list and describe the Group Policy settings that are associated with domains and forests.

**Group Policy Settings Associated with Data Store**

| Group Policy Setting | Description |
|---|---|
| Audit Directory Services Access | When it is enabled, this Group Policy setting causes successful and failed directory access events to be logged in the Directory Service event log. |

**Group Policy Settings Associated with Active Directory Searches**

| Group Policy Setting | Description |
| --- | --- |
| Maximum size of Active Directory searches | Specifies the maximum number of objects that the system displays in response to a command to browse or search Active Directory.<br><br>This policy affects all browse displays that are associated with Active Directory, such as those in Local Users and Groups, Active Directory Users and Computers, and dialog boxes that are used to set permissions for user or group objects in Active Directory.<br><br>If you enable this policy, you can use it to limit the number of objects that are returned from an Active Directory search.<br><br>If you disable this policy or if you do not configure it, the system displays up to 10,000 objects. |
| Enable filter in Find dialog box | Displays the filter bar above the results of an Active Directory search. The filter bar consists of buttons for applying additional filters to search results.<br><br>If you enable this policy, the filter bar appears when the Active Directory **Find** dialog box opens, but users can hide it. |
| Hide Active Directory folder | Hides the Active Directory folder in My Network Places.<br><br>The Active Directory folder displays Active Directory objects in a browse window.<br><br>If you enable this policy, the Active Directory folder does not appear in the My Network Places folder.<br><br>If you disable this policy or if you do not configure it, the Active Directory folder appears in the My Network Places folder. |

**Group Policy Settings Associated with Global Catalogs**

| Group Policy Setting | Description |
| --- | --- |
| Automated Site Coverage by the DC Locator DNS SRV Records | Determines whether domain controllers dynamically register DC Locator site-specific SRV resource records for the closest sites where no domain controller for the same domain exists (or no global catalog server for the same forest exists). These DNS records are dynamically registered by the Net Logon service, and they are used to locate domain controllers. |
| Sites Covered by the GC Locator DNS SRV Records | Specifies the sites for which global catalog servers should register the site-specific GC Locator SRV resource records in DNS. These records are registered in addition to the site-specific SRV resource records registered for the site where the global catalog server is located and, if the global catalog server is appropriately configured, for the sites without a global catalog server in the same forest for which this global catalog server is the closest global catalog server. These records are registered by Net Logon service.<br><br>If this policy is not configured, it is not applied to any global catalog servers and global catalog servers use their local configuration. |

**Group Policy Settings Associated with Replication**

| Group Policy Setting | Description |
| --- | --- |

| | |
|---|---|
| Account Lockout Policy:<br>• Account lockout duration<br>• Account lockout threshold<br>• Reset account lockout counter after | Changes to these settings in the Domain Security Policy trigger urgent replication. |
| Password Policy:<br>• Enforce password history<br>• Maximum password age<br>• Minimum password age<br>• Minimum password length<br>• Password must meet complexity requirements<br>• Store passwords using reversible encryption | Changes to these settings in the Domain Security Policy trigger urgent replication. |
| Contact PDC on logon failure | Account lockout and domain password changes rely on contacting the primary domain controller (PDC) emulator urgently to update the PDC emulator with the change. If Contact PDC on logon failure is disabled, replication of password changes to the PDC emulator occurs non-urgently. |

**Group Policy Settings Associated with Interactive Logon**

| Group Policy Setting | Description |
|---|---|

| | |
|---|---|
| Password Policy:<br>• Enforce password history<br>• Maximum password age<br>• Minimum password age<br>• Minimum password length<br>• Password must meet complexity requirements<br>• Store password using reversible encryption for all users in the domain | Changes to the Password Policy settings control:<br>• The strength and complexity required of the password of every user |
| Audit Policy:<br>• Audit account logon events<br>• Audit account management<br>• Audit logon events | Changes to the Audit Policy settings control:<br>• Auditing of logons and log offs<br>• Auditing of password and permissions changes |
| User Rights Assignment:<br>• Access the computer from the network<br>• Deny logon as a batch job<br>• Deny logon as a service<br>• Deny logon locally<br>• Deny logon through terminal services<br>• Logon as a batch job<br>• Logon as a service<br>• Logon locally | Changes to the User Rights Assignment settings control:<br>• Which users are allowed or disallowed to log on to perform different tasks, including logging on as a batch job and a service<br>• Which users are allowed or disallowed to logon locally or through terminal services, as well as who can or cannot access the computer from the network |

| | |
|---|---|
| Security Options:<br>• Accounts: Limit local accounts use of blank passwords to console logon only<br>• Domain member: Maximum machine account password age<br>• Domain member: Require strong (in Windows 2000 or later) session key<br>• Interactive logon: Do not display last user name<br>• Interactive logon: Do not require CTRL+ALT+DEL<br>• Interactive logon: Message Text for users attempting to log on<br>• Interactive logon: Message title for users attempting to log on<br>• Interactive logon: Number of previous logons to cache (in case the domain controller is not available)<br>• Interactive logon: Require domain controller authentication to unlock workstation<br>• Interactive logon: Smart card removal behavior<br>• Recovery console: Allow automatic administrative logon<br>• Shutdown: Allow system to be shut down without having to log on | Changes to the Security Options settings control:<br>• Message text and title displayed by the GINA during an interactive logon<br>• Domain member settings<br>• Authentication settings, including allowing or disallowing blank passwords and password age |

**Group Policy Settings Associated with Access Tokens**

| **Group Policy Setting** | **Description** |
|---|---|
| User Rights Assignment:<br>• Create a token object<br>• Replace a process level token | Changes to these settings control:<br>• Calling APIs to create tokens.<br>• Whether a process can replace a token. |

| | |
|---|---|
| Audit Policy:<br>• Audit policy change<br>• Audit privilege use<br>• Audit process tracking | Changes to this setting will:<br>• Generate audits when rights are assigned with one of the tools discussed earlier.<br>• Enable audit privilege use. Will log when SeAssignPrimaryTokenPrivilege was used.<br>• Create an audit for assigning a primary token that contains the two processes involved and the identity of the token assigned. |
| Security Options:<br>• Network access: Let Everyone permissions apply to anonymous users | Changes to this setting affect whether Everyone is in the token for anonymous users. |

**Group Policy User Rights Assignment Settings Associated with Kerberos**

| Group Policy Setting | Description |
|---|---|
| Impersonate a client after authentication | Windows 2000 security setting that was first introduced in Windows 2000 SP4. When you assign this user right to a user, you permit programs that run on behalf of that user to impersonate a client. This security setting helps to prevent unauthorized servers from impersonating clients that connect to it through methods such as remote procedure calls (RPC) or named pipes.<br>By default, members of the Administrators group and the System account are assigned this user right. The following components also are assigned this user right:<br>• Services that are started by the Service Control Manager<br>• Component Object Model (COM) servers that are started by the COM infrastructure and that are configured to run under a specific account |

**Group Policy Kerberos Policy Settings Associated with Kerberos**

| Group Policy Setting | Description |
|---|---|
| Enforce user logon restrictions | Determines whether the KDC validates every request for a session ticket against the user rights policy on the target computer. When this policy is enabled, the user requesting the session ticket must have the right to either Log on locally or to Access this computer from network. Validation of each request is optional because the extra step takes time and might slow network access to services. By default, this policy is enabled. |
| Maximum lifetime for service ticket | Determines the maximum amount of time (in minutes) that a ticket granted for a service (that is, a session ticket) can be used to access the service. If the setting is zero minutes, the ticket never expires. Otherwise, the setting must be greater than ten minutes and less than the setting for Maximum lifetime for user ticket. By default, the setting is 600 minutes (10 hours). |
| Maximum lifetime for user ticket | Determines the maximum amount of time (in hours) that a ticket-granting ticket (TGT) for a user can be used. When a TGT expires, a new one must be requested or the existing one must be renewed. By default, the setting is ten hours. |
| Maximum lifetime for user ticket renewal | Determines the longest period of time (in days) that a TGT can be used if it is repeatedly renewed. By default, the setting is seven days. |

| | |
|---|---|
| Maximum tolerance for computer clock synchronization | Determines the maximum difference (in minutes) that the Kerberos V5 protocol tolerates between the clock time on a client and the clock time on a server while still considering the two clocks synchronous. By default, the setting is five minutes. |

To find more information about these Group Policy settings, see Group Policy Settings Reference in the "Tools and Settings Collection" or see "Account Policy Settings."

**Domains and Forests WMI Classes**

Windows Management Instrumentation (WMI) provides access to information about certain objects in a Windows 2000 Server or Windows Server 2003 operating system. WMI providers and classes represent the managed resources on a computer and are used by administrators and developers for scripting and monitoring purposes. The following tables list and describe the WMI classes that are associated with Active Directory domains and forests.

**WMI Classes Associated with Data Store, Service Principal Names (SPNs) and Active Directory Searches**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| **rootDSE** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2003<br>• Windows 2000 Server |
| **DS_LDAP_Class_Containment** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2003<br>• Windows 2000 Server |
| **DS_LDAP_Instance_Containment** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2003<br>• Windows 2000 Server |

**WMI Classes Associated with Active Directory Replication**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| **MSAD_DomainController** | \\root\MicrosoftActiveDirectory | Domain controllers running:<br>• Windows Server 2003<br>• Windows 2000 Server |
| **MSAD_NamingContext** | \\root\MicrosoftActiveDirectory | Domain controllers running:<br>• Windows Server 2003<br>• Windows 2000 Server |

| | | |
|---|---|---|
| **MSAD_ReplNeighbor** | \\root\MicrosoftActiveDirectory | Domain controllers running:<br>• Windows Server 2003<br>• Windows 2000 Server |
| **MSAD_ReplCursor** | \\root\MicrosoftActiveDirectory | Domain controllers running:<br>• Windows Server 2003<br>• Windows 2000 Server |
| **MSAD_ReplPendingOp** | \\root\MicrosoftActiveDirectory | Domain controllers running:<br>• Windows Server 2003<br>• Windows 2000 Server |

**WMI Classes Associated with Trusts**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| **Microsoft_TrustProvider** | root\microsoftactivedirectory | Domain controllers running Windows Server 2003 |
| **Microsoft_DomainTrustStatus** | root\microsoftactivedirectory | Domain controllers running Windows Server 2003 |
| **Microsoft_LocalDomainInfo** | root\microsoftactivedirectory | • Domain controllers running Windows Server 2003 |

**WMI Classes Associated with Interactive Logon**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| **Win32_LogonSession** | \root\cimv2 | Computers running:<br>• Windows Server 2003<br>• Windows XP |
| **Win32_LogonSessionMappedDisk** | \root\cimv2 | Computers running:<br>• Windows Server 2003<br>• Windows XP |
| **Win32_NetworkLoginProfile** | \root\cimv2 | Computers running:<br>• Windows Server 2003<br>• Windows XP |

**WMI Classes Associated with Access Tokens**

| Class Name | Namespace | Version Compatibility |
|---|---|---|

| | | |
|---|---|---|
| **Win32_TokenGroups** | \root\cimv2 | Computers running:<br>• Windows Server 2003<br>• Windows XP |
| **Win32_TokenPrivileges** | \root\cimv2 | Computers running:<br>• Windows Server 2003<br>• Windows XP |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by Domains and Forests**

The following tables list the network ports associated with domains and forests.

**Port Assignments for Raising Active Directory Functional Levels**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP SSL | N/A | 636 |

**Port Assignments for Data Store**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP SSL | N/A | 636 |
| RPC Endpoint Mapper | 135 | 135 |
| Global Catalog LDAP | N/A | 3268 |
| Global Catalog LDAP SSL | N/A | 3269 |

**Port Assignments for Service Publication and SPNs**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP SSL | N/A | 636 |
| RPC Endpoint Mapper | 135 | 135 |
| Global Catalog LDAP | N/A | 3268 |
| Global Catalog LDAP SSL | N/A | 3269 |
| Kerberos | 88 | 88 |

**Port Assignments for Raising Active Directory Searches**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP SSL | N/A | 636 |

| | | |
|---|---|---|
| Global Catalog LDAP | N/A | 3268 |
| Global Catalog LDAP SSL | N/A | 3269 |

**Port Assignments for Global Catalogs**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | N/A | 3268 |
| LDAP | N/A | 3269 (global catalog Secure Sockets Layer [SSL]) |
| LDAP | 389 | 389 |
| LDAP | N/A | 686 (SSL) |
| RPC/REPL | 135 | 135 (endpoint mapper) |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

**Port Assignments for Replication**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | N/A | 686 (SSL) |
| RPC/REPL | N/A | 135 (endpoint mapper) |
| LDAP | N/A | 3268 |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

**Port Assignments for Operations Masters**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | N/A | 686 (SSL) |
| RPC/REPL | N/A | 135 (endpoint mapper) |
| Netlogon | N/A | 137 |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

**Port Assignments for Interactive Logon**

| Service Name | UDP | TCP |
|---|---|---|

| | | |
|---|---|---|
| Kerberos | 88 | 88 |
| Local Security Authority (LSA) RPC | Dynamic RPC | Dynamic RPC |
| NTLM | Dynamic | Dynamic |

**Port Assignments for Kerberos V5 Protocol**

| Service Name | UDP | TCP |
|---|---|---|
| DNS | 53 | 53 |
| Kerberos | 88 | 88 |

**Port Assignment for DC Locator**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |

The following table shows the list of ports that might need to be opened before you establish trusts.

**Ports Required for Trusts**

| Task | Outbound Ports | Inbound Ports | From–To |
|---|---|---|---|
| Set up trusts on both sides from the internal forest | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution — portmapper (135 TCP) Net Logon fixed port | N/A | Internal domain domain controllers–External domain domain controllers (all ports) |
| Trust validation from the internal forest domain controller to the external forest domain controller (outgoing trust only) | LDAP (389 UDP)<br><br>Microsoft SMB (445 TCP)<br><br>Endpoint resolution — portmapper (135 TCP) Net Logon fixed port | N/A | Internal domain domain controllers–External domain domain controllers (all ports) |
| Use Object picker on the external forest to add objects that are in an internal forest to groups and DACLs | N/A | LDAP (389 UDP and TCP)<br><br>Windows NT Server 4.0 directory service fixed port<br><br>Net Logon fixed port<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution portmapper (135 TCP) | External server–Internal domain PDCs (Kerberos)<br><br>External domain domain controllers–Internal domain domain controllers (Net Logon) |

| | | | |
|---|---|---|---|
| Set up trust on the external forest from the external forest | N/A | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP) | External domain domain controllers–Internal domain domain controllers (all ports) |
| Use Kerberos authentication (internal forest client to external forest) | Kerberos (88 UDP) | N/A | Internal client–External domain domain controllers (all ports) |
| Use NTLM authentication (internal forest client to external forest) | N/A | Endpoint resolution – portmapper (135 TCP) Net Logon fixed port | External domain domain controllers–Internal domain domain controllers (all ports) |
| Join a domain from a computer in the internal network to an external domain | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution — portmapper (135 TCP) Net Logon fixed port<br><br>Windows NT Server 4.0 directory service fixed port | N/A | Internal client–External domain domain controllers (all ports) |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Windows Support Tools
- Command-Line References in the Tools and Settings Collection for information about DSQuery and Ntdsutil.
- Microsoft Platform SDK on MSDN for more information about many WMI classes that are associated with the DNS Server service.
- Group Policy Settings Reference in the Tools and Settings Collection for information about Group Policy settings that are associated with the DNS Client service.
- Registry Reference in the Tools and Settings Collection for information about registry entries that are associated with DNS.

# Active Directory Schema Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

The schema is the component of the Active Directory directory service that defines all the objects and attributes that the directory service uses to store data. You can combine some objects in the schema to create more-complex definitions if objects of greater complexity are required. You can also add new definitions to the schema to support new types of objects in the directory.

**In this subject**


- What Is the Active Directory Schema?
- How the Active Directory Schema Works
- Active Directory Schema Tools and Settings

# What Is the Active Directory Schema?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**What Is the Active Directory Schema?**

**In this section**


- Using Objects to Store Data
- Building the Schema
- Related Information


The schema is the Active Directory component that defines all the objects and attributes that the directory service uses to store data.

Active Directory stores and retrieves information from a wide variety of applications and services. So that it can store and replicate data from a potentially infinite variety of sources, Active Directory standardizes how data is stored in the directory. By standardizing how data is stored, the directory service can retrieve, update, and replicate data while ensuring that the integrity of the data is maintained.

The directory service uses objects as units of storage. All objects are defined in the schema. Each time that the directory handles data, the directory queries the schema for an appropriate object definition. Based on the object definition in the schema, the directory creates the object and stores the data.

Object definitions control the types of data that the objects can store, as well as the syntax of the data. Using this information, the schema ensures that all objects conform to their standard definitions. As a result, Active Directory can store, retrieve, and validate the data that it manages, regardless of the application that is the original source of the data. Only data that has an existing object definition in the schema can be stored in the directory. If a new type of data needs to be stored, a new object definition for the data must first be created in the schema.

**Using Objects to Store Data**

Active Directory uses objects to store information. Objects are data structures that consist of multiple attributes that store both data and its related metadata. Metadata is data that describes the properties of other data. For example, an object that stores a user account has many attributes, including attributes that contain the user's logon name, first name, last name, and password. Each of those attributes has additional attributes that contain metadata about the information that the attribute stores. The logon name attribute, for example, has multiple attributes of its own. One attribute that is associated with the logon name specifies that the logon name is a required attribute, which means that the user object is not valid unless it contains the logon name attribute. Another attribute that is associated with the logon name specifies the syntax of the value that is stored in the logon name attribute. This ensures that the value that the logon name attribute contains is in a valid format. Both of these attributes contain metadata for the logon name attribute; that is, they define the characteristics of the logon name attribute.

The object definitions in the schema list all the object attributes and define how these attributes relate to each other. Some objects are simple and contain only a few attributes, while other objects are quite complex and contain hundreds of attributes. Attributes themselves are objects, and the schema contains a definition for each one. To define new objects, smaller objects are associated with one another to define the necessary attributes of the new objects.

For a user object, the user's logon name attribute is a smaller object that contains a number of attributes of its own. Among them are attributes that define the syntax of the logon name and specify whether or not the logon name attribute is optional or required. The first name and last name attributes are also smaller objects whose definitions can be found in the schema. The object definition that defines the user object lists the logon name attribute object, the first name and last name attribute objects, and many other attribute objects, and it defines how these objects relate to each other to store the data that represents a user account.

Defining objects and attributes in this manner gives the schema the ability to efficiently define many different types of objects while retaining the ability to add new types of objects when necessary. Many objects have some attributes in common. For example, many objects have a security descriptor to define who is allowed to access and change their contents. Rather than create a separate security descriptor definition for each object definition, the schema defines a single security descriptor object, and all other object definitions refer to the single security descriptor definition. This makes it possible for every object that

needs a security descriptor to have one security descriptor while keeping only one definition for the security descriptor in the schema.

**Building the Schema**

The Active Directory installation process that creates the forest also generates the default schema. Thereafter, the default schema replicates to each new domain controller during the installation of the directory on that new domain controller. The default schema contains all the standard object definitions that are necessary for Active Directory to function in a standard deployment.

Active Directory uses a multimaster replication topology, which means that any domain controller in a forest can write a change to the directory database and then replicate that change to other domain controllers in the same forest. For a domain controller to create a new object and write it to the directory, the domain controller must have access to the object definition that is needed to create the new object. Every domain controller in a forest maintains a copy of the schema, which makes it possible for domain controllers to have access to the object definitions that they need to store and retrieve information in the directory.

In some situations, the default attributes and object definitions in the schema are insufficient to create new object types that are required by some applications or services that interoperate with the directory. In these situations, it is possible to customize the schema by adding new object definitions to it. The process of adding definitions to the schema is referred to as "extending the schema."

It is important to plan the deployment of schema extensions carefully. The directory stores the schema and replicates schema changes to every domain controller throughout the forest. Therefore, extending the schema creates replication traffic, which can briefly affect network traffic. For more information about extending the schema, see "How the Active Directory Schema Works."

**Related Information**

- How the Active Directory Schema Works

# How the Active Directory Schema Works

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**How the Active Directory Schema Works**

**In this section**


- Active Directory Schema Architecture
- Active Directory Schema Physical Structure
- Active Directory Schema Processes and Interactions
- Related Information


The schema is the Active Directory component that defines all the objects and attributes that the directory service uses to store data. The physical structure of the schema consists of the object definitions. The schema itself is stored in the directory.

The schema is stored in its own partition (the schema partition) in the directory. The schema is replicated among all the domain controllers in the forest, and any change that is made to the schema is replicated to every domain controller in the forest. Because the schema dictates how information is stored, and because any changes that are made to the schema affect every domain controller, changes to the schema should be made only when necessary — through a tightly controlled process — after testing has been performed to ensure that there will be no adverse effects on the rest of the forest.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

**Active Directory Schema Architecture**

In Active Directory, the schema defines the following:

- Objects that are used to store data in the directory
- The rules that govern the structure of those objects
- The structure and the content of the directory itself

These definitions consist of objects, attributes, and classes, which are described in the following section.

**Schema Components**

Objects, attributes, and classes are the basic components that are used to build object definitions in the schema.

**Objects**

Objects are structures that store both data that the objects represent and data that controls the content and structure of the objects. For example, a user account object contains a user's logon name as well as data that indicates the proper syntax for storing the user's logon name in the user object.

Active Directory uses objects to store data while the data is maintained in the directory. When the directory stores an object, some associated data is also stored along with the object. The extra data is stored in the attributes of the object.

**Attributes**

Attributes contain data that defines the information that is stored in an object or in another attribute. For example, a user account object has attributes that store user information, such as the user's first name, last name, password, office number, and telephone number. Different types of objects have different attributes. A printer object has attributes that store information

that is related to printers, such as the printer model, the number of paper trays, the current location of the printer, and the port to which the printer is attached.

Some attributes contain information that relates to other attributes, such as syntax information or flags that label the attribute as optional or required. Syntax attributes define the format that is used to store data in other attributes. For example, because a telephone number comprises only digits, a syntax attribute might specify that the phone number attribute can only store digits 0 through 9, and letters of the alphabet are prohibited. Active Directory uses syntax attributes to ensure that information is stored in a legitimate format and that the information is a valid data type.

Attributes can be required or optional. A user account object requires a user name attribute and a password attribute, although the office number attribute is optional. That is, a user account cannot be used without a user name; however, it can be used without an office number. The office number attribute is included for convenience, and it is considered an optional attribute.

An object definition is really an association of various attributes that are used to describe the characteristics of an object that stores specific pieces of data. The kind of data that the object stores determines which attributes are needed to define the object.

Large objects are made up of many smaller objects. In the user account object example, the user's logon name attribute is a smaller object that contains a number of attributes of its own, including attributes that define the syntax of the logon name and that specify whether or not the logon name attribute is optional or required. The first name and last name attributes are also smaller objects that are defined in the schema. The object definition for the user account object lists the logon name attribute and the first name and last name attributes, as well as many other attributes, and it defines how these attributes relate to each other to store the data that represents a user account.

Defining objects and attributes this way gives the schema the ability to efficiently define many different types of objects. Many objects have some attributes in common. For example, many objects have a security descriptor to define who is allowed to access and make changes to the contents of the object. Rather than create a separate security descriptor definition for each object definition, the schema defines a single security descriptor object, and all other object definitions refer to the single security descriptor definition. This makes it possible for every object that needs a security descriptor to have one, while keeping only one definition for the security descriptor in the schema.

**Classes**

Object definitions are categorized into groups that are called classes. Classes act as blueprints that can be used each time a new object is created. When a new object is created in the directory, the object's class determines the attributes that are associated with the new object, including which attributes are required and which attributes are optional.

Active Directory has predefined classes that define all of the different object type's that the directory needs to function properly. For example, when a new user account object is created in the directory, its definition comes from the **classUser** class. The class dictates that the new account object is required to have a user name attribute and a password attribute, and optionally it might have an office number attribute.

Object definitions are created by nesting classes inside one another. Nesting classes produces a parent/child relationship between the classes. Classes that are nested inside another class are referred to as subclasses. The parent of a subclass is referred to as a superclass. When one class is nested inside another, the nested class inherits the properties of the parent superclass. When various classes that contain particular attributes are nested inside another object class, a new object definition is created. Which classes are nested depends on which attributes are needed to define the new object type.

The schema stores class information, but it does not store the actual objects that are derived from a class. For example, when a new user account object is created, it is not stored in the schema. Active Directory looks up the user *class* in the schema. Active Directory then retrieves information regarding the object type and its associated attributes from the user class in the schema and uses that information to create the new user account object. When the new account is created, the data is stored in the new object, and Active Directory then writes the new account information into the directory database.

The schema is the master list of all classes and attributes that can be used in the directory. During the installation of Active Directory, the Schema.ini file is used to build the default schema on the first domain controller in the forest. The default schema provides all the necessary classes that Active Directory needs to function. Administrators or developers might want to add their own classes or add their own attributes to an existing object type. They can do this through the process of extending the schema. The schema should only be extended in special situations. For more information about extending the schema, see "Modifying the Schema" later in this section.

**Schema Objects**

Active Directory uses objects to store and reference data in the directory. The schema defines the types of objects that are available to the directory service. The schema is stored in the schema partition, which is also defined as an object in the directory. The attributes and classes in Active Directory are stored in the schema partition as directory objects that are called schema objects. The schema partition itself is represented in Active Directory by an object that is an instance of the Directory

Management Domain (**dMD**) class.

Storing the schema in the directory has many advantages. For example, when user applications locate the schema in the directory, they can read the schema to discover what types of objects and properties are available. Because the schema is stored in the directory — like the rest of the data in the directory — applications can locate the schema by using the same process that they use to locate any other object in the schema.

A schema object, called a classSchema object, defines each class in the schema. Another schema object, called an attributeSchema object, defines each attribute in the schema. Therefore, every class is actually an instance of the **classSchema** class, and every attribute is an instance of the **attributeSchema** class, as shown in the following figure.

**Schema Objects**




Schema objects are used to define classes and attributes in the schema. Classes in the schema are used to define objects in the directory.

**attributeSchema Objects**

Attributes describe the classes that are defined in the schema. Attributes are defined in the schema separately from classes, which enables a single attribute definition to be applied to many classes.

Attributes are attributeSchema objects. Each attributeSchema object is an instance of the **attributeSchema** class. Like any other object, the attributeSchema object has its own attributes. The attributes for the attributeSchema object store, among other things, the following information:

- The Lightweight Directory Access Protocol (LDAP) display name of the attribute
- The object identifier for the attribute
- The globally unique identifier (GUID) for the attribute
- The syntax of the attribute
- The range for the attribute. For integers, range defines the minimum and maximum value; for strings, range defines the minimum and maximum length.
- Whether the attribute is a multivalue attribute. Note that multivalue attributes hold a set of values with no particular order. Multivalue attributes might not be returned in the order in which they were stored (or in any other order).
- Whether and how the attribute is indexed

**Attributes for the attributeSchema class**

Attributes for the **attributeSchema** class are described in the following table.

**Attributes for the attributeSchema Class**

| Attribute | Syntax | Mandatory? | Multivalue? | Description |
|---|---|---|---|---|
| **cn** | Unicode | Yes | No | Descriptive relative distinguished name for the schema object |

| | | | | |
|---|---|---|---|---|
| **attributeID** | Object identifier | Yes | No | Object identifier that uniquely identifies this attribute |
| **lDAPDisplayName** | Unicode | Yes, but automatic | No | Name by which LDAP clients identify this attribute |
| **schemaIDGUID** | String(Octet) | Yes | No | GUID that uniquely identifies this attribute |
| **mAPIID** | Integer | No | No | Integer by which Messaging API (MAPI) clients identify this attribute |
| **attributeSecurityGUID** | GUID | No | No | GUID by which the security system identifies the property set of this attribute |
| **attributeSyntax** | Object identifier | Yes | No | Syntax object identifier of this attribute |
| **oMSyntax** | Integer | Yes | No | Syntax of this attribute as defined by the XAPIA X/Open Object Model (XOM) specification |
| **isSingleValue** | BOOL | Yes | No | Indicates whether this attribute is a single-value or multivalue attribute<br><br>Note that multivalue attributes hold a set of values with no particular order. Multivalue attributes might not be returned in the order in which they were stored (or in any other order). |
| **extendedCharsAllowed** | BOOL | No | No | Indicates whether extended characters are allowed in the value of this attribute<br><br>Applies only to attributes of syntax String(teletex). |
| **rangeLower** | Integer | No | No | Lower range of values that are allowed for this attribute |
| **rangeUpper** | Integer | No | No | Upper range of values that are allowed for this attribute |
| **systemFlags** | Integer | No | No | Flags that determine specific system operations<br><br>**Note:** This attribute cannot be set or modified.<br><br>The following **systemFlags** attributes are relevant to the schema objects:<br><br>Attribute is required to be a member of the partial set = 0x00000002<br><br>Attribute is not replicated = 0x00000001<br><br>Attribute is a constructed attribute = 0x00000004 |

| | | | | |
|---|---|---|---|---|
| **searchFlags** | Integer | No | No | The **searchFlags** property of each property's attributeSchema object defines whether a property is indexed and other behavior.<br><br>The seven currently defined bits for this attribute are:<br><br>1 = Index over attribute only<br><br>2 = Index over container and attribute<br><br>4 = Add this attribute to the ambiguous name resolution (ANR) set (should be used in conjunction with 1)<br><br>8 = Preserve this attribute on logical deletion (that is, make this attribute available on tombstones)<br><br>16 = Include this attribute when copying a user object<br><br>32 = Create a Tuple index for the attribute to improve medial searches<br><br>64 = Reserved for future use; value should be 0.<br><br>128 = Available in Windows Server 2003 Service Pack 1 (SP1) only. Mark the attribute confidential (CONTROL_ACCESS is required to read it). |
| **isMemberof PartialAttributeSet** | BOOL | No | No | A Boolean value that defines whether the attribute is replicated to the global catalog<br><br>A value of TRUE means that the attribute is replicated to the global catalog. In environments where domain controllers are running Windows 2000 Server, changing this value might cause full synchronization of the Partial Attribute Set. For more information about the conditions under which full synchronization occurs, see How the Global Catalog Works on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=47817. |
| **systemOnly** | BOOL | No | No | Attributes on which Windows Server 2003 and Active Directory depend for normal operations<br><br>If TRUE, only the system can modify this attribute. No user-defined attribute must ever have the **systemOnly** flag set. |
| **objectClass** | Object identifier | Yes | Yes | Class of this object, which is always attributeSchema |
| **nTSecurityDescriptor** | NT-Sec-Des | Yes | No | Security descriptor on the attributeSchema object itself |
| **oMObjectClass** | String(Octet) | No | No | For object-syntaxed attributes (OM-syntax = 127), the Basic Encoding Rules (BER) encoded object identifier of the XOM object class<br><br>For more information about BER encoding, see Request for Comments (RFC) 2251 in the IETF RFC Database. |
| **LinkID** | Integer | No | No | Value that determines whether the attribute is a linked attribute. An even integer denotes a forward link; an odd integer denotes a back link.<br><br>A forward link references a target object in the directory; a back link refers back to the source object that has a forward link to it. |

**Single-value and multivalue attributes**
Some attributes contain a single value, and other attributes can contain multiple values. For example, the **password** attribute on a user object contains a single value: the password that is associated with the user account. The **memberOf** attribute contains multiple values: a list of the various groups of which the user account is a member. Each item in the list is stored as a separate value in the **memberOf** attribute. Attributes that can store more than a single value are referred to as "multivalue" attributes.

Whether or not attributes are single-value or multivalue is defined by the **singleValue** attribute as a Boolean value. The Active Directory Schema snap-in reports this attribute as "single-value" or multivalue rather than as the attribute-value pair. (Active Directory Schema is a Microsoft Management Console (MMC) graphical interface tool in Windows Server 2003 that administrators can use to manage the schema.)

All values that are defined for a multivalue attribute must have a uniform syntax. Multivalue attributes that are not linked can store a maximum of 1,300 entries. This limit is based on the size and type of the values that are stored.

### Note

- When it retrieves data, LDAP reads a multivalue attribute as a single entity. This can be inconvenient or even impossible when the number of values in a multivalue attribute becomes large. A property called **Range** can be specified as part of an attribute description to retrieve the values of a multivalue attribute incrementally. Servers might or might not honor the **Range** property. Servers that support the **Range** property include the object identifier 1.2.840.113556.1.4.802 in the **supportedControls** operational attribute on the **rootDSE**. Clients must not use the **Range** property unless this object identifier is present. The **Range** property is a constant, case-insensitive string value (Range=), followed by a range specifier that lists the initial value and terminal value in the range.

**Linked attributes**
Linked attributes make it possible to associate one object with another object. A linked attribute represents an interobject, distinguished-name reference. Linked attributes occur in pairs: forward link and backward link (back link). Linked attribute pairs exist on source objects and target objects, as follows:

- The forward link is a linked attribute on a source object (for example, the **member** attribute on the group object). It can be a single-valued attribute or a multi-valued attribute.

- As an option, a back link can be defined on a target object (for example, the **memberOf** attribute on the user object). A back-link attribute should be created as a multi-valued attribute, and it cannot exist without a corresponding forward link. The back-link attribute cannot be updated directly. Instead, it is automatically calculated when it is queried, based on the corresponding forward link. A back-link value on any instance of an object consists of the distinguished names of all source objects that have the target object's distinguished name in their corresponding forward link.

A linked pair is identified by the **linkID** values of two attributeSchema objects. The **linkID** of the forward link is an even, positive, nonzero value, and the **linkID** of the associated back link is the value of the forward **linkID**, incremented by one. For example, the value of **linkID** for **member** is 2 and the value of **linkID** for **memberOf** is 3.

**Automated links**
The **linkIDs** must be unique within the forest. In Windows 2000 Server, **linkIDs** are allocated by the applications that used them, and it is up to application developers to make sure that the **linkIDs** for the applications do not conflict with **linkIDs** that are already in use. Windows Server 2003 can generate link values automatically to ensure that there are no conflicts within the forest. When new objects that use linked attributes are created, Active Directory automatically generates the necessary **linkIDs**.

**Indexed attributes**
Directory searches for attributes that are indexed are more efficient than searches for attributes that are not indexed. Attributes are indexed when the least significant bit in their **searchFlags** attribute is set to the value 1. Changing the value of the bit to 1 dynamically builds an index; changing the value to 0 or deleting it removes an index for the attribute in question. The index is built automatically by a background thread on the directory server.

The values for indexed attributes are stored in a sorted list. This makes searching much more efficient because the system needs to search only until it locates the area in the list where the value should be, based on the sort. If the value is not there, the system can assume it will not find the value anywhere else in the list, and it can terminate the search. When attributes are not indexed, the entire list must be searched to determine whether or not a particular value actually exists.

Indexing requires more storage to maintain the lists, but it makes searching more efficient. Nonindexed attributes are less efficient to search, but they require less storage to maintain. With this in mind, only attributes that are frequently referenced

should be indexed. Ideally, indexed attributes are single-value attributes with unique values that are evenly distributed across the set of instances. Multivalue attributes can be indexed, but building the index requires more storage and updating.

**Confidential attributes**

On domain controllers running Windows Server 2003 SP1 and later, attributes whose **attributeSchema** objects are marked with a specific search flag are interpreted as confidential. Attributes that are so marked can be read only by users to whom the following access is granted:

- READ_PROPERTY
- CONTROL_ACCESS

The ability to mark attributes as confidential allows administrators to protect attributes from the read access that is granted by default to most users. Rather than changing the defaults that are expected by existing applications, administrators can create new attributes that can be read only by administrators or those to whom access is specifically granted. Because domain controllers must be running Windows Server 2003 SP1 or later to recognize the flag that marks the attribute as confidential, protection of marked attributes is fully effective only in environments where all domain controllers have been upgraded to Windows Server 2003 SP1 or later. In particular, the schema operations master (also known as the schema flexible single-master operations (FSMO) role holder) must be upgraded to Windows Server 2003 SP1 or later.

By default, Active Directory performs a read-access check on an object in the following cases:

- When evaluating whether the object matches the search filter
- When returning attributes of an object that matches the search filter

Default security allows read access to the Pre-Windows 2000 Compatible Access group, which usually contains Authenticated Users. Authenticated Users is a security group that includes users whose identities can be authenticated by the server or by a trusted security authority. Because the Pre-Windows 2000 Compatible Access group is granted Read_Property permissions on **user** objects (including **inetOrgPerson**) and **group** objects, it is difficult to introduce a new attribute on these objects that is protected from reading by most users. For example, you might want to add a social security number attribute to the **user** class.

To solve this problem, Windows Server 2003 SP1 and later provides a way to mark an attribute as confidential by setting a **searchFlags** value on the attributeSchema object in the schema. The search flag contains multiple bits representing various properties of an attribute. In the schemas, the **searchFlags** attribute has six bits that can be set. (See the bits that are defined for the **searchFlags** attribute in "Attributes for the attributeSchema class" earlier in this section.) For example, a value of 1 in the least significant bit means that the attribute is indexed. The binary form of this flag value is 00000001. A new bit is added in Windows Server 2003 SP1 or later: A value in bit 128 (10000000) designates the attribute as confidential. Therefore, on a domain controller running Windows Server 2003 SP1 or later, you can designate an attribute as confidential by setting bit 128 in the **searchFlags** attribute of the respective attributeSchema object.

| **Note** |
| --- |
| You cannot set this flag on base-schema attributes, such as common-name. |

When an attribute is so designated, in addition to READ_PROPERTY permission (for that attribute or all attributes on that object), CONTROL_ACCESS permission (for that attribute or a property set that contains that attribute) is required to read the attribute.

| **Note** |
| --- |
| When you assign CONTROL_ACCESS permissions to a user or group, you must specify a GUID — in this case, the GUID of the attribute or property set that contains the attribute. |

To set the bit flag when no values are set, you can simply add the value 128 to the **searchFlags** attribute in the properties of the attributeSchema object for the confidential attribute. If a value is already present in the attribute, perform a bitwise OR operation to add 10000000 to the existing binary value, and then convert to decimal. For example, if the bit is set for indexing and containerized search (00001001), the combined values 10000000 OR 00001001 equal 10001001, translating to a decimal value of 137.

By default, only members of the Administrators built-in group are granted CONTROL_ACCESS to all objects. Therefore, only

administrators can read confidential attributes. You can delegate the CONTROL_ACCESS right to any specific group by using the Dsacls command-line tool or by scripting. Dsacls is included in Windows Support Tools. To grant access to a confidential attribute, use Dsacls to grant the CONTROL_ACCESS right (CA) to the required group or user. Doing so introduces a way to impose additional security checks that control read access to selected attributes. For more information about Dsacls, see Dsacls.exe in the Active Directory Management Support Tools section of Windows Support Tools on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=43235. For more information about property sets and CONTROL_ACCESS permissions, see the Security Descriptors and Access Control Lists Technical Reference.

**Syntaxes**

The syntax for an attribute defines the storage representation, byte ordering, and matching rules for comparisons of property types. The syntax defines whether the attribute value must be a string, a number, or a unit of time. Each attribute of every object is associated with exactly one syntax. The syntaxes are not represented as objects in the schema, but they are programmed to be understood by Active Directory. The allowable syntaxes in Active Directory are predefined. You cannot add new syntaxes.

When you define a new attribute, you must specify both the **attributeSyntax** and the **oMSyntax** numbers of the syntax that you want for that attribute. The **attributeSyntax** number is an object identifier, and the **oMSyntax** number is an integer. **oMSyntax** is defined by the XOM specification. Using this model, the syntax can provide detailed syntax definitions. For example, distinct **oMSyntax** attributes distinguish several types of printable strings, according to such factors as the supported character set and whether case is significant. The following table lists the valid syntaxes for attributes in the Active Directory schema.

**Valid Syntaxes for Attributes in the Active Directory Schema**

| Syntax | attributeSyntax | oMSyntax | ASN 1-Encoded Object Identifier | Description |
|---|---|---|---|---|
| Undefined | 2.5.5.0 | | \x550500 | Not a legal syntax |
| Object(DN-DN) | 2.5.5.1 | 127 | \x550501 | The fully qualified name of an object in the directory |
| String(Object-Identifier) | 2.5.5.2 | 6 | \x550502 | The object identifier |
| Case-Sensitive String | 2.5.5.3 | 27 | \x550503 | General string — differentiates uppercase and lowercase |
| CaseIgnoreString(Teletex) | 2.5.5.4 | 20 | \x550504 | Teletex — does not differentiate uppercase and lowercase |
| String(Printable), String(IA5) | 2.5.5.5 | 19, 22 | \x550505 | Printable string or IA5-String<br><br>Both character sets are case sensitive. |
| String(Numeric) | 2.5.5.6 | 18 | \x550506 | A sequence of digits |
| Object(DN-Binary) | 2.5.5.7 | 127 | \x550507 | A distinguished name plus a binary large object |
| Boolean | 2.5.5.8 | 1 | \x550508 | TRUE or FALSE values |
| Integer, Enumeration | 2.5.5.9 | 2, 10 | \x550509 | A 32-bit number or enumeration |
| String(Octet) | 2.5.5.10 | 4 | \x55050A | A string of bytes |
| String(UTC-Time), String(Generalized-Time) | 2.5.5.11 | 23, 24 | \x55050B | UTC time or generalized-time |
| String(Unicode) | 2.5.5.12 | 64 | \x55050C | Unicode string |
| Object(Presentation-Address) | 2.5.5.13 | 127 | \x55050D | Presentation address |
| Object(DN-String) | 2.5.5.14 | 127 | \x55050E | A DN-String plus a Unicode string |

| | | | | |
|---|---|---|---|---|
| String(NT-Sec-Desc) | 2.5.5.15 | 66 | \x55050F | A Windows NT security descriptor |
| LargeInteger | 2.5.5.16 | 65 | \x550510 | A 64-bit number |
| String(Sid) | 2.5.5.17 | 4 | \x550511 | Security identifier (SID) |

**Object identifiers**

Object identifiers are unique numeric values that are granted by various issuing authorities to identify data elements, syntaxes, and other parts of distributed applications. Because they are globally unique, object identifiers ensure that the objects that are defined by these issuing authorities do not conflict with one another when different directories, such as Active Directory and Novell Directory Services, are used concurrently within a global directory namespace.

Object identifiers are found in Open Systems Interconnection (OSI) applications, X.500 directories, applications that use Simple Network Management Protocol (SNMP), and other applications in which uniqueness is important. Object identifiers are based on a tree structure in which a superior issuing authority allocates a branch of the tree to a subordinate authority, which in turn allocates subbranches of the tree.

LDAP requires a directory service, such as Active Directory, to identify object classes and attributes with an object identifier syntax. The object identifier is the value for the **governsID** attribute in a classSchema object and for the **attributeID** attribute in an attributeSchema object. These are required attributes; therefore, object identifiers are necessary when you create new classes or attributes.

Object identifiers in the Active Directory base schema include some object identifiers that are issued by the International Organization for Standardization (ISO) for X.500 classes and attributes and some object identifiers that are issued by Microsoft. Object identifier notation is a dotted string of nonnegative numbers, for example, 1.2.840.113556.1.5.4. The components of this object identifier are shown in the following table.

**Components of an Object Identifier (1.2.840.113556.1.5.4)**

| Numerical Values of the Object Identifier | What the Numerical Values Mean |
|---|---|
| 1 | ISO (“root” authority) issued 1.2 to ANSI. |
| 2 | ANSI issued 1.2.840 to the USA. |
| 840 | USA issued 1.2.840.113556 to Microsoft. |
| 113556 | Microsoft internally manages several object identifier branches under 1.2.840.113556 that include the following three identifiers. |
| 1 | A branch called Active Directory that includes the following two identifiers. |
| 5 | A branch called Classes that includes the following identifier. |
| 4 | A class called Builtin-Domain. |

Object identifiers ensure that every object is interpreted appropriately, for example, that a telephone number is not mistaken for an employee number. A series of widely used objects and attributes is standardized for use in object identifiers. New object identifiers are issued by standards authorities, and they form a hierarchy below which new object identifiers can be managed internally.

Most countries and regions in the world have an identified National Registration Authority (NRA) that is responsible for issuing object identifiers to enterprises. In the United States, the NRA is the American National Standards Institute (ANSI). The NRA issues root object identifiers. An enterprise can register a name for the object identifier as well. A fee is associated with registering root object identifiers and registered names. Contact the NRA for your country or region for details. The ISO recognizes NRAs and maintains a list of contacts on its Web site.

The issuing authority assigns an object identifier space that is a branch of the ISO-International Telecommunication Union (ITU) object identifier tree. For example, your organization might be assigned the space 1.2.840.111111. You can extend this space internally within the constraints of the structure of an object identifier. You can subdivide this space further (by appending dotted decimals to the object identifier root) and assign these subspaces to various divisions within your organization. Each division, in turn, can further subdivide the subspace that is allotted to it.

For the example object identifier 1.2.840.111111, your organization might use the subspace 1.2.840.111111.1.4 for attributes

and the subspace 1.2.840.111111.1.5 for classes. An internal issuing authority in your organization, using an Administrator account, might then allocate object identifiers from this space when requested. The **governsID** attribute on every classSchema object and the **attributeID** attribute on every attributeSchema object are mandatory attributes that contain an object identifier string. In this example, all of your organization-created classSchema objects have a **governsID** attribute of the form 1.2.840.111111.1.5.*x*, where *x* is a decimal number. Similarly, all of your organization-created attributeSchema objects have an **attributeID** attribute of the form 1.2.840.111111.1.4.*x*.

Microsoft also offers a free object identifier registration service.

#### classSchema Objects

The classSchema object specifies the various attributes of the class with which it is associated. Among other things, it defines the following constraints for objects that are instances of the class:

- **mustContain** attributes, which include mandatory attributes that must be present on any object that is an instance of this class
- **mayContain** attributes, which include optional attributes that may be found on an object that is an instance of this class
- Hierarchy rules that determine the possible parents in the directory tree of an object that is an instance of the class

An object can have attributes that belong only to either the **mustContain** list or the **mayContain** list for the class.

The classSchema object is essentially a template that contains the rules for creating objects in an Active Directory class. When a new object is created in a class, the classSchema object ensures that this new object has the same attributes as all other objects in the class.

The classSchema object contains, among other things, the following information:

- The LDAP display name of the class
- The object identifier for the class
- The GUID for the class
- The attributes that must be present for an instance of the class
- Other attributes that can be present for an instance of the class
- The classes to which the parent of instances of this class may belong
- The superclass from which this class inherits characteristics
- Other **auxiliary** classes from which this class inherits attributes
- The type of class (**abstract**, **structural**, or **auxiliary**)
- The default hiding state for the class. If you do not want instances of the class to be displayed by the end-user user interface (UI), you can define the class as hidden by default.

#### Attributes for classSchema objects

The following table lists the attributes that a classSchema object can have.

**Attributes for a classSchema Object**

| Attribute | Syntax | Mandatory? | Multivalue? | Description |
|---|---|---|---|---|
| **cn** | Unicode | Yes | No | Descriptive relative distinguished name for the schema object |

| | | | | |
|---|---|---|---|---|
| **governsID** | Object identifier | Yes | No | The object identifier that uniquely identifies this class |
| **lDAPDisplayName** | Unicode | Yes | No | The name by which LDAP clients identify this class |
| **schemaIDGUID** | String(Octet) | Yes, but defaulted | No | The GUID that uniquely identifies this class |
| **rDNAttID** | Object identifier | No | No | The relative distinguished name type of instances of this class (OU, CN) |
| **subClassOf** | Object identifier | Yes | No | The class from which this object inherits attributes |
| **systemMustContain** | Object identifier | No | Yes | The list of mandatory attributes for instances of this class<br><br>This list cannot be changed. |
| **mustContain2** | Object identifier | No | Yes | The mandatory attributes for instances of this class |
| **systemMayContain** | Object identifier | No | Yes | The optional attributes for instances of this class |
| **mayContain2** | Object identifier | No | Yes | The optional attributes for instances of this class |
| **systemPossSuperiors2** | Object identifier | No | Yes | The classes that can be parents of this class in the directory hierarchy<br><br>After the class is created, this property cannot be changed. |
| **possSuperiors2** | Object identifier | No | Yes | The classes that can be parents of this class in the directory hierarchy<br><br>For an existing classSchema object, values can be added to this property but not removed. |
| **systemAuxiliaryClass2** | Object identifier | No | Yes | The **auxiliary** classes from which this class inherits its optional (**mayContain**) and mandatory (**mustContain**) attributes<br><br>After creation of the class, this property cannot be changed. |
| **auxiliaryClass2** | Object identifier | No | Yes | The **auxiliary** classes from which this class inherits its optional (**mayContain**) and mandatory (**mustContain**) attributes<br><br>This is a multivalue property that specifies the auxiliary classes that this class inherits from. For an existing classSchema object, values can be added to this property but not removed. |
| **defaultHidingValue** | BOOL | No | No | The default hiding state for the class<br><br>If you do not want instances of the class displayed in the UI for Active Directory admin tools **New** menus, you can define the class as hidden. |

| | | | | |
|---|---|---|---|---|
| **defaultSecurityDescriptor** | String(Octet) | No | No | The default security descriptor that is assigned to new instances of this class if no security descriptor is specified during creation of the class or that is merged into a security descriptor if a security descriptor is specified |
| **objectClassCategory** | Integer | Yes | No | Class types are defined as follows:<br>88 Class = 0<br>Structural = 1<br>Abstract = 2<br>Auxiliary = 3<br>The 88 class type is included for compatibility with older directories. Active Directory does not use 88 class type objects, and they should not be used in defining new classes for Active Directory. |
| **systemOnly** | BOOL | No | No | An attribute of a classSchema object<br>If it is set to TRUE, only the system can create and modify instances of this class. If it is set to FALSE, modifications can also be made by users who have appropriate permissions. |
| **objectClass** | Object Identifier | Yes | Yes | This object's class, which is always **classSchema** |
| **nTSecurityDescriptor** | NT-Sec-Desc | Yes | No | The security descriptor on the classSchema object |
| **defaultObjectCategory** | Distinguished name | Yes | No | The default object category of new instances of this class<br>If none has been specified, the **objectClass** value is used. |

The attributes in a classSchema object's **mustContain** attribute list are not the complete set of attributes that must be present for an instance of a class to exist. For example, in a class A, the classSchema object B specifies a list of **mustContain** attributes that an instance of A must have through the **systemMustContain** and **mustContain** attributes. However, because mandatory attributes are also inherited, the complete list of attributes for an instance of class A includes the inherited **mustContain** attributes from all classes from which B inherits, that is, all classes in the **subClassOf** and **auxiliaryClass** lists for the classSchema object B. The **mayContain** attributes for object B are also defined this way. The **possSuperiors** attributes are defined this way as well, except that **possSuperiors** attributes are inherited only from classes in the **subClassOf** list, not from the classes in the **auxiliaryClass** list.

#### Mandatory attributes

Mandatory attributes are object attributes for which you must specify values. If you do not specify a value for a mandatory attribute, the attribute receives a default value or the object is not created until you specify a value for the attribute. The object's mandatory attributes are determined by the class to which the object belongs.

Some mandatory attributes are inherited. Because every classSchema object belongs to a superclass called **Top** in the class hierarchy, each classSchema object inherits the mandatory attributes that belong to **Top**. The following table lists the mandatory attributes that every object inherits from **Top**. To see a graphical representation of the Active Directory class hierarchy, see the figure under "Object class hierarchy" later in this section.

**Mandatory Attributes That All classSchema Objects Inherit**

| Inherited Mandatory Attribute | Default Status |
|---|---|
| **nTSecurityDescriptor** | Set to a default value if not specified. The default value depends on the default security descriptor for the classSchema object. |

| | |
|---|---|
| **objectCategory** | Automatically set to the value of the default object category of the class, which is usually the class itself. Can be changed after the class is created. |
| **objectClass** | No default. An administrator must specify the class. |

You can view an object's mandatory attributes by using the Active Directory Schema snap-in. The attributes are displayed on the **Attributes** tab in the properties dialog box for the object. Because some of an object's mandatory attributes are inherited from its parent class, you might need to view the attributes of the parent class to identify all the mandatory attributes of the object.

**System and changeable attribute pairs**

Some properties of a class-definition object are contained in pairs of attributes, in which the value of one of these attributes can be changed by administrators and the value of the other attribute cannot be changed. These attribute pairs are as follows:

- **mustContain/systemMustContain**
- **mayContain/systemMayContain**
- **possSuperiors/systemPossSuperiors**
- **auxiliaryClass/systemAuxiliaryClass**

In each of these pairs, the value of the attribute that begins with the word "system" cannot be changed by administrators. This enables Active Directory to protect certain key attributes of certain classes and to ensure that the schema stays consistent and usable. System-only properties can be changed only by the Directory System Agent (DSA). System-only properties are those properties on which Active Directory depend for normal operations. For example, the **attributeID** attribute and the **governsID** attribute in the schema are system-only attributes. The values of the other (nonsystem) attributes in each pair can be changed by administrators.

**Categories of Object Classes**

Different categories of object classes make it possible to define structure in the directory. The 1993 X.500 specification requires that object classes be assigned to one of three categories:

- Structural
- Abstract
- Auxiliary

A fourth category, which is referred to as 88 classes, is associated with object classes, although this category is not part of the 1993 specification. The 88 class category is in the specification that existed before 1993, and it should not be used for adding classes to Active Directory.

**Structural classes**

Structural classes are the only classes that can have instances in the directory. That is, you can create directory objects whose class is one of the structural classes. A structural class:

- Can be used in defining the structure of the directory.
- Is derived from either an abstract class or another structural class.
- Can include any number of auxiliary classes in its definition.

This type of class is specified by a value of 1 in the **objectClassCategory** attribute.

**Abstract classes**

Abstract classes are templates that are used to derive new structural classes. Abstract classes cannot be instantiated in the directory. This means that no object can belong only to an abstract class; each object of an abstract class also belongs to some structural subclass of that class. A new abstract class can be derived from an existing abstract class. This type of class is specified by a value of 2 in the **objectClassCategory** attribute. Abstract classes only provide attributes for subordinate classes,

which are called subclasses. A subclass contains all mandatory and optional attributes of the class from which it is derived (its superclass) in addition to those attributes that are specific to the class itself. Likewise, the subclass of that class contains all attributes of both superclasses, and so forth.

**Auxiliary classes**

Auxiliary classes are like include files; they contain a list of attributes. Adding an auxiliary class to the definition of a structural class or an abstract class' adds the auxiliary classs attributes to the definition. An auxiliary class cannot be instantiated in the directory, but new auxiliary classes can be derived from existing auxiliary classes. This type of class is specified by a value of 3 in the **objectClassCategory** attribute. For example, the **securityPrincipal** class is an auxiliary class, and it derives its attributes from the parent abstract class called **Top**. Although you cannot create a security principal object in the directory (because auxiliary classes cannot have instances), you can create an object of the structural class **user**, which has the **securityPrincipal** class as an auxiliary class. The attributes of the **securityPrincipal** class help the system recognize the user object as a security account. Similarly, the **group** class has **securityPrincipal** as an auxiliary class.

The behavior of auxiliary classes has changed in Windows Server 2003 and later. In Windows 2000 Server, changes that are made to an auxiliary class affect its parent class as well as all instances of the parent object. For example, adding an auxiliary class called **pager** to the structural class **user** affects all instances of **user**, which is all of the user accounts that are created with the **user** class.

In Windows Server 2003 and later, auxiliary classes can be assigned dynamically to individual instances of classes, rather than being applied automatically to all instances. For example, you can assign the **pager** auxiliary class to only those users who need it.

**88 classes**

Classes that were defined before 1993 are not required to fall into one of the preceding categories; assigning classes to categories was not required in the X.500 1988 specification. Classes that were defined before the X.500 1993 standards are assigned to the 88 class. This type of class is specified by a value of 0 in the **objectClassCategory** attribute. *Do not use 88 classes or define new 88 classes.*

**Object class hierarchy**

Inheritance, which is also referred to as derivation, is the ability to build new object classes from existing object classes. A new object is defined as a subclass of its parent object. A subclass is a class that inherits from some other class; for example, a subclass inherits structure and content rules from the parent class. The parent object becomes a superclass of the new object. A superclass is a class from which one or more other classes inherit information. The inherited information includes mandatory and optional attributes (**systemMustContain**, **mustContain**, **systemMayContain**, and **mayContain**) and the parent classes in the directory hierarchy (**systemPossSuperiors** and **possSuperiors**). This relationship between the superclasses and their subclasses represents the object class hierarchy, which is illustrated in the following figure.

**Object Class Hierarchy**




All structural object classes are subclasses, directly or indirectly, of a single abstract object class, which is called **Top**. Every object that is represented in the directory belongs to **Top**, and as a result every entry must have an **objectClass** attribute. When you create a new class, you must specify the superclass. If you do not create a subclass of an existing class, the new class is a subclass of **Top**.

A new class can inherit mandatory and optional attributes from more than one existing class. However, any additional classes must be specified by the **auxiliaryClass** attribute.

### Note

- If you later add another attribute to a class that has subclasses or auxiliary subclasses, the new attribute is automatically added to the subclasses after the schema cache has been updated.

#### Structure and Content Rules

The schema enforces rules that govern both the structure and the content of Active Directory. These rules validate changes to objects to ensure the integrity of the directory.

#### Structure rules

Structure rules define the possible tree structures. When you create a new object, structure rules determine the validity of the object class to which you designate the new object. You cannot create an object that belongs to a nonexistent class. You must first create the new class.

Conversely, these rules do not allow you to delete or modify an object that has already been deleted. In Active Directory, the structure rules are completely expressed by the **possSuperiors** and **systemPossSuperiors** attributes that are present on each classSchema object. These attributes specify the possible classes that can be parents of an object instance of the class. In other words, the **possSuperiors** and **systemPossSuperiors** attribute values determine the object classes and, therefore, the location in the directory tree under which objects of the class can be instantiated.

#### Content rules

Content rules determine the mandatory and optional attributes of the class instances that are stored in the directory. New objects must contain all the mandatory attributes that are specified by the classSchema object in the schema. New objects can contain any of the optional attributes. In Active Directory, the content rules are completely expressed by the **mustHave**, **mayHavemayContain**, **systemMustContain**, and **systemMayContain** attributes of the schema definitions for each class. In addition, there are some operational attributes that are read-only. These are identified by the first bit of the **systemFlags** property for an attribute definition that is set to 1, for example, **lastLogon**.

For more information about mandatory and optional attributes, see Active Directory Schema in the Microsoft Platform SDK on MSDN.

#### Active Directory Schema Physical Structure

The schema exists as a set of directory objects, and it is stored in the directory. Active Directory partitions the information in the directory to facilitate more efficient replication. The schema is stored in its own directory partition so that it can replicate independently of other data that is stored in the directory.

#### Schema Location

The objects that are stored in Active Directory are arranged in a logical hierarchy called the directory tree. The directory tree is divided into directory partitions. A directory partition is a tree of objects that forms a unit of replication in Active Directory. The schema has its own directory partition to prevent potential dependency problems that can arise when new schema classes and new instances of the class are replicated simultaneously. Because the schema has its own tree, it is possible to replicate new schema changes to other domain controllers before replicating any new objects that may have been created based on those changes. This is important because the other domain controllers must have access to the object definitions that are stored in the schema before those domain controllers can properly store any new objects that are created by using those definitions.

Active Directory includes a preconfigured database, commonly referred to as the base directory information tree (DIT), which contains the information that is required to install and run Active Directory. The base DIT is created during a fresh installation of a Windows Server 2003 or later domain controller. The base DIT is contained in a file named Ntds.dit. One section of the base DIT is the base schema.

#### Schema Head

The schema head is the topmost object of the schema directory partition. Conceptually, the schema head functions like other containers in the directory tree, which means that it contains all of the schema information. However, the schema head is not a container in the sense of a special Active Directory object that contains other objects; rather, it is a special-purpose object class. Therefore, it is referred to as the schema head rather than the Schema container. The schema head contains all of the class definitions and attribute definitions that are required to locate objects in Active Directory and create new objects.

The relationships of the schema head, Configuration container, and Domain container are illustrated in the following figure.

### Schema Head




Every Active Directory object can be referenced by a unique and unambiguous name known as the distinguished name. The distinguished name identifies the domain that holds the object as well as the complete path through the container hierarchy by which the object is reached. The distinguished name of the schema head can be expressed as follows:

cn=schema,cn=configuration,dc=*ForestRootDomainName*

You can view the contents of the schema head by using the Active Directory Schema snap-in. You can also bind to the schema directory partition and view schema objects by using the Active Directory Service Interfaces (ADSI) Edit snap-in or the Ldp tool.

**Note**

- ADSI Edit is not one of the default MMC snap-ins provided. To use ADSI Edit and Ldp, install the support tools that are located in the Support\Tools folder on the Windows Server 2003 or later operating system CD.

You can locate the schema head without knowing the domain name. Installation scripts and other applications can gain access to the schema because they bind to a special entry at the top of the logical namespace, called the root DSA-specific Entry (**rootDSE**), which provides the schema location. **rootDSE** represents the top of the logical namespace and, therefore, the top of the LDAP search tree. The attributes of **rootDSE** identify, among other things, the directory partitions — that is, the domain, schema, and configuration directory partitions — as well as the forest root domain directory partition. One attribute, **schemaNamingContext**, provides the location of the schema so that applications that connect to any domain controller can find and read the schema.

**subSchema Subentry**

**rootDSE** also carries a mandatory attribute called **subSchemaSubEntry**. The value of this attribute is the distinguished name of a subSchema object in the directory that defines the attributes (in **attributeTypes**) and classes (in **objectClasses**) that make up the Active Directory schema. This special object, which is an instance of the unique **subSchema** class, is used for administering information about the schema, in particular, the object classes and attribute types that are supported. This enables client applications to retrieve the schema information by querying the **subSchema** entry. Clients must only retrieve attributes from a **subSchema**entry by requesting a base object search of the entry, where the LDAP search filter is (objectClass=subSchema). The following distinguished name describes the location of the SubSchemaSubEntry container:

CN=Aggregate,CN=Schema,CN=Configuration,DC=*DomainName*,DC=*DomainRoot*

For more information about LDAP searches, distinguished names, and containers, see "Data Store Technical Reference."

**Schema Files**

Active Directory data is distributed among all domain controllers in the forest. No single domain controller stores all Active Directory data for the entire forest, but every domain controller does hold a copy of the schema. The database that stores the Active Directory data on a particular domain controller is in a file named Ntds.dit. The location of the Ntds.dit file is an option that you can set during the domain controller promotion process while you create the directory. The default location for Ntds.dit and the database log files is *systemroot*\Ntds. For more information about the Ntds.dit file, see "Data Store Technical Reference."

Another file, the Schema.ini initialization file, contains the information that is necessary for creating the default directory objects and the default security for the directory tree, as well as the Active Directory display specifiers. Although this file is named Schema.ini, the schema itself is actually stored in the Active Directory database, Ntds.dit. Schema.ini is only used by the Active Directory Installation Wizard to build the initial schema structure in the directory during the domain controller promotion process. The Schema.ini file must not be modified.

**Active Directory Schema Processes and Interactions**

Normally, you do not interact directly with the schema on a daily basis. Active Directory uses the schema to help create objects that are stored in the directory. You interact with those objects, not the schema. You interact directly with the schema when you make modifications to the schema by adding definitions to it or by modifying existing definitions.

Schema changes can affect the entire directory. Before you make any changes to the schema, you should thoroughly test those changes in an isolated environment to ensure that the directory continues to function as planned after the changes have been deployed. Only members of the Schema Admins group can make changes to the schema.

The two most common reasons for modifying the schema are:

- Installation of an application that needs to add customized object definitions so that it can store information in the directory, for example, an e-mail program that stores the user's e-mail names in the directory

- Testing of the development of applications that use the directory for data storage in which customized object definitions need to be added to the schema and modified throughout their lifetimes as the development process proceeds

**Modifying the Schema**

When the existing class and attribute definitions in the schema do not meet the needs of your organization, you can add or modify schema objects to extend the schema. You modify an existing attribute or add a new class or attribute to the schema to store a new type of information in the directory. The process of modifying or updating the schema is often referred to as extending the schema.

Before you extend the schema, you must take steps to ensure that the extension does not cause problems in the directory. Do not extend the schema in your production forest without testing the extension in a private forest. First, make changes in a test environment, and check that the changes behave as expected and that they meet your needs. After verifying the changes, you can use various utilities, such as Ldifde, and scripts or customized applications to export the extensions from the test environment and import them into the production environment. You perform most schema extensions by using applications or scripts that are written to extend the schema.

**Note**

- As is true for every object in Active Directory, schema objects are protected by access control lists (ACLs). Therefore, only authorized users can alter the schema.

To add or modify a class definition or attribute definition, you add or modify the corresponding **classSchema** object or **attributeSchema** object. This process is similar to adding or modifying any object in Active Directory, except that additional checks are performed to ensure that changes do not cause inconsistencies or problems in the schema.

**Adding Information to the Schema**

Extending the schema is a major change, with implications throughout the directory. Extend the schema only when it is absolutely necessary. Many schema modifications cannot be reversed; therefore, you must thoroughly plan and test changes before you deploy them in your production forest.

**Planning to extend the schema**

Planning to extend the schema involves examining the default schema that comes with Active Directory to verify that there is no way to use the existing classes or attributes for your needs. It also involves understanding the types of modifications that can and cannot be made.

For more information about modifying the schema, see "Active Directory Schema" in the Microsoft Platform SDK on MSDN.

**Enabling schema extensions**

After you decide to make changes to the schema and you carefully plan the types of changes that you are going to make, you can proceed by implementing and testing the schema extensions in a test environment that is isolated from your production network.

Because extending the schema is a significant operation, Active Directory has the following safety features, or interlocks, that restrict schema modifications:

- The Active Directory Schema snap-in must be registered. Active Directory Schema is not listed with the default MMC snap-ins. It must be registered before it can be made available. This means that Schmmgmt.dll must be registered by using Regsvr32.exe.

- Changes to the schema must be written only on the schema master. Although all domain controllers have a copy of the schema in their Active Directory database, only the domain controller that holds the schema operations master role (also known as flexible single master operations (FSMO)) is allowed to write changes to the schema.

- Administrators must have specific access rights. Only members of the schema administrators group (Schema Admins) or another administrator with explicit permission can make changes to the schema.

- Schema modifications must be enabled. By default, schema modification is disabled on all domain controllers, including

the domain controller that hosts the schema operations master role.

**Identifying or transferring the schema operations master role**
Active Directory performs schema updates in a single-master fashion to prevent conflicts such as conflicts resulting from simultaneous schema updates on two different computers. The one domain controller in the forest that is allowed to perform schema updates at any specific time is referred to as the schema master. The schema master is the domain controller that holds the schema operations master role. After the schema master completes an update, it replicates the changes to all other domain controllers by normal replication channels. In this way, changes to the schema are distributed throughout the forest.

### Note

- To update the schema, the domain controller that holds the schema operations master role must be available on the network. All schema modifications must be made to the domain controller that holds the schema operations master role. If you attempt to modify the schema from a domain controller that does not hold that role, the domain controller generates a referral to the current schema master to process the modifications.

Because the schema master must be used to extend the schema, the domain controller that currently holds the schema operations master role in the forest must be identified. As the forest grows and changes, it may also become necessary to assign the schema operations master role to a different domain controller. You can accomplish both of these tasks by using Active Directory Schema or the command-line tool Ntdsutil.

You can change the domain controller that serves as the schema master at any time. The current schema master in the enterprise is identified by the value of the **FSMORoleOwner** attribute on the schema head of the domain. By default, the first domain controller that is installed in the forest is the initial schema master.

**Granting access rights to make schema changes**
To modify the schema, you must use an account that is a member of the Schema Admins group. By default, the only member in the Schema Admins group is the Administrator account in the root domain of the enterprise. You must explicitly add other accounts.

You can use Active Directory Users and Computers in MMC to verify that an account is a member of the Schema Admins group. Restrict membership in the Schema Admins group to prevent unauthorized access to the schema. Improper modification of the schema can have serious consequences.

By default, only members of the Schema Admins group have permission to write to the schema. You can grant explicit permissions to use Active Directory Schema to specific users. However, this is not recommended.

**Enabling schema modifications on a domain controller**
If you want to allow the domain controller that holds the schema operations master role to modify the schema, use Active Directory Schema to enable schema modifications.

**Removing Information from the Schema**
It is not possible to remove object definitions from the schema. However, object definitions can be rendered unusable through the process of deactivation. When an object definition is deactivated, it can no longer be used to create new objects in the directory. Objects whose definitions have been deactivated in the schema are referred to as defunct.

**Deactivating schema objects**
You cannot deactivate schema objects that are part of the default schema that ships with Active Directory. You can freely deactivate schema objects that have been added to the default schema.

As a result of replication latency, it is not possible to accurately determine if any objects have been created by using a given schema definition or to predict if the objects may be restored from backup media. Therefore, Active Directory does not support the actual deletion of schema objects. Rather, it provides a mechanism for deactivating schema objects in such a way that they become unavailable for use in the directory. Although a schema object still physically exists in the directory after it has been deactivated, new instances of it cannot be created in the directory. Any existing instances of data that are associated with the deactivated schema object continue to exist in the directory; however, there is no way to modify these data instances other than to delete them. This behavior makes the schema object appear to be deleted from the schema.

You can use Active Directory Schema or ADSI Edit to deactivate a class or an attribute by setting the attribute **isDefunct** to the Boolean value of TRUE on the schema object.

You can use Active Directory Schema to view defunct schema objects: on the **View** menu, make sure that **Defunct Objects** is selected. In addition, you can write programs and scripts to search for all schema objects that have the attribute **isDefunct** set

to TRUE. You can also use the Search function of the Ldp tool to search the schema with a filter that is set to isDefunct=TRUE.

As with the addition or modification of classes or attributes, some special validation checks are performed on the deactivation of classes or attributes to ensure consistency of the schema. In particular, when you deactivate a class, Active Directory verifies that the class is not used in the **subClassOf**, **auxiliaryClass**, or **possSuperiors** attributes of any existing active class. Similarly, when you deactivate an attribute, Active Directory checks that the attribute is not used in the **mustContain** or **mayContain** attributes of any existing active class.

**Deactivating existing classes and attributes**

Deactivation of schema classes and attributes is subject to the following restrictions:

- You cannot deactivate a class or attribute that appears in the base Active Directory schema. You can only deactivate schema extensions of the base schema.

- You cannot deactivate an attribute that is referenced in the **mustContain**,**systemMustContain**, **mayContain**, **systemMayContain**, or **rdnAttId** properties of an active class.

- You cannot deactivate a class that is referenced in the **subClassOf**, **auxiliaryClass**, or **possibleSuperior** properties of an active class.

To deactivate an attribute, set the **isDefunct** attribute of itsattributeSchema object to TRUE. When an attribute is deactivated, it can no longer be added to new class definitions. To reactivate an attribute, set the **isDefunct** attribute to FALSE.

To deactivate a class, set the **isDefunct** attribute of its classSchema object to TRUE. When a class is deactivated, new object instances of the class can no longer be created. To reactivate a class, set the **isDefunct** attribute to FALSE.

**Effects of deactivating a schema object on schema updates**

After a class **salesUser** is deactivated, any subsequent addition or modification of instances of **salesUser** fails as if **salesUser** has been deleted from the system; the same error codes are returned as if **salesUser** never existed at all. For example, creating a new instance of **salesUser** fails, and trying to modify or rename an existing instance of **salesUser** fails. Similarly, if an attribute **hasPager** is deactivated, **hasPager** is treated as nonexistent for new object creations, and attempting to modify (add or replace) the value of **hasPager** in an existing object fails.

However, you can still search for and delete existing instances of deactivated schema objects. For example, you can still search for all existing instances of **salesUser** and delete them, if necessary. Note that what you are doing is searching for and deleting objects that were created in the directory using the deactivated class, **salesUser**, not deleting the actual class definition from the schema. Similarly, you can search for all instances that have a value for the attribute **hasPager** and delete **hasPager** from the existing objects. This facilitates cleanup after a schema object is deactivated. If you decide that a class is not needed anymore, you can deactivate it so that no one can use it for any modifications. You can then clean up the existing instances of the class by searching for all instances and deleting them. Active Directory does not perform any automatic cleanup of data instances after a schema object is deactivated.

Similarly, you can deactivate an attribute and clean up all its instances. Note that, for a deactivated attribute, you can delete only the entire attribute from an object, not certain values of the attribute. For example, if **hasPager** is a multivalue attribute and an object has more than one value for **hasPager**, you cannot delete only one value of **hasPager** from the object.

In addition to affecting instances of the schema object, deactivating a schema object also affects schema updates, because schema object updates are subject to special validation checks. For example, if you deactivate the class **salesUser** or the attribute **hasPager**, subsequent schema object updates behave as follows:

- Deactivated classes and attributes can be renamed in the schema. They can also be modified to make other changes to the corresponding classSchema and attributeSchema objects. Note that this behavior is different from the behavior of the Windows 2000 Active Directory schema, in which no modifications are allowed on defunct classes or attributes, except modifications of the **isDefunct** attribute to reactivate the deactivated class or attribute.

- Validation checks that are performed when you add a new class or attribute or when you modify an existing nondefunct class or attribute treat **salesUser** as nonexistent. For example, if **hasPager** is a defunct attribute, trying to modify an existing nondefunct class by adding mayContain=hasPager fails because the validation checks that are performed at schema modification time fail as if **hasPager** does not exist. Or, in the case of a defunct class **salesUser**, trying to add a new class with subClassOf=salesUser fails, because **salesUser** is treated as nonexistent by the validation checks that are performed during the addition of the class.

There is an exception to the rule of defunct classes being treated as nonexistent in Windows 2000 Active Directory. When you try to add or modify a class or attribute so that it has the same distinguished name; object identifier (**attributeID**, **governsID**); **lDAPDisplayName**; **mAPIID**; or **schemaIDGUID** as the defunct class or attribute, the operation fails. In these cases, the class or attribute is treated as an active schema object from the standpoint of schema consistency checks during schema update operations. In Windows Server 2003 or later, this behavior has changed. Defunct objects are still left in the directory. However, it is possible to create a new object that uses the same identification attribute values (that is, **attributeID**, **governsID**, **lDAPDisplayName**, **mAPIID**, or **schemaIDGUID**) as a defunct schema object, as long as the new object's distinguished name is unique. This makes it possible to deactivate a schema object and then create an entirely new schema object — reusing the same values in its definition as the deactivated object — as if the old object were actually deleted. For example, you can change the canonical name (CN) of a defunct object so that you can reuse its old name in the definition of a new schema object. This is very valuable for developers who must constantly modify various classes and attributes for testing purposes. It is possible now to mark an old version of an object definition as defunct and install a new version of the object that uses the same identifiers, making it much easier to develop and test applications that interact with the directory.

This ability to make schema objects defunct can be very useful in different ways in production environments. You can clean up schema objects that are no longer needed by making them defunct. Defunct schema objects are not visible in Active Directory Schema unless you specifically select them on the **View** menu. Then, you can delete existing instances of those classes or attributes if you want to. At the same time, if you need a schema object later, you can reactivate it by modifying the **isDefunct** attribute on the object. This also protects against the accidental removal of a schema object by making it defunct. Making a schema object defunct can be reversed easily with no side effects. Note that because Active Directory does not do any cleanup of data instances after a schema object is made defunct, all instances of the schema object that are made defunct by mistake remain and become valid data when the defunct schema object is made active again.

**Reactivating schema objects**

You can reactivate a defunct schema object either by removing the **isDefunct** attribute from the object or by setting the value of the **isDefunct** attribute to FALSE. You can perform both operations by using either Active Directory Schema or ADSI Edit. Because reactivating a defunct schema object requires schema updates that are similar to adding a new schema object, Active Directory performs the same validation checks as it does when you add a new schema object.

The following validation checks are performed when a defunct schema object is reactivated:

- A defunct class cannot be reactivated unless the attributes that are referenced in its **mustContain**, **systemMustContain**, **mayContain**, **systemMayContain**, or **rdnAttId** attributes — as well as the classes that are referenced in its **subClassOf**, **auxiliaryClass**, and **possibleSuperior** attributes — are already active.

- A defunct class cannot be reactivated if the value of any of the attributes **lDAPDisplayName**, **governsID**, or **schemaIDGUID** is the same as the value that is stored in an already active class or attribute.

- A defunct attribute cannot be reactivated if the value of any of the attributes **lDAPDisplayName**, **attributeID**, **schemaIDGUID**, **mAPIID** is already in use by an active class or attribute.

A defunct schema object can be reactivated at any time. The only restriction is that the **isDefunct** attribute should be the only attribute that is present in the modify call. This helps prevent any ambiguity problems while schema consistency checks are performed.

**Updating the Schema Cache**

When changes are made to Active Directory, they are validated against the schema, which can affect domain controller performance.

The schema is stored in the directory database. Because all changes are validated against the schema, they result in queries of the schema in the directory database, which can increase the workload on a domain controller. To make the operation more efficient, domain controllers cache the schema in memory. Anytime the schema is updated, the schema cache is also updated automatically.

After the domain controller is started, the schema cache is loaded from the schema information in the underlying database and updated automatically whenever the schema is updated. When changes are made, the schema cache is updated automatically within five minutes after the first change is applied. During the interval before the schema updates are copied to the schema cache, objects that reference a new or modified class or attribute cannot be added. If another domain controller attempts to use the new schema modifications, a replication interval must pass before the change becomes available. This behavior keeps the cache consistent, but it can be confusing because changes are not apparent until the cache is updated, even though they are applied to the directory database. When a change is replicated to a domain controller from a replication partner, the cache is updated immediately, without the five-minute delay.

You can update the schema cache by adding the **schemaUpdateNow** attribute to **rootDSE** with a value of 1. The value is not used; it acts as a trigger or operational attribute. You can write this attribute to start a cache reload.

Adding the **schemaUpdateNow** attribute causes a schema cache update to start immediately. The update operation is "blocking," which means that no other modifications can be made until this operation is complete. This prevents multiple changes from being made to the schema simultaneously. If the operation finishes with no errors, the cache is updated and all schema updates are ready to be used. The return of an error, however, indicates that the cache update is not successful. Applications that add the **schemaUpdateNow** attribute should be designed to accommodate the blocking write, particularly to provide feedback if the program or script runs interactively.

You can also update the schema cache immediately by using Active Directory Schema: in the console tree, right-click **Active Directory Schema** and then click **Reload the Schema**.

**Note**

- If you must make multiple changes to the schema, complete all changes before forcing an immediate schema cache update, rather than forcing an update after each change. Multiple cache loads can result in increased workload on the server.

#### Default Security Settings for Active Directory Objects

The default security descriptor for an Active Directory object is specified in the schema. When a new object is created, Active Directory configures the default access rights for that new object. The security descriptor contains the settings that are used to configure the default access rights, and the security descriptor is stored in the schema as part of that object types definition.

**Note**

- There are special cases in which default security is not applied on newly created objects. For more information about access rights, see "Active Directory Schema" in the Microsoft Platform SDK on MSDN.

#### Default security settings for the domain directory partition

The domain directory partition object is derived from the object class **domainDNS**. Therefore, the default security of the domain directory partition object is equivalent to the default security for **domainDNS**.

The default security descriptor for the domain directory partition includes the following permissions:

- Full Control to the Domain Admins group and the System group and Read to the Authenticated Users group.
- Read on all properties to the Everyone group. This permission provides backward compatibility for application programming interfaces (APIs).
- Replicating Directory Changes, Replication Synchronize, and Manage Replication Topology permissions to the Enterprise Domain Controllers group. These permissions allow members of the Enterprise Domain Controllers group to manage replication automatically.
- Replicating Directory Changes, Replication Synchronize, and Manage Replication Topology permissions to the Builtin Administrators group. Administrators of individual domain controllers can use these permissions to troubleshoot replication problems.
- Inheritable Full Control to the Enterprise Admins group. Enterprise Admins, by definition, have complete control of each domain.
- Inheritable List Contents to the Pre–Windows 2000 Compatible Access group.
- Inheritable Read Property on Remote Access Service (RAS) Information, General Information, Membership, User Account Restrictions, and User Logon on all User Objects permissions to the Pre–Windows 2000 Compatible Access group.
- Inheritable Read on all Group objects.
- Inheritable Auditing successful/failed writes to the Everyone group. Activating the auditing policy ensures that writes that are performed on any object in the directory are audited immediately, without the need for extra user intervention.

Inheritable access control entries (ACEs) provide a convenient way to remove auditing policy.

**Default security settings for the configuration directory partition**
The default security descriptor for the configuration directory partition includes the following permissions:

- Full Control to the Domain Admins group and System and Read to the Authenticated Users group.
- Replicating Directory Changes, Replication Synchronize, and Manage Replication Topology permissions to the Enterprise Domain Controllers group. These permissions enable domain controllers in the forest to replicate from each other and automatically reconfigure the replication topology on the basis of replication delays and latency for the configuration directory partition.
- Replicating Directory Changes, Replication Synchronize, and Manage Replication Topology permissions to the Builtin Administrators group. These permissions enable administrators from individual domain controllers to synchronize replication and topology management for the configuration directory partition.
- Enable Inheritable Full Control to the Enterprise Admins group. This permission allows members of the Enterprise Admins group exclusive control over the Configuration container. The Enable Inheritable Full Control permission is required to control the Configuration container throughout the forest.
- Enable Inheritable Auditing to the Writes by the Everyone group. Activating the auditing policy ensures that writes that are performed on any object in the directory are audited immediately, without the need for extra user intervention. Inheritable ACEs provide a convenient way to remove auditing policy.

**Default security settings for the schema directory partition**
The default security descriptor for the schema directory partition includes the following permissions:

- Write access to the schema head to the Schema Admins group.
- Change Schema Master control permission to the Schema Admins group. This permission enables members of the Schema Admins group to change which domain controller holds the schema operations master role.
- Inheritable Full Control permission to the Schema Admins group. By default, the Schema Admins group is the only group that has Write access to the entire schema head. A schema object does not have any exclusive control over its own security; therefore, the object inherits its security from the schema head.
- Replicating Directory Changes, Replication Synchronize, and Manage Replication Topology to the Enterprise Domain Controllers group. These permissions enable the members of the Enterprise Domain Controllers group to manage replication of the schema in the forest automatically.
- Replicating Directory Changes, Replication Synchronize, and Manage Replication Topology permissions to the Builtin Administrators group. These permissions enable the administrators of domain controllers to resolve replication issues.
- Read permissions to the Authenticated Users group. This permission gives the members of the Authenticated Users group the right to read the schema.
- Audit successful/failed writes to the Everyone group. Activating the auditing policy ensures that writes that are performed on any object in the directory are audited immediately without the need for extra user intervention. Inheritable ACEs provide a convenient way of removing auditing policy.

**Default security settings for attributes and classes**
All attributes and classes inherit security from the ACLs on the schema head. This ensures that security for the entire schema is consistent.

### Note

- The default security settings allow Write access to the schema head only to the Schema Admins group.

**Issues Related to Modifying the Schema**
Three main issues relate to modifying the schema:

- The impact on replication
- Managing concurrent schema modification operations
- Handling invalid object instances

**Impact on replication**
Because the schema is replicated across all domain controllers in the forest, a schema update that is performed at one domain controller is propagated throughout the forest. This guarantees that the schema is consistent across the forest. However, because of replication latencies, there can be temporary inconsistencies.

Active Directory solves this problem by explicitly replicating the schema head from the originating server when failures occur. Additionally, the replication of the schema head triggers an immediate schema cache update on the target server. Active Directory then replicates the failed object again.

**Managing concurrent schema modification operations**
Programs that modify the schema should not be run concurrently unless the programs include provisions to check that schema modifications that are made by one program will not conflict with schema modifications that are made by the other programs. Active Directory must ensure that different program threads do not perform simultaneous, conflicting schema updates, such as one thread deleting an attribute when another thread is adding the attribute to the **mayContain** list of a class.

To ensure that different program threads do not perform simultaneous, conflicting schema updates, any thread that attempts to perform a schema update also writes a special attribute on the schema head automatically as part of the transaction. Active Directory automatically causes the thread to write the attribute so that you do not have to do so in your program code. Only one thread can write this attribute at any one time. This method guarantees schema consistency, but it does not guarantee which of the updates is successful. This is an important consideration when schema updates are made in a batch, such as during the installation of directory-enabled applications.

For example, consider a scenario in which two programs, program A and program B, are being installed simultaneously. Both programs are designed to create several new schema objects. Because Active Directory creates one thread per object update, it is possible that *some* of the objects in program A and *some* of the objects in program B are created (if the internal threads do not overlap). Then, one of the installations fails because a thread for a schema object creation for program A overlaps with a thread for a schema object creation for program B.

Assume that program A fails because of the thread conflict and the resulting failure to create some of the needed schema objects. Attempting to reinstall program A again does not work because some of the objects that program A created are already in the schema, and trying to create them again in the second run returns an error.

**Handling invalid object instances**
Schema updates can make an existing instance of an object invalid. For example, suppose a user account object, JohnDoe, is an instance of **classSalesEmployee**. The **classSalesEmployee** class has an attribute, **POBox**, in its **mayContain** list. Therefore, because JohnDoe is an instance of **classSalesEmployee**, JohnDoe can have the **POBox** attribute defined in it. Assume that JohnDoe does have this attribute currently defined in it. A schema update is performed that modifies **classSalesEmployee** by deactivating the attribute **POBox** from its **mayContain** list. Note that this change makes the instance of JohnDoe invalid because JohnDoe now has an attribute, **POBox**, that it is not allowed to have according to the class definition of **classSalesEmployee** (of which JohnDoe is an instance). Active Directory allows the now-invalid objects to remain in the directory and ensures that they do not cause problems in the rest of the schema. Active Directory does not automatically clean up invalid objects, but invalid objects and attributes appear in searches and can be deactivated manually.

**Interoperating with Other Directories**
To provide more support for industry standards in Active Directory, the schema supports the **inetOrgPerson** object class and the use of a standard file format for importing and exporting schema modifications.

**inetOrgPerson Object Class**
The Active Directory schema in Windows Server 2003 supports the **inetOrgPerson** object class. The **inetOrgPerson** object class and its associated attributes are defined in RFC 2798. The **inetOrgPerson** object class is used in several non-Microsoft LDAP and X.500 directory services to represent users in the directory.

Support for inetOrgPerson makes migrations from other LDAP directories to Active Directory more efficient. The

**inetOrgPerson** class was added by adding new attributes to the schema and deriving inetOrgPerson from the **user** class and the new attributes. As a result, inetOrgPerson can be used as a security principal.

**inetOrgPerson in Windows Server 2003 and Later**
To help organizations migrate their applications and directory data to Active Directory, the **inetOrgPerson** class has been added to the base schema.

The definition of the **inetOrgPerson** class is based on RFC 2798. However, some issues have caused the Active Directory definition of this class to differ from the RFC definition, as follows:

- The **inetOrgPerson** class is derived from the **user** class, not the **organizationalPerson** class. This makes it possible for inetOrgPerson to operate as a security principal. Accounts that are created as inetOrgPerson objects can log in to Active Directory.

- Attributes that are already defined in the base schema are not changed.

  Several attributes in the base schema are defined as single-value instead of multivalue, or the reverse, as defined in the RFC 2798. The object identifier of some attributes differs from the definition in the RFC. Changing the definition of these attributes in the base schema to match the RFC causes backward compatibility problems with the Windows 2000 Active Directory schema and may result in application problems. For example, Microsoft Outlook will not display the data for attributes that have an unexpected value for the **isSingleValued** attribute. Active Directory management tools may treat the attributes incorrectly and inadvertently overwrite data or fail to update them at all.

- The **inetOrgPerson** class inherits two mandatory attributes:

  - The **cn** attribute is used to name a new instance of the class. The **samAccountName** attribute is the user's login ID.

  - The **objectCategory** attribute for inetOrgPerson is set to **Person**. This has the effect of returning all user, computer, and inetOrgPerson objects when the filter in a query is objectCategory=Person.

**Importing and Exporting Directory Information**
Active Directory supports the use of files that are formatted with LDAP Data Interchange Format (LDIF) for importing and exporting information in the directory. This includes information that is stored in the schema, such as schema modifications.

LDIF is an industry-standard, ASCII, text-based file format that can be used to define data in a text file that is used to import information, such as schema modifications, into LDAP-based directories. LDIF is defined in RFC 2849.

After an LDIF file is created, a tool such as Ldifde.exe, which is available in Windows Server 2003 or later, performs the import operation using the data file for input. Ldifde.exe can also be used to export data from the directory. Exported data is also stored using the LDIF file format to make it easier to import the data into another LDAP-based directory. For more information about LDAP, see "Data Store Technical Reference."

**Related Information**
The following resources contain additional information that is relevant to this section.

- RFC 2251 in the IETF RFC Database

- Microsoft Platform SDK on MSDN for more information about mandatory and optional attributes, modifying the schema, and access rights (in Active Directory Schema)

- Data Store Technical Reference

# Active Directory Schema Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Active Directory Schema Tools and Settings**

**In this section**


- Active Directory Schema Tools
- Related Information


When existing class and attribute definitions in the Active Directory schema do not meet the needs of your organization, you can use schema-based administrative tools to modify or add schema objects. You can modify an existing attribute or add a new class or attribute to the schema to store a new type of information in the directory. The process of modifying or updating the schema is often referred to as "extending the schema." In addition to using schema tools to extend the schema, you can perform most schema extensions by using customized applications or Active Directory Service Interfaces (ADSI) scripts.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

**Note**

- Extending the schema is a major change with implications for the entire directory. Extend the schema only when it is absolutely necessary. Many schema modifications cannot be reversed; therefore, you must thoroughly plan and test changes in an isolated environment before you deploy them in your production forest.

This section contains information about the tools that are associated with the Active Directory schema.

**Active Directory Schema Tools**

Normally, you do not interact directly with the schema on a daily basis. Active Directory uses the schema to create objects that are stored in the directory. You interact with those objects, not with the schema. You interact directly with the schema when you make modifications to the schema by adding definitions to it or by modifying existing definitions.

Only members of the Schema Admins group can make changes to the schema. The two most common scenarios for modifying the schema are as follows:

- You install an application that adds customized object definitions so that it can store information in the directory; for example, you install an e-mail program that stores user e-mail names in the directory.
- You test the development of applications that use the directory for data storage. In this scenario, you add customized object definitions to the schema and modify them throughout their lifetimes as the development process proceeds.

**Note**

- Changes to the schema must be written only on the schema master. Although all domain controllers have a copy of the schema in their Active Directory database, only the domain controller that holds the schema operations master role (also known as flexible single master operations (FSMO)) is allowed to write changes to the schema.

The following tools are associated with the Active Directory schema.

**Adsiedit.exe: ADSI Edit**

**Category**

ADSI Edit is included when you install Support Tools for Windows Server 2003 and later.

**Version Compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

ADSI Edit is a Microsoft Management Console (MMC) snap-in that uses ADSI, which uses the Lightweight Directory Access Protocol (LDAP). You can use ADSI Edit to view and modify directory objects in the Active Directory database. You can also use it to view schema directory partition objects and properties. When you open ADSI Edit, the Schema container is displayed by default. You can expand the container to view schema classes and attributes.

To find more information about ADSI Edit, see "Support Tools Help" in Tools and Settings Collection.

**Csvde.exe: Csvde**

**Category**
Csvde is a command-line tool that ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

The comma-separated value (CSV) file format is a simple format whose primary benefit is ease of use. In the CSV file format, each line represents a discrete object in the directory, and the object's attributes are separated by commas. The first line of the file always contains all of the attribute names. Each subsequent line represents a different entry in the directory. Values for multivalue attributes can also be specified, and they are delimited by semicolons (;).

Because this format is compatible with the Microsoft Excel CSV format, you can use Csvde.exe to export directory information to an Excel spreadsheet or to import data from a spreadsheet into Active Directory. You can use this format only for additions to the directory. Csvde.exe cannot be used to modify or delete objects. Csvde.exe also supports batch operations that are based on CSV.

The parameters that are used for the Csvde.exe tool are the same as the parameters that are used for the Ldifde.exe tool. However, unlike Ldifde.exe, Csvde.exe can export data from Active Directory into files that can be read by certain applications. For example, if you want to view all Active Directory users in an Excel report, you can use Csvde.exe to export the directory data into the CSV file format, which you can then read in Excel.

To find more information about Csvde.exe, see "Command-Line References" in Tools and Settings Collection.

**Dsa.msc: Active Directory Users and Computers**

**Category**
Active Directory Users and Computers is an MMC snap-in in Administrative Tools that is installed automatically on all domain controllers running Windows Server 2003.

### Version compatibility

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Active Directory Users and Computers is a graphical user interface (GUI) tool that you can use to manage users and computers in Active Directory domains. To modify the schema, you must use an account that is a member of the Schema Admins group. By default, the only member in the Schema Admins group is the Administrator account in the root domain of the enterprise. You must explicitly add other accounts.

You can use Active Directory Users and Computers to verify that an account is a member of the Schema Admins group. Restrict membership in the Schema Admins group to prevent unauthorized access to the schema. Improper modification of the schema can have serious consequences.

By default, only members of the Schema Admins group have permission to write to the schema. You can assign explicit permissions to use the Active Directory Schema snap-in to specific users; however, this is not recommended.

**Ldifde.exe: Ldifde**

**Category**
Ldifde is a command-line tool that ships with Windows Server 2003.

### Version compatibility

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Active Directory supports the use of files that are formatted with the LDAP Data Interchange Format (LDIF) for importing and exporting information in the directory. This includes information that is stored in the schema, such as schema modifications. After an LDIF file is created, a tool such as Ldifde.exe performs the import operation by using the LDIF file for input. You can also use Ldifde.exe to add, modify, and delete directory objects; export Active Directory user and group information to other applications or services; and populate Active Directory with data from other directory services.

To find more information about Ldifde.exe, see "Command-Line References" in Tools and Settings Collection.

**Ntdsutil.exe: Ntdsutil**

**Category**
Ntdsutil is a command-line tool that ships with Windows Server 2003.

**Version Compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition |

Ntdsutil.exe provides advanced management capabilities for Active Directory. For the Active Directory schema, you can use Ntdsutil.exe to identify, transfer, or seize the schema operations master role. This tool is intended for use by experienced administrators.

To find more information about Ntdsutil, see "Command-Line References" in Tools and Settings Collection.

**Schmmgmt.msc: The Active Directory Schema snap-in**

**Category**
The Active Directory Schema snap-in is an MMC snap-in in Administrative Tools that is installed automatically on all domain controllers running Windows Server 2003. However, you must register it manually before you use it for the first time.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
| --- | --- |
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

The Active Directory Schema snap-in is a GUI tool that members of the Schema Admins group can use to manage Active Directory objects and their associated attributes. You can use this tool to create and modify classes and attributes. You can also use it to specify what attributes are indexed and what attributes are replicated to the global catalog.

The Active Directory Schema snap-in is not one of the default MMC snap-ins that is provided with Windows Server 2003. To make it appear in the list of available snap-ins, install the Windows Server 2003 Administration Tools Pack (Adminpak.msi). To register the Active Directory Schema snap-in, run **Regsvr32 Schmmgmt.dll** from the command prompt or from the **Run** command on the **Start** menu.

**ADSI and Visual Basic Scripts**

Active Directory provides a set of interfaces that you can use programmatically to gain access to directory objects, including schema objects. ADSI conforms to the Component Object Model (COM), and it supports standard COM features. ADSI defines a directory service model and a set of COM interfaces that you can easily use with a variety of programming languages. With Microsoft Visual Basic, Scripting Edition and ADSI, you can write scripts to modify the directory in various ways, including extending the schema.

For more information about using ADSI and scripting to modify the schema, see Using Active Directory Service Interfaces in the Microsoft Platform SDK on MSDN.

**Related Information**

The following resources contain additional information that is relevant to this section:

- "Support Tools Help" in Tools and Settings Collection for information about ADSI Edit
- "Command-Line References" in Tools and Settings Collection for information about Csvde, Ldifde, and Ntdsutil
- Microsoft Platform SDK on MSDN for information about using ADSI and scripting to modify the schema (in Using Active Directory Service Interfaces)

# Data Store Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

In the Active Directory directory service, the data store contains database files and processes that store and manage directory information for users, services, and applications. A copy of the data store runs on each domain controller in the forest. The data store provides access to directory information through a process called the Directory System Agent (DSA). The data store holds directory information in a database file called Ntds.dit. The Active Directory data store is often referred to as the directory.

**In this subject**


- What Is the Data Store?
- How the Data Store Works
- Data Store Tools and Settings

# What Is the Data Store?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- The Business Need
- The Data Store Solution
- Data Store Scenarios
- Directories and Databases
- Data Store Dependencies
- Related Information


In Active Directory, the data store contains database files and processes that store and manage directory information for users, services, and applications. A copy of the data store runs on each domain controller in the forest. The Active Directory data store is often referred to as the directory.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

**The Business Need**

Organizations generate and store data about many different kinds of entities (computers, users, services, and so on) that exist in their information technology (IT) infrastructure. Historically, organizations maintain this information in multiple locations and formats. For example, an organization might store employee information in a human resources database, application data in an application-specific database, Quality of Service (QoS) rules on routers, and security information (user name and password) on a network server. Storing, managing, and making available many disparate data sets is expensive and time consuming.

**The Data Store Solution**

The Active Directory data store solves the problem of managing disparate data sets across an organization. The data store is a single, general-purpose database that can hold many types of data and distribute that data to users anywhere on the network. The following figure illustrates the use of the Active Directory data store as a single, general-purpose database for all of an organization's data.

**Data Store**




The Active Directory data store achieves the following major design goals:

- To provide distributed access to the directory by users and applications, a copy of the data store resides on all domain controllers in a domain or forest.
- To provide standard interfaces for accessing data, the data store includes a Lightweight Directory Access Protocol (LDAP) interface and a Messaging API (MAPI) interface.
- To support quick searches of directory data, the data store includes efficient query and index mechanisms.
- To ensure scalability, the data store uses a hierarchical, or tree, model that supports partitioning. The data store also ensures scalability by automatically managing database growth and increasing the database file size when necessary.
- To support both full and incremental restoration of data, the data store uses a transactional model, logging uncommitted transactions in log files.
- To ensure data consistency, the data store enforces a set of extensible data type and format constraints, called the schema.

**Data Store Scenarios**
Active Directory is required for the deployment and operation of Windows Server 2003 and later domains and forests. The data store is a key and required part of any Active Directory deployment. Therefore, the data store exists in any Windows Server 2003 and later domain or forest. Use of the data store can be further delineated into two main scenarios, based on the type of data that is stored. These scenarios are the network operating system (NOS) directory scenario and the NOS and application directory scenario.

**NOS Directory Scenario**
In the NOS directory scenario, an organization uses the data store only for storage of NOS-related information, including information about users, computers, services, printers, and so on. This scenario applies to organizations that are not running directory-enabled applications (applications that can communicate with and store data in a directory data store) or to organizations that are using a different directory data store than Active Directory for storing application data. In this scenario, you can manage the data in the data store by using the command-line and graphical user interface (GUI) tools that are provided with Windows Server 2003 and later.

**NOS and Application Directory Scenario**
In the NOS and application directory scenario, directory-enabled applications that are not a part of the NOS use the data store to hold application-specific data. Examples of directory-enabled applications include messaging, customer relationship management (CRM), enterprise resource planning (ERP), and document management applications. In this scenario, the schema (the set of rules that define the type and formatting of data) is generally extended to include data definitions that are specific to the directory-enabled applications that use the data store.

**Directories and Databases**
In the context of directory services, questions often arise regarding the differences between directories and databases. The Active Directory data store shares many aspects in common with databases, including the storage of data in rows and columns and the use of a hierarchical data model.

A directory differs from a database primarily in its intended use. A directory is optimized for read operations, while a database is optimized for write and change operations. Therefore, any data that is read many more times than it is written or modified is a good candidate for storage in a directory.

A directory also typically differs from a database in the protocols that it uses to access information in the data store. A directory, such as Active Directory, is most often accessed with LDAP. A database is most often accessed with an interface such as Structured Query Language (SQL).

**Data Store Dependencies**
The Active Directory data store depends on several related technologies and resources for its proper functioning. The following sections describe these technologies and resources.

**NTFS**
The data store requires a disk volume that is formatted with the NTFS file system. NTFS is more powerful than FAT16 or FAT32, and NTFS includes features that are required for hosting Active Directory. For more information about NTFS, see "NTFS Technical Reference."

**Network Connectivity**

So that users and computers can access the data store, and to support replication of the data store between domain controllers, network connectivity with TCP/IP is required. For more information about networking and TCP/IP, see "TCP/IP Technical Reference."

**Active Directory Replication**

Except for very small networks, directory data must reside in more than one place on the network to be equally useful to all users. Through the process of replication, Active Directory maintains copies, or replicas, of the data store on each domain controller in the forest, which helps to ensure data availability and performance for all users. For more information about replication, see "Active Directory Replication Model Technical Reference."

**FRS**

The File Replication service (FRS) is a process that is associated with the Distributed File System (DFS). FRS is used by Active Directory to replicate the system volume (SYSVOL) between domain controllers. SYSVOL holds certain kinds of Active Directory information, including Group Policy objects (GPOs) and scripts. For more information about FRS, see "FRS Technical Reference."

**Disk Space**

There are no practical limits to the number of objects that can be stored in the Active Directory data store. The Active Directory directory database has been tested for up to 40 million objects. Performance tests show logon performance for a single LDAP client to be the same with 10,000 objects, 100,000 objects, and 1 million objects; that is, the directory service does not slow measurably when the size of the database increases.

The minimum free disk space requirements for the directory database and log files depend on the size of the database. On partitions that hold the database only (and not the database log files), free disk space must not fall below the greater of 20 percent of the disk space that is consumed by the database or 500 megabytes (MB). On disk volumes that hold the database and the database log files, free disk space must not fall below the greater of 20 percent of the combined size of the Ntds.dit and log files or 1 gigabyte (GB). For more information about the Ntds.dit and log files, see "How the Data Store Works."

**Related Information**

The following resources contain additional information that is relevant to this section:

- NTFS Technical Reference
- TCP/IP Technical Reference
- Active Directory Replication Model Technical Reference
- FRS Technical Reference
- How the Data Store Works

# How the Data Store Works

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Data Store Architecture
- Data Store Protocols
- Data Store Interfaces
- Data Store Logical Structure
- Data Store Physical Structure
- Data Store Processes and Interactions
- Network Ports Used by the Data Store
- Related Information


In Active Directory, the data store contains database files and processes that store and manage directory information for users, services, and applications. A copy of the data store runs on each domain controller in the forest. The Active Directory data store is often referred to as the directory.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

The ideal environment for the data store includes the following:

- A domain controller running an operating system in the Windows Server 2003 or later family and containing hardware that meets the minimum hardware requirements of the edition of the operating system.
- For environments consisting of multiple domain controllers, the presence of a fully functioning Active Directory replication topology
- For environments consisting of multiple domain controllers, the presence of a fully functioning File Replication Service (FRS) topology
- A regular backup schedule
- Regular monitoring of Active Directory, either through manual review of event logs or through an automated monitoring solution, such as Microsoft Operations Manager (MOM)

This section describes the elements of the Active Directory data store, including its architecture, protocols, interfaces, logical structure, physical structure, processes and interactions, and network ports.

**Data Store Architecture**

The Active Directory data store consists of several components that together provide directory services to directory clients and to other directory servers. These components include three service components, four interfaces, and the directory database

where data is actually stored.

The following figure illustrates the architecture of the data store.

**Data Store Architecture**




The following table describes the components of the data store.

**Data Store Components**

| Component | Description |
|---|---|
| Interfaces: Lightweight Directory Access Protocol (LDAP), replication (REPL) and domain controller management interface, Messaging API (MAPI), Security Accounts Manager (SAM)) | The data store interfaces provide a way for directory clients and other directory servers to communicate with the data store. |
| Directory System Agent (DSA) (Ntdsa.dll) | The DSA, which runs as Ntdsa.dll on each domain controller, provides the interfaces through which directory clients and other directory servers gain access to the directory database. In addition, the DSA enforces directory semantics, maintains the schema, guarantees object identity, and enforces data types on attributes. |
| Database layer | The database layer is an application programming interface (API) that resides in Ntdsa.dll and provides an interface between applications and the directory database to protect the database from direct interaction with applications. Calls from applications are never made directly to the database; they go through the database layer. In addition, because the directory database is flat, with no hierarchical namespace, the database layer provides the database with the abstraction of an object hierarchy. |
| Extensible Storage Engine (ESE) (Esent.dll) | The ESE, which runs as Esent.dll, manages the tables of records — each with one or more columns — that make up the directory database. |
| Database files | The data store stores directory information in a single database file called Ntds.dit. In addition, the data store also uses log files, to which it temporarily writes uncommitted transactions. |

The DSA, database layer, and ESE are described in the following sections. For information about the interfaces, see "Data Store Interfaces" later in this section. For information about the database files, see "Data Store Physical Structure" later in this section.

**DSA**

The DSA runs on every domain controller as Ntdsa.dll, and it provides access to the directory database. The DSA runs as part of the Local Security Authority (LSA) process (Lsass.exe), which manages authentication packages and authenticates users and services. Running in Lsass.exe enables Active Directory to securely manage sensitive information, such as account passwords.

Clients can use one of the supported interfaces to connect (bind) to the DSA and then search for, read, and write to Active Directory objects and their attributes.

A forest-wide object in the directory, the NTDS Settings object (class **nTDSDSA**), represents the DSA on a domain controller, and this object contains configuration information about the DSA on that domain controller.

In addition to providing the interfaces through which directory clients gain access to directory data, the DSA provides the following functionality.

**Object identification**
Every object in Active Directory has a permanent globally unique identifier (GUID), which is associated with several string forms of the object name (SAMAccountName, user principal name, and distinguished name), as well as a security identifier (SID). The object names and the SID are not permanent; that is, they can be changed. All permanent references to the object are kept in terms of the GUID. The object name is used for hierarchy navigation and display purposes, and the SID is used for access control. The DSA maintains the GUID association with an object when the object's string name or SID changes, for example, when the object is moved to a different folder (the string name changes) or when the object is moved to a different domain (the string name and the SID change).

**Schema enforcement**
The DSA ensures that data in the directory adheres to the data definitions that are provided by the directory schema. The schema is the set of rules that determines what kind of data the directory can hold.

### Note

- In forests, if an update does not produce a conflict with the schema at the originating replica, the update is considered acceptable at all replicas. Therefore, replicated updates do not perform schema checks, and you do not have to wait until the schema replicates before creating instances of a new object or attribute. For more information about replication, see "Active Directory Replication Model Technical Reference".

**Access control enforcement**
The DSA enforces security limitations in the directory by reading SIDs on the access token.

**Support for replication**
The API that is called to initiate replication is implemented in the DSA.

**Referrals**
The DSA manages the directory hierarchy information (referred to as knowledge), which it receives from the database layer. The DSA is responsible for cross-references of Active Directory domain objects.

**Functional levels**
Beginning with Windows Server 2003 and later domain controllers, the DSA has a built-in numeric value that identifies its operating system version — and therefore its functional capabilities — to other services. Services that rely on the DSA can use this numeric value to determine which of their service features to enable.

**DSA GUID and Invocation ID**
Both the DSA and the Active Directory database are represented uniquely and have their own respective GUIDS. The DSA GUID is the GUID of the NTDS Settings object (class **nTDSDSA**). The value of the DSA GUID is stored in the **objectGUID** attribute of the NTDS Settings object of the given domain controller server object, which resides in the Sites container in the configuration directory partition.

Domain controllers use the DSA GUID to locate replication partners. Replication uses a special Domain Name System (DNS) name that contains the DSA GUID. This GUID-based DNS name is an alias for the local computer name. The Net Logon service registers this alias resource record in DNS in the form of the CNAME (canonical name) resource record.

The DSA GUID is created when the domain controller is initially promoted, that is, when Active Directory is installed. The DSA GUID is destroyed only if Active Directory is removed from the domain controller. The DSA GUID ensures that the DSA remains recognizable when a domain controller is renamed. The DSA GUID is also not affected by the Active Directory backup and restore process.

The Active Directory database has its own GUID, which the DSA uses to identify the database instance (the version of the database). The database GUID is stored in the **invocationId** attribute on the nTDSDSA object. Unlike the DSA GUID, which never changes for the lifetime of the domain controller, the invocation ID changes during the Active Directory restore process to ensure the consistency of the replication process. On domain controllers running Windows Server 2003 or later, the invocation ID also changes when an application directory partition is removed from or added to the domain controller.

**Database Layer**
The database layer is an API that resides in Ntdsa.dll and provides an object view of the directory database, making the data

accessible to the DSA as a set of hierarchical containers. By applying schema semantics to database records, the database layer isolates the upper components of the directory service from the underlying database system. The database layer is an internal interface. No database access calls are made directly to the ESE; instead, all database access is routed through the database layer.

In the directory database, each object is identified by its relative distinguished name, which is unique in the object's parent container. The relative distinguished name and the chain of successive parent object names make up the object's distinguished name. The database stores the relative distinguished name for each object, as well as a reference to the parent object. The database layer follows these parent references and concatenates the successive relative distinguished names to form distinguished names, thereby defining the object hierarchy.

The database layer is also responsible for the creation, retrieval, and deletion of individual records (objects), attributes within records, and values within attributes. To carry out these functions, the database layer uses the schema cache (an in-memory structure in the DSA) to get the information about the attributes that it needs.

**ESE**

The Extensible Storage Engine (ESE) is a Windows component that is used by Active Directory, as well as by several other Windows components, as an interface to the data that is stored in an indexed and sequential access method (ISAM) database. (The Active Directory database is an ISAM database.) The ESE is responsible for indexing the data in the database file and for transferring the data in and out of the database. Its purpose is to enable applications to store and retrieve data by using the ISAM. The ESE provides applications with a consistent data state by means of transacted data update and retrieval. A crash recovery mechanism maintains data consistency, even in the event of a system crash. Transactions in the ESE are highly concurrent, making the ESE suitable for server applications. The ESE caches data intelligently to ensure high-performance access to data. In addition, the ESE is resource efficient, making it suitable for applications that perform auxiliary roles.

The version of the ESE that runs on domain controllers is implemented in Esent.dll.

The following characteristics of the ESE make it well suited to the storage needs of Active Directory. The ESE:

- Supports databases of up to 16 terabytes (TB) in size, and it can hold many millions of objects per domain.
- Supports indexing.
- Supports multivalued attributes.
- Supports update operations that are transacted for stability and integrity across system failures.
- Can be backed up while the domain controller is online.
- Handles sparsely populated objects well; that is, space in the database is not reserved for attributes that do not have values.

**Note**

- The encrypting file system (EFS) enables users to encrypt individual files, folders, or entire data drives. Because EFS initialization and domain controller startup occur in parallel, EFS recovery operations can interfere with the ability of the ESE to start Active Directory and cause domain controller startup to fail.

The Active Directory schema defines all the attributes that are required and allowed for a given object. However, the ESE reserves storage only for the space that is used — that is, only for the attributes that have values, not for all possible attributes. For example, if a user object has 50 attributes defined in the schema and you create a user with values for only 4 attributes, storage space is allocated only for those 4 attributes. If more attributes are added later, more storage is allocated for them.

The ESE implements the search and retrieval functionality of the underlying database. Also, the ESE is able to store attributes that can have multiple values. For example, the database can store multiple phone numbers for a single user without requiring a different phone number attribute for each phone number.

ESE provides transactional views of the database. The cost of providing these views is that any object that is modified in a transaction has to be temporarily copied so that two views of the object can be provided: one to the thread inside that transaction and one to threads in other transactions. This copy must remain as long as any two transactions in the process have different views of the object. The repository that holds these temporary copies is called the version store. Because the

version store requires contiguous virtual address space, it has a size limit. If a transaction is open for a long time while changes are being made (either in that transaction or in others), eventually the version store can be exhausted. At this point, no further database updates are possible.

**Note**

- The percentage of version store space that is available for processing has been significantly increased in Windows Server 2003 and later.

**Data Store Protocols**

The primary protocol that is used by the Active Directory data store is Lightweight Directory Access Protocol (LDAP), which runs on top of TCP/IP. In addition, the data store uses remote procedure call (RPC) for MAPI, replication, domain controller management, and SAM-related operations. And, although it is not widely used, the data store also supports the use of Simple Mail Transfer Protocol (SMTP) for replication between domain controllers where "always on" network connections do not exist.

The following figure illustrates the protocols and interfaces that are used by the data store. For more information about the data store protocols, see the table that follows this figure. For more information about the data store interfaces, see "Data Store Interfaces" later in this section.

**Data Store Protocols**




The following table describes the protocols that are used by the data store.

**Data Store Protocols**

| Protocol | Description |
|---|---|

| | |
|---|---|
| LDAP | LDAP is a directory service protocol that specifies directory communications. It runs directly over TCP/IP, and it can also run over User Datagram Protocol (UDP) connectionless transports. Clients can use LDAP to query, create, update, and delete information that is stored in a directory service over a TCP connection through the TCP default port 389. Active Directory supports LDAP v2 (RFC 1777) and LDAP v3 (RFC 2251). LDAP v3 is an industry standard that can be used with any directory service that implements the LDAP protocol. LDAP is the preferred and most common way of interacting with Active Directory.<br><br>Historically, LDAP is a simplified ("lightweight") version of Directory Access Protocol (DAP), which is the original protocol that was used to interact with X.500 directories. X.500 defines an earlier set of standards that was developed by the International Organization for Standardization (ISO). LDAP is simpler than DAP in two key ways:<br><br>• Rather than using its own protocol stack according to the Open Systems Interconnection (OSI) networking model, LDAP communicates over Internet Protocol (IP) by using either UDP or TCP.<br><br>• LDAP syntax is easier to use than DAP syntax.<br><br>For these reasons, LDAP is widely used and accepted as the standard protocol for directory service access.<br><br>The following key aspects characterize LDAP:<br><br>• The protocol is carried directly over TCP for connection-oriented transport (receipt of data is acknowledged) and over UDP for connectionless transport (sent or received data is not acknowledged).<br><br>• Most protocol data elements can be encoded as ordinary strings, for example, as distinguished names.<br><br>• Referrals to other servers can be returned to the client.<br><br>• Simple Authentication and Security Layer (SASL) mechanisms can be used with LDAP to provide associated security services. SASL Digest is supported in Windows Server 2003 and later as the authentication standard for LDAP.<br><br>• Attribute values and distinguished names can be internationalized through the use of the ISO 10646 character set.<br><br>• The protocol can be extended to support new operations, and controls can be used to extend existing operations.<br><br>• The schema is published through an attribute on the directory root object (rootDSE) for use by clients. |
| RPC | The data store uses RPC for replication (REPL) and domain controller management communications, MAPI communications, and SAM-related communications. RPC is a powerful, robust, efficient, and secure interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality residing in a different process. That different process can be on the same computer, on the local area network (LAN), or across the Internet. |
| SMTP | The data store can use Internet-standard SMTP as the protocol for replication communications when a permanent, "always on" network connection does not exist between two domain controllers. SMTP is used to transport and deliver messages based on specifications in Request for Comments (RFC) 821 and RFC 822. SMTP also includes enhancements that build on the basic delivery functions of the protocol. There are SMTP options available that provide control over the routing and delivery of messages and that provide secure communications. |

**Data Store Interfaces**

As shown in the "Data Store Architecture" figure earlier in this section, network clients and other directory servers obtain access to Active Directory through one of the interfaces that are described in the following table.

**Data Store Interfaces**

| Interface | Description |
|---|---|

| | |
|---|---|
| LDAP | LDAP is the primary interface for the data store. Directory clients use LDAP v3 to connect to the DSA through the LDAP interface. The LDAP interface is part of Wldap32.dll. LDAP v3 is backward compatible with LDAP v2. |
| REPL | The REPL management interface provides functionality for finding data about domain controllers, converting the names of network objects between different formats, manipulating service principal names (SPNs) and DSAs, and managing replication of servers. |
| MAPI | Messaging clients gain access to the Microsoft Exchange Server directory service by using MAPI address book providers. For compatibility with existing messaging clients, Active Directory supports the MAPI-RPC address book provider, which provides access to Active Directory, for example, to find the telephone number of a user. |
| SAM | SAM is a proprietary interface for connecting to the DSA on behalf of clients that use Windows NT 4.0 or earlier. These clients use Windows NT 4.0 networking APIs to connect to the DSA through SAM. Replication with Windows NT 4.0 backup domain controllers (BDCs) goes through the SAM interface as well. |

**Note**

- The MAPI interface is provided only for support of legacy Microsoft Outlook clients. Development against the MAPI interface is no longer supported.

**Data Store Logical Structure**

The Active Directory data store uses the LDAP information model as its model for organizing data. In addition, the data store is organized into logical partitions that can each be replicated independently, so that individual domain controllers only need to store data that is pertinent to the domain in which they reside.

**LDAP Information Model**

The LDAP information model is based on the entry, which contains information about some object, for example, a person or computer. In Active Directory, an LDAP entry is referred to as an object. Active Directory objects are composed of attributes, which have a type and one or more values.

The implementation of the information model is called the schema, which is a set of objects that defines the structure and content of every object that can be created in a directory service.

**Directory Tree**

Active Directory organizes objects into a hierarchical tree structure. The directory tree represents the hierarchy of Active Directory objects for a given forest. The structure of the hierarchy is derived from the schema, and the directory service implements the hierarchy. The hierarchy provides the basis both for using names for navigation and for defining the scope of search requests.

Every object in Active Directory has exactly one parent, and a reference to the parent object is stored with the object. As a result of these parent references, the hierarchy of objects that are managed by Active Directory forms a tree structure. The objects that populate the directory create this tree structure according to the rules of the schema. For example, the schema might dictate that a given class of object can be the child of one class but not the child of another class.

The architectural restrictions and requirements in the directory tree are as follows:

- Domain objects, which are containers, can be children only of other domain objects. For example, a domain cannot be the child of an organizational unit (OU).

- The root object of the directory tree is called the DSA-specific Entry (DSE), or rootDSE. The rootDSE object has no hierarchical name or schema class, but it does have a set of attributes that identify the contents of the domain controller on which it resides. Therefore, rootDSE constitutes the root of the directory tree from the perspective of the domain controller to which you are connected.

- Below the rootDSE, every directory has a root domain, which is the first domain that is created in a forest. This domain always has a child container called Configuration, which contains configuration data for the forest.

**rootDSE**

In the LDAP information model, one or more directory servers can jointly provide access to a directory tree. The servers that host directories are known in LDAP terminology as Directory System Agents (DSAs). In Active Directory, the DSAs are always domain controllers.

At the root of a directory tree is a DSE, which is not part of any directory partition. TherootDSE represents the top of the logical namespace for one domain controller. The rootDSE attributes contain information about the directory server, including its capabilities and configuration.

There is only one root for a given directory, but the information that is stored in the root is specific to the domain controller to which you connect. Among other things, the attributes of rootDSE identify the following key information:

- The directory partitions (the domain, schema, and configuration directory partitions) that are specific to one domain controller.
- The forest root domain directory partition.

In this way, the rootDSE provides a "table of contents" for a given domain controller.

**rootDSE operational attributes**
The rootDSE publishes information through attributes on the rootDSE object. RFC 2251 and RFC 2252 define a set of seven rootDSE operational attributes that LDAP v3 servers are expected to publish, and these attributes are described in the following list. All LDAP servers recognize these attribute names, but when the attribute corresponds to a feature that the server does not implement, the attribute is absent.

**supportedLDAPVersion**
LDAP versions that the server implements. Domain controllers running Windows 2000 Server and later support LDAP v2 and LDAP v3.

**namingContexts**
Naming contexts (directory partitions) that this server masters (stores as a writable replica) or shadows (stores as a read-only replica). A client uses this attribute to choose suitable base objects for searching when the client contacts a server.

**subschemaSubentry**
The name of a subschema entry, which is used to administer information about the schema, in particular, the object classes and attribute types that are supported.

**altServer**
Uniform Resource Locators (URLs) of other servers that can be contacted when this server becomes unavailable. If the server does not know of any other servers, this attribute is absent. Clients can cache this information in case their preferred LDAP server becomes unavailable. This attribute is absent by default for Active Directory servers.

**supportedExtension**
Object identifiers that identify the supported extended LDAP operations that the server supports. LDAP extensions include operations such as different query types, paging operations, and sorting methods, which each have a different object identifier. This attribute is absent by default for Active Directory servers.

### Note

- A new extended LDAP operation on domain controllers running Windows Server 2003 enables client refresh of a dynamic entry in the directory. Object identifier = 1.3.6.1.4.1.1466.101.119.1 is defined and published in the **supportedExtension** attribute of the rootDSE object.

**supportedControl**
Object identifiers that identify the LDAP controls that the server supports. If the server does not support any controls, this attribute is absent. Server controls extend LDAP functionality; examples include a control to move objects across domains and a control to delete an entire subtree of a container object.

**supportedSASLMechanisms**
The names of the SASL mechanisms that the server supports. SASL is a standard for negotiating an authentication mechanism and, optionally, an encryption mechanism to provide client authorization for accessing Active Directory. By default, Generic Security Service API (GSSAPI) is supported.

**rootDSE informational attributes**
In addition to the previous operational attributes, Active Directory also supports the following informational attributes.

**currentTime**
The current time in the generalized time format.

**dsServiceName**
The distinguished name of the NTDS settings object that represents this domain controller in the configuration directory partition.

**defaultNamingContext**
The default naming context (directory partition) for a particular server. This value is the distinguished name of the domain directory partition for which this domain controller is authoritative.

**schemaNamingContext**
The naming context (directory partition) for the forest schema.

**configurationNamingContext**
The naming context (directory partition) for the forest Configuration container.

**rootDomainNamingContext**
The distinguished name for the domain naming context (directory partition) of the first domain that is created in this forest. This domain functions as the forest root domain.

**supportedLDAPPolicies**
Supported LDAP administrative query policies, such as the maximum size of a result set (**MaxResultSetSize**) and maximum page size (**MaxPageSize**).

**highestCommittedUsn**
Highest update sequence number (USN) that is committed to the database on this domain controller.

**dnsHostName**
The DNS name of this domain controller.

**serverName**
The fully qualified distinguished name for this domain controller.

**supportedCapabilities**
The object identifier value (1.2.840.113556.1.4.800) that indicates the additional capabilities of an Active Directory server, such as dynamic update, integrated DNS zones, and LDAP policies.

**LdapServiceName**
The SPN for the LDAP server, which is used for mutual authentication.

**isSynchronized**
A Boolean indicator for whether the domain controller has completed its initial synchronization with replica partners.

**isGlobalCatalogReady**
A Boolean indicator for whether the domain controller is prepared to advertise itself as a global catalog.

**domainControllerFunctionality**
The numeric level of Active Directory functionality for the domain controller.

**Distribution of Directory Data**
So that it can scale to tens of millions of objects, the Active Directory data store is logically partitioned in such a way that each domain controller does not store the entire directory. To accomplish logical partitioning, the data is categorized according to a naming scheme: object names group objects into logical categories so that the objects can be managed and replicated appropriately. In Active Directory, the largest of these logical categories is called a directory partition.

Every domain controller holds at least one directory partition that stores domain data, such as users, groups, and OUs. Every domain controller also stores two nondomain directory partitions that store forest-wide data, which includes the schema and configuration data.

The data store holds data for a single forest. Although there is a single directory, some directory data is distributed within domains, while other data is distributed throughout the forest without regard for domain boundaries. In Windows Server 2003 and later, data can also be distributed to domain controllers according to applications that use the data, where the scope of distribution is set by the application. The following table describes the three types of data that are stored in the Active Directory data store.

### Types of Active Directory Data

| Data Type | Description |
| --- | --- |
| Domain-wide data | • Domain-specific data is stored in a domain directory partition.<br>• A full, writable replica of the domain directory partition is replicated to every domain controller in the domain. |
| Forest-wide data | • Forest-wide data is stored in two directory partitions: the configuration directory partition and the schema directory partition. The Configuration container is the topmost object of the configuration directory partition; the Schema container is the topmost object of the schema directory partition.<br>• A full, writable replica of the configuration directory partition is replicated to every domain controller in the forest.<br>• A read-only replica of the schema directory partition is replicated to every domain controller in the forest. The schema is writable on only the domain controller that holds the schema operations master role, and writing to the schema requires first adding a registry entry on the schema master. Schema updates replicate to every domain controller in the forest.<br>• In addition to a full, writable replica of a single domain (the domain for which the domain controller is authoritative), special domain controllers that are designated as global catalog servers also store partial, read-only replicas of every other domain directory partition in the forest. (The read-only replicas in the global catalog are "partial" because they store only some of the attributes for each object.) A domain controller that is a global catalog server can be queried to find any object in the forest. |
| Application data | • Applications can use a new type of directory partition in Windows Server 2003 and later Active Directory, called an application directory partition, to store application-specific data that has a scope of interest that is smaller than the entire forest or domain. This data can be characterized as either changing frequently (dynamic) or having a short useful lifetime (volatile). For example, Windows Server 2003 and later DNS can use application directory partitions to store dynamically updated DNS zone data on only those domain controllers that are DNS servers rather than on all domain controllers in the domain, as is required for Windows 2000 Active Directory–integrated zones. |

**Directory Partition Hierarchy**

In Active Directory, a directory partition is a portion of the directory namespace. Each directory partition contains a subtree of the directory objects in the directory tree. The same directory partition can be stored as a replica on many domain controllers, and the replicas are updated through directory replication.

There is an important distinction between the physical storage of a directory partition and its logical position in the directory tree. Physically, all objects are stored in a single database table, regardless of the directory partition to which they are assigned because of their object names. Logically, the head of a directory partition appears in the naming hierarchy as the topmost object; that is, the Domain container, the Configuration container, and the Schema container each has a distinguished name that identifies its position in the hierarchy.

Every domain controller stores a replica of a domain directory partition, the configuration directory partition, and the schema partition. For more information about these partitions, see "Directory Partitions" later in this section.

**Note**

- Although the schema directory partition is replicated to every domain controller in the forest, it can be updated only on the domain controller that holds the schema operations master role.

The following figure is a conceptual diagram of the directory tree hierarchy, including the directory root (rootDSE) and the default directory partitions below the directory root. TherootDSE represents the top of the logical namespace for one domain controller and, as such, it represents the top of the LDAP search tree. There is only one root for a given directory, but the information that is stored in the root is specific to the domain controller to which you connect.

In any Active Directory forest, the first domain directory partition that is created in the forest (the forest root domain), the configuration directory partition, and the schema directory partition always form the hierarchy that is illustrated in the

following figure.

**Default Active Directory Partitions**




Additional directory partitions are possible in the form of application directory partitions, but these partitions are not stored by default on every domain controller.

**Directory Partition Cross-Reference Objects**

When Active Directory is installed to create the first domain controller in a new forest, the three directory partitions that are shown in the previous figure are created on the domain controller. At this time, a cross-reference object (class **crossRef**) is created for each directory partition in the Partitions container in the configuration directory partition (CN=partitions,CN=configuration,DC=*forestRootDomain*). Creation of each subsequent directory partition in the forest, either by installing Active Directory to create a new domain or by creating a new application directory partition on an existing domain controller, initiates the creation of an associated cross-reference object in the Partitions container.

**Note**

- You can also manually create a cross-reference object for an application directory partition.

A cross-reference object identifies the name and server location of each directory partition in the forest. The replication system uses this information to identify servers that store the same directory partitions. LDAP queries use cross-reference objects to create referrals to different domains.

**Forest Root Domain**

Because the forest root domain is the first domain that is created in a forest, it is the root name in the domain namespace hierarchy. In naming only, the topmost object of the configuration directory partition — the Configuration container — is the child of the forest root domain object in the hierarchy. The LDAP distinguished name of the Configuration container (CN=configuration,DC=*forestRootDomain*) reflects this naming hierarchy, which links the configuration directory partition to the forest.

Similarly, the topmost object in the schema directory partition — the Schema container — is the child of the Configuration container. The distinguished name of the Schema container (CN=schema,CN=configuration,DC=*forestRootDomain*) links the schema to the forest.

**Directory Partitions**

The Active Directory data store holds three default directory partitions:

- The configuration directory partition
- The schema directory partition
- The domain directory partition

Optionally, the data store may also hold one or more application directory partitions.

**Configuration Directory Partition**

The configuration directory partition is created initially when the first domain of a forest is created during the installation of Active Directory. Thereafter, it is replicated to every new domain controller that is added to the forest. The configuration directory partition holds information of global interest, for example, the default configuration and policy information for all instances of a given service in the forest.

**Note**

- Global information should be stored in one of two places: in a child of the Services container or in a child of a site object.

**Contents of the configuration directory partition**
In Active Directory tools, such as ADSI Edit, the configuration directory partition is represented by a Configuration container. When you open ADSI Edit (Adsiedit.msc) from the **Run** dialog box, the Configuration container for the forest of the domain to which you are connected is displayed, along with the current domain directory partition and the schema directory partition.

### Note

- If you open ADSI Edit with Microsoft Management Console (MMC) (as opposed to opening ADSI Edit from the **Run** dialog box), you must connect to the directory partition to view it.

The following objects are child containers in the Configuration container.

**DisplaySpecifiers**
Contains the objects that define different user interfaces (UIs) for each object class in the schema that requires a graphical user interface (GUI), for example, context menus and property pages. The display specification system uses the information that is stored in the display specifiers to form different UIs for administrators and end users. One set of elements can be associated with administrative applications, and a different set of elements can be associated with end-user applications. What you see and what the user sees are different, even though in both cases the UI references the same objects. The display specification system stores information for property sheets, context menus, icons, creation wizards, and localized class and attribute names.

This UI specification occurs at the level of each object class as defined in the schema in classSchema objects. Each classSchema object can have UI display specification information that is uniquely associated with it.

The DisplaySpecifiers container stores other containers that correspond to each locale that is supported. A locale is either a language or a language in combination with a country/region. Windows supports more than 150 locales, such as French (Belgium), Arabic (Saudi Arabia), and so forth. The names of locale containers are the hexadecimal representations of the locale identifiers (LCIDs). For example, the English (United States) locale container is 409.

Display-specifier objects (class **displaySpecifier**) are named by appending the LDAP Display Name of the class object with the string -"-Display.". For example, the **user** class has a corresponding display-specifier object called user-Display. Therefore, when an Active Directory administrative tool displays an object of a particular class, the object is displayed according to information that is contained in the display-specifier object whose name contains the same name as the respective class in the container for the current locale.

Because the Active Directory schema can be modified by creating new classes and attributes or by modifying existing classes, display-specifier objects can be modified to reflect any new UI elements that schema modifications require.

**Extended-rights**
Stores objects of the class **controlAccessRight** that can be used by applications to extend standard access control.

**ForestUpdates**
Stores operation objects that are generated by forest preparation tasks (when you run **adprep /forestprep**) so that the system can check for the tasks that have and have not been completed when you upgrade the first domain controller in the forest to Windows Server 2003 or later. The child object CN=Operations contains the objects that represent each update operation. These objects are named for the GUID of the operation. The child object CN=Windows2003Update is created to indicate that all **adprep** operations have run.

**LostAndFoundConfig**
Provides storage for configuration directory partition objects that are being created in containers that are simultaneously being deleted on another domain controller in the same forest. If, before replication, an object is created in or moved to a location that no longer exists after replication, the "lost" object is added to the LostAndFoundConfig container. A LostAndFound container in each domain directory partition serves the same purpose for domain-specific objects.

**NTDS Quotas**
Stores objects (class **msDS-QuotaControl**) that contain object ownership quota assignments for the configuration directory partition. Quotas limit the number of objects that a user (including inetOrgPerson), group, computer, or service can own in a domain directory partition, configuration directory partition, or application directory partition.

**Partitions**
Stores the cross-references to every directory partition in the forest, including the configuration partition, the schema partition, and all domain directory partitions. These cross-references to directory partitions make referrals to other domains possible during LDAP searches.

**Physical locations**

Not used in this version of Windows but reserved for future use.

**Services**
Stores data for various networking services and applications, such as Message Queuing (**MsmqServices**), public key services, and Routing and Remote Access. In the Windows NT container, the Directory Service object (**nTDSService**) has attributes that you can use to manage garbage collection intervals and dynamic object Time to Live (TTL). For the most part, however, objects in the Services container do not require administration. Specifically, do not publish application services in this container. Such services are best published in the Program Data container in the domain directory partition.

**Sites**
Identifies all of the sites in the network, the domain controllers in those sites, and the replication topology. The contents take the form of replication transports, subnets, and the first site and site link that are created, called Default-First-Site-Name and Default-First-Site-Link, respectively.

**Well-known security principals**
Contains the special identities that are defined by the security system, such as Everyone, Local System, Principal Self, Authenticated User, and Creator Owner.

**Schema Directory Partition**
The Active Directory schema is stored in the Schema container in the schema directory partition. The schema consists of a set of object classes, attributes, and syntaxes. It also defines rules that ensure that objects are created and modified with consistency. Active Directory contains a default set of classes and attributes that cannot be modified. However, if you have Schema Admins credentials and if schema modification is enabled for the domain controller, you can extend the schema by adding new attributes and classes to represent application-specific classes. These changes are targeted at the domain controller that is the schema master for the forest. Only the schema master stores a writable copy of the schema.

**Domain Directory Partition**
When you create a new domain, a domain directory partition is created in Active Directory as an instance of the class **domainDns**. A cross-reference object is added for the domain in the Partitions container to advertise the domain's location in the directory.

**Contents of a domain directory partition**
The topmost object in each domain directory partition is a domainDns object that is named with the DNS name of the domain.

A domain directory partition is represented by a domain container, and a domain directory partition has the following child containers.

**Builtin**
Default local group accounts, such as Administrators and Users, that are installed whenever a new workstation, server, or domain controller is set up.

**Computers**
A default storage area for "new" computer objects that were originally created through legacy APIs. When a Windows NT 4.0 domain or a Windows NT 3.51 domain is upgraded to Windows 2000, or when a Windows NT 4.0 domain is upgraded to Windows Server 2003 or later, the computer accounts are moved automatically to the Computers container. The Computers container also supports the Windows NT 4.0 tool User Manager (Usrmgr). This container cannot be renamed, but when a Windows Server 2003 or later domain has a domain functional level of Windows Server 2003 or later, the default location for computer accounts can be specified as an OU to redirect the location of newly created computer objects.

**Domain Controllers**
The default container for new domain controllers. The Domain Controllers container cannot be renamed.

**ForeignSecurityPrincipals**
Proxy objects for security principals that are from Windows NT 4.0 domains or Windows NT 3.51 domains, or that are from different forests, and that have been added to Windows 2000 or later groups.

**LostAndFound (advanced features)**
A storage area for new domain objects whose containers were deleted elsewhere at the same time that the object was created. If an object has been created in or moved to a location that is missing after replication, the "lost" object is added to the LostAndFound container. The LostAndFoundConfig container in the configuration directory partition serves the same purpose for forest-wide objects.

**NTDS Quotas (advanced features)**
A storage area for objects of the class **msDS-QuotaControl** that contain object ownership quotas for the domain directory partition. Quotas limit the number of objects that a user, group, computer, or service can create in a directory partition.

**Program Data (advanced features)**
An empty container that is available for applications to store application-specific data in the domain directory partition.

**System (advanced features)**
Built-in system settings for the various system service containers and objects.

**Users**
A default storage area for new user accounts created through legacy APIs that are not Active Directory–aware. When a Windows NT 4.0 domain or a Windows NT 3.51 domain is upgraded to Windows 2000 or later, or when a Windows NT 4.0 domain is upgraded to Windows Server 2003 or later, the user accounts and groups are moved automatically to the Users container. The Users container also supports the Windows NT 4.0 tool User Manager (Usrmgr). This container cannot be renamed, but when a Windows Server 2003 or later domain has a domain functional level of Windows Server 2003 or later, the default location for user and group accounts can be specified as an OU to redirect the location of newly created user objects.

### Note

- The Users and Computers containers and several other special containers, called "well-known" containers, can be located dependably by applications.

**Deleted Objects**
A special container, which is not visible in the UI, to which objects are moved when they are deleted. The deleted objects are stored as tombstones, which are eventually removed by garbage collection. The contents of the Deleted Objects container are visible if you search by using the 1.2.840.113556.1.4.417 control, which enables you to see deleted objects.

**Infrastructure**
An object of class **infrastructureUpdate** that identifies the NTDS Settings object of the domain controller that holds the infrastructure operations master role for the domain.

**Contents of the System container**
The System container stores per-domain operational information, which includes the default local security policy, file link tracking, Microsoft NetMeeting network meeting objects, objects representing other trusted domains, and containers for RPC and Windows Sockets (Winsock) connection points.

The System container has the following child containers.

**AdminSDHolder**
Administrator security descriptor holder. Windows Server 2003 and later implements protection of administrative groups through a background task that computes the set of memberships and checks whether their security descriptors are well-known protected security descriptors. If the security descriptor of any administrative account does not match that of AdminSDHolder, the security descriptor is overwritten with the security descriptor of AdminSDHolder. This task is executed only on the domain controller that holds the primary domain controller (PDC) emulator operations master role.

**COMPartitions**
A namespace that is used by Component Services to allow multiple versions of the same COM+ application to exist on the same physical computer.

**ComPartitionSets**
A conceptual collection of COM+ partitions.

**Default Domain Policy**
The common domain location for the AppCategories container (class **classStore**), which holds objects of the class **categoryRegistration**. These objects are managed by the Group Policy Software Installation MMC extension, and they can be associated with applications that are deployed through Group Policy in the domain.

### Note

- The Group Policy object (GPO) for the default domain policy is stored in the Policies container.

**Dfs-Configuration**
Lists the fault-tolerant Distributed File System (DFS) configuration and DFS volume information.

**DomainUpdates**
Stores operation objects that are generated by domain preparation tasks (when you run **adprep /domainprep**) so that the

system can check for tasks that have and have not been completed when you upgrade the first domain controller in the domain to Windows Server 2003 or later. The child Operations object contains the objects that represent each update operation. These objects are named for the GUID of the operation. The child Windows2003Update object is created to indicate that all **adprep** operations have run.

**File Replication Service**
Lists the Domain System Volume (SYSVOL) and provides a replication schedule from Sunday through Saturday, 12:00 A.M. to 12:00 A.M.

**FileLinks**
Used by the Distributed Link Tracking Server service (TrkSvr) to store information about linked files that have moved across NTFS volumes. Includes **ObjectMoveTable**, which tracks moved files, and **VolumeTable**, which maps volume IDs to computer IDs.

**IP Security**
Contains the Internet Protocol security (IPSec) policies that are applied to local computers, domain member servers, domains, OUs, or any GPO in Active Directory. Depending on your organization's guidelines, IPSec policies can store multiple security specifications, called rules, so that one policy can be applied to multiple computers. These rules apply to all users who log on to the computer

**Meetings**
A folder that is used by Microsoft NetMeeting to publish network meeting objects.

**MicrosoftDNS**
A container in which Active Directory–integrated zone database records are created when the scope of replication is all domain controllers in the domain. When DNS data is stored in Active Directory, each DNS zone is an Active Directory container object (dnsZone). The dnsZone object contains a dnsNode object for every unique name in that zone. The dnsNode object has a **dnsRecord** multivalue attribute that contains a value for every resource record that is associated with an object's name. When the scope of replication of DNS data is all DNS servers in the domain or all DNS servers in the forest, DNS zone data is stored in application directory partitions, not in the domain directory partition.

**Policies**
Contains GPOs, which specify user and computer configurations for groups of users and computers. This container stores the default domain policy (lists the security groups and default permissions for the domain); default domain controller policy; and policy for passwords, lockouts, Kerberos, EFS data recovery, and trusted root certificates. Each GPO in the Policies container is identified by a GUID, and each GPO includes the following:

- Version information that is used to ensure that information is synchronized with Group Policy template information
- Status information that indicates whether the GPO is enabled or disabled
- A list of components, or extensions, that have settings in the GPO

In addition to being stored in the Policies container, GPOs are also stored in a Group Policy template, and they are identified by GUID. The Group Policy template is located in the SYSVOL, and it is used to store file type data for the GPOs.

**RAS and IAS Servers Access Check**
Verifies permissions for the RAS and IAS Servers domain local security group to access user account properties. When Routing and Remote Access (configured for Windows authentication) or Internet Authentication Service (IAS) queries a domain controller to validate user account credentials and receives a failure notification, there is no indication why the failure occurred. It might have occurred because the user's credentials were invalid or because the server running Routing and Remote Access or IAS did not have the appropriate permissions to access user account properties. Routing and Remote Access and IAS servers obtain access to user account properties by adding the computer account of the server to the RAS and IAS Servers group. The RAS and IAS Servers group has Read permission to the RAS and IAS Servers Access Check container. By attempting access to the RAS and IAS Servers Access Check container, the Routing and Remote Access or IAS server determines whether or not it has permissions to access user account properties.

**RpcServices**
Includes the RPC name service lookup for domains using versions of Windows earlier than Windows 2000.

**WinsockServices**
Winsock services that publish themselves by using the registration and resolution (RnR) APIs are published in this container.

**WMIPolicy**
Stores Windows Management Instrumentation (WMI) filters. A WMI filter is a query based on configuration and event information that is made available by WMI providers. By associating a WMI filter with a specific GPO, that WMI filter determines which computers process policy settings in that GPO. Computers that meet all of the conditions that are specified in the WMI filter return a value of TRUE and then process the GPO.

**Application Directory Partitions**
Application directory partitions are special partitions that can be created on domain controllers running Windows Server 2003 or later to provide LDAP storage for dynamic and volatile data. In Windows 2000 Server, nondomain data is limited to configuration and schema data, which is replicated to every domain controller in the forest. In a Windows Server 2003 or later forest, application directory partitions provide storage for nondomain, application-specific data that can be replicated to any arbitrary set of domain controllers in the forest — as few or as many as needed by the application that uses the data.

### Note

- Applications must be specifically programmed to use application directory partitions.

**Characteristics and requirements**
Application directory partitions have characteristics that make them suitable for storing dynamic data (data that is updated frequently) or volatile data (data that has a specified useful lifetime), as opposed to more static data that domain directory partitions and other directory partitions store. Although application directory partitions are represented as instances of the **domainDNS** class, as are domain directory partitions, they have characteristics that differentiate their behavior from domain directory partitions.

### Note

- Applications can create application directory partitions on Windows Server 2003 or later based domain controllers by explicitly creating a domainDNS object. Only Windows Server 2003–based and later domain controllers allow creation of domainDNS objects by means other than Active Directory installation. Because Windows 2000 Server and earlier domain controllers do not recognize the request to create a domainDNS object, no special instance type is required to differentiate between application directory partitions and domain directory partitions.

**Similarities to domain directory partitions**
The following characteristics of application directory partitions are the same as the characteristics of domain directory partitions:

- Naming: Application directory partitions can be named in the same manner as domains, and they can be attached anywhere in the Active Directory namespace where a domain can be attached. Therefore, an application directory partition can be a child of a domain or of another application directory partition.

- DNS location: Application directory partitions can be discovered through DNS.

- Cross-reference object: A cross-reference object is created in the CN=partitions,CN=configuration,DC=*ForestRootDomain* container for each application directory partition.

- Change notification for replication: The time intervals that control the latency of the initiation of an originating change notification to replication partners within a site can be configured. These configuration changes are applied to all directory partitions, but they are most useful for application directory partition replication of dynamic data.

- Access control: The security and access control model for the application directory partition is the same as that for other partitions in Active Directory.

- TTL: For objects that are stored in an application directory partition, you can assign a TTL value that is configurable. The only requirement is that the domain controller is running Windows Server 2003.

  For objects that are stored in a domain directory partition replica on a Windows Server 2003–based or later domain controller in a forest that has a forest functional level of Windows Server 2003 or later, you can assign a TTL. For all other functional levels, however, TTL values cannot be assigned to objects in domain directory partitions.

**Differences from domain directory partitions**
The following characteristics are different from domain directory partitions and specific to application directory partitions:

- Credentials for creation: Creating an application directory partition requires Domain Admins credentials in the forest root domain or Enterprise Admins credentials.

- Replication scope: Whereas domain directory partitions are replicated to every domain controller in a domain, application directory partitions can be replicated to any selected set of domain controllers in the forest. Because application directory partitions do not have domain or site affiliations, either an administrator or the application that creates the directory partition must explicitly define the scope of replication.

  **Note**

  - It is best to store replicas of dynamic data on domain controllers in the same site because intersite replication latency might not be acceptable for replica consistency.

- NetBIOS name: Application directory partition objects do not have network basic input/output system (NetBIOS) names. For this reason, their cross-reference objects, which are usually identified by the unique NetBIOS names of their respective directory partitions, are named by GUIDs.

- Contents: Security principals (users, groups, services, and computers) cannot be stored in application directory partitions.

- Global catalog: Application directory partitions are not replicated to the global catalog.

- Volatility: Whereas domain data should be relatively static, application directory partition data can be volatile, such as the real-time data that is used by Microsoft Telephony API (TAPI) voice conferencing applications.

**Stored objects**
Any type of object, with the exception of security principals, can be stored in an application directory partition. In addition, objects that reside in an application directory partition:

- Can maintain distinguished-name references to other objects in the same application directory partition, objects in the configuration and schema directory partitions, and any directory partition root object.

- Cannot maintain distinguished-name references to objects in other application directory partitions or in domain directory partitions.

- Cannot be moved to other Active Directory partitions outside the application directory partition in which they were created.

Applications can be programmed to store their dynamic data in application directory partitions. The following are examples of applications and the kinds of dynamic objects that they store:

- Real-time communication applications: These types of applications typically maintain rendezvous information about active users (user name, computer, and IP address) and conferences in progress (conference name, description, time, and originator).

- Network services: Routing and Remote Access, Remote Authentication Dial-In User Service (RADIUS), Dynamic Host Configuration Protocol (DHCP), and Common Open Policy Service (COPS) servers need to store dynamically changing forms of data in a directory so that they can be accessed using one uniform method: LDAP. This data generally does not require global replication, and it is also shared by many applications in the proximity of a single directory. (In this case, the data is shared between RAS, RADIUS, DHCP, and COPS servers in a particular site.) The object types include information about user sessions and bandwidth used.

- Directory-enabled networking: For implementing policy-based networking, there is a large amount of nonstatic, or volatile, information that must be maintained, including information about the state of network links, the flows that are active through a router, and the data rate for each flow. Policy servers might need access to such information to make a

decision based on policies that are predicated on dynamic network state in addition to static data that is stored in the directory.

The application directory partition, with its controlled replication scope, can accommodate transient data needs in Active Directory by providing network applications with uniform and consistent access through LDAP to both static and dynamic data.

**Security descriptor reference domains**

When an object is created in Active Directory, it is assigned a security descriptor that contains default access control entries (ACEs) that are taken from the **defaultSecurityDescriptor** attribute of the class of which the object is an instance. The default ACEs can include domain-relative, well-known SIDs. In the case of objects in a domain directory partition, the reference domain by which the system interprets the domain-relative SIDs is the domain that is represented by the domain directory partition.

In the case of an application directory partition, there is no security association between the application directory partition and the domain of the domain controller on which it is created. In a Windows Server 2003 or later forest, a new, single-valued attribute on the cross-reference object — **msDS-SDReferenceDomain** — stores the name of the reference domain that is used to interpret the domain-relative SIDs in default ACEs for new objects in an application directory partition.

When an application directory partition is created, the value of **msDS-SDReferenceDomain** on the respective cross-reference object is determined as follows:

- If the parent object of the application directory partition is another application directory partition, the value of **msDS-SDReferenceDomain** is set to the value of **msDS-SDReferenceDomain** on the cross-reference object of the parent application directory.
- If the parent object is a domain directory partition, the value of **msDS-SDReferenceDomain** is set to the parent domain.
- If there is no parent object, the value of **msDS-SDReferenceDomain** is set to the forest root domain.

**Note**

- It is recommended that domain local groups not be used in the access control lists (ACLs) of objects in application directory partitions. If an application directory partition has replicas on domain controllers in two different domains, the membership of such domain local groups in ACLs can expand to yield different results, depending on the domain controller to which a user connects.

**GUID names for cross-reference objects**

Domains in Active Directory must have NetBIOS names that are unique in the forest so that they are backward compatible with computers that are running earlier versions of Windows. Domains must also have DNS names that are unique. The relative distinguished name of the cross-reference object for each domain directory partition is taken from its NetBIOS name, which guarantees that each cross-reference object is unique within the Partitions container. For example, if a domain has a NetBIOS name of CONCORP and a fully qualified DNS domain name of concorp.contoso.com, the distinguished name of its cross-reference object is CN=concorp,CN=partitions,CN=configuration,DC=contoso,DC=com. The domain component (DC) of the cross-reference object distinguished name is always the forest root domain, which is always the parent of the configuration directory partition. If there were another domain in the same forest named concorp.avionics.contoso.com, this domain would necessarily be given a NetBIOS name other than CONCORP (for example, CONCORPAV).

In the case of application directory partitions, there is no associated NetBIOS name, only the DNS name. Given the nature of DNS hierarchical naming, which allows identical left-most name components, a cross-reference object cannot use the left-most DNS name component as its relative distinguished name. For this reason, cross-reference objects that reference application directory partitions use their object GUID as their relative distinguished names.

To associate a cross-reference object with its application directory partition, you can read the **nCName** attribute of the cross-reference object, which gives the distinguished name of the directory partition.

**Data Store Physical Structure**

The physical structure of the data store consists of the Active Directory database file (Ntds.dit) and its associated log and temporary files. The data that is held in the database file includes both static data and dynamic data. The following figure illustrates the physical structure of the data store.

### Data Store Physical Structure




The following table describes the physical components of the data store.

**Data Store Physical Structure Components**

| Component | Description |
|---|---|
| NTDS.DIT | The physical database file in which all directory data is stored. This file consists of three internal tables: the data table, link table, and security descriptor (SD) table. |
| EDB.LOG | The log file into which directory transactions are written before being committed to the database file. |
| EDB.CHK | The file that is used to track the point up to which transactions in the log file have been committed. |
| RES1.LOG, RES2.LOG | Files that are used to reserve space for additional log files if EDB.LOG becomes full. |

**Active Directory Database**

Active Directory data is stored in the Ntds.dit database file. The Active Directory database (Ntds.dit) contains three internal tables, the data table, link table, and SD table, which are described in the following sections.

Two copies of Ntds.dit are present in separate default locations on a domain controller, *systemroot*\NTDS and *systemroot*\System32:

- *systemroot*\NTDS\Ntds.dit stores the database that is in use on a domain controller. It contains the values for the domain and a replica of the values for the forest (the Configuration container data).

- *systemroot*\System32\Ntds.dit is the distribution copy of the default directory that is used when you install Active Directory on a server running Windows Server 2003 or later to create a domain controller. Because this file is available, you can run the Active Directory Installation Wizard without having to use the server operating system CD.

**Data Table**

The data table contains all the information in the Active Directory data store: users, groups, application-specific data, and any other data that is stored in Active Directory after its installation. The data table can be thought of as having rows (each representing an instance of an object, such as a user) and columns (each representing an attribute in the schema, such as **GivenName**). For each attribute in the schema, the table contains a column, called a field. Field sizes can be fixedor variable. Fixed-size fields contain an integer or long integer as the data type. Variable-size fields typically hold string types, for example, Unicode strings. The database allocates only as much space as a variable-size field needs: 16 bits for a 1-character Unicode string, 160 bits for a 10-character Unicode string, and so on.

The database space that is used to store an object depends on the number of attributes for which values are set and the size of the values. For example, if the administrator creates two user objects (User1 and User2), sets only the minimum attributes on them, and then later adds a 10-character description to User2, the User2 space is approximately 80 bytes bigger than the User1 space (20 bytes for the 10 characters, plus metadata on the newly generated attribute).

Database records cannot span database pages; therefore, each object is limited to 8 kilobytes (KB). However, some attribute values of an object do not count fully against this limit. Long, variable-length values can be stored on a different page than the object record, leaving behind only a 9-byte reference. In this way, an object and all its attribute values can be much larger than 8 KB.

**Link Table**

The link table contains data that represents linked attributes, which contain values that refer to other objects in

Active Directory. An example is the **MemberOf** attribute on a user object, which contains values that reference groups to which the user belongs. The link table is much smaller than the data table.

**SD Table**
The SD Table contains data that represents inherited security descriptors for each object. With the introduction of the SD table in Windows Server 2003 or later, inherited security descriptors no longer have to be duplicated on each object that inherits security descriptors. Instead, inherited security descriptors are stored in the SD table and linked to the appropriate objects.

**Log Files**
The Active Directory data store includes the log files that are described in the following table. For information about how these log files are used, see "Logging Transactions to Enable Log-based Recovery" later in this section.

### Log Files

| File | Description |
|---|---|
| Edb.log | This file stores directory transactions before they are written to the Ntds.dit directory database file. This file has a fixed size of 10 megabytes (MB). |
| Edb00001.log, Edb00002.log, Edb00003.log, and so on | Active Directory creates additional log files as necessary as existing log files fill up. Each log file has a fixed size of 10 MB. |
| Edb.chk | This file keeps track of the point up to which transactions in the log file have been committed. |

**Minimum Space Requirements**
The minimum space requirements for the directory database and log files — and the requirements for which monitoring is essential — are free disk space on the partition or partitions that store the directory database and logs, as follows:

- Ntds.dit partition and log file partition: On either partition, free disk space must not fall below the greater of 500 MB or 20 percent of the combined Ntds.dit and log file sizes.

- Ntds.dit and log files on the same volume: Free disk space must not fall below the greater of 1 gigabyte (GB) or 20 percent of the combined Ntds.dit and log file sizes.

**Objects and Capacity**
There are no practical limits to the number of objects that can be stored in Active Directory. The Active Directory database has been tested for up to 60 million objects. Performance tests show logon performance for a single LDAP client to be the same with 10,000 objects, 100,000 objects, and 1 million objects; that is, the directory service does not slow measurably when the size of the database increases.

### Note

- In a mixed environment in which backup domain controllers (BDCs) are running Windows NT 4.0, the recommended limit for the number of security principal objects per domain is 40,000 (the sum of users, groups, and computers). This limit is based on the Windows NT 4.0 SAM database storage capacity.

**Maximum Database Record Size**
The maximum size of a database record is 8110 bytes, based on an 8-kilobyte (KB) page size. Because of variable overhead requirements and the variable number of attributes that an object might have, it is impossible to provide a precise limit for the maximum number of multivalues that an object can store in its attributes. For all practical purposes, the size of objects is not significant in Active Directory if you use the recommended guidelines described in "Static Data" later in this chapter.

The only value that can actually be computed is the maximum number of values in a nonlinked, multivalued attribute when the object has only one attribute (which is impossible). In Windows 2000 Active Directory, this number is computed at 1575 values. From this value, taking various overhead estimates into account and generalizing about the other values that the object might store, the practical limit for number of multivalues stored by an object is estimated at 800 nonlinked values per object across all attributes.

| Note |
|---|
| Attributes that represent links do not count in this value. For example, the members linked, multivalued attribute of a group object can store many thousands of values because the values are links only. |

The practical limit of 800 nonlinked values per object is increased in Windows Server 2003 and later. When the forest has a functional level of Windows Server 2003 or higher, for a theoretical record that has only one attribute with the minimum of overhead, the maximum number of multivalues possible in one record is computed at 3937. Using similar estimates for overhead, a practical limit for nonlinked multivalues in one record is approximately 1200. These numbers are provided only to point out that the maximum size of an object is somewhat larger in Windows Server 2003 and later.

**Active Directory Data**

The key characteristics of data that is stored by any directory service are size and latency. A directory service should store objects that are not so large that they hamper replication. Therefore, large unstructured data sets are not appropriate for storage in Active Directory. Domain data and configuration data are typically static in nature, having a useful lifetime at least as long as a single period of replication latency.

In a Windows 2000 Server forest, Active Directory does not support storage of nonstatic (dynamic or volatile) data. In a Windows Server 2003 or later forest, however, Active Directory can store nonstatic data in application directory partitions, which can be configured for limited replication or used as real-time data stores for data that is too volatile for replication.

**Static Data**

Active Directory is appropriate for the storage of static data that has the following characteristics:

- The data is globally useful information in the domain that needs to be replicated to each Active Directory domain controller.

- The data has well-defined object attributes and semantics.

- The data has a useful life that is at least two times the maximum replication latency for the forest, including replication of data that is marked to replicate to the global catalog. In general, if data can become outdated before the completion of a replication cycle or shortly thereafter, it should not be stored in Active Directory. Clients should be able to tolerate the inability to update data for at least as long as it takes for the data to be replicated throughout the domain.

- The data-per-attribute value is not so large that it affects performance. Most attribute values are replicated as a single block of data; therefore, an attribute that is *x* MB in size requires an equivalent amount of buffer space in the sending domain controller and in the receiving domain controllers. If the amount of required buffer space is large, the performance of the domain controllers can be affected adversely.

  Attributes can have large values when the attribute is multivalued and linked. For these attributes, only the values that change are replicated, not the entire attribute.

**Dynamic Data**

An object that has a TTL value is considered to be dynamic data. Windows Server 2003 introduces a new auxiliary class, **dynamicObject**, that can be attached to object instances for the storage of TTL values. For an object containing a TTL value, when the time that is specified by the TTL value is reached, the object is automatically removed from Active Directory by garbage collection. Dynamic data typically changes faster than the replication latency that is required for propagating the change to all replicas of the data.

The following characteristics apply to dynamic data:

- A dynamic object that is deleted as a result of expiration of its TTL does not leave a tombstone.

- Dynamic objects are treated in the same manner as nondynamic objects during search, compare, add, delete, and modify (including rename) operations.

- A nondynamic object, after it is created, cannot be changed to a dynamic object. Likewise, a dynamic object, after it is created, cannot be changed to a nondynamic object.

- A nondynamic object cannot be added in a subordinate position to a dynamic object.

- Dynamic objects with TTL values are supported in the following directory partitions:

  - Domain directory partitions at the Windows Server 2003 or higher forest functional level

- Application directory partitions on a Windows Server 2003 or higher–based domain controller, irrespective of domain or forest functional level

- Dynamic objects are not supported in configuration directory partitions or schema directory partitions.

**TTL for dynamic objects**
The TTL for a dynamic object is set when the object is created, and it is automatically decremented thereafter. When its TTL expires, the dynamic object disappears without leaving a tombstone. However, a client application can programmatically cause a dynamic object to remain longer than its current remaining lifetime by refreshing (modifying) its TTL value, as stored in the **entryTTL** attribute on the object.

The Directory Service object (nTDSService) in the Configuration container (CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,DC= *forestRootDomain*) stores the default and minimum dynamic object TTL values for the forest in the multivalue attribute **msDS-Other-Settings**.

The default values for the two configurable TTL parameters are as follows:

- Default TTL value = 86400 seconds (1 day)

- Minimum TTL value = 900 seconds (15 minutes)

The default TTL value is assigned to a newly created dynamic object if a TTL is not explicitly specified by the application that creates the object. The minimum TTL overrides any client-specified TTL value that is lower than the configured minimum. No configurable maximum TTL value needs to be provided, because a maximum is imposed by the attributeSchema object.

The syntax for the two configurable TTL parameters is as follows, where NNNN is expressed in seconds:

```
                DynamicObjectDefaultTTLSeconds=NNNN
DynamicObjectMinTTLSeconds=NNNN
```

By configuring these values, you can set TTL values that enforce a low level of refresh traffic or, conversely, provide a highly up-to-date directory.

**Data Store Processes and Interactions**
The Active Directory data store performs a number of processes and interactions that are important to the functioning of Active Directory. The following sections describe these processes and interactions.

These descriptions make the following assumptions regarding the environment in which the Active Directory data store is running:

- The disk volume on which the directory database is stored contains ample free space.

- Network connectivity on the domain controller is functioning properly.

**Database Operations**
Data operations (such add, modify, and delete) are committed to the Active Directory database as transactions, which are the units of work that are performed by a database. Transactions are atomic; that is, they are either completed in full or not applied at all. If for any reason an error occurs and a transaction is unable to complete all of its steps, the system is returned to the state that existed before the transaction began. An example of an atomic transaction is an account transfer transaction in which money is removed from account A and placed into account B. If the system fails after it removes the money from account A and before it places the money into account B, the transaction processing system puts the money back into account A and returns the system to its original state; that is, the transaction processing system rolls back the transaction.

In Active Directory, operations on a single object cannot be applied across multiple objects. Active Directory writes a transaction synchronously to the transaction log file and then to the database. First, a change is made to an in-memory copy of the object. Then, the change is written to the log file, which ensures that the change takes effect, even if the database shuts down after that point. The database engine continually updates the database file with recent changes. The database update works from memory — not from the log files — and it keeps pace with the updates, rather than waiting for the server to be available. This method of performing updates is referred to as “advancing the checkpoint,” where the checkpoint is the point in time at which all changes that have been made so far have been fully written to the database.

**LDAP Functional Model**
Database operations in Active Directory are based on the LDAP functional model. The LDAP functional model includes support for common database operations, such as connecting to a directory database, searching for data, and modifying data. An LDAP client connects to Active Directory and requests information or performs an operation. Active Directory performs the operation or provides the information, or it refers the client to another LDAP server that might be able to do so.

Nine operations form the LDAP functional model. These operations are grouped into three areas:

- Authentication. Enables the client to prove its identity to the DSA:
    - Bind
    - Unbind
    - Abandon
- Interrogation. Enables the client to interrogate the directory and retrieve information:
    - Search
    - Compare
- Update. Enables the client to update information in the directory tree:
    - Add
    - Delete
    - Modify
    - Modify relative distinguished name

A typical LDAP client application interacts with Active Directory through the operations that are described in the following sections.

For more information about LDAP, see "Lightweight Directory Access Protocol (LDAP)" in the Microsoft Platform SDK on MSDN.

**Establishing a connection**
When an LDAP client connects to Active Directory, an LDAP session is established. Options are available to affect the way in which the connection is established. These options include setting a time-out value, connecting to a secure LDAP server, and verifying that a server is available.

For the purposes of protocol exchanges, all protocol operations are encapsulated in a common envelope. LDAPMessageis encapsulated in the Protocol Data Unit (PDU) format. LDAPMessage consists of protocol operations, such as LDAP Bind Request, LDAP Bind Response, LDAP Search Request, and LDAP Search Response operations.

LDAPMessage PDUs are mapped directly to the TCP data stream. Active Directory clients use the following LDAP ports:

- Port 389. In accordance with RFC 2251, Active Directory uses port 389 as the default port for domain controller communications.
- Port 636. Active Directory supports port 636 for LDAP Secure Sockets Layer (SSL) communications.
- Port 3268 and port 3269. The global catalog monitors port 3268 for LDAP communications and port 3269 for LDAP SSL communications.

**Modifying a directory object**
The LDAP API contains functions for adding and deleting directory objects and for comparing and modifying attribute values in

existing objects. LDAP v3 provides extensions to add, delete, and modify functions that enable the use of controls to perform these operations. Controls are described in RFC 2251 as a mechanism to extend the functionality of LDAP. Windows Server 2003 and later supports several extension controls that go beyond the controls that are identified by LDAP v3.

**Searching the directory**
Searching is the most common directory activity, and the LDAP APIs provide a variety of search criteria and result-retrieval methods. The client searches the LDAP server by passing it a set of parameters that describe the information of interest. These parameters describe where and how to search, and they define the search criteria that a client needs. The client uses a search filter to describe the objects that it wants. Search filters are defined in RFC 2254. Extensions to the base LDAP API, in the form of LDAP v3 controls, provide the ability to sort results and set various limits on the search operation. Search results can be processed by paging and by sorting. Paging and sorting are supported in Windows 2000 Server and later as new LDAP v3 control extensions for processing search results on the server.

**Closing the connection (unbinding)**
Unbinding closes the connection between the directory client and the DSA and disposes of the session handle. Applications call the unbind function when an LDAP client finishes communicating with a server. There is no server response to an unbind request.

**Logging Transactions to Enable Log-based Recovery**
To maximize system efficiency and ensure data integrity, Active Directory uses “write ahead” logging. With “write ahead” logging, transactions are written as quickly as possible to a transaction log file, and they are then written to the Ntds.dit database file as the system has time or at system shutdown.

The log file that is used by Active Directory to record transactions is called Edb.log, which has a fixed size of 10 MB. When this log file becomes full, Active Directory automatically creates a new log file, for example, Edb00001.log, which also has a fixed size of 10 MB. Active Directory continues to create additional log files as they become necessary. Through circular logging, the oldest log file is purged when all transactions in that log have been written to the Ntds.dit database. In addition to the active log files, two reserve log files exist, called Res1.log and Res2.log, ,also with a fixed size of 10 MB. These files reserve the last 20 MB of disk space on the drive to provide room for a graceful shutdown if all other disk space is consumed.

Active Directory also maintains a checkpoint file (Edb.chk) that records which transactions in the log have been committed to the database. In circular logging, log files are deleted when the checkpoint advances past the last transaction that is included in that log file. The checkpoint cannot be advanced past the beginning of any transaction that is still open.

If the computer stops responding, the ESE can detect an improper shutdown by checking the last log recorded. If the last record is not a “shutdown” record, the ESE reprocesses the logs from the checkpoint that is stored in Edb.chk. This event occurs at the first reboot after the system is shut down improperly. If the checkpoint file is missing for any reason, every transaction within the log file is reprocessed.

The number of log files that exist at any point in time is related to the number and size of transactions that are running in the database. If a system is idle, you can expect to see one log file. If a domain controller is making large changes (for example, replicating a 5,000-member group), you can expect to see several log files. If you just edited the schema to define an index on an attribute that has many large data values, you might see hundreds of log files. (The checkpoint cannot be advanced until the index creation is completed.)

**Growth and Space Management**
The directory database is a self-maintained system in which growth is managed by periodically compacting data and reorganizing free space into contiguous pages. Other than regular backup, the directory database requires no maintenance during ordinary operation. The only requirement is adequate space to accommodate normal growth.

The directory database grows when more data is added to it than free space is available to accommodate the new values. Free space, also called “white space,” is managed automatically by the directory database. Conditions under which free space is made available to the database can affect database size, sometimes in ways that are not expected. For this reason, an understanding of how free space is managed can help provide correct interpretation of changes in the size of the directory database file. The following sections describe space allocation and free space reclamation.

**Space Allocation**
Space is allocated hierarchically in the directory database, with the database at the top of the hierarchy. Object tables are children of the database. Each table has a child long-value table, and each table can have child indexes. The long-value table stores values that are too large for the main object table. Security descriptors are examples of values that are often large enough to require storage in the long-value table.

As objects are deleted and values are replaced with smaller values, free space is returned to the database. However, there is no optimum or target amount of white space to be maintained. The amount of unused space in the directory database is of concern only when monitoring indicates that free disk space is becoming limited. You can configure a domain controller to log

an event that reports the amount of unused space in the directory database that can be reclaimed by offline defragmentation.

**Requesting space**
The database can use free space if the space is available to the table or index that needs it and if the space occurs in blocks that are large enough to accommodate the transaction. When an index or table requires additional space, it can request space only from its parent. For example, the available space in an index cannot be used by another index. If the parent table requires space, it must request space from the database. If space is not available at the database level to complete the request, the database increases its size to accommodate the transaction.

**Space consumption by indexes**
The estimated space that is required for adding an index to an attribute is rarely greater than 1 percent of database size. However, Tuple indexes, which index an attribute in a way that allows users to search the attribute by using partial text strings, can be quite large. To provide partial text string searching, a Tuple index includes one entry for each substring of a **Text** or **Long Text** column, with a minimum and maximum length allowed for the substrings that are indexed. Although Tuple indexes can dramatically speed queries of the form "find all records that contain *string value*," they can potentially increase database size by as much as 20 percent, and they should be used with caution.

**Storing data**
Objects are stored in records that are written to 8-KB data pages. The maximum record size is 8110 bytes, with the exception of long-value columns, which can be up to 2 GB. As values are deleted, space that is freed by transactions is returned to the table that owns the space. Free space is returned gradually in small segments, and it is periodically reorganized into contiguous pages through an automatic process called online defragmentation.

Data pages are filled with records in an ordered fashion. For example, all records beginning with A must be stored in the page that is designated as storing the values that range from A to some other value. For example, given pages A-E and F-J, if a value corresponds to the A-E page and that page is full, the value cannot be stored in the F-J page, even if space is available. Instead, one or more new pages are added. When a table grows, it grows in increments equal to one-fourth its initially allocated size.

In this way, Active Directory data storage optimizes lookup speed as opposed to storage space. However, because of the hierarchical means of allocating free space and the order that is required for filling data pages, free space that is present in the database is not always available for use.

**Free Space Reclamation**
The version store is the area of the database that manages version control. When a transaction is committed to the database, a cleanup process returns space that is freed by modify and delete transactions to the database. For each modify or delete operation, the existing version of the record is written to the version store so that the database maintains a copy of the old version until the new version is written to the database. After the transaction is committed to the database, any space that is freed from deleted records and long values is returned to the table or index that owns the space. Until the change is committed to the database, requests for the object continue to access the old version. If the transaction is rolled back, the version store record is used to undo the transaction.

The version store has a size limit that is the lesser of the following: one-fourth of total random access memory (RAM) or 100 MB.

Because most domain controllers have more than 400 MB of RAM, the most common version store size is the maximum size of 100 MB. If too many large changes or deletions occur simultaneously, it is possible for the version store to run out of processing space. In this event, cleanup of free space is suspended temporarily. On domain controllers running Windows 2000 Server, the most common cause of version store overload is large-scale bulk deletions.

**Bulk deletions and database growth in Windows 2000**
Delete operations are the most CPU-intensive operations that the version store processes. On domain controllers running Windows 2000 Server, bulk deletions, such as the deletion of an entire tree of objects at one time, can cause a temporary condition in which free space cannot be returned to the database in a timely fashion because the cleanup process cannot keep up with the deletions. Event ID 602 is logged in the Directory Services event log to indicate this condition.

During the time that pages are being skipped by the cleanup process, free space is not released to the database, and space is not reclaimed until the next scheduled online defragmentation occurs. In the meantime, processing requirements can cause the database to grow. In particular, when bulk deletions or other bulk changes coincide with database additions, significant growth can occur. In addition, space from the deletion of long values is not returned to the database by online defragmentation. As a result of these conditions, the directory database on domain controllers running Windows 2000 Server can actually increase in size following a bulk deletion.

On domain controllers running Windows Server 2003 or later, the effects of these conditions are greatly reduced by improvements in version store cleanup and online defragmentation. However, if event ID 602 is logged in the Directory Services event log, running online defragmentation manually can alleviate the problem. On domain controllers running

Windows 2000 Server, the only way to prompt online defragmentation is to change the garbage collection interval to the minimum value of one hour to force garbage collection and online defragmentation to occur as soon as possible.

**Improved space processing in Windows Server 2003 and Later**
Two improvements in the Windows Server 2003 and later processing of free space eliminate the database growth problems that can result from large-scale bulk deletions:

- The threshold at which the database begins skipping cleanup operations is increased from 5 percent to 90 percent.
- Space is reclaimed from long-value deletions.

The threshold of maximum pages that can be processed by the version store is the limiting factor in whether the cleanup process can keep pace with deletions. The version store cleanup process can take place only as long as the version store has sufficient space. With a maximum version store size of 100 MB, only 5 MB (5 percent) is available in Windows 2000 Server, and this low threshold is responsible for early suspension of the cleanup process. The threshold of 90 MB (90 percent) in Windows Server 2003 and later eliminates this problem. For this reason, large-scale bulk deletions that can be problematic on domain controllers running Windows 2000 Server present no significant growth concerns on domain controllers running Windows Server 2003 or later.

In addition, online defragmentation on domain controllers running Windows Server 2003 or later returns the space that is freed by long values to the long-value table, which further optimizes the availability of space in the database.

**Bulk Changes and Inheritable Security Descriptors**
In addition to bulk deletions, large-scale bulk changes can cause database growth, particularly on domain controllers running Windows 2000 Server. The most common cause of database growth on domain controllers running Windows 2000 Server is the propagation of inherited security descriptor changes in a large directory tree.

Every object in Active Directory has a security descriptor attribute that contains the identification and authorization values that apply to that object. Authorization information takes the form of discretionary access control lists (DACLs) (access information) and system access control lists (SACLs) (auditing information). Because every object has a security descriptor, a security descriptor change that affects every object in a large tree (for example, an entire domain) can significantly affect database size.

**How Inheritable Security Descriptors Are Written**
The value of the security descriptor attribute changes when a new permission, right, or auditing setting is added or an existing permission, right, or auditing setting is removed from an object's DACL or SACL by adding, removing, or modifying an ACE. When an inheritable ACE is applied to an object in a directory tree, you can specify the type of object that is affected by the ACE. However, the change is written to every object in the tree, not just to the object type to which the permission, right, or auditing setting applies. The ACE applies to all objects, but it is effective only on objects whose object type matches the object GUID that is specified in the ACE.

This system ensures that the appropriate security descriptor can always be inherited from a parent object. For example, suppose you apply an inheritable ACE at the domain level that applies to all computers. Computers exist in container objects, to which the ACE does not directly apply. However, when you create a computer object in the container, the computer object receives its inherited security descriptor from its parent container. Every object that is within the scope of the inheritance receives the update to the security descriptor.

**Security Descriptors in Windows 2000 Server**
On domain controllers running Windows 2000 Server, the full security descriptor value is stored for every object. The average size of a security descriptor is 5 KB, and as a result security descriptors can account for 40 percent or more of database size. On domain controllers running Windows 2000 Server, an ACL change that is inherited by a large directory tree (for example, a domain) that contains hundreds of thousands or millions of objects can cause significant database growth. For example, suppose you add 30 ACEs to an inheritable SACL or DACL on the domain object of a domain that has 1,000,000 objects. At an estimated 52 bytes per ACE, each object grows by approximately 1.5 KB. The propagation of this change to 1 million objects results in database growth of 1.4 GB. In the worst-case scenario, the database grows beyond its storage limits, and Active Directory cannot function.

You can use Active Directory administrative tools to manage security settings on Active Directory objects to a high degree of specificity. Significant and unexpected changes to security descriptors can occur during the management of access control and auditing if an administrator does not understand the effects of the changes. The UI provides the ability to select general ACE categories that have prepopulated permissions, but the UI also provides the ability to modify the properties of a general ACE. The effect of editing an existing general ACE on a DACL or SACL by removing one of the properties is the replacement of a single ACE with multiple ACEs, one for each of the permissions that is selected in the modified ACE. On domain controllers running Windows 2000 Server, such a change can cause significant database growth if the modified ACE is inherited to a large tree of objects. For example, if directory services access auditing is enabled and an administrator edits the access to specify

only a subset of the activity and objects to be audited, the result is propagation of a separate ACE for each selection to every object in the tree. Given the default settings for the Everyone group in Windows 2000 Server, such a change can result in significant database growth.

Default settings in Windows Server 2003 and later, and the storage of only unique security descriptors greatly decrease the likelihood that such a change can cause unmanageable growth.

**Single-Instance Security Descriptors in Windows Server 2003 and Later**
On domain controllers running Windows Server 2003 or later, the potential for database growth as a result of the volume of changes that might be caused by security descriptor inheritance is mitigated by a new method of storing single instances of unique security descriptors.

On domain controllers running Windows Server 2003 or later, a unique security descriptor is stored only once rather than being stored for every object that inherits it. For every object that inherits the security descriptor or otherwise stores an identical security descriptor, only a pointer is stored. This change eliminates data redundancy and the database growth that can result from changes to inheritable ACEs.

The object itself stores an 8-byte pointer to the appropriate unique security descriptor. Unique security descriptors are stored in memory in a relatively small table. In addition, the information that is required to map these pointers to the respective security descriptor is cached on domain controllers. By using this cached information and the in-memory table of unique security descriptors, domain controllers can match pointers to security descriptors without having to access the database itself. Therefore, checks to establish whether a new security descriptor matches an existing one have no significant impact on domain controller performance.

### Note

- When you upgrade a domain controller from Windows 2000 Server to Windows Server 2003, the conversion of duplicate security descriptors to pointers can result in +40 percent of the database size being converted to unused space. Performing offline defragmentation after the upgrade decreases the database size by removing that unused space.

**Effect of Object Ownership on Unique Security Descriptors**
Having different owners causes objects that would otherwise have the same security descriptor to require several unique security descriptors. For example, if a member of the Domain Admins group creates a container object and sets permissions on the container that are inherited by all child objects, all child objects have the same security descriptor. In this case, only one security descriptor is actually stored in Active Directory; the rest of the objects simply reference that unique security descriptor. However, if objects that are created in the parent container are created by individuals who are not members of the Domain Admins group, those objects have security descriptors that are different from the security descriptors of the objects that are created by members of the Domain Admins group.

Despite the fact that having different object owners requires unique security descriptors, most Active Directory databases contain a relatively small number of user objects in comparison to the number of nonuser objects. In addition, most objects are created by the system and not by individual users. Because duplication of security descriptors far outweighs the potential for different object owners, the positive effect of single-instance security descriptors is not usually lessened to a noticeable degree by multiple object owners.

**Storage and Removal of Object Deletions**
When data is deleted from Active Directory, the data cannot simply disappear from the directory because the deletion must be replicated. Therefore, instead of deleting an object physically from the database, the directory service removes most of the attributes and then tags the object as a tombstone by setting the **isDeleted** attribute value to TRUE, which means that the object has been logically deleted from the directory but not yet completely removed. Tombstones are replicated to communicate object deletions. The **isDeleted** attribute value alerts replication partners that the object has been deleted. Objects that are identified as tombstones are moved to the hidden Deleted Objects container of their respective directory partition. Tombstones remain in the directory for a default period of 60 days, which is referred to as the tombstone lifetime.

**Garbage Collection: Permanent Removal of Expired Tombstones**
Garbage collection is a housekeeping process that runs on every domain controller to permanently remove expired tombstones from the directory database. Although they represent deleted objects, tombstones take up space in every directory partition replica. Eventually, the tombstones themselves must be deleted to keep the directory database from growing without limit. At regular intervals, objects that are no longer needed by the directory service are deleted as “garbage.”

The garbage collection process:

- Deletes tombstones.

- Defragments the database file to compact data and optimize free space (triggers online defragmentation).

Garbage collection runs independently on each domain controller. When the garbage collection process occurs, the process finds the set of tombstones whose originating deletion occurred more than a tombstone lifetime ago, and then it deletes each tombstone in that set.

Two attributes of the Directory Service object (nTDSService) in the configuration container (CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration, DC=*forestRootDomain*) control how garbage collection runs and what it removes, as follows:

- The tombstone lifetime determines the amount of time that a deleted object lives as a tombstone in the directory before being collected as garbage. This amount of time is set in the **tombstoneLifetime** attribute. The minimum setting is 2 days. The default setting for this attribute varies according to operating system and forest creation or upgrade, as follows:

    - Forests that were created on a domain controller running Windows Server 2003 with no service pack installed or any version of Windows 2000 Server: 60 days.

    - Forests that were upgraded from Windows Server 2003 with no service pack installed or any version of Windows 2000 Server: 60 days.

    - Forests that were created on a domain controller running Windows Server 2003 with Service Pack 1 (SP1) or later: 180 days.

    **Note**

    - You cannot purge tombstones before the expiration of the tombstone lifetime.

- The garbage collection interval determines how often a domain controller examines its database for expired tombstones that can be collected. This interval is set in the **garbageCollPeriod** attribute. The default setting is 12 hours, and the minimum setting is 1 hour.

    **Note**

    - The default value for each of these two attributes applies if the attribute is not set (which is the initial state of the system). The minimum value applies if the attribute is set to a value that is below the minimum (that is, if the minimum is not declared in the schema).

The maximum garbage collection interval is one-third of the tombstone lifetime (in hours). For example, if you set **tombstoneLifetime** to 30 days and **garbageCollPeriod** to 300 hours, the actual garbage collection period is only 10 days or 240 hours.

**Tombstone Lifetime and Active Directory Backup and Restore**

Active Directory does not allow a restore from a directory backup that is older than the tombstone lifetime. A restore from backup creates a directory partition replica that has not performed replication since the time of backup (or earlier). If the backup is taken more than a tombstone lifetime before the restore, objects that are deleted in the meantime have no tombstones; therefore, a new directory partition replica that is created by the restore operation never receives these deletions. For this reason, a restore procedure will not restore a backup that is taken more than one tombstone lifetime before the time of the restore. Therefore, it is a recommended best practice to back up Active Directory at least twice during a tombstone lifetime. On domain controllers running Windows Server 2003 with SP1 and later, event ID 2089 in the Directory Service event log provides notification when any directory partition, including application directory partitions and Active Directory Application Mode (ADAM) partitions, has not been backed up since the beginning of a specified time period. By default, this period is half the tombstone lifetime interval, as defined in the registry entry **Backup latency interval (days)** in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics**.

It is likewise important that the tombstone lifetime be substantially longer than the expected replication latency. The default setting of 60 days or 180 days for the tombstone lifetime is generous enough to accommodate unforeseen circumstances. However, monitoring domain controller operation is essential for ensuring that a domain controller does not remain offline without detection.

Because the expiration of a tombstone lifetime is based on the time when an object is deleted logically — rather than the time when a particular server receives that tombstone through replication — an object's tombstone is collected as garbage on all servers at approximately the same time. If the tombstone has not yet replicated to a particular server, that server never records the deletion. A domain controller that becomes outdated by being disconnected from the replication topology for longer than a tombstone lifetime cannot be restored from backup. However, an object that remains on an outdated domain controller is retained if the domain controller is reconnected to the replication topology. An object that remains after its tombstone is deleted is called a lingering object. Such objects create inconsistency in Active Directory and should be removed. For information about removing lingering objects, see "Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042)" in "Troubleshooting Active Directory Replication Problems" in the Windows Server 2003 Operations Guide at http://go.microsoft.com/fwlink/?LinkId=44131.

**Garbage Collection Scheduling Enhancements**

The process for completing garbage collection has changed in Windows Server 2003 and later to improve storage conditions in the directory database. Garbage collection removes a maximum of 5,000 objects per pass to avoid indefinitely delaying other directory service tasks. However, the rate at which remaining tombstones are deleted when more than 5,000 tombstones have expired has increased from Windows 2000 Server to Windows Server 2003 or later, as follows:

- Windows 2000 Server: If collection stops because of the 5,000-object limit (rather than by running out of objects to collect), the next garbage collection pass is scheduled for half the normal garbage collection interval (by default, every 6 hours instead of 12 hours). Garbage collection continues running at this accelerated pace until all objects have been collected.

- Windows Server 2003 and later: Rather than waiting a set time to remove a subsequent set of 5,000 tombstones, a domain controller continues deleting tombstones according to CPU availability. If no other process is using the CPU, garbage collection proceeds. Removing tombstones in this way keeps the database size from increasing inordinately as a result of the inability of garbage collection to fully complete removal of all tombstones during a garbage collection interval.

**Garbage Collection and Linked Attributes**

When an object that has a linked attribute is deleted, although the object itself becomes a tombstone (for example, when a user object is deleted), the referent object (for example, a group object to which the user belonged) does not reference the tombstone of the deleted object. Rather, the group-user link value is simply deleted from the database, and the group object shows no evidence of the former user's membership.

However, when an object to which a link refers is removed as a referent (for example, when a user is removed from a group), the user value is treated differently by Active Directory in Windows 2000 Server and in Windows Server 2003 and later. In Windows 2000 Server, the entire group **member** attribute is replicated. In Windows Server 2003 and later, the link value for the removed user is marked "absent" and then replicated, much like a tombstone.

**Reanimating Tombstones — Restoring Individual Objects Online**

The Windows Server 2003 and later directory database supports an LDAP API that reanimates the tombstone of a single object (undeletes the object) to avoid the necessity for an offline restore process in the event that an object is deleted unintentionally. This API is available for creating applications to restore the attributes that are preserved on tombstones, which include the object SID, GUID, and security descriptor, as well as any indexed attributes.

| **Note** |
| --- |
| When the deletion is performed on a domain controller that is running Windows Server 2003 with SP1 or later, the **sIDHistory** attribute is also retained. |

Only attributes that are retained on the tombstone are restored; all other data must be recreated. Therefore, to restore an entire deleted container or a set of multiple objects, authoritative restore is still the best option.

**Stale Account Detection**

Stale account detection is required so that unused computer and user accounts can be removed from Active Directory. On domain controllers running Windows Server 2003 and later, these accounts can be identified by the value in the **lastLogonTimeStamp** attribute of user and computer objects. This attribute identifies the last time that a user or computer successfully logged on to a domain. The attribute value is replicated to all domain controllers in the domain. **lastLogonTimeStamp** is new in Windows Server 2003, and its use is enhanced in Windows Server 2003 with SP1.

**Scenarios that generate stale accounts**

The following common scenarios result in the proliferation of computer and user accounts that are often abandoned by the

users who create the accounts:

- Delegated computer account creation, which is common in corporate environments, provides the freedom to create multiple accounts but no requirement for deleting them.
- Web applications allow external partners to create and manage user accounts that are provided by the host company.
- Employees leave the company.

**Detection of successful password verification**

In earlier versions of Windows, the **pwdLstSet** attribute was used to identify stale accounts. However, many conditions interfere with the accuracy of this method, such as users who are on sabbatical or otherwise temporarily not using an account. In addition, password time limits can vary throughout an organization. The most efficient way to identify an account that is stale is by evaluating the last time that the user successfully logged on to the domain, rather than the last time the password was reset.

The introduction of the **lastLogonTimeStamp** attribute in Windows Server 2003 and later provides the ability to mark an account each time that the system is asked to validate the account password (that is, at logon) for NTLM and Kerberos authentication. These two security protocols update the **lastLogonTimeStamp** attribute during successful authentication.

**Implementation in the domain**

Updates to **lastLogonTimeStamp** are replicated in the domain directory partition. Therefore, the value of this attribute can be checked on any domain controller in the domain. To minimize the effect of domain-wide replication every time that a user or computer logs on, **lastLogonTimeStamp** is updated only periodically. The period of update is configurable, as follows:

- Object: DC=*DomainName*
- Attribute: **msDS-LogonTimeSyncInterval**
- Default value: 14 days

Because the **lastLogonTimeStamp** attribute is available only when the Windows Server 2003 and later schema is in effect on all domain controllers in the domain, the feature is activated when the domain functional level is raised to Windows Server 2003 or higher. After the domain functional level increase, the first time that a successful logon occurs, the **lastLogonTimeStamp** attribute is activated on the user or computer object. This method of activation can result in loopholes for certain accounts, as follows:

- Accounts that are created before the domain functional level increase but that are not used subsequently
- Accounts that are created after the domain functional level increase but that are not used

In these cases, **lastLogonTimeStamp** is not present on the account object. To identify these accounts, you can query for user and computer objects that have no **lastLogonTimeStamp** attribute and that have a value in the **whenCreated** attribute that is earlier than or equal to a specified date.

For example, suppose a user account is created on 1/1/2005. The user leaves the company on 5/1/2005. The domain functional level is raised on 6/1/2005. On 9/1/2005, you query for accounts that do not have the **lastLogonTimeStamp** attribute. The query returns the account of the user who left the company before the raising of the functional level. However, you need to know if the account was created recently and has not yet been used or if the account was created before the functional level increase and is legitimately out of use. To answer this question, you query for accounts that do not have the **lastLogonTimeStamp** and that have a **whenCreated** value of 8/15/2005 or earlier. You select this date on the assumption that any account that was created within the last two weeks will have been used and any account that was created between June 1 and August 15 will certainly have been used. Accounts that are returned by the query are therefore legitimately out of use and can be deleted.

**Randomization mechanism for initial use of lastLogonTimeStamp**

To avoid network overloads due to thousands of concurrent **lastLogonTimeStamp** updates when the domain functional level is raised, a five-day window is used to calculate whether the **lastLogonTimeStamp** value should be updated. The following randomization mechanism is applied to control updates after the domain functional level is raised:

1. At the time that the functional level is raised, the default value of 14 days for **msDS-LogonTimeSyncInterval** is used, regardless of whether a different value is set in the attribute.

2. The value in **lastLogonTimeStamp** is used to determine the number of days since the last valid logon of the account.

3. A random percentage of the value 5 is taken and then subtracted from the value in **msDS-LogonTimeSyncInterval**.

4. If the result is greater than or equal to the value in **lastLogonTimeStamp**, the value in **lastLogonTimeStamp** is updated.

For example, suppose that the value in **lastLogonTimeStamp** indicates 12 days since the last successful logon. Suppose also that the random percentage is 80 percent of 5 = 4. With the default value of 14 days in effect for **msDS-LogonTimeSyncInterval**, the calculation is 14 - 4 = 10. Because 12 > 10, **lastLogonTimeStamp** is updated. In this example, in all cases where the value is less than 10, **lastLogonTimeStamp** is not updated until the first successful logon that occurs after the **msDS-LogonTimeSyncInterval** value is reached.

| **Note** |
| --- |
| If the value in **msDS-LogonTimeSyncInterval** is set to **0**, the stale account detection feature is disabled and **lastLogonTimeStamp** is not updated. |

**Security protocols that update lastLogonTimeStamp in Windows Server 2003**

In the initial implementation of **lastLogonTimeStamp** in Windows Server 2003, **lastLogonTimeStamp** is updated by the following authentication protocols:

- NTLM
- Kerberos

Security protocols notify the directory database of a successful logon so that the account can be marked appropriately in the **lastLogon** and **lastLogonTimeStamp** attributes. Therefore, to use stale account detection effectively, it is important that administrators are aware of the security protocols that are in use in their environments.

For example, in Windows Server 2003 environments, the following security protocols request that the system validate a user password:

- Secure channel setup (computer accounts only)
- NTLM authentication (both interactive and network)
- Kerberos authentication
- Digest authentication
- Third-party Security Support Provider Interfaces (SSPIs) when using the LSA APIs

The following security protocols can reference an account for a security operation:

- Schannel, when mapped to an account
- Kerberos Service-for-User (S4U) logons
- AuthZ APIs in non-S4U mode

The obvious problem is that some accounts might be authenticated without updating the **lastLogonTimeStamp**, and if this detection method is relied on for stale account cleanup, accounts might be deleted that are still in use.

For example, the following issues occur in Windows Server 2003 implementations:

- The NTLM authentication protocol does not update the **lastLogonTimeStamp** attribute for all network logons. NTLM is therefore not reliable for stale account detection.

- The Kerberos S4U authentication protocol, which is used for Web single sign on (SSO) provisioning, does not update **lastLogonTimeStamp**. Specifically, when extranet users log on to servers running Microsoft Internet Information Services (IIS), Active Directory user accounts are mapped to certificates that are trusted by the IIS server. These certificates are distributed to the client computer so that users do not have to specify credentials. (That is, they are not directly authenticated by a security protocol.)

#### lastLogonTimeStamp improvements in Windows Server 2003 SP1

Windows Server 2003 SP1 corrects the problems of some protocols not updating **lastLogonTimeStamp**, as follows:

- The inconsistency of NTLM updates of **lastLogonTimeStamp** is corrected.

- Updates to the Key Distribution Center (KDC) ensure that **lastLogonTimeStamp** is updated in Active Directory when a user logs on to a protected IIS Web site.

#### Scripting stale account detection

You can implement a scripting solution for managing stale account detection and cleanup. For more information about **lastLogonTimeStamp** and instructions for creating scripts, see "Dandelions, VCR Clocks, and Last Logon Times: These are a Few of Our Least Favorite Things" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=47965.

#### Database Defragmentation

As described earlier in this section, the database system uses the quickest way to fill database pages. Although this system is efficient in searching and updating the database quickly, it does not make the most efficient use of space in the database. As write operations occur and data pages fill, a certain amount of space often remains in a page that cannot be used by other transactions because the space is not large enough. In addition, space that is freed by version store cleanup is returned to the owning indexes or tables in noncontiguous chunks. Therefore, free space that exists in the database is not always available for use because it is fragmented.

Defragmentation rearranges how the data is stored in the database, compressing the data and making free space available in contiguous pages. Free space is managed automatically during database defragmentation.

Database defragmentation can take place online (while the computer is running as a domain controller) or offline (while the computer is running in Directory Services Restore Mode). Defragmentation has different effects on database size, depending on whether defragmentation is run while the domain controller is online or offline:

- Online defragmentation: Occurs automatically by default. Online defragmentation compacts data and optimizes free space for database use, thereby maintaining the file size of the database.

- Offline defragmentation: Occurs only manually in Directory Services Restore Mode. Offline defragmentation compacts data and returns free space to the file system, thereby decreasing the file size of the database.

#### Online Defragmentation

During online defragmentation, space in partially filled data pages is reorganized to consolidate it into full 8-KB pages. In addition, any pages that have been skipped by version store cleanup are reclaimed for use by the database. Online defragmentation returns a freed page directly to the part of the tree (database, object table, index, or long-value table) that owns the page that is being freed. When a sufficient number of contiguous free pages (approximately 16) exist in a given child tree, the space is released to the parent tree.

### Note

- On domain controllers running Windows 2000 Server, online defragmentation does not reclaim free space that is deleted from long-value tables.

#### Automatic online defragmentation

Windows 2000 Server and later based domain controllers perform online defragmentation automatically. The ESE invokes online defragmentation after each garbage collection, which occurs every 12 hours by default.

**Manual online defragmentation**
On domain controllers running Windows Server 2003 and later, you can prompt online defragmentation to run manually. Although online defragmentation occurs automatically following garbage collection and the garbage collection interval is known and configurable, the duration of the garbage collection process is variable, depending on how many tombstones have expired and the load requirements of the domain controller at the time of garbage collection. The variable nature of garbage collection can cause unpredictable performance between the time garbage collection begins and online defragmentation completes.

By performing online defragmentation during off-peak hours, you can ensure more-even directory performance during peak hours. To control when online defragmentation occurs, you can configure domain controllers running Windows Server 2003 and later as follows:

- To start online defragmentation manually at any time, irrespective of garbage collection intervals, you can add the operational attribute **doOnlineDefrag** to the rootDSE object. The value of this attribute dictates the duration (in seconds) for which online defragmentation runs. After setting this attribute value, you can create a script that takes the duration as input so that online defragmentation can be triggered by the script at any time for any duration of runtime.

- If you want online defragmentation to always occur only on demand, you can configure the registry so that automatic online defragmentation is disabled.

  **Note**

  - The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

You do not have to disable automatic online defragmentation to manually start online defragmentation. After you add the **doOnlineDefrag** attribute, you can start manual online defragmentation at any time that conditions warrant it, while allowing automatic online defragmentation to continue as well.

**Offline Defragmentation**
Offline defragmentation rebuilds the directory database and releases unneeded space back to the file system. Offline defragmentation must be performed in Directory Services Restore Mode, which restarts the computer as a stand-alone server, that is, as a computer that is not acting as a domain controller. In Directory Services Restore Mode, you can use the **ntdsutil** command-line tool to defragment the Ntds.dit file. Offline defragmentation produces a defragmented version of the database file in a separate directory. You can archive the original Ntds.dit file and move the defragmented file into the current directory.

Only offline defragmentation provides a clear picture of the amount of space that is consumed by the database file. You can use offline defragmentation to test database growth by comparing the defragmented version of the file with the fragmented version. For example, on a newly installed domain controller, if you perform a bulk load of objects and then defragment the database file offline, the difference between the two files is the space that is occupied by the new objects. There is no standard or optimal value for the percentage of free space that might exist in the database at any given time. If you are concerned about database size, set the DSA to log a message in the Directory Service event log during garbage collection that states how much disk space could be freed up through offline defragmentation.

**Linked Attribute Management**
Attributes of certain objects can be linked in the directory database to attributes of other objects to provide interobject references from one object back to the other object for either usability or administrative purposes. For example, a user object can be a member of a group, and a group can have users and groups as members. To define this relationship, the group object has a **member** attribute that can have multiple values of the distinguished names of every user, computer, or group that has membership in the group. Similarly, the user, computer, and group objects have a multivalued **memberOf** attribute that identifies all the groups to which a user, computer, or group belongs. These attributes are “linked” in the Active Directory database so that you can, for example, look at the **member** attribute on the object GroupA and determine that UserB is a member of GroupA. Likewise, you can look at the **memberOf** attribute on the object UserB and determine that UserB belongs to GroupA.

The GUID of an object unambiguously indicates which object the linked value applies to, even if the object has been renamed, deleted, or garbage-collected (permanently removed from the database). Individual values of multivalued attributes do not have an associated GUID. Instead, a link-value identifier serves the purpose of distinguishing existing values (including value

tombstones) from values that have been garbage-collected and then recreated. This identifier, called the link ID, remains fixed from the creation of the value until the value is garbage-collected. When the originating server creates a new value, an ID is assigned to the value. Linked values replicate with the GUID of the object that contains them.

**Effect of Moving Linked Objects**

Active Directory maintains referential integrity between objects that reference each other by distinguished name so that when one object is moved in the directory tree, the referencing object can track the referenced object. For example, an object UserB is a member of GroupA. If you move the object UserB from the CN=Users,DC=concorp,DC=contoso,DC=com container to the OU=Sales,DC=concorp,DC=contoso,DC=com OU, the distinguished name value for UserB in the **member** attribute on the GroupA object automatically changes as follows:

- Before the move: CN=UserB,CN=users,DC=concorp,DC=contoso,DC=com
- After the move: CN=UserB,OU=sales,DC=concorp,DC=contoso,DC=com

This change occurs as soon as the move occurs. In this way, linked attributes effectively share their values.

**Identifying Link Pairs**

The sharing of values by linked attributes is achieved when two attributes are marked in the schema as having the same link-pair identifier — one is marked as the forward link (it points to another object in the directory) and the other as the back-link (it points back to the object that has a forward link to it). Changing the forward link attribute on one object affects the back-link on the referenced object. In the current example, the **member** attribute is the forward link and the **memberOf** attribute points back to the value in **member**. For reasons that relate to security and replication, only the forward-linked attribute can be modified.

### Note

- The decision to link two objects must be made at the time that the objects are added to the schema.

To find all the objects that are members of GroupA, links are examined for all records in which the link pair is **member**/**memberOf** and the back-linked attribute identifies GroupA. The link pairs of those records provide the database identifiers of all the records (objects) that are members of GroupA.

The **member** and **memberOf** example uses a multivalued forward link and a multivalued back-link, but there is no requirement that the forward link be a multivalued link. For example, each group has one manager, but a manager can manage multiple groups. This relationship is identified in the database by the **manager**/**directReports** link pair, where **manager** is a single-valued attribute of a user object and **directReports** (common name Reports) is a multivalued attribute of a user object. The back-linked attribute must always be multivalued because it is impossible to restrict who creates links to various objects.

**Deleting Linked Objects**

When an object that is linked through a multivalued attribute is deleted, all of its linked attribute values are removed from the respective linked objects. In the preceding example, if a user manages multiple users and one of the users is deleted, the **directReports** multivalued attribute on the user object of the user that is the manager suddenly (and with no change to any replication-related metadata) loses a value. Similarly, if the user object for the manager is deleted, the value of the **manager** attribute on the user object of the user who was a direct report is suddenly blank. Nothing about the nondeleted object changes in either case, except that the attribute value is gone.

**Restoring Linked Objects and Their Back-Links**

When objects are deleted by mistake, you can restore them in Active Directory by performing an authoritative restore of the deleted objects. You perform authoritative restore by restoring the domain controller from its latest backup and then marking the deleted objects as authoritative so that they are not removed by replication.

However, when an object that has back-links is restored on a domain controller running Windows 2000 Server, its back-links are not restored as part of the authoritative restore process. For example, when user objects or group objects are deleted and then authoritatively restored, the **memberOf** values must be restored manually. You must add a restored user or group back to the groups to which it belonged at the time that its account was deleted.

On domain controllers running Windows Server 2003 or later in a forest that has a forest functional level of Windows Server 2003 or higher, back-links are restored automatically for authoritatively restored objects under the following conditions:

- The link was added after the forest functional level was raised to Windows Server 2003 or higher (linked-value replication was in effect).

- Replication of a restored user object precedes replication of a restored group to which it belongs when the user and group are both authoritatively restored at the same time.

If the domain controller on which you are authoritatively restoring objects is running Windows Server 2003 with SP1 or later, you can use the Windows Server 2003 SP1 or later version of Ntdsutil to restore back-links for any authoritatively restored objects that have links, including objects in different domains. For more information about how to restore back-links on domain controllers running Windows Server 2003 with SP1 or later, see "Performing an Authoritative Restore of Active Directory Objects" in "Administering Active Directory Backup and Restore" in the Windows Server 2003 Operations Guide at http://go.microsoft.com/fwlink/?LinkId=44194.

For more information about replication of deleted linked objects, see "Active Directory Replication Model Technical Reference".

**Phantom Records for Interdomain Object References**
An attribute that has a distinguished name as a value references (points to) the named object. When the referenced object does not actually exist in the local directory database because it is in a different domain, a placeholder record called a phantom is created in that database as the object reference. Because there is a reference to it, the referenced object must exist in some form, either as the full object (if the domain controller stores the respective domain directory partition) or as an object reference (when the domain controller does not store that domain).

A phantom record contains the GUID and the distinguished name of the object that is being referenced. In the case of references to security principals, the phantom also contains the SID. A common example of a distinguished name value that references an object in a different domain is the **nCName** attribute of a cross-reference (crossRef) object. Every domain controller stores a cross-reference object (in the configuration directory partition) for every other domain directory partition in the forest. Therefore, the **nCName** attribute value for every cross-reference object necessarily references a phantom in the local directory database.

**Infrastructure Master and Phantom Records**
The infrastructure master is a single domain controller in each domain that tracks name changes of referenced objects and updates the references on the referencing object. When a referenced object is moved to a different domain (which effectively renames the object), the infrastructure master updates the distinguished name of the phantom. (The infrastructure master finds phantom records by using a database index that is created only on domain controllers that hold the infrastructure operations master role.) When the reference count of the phantom falls to zero (no objects are referencing the object that the phantom represents), garbage collection on each domain controller removes the phantom.

### Note

- Because objects can reference objects in different domains, the infrastructure operations master role is not compatible with global catalog server status if there is more than one domain in the forest. If a global catalog server holds the infrastructure operations master role, phantom records are never created because the referenced object is always located in the directory database on the global catalog server.

**Ownership Quotas on Directory Objects**
On domain controllers running Windows Server 2003 or later, you can specify quotas that limit the number of objects that a security principal (user, group, computer, or service) can own in a domain, configuration, or application directory partition. By default, the security principal (or administrative group in some cases) that creates an object is the object owner, although ownership can be transferred. With Active Directory quotas, no one can create unlimited numbers of objects in a directory partition, which is a method that can be used to launch denial-of-service attacks.

### Note

- The term "user" in this discussion of security principals indicates both the **user** and **inetOrgPerson** classes of objects.

By default, quotas are not set, so there are no limits to the number of objects that can be owned by any security principal. For each target directory partition, you can set different limits for different security principals or you can set limits that apply to all security principals, or you can do both.

Quotas can be set per directory partition, except for the schema directory partition, which does not support quotas. You can set quotas to apply to security principals for a directory partition by using the command line. Use the following commands to manage quotas:

- **dsadd quota** to set quotas for a directory partition

- **dsmod quota** to modify existing quotas
- **dsmod partition** to modify default quota settings for a partition
- **dsquery quota** to search for quota assignments and quota consumption

For more information about directory quotas and how to use **dsadd**, **dsmod**, and **dsquery** commands to manage directory quotas, see "Data Store Tools and Settings".

**Quota Assignment Objects**

Quota assignments for security principals are represented in a directory partition as instances of the object class **msDS-QuotaControl**. Each object in this class has the following attributes:

- **Common-Name**: The relative distinguished name (also known as RDN) of the security principal object. This mandatory value should be a friendly name (or **sAMAccountName**) of the security principal whose quota is being specified.
- **msDS-QuotaTrustee**: The SID of the security principal (user, computer, or group) for which the quota is being assigned.
- **msDS-QuotaAmount**: The mandatory value of the assigned quota (number of objects owned in the database) for the security principal. A value of `-1` for this attribute denotes an unlimited quota.
- **Flags**: This optional attribute is reserved for future use.

Instances of **msDS-QuotaControl** objects are stored in a well-known container called NTDS Quotas in the directory partition to which the quota applies.

**Quota Containers**

The NTDS Quotas container (class **msDS-QuotaContainer**) is a special system container that the quota system uses. The NTDS Quotas container has the following characteristics:

- It is a child object of the root of a parent directory partition (including domain, configuration, and application directory partitions).
- It cannot be deleted, but it can be renamed.
- It is tracked by the parent container through the **wellKnownObjects** attribute on the parent container.

Each directory partition has exactly one well-known NTDS Quotas container that stores explicit quota assignment objects for the partition.

Two attributes of this well-known container specify a default quota value and how quotas are computed for that directory partition:

- **msDS-DefaultQuota**: A default quota that applies to any security principal that creates objects in the directory partition and for which no explicit quota assignment exists. If no value is set for this attribute or if the attribute has the value -1, there is no default quota for the directory partition. In this case, security principals that do not have explicit quotas assigned have the ability to create unlimited objects in that directory partition, subject to access control. The default value for **msDS-DefaultQuota** on a directory partition is unlimited (not set). To manage this value, use the **dsmod partition** command.
- **msDS-TombstoneQuotaFactor**: The percentage factor by which tombstone object count should be reduced for the purpose of quota accounting. By default, the setting is 100, which means that a tombstone (a deleted object) that was originally created by a security principal is equal to 1 object of that user's quota. If the percentage factor were set to 50, each tombstone would represent only half an object rather than an entire object. Likewise, a tombstone quota factor of 25 means that four tombstones are the equivalent of 1 owned object. In this way, the quota for a security principal is computed relative to objects and tombstones. To manage this value, use the **dsmod partition** command.

**Tracking Object Ownership**
To enforce compliance with Active Directory quotas, object ownership is tracked by each domain controller running Windows Server 2003 or later. After a domain controller is upgraded to Windows Server 2003 or later, the system determines the number of objects that are owned by each security principal and for each directory partition that the domain controller hosts, and the system generates a tracking table for internal object ownership to track quota consumption. Thereafter, when objects are created, deleted, or reanimated or their ownership is transferred in a directory partition that the domain controller hosts, the quota tracking data is suitably updated.

**Enforcing Quotas**
Quotas are enforced only on originating updates and only by domain controllers running Windows Server 2003 or later. When an update originates on a Windows Server 2003 or later domain controller, quotas are enforced. At the time of the operation, if the requestor is not exempt from quotas, the effective quota limit for the requestor is computed on the basis of the default quota setting for the directory partition or any explicit quota assignments. Further, the requestor's current quota consumption is computed using the object ownership tracking data and the tombstone quota factor. If the pending object transaction (create, delete, reanimate, or transfer of ownership) causes the quota limit to be exceeded, the transaction fails.

If an originating update occurs on a domain controller running Windows 2000 Server, quotas are not enforced. However, as Windows Server 2003 or later domain controllers receive replicated updates, regardless of their origin, they update object ownership tracking data for the respective security principal.

Thus, although Windows Server 2003 and later domain controllers can always update quota tracking data, quotas are effectively enforced at the domain level only when all domain controllers in the domain are running Windows Server 2003 or later and at the forest level (for the configuration directory partition) only when all domain controllers in the forest are running Windows Server 2003 or higher.

**Effective quota**
Multiple quota assignments can exist for a given security principal through an individual quota assignment or through quota assignments for one or more security groups of which the principal is a member, or both. When multiple quota assignments exist for a security principal, the effective quota is computed to be the maximum of the applicable quota assignments that are specified. For example, if a user belongs to two groups that have different quotas specified for the same directory partition, the higher of the two quota values applies to the user.

If no quota assignments are specified for a security principal or for any of the security groups of which the security principal is a member, the effective quota is the default quota that is specified for the directory partition in the **msDS-DefaultQuota** attribute on the NTDS Quotas container.

**Quota exemptions**
Members of the Domain Admins and Enterprise Admins groups are exempted from all quota-imposed limits.

**Using Quotas**
You can use quotas in combination with security groups to manage objects that are created in Active Directory. The following examples illustrate using a default quota, with exceptions to set explicit quota assignments for those groups that need the ability to create large numbers of objects.

**Managing DNS data in application directory partitions**
When you use Active Directory–integrated DNS, zone data can be stored in the domain directory partition, if you want DNS data to replicate to all domain controllers, or in application directory partitions, if you want DNS data to replicate only to domain controllers that are DNS servers.

To decrease domain-wide replication and avoid replicating zone data to the global catalog, you can take advantage of application directory partitions on DNS servers running Windows Server 2003 or later. By default, members of the Authenticated Users group have the ability to create resource records on DNS servers that are in the domain of the client computer. This ability is required to allow all computers to dynamically update DNS zone data. A typical authenticated user needs to register a maximum of 10 records in DNS. To ensure that malicious users or applications do not create inappropriate resource records, you can set a default quota on the application directory partitions ForestDnsZones and DomainDnsZones, which are replicated to every DNS server in the forest and domain, respectively. A default limit of 10 objects ensures that all computers can update DNS appropriately but cannot launch denial-of-service attacks.

**Explicit quotas for domain controllers and other servers**
The number of resource records that must be registered in DNS is significantly greater for domain controllers and for multihomed computers that register adapter-specific records for all of their connections. To accommodate the resource record requirements of atypical authenticated users such as domain controllers, multihomed computers, and multiadapter servers, you can create special groups for these specific types of computers. For example, you might create the Enterprise Domain Controllers group that has all domain controllers as members. To ensure that this group can create the maximum DNS resource records that might be required, set a quota of 300 to 400 objects for this group on the DomainDnsZones and

ForestDnsZones application directory partitions.

The same rule applies to servers such as virtual private network (VPN) or Web servers, which are required to register a large number of resource records in DNS. Depending on how many adapter-specific resource records your VPN or Web server must register, you can create a group for those servers and assign the appropriate quota to ensure that the Authenticated Users default quota does not apply.

**Explicit quotas for DHCP servers**
To allow DHCP servers to have essentially unlimited ability to create resource records when clients depend on DHCP-assigned IP addresses, you can create a group that contains all DHCP servers in the domain or forest and then set an explicit quota assignment for that group on the DomainDnsZones and ForestDnsZones application directory partitions, respectively. Set a quota that is appropriate to the number of clients for which your DHCP server registers records, multiplied by the upper limit of the number of records registered per computer. In this way, DHCP servers can update DNS with client resource records as needed.

**Managing print objects in domain directory partitions**
When Active Directory is used to publish shared printing resources, print servers require the ability to create potentially large numbers of objects in Active Directory. Print servers publish printers in Active Directory as print queue objects (class **printQueue**), which are child objects of the print server computer object and which contain a subset of the information that is stored on the print server for a printer. When print servers go offline for any reason, they must republish their print queues when they return to service. For this reason, print servers require the ability to create large numbers of objects in the domain directory partition.

To ensure that print servers can create appropriate numbers of objects in the domain, you can create a group that contains all print servers in the domain and then specify an explicit quota assignment on the domain that is higher than the default assignment.

**Network Ports Used by the Data Store**
The network ports that are used by the data store are listed in the following table.

### Port Assignments for the Data Store

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | None | 389 |
| LDAP SSL | None | 636 |
| RPC Endpoint Mapper | 135 | 135 |
| Global Catalog LDAP | None | 3268 |
| Global Catalog LDAP SSL | None | 3269 |

**Related Information**
The following resources contain additional information that is relevant to this section.

- Active Directory Replication Model Technical Reference
- Microsoft Platform SDK on MSDN for more information about LDAP (in “Lightweight Directory Access Protocol”) (http://go.microsoft.com/fwlink/?linkid=20540)
- Data Store Tools and Settings

# Data Store Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Data Store Tools
- Data Store Registry Entries
- Data Store Group Policy Settings
- Data Store WMI Classes
- Network Ports Used by the Data Store


This section contains information about the tools, registry entries, Group Policy settings, Windows Management Instrumentation (WMI) classes, and network ports that are associated with the data store.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

**Data Store Tools**

The following tools are associated with the data store.

**Adsiedit.msc: ADSI Edit**

**Category**

This tool ships with Support Tools for Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
| --- | --- |

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

ADSI Edit is a Microsoft Management Console (MMC) tool that you can use to view and modify directory objects.

To find more information about ADSI Edit, see Support Tools Help in Tools and Settings Collection.

**Csvde.exe: Csvde**

**Category**
This tool ships with Windows Server 2003.

**Version compatibility**

| **Can Be Run From** | **Can Be Run Against** |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Csvde to import and export data from Active Directory by using files that store data in the comma-separated value (CSV) file format standard. Csvde also supports batch operations that are based on CSV.

To find more information about Csvde, see "Command-Line References" in Tools and Settings Collection.

**Dsadd.exe: Dsadd**

**Category**
This tool ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsadd to add specific types of objects to the directory.

To find more information about Dsadd, see "Command-Line References" in Tools and Settings Collection.

**Dsget.exe: Dsget**

**Category**
This tool ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsget to display the selected properties of a specific object in the directory.

To find more information about Dsget, see "Command-Line References" in Tools and Settings Collection.

**Dsmod.exe: Dsmod**

**Category**
This tool ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsmod to modify an existing object of a specific type in the directory.

To find more information about Dsmod, see "Command-Line References" in Tools and Settings Collection.

**Dsmove.exe: Dsmove**

**Category**
This tool ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsmove to move a single object, within a domain, from its current location in the directory to a new location. You can also use Dsmove to rename a single object without moving it in the directory tree.

To find more information about Dsmove, see "Command-Line References" in Tools and Settings Collection.

**Dsquery.exe: Dsquery**

**Category**
This tool ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsquery to query Active Directory according to specified criteria.

To find more information about Dsquery, see "Command-Line References" in Tools and Settings Collection.

**Dsrm.exe: Dsrm**

**Category**
This tool ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsrm to delete an object of a specific type, or any general object, from the directory.

To find more information about Dsrm, see "Command-Line References" in Tools and Settings Collection.

**Ldifde.exe: Ldifde**

**Category**
This tool ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Ldifde to create, modify, and delete directory objects on domain controllers. You can also use Ldifde to extend the schema, export Active Directory user and group information to other applications or services, and populate Active Directory with data from other directory services.

To find more information about Ldifde, see "Command-Line References" in Tools and Settings Collection.

**Ldp.exe: Ldp**

**Category**
This tool ships with Support Tools for Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Ldp is a Lightweight Directory Access Protocol (LDAP) graphical user interface (GUI) tool that you can use to perform operations such as connect, bind, search, modify, add, and delete against any LDAP-compatible directory, such as Active Directory. You can also use Ldp to view objects that are stored in Active Directory, along with their metadata, for example, security descriptors and replication metadata.

You can use the online **dbdump** feature in Ldp to view values that are stored in the database while the domain controller is running. You can trigger **dbdump** by modifying the **dumpDatabase** attribute on the **rootDSE**.

To find more information about Ldp, see "Support Tools Help" in Tools and Settings Collection.

**Ntdsutil.exe: Ntdsutil**

**Category**
This tool ships with Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition |

You can use Ntdsutil to perform Active Directory database maintenance, manage and control single master operations, and remove metadata left behind by domain controllers that are removed from the network without being properly uninstalled. This tool is intended for use by experienced administrators.

To find more information about Ntdsutil, see "Command-Line References" in Tools and Settings Collection.

**Data Store Registry Entries**
The following registry entries are associated with the data store.

The registry entries under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics** control the logging level for the component or process that is specified in the entry name. The value for each entry is set to an integer from and including 0 (no logging) through 5 (most verbose logging).

The registry entries under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** control or contain information about the configuration of the data store.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics**
The following registry entries are located under
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics**.

**3 ExDS Interface Events**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that result from communication between Active Directory and Microsoft Exchange clients.

**4 MAPI Interface Events**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that result from communication between Active Directory and Microsoft Exchange clients.

**6 Garbage Collection**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are generated when objects that are marked for deletion are actually deleted.

**7 Internal Configuration**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of internal operations.

**8 Directory Access**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**

Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of read and write operations to directory objects from all sources.

**9 Internal Processing**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are related to internal directory service operations.

**11 Initialization/Termination**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are generated by starting and stopping Active Directory.

**12 Service Control**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of Active Directory service events.

**13 Name Resolution**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are generated by the resolution of addresses and Active Directory names.

**14 Backup**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are related to backing up Active Directory. Specifically, controls the logging of events that occur when Extensible Storage Engine (ESE) database records are read or written during backup.

**16 LDAP Interface Events**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are related to LDAP.

**18 Global Catalog**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are related to the global catalog.

**22 DS RPC Client**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are related to communication between Directory System Agents (DSAs). Examples of such communication include replication and the forwarding of look-ups to a global catalog. Examples of logged events include remote procedure call (RPC) errors, canceled calls, Domain Name System (DNS) resolution failures, and service principal name (SPN)–related operations.

**23 DS RPC Server**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are related to a DSA acting as an RPC server. A DSA acts as an RPC server, for example, during outbound replication, replication setup operations, cross-domain moves, and membership queries or when a client makes a look-up call.

**24 DS Schema**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls the logging of events that are related to schema errors and operations. Such errors and operations include schema additions, deletions, modifications, look-up errors, look-up failures, and caching errors.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**
The following registry entries are located under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

**Configuration NC**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Contains the distinguished name of the configuration directory partition.

**Database Backup Path**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Determines the directory that is used as the target directory when online backups of the directory database are performed.

**Database Log Files Path**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Determines the directory path that is used to store Active Directory log files.

**Database Logging/Recovery**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Controls a Microsoft Jet database engine parameter called *JET_paramRecovery* that determines whether database operations are logged.

**DS Drive Mappings**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Used to track local drive mapping names, so that the database file can be located if drive mappings are modified.

**DSA Database File**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server;

Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Determines the file that is used by the domain controller for storing Active Directory objects.

**DSA Working Directory**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Determines the working directory of Active Directory.

**Hierarchy Table Recalculation Interval (minutes)**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Determines how frequently the hierarchy table for the directory database is built. The hierarchy table is used only by the Messaging API (MAPI) interface.

**Ldapserverintegrity**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Determines whether connection integrity is required (by means of checksum-signed packets) for LDAP connections.

**Machine DN Name**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Contains the distinguished name of the computer on which the domain controller is running.

**Root Domain**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Contains the distinguished name of the root domain of the Active Directory forest to which the domain controller is connected.

**Schema Version**

**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows 2000 Server; Windows 2000 Advanced Server; Windows 2000 Datacenter Server; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Contains the schema version for which a particular operating system is configured.

**System Schema Version**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Domain controllers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

Contains the version of the schema at the time that a backup is taken. This value is used to prevent an incompatible schema version from being restored from backup.

**Data Store Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with the data store.

### Group Policy Settings Associated with the Data Store

| Group Policy Setting | Description |
|---|---|
| Audit Directory Services Access | When it is enabled, this Group Policy setting causes successful and failed directory access events to be logged in the Directory Service event log. |

**Data Store WMI Classes**
The following table lists and describes the WMI classes that are associated with the data store.

### WMI Classes Associated with the Data Store

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| **rootDSE** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

| | | |
|---|---|---|
| **DS_LDAP_Class_Containment** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |
| **DS_LDAP_Instance_Containment** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

For more information about these WMI classes, see "Mapping Active Directory to WMI" in the WMI SDK documentation on MSDN.

**Network Ports Used by the Data Store**

The network ports that are used by the data store are listed in the following table.

**Port Assignments for the Data Store**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | None | 389 |
| LDAP SSL | None | 636 |
| RPC Endpoint Mapper | 135 | 135 |
| Global Catalog LDAP | None | 3268 |
| Global Catalog LDAP SSL | None | 3269 |

# DNS Support for Active Directory Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

Active Directory is the directory service for the Windows Server 2003 operating system; DNS is the primary name resolution service for Windows Server 2003, and a core component of Windows Server 2003 TCP/IP networking.

Strictly speaking, DNS is not a component of Active Directory. However, knowledge of how Active Directory depends on DNS is necessary to understand how Active Directory components are able to perform their assigned operations. DNS is used by Active Directory for domain controller location and DNS naming is leveraged by Active Directory for domain naming. Because Active Directory depends on DNS for domain controller location and DNS influences Active Directory domain naming, DNS is discussed in this technical reference subject as a component of the Windows Server 2003 directory service solution.

**In this subject**


- What Is DNS Support for Active Directory?
- How DNS Support for Active Directory Works
- DNS Support for Active Directory Tools and Settings

# What Is DNS Support for Active Directory?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**What Is DNS Support for Active Directory?**
DNS is the primary name resolution service for Windows Server 2003 and later. Active Directory depends on DNS for domain controller location, and DNS influences Active Directory domain naming. Thus, to fully understand Active Directory, it helps to understand how DNS acts as an integral component in the design of Active Directory.

Active Directory requires:

- A name resolution service that enables network hosts and services to locate Active Directory domain controllers.
- A naming structure that enables an enterprise to reflect its organizational structure in the names of its directory service domains.

DNS provides Active Directory with both a name resolution service for domain controller location and a hierarchical design that Active Directory leverages to provide a naming convention that can reflect organizational structure.

Typically, a DNS domain namespace deployed to accommodate the Active Directory mirrors the Active Directory domain namespace. In cases where there is an existing DNS namespace prior to Active Directory deployment, the DNS namespace is typically partitioned for Active Directory, and a DNS subdomain and delegation for the Active Directory forest root is created. Additional DNS domain names are then added for each Active Directory child domain.

DNS data is used to support the location of Active Directory domain controllers also. During or after the creation of the DNS zones used to support Active Directory domains, the zones are populated with DNS resource records that enable network hosts and services to locate Active Directory domain controllers.

**Related Information**
The following resource contains additional information that is relevant to this section.

- How DNS Works

# How DNS Support for Active Directory Works

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**How DNS Support for Active Directory Works**

**In this section**


- DNS Support for Active Directory Architecture
- DNS Physical Structure in Support of Active Directory
- DNS Support for Active Directory Processes and Interactions
- Network Ports Used by DNS in Support of Active Directory
- Related Information


Active Directory uses DNS as its domain controller location mechanism and leverages the namespace design of DNS in the design of Active Directory domain names. As a result, DNS is positioned within the discoverability and logical structure components of Active Directory technology components.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

Typically, a Windows Server 2003 and later DNS namespace is deployed to mirror an Active Directory forest and domain infrastructure. In such a deployment, a partition of the DNS namespace is set aside for Active Directory, where a DNS domain name such as corp.contoso.com is used support the Active Directory forest root domain, and then subdomains of this name are created to suit additional Active Directory domains as needed.

**DNS Support for Active Directory Architecture**

Active Directory is dependent on DNS as a domain controller location mechanism and uses DNS domain naming conventions in the architecture of Active Directory domains. There are three components in the dependency of Active Directory on DNS:

- Domain controller locator (Locator)
- Active Directory domain names in DNS
- Active Directory DNS objects

**DNS Support for Active Directory Components**

| Component | Description |
| --- | --- |
| Domain controller locator (Locator) | The Windows Server 2003 or later domain controller locator, implemented in the Net Logon service, enables a client to locate a domain controller. The component contains the DNS–compatible and the Windows NT 4.0–compatible locators that provide interoperability in a mixed Windows Server 2003 or later – and Windows NT 4.0–based environment. |

| | |
|---|---|
| Active Directory domain names in DNS | Every Windows Server 2003 or later Active Directory domain has a DNS domain name (for example, contoso.com), and every Windows Server 2003 or later based computer has a DNS name (for example, win2kserver.contoso.com). Architecturally, domains and computers are represented both as objects in Active Directory and as nodes in DNS.<br><br>Because DNS domains and Active Directory domains share identical domain names, it is easy to confuse their roles. The two namespaces, although typically sharing an identical domain structure, store different data and, therefore, manage different objects:<br><br>• DNS stores zones and resource records, and Active Directory stores domains and domain objects. Both systems use a database to resolve names.<br><br>• DNS resolves domain names and computer names to resource records through requests received by DNS servers as DNS queries to the DNS database.<br><br>• Active Directory resolves domain object names to object records through requests that are received by domain controllers either as LDAP search requests or as modify requests to the Active Directory database.<br><br>Thus, the Active Directory domain computer account object is in a different namespace from the DNS host record that represents the same computer in the DNS zone. |
| Active Directory DNS objects | When DNS data is stored in Active Directory, each DNS zone is an Active Directory container object (class **dnsZone**). The dnsZone object contains a DNS node object (class **dnsNode**) for every unique name within that zone. These unique names include the variations assigned to a specific host computer when it functions, for example, as a primary domain controller or as a global catalog server. The dnsNode object has a dnsRecord multivalue attribute that contains a value for every resource record that is associated with an object's name.<br><br>For more information about Active Directory DNS objects, see How DNS Works. |

**Active Directory and DNS Domain Names**

Active Directory domains have two types of names: DNS names and NetBIOS names. In general, both names are visible to end users. The DNS names of Active Directory domains include two parts, a prefix and a suffix. The DNS prefix is the first label in the DNS name of the domain. The suffix is the name of the Active Directory forest root domain.

When a Windows NT 4.0 master user domain (MUD) is upgraded, the decision is made whether or not to use the current NetBIOS name of the domain as a DNS prefix. If the name is appropriate to represent the organizational structure of the enterprise and satisfies the prefix naming rules in the table below, the name is typically kept. In this case, the NetBIOS name of the domain is the same as the DNS prefix of the domain.

**Registered DNS Name Prefix Rules**

| Rule | Explanation |
|---|---|
| Select a prefix that is not likely to become outdated. | Avoid names such as a business line or operating system that might change in the future. Geographical names are recommended. |
| Select a prefix that includes Internet standard characters only. | A-Z, a-z, 0-9, and (-), but not entirely numeric. |
| Include 15 characters or less in the prefix. | If you choose a prefix length of 15 characters or less, then the NetBIOS name is the same as the prefix. |

If the current NetBIOS name of the domain is inappropriate to represent the organizational structure of the enterprise or if the current name fails to satisfy the prefix naming rules, a new prefix is used. In this case, the NetBIOS name of the domain will be different from the DNS prefix of the domain.

For each new Active Directory domain, a prefix that is appropriate for the organizational structure of the enterprise and that satisfies prefix naming rules is used. Typically, the NetBIOS name of the domain will be the same as the name of the DNS prefix.

The Active Directory forest root domain is also the name of the Active Directory forest. The forest root name is a DNS name that consists of a prefix and a suffix in the form of prefix.suffix. For example, an organization might have the forest root name corp.contoso.com. In this example, corp is the prefix and contoso.com is the suffix.

The suffix is selected from a list of existing DNS names on the network. For the prefix, a new name that has not been used on the network previously is selected. By attaching a new prefix to an existing suffix, a unique namespace is created. Creating a new namespace for Active Directory ensures that any existing DNS infrastructure does not need to be modified to accommodate Active Directory.

Typically, the DNS names that are used in the Active Directory namespace are registered with an Internet authority. Only registered names are guaranteed to be globally unique. If two organizations register the same DNS domain name, or if an organization merges with, acquires, or is acquired by another company that uses the same DNS name, then the two infrastructures cannot interact with one another.

Active Directory domain names should be fully qualified, such as corp.contoso.com or contoso.com. Commonly accepted top level domains for AD domains can consist of traditional three-letter top level domains such as .com, .biz, or .gov and two-letter country code domains such as .uk or .jp. For more information about top-level domain names and two-letter country code domains, see Root Zone Database.

Single-label names, such as contoso, can never be registered on the Internet. They require additional configuration, causing higher total cost of ownership (TCO). Furthermore, some Active Directory-enabled applications are incompatible with single-label domain names, or they have announced intentions to drop support for them in the future.

If you create an Active Directory domain with a single-label name, it can be very difficult to rename or migrate to a fully qualified domain name. Beginning with Windows Server 2008 R2, the Active Directory Installation Wizard does not allow creation of new domains with single label names. Both Windows Server 2008 and Windows Server 2008 R2 continue to allow installation of additional domain controllers to existing domains that have a single-label DNS name. In-place upgrades of domain controllers in a domain that has a single-label DNS name is also allowed, subject to the upgrade compatibility matrix (for example, and x86-based domain controller cannot be in-place upgraded to x64-based operating system).

For more information, see Microsoft Knowledge Base article 300684.

**DNS Physical Structure in Support of Active Directory**

The physical structure of Active Directory information in DNS is represented in DNS zones and resource records, which, in turn, are typically stored in Active Directory as Active Directory–integrated DNS zones. The DNS zones that support Active Directory domains can also be stored in standard, file-based, DNS zones. In addition, the DNS dynamic update protocol is utilized by Active Directory in order to make the registration of domain controller DNS resource records automatic.

The following diagram illustrates the physical storage of Active Directory domain information in DNS zones hosted on domain controllers, and in standard DNS zones hosted on member servers.

**DNS Physical Structure in Support of Active Directory**




The following table describes the physical components in the Active Directory/DNS physical structure.

**DNS Physical Structure in Support of Active Directory Components**

| Active Directory/DNS Physical Structure Component | Description |
|---|---|
| Active Directory/DNS physical structure requirements | The DNS server used to support Active Directory must support the SRV resource record and, ideally, the DNS dynamic update protocol. If the DNS server does not support the DNS dynamic update protocol, the SRV resource records required to locate Active Directory domain controllers can be added to DNS manually. |
| _msdcs DNS subdomain | The Microsoft-specific subdomain enables location of domain controllers that have specific roles in the Active Directory domain or forest. Resource records for the DNS root domain of a new Active Directory forest are stored in a _msdcs zone instead of a subdomain, and that zone is stored in the forest-wide application directory partition. |
| Domain controller SRV resource records registered in DNS | There are multiple DNS SRV resource records registered on behalf of domain controllers. A description of each resource record is below. |
| DNS application directory partitions | In Windows Server 2003 and later, the Active Directory–integrated DNS zones that support Active Directory domains can be stored in Active Directory application directory partitions. By default, the DNS root domain of a new Active Directory forest is stored in the domain-wide application directory partition for the forest root domain. |

**DNS Support for Active Directory Physical Structure Requirements**

In order for a DNS server to be able to support Active Directory, the server is required to support the service (SRV) resource record type and the dynamic update protocol, as described in the RFC 2136. Active Directory uses DNS as the location mechanism for domain controllers, enabling computers on the network to obtain IP addresses of domain controllers. During the installation of Active Directory, the service (SRV) and address (A) resource records are dynamically registered in DNS. Both types of records are necessary for the functionality of the domain controller locator (Locator) mechanism.

To find domain controllers in a domain or forest, a client queries DNS for the SRV and A DNS resource records of the domain controller. The resource records provide the client with the names and IP addresses of the domain controllers. In this context, the SRV and A resource records are referred to as Locator DNS resource records.

When a domain controller is added to a forest, a DNS zone hosted on a DNS server is updated with the Locator DNS resource records for that domain controller. For this reason, the DNS zone must allow dynamic updates (RFC 2136), and the DNS server hosting that zone must support the SRV resource records (RFC 2782) to advertise the Active Directory directory service.

At the very least, the DNS server must support the SRV resource record; but the SRV resource records can be added to DNS manually. After installing Active Directory, these records can be found on the domain controller in the following location: *systemroot*\System32\Config\Netlogon.dns.

**Note**

- The configuration of reverse lookup zones is not based on the Windows Server 2003 and later domain structure; instead, it is based on the range of IP addresses assigned to a company. If a company is assigned B class IP addresses such as 172.56.X.Y., then a reverse lookup zone of 56.172.in-addr.arpa. will be created. It might contain delegations to some other domains such as 1.56.172.in-addr.arpa., 2.56.172.in-addr.arpa., or others. It is also possible to configure classless reverse lookup zones as described in the RFC Best Current Practice paper “Classless IN_ADDR.ARPA delegation.”

**_msdcs Subdomain**

To facilitate locating Windows Server 2003 or later based domain controllers, in addition to the standard *_Service*.*_Protocol*.*DnsDomainName* format, the Net Logon service registers SRV records that identify the well-known server-type pseudonyms “dc” (domain controller), gc (global catalog), pdc (primary domain controller), and “domains” (globally unique identifier, or GUID) as prefixes in the _msdcs subdomain. This Microsoft-specific subdomain enables location of domain controllers that have Windows Server 2003 or laterspecific roles in the domain or forest, as well as the location by GUID when a domain has been renamed. To accommodate locating domain controllers by server type or by GUID (abbreviated “dctype”), Windows Server 2003 and later based domain controllers register SRV records in the following form:

*_Service._Protocol.DcType.* **_msdcs.** *DnsDomainName*

**Note**

- There are also site–specific DNS resource records that use a slightly different format. For more information about site discovery, see "Domain Controller Location in the Closest Site" later in this document.

The subdomain _msdcs.*DnsDomainName* is used to find an LDAP server that is running TCP and also functioning in a particular Windows Server role. The name _msdcsis reserved for locating domain controllers and Kerberos servers. The single keyword _msdcswas chosen to avoid cluttering the DNS namespace unnecessarily. Other constant, well-known names (pdc, dc, and gc) were kept short to avoid exceeding the maximum length allowed for a DNS name.

**Domain Controller SRV Resource Records**

When a Windows Server 2003 or later based domain controller starts up, the Net Logon service uses dynamic updates to register SRV and A resource records in the DNS database, as described in Internet Engineering Task Force (IETF) RFC 2782, "A DNS RR for specifying the location of services (DNS SRV)." Windows Server 2003 also uses secure dynamic update using the GSS-TSIG algorithm, as described in RFC 3645, "Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG)."

The SRV record is used to map the name of a service (in this case, the LDAP service) to the DNS computer name of a server that offers that service. In a Windows Server 2003 network, an LDAP resource record locates a domain controller.

When a workstation logs on to a Windows Server 2003 or later domain, it queries DNS for SRV records in this general form:

*_Service* **.**_ *Protocol* **.** *DnsDomainName*

Because Active Directory servers offer the LDAP service over the TCP protocol, clients find an LDAP server by querying DNS for a record of the form:

**_ldap._tcp.** *DnsDomainName*

**Note**

- The service and protocol strings require an underscore (_) prefix to prevent potential collisions with existing names in the namespace.

Active Directory site-specific resource records are also used during logon. For more information, see "Domain Controller Location in the Closest Site" later in this document.

**SRV Records Registered by Net Logon**

The list in the following table provides the definitions of the names associated with registered SRV records. It also describes the lookup criteria supported by each record and the checks performed by Net Logon as each record is registered.

In the descriptions of registered SRV records, *DnsDomainName* refers to the DNS domain name that is used during creation of the domain controller when the domain tree is joined or created (that is, when the computer runs the Active Directory Installation Wizard). *DnsForestName* refers to the DNS domain name of the forest root domain.

An owner name is the name of the DNS node to which a resource record pertains. These resource records are used by domain controller Locator. The following table lists the owner names of the SRV records that are registered by Net Logon.

**SRV Records That Are Registered by Net Logon**

| SRV Resource Record | Description |
|---|---|
| **_ldap._tcp.** *DnsDomainName* **.** | Enables a client to locate a server that is running the LDAP service in the domain named *DnsDomainName*. The server is not necessarily a domain controller — that is, the only assumption that can be made about the server is that it supports the LDAP application programming interface (API). All Windows Server 2003–based domain controllers register this SRV record (for example, _ldap._tcp.contoso.com.). |
| **_ldap._tcp.** *SiteName* **. _sites.** *DnsDomainName* **.** | Enables a client to locate a server that is running the LDAP service in the domain named *DnsDomainName* in the site named *SiteName*. *SiteName* is the relative distinguished name of the site object that is stored in the Configuration container in Active Directory. All Windows Server 2003 and later based domain controllers register this SRV record (for example, _ldap._tcp.charlotte._sites.contoso.com.). |

| | |
|---|---|
| **_ldap._tcp.dc._msdcs.** *DnsDomainName* **.** | Enables a client to locate a domain controller (dc) of the domain named *DnsDomainName*. All Windows Server 2003 and later based domain controllers register this SRV record. |
| **_ldap._tcp.** *SiteName* **. _sites.dc._msdcs.** *DnsDomainName* **.** | Enables a client to locate a domain controller for the domain named *DnsDomainName* and in the site named *SiteName*. All Windows Server 2003 and laterbased domain controllers register this SRV record. |
| **_ldap._tcp.pdc._msdcs.** *DnsDomainName* **.** | Enables a client to locate the server that is acting as the primary domain controller (PDC) in the mixed-mode domain named *DnsDomainName*. Only the PDC emulator master of the domain (the Windows Server 2003 and laterbased domain controller that advertises itself as the primary domain controller to computers that need a primary domain controller) registers this SRV record. |
| **_ldap._tcp.gc._msdcs.** *DnsForestName* **.** | Enables a client to locate a global catalog (gc) server for this forest. Only domain controllers that are functioning as gc servers for the forest named in *DnsForestName* register this SRV record (for example, _ldap._tcp.gc._msdcs.contoso.com.). |
| **_ldap._tcp.** *SiteName* **. _sites.gc._msdcs.** *DnsForestName* **.** | Enables a client to locate a global catalog (gc) server for this forest in the site named in SiteName. Only domain controllers that are serving as gc servers for the forest named in DnsForestName register this SRV record (for example, _ldap._tcp.charlotte._sites.gc._msdcs.contoso.com.). |
| **_gc._tcp.***DnsForestName***.** | Enables a client to locate a global catalog (gc) server for this domain. The server is not necessarily a domain controller. Only a server that is running the LDAP service and functioning as the GC server for the forest named *DnsForestName* registers this SRV record (for example, _gc._tcp.contoso.com.). In Windows Server 2003 and later, a GC server is a domain controller. Other implementations of directory services (that are not Windows Server 2003 or later implementations) can also register servers as GC servers. |
| **_gc._tcp.***SiteName***. _sites.***DnsForestName***.** | Enables a client to locate a global catalog (gc) server for this forest in the site named *SiteName*. The server is not necessarily a domain controller. Only a server that is running the LDAP service and functioning as the GC server for the forest named *DnsForestName* registers this SRV record (for example, _gc._tcp.charlotte._sites.contoso.com.). |
| **_ldap._tcp.** *DomainGuid* **. domains._msdcs.** *DnsForestName* **.** | Enables a client to locate a domain controller in a domain on the basis of its GUID. A GUID is a 128-bit number that is automatically generated for referencing objects in Active Directory — in this case, the domain object. This operation is expected to be infrequent; it occurs only when the DnsDomainName of the domain has changed, the DnsForestName is known, and DnsForestName has not also been renamed (for example, _ldap._tcp.4f904480-7c78-11cf-b057-00aa006b4f8f.domains. _msdcs.contoso.com.).<br><br>All domain controllers register this SRV record. |
| **_kerberos._tcp.** *DnsDomainName* **.** | Enables a client to locate a server that is running the Kerberos KDC service for the domain that is named in *DnsDomainName*. The server is not necessarily a domain controller. All Windows Server 2003 and later based domain controllers that are running an RFC 1510–compliant Kerberos KDC service register this SRV record. |
| **_kerberos._udp.** *DnsDomainName* **.** | Same as **_kerberos._tcp.***DnsDomainName*, except that UDP is implied. |
| **_kerberos._tcp.** *SiteName* **. _sites.** *DnsDomainName* **.** | Enables a client to locate a server that is running the Kerberos KDC service for the domain that is named *DnsDomainName* and is also in the site named *SiteName*. The server is not necessarily a domain controller. All Windows Server 2003 and later based domain controllers that are running an RFC 1510–compliant Kerberos KDC service register this SRV record. |
| **_kerberos._tcp.dc._msdcs.** *DnsDomainName* **.** | Enables a client to locate a domain controller that is running the Windows Server 2003 or later implementation of the Kerberos KDC service for the domain named in *DnsDomainName*. All Windows Server 2003 and later based domain controllers that are running the KDC service (that is, that implement a public key extension to the Kerberos v5 protocol Authentication Service Exchange subprotocol) register this SRV record. |

| | |
|---|---|
| **_kerberos.tcp.** *SiteName* **. _sites.dc._msdcs.** *DnsDomainName* **.** | Enables a client to locate a domain controller that is running the Windows Server 2003 implementation of the Kerberos KDC service for the domain that is named *DnsDomainName* and that is also in the site named *SiteName*. All Windows Server 2003 and later based domain controllers that are running the KDC service (that is, that implement a public key extension to the Kerberos v5 protocol Authentication Service Exchange subprotocol) register this SRV record. |
| **_kpasswd._tcp**.*DnsDomainName***.** | Enables a client to locate a Kerberos Password Change server for the domain. All servers that provide the Kerberos Password Change service (which includes all Windows Server 2003 and later based domain controllers) register this name. This server must at least conform to the Kerberos Change Password Protocol. (For more information about this draft, see the Microsoft Platform SDK.) The server is not necessarily a domain controller. All Windows Server 2003 and later based domain controllers that are running an RFC 1510–compliant Kerberos KDC service register this SRV record. |
| **_kpasswd._udp**.*DnsDomainName***.** | Same as **_kpasswd._tcp.***DnsDomainName*, except that UDP is implied. |

In addition to the SRV records listed in the table above, Net Logon also registers a DNS alias (CNAME) record for use by Active Directory replication, in the format: *DsaGuid***._msdcs.***DnsForestName*. The Locator does not use this record. This record enables a client to locate any domain controller in the forest by looking up an A record. The only information that is known about the domain controller is the GUID of the directory system agent (DSA) object for the domain controller and the name of the forest in which the domain controller is located. This record is used to facilitate renaming a domain controller.

If multiple domain controllers have the same criteria, multiple records exist with the same owner name. A client that is looking for a domain controller with specific criteria would receive all the applicable records from the DNS server. The client would pick one of the returned records to select a domain controller, as described in RFC 2782.

**Other SRV Record Content**

The following information is also included in an SRV record.

| SRV Record Field | Description |
|---|---|
| Priority | The priority of the server. Clients attempt to contact the server with the lowest priority. |
| Weight | A load-balancing mechanism that is used when selecting a target host from those that have the same priority. Clients randomly choose SRV records that specify target hosts to be contacted, with probability proportional to the weight. |
| Port number | The port where the server is listening for this service. |
| Target | The fully qualified domain name of the host computer. |

The algorithm by which clients interpret and select among SRV resource records is defined in RFC 2782, "A DNS RR for specifying the location of services (DNS SRV)."

**Host Records for Non-SRV-Aware Clients**

Net Logon registers the following DNS A records for the use of LDAP clients that do not support DNS SRV records (that is, clients that are non-SRV-aware). Locator does not use these records.

### Host (A) Records Registered by Net Logon

| Host (A) Resource Record | Description |
|---|---|
| *DnsDomainName* **.** | Enables a non-SRV-aware client to locate any domain controller in the domain by looking up an A record. A name in this form is returned to the LDAP client through an LDAP referral. A non-SRV-aware client looks up the name; an SRV-aware client looks up the appropriate SRV resource record. |
| **gc._msdcs.** *DnsForestName* **.** | Enables a non-SRV-aware client to locate any global catalog server in the forest by looking up an A record. A name in this form is returned to the LDAP client through an LDAP referral. A non-SRV-aware client looks up this name; an SRV-aware client looks up the appropriate SRV resource record. |

**Example of Registered Resource Records**
The following example illustrates the combined information that is contained in A resource records and SRV resource records.

A domain controller named Phoenix in the domain contoso.com has an IP address of 157.55.81.157. It registers the following A records and SRV records with DNS:

```
phoenix.contoso.com   A   157.55.81.157
_ldap._tcp.contoso.com    SRV  0 0 389 phoenix.contoso.com
_kerberos._tcp.contoso.com   SRV  0 0 88 phoenix.contoso.com
_ldap._tcp.dc._msdcs.contoso.com  SRV  0 0 389 phoenix.contoso.com
_kerberos._tcp.dc._msdcs.contoso.com  SRV  0 0 88 phoenix.contoso.com.
```

When the appropriate SRV records and A records are in place, a DNS query for _ldap._tcp.dc._msdcs.contoso.com returns the names and addresses of all domain controllers in the domain.

For more information about A records, SRV records, DNS, and dynamic updates, see "DNS Protocol" and "DNS Physical Structure" in How DNS Works.

**DNS Application Directory Partitions**
DNS zones stored in Active Directory replicate to Active Directory domain controllers according to different replication scopes. In Windows 2000 Server, a DNS zone stored in Active Directory is replicated to all domain controllers in the Active Directory domain. Windows Server 2003 and later have added application directory partitions, which enable the DNS zone to be stored in different replication scopes. The following table describes all of the replication scopes available to a DNS zone stored in Active Directory.

**Windows Server 2003 and later Active Directory Storage Options**

| Active Directory Storage Option | Replication Scope |
|---|---|
| Domain partition | Active Directory domain partition for each domain in the forest. DNS zones stored in this partition are replicated to all domain controllers in the domain. This is the only Active Directory storage option for DNS zones that are replicated to domain controllers running Windows 2000 Server. |
| Forest-wide DNS application directory partition | DNS application directory partition for the entire forest. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controllers in the forest. This DNS application directory partition is created when you install the DNS Server service on the first Windows Server 2003 or later domain controller in the forest. |
| Domain-wide DNS application directory partition | DNS application directory partition for each domain in the forest. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controllers in the domain. For the forest root domain, this DNS application directory partition is created when you first install the DNS Server service on a Windows Server 2003 or later domain controller in the forest.<br><br>For each new domain in the forest (child domain), this DNS application directory partition is created when you first install the DNS Server service on a Windows Server 2003 or later domain controller for the new domain. |
| Custom DNS application directory partition | DNS application directory partition for any domain controller that is enlisted in its replication scope. This type of DNS application directory partition does not exist by default and must be created. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controller that enlist in the partition. |

As stated earlier, the Locator DNS resource records for an Active Directory domain are stored in the _msdcs.*DnsDomainName* subdomain for the Active Directory domain. In Windows Server 2003 or later, when the DNS root domain of a new Active Directory forest is created on a Windows Server 2003 or later domain controller, a DNS zone is automatically created for the _msdcs.*DnsForestName* and stored in the forest-wide DNS application directory partition, which replicates to all Windows Server 2003 and later domain controllers in the forest running the DNS Server service.

**Note**

- DNS zones stored in application directory partitions cannot be accessed by Windows 2000 Server domain controllers.

**DNS Support for Active Directory Processes and Interactions**

When a Windows Server 2003 or later member server is promoted to an Active Directory domain controller by installing Active Directory, the Net Logon service registers the DNS resource records necessary for network hosts and services to be able to locate the domain controller on the network. When network hosts and services attempt to perform an operation (such as joining a domain, for example) that requires an Active Directory domain controller, the Locator mechanism is used to locate the domain controller through DNS. The following table describes the processes and interactions involved in the registration and location of domain controllers in DNS.

### Active Directory and DNS Processes and Interactions

| Process or Interaction | Description |
|---|---|
| Domain controller DNS name registration | The Net Logon service registers DNS resource records on behalf of the Active Directory domain controller in the DNS zone with the same name as the Active Directory domain hosted by the domain controller. |
| DNS delegation, forwarders | DNS delegation resource records are created in the zone that is a parent of the zone supporting the Active Directory domain. The delegation enables the DNS name of the domain controller to be resolved downward from the root of the DNS hierarchy.<br><br>DNS forwarders are another DNS feature that enable domain controller location, and are commonly used for an Active Directory client in one domain to locate a domain controller in another domain. |
| DNS domain controller location | Network hosts and services use the DNS Locator mechanism to locate domain controllers in the Active Directory forest. |

**Domain Controller Name Registration**

Every Windows Server 2003 or laterbased domain controller registers two types of names at startup:

- A DNS domain name with the DNS service (for example, noam.contoso.com).
- A NetBIOS name with Windows Internet Name Service (WINS) or another transport-specific service (for example, noam).

When a user starts a computer and logs on to a domain, the computer must do one of two things:

- If the name of the logon domain is a DNS name, the computer must query DNS to find a domain controller with which to authenticate.
- If the name of the logon domain is a NetBIOS name, the computer must send a mailslot message to find a domain controller for the specified domain.

After the computer has found a domain controller, the information is cached so that a new query is not required for subsequent logon sessions.

**DNS Domain Name Registration**

Active Directory supports dynamic registration of domain controller addresses in DNS. After Active Directory has been installed during domain controller creation, the Net Logon service dynamically creates records in the DNS database that are used to locate the server. Dynamic update (described in RFC 2136) is a recent addition to the DNS standard; this addition to the standard defines a protocol for dynamically updating a DNS server with new or changed resource record values. Before the advent of this new protocol, administrators had to manually create the records that are stored on DNS servers. The implementation of DNS server that is included with Windows Server 2003 and later supports dynamic updates, as does the Berkeley Internet Name Domain (BIND) version 8.*x* implementation of DNS.

### Note

- By default, the Windows Server 2003 or later DNS Server service running on a domain controller is configured to accept secure dynamic update only.

Every Windows Server 2003 or later based domain controller dynamically registers SRV records in DNS. The SRV records enable servers to be located by service type (for example, LDAP) and protocol (for example, TCP). Because domain controllers

are LDAP servers that communicate over TCP, SRV records can be used to find the DNS computer names of domain controllers. In addition to registering LDAP-specific SRV records, Net Logon also registers Kerberos v5 authentication protocol–specific SRV records to enable locating servers that run the Kerberos Key Distribution Center (KDC) service.

Every Windows Server 2003 or later based domain controller also dynamically registers a single host (A) resource record that contains the name of the domain (*DnsDomainName*) where the domain controller is and the IP address of the domain controller. The host (A) resource record makes it possible for clients that do not recognize service (SRV) resource records to locate a domain controller by means of a generic host lookup.

**NetBIOS Domain Name Registration**
A domain controller registers its NetBIOS name (*DomainName*[1C]) by broadcasting or directing a NetBIOS name registration request to a NetBIOS name server, such as a WINS server. Registering the NetBIOS name makes it possible for Windows-based clients that are not DNS-enabled to find the domain controllers are running Windows NT 3.51 or later. In this case, the client finds the domain controller by sending a Net Logon mailslot request that is based on the NetBIOS domain name.

### Note

- NetBIOS recognizes domain controllers by the [1C] registration.

**DNS Delegation, Forwarders**
To fully support an Active Directory domain, the DNS infrastructure must have the DNS delegations that are necessary to enable name resolution during domain controller location. When an Active Directory domain is created, a DNS delegation entry must exist in the DNS zone that is the parent of the zone supporting the Active Directory domain. The delegation enables the name of the domain controller hosting the domain to be resolved by any host or service in the DNS namespace. The delegation resource records must be added by the network administrator who administers the DNS server hosting the parent zone. Alternately, the parent zone could host the DNS resource records for the domain name and, in this case, the delegation is unnecessary.

For example, if there is a DNS zone supporting the domain corp.contoso.com, then a delegation for this name must exist in the parent zone contoso.com. When a DNS query for the name is sent to the root of the DNS namespace, delegations can be followed until the DNS server hosting the zone for corp.contoso.com is identified. Without this delegation, only those network hosts and services configured with the IP address of the DNS server hosting the zone for corp.contoso.com will be able to resolve its DNS name. All other network hosts and services will be unable to resolve the name, and the domain controller will not be available to them for such Active Directory operations as joining a domain, logging on to a domain, or searching Active Directory.

When the Active Directory domain name is specified in the Active Directory Installation Wizard, the wizard reads the delegation entry in DNS, prompts the user to install the DNS Server service locally and configures the corresponding DNS zone automatically. Also, if the computer is already configured with a preferred DNS server, the wizard will configure a forwarder on the DNS Server service with the IP address of the prior preferred DNS server.

Forwarders are also commonly configured on the DNS server hosting a zone in support of an Active Directory domain. Forwarders are used by a DNS server to forward queries for a domain name about which it has no local data. If an Active Directory client in one domain needs to access a resource in another Active Directory domain, it will need to locate a domain controller in that domain. If the DNS server used by that client is unable to locate a domain controller for the other domain using its local data, it can forward the client request to another DNS server, such as the DNS server that hosts the zone for the Active Directory forest root domain.

In summary, delegation enables name resolution in a descending direction from the root of the hierarchy, and forwarders enable name resolution in a ascending direction toward the root of the hierarchy.

Windows Server 2003 and later introduces an enhancement to forwarding called conditional forwarders. When you use conditional forwarding, you can configure your DNS servers to forward queries to different servers based on the domain name specified in the query. This eliminates steps in the standard forwarding chain and reduces network traffic. When conditional forwarding is applied, the DNS server hosting a zone in support of the one Active Directory domain can forward queries to DNS servers hosting zones in support of the Active Directory domain name specified in the client query.

Although root hints (resource records that list the DNS servers hosting the DNS root zone) can also be used to facilitate domain controller location in place of forwarders, DNS deployments commonly use forwarders for remote domain controller location to reduce the complexity of administering root hints for both internal and Internet resolution.

For more information about delegation, forwarders, and root hints, see "How DNS Works."

**Domain Controller Locator**
The domain controller locator (Locator) algorithm consists of two main parts:

- Locator finds which domain controllers are registered with a DNS server.
- Locator submits a DNS query to the DNS server to locate a domain controller in the specified domain.

After this query is resolved, an LDAP User Datagram Protocol (UDP) lookup is sent to one or more of the domain controllers listed in the response to the DNS query to ensure their availability. Finally, the Net Logon service caches the discovered domain controller to aid in resolving future requests.

**Domain Controller Locator Process**

Each Windows Server 2003 or later based domain controller registers its DNS domain name on the DNS server and registers its NetBIOS name by using a transport-specific mechanism (for example, WINS). Thus, a DNS client locates a domain controller by querying DNS, and a NetBIOS client locates a domain controller by querying the appropriate transport-specific name service. Because the code for the DNS-compatible Locator and the Windows NT 4.0–compatible Locator is shared, both DNS clients and NetBIOS clients are supported.

The process that the Locator follows can be summarized as follows:

1. On the client (the computer that is locating the domain controller), the Locator is initiated as a remote procedure call (RPC) to the local Net Logon service. The Locator API (DsGetDcName) is implemented by the Net Logon service.
2. The client collects the information that is needed to select a domain controller and passes the information to the Net Logon service by using the DsGetDcName API.
3. The Net Logon service on the client uses the collected information to look up a domain controller for the specified domain in one of two ways:
    - For a DNS name, Net Logon queries DNS by using the DNS-compatible Locator — that is, DsGetDcName calls DnsQuery to read the SRV records and A records from DNS after it appends an appropriate string to the front of the domain name that specifies the SRV record.
    - For a single label name, Net Logon performs domain controller discovery by using the Windows NT 4.0–compatible Locator — that is, by using the transport-specific mechanism (for example, WINS).

    **Note**

In Windows NT 4.0 and earlier, "discovery" is a process for locating a domain controllerfor authentication in either the primary domain or a trusted domain.

1. The Net Logon service sends a datagram to the discovered domain controllers that register the name. For NetBIOS domain names, the datagram is implemented as a mailslot message. For DNS domain names, the datagram is implemented as an LDAP UDP search.
2. Each available domain controller responds to the datagram to indicate that it is currently operational and then returns the information to DsGetDcName.
3. The Net Logon service returns the information to the client from the domain controller that responds first.
4. The Net Logon service caches the domain controller information so that it is not necessary to repeat the discovery process for subsequent requests. Caching this information encourages the consistent use of the same domain controller and, thus, a consistent view of Active Directory.

**Domain Controller Location in the Closest Site**

During a search for a domain controller, the Locator attempts to find a domain controller in the site closest to the client. When the domain that is being sought is a Windows Server 2003 or later domain, the domain controller uses the information stored in Active Directory to determine the closest site. When the domain being sought is a Windows NT 4.0 domain, domain controller discovery occurs when the client starts and uses the first domain controller that it finds.

Each Windows Server 2003 or later based domain controller registers DNS records that indicate the site where the domain

controller is located. The site name (the relative distinguished name of the site object in Active Directory) is registered in several records so that the various roles the domain controller might perform (for example, global catalog server or Kerberos server) can be associated with the domain controller's site. When DNS is used, the Locator searches first for a site-specific DNS record before it begins to search for a DNS record that is not site-specific (thereby preferentially locating a domain controller in that site).

A client computer stores its own site information in the registry, but the computer is not necessarily located physically in the site associated with its IP address. For example, a portable computer that was moved to a new location could contact a domain controller in its home site, which is not the site to which the computer is currently connected. In this situation, the domain controller looks up the client site on the basis of the client IP address by comparing the address to the sites that are identified in Active Directory, and then returns the name of the site that is closest to the client. The client then updates the information in the registry.

The domain controller stores site information for the entire forest in the Configuration container. The domain controller uses the site information to check the IP address of the client computer against the list of subnets in the forest. In this way, the domain controller ascertains the name of the site in which the client is assumed to be located, or the site that is the closest match, and returns this information to the client.

**Active Directory Site and Subnet Objects**
A site is a collection of subnets that have high-speed connections. In Active Directory, a site is defined by a site object in the cn=Sites,cn=Configuration,dc=*ForestRootDomain* container. A subnet is an addressed segment within a site and is represented by an object in the cn=Subnets,cn=Sites,cn=Configuration,dc=*ForestRootDomain* container.

The site in which a domain controller is located is identified in the Configuration container by the domain controller object that is located within the cn=Servers container beneath the site object for a particular site. A domain controller can identify the site of a client by using the subnet object in the Sites container. Each subnet object has a *siteObject* property ("attribute") that links it to a site object; the value of the *siteObject* property is the distinguished name of the site object. This link enables a domain controller to identify clients that have an IP address in the specified subnet as being in the specified site.

Subnet names in Active Directory take the form "network/bits masked" (for example, the subnet object 172.16.72.0/22 has a subnet of 172.16.72.0 and a 22-bit subnet mask). If this subnet had a *siteObject* property value that contained the distinguished name of the Seattle site object, all IP addresses in the 172.16.72.0/22 subnet would be considered to be in the Seattle site. The *siteObject* property is a single value, which implies that a single subnet maps to a single site. However, multiple subnet objects can be linked to the same site object. The directory administrator manually creates subnet objects and, hence, the *siteObject* property value.

The Configuration container (including all of the site and subnet objects in it) is replicated to all domain controllers in the forest. Therefore, any domain controller in the forest can identify the site in which a client is located, compare it to the site in which the domain controller is located, and indicate to the client whether that domain controller's site is the closest site to the client.

For more information about networks, subnets, and subnet masks, see "How TCP/IP Works."

**IP Addresses Mapped to Site Names**
During Net Logon startup, the Net Logon service on each domain controller enumerates the site objects in the Configuration container. Net Logon on each domain controller is also notified of any changes made to the site objects. Net Logon uses the site information to build an in-memory structure that is used to map IP addresses to site names.

When a client that is searching for a domain controller receives the list of domain controller IP addresses from DNS, the client begins querying the domain controllers in turn to find out which domain controller is available and appropriate. Active Directory intercepts the query, which contains the IP address of the client, and passes it to Net Logon on the domain controller. Net Logon looks up the client IP address in its subnet-to-site mapping table by finding the subnet object that most closely matches the client IP address and then returns the following information:

- The name of the site in which the client is located, or the site that most closely matches the client IP address.
- The name of the site in which the current domain controller is located.
- A bit that indicates whether the found domain controller is located (bit is set) or not located (bit is not set) in the site closest to the client.

The domain controller returns the information to the client. The response also contains various other pieces of information that describe the domain controller.

The client inspects the information to determine whether to try to find a better domain controller. The decision is made as follows:

- If the returned domain controller is in the closest site (the returned bit is set), the client uses that domain controller.
- If the client has already tried to find a domain controller in the site in which the domain controller claims the client is located, the client uses that domain controller.
- If the domain controller is not in the closest site, the client updates its site information and sends a new DNS query to find a new domain controller in the site. If the second query is successful, the client uses the new domain controller. If the second query fails, the client uses the original domain controller.
- If the domain that is being queried by a computer is the same as the domain to which the computer is joined, the site in which the computer resides (as reported by a domain controller) is stored in the computer registry. The client stores this site name in the **DynamicSiteName** registry entry in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters. Therefore, the DsGetSiteName API returns the site in which the computer is located.

You can override the dynamically determined value using the registry, but you should never change dynamically determined values. To override the dynamic site name, add the **SiteName** entry with the REG_SZ data type in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters. When a value is present for the **SiteName** entry, the **DynamicSiteName** entry is not used.

**Note**

- Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Windows Server 2003 or later. There are programs available in Control Panel or Microsoft Management Console (MMC) for performing most administrative tasks. These programs provide safeguards that prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Registry editors bypass the standard safeguards that are provided by these administrative tools. Modifying the registry is recommended only when no administrative tool is available. Before you make changes to the registry, it is recommended that you back up any valuable data on the computer. For instructions about how to edit registry entries, see Help for the registry editor that you are using.

If the domain being located is the same as the domain to which the computer is joined and the computer has not physically moved to a different site since the last query, the dynamically determined site name in the registry is the actual site in which the computer is located. Thus, the client finds a domain controller in the correct site without having to retry the operation. On the other hand, if the site name in the registry is not the current site of the computer (for example, if the computer is portable), the domain controller location process serves to update the site information in the registry.

**Automatic Site Coverage**

There is not necessarily a domain controller in every site. For various reasons, it is possible that no domain controller exists for a particular domain at the local site. By default, each domain controller checks all sites in the forest and then checks the replication cost matrix. A domain controller advertises itself (registers a site-related SRV record in DNS) in any site that does not have a domain controller for that domain and for which its site has the lowest-cost connections. This process ensures that every site has a domain controller that is defined by default for every domain in the forest, even if a site does not contain a domain controller for that domain. The domain controllers that are published in DNS are those from the closest site (as defined by the replication topology).

For example, given one domain and three sites, a domain controller for that domain might be located in two of the sites, but there might be no domain controller for the domain in the third site. Replication to the domain that does not have a domain controller in the third site might be too expensive in terms of cost or replication latency. To ensure that a domain controller can be located in the site closest to a client computer, if not the same site, Windows Server 2003 and later automatically attempts to register a domain controller in every site. The algorithm that is used to accomplish automatic site coverage determines how one site can cover another site when no domain controller exists in the second site.

**Determining Site Coverage on the Basis of Site-link Cost**

Site coverage is determined according to site-link costs, and domain controllers register themselves in sites accordingly. For example, given one domain and sites A, B, and C, site A has no domain controllers for the domain. If a client in site A attempts to locate a domain controller, which domain controller should be returned? The answer depends on which site covers site A for the domain.

In the example, a site link exists between site A and both of the other sites — that is, the connections between domain controllers in site A, site B, and site C are configured for replication over site links in Active Directory Sites and Services. Costs are associated with site links based on the expense of transferring data over the connections. The administrator uses the speed of the connection between sites to assign a cost to the communication link, and replication uses the cost to establish the least expensive route for replication traffic.

Site A and site B are connected by site link AB. Site A and site C are connected by site link AC, with the following costs:

Site link AB cost = 50.

Site link AC cost = 100.

The link between site A and site C has a much higher cost than the link between site A and site B. The administrator configured this cost based on the expensive Integrated Services Digital Network (ISDN) line that connects site A and site C, and the administrator would prefer that resources in site B be used when possible. The site coverage algorithm (described in the next sub-section) ensures that a domain controller in site B registers itself as a domain controller for site A. In this way, clients in Site A that are looking for a domain controller find one from site B, instead of possibly finding one from site C.

**Site Coverage Algorithm**

During registration of SRV records in DNS, the following algorithm determines if the domain controllers should register site SRV records to designate themselves as preferred domain controllers in Active Directory sites where no domain controller exists for a particular domain.

For every domain controller in the forest, this procedure is followed:

- Build a list of target sites — sites that have no domain controllers for this domain (the domain of the current domain controller).
- Build a list of candidate sites — sites that have domain controllers for this domain.
- For every target site, follow these steps:
    1. Build a list of candidate sites of which this domain is a member. (If none, do nothing.)
    2. Of these, build a list of sites that have the lowest site link cost to the target site. (If none, do nothing.)
    3. If more than one, break ties (reduce this list to one candidate site) by choosing the site with the largest number of domain controllers.
    4. If more than one, break ties by choosing the site that is first alphabetically.
    5. Register target-site-specific SRV records for the domain controllers for this domain in the selected site.

**Cache Time-out and Closest Site**

If a domain member computer requests a domain controller while all domain controllers in its site are offline, the Locator necessarily returns a domain controller in a different site. The location of this domain controller is stored in the client cache. The cache lifetime is controlled by the **CloseSiteTimeout** entry in the registry.

In addition, the domain controller performs authentication, and a secure channel is set up. On subsequent location attempts, the lifetime of the cache and the lifetime of the secure channel are secondary to the location of a domain controller in the closest site.

If the domain controller that is stored in the client cache is not in a site that is close to the client, Net Logon attempts to find a close domain controller when either of the following events occurs:

- An interactive logon process uses pass-through authentication on the secure channel.
- The value in the CloseSiteTimeout registry entry has elapsed since the last attempt, and any other attempt is made to use the secure channel (for example, pass-through authentication of network logons).

Thus, Net Logon attempts to find a close domain controller only on demand. The default value of the **CloseSiteTimeout**

period is 15 minutes; the maximum value is 49 days; the minimum value is 60 seconds. The implications of these settings are:

- If the time-out value is too large, a client never tries to find a close domain controller if there is not one available at startup.

- If the time-out value is too small, secure channel traffic is unnecessarily slowed down by discovery attempts.

**Clients with No Apparent Site**
Sometimes the client pings a domain controller and the client IP address cannot be found in the subnet-to-site mapping table. In this case, the domain controller returns a NULL site name, and the client uses the returned domain controller.

**Types of Locators**
On the basis of parameters passed to Net Logon in the DsGetDcName API, the process of locating a domain controller proceeds in one of two ways:

- The DNS-compatible Locator is used if the domain name passed to DsGetDcName is a DNS-compatible name. The Net Logon service on the client looks up the name in DNS (by calling DnsQuery) after it appends an appropriate string to the front of the domain name. The DNS service supports a query for determining the set of domain controllers. If the client site name is known, the client DNS query specifies the site. DNS returns the IP addresses of domain controllers that match the DNS query. The client Net Logon service sends an LDAP UDP message to one or more of the domain controllers that have been returned by DNS in order to determine whether any of the specified domain controllers are running and support the specified domain.

- The Windows NT 4.0–compatible Locator is used if the domain name passed to DsGetDcName is a NetBIOS name. The Net Logon service on the client sends a transport-specific logon request query to locate a domain controller in a particular domain and then sends a mailslot message to one or more of the domain controllers to determine whether any of the domain controllers it found are running and support the specified domain.

**DNS Client Service Configuration Elements**
The DNS configuration on Active Directory client computers follows a specific computer naming scheme and specifies how these clients will locate DNS servers. The following table lists the configurations for DNS configuration elements.

**DNS Configuration for Client Computers**

| DNS Configuration Element | Configuration |
|---|---|
| Computer naming | Use default naming. When a Windows 2000, Windows XP, or Windows Server 2003 or later computer joins a domain, the computer assigns itself a primary DNS name comprised of the host name of the computer and the name of the domain. |
| Client resolver configuration | Configure client computers to point to any DNS server on the network. |

**Note**

- Active Directory clients and domain controllers can dynamically register their DNS names using a DNS server that is not authoritative for the DNS name of the Active Directory domain. The DNS server used by these clients will refer the registration to the DNS server that is authoritative for the name of the Active Directory domain.

A computer might have an existing DNS name if the organization previously statically registered the computer in DNS or the organization previously deployed an integrated DHCP solution. If client computers already have a registered DNS name, when the domain to which they are joined is upgraded to Windows Server 2003 or later Active Directory, the client computers will have two names: the existing DNS name, and the new primary name.

Clients can be located by either name. Any existing DNS, DHCP, or integrated DNS/DHCP solution is left intact. The new primary names are created automatically and updated by means of dynamic update. The old names are cleaned up automatically by means of DNS scavenging.

To take advantage of Kerberos authentication when connecting to a server running Windows 2000, Windows XP, Windows

Server 2003 or later, the client must connect to the server by using the primary name.

**Network Ports Used by DNS in Support of Active Directory**

The network ports used by DNS are documented in the DNS Technical Reference. For more information, see "Network Ports used by DNS" in How DNS Works.

**Related Information**

The following resource contains additional information that is relevant to this section.

- How DNS Works

# DNS Support for Active Directory Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**DNS Support for Active Directory Tools and Settings**

**In this section**


- DNS Support for Active Directory Tools
- DNS Support for Active Directory Registry Entries
- DNS Support for Active Directory Group Policy Settings
- DNS Support for Active Directory WMI Classes
- Related Information


**DNS Support for Active Directory Tools**

The following tools are associated with the Domain Name System (DNS) support for Active Directory. For more information about DNS tools and settings in a Windows Server 2003 network, see "DNS Tools and Settings" in DNS Technical Reference.

| Note |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to Active Directory Domain Services. |

**Dcdiag.exe: Domain Controller Diagnostic Tool**

**Category**

This tool is included in the Windows 2000 Server and later Support Tools.

**Version compatibility**

This tool runs on Windows XP, and the Windows 2000 Server and later operating systems.

You can use Domain Controller Diagnostic Tool to verify that there are sufficient resources for the DNS infrastructure when deploying the Windows 2000 Server or later Active Directory directory service. This tool analyzes the state of domain controllers in a forest or enterprise and reports any problems, to assist in troubleshooting. As an end-user reporting program, Domain Controller Diagnostic Tool queries the directory service infrastructure and uses the results to identify abnormal behavior in the system. Domain Controller Diagnostic Tool provides a framework for executing tests and a series of tests to verify different functional areas of the system. This framework selects which domain controllers are tested according to scope directives from the user, such as enterprise, site, or single server.

**Dnscmd.exe: Dnscmd**

**Category**

This tool is included in the Windows Server 2003 and later Support Tools.

**Version compatibility**

This tool runs on the Windows 2000 Server and later operating systems.

Dnscmd is used to view the properties of DNS servers, zones, and resource records. In addition, Dnscmd is used to modify all aspects of the DNS Server service, including creating and deleting zones and resource records, and forcing replication events between DNS server physical memory and DNS databases. Dnscmd can also be useful for developing scripts for configuring a DNS server.

**Dnslint.exe: DNSLint**

**Category**

This tool is available for download on the Microsoft Web site.

**Version compatibility**
This tool runs on Windows XP, and the Windows 2000 Server and later operating systems.

DNSLint is a Microsoft Windows utility that can be used to help diagnose common DNS name resolution issues. It can be targeted to look for specific DNS record sets and ensure that they are consistent across multiple DNS servers. It can also be used to verify that DNS records used specifically for Active Directory replication are correct.

**Dnsmgmt.msc: DNS console**

**Category**
This tool is included in the Windows 2000 Server and later operating systems, and is installed when the DNS Server service is installed. This tool is also installed with the Windows 2000 Server or later Administration Tools Pack (Adminpak.msi).

**Version compatibility**
This tool runs on the Windows 2000 Server and later operating systems. When installed from one of the administration tools packs, this tool can also run on Microsoft Windows 2000 Professional and Windows XP.

The DNS console is used to administer the DNS Server service. It can be used to modify all aspects of the DNS Server service, including creating and deleting zones and resource records, and forcing replication events between DNS server physical memory and DNS databases. The DNS console can also be used to perform diagnostics on the DNS infrastructure of a network.

**Eventvwr.exe: Event Viewer**

**Category**
This tool is included in all Windows server and client operating systems.

**Version compatibility**
This tool runs on Windows XP, and the Windows 2000 Server and later operating systems.

You can use Event Viewer to monitor events recorded in event logs. Typically, a computer stores the Application, Security, and System logs. It could also contain other logs, depending on the role of the computer and the applications that are installed on it. For example, DNS servers write DNS-related events (such as errors that occur when the DNS Server service is invoked) to log files which can be read by using Event Viewer.

**Ipconfig.exe: Ipconfig**

**Category**
This tool is included in all Windows server and client operating systems.

**Version compatibility**
This tool runs on all Windows server and client operating systems.

Ipconfig displays all current TCP/IP network configuration values and refreshes Dynamic Host Configuration Protocol (DHCP) and DNS settings. Used without parameters, Ipconfig displays IPv6 addresses or the IPv4 address, subnet mask, and default gateway for all adapters.

**Netdiag.exe: Network Connectivity Tester**

**Category**
This tool is included in the Windows 2000 and later Support Tools.

**Version compatibility**
This tool runs on Windows XP, and the Windows 2000 Server and later operating systems.

You can use Network Connectivity Tester to help isolate networking and connectivity problems, by performing a series of tests to determine the state of your network client, and whether it is functional. These tests, and the key network status information they expose, give network administrators and support personnel a more direct means of identifying and isolating network problems. Moreover, because this tool does not require that parameters or switches be specified, support personnel and network administrators can focus on analyzing the output rather than on training users how to use the tool.

**Netmon.exe: Network Monitor**

**Category**
This tool is installed with the Windows Server 2003 and later Administration Tools Pack.

**Version compatibility**
This tool runs on all Windows operating systems.

Network Monitor captures data about the packets on a network and logs them for subsequent analysis. The monitored data can be filtered using many criteria, including protocol, ports, physical addresses, and logical addresses. Network Monitor can be useful in many situations, such as when you are troubleshooting an environment that has a firewall between a DNS server and a client, or between two DNS servers.

There are two versions of Network Monitor: the Network Monitor that is provided as part of the Windows Server 2003 operating system, and the Network Monitor that is part of Microsoft Systems Management Server (SMS). The version of Network Monitor that is included with the Windows Server 2003 operating system captures only data about network packets that are sent to or from the server on which you run Network Monitor; it also captures data about network broadcasts that are received. The Network Monitor that is included with SMS can monitor all network packets on a network segment regardless of their source or destination.

**Nslookup.exe: Nslookup**

**Category**
This tool is included in all Windows server and client operating systems.

**Version compatibility**
This tool runs on all Windows server and client operating systems.

Nslookup is used to query DNS servers and to obtain detailed responses. The information obtained using Nslookup can be used to diagnose and solve name resolution problems, verify that resource records are added or updated correctly in a zone, and debug other server-related problems.

**DNS Support for Active Directory Registry Entries**
Active Directory uses DNS to enable client computers to locate domain controllers and to enable domain controllers to locate each other. Domain controllers register SRV records in DNS, and clients and other domain controllers query for these records. Which records are registered in DNS and how they are registered depends on settings in the Windows registry. The following registry entries are associated with DNS support for Active Directory.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, modify settings though the use of policy settings or through dedicated management tools, such as the DNS snap-in for the Microsoft Management Console (MMC), to accomplish tasks, rather than editing the registry directly. If you must edit the registry, use extreme caution.

**DNS\Parameters**
The following registry entries are used to configure various DNS options that control how DNS interacts with the Active Directory environment.

**EnableDirectoryPartitions**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

**Version**
Windows Server 2003 and later

If you do not want the default DNS application directory partitions to be created automatically, you must disable the **EnableDirectoryPartitions** registry key. The values for this key are **0x0** (disable) and **0x1** (enable).

For more information about this registry entry, see the "Registry Reference" in the Tools and Settings Collection.

**DnsAvoidRegisterRecords**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters

**Version**
Windows Server 2003 and later

Specifies the list of data corresponding to the DNS resource records that should not be registered for a domain controller by the Net Logon service. Restarting the Net Logon service is not required to make the changes to this value effective. If the

**DnsAvoidRegisterRecords** registry key is created or modified within the first 15 minutes after the Net Logon service is started, there might be a short delay before the appropriate DNS updates appear and are replicated to the other DNS servers. If the modifications are made while the Net Logon service is stopped there is a short delay before the DNS updates appear after the Net Logon service is restarted.

In this value, list the data corresponding to the DNS resource records that should not be registered for this domain controller by the Net Logon service. The list of data includes:

| Data Value | Resource Record Type | DNS Resource Record |
|---|---|---|
| LdapIpAddress | A | <*DnsDomainName*> |
| Ldap | SRV | _ldap._tcp.<*DnsDomainName*> |
| LdapAtSite | SRV | _ldap._tcp.<*SiteName*>._sites.<*DnsDomainName*> |
| Pdc | SRV | _ldap._tcp.pdc._msdcs.<*DnsDomainName*> |
| Gc | SRV | _ldap._tcp.gc._msdcs.<*DnsForestName*> |
| GcAtSite | SRV | _ldap._tcp.<*SiteName*>._sites.gc._msdcs.<*DnsForestName*> |
| DcByGuid | SRV | _ldap._tcp.<*DomainGuid*>.domains._msdcs.<*DnsForestName*> |
| GcIpAddress | A | _gc._msdcs.<*DnsForestName*> |
| DsaCname | CNAME | <*DsaGuid*>._msdcs.<*DnsForestName*> |
| Kdc | SRV | _kerberos._tcp.dc._msdcs.<*DnsDomainName*> |
| KdcAtSite | SRV | _kerberos._tcp.dc._msdcs.<*SiteName*>._sites.<*DnsDomainName*> |
| Dc | SRV | _ldap._tcp.dc._msdcs.<*DnsDomainName*> |
| DcAtSite | SRV | _ldap._tcp.<*SiteName*>._sites.dc._msdcs.<*DnsDomainName*> |
| Rfc1510Kdc | SRV | _kerberos._tcp.<*DnsDomainName*> |
| Rfc1510KdcAtSite | SRV | _kerberos._tcp.<*SiteName*>._sites.<*DnsDomainName*> |
| GenericGc | SRV | _gc._tcp.<*DnsForestName*> |
| GenericGcAtSite | SRV | _gc._tcp.<*SiteName*>._sites.<*DnsForestName*> |
| Rfc1510UdpKdc | SRV | _kerberos._udp.<*DnsDomainName*> |
| Rfc1510Kpwd | SRV | _kpasswd._tcp.<*DnsDomainName*> |
| Rfc1510UdpKpwd | SRV | _kpasswd._udp.<*DnsDomainName*> |

**DNS Support for Active Directory Group Policy Settings**

Active Directory uses DNS to enable client computers to locate domain controllers and to enable domain controllers to locate each other. Domain controllers register SRV records in DNS, and clients and other domain controllers query for these records. Which records are registered in DNS and how they are registered depends on specific Group Policy settings The following tables list and describe the Group Policy settings that are associated with DNS support for Active Directory.

**Net Logon Group Policy Settings Associated with DNS Support for Active Directory**

| Group Policy Setting | Description |
|---|---|
| Site Name | Specifies the Active Directory site to which computers belong. An Active Directory site is one or more well-connected TCP/IP subnets that allow administrators to configure Active Directory access and replication. |

| | |
|---|---|
| Negative DC Discovery Cache Setting | Specifies the amount of time (in seconds) the DC locator retains that a domain controller (DC) could not be found in a domain. When a subsequent attempt to locate the DC occurs within the time set in this setting, DC Discovery immediately fails, without attempting to find the DC. |

**Domain Controller Locator Group Policy Settings Associated with DNS Support for Active Directory**

| Group Policy Setting | Description |
|---|---|
| Dynamic Registration of the DC Locator DNS Records | Determines if Dynamic Registration of the DC locator DNS resource records is enabled. These DNS records are dynamically registered by the Net Logon service and are used by the Locator algorithm to locate the DC. |
| DC Locator DNS records not registered by the DCs | Determines which DC Locator DNS records are not registered by the Netlogon service. |
| Refresh Interval of the DC Locator DNS Records | Specifies the Refresh Interval of the DC Locator DNS resource records for DCs to which this setting is applied. These DNS records are dynamically registered by the Net Logon service and are used by the Locator algorithm to locate the DC. This setting can be applied only to DCs by using dynamic update. |
| Weight Set in the DC Locator DNS SRV Records | Specifies the Weight field in the SRV resource records registered by the DCs to which this setting is applied. These DNS records are dynamically registered by the Net Logon service and are used to locate the DC. |
| Priority Set in the DC Locator DNS SRV Records | Specifies the Priority field in the SRV resource records registered by DCs to which this setting is applied. These DNS records are dynamically registered by the Net Logon service and are used to locate the DC. |
| TTL Set in the DC Locator DNS SRV Records | Specifies the value for the Time-To-Live (TTL) field in Net Logon registered SRV resource records. These DNS records are dynamically registered by the Net Logon service and are used to locate the domain controller (DC). |
| Automated Site Coverage by the DC Locator DNS SRV Records | Determines whether DCs dynamically register DC Locator site-specific SRV records for the closest sites where no DC for the same domain exists (or no Global Catalog (GC) for the same forest exists). These DNS records are dynamically registered by the Net Logon service and are used to locate the DC. |
| Sites Covered by the DC Locator DNS SRV Records | Specifies the sites for which the DCs register the following:<br>• Site-specific DC Locator DNS SRV resource records<br>• Site-specific SRV records registered for the site where the DC is located<br>• Records registered by a DC configured to register DC Locator DNS SRV records for those sites without a DC that are closest to it |
| Sites Covered by the GC Locator DNS SRV Records | Specifies the sites for which the GCs should register the following:<br>• Site-specific GC locator DNS SRV resource records<br>• Site-specific SRV records registered for the site where the GC is located<br>• Records registered by a GC configured to register GC Locator DNS SRV records for those sites without a GC that are closest to it |

| | |
|---|---|
| Sites Covered by the Application Directory Partition Locator DNS SRV Records | Specifies the sites for which the DCs hosting the application directory partition should register the following:<br>• Site-specific, application directory partition-specific DC Locator DNS SRV resource records<br>• Site-specific SRV records registered for the site where the DC is located<br>• Records registered by a DC configured to register DC Locator DNS SRV records for those sites without a DC that are closest to it. |
| Location of the DCs hosting a domain with a single label DNS name | Specifies whether the computers to which this setting is applied attempt DNS name resolution of single-label domain names. |

For more information about Group Policy settings, see the "Group Policy Settings Reference" in the Tools and Settings Collection.

**DNS Support for Active Directory WMI Classes**

The WMI classes that are associated with the DNS support for Active Directory are the DNS WMI classes. The following table lists and describes the WMI classes that are associated with the DNS Server service.

**WMI Classes Associated With DNS Server Service**

| Microsoft DNS WMI Class | Description |
|---|---|
| MicrosoftDNS_Server | Describes a DNS server. Every instance of this class might be associated with one instance of class MicrosoftDNS_Cache, one instance of class MicrosoftDNS_RootHints, and multiple instances of class MicrosoftDNS_Zone. |
| MicrosoftDNS_Domain | Represents a domain in a DNS hierarchy tree. |
| MicrosoftDNS_Zone | Describes a DNS zone. Every instance of the class MicrosoftDNS_Zone must be assigned to exactly one DNS Server. Zones might be associated with multiple instances of the classes MicrosoftDNS_Domain and MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_Cache | Describes a cache existing on a DNS server (do not confuse this with a cache file that contains root hints). This class simplifies visualizing the containment of DNS objects, rather than representing a real object. The class, MicrosoftDNS_Cache, is a container for the resource records cached by the DNS server.<br>Every instance of the class MicrosoftDNS_Cache must be assigned to exactly one DNS server. It might be associated with multiple instances of MicrosoftDNS_Domain and MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_RootHints | Describes the RootHints stored in a cache file on a DNS server. This class simplifies visualizing the containment of DNS objects, rather than representing a real object. Class MicrosoftDNS_RootHints is a container for the resource records stored by the DNS server in a cache file.<br>Every instance of the class MicrosoftDNS_RootHints must be assigned to exactly one DNS server. It might be associated with multiple instances of the MicrosoftDNS_ResourceRecord class. |
| MicrosoftDNS_Statistic | Represents a single DNS server statistic. |
| MicrosoftDNS_ServerDomainContainment | Every instance of the class MicrosoftDNS_ServerDomainContainment might contain multiple instances of the class MicrosoftDNS_Domain. |

| | |
|---|---|
| MicrosoftDNS_DomainDomainContainment | Every instance of the MicrosoftDNS_DomainDomainContainment class might contain multiple other instances of MicrosoftDNS_Domain. |
| MicrosoftDNS_DomainResourceRecordContainment | Every instance of the class MicrosoftDNS_DomainResourceRecordComtainment might contain multiple instances of the MicrosoftDNS_ResourceRecord class. |
| MicrosoftDNS_ResourceRecord | Represents the general properties of a DNS RR. |
| MicrosoftDNS_AAAAType | Represents an IPv6 Address (AAAA), often pronounced quad-A, RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_AFSDBType | Represents an Andrew File System Database Server (AFSDB) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_ATMAType | Represents an ATM Address-to-Name (ATMA) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_AType | Represents an Address (A) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_CNAMEType | Represents a Canonical Name (CNAME) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_HINFOType | Represents a Host Information (HINFO) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_ISDNType | Represents an ISDN RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_KEYType | Represents a KEY RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MBType | Represents a Mailbox (MB) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MDType | Represents a Mail Agent for Domain (MD) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MFType | Represents a Mail Forwarding Agent for Domain (MF) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MGType | Represents an MG RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MINFOType | Represents an Mail Information (MINFO) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MRType | Represents a Mailbox Rename (MR) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MXType | Represents a Mail Exchanger (MX) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_NSType | Represents a Name Server (NS) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_NXTType | Represents a Next (NXT) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_PTRType | Represents a Pointer (PTR) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_RPType | Represents a Responsible Person (RP) RR. Subclass of MicrosoftDNS_ResourceRecord. |

| | |
|---|---|
| MicrosoftDNS_RTType | Represents a Route Through (RT) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_SIGType | Represents a Signature (SIG) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_SOAType | Represents a Start Of Authority (SOA) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_SRVType | Represents a Service (SRV) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_TXTType | Represents a Text (TXT) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_WINSRType | Represents a WINS-Reverse (WINSR) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_WINSType | Represents a WINS RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_WKSType | Represents a Well-Known Service (WKS) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_X25Type | Represents an X.25 (X25) RR. Subclass of MicrosoftDNS_ResourceRecord. |

For more information about many WMI classes, see the WMI SDK documentation on MSDN.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Microsoft Platform SDK" on MSDN for more information about many WMI classes that are associated with the DNS Server service.
- "Group Policy Settings Reference" in the Tools and Settings Collection for information about Group Policy settings that are associated with the DNS Client service.
- "Registry Reference" in the Tools and Settings Collection for information about registry entries that are associated with DNS.

# Domain Controller Roles

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

A domain controller is a server that is running a version of the Windows Server® operating system and has Active Directory® Domain Services installed.

| **Note** |
| --- |
| In Windows Server 2003 and Microsoft Windows® 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008, the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

When you install Windows Server on a computer, you can choose to configure a specific server role for that computer. When you want to create a new forest, a new domain, or an additional domain controller in an existing domain, you configure the server with the role of domain controller by installing AD DS.

By default, a domain controller stores one domain directory partition consisting of information about the domain in which it is located, plus the schema and configuration directory partitions for the entire forest. A domain controller that runs Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 can also store one or more application directory partitions. There are also specialized domain controller roles that perform specific functions in an AD DS environment. These specialized roles include global catalog servers and operations masters.

**Global Catalog Servers**

Every domain controller stores the objects for the domain in which it is installed. However, a domain controller designated as a global catalog server stores the objects from all domains in the forest. For each object that is not in the domain for which the global catalog server is authoritative as a domain controller, a limited set of attributes is stored in a partial replica of the domain. Therefore, a global catalog server stores its own full, writable domain replica (all objects and all attributes) plus a partial, read-only replica of every other domain in the forest. The global catalog is built and updated automatically by the AD DS replication system. The object attributes that are replicated to global catalog servers are the attributes that are most likely to be used to search for the object in AD DS. The attributes that are replicated to the global catalog are identified in the schema as the partial attribute set (PAS) and are defined by default by Microsoft. However, to optimize searching, you can edit the schema by adding or removing attributes that are stored in the global catalog.

The global catalog makes it possible for clients to search AD DS without having to be referred from server to server until a domain controller that has the domain directory partition storing the requested object is found. By default, AD DS searches are directed to global catalog servers.

The first domain controller in a forest is automatically created as a global catalog server. Thereafter, you can designate other domain controllers to be global catalog servers if they are needed.

**Operations Masters**

Domain controllers that hold operations master roles are designated to perform specific tasks to ensure consistency and to eliminate the potential for conflicting entries in the Active Directory database. AD DS defines five operations master roles: the schema master, domain naming master, relative identifier (RID) master, primary domain controller (PDC) emulator, and infrastructure master.

The following operations masters perform operations that must occur on only one domain controller in the forest:

- Schema master
- Domain naming master

The following operations masters perform operations that must occur on only one domain controller in a domain:

- Primary Domain Controller (PDC) emulator
- Infrastructure master

- Relative ID (RID) master

# Global Catalog Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

The global catalog is a distributed data repository that facilitates searching and logons in an Active Directory® Domain Services (AD DS) directory service forest. The global catalog allows users and applications to locate objects in any domain in the forest by searching on an important attribute of the object (for example, search on the user's last name to locate the user object or search on a building name to find a printer in that building). The global catalog facilitates logons by ensuring that membership in universal groups from all domains are represented in the user's access credentials. One or more domain controllers are designated as global catalog servers because they host the global catalog, which is built automatically by the AD DS replication system.

**In this subject**


- What Is the Global Catalog?
- How the Global Catalog Works
- Global Catalog Tools and Settings

# What Is the Global Catalog?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Common Global Catalog Scenarios
- Global Catalog Dependencies and Interactions
- Related Information


The global catalog is a distributed data repository that contains a searchable, partial representation of every object in every domain in a multidomain Active Directory Domain Services (AD DS) forest. The global catalog is stored on domain controllers that have been designated as global catalog servers and is distributed through multimaster replication. Searches that are directed to the global catalog are faster because they do not involve referrals to different domain controllers.

| **Note** |
| --- |
| In Windows Server® 2003 and Microsoft Windows® 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008, the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

In addition to configuration and schema directory partition replicas, every domain controller in a forest stores a full, writable replica of a single domain directory partition. Therefore, a domain controller can locate only the objects in its domain. Locating an object in a different domain would require the user or application to provide the domain of the requested object.

The global catalog provides the ability to locate objects from any domain without having to know the domain name. A global catalog server is a domain controller that, in addition to its full, writable domain directory partition replica, also stores a partial, read-only replica of all other domain directory partitions in the forest. The additional domain directory partitions are partial because only a limited set of attributes is included for each object. By including only the attributes that are most used for searching, every object in every domain in even the largest forest can be represented in the database of a single global catalog server.

| **Note** |
| --- |
| A global catalog server can also store a full, writable replica of an application directory partition, but objects in application directory partitions are not replicated to the global catalog as partial, read-only directory partitions. |

The global catalog is built and updated automatically by the AD DS replication system. The attributes that are replicated to the global catalog are identified in the schema as the partial attribute set (PAS) and are defined by default by Microsoft. However, to optimize searching, you can edit the schema by adding or removing attributes that are stored in the global catalog.

In Windows 2000 Server environments, any change to the PAS results in full synchronization (update of all attributes) of the global catalog. Later versions of Windows Server reduce the impact of updating the global catalog by replicating only the attributes that change.

In a single-domain forest, a global catalog server stores a full, writable replica of the domain and does not store any partial replica. A global catalog server in a single-domain forest functions in the same manner as a non-global-catalog server except for the processing of forest-wide searches.

**Common Global Catalog Scenarios**

The following events require a global catalog server:

- **Forest-wide searches.** The global catalog provides a resource for searching an AD DS forest. Forest-wide searches are identified by the LDAP port that they use. If the search query uses port 3268, the query is sent to a global catalog server.
- **User logon.** In a forest that has more than one domain, two conditions require the global catalog during user

authentication:

- In a domain that operates at the Windows 2000 native domain functional level or higher, domain controllers must request universal group membership enumeration from a global catalog server.
- When a user principal name (UPN) is used at logon and the forest has more than one domain, a global catalog server is required to resolve the name.

- Universal Group Membership Caching: In a forest that has more than one domain, in sites that have domain users but no global catalog server, Universal Group Membership Caching can be used to enable caching of logon credentials so that the global catalog does not have to be contacted for subsequent user logons. This feature eliminates the need to retrieve universal group memberships across a WAN link from a global catalog server in a different site.

| **Note** |
| --- |
| Universal groups are available only in a domain that operates at the Windows 2000 native domain functional level or higher. |

- Exchange Address Book lookups. Servers running Microsoft Exchange Server rely on access to the global catalog for address information. Users use global catalog servers to access the global address list (GAL).

**Search Requests**
Because a domain controller that acts as a global catalog server stores objects for all domains in the forest, users and applications can use the global catalog to locate objects in any domain within a multidomain forest without a referral to a different server.

When a forest consists of a single domain, every domain controller has a full, writable copy of every object in the domain and forest. However, it is important to retain the global catalog on at least one domain controller because many applications use port 3268 for searching. For example, if you do not have any global catalog servers, the **Search** command on the **Start** menu cannot locate objects in AD DS.

The replicas that are replicated to the global catalog also include the access permissions for each object and attribute. If you are searching for an object that you do not have permission to access, you do not see the object in the list of search results. Users can find only objects to which they are allowed access.

**User Logon Support**
In addition to its role as a search provider, in a forest that has more than one domain, the global catalog has a role as an identity source during the user logon process. Universal groups can provide access to resources outside of the user's domain. User principal names (UPNs) can specify a domain other than the domain of the user. By making universal group membership and UPN domain-user mapping information available on all global catalog servers, the global catalog provides the definitive source for groups that are capable of providing access in more than one domain and names that do not unequivocally identify the domain of the user.

**Universal Group Membership**
During the domain logon process, the user must be authenticated. During the authentication process, the user is validated (the domain controller verifies the identity of the user) and the user receives authorization data for access to resources. To provide authorization data of a user, the authenticating domain controller retrieves the security identifiers (SIDs) for all security groups of which the user is a member and adds these SIDs to the user's access token. In a forest that has more than one domain, the global catalog is the only location where memberships of all universal groups in that forest can be ascertained. For this reason, access to a global catalog server is required for successful authentication in a domain that can have universal groups.

The global catalog stores the membership (the **member** attribute) of only universal groups. The membership of other groups can be ascertained at the domain level.

Because a universal group can have members from domains other than the domain where the group object is stored and can be used to provide access to resources in any domain, only a global catalog server is guaranteed to have all universal group memberships that are required for authentication.

For example, a user might be a member of a universal group that has its group object stored in a different domain but provides access to resources in the user's domain. To ensure that the user can be authorized to access resources appropriately in this domain, the domain controller must have access to the membership of all universal groups in the forest.

If a global catalog server is not available, the user logon fails.

**User Principal Name**
A user principal name (UPN) is a logon name that takes the form of an e-mail address. A UPN specifies the user ID followed by a DNS domain name, separated by an "@" character (for example, jsmith@contoso.com). UPNs allow administrative management of the UPN suffix to provide logon names that:

- Match the user's e-mail name.
- Do not reveal the domain structure of the forest.

When a user account is created, the UPN suffix is generated by default as *userName@ DnsDomainName*, but it can be changed administratively. For example, in a forest that has four domains, the UPN suffix might be configured to map to the external DNS name for the organization. The **userPrincipalName** attribute of the user account identifies the UPN and is replicated to the global catalog.

When you use a UPN to log on to a domain, your workstation contacts a global catalog server to resolve the name because the UPN suffix is not necessarily the domain for which the contacted domain controller is authoritative. If the DNS domain name in the UPN suffix is not a valid DNS domain, the logon fails. Assuming the UPN suffix is a valid DNS name, the global catalog server returns the name of the AD DS domain to your workstation, which then queries DNS for a domain controller in that domain.

If a company has more than one forest and uses trust relationships between the domains in the different forests, a UPN cannot be used to log on to a domain that is outside the user's forest because the UPN is resolved in the global catalog of the user's forest.

**Universal Group Membership Caching**
Universal Group Membership Caching eliminates the need for a domain controller in a multidomain forest to contact a global catalog server during the logon process in domains where universal groups are available. Caching group membership reduces WAN traffic, which helps in sites where updating the cached group membership of security principals, including user and computer accounts, generates less traffic than replicating the global catalog to the site.

Use the following criteria to determine if a site is a good candidate for Universal Group Membership Caching:

- Number of users and computers in the site: The site has less than 500 combined users and computers, including transient users who log on occasionally but not on a regular basis. The cache of a user who logs on once continues to be updated periodically for 180 days after the first logon. A general limit of 500 membership caches can be updated at a time. If greater than 500 security principals have cached group memberships, some caches might not be updated.
- Number of domain controllers: Each domain controller performs a refresh on every user in its site once every eight hours. Depending on the number of domains in the forest, 500 security principals and two domain controllers could generate more WAN traffic than placing a global catalog server in the site. Therefore, you need to rationalize the WAN costs when exceeding 500 security principals and two domain controllers.
- Tolerance for high latency in group updates. Because domain controllers in the site where Universal Group Membership Caching is enabled update the membership caches every eight hours, and because credentials are always taken from the cache, updates to group memberships are not reflected in the security principal's credentials for up to eight hours.

**Address Book Lookups**
Exchange Server uses the global catalog to store mail recipient data that enables clients in a forest to send and receive e-mail messages.

**Global Catalog Dependencies and Interactions**
Global catalog servers have the following dependencies and interactions with other Windows Server technologies:

- AD DS installation. When AD DS is installed on the first domain controller in a forest, the installation application creates that domain controller as a global catalog server.
- AD DS replication. The global catalog is built and maintained by AD DS replication:
    - Subsequent to forest creation, when a domain controller is designated as a global catalog server, AD DS

replication automatically transfers PAS replicas to the domain controller, including the partial replica of every domain in the forest other than the local domain.

    - To facilitate intersite replication of global catalog server updates, AD DS replication selects global catalog servers as bridgehead servers whenever a global catalog server is present in a site and domains that are not present in the site exist in other sites in the forest.

- Domain Name System (DNS). Global catalog server clients depend on DNS to provide the IP address of global catalog servers. DNS is required to advertise global catalog servers for domain controller location.

- Net Logon service. Global catalog advertisement in DNS depends on the Net Logon service to perform DNS registrations. When replication of the global catalog is complete, or when a global catalog server starts, the Net Logon service publishes service (SRV) resource records in DNS that specifically advertise the domain controller as a global catalog server.

- Domain controller Locator: When a global catalog server is requested (by a user or application that launches a search over port 3268, or by a domain controller that is authenticating a user logon), the domain controller Locator queries DNS for a global catalog server.

In the following diagram, global catalog interactions include tracking a global catalog server through the following interactions, which are indicated by boxes:

- **Active Directory installation of a new forest:** Global catalog creation occurs during AD DS installation of the first domain controller in the forest.

- **Net Logon registration:** Resource records are registered in DNS to advertise the domain controller as a global catalog server.

- **AD DS replication:**

    - When a new domain controller (DC2) is created and an administrator designates it as a global catalog server, replication of the PAS from DC1 occurs.

    - DC1 in DomainA replicates changes for DomainA to DC2, and DC2 replicates updates to data for DomainB to DC1.

- **DC location:** The dotted lines enclose the processes whereby two clients locate a global catalog server by querying DNS:

    - A through C: (A) ClientX sends a query to the global catalog, which prompts (B) a DNS query to locate the closest global catalog server, and then (C) the client contacts the returned global catalog server DC2 to resolve the query.

    - 1 through 5: (1) ClientY logs on to the domain, which prompts (2) a DNS query for the closest domain controllers. (3) ClientY contacts the returned domain controller DC3 for authentication. (4) DC3 queries DNS to find the closest global catalog server and then (5) contacts the returned global catalog server DC2 to retrieve the universal groups for the user.

**Interactions with Other Windows Technologies**




The global catalog solves the problem of how to locate domain data that is not stored on a domain controller in the domain of the client that requires the information. By using different ports for standard LDAP queries (port 389) and global catalog queries (port 3268), AD DS effectively separates forest-wide queries that require a global catalog server from local, domainwide queries that can be serviced by the domain controller in the user's domain.

**Related Information**

- Active Directory Replication Topology Technical Reference
- DNS Support for Active Directory Technical Reference
- Active Directory Schema Technical Reference

# How the Global Catalog Works

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Global Catalog Architecture
- Global Catalog Protocols
- Global Catalog Interfaces
- Global Catalog Physical Structure
- Global Catalog Processes and Interactions
- Network Ports Used by the Global Catalog
- Related Information


In a multidomain Active Directory® Domain Services (AD DS) forest, the global catalog provides a central repository of domain information for the forest by storing partial replicas of all domain directory partitions. These partial replicas are distributed by multimaster replication to all global catalog servers in a forest.

| **Note** |
| --- |
| In Windows Server® 2003 and Microsoft Windows® 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008 and, the directory service is named Active Directory Domain Services . The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

The global catalog makes the directory structure within a forest transparent to users who perform a search. For example, if you search for all printers in a forest, a global catalog server processes the query in the global catalog and then returns the results. Without a global catalog server, this query would require a search of every domain in the forest.

During an interactive domain logon, the domain controller authenticates the user by verifying the user's identity, and also provides authorization data for the user's access token by determining all groups of which the user is a member. Because the global catalog is the forestwide location of the membership of all universal groups, access to a global catalog server is a requirement for authentication in a multidomain forest. A global catalog server is also required for applications such as Microsoft Exchange Server.

For more information about other reasons why a global catalog server is required, see the FAQ on the Ask the Directory Services Team blog (http://blogs.technet.com/b/askds/archive/2011/09/30/friday-mail-sack-super-slo-mo-edition.aspx#gc).

An ideal distribution of the global catalog is to have at least one global catalog server in each AD DS site. When a global catalog server is available in a site, the authenticating domain controller is not required to communicate across a WAN link to retrieve global catalog information. It is recommended make all domain controllers be global catalog servers.

In branch office scenarios, it is often not feasible to deploy a global catalog server in every branch site. At the same time, it is not cost effective to contact a global catalog server over a WAN link for every logon that occurs in the site. On domain controllers that are running Windows Server 2003 or later, universal group membership can be cached so that the domain controller must connect to a global catalog server across a WAN link only for initial logons in the site; thereafter, universal group membership can be checked from a local cache. Search clients, however, must always connect to global catalog servers across the WAN if no global catalog server exists in the client's site.

This subject describes the functionality of the global catalog and the replication of objects to global catalog servers in an AD DS forest.

**Global Catalog Architecture**

Global catalog server architecture differs from non-global catalog server architecture in its use of the nonstandard LDAP port 3268, which directs queries to the global catalog. Queries over this port are formed the same way as any LDAP query, but AD DS varies the search behavior according to the port that is used: queries over port 3268 target the global catalog directory partitions (including the read-only domain directory partitions and the one writable domain directory partition for which the server is authoritative), and queries over port 389 target only the writable domain, configuration, application, and schema directory partition replicas stored by the global catalog server in its role as a domain controller. In addition, domain controllers use the proprietary replication interface when they contact global catalog servers to retrieve universal group membership during client logons.

Search clients include Exchange Address Book clients, which use the client MAPI provider Emsabp32.dll to look up e-mail addresses in the global catalog. The client-side MAPI provider communicates with the server through the proprietary Name Service Provider Interface (NSPI) RPC interface.

Windows NT clients use Net APIs to communicate with the Security Accounts Manager (SAM) on the primary domain controller (PDC) emulator. The PDC emulator, a domain controller operations master role in AD DS domains, manages search and replication communication with clients that are running Windows NT.

The relationships between these architectural components are shown in the following diagram. Descriptions for the major components are provided in the subsequent table.

**Global Catalog Architecture**




For a more detailed description of LDAP and replication client-server architecture, see "How the Active Directory Replication Model Works."

**Global Catalog Architecture Components**

| Component | Description |
|---|---|
| Clients | Global catalog clients, including search clients and Address Book clients, as well as domain controllers performing replication and universal group security identifier (SID) retrieval during logon in a multidomain forest. |
| Network | The physical IP network. |
| Interfaces | LDAP over port 389 for read and write operations and LDAP over port 3268 for global catalog search operations. NSPI and replication (REPL) use proprietary RPC protocols. Retrieval of universal group membership occurs over RPC as part of the replication RPC interface. Windows NT 4.0 clients and backup domain controllers (BDCs) communicate with AD DS through the Security Accounts Manager (SAM) interface. |
| Directory System Agent (DSA) | The directory service component that runs as Ntdsa.dll on each domain controller, providing the interfaces through which services and processes gain access to the directory database. |

| | |
|---|---|
| Extensible Storage Engine (ESE) | The directory service component that runs as Esent.dll. ESE manages the tables of records that comprise the directory database. |
| Ntds.dit database file | The AD DS data store. |

#### Global Catalog Protocols

The following diagram shows the four interfaces into AD DS and the protocols that package the data according to their specific applications. These protocols and interfaces are the same for all domain controllers and are not specific to global catalog servers. The significance for the global catalog server is that domain controllers use the proprietary RPC replication protocol not only for replication, but also to contact the global catalog server when retrieving universal group membership information and when updating the group membership cache when Universal Group Membership Caching is enabled.

**Global Catalog Protocols**




The protocols are described in the following table.

**Global Catalog Protocols**

| Protocol | Description |
|---|---|
| Lightweight directory access protocol (LDAP) | The primary directory service protocol that specifies directory communications. It runs directly over TCP/IP, and it can also run over User Datagram Protocol (UDP) connectionless transports (UDP access is primarily used by the domain controller Locator process and can also be used to query the rootDSE). Clients use LDAP to query, create, update, and delete information that is stored in a directory service over a TCP connection through the TCP default port 389. Global catalog clients can use LDAP to query AD DS over a TCP connection through the TCP port 3268. AD DS supports LDAP v2 (RFC 1777) and LDAP v3 (RFC 2251). LDAP v3 is an industry standard that can be used with any directory service that implements the LDAP protocol. LDAP is the preferred and most common way of interacting with AD DS. |
| Remote procedure call (RPC) | Protocol for replication (REPL) and domain controller management communications (including global catalog server interactions), NSPI address book communications, and SAM-related communications. RPC is a powerful, robust, efficient, and secure interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality residing in a different process. That different process can be on the same computer, on the local area network (LAN), or across the Internet. |
| Simple mail transfer protocol (SMTP) | Protocol for replication communications when a permanent, "always on" network connection does not exist between two domain controllers. SMTP is used to transport and deliver messages based on specifications in Request for Comments (RFC) 821 and RFC 822. SMTP can replicate only the configuration and schema directory partitions and global catalog read-only replicas (not writable domain data). |

For more information about AD DS protocols, see "How the Data Store Works."

#### Global Catalog Interfaces

Interfaces for global catalog servers are the AD DS data store interfaces, shown in the previous figure and described in the following table.

**Global Catalog Data Store Interfaces**

| Interface | Description |
|---|---|

| | |
|---|---|
| LDAP | The primary interface for AD DS access. Directory clients use LDAP v3 to connect to the DSA through the LDAP interface. The LDAP interface is part of Wldap32.dll. LDAP v3 is backward compatible with LDAP v2. |
| REPL | The replication management interface that provides functionality for finding data about domain controllers, converting the names of network objects between different formats, manipulating service principal names (SPNs) and DSAs, and managing replication of servers. |
| NSPI/MAPI | Name Service Provider Interface (NSPI) by which Messaging API (MAPI) clients access AD DS. Messaging clients gain access to AD DS by using MAPI address book providers. For compatibility with existing messaging clients, AD DS supports the NSPI/RPC address book provider, which provides directory access, for example, to find the telephone number of a user. |
| SAM | Proprietary interface for connecting to the DSA on behalf of clients that run Windows NT 4.0 or earlier. These clients use Windows NT 4.0 networking APIs to connect to the DSA through SAM. Replication with Windows NT 4.0 backup domain controllers (BDCs) occurs through the SAM interface as well. |

**Note**

The NSPI (MAPI) interface is provided only for support of legacy Microsoft Outlook clients. Development against this interface is no longer supported.

For more information about AD DS data store interfaces, see "How the Data Store Works."

**Global Catalog Physical Structure**

AD DS is a distributed directory service in which data is stored as replicas on multiple domain controllers to provide a virtual database that maintains consistency through AD DS replication. Domain controllers provide the domainwide distribution of directory data. Global catalog servers provide the forestwide distribution of directory data in a multidomain forest.

**Global Catalog Partial Attribute Set**

In its role as a domain controller, a global catalog server stores one domain directory partition that has writable objects with a full complement of writable attributes. In its role as global catalog server, it also stores the objects of all other domain directory partitions in a multidomain forest as read-only objects with a partial set of attributes. The set of attributes that are marked for inclusion in the global catalog are called the partial attribute set (PAS). An attribute is marked for inclusion in the PAS as part of its schema definition.

Objects in the schema that define an attribute are **attributeSchema** objects, which themselves have an attribute **isMemberOfPartialAttributeSet**. If the value of that attribute is TRUE, the attribute is replicated to the global catalog. The replication topology for the global catalog is generated automatically by the Knowledge Consistency Checker (KCC), a built-in process that implements a replication topology that is guaranteed to deliver the contents of every directory partition to every global catalog server.

The attributes that are replicated to the global catalog by default include a base set that have been defined by Microsoft as the attributes that are most likely to be used in searches. Administrators can use the Microsoft Management Console (MMC) Active Directory Schema snap-in to specify additional attributes to meet the needs of their installation. In the Active Directory Schema snap-in, you can select the **Replicate this attribute to the global catalog** check box to designate an **attributeSchema** object as a member of the PAS, which sets the value of the **isMemberOfPartialAttributeSet** attribute to TRUE.

**Domain Controller and Global Catalog Server Structure**

The physical representation of global catalog data is the same as all domain controllers: the Ntds.dit database stores object attributes in a single file. On a domain controller that is not a global catalog server, the Ntds.dit file contains a full, writable replica of every object in one domain directory partition for its own domain, plus the writable configuration and schema directory partitions.

**Note**

The schema directory partition is writable only on the domain controller that is the schema operations master for the forest.

The following diagram shows the physical representations of the global catalog as a forestwide resource that is distributed as a database on global catalog servers.

**Global Catalog Physical Structure**




As shown in the preceding diagram, a global catalog server stores a replica of its own domain (full and writable) and a partial, read-only replica of all other domains in the forest. All directory partitions on a global catalog server, whether full or partial, are stored in the directory database file (Ntds.dit) on that server. That is, there is not a separate storage area for global catalog attributes; they are treated as additional information in the directory database of the global catalog server.

The following table describes the physical components of the diagram.

**Global Catalog Server Physical Components**

| Physical Component | Description |
|---|---|
| Active Directory forest | The set of domains that comprise the AD DS logical structure and that are searchable in the global catalog. |
| Domain controller | Server that stores one full, writable domain directory partition plus forestwide configuration and schema directory partitions. Global catalog servers are always domain controllers. |
| Global catalog server | Domain controller that stores one full, writable domain plus forestwide configuration and schema directory partitions, as well as a partial, read-only replica of all other domains in the forest. |
| Ntds.dit | Database file that stores replicas of the AD DS objects held by any domain controller, including global catalog servers. |

**Global Catalog Processes and Interactions**

In addition to its activities as a domain controller, the global catalog server supports the following special activities in the forest:

- User logon: In a multidomain forest, domain controllers must contact a global catalog server to retrieve any SIDs of universal groups that the user is a member of. Additionally, if the user specifies a logon name in the form of a UPN, the domain controller contacts a global catalog server to retrieve the domain of the user.

- Universal and global group caching and updates: In sites where Universal Group Membership Caching is enabled, domain controllers that are running Windows Server 2003 or later cache group memberships and keep the cache updated by contacting a global catalog server.

- Global catalog searches: Clients can search the global catalog by specifying port 3268 or by using search applications that use this port. Search activities include:

    - Validation of references to non-local directory objects. When a domain controller holds a directory object with an attribute that references an object in another domain, this reference is validated by contacting a global catalog server.

    - Exchange Address Book lookups: Exchange Server uses AD DS as the address book store. Outlook clients query the global catalog to locate Address Book information.

- Global catalog server creation and advertisement: Global catalog servers register global-catalog-specific service (SRV) resource records in DNS so that clients can locate them according to site. If no global catalog server is available in the site of the user, a global catalog server is located in the next closest site, according to the cost matrix that is generated by the KCC from site link cost settings.

- Global catalog replication: Global catalog servers must either have replication partners for all domains or be able to replicate with another global catalog server. When changes to the PAS occur on, and are replicated between, domain controllers that are running Windows Server 2003 or later, only the updated attributes are replicated. Changes to the PAS that occur on domain controllers that are running Windows 2000 Server prompt a full synchronization of the entire global catalog (all attributes in the PAS are replicated anew to all global catalog servers). For more information about PAS replication, see Global Catalog Replication.

**User Logon**

When a domain user logs on interactively to a domain, the contacted domain controller must retrieve information from a global catalog server under the following conditions:

- The user's domain is Windows 2000 native domain functional level or higher. In this case, the user might belong to a universal group whose object is stored in a different domain.

- The user's logon name is a user principal name (UPN), which has the format *sAMAccountName@DNSDomainName*. In this case, the DNS domain suffix is not necessarily the user's domain and the identity of the user's domain must be retrieved from a global catalog server.

**Universal Group SID Retrieval**

A universal group is a security group that is available at the Windows 2000 native domain functional level or higher. During interactive user logon, the authenticating domain controller retrieves the SIDs that the user's workstation requires to build the access token for the user. To retrieve the SIDs of all universal groups to which the user belongs, the authenticating domain controller must contact a global catalog server. If a global catalog server is not available in the site when a user logs on to a domain in which universal groups are available, the computer will use cached credentials to log the user on if the user has previously logged on to the domain from the same workstation. If the user has not previously logged on to the domain from the same computer, the user can log on to only the local computer.

Of the three group types that are used to allow access to resources in a forest (domain local, global, and universal), only universal groups have their membership replicated to the global catalog. The values of the **member** attribute of universal group objects that exist in all domains must be available to an authenticating domain controller because:

- Universal groups can contain members, including individual user accounts and global groups, from any domain in the forest. Therefore, the user who is logging on might have a membership in a universal group that exists in a different domain.

- Universal groups can be added to access control lists in any domain in the forest. Therefore, the user might have permissions on objects in this domain by virtue of membership in a universal group that exists in a different domain.

Contrast this requirement with the domain local group membership. Domain local groups can also have members from other domains; however, domain local groups can be added to access control lists in only the domain in which they are created. Therefore, a domain controller can always determine a user's membership in all domain local groups that are required for authorization in its own domain.

For global groups, although these groups can be added to access control lists in any domain, they can contain accounts from only the domain in which they are created. Therefore, the global group membership of the user who is logging on can always

be checked locally. However, global groups can be members of universal groups that exist in different domains. When group nesting is in effect (which has the same availability criteria as availability of universal groups), being a member of a global group that is itself a member of a universal group can give the user access to resources other than those allowed by membership in the global group alone.

During the logon process, the authenticating domain controller retrieves a list of global group SIDs from the user's domain. If universal groups are available in the user's domain, the domain controller passes the list of global group SIDs to the nearest global catalog server. The global catalog server responds as follows:

- Enumerates the **member** attribute of all universal groups in the forest and adds universal group SIDs to the user's SID list as follows:
    - All universal groups that contain the user's SID.
    - All universal groups that contain the SID of any of the global groups in the user's SID list.
- Returns the list, including both universal group SIDs and global group SIDs, to the domain controller.

The authenticating domain controller sends the list of SIDs that is returned from the global catalog server to the user's computer, along with domain local group SIDs from the user's domain. The user's local computer, which creates the access token for the user, adds the returned SIDs to the access token.

For more information about how domain controllers cache group membership, see Universal Group Membership Caching. For more information about how SIDs are retrieved and used in access tokens, see "How Access Tokens Work." For more information about the authentication process, see "How the Kerberos Version 5 Authentication Protocol Works." For more information about the logon process, see "How Interactive Logon Works."

**UPN Logon**

A global catalog server resolves UPNs when the authenticating domain controller does not have knowledge of the account. For example, if a users account is located in noam.corp.contoso.com and the user logs on with a UPN of JSmith@contoso.com, the domain name in the UPN suffix does not match the user's domain. In the Windows Server 2003 and Windows 2000 Server logon screen, you can either type your user name (**sAMAccountName**) and select the domain name from the drop-down list, or you can use a UPN. If you use a UPN, as soon as you type the @ sign, the domain list becomes unavailable. In this case, domain controllers running Windows Server 2003 or Windows 2000 Server take the domain name from the UPN suffix. The UPN suffix is often (usually) different from the home domain of the user. In this case, the responding domain controller must contact a global catalog server to determine the domain of the user.

The following diagram shows the general sequence that occurs when a user logs on to a domain by using a UPN.

**UPN Logon and Global Catalog Interaction**




1. Because the domain of the user is not necessarily the same as the UPN suffix, the domain controller Locator contacts the nearest domain controller according to the site of the client computer.
2. The contacted domain controller determines whether the DNS name in the UPN suffix is the domain for which the domain controller is authoritative.

- If the domain name in the UPN suffix matches the domain of the domain controller, the domain controller attempts to process the client authentication. If the user name is not found, the domain controller contacts a global catalog server.

- If the DNS name in the UPN suffix does not match the domain of the domain controller, the domain controller contacts a global catalog server.

3. The global catalog server uses the **userPrincipalName** attribute of the user object to look up the distinguished name of the user object, and returns this value to the domain controller.

4. The domain controller extracts the domain name from the distinguished name of the user and returns this value to the client.

5. The client requests a domain controller for its domain.

If a forest trust exists between two forests, the default form of a UPN (*sAMAccountName@DnsDomainName*) is used for authentication in a different forest. If you create a forest trust and the second-level DNS domain name (for example, contoso.com) exists in both forests, the New Trust Wizard detects the conflict and only one forest is authoritative for that name suffix.

If NTLM-based trust relationships are created between the domains in different forests—as is required for a trust relationship between a domain in an Active Directory forest and a Windows NT 4.0 domain, between domains in two Windows 2000 Server forests, or between a domain in a Windows 2000 Server forest and a domain in a Windows Server 2003 forest—a UPN cannot be used to log on to a trusting domain that is outside the forest because the UPN is resolved in the global catalog of only one forest.

**Logon Process in a Single-Domain Forest**
In a single-domain forest, all domain controllers can service all logon requests, including UPN logons, without requiring a global catalog server. However, only domain controllers that are configured as global catalog servers can respond to LDAP traffic over port 3268.

For more information about the logon process, see "Interactive Logon Technical Reference." For more information about forest trust relationships, see "Domain and Forest Trusts Technical Reference."

**Universal Group Membership Caching**
In multidomain forests where remote sites do not have a global catalog server, the need to contact a global catalog server over a potentially slow WAN connection can be problematic. On domain controllers that are running Windows Server 2003 or later, the Universal Group Membership Caching feature is available by default (does not require a specific functional level or domain mode), although it must be enabled on a per-site basis.

When enabled, this feature allows a domain controller to cache global group SIDs and universal group SIDs that it retrieves from a global catalog server so that future logons do not require contacting a global catalog server. This storage is referred to as "caching," but the memberships are actually stored in a non-volatile AD DS value. The memberships that are written to this value are not lost as a result of a restart or power outage. For the purposes of this discussion, the term "cache" refers to this value. Group membership is cached for user accounts and computer accounts.

Caching group memberships in branch site locations has the following potential benefits:

- Faster logon times because authenticating domain controllers no longer need to contact a global catalog server to obtain universal group membership.

- Higher availability because logon is still possible if the WAN link to the site of the global catalog server is unavailable.

- No need to upgrade the hardware of existing domain controllers to handle the extra system requirements necessary for hosting the global catalog.

- Minimized network bandwidth usage because a branch site domain controller does not have to replicate all of the objects located in the global catalog.

**Enabling Universal Group Membership Caching**

Universal Group Membership Caching can be enabled for a site by using the Active Directory Sites and Services MMC snap-in to edit the properties of the NTDS Site Settings object (CN=NTDS Site Settings,CN=*TargetSiteName*,CN=Sites,CN=Configuration,CN=*ForestRootDomain*). In Active Directory Sites and Services, if you click a site object, the NTDS Site Settings object for the site is visible in the details pane. Right-click the **NTDS Site Settings** object and then click **Properties**. In the **NTDS Site Settings Properties** dialog box, click **Enable Universal Group Membership Caching**.

| **Note** |
| --- |
| The **options** attribute of the NTDS Site Settings object, which controls this feature, has a default value of 0. When only the Universal Group Membership Caching option is enabled, the attribute value is 32. However, this attribute is a bit field, so its full functionality is derived from computing a bitwise AND of all of the bits that are set. |

When the feature is enabled for a site, domain controllers in the site cache both universal group membership and global group membership for first-time logons and keep the cache updated thereafter. The feature allows specifying the site from which to retrieve group membership. In the **NTDS Site Settings Properties** dialog box, you can use the **Refresh cache from** list to specify the site to use. The **msDS-Preferred-GC-Site** attribute stores the distinguished name of the specified site and controls this setting.

If no site is specified, the closest-site mechanism uses the cost setting on the site link to determine which site has the least-cost connection to contact a global catalog server.

If the user has not logged on to the domain previously and a global catalog server is not available, the user can log on to only the local computer.

| **Note** |
| --- |
| If a user is the Administrator in the domain (Builtin Administrator account), the user can always log on to the domain, even when a global catalog server is not available. |

For more information about closest-site mechanisms, see "How Active Directory Replication Topology Works" and "How DNS Support for Active Directory Works."

**Global Group and Universal Group SID List**

Although the feature is named Universal Group Membership Caching, in fact the domain controller caches both universal group SIDs and global group SIDs. As described in "Universal Group SID Retrieval" earlier in this subject, the authenticating domain controller retrieves the list of universal group SIDs from the global catalog server. The domain controller sends the list of global group SIDs to the global catalog server so that universal group membership can be ascertained for these groups as well as for the user or computer account itself. Therefore, both global group SIDs and universal group SIDs are included in the list of SIDs that is returned from the global catalog server. The domain controller caches the entire list of SIDs. When the account attempts subsequent logons in the site, the universal group and global group SIDs for the account are obtained from the cache. Domain local group SIDs are always obtained from the local directory database.

After a security principal has logged on in a site that has Universal Group Membership Caching enabled, the group cache for the account on the authenticating domain controller is immediately populated. However, it can take up to eight hours for other domain controllers in the same site to populate the group cache. During this time, if the account is authenticated by a domain controller that has not populated the account's cache, a global catalog server must be contacted for the logon to proceed. After eight hours, all domain controllers that are running Windows Server 2003 or later in the site can process all subsequent logons by using the cached membership.

**Group Cache Communication Between Domain Controllers**

Although the cache itself is not replicated between domain controllers, knowledge that an account has logged on in the site is replicated to all other domain controllers in the domain by means of a site affinity attribute (**msDS-Site-Affinity**) of the respective user or computer object. This multivalued attribute identifies the sites in which the account has logged on and includes a timestamp for each site. The domain controllers in the domain use the replicated site affinity attribute to determine which account memberships are cached for their site, and then independently populate their copy of each account's cache by contacting a global catalog server. For more information about how the cache is populated, see How the Cache is Populated at First Logon and How the Cache is Refreshed. For more information about how this attribute is updated, see Group Cache Storage.

**Cache Refresh and the Availability of Group Changes**

The caching of group SIDs in this way, including both universal group and global group SIDs, can lead to unexpected behavior when an administrator modifies the universal group or global group membership for an account and expects the change to be reflected on all domain controllers in the domain after the first replication cycle following the change. Even if the change is

made on the domain controller that authenticates the account or has been received through replication, the membership in the cache is used instead of the value of the object **member** attribute. By default, domain controllers update the membership cache for accounts in the site every eight hours. As a result, changes to the global group or universal group membership of an account can take up to eight hours to be reflected on domain controllers in a site where Universal Group Membership Caching is enabled.

To refresh the cache, domain controllers running Windows Server 2003 or later send a universal group membership confirmation request to a global catalog server. There is no limit to the number of accounts that can be cached, but a maximum of 500 account caches can be updated during any cache refresh.

| **Note** |
| --- |
| Universal Group Membership Caching can be enabled in a site that has domain controllers that are running Windows 2000 Server. If Universal Group Membership Caching is enabled in such a site, users might experience inconsistent group membership, depending on which domain controller authenticates them. For this reason, it is recommended that you either upgrade all domain controllers that are running Windows 2000 Server to Windows Server 2003 or later when group caching is enabled in a site, or remove them. |

Because the group memberships are cached, there is a period of latency before group membership changes are reflected in an account's access token. When group membership changes, the changes are not reflected in the access token until the following events have occurred (in order):

1. The changes are replicated to the global catalog server that is used for the refresh of the cache.

2. The cache on the domain controllers in the account's site is refreshed. Although the cache refresh is not a replication event, the process uses the site link schedule. Therefore, a closed site link schedule postpones the cache refresh until the schedule opens.

3. The user has logged off and back on again (user account is authenticated) or the computer has restarted (computer account is authenticated).

When an access token is created during logon, the token contents are static until that user logs off and logs on again. Furthermore, as long as Universal Group Membership Caching is enabled, an account's memberships are cached, and the cache entry has not expired, the cache entry is used during logon. If changes have been made to group membership and the refresh task has not run, the changes are not reflected until either the cache entry expires or the refresh task runs and processes the cache entry.

The length of the latency period depends on when the next refresh task is scheduled to run. The refresh task reschedules itself for its next refresh during each current refresh run, as follows:

- Beginning with the current time plus the registry-configured refresh interval, the domain controller consults the replication schedule on the site link that connects its site to the site of the closest (or designated) global catalog server.

- If the site link schedule allows replication at the projected time, the refresh task is scheduled to run at this time.

- If the site link schedule does not allow replication at the projected time, the scheduling algorithm steps forward one minimum replication interval (15 minutes) and checks the schedule again.

- This process is repeated until an open replication interval is found. If no open interval is found in the seven-day schedule, the refresh task is scheduled to run by using a time calculated as the current time plus the registry-configured refresh interval. In this case, event ID 1671 is logged as a warning message that indicates the group membership cache refresh task was unable to find the next available time slot of connectivity to the site of the global catalog server.

If faster updates are required, an administrator can initiate a cache refresh manually on the domain controllers in the user's site. For more information about refreshing the user cache, see Registry Settings that Affect Cache Refresh and Site Affinity Limits.

**Determining the Site to Use for Populating and Refreshing the Cache**

You can designate a site from which to initially populate and subsequently refresh the group membership cache. The Universal Group Membership Caching feature user interface (UI) contains an option to select a site from the list of existing sites. When a

site has been selected and the cache on a domain controller must be populated for the first time or updated, the domain controller contacts a global catalog server in the designated site. If no site is designated, site link costs are evaluated to determine the lowest-cost site that contains a global catalog server. The site link cost matrix is supplied by the Intersite Messenger (ISM) service.

The UI that you use to designate a preferred site for cache refresh does not check for the presence of a global catalog server in the selected site. Therefore, it is possible to designate a refresh site that does not contain a global catalog server. In this case, or in any case where a refresh site is designated but a global catalog server does not respond, the domain controller uses the site link cost matrix and logs event ID 1668 in the Directory Service log in Event Viewer, which indicates that the group membership cache refresh task did not locate a global catalog in the preferred site, but was able to find a global catalog in the following available site. The event lists the named preferred site and the actual site that was used.

**Group Cache Storage**

Cached group membership is stored as additional attributes of user and computer objects. Three **attributeSchema** objects in the AD DS schema for the **user** object class (and inherited by the **computer** object class) support this feature:

- **msDS-Cached-Membership**: (cached membership) A binary large object (BLOB) that contains both universal and global group memberships (the group SIDs) for the user. This attribute has the following characteristics:
    - Is single valued.
    - Is not indexed.
    - Can be deleted.
    - Cannot be written.
    - Is not replicated.
- **msDS-Cached-Membership-Time-Stamp**: (last refresh time) Contains the time that the cached membership was last updated, either by the first logon or by a refresh. This attribute is used for the “staleness” check. The maximum period that is tolerated when using a cached group membership is called the staleness interval. The staleness interval, set in the registry as 7 days, is measured against the current time and the last refresh time. If the timestamp indicates that the cache is older than the staleness interval allows, the cached membership is invalidated and a global catalog server is required for logon. This attribute has the following characteristics:
    - Is large integer, time valued.
    - Is indexed.
    - Can be deleted.
    - Cannot be written.
    - Is not replicated.
- **msDS-Site-Affinity**: Identifies the site(s) where the account has logged on plus a timestamp that indicates the start time for the cached logon in the respective site. Presence of a value in this attribute causes automatic population of group memberships and refresh at every refresh interval. When a domain controller refreshes its cached memberships (every 8 hours by default), the timestamp is used for removing accounts from the cache that have not logged on within the site affinity time limit (the cache expires). To avoid replication of this attribute every time the account logs on, the timestamp is updated only when the age exceeds 50 percent of the age limit that is set in the registry (180 days, by default) and one of the following actions occurs:
    - The account logs on and is authenticated by a domain controller.
    - A user changes his or her account password. This update ensures that users who go for extended periods without logging on have their site affinity values updated.

This attribute has the following characteristics:

- Is multivalued.
- Is indexed.
- Can be deleted.
- Can be written.
- Is replicated.

| Note |
| --- |
| You can use ADSI Edit in Windows Support Tools to clear the cached entries for an account by deleting the values in **msDS-Cached-Membership** and **msDS-Cached-Membership-Time-Stamp** from the user or computer object. The attribute values are repopulated at the next logon or cache refresh, whichever comes first. |

**Registry Settings that Affect Cache Refresh and Site Affinity Limits**

Registry settings on each domain controller determine the limits that are imposed on group membership caches. The entries in the following table are under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**, unless otherwise noted, and they can be used to manage the cache.

Changes to these registry settings take effect the next time the refresh task runs.

| Note |
| --- |
| In the following table, some of the entry names contain the string "(minutes)". Note that this string is part of the entry name and must be included when creating the entry. For example:<br>• The value name **Cached Membership Refresh Interval (minutes)** is correct.<br>• The value name **Cached Membership Refresh Interval** is incorrect. |

- 
- 
- 

**Registry Entries Used to Configure Caching Behavior**

| Registry Entry | Type | Default Value | Notes |
| --- | --- | --- | --- |
| **Cached Membership Site Stickiness (minutes)** | DWORD | 172800<br><br>(Value is in minutes. This setting equals 180 days) | Defines how long the site affinity will remain in effect. The site affinity value is updated when half of the period defined by this value has expired. If an account has not logged on with a domain controller for a period of one half of this value or longer, the account is removed from the list of accounts whose memberships are being refreshed. The default value is recommended. |

| | | | |
|---|---|---|---|
| **Cached Membership Staleness (minutes)** | DWORD | 10080<br><br>(Value is in minutes. This setting equals 7 days) | Determines the maximum staleness value when using cached group membership. The account cannot log on if the cached membership list is older than the staleness value and if no global catalog server is available. The default value is recommended. |
| **Cached Membership Refresh Interval (minutes)** | DWORD | 480<br><br>(Value is in minutes. This setting equals 8 hours) | Defines the length of time between group membership cache refreshes. This value should be changed to synchronize once a day (1440 minutes). For dial-up connections, you might want a higher value than 24 hours. Lowering the value to increase the frequency of cache refresh is not recommended because it causes increased WAN traffic, potentially defeating the purpose of Universal Group Membership Caching. |
| **Cached Membership Refresh Limit** | DWORD | 500 | Defines the maximum number of user and computer accounts that are refreshed. Increase this setting only if event ID 1669 occurs in the event log or you have more than 500 users and computers in a branch. If the number of users and computers in a branch exceeds 500, a general recommendation is to either place a global catalog server in the branch or increase the Cached Membership Refresh Limit above 500. Be aware that increasing the limit might incur more WAN traffic than that caused by global catalog update traffic. |
| **SamNoGcLogonEnforceKerberosIpCheck**<br><br>This setting appears under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa**. | DWORD | 0 | By default, allows site affinity to be tracked for Kerberos logons that originate outside the site. A value of 1 configures SAM so it does not give site affinity to Kerberos logon requests that originate outside the current site. This option should be configured to 1 on domain controllers in the branch-sites to prevent logon requests from outside the site being given affinity for the local site. This setting prevents an account's affinity from being changed during the logon process when connecting to a Kerberos key distribution center (KDC) outside of the account's site. This setting applies only to Kerberos logons; it does not affect site affinity caching for NTLM logons from different sites. |

| | | | |
|---|---|---|---|
| **SamNoGcLogonEnforceNTLMCheck**<br><br>This setting appears under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa**. | DWORD | 0 | A value of 1 configures SAM to not give site affinity to NTLM logon requests that are network logon requests; it may not prevent caching for other logon types. This setting reduces the number of accounts with site affinity by excluding those that simply accessed network resources (by using NTLM). This option should not be enabled if you use older clients that must authenticate from outside the branch to local resources in the branch. |

**Note**

To minimize the probability that remote users will use the UGC-enabled branch domain controller for logon, remove the domain controller for site-less SRV records from DNS by using the **DnsAvoidRegisterRecords** parameter for **Netlogon**. For more information about using **DnsAvoidRegisterRecords**, see Microsoft Knowledge Base article 306602 (http://go.microsoft.com/fwlink/?LinkId=126645).

**Methods of Refreshing the Cached Memberships**

You can refresh cached memberships on a single domain controller.

For a one-time, immediate cache refresh:

- Use Ldp.exe to modify the operational attribute **updateCachedMemberships** on the rootDSE with a value of **1**. Adding a value of 1 to this attribute instructs the local domain controller to perform the update. If the site link schedule allows replication at the time you modify the attribute, this update occurs right away. This method is the preferred method for updating a single domain controller because it does not require restarting the domain controller. For information about using Ldp to modify this attribute, see the Note below. Ldp.exe can be installed from Windows Support Tools in Windows Server 2003 and Windows 2000 Server, and it is installed by default on domain controllers that run Windows Server 2008 and Windows Server 2008 R2.

  -or-

- Restart the domain controllers in the site to restart the cache refresh interval, which triggers a cache refresh.

- Use the following procedure to modify the **updateCachedMemberships** operational attribute. To perform this operation, the user needs the control access right "Refresh Group Cache for Logons" on the NTDS Settings object for the domain controller. Default access is granted to System, Domain Admins, and Enterprise Admins.

1. Start Ldp.exe and bind to the target domain controller where the cache reset is to be performed. (Do not select **Tree view** in the **View** menu.) When first binding to a domain controller with Ldp, the default location is rootDSE. You can view the attributes for rootDSE in the details pane. However, operational attributes are not listed.

2. On the **Browse** menu, click **Modify**.

3. In the **Modify** dialog box, in the **Edit Entry Attribute** box, type **updateCachedMemberships**. In the **Values** box, type **1**. Be sure to leave the **Dn** box blank.

4. Click **Enter**. The command should appear in the entry list.

5. Click **Run**. If the operation was successful, Ldp will report “Modified” in the output.

**Method of Clearing the Cached Memberships**

You can clear all cached memberships on all domain controllers in a site. However, doing so can affect performance. The need to clear the cached memberships might arise due to too many cached accounts, causing inability to refresh all account caches during a single cache refresh. For example, sites that have many transient accounts might exceed the 500-account refresh limit.

If you have more than 500 accounts cached and you want to clear them for all domain controllers in the site, you can do so by editing the registry.

| **Note** |
| --- |
| If you must edit the registry, use extreme caution and be sure that you back it up first. Registry information is provided here as a reference for use by only highly skilled directory service administrators. Do not directly edit the registry unless, as in this case, no Group Policy setting or other Windows tools can accomplish the task. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. Storage of incorrect values can result in unrecoverable errors in the system. |

On one domain controller, you can set the **Cached Membership Site Stickiness (minutes)** registry entry to 0 and then initiate a cache refresh by using the operational attribute method on that domain controller, as described in "Methods of Refreshing the Cached Memberships" earlier in this subject. The 0 value in the setting causes the cache to be purged—values in all three attributes (**msDS-Cached-Membership**, **msDS-Cached-Membership-Time-Stamp**, and **msDS-Site-Affinity**) are cleared. After the site stickiness attribute deletion has replicated within the site, you can then initiate cache refresh on the other domain controllers in the site. Remember to return the value in **Membership Site Stickiness (minutes)** to its default value of 180.

**Diagnostic Logging Levels and Events**

Diagnostic logging for Universal Group Membership Caching can be set in the registry entry **20 Group Caching** under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics**

Data type: REG_DWORD

Value range: 0-5

Default value: 0

Significant events are reported at logging level 2 with many additional events reported at logging level 5. For troubleshooting, set the logging level to 5.

**Sample Events at Logging Level 0**

Event ID 1667: The group membership cache refresh task detected that the following site in which a global catalog was found is not one of the cheapest sites, as indicated by the published site link information.

Event ID 1668: The group membership cache refresh task did not locate a global catalog server in the preferred site, but was able to find a global catalog server in the following available site. Preferred site: <site name> Available site: <site name>

Event ID 1669: The group membership cache refresh task has reached the maximum number of users for the local domain controller.

Event ID 1670: The group membership cache refresh task is behind schedule. Consider forcing a group membership cache update.

**Sample Events at Logging Level 2**

Event ID 1776 internal event: The group membership cache task is starting.

Event ID 1777 internal event: The group membership cache task has finished. The completion status was 0, and the exit Internal ID was ######.

Event ID 1779 internal event: The Global Catalog Domain Controller <dcname> in site <site name>, domain <domain name> will be used to update the group memberships.

Event ID 1781 internal event: By examining the published connectivity information, the group membership cache task has determined site <site name> is a site with a low network cost to contact. The task will schedule itself based on the schedule of network connectivity to this site.

Event ID 1782 internal event: By examining the published connectivity information, the group membership cache task cannot find an efficient site to obtain group membership information. The task will run using the global catalog server that is closest, as determined by the Net Logon locator, and will schedule itself based on a fixed period.

Event ID 1842 internal event: The following site link will be used to schedule the group membership cache refresh task. Site link: <distinguished name of site link>

**Sample Events at Logging Level 5**
Event ID 1778 internal event: The group membership cache task will run again in xx minutes.

Event ID 1784 internal event: The group membership cache task determined that site <distinguished name of site> does not have a global catalog server.

**How the Cache is Populated at First Logon**
By default, the caching attributes on the user and computer objects are not populated. The following diagram shows how the domain controller builds the list of SIDs to be cached and where in the process the caching attributes are populated during the user's first logon in the site. This example assumes that the user is in a site that has Universal Group Membership Caching enabled, the domain of the client workstation is the same as the domain of the user, and the domain has a functional level that allows universal groups.

**Universal Group Membership Caching Process at First Logon**




The following events occur at each step in the preceding diagram:

1. A user logs on in a site where Universal Group Membership Caching is enabled. The user is authenticated by the domain controller as being the requesting user.

2. The domain controller checks the values of the three caching attributes of the user object.

3. Finding that the attributes are not populated, the domain controller checks its local directory and retrieves the SID of the user (including SID history, if available) and the SIDs of all global groups to which the user belongs.

4. The domain controller sends this list of SIDs to the global catalog server. The global catalog server checks the universal group memberships of the user and all global groups in the list. The global catalog server returns the list of combined universal group and global group SIDs to the domain controller.

5. The user's cache attributes are populated as follows:

    a. The combined list of global group and universal group SIDs is recorded in the **msDS-Cached-Membership** attribute.

    b. The time is recorded in the **msDS-Cached-Membership-Time-Stamp** attribute (this time indicates the last time the cache was updated; on the first logon, it also happens to be the time the user logged on).

    c. If **SamNoGcLogonEnforceNTLMCheck** or **SamNoGcLogonEnforceKerberosIpCheck**, or both, are enabled on the domain controller, the **msDS-Site-Affinity** attribute is ignored.

    d. If the GUID for the local site exists in the **msDS-Site-Affinity** attribute and the settings in step c. are not enabled, the timestamp value in **msDS-Site-Affinity** is evaluated as follows: If the value indicates an age that is less than half the value in **Cached Membership Site Stickiness (minutes)**, the logon proceeds. If the timestamp indicates

an age that is greater than half the value in **Cached Membership Site Stickiness (minutes)**, or if the attribute is not populated, the site GUID and time are written to the **msDS-Site-Affinity** attribute, and the logon proceeds.

6. The domain controller checks its local directory for any domain local groups to which the user belongs and adds domain local group SIDs to its list of global group and universal group SIDs.

| **Note** |
| --- |
| The process for accomplishing Step 6 differs depending on whether the domain of the client computer is the same as the domain of the user and, if not, whether the client computer is joined to a domain that has a mixed domain mode or functional level, or a native domain mode or functional level. For more information about how SIDs are retrieved and added to access tokens, see "Access Tokens Technical Reference." |

7. The domain controller sends the entire list of SIDs to the client computer, where the LSA retrieves SIDs of the user's built-in group memberships and constructs the user's access token.

| **Note** |
| --- |
| Global catalog servers in a site where caching is enabled do not populate the **msDS-Cached-Membership** and **msDS-Cached-Membership-Time-Stamp** attributes of users they authenticate. Because global catalog servers are already aware of universal group memberships throughout the forest and global group memberships for the domain, there is no need for them to use these attributes. |

**How the Cache is Used for Subsequent Logons**

When Universal Group Membership Caching is enabled in the site, the following sequence occurs during account logon:

1. The account is authenticated by the contacted domain controller.

2. The domain controller checks for the presence of values in the caching attributes of the respective user or computer object. If the attribute values are present, the domain controller updates the values as follows:

   a. If the value in the **msDS-Cached-Membership-Time-Stamp** attribute indicates an age that is less than the staleness interval (**Cached Membership Staleness (minutes)**, default seven days), the domain controller reads the group SIDs from the **msDS-Cached-Membership** attribute and the logon proceeds.

   b. If the value in **msDS-Cached-Membership-Time-Stamp** indicates an age of greater than the staleness interval, the domain controller contacts a global catalog server to request the universal group membership. If a global catalog server cannot be contacted, the logon is denied.

   c. If the value in the timestamp in **msDS-Site-Affinity** is equal to or older than 50 percent of the site stickiness setting, the timestamp is updated with the current time.

3. The domain controller returns the group SIDs from the cache plus any domain local group SIDs to the client computer and the logon proceeds.

| **Note** |
| --- |
| At no time during a successful logon does the local domain controller check with a global catalog server to see if the account's group membership has changed. Changes to an account's group membership are not reflected in the account's access token until the local domain controller performs a cache refresh. The default amount of time between cache refreshes is eight hours. This interval could result in an inconsistent view of group membership if the account was authenticated by a domain controller in a different site. This discrepancy might also confuse administrators who are unfamiliar with how universal group membership caching works. |

**How the Cache is Refreshed**

The cache refresh process occurs automatically on every domain controller that is running Windows Server 2003 or later and has received replication of the **msDS-Site-Affinity** attribute update for a user or computer object or has already cached group

memberships. The refresh operation occurs differently depending on whether a site is selected for the preferred refresh site.

**Cache Refresh Process When a Preferred Refresh Site is Not Selected**

When the refresh interval expires, the domain controller proceeds as follows:

1. Lists all the site links that connect the domain controller's site to a site that hosts a global catalog server in increasing order of cost values on the site link objects.

2. Selects the lowest-cost site link and schedules the refresh by using the site link schedule. If no site link schedule is set, then replication is always available. Depending on the schedule, the refresh proceeds as follows:

   - If the schedule currently allows replication, the domain controller begins the refresh.

   - If the schedule does not currently allow replication, the domain controller schedules the refresh to begin when the schedule allows replication.

     | **Note** |
     | --- |
     | When the refresh is postponed according to the site link schedule, a random stagger in the range of 0-15 minutes is added to the computed start time. Schedule staggering has the effect of ensuring that domain controllers begin their refresh at slightly different times, thereby improving load balancing on the global catalog server. |

   - When the schedule allows replication, begins the refresh by locating and binding to a global catalog server in the next closest site.

   - Removes accounts that have a populated cache but no site affinity. Cached entries that do not include a populated **msDS-Site-Affinity** value are purged at this time. A maximum of 64 entries are removed at a time. If more entries need to be removed, they are removed during subsequent refreshes.

   - Removes any account whose site affinity matches the local site, but whose site affinity time period has expired. In this case, the values in the three cache attributes are deleted and this account no longer has a group membership cache on the domain controller.

   - Builds a list of accounts by querying the global catalog for all accounts that have GUIDs in their **msDS-Site-Affinity** attribute that match the GUID of the domain controller's site.

   - Updates cache attributes of the accounts in the list by querying the global catalog for each account's group membership, as follows:

     - Update the **msDS-Cached-Membership** attribute with the account's group membership SIDs.

     - Update the **msDS-Cached-Membership-Time-Stamp** attribute with the time of refresh.

   - Repeats the process for each account until all accounts are updated or until the refresh limit of 500 accounts is reached. If the refresh limit is reached, the domain controller logs event ID 1669 in the Directory Service event log, and the refresh stops.

   - Checks to ensure that the refresh task has not fallen behind in terms of the maximum period allowed for an account's cached membership list to be valid for logons. This step is accomplished by locating the record with the oldest **msDS-Cached-Membership-Time-Stamp** value and comparing the timestamp value to the staleness interval (seven days by default). If the entry is more than seven days old, the domain controller logs event ID 1670, indicating that the refresh task has fallen behind.

     | **Note** |
     | --- |

When the domain controller encounters the refresh limit, it stops updating cache entries. Because the order in which the updates occur cannot be predicted, there is no way to ensure that the caches of the most recent accounts are updated. The staleness check in step 9 checks all cached entries, even those excluded due to exceeding the refresh limit. After about one week, the non-updated cache entries will become stale and cause the falling behind error to be reported in the event log.

**Note**

- When the domain controller encounters the refresh limit, it stops updating cache entries. Because the order in which the updates occur cannot be predicted, there is no way to ensure that the caches of the most recent accounts are updated. The staleness check in step 9 checks all cached entries, even those excluded due to exceeding the refresh limit. After about one week, the non-updated cache entries will become stale and cause the falling behind error to be reported in the event log.

**Cache Refresh Process When A Preferred Refresh Site is Selected**

When a site is selected to always be used for refreshing the group membership cache, the domain controller does not need to order the site links according to cost, but simply contacts a global catalog server in the specified site. However, if no global catalog server is available at the time the refresh is attempted, the domain controller logs event ID 1782, indicating that a domain controller could not be found in the preferred site, and uses the site link cost to locate a global catalog server in the next closest site.

**Inconsistent Access to Domain-Based Objects on Global Catalog Servers**

When specifying Read or List permissions for domain data that is also replicated to the global catalog, use global groups rather than domain local groups because the access token that is created for the user by the global catalog server does not necessarily contain information about domain local groups to which the user belongs.

When a user connects to a global catalog server, an access token is created for the user that is used in subsequent access decisions on the server. If the user is a member of a domain other than the domain of the global catalog server, the global catalog server contacts a domain controller in the user's domain to authenticate the user and retrieve authorization data. The domain controller returns information about the user, including the SIDs of global groups in the user's domain to which the user belongs. The domain controller from the user's domain *does not return domain local group SIDs to the global catalog server*.

Universal group membership is retrieved from the global catalog, and the global catalog server looks to its own domain (which is not necessarily the domain of the user) for any domain local groups to which the user belongs. Thus the access token for the user on the global catalog server contains the global groups and universal groups to which the user belongs, as well as any domain local groups to which the user belongs in the domain of the global catalog server.

The effect of a missing domain local group SID in the user's access token is that the user's access to global catalog data is unpredictable. For example, if access to the **homePhone** attribute of a user object is restricted by a domain local group in the user's domain, and the user is a member of that group, the user is able to view that attribute in the global catalog when both of the following conditions are true:

- The **homePhone** attribute is replicated to the global catalog.
- The global catalog server to which the user connects does not host a writable copy of the user's domain.

Similarly, if the user is a member of a domain local group in a domain other than the domain hosted by the global catalog server, and that group is granted read access to the **homePhone** attribute, the user cannot view that attribute in the global catalog.

**Global Catalog Searches**

The location of an object in AD DS is provided by the distinguished name of the object, which includes the full path to a replica of the object, culminating in the directory partition that holds the object. However, the user or application does not always have the distinguished name of the target object, or even the domain of the object. To locate objects without knowing the domain, the most commonly used attributes of the object are replicated to the global catalog. By using these object attributes and directing the search to the global catalog, requesters can find objects of interest without having to know their directory location. For example, to locate a printer, you can search according to the building of the printer. To locate a person, you can provide the name of the person. To locate all people who are managed by someone, you can provide the manager's name.

**LDAP Search Ports**
AD DS uses the Lightweight Directory Access Protocol (LDAP) as its access protocol. LDAP search requests can be sent and received by AD DS on port 389 (the default LDAP access port) and port 3268 (the default global catalog port). LDAP traffic that uses the Secure Sockets Layer (SSL) authentication protocol accesses ports 636 and 3269, respectively. In this discussion, search behavior that applies to ports 389 and 3268 also apply to the respective behavior of LDAP queries over ports 636 and 3269.

When a search request is sent to port 389, the search is conducted on a single domain directory partition. If the object is not found in that domain or the schema or configuration directory partitions, the domain controller refers the request to a domain controller in the domain that is indicated in the distinguished name of the object.

When a search request is sent to port 3268, the search includes all directory partitions in the forest — that is, the search is processed by a global catalog server. If the request specifies attributes that are part of the PAS, the global catalog can return results for objects in any domain without generating a referral to a domain controller in a different domain. Only global catalog servers receive LDAP requests through port 3268. Certain LDAP client applications are programmed to use port 3268. Even if the data that satisfies a search request is available on a regular domain controller, if the application launching the search uses port 3268, the search necessarily goes to a global catalog server.

**Search Criteria That Target the Global Catalog**
Searches are directed to a global catalog server under the following conditions:

- You specify port 3268 or 3269 in an LDAP search tool.
- You select **Entire Directory** in a search-scope list in an AD DS snap-in or search tool, such as Active Directory Users and Computers or the **Search** command on the **Start** menu.
- You write the distinguished name as an attribute value, where the distinguished name represents a nonlocal object. For example, if you are adding a member to a group and the member that you are adding is from a different domain, a global catalog server verifies that the user object represented by the distinguished name exists and obtains its GUID.

**Characteristics of a Global Catalog Search**
The following characteristics differentiate a global catalog search from a standard LDAP search:

- Global catalog queries are directed to port 3268, which explicitly indicates that global catalog semantics are required. By default, ordinary LDAP searches are received through port 389. If you bind to port 389, even if you bind to a global catalog server, your search includes a single domain directory partition. If you bind to port 3268, your search includes all directory partitions in the forest. If the server you attempt to bind to over port 3268 is not a global catalog server, the server refuses the bind.
- Global catalog searches can specify a non-instantiated search base, indicated as "com" or " " (blank search base).
- Global catalog searches cross directory partition boundaries. The extent of the standard LDAP search is the directory partition.
- Global catalog searches do not return subordinate referrals. If you use port 3268 to request an attribute that is not in the global catalog, you do not receive a referral to it. Subordinate referrals are an LDAP response; when you query over port 3268, you receive global catalog responses, which are based solely on the contents of the global catalog. If you query the same server by using port 389, you receive referrals for objects that are in the forest but whose attributes are not referenced in the global catalog.

| **Note** |
| --- |
| A referral to a directory that is external to AD DS can be returned by the global catalog if a base-level search for an external directory is submitted and if the distinguished name of the external directory uses the domain component (dc=) naming attribute. This referral is returned according to the ability of AD DS to construct a Domain Name System (DNS) name from the domain components of the distinguished name and is not based on the presence of any cross-reference object. The same referral is returned by using the LDAP port; it is not specific to the global catalog. |

Because the **member** attribute is not replicated to the global catalog for all group types, and because the **memberOf** attribute derives its value by referencing the **member** attribute (called back links and forward links, respectively), the results of searches

for members of groups and groups of which a member belongs can vary, depending on whether you search the global catalog (port 3268) or the domain (port 389), the kind of groups that the user belongs to (global groups or domain local groups), and whether the user belongs to universal groups outside the local domain.

For more information about global catalog searches and the implications of searching on back links and forward links, see "How Active Directory Searches Work."

**The Infrastructure Master and Phantom Records**

An attribute that has a distinguished name as a value references (points to) the named object. When the referenced object does not actually exist in the local directory database because it is in a different domain, a placeholder record called a phantom is created in that database as the object reference. Because there is a reference to it, the referenced object must exist in some form, either as the full object (if the domain controller stores the respective domain directory partition) or as an object reference (when the domain controller does not store that domain).

The infrastructure master is a single domain controller in each domain that tracks name changes of referenced objects and updates the references on the referencing object. When a referenced object is moved to a different domain (which effectively renames the object), the infrastructure master updates the distinguished name of the phantom. The infrastructure master finds phantom records by using a database index that is created only on domain controllers that hold the infrastructure operations master role. When the reference count of the phantom falls to zero (no objects are referencing the object that the phantom represents), garbage collection on each domain controller removes the phantom.

Because objects can reference objects in different domains, the infrastructure operations master role is not compatible with global catalog server status if more than one domain is in the forest. If a global catalog server holds the infrastructure operations master role, phantom records are never created because the referenced object is always located in the directory database on the global catalog server.

For more information about the infrastructure operations master role, see "How Operations Masters Work."

**Exchange Address Book Lookups**

The Exchange Server directory service is integrated with AD DS. When mail users want to find a person within their organization, they usually search the global address book (GAL), which is an aggregation of all messaging recipients in the enterprise, including mailbox-enabled users, mail-enabled users, groups, and contacts. The GAL is a virtual linked list of pointers to the mail recipient objects that comprise it. Mail recipients can be user accounts (both enabled and disabled accounts), contacts, distribution lists, security groups, and folders. The GAL is automatically populated by a service on the Exchange Server, and user's can create customized address lists. Exchange Server uses the global catalog to generate the GAL.

Every Outlook client is configured with the name of an Exchange server. Exchange servers use AD DS and DNS to locate a global catalog server. When an Outlook client user opens the Address Book, or when a user composes a message and types a name or an address in the To: field, the Outlook client uses the global catalog server that is specified by its Exchange server to search the contents of the GAL or other address lists.

For more information about how Outlook clients locate address information in the global catalog, see "How Active Directory Searches Work."

**Global Catalog Server Creation and Advertisement**

You can designate a domain controller as a global catalog server by simply selecting the **Global catalog** check box in the properties of the NTDS Settings object, located beneath each server object in Active Directory Sites and Services. On servers that run Windows Server 2008 R2 or Windows Server 2008, you can also designate a domain controller as a global catalog server during AD DS installation. However, for search clients and other domain controllers in the forest to locate it, the global catalog server must register itself in DNS.

The Net Logon service on a domain controller registers service (SRV) resource records in DNS that identify the domain controller so that it can be located. In the case of a global catalog server, additional SRV resource records are registered to identify the server as a global catalog server. These SRV resource records contain the server name, the forest name, and the site name for the global catalog server. DNS queries for these records return all global catalog servers in the requested site. When a client requests a global catalog server by launching an LDAP search over port 3268, the domain controller Locator on the client queries DNS for a domain controller that hosts the global catalog. For more information about how domain controllers and global catalog servers are located, see "How DNS Support for Active Directory Works."

By default, before a domain controller advertises itself as a global catalog server in DNS, the global catalog contents must be replicated to the server. This process involves replication of a partial, read-only replica of every domain in the forest except for the domain for which the new global catalog server is authoritative. The duration of this process depends on how many domains the forest contains, the size of the domains, and the relative locations of source and destination domain controllers. If multiple domains are in the forest and if source domain controllers are located only in distant sites, the process takes longer than if all domains are in the same site or in only a few sites. When replication must occur between sites to create the global

catalog, replication occurs according to the site link schedule.

When creating a new global catalog server, the process can be delayed by several conditions, including the following:

- The KCC could not reach a source domain controller from which to replicate a directory partition.
- Replication cannot begin until the scheduled time.
- Replication of the directory partition is in progress but has not yet completed. This delay might occur if the directory partition is very large. In addition, the replication priority queue prioritizes addition of new directory partitions at a lower priority than incremental replication of existing partitions.
- The source domain controller for a directory partition has gone offline or is unavailable due to network problems.

These conditions are logged in the Directory Service log when the logging level is set to 0 (the default setting) in the **Global Catalog** entry under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics**. If you want to receive more information, you can increase the logging level by editing this entry.

| **Warning** |
| --- |
| If you must edit the registry, use extreme caution and be sure that you back it up first. Registry information is provided here as a reference for use by only highly skilled directory service administrators. Do not directly edit the registry unless, as in this case, no Group Policy setting or other Windows tools can accomplish the task. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. Storage of incorrect values can result in unrecoverable errors in the system. |

**Requirements for Global Catalog Readiness**

By default, a global catalog server is not considered "ready" (the server advertises itself in DNS as a global catalog server) until all read-only directory partitions have been fully replicated to the new global catalog server. The **Global Catalog Partition Occupancy** registry entry under **HKEY_Local_Machine\System\CurrentControlSet\Services\NTDS\Parameters** determines the requirements for how many read-only directory partitions must be present on a domain controller for it to be considered a global catalog server, from no partitions (0) to all partitions (6).

For domain controllers that run Windows Server 2003 or later, the default occupancy value requires that all read-only directory partitions be replicated to the global catalog server before the Net Logon service registers SRV resource records in DNS. For most conditions, this default provides the best option for ensuring that a global catalog server provides a consistent view of the directory. In less common circumstances, however, it might be useful to make the global catalog server available with an incomplete set of partial domain directory partitions—for example, when delay of replication of a domain that is not required by users is jeopardizing their ability to log on.

The **Global Catalog Partition Occupancy** entry can have the values shown in the following table.

**Global Catalog Partition Occupancy Level Values**

| Value | Description |
| --- | --- |
| 0 | No occupancy requirement. Removing the occupancy level requirement might be useful in a scenario where domain controllers are being staged for deployment but are not yet in production. |
| 1 | At least one read-only directory partition in the site has been added by the KCC. This level, as well as level 3 and level 5, provide the ability to distinguish between a source for the directory partition being reachable (at least one object has been returned) and the entire directory partition having been replicated (as in levels 2, 4, and 6).<br><br>When the KCC can reach the first object, it can create a replica link, which is the agreement between the source and destination domain controllers to replicate to the destination. If the KCC cannot reach a source, the KCC logs event ID 1558 in the Directory Service log, which indicates the distinguished name of the directory partition that has not been fully synchronized. In this case, the KCC continues to try to replicate the partition each time it runs (every 15 minutes by default).<br><br>When the KCC succeeds in creating the replica link, it passes responsibility for retrying and completing the synchronization to the replication engine. The KCC then stops logging events, after which the replication status can be checked by using the **repadmin/showrepl** command. |

| | |
|---|---|
| 2 | At least one read-only directory partition in the site has been fully synchronized. |
| 3 | All read-only directory partitions in the site have been added by the KCC (at least one has been fully synchronized). In this case, the KCC has been able to contact one source for every directory partition in the site. This level is useful when you want to advertise a global catalog server as soon as possible with a high likelihood of success. |
| 4 | All read-only directory partitions in the site have been fully synchronized. With this setting, if a source for any directory partition is not available, DNS registrations cannot occur. On domain controllers that are running Windows 2000 Server with Service Pack 1 (SP1) or Windows 2000 Server with Service Pack 2 (SP2), this occupancy level is the default requirement before the global catalog server is advertised in DNS. |
| 5 | All read-only directory partitions in the forest have been added by the KCC (at least one has been fully synchronized). |
| 6 | All read-only directory partitions in the forest have been fully synchronized. On domain controllers that are running Windows Server 2003 or Windows 2000 Server with SP3 or later, this occupancy level is the default requirement before the global catalog server is advertised in DNS. This setting ensures the highest level of consistency. |

Event ID 1578 reports the level that is required and the level that the domain controller has achieved.

**Advertising a Global Catalog Server Prior to Full Synchronization**

By default, a domain controller checks every 30 minutes to see whether it has received all of the read-only directory partitions that are required to be present before the server advertises itself in DNS as a global catalog server. Event ID 1110 indicates that the promotion is being delayed because the required directory partitions have not all been synchronized.

This delay is controlled by the **Global Catalog Delay Advertisement (sec)** registry entry under **HKEY_Local_Machine\System\CurrentControlSet\Services\NTDS\Parameters**. If you set a value for **Global Catalog Delay Advertisement (sec)**, it overrides the requirements set in **Global Catalog Partition Occupancy** and allows global catalog advertisement without requiring full synchronization of all read-only directory partitions.

When conditions preclude the successful synchronization of the new global catalog server, you can force advertisement of the global catalog server and then remove the global catalog from the server. Until the global catalog server is successfully advertised, you are unable to remove it.

**Replication Process for Global Catalog Creation**

When you designate a domain controller to be a global catalog server, the Knowledge Consistency Checker (KCC) on the domain controller runs immediately and updates the replication topology. When the KCC runs, it checks to see whether the Global Catalog option is selected for any domain controllers, and creates the replication topology accordingly. The KCC configures the newly selected global catalog server to be the destination server for a read-only replica of every domain directory partition in the forest except for the writable domain directory partition that the server already holds. The KCC on the global catalog server must be able to reach a server that will be the source of each read-only directory partition.

When the KCC locates an available source domain controller, it creates an inbound connection on the new global catalog server and replication of that read-only partition takes place. If the source is within the site, replication begins immediately. If the source is in a different site, replication begins when it is next scheduled. Replication of all objects in the partial directory partition must complete successfully before the directory partition is considered to be present on the global catalog server.

**Successful Completion of Global Catalog Creation**

When all directory partitions are present, the domain controller sets its rootDSE **isGlobalCatalogReady** attribute to TRUE and the Net Logon service on the domain controller registers SRV resource records that specifically advertise the global catalog server in DNS. At this point, the global catalog is considered to be available, and event ID 1119 is logged in the Directory Service log.

**Global Catalog Replication**

Although read-only directory partitions on global catalog servers cannot be updated directly by administrative LDAP writes, they are updated through AD DS replication when changes occur to their respective writable replicas.

The following diagram represents the AD DS database on a global catalog server in the corp.contoso.com forest root domain. A global catalog server has a single directory database. However, to represent the different logical directory partitions in the forest, the diagram shows database divided into segments. The top three segments represent directory partitions that are writable replicas for the domain controller (the domain, configuration, and schema directory partitions). The bottom three segments represent directory partitions that are read-only replicas of the other domains in the forest.

**Writable and read-only replicas in the AD DS database on a global catalog server**




The source domain controller for replication of a given directory partition to a global catalog server can be either a non-global catalog domain controller or another global catalog server. In the following diagram, each directory partition on the global catalog server is being updated by a non-global catalog domain controller. The writable replicas on the global catalog server are updated by a domain controller that is authoritative for the same domain, Corp.contoso.com. The replication for the Corp.contoso.com domain and the configuration and schema directory partitions is two-way because the replicas are all writable.

Each of the read-only replicas is updated by a source domain controller that is authoritative for the respective directory partition. The replication is one-way because read-only replicas never update writable replicas.

**Direction of directory partition replica updates between a global catalog server and other domain controllers**




**Replication Between Global Catalog Servers**

As is true for all domain controllers, a global catalog server uses a single topology to replicate the schema and configuration directory partitions, and it uses a separate topology, if needed, for each domain directory partition. However, when a two-way connection exists between the servers, either for replication of the schema and configuration directory partitions or for replication in opposite directions of the two writable domain directory partitions, all replicas on each global catalog server use the same connection to update their common replicas when changes are available.

The diagram below shows the directions of replication between directory partitions on two global catalog servers that are in different sites and are authoritative for different domains. The writable replicas of soam.corp.contoso.com and corp.contoso.com update the respective read-only replicas in one direction only (a writable replica is never updated by a read-only replica). Because neither domain controller is authoritative for the noam.corp.contoso.com and eur.corp.contoso.com domain replicas, the global catalog servers can be sources for replication of these partial read-only replicas. This replication is shown as two-way because a two-way connection already exists and these replicas are each capable of updating the other.

**Direction of directory partition replica updates between two global catalog servers in different domains**




In the preceding diagram, the read-only replicas can also be updated from other domain controllers. Unless the forest functional level is Windows 2000 Server, the intersite KCC algorithm avoids creating redundant connection objects by implementing one-way replication where possible. For example, if the schema and configuration writable replicas and the Corp and Eur read-only domain replicas on GC1 are all updated by a domain controller other than GC2, replication of the Corp and Eur read-only replicas from GC1 to GC2 occurs in one direction if it occurs. In this case, GC1 might not generate a connection object for replication from GC2.

**Global Catalog Replication of Additions to the Partial Attribute Set**

Each global catalog server in an AD DS forest hosts a copy of every existing object in that forest. For the objects of its own domain, a global catalog server has information related to all attributes that are associated with those objects. For the objects in domains other than its own, a global catalog server has only information that is related to the set of attributes that are marked in the AD DS schema to be included in the partial attribute set (PAS). As described earlier, the PAS is defined by

Microsoft as those attributes that are most likely to be used for searches. These attributes are replicated to every global catalog server in an AD DS forest.

If you want to add an attribute to the PAS, you can mark the attribute by using the Active Directory Schema snap-in to edit the **isMemberOfPartialAttributeSet** value on the respective **attributeSchema** object. You mark the attribute by placing a checkmark next to **isMemberOfPartialAttributeSet**. If the **isMemberOfPartialAttributeSet** value is checked (set to TRUE), the attribute is replicated to the global catalog. If the value is not checked (set to FALSE), the attribute is not replicated to the global catalog.

When the PAS is updated on the domain controller that has the schema operations master role, the writable schema directory partition replicates normally to all domain controllers. When a global catalog server attempts to update its read-only directory partition from a source replication partner whose schema has been udpated with the PAS change, the destination global catalog server receives the information that the PAS has been updated.

Depending on the version of Windows that is running on the replication partners, a global catalog server responds to an update to the PAS either by initiating replication of all attributes in the read-only directory partitions in the global catalog (full synchronization), or only the attributes that were added to the PAS (updates only), as follows:

- **Updates only—both replication partners are running Windows Server 2003 or later:** Only the added attributes are replicated and there is no replication impact. In this case, when the **isMemberOfPartialAttributeSet** value of a new **attributeSchema** object in the schema directory partition is set to TRUE and this change is either made on a global catalog server or replicated in by a global catalog server, the global catalog server records unique NTDS Replication informational events in the Directory Service log for each read-only directory partition that it hosts, as follows:

    - Event ID 1704, stating that in response to addition of one or more attributes to the PAS, the global catalog server has initiated replication of the PAS for the specified directory partition from the specified domain controller.

    - Event ID 1702, stating that synchronization of the PAS has completed successfully.

- **Full synchronization—both replication partners are running Windows 2000 Server:** The global catalog server initiates replication of all PAS attributes of all objects in all read-only domain directory partitions. In this case, all replication partner metadata and up-to-dateness metadata for the directory partition is discarded and the global catalog server builds a new list of replication partners and up-to-dateness information relative to those partners. Building its partner information anew, it requests replication of all attributes, filtering out any attribute that is not marked for inclusion in the PAS. If the attributes can be replicated over an RPC connection, the domain controller attempts a full replication over the RPC connection before it uses a configured SMTP connection. If full replication is completed, the up-to-dateness vector that is created ensures that later replication requests on other connections do not include data that has been received during the initial full replication.

  Full replication of all attributes in the PAS to bring all global catalog servers up-to-date causes increased network traffic while it is in progress and can take from several minutes to hours, depending on the size of the directory. Although interruption of service does not occur, this replication causes higher bandwidth consumption than is required for usual day-to-day replication. The resulting bandwidth consumption for each global catalog server is equivalent to that caused by assigning the global catalog role to a domain controller.

  In this case, when the **isMemberOfPartialAttributeSet** value of an additional **attributeSchema** object in the schema directory partition is set to TRUE and this change is either made on a global catalog server or replicated in by a global catalog server, a unique NTDS Replication informational event ID 1575 is recorded in the Directory Service log for each read-only directory partition that is hosted by the global catalog server. Event ID 1575 states that one or more new attributes has been added to the PAS for the specified directory partition and that full synchronization of all objects and all attributes in the specified directory partition will be performed from the source domain controller on the next replication cycle.

- **Full synchronization—a global catalog server that is running Windows Server 2003 or later sources replication of a partial, read-only directory partition from a global catalog server or domain controller that is running Windows 2000 Server:** The global catalog server that is running Windows Server 2003 or later reverts to Windows 2000 Server behavior and, as described above, initiates full replication of all objects and their attributes in all read-only directory partitions. For each read-only directory partition that it hosts, the global catalog server that is running Windows Server 2003 or alter records unique NTDS Replication informational events in the Directory Service log, as follows:

- Event ID 1575, stating that full synchronization of the PAS will be initiated on the next replication cycle.
- Event ID 1703, stating that synchronization of the PAS has been initiated.
- Event ID 1702, stating that synchronization of the PAS has completed successfully.

| **Note** |
| --- |
| The AD DS schema in Windows Server 2003 and later contains attributes that are marked for inclusion in the PAS as soon as the forest functional level is raised to Windows Server 2003 or higher. Replication of the additions to the PAS to global catalog servers is triggered by raising the forest functional level to Windows Server 2003 or higher. Therefore, upgrading the schema has no impact on Windows 2000–based global catalog servers because the global catalog is updated only when all domain controllers are running Windows Server 2003 or later (the requirement for raising the forest functional level to at least Windows Server 2003). For more information about functional levels, see "How Active Directory Functional Levels Work." |

For more information about replication metadata, including up-to-dateness vector, see "How the Active Directory Replication Model Works."

**Global Catalog Response to Deletions from the Partial Attribute Set**
Removing an attribute from the PAS does not involve replication of a deletion, but is handled locally. If you set the **isMemberOfPartialAttributeSet** value to FALSE in the schema, the attribute is removed from the directory of each global catalog server immediately after receiving the schema update. This behavior is the same on global catalog servers running any version of Windows Server.

**Removing the Global Catalog from a Domain Controller**
When you add the global catalog to a domain controller, a partial replica of each domain directory partition, other than the domain for which the domain controller is authoritative, is copied to the domain controller as described in Replication Process for Global Catalog Creation. This replication occurs immediately within a site or at the next scheduled replication between sites. However, when you remove the global catalog from a domain controller, the KCC begins removing the read-only replicas one at a time by means of an asynchronous process that removes objects gradually over time until no read-only objects remain.

The Windows 2000 KCC removes approximately 500 objects per attempt. Each time the KCC runs (every 15 minutes by default), it attempts the removal of the read-only replica until there are no remaining objects. At an estimated 2000 objects per hour, complete removal of the global catalog from the domain controller can take from several hours to days, depending on the size of the directory.

The KCC in Windows Server 2003 or later initially instructs the directory to remove each replica, but once the instruction is received, the directory itself keeps track of the progress of replica removal and reschedules the work accordingly. Rather than removing a fixed number of objects per removal attempt, AD DS continues removing objects until either the replica is gone or a higher priority replication operation is in the queue. Read-only replica removal receives the lowest possible priority, meaning that any other replication work will interrupt it. Thus, removal work is pre-empted for the other work and then resumed later. On domain controllers running Windows Server 2003 or later, event log messages record when removal of a partial replica is starting, being resumed, and finishing.

KCC events that are logged in the Directory Service log during global catalog removal include:

- Event ID 1744: Local domain controller is no longer a global catalog server.
- Event ID 1659: Removal of a directory partition from the local database has been resumed, including the approximate number of objects remaining to be removed and number of link values of attributes remaining to be removed.
- Event ID 1660: A specified directory partition has been removed from the local domain controller.

For a normally-to-lightly loaded system, read-only replica removal occurs as fast as possible in the background. On a server that is very busy with replication (for example, a hub domain controller), estimating the time required for global catalog removal is difficult.

As soon as you set the domain controller to not be a global catalog server by clearing the **Global Catalog** check box on the NTDS Settings object properties page, Net Logon deregisters the SRV resource records that specifically advertise the global

catalog server in DNS. Therefore, although the read-only domain replicas might take hours or days to remove, the domain controller immediately stops advertising itself as a global catalog server in DNS and immediately stops accepting LDAP requests over port 3268.

| **Note** |
| --- |
| On a domain controller running Windows 2000 Server, if you check the **Global Catalog** box again before a partial replica has been completely removed, the KCC begins the process of replicating in the read-only replicas as follows: If a given replica is completely gone, the KCC adds the replica back. If the replica is still in the process of being removed, the KCC does not add it again until the initial removal is complete. Thus, during the removal period, there can conceivably be a mix of new replicas from the most recent global catalog instance, and some old replicas in the process of being removed from the previous instance. This condition is temporary and of no consequence to users because the domain controller is not being advertised as a global catalog server. |

**Network Ports Used by the Global Catalog**

The following ports are used by global catalog servers:

**Port Assignments for Global Catalog Servers**

| Service Name | UDP | TCP |
| --- | --- | --- |
| LDAP | | 3268 (global catalog) |
| LDAP | | 3269 (global catalog Secure Sockets Layer [SSL]) |
| LDAP | 389 | 389 |
| LDAP | | 636 (SSL) |
| RPC/REPL | | 135 (endpoint mapper) |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- RFC 2247 (http://go.microsoft.com/fwlink/?LinkId=192755) in the Internet Engineering Task Force (IETF) RFC Database
- Domain Controller Roles
- How Active Directory Searches Work
- Active Directory Collection
- Windows Security Collection
- Active Directory Replication Model Technical Reference
- Active Directory Replication Topology Technical Reference
- Domains and Forests Technical Reference
- Interactive Logon Technical Reference
- Access Tokens Technical Reference

# Global Catalog Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Global Catalog Tools
- Global Catalog Registry Entries
- Global Catalog Group Policy Settings
- Network Ports Used by Global Catalog Servers
- Related Information


**Global Catalog Tools**
Tools that are associated with a global catalog server are the same tools that you use to manage any domain controller.

The following tools have commands that are specific to managing global catalog servers.

**Adsiedit.msc: ADSI Edit**

**Category**
A Microsoft Management Console (MMC) snap-in that is available in Windows Support Tools in Windows Server® 2003 and Microsoft Windows® 2000 Server. It is built into Windows Server 2008 R2 and Windows Server 2008 and available if you have the Active Directory® Domain Services (AD DS) or the Active Directory Lightweight Directory Services (AD LDS) server role installed. It is also available if you install the Active Directory Domain Services Tools that are part of the Remote Server Administration Tools (RSAT). For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

| **Note** |
| --- |
| In Windows Server 2003 and Windows 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008, the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

**Version compatibility**

| Can Be Run From | Can Be Run Against |
| --- | --- |

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows 7 with Remote Server Administration Tools (RSAT) installed.<br>• Windows XP Professional<br>• Windows Vista with Remote Server Administration Tools (RSAT) installed. | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

ADSI Edit is an MMC snap-in that you can use to view and modify attributes of directory objects as well as the root DSA-specific entry (DSE) (rootDSE) attributes for the domain controller.

To find more information about ADSI Edit, see "Support Tools Help" in Tools and Settings Collection.

**Dssite.msc: Active Directory Sites and Services**

**Category**
Administrative Tools, MMC snap-in

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows 7 with Remote Server Administration Tools (RSAT) installed.<br>• Windows Vista with Remote Server Administration Tools (RSAT) installed.<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Active Directory Sites and Services to create, modify, and delete site configuration objects in Active Directory, including sites, subnets, connection objects, and site links. You can also use Active Directory Sites and Services to create the intersite topology, including mapping subnet addresses to sites, creating and configuring site links, creating manual connection objects, forcing replication over a connection, setting a domain controller to be a global catalog server, and selecting preferred bridgehead servers.

**Repadmin.exe: Repadmin**

**Category**
Windows Support Tools, command-line

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows 7 with Remote Server Administration Tools (RSAT) installed.<br>• Windows Vista with Remote Server Administration Tools (RSAT) installed.<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Repadmin is used to view the replication information on domain controllers. You can determine the last successful replication of all directory partitions, identify inbound and outbound replication partners, identify the current bridgehead servers, view object metadata, and generally manage Active Directory replication topology. You can use Repadmin to force replication of an entire directory partition or of a single object. You can also list domain controllers in a site.

Repadmin is extended in Windows Server 2003 to enable commands to target sets of domain controllers. For example, you can target all domain controllers in a site or domain, or all domain controllers that are global catalog servers. In Windows 2000 Server, Repadmin can report information about only one domain controller at a time.

For more information about Repadmin, see "Support Tools Help" in Tools and Settings Collection.

**Ldp.exe: Ldp**

**Category**
Windows Support Tools, GUI

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows 7 with Remote Server Administration Tools (RSAT) installed.<br>• Windows Vista with Remote Server Administration Tools (RSAT) installed.<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Ldp is a Lightweight Directory Access Protocol (LDAP) graphical user interface (GUI) tool that you can use to perform operations such as connect, bind, search, modify, add, and delete against any LDAP-compatible directory, such as AD DS. You can also use Ldp to view objects that are stored in AD DS, along with their metadata, for example, security descriptors and replication metadata.

You can use Ldp to view and edit the **updateCachedMemberships** operational attribute on the **rootDSE**.

For more information about Ldp, see "Support Tools Help" in Tools and Settings Collection.

**Global Catalog Registry Entries**

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following registry entries are associated with the global catalog.

**NTDS Parameters**

The following registry entries under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** control or contain information about the configuration of the global catalog.

**Global Catalog Promotion Complete**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2003:

Used for Install From Media. This entry is set in conjunction with the domain controller setting its rootDSE attribute **isGlobalCatalogReady** to TRUE, the Net Logon service on the domain controller registering SRV resource records that specifically advertise the global catalog in DNS, and the domain controller beginning to listen on port 3268.

**Global Catalog Partition Occupancy**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003:

The requirement for read-only replicas that must be added (replication partner established) or synchronized (replication completed), or both, on the global catalog server before the server is advertised in DNS. Lower occupancy levels specify varying levels of replication completeness, including advertising in DNS when all read-only replicas of only those domains represented in the domain controller's site are synchronized.

**Version**
Windows 2000 Server with SP3 and later:

The occupancy level requires full synchronization of all read-only replicas.

**Version**
Windows 2000 Server with Service Pack (SP) 2 and earlier:

The occupancy level requires only the replicas of domains in the site.

**Global Catalog Delay Advertisement**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server

Overrides the requirements set in **Global Catalog Partition Occupancy** entry and allows global catalog advertisement without requiring full synchronization of all read-only replicas. If you do not set this value, checking for synchronized read-only partitions continues to occur at 30-minute intervals until the requirements are met.

**Cached Membership Site Stickiness (minutes)**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003

The maximum time during which an account's cached membership can be refreshed automatically without the account having to log on in this site. The default value is one-half the value of the account's site affinity setting, which is 180 days by default. If the account has not logged on in more than 90 days, the account's group membership cache expires and must be reinstated at the next logon by contacting a global catalog server.

**Cached Membership Staleness (minutes)**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003

The maximum staleness to tolerate when using a cached group membership. The default value is one week. If the cached membership age is greater than this interval and no global catalog server is available, the logon fails. If no value is present, the

default value is used.

**Cached Membership Refresh Interval (minutes)**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003

The frequency of automatic cache refresh. The default value is eight hours. If no value is present, the default value is used.

**Cached Membership Refresh Limit**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003

The maximum number of user and computer accounts that are refreshed during a group membership cache refresh.

**SamNoGcLogonEnforceKerberosIpCheck**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa**

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003

By default, allows site affinity to be tracked for Kerberos logons that originate outside the site. This setting only applies to Kerberos logons; it will not affect site affinity caching for NTLM logons from different sites.

**SamNoGcLogonEnforceNTLMCheck**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa**

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003

A value of 1 configures Security Accounts Manager (SAM) to not give site affinity to NTLM logon requests that are network logon requests; it may not prevent caching for other logon types.

**NTDS Diagnostics**
The following registry entry under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics** control the logging level for the component or process that is specified in the entry name.

**Global Catalog**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server

The logging level for the global catalog on the domain controller. The value is set to an integer from 0 (no logging) through 5 (most verbose logging).

**20 Group Caching**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003

The logging level for Universal Group Membership Caching on a domain controller in a site where this feature is enabled. The value is set to an integer from 0 (no logging) through 5 (most verbose logging). Significant events are reported at logging level

2. with many additional events reported at logging level 5.

**Global Catalog Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with global catalog servers.

**Group Policy Settings Associated with the Global Catalog**

| Group Policy Setting | Description |
|---|---|
| Automated Site Coverage by the DC Locator DNS SRV Records | Determines whether domain controllers will dynamically register DC Locator site-specific SRV resource records for the closest sites where no domain controller for the same domain exists (or no global catalog server for the same forest exists). These DNS records are dynamically registered by the Net Logon service, and they are used to locate domain controllers. |
| Sites Covered by the GC Locator DNS SRV Records | Specifies the sites for which the global catalog servers should register the site-specific GC Locator SRV resource records in DNS. These records are registered in addition to the site-specific SRV resource records registered for the site where the global catalog server resides and, if the global catalog server is appropriately configured, for the sites without a global catalog server in the same forest for which this global catalog server is the closest global catalog server. These records are registered by Net Logon service.<br><br>If this policy is not configured, then it is not applied to any global catalog servers and global catalog servers use their local configuration. |

**Network Ports Used by Global Catalog Servers**

The following table shows the network ports that are used by global catalog servers.

**Port Assignments for Global Catalog Servers**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | | 3268 (global catalog) |
| LDAP | | 3269 (global catalog Secure Sockets Layer [SSL]) |
| LDAP | 389 | 389 |
| LDAP | | 636 (SSL) |
| RPC/REPL | | 135 (endpoint mapper) |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- What Is the Global Catalog?
- How the Global Catalog Works

# Operations Masters Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Operations Masters Technical Reference**

Operations masters prevent conflicting updates from being made to the Active Directory® Domain Services (AD DS) directory service database by performing specific tasks that no other domain controllers are permitted to perform. Most updates can be made to AD DS on any domain controller and because they are replicated by using a multimaster replication model, all other domain controllers receive the updates.

| **Note** |
| --- |
| In Windows Server® 2003 and Microsoft Windows® 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008 , the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

However, specific updates, such as password changes, the addition or deletion of a domain or application directory partition, or writing to the AD DS schema, are performed on domain controllers that hold one of five operations master roles to prevent database inconsistencies.

When you install AD DS on the first domain controller in a forest, it automatically assigns five operations master roles to that domain controller. Each of the roles can be transferred to another domain controller at any time.

This subject explains what operations masters are, and how each operations master role functions in Active Directory.

**In this subject**


- What are Operations Masters?
- How Operations Masters Work
- Operations Masters Tools and Settings

# What are Operations Masters?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Operations Master Roles
- Operations Master Dependencies
- Related Information


Active Directory Domain Services (AD DS) defines five operations master roles: the schema master, domain naming master, relative identifier (RID) master, primary domain controller (PDC) emulator, and infrastructure master. The domain controllers that hold operations master roles are designated to perform specific tasks to ensure consistency and to eliminate the potential for conflicting entries in the AD DS database.

| **Note** |
| --- |
| In Windows Server® 2003 and Microsoft Windows® 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008 , the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

AD DS is a multimaster enabled database, which provides the flexibility of allowing changes to occur at any domain controller in the forest. However, because it is multimaster enabled, it can also allow conflicting updates that can potentially lead to problems when data is replicated throughout the domain or forest.

The general approach to resolving AD DS replication conflicts is to order all update operations (Add, Modify, Move, and Delete) by assigning a globally unique stamp to the originating update. Each replicated attribute value (or multivalue) is stamped during the originating update and this stamp is replicated with the value. The stamp that is applied during an originating write consists of a version number, a time stamp of when the originating write occurred, and the originating domain controller. Conflicts are resolved by comparing the version number. If two stamps have the same version number, the originating time almost always breaks the tie. In the extremely rare event that the same attribute is updated on two different domain controllers during the same second, the originating domain controller breaks the tie in an arbitrary fashion.

Although this resolution method is acceptable, some changes are too difficult to resolve by using the stamp of the originating update. In such cases, it is best to prevent the conflict from occurring rather than to try to resolve it after it has occurred.

When changes such as the addition or removal of domains to a forest or password changes are made, AD DS performs them in a single-master fashion to prevent conflicting updates from occurring. In a single-master update model, only one domain controller in the entire directory is allowed to process the update. This is similar to the role of a Windows NT PDC, in which the PDC is responsible for processing all updates in a given domain.

AD DS extends the single-master model to include multiple roles that are responsible for different types of updates. AD DS also provides the ability to transfer an operations master role to another domain controller.

- By designating a single domain controller to manage specific tasks, AD DS enhances your ability to avoid conflicts in the directory, ensure consistency of the schema, and to add a domain to, or remove a domain from a forest. Operations masters also maintain interaction between earlier versions of Windows operating systems that do not include AD DS or Active Directory and they maintain consistent group-to-user references across domains.

**Operations Master Roles**

The five operations master roles are assigned automatically when the first domain controller in a given domain is created. Two forest-level roles are assigned to the first domain controller created in a forest and three domain-level roles are assigned to the first domain controller created in a domain.

**Forestwide Operations Master Roles**

The schema master and domain naming master are forestwide roles, meaning that there is only one schema master and one

domain naming master in the entire forest.

**Schema Master**
The schema master is responsible for performing updates to the AD DS schema. The schema master is the only domain controller that can perform write operations to the directory schema. Those schema updates are replicated from the schema master to all other domain controllers in the forest. Having only one schema master for each forest prevents any conflicts that would result if two or more domain controllers attempt to concurrently update the schema.

**Domain Naming Master**
The domain naming master manages the addition and removal of all domains and directory partitions, regardless of domain, in the forest hierarchy. The domain controller that has the domain naming master role must be available in order to perform the following actions:

- Add new domains or application directory partitions to the forest.
- Remove existing domains or application directory partitions from the forest.
- Add replicas of existing application directory partitions to additional domain controllers.
- Add or remove cross-reference objects to or from external directories.
- Prepare the forest for a domain rename operation.

**Domainwide Operations Master Roles**
The other operations master roles are domainwide roles, meaning that each domain in a forest has its own RID master, PDC emulator, and infrastructure master.

**RID Master**
The relative identifier (RID) operations master allocates blocks of RIDs to each domain controller in the domain. Whenever a domain controller creates a new security principal, such as a user, group, or computer object, it assigns the object a unique security identifier (SID). This SID consists of a domain SID, which is the same for all security principals created in the domain, and a RID, which uniquely identifies each security principal created in the domain.

**PDC Emulator**
**The primary domain controller (PDC) emulator operations master**. The PDC emulator receives preferential replication of password changes that are performed by other domain controllers in the domain, and it is the source for the latest password information whenever a logon attempt fails as a result of a bad password. It is a preferred point of administration for services (examples are Group Policy and Distributed File System, DFS). For this reason, of all operations master roles, the PDC emulator operations master role has the highest impact on the performance of the domain controller that hosts that role. The PDC emulator in the forest root domain is also the default Windows Time service (W32time) time source for the forest.

The PDC emulator operations master also processes all replication requests from Windows NT Server 4.0 backup domain controllers (BDCs). It processes all password updates for clients not running Active Directory–enabled client software, plus any other directory write operations.

**Infrastructure Master**
The infrastructure operations master is responsible for updating object references in its domain that point to the object in another domain. The infrastructure master updates object references locally and uses replication to bring all other replicas of the domain up to date. The object reference contains the object's globally unique identifier (GUID), distinguished name and possibly a SID. The distinguished name and SID on the object reference are periodically updated to reflect changes made to the actual object. These changes include moves within and between domains as well as the deletion of the object. If the infrastructure master is unavailable, updates to object references are delayed until it comes back online.

**Operations Master Dependencies**
Domain controllers designated as operations masters have the following dependencies:

**Operations master placement**
Because operations masters are critical to the long-term performance of the directory, they must be available to all domain controllers and desktop clients that require their services. Careful placement of your operations masters becomes more important as you add more domains and sites to build your forest.

By improperly placing operations master role holders, you might prevent clients running Windows NT Workstation 4.0, Windows 95, or Windows 98 without the Active Directory client installed from changing their passwords, or be unable to add

domains and new objects, such as users and groups. You might also be unable to make changes to the schema. In addition, name changes might not properly appear within group memberships that are displayed in the user interface.

As your environment changes, you must avoid the problems associated with improperly placed operations master role holders. Eventually, you might need to reassign the roles to other domain controllers.

Although you can assign the operations master roles to any domain controller, follow these guidelines to minimize administrative overhead and ensure the performance of Active Directory:

- Leave the two forestwide roles on a domain controller in the forest root domain.
- Place the two forestwide roles on a global catalog server.
- Place the three domainwide roles on the same domain controller.
- In a forest that contains multiple domains, do not place the domainwide roles on a global catalog server unless all domain controllers in the domain are also global catalog servers.
- Place the domainwide roles on a higher performance domain controller.
- Adjust the workload of the operations master role holder, if necessary.

**Active Directory replication**
Operations masters replicate changes made on them throughout the domain or forest, depending on whether they hold a domain or forest role. AD DS replication must be working properly in order for the other domain controllers to receive these changes.

**Domain Name System (DNS)**
AD DS requires that DNS is properly designed and deployed so that domain controllers can correctly resolve DNS names of replication partners. If DNS is not working properly, operations masters cannot be contacted to perform their specific domain or forest functions.

**Security**
User rights for designating operations master roles can be set for groups or users in a forest. This allows you to limit or add to the group of default users that can change operations master role holders in a forest or domain. The following user rights are required to change operations master role holders:

- The **Change Schema Master** right is required to transfer or seize the schema master. By default, only members of the Schema Administrators group are assigned this right.
- The **Change Domain Master** right is required to transfer or seize the domain naming master role. Be default, only members of the Enterprise Admins group are assigned this right.
- The **Change PDC** right is required to transfer or seize the PDC emulator role. By default, only members of the Domain Admins group are assigned this right.
- The **Change Infrastructure Master** right is required to transfer or seize the infrastructure master. By default, only members of the Domain Admins group are assigned this right.
- The **Change RID Master** right is required to transfer or seize the RID master role. By default, only members of the Domain Admins group are assigned this right.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Active Directory Replication Technologies
- DNS Support for Active Directory Technical Reference

# How Operations Masters Work

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Operations Master Protocols
- Operations Master Roles and Functionality
- Role Transfer and Seizure
- Network Ports Used by Operations Masters
- Related Information


Active Directory Domain Services (AD DS) is a multimaster enabled database, which provides the flexibility of allowing changes to the directory to occur at any domain controller. However, some changes, if performed in a multimaster fashion, can cause errors in the AD DS database. For these changes, one domain controller, called an operations master, is assigned to accept requests for such changes. Operations masters perform updates to the directory in a single-master fashion, meaning that only the domain controller assigned to hold the operations master role is allowed to process the update. This single-master update model prevents conflicting updates from being made to the AD DS database.

| **Note** |
| --- |
| In Windows Server® 2003 and Microsoft Windows® 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008 , the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

Five operations masters are assigned to perform specific tasks in an AD DS environment. The schema master and the domain naming master are forestwide roles, meaning that only one of each of these types of operations masters is in a forest. The relative identifier (RID) master, infrastructure master, and primary domain controller (PDC) emulator master are domainwide roles, meaning one of each of these types of operations masters is in each domain in a forest.

Because operation masters are used to perform specific tasks and are important to the performance of the directory, they must be available to all domain controllers and directory clients that require their services.

This section describes the functionality and interactions of each operations master.

**Operations Master Protocols**

Operations masters use the same protocols as other domain controllers. The protocols that package the data sent to and from a domain controller are described in the following table.

**Operations Master Protocols**

| Protocol | Description |
| --- | --- |
| Lightweight directory access protocol (LDAP) | The primary directory service protocol that specifies directory communications. It runs directly over TCP/IP, and it can also run over UDP connectionless transports (UDP access is primarily used by the domain controller Locator process). Clients use LDAP to query, create, update, and delete information that is stored in a directory service over a TCP connection through the TCP port 389. AD DS supports LDAP v2 (RFC 1777) and LDAP v3 (RFC 2251). LDAP v3 is an industry standard that can be used with any directory service that implements the LDAP protocol. LDAP is the preferred and most common way of interacting with AD DS. |

| | |
|---|---|
| Remote procedure call (RPC) | Protocol for replication (REPL), domain controller management communications, and SAM-related communications. RPC is a powerful, robust, efficient, and secure interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality residing in a different process. That different process can be on the same computer, on the local area network (LAN), or across the Internet. |
| Simple mail transfer protocol (SMTP) | Protocol for replication communications when a permanent, "always on" network connection does not exist between two domain controllers. SMTP is used to transport and deliver messages based on specifications in Request for Comments (RFC) 821 and RFC 822. SMTP can replicate configuration, schema, and global catalog replicas only (not writable domain data). |

For more information about AD DS protocols, see "How the Data Store Works."

**Operations Master Roles and Functionality**

Five operations master roles manage single-master operations in AD DS.

Two operations master roles exist in each forest:

- **The schema master**, which governs all changes to the schema.
- **The domain naming master**, which adds and removes domains to and from the forest.

In addition to the two forestwide operations master roles, three operations master roles exist in each domain:

- **The primary domain controller (PDC) emulator**. The PDC emulator processes all replication requests from Microsoft Windows NT 4.0 backup domain controllers and processes all password updates for clients that are not running Active Directory–enabled client software.
- **The relative identifier (RID) master**. The RID master allocates RIDs to all domain controllers to ensure that all security principals have a unique identifier.
- **The infrastructure master**. The infrastructure master for a given domain maintains a list of the security principals from other domains that are members of groups within its domain.

**Schema Master**

The schema master controls all originating updates to the schema. The schema contains the master list of object classes and attributes that are used to create all AD DS objects, such as computers, users, and printers. The domain controller that holds the schema master role is the only domain controller that can perform write operations to the directory schema. These schema updates are replicated from the schema operations master to all other domain controllers in the forest.

**Domain Naming Master**

The domain naming master is a forestwide operations master role because it manages the addition and removal of all directory partitions, regardless of domain, in the forest hierarchy. The domain controller that has the domain naming master role must be available in order to perform the following actions:

- Add or remove domains.
- Add or remove application directory partitions.
- Add or remove cross-reference objects.
- Validate domain rename instructions.

**Add or Remove Domains**

Only the domain controller holding the domain naming master role has the authority to add a new domain. The domain naming master manages this process, preventing multiple domains from joining the forest with the same domain name. When you use the Active Directory Domain Services Installation Wizard to create or remove a child domain, it contacts the domain naming master and requests the addition or deletion. The domain naming master is responsible for ensuring that domain names are unique. If the domain naming master is unavailable, you cannot add domains or remove them from the forest.

**Add or Remove Application Directory Partitions**

Application directory partitions are special partitions that can be created on domain controllers running Windows Server 2003 or later to provide LDAP storage for dynamic data. In Windows 2000 Server, nondomain data is limited to configuration and schema data, which is replicated to every domain controller in the forest. For domain controllers that run Windows Server 2003 or later, application directory partitions provide storage for nondomain, application-specific data that can be replicated to any arbitrary set of domain controllers in the forest — as few or as many as needed by the application that uses the data.

The Domain Name System (DNS) creates and uses application directory partitions by default when it is installed in a forest with domain controllers that run Windows Server 2003 or later. DNS automatically creates two default DNS application directory partitions below the forest root domain in the domain hierarchy, one for the forest (ForestDnsZones) and one for the domain (DomainDnsZones). Thereafter, when you install AD DS to create a new domain in the forest and configure that server to be a DNS server, DNS creates one default application directory partition below the new domain directory partition in the domain hierarchy.

If the domain controller hosting the domain naming operations master role is not running Windows Server 2003 or later, or if it is unavailable, you cannot add or remove application directory partitions from the forest.

**Add or Remove Cross-Reference Objects**

When AD DS is installed to create the first domain controller in a new forest, the schema, configuration, and domain directory partitions are created on the domain controller. At this time, a cross-reference object (class **crossRef**) is created for each directory partition in the Partitions container in the configuration directory partition (CN=partitions,CN=configuration,DC=*forestRootDomain*). Creation of each subsequent directory partition in the forest, either by installing AD DS to create a new domain or by creating a new application directory partition on an existing domain controller, initiates the creation of an associated cross-reference object in the Partitions container.

| **Note** |
| --- |
| You can also manually create a cross-reference object for an application directory partition. |

A cross-reference object identifies the name and server location of each directory partition in the forest. The replication system uses this information to identify servers that store the same directory partitions. LDAP queries use cross-reference objects to create referrals to different domains. If the domain naming master is unavailable, you cannot add or remove cross-reference objects.

**Validate Domain Rename Instructions**

When you use the domain rename tool, rendom.exe, to rename an AD DS domain, the tool must be able to access the domain naming operations master. The domain naming master is the domain controller responsible for validating the instructions that the domain rename tool has generated for the new forest.

Certain changes occur on the domain naming master in preparation for the actual execution of a domain rename. On the domain naming master, the XML-encoded script containing the domain rename instructions is written to the **msDS-UpdateScript** attribute on the Partitions container object (cn=partitions,cn=configuration,dc=*ForestRootDomain*) in the configuration directory partition. The Partitions container can only be updated on the domain controller holding the domain naming operations master role for the forest. Therefore, the **msDS-UpdateScript** attribute must be changed on the domain naming master.

In addition to the **msDS-UpdateScript** attribute value being written to the Partitions container, the new DNS name of each domain being renamed is also written by rendom.exe to the **msDS-DnsRootAlias** attribute on the cross-reference object (class **crossRef**) corresponding to that domain. Again, because cross-reference objects are stored in the Partitions container and the Partitions container can only be updated on the domain naming master, the **msDS-DnsRootAlias** attribute can only be changed on the domain naming operations master.

The domain naming master replicates the script stored in the **msDS-UpdateScript**attribute and the DNS name in the **msDS-DnsRootAlias** attribute to all domain controllers in the forest through scheduled replication of the configuration directory partition.

In preparation of a domain rename operation; rendom.exe uses the domain naming operations master to:

- Retrieve all information needed to compute the list of rename instructions for updating the configuration and schema directory partitions.
- Write the resulting script to the **msDS-UpdateScript** attribute of the Partitions container.
- Change the **msDS-DnsRootAlias** attribute of all cross-reference objects for domains that are being renamed.

- Validate the forest description against the current state of the forest. If any of these validity checks fails, the command fails. The following requirements are verified:
    - Each existing domain is part of the new forest.
    - The new forest is well formed.
    - The new forest does not re-assign domain names that are being relinquished as part of the current domain rename operation.

For more information about domain rename, see "Domain Rename Technical Reference."

**Relative Identifier Master**

The relative identifier (RID) master is a domainwide operations master role. The RID master is responsible for allocating sequences of unique RIDs to each domain controller in its domain and for moving objects from one domain to another.

You can create a new security principal object (user, group, or computer) on any domain controller. When you create a security principal object, the domain controller attaches a unique Security ID (SID) to the object. There are four elements of a domain SID, one of which is the RID for the domain.

The following table describes the elements of the domain **SID: S-1-5-Y1-Y2-Y3-Y4**.

**Elements of a Security Identifier (SID)**

| SID element | Description |
|---|---|
| S-1 | Indicates a revision 1 SID. At this time 1 is the only SID revision in use. |
| 5 | Indicates the issuing agency or issuing authority. A 5 always indicates Windows NT or AD DS domains. Note that a well-known SID may use a 1 or 0 as the issuing identity to signify that it is a well known SID. |
| Y1-Y2-Y3 | The domain identifier portion of the SID. This is the same for every security principal object created in that domain. |
| Y4 | The relative ID (RID) for the domain, which represents a user name or group. This is obtained from the RID pool on a domain controller at the time the object is created. |

**RID Allocation**

Domain controllers running AD DS have a shared RID pool. The RID operations master is responsible for maintaining a pool of RIDs to be used by the domain controllers in its domain and for providing groups of RIDs to each domain controller when necessary. When a new AD DS domain controller is added to the domain, the RID master allocates a batch of approximately 500 RIDs from the domain RID pool to that domain controller. Each time a new security principal is created on a domain controller, the domain controller draws from its local pool of RIDs and assigns one to the new object. When the number of RIDs in a domain controller's RID pool falls below approximately 100, that domain controller submits background requests (by means of RPC) for additional RIDs from the domain's RID master. The RID master allocates a block of approximately 500 RIDs from the domain's RID pool to the pool of the requesting domain controller.

The RID master does not actually maintain a pool of numbers. Rather, it maintains the highest value of the last range it allocated. When a new request is received, it increments that value by one to establish the low value in the new RID pool and then adds four hundred and ninety nine to establish the new maximum value. It sends these two values to the requesting domain controller to use as its next allocation of RIDs.

If a domain controller's local RID pool is empty, and it cannot contact the domain's RID master to request additional RIDs, the domain controller will log event ID 16645, indicating that the maximum account identifier allocated to the domain controller has been assigned and the domain controller has failed to obtain a new identifier pool from the RID master. Likewise, when attempting to add new objects in AD DS, such as users, computers, or domain controllers, you might notice event ID 16650 in the System log indicating that the object cannot be created because the directory service was unable to allocate a relative identifier. Network connectivity to the RID master might have been lost or the RID master might have been removed from the network. In any case, you cannot create new security principal objects on the domain controller until RID pool acquisition is successful.

**Cross-Domain Moves**
Migrating AD DS objects from one domain to another requires the availability of the RID master. You can only move an object out of its domain if the domain's RID master can be contacted. Requiring that objects be moved from one domain to another by using the RID master prevents AD DS from creating two objects in different domains with the same unique identifier. This might happen if one object is moved from two domain controllers simultaneously to two different domains.

You can use the Active Directory Migration Tool (ADMT) to perform intraforest migrations of domain objects from one domain (the source domain) to another (the target domain). The RID master must be online and available in the source domain for ADMT to migrate successfully. If the RID master is unavailable, cross-domain migration of AD DS objects will fail.

**RID Attributes in AD DS**
The following are RID-related attributes in AD DS:

**FsmoRoleOwner**
DN path: CN=RID Manager$,CN=System,DC=<domain>,DC=com

This attribute points to the Domain Name path for the current RID master's NTDS Settings object according to the domain controller that is being queried.

**RidAvailablePool**
DN path: CN=RID Manager$,CN=System,DC=<domain>,DC=com

This attribute defines the global RID space from which RID pools are allocated to the RID master.

**RidAllocationPool**
DN Path: CN=Rid Set,CN=<computername>,OU=domain controllers,DC=<domain>,DC=com

Each domain controller has two RID pools: the one that they are currently allocating from, and the pool that they will use next. This attribute defines the RID pool that will be used in the domain when the current RID pool is exhausted.

**RidNextRid**
DN Path: CN=Rid Set,CN=<computername>,OU=domain controllers,DC=<domain>,DC=com

This attribute defines the next free RID in the current allocation pool that is assigned to the next security principal created on the local domain controller. **RidNextRid** is a non-replicated value in AD DS.

**RidPreviousAllocationPool**
DN Path: CN=Rid Set,CN=<computername>,OU=domain controllers,DC=<domain>,DC=com

This attribute defines the RID pool from which the RID master actually allocates from. The value for **RidNextRid** is implicitly a member of this pool.

**RidUsedPool**
DN Path: CN=Rid Set,CN=<computername>,OU=domain controllers,DC=<domain>,DC=com

This attributed defines the RID pools that have been used by a domain controller.

**NextRid**
DN Path: DC=<domain>,DC=com

This attribute defines the next RID field used by the RID master.

**Primary Domain Controller (PDC) Emulator**
The PDC emulator serves as primary domain controller for compatibility with earlier Windows operating systems. Windows 2000 Server and Windows Server 2003 interoperate with Windows NT workstations, member servers, and backup domain controllers. Windows Server 2008 R2 and Windows Server 2008 domain controllers cannot interoperate with Windows NT domain controllers.

The primary domain controller (PDC) in a Windows NT 3.51 or Windows NT 4.0 domain is responsible for the following:

- Processing password changes from both users and computers
- Replicating updates to backup domain controllers
- Running the Domain Master Browser

Active Directory uses multimaster replication for most directory updates, including password changes. Therefore, if the PDC emulator becomes unavailable, the impact is small. However, in a mixed environment with Windows NT–based domain controllers and Active Directory, the unavailability of the PDC emulator has the following impact:

- When a user of a computer running Windows NT Workstation 4.0, Windows 95, or Windows 98 without the Active Directory client installed attempts a password change, the user sees a message similar to the following: "Unable to change password on this account. Please contact your system administrator."

- In a mixed domain, the event logs of Windows NT 4.0 BDCs contain entries showing failed replication attempts.

- In a mixed domain, when a user tries to start User Manager on a Windows NT 4.0 backup domain controller, it results in a "domain unavailable" error message. If User Manager is already running, you will see an "RPC server unavailable" message. Attempting to create an account by using the **net user /add** command results in a "could not find domain controller for this domain" message. When you run Server Manager, you will see a message similar to the following: "Cannot find the primary domain controller for <*domain name*>. You may administer this domain, but certain domain-wide operations will be disabled."

As operating systems are upgraded, either to Windows 7, Windows Vista, Windows XP, Windows 2000, Windows Server 2003, or (for Windows NT Workstation 4.0, Windows 95, and Windows 98) by installing the Active Directory client, they cease to rely on the PDC and, instead, behave in the following manner:

- Clients do not make password changes at the PDC emulator. Instead, clients update passwords at any domain controller in the domain.

- The PDC emulator does not receive Windows NT 4.0 replication requests after all Windows NT 4.0 BDCs in a domain are upgraded to AD DS.

- Clients use Active Directory to locate network resources. They do not require the Computer Browser service.

In addition, the domain controller holding the PDC emulator role performs the following functions:

- When password changes are performed by other domain controllers in the domain, they are sent to the PDC emulator by using higher priority replication.

- When an authentication fails with an invalid password at other domain controllers in the domain, the authentication request is retried at the PDC emulator before failing. If a recent password update has reached the PDC emulator, the retried authentication request should succeed.

- When an authentication succeeds for an account for which the most recent authentication attempt at the domain controller failed, the domain controller communicates this fact ("zero lockout count") to the PDC emulator. This resets the lockout counter at the PDC emulator in case another client attempts to validate the same account by using a different domain controller.

Therefore, when the PDC emulator is unavailable, you might experience an increase in support requests regarding password difficulties. However, unlike the Windows NT 4.0 environment, where the PDC was the only computer that wrote the updated password to the domain, in AD DS, any domain controller can write the password update to the directory and normal replication will ensure that all domain controllers in the domain eventually become aware of the change. This will occur even if the PDC emulator operations master remains unavailable.

**Windows Server 2003 Well Known Security Principals**

The PDC emulator operations master is responsible for managing well-known security principals. For more information about well-known security principals that are added in Windows Server 2008, see Appendix A: Background Information for Upgrading Active Directory Domains (http://go.microsoft.com/fwlink/?LinkId=192800). The rest of this section covers information about well-known security principals that are added in Windows Server 2003.

When you upgrade your environment from Windows 2000 Server, it is important that you upgrade the PDC emulator early in the upgrade process so that the "CN=WellKnown Security Principals,CN=Configuration,DC=<YourDomain>" container is updated with the well-known security principals that are introduced in Windows Server 2003.

After upgrading the Windows 2000–based domain controller holding the role of the PDC emulator in each domain in the forest to Windows Server 2003, several new well-known and built-in groups are created and some new group memberships are established. If you transfer the PDC emulator role to a Windows Server 2003–based domain controller instead of upgrading it, these groups will be created when the role is transferred. The new well-known and built-in groups are:

- Builtin\Remote Desktop Users
- Builtin\Network Configuration Operators
- Performance Monitor Users
- Performance Log Users
- Builtin\Incoming Forest Trust Builders
- Builtin\Performance Monitoring Users
- Builtin\Performance Logging Users
- Builtin\Windows Authorization Access Group
- Builtin\Terminal Service License Server

The newly established group memberships are:

- If the Everyone group is in the Pre-Windows 2000 Compatible Access group, the Anonymous Logon group and Authenticated Users group are also added to the Pre-Windows 2000 Compatible Access group.
- The Network Servers group is added to the Performance Monitoring Users group.
- The Enterprise Domain Controllers group is added to the Windows Authorization Access group.

In addition, when upgrading the Windows 2000–based domain controller that holds the role of the PDC emulator in the forest root domain, the following additional security principals are created:

- LocalService
- NetworkService
- NTLM Authentication
- Other Organization
- Remote Interactive Logon
- SChannel Authentication
- This Organization

**AdminSDHolder**

The Administrator security descriptor holder protects administrative groups through a background task that computes the set of memberships and checks whether their security descriptors are well-known protected security descriptors. If the security descriptor of any administrative account does not match that of AdminSDHolder, the security descriptor is overwritten with the security descriptor of AdminSDHolder. This task is executed only on the domain controller that holds the primary domain controller (PDC) emulator operations master role.

**DNS and the PDC**

The first domain controller deployed in every domain is automatically assigned the PDC emulator role. During the AD DS installation on this server, Net Logon registers the _ldap._tcp.pdc._msdcs.DnsDomainName DNS SRV resource record that enables clients to locate the PDC emulator. Only the PDC emulator of the domain registers this SRV resource record.

**DFS and the PDC**

When a change is made to a Distributed File System (DFS) domain-based namespace, the change is made on the domain controller that acts as the PDC emulator master. DFS root servers that host domain-based namespaces poll the PDC emulator periodically to obtain an updated version of the DFS metadata, and then they store this metadata in memory. Therefore, if the domain controller that holds the PDC emulator role is not available, DFS will not function properly.

For more information about DFS and the role of the PDC emulator, see "How DFS Works" on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=48186.

**Raising the Domain Functional Level**

The functional level of a domain or forest defines the set of advanced AD DS features that are available in that domain or forest. The functional level of a domain or forest also defines the set of Windows Server operating systems that can run on the domain controllers in that domain or forest.

The PDC emulator is the domain controller targeted when you raise the domain functional level. Therefore, the PDC emulator must be available to raise the domain functional level. For more information domain and forest functional levels, see "Active Directory Functional Levels Technical Reference."

**Infrastructure Master**

The infrastructure master is responsible for updating the group-to-user references whenever the members of groups are renamed or changed within a domain.

For example, suppose you use Active Directory Users and Computers to add a user to a group within a single domain. While still connected to the same domain controller, you can view the group's membership and see the user you just added. If you then rename the user object (that is, change its **CN** attribute) and then display the group membership, you instantly see the user's new name in the list of group members.

However, when the user and group are in different domains, there is a lag between the time when you rename a user object and when a group containing that user displays the user's new name. This time lag is inevitable in a distributed system where sites function independently and must rely on replication for changes to be distributed to all sites.

The domain controller holding the infrastructure master role for the group's domain is responsible for updating the cross-domain group-to-user reference to reflect the user's new name. Periodically, the infrastructure master scans the domain accounts and verifies the membership of the groups. If a user account is moved to a new domain, the infrastructure master identifies the user account's new domain and updates the group accordingly. After the infrastructure master updates these references locally, it uses replication to update all other replicas of the domain. If the infrastructure master is unavailable, these updates are delayed.

**Phantom Records**

Whenever a domain controller contains a reference to an object in another domain, the domain controller's local database contains a phantom record for that object. Within the local database, a reference to the object is represented as the database pointer to (in fact the primary key of) the phantom record.

For example, if a user in domain A is a member of a group in domain B, the local database of a domain controller in domain B, holding the group, contains a phantom record for that user. The group references the user by pointing to the phantom. The infrastructure master is responsible for updating the phantom records in its local database. If the distinguished name of the user in domain A changes, the infrastructure master in domain B is responsible for updating its reference to the user.

Each phantom record in the database contains the referenced object's GUID, its SID (for references to security principals), and its distinguished name. If the referenced object is renamed or moved, its GUID does not change; its SID changes if the move is cross-domain, and its distinguished name always changes.

The infrastructure master periodically examines the phantom records in its local database. It compares the data in a phantom record to the corresponding data on the replica of the referenced object contained in the global catalog. Because global catalog servers receive regular updates for all objects in the forest through replication, the global catalog's data is – allowing for replication latency – always up to date. If the SID or the distinguished name in a phantom record does not match the SID or the distinguished name of the object in the global catalog, then the infrastructure master updates its local database with the values from the global catalog. The infrastructure master then replicates the new values to the other domain controllers in its domain.

A global catalog server holds a complete replica of its own domain and a partial replica of every other domain in the forest. The complete replica of its own domain includes the GUID, SID, and distinguished name of the domain's objects. The partial replica of another domain includes the GUIDs, SIDs, and distinguished names of all the objects in that domain. Because this data exists

in the global catalog server's local database, there are no phantom records representing cross-domain object references on a global catalog server.

Cross-domain object references on a global catalog server are represented in the same way as intradomain references on any domain controller – that is, as a direct pointer to the referenced object. Therefore, if the infrastructure master is running on a global catalog server, it never finds any phantom records representing cross-domain references in its local database. Consequently, the infrastructure master is unable to replicate any updated references to cross-domain objects to any other domain controller in its domain. For this reason, the infrastructure master should never run on a global catalog server in a forest that contains multiple domains.

If every domain controller in a domain is a global catalog server, then no domain controller in the domain has any phantom records that need to be updated. In this case, the infrastructure master has no work to do and remains in a dormant state.

For more information about phantom records, see "How the Data Store Works" in this collection.

**Role Transfer and Seizure**

You can reassign an operations master role by transfer or, as a last resort, by seizure.

Role transfer is the preferred method to move an operations master role from one domain controller to another. When an operations master role is transferred, the previous role holder replicates its most recent updates to the new role holder to ensure that no information is lost. After the transfer completes, the previous role holder reconfigures itself so that it no longer attempts to perform as the operations master while the new domain controller assumes those duties. This prevents the possibility of duplicate operations masters existing on the network at the same time, which can lead to inconsistencies in the directory.

Role seizure should be used only as a last resort to forcibly remove an operations master role from one domain controller and assign the role to a different domain controller. Use this process only when the previous operations master fails and remains out of service for an extended period of time. For more information about the ramifications of roles seizure, see Seizing Operations Master Roles.

During a role seizure, the domain controller that assumes the operations master role does not verify that replication is updated, so recent changes can be lost. Because the previous role holder is unavailable during the role seizure, it cannot know that a new role holder exists. If the previous role holder comes back online, it assumes that it is still the operations master. This can result in duplicate operations master roles on the network, which can lead to corruption of data in the directory.

**Transferring Operations Master Roles**

To transfer a role to a new domain controller, ensure that replication is up to date and functioning properly between the domain controller you are transferring the role from and the domain controller assuming the new role. This minimizes the time required to complete the role transfer. If replication is sufficiently out of date, the transfer process will take longer.

When an operations master role is transferred from one domain controller to another, the original role holder checks that all recent updates have been replicated to the new role holder. In the event that you must transfer an operations master role, it is best to transfer the role to a domain controller located in the same site as the original role holder. Changes made by the original role holder will replicate during the next replication cycle. Therefore, domain controllers located in the same site are likely to have the most recent information from the original operations master role holder.

If the new role holder is not located in the same site, changes might still need to replicate before the role transfer completes (this occurs during the role transfer process). If the new role holder is in the same site, it is more likely that the changes are already replicated, eliminating the need to replicate a large amount of changes during the role transfer process. This can significantly decrease the time required to transfer the role.

**Operational Attributes Used to Transfer Roles**

Not all attribute values are stored in a directory service. Instead, attribute values that are not contained in the directory can be calculated when a request for the attribute is made. This type of attribute is called an operational attribute. Note that this type of attribute is defined in the schema but it does not contain a value in the directory. Instead, the domain controller that processes a request for an operational attribute calculates the attribute's value to answer the client request.

The following operational attributes are used to transfer operations master roles and are located on the rootDSE (or root DSA-specific entry, the root of the AD DS tree for a given domain controller where specific information about the domain controller is kept):

- becomeRidMaster
- becomeSchemaMaster

- becomeDomainMaster
- becomePDC
- becomeInfrastructureMaster

During the operation of writing to the appropriate operational attribute on the domain controller that is receiving the operations master role, the role is removed from the old domain controller and is assigned to the new domain controller automatically. No manual intervention is required.

**Decommissioning a Role Holder**
When you use the Active Directory Domain Services Installation Wizard to decommission a domain controller that currently holds one or more operations master roles, the wizard reassigns the roles to a different domain controller. When the wizard is run, it determines whether the domain controller currently holds any operations master roles. If it detects any operations master roles, it queries the directory for other eligible domain controllers and transfers the roles to a new domain controller. A domain controller is eligible to hold a domainwide role if it is a member of the same domain. A domain controller is eligible to hold a forestwide role if it is a member of the same forest.

You cannot control which domain controller the wizard chooses and the wizard does not indicate which domain controller receives the roles. Because of this behavior, it is best to transfer the roles prior to decommissioning the role holder so that you have control over which new domain controller will hold the role.

**Operational Attributes Used to Decommission a Role Holder**
If you do not transfer operations master roles before demoting a domain controller, the **GiveAwayAllFsmoRoles** operational attribute is written, which triggers the domain controller to locate other domain controllers to transfer any roles it currently owns. AD DS determines which roles the domain controller being decommissioned currently owns and locates a suitable domain controller to assume these roles by following these rules:

- Locate a server in the same site
- Locate a server to which there is RPC connectivity
- Use a server over an asynchronous transport such as SMTP

In all role transfers, if the role is a domain-specific role, the role can be moved only to another domain controller in the same domain. Otherwise, any domain controller in the forest is a candidate.

**Seizing Operations Master Roles**
Role seizure is the act of assigning an operations master role to a new domain controller without the cooperation of the current role holder. During role seizure, a new domain controller assumes the operations master role without communicating with the current role holder. Because this can adversely affect your AD DS environment, seize operations master roles only as a last resort and if the current role owner is offline and unlikely to return to service.

Role seizure can create two conditions that can cause problems in the directory. First, the new role holder starts performing its duties based on the data located in its current directory partition. If replication did not complete prior to the time when the original role holder went offline, the new role holder might not receive changes that were made to the previous role holder. This can cause data loss or data inconsistency in the directory database.

To minimize the risk of losing data to incomplete replication, do not perform a role seizure until enough time has passed to complete at least one complete end-to-end replication cycle across your network. Allowing enough time for complete end-to-end replication ensures that the domain controller that assumes the role is as up-to-date as possible.

Second, the original role holder is not informed that it is no longer the operations master role holder, which is not a problem if the original role holder stays offline. However, if it comes back online (for example, if the hardware is repaired or the server is restored from a backup), it might try to perform the operations master role that it previously owned. This can result in two domain controllers performing the same operations master role simultaneously. Depending on the role that was seized, the severity of duplicate operations master roles varies from no visible effect to potential corruption of the AD DS database. Seize the operations master role to a domain controller that has the most recent updates from the current role holder to minimize the impact of the role seizure.

In Windows 2000 Server with Service Pack 3 (SP3) or later and Windows Server 2003 or later, if an operations master that has been taken offline is intentionally or unintentionally returned to service, it must successfully replicate inbound changes on the

directory partition that replicates and maintains that operations master state before resuming its previously held role. The directory partitions that maintain each operations master are defined in the following table.

**Operations Masters and Corresponding Directory Partitions**

| Operations Master Role | Directory Partition |
|---|---|
| Schema Master | Schema partition |
| Domain Naming Master | Configuration partition |
| Primary Domain Controller (PDC) Emulator | Domain partition for the operations master role owner's domain |
| RID Master | Domain partition for the operations master role owner's domain |
| Infrastructure Master | Domain partition for the operations master role owner's domain |

By waiting a full replication cycle, the domain controller can determine if another operations master exists before it brings itself back online. Successful replication must occur before dependent operations can be performed. When the previous role holder receives the current environmental state from another replica through inbound replication, it will update its copy of the database so that it no longer hosts the role in question. This is to prevent more than one domain controller from holding the same operations master role in each domain or forest.

Operations master roles that are hosted by a single domain controller in that role's domainwide or forestwide replication scope do not have to satisfy the initial synchronization requirement because the domain controller has no replication partners. Initial synchronization requirements only exist when a current operations master role owner's **hasMasterNC** attribute contains references to more than one domain controller. The **hasMasterNC** attribute is part of a domain controller's NTDS settings object in the CN=configuration partition of an operations master.

**Ramifications of Role Seizure**

If a role is seized, the new role holder is configured to host the operations master role with the assumption that you do not intend to return the previous role holder to service. Use role seizure only when the previous role holder is not available and you need the operations master role to keep the directory functioning. Because the previous role holder is not available during a seizure, you cannot reconfigure the previous role holder and inform it that another domain controller is now hosting the operations master role.

To reduce risk, perform a role seizure only if the missing operations master role unacceptably affects performance of the directory. Calculate the effect by comparing the impact of the missing service provided by the operations master to the amount of work that is needed to bring the previous role holder safely back online after you perform the role seizure.

AD DS continues to function when the operations master roles are not available. If the role holder is offline only for a short period, you might not need to seize the role to a new domain controller. Remember that returning an operation master to service after the role is seized can have dire consequences if it is not done properly. The following table shows the consequences of an unavailable operations master role and restoring a role holder after the role has been seized.

**Operations Master Role Functionality Risk Assessment**

| Operations Master Role | Consequences if Role is Unavailable | Risk of Improper Restoration | Recommendation for Returning to Service After Seizure |
|---|---|---|---|
| Schema master | You cannot make changes to the schema. | Conflicting changes can be introduced to the schema if both schema masters attempt to modify the schema at the same time. This can result in a fragmented schema. | Not recommended. Can lead to a corrupted forest. |
| Domain naming master | You cannot add or remove domains from the forest, add or remove application directory partitions, or perform domain rename operations. | You cannot add or remove domains or application directory partitions, or clean-up metadata. Domains and application directory partitions might appear as though they are still in the forest even though they are not. | Not recommended. Can lead to data corruption. |

| | | | |
|---|---|---|---|
| PDC emulator | You cannot change passwords on clients that do not have Active Directory client software installed. No replication to Windows NT 4.0 backup domain controllers. | Password validation can randomly pass or fail. Password changes take much longer to replicate throughout the domain. | Allowed. User authentication can be erratic for a time, but no permanent damage occurs. |
| Infrastructure master | Delays displaying updated group membership lists in the user interface when you move users from one group to another within a single domain. | Displays incorrect user names in group membership lists in the user interface after you move users from one group to another. | Allowed. May impact the performance of the domain controller hosting the role, but no damage occurs to the directory. |
| RID master | Eventually, domain controllers cannot create new directory objects as each of their individual RID pools is depleted. | Duplicate RID pools can be allocated to domain controllers, resulting in data corruption in the directory. This can lead to security risks and unauthorized access. | Not recommended. Can lead to data corruption. |

**Operational Attributes Used to Seize Roles**

When you seize an operations master role from an existing computer, the **fsmoRoleOwner** attribute is modified on the object that represents the root of the data directly, bypassing synchronization of the data and graceful transfer of the role.

The **fsmoRoleOwner** attribute of each of the following objects is written with the distinguished name of the NTDS Settings object (the data in AD DS that defines a computer as a domain controller) of the domain controller that is taking owner'ship of that role:

**Schema master**

LDAP://CN=Schema,CN=Configuration,DC=Contoso,DC=Com

**Domain naming master**

LDAP://CN=Partitions,CN=Configuration,DC=Contoso,DC=Com

**RID master**

LDAP://CN=Rid Manager$,CN=System,DC=Contoso,DC=Com

**PDC emulator**

LDAP://DC=Contoso,DC=Com

**Infrastructure master**

LDAP://CN=Infrastructure,DC=Contoso,DC=Com

As replication of this change starts to spread, other domain controllers learn of the operations master role change.

For example, if Server1 is the PDC emulator in the Contoso.com domain and is decommissioned and the administrator is unable to decommission the computer properly, Server2 needs to be forcibly assigned the PDC emulator role. After the role is seized, the value

CN=NTDS Settings,CN=Server2,CN=Servers,CN=Default-First-Site-Name,CN=Sites,CN=Configuration,DC=Contoso,DC=Com

is present on the LDAP://DC=Contoso,DC=com object.

**Network Ports Used by Operations Masters**

The following ports are used by all domain controllers.

### Port Assignments for Operations Masters

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | | 686 (Secure Sockets Layer [SSL]) |
| RPC/REPL | | 135 (endpoint mapper) |
| NetLogon | | 137 |

| | | |
|---|---|---|
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- What are Operations Masters?
- How the Data Store Works
- Active Directory Functional Levels Technical Reference

# Operations Masters Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Operations Masters Tools
- Network Ports Used by Operations Masters
- Related Information


**Operations Masters Tools**

Tools that are associated with operations masters are the same tools that you use to manage any domain controller.

The following tools have commands that are specific to managing operations masters.

**Domain.msc: Active Directory Domains and Trusts**

**Category**

An Active Directory Domain Services (AD DS) Administrative Tools Microsoft Management Console (MMC) snap-in that is automatically installed on all domain controllers running Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003. It is also available if you install the Active Directory Domain Services Tools that are part of the Remote Server Administration Tools (RSAT). For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

| **Note** |
| --- |
| In Windows Server® 2003 and Microsoft Windows® 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008, the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

**Version compatibility**

| Can Be Run From | Can Be Run Against |
| --- | --- |

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows 7 with Remote Server Administration Tools (RSAT) installed.<br>• Windows Vista with Remote Server Administration Tools (RSAT) installed.<br>• Windows XP Professional with Adminpak.msi installed[1] | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

1. To run the Windows Server 2003 Administration Tools Pack (Adminpak.msi) on a computer that is running Windows XP Professional, Windows XP Home Edition, or Windows XP 64-Bit Edition Version 2003, Windows XP Service Pack 1 (SP1) must be installed.

Active Directory Domains and Trusts provides a graphical interface in which you can view and manage all domains in the forest.

Use this tool to transfer the role of the domain naming master to another computer in the forest. You must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group, or you must have been delegated the appropriate authority to transfer this operations master role. For more information about the specific rights that are required to transfer specific operations master roles, see "What are Operations Masters?."

**Dsa.msc: Active Directory Users and Computers**

**Category**
An Active Directory Administrative Tools Microsoft Management Console (MMC) snap-in that is automatically installed on all domain controllers running Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003. It is also available if you install the Active Directory Domain Services Tools that are part of the Remote Server Administration Tools (RSAT). For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

**Version compatibility**

| Can Be Run From | Can Be Run Against |
| --- | --- |
| Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows 7 with Remote Server Administration Tools (RSAT) installed.<br>• Windows Vista with Remote Server Administration Tools (RSAT) installed.<br>• Windows XP Professional with Adminpak.msi installed[1] | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

1. To run the Windows Server 2003 Administration Tools Pack (Adminpak.msi) on a computer that is running Windows XP Professional, Windows XP Home Edition, or Windows XP 64-Bit Edition Version 2003, Windows XP Service Pack 1 (SP1) must be installed.

Active Directory Users and Computers provides a graphical user interface that can be used to manage users and computers in AD DS domains.

Use this tool to transfer the role of the RID master, the PDC emulator, and the Infrastructure master to another computer in the domain. You must be a member of the Domain Admins group or the Enterprise Admins group, or you must have been delegated the appropriate authority to transfer these operations master roles. For more information about the specific rights that are required to transfer specific operations master roles, see "What are Operations Masters?."

**schmmgmt.msc: Active Directory Schema**

**Category**
An Active Directory Administrative Tools Microsoft Management Console (MMC) snap-in that is automatically installed on all domain controllers running Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003. It is also available if you install the Active Directory Domain Services Tools that are part of the Remote Server Administration Tools (RSAT). For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

**Version compatibility**

| Can Be Run From | Can Be Run Against |
| --- | --- |
| Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows 7 with Remote Server Administration Tools (RSAT) installed.<br>• Windows Vista with Remote Server Administration Tools (RSAT) installed.<br>• Windows XP Professional with Adminpak.msi installed[1] | The schema master running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

1. To run the Windows Server 2003 Administration Tools Pack (Adminpak.msi) on a computer that is running Windows XP Professional, Windows XP Home Edition, or Windows XP 64-Bit Edition Version 2003, Windows XP Service Pack 1 (SP1) must be installed.

Active Directory Schema is a graphical user interface that can be used to manage AD DS object classes and their associated attributes.

Use this tool to transfer the role of the schema master to another computer in the forest. You must be a member of the Schema Admins group, or you must have been delegated the appropriate authority to transfer this operations master role. For more information about the specific rights that are required to transfer specific operations master roles, see "What are Operations Masters?."

Before the snap-in can be used, it must be registered so that it appears as an available snap-in for the Microsoft Management Console. For more information about registering Active Directory Schema snap-in, see "Active Directory Schema Technical Reference."

**Ntdsutil.exe: Ntdsutil**

**Category**
This tool ships with Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003. It is also available if you install the Active Directory Domain Services Tools that are part of the Remote Server Administration Tools (RSAT). For more

information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition |

You can use Ntdsutil to perform AD DS database maintenance; including transferring and seizing all operations master roles. This tool is intended for use by experienced administrators. For more information, see Ntdsutil (http://go.microsoft.com/fwlink/?LinkId=192808).

**Dsquery.exe: Dsquery**

**Category**
Dsquery is a command-line tool that ships with Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003. It is also available if you install the Active Directory Domain Services Tools that are part of the Remote Server Administration Tools (RSAT). For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

**Version Compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows 7 with Remote Server Administration Tools (RSAT) installed.<br>• Windows Vista with Remote Server Administration Tools (RSAT) installed.<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Use **Dsquery server -hasfsmo {schema | name | infr | pdc | rid}**to search for the domain controller that holds the requested operations master role. For more information, see Dsquery server (http://go.microsoft.com/fwlink/?LinkId=192809).

**Network Ports Used by Operations Masters**
The following ports are used by all domain controllers.

**Port Assignments for Operations Masters**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | | 636 (Secure Sockets Layer [SSL]) |
| RPC/REPL | | 135 (endpoint mapper) |
| Net Logon | | 137 |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

**Related Information**
The following resources contain additional information that is relevant to this section.

- Data Store Tools and Settings
- Active Directory Replication Tools and Settings

# Active Directory Replication Technologies

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

Active Directory is a distributed directory service that stores objects that represent real-world entities such as users, computers, services, and network resources. Objects in the directory are distributed among all domain controllers in a forest, and all domain controllers can be updated directly. Active Directory replication is the process by which the changes that originate on one domain controller are automatically transferred to other domain controllers that store the same data.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**Active Directory Replication Model**

The replication model comprises the mechanisms that support the multimaster update capabilities of Active Directory domain controllers. To ensure that replication data is transferred efficiently in the multimaster system, domain controllers track the changes that they have received and request only the updates that have occurred since the last replication. The update tracking is based on the state of the data as it exists on a replicating pair of domain controllers at the time of replication. Update tracking ensures that:

- Only changes that have not been received are replicated to a destination.
- Conflicts are resolved according to the last change that occurred, even when individual domain controller clocks are not synchronized or when administrators at different domain controllers make changes to the same object.

The replication model also accommodates multimaster updates by enabling replicated changes to be stored on destination domain controllers and forwarded to other domain controllers. This store-and-forward capability removes the need for every domain controller on which updates originate to contact every other domain controller that requires the updates.

**Active Directory Replication Topology**

The replication topology is the current set of Active Directory connections by which domain controllers in a forest communicate over local area networks (LANs) and wide area networks (WANs) to synchronize the directory partition replicas that they have in common. The replication topology ensures the transfer of changes to all directory partition replicas in the forest without redundancy. Replication topology generation is dynamic and adapts to network conditions and availability of domain controllers.

To ensure a consistent replication topology, domain controllers use global configuration data to arrive at the same view of domain controller data. They apply the same algorithm to this data to arrive at an identical replication topology. Operating independently, each domain controller contributes to a uniform and efficient replication topology.

Replication topology generation is optimized for speed within sites and for cost between sites. Replication between domain controllers in the same site occurs automatically in response to changes and does not require administrative management. Replication within a site is sent uncompressed to reduce processing time. Replication between domain controllers in different sites can be managed to control the scheduling and routing of replication over WAN links. Replication between sites is compressed so that it uses less bandwidth when sent across WAN links, thereby reducing the cost.

# Active Directory Replication Model Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Active Directory Replication Model Technical Reference**

The Active Directory directory service replication model ensures that Active Directory data on one domain controller is consistent with the replica of the same data on other domain controllers in the same domain. The Active Directory replication model determines how changes to Active Directory data are propagated and tracked automatically between domain controllers. The replication model allows directory data on each domain controller to be updated directly. Each domain controller maintains replication metadata that indicates the update status both for itself and relative to other domain controllers. In addition, the replication model allows each domain controller to request (pull) only changes that need to be replicated and to forward changes to other domain controllers that need them.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**In this subject**


- What Is the Active Directory Replication Model?
- How the Active Directory Replication Model Works
- Active Directory Replication Tools and Settings

# What Is the Active Directory Replication Model?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**What Is the Active Directory Replication Model?**

**In this section**


- Replication Model Components
- Technologies Related to Active Directory Replication
- Active Directory Replication Dependencies
- Related Information


Active Directory replication is the means by which changes to directory data are transferred between domain controllers in an Active Directory forest. The Active Directory replication model defines the mechanisms that allow directory updates to be transferred automatically between domain controllers to provide a seamless replication solution for the Active Directory distributed directory service.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**Note**

- This discussion of the replication model and related mechanisms for transferring directory data between domain controllers does not include the topic of replication topology. Replication topology is the set of connections that are generated by the Knowledge Consistency Checker (KCC) to enable replication to take place between domain controllers.

Active Directory is distributed by means of directory partitions. In addition to directory partitions that store forest-wide data, each domain controller stores a replica of a single domain directory partition, which contains data that is specific to one or more closely aligned business units—the users, computers, organizational units, and network resources that are managed by the same set of service and data administrators. Because each domain controller stores only one domain directory partition, Active Directory can scale to hundreds or thousands of domains storing millions of objects.

To efficiently synchronize data between domain controllers that store the same domain, Active Directory replication transfers updates according to directory partition. Each domain controller receives directory updates to the data that is stored in its domain only, as well as updates that are stored in the two directory partitions that store configuration and schema data for the forest.

**Note**

- In Windows Server 2003 forests, domain controllers can also store application directory partitions, which store application data that can be replicated to only the domain controllers that store the directory partition, irrespective of domain.

Active Directory replication manages the transfer of these updates to the appropriate domain controllers automatically, keeping domain data up-to-date among all domain controllers in the domain, regardless of location. In the process, all domain controllers in the forest are also updated with changes to forest-wide data.

**Replication Model Components**

To globally distribute the directory service, the Active Directory replication model incorporates the components in the following table.

**Replication Model Components and Advantages**

| Component | Description | Advantage |
|---|---|---|
| Multimaster replication | Every domain controller can receive originating updates to data for which it is authoritative, rather than having a single domain controller that receives all original updates (single-master replication, such as Windows NT 4.0 replication). | Provides fault tolerance, eliminating the dependency on a single domain controller to maintain directory operations. |
| Pull replication | Domain controllers request (pull) changes rather than send (push) changes that might not be needed. | Reduces unnecessary network traffic. |
| Store-and-forward replication | Each domain controller communicates with a subset of domain controllers to transfer replication changes, rather than one domain controller being responsible for communicating with every other domain controller that requires the change. | Balances the replication load among many domain controllers. |
| State-based replication | Each domain controller tracks the state of replication updates. | Conflicts and unnecessary replication are reduced. |

The Active Directory replication model ensures:

- **Domain controller availability.** Multimaster replication ensures that all domain controllers are available for updates, eliminating the potential for slow service if only a single updatable domain controller were available.

- **Efficient transfer of data.** State-based and pull replication ensures the minimum replication traffic and the maximum efficiency to retrieve only the changes that are needed.

- **Reliable consistency.** Directory consistency is guaranteed within the same period of replication latency.

- **Conflict resolution.** Even if two administrators change the same attribute on different domain controllers at the same time, conflict resolution ensures that only one of the values is replicated to all domain controllers.

**Replication Latency**

Multimaster replication involves *latency* — the period of time for an update that occurs on the originating domain controller to reach all other domain controllers that need it. To address replication latency, multimaster replication ensures *loose consistency* with *convergence*, as follows:

- *Loose consistency* means that the replicas are not guaranteed to be consistent with each other at any particular point in time because changes can originate from any replica at any time.

- *Convergence* means that if the system is allowed to reach a steady state in which no new updates are occurring and all previous updates have been completely replicated, all replicas of the same directory partition are guaranteed to converge on the same set of values.

With multimaster replication, it is not necessary for every domain controller to replicate with every other domain controller. Instead, the system implements a robust set of connections that determines which domain controllers replicate to which other domain controllers to ensure that networks are not overloaded with replication traffic and that replication latency is not so long that it inconveniences users. The set of connections through which changes are replicated to domain controllers in an enterprise is called the *replication topology*.

Although it involves latency, multimaster update capability provides high availability of write access to directory objects because several servers can contain writable copies of an object. Each domain controller in the domain can accept updates independently, without communicating with other domain controllers. Active Directory replication resolves any conflicts that occur when multiple updates are made to a single directory object.

**State-based Vs. Log-based Replication**

In state-based replication, each domain controller (master) in the multimaster system applies updates to its replica as they arrive, without maintaining a change log file. In a typical log-based replication system (also called “change-based”), each master keeps a log of the updates that it originated and communicates its log to every other replica. After a log has arrived at a replica, the replica applies the log, bringing itself more up-to-date. In this process, the destination receives and stores a record

of all changes, not just the changes it needs.

Active Directory replication relies on the current "state" (the current values of all objects) of the source replica instead of logs. The current state includes metadata that is used to resolve conflicts and to avoid sending the full replica on each replication cycle.

Generally speaking, a directory partition replica maintains all of its objects in a list ordered by last modification. This list is a log of sorts, but one whose size is a tiny fraction of the size of the replica itself. A typical replication request can be satisfied by examining only the last few objects on the list because the replication destination server is aware of how much of its replication source's list of changes have already been processed.

**Multimaster Vs. Single-master Replication**
Although a single-master model is adequate for a directory that has a small number of replicas and for an environment where all of the changes can be applied centrally, this approach does not scale beyond small organizations nor does it address the needs of decentralized organizations.

Multimaster replication provides the following advantages over single-master replication:

- If one domain controller becomes inoperable, other domain controllers can continue to update the directory. In single-master replication, if the master becomes inoperable, directory updates cannot take place. For example, if the failed server holds your password and your password has expired, you cannot reset your password and therefore you cannot log on to the domain.

- Servers that are capable of making changes to the directory can be distributed across the network and can be deployed in multiple locations.

- By creating multiple replicas of the directory and keeping the replicas consistent, the directory service can handle more queries per second. Directory services must handle a large number of queries compared to the number of updates they must process. A typical ratio of queries to updates is 99:1.

**Pull Vs. Push Replication**
In push replication, a source domain controller sends unsolicited information to update a destination domain controller. Push replication is problematic because it is difficult for the source to know what information the destination needs. The destination can receive the same information from another source. Therefore, a source can send unnecessary information to a destination.

**Technologies Related to Active Directory Replication**
File Replication service (FRS) is related to Active Directory replication because it requires the Active Directory replication topology. FRS is a multimaster replication service that is used to replicate files and folders in the System Volume (SYSVOL) shared folder on domain controllers and in Distributed File System (DFS) shared folders. FRS works by detecting changes to files and folders and then replicating the updated files and folders to other replica members, which are connected in a replication topology.

FRS uses the replication topology that is generated by the KCC to replicate the SYSVOL files to all domain controllers in the domain. SYSVOL files are required by all domain controllers for Active Directory to function. For more information about FRS and how it uses the Active Directory replication topology, see "FRS Technical Reference." For more information about SYSVOL, see "Data Store Technical Reference." For more information about DFS, see "DFS Technical Reference."

**Active Directory Replication Dependencies**
Active Directory replication has the following dependencies:

- **DNS**. The Domain Name System (DNS) that resolves DNS names to IP addresses. Active Directory requires that DNS is properly designed and deployed so that domain controllers can correctly resolve DNS names of replication partners.

- **Remote procedure call (RPC)**. Active Directory replication requires IP connectivity and the Remote Procedure Call (RPC) to transfer updates between replication partners.

- **Kerberos v5 authentication**. The authentication protocol for both authentication and encryption that is required for all Active Directory RPC replication.

- **LDAP protocol**. The primary access protocol for Active Directory. Replication of an entire replica of an Active Directory domain, as occurs when Active Directory is installed on an additional domain controller in an existing domain, uses LDAP communication rather than RPC.

The following diagram shows the interaction of these components within the replication process.

**Replication Interactions with Other Technologies**




**Related Information**

The following resources contain additional information that is relevant to this section.

- How the Active Directory Replication Model Works
- Active Directory Replication Topology Technical Reference
- DNS Support for Active Directory Technical Reference
- Data Store Technical Reference
- FRS Technical Reference
- DFS Technical Reference

# How the Active Directory Replication Model Works

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**How the Active Directory Replication Model Works**

**In this section**


- Active Directory Replication Model Architecture
- Active Directory Replication Model Physical Structure
- Active Directory Data Updates
- Domain Controller Notification of Changes
- Identifying and Locating Replication Partners
- Urgent Replication
- Network Ports Used by Active Directory Replication
- Related Information


Active Directory data takes the form of objects that have properties, or attributes. Each object is an instance of an object class, and object classes and their respective attributes are defined in the Active Directory schema. The values of the attributes define the object, and a change to a value of an attribute must be transferred from the domain controller on which it occurs to every other domain controller that stores a replica of that object.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

Thus, Active Directory replicates directory data updates at the attribute level. In addition, updates from the same directory partition are replicated as a unit to the corresponding replica on the destination domain controller over the same connection to optimize network usage.

The information in this section applies to organizations that are designing, deploying, or operating an Active Directory infrastructure that satisfies the following requirements:

- A Domain Name System (DNS) infrastructure is in place that manages the name resolution for domain controllers in the forest. Active Directory-integrated DNS is assumed, wherein DNS zone data is stored in Active Directory and is replicated to all domain controllers that are DNS servers.
- All Active Directory sites have local area network (LAN) connectivity.
- IP connectivity is available between all datacenter locations and branch sites.

The limits for data that can be replicated in one replication cycle are as follows:

- Values that can be transferred in one replication cycle (replication of the current set of updates between a source and destination domain controller): no limit.
- Values that can be transferred in one replication packet: approximately 100. Replication exchanges continue during the

course of one replication cycle until no values are left to send.

- Values that can be written in a single database transaction: 5,000. The effect of this limit depends on the forest functional level:

    - Windows 2000 forest functional level: The minimum unit of replication at this level is the entire attribute. Therefore, changes to any value in the linked, multivalued **member** attribute results in replicating the entire attribute. For this reason, the supported size of group membership is limited to 5,000.

    - Windows Server 2003 or Windows Server 2003 interim forest functional level: The minimum unit of replication is a single value of a linked, multivalued attribute. Therefore, the limitation on group membership is effectively removed.

This section covers the interactions that take place between individual domain controllers to synchronize directory data in an Active Directory forest.

**Active Directory Replication Model Architecture**
Active Directory replication operates within the directory service component of the security subsystem. The directory service component, Ntdsa.dll, is accessed through the Lightweight Directory Access Protocol (LDAP) network protocol and LDAP C application programming interface (API) for directory service updates, as implemented in Wldap32.dll. The updates are transported over Internet Protocol (IP) as packaged by the replication remote procedure call (RPC) protocol. Simple Mail Transfer Protocol (SMTP) can also be used to prepare non-domain updates for Transmission Control Protocol (TCP) transport over IP.

The Directory Replication System (DRS) client and server components interact to transfer and apply Active Directory updates between domain controllers.

When SMTP is used for the replication transport, Ismserv.exe on the source domain controller uses the Collaborative Data Object (CDO) library to build an SMTP file on disk with the replication data as the attached mail message. The message file is placed in a queue directory. When the mail is scheduled for transfer by the mail server application, the SMTP service (Smtpsvc) delivers the mail message to the destination domain controller over TCP/IP and places the file in the drop directory on the destination domain controller. Ismserv.exe applies the updates on the destination.

The following diagram shows the client-server architecture for replication clients and LDAP clients.

**Replication and LDAP Client-Server Architecture**




The following table describes the replication architecture components.

**Replication Architecture Components**

| Component | Description |
|---|---|
| Ntdsapi.dll | Manages communication with the directory service over RPC. |
| Private DRS client | Private version of Ntdsapi.dll that runs on domain controllers to make RPC calls for replication. |
| Wldap32.dll | Client library with APIs for access to directory service. |
| Asn.1 | Encodes and decodes LDAP requests for transport over TCP/IP or UDP/IP. |
| Drs.idl | Set of functions for replication (for example, Get-Changes) and maintenance (for example, Get replication status) |
| MAPI (address book) | Entry protocol for address book applications such as Microsoft Outlook. |
| Domain rename | Carries out domain rename instructions. |
| Ntdsa.dll | The directory service module, which supports the Windows Server 2003 and Windows 2000 replication protocol and LDAP, and manages partitions of data. |
| ISMServ.exe | Prepares replication data in e-mail message format for SMTP protocol transport. |
| CDO library | Used by Ismserv.exe to package replication data into a mail message. |
| Smtpsvc | SMTP service. |

**Note**

- If Windows NT 4.0 backup domain controllers (BDCs) are operating in the forest, Windows NT4 Net APIs provide an entry to the security accounts manager (SAM) on the primary domain controller (PDC) emulator.

The protocols that are used by Active Directory replication are described in the following table. RPC and SMTP are the replication transport protocols. LDAP is a directory access protocol, and IP is a network wire protocol.

**Active Directory Access and Replication Protocols**

| Protocol | Description |
|---|---|
| LDAP | The primary directory access protocol for Active Directory. Windows Server 2003 family, Windows XP, Windows 2000 Server family, and Windows 2000 Professional clients, as well as Windows 98, Windows 95, and Windows NT 4.0 clients that have the Active Directory client components installed, use LDAP v3 to connect to Active Directory. |
| IP | Routable protocol that is responsible for the addressing, routing, and fragmenting of packets by the sending node. IP is required for Active Directory replication. |
| Replication RPC | The Directory Replication Service (Drsuapi) RPC protocol, used in the enabling of administration and monitoring of Active Directory replication, to communicate replication status and topology and network topology from a client running administrative tools to a domain controller. RPC is required by Active Directory replication. |
| Replication Simple Mail Transfer Protocol (SMTP) | Replication protocol that can be used by Active Directory replication over IP network transport for message-based replication between sites only and for non-domain replication only. |

**Replication Subsystem**

Within the directory service component of the Active Directory architecture, the replication subsystem interacts with the operational layer to implement replication changes on the destination domain controller. The replication subsystem also determines the changes that a replication partner already has or those that are needed.

The database layer manages the database capability of the directory service. The extensible storage engine (Esent.dll) communicates directly with individual records in the directory data store.

The following diagram shows the replication subsystem components.

**Replication Subsystem Components**




The components of the replication subsystem are described in the following table.

**Replication Subsystem Components**

| Component | Description |
|---|---|
| Ntdsa.dll | Directory system agent (DSA), which provides the interfaces through which directory clients and other directory servers gain access to the directory database. |
| Replication | Directory Replication System (DRS) interface, which communicates with the database through RPC. |
| Operational layer | Performs low-level operations on the database without regard for protocol. |
| Database layer | API that resides within Ntdsa.dll and provides an interface between applications and the directory database to protect the database from direct interaction with applications. |

| | |
|---|---|
| Extensible storage engine (ESE) | Manages the tables of records, each with one or more columns that comprise the directory database. |
| Ntds.dit | The directory database file. |

**Active Directory Replication Model Physical Structure**

The Active Directory replication model components that determine how Active Directory replication functions between domain controllers are associated with mechanisms that effect automatic transfer of changes between replicating domain controllers, as described in the following table.

**Active Directory Replication Model Components and Related Mechanisms**

| Component | Description | Related Mechanisms |
|---|---|---|
| Multimaster replication | All domain controllers accept LDAP requests for changes to attributes of Active Directory objects for which they are authoritative, subject to security constraints that are in place. Each originating update is replicated to one or more other domain controllers, which record it as a replicated update. | LDAP update<br><br>Directory partitions<br><br>Change notification<br><br>Change tracking<br><br>Conflict resolution |
| Pull replication | When an update occurs on a domain controller, it notifies its replication partner. The partner domain controller responds by requesting (pulling) the changes from the source domain controller. | DNS name resolution<br><br>Kerberos authentication<br><br>Change tracking<br><br>Change notification<br><br>Change request |
| Store-and-forward replication | Domain controllers store changes received from replication partners and forward the changes to other domain controllers so that the originating domain controller for each change is not required to transfer changes to every other domain controller that requires the change. | Change tracking<br><br>Kerberos authentication<br><br>DNS name resolution<br><br>Change notification<br><br>Change request |

| | | |
|---|---|---|
| State-based replication | Active Directory replication is driven by the difference between the current "state" (the current values of all attributes) of the directory partition replica on the source and destination domain controllers. This state includes metadata that is used to resolve conflicts and to avoid sending the full replica on each replication cycle. | Change-tracking metadata:<br>• Update sequence number (USN) counter<br>• Up-to-dateness vector<br>• High-watermark |

These mechanisms are implemented by the replication system in a sequence of events that occurs between two domain controllers.

The following diagram shows a simplified version of the sequence between source and destination domain controllers when the source initiates replication by sending a change notification.

**Replication Sequence**




**Active Directory Data Updates**

When a change is made to an object in a directory partition, the value of the changed attribute or attributes must be updated on all domain controllers that store a replica of the same directory partition. Domain controllers communicate data updates automatically through Active Directory replication. Their communication about updates is always specific to a single directory partition at a time.

Active Directory data is logically partitioned so that all domain controllers in the forest do not store all objects in the directory. Active Directory objects are instances of schema-defined classes, which consist of named sets of attributes. Schema definitions determine whether an attribute can be administratively changed. Attributes that cannot be changed are never updated and therefore never replicated. However, most Active Directory objects have attribute values that can be updated.

Different categories of data are stored in replicas of different directory partitions, as follows:

- Domain data that is stored in domain directory partitions:
    - Every domain controller stores one writable domain directory partition.
    - A domain controller that is a global catalog server stores one writable domain directory partition and a partial, read-only replica of every other domain in the forest. Global catalog read-only replicas contain a partial set of attributes for every object in the domain.
- Configuration data: Every domain controller stores one writable configuration directory partition that stores forest-wide data controlling site and replication operations.
- Schema data: Every domain controller stores one writable schema partition that stores schema definitions for the forest. Although the schema directory partition is writable, schema updates are allowed on only the domain controller that holds the role of schema operations master.
- Application data: Domain controllers that are running Windows Server 2003 can store directory partitions that store application data. Application directory partition replicas can be replicated to any set of domain controllers in a forest, irrespective of domain.

**Changes to Attributes**

Active Directory updates originate on one domain controller (originating updates) and the same update is subsequently made on other domain controllers during the replication process (replicated updates).

Object update behavior is consistent and predictable: when a set of changes is made to a specific directory partition replica, those changes will be propagated to all other domain controllers that store replicas of the directory partition. How soon the changes are applied depends on the distance between the domain controllers and whether the change must be sent to other sites.

The following key points are central to understanding the behavior of Active Directory updates:

- Changes occur at the attribute level; only the changed attribute value is replicated, not the entire object.
- At the time of replication, only the current value of an attribute that has changed is replicated. If an attribute value has changed multiple times between replication cycles (for example, between scheduled occurrences of intersite replication), only the current value is replicated.
- The smallest change that can be replicated in Windows 2000 Active Directory is an entire attribute; even if the attribute is linked and multivalued, all values replicate as a single change. The smallest change that can be replicated in Windows Server 2003 Active Directory is a separate value in a multivalued attribute that is linked. This Windows Server 2003 feature is called linked-value replication.
- The individual values of a multivalued attribute are replicated separately under the following conditions:
    - The forest functional level is Windows Server 2003 interim or Windows Server 2003.
    - The attribute is linked. Linked attributes have the following characteristics:

      The attribute has distinguished name syntax.

      The attribute is marked as linked in the schema.
- An attribute is available for replication as soon as it is written.
- Replication is store-and-forward and moves sequentially through a set of connected domain controllers that host directory partition replicas.
- Multimaster conflict resolution is effective without depending on clock synchronization.

- Originating updates to a single object are written to the database in the same transaction, so partially written objects are not possible and a consistent view of the object is maintained.

- For replicated updates to large numbers of values in linked multivalued attributes, such as the **member** attribute of a group, updates are not always guaranteed to be applied in the same transaction. In this case, the updates are guaranteed to be applied in one or more subsequent transactions in the same replication cycle (all updates from one source are applied at the destination).

- After a replication cycle is initiated, all available changes to a directory partition on the source domain controller are sent to the destination domain controller, including changes that occur while the replication cycle is in progress.

**Effect of Schema Changes on Replication**

Attribute definitions are stored in **attributeSchema** objects in the schema directory partition. Changes to **attributeSchema** objects block other replication until the schema changes are performed. During replication of any directory partition other than the schema directory partition, the replication system first checks to see whether the schema versions of the source and the destination domain controllers are in agreement. If the versions are not the same, the replication of the other directory partition is rescheduled until the schema directory partition is synchronized.

Prior to upgrading a domain controller from Windows 2000 Server to Windows Server 2003, you must update the schema to be compatible with Windows Server 2003. When you run Adprep.exe, Windows Server 2003 schema is installed in the forest. This process upgrades the schema on each Windows 2000–based domain controller. Thereafter, you can begin upgrading domain controllers to Windows Server 2003.

### Note

- The Windows Server 2003 schema update adds 25 indexed attributes to the schema directory partition. An update of this size can cause replication delays in a large database. For this reason, domain controllers that are running Windows 2000 Server must be running at a minimum Windows 2000 Service Pack 2 (SP2) plus all additional Windows updates. However, it is highly recommended that you install Windows 2000 Service Pack 3 (SP3) on all domain controllers prior to preparing your infrastructure for upgrade to the Windows Server 2003 operating system.

**Effect of Raising the Forest Functional Level on Existing Linked, Multivalued Attributes**

Existing linked, multivalued attributes are not directly affected when you raise the forest functional level to enable linked-value replication. These attribute values are converted to replicate as single values only when they are modified. This design avoids the performance effects that would potentially result from rewriting the existing **member** attribute values of all group objects in the forest at the same time.

Because the **member** attribute is not converted until it is modified, a group that exceeded the 5,000-member limit in Windows 2000 continues to represent a replication issue because the original set of members continues to replicate as a unit under the new forest functional level. New members that are added and any member values that are updated replicate separately thereafter. Therefore, if the groups that were created in Windows 2000 do not exceed the 5,000-member limit, no replication issues are associated with the group.

Raising the forest functional level does affect the ability of an authoritative restore process to restore the back-links of a restored object that has one or more single-valued or multivalued linked attributes. The version of Ntdsutil that is included with Windows Server 2003 provides this functionality automatically when the forest has a functional level of Windows Server 2003 interim or Windows Server 2003. The version of Ntdsutil that is included with Windows Server 2003 SP1 also provides the ability to restore back-links that were created before the forest functional level was raised to Windows Server 2003 interim or Windows Server 2003. For more information about restoring back-linked attribute values, see "Authoritative Restore" later in this section.

For more information about how linked, multivalued attributes are stored, see "How the Data Store Works."

**Originating Updates: Initiating Changes**

As a Lightweight Directory Access Protocol (LDAP) directory service, Active Directory supports the following four types of update requests:

- Add an object to the directory.

- Modify (add, delete, or replace) attribute values of an object in the directory.

- Move an object by changing the name or parent of the object.

- Delete an object from the directory.

Each LDAP request generates a separate write transaction. An LDAP directory service processes each write request as an atomic transaction; that is, the transaction is either completed in full or not applied at all. The practical limit to the number of values that can be written in one LDAP transaction is approximately 5,000 values added, modified, or deleted at the same time.

A write request either commits and all its effects are durable, or it fails before completion and has no effect. A write request that commits is called an originating update. An originating update is initiated and committed at a specific replica. The absolute success or failure of an update applies even for requests that might affect several attributes of a single object, such as Add or Modify. In this case, if one attribute update fails, they all fail and the object is not updated.

An originating update enforces schema restrictions, including allowable parent object types and syntax for mandatory and optional attributes for an object. The restrictions are enforced according to the schema that exists on the domain controller at the moment of the update.

**Originating Add**
An Add request makes a new object with a unique **objectGUID** attribute. The values of all replicated attributes that are set by the Add request are stamped Version = 1.

The Add request fails immediately if the parent object does not exist or if the originating domain controller does not contain a writable replica of the parent object's directory partition.

**Originating Modify**
All Modify operations replace the current value of an attribute with a new value. A modify request can specify one of the following:

- That an attribute be deleted from the object. Attribute deletion is best thought of as replacing the attribute value with NULL. The NULL value occupies no storage of its own but does carry a stamp, as does any value that is stored as a directory attribute.

- That a value be added to the current value of an attribute, as when modifying an attribute that can have multiple values. The effect is to replace the current values with the current values plus the added value.

For each attribute in the request, a Modify request compares the new value in the request with the existing value. If the values are the same, the request to modify that attribute is ignored. If the resulting Modify request does not change any attributes of the object, the entire request is ignored.

Otherwise, a Modify request computes a stamp in the metadata for each new replicated attribute value by reading the version from the existing value (version=0 for an attribute that has never been written) and then adding 1 to this value. The Modify request replaces the old stamp values with new stamp values.

**Originating Move**
A Move request is essentially a special Modify request for a single attribute, the **name** attribute. The operation proceeds as described for the Modify request.

**Originating Delete**
A Delete request is essentially a special Modify request that does the following series of operations:

1. Sets the **isDeleted** attribute to TRUE, which marks the object as a tombstone (an object that has been deleted but not fully removed from the directory).

2. Changes the relative distinguished name of the object to a value that cannot be set by an LDAP application (a value that is impossible).

3. Strips all attributes that are not needed by Active Directory. A few important attributes (including **objectGUID**, **objectSid**, **distinguishedName**, **nTSecurityDescriptor**, and **uSNChanged**) are preserved on the tombstone.

   **Note**

- Because these attributes are preserved, tombstones can be restored (reanimated) by applications that use the LDAP API for undeleting an object.

- On domain controllers running Windows Server 2003 with SP1, the **sIDHistory** attribute is also retained on tombstone objects.

4. Moves the tombstone to the Deleted Objects container, which is a hidden container within those directory partitions that allow deletions.

On domain controllers running Windows Server 2003, you can restore tombstones in the Deleted Objects container to an active state.

**Configuration Objects Protected from Deletion**
Certain objects in the configuration directory partition are critical to the functioning of a domain controller. These objects and their child objects are protected from deletion.

**Reanimination of Protected Objects**
Each domain controller protects the following objects from deletion:

- The cross-reference (class **crossRef**) objects that represent the writable directory partitions that are stored on the domain controller.

- The RID object for the domain controller.

These objects are protected at that domain controller, as follows:

- The domain controller rejects an originating deletion of a protected object.

- The domain controller does not carry out a replicated deletion of a protected object. Instead, the threatened protected object is revived by updating its replication metadata as if each attribute had just been updated. The update is then replicated out, thereby reanimating the deleted object.

**Reanimation of the NTDS Settings Object**
The NTDS Settings (class **nTDSDSA**) object is also protected from deletion. On domain controllers running Windows 2000 Server or Windows Server 2003 with no service pack installed, the NTDS Settings object is reanimated and partner domain controllers continue to attempt to replicate with it. However, because the NTDS Settings object represents a domain controller in the replication topology, preserving it as a replication source when a domain controller has been removed from service is counterproductive and represents a security risk. For example, if a domain controller is demoted (that is, Active Directory is removed by running Dcpromo.exe), another server with the same name can be installed and mistaken by a replication partner for the demoted domain controller. To eliminate the possibility of improper replication attempts, domain controllers running Windows Server 2003 with SP1 disable the ability of the preserved NTDS Settings object to receive replication requests. Although the object is preserved on the domain controller that deleted it, replication attempts with the server that is represented by the deleted NTDS Settings object are discontinued. For information about how NTDS Settings replication metadata is preserved, see "Preservation of Replication Metadata" later in this section.

**Replicated Update Tracking by Domain Controllers**
A replicated update is performed on one domain controller as a result of receiving replication of an originating update that was performed at another domain controller. There is not necessarily a one-to-one correspondence between originating and replicated updates. A single replicated update might reflect a set of originating updates (even updates originating at different domain controllers) to the same object.

For example, the manager of a user object can be changed at one domain controller at the same time the address of the same user is changed at another domain controller. A third domain controller might receive these changes separately to the user object and in turn replicate the changes to a fourth domain controller in a single replicated update.

To avoid endless replication of the same update and reapplication of an update that is received from different replication partners, a domain controller must be able to recognize replicated updates that it has already received as opposed to those that it has not. Some directory services use timestamps to determine what changes need to be propagated, on the basis of preserving the last write. But keeping time closely synchronized in a large network is difficult. When the latest time of a directory write is the only means of determining which of two changes is recorded and replicated, skewed time on a domain

controller can result in data loss or directory corruption.

Active Directory replication does not primarily depend on time to determine what changes need to be propagated. Instead it uses update sequence numbers (USNs) that are assigned by a counter that is local to each domain controller. Because these USN counters are local, it is easy to ensure that they are reliable and never run backward (that is, they cannot decrease in value).

**Note**

- Replication uses Kerberos v 5 authentication protocol for security, which does require that the time services on domain controllers are synchronized.

When a conflict occurs, instead of using timestamps as the primary mechanism to determine what updates are preserved, Active Directory uses volatility (number of changes) as the first element of the per-attribute stamps that are compared during conflict resolution. The second element is a timestamp. Therefore, if an attribute is updated once on domain controller A and once on domain controller B, the last writer's update is preserved. But if the attribute is updated on domain controller A, then on domain controller B, and then again on domain controller A, the update of domain controller A is preserved even if the clock of domain controller B is set forward from that of domain controller A. With Active Directory, clock skew can never prevent a value from being overwritten.

Domain controllers use USNs to simplify recovery after a failure. When a domain controller is restored following a failure, it queries its replication partners for changes with USNs greater than the USN of the last change it received from each partner.

You do not need to be concerned with USNs and their implications unless you experience problems with replication that require troubleshooting.

**Server Object GUID (DSA GUID) and Server Database GUID (Invocation ID)**

The server object that represents a domain controller in the Sites container of the configuration directory partition has a globally unique identifier (GUID) that identifies it to the replication system as a domain controller. This GUID, called the DSA (Directory System Agent) GUID, is used in USNs to track originating updates. It is also used by domain controllers to locate replication partners. The DSA GUID is the GUID of the NTDS Settings object (class **nTDSDSA**), which is a child object of the server object. Its value is stored in the **objectGUID** attribute of the NTDS Settings object.

The DSA GUID is created when Active Directory is initially installed on the domain controller and destroyed only if Active Directory is removed from the domain controller. The DSA GUID ensures that the DSA remains recognizable when a domain controller is renamed. The DSA GUID is not affected by the Active Directory restore process.

The Active Directory database has its own GUID, which the DSA uses to identify the database instance (version of the database). The database GUID is stored in the **invocationId** attribute on the NTDS Settings object. Unlike the DSA GUID, which never changes for the lifetime of the domain controller, the invocation ID is changed during an Active Directory restore process to ensure replication consistency. For more information about replication following a restore process, see "Active Directory Replication on a Restored Domain Controller" later in this section.

On domain controllers that are running Windows Server 2003, the invocation ID also changes when an application directory partition is removed from or added to the domain controller.

**Determining Changes to Replicate: Update Sequence Numbers**

A source domain controller uses USNs to determine what changes have already been received by a destination domain controller that is requesting changes. The destination domain controller uses USNs to determine what changes it needs to request.

The current USN is a 64-bit counter that is maintained by each Active Directory domain controller as the **highestCommittedUsn** attribute on the rootDSE object. At the start of each update transaction (originating or replicated), the domain controller increments its current USN and associates this new value with the update request.

**Note**

- The rootDSE (DSA-specific Entry) represents the top of the logical namespace for one domain controller. RootDSE has no hierarchical name or schema class, but it does have a set of attributes that identify the contents of a given domain controller.

The current USN value is stored on an updated object as follows:

- Local USN: The USN for the update is stored in the metadata of each attribute that is changed by the update as the local

USN of that attribute (originating and replicated writes). As the name implies, this value is local to the domain controller where the change occurs.

- **uSNChanged**: The maximum local USN among all of an object's attributes is stored as the object's **uSNChanged** attribute (originating and replicated writes). The **uSNChanged** attribute is indexed, which allows objects to be enumerated efficiently in the order of their most recent attribute write.

  **Note**

  - When the forest functional level is Windows Server 2003 or Windows Server 2003 interim, discrete values of linked multivalued attributes can be updated individually. In this case, there is a **uSNChanged** associated with each link in addition to the **uSNChanged** associated with each object. Therefore, updates to individual values of linked multivalued attributes do not affect the local USN, only the **uSNChanged** attribute on the object.

- Originating USN: For an originating write only, the update's USN value is stored with each updated attribute as the originating USN of that attribute. Unlike the local USN and **uSNChanged**, the originating USN is replicated with the attribute's value.

**Tracking Object Creation, Replication, and Change**

The following series of diagrams illustrates the replication-related data for a single object and one of its attributes as it goes from creation through replication.

The following diagram shows the replication-related data for the user object when it is first created on domain controller DC1. Before the user object is created, the current USN for the domain controller is 4710. When the object is created, the local USN of 4711 is assigned to each attribute of the user object, and the current USN for the domain controller increments from 4710 to 4711. Because the object has not yet changed, the value of its **uSNChanged** attribute is the same as its **uSNCreated** attribute 4711.

### 1. Replication-related Data on DC1 When a User Object is Created

| Property | Value | Local USN | Version | Originating Time | Originating DC | Originating USN |
|---|---|---|---|---|---|---|
| cn | Jeff Smith | 4711 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| userPassword | 6Be8W5q- | 4711 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| sAMAccountName | JSmith | 4711 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| userPrincipalName | JSmith@contoso.com | 4711 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |

The next diagram shows the change to the destination domain controller when the new user object is replicated. The object is created as a replicated update on DC2. Notice that the per-attribute originating USN and stamp (version, originating time, originating DC) are replicated from DC1 to DC2, but the per-attribute local USN and per-object **uSNChanged** are unique to DC2.

### 2. Replication-related Data on DC2 When a New User Object is Replicated From DC1




| Property | Value | Local USN | Version | Originating Time | Originating DC | Originating USN |
|---|---|---|---|---|---|---|
| cn | Jeff Smith | 1746 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| userPassword | 6Be8W5q- | 1746 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| sAMAccountName | JSmith | 1746 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| userPrincipalName | JSmith@contoso.com | 1746 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |

The following information is transferred in the metadata of an updated attribute value from the source domain controller to the destination domain controller:

- The originating USN value for the updated attribute, which is the USN assigned by the domain controller on which the update was made.

- The stamp, which is used to resolve conflicts.

The next diagram illustrates the change in the replicated object on DC2 when someone changes the password (the **userPassword** property in the diagram) of the object on that domain controller. By this time, the current USN on DC2 has increased from 1746 to 2453. The update request changes the password and increments the current USN to 2453 on DC2. The request also sets the password attribute's originating USN and local USN to 2453 and creates a new stamp for the password value. The version number of this password's stamp is 2, which is one version number higher than the version of the previous password.

**3. Replication-related Data on DC2 After the User Password Value Has Been Changed on DC2**




| Property | Value | Local USN | Version | Originating Time | Originating DC | Originating USN |
|---|---|---|---|---|---|---|
| cn | Jeff Smith | 1746 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| userPassword | sEP3569?@2 | 2453 | 2 | 2003-09-10 11:53.29 | <DC2_GUID> | 2453 |
| sAMAccountName | JSmith | 1746 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| userPrincipalName | JSmith@contoso.com | 1746 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |

In the next diagram, the changed password is now replicated back to the original domain controller, whose current USN has increased to 5039. The replicated update increments the current USN of DC1 to 5040. The per-attribute originating USN and stamp (version, originating time, originating DC) are replicated from DC2 to DC1, and the per-attribute local USN and per-object **uSNChanged** values are set to 5040.

**4. Replication-related Data on DC1 After the Password Change Has Replicated to DC1**




| Property | Value | Local USN | Version | Originating Time | Originating DC | Originating USN |
|---|---|---|---|---|---|---|
| cn | Jeff Smith | 4711 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| userPassword | sEP3569?@2 | 5040 | 2 | 2003-09-10 11:53.29 | <DC2_GUID> | 2453 |
| sAMAccountName | JSmith | 4711 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |
| userPrincipalName | JSmith@contoso.com | 4711 | 1 | 2003-09-10 10:49.03 | <DC1_GUID> | 4711 |

**Replication Request Filtering**

Destination domain controllers use the originating USN to track changes they have received from other domain controllers with which they replicate. When requesting changes from a source domain controller, the destination informs the source of the updates it has already received so that the source never replicates changes that the destination does not need.

Two values are used by source and destination domain controllers to filter updates when the destination requests changes from the source replication partner:

- **Up-to-dateness vector**. The current status of the latest originating updates to occur on all domain controllers that store a replica of a specific directory partition.

- **High-watermark (direct up-to-dateness vector)**. The latest update (originating or replicated) to a specific directory partition that has been received by a destination from a specific source replication partner during the current replication cycle.

Both of these values specify the invocation ID of the source domain controller.

**Attributes to Send for Replication: Up-to-Dateness Vector**

The up-to-dateness vector is a value that the destination domain controller maintains for tracking the originating updates that are received from all source domain controllers. When a destination domain controller requests changes for a directory partition, it provides its up-to-dateness vector to the source domain controller. The source domain controller uses this value to reduce the set of attributes that it sends to the destination domain controller.

The up-to-dateness vector contains an entry for each domain controller that holds a full replica of the directory partition. The up-to-dateness vector values include the database GUID (invocation ID) of the source domain controller and the highest originating write (based on the USN) received from the source domain controller. If the up-to-dateness entry that corresponds to source domain controller X contains the USN *n*, the destination domain controller guarantees that it holds all updates to a specific directory partition that originated at domain controller X and that have an originating USN value of less than or equal

to *n*.

If the destination already has an up-to-date value, the source domain controller does not send that attribute. If the source has no attributes to send for an object, it sends no information at all about that object.

At the completion of a successful replication cycle between two replication partners, the source domain controller returns its up-to-dateness vector to the destination, including the highest originating USN on the source domain controller. The destination merges this information into its up-to-dateness vector. In this way, the destination tracks the latest originating update it has received from each partner, as well as the status of every other domain controller that stores a replica of the directory partition.

**Timestamp on Up-to-Dateness Vector in Windows Server 2003**
On domain controllers that are running Windows Server 2003, the up-to-dateness vector includes a timestamp that represents the last time the local (destination) domain controller has completed a full replication cycle with the source domain controller. The replication cycle may have occurred directly (direct replication partner) or indirectly (transitive replication partner). The timestamp is recorded whether or not the local domain controller actually received any changes from the partner.

By examining the timestamps, a domain controller can quickly identify other domain controllers that are not replicating. Warning messages are posted to the event log on each domain controller when non-replicating partners are discovered (Event ID 1864 in the Directory Service event log).

**Objects to Consider for Replication: High-Watermark (Direct Up-to-Dateness Vector)**
The high-watermark, or direct up-to-dateness vector, is a value that the destination domain controller maintains during replication to keep track of the most recent attribute change that it has received from a specific source domain controller for an object in a specific directory partition. When sending changes to a destination domain controller, the source domain controller provides the changes in increasing order of **uSNChanged**. Although the **uSNChanged** values from the source domain controller are not stored on objects at the destination domain controller, the destination domain controller keeps track of the **uSNChanged** value of the most recent object that was successfully updated from the source domain controller for a specific directory partition. This USN is called the destination's high-watermark with respect to the directory partition and the source domain controller.

When requesting changes during a replication cycle, the destination provides the high-watermark value with each request to the source domain controller, which in turn uses this value to filter the objects that it considers for continuing replication to the destination. If the **uSNChanged** value of an object on the source domain controller is less than or equal to the high-watermark value of the destination domain controller, the object update has already been received by the destination domain controller and is therefore not replicated. The high-watermark serves to decrease the CPU time and number of disk I/O operations that would otherwise be required.

The up-to-dateness vector and the high-watermark are complementary filter mechanisms that work together to decrease replication latency. Whereas the high-watermark prevents irrelevant objects from being considered by the source domain controller with respect to a single destination, the up-to-dateness vector helps the source domain controller to filter irrelevant attributes (and entire objects if all attributes are filtered) on the basis of the relationships between all sources of originating updates and a single destination.

For a specific directory partition, a domain controller maintains a high-watermark value for only those domain controllers from which it requests changes, but it maintains an up-to-dateness vector entry for every domain controller that has ever performed an originating update, which is typically every domain controller that holds a full replica of the directory partition.

**Multiple Paths Without Redundant Replication**
Multiple replication paths can exist between a pair of domain controllers. Multiple paths provide fault tolerance and can reduce latency. However, when multiple paths exist, you might expect the same change to be sent along each path to a specific domain controller or that a change might replicate in an endless loop. Active Directory prevents these potential problems with multiple paths by using the up-to-dateness vector. The ability to eliminate redundancy is called "propagation dampening."

The following is an example of how replication ordinarily occurs:

1. DC A updates a password attribute. In this example, the originating USN of the attribute is set to 3.

2. Destination DC B requests changes from source DC A and sends its high-watermark and up-to-dateness vector to DC A.

3. According to the high-watermark that was passed by DC B, source DC A examines one or more objects, one of which contains the changed password. When DC A encounters the changed password attribute, it proceeds as follows:

    a. First, DC A finds that the originating directory system agent (DSA) of the password attribute is DC A.

| Note |
| --- |
| The DSA is the server-side process that creates an instance of a directory service. The DSA provides access to the physical store of directory information located on a hard disk. |

b. Therefore, DC A reads the up-to-dateness vector supplied by DC B and finds that DC B is guaranteed to be up-to-date with updates that originated at DC A and that have an originating USN of less than or equal to 2.

c. DC A then finds that the originating USN of the password attribute is 3.

d. Because 3 is greater than 2, DC A sends the changed password attribute to DC B.

To illustrate propagation dampening, suppose that DC B had already received the password update from DC C, which had received it from DC A. In this case, the entry in the up-to-dateness vector of DC B for DC A would contain the USN value 3, not 2. Therefore, DC A would not send the changed password to DC B.

For information about viewing the replication metadata, see "Active Directory Replication Tools and Settings".

**Multimaster Conflict Resolution Policy**
Active Directory must ensure that all domain controllers agree on the value of the updated attribute after replication occurs. The general approach to resolving conflicts is to order all update operations (Add, Modify, Move, and Delete) by assigning a globally unique (per-object and per-attribute) stamp to the originating update. Thus each replicated attribute value (or multivalue) is stamped during the originating update and this stamp is replicated with the value.

**Conflict Resolution Stamp**
The stamp that is applied during an originating write has the following three components:

- The version is a number that is incremented for each originating write. The version of the first originating write is 1. The version of each successive originating write is increased by 1.
- The originating time is the time of the originating write, to a one-second resolution, according to the system clock of the domain controller that performed the write.
- The originating DC is the invocationID of the domain controller that performed the originating write.

When stamps are compared, the version is the most significant, followed by the originating time and then the originating DC. If two stamps have the same version, the originating time almost always breaks the tie. In the extremely rare event that the same attribute is updated on two different domain controllers during the same second, the originating DC breaks the tie in an arbitrary fashion.

Two different originating writes of a specific attribute of a particular object cannot assign the same stamp because each originating write advances the version at a specified originating domain controller. The originating time does not contribute to uniqueness. Replicated writes cannot decrease the version because values with smaller versions lose during conflict resolution. You can see all three components of the stamp in the output of the **repadmin /showobjmeta** command.

In the case of a conflict, the ordering of stamps allows a consistent resolution, as described in the following table.

**Replication Conflict Resolution**

| Replication Conflict | Description | Resolution |
| --- | --- | --- |
| Attribute value conflict | A Modify operation sets the value of an attribute. Concurrently, at another domain controller, a Modify operation sets the value of the same attribute to a different value. | The attribute value at all domain controllers is the value with the larger version number in its stamp. |

| | | |
|---|---|---|
| Add or Move under deleted parent, Delete non-leaf object | An Add or Move operation makes an object a child of parent object. Concurrently, at another domain controller, a Delete operation deletes the parent object. | At all domain controllers, the parent object is deleted and the child object is a child of the special LostAndFound container in the directory partition. Stamps are not involved in the resolution. |
| Relative distinguished name conflict | An Add or Move operation names a child object below a parent object. Concurrently, at another domain controller, an Add or Move operation names a different child of the same parent with the same child name, resulting in two child objects with identical relative distinguished name values below the same parent object. | The child object whose naming attribute has the larger version number in its stamp keeps its given name. The child object whose relative distinguished name attribute (for example, CN for most objects, OU for organizational units, DC for domain components) has the smaller version number in its stamp is named by the following convention: At all domain controllers, a system-assigned value that is unique to the conflicting name and cannot conflict with any client-assigned value is assigned to the child object. For example, if the relative distinguished name of a child object was "CN=ABC" before conflict resolution, its relative distinguished name after resolution is "CN=ABC*CNF:<*GUID*>," where "*" represents a reserved character, "CNF" is a constant that indicates a conflict resolution, and "<*GUID*>" represents a printable representation of the **objectGUID** attribute value. |

**Preservation of Replication Metadata**

On domain controllers running Windows Server 2003 with no service pack installed, replication metadata of the NTDS Settings object is maintained by default after Active Directory is removed from a domain controller. The mechanism for preserving replication metadata is disabled in Windows Server 2003 SP1.

**How Replication Metadata is Preserved in Windows Server 2003**

The period of time during which the replication metadata of the NTDS Settings object is maintained after Active Directory is removed from the respective domain controller is determined by an attribute of the Directory Service object (cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration,dc=*ForestRootDomainName*). This attribute, **replTopologyStayOfExecution**, has a default value of 14 days and a maximum value of half the tombstone lifetime.

For example, if you set the tombstone lifetime to 30 days and the **replTopologyStayOfExecution** value to 20 days, the actual stay-of-execution value is 15 days.

The root object of each directory partition replica has a multivalued attribute **repsFrom** that contains configuration and persistent state information associated with inbound replication from each source replica of that directory partition. A destination domain controller uses this information to poll source domain controllers for replication changes. The Knowledge Consistency Checker (KCC) is the replication topology component that manages creation of connection objects. When the KCC detects a **repsFrom** value with no corresponding connection object, as is the case when the NTDS Settings object has been deleted, the KCC checks the tombstone for the NTDS Settings object for the time it was deleted and compares that time to the interval in **replTopologyStayOfExecution**. If the NTDS Settings object has been deleted for a time that is equal to or greater than the value in **replTopologyStayOfExecution**, then the KCC removes the **repsFrom** value and replication is no longer attempted with the deleted server. For more information about the data that is stored in **repsFrom**, see "Destination Identification of Source Replication Partners" later in this section.

Prior to the stay-of-execution lifetime, you can see evidence of failed replication attempts in the Directory Service event log and in the output of Windows Support tools such as Repadmin.exe. When replication attempts are from a server designated as <unknown>, they are likely due to the stay-of-execution mechanism.

**Disabled replTopologyStayOfExecution in Windows Server 2003 SP1**

On domain controllers running Windows Server 2003 with SP1, the stay-of-execution mechanism is disabled so that replication attempts do not continue after a domain controller is demoted. Without the **repsFrom** metadata, destination domain controllers no longer poll for replication from the deleted source. Eliminating this mechanism ensures that computers attempting to maliciously spoof the deleted domain controller cannot respond to polling for replication changes. However, if you provide a value for **replTopologyStayOfExecution**, the stay-of-execution mechanism takes effect according to that value.

**Replication of Linked and Nonlinked Attributes**

Attributes are replicated differently depending on whether they are linked or nonlinked. Understanding the differences in how these attributes operate is helpful to understanding their effect on replication.

**Linked Attributes**

A linked attribute represents an interobject, distinguished-name reference. Linked attributes occur in pairs — forward link and backward link (back-link). The forward link is the linked attribute on the source object (for example, the **member** attribute on

the group object), while the backward link is the linked attribute on the target object (for example, the **memberOf** attribute on the user object). A back-link value on any object instance consists of the distinguished names of all the objects that have the object's distinguished name set in their corresponding forward link. For example, **manager** and **directReports** are a pair of linked attributes, where **manager** is the forward link and **directReports** is the back-link. If Bill is Joe's manager and the distinguished name of Bill's user object is stored in the **manager** attribute of Joe's user object, then the distinguished name of Joe's user object appears in the **directReports** attribute of Bill's user object.

A linked attribute can have either single or multiple values. For example, the **manager** attribute identifies the distinguished name of a single manager of the object or objects that are managed. The **directReports** attribute of a user object can have multiple values of user names.

The relationships between linked attributes are stored in a separate table in the directory database as link pairs. The matching pair of Link IDs (defined as any two numbers 2N and 2N+1) tie the attributes together. For example, the **member** attribute has a link ID of 2 and the **memberOf** attribute has a link ID of 3. Because the **member** and the **memberOf** attributes are linked in the database and indexed for searching, the directory can be examined for all records in which the link pair is **member**/**memberOf** and the **memberOf** attribute identifies the group (for example, "What user objects have group X as a value in their **memberOf** attribute?").

Attributes are marked in the schema as being linked. Attributes with the distinguished name syntax **Object(DS-DN)**, **Object(DN-String)**, or **Object(DN-Binary)** can be linked, but not all such attributes are linked.

**Nonlinked Attributes**

Nonlinked, distinguished-name attributes reference other objects in the same way as linked attributes do except that their behavior is different when a referred-to object is deleted, as described in "Replication of Deletion Updates" later in this section. In addition, nonlinked, multivalued attributes have an approximate limit of 1200 values (increased from the Windows 2000 Server limit of approximately 800 values). This limit is based on an approximate maximum page size of 8 kilobytes. For attributes of this maximum size, there are no storage or replication drawbacks or limitations. For more information about data size requirements, see "How the Data Store Works."

**USNs and Database Transactions**

Originating and local USNs effectively identify the database transactions that record originating and replicated changes. Local USNs represent a local database transaction. Therefore, local USNs that are the same represent the same database transaction. For example, when two updates to the same object originate on different systems, they are assigned different originating GUIDs and originating USN values. By definition, the local USN for each originating change is the same as the respective originating USN. Then, these changes replicate to a third domain controller independently. The originating USNs stay the same, but the local USNs now change to reflect the two separate transactions. When this set of changes replicates to a fourth domain controller, the two changes travel together. When they are applied at the fourth system, the originating USNs are still different, but the local USNs are now the same, reflecting their application in the same transaction.

Individual values of a multivalued attribute have slightly different rules. On origination, each unique value is assigned its own originating USN. The values replicate out independently, in roughly increasing USN order; however, atomicity between values is not guaranteed. When a set of values is applied on the destination domain controller, they are applied in batches to the same object. The batch size is approximately 100 multivalues. Therefore, when they are written on the destination domain controller, their originating USNs are still different, but the values on the same object, in groups of 100, have the same local USN.

**Atomicity of Linked and Nonlinked Attributes**

Atomicity is a guarantee by a database system that a grouping is applied in a single transaction. "Atomic" can be defined as "indivisible." Atomicity of a transaction means that the transaction occurs in total, or not at all. This guarantee ensures that objects are always in a complete state of either pre- or post-update. Nothing in between is possible.

As evidenced by the application of the local USN, write transactions are performed according to the way replication data is packaged and received. Rules that affect whether attribute values replicate together or separately differ for linked and nonlinked attributes.

- Nonlinked attributes: All changes to nonlinked attributes of an object replicate together and are assigned the same local USN. Therefore, they are committed to the database in the same write transaction.

- Linked attributes: Changes to linked attributes are not guaranteed to replicate together. Therefore, they are not guaranteed to be written in the same database transaction.

The issue for an update transaction, then, is at what points in the replication stream the application of an update can be interrupted and the effect of interruption on the end state. For inbound replication, the unit of transaction is the entire object for nonlinked attributes. For linked attributes, the unit of transaction is the batch of values for the same object. Therefore, it is

possible for a replicated update to a single object that has linked values to occur in more than one transaction. However, transactions to apply all updates to the same object are guaranteed to occur during the same replication cycle.

Atomicity provides the following assurances:

- Inferences can be based on the relationships of multiple attributes: If you always update attribute 1 and attribute 2 together and you see attribute 1, you know that you can make a decision based on attribute 2.

- Even if the system crashes in the middle of applying the replicated update, the database system guarantees that all changes are applied or that no changes are applied, and nothing in between.

**Replication Order of Linked and Nonlinked Attribute Updates**
To accommodate the guarantee that changes that are committed together on the originating domain controller appear in the same write transaction on the destination domain controller, updates to nonlinked attributes are prioritized more highly than updates to linked attributes. For this reason, although originating USNs are applied as they occur in order of the update and outbound replication changes are assembled in this update order, nonlinked updates are packaged for transmission ahead of linked updates. Therefore, when an object is updated and some of the updated attributes are linked and some are not linked, the replication packet orders the nonlinked values first, splitting the order of values between linked and nonlinked. For example, if two objects (obj1 and obj2) receive updates simultaneously to one linked attribute and one nonlinked attribute on each object, the order of the attribute values in the replication data packet are as follows:

obj1.nonlinked; obj2.nonlinked; obj1.linked; obj2.linked

Now, suppose that the number of nonlinked changes is sufficiently large that the values fill an entire replication packet. When both linked and nonlinked attributes are replicated, this scenario results in a set of changes that are made to one of the objects spanning multiple replication packets. All nonlinked values for a single object are included in one packet, but linked values for the same object can span multiple packets. Therefore, the write transaction at the destination can occur out of update order, and replicated updates to a single object can require multiple transactions.

**Transaction Order of Linked and Nonlinked Attribute Updates**
For linked attributes, although objects are transmitted in increasing USN order, object transactions might be applied out of order as a result of parent-child relationships. This condition guarantees that object creation precedes the application of any links to that object.

In addition, when the object already exists at the destination, although replication of nonlinked attributes occurs first, transaction of nonlinked attribute updates before linked attribute updates is not guaranteed. However, nonlinked updates are generally applied before linked updates at the destination.

For applications that depend on updates being written in update order, the prioritized implementation of replicated linked and nonlinked attribute updates can occasionally result in unexpected behavior.

**Group Membership and Linked-Value Replication**
The replication of linked, multivalued attributes is especially important for group objects. Potentially, the linked, multivalued **member** attribute of a group object can have thousands of values. Linked-value replication in Windows Server 2003 enables individual values to replicate separately. Linked-value replication requires a forest functional level of Windows Server 2003 or Windows Server 2003 interim. When it is in effect, linked-value replication solves the problem of replication delays that are due to the inability to write an entire **member** attribute in a single database transaction. Linked-value replication also makes restoring group membership back-links possible when a user or group object is authoritatively restored.

**Group Membership Replication in Windows 2000 Forests**
Linked-value replication is not available in Windows 2000 Server forests. Because an originating update must be written in a single database transaction, and because the practical limit for a single transaction is 5,000 values, membership of more than 5,000 values is not supported in Windows 2000 Server Active Directory. A group of this size represents a limitation both in terms of the database write operation that is required to record a change to an attribute of that size and the transfer of that much data over the network.

These conditions have the following impacts on replication, most notably for group and distribution list objects:

- Lost changes: If values of the same multivalued attribute are updated on two different domain controllers during a period of replication latency, the most recently changed replica of the attribute with all its multiple values is replicated and any earlier changes are lost. Changes to the separate values are not merged.

  **Note**

    - Because all changes to an object must be written in the same database transaction, multiple changes to a single group object can take a relatively long time to be written, which increases the likelihood of another change occurring to the same object prior to the completion of the original write.

- Excessive network bandwidth consumption: For example, when one member is added to a group of 3,000 members, the **member** attribute with all 3,001 values is transmitted between domain controllers. Transmission of all values to apply a change to only the updated value or values is inefficient use of network resources.

These limitations are effectively removed in a forest that has a forest functional level of Windows Server 2003 or Windows Server 2003 interim. At these levels, linked-value replication accommodates replication of individually updated **member** values.

**Group Membership Replication in Windows Server 2003 Forests**
In a Windows Server 2003 forest that has a forest functional level of Windows Server 2003 or Windows Server 2003 interim, linked-value replication provides the following benefits:

- Removed likelihood of losing entire sets of changes to the same group membership made on different domain controllers.

- Greatly reduced likelihood of update collisions, where the same **member** value is changed on different domain controllers at the same time and one update is lost.

- Improved network efficiency by transmitting only updated values and not the entire set of attribute values, which can include many thousands of values.

Although replication of many thousands of individual membership updates can be accommodated in a Windows Server 2003 forest, LDAP writes have a practical limit of approximately 5,000 updates in a single transaction. Because originating updates are required to complete in a single transaction, a practical limit of approximately 5,000 updates to a single object is recommended.

**Note**

- Only originating updates must be applied in the same database transaction. Replicated updates can be applied in more than one database transaction.

**Replication of Deletion Updates**
Object deletions are replicated by replicating tombstones. After an object is deleted but before it is removed from the directory, object references that formerly referred to the object now refer to the deleted object's tombstone. The **isDeleted** attribute, which has a value of TRUE when an object is a tombstone, indicates the object deletion to other domain controllers. Deleted objects are stored in the Deleted Objects hidden container. Every directory partition has a Deleted Objects container.

**Note**

- Objects are moved to the Deleted Objects container according to system flags. Certain protected configuration objects, such as NTDS Settings, do not move to the Deleted Objects container when they are deleted. In addition, dynamic objects that are deleted are removed from Active Directory according to a Time-to-Live (TTL) setting. They are not moved to the Deleted Objects container as tombstones. DNS resource records that are removed by scavenging are also not stored as tombstones when DNS is Active Directory–integrated. For more information about storage of object deletions and dynamic data, see "Storage and Removal of Object Deletions" and "Dynamic Data," respectively, in How the Data Store Works. For more information about how DNS data is removed from Active Directory, see "Understanding Aging and Scavenging" in How DNS Works.

By default, tombstones have a lifetime of 60 days (180 days in a forest that was created on a server running Windows Server 2003 with SP1), after which they are permanently removed from the directory database through a process called garbage collection.

The tombstone lifetime can be changed, but it is important to ensure that the tombstone lifetime is larger than the worst possible replication latency for any directory partition so that a tombstone cannot be deleted before it has replicated to every

directory partition replica. In addition, Active Directory does not allow data to be restored from a backup image that is older than the tombstone lifetime.

**Note**

- A tombstone is invisible to normal LDAP searches. However, a tombstone is visible to searches that use the special LDAP control 1.2.840.113556.1.4.417.

**Effect of Deletion on Linked Values**

When an object is deleted, two types of automatic cleanup occur with regard to linked values:

- All of the forward links that are held by the deleted object are removed; that is, if a group object is deleted, all of the **member** links on the group object are removed. The target of each link (for example, the user object that is named in the **member** attribute) is not affected by the removal of the group and its forward link, except that its **memberOf** back-link attribute loses the value that corresponds to the deleted group.

- All the back-links that are held by the deleted object are removed; that is, if a user is deleted, its back-links include all the group objects that have forward links to the user object. For example, if a user is a member of multiple groups and the user is deleted, the user's distinguished name value is removed from the **member** attributes of each group object that is named in the **memberOf** attribute of the deleted user. The group object is not changed in any way other than its membership.

During replication, if a linked value replicates to a domain controller, the linked value is silently dropped (that is, it disappears from the database with no history) under the following conditions:

- The forward link holder object that is referenced by the link value has been deleted.

- The backward link holder that is referenced by the link value (the target of the back-link) has been deleted.

When a target and source object are authoritatively restored together on a domain controller running Windows Server 2003, whether the links are restored on the destination domain controller depends on which restored object replicates out first:

- If the restored group object replicates to a destination domain controller before the restored user object, the group membership of the user is not restored because the user object does not exist on the destination.

- If the user object replicates out ahead of the group object and if the functional level of the forest was Windows Server 2003 or Windows Server 2003 interim when the back-link was created (that is, when the user was added to the group), the group membership of the user is restored in the **member** and **memberOf** attributes on the respective objects.

For more information about restoring back-links on authoritatively restored objects, see "Authoritative Restore" later in this section.

**Replication of Absent Linked Object References**

A different condition occurs when the object that the link refers to is simply removed as a referent, as it is when a user is removed from a group.

In a Windows 2000 forest, this condition results in replication of the entire linked attribute. In a forest with a functional level of Windows Server 2003 or Windows Server 2003 interim, this condition is replicated on a per-value basis. For example, if a user is removed from a group, the user value in the **member** attribute is marked “absent” in the database so that its condition can be replicated, much like a tombstone.

The deletion of the user object from the group replicates as an absent value in the **member** attribute of the group. After a tombstone lifetime, absent values are physically removed from the database. Until that time, the link value remains in the database. Such absent values can be restored prior to the end of the tombstone lifetime. If you add the user back to the group prior to the end of a tombstone lifetime, the link value is restored in the database by the value being marked "present."

**Deleting Group Objects in a Windows Server 2003 Forest**

In a forest that has a forest functional level of Windows Server 2003 or Windows Server 2003 interim, you can delete a group object of any size and its multivalued **member** attribute values are cleaned up in the background. For this reason, it is not

required for the deletion to complete in one transaction, as it is in a Windows 2000 forest. However, if you delete more than 5,000 members in the same transaction without deleting the group itself, the same long-running transaction conditions can occur as described in “Group Membership Replication in Windows 2000 Forests” earlier in this section.

Therefore, if you must add, modify, or delete more than 5,000 members of the same group, do so in blocks of less than 5,000 members to avoid long-running transactions and excessive network bandwidth consumption.

For more information about linked and nonlinked values, garbage collection, and tombstones, see "Storage and Removal of Object Deletions" in How the Data Store Works.

**Deletions on Nonreplicating Domain Controllers**

If a domain controller fails to replicate for a number of days that exceeds the tombstone lifetime, replicas of objects that have been deleted from a writable partition might remain in that domain controller’s directory. Because the tombstones of the deleted objects are permanently removed from the directory at the end of the tombstone lifetime, a domain controller that fails to replicate changes for tombstoned objects never deletes them.

This condition can occur for a variety of reasons, including:

- Prolonged misconfigurations.
- Prolonged errors in name resolution, authentication, or the replication engine that block inbound replication.
- Turning on a domain controller that has been offline for longer than the tombstone lifetime.
- Advancing system time or reducing tombstone lifetime values in an attempt to accelerate garbage collection before end-to-end replication has taken place for all directory partitions in the forest.

The condition of outdated objects can also occur due to hardware and software problems that render the domain controller unreachable. Regardless of the reason, a deleted object can remain on a domain controller any time the domain controller goes offline before receiving a deletion and remains offline for longer than the tombstone lifetime of that deletion.

These outdated objects, called lingering objects, create inconsistency in the directory. If a change is made to an outdated object on the reconnected domain controller, it is possible for the object to be recreated in the directory under certain conditions. To avoid this situation, replication of an outdated object is prohibited by default in newly created Windows Server 2003 forests.

**Backup Latency Interval**

On domain controllers running Windows Server 2003 with SP1, a new NTDS Replication event provides a warning to administrators when a domain controller has not been backed up. Event ID 2089 provides the backup status of each directory partition that a domain controller stores, including application directory partitions. Specifically, event ID 2089 is logged in the Directory Service event log when partitions in the Active Directory forest are not backed up within a backup latency interval. The value for the backup latency interval is stored as a **REG_DWORD** value in the **Backup Latency Threshold (days)** entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

By default, the value of **Backup Latency Threshold (days)** is half the value of the tombstone lifetime of the forest. If halfway through the tombstone lifetime a directory partition has not been backed up, event ID 2089 is logged in the Directory Service event log and continues daily until the directory partition is backed up.

This event serves as a warning to administrators and monitoring applications to make sure that domain controllers are backed up before the tombstone lifetime expires. However, it is recommended that you take backups at a much higher frequency than the default value of **Backup Latency Threshold (days)**.

**Replication Consistency Setting**

If the attributes on a lingering object never change, the object is never considered for replication. However, if an attribute changes, the attribute is considered for outbound replication. Because the destination domain controller does not hold the object for the attribute that is being replicated, an update cannot be performed. How this condition is resolved depends on the replication consistency setting on the domain controller.

A registry setting on domain controllers that are running Windows Server 2003 or Windows 2000 Server with SP3 provides a consistency value that determines whether a domain controller replicates and reanimates an updated object that has been deleted from all other replicas, or whether replication of such objects is blocked. The default settings are different on domain controllers that are running Windows 2000 Server with SP3 and Windows Server 2003.

**Strict Replication Consistency**

To avoid problems with reanimating objects that have been deleted, a domain controller that is running Windows Server 2003

in a newly created (not upgraded) Windows Server 2003 forest blocks inbound replication by default when it receives an update to an object that it does not have.

**Note**

- Active Directory replication uses update tracking to differentiate between replicating a newly created object and updating an attribute for an existing object. Replication of a lingering object is an attempt to update an attribute or attributes of an object that the destination domain controller cannot update because the object does not exist.

Replication is halted in the directory partition for the object until the lingering object is removed from the source domain controller or the strict replication consistency setting is disabled. For information about how lingering objects are removed, see "Lingering Object Removal" later in this section.

#### Loose Replication Consistency

When strict replication consistency is disabled, the effect is called "loose" consistency. By using "loose" consistency, the destination domain controller detects that it does not have the object for the attribute that is being replicated. The destination domain controller requests the entire object from the source partner, and thereby reanimates the object in its copy of the directory. The same process repeats on all domain controllers that do not have a copy of the object.

This mechanism can be used to cause lingering objects to be reanimated across the entire forest. If a lingering object is discovered and its presence is intended, then perform any update to the object. As long as replication consistency is set to loose (strict replication consistency is disabled) on all domain controllers, the object will be reanimated as it replicates around the forest.

Loose replication consistency is the default setting for domain controllers that are running Windows 2000 Server with SP3 or later. The Windows 2000 Server default is not changed by upgrading to Windows Server 2003; strict replication consistency remains disabled and replication is allowed to proceed. Keeping the Windows 2000 Server setting is required to ensure that the upgraded domain and forest are consistent with Windows 2000 Server functionality. You must change the setting manually following the upgrade.

#### Storage for Consistency Setting

The setting for replication consistency is in the registry on each domain controller.

The value for the consistency setting is stored in the **Strict Replication Consistency** entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**.

The values are as follows:

- Value: **1** (Set to **0** to disable)
- Default: **1** (enabled) in a new Windows Server 2003 forest, otherwise **0**.
- Data type: **REG_DWORD**

**Note**

- Having **Strict Replication Consistency** set to **0** or unset is equivalent to the Windows 2000 Server setting applied by the **Correct Missing Objects** registry entry. However, the semantics for **Correct Missing Objects** are the opposite of **Strict Replication Consistency**: **Correct Missing Objects**=1 is equivalent to **Strict Replication Consistency**=0 or unset.

#### Lingering Object Removal

On domain controllers that are running Windows Server 2003, you can use Repadmin to analyze and remove lingering objects from a domain controller that you suspect or know has not replicated for a tombstone lifetime or longer. If strict replication consistency is in effect and replication fails on the destination domain controller, event ID 1988 is logged in the Directory Service event log on the destination domain controller. Event ID 1988 indicates that the local domain controller has attempted to replicate an object from the source domain controller and the object is not present on the local domain controller because it might have been deleted and its tombstone already garbage collected. Event ID 1988 provides the GUID-based distinguished name of the source domain controller as well as the distinguished name and GUID of the outdated object. Replication of the directory partition containing the outdated object does not continue with the source domain controller until the situation has been resolved.

In this event, you can use an up-to-date domain controller as the authority against which to compare the objects on the source replication partner suspected of harboring lingering objects. This domain controller acts as the authoritative directory replica to reveal outdated objects in the suspect directory database on the destination.

In the Repadmin command-line arguments that remove lingering objects, the roles of source and destination are switched. The **repadmin /removelingeringobjects** command compares the directories of two domain controllers, as follows:

- A "source" domain controller that you designate as the authoritative reference server.
- A "destination" domain controller, which is the source replication partner that has attempted to replicate an outdated object.

To use the **repadmin /removelingeringobjects** command, both source and destination domain controllers must be running Windows Server 2003.

If the comparison reveals any mismatched objects, Repadmin either displays or removes the objects that are found on the destination but not on the source, depending on the arguments that you use. The advisory mode argument allows you to view the results of the command before you take action to remove any objects from the directory.

**Repadmin /RemoveLingeringObjects Command Syntax**

The command **repadmin /removelingeringobjects** has the following syntax:

```
                /removelingeringobjects <Dest_DC_LIST> <Source DC GUID> <NC>
[/ADVISORY_MODE]
```

where

- <Dest_DC_List> is the DNS or NetBIOS name of one or more domain controllers that you suspect of harboring lingering objects. <DC_List> provides the ability to target specific domain controllers, such as all domain controllers in a site, all global catalog servers, or domain controllers that hold specific operations master roles. To see the syntax for DC_List, type **repadmin /listhelp** at the command prompt.
- <Source DC GUID> is the GUID you obtained by running **repadmin /showrepl** against the source domain controller that you are using as the authoritative server.
- <NC> is the distinguished name of the directory partition that contains the lingering object.
- [/ADVISORY_MODE] is an optional switch that specifies that no deletions are performed on the destination domain controller, but are displayed (logged) only. Using this switch is recommended prior to allowing Repadmin to remove any objects.

To use this command, you must first obtain the GUID of the authoritative source domain controller by running **repadmin /showrepl <source_server>**, where <source_server> is the name of the domain controller that has a writable copy of the directory partition that will serve as the authoritative replica. The output of this command provides the "DC GUID," which the **/removelingeringobjects** command requires to identify the authoritative source.

**RemoveLingeringObjects Implementation**

When you run **repadmin /removelingeringobjects**, the tool performs the following steps to compare the directories of the source and destination domain controllers and log (or remove) any found lingering objects:

1. Check to ensure that the directory partition and the source domain controller are valid.
2. Verify that the user has the DS-Replication-Manage-Topology extended right on the directory partition container object specified in <NC>. This extended right is required to verify object state between two domain controllers. Members of the Domain Admins group have this right by default.
3. Ensure that both source and destination use the same objects for comparison by merging the up-to-dateness vectors to filter out any objects that have not replicated from the source to the destination or from the destination to the source. This check rules out a lingering object on the destination if the destination has not received the tombstone from the

source, and vice versa. Any such nonreplicated objects are removed from the comparison.

4. Create the list of object GUIDs for each domain controller to be compared. Examine the metadata of each object and use the merged up-to-dateness vector to determine whether the object should be present on both source and destination.

5. For each GUID that is in the list for the destination, determine if it is in the list of GUIDs for the source.

6. If a GUID is not found on the source, the object is identified to be outdated on the destination and is either displayed or deleted on the destination server. If advisory mode has been specified, the GUID is displayed only.

For more information about up-to-dateness vectors, see "Replicated Update Tracking by Domain Controllers" later in this section. For more information about removing lingering objects, see "Troubleshooting Active Directory Replication Problems" in the Active Directory Operations Guide.

**Active Directory Replication on a Restored Domain Controller**
All domain controllers must be backed up routinely to ensure directory integrity. In case of failure on a domain controller, the backup media can be used to restore the domain controller to its state at the time of the backup. When a domain controller is restored from a backup, it can then be brought up-to-date by normal replication.

There are two general methods for restoring Active Directory from backup media, each of which has different replication consequences:

- Nonauthoritative restore. Replication brings the domain controller up-to-date from its state at the time of backup, including updating deletions that have occurred since the time of the backup.

- Authoritative restore. Objects that were deleted can be reinstated.

**Nonauthoritative Restore**
Nonauthoritative restore is the default method of performing a restore of Active Directory, and it is used in the majority of restore situations, such as domain controller hard disk failure. Nonauthoritative restore is performed by using the backup tool that you used to create the backup file. A nonauthoritative restore returns Active Directory on the domain controller to a state that is consistent with the time that the backup was taken. When the domain controller restarts following the restore process, it requests changes from its replication partners. Through the normal replication process, the restored domain controller receives any directory changes that have occurred since the time of the backup.

Because the restore process does not restore any previously deleted data to Active Directory, it is described as nonauthoritative.

The following example illustrates the process of replication between a set of domain controllers before and after nonauthoritative restore, including the metadata changes that determine the effect of the restore operation. In this example, three domain controllers (DC1, DC2, and DC3) exist as part of the Active Directory replication topology. The scenario includes the following steps:

**Step 1**
Prior to backup, DC1 maintains replication information about itself and its replication partners:

- Its own invocation ID (database GUID)

- Highest USN that it has committed for changes to the local database.

- Replication metadata for its two replication partners, DC2 and DC3:

    - Invocation ID of the local Active Directory database of each domain controller.

    - The up-to-dateness vector, as shown in the table below.

**Step 2**
Some changes have been made to the Active Directory database on DC1 since Step 1, which result in advancing the current USN on DC1 to 35. At this point, DC1 is backed up.

**Step 3**
A total of 14 additional change changes have been made to the local Active Directory database on DC1, advancing the highest committed USN on DC1 to 49. DC1 now begins replicating these changes to DC2. After replicating 12 of the 14 changes to DC2, DC1 becomes unavailable due to hard drive corruption. As a result, DC2 has USN 47 as the last change it received from DC1. At this point, the up-to-dateness vector on DC2 shows a highest committed USN of 47 for DC1.

**Step 4**
A nonauthoritative restore has been completed and DC1 is restarted. During the restore process a new invocation ID has been assigned to the local Active Directory database on DC1. For replication purposes, DC1 adds its old invocation ID (GUID1) and highest committed USN of 35 to its up-to-dateness vector, effectively restoring it to its state as of Step 2.

**Step 5**
DC1 replicates in the 12 changes that were saved to DC2 before the restore (which corresponds to Step 3 in the following table). The last two changes that occurred as originating updates on DC1 prior to going offline are lost.

**DC1 Replication Metadata Values Pre- and Post-Nonauthoritative Restore**

| | **Step 1: Prior to backup** | **Step2: Immediately post-backup** | **Step 3: Offline status before restore** | **Step 4: Immediately post-Restore** | **Step 5: After post-restore replication** |
|---|---|---|---|---|---|
| **Invocation ID of DC1** | DC1-GUID1 | DC1-GUID1 | DC1-GUID1 | DC1-GUID4 | DC1-GUID4 |
| **Highest committed USN on DC1** | 20 | 35 | 49 | 35 | 47 |
| **Up-to-dateness vector on DC1** | DC2-GUID2, USN=10;<br>DC3-GUID3, USN=45 | DC2-GUID2, USN=25;<br>DC3-GUID3, USN=60 | DC2-GUID2, USN=37;<br>DC3-GUID3, USN=60 | DC2-GUID2, USN=25;<br>DC3-GUID3, USN=60;<br>DC1-GUID1, USN=35 | DC2-GUID2, USN=37;<br>DC3-GUID3, USN=60;<br>DC1-GUID1, USN=47 |

**Authoritative Restore**
The primary purpose of an authoritative restore is to reinstate objects that were deleted from Active Directory. To reinstate objects that were intentionally or accidentally deleted, a nonauthoritative restore must be completed and followed by an authoritative restore. Authoritative restore is performed by using Ntdsutil.exe.

The nonauthoritative restore process cannot reinstate deleted objects from a backup image because the backup media that was used to restore the domain controller contains an image of Active Directory that was created before the objects were deleted. In this case, the deletions would simply be replicated in from the up-to-date replication partner and applied by the restored domain controller. To reinstate a deleted object, an authoritative restore is required.

The authoritative restore process works as follows:

1. The domain controller is restarted in Directory Services Restore Mode, and a nonauthoritative restore of Active Directory is performed by using backup media that was created before the object was deleted.

2. Following the nonauthoritative restore but prior to restarting, the object metadata is altered by Ntdsutil.exe so that it has a higher USN than any other possible version of the object (by default, the version number is increased by 100,000). The effect is to render the object or objects as authoritative and reinstates them in Active Directory.

The authoritative restore process results in a merge of authoritatively restored objects and attributes and existing objects and attributes. The authoritative restore process does not affect objects that were created after the backup was created. This merge functionality applies to objects with linked attributes and nonlinked attributes alike. For example, if you are restoring an object that has attribute A and attribute B in the backup version and has attributes A′, B′, and C in the current directory, attribute C is retained after authoritative restore. Therefore, a group object that has the **member** value of User1 in the backup and has both User1 and User2 in the current directory includes both of those memberships after authoritative restore of the group object. Any post-backup **memberOf** or **member** attribute values that were added to a user or group, respectively, are not affected by replication updates after the restore procedure.

**Note**

- Only the domain and configuration directory partitions can be marked as authoritative. The schema directory partition cannot be authoritatively restored.

**Restoring Back-links for Authoritatively Restored Objects**

When authoritative restore is performed on domain controllers running Windows Server 2000, the procedure recovers objects that have been deleted, but it does not restore the back-links for any objects that have linked attributes. The effect of not restoring back-links for the restored objects is particularly problematic for group memberships, which must be restored manually.

In a forest that has a forest functional level of Windows Server 2003 or Windows Server 2003 interim, the procedure for performing authoritative restore automatically restores back-links for multivalued, linked attributes. For example, the **member** attribute of groups to which a restored user object belongs are updated. This restoration applies to only those links that were created after the functional level was raised. For example, if you added a user to a group before raising the forest functional level, the user's membership in that group is not restored if you delete the user and then authoritatively restore the user. Automatic restore of back-links requires the raised forest functional level because link restoration is made possible by linked-value replication.

**Restoring Back-links Created Before Linked-Value Replication**

An updated version of Ntdsutil that is included with Windows Server 2003 SP1 makes it possible to also restore back-links that were created before implementation of linked-value replication. On domain controllers that have this updated version of Ntdsutil, the **authoritative restore** option generates an LDAP Data Interchange Format (LDIF) file that can be used to restore any back-links that are not restored automatically. In addition, Ntdsutil generates a text file that you can use to create an LDIF file for restoring back-links for groups in other domains. The LDIF file can be used to restore back-links on domain controllers running Windows 2000 Server or Windows Server 2003, and it does not depend on forest functional level. This method also resolves the problem of links not being restored if linked user and group objects are authoritatively restored together and the restored group object replicates out before the restored user object. For more information about restoring back-links, see "Managing Active Directory Backup and Restore" in the Active Directory Operations Guide.

**Domain Controller Notification of Changes**

Replication within a site occurs as a response to changes. On its NTDS Settings object, the source domain controller stores a **repsTo** attribute that lists all servers in the same site that pull replication from it. When a change occurs on a source domain controller, it notifies its destination replication partner, prompting the destination domain controller to request the changes from the source domain controller. The source domain controller either responds to the change request with a replication operation or places the request in a queue if requests are already pending. Replication occurs one request at a time until all requests in the queue are processed.

| **Note** |
| --- |
| Replication between sites occurs according to a schedule, where the destination requests changes at the specified time. By default, change notification is not enabled on site links. |

**Notification Delay Values**

When a change occurs on a domain controller within a site, two configurable intervals determine the delay between the change and subsequent events:

- **Initial notification delay**. The delay between the change to an attribute and notification of the first partner. This interval serves to stagger network traffic caused by intrasite replication. When a domain controller makes a change (originating or replicated) to a directory partition, it starts the timer for the interval; when the timer expires, the domain controller begins to notify its replication partners (for that directory partition and within its site) that it has changes.

- **Subsequent notification delay**. The delay between notification of the first partner and notification of each subsequent partner. A domain controller does not notify all of its replication partners at one time. By delaying between notifications, the domain controller distributes the load of responding to replication requests from its partners.

The storage location and default values for the initial and subsequent notification delay intervals vary depending on the version of the operating system (Windows Server 2003 or Windows 2000 Server) that is running on the domain controller and whether the domain controller has been upgraded. Application of notification delay values is also potentially affected by their default status and by the forest functional level.

**Storage of Intrasite Notification Delay Values**
On a domain controller that is running Windows Server 2003, intrasite notification delay values are specific to each directory partition and are stored in two attributes of the cross-reference object for each directory partition.

**Windows Server 2003 Storage of Notification Delay Values**
On domain controllers that are running Windows Server 2003, the attributes that store the change notification values are located on each cross-reference object in the Partitions container within the configuration directory partition:

- The value for initial change notification delay is stored in the **msDS-Replication-Notify-First-DSA-Delay** attribute.

  The default value is 15 seconds.

- The value for subsequent notification delay is stored in the **msDS-Replication-Notify-Subsequent-DSA-Delay** attribute.

  The default value is 3 seconds.

These attributes do not exist in Windows 2000 Server.

Although the attribute values are present on all domain controllers that are running Windows Server 2003, the default values of 15 seconds for initial notification delay and 3 seconds for subsequent notification delay are in effect only under the conditions described in "Application of Notification Delay Values by Domain Controllers" later in this section.

**Windows 2000 Server Storage of Notification Delay Values**
On domain controllers that are running Windows 2000 Server, intrasite notification delay values affect all directory partitions and are stored in registry entries on each domain controller. The registry entries are as follows:

- The value for the delay before the first change notification is stored in the **Replicator notify pause after modify (secs)** entry in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters.

  The default value is 300 seconds.

- The value for the delay before each subsequent change notification is stored in the **Replicator notify pause between DSAs (secs)** entry in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters.

  The default value is 30 seconds.

**Effect of Upgrade and Forest Functional Level on Registry Notification Delay Values**
Notification delay values for first and subsequent change notification are preserved as registry settings during upgrades from Windows 2000 Server to Windows Server 2003.

Raising the forest functional level to Windows Server 2003 Interim or Windows Server 2003 affects notification delay values for first and subsequent change notification, as follows:

- If a registry entry has the Windows 2000 Server default value of 300 seconds for initial notification delay or 30 seconds for subsequent notification delay at the time the forest functional level is raised, the registry entry that has the default value is removed and the Windows Server 2003 default value takes effect on the respective attribute of each cross-reference object.

- If a registry entry has a nondefault value, the registry value is preserved.

| **Note** |
| --- |
| One or both notification delay registry entries can be added to the registry to override the cross-reference values, if needed, to control notification by a specific domain controller running Windows Server 2003. |

**Application of Notification Delay Values by Domain Controllers**
To accommodate both the registry and cross-reference object locations of notification delay information, the process that is used to determine which change notification delay values to apply favors the settings in the registry if any exist. When replication partners send notification of changes, notification delay values are checked according to the operating systems that

are running on the partners, as follows:

- Windows Server 2003 only:
    - Check the registry for the presence of initial and subsequent notification delay values and use those values if they exist.
    - If no registry values exist, check the cross-reference object for the directory partition to which the change has occurred. If values are set, use those values.
    - Otherwise, use the default values of 15 seconds for initial notification delay and 3 seconds for subsequent notification delay.
- Windows Server 2003 and Windows 2000 Server:
    - Check the registry for the presence of registry values. If a value is set for an entry, use the value that is set for all relevant directory partitions.
    - If no value is set on a registry entry (the default is in effect), use the respective default values of 300 seconds for initial notification delay and 30 seconds for subsequent notification delay.

**Identifying and Locating Replication Partners**
Replication occurs in one direction between two domain controllers at a time. The replication topology determines the replication partnerships between source and destination domain controllers. As a replication source, the domain controller must determine the replication partners it must notify when changes occur. As a replication destination, the domain controller participates in replication either by responding to notification of changes from a source, or by requesting changes to initiate replication when it starts up or in response to a schedule.

**Destination Identification of Source Replication Partners**
When the KCC creates the replication topology, it creates connection objects on destination domain controllers that represent the inbound connection from the replication source domain controller. For each source domain controller that is represented by an inbound connection object, the KCC writes information to the **repsFrom** attribute of the directory partition object for each directory partition that the destination domain controller has in common with the source domain controller. This information is local to the destination domain controller and is not replicated. For a destination domain controller, the **repsFrom** attribute contains the following information about each source domain controller:

- Directory partition name.
- NTDS Settings object distinguished name.
- GUID-based DNS name (CNAME alias).
- Whether replication is intrasite, intersite using IP (RPC), or intersite using SMTP.
- Flags that govern the options, which indicate whether notification is allowed (no value is present) or not allowed (NEVER_NOTIFY is present).
- Invocation ID, for the purpose of detecting a change due to a restore from backup.
- Direct USN vector (high-watermark), which is the USN pair of object USN (OU) and property USN (PU).
- Time of last attempt and time of last success.
- Last error code.
- Count of consecutive failures, if any.

If notification is enabled on the destination domain controller, the destination responds to notification messages from source domain controllers by sending a request for changes. If notification is not enabled, the destination responds to notification requests with an error.

Notification is enabled on a destination domain controller except under the following conditions, in which the KCC sets the option NEVER_NOTIFY in **repsFrom**:

- The connection object that the KCC creates is inbound from a server in a different site. Such a destination is designated as a bridgehead server for the directory partition. Bridgehead servers pull replication according to a schedule on the site link object.

- The connection object is created to identify the source from which this domain controller replicates to add Active Directory as a new domain controller.

#### Source Identification of Destination Replication Partners

The source domain controller keeps track of the replication partners that pull changes from it and uses the information to locate partners for change notification. This information is not provided by the KCC, but rather by the source domain controller itself during a replication cycle. The first time a domain controller receives a request for changes from a new destination, the source creates an entry for the destination in the **repsTo** attribute on the respective directory partition object.

Whenever the source has changes, it sends a notification to all replication partners that are identified in the **repsTo** value for the respective directory partition. Like the **repsFrom** data, this information is stored locally on the domain controller and is not replicated. When updates occur, the source domain controller checks the **repsTo** attribute to determine the identities of its destination replication partners. The source domain controller notifies them one by one that changes are available.

### Note

- The output of the command **repadmin /showrepl /v /all** shows the contents of the **repsFrom** value (INBOUND NEIGHBORS) and the **repsTo** value (OUTBOUND NEIGHBORS).

If a destination domain controller has the NEVER_NOTIFY option set, the destination responds to a notification message with an error. When the error is received by the source domain controller, it removes the entry for that destination from its **repsTo** data.

If a destination domain controller moves to a different site, the KCC adjusts the notify setting to indicate that notification is not needed.

#### Locating Replication Partners

The partner that initiates replication locates its partner by querying DNS to look up the IP address of the partner according to the GUID of the NTDS Settings object (class **nTDSDSA**). The NTDS Settings object represents the directory system agent (DSA) on the domain controller. Its GUID uniquely identifies the domain controller and is guaranteed to find the correct server, even if the name of the server has been changed.

The GUID of the NTDS Settings object is stored in the **objectGUID** attribute. The DSA GUID is created when Active Directory is installed on the domain controller, and is destroyed only if Active Directory is removed from the domain controller to create a member server.

### Note

- The Active Directory database also has a GUID, which is used by the DSA to identify the specific versions of the database when a database has been restored. This GUID is stored in the **invocationId** attribute on the NTDS Settings object.

As part of the DNS registration process, the Net Logon service on every domain controller registers a canonical name (CNAME) resource record, or alias record, which is constructed using the DSA GUID as the alias and maps to the DNS fully qualified domain name (FQDN) of the respective domain controller. The format of CNAME alias is:

*DsaGuid*._msdcs.*DNSForestName*.

To initiate replication, the domain controller retrieves the GUID-based DNS name that is stored in **repsTo** (for a source domain controller that is sending a change notification) and **repsFrom** (for a destination domain controller that is initiating replication) and queries DNS for the CNAME resource record. DNS responds by returning both the CNAME resource record and the host (A) resource record, which contains the IP address of the destination domain controller. The domain controller uses

information in the CNAME resource record to authenticate to the replication partner. Therefore, by using the CNAME and A resource record data, the domain controller can initiate replication.

**Note**

- The _msdcs.*DnsForestName* DNS zone contains a variety of forest-wide service (SRV) resource records that are used to locate special servers, such as domain controllers and global catalog servers, and to facilitate replication. In the context of this discussion, it is important to note that if a DNS server that is authoritative for the _msdcs.*DnsForestName* zone is not available, replication between domain controllers cannot occur.

For more information about the DSA and application directory partitions, see "How the Data Store Works." For more information about SRV resource records and the _msdcs subdomain, see "DNS Support for Active Directory Technical Reference".

**Urgent Replication**

Certain important events trigger replication immediately, overriding existing change notification. Urgent replication is implemented immediately by using RPC/IP to notify replication partners that changes have occurred on a source domain controller. Urgent replication uses regular change notification between destination and source domain controller pairs that otherwise use change notification, but notification is sent immediately in response to urgent events instead of waiting the default period of 15 seconds (or 300 seconds on domain controllers that are running Windows 2000).

**Events That Trigger Urgent Replication**

Urgent Active Directory replication is always triggered by certain events on all domain controllers within the same site, regardless of what operating system is running on the domain controller. When you have enabled change notification between sites, these triggering events also replicate immediately between sites.

Between domain controllers in the same site, immediate notification is caused by the following events:

- Assigning an account lockout, which a domain controller performs to prohibit a user from logging on after a certain number of failed attempts.

  | **Note** |
  | --- |
  | An account unlock is not urgently replicated. |

- Changing the account lockout policy.
- Changing the domain password policy.
- Changing a Local Security Authority (LSA) secret, which is a secure form in which private data is stored by the LSA (for example, the password for a trust relationship).
- Changing the password on a domain controller computer account.
- Changing the relative identifier (known as a "RID") master role owner, which is the single domain controller in a domain that assigns relative identifiers to all domain controllers in that domain.

**Urgent Replication of Account Lockout Changes**

Account lockout is a security feature that sets a limit on the number of failed authentication attempts that are allowed before the account is "locked out" from a further attempt to log on, in addition to a time limit for how long the lockout is in effect.

The PDC emulator receives urgent replication of account lockouts. In Active Directory domains, a single domain controller in each domain holds the role of PDC emulator, which simulates the behavior of a Windows NT version 3.*x*–based or Windows NT 4.0–based PDC. In Windows NT domains, the only domain controller that can accept updates is the PDC. If authentication fails at a BDC, the authentication request is passed immediately to the PDC, which is guaranteed to have the current password.

| **Note** |
| --- |

Because Active Directory Lightweight Directory Services (AD LDS) does not have a PDC emulator, urgent replication of account lockouts does not occur. An account lockout is replicated according to the replication frequency schedule. In order to ensure account lockouts are replicated as soon as possible in AD LDS, configure the replication frequency for every 15 minutes, or to another value that is suitable to your environment. The lowest value that can be set is 15 minutes. For more information, see Setting replication frequency (http://go.microsoft.com/fwlink/?LinkId=212626).

An account lockout is urgently replicated to the PDC emulator and is then urgently replicated to the following:

- Domain controllers in the same domain that are located in the same site as the PDC emulator.
- Domain controllers in the same domain that are located in the same site as the domain controller that handled the account lockout.
- Domain controllers in the same domain that are located in sites that have been configured to allow change notification between sites (and, therefore, urgent replication) with the site that contains the PDC emulator or with the site where the account lockout was handled. These sites include any site that is included in the same site link as the site that contains the PDC emulator or in the same site link as the site that contains the domain controller that handled the account lockout.

In addition, when authentication fails at a domain controller other than the PDC emulator, the authentication is retried at the PDC emulator. For this reason, the PDC emulator locks the account before the domain controller that handled the failed-password attempt if the bad-password-attempt threshold is reached.

**Note**

- When a bad password is used in an attempt to change a password, the lockout count is incremented on that domain controller only and is not replicated. As such, an attacker could try (# of domain controllers)*(lockout threshold -1) + 1 guesses before the account is locked out. Although this scenario has a relatively small impact on account lockout security, domains with an exceptionally high number of domain controllers represent a significant increase in the total number of guesses available to an attacker. Because a user cannot specify the domain controller on which the password change is attempted, an attack of this type requires an advanced tool.

**Replication of Password Changes**

Password changes are replicated differently than both normal (non-urgent) replication and urgent replication. Changes to security account passwords present a replication latency problem wherein a user's password is changed on domain controller A and the user subsequently attempts to log on, being authenticated by domain controller B. If the password has not replicated from A to B, the attempt to log on fails. Active Directory replication remedies this situation by forwarding password changes immediately to a single domain controller in the domain, the PDC emulator.

In Active Directory, when a user password is changed at a domain controller, that domain controller attempts to update the respective replica at the domain controller that holds the PDC emulator role. Update of the PDC emulator occurs immediately, without respect to schedules on site links. The updated password is propagated to other domain controllers by normal replication within a site.

When the user logs on to a domain and is authenticated by a domain controller that does not have the updated password, the domain controller refers to the PDC emulator to check the credentials of the user name and password rather than denying authentication based on an invalid password. Therefore, the user can log on successfully even when the authenticating domain controller has not yet received the updated password. On domain controllers that are running Windows Server 2003 or Windows 2000 Server with SP4, if the authentication is successful at the PDC emulator, the PDC emulator replicates the password immediately to the requesting domain controller to prevent that domain controller from having to check the PDC emulator again.

If the update at the PDC emulator fails for any reason, the password change is replicated non-urgently by normal replication.

For clients that are running Windows NT 4.0 or clients that are running Windows 95 or Windows 98 without the Directory Service Client Pack, the client attempts to contact the PDC emulator. If the client has the Directory Service Client Pack installed, the client contacts any domain controller and the contacted domain controller then attempts to contact the PDC emulator.

**Note**

- The Group Policy setting "Contact PDC on logon failure" can be disabled to keep a domain controller from contacting the

PDC emulator if the PDC emulator role owner is not in the current site. If this setting is disabled, the password change reaches the PDC emulator non-urgently through normal replication.

**Network Ports Used by Active Directory Replication**

By default, RPC-based replication uses dynamic port mapping. When connecting to an RPC endpoint during Active Directory replication, the RPC run time on the client contacts the RPC endpoint mapper on the server at a well-known port (port 135). The server queries the RPC endpoint mapper on this port to determine what port has been assigned for Active Directory replication on the server. This query occurs whether the port assignment is dynamic (the default) or fixed. The client never needs to know which port to use for Active Directory replication.

For more information about how to restrict Active Directory replication traffic and RPC traffic to a specific port, see article 224196 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=111871).

**Note**

- An endpoint comprises the protocol, local address, and port address.

In addition to the dynamic port 135, other ports that are required for replication to occur are listed in the following table.

**Port Assignments for Active Directory Replication**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | | 636 (Secure Sockets Layer [SSL]) |
| LDAP | | 3268 (global catalog) |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |

Replication within a domain also requires File Replication Service (FRS) using a dynamic RPC port. For more information about how to restrict FRS replication to a specific port, see article 319553 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=111872).

**Related Information**

The following resources contain additional information that is relevant to this section.

- How the Data Store Works
- Active Directory Replication Topology Technical Reference
- Active Directory Schema Technical Reference
- DNS Support for Active Directory Technical Reference

# Active Directory Replication Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Active Directory Replication Tools and Settings**

**In this section**


- Active Directory Replication Tools
- Active Directory Replication Registry Entries
- Active Directory Replication Group Policy Settings
- Active Directory Replication WMI Classes
- Network Ports Used by Active Directory Replication
- Related Information


**Active Directory Replication Tools**

The following tools are associated with Active Directory replication.

| Note |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**Dssite.msc: Active Directory Sites and Services**

**Category**

Active Directory Administrative Tools Microsoft Management Console (MMC) snap-in. This tool is installed automatically when you install Active Directory, and is available on the Start menu under Programs\Administrative Tools. This tool also ships with the Administration Tools Pack (Adminpak.msi).

**Version compatibility**

Active Directory Sites and Services provides a view into the Sites container of the configuration directory partition. Use Active Directory Sites and Services to manage Active Directory replication topology. The following objects and their properties can be managed by using this tool:

- Sites container: Add new sites.
- Site objects: Add new servers to a site.
- NTDS Site Settings object: For each site, view the connection object schedule and enable Universal group membership caching.
- Server object: View the NTDS Settings object and designate the server as a bridgehead server.
- NTDS Settings object: View inbound connections for the server. View the connection object schedule and change the source server for the connection.
- Inter-Site Transports container: Manage IP and SMTP site links.
- Site link objects: Manage the site link properties for a set of sites.

- Subnets container: Add, remove, and configure subnets with IP addresses. Associate subnets with sites.

**Repadmin.exe: Repadmin**

**Category**
Command-line tool.

**Version compatibility**
Repadmin is used to view the replication information on domain controllers. You can determine the last successful replication of all directory partitions, identify inbound and outbound replication partners, identify the current bridgehead servers, view object metadata, and generally manage Active Directory replication topology. You can use Repadmin to force replication of an entire directory partition or of a single object. You can also list domain controllers in a site.

Repadmin is extended to enable commands to target sets of domain controllers. For example, you can target all domain controllers in a site or domain, or all domain controllers that are global catalog servers.

Repadmin also includes the **RemoveLingeringObjects** command, which removes objects that are outdated (do not exist in a replica of the same directory partition on the source domain controller).

For more information about removing lingering objects, see "Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042)" in the Windows Server 2003 Operations Guide at http://go.microsoft.com/fwlink/?LinkId=44131. For more information about Repadmin, see Repadmin Overview.

**Ntdsutil.exe: Ntdsutil**

**Category**
Command-line tool.

**Version compatibility**
Ntdsutil.exe provides management capabilities for Active Directory. You can use Ntdsutil.exe to perform Active Directory database maintenance, manage and control single-master operations, and remove replication metadata left behind by domain controllers that are removed from the network without uninstalling Active Directory. The version of Ntdsutil that is included with Windows Server 2003 SP1 removes File Replication service (FRS) metadata in addition to Active Directory replication metadata. You can also use Ntdsutil to create application directory partitions and perform authoritative restore operations. This tool is intended for use by experienced administrators.

**Active Directory Replication Registry Entries**
The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following registry settings cannot be modified by using Group Policy or other Windows tools.

**NTDS Parameters Registry Settings**
The following registry entries are associated with Active Directory replication.

**Replicator notify pause after modify (secs)**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows 2000 Server.

**Default value**
Windows 2000 Server: 300 seconds.

The value for the delay between an originating update on a domain controller and the first change notification. On domain controllers running Windows Server 2003 or higher, the value for initial change notification delay is stored in the **msDSReplicationNotifyFirstDSADelay** attribute on the cross-reference object for each directory partition in the Configuration container. The default value in Windows Server 2003 and higher operating systems is decreased to 15 seconds

when the forest functional level is Windows Server 2003 or higher.

**Replicator notify pause between DSAs (secs)**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows 2000 Server.

**Default value**
Windows 2000 Server: 30 seconds

The value for the delay before each subsequent change notification. On domain controllers running Windows Server 2003, the value for subsequent notification delay is stored in the **msDSReplicationNotifySubsequentDSADelay** attribute on the cross-reference object for each directory partition in the Configuration container. The default value in Windows Server 2003 is decreased to 3 seconds when the forest functional level is Windows Server 2003.

**RPC Replication Timeout (mins)**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
Windows 2000 Server: 45 minutes; Windows Server 2003 and higher server operating systems: 5 minutes.

The number of minutes between initiation of Active Directory replication and the RPC timeout. The domain controller must be restarted before the change takes effect.

**Strict replication consistency**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server with SP3.

**Default value**
Windows 2000 Server with SP3: off (0); Windows Server 2003 and higher server operating systems: on (1)

The value that determines the treatment of replication of outdated objects that exist on reconnected domain controllers that have not replicated in longer than a tombstone lifetime. If the destination domain controller has strict replication consistency enabled, inbound replication of an outdated object is blocked. If the destination domain controller has strict replication disabled, inbound replication of the full object occurs.

**Replicator intra site packet size (objects)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1/1,000,000th the size of RAM, with a minimum of 100 objects and a maximum of 1,000 objects.

The maximum number of objects per packet for RPC replication within a site.

**Replicator intra site packet size (bytes)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1/100th the size of RAM, with a minimum of 1 megabyte (MB) and a maximum of 10 MB.

The maximum size of objects per packet for RPC replication within a site.

**Replicator inter site packet size (objects)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server.

**Default value**
1/1,000,000th the size of RAM, with a minimum of 100 objects and a maximum of 1,000 objects.

The maximum number of objects per packet for RPC replication between sites.

**Replicator inter site packet size (bytes)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

The maximum size of objects per packet for RPC replication between sites.

**Default value**
1/100th the size of RAM, with a minimum of 1 MB and a maximum of 10 MB.

**Replicator async inter site packet size (objects)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server.

**Default value**
1/1,000,000th the size of RAM, with a minimum of 100 objects and a maximum of 1,000 objects.

The maximum number of objects per packet for SMTP replication between sites.

**Replicator async inter site packet size (bytes)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1 MB.

The maximum size of objects per packet for SMTP replication between sites.

**Replicator compression algorithm**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003.

**Default value**
For Windows 2000 Server compression, change the value to 2.

Determines the compression algorithm that is used on a site link

**Repl topology update delay (secs)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
300 seconds.

Number of seconds to wait between the time Active Directory starts and the KCC performs the first topology check.

To find more information about **Repl topology update delay (secs)**, see "Registry Reference" in Tools and Settings Collection.

**Repl topology update period (secs)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server.

**Default value**
900 seconds.

Interval between KCC replication topology checks.

To find more information about **Repl topology update period (secs)**, see "Registry Reference" in Tools and Settings Collection.

**IntersiteFailuresAllowed**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1.

Number of failed replication attempts prior to excluding nonresponding servers from the intersite topology.

**MaxFailureTimeForIntersiteLink (sec)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
7200 seconds (2 hours).

Time in seconds that must elapse prior to excluding nonresponding servers from the intersite topology.

**NonCriticalLinkFailuresAllowed**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1.

Number of failed replication attempts prior to excluding nonresponding servers from the intrasite topology.

**MaxFailureTimeForNonCriticalLink (sec)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server.

**Default value**
43200 seconds (12 hours).

Time in seconds that must elapse prior to excluding nonresponding servers from the intrasite topology.

**CriticalLinkFailuresAllowed**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
0.

Number of failed replication attempts prior to excluding nonresponding servers for immediate neighbor connections within a site.

**MaxFailureTimeForCriticalLink (sec)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
7200 seconds (2 hours).

Time in seconds that must elapse prior to excluding nonresponding servers for immediate neighbor connections within a site.

**TCP/IP Port**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
135.

TCP port that the directory service uses instead of using dynamic port 135. The domain controller must be restarted before the change takes effect.

**Backup Latency Threshold (days)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003 with SP 1

**Default value**
Half the value of the tombstone lifetime of the forest.

When the value is reached, logs event ID 2089 in the Directory Service event log, warning administrators and monitoring applications to make sure that domain controllers are backed up before the tombstone lifetime expires.

**Active Directory Replication Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with Active Directory replication updates.

**Group Policy Settings Associated with Active Directory Replication**

| Group Policy Setting | Description |
|---|---|
| Account Lockout Policy:<br>• Account lockout duration<br>• Account lockout threshold<br>• Reset account lockout counter after | Changes to these settings in the Domain Security Policy trigger urgent replication. |
| Password Policy:<br>• Enforce password history<br>• Maximum password age<br>• Minimum password age<br>• Minimum password length<br>• Password must meet complexity requirements<br>• Store passwords using reversible encryption | Changes to these settings in the Domain Security Policy trigger urgent replication. |
| Contact PDC on logon failure | Account lockout and domain password changes rely on contacting the primary domain controller (PDC) emulator urgently to update the PDC emulator with the change. If Contact PDC on logon failure is disabled, replication of password changes to the PDC emulator occurs non-urgently. |

To find more information about these Group Policy settings, see “Group Policy Settings Reference” in Tools and Settings

Collection.

**Active Directory Replication WMI Classes**

The following table lists and describes the WMI classes that are associated with Active Directory replication. These classes are shipped with Windows Server 2003 or later server operating systems, but are also compatible with Windows 2000 Server.

**WMI Classes Associated with Active Directory Replication**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| MSAD_DomainController | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |
| MSAD_NamingContext | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |
| MSAD_ReplNeighbor | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |
| MSAD_ReplCursor | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |
| MSAD_ReplPendingOp | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by Active Directory Replication**

By default, RPC-based replication uses dynamic port mapping. When connecting to an RPC endpoint during Active Directory replication, the RPC run time on the client contacts the RPC endpoint mapper on the server at a well-known port (port 135). The server queries the RPC endpoint mapper on this port to determine what port has been assigned for Active Directory replication on the server. This query occurs whether the port assignment is dynamic (the default) or fixed. The client never needs to know which port to use for Active Directory replication.

**Note**

- An endpoint comprises the protocol, local address, and port address.

In addition to the dynamic port 135, other ports that are required for replication to occur are listed in the following table.

**Port Assignments for Active Directory Replication**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | | 636 (Secure Sockets Layer [SSL]) |
| LDAP | | 3268 (global catalog) |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

Replication within a domain also requires FRS using a dynamic RPC port.

**Related Information**

The following resources contain additional information that is relevant to this section.

- How the Active Directory Replication Model Works
- How Active Directory Replication Topology Works

# Active Directory Replication Topology Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Active Directory Replication Topology Technical Reference**

The replication topology of Active Directory directory service provides the network of connections between domain controllers in a forest according to their location in Active Directory sites. A site is an Active Directory object that you create and configure to represent an area of good network connectivity, typically corresponding to a local area network (LAN). The site object is associated with a set of one or more subnets, which are objects that identify a range of IP addresses. Each domain controller has an IP address that maps to a subnet, and that mapping in turn identifies the site of the domain controller. By recognizing domain controllers according to site locations, the replication system ensures that each domain controller is updated with directory changes in the most efficient and timely manner possible, given network conditions and directory service configuration. The replication topology is generated automatically at regular intervals to accommodate network and configuration changes, and is designed to ensure that all domain controllers are connected without redundancy and with minimum cost.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**In this subject**


- What Is Active Directory Replication Topology?
- How Active Directory Replication Topology Works
- Active Directory Replication Tools and Settings

# What Is Active Directory Replication Topology?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**What Is Active Directory Replication Topology?**

**In this section**


- Replication Within and Between Sites
- Technologies Related to Active Directory Replication Topology
- Active Directory Replication Topology Dependencies
- Related Information


Replication of updates to Active Directory objects are transmitted between multiple domain controllers to keep replicas of directory partitions synchronized. Multiple domains are common in large organizations, as are multiple sites in disparate locations. In addition, domain controllers for the same domain are commonly placed in more than one site.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

Therefore, replication must often occur both within sites and between sites to keep domain and forest data consistent among domain controllers that store the same directory partitions. Site objects can be configured to include a set of subnets that provide local area network (LAN) network speeds. As such, replication within sites generally occurs at high speeds between domain controllers that are on the same network segment. Similarly, site link objects can be configured to represent the wide area network (WAN) links that connect LANs. Replication between sites usually occurs over these WAN links, which might be costly in terms of bandwidth. To accommodate the differences in distance and cost of replication within a site and replication between sites, the intrasite replication topology is created to optimize speed, and the intersite replication topology is created to minimize cost.

The Knowledge Consistency Checker (KCC) is a distributed application that runs on every domain controller and is responsible for creating the connections between domain controllers that collectively form the replication topology. The KCC uses Active Directory data to determine where (from what source domain controller to what destination domain controller) to create these connections.

**Replication Within and Between Sites**

The KCC creates separate replication topologies to transfer Active Directory updates within a site and between all configured sites in the forest. The connections that are used for replication within sites are created automatically with no additional configuration. Intrasite replication takes advantage of LAN network speeds by providing replication as soon as changes occur, without the overhead of data compression, thus maximizing CPU efficiency. Intrasite replication connections form a ring topology with extra shortcut connections where needed to decrease latency. The fast replication of updates within sites facilitates timely updates of domain data. In deployments where large datacenters constitute hub sites for the centralization of mission-critical operations, directory consistency is critical.

Replication between sites is made possible by user-defined site and site link objects that are created in Active Directory to represent the physical LAN and WAN network infrastructure. When Active Directory sites and site links are configured, the KCC creates an intersite topology so that replication flows between domain controllers across WAN links. Intersite replication occurs according to a site link schedule so that WAN usage can be controlled, and is compressed to reduce network bandwidth requirements. Site link settings can be managed to optimize replication routing over WAN links. The connections that are created between sites form a spanning tree for each directory partition in the forest, merging where common directory partitions can be replicated over the same connection.

In remote branch locations, replication of updates from the hub sites is optimized for network availability. Thus, because intrasite replication is optimized for speed, branch locations across WAN links can be assured of receiving data from hub sites

that is up-to-date and reliable; but because intersite replication is scheduled, branch sites receive this replication only at intervals that are deemed appropriate and cost-effective for remote operations.

#### Technologies Related to Active Directory Replication Topology

The following technologies interact with Active Directory replication.

#### File Replication Service

File Replication service (FRS) is related to Active Directory replication because it requires the Active Directory replication topology. FRS is a multimaster replication service that is used to replicate files and folders in the system volume (SYSVOL) shared folder on domain controllers and in Distributed File System (DFS) shared folders. FRS works by detecting changes to files and folders and then replicating the updated files and folders to other replica members, which are connected in a replication topology.

FRS uses the replication topology that is generated by the KCC to replicate the SYSVOL files to all domain controllers in the domain. SYSVOL files are required by all domain controllers for Active Directory to function. For more information about FRS and how it uses the Active Directory replication topology, see "FRS Technical Reference". For more information about SYSVOL, see "Data Store Technical Reference."

#### SMTP

Simple Mail Transfer Protocol (SMTP) is a packaging protocol that can be used as an alternative to the remote procedure call (RPC) replication transport. SMTP can be used to transport nondomain replication over IP networks in mail-message format. Where networks are not fully routed, e-mail is sometimes the only transport method available.

#### Active Directory Replication Topology Dependencies

Active Directory replication topology has the following dependencies:

- **Routable IP infrastructure**. The replication topology is dependent upon a routable IP infrastructure from which you can map IP subnet address ranges to site objects. This mapping generates the information that is used by client workstations to communicate with domain controllers that are close by, when there is a choice, rather than those that are located across WAN links.

- **DNS**. The Domain Name System (DNS) resolves DNS names to IP addresses. Active Directory replication topology requires that DNS is properly designed and deployed so that domain controllers can correctly resolve the DNS names of replication partners.

  DNS also stores service (SRV) resource records that provide site affinity information to clients searching for domain controllers, including domain controllers that are searching for replication partners. Every domain controller registers these records so that they can be located according to site.

- **Net Logon service**. Net Logon is required for DNS registrations.

- **RPC**. Active Directory replication requires IP connectivity and RPC to transfer updates between replication partners within sites. RPC is required for replication between two sites containing domain controllers in the same domain, but SMTP is an alternative where RPC cannot be used and domain controllers for the same domain are all located in one site so that intersite replication of domain data is not required.

- **Intersite Messaging**. Intersite Messaging is required for SMTP intersite replication and for site coverage calculations. If the forest functional level is Windows 2000, Intersite Messaging is also required for intersite topology generation.

The following diagram shows the interaction of these technologies with the replication topology, which is indicated by the two-way connections between each set of domain controllers.

**Replication Topology and Dependent Technologies**




**Related Information**

The following resources contain additional information that is relevant to this section:

- How Active Directory Replication Topology Works
- TCP/IP Technical Reference

# How Active Directory Replication Topology Works

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**How Active Directory Replication Topology Works**

**In this section**


- Active Directory KCC Architecture and Processes
- Replication Topology Physical Structure
- Performance Limits for Replication Topology Generation
- Goals of Replication Topology
- Topology-Related Objects in Active Directory
- Replication Transports
- Replication Between Sites
- KCC and Topology Generation
- Network Ports Used by Replication Topology
- Related Information


Active Directory implements a replication topology that takes advantage of the network speeds within sites, which are ideally configured to be equivalent to local area network (LAN) connectivity (network speed of 10 megabits per second [Mbps] or higher). The replication topology also minimizes the use of potentially slow or expensive wide area network (WAN) links between sites.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

When you create a site object in Active Directory, you associate one or more Internet Protocol (IP) subnets with the site. Each domain controller in a forest is associated with an Active Directory site. A client workstation is associated with a site according to its IP address; that is, each IP address maps to one subnet, which in turn maps to one site.

Active Directory uses sites to:

- Optimize replication for speed and bandwidth consumption between domain controllers.
- Locate the closest domain controller for client logon, services, and directory searches.
- Direct a Distributed File System (DFS) client to the server that is hosting the requested data within the site.
- Replicate the system volume (SYSVOL), a collection of folders in the file system that exists on each domain controller in a domain and is required for implementation of Group Policy.

The ideal environment for replication topology generation is a forest that has a forest functional level of at least Windows

Server 2003. In this case, replication topology generation is faster and can accommodate more sites and domains than occurs when the forest has a forest functional level of Windows 2000. When at least one domain controller in each site is running Windows Server 2003, more domain controllers in each site can be used to replicate changes between sites than when all domain controllers are running Windows 2000 Server.

In addition, replication topology generation requires the following conditions:

- A Domain Name System (DNS) infrastructure that manages the name resolution for domain controllers in the forest. Active Directory–integrated DNS is assumed, wherein DNS zone data is stored in Active Directory and is replicated to all domain controllers that are DNS servers.
- All physical locations that are represented as site objects in Active Directory have LAN connectivity.
- IP connectivity is available between each site and all sites in the same forest that host operations master roles.
- Domain controllers meet the hardware requirements for Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.
- The appropriate number of domain controllers is deployed for each domain that is represented in each site.

This section covers the replication components that create the replication topology and how they work together, plus the mechanisms and rationale for routing replication traffic between domain controllers in the same site and in different sites.

**Active Directory KCC Architecture and Processes**

The replication topology is generated by the Knowledge Consistency Checker (KCC), a replication component that runs as an application on every domain controller and communicates through the distributed Active Directory database. The KCC functions locally by reading, creating, and deleting Active Directory data. Specifically, the KCC reads configuration data and reads and writes connection objects. The KCC also writes local, nonreplicated attribute values that indicate the replication partners from which to request replication.

For most of its operation, the KCC that runs on one domain controller does not communicate directly with the KCC on any other domain controller. Rather, all KCCs use the knowledge of the common, global data that is stored in the configuration directory partition as input to the topology generation algorithm to converge on the same view of the replication topology.

Each KCC uses its in-memory view of the topology to create inbound connections locally, manifesting only those results that apply to itself. The KCC communicates with other KCCs only to make a remote procedure call (RPC) request for replication error information. The KCC uses the error information to identify gaps in the replication topology. A request for replication error information occurs only between domain controllers in the same site.

**Note**

- The KCC uses only RPC to communicate with the directory service. The KCC does not use Lightweight Directory Access Protocol (LDAP).

One domain controller in each site is selected as the Intersite Topology Generator (ISTG). To enable replication across site links, the ISTG automatically designates one or more servers to perform site-to-site replication. These servers are called bridgehead servers. A bridgehead is a point where a connection leaves or enters a site.

The ISTG creates a view of the replication topology for all sites, including existing connection objects between all domain controllers that are acting as bridgehead servers. The ISTG then creates inbound connection objects for servers in its site that it determines will act as bridgehead servers and for which connection objects do not already exist. Thus, the scope of operation for the KCC is the local server only, and the scope of operation for the ISTG is a single site.

Each KCC has the following global knowledge about objects in the forest, which it gets by reading objects in the Sites container of the configuration directory partition and which it uses to generate a view of the replication topology:

- Sites
- Servers
- Site affiliation of each server

- Global catalog servers
- Directory partitions stored by each server
- Site links
- Site link bridges

Detailed information about these configuration components and their functionality is provided later in this section.

The following diagram shows the KCC architecture on servers in the same forest in two sites.

**KCC Architecture and Processes**




The architecture and process components in the preceding diagram are described in the following table.

**KCC Architecture and Process Components**

| Component | Description |
|---|---|
| Knowledge Consistency Checker (KCC) | The application running on each domain controller that communicates directly with the Ntdsa.dll to read and write replication objects. |
| Directory System Agent (DSA) | The directory service component that runs as Ntdsa.dll on each domain controller, providing the interfaces through which services and processes such as the KCC gain access to the directory database. |
| Extensible Storage Engine (ESE) | The directory service component that runs as Esent.dll. ESE manages the tables of records, each with one or more columns. The tables of records comprise the directory database. |
| Remote procedure call (RPC) | The Directory Replication Service (Drsuapi) RPC protocol, used to communicate replication status and topology to a domain controller. The KCC also uses this protocol to communicate with other KCCs to request error information when building the replication topology. |

| | |
|---|---|
| Intersite Topology Generator (ISTG) | The single KCC in a site that manages intersite connection objects for the site. |

The four servers in the preceding diagram create identical views of the servers in their site and generate connection objects on the basis of the current state of Active Directory data in the configuration directory partition. In addition to creating its view of the servers in its respective site, the KCC that operates as the ISTG in each site also creates a view of all servers in all sites in the forest. From this view, the ISTG determines the connections to create on the bridgehead servers in its own site.

**Note**

- A connection requires two endpoints: one for the destination domain controller and one for the source domain controller. Domain controllers creating an intrasite topology always use themselves as the destination end point and must consider only the endpoint for the source domain controller. The ISTG, however, must identify both endpoints in order to create connection objects between two other servers.

Thus, the KCC creates two types of topologies: intrasite and intersite. Within a site, the KCC creates a ring topology by using all servers in the site. To create the intersite topology, the ISTG in each site uses a view of all bridgehead servers in all sites in the forest. The following diagram shows a high-level generalization of the view that the KCC sees of an intrasite ring topology and the view that the ISTG sees of the intersite topology. Lines between domain controllers within a site represent inbound and outbound connections between the servers. The lines between sites represent configured site links. Bridgehead servers are represented as BH.

**KCC and ISTG Views of Intrasite and Intersite Topology**




**Replication Topology Physical Structure**

The Active Directory replication topology can use many different components. Some components are required and others are not required but are available for optimization. The following diagram illustrates most replication topology components and their place in a sample Active Directory multisite and multidomain forest. The depiction of the intersite topology that uses multiple bridgehead servers for each domain assumes that at least one domain controller in each site is running at least Windows Server 2003. All components of this diagram and their interactions are explained in detail later in this section.

**Replication Topology Physical Structure**




In the preceding diagram, all servers are domain controllers. They independently use global knowledge of configuration data to generate one-way, inbound connection objects. The KCCs in a site collectively create an intrasite topology for all domain controllers in the site. The ISTGs from all sites collectively create an intersite topology. Within sites, one-way arrows indicate the inbound connections by which each domain controller replicates changes from its partner in the ring. For intersite replication, one-way arrows represent inbound connections that are created by the ISTG of each site from bridgehead servers (BH) for the same domain (or from a global catalog server [GC] acting as a bridgehead if the domain is not present in the site) in other sites that share a site link. Domains are indicated as D1, D2, D3, and D4.

Each site in the diagram represents a physical LAN in the network, and each LAN is represented as a site object in Active Directory. Heavy solid lines between sites indicate WAN links over which two-way replication can occur, and each WAN link is represented in Active Directory as a site link object. Site link objects allow connections to be created between bridgehead servers in each site that is connected by the site link.

Not shown in the diagram is that where TCP/IP WAN links are available, replication between sites uses the RPC replication transport. RPC is always used within sites. The site link between Site A and Site D uses the SMTP protocol for the replication transport to replicate the configuration and schema directory partitions and global catalog partial, read-only directory partitions. Although the SMTP transport cannot be used to replicate writable domain directory partitions, this transport is required because a TCP/IP connection is not available between Site A and Site D. This configuration is acceptable for replication because Site D does not host domain controllers for any domains that must be replicated over the site link A-D.

By default, site links A-B and A-C are transitive (bridged), which means that replication of domain D2 is possible between Site B and Site C, although no site link connects the two sites. The cost values on site links A-B and A-C are site link settings that determine the routing preference for replication, which is based on the aggregated cost of available site links. The cost of a direct connection between Site C and Site B is the sum of costs on site links A-B and A-C. For this reason, replication between Site B and Site C is automatically routed through Site A to avoid the more expensive, transitive route. Connections are created between Site B and Site C only if replication through Site A becomes impossible due to network or bridgehead server conditions.

**Performance Limits for Replication Topology Generation**

Active Directory topology generation performance is limited primarily by the memory on the domain controller. KCC performance degrades at the physical memory limit. In most deployments, topology size will be limited by the amount of domain controller memory rather than CPU utilization required by the KCC.

Scaling of sites and domains is improved in Windows Server 2003 by improving the algorithm that the KCC uses to generate

the intersite replication topology. Because all domain controllers must use the same algorithm to arrive at a consistent view of the replication topology, the improved algorithm has a forest functional level requirement of Windows Server 2003 or Windows Server 2003 interim.

KCC scalability was tested on domain controllers with 1.8 GHz processor speed, 512 megabytes (MB) RAM, and small computer system interface (SCSI) disks. KCC performance results at forest functional levels that are at least Windows Server 2003 are described in the following table. The times shown are for the KCC to run where all new connections are needed (maximum) and where no new connections are needed (minimum). Because most organizations add domain controllers in increments, the minimum generation times shown are closest to the actual runtimes that can be expected in deployments of comparable sizes. The CPU and memory usage values for the Local Security Authority (LSA) process (Lsass.exe) indicate the more significant impact of memory versus percent of CPU usage when the KCC runs.

**Note**

- Active Directory runs as part of the LSA, which manages authentication packages and authenticates users and services.

**Minimum and Maximum KCC Generation Times for Domain-Site Combinations**

| Domains | Sites | Connections | KCC Generation Time (seconds) | Lsass.exe Memory Usage (MB) | Lsass.exe CPU Usage (%) |
|---|---|---|---|---|---|
| **1** | 500 | Maximum | 43 | 100 | 39 |
| | | Minimum | 1 | 100 | 29 |
| | 1,000 | Maximum | 49 | 149 | 43 |
| | | Minimum | 2 | 149 | 28 |
| | 3,000 | Maximum | 69 | 236 | 46 |
| | | Minimum | 2 | 236 | 63 |
| **5** | 500 | Maximum | 70 | 125 | 29 |
| | | Minimum | 2 | 126 | 71 |
| | 1,000 | Maximum | 77 | 237 | 28 |
| | | Minimum | 3 | 237 | 78 |
| | 2,000 | Maximum | 78 | 325 | 43 |
| | | Minimum | 5 | 325 | 77 |
| | 3,000 | Maximum | 85 | 449 | 52 |
| | | Minimum | 6 | 449 | 75 |
| | 4,000 | Maximum | 555 | 624 | 46 |
| | | Minimum | 34 | 624 | 69 |
| **20** | 1,000 | Maximum | 48 | 423 | 65 |
| | | Minimum | 5 | 423 | 81 |
| **40** | 1,000 | Maximum | 93 | 799 | 56 |
| | | Minimum | 12 | 799 | 96 |
| | 2,000 | Minimum | 38 | 874 | 71 |

These numbers cannot be used as the sole guidelines for forest and domain design. Other limitations might affect performance and scalability. A limitation of note is that when FRS is deployed, a limit of 1,200 domain controllers per domain is recommended to ensure reliable recovery of SYSVOL.

For more information about FRS limitations, see "FRS Technical Reference." For more information about the functional level requirements for the intersite topology generation algorithm, see "Automated Intersite Topology Generation" later in this section.

**Goals of Replication Topology**
The KCC generates a replication topology that achieves the following goals:

- Connect every directory partition replica that must be replicated.
- Control replication latency and cost.
- Route replication between sites.
- Effect client affinity.

By default, the replication topology is managed automatically and optimizes existing connections. However, manual connections created by an administrator are not modified or optimized.

**Connect Directory Partition Replicas**
The total replication topology is actually composed of several underlying topologies, one for each directory partition. In the case of the schema and configuration directory partitions, a single topology is created. The underlying topologies are merged to form the minimum number of connections that are required to replicate each directory partition between all domain controllers that store replicas. Where the connections for directory partitions are identical between domain controllers — for example, two domain controllers store the same domain directory partition — a single connection can be used for replication of updates to the domain, schema, and configuration directory partitions.

A separate replication topology is also created for application directory partitions. However, in the same manner as schema and configuration directory partitions, application directory partitions can use the same topology as domain directory partitions. When application and domain directory partitions are common to the source and destination domain controllers, the KCC does not create a separate connection for the application directory partition.

A separate topology is not created for the partial replicas that are stored on global catalog servers. The connections that are needed by a global catalog server to replicate each partial replica of a domain are part of the topology that is created for each domain.

The routes for the following directory partitions or combinations of directory partitions are aggregated to arrive at the overall topology:

- Configuration and schema within a site.
- Each writable domain directory partition within a site.
- Each application directory partition within a site.
- Global catalog read-only, partial domain directory partitions within a site.
- Configuration and schema between sites.
- Each writable domain directory partition between sites.
- Each application directory partition between sites.
- Global catalog read-only, partial domain directory partitions between sites.

Replication transport protocols determine the manner in which replication data is transferred over the network media. Your network environment and server configuration dictates the transports that you can use. For more information about transports, see "Replication Transports" later in this section.

**Control Replication Latency and Cost**
Replication latency is inherent in a multimaster directory service. A period of replication latency begins when a directory

update occurs on an originating domain controller and ends when replication of the change is received on the last domain controller in the forest that requires the change. Generally, the latency that is inherent in a WAN link is relative to a combination of the speed of the connection and the available bandwidth. Replication cost is an administrative value that can be used to indicate the latency that is associated with different replication routes between sites. A lower-cost route is preferred by the ISTG when generating the replication topology.

Site topology is the topology as represented by the physical network: the LANs and WANs that connect domain controllers in a forest. The replication topology is built to use the site topology. The site topology is represented in Active Directory by site objects and site link objects. These objects influence Active Directory replication to achieve the best balance between replication speed and the cost of bandwidth utilization by distinguishing between replication that occurs within a site and replication that must span sites. When the KCC creates replication connections between domain controllers to generate the replication topology, it creates more connections between domain controllers in the same site than between domain controllers in different sites. The results are lower replication latency within a site and less replication bandwidth utilization between sites.

Within sites, replication is optimized for speed as follows:

- Connections between domain controllers in the same site are always arranged in a ring, with possible additional connections to reduce latency.
- Replication within a site is triggered by a change notification mechanism when an update occurs, moderated by a short, configurable delay (because groups of updates frequently occur together).
- Data is sent uncompressed, and thus without the processing overhead of data compression.

Between sites, replication is optimized for minimal bandwidth usage (cost) as follows:

- Replication data is compressed to minimize bandwidth consumption over WAN links.
- Store-and-forward replication makes efficient use of WAN links — each update crosses an expensive link only once.
- Replication occurs at intervals that you can schedule so that use of expensive WAN links is managed.
- The intersite topology is a layering of spanning trees (one intersite connection between any two sites for each directory partition) and generally does not contain redundant connections.

**Route Replication Between Sites**

The KCC uses the information in Active Directory to identify the least-cost routes for replication between sites. If a domain controller is unavailable at the time the replication topology is created, making replication through that site impossible, the next least-cost route is used. This rerouting is automatic when site links are bridged (transitive), which is the default setting.

Replication is automatically routed around network failures and offline domain controllers.

**Effect Client Affinity**

Active Directory clients locate domain controllers according to their site affiliation. Domain controllers register SRV resource records in the DNS database that map the domain controller to a site. When a client requests a connection to a domain controller (for example, when logging on to a domain computer), the domain controller Locator uses the site SRV resource record to locate a domain controller with good connectivity whenever possible. In this way, a client locates a domain controller within the same site, thereby avoiding communications over WAN links.

Sites can also be used by certain applications, such as DFS, to ensure that clients locate servers that are within the site or, if none is available, a server in the next closest site. If the ISTG is running Windows Server 2003 or later server operating systems, you can specify an alternate site based on connection cost if no same-site servers are available. This DFS feature, called "site costing," is new in Windows Server 2003.

For more information about the domain controller Locator, see "DNS Support for Active Directory Technical Reference." For more information about DFS site costing, see "DFS Technical Reference."

**Topology-Related Objects in Active Directory**

Active Directory stores replication topology information in the configuration directory partition. Several configuration objects define the components that are required by the KCC to establish and implement the replication topology.

Active Directory Sites and Services is the Microsoft Management Console (MMC) snap-in that you can use to view and manage the hierarchy of objects that are used by the KCC to construct the replication topology. The hierarchy is displayed as the contents of the Sites container, which is a child object of the Configuration container. The Configuration container is not identified in the Active Directory Sites and Services UI. The Sites container contains an object for each site in the forest. In addition, Sites contains the Subnets container, which contains subnet definitions in the form of subnet objects.

The following figure shows a sample hierarchy, including two sites: Default-First-Site-Name and Site A. The selected NTDS Settings object of the server MHDC3 in the site Default-First-Site-Name displays the inbound connections from MHDC4 in the same site and from A-DC-01 in Site A. In addition to showing that MHDC3 and MHDC4 perform intrasite replication, this configuration indicates that MHDC3 and A-DC-01 are bridgehead servers that are replicating the same domain between Site A and Default-First-Site-Name.

**Sites Container Hierarchy**




**Site and Subnet Objects**

Sites are effective because they map to specific ranges of subnet addresses, as identified in Active Directory by subnet objects. The relationship between sites and subnets is integral to Active Directory replication.

**Site Objects**

A site object (class **site**) corresponds to a set of one or more IP subnets that have LAN connectivity. Thus, by virtue of their subnet associations, domain controllers that are in the same site are well connected in terms of speed. Each site object has a child NTDS Site Settings object and a Servers container. The distinguished name of the Sites container is CN=Sites,CN=Configuration,DC=*ForestRootDomainName*. The Configuration container is the topmost object in the configuration directory partition and the Sites container is the topmost object in the hierarchy of objects that are used to manage and implement Active Directory replication.

When you install Active Directory on the first domain controller in the forest, a site object named Default-First-Site-Name is created in the Sites container in Active Directory.

**Subnet Objects**

Subnet objects (class **subnet**) define network subnets in Active Directory. A network subnet is a segment of a TCP/IP network to which a set of logical IP addresses is assigned. Subnets group computers in a way that identifies their physical proximity on the network. Subnet objects in Active Directory are used to map computers to sites. Each subnet object has a **siteObject** attribute that links it to a site object.

**Subnet-to-Site Mapping**

You associate a set of IP subnets with a site if they have high-bandwidth LAN connectivity, possibly involving hops through high-performance routers.

**Note**

- LAN connectivity assumes high-speed, inexpensive bandwidth that allows similar and reliable network performance, regardless of which two computers in the site are communicating. This quality of connectivity does not indicate that all servers in the site must be on the same network segment or that hop counts between all servers must be identical. Rather, it is the measure by which you know that if a large amount of data needs to be copied from one server to another, it does not matter which servers are involved. If you find that you are concerned about such situations, consider creating another site.

When you create subnet objects in Active Directory, you associate them with site objects so that IP addresses can be localized according to sites. During the process of domain controller location, subnet information is used to find a domain controller in the same site as, or the site closest to, the client computer. The Net Logon service on a domain controller is able to identify the site of a client by mapping the client's IP address to a subnet object in Active Directory. Likewise, when a domain controller is installed, its server object is created in the site that contains the subnet that maps to its IP address.

You can use Active Directory Sites and Services to define subnets, and then create a site and associate the subnets with the site. By default, only members of the Enterprise Admins group have the right to create new sites, although this right can be delegated.

In a default Active Directory installation, there is no default subnet object, so potentially a computer can be in the forest but have an IP subnet that is not contained in any site. For private networks, you can specify the network addresses that are provided by the Internet Assigned Numbers Authority (IANA). By definition, that range covers all of the subnets for the organization. However, where several class B or class C addresses are assigned, there would necessarily be multiple subnet objects that all mapped to the same default site.

To accommodate this situation, use the following subnets:

- For class B addresses, subnet 128.0.0.0/2 covers all class B addresses.

- For class C addresses, subnet 192.0.0.0/3 covers all class C addresses.

**Note**

- The Active Directory Sites and Services MMC snap-in neither checks nor enforces IP address mapping when you move a server object to a different site. You must manually change the IP address on the domain controller to ensure proper mapping of the IP address to a subnet in the appropriate site.

**Server Objects**

Server objects (class **server**) represent server computers, including domain controllers, in the configuration directory partition. When you install Active Directory, the installation process creates a server object in the Servers container within the site to which the IP address of the domain controller maps. There is one server object for each domain controller in the site.

A server object is distinct from the computer object that represents the computer as a security principal. These objects are in separate directory partitions and have separate globally unique identifiers (GUIDs). The computer object represents the domain controller in the domain directory partition; the server object represents the domain controller in the configuration directory partition. The server object contains a reference to the associated computer object.

The server object for the first domain controller in the forest is created in the Default-First-Site-Name site. When you install Active Directory on subsequent servers, if no other sites are defined, server objects are created in Default-First-Site-Name. If other sites have been defined and subnet objects have been associated with these sites, server objects are created as follows:

- If additional sites have been defined in Active Directory and the IP address of the installation computer matches an existing subnet in a defined site, the domain controller is added to that site.

- If additional sites have been defined in Active Directory and the new domain controller's IP address does not match an existing subnet in one of the defined sites, the new domain controller's server object is created in the site of the source domain controller from which the new domain controller receives its initial replication.

When Active Directory is removed from a server, its NTDS Settings object is deleted from Active Directory, but its server object remains because the server object might contain objects other than NTDS Settings. For example, when Microsoft Operations Manager or Message Queuing is running on a domain controller, these applications create child objects beneath the server object.

**NTDS Settings Objects**
The NTDS Settings object (class **nTDSDSA**) represents an instantiation of Active Directory on that server and distinguishes a domain controller from other types of servers in the site or from decommissioned domain controllers. For a specific server object, the NTDS Settings object contains the individual connection objects that represent the inbound connections from other domain controllers in the forest that are currently available to send changes to this domain controller.

### Note

- The NTDS Settings object should not be manually deleted.

The **hasMasterNCs** multivalued attribute (where "NC" stands for "naming context," a synonym for "directory partition") of an NTDS Settings object contains the distinguished names for the set of writable (non-global-catalog) directory partitions that are located on that domain controller, as follows:

- DC=Configuration,DC=*ForestRootDomainName*
- DC=Schema,DC=Configuration,DC=*ForestRootDomainName*
- DC=*DomainName*,DC=*ForestRootDomainName*

The **msDSHasMasterNCs** attribute is new attribute introduced in Windows Server 2003, and this attribute of the NTDS Settings object contains values for the above-named directory partitions as well as any application directory partitions that are stored by the domain controller. Therefore, on domain controllers that are DNS servers and use Active Directory–integrated DNS zones, the following values appear in addition to the default directory partitions:

- DC=ForestDNSZones,DC=*ForestRootDomainName* (domain controllers in the forest root domain only)
- DC=DomainDNSZones,DC=*DomainName*,DC=*ForestRootDomainName* (all domain controllers)

Applications that need to retrieve the list of all directory partitions that are hosted by a domain controller can be updated or written to use the **msDSHasMasterNCs** attribute. Applications that need to retrieve only domain directory partitions can continue to use the **hasMasterNCs** attribute.

For more information about these attributes, see Active Directory in the Microsoft Platform SDK on MSDN.

**Connection Objects**
A connection object (class **nTDSConnection**) defines a one-way, inbound route from one domain controller (the source) to the domain controller that stores the connection object (the destination). The KCC uses information in cross-reference objects to create the appropriate connection objects, which enable domain controllers that store the same directory partitions to replicate with each other. The KCC creates connections for every server object in the Sites container that has an NTDS Settings object.

The connection object is a child of the replication destination's NTDS Settings object, and the connection object references the replication source domain controller in the **fromServer** attribute on the connection object — that is, it represents the inbound half of a connection. The connection object contains a replication schedule and specifies a replication transport. The connection object schedule is derived from the site link schedule for intersite connections. For more information about intersite connection schedules, see "Connection Object Schedule" later in this section.

A connection is unidirectional; a bidirectional replication connection is represented as two inbound connection objects. The KCC creates one connection object under the NTDS Settings object of each server that is used as an endpoint for the connection.

Connection objects are created in two ways:

- Automatically by the KCC.
- Manually by a directory administrator by using Active Directory Sites and Services, ADSI Edit, or scripts.

Intersite connection objects are created by the KCC that has the role of intersite topology generator (ISTG) in the site. One domain controller in each site has this role, and the ISTG role owners in all sites use the same algorithm to collectively generate

the intersite replication topology.

**Ownership of Connection Objects**

Connections that are created automatically by the KCC are "owned" by the KCC. If you create a new connection manually, the connection is not owned by the KCC. If a connection object is not owned by the KCC, the KCC does not modify it or delete it.

**Note**

- One exception to this modification rule is that the KCC automatically changes the transport type of an administrator-owned connection if the **transportType** attribute is set incorrectly (see "Transport Type" later in this section).

However, if you modify a connection object that is owned by the KCC (for example, you change the connection object schedule), the ownership of the connection depends on the application that you use to make the change:

- If you use an LDAP editor such as Ldp.exe or Adsiedit.msc to change a connection object property, the KCC reverses the change the next time it runs.

- If you use Active Directory Sites and Services to change a connection object property, the object is changed from automatic to manual and the KCC no longer owns it. The UI indicates the ownership status of each connection object.

In most Active Directory deployments, manual connection objects are not needed.

If you create a connection object, it remains until you delete it, but the KCC will automatically delete duplicate KCC-owned objects if they exist and will continue to create needed connections. Ownership of a connection object does not affect security access to the object; it determines only whether the KCC can modify or delete the object.

**Note**

- If you create a new connection object that duplicates one that the KCC has already created, your duplicate object is created and the KCC-created object is deleted by the KCC the next time it runs.

**ISTG and Modified Connections**

Because connection objects are stored in the configuration directory partition, it is possible for an intersite connection object to be modified by an administrator on one domain controller and, prior to replication of the initial change being received, to be modified by the KCC on another domain controller. Overwriting such a change can occur within the local site or when a connection object changes in a remote site.

By default, the KCC runs every 15 minutes. If the administrative connection object change is not received by the destination domain controller before the ISTG in the destination site runs, the ISTG in the destination site might modify the same connection object. In this case, ownership of the connection object belongs to the KCC because the latest write to the connection object is the write that is applied.

**Manual Connection Objects**

The KCC is designed to produce a replication topology that provides low replication latency, that adapts to failures, and that does not need modification. It is usually not necessary to create connection objects when the KCC is being used to generate automatic connections. The KCC automatically reconfigures connections as conditions change. Adding manual connections when the KCC is employed potentially increases replication traffic by adding redundant connections to the optimal set chosen by the KCC. When manually generated connections exist, the KCC uses them wherever possible.

*Adding extra connections does not necessarily reduce replication latency*. Within a site, latency issues are usually related to factors other than the replication topology that is generated by the KCC. Factors that affect latency include the following:

- Interruption of the service of key domain controllers, such as the primary domain controller (PDC) emulator, global catalog servers, or bridgehead servers.

- Domain controllers that are too busy to replicate in a timely manner (too few domain controllers).

- Network connectivity issues.

- DNS server problems.

- Inordinate amounts of directory updates.

For problems such as these, creating a manual connection does not improve replication latency. Adjusting the scheduling and costs that are assigned to the site link is the best way to influence intersite topology.

**Site Link Objects**
For a connection object to be created on a destination domain controller in one site that specifies a source domain controller in another site, you must manually create a site link object (class **siteLink**) that connects the two sites. Site link objects identify the transport protocol and scheduling required to replicate between two or more sites. You can use Active Directory Sites and Services to create the site links. The KCC uses the information stored in the properties of these site links to create the intersite topology connections.

A site link is associated with a network transport by creating the site link object in the appropriate transport container (either IP or SMTP). All intersite domain replication must use IP site links. The Simple Mail Transfer Protocol (SMTP) transport can be used for replication between sites that contain domain controllers that do not host any common domain directory partition replicas.

**Site Link Properties**
A site link specifies the following:

- Two or more sites that are permitted to replicate with each other.
- An administrator-defined cost value associated with that replication path. The cost value controls the route that replication takes, and thus the remote sites that are used as sources of replication information.
- A schedule during which replication is permitted to occur.
- An interval that determines how frequently replication occurs over this site link during the times when the schedule allows replication.

For more information about site link properties, see "Site Link Settings and Their Effects on Intersite Replication" later in this section.

**Default Site Link**
When you install Active Directory on the first domain controller in the forest, an object named DEFAULTIPSITELINK is created in the Sites container (in the IP container within the Inter-Site Transports container). This site link contains only one site, Default-First-Site-Name.

**Site Link Bridging**
By default, site links for the same IP transport that have sites in common are bridged, which enables the KCC to treat the set of associated site links as a single route. If you categorically do not want the KCC to consider some routes, or if your network is not fully routed, you can disable automatic bridging of all site links. When this bridging is disabled, you can create site link bridge objects and manually add site links to a bridge. For more information about using site link bridges, see "Bridging Site Links Manually" later in this section.

**NTDS Site Settings Object**
NTDS Site Settings objects (class **nTDSSiteSettings**) identify site-wide settings in Active Directory. There is one NTDS Site Settings object per site in the Sites container. NTDS Site Settings attributes control the following features and conditions:

- The identity of the ISTG role owner for the site. The KCC on this domain controller is responsible for identifying bridgehead servers. For more information about this role, see "Automated Intersite Topology Generation" later in this section.
- Whether domain controllers in the site cache membership of universal groups and the site in which to find a global catalog server for creating the cache.
- The default schedule that applies to connection objects. For more information about this schedule, see "Connection Object Schedule" later in this section.

**Note**

- To allow for the possibility of network failure, which might cause one or more notifications to be missed, a default schedule of once per hour is applied to replication within a site. You do not need to manage this schedule.

**Cross-Reference Objects**

Cross-reference objects (class **crossRef**) store the location of directory partitions in the Partitions container (CN=Partitions,CN=Configuration,DC=*ForestRootDomainName*). The contents of the Partitions container are not visible by using Active Directory Sites and Services, but can be viewed by using Adsiedit.msc to view the Configuration directory partition.

Active Directory replication uses cross-reference objects to locate the domain controllers that store each directory partition. A cross-reference object is created during Active Directory installation to identify each new directory partition that is added to the forest. Cross-reference objects store the identity (**nCName**, the distinguished name of the directory partition where "NC" stands for "naming context," a synonym for "directory partition") and location (**dNSRoot**, the DNS domain where servers that store the particular directory partition can be reached) of each directory partition.

### Note

- Starting in Windows Server 2003 Active Directory, a special attribute of the cross-reference object, **msDS-NC-Replica-Locations**, identifies application directory partitions to the replication system. For more information about how application directory partitions are replicated, see "Topology Generation Phases" later in this section.

**Replication Transports**

Replication transports provide the wire protocols that are required for data transfer. There are three levels of connectivity for replication of Active Directory information:

- Uniform high-speed, synchronous RPC over IP within a site.
- Point-to-point, synchronous, low-speed RPC over IP between sites.
- Low-speed, asynchronous SMTP between sites.

The following rules apply to the replication transports:

- Replication within a site always uses RPC over IP.
- Replication between sites can use either RPC over IP or SMTP over IP.
- Replication between sites over SMTP is supported for only domain controllers of different domains. Domain controllers of the same domain must replicate by using the RPC over IP transport. Therefore, replication between sites over SMTP is supported for only schema, configuration, and global catalog replication, which means that domains can span sites only when point-to-point, synchronous RPC is available between sites.

The Inter-Site Transports container provides the means for mapping site links to the transport that the link uses. When you create a site link object, you create it in either the IP container (which associates the site link with the RPC over IP transport) or the SMTP container (which associates the site link with the SMTP transport).

For the IP transport, a typical site link connects only two sites and corresponds to an actual WAN link. An IP site link connecting more than two sites might correspond to an asynchronous transfer mode (ATM) backbone that connects, for example, more than two clusters of buildings on a large campus or connects several offices in a large metropolitan area that are connected by leased lines and IP routers.

**Synchronous and Asynchronous Communication**

The RPC intersite and intrasite transport (RCP over IP within sites and between sites) and the SMTP intersite transport (SMTP over IP between sites only) correspond to synchronous and asynchronous communication methods, respectively. Synchronous communication favors fast, available connections, while asynchronous communication is better suited for slow or intermittent connections.

**Synchronous Replication Over IP**

The IP transport (RPC over IP) provides synchronous inbound replication. In the context of Active Directory replication, synchronous communication implies that after the destination domain controller sends the request for data, it waits for the

source domain controller to receive the request, construct the reply, and send the reply before it requests changes from any other domain controllers; that is, inbound replication is sequential. Thus in synchronous transmission, the reply is received within a short time. The IP transport is appropriate for linking sites in fully routed networks.

**Asynchronous Replication Over SMTP**
The SMTP transport (SMTP over IP) provides asynchronous replication. In asynchronous replication, the destination domain controller does not wait for the reply and it can have multiple asynchronous requests outstanding at any particular time. Thus in asynchronous transmission, the reply is not necessarily received within a short time. Asynchronous transport is appropriate for linking sites in networks that are not fully routed and have particularly slow WAN links.

### Note

- Although asynchronous replication can send multiple replication requests in parallel, the received replication packets are queued on the destination domain controller and the changes applied for only one partner and directory partition at a time.

**Replication Queue**
Suppose a domain controller has five inbound replication connections. As the domain controller formulates change requests, either by a schedule being reached or from a notification, it adds a work item for each request to the end of the queue of pending synchronization requests. Each pending synchronization request represents one <source domain controller, directory partition> pair, such as "synchronize the schema directory partition from DC1," or "delete the ApplicationX directory partition."

When a work item has been received into the queue, notification and polling intervals do not apply — the domain controller processes the item (begins synchronizing from that source) as soon as the item reaches the front of the queue, and continues until either the destination is fully synchronized with the source domain controller, an error occurs, or the synchronization is pre-empted by a higher-priority operation.

**SMTP Intersite Replication**
When sites are on opposite ends of a WAN link (or the Internet), it is not always desirable — or even possible — to perform synchronous, RPC-based directory replication. In some cases, the only method of communication between two sites is e-mail. When connectivity is intermittent or when end-to-end IP connectivity is not available (an intermediate site does not support RPC/IP replication), replication must be possible across asynchronous, store-and-forward transports such as SMTP.

In addition, where bandwidth is limited, it can be disadvantageous to force an entire replication cycle of request for changes and transfer of changes between two domain controllers to complete before another can begin (that is, to use synchronous replication). With SMTP, several cycles can be processing simultaneously so that each cycle is being processed to some degree most of the time, as opposed to receiving no attention for prolonged periods, which can result in RPC time-outs.

For intersite replication, SMTP replication substitutes mail messaging for the RPC transport. The message syntax is the same as for RPC-based replication. There is no change notification for SMTP–based replication, and scheduling information for the site link object is used as follows:

- By default, SMTP replication ignores the **Replication Available** and **Replication Not Available** settings on the site link schedule in Active Directory Sites and Services (the information that indicates when these sites are connected). Replication occurs according to the messaging system schedule.

- Within the scope of the messaging system schedule, SMTP replication uses the replication interval that is set on the SMTP site link to indicate how often the server requests changes. The interval (**Replicate every ____ minutes**) is set in 15-minute increments on the **General** tab in site link **Properties** in Active Directory Sites and Services.

The underlying SMTP messaging system is responsible for message routing between SMTP servers.

**SMTP Replication and Intersite Messaging**
Intersite Messaging is a component that is enabled when Active Directory is installed. Intersite Messaging allows for multiple transports to be used as add-ins to the Intersite Messaging architecture. Intersite Messaging enables messaging communication that can use SMTP servers other than those that are dedicated to processing e-mail applications.

When the forest has a functional level of at least Windows 2000, Intersite Messaging also provides services to the KCC in the form of querying the available replication paths. In addition, Net Logon queries the connectivity data in Intersite Messaging when calculating site coverage. By default, Intersite Messaging rebuilds its database once a day, or when required by a site link change.

When the forest has a functional level of at least Windows Server 2003, the KCC does not use Intersite Messaging for

calculating the topology. However, regardless of forest functional level, Intersite Messaging is still required for SMTP replication, DFS, universal group membership caching, and Net Logon automatic site coverage calculations. Therefore, if any of these features are in use, do not stop Intersite Messaging.

For more information about site coverage and how automatic site coverage is calculated, see "How DNS Support for Active Directory Works." For more information about DFS, see "DFS Technical Reference."

**Requirements for SMTP Replication**

The KCC does not create connections that use SMTP until the following requirements are met:

- Internet Information Services (IIS) is installed on both bridgehead servers.
- An enterprise certification authority (CA) is installed and configured on your network. The CA signs and encrypts SMTP messages that are exchanged between domain controllers, ensuring the authenticity of directory updates. Specifically, a domain controller certificate must be present on the replicating domain controllers. The replication request message, which contains no directory data, is not encrypted. The replication reply message, which does contain directory data, is encrypted using a key length of 128 bits.
- The sites are connected by SMTP site links.
- The site link path between the sites has a lower cost than any IP/RPC site link that can reach the SMTP site.
- You are not attempting to replicate writable replicas of the same domain (although replication of global catalog partial replicas is supported).
- Each domain controller is configured to receive mail.

You must also determine if mail routing is necessary. If the two replicating domain controllers have direct IP connectivity and can send mail to each other, no further configuration is required. However, if the two domain controllers must go through mail gateways to deliver mail to each other, you must configure the domain controller to use the mail gateway.

**Note**

- RPC is required for replicating the domain to a new domain controller and for installing certificates. If RPC is not available to the remote site, the domain must be replicated and certificates must be installed over RPC in a hub site and the domain controller then shipped to the remote site.

**Comparison of SMTP and RPC Replication**

The following characteristics apply to both SMTP and RPC with respect to Active Directory replication:

- For replication between sites, data that is replicated through either transport is compressed.
- Active Directory can respond with only a fixed (maximum) number of changes per change request, based on the size of the replication packet. The size of the replication packet is configurable. For information about configuring the replication packet size, see "Replication Packet Size" later in this section.
- Active Directory can apply a single set of changes at a time for a specific directory partition and replication partner.
- The response data (changes) are transported in one or many frames, based on the total number of changed or new values.
- TCP transports the data portion by using the same algorithm for both SMTP and RPC.
- If transmission of the data portion fails, complete retransmission is necessary.

Point-to-point synchronous RPC replication is available between sites to allow the flexibility of having domains that span multiple sites. RPC is best used between sites that are connected by WAN links because it involves lower latency. SMTP is best used between sites where RPC over IP is not possible. For example, SMTP can be used by companies that have a network

backbone that is not based on TCP/IP, such as companies that use an X.400 backbone.

Active Directory replication uses both transports to implement a request-response mechanism. Active Directory issues requests for changes and replies to requests for changes. RPC maps these requests into RPC requests and RPC replies. SMTP, on the other hand, actually uses long-lived TCP connections (or X.400-based message transfer agents in non-TCP/IP networks) to deliver streams of mail in each direction. Thus, RPC transport expects a response to any request immediately and can have a maximum of one active inbound RPC connection to a directory partition replica at a time. The SMTP transport expects much longer delays between a request and a response. As a result, multiple inbound SMTP connections to a directory partition replica can be active at the same time, provided the requests are all for a different source domain controller or, for the same source domain controller, a different directory partition. For more information, see "Synchronous and Asynchronous Communication" earlier in this section.

**Replication Packet Size**

Replication packet sizes are computed on the basis of memory size unless you have more than 1 gigabyte (GB). By default, the system limits the packet size as follows:

- The packet size in bytes is 1/100th the size of RAM, with a minimum of 1 MB and a maximum of 10 MB.
- The packet size in objects is 1/1,000,000th the size of RAM, with a minimum of 100 objects and a maximum of 1,000 objects. For general estimates when this entry is not set, assume an approximate packet size of 100 objects.

There is one exception: the value of the **Replicator async inter site packet size (bytes)** registry entry is always 1 MB if it is not set (that is, when the default value is in effect). Many mail systems limit the amount of data that can be sent in a mail message (2 MB to 4 MB is common), although most Windows-based mail systems can handle large 10-MB mail messages.

Overriding these memory-based values might be beneficial in advanced bandwidth management scenarios. You can edit the registry to set the maximum packet size.

**Note**

- If you must edit the registry, use extreme caution. Registry information is provided here as a reference for use by only highly skilled directory service administrators. It is recommended that you do not directly edit the registry unless, as in this case, there is no Group Policy or other Windows tools to accomplish the task. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. Storage of incorrect values can result in unrecoverable errors in the system.

Setting the maximum packet size requires adding or modifying entries in the following registry path with the REG_DWORD data type: HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters. These entries can be used to determine the maximum number of objects per packet and maximum size of the packets. The minimum values are indicated as the lowest value in the range.

- For RPC replication within a site:
    - **Replicator intra site packet size (objects)**

      Range: >=2
    - **Replicator intra site packet size (bytes)**

      Range: >=10 KB
- For RPC replication between sites:
    - **Replicator inter site packet size (objects)**

      Range: >=2
    - **Replicator inter site packet size (bytes)**

      Range: >=10 KB

- For SMTP replication between sites:

    - **Replicator async inter site packet size (objects)**

      Range: >=2

    - **Replicator async inter site packet size (bytes)**

      Range: >=10 KB

**Transport Type**
The **transportType** attribute of a connection object specifies which network transport is used when the connection is used for replication. The transport type receives its value from the distinguished name of the container in the configuration directory partition that contains the site link over which the connection occurs, as follows:

- Connection objects that use TCP/IP have the **transportType** value of CN=IP,CN=Inter-Site Transports,CN=IP,DC=Configuration,DC=*ForestRootDomainName*.

- Connection objects that use SMTP/IP have the **transportType** value of CN=SMTP,CN=Inter-Site Transports,CN=IP,DC=Configuration,DC=*ForestRootDomainName*.

- For intrasite connections, **transportType** has no value; Active Directory Sites and Services shows the transport of "RPC" for connections that are from servers in the same site.

If you move a domain controller to a different site, the connection objects from servers in the site from which it was moved remain, but the transport type is blank because it was an intrasite connection. Because the connection has an endpoint outside of the site, the local KCC in the server's new site does not manage the connection. When the ISTG runs, if a blank transport type is found for a connection that is from a server in a different site, the **transportType** value is automatically changed to IP. The ISTG in the site determines whether to delete the connection object or to retain it, in which case the server becomes a bridgehead server in its new site.

**Replication Between Sites**
Replication between sites transfers domain updates when domain controllers for a domain are located in more than one site. Intersite replication of configuration and schema changes is always required when more than one site is configured in a forest. Replication between sites is accomplished by bridgehead servers, which replicate changes according to site link settings.

**Bridgehead Servers**
When domain controllers for the same domain are located in different sites, at least one bridgehead server per directory partition and per transport (IP or SMTP) replicates changes from one site to a bridgehead server in another site. A single bridgehead server can serve multiple partitions per transport and multiple transports. Replication within the site allows updates to flow between the bridgehead servers and the other domain controllers in the site. Bridgehead servers help to ensure that the data replicated across WAN links is not stale or redundant.

Any server that has a connection object with a "from" server in another site is acting as a destination bridgehead. Any server that is acting as a source for a connection to another site acts as a source bridgehead.

### Note

- You can identify a KCC-selected bridgehead server in Active Directory Sites and Services by viewing connection objects for the server (select the NTDS Settings object below the server object); if there are connections from servers in a different site or sites, the server represented by the selected NTDS Settings object is a bridgehead server. If you have Windows Support Tools installed, you can see all bridgehead servers by using the command **repadmin /bridgeheads**.

KCC selection of bridgehead servers guarantees bridgehead servers that are capable of replicating all directory partitions that are needed in the site, including partial global catalog partitions. By default, bridgehead servers are selected automatically by the KCC on the domain controller that holds the ISTG role in each site. If you want to identify the domain controllers that can act as bridgehead servers, you can designate preferred bridgehead servers, from which the ISTG selects all bridgehead servers. Alternatively, if the ISTG is not used to generate the intersite topology, you can create manual intersite connection objects on

domain controllers to designate bridgehead servers.

In sites that have at least one domain controller that is running Windows Server 2003, the ISTG can select bridgehead servers from all eligible domain controllers for each directory partition that is represented in the site. For example, if three domain controllers in a site store replicas of the same domain and domain controllers for this domain are also located in three or more other sites, the ISTG can spread the inbound connection objects from those sites among all three domain controllers, including those that are running Windows 2000 Server.

In Windows 2000 forests, a single bridgehead server per directory partition and per transport is designated as the bridgehead server that is responsible for intersite replication of that directory partition. Therefore, for the preceding example, only one of the three domain controllers would be designated by the ISTG as a bridgehead server for the domain, and all four connection objects from the four other sites would be created on the single bridgehead server. In large hub sites, a single domain controller might not be able to adequately respond to the volume of replication requests from perhaps thousands of branch sites.

For more information about how the KCC selects bridgehead servers, see “Bridgehead Server Selection” later in this section.

**Compression of Replication Data**
Intersite replication is compressed by default. Compressing replication data allows the data to be transferred over WAN links more quickly, thereby conserving network bandwidth. The cost of this benefit is an increase in CPU utilization on bridgehead servers.

By default, replication data is compressed under the following conditions:

- Replication of updates between domain controllers in different sites.
- Replication of Active Directory to a newly created domain controller.

A new compression algorithm is employed by bridgehead servers that are running at least Windows Server 2003. The new algorithm improves replication speed by operating between two and ten times faster than the Windows 2000 Server algorithm.

**Windows 2000 Server Compression**
The compression algorithm that is used by domain controllers that are running Windows 2000 Server achieves a compression ratio of approximately 75% to 85%. The cost of this compression in terms of CPU utilization can be as high as 50% for intersite Active Directory replication. In some cases, the CPUs on bridgehead servers that are running Windows 2000 Server can become overwhelmed with compression requests, compounded by the need to service outbound replication partners. In a worst case scenario, the bridgehead server becomes so overloaded that it cannot keep up with outbound replication. This scenario is usually coupled with a replication topology issue where a domain controller has more outbound partners than necessary or the replication schedule was overly aggressive for the number of direct replication partners.

### Note

- If a bridgehead server has too many replication partners, the KCC logs event ID 1870 in the Directory Service log, indicating the current number of partners and the recommended number of partners for the domain controller.

**Windows Server 2003 Compression**
On domain controllers that are running Windows Server 2003, compression quality is comparable to Windows 2000 but the processing burden is greatly decreased. The Windows Server 2003 algorithm produces a compression ratio of approximately 60%, which is slightly less compression than is achieved by the Windows 2000 Server ratio, but which significantly reduces the processing load on bridgehead servers. The new compression algorithm provides a good compromise by significantly reducing the CPU load on bridgehead servers, while only slightly increasing the WAN traffic. The new algorithm reduces the time taken by compression from approximately 60% of replication time to 20%.

The Windows Server 2003 compression algorithm is used only when both bridgehead servers are running Windows Server 2003. If a bridgehead server that is running Windows Server 2003 replicates with a bridgehead server that is running Windows 2000 Server, then the Windows 2000 compression algorithm is used.

**Reverting to Windows 2000 Compression**
For slow WAN links (for example, 64 KB or less), if more compression is preferable to a decrease in computation time, you can change the compression algorithm to the Windows 2000 algorithm. The compression algorithm is controlled by the REG_DWORD registry entry HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters**Replicator compression algorithm**. By editing this registry entry, you can change the algorithm that is used for compression to the

Windows 2000 algorithm.

**Note**

- If you must edit the registry, use extreme caution. Registry information is provided here as a reference for use by only highly skilled directory service administrators. It is recommended that you do not directly edit the registry unless, as in this case, there is no Group Policy or other Windows tools to accomplish the task. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. Storage of incorrect values can result in unrecoverable errors in the system.

The default value is 3, which indicates that the Windows Server 2003 algorithm is in effect. By changing the value to 2, the Windows 2000 algorithm is used for compression. However, switching to the Windows 2000 algorithm is not recommended unless both bridgehead domain controllers serve relatively few branches and have ample CPU (for example, > dual processor 850 megahertz [MHz]).

**Site Link Settings and Their Effects on Intersite Replication**

In Active Directory Sites and Services, the **General** tab of the site link **Properties** contains the following options for configuring site links to control the replication topology:

- A list of two or more sites to be connected.

- A single numeric cost that is associated with communication over the link. The default cost is 100, but you can assign higher cost values to represent more expensive transmission. For example, sites that are connected by low-speed or dial-up connections would have high-cost site links between them. Sites that are well connected through backbone lines would have low-cost site links. Where multiple routes or transports exist between two sites, the least expensive route and transport combination is used.

- A schedule that determines days and hours during which replication can occur over the link (the link is available). For example, you might use the default (100 percent available) schedule on most links, but block replication traffic during peak business hours on links to certain branches. By blocking replication, you give priority to other traffic, but you also increase replication latency.

  **Note**

  - Scheduling information is ignored by site links that use SMTP transports; the mail is stockpiled and then exchanged at the times that are configured for your mail infrastructure.

- An interval in minutes that determines how often replication can occur (default is every 180 minutes, or 3 hours). The minimum interval is 15 minutes. If the interval exceeds the time allowed by the schedule, replication occurs once at the scheduled time.

A site can be connected to other sites by any number of site links. For example, a hub site has site links to each of its branch sites. Each site that contains a domain controller in a multisite directory must be connected to at least one other site by at least one site link; otherwise, it cannot replicate with domain controllers in any other site.

The following diagram shows two sites that are connected by a site link. Domain controllers DC1 and DC2 belong to the same domain and are acting as partner bridgehead servers. When topology generation occurs, the ISTG in each site creates an inbound connection object on the bridgehead server in its site from the bridgehead server in the opposite site. With these objects in place, replication can occur according to the settings on the SB site link.

**Connections Between Domain Controllers in Two Sites that Are Connected by a Site Link**




**Site Link Cost**

The ISTG uses the cost settings on site links to determine the route of replication between three or more sites that replicate the same directory partition. The default cost value on a site link object is 100. You can assign lower or higher cost values to site links to favor inexpensive connections over expensive connections, respectively. Certain applications and services, such as

domain controller Locator and DFS, also use site link cost information to locate nearest resources. For example, site link cost can be used to determine which domain controller is contacted by clients located in a site that does not include a domain controller for the specified domain. The client contacts the domain controller in a different site according to the site link that has the lowest cost assigned to it.

Cost is usually assigned not only on the basis of the total bandwidth of the link, but also on the availability, latency, and monetary cost of the link. For example, a 128-kilobits per second (Kbps) permanent link might be assigned a lower cost than a dial-up 128-Kbps dual ISDN link because the dial-up ISDN link has replication latency-producing delay that occurs as the links are being established or removed. Furthermore, in this example, the permanent link might have a fixed monthly cost, whereas the ISDN line is charged according to actual usage. Because the company is paying up-front for the permanent link, the administrator might assign a lower cost to the permanent link to avoid the extra monetary cost of the ISDN connections.

The method used by the ISTG to determine the least-cost path from each site to every other site for each directory partition is more efficient when the forest has a functional level of at least Windows Server 2003 than it is at other levels. For more information about how the KCC computes replication routes, see "Automated Intersite Topology Generation" later in this section. For more information about domain controller location, see "How DNS Support for Active Directory Works."

**Transitivity and Automatic Site Link Bridging**

By default, site links are transitive, or "bridged." If site A has a common site link with site B, site B also has a common site link with Site C, and the two site links are bridged, domain controllers in site A can replicate directly with domain controllers in site C under certain conditions, even though there is no site link between site A and site C. In other words, the effect of bridged site links is that replication between sites in the bridge is transitive.

The setting that implements automatic site link bridges is **Bridge all site links**, which is found in Active Directory Sites and Services in the properties of the IP or SMTP intersite transport containers. The default bridging of site links occurs automatically and no directory object represents the default bridge. Therefore, in the common case of a fully routed IP network, you do not need to create any site link bridge objects.

**Transitivity and Rerouting**

For a set of bridged site links, where replication schedules in the respective site links overlap (replication is available on the site links during the same time period), connection objects can be automatically created, if needed, between sites that do not have site links that connect them directly. All site links for a specific transport implicitly belong to a single site link bridge for that transport.

Site link transitivity enables the KCC to re-route replication when necessary. In the next diagram, a domain controller that can replicate the domain is not available in Seattle. In this case, because the site links are transitive (bridged) and the schedules on the two site links allow replication at the same time, the KCC can re-route replication by creating connections between DC3 in Portland and DC2 in Boston. Connections between domain controllers in Portland and Boston might also be created when a domain controller in Portland is a global catalog server, but no global catalog server exists in the Seattle site and the Boston site hosts a domain that is not present in the Seattle site. In this case, connections can be created between Portland and Boston to replicate the global catalog partial, read-only replica.

### Note

- Overlapping schedules are required for site link transitivity, even when **Bridge all site links** is enabled. In the example, if the site link schedules for SB and PS do not overlap, no connections are possible between Boston and Portland.

**Transitive Replication when Site Links Are Bridged, Schedules Overlap, and Replication Must Be Rerouted**




In the preceding diagram, creating a third site link to connect the Boston and Portland sites is unnecessary and counterproductive because of the way that the KCC uses cost to route replication. In the configuration that is shown, the KCC uses cost to choose either the route between Portland and Seattle or the route between Portland and Boston. If you wanted the KCC to use the route between Portland and Boston, you would create a site link between Portland and Boston instead of the site link between Portland and Seattle.

**Aggregated Site Link Cost and Routing**
When site links are bridged, the cost of replication from a domain controller at one end of the bridge to a domain controller at the other end is the sum of the costs on each of the intervening site links. For this reason, if a domain controller in an interim site stores the directory partition that is being replicated, the KCC will route replication to the domain controller in the interim site rather than to the more distant site. The domain controller in the more distant site in turn receives replication from the interim site (store-and-forward replication). If the schedules of the two site links overlap, this replication occurs in the same period of replication latency.

The following diagram illustrates an example where two site links connecting three sites that host the same domain are bridged automatically (**Bridge all site links** is enabled). The aggregated cost of directly replicating between Portland and Boston illustrates why the KCC routes replication from Portland to Seattle and from Seattle to Boston in a store-and-forward manner. Given the choice between replicating at a cost of 4 from Seattle or a cost of 7 from Boston, the ISTG in Portland chooses the lower cost and creates the connection object on DC3 from DC1 in Seattle.

**Bridged Site Links Routing Replication Between Three Sites According to Cost**




In the preceding diagram, if DC3 in Portland needs to replicate a directory partition that is hosted on DC2 in Boston but not by any domain controller in Seattle, or if the directory partition is hosted in Seattle but the Seattle site cannot be reached, the ISTG creates the connection object from DC2 to DC3.

**Significance of Overlapping Schedules**
In the preceding diagram, to replicate the same domain that is hosted in all three sites, the Portland site replicates directly with Seattle and Seattle replicates directly with Boston, transferring Portland's changes to Boston, and vice versa, through store-and-forward replication. Whether the schedules overlap has the following effects:

- PS and SB site link schedules have replication available during at least one common hour of the schedule:
    - Replication between these two sites occurs in the same period of replication latency, being routed through Seattle.
    - If Seattle is unavailable, connections can be created between Portland and Boston.
- PS and SB site link schedules have no common time:
    - Replication of changes between Portland and Boston reach their destination in the next period of replication latency after reaching Seattle.
    - If Seattle is unavailable, no connections are possible between Portland and Boston.

    **Note**

If **Bridge all site links** is disabled, a connection is never created between Boston and Portland, regardless of schedule overlap, unless you manually create a site link bridge.

**Site Link Changes and Replication Path**
The path that replication takes between sites is computed from the information that is stored in the properties of the site link objects. When a change is made to a site link setting, the following events must occur before the change takes effect:

- The site link change must replicate to the ISTG of each site by using the previous replication topology.

- The KCC must run on each ISTG.

As the path of connections is transitively figured through a set of site links, the attributes (settings) of the site link objects are combined along the path as follows:

- Costs are added together.
- The replication interval is the maximum of the intervals that are set for the site links along the path.
- The options, if any are set, are computed by using the AND operation.

  **Note**

  - Options are the values of the **options** attribute on the site link object. The value of this attribute determines special behavior of the site link, such as reciprocal replication and intersite change notification.

Thus the site link schedule is the overlap of all of the schedules of the subpaths. If none of the schedules overlap, the path is not used.

**Bridging Site Links Manually**

If your IP network is composed of IP segments that are not fully routed, you can disable **Bridge all site links** for the IP transport. In this case, all IP site links are considered nontransitive, and you can create and configure site link bridge objects to model the actual routing behavior of your network. A site link bridge has the effect of providing routing for a disjoint network (networks that are separate and unaware of each other). When you add site links to a site link bridge, all site links within the bridge can route transitively.

A site link bridge object represents a set of site links, all of whose sites can communicate through some transport. Site link bridges are necessary if both of the following conditions are true:

- A site contains a domain controller that hosts a domain directory partition that is not hosted by a domain controller in an adjacent site (a site that is in the same site link).
- That domain directory partition is hosted on a domain controller in at least one other site in the forest.

**Note**

- Site link bridge objects are used by the KCC only when the **Bridge all site links** setting is disabled. Otherwise, site link bridge objects are ignored.

Site link bridges can also be used to diminish potentially high CPU overhead of generating a large transitive replication topology. In very large networks, transitive site links can be an issue because the KCC considers every possible connection in the bridged network, and selects only one. Therefore, in a Windows 2000 forest that has a very large network or a Windows Server 2003 or higher forest that consists of an extremely large hub-and-spoke topology, you can reduce KCC-related CPU utilization and run time by turning off **Bridge all site links** and creating manual site link bridges only where they are required.

**Note**

- Turning off **Bridge all site links** affects the ability of DFS clients to locate DFS servers in the closest site. If the ISTG is running at least Windows Server 2003, the **Bridge all site links** must be enabled to generate the intersite cost matrix that DFS requires for its site-costing functionality. If the ISTG is running at least Windows Server 2003 with Service Pack 1 (SP1), you can enable **Bridge all site links** and then run the **repadmin /siteoptions W2K3_BRIDGES_REQUIRED** command on each site where you need to accommodate the DFS site-costing functionality. This command disables automatic site link bridging for the KCC but allows default Intersite Messaging options to enable the site-costing calculation to occur for DFS. For more information about turning off this functionality while accommodating DFS, see "DFS Site Costing and Windows Server 2003 SP1 Site Options" later in this section. For more information about site link cost and DFS, see "DFS Technical Reference."

You create a site link bridge object for a specific transport by specifying two or more site links for the specified transport.

**Requirements for manual site link bridges**
Each site link in a manual site link bridge must have at least one site in common with another site link in the bridge. Otherwise, the bridge cannot compute the cost from sites in one link to the sites in other links of the bridge. If bridgehead servers that are capable of the transport that is used by the site link bridge are not available in two linked sites, a route is not available.

**Manual site link bridge behavior**
Separate site link bridges, even for the same transport, are independent. To illustrate this independence, consider the following conditions:

- Four sites have domain controllers for the same domain: Portland, Seattle, Detroit, and Boston.
- Three site links are configured: Portland-Seattle (PS), Seattle-Detroit (SD), and Detroit-Boston (DB).
- Two separate manual site link bridges link the outer site links PS and DB with the inner site link SD.

The presence of the PS-SD site link bridge means that an IP message can be sent transitively from the Portland site to the Detroit site with cost 4 + 3 = 7. The presence of the SD-DB site link bridge means that an IP message can be sent transitively from Seattle to Boston at a cost of 3 + 2 = 5. However, because there is no transitivity between the PS-SB and SB-DB site link bridges, an IP message cannot be sent between Portland and Boston with cost 4 + 3 + 2 = 9, or at any cost.

In the following diagram, the two manual site link bridges means that Boston is able to replicate directly only with Detroit and Seattle, and Portland is able to replicate directly only with Seattle and Detroit.

**Note**

- If you need direct replication between Portland and Detroit, you can create the PS-SB-DB site link bridge. By excluding the PS site link, you ensure that connections are neither created nor considered by the KCC between Portland and Detroit.

**Two Site Link Bridges that Are Not Transitive**




In the diagram, connection objects are not possible between DC4 in Detroit and DC3 in Portland because two site link bridges are not transitive. For connection objects to be possible between DC3 and DC4, the site link DB must be added to the PS-SB site link bridge. In this case, the cost of replication between DC3 and DC4 is 9.

**Note**

- Cost is applied differently to a site link bridge than to a site link that contains more than two sites. To use the preceding example, if Seattle, Boston, and Portland are all in the same site link, the cost of replication between any of the two sites is the same.

Bridging site links manually is generally recommended for only large branch office deployments. For more information about using manual site link bridging, see the "Windows Server 2003 Active Directory Branch Office Deployment Guide."

**Site Link Schedule**
Replication using the RPC transport between sites is scheduled. The schedule specifies one or many time periods during which replication can occur. For example, you might schedule a site link for a dial-up line to be available during off-peak hours (when telephone rates are low) and unavailable during high-cost regular business hours. The **schedule** attribute of the site link object specifies the availability of the site link. The default setting is that replication is always available.

**Note**

- The **Ignore schedules** setting on the IP container is equivalent to replication being always available. If **Ignore schedules** is selected, replication occurs at the designated intervals but ignores any schedule.

If replication goes through multiple site links, there must be at least one common time period (overlap) during which replication is available; otherwise, the connection is treated as not available. For example, if site link AB has a schedule of 18:00 hours to 24:00 hours and site link BC has a schedule of 17:00 hours to 20:00 hours, the resulting overlap is 18:00 hours through 20:00 hours, which is the intersection of the schedules for site link AB and site link BC. During the time in which the schedules overlap, replication can occur from site A to site C even if a domain controller in the intermediate site B is not available. If the schedules do not overlap, replication from the intermediate site to the distant site continues when the next replication schedule opens on the respective site link.

**Note**

- Cost considerations also affect whether connections are created. However, if the site link schedules do not overlap, the cost is irrelevant.

**Scheduling across time zones**

When scheduling replication across time zones, consider the time difference to ensure that replication does not interfere with peak production times in the destination site.

Domain controllers store time in Coordinated Universal Time (UTC). When viewed through the Active Directory Sites and Services snap-in, time settings in site link object schedules are displayed according to the local time of the computer on which the snap-in is being run. However, replication occurs according to UTC.

For example, suppose Seattle adheres to Pacific Standard Time (PST) and Japan adheres to Japan Standard Time (JST), which is 17 hours later. If a schedule is set on a domain controller in Seattle and the site link on which the schedule is set connects Seattle and Tokyo, the actual time of replication in Tokyo is 17 hours later.

If the schedule is set to begin replication at 10:00 PM PST in Seattle, the conversion can be computed as follows:

- Convert 10:00 PM PST to 22:00 PST military time.
- Add 8 hours to arrive at 06:00 UTC, the following day.
- Add 9 hours to arrive at 15:00 JST.
- 15:00 JST converts to 3:00 PM.

Thus, when replication begins at 10:00 o'clock at night in Seattle, it is occurring in Tokyo at 3:00 o'clock in the afternoon the following day. By scheduling replication a few hours later in Seattle, you can avoid replication occurring during working hours in Japan.

**Schedule implementation**

The times that you can set in the **Schedule** setting on the site link are in one-hour increments. For example, you can schedule replication to occur between 00:00 hours and 01:00 hours, between 01:00 hours and 02:00 hours, and so forth. However, each block in the actual connection schedule is 15 minutes. For this reason, when you set a schedule of 01:00 hours to 02:00 hours, you can assume that replication is queued at some point between 01:00 hours and 01:14:59 hours.

**Note**

- RPC synchronous inbound replication is serialized so that if the server is busy replicating this directory partition from another source, replication from a different source does not begin until the first synchronization is complete. SMTP asynchronous replication is processed serially by order of arrival, with multiple replication requests queued simultaneously.

Specifically, a replication event is queued at time $t + n$, where $t$ is the replication interval that is applied across the schedule and $n$ is a pseudo-random number from 1 minute through 15 minutes. For example, if the site link indicates that replication can occur from 02:00 hours through 07:00 hours, and the replication interval is 2 hours (120 minutes), $t$ is 02:00 hours, 04:00 hours, and 06:00 hours. A replication event is queued on the destination domain controller between 02:00 hours and 02:14:59 hours, and another replication event is queued between 04:00 hours and 04:14:59 hours. Assuming that the first

replication event that was queued is complete, another replication event is queued between 06:00 hours and 06:14:59 hours. If the synchronization took longer than two hours, the second synchronization would be ignored because an event is already in the queue.

Replication can extend beyond the end of the schedule. A period of replication latency that starts before the end of the schedule runs until completion, even if the period is still running when the schedule no longer allows replication to be available.

**Note**

- The replication queue is shared with other events, and the time at which replication takes place is approximate. Duplicate replication events are not queued for the same directory partition and transport.

**Connection object schedule**

Each connection object has a schedule that controls when (during what hours) and how frequently (how many times per hour) replication can occur:

- None (no replication)
- Once per hour (the default setting)
- Twice per hour
- Four times per hour

The connection object schedule and interval are derived from one of two locations, depending on whether it is an intrasite or intersite connection:

- Intrasite connections inherit a default schedule from the **schedule** attribute of the NTDS Site Settings object. By default, this schedule is always available and has an interval of one hour.
- Intersite connections inherit the schedule and interval from the site link.

Although intrasite replication is prompted by changes, intrasite connection objects inherit a default schedule so that replication occurs periodically, regardless of whether change notification has been received. The connection object schedule ensures that intrasite replication occurs if a notification message is lost, or if notification does not take place, due to network problems or a domain controller becomes unavailable. The NTDS Site Settings schedule has a minimum replication interval of 15 minutes. This minimum replication interval is not configurable and determines the smallest interval that is possible for both intrasite and intersite replication (on a connection object or a site link, respectively).

For intersite replication, the schedule is configured on the site link object, but the connection object schedule actually determines replication; that is, the connection object schedule for an intersite connection is derived from the site link schedule, which is applied through the connection object schedule. Scheduled replication occurs independently of change notification.

**Note**

- You do not need to configure the connection object schedule unless you are creating a manual intersite replication topology that does not use the KCC automatic connection objects.

The KCC uses a two-step process to compute the schedule of an intersite connection.

1. The schedules of the site links traversed by a connection are merged together.
2. This merged schedule is modified so that it is available at only certain periods. The length of those periods is equal to the maximum replication interval of the site links traversed by this connection.

By using Active Directory Sites and Services, you can manually revise the schedule on a connection object, but such an override is effective for only administrator-owned connection objects.

**Replication Interval**
For each site link object, you can specify a value for the replication interval (frequency), which determines how often replication occurs over the site link during the time that the schedule allows. For example, if the schedule allows replication between 02:00 hours and 04:00 hours, and the replication interval is set for 30 minutes, replication can occur up to four times during the scheduled time.

The default replication interval is 180 minutes, or 3 hours. When the KCC creates a connection between a domain controller in one site and a domain controller in another site, the replication interval of the connection is the maximum interval along the minimum-cost path of site link objects from one end of the connection to the other.

**Interaction of Replication Schedule and Interval**
When multiple site links are required to complete replication for all sites, the replication interval settings on each site link combine to affect the entire length of the connection between sites. In addition, when schedules on each site link are not identical, replication can occur only when the schedules overlap.

Suppose that site A and site B have site link AB, and site B and site C have site link BC. When a domain controller in site A replicates with a domain controller in site C, it can do so only as often as the maximum interval that is set for site link AB and site link BC allows. The following table shows the site link settings that determine how often and during what times replication can occur between domain controllers in site A, site B, and site C.

**Replication Interval and Schedule Settings for Two Site Links**

| Site Link | Replication Interval | Schedule |
|---|---|---|
| AB | 30 minutes | 12:00 hours to 04:00 hours |
| BC | 60 minutes | 01:00 hours to 05:00 hours |

Given these settings, a domain controller in site A can replicate with a domain controller in site B according to the AB site link schedule and interval, which is once every 30 minutes between the hours of 12:00 and 04:00. However, assuming that there is no site link AC, a domain controller in site A can replicate with a domain controller in site C between the hours of 01:00 and 04:00, which is where the schedules on the two site links intersect. Within that timespan, they can replicate once every 60 minutes, which is the greater of the two replication intervals.

**KCC and Topology Generation**
The Knowledge Consistency Checker (KCC) is a dynamic-link library (DLL) that runs as a distributed application on every domain controller. The KCC on each domain controller modifies data in its local instance of the directory in response to forest-wide changes, which are made known to the KCC by changes to data in the configuration directory partition.

The KCC generates and maintains the replication topology for replication within sites and between sites by converting KCC-defined and administrator-defined (if any) connection objects into a configuration that is understood by the directory replication engine. By default, the KCC reviews and makes modifications to the Active Directory replication topology every 15 minutes to ensure propagation of data, either directly or transitively, by creating and deleting connection objects as needed. The KCC recognizes changes that occur in the environment and ensures that domain controllers are not orphaned in the replication topology.

Operating independently, the KCC on each domain controller uses its own view of the local replica of the configuration directory partition to arrive at the same intrasite topology. One KCC per site, the ISTG, determines the intersite replication topology for the site. Like the KCC that runs on each domain controller within a site, the instances of the ISTG in different sites do not communicate with each other. They independently use the same algorithm to produce a consistent, well-formed spanning tree of connections. Each site constructs its own part of the tree and, when all have run, a working replication topology exists across the enterprise.

The predictability of all KCCs allows scalability by reducing communication requirements between KCC instances. All KCCs agree on where connections will be formed, ensuring that redundant replication does not occur and that all parts of the enterprise are connected.

The KCC performs two major functions:

- Configures appropriate replication connections (connection objects) on the basis of the existing cross-reference, server, NTDS settings, site, site link, and site link bridge objects and the current status of replication.

- Converts the connection objects that represent inbound replication to the local domain controller into the replication agreements that are actually used by the replication engine. These agreements, called replica links, accommodate replication of a single directory partition from the source to the destination domain controller.

#### Intervals at Which the KCC Runs

By default, the KCC runs its first replication topology check five minutes after the domain controller starts. The domain controller then attempts initial replication with its intrasite replication partners. If a domain controller is being used for multiple other services, such as DNS, WINS, or DHCP, extending the replication topology check interval can ensure that all services have started before the KCC begins using CPU resources.

You can edit the registry to modify the interval between startup and the time the domain controller first checks the replication topology.

**Note**

- If you must edit the registry, use extreme caution. Registry information is provided here as a reference for use by only highly skilled directory service administrators. It is recommended that you do not directly edit the registry unless, as in this case, there is no Group Policy or other Windows tools to accomplish the task. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. Storage of incorrect values can result in unrecoverable errors in the system.

Modifying the interval between startup and the time the domain controller first checks the replication topology requires changing the **Repl topology update delay (secs)** entry in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters as appropriate:

- **Value:** Number of seconds to wait between the time Active Directory starts and the KCC runs for the first time.
- **Default:** 300 seconds (5 minutes)
- **Data type:** REG_DWORD

Thereafter, as long as services are running, the KCC on each domain controller checks the replication topology every 15 minutes and makes changes as necessary.

Modifying the interval at which the KCC performs topology review requires changing the **Repl topology update period (secs)** entry in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters as appropriate:

- **Value:** Number of seconds between KCC topology updates
- **Default:** 900 seconds (15 minutes)
- **Data type:** REG_DWORD

#### Objects that the KCC Requires to Build the Replication Topology

The following objects, which are stored in the configuration directory partition, provide the information required by the KCC to create the replication topology:

- **Cross-reference.** Each directory partition in the forest is identified in the Partitions container by a cross-reference object. The attributes of this object are used by the replication system to locate the domain controllers that store each directory partition.
- **Server.** Each domain controller in the forest is identified as a server object in the Sites container.
- **NTDS Settings.** Each server object that represents a domain controller has a child NTDS Settings object. Its presence identifies the server as having Active Directory installed. The NTDS Settings object must be present for the server to be considered by the KCC for inclusion in the replication topology.
- **Site.** The presence of the above objects also indicates to the KCC the site in which each domain controller is located for replication. For example, the distinguished name of the NTDS Settings object contains the name of the site in which the server object that represents the domain controller exists.

- **Site link.** A site link must be available between any set of sites and its schedule and cost properties evaluated for routing decisions.

- **Site link bridge.** If they exist, site link bridge objects and properties are evaluated for routing decisions.

If the domain controller is physically located in one site but its server object is configured in a different site, the domain controller will attempt intrasite replication with a replication partner that is in the site of its server object. In this scenario, the improper configuration of servers in sites can affect network bandwidth.

If a site object exists for a site that has no domain controllers, the KCC does not consider the site when generating the replication topology.

**Topology Generation Phases**

The KCC generates the replication topology in two phases:

- **Evaluation.** During the evaluation phase, the KCC evaluates the current topology, determines whether replication failures have occurred with the existing connections, and constructs whatever new connection objects are required to complete the replication topology.

- **Translation.** During the translation phase, the KCC implements, or "translates," the decisions that were made during the evaluation phase into agreements between the replication partners. During this phase, the KCC writes to the **repsFrom** attribute on the local domain controller (for intrasite topology) or on all bridgehead servers in a site (for intersite topology) to identify the replication partners from which each domain controller pulls replication. For more information about the information in the replication agreement, see "How the Active Directory Replication Model Works."

**KCC Modes and Scopes**

Because individual KCCs do not communicate directly to generate the replication topology, topology generation occurs within the scope of either a single domain controller or a single site. In performing the two topology generation phases, the KCC has three modes of operation. The following table identifies the modes and scope for each mode.

**Modes and Scopes of KCC Topology Generation**

| KCC Mode | Performing Domain Controllers | Scope | Description |
|---|---|---|---|
| Intrasite | All | Local server | Evaluate all servers in a site and create connection objects locally on this server from servers in the same site that are adjacent to this server in the ring topology. |
| Intersite | One domain controller per site that has the ISTG role | Local site | Evaluate the servers in all sites and create connection objects both locally and on other servers in the site from servers in different sites. |
| Link translation | All | Local server | Translate connection objects into replica links (partnerships) for each server relative to each directory partition that it holds. |

**Topology Evaluation and Connection Object Generation**

The KCC on a destination domain controller evaluates the topology by reading the existing connection objects. For each connection object, the KCC reads attribute values of the NTDS Settings object (class **nTDSDSA**) of the source domain controller (indicated by the **fromServer** value on the connection object) to determine what directory partitions its destination domain controller has in common with the source domain controller.

**Topology evaluation for all domain controllers**

To determine the connection objects that need to be generated, the KCC uses information stored in the attributes of the NTDS Settings object that is associated with each server object, as follows:

- For all directory partitions, the multivalued attribute **hasMasterNCs** stores the distinguished names of all directory partitions that are stored on that domain controller.

- For all domain controllers, the value of the **options** attribute indicates whether that domain controller is configured to host the global catalog.

- The **hasPartialReplicaNCs** attribute contains the set of partial-replica directory partitions (global catalog read-only

domain partitions) that are located on the domain controller that is represented by the server object.

**Topology evaluation for domain controllers running Windows Server 2003**

For all domain controllers that are running Windows Server 2003, the **msDS-HasDomainNCs** attribute of the NTDS Settings object contains the name of the domain directory partition that is hosted by the domain controller.

In forests that have the forest functional level of at least Windows Server 2003 or Windows Server 2003 interim, the following additional information is used by the KCC to evaluate the topology for application directory partitions and to generate the needed connections:

- The linked multivalued attribute **msDS-NC-Replica-Locations** on cross-reference objects stores the distinguished names of NTDS Settings objects for all domain controllers that are configured to host a replica of the corresponding application directory partition.

  **Note**

  - When you remove Active Directory from a server that hosts an application directory partition, its corresponding entry in this multivalued attribute is automatically dropped because **msDS-NC-Replica-Locations** is a linked attribute.

- Application directory partition replica locations are determined by matching the values of the **hasMasterNCs** attribute with the values of the **msDS-NC-Replica-Locations** linked multivalued attribute of cross-reference objects. The **msDS-NC-Replica-Locations** attribute holds distinguished name references to the NTDS Settings objects for domain controllers that have been configured to store replicas of the application directory partition. The **msDS-NC-Replica-Locations** attribute facilitates the enumeration of *existing* replicas for a given application directory partition. Connection objects can then be created between the domain controllers that hold matching replicas.

Be aware that due to replication latency, the configuration of replicas in attribute values does not guarantee the existence of the replica on a given server. For example, you can designate a domain controller as a global catalog server by clicking the **Global Catalog** check box on the NTDS Settings object properties in Active Directory Sites and Services. However, until all of the partial domain directory partitions have replicated to that domain controller and global-catalog-specific SRV records are registered in DNS, it is not a functioning global catalog server (does not advertise as a global catalog server in DNS). Similarly, observing the NTDS Settings name for a server in the **msDS-NC-Replica-Locations** attribute on the cross-reference object does not indicate that the replica has necessarily been fully replicated to that server.

**Connection Translation**

All KCCs process their connection objects and translate them into connection agreements, also called "replica links," between pairs of domain controllers. At specified intervals, Active Directory replicates data from the source domain controller to the destination for directory partitions that they have in common. These replication agreements do not appear in the administrative tools; the replication engine uses them internally to track the directory partitions that are to be replicated from specified servers.

For each directory partition that two domain controllers have in common and that matches the full and partial characteristics of a replication source, the KCC creates (or updates) a replication agreement on the destination domain controller. Replication agreements take the form of entries for each source domain controller in the **repsFrom** attribute on the topmost object of each directory partition replica. This value is stored and updated locally on the domain controller and is not replicated. The KCC updates this attribute each time it runs.

For example, suppose a connection object is created between two domain controllers from different domains. Assuming that neither of these domain controllers is a global catalog server and neither stores an application directory partition, the KCC identifies the only two directory partitions that the domain controllers have in common — the schema directory partition and the configuration directory partition. If a connection object links domain controllers in the same domain, at least three directory partitions are replicated: the schema directory partition, the configuration directory partition, and the domain directory partition.

In contrast, if the connection object that is created establishes replication between two domain controllers that are global catalog servers, then in addition to the directory partitions the domain controllers have in common, a partial replica of each additional domain directory partition in the forest is also replicated between the two domain controllers over the same connection.

For more information about replication agreements, see "How the Active Directory Replication Model Works."

**Read-only and Writable Replicas**
When computing the replication topology, the KCC must consider whether a replica is writable or read-only. For each potential set of replication partners in the topology, the considerations are as follows:

- A writable replica can receive updates from a corresponding writable replica.
- A read-only replica can receive updates from a corresponding writable replica.
- A read-only replica can receive updates from a corresponding read-only replica.
- A writable replica *cannot* receive updates from a corresponding read-only replica.

In Windows 2000 forests, for any one domain directory partition, the KCC calculates two topologies: one for the writable replicas and one for the read-only replicas. This calculation allows redundant connections for read-only replicas under certain conditions.

The improved KCC spanning tree algorithm eliminates redundancy that can occur in Windows 2000. The algorithm computes only one topology with slightly different behavior for replicating the global catalog. The KCC on a domain controller that is not a global catalog server does not consider global catalog servers in its calculations for read-only domain replicas because it never replicates read-only data *from* a global catalog server.

**Automated Intrasite Topology Generation**
For replication within a site, a topology is generated and then optimized to minimize the number of hops to three. The means by which the three-hop minimum is achieved varies according to the number of domain controllers that are hosted in the site as well as the presence of global catalog servers. Generally, the intrasite topology is formed in a ring. The topology becomes more complex as the number of servers increases. However, the KCC can accommodate thousands of domain controllers in a site.

**Simplified Ring Topology Generation**
A simplified process for creating the topology for replication within a site begins as follows:

- The KCC generates a list of all servers in the site that hold that directory partition.
- These servers are connected in a ring.
- For each neighboring server in the ring from which the current domain controller is to replicate, the KCC creates a connection object if one does not already exist.

This simple approach guarantees a topology that tolerates a single failure. If a domain controller is not available, it is not included in the ring that is generated by the list of servers. However, this topology, with no other adjustments, accommodates only seven servers. Beyond this number, the ring would require more than three hops for some servers.

The simplest case scenario — seven or fewer domain controllers, all in the same domain and site — would result in the replication topology shown in the following diagram. The only directory partitions to replicate are a single domain directory partition, the schema directory partition, and the configuration directory partition. Those topologies are generated first, and at that point, sufficient connections to replicate each directory partition have already been created.

In the next series of diagrams, the arrows indicate one-way or two-way replication of the type of directory partitions indicated in the Legend.

### Simple Ring Topology that Requires No Optimization







Because a ring topology is created for each directory partition, the topology might look different if domain controllers from a second domain were present in the site. The next diagram illustrates the topology for domain controllers from two domains in the same site with no global catalog servers defined in the site.

**Ring Topology for Two Domains in a Site that Has No Global Catalog Server**







The next diagram illustrates replication between a global catalog server and three domains to which the global catalog server does not belong. When a global catalog server is added to the site in DomainA, additional connections are required to replicate updates of the other domain directory partitions to the global catalog server. The KCC on the global catalog server creates connection objects to replicate from domain controllers for each of the other domain directory partitions within the site, or from another global catalog server, to update the read-only partitions. Wherever a domain directory partition is replicated, the KCC also uses the connection to replicate the schema and configuration directory partitions.

**Note**

- Connection objects are generated independently for the configuration and schema directory partitions (one connection) and for the separate domain and application directory partitions, unless a connection from the same source to destination domain controllers already exists for one directory partition. In that case, the same connection is used for all (duplicate connections are not created).

**Intrasite Topology for Site with Four Domains and a Global Catalog Server**




**Expanded Ring Topology Within a Site**

When the number of servers in a site grows beyond seven, the KCC estimates the number of connections that are needed so that if a change occurs at any one domain controller, there are as many replication partners as needed to ensure that no domain controller is more than three replication hops from another domain controller (that is, a change takes no more than three hops before it reaches another domain controller that has not already received the change by another path). These optimizing connections are created at random and are not necessarily created on every third domain controller.

The KCC adds connections automatically to optimize a ring topology within a site, as follows:

- Given a set of nodes in a ring, create the minimum number of connections, *n*, that each server must have to ensure a path of no more than three hops to another server.

Given *n*, topology generation proceeds as follows.

- If the local server does not have *n* extra connections, the KCC does the following:
    - Chooses *n* other servers randomly in the site as source servers.
    - For each of those servers, creates a connection object.

This approach approximates the minimum-hop goal of three servers. In addition, it grows well, because as the site grows in server count, old optimizing connections are still useful and are not removed. Also, every time an additional 9 to 11 servers are added, a connection object is deleted at random; then a new one is created, ideally having one of the new servers as its source. This process ensures that, over time, the additional connections are distributed well over the entire site.

The following diagram shows an intrasite ring topology with optimizing connections in a site that has eight domain controllers in the same domain. Without optimizing connections, the hop count from DC1 to DC2 is more than three hops. The KCC creates optimizing connections to limit the hop count to three hops. The two one-way inbound optimizing connections accommodate all directory partitions that are replicated between the two domain controllers.

**Intrasite Topology with Optimizing Connections**




**Excluded Nonresponding Servers**

The KCC automatically rebuilds the replication topology when it recognizes that a domain controller has failed or is unresponsive.

The criteria that the KCC uses to determine when a domain controller is not responsive depend upon whether the server computer is within the site or not. Two thresholds must be reached before a domain controller is declared "unavailable" by the KCC:

- The requesting domain controller must have made *n* attempts to replicate from the target domain controller.
    - For replication between sites, the default value of *n* is 1 attempt.
    - For replication within a site, the following distinctions are made between the two immediate neighbors (in the ring) and the optimizing connections:

      For immediate neighbors, the default value of *n* is 0 failed attempts. Thus, as soon as an attempt fails, a new server is tried.

      For optimizing connections, the default value of *n* is 1 failed attempt. Thus, as soon as a second failed attempt occurs, a new server is tried.
- A certain amount of time must have passed since the last successful replication attempt.
    - For replication between sites, the default time is 2 hours.
    - For replication within a site, a distinction is made between the two immediate neighbors (in the ring) and the optimizing connections:

      For immediate neighbors, the default time is 2 hours.

      For optimizing connections, the default value is 12 hours.

You can edit the registry to modify the thresholds for excluding nonresponding servers.

**Note**

- If you must edit the registry, use extreme caution. Registry information is provided here as a reference for use by only highly skilled directory service administrators. It is recommended that you do not directly edit the registry unless, as in this case, there is no Group Policy or other Windows tools to accomplish the task. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. Storage of incorrect values can result in unrecoverable errors in the system.

Modifying the thresholds for excluding nonresponding servers requires editing the following registry entries in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters, with the data type REG_DWORD. You can modify these values to any desired value as follows:

For replication between sites, use the following entries:

- **IntersiteFailuresAllowed**

  Value: Number of failed attempts

  Default: 1

- **MaxFailureTimeForIntersiteLink (sec)**

  Value: Time that must elapse before being considered unavailable, in seconds

  Default: 7200 (2 hours)

For optimizing connections within a site, use the following entries:

- **NonCriticalLinkFailuresAllowed**

  Value: Number of failed attempts

  Default: 1

- **MaxFailureTimeForNonCriticalLink**

  Value: Time that must elapse before considered unavailable, in seconds

  Default: 43200 (12 hours)

For immediate neighbor connections within a site, use the following entries:

- **CriticalLinkFailuresAllowed**

  Value: Number of failed attempts

  Default: 0

- **MaxFailureTimeForCriticalLink**

  Value: Time that must elapse before considered unavailable, in seconds

  Default: 7200 (2 hours)

When the original domain controller begins responding again, the KCC automatically restores the replication topology to its pre-failure condition the next time that the KCC runs.

#### Fully Optimized Ring Topology Generation

Taking the addition of extra connections, management of nonresponding servers, and growth-management mechanisms into account, the KCC proceeds to fully optimize intrasite topology generation. The appropriate connection objects are created and deleted according to the available criteria.

**Note**

- Connection objects from nonresponding servers are not deleted because the condition is expected to be transient.

#### Automated Intersite Topology Generation

To produce a replication topology for hundreds of domains and thousands of sites in a timely manner and without compromising domain controller performance, the KCC must make the best possible decision when confronted with the question of which network link to use to replicate a given directory partition between sites. Ideally, connections occur only

between servers that contain the same directory partition(s), but when necessary, the KCC can also use network paths that pass through servers that do not store the directory partition.

Intersite topology generation and associated processes are improved in Windows Server 2003 in the following ways:

- Improved scalability: A new spanning tree algorithm achieves greater efficiency and scalability when the forest has a functional level of Windows Server 2003. For more information about this new algorithm, see "Improved KCC Scalability in Windows Server 2003 Forests" later in this section.

- Less network traffic: A new method of communicating the identity of the ISTG reduces the amount of network traffic that is produced by this process. For more information about this method, see "Intersite Topology Generator" later in this section.

- Multiple bridgehead servers per site and domain, and initial bridgehead server load balancing: An improved algorithm provides random selection of multiple bridgehead servers per domain and transport (the Windows 2000 algorithm allows selection of only one). The load among bridgehead servers is balanced the first time connections are generated. For more information about bridgehead server load balancing, see "Windows Server 2003 Multiple Bridgehead Selection" later in this section.

**Factors Considered by the KCC**

The spanning tree algorithm used by the KCC that is running as the ISTG to create the intersite replication topology determines how to connect all the sites that need to be connected with the minimum number of connections and the least cost. The algorithm must also consider the fact that each domain controller has at least three directory partitions that potentially require synchronization with other sites, not all domain controllers store the same partitions, and not all sites host the same domains.

The ISTG considers the following factors to arrive at the intersite replication topology:

- Location of domain directory partitions (calculate a replication topology for each domain).

- Bridgehead server availability in each site (at least one is available).

- All explicit site links.

- With automatic site link bridging in effect, consider all implicit paths as a single path with a combined cost.

- With manual site link bridging in effect, consider the implicit combined paths of only those site links included in the explicit site link bridges.

- With no site link bridging in effect, where the site links represent hops between domain controllers in the same domain, replication flows in a store-and-forward manner through sites.

**Improved KCC Scalability in Windows Server 2003 Forests**

KCC scalability is greatly improved in Windows Server 2003 forests over its capacity in Windows 2000 forests. Windows 2000 forests scale safely to support 300 sites, whereas Windows Server 2003 forests have been tested to 3,000 sites. This level of scaling is achieved when the forest functional level is Windows Server 2003. At this forest functional level, the method for determining the least-cost path from each site to every other site for each directory partition is significantly more efficient than the method that is used in a Windows 2000 forest or in a Windows Server 2003 forest that has a forest functional level of Windows 2000.

**Windows 2000 Spanning Tree Algorithm**

The ability of the KCC to generate the intersite topology in Windows 2000 forests is limited by the amount of CPU time and memory that is consumed when the KCC computes the replication topology in large environments that use transitive (bridged) site links. In a Windows 2000 forest, a potential disadvantage of bridging all site links affects only very large networks (generally, greater than 100 sites) where periods of high CPU activity occur every 15 minutes when the KCC runs. By default, the KCC creates a single bridge for the entire network, which generates more routes that must be processed than if automatic site link bridging is not used and manual site link bridges are applied selectively.

In a Windows 2000 forest, or in a Windows Server 2003 forest that has a forest functional level of Windows 2000, the KCC reviews the comparison of multiple paths to and from every destination and computes the spanning tree of the least-cost path. The spanning tree algorithm works as follows:

- Computes a cost matrix by identifying each site pair (that is, each pair of bridgehead servers in different sites that store the directory partition) and the cost on the site link connecting each pair.

  **Note**

  - This matrix is actually computed by Intersite Messaging and used by the KCC.

- By using the costs computed in the matrix, builds a spanning tree between sites that store the directory partition.

This method becomes inefficient when there are a large number of sites.

**Note**

- CPU time and memory is not an issue in a Windows 2000 forest as long as the following criteria apply:
- D is the number of domains in your network
- S is the number of sites in your network
- (1 + D) * S^2 <= 100,000

**Windows Server 2003 Spanning Tree Algorithm**

A more efficient spanning tree algorithm improves efficiency and scalability of replication topology generation in Windows Server 2003 forests. When the forest functional level is either Windows Server 2003 or Windows Server 2003 interim, the improved algorithm takes effect and computes a minimum-cost spanning tree of connections between the sites that host a particular directory partition, but eliminates the inefficient cost matrix. Thus, the KCC directly determines the lowest-cost spanning tree for each directory partition, considering the schema and configuration directory partitions as a single tree. Where the spanning trees overlap, the KCC generates a single connection between domain controllers for replication of all common directory partitions.

In a Windows Server 2003 forest, both versions of the KCC spanning tree algorithms are available. The algorithm for Windows 2000 forests is retained for backwards compatibility with the Windows 2000 KCC. It is not possible for the two algorithms to run simultaneously in the same enterprise.

**DFS Site Costing and Windows Server 2003 SP1 Site Options**

When the forest functional level is Windows Server 2003 or Windows Server 2003 interim and the ISTG does not use Intersite Messaging to calculate the intersite cost matrix, DFS can still use Intersite Messaging to compute the cost matrix for its site-costing functionality, provided that the **Bridge all site links** option is enabled. In branch office deployments, where the large number of sites and site links makes automatic site link bridging too costly in terms of the replication connections that are generated, the **Bridge all site links** option is usually disabled on the IP container (CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=*ForestRootDomain*). If the **Bridge all site links** option is disabled, DFS is unable to use Intersite Messaging to calculate site costs.

When the forest functional level is Windows Server 2003 or Windows Server 2003 interim and the ISTG in a site is running Windows Server 2003 with SP1, you can have the **Bridge all site links** option enabled, and then use a site option to turn off automatic site link bridging for KCC operation without hampering the ability of DFS to use Intersite Messaging to calculate the cost matrix. This site option is set by running the command **repadmin /siteoptions W2K3_BRIDGES_REQUIRED**. This option is applied to the NTDS Site Settings object (CN=NTDS Site Settings,CN=*SiteName*,CN=Sites,CN=Configuration,DC=*ForestRootDomain*). When this method is used to disable automatic site link bridging (as opposed to disabling **Bridge all site links**), default Intersite Messaging options enable the site-costing calculation to occur for DFS.

| **Note** |
| --- |
| The site option on the NTDS Site Settings object can be set on any domain controller, but it does not take effect until replication of the change reaches the ISTG role holder for the site. |

**Intersite Topology Generator**

The KCC on the domain controller that has the ISTG role creates the inbound connections on all domain controllers in its site that require replication with domain controllers in other sites. The sum of these connections for all sites in the forest forms the

intersite replication topology.

A fundamental concept in the generation of the topology within a site is that each server does its part to create a site-wide topology. In a similar manner, the generation of the topology between sites depends on each site doing its part to create a forest-wide topology between sites.

**ISTG Role Ownership and Viability**

The owner of the ISTG role is communicated through normal Active Directory replication. Initially, the first domain controller in the site is the ISTG role owner. It communicates its role ownership to other domain controllers in the site by writing the distinguished name of its child NTDS Settings object to the **interSiteTopologyGenerator** attribute of the NTDS Site Settings object for the site. As a change to the configuration directory partition, this value is replicated to all domain controllers in the forest.

The ISTG role owner is selected automatically. The role ownership does not change unless:

- The current ISTG role owner becomes unavailable.
- All domain controllers in the site are running Windows 2000 and one of them is upgraded to Windows Server 2003.

If at least one domain controller in a site is running Windows Server 2003, the ISTG role is assumed by a domain controller that is running Windows Server 2003.

The viability of the current ISTG is assessed by all other domain controllers in the site. The need for a new ISTG in a site is established differently, depending on the forest functional level that is in effect.

- Windows 2000 functional level: At 30-minute intervals, the current ISTG notifies every other domain controller of its existence and availability by writing the **interSiteTopologyGenerator** attribute of the NTDS Site Settings object for the site. The change replicates to every domain controller in the forest. The KCC on each domain controller monitors this attribute for its site to verify that it has been written. If a period of 60 minutes elapses without a modification to the attribute, a new ISTG declares itself.
- Windows Server 2003 or Windows Server 2003 interim functional level: Each domain controller maintains an up-to-dateness vector, which contains an entry for each domain controller that holds a full replica of any directory partition that the domain controller replicates. On domain controllers that are running Windows Server 2003, this up-to-dateness vector contains a timestamp that indicates the times that it was last contacted by its replication partners, including both direct and indirect partners (that is, every domain controller that replicates a directory partition that is stored by this domain controller). The timestamp is recorded whether or not the local domain controller actually received any changes from the partner. Because all domain controllers store the schema and configuration directory partitions, every domain controller is guaranteed to have the ISTG for its site among the domain controllers in its up-to-dateness vector.

  This timestamp eliminates the need to receive periodic replication of the updated **interSiteTopologyGenerator** attribute from the current ISTG. When the timestamp indicates that the current ISTG has not contacted the domain controller in the last 120 minutes, a new ISTG declares itself.

The Windows Server 2003 method eliminates the network traffic that is generated by periodically replicating the **interSiteTopologyGenerator** attribute update to every domain controller in the forest.

**ISTG Eligibility**

When at least one domain controller in a site is running Windows Server 2003, the eligibility for the ISTG role depends on the operating system of the domain controllers. When a new ISTG is required, each domain controller computes a list of domain controllers in the site. All domain controllers in the site arrive at the same ordered list. Eligibility is established as follows:

- If no domain controllers in the site are running Windows Server 2003, all domain controllers that are running Windows 2000 Server are eligible. The list of eligible servers is ordered by GUID.
- If at least one domain controller in the site is running Windows Server 2003, all domain controllers that are running Windows Server 2003 are eligible. In this case, the entries in the list are sorted first by operating system and then by GUID. In a site in which some domain controllers are running Windows 2000 Server, domain controllers that are running Windows Server 2003 remain at the top of the list and use the GUID in the same manner to maintain the order.

The domain controller that is next in the list of servers after the current owner declares itself the new ISTG by writing the **interSiteTopologyGenerator** attribute on the NTDS Site Settings object.

If the current ISTG is temporarily disconnected from the topology, as opposed to being shut down, and a new ISTG declares itself in the interim, then two domain controllers can temporarily assume the ISTG role. When the original ISTG resumes replication, it initially considers itself to be the current ISTG and creates inbound replication connection objects, which results in duplicate intersite connections. However, as soon as the two ISTGs replicate with each other, the last domain controller to write the **intersiteTopologyGenerator** attribute continues as the single ISTG and removes the duplicate connections.

**Bridgehead Server Selection**

Bridgehead servers can be selected in the following ways:

- Automatically by the ISTG from all domain controllers in the site.
- Automatically by the ISTG from all domain controllers that are identified as preferred bridgehead servers, if any preferred bridgehead servers are assigned. Preferred bridgehead servers must be assigned manually.
- Manually by creating a connection object on a domain controller in one site from a domain controller in a different site.

By default, when at least one domain controller in a site is running Windows Server 2003 (regardless of forest functional level), any domain controller that hosts a domain in the site is a candidate to be an ISTG-selected bridgehead server. If preferred bridgehead servers are selected, candidates are limited to this list. The connections from remote servers are distributed among the available candidate bridgehead servers in each site. The selection of multiple bridgehead servers per domain and transport is new in Windows Server 2003. The ISTG uses an algorithm to evaluate the list of domain controllers in the site that can replicate each directory partition. This algorithm is improved in Windows Server 2003 to randomly select multiple bridgehead servers per directory partition and transport. In sites containing only domain controllers that are running Windows 2000 Server, the ISTG selects only one bridgehead server per directory partition and transport.

When bridgehead servers are selected by the ISTG, the ISTG ensures that each directory partition in the site that has a replica in any other site can be replicated to and from that site or sites. Therefore, if a single domain controller hosts the only replica of a domain in a specific site and the domain has domain controllers in another site or sites, that domain controller must be a bridgehead server in its site. In addition, that domain controller must be able to connect to a bridgehead server in the other site that also hosts the same domain directory partition.

### Note

- If a site has a global catalog server but does not contain at least one domain controller of every domain in the forest, then at least one bridgehead server must be a global catalog server.

**Preferred Bridgehead Servers**

Because bridgehead servers must be able to accommodate more replication traffic than non-bridgehead servers, you might want to control which servers have this responsibility. To specify servers that the ISTG can designate as bridgeheads, you can select domain controllers in the site that you want the ISTG to always consider as preferred bridgehead servers for the specified transport. These servers are used exclusively to replicate changes collected from the site to any other site over that transport. Designating preferred bridgehead servers also serves to exclude those domain controllers that, for reasons of capability, you do *not* want to be used as bridgehead servers.

Depending on the available transports, the directory partitions that must be replicated, and the availability of global catalog servers, multiple bridgehead servers might be required to replicate full and partial copies of directory data from one site to another.

The ISTG recognizes preferred bridgehead servers by reading the **bridgeheadTransportList** attribute of the server object. When this attribute has a value that is the distinguished name of the transport container that the server uses for intersite replication (IP or SMTP), the KCC treats the server as a preferred bridgehead server. For example, the value for the IP transport is CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,DC=*ForestRootDomainName*. You can use Active Directory Sites and Services to designate a preferred bridgehead server by opening the server object properties and placing either the IP or SMTP transport into the preferred list, which adds the respective transport distinguished name to the **bridgeheadTransportList** attribute of the server.

The **bridgeheadServerListBl** attribute of the transport container object is a backlink attribute of the **bridgeheadTransportList** attribute of the server object. If the **bridgeheadServerListBl** attribute contains the distinguished name of at least one server in a site, then the KCC uses only preferred bridgehead servers to replicate changes for that site,

according to the following rules:

- If at least one server is designated as a preferred bridgehead server, updates to the domain directory partition hosted by that server can be replicated only from a preferred bridgehead server. If at the time of replication no preferred bridgehead server is available for that directory partition, replication of that directory partition does not occur.

- If any bridgehead servers are designated but no domain controller is designated as a preferred bridgehead server for a specific directory partition that has replicas in another site or sites, the KCC selects a domain controller to act as the bridgehead server, if one is available that can replicate the directory partition to the other site or sites.

Therefore, to use preferred bridgehead servers effectively, be sure to:

- Assign at least two or more bridgehead servers for each of the following:

    - Any domain directory partition that has a replica in any other site.

    - Any application directory partition that has a replica in another site.

    - The schema and configuration directory partitions (one bridgehead server replicates both) if no domains in the site have replicas in other sites.

- If the site has a global catalog server, select the global catalog server as one of the preferred bridgehead servers.

**Windows 2000 Single Bridgehead Selection**
In a Windows 2000 forest or in a Windows Server 2003 forest that has a forest functional level of Windows 2000, the ISTG selects a single bridgehead server per directory partition and transport. The selection changes only when the bridgehead server becomes unavailable. The following diagram shows the automatic bridgehead server (BH) selection that occurs in the hub site where all domain controllers host the same domain directory partition and multiple sites have domain controllers that host that domain directory partition.

**Windows 2000 Single Bridgehead Server in a Hub Site that Serves Multiple Branch Sites**




**Windows Server 2003 Multiple Bridgehead Selection**
When at least one domain controller in a site is running Windows Server 2003 (and thereby becomes the ISTG), the ISTG begins performing random load balancing of new connections. This load balancing occurs by default, although it can be disabled.

When creating a new connection, the KCC must choose endpoints from the set of eligible bridgeheads in the two endpoint sites. Whereas in Windows 2000 the ISTG always picks the same bridgehead for all connections, in Windows Server 2003 it picks one randomly from the set of possible bridgeheads.

Assuming two out of three of the domain controllers have been added to the site since the ISTG was upgraded to Windows Server 2003, the following diagram shows the connections that might exist on domain controllers in the hub site to

accommodate the four branch sites that have domain controllers for the same domain.

**Random Bridgehead Server Selection in a Hub Site in which the ISTG Runs Windows Server 2003**




If one or more new domain controllers are added to the hub site, the inbound connections on the existing bridgehead servers are not automatically re-balanced. The next time it runs, the ISTG considers the newly added server(s) as follows:

- If preferred bridgehead servers are not selected in the site, the ISTG considers the newly added servers as candidate bridgehead servers and creates new connections randomly if new connections are needed. It does not remove or replace the existing connections.

- If preferred bridgehead servers are selected in the site, the ISTG does not consider the newly added servers as candidate bridgehead servers unless they are designated as preferred bridgehead servers.

The initial connections remain in place until a bridgehead server becomes unavailable, at which point the KCC randomly replaces the connection on any available bridgehead server. That is, the endpoints do not change automatically for the existing bridgehead servers. In the following diagram, two new domain controllers are added to the hub site, but the existing connections are not redistributed.

**New Domain Controllers with No New Connections Created**




Although the ISTG does not rebalance the load among the existing bridgehead servers in the hub site after the initial connections are created, it does consider the added domain controllers as candidate bridgehead servers under either of the following conditions:

- A new site is added that requires a bridgehead server connection to the hub site.

- An existing connection to the hub site becomes unavailable.

The following diagram illustrates the inbound connection that is possible on a new domain controller in the hub site to accommodate a failed connection on one of the existing hub site bridgehead servers. In addition, a new branch site is added and a new inbound connection can potentially be created on the new domain controller in the hub site. However, because the selection is random, there is no guarantee that the ISTG creates the connections on the newly added domain controllers.

**Possible New Connections for Added Site and Failed Connection**




**Using ADLB to Balance Hub Site Bridgehead Server Load**

In large hub-and-spoke topologies, you can implement the redistribution of existing bridgehead server connections by using the Active Directory Load Balancing (ADLB) tool (Adlb.exe), which is included with the "Windows Server 2003 Active Directory Branch Office Deployment Guide." ADLB makes it possible to dynamically redistribute the connections on bridgehead servers. This application works independently of the KCC, but uses the connections that are created by the KCC. Connections that are manually created are not moved by ADLB. However, manual connections are factored into the load-balancing equations that ADLB uses.

The ISTG is limited to making modifications in its site, but ADLB modifies both inbound and outbound connections on eligible bridgehead servers and offers schedule staggering for outbound connections.

For more information about how bridgehead server load balancing and schedule staggering work with ADLB, see the "Windows Server 2003 Active Directory Branch Office Planning and Deployment Guide" on the Web at http://go.microsoft.com/fwlink/?linkID=28523.

**Network Ports Used by Replication Topology**

By default, RPC-based replication uses dynamic port mapping. When connecting to an RPC endpoint during Active Directory replication, the RPC run time on the client contacts the RPC endpoint mapper on the server at a well-known port (port 135). The server queries the RPC endpoint mapper on this port to determine what port has been assigned for Active Directory replication on the server. This query occurs whether the port assignment is dynamic (the default) or fixed. The client never needs to know which port to use for Active Directory replication.

**Note**

- An endpoint comprises the protocol, local address, and port address.

In addition to the dynamic port 135, other ports that are required for replication to occur are listed in the following table.

**Port Assignments for Active Directory Replication**

| Service Name | UDP | TCP |
|---|---|---|

| | | |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | | 636 (Secure Sockets Layer [SSL]) |
| LDAP | | 3268 (global catalog) |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

Replication within a domain also requires FRS using a dynamic RPC port.

**Setting Fixed Replication Ports Across a Firewall**

For each service that needs to communicate across a firewall, there is a fixed port and protocol. Normally, the directory service and FRS use dynamically allocated ports that require a firewall to have a wide range of ports open. Although FRS cannot be restricted to a fixed port, you can edit the registry to restrict the directory service to communicate on a static port.

### Note

- If you must edit the registry, use extreme caution. Registry information is provided here as a reference for use by only highly skilled directory service administrators. It is recommended that you do not directly edit the registry unless, as in this case, there is no Group Policy or other Windows tools to accomplish the task. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. Storage of incorrect values can result in unrecoverable errors in the system.

Restricting the directory service to using a fixed port requires editing the **TCP/IP Port** registry entry (REG_DWORD), located in:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

Changing this registry entry on a domain controller and restarting it causes the directory service to use the TCP port named in the registry entry. For example, port 49152 is DWORD=0000c000 (hexadecimal).

**Related Information**

The following resources contain additional information that is relevant to this section.

- Active Directory Replication Model Technical Reference
- TCP/IP Technical Reference
- FRS Technical Reference
- DNS Support for Active Directory Technical Reference
- DFS Technical Reference

# Active Directory Replication Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Active Directory Replication Tools and Settings**

**In this section**


- Active Directory Replication Tools
- Active Directory Replication Registry Entries
- Active Directory Replication Group Policy Settings
- Active Directory Replication WMI Classes
- Network Ports Used by Active Directory Replication
- Related Information


**Active Directory Replication Tools**

The following tools are associated with Active Directory replication.

| Note |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**Dssite.msc: Active Directory Sites and Services**

**Category**

Active Directory Administrative Tools Microsoft Management Console (MMC) snap-in. This tool is installed automatically when you install Active Directory, and is available on the Start menu under Programs\Administrative Tools. This tool also ships with the Administration Tools Pack (Adminpak.msi).

**Version compatibility**

Active Directory Sites and Services provides a view into the Sites container of the configuration directory partition. Use Active Directory Sites and Services to manage Active Directory replication topology. The following objects and their properties can be managed by using this tool:

- Sites container: Add new sites.
- Site objects: Add new servers to a site.
- NTDS Site Settings object: For each site, view the connection object schedule and enable Universal group membership caching.
- Server object: View the NTDS Settings object and designate the server as a bridgehead server.
- NTDS Settings object: View inbound connections for the server. View the connection object schedule and change the source server for the connection.
- Inter-Site Transports container: Manage IP and SMTP site links.
- Site link objects: Manage the site link properties for a set of sites.

- Subnets container: Add, remove, and configure subnets with IP addresses. Associate subnets with sites.

**Repadmin.exe: Repadmin**

**Category**
Command-line tool.

**Version compatibility**
Repadmin is used to view the replication information on domain controllers. You can determine the last successful replication of all directory partitions, identify inbound and outbound replication partners, identify the current bridgehead servers, view object metadata, and generally manage Active Directory replication topology. You can use Repadmin to force replication of an entire directory partition or of a single object. You can also list domain controllers in a site.

Repadmin is extended to enable commands to target sets of domain controllers. For example, you can target all domain controllers in a site or domain, or all domain controllers that are global catalog servers.

Repadmin also includes the **RemoveLingeringObjects** command, which removes objects that are outdated (do not exist in a replica of the same directory partition on the source domain controller).

For more information about removing lingering objects, see "Fixing Replication Lingering Object Problems (Event IDs 1388, 1988, 2042)" in the Windows Server 2003 Operations Guide at http://go.microsoft.com/fwlink/?LinkId=44131. For more information about Repadmin, see Repadmin Overview.

**Ntdsutil.exe: Ntdsutil**

**Category**
Command-line tool.

**Version compatibility**
Ntdsutil.exe provides management capabilities for Active Directory. You can use Ntdsutil.exe to perform Active Directory database maintenance, manage and control single-master operations, and remove replication metadata left behind by domain controllers that are removed from the network without uninstalling Active Directory. The version of Ntdsutil that is included with Windows Server 2003 SP1 removes File Replication service (FRS) metadata in addition to Active Directory replication metadata. You can also use Ntdsutil to create application directory partitions and perform authoritative restore operations. This tool is intended for use by experienced administrators.

**Active Directory Replication Registry Entries**
The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following registry settings cannot be modified by using Group Policy or other Windows tools.

**NTDS Parameters Registry Settings**
The following registry entries are associated with Active Directory replication.

**Replicator notify pause after modify (secs)**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows 2000 Server.

**Default value**
Windows 2000 Server: 300 seconds.

The value for the delay between an originating update on a domain controller and the first change notification. On domain controllers running Windows Server 2003 or higher, the value for initial change notification delay is stored in the **msDSReplicationNotifyFirstDSADelay** attribute on the cross-reference object for each directory partition in the Configuration container. The default value in Windows Server 2003 and higher operating systems is decreased to 15 seconds

when the forest functional level is Windows Server 2003 or higher.

**Replicator notify pause between DSAs (secs)**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows 2000 Server.

**Default value**
Windows 2000 Server: 30 seconds

The value for the delay before each subsequent change notification. On domain controllers running Windows Server 2003, the value for subsequent notification delay is stored in the **msDSReplicationNotifySubsequentDSADelay** attribute on the cross-reference object for each directory partition in the Configuration container. The default value in Windows Server 2003 is decreased to 3 seconds when the forest functional level is Windows Server 2003.

**RPC Replication Timeout (mins)**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
Windows 2000 Server: 45 minutes; Windows Server 2003 and higher server operating systems: 5 minutes.

The number of minutes between initiation of Active Directory replication and the RPC timeout. The domain controller must be restarted before the change takes effect.

**Strict replication consistency**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server with SP3.

**Default value**
Windows 2000 Server with SP3: off (0); Windows Server 2003 and higher server operating systems: on (1)

The value that determines the treatment of replication of outdated objects that exist on reconnected domain controllers that have not replicated in longer than a tombstone lifetime. If the destination domain controller has strict replication consistency enabled, inbound replication of an outdated object is blocked. If the destination domain controller has strict replication disabled, inbound replication of the full object occurs.

**Replicator intra site packet size (objects)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1/1,000,000th the size of RAM, with a minimum of 100 objects and a maximum of 1,000 objects.

The maximum number of objects per packet for RPC replication within a site.

**Replicator intra site packet size (bytes)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1/100th the size of RAM, with a minimum of 1 megabyte (MB) and a maximum of 10 MB.

The maximum size of objects per packet for RPC replication within a site.

**Replicator inter site packet size (objects)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server.

**Default value**
1/1,000,000th the size of RAM, with a minimum of 100 objects and a maximum of 1,000 objects.

The maximum number of objects per packet for RPC replication between sites.

**Replicator inter site packet size (bytes)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

The maximum size of objects per packet for RPC replication between sites.

**Default value**
1/100th the size of RAM, with a minimum of 1 MB and a maximum of 10 MB.

**Replicator async inter site packet size (objects)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server.

**Default value**
1/1,000,000th the size of RAM, with a minimum of 100 objects and a maximum of 1,000 objects.

The maximum number of objects per packet for SMTP replication between sites.

**Replicator async inter site packet size (bytes)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1 MB.

The maximum size of objects per packet for SMTP replication between sites.

**Replicator compression algorithm**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003.

**Default value**
For Windows 2000 Server compression, change the value to 2.

Determines the compression algorithm that is used on a site link

**Repl topology update delay (secs)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
300 seconds.

Number of seconds to wait between the time Active Directory starts and the KCC performs the first topology check.

To find more information about **Repl topology update delay (secs)**, see "Registry Reference" in Tools and Settings Collection.

**Repl topology update period (secs)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server.

**Default value**
900 seconds.

Interval between KCC replication topology checks.

To find more information about **Repl topology update period (secs)**, see "Registry Reference" in Tools and Settings Collection.

**IntersiteFailuresAllowed**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1.

Number of failed replication attempts prior to excluding nonresponding servers from the intersite topology.

**MaxFailureTimeForIntersiteLink (sec)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
7200 seconds (2 hours).

Time in seconds that must elapse prior to excluding nonresponding servers from the intersite topology.

**NonCriticalLinkFailuresAllowed**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
1.

Number of failed replication attempts prior to excluding nonresponding servers from the intrasite topology.

**MaxFailureTimeForNonCriticalLink (sec)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows 2000 Server.

**Default value**
43200 seconds (12 hours).

Time in seconds that must elapse prior to excluding nonresponding servers from the intrasite topology.

**CriticalLinkFailuresAllowed**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
0.

Number of failed replication attempts prior to excluding nonresponding servers for immediate neighbor connections within a site.

**MaxFailureTimeForCriticalLink (sec)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
7200 seconds (2 hours).

Time in seconds that must elapse prior to excluding nonresponding servers for immediate neighbor connections within a site.

**TCP/IP Port**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server.

**Default value**
135.

TCP port that the directory service uses instead of using dynamic port 135. The domain controller must be restarted before the change takes effect.

**Backup Latency Threshold (days)**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters

**Version**
Windows Server 2008 R2, Windows Server 2008, Windows Server 2003 with SP 1

**Default value**
Half the value of the tombstone lifetime of the forest.

When the value is reached, logs event ID 2089 in the Directory Service event log, warning administrators and monitoring applications to make sure that domain controllers are backed up before the tombstone lifetime expires.

**Active Directory Replication Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with Active Directory replication updates.

**Group Policy Settings Associated with Active Directory Replication**

| Group Policy Setting | Description |
|---|---|
| Account Lockout Policy:<br>• Account lockout duration<br>• Account lockout threshold<br>• Reset account lockout counter after | Changes to these settings in the Domain Security Policy trigger urgent replication. |
| Password Policy:<br>• Enforce password history<br>• Maximum password age<br>• Minimum password age<br>• Minimum password length<br>• Password must meet complexity requirements<br>• Store passwords using reversible encryption | Changes to these settings in the Domain Security Policy trigger urgent replication. |
| Contact PDC on logon failure | Account lockout and domain password changes rely on contacting the primary domain controller (PDC) emulator urgently to update the PDC emulator with the change. If Contact PDC on logon failure is disabled, replication of password changes to the PDC emulator occurs non-urgently. |

To find more information about these Group Policy settings, see “Group Policy Settings Reference” in Tools and Settings

Collection.

**Active Directory Replication WMI Classes**

The following table lists and describes the WMI classes that are associated with Active Directory replication. These classes are shipped with Windows Server 2003 or later server operating systems, but are also compatible with Windows 2000 Server.

**WMI Classes Associated with Active Directory Replication**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| MSAD_DomainController | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |
| MSAD_NamingContext | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |
| MSAD_ReplNeighbor | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |
| MSAD_ReplCursor | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |
| MSAD_ReplPendingOp | \\root\MicrosoftActiveDirectory | Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003<br>Windows 2000 Server |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by Active Directory Replication**

By default, RPC-based replication uses dynamic port mapping. When connecting to an RPC endpoint during Active Directory replication, the RPC run time on the client contacts the RPC endpoint mapper on the server at a well-known port (port 135). The server queries the RPC endpoint mapper on this port to determine what port has been assigned for Active Directory replication on the server. This query occurs whether the port assignment is dynamic (the default) or fixed. The client never needs to know which port to use for Active Directory replication.

**Note**

- An endpoint comprises the protocol, local address, and port address.

In addition to the dynamic port 135, other ports that are required for replication to occur are listed in the following table.

**Port Assignments for Active Directory Replication**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | 389 | 389 |
| LDAP | | 636 (Secure Sockets Layer [SSL]) |
| LDAP | | 3268 (global catalog) |
| Kerberos | 88 | 88 |
| DNS | 53 | 53 |
| SMB over IP | 445 | 445 |

Replication within a domain also requires FRS using a dynamic RPC port.

**Related Information**

The following resources contain additional information that is relevant to this section.

- How the Active Directory Replication Model Works
- How Active Directory Replication Topology Works

# Active Directory Search and Publication Technologies

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

Active Directory exists so that users, services, and applications can search for and publish useful information in the directory. The operations that users, services and applications perform against the directory include the following:

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

- Performing searches against the data
- Finding (or in the case of services, publishing) information related to services that are available on the network

**Active Directory Search and Publication Architecture**

The Active Directory architecture that supports search and service publication can be divided into two functional areas:

- Search
- Service publication

The following table describes these Active Directory functional areas.

**Active Directory Search and Service Publication**

| Functional Area | Description |
| --- | --- |
| Active Directory search | Directory clients and services need a way to find data that is stored in the directory. Requests for directory objects are carried out either through the Active Directory Service Interfaces (ADSI) Lightweight Directory Access Protocol (LDAP) provider or through the LDAP application programming interface (API). |
| Active Directory service publication | Service publication in Active Directory enables services to provide information about themselves in the directory, and it enables directory clients to search for available services on the network. In addition, Active Directory supports service principal names (SPNs) as the means by which client applications can authenticate the services that they use. |

**Active Directory Search**

The primary components of the architecture for the Active Directory search function include the directory client applications that search the directory; LDAP, which is used for searching and retrieving directory information; and the Active Directory database against which the directory client applications search.

The following table describes the Active Directory search components.

**Active Directory Search Components**

| Search Component | Description |
| --- | --- |
| Directory client application | A directory client application is any application that is capable of searching for information that is stored in Active Directory. |

| | |
|---|---|
| LDAP | LDAP is a directory service protocol that specifies directory communications. It runs directly over TCP/IP, and it can also run over user datagram protocol (UDP) connectionless transports. LDAP enables clients to query, create, update, and delete information that is stored in a directory service over a TCP connection. LDAP is the preferred and most common means of interacting with Active Directory. |
| Active Directory database | The Active Directory database is the structured data store that Active Directory uses to store information about objects on the network, including users, user groups, computers, services, applications, application data, shared files, and distribution lists. A copy (or replica) of the Active Directory database resides on every domain controller in an Active Directory forest. |

**Active Directory Service Publication**

The primary components of the architecture for Active Directory service publication are the services that publish information about themselves and the client applications that search the directory to find and authenticate services. Active Directory provides the storage and distribution mechanism for published service information and for the SPN attributes that are used in mutual authentication. The Key Distribution Center (KDC) provides the mechanism for authenticating services, using SPNs that are constructed by the client applications.

The following table describes the Active Directory service publication and SPN components.

**Active Directory Service Publication and SPN Components**

| Service Publication Component | Description |
|---|---|
| Service | An application that makes data or operations available to client applications. |
| Client application | An application, which runs on a workstation (or on a server), that makes use of a service. |
| KDC | A service, which runs on every domain controller, that provides authentication services for clients as well as for servers and services. |
| Connection point object | An object in Active Directory that contains information about a service. |
| Service account object | An object in Active Directory that represents the account in whose security context a service runs and on which an SPN attribute resides. |
| SPN attribute | An attribute that contains a unique name that identifies an instance of a service and that is associated with the logon account under which the instance of the service runs. |

**Active Directory Search and Publication Scenarios**

Scenarios that rely on Active Directory search and publication include performing directory searches, advertising available services, finding available services, and authenticating services.

**Performing Directory Searches**

Searching the directory is a common Active Directory scenario in which directory clients use LDAP to query the directory and find information. Clients search the directory for a wide variety of information, including address book information, information about shared resources, and information related to a specific directory-enabled application. This scenario requires directory clients, LDAP, and the Active Directory database.

**Advertising Available Services**

A service that has services to offer client applications can use Active Directory as a way of advertising its services. In this scenario, a network service (at the time when it is installed) publishes a special object in the directory, called a connection point object. The connection point object holds information about the service, including binding information that a client application can use to connect to the service.

**Finding Available Services**

In large, distributed networks, directory clients must be able to find the network services that they need, regardless of where those services reside on the network. In this scenario, client applications search the directory for connection point objects that contain information about specific services that are available on the network. Client applications can then use this information to connect to the services that they need.

**Authenticating Services**

In large, distributed networks, a client application must be able to authenticate a service before the client application uses the

service. The process of authenticating a service protects client applications from malicious or accidental damage or breaches of security that can be caused by an unauthorized, or rogue, service. In this scenario, a client application requests authentication of an SPN representing the service that the client application wants to use. If the service can authenticate against a domain controller by using the SPN, the client application can safely use that service.

# Active Directory Searches Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

In the Active Directory directory service, directory clients use Lightweight Directory Access Protocol (LDAP) to perform searches against the directory. LDAP is defined by Request for Comments (RFC) 3377, "Lightweight Directory Access Protocol (v3): Technical Specification." LDAP is a standards-based protocol that makes it possible for users to query and update information in an LDAP-based directory service, such as Active Directory.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**In this subject**


- What Are Active Directory Searches?
- How Active Directory Searches Work
- Active Directory Searches Tools and Settings

# What Are Active Directory Searches?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**What Are Active Directory Searches?**

**In this section**


- The Business Need
- The Active Directory Searches Solution
- LDAP and Security
- Related Information


In Active Directory, directory clients use Lightweight Directory Access Protocol (LDAP) to perform searches against the directory. LDAP is defined by Request for Comments (RFC) 3377, "Lightweight Directory Access Protocol (v3): Technical Specification." LDAP is a standards-based protocol that makes it possible for users to query and update information in an LDAP-based directory service, such as Active Directory. To perform Active Directory searches, LDAP-compliant directory clients compose a search query using LDAP syntax and then submit the query to Active Directory for processing.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

Active Directory searches are an important part of many directory activities. Active Directory searches are commonly used by:

- Directory clients to find a domain controller.
- Directory clients to find network resources, such as printers and shared folders.
- Directory clients to find address book information, such as user names, e-mail addresses, and phone numbers.
- Non-Microsoft applications, such as Web portal, e-commerce, and white pages applications, to find and modify application-specific data that is stored in Active Directory.

**The Business Need**

Active Directory can be used to store all kinds of information, including identity, resource, infrastructure, and application-specific information. For this information to be useful, directory clients must be able to query and retrieve the information from Active Directory with a commonly understood protocol. This protocol must interact with a variety of directory client applications — and with a variety of client operating systems, because many organizations deploy more than one client operating system.

**The Active Directory Searches Solution**

The Active Directory searches solution consists of the following:

- Directory clients that submit LDAP search requests
- Domain controllers (or computers running Active Directory Application Mode (ADAM)) that accept and respond to the LDAP search requests

Any directory client that supports LDAP v2 (RFC 1777, "Lightweight Directory Access Protocol") or LDAP v3 (RFC 3377) can submit LDAP queries to Active Directory. The following figure illustrates the Active Directory searches solution.

**Active Directory Searches**




**Note**

- In the Windows Server 2003 family, a server running Windows Server 2003, Web Edition cannot be a domain controller.

Windows operating systems that can be used as a directory client include the following:

- Windows 95 (requires Active Directory Client Extensions)
- Windows 98 (requires Active Directory Client Extensions)
- Windows NT 4.0 (requires Active Directory Client Extensions)
- Windows 2000 Professional, Windows 2000 Server family
- Windows XP Home Edition, Windows XP Professional
- Windows Server 2003 family

For more information about Active Directory Client Extensions, see "ADSI WinNT Provider" in the Microsoft Platform SDK on MSDN.

Active Directory supports both LDAP v2 and LDAP v3. For more information about LDAP v2, see RFC 1777 in the IETF RFC Database. For more information about LDAP v3, see RFC 3377 in the IETF RFC Database.

**LDAP and Security**

Active Directory is designed to be secure by default, which means that directory clients must authenticate themselves before they can search Active Directory. Directory clients authenticate themselves against Active Directory by using the security mechanisms that are built into Windows Server 2003, including the Simple Authentication and Security Layer (SASL) protocol, the Kerberos version 5 (V5) authentication protocol, and NTLM. If necessary, Active Directory can also be configured to accept search requests from anonymous (unauthenticated) directory clients. For more information about search requests from anonymous directory clients, see "Anonymous queries" in "How Active Directory Searches Work." For more information about authentication methods, see "Logon and Authentication Technologies."

**Related Information**

The following resources contain additional information that is relevant to this section:

- "ADSI WinNT Provider" in the Microsoft Platform SDK on MSDN
- RFC 1777 in the IETF RFC Database
- RFC 3377 in the IETF RFC Database
- How Active Directory Searches Work
- Logon and Authentication Technologies
- DNS Support for Active Directory Technical Reference

# How Active Directory Searches Work

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**How Active Directory Searches Work**

**In this section**


- Active Directory Searches Architecture
- Active Directory Searches Protocols
- Active Directory Searches Interfaces
- The LDAP Standard
- Active Directory Searches Processes and Interactions
- Network Ports Used by Active Directory Searches
- Related Information


In Active Directory, directory clients use Lightweight Directory Access Protocol (LDAP) to perform searches against the directory. LDAP is defined by Request for Comments (RFC) 3377, "Lightweight Directory Access Protocol (v3): Technical Specification." LDAP is a standards-based protocol that makes it possible for users to query and update information in an LDAP-based directory service, such as Active Directory. To perform Active Directory searches, LDAP-compliant directory clients compose a search query using LDAP syntax and then submit the query to Active Directory for processing.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

The ideal environment for Active Directory searches includes network connections between directory clients and Active Directory that are available and working properly, domain controllers that are running and replicating without errors, and directory clients that support LDAP v2 or LDAP v3. This section describes the elements of Active Directory searches, including architecture, protocols, interfaces, LDAP, processes and interactions, and network ports.

**Active Directory Searches Architecture**

The architecture for Active Directory searches includes both client and server components. On the client side, a directory client application constructs LDAP requests to be sent to Active Directory. These requests can be one of several types, including connect, bind (authentication), modify, and unbind. Depending on how a directory client application is written, one of three different application programming interfaces (APIs) is used to submit requests. The LDAP requests are received and processed by the Directory System Agent (DSA), which is represented by Ntdsa.dll on a domain controller. Ntdsa.dll runs as a part of the local security authority (LSA), which runs as Lsass.exe. The architecture for Active Directory searches is shown in the following figure.

**Active Directory Searches Architecture**




The components of the architecture for Active Directory searches are described in the following table.

**Active Directory Searches Components**

| Component | Description |
|---|---|
| Directory client application | Any application that is used to send LDAP messages to an LDAP directory service |
| LDAP APIs | Interfaces through which a directory client application can submit an LDAP request. For directory clients running a Windows operating system, one of three LDAP APIs can be used. For more information, see "Active Directory Searches Interfaces" later in this section. |
| DSA (Ntdsa.dll) | The DSA, which runs on each domain controller, accepts LDAP requests through its LDAP protocol head and issues LDAP responses based on those requests. |
| LDAP protocol head | The LDAP interface on the DSA through which LDAP requests are accepted |
| LDAP request | A message that is constructed with LDAP syntax and sent by a directory client application to an LDAP directory service, such as Active Directory |
| LDAP response | A message that is constructed with LDAP syntax and sent by an LDAP directory service, such as Active Directory, to a directory client application |

**Active Directory Searches Protocols**

LDAP is the standards-based protocol that directory client applications use to submit LDAP directory requests (for example, bind, search, modify, and unbind) to Active Directory. Active Directory supports both LDAP v2, which is defined in RFC 1777, "Lightweight Directory Access Protocol," and LDAP v3, which is defined in RFC 3377. LDAP is designed to be connection oriented, and it uses Transmission Control Protocol (TCP) as its transport layer protocol. LDAP uses Internet Protocol (IP) as its network layer protocol. The Secure Sockets Layer (SSL) protocol can be used for encrypting LDAP communications.

The protocols that are used by Active Directory searches — and their relationship to the International Organization for Standardization (ISO) Open Systems Interconnection (OSI) 7-layered network protocol model — are shown in the following figure.

**Note**

- The ISO OSI model is a layered architecture that standardizes levels of service and types of interaction for computers that exchange information through a communications network. The ISO OSI model separates computer-to-computer communications into seven layers, or levels, with each level building on the standards that are contained in the levels below it.

**Active Directory Searches Protocols**




The following table describes the protocols that are used in Active Directory searches.

**Protocols Used in Active Directory Searches**

| Component | Description |
|---|---|
| LDAP | A standards-based protocol that is used for communication between directory clients and a directory service. LDAP is the primary directory access protocol for Active Directory. |
| SSL | A standards-based protocol for encrypting private communications on the Internet. When a Certificate Services infrastructure is in place (that is, when you have installed a certification authority (CA)), specifying server port 636 in an LDAP bind call results in an attempt to establish an SSL session. The CA serves as the security authority for the secure channel (Schannel) authentication protocol, which includes SSL. For this reason, some form of certificate services must be in place to use LDAP over SSL. For more information about Certificate Services, see "Certificate Services Technical Reference." |
| TCP | The connection-oriented session protocol on which LDAP runs |
| IP | A routable protocol that is responsible for the addressing, routing, and fragmenting of packets by the sending node. IP is required for Active Directory replication. |

**Active Directory Searches Interfaces**

On computers running a Windows operating system, directory client applications can use one of three different LDAP APIs. As shown in the following figure, all LDAP requests from a directory client application ultimately go through the native LDAP C API, which runs as Wldap32.dll.

**Active Directory Searches Interfaces**




The following table describes the Windows LDAP APIs that are available for Active Directory searches.

**Active Directory Searches APIs**

| API | Description |
|---|---|

| | |
|---|---|
| System.DirectoryServices<br><br>(ADSI for .NET Framework) | A namespace in the Microsoft .NET Framework that provides simple programming access to LDAP directories, such as Active Directory. This API requires the Microsoft .NET Framework to be installed. |
| Active Directory Service Interfaces (ADSI) (Ads*.dll) | A set of Component Object Model (COM) interfaces that abstracts the capabilities of directory services from different network providers (such as LDAP) in a distributed computing environment to present a single set of directory service interfaces for managing network resources. |
| Native LDAP API (Wldap32.dll) | Provides functions that enable directory client applications to search for and retrieve information from an LDAP directory service, as well as functions for modifying directory entries where such modifications are permitted. There are also functions that provide access control for servers by allowing clients to authenticate themselves. |

**The LDAP Standard**
To understand how Active Directory searches work, it is important to understand the LDAP standard. Active Directory structure is based on LDAP v3 requirements, which are defined in a series of Internet Engineering Task Force (IETF) RFCs, as summarized in RFC 3377.

LDAP features and functionality are described in four models:

- The information model, or schema, which describes the structure of information in a directory information tree
- The naming model, which describes how information is organized and referenced
- The functional model, which describes what can be done with the information
- The security model, which describes how information is protected in the directory information tree

Standardization of LDAP behavior through these models makes LDAP the preferred access protocol for all directory services.

**Differences Between LDAP v2 and LDAP v3**
Two versions of the LDAP specification exist: LDAP v2 and LDAP v3. LDAP v3 includes the following features that are not supported in LDAP v2:

- The use of UCS Transformation Format (UTF)-8 for all text string attributes to support extended character sets. Active Directory sends all responses in UTF-8-encoded form.
- Operational attributes that the directory maintains for its own use, for example, to log the date and time when another attribute is modified.
- Referrals, which enable a server to direct a client to another server. An LDAP server can return a referral to an LDAP client when the operation that is presented by the client cannot be serviced locally but the LDAP server has information about other LDAP servers that can handle the operation.
- Schema publishing with the directory, so that a client can discover what object classes and attributes a server supports.
- Extended search operations that facilitate paging and sorting of results and client-defined searching and sorting controls.
- Stronger security through the Simple Authentication Security Layer (SASL) authentication mechanism.
- Extended operations, which provide additional functionality without changing the protocol version.

**Note**

- LDAP v3 is backward compatible with LDAP v2. One requirement of an LDAP v3 server is that an LDAP v2 client should be able to connect to the server. Active Directory is compliant with both LDAP v3 and LDAP v2 standards.

**LDAP Information Model**
The LDAP information model describes the structure of information in a directory. The LDAP information model is based on the entry, which contains information about some object, for example, a person or computer. In Active Directory, an LDAP entry is referred to as an object. Entries are composed of attributes, which have a type and one or more values.

The LDAP information model organizes objects into a hierarchical tree structure. The implementation of the model is called the schema, which is a set of objects that defines the structure and content of every object that can be created in a directory service.

A default (base) schema is installed with Active Directory. The Active Directory schema is stored in the schema directory partition, and the schema is replicated to every domain controller in the Active Directory forest.

For more information about the Active Directory schema, see "Active Directory Schema Technical Reference."

**Classes and Attributes**
Classes and attributes are defined in the schema by classSchema objects (object classes) and attributeSchema objects (object attributes), as follows:

- Object classes are categories of objects that can be created in the directory. For example, users, computers, and printers are classes of objects. Every object in the directory is created as an instance of some class according to the definition that is stored in the classSchema object for the respective class.

- Object attributes are the characteristics of the object. (These characteristics are also called properties). An attribute can hold a value or values that represent some property of the object. For example, given name, surname, and e-mail address are attributes of every object of the user class, and their values can be created only as character strings. The schema specifies the attributes that are required to have values and the attributes that can have values as an option. In Active Directory, only attributes that have values assigned to them actually use storage space in the database.

- Each attribute has a syntax that is specified in the schema that determines the kind of values that are allowed in the attribute. Examples of attribute syntaxes are Unicode string, binary, and integer.

New object classes and attributes can be added to the schema, and existing objects can be modified by adding or modifying classSchema and attributeSchema objects.

Child classes inherit attributes from their parent classes. Therefore, each class builds on the attribute set of its parent class. The position in the directory tree of one object relative to another is also defined in the schema.

**Categories of Object Classes**
Different categories of object classes make it possible to define structure in the directory. The 1993 X.500 specification requires that object classes be assigned to one of three categories:

- Structural

- Abstract

- Auxiliary

**Structural classes**
Structural classes are the only classes that can have instances in the directory. That is, you can create directory objects whose class is one of the structural classes. A structural class:

- Can be used in defining the structure of the directory.

- Is derived from either an abstract class or another structural class.

- Can include any number of auxiliary classes in its definition.

**Abstract classes**
Abstract classes are templates that are used to derive new structural classes. Abstract classes cannot be instantiated in the directory. This means that no object can belong only to an abstract class; each object of an abstract class also belongs to some

structural subclass of that class. A new abstract class can be derived from an existing abstract class. Abstract classes only provide attributes for subordinate classes, which are called subclasses. A subclass contains all mandatory and optional attributes of the class from which it is derived (its superclass) in addition to those attributes that are specific to the class itself. Likewise, the subclass of that class contains all attributes of both superclasses, and so forth.

**Auxiliary classes**

Auxiliary classes are like include files; they contain a list of attributes. Adding an auxiliary class to the definition of a structural class or an abstract class adds the auxiliary class's attributes to the definition of the structural class or the abstract class. An auxiliary class cannot be instantiated in the directory, but new auxiliary classes can be derived from existing auxiliary classes. For example, the **securityPrincipal** class is an auxiliary class, and it derives its attributes from the parent abstract class called **Top**. Although you cannot create a security principal object in the directory (because auxiliary classes cannot have instances), you can create an object of the structural class user, which has the **securityPrincipal** class as an auxiliary class. The attributes of the **securityPrincipal** class help the system recognize the user object as a security account. Similarly, the **group** class has **securityPrincipal** as an auxiliary class.

**Windows 2000 Server vs. Windows Server 2003 auxiliary classes**

The behavior of auxiliary classes has changed in Windows Server 2003. In Windows 2000 Server, changes that are made to an auxiliary class affect its parent class as well as all instances of the parent object. For example, adding an auxiliary class called **pager** to the structural class **user** affects all instances of **user**, which are all of the user accounts that are created with the **user** class.

In Windows Server 2003, auxiliary classes can be assigned dynamically to individual instances of classes, rather than being applied automatically to all instances. For example, you can assign the **pager** auxiliary class to only those users who need it.

**Object Identifiers**

Object identifiers (also known as OIDs) are hierarchical, dotted-decimal numeric values that uniquely identify entries in a data model. Object identifiers are found in OSI applications, X.500 directories, Simple Network Management Protocol (SNMP), and other applications in which uniqueness is required. Object identifiers are based on a tree structure in which a designated issuing authority (such as the ISO) allocates a branch of the tree to a subauthority, which in turn can allocate subbranches. The Active Directory schema identifies the object identifier for each class, attribute, and syntax.

**Schema Versions**

The schema directory partition has an attribute called **objectVersion** that stores the schema version number for a forest. To facilitate upgrading your forest from Windows 2000 Server to Windows Server 2003, Windows Server 2003 includes a tool (Adprep.exe) that automatically upgrades the schema to the version that is required by Windows Server 2003. Before you upgrade the first domain controller, you must use this tool to prepare the domain and the forest so that the schema is upgraded before the operating system is upgraded. Running Adprep.exe prepares the forest and the domain for the Active Directory upgrade by making certain data changes, including an extension of the schema. This schema extension merges existing schema information with new schema information that is supplied by Adprep.exe. This merging of schema information does not affect any schema modifications that you may have already made in your existing environment.

**Leaf Objects and Container Objects**

Objects in Active Directory are either leaf objects or container objects. A leaf object is an object that has no child objects. The term "container" refers to one of two things:

- An object of the **container** structural class
- An object that has child objects

In the schema, a structural class defines objects that can be created as instances of the class in Active Directory. To have child objects, an object must be an instance of a class that is defined by the schema as being a possible superior of those child objects.

### Note

- A new attribute in the Windows Server 2003 schema, **msDsApproxImmedSubordinates**, stores the approximate count of immediate children of a container. You can use the ADSI Edit snap-in in Microsoft Management Console (MMC) to view the value of this attribute for any container object. Windows 2000 Server containers do not have this attribute.

**Directory Tree**

The directory tree represents the hierarchy of Active Directory objects for a given forest. The structure of the hierarchy is derived from the schema, and the directory service implements the hierarchy. The hierarchy provides the basis both for using names for navigation and for defining the scope of search requests.

Every object in Active Directory has exactly one parent, and a reference to the parent object is stored with the object. By virtue of these parent references, the hierarchy of objects that are managed by Active Directory forms a tree structure. The objects that populate the directory create this tree structure according to the rules of the schema. For example, the schema might dictate that a given class of object can be the child of one class but not the child of another class.

The following are architectural restrictions and requirements in the directory tree:

- Domain objects, which are containers, can be children only of other domain objects. For example, a domain cannot be the child of an organizational unit (OU).

- The root of the directory tree is a DSA-Specific Entry (DSE) called the rootDSE, or the directory root. The rootDSE is an object that has no hierarchical name or schema class, but it does have a set of attributes that identify the contents of a given domain controller. Therefore, the rootDSE constitutes the root of the directory tree from the perspective of the domain controller to which you are connected.

- Below the rootDSE, every directory has a root domain, which is the first domain that is created in a forest. This domain always has a child container called the Configuration container, which contains configuration data for the forest.

**rootDSE**
An LDAP directory service is also referred to as a DSA. In Active Directory, servers that host DSAs are domain controllers.

At the root of an LDAP directory tree is a DSE (the rootDSE), which is not part of any directory partition. TherootDSE represents the top of the logical namespace for one domain controller. The rootDSE attributes contain information about the directory server, including its capabilities and configuration.

There is only one root for a given DSA, but the information that is stored in the root is specific to the domain controller to which you are connected. Among other things, the attributes of the rootDSE identify the following key information:

- The directory partitions (the domain, schema, and configuration directory partitions) that are specific to one domain controller

- The forest root domain directory partition

In this way, the rootDSE provides a "table of contents" for a given domain controller. For more information about rootDSE attributes, see "Data Store Technical Reference."

**LDAP Naming Model**
The LDAP naming model describes how information is organized and referenced in an LDAP directory. LDAP requires that names of directory objects be formed according to RFC 1779, "A String Representation of Distinguished Names," and RFC 2247, "Using Domains in LDAP/X.500 Distinguished Names."

For more information about the LDAP naming model, see RFC 1779 and RFC 2247 in the IETF RFC Database.

**Distinguished Name**
Every object in Active Directory has a distinguished name (also known as DN). A distinguished name uniquely identifies an object by using the name of the object, plus the names of the container objects and domains that contain the object. Therefore, the distinguished name identifies the object as well as its location in a tree. The distinguished name is unambiguous (that is, it identifies one object only) and unique (that is, no other object in the directory has this name). It contains enough information for an LDAP client to retrieve the object's information from the directory.

For example, a user named Jeff Smith works in the marketing department of a company as a promotions coordinator. His user account is created in an OU that stores the accounts for marketing department employees who are engaged in promotion activities. The root domain of the company is proseware.com, and the local domain is noam.proseware.com. The distinguished name for this user object is:

cn=Jeff Smith,ou=promotions,ou=marketing,dc=noam,dc=proseware,dc=com

**Relative Distinguished Name**
The relative distinguished name (also known as the RDN) of an object is the part of the distinguished name that is an attribute of the object itself — the part of the object name that identifies this object as unique within a container. For the example in the previous paragraph, the relative distinguished name of the user object:

cn=Jeff Smith,ou=promotions,ou=marketing,dc=noam,dc=proseware,dc=com

is cn=Jeff Smith.

The following figure illustrates the relative distinguished names that make up the distinguished name of the user object Jeff Smith.

**Relative Distinguished Names That Make Up a Distinguished Name**




The maximum length that is allowed for a relative distinguished name is 255 characters, but attributes have specific limits that are imposed by the directory schema. For example, in the case of the common name (cn), which is the attribute type that is often used for naming the relative distinguished name, the maximum number of characters that is allowed is 64.

Active Directory relative distinguished names are unique within a container; that is, Active Directory does not permit two objects with the same relative distinguished name under the same parent container. However, two objects can have identical relative distinguished names but still be unique in the directory because, within their respective parent containers, their distinguished names are not the same. For example, the object cn=Jeff Smith,cn=users,dc=noam,dc=proseware,dc=com is recognized by LDAP as being different from cn=Jeff Smith,ou=marketing,dc=noam,dc=proseware,dc=com.

The relative distinguished name for each object is stored in the Active Directory database. Each object in the directory contains a reference to the parent of the object. An LDAP operation can construct the entire distinguished name by following these references to the root.

**Naming Attributes**

Each portion of the distinguished name is expressed as attribute_type=value. The attribute type that is used to describe the object's relative distinguished name (in the Jeff Smith example, cn) is called the naming attribute. In Active Directory, instances of classes have a default mandatory naming attribute that is defined in the schema. For example, part of the definition of the class **user** is the attribute **cn** (**Common-Name**) as the naming attribute. Therefore, the relative distinguished name for user Jeff Smith is expressed as cn=Jeff Smith.

Classes that do not define a naming attribute inherit the naming attribute from their parent class. If you create a new class in the Active Directory schema (that is, if you create a new classSchemaobject), you can use the optional **rDNAttID** attribute to specify the naming attribute for the class.

The following table shows the naming attributes that are used in Active Directory.

**Default Active Directory Naming Attributes**

| Object Class | Naming Attribute Display Name | Naming Attribute LDAP Name |
|---|---|---|
| **user** | **Common-Name** | cn |
| **organizationalUnit** | **Organizational-Unit-Name** | ou |
| **domain** | **Domain-Component** | dc |

Other naming attributes that are described in RFC 2253, "Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names," such as o= for organization name and c= for country/region name, are not used in Active Directory.

The use of distinguished names, relative distinguished names, and naming attributes is required when you are programming

for LDAP and using ADSI or other scripting or programming languages. Active Directory tools, such as Active Directory Users and Computers, do not require you to enter such values, nor do they display these values. However, LDAP editors, such as ADSI Edit in Support Tools, require input and display output in the LDAP distinguished name format.

**Using different naming attributes for users to avoid naming collisions**

To ensure data integrity, Active Directory requires that relative distinguished names be unique in a container. By default, the **user** class uses **Common-Name** (**cn**) as the naming attribute, which ties the test for uniqueness to the user name. The combination of these two restrictions can result in naming collision problems in large deployments. For example, a very large company might want to create user accounts in the same OU where, as a result of the high incidence of certain common names, many user objects have identical first and last names and, therefore, identical relative distinguished names. In this scenario, it is helpful to be able to use a different naming attribute that guarantees uniqueness, such as an employee ID that is created by the human resources department.

The **inetOrgPerson** object class is a general-purpose object class that holds attributes about people, and it is defined in RFC 2798, Definition of the inetOrgPerson LDAP Object Class. A solution is provided in the Windows Server 2003 schema so that administrators can delete **inetOrgPerson** (which uses **cn** as the naming attribute in the default schema) and re-create it using any attribute as the naming attribute. For example, instead of **cn**, the attribute **emplID** can be used as the naming attribute. You can choose the attribute and select one that will guarantee that there are no naming collisions.

For more information about **inetOrgPerson**, see Active Directory Schema Technical Reference.

**Object identity**

In addition to its distinguished name, every object in Active Directory has an identifier that is unique. This identifier is called the globally unique identifier (GUID). A GUID is a unique 128-bit number that is assigned by the DSA when the object is created. Objects might be moved or renamed within a forest, but their GUID never changes. The GUID is stored in an attribute, **objectGUID**, which is present on every object. The **objectGUID** attribute is protected so that it cannot be altered or removed. When you store a reference to an Active Directory object in an external store (for example, in a database, such as Microsoft SQL Server), you should use the **objectGUID** value to represent the object uniquely.

**Active Directory Name Formats**

In addition to distinguished names, other formats for providing object names are supported by Active Directory. These formats accommodate the different forms that a name can take, depending on its application of origin.

The LDAP distinguished name in the example earlier in this section can be used to illustrate the other Active Directory name formats. This example distinguished name identifies the account of user Jeff Smith in the OU Promotions, in the OU Marketing, in the domain noam.proseware.com:

cn=Jeff Smith,ou=promotions,ou=marketing,dc=noam,dc=proseware,dc=com

The following formats are supported by Active Directory, and the format examples are based on the LDAP distinguished name of Jeff Smith:

- LDAP Uniform Resource Locator (URL). Active Directory supports access through the LDAP protocol from any LDAP-enabled client. LDAP URLs are used in scripting. An LDAP URL names the server that holds Active Directory and the attributed name of the object (the distinguished name), for example:

  ldap://server1.proseware.com/cn=Jeff Smith,ou=promotions,ou=marketing,dc=noam,dc=proseware,dc=com

- Active Directory canonical name. By default, the user interface (UI) for certain Active Directory administrative tools displays object names as canonical names. When you view the properties of objects in Active Directory Users and Computers and Active Directory Sites and Services, the **Object** tab shows the canonical name of the selected object. The canonical name lists the relative distinguished names from the root downward, without the naming attribute descriptors, and it uses the Domain Name System (DNS) domain name, as follows:

  noam.proseware.com/marketing/promotions/Jeff Smith

### Note

- If the name of a container includes a forward slash character (/), the system requires an escape character in the form of a backslash (\) to distinguish between forward slashes that separate elements of the canonical name and the forward slash that is part of the container name. For example, if the name of an OU is Promotions/Northeast and the name of the domain is proseware.com, the canonical name is displayed as proseware.com/Promotions\/Northeast.

**DNS-to-LDAP Distinguished Name Mapping**
Although DNS domain names match Active Directory domain names, they are not the same thing. Active Directory names have a different format, which is required by LDAP to identify directory objects. Therefore, DNS domain names are mapped to Active Directory domain names (and back again) as described in RFC 2247.

All access to Active Directory is carried out through LDAP, and every object in Active Directory has an LDAP distinguished name. An algorithm automatically provides an LDAP distinguished name for each DNS domain name.

The algorithm provides a domain component (dc) attribute-type label for each DNS label in the DNS domain name. Each DNS label corresponds to the relative distinguished name of a DNS domain. For example, the DNS domain noam.proseware.com is translated to the LDAP distinguished name that has the form dc=noam,dc=proseware,dc=com.

**LDAP Functional Model**
The LDAP functional model describes what can be done with directory information. The functional model consists of nine operations in the following three areas:

- *Authentication* provides a method for the client to prove its identity to the DSA.
- *Interrogation* provides a method for the client to interrogate the directory information tree.
- *Update* provides a method for the client to add or modify information in the directory information tree.

The following table shows the nine basic LDAP operations according to their functional areas.

**LDAP Functional Model Operations**

| Functional Area | Operations |
|---|---|
| Authentication | 1. Bind<br>2. Unbind<br>3. Abandon |
| Interrogation | 1. Search<br>2. Compare |
| Update | 1. Add<br>2. Modify<br>3. Modify RDN<br>4. Delete |

**Authentication**
LDAP authentication consists of the following operations:

- Bind. Initiates a protocol session to the DSA. After a session is established, a method of authentication is negotiated between the DSA and the client. When the client is authenticated by the DSA, the DSA returns a bind response to the client.
- Unbind. Terminates an LDAP session between the client and the DSA.
- Abandon. Issued by the client to stop obtaining the results of a previously initiated operation.

**Interrogation**
LDAP interrogation consists of the following operations:

- Search. Used to select entries from a specific region of the directory information tree based on customized criteria called a search filter. The search operation includes several arguments, which are described in "Performing an LDAP Search" later in this section.

- Compare. Returns a Boolean response based on a comparison of an entry's attribute value with a user-supplied value.

**Update**
LDAP update consists of the following operations:

- Add. Creates an object in the directory information tree based on information that is provided by the client to the DSA. The information that is passed to the DSA must meet the conditions that are imposed on entry creation through the classes that are defined in the schema.

- Modify. Allows the client to modify an entry's attributes. Modification of attributes includes creating, modifying, and deleting the attributes.

- Modify RDN. Provides a mechanism for an entry to be moved in the directory. Modifying the relative distinguished name components of an entry effectively moves the entry to a new container in the directory.

- Delete. Provides a method for a client to remove an entry from the directory. The success of an entry modification depends on the schema's constraints and the access permission of the client.

**Extended LDAP Controls**
Active Directory supports several LDAP controls that extend the functionality of LDAP v3 beyond the standard set of nine LDAP operations. Microsoft has defined these LDAP controls to provide functionality in Active Directory that is not provided by current IETF RFCs. The rootDSE indicates all controls that are in effect for the contacted server through the object identifier values in the **supportedControl** attribute.

Extended LDAP control functionality is useful to programmers who are using LDAP to perform directory operations. Some of the operations that can be implemented by using extended controls are deleting trees, paging and sorting search results, and showing deleted objects. For more information about these operations, see "Optimizing Searches" later in this section.

**LDAP Security Model**
The LDAP security model specifies how to access information in the directory in a secure manner. RFC 2251, "Lightweight Directory Access Protocol (v3)," states that SASL mechanisms may be used with LDAP to provide security services. RFC 2222, "Simple Authentication and Security Layer (SASL)," defines SASL.

The rootDSE includes an attribute called **supportedSASLMechanisms** that lists the supported SASL security features.

RFC 2222 describes a method for providing authentication support in connection-based protocols. To use this specification, protocols have to include a command to identify and authenticate users to a server and, optionally, to negotiate protection of subsequent protocol interactions. If protection is negotiated, a security layer is inserted between the protocol and the connection. RFC 2222 explains how a protocol specifies such a command, and it defines several mechanisms for use by the command, as well as by the protocol, that carry a negotiated security layer over the connection.

Active Directory supports SASL mechanisms, including Kerberos V5 and NTLM. To view the supported SASL mechanisms, you can query the rootDSE. When you perform the query, Active Directory returns the **supportedSASLMechanisms** attribute and other attributes. The **rootDSE** attribute value is **supportedSASLMechanisms: GSSAPI; GSS-SPNEGO**, where GSSAPI = Kerberos and GSS-SPNEGO = NT Negotiate (Kerberos, NTLM, and so on).

The Generic Security Service Application Program Interface (GSS-API), which is defined in RFC 2078, Generic Security Service Application Program Interface, Version 2, provides security services to callers in a generic manner that can be supported with a range of underlying mechanisms and technologies. This enables applications (at the source level) to be ported to different environments. RFC 2078 defines GSS-API services at a level that is independent of the underlying mechanism and programming language environment.

RFC 1510, "The Kerberos Network Authentication Service (V5)," defines an authentication process that provides a method for verifying the identities of principals, such as workstation users and network servers, on an open network. For authentication purposes, clients use Kerberos tickets, which represent the clients' network credentials. Clients obtain the tickets from the

Kerberos Key Distribution Center (KDC), and they present these tickets when a network connection is established. Kerberos represents the identity of a client by using its domain name, user name, and password.

The Active Directory security infrastructure also supports the following primary security protocols:

- The NTLM authentication protocol is provided to support Windows NT version 4.0 and earlier. NTLM will continue to be supported and used for pass-through network authentication, remote file access, and authenticated Remote Procedure Call (RPC) connections to previous versions of Windows NT.

- Distributed Password Authentication (DPA) is the authentication protocol that is used by many Internet membership organizations, such as MSN and CompuServe. This authentication protocol is part of Microsoft Commercial Internet System (MCIS) services, and it is designed specifically to enable users to use the same Internet membership password to connect to various Internet sites that are part of the same membership organization. The Internet content servers use MCIS as an authentication service, and users can connect to multiple sites without reentering their passwords.

The following table lists the LDAP authentication methods that are supported by Active Directory.

**LDAP Authentication Methods**

| **LDAP Authentication Method** | **Description** | **Credential** |
|---|---|---|
| SIMPLE | Authentication with a simple plaintext password | A string containing the user's password |
| NTLM | An authentication protocol that is used on stand-alone systems and on networks that include systems running the Windows NT 4.0 operating system | An array of strings containing the domain name, user name, and encrypted password |
| DPA | Distributed password authentication, which is used by the Microsoft Membership System component of the MCIS | Encrypted password |
| NEGOTIATE (GSS-API) | Does not provide any authentication services; instead, it chooses the most appropriate authentication method from a list of available services and passes all authentication information on to that service. | To log on as the current user, set the *dn* and *cred* parameters to NULL. To log on as another user, pass a pointer to a SEC_WINNT_AUTH_IDENTITY structure with the appropriate user name and password. |
| SSPI | This constant (Security Support Provider Interface (SSPI)) is obsolete, and it is included for backward compatibility only. Using this constant selects GSS-API. | Same as NEGOTIATE |

**Note**

- If a directory client attempts to bind to a directory server without specifying credentials, an anonymous bind is attempted. Active Directory does not accept anonymous binds by default. However, Active Directory can be configured to accept anonymous binds. For more information, see "Anonymous queries" later in this section.

Although the LDAP security model does not include mechanisms for access control, Active Directory provides access control in the form of access control lists (ACLs) on directory objects.

**LDAP Encoding**

In accordance with RFC 2251, LDAP protocol elements are encoded on the wire using the Basic Encoding Rules (BER) of Abstract Syntax Notation One (ASN.1).

**Active Directory Searches Processes and Interactions**

Performing an Active Directory search involves several processes, including finding a server, establishing a connection, authenticating, and actually performing the search. In addition, performing an Active Directory search may involve the use of an LDAP referral, which is a reference pointing to a location where data that is searched for might be found. When a search is performed, several features may be used to optimize the search, including indexing, LDAP query policy, the global catalog, query limits, and paging.

**Searching Active Directory**
Searching is the most common directory activity. To complete an LDAP search, a directory client must complete a series of steps, as follows:

- Find an LDAP directory server
- Establish a connection
- Authenticate against (bind to) the LDAP directory server
- Perform a search

Windows Server 2003 provides several directory client tools for searching Active Directory, including the following:

- Active Directory Users and Computers (a Microsoft Management Console (MMC) snap-in that is installed by default on domain controllers)
- ADSI Edit (an MMC snap-in that is installed with Windows Support Tools)
- Dsquery (a command-line tool that is installed by default on Windows Server 2003)
- Ldp (an LDAP graphical user interface (GUI) tool that is installed with Windows Support Tools)
- Search Companion (which appears on the **Start** menu as the **Search** command)

In addition to these tools, many non-Microsoft applications are available for searching Active Directory as well as other LDAP directory servers. The information that you provide the search tool to complete the steps that are described in the following sections depends on the tool that you are using. For example, some directory client tools and applications find an LDAP directory server and establish a connection automatically.

**Finding a Server**
The first step that a directory client must take in conducting an Active Directory search is to find an LDAP directory server (in other words, a domain controller) to search against. To find a domain controller, directory clients rely on DNS. When a domain controller starts up, it registers service (SRV) records in DNS that indicate that the domain controller provides LDAP directory services. To locate a domain controller, a directory client performs a DNS query for SRV records of hosts that provide LDAP directory services. For more information about how DNS is used to locate domain controllers, see "DNS Support for Active Directory Technical Reference."

**Establishing a Connection**
After it finds a directory server, a directory client must next connect to the server. A directory client can connect to an LDAP directory server by opening a session on a TCP port number on which the LDAP directory server is listening. The string that is used to establish the connection includes the fully qualified domain name (FQDN) of the LDAP directory server, along with the TCP port number of the directory server. For standard LDAP searches, directory clients connect to TCP port 389. To search the Active Directory global catalog, directory clients must instead connect to TCP port 3268.

For more information about the global catalog, see "Global Catalog Searches" later in this section. For information about TCP ports that are used by LDAP directory servers and clients, see "Network Ports Used by Active Directory Searches" later in this section.

**Authenticating (Binding)**
After a directory client establishes a communications path to the domain controller, it must bind to the domain controller to establish the logon and authentication credentials and, if necessary for Windows-based computers, set up a secure channel. (A client can also attempt to bind to a domain controller without first establishing a connection.) The bind operation identifies the connecting person, device, or application to the server by providing a distinguished name and some type of authentication credential, such as a password. The exact credentials depend on the authentication method that is being used.

LDAP v3 enables the client to negotiate with the LDAP server to determine the best available security package. If no security package is available, the bind is a simple bind that uses a plaintext password. The Microsoft implementation of the LDAP API uses the NEGOTIATE flag so that the client can discover the best security package that is available. For example, a SASL mechanism, such as Kerberos V5 or NTLM, might be used. A domain controller can be also be configured to accept anonymous connections.

**LDAP bind request**
The **bind** command initiates a protocol session to the domain controller. After a session is established, a method of authentication is negotiated between the domain controller and the client. By default, Kerberos is used, but other methods can also be used. Finally, the domain controller returns a bind response to the client when the client is authenticated.

According to RFC 2251, the LDAP bind request has the following parameters:

- *Version*. A version number that indicates the version of the protocol to be used in this protocol session. Note that there is no version negotiation, and the client sets this parameter to the appropriate version.

- *Name*. The name of the directory object to which the client wants to bind. This field can take on a null value (a zero-length string) for the purposes of anonymous binds, when authentication has been performed at a lower layer, or when the client is using credentials that are established through a SASL mechanism.

- *Authentication*. Information that is used to authenticate the name, if any, that is provided in the bind request.

When a server receives a bind request, it authenticates the requesting client, if necessary. The server then returns a bind response to the client indicating the status of the authentication.

In the following Network Monitor example of an LDAP bind request, the LDAP version is 3, the Name is null, and the authentication type is SASL:

```
LDAP: ProtocolOp: BindRequest (0)
  LDAP: MessageID = 11 (0xB)
  LDAP: ProtocolOp = BindRequest
    LDAP: Version = 3 (0x3)
    LDAP: Name =
    LDAP: Authentication Type = Sasl
      LDAP: Sasl Mechanism = GSS-SPNEGO
      LDAP: Sasl Credentials
```

**LDAP authentication methods**
The Windows Server 2003 and Windows 2000 Server implementations of LDAP include support for the authentication methods that are described in "LDAP Security Model" earlier in this section.

### Note

- Authentication within and between Windows Server 2003 and Windows 2000 Server domains is performed by using either the Kerberos V5 authentication protocol (the default method) or NTLM (for Windows NT 40 and earlier). Other methods are available to other clients and external users who connect over the Internet. For more information about authentication, see "Logon and Authentication Technologies."

**LDAP bind response**
An LDAP bind response is an indication from the server regarding the status of a request for authentication of the client. If the bind is successful, the result code is "success." Otherwise, an error is reported in accordance with RFC 2251.

The following is a Network Monitor example of a successful LDAP bind response:

```
LDAP: ProtocolOp: BindResponse (1)
  LDAP: MessageID = 18 (0x12)
  LDAP: ProtocolOp = BindResponse
    LDAP: Result Code = Success
    LDAP: Matched DN =
    LDAP: Error Message =
    LDAP: Sasl Mechanism = GSSAPI
    LDAP: Sasl Credentials
```

The following is a Network Monitor example of an unsuccessful LDAP bind response:

```
  LDAP: ProtocolOp: BindResponse (1)
  LDAP: MessageID = 8 (0x8)
  LDAP: ProtocolOp = BindResponse
    LDAP: Result Code = Invalid Credentials
    LDAP: Matched DN =
```

```
LDAP: Error Message =
```

**SPNEGO**
The client and server can use simple protected negotiation (SPNEGO) to negotiate either through Kerberos V5 or NTLM, depending on the authentication mechanisms that are available to the particular client and server. With SPNEGO, both the client and server negotiate on a common secure authentication mechanism (Kerberos V5 or NTLM). For more information about SPNEGO, see "Logon and Authentication Technologies.

**Concurrent bind for Web site logons**
Internet service providers (ISPs) can use LDAP authentication for Active Directory user accounts. In this case, users do not actually have (or need) direct access to Active Directory objects. An application on the Web server uses LDAP to request user name and password verification of credentials that are presented by the site user. The application then passes the user information to Active Directory, which authenticates the user.

Domain controllers running Windows Server 2003 can perform concurrent binds. Unlike normal binds that are serialized on a single connection, concurrent binds enable applications to request multiple binds over a single LDAP connection. Concurrent binds are faster than normal binds, which require construction of a full security token for later use in access checks. In the Web site case, the user has no need for the security token.

Unlike a normal LDAP bind, a concurrent bind does not determine a user's group association or build a token. The concurrent bind determines only whether the user has a valid user name and password and that the account is enabled. Eliminating the construction of the access token permits the concurrent bind to complete in a fraction of the time that it takes a normal bind to complete.

**Determining if concurrent bind exists**
To determine if your version of Active Directory supports concurrent binds, view the values in the **supportedExtension** attribute of the rootDSE. If concurrent binds are supported, the returned attribute includes the value 1.2.840.113556.1.4.1781.

For information about concurrent binding, see "Using Concurrent Binding" in the Microsoft Platform SDK on MSDN.

**Anonymous queries**
By default, anonymous LDAP operations to Active Directory, other than rootDSE searches and binds, are not permitted in Windows Server 2003. (Active Directory in Windows 2000 Server accepts anonymous requests; a successful result depends on objects having correct user permissions in Active Directory.)

To enable anonymous binding to Active Directory in Windows Server 2003, you must change the seventh character of the **dsHeuristics** attribute on the following directory object:

CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,Root *domain in forest*

Valid values for the **dsHeuristics** attribute are **0** and **2**. By default, the **dsHeuristics** attribute does not exist, but its internal default is **0**. If you set the seventh character to **2**, anonymous clients can perform any operation that is permitted by the access control list (ACL). If the attribute is already set, do not modify any bits in the **dsHeuristics** string other than the seventh bit. If the value is not set, make sure that you provide the leading zeros up to the seventh bit. You can use Adsiedit.msc to make the change to the **dsHeuristics** attribute.

After you set the **dsHeuristics** attribute, if you want anonymous users to be able to query Active Directory, you can enable anonymous access to specific directory objects. Users gain anonymous access to Active Directory objects through Anonymous Logon, which is a special security identifier (SID) that is used to represent anonymous network callers that perform an LDAP bind with NULL credentials.

**Performing an LDAP Search**
LDAP searches are the most common LDAP operations that are performed against an Active Directory domain controller. An LDAP search retrieves information about all objects within a specific scope that have certain characteristics, for example, the telephone number of every person in a department.

The following parameters are used in LDAP to accomplish an LDAP search:

- *Search base* (the distinguished name of the search base object). Defines the location in the directory from which the LDAP search begins.

- *Search scope*. Defines how deep to search within the search base:

    - *Base* (or *zero level*). Indicates a search of the base object only.

- *One level*. Indicates a search of objects immediately subordinate to the base object but not the base object itself.
- *Subtree*. Indicates a search of the base object and the entire subtree of which the base-object distinguished name is the topmost object.

- *Filter*. Allows certain entries in the subtree and excludes others.
- *Selection*. Indicates what attributes to return from objects that match the filter criteria.
- Optional controls that affect how the search is processed.

The following figure shows the base distinguished name and the possible search scopes of an LDAP search.

**LDAP Search Base and Three Search Scopes**




The following figure shows the base distinguished name for a container object that is the search base for a search of an OU structure in Active Directory. The entire distinguished name is constructed from the relative distinguished name of an OU named Sales.

**Base Distinguished Name for an LDAP Search**




**Search filters**

By using search filters, you can define search criteria to achieve more effective and efficient searches. For example, you might be interested in all the user's whose surname is Smith, or you might want to find all the team members who report to a manager named Jeff Smith. ADSI supports LDAP search filters as defined in RFC 2254, "The String Representation of LDAP Search Filters." These search filters are represented by UTF-8 strings. The following table lists some commonly used LDAP search filters.

**Common LDAP Search Filters**

| Filter | Description |
|---|---|
| (objectCategory=*) | All objects |
| (&(objectClass=user)(!(cn=susan))) | All user objects except susan |
| (sn=sm*) | All objects with a surname that starts with sm |
| (&(objectClass=contact)(\|(sn=Smith) (sn=Johnson))) | All contacts with a surname equal to Smith or Johnson |

The search filters in the previous table use one of the following formats:

(*attribute operator value*)

or

(*operator* (*filter1*) (*filter2*))

The following table lists some of the most frequently used LDAP search filter operators.

**Commonly Used LDAP Search Filter Operators**

| Operator | Description |
|---|---|
| = | Equal to |
| ~= | Approximately equal to |
| <= | Lexicographically less than or equal to |
| >= | Lexicographically greater than or equal to |
| & | AND |
| \| | OR |
| ! | NOT |

For more information about LDAP search filter syntax, see "Search Filter Syntax" in the Microsoft Platform SDK on MSDN.

**objectClass vs. objectCategory in a search filter**

Sometimes, you might want to search Active Directory for objects of a particular type. For example, you might want to search for Active Directory objects that represent users. You can do this by searching for objects of a particular object class, using the **objectClass** attribute (objectClass=user). Or, you can search for objects of a particular category, using the **objectCategory** attribute (objectCategory=user). Because of the class inheritance hierarchy in the schema, every object in Active Directory is in fact a member of many classes — four or five, on average. For this reason, searches that use **objectClass** can be less accurate than searches that use **objectCategory**.

For example, a search filter of objectClass=user returns both user and computer objects. On the other hand, **objectCategory** usually refers to the most specific class in the object's class hierarchy. Every Active Directory object has an **objectCategory** attribute whose value is a classSchema object. For example, a search filter of objectCategory=user returns only user objects. Although **objectClass** can have multiple values, the attribute **objectCategory** has only one value.

**Note**

- Every classSchema object has an attribute called **defaultObjectCategory**, which is the object category of an instance of the class if no object category is specified by the user. For most classes, the **defaultObjectCategory** value is the class itself. In the search filter, you can specify objectCategory=X, where X is the **ldapDisplayName** of a class and LDAP automatically expands the filter to objectCategory=*defaultObjectCategory of class X*. The **objectCategory** attribute has a syntax of distinguished name, and LDAP automatically converts the value for **objectCategory** to the distinguished name format. For example, if you use objectCategory=contact in the filter, the filter changes to objectCategory=cn=Person,cn=Schema,cn=Configuration,dc=*ForestRootDomain*. (**Person** is the **defaultObjectCategory** for the class **Contact**).

For more information about class inheritance, see "Active Directory Schema Technical Reference."

**LDAP Referrals**

When a requested object exists in the directory but is not present on the contacted domain controller, resolution of the object name depends on information that is stored on that domain controller about how the directory is partitioned. In a partitioned directory, by definition, the entire directory is not always available on any one domain controller.

In its Configuration container, every domain controller has information about the other domains in the forest. The objects in cn=Partitions,cn=Configuration,dc=*ForestRootDomain* are cross-reference objects that contain information that Active Directory uses to construct the directory tree hierarchy. When an operation in Active Directory requires action on objects that might exist in the forest but that are not located in the particular domain that is stored on a domain controller, that domain controller must send the client a message that describes where to continue this action — that is, the client is "referred" to a domain controller that is presumed to hold the requested object.

An LDAP referral is a domain controller's way of indicating to a client application that the domain controller does not have a copy of a requested object (or, more precisely, that the domain controller does not hold the section of the directory tree where that object would be if, in fact, it exists) and of providing the client with a name of a server that is more likely to hold the object. The client uses the name of the server that is provided in the LDAP referral as the basis for a DNS search for a domain

controller. Ideally, referrals always reference a domain controller that does indeed hold the object. However, it is possible for the referred-to domain controller to generate yet another referral, although it usually does not take long to discover that the object does not exist and to inform the client that the object does not exist. Active Directory returns referrals in accordance with RFC 2251.

Clients do not have to know the name or location of a child domain to contact a domain controller in that domain. They can query the root domain and reach the appropriate domain controller by being referred there. Two situations generate this type of domain controller response:

- An external referral, in which the base distinguished name of the requested object is not in this directory, but the domain controller holds information about another LDAP directory where the requested object might be found.

- A subordinate referral, in which the base distinguished name of the requested object is in this directory, but the directory partition that contains the requested object is not stored locally.

Every domain controller contains information about how the directory is partitioned, and this information can be used in conjunction with DNS to find the Active Directory domain that contains a particular object.

**Knowledge References**
Active Directory stores information about the existence and location of directory partitions in a forest, including the names of the directory partitions and the names of domain controllers that hold replicas of those directory partitions. Active Directory uses this information (known as knowledge references) to generate referrals to other domain controllers.

Active Directory uses three kinds of knowledge references to generate referrals to other domain controllers:

- A subordinate reference, which is knowledge of a directory partition (or partitions) directly adjacent in the naming hierarchy to a directory partition that is held by the domain controller.

- A cross-reference, which is knowledge of one directory partition and which is stored in a cross-reference object in the Partitions container. On a specific domain controller, the combination of all cross-references provides knowledge of all directory partitions in the forest, regardless of the locations of those directory partitions in the directory tree.

  **Note**

  - The state of cross-reference knowledge at any specific time is subject to the effects of directory replication latency.

- A superior reference, which is knowledge of a specifically designated location that is used whenever the domain controller has no knowledge of the search base that is used in a search.

Knowledge references are the glue that holds the pieces of the distributed directory together. Because Active Directory is logically partitioned and directory partitions are the discrete components of the directory that replicate between domain controllers, either all objects in a directory partition are present on a particular domain controller or no objects in the directory partition are present on the domain controller. References have the effect of linking directory partitions together, which enables operations such as searches to span multiple domain directory partitions.

In Active Directory, referrals are generated when a client requests that the directory locate an object where, based on the position at which the search begins, no copy exists in a local directory partition. When Active Directory can determine definitively from the information it holds that no such object exists in the directory — rather than that it might exist somewhere else, even though no copy exists locally — instead of sending a referral, the directory returns an error message to the client that no such object exists in the forest.

For information about replication of directory partitions, see "Active Directory Replication Model Technical Reference."

**Subordinate References**
When a client requests a search, the domain controller searches all objects at or below the search base in the directory partition that the domain controller holds. If a subtree search has a search base that includes child partitions, the domain controller uses information that is provided by subordinate references to return referrals (called subordinate referrals) to these partitions on other domain controllers.

Subordinate referrals are returned as part of the data that is returned from the distinguished name of the directory partition

that is named in the search base. The referral contains the distinguished name of the subordinate directory partition and the access point to which queries can be referred. An access point consists of a DNS name and a port number, which is the information that is required to contact a specific LDAP server. Access points are generated from information that is contained in the cross-reference object.

**Cross-References**

Cross-references are stored as directory objects of the class **crossRef** that identify the existence and location of all directory partitions, irrespective of their location in the directory tree. Cross-reference objects are used to generate referrals to other directory partitions in the forest and to external directories.

Cross-references enable every domain controller to be aware of all directory partitions in a forest, rather than only the directory partitions that a particular domain controller holds. Because cross-reference objects are stored in the Configuration container, they are replicated to every domain controller in the forest.

Cross-reference objects store the following information:

- **nCName** . The distinguished name of the directory partition that the crossRef object references. ("nC" stands for "naming context," which is a synonym for "directory partition.") The combination of all of the **nCName** properties in the forest defines the entire directory tree, including the subordinate and superior relationships between directory partitions.

- **dNSRoot**. The DNS name of the domain where servers that store the particular directory partition can be reached. This value can also be a DNS host name.

Cross-reference objects are created in two ways:

- Internally, by the system, to refer to known locations that are within the forest.

- Externally, by administrators, to refer to locations that are external to the forest.

**Internal cross-references**

An internal cross-reference is a cross-reference object that is created by the system. For every directory partition in a forest, there is an internal cross-reference object in the Partitions container (cn=Partitions,cn=Configuration,dc=*ForestRootDomain*). When you create a new forest, the Active Directory Installation Wizard creates the following three directory partitions:

- First domain directory partition

- Configuration directory partition

- Schema directory partition

### Note

- On a domain controller running Windows Server 2003, one or more application directory partitions might be created in addition to the domain, configuration, and schema directory partitions.

For each of these partitions, a cross-reference object is created automatically. Thereafter, when a new domain or application directory partition is created in the forest, another directory partition is created and the respective cross-reference object is created.

### Note

- Whenever you create a new domain or application directory partition, the appropriate cross-reference objects are created automatically.

Because these cross-reference objects are located in the Configuration container, they are replicated to every domain controller in the forest, and therefore every domain controller has knowledge of the name of every partition in the forest (as well as their superior and subordinate relationships to each other). By virtue of this knowledge, any domain controller can generate referrals to any other domain in the forest, as well as to the schema and configuration directory partitions.

**External cross-references**
An external cross-reference is a cross-reference object that you create manually to provide the location of an object that is not stored in the forest. If your LDAP clients submit operations for an external LDAP namespace against servers in your forest and you want your forest's servers to always refer these clients to the correct location, you can create a cross-reference object for that directory in the Partitions container.

There are two ways that external cross-references are used:

- To reference external directories by their disjoint directory name, that is, a name that is not contiguous with the name of the local directory tree. In this case, when you create the cross-reference, you create a reference to a location that is not a child of any object in the local directory.

- To reference external directories by a name that is within the Active Directory namespace, that is, a name that is contiguous with the name of this directory tree. In this case, when you create the cross-reference, you create a referenceto a location that is a child of an existing directory object.

### Note

- An external directory that is stored on a domain controller running Windows Server 2003 or Windows 2000 Server does not require an explicit cross-reference object. Because the domain component (dc=) portions of the distinguished names of all Windows Server 2003 domains match their DNS addresses and because DNS is the worldwide namespace, all domain controllers running Windows Server 2003 or Windows 2000 Server can generate external referrals to each other automatically, assuming that the appropriate credentials and trust relationships exist.

**Creating External Cross-References**
The only time that you have to create a cross-reference object is when you want to extend a search to a directory outside the forest that is a non-Windows LDAP directory service. In this case, you can use an LDAP editor such as ADSI Edit to create objects of the class **crossRef** (in the Partitions container) that reference external directories.

When you create a cross-reference object, you must provide the values for three attributes:

**cn**
The name that describes the directory. For example, for the domain noam.proseware.com, your **cn** value might be noam or something else that describes that domain, such as NorthAmerica. This value is for administrative convenience; it has no impact on the system.

**nCName**
The distinguished name of the domain directory partition to which your cross-reference refers. If the domain name is noam.reskit.com, the value of **nCName** is dc=noam,dc=proseware,dc=com.

**dnsRoot**
The DNS host name of an LDAP server in the domain that is identified by **nCName** (for example, server1.noam.proseware.com). The value of **dnsRoot** can also be the domain name if you do not want to specify a server.

### Note

- You must be able to resolve (use the **ping** command to test) the name in **dnsRoot**, which does not necessarily name another system running Windows Server 2003; it might be the DNS address of an LDAP server instead of a domain controller. If the directory partition is a Windows Server 2003 domain in another forest, automatically generated knowledge is usually sufficient and no external cross-reference is required.

ADSI Edit, an MMC snap-in and Active Directory administrative tool, provides a convenient GUI for creating cross-reference objects.

To make use of cross-references, clients must be able to follow ("chase") referrals that are returned. Windows Address Book chases referrals by default. When you are using ADSI programmatically (for example, by using Active Data Objects (ADO) to search), you must specify whether to chase referrals.

### Note

- When you set the LDAP referral chasing option (**LDAP_OPT_REFERRALS**) to **ON**, the LDAP client (Wldap32) ensures that plaintext passwords are not sent over the network and that the same level of security that is required of the originating

query server is achievable from the referral server. If the referral server does not support the security package that was used to bind to the query server or if the referral server does not support the same level of integrity or privacy as the query server, Wldap32 does not chase the referral.

- If no bind is performed on the originating query server, or if a simple (plaintext) bind is performed on the originating server without the use of SSL encryption to protect the passwords, the LDAP client performs an anonymous bind to the referral server. The latter behavior protects passwords from being sent inadvertently over the network in plaintext as the result of chasing a referral.

**Creating an external cross-reference for an external location**
To create a cross-reference to an external directory by referencing an external location, give the **nCName** attribute a value that is the name of the actual external directory. For example, an external LDAP directory might use X.500 naming (such as o=*Organization Name*,c=*Country/Region*), which can be used for the value of the **nCName** attribute. Queries for this directory must specify the external object by name in the search base distinguished name. A request for a referral to such a location might come in the form of an LDAP URL that is embedded in an e-mail message or from an application that specifically names the directory distinguished name.

**Creating an external cross-reference for an internal location**
If you want a subtree search of a portion of your directory to always include an external LDAP directory that is not a Windows Server 2003 or Windows 2000 Server directory service, you can create a cross-reference to the external directory for an internal location. To create an internal location that references an external directory, give the **nCName** attribute of the cross-reference object a value that is an immediate child object of an existing directory object and that also matches the distinguished name of the external directory. Choose the location according to where you want the external directory to be locatable in Active Directory.

This type of cross-reference is especially useful for integrating dynamic directories smoothly. For example, you might use an instant messaging application such as Microsoft NetMeeting conferencing software to publish a list of current or planned conference calls. LDAP is an effective protocol for querying such a published list. However, short-lived, highly volatile data is inappropriate for Active Directory storage. Therefore, you might use an in-memory, nonreplicated LDAP server (one that can store volatile data) at an arbitrary point in the namespace. This "volatile" directory service can then be configured to inhabit a name inside your organization's Active Directory namespace, and it can be made available to users through a cross-reference for an internal location in Active Directory. For example, users can use this location to find the list of instant messaging conversations by directory tree navigation.

**Note**

- On Windows Server 2003 domain controllers, application directory partitions can be used to provide directory storage for dynamic data. For more information about application directory partitions, see "Data Store Technical Reference."

Suppose your domain name is proseware.com and you install your messaging application on an LDAP server that does not use Active Directory, which is named vds.it.proseware.com. On that server, you create a directory for your volatile data, such as cn=conversations,dc=proseware,dc=com. Then, on your Active Directory domain controller, you create a cross-reference object and use the following attribute values:

- cn=conversation server
- nCName=cn=conversations,dc=proseware,dc=com
- dnsRoot=vdserver.it.proseware.com

When a user performs a subtree search of dc=proseware,dc=com, the client receives results from the domain controllers running Windows 2000 Server or Windows Server 2003 and also a subordinate referral to the volatile directory service server at vdserver.it.proseware.com, which instructs the client to continue the LDAP search from cn=conversations,dc=proseware,dc=com and below on that server.

**Superior References**
A superior reference is the fully qualified DNS name of a directory partition that is stored in the **superiorDNSRoot** attribute on the crossRef object for the forest root domain (the first domain that is created in the forest). A domain controller uses its superior reference to construct a referral only when a search base does not match any directory partition that is defined by the cross-reference objects. A superior reference contains no directory tree information; it consists of only an access point to which

otherwise unanswerable queries can be referred.

By default, **superiorDNSRoot** does not store a value, but the directory uses the dc= components of the search base distinguished name to construct the equivalent of a superior referral. You can use the **superiorDNSRoot** attribute to define a location to which all queries that cannot be resolved are sent.

**Optimizing Searches**

Because searches against Active Directory require CPU and memory resources, searches should be performed in ways that minimize their performance impact on domain controllers. You can use several techniques to reduce the impact of directory searches on domain controllers running Windows Server 2003.

Extensions to the base LDAP API, in the form of LDAP v3 controls, make it possible to sort results and set various limits on search operations. You can process search results by paging and by sorting. Paging and sorting are supported in Windows Server 2003 as new LDAP v3 control extensions.

**Effects of Indexing on Search Behavior**

You can improve search behavior by indexing the attributes that are most likely to be searched. The **searchFlags** attribute of an attributeSchema object indicates whether the attribute is indexed and the nature of the indexing.

The **searchFlags** attribute is an integer value whose least significant bits indicate whether the attribute is indexed and in what way. Use the OR operator to combine the bit values to achieve the appropriate indexing behavior.

Values in the **searchFlags** attribute have the following effects:

- 1 = Index over attribute only. Bit 0 must be set (0001). You can use the Active Directory Schema MMC snap-in to set this flag.
- 2 = Index the attribute in each container (this value is used in conjunction with 1). Bit 0 and bit 1 must be set (0011). You can use Active Directory Schema to set this flag.
- 4 = Add this attribute to the ANR set (this value is used in conjunction with 1). Bit 0 and bit 2 must be set (0101).
- 8 = Preserve this attribute on logical deletion (that is, make this attribute available on tombstones). Bit 3 must be set (1000).
- 16 = Copy the value for this attribute when the user object is copied. You can use this value to set up a user object template so that you can create multiple users by making copies of an original and providing values for only those attributes that are unique to each user. Bit 4 must be set (10000). For example, set this flag for attributes, such as **company** and **department**, that you want to copy, but do not set attributes that must be unique, such as **objectSid** and **sAMAccountName**.
- 32 = Index this attribute for medial-string search (this value is used in conjunction with 1). Bit 5 and bit 0 must be set (100001). This indexing capability is provided on domain controllers running Windows Server 2003 when the schema has been updated for Windows Server 2003 functionality.

**Optimizing container searches**

An attribute can be indexed for the attribute only or for both the attribute and the containers that are on the same search level. The value of the **searchFlags** attribute on the attributeSchema object indicates how the attribute is indexed.

One-level searches are confined to the immediate children of the base object of the search. An index on containers, as well as the attribute, optimizes the search where values are expected to be common across containers. For example, suppose you store phone book data for the United States in Active Directory, organized by state and then by region. If you want to search for a Jeff Dulong in northern Idaho, you go to the Northern Idaho container and do a one-level search for (&(givenName=Jeff)(sn=Dulong)). If the **givenName** or **sn** attributes are indexed on the attribute only (flag 1 on the **searchFlags** attribute), the search must examine all the Dulongs in the United States, all the Jeffs in the United States, or all the people in northern Idaho. With the **searchFlags** value of 2 set (index the container and the attribute), the search examines only the Dulongs or only the Jeffs in northern Idaho, which results in a substantial improvement in the time that the search takes to complete.

**Note**

- Active Directory Schema provides an option to index an attribute for container searches.

**ANR**
Ambiguous name resolution (ANR) is the process of searching for a string value in a set of attributes by using one filter of the form (ANR=*string*).

**ANR attribute set**
Attributes are designated as members of the ANR attribute set according to the value of the **searchFlags** attribute on the attributeSchema object that defines the attribute.

By default, the following attributes are members of the ANR attribute set, and they are evaluated when you enter an ANR search string in an LDAP filter:

- **givenName** (first name)
- **sn** (surname, or last name)
- **displayName** (the name that is given to the object when it is created)
- **RDN** (the relative distinguished name of the object)
- **legacyExchangeDN** (for enterprises that have upgraded a Microsoft Exchange installation to a later version of Exchange that is synchronized with Active Directory, the distinguished name of the old Exchange mailbox that corresponds to the user in Active Directory)
- **physicalDeliveryOfficeName** (for example, Building A, Suite 1234)
- **proxyAddresses** (the collection of e-mail addresses over all e-mail address spaces that the Exchange server knows about)

When the (ANR=*string*) filter is encountered by Active Directory, the filter is expanded to include a search of every attribute in the ANR set. For example, an ANR search is implemented when you use the **Check Names** option in Microsoft Outlook that searches against a Microsoft Exchange mail server. Because **physicalDeliveryOfficeName** is in the default ANR attribute set, you can type the building number and room number as it appears in the attribute value into the **To** line of a mail message, click **Check Names**, and the Address Book responds with the user name of any account that contains that value in the **physicalDeliveryOfficeName** attribute of a user object.

In Active Directory Users and Computers, you can use the **Filter** or **Find** options on the toolbar and select the **Custom** and **Advanced** options to enter an ANR filter. You can also use ADSI Edit to perform filtered LDAP searches.

**ANR matching of an embedded space**
For the **givenName** and **sn** attributes, if a space is embedded in the string that is presented in an ANR filter, the string is split at the first such space, and each piece of the string is evaluated separately. You can use this feature to search for a user object by providing the first few characters of the first name (**givenName**) and the first few characters of the last name (**sn**). For example, the filter (ANR=mar ha) finds all objects that have a **givenName** attribute value that begins with "mar" and an **sn** attribute value that begins with "ha". In this example, the filter returns a user who has a **givenName** attribute of Mark and an **sn** attribute of Hanson, as well as a user who has a **givenName** attribute of Mark and an **sn** attribute of Harrington.

**First/last and last/first functionality**
For the **givenName** and **sn** attributes, if a space is embedded in the string that is presented in an ANR filter, the filter is expanded to find the values in both respective positions. For example, the filter (ANR=c d) finds both the user David Campbell and the user Craig Dewer.

**Expanded ANR filter**
When an ANR filter is encountered by Active Directory in an LDAP search, the filter is expanded to construct an OR operation on the string for every attribute in the ANR set. For the **givenName** and **sn** attributes, the first/last and last/first matching are also applied.

The ANR filter of the form (anr= *xxxyyy*) is expanded in the following manner:

```
(| (displayName=xxx yyy*)
(givenName=xxx yyy*)
(physicalDeliveryOfficeName=xxx yyy*)
(proxyAddresses=xxx yyy*)...
(sn=xxx yyy*)
```

```
(& (givenName=xxx*)(sn=yyy*))
(& (givenName=yyy*)(sn=xxx*)))
```

The last three lines of the expanded filter are the portion of the filter that evaluates first name and last name.

**Suppressing first/last and last/first functionality**
You can turn off first/last and last/first functionality if you do not need this type of evaluation and you do not want CPU cycles to be used for it. For example, in the filter (ANR=c d), which finds both the user David Campbell and the user Craig Dewer, you can specify that ANR will always find only Craig Dewer, perhaps because names in your database never occur in any order other than first name/last name and no searches are likely to be looking for anything other than first name/last name. Likewise, if names are always represented in last/first format, you can turn off the first/last functionality. If you do not have values that apply to the first names and last names of people, you can disable both first/last and last/first functionality.

The **dsHeuristics** attribute on the Directory Service object (cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration,dc=*ForestRootDomain*) contains a string value that governs the use of first/last and last/first functionality in the first two character positions. The default value of **dsHeuristics** is 00, which indicates that both functions are enabled. (For all positions, 0 means "perform the default behavior.") The first character in the string governs first/last functionality; the second character in the string governs last/first functionality.

You can modify the first two characters of the string to suppress either one or both functionalities as follows:

- 10 = Suppress first/last functionality. (This value can also be written 1 or 10000 because both mean that only the first character's behavior must be something other than the default.)
- 01 = Suppress last/first functionality.
- 11 = Suppress last/first and first/last functionality.

**Adding attributes to the ANR set**
You can add attributes to the default ANR set by setting a flag on the attributeSchema object. By using ADSI Edit, connect to the schema, and then open the properties on the attribute that you want to add. Set the **searchFlags** attribute value to a value that represents a bitwise OR operation of 4 and 1 to the existing value. The value 4 adds the attribute to the ANR set; the value 1 indexes the attribute.

### Note

- Limit the number of attributes that you add to the ANR set as much as possible, because adding an attribute to the ANR set affects the performance of any search that is performed against that attribute on any domain controller.

**Searching for deleted objects**
When an Active Directory object is deleted, it is stored as a tombstone in the Deleted Objects container of the respective directory partition, with the exception of some objects from the configuration partition, which remain in their original location, even after they are deleted. When an object is marked as a tombstone, the attributes are stripped from the object, with the exception of **objectGuid**, **objectSid**, **distinguishedName**, **nTSecurityDescriptor**, and **usnChanged**, and any additional attributes that are marked in the schema as "preserve on delete." For more information about tombstones, see "Active Directory Replication Model Technical Reference."

Tombstones are stored for a configurable period of time to allow replication of the deletion to occur. All domain directory partitions, as well as the configuration directory partition, have a hidden Deleted Objects container that stores tombstones until they are permanently removed from the directory by garbage collection. By using the **Show Deleted Object** control (controlType = 1.2.840.113556.1.4.417) in conjunction with search commands, you can view Active Directory objects that have been deleted but not yet garbage collected.

To retrieve tombstone objects, use Ldp.exe to list the contents of the Deleted Objects container for a directory partition.

**Medial-string searches**
In Windows 2000 Active Directory, efficient substring searches are supported for query filters with initial strings (for example, cn=abc*). However, medial-string searches (for example, cn=*def*) and final-string searches (for example, cn=*ghi), although supported, are slow as a result of the lack of appropriate indexing. Active Directory in Windows Server 2003 provides support for selectively indexing attributes that have **DirectoryString** syntax for efficient medial-string and final-string searching. For example, if you index the surname (**sn**) attribute for medial-string searching, you can find all users named McLaughlin, MacLaughlin, and McLaughlan by using the search string cn=*claughl*. You can use the **searchFlags** attribute with a value

of 33 (1 AND 32) to implement both attribute-level and medial-string indexing.

**Note**

- The **searchFlags** value of 32 makes only medial-string searching efficient. To ensure efficiency for initial-string searches as well, use the value 33.

This new type of indexing, called Tuple indexing, is available only on domain controllers running Windows Server 2003, and it is available at any forest functional level. With Tuple indexing, you can see improved performance of medial-string searches on strings that are *at least 3 characters* in length. For example, assuming that the attribute **cn** has been indexed for medial-string search, the performance of searches by filters such as cn=*mic* or cn=*mich* is substantially faster than searches that are run without indexing for medial-string search. However, searches on medial strings that have a shorter length, such as cn=*m* or cn=*mi*, might not see any improvement in performance compared to searches that are not indexed for medial search.

Database indexes are generated separately on each domain controller. Therefore, Tuple indexing only affects domain controllers running Windows Server 2003, and it does not affect any Windows 2000 Server–based domain controllers that are present in the domain or forest. Because searches of the type cn=*def* are possible on Windows 2000 Server–based domain controllers, there are no backward compatibility issues with Tuple indexing and Windows 2000 Server domain controllers. You can perform medial-string searches on both domain controllers running Windows 2000 Server (without the benefit of Tuple indexing) and domain controllers running Windows Server 2003 (with the benefit of Tuple indexing).

**Performance considerations**

Tuple indexing is a resource-intensive operation for the Active Directory database, which means that updating an attribute that has Tuple indexing applied takes longer than normal. For this reason, by default, no attributes are indexed for this type of search in the Windows Server 2003 schema.

**Note**

- Only attributes that are often used to perform medial-string searches should be marked as indexed for medial-string search. Unnecessary indexing, especially of this type, can cause update performance to decrease.

**Deleted objects exception to medial-string search**

Because Tuple indexes can become quite large, values of Tuple-indexed attributes on deleted objects (called tombstones) are not included in the index. Therefore, if the special LDAP control to *show deleted objects* is specified in the search operation, the search cannot use Tuple indexes to speed up the search.

**Indexing an attribute for medial-string search**

You can use ADSI Edit to mark an attribute for Tuple indexing. To use ADSI Edit, you must have Windows Support Tools installed.

The process for marking an attribute for Tuple indexing involves changing the value of the **searchFlags** attribute on the corresponding attributeSchema object that defines the attribute that you want to index.

**Efficient listing of large search results in Windows Server 2003**

New LDAP v3 controls are introduced in Windows Server 2003 that improve the ability of applications to display large lists of search results. Two new controls extend the LDAP protocol to provide a scalable solution for managing large directory deployments by improving the paging and sorting of search lists, as follows:

- Virtual list view (VLV). Allows only the viewed portion of the list to be required in memory.

- Attribute scoped query (ASQ). Applies VLV functionality to the listing of search results on large, multivalue attributes.

These two controls enable a client to browse large lists of objects without a negative impact to the client. VLV and ASQ lower network traffic and ease the client's burden of holding a large set of unneeded data in memory.

VLV and ASQ can also be used to extend other applications that retrieve information from Active Directory. For more information about programming applications that implement VLV and ASQ, see "Extended Controls" in the Microsoft Platform SDK on MSDN.

**VLV**

A virtual list is a GUI technique that is employed when ordered lists containing a large number of entries need to be displayed. When the LDAP protocol is extended to use VLV, a window that contains a small number of visible list entries is drawn. The

visible portion of the list can be relocated to different points in the list by means of scrolling, slider bars, and cursor keys as well as PAGE UP and PAGE DOWN keys. The user experience is that the full list can be browsed at will, even though it might contain millions of entries. In fact, the complete list contents are never required at any one time. Rather than retrieve the complete list from wherever it is stored (typically, from disk or a remote server), the server retrieves only the information that is required to display the part of the list that is currently in view on the client.

When you have indexed appropriate attributes to use VLV, you can use Ldp (or other VLV-enabled applications) to retrieve large lists from Active Directory without requiring the server to return the entire contents of the results set (improving response speed) and without requiring the search client to store the results set in memory (improving the speed at which results are listed).

**Configuring VLV**
For the directory service to efficiently browse a container that is sorted on a particular attribute, special indexes have to exist in the directory. These indexes are available in all the languages that are installed on the particular server.

You can define the views that the directory supports by modifying the attributeSchema object for the particular attribute on which you want to search. More specifically, you can modify the **searchFlags** attribute on the attributeSchema object to set it for VLV indexing.

To index an attribute for VLV, set the **searchFlags** value as follows:

*current value* OR 2

For example, if you want to add support for VLV container browsing using the **name** attribute, modify the attributeSchema object cn=RDN,cn=Schema,cn=Configuration,dc=*DomainName*.

The default value of **searchFlags** for this object is 13 (1101); therefore, in this case it becomes 15 (1111). If the attribute has a different value, simply perform an OR operation on that value with 2 (0010).

Active Directory does not implement any specific restrictions on the use of the VLV control. However, if a search operation cannot be processed by the server before the time-out (to build the required table that is specified by the search filter and sort key), the server returns a time-out.

**ASQ**
The ASQ LDAP control, which is new in Windows Server 2003, makes it possible to apply the query scope to a linked-value attribute instead of to a container. In this way, ASQ extends the VLV feature to apply to large, multivalue attributes with distinguished name values (the type of attributes that can be linked).

With ASQ, the VLV control can be run on an attribute with distinguished name syntax to:

- Receive a sorted list of values in a multivalue attribute that has a large number of distinguished name values, such as the **member** attribute on a group object.

- Perform a nested query on a multivalue attribute that has a large number of distinguished name values and return the values in a single operation.

For example, if an ASQ search is performed on the **member** attribute of a group object, ADSI enumerates the distinguished names in the **member** attribute and searches each of the members of the group for the specified search criteria.

**LDAP Query Policy**
LDAP v3 supports the definition of client query policy. By default, limits are placed on the server resources that are available to clients requesting LDAP queries, paged result sets, and sorted result sets. These limits constitute the LDAP query policy.

The query policy is stored as a multivalue attribute (**lDAPAdminLimits**) of the Default Query Policy object in the configuration directory partition (cn=Default Query Policy,cn=Query-Policies,cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration,dc=*Forest RootDomain*).

**Note**

- LDAP query policy is not related to, and is not a subset of, Group Policy.

Because the workload and resources of a given server can vary, the LDAP query policy is configurable at the server level.

LDAP query policy applies to the following LDAP query–related operations:

- Search. The basic query operation. An LDAP search might cover a small part of a single domain, or it might span every directory partition in the forest. A search can generate a significant amount of disk activity, take a long time, and return a large volume of data.

- Search with Paged Results. Because a search can return a large volume of data, the client can ask the server to hold the result set and return it in "pages" of a specified length. The server must hold the result set until the client releases it or unbinds.

- Search with Sorted Results. A client can request a result set in a particular order. Sorting requires storage and CPU cycles at the server. The resources consumed are directly proportional to the size of the result set.

- Search with Replication. The administrator can specify the maximum number of attribute values that can be returned per request.

- Change Notify. A client can request change notification on particular objects in the directory. The mechanism that is used to post a Change Notify request is the asynchronous LDAP query.

In the absence of any other assigned policy, all domain controllers use the default query policy. If a site policy is assigned, the domain controller uses the site policy. If a specific policy has been assigned to a domain controller, this policy takes precedence over any site policy.

**Administrative query policies: attribute values**
The following table shows the administrative limits for default query policy.

**Default Values for LDAP Administrative Limits**

| LDAP Administrative Limits | Default Value | Description and Search Behavior |
|---|---|---|
| **MaxConnIdleTime** | 900 | Maximum Connection Idle Time. The maximum time (in seconds) that the client is allowed to be idle before the connection is closed. |
| **MaxActiveQueries** | 20 | Maximum Active Queries. The maximum number of concurrent search operations allowed on the server. When the stated limit is reached, the LDAP server returns a busy notification. |
| **MaxNotificationPerConn** | 5 | Maximum Notifications per Connection. The maximum number of concurrent notification requests allowed per connection on the server. When the stated limit is reached, the server returns a busy notification. |
| **MaxPageSize** | 1000 | Maximum Page Size. The largest page size allowed by the server (in number of rows). The server returns the number of rows that are specified by **MaxPageSize**. If paged results are requested, the client can retrieve additional pages until all results are returned. |
| **MaxQueryDuration** | 120 | Maximum Query Duration. The maximum elapsed time (in seconds) that is allowed for a query to complete. If paged results are requested, the client can continue the query if the timer expires before the query completes. When the stated limit is reached, the server returns the **timeLimitExceeded** error. |
| **MaxReceiveBuffer** | 10485760 | Maximum Receive Buffer. The maximum LDAP request size (in bytes) that the server attempts to process. If the server receives a request that is larger than this value, it closes the connection. |
| **MaxTempTableSize** | 10000 | Maximum Temporary Table Size. The upper limit (in candidate objects) on the temporary table. If the temporary table maximum limit is reached by an "OR" query optimization, the optimization is abandoned and replaced with a direct table scan. |
| **MaxResultSetSize** | 262144 | Maximum Result Set Storage. The maximum storage (in kilobytes (KB)) that the server can hold for all paged result sets. If the stated limit is reached, the oldest result sets are discarded. |

| | | |
|---|---|---|
| **MaxPoolThreads** | 4 | Per Processor Asynchronous Thread Queue (ATQ) Threads. The number of threads that are allocated by ATQ per processor. This value is sent as an advisory notification to ATQ. ATQ decides whether to use it or not.<br><br>**Note** If it takes a long time to bind, increase the count to 6 or 8. |
| **MaxDatagramRecv** | 1024 | Maximum Receive Datagram Size. The maximum size of datagrams (in bytes) that can be received by the server. The server preallocates datagram buffers and cannot receive datagrams with a size that is larger than the stated limit. |
| **InitRecvTimeout** | 120 | The maximum time (in seconds) that the server waits for the initial request before the connection is dropped. |
| **MaxConnections** | 5000 | The maximum number of concurrent LDAP connections allowed on the server. User Datagram Protocol (UDP) connections do not count toward this limit. If the limit is reached, the LDAP server sends back an LDAP disconnect notification and closes down the connection. |
| **MaxValRange** | 1500 | Maximum Value Range. Controls the threshold at which the server will start returning the range option for attributes with a large number of values. The minimum value for this policy is 30. |

**Query Limits**

To improve the query response time for searches for Active Directory objects, searches are limited to 1,000 objects by default. However, you may want to increase this limit as your organization grows. You can control the buffer size that is allocated for storing the number of objects that are returned by a query search. To control the buffer size, you can either modify the registry on the search client or use Group Policy to set the buffer size on all computers in a domain, site, or OU.

**Modifying the registry to change the maximum query limit**

You can increase the number of objects that are returned by an Active Directory search on an individual search client by setting the limit in the registry. The **REG_DWORD** value named **QueryLimit** on the registry key **HKEY_CURRENT_USER\Software\Policies\Microsoft\Windows\Directory** controls the maximum query limit on the local computer.

**Using Group Policy to change the maximum query limit**

You can use Group Policy to increase the maximum query limit on objects that are returned in response to a command to browse or search Active Directory when the command is issued from an operating system in the Windows 2000 Server family or the Windows Server 2003 family. In the User Configuration section of Group Policy, the setting is found in \Administrative Templates\Desktop\Active Directory\Maximum size of Active Directory searches. You can edit this setting with Group Policy Management Console (GPMC).

### Note

- Because this query limit policy is a user-based policy, it is not enforced until the user has logged off the system and then logged on again. Also, the policy only applies to default Active Directory tools, such as Active Directory Users and Computers. Any query limits that are set through Group Policy do not apply to LDAP queries that are performed programmatically.

**Paging**

When an LDAP client is accessing a server across a slow connection or if the result set from a given query may be very large, the client can retrieve a result set in small pieces. The Simple Paged Result extended control provides for this type of retrieval. Options for this control enable the client to set the initial page size and reset the page size with each subsequent request to the server.

The Simple Paged Result control is also used to access all of a large results set when there is a server-side administrative limit to the number of items that are returned from a query. For example, domain controllers running Active Directory have a default server-side limit of 1,000 entries as the maximum number of results that are returned in a single request. If the results of a query exceed this limit, the Paged Results control is used, with a page size equal to or less than the server-side limit, to retrieve all of the results of the query.

For more information about paging, see “Paging Search Results” in the Microsoft Platform SDK on MSDN.

**Multivalue Attributes**

Attributes can be single value or multivalue. Single-value and multivalue attributes are defined by the **singleValued** attribute

being set to **TRUE** or **FALSE**. A multivalue attribute can contain multiple values, all of uniform syntax. Note that multivalue attributes hold a set of values with no particular order. There is no guarantee that multivalue properties are ever going to be returned in the order in which they were stored (or in any other order).

Retrieving the contents of a multivalue attribute, such as a distribution list, can often result in a large number of returned values. LDAP servers often place limits on the maximum number of attribute values that can be retrieved in a single query. If an attribute has more members than can be returned by the server in a single call, the only way to enumerate all of the attribute values is through the use of the range option.

Range retrieval involves requesting a limited number of attribute values in a single query. The number of values that are requested must be less than, or equal to, the maximum number of values that are supported by the server. To reduce the number of times that the query must contact the server, the number of requested values should be as close to this maximum as possible.

Active Directory servers set a limit on the maximum number of attribute values that are returned. The version of the operating system for the server that supplies the requested data determines the maximum number of values that can be retrieved in a single query. The following table lists the operating system version and the maximum number of values that can be retrieved in a single query.

**Maximum Query Size for Multivalue Attribute Searches**

| Operating system version | Maximum values returned |
|---|---|
| Domain controllers running Windows 2000 Server | 1,000 |
| Domain controllers running Windows Server 2003 | 1,500 |

**Note**

- In the Windows Server 2003 family, a server running Windows Server 2003, Web Edition cannot be a domain controller.

#### Attribute range option

The range option can be specified as part of an LDAP search to retrieve the values of a multivalue attribute incrementally. An LDAP search on an attribute includes an attribute type (for example, **member**) and a list of options, one of which can be the range option. When it is presented in a **searchRequest** message, the range option specifies a zero-relative range of elements (for example, 0 through 9) to be retrieved. When the range option is specified, followed by a range specifier, only the number of values in that range is retrieved. For more information about the range option, see "Searching Using Range Retrieval" in the Microsoft Platform SDK on MSDN.

**Note**

- You can use the range option in multivalue attribute searches by using Ldp.exe.

#### Linked-value attributes

Linked-value attributes are a new feature in Windows Server 2003. Linked-value attributes are pairs of attributes in which the system calculates the values of one attribute (the back link) based on the values that are set on the other attribute (the forward link) throughout the forest. A back-link value on any object instance consists of the distinguished names of all the object's that have the object's distinguished name set in the corresponding forward link. A back link is always a multivalue attribute. A forward link can be a multivalue attribute or a single-value attribute.

For example, **Manager** and **Reports** are a pair of linked-value attributes, where **Manager** is the forward link and **Reports** is the back link. For example, suppose that Bill is Joe's manager. If you store the distinguished name of Bill's user object in the **Manager** attribute of Joe's user object, the distinguished name of Joe's user object shows up in the **Reports** attribute of Bill's user object.

A forward-link/back-link pair is identified by the **linkID** values of two **attributeSchema** definitions. The **linkID** of the forward link is an even, positive, nonzero value, and the **linkID** of the associated back link is the forward **linkID** plus one. For example, the **linkID** for **Manager** is 42, and the **linkID** for **Reports** is 43.

The following is a list of guidelines for defining a new pair of linked-value attributes:

- The **linkID** values must be unique among all attributeSchema objects. To avoid conflicts, you can obtain values for **linkID** attributes from the Link ID Generator on MSDN.

- A back link must have a corresponding forward link; that is, the forward link must exist before a corresponding back-link attribute can be created.

- The **attributeSchema** value of a forward link must be **2.5.5.1**, **2.5.5.7**, or **2.5.5.14**. These values correspond to syntaxes that contain a distinguished name, such as the Object(DS-DN) syntax.

- The **attributeSchema** value of a back link must be **2.5.5.1**, which is the Object(DS-DN) syntax.

- By convention, back-link attributes are added to the **mayContain** value of the top abstract class. This enables the back-link attribute to be read from objects of any class, because they are not actually stored with the object but calculated based on the forward-link values.

For information about the replication of linked-value attributes, see "How the Active Directory Replication Model Works."

**Global Catalog Searches**
The global catalog enables searches for Active Directory objects in any domain in the forest, without the need for subordinate referrals. Users can find an object of interest quickly without having to know which domain holds the object.

**Global catalog servers**
A global catalog server is a domain controller that stores extra information. The database on the global catalog server stores information about every object in the forest, instead of information for the objects in only one domain. The information about objects that occur in directory partitions for domains other than the local domain includes only a subset of attributes for each object. In this way, the global catalog enables forest-wide searches without requiring replication of the entire contents of Active Directory to every domain controller. The Knowledge Consistency Checker (KCC) process creates a replication topology that ensures delivery of the partial contents of every directory partition to every global catalog server in the forest.

### Note

- A global catalog server stores full copies of the schema and configuration directory partitions, which is the same as for any domain controller.

By default, the server on which you install Active Directory to create the first domain in a new forest is a global catalog server. Thereafter, you must designate additional global catalog servers if you need them.

**Searching the domain vs. searching the global catalog**
The decision whether to search the domain or the global catalog is based on the scope of the search:

- When the scope of a search is the domain or an OU, the query can be resolved in the domain directory partition by using an LDAP search.

- When the scope of a search is the forest and the attributes that are being searched against are part of the global catalog, the query can be resolved in any directory partition by using a global catalog search.

**Searches that use the global catalog by default**
To search the global catalog, instead of the local domain partition, you must specify port 3268 in the search tool that you are using, instead of port 389, the standard LDAP port. Anytime that you specify port 3268, you are searching in the global catalog. In addition, the global catalog is searched by default under the following conditions:

- During the logon process, when a user principal name (UPN) is presented. The global catalog is searched to find the domain and account name on the basis of the UPN.

- During the logon process, to expand universal groups. Universal group membership can span domains. It is possible, therefore, that a user has a membership in a universal group that is not in the logon domain. For this reason, the global catalog is contacted to search the membership of universal groups. If a membership is found, the group is attached to the user's logon credentials.

- When you select **Entire Directory** in a search-scope list in an Active Directory tool, such as Active Directory Users and Computers.

- When you write the distinguished name value for a property, where the distinguished name represents a nonlocal object. For example, if the member that you are adding is from a different domain, the global catalog is used to verify that the user object that is represented by the distinguished name actually exists.

**Global catalog search base**
For an LDAP search, you must supply a valid search base. For a global catalog search, the search base can be any value, including the value NULL ( ). A search base of NULL effectively scopes the search on the search computer to the global catalog. If you use a NULL search base with a scope of one level or subtree and specify port 389 (the default LDAP port), the search fails. Therefore, if you submit a NULL search to the global catalog port and then change the port to the LDAP port, you must change the search base for the search to succeed.

**Characteristics of a global catalog search**
The following additional characteristics differentiate a global catalog search from a standard LDAP search:

- A global catalog search crosses directory partition boundaries. The extent of an LDAP search is the directory partition.

- A global catalog search does not return subordinate referrals. If you use port 3268 to request an attribute that is not in the global catalog, you do not receive a referral to it. Subordinate referrals are an LDAP response. When you query a server over port 3268, you receive global catalog responses, which are based solely on the contents of the global catalog. If you query the same server over port 389, you receive referrals for objects that are in the forest but whose attributes are not referenced in the global catalog.

### Note

- An external referral can be returned by the global catalog if a base-level search for an external directory is submitted and if the distinguished name of the external directory uses the domain component (dc=) naming attribute. This referral is returned according to the ability of Active Directory to construct a DNS name from the domain components of the distinguished name, and it is not based on the presence of any cross-reference object. The same referral is returned by using the LDAP port; it is not specific to the global catalog. For more information about constructing a DNS name from the domain components, see “Superior References” earlier in this section.

When a forest consists of a single domain, there is no need for a global catalog because every domain controller has a full, writable copy of every object in the domain and forest. However, it is important to retain the global catalog on at least one domain controller because many applications use port 3268 for searching. For example, if you do not have any global catalog servers, the **Search** command on the **Start** menu of Windows 2000 Professional, Windows 2000 Server, Windows XP Professional, and Windows Server 2003 cannot locate objects in Active Directory.

**Searching on back links**
Some Active Directory attributes cannot be located specifically in the directory database. A back link is an attribute that can be computed only by referencing another attribute, called a forward link. An example of a back-link attribute is the **memberOf** attribute on a user object, which relies on the group attribute **member** to derive its values. For example, if you request the groups of which a specific user is a member, the forward link **member** — an attribute of the group object — is searched to find values that match the user name that you specify.

When you request the value of a back link on a particular object (for example, “What objects are managed by UserB?”), the system searches for all objects whose corresponding forward link names the original object (that is, " What objects have the distinguished name of UserB as the value in their **managedBy** attribute?"). The results of that search, and therefore the apparent contents of the back-link attribute, depend on the LDAP port to which the client is bound. That is, the results can differ, depending on whether the client binds to the local domain (LDAP port 389) or to the global catalog (LDAP port 3268).

**Effect of the global catalog on searches of back links and forward links**
Because of the way that groups are enumerated by a global catalog server, the results of a back-link search can vary, depending on whether you search the global catalog (port 3268) or the domain (port 389), the kind of groups that the user belongs to (global groups or domain local groups), and whether the user belongs to universal groups outside the local domain.

### Note

- Universal groups are available only in domains where domain controllers are operating in Windows 2000 native mode (in a Windows 2000 domain), in domains that have a Windows 2000 native domain functional level (in a

Windows Server 2003 domain), or in forests that have a Windows Server 2003 forest functional level.

Domain controllers store a full replica of the domain global groups, domain local groups, and any universal groups that are created in that domain. Therefore, if a query occurs over port 389, the global group and domain local group memberships for users in the local domain can be found, as well as memberships in any universal groups that are stored in that domain.

Global catalog servers store a partial replica of every group in the forest. The **member** attribute of a group is replicated to the global catalog for universal group objects but not for global groups or domain local groups. Therefore, if a query occurs over port 3268 and the global catalog server that is contacted is authoritative for the user's domain, that user's domain local memberships and global group memberships can be enumerated because full replicas of those objects exist on the server. Otherwise, only the user's universal group memberships are available.

Connecting to the global catalog always locates the user's membership in universal groups. However, connecting to the global catalog does not locate the users membership in global groups or domain local groups, unless the global catalog server happens to be an authoritative domain controller for the user's domain.

**Network Ports Used by Active Directory Searches**
The network ports that are used by Active Directory searches are listed in the following table.

**Port Assignments for Active Directory Searches**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | None | 389 |
| LDAP SSL | None | 636 |
| Global Catalog LDAP | None | 3268 |
| Global Catalog LDAP SSL | None | 3269 |

**Related Information**
The following resources contain additional information that is relevant to this section:

- Certificate Services Technical Reference
- Active Directory Schema Technical Reference
- Data Store Technical Reference
- DNS Support for Active Directory Technical Reference
- Logon and Authentication Technologies
- RFC 1777 in the IETF RFC Database
- RFC 3377 in the IETF RFC Database
- Microsoft Platform SDK on MSDN for more information about concurrent binding (in "Using Concurrent Binding")
- Microsoft Platform SDK on MSDN for more information about LDAP search filter syntax (in "Search Filter Syntax")
- Active Directory Replication Model Technical Reference
- Microsoft Platform SDK on MSDN for more information about programming applications that implement VLV and ASQ (in "Extended Controls")
- Microsoft Platform SDK on MSDN for more information about paging (in "Paging Search Results")

- Microsoft Platform SDK on MSDN for more information about the range option (in “Searching Using Range Retrieval”)
- Link ID Generator on MSDN for values for **linkID** attributes

# Active Directory Searches Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Active Directory Searches Tools and Settings**

**In this section**


- Active Directory Searches Tools
- Active Directory Searches Group Policy Settings
- Active Directory Searches WMI Classes
- Network Ports Used by Active Directory Searches
- Related Information


This section contains information about the tools, Group Policy settings, Windows Management Instrumentation (WMI) classes, and network ports that are associated with Active Directory searches.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**Active Directory Searches Tools**

The following tools are associated with Active Directory searches.

**Adsiedit.msc: ADSI Edit**

**Category**

This tool ships with Support Tools for Windows Server 2003.

**Version Compatibility**

| Can Be Run From | Can Be Run Against |
| --- | --- |

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br>• Windows Server 2008<br>• Windows Server 2008 R2<br><br>Computers running:<br>• Windows XP Professional<br>• Windows Vista<br>• Windows 7 | Domain controllers running:<br>• Windows Server 2008 R2<br>Windows Server 2008<br>Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server<br>• |

ADSI Edit is a Microsoft Management Console (MMC) tool that you can use to view and modify directory objects.

To find more information about ADSI Edit, see "Support Tools Help" in Tools and Settings Collection.

**Dsquery.exe: Dsquery**

**Category**
This tool ships with Windows Server 2003.

**Version Compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

You can use Dsquery to perform searches against Active Directory according to specified criteria.

To find more information about Dsquery, see "Command-Line References" in Tools and Settings Collection.

**Ldp.exe: Ldp**

**Category**
This tool ships with Support Tools for Windows Server 2003.

**Version Compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Ldp is a Lightweight Directory Access Protocol (LDAP) graphical user interface (GUI) tool that you can use to perform

operations, such as connect, bind, search, modify, add, and delete, against any LDAP-compatible directory, such as Active Directory.

To find more information about Ldp, see "Support Tools Help" in Tools and Settings Collection.

**Ntdsutil.exe: Ntdsutil**

**Category**
This tool ships with Windows Server 2003.

**Version Compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition |

You can use Ntdsutil to perform Active Directory database maintenance, manage and control single master operations, and remove metadata left behind by domain controllers that are removed from the network without being properly uninstalled. This tool is intended for use by experienced administrators.

To find more information about Ntdsutil, see "Command-Line References" in Tools and Settings Collection.

**Active Directory Searches Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with Active Directory searches.

**Group Policy Settings Associated with Active Directory Searches**

| Group Policy Setting | Description |
|---|---|
| **Maximum size of Active Directory searches** | Specifies the maximum number of objects that the system displays in response to a command to browse or search Active Directory.<br><br>This policy affects all browse displays that are associated with Active Directory, such as those displays in Local Users and Groups, Active Directory Users and Computers, and dialog boxes that are used to set permissions for user or group objects in Active Directory.<br><br>If you enable this policy, you can use it to limit the number of objects that are returned from an Active Directory search.<br><br>If you disable this policy or if you do not configure it, the system displays up to 10,000 objects. |
| **Enable filter in Find dialog box** | Displays the filter bar above the results of an Active Directory search. The filter bar consists of buttons for applying additional filters to search results.<br><br>If you enable this policy, the filter bar appears when the Active Directory **Find** dialog box opens, but users can hide it. |

| | |
|---|---|
| **Hide Active Directory folder** | Hides the Active Directory folder in My Network Places.<br><br>The Active Directory folder displays Active Directory objects in a browse window.<br><br>If you enable this policy, the Active Directory folder does not appear in the My Network Places folder.<br><br>If you disable this policy or if you do not configure it, the Active Directory folder appears in the My Network Places folder. |

To find more information about Group Policy settings, see the "Group Policy Settings Reference" in Tools and Settings Collection.

**Active Directory Searches WMI Classes**

The following table lists and describes the WMI classes that are associated with Active Directory searches.

**WMI Classes Associated with Active Directory Searches**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| **rootDSE** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |
| **DS_LDAP_Class_Containment** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

| | | |
|---|---|---|
| **DS_LDAP_Instance_Containment** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

For more information about these WMI classes, search for "Mapping Active Directory to WMI" in the WMI SDK documentation on MSDN.

**Network Ports Used by Active Directory Searches**

The network ports that are used by Active Directory searches are listed in the following table.

**Port Assignments for Active Directory Searches**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | None | 389 |
| LDAP SSL | None | 636 |
| Global Catalog LDAP | None | 3268 |
| Global Catalog LDAP SSL | None | 3269 |

**Related Information**

The following resources contain additional information that is relevant to this section:

- "Support Tools Help" in Tools and Settings Collection for information about ADSI Edit and Ldp
- "Command-Line References" in Tools and Settings Collection for information about Dsquery and Ntdsutil
- "Group Policy Settings Reference" in Tools and Settings Collection for more information about Group Policy settings for Active Directory searches
- Microsoft Platform SDK on MSDN for more information about WMI classes for Active Directory searches (in "Mapping Active Directory to WMI")

# Service Publication and Service Principal Names Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

Services use service publication in the Active Directory directory service to provide information about themselves in the directory for easy discovery by client applications and other services. Service publication occurs when the installation program for a service publishes information about the service, including binding and keyword information, to the directory. In addition, Active Directory supports service principal names (SPNs) as a means by which client applications can identify and authenticate the services that they use. Service publication is accomplished through the creation of service objects (also called connection point objects) in Active Directory. Service authentication is accomplished through Kerberos authentication of SPNs.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**In this subject**


- What Are Service Publication and Service Principal Names?
- How Service Publication and Service Principal Names Work
- Service Publication and Service Principal Names Tools and Settings

# What Are Service Publication and Service Principal Names?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- The Business Need
- The Service Publication and SPN Solution
- Service Publication and SPN Scenarios
- Related Information


Services use service publication in Active Directory to provide information about themselves in the directory for easy discovery by client applications and other services. Service publication occurs when the installation program for a service publishes information about the service, including binding and keyword information, to the directory. In addition, Active Directory supports service principal names (SPNs) as a means by which client applications can identify and authenticate the services that they use. Service publication is accomplished through the creation of service objects (also called connection point objects) in Active Directory. Service authentication is accomplished through Kerberos authentication of SPNs.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**The Business Need**

As organizations come to rely more heavily on distributed services for getting work accomplished, service consumers (that is, users and client applications) in those organizations need a way to search for and discover services on the network quickly and easily. At the same time, administrators need to protect service consumers from insecure or unidentified services and to keep track of valid services that are running on their networks.

**The Service Publication and SPN Solution**

Service publication and SPNs are two separate, independently deployable Active Directory technologies that together provide a complete solution for the service discovery and authentication needs in an organization.

Service publication in Active Directory enables users and applications to find services on the network — without having to know on which server the service resides — through the use of connection point objects in the directory. A connection point object is created by a service at the time that the service is installed. The service uses this object to publish information about itself, including binding information and searchable keywords. Clients of the service can search against the data that is stored in connection point objects in the directory to find out about services on the network.

Service consumers can use SPNs in Active Directory to authenticate the identity of services before using the services. An SPN is a unique attribute that resides on the account object in whose security context a given service runs. SPN structures generally follow Internet Engineering Task Force (IETF) naming conventions, and they often include the name of the computer on which the service is running. SPNs may be used to request Kerberos tickets, and they are required for mutual authentication between clients and services. For more information about SPNs, see "How Service Publication and Service Principal Names Work."

Before connecting to a service, a service consumer can construct an SPN string that matches the SPN of the service by using information that is gathered from a connection point object or some other source. The service consumer can then present this SPN string to the Key Distribution Center (KDC) and request that the KDC authenticate the service that is represented by the SPN.

The following figure illustrates the relationship between services, client applications, and Active Directory.

**Service Publication and SPNs in Active Directory**




**Service Publication and SPN Scenarios**
The following sections describe the scenarios to which service publication and SPNs most commonly apply.

**Standardized Service Publication**
Service publication through connection point objects is most commonly used by organizations to establish a single, standard mechanism for publishing service information throughout the organization. This scenario takes advantage of Active Directory as the single, pervasive directory service that is available to all clients throughout the network. This scenario involves services that can write connection point objects to the directory. As an alternative, administrators can create and maintain connection point objects manually.

**Protecting Clients from Unauthorized Services**
In this scenario, organizations that are concerned about protecting clients from unauthorized services use SPNs as a way to enable client authentication of services. This scenario involves client applications that can request authentication of SPNs.

**Related Information**

- How Service Publication and Service Principal Names Work
- "Microsoft Platform SDK" on MSDN for more information about service publication (in "Service Publication")
- "Microsoft Platform SDK" on MSDN for more information about service authentication (in "Mutual Authentication Using Kerberos")

# How Service Publication and Service Principal Names Work

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Service Publication and SPN Architecture
- Service Publication and SPN Protocols
- Service Publication and SPN Physical Structure
- Service Publication and SPN Processes and Interactions
- Related Information


Services use service publication in Active Directory to provide information about themselves in the directory for easy discovery by client applications and other services. In addition, Active Directory supports service principal names (SPNs) as a means by which client applications can identify and authenticate the services that the client applications use. Service publication is accomplished through the creation of service objects (also called connection point objects) in Active Directory. Service authentication is accomplished through Kerberos authentication of SPNs.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

Service publication and SPNs work through special objects and attributes that are stored in Active Directory. These objects and attributes are published to the directory by the respective services that the objects and attributes represent. Service publication and SPNs are optional capabilities of Active Directory, and they are not required for a properly functioning Active Directory infrastructure. However, a properly functioning Active Directory infrastructure is required for service publication and SPNs to work.

In addition to a properly functioning Active Directory infrastructure, an ideal environment for service publication and SPNs requires that the following components be present and functioning properly:

- Kerberos and Key Distribution Center (KDC)
- The services that use Active Directory to publish information about themselves
- The client applications that authenticate and use the published services
- Network connectivity, including Internet Protocol (IP), among all components

This section describes the elements of Active Directory service publication and SPNs, including their architecture, protocols, physical structure, and processes and interactions.

**Service Publication and SPN Architecture**

The primary components of the service publication and SPN architecture are the services that publish information about themselves and the client applications that search the directory to discover the services. In addition, Active Directory plays a key role in the architecture, both as the storage and distribution mechanism for published service information and as a source for the SPN attributes that are used in mutual authentication. Finally, the KDC provides the mechanism for authenticating services by means of the SPNs that are constructed by the client applications. The relationships of these components are shown in the following figure.

**Service Publication and SPN Architecture**




The following table describes the service publication and SPN components.

**Service Publication and SPN Components**

| Component | Description |
|---|---|
| Service | An application that makes data or operations available to network clients. |
| Client application | An application running on a workstation or on a server that makes use of a service. |
| KDC | A service running on every domain controller that provides authentication services for clients, as well as for servers and services. A client application requests the authentication of an SPN through a KDC. |
| Active Directory | The directory that is used to store connection point objects and service account objects. |
| Connection point object | An object in Active Directory that contains information about a service. Client applications search for these objects to find out about services that are available on the network. Typically, connection point objects are created on the computer object of the computer on which a given service is running. |
| SPN attribute | An attribute that contains a unique name that identifies an instance of a service. An SPN attribute resides on the service account object in whose security context a given service instance runs. |

**Service Publication and SPN Protocols**

Service publication and SPNs rely on the IP network protocol as well as on transport protocols, including Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). Client applications perform service discovery against Active Directory by using the standard operations of the Lightweight Directory Access Protocol (LDAP). Client applications request mutual authentication of services by using the Kerberos V5 authentication protocol. The protocols that are used in service publication and authentication are shown in the following figure.

**Service Publication and SPN Protocols**




The following table describes the service publication and SPN protocols.

**Service Publication and SPN Protocols**

| Protocol | Description |
|---|---|
| LDAP | The core protocol that is used to communicate with Active Directory. LDAP is described in RFC 2551 (LDAP v3) and RFC 1777 (LDAP v2). |
| Kerberos | The authentication protocol that is used by clients and services to communicate with the KDC. |
| TCP | The connection-oriented session protocol on which LDAP runs. |
| UDP | The connectionless session protocol on which Kerberos runs. |
| IP | The standard network protocol that is used by Windows Server 2003. |

**Service Publication and SPN Physical Structure**

In Active Directory, the physical structure for service publication consists of connection point objects and the containers in which those objects reside. Typically, connection point objects are created on the computer object of the computer on which a given service is running.

Microsoft SQL Server 2000 is an example of an application that supports the publication of information about itself in Active Directory by using connection point objects. SQL Server 2000 supports the registration of instances of the SQL Server relational database engine, databases, replication publications, and analysis servers in Active Directory.

For example, SQL Server 2000 can be installed on a computer for which the distinguished name in Active Directory is the following:

CN=dbserver,OU=ComputersOU,DC=corp,DC=cohovineyard,DC=com

After installing SQL Server 2000, an administrator can register an instance of the SQL Server relational database engine in Active Directory by using the tools that are provided by SQL Server 2000. After the database engine is registered in Active Directory, an **MS-SQL-SQLServer** connection point object is added to the container. This new connection point object has the following distinguished name:

CN=MS-SQL-SQLServer,CN=dbserver,OU=ComputersOU,DC=corp,DC=cohovineyard,DC=com

For more information about SQL Server 2000 service publication, see “Active Directory Integration” in the “Microsoft Platform SDK” on MSDN.

**Note**

- You can search Active Directory for existing connection point objects, using the **dsquery** command-line tool, by typing the following at a command prompt:
- **dsquery * forestroot -filter (objectclass=serviceconnectionpoint)**
- For more information about **dsquery**, at a command-line prompt, type:
- **dsquery /?**

The physical structure for client authentication of services consists of the service account objects of the services to be authenticated, the SPN attributes that reside on those objects, and the containers in which those objects reside. The service account under which a service runs determines the security identity of the service at run time, that is, it determines the service’s primary security context. The security context determines the service’s ability to access local and network resources. Generally, the service installer specifies the service account when it calls the **CreateService** function to install the service on a host computer. In addition, if the service is to support mutual authentication by means of SPNs, the service installer registers SPNs on the service account object as an attribute. If the service runs under a domain user account, the SPN attribute is created on the domain user account under which the service runs. If the service runs under the Local System account, the service’s logon account is the computer account of the host computer.

**Connection Point Objects**

A connection point object is used in Active Directory to represent a service and to store information about that service. Connection point objects are derived from the **connectionPoint** schema class, which is the base class from which all objects representing connectable resources are derived. The following figure shows some of the object classes that are derived from the **connectionPoint** object class.

**connectionPoint Class and Derived Classes**




The following items are examples of objects that are derived from the **connectionPoint** class and that are used by various services to publish their availability on the network.

**printQueue and volume**

The **printQueue** and **volume** objects are the standard connection points that are used by print services and file services, respectively. These objects are used to help directory clients find available printers and shared volumes on the network.

**rpcEntry**

The **rpcEntry** object class is used by services that publish information about themselves by using the RPC Name Service (RpcNs) application programming interfaces (APIs). Typically, objects of this class can be found in a container in the domain directory partition whose distinguished name is similar to the following:

CN=RpcServices,CN=System,DC=cohovineyard,DC=com

**serviceInstance**

The **serviceInstance** object class is used by Windows Sockets (Winsock) services that publish information about themselves by using registration and resolution (RnR).

Typically, objects of this class can be found in a container in the domain directory partition whose distinguished name is similar to the following:

CN=WinsockServices,CN=System,DC=cohovineyard,DC=com

**serviceConnectionPoint**

The **serviceConnectionPoint** object class is most commonly used by application-oriented (rather than infrastructure) services that publish information about themselves in Active Directory. Examples of services that publish information about themselves by using service connection points (SCPs) include SQL Server 2000, Remote Installation Services (RIS), and Active Directory Application Mode (ADAM).

Active Directory does not define how the attributes of the **serviceConnectionPoin**t class are to be used; therefore, clients of a service that is published as an SCP must be able to interpret and use the data that is published in the SCP. Services that must publish additional data about themselves can extend the Active Directory schema by creating a subclass of the **serviceConnectionPoint** class and giving the subclass a distinct name.

The most important attributes of the **serviceConnectionPoint** class are **keywords**, **serviceDNSName**, **serviceDNSNameType**, **serviceClassName**, and **serviceBindingInformation**. Client applications search the directory for **keywords** values to locate a connection point object. When a connection point object is found, clients can read other attributes to retrieve service data. The most important attributes of the **serviceConnectionPoint** class are described in the following table.

**serviceConnectionPoint Attributes**

| Attribute | Description |
|---|---|

| | |
|---|---|
| **keywords** | The **keywords** attribute can contain multiple string values that identify a given service. This attribute is included in the global catalog, which means that clients in any domain of a forest can search the global catalog for keywords that are associated with the service. This attribute is also indexed, which improves query performance. The installer that creates the SCP object sets the values of the keywords attribute. Typically, these values are not modified by the active service.<br><br>The exact keywords that should be included in your SCP depend on how clients search for the service. One type of keyword that is important to include is the GUID string, because globally unique identifiers (GUIDs) are guaranteed to be unique in a forest. Use the GUID string format that is returned by the **GuidToString** function in the remote procedure call (RPC) library. Also important to include are human-readable names for clients to search for a given service. The keywords in an SCP object should include GUID strings and names that identify the following data about the service:<br><br>• The company or organization, for example, Coho Vineyard.<br><br>• The product or service, for example, SQL Server. This enables client applications to find SCPs for services of that type.<br><br>• The specific version of the product or service, for example, 7.5.<br><br>For SCP objects that publish a specific set of data or capabilities for a type of service, a GUID string or name should be included that identifies the specific instance. For example, a database service might publish an SCP object for a specific database. In this case, the SCP object should include a product GUID to identify the service and another GUID to identify the database. |
| **serviceDNSName** and **serviceDNSNameType** | Client applications use the **serviceDNSName** and **serviceDNSNameType** attributes to determine the service's host computer. These attributes are not indexed, and they are not replicated as part of the global catalog.<br><br>The **serviceDNSNameType** value indicates the type of DNS name that is specified by **serviceDNSName —** usually, "A" if **serviceDNSName** contains a host name or "SRV" if **serviceDNSName** contains an SRV record name.<br><br>The **serviceDNSName** value is typically the DNS name of the service's host computer. Your service installer can call the **GetComputerNameEx** function to get the DNS name of the local computer.<br><br>For services that have DNS SRV records, **serviceDNSName** can be the name of the SRV record. A client application uses the DNS APIs to retrieve all the SRV records that match this name. The client then retrieves the DNS host name from one of the SRV records. This technique is useful for replicated services, because SRV records also include data that enables the client to select the best replica. |
| **serviceClassName** | This is a single-value property that identifies the class of service that is represented by the SCP. This attribute is not indexed, and it is not replicated as part of the global catalog.<br><br>This is a descriptive string that is specific to the service that published the SCP; for example SqlServer. For services that support mutual authentication, clients can use this property, along with the DNS name of the service's host computer, to form an SPN. |

| | |
|---|---|
| **serviceBindingInformation** | This is a multivalue property that contains string values that store data that is required to bind to a service. This property is indexed, and it is replicated to the global catalog.<br><br>The content of **serviceBindingInformation** is specific to the service that publishes the SCP; clients must interpret the binding data. In the most common case, the binding data consists of a port number on the service host computer.<br><br>Examples of service binding information include the following:<br><br>• File Service: A Universal Naming Convention (UNC) name for a share.<br><br>Example: \\MyServer\MyshareName<br><br>• Web Service: A Uniform Resource Locator (URL).<br><br>Example: http://www.cohovineyard.com<br><br>• RPC Service: RPC binding, encoded information that is used to connect to the RPC server. RPC bindings can be converted to and from strings with the RPC APIs.<br><br>Example: ncacn_ip_tcp:server.cohovineyard.com |

**SPNs**

A service principal name (SPN) is the name by which a client uniquely identifies an instance of a service. SPN structures generally follow Internet Engineering Task Force (IETF) naming conventions, and they often include the name of the computer on which the service is running. SPNs may be used to request Kerberos tickets, and they are required for mutual authentication. An example of an SPN is MyDBService/host1.cohovineyard.com, where MyDBService is the name of the service and host1.cohovineyard.com is the name of the host on which the service resides. If multiple instances of a service are installed on computers throughout a forest, each instance must have its own SPN. A given service instance can have multiple SPNs if there are multiple names that clients might use for authentication. For example, an SPN always includes the name of the host computer on which the service instance is running; therefore, a service instance might register an SPN for each name or alias of its host.

Before Kerberos can use an SPN to authenticate a service, the SPN must be registered on the account object that the instance of the service uses to log on. A given SPN can be registered on only one account. For Win32 services, a service installer specifies the logon account when an instance of the service is installed. The installer then composes the SPNs and writes them as a property of the account object in Active Directory. If the logon account of a service instance changes, the SPNs must be manually registered again under the new account.

When a client attempts to connect to a service, it locates an instance of the service, composes an SPN for that service instance, connects to the service, and presents the SPN for the service to authenticate.

**Name Formats for Unique SPNs**

An SPN must be unique in the forest in which it is registered. If it is not unique, authentication will fail. The SPN syntax has four elements: two required elements and two additional elements that can be used, if necessary, to produce a unique name, as follows:

*ServiceClass/Host:Port/ServiceName*

Each of these SPN elements is described in the following table.

**SPN Elements**

| Element | Description |
|---|---|
| *ServiceClass* | A string that identifies the general class of service, for example, SqlServer. There are well-known service class names, such as www for a Web service or ldap for a directory service. In general, this can be any string that is unique to the service class. Be aware that the SPN syntax uses a forward slash (/) to separate elements; therefore, the forward slash cannot appear in a service class name. |

| | |
|---|---|
| *Host* | The name of the computer on which the service is running. This can be a fully-qualified DNS name or a network basic input/output system (NetBIOS) name. Be aware that NetBIOS names are not guaranteed to be unique in a forest, and Active Directory does not enforce uniqueness for the host component of an SPN. |
| *Port* | An optional port number to differentiate between multiple instances of the same service class on a single host computer. This component can be omitted if the service uses the default port number that is reserved for its service class (port number 389 for LDAP, for example). |
| *ServiceName* | An optional name that is used in the SPN of a replicable service to identify the data or services that are provided by the service or to identify the domain that is served by the service. This component can have one of the following formats (each of which are guaranteed to be unique within Active Directory):<br>• The distinguished name or objectGUID of an object in Active Directory, such as a connection point object.<br>• The DNS name of the domain for a service that provides a specified service for a domain as a whole.<br>• The DNS name of an SRV or mail exchange (MX) record. |

The components that are present in a service's SPNs depend on how the service is identified and replicated. There two types of services: host-based services and replicable services.

**Host-based services**
For a host-based service, the *ServiceName* component is omitted because the service is uniquely identified by the service class and by the name of the host computer on which the service is installed, as follows:

*ServiceClass/Host*

The service class alone is sufficient to identify for clients the features that the service provides. You can install instances of the service class on many computers, and each instance provides services that are identified with its host computer. File Transfer Protocol (FTP) and Telnet are examples of host-based services. The SPNs of a host-based service instance can include the port number if the service uses a nondefault port or if there are multiple instances of the service on the host, as follows:

*ServiceClass/Host:Port*

**Replicable services**
For a replicable service, there can be one or many instances of the service (replicas), and clients do not differentiate which replica they connect to because each replica provides the same service. The SPNs for each replica have the same *ServiceClass* and *ServiceName* components, where *ServiceName* identifies more specifically the features that are provided by the service. Only the *Host* and optional *Port* components vary from SPN to SPN.

An example of a replicable service is an instance of a database service that provides access to a specified database. In this case, *ServiceClass* identifies the database application, and *ServiceName* identifies the specific database. *ServiceName* might be the distinguished name of an SCP object that contains connection data for the database. As shown in the following examples, only the value for *host* varies among the service instances:

MyDBService/*host1.cohovineyard.com*/CN=hrdb,OU=mktg,DC=*cohovineyard*,DC=com

MyDBService/*host2.cohovineyard.com*/CN=hrdb,OU=mktg,DC=*cohovineyard*,DC=com

MyDBService/*host3.cohovineyard.com*/CN=hrdb,OU=mktg,DC=*cohovineyard*,DC=com

If clients use the NetBIOS name to compose a service's SPN, each replica must also register an SPN that contains the NetBIOS name. The following examples show SPNs that are constructed by using NetBIOS host names:

MyDBService/*host1*/CN=hrdb,OU=mktg,DC=example,DC=com

MyDBService/*host2*/CN=hrdb,OU=mktg,DC=example,DC=com

MyDBService/*host3*/CN=hrdb,OU=mktg,DC=example,DC=com

Another example of a replicable service is one that provides services to an entire domain. In this case, the *ServiceName* component is the DNS name of the domain being served. A Kerberos KDC is an example of this type of replicable service.

Be aware that if the DNS name of a computer changes, that computer updates the *Host* element for all registered SPNs for that computer in the forest.

With Active Directory you can find and connect to a directory-enabled service by using attributes other than the computer name or the IP address of the computer on which the service exists. Active Directory uses other object attributes, such as the service display name (commonly called the friendly name) to find services like Dynamic Host Configuration Protocol (DHCP).

**Service Publication and SPN Processes and Interactions**
This section describes the processes and interactions that are related to service publication and SPNs.

**How a Service Publishes Its Service Information**
Typically, a setup program for a service (or an installation option for the executable routine of the service) installs the service and creates a service account, if necessary, and the SCP. At service startup, the service itself checks to ensure that the SCP contains the correct binding information.

To create an object of any kind, the process that creates the object must have permission to create a child in the container in which the object is to be created. For example, to create the object CN=MS-SQL-SQLServer,CN=dbserver,OU=ComputersOU,DC=corp,DC=cohovineyard,DC=com, the process must have permission to create a child in the CN=dbserver,OU=ComputersOU,DC=corp,DC=cohovineyard,DC=com container. To remove an object, the process that deletes the object must have permission to delete a child from the container from which the object is to be deleted. Or, the process must have permission to delete the object itself. To update a connection point object, the process performing the operation must have permission to write to the properties of the object that needs to be updated.

These permissions to create objects might not be available to a service account or computer account during the installation of a service, but an administrator, or anyone with the necessary permissions, can perform the installation. When you install a published service, you can set the protection on any connection point object to allow updates by the service at run time. This way, the service can run with a minimally privileged account as a security precaution and still be able to maintain its connection point.

**Where Services Are Published**
In the domain hierarchy of Active Directory, service-related objects can be created in three places:

- Computer object
- Organizational Unit container
- System container (or one of its children)

The following sections describe these locations in Active Directory for service-related objects.

**Computer object**
When a service that supports service publication is installed on a computer, the service generally (but not always) publishes its connection point object as a child object of the Computer object on which the service is installed.

The Computer object where the service is installed is a convenient and intuitive location for these objects, because when a Computer object is moved to a different location, Active Directory also moves any service objects on the Computer object to the new location. In addition, when a Computer object is deleted or when it is joined to a different domain, the service objects are also deleted or moved; therefore, there is less likelihood that a service will become an orphan.

### Note

- A computer must be a domain member to have a valid Computer object in the directory.

**Organizational Unit container**
An Organizational Unit container is a general-purpose container that can be used to group together and organize objects, including computer objects, for administrative purposes. Organizational Unit containers can also be used to group connection point objects together for administrative purposes.

All of the object classes that are derived from the **connectionPoint** class are valid children of the **organizationalUnit** object class, from which Organizational Unit containers are derived. This means that a service connection point object can be created in any container.

**System container**
The System container stores operational information, which is organized by domain. This information includes the default local security policy, file link tracking, network meetings, and containers for connection points to RPC and Winsock. By default, the System container is hidden, and it provides a convenient place for storing objects that are of interest to administrators but not

to users.

Services that use the Winsock RnR or RpcNs APIs to advertise themselves automatically create connection point objects in the WinsockServices and RpcServices containers.

Services that create connection point objects in the System container perform the following operations when they create a connection point object:

- They create a subcontainer in the System container and give it a name that clearly relates it to the service.

- They publish the connection point objects in the System container. For example, Microsoft NetMeeting uses the Meetings container to publish network meeting objects. When multiple products from a single vendor publish objects in this manner, the containers for each service must be created as children of a container (an object of the class **container**), and the container must be given a name that clearly relates to the particular vendor. The vendor-specific container must be created as a child of the System container.

**Publishing with the RPC Name Service**
RPC services use the RpcNs API to publish themselves in a namespace; the RpcNs APIs publish the RPC entries in Active Directory. Services create RPC bindings and publish them in the Active Directory namespace as named RPC Server entries with attributes. The attributes include the unique interface ID and a GUID that are known to clients. Clients can then search for RPC servers that offer the desired interface, import the binding, and connect to the server.

**Publishing with Winsock RnR**
Winsock services can use the RnR APIs to publish services and to locate those services. RnR publication occurs in two steps. The first step installs a service class that associates a GUID with a name for the service. The service class can hold service-specific configuration information. Services can then publish themselves as instances of the service class. After the services are published, clients can query Active Directory for instances of a given class, using the RnR APIs, and then select an instance to which to bind.

**How a Service Registers Its SPNs**
Before a client can use an SPN to authenticate an instance of a service, the SPN must be registered on the user or computer account that the service instance will use to log on. Typically, SPN registration is performed by the service installation program that installs the service.

The service installation program that installs a service instance on a host computer typically performs the following steps:

- Calls the **DsGetSpn** function to create one or more unique SPNs for the service instance.

- Calls the **DsWriteAccountSpn** function to register the names on the service's logon account.

**DsWriteAccountSpn** registers SPNs as a property of a user or computer account object in the directory. User and Computer objects have a **servicePrincipalName** attribute, which is a multivalued attribute for storing all the SPNs that are associated with a user or computer account. If the service runs under a user account, the SPNs are stored in the **servicePrincipalName** attribute of that account. If the service runs in the Local System account, the SPNs are stored in the **servicePrincipalName** attribute of the account of the service's host computer. The **DsWriteAccountSpn** caller must specify the distinguished name of the account object under which the SPNs are stored.

To ensure that registered SPNs are secure, the **servicePrincipalName** attribute cannot be written directly; it can only be written by calling **DsWriteAccountSpn**. The caller must have Write access to the **servicePrincipalName** attribute of the target account. Typically, Write access is granted by default only to domain administrators. However, there is a special case in which Active Directory allows a service running under the Local System account to register its own SPNs on the computer account of the service's host. In this case, the SPN being written must have the form *ServiceClass*/*Host*, and *Host* must be the DNS name of the local computer.

**DsWriteAccountSpn** can also remove SPNs from an account. An operation parameter indicates whether the SPNs are to be added to the account, removed from the account, or used to completely replace all current SPNs for the account.

For more information about how a service registers its SPNs, see the "Microsoft Platform SDK" on MSDN.

Host-based services that use the simple SPN format *ServiceClass*/*Host* have the option of using the **DsServerRegisterSpn** function, which both creates and registers SPNs for a service instance. **DsServerRegisterSpn** is a helper function that calls **DsGetSpn** and **DsWriteAccountSpn**.

**How a Client Application Finds a Service**
Clients find service objects by querying the directory. The client can limit the scope of the query to its domain, or it can search the entire forest by using the global catalog. In either case, the client does not need information about the location of the service to perform the search.

If the connection point objects that are being searched for are direct instances of the **serviceConnectionPoint** class, a client application can locate published services by searching for any object where **objectCategory** is equal to **serviceConnectionPoint** and **objectClass** is equal to **serviceConnectionPoint**. The **keywords**attribute contains the vendor-specific and application-specific GUID.

For more information about how a client application finds a service, see the "Microsoft Platform SDK" on MSDN.

**How a Client Composes an SPN**
To mutually authenticate a service, a client application composes an SPN for the service instance to which it wants to connect and then presents this SPN to the KDC for authentication. The client application can use **DsMakeSpnfunction** to compose an SPN. The client specifies the components of the SPN by using known information or information that is retrieved from sources other than the service itself.

The form of an SPN is as follows, where *ServiceClass* and *Host* are required and *Port* and *ServiceName* are optional:

*ServiceClass/Host:Port/ServiceName*

A client application can retrieve components of the SPN from sources such as a connection point object, user input, or hard-coded strings that are contained within the client application. For example, the client can read the **serviceDNSName** attribute of a service's connection point object to get the *Host* component. The **serviceDNSName** attribute contains either the DNS name of the server on which the service instance is running or the DNS name of SRV records that contain the host information for service replicas. The *ServiceName* component, which is used only for replicable services, can be the distinguished name of the service's connection point object, the DNS name of the domain that is served by the service, or the DNS name of SRV or MX records.

**How a Client Authenticates a Service**
The means by which a client authenticates a service ensures the authenticity of the service in two ways. First, because an SPN is constructed from several elements and must be unique, a rogue service has no way of publishing a valid SPN to the directory. Next, a given service can only be authenticated if that service can successfully exchange encrypted information with a KDC.

The Kerberos mechanism by which a client authenticates a service works as follows. When a service is installed, a service installer, running with administrator privileges, registers one or more unique SPNs for each service instance. The names are registered in Active Directory on the user or computer account object that the service instance will use to log on. When a client requests a connection to a service, it composes an SPN for a service instance, using known data or data that is provided by the user. The client then uses the Security Support Provider Interface (SSPI) negotiate package to present the SPN to the KDC for the client domain account. The KDC searches the forest for a user or computer account on which that SPN is registered. If the SPN is registered on more than one account, the authentication fails (which prevents a rogue service from attempting to impersonate a valid service). Otherwise, the KDC encrypts a message using the password of the account on which the SPN was registered. The KDC passes this encrypted message to the service instance. The service uses the SSPI negotiate package to decrypt the message (using the password that only the service and the KDC possess), which it passes back to the KDC. The KDC authenticates the service if the decrypted message matches the original message.

**Related Information**

- "Microsoft Platform SDK" on MSDN for more information about SQL Server 2000 service publication (in "Active Directory Integration")
- "Microsoft Platform SDK" on MSDN for more information about how a service registers its SPNs (in "How a Service Registers its SPNs")
- "Microsoft Platform SDK" on MSDN for more information about how a client application finds a service (in "Service Publication")
- "Microsoft Platform SDK" on MSDN for more information about the use of SPNs in authentication (in "About Mutual Authentication Using Kerberos")

# Service Publication and Service Principal Names Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Service Publication and SPN Tools
- Service Publication and SPN WMI Classes
- Network Ports Used by Service Publication and SPNs


This section contains information about the tools, Windows Management Instrumentation (WMI) classes, and network ports that are associated with service publication and service principal names (SPNs).

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to Active Directory, but the information is also applicable to AD DS. |

**Service Publication and SPN Tools**

The following tools are associated with service publication and SPNs.

**Adsiedit.msc: ADSI Edit**

**Category**

This tool ships with Support Tools for Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
| --- | --- |

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

ADSI Edit is a Microsoft Management Console (MMC) tool that you can use to view and modify directory objects and attributes, including connection point objects, and SPN attributes.

To find more information about ADSI Edit, see "Support Tools Help" in Tools and Settings Collection.

**Repadmin.exe: Repadmin**

**Category**
This tool ships with Support Tools for Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Administrators can use Repadmin to monitor and manage replication between domain controllers. To find more information about Repadmin, at a command prompt type **repadmin /?** or see "Command-Line References" in Tools and Settings Collection.

**Setspn.exe: Setspn**

**Category**
This tool ships with Support Tools for Windows Server 2003.

**Version compatibility**

| Can Be Run From | Can Be Run Against |
|---|---|

| | |
|---|---|
| Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Administrators can use this command-line tool to read, modify, and delete values in the **servicePrincipalNames** attribute on an Active Directory service account object.

To find more information about Setspn, see "Support Tools Help" in Tools and Settings Collection.

**Service Publication and SPN WMI Classes**

The following table lists and describes the WMI classes that are associated with service publication and SPNs.

**WMI Classes Associated with Service Publication and SPNs**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| **DS_LDAP_Class_Containment** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

| | | |
|---|---|---|
| **DS_LDAP_Instance_Containment** | root\directory\LDAP | Domain controllers running:<br>• Windows Server 2008 R2<br>• Windows Server 2008<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

For more information about these WMI classes, see "Mapping Active Directory to WMI" in the WMI SDK documentation on MSDN.

**Network Ports Used by Service Publication and SPNs**

The network ports that are used by service publication and SPNs are listed in the following table.

**Port Assignments for Service Publication and SPNs**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | None | 389 |
| LDAP SSL | None | 636 |
| RPC Endpoint Mapper | 135 | 135 |
| Global Catalog LDAP | None | 3268 |
| Global Catalog LDAP SSL | None | 3269 |
| Kerberos | 88 | 88 |

# Active Directory Installation, Upgrade, and Migration Technologies

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Active Directory Installation Wizard
- Active Directory Preparation Tool
- Active Directory Migration Tool
- Active Directory Functional Levels
- Domain Rename


The Active Directory directory service is used in server operating system environments to manage networks comprising Microsoft Windows clients, Windows 2000 Server and Windows Server 2003 servers and domain controllers, and Windows-compatible applications and services. There are many benefits to deploying Active Directory, including centralized management, scalability, and the ability to easily delegate administrative authority.

Some of the tools that are available to help you install and manage Active Directory include the following:

- The Active Directory Installation Wizard (Dcpromo.exe), which is used to install Active Directory
- The Active Directory Preparation Tool (Adprep.exe), which is used to prepare an existing Active Directory environment for a newer version of the directory service
- The Active Directory Migration Tool (ADMT), which is used to restructure Microsoft Windows NT 4.0 or Active Directory domains
- The domain rename tools (Rendom.exe and Gpfixup.exe), which are used to rename and restructure domains and adjust Group Policy settings to accommodate the new structure

Finally, new Active Directory features are enabled in a Windows Server 2003 environment by increasing domain and forest functional levels.

**Active Directory Installation Wizard**

The Active Directory Installation Wizard (Dcpromo.exe) configures a server to be a domain controller by installing Active Directory. Using the Active Directory Installation Wizard you can either add a domain controller to an existing domain or create the first domain controller in a new domain. The wizard provides the option to do this.

You can run the Active Directory Installation Wizard from the command line, or from the Configure Your Server Wizard. You can also install Active Directory using an unattended setup script called an answer file.

When running the wizard from the command line, you can append the **/adv** switch to the **dcpromo** command to populate the directory using a backup of system state data from another domain controller in the same domain. Installing from backup media reduces the amount of data that must be replicated over the network, thus reducing the time required to install Active Directory.

Before installing Active Directory, the wizard verifies that the server is eligible to run Active Directory by checking certain criteria. The following are examples of some of the criteria that the wizard checks before allowing an Active Directory installation to proceed:

- There is sufficient disk space on the computer to store the Active Directory database.

- You have sufficient privileges to install Active Directory on the computer.
- The computer is running a correct version of the operating system.

After all prerequisites have been met, a user interface is used to gather information specific to the environment in which Active Directory will be installed, such as the Domain Name System (DNS) name and the storage locations for the Active Directory database and the SYSVOL shared folder. If you are installing Active Directory using an answer file, you will prepopulate the script with this information. Finally, the wizard configures Active Directory and makes the server a domain controller.

**Active Directory Preparation Tool**
The Active Directory Preparation Tool (Adprep.exe) prepares your infrastructure for a new version of Active Directory. Adprep.exe is located on the Windows Server 2003 operating system CD and can only be run from the command line.

To prepare your Windows 2000 forest and domains for upgrade to Windows Server 2003 Active Directory, or to prepare for new Windows Server 2003–based domain controllers, you must:

- Run **adprep /forestprep** once on the schema master to prepare the forest.
- Run **adprep /domainprep** once on the infrastructure master in each domain in which you plan to install a Windows Server 2003–based domain controller.

Adprep.exe performs a collection of operations that prepare the forest and domains for an Active Directory upgrade, including:

- Extending the current schema with new schema information, while preserving previous schema modifications in your environment.
- Allowing applications to access data in Active Directory after security settings have been improved.

**Active Directory Migration Tool**
The Active Directory Migration Tool (ADMT) is used to consolidate Windows NT 4.0 or Active Directory domains to improve manageability, scalability, and security. You can use ADMT to restructure your Windows NT 4.0 domains to an Active Directory forest or to restructure Windows Server 2003 Active Directory domains within and between Active Directory forests.

When you use ADMT to migrate a user or group from either a Windows NT 4.0 domain to an Active Directory forest or between two Active Directory forests, you can use security identifier (SID) history to maintain resource permissions during migration so that access to resources is not interrupted.

ADMT includes wizards that automate common migration tasks, such as copying users, groups, and service accounts; moving computers; migrating trusts; and performing security translation. ADMT is available on the Windows Server 2003 operating system CD, in the Admigration.msi file in the \i386\admt directory. You can perform ADMT tasks by using the ADMT console, by using a command-line procedure, or by using a script. When running ADMT from the command line, you can use an option file to specify command-line options.

**Active Directory Functional Levels**
Although most Active Directory features are available by default, certain new features of Windows Server 2003 Active Directory, such as efficient group membership replication and domain rename, cannot be enabled until all domain controllers within a given scope (domain or forest) are running Windows Server 2003.

Windows Server 2003 Active Directory does not automatically enable such a feature even when all domain controllers within a forest are running Windows Server 2003. Instead, the responsibility for the decision is administrative so that you can still add domain controllers that are running earlier versions of Windows. By advancing the functional level of a Windows Server 2003 domain or forest, you enable new features within that scope.

Functional levels protect against incompatibility. Before you raise a functional level, Active Directory verifies that all domain controllers are running the correct version of the operating system. After you raise the functional level, Active Directory no longer allows the introduction of a domain controller with an incompatible version of Windows.

**Domain Rename**
Domain rename capabilities are available in a Windows Server 2003 forest that has a forest functional level of

Windows Server 2003. These capabilities are not available in Windows 2000 Server operating systems.

The structure of an Active Directory forest is the result of the order in which you create domains and the hierarchical names of those domains. Beginning with the forest root domain, all child domains derive their distinguished names and default DNS names from the forest root domain name. The same is true of every additional tree in the forest. The way to change the hierarchical structure of an existing domain tree is to rename the domains. For example, you can rename a child domain to have a different parent or rename a child domain to be a new tree-root domain. In each case, you reposition an existing domain to create a different domain-tree structure. Alternatively, you can rename domains without affecting the structure. For example, if you rename a root domain, the names of all child domains below it are also changed, but you have not created a different domain-tree structure.

The ability to rename domains provides you with the flexibility to make important name changes and forest structural changes as the needs of your organization change. Using domain rename, you can not only change the name of a domain, but you can change the structure of the domain hierarchy and change the parent of a domain or move a domain residing in one domain tree to another domain tree. The domain rename process can accommodate scenarios involving acquisitions, mergers, or name changes in your organization, but it is not designed to accommodate forest mergers or the movement of domains between forests.

The domain rename process involves making basic changes independently at each domain controller in a forest. You set up an administrative computer from which you issue commands that are executed remotely at each domain controller. These commands update the directory database at each domain controller individually with the changes that are necessary for renaming the domains; that is, the updates that rename the domains do not spread across the forest through Active Directory replication.

Rendom.exe is a new tool that is included with Windows Server 2003. You use the Rendom tool to carry out the following multiple steps in the domain rename process:

- Freeze the current state of the forest so that no changes can occur while a domain rename operation is being performed.
- Prepare the contents of the forest for a domain rename operation. Rendom runs multiple scripts that perform this preparation.
- Execute a domain rename operation.
- Clean up old domain names.

Another tool, Gpfixup, is provided to reinstate Group Policy from the original domains into the newly named domains in the forest.

In Windows Server 2003, domain rename is intended to be a supported method for renaming domains when domain renames are necessary; *it is not intended to make domain rename a routine operation*. The domain rename process is complex and requires a great deal of care in planning and execution. In addition, the time that is required for a complete domain rename operation is directly proportional to the size of an Active Directory forest in terms of its number of domains, domain controllers, and member computers. Therefore, although domain rename is possible in Windows Server 2003, it should not be undertaken lightly.

# Active Directory Functional Levels Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Active Directory Functional Levels Technical Reference**

Domain and forest functional levels provide a way to enable advanced Active Directory features that are appropriate to the Microsoft Windows Server operating systems running on domain controllers on your network. Different functional levels exist for both domains and forests so that your organization can maximize the capabilities of Active Directory Domain Services (AD DS) within each domain or forest that you maintain.

This subject explains the advanced features that are provided at each domain and forest functional level and how to use Active Directory administration tools to raise the functional level.

| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

**In this subject**


- What Are Active Directory Functional Levels?
- How Active Directory Functional Levels Work
- Active Directory Functional Levels Tools and Settings

# What Are Active Directory Functional Levels?

Updated: November 19, 2014

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Active Directory Functional Level Scenarios
- Active Directory Functional Level Dependencies and Rollback Options


| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

In Active Directory Domain Services (AD DS), domain controllers can run different versions of Windows Server operating systems. The functional level of a domain or forest depends on which versions of Windows Server operating systems are running on the domain controllers in the domain or forest. The functional level of a domain or forest controls which advanced features are available in the domain or forest.

Ideally, all servers in an organization could run the latest version of Windows and take advantage of all the advanced features that are available with the newest software. But organizations often have a mixture of systems, generally running different versions of operating systems, which are migrated to the latest version only as organizational requirements demand additional functionality, either for the entire organization or for a specific area of the organization.

AD DS supports phased implementation of new versions of Windows Server and advanced features on domain controllers by providing multiple functional levels, each of which is specific to the versions of Windows Server operating systems that are running on the domain controllers in the environment. These functional levels provide configuration support for the AD DS features and ensure compatibility with domain controllers running earlier versions of Windows Server.

AD DS does not automatically enable advanced features, even if all domain controllers within a forest are running the same version of Windows Server. Instead, an administrator raises a domain or forest to a specific functional level to safely enable advanced features when all domain controllers in the domain or forest are running an appropriate version of Windows Server. When an administrator attempts to raise the functional level, AD DS checks whether all domain controllers are running an appropriate Windows Server operating system to ensure the proper environment for enabling new Active Directory features.

Raising the functional level allows the introduction of advanced features but also limits the versions of Windows Server that can run on domain controllers in the environment. AD DS has two types of functional levels:

- **Domain functional level**. Six domain functional levels are available:

  Windows 2000 mixed (the default in Windows Server 2003)

  Windows 2000 native

  Windows Server 2003 interim

  Windows Server 2003

  Windows Server 2008

  Windows Server 2008 R2

  Setting the functional level for a domain enables features that affect the entire domain and that domain only. If all domain controllers in a domain are running Windows Server 2008 R2 and the functional level is set to Windows Server 2008 R2, all domain-wide features are available.

> **Note**
> The default domain functional level for Windows Server 2008 and Windows Server 2008 R2 is set to the existing value of the forest functional level.

- **Forest functional level**. Five forest functional levels are available:

  Windows 2000 (the default in Windows Server 2003 and Windows Server 2008)

  Windows Server 2003 interim

  Windows Server 2003 (the default in Windows Server 2008 R2)

  Windows Server 2008

  Windows Server 2008 R2

  Setting the functional level for a forest enables features across all the domains within a forest. If all domain controllers in a forest are running Windows Server 2008 R2 and the functional level is set to Windows Server 2008 R2, all forest-wide features are available.

> **Note**
> In addition to the Windows Server 2008 R2 forest functional level requirement, the Active Directory Recycle Bin feature must be enabled before it becomes available for use. For more information, see the Active Directory Recycle Bin Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkID=178657).

When domain controllers running earlier versions of Windows Server are included in your domain or forest with domain controllers running later versions of Windows Server, advanced Active Directory features are limited.

The concept of enabling additional functionality in AD DS exists in Windows 2000 with mixed and native modes. Mixed-mode domains can contain Windows NT 4.0 backup domain controllers and cannot use Universal security groups, group nesting, and security identifier (SID) history capabilities. When the domain is set to native mode, Universal security groups, group nesting, and security identifier (SID) history capabilities are available. Domain controllers running Windows NT 4.0 or Windows 2000 Server are not aware of Windows Server 2003 or higher domain and forest functional levels.

**Active Directory Functional Level Scenarios**

Functional levels provide a method for any organization to introduce the latest Active Directory functionality. The advanced functionality can be optimized immediately, or phased in over time, depending on the needs of each organization. The following scenarios describe how functional levels can meet a range of deployment needs:

- A small organization running only Windows Server 2008 R2 on domain controllers. In this scenario, a small organization might be creating an AD DS infrastructure for the first time, using only two domain controllers in a forest with a single domain. By installing Windows Server 2008 R2 on both domain controllers and selecting the Windows Server 2008 R2 forest functional level during the AD DS installation on the first domain controller, the organization can immediately take advantage of all of the advanced Active Directory features that are available.

> **Note**
> The option to select the domain or forest functional level during AD DS installation is available only in Windows Server 2008 and Windows Server 2008 R2.

- A mid-size organization with accelerated upgrade requirements. In this scenario, a mid-size organization might have a native mode Windows Server 2003 domain, and would now like to start upgrading domain controllers to Windows Server 2008. Upgrading to Windows Server 2008 on some domain controllers, while maintaining some previously installed domain controllers running Windows Server 2003, enables the mid-size organization to take advantage of much of the latest functionality, even though all domain controllers are not running Windows Server 2008. Later, after all domain controllers running Windows Server 2003 have been upgraded to Windows Server 2008, the organization can raise the forest functional level to take advantage of all advanced Active Directory features.

- A large organization with phased upgrade requirements. In this scenario, an enterprise with a mature and complex forest with multiple domains, some of which are still running Windows 2000 Server, might now be ready to install one or more new domain controllers running Windows Server 2008. Members of the organization might want to raise functional levels in a phased manner (one domain at a time). They could target certain domains that are running Windows 2000 domain controllers and upgrade them first to Windows Server 2003 and then to Windows Server 2008. After they upgrade all domain controllers in the domain, they could raise the functional level of only that domain to Windows Server 2008 to take advantage of additional Active Directory features.

  After they are satisfied with the results of that domain, members of the organization can carry out a similar exercise in other domains. After upgrading every domain in the forest to Windows Server 2008, they can finally raise the functional level of the forest to Windows Server 2008 to take advantage of the advanced Active Directory features for Windows Server 2008.

**Active Directory Functional Level Dependencies and Rollback Options**

Active Directory domain and forest functionality has the following dependencies:

- After all domain controllers are running an appropriate version of Windows Server, the domain or forest must be configured to support the appropriate domain or forest functional level. That is, to provide support in a domain or forest for advanced Active Directory features, an administrator must raise the domain functional level or forest functional level, which can only be done if the domain controllers are each running the appropriate version of Windows Server.

- After the domain functional level is raised, domain controllers running earlier versions of Windows Server cannot be introduced into the domain. After the forest functional level is raised, domain controllers running earlier versions of Windows Server cannot be introduced into the forest.

- With versions of Windows Server that are earlier than Windows Server 2008 R2, you cannot roll back or lower a functional level under any circumstances. If you have to revert to a lower functional level with a version of Windows Server that is earlier than Windows Server 2008 R2, you must rebuild the domain or forest or restore it from a backup copy.

- After you set the domain functional level to a certain value in Windows Server 2008 R2, you cannot roll back or lower the domain functional level, with one exception: when you raise the domain functional level to Windows Server 2008 R2 and if the forest functional level is Windows Server 2008 or lower, you have the option of rolling the domain functional level back to Windows Server 2008. You can lower the domain functional level only from Windows Server 2008 R2 to Windows Server 2008. If the domain functional level is set to Windows Server 2008 R2, it cannot be rolled back, for example, to Windows Server 2003.

- After you set the forest functional level to a certain value in Windows Server 2008 R2, you cannot roll back or lower the forest functional level, with one exception: when you raise the forest functional level to Windows Server 2008 R2 and if the Active Directory Recycle Bin is not enabled, you have the option of rolling the forest functional level back to Windows Server 2008. For more information about the Active Directory Recycle Bin, see What's New in AD DS: Active Directory Recycle Bin (http://go.microsoft.com/fwlink/?LinkId=141392). You can lower the forest functional level only from Windows Server 2008 R2 to Windows Server 2008. If the forest functional level is set to Windows Server 2008 R2, it cannot be rolled back to Windows Server 2003, for example.

# How Active Directory Functional Levels Work

Updated: November 19, 2014

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- The Architecture of Active Directory Functional Levels
- Physical Structure for Active Directory Functional Levels
- Processes and Interactions with Active Directory Functional Levels
- Related Information


| **Note** |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

Different Active Directory features are available at different functional levels. Raising domain and forest functional levels is required to enable certain new features as domain controllers are upgraded from earlier versions of Windows Server to later versions. Even for a new installation of Windows Server 2003 Active Directory (as opposed to an upgrade), an administrator must raise functional levels to obtain the full capabilities of Windows Server 2003 Active Directory. However, in Windows Server 2008 and Windows Server 2008 R2, you can select a domain or forest functional level during the installation of Active Directory Domain Services (AD DS).

**The Architecture of Active Directory Functional Levels**

Active Directory functional levels are of two types, as shown in the following table.

**Domain and Forest Functional Levels**

| Type of Functional Level | Description |
| --- | --- |
| Domain functional level | Domain functional levels enable features that affect the entire domain and that domain only. They also control which Windows Server operating systems can be run on domain controllers in the domain. |
| Forest functional level | Forest functional levels enable features across all domains within a forest. They also control which Windows Server operating systems can be run on domain controllers in all domains in the forest. |

The following table shows the relationship and dependencies between domain and forest functional levels.

**Relationship Between Domain and Forest Functional Levels**

| Domain Functional Level | Forest Functional Level |
| --- | --- |
| Windows 2000 mixed | Windows 2000 |
| Windows 2000 native | Windows 2000 |
| Windows Server 2003 | Windows Server 2003 |
| Windows Server 2008 | Windows Server 2008 |

| | |
|---|---|
| Windows Server 2008 R2 | Windows Server 2008 R2 |

Another domain functional level, Windows Server 2003 interim, is available for upgrades directly from Windows NT 4.0 to Windows Server 2003.

**Default Windows Server 2003 Active Directory Features Available with All Functional Levels**

Although many new features in Windows Server 2003 Active Directory have operational requirements that are available only at specific functional levels, most Windows Server 2003 Active Directory features are available by default at all functional levels. When the first Windows Server 2003–based domain controller is deployed in a domain or forest, a set of default Active Directory features becomes available. The following table summarizes the Active Directory features that are available by default on any domain controller running Windows Server 2003.

**Default Windows Server 2003 Active Directory Features**

| Feature | Functionality |
|---|---|
| Multiple selection of user objects | Allows you to modify common attributes of multiple user objects at one time. |
| Drag and drop functionality | Allows you to move Active Directory objects from container to container by dragging one or more objects to a location in the domain hierarchy. You can also add objects to group membership lists by dragging one or more objects (including other group objects) to the target group. |
| Efficient search capabilities | Search functionality is object-oriented and provides an efficient search that minimizes network traffic associated with browsing objects. |
| Saved queries | Allows you to save commonly used search parameters for reuse in Active Directory Users and Computers. |
| Active Directory command-line tools | Allows you to run new directory service commands for administration scenarios. |
| **InetOrgPerson** class | The **InetOrgPerson** class has been added to the base schema as a security principal. It can be used in the same manner as the user class. |
| Application directory partitions | Allows you to configure the replication scope for application-specific data among domain controllers. For example, you can control the replication scope of Domain Name System (DNS) zone data stored in Active Directory so that only specific domain controllers in the forest participate in DNS zone replication. |
| Ability to add additional domain controllers by using backup media | Reduces the time it takes to add an additional domain controller in an existing domain by using backup media. |
| Universal group membership caching | Prevents the need to locate a global catalog across a wide area network (WAN) when logging on by storing universal group membership information on an authenticating domain controller running Windows Server 2003. |
| Secure Lightweight Directory Access Protocol (LDAP) traffic | Active Directory administrative tools sign and encrypt all LDAP traffic by default. Signing LDAP traffic guarantees that the packaged data comes from a known source and that it has not been tampered with. |
| Partial synchronization of the global catalog | Provides improved replication of the global catalog when schema changes add attributes to the global catalog partial attribute set. Only the new attributes are replicated, not the entire global catalog. |
| Active Directory quotas | Quotas can be specified in Active Directory to control the number of objects a user, group, or computer can own in a given directory partition. Members of the Domain Admins and Enterprise Admins groups are exempt from quotas. |

For more information about new AD DS features in Windows Server 2008, see What's New in AD DS in Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkId=180673). For more information about new AD DS features in Windows Server 2008 R2, see What's New in Active Directory Domain Services (http://go.microsoft.com/fwlink/?LinkId=139655).

**Domain Functional Level**
Domain functional levels represent the extension of the mixed and native domain modes that were available in Windows 2000:

- In Windows 2000 Active Directory, certain features are enabled only when all domain controllers in a domain are running Windows 2000 and the domain mode is set to native mode. Otherwise, the domain is running in mixed mode, where only features that are compatible with Windows NT 4.0–based domain controllers are operational. The domain administrator must change the mode from mixed to native and ensure that Windows NT 4.0–based domain controllers are no longer required in the domain.

- In AD DS, the domain functional levels that correspond to mixed mode and native mode are the Windows 2000 mixed and Windows 2000 native functional levels, respectively. Raising the domain functional level from Windows 2000 mixed to Windows 2000 native is still based on the discretion of the administrator, and Active Directory does not verify that the domain has no Windows NT 4.0–based domain controllers in the domain.

Before a domain can be raised to other functional levels, all domain controllers in the domain must be running a corresponding or later version of Windows Server. For example, before a domain can be raised to the Windows Server 2008 domain functional level, all domain controllers must be running Windows Server 2008 or Windows Server 2008 R2. Before allowing the administrator to raise a domain to this functional level, AD DS verifies that all domain controllers in the domain are running an appropriate version of Windows Server.

For those organizations that are upgrading from Windows NT 4.0 directly to Windows Server 2003, another domain functional level called Windows Server 2003 interim accommodates only domain controllers that are running Windows NT 4.0 and Windows Server 2003 and it excludes Windows 2000–based domain controllers. For more information about this interim domain functional level, see "Functional Levels When Upgrading Windows NT 4.0 Domains" later in this section.

Just as in Windows 2000 where you could not revert from a native mode to a mixed mode domain, raising the domain functional level is not reversible, except for specific situations in Windows Server 2008 R2. For more information, see Active Directory Functional Level Dependencies and Rollback Options. The only changes that occur in an environment when raising the functional level is the availability of advanced Active Directory features.

The default domain functional level for a new Windows Server 2003 installation or upgrade is Windows 2000 mixed in all cases, except when you are upgrading a Windows 2000 domain that is in native mode. In this case, the default domain functional level is Windows 2000 native.

For new domain installations in Windows Server 2008 and Windows Server 2008 R2, the default domain functional level is automatically computed and set to the existing forest functional level or the value that is set for the forest functional level when you run Dcpromo.exe.

**Features Enabled by the Domain Functional Level**
Various features are enabled at different domain functional levels, and only certain operating systems are supported on domain controllers at each level. For more information, see Understanding AD DS Functional Levels (http://go.microsoft.com/fwlink/?LinkId=180702)

**Forest Functional Level**
Forest functional levels are new in Windows Server 2003. Features that become available at forest functional levels must be compatible with all domain controllers in the forest.

**Features Enabled at Forest Functional Levels**
Various features are enabled at different forest functional levels, and only certain operating systems are supported on domain controllers at each level. For more information, see Understanding AD DS Functional Levels (http://go.microsoft.com/fwlink/?LinkId=180702)

As a side effect of raising the forest functional level to Windows Server 2003, the following new attributes are replicated to the partial attribute set of the global catalog:

- **Ms-DS-Trust-Forest-Trust-Info**

- **Trust-Direction**

- **Trust-Attributes**

- **Trust-Type**

- **Trust-Partner**
- **Security Identifier**
- **Ms-DS-Entry-Time-To-Die**
- **MSMQ-Secured-Source**
- **MSMQ-Multicast-Address**
- **Print-Memory**
- **Print-Rate**
- **Print-Rate-Unit**
- **MS-DRM-Identity-Certificate**

When the forest functional level is Windows Server 2003 or Windows Server 2003 interim, the global catalog is present on Windows Server 2003 domain controllers. Therefore, replication traffic due to the addition of these attributes to the partial attribute set is not significant. This is because of the feature, new in Windows Server 2003, that allows partial synchronization of the global catalog.

**Functional Level for Windows Server 2008 and Windows Server 2008 R2 Deployments**
In a new deployment of AD DS that does not involve upgrading an existing domain, if the first domain controller in the forest is running Windows Server 2008 or Windows Server 2008 R2, you can choose the forest and domain functional levels during the AD DS installation of the first domain controller in the forest root domain. You can choose the domain functional level during the AD DS installation of the first domain controller for each additional domain in the forest. You can choose different functional level values for the forest and for each domain that you create within the forest, but the domain functional level can never be lower than the forest functional level.

If you do not specify a higher value when you create the forest root domain for Windows Server 2008, the default forest functional level is Windows 2000. If you do not specify a higher value when you create the forest root domain for Windows Server 2008 R2, the default forest functional level is Windows Server 2003.

If you do not specify a higher value when you create a domain in Windows Server 2008 or Windows Server 2008 R2, the default domain functional level is set to the current value of the forest functional level or, in the case of the forest root domain, the value that you choose for the forest functional level during the installation of AD DS.

**Functional Level for Windows Server 2003 Deployments**
In a new deployment of Windows Server Active Directory that does not involve upgrading an existing domain, the first domain controller is created by installing Active Directory on a server that is running Windows Server 2003 and it becomes the domain controller for the forest root domain. The forest functional level of the newly created forest is Windows 2000 and the domain functional level is Windows 2000 mixed. Immediately raising the domain level and then the forest level to Windows Server 2003 provides full Windows Server 2003 functionality with no further intervention. All domains created subsequently in the forest are created at the Windows Server 2003 functional level.

**Functional Level When Upgrading Windows NT 4.0 Domains**
Active Directory provides an intermediate functional level called Windows Server 2003 interim for user environments that currently have only Windows NT 4.0-based domain controllers and that have the following domain upgrade plans:

- Plan to upgrade to a Windows Server 2003 forest directly from Windows NT 4.0.
- Do not plan to install domain controllers running Windows 2000 at any time.

Moving to the Windows Server 2003 interim forest functional level is especially useful for large groups (greater than 5,000 members) in an existing deployment for which upgrading to Windows 2000 would require breaking those groups into smaller groups. Windows 2000 replication limits the size of Windows 2000 groups; Windows Server 2003 replication supports large groups (greater than 5,000 members) because of the link-valued replication feature, which is available only if the forest functional level is Windows Server 2003 interim or higher.

**Note**

- The Windows Server 2003 interim functional level does not support domain controllers that are running Windows 2000, Windows Server 2008, or Windows Server 2008 R2. For this reason, if domain controllers that run Windows 2000, Windows Server 2008, or Windows Server 2008 R2 are in your current deployment or there are plans to install them in the future, do not use the Windows Server 2003 interim functional level.

**Physical Structure for Active Directory Functional Levels**

Each functional level has requirements that are enforced by AD DS. AD DS includes an attribute that enables domain controllers to recognize the current functional level of any domain and of the forest. A numbering system identifies the different functional levels. The Windows Server 2003, Windows Server 2008, Windows Server 2008 R2 operating systems also use these level numbers as built-in indicators that enable domain controllers to identify the highest and lowest supported levels.

**Internal Representation of Functional Levels**

Internally, functional levels are represented as positive integers. The following table lists the domain and forest functional levels that are available in AD DS and their corresponding integer values.

**Active Directory Domain and Forest Functional Levels**

| Domain Functional Levels | Forest Functional Levels |
|---|---|
| 0 — Windows 2000 mixed | 0 — Windows 2000 |
| 0 — Windows 2000 native | 0 — Windows 2000 |
| 1 — Windows Server 2003 interim | 1 — Windows Server 2003 interim |
| 2 — Windows Server 2003 | 2 — Windows Server 2003 |
| 3 —Windows Server 2008 | 3 —Windows Server 2008 |
| 4 —Windows Server 2008 R2 | 4 —Windows Server 2008 R2 |

Both Windows 2000 mixed and Windows 2000 native domain functional levels are internally represented by the integer 0.

**Storage of Functional Levels**

In an AD DS forest, the value of the functional levels of domain controllers, domains, and forests are stored in an **msDSBehaviorVersion** attribute. This attribute is present on three different objects in the directory, each of which is authoritative for identifying the functional level of one of these three respective scopes (domain controller, domain, or forest). When enforcing domain and forest functional levels, Active Directory uses the value of **msDSBehaviorVersion** on the appropriate object to detect the current functional level.

**Note**

- The functional level of a domain controller cannot be raised. Its functional level is fixed by the operating system.

The following table summarizes the storage of functional levels in the **msDSBehaviorVersion** attribute on the three defining objects and directory partition locations for the three respective scopes.

**Storage of Functional Levels in msDSBehaviorVersionAttribute**

| Functional Level Scope | Object | Directory Partition Where Object Is Located |
|---|---|---|
| Domain controller | NTDS Settings object (class **nTDSDSA**) | Configuration |
| Domain | Domain object (class **domainDNS**) | Domain |
| Forest | Partitions container object (class **crossRefContainer**) | Configuration |

**Storage of Domain Controller Functional Levels**

Beginning with Windows Server 2003, the operating system has two built-in, fixed values that identify the functional levels of

AD DS with which it is compatible, as follows:

- A fixed value that indicates the highest functional level that is supported by the current product release. Windows Server 2003 operating systems have a value of 2, indicating that Windows Server 2003 is the highest functional level supported. Because Windows 2000 and Windows NT 4.0 do not recognize functional levels, a domain controller running one of these operating systems has an implied highest functional level value of 0.

- A fixed value that indicates the lowest functional level that is supported by the current product release. Windows Server 2003 operating systems have a value of 0, indicating that they support both Windows 2000 and Windows NT 4.0. Windows 2000 and earlier operating systems have an implied lowest functional level value of 0.

AD DS uses these values to prohibit introduction of domain controllers that are running incompatible operating systems into the domain or forest.

When AD DS is installed to create a domain controller, the highest fixed value that is present in the operating system is stored in the **msDSBehaviorVersion** attribute of the NTDS Settings object (class **nTDSDSA**) associated with the domain controller. This value indicates the functional level of the domain controller. The NTDS Settings object is stored in the configuration directory partition (cn=NTDS Settings,cn=*ServerName*,cn=servers,cn=sites,cn=configuration,dc=*ForestRootDomainName*). Because every domain controller holds a copy of the configuration directory partition, any domain controller can discover the functional level of any other domain controller in the forest.

Absence of this attribute on the NTDS Settings object indicates that the domain controller is running Windows 2000, and is treated as having a value of 0. Windows NT 4.0–based domain controllers do not have an NTDS Settings object.

**Storage of Domain Functional Levels**
The functional level of a domain is identified by the value in the **msDSBehaviorVersion** attribute on the domain object (class **domainDNS**) in the domain directory partition (dc=*DomainName*,dc=*ForestRootDomainName*). On the basis of this value, certain features are enabled that have an impact in the scope of a single domain only. Absence of this attribute on a domain object is equivalent to domain functional level 0.

**Storage of Forest Functional Levels**
The functional level of a forest is identified by the value in the **msDSBehaviorVersion** attribute on the Partitions container object (class **crossRefContainer**) in the configuration directory partition (cn=partitions,cn=configuration,dc=*ForestRootDomainName*). On the basis of this value, certain features are enabled that affect the entire forest. Absence of this attribute on the Partitions container is interpreted as forest level 0.

**Processes and Interactions with Active Directory Functional Levels**
Functional levels protect against incompatibility. Before raising a functional level, AD DS verifies that all domain controllers are running the correct version of the operating system. After raising the functional level, AD DS no longer allows the introduction of a domain controller that is running an incompatible version of Windows Server.

Forest and domain functional levels are restrictive for domain controllers only. Clients can interact with Active Directory regardless of the domain or forest functional level that is in effect. In a domain that has a functional level of Windows Server 2008 R2, domain controllers can still authenticate computers that are running earlier versions of Windows.

**Raising Domain Functional Levels**
The domain functional level can be raised by using Active Directory Domains and Trusts or Active Directory Users and Computers, which are Active Directory administrative tools that are available on each domain controller. The ability to raise the domain functional level depends on the operating systems that are running on domain controllers in the environment:

- If an environment has domain controllers running Windows NT 4.0 as well as domain controllers running Windows 2000 in the domain, Windows Server 2003–based domain controllers can be introduced, but the domain functionality is essentially that of a Windows 2000 mixed-mode domain. Domain controllers that run Windows Server 2008 or Windows Server 2008 R2 cannot be installed in a domain that has domain controllers running Windows NT 4.0.

  It is possible to raise the domain functional level directly to Windows Server 2003 from Windows 2000 mixed after all domain controllers in the domain are upgraded to Windows Server 2003.

- After all Windows NT 4.0–based domain controllers in an environment are upgraded to either Windows 2000 or Windows Server 2003, the domain functional level can be raised to Windows 2000 native, which provides functionality that is essentially that of a Windows 2000 native-mode domain.

- After all domain controllers in the domain are upgraded to Windows Server 2003, the domain functional level can be raised to Windows Server 2003, which provides all Windows Server 2003 domain-level features.

- After all domain controllers in the domain are upgraded to Windows Server 2008, the domain functional level can be raised to Windows Server 2008, which provides all Windows Server 2008 domain-level features.

- After all domain controllers in the domain are upgraded to Windows Server 2008 R2, the domain functional level can be raised to Windows Server 2008 R2, which provides all Windows Server 2008 R2 domain-level features.

**Note**

- Raising the domain functional level cannot be reversed, except for specific situations in Windows Server 2008 R2. In all other cases, after the domain functional level is raised, domain controllers that are not supported at that level can no longer be introduced. For more information, see Active Directory Functional Level Dependencies and Rollback Options.

The following requirements apply to raising all domain functional levels:

- Raising the domain functional level requires Domain Admins credentials.

- The domain functional level can be raised on only the primary domain controller (PDC) emulator operations master. AD DS administrative tools (Active Directory Domains and Trusts and Active Directory Users and Computers) automatically target the PDC emulator when the domain functional level is raised.

**Raising Forest Functional Levels**

When Windows Server 2003 is installed on the first domain controller in the forest, the forest functional level is set to Windows 2000. The forest retains this level until an administrator changes it. Raising the forest level to Windows Server 2003 has the following requirements:

- All domain controllers in the forest are running Windows Server 2003.

- All domains in the forest have either the Windows 2000 native or the Windows Server 2003 functional level.

Therefore, if any domains are operating at the Windows 2000 mixed domain functional level, even if all domain controllers are running Windows Server 2003, AD DS does not allow the forest functional level to be raised to Windows Server 2003.

If all domain controllers in the forest are running Windows Server 2003, but some or all domains in the forest are still operating at the Windows 2000 native domain functional level, AD DS automatically raises the functional level of those domains to Windows Server 2003 when the forest functional level is raised to Windows Server 2003. This same behavior occurs at when the forest functional level is raised to higher values. In other words, if you raise the forest functional level to Windows Server 2008 R2, if all domains have domain controllers running Windows Server 2008 R2, they are automatically raised to the Windows Server 2008 R2 domain functional level.

Active Directory Domains and Trusts can be used to raise the forest functional level.

**Note**

- Raising the forest functional level cannot be reversed, except for specific situations in Windows Server 2008 R2. In all other cases, after the forest functional level is raised, domain controllers that run earlier versions of Windows Server can no longer be introduced into the forest. For more information, see Active Directory Functional Level Dependencies and Rollback Options.

The following requirements apply to raising all forest functional levels:

- Raising the forest functional level requires Enterprise Admins credentials.

- The forest functional level can be raised on only the schema operations master. Active Directory Domains and Trusts automatically targets the schema operations master when the forest functional level is raised.

**Enforcing Functional Levels**
To ensure consistent functionality across the domain controllers in a domain or forest, AD DS checks for functional levels, as follows:

- At domain controller creation.
- After a directory restore.
- At domain controller startup.
- When raising functional levels.

The Active Directory Domain Services Installation Wizard, the application that installs AD DS when a new domain controller is created, ensures that a domain controller is not created in a domain or forest that cannot support it. The wizard verifies whether the current domain or forest functional level falls within the highest and lowest fixed values for the operating system on the new domain controller. If the values are not compatible, the AD DS installation fails. For example, if the domain has the functional level Windows Server 2003 (domain level 2), an attempt to install AD DS on a Windows 2000–based server to create an additional domain controller fails because the highest and lowest fixed value supported by Windows 2000 is 0. Similarly, if the forest has a functional level of Windows Server 2003 (forest level 2), an attempt to install AD DS to create a new domain on a Windows 2000–based server fails.

AD DS likewise prevents restoring a domain controller into a domain or forest that has an incompatible functional level.

When a domain controller starts, an additional test is performed to verify whether the highest fixed value of the operating system matches the value of the **msDSBehaviorVersion** attribute on the NTDS Settings object of the domain controller.

**Checks for Raising Domain and Forest Functional Levels**
AD DS determines whether the functional level can be raised, as follows:

- The ability to raise the domain functional level depends on the functional levels of all domain controllers in the domain. In general, the functional level of the domain can be no higher than the lowest functional level of any domain controller in the domain.
- The ability to raise the forest functional level depends on the functional level of all domain controllers in the forest. In general, the forest level can be no higher than the lowest functional level of any domain controller in the forest.

### Note

- The ability to raise domain and forest functional levels is provided in the administrative tools Active Directory Users and Computers (raise domain functional level) and Active Directory Domains and Trusts (raise domain and forest functional levels).

**Verifying the Minimum Required Functional Level of Domain Controllers**
When an administrator attempts to raise the functional level of the forest or domain, AD DS checks the value of **msDSBehaviorVersion** on the NTDS Settings object of every domain controller within the scope (forest or domain) to determine if that value is equal to or greater than the desired functional level. If it is not, AD DS prevents raising the functional level.

For example, if the request is received to raise the domain level to 2 (Windows Server 2003), AD DS checks that the value of **msDSBehaviorVersion** on the NTDS Settings object on all domain controllers in the domain is 2 or higher before making the change. If all domain controllers are running Windows Server 2003, the value of **msDSBehaviorVersion** on the NTDS Settings object of each domain controller is 2, and thus AD DS allows raising the functional level of the domain to Windows Server 2003.

If any domain controller is running Windows 2000, the fact that **msDSBehaviorVersion** attribute is missing on the NTDS Settings object indicates Windows 2000 and an implied value of 0, and AD DS will prevent raising the domain functional level to Windows Server 2003. However, Windows NT 4.0–based domain controllers do not have an NTDS Settings object, so AD DS cannot check these computers directly for functional level. For this reason, AD DS performs an additional check to rule out the existence of Windows NT 4.0–based domain controllers.

**Ruling out Windows NT 4.0–Based Domain Controllers When Raising the Forest Level to Windows Server 2003**

When a member of the Enterprise Admins group attempts to raise the forest functional level to Windows Server 2003 (level 2), AD DS checks for the absence of Windows NT 4.0–based domain controllers by checking the **ntMixedDomain** attribute of the domain object. The **ntMixedDomain** attribute identifies a Windows 2000 domain as being in mixed mode or native mode.

The **ntMixedDomain** attribute on the domain object indicates the status of mixed or native domain mode as follows:

- **ntMixedDomain** value of 1 indicates mixed mode (Windows NT 4.0–based domain controllers can be present in the domain).
- **ntMixedDomain** value of 0 indicates native mode (no Windows NT 4.0–based domain controllers can be present in the domain).

To rule out the existence of Windows NT 4.0–based domain controllers when raising the forest functional level to Windows Server 2003 (level 2), AD DS verifies that the **ntMixedDomain** attribute has a value of 0 for every domain in the forest.

**Ruling out Windows NT 4.0 Domain Controllers When Raising the Domain Level to Windows Server 2003**

Unlike the behavior when raising the forest functional level, AD DS does not check for the absence of Windows NT 4.0–based domain controllers at the time an administrator raises the domain functional level. Just as in Windows 2000, a domain administrator is responsible for ensuring that no Windows NT 4.0–based domain controllers are running in the domain before raising the domain functional level from Windows 2000 mixed to Windows 2000 native or Windows Server 2003.

**Note**

- If an administrator changes the mode of a Windows 2000 domain from mixed to native mode or raises the domain functional level from Windows 2000 mixed to Windows 2000 native (or Windows Server 2003) when Windows NT 4.0–based domain controllers are present in the domain, the remaining Windows NT 4.0–based domain controllers no longer participate in the Active Directory forest. Therefore, the domain administrator must be absolutely sure that no Windows NT 4.0–based domain controllers are operating in the domain before raising the domain functional level to Windows 2000 native or Windows Server 2003.

When the domain functional level is raised to Windows Server 2003, AD DS not only raises the domain functional level to 2, but also changes the **ntMixedDomain** value from 1 to 0. When the domain functional level is raised to Windows 2000 native, AD DS only changes the **ntMixedDomain** value from 1 to 0.

**Checks for Raising the Functional Level to Windows Server 2003 Interim (Level 1)**

When attempting to raise the forest functional level to Windows Server 2003 interim, AD DS verifies that all domain controllers in the forest are at level 1 or higher (by checking the **msDSBehaviorVersion** attribute on the NTDS Settings object for each domain controller in the forest, to ensure that it is at level 1 or higher). However, it does not verify that all domains have a Windows 2000 native functional level (the **ntMixedDomain** attribute value of the domain object is 0).

The absence of this check allows Windows NT 4.0–based domain controllers to coexist with Windows Server 2003–based domain controllers, but it precludes future introduction of Windows 2000–based domain controllers into the domain or forest. This is because the highest and lowest functional level supported by Windows 2000 is 0 (which is less than 1), and when you attempt to install AD DS on a Windows 2000 server, the Active Directory Domain Services Installation Wizard will determine that the forest is at level 1 and it will prevent the installation.

When the forest functional level is raised to Windows Server 2003 interim, AD DS raises the forest functional level to 1. Consequently, every domain also raises its functional level (the domain functional level) to 1 (and therefore every domain is also at Windows Server 2003 interim level), but leaves the **ntMixedDomain** value at 1, thus allowing Windows NT 4.0-based domain controllers to coexist with Windows Server 2003–based domain controllers.

There is no administrative user interface (UI) for raising the forest and domain functional levels to Windows Server 2003 interim. An LDAP application, such as ADSI Edit or Ldp in Windows Support Tools, can be used, or a script can be used. In either case, the value of the **msDSBehaviorVersion** attribute should be modified to 1.

**Related Information**

The following resources contain additional information that is relevant to this section:

- How DNS Support for Active Directory Works
- "Support Tools Help" in the Tools and Settings Collection for information about ADSI Edit

- “Support Tools Help” in the Tools and Settings Collection for information about Ldp

# Active Directory Functional Levels Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Tools for Managing Active Directory Functional Levels
- Network Ports Used to Raise Active Directory Functional Levels
- Related Information


| Note |
| --- |
| In Windows 2000 Server and Windows Server 2003, the directory service is named Active Directory. In Windows Server 2008 and Windows Server 2008 R2, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

This section contains information about the tools that are associated with Active Directory functional levels.

**Tools for Managing Active Directory Functional Levels**
The following tools are associated with Active Directory functional levels.

**Domain.msc: Active Directory Domains and Trusts**

**Category**
An Active Directory Administrative Tools Microsoft Management Console (MMC) snap-in that is automatically installed on all domain controllers.

**Version compatibility**
This tool is compatible with domain controllers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. This tool can also be run on computers running Windows XP Professional. Updated versions of this tool can also be run on later versions of Windows. For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

Active Directory Domains and Trusts provides a graphical user interface (GUI) that can be used to manage Active Directory forests, domains, and trusts. Specific to functional levels, this tool can be used for the following:

- To view the current domain functional level by viewing the properties of the domain object.
- To view the current forest functional level by viewing the properties of the **Active Directory Domains and Trusts** node.
- To raise a domain functional level.
- To raise a forest functional level.

**Dsa.msc: Active Directory Users and Computers**

**Category**
An Active Directory Administrative Tools MMC snap-in that is automatically installed on all domain controllers.

**Version compatibility**
This tool is compatible with domain controllers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. This tool can also be run on computers running Windows XP Professional. Updated versions of this tool can also be run on later versions of Windows. For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?

LinkID=177813).

Active Directory Users and Computers provides a GUI that can be used to manage users and computers in Active Directory domains. Specific to functional levels, this tool can be used for the following:

- To view the current domain functional level by viewing the properties of the domain object.
- To raise a domain functional level.

In addition, LDAP Query can be used in this tool for the following:

- To identify domain controllers running Windows NT 4.0.
- To connect to a domain.

**Adsiedit.exe: ADSI Edit**

**Category**
This tool is included with Support Tools for Windows Server 2003, and it is automatically installed on domain controllers that run Windows Server 2008 or Windows Server 2008 R2.

**Version Compatibility**

| Can be run from | Can be run against |
|---|---|
| Computers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server |

Updated versions of this tool can also be run on later versions of Windows. For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

ADSI Edit is an MMC tool that uses Active Directory Service Interfaces (ADSI), which ultimately uses the Lightweight Directory Access Protocol (LDAP) protocol. This tool can be used to view and modify directory objects in the Active Directory database. Specific to functional levels, this tool can be used to edit the value of the **msDS-Behavior-Version** attribute of the Partitions container to raise the forest functional level to Windows Server 2003 interim, which cannot be done by using the Active Directory MMC administrative tools.

The **msDS-Behavior-Version** attribute is set in the ADSI Edit console on the Partitions container object (class **crossRefContainer**) in the configuration directory partition (cn=partitions,cn=configuration,dc=*ForestRootDomainName*). Raising the forest functional level to Windows Server 2003 interim requires changing the default value of the attribute from 0 to a value of 1.

To find more information about ADSI Edit, see “Support Tools Help” in Tools and Settings Collection.

**Ldp.exe: Ldp**

**Category**
This tool is included with Support Tools for Windows Server 2003, and it is automatically installed on domain controllers that run Windows Server 2008 or Windows Server 2008 R2.

**Version Compatibility**

| Can be run from | Can be run against |
|---|---|
| Computers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br>• Windows XP Professional | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server |

Updated versions of this tool can also be run on later versions of Windows. For more information, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

Ldp is a GUI tool that can be used to perform LDAP operations (such as connect, bind, search, modify, add, or delete) against any LDAP-compatible directory, such as AD DS.

Specific to functional levels, Ldp can be used as an alternative to ADSI Edit to modify the value of **msDS-Behavior-Version** attribute of the Partitions container object (class **crossRefContainer**) in the configuration directory partition (cn=partitions,cn=configuration,dc=*ForestRootDomainName*) and raise the forest functional level to Windows Server 2003 interim, which cannot be done by using the Active Directory administrative MMC tools.

The **msDS-Behavior-Version** attribute is set in Ldp by using the **Replace** operation in the **Modify** dialog box to change the default value of the attribute from 0 to a value of 1, which raises the forest functional level to Windows Server 2003 interim.

To find more information about Ldp, see "Support Tools Help" in Tools and Settings Collection.

#### Dcpromo.exe: Active Directory Domain Services Installation Wizard

**Category**
A wizard that is included with Windows Server. It is available from the command line or from the Configure Your Server Wizard on any computer running Windows Server 2003 or from Server Manager on servers that run Windows Server 2008 or Windows Server 2008 R2.

**Version compatibility**
This tool is compatible with computers running all editions of Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2, except Web Edition.

The Active Directory Domain Services Installation Wizard provides a GUI for setting up a domain controller by installing AD DS and, optionally, DNS. Specific to functional levels, the wizard can be used on a Windows NT 4.0 PDC when you are upgrading it to Windows Server 2003 and forming a new forest to raise the forest functional level to Windows Server 2003 interim, if appropriate.

**Network Ports Used to Raise Active Directory Functional Levels**
The following table shows the network ports that are used to raise functional levels.

### Port Assignments for Raising Active Directory Functional Levels

| Service Name | TCP |
|---|---|
| LDAP | 389 |
| LDAP SSL | 636 |

In addition, replication to all domain controllers requires the ports that are used for replication. For more information about the ports that are used for replication, see Active Directory Replication Tools and Settings in this collection.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Data Store Tools and Settings

- “Support Tools Help” in the Tools and Settings Collection Collection

# Active Directory Installation Wizard Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Active Directory Installation Wizard Technical Reference**
The Active Directory Installation Wizard is used to install the Active Directory directory service on a computer running Microsoft Windows 2000 Server and Microsoft Windows Server 2003. The wizard provides a user interface for gathering information about your computing environment. It also checks prerequisites for installing Active Directory, such as available disk space and operating system version, to ensure that the server is capable of becoming a domain controller. If all prerequisites are met, the Active Directory Installation Wizard uses the information gathered to configure Active Directory.

This subject explains what the Active Directory Installation Wizard is and how it works for Windows Server 2003 and Windows 2000 Server. For information about using the Active Directory Domain Services Installation Wizard in Windows Server 2008 R2 and Windows Server 2008, see AD DS Installation and Removal Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkID=139657).

**In this subject**


- What Is the Active Directory Installation Wizard?
- How the Active Directory Installation Wizard Works

# What Is the Active Directory Installation Wizard?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**What Is the Active Directory Installation Wizard?**

**In this section**


- Active Directory Installation Wizard Scenarios
- Active Directory Installation Wizard Dependencies
- Related Information


The Active Directory Installation Wizard is used to install and configure Active Directory in Windows Server 2003 and Windows 2000 Server, and to remove Active Directory when a domain controller is decommissioned. Understanding the processes that occur when the Active Directory Installation Wizard runs can help you recognize and troubleshoot issues that might arise during the installation or removal of a domain controller.

| **Note** |
| --- |
| This topic explains what the Active Directory Installation Wizard is for Windows Server 2003 and Windows 2000 Server. For information about using the Active Directory Domain Services Installation Wizard in Windows Server 2008 R2 and Windows Server 2008, see AD DS Installation and Removal Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkID=139657). |

Active Directory is not installed by default on a computer running Windows Server 2003. When you install Windows Server 2003, the computer assumes the role of a standalone server or a member server that is a part of a domain. For the computer to become a domain controller, you must use the Active Directory Installation Wizard to install and configure Active Directory.

**Active Directory Installation Wizard Scenarios**

You can create two types of domain controllers by using the Active Directory Installation Wizard:

- Domain controller for a new domain
- Additional domain controller for an existing domain

When you create a domain controller for a new domain, the domain can be one of the following types:

- Domain in a new forest.

  Select this domain type if you are creating the first domain in your organization or if you want the new domain to be independent of your existing forests. This first domain is the forest root domain.

- Child domain in an existing domain tree.

  Select this domain type if you want the new domain to be a child of an existing domain.

- Additional domain tree in an existing forest.

  Select this domain type if you want to create a domain tree that is separate from any existing domain trees.

  For more information about the Active Directory logical structure, see "Domains and Forests Technical Reference."

**Active Directory Installation Wizard Dependencies**

Active Directory installation has the following dependencies.

**Network Configuration**
Unstable or problematic network configurations are either recognized during the installation of Active Directory and prohibit you from continuing, or they create complications in your environment after the installation is complete. Before installing Active Directory, ensure that your network is properly configured.

**DNS Configuration**
Before installing Active Directory, plan your Domain Name System (DNS) namespace to support Active Directory and determine which servers will host the DNS zones. Consider the following DNS configuration options based on your environment:

1. Allow the Active Directory Installation Wizard to configure the DNS server.

   Consider this option if you do not have DNS configured in your environment or you have a DNS server hosted by an Internet Service Provider (ISP). During the Active Directory installation, the wizard installs and configures DNS.

   If you do not have an existing DNS configuration, the user interface portion of the Active Directory Installation Wizard gives you the opportunity to either configure the DNS client or install the DNS Server service. If the server on which Active Directory is being installed is pointing to an existing DNS server that is hosted by an ISP, that DNS server is added as a forwarder.

2. Point to an existing DNS server.

   Configure the server on which you are installing Active Directory to point to an existing DNS server in your environment before running the Active Directory Installation Wizard.

3. Configure DNS yourself.

   Only configure DNS yourself if you have planned your DNS namespace to support Active Directory and understand the process for configuring DNS. If you choose to configure DNS, use the following guidelines:

   - Configure the server with a static IP address.
   - Configure the DNS server settings on the TCP/IP properties page to use the local computer as the preferred DNS server. All DNS server addresses should belong to DNS servers in the forest. Do not configure DNS servers in your forest to point to ISP DNS servers. To view the current IP configuration and verify that the local computer is pointing to itself for its preferred DNS server, type **ipconfig /all** at the command prompt.
   - Configure the DNS Forward Lookup zone to accept dynamic updates. If dynamic updates are not allowed, all DNS record registration must be completed manually.
   - Configure the DNS server to forward DNS requests to the appropriate ISP or corporate DNS servers to ensure that names outside of the DNS namespace can be resolved.
   - If the **Enable Forwarders** check box is unavailable, the DNS server is attempting to host a root zone, identified by a zone named only with a period, or dot (.). This root zone is created by default only on Windows 2000–based DNS servers. It is not created by default on DNS servers that are running Windows Server 2003. You must delete the root zone to enable the DNS server to forward DNS requests.

**Disk Space Recommendations**
At minimum, a domain controller requires available free disk space for the Active Directory database, Active Directory log files, SYSVOL, and the operating system. Use the following guidelines to determine how much disk space to allot for your Active Directory installation:

- On the drive that will contain the Active Directory database, Ntds.dit, provide 0.4 gigabytes (GB) of storage for each 1,000 users. For example, for a forest with two domains (domain A, domain B), with 10,000 and 5,000 users respectively, provide a minimum of 4 GB of disk space for each domain controller that hosts domain A and a minimum of 2 GB of disk space for each domain controller that hosts domain B.

- On the drive containing the Active Directory log files, provide at least 500 MB of available space.
- On the drive containing the SYSVOL shared folder, provide at least 500 MB of available space.
- On the drive containing the Windows Server 2003 operating system files, to run setup, provide at least 1.25 GB to 2 GB of available space.

To prevent single disk failures, many organizations use a redundant array of independent disks (RAID). For domain controllers that are accessed by fewer than 1,000 users, you can locate all four components on a single RAID 1 array. For domain controllers that are accessed by more than 1,000 users, place the log files on one RAID array and keep the SYSVOL shared folder and the database together on a separate RAID array.

The above figures are recommendations only and should be tested in a lab environment before deploying Active Directory and monitored closely in your production environment. It is always a best practice to allow for future growth of objects in Active Directory. For more information about object growth and database capacity in Active Directory, see "Data Store Technical Reference."

**Client Operations**
Windows Server 2003–based domain controllers implement security settings that require clients and other servers to communicate with those domain controllers in a more secure way. Windows 95–based clients and clients running Windows NT 4.0 with service pack 3 or earlier cannot meet the logon or resource access requirements of a Windows Server 2003–based domain controller. Prior to installing Active Directory, take an inventory of your client operating systems and take the necessary actions to ensure continued access to domain controllers in your environment.

For more information about new security settings that are implemented by domain controllers running Windows Server 2003, see "How the Active Directory Installation Wizard Works."

**Administrative Rights**
You must be an Administrator on the local computer to run the Active Directory Installation Wizard if you are creating the first domain controller in a forest. If you are creating a child domain or a new tree domain in an existing forest, you must provide the credentials of a member of the Enterprise Admins group.

**Availability of the Domain Naming Master**
If you are creating an additional domain in an existing forest or removing a domain from a forest, the Active Directory Installation Wizard must access the domain naming operations master or it generates an error. Verify that the domain controller holding the role of domain naming master for the forest is online. If the domain naming master has recently been offline, ensure that inbound replication of the configuration partition has completed since the computer was restarted.

**Forest and Domain Preparation**
If you are upgrading an existing Windows 2000 domain to Windows Server 2003 or adding a Windows Server 2003–based domain controller to an existing Windows 2000 domain, you must prepare the forest and each domain in the forest by using the Active Directory Preparation Tool (ADPrep.exe). The tool performs the following tasks:

- Tightens security on resources that use the Everyone group to grant access by:
    - Improving default security descriptors.
    - Changing group memberships; for example, the Anonymous Logon group is no longer a member of the Everyone group.
    - Improving security settings by adding access control entries (ACE) to the following pre-existing, non-schema objects:

      CN=AIA,CN=Public Key Services,CN=Services,CN=Configuration,DC=X container

      CN=CDP,CN=Public Key Services,CN=Services,CN=Configuration,DC=X container

      CN=Configuration,DC=*ForestRootDomain* container
- Creates new objects used by individual applications.

- Creates new containers that can be used to verify that the preparation was successful.
- Updates the Active Directory schema.

  The Active Directory Preparation tool merges your current schema with new schema information provided by the tool. Previous schema modifications in your environment are not affected by this operation.

Adprep includes the following two switches: **Adprep /forestprep** to prepare the forest, and **Adprep /domainprep** to prepare each domain.

Run **Adprep /forestprep** on the domain controller holding the schema operations master role for the forest before you install Active Directory on a Windows Server 2003-based member server or standalone server to create an additional domain controller in an existing forest during Active Directory installation. After **Adprep /forestprep** runs successfully, you can run **Adprep /domainprep**. However, before running **Adprep /domainprep** in an existing domain, changes made to the forest after running **Adprep /forestprep** must have replicated to the infrastructure operations master in that domain.

You must successfully run **Adprep /domainprep** in each domain before you join a Windows Server 2003–based member server or standalone server to a domain as an additional domain controller.

**Adprep /domainprep** must run successfully in a domain before you can upgrade existing Windows 2000–based domain controllers to Windows Server 2003.

**Domain Controller Configuration**

Before installing Active Directory, consider the status of the computer that you will configure as a domain controller. If the computer is running Windows NT 4.0 or Windows 2000, determine whether or not the computer can be upgraded to Windows Server 2003. If you are upgrading, determine if your environment should be prepared for the upgrade and whether it will be necessary to complete the Active Directory Installation Wizard. These decision points are outlined in the following table.

**Decision Points for Domain Controller Configuration**

| Operating System | Upgradeable to Windows Server 2003? | Use the Active Directory Preparation Tool? | Install Active Directory? |
|---|---|---|---|
| Windows NT 4.0 Workstation | NO | N/A | N/A |
| Windows NT 4.0 Member Server | YES | It is not necessary to run ADPrep. | Only if the computer is to become a domain controller. |
| Windows NT 4.0<br><br>PDC or BDC | YES | It is not necessary to run ADPrep. | YES - Active Directory must be installed. |
| Windows 2000 Professional | NO | N/A | N/A |
| Windows 2000 Member Server (SP4) | YES | YES - The upgrade cannot proceed until you run ADPrep. | Only if the computer is to become a domain controller. |
| Windows 2000 domain controller | YES | YES - The upgrade cannot proceed until you run ADPrep. | NO - Active Directory is installed. |

**Windows Time Service Configuration**

After installing Active Directory on the first domain controller in a new forest, it is important to correctly configure the Windows Time service. The Windows Time service provides time synchronization to peers and clients, ensuring that time is consistent throughout an enterprise. If the Windows Time service is not properly configured, Kerberos authentication will fail, adversely affecting Active Directory replication, user logon, and shared folder access.

By default, the first domain controller deployed holds the PDC emulator operations master role, and should be set to synchronize time with a valid Network Time Protocol (NTP) source. If no source is configured, the service will log a message to the event log, and use the local clock when providing time to clients. Although internet NTP sources are valid for this configuration, we recommend that a dedicated hardware device, such as a GPS, or radio clock be employed in the interest of security.

**Related Information**

The following resources contain additional information that is relevant to this section.

- DNS Support for Active Directory Technical Reference
- Active Directory Functional Levels Technical Reference

# How the Active Directory Installation Wizard Works

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**How the Active Directory Installation Wizard Works**

**In this section**


- Active Directory Installation Wizard Architecture
- Active Directory Installation Methods
- Active Directory Installation Wizard Prerequisites
- Active Directory Installation Wizard User Interface
- Active Directory Installation Wizard Directory Configuration Process
- Active Directory Removal
- Related Information


This section details the architecture of the Active Directory Installation Wizard for Windows Server 2003 and Windows 2000 Server, methods for invoking the wizard, and the directory configuration process. This section also details the removal of Active Directory from a domain controller.

| **Note** |
| --- |
| This topic explains how the Active Directory Installation Wizard works for Windows Server 2003 and Windows 2000 Server. For information about using the Active Directory Domain Services Installation Wizard in Windows Server 2008 R2 and Windows Server 2008, see the AD DS Installation and Removal Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkID=139657). |

You can use the Active Directory Installation Wizard to install Active Directory on a computer, making that computer a domain controller. Before installing Active Directory to create the first domain controller in a forest, ensure that your network is online and functioning properly and that you have designed your Domain Name System (DNS) namespace and have a plan for installing and configuring DNS.

**Active Directory Installation Wizard Architecture**

The Active Directory Installation Wizard (Dcpromo.exe) architecture verifies that certain prerequisites are met prior to an Active Directory installation, gathers information specific to your operating environment, and configures Active Directory. It works differently depending on the installation scenario you choose and the installation method you are using. The following figures illustrate how the Active Directory Installation Wizard works when creating the first domain controller in a forest, an additional domain controller in an existing domain, and the installation process when installing Active Directory from media, which is a new feature in Windows Server 2003 Active Directory.

### Active Directory Installation on the First Domain Controller in a Forest




When you install Active Directory on a computer that is the first domain controller in a forest, the Active Directory Installation Wizard checks prerequisites and gathers information from you through the wizard user interface. It invokes the Active Directory installer (Ntdsetup.dll) through a remote procedure call (RPC) interface, exposed by the Lsass.exe security subsystem, to install the first domain controller in the forest. As the installer works, directory information is written to the Ntds.dit file by Esent.dll. If you select the option to install DNS, the Active Directory Installation Wizard contacts the DNS Server service, which in turn accesses Sysocmgr.exe and Dns.exe, the components that install and start DNS. After the Active Directory installation is complete and the computer restarts, the DNS Server service configures the Active Directory–integrated DNS zones and sends the DNS configuration information back to Active Directory.

**Active Directory Installation on an Additional Domain Controller in an Existing Domain**




When you install an additional domain controller in an existing forest, the Active Directory Installation Wizard checks prerequisites, gathers information from you through the wizard user interface, and invokes the Active Directory installer (Ntdsetup.dll) through an RPC interface exposed by the Lsass.exe security subsystem. The installer contacts an existing domain controller. The existing domain controller replicates all domain, forest, and directory partition information back to the new domain controller. As the installer works, directory information is written to the Ntds.dit file by Esent.dll. If you select the option

to install DNS, the Active Directory Installation Wizard contacts the DNS Server service, which in turn accesses Sysocmgr.exe and Dns.exe, the components that install and start DNS. After the Active Directory installation is complete and the computer restarts, the DNS Server service configures the Active Directory–integrated DNS zones and sends the DNS configuration information back to Active Directory.

**Active Directory Installation from Media**




When you install Active Directory on an additional domain controller from backup media, the Active Directory Installation Wizard checks prerequisites, gathers information from you through the wizard user interface, and invokes the Active Directory installer (Ntdsetup.dll) through an RPC interface exposed by the Lsass.exe security subsystem. As the installer works, directory information is written to the Ntds.dit file by Esent.dll. When you install from backup media, all domain information is backed up from a domain controller in the same domain and restored to the local computer. The installer accesses the restored files to configure the directory service. A domain controller is contacted over the network to replicate any changes that have been made since the backup media was created. The result is less information being sent over the network and a faster Active Directory installation. If you select the option to install DNS, the Active Directory Installation Wizard contacts the DNS Server service, which in turn accesses Sysocmgr.exe and Dns.exe, the components that install and start DNS. After the Active Directory installation is complete and the computer restarts, the DNS Server service configures the Active Directory–integrated DNS zones and sends the DNS configuration information back to Active Directory.

The architectural components in each of the preceding diagrams are explained in the following table.

**Active Directory Installation Wizard Architectural Components**

| Component | Description |
|---|---|
| Dcpromo.exe | The Active Directory Installation Wizard. |
| Lsass.exe | The LSA security subsystem, which provides services to both the kernel mode and the user mode for validating access to objects, checking user privileges, and generating audit messages. |
| Ntdsetup.dll | The Active Directory installer. |
| Esent.dll | Extensible storage engine (ESE). Communicates directly with individual records in the directory data store on the basis of the object's relative distinguished name attribute. |
| Ntdsa.dll | The directory service module that supports the Windows Server 2003 replication protocol and Lightweight Directory Access Protocol (LDAP), and manages partitions of data. |

| | |
|---|---|
| Ntds.dit | Directory database or data store. Requires NTFS formatting and is manipulated only by the ESE database engine. Encryption of the Ntds.dit file or related log files, their parent folders, or storage drive is not recommended and not supported. |
| Schema.ini | Contains the information necessary for creating the default directory objects and the default security for the Ntds.dit, as well as the Active Directory display specifiers. |
| Dns.exe | The executable from which the DNS Server service is run. |
| Sysocmgr.exe | The System Optional Component Installation Manager interprets INF files and installs optional components such as the DNS Server service. |

**Active Directory Installation Methods**
You can install Active Directory using any of the following four methods:

- Install from the command line
- Install from the Configure Your Server Wizard
- Install unattended by using an answer file
- Install from media

**Installing from the Command Line**
You can initiate an Active Directory installation from the command line any time after Windows Server 2003 has been installed. To install Active Directory by using the command line, in the Run dialog box, type **dcpromo** and click **OK**.

**Installing from the Configure Your Server Wizard**
You can install Active Directory at any time by using the Configure Your Server Wizard. The Manage Your Server dialog box automatically appears when you restart the server after installing Windows Server 2003. To invoke the Configure Your Server Wizard, select **Add or remove a role**. Select the **Domain Controller (Active Directory)** role in the Configure Your Server Wizard, and then click **Next**. The Active Directory Installation Wizard starts.

**Installing Unattended by Using an Answer File**
You can run the Active Directory Installation Wizard without having to be present to answer the questions by using an answer file. An answer file is a text file that you populate with the parameters that the wizard needs to install Active Directory. You can use an answer file to install Windows Server 2003 and include the options necessary to subsequently install Active Directory. Alternatively, you can create an answer file that contains only the options necessary for installing Active Directory. These parameters include the domain controller type (additional domain controller for an existing domain or a new domain controller for a new domain), the configuration of the domain that is being created (new forest, new tree root, or new child) and Active Directory forest and domain functional levels.

The answer file options for installing Active Directory are in the DCInstall section of the Unattend.txt file. You can find instructions for creating an answer file for an Active Directory installation in the Deploy.cab file in the Support\Tools folder on the Windows Server 2003 product CD. Inside the Deploy.cab file, open Ref.chm to access the Unattend.txt file. In the left pane, expand **Unattend.txt**, and click **DCInstall**. In Windows 98, Windows 2000, and Windows Server 2003, use Windows Explorer to extract this document. In Windows 95 and earlier, or from MS-DOS, use the **Extract** command to open the file. You can rename the Unattend.txt file and save it in any location. The Unattend.txt file also contains sample answer files. After creating the answer file, append the **/answer** switch to the **dcpromo** command when installing Active Directory from the command line.

You can run this answer file after Windows Server 2003 setup is complete and after you have logged on to the system.

**Installing from Media**
When you install from media, you can pre-populate Active Directory with system state backup data from an existing Windows Server 2003–based domain controller that belongs to the same domain in which this server will become an additional domain controller. This backup can be stored on a local CD, DVD, or hard disk partition. Then, the system state backup must be restored locally on the server you are installing Active Directory on. When the data is restored, proceed with the Active Directory installation.

To install Active Directory from media, append the **/adv** switch to the **dcpromo** command when installing from the command line to enable install from media options in the Active Directory Installation Wizard.

When you select the option to copy domain information from backup files, you must identify the location of the files. If the domain controller from which you restore the system state data is a global catalog server, you have the option make this new domain controller a global catalog server. Once you have selected an option, the wizard proceeds with the installation.

The install from media feature in Windows Server 2003 reduces the time required to install directory information by reducing the amount of data that is replicated over the network. Installing from media is most beneficial in environments with very large domains or for installing new domain controllers that are connected by a slow network link. If the computer on which Active Directory is being installed is located in close proximity to another domain controller and the network connections are optimal, installing from media will only be slightly faster than copying the information over the network.

When you install Active Directory on Windows 2000 Server, the computer has to populate its entire directory service with directory information replicated from another domain controller, which requires time for the replication to occur across the network.

**Active Directory Installation Wizard Prerequisites**
The Active Directory Installation Wizard confirms several configuration and security parameters before it begins collecting information from you. The following prerequisites must be met before an Active Directory installation proceeds:

- An installation or removal operation of Active Directory must not already be in progress.
- Command line parameters must be correct if you are installing from media or using an answer file.
- You must be logged on to the local computer as a member of the Administrators group. If the credentials provided at logon are insufficient, you cannot continue.
- Certificate Services must not be installed on the computer. Changing the computer name or domain membership of a certification authority (CA) invalidates the certificates that are issued from that CA. Uninstall Certificate Services and reinstall after the Active Directory installation is complete.
- The computer must be running in normal mode. If the computer is running in safe mode, an error message advises you to restart in normal mode.
- If you have previously removed Active Directory from the computer, and you are attempting a new Active Directory installation, you must restart the computer before the new installation proceeds.
- The computer must be running Windows 2000 Server, Windows Server 2003 Standard Server, Windows Server 2003 Enterprise Server, or Windows Server 2003 Datacenter Server.
- At least one logical disk drive must be formatted with the NTFS file system. You can only install the SYSVOL shared folder on a drive that is formatted with NTFS and we also recommend installing the Active Directory database on an NTFS formatted drive.
- There must be sufficient free disk space on the computer. The Active Directory Installation Wizard requires at least 20 MB of free disk space on the computer or it does not proceed. For more information about the disk space requirements for installing Active Directory, see "What Is the Active Directory Installation Wizard?."
- If the computer's name has been changed, you must restart it after the name change for the wizard to proceed. The wizard also checks to see that the computer name conforms to DNS naming requirements.
- If Terminal Server is running on the computer, a warning indicates that installing Active Directory on a computer running Terminal Server changes the local security policy, allowing only administrators to log on to the computer. You can cancel the wizard and remove Terminal Server.

After all installation prerequisites are met, the Active Directory Installation Wizard continues with the user interface portion of the installation process.

**Active Directory Installation Wizard User Interface**
Use the user interface portion of the wizard to select the type of domain controller you want to create and to gather specific information about your environment for the proper configuration of that domain controller.

**Verifying Client Operating System Compatibility**
Domain controllers running Windows Server 2003 have improved security settings that might affect clients that are running older versions of Windows. Server message block (SMB) packet signing and secure channel signing are security policies that are enabled by default on domain controllers running Windows Server 2003.

Clients running Windows NT 4.0 with Service Pack 2 or earlier, and clients running Windows 95 without the Directory Service Client Pack, do not support SMB packet signing and cannot authenticate to a domain controller running Windows Server 2003.

Clients running Windows NT 4.0 with Service Pack 3 or earlier do not support secure channel signing. These clients cannot establish communications with a Windows Server 2003–based domain controller. To ensure successful communication, upgrade these clients to a later version of the Windows operating system or service pack. However, if you cannot upgrade your clients, you must disable secure channel signing on all domain controllers running Windows Server 2003 so that the traffic passing through the secure channel does not need to be signed or encrypted.

### Note

- Secure channel signing does not affect Windows 95 clients.

**Choosing Additional Domain Controller or Member Server**
When you upgrade an existing Windows NT 4.0 backup domain controller (BDC) to Windows Server 2003, the Active Directory Installation Wizard executes immediately when the operating system upgrade is complete and gives you the following options:

- Install Active Directory, creating an additional domain controller in the domain.
- Convert the BDC to a member server joined to the domain.

### Note

- You must upgrade and install Active Directory on the Windows NT 4.0 primary domain controller (PDC) in a domain before you can install Active Directory on a Windows NT 4.0 BDC.

**Copying Domain Information**
If you are creating an additional domain controller in an existing domain and you are installing Active Directory by using the install from media installation method, the Copying Domain Information page appears in the wizard. This wizard page, rather than the /adv switch used with the dcpromo command, determines how the Active Directory installation is performed. This wizard page presents you with the following options for copying domain information:

- Over the network from a domain controller
- From these restored backup files

The **Over the network from a domain controller** option simply performs the standard Active Directory installation without using any backup media. If you want to install from media, you must select **From these restored backup files**. Copying the restored information from local media is the recommended option and will result in a faster Active Directory installation.

**Providing the System Key**
If you are installing from media and system key protection was enabled on the domain controller from which you created the system state backup, the Active Directory Installation Wizard looks for the system key in the restored Active Directory files and prompts you accordingly:

- If the system key was stored on the domain controller's hard disk, the system key is stored in the backup files and available to the wizard and you are not prompted for it.
- If a system key password was set on the domain controller, then the same password is set on the system key in the backup files and the wizard prompts you for the password.
- If the system key for the domain controller was stored on a floppy disk, then the system key is not present in the backup files. The wizard prompts you for the floppy disk that contains the system key.

**Configuring the Global Catalog**
If you are installing from media by using the system state backup of a global catalog server, you will be asked if you want to configure the new domain controller as a global catalog server as well.

You can designate a domain controller as a global catalog server at any time by using the Active Directory Sites and Services console. Therefore, you do not need to specify whether a domain controller will be a global catalog server during the Active Directory installation. However, if a global catalog server is required in the site in which you are installing the domain controller, designating the domain controller as a global catalog server during the Active Directory installation eliminates the additional configuration step.

**Contacting an Active Directory Domain Controller**
If you upgrade a Windows NT 4.0 BDC to Windows Server 2003 and then install Active Directory, the domain controller must be able to contact an Active Directory domain controller for the domain of which it was previously a member. If this is the first Windows NT 4.0 BDC being upgraded, the PDC must have already been upgraded and the BDC must be able to contact it. If this is an additional BDC, it must be able to contact either the PDC or another domain controller from its domain. If the BDC successfully contacts an Active Directory domain controller for the domain, the wizard proceeds without any notification.

If the domain controller is not able to successfully contact an Active Directory domain controller, but is able to contact a non-directory service (pre-Windows 2000) domain controller, it might be for one the following reasons:

- The primary domain controller (PDC) for the domain of which the BDC was a member has not been upgraded to Windows Server 2003.

- Network configuration or connectivity problems are preventing a domain controller from being contacted.

If the BDC cannot contact the PDC for any of the above reasons, the Cannot Contact an Active Directory Domain Controller wizard page will advise you of the issue. You can either stop to investigate and correct the problem before continuing, or choose the option to make the BDC a member server in the Active Directory domain.

**Providing Network Credentials**
To create a new domain controller in an existing forest, the Active Directory Installation Wizard requires that you provide valid network credentials in the form of a user name, password, and domain name. The wizard accepts the user name in the form of a logon user identifier – for example, johnSmith – or in the form of domain\user name. We do not recommend using a user principal name (UPN) formed as johnSmith@*domainName*.

- To create an additional domain controller for an existing domain, you must provide credentials for a member of either the Enterprise Admins group or the Domain Admins group for the domain to which you are adding the additional domain controller.

- To create a child domain in an existing domain tree, you must provide credentials for either a member of the Enterprise Admins group or an account that has been delegated the authority to create the new child domain by a member of the Enterprise Admins group.

- To create a new tree-root domain, you must you must provide credentials for either a member of the Enterprise Admins group or an account that has been delegated the authority to create the new tree-root domain by a member of the Enterprise Admins group.

### Note

- Network credentials are not required to create a new forest. The Active Directory Installation Wizard skips this step and immediately prompts you to enter a valid DNS name. All other configurations require you to provide network credentials before entering a valid DNS name.

**Providing a New Domain Name**
To create a domain controller in a new domain, you must provide a DNS name for the new domain. Acceptable naming conventions for domain names include the letters A through Z, numerals 0 through 9, and a hyphen (-). A dot (.) separates the discrete parts of a domain name, commonly known as labels. Each domain label cannot be longer than 63 bytes. Fully qualified domain names cannot:

- Have a single label.

Although Dcpromo.exe in Windows Server 2008 and Windows Server 2003 allows you to create a single-label DNS domain name, you should not use a single-label DNS name for a domain for several reasons. In Windows Server 2008 R2, Dcpromo.exe does not allow you to create a single-label DNS name for a domain. For more information, see http://go.microsoft.com/fwlink/?LinkId=92467.

| **Note** |
| --- |
| DNS names that do not contain a dot (.) cannot be registered on the internet and require additional configuration to be dynamically registered by DNS clients. |

- Exceed 64 characters.
- Contain embedded spaces.
- Be solely numeric.

To create a new child domain, you must provide the full DNS name of a parent domain. The wizard verifies that the parent domain exists. If the wizard does not recognize the DNS name of the parent domain, an error message states that the specified domain could not be contacted and that you should verify that the DNS name is typed correctly. If the DNS name of the parent domain is correct, there might be a problem with your DNS configuration.

The Active Directory Installation Wizard uses the domain controller Locator to verify that the DNS name is unique in the forest.

**Generating a NetBIOS Domain Name**

The Active Directory Installation Wizard generates a NetBIOS name from a valid DNS domain name, except when upgrading a Windows NT 4.0 PDC. When you upgrade a PDC running Windows NT 4.0, the NetBIOS name by default is the previous NetBIOS name of the domain and you cannot select a new name.

NetBIOS-created names have a 15-character limit. If your existing Windows NT 4.0 domain name is 15 characters in length, and your new domain name is more than 15 characters, the NetBIOS name truncates at the 15-character limit. If your domain name is truncated this way, a new NetBIOS name is created that is identical to your previous Windows NT 4.0 domain name, and domain creation fails. Domain creation also fails if your previous NetBIOS name contains embedded spaces; however, the Active Directory Installation Wizard ignores leading and trailing spaces.

In all other installation scenarios, the Active Directory Installation Wizard prompts you to either change or accept the NetBIOS name that is derived from the domain name. The wizard then uses the domain controller Locator to verify that the NetBIOS name that you have accepted is unique on the network.

**Using DNS Registration Diagnostics**

If you are installing Active Directory to create the first domain controller in a new forest, DNS must be installed on your computer. If you are installing an additional domain controller in an existing domain, DNS must be configured properly on your network. If DNS is not installed on the computer, the wizard prompts you to configure the DNS client and does not proceed until the configuration is complete. After the wizard verifies that the DNS client is installed, it proceeds with the Active Directory installation. During the Active Directory installation, DNS diagnostic tests verify that your DNS configuration supports Active Directory. The tests validate the full DNS name, verify that the appropriate DNS SRV resource records exist, and report the most common DNS misconfiguration scenarios.

For more information about DNS SRV resource records, see "How DNS Support for Active Directory Works."

When you install Active Directory on the first domain controller in a new domain, the Active Directory Installation Wizard performs a DNS query for the DNS name of the Active Directory domain specified in the wizard. The DNS query ensures that network hosts and services can locate the domain controller hosting the Active Directory domain. The wizard queries the DNS servers configured for its network connection for the start of authority (SOA) resource record for the zone that contains the DNS name of the Active Directory domain. After the SOA resource record is received and the name of the primary server is extracted from the response, the new domain controller queries DNS for the address (A) resource record for the primary DNS server name. After the new domain controller determines the IP address of the primary DNS server, it sends a dynamic update request to determine whether or not the DNS server exists and is capable of accepting DNS dynamic updates.

When you install Active Directory on an additional domain controller in an existing domain, a DNS server running on the network is assumed and a search for an authoritative DNS server is not performed.

The results of these diagnostic tests are displayed on the DNS Registration Diagnostics page of the Active Directory Installation Wizard. If the wizard detects a problem with the DNS configuration, it provides details about the problem. You can use the

information provided on this page to repair the problem and run the wizard again.

If DNS is not installed on the computer, you are given the option to install it now or after the Active Directory installation has completed. If you choose to install DNS during the Active Directory installation, a DNS server is installed locally and configured with a primary DNS zone that matches the name of the new Active Directory domain. The wizard configures DNS to accept dynamic updates and to modify the computer's IP configuration so that the domain controller points to its local DNS server as the preferred DNS server.

**Verifying Active Directory Forest and Domain Functional Levels**
Windows Server 2003 forest and domain functional levels provide a system for safely enabling new Windows Server 2003 features. Before you can install Active Directory, the Active Directory Installation Wizard ensures that the domain or forest can support the domain controller.

For example, if the forest functional level is Windows Server 2003, you cannot create a new domain on a Windows 2000–based server. Similarly, if the domain functional level is Windows Server 2003, you cannot install Active Directory on a Windows 2000–based server to create an additional domain controller in an existing domain. The system also prevents restoring a domain controller with information from a domain controller that has an incompatible functional level.

When upgrading a PDC running Windows NT 4.0 and configuring it as the first domain controller in a new forest, the Active Directory Installation Wizard prompts you to raise the forest functional level to Windows Server 2003 interim, because it is assumed that the forest does not contain Windows 2000 domain controllers. The Windows Server 2003 interim forest functional level allows only Windows Server 2003–based and Windows NT 4.0–based domain controllers to co-exist in a forest.

After all Windows NT 4.0 PDCs and BDCs are upgraded to Windows Server 2003, the forest functional level can be raised to Windows Server 2003.

When you upgrade from Windows 2000 to Windows Server 2003, the Active Directory Installation Wizard does not prompt you to raise the forest functional level. After an upgrade or clean installation of a Windows Server 2003–based domain controller, the default domain functional is Windows 2000 mixed and the default forest functional level is Windows 2000. These functional levels allow for the existence of domain controllers running Windows Server 2003, Windows 2000 or Windows NT 4.0. Functional levels can be raised manually when conditions in the domain or forest are appropriate.

For more information about functional levels, see "Active Directory Functional Levels Technical Reference."

**Selecting a Location for Database and Log Folders**
The Active Directory Installation Wizard displays the default location for the Active Directory database (Ntds.dit), and log files and enables you to specify an alternative location or accept the location provided.

During Windows Server 2003 installation, the Ntds.dit Active Directory database template is placed in a default location, the %*systemroot*%\System32 directory. In this location, the Ntds.dit database template does not function as the directory database; it exists as a template that is used to create the user's Active Directory database files when Active Directory is installed. The Ntds.dit template enables you to install Active Directory without using the Windows Server 2003 product CD.

We recommend that you install the Ntds.dit database template on an NTFS volume. For optimal domain controller performance, select separate physical hard disks for the Ntds.dit database template and the Active Directory log files.

**Selecting a Location for the Shared System Volume**
The Active Directory Installation Wizard displays the default location for the system volume (SYSVOL) and enables you to specify an alternative location or accept the location provided. An NTFS volume is required for the SYSVOL shared folder.

For optimal performance on domain controllers that are accessed by more than 1,000 users, place the log files on one physical hard disk and keep the SYSVOL shared folder and the Ntds.dit database together on a separate physical hard disk.

**Assigning Permissions to User and Group Objects**
When you create a domain controller for a new domain, the Active Directory Installation Wizard prompts you to select default permissions for user and group objects. It provides you with the following two options:

- **Permissions compatible with pre-Windows 2000 Server operating systems**
- **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**

**Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems** is the default selection and adds the Everyone security identifier (SID) and the Anonymous Logon SID to the Pre-Windows 2000 Compatible Access local group. The Pre-Windows 2000 Compatible Access local group has access to the user and group object attributes that

existed in Windows NT 4.0 and that are required by certain applications to function with Active Directory.

If you include the Everyone SID in an ACL or group membership in Windows 2000 or earlier operating systems, you allow authenticated users, guest users, and anyone with an anonymous logon to gain access to many resources. In Windows Server 2003, the Everyone SID no longer allows access to anonymous users. This change was made to disallow inappropriate access to resources by anonymous users and to avoid granting accidental access to anonymous users.

**Note**

- The Everyone SID contains an access token for every user account in the forest, including the Guest account. The Anonymous Logon SID contains only an access token for anonymous users. Select the **Permissions Compatible with pre-Windows 2000 servers** option to populate the Pre-Windows 2000 Compatible Access group with both SIDs.

In Windows Server 2003, the Anonymous Logon SID was created for backward compatibility with any pre-Windows 2000 application requiring less strict permissions than those granted by Windows 2000–based and Windows Server 2003–based domain controllers. For example, if you have Windows NT 4.0–based Remote Access service servers that are running on Windows NT 3.*x* or Windows NT 4.0–based computers, or if these applications are running on Windows 2000 or Windows Server 2003 computers that are located in a Windows NT domain, choose the **Permissions compatible with pre-Windows 2000 servers** option to add the Anonymous Logon SID to the pre-Windows 2000 Compatible Access local group so that the user and group object attributes that existed in Windows NT 4.0 can be accessed.

**Choose Permissions compatible only with Windows 2000 or Windows Server 2003** only if all of your server-based applications are running on Windows 2000–based or Windows Server 2003–based servers that are members of Windows 2000 or Windows Server 2003 domains. This option prevents the Pre-Windows 2000 Compatible Access group from being populated. This group remains empty when high-security is required in pure Windows 2000 and Windows Server 2003 environments.

**Providing an Administrator Password**

When you upgrade a Windows NT 4.0 BDC and select the member server option to make the computer a member server in the new Active Directory domain, you are electing to place the computer in a new role as a member server in the new Active Directory domain. As part of this process, all users and groups that were previously created on the computer are deleted and all built-in groups are recreated. After this conversion process completes, the previous Windows NT 4.0 administrator account is no longer valid, and a new password is necessary to access the new administrator account. You must assign a new server administrator password that will have access to the newly created groups.

**Providing a Directory Services Restore Mode Administrative Password**

On every new domain controller, whether it has been upgraded from a Windows NT 4.0–based domain controller or it was made a domain controller after a clean installation of Windows Server 2003, the Active Directory Installation Wizard prompts you for an Administrator account password that is to be used for authenticating to the Security Accounts Manager (SAM) database when the computer is started in Directory Services Restore Mode.

Starting a domain controller in Directory Services Restore Mode causes the domain controller to temporarily operate as a stand-alone server. When operating in this mode, the SAM database uses a minimal set of user and group definitions stored in the registry.

**Summary**

The Active Directory Installation Wizard enables you to review and confirm the configuration options you selected and enables you to move back through the wizard and make corrections. After you verify your Active Directory configuration settings, click **Next** to complete the interactive portion of the Active Directory Installation Wizard and begin the directory service configuration.

**Active Directory Installation Wizard Directory Configuration Process**

After gathering all necessary information from the user, the Active Directory Installation Wizard begins the configuration of the directory.

**Creating the SYSVOL Shared Folder**

When you create a domain, the Active Directory Installation Wizard also creates the SYSVOL folder, which contains the shared system volume, SYSVOL. SYSVOL provides a default location for files that must be shared throughout an Active Directory domain. SYSVOL creation requires an NTFS volume. Active Directory does not function if the SYSVOL folder is not properly configured and shared. When it is created, the SYSVOL folder is shared automatically by the Net Logon service.

The SYSVOL folder also includes the NETLOGON shared folder. The NETLOGON shared folder includes system policies and user-based logon and logoff scripts for network clients that do not run Windows 2000 or Windows Server 2003. On servers that are upgraded from Windows NT 4.0, files that were located in the original NETLOGON shared folder

(\\%systemroot%\SYSTEM32\REPL\Import\Scripts) are moved to the \\<DomainName>\SYSVOL\<FullyQualifiedDomainName>\Scripts folder in the SYSVOL tree.

The SYSVOL folder also includes:

- File system junctions.
- User logon, logoff, startup, and shutdown scripts.
- Windows 2000 and Windows Server 2003 Group Policy objects.
- File Replication service (FRS) staging directories and files.

  FRS replicates the SYSVOL shared folder. After Active Directory is installed and the domain controller is restarted, FRS actually creates the system volume objects in the system volume directory and enables replication of the SYSVOL shared folder on the domain controller. The new domain controller advertises itself in DNS as a domain controller only after SYSVOL is replicated to all domain controllers in the forest.

**Creating Directory Partitions**

The Active Directory Installation Wizard copies the Ntds.dit Active Directory database template from its location in the *%systemroot%*\System32 directory to the destination specified in the wizard and configures the local server to host the directory service. This process includes the creation of the directory partitions and the default domain security principals.

The Ntds.dit Active Directory database template includes a default copy of the schema, and the schema.ini file, which specifies the default structure of the configuration and domain directory partitions. The Ntds.dit template is used in the following manner:

- When you install Active Directory on a computer that is going to be the root of a forest, the Active Directory Installation Wizard uses the default copy of the schema and the information in the schema.ini file to create the new Active Directory database.
- When you install the first domain controller in a new domain that is not the root of the forest, the wizard synchronizes the schema and configuration directory partitions with the schema and configuration directory partitions of the forest root. From the Ntds.dit template, only the information in the schema.ini file that specifies the structure of the domain directory partition is used.

The following directory partitions are created as default partitions on the first domain controller in a forest, and they are updated through replication on every subsequent domain controller that is created in the forest:

- The schema directory partition is created as cn=schema,cn=configuration,dc=*forestRootDomain*. The schema.ini file is used to create default directory objects, and display specifiers, which are the containers of objects that describe how your Active Directory Users and Computers administrative console appears. Schema.ini implements default security on the directory database.
- The configuration directory partition is created as cn=configuration,dc=*forestRootDomain*.
- The domain directory partition is created as dc=*domainName* and contains the security principals for the domain.

  If you install a new domain controller from a backup of an existing domain controller by using **dcpromo /adv**, the domain directory partition is restored from the backup media onto the new domain controller.
- Two DNS application directory partitions below the forest root domain are automatically created by the DNS Server service when the computer restarts after the Active Directory Installation Wizard has finished. One application directory partition is created for the forest, ForestDnsZones, and one for the domain, DomainDnsZones. You can use the DNS administrative tool or the dnscmd.exe command-line tool, located in the \Support\Tools directory on the Windows Server 2003 product CD, to use these application directory partitions for DNS zone storage.

  **Note**

- If you are installing an additional domain controller in an existing forest, the domain controller holding the domain naming operations master role must be online, available, and running Windows Server 2003 for these application directory partitions to be created. If the domain naming master is unavailable or is running Windows 2000, the DNS Server service will attempt to create the application directory partitions again at a later time.

**Replicating Directory Partitions**
When you create a new domain in an existing forest, the schema and configuration directory partitions are updated on the new domain controller through replication from a source domain controller and a new domain directory partition containing all of the default domain objects is created.

**Note**

- All data replicated during the installation of Active Directory is compressed. After you have completed the wizard and the domain controller is restarted, the Knowledge Consistency Checker (KCC) reconfigures compression according to site boundaries.

When you install an additional domain controller in an existing domain, all three directory partitions are updated through replication from a source domain controller. Because of the large amount of data that must be synchronized when you install an additional domain controller, you might opt to delay full replication until the computer is restarted. To stop the replication process, click the **Finish Replication Later** button when it appears. Clicking the **Finish Replication Later** button causes only critical information, such as the schema directory partition, the configuration directory partition, and some objects from the domain directory partition, to replicate. After these critical items replicate and the computer restarts, replication continues as part of the normal replication activity of the domain controller. The domain controller does not advertise itself as such until replication completes.

The computer on which you install Active Directory uses the domain controller Locator to find a domain controller in the parent domain (for a new child domain) or in its own domain (for an additional domain controller in an existing domain) to act as the source domain controller for replication. The computer queries the source domain controller for the distinguished names of the Configuration container and the Schema container by posting an LDAP query that is based on the NULL distinguished name, which retrieves the root attributes of the Active Directory tree. The source domain controller replicates the schema directory partition and configuration directory partition (in that order), referenced only by their distinguished names. After the directory partitions have been replicated to the computer on which you are installing Active Directory, the GUIDs of the containers are established from the replicated data, although the directory partitions continue to be referenced solely by the distinguished name string for the duration of the installation process.

Failure to fully replicate any of the directory partitions results in the failure to install Active Directory. To ensure complete synchronization, there is a critical point in the replication process beyond which the process cannot be terminated: Prior to replication of the attributes from the domain directory partition, you can cancel the installation process. After the domain directory partition is replicated, you cannot cancel the installation process.

**Creating the Domain**
Regardless of the type of domain that you create, the Active Directory Installation Wizard performs the following operations during the installation process:

- Sets the computer Domain Name System (DNS) root domain name to the name of the new domain by using this format:

  *<computerName>.<domainName><forestRootDomainName>*

- Determines whether the server is joined to a domain. If the computer is a member server of a domain and if you are logged on as an administrator, the wizard removes the computer from the domain and reuses the account. If you are not logged on as an administrator, the wizard alerts you that an administrator must remove the computer account for the server from the domain.

- Creates a computer account in the Domain Controllers container in the new domain. The account is added to the Domain Controllers global group in the Users container. This account allows the computer to authenticate to other domain controllers.

- Applies the password you have provided for the administrator account that will be used when you start the domain controller in Directory Services Restore Mode.

- Creates a cross-reference object in the Configuration container. When the configuration directory partition is replicated to the new domain controller, a cross-reference object is created on the domain naming master and is then replicated throughout the forest. This object is used by LDAP to locate resources in other domains.

- Modifies the Local Security Authority (LSA) membership policy to distinguish the computer as a domain controller.

- Removes the **Start** menu shortcut to the local security settings and adds two new shortcuts to the following Group Policy security setting nodes:

    - Domain security settings for all users and computers.

    - Security settings that are specifically targeted at domain controllers.

**Creating a Forest Root Domain**

The following operations occur when you create the forest root domain:

- Creation of the Schema container and the Configuration container.

- Assignment of the PDC emulator, RID master, domain naming master, schema master, and infrastructure master roles to the domain controller.

### Note

- Creating a new forest does not affect any existing domain and, therefore, does not use a source domain controller during the installation of Active Directory.

**Creating a New Child Domain**

The following operations occur when you create a child domain in an existing domain tree:

- Verification of the name that you provide as a valid child domain name.

- Location of a source domain controller in the parent domain.

- Synchronization of the system time of the child domain with the system time of the source domain controller.

- Creation of parent-child trust objects in the System folder on both the parent domain and the child domain. These objects (class **trustedDomain**) identify two-way transitive trust relationships between the child domain and the parent domain.

- Replication of the Schema container and the Configuration container from the source domain controller in the parent domain.

- Assignment of the PDC emulator, RID master, and infrastructure master roles to the domain controller.

**Creating a New Tree-Root Domain in an Existing Forest**

When you create a new tree-root domain, its creation depends on other domains in the forest. The Active Directory Installation Wizard creates various new accounts, trust relationships, and cross-reference objects to incorporate the new tree-root domain into the forest.

The following operations occur when you create a new tree-root domain in an existing forest:

- Location of a source domain controller in the forest root domain.

- Synchronization of domain system time with the system time of the source domain controller.

- Creation of a tree-root trust relationship between the tree root domain and the forest root domain, and creation of a trust object (class **trustedDomain**) in both domains. The tree-root trust relationship is two-way and transitive.

- Assignment of the PDC emulator, RID master, and infrastructure master roles to the domain controller.

**Creating an Additional Domain Controller in an Existing Domain**
To add an additional domain controller to an existing domain, install Active Directory on a Windows Server 2003-based member server. The same verification and configuration processes occur as during the creation of a new domain. There are no specific namespace or TCP/IP checks and there are no operations master roles assigned.

However, if any of the following operations fail, the installation of Active Directory cannot proceed.

- Joining the computer to the domain. If the additional domain controller is already a member of a domain prior to the Active Directory installation, the computer is removed from the previous domain and rejoined to the new domain for which it is a new domain controller. Changing the join state might change the computer name as well. If the Active Directory Installation Wizard changes the join state of the computer and the Active Directory installation fails elsewhere in the wizard, you must restart the computer before attempting the installation again.

- Forced synchronization from the source server to the RID master, which ensures that a relative identifier pool is quickly provided to the new domain controller. The RID master does not have to be available during the installation of Active Directory, but it must be available at some point after the installation to transfer relative identifiers to the new domain controller.

After these operations succeed, the replication process begins.

**Setting Services to Start Automatically**
During your Active Directory installation, the following services are configured to start automatically:

- The RPC Locator, which enables distributed applications to use the Microsoft RPC name service. The RPC Locator manages the RPC name service database. For more information about the RPC Locator, see the Software Development Kit (SDK) on MSDN.

- The Net Logon service, which runs the Locator algorithm. Net Logon also is responsible for creating a secure channel between clients and domain controllers during the logon process, registering service (SRV) resource records in DNS, and supporting the Windows NT 4.0 replication protocol (LMRepl).

- The Key Distribution Center (KDC) service, which runs on a physically secure server and maintains a database with account information for all security principals in its realm — the Kerberos v5 authentication protocol equivalent of a Windows Server 2003 domain.

- IsmServ (Intersite Messaging [ISM] service), which is used for mail-based replication between sites and for calculating the least expensive routes between sites. Active Directory includes support for replication between sites by using Simple Mail Transfer Protocol (SMTP) over the IP transport. SMTP support is provided by the SMTP service, which is a component of Internet Information Services (IIS). The set of transports that are used for communication between sites must be extensible; therefore, each transport is defined in a separate add-in dynamic link library (DLL). These add-in DLLs are loaded into the ISM service, which runs on all domain controllers that are candidates for performing communication between sites. The ISM service directs send requests and receive requests to the appropriate transport add-in DLLs, which then route the messages to the ISM service on the destination computer.

- TrkSvr (Distributed Link Tracking Server service), which runs on each domain controller in a domain. This service enables client applications to track linked documents that have been moved to a location on another NTFS volume in the same domain, in another domain, or in a workgroup. The Distributed Link Tracking Server service helps resolve shortcuts and OLE links to NTFS-resident files that have undergone a name change, a path change, or both.

- W32time (Windows Time service), which synchronizes clocks between clients and servers that run Windows Server 2003. Time synchronization is automatic. For more information about the Windows Time service, see "How the Windows Time Service Works."

**Setting Security**
During the installation of Active Directory, security is enabled on directory service and file replication directories for access

control, and actions allowed on domain objects are set through Group Policy.

**Access Control**
Default access control lists are configured on file and directory objects. Access control lists are also configured for the following registry keys and file system objects, including all child objects:

- HKEY_LOCAL_MACHINE\SOFTWARE
- HKEY_LOCAL_MACHINE\SYSTEM
- HKEY_USERS\DEFAULT
- PROGRAM FILES
- %WINDIR%

For more information about access control, see "How Security Descriptors and Access Control Lists Work."

**Group Policy**
Group Policy is replicated from only the first domain controller in a domain to all additional domain controllers. In the case of the first domain controller, the following security templates in the %Windir%\Inf directory are used to configure default Group Policy:

- DCFirst.inf is used to define the default Password, Lockout, and Kerberos Group Policy settings for the default Domain Group Policy object.
- DefltDC.inf is used to define the Audit and User Rights Group Policy settings for the default Domain Controllers Group Policy object.
- DCUp.inf is used to define Windows Server 2003–specific settings during the upgrade of a Windows NT 4.0–based domain controller.

### Note

- To verify that Group Policy has replicated and been properly configured and applied, look in the Application log in Event Viewer for event ID 1704. If event ID 1704 is present, policy propagation succeeded.

Domains and domain controllers have default security policies. The domain controller security policy has precedence over the domain security policy. For example, if you want to grant the Add Workstation to Domain privilege to a user, you modify the default domain controller security policy rather than the default domain security policy.

For more information about Group Policy, see "Core Group Policy Technical Reference."

**Security Accounts Manager (SAM) Database**
When a Windows NT 4.0 PDC is upgraded to Windows 2000 or Windows Server 2003, the Active Directory Installation Wizard initializes immediately at the end of setup. When Active Directory is installed, local accounts located in the Windows NT 4.0 registry-based SAM database are migrated to Active Directory as domain accounts. The existing SAM is deleted, and a new, smaller registry-based SAM is created and used for starting the new domain controller in Directory Services Restore Mode for system repair. In Directory Services Restore Mode, the operating system is running without Active Directory and all user validation occurs through the SAM database in the registry.

### Note

- When you upgrade a Windows NT 4.0 PDC to a Windows 2000–based or Windows Server 2003–based domain controller, the previous SAM database is deleted to prevent password attacks.

If Active Directory is removed from the server, the new SAM is available for local user and group accounts on the member server. The computer SID does not change during the installation or removal of Active Directory.

**Site Determination**

The Active Directory Installation Wizard determines the site to add the new domain controller. It first checks the existing site of the source domain controller to find the subnet of the computer that you are installing. If the subnet is not found, you can add the computer to an existing site. You can create a new site for this domain controller after the Active Directory installation is complete. You can then move the domain controller from the installation site to the new site.

**Note**

- To automatically place the domain controller in the correct site, all domain controllers in the forest must have the correct site and subnet information before installing Active Directory.

During site determination, the wizard attempts to use the domain controller Locator to find the site you specified or the site in which the computer is currently located. DNS checks the site name to test that it is a valid label. A valid DNS label must contain at least one non-numeric character to distinguish it from an IP address. If the domain controller Locator does not return a site for the computer (that is, the computer's subnet is not associated with a site), the site of the source domain controller is used as the site for the new domain controller.

When you are installing the first domain in a new forest, the default site, Default-First-Site-Name, is used.

**Note**

- When you are performing an unattended installation, you can specify a site with the SiteName parameter in the answer file. For more information about unattended installation of Active Directory, see "Active Directory Installation Methods" earlier in this subject.

After the correct site is determined, the wizard verifies that a site object exists in Active Directory for that site and that a server object exists for the additional domain controller. If a server object does not exist, one is created. If the server object exists, the associated NTDS Settings object is deleted and then recreated for the new domain controller. An NTDS Settings object is created for each domain controller in the forest. If the NTDS Settings object already exists, the wizard performs as if this domain controller is being reinstalled.

**Note**

- The NTDS Settings object is always created on a remote server (the source domain controller) for an additional domain controller. For a new domain, the NTDS Settings object is created on the computer that has the domain naming master role. During replication of the Configuration container, the NTDS Settings object is replicated to all domain controllers.

**Active Directory Removal**

The Active Directory Installation Wizard also removes Active Directory. When you start the wizard on a domain controller, the computer is identified as a server that contains Active Directory, and the wizard prompts you for the information that is required to remove Active Directory.

When you remove Active Directory, you remove all short cuts that were added when Active Directory was installed, and you restore the local policy shortcut. You also restore the shortcut on the **Administrative Tools** menu that provides access to the local security settings for the member server or for the standalone server.

If the Active Directory Installation Wizard recognizes that the computer is a global catalog server, it advises you so that you can make sure that other global catalog servers are accessible to users. The wizard then asks you to indicate whether the domain controller you are decommissioning is the last domain controller in the domain.

**Providing Administrative Credentials to Remove Active Directory**

To remove Active Directory, you must provide the following administrative credentials:

- To remove Active Directory from a domain controller that is the last domain controller in a child domain, you must provide credentials for a member of the Enterprise Admins group.

- To remove Active Directory from a domain controller that is the last domain controller in a tree-root domain, you must provide credentials for or be logged on as a member of the Enterprise Admins group.

- To remove Active Directory from a domain controller that is the last domain controller in the forest, you do not have to provide credentials. However, you must log on to the domain as an administrator or as a member of the Domain Admins group.

- To remove Active Directory from a domain controller that is not the last domain controller in the domain, you do not have to provide credentials. However, you must be logged on as a member of either the Domain Admins group or the Enterprise Admins group.

In many cases, removing Active Directory requires Enterprise Admins credentials. The authority to create or remove a domain can be delegated, in which case the credentials of the delegated user suffice.

**Removing Active Directory from an Additional Domain Controller or the Last Domain Controller**
When you remove Active Directory from either an additional domain controller or from the last domain controller in the domain, the Active Directory Installation Wizard performs the following operations:

- Replication of changes to the configuration directory partition and the schema directory partition. For an additional domain controller, it replicates changes to the configuration, schema, and domain directory partitions.

- Transfer of any operations master roles that the domain controller is holding to another domain controller. While it is possible to allow the Active Directory Installation Wizard to do this, it is not recommended that you rely on the wizard to perform this action. Controlling where operations masters are placed ensures that you can locate and administer each operations master role as necessary.

- Removal of the system volume objects from the directory database; the system volume objects from the NtFrs database; and deletion of the SYSVOL directory hierarchy. NtFrs requests that the Net Logon service remove the share from the system volume.

- Removal of the NTDS Settings object and cross-reference objects.

- Update of DNS to remove the Locator records. (When the NTDS Settings object is deleted, the Directory System Agent (DSA) notifies the Net Logon service, and the Net Logon service removes the records).

- Creation of the local SAM database in the same manner as during a new installation, including creation of the local administrator account and setting the password.

- Modification of the LSA membership policy to distinguish whether the computer is a standalone server or a member server.

- Stop of the Net Logon service and other services. The same services that were started during the installation of Active Directory procedure are stopped. Services that relate only to the directory service are configured to not start automatically.

**Removing Active Directory from an Additional Domain Controller**
When you remove an additional domain controller, The Active Directory Installation Wizard:

- Locates a source domain controller in the same domain where the additional domain controller account exists and replicates changes to it.

- Sets the computer account type to member server and moves the computer account for the additional server from the Domain Controllers container to the Computers container.

**Removing Active Directory from the Last Domain Controller**
When you remove the last domain controller in the domain, The Active Directory Installation Wizard:

- Verifies that no child domains exist.

- Locates a source domain controller in the parent domain and replicates changes to it.

- Removes Active Directory objects specific to this domain from the forest. The wizard contacts the domain naming master and removes the NTDS Settings and cross-reference objects.

- Removes trust objects on the parent server. The trustedDomain objects in the System folder are deleted.
- Places the server in a workgroup called "Workgroup."

If the NTDS Settings object is not removed successfully from Active Directory (for example, if a server fails during the removal of Active Directory), remove the object manually by using ntdsutil.exe.

**Removing Active Directory from a Domain Controller that Hosts an Application Directory Partition**
In Windows Server 2003, non-domain data that is not of global interest can be stored on and replicated between designated domain controllers located in different domains throughout the forest by using application directory partitions. For example, you can store data that is specific to a single application in an application directory partition.

When you are removing a domain controller on which an application directory partition is present, the Active Directory Installation Wizard blocks the removal of that domain controller if it determines that the domain controller hosts the last copy of the application directory partition in the forest. If the removal of Active Directory is blocked, a dialog box lists the application directory partitions existing on the computer. Click **Next** to remove the application directory partitions from the domain controller and continue with the removal of Active Directory.

If at least one other replica of the application directory partition is located in the domain, the wizard skips the dialog box mentioned above and proceeds with the removal of Active Directory. Any changes to the application directory partition are replicated to another replica before it is destroyed.

**Provide a New Administrator Password**
When Active Directory is removed from a computer, all users and groups that were created on the computer are deleted and all built-in groups are recreated. You must assign a new server administrator password that will have access to the newly created groups.

**Related Information**
The following resources contain additional information that is relevant to this section.

- DNS Support for Active Directory Technical Reference
- Active Directory Functional Levels Technical Reference
- Core Group Policy Technical Reference

# Domain Rename Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Domain Rename Technical Reference**
Microsoft Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003 operating systems provide the capabilities to rename domains in an Active Directory Domain Services (AD DS) directory service forest after the forest structure is in place. These domain rename capabilities are available in a forest that has a forest functional level of Windows Server 2003 or higher. These capabilities are not available in Microsoft Windows 2000 Server operating systems.

| **Note** |
| --- |
| In Windows Server® 2003 and Microsoft Windows® 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008, the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

You can use the domain rename process to change the names of your domains, and you can also use it to change the structure of the domain trees in your forest. This process involves updating the Domain Name System (DNS) and trust infrastructures as well as Group Policy and service principal names (SPNs).

The structure of an Active Directory forest is the result of the order in which you create domains and the hierarchical names of those domains. Beginning with the forest root domain, all child domains derive their distinguished names and default DNS names from the name of the forest root domain. The same is true of every additional tree in the forest. The way to change the hierarchical structure of an existing domain tree is to rename the domains. For example, you can rename a child domain to have a different parent or rename a child domain to be a new tree-root domain. In each case, you reposition an existing domain to create a different domain-tree structure.

Alternatively, you can rename domains without affecting the structure of the forest. For example, if you rename a root domain, the names of all child domains below it are also changed, but you have not created a different domain-tree structure.

This technical reference subject does not provide procedural information that you need to perform a domain rename operation. Rather, it provides essential background information that you must thoroughly understand before undertaking a domain rename operation in a production environment. This technical reference subject does include information about where to obtain the tools and procedures for a domain rename operation.

**In this subject**


- What Is Domain Rename?
- How Domain Rename Works

# What Is Domain Rename?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**What Is Domain Rename?**

**In this section**


- Domain Rename Constraints and Capabilities
- Core Scenarios for Domain Rename
- Domain Rename Dependencies and Interactions with Other Technologies
- Related Information


You can use the domain rename process to change the names of your domains, and you can also use it to change the structure of the domain trees in your forest. This process involves updating the Domain Name System (DNS) and trust infrastructures as well as Group Policy and service principal names (SPNs).

The ability to rename domains provides you with the flexibility to make important name changes and forest structural changes as the needs of your organization change. Using domain rename, you cannot only change the name of a domain, but you can change the structure of the domain hierarchy and change the parent of a domain or move a domain residing in one domain tree to another domain tree. The domain rename process can accommodate scenarios involving acquisitions, mergers, or name changes in your organization, but it is not designed to accommodate forest mergers or the movement of domains between forests.

**Note**

Domain rename is intended to be a supported method for renaming domains when domain renames are necessary; *it is not intended to make domain rename a routine operation*. The domain rename process is complex, and it requires a great deal of care in planning and execution. In addition, the time that is required for a complete domain rename operation is directly proportional to the size of an Active Directory forest in terms of its number of domains, domain controllers, and member computers. Therefore, although domain rename is possible, it should not be undertaken lightly.

The domain rename operation is not supported in Microsoft Exchange Server 2007 or Exchange Server 2010. DNS domain rename is supported in Exchange Server 2003. However, renaming of the NetBIOS domain name is not supported in any version of Exchange Server. Other non-Microsoft applications might also not support domain rename. For more information about other Microsoft applications that are incompatible with domain rename, see article 300684 (http://go.microsoft.com/fwlink/?LinkId=185229) in the Microsoft Knowledge Base.

In Windows Server® 2003 and Microsoft Windows® 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008, the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory.

For information about methods of merging forests, see “Windows 2000/2003: Multiple Forests Considerations White Paper” on the Microsoft Web site.

**Domain Rename Constraints and Capabilities**

The domain rename capabilities provide solutions to some of the problems that are not addressed in Windows 2000 Server. In a Windows 2000 Server forest, you cannot rename domains after the forest structure is in place without moving domain contents or recreating them.

Although you can rename domains in forests that do not include domain controllers that run Windows 2000, there are some important constraints on these operations. The constraints on domain rename in forests with and without Windows 2000 Server are described in the next two sections, followed by a description of the domain rename capabilities in Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2.

**Domain Rename Constraints in Windows 2000 Server**

The constraints that are associated with changing domain names or restructuring domain trees in Windows 2000 Server Active Directory are prohibitive.

In a Windows 2000 Server forest, you cannot:

- Change the DNS name or the network basic input/output system (NetBIOS) name of a domain. Although you cannot rename a domain, you can achieve the same results by moving its contents into a new domain that has the name that you want the existing domain to have. You can use Active Directory Object Manager (MoveTree) in Windows 2000 Server family Support Tools to move directory objects between domains.

- Move a domain in a forest in a single operation. You can make copies of items in a domain and move items from a domain, but you cannot move the entire domain itself within a forest.

- Split a domain into two domains in a single operation. To split a domain, you must create a new domain and then move users and resources from the existing domain into the new domain.

- Merge two domains into a single domain in a single operation. To merge domains, you must move all the contents from one of the domains into the other domain and then demote all domain controllers in the empty domain and decommission it.

In a Windows 2000 Server forest, significant administrative overhead is associated with performing manual move operations to rename one or more domains or to restructure a domain tree.

**Domain Rename Constraints in Windows Server 2003 and later**
Domain rename is not a trivial operation, and there are important constraints on the domain rename operation in a forest that has domain controllers that run Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003. When you decide whether to rename or restructure domains in an existing forest, be sure to consider what you *cannot* do with domain rename and restructuring. Although a forest has domain rename and restructuring capabilities, certain types of structural changes are not supported.

In a forest with domain controllers that run Windows Server 2003 or later, you cannot:

- Change which domain is the forest root domain. Changing the DNS name or the NetBIOS name (or both) of the forest root domain is supported.

- Drop domains from the forest or add domains to the forest. The number of domains in the forest before and after the domain rename and restructuring operation must remain the same.

- Rename a domain with a name that was taken from another domain in a single domain rename and restructuring operation.

**Domain Rename Capabilities in Windows Server 2003 and later**
Tools are provided that you can use to safely rename domains to restructure a forest, reapply preexisting Group Policy to the new forest structure, and clean up the old domain names that you will no longer use.

The domain rename process involves making basic changes independently at each domain controller in a forest. You set up an administrative computer from which you issue commands that are executed remotely at each domain controller. These commands update the directory database at each domain controller individually with the changes that are necessary for renaming the domains; that is, the updates that rename the domains do not spread across the forest through AD DS replication.

You use the Rendom tool to carry out the following multiple steps in the domain rename process:

- Freeze the current state of the forest so that no changes can occur while a domain rename operation is being performed.

- Prepare the contents of the forest for a domain rename operation. Rendom runs multiple scripts that perform this preparation.

- Execute a domain rename operation.

- Clean up old domain names.

Another tool, Gpfixup, is provided to reinstate Group Policy from the original domains into the newly named domains in the forest.

The Rendom tool, the Gpfixup tool, and the instructions for using them are available on the Windows Server 2003 operating system CD. These tools are also built into domain controllers that run Windows Server 2008 R2 or Windows Server 2008, and they are available in Remote Server Administration Tools (RSAT). You can also download these tools and instructions from Windows Server 2003 Domain Rename Tools on the Microsoft Web site.

| **Note** |
| --- |
| The versions of the domain rename tools that are available at the Microsoft Web site are updated to perform domain rename operations in a forest that has Exchange Server 2003 with Service Pack 1 (SP1) deployed. Versions of the tools that are available on the Windows Server 2003 operating system CD do not have this capability. |

**Core Scenarios for Domain Rename**

By using domain rename capabilities, you can make several kinds of changes to an existing forest. For example, you can perform domain rename to:

- Rename domains without repositioning any domains in the forest structure.
- Create a new domain-tree structure by repositioning domains within a tree.
- Create a new tree root.
- Create a new domain-tree structure by repositioning domains to a different tree.
- Reuse a domain name.

**Domain Rename Without Repositioning**

You can rename domains without restructuring the forest in terms of the parent-child relationships between existing domains. For example, suppose that an existing cohovineyard.com forest for the Coho Vineyard company has four domains with the following names:

- cohovineyard.com (root)
- eu.cohovineyard.com
- hr.cohovineyard.com
- sales.cohovineyard.com

The company decides to expand into wine bottling and distribution and wants to change its name from Coho Vineyard to Coho Winery. The AD DS domain names must now reflect the new company name. As shown in the following figure, the target forest still has four domains with the following new names:

- cohowinery.com (root)
- eu.cohowinery.com
- hr.cohowinery.com
- sales.cohowinery.com

By renaming the forest root domain, a condition is created in which all child domains in the tree must be renamed to preserve

the original forest structure, as shown in the following figure.

(In the following series of figures, two-way arrows indicate two-way, transitive trust relationships.)

**Domain Rename of Four Domains Without Repositioning Domains**




**Domain Rename with Repositioning in the Same Tree**

You can change the structure of a domain tree by renaming a child domain to appear in a different location in the tree. For example, in the cohowinery.com forest, the products.sales.cohowinery.com domain is currently a child of the sales.cohowinery.com domain, placing it two levels below the forest root domain. If the Products division is no longer a subdivision of the Sales organization as a result of an internal reorganization, the company might want to change the domain structure to put the Products organization at the same level as the Sales organization. The following figure shows how changing the parent of products.sales.cohowinery.com results in a restructured domain tree.

**Domain Rename to Change the Parent of a Child Domain**




**Domain Rename with Creation of a New Tree Root**

When you restructure a forest, you can move a domain (except the forest root domain) anywhere within the forest in which the domain resides. You can even move a domain so that it becomes the root of its own domain tree. For example, in the cohowinery.com forest, the domain for the European branch of the company, named eu.cohowinery.com, is a child of the forest root domain. Company management determines that the European division's internal domain name should better reflect its Internet DNS name, cohovineyardandwinery.com. In the target forest structure shown in the following figure, the eu.cohowinery.com domain is moved so that it becomes its own tree-root domain named cohovineyardandwinery.com.

**Domain Rename to Create a New Tree Root**




**Domain Rename with Repositioning to a Different Tree**

By renaming domains, you can effectively move a child domain to a different parent, even if the parent is in a different tree. For example, in the current example forest structure, the domain for the Human Resources (hr) organization is a child of cohowinery.com. This domain has domain controllers in the United States. However, changes in the company have prompted the Human Resources organization to move its location to Europe. Company management wants to move the hr.cohowinery.com domain so that it becomes a child of the domain cohovineyardandwinery.com, residing in another domain tree. As shown in the following figure, the hr.cohowinery.com domain is renamed to hr.cohovineyardandwinery.com.

**Domain Rename to Move a Domain to a Different Tree**




**Reusing a Domain Name**

As described in Domain Rename Constraints in Windows Server 2003 and later, the domain rename operation cannot rename two or more domains so that one domain gives up its name and another domain assumes the same name in a single forest restructuring operation. For example, in the Current Forest configuration in the preceding figure, you cannot use a single domain rename operation to restructure the current forest so that the cohovineyardandwinery.com domain is named something else and the hr.cohowinery.com domain assumes the name cohovineyardandwinery.com.

However, you can accomplish the desired result by first performing the domain rename operation to rename the cohovineyardandwinery.com domain to something else. When you are absolutely sure that the first domain rename operation is complete, you can then perform the domain rename operation again so that hr.cohowinery.com assumes the domain name cohovineyardandwinery.com.

**Domain Rename Dependencies and Interactions with Other Technologies**

So that they can interact in a forest structure, AD DS domains require proper configuration and functioning of the following technologies:

- DNS: AD DS clients, including domain controllers, use DNS zone data to locate resources in a forest. Without a proper DNS infrastructure, AD DS security and replication services cannot function. In a domain name operation, DNS names of

domain controllers and member computers must also change, as follows:

  - The DNS names of domain controllers must change to match the new domain names by changing the primary DNS suffix.

  - The DNS names of member computers change automatically in one of two ways: by configuring the primary DNS suffix to change when the domain name change replicates (with a potentially significant replication impact) or by applying Group Policy to configure the primary DNS suffix change before the domain rename.

- SPNs: SPNs are used for mutual authentication between domain controllers during replication. To ensure that replication can occur following a domain rename operation, the SPN values must change on the domain controllers.

- Trust: Two-way, transitive trust relationships between parent and child domains provide the security infrastructure that is required for resource sharing between domains in the same forest and for delegating management of AD DS objects. To change the structure of a domain tree, you must manually create the trust relationships that enable parent-child relationships in the new structure.

In addition to these basic requirements, when the following features are in effect, they require appropriate changes in the forest in preparation for the domain rename operation:

- Distributed File System (DFS): Folder redirection and roaming user profiles can be used to locate the user's home directory and user profile, respectively, in a network location. When these features are applied using domain-based DFS paths, these paths must change to reflect the new domain names.

- Certification authorities (CAs): Management of enterprise certificates through a domain rename procedure requires that CAs not be installed on domain controllers and that they be configured with appropriate Uniform Resource Locators (URLs).

**Related Information**

The following resources contain additional information that is relevant to this section:

- "Windows 2000/2003: Multiple Forests Considerations White Paper" on the Microsoft Web site

- Windows Server 2003 Domain Rename Tools (http://go.microsoft.com/fwlink/?linkid=5585) on the Microsoft Web site

- How Domain Rename Works

- Domains and Forests Technical Reference

- DNS Support for Active Directory Technical Reference

- Core Group Policy Technical Reference

- Folder Redirection Extension Technical Reference

- DFS Technical Reference

- "Step-by-Step Guide to Implementing Domain Rename" in Windows Server 2003 Domain Rename Tools on the Microsoft Web site

# How Domain Rename Works

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**How Domain Rename Works**

**In this section**


- Domain Rename Processes and Interactions
- Related Information


You can use the domain rename process to change the names of your domains, and you can also use it to change the structure of the domain trees in your forest. This process involves updating the Domain Name System (DNS) and trust infrastructures as well as Group Policy and service principal names (SPNs).

Because the domain rename process involves updating the DNS and trust infrastructures as well as Group Policy and SPNs, a domain rename operation affects every domain controller in the forest. Domain rename is a multistep process that results in updates to the directory and in other side effects. This section provides details about the domain rename process and its interactions with Active Directory Domain Services (AD DS), DNS, Group Policy, and security.

| **Note** |
| --- |
| In Windows Server 2003 and Microsoft Windows 2000 Server, the directory service is named Active Directory. In Windows Server 2008 R2 and Windows Server 2008, the directory service is named Active Directory Domain Services. The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

It is imperative that you not attempt a domain rename operation until you read and understand the contents of this technical reference. The domain rename operation is not supported in Microsoft Exchange Server 2007 or Exchange Server 2010. DNS domain rename is supported in Exchange Server 2003. However, renaming of the NetBIOS domain name is not supported in any version of Exchange Server. Other non-Microsoft applications might also not support domain rename. For more information about other Microsoft applications that are incompatible with domain rename, see article 300684 (http://go.microsoft.com/fwlink/?LinkId=185229) in the Microsoft Knowledge Base.

**Rules for a Well-Formed Forest**

The domain rename capabilities that result in the restructure of a forest support any set of changes to the DNS names and network basic input/output system (NetBIOS) names of the domains in a forest that results in the forest being “well-formed.”

In a well-formed forest, the following conditions are true:

- The DNS names of the domains in the forest form one or more trees.
- The forest root domain is the root of one of these trees.
- An application directory partition cannot have a domain directory partition as a child.

**Domain Rename Conditions and Effects on Service**

An understanding of the conditions and effects of the domain rename process, and a willingness to fully accommodate them, are important prerequisites for undertaking the process.

The following conditions and effects are inherent in the domain rename process:

- Domain rename is supported in a forest in which Exchange Server 2003 with Service Pack 1 (SP1) is deployed. However, domain rename is not supported in an Active Directory forest in which Exchange 2000 Server is deployed. When the domain rename tool detects this condition, it will not proceed with the domain rename process.

  Domain rename is also not supported in a forest in which Exchange Server 2007 or Exchange Server 2010 is deployed.

- Remote management (also known as headless management) is helpful in the domain rename process. Each domain controller in the forest is contacted individually with the domain rename changes; the changes do not spread across the forest through AD DS replication. This condition does not imply that each domain controller in the forest has to be visited physically by an administrator. However, if you want to rename a domain in a large forest, it is highly recommended that you implement remote management of the domain controllers in the forest. In the event that some domain controllers do not respond during the domain rename process, remote management greatly improves your ability to troubleshoot the problem.

- The forest is out of service for a short period of time. Forest service is interrupted during the time it takes for each domain controller to perform the directory database updates that are necessary for the domain rename and to then reboot. The time that the forest is out of service is proportional to the number of domain controllers in the forest. This relatively short service interruption is preferable to the alternative of having the forest in odd "in-between" states for a much longer period of time.

- All domain controllers must either complete the domain rename operation successfully or be eliminated from the forest. The domain rename takes effect even if it is impossible to update some domain controllers in the forest. For the domain rename operation to be complete, every domain controller in the forest must be contacted and updated. If you declare your domain rename operation complete without updating some number of domain controllers because you could not contact them, you must remove all the domain controllers that you could not contact from the forest.

- Each member computer that is joined to a renamed domain must be rebooted twice after all domain controllers are updated. Computers running Windows NT 4.0 must be unjoined and then rejoined to the renamed domain instead of being rebooted.

- If you want DNS host names of domain controllers to match a new domain name, you must perform domain controller rename procedures after the domain rename operation is complete. The DNS host names of *domain controllers* are not changed automatically by the domain rename operation to reflect the new domain name. In other words, the primary DNS suffix of a domain controller will not match the new domain DNS name after the domain has been renamed. Having the host name of a domain controller decoupled from its domain name has no impact on forest service. However, domain controller rename requires a separate, multistep procedure after the domain rename operation is complete.

- The DNS suffix of host names for member computers in a domain that is being renamed might not match the new DNS name of the domain for a period of time. By default, the DNS suffix portion of *member computer* names is updated automatically when the domain to which the computers are joined changes (as happens when you rename a domain). In general, the period of time during which the DNS name of the domain does not match the DNS suffix of member computer names is proportional to the number of computers in the domain. In some cases, you might want to configure the computers to keep the computer names from being updated automatically.

**Domain Rename Processes and Interactions**

Domain rename is implemented in a monitored, step-by-step process that ensures that every domain controller in the forest completes its changes one step at a time — that is, the next step in the process cannot occur until the current step is complete at every domain controller in the forest.

**The Domain Rename Tool (Rendom)**

Rendom.exe is the command-line utility for renaming domains. Rendom is included on the Windows Server 2003 operating system CD. However, an updated version of Rendom is available for download in Windows Server 2003 Domain Rename Tools on the Microsoft Web site. This version of Rendom makes domain rename possible in forests that have Exchange Server 2003 with SP1 deployed.

Rendom.exe is built into domain controllers that run Windows Server 2008 R2 and Windows Server 2008. It is also available in Remote Server Administration Tools (RSAT).

**The Domain Rename State File**

When you issue the first command to begin the domain rename process, Rendom generates an XML-structured text file, called a state file, which contains a list of all the domain controllers in the forest. As domain controllers progress through the various steps in the procedure, Rendom updates the state file to track the state of each domain controller relative to the completion of the domain rename process.

As you perform each step in the domain rename operation, Rendom automatically updates the state file. By using the state file to monitor the state of completion of each domain controller in the state file, you receive the information that you need to

issue the next Rendom command in the sequence.

**Domain Controller States**
Rendom records four states of completion for each domain controller in the state file:

- Initial: Each domain controller that is reachable during the domain rename procedure starts out from the Initial state.
- Prepared: When the domain rename instructions that are written by Rendom have been verified by a domain controller independently, it advances to the Prepared state.
- Final: From the Prepared state, a domain controller advances to one of two Final states. The domain rename process stops when every domain controller in the forest has reached either of the following states:
    - Done: This state signifies that the domain rename is complete at that domain controller.
    - Error: This state implies that some irrecoverable error has occurred, and further progress on the domain rename is deemed to be impossible at that domain controller.

The steps in the domain rename procedure that attempt to take a domain controller from the Initial state to the Prepared state and from the Prepared state to a Final state can be executed only after every domain controller in the forest has reached the required state. A step can be executed multiple times for any domain controllers that cannot be reached in an initial attempt. Each such additional execution of the same step attempts to contact only those domain controllers that have not achieved the required state.

For more information about the contents of the state file, see Domain Rename Script and State File.

**General Steps in the Domain Rename Process**
The following set of steps is a very high-level representation of what happens during the domain rename process. A more detailed explanation of each step is provided later in this section, beginning with "Specifying the Target Forest Structure."

| **Note** |
| --- |
| Do not try to complete the following general steps or any steps that are described in this technical reference. Do not undertake any domain rename procedures until you read and understand all information in this technical reference. The fully documented set of procedures, including all the tasks that you perform, is provided in Administering Active Directory Domain Rename (http://go.microsoft.com/fwlink/?LinkID=187639). |

The general steps in the domain rename procedure are as follows:

1. Before beginning the domain rename process, prepare a list of domains in the forest: Specify the new forest structure that will be represented by the set of changed domain names in the forest.

   Be sure to avoid any possible name conflicts with the new names that you choose. Name conflicts can cause unpredictable and severe results. For example, a conflict with the NetBIOS name can render a domain controller unusable because you might not be able to properly remove AD DS from it.

2. To begin the domain rename procedure, generate a script that contains the instructions for renaming domains in the forest: Generate domain rename instructions that are encoded as a special script based on the specified new forest structure and transfer it to every domain controller in the forest.

3. Verify that all domain controllers are adequately prepared to make the necessary updates to rename the domains: Verify the validity of the domain rename instructions (in the script) at every domain controller, and verify that every domain controller is ready to execute those instructions.

4. Execute the actual domain rename instructions: Execute the domain rename instructions at every domain controller in the forest. At this step, a brief interruption in the forest service may occur.

5. Fix up Group Policy: Update metadata in the directory so that policy settings can continue to be applied after the domain rename.

6. Clean up all domain rename–related metadata that is written to the directory so that the directory is ready for another round of the domain rename operation, if necessary: After the domain rename procedure is complete, remove all metadata that the domain rename operation writes to the directory.

**Requirements for Domain Rename**
Before a domain rename operation begins, the following requirements must be met:

- The forest functional level must be Windows Server 2003 or higher.
- If the position of domains will change, trust relationships must be created to provide trust between any domain that will be renamed (and therefore repositioned) and the domain that is to be its parent in the new structure.
- DNS zones must exist for the new domains.
- Domain-based Distributed File System (DFS) folder redirection paths must be redirected to a server-based path.
- Domain-based roaming user profiles must be relocated to a server-based share or stand-alone DFS path.
- Computers in the to-be-renamed domains must be configured to change their host names to reflect the new domain names.
- Certification authority (CA) requirements must be met.

**Forest Functional Level Requirement**
The domain rename operation is supported in an AD DS forest if—and only if—no domain controllers in the forest are running Windows 2000 Server and the forest functional level has been raised to Windows Server 2003 or higher. Raising the forest functional level to Windows Server 2003 requires that all domains in the forest have a domain functional level of at least Windows 2000 native, regardless of what domains you are renaming.

For more information about forest and domain functional levels, see "Active Directory Functional Levels Technical Reference."

**Trust Relationship Requirement**
You can use the domain rename process to reposition any domain in the domain tree hierarchy of a forest, with the exception of the forest-root domain. Remember that although you can rename the forest root domain (you can change its DNS and NetBIOS names), you cannot reposition it in such a way that you designate a different domain to become the new forest root domain. If your domain rename operation involves restructuring the forest through repositioning of the domains in the domain tree hierarchy, as opposed to simply changing the names of the domains in place, you must first create the necessary shortcut trust relationships between domains so that the new forest structure has two-way, transitive trust paths between every pair of domains in the target forest, just as your current forest does.

**Types of trust relationships**
A hierarchy of AD DS domains is implemented by trust relationships between domains. The following types of trust relationships are established between domains in an AD DS forest:

- Parent-child: The trust that is established when you create a new domain in an existing tree in the forest. The AD DS installation process creates a transitive, two-way trust relationship automatically between the new domain (the child domain) and the domain that immediately precedes it in the namespace hierarchy (the parent domain).
- Tree-root: The trust that is established when you add a new domain tree to the forest. The AD DS installation process creates a transitive, two-way trust relationship automatically between the domain that you are creating (the new tree-root domain) and the forest root domain.
- Shortcut: A manually created, one-way or two-way, transitive trust relationship between any two domains in the forest that is created to shorten the trust path.

The effect of the transitive, two-way trust relationships that are created automatically by the AD DS installation process is that there is complete trust between all domains in a forest — every domain has a transitive trust relationship with its parent domain, and every tree-root domain has a transitive trust relationship with the forest root domain. If you use the domain rename process to restructure an existing forest by altering the domain tree hierarchy, the necessary trust relationships are not

created automatically. Therefore, as part of the preparation phase of the domain rename process, you must manually create the trust relationships that are required to preserve complete trust between all domains in your new forest after it is restructured.

**Precreating trust relationships for a restructured forest**
If you plan to use the domain rename process to reposition one or more domains in the domain tree hierarchy, for each domain that you plan to reposition, you must create the necessary shortcut trust relationships between the domain that you want to reposition and its new parent domain (or the forest root domain if the repositioned domain becomes a tree root). These newly created trust relationships substitute for the required tree-root or parent-child trust relationships that will be missing in the restructured forest.

**Precreating a parent-child trust relationship**
For example, suppose that the Coho Winery company wants to restructure the cohowinery.com forest so that the products.sales.cohowinery.com domain becomes a child of the cohowinery.com domain. Before the domain rename operation is performed to carry out this restructure, a two-way, transitive, shortcut trust relationship between products.sales.cohowinery.com and cohowinery.com must be created first. This trust relationship establishes the basis for the two-way, parent-child trust relationship that will be required for the targeted parent and child domains.

The following figure shows the before and after domain structures and the shortcut trust relationship that will serve as a parent-child trust relationship in the target forest.

**Shortcut Trust Relationship That Becomes a Parent-Child Trust**




**Precreating multiple parent-child trust relationships**
For scenarios in which you want to restructure a domain that is both a child domain and a parent domain, you might need to create shortcut trust relationships in two places. For example, suppose the Coho Winery company wants to restructure the cohowinery.com forest to move the hr.sales.cohowinery.com domain so that it becomes a child of the eu.cohowinery.com domain. At the same time, the company wants to make its child domain, payroll.hr.sales.cohowinery.com, a direct child of its current grandparent domain, sales.cohowinery.com. To perform this restructure operation, shortcut trust relationships must be created that become the parent-child trust relationships for the new forest after the restructure, as follows:

- A two-way, transitive, shortcut trust relationship is created between the eu.cohowinery.com and hr.sales.cohowinery.com domains, which will effect a two-way, transitive, parent-child trust relationship between eu.cohowinery.com and hr.eu.cohowinery.com after the restructure.

- A two-way, transitive, shortcut trust relationship is created between the sales.cohowinery.com domain and the payroll.hr.sales.cohowinery.com domain, which will effect a two-way, transitive, parent-child trust relationship between sales.cohowinery.com and payroll.sales.cohowinery.com after the restructure.

The following figure shows the before and after domain structures and the shortcut trust relationships that will serve as parent-child trust relationships in the target forest.

**Shortcut Trust Relationships to Reposition Two Domains**




**Precreating a tree-root trust relationship with the forest root domain**

When a domain is renamed to become a new tree root, the new tree-root domain must have a two-way, transitive trust relationship with the forest root domain. For this scenario, you create a two-way, shortcut trust relationship between the domain that you want to rename to become a new tree-root domain and the forest root domain.

For example, suppose that there is a deep tree in the cohowinery.com forest and the Coho Winery company wants to create a new tree by moving the lowest-level domain, eu.sales.cohowinery.com, to become a tree-root domain. The following figure shows the two-way, shortcut trust relationship that is created, and the tree-root trust relationship that it provides after restructure, when the eu.sales.cohowinery.com domain is renamed to create the tree-root domain cohovineyardandwinery.com.

**Shortcut Trust Relationship Providing a Tree-Root Trust Relationship**




**DNS Zones Requirement**

When an application or client requests access to AD DS, the domain controller locator (Locator) mechanism finds an AD DS server (that is, a domain controller).

In response to client requests for AD DS services, the Locator uses service (SRV) resource records in DNS to locate domain controllers. When DNS SRV resource records are not available, directory clients experience failures when they try to access AD DS. Therefore, before you rename an AD DS domain, you must be sure that the appropriate DNS zones exist for the forest and for each domain. If the appropriate DNS zones do not exist, you must create the DNS zone or zones that will contain the SRV resource records for the renamed domains. It is also highly recommended that you configure the DNS zone or zones to allow secure dynamic updates. This DNS zone requirement applies to each domain that is renamed as a part of the domain rename operation.

The DNS requirements for renaming an AD DS domain are identical to the DNS requirements for supporting an existing AD DS domain. Your current DNS infrastructure already provides the necessary support for your AD DS domain's use of its current name. Usually, you simply mirror the existing DNS infrastructure to add support for the new name of your domain.

For example, suppose the Coho Vineyard company wants to rename its existing AD DS domain sales.cohovineyard.com to marketing.cohovineyard.com. If the SRV resource records that are registered by the domain controllers in the AD DS domain named sales.cohovineyard.com are registered in the DNS zone named sales.cohovineyard.com, a new DNS zone called marketing.cohovineyard.com must be created, corresponding to the new name of the AD DS domain.

For more information about how DNS provides support for AD DS, see DNS Support for Active Directory Technical Reference. For more information about the Locator, see Locator mechanism.

**Folder Redirection and Roaming User Profile Requirement**

If you are using Folder Redirection or roaming user profiles, you might need to relocate the network paths for them before the domain rename operation. If you place Folder Redirection or roaming user profiles (and the home directory) on a network location by using a domain-based, Distributed File System (DFS) namespace, renaming the domain invalidates the path to this domain-based namespace, and Folder Redirection or roaming user profiles that use this path stop working. For roaming user profiles, this prevents user logons until the paths are fixed. However, if the NetBIOS name of a domain is used to connect to a domain-based DFS namespace and the NetBIOS name of the domain is not changed during a domain rename operation, the path to the domain-based DFS namespace continues to be valid. The same is true for the fully qualified domain name (FQDN); if the FQDN name of a domain is used to connect to a domain-based DFS namespace and the FQDN name of the domain is not changed during a domain rename operation, the path to the domain-based DFS namespace continues to be valid. If you are not using DFS namespaces at all and paths are not specified to server names using FQDNs, the paths will continue to be valid.

For more information, see Relocate Folder Redirection and Roaming User Profiles (http://go.microsoft.com/fwlink/?LinkId=205684).

**Computer Host Name Requirement**

By default, the primary DNS suffix of a member computer of an AD DS domain is configured to change automatically when domain membership of the computer changes. The same default behavior is true when the DNS name of the domain to which a computer is joined changes. Therefore, a rename of an AD DS domain can cause modification of the primary DNS suffix and, consequently, of the full DNS host names of the computers that are the members of the renamed domain.

For example, if the sales.cohowinery.com domain is renamed to marketing.cohowinery.com, the primary DNS suffix of the member computers of this domain might also change from sales.cohowinery.com to marketing.cohowinery.com, depending on whether the default behavior is in effect. If the default behavior is in effect, the full DNS host name of a computer in the renamed domain will change from *hostName*.sales.cohowinery.com to *hostName*.marketing.cohowinery.com.

**Conditions for automatic computer name change**

The primary DNS suffix, and therefore the full DNS name of a member computer in an AD DS domain, changes when the domain is renamed if both of the following conditions are true:

- The primary DNS suffix of the computer is configured to be updated when domain membership changes.
- No Group Policy setting specifies that a primary DNS suffix is applied to the member computer.

| **Note** |
| --- |
| The DNS host names of domain controllers in a renamed domain do not change automatically to use the new domain DNS name as the primary DNS suffix, regardless of the primary DNS suffix configuration. In other words, unlike the names of member computers, the DNS names of domain controllers in a renamed domain remain unchanged. You can rename the domain controllers in a separate step, using a special domain controller rename procedure, after the domain rename operation is complete. |

**Replication effects of renaming large numbers of computers**
If the conditions that prompt automatic update of the primary DNS suffix for all computers in the domain are true, and if there are a large number of member computers in the domain that are being renamed, replication of that many changes might cause excessive traffic on your network. The replication traffic that results from member computer names being updated is a function of the number of computers that are members of the renamed domain. This replication "storm" that can result from thousands of computer names being changed at approximately the same time is a problem for only the largest deployments. You can avoid this condition by changing the setting the prompts automatic update of the primary DNS suffix. If you do not think that the resulting replication traffic poses any problem to your infrastructure (that is, a risk of network congestion or saturation), this preparatory step is not required.

A computer name change triggers update of the **dnsHostName** and **servicePrincipalName** attributes on the corresponding computer account in AD DS. These attributes are typically updated during a reboot of the member computer, as required by the domain rename procedure after the domain is renamed. Update of these attributes by a large number of computers in a short period of time might trigger replication activity that saturates the network. Moreover, a computer name change triggers updates of the host address (A) and pointer (PTR) DNS resource records in the DNS database. Such updates also cause additional replication traffic, regardless of whether DNS zones are stored in AD DS or in some other DNS store. For these reasons, you should prepare for the domain rename operation in advance by reconfiguring the default behavior that changes the primary DNS suffix on member computers when a domain is renamed.

**Group Policy revision of the primary DNS suffix before a domain rename**
To avoid replication and other update traffic that is caused by automatic update of the primary DNS suffix on all member computers following a domain rename, you can use Group Policy to revise the primary DNS suffix to the new domain name before the domain rename operation. When you use Group Policy for this purpose, member names are not automatically updated, but they have the correct primary DNS suffix when you perform the domain rename.

Rather than going to each computer and changing the setting so that the computer name is not updated automatically when the domain name changes, you can use Group Policy to establish the primary DNS suffix for all computers in the domain as the target DNS domain name. Using Group Policy to establish the primary DNS suffix has the following beneficial effects:

- You disable the default behavior of updating the DNS suffix when the DNS domain name changes, thereby avoiding excessive replication traffic.

- The computer names are preconfigured to have a primary DNS suffix that matches the target domain name after the domain rename operation.

| **Note** |
| --- |
| The primary DNS suffix does not have to match the new domain name. If your current implementation uses a primary DNS suffix that does not match the DNS name of the domain to which the member computers are joined and if you do not want the DNS suffix to change after the domain rename, you can ensure that these computers are not renamed after the domain rename by verifying that at least one of the two conditions in "Conditions for automatic computer name change" earlier in this section is not satisfied. |

You can use the Group Policy setting **Primary DNS Suffix** to establish the primary DNS suffix for the domain as the new DNS domain name. When this Group Policy setting is in effect, it overrides the default behavior of changing the primary DNS suffix when the DNS name of the domain changes. In this case, the computer names remain the same when you rename the domain, and replication of name changes does not occur.

However, just as automatic change of the primary DNS suffix causes replication, setting **Primary DNS Suffix** causes replication as well. For this reason, Group Policy should be applied in stages to member computers. To apply Group Policy to all member computers in the domain or domains being renamed — while also avoiding replication on a large scale, divide computer objects among several locations in AD DS, either organizational units (OUs) or sites or both.

**Configuration requirements before application of Group Policy**
When you apply the Group Policy setting **Primary DNS Suffix**, the DNS suffix of member computers no longer matches the DNS name of the domain of which they are members. To allow the member computers of a domain to have a primary DNS suffix that does not match the DNS domain name, you must first configure the domain to accept the names that the DNS suffix can have. This configuration must be in place before you can set Group Policy to apply to a set of computers.

The **msDS-AllowedDNSSuffixes** multivalued attribute of the domain object (the domainDns object for the domain) contains a list of DNS suffixes that member computers of the AD DS domain can have. The Group Policy setting **Primary DNS Suffix** uses the values in this attribute. By editing this attribute to add the primary DNS suffix of the new domain, you make it possible

for Group Policy to set the new domain name as the primary DNS suffix.

**CA Requirements**

Management of enterprise certificates can continue during a domain rename procedure when the following requirements are in effect before domain rename:

- The CAs are not installed on domain controllers.
- As a best practice, all the CAs should include both Lightweight Directory Access Protocol (LDAP) and Hypertext Transfer Protocol (HTTP) Uniform Resource Locators (URLs) in their Authority Information Access (AIA) and certificate revocation list (CRL) distribution point extensions.

| **Note** |
| --- |
| If any certificate that is issued by a CA has only one of these URL types, the certificate may or may not work. Depending on the complexity of your domain configuration, the steps described in the "Step-by-Step Guide to Implementing Domain Rename" (in Windows Server 2003 Domain Rename Tools on the Microsoft Web site) might not be sufficient for proper management of CAs after the domain rename operation. Anyone who undertakes domain rename in an environment that uses certificates must have considerable expertise in managing Microsoft CAs. |

If one or more of the following conditions exist at the time of domain rename, CA management is not supported:

- The CA is configured to have only LDAP URLs for its CRL distribution point or AIA. Because the old LDAP extensions are invalid after the domain rename operation, all the certificates that are issued by the CA are no longer valid. As a workaround, you have to renew the existing CA hierarchy and all issued End Entity certificates.
- An interdomain trust relationship is based on cross-certification with name constraints. After the domain rename operation, the name constraints might not be valid. As a workaround, you have to reissue cross-certificates with appropriate name constraints.
- An e-mail name in the style of Request for Comments (RFC) 822, "Standard for the Format of ARPA Internet Text Messages," is used in the AD DS user account. If the CA (or the certificate template) is configured to include RFC 822–type e-mail names and this e-mail name style is used in the certificates that are issued, these certificates will contain an incorrect e-mail name after a domain rename operation. You should change any such user accounts before any certificates are issued.

For more information about certificates and CAs, see "Certificates Technical Reference" and "CA Certificates Technical Reference."

**Specifying the Target Forest Structure**

After you plan your domain rename, you begin the domain rename process by using the domain rename tool (Rendom) to make a list of the domain names, including application directory partition names, in the current forest structure. After you use Rendom to create this list, you create the new forest structure by editing the names in the list. In accordance with the DNS hierarchical naming scheme, the structure of the forest is implicit in the set of DNS names for the domains and the application directory partitions that make up the forest. In addition to specifying DNS name changes for domains and application directory partitions, you can also change the NetBIOS name of any domain. Changes to the DNS and NetBIOS names of the domains and to the DNS names of the application directory partitions that constitute the forest are supported, subject to the constraints outlined in "Domain Rename Constraints" in What Is Domain Rename?.

**Current Domain Names: Generating the Forest Description File**

The **rendom /list** command generates the current forest description and writes it to an output file (which has the default name Domainlist.xml) using an XML-encoded structure. This file contains a list of all domains and application directory partitions in the forest, along with their corresponding DNS and NetBIOS names. (Application directory partitions do not have NetBIOS names.) Each domain and application directory partition is also identified by a globally unique identifier (GUID), which does not change with a domain rename. To simplify specifying the new forest structure, Rendom gathers and compiles the current forest structure automatically in such a way that the new forest structure can be overlaid on top of the current forest structure.

The following figure shows an example forest description file that is generated for a forest that has three domains named cohovineyard.com (the forest root domain), sales.cohovineyard.com, and hr.sales.cohovineyard.com, as well as four application directory partitions named DomainDnsZones.hr.sales.cohovineyard.com, DomainDnsZones.sales.cohovineyard.com, DomainDnsZones.cohovineyard.com, and ForestDnsZones.cohovineyard.com that are used by DNS. In the example,

nonvariable values are designated in bold text.

**Forest Description File (Domainlist.xml) from the rendom /list Command**

```
             <Forest>
<Domain><!-- PartitionType:Application --><GUID>78438a56-f4a7-383a-5c82-
fe05a76ed464</GUID><DNSname>DomainDnsZones.hr.sales.cohovineyard.com</DNSname>
    <NetBiosName></NetBiosName><DcName></DcName></Domain><Domain>
    <GUID>89cf8ae3-f4a3-453b-ac5c-cb05a76bca40</GUID>
    <DNSname>hr.sales.cohovineyard.com</DNSname><NetBiosName>HR</NetBiosName><DcName>
</DcName></Domain><Domain><!-- PartitionType:Application --><GUID>b9748a56-c4a7-385a-
5d84-7490e76ba484</GUID><DNSname>DomainDnsZones.sales.cohovineyard.com</DNSname>
<NetBiosName></NetBiosName><DcName></DcName></Domain><Domain><GUID>89238a56-f3a1-343b-
bc5c-cb05a76bc341</GUID><DNSname>sales.cohovineyard.com</DNSname>
    <NetBiosName>SALES</NetBiosName><DcName></DcName></Domain><Domain><!--
PartitionType:Application --><GUID>ea658a56-f4a7-383a-5c82-cb05a76bdf35</GUID>
<DNSname>DomainDnsZones.cohovineyard.com</DNSname><NetBiosName></NetBiosName><DcName>
</DcName></Domain><Domain><!-- PartitionType:Application --><GUID>ea658a56-f4a7-383a-
5c82-cb05a76bd461</GUID>
    <DNSname>ForestDnsZones.cohovineyard.com</DNSname><NetBiosName></NetBiosName>
<DcName></DcName></Domain><Domain><! — ForestRoot --><GUID>34fg8ae3-f4a3-453b-ac5c-
3ce5a76bca42</GUID><DNSname>cohovineyard.com</DNSname>
<NetBiosName>COHOVINEYARD</NetBiosName><DcName></DcName></Domain></Forest>
```

**Target Domain Names: Editing the Forest Description File**

The current forest description file that is generated by the **rendom /list** command is a text file that you can edit to express the target forest structure as new domain names. To express the new structure, you simply modify the **<DNSname>** and **<NetBiosName>** fields to contain the new names where they are needed.

**Note**

- The domain GUID value in the <GUID> field represents a fixed name for a domain, and it cannot be modified. The GUID provides a permanent identity that can be used to identify a renamed domain in the new forest structure.

For example, using the preceding forest description file, if you want to change the DNS name of the hr.sales.cohovineyard.com domain to payroll.sales.cohovineyard.com and change its NetBIOS name from HR to PAYROLL, you replace the variable values in the appropriate lines in the forest description file as follows. (Nonvariable values are designated in bold text.)

- Current DNS and NetBIOS names:

  <**DNSname**>hr.sales.cohovineyard.com**</DNSname>**

  <**NetBiosName**>HR<**/NetBiosName**>

- Modified DNS and NetBIOS names; the two lines above are modified as follows:

  <**DNSname**>payroll.sales.cohovineyard.com<**/DNSname**>

  <**NetBiosName**>PAYROLL<**/NetBiosName**>

Based on the new forest structure, Rendom reads information for each domain from the directory and then generates a set of directory update instructions. By default, the information that is collected by Rendom is retrieved from any available domain controller in each domain. However, you have the option to specify a particular domain controller in each domain from which to retrieve the information that is required to generate the domain rename update instructions. To use this option, simply add the DNS host name of the desired domain controller in the <**DcName**><**/DcName**> field of the forest description file.

After you edit the forest description file, you can use the **rendom /showforest** command to display the new forest structure that is contained in the file. (This command does not have any effect on the directory itself.) The user-friendly format uses indentations to reflect the domain hierarchy in the forest.

**Transferring Domain Rename Instructions to AD DS**

After you create the new forest structure by editing the forest description file, the next step in the domain rename process involves translating the new forest structure into a sequence of directory update instructions to be executed individually on each domain controller in the forest. This translation occurs when you issue the **rendom /upload** command, and the resultant

domain rename instructions are uploaded to the configuration directory partition at the domain controller that is currently the domain naming operations master for the forest. The instructions are then replicated to all other domain controllers in the forest through normal replication of the configuration directory partition. Only after these rename instructions replicate to every domain controller in the forest and the required conditions are verified at each domain controller can each domain controller proceed with executing the rename instructions.

The following sections describe the details of the changes that are produced by the **rendom /upload** command. The actual sequence of actions that occur in response to the command are described in "Actions Performed by Rendom in Response to the **rendom /upload** Command" later in this section.

**Domain Rename Script and State File**

To transform the current forest structure into the target structure, Rendom translates the new forest description to a set of update instructions that the domain controller uses to update its replica directory partitions to make the name changes effective. The **rendom /upload** command generates the required domain rename instructions, which are encoded in a special script format that is private to the implementation. The **rendom /upload** command also generates the state file that is used to track the progress of the domain rename operation.

### Note

- The actions that are performed by the **rendom /upload** command and the resultant changes that are made to the directory are in preparation for the rename operation. No domain name changes occur at this step.

**Domain rename script for update instructions**

The script that is generated by the **rendom /upload** command is private to the domain rename procedure, and it does not create the directory changes themselves. Rather, the script provides instructions for domain controllers to execute in response to Rendom commands. The script is an XML-encoded document that has three components:

- Test: a read transaction to validate that the directory replica at the domain controller is in an appropriate state to perform an action.

- Action: an update transaction to be performed on the directory database at the domain controller.

- Signature: a cryptographic hash that proves to the domain controller that the script was prepared by the Rendom utility (and, therefore, that it is authentic and correct) and not manually by someone acting as an administrator.

**State file for tracking**

While translating the forest description to the update instruction script and transferring the script containing the description to the directory, Rendom also generates the state file (which has the default name DClist.xml). This file is structured as an XML-text document to facilitate tracking the state of each domain controller as it progresses through the various steps of the domain rename process. This file contains an entry for every domain controller in the forest. Each domain controller is identified by its DNS host name, along with fields for its current state and the last error encountered on the domain controller while it processed the domain rename instructions.

**Preparing Domain Controllers for Domain Rename**

When you run the **rendom /upload** command, certain changes occur on the domain naming operations master in preparation for the actual execution of domain rename. On the domain naming master, the XML-encoded script that contains the domain rename instructions is written to the single-valued, octet-string attribute **msDS-UpdateScript** on the Partitions container object (cn=partitions,cn=configuration,dc=*ForestRootDomain*) in the configuration directory partition. The Partitions container can be updated only on the domain controller that is the domain naming master for the forest; therefore, the **msDS-UpdateScript** attribute is necessarily changed on the domain controller that holds the domain naming operations master role (also known as the flexible single master operations (FSMO) role). From this source domain controller, the script that is stored in the **msDS-UpdateScript** attribute replicates to all domain controllers in the forest through normal replication of the configuration directory partition.

**Establishing a DNS Alias for a New Domain Name**

In addition to the **msDS-UpdateScript** value being written to the Partitions container, the new DNS name of each domain being renamed is also written by Rendom to the single-valued, Unicode-string attribute **msDS-DnsRootAlias** on the cross-reference object (class **crossRef**) that corresponds to that domain. Again, because cross-reference objects are stored in the Partitions container and the Partitions container can be updated only on the domain naming master, the **msDS-DnsRootAlias** attribute can be changed only on the domain controller that holds the domain naming operations master role. From this source domain controller, the DNS name in the **msDS-DnsRootAlias** attribute replicates to all domain controllers in the forest through normal replication of the configuration directory partition.

**Locator mechanism**

AD DS clients require DNS so that they can locate domain controllers. The Locator is an algorithm that runs in the context of the Net Logon service. The Locator finds the best domain controller for a request, based on network proximity, by using the resource records that are registered by domain controllers in DNS. During the installation of AD DS, the SRV and A resource records are dynamically registered in DNS by the Net Logon service. Both types of records are necessary for the functionality of the Locator mechanism. To find domain controllers in a domain or forest, a client queries DNS for the SRV and A DNS resource records of the domain controller. The resource records provide the client with the names and IP addresses of the domain controllers. In this context, the SRV and A resource records are referred to as Locator DNS resource records. Every domain controller also registers a canonical name (CNAME) record for use by AD DS replication.

The general format for DNS resource records includes several fields, including, but not limited to, the following:

```
Owner name    time to live    record class    record type    data/host name
```

According to the client request, the Locator can use various specified criteria that map to values in the SRV resource records to locate domain controllers with specific roles or capabilities. For example, the Locator can find a domain controller that has a full, writable replica of a domain directory partition; a domain replica in a specified site; or a domain controller that is a global catalog server. The owner name in an SRV resource record provides the specific information about the type of domain controller to be located. For example, a domain controller can be located in a specific site by using the following format:

```
_Service._Protocol.SiteName._sites.DnsDomainName
```

The following owner name in an SRV resource record specifies a domain controller (identified by the LDAP service running over the TCP protocol) in the East site and the Cohowinery.com domain:

```
_ldap._tcp.east._sites.cohowinery.com
```

For more information about the Locator and the DNS resource records that are registered by a domain controller, see "How DNS Support for Active Directory Works."

**Publishing two sets of Locator SRV resource records in DNS**

The directory service is extended in Windows Server 2003 and later forests to use the **msDS-DnsRootAlias** attribute to support the concept of a DNS alias for a domain. The presence of this attribute prompts the Net Logon service running on the domain controller to register the DNS domain name value of this attribute in DNS.

With the addition of the **msDS-DnsRootAlias** attribute, the Net Logon service has the ability to register not one but two domain names in DNS. The identities of the real (original) domain name and the alias (target) domain name are established by publishing two sets of resource records in DNS, as follows:

- The real domain name is published in DNS. Each domain normally has a single DNS name. The domain-specific DNS records that are published by the Net Logon service are derived from the **dnsRoot** attribute on the cross-reference object for the domain, which holds the real DNS name of the domain (as opposed to an alias). The forest-specific DNS records are derived from the cross-reference object for the forest root domain, which holds the real name of the forest root domain. The Net Logon service running on a domain controller also responds to Locator pings for the domain name and performs secure channel setup between a computer that is joined to the domain and the domain controller on which the Net Logon service runs.

- The alias domain name is published in DNS. Soon after the **msDS-DnsRootAlias** attribute value on a cross-reference object is set on a domain controller, either by an originating write or by a replicated write, the Net Logon service on the domain controller additionally publishes a parallel set of DNS Locator records for the domain alias name that is held in the **msDS-DnsRootAlias** attribute. The domain-specific DNS records are derived from the **msDS-DnsRootAlias** attribute on the cross-reference object for the domain, and the forest-specific DNS records are derived from the **msDS-DnsRootAlias** attribute on the cross-reference object for the forest root domain. Further, on domain controllers for a domain whose **msDS-DnsRootAlias** value is set, the Net Logon service responds to the Locator pings for the DNS alias as if the value in **msDS-DnsRootAlias** were the real domain name. Also, secure channel setup from a domain member that connects to a domain controller for a domain that is named in the **msDS-DnsRootAlias** attribute succeeds as if it were the real domain name.

Note that in the published SRV resource record corresponding to the alias domain name, the owner field reflects the new domain name, whereas the host name field reflects the actual DNS name of the domain controller. For example, if the domain cohovineyard.com is being renamed to cohowinery.com, the Net Logon service publishes the following two SRV resource records for the LDAP service on a domain controller named DC01 in that domain:

```
_ldap._tcp.cohovineyard.com  SRV  0  0  389  dc01.cohovineyard.com
_ldap._tcp.cohowinery.com  SRV  0  0  389  dc01.cohovineyard.com
```

Note that in the second SRV record above, the owner field corresponds to the new name to which the domain will be renamed, and the host name field reflects the true DNS name of the domain controller.

**Prepublishing CNAME Records for Replication**

The set of domain-specific resource records that are published in DNS by the Net Logon service running on a domain controller includes a DNS CNAME record that AD DS uses for replication.

The owner name of the CNAME record is provided by:

```
<DsaGuid>._msdcs.<DnsForestName>.
```

which allows a domain controller to locate a replication partner in the forest.

To find its replication partner, the domain controller requires knowledge of only the GUID of the directory system agent (DSA) object for that domain controller (as represented by **DsaGuid** in the CNAME record). The DSA GUID is the GUID of the NTDS Settings object (class **nTDSDSA**). Its value is stored in the **objectGUID** attribute of the NTDS Settings object, which is a child of the domain controller server object. These objects reside in the Sites container in the configuration directory partition. If the forest root domain is being renamed, it is essential that this CNAME record is published beforehand in DNS so that replication continues to work after a domain controller is updated to change the forest root domain name. If the DNS CNAME records for the new forest root domain name are not published beforehand for each domain controller, and records for the old forest root name are replicated among DNS servers that use AD DS to store DNS zones, replication is interrupted as a result of a cyclic condition. This condition generates the following error: "Cannot replicate because the CNAME record cannot be read from the local DNS replica; the CNAME record is not present in the local DNS replica because it is not replicating." Prepublication of the DNS records shortens the period during which the directory service is unavailable during the forest restructuring.

Note that in the prepublished CNAME resource record that corresponds to the new forest name, the owner field reflects the new forest name and the host name field reflects the actual DNS name of the domain controller. For example, if the forest cohovineyard.com is renamed to cohowinery.com, the following two CNAME resource records are published by the Net Logon service for a domain controller named DC01 in that domain:

```
<DsaGuid>.cohovineyard.com  IN  CNAME  dc01.cohovineyard.com
<DsaGuid>.cohowinery.com  IN  CNAME  dc01.cohovineyard.com
```

Note that in the second SRV record above, the owner field corresponds to the new name to which the forest will be renamed, and the host name field reflects the true DNS name of the domain controller.

Before the domain rename process can continue to the next step, the prepublished CNAME records corresponding to the value in **msDS-DnsRootAlias** must replicate to all DNS servers that are authoritative for those records.

**Establishing SPNs for a New Domain Name**

The DSA on every domain controller writes a set of SPNs to the **servicePrincipalName** attribute on the domain controller computer object in the domain directory partition. Among other things, SPNs are used for mutual authentication between domain controllers during AD DS replication. The specific SPN that is used for mutual authentication between replication partners has the following three-part format:

```
E3514235-4B06-11D1-AB04-00C04FC2DCD2/<DsaGuid>/<DnsForestName>
```

where the first part is the remote procedure call (RPC) interface GUID for AD DS replication, the second part is the GUID of the DSA object for the domain controller, and the third part is the DNS name of the forest root domain.

**Prepublishing Two Sets of SPNs**

Soon after the value for the **msDS-DnsRootAlias** attribute on a cross-reference object is set on a domain controller, either by an originating write or by a replicated write, the DSA on the domain controller rewrites the SPNs on the domain controller computer object so that each SPN that includes the domain name (or the forest root name) is present in two versions — one for the actual domain (or forest root) name and one for the alias that is held in the **msDS-DnsRootAlias** attribute of the domain (or forest root) cross-reference object. The SPN values that correspond to the domain name alias require prepublication for the same reason that the DNS CNAME records do, that is, to shorten a cyclic condition that produces the following error: "Cannot replicate because the SPN needed for mutual authentication cannot be read from the directory replica; the required SPN value is not present in the directory replica because it is not replicating." Prepublication of the SPNs shortens the period during which the directory service is unavailable during the forest restructuring.

Before the domain rename process can continue to the next step, the prepublished SPNs that correspond to the value in **msDS-DnsRootAlias** must replicate to all domain controllers in the domain as well as to all global catalog servers in the

forest.

**Actions Performed by Rendom in Response to the rendom /upload Command**
The following sequence of actions occurs in response to the **rendom /upload** command:

- Rendom connects to the domain controller that holds the domain naming operations master role for the forest and validates the forest description against the current state of the forest. If any of these validity checks fails, the command fails now. The following requirements are verified:
    - Each existing domain is part of the new forest.
    - The new forest is well formed.
    - The new forest does not reassign domain names that are being relinquished as part of the current domain rename operation.
- While it is still connected to the domain naming master, Rendom retrieves all the information that is needed to compute the list of rename instructions for updating the configuration and schema directory partitions.
- Rendom connects to a randomly selected domain controller (or the domain controller that is specified in the **<DcName></DcName>** field of the domain entry in the forest description file) for each domain, one by one, and retrieves all the information that is needed to compute the list of rename instructions for updating each domain.
- Rendom computes the full list of domain rename instructions for updating the entire forest (creates the script) and computes a signature to include in the script so that the authenticity of the script can be proven to a domain controller.
- While it is still connected to the domain naming master, Rendom writes the resulting script to the **msDS-UpdateScript** attribute of the Partitions container.
- While it is still connected to the domain naming master, Rendom writes the **msDS-DnsRootAlias** attribute of all cross-reference objects for domains that are being renamed.

On the control station (the computer from which the Rendom commands are issued), Rendom writes a new state file to track the progress of every domain controller in the forest. The state file contains an entry for every domain controller in the forest, with each domain controller entry marked to be in the Initial state.

**Note**

- Because the **msDS-UpdateScript** and **msDS-DnsRootAlias** attributes are first written to the directory on the domain controller holding the domain naming operations master role and then replicated to the remaining domain controllers in the forest, if the domain naming master is unavailable during the **rendom /upload** operation, the process cannot continue.

**Side effect of the rendom /upload command on forest configuration**
If certain types of changes in the forest are made after the domain rename operation begins, these changes can affect the outcome of the rename operation and cause it to fail in such a way that some of the domain controllers can never complete the domain rename process. To prevent this situation, successful execution of the **rendom /upload** command freezes the forest configuration with respect to these types of changes after the domain rename operation begins. The following operations cannot be performed in a forest after the **rendom /upload** command has completed successfully:

- The addition or removal of a domain or application partition
- The addition or removal of a domain controller into or from an existing domain
- The addition or removal of trusts, including cross-forest trusts. (Note that attributes of an existing trust can be changed during this frozen configuration. For example, a unidirectional trust can be converted to a bidirectional trust.)

The forest configuration remains frozen as long as the **msDS-UpdateScript** attribute of the Partitions container has a value

that indicates that a domain rename operation is in progress. The forest configuration becomes unfrozen at the conclusion of the domain rename operation through execution of the **rendom /end** command.

#### Verifying Domain Controller Readiness

This step of the domain rename process verifies that the directory database at each domain controller in the forest is adequately prepared to perform the directory modifications that are dictated by the script in the **msDS-UpdateScript** attribute. In response to the **rendom /prepare** command, the necessary verification is performed on each domain controller through execution of the test component of the script in **msDS-UpdateScript** as a read-only transaction on the local directory database. The test checks for the following conditions:

- The correct script in **msDS-UpdateScript** replicating to this domain controller
- Any trust relationships that are required by the new forest structure
- Prepublished SPNs for all domain controllers
- Name conflicts that are attributable to administrative errors since the time the script in **msDS-UpdateScript** was created. For example, an administrator might have mistakenly removed a trust that is required in the new forest or created a computer account whose **SamAcountName** equals the new name of some trusted domain, creating a name conflict with an interdomain trust account.

In addition, the test includes an authorization check on each domain controller to determine whether or not the user running the **rendom /prepare** command is authorized to execute the domain rename instructions that are contained in **msDS-UpdateScript**. The authorization check verifies that the user has Write permission on the **msDS-UpdateScript** attribute on the Partitions container. If the user is not authorized, the command fails with an error.

#### Actions Performed in Response to the rendom /prepare Command

The following sequence of actions occurs in response to the **rendom /prepare** command:

- Rendom issues a special RPC (for internal use only) to every domain controller in the forest in turn, requesting authorization and verification of readiness as encoded in the test component of the script in **msDS-UpdateScript**. Rendom issues the RPC to multiple domain controllers at a time with a high degree of concurrency, while ensuring that resource limits on the computer that is executing the command are not exceeded. The RPC request and response are signed and sealed for integrity and privacy.
- In response to the RPC, each domain controller performs the authorization check, ensures the authenticity of the script in **msDS-UpdateScript** by validating the signature, and verifies the test component of the domain rename script before responding. If any of these checks fails on a domain controller, the RPC returns an error for that domain controller.
- Rendom updates the state file with the state of each domain controller that was successfully contacted and verified. The state of each successfully verified domain controller is updated from the Initial state to the Prepared state. The Prepared state indicates that the domain controller authorized the execution of the restructure and that the contents of its directory database are consistent with the rename instructions that are contained in **msDS-UpdateScript**. If a domain controller cannot be contacted or if it fails any of the checks, its corresponding state in the state file remains as Initial with an appropriate error code and message to indicate the cause of failure.

#### Executing Domain Rename Instructions

In the final step of the domain rename process, the directory database at each domain controller in the forest is updated individually to implement the new forest structure. This process does not rely on AD DS replication. Rather, the action component of the script in **msDS-UpdateScript** performs the required modifications to the directory database locally on each domain controller as a single update transaction. The action component of the script is the actual update of the domain name. The **rendom /execute** command performs this final update step.

### Note

- The actions that are performed at this step make the actual domain name changes effective at each domain controller. This step causes a brief interruption in service. Before this point in the process, forest service is not disrupted.

#### Single-User Mode

To perform the action component of the script in **msDS-UpdateScript**, the directory service on each domain controller enters a special mode called single-user mode. In single-user mode, AD DS refuses service to all normal clients of the directory service, including LDAP, Messaging API (MAPI), replication, Security Accounts Manager (SAM), Kerberos, and other directory service RPCs. In this mode, only the directory service itself can read and write the local directory database, using a single thread. The single-user mode is needed because the domain rename updates that the directory service performs invalidate internal directory data structures. After the directory service performs these updates in single-user mode, the domain controller reboots. After the reboot, the internal data structures are rebuilt to a consistent state.

**Update Transaction**

All the directory updates that are specified by the action component of the script in **msDS-UpdateScript** are performed within an Extensible Storage Engine (ESE) database transaction. If no errors occur, the transaction is committed; otherwise, the entire transaction is rolled back. Domain controllers that commit the update transaction in single-user mode and then reboot are considered to have completed the domain rename.

**Entering single-user mode**

As part of a successful update transaction, the nonreplicated, single-valued integer attribute **msDS-ReplicationEpoch** on the NTDS Settings object for the domain controller is updated to a new value through incrementation of its current value. If two domain controllers have different **msDS-ReplicationEpoch** values, no directory replication RPC interaction is allowed between them. In addition to replication, nested group membership evaluation and global catalog lookups are also discontinued. Because the domain rename procedure involves making these updates independently at all domain controllers in the forest, it is impossible to modify these domain controllers simultaneously. The goal of the **msDS-ReplicationEpoch** attribute is to minimize potentially complex interactions, including replication, between those domain controllers that have completed the domain rename and those domain controllers that have not yet completed the domain rename.

**Switching real and alias DNS names**

Further, as part of a successful update transaction, the values of the **dnsRoot** and **msDS-DnsRootAlias** attributes on the cross-reference objects of renamed domains are interchanged so that the new domain name, formerly stored in **msDS-DnsRootAlias**, becomes the effective domain name that is stored in **dnsRoot**. The old domain name is saved as a DNS alias until a subsequent step removes it.

**Actions Performed in Response to the rendom /execute Command**

The following sequence of actions occurs in response to the **rendom /execute** command:

- Rendom checks to ensure that all domain controllers in the state file are marked with a current state of Prepared. If the state of any domain controller is not Prepared, Rendom reports an error and the process cannot continue.

- When all domain controllers in the state file are in the Prepared state, Rendom issues a special RPC (for internal use only) to every domain controller in the forest requesting execution of the directory updates that are encoded in the action component of the script in **msDS-UpdateScript**. Rendom issues the RPC to multiple domain controllers at a time with a high degree of concurrency, while ensuring that resource limits on the computer that is executing the command are not exceeded. The RPC request and response are signed and sealed for integrity and privacy.

- In response to the RPC, each domain controller first verifies the test component of the domain rename script in **msDS-UpdateScript** (just as in the previous step that checks for domain controller readiness). The domain controller then performs the action component of the script in an update transaction, as described in "Update Transaction" earlier in this section. If any check fails on a domain controller or if the update transaction cannot be successfully committed, the RPC returns an error for that specific domain controller.

- Rendom updates the state file with the state of each domain controller that was successfully contacted. For each domain controller that successfully completed the update, Rendom changes the state from Prepared to the final state of Done. If a domain controller cannot be contacted or if it fails any of the checks, its corresponding state in the state file remains Prepared, with an appropriate error code and message to indicate the cause of failure. If the RPC returns with an error that is deemed irrecoverable, the corresponding state for the domain controller is updated to the final state of Error, with an appropriate error code and message to indicate the cause of failure.

**Determining Domain Rename Completion**

In a large forest, some number of domain controllers might prove impossible to contact during the domain rename process. These domain controllers never reach the final state of Done. You must decide how long to continue to try to reach domain controllers that are unreachable and how long to retry failed update attempts on domain controllers that have reached the Error state. When further progress in the forest is not possible, declare the domain rename process to be complete. All domain controllers that did not reach the final Done state, because they were either unreachable or finished in the Error state, must be

demoted or removed from service, if demotion is not possible. For a forest to function without problems after a domain rename, only domain controllers that have reached the Done state can exist in the forest.

**Reconciling Group Policy After a Domain Rename**

When the DNS name of a domain changes, any references to Group Policy objects (GPOs) in the renamed domain through Group Policy links (the **gpLink** attribute) on sites, domains, and OUs are rendered invalid because they are based on the old domain name. Further, the optional attribute **gpcFileSysPath** on a GPO that holds a uniform naming convention (UNC) path to a Group Policy templates folder located in the SYSVOL of the renamed domain is also rendered invalid because the path uses the old domain DNS name. To correct the severed Group Policy links and the invalid UNC paths in GPOs in the renamed domain, you can use a Group Policy tool, Gpfixup.exe, to refresh the Group Policy links and the UNC paths in GPOs that are based on the new domain name. The Gpfixup tool is available for download in Windows Server 2003 Domain Rename Tools on the Microsoft Web site. The tool is built into domain controllers that run Windows Server 2008 R2 and Windows Server 2008. It is also available in Remote Server Administration Tools (RSAT). Run Gpfixup once for every renamed domain soon after the actual domain rename operation is complete and before another domain rename operation is performed.

### Note

- Gpfixup refreshes all intradomain GPO references or links (that is, where the link and the target GPO are in the same domain) in the renamed domain. However, cross-domain references to GPOs in the renamed domain (that is, where the link is in a different domain from the domain containing the GPO) are not rebuilt automatically by this tool. For these cross-domain references to work, you must repair the cross-domain links manually by deleting the old Group Policy links and re-establishing new links.

**Removing Old Domain Names After a Domain Rename**

Following completion of the domain rename process for a forest, you use the **rendom /clean** command to remove the old domain names from AD DS. This cleanup step removes all values of **msDS-DnsRootAlias** from the domain naming operations master, and removal of this value is replicated to all domain controllers in the forest. Because this attribute holds the old domain name after the domain rename is completed, replication of the removal of the **msDS-DnsRootAlias** attribute value to all domain controllers in the forest prompts the Net Logon service on the domain controller to deregister the DNS Locator resource records for the old domain name. In the course of this deregistration, the DNS CNAME resource records, which are used by AD DS replication, are also removed for the old domain name. In addition, the cleanup procedure removes the **msDS-UpdateScript** attribute value from the Partitions container on the domain naming master.

After you run the **rendom /clean** command successfully, the new forest is ready for another forest restructuring operation.

**Related Information**

The following resources contain additional information that is relevant to this section:

- Windows Server 2003 Domain Rename Tools on the Microsoft Web site
- Active Directory Functional Levels Technical Reference
- Domains and Forests Technical Reference
- DNS Support for Active Directory Technical Reference
- Designing the Active Directory Logical Structure in *Designing and Deploying Directory and Security Services* in the Windows Server 2003 Deployment Kit
- Folder Redirection Extension Technical Reference
- "Step-by-Step Guide to Implementing Domain Rename" in Windows Server 2003 Domain Rename Tools on the Microsoft Web site
- Certificates Technical Reference
- CA Certificates Technical Reference
- What Is Domain Rename?

- How DNS Support for Active Directory Works
- Core Group Policy Technical Reference

# Core Operating System Collection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Core Operating System Collection**
The Microsoft Windows Server 2003 operating systems contain core technologies that provide basic functionality such as printing, device management, data recovery and troubleshooting.

**Recovery and Reliability**
Recovery and reliability technologies help you to prevent problems such as data loss and system instability during all phases of operation — from startup through shutdown. In the event of a problem, these technologies help you to recover and repair your data and operating system.

**Device Management Technologies**
Device management provides a way for Microsoft Windows Server 2003 to uniformly notify all applications and system components of changes that may affect their operation or access to resources.

Applications and system components use and process device messages between the operating system and devices to take advantage of new resources when they become available and to prevent loss of data when existing resources become unavailable. Device messages notify applications of device change events, such as the removal of a hardware device; power messages notify applications of power management events. Applications and device drivers can also define and use custom messages to enable notification of other types of events.

**Printing Technologies**
Network printing is the collection of software components in the Microsoft Windows Server 2003 operating systems that provide network printing services for client computers. In addition to providing services to clients running versions of the Windows operating system, Windows Server 2003 can provide network printing services to client computers running other operating systems when the appropriate components are installed.

In some small offices, it might be practical to have a printer attached to every computer, but in an organization consisting of more than a few computers, sharing printing resources is desirable. Larger organizations with many users might be spread out over a large area and consist of multiple computing platforms. Windows Server 2003 gives you options for printing to remote printers in a variety of network configurations and from a variety of computing platforms.

# Printing Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Network printing is the collection of software components in the Microsoft Windows Server 2003 operating systems that provide network printing services for client computers. In addition to providing services to clients running versions of the Windows operating system, Windows Server 2003 can provide network printing services to client computers running other operating systems when the appropriate components are installed.

In some small offices, it might be practical to have a printer attached to every computer, but in an organization consisting of more than a few computers, sharing printing resources is desirable. Larger organizations with many users might be spread out over a large area and consist of multiple computing platforms. Windows Server 2003 gives you options for printing to remote printers in a variety of network configurations and from a variety of computing platforms.

# Network Printing Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Network printing is the collection of software components in the Microsoft Windows Server 2003 operating systems that provide network printing services for client computers. In addition to providing services to clients running versions of the Windows operating system, Windows Server 2003 can provide network printing services to client computers running other operating systems when the appropriate components are installed.

**In this subject**


- What Is Network Printing?
- How Network Printing Works
- Network Printing Tools and Settings

# What Is Network Printing?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Network Printing?**

**In this section**


- Common Printing Scenarios
- Dependencies and Interactions


Network printing is the collection of software components in Windows Server 2003 that provide network printing services for client computers. In addition to providing services to clients running versions of the Windows operating system, Windows Server 2003 can provide network printing services to client computers running other operating systems when the appropriate components are installed.

In some small offices, it might be practical to have a printer attached to every computer, but in an organization consisting of more than a few computers, sharing printing resources is desirable. Larger organizations with many users might be spread out over a large area and consist of multiple computing platforms. Windows Server 2003 gives you options for printing to remote printers in a variety of network configurations and from a variety of computing platforms.

Running on a variety of client platforms, applications send print jobs to printers that are connected to a print server running Windows Server 2003. The printer can be connected to the print server through internal network adapters, external network adapters, or another server. This subject discusses the nature and action of components involved in providing network printing services on Windows Server 2003.

**Common Printing Scenarios**

With network printing services in Windows Server 2003, you can provide printing services for a variety of client operating systems, including non-Windows operating systems and older versions of Windows. The following scenarios describe the printing services provided by Windows Server 2003 to various client platforms.

**Printing From Client Computers Running Windows 2000 and Windows XP**

Client computers running Windows 2000 or Windows XP are able to send print jobs to print servers running Windows Server 2003 with no additional software. They are also able to use the Active Directory directory service to locate print devices with the capabilities and location that the user desires.

**Printing From Client Computers Running Windows 95, Windows 98, and Windows NT**

Client computers running Windows 95, Windows 98, and Windows NT can use Windows Server 2003 network printing services without adding software. However, to use Active Directory to search for print devices, these clients must have Active Directory Client Extensions installed.

In Windows Server 2003, printer drivers for client computers running Windows 95, Windows 98, or Windows Millennium Edition can be installed on the print server. When a client then connects to the server, it downloads the drivers from the server and begins printing.

**Printing From Client Computers Running Other Microsoft Operating Systems**

Client computers running Windows 3*.x* and MS-DOS operating systems use their own printer drivers. Clients with these systems can get a list of all available printers in a domain but cannot use Active Directory. These clients can print to a print server running Windows Server 2003 if a local printer is installed and the local port is redirected to the shared printer on the print server running Windows Server 2003.

**Printing From UNIX-Based Clients**

Computers running UNIX-based operating systems, and other systems that use the client program line printer remote (LPR), can send printing jobs to the receiving utility, line printer daemon (LPD), on a print server running Windows Server 2003. LPR clients must comply with Request for Comments (RFC) 1179. For more information about LPR, see RFC 1179 in the IETF RFC Database.

The combination of the LPR and LPD was developed for UNIX-based computers, but it is widely used for many operating systems. Both utilities are included in Print Services for UNIX.

**Printing From NetWare Clients**

With File and Print Services for NetWare installed on a print server running Windows Server 2003, the server appears to a NetWare client as a NetWare 3.*x*–compatible file and print server. File and Print Services presents the same dialog boxes to the client as a NetWare-based server uses to process a print job from a client. A user can display and search for printers on the print server running Windows Server 2003 just like in a NetWare environment.

**Printing From Macintosh Clients**

Macintosh clients can send print jobs to a print server that is running Windows Server 2003 when Print Server for Macintosh is installed on the server. To the Macintosh-based client, the server appears to be an AppleTalk printer on the network, and no reconfiguration of the client is necessary.

**Dependencies and Interactions**

The following software components might be required to provide print services to clients, depending on the client platform.

- **Active Directory Client Upgrade.** Clients that are running Windows 95, Windows 98, and Windows Millennium Edition can use Windows Server 2003 network printing services without adding software. However, to use Active Directory for a printer search, these clients must have Active Directory Client Extensions installed.

- **Network Redirector.** Computers running MS-DOS, Windows 3.0, and Windows 3.1 send documents for printing to local ports. To print to network print servers, clients running these operating systems require a network redirector for printing. A redirector is not necessary for Microsoft Windows for Workgroups because the redirector is built in. In these cases, install Microsoft Network Client 3.0 for MS-DOS or Windows 3.*x*.

- **Print Services for UNIX.** To accept printing jobs from LPR clients, Windows Server 2003 must have the LPD service installed. LPD is provided by Print Services for UNIX.

- **File and Print Services for NetWare.** To accept printing jobs from NetWare clients, Windows Server 2003 must have the File and Print Services for NetWare installed.

- **Print Server for Macintosh.** To accept printing jobs from AppleTalk clients, like a Macintosh computer, Windows Server 2003 must have the Print Server for Macintosh installed.

# How Network Printing Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Network Printing Works**
**In this section**


- Terms and Definitions
- Network Printing Architecture
- Network Printing Protocols
- Network Printing Processes and Interactions
- Network Ports used by Network Printing
- Related Information


Network Printing enables users in locations geographically separate from each other and from their print devices to produce documents for themselves and others. Print servers enable multiple clients to share one or more print devices. By using the Active Directory directory service in the Windows 2000 Server and Windows Server 2003 operating systems, users can locate printers in a target location that supports the features necessary to produce their documents.

This section describes the components and processes of network printing with print servers running Windows Server 2003.

**Terms and Definitions**
The following three terms are data types associated with Network Printing.

**RAW**
Data type used in a page description language that is ready to be sent to a print device; for example, PCL or PostScript. This is the default data type for print jobs on computers that are not running Windows XP Professional or Windows 2000. These RAW files are device-dependent; that is, the spooled data is destined and formatted for a particular device.

**EMF**
Data type used between print clients running Windows XP Professional and print servers running Windows Server 2003. With an enhanced metafile (EMF), the Graphics Device Interface (GDI) releases console control after generating the EMF. The EMF data is then interpreted in the background on a spooler thread and sent to the printer driver. Splitting the rendering of a print job in this way is especially useful for very large documents because the application is not tied up for the entire rendering time.

EMF spool files are encoded to provide greater printer independence. For example, a graphic measuring 2 inches by 4 inches on a video graphics adapter (VGA) display and stored in an EMF maintains those dimensions whether it is printed on a 300-dpi laser printer or a 75-dpi dot-matrix printer. The EMF data type also ensures that the print server uses the fonts you specified.

**TEXT**
Data type used to send a simple-text print job to a printer (such as a PostScript device) that cannot interpret simple text. The spooler creates a new print job, embedding the text in print instructions that are derived from the printer's defaults for font, form, and orientation. Several character sets are in common use, and text files do not indicate which one to use. The TEXT data type uses the ANSI character set, so it might print some characters incorrectly if the application uses a different set.

**Network Printing Architecture**
Network printing components include a server spooler on a print server running a product in the Windows Server 2003 family and a client spooler running on a client computer. Because print jobs usually originate on a client computer, this section first describes the architecture of the client spooler.

**Client Spooler**
The client spooler (Winspool.drv) is the component on Windows operating systems that delivers print jobs from applications on the client computer to print servers. On client computers running Windows XP Professional or Windows 2000 operating systems, print jobs are delivered directly to the client spooler; on other Windows operating systems, a redirector is necessary to deliver print jobs to the client spooler, and on non-Windows operating systems, other mechanisms are required to complete

a print job.

The following figure shows the components involved in generating and sending print jobs to the print server.

**Components of Client Computers That Generate and Send Print Jobs**




**Note**

- The preceding figure is a simplification of network printing processes. Specific platforms use different components and processes. For information about printing processes specific to your environment, you might need to refer to documentation specific to your platform.

Windows clients and non-Windows clients include the following operating systems:

**Windows Clients**
Client computers running Windows 95, Windows 98, Windows Millennium Edition, Windows 2000, Windows XP, or Windows Server 2003.

**Non-Windows Clients**
Client computers running UNIX, Apple Macintosh, or MS-DOS. This group also includes NetWare clients.

The following table defines the client components that generate and send print jobs.

**Components of Client Computers That Generate and Send Print Jobs**

| Component | Description |
|---|---|
| Windows Application | Any program that creates printable data. |
| Non-Windows Application | Any program that creates printable data. |
| Client Spooler | Processes and sends print jobs from the client to the print server. |
| Graphics Device Interface (GDI) | Provides services to the printer driver, including caching, client-server communications, and character conversion. |
| Printer Driver | Communicates with print devices to generate output from data supplied by the GDI. |

| | |
|---|---|
| Line Printer Remote (LPR) | Protocol that sends print jobs to a server running the server-side protocol, Line Printer Daemon (LPD). |
| AppleTalk | Protocol used by Macintosh clients and Windows servers, when Services for Macintosh is installed on the servers. |
| NetWare | Protocol used by NetWare clients for file and print services and by Windows servers when NWLink is installed on the server. |
| Local Print Device | Prints device-specific data. This device is attached to a local port on the server. |
| Remote Print Device | Prints device-specific data. This device is directly connected to the network and managed by this print server. |
| Remote Print Server | Accepts print jobs for print devices it manages. |

**Server Spooler**

Windows Server 2003 accepts and delivers print jobs using a variety of protocols to serve a diversity of client computers and print devices. The server spooler is the primary component of network printing architecture on Windows Server 2003. The server spooler and its components are consolidated into a single architecture, providing smooth background printing by using background thread processing. Because the thread processing for one print job must complete before the processing begins for the next print job, the spooler passes data to the printer only when the printer is ready to receive more data.

The following figure illustrates the server spooler and how it processes a print job.

**Components of Windows Server 2003 Print Server**




The following table defines the components of network printing on Windows Server 2003.

**Components of Network Printing on Windows Server 2003**

| Component | Description |
|---|---|
| Print Spooler Service | Manages the printing process, which includes retrieving the location of the correct printer driver, loading the driver, spooling high-level function calls into a print job, scheduling the print job for printing, and so on. |
| Print Router | Routes print jobs. The print router determines which print provider to call based on a printer name or other information supplied with the print job. |
| Local Print Provider | Provides job control and printer management capabilities for all printers that are accessed through the print monitors of the local print provider. |
| Print Monitor | Directs print jobs to the language monitor and port monitor. |
| Language Monitor | Provides bidirectional communication between the print device and the client. |
| Local and Remote Port Monitors | Sends print jobs to print devices that are attached to local and remote ports. |
| Graphics Device Interface (GDI) | Provides services to the printer driver, including caching, client-server communications, and character conversion. |
| Printer Driver | Communicates with print devices to generate output from data supplied by the GDI. |
| Remote Print Provider | Directs print jobs to remote print servers. |
| Local Print Device | Prints device-specific data. This device is attached to a local port on the server. |
| Remote Print Device | Prints device-specific data. This device is directly connected to the network and managed by this print server. |
| Remote Print Server | Accepts print jobs for print devices it manages. |

**Network Printing Protocols**
The following protocols are used by Windows Server 2003 when providing network print services. Note that some protocols might not be used, depending on the needs of the client computers.

**Server Message Block (SMB)**
A file-sharing protocol designed to allow network computers to transparently access files that reside on remote systems over a variety of networks. The SMB protocol defines a series of commands that pass information between computers. SMB uses four message types: session control, file, printer, and message.

**Line Printer Remote (LPR)**
A connectivity tool that runs on client computers and that is used to print files to a computer running a Line Printer Daemon (LPD) server.

**Line Printer Daemon (LPD)**
A service on a print server that receives print jobs from Line Printer Remote (LPR) tools that are running on client computers.

**remote procedure call (RPC)**
A message-passing facility that allows a distributed application to call services that are available on various computers on a network. Used during remote administration of computers.

**Internetwork Packet Exchange (IPX)**
A network protocol native to NetWare that controls addressing and routing of packets within and between local area networks (LANs). IPX does not guarantee that a message will be complete (no lost packets).

**Internet Printing Protocol (IPP)**
The protocol that uses the Hypertext Transfer Protocol (HTTP) to send print jobs to printers throughout the world. Windows 2000, Windows XP, and Windows Server 2003 support Internet Printing Protocol (IPP) version 1.0.

**Transmission Control Protocol/Internet Protocol (TCP/IP)**
A set of networking protocols widely used on the Internet that provides communications across interconnected networks of computers with diverse hardware architectures and various operating systems. TCP/IP includes standards for how computers communicate and conventions for connecting networks and routing traffic.

**AppleTalk**
The set of network protocols on which AppleTalk network architecture is based. The AppleTalk Protocol is installed with Services for Macintosh to help users' access resources on a network.

**Simple Network Management Protocol (SNMP)**
A network protocol used to manage TCP/IP networks. In Windows, the SNMP service is used to provide status information about a host on a TCP/IP network.

**Network Printing Processes and Interactions**
This section briefly describes the network printing process on the client and then the server, followed by a more detailed description of the components and their activities. This section describes a fully functional network with all required software installed on the server and the clients. It also assumes that the print server is running Windows Server 2003.

**Print Client Processing**
The components and processes involved in printing vary depending on the client and the operating system, even within the Windows family of operating systems. Because of differences between print clients, the components used in the server process can also vary. Since the following description of the printing process is simplified, the print clients can be grouped into Windows and non-Windows platforms.

**Client Computers Running a Windows Operating System**
On client computers running Windows 2000 or Windows XP operating systems, the client spooler sends an enhanced metafile (EMF) through the client spooler to the spooler on the print server.

The client spooler forwards information and data from the GDI and the printer driver as an EMF to the spooler on the print server. The client spooler does this by making a remote procedure call (RPC) to the server spooler (Spoolsv.exe), which makes a direct application programming interface (API) call to the print router (Spoolss.dll). The print router polls the remote print providers, and the remote print provider (Win32spl.dll) makes an RPC to the server spooler, which receives the print job over the network.

On clients running Windows 95, Windows 98, or Windows Millennium Edition, a Windows-based application creates an EMF, which is then rendered to a page description language (PDL). Next, the print job is sent to a local printer port, where it would go if printing to a local printer, by the Server Message Block (SMB) redirector. The redirector delivers the job to the print server service running on the print server.

On clients running Windows NT 4.0, applications create output in a page-description language (PDL). Next, the PDL data is sent with an RPC over SMB to the server spooler.

**Client Computers Running a Non-Windows Operating System**
On client computers running a non-Windows operating system, or using an application that is designed for a non-Windows operating system, a software component performs functions similar to that of the Graphics Device Interface (GDI). Next, the device information and print data are sent to the print server. Windows Server 2003 can accept print jobs sent from a client using Line Printer Remote (LPR), AppleTalk, and NetWare when the server is running, respectively, the Line Printer Daemon (LPD) service, Print Server for Macintosh, or File and Print Services for NetWare. Typically, such clients would be running UNIX, Apple Macintosh, and NetWare, though additional software could enable them to use a protocol associated with a different system. Also, a client might send a print job by using Internet Printing Protocol (IPP).

**Print Server Processing**
On the print server, the service that initially receives incoming print jobs depends on the protocol used by the client to submit the print job. Print jobs sent by way of AppleTalk are received by the Service for Macintosh component, print jobs sent by way of LPR are received by the LPD service, and print jobs sent by way of SMB are received by the Server service.

However, print jobs from most client computers running Windows are received by the Spooler service by way of RPC. This includes print jobs from client computers running Windows 2000 and Windows XP, and it is print jobs from such clients that this section describes.

The Spooler service, Spoolsv.exe, is an executable file that is installed as a service. The spooler is loaded when the operating system starts, and it continues to run until the operating system is shut down. The Spooler service manages the printing process, which includes such tasks as retrieving the location of the correct printer driver, loading that driver, spooling high-level function calls into a print job, and scheduling the print job for printing. When the tasks for a particular print job are complete, the Spooler service passes the job to the router.

The spooler router, Spoolss.dll, determines which print provider to call, based on a printer name or other information provided with the print job. If the print job is for a printer served by the local server, the provider is the local print provider, Localspl.dll; if the print job is for a printer served by a remote server, the job is routed to a remote print provider.

The print provider is the software that supports the specified print device. The local print provider, Localspl.dll, actually calls the printer driver to render the print job. The print provider is also responsible for spooling the print job.

After a print job is spooled, it is the Print Monitor that directs the output to the appropriate port monitor. Port monitors provide a communications path between the print spooler and the port drivers that access the port hardware. Port monitors serve both local print devices connected directly to the server and remote print devices connected directly to the network.

In addition to the port monitor, when the target print device supports two-way communication, a language monitor provides a communication path for configuration and status information. The language monitor supports delivery of unsolicited information from the printer, for example, notice that it has run out of paper.

**Printing Components in Detail**
The following sections cover in greater detail some of the components on the client and server involved in network printing.

**Remote Print Provider**
When the server spooler receives a print job from the client spooler, the print router queries the remote print providers, and gives control of the print job to the first provider that recognizes the destination printer. Remote print providers can be created by other companies; Microsoft supplies the following three remote print providers:

**Win32spl.dll**
Win32spl.dll is the remote print provider on a Windows-based network. Win32spl.dll sends jobs to print servers running Windows NT 4.0 or Windows for Workgroups. For a print server running Windows Server 2003, Win32spl.dll makes an RPC to the spooler on the server (Spoolsv.exe), which makes a call to the print router (Spoolss.dll). The print router receives the job over the network and passes it to the local print provider as if a local client had submitted it.

**Inetpp.dll**
Inetpp.dll performs specific processing on Internet print servers. For an Internet print server, Inetpp.dll recognizes the printer name by its URL and submits the print job to the appropriate printer by using IPP.

**Nwprovau.dll**
Nwprovau.dll performs specific processing on NetWare print servers. For a NetWare print server, Nwprovau.dll recognizes the printer name when the print router on the print server running Windows Server 2003 polls it. Nwprovau.dll takes control of the print job and sends a message to the NetWare Workstation Service, which in turn passes control to the NetWare redirector. The NetWare redirector sends the print job to the NetWare print server.

**Print Router**
When submitting a document for printing, the application communicates with the client side of the spooler (Winspool.drv). Winspool.drv makes an RPC connection to the server side of the spooler (Spoolsv.exe), which makes a direct API call to the print router (Spoolss.dll), which is also on the server side of the spooler. The print router passes the print job to the local print provider.

**Graphics Device Interface (GDI)**
The GDI calls the printer driver and provides information about the type of printer needed and the data type used. The printer driver responds by sending printer data to the GDI, which the GDI passes to the spooler.

Computers running Windows 2000, Windows XP Professional, and any 32-bit version of Windows Server 2003 can run 16-bit applications through a process called Windows on Windows (WOW). WOW allows 16-bit applications to run on an x86-based computer.

WOW interacts with the applications through the GDI and with printer drivers through the device driver interface (DDI). The GDI also provides services to the printer driver, including caching, client and server communications, and conversion from American National Standards Institute (ANSI) character set to Unicode.

### Note

- The 64-bit versions of Windows XP Professional and Windows Server 2003 do not support 16-bit applications. However, WOW 64 does allow 32-bit applications to run on 64-bit versions of Windows XP Professional and Windows Server 2003.

**Printer Drivers**
A printer driver is a software program that understands how to communicate with printers and plotters (devices used to draw charts, diagrams, and other line-based graphics). Printer drivers translate the information that an application sends through the

GDI into drawing commands that the printer understands. These drawing commands are for creating text and graphics. Various drivers must be installed on the print server to support different hardware and operating systems. For example, an administrator who uses a computer that is running Windows Server 2003 and who shares a printer with clients that are running Windows 95, Windows 98, and Windows Millennium Edition, might need to install the appropriate drivers so that clients that are running these versions of Windows are not prompted to install the missing drivers.

The printer driver sends information about the printer settings, including the specifications that are needed to produce each character of the document, to the GDI. The printer driver also transmits helper services or utilities required to make the output print correctly.

Some kernel-mode printer drivers might cause the system to stop responding. Windows Server 2003 gives you the ability to control the installation of printer drivers. By default, all editions of Windows Server 2003 block the installation of kernel-mode drivers, but you can allow their installation by disabling the Group Policy setting **Disallow installation of printers using Kernel Mode drivers**. This setting affects only the installation of new kernel-mode drivers. If you are upgrading your system, existing kernel-mode drivers are unaffected.

Windows Server 2003 provides the following three generic printer drivers:

- Universal (Unidriver)
- PostScript
- HP-GL/2 plotter

**Universal printer driver (Unidriver)**
The Unidriver is also called the raster driver because it supports raster (bitmap) graphics printing and is compatible with many printers. This driver supports the following features:

- Color printing at various depths, such as 4 bits per pixel (bpp), 8 bpp, and 24 bpp.
- Scalable TrueType and OpenType fonts, device fonts (including double-byte), grayscale printing, font substitution, run length encoding (RLE), Tag Image File Format (TIFF) version 4.0, and Delta Row Compression (DRC).
- An extension interface that allows printer manufacturers to customize the driver for specific models.
- Vector graphics languages HP-GL/2 and PCL-XL, monochrome, and color. This is a change from the version of Unidriver that shipped with Windows 2000. (The added HP-GL/2 support, which is used with normal PCL5e printers, is different from the HP-GL/2 for plotters, which is used only for plotters.)

The Unidriver contains the following component files:

- Unidrv.dll is the printer graphics driver file for printer languages based on raster images, including most inkjet and dot-matrix printer languages.
- Unidrvui.dll is the configuration file. It displays the user interface for Unidrv.dll.
- Unires.dll, Stdnames.gpd, Ttfsub.gpd and some other support files.
- Raster minidriver, which has the file name extension .gpd, is the data file, and is also called the characterization file. The file name depends on the printer or printer family.

**PostScript printer driver**
The Windows Server 2003 PostScript driver sets up the print jobs for printers that have an onboard PostScript interpreter. This driver automatically supports key features — including binary transfer compression, resolution, and paper source — for printers reporting as PostScript levels 1, 2, and 3. Its output is fully compliant with Document Structuring Convention (DSC) 3.0, a structural specification for PostScript documents.

This driver contains the following component files:

- Pscript5.dll is the driver file.

- Ps5ui.dll is the configuration file. It displays the user interface for the driver file, reports device capabilities to applications, and handles DevMode settings. Ps5ui.dll also allows you to enable TrueType and OpenType font substitution and image color matching and to determine whether color matching needs to be done on the host or on the printer.

- The data file, which has the file name extension .ppd, is also known as a characterization file. The file name preceding the .ppd extension varies and is typically determined by the printer manufacturer. PPD files are the only printer driver files that are generally binary-compatible across processors and platforms.

**HP-GL/2 plotter driver**

The Windows Server 2003 plotter driver supports diverse plotters that use the HP-GL/2 language; however, Windows Server 2003 does not support HP-GL. The output from an HP-GL/2 driver requires a plotting device that can process all the enhancements built into the HP-GL/2 language.

This driver contains the following component files:

- Plotter.dll is the driver file.

- Plotui.dll is the configuration file. It displays the user interface for Plotter.dll.

- The driver contains a data file, which has the file name extension .pcd. The file name preceding the .pcd extension varies and is typically determined by the plotter manufacturer.

**Local Print Provider**

The local print provider, Localspl.dll, writes the job data to a spool (.spl) file. It also tracks administrative information, such as user name, document name, and data type, in a shadow (.shd) file. Spooling protects a print job by saving it on disk. If a power failure or other disaster occurs before all jobs in the queue are printed, the .spl and .shd files preserve the documents and prevent loss of data after processing resumes.

### Note

- If the print job is only partially saved to disk, the spooler might recognize the print job as corrupt and delete it upon restart.

Next, Localspl.dll polls the print processors for recognition of the document data type. If no data type has been set, Localspl.dll receives the job and uses the default data type.

By default, .spl and .shd files are written to the folder *systemroot*\System32\Spool\Printers, and they are deleted after the job prints. If the hard disk partition containing Windows Server 2003 does not have enough space for these files, you can change the location of the folder. You might also move the spool files if you want to move the spooling load to a separate disk array.

### Note

- Do not attempt to spool to the root of a disk, such as D:\. This causes the files to spool to the old default spool folder.

**Print Processor**

The print processor is a part of the spooler process. If the job type is EMF, the print processor works with GDI to move the spooled print jobs from the hard disk to the printer driver. If the data type is RAW, the print processor helps in moving the job directly to the printer. The print processor controls certain features of the print job, depending on the kind of job. For example, if the data type of the spooled job is EMF, the print processor can help in printing the document in reverse order or it can print it in booklet form instead of the default one-page-per-sheet form. Winprint.dll is the name of the default print processor on Windows Server 2003. Winprint.dll is part of Localspl.dll.

Software vendors can develop their own print processors for custom data types. A printer vendor can also develop a custom print processor if the vendor has its own printer driver or supports a data type other than the three that WinPrint.dll supports. Also, a vendor can write a custom print processor if the vendor wants to add functionality that Winprint.dll does not provide. Usually, the print processor is installed when the printer driver is installed.

**Print Monitors**

Windows Server 2003 supports two kinds of print monitors:

- Language monitors
- Port monitors

  Port monitors are further subdivided into:

  - Local port monitors
  - Remote port monitors

**Language monitor**
The language monitor provides the common language that is needed for the client and printer to understand each other in bidirectional communication so that you can configure the printer and monitor printer status. You can request configuration and status from the printer, and the printer can send unsolicited status information (such as "Paper tray empty") to the client.

Windows Server 2003 includes Pjlmon.dll, a language monitor that uses Printer Job Language (PJL). PJL is the language that implements the bidirectional communication; it does this by using a parallel port connection, such as one between a Hewlett-Packard LaserJet 5Si and the computer. Any bidirectional printer that uses PJL can use Pjlmon.dll.

If a printer uses a different printer language, the vendor can develop a language monitor for it. A vendor might also develop a language monitor to add data, such as printer-specific control information, to the print stream that is going to a unidirectional printer.

To take advantage of bidirectional printing, you need a bidirectional printer, and a correctly configured bidirectional port and port monitor. The standard port monitor supports bidirectional communication by default.

**Local port monitor**
The local port monitor controls parallel and serial input/output (I/O) ports where a printer might be attached. The local port monitor sends print jobs to local devices, including those on familiar ports such as LPT1, COM1, or FILE (Print to File).

A second local port monitor, Usbmon.dll, controls communication through universal serial bus (USB) ports. Usbmon.dll installs automatically whenever you plug a USB printer into the correct physical port on your computer. If USB has been enabled in the BIOS, Windows XP Professional detects the device and displays its settings on the screen.

Installation opens the USB Root Hub and a generic USB printer port (for example, USB001, USB002), and copies the printer Plug and Play driver.

**Remote port monitors**
All other port monitors that are supplied with Windows Server 2003 are remote monitors that enable printing to remote printers. An example is the standard TCP/IP port monitor.

**Network Ports Used by Network Printing**
The following table lists services used in network printing and the protocols they provide in the network printing process. Not all services are installed or active at all times, and the protocols used in network printing depend on the type of clients submitting jobs to the print server.

### Port Assignments for Network Printing

| Service Name | UDP | TCP |
|---|---|---|
| Services for Macintosh (uses AppleTalk) | 201, 202, 204, 206 | 201, 202, 204, 206 |
| Internet Printing Protocol (uses HTTP) | 80 | 80 |
| NWLink (uses IPX/SPX) | 213 | 213 |
| LPDSVC (accepts LPR) | 515 | 515 |
| LPR Port (sends LPR) | 721-731, >1023 | 721-731, >1023 |
| Server (using SMB) | 137, 138 | 139 |
| Server (using RPC) | 530 | 530 |

| | | |
|---|---|---|
| Standard Port Monitor (using SNMP) | 161, 162 | 161, 162 |
| Standard Port Monitor (using LPR) | > 1023 | > 1023 |
| Standard Port Monitor (using RAW) | 2000, 2501-2503, 3001, 6869, >=9100 | 2000, 2501, 2503, 3001, 6869, >=9100 |

Standard LPR, according to RFC 1179 from the Internet Engineering Task Force, uses ports 721 through 731 to send data. However, Windows Server 2003 can be configured on a per-port or per-server basis to use port number higher than 1023 for sending data with LPR.

Also different from RFC 1179 is the standard port monitor, which sends documents to a printer by using either the RAW or LPR printing protocol. Together, these protocols support most current TCP/IP printers.

The RAW protocol is the default print protocol for most print devices. To send a job formatted with RAW, the print server opens a TCP stream to the printer's port 9100 or another port number and selects the connections to multiport external devices. For example, on certain print devices port 9101 goes to the first parallel port, 9102 goes to the second parallel port, and so on.

The standard port monitor uses the LPR protocol when you specify LPR protocol during port installation or reconfiguration, or when the RAW protocol cannot be established.

The standard port monitor deviates from the LPR standard in the following ways:

- The standard port monitor does not conform to the RFC 1179 requirement that the source TCP port must lie between port 721 and port 731. The standard port monitor uses ports from the general, unreserved pool of ports (1024 and above).

- The standard port monitor sends a print job to the spooler without determining the actual job size and designates a default job size, regardless of the actual size of the job. The LPR standard states that print jobs must include information about the size of the job the port monitor sends. Sending a print job that contains job size information requires that the port monitor spool the job twice — once to determine size, and once to send the job to the spooler. Spooling the job once, by using the standard port monitor, improves printing performance.

**Related Information**

The following resource contains additional information that is relevant to this section.

- RFC 1179 in the IEFC RFC Database

# Network Printing Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Printing Tools and Settings**

**In this section**


- Network Printing Tools
- Network Printing Registry Entries
- Network Printing Group Policy Settings
- Network Printing WMI Classes
- Network Ports Used by Network Printing
- Related Information


**Network Printing Tools**

The following tools are associated with Network Printing.

**Prndrvr.vbs**

**Category**

Operating system tool.

**Version compatibility**

Runs on Windows Server 2003. Cannot run on or target computers running earlier versions of Windows.

You can use the Prndrvr.vbs script to administer printer drivers, for example, use this script to add, delete, and list drivers, either locally or remotely.

Prndrvr.vbs is written in VBScript and uses the Windows Management Instrumentation (WMI) print provider to administer the printing sub-system.

**Prnjobs.vbs**

**Category**

Operating system tool.

**Version compatibility**

Runs on Windows Server 2003. Cannot run on or target computers running earlier versions of Windows.

The Prnjobs.vbs script controls print jobs. With this script, you can pause, resume, cancel, and list print jobs, either locally or remotely.

Prnjobs.vbs is written in VBScript and uses the Windows Management Instrumentation (WMI) print provider to administer the printing sub-system.

**Prncnfg.vbs**

**Category**

Operating system tool.

**Version compatibility**

Runs on Windows Server 2003. Cannot run on or target computers running earlier versions of Windows.

You can use the Prncnfg.vbs script to get and set printer configuration or to rename a printer, either locally or remotely.

Prncnfg.vbs is written in VBScript and uses the Windows Management Instrumentation (WMI) print provider to administer the printing sub-system.

**Prnqctl.vbs**

**Category**

Operating system tool.

**Version compatibility**
Runs on Windows Server 2003. Cannot run on or target computers running earlier versions of Windows.

The Prnqctl.vbs script is used to print a test page, pause or resume a printer, and clear a printer queue, either locally or remotely.

Prnqctl.vbs is written in VBScript and uses the Windows Management Instrumentation (WMI) print provider to administer the printing sub-system.

**Prnmngr.vbs**
**Category**
Operating system tool.

**Version compatibility**
Runs on Windows Server 2003. Cannot run on or target computers running earlier versions of Windows.

With the Prnmngr.vbs script, you can add and delete printers, and list printers and printer connections. You can also use it for getting and setting the default printer.

Prnmngr.vbs is written in VBScript and uses the Windows Management Instrumentation (WMI) print provider to administer the printing sub-system.

**Prnport.vbs**
**Category**
Operating system tool.

**Version compatibility**
Runs on Windows Server 2003. Cannot run on or target computers running earlier versions of Windows.

Use the Prnport.vbs script to add, delete, and list Standard TCP/IP ports. You can also use it for getting and setting the port configuration.

Prnport.vbs is written in VBScript and uses the Windows Management Instrumentation (WMI) print provider to administer the printing sub-system.

**PrintUI.dll**
**Category**
Operating system tool.

**Version compatibility**
Runs on Windows 2000, Windows XP, and Windows Server 2003.

PrintUI.dll allows you to modify or delete local printers. It also allows you to add printer connections for users based on their computer names rather than by having users add their own printers. As a result, you can control most of the end user's printing experience by running a local or remote batch file or a logon script that contains PrintUI.dll commands.

You are encouraged to use the VBScripts listed in the WMI Print Provider or the PrnAdmin sections instead of using this utility.

**PrnAdmin**
**Category**
Resource Kit tool.

**Version compatibility**
The PrnAdmin utility lets you use scripts that perform administrative printer functions. You can create your own scripts or modify one of the sample scripts that are provided with PrnAdmin. Either way, you can manage both local and remote printers by using PrnAdmin.

PrnAdmin is written in VBScript, and its file name is Prnmgr.vbs.

PrnAdmin performs these functions:

- Adds or deletes the following:
    - Local or remote printers.
    - Printer connections.

    - Local or remote forms.
    - Ports on local or remote servers.
    - Printer drivers on local or remote servers.
- Enumerates ports, drivers, printers, or forms on local or remote printers.
- Pauses, resumes, or clears the print jobs of a local or remote printer.
- Configures a local or remote printer.

To find more information about PrnAdmin, see "Resource Kit Tools Help" in the Tools and Settings Collection.

**Network Printing Registry Entries**
The Windows platforms that use these registry entries and subkeys are noted in the **Version** portion of the listings. Many registry elements that are noted as being for Windows XP and Windows Server 2003 might also work on older versions of Windows; check documentation specifically for those versions for more details.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following registry entries and subkeys are associated with Network Printing.

For more information about registry entries, see the "Registry Reference" in the Tools and Settings Collection.

**CurrentControlSet\Control\Print**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Print.

**BrowsePrintWorkstations**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Version**
Windows XP and Windows Server 2003.

Used to calculate the number of clients in whose browse list the printer appears.

**EnableBroadcastSpoolerStatus**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Version**
Windows XP and Windows Server 2003.

Specifies whether the system sends spooler and job status messages when the job count changes.

**MaxRpcSize**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Version**
Windows XP and Windows Server 2003.

Specifies the maximum size, in bytes, of incoming data blocks from a remote procedure call (RPC).

**Monitors**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Version**
Windows XP and Windows Server 2003.

The Monitors subkey stores entries for print monitors installed on Windows Server 2003. The Monitors subkey contains a subkey for each port monitor that stores configuration data for the port monitor installed on the system.

**NetPrinterDecayPeriod**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Version**
Windows XP and Windows Server 2003.

Specifies how long a network printer remains in the browse list.

**PendingUpgrades**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Version**
Windows XP and Windows Server 2003.

The PendingUpgrades subkey is used by Windows Server 2003 during setup.

**Providers**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Version**
Windows XP and Windows Server 2003.

The Providers subkey contains information about the network printers. The printers listed in this subkey are not hosted locally, but they are printers to which the local computer is connected.

**ServerThreadTimeout**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Version**
Windows XP and Windows Server 2003.

Specifies the amount of time that the server thread is idle before it notifies other print servers of its printers.

**SpoolerPriority**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print

**Version**
Windows XP and Windows Server 2003.

Specifies the priority class of the print spooler process. The priority of the print spooler process establishes a range of priority for print spooler threads.

**environment-name**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Environments

**Version**
Windows XP and Windows Server 2003.

The *environment-name* subkey represents a computing platform environment and stores entries that apply only to that environment.

**Version-number**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Environments*environment-name*\Drivers

**Version**
Windows XP and Windows Server 2003.

The Version-*number* subkey contains information about the drivers for a print environment. There is a different number for each version of the printer driver. Version-2 drivers are for Windows NT 4.0; Version-3 drivers are for Windows 2000 and Windows Server 2003.

**driver-name**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Environments*environment-name*\Drivers\Version-*number*

**Version**
Windows XP and Windows Server 2003.

The *driver-name* subkey contains information about necessary files and settings for the specified printer driver.

**Version**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Environments*environment-name*\Drivers\Version-*number**driver-name*

**Version**
Windows XP and Windows Server 2003.

Stores the version number of the installed printer driver specified by *driver-name* in the registry path.

**monitor**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Monitors

**Version**
Windows XP and Windows Server 2003.

The *monitor* subkey contains information about the status and drivers of the monitors used by the computer, and the actual name of the subkey corresponds to a unique monitor.

**Ports**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Monitors*monitor*

**Version**
Windows XP and Windows Server 2003.

The Ports subkey contains information about the current availability of the ports.

**StatusUpdateEnabled**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Monitors*monitor*\Ports

**Version**
Windows XP and Windows Server 2003.

Specifies whether the port monitor is able to update the spooler with the port status.

**StatusUpdateInterval**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Monitors*monitor*\Ports

**Version**
Windows XP and Windows Server 2003.

Specifies how often the port monitor updates the spooler with port status.

**EventLog**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers

**Version**
Windows XP and Windows Server 2003.

Specifies the level of detail of events that the print spooler logs for local print jobs.

**NetWare or Compatible Network**
**Registry path**

HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers

**Version**
Windows XP and Windows Server 2003.

Stores options for printing when using print resources on NetWare servers.

**RestartJobOnPoolEnabled**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers

**Version**
Windows XP and Windows Server 2003.

Specifies whether a print job to a local printer can be restarted after an error.

**RestartJobOnPoolError**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers

**Version**
Windows XP and Windows Server 2003.

Specifies how long the print monitor waits after an error occurs before starting a print job on another port in the print pool.

**DisplayName**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers\NetWare or Compatible Network

**Version**
Windows XP and Windows Server 2003.

Specifies the localized name of the print service provider displayed in the Print Manager and Printers main window.

**Name**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers\NetWare or Compatible Network

**Version**
Windows XP and Windows Server 2003.

Specifies the dynamic-link library (DLL) for the NetWare or Compatible Print Service provider. This DLL is stored in *systemroot*\Winnt\System32.

**Timeout**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Control\Print\Providers\NetWare or Compatible Network

**Version**
Windows XP and Windows Server 2003.

Specifies how often Windows Server 2003 updates the print queue of a NetWare server. The value of this entry specifies the time interval between automatic refresh updates.

**Microsoft\Windows NT\CurrentVersion\Print**
The following registry entries are located under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print.

**Printers**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print

**Version**
Windows XP and Windows Server 2003.

The Printers subkey contains information about the local printers.

**Providers**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print

**Version**
Windows XP and Windows Server 2003.

The Providers subkey contains information about the network printers. The printers listed in this subkey are not hosted locally, but the local computer is connected to them.

**printer-name**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers

**Version**
Windows XP and Windows Server 2003.

The *printer-name* subkey stores information about printers connected to the print server.

**DefaultSpoolDirectory**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers

**Version**
Windows XP and Windows Server 2003.

Specifies the folder holding spooled files that the system is printing on a local printer.

**PrinterDriverData**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers*printer-name*

**Version**
Windows XP and Windows Server 2003.

The PrinterDriverData subkey contains configuration data for the specified printer.

**SpoolDirectory**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers*printer-name*

**Version**
Windows XP and Windows Server 2003.

Stores a printer-specific location for spool files. The value of this entry takes precedence over the global default stored in the **DefaultSpoolDirectory** entry.

**Transparency**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers\printer-name\PrinterDriverData

**Version**
Windows XP and Windows Server 2003.

Specifies the background mode when the value of the registry entry **Winprint_TextNoCRTranslation** is set to 1. If the background mode is set to transparent, then the driver will not print the background of the document, regardless of color.

**Winprint_TextNoCRTranslation**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers*printer-name*\PrinterDriverData

**Version**
Windows XP and Windows Server 2003.

Enables or disables carriage return translation. By default, Windows Server 2003 translates a carriage return to a carriage return plus a line feed.

**Winprint_TextNoTranslation**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers*printer-name*\PrinterDriverData

**Version**
Windows XP and Windows Server 2003.

Enables or disables carriage return and line feed translation. By default, Windows Server 2003 translates a carriage return or a line feed to a carriage return plus a line feed.

**server-name**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers

**Version**
Windows XP and Windows Server 2003.

The *server-name* subkey stores information about printers connected to the print server specified by *server-name*.

**Printers**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*

**Version**
Windows XP and Windows Server 2003.

The Printers subkey contains information about the individual network printers to which the local computer is connected.

**Providers**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*

**Version**
Windows XP and Windows Server 2003.

The Providers subkey stores configuration data used by the specified server.

**EventLog**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*\Providers

**Version**
Windows XP and Windows Server 2003.

Specifies the level of detail of events that the print spooler logs for the specified server.

**RestartJobOnPoolEnabled**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*\Providers

**Version**
Windows XP and Windows Server 2003.

Specifies whether a print job on the print server, specified in the registry path by *server-name*, can be restarted after an error.

**RestartJobOnPoolError**
**Registry path**
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Providers\LanMan Print Services\Servers*server-name*\Providers

**Version**
Windows XP and Windows Server 2003.

Specifies how long the print monitor for the printer server, specified in the registry path by *server-name*, waits after an error occurs before starting a print job on another port in the print pool.

**Network Printing Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with Network Printing.

### Group Policy Settings Associated with Network Printing

| Group Policy Setting | Description |
|---|---|

| | |
|---|---|
| Allow printers to be published | Determines whether the computer's shared printers can be published in Active Directory. |
| Disallow installation of printers that use kernel-mode drivers | Determines whether printers using kernel-mode drivers can be installed on the local computer. If enabled, prevents a printer that is using kernel-mode drivers from being installed on the local computer. |
| Automatically publish new printers in Active Directory | Determines whether the Add Printer Wizard automatically publishes the computer's shared printers in Active Directory. |
| Allow pruning of published printers | Determines whether the domain controller can prune (delete from Active Directory) the printers published by this computer. |
| Printer browsing | Announces the presence of shared printers to print browse master servers when Active Directory is not available. |
| Prune printers that are not automatically republished | Determines whether the system prunes printers that are not automatically republished. (Applies to printers running on operating systems other than Windows 2000 or Windows Server 2003. Also applies to printers running on Windows 2000 or Windows Server 2003 that are published outside of their domain.) |
| Directory pruning interval | Specifies how often the pruning service on a domain controller contacts computers to verify that their printers are operational. |
| Directory pruning retry | Specifies how many times the pruning service on a domain controller repeats its attempt to contact a computer before pruning the computer's printers. |
| Directory pruning priority | Sets the priority of the pruning thread. |
| Check published state | Directs the system to periodically verify that printers published by this computer still appear in Active Directory. Also specifies how often the system repeats verification. |
| Web-based printing | Determines whether Internet printing is supported on this server. (If you disable this policy, Internet printing is not supported.) |
| Custom support URL in the Printers and Faxes folder | Adds a customized Web page link to the Printers and Faxes folder. By default, the Printers and Faxes folder includes a link to the Microsoft Support Web page or to the Web page for the vendor of the currently selected printer. You can use this policy to replace the default link with a link to your organization's Web page. |
| Computer location | Specifies the default location criteria used when searching for printers. |
| Pre-populate printer search location text | Enables the physical Location Tracking support feature of Windows Server 2003 printers. If you disable this policy or do not configure it, Location Tracking is disabled. Printer proximity is estimated based on IP address and subnet mask. |
| Allow print spooler to accept client connections | Determines how clients access the print server over the network. By default, this policy setting automatically enables connections upon creation of the first shared printer or clustered server spooler resource. |
| Point and Print restrictions | Restricts a client computer's ability to connect to and install a printer driver from print servers. By default, this policy setting allows managed clients to use Point and Print only with servers within their forest. An administrator can use this setting to add servers to the list of trusted print servers. Disabling this setting enables connections to any accessible print server. |

For more information about Group Policy settings, see the "Group Policy Settings Reference" in the Tools and Settings Collection.

**Network Printing WMI Classes**

The following table lists and describes the WMI classes that are associated with Network Printing.

**WMI Classes Associated with Network Printing**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_Printer | \root\cimv2 | Windows XP and Windows Server 2003 ; supported as read-only in Windows 2000 |
| Win32_PrinterConfiguration | \root\cimv2 | Windows XP and Windows Server 2003 |
| Win32_PrinterController | \root\cimv2 | Windows XP and Windows Server 2003 |
| Win32_PrinterDriver | \root\cimv2 | Windows XP and Windows Server 2003 |
| Win32_PrinterDriverDll | \root\cimv2 | Windows XP and Windows Server 2003 |
| Win32_PrinterSetting | \root\cimv2 | Windows XP and Windows Server 2003 |
| Win32_PrintJob | \root\cimv2 | Windows XP and Windows Server 2003 ; supported as read-only in Windows 2000 |
| Win32_TCPIPPrinterPort | \root\cimv2 | Windows XP and Windows Server 2003 |

For more information that describes many WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by Network Printing**

The following table lists services used in network printing and the protocols they provide in the network printing process. Not all services are installed or active at all times, and the protocols used in network printing depend on the type of clients submitting jobs to the print server.

**Port Assignments for Network Printing**

| Service Name | UDP | TCP |
|---|---|---|
| Services for Macintosh (uses AppleTalk) | 201, 202, 204, 206 | 201, 202, 204, 206 |
| Internet Printing Protocol (uses HTTP) | 80 | 80 |
| NWLink (uses IPX/SPX) | 213 | 213 |
| LPDSVC (accepts LPR) | 515 | 515 |
| LPR Port (sends LPR) | 721-731, >1023 | 721-731, >1023 |
| Server (using SMB) | 137, 138 | 139 |
| Server (using RPC) | 530 | 530 |
| Standard Port Monitor (using SNMP) | 161, 162 | 161, 162 |
| Standard Port Monitor (using LPR) | > 1023 | > 1023 |
| Standard Port Monitor (using RAW) | 2000, 2501-2503, 3001, 6869, >=9100 | 2000, 2501, 2503, 3001, 6869, >=9100 |

Standard LPR, according to RFC 1179 from the Internet Engineering Task Force, uses ports 721 through 731 to send data. However, Windows Server 2003 can be configured on a per-port or per-server basis to use port number higher than 1023 for sending data with LPR.

Also different from RFC 1179 is the standard port monitor, which sends documents to a printer by using either the RAW or LPR printing protocol. Together, these protocols support most current TCP/IP printers.

The RAW protocol is the default print protocol for most print devices. To send a job formatted with RAW, the print server opens a TCP stream to the printer's port 9100 or another port number and selects the connections to multiport external

devices. For example, on certain print devices port 9101 goes to the first parallel port, 9102 goes to the second parallel port, and so on.

The standard port monitor uses the LPR protocol when you specify LPR protocol during port installation or reconfiguration, or when the RAW protocol cannot be established.

The standard port monitor deviates from the LPR standard in the following ways:

- The standard port monitor does not conform to the RFC 1179 requirement that the source TCP port must lie between port 721 and port 731. The standard port monitor uses ports from the general, unreserved pool of ports (1024 and above).

- The standard port monitor sends a print job to the spooler without determining the actual job size and designates a default job size, regardless of the actual size of the job. The LPR standard states that print jobs must include information about the size of the job the port monitor sends. Sending a print job that contains job size information requires that the port monitor spool the job twice — once to determine size, and once to send the job to the spooler. Spooling the job once, by using the standard port monitor, improves printing performance.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Registry Reference" in the Tools and Settings Collection

- "Resource Kit Tools" in the Tools and Settings Collection

# Recovery and Reliability Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The recovery and reliability technologies in the Microsoft Windows Server 2003 family help you to prevent problems such as data loss and system instability during all phases of operation —from startup through shutdown. In the event of a problem, these technologies help you to recover and repair your data and operating system.

**Recovery and Reliability Technologies Components**
The recovery and reliability technologies can help you to protect your data, prevent server instability and diagnose performance problems.

**Volume Shadow Copy Service**
The Volume Shadow Copy service provides an infrastructure for creating consistent point-in-time copies known as shadow copies. The Volume Shadow Copy service can produce consistent shadow copies by coordinating with business applications, file-system services, backup applications, and storage hardware. Several features in Windows Server 2003 use the Volume Shadow Copy service, including Shadow Copies for Shared Folders and Backup.

**Backup**
The Backup utility helps you protect your data in the event your hard disk fails or files are accidentally erased due to hardware or storage media failure. By using Backup you can create a duplicate copy of the data on your hard disk and then archive it on another storage device, such as a hard disk or a tape. Backup takes advantage of the Volume Shadow Copy service to create an accurate copy of the contents of your hard drive, including any open files or files that are being used by the system.

**Shadow Copies for Shared Folders**
Shadow Copies for Shared Folders provides you with point-in-time copies of files that are located on shared resources, such as a file server. With Shadow Copies for Shared Folders, you can quickly recover deleted or changed files that are stored on the network without administrator assistance, which can increase productivity and reduce costs.

**Recovery Console**
Recovery Console is a character-mode environment that you can run from the Windows Server 2003 operating system CD or install on x86-based computers as a startup option. Using Recovery Console, you can enable and disable services, format drives, read and write data on a local drive (including drives formatted to use NTFS), and perform many other administrative tasks. Recovery Console is particularly useful if you need to repair your system by copying a file from a floppy disk or CD-ROM to your hard disk, or if you need to reconfigure a service that is preventing your computer from starting properly.

**Emergency Management Services**
Emergency Management supports remote management and system recovery for servers that are not accessible through a standard network connection, such as a Local Area Network (LAN) or the Internet. By using Emergency Management Services in conjunction with other software and hardware components, you can remotely perform tasks such as recovering unresponsive servers or powering servers off and on, even if you cannot access the servers through a conventional connection. Emergency Management Services can also help you to save space, energy, and the cost of hardware because it supports servers that are configured without a locally attached keyboard, monitor, or mouse. Emergency Management Services helps you to return servers to a state where it is possible to manage them using standard tools — Windows technologies that rely on a standard network connection.

**Shutdown Event Tracker**
Shutdown Event Tracker is a technology that provides you with a way to consistently track why IT professionals restart or shut down servers. It is enabled by default, and is a routine part of the computer shutdown process. Shutdown Event Tracker gathers the reasons that users give for restarting or shutting down servers. This data can help you to create a comprehensive picture of the performance and reliability of your system environment.

# Recovery Console Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovery Console Technical Reference**
Recovery Console is a character-mode environment that you can run from the Microsoft Windows Server 2003 operating system CD or install on x86-based computers as a startup option. You should use Recovery Console only when other startup options have failed to start your system. Recovery Console is particularly useful if you need to repair your system by copying a file from a floppy disk or CD-ROM to your hard disk, or if you need to reconfigure a service that is preventing your computer from starting properly.

In contrast to normal or Safe Mode, the Windows graphical user interface (GUI) is not available within Recovery Console. You should only use Recovery Console if you are an advanced user who can use basic commands to identify and locate problem drivers and files.

**In this subject**


- What Is Recovery Console?
- Recovery Console Tools and Settings

# What Is Recovery Console?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Recovery Console?**

**In this section**


- Related Information


Recovery Console is a character-mode environment that you can run from the Microsoft Windows Server 2003 operating system CD or install on x86-based computers as a startup option. You should use Recovery Console only when other startup options have failed to start your system.

In contrast to normal or Safe Mode, the Windows graphical user interface (GUI) is not available within Recovery Console. You should only use Recovery Console if you are an advanced user who can use basic commands to identify and locate problem drivers and files. In addition, you will need the password for the built-in Administrator account to use Recovery Console.

With Recovery Console, you can enable and disable services, format drives, read and write data on a local drive (including drives formatted to use the NTFS file system), and perform many other administrative tasks. Recovery Console is particularly useful if you need to repair your system by copying a file from a floppy disk or CD-ROM to your hard disk, or if you need to reconfigure a service that is preventing your computer from starting properly.

You can use Recovery Console to help you recover from the following problems:

- Corrupted or deleted startup files caused by incompatible software, user error, or virus activity.
- Disk problems related to damage to the master boot record (MBR), partition table, or boot sector on x86-based systems.
- A partition boot sector that has been overwritten by the setup program.

If critical system files, such as Ntldr.exe or Ntoskrnl.exe, are missing or corrupted, you can restore them by starting Recovery Console and copying undamaged files from the Windows Server 2003 operating system CD or other removable disk media.

Boot sector damage can be caused by incompatible software, hardware problems, virus activity, or attempting to configure your computer as a multiple-boot system. For example, setup programs for other operating systems might not be compatible with Windows Server 2003 and might attempt to overwrite the boot sector or startup files. To prevent problems, avoid configuring multiple-boot configurations on computers running Windows Server 2003.

### Note

- Before you install Recovery Console, you need to be aware of certain disk and file system limitations. Recovery Console is sensitive to file-system changes. If you install Recovery Console to a hard disk that uses the FAT file system, converting to the NTFS file system causes Recovery Console to stop functioning. You must reinstall Recovery Console after converting to NTFS.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Backup Technical Reference
- Emergency Management Services Technical Reference

# Recovery Console Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovery Console Tools and Settings**

**In this section**


- Recovery Console Tools
- Recovery Console Group Policy Settings
- Related Information


This section lists the tools and the syntax that you can use with Recovery Console. Recovery Console is a character-mode environment that you can run from the Windows Server 2003 operating system CD or install on x86-based computers as a startup option. You should use Recovery Console only when other startup options have failed to start your system. You will need the password for the built-in Administrator account to use Recovery Console

You can use the Recovery Console Group Policy settings to enable the **set** command to customize Recovery Console or to enable logon to the console without the administrator password.

**Recovery Console Tools**

The following tools are associated with Recovery Console. When using Recovery Console, you can view and reuse previous commands by pressing the UP ARROW and DOWN ARROW keys, which move you forward or backward through your command history.

For the list of Recovery Console commands that follow, brackets ([]) enclose optional parameters and a pipe (|) separates mutually-exclusive choices. Recovery Console commands and parameters are not case sensitive.

**Attrib**

Use the **attrib** command to change the file attributes for a single file or folder. Use the following syntax:

```
attrib
-|+[c][h][r][s] [drive:][path]filename
```

| Parameter | Description |
|---|---|
| + | Sets an attribute. |
| - | Clears an attribute. |
| c | Sets or clears a compressed file attribute. |
| h | Sets or clears a hidden file attribute. |
| r | Sets or clears a read-only file attribute. |
| s | Sets or clears a system file attribute. |
| drive: | Specifies the drive letter to use. |
| path: | Specifies the directory path to use. |

You should also follow these guidelines when you use the **attrib** command:

- You must set or clear at least one attribute.
- You can change attributes for only one file or directory at a time.
- You can view attributes by using the **dir** command.

- Do not separate attribute parameters with spaces.
- You can set multiple attributes simultaneously.
    - To change multiple attributes in the same way, use either the set or clear parameter (+ or –). Include all the attribute options to be changed, and do not separate them with spaces. For example, to set the compressed, hidden, and read-only attributes for a single file, type the following:

      ```
      attrib +chr filename
      ```

    - To change multiple attributes in different ways, use the set parameter (+) and include all the attribute letters to be set, followed without a space by the clear parameter (–) and all the attribute letters to be cleared. For example, to set the compressed and hidden file attributes and to clear the read-only file attribute, type:

      ```
      attrib +ch-r filename
      ```

**Batch**

Use the **batch** command to run the commands specified in a text file. Use the following syntax:

```
batch
inputfile [outputfile]
```

| Parameter | Description |
|---|---|
| `inputfile` | Specifies the text file (by using [*drive*:][*path*][*filename*] format) that contains the list of commands you want to carry out. |
| `outputfile` | If specified, stores the output of the **batch** command in the specified file. If you do not specify a value for *outputfile*, the **batch** command displays its output on the screen. Specify *outputfile* by using [*drive*:][*path*][*filename*] format |

The **batch** command cannot call itself recursively. Do not include the **batch** command in the file specified by the *inputfile* parameter.

**Bootcfg**

For x86-based systems, use the **bootcfg** command to scan your hard disks and use the information to modify the contents of the Boot.ini file or rebuild a new copy. Always back up the Boot.ini file before modifying it. Use the following syntax:

```
bootcfg [/add] [/default]| [/list] [/rebuild] [/scan]
```

| Parameter | Description |
|---|---|
| `/add` | Adds a Windows installation to the operating system boot menu list. |
| `/default` | Sets the default boot menu. |
| `/disableems` | Disables redirection in the Windows loader |
| `/ems` | Enables redirection in the Windows loader for a specific configuration. |
| `/list` | Lists the entries already in the boot menu list. |
| `/rebuild` | Scans hard disks for Windows installations and selects which to add. |
| `/scan` | Scans all disks for Windows installations and displays the results. |

**Note**

- The Recovery Console **bootcfg** command is not the same as the Windows Server 2003 Bootcfg.exe command-line tool. Bootcfg.exe resides in the *systemroot*\System32 folder and is a stand-alone command-line tool that you cannot use in

**Cd or Chdir**

Use the **cd** or **chdir** command to display the name of the current volume or directory, or to change to the folder specified. The **chdir** command only operates within the system directories of the current Windows installation, removable media, the root directory of any hard disk partition, or the local installation sources. Use the following syntax:

```
cd [path]|[..]|[drive:]
```

–or–

```
chdir [path]|[..]|[drive:]
```

If you want to display the current volume and folder, use the **cd** or **chdir** command without parameters.

| Parameter | Description |
|---|---|
| `path` | Specifies the directory that you want to change to. |
| `..` | Displays the parent folder. |
| `drive:` | Specifies the drive that you want to change to. |

These commands treat spaces as delimiters, which require that a space is inserted before all arguments, including double periods. Use quotation marks to enclose a path or file name that contains a space.

**Chkdsk**

Use the **chkdsk** command to check a volume, and if needed, to repair the volume. Also, use **chkdsk** to recover and move readable information before marking bad sectors as unusable. Use the following syntax:

```
chkdsk [drive:] [/p]|[/r]
```

You can use the **chkdsk** command without parameters. When you do not specify a volume, **chkdsk** runs on the current volume.

| Parameter | Description |
|---|---|
| `drive:` | Specifies the volume that you want **chkdsk** to check. |
| `/p` | Performs an exhaustive volume check. This parameter does not make any changes to the volume. |
| `/r` | Locates bad sectors and recovers readable information before marking them as unusable. Implies /p. |

Chkdsk requires the file Autochk.exe. If **chkdsk** cannot find Autochk.exe in the *systemroot*\System32 directory, it attempts to locate Autochk.exe on the Windows Server 2003 installation CD.

**Cls**

Use the **cls** command to clear the screen and redisplay the command prompt. Use the following syntax:

**cls**

**Copy**

Use the **copy** command to copy a single file to a specified location. Use the following syntax:

```
copy
source
destination
```

| Parameter | Description |
|---|---|
| `source` | Specifies the file (by using [*drive*:][*path*][*filename*] format) that you want copied. |

| | |
|---|---|
| `destination` | Specifies the destination (to by using [*drive*:][*path*][*filename*] format) that you want the file *source* file copied format. |

You should also follow these guidelines when you use the **copy** command:

- You cannot use wildcard characters (for example: * and ?) with the **copy** command.
- If you do not specify a destination directory, the **copy** command uses the current folder by default.
- If you do not specify a destination file name, the **copy** command uses the existing file name by default.
- If the *destination* file name already exists, you are warned before overwriting it.
- Compressed files from the Windows Server 2003 operating system CD are automatically expanded as they are copied.
- The source can be removable media, any directory within the system directories of the current Windows installation, the root of any drive, the local installation sources, or the Cmdcons directory.
- The destination can be any of the same locations as the source except removable media. If a destination is not specified, it defaults to the current directory.

**Del or Delete**

Use the **del** or **delete** command to delete a file or folder. Use the following syntax:

```
del [drive:][path]filename
```

–or–

```
delete [drive:][path]filename
```

| Parameter | Description |
|---|---|
| `drive:` | Specifies the volume of the file that you want to delete. |
| `path` | Specifies the directory of the file that you want to delete. |
| `filename` | Specifies the file that you want to delete. |

You cannot use wildcard characters with this command. The **delete** command only operates within the system directories of the current Windows installation, removable media, the root directory of any hard disk partition, or the local installation sources.

**Dir**

Use the **dir** command to display a list of the files and folders in a directory and to specify a particular file or group of files for which you want to see a listing. Multiple file names can be used. File names can be separated by spaces, commas, or semicolons. You can use wildcard characters (? and *) in the *filename* parameter to display a group of files.

**Dir** also lists the disk's volume label and serial number, and the total number of files listed, their cumulative size, and the free space (in bytes) remaining on the disk. For each file and subdirectory, **dir** lists the file name extension, the file size in bytes, the date and time the file was last modified, and the following file attributes, if applicable:

Use the following syntax:

```
dir [drive:][path][filename]
```

| Parameter | Description |
|---|---|
| `drive:` | Specifies the volume of the directory for which you want a listing. |

| | |
|---|---|
| `path` | Specifies the directory for which you want a listing. |
| `filename` | Specifies the file for which you want a listing. |

In Recovery Console, the **dir** command functions differently that the command line because it lists all folders and files, including those with hidden and system attributes set. For each file and subdirectory, the **dir** command lists its attributes (if they apply) by using the following abbreviations.

**a**
Archive

**c**
Compressed

**d**
Directory

**e**
Encrypted

**h**
Hidden

**p**
Reparse point

**r**
Read-only

**s**
System file

**Disable**
Use the **disable** command to disable a service or driver. Use the following syntax:

```
disable
servicename
```

| Parameter | Description |
|---|---|
| `servicename` | Specifies the service or driver that you want to disable. |

Use a related command, **listsvc,** to view a list of service and driver names for your system. The **disable** command displays the previous start type of the service before changing it to SERVICE_DISABLED. Record this value so that you can restore the original state of the service after troubleshooting a problem.

There are five startup types. The first three, SERVICE_AUTO_START, SERVICE_DISABLED, and SERVICE_DEMAND_START, correspond to the standard startup types, Automatic, Disabled, and Manual, which you typically configure using the Services snap-in in the Computer Management administrative tool. The last two, SERVICE_BOOT_START and SERVICE_SYSTEM_START, are commonly used to configure the way device drivers load, for example, when Windows starts.

**Diskpart**
Use the **diskpart** command to manage the partitions on your hard disk. For example, you can use **diskpart** to create or delete disk partitions. Use the following syntax:

```
diskpart[/add|/delete] [device-name|drive-name|partition-name] [size]
```

| Parameter | Description |
|---|---|
| `/add` | Creates a new disk partition. |
| `/delete` | Deletes an existing partition. |

| | |
|---|---|
| device-name | Specifies the name of the device for which you want to create or delete a partition, for example, \Device\HardDisk0. To obtain the name of a device, view the output of the **map** command. |
| drive-name | Specifies the drive letter of the partition that you want to delete, for example, D:. Use only with **/delete**. |
| partition-name | Specifies the partition that you want to delete and can be used in place of the *drive-name* parameter. For example,\Device\HardDisk0. Use only with **/delete**. |
| size | Specifies the size, in megabytes, of the partition you want to create. Use only with **/add**. |

If you do not use a parameter, a user interface for managing your partitions appears.

**Note**

- This command can damage your partition table if the disk has been upgraded to a dynamic disk. Do not modify the structure of dynamic disks unless you are using the Disk Management snap-in.

**Enable**

Use the **enable** command to enable or change the startup type of a service or driver. Use the following syntax:

```
enable
servicename [start_type]
```

| Parameter | Description |
|---|---|
| servicename | Specifies the service or driver that you want to enable. |
| start_type | Specifies the startup type for a service or driver. Valid values are:<br>• SERVICE_BOOT_START<br>• SERVICE_SYSTEM_START<br>• SERVICE_AUTO_START<br>• SERVICE_DEMAND_START |

Use the related command, **listsvc,** to view a list of service and driver names for your system. The **enable** command displays the previous start type of the service before changing it. Record this value so that you can restore the original state of the service after troubleshooting a problem.

If you do not specify a new start type, the enable command displays the previous start type.

**Exit**

Use the **exit** command to close Recovery Console and restart your computer. Use the following syntax:

**exit**

**Expand**

Use the **expand** command to expand a compressed file that is stored on the Windows Server 2003 operating system CD or from within a cabinet (.cab) file, and copy it to a specified destination. If the destination is not specified, the file is copied to the current directory. Use the following syntax:

```
expand
source [/f:filespec] [target] [/y]
expandsource [/f:filespec] /d
```

| Parameter | Description |
|---|---|
| source | Specifies the file that you want to expand (by using [*drive*:][*path*][*filename*] format). You cannot use wildcard characters (* and ?). |

| | |
|---|---|
| `target` | Specifies the destination folder and/or file name for the new file by using [*drive*:][*path*][*filename*] format. |
| `/f:filespec` | Specifies the specific file(s) that you want to expand if the source contains more than one file. Wildcards are optional. |
| `/y` | Specifies that the confirmation prompt (that appears when attempting to overwrite an existing file) is not required. |
| `/d` | Specifies that files display, but does not expand the files in the .cab file. |

You cannot specify removable media such as a floppy disk drive or a CD-ROM as the destination.

**Fixboot**

Use the **fixboot** command to rewrite the boot sector code to a hard disk. This is useful for repairing a corrupted boot sector on x86-based systems. If you need to replace the boot sector of a volume that is not the system volume, then you must specify the appropriate drive letter. Use the following syntax:

```
fixboot [drive:]
```

| Parameter | Description |
|---|---|
| `drive:` | Specifies the volume drive letter on which to rewrite a new boot sector. |

If you do not specify a drive, the default is the system boot volume.

**Fixmbr**

Use the **fixmbr** command to rewrite the master boot code of the master boot record (MBR) of the startup hard disk. This command is useful for repairing corrupted MBRs. Use the following syntax:

```
fixboot [device-name]
```

| Parameter | Description |
|---|---|
| `device-name` | Specifies the name of the device that needs a new MBR, for example, \Device\HardDisk1. |

If you do not specify a device, the default is disk 0. If disk 0 is not the device that needs repairing, you can obtain the device name of other disks by using the **map** command.

If the **fixmbr** command detects an invalid or nonstandard partition table signature, it prompts you for permission before rewriting the MBR.

Use this command with care because it can damage your partition table if any of the following applies:

- A virus is present and a non-Microsoft operating system is installed on the same computer.
- A nonstandard MBR is installed by a non-Microsoft disk utility.
- A hardware problem exists.

Always run antivirus software before using this command.

When you run the **fixmbr** command, it only overwrites the master boot code and leaves the existing partition table intact. If corruption in the MBR affects the partition table, running the **fixmbr** command might not resolve the problem.

**Format**

Use the **format** command to format the specified volume to the specified file system. Use the following syntax:

```
format [drive:] [/q] [/fs:file_system]
```

| Parameter | Description |
|---|---|

| | |
|---|---|
| `drive:` | Specifies the drive letter for the volume you want to format. |
| `/q` | Specifies a quick format (clears only the table of contents). |
| `/fs:file-system` | Specifies the file system you want to use. Valid values for *file-system* include FAT, FAT32, and NTFS. |

Consider the following guidelines before using the **format** command:

- If a file system is not specified, the **format** command defaults to the NTFS file system.
- If you choosing FAT for the file system, this command formats a volume as FAT16. FAT16 volumes cannot be larger than 4 GB. Limit FAT16 partitions to2 GB to increase storage efficiency and to maintain compatibility with MS-DOS, Windows 95, Windows 98, and Windows Millennium Edition.
- Windows Server 2003 can format FAT32 volumes up to 32 GB in size. For larger volumes, use NTFS.

**Help**

Use the **help** command to view Help information for Recovery Console commands. Use the following syntax:

```
help [command]
```

| Parameter | Description |
|---|---|
| `Command` | Specifies the drive letter for the volume you want to format. |

Use the *command* parameter to specify a name of any Recovery Console command.

If you do not specify a parameter, Help lists information about all the supported commands.

**Listsvc**

Use the **listsvc** command to view details about the services and drivers on your system, including service start types. Use the following syntax:

```
listsvc
```

Use the **listsvc** command together with the **disable** and **enable** commands. The information displayed is extracted from the registry file System that is located in the *systemroot*\System32\Config folder. If the file System is damaged or missing, the information displayed might be inaccurate.

**Logon**

Use the **logon** command to detect and logon to Windows installations. Use the following syntax:

```
logon
```

You must correctly enter the local administrator password within three attempts or the computer will restart.

**Map**

Use the **map** command to list all drive letters, file system types, volume sizes, and mappings to physical devices that are currently active. Use the following syntax:

```
map [arc]
```

| Parameter | Description |
|---|---|
| `Arc` | Use the **arc** parameter to force the use of the Advanced RISC Computing (ARC) specification format to describe paths instead of using device paths. You can use this information to create or repair the Boot.ini file. |

If you do not use the **arc** parameter, the **map** command displays device names. The **map** command also displays the type of file system and the size of each disk in megabytes (MB).

The **map** command might not work correctly with systems that are using dynamic disk features.

**Md or Mkdir**

Use the **md** or **mkdir** command to create a new directory or subdirectory. Use the following syntax:

```
md [drive:]pathmkdir [drive:]path
```

| Parameter | Description |
|---|---|
| drive: | Specifies the volume on which to create a folder. |
| path | Specifies the name of the folder to create. |

The **md** or **mkdir** command has the following limitations:

- You cannot use wildcard characters with this command.
- The command only operates within the system directories of the current Windows installation, removable media, the root directory of any hard disk partition, or the local installation sources.
- This command might not display all the volumes on a disk or the correct volume sizes on dynamic disks.

**More or Type**

Use the **more** or **type** command to display the contents of a text file. Use the following syntax:

```
more [path\]filenametype [path\]filename
```

| Parameter | Description |
|---|---|
| filename | Specifies the file name to view. |
| path | Specifies the folder where the file is located. |

If a text file is too large to fit on one screen, use the following page viewing options:

- ENTER Scroll down one line at a time.
- SPACEBAR Scroll down one page at a time.
- ESC to quit viewing the text file.

**Rd or Rmdir**

Use the **rd** or **rmdir** command to delete a directory or subdirectory. Use the following syntax:

```
rd [drive:]pathrmdir [drive:]path
```

| Parameter | Description |
|---|---|
| drive: | Specifies the volume on which to delete a folder. |
| Path | Specifies the name of the folder to delete. |

The **rd** or **rmdir** command has the following limitations:

- You cannot use wildcard characters with this command.
- The path must be empty or the command will fail.
- The command operates within the system directories of the current Windows installation, removable media, the root directory of any hard disk partition, or the local installation sources.

**Ren or Rename**

Use the **ren** or **rename** command to rename a file or directory. Use the following syntax:

```
ren [drive:][path]name1 name2rename [drive:][path]name1 name2
```

| Parameter | Description |
|---|---|
| drive: | Specifies the volume drive letter on which the file to be renamed resides. |
| path | Specifies the path to the file or folder to be renamed. |
| name1 | Specific the file or folder to be renamed. |
| name2 | Specifies the new name for the file or folder. |

You cannot use wildcard characters with this command.

**Set**

Use the **set** command to set Recovery Console environment variables. Recovery Console disables the **set** command by default. To enable the **set** command, you must enable the **Allow floppy copy and access to all drives and all folders** Group Policy setting. For more information about the Group Policy settings for Recovery console, see "Recovery Console Group Policy Settings" later in this topic.

Use the following syntax:

```
set [variable = [TRUE|FALSE]]
```

| Environment Variable | Description |
|---|---|
| AllowWildCards | Set to TRUE to enable wildcard character (* and ?) support for some commands, such as DEL, that do not otherwise support them. |
| AllowAllPaths | Set to TRUE to allow access to all files and folders on the computer. |
| AllowRemovableMedia | Set to TRUE to allow files to be copied to removable media, such as floppy disks. |
| NoCopyPrompt | Set to TRUE to suppress the confirmation prompt that appears when overwriting a file. |

To display the list of current environment variables, use the **set** command without specifying a parameter. All Recovery Console environmental variables are set to FALSE by default.

**Systemroot**

Use the **systemroot** command to set the current directory to the *systemroot* directory of the Windows Server 2003 installation with which you are currently working. Use the following syntax:

```
systemroot
```

**Recovery Console Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with Recovery Console.

**Group Policy Settings Associated with Recovery Console**

| Group Policy Setting | Description |
|---|---|
| **Allow floppy copy and access to all drives and all folders** | Must be enabled in order to use the **set** command to customize Recovery Console |

| | |
|---|---|
| **Recovery Console: Allow automatic administrative logon** | Allows you to to bypass the logon process when Recovery Console starts. Activating this policy eliminates a security barrier intended to protect your computer against unauthorized users. Therefore, It is important that you enable this policy only on systems that have secure consoles, such as those in locked rooms. |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Backup Technical Reference
- Emergency Management Services Technical Reference

# Emergency Management Services Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Emergency Management Services Technical Reference**
Emergency Management Services is a feature of the Microsoft Windows Server 2003 operating systems that supports remote management and remote system recovery for servers that are not accessible through a standard network connection, such as a local area network (LAN) or the Internet. By using Emergency Management Services in conjunction with other software and hardware components, you can remotely perform tasks such as recovering unresponsive servers or powering servers off and on, even if you cannot access the servers through a conventional connection. Emergency Management Services can also help you to save space, energy, and the cost of hardware because it supports servers that are configured without a locally attached keyboard, monitor, or mouse.

Emergency Management Services helps you to return servers to a state where you can manage them using standard tools — Windows technologies that rely on a standard network connection. These tools include Microsoft Management Console (MMC), Systems Management Server, Telnet, Terminal Services Remote Desktop for Administration, Windows Management Instrumentation (WMI), and the Remote Command Service (Rcmd.exe and Rcmdsvc.exe).

**In this subject**


- What Is Emergency Management Services?
- How Emergency Management Services Works
- Emergency Management Services Tools and Settings

# What Is Emergency Management Services?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Emergency Management Services?**

**In this section**


- Common Emergency Management Services Scenarios
- Technologies Related to Emergency Management Services
- Emergency Management Services Dependencies
- Related Information


Before the introduction of Windows Server 2003, you could remotely manage servers running the Windows operating system only by using a conventional in-band connection — one that uses a computer's standard communication channel, typically a network connection. If a server stopped responding, you had to manage it locally. With Windows Server 2003, you can use Emergency Management Services to manage servers remotely even when they are not responding over the network. Emergency Management Services is included with Windows Server 2003, Standard Edition, Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, and Windows Server 2003, Web Edition.

Emergency Management Services supports out-of-band connections — those using an alternate communication channel, most commonly a serial port. Unlike in-band connections, out-of-band connections do not rely on operating system network drivers to function. If you use Emergency Management Services with the appropriate hardware, you can remotely perform almost all management tasks. This capability can help you to efficiently manage large numbers of servers with high availability requirements, especially servers configured without a locally attached keyboard, monitor, and mouse. Some computer firmware may also allow you to configure a server without a keyboard controller and video adapter. Because these components are optional, Emergency Management Services can present significant savings in space, energy consumption, and hardware costs.

**Common Emergency Management Services Scenarios**

You can use Emergency Management Services to remotely troubleshoot problems that arise during any of the following system states:

**Powering on or off**

The computer is powering up or is in the process of shutting down.

**Firmware initializing the system**

The computer is running the Power on Self Test (POST), performing rudimentary hardware checks, and determining the devices that are available.

**Windows loading**

This is the startup phase in which the Windows loader and kernel load and initialize system components such as drivers and services. For multiple-boot systems, the Windows loader might display a boot menu listing the operating systems that are installed. Based on your selection, the appropriate system files are loaded and initialized.

**Windows running**

Windows Server 2003 starts without problems and transitions into one of the following states:

- Text-mode setup phase
- GUI-mode setup phase
- Normal operation

**Windows experiencing problems**

Windows Server 2003 is experiencing software or hardware problems that might adversely affect system responsiveness or might interrupt network connectivity. Certain processes might have depleted system resources, a situation that prevents you

from accessing the remote server using an in-band connection.

**Windows unresponsive**
If Windows Server 2003 is unresponsive to normal in-band connection methods, you can use Emergency Management Services to attempt recovery. For example, if a Stop error occurs, you can use Emergency Management Services to view Stop message information and restart the system. However, if software or hardware problems are sufficiently severe, Emergency Management Services might become unavailable. In these situations, the ability to remotely recover from problems depends on the hardware features built into the system firmware, motherboard, service processor, and uninterruptible power supply (UPS).

**Technologies Related to Emergency Management Services**
When servers can be accessed over the standard network, you should manage them using an in-band connection. Examples of Windows Server technologies that use an in-band connection include:

- Group Policy
- Microsoft Management Console (MMC)
- Systems Management Server
- Telnet
- Terminal Services Remote Desktop for Administration
- Windows Management Instrumentation Command-line (WMIC)
- Windows Script Host (WSH)

The purpose of Emergency Management Services is to help you return a server to a state where in-band management is possible. Unexpected events, however, can interrupt in-band connectivity. For example, a network adapter problem or system instability can cause a server to become unresponsive to management attempts from client computers that use an in-band connection. Because Emergency Management Services removes dependence on operating system network drivers, client computers can communicate with servers through the out-of-band port in situations such as these, so you can attempt to bring the servers back into service remotely. Emergency Management Services is designed to be a last resort option for remote management. You should first try in-band tools.

**Emergency Management Services Dependencies**
Emergency Management Services has the following dependencies:

- **Windows loader or Windows kernel (Ntoskrnl.exe).** Emergency Management Services requires that the loader or the kernel of the Windows Server 2003 operating system be at least partially functioning.

- **Serial ports.** To perform out-of-band management, you need to establish a secure connection through a serial port, phone line, or additional network connection. The serial port, also known as a communication (COM) port, is the most common out-of-band interface. It is the default out-of-band device for Emergency Management Services. The serial port used for Emergency Management Services must be a 16550 Universal Asynchronous Receiver/Transmitter (UART). If your system firmware is compatible with Emergency Management Services, we recommend that you configure it to use the same serial port settings.

  It is important to note that Emergency Management Services and the Windows debugger cannot share the same COM port. The system can have only a single out-of-band management port. Windows Server 2003 does not support a configuration in which one out-of-band port sends outgoing communication and a different out-of-band port (on a second device) receives incoming data.

- **Client terminal software.** You can access all Emergency Management Services output by using terminal emulation software such as Telnet and HyperTerminal. The software you choose should support serial port and terminal definition settings that are compatible with Emergency Management Services, as well as with your system firmware. If possible, use software that supports the VT-UTF8 protocol because it provides for multilingual versions of Windows. If English is the only language you need to support, the VT100+ terminal definition is sufficient. At a minimum, you can use the VT100 terminal definition.

**Related Information**

The following resources contain additional information that is relevant to this section.

- For more information about Emergency Management Services and Group Policy, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.
- For more information about Microsoft Management Console (MMC), click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.
- For more information about Systems Management Server, see MSDN and type the appropriate key words in the "Search MSDN For" text box.
- For more information about Telnet, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.
- For more information about Windows Management Instrumentation Command-line (WMIC), click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.
- For more information about Windows Script Host (WSH), click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.

# How Emergency Management Services Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Emergency Management Services Works**

**In this section**


- Terms and Definitions
- Emergency Management Services Architecture and Interfaces
- Emergency Management Services Physical Structures
- Emergency Management Services Processes and Interactions
- Ports Used by Emergency Management Services
- Related Information


Emergency Management Services is a technology in Windows Server 2003 that supports remote management and system recovery for servers that are not accessible through an in-band connection. An in-band connection is a connection between two computers that relies on a standard network, such as a local area network (LAN) or the Internet, and on standard remote administration tools, such as Remote Desktop or Telnet. You can use this type of connection to remotely manage computers only if both the local and remote computers are in a functional state and accessible on the network.

Emergency Management Services redirects text output to the out-of-band connection. An out-of-band connection is a nonstandard connection between two computers, such as a serial port connection, and is useful when a remote server cannot access the network or is not fully functional. Emergency Management Services provides a command-line environment for managing a server through the out-of-band port. The capability of redirecting text output is also known as console redirection.

Emergency Management Services consists of the following three components:

- Console Redirection
- Special Administration Console (SAC) environment
- !Special Administration Console (!SAC) environment

Emergency Management Services also supports headless operation for computers that have the appropriate out-of-band components. (Headless operation refers to computers that function without a local display or input devices, such as a keyboard or mouse.)

The following sections provide an in-depth view of how Emergency Management Services works in an optimal environment. An optimal environment for Emergency Management Services is defined as follows:

- Emergency Management Services is one part of a complete management solution that includes in-band as well as out-of-band technologies. To support this solution, you have selected and configured the hardware that will allow you to access and manage the server at all times.
- Your system firmware supports console redirection and the Serial Port Console Redirection (SPCR) table. If it does not, you must enable and configure Emergency Management Services manually.
- You have configured your hardware and Emergency Management Services to use the highest baud rate possible. This gives you the best performance, and also prevents a slow text-mode setup process.
- If you are using Emergency Management Services with a serial port, you are using a null modem cable with carrier

detect.

- If you choose to use a service processor, purchase one that supports Emergency Management Services. A service processor can provide additional management functionality that complements that of Emergency Management Services.

- You have configured your server to restart automatically if a Stop error occurs. This allows the server to restart without requiring you to use Emergency Management Services or the power switch to reboot it manually.

- Your hardware and firmware support the same terminal type that Emergency Management Services uses, preferably VT-UTF8 or VT100+. To prevent incompatibility, you must at least select hardware and firmware that support VT100. If the terminal types are not the same, it can be extremely difficult to send the proper command escape sequences.

- You have limited physical access to your server. You protect sensitive areas such as racks, computer labs, server farms, and data centers by locking them, choosing hardware with built-in security features, and setting up a separate network dedicated exclusively to traffic management.

**Terms and Definitions**

Before you review the Emergency Management Services architecture, it is helpful to understand the terms associated with this technology. The following is a brief introduction to the three components of Emergency Management Services and a glossary of terms associated with each component.

**Console Redirection**

Emergency Management Services provides console redirection, or the ability to send and receive character-mode output or input by using the out-of-band port. This functionality enables you to manage a remote server that would otherwise be inaccessible. For example, a server that supports console redirection is able to receive keyboard input from a managing client computer and to send character-mode output back to the client computer by using the out-of-band port. If you enable console redirection, you can still use local input and display devices. Therefore, it is possible for a computer running Windows Server 2003 to process in-band, out-of-band, and local console input and output (I/O) simultaneously.

Console redirection uses a character-mode display instead of a graphical user interface (GUI) for the following reasons:

- **Hardware compatibility**. A character-mode display provides compatibility with different types of communication media, from relatively slow serial port connections to higher speed technologies. The serial port is the most commonly used out-of-band device.

- **Software compatibility**. A character-mode display simplifies interoperability issues with different types of terminals and terminal emulation software, including those running on other operating systems such as UNIX.

You can use Emergency Management Services in conjunction with hardware that increases your ability to access a remote server through out-of-band means. Windows Server 2003, hardware integrated into the motherboard, or add-in devices can provide the following types of console redirection:

- **Firmware console redirection**. System firmware that supports console redirection provides for remote management before Windows Server 2003 starts. Types of firmware include basic input/output system (BIOS) for x86-based computers and Extensible Firmware Interface (EFI) for Itanium-based systems.

- **Emergency Management Services console redirection**. Windows Server 2003 components, such as the Windows loader (Ntldr), support console redirection.

- **Service processor console redirection**. If your server includes a specialized hardware component called a service processor, the service processor can also provide console redirection. For more information about service processors, see “Service Processors” later in this section.

**Firmware Console Redirection**

Your computer’s firmware might support console redirection and can work in conjunction with Emergency Management Services to give you greater control over a remote server. Typically, firmware for Itanium-based computers supports console redirection. For x86-based computers, contact your computer manufacturer for information about firmware console redirection functionality. If your x86-based computer does not support firmware console redirection, you might be able to

obtain a firmware upgrade.

You can use firmware console redirection to remotely manage a server from restart until the initial load of Windows Server 2003. Without firmware console redirection, you cannot manage your computer remotely until Windows Server 2003 starts.

**Note**

- A service processor can provide console redirection before Windows Server 2003 starts, even without firmware console redirection.

With firmware console redirection, you can perform the following remote management tasks before Windows Server 2003 starts.

- **Remotely view startup progress.** You can remotely view computer status before Windows Server 2003 starts, such as power-on self-test (POST) information, disk-related error messages, and other text displayed by the computer's firmware. A computer that has firmware console redirection typically allows the POST to complete successfully without an attached keyboard, mouse, or display device.

- **Remotely view and edit firmware settings.** You can remotely access the configuration program provided by the computer firmware to change hardware settings. For example, you can remotely disable integrated peripherals or change the boot device order. Without firmware console redirection, you must make these types of changes locally.

- **Remotely view and respond to Pre-Boot eXecution Environment prompts.** If a server's firmware supports the Pre-Boot eXecution Environment (PXE) standard, you can use firmware console redirection to remotely view and respond to the F12 network boot prompt.

- **Remotely view and respond to the boot from CD prompt.** With firmware console redirection, you can remotely respond to the **Press Any Key to Boot from CD** prompt when starting your server by using the CD for Windows Server 2003.

#### Emergency Management Services Console Redirection

If Emergency Management Services is enabled when Windows Server 2003 starts, the operating system assumes the responsibility for console redirection from the firmware. The ability to redirect console I/O to the out-of-band port is built into the Windows Server 2003 components listed in the following table.

**Components that Support Emergency Management Services Console Redirection**

| Windows Server 2003 Component or Feature | Description |
|---|---|
| Windows loader for x86-based computers (Ntldr) | The Windows Server 2003 operating system loader. When Ntldr is running, you can remotely view and select the Recovery Console or, on x86-based multiple-boot computers, which operating system to start. |
| Windows kernel (Ntoskrnl.exe) | The core (also called the kernel) of the Windows Server 2003 operating system. Code that runs as part of the kernel does so in privileged processor mode and has direct access to system data and hardware. The Windows kernel supports console redirection, which enables you to remotely view system information during normal operation, or remotely view Stop message text when a system problem occurs. |
| Recovery Console | A command-line environment that enables you to perform advanced troubleshooting and maintenance tasks, such as disabling a driver that you suspect is causing a startup problem. |
| Command Prompt (Cmd.exe) | A character mode command shell user interface that provides an environment for running commands and applications. |

| | |
|---|---|
| Text-mode Setup (including the CD-ROM Setup loader) | The initial portion of the Windows Server 2003 installation process when Setup displays character mode prompts and status information text. During text-mode setup, files are copied from the distribution folder to the local hard disk. |
| Startrom.com at 9600 baud for x86-based computers | Starts the x86-based Remote Installation Services (RIS) process. This file is downloaded and run by the RIS client to initiate the operating system installation procedure. Only special versions of Startrom.com that use 9600 baud support Emergency Management Services console redirection. For more information about Startrom.com, see "Remote Installation Services" later in this chapter. |

Emergency Management Services does not provide console redirection for the components listed in the following table.

**Components that Do Not Support Emergency Management Services Console Redirection**

| Windows Server 2003 Component | Description |
|---|---|
| The **Press Any Key to Boot from CD** prompt | This prompt appears when you start your system by using the Windows Server 2003 operating system CD. |
| GUI-mode Setup | Emergency Management Services processes only character-mode input and output. Emergency Management Services cannot display GUI-mode graphics, such as windows and dialog boxes, and it cannot redirect mouse input. You can, however, perform Special Administration Console (SAC) commands and monitor setup logs during GUI-mode setup. |
| Startrom.com at baud rates greater than 9600 | Startrom.com supports Windows Emergency Management Services console redirection only at 9600 baud. |
| Operating system boot menu for Itanium-based computers | For Itanium computers, the EFI boot manager (IA64ldr), not the Windows loader, displays the list of operating systems installed on the system. To remotely view and select boot manager menu items on Itanium-based systems, you must enable firmware console redirection. |

**Special Administration Console**

Special Administration Console (SAC), a kernel-mode component, is the primary Emergency Management Services command-line environment. You can use SAC to remotely manage a server through its out-of-band port when Windows runs in GUI mode.

SAC is available early in the startup process. You can use it to manage a server during normal system operation, safe mode, and the GUI phase of Windows Server 2003 setup. When Emergency Management Services is enabled, SAC is always active as long as the kernel is running. SAC is a separate environment from the Windows Server 2003 command prompt (Cmd.exe). You can remotely monitor the status of a server or troubleshoot problems by using terminal software to access SAC.

You can use SAC to perform the following remote tasks when the operating system is functioning normally:

- Gather server information, such as computer name and IP address.
- Obtain a list of the services and applications running on the computer.
- Restart or shut down a server as part of planned maintenance.
- Change the system time and date.

You can use SAC to perform the following troubleshooting tasks:

- Raise or lower the priority of a process, or end a process that is consuming excessive processor time or other system resources. By establishing an out-of-band connection to SAC, you can continue to manage a server even when low

system resources have caused it to become unresponsive to in-band management tools.

- Restart or shut down a server that has stopped responding to in-band commands and tools.
- View or change a server's Transmission Control Protocol/Internet Protocol (TCP/IP) networking information to resolve issues, such as problems caused by a duplicate IP address.
- Maintain communication with a server during network outages. For example, you might want to verify that a connectivity issue is caused by a problem that affects the primary network rather than by a problem with the server hardware or a driver.
- Generate a list of drivers for diagnostic purposes.
- View setup logs to determine the cause of a problem that prevents Windows Server 2003 installation from completing.

**Note**

- SAC usage is not safeguarded by password and logon requirements. You must secure physical access to computers running Emergency Management Services by placing the appropriate restrictions on all connecting terminals.

**!Special Administration Console**

!Special Administration Console (!SAC) is an abbreviated version of SAC that accepts input and sends output through the out-of-band port. !SAC is a separate entity from SAC and the Windows Server 2003 command prompt. !SAC can become available while a computer is running Windows Server 2003 in normal mode, safe mode, or the GUI-mode phase of Windows Server 2003 setup. Unlike SAC, you cannot invoke !SAC directly. After a specific failure point is reached, Emergency Management Services components make the transition from SAC to !SAC.

!SAC provides a subset of SAC commands and enables you to perform a limited set of tasks, including the following:

- Remotely view Stop message text.
- Restart the computer.
- View an abbreviated log of loaded drivers and some kernel events.
- Obtain computer identification information.

**Glossary of Emergency Management Services Terms**

The following terms are used to describe the components and processes of Emergency Management Services.

**Bootcfg**

Configures, queries, or changes Boot.ini file settings.

**Console redirection**

Sending and receiving character-mode output or input by using the out-of-band port. A server that supports console redirection is able to receive keyboard input from a managing client computer and send character-mode output back to the client computer by using the out-of-band port. Enabling console redirection does not prevent you from using local input and display devices. Therefore, it is possible for a computer running Windows Server 2003 to process in-band, out-of-band, and local console input and output (I/O) simultaneously.

**GUI mode**

The portion of Setup that uses a graphical user interface (GUI).

**in-band remote management tools**

Microsoft Management Console (MMC), Systems Management Server, Telnet, Terminal Services Remote Desktop for Administration, Windows Management Instrumentation (WMI)

**Service processor**

A micro controller that operates independently of the computer processor(s) and operating system and uses its own custom firmware. Service processors are typically integrated into the motherboard or a PCI adapter.

**Setup logs**
A program used by Remote Installation Services (RIS) that is the first file downloaded to the client using Trivial File Transfer Protocol (TFTP). Startrom.com is a small program that displays the **Press F12 for Network Service Boot** prompt. If F12 is pressed within three seconds, the Client Installation Wizard (OSChooser) is downloaded to begin the remote installation process.

**Stop error**
A serious error that affects the operating system and that could place data at risk. The operating system generates an obvious message, a screen with the Stop error, rather than continuing on and possibly corrupting data. Also called a fatal system error.

**terminal concentrator**
A network-capable management device that consolidates serial access to multiple servers. Terminal concentrators have at least one Ethernet port that you can connect to a primary or secondary network and multiple serial ports for connecting to the managed servers. The remote management client connects to the terminal concentrator through the Ethernet port and the managed servers connect to the terminal concentrator with null modem cables.

**VT-UTF8**
A character set for protocols evolving beyond the use of ASCII. The UTF-8 protocol provides for support of extended ASCII characters and translation of UCS-2, an international 16-bit Unicode character set. UTF-8 enables a far greater range of names than can be achieved using ASCII or extended ASCII encoding for character data.

**unattend.txt and Winnt.sif**
The Unicode-aware terminal type, based off of a VT-100 terminal type, that is used for transmitting data over an EMS connection.

**Universal Asynchronous Receiver/Transmitter (UART)**
An integrated circuit (silicon chip) that is commonly used in microcomputers to provide asynchronous communication. The UART provides parallel-to-serial conversion of data to be transmitted and serial-to-parallel conversion of data received.

**universal serial bus (USB)**
An external bus that supports Plug and Play installation. Using USB, you can connect and disconnect devices without shutting down or restarting your computer. You can use a single USB port to connect up to 127 peripheral devices, including speakers, telephones, CD-ROM drives, joysticks, tape drives, keyboards, scanners, and cameras. A USB port is usually located on the back of your computer near the serial port or parallel port.

**Emergency Management Services Architecture and Interfaces**
The following table describes the three components of Emergency Management Services.

### Emergency Management Services Components

| Component | Description |
|---|---|
| Special Administration Console (SAC) | The primary Emergency Management Services command-line environment hosted by Windows Server 2003 operating systems. It is separate from the Windows Server 2003 command-line environment and provides different functionality. |
| !Special Administration Console (!SAC) | An auxiliary Emergency Management Services command-line environment that is hosted by Windows Server 2003 operating systems. Like SAC it accepts input and sends output through the out-of-band port. !SAC is a separate entity from both SAC and the Windows Server 2003 command-line environment. |
| Special Administration Console Helper service (Sacsvr) | A user-mode service you can use to create Windows command-prompt channels when you type cmd at the SAC prompt. With the command-prompt channels, you can use in-band tools on a remotely administered server when an in-band connection is not available. |

**Emergency Management Services Physical Structures**
This section describes the types of hardware and software you might choose to implement Emergency Management Services on a remote server.

**Out-of-Band Components**
Depending on the out-of-band management tasks you must perform, you might need to use Emergency Management Services in conjunction with out-of-band hardware components. You can manage nearly every operating system state by using only Emergency Management Services and a serial port. However, if you need to manage many servers or need to manage servers in situations when the operating system is not functioning, you need supporting firmware and hardware components. These components can support situations ranging from power up to operating system initialization to unresponsive servers.

You might need any combination of the following components.

#### Out-of-Band Ports

Serial ports are the most common out-of-band hardware interface, but service processors sometimes provide alternative methods. For example, some service processors provide an RJ-45 Ethernet port for out-of-band connections. In addition, advances in technology are opening up new possibilities for out-of-band ports. For more information about types of out-of-band ports, see "Service Processors" later in this section.

To be compatible with Emergency Management Services, out-of-band ports must meet the following requirements.

- **Exclusive access.** The Windows Server 2003 operating system must have exclusive access to the out-of-band port. Instead of requesting input and output operations from the Windows I/O Manager, Emergency Management Services writes information directly to Universal Asynchronous/Receive Transmit (UART) hardware.

- **Single out-of-band port.** Emergency Management Services supports only one out-of-band management port. Windows Server 2003 does not support a configuration in which one out-of-band port is used to send data and another one is used to receive incoming data.

- **Continuous availability.** Emergency Management Services requires that the out-of-band port be available at all times, including when the Windows loader is running, the operating system is fully loaded, and when a Stop error occurs. Furthermore, the device must not be powered off while the system is running.

#### Serial Ports

The serial port, which is also referred to as a COM port, is the default out-of-band device for Emergency Management Services. This is because serial ports have the following characteristics.

- **Universality.** At least one serial port exists on most computers.

- **Simplicity and cost effectiveness.** Traditional serial ports use common 16550 or 16450 UART integrated circuits. UART interfaces are standard equipment on most motherboards. Hardware manufacturers can also incorporate inexpensive 16550 or 16450 UART interfaces into add-on adapters that provide out-of-band ports.

- **Hardware flexibility.** UARTs can be logically located at I/O addresses other than those traditionally reserved for COM1 or COM2 (provided that Emergency Management Services is properly configured to use the corresponding hardware settings). Serial ports enable external connectivity to other serial devices such as terminal concentrators, which act as network interfaces to out-of-band hardware.

- **Interpretability.** Several well-defined serial port communication standards exist, enabling interoperability with terminal hardware and software running on operating systems other than Windows, such as UNIX. The computer's firmware initializes serial ports early in the startup process and serial ports are detected and usable before Windows Server 2003 starts.

### Note

- When you use Emergency Management Services with a serial port, use null mode cables that support the Carrier Detect (CD) signal. Cables with the CD pin directly connected to the Request to Send pin also function correctly.

#### Requirements for Serial Port Compatibility

In addition to the general requirements described earlier in this section, serial ports must meet a number of other requirements to be compatible with Emergency Management Services.

- **UART device standards.** The UART control register must function as a standard 16450 or 16550 UART device. Before using a UART device for Emergency Management Services, Windows Server 2003 tests for this compliance.

- **UART device availability.** The UART device must be available before the Windows loader starts. As a rule, out-of-band devices must be configured by the firmware as part of the POST process.

- **UART device I/O address.** The UART device I/O address must not change. Emergency Management Services does not support Plug and Play. If you change UART hardware resource settings without making corresponding changes to

Emergency Management Services settings, data might be written to an invalid system address.

**Technology Trends for Out-of-Band Ports**

A recent industry trend is the movement towards legacy-free computers, that is, those that do not use super I/O chips or serial ports located on the Industry Standard Architecture (ISA) bus. (A super I/O chip is a single chip that contains integrated drive electronics (IDE) controllers, a floppy controller, serial ports, and a parallel port). The PC 99 and PC 2001 System Design Guides strongly recommend alternatives to super I/O chips for such legacy-free designs. In these systems, a UART placed on a Peripheral Component Interconnect (PCI) adapter is the preferred mechanism for implementing out-of-band ports.

Additionally, hardware manufacturers might develop proprietary out-of-band solutions that combine custom hardware and software solutions. Other possibilities for out-of-band external interfaces include connectivity based on industry standard Universal Serial Bus (USB), IEEE1394 Firewire ports, or other proprietary connections. Contact your hardware vendor for more information about technologies compatible with Emergency Management Services.

**Service Processors**

Emergency Management Services requires that the Windows loader or kernel be at least partially functioning. If you need to communicate with a server when this is not the case — such as when the computer is completely unresponsive due to a severe error — you might need a remote management channel that is independent of the operating system. Service processors can fill this need.

Service processors are micro controllers that operate independently of the computer processor(s) and operating system and use their own custom firmware. Service processors are typically integrated into the motherboard or a PCI adapter.

You can communicate directly with a service processor by using an out-of-band connection, regardless of the state of the computer. Depending on the manufacturer, a service processor might use any of a variety of out-of-band communication channels. The most common out-of-band channels are serial ports and RJ-45 Ethernet ports. However, if the out-of-band communication channel is a serial port, you can use only one tool at a time — either Emergency Management Services or the service processor.

**Serial port connections to service processors**

When a service processor uses a serial port for out-of-band connections, it passively monitors the traffic between the UART interface and the serial port. This behavior is sometimes referred to as pass-through mode. When the service processor detects information in the data stream that applies to it, it takes control of the out-of-band port. The service processor then presents a remote management user interface. The service processor must release control of the serial port when either of the following occurs:

- The service processor detects control sequences in the serial data stream instructing it to immediately release control of the out-of-band port.
- The firmware POST process completes and the Windows Server 2003 loader starts.

The following figure illustrates how Emergency Management Services, the UART, a service processor, and a serial port fit together.

**Serial Port Null Modem and Pass-Through Service Processor Configuration**




**Ethernet port connections to service processors**

Some service processors include an integrated RJ-45 Ethernet port that is dedicated for out-of-band use. These network-

enabled service processors typically use an Internet Protocol (IP) address different from the host computer and can be connected to a secondary management network for added security and redundancy.

**Service processor features**
Manufacturers vary in the management tools, features, and client interfaces they provide. Management tools and client interfaces range from simple Telnet consoles to complex Web browser environments. Some service processors also include their own power supplies to enhance reliability. Service processor features that you might want to use with Emergency Management Services include the following:

- Remote reset
- Console redirection
- Remote power on
- Remote power off

**Firmware Console Redirection**
Firmware console redirection works with Emergency Management Services to provide out-of-band support for any operating state. If you do not have firmware console redirection — or a service processor that provides console redirection — you cannot manage your server's remotely until Windows Server 2003 begins loading.

For a seamless transition from firmware console redirection to Emergency Management Services console redirection, the firmware or service processor must completely release control of the out-of-band port when the Windows loader starts. For best results, terminal software running on the client computer needs to use serial port and terminal definition settings that are compatible with firmware console redirection and Emergency Management Services console redirection.

**Terminal Concentrators**
A terminal concentrator is a network-capable management device that consolidates serial access to multiple servers. Terminal concentrators have at least one Ethernet port that you can connect to a primary or secondary network and multiple serial ports for connecting to the managed servers. The remote management client connects to the terminal concentrator through the Ethernet port and the managed servers connect to the terminal concentrator with null modem cables.

By using a terminal concentrator, you can establish an in-band connection from a remote management computer, typically by using the Telnet protocol, to monitor many servers simultaneously. By sending specific commands to the terminal concentrator, you can switch from one serial port to another, and therefore from one server to another.

Although the in-band connection from the management client eliminates physical security concerns, it does introduce logical security concerns. Some terminal concentrators provide authentication and encryption security for the in-band connection. For example, some terminal concentrators support Secure Shell (SSH), which is a secure alternative to Telnet. Because security features for terminal concentrators are not standardized, however, you might need to provide your own logical security. For example, you might set up a secondary private management network or use a secure router. It is important to ensure that only authorized users can access a terminal concentrator, because it provides access to the credentials for multiple servers.

**Intelligent Uninterruptible Power Supplies and Intelligent Power Switches**
An Uninterruptible Power Supply (UPS) is an optional component that provides a backup power source to a server in the event of a power failure. Some UPS units, known as intelligent UPSs, enable you to remotely cycle power, providing rudimentary remote management. For Emergency Management Services systems with firmware redirection enabled, an intelligent UPS can enhance remote management capabilities by responding to remote commands sent to it.

When used with Emergency Management Services, an intelligent UPS must passively monitor the serial port traffic and support the key sequences that apply to it and must not affect communication between the managing client and the server running Windows Server 2003. If you also use a service processor, the intelligent UPS must be able to differentiate the key sequences directed to itself, the service processor, and Emergency Management Services. The UPS must also be able to inform the system of a planned power shutdown.

An intelligent power switch is an alternative to an intelligent UPS. An intelligent power switch can provide features similar to an intelligent UPS, except that it does not provide a backup source of power.

### Note

- To take full advantage of intelligent UPS or intelligent power switch capabilities, you need to configure servers to automatically start up when power is applied. Your servers firmware might provide an option to enable this feature.

**Emergency Management Services Processes and Interactions**
The following section describes certain tasks that you can perform using Emergency Management Services.

**Remote Tasks Supported During Normal Operation**
With Emergency Management Services console redirection, you can perform the following remote management tasks when the operating system is functioning normally.

**Verify that Windows has started**
When the Windows loader starts, it displays the following message by using the out-of-band port:

```
Windows is now starting
```

The preceding message indicates that the Windows Server 2003 startup process has begun.

When the Windows kernel is initialized, it displays the following prompt by using the out-of-band port, indicating that Emergency Management Services is enabled:

```
SAC>
```

**Choose an item on the operating system boot menu (forx86-based computers)**
For multiple-boot systems that have Windows Server 2003 installed along with Recovery Console or another Microsoft operating system, the Windows loader (Ntldr) automatically displays an operating boot menu, both to the local display and to the out-of-band port. By using this operating system boot menu, you can start the Windows Server 2003 operating system or Recovery Console. You also have the option of viewing the Windows Advanced Options Menu by pressing the F8 key. The following is a typical operating system boot menu.

```
Please select the operating system to start:

Microsoft Windows Server 2003
Recovery Console

Use the up and down arrow keys to move the highlight to your choice.
Press Enter to choose.
Seconds until highlighted choice will be started automatically: 29

For troubleshooting and advanced startup options for Windows, press F8.
```

**Monitor the progress of Windows Server 2003 text-mode setup**
The text-mode setup phase displays character-based information, which supports Windows Emergency Management Services console redirection. Therefore, you can remotely view and respond to text-mode setup prompts.

**View and respond to PXE environment prompts**
If the server's firmware supports the Pre-Boot eXecution Environment (PXE) standard but does not support firmware console redirection, an alternate version of the Startrom.com file provides Emergency Management Services console redirection.

**Remote Tasks Supported for Problem Resolution**
With Emergency Management Services console redirection, you can perform the following remote management tasks to troubleshoot problems.

**Use Last Known Good Configuration to troubleshoot problems**
Last Known Good Configuration is a troubleshooting option you can use when a driver or application installation problem prevents a computer from starting in normal mode. By using Last Known Good Configuration, you can recover from a problem by reversing the driver and registry changes made since you last started Windows Server 2003. You can invoke Last Known Good Configuration from the Windows Advanced Options Menu. You can display the Windows Advanced Options Menu by pressing F8 during system startup or at the operating system boot menu, if it appears.

**Use Recovery Console to troubleshoot problems**
Recovery Console is a command-line environment that you can use to remotely perform advanced troubleshooting operations, such as manually replacing system files. You can invoke Recovery Console by installing it on your computer and then selecting it from the operating system boot menu or by starting it from the Windows Server 2003 installation CD.

**Ports Used by Emergency Management Services**
Emergency Management Services redirects text through an out-of-band connection, typically a serial port connection. Serial ports are also known as COM ports. The following table lists the COM ports used by Emergency Management Services and their settings.

**Typical Serial Port Resource Settings**

| Friendly Name | Typical Hexadecimal I/O Address | Interrupt Request (IRQ) Value |
|---|---|---|
| **COM1** | Emergency Management Services supports the industry standard 3F8 (this is typically the firmware default or AUTO value for COM1). Values of 2F8, 3F8, On, and Off might also be available. | 4 (this is typically the default or AUTO value) |
| **COM2** | Emergency Management Services supports the industry standard 2F8 (this is typically the firmware default or AUTO value for COM2). Values of 2F8, 3F8, On, and Off might also be available. | 3 (this is typically the default or AUTO value) |

**Note**

- It is important to realize that EMS does not use both COM1 and COM2 at the same time: it uses one or the other. In addition, it can use a non-legacy serial port address provided the system designer builds an EMS-aware design.

**Related Information**

The following resources contain additional information that is relevant to this section.

- For more information about Emergency Management Services, click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

# Emergency Management Services Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Emergency Management Services Tools and Settings**

**In this section**


- Emergency Management Services Tools
- Emergency Management Services Registry Entries
- Emergency Management Services Group Policy Settings
- Emergency Management Services WMI Classes
- Ports Used by Emergency Management Services
- Related Information


Emergency Management Services is a new technology in Windows Server 2003 that administrators can use to remotely manage and recover servers that are not accessible through an in-band connection.

An in-band connection is a connection between two computers that relies on standard network connectivity, such as a local area network (LAN) or the Internet, and on standard remote administration tools, such as Remote Desktop or Telnet. You can use this type of connection to remotely manage computers only if both the local and remote computers are in a functional state and accessible on the network.

Emergency Management Services uses a nonstandard or out-of-band connection, such as a serial port connection, to connect two computers. This arrangement is useful when you are managing a particular computer that cannot access the network or is not fully operational and, as a result, must rely on the other computer for critical management functionality. In order to provide this functionality, Emergency Management Services makes available three tools: console redirection, Special Administration Console (SAC), and !Special Administration Console (!SAC).

Emergency Management Services also supports remotely administered serversthat do not have a locally attached keyboard, video, or mouse (also referred to as *headless* servers). Emergency Management Services also supports server configurations that have no keyboard controller or video controller, if the system firmware supports this. Using servers that do not have a locally attached keyboard, video or mouse in, for example, a datacenter or similar facility has several benefits. It can result in a significant savings in space, energy consumption, and hardware costs.

Emergency Management Services is only enabled by default if the Advanced Configuration and Power Interface (ACPI) Serial Port Console Redirection (SPCR) table in the system BIOS is set to indicate the out-of-band connecting settings. Otherwise, Emergency Management Services is not enabled by default, but can be enabled at any of the following times:

- During a new Windows Server 2003 installation (including RIS-based and image-based setups)
- During a Windows Server 2003 upgrade (for x86-based systems only)
- After you complete a Windows Server 2003 installation

Emergency Management Services functionality is built into Windows Server 2003. It is not necessary to copy additional files from the Windows Server 2003 installation CD.

**Note**

- The SPCR table, found on some computers that are equipped with ACPI-compliant firmware, provides information about how the system firmware and the service processor, if available, use the out-of-band management port. Emergency Management Services uses the information in the SPCR table to ensure a smooth transition between the output of the system firmware and the output of Emergency Management Services.

The purpose of Emergency Management Services is to help you return a server to a state where you can manage it using an in-band connection — the preferred way to manage a server. Unexpected events, however, can interrupt in-band connectivity. For example, a network adapter problem or system instability can cause a server to become unresponsive to management attempts from client computers that use an in-band connection. Because Emergency Management Services does not rely on operating system network drivers, client computers can communicate with servers through the out-of-band port in such situations and can attempt to bring the servers back into service remotely.

**Emergency Management Services Tools**
The following tools are associated with Emergency Management Services:

- Console redirection
- Special Administration Console (SAC)
- !Special Administration Console (!SAC)

Emergency Management Services console redirection is able to utilize a number of standard Windows features, such as Recovery Console, by redirecting text output from these features through the out-of-band port to a remote computer running terminal emulation software. This can be of vital importance when the remotely managed computer is not accessible, either locally or over the network, and the only means of communicating with the computer is through the out-of-band connection. Emergency Management Services console redirection is only available when the operating system is functioning normally.

SAC and !SAC are tools that are specific to Emergency Management Services. These tools are entirely separate from other Windows tools or features and provide their own functionality. Like Emergency Management Services console redirection, SAC is only available when the operating system is functioning normally. Unlike Emergency Management Services console redirection and SAC, !SAC only becomes available when the system is not functioning normally and SAC itself has failed to load or ceased to operate.

**Console Redirection**
**Category**
Windows Server 2003 operating system tool

**Version compatibility**
Supported on all Windows Server 2003 operating systems

Emergency Management Services console redirection enables you to use a number of standard Windows features, such as Recovery Console, even when the computer that you are managing is not otherwise accessible.

If Emergency Management Services is enabled when Windows Server 2003 starts, the operating system assumes responsibility for console redirection from the firmware. Up to this point, the firmware has provided a certain level of functionality independent of the operating system, and, in a manner similar to Emergency Management Services, has redirected text output through the out-of-band port. (Firmware is defined as software routines and low-level input/output instructions stored in read-only memory [ROM]. Unlike random access memory [RAM], read-only memory stays intact even in the absence of electrical power.)

The following table lists the Windows Server 2003 components and features that support Emergency Management Services console redirection after the operating system starts.

**Components That Work With Emergency Management Services Console Redirection**

| Windows Server 2003 Component or Feature | Description |
|---|---|
| Windows loader forx86-based computers (Ntldr) | The Windows Server 2003 operating system loader. When Ntldr is running, you can remotely view and select the Recovery Console or, on x86-based multiple-boot computers, which operating system to start. |

| | |
|---|---|
| Windows kernel (Ntoskrnl.exe) | The core (also called the kernel) of the Windows Server 2003 operating system. Code that runs as part of the kernel does so in privileged processor mode and therefore has direct access to system data and hardware.<br><br>The Windows kernel supports console redirection, which enables you to remotely view system information during normal operation, or remotely view Stop message text when a system problem occurs. |
| Recovery Console | A command-line environment that you can use to perform advanced troubleshooting and maintenance tasks, such as disabling a driver that you suspect is causing a startup problem. |
| Command Prompt (Cmd.exe) | A character-mode, command shell user interface that provides an environment for running commands and applications. |
| Text-mode Setup (including the CD-ROM Setup loader) | The initial portion of the Windows Server 2003 installation process when Setup displays character-mode prompts and status information text. During text-mode Setup, files are copied from the distribution folder to the local hard disk. |
| Startrom.com at 9600 baud for x86-based computers | Starts the x86-based Remote Installation Services (RIS) process. This file is downloaded and run by the RIS client to initiate the operating system installation procedure. Only special versions of Startrom.com that use 9600 baud support Emergency Management Services console redirection.<br><br>RIS includes software services that an administrator can use to set up new client computers remotely, without having to visit each client. The destination clients must support remote booting. |

Emergency Management Services does not provide console redirection for the components listed in the following table.

**Components That Do Not Work With Support Emergency Management Services Console Redirection**

| **Windows Server 2003 Component** | **Description** |
|---|---|
| The **Press Any Key to Boot from CD** prompt | This prompt appears when you start your system using the Windows Server 2003 operating system CD. |
| GUI-mode Setup | Emergency Management Services processes only character-mode input and output. Emergency Management Services cannot display GUI-mode graphics, such as windows and dialog boxes, or redirect mouse input. Using the Special Administration Console (SAC), you can, however, view Setup “channels” that allow you, in turn, to view Setup log files as they are created so that you can monitor the progress of Setup. |
| Startrom.com at baud rates greater than 9600 | Startrom.com supports Emergency Management Services console redirection only at 9600 baud. |
| Operating system boot menu for Itanium-based computers | For Itanium-based computers, the EFI boot manager (IA64ldr), not the Windows loader, displays the list of operating systems installed on the system. To remotely view and select boot manager menu items on Itanium-based systems, you must enable firmware console redirection. |

**Remote Management Tasks Supported by Console Redirection during Normal Operation**

With Emergency Management Services console redirection, you can perform the following remote management tasks when the operating system is functioning normally.

- **Verify that Windows has started**

  When the Windows loader starts, it displays the following message using the out-of-band port:

```
Windows is now starting
```

The preceding message indicates that the Windows Server 2003 startup process has begun.

When the Windows kernel initializes, it displays the following prompt by using the out-of-band port, indicating that Emergency Management Services is enabled:

```
SAC>
```

- **Choose an item on the operating system boot menu (for x86-based computers)**

  For multiple-boot systems that have Windows Server 2003 installed along with Recovery Console or another Microsoft operating system, Ntldr automatically displays an operating boot menu, both to the local display and to the out-of-band port. By using this operating system boot menu, you can start the Windows Server 2003 operating system or Recovery Console. You also have the option of viewing the Windows Advanced Options Menu by pressing the F8 key. The following is a typical operating system boot menu.

```
Please select the operating system to start:

Microsoft Windows Server 2003
Recovery Console

Use the up and down arrow keys to move the highlight to your choice.
Press Enter to choose.
Seconds until highlighted choice will be started automatically:  29

For troubleshooting and advanced startup options for Windows, press F8.
```

- **Monitor the progress of Windows Server 2003 text-mode Setup**

  The text-mode Setup phase displays character-based information, which supports Emergency Management Services console redirection. You can remotely view and respond to text-mode Setup prompts.

- **Monitor the progress of Windows Server 2003 GUI-mode Setup**

  Emergency Management Services processes only character-mode input and output. Emergency Management Services cannot display GUI-mode graphics, such as windows and dialog boxes, or redirect mouse input. Using Special Administration Console (SAC), you can, however, view Setup "channels" that allow you, in turn, to view Setup log files as they are created so that you can monitor the progress of Setup.

- **View and respond to Preboot Execution Environment (PXE) prompts**

  If the server firmware supports the PXE standard but does not support firmware console redirection, an alternate version of the Startrom.com file provides Emergency Management Services console redirection.

**Remote Management Tasks Supported by Console Redirection for Troubleshooting**

With Emergency Management Services console redirection, you can perform the following remote management tasks to troubleshoot problems.

- **Use Last Known Good Configuration to troubleshoot problems**

  Last Known Good Configuration is a startup option that you can use to troubleshoot when a driver or application installation problem prevents a computer from starting in normal mode. By using Last Known Good Configuration, you can recover from a problem by reversing the driver and registry changes that were made since you last started Windows Server 2003. You can access Last Known Good Configuration from the Windows Advanced Options Menu. The Windows Advanced Options Menu can be displayed by pressing F8 during system startup or at the operating system boot menu, if available.

- **Use Recovery Console to troubleshoot problems**

  Recovery Console is a command-line environment that you can use to remotely perform advanced troubleshooting operations, such as manually replacing system files. You can use Recovery Console by installing it on your computer and then selecting it from the operating system boot menu. You can also start it from the Windows Server 2003 installation CD.

**Sacdrv.sys: Special Administration Console**
**Category**
Windows Server 2003 operating system tool

**Version compatibility**
Supported on all Windows Server 2003 operating systems

Special Administration Console (SAC) is the primary Emergency Management Services command-line environment. You can access this console through an out-of-band remote management port, such as a COM port or the redirection port specified by the hardware platform, by using terminal software that supports VT-UTF8, or VT100. SAC is a kernel-mode component.

### Note

- Command-line functionality for Emergency Management Services requires mutual dependence across three components: Sacdrv.sys, Sacsvr.dll, and Sacsess.exe. The SAC driver provides the core channel functionality (as well as basic SAC> prompt features), while the SAC service and the SAC session register with the SAC driver to complete the functional loop. The service is responsible for responding to SAC driver requests to start sessions, and the session is responsible for creating a new SAC channel and routing **STDIO** between the channel and its child Cmd.exe process. At the same time, it is important to note that Sacdrv.sys was designed to allow Special Administration Console to continue to operate successfully, even if the system is in a state where Sacsvr.dll and Sacsess.exe cannot operate.

SAC is available early in the startup process, and you can use it to manage the server during normal system operation, Safe Mode, and the GUI phase of Windows Server 2003 Setup. When Emergency Management Services is enabled, SAC is always active as long as the kernel is running. SAC is a separate environment from the Windows Server 2003 command prompt (Cmd.exe). You can remotely monitor the status of a server or troubleshoot problems by using terminal software on a remote computer to access SAC.

**Remote Management Tasks Supported by SAC during Normal Operation**
You can use SAC to perform the following remote tasks when the operating system is functioning normally:

- Gather server information, such as computer name and IP address.
- Obtain a list of the services and applications that are running on the computer.
- Restart or shut down a server as part of planned maintenance.
- Change the system time and date.

**Remote Management Tasks Supported by SAC for Troubleshooting**
You can use SAC to perform the following troubleshooting tasks:

- Raise or lower the priority of a process, or end a process that is consuming excessive processor time or other system resources. By establishing an out-of-band connection to SAC, you can continue to manage a server even when low system resources have caused it to become unresponsive to management tools that rely on an in-band connection.
- Restart or shut down a server that has stopped responding to in-band commands and tools.
- View or change TCP/IP networking information for a server to resolve issues, such as problems caused by a duplicate IP address.
- Maintain communication with a server during network outages. For example, you might want to verify that a connectivity issue is caused by a problem that affects the primary network rather than a problem with the server hardware or a driver.

- Generate a list of drivers for diagnostic purposes.
- View Setup logs to determine the cause of a problem that prevents Windows Server 2003 installation from completing.

**Note**

- Use of SAC is not safeguarded by password and logon requirements. You must secure physical access to computers running Emergency Management Services by placing appropriate restrictions on all connecting terminals.

**Using SAC**

SAC is an operating system component that runs in kernel mode and displays the following prompt when you establish an out-of-band connection to a computer that is running Emergency Management Services:

```
SAC>
```

In addition to providing a variety of commands that support remote management, SAC provides access to the Windows command prompt and to Setup logs. When you use SAC, you can create multiple user sessions or *channels* and switch among them so that you can use SAC commands while concurrently running command-line commands or viewing Setup logs.

**SAC Commands**

The following table lists the commands that are available in the SAC environment. The commands are not case-sensitive. For some commands, you will have to specify a process identifier (PID)number. You can obtain PID numbers by using the **t** command, which causes SAC to list active processes and the PIDs that are assigned to them.

**Table 10.11 SAC Commands**

| SAC Command | Description |
|---|---|
| **?** or **help** | Lists the available SAC commands. |
| **ch** | Lists all available channels. See Table 10.12 for a complete list of channel management commands. |
| **cmd** | Creates a Windows command prompt channel once logon credentials are validated. |
| **Crashdump** | Manually generates a Stop message and forces a memory dump file to be created. |
| **D** | Dumps the current kernel log. |
| f | Switches the information output of the **tlist** command between showing processes only or showing both processes and threads. |
| I | Configures IP parameters by providing network number, IP address, and subnet information using the following format.<br><br>*<network#> <iPaddress> <subnet>*<br><br>If no parameters are passed, this command lists IP information.<br><br>**Note**<br>Please note that the IP address will be restored to the original value when the server is rebooted and any changes made by using this command will be lost. |
| **id** | Displays computer identification information. |
| **k** *<PID>* | Ends the process specified by its process ID (PID). |
| **l** *< PID>* | Lowers the priority of a process (and any associated child processes) to the lowest possible level. |
| **lock** | Locks access to command prompt channels. |

| | |
|---|---|
| **m** < *PID*> <*MB-allow*> | Limits the memory usage of a process (and any child processes) to the specified number of megabytes. |
| p | Causes **tlist** command output to pause after displaying one screen of information. |
| **r** < *PID*> | Raises the priority of a process and any associated child processes by one level. |
| **restart** | Restarts the computer immediately. |
| **s** | Displays the current date (using 24-hour clock format) if you provide no parameters.<br>Sets the system time if you provide the date and, optionally, the time in the following format:<br>*mm/dd/yyyy hh:mm* |
| **shutdown** | Shuts down the computer. |
| **t** | Lists the processes and threads that are currently running. The output is referred to as a *task list* or *tlist*. |

**Note**

- The SAC prompt (SAC>) might not immediately appear after a connection is established to a remote server running Emergency Management Services. If the SAC prompt does not appear, Press ENTER several times until the prompt does appear.

**SAC Channels**

SAC supports three types of channels: SAC, Windows command prompt, and Setup logs.

**SAC Channel**

The SAC channel, which is always active, is assigned a channel name of SAC and a channel number of 0 (zero).

**Windows Command Prompt Channels**

The Special Administration Console Helper service (Sacsvr) is a user-mode service that you can use to create command prompt channels when you type the **cmd** command from the SAC prompt. To use a Windows command prompt channel, you must log on to the server with administrative credentials. You must log on separately to each command prompt session. SAC assigns a name of "Cmd000*x*" to all command prompt channels, where *x* represents the next available channel number.

When using a Windows command prompt channel, you can run text-based tools as though you were logged on locally. These tools include command-shell commands, such as **dir**, and text-based console programs, such as the boot configuration tool (Bootcfg.exe).

As with other services, the Special Administration Console Helper service requires initialization time when Windows Server 2003 starts. Therefore, command prompt channels might not be immediately available, and you might have to wait until the operating system startup process is nearly complete.

**Note**

- The Windows command prompt might not be available or might become unavailable if system resources become low. In these situations, you can typically continue to manage the server by using SAC.

**Setup Log Channels**

You can access the Setup log channels only during the GUI-mode phase of Windows Server 2003 Setup. By viewing the following Setup log files, which contain the same information, you can monitor the installation process or diagnose problems that might cause Setup to stop or pause. These files are created automatically.

- **Setuplog.txt.** Provides information about files copied to your system by Windows Setup.
- **Setupact.log.** Provides information about actions performed by Windows Setup.
- **Setuperr.log.** Provides Setup-related warning and error information.

The following text shows output that is generated by the SAC **ch** command during the GUI-mode phase of Windows Server 2003 Setup.

```
SAC>ch
Channel List

(Use "ch -?" for information on using channels)

# Status  Channel Name
0 (AV)    SAC
1 (AR)    setuplog.txt
2 (AR)    setupact.log
3 (AR)    setuperr.log
```

**SAC Channel Commands**

You can create and switch among up to ten concurrent channels. The following table provides a list of channel commands.

**SAC Channel Commands**

| Channel Command | Description |
|---|---|
| **ch** | Lists all channels. |
| **ch -si** <#> | Switches to a channel by number. Press ENTER to confirm your channel selection. |
| **ch -sn** <*name*> | Switches to a channel by name. Press ENTER to confirm your channel selection. |
| **ch -ci** <#> | Closes a channel by number. Press ENTER to confirm your channel selection. |
| **ch -cn** <*name*> | Closes a channel by name. Press ENTER to confirm your channel selection. |
| **ESC** followed by **TAB** | Cycles through the available channels. Press ENTER to confirm your channel selection. |
| **ESC** then **TAB** followed by **0** | Returns to channel 0 (the SAC channel). Press ENTER to confirm your channel selection. |

SAC returns channel status information in the following format:

<*Channel #*> <*Two-character status*> <*Channel name*>

The following illustrates output generated by the SAC **ch** command:

```
SAC>ch
Channel List

(Use "ch -?" for information on using channels)

# Status  Channel Name
0 (AV)    SAC
1 (AV)    Cmd0001
2 (AV)    Cmd0002
```

To switch from SAC (Channel 0) to the next available channel (Channel 1), press ESC followed by TAB, and then press ENTER. To return to SAC, press ESC followed by TAB, type 0, and then press ENTER.

The following table describes the information provided by the two-character channel status code.

**SAC Channel Status Information**

| Status Code (AB) | Description |
|---|---|
| **A (First Character)** | Indicates channel status.<br><br>A = Active Channel, I = Inactive Channel |
| **B (Second Character)** | Channel terminal emulation type.<br><br>V = VT-UTF8, R = Raw (no emulation) |

**Note**

- Although you can create and use multiple channels to manage a computer running Windows Server 2003, you can only

access a single channel at a time. Therefore, access to SAC by multiple users is not possible.

**!Special Administration Console (resides in Ntoskrnl.exe)**
**Category**
Windows Server 2003 operating system tool

**Version compatibility**
Supported on all Windows Server 2003 operating systems

The !Special Administration Console (!SAC) is an auxiliary Emergency Management Services command-line environment that is hosted on Windows Server 2003. It also accepts input and sends output through the out-of-band port. !SAC is a separate entity from both SAC and Windows Server 2003 family command-line environments. After a specific failure point is reached, Emergency Management Services components determine when the shift should be made from SAC to !SAC. !SAC becomes available automatically if SAC fails to load or is not functioning.

!SAC provides a subset of SAC commands that you can use to restore system functionality during a failure. Its two primary functions are:

- To redirect Stop error message text.
- To restart the computer if SAC becomes unavailable.

**Using !SAC**
!SAC provides a command-line environment that you can use if a computer running Windows Server 2003 stops responding or SAC has failed to load or ceased to operate.

When !SAC becomes available, the following prompt appears:

```
!SAC>
```

The !SAC prompt might be preceded by a Stop message that contains information about the problem that caused the error condition to occur.

**!SAC Commands**
The following table lists available !SAC commands.

### !SAC Commands

| !SAC Command | Description |
|---|---|
| **?** or **help** | Lists available commands. |
| **D** | Displays all log entries (screen pauses at each page of information). |
| **Id** | Displays computer identification information. |
| **Restart** | Restarts the computer. |

Like SAC, it is not possible for multiple users to access !SAC simultaneously.

### Note

- The !SAC prompt might not immediately appear after you establish a connection to a remote server running Emergency Management Services. If the !SAC> prompt does not appear, press ENTER several times until it does appear.

**Preventing Automatic Restarts after a Stop Message**
An **Automatically reboot** setting in Control Panel controls whether your system restarts after a Stop message occurs. This option is enabled by default. It causes your computer to restart after a Stop message instead of enabling !SAC. To use !SAC after a Stop message occurs, you must disable the **Automatically reboot** setting. This setting might be important for troubleshooting, depending on the nature of the problem (for example, the computer is repeatedly failing after restart), because it gives you time to remotely view and record Stop message text for later reference while searching through technical resources, such as Microsoft Knowledge Base articles. In most cases, however, the **Automatically reboot** setting is preferred for computers in remote locations.

**Emergency Management Services Registry Entries**
The registry entries in this section are associated with Emergency Management Services.

The information here provided is a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

For more information about the Emergency Management Services registry entries, see the Registry Reference for Windows Server 2003.

**The following registry entries govern core Emergency Management Services Special Administration Console functionality, which includes the SAC command set as well as virtual channel support.**

**ComputerName**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\ComputerName\ComputerName\

**Version**
Windows Server 2003

Access required: Read

Function: Gets computer name for computer information.

**DisableCmdSessions**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SacDrv\

**Version**
Windows Server 2003

Access required: Read

Function: Prevents the **cmd** command from starting a SAC session.

**PROCESSOR_ARCHITECTURE**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Session Manager\Environment\

**Version**
Windows Server 2003

Access required: Read

Function: Gets processor architecture for computer information.

**Start**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SacSvr\

**Version**
Windows Server 2003

Access required: Write

Function: Forces the SAC service attribute to automatic if the mode is manual.

**SystemSetupInProgress**
**Registry path**
HKEY_LOCAL_MACHINE\System\Setup\

**Version**
Windows Server 2003

Access required: Read

Function: Determines if Setup is in progress. If Setup is in progress, then session locking is disabled – there is no password for

the administrator during GUI mode Setup

**The following registry entries govern Cmd.exe functionality by means of a SAC channel. That is, the session implements a SAC channel and routes Cmd.exe I/O to the channel so that it can be accessed by means of remote users.**

**BypassAuthentication**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Services\SacSvr\

**Version**
Windows Server 2003

Access required: Read

Function: Deliberately disables session authentication.

Who accesses the setting: Administrators

**LoadProfilesDisabled**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SacSvr\Parameters\

**Version**
Windows Server 2003

Access required: Read

Function: Deliberately disables the loading of user profiles upon successful session authentication.

Who accesses the setting: Administrators

**Netsvcs**
**Registry path**
HKEY_LOCAL_MACHINE\ Software\Microsoft\Windows NT\CurrentVersion\Svchost

**Version**
Windows Server 2003

Access required: Read/Write

Function: States whether the SAC service is a member of the Netsvcs group. If it is not, it should be added.

What accesses the setting: Service Control Manager

**ServiceDll**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SacSvr\Parameters\

**Version**
Windows Server 2003

Access required: Write/Create

Function: Specifies the location of the SAC service dll binary file.

What accesses the setting: Service Control Manager

**TimeOutDisabled**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SacSvr\Parameters\

**Version**
Windows Server 2003

Access required: Read

Function: Deliberately disables the time-out (auto-session lock) feature.

Who accesses the setting: Administrators

**TimeOutInterval**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SacSvr\Parameters\

**Version**
Windows Server 2003

Access required: Read

Function: Sets the time-out (auto-session lock) duration

Who accesses the setting: Administrators

**The following registry entries enable the Service Control Manager to successfully start the SAC service. The purpose of the SAC service (Special Administration Console Helper service) is to provide a mechanism for the SAC driver (sacdrv.sys) to start user-mode instances of SAC sessions (Sacsess.exe). Note that the Start value is set to offstate (4) when the service exits.**

**SacSvr**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\

**Version**
Windows Server 2003.

Access required: Read

Function: Populates the SAC service parameters so that the Service Control Manager is properly informed and can start the service correctly.

What accesses the setting: Service Control Manager

**Description**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SacSvr\

**Version**
Windows Server 2003.

Access required: Read

Function: This value of the SacSvr key populates the SAC service parameters so that the Service Control Manager is properly informed and can start the service correctly.

What accesses the setting: Service Control Manager

**DisplayName**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ SacSvr\

**Version**
Windows Server 2003.

Access required: Read

Function: This value of the SacSvr key populates the SAC service parameters so that the Service Control Manager is properly informed and can start the service correctly.

What accesses the setting: Service Control Manager

**ErrorControl**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ SacSvr\

**Version**
Windows Server 2003.

Access required: Read

Function: This value of the SacSvr key populates the SAC service parameters so that the Service Control Manager is properly

informed and can start the service correctly.

What accesses the setting: Service Control Manager

**ImagePath**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ SacSvr\

**Version**
Windows Server 2003.

Access required: Read

Function: This value of the SacSvr key populates the SAC service parameters so that the Service Control Manager is properly informed and can start the service correctly.

What accesses the setting? Service Control Manager

**ObjectName**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ SacSvr\

**Version**
Windows Server 2003.

Access required: Read

Function: This value of the SacSvr key populates the SAC service parameters so that the Service Control Manager is properly informed and can start the service correctly.

What accesses the setting: Service Control Manager

**Start**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ SacSvr\

**Version**
Windows Server 2003.

Access required: Read

Function: This value of the SacSvr key populates the SAC service parameters so that the Service Control Manager is properly informed and can start the service correctly.

What accesses the setting: Service Control Manager

**Type**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ SacSvr\

**Version**
Windows Server 2003.

Access required: Read

Function: This value of the SacSvr key populates the SAC service parameters so that the Service Control Manager is properly informed and can start the service correctly.

What accesses the setting: Service Control Manager

**Emergency Management Services Group Policy Settings**
There are no Group Policy settings associated with Emergency Management Services.

**Emergency Management Services WMI Classes**
There are no WMI classes associated with Emergency Management Services.

**Ports Used by Emergency Management Services**
**Network Ports**
As a technology that uses an out-of-band connection, Emergency Management Services does not directly utilize any network ports.

**Serial Ports**
Emergency Management Services redirects text output from certain tools and functions through an out-of-band connection, typically a serial port connection. However, this can also be a connection that is defined in the Serial Port Console Redirection (SPCR) table, if this table is supported by the system firmware. If Emergency Management Services is configured to use an SPCR table (including all Itanium-based systems), the following settings are available.

- Any COM port
- Any I/O port in memory mapped I/O space
- Other SPCR table information
- The serial port EFI console device path (Itanium-based systems only)

**Notes**

- The SPCR table is only available on computers with ACPI-compliant firmware. However, not all ACPI-compliant computers provide an SPCR table.
- The SPCR table is typically enabled if firmware console redirection is enabled or a service processor that supports Emergency Management Services is present. See your system documentation to determine whether your computer supports the SPCR table and whether you need to follow special requirements to enable it.

To find more information about the SPCR table, see MSDN and type the appropriate key words in the "Search for" text box.

Serial ports are also known as COM ports. The following table lists the COM ports used by Emergency Management Services and their settings. In older systems without firmware that supports the SPCR table, the user must configure these ports manually. However, it is important to emphasize that in order to achieve a higher degree of integration between Emergency Management Services and the system in general, it is preferable to define these connections in the SPCR table, if available.

**Typical Serial Port Resource Settings**

| Friendly Name | Typical Hexadecimal I/O Address | Interrupt Request (IRQ) Value |
|---|---|---|
| **COM1** | Emergency Management Services supports the industry standard 3F8 (this is typically the firmware default or AUTO value for COM1). Values of 2F8, 3F8, On, and Off might also be available. | N/A |
| **COM2** | Emergency Management Services supports the industry standard 2F8 (this is typically the firmware default or AUTO value for COM2). Values of 2F8, 3F8, On, and Off might also be available. | N/A |

**Notes**

- It is important to realize that Emergency Management Services does not use both COM1 and COM2 at the same time. It uses one port or the other. In addition, it can use a nonlegacy serial port address if a developer builds a feature or component that is written to the Emergency Management Services application programming interfaces (APIs) and specifies such an address.
- The Windows Server 2003 operating system must have exclusive access to the out-of-band port. Instead of requesting input and output operations from the Windows I/O Manager, Emergency Management Services writes information directly to Universal Asynchronous/Receive Transmit (UART) hardware. fallback

**Related Information**
The following resource contains additional information that is relevant to this section.

- For more information about Emergency Management Services, Group Policy and Telnet, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.

# Shutdown Event Tracker Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shutdown Event Tracker Technical Reference**
Shutdown Event Tracker is a feature of the Microsoft Windows Server 2003 operating systems that provides a way for IT professionals to consistently track why users restart or shut down their computers. Shutdown Event Tracker captures the reasons users give for restarts and shutdowns to help create a comprehensive picture of an organization's system environment. It does not document why users choose other options, such as **Log off** and **Hibernate**.

In a more active sense, Shutdown Event Tracker also provides IT professionals with a specific tool, Remote Shutdown (Shutdown.exe), for restarting or shutting down both local and remote computers, while at the same time supplying reasons for doing so. In addition, users can employ Remote Shutdown to hibernate a local computer and abort delayed shutdowns.

Shutdown Event Tracker is enabled by default and is a routine part of the computer shutdown process.

**In this subject**


- What Is Shutdown Event Tracker?
- How Shutdown Event Tracker Works
- Shutdown Event Tracker Tools and Settings

# What Is Shutdown Event Tracker?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Shutdown Event Tracker?**

### In this section


- Common Scenarios for Shutdown Event Tracker
- Shutdown Event Tracker Dependencies on or Interactions with Other Technologies
- Shutdown Event Tracker Logical Diagrams
- How Shutdown Event Tracker Differs from Similar Technologies
- How Shutdown Event Tracker Works


Shutdown Event Tracker is a feature of the Microsoft Windows Server 2003 operating systems that provides a way for IT professionals to consistently track why users restart or shut down their computers. Shutdown Event Tracker captures the reasons users give for restarts and shutdowns to help create a comprehensive picture of an organization's system environment. It does not document why users choose other options, such as **Log off** and **Hibernate**.

In a more active sense, Shutdown Event Tracker also provides IT professionals with a specific tool, Remote Shutdown (Shutdown.exe), for restarting or shutting down both local and remote computers, while at the same time supplying reasons for doing so. In addition, users can employ Remote Shutdown to hibernate a local computer and cancel delayed shutdowns.

Shutdown Event Tracker is enabled by default and is a routine part of the computer shutdown process.

In first considering Shutdown Event Tracker, it is important to understand the difference between expected and unexpected restarts and shutdowns.

**Expected restarts and shutdowns**

Expected restarts and shutdowns can be either planned or unplanned.

A planned shutdown is one in which both the user and the computer fully anticipate the shutdown. For example, as a matter of either policy or habit, a user might shut down his or her computer at the end of each day. When a user has control over the timing of a restart or shutdown, the task is planned.

An unplanned shutdown is one in which the user does not anticipate the shutdown, but has time to perform the shutdown in a normal manner. For example, if an application becomes unresponsive, the user might be forced to restart or shut down the computer. When a user does not have control over the exact timing of a restart or shutdown, the task is unplanned.

### Note

- For the purpose of this topic, shutting down the system in a "normal manner" is defined as the user clicking **Start** and then **Shutdown**, or pressing CTRL+ALT+DELETE and then clicking **Shutdown**, or utilizing Shutdown.exe, a tool specific to Shutdown Event Tracker. These methods call the InitiateSystemShutdownEx application programming interface (API), which in turn spawns a number of related events. To find more information about these events, see "How Shutdown Event Tracker Works" in "How Shutdown Event Tracker Works."

**Unexpected restarts and shutdowns**

Unexpected shutdowns are shutdowns that the computer does not anticipate, and that the user may or may not anticipate. For example, the computer abruptly loses power and immediately shuts down. In this instance, neither the computer nor the user could have anticipated the shutdown. However, in another example, the user shuts down the computer by holding down the power button. In this instance, the user did, in fact, anticipate the shutdown but the computer, as in the first example, did not. It is important to understand that in both instances shutdown did not occur in a normal manner and that therefore the two shutdowns must be classified as "unexpected."

**Common Scenarios for Shutdown Event Tracker**

Shutdown Event Tracker is commonly used in the following scenarios:

**Expected restarts and shutdowns**
In this scenario, the user has complete control over the restart or shutdown process, and when the user performs the shutdown in a normal manner, the operating system has time to complete its usual shutdown routine. Moreover, because there is time, the user can provide a reason for the restart or shutdown in the dialog box presented to the user prior to shutdown.

In addition, it is important to understand that both planned and unplanned shutdowns form separate parts of the overall expected restarts and shutdowns scenario. They are described as follows:

**Planned restarts and shutdowns**
In this scenario, both the user and the computer fully anticipate the restart or shutdown: the user because he or she deliberately initiates the process, and the computer because the process is initiated in a normal manner, either by the user clicking **Start** and then **Shutdown**, or by pressing CTRL+ALT+DELETE and then clicking **Shutdown**,or by utilizing Shutdown.exe. When a restart or shutdown is planned, and shutdown occurs in a normal manner, the user is given the opportunity to provide a reason for the event at the time the event occurs.

**Unplanned restarts and shutdowns**
In this scenario, the user deliberately shuts down the computer, but in response to an unexpected event, such as an application failing to respond. The shutdown is therefore unplanned. If the user shuts down the computer in a normal manner, the operating system has time to execute the shutdown in a normal manner as well, and the user is given the opportunity to provide a reason for an "unplanned" shutdown. Conversely, if in response to the same unexpected event the user does not shut down the computer in a normal manner (for example, if the user unplugs the computer), the shutdown is no longer viewed as expected but unplanned, but as simply "unexpected."

**Unexpected restarts and shutdowns**
In this scenario, the user has no control over the restart or shutdown process, and the computer does not have time to complete its usual shutdown routine. For example, in the event of a power loss, the computer will immediately shut down unless there is an uninterruptible power supply (UPS) or other fault-tolerant strategy in place. As a result, the user will not be given the opportunity to provide a reason for the shutdown at the time the shutdown occurs. However, the **Shutdown Event Tracker** (unexpected shutdown) dialog box appears to the first person with the **Shutdown the system** user right or with administrative credentials who logs on to the computer after the event, and the user can, at that time, provide a reason for the shutdown.

### Note

- It is important to realize that a person with administrative credentials automatically has the **Shutdown the system** user right, but that it is possible for another person to have only this one right and not have administrative credentials.

Other reasons for unexpected shutdowns include (but are not limited to):

- A hardware failure
- A system error (also referred to as a bugcheck, system crash, fatal system error, or stop error)
- The user holding down the computer's power button for five seconds or longer
- The user pressing CTRL+ALT+DELETE and then holding down the CTRL key while clicking **Shutdown**

**Local restarts and shutdowns**
In this scenario, Remote Shutdown (Shutdown.exe) enables users to restart or shut down a local computer by either of two means: 1) the graphical user interface (GUI), invoked by typing **Shutdown /i** at the command prompt, or 2) the same **Shutdown** command used in combination with various *other* command-line parameters (for example, **Shutdown /s**, which causes the computer to shut down after a short interval).

**Remote restarts and shutdowns**
In this scenario, Remote Shutdown (Shutdown.exe) enables users to restart or shut down one or more remote computers by either of two means: 1) the graphical user interface (GUI), invoked by typing **Shutdown /i** at the command prompt, or 2) the same **Shutdown** command used in combination with various *other* command-line parameters (for example, **Shutdown -m** [*ComputerName*], which specifies the computer that the user wants to shut down).

### Note

- Although Remote Shutdown can be used to restart or shut down both local and remote computers, its primary purpose is to control the shutdown behavior of remote computers.

In addition, IT professionals can use this tool to perform remote bulk annotations of unexpected shutdowns, an alternative to the time-consuming task of logging on to each computer to record a reason for an unexpected shutdown. For example, a thousand computers in a datacenter all shut down at the same time because there is a catastrophic loss of power to the entire facility. In this circumstance, a user with administrative credentials can later record the same shutdown reason (power loss) in a single place (Remote Shutdown) for all one thousand computers, and is not required to log on to each computer to perform the same function over and over again.

To find more information about Shutdown Event Tracker command-line parameters, see Command Line References in Tools and Settings Collection.

**Shutdown Event Tracker Dependencies on or Interactions with Other Technologies**
Shutdown Event Tracker depends on, or interacts with, the following technologies:

- Windows Base Services
- Event log
- Registry
- Group Policy
- Custom Reason Editor
- System State Data feature
- System State Data Formatter (SSDFormat)
- Windows Error Reporting (WER)
- Poolmon

To find more information about this topic, see "How Shutdown Event Tracker Works" in "How Shutdown Event Tracker Works."

**Shutdown Event Tracker Logical Diagrams**
The following figures show the logical interactions between Shutdown Event Tracker components and between Shutdown Event Tracker and other Windows components.

**Shutdown Event Tracker Logical Diagram — Expected Shutdown**




**Shutdown Event Tracker Logical Diagram — Computer Startup After Expected or Unexpected Shutdown**




**Shutdown Event Tracker Logical Diagram — User Logon After Expected or Unexpected Shutdown**




To find more information about the terms used in these diagrams (for example, Savedump.exe and Event ID 6008), click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection. See also MSDN and type the appropriate key words in the "Search for" text box.

**How Shutdown Event Tracker Differs from Similar Technologies**

Other technologies, procedures, and actions in addition to Shutdown Event Tracker can be used to shut down Windows Server 2003. These include:

- Tssshutdn.exe
- Emergency Management Services (EMS)
- Windows Management Instrumentation (WMI)
- Computer Management MMC (Microsoft Management Console)
- Earlier versions of Shutdown.exe
- Pressing CTRL+ALT+DELETE, and then holding down the CTRL key while clicking Shutdown

**Note**

- This action may cause permanent data loss and should therefore be used with caution.

However, unlike Shutdown Event Tracker, none of these technologies enable the user to provide a reason for the shutdown.

**Related Information**

The following resources contain additional information that is relevant to this section.

- To find more information about Custom Reason Editor, click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection. After downloading Windows Resource Kit Tools, click **Start**, **All Programs**, **Windows Resource Kit Tools**, and then **Windows Resource Kit Tools Read Me**.

- To find more information about System State Formatter (SSDFormat), click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

- To find more information about Windows Error Reporting (WER), click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

- To find more information about Poolmon.exe, click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

- To find more information about InitiateSystemShutdown, InitiateSystemShutdownEx, and ExitWindowsEx (the shutdown APIs), see MSDN and type the appropriate key words in the "Search for" text box.

- To find more information about the terms used in the logical diagrams (for example, Savedump.exe and Event ID 6008), click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

- To find more information about Shutdown Event Tracker command-line parameters, see Command Line References in Tools and Settings Collection.

# How Shutdown Event Tracker Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Shutdown Event Tracker Works**

**In this section**


- Terms and Definitions
- Shutdown Event Tracker Architecture
- Shutdown Event Tracker Protocols
- Shutdown Event Tracker Interfaces
- Shutdown Event Tracker Processes and Interactions
- Network Ports Used by Shutdown Event Tracker
- Related Information


Shutdown Event Tracker is a feature of the Microsoft Windows Server 2003 operating systems that provides a way for IT professionals to consistently track why users restart or shut down their computers. Shutdown Event Tracker captures the reasons users give for restarts and shutdowns to help create a comprehensive picture of an organization's system environment. It does not document why users choose other options, such as **Log off** and **Hibernate**.

In a more active sense, Shutdown Event Tracker also provides IT professionals with a specific tool, Remote Shutdown (Shutdown.exe), for restarting or shutting down both local and remote computers, while at the same time supplying reasons for doing so. In addition, users can employ Remote Shutdown to hibernate a local computer and cancel delayed shutdowns.

Shutdown Event Tracker is enabled by default and is a routine part of the computer shutdown process.

The following sections provide an in-depth view of how Shutdown Event Tracker works in an optimal environment. An optimal environment for Shutdown Event Tracker is defined as follows:

- Windows Server 2003 is correctly installed and Shutdown Event Tracker is enabled.
- All other dependencies are in place, properly designed and deployed, and functioning normally.
- Server hardware is sized appropriately and there are no disk, CPU, memory, or network bottlenecks that affect the performance of the technology.

**Terms and Definitions**

The following terms are associated with Shutdown Event Tracker.

**Expected Restarts and Shutdowns**

Expected restarts and shutdowns can be either planned or unplanned. These shutdown types are explained as follows:

A *planned shutdown* is one in which both the user and the computer fully anticipate the shutdown. For example, as a matter of either policy or habit, a user might shut down his or her computer at the end of each day. When a user has control over the timing of a restart or shutdown, the task is planned.

An *unplanned shutdown* is one in which the user does not anticipate the shutdown, but has time to perform the shutdown in a normal manner. For example, if an application becomes unresponsive, the user might be forced to restart or shut down the computer. When a user does not have control over the exact timing of a restart or shutdown, the task is unplanned.

**Note**

- For the purpose of this topic, shutting down the system in a "normal manner" is defined as the user clicking **Start** and then **Shutdown**, or pressing CTRL+ALT+DELETE and then clicking **Shutdown**, or utilizing Remote Shutdown (Shutdown.exe), a tool specific to Shutdown Event Tracker. These methods call the InitiateSystemShutdownEx application programming interface (API), which in turn spawns a number of related events. To find more information about these events, see "Shutdown Event Tracker Architecture."

#### Unexpected Restarts and Shutdowns

Unexpected shutdowns are shutdowns that the computer does not anticipate, and that the user may or may not anticipate. For example, the computer abruptly loses power and immediately shuts down. In this instance, neither the computer nor the user could have anticipated the shutdown. However, in another example, the user shuts down the computer by holding down the power button. In this instance, the user did, in fact, anticipate the shutdown but the computer, as in the first example, did not. It is important to understand that in both instances shutdown did not occur in a normal manner and that therefore the two shutdowns must be classified as "unexpected."

#### Shut Down Windows (Expected Shutdown) Dialog Box

The **Shut Down Windows** (expected shutdown) dialog box appears when users restart or shut down the operating system in a normal manner. This dialog box differs from other Windows shutdown dialog boxes (for example, in earlier versions of Windows, or when Shutdown Event Tracker has been disabled) in that it prompts users to supply a reason and a comment to explain the action. In addition, since both planned and unplanned shutdowns are performed in a normal manner, the **Shut Down Windows** (expected shutdown) dialog box appears in either instance.

### Note

- Though titled as **Shut Down Windows**, this dialog box functions, in essence, as the expected shutdown dialog box.

#### Shutdown Event Tracker (Unexpected Shutdown) Dialog Box

When an unexpected shutdown occurs the user is not given the opportunity to perform the shutdown in a normal manner, and the **Shut Down Windows** (expected shutdown) dialog box does not appear at the time of shutdown. Instead, the **Shutdown Event Tracker** (unexpected shutdown) dialog box appears to the first person with the **Shutdown the system** user right or with administrative credentials who logs on to the computer after the restart or shutdown. Like the **Shut Down Windows** (expected shutdown) dialog box, it prompts this user to supply a reason and a comment for the shutdown.

### Notes

- Though titled as **Shutdown Event Tracker**, this dialog box functions, in essence, as the unexpected shutdown dialog box.

- It is important to realize that a person with administrative credentials automatically has **the Shutdown the system** user right, but that it is possible for another person to have only this one right and not have administrative credentials.

#### Shutdown Event Tracker Architecture

From an architectural standpoint, Shutdown Event Tracker is divided into three areas:

- The Shutdown Event Tracker interface (the expected shutdown dialog and the unexpected shutdown dialog)

- Other Windows technologies with which Shutdown Event Tracker interacts (Windows Base Services; Event Log; Registry; Group Policy; System State Data feature; Windows Error Reporting [WER]; Poolmon.exe or Memory Pool Monitor)

- Tools specific to Shutdown Event Tracker (Remote Shutdown [Shutdown.exe]; SSDFormat or System State Data Formatter; Custom Reason Editor)

The following figure shows the components within these three areas and how the areas and components interact.

### Shutdown Event Tracker Architectural Diagram




The following table lists the components in the first of the three areas described previously.

**Shutdown Event Tracker Components: Interface**

| Component | Description |
|---|---|
| Shut Down Windows (expected shutdown) dialog box | Appears at shutdown if shutdown occurs in a normal manner.<br><br>To find more information about the **Shut Down Windows** (expected shutdown) dialog box, see "Terms and Definitions" earlier in this section. |
| Shutdown Event Tracker (unexpected shutdown) dialog box | Appears at logon if DirtyShutdown key is present in the registry.<br><br>To find more information about the **Shutdown Event Tracker** (unexpected shutdown) dialog box, see "Terms and Definitions" earlier in this section. |

The following table lists the components in the second of the three areas described previously.

**Shutdown Event Tracker Components: Associated Windows Technologies**

| Technology | Description |
|---|---|

| | |
|---|---|
| Windows Base Services | The base services functions give applications access to the resources of the computer and the features of the underlying operating system, such as memory, file systems, devices, processes, and threads. An application uses these functions to manage and monitor the resources it needs to complete its work.<br><br>For Shutdown Event Tracker, Windows Base Services hosts components that log events to the event log; reads and interprets configuration information; and provides the public interface for initiating shutdowns. |
| Event log | The repository to which shutdown events are written by the expected shutdown dialog, the unexpected shutdown dialog, and Shutdown.exe. |
| Registry | Stores configuration data associated with Shutdown Event Tracker. Interacts with Custom Reason Editor, expected shutdown dialog (in providing custom reasons), unexpected shutdown dialog (in determining if a previous shutdown was unexpected), Shutdown.exe, and the System State Data feature. |
| Group Policy | Specifies configurations for groups of users and computers. These groups include software policies, scripts, user documents and configurations, application deployment, and security configurations.<br><br>To find more information about how Shutdown Event Tracker uses Group Policy, see "Shutdown Event Tracker Tools and Settings" in "Shutdown Event Tracker Tools and Settings." |
| System State Data feature | Gathers information for root-cause analysis of unplanned shutdowns. System state data is recorded in a log file when a user who has the **Shutdown the system** user right or administrative credentials specifies an unplanned reason for shutting down a computer. |
| Windows Error Reporting (WER) | WER is a set of technologies that captures product failure (also known as crash) data, enables end users to report failure information, and enables software and hardware vendors to analyze and respond to these problems.<br><br>WER is used by the System State Data feature to transmit shutdown data to Microsoft or other designated recipient.<br><br>To find more information about Windows Error Reporting, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection. |
| Memory Pool Monitor (Poolmon.exe) | Displays data that the operating system collects about memory allocations from the system paged and nonpaged kernel pools. This information can be used to find kernel mode memory leaks.<br><br>Poolmon can also be used to confirm findings made by the System State Data feature that are suggestive of memory leaks.<br><br>To find more information about Poolmon.exe, see the Microsoft Knowledge Base. |

The following table lists the components in the third of the three areas described previously.

**Shutdown Event Tracker Components: Tools Specific to the Feature**

| Tool | Description |
|---|---|
| Remote Shutdown (Shutdown.exe) | Enables users to restart or shut down a local computer or one or more remote computers by either of two means: 1) the graphical user interface (GUI), invoked by typing **Shutdown /i** at the command prompt, or 2) the same **Shutdown** command used in combination with various other command-line parameters (for example, **Shutdown /s**, which causes the computer to shut down after a short interval).<br><br>Users can also employ Remote Shutdown to hibernate local computers and to cancel delayed shutdowns from the command prompt.<br><br>In addition, IT professionals can perform remote bulk annotations of unexpected shutdowns.<br><br>To find more information about Shutdown Event Tracker command-line parameters, see "Command Line References" in Tools and Settings Collection.<br><br>To find more information about bulk annotations, see, "Shutdown Event Tracker Processes and Interactions" later in this section. |

| | |
|---|---|
| System State Data Formatter (SSDFormat.exe) | A command-line tool that creates a formatted copy of a system state data log file. SSDFormat.exe opens a system state data log file, adds an XSL file header to format the data, fixes any characters from the original file that are not legitimate XML characters, and saves the changes to a new XML file. When the user opens the XML file created by SSDFormat in any XML-capable viewer, such as Microsoft Internet Explorer, the XSL style sheet formats the data into tabular form.<br><br>To find more information about SSDFormat, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection. |
| Custom Reason Editor (CustomReasonEdit.exe) | A command-line and graphical user interface (GUI) tool that allows users to add, modify, and delete custom shutdown reasons for Shutdown Event Tracker.<br><br>To find more information about Custom Reason Editor, click "Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection." After downloading Windows Resource Kit Tools, click **Start**, **All Programs**, **Windows Resource Kit Tools**, and then **Windows Resource Kit Tools Read Me**. |

**Shutdown Event Tracker Protocols**

Shutdown Event Tracker uses the following protocols:

- RPC (remote procedure call), when communicating with a remote server
- HTTPS (Hypertext Transfer Protocol over Secure Sockets Layer), which is used by Windows Error Reporting (WER)

Both RPC and HTTPS reside at the application layer of the OSI (Open Systems Interconnection) model.

**Shutdown Event Tracker Interfaces**

The following APIs are associated with Shutdown Event Tracker.

### Shutdown Event Tracker APIs

| API Name | Description |
|---|---|
| InitiateSystemShutdown | Initiates a shutdown and optional restart of the specified computer. Does not support Shutdown Event Tracker and the use of shutdown reason options. Should not be used by a developer who wishes to utilize Shutdown Event Tracker functionality. |
| InitiateSystemShutdownEx | Initiates a shutdown and optional restart of the specified computer.<br><br>Supports Shutdown Event Tracker and the use of shutdown reason options. Should be used by a developer who wishes to utilize Shutdown Event Tracker functionality. |
| ExitWindowsEx | Either logs off the current user and shuts down the system, or shuts down and restarts the system. It sends the WM_QUERYENDSESSION message to all applications to determine if they can be terminated.<br><br>In the unlikely event that ExitWindowsEx fails to shut down the system, or the shutdown is cancelled by the user, it writes an event to the event log. |

To find more information about Shutdown Event Tracker interfaces, see MSDN and type the appropriate key words in the "Search for" text box.

**Shutdown Event Tracker Processes and Interactions**

The following is a general description of how Shutdown Event Tracker functions at a basic level.

As stated earlier, Shutdown Event Tracker provides a simple and standard mechanism for IT professionals to consistently document the reasons for shutting down or restarting computers. The information captured can then be used to analyze the root causes of shutdowns and to develop a more complete understanding of the system environment.

To record a restart or shutdown reason, users enter the reason into either the **Shut Down Windows** (expected shutdown) dialog box or the **Shutdown Event Tracker** (unexpected shutdown) dialog box. The **Shut Down Windows** dialog box is displayed when shutdown occurs in a normal manner (an expected shutdown); the **Shutdown Event Tracker** dialog box appears when shutdown does not occur in a normal manner (an unexpected shutdown). In each instance, Windows provides

users with a list of predefined shutdown reasons from which to choose. Users can also employ reason options that they have created using the Custom Reason Editor. The information that users provide is recorded in the system log in Event Viewer.

For these processes and interactions to work as described, Shutdown Event Tracker should reside in an optimal environment. An optimal environment for Shutdown Event Tracker is defined as follows:

- Windows Server 2003 is correctly installed and Shutdown Event Tracker is enabled.
- All other dependencies are in place, properly designed and deployed, and functioning normally.
- Server hardware is sized appropriately and there are no disk, CPU, memory, or network bottlenecks that affect the performance of the technology.

**Shutdown Event Tracker Technologies**

The following is a more detailed description of specific technologies that either form a part of, or are closely identified with, Shutdown Event Tracker.

**Remote Shutdown**

Remote Shutdown (Shutdown.exe) enables users to restart or shut down a local computer or one or more remote computers by either of two means: 1) the graphical user interface (GUI), invoked by typing **Shutdown /i** at the command prompt, or 2) the same **Shutdown** command used in combination with various other command-line parameters (for example, **Shutdown -m** [*ComputerName*], which specifies the computer that the user wants to shut down).

### Note

- Although Remote Shutdown can be used to restart or shut down both local and remote computers, its primary purpose is to control the shutdown behavior of remote computers.

IT professionals can also use this tool to perform remote bulk annotations of unexpected shutdowns, an alternative to the time-consuming task of logging on to each computer to record a reason for an unexpected shutdown. For example, a thousand computers in a datacenter all shut down at the same time because there is a catastrophic loss of power to the entire facility. In this circumstance, a user with administrative credentials can later record the same shutdown reason (power loss) in a single place (Remote Shutdown) for all one thousand computers, and is not required to log on to each computer to perform the same function over and over again.

In addition, users can employ Remote Shutdown to hibernate local computers and to cancel delayed shutdowns from the command prompt.

To find more information about Shutdown Event Tracker command-line parameters, see "Command Line References" in Tools and Settings Collection.

**System State Data feature**

The System State Data feature gathers information for root-cause analysis of expected but unplanned shutdowns, and of unexpected shutdowns.

System state data is written to a log file when a user who has the **Shutdown the system** user right or administrative credentials specifies an "unplanned" reason for shutting down the computer in the case of an expected, but unplanned, shutdown. This file is stored in the *%windir%\system32\LogFiles\Shutdown* directory. The first user with administrative credentials who logs on to the computer after the shutdown will see a dialog box notification that reads, "The system has restarted after an unplanned shutdown. A log of this event has been created." From this dialog box, the administrator can navigate to a more detailed description of the system state data file, its contents, and the Microsoft privacy policy for data collection on the Web. The administrator can choose whether to send the system state data file to Microsoft by clicking either the **Send Error Report** or **Don't Send** button.

**System State Data Formatter**

System State Data Formatter (SSDFormat.exe) is a command-line tool that creates a formatted copy of a system state data log file. SSDFormat.exe opens a system state data log file, adds an XSL file header to format the data, fixes any characters from the original file that are not legitimate XML characters, and saves the changes to a new XML file. When the user opens the XML file created by System State Data Formatter in any XML-capable viewer, such as Internet Explorer, the XSL style sheet formats the data into tabular form.

To find more information about System State Data Formatter, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.

**Custom Reason Editor**
Custom Reason Editor (CustReasonEdit.exe) is a command-line and GUI tool that enables users to add, modify, and delete custom shutdown reasons. Users can employ command-line functionality to perform basic importing and exporting of custom reasons to a registry file, or they can use the GUI to perform all other types of custom reason editing. CustReasonEdit includes a set of sample reasons as a guide to new users.

To find more information about Custom Reason Editor, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection. After downloading Windows Resource Kit Tools, click **Start**, **All Programs**, **Windows Resource Kit Tools**, and then **Windows Resource Kit Tools Read Me**.

**Other Processes and Interactions**
The following is a brief description of other processes and interactions associated with Shutdown Event Tracker.

**Enabling or Disabling Shutdown Event Tracker**
The user can employ either Group Policy or the registry to enable or disable Shutdown Event Tracker on the user's system.

When Shutdown Event Tracker is enabled, the **Shut Down Windows** (expected shutdown) dialog box appears when the user clicks **Start** and then clicks **Shutdown**, or when the user presses CTRL+ALT+DELETE and then clicks **Shutdown**. If an unexpected restart or shutdown occurs, the **Shutdown Event Tracker** (unexpected shutdown) dialog box will appear to the first person with the **Shutdown the system** user right or with administrative credentials who logs on to the computer.

When Shutdown Event Tracker is disabled, the standard Windows shutdown dialog box appears. It provides a way to restart or shut down the users computer, although it does not provide a way to record whether or not the event is planned or unplanned, or a way to record a reason for it. If an unexpected restart or shutdown occurs under these circumstances, the **Shutdown Event Tracker** (unexpected shutdown) dialog box will *not* appear to the first person with the **Shutdown the system user** right who logs on to the computer.

To find more information about enabling or disabling "Shutdown Event Tracker, see "Shutdown Event Tracker Tools and Settings" in "Shutdown Event Tracker Tools and Settings." See also "Shutdown Event Tracker Tools and Settings" in "Shutdown Event Tracker Tools and Settings."

**Using Shutdown Event Tracker with Non-English Versions of Windows Server 2003**
Users can configure Shutdown Event Tracker to display standard and custom shutdown reasons in languages other than English by modifying the registry.

For example, to display the *standard* shutdown reasons in Romanian, the system version must be Romanian, and the user locale must be set to Romanian as well.

To display *custom* shutdown reasons in Romanian (or another language), users must employ the Custom Reason Editor, which writes custom shutdown reasons to the registry.

To find more information about Custom Reason Editor, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection. After downloading Windows Resource Kit Tools, click **Start**, **All Programs**, **Windows Resource Kit Tools**, and then **Windows Resource Kit Tools Read Me**.

**Network Ports Used by Shutdown Event Tracker**
The following table lists the port assignments for transmitting data collected by Shutdown Event Tracker.

**Port Assignments for Shutdown Event Tracker**

| Service Name | User Datagram Protocol (UDP) | Transmission Control Protocol (TCP) |
|---|---|---|
| Windows Error Reporting (WER) | Not applicable | 443 |
| Remote Procedure Call (RPC) | Not applicable | 445 and 139 |

Windows Error Reporting (WER) is a set of technologies built into Windows Server 2003 that captures product failure (also know as crash) data, enables end users to report failure information, and enables software and hardware vendors to analyze and respond to these problems.

When the **Report unplanned shutdown events** policy setting is enabled, error reporting will include unplanned shutdown events. When this policy setting is disabled, unplanned shutdown events will not be included in error reporting.

If this policy setting is not configured, the user is able to control unplanned shutdown reporting using Control Panel, which is set to **upload unplanned shutdown events** by default.

The System State Data feature uses Windows Error Reporting to transmit shutdown data to Microsoft or other designated

recipient.

**Note**

- Data will not be submitted to the designated recipient unless users click **Send Report** when presented with the **Windows Error Reporting** dialog box.

**Related Information**

The following resources contain additional information that is relevant to this section.

- To find more information about Custom Reason Editor, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection. After downloading Windows Resource Kit Tools, click **Start**, **All Programs**, **Windows Resource Kit Tools**, and then **Windows Resource Kit Tools Read Me**.

- To find more information about System State Formatter (SSDFormat), click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.

- To find more information about Windows Error Reporting (WER), click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.

- To find more information about Poolmon.exe, click "Windows Server 2003 Resource Kit Tools Help" in Tools and Settings Collection.

- To find more information about InitiateSystemShutdown, InitiateSystemShutdownEx, and ExitWindowsEx (the shutdown APIs), see MSDN and type the appropriate key words in the "Search for" text box.

- To find more information about Shutdown Event Tracker command-line parameters, see "Command Line References" in Tools and Settings Collection.

- To find more information about Shutdown Event Tracker Group Policy settings, click "Group Policy Settings Reference for Windows Server 2003" in Tools and Settings Collection.

# Shutdown Event Tracker Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shutdown Event Tracker Tools and Settings**

### In this section


- Shutdown Event Tracker Tools
- Shutdown Event Tracker Registry Entries
- Shutdown Event Tracker Group Policy Settings
- Network Ports Used by Shutdown Event Tracker


Shutdown Event Tracker is a feature of the Microsoft Windows Server 2003 operating systems that provides a way for IT professionals to consistently track why users restart or shut down their computers. Shutdown Event Tracker captures the reasons users give for restarts and shutdowns to help create a comprehensive picture of an organization's system environment. It does not document why users choose other options, such as **Log off** and **Hibernate**.

In a more active sense, Shutdown Event Tracker also provides IT professionals with a specific tool, Remote Shutdown (Shutdown.exe), for restarting or shutting down both local and remote computers, while at the same time supplying reasons for doing so. In addition, users can employ Remote Shutdown to hibernate a local computer and cancel delayed shutdowns.

Shutdown Event Tracker is enabled by default and is a routine part of the computer shutdown process.

**Shutdown Event Tracker Tools**

The following tools are associated with Shutdown Event Tracker.

**Shutdown.exe: Remote Shutdown**

**Category**

Remote Shutdown (Shutdown.exe) is part of Windows Server 2003.

**Version compatibility**

Remote Shutdown can also be used to control restarts or shutdowns on computers running Windows NT 4.0 and Windows 2000.

Remote Shutdown (Shutdown.exe) enables users to restart or shut down a local computer or one or more remote computers by either of two means: 1) the graphical user interface (GUI), invoked by typing **Shutdown /i** at the command prompt, or 2) the same **Shutdown** command used in combination with various *other* command-line parameters (for example, **Shutdown /s**, which causes the computer to shut down after a short interval).

IT professionals can also use this tool to perform remote bulk annotations of unexpected shutdowns, an alternative to the time-consuming task of logging on to each computer to record a reason for an unexpected shutdown.

In addition, users can employ Remote Shutdown to hibernate local computers and to cancel delayed shutdowns from the command prompt.

### Note

- Although Remote Shutdown can be used to restart or shut down both local and remote computers, its primary purpose is to control the shutdown behavior of remote computers.

To find more information about Shutdown Event Tracker command-line parameters, see Command Line References in Tools and Settings Collection.

To find more information about bulk annotations, see "Shutdown Event Tracker Processes and Interactions" in "How Shutdown Event Tracker Works."

**System State Data: System State Data feature**

**Category**

The System State Data feature is part of Windows Server 2003.

**Version compatibility**
The System State Data feature runs on and targets all Windows Server 2003 operating systems.

System State Data gathers information for root-cause analysis of expected but unplanned shutdowns, and of unexpected shutdowns.

System state data is written to a log file when a user who has the **Shutdown the system** user right or administrative credentials specifies an "unplanned" reason for shutting down the computer. This file is stored in the *%windir%\system32\LogFiles\Shutdown* directory. The first user with administrative credentials who logs on to the computer after the shutdown will see a dialog box notification that reads: "The system has restarted after an unplanned shutdown. A log of this event has been created." From this dialog box, the administrator can navigate to a more detailed description of the system state data file, its contents, and the Microsoft privacy policy for data collection on the Web. The administrator can choose whether to send the system state data file to Microsoft by clicking either the **Send Error Report** or **Don't Send** button.

**Windows Error Reporting: WER**
**Category**
Windows Error Reporting is part of Windows Server 2003.

**Version compatibility**
Windows Error Reporting runs on and targets all Windows Server 2003 operating systems.

WER is a set of technologies that captures product failure (also known as crash) data, allows end users to report failure information, and allows software and hardware vendors to analyze and respond to these problems.

When the **Report unplanned shutdown events** policy setting is enabled, error reporting will include unplanned shutdown events. When this policy setting is disabled, unplanned shutdown events will not be included in error reporting.

If this setting is not configured, the user is able to control unplanned shutdown reporting using Control Panel, which is set to **upload unplanned shutdown events** by default.

The System State Data feature uses Windows Error Reporting to transmit shutdown data to Microsoft or to another designated recipient.

### Note

- Data will not be submitted to the designated recipient unless users click **Send Error Report** when presented with the **Windows Error Reporting** dialog box.

To find more information about Windows Error Reporting, click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

**SSDFormat.exe: System State Data Formatter**
**Category**
System State Data Formatter (SSDFormat.exe) is available in Resource Kit Tools in Tools and Settings Collection.

**Version compatibility**
System State Data Formatter runs on all Windows Server 2003 operating systems.

SSDFormat.exe is a command-line tool that creates a formatted copy of a system state data log file. SSDFormat.exe opens a system state data log file, adds an XSL file header to format the data, fixes any characters from the original file that are not legitimate XML characters, and saves the changes to a new XML file. When the user opens the XML file created by SSDFormat in any XML-capable viewer, such as Microsoft Internet Explorer, the XSL style sheet formats the data into tabular form.

To find more information about System State Data Formatter, click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

**CustReasonEdit.exe: Custom Reason Editor**
**Category**
Custom Reason Editor (CustomReasonEdit.exe) is available in Resource Kit Tools in Tools and Settings Collection. After downloading Windows Resource Kit Tools, click **Start**, **All Programs**, **Windows Resource Kit Tools**, and then **Windows Resource Kit Tools Read Me** to locate the tool.

**Version compatibility**
Custom Reason Editor runs on all Windows Server 2003 operating systems.

CustomReasonEdit.exe enables users to add, modify, and delete custom shutdown reasons for Shutdown Event Tracker.

Custom Reason Editor requires that the user have administrative credentials on the target system to function correctly. The tool has two interfaces: the command-line interface allows for basic importing and exporting functionality, and the graphical user interface allows for complete custom reason editing. Finally, a set of sample reasons is included in the package.

To find more information about Custom Reason Editor, click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection. After downloading Windows Resource Kit Tools, click **Start**, **All Programs**, **Windows Resource Kit Tools**, and then **Windows Resource Kit Tools Read Me**.

**Poolmon.exe: Memory Pool Monitor**
**Category**
Memory Pool Monitor (Poolmon.exe) is available in the Windows NT 4.0 Resource Kit and in the \Support\Tools folder of Windows 2000, and on the Windows XP and Windows Server 2003 CD-ROMs.

**Version compatibility**
Memory Pool Monitor runs on Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003.

Poolmon.exe displays data that the operating system collects about memory allocations from the system paged and nonpaged kernel pools, and the memory pools used for Terminal Services sessions. The data is grouped by pool allocation tag. This information can be used by Microsoft Product Support Services to find kernel-mode memory leaks.

A memory leak is caused by an application or a process that allocates memory for use but does not free it up when finished. The result is that available memory is completely used over time, often causing the system to stop functioning properly.

Poolmon.exe is used to verify, or to further substantiate, memory usage information captured by the System State Data feature. In this sense, Memory Pool Monitor should be viewed as a non-interactive supplement to System State Data.

To find more information about Poolmon.exe, see MSDN and type the appropriate key words in the "Search for" text box.

**Shutdown Event Tracker Registry Entries**
Shutdown Event Tracker interacts with the registry in the following ways:

- The expected shutdown dialog reads custom shutdown reasons from the registry.
- Remote Shutdown (Shutdown.exe) reads custom shutdown reasons from the registry. It also writes bulk annotations to the registry and deletes keys from the registry.
- Custom Reason Editor (CustReasonEdit.exe) writes custom shutdown reasons to the registry.
- The unexpected shutdown dialog reads from the registry to determine if the previous shutdown was unexpected.
- The Event Log service writes the Shutdown Event Tracker heartbeat to the registry and then deletes it just before a normal shutdown occurs. On restart it verifies whether the heartbeat is present and, if so, writes the DirtyShutdown key to the registry.

### Note

- Used in this context, heartbeat is a defined as a time stamp interval, written once a minute, that tells the registry, and therefore the system, that Shutdown Event Tracker is still enabled.
- The System State Data feature reads system configuration information from the registry.

To find more information about these interactions, see How Shutdown Event Tracker Works.

The information here provided is a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following registry entries are associated with Shutdown Event Tracker.

**BugcheckString**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Reliability\

**Version**
Windows Server 2003.

This key contains the bug check string information that is used to fill in the unexpected shutdown dialog box comment field (which appears at logon after an unexpected shutdown) if the previous shutdown was caused by a system failure (also know as a system crash).

**DirtyShutdown**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Reliability\

**Version**
Windows Server 2003.

This key is set during event log startup. It indicates whether a previous shutdown was expected.

**LastAliveStamp**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Reliability\

**Version**
Windows Server 2003.

This key is cleared during shutdown. It indicates the date and time of the previous unexpected shutdown if it is present during startup.

**ReliabilityGUID**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Reliability\

**Version**
Windows Server 2003.

This key enables a GUID (globally unique identifier) to be written to the system state data file in order to uniquely identify the computer this file came from. It is not possible to physically identify the computer itself using this GUID, but it is possible to see how many different computers sent files and how many distinct reports were submitted by each computer. If the GUID is deleted from the registry, a new GUID is generated when a new system state data (.xml) file is created in the *%windir%\system32\LogFiles\Shutdown* directory at the time of an unplanned shutdown.

**ShutdownIgnorePredefinedReasons**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Reliability\

**Version**
Windows Server 2003.

This key prevents the predefined or built-in shutdown reasons from being displayed. If at least one custom reason is defined in the registry and this key is set to "1," the built-in reasons are not displayed.

**TimeStampInterval**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Reliability\

**Version**
Windows Server 2003 and Windows 2000.

This key defines how often LastAliveStamp (or heartbeat) is written to the registry. By default it is written every minute in Windows Server 2003 and every five minutes in Windows 2000.

**UserDefined**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Reliability\

**Version**
Windows Server 2003.

This key contains custom reasons stored as values. To add custom reasons, the user must define one key value for each reason. Each reason has a major and minor code that uniquely identifies the reason.

To find more information about Custom Reason Editor, click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection. After downloading Windows Resource Kit Tools, click **Start**, **All Programs**, **Windows Resource Kit Tools**, and then **Windows Resource Kit Tools Read Me**.

**HKEY_LOCAL_MACHINE\SOFTWARE**
The following registry entries are located under HKEY_LOCAL_MACHINE\SOFTWARE\.

**ShutdownReasonUI**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows NT\Reliability\

**Version**
Windows Server 2003.

This key is present if the **Display Shutdown Event Tracker** policy setting has been enabled.

**ShutdownReasonUI**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Reliability\

**Version**
Windows Server 2003.

Shutdown Event Tracker references the Group Policy key first (see the previous entry). If the Group Policy key is not present, then this key can be configured as “0” (off), or “1” (on). If the Group Policy key is not present and this key is invalid or missing, then Shutdown Event Tracker is off.

**Shutdown Event Tracker Group Policy Settings**
The following tables list and describe the Group Policy settings that are associated with Shutdown Event Tracker.

### Group Policy Settings Associated with Display Shutdown Event Tracker

| Group Policy Setting | Description |
|---|---|
| Not Configured | No choices are given. Specifies that no change is made to the registry for this setting. |
| Enabled | Shutdown Event Tracker is displayed. Specifies that the registry indicate that the policy setting applies to users or computers that are subject to this Group Policy object. |
| Disabled | Shutdown Event Tracker should not be displayed. Specifies that the registry indicate that the policy setting does not apply to users or computers that are subject to this Group Policy object. |

**Note**

- The path to Display Shutdown Event Tracker is Local Computer Policy \Computer Configuration\ Administrative Templates\System.

### Group Policy Settings Associated with Activate Shutdown Event Tracker System State Data Feature

| Group Policy Setting | Description |
|---|---|
| Not Configured | Specifies that no change is made to the registry for this policy setting. |
| Enabled | Specifies that the registry indicate that the policy setting applies to users or computers that are subject to this Group Policy object. |
| Disabled | Specifies that the registry indicate that the policy setting does not apply to users or computers that are subject to this Group Policy object. |

**Note**

- The path to Activate Shutdown Event Tracker System State Data feature is Local Computer Policy\Computer Configuration\Administrative Templates\System.

**Group Policy Settings Associated with Enable Persistent Time Stamp**

| Group Policy Setting | Description |
|---|---|
| Not Configured | Specifies that no change is made to the registry for this policy setting. |
| Enabled | Specifies that the registry indicate that the policy setting applies to users or computers that are subject to this Group Policy object.<br><br>The policy setting enables the user to customize how often the Persistent System Time Stamp is written to disk. The range is from 1 through 86400 seconds (1 day). |
| Disabled | Specifies that the registry indicate that the policy setting does not apply to users or computers that are subject to this Group Policy object. |

**Note**

- The path to Enable Persistent Time Stamp is Local Computer Policy\Computer Configuration\Administrative Templates\System.

**Group Policy Settings Associated with Report Unplanned Shutdown Events**

| Group Policy Setting | Description |
|---|---|
| Not Configured | Specifies that no change is made to the registry for this policy setting. |
| Enabled | Specifies that the registry indicate that the policy setting applies to users or computers that are subject to this Group Policy object.<br><br>When this policy setting is enabled, error reporting will include unplanned shutdown events. |
| Disabled | Specifies that the registry indicate that the policy setting does not apply to users or computers that are subject to this Group Policy object.<br><br>When this policy setting is disabled, unplanned shutdown events will not be included in error reporting. |

**Note**

- The path to Report unplanned shutdown events is Local Computer Policy \Computer Configuration\Administrative Templates\System\Error Reporting\Advanced Error Reporting settings.

To find more information about these Group Policy settings, click Group Policy Settings Reference for Windows Server 2003 in Tools and Settings Collection.

**Network Ports Used by Shutdown Event Tracker**

The following table lists the port assignments for transmitting data collected by Shutdown Event Tracker.

**Port Assignments for Shutdown Event Tracker**

| Service Name | User Datagram Protocol (UDP) | Transmission Control Protocol (TCP) |
|---|---|---|
| Windows Error Reporting (WER) | Not applicable | 443 |
| Remote Procedure Call (RPC) | Not applicable | 445 and 139 |

Windows Error Reporting (WER) is a set of technologies built into Windows Server 2003 that captures failure (also known as crash) data, enables end users to report failure information, and enables software and hardware vendors to analyze and respond to these problems.

When the **Report unplanned shutdown events** policy setting is enabled, error reporting will include unplanned shutdown events. When this policy setting is disabled, unplanned shutdown events will not be included in error reporting.

If this policy setting is not configured, the user is able to control unplanned shutdown reporting using Control Panel, which is set to **upload unplanned shutdown events** by default.

The System State Data feature uses Windows Error Reporting to transmit shutdown data to Microsoft or other designated recipient.

### Note

- Data will not be submitted to the designated recipient unless users click **Send Report** when presented with the **Windows Error Reporting** dialog box.

#### Related Information

The following resources contain additional information that is relevant to this section.

- To find more information about Custom Reason Editor, click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection. After downloading Windows Resource Kit Tools, click **Start**, **All Programs**, **Windows Resource Kit Tools**, and then **Windows Resource Kit Tools Read Me**.

- To find more information about System State Formatter (SSDFormat), click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

- To find more information about Windows Error Reporting (WER), click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

- To find more information about Poolmon.exe, click Windows Server 2003 Resource Kit Tools Help in Tools and Settings Collection.

- To find more information about InitiateSystemShutdown, InitiateSystemShutdownEx, and ExitWindowsEx (the shutdown APIs), see MSDN and type the appropriate key words in the "Search for" text box.

- To find more information about the Group Policy settings listed in this section, click Group Policy Settings Reference for Windows Server 2003 in Tools and Settings Collection.

- To find more information about Shutdown Event Tracker command-line parameters, see Command Line References in Tools and Settings Collection.

# Backup Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backup Technical Reference**
The Backup utility in the Microsoft Windows Server 2003 operating systems helps you protect your data if your hard disk fails or if files are accidentally erased due to hardware or storage media failure. By using Backup, you can create a duplicate copy of the data on your hard disk and then archive it on another storage device, such as a hard disk or a tape.

If the original data on your hard disk is accidentally erased or overwritten, or becomes inaccessible because of a hard-disk malfunction, you can easily restore it from the disk or archived copy.

**In this subject**


- What Is Backup?
- How Backup Works
- Backup Tools and Settings

# What Is Backup?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Backup?**
The Backup utility in the Microsoft Windows Server 2003 operating systems helps you protect your data if your hard disk fails or files are accidentally erased due to hardware or storage media failure. By using Backup, you can create a duplicate copy of the data on your hard disk and then archive it on another storage device, such as a hard disk or a tape.

If the original data on your hard disk is accidentally erased or overwritten, or becomes inaccessible because of a hard-disk malfunction, you can easily restore it from the disk or archived copy.

Backup uses the Volume Shadow Copy service to create an accurate copy of the contents of your hard drive, including any open files or files that are being used by the system. Users can continue to access the system while Backup is running, without risking loss of data. For more information about the Volume Shadow Copy service, see "Volume Shadow Copy Service Technical Reference."

Using Backup, you can:

- Archive selected files and folders on your hard disk.
- Restore the archived files and folders to your hard disk or any other disk you can access.
- Make a copy of your computer's System State data.
- Use Automated System Recovery (ASR) to create a backup set that contains the System State data, system services, and all disks associated with the operating system components. ASR also creates a floppy boot disk that contains information about the backup, the disk configurations (including basic and dynamic volumes), and how to restore your system.
- Make a copy of any Remote Storage data and any data stored in mounted drives.
- Create a log of what files were backed up and when the backup was performed.
- Make a copy of your computer's system partition, boot partition, and the files needed to start up your system in case of a computer or network failure.
- Schedule regular backups to keep your archived data up-to-date.

Backup also performs simple media management functions such as formatting. You can perform more advanced management tasks such as mounting and dismounting a tape or disk by using Removable Storage, which is a feature in Windows Server 2003.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Backup Technical Reference
- How Backup Works
- Volume Shadow Copy Service Technical Reference

# How Backup Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Backup Works**

**In this section**


- Types of Backup
- Volume Shadow Copy Service
- Files Skipped During Backup
- System State Data
- Permissions and User Rights Required to Back Up
- Automated System Recovery
- Restoring File Security Settings
- Restoring Distributed Services
- Related Information


The Backup utility in Windows Server 2003, which uses the Volume Shadow Copy service, helps you protect your data if your hard disk fails or files are accidentally erased due to hardware or storage media failure. By using Backup, you can create a duplicate copy of the data on your hard disk and then archive it on another storage device, such as a hard disk or a tape.

If the original data on your hard disk is accidentally erased or overwritten, or becomes inaccessible because of a hard-disk malfunction, you can easily restore it from the disk or archived copy.

**Types of Backup**

The Backup utility supports five methods of backing up data on your computer: a copy backup, daily backup, differential backup, incremental backup, and normal backup.

**Types of Backup**

| Type | Description |
|---|---|
| Copy backup | Copies all the files that you select, but does not mark each file as having been backed up (in other words, the archive attribute is not cleared). Copying is useful if you want to back up files between normal and incremental backups because copying does not affect these other backup operations. |
| Daily backup | Copies all the files that you select that have been modified on the day that the daily backup is performed. The backed-up files are not marked as having been backed up (in other words, the archive attribute is not cleared). |
| Differential backup | Copies files that have been created or changed since the last normal or incremental backup. It does not mark files as having been backed up (in other words, the archive attribute is not cleared). If you are performing a combination of normal and differential backups, you must have the last normal as well as the last differential backup to restore files and folders. |
| Incremental backup | Backs up only those files that have been created or changed since the last normal or incremental backup. It marks files as having been backed up (in other words, the archive attribute is cleared). If you use a combination of normal and incremental backups, you will need to have the last normal backup set as well as all incremental backup sets to restore your data. |

| | |
|---|---|
| Normal backup | Copies all the files that you select and marks each file as having been backed up (in other words, the archive attribute is cleared). With normal backups, you only need the most recent copy of the backup file or tape to restore all of the files. You usually perform a normal backup the first time you create a backup set. |

If you want the quickest backup method that requires the least amount of storage space, you should back up your data using a combination of normal backups and incremental backups. However, recovering files from this combination of backups can be time-consuming and difficult because the backup set might be stored on several disks or tapes.

If you want to restore your data more easily, you should back up your data using a combination of normal backups and differential backups. This backup set is usually stored on only a few disks or tapes. However, this combination of backups is more time-consuming.

**Volume Shadow Copy Service**

Backup uses the Volume Shadow Copy service to create a volume shadow copy, which is an accurate copy of the contents of your hard drive, including any open files, files that are being used by the system, and databases that are held open exclusively.

Backup uses the Volume Shadow Copy service to ensure that:

- Applications can continue to write data to the volume during a backup.
- Files that are open are no longer omitted during a backup.
- Backups can be performed at any time, without locking out users.

If you choose to disable the volume shadow copy using advanced options or if the service fails, Backup will revert to creating a backup without the Volume Shadow Copy service technology. If this occurs, Backup skips files that are open or in use by other applications at the time of the backup.

**Note**

- Some applications manage storage consistency differently while files are open, which can affect the consistency of the files in a backup. For critical applications, consult the application documentation or your provider for information about the recommended backup method. When in doubt, close the application before performing a backup.

For more information, see "Volume Shadow Copy Service Technical Reference."

**Files Skipped During Backup**

Backup skips certain files by default, including files in the following categories:

- **Files that the person performing the backup does not have permission to read.** Only users with backup rights can copy files that they do not own. Members of the Backup Operators and Administrators groups have these permissions by default.
- **Temporary files.** For example, Pagefile.sys, Hiberfil.sys, Win386.swp, 386spart.par, Backup.log, and Restore.log. These files are not backed up or restored by Backup. The list of skipped temporary files is embedded into Backup and cannot be changed.
- **Registry files on remote computers.** Windows Server 2003 backs up only local registry files.

**System State Data**

With Backup, you can back up the System State data for your computer. System State data includes the registry, the COM+ Class Registration database, files under Windows File Protection, and system boot files. Depending on the configuration of the server, other data might be included in the System State data. For example, if the server is a certificate server, the System State data also contains the Certificate Services database. If the server is a domain controller, the Active Directory directory service database and the SYSVOL directory are included in the System State data.

With Backup, the following system components might be included in a backup of the System State:

| Component | When this component is included in System State? |
|---|---|

| | |
|---|---|
| Registry | Always |
| COM+ Class Registration database | Always |
| Boot files, including the system files | Always |
| Certificate Services database | If it is a Certificate Services server |
| Active Directory database | If it is a domain |
| SYSVOL directory | Only if it is a domain controller |
| Cluster service information | If it is within a cluster |
| IIS Metadirectory | If it is installed |
| System files that are under Windows File Protection | Always |

When you choose to back up or restore the System State data, all of the System State data that is relevant to your computer is backed up or restored. You cannot choose to back up or restore individual components of the System State data because of dependencies among the System State components. However, you can restore the System State data to an alternate location. If you do this, only the registry files, SYSVOL directory files, cluster database information files, and system boot files are restored to the alternate location. The Active Directory database, Certificate Services database, and COM+ Class Registration database are not restored if you designate an alternate location when you restore the System State data.

**Files Under Windows File Protection**

Backup works together with the catalog file for the Windows File Protection service when backing up and restoring boot and system files. System files are backed up and restored as a single entity. The Windows File Protection service catalog file, located in the folder *systemroot*\system32\catroot\{F750E6C3-38EE-11D1-85E5-00C04FC295EE}, is backed up with the system files.

In Windows NT 4.0 and earlier, backup programs could selectively back up and restore operating system files. However, Windows 2000 Server and Windows Server 2003 do not allow incremental restores of operating system files.

There is an Advanced Backup option that automatically backs up protected system files with the System State data. All of the system files that are in the *systemroot*\ directory and the startup files that are included with the System State data are backed up when you use this option.

**Permissions and User Rights Required to Back Up**

You must have certain permissions and user rights to back up files and folders using Backup. If you are an administrator or a backup operator in a local group, you can back up any file and folder on the local computer to which the local group applies. Similarly, if you are an administrator or backup operator on a domain controller, you can back up any file and folder locally on any computer in the domain or any computer on a domain with which you have a two-way trust relationship. However, if you are not an administrator or a backup operator and you want to back up files, then you must be the owner of the files and folders that you want to back up, or you must have one or more of the following permissions for the files and folders you want to back up: Read, Read and Execute, Modify, and Full Control.

You must also be certain that there are no disk-quota restrictions that might restrict your access to a hard disk. These restrictions make it impossible for you to back up data. You can check whether you have any disk-quota restrictions by right-clicking the disk you want to save data to, clicking **Properties**, and then clicking the **Quota** tab.

You can also restrict access to a backup file by selecting **Allow only the owner and the Administrator access to the backup data** in the **Backup Job Information** dialog box. If you select this option, only an administrator or the person who created the backup file will be able to restore the files and folders.

**Automated System Recovery**

Automated System Recovery (ASR) is a part of Backup that you can use to recover a system that will not start. With ASR, you can create ASR sets on a regular basis as part of an overall plan for system recovery in case of system failure. You should use ASR as a last resort in system recovery, only after you have exhausted other options such as the startup options Safe Mode and Last Known Good Configuration.

ASR is a recovery option that has two parts: ASR backup and ASR restore. You can access the backup portion through the **Automated System Recovery Preparation Wizard** located in Backup. The **Automated System Recovery Preparation Wizard** creates an ASR set, which is a backup of the System State data, system services, and all disks associated with the operating system components. It also creates a floppy disk, which contains information about the backup, the disk

configurations (including basic and dynamic volumes), and how to restore your system.

You can access the restore part of ASR by pressing **F2** when prompted in the text mode portion of Setup. ASR reads the disk configurations from the floppy disk and restores all of the disk signatures, volumes and partitions on the disks that are required to start your computer (at a minimum). It will attempt to restore all of the disk configurations, but under some circumstances it might not be able to. ASR then installs a simple installation of Windows and automatically starts to restore from backup using the backup ASR set.

**Note**

- ASR does not include data files. You should back up data files separately on a regular basis and restore them after the system is working.
- ASR only supports FAT16 volumes up to 2.1 gigabytes (GB). ASR does not support 4-GB FAT16 partitions that use a cluster size of 64 K. If your system contains 4-GB FAT16 partitions, convert them from FAT16 to NTFS before using ASR.

**Restoring File Security Settings**

Backup preserves permissions, ownership, and audit flags on files restored to NTFS volumes, but not on files restored to FAT volumes. It is not possible to secure data on FAT volumes.

When you restore files to a new computer or hard disk, you do not have to restore security information. The files inherit the permissions of the directory in which they are placed. If the directory has no permissions, the file retains its previous permissions, including ownership.

**Restoring Distributed Services**

In Backup, you can restore distributed services data that is part of the System State data, such as the Active Directory database, using one of three restore methods:

- Primary restore
- Normal (nonauthoritative) restore
- Authoritative restore

To understand how each restore method works, it is important to understand how the Backup utility backs up data for distributed services. When you back up the System State data on a domain controller, you are backing up all Active Directory data that exists on that server (along with other system components such as the SYSVOL directory and the registry). To restore these distributed services to that server, you must restore the System State data. However, the number and configuration of domain controllers in your system will dictate the type of restore method you choose.

For example, if you need to roll back replicated Active Directory changes, but have more than one domain controller in your organization, you will need to perform an authoritative restore to ensure that your restored data gets replicated to all of your servers. However, if you need to restore Active Directory data on a stand-alone domain controller or on the first of several domain controllers, you will need to perform a primary restore. If you need to restore Active Directory data on just one domain controller in a system where Active Directory data is replicated across several domain controllers, you can use a normal restore if your restored data does not have to be replicated to all your servers.

**Primary restore**

Use this type of restore when the server you are trying to restore is the only running server of a replicated data set (for example, the SYSVOL and File Replication service are replicated data sets). Select primary restore only when restoring the first replica set to the network. Do not use a primary restore if one or more replica sets have already been restored. Typically, perform a primary restore only when all the domain controllers in the domain have failed, and you are trying to rebuild the domain from backup.

| Distributed Data | Reason for Using Primary Restore of System State Data |
|---|---|
| Active Directory | Restoring a stand-alone domain controller.<br><br>Restoring the first of several domain controllers. |

| | |
|---|---|
| SYSVOL | Restoring a stand-alone domain controller.<br><br>Restoring the first of several domain controllers |
| Replica sets | Restoring the first replica set. |

**Normal restore**

During a normal restore operation, Backup operates in nonauthoritative restore mode. That is, any data that you restore, including Active Directory objects, will have their original update sequence number. The Active Directory replication system uses this number to detect and propagate Active Directory changes among the servers in your organization. Because of this, any data that is restored nonauthoritatively will appear to the Active Directory replication system as though it is old, which means the data will never get replicated to your other servers. Instead, if newer data is available from your other servers, the Active Directory replication system will use this to update the restored data. To replicate the restored data to the other servers, you must use an authoritative restore.

| Distributed Data | Reason for Using Normal Restore of System State Data |
|---|---|
| Active Directory | Restoring a single domain controller in a replicated environment. |
| SYSVOL | Restoring a single domain controller in a replicated environment. |
| Replica sets | Restoring all but the first replica sets (that is, sets 2 through n, for n replica sets). |

**Authoritative restore**

To authoritatively restore Active Directory data, you need to run the Ntdsutil utility after you have restored the System State data but before you restart the server. The Ntdsutil utility lets you mark Active Directory objects for authoritative restore. When an object is marked for authoritative restore its update sequence number is changed so that it is higher than any other update sequence number in the Active Directory replication system. This will ensure that any replicated or distributed data that you restore is properly replicated or distributed throughout your organization.

For example, if you inadvertently delete or modify objects stored in Active Directory, and those objects are replicated or distributed to other servers, you will need to authoritatively restore those objects so they are replicated or distributed to the other servers. If you do not authoritatively restore the objects, they will never get replicated or distributed to your other servers because they will appear to be older than the objects currently on your other servers. Using the Ntdsutil utility to mark objects for authoritative restore ensures that the data you want to restore gets replicated or distributed throughout your organization. On the other hand, if your system disk has failed or the Active Directory database is corrupted, then you can simply restore the data nonauthoritatively without using the Ntdsutil utility.

You can run the Ntdsutil command-line utility from the command prompt. Help for the Ntdsutil utility is available through the command prompt by typing **ntdsutil /?**.

| Distributed Data | Reason for Using Authoritative Restore of System State Data |
|---|---|
| Active Directory | Rolling back or undoing changes |
| SYSVOL | Resetting data |
| Replica sets | Rolling back or undoing changes |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Volume Shadow Copy Service Technical Reference
- Shadow Copies for Shared Folders Technical Reference

# Backup Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backup Tools and Settings**
**In this section**


- Backup Tools
- Backup Registry Entries
- Related Information


**Backup Tools**
The following tools are associated with Backup.

**Ntbackup.exe: Backup**
**Category**
Windows Server 2003 operating system tool

**Version compatibility**
Runs on computers running the Windows XP or Windows Server 2003 operating systems.

Performs backup operations at a command prompt or through a batch file.

For more information about Ntbackup.exe, see the Windows Server 2003 online Help.

**VSSadmin.exe: Volume Shadow Copy Service Administrative Tool**
**Category**
Windows Server 2003 operating system tool

**Version compatibility**
Runs on computers running the Windows XP or Windows Server 2003 operating systems.

Displays current volume shadow copies and all installed shadow copy writers and providers in the command window. The Windows Server 2003 version of this command contains more parameters than the one that is in Windows XP.

To find more information about VSSadmin.exe, see “Command Line References” in the Tools and Settings Collection.

**Backup Registry Entries**
The following registry entries are associated with Backup.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**Control**
The following registry entry is located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control.

**BackupRestore**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore

**Version**
Windows Server 2003

Contains subkeys that determine which files the system does not back up and which registry subkeys and entries the system does not restore.

**Control**

The following registry entry is located under HKEY_CURRENT_USER\SYSTEM\CurrentControlSet\Control.

**BackupRestore**
**Registry path**
HKEY_CURRENT_USER\SYSTEM\CurrentControlSet\Control\BackupRestore

**Version**
Windows Server 2003

Lists the files that are owned by the current user that should not be backed up.

**Microsoft**
The following registry entry is located under HKEY_CURRENT_USER\Software\Microsoft.

**Ntbackup**
**Registry path**
HKEY_CURRENT_USER\Software\Microsoft\Ntbackup

**Version**
Windows Server 2003

Contains subkeys that store configuration data for Windows Backup.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Volume Shadow Copy Service Technical Reference
- Shadow Copies for Shared Folders Technical Reference

# Shadow Copies for Shared Folders Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shadow Copies for Shared Folders Technical Reference**
Shadow Copies for Shared Folders is a new file-storage technology in the Microsoft Windows Server 2003 operating systems. Shadow Copies for Shared Folders provides point-in-time copies of files that are located on a shared network resource, such as a file server. With Shadow Copies for Shared Folders, users can quickly recover deleted or changed files that are stored on the network without administrator assistance, which can increase productivity and reduce administrative costs.

**In this subject**


- What Is Shadow Copies for Shared Folders?
- How Shadow Copies for Shared Folders Work
- Shadow Copies for Shared Folders Tools and Settings

# What Is Shadow Copies for Shared Folders?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Shadow Copies for Shared Folders?**

Shadow Copies for Shared Folders is a new file-storage technology in the Microsoft Windows Server 2003 operating systems. Shadow Copies for Shared Folders uses the Volume Shadow Copy service to provide point-in-time copies of files that are located on a shared network resource, such as a file server. With the Previous Versions client for Shadow Copies for Shared Folders, users can view shared files and folders as they existed at points of time in the past, without administrator assistance. Accessing previous versions of files, or shadow copies, is useful because users can:

- **Recover files that were accidentally deleted**. If users accidentally delete a file, they can open a previous version and copy it to a safe location.
- **Recover from accidentally overwriting a file**. If users accidentally overwrite a file, they can recover a previous version of the file.
- **Compare different versions of a file while working**. Users can use previous versions when they want to check what has changed between two versions of a file.

A common scenario for recovering lost or corrupted files occurs when an end user submits an urgent request to the IT help desk to find an archived version of a file. If the organization has an archiving system in place, this request usually requires a costly and time-intensive search of archived media, which in many instances is a tape backup. This situation creates several problems:

- Potential loss of business if the lost document is time-sensitive or labor-intensive to replace
- Decreased productivity for the end user
- Increased cost to help-desk and IT-support services

Because end users can access previous versions of files by themselves, using Shadow Copies of Shared Files for routine file recovery scenarios can help to:

- Reduce demand on busy administrators, for example, by reducing requests to restore files from tape.
- Reduce the cost of recovering single or multiple files.

### Note

- Shadow copies do not eliminate the need to perform regular backups or provide protection from media failure. In addition, shadow copies are not permanent. As new shadow copies are taken, old shadow copies are deleted when the storage limit is reached, when you run out of disk space, or when the number of shadow copies reaches 64, whichever occurs first. Therefore, shadow copies might not be present for as long as end users expect them to be.

**Related Information**

The following resources contain additional information that is relevant to this section.

- How Shadow Copies for Shared Folders Work
- Backup Technical Reference
- Volume Shadow Copy Service Technical Reference

# How Shadow Copies for Shared Folders Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Shadow Copies for Shared Folders Works**

**In this section**


- Shadow Copies for Shared Folders Architecture
- Source Volume
- Storage Volume
- Previous Versions Clients
- File Permissions
- Related Information


Shadow Copies for Shared Folders is a new file-storage technology in Microsoft Windows Server 2003. Shadow Copies for Shared Folders provides point-in-time copies of files that are located on shared network resources, such as a file server. With the Previous Versions client for Shadow Copies for Shared Folders, users can view shared files and folders as they existed at points of time in the past. Accessing previous versions of your files, or shadow copies, is useful because users can:

- **Recover files that were accidentally deleted**. If users accidentally delete a file, they can open a previous version and copy it to a safe location.
- **Recover from accidentally overwriting a file**. If users accidentally overwrite a file, they can recover a previous version of the file.
- **Compare different versions of a file while working**. Users can use previous versions when they want to check what has changed between two versions of a file.

**Shadow Copies for Shared Folders Architecture**

There are three main components that make up Shadow Copies for Shared Folders: the source volume, the storage volume, and the Previous Versions Client.

**Shadow Copies for Shared Folders Components**

| Component | Description |
|---|---|
| Source volume | The volume that is being copied. This volume is located on a server running Windows Server 2003 and has Shadow Copies for Shared Folders enabled. This volume contains several folders that are shared on the network. This is typically a volume on a file server. |
| Storage volume | The volume where shadow copies are stored. |
| Previous Versions Client | The interface that is used to view shadow copies. The client is available for Windows Server 2003, Windows XP, and Windows 2000 operating systems. For more information, see "Previous Versions Clients." |

The following diagram shows the main components of Shadow Copies for Shared Folders and how they interact.

**Shadow Copies for Shared Folders Architecture**




**Source Volume**
The source volume is the volume that is being copied, typically a volume on a file server. Shadow Copies for Shared Folders is enabled on a per-volume basis. That is, you can only make shadow copies of an entire volume. You cannot select specific shared folders and files on a volume to be copied.

The volume must reside on a server running Windows Server 2003 and must be formatted using the NTFS file system. Shadow Copies for Shared Folders is built upon the Volume Shadow Copy service technology, which provides a way to make copies of open files. For more information, see "Volume Shadow Copy Service Technical Reference."

The mounted drive will not be included when shadow copies are taken. You should enable shadow copies only on volumes without mount points or only when you do not want the shared resources on the mounted volume to be copied.

You can access the server portion of Shadow Copies for Shared Folders through the **Shadow Copies** tab of the **Local Disk Properties** dialog box.

### Note

- When you enable Shadow Copies for Shared Folders on a volume, a default scheduled task is also created. The default schedule for copies is twice a day at 7:00 A.M. and 12:00 noon, Monday through Friday.

**Storage Volume**
The storage volume is where shadow copies are stored. Shadow copies can be stored on any NTFS-formatted volume that is available to the server, including mounted volumes. Because of the high I\O involved in creating the copies, we recommend that you store the shadow copies on a volume on a separate disk from the disk that contains the source volume.

If you use Shadow Copies for Shared Folders on mounted volumes in a cluster, do not place the storage volume on a volume that is mounted to the source volume. In addition, do not mount the source volume to the storage volume. Otherwise, the cluster dependency between the Physical Disk resources of the mount point volume and the volume it is mounted to will interfere with the cluster dependency that is introduced by VSS between the source and storage volumes.

Shadow Copies for Shared Folders works by making a block-level copy of any changes that have occurred to a file since the last shadow copy was created. The file changes are copied and stored as blocks, or units of data. Generally, the entire file is not copied. Only the previous values of the changed blocks are copied to the storage area. As a result, previous versions of files do not usually take up as much disk space as the current file.

However, the amount of disk space that is used for changes can vary, depending on the application that changed the file. For example, some applications rewrite the entire file when a change is made, but other applications add changes to the existing file. If the entire file is rewritten to disk when a change is made, then the shadow copy contains the entire file.

The minimum amount of storage space that you can specify to be used for storing shadow copies on the storage volume is approximately 100 megabytes (MB) for computers running Windows XP with Service Pack 2 or Windows Server 2003. It is approximately 300 MB for Windows Server 2003 with hotfix 833167 applied. The default storage size is 10% of the source volume (the volume being copied). When the storage limit is reached, older versions of the shadow copies will be deleted and cannot be restored. There is also a limit of 64 shadow copies per volume that can be stored. When this limit is reached, the oldest shadow copy will be deleted and cannot be retrieved.

### Note

- Shadow copies are read-only. You cannot edit the contents of a shadow copy.

**Previous Versions Clients**

There are two clients for Shadow Copies for Shared Folders:

- **Previous Versions Client for Windows XP and Windows Server 2003.** This client can be installed on Windows XP from the Windows Server 2003 installation media or it can be installed using the .msi package that is available at the Microsoft Download Center. Previous Versions Client is installed by default on Windows Server 2003 and is included when you install Windows XP Service Pack 2.

- **Previous Versions Client for Windows 2000.** This client is only available from the .msi package available at the Microsoft Download Center. In order for Windows 2000 clients to view shadow copies, a registry key must be set on the computer running Windows Server 2003 that enables access by Windows 2000 clients. The .msi package that contains the Windows 2000 client also contains an installation package for the server that will set this registry key.

You can access the client view of shadow copies through the **Previous Versions** tab of the **Properties** dialog box of the shared file or folder.

You cannot view previous versions of files on a local shared resource. You must have a working network connection and be viewing the volume as a shared network resource. On a computer running Windows Server 2003, in order to view shadow copies of a local shared folder, you must create a connection to the shared resource on your local computer as if it were in another location on the network. For example: To view previous versions of files in the shared folder ***\Documents***on the local computer ***\\Server1***, you need to create a network connection to ***\\Server1\Documents*** and access the previous versions as if they were on a remote computer.

**File Permissions**
When you restore a file to a previous version, the file permissions will not be changed. File permissions remain the same as they were before the file was restored. When you recover a file that was accidentally deleted, the file permissions are set to the default permissions for the directory the file is in. This directory might have different permissions than the file.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Volume Shadow Copy Service Technical Reference

- Backup Technical Reference

# Shadow Copies for Shared Folders Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shadow Copies for Shared Folders Tools and Settings**

**In this section**


- Shadow Copies for Shared Folders Tools
- Shadow Copies for Shared Folders Registry Entries
- Shadow Copies for Shared Folders WMI Classes
- Network Ports Used by Shadow Copies for Shared Folders
- Related Information


**Shadow Copies for Shared Folders Tools**

The following tools are associated with Shadow Copies for Shared Folders.

**VSSadmin.exe: Volume Shadow Copy Service Administrative Tool**

**Category**

Windows Server 2003 operating system tool

**Version compatibility**

Runs on computers running the Windows XP or Windows Server 2003 operating systems.

Displays current volume shadow copies and all installed shadow copy writers and providers in the command window. The Windows Server 2003 version of this command contains more parameters than the one that is in Windows XP.

To find more information about VSSadmin.exe, see "Command Line References" in the Tools and Settings Collection.

**VolPerf.exe: Shadow Copy Performance Counters**

**Category**

Windows Server 2003 Resource Kit tool

**Version compatibility**

Runs on computers running any of the Windows Server 2003 operating systems.

Enables administrators to use System Monitor to monitor their shadow copies.

To find more information about Volperf.exe, see "Resource Kit Tools" in the Tools and Settings Collection.

**VolRest.exe: Shadow Copies for Shared Folders Restore Tool**

**Category**

Windows Server 2003 Resource Kit tool

**Version compatibility**

Runs on computers running the Windows XP or Windows Server 2003 operating systems.

Searches for and restores previous file versions from shared folders that are protected by the Shadow Copies for Shared Folders feature. With VolRest, administrators can browse or restore previous versions of one or multiple files that are stored on a server where Shadow Copies for Shared Folders has been enabled.

To find more information about Volrest.exe, see "Resource Kit Tools" in the Tools and Settings Collection.

**Shadow Copies for Shared Folders Registry Entries**

The following registry entries are associated with Shadow Copies for Shared Folders.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This

can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**Parameters**
The following registry entry is located under \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters.

**DisableDownLevelTimewarp**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\DisableDownLevelTimewarp

**Version**
Runs on computers running the Windows Server 2003 operating system.

This registry entry enables or disables support for Shadow Copies for Shared Folders on clients running the Windows 2000 operating system.

Set the value of the DisableDownLevelTimewarp registry key to 0 (zero) to enable Shadow Copies for Shared Folders from clients running Windows 2000. Set the value of the DisableDownLevelTimewarp registry key to 1 (one), the default, to disable Shadow Copies for Shared Folders from clients running Windows 2000.

**VOLSNAP**
The following registry entry is located under \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VOLSNAP.

**MinDiffAreaFileSize**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VOLSNAP\MinDiffAreaFileSize

**Version**
Windows Server 2003

Controls the initial size of the shadow copy storage area for each shadow copy. The MinDiffAreaFileSize registry key uses the REG_DWORD data type.

The MinDiffAreaFileSize registry key specifies the minimum size of the shadow copy storage area. The default size of the shadow copy storage area is 300 MB, and the maximum size is 3 gigabytes (GB). For an accurate setting, specify a value that is a multiple of 300 MB. Otherwise, the next multiple of 300 MB will be used. A value of 300 is equal to 300 MB, and a value of 3,000 is equal to 3 GB.

Shadow copies might not function correctly if the value of the MinDiffAreaFileSize registry key is larger than the maximum size of the shadow copy storage area. Before you change this registry entry, run the **vssadmin list shadowstorage** command. Make sure that the shadow copy storage area is larger than 300 MB or that the shadow copy storage area is larger than the value that is set in the registry. Also, make sure that you have at least 300 MB of free space on the shadow copy storage volume.

**Shadow Copies for Shared Folders WMI Classes**
The following table lists and describes the WMI classes that are associated with Shadow Copies for Shared Folders.

### WMI Classes Associated with Shadow Copies for Shared Folders

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_ShadowBy | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowContext | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowCopy | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowDiffVolumeSupport | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowFor | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowOn | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowStorage | \root\cimv2 | Windows Server 2003 family |

| | | |
|---|---|---|
| Win32_ShadowVolumeSupport | \root\cimv2 | Windows Server 2003 family |
| Win32_Volume | \root\cimv2 | Windows Server 2003 family |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by Shadow Copies for Shared Folders**

Shadow Copies for Shared Folders uses the following network ports.

**Network Ports Used by Shadow Copies for Shared Folders**

| Service Name | Description | Port Used |
|---|---|---|
| Server Message Block (SMB) | Required for Shadow Copies for Shared Folders | TCP 138, 139, 445 |
| DCOM/Remote Procedure Call (RPC) | Required for remote management using WMI. | TCP 135, dynamic |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Backup Technical Reference
- Volume Shadow Copy Service Technical Reference

# Volume Shadow Copy Service Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Volume Shadow Copy Service Technical Reference**
The Volume Shadow Copy Service provides the backup infrastructure for the Microsoft Windows XP and Microsoft Windows Server 2003 operating systems, as well as a mechanism for creating consistent point-in-time copies of data known as shadow copies.

The Volume Shadow Copy Service can produce consistent shadow copies by coordinating with business applications, file-system services, backup applications, fast-recovery solutions, and storage hardware. Several features in the Windows Server 2003 operating systems use the Volume Shadow Copy Service, including Shadow Copies for Shared Folders and Backup.

Shadow copies can be used for a number of purposes such as:

- Creating consistent backups of open files and applications
- Creating shadow copies for shared folders
- Quickly recovering and restoring files and data
- Creating transportable shadow copies using a hardware provider — for backup, testing and data mining scenarios

**In this subject**


- What Is Volume Shadow Copy Service?
- How Volume Shadow Copy Service Works
- Volume Shadow Copy Service Tools and Settings

# What Is Volume Shadow Copy Service?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Volume Shadow Copy Service?**

**In this section**


- Volume Shadow Copy Service Scenarios
- Related Information


The Volume Shadow Copy Service provides the backup infrastructure for the Microsoft Windows XP and Microsoft Windows Server 2003 operating systems, as well as a mechanism for creating consistent point-in-time copies of data known as shadow copies.

Previous to the Volume Shadow Copy Service and its standard set of extensible application programming interfaces (APIs), there was no standard way to produce "clean" (uncorrupted) snapshots of a volume. Snapshots often contained corruptions due to "torn writes" that required the use of utilities such as Chkdsk.exe to repair. Torn writes occur when an unplanned event (such as a power failure) prevents the system from completely writing a block of data to disk. The Volume Shadow Copy Service APIs prevent torn writes by enabling applications to flush partially committed data from memory.

The Volume Shadow Copy Service has native support for creating consistent shadow copies across multiple volumes, regardless of the snapshot technology or application. The Volume Shadow Copy Service can produce consistent shadow copies by coordinating with business applications, file-system services, backup applications, fast recovery solutions, and storage hardware. Several features in the Windows Server 2003 operating systems use the Volume Shadow Copy Service, including Shadow Copies for Shared Folders and Backup.

**Volume Shadow Copy Service Scenarios**

Shadow copies can be used for a number of purposes such as:

- Creating consistent backups of open files and applications
- Creating shadow copies for shared folders
- Creating transportable shadow copies using a hardware provider — for backup, testing and data mining scenarios
- Quickly recovering and restoring files and data

The Volume Shadow Copy Service is commonly used in the following scenarios.

**Consistent backups of open files and applications**

Applications that are running often keep their files open continuously. For backup, this can present a problem because this prevents backup applications from accessing and copying these files to backup media. Even if an application does not have its files open, it is possible—because of the finite time needed to open, back up, and close a file—that files copied to the backup media might not all reflect the same application state at the same point in time.

Additionally, backing up servers that are running critical applications such as databases or messaging services presents a unique challenge. These applications run in a volatile state as a result of extensive optimizations that deal with huge flows of transactions and messages. Because these applications keep their data in a constant flux between memory and disk, it is difficult to pinpoint the data that needs to be archived. The most straightforward solution is to interrupt the application during backup, which puts the data into a stable state, but might result in unacceptable amounts of downtime, particularly if the applications are large.

For both problems, the Volume Shadow Copy Service provides a solution by enabling a snapshot of the data at a given point in time, while minimizing the interruption to applications. This can involve cooperation from the applications, so that they can notify the operating system that they are briefly quiescing their data for archival. During quiescence, applications make data on the disk consistent. For example, an application might flush its buffers to disk or write out in-memory data to disk.

The Volume Shadow Copy Service provides interfaces for storage-management applications (called requestors), business

applications (called writers) and storage devices (with components called providers) to ensure that volume shadow copies are in a consistent state.

**Shadow Copies for Shared Folders**
Shadow Copies for Shared Folders uses the Volume Shadow Copy Service to provide point-in-time copies of files that are located on a shared network resource, such as a file server. With Shadow Copies for Shared Folders, users can quickly recover deleted or changed files that are stored on the network without administrator assistance, which can increase productivity and reduce administrative costs.

For more information about Shadow Copies for Shared Folders, see "Shadow Copies for Shared Folders Technical Reference."

**Shadow Copy Transport**
With a hardware provider that is designed for use with the Volume Shadow Copy Service, you can create transportable shadow copies that can be imported onto servers within the same subsystem (for example, a storage area network (SAN)). While multiple servers can access the same storage array, they do not share access to the same storage unit or LUN. For servers to share the same data, the traditional solution is to first copy the data on one server and then restore it to a second server. This process can be quite lengthy.

With the Volume Shadow Copy Service and a storage array with a hardware provider designed for use with the Volume Shadow Copy Service, it is possible to create a shadow copy of the source data volume on one server and then to import the shadow copy onto another server (or back to the same server). This process is accomplished in a few minutes, regardless of the size of the data. The transport process is accomplished through a series of steps using a shadow copy requestor (storage-management application) that supports transportable shadow copies. The following is a generalized version of those steps:

1. Create a transportable shadow copy of the source data on a server.

2. Import the shadow copy to a server connected to the SAN (you can import to a different server or the same server).

3. The data is now ready to be used.

Note that shadow copies are read-only. If you want to convert a shadow copy to a read/write LUN, you can use a storage-management application (including some requestors) that is compatible with the Virtual Disk Service, in addition to the Volume Shadow Copy Service. Using this application, you can remove the shadow copy from Volume Shadow Copy Service management and convert it to a read/write LUN.

Volume Shadow Copy Service transport is an advanced solution on computers running Windows Server 2003, Enterprise Edition or Windows Server?2003, Datacenter Edition that works only if there is a hardware provider on the storage array. Shadow copy transport can be used for a number of purposes, including tape backups, data mining, and testing.

**Transport for Tape Backups**
Backing up data to tape is perhaps the most I/O intensive operation that an application can perform. On a production local area network (LAN), this is problematic because the data must travel across two systems: from the storage on the server across to the network to the backup system, and from the backup system to the tape device. This high traffic process can result in considerable congestion on the production network, degrading application performance.

One way to alleviate this problem is to transport shadow copies from the production server onto a backup server, where the shadow copy volumes can be backed up to tape. This option removes backup traffic from the production server. While some backup applications might be designed with the hardware provider software that enables transport, others are not. The administrator should determine whether or not this functionality is included in the backup application.

**Transport for Data Mining**
The data in use by a particular production server is often useful to different groups or departments within an organization. Rather than add additional load to the production server, a shadow copy of the data can be made available through transport to another server. The shadow copy can then be processed for different purposes, without any performance impact on the original server.

**Transport for Testing**
Shadow copies can be particularly useful for software developers, who often need data against which to test new versions of an application. Developers can make a shadow copy of the source data and transport it to a different server in order to test an application. If an application corrupts the shadow copy of the data, there is no impact on the original data.

**Fast Recovery**
Fast recovery uses shadow copies as a complement to traditional tape-based recovery solutions. Using point-in-time shadow copies with Active Directory and Exchange Server configurations, for example, helps you to quickly recover from a number of

specific system problems, including:

- Hardware failure
- Storage failure
- Software corruption

Because this mirroring (or "cloning") process is fast and nondisruptive to system performance, shadow copies can be made more frequently than tape backups. Shadow copies that are kept locally on a storage area network can be quickly accessed. With the appropriate hardware provider, they can be transported to a backup server, backed up to tape, and sent to off-site storage for archiving.

With the Volume Shadow Copy Service and the Virtual Disk Service, Windows Server 2003 contains the functionality to enable fast data restores that only take a few minutes — compared to the hours it can take with tape backups.

As with tape-based recovery of data, recovery using the Volume Shadow Copy Service and Virtual Disk Service requires planning well ahead of the potential server failure. The difference is that the data recovery from tape is far slower and more labor intensive than when using Windows Server 2003 technology and storage area networks.

**Mechanisms Underlying Fast Recovery**
The following sections briefly outline the conceptual aspects of each of these fast recovery solution components.

**Shadow Copy Creation**
The process of creating shadow copies begins when the backup application (requestor) contacts the Volume Shadow Copy Service to request a shadow copy. The Volume Shadow Copy Service, acting as a coordinator, notifies each of the writers to prepare for shadow copy creation. Once the writers indicate they are prepared, the requestor instructs the Volume Shadow Copy Service to create the shadow copy. The Volume Shadow Copy Service informs the writers to quiesce their stores and briefly freeze (stop) writes to disk for the few seconds it takes the provider to create the point-in-time shadow copies of the volumes.

Once a shadow copy of one or more volumes is made, the Volume Shadow Copy Service informs the writers they can thaw (resume writing to disk) and the original data continues to be used in production. At this point, the read-only shadow copy resides in the storage array until a need arises.

**Fast Recovery Process**
After a failure event, a management application using both the Volume Shadow Copy Service and the Virtual Disk Service can remove the failed LUNs from the system. The shadow copies of the LUNs can be converted from read-only to read/write. The LUNs can then be assigned to another server.

A generalized example of shadow copy creation and transport between two servers is shown in the following diagram. A single point-in-time shadow copy is made of the three LUNs assigned to Server 1. This shadow copy set can be accessed by Server 2 through the process of first unmasking, then mounting the LUNs. The complete process is called "shadow copy transport."

**Shadow Copy Creation and Transport Between Two Servers**




Note that this is a general example of transporting shadow copies of LUNs between two servers.

**Related Information**

The following resources contain additional information that is relevant to this section.

- How Volume Shadow Copy Service Works
- Backup Technical Reference
- Shadow Copies for Shared Folders Technical Reference

# How Volume Shadow Copy Service Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Volume Shadow Copy Service Works**

**In this section**


- Methods for Creating Shadow Copies
- Volume Shadow Copy Service Architecture
- Components Used for Creating Shadow Copies
- How Shadow Copies Are Created
- Related Information


The Volume Shadow Copy Service provides the backup infrastructure for the Microsoft Windows XP and Microsoft Windows Server 2003 operating systems, as well as a mechanism for creating consistent point-in-time copies of data known as shadow copies.

The Volume Shadow Copy Service can produce consistent shadow copies by coordinating with business applications, file-system services, backup applications, fast-recovery solutions, and storage hardware. Several features in the Windows Server 2003 operating systems use the Volume Shadow Copy Service, including Shadow Copies for Shared Folders and Backup.

**Methods for Creating Shadow Copies**

There are two methods for creating shadow copies: making either a complete copy (a full copy or clone) or copying only the changes to the volume (a differential copy or copy-on-write). Each method results in two data images — the original volume and the shadow copy volume. The functional difference between the two is that the original volume maintains full read/write capabilities, whereas the shadow copy volume is read-only. This read-only status ensures that the shadow copy volume remains a point-in-time copy until its status is changed by the administrator for a specific purpose.

**Clone (Full Copy/Split Mirror):**

A clone is a full copy of the original data on a volume. You can create a clone through either software or hardware mirroring. Clones remain synchronized until the mirror connection is broken for the shadow copy. From this point forward, the source data and the shadow copy volume are independent. The original volume continues to take application changes, while the shadow copy volume remains an exact read-only copy of the original data at the time of the break.

Hardware vendors offer different hardware-based implementations (sometimes called split mirrors, snapshot mirrors, or clones) for creating identical images of volumes that can be used for online backup, application development, and testing.

**Copy-on-Write (Differential Copy)**

The copy-on-write method creates shadow copies that are differential rather than full copies of the original data. Like the clone method of creating shadow copies, the copy-on-write method can produce shadow copies using either software or hardware solutions. This method makes a copy of the original data before it is overwritten with new changes, as shown in the following table. When a change to the original volume occurs, but before it is written to disk, the block about to be modified is read and then written to a "differences area", which preserves a copy of the data block before it is overwritten with the change. Using the blocks in the differences area and unchanged blocks in the original volume, a shadow copy can be logically constructed that represents the shadow copy at the point in time in which it was created.

**The Copy-on-Write Method of Creating Shadow Copies**

| Time | Source Data | Contents | Shadow Copy | Contents |
|---|---|---|---|---|
| T0 | Original data | 1 2 3 4 5 | No copy | — |
| T1 | Original data overwritten | 1 2 3 4 5 | Differences and index stored on shadow copy | 3 |

The advantage of the copy-on-write method is that it creates shadow copies very rapidly because it is only writing the changes

to disk. The disadvantage is that in order to fully restore the data, the original data must still be available. Without the original data, the shadow copy is incomplete and cannot be used. Another disadvantage is that the performance of copy-on-write implementations can affect the performance of the original volume.

**Volume Shadow Copy Service Architecture**

The following diagram and table describe how the Volume Shadow Copy Service coordinates with various components to create a shadow copy of a volume.

**Volume Shadow Copy Service Architecture Diagram**




**Volume Shadow Copy Service Components**

| Component | Description |
|---|---|
| Volume Shadow Copy Service | A service that coordinates various components to create consistent shadow copies of one or more volumes. |
| Requestor | An application that requests that a volume shadow copy be taken. A backup application is an example. |
| Writer | A component of an application that stores persistent information on one or more volumes that participate in shadow copy synchronization. Typically, this is a database application like SQL Server or Exchange Server, or a system service like Active Directory. |
| Provider | A component that creates and maintains the shadow copies. Examples are the system provider included with the operating system and the hardware providers included with storage arrays. |
| Source volume | The volume that contains the data to be shadow copied. |
| Storage volume | The volume that holds the shadow copy storage files for the system copy-on-write software provider. |

**Components Used for Creating Shadow Copies**

This section outlines the requestors, writers, and providers that are necessary for creating consistent shadow copies. The Volume Shadow Copy Service provides coordination among these components.

**Requestors: Initiating Shadow Copy Creation**

The Volume Shadow Copy Service is invoked by the requestor, which is typically a backup application that creates shadow copy volumes to back up data while the source volume continues to operate in production. Requestors can also be management applications that manage shadow copy creation and usage, or fast recovery solutions which are specific products that reduce service level agreement (SLA) times for specific applications.

The requestor also communicates with the writers to gather information about what should be backed up and how it should be backed up.

**Writers: Preventing Data Inconsistencies**

Writers are software that is included in applications and services that help provide consistent shadow copies. Writers serve two main purposes:

- When applications and services are running, the writer responds to signals provided by the Volume Shadow Copy Service interface to allow applications to prepare and quiesce their data stores for shadow copy creation and to ensure that no writes occur on the volume while the shadow copy is being created. (During quiescence, applications make data on the disk consistent. For example, an application might flush its buffers to disk or write out in-memory data to disk.)

- The writer also provides information about the application name, icons, files to include or exclude, and a strategy to

restore the files.

If an application or service is not running, and the writer cannot respond to the requesting backup application, it is assumed that all data on the volume is consistent, the databases are closed, and no additional effort is required to perform the backup.

A writer is associated with one or more components. A component is a group of files that must be backed up as a unit. For example, a database and set of log files. For a backup to be successful, all files associated with the component must be backed up. Writers also provide information about restoring the data on a component-by-components basis.

If an application has no writer, the shadow copy will still occur and all of the data, in whatever form it is in at the time of the copy, will be included in the shadow copy. This means that there might be inconsistent data that is now contained in the shadow copy. This data inconsistency is caused by incomplete writes, data buffered in the application that is not written, or open files that are in the middle of a write operation. Even though the file system flushes all buffers prior to creating a shadow copy, the data on the disk can only be guaranteed to be crash-consistent if the application has completed all transactions and has written all of the data to the disk. (Data on disk is "crash-consistent" if it is the same as it would be after a system failure or power outage.)

Under this design, the responsibility for data consistency has been shifted from the requestor application to the production application. The advantage of this approach is that application developers — those most knowledgeable about their applications — can ensure, through development of their own writers, the maximum effectiveness of the shadow copy creation process.

**Applications that are not shadow copy–enabled**
If the computer has no applications that are enabled for the Volume Shadow Copy Service, the data in a shadow copy is considered to be in a "crash consistent" state. All files that were open will still exist, but are not guaranteed to be free of incomplete I/O operations or data corruption.

While the crash-consistent state does not fully deal with all the issues associated with defining a stable backup set, it has several advantages over the backup set that conventional backup operations would have to use.

- For example, a shadow copy of a volume, even in crash-consistent state, still contains all files. A backup set created without a shadow copy would not contain all files open exclusively at the time of the backup. Files held open at the time of the backup operation are excluded from the backup.

The shadow copy of the volume is created at a single point in time and is synchronized across the whole volume set. In order to avoid inconsistencies, shadow copies are not taken file by file.

**Providers: Creating Shadow Copies**
In the context of shadow copy creation, a provider is a component that serves as an interface to the point-in-time imaging capabilities — either on the storage array (hardware-based) or in the operating system (software based). Providers manage running volumes and create shadow copies of them on demand. In response to a request from a requestor, the Volume Shadow Copy Service signals applications that a shadow copy is about to be created, then signals the provider to create and maintain that shadow copy until it is no longer needed. Hardware providers are implemented as a user-mode component which communicates with the hardware that will expose the shadow copy data. Windows Server 2003 includes a system software provider with shadow copy functionality.

Alternatively, other hardware and software vendors can develop their own hardware or software providers to provide point-in-time imaging functionality. Windows Server 2003 supports multiple hardware and software providers that can be used in combination to solve many different IT operational scenarios.

The Volume Shadow Copy Service uses the following hierarchy to select the provider to use during shadow copy creation:

1. Hardware provider

2. Software provider

3. System software provider.

However, if a specific IT operational problem requires it, the requestor can override this hierarchy.

**Hardware-based Providers**
Hardware-based shadow copy providers act as an interface between the Volume Shadow Copy Service and the hardware level

by working in conjunction with a hardware storage adapter or controller. The work of creating the shadow copy is performed by a host adapter, storage appliance, or RAID controller outside of the operating system.

Hardware providers always take the shadow copy of an entire LUN, but the Volume Shadow Copy Service will only expose the shadow copy of the volume or volumes that were requested.

While a hardware-based shadow copy provider makes use of the Volume Shadow Copy Service functionality that defines the point in time, allows data synchronization, manages the shadow copy, and provides a common interface with backup applications, the Volume Shadow Copy Service does not specify the underlying mechanism by which the hardware-based provider produces and maintains shadow copies.

**Software-based Providers**
Software-based shadow copy providers typically intercept and process I/O requests in a software layer between the file system and the volume manager software.

These providers are implemented as a user-mode DLL component and at least one kernel-mode device driver, typically a storage filter driver. The work of creating these shadow copies is done in software.

A software-based shadow copy provider must maintain a "point-in-time" view of a volume by having access to a data set that can be used to recreate volume status prior to the shadow copy. An example of this is the copy-on-write technique of the system provider. However, the Volume Shadow Copy Service places no restrictions on what technique software-based providers use to create and maintain shadow copies.

A software provider will be applicable to a wider range of storage platforms than a hardware-based provider and should be able to work with basic disks or logical volumes equally well.

This implementation sacrifices the performance that might be available by implementing shadow copies in hardware and does not make use of any vendor-specific features.

**System Provider**
One shadow copy provider, the system provider, is supplied as a default part of the Windows Server 2003 operating system. While a default provider is supplied as part of Windows, other vendors are free to supply their own implementations that are optimized for their own storage hardware and software applications.

To maintain the "point in time" view of a volume contained in a shadow copy, the system provider uses a copy-on-write technique. Copies of the blocks on volume that have been modified since the beginning of the shadow copy creation are stored in a shadow copy storage area.

The system provider can expose the production volume, which can be written to and read from normally. When the shadow copy is needed, it logically applies the differences to data on the production volume to expose the complete shadow copy.

For the system provider, the shadow copy storage area must be on an NTFS volume. The volume to be shadow copied does not need to be an NTFS volume, but at least one volume mounted on the system must be an NTFS volume.

**How Shadow Copies Are Created**
The various roles of the requestor, writer and provider are put into context in this section, which lists the steps that need to be taken to create a shadow copy. Overall coordination of the requestor, writer, and provider is controlled by the Volume Shadow Copy Service, as shown in the following diagram.

**Shadow Copy Creation Process**




1. The requestor asks the Volume Shadow Copy Service to enumerate the writers, gather the writer metadata, and prepare

for shadow copy creation.

2. The writer creates an XML description of the backup components to the Volume Shadow Copy Service, and defines the restore method. The Volume Shadow Copy Service notifies the application-specific writer to prepare its data for making a shadow copy.

3. The writer prepares the data in whatever way is appropriate, such as completing all open transactions, rolling transaction logs, and flushing caches. When the data is prepared for shadow copy creation, the writer notifies the Volume Shadow Copy Service.

4. The Volume Shadow Copy Service initiates the "commit" shadow copy phase.

5. The Volume Shadow Copy Service tells the writers to quiesce their data and temporarily freeze requestor (application) I/O write requests (I/O read requests are still possible) for the several seconds required to create the shadow copy of the volume or volumes. The application freeze is not allowed to take longer than 60 seconds. The Volume Shadow Copy Service flushes the file system buffer and then freezes the file system, which ensures that file system metadata is written and that the data is written in a consistent order.

6. The Volume Shadow Copy Service tells the provider to create the shadow copy (a maximum of 10 seconds).

7. The Volume Shadow Copy Service thaws the file system. After the shadow copy is created, the Volume Shadow Copy Service releases the writers from their temporary inactive phase and all queued write I/Os are completed.

8. The Volume Shadow Copy Service queries the writers to confirm that write I/Os were successfully held during shadow copy creation.

9. If the writes were not successfully held (meaning that the shadow copy data is potentially inconsistent), the shadow copy is deleted and the requestor is notified.

10. The requestor can retry the process (go back to step 1) or notify the administrator to retry at a later time.

11. If the copy is successful, the Volume Shadow Copy Service gives the location information for the shadow copy back to the requestor.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Shadow Copies for Shared Folders Technical Reference

- Backup Technical Reference

# Volume Shadow Copy Service Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Volume Shadow Copy Service Tools and Settings**

**In this section**


- Volume Shadow Copy Service Tools
- Volume Shadow Copy Service Registry Entries
- Volume Shadow Copy Service WMI Classes
- Related Information


**Volume Shadow Copy Service Tools**

The following tools are associated with the Volume Shadow Copy Service.

**DiskRaid.exe: RAID Configuration Tool**

**Category**

Windows Server 2003 Resource Kit tool

**Version compatibility**

Runs on computers running the Windows Server 2003 operating systems.

Enables configuration and management of redundant array of independent (or inexpensive) disks (RAID) storage subsystems.

To find more information about DiskRaid.exe, see "Resource Kit Tools" in the Tools and Settings Collection.

**VSSadmin.exe: Volume Shadow Copy Service Administrative Tool**

**Category**

Windows Server 2003 operating system tool

**Version compatibility**

Runs on computers running the Windows XP or Windows Server 2003 operating systems.

Displays current volume shadow copies and all installed shadow copy writers and providers in the command window. The Windows Server 2003 version of this command contains more parameters than the one that is in Windows XP.

To find more information about VSSadmin.exe, see "Command Line References" in the Tools and Settings Collection.

**VolPerf.exe: Shadow Copy Performance Counters**

**Category**

Windows Server 2003 Resource Kit tool

**Version compatibility**

Runs on computers running any of the Windows Server 2003 operating systems.

Enables administrators to use System Monitor to monitor their shadow copies.

To find more information about Volperf.exe, see "Resource Kit Tools" in the Tools and Settings Collection.

**VolRest.exe: Shadow Copies for Shared Folders Restore Tool**

**Category**

Windows Server 2003 Resource Kit tool

**Version compatibility**

Runs on computers running the Windows XP or Windows Server 2003 operating systems.

Searches for and restores previous file versions from shared folders that are protected by the Shadow Copies for Shared Folders feature. With VolRest, administrators can browse or restore previous versions of one or multiple files that are stored on a server where Shadow Copies for Shared Folders has been enabled.

To find more information about Volrest.exe, see “Resource Kit Tools” in the Tools and Settings Collection.

**Vrfydsk.exe: Verify Disk**
**Category**
Windows Server 2003 Resource Kit tool

**Version compatibility**
Runs on computers running Windows Server 2003 operating systems.

Checks the consistency and integrity of the file system (FAT or NTFS) on a specified volume. It creates and displays a status report for the volume based on the file system and also lists the errors on that volume, if any. However, if any errors are discovered, VrfyDsk cannot correct them. To correct these errors, run Chkdsk in fix mode.

To find more information about Vrfydsk.exe, see “Resource Kit Tools” in the Tools and Settings Collection.

**Volume Shadow Copy Service Registry Entries**
The following registry entries are associated with Volume Shadow Copy Service.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**VSS**
The following registry entry is located under \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VSS.

**VssAccessControl**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VSS\VssAccessControl

**Version**
Windows Server 2003

Enables or prevents a writer from using a specific user account.

In order for any writer to use the VSS infrastructure, the writer must run under an account that is a member of the local Administrators or Backup Operators group on the local computer. For example, a writer running under the Local System account meets this requirement. This is true for the vast majority of writers.

You can, however, enable a writer to use a specific user account by adding the appropriate registry entry. You can also prevent a writer from using a specific user account.

To enable a writer to use a specific user account, add a REG_DWORD entry with the name equal with the user name. Set the value of the registry entry to 1 (one).

To prevent a writer from using a specific user account, add a REG_DWORD entry with the name equal with the user name. Set the value of the registry entry to 0 (zero).

**VOLSNAP**
The following registry entry is located under \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VOLSNAP.

**MinDiffAreaFileSize**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VOLSNAP\MinDiffAreaFileSize

**Version**
Windows Server 2003

Controls the initial size of the shadow copy storage area for each snapshot. The MinDiffAreaFileSize registry key uses the REG_DWORD data type.

The MinDiffAreaFileSize registry key specifies the minimum size of the shadow copy storage area. The default size of the shadow copy storage area is 300 MB, and the maximum size is 3 gigabytes (GB). For an accurate setting, specify a value that is a multiple of 300 MB. Otherwise, the next multiple of 300 MB will be used. A value of 300 is equal to 300 MB, and a value of 3,000 is equal to 3 GB.

Shadow copies might not function correctly if the value of the MinDiffAreaFileSize registry key is larger than the maximum size

of the shadow copy storage area. Before you change this registry entry, run the **vssadmin list shadowstorage** command. Make sure that the shadow copy storage area is larger than 300 MB or that the shadow copy storage area is larger than the value that is set in the registry. Also, make sure that you have at least 300 MB of free space on the shadow copy storage volume.

**Volume Shadow Copy Service WMI Classes**

The following table lists and describes the WMI classes that are associated with the Volume Shadow Copy Service.

**WMI Classes Associated with Volume Shadow Copy Service**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_ShadowBy | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowContext | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowCopy | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowDiffVolumeSupport | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowFor | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowOn | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowProvider | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowStorage | \root\cimv2 | Windows Server 2003 family |
| Win32_ShadowVolumeSupport | \root\cimv2 | Windows Server 2003 family |
| Win32_Volume | \root\cimv2 | Windows Server 2003 family |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Backup Technical Reference
- Shadow Copies for Shared Folders Technical Reference

# Scalability Factors for Shadow Copies

Applies To: Windows Server 2003, Windows Server 2003 R2

With the increased adoption of the Volume Shadow Copy Service (VSS), we have had several requests from partners and customers for guidance on provisioning server and storage resources in their deployments. This document is intended to provide general information that a partner or customer needs to make appropriate choices for scalability when deploying Shadow Copies of Shared Folders.

This information refers to VSS deployments on Windows Server® 2003 with Service Pack 1 and Windows Server 2003 R2. There are three significant parameters to consider when planning for scalability: availability of paged pool memory, nonpaged pool memory, and system virtual memory.

**Paged pool memory and snapshot limitations**
The NTFS file system requires 1 to 2 MB of paged pool memory per million files for snapshots and volumes mounted on the system. The usage is independent of actual number of snapshots on the system. This is an approximation—the actual usage would vary depending on the file access patterns.

**Example:** Mounting 64 snapshots of a volume with 2 million files is equivalent to mounting 8 snapshots of a volume with 16 million files.

**Impact of file access patterns on paged pool memory usage**
NTFS maps file records in 64-MB chunks. Each 64-MB mapping of file records consumes 64 KB of paged pool memory. Each 64-MB mapping of file records represents 65,536 files, where each file is represented by a 1-KB file record. NTFS maps a new 64-MB range of its file records whenever one of those file records is accessed. If all the file records are mapped, this translates to 1 MB of paged pool memory used for every 1 million files. Besides the file records, there are many other NTFS disk data structures that need to be mapped—and those mappings also consume paged pool memory. That is why we estimate the paged pool memory usage as being between 1 MB and 2 MB per million files.

**Mitigation:**

- Split the volumes among more servers.
- Install the KB 891957 update described in the Microsoft Knowledge Base article 913648 (http://go.microsoft.com/fwlink/?LinkId=89811).

  Based on heuristics, this update dismounts all mounted shadow copies if:

  - More than 128 MB of paged pool memory is being used by the mounted shadow copies.

    Or

  - The available paged pool memory drops below 50 MB.

VSS attempts to estimate the amount of paged pool memory actually being used by the snapshots by monitoring the change in paged pool memory allocation whenever a snapshot is deleted or dismounted. The dismounts occur during shadow copy creation.

**Scope:**
This limitation is significant only on 32-bit systems, which may have a maximum paged pool memory of 650 MB. 64-bit systems support 128 GB of paged pool memory.

**Event:**
A lack of paged pool memory leads to events such as Event 2020.

**System virtual memory limitation**
For every 16 KB of diff area on a system, VSS consumes about 40 bytes of virtual memory on the system. However, optimizations used by VSS can reduce the usage of virtual memory, so a production system will most likely have better usage.

The total diff area on a system is configured by the shadow copy storage space setting and is often set at 10% of the volume to be shadow copied.

**Example:** To snapshot a 10 TB volume, we prescribe a 1 TB diff area, which requires about 2 GB of virtual memory after accounting for the optimizations used by VSS.

**Scope**
This limitation is significant only on 32-bit systems which have 2 GB of addressable space. 64-bit systems have 8 TB of virtual address space.

**Event**
A lack of virtual memory leads to Event 6.

**Nonpaged pool memory limitation**
The amount of nonpaged pool memory in the system is proportional to the amount of physical memory, but the amount of nonpaged pool memory on a 32-bit system is never greater than 256 MB—regardless of the amount of physical memory available. VSS consumes 1 bit (16 KB) of the volume to be shadow copied.

**Example:** A 1-TB volume that is shadow copied consumes 8 MB of nonpaged pool memory.

**Snapshot creation:**
VSS consumes 5 bits per volume to be shadow copied only during snapshot creation. This usage is pruned once the snapshot is created.

**Example:** A 5-TB volume consumes 200 MB of nonpaged pool memory during the snapshot creation phase and then drops to using 40 MB after that.

**Scope**
While this limitation applies to both 32-bit and 64-bit systems, it is less significant on 64-bit systems. On 64-bit systems, nonpaged pool memory is capped at 7/8th of physical memory and can never exceed 128 GB. 32-bit systems are limited to 256 MB of nonpaged pool memory, and 128 MB of nonpaged pool memory if the system is booted with the **/3GB** switch.

**Event:**
A lack of nonpaged pool memory usually results in Event 5.

**Additional resources**
For memory limits for specific Windows releases, see http://go.microsoft.com/fwlink/?LinkId=89378.

# Device Management Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this subject**


- Device Management Architecture
- Device Management Components
- Device Management Scenarios
- New Storage Device Drivers in Windows Server 2003


Device management provides a way for Microsoft Windows Server 2003 to uniformly notify all applications and system components of changes that may affect their operation or access to resources.

Applications and system components use and process device messages between the operating system and devices to take advantage of new resources when they become available and to prevent loss of data when existing resources become unavailable. Device messages notify applications of device change events, such as the removal of a hardware device; power messages notify applications of power management events. Applications and device drivers can also define and use custom messages to enable notification of other types of events.

Windows Server 2003 uses the Device Manager snap-in to dynamically track and manage devices and drivers on a system. For example, you can view device status and set device properties through Device Manager. If you request to update a driver, Device Manager calls the Hardware Update wizard.

Device Manager displays all devices installed in the system, as illustrated in the following figure.

**Device Manager**




The Device Manager display represents the computer's hardware configuration. The Device Manager display is recreated each time the computer is started and whenever a dynamic change to the computer configuration occurs, such as the addition of a new device while the system is running.

Device Manager allows you to do the following:

- Determine whether the hardware on your computer is working properly

- Change hardware configuration settings
- Identify the device drivers that are loaded for each device and obtain information about each device driver
- Change advanced settings and properties for devices
- Install updated device drivers
- Disable, enable, and uninstall devices
- Roll back to the previous version of a driver
- Print a summary of the devices that are installed on the computer
- Change resources, such as interrupt requests (IRQs), assigned to devices

Device Manager can be used to check the hardware status and update device drivers on a computer. Device Manager cannot be used to change resource allocation settings; resources are automatically allocated by the system during hardware setup. In order to manage device drivers through Device Manager, you must be a member of the Administrators group.

Some hardware cannot handle dynamic resource assignments. However, these resource assignments can be assigned through Device Manager.

Device Manager manages devices on a local computer only. On a remote computer, Device Manager works only in read-only mode.

The following table lists some of the major device classes that Device Manager tracks. Each of these classes has a corresponding class GUID value in the Windows registry.

**Device Manager Class**

| Class | Description |
|---|---|
| Battery Devices | Batteries and UPS devices. |
| Biometeric Device | Biometric-based personal identification devices. |
| Bluetooth Devices | Bluetooth devices. |
| CD-ROM Drives | CD-ROM drives, including SCSI CD-ROM drives. |
| Disk Drives | Hard disk drives. |
| Display Adapters | Video adapters installed on the system. |
| DVD/CD-ROM drives | CD-ROM drives, including SCSI CD-ROM drives, and DVD drives. |
| Floppy Disk Controllers | Floppy disk drive controllers. |
| Floppy Disk Drives | Floppy disk drives. |
| Hard Disk Controllers | Hard disk controllers, including ATA/ATAPI controllers (not SCSI or RAID disk controllers, which have their own Class). |

| | |
|---|---|
| Human Interface Devices | Interactive input devices not operated by the system-supplied HID class driver, which includes USB devices that comply with USB HID Standard and non-USB devices that use a HID minidriver. |
| Imaging Device | Still-image-capture devices, digital cameras, and scanners. |
| IrDA Devices | Infrared devices, including Serial-IR and Fast-IR NDIS miniports. The Network Adapter class displays other NDIS NIC miniports. |
| Keyboards | Keyboard device. |
| Media Changers | SCSI media changer devices. |
| Memory Technology Driver | Memory devices, such as flash memory cards. |
| Monitors | Each of the monitors. Specifies the features of a particular monitor to be stored in the registry for use by drivers of video adapters. |
| Mouse | Mouse devices and other pointing devices, such as trackballs. |
| Multifunction Devices | Combination cards, such as a PCMCIA modem and network card adapter. |
| Multimedia | Audio and DVD multimedia devices, joystick ports, and full-motion video capture devices. |
| Network Adapters | Network adapters, including NDIS NIC miniport drivers, excluding Fast-IR miniport drivers, NDIS intermediate drivers (of virtual adapters), and CoNDIS MCM miniport drivers. |
| PCMCIA Adapters | PCMCIA and CardBus host controllers, but not PCMCIA or CardBus peripherals. |
| Ports (COM & LPT) | Serial and parallel ports. |
| Printers | All printers. |
| Processors | Processor types. |
| SCSI and RAID controllers | SCSI HBAs (Host Bus Adapters) and disk-array controllers. |
| Smart Card Readers | Smart card readers. |
| Storage Volumes | Storage volumes defined by the system-supplied logical volume manager and class drivers that create device objects to represent storage volumes, such as the system disk class driver. |
| System Devices | HALs, system buses, system bridges, the system ACPI driver, and the system volume manager driver. |
| USB | USB host controllers and USB hubs, but not USB peripherals. |
| Windows CE USB ActiveSync Devices | Windows CE ActiveSync devices. |

For a complete list of the system-supplied device setup classes, along with their class GUID entries in the registry, see "Microsoft Windows Driver Development Kits" on the Microsoft Web site.

**Device Management Architecture**

The following diagram shows an overview of how device drivers and device management components are organized on a system.

**Device Management Architecture**




**Device Management Components**

The following table describes the components used for managing devices.

**Device Management Architecture Components**

| Component | Description |
|---|---|
| Hardware Abstraction Layer (HAL) | A component based on the Windows NT operating system that provides platform-specific support for the Windows NT kernel, I/O Manager, kernel-mode debuggers, and the lowest-level device drivers. The HAL exports routines that abstract platform-specific hardware details about caches, I/O buses, interrupts, and so forth. The HAL provides an interface between the platform's hardware and the system software. For example, the HAL implements a routine to map each device driver's bus-relative device interrupt vector to a system-assigned vector with a corresponding platform-specific hardware priority (DIRQL). The HAL also provides several routines that provide DMA transfer support. |
| Device Drivers | Files used to communicate between devices and Windows Server 2003. Drivers are structured under the Windows Device Model (WDM). The three WDM driver types are bus driver, function driver, and filter driver. |
| I/O Manager | Subsystem that controls and interacts with all devices on the system and provides routines that drivers can call to have the I/O Manager insert IRPs into the associated device queue. |
| Plug and Play Manager | Component responsible for determining the hardware resources requested by each device (for example, input/output ports [I/O], interrupt requests [IRQs], direct memory access [DMA] channels, and memory locations) and assigning hardware resources appropriately. Plug and Play Manager reconfigures resource assignments when necessary, such as when a new device added to the system requires resources already in use. |
| Power Manager | Component responsible for managing power usage for the system. Power Manager administers the system-wide power policy and tracks the path of power IRPs through the system. |
| Device Manager | Interface displaying the computer's hardware configuration information. The Device Manager display is recreated each time the computer is started or whenever a dynamic change to the computer configuration occurs, such as addition of a new device while the system is running. Device Manager is used to enable or disable devices, troubleshoot devices, update drivers, use driver rollback, and change resources such as interrupt requests (IRQs) assigned to devices. |

**Device Management Scenarios**

This section contains descriptions of some common scenarios for device management.

**Viewing system devices**

Device Manager provides a quick view of a system and all the hardware devices installed. Device Manager also provides information about the version of the driver installed, the provider, and digital signature, if any.

**Troubleshooting driver problems**

Device Manager automatically flags the status of a device so that you can readily detect where hardware problems or conflicts occur.

**Updating device drivers**

Device Manager allows you to manage the devices on your system by allowing you to check a specific location, such as Windows Update or a CD, for updated device drivers. Updated device drivers are then automatically installed and resources reallocated.

**New Storage Device Drivers in Windows Server 2003**
New and updated storage device drivers are available for Windows Server 2003, including drivers for iSCSI and Multipath I/O (MPIO) technologies. Windows Server 2003 also includes a new port driver model, Storport.

### Note

- The Microsoft Storage Server 2003 Web site has the most recent information about storage devices and drivers, as well as storage architecture.

Storage architecture consists of a series of layered drivers, as shown in the following figure. When a storage device is added, each layer is responsible for ensuring that the disk is functional.

### Windows Storage Stacks




The following table describes the components of the previous diagram.

### Windows Storage Stacks Components

| Component | Description |
|---|---|
| I/O Subsystem | Sends I/O requests to the file system. |
| File System | Structure of a disk, either FAT or NTFS. File systems access and manage all data on the disk, including information about the structure of the disk itself. |
| Volume Snapshot | Manages software snapshots. |
| Volume Manager | Presents volumes to users; supports basic and dynamic disks. |
| Partition Manager | Manages disk partitions. |
| Class | Manages specific devices types, such as disks and tapes. |
| Port | Manages specific transports, such as SCSIport or Storport. |
| Miniport | Vendor-supplied functionality linked to a specific port driver. Manages hardware-specific details. |
| Disk Subsystem | Disk processes such as read/write. |
| Device Stack | Stack layers that manage storage devices. |

| | |
|---|---|
| Driver Stack | Chain of drivers that support the operations of a device. |

**Storport**

The Storport driver is new to Windows Server 2003. Storport implements an architecture designed for better performance for RAID systems and in Storage Area Network (SAN) environments. Storport delivers higher I/O throughput, enhanced manageability, and an improved miniport interface. Storport better utilizes faster adapters through the use of reduced Delay Procedure Call (DPC) and improved queue management.

**iSCSI**

Internet Small Computer System Interface (iSCSI) is an alternative to Fibre Channel protocols. iSCSI is used to transfer data storage over long distances, making it a key contributor to the advancement and growth of storage area networks (SANs). iSCSI is also used to transmit data over local area networks (LANs), wide area networks (WANs), or the Internet.

Microsoft released the Microsoft iSCSI Software Initiator in June 2003 for Windows 2000, Windows XP, and Windows Server 2003. The Microsoft iSCSI Software Initiator is not included with Windows Server 2003, but you can download the driver from the Microsoft Download site. Customers should ensure that the hardware included in their iSCSI solution has passed the Designed for Windows Logo Program qualification tests.

**Multipath I/O**

Microsoft Multipath I/O architecture allows hardware manufacturers to develop highly available storage devices that can be added easily to a computer running Windows 2000 or Windows Server 2003. The Multipath I/O (MPIO) DDK consists of three standard drivers plus a fourth, device-specific driver that a storage vendor will use to create a Microsoft MPIO solution that works with the specific storage array it supports.

Microsoft MPIO enables the use of multiple physical paths from a server to a storage device. A physical path includes adapters, cables, and switches that connect the storage device to the server. In the event that one of the components in a path fails, all subsequent I/O is routed over an alternate path, providing applications uninterrupted access to the data on the storage device. In addition, Microsoft MPIO can be used to balance I/O operations over these paths using methods and algorithms defined by the storage vendor.

# Device Driver Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft Windows Server 2003 uses device drivers to interact with hardware devices, such as displays and keyboards. Rather than access devices directly, Windows loads device drivers and calls functions in the drivers to carry out actions on the devices. Each device driver exports a set of functions; Windows calls these functions to complete an action, such as drawing a circle on a display or translating a keyboard scan code. The driver functions contain the device-specific code needed to carry out actions on the device.

Windows requires device drivers for the display, keyboard, and communication ports. Other drivers may also be required if the user adds optional devices to the system.

**In this subject**


- What Is a Device Driver?
- How Device Drivers Work
- Device Driver Tools and Settings

# What Is a Device Driver?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is a Device Driver?**

**In this section**


- Device Driver Dependencies and Interactions
- Related Information


A device driver is a component that Windows uses to provide I/O services for and interact with an underlying device, such as a modem or network adapter. Rather than access the device directly, Windows loads device drivers and calls functions in the drivers to carry out actions on the device. The driver functions contain the device-specific code needed to carry out actions on the device. Device drivers for Windows Server 2003 are usually library files with .sys filename extensions.

Windows Server 2003 installs and configures device drivers automatically. This saves you from needing to allocate resources manually or check for conflicts in devices.

Windows Server 2003 checks for a digital signature during device driver installation and issues a message if the driver is not signed. Windows Server 2003 checks a database of known problem drivers when a computer is started or when a device driver is loaded; if a driver is located in the database of known problem drivers, that driver cannot be loaded, installed, or used on the computer.

**Device Driver Dependencies and Interactions**

**Plug and Play**

For automatic configuration of device drivers to occur, the system should have a BIOS compatible with Plug and Play. Plug and Play technology provides software and hardware support that enables the Windows operating system to detect and configure the hardware with little or no user involvement. This makes it easier to add and configure hardware on a computer running Windows Server 2003 without needing special knowledge about hardware configurations.

**Power Management**

Windows Server 2003 offers enhanced power-management features for desktop and mobile computers. The operating system supports the Advanced Configuration and Power Interface (ACPI) specification, which provides reliable power management and system configuration features. Devices and drivers designed in compliance with the ACPI specification work with the operating system to respond to or request changes in the system power state.

The following illustration shows a system-wide overview of how device drivers interact with other components of a system running Windows Server 2003.

**Device Drivers**




The following table describes the components used for installing and managing devices and drivers.

**Device Driver Components**

| Component | Description |
|---|---|

| | |
|---|---|
| Hardware Abstraction Layer (HAL) | A component based on the Windows NT operating system that provides platform-specific support for the Windows NT kernel, I/O Manager, kernel-mode debuggers, and the lowest-level device drivers. The HAL exports routines that abstract platform-specific hardware details about caches, I/O buses, interrupts, and so forth. The HAL provides an interface between the platform's hardware and the system software. For example, the HAL implements a routine to map each device driver's bus-relative device interrupt vector to a system-assigned vector with a corresponding platform-specific hardware priority (DIRQL). The HAL also provides several routines that provide DMA transfer support. |
| Device Drivers | Files used to communicate between devices and Windows Server 2003. Drivers are structured under the Windows Device Model (WDM). The three WDM driver types are bus driver, function driver, and filter driver. |
| I/O Manager | Subsystem that controls and interacts with all devices on the system and provides routines that drivers can call to have I/O Manager insert IRPs into the associated device queue. |
| Plug and Play Manager | Component responsible for determining the hardware resources requested by each device (for example, input/output ports [I/O], interrupt requests [IRQs], direct memory access [DMA] channels, and memory locations) and assigning hardware resources appropriately. Plug and Play Manager reconfigures resource assignments when necessary, such as when a new device added to the system requires resources already in use. |
| Power Manager | Component responsible for managing power usage for the system. Power Manager administers the system-wide power policy and tracks the path of power IRPs through the system. |
| Device Manager | Interface displaying the computer's hardware configuration information. The Device Manager display is recreated each time the computer is started or whenever a dynamic change to the computer configuration occurs, such as addition of a new device while the system is running. Device Manager is used to enable or disable devices, troubleshoot devices, update drivers, use driver rollback, and change resources such as interrupt requests (IRQs) assigned to devices. |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Plug and Play Technical Reference
- "Microsoft Windows Driver Development Kits" on the Microsoft Web site

# How Device Drivers Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Device Drivers Work**

**In this section**


- Device Driver Architecture
- Device Driver Physical Structure
- Device Driver Processes and Interactions
- Related Information


A device driver is a component that Windows uses to interact with a hardware device, such as a modem or network adapter. Rather than access the device directly, Windows loads device drivers and calls functions in the drivers to carry out actions on the device. Each device driver exports a set of functions; Windows calls these functions. The driver functions contain the device-specific code needed to carry out actions on the device.

**Driver Signing and Digital Signatures**

Windows Server 2003 includes many features that help to ensure that the device drivers installed on a computer are reliable and up to date.

Drivers are signed by Microsoft after they pass a series of tests for functionality and reliability. For a driver to earn this certification, it must pass a series of tests administered by the Windows Hardware Quality Labs (WHQL). Due to stringent WHQL standards, using signed drivers typically results in a more stable system. Microsoft digitally signs drivers that pass the WHQL tests and Windows Server 2003 performs signature detection for signed drivers for various device types, which include the following:

- Keyboards
- Hard disk controllers
- Modems
- Mouse devices
- Multimedia devices
- Network adapters
- Printers
- SCSI adapters
- Smart card readers
- Video adapters

**Note**

- During installation of some hardware devices, you may be prompted to choose **Ignore** on the driver signing dialog box. This indicates that the driver has not gone through digital signing.

A digital signature indicates that a driver file is an original, unaltered system file that Microsoft has approved for use with Windows Server 2003.

Windows Server 2003 checks for a digital signature during device driver installation and issues a message if the driver is not signed. Windows Server 2003 checks a database of known problem drivers when a computer is started or when a device driver is loaded; if a driver is located in the database of known problem drivers, that driver cannot be loaded, installed, or used on the computer.

Updated versions of signed drivers are available for download from the Windows Update Web site.

**Device Driver Architecture**
The Windows Server 2003 I/O system provides a layered architecture for serial device drivers, consisting of the drivers and their corresponding device objects.

**Plug and Play Drivers**
Plug and Play supports three types of drivers:

- Function drivers, which drive individual devices
- Filter drivers, which filter I/O requests for devices, classes of devices, and buses
- Bus drivers, which service bus controllers, adapters, and bridges

The most common bus driver types are PCI, PnP ISA, SCSI, and USB. During enumeration, a bus driver identifies the devices on its bus and creates device objects for them.

**Device Objects**
A driver creates a device object (DO) for each device it controls; the device object represents the device to the driver. Plug and Play recognizes three types of device objects:

- Physical device objects (PDOs), which represent individual devices on a bus to a bus driver
- Functional device objects (FDOs), which represent individual devices to a function driver
- Filter device objects (filter DOs), which represent individual devices to a filter driver

The following figure shows a chain of attached device objects and drivers for a Plug and Play device.

**Device Objects in a Device Stack**




The device stack includes a device object for each driver involved in handling the I/O to a particular device.

**Components for Device Objects in a Device Stack**

| Component | Description |
|---|---|
| Lower-level filter driver | Optional filter driver, loaded after the bus driver and any bus driver filters. There are two types of filters, device-specific and per-class. A lower-level device-specific filter monitors and/or modifies I/O requests to a particular device; typically, such filters redefine hardware behavior to match specifications. A lower-level per-class filter monitors and/or modifies I/O requests for a class of devices; for example, a lower-level class filter driver for mouse devices could provide acceleration, performing a nonlinear conversion of mouse movement data. |

| | |
|---|---|
| Upper-level filter driver | Optional filter driver, loaded after the function driver, which attaches its device object above the FDO. There are two types of upper-level filters, device-specific and per-class. An upper-level device filter driver adds value for a particular device; for example, an upper-level device filter for a keyboard could enforce additional security checks. An upper-level per-class filter adds value for all devices of a particular class. |

For Plug and Play devices, the device object at the bottom of the device stack is the PDO created by the underlying bus driver. The FDO is created by the function driver for the device. Filter DOs are attached above and below the FDO.

The PDO is created by the underlying bus driver when it enumerates the device. The FDO and filter DOs are created by FDOs or filter DOs for devices enumerated by the Plug and Play Manager

**Device Driver Physical Structure**

The following figure shows how device drivers operate in both kernel mode and user mode.

**Device Driver Physical Structure**




The following table describes the components of the preceding Device Driver Physical Structure diagram.

**Components for Device Drivers**

| Component | Description |
|---|---|
| Win32 API | 32-bit application programming interface for Windows operating systems that initiates I/O operations through Ntdll |
| Ntdll | Function library that exposes stubs that invoke Windows system calls (like NtReadFile below) as exported functions |
| NtReadFile | I/O Manager system call that creates an I/O request packet and directs it to the appropriate driver |
| I/O Manager | Subsystem that controls and interacts with all devices on the system and provides routines that drivers can call to have the I/O Manager insert IRPs into the associated device queue |
| IRP | I/O Request Packet sent by I/O Manager to request services from a driver |
| IoCallDriver | Routine that sends an IRP to the driver associated with a specified device object |

Once the request has been processed, the outcome is one of three results:

**Queue for later process**

The device driver defers the request processing to a later time; for example, when necessary resources are available, or the device is not busy.

**Send to hardware port**
The device driver programs the hardware to perform the operation specified in the request.

**Send to another device driver**
The device driver sends the request to another device driver; for example, if the request parameters change.

**Device Driver Processes and Interactions**
This section describes device driver processes and interactions for computers running Windows Server 2003.

**How Windows Server 2003 Installs Device Drivers**
Windows Server 2003 automatically configures a device so that it will work properly with other devices that are installed on the computer. As part of the configuration process, Windows assigns a unique set of system resources to each device so that it can function with little or no user input.

Windows Server 2003 device drivers included on the Setup CD are stored in a single cabinet file, Driver.cab. This file is used by Setup and other system components as a driver file source. Information files (.inf files) are searched when Windows Server 2003 starts or new hardware is detected. These text files provide the names and locations of driver-related files and the initial settings required for new devices to work. During setup, Driver.cab is copied from the installation CD to the WINNT\Driver Cache*platform* directory of the local hard disk. The variable *platform* here is a placeholder for a specific value that reflects the architecture of the system: for example, i386.

Device drivers are installed with no user intervention if certain conditions are met:

- Installing the driver does not require showing a user interface.
- The driver package contains all files needed to complete the installation.
- The driver package is available on the system in the Driver.cab file or was previously installed.
- The driver package is digitally signed.
- No errors occur during installation.

If any of the conditions is not met, the device installation restarts and the user might need to respond to dialog boxes or messages. Manual installation of a driver requires that the person installing the driver be a member of the Administrators group.

Windows Server 2003 determines which device driver to load for a device using the following criteria.

- Driver ranking schemes
- Driver search location policies
- Windows Driver Protection
- Windows Update

To facilitate device installation, devices that are set up and configured in the same way are grouped into a device setup class. For example, SCSI media changer devices are grouped into the MediumChanger device setup class.

**How Device Drivers Are Processed**
Microsoft uses a multistage process to test device drivers. Drivers are subjected to compatibility, reliability, and functionality tests administered by the Windows Hardware Quality Lab (WHQL). Drivers that successfully complete the process are digitally signed. Because of this testing, signed drivers are typically more robust and reliable. Once a driver is digitally signed, Windows Server 2003 recognizes it when it is loaded. Windows Server 2003 will notify a user if a driver is not signed or if a driver file has been changed since its inclusion on the Microsoft Hardware Compatibility List (HCL).

The digital signature is associated with an individual driver version and certifies that the driver provided with the device is identical to the driver that was tested.

**Driver Signature Checking**

Driver Signature Checking enables you to detect unsigned drivers before you install them. You can set verification levels for driver signature checking to ensure that Windows Server 2003 inspects files for digital signatures whenever you install or update drivers.

Signature Checking Levels are set to one of three categories, as described in the following table. These levels apply to processing both signed and unsigned drivers.

**Signature Checking Levels**

| Level | Description |
|---|---|
| Level 0 (Ignore) | Disables digital signature checking. Setup silently checks the signature on the driver, logs any unsigned driver files to a log file, and allows installation of the driver even if the driver is unsigned. |
| Level 1 (Warn) | Checks the signature on the driver before installation by checking whether the driver has passed WHQL testing. A message appears whenever a user tries to install a driver that fails a signature check. The driver can still be installed; however, this is not recommended. |
| Level 2 (Block) | Blocks installation of a driver that fails a signature check. The user is notified that Windows Server 2003 cannot install the unsigned driver. |

**Note**

- Warn is the default setting. You can change the signature-checking policy for a user without administrator permissions, but you must have administrator permissions to change the driver-signing policy for a computer.

**Verifying File Signatures**

The File Signature Verification Tool (Sigverif.exe) detects signed files and allows you to do the following:

- View the certificate of a signed files to verify that the file has not been tampered with after being digitally signed
- Search for signed files in a specific location
- Search for unsigned files in a specific location

**How Drivers are Ranked**

Windows Server 2003 uses driver ranking schemes to determine which driver to load when multiple drivers are available for a device. Drivers are ranked by whether they are signed and how closely their Plug and Play ID matches the device's Plug and Play ID. The Plug and Play ID of a driver or device consists of hardware IDs and compatible IDs. If the hardware ID of the driver exactly matches one of the hardware IDs of the device, there is a hardware match. If some other match occurs (for example, device compatible ID to driver compatible ID) there is a compatible match. Driver rank also depends on whether the device information file (.inf file) for the device includes information specifically for installations in a Windows Server 2003 environment. If multiple drivers for a device exist, the lowest ranking driver — the driver with the highest chance of optimum functionality — is installed.

The following table summarizes the driver ranking scheme for Windows Server 2003 from lowest (best match) to highest.

**Windows Server 2003 Driver Ranking Scheme**

| Rank | Description |
|---|---|
| 1 | Signed driver with a hardware match to a device |
| 2 | Signed driver with a compatible match to a device |
| 3 | Unsigned driver with a hardware match to the device, with Windows Server 2003 targeted in the INF section |
| 4 | Unsigned driver with a compatible match to the device, with Windows Server 2003 targeted in the INF section |
| 5 | Unsigned driver with a hardware match to the device, without Windows Server 2003 targeted in the INF section |

| | |
|---|---|
| 6 | Unsigned driver with a compatible match to the device, without Windows Server 2003 targeted in the INF section |

**How Driver Protection Affects Applications**

Windows Driver Protection prevents the installation of drivers known to cause instabilities when installed under Windows Server 2003. Windows Server 2003 includes a list of such drivers, known as the Driver Protection List, which is checked at the following times.

- During installation of Windows Server 2003
- Each time the system is rebooted
- Each time a device or application is installed

Driver Protection compares each driver that attempts to load under Windows Server 2003 against the Driver Protection List before loading that driver. The comparison and matching happen using the criteria in the following table.

**File Attributes Matched by Driver Protection**

| Key | Meaning |
|---|---|
| SYS_NAME | File name of the driver file. |
| BIN_PRODUCT_VERSION | Product version of the driver file. |
| UPTO_BIN_PRODUCT_VERSION | Superset of the BIN_PRODUCT_VERSION. Represents all product versions of the driver file that are less than or equal to the version specified. |
| LINK_DATE | Link date of the driver file. LINK_DATE appears to the operating system as a more accurate and consistent value than BIN_PRODUCT_VERSION, primarily due to the various ways versioning information can be stored. LINK_DATE is consistently set by the linker. |
| UPTO_LINK_DATE | Superset of LINK_DATE. Represents all link dates of the driver file that are less than or equal to the date specified. |

If Driver Protection finds a match for a driver in the Driver Protection List, the following occurs:

- The driver is not loaded under Windows Server 2003.
- The operating system logs an event in the Event Log.
- A notification is displayed to the user.

Windows Driver Protection is a low-level kernel feature that was designed to be the last resort for protecting a Windows Server 2003 system from errant drivers. Driver Protection will block known problem drivers from loading, but it does not block any associated applications that might depend on those drivers. Therefore, the behavior of applications that depend on blocked drivers will vary depending on the implementation of the application.

Some applications — such as antivirus programs — install drivers in order to provide their core functionality. Other applications — such as CD-burning applications — use drivers for portions of their feature sets. If a driver does not load, the former category of application might not work at all, while the latter might work for features that do not depend on the driver. Because Microsoft works closely with the driver and application vendors to implement fixes, in most cases the application vendor will have an updated version compatible with Windows Server 2003 available on the vendor's Web site for download.

Most applications that ship drivers on the Driver Protection List will already be in the Windows Server 2003 application compatibility database. Users will get warnings whenever these applications are installed or run. However, because many drivers are sublicensed to multiple application vendors, Microsoft may not have information about every application that includes the driver. Therefore, for some applications, only the Driver Protection warning will appear.

The Driver Protection List database is updated through Microsoft Windows Update. Drivers are added to the Driver Protection List based upon feedback from users about problems that can be reproduced and confirmed by Microsoft. Whenever a system connects to the Windows Update Web site, the system's driver protection list is updated. If a driver on the list was previously installed, the driver will be disabled and the user will be told how to obtain a new driver.

**How Drivers are Located**
Driver search location policies are Group Policies that specify where to search for new or updated device drivers. Group Policy can be configured to search a local machine and/or Windows Update, as well as turning off this policy.

**Windows Update**
The Windows Update Web site provides a central location for product enhancements, such as device drivers.
Windows Server 2003 users can install or update drivers from the Windows Update Web site. When a user accesses the Windows Update Web site, Windows Update compares the drivers installed on a user's system with the latest updates available. If newer drivers are found, Windows Update offers the list of applicable drivers to the user.

Drivers are included on Windows Update only if they are digitally signed, they have passed the testing requirements for the Windows Logo Program, and the vendor has given Microsoft redistribution rights for those drivers. This helps to ensure that the drivers offered from Windows Update are of high quality and reliability.

Devices each have a hardware ID that uniquely identifies the device, which is identified in the .inf file for the driver. Drivers that have a hardware or compatible match with the device, such as a device hardware ID to driver hardware ID, are candidates for download and installation. If a hardware or compatible match exists, Windows Update determines whether the driver on Windows Update is newer than the installed one. If the driver on Windows Update is newer, the driver is presented to the user for potential download. Also, if the hardware ID for the driver on Windows Update is a better match than the installed one, Windows Update offers that driver to the user. If the user chooses to install the offered driver, the file is downloaded, and the Windows Update ActiveX control points Device Manager to the .inf file for installation.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Plug and Play Technical Reference
- "Microsoft Windows Driver Development Kits" on the Microsoft Web site

# Device Driver Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Device Driver Tools and Settings**

**In this section**


- Device Driver Tools
- Device Driver Registry Entries
- Device Driver Group Policy Settings
- Device Driver WMI Classes
- Related Information


This section describes tools, registry entries, Group Policy settings, and Windows Management Instrumentation (WMI) classes associated with device drivers.

**Device Driver Tools**

The following tools are associated with device drivers.

**ChkINF.exe: INF File Syntax Checker**

**Category**

ChkINF.exe ships with the Windows Driver Development Kit.

**Version compatibility**

Windows 2000, Windows XP, and Windows Server 2003

ChkINF checks the structure and syntax of Windows Server 2003 driver .inf files. Results are presented in HTML format and include a list of the errors and warnings detected in each .inf file, with each error and warning shown next to the errant .inf file entry.

**DevCon.exe: Device Console**

**Category**

Device Console ships with the Windows Driver Development Kit.

**Version compatibility**

Windows 2000, Windows XP, and Windows Server 2003

Device Console (Devcon) is a command-line tool that displays detailed information about devices and lets you search for and manipulate devices from the command line. Devcon enables, disables, installs, configures, and removes devices on the local computer and displays detailed information about devices on local and remote computers.

For more information about Devcon.exe and the syntax used to carry out commands, see the Windows Driver Development Kit.

**GenINF.exe: INF File Creation Wizard**

**Category**

GenINF ships with the Windows Driver Development Kit.

**Version compatibility**

GenINF generates .inf files for Windows 2000, Windows XP, and Windows Server 2003.

GenINF is a GUI-based tool that guides users through the process of creating device driver .inf files. GenINF does not generate fully valid .inf files. However, the skeleton .inf files that it generates should help developers complete their device driver .inf files.

GenINF is located in the \Tools subdirectory of the DDK installation directory. Because GenINF is GUI-based, you can execute it from within the Windows Server 2003 user interface or from the command prompt. The command to start GenINF takes no arguments. GenINF supports all device classes but includes specific support for only certain device classes. For a list of these classes, see the Geninf.htm file located in the Tools\geninf subdirectory of the DDK.

For more information about GenINF.exe, see the Windows Driver Development Kit.

**SigVerif.exe: File Signature Verification**
**Category**
Operating system tool

**Version compatibility**
Windows XP, Windows 2000, and Windows Server 2003

File Signature Verification checks to see which system files and device driver files are digitally signed and displays its findings. If you have enabled logging, the search results are also written to a log file.

**Device Driver Registry Entries**
The following registry entries are associated with device drivers.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools — such as Microsoft Management Console (MMC), Device Manager, or the software that accompanies a device — to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

### Note

- Because many of the device driver subkeys and registry entries are hardware-specific, the following section contains only a partial list of registry keys and entries.

**Enum**
This registry key is located at HKEY_LOCAL_MACHINE\SOFTWARE\Enum.

This class key stores the enumeration information for Plug and Play devices. Under Enum, each enumerator gets its own branch, with one child under the enumerator for each device it has ever enumerated and set up. There is a special enumerator branch, Enum\Root; this branch contains registry entries for legacy hardware that is not Plug and Play for which there is no enumerator but that can be detected using other means. Device Installer creates new registry entries when an enumerator detects a device for which there is no registry entry.

**Class**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class.

The Class subkey stores configuration data for classes of hardware devices displayed in Device Manager.

Each subkey in the Class subkey represents a device class and is named by the globally unique identifier (GUID) for the class: a 16-byte value generated from the unique identifier on a device, the current date and time, and a sequence number. A GUID is used to identify a particular device or component.

**{Class-GUID}**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class

**Version**
Microsoft Windows NT 4.0 Server and Windows NT 4.0 Workstation and later

The {*Class-GUID*} subkey represents the component associated with the class-GUID.

Each {*Class-GUID*} subkey in the Class subkey represents a device class and is named by the GUID for the class, which is a 16-byte value generated from the unique identifier on a device, the current date and time, and a sequence number. A GUID identifies a particular device or component.

It appears as a 16-digit alphanumeric string within curly braces. These {*Class-GUID*} subkeys store configuration data that apply to the entire class of devices, such as the class of all modems.

Within the {*Class-GUID*} subkeys are 4-digit numbered subkeys, which represent particular devices that are instances of the class. These subkeys store data that apply only to that device.

Because the Class subkey is hardware-specific, its contents vary widely and it is dependent upon the devices installed on a system.

**Services**
This registry key is located at HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services.

Services registry keys and entries define background operations by device drivers or other tasks. The Services key includes a number of subkeys for such services as ACPI or Plug and Play. Registry entries for each subkey vary, but each subkey may include some of the following registry entries.

**Description**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services*Subkey*

**Version**
Windows XP, Windows 2000, Windows Server 2003

Text field that displays the description of the service.

**DisplayName**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services*Subkey*

**Version**
Windows XP, Windows 2000, Windows Server 2003

Text field that displays the default service name.

**ErrorControl**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services*Subkey*

**Version**
Windows XP, Windows 2000, Windows Server 2003

Tells the system how to boot if a device driver fails. The following table lists the accepted values for this registry entry.

### ErrorControl Values

| Value | Description |
|---|---|
| 00 | Ignore. If the driver fails to load or initialize, then startup continues as normal and no warning message appears. |
| 01 | Normal. If the driver fails to load or initialize, then startup proceeds, but a warning message appears. |
| 02 | Severe. If the driver fails to load or initialize, then startup is declared to have failed and restarts by using the LastKnownGood control set. If startup is already using the LastKnownGood control set, then startup continues |
| 03 | Critical. If the driver fails to load or initialize, then startup is declared to have failed and restarts by using the LastKnownGood control set. If startup is already using the LastKnownGood control set, then startup stops and enters Debugging Mode. |

**Group**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services*Subkey*

**Version**
Windows XP, Windows 2000, Windows Server 2003

Specifies which group the service belongs to when a service is started

**ImagePath**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services*Subkey*

**Version**
Windows XP, Windows 2000, Windows Server 2003

Location of Services Control Manager, which uses the information located in the subkeys to initialize the services. The default path is %SystemRoot%\system32\services.exe.

**Start**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services*Subkey*

**Version**
Windows XP, Windows 2000, Windows Server 2003

Determines when a service should be started on a system. The following table lists the accepted values for this registry entry.

### Start Values

| Value | Description |
|---|---|
| 00 | Boot. Starts the service during the boot sequence. |
| 01 | System. Starts the service when loaded by the I/O subsystem in the kernel. |
| 02 | Automatic. Starts the service automatically. |
| 03 | Manual. Starts the service when a user initiates it from the control panel or a command line. |
| 04 | Disable. Loads the service, but neither the system nor a user can start it. |

**Type**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services*Subkey*

**Version**
Windows XP, Windows 2000, Windows Server 2003

Identifies the type of service. The following table lists the accepted values for this registry entry.

### Type Values

| Value | Description |
|---|---|
| 01 | Kernel-mode device driver |
| 02 | File system driver |
| 04 | Network adapter |
| 10 | Win32 service that can run as a standalone process |
| 20 | Win32 service that shares address space with other services at the same time |

**Device Driver Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with device drivers.

### Group Policy Settings Associated with Device Drivers

| Group Policy Setting | Description |
|---|---|
| Administrator_Power_Policy | Defines limits to certain power policy values that are applied globally to all users' power schemes. |
| Global_Machine_Power_Policy | Contains global computer power policy settings that apply to all power schemes for all users. |
| Global_Power_Policy | Contains global power policy settings that apply to all power schemes. |
| Global_User_Power_Policy | Contains global user power policy settings that apply to all power schemes. |
| Machine_Power_Policy | Contains computer power policy settings that are unique to each power scheme on the computer. |

| | |
|---|---|
| Machine_Processor_Power_Policy | Contains processor power policy settings that apply while the system is running on AC power or battery power. |
| Power_Action_Policy | Contains information used to set the system power state. |
| Power_Policy | Contains power policy settings that are unique to each power scheme. |
| Processor_Power_Policy | Contains information about processor performance control. |
| Remove Computer from Docking Station | Determines whether a user can undock a portable computer from its docking station without logging on. If this policy is set, the user must log on before removing the portable computer from its docking station. If this policy is not set, the user can remove the portable computer from its docking station without logging on. |
| Allow undock without having to log on | Determines whether a user must log on to request that a portable computer be removed from a docking station. If this policy is not set, the user must log on to request an undock, at which time the user must have been granted the Remove Computer from Docking Station privilege. If this policy is set, a user can request an undock by pressing the portable computer's physical eject button. |
| Allowed to format and eject removable media | Determines who is allowed to format and eject removable NTFS media. This capability can be given to Administrators, Administrators and Power Users, or Administrators and Interactive Users. |
| Prevent users from installing printer drivers | Determines who is allowed to install a printer driver as part of adding a network printer. For a computer to print to a network printer, the driver for that network printer must be installed on the local computer. If this policy is set, only Administrators and Power Users can install a printer driver as part of adding a network printer. If this policy is not set, any user can install a printer driver as part of adding a network printer. You can use this setting to prevent unprivileged users from downloading and installing an untrusted printer driver. |
| Restrict CD-ROM access to locally logged-on user | Determines whether a CD-ROM is accessible to both local and remote users simultaneously. If this policy is set, it allows only the interactively logged-on user to access removable CD-ROM media. If this policy is set and no one is logged on interactively, the CD-ROM can be accessed over the network. |
| Restrict floppy access to locally logged-on user only | Determines whether removable floppy media are accessible to both local and remote users simultaneously. If this policy is set, it allows only the interactively logged-on user to access removable floppy media. If this policy is set and no one is logged on interactively, the floppy can be accessed over the network. |
| Unsigned driver installation behavior | Determines what happens when an attempt is made to install a device driver (by means of Setup API) that has not been certified by the Windows Hardware Quality Lab (WHQL). |

**Device Driver WMI Classes**

The Windows Server 2003 Driver Development Kit (DDK) covers WMI classes that are associated with device drivers. For more information about the Windows Server 2003 device driver classes, see "Microsoft Windows Driver Development Kits" on the Microsoft Web site.

### WMI Classes Associated with Device Drivers

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| WMIBinaryMofResource | \root\wmi | Windows XP, Windows 2000 Professional, Windows Server 2000, Windows NT Server 4.0 SP4 and later, Windows Server 2003 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Plug and Play Technical Reference"
- "Registry Reference" in the Tools and Settings Collection

- "Microsoft Windows Driver Development Kits" on the Microsoft Web site

# Plug and Play Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Plug and Play technology provides a combination of hardware and software support that enables applications and the system to recognize and adapt to changes in hardware configuration. With Plug and Play in Microsoft Windows Server 2003, you can add devices while the system is in use and leave the job of configuring and starting the hardware to the operating system.

In order to use Plug and Play to install new devices without user intervention, the computer hardware, the device drivers, and the BIOS must be compliant with Plug and Play specifications. If any of these components are not designed to support Plug and Play, devices can still be added to a computer, but they must be manually configured to work properly.

**In this subject**


- What Is Plug and Play?
- How Plug and Play Works
- Plug and Play Tools and Settings

# What Is Plug and Play?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Plug and Play?**

**In this section**


- Common Plug and Play Scenarios
- Technologies Related to Plug and Play
- Related Information


Plug and Play technology provides a combination of software and hardware support that enables the Windows operating system to detect and configure hardware with little or no user involvement. Plug and Play makes it easier to add and configure hardware on a computer running Microsoft Windows Server 2003 without the user needing special knowledge about hardware configurations. Plug and Play in Windows Server 2003 supports a wide range of device types.

The Plug and Play framework defines several base components that can be implemented in any environment. Buses, devices, and resources are abstracted from the base components via hardware-specific enumerators, arbitrators, and device drivers. This ensures that the framework is extensible not only to other environments but to future bus and device architectures.

In Windows Server 2003, Plug and Play provides the following services.

- Detection of a Plug and Play device
- Determination of hardware resource requirements
- Determination of device identification number
- Allocation of hardware resources
- Dynamic loading, initialization, and unloading of drivers
- Notification to other drivers and applications when a new device is available
- Coordination with power management components to install and remove devices

Some bus types, such as Peripheral Component Interconnect (PCI) and USB, take full advantage of Plug and Play. Older bus types, such as ISA, do not take full advantage of Plug and Play and require more user interaction to ensure that devices are correctly installed.

After Windows Server 2003 detects a Plug and Play device, the device driver is configured and loaded dynamically, typically without requiring user input.

The following diagram shows the Plug and Play states of a device after initialization.

**Plug and Play Device States**




The following table describes the Plug and Play states for a device.

**Plug and Play Device States**

| State | Description |
|---|---|
| Start State | After it has assigned all necessary hardware resources to the device, the Plug and Play Manager sends an I/O request packet (IRP) indicating that all drivers for the device have performed their start operations. The device may have been recently enumerated and is being started for the first time, or the device may be restarting after being stopped for resource rebalancing. |
| Remove-Pending State | The Plug and Play Manager sends a requests to the drivers to remove a device's software representation (device objects, and so forth). The Plug and Play Manager sends this IRP when a device has been removed in an orderly fashion (for example, initiated by a user in the Unplug or Eject Hardware program), when a device has been removed by surprise (a user pulls the device from its slot without prior warning), or when the user makes a request to update drivers. |
| Surprised-Removed State | The Plug and Play Manager sends a request to notify drivers that a device is no longer available for I/O operations and has probably been unexpectedly removed from the machine. In response to this request, the drivers power off the device and perform additional removal processing. |
| Stop-Pending State | The Plug and Play Manager sends a request asking whether the drivers for a device can stop the device. If all the drivers in the device stack return a confirmation, the drivers put the device into a stop-pending state, from which the device can be stopped quickly. |
| Stopped State | Plug and Play Manager requests that drivers for a device driver hold IRPs that require access to the device. |

**Common Plug and Play Scenarios**

Plug and Play works automatically by default on computers running Windows Server 2003. The following are the most common scenarios in which Plug and Play is used.

**Adding new hardware to a computer**

In Windows Server 2003, Plug and Play installs most hardware devices that can be loaded at a time other than during the system startup process. If the hardware is compliant with Plug and Play, Windows Server 2003 will automatically detect and configure the new hardware once it is connected to the system. Plug and Play devices are capable of identifying themselves and declaring their services and resource requirements to the operating system. This information enables Windows Server 2003 to determine and set a working configuration, then load the appropriate device drivers without user intervention.

**Docking a mobile computer**

The Plug and Play framework allows docking of mobile computers and insertion of devices while the computers are running. When such a connection occurs, Windows Server 2003 will recognize that a device is connected, recognize the services and requirements, and load the necessary drivers. Applications are notified about dynamic events so that they can take advantage of new functionality or start using available devices.

**Technologies Related to Plug and Play**

In Windows Server 2003, Plug and Play support is optimized for computers that include an Advanced Configuration and Power Interface (ACPI) BIOS. ACPI devices are defined by the Advanced Configuration and Power Interface (ACPI) Specification, a hardware and software interface specification that combines and enhances the Plug and Play and Advanced Power Management (APM) standards. ACPI devices include low-level system devices such as batteries, thermal zones, and other devices defined in a system's *ACPI namespace,* a hierarchical namespace that an ACPI BIOS uses to reference objects. ACPI shifts many power management tasks to the operating system and provides the following features.

- Automatic and dynamic recognition of installed hardware
- Hardware resource allocation and reallocation
- Loading of appropriate drivers
- Interface for driver interaction with the Plug and Play system
- Dynamic handling of device management and power management events

### Note

- For all Plug and Play features to work, a system must include an ACPI-compliant BIOS that works with the ACPI to provide reliable power management and system configuration.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Device Driver Technical Reference
- "Microsoft Windows Driver Development Kits" on the Microsoft Web site

# How Plug and Play Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Plug and Play Works**

**In this section**


- Plug and Play Architecture
- Plug and Play Processes and Interactions
- Network Ports Used by Plug and Play
- Related Information


With Plug and Play under Microsoft Windows Server 2003, you can connect a hardware device to your system and leave the job of configuring and starting the device to the operating system. If the device and drivers are not designed to take advantage of Plug and Play, Windows Server 2003 will not be able to automatically configure and start the device. Plug and Play in Windows Server 2003 supports a wide range of devices.

In Windows Server 2003, Plug and Play support is optimized for computers that include an Advanced Configuration and Power Interface (ACPI) BIOS. ACPI devices are defined by the Advanced Configuration and Power Interface (ACPI) Specification, a hardware and software interface specification that combines and enhances the Plug and Play and Advanced Power Management (APM) standards. ACPI devices include low-level system devices such as batteries and thermal zones.

On x86-based computers, the interaction between the BIOS and Plug and Play is dependent on whether the system BIOS or the operating system configures the hardware.

The following table describes the recommended BIOS settings on x86-based systems.

**Recommended Plug and Play BIOS Settings on x86-based Systems**

| Condition | Recommended BIOS Setting |
|---|---|
| Fully compliant ACPI system (ACPI BIOS is present; ACPI Hardware Abstraction Layer (HAL) is installed.) | Windows Server 2003 assigns device resources and ignores BIOS settings. This includes assigning IRQ, DMA, and I/O resources and arbitrating conflicts for all PCI devices. |
| Non-compliant ACPI system (ACPI BIOS is present; compliance problems prevented ACPI HAL installation.) | The system BIOS assigns device resources prior to loading of the operating system, and the Plug and Play BIOS settings must be set to **No/Disabled**. If your devices have a static configuration, you must turn the computer off before removing or attaching these devices. |
| Non-ACPI systems | The system BIOS assigns device resources prior to the loading of the operating system, and the Plug and Play BIOS setting must be set to **No/Disabled**. If your devices have a static configuration, you must turn the computer off before removing or attaching these devices |
| Dual-boot Windows Server 2003, Windows Server 2003, and Microsoft Windows 98, or Windows Me operating systems | The Plug and Play BIOS setting must be set to **No/Disabled**. Disabling Plug and Play in the BIOS is the recommended method to prevent errors that may otherwise occur. For example, if the system check for Plug and Play on a Windows 98 ACPI system passes, the system check for Plug and Play might fail on a Windows Server 2003 system. |

Plug and Play detection runs with the logon process and relies on system firmware, hardware, device drivers, and operating system features to detect and enumerate new devices. ACPI firmware provides enhanced features, such as hardware resource sharing. When Plug and Play components are coordinated, Windows Server 2003 can detect new devices, allocate system resources, and install or request drivers with minimal user intervention.

**Plug and Play Architecture**

The following diagram shows the components that work together to support Plug and Play.

**Plug and Play Architecture**




The kernel-mode Plug and Play Manager maintains central control, directing bus drivers to perform enumeration and configuration and directing device drivers to add a device, start a device, and so on.

For example, Plug and Play Manager can send a request to determine whether a device can be safely paused or removed, and to give the device driver a chance to synchronize outstanding I/O requests to the incoming request. Plug and Play in kernel mode coordinates with Plug and Play in user mode to pause or remove devices that are available for such actions.

Drivers for Plug and Play devices do not assign resources; instead, the requested resources for a device are identified when the device is enumerated. When a device is enumerated, Plug and Play Manager retrieves the requirements for each device during resource allocation. Resources are not dynamically configurable for legacy devices, so Plug and Play Manager assigns resources to legacy devices first.

I/O Manager provides core services for device drivers. It is the kernel-mode component that translates user-mode read and write commands into read and write IRPs. I/O Manager manages all other main operating system IRPs.

The following table describes the components from the previous illustration, Plug and Play Architecture.

**Plug and Play Components**

| Component | Description |
|---|---|
| Hardware Abstraction Layer (HAL) | A component based on the Windows NT operating system that provides platform-specific support for the Windows NT kernel, the I/O Manager, kernel-mode debuggers, and the lowest-level device drivers. The HAL exports routines that abstract platform-specific hardware details about caches, I/O buses, interrupts, and so forth. The HAL provides an interface between the platform's hardware and the system software. For example, the HAL implements a routine to map each device driver's bus-relative device interrupt vector to a system-assigned vector with a corresponding platform-specific hardware priority (DIRQL). The HAL also provides several routines that provide DMA transfer support. |
| Windows Driver Model (WDM) | Driver files used to communicate between devices and Windows Server 2003 are structured under the Windows Device Model (WDM). The three WDM driver types are bus driver, function driver, and filter driver. WDM drivers must support Plug and Play, Power Management, and Windows Management Instrumentation (WMI). |

| | |
|---|---|
| Plug and Play Manager (Kernel) | Kernel-mode component responsible for determining the hardware resources requested by each device) and assigns resources appropriately. Reconfigures resource assignments when necessary, such as when a new device added to the system requires resources already in use.<br><br>The kernel-mode Plug and Play Manager interacts with the user-mode Plug and Play Manager by receiving control requests and notifying the user-mode Plug and Play Manager that a new device is present on the system that needs to be started. |
| I/O Manager | Subsystem that controls and interacts with all devices on the system and provides routines that drivers can call to have the I/O Manager insert IRPs into the associated device queue. |
| Power Manager | Component responsible for managing power usage for the system. Power Manager administers the system-wide power policy and tracks the path of power IRPs through the system. |
| Executive Components | The collection of kernel-mode operating system components. Executive components include executive support, the kernel, Memory Manager, Cache Manager, process structure, interprocess communication (LPC and RPC), Object Manager, I/O Manager, Configuration Manager, and Security Reference Monitor. Each executive component, with the exception of Cache Manager, supplies a specialized set of native system services. Every executive component also exports a set of kernel-mode support routines for use by other executive components. |
| Plug and Play Manager (User) | User-mode component that interacts with user-mode setup components, such as Class Installers, to configure and install devices, as well as interacting with applications to register notification of device changes or to notify the application when a device event occurs.<br><br>The user-mode Plug and Play Manager interacts with the kernel-mode Plug and Play Manager by sending control requests. |
| Win32 Services | Windows Server 2003 functionality that uses the Win32 API to manage resources on a computer. |

**Plug and Play Processes and Interactions**
The following section describes Plug and Play processes and interactions.

**How Plug and Play Detects and Installs a Device**
The following diagram shows how Plug and Play detects and installs a device.

**Plug and Play Driver Installation Process**




When a hardware device is connected — as when you plug a USB camera into a USB port — Plug and Play Manager goes through the following steps to install the device.

1. After receiving an insertion notification, Plug and Play Manager checks what hardware resources the device needs —

such as memory ranges, I/O ranges, and DMA channels. Plug and Play Manager then assigns those resources.

2. Plug and Play Manager checks the hardware identification number of the device. Plug and Play Manager then checks the hard drive, floppy drives, CD-ROM drives, and Windows Update for drivers that match the hardware identification number of the device.

3. If multiple drivers are found, Plug and Play Manager chooses the driver that is the best match by looking for the closest hardware ID or compatible ID match, driver signatures, and other driver features.

4. Plug and Play Manager then installs the best-match driver and the operating system starts the device.

When multiple drivers are available for a given device, Windows Server 2003 uses driver-ranking schemes to determine the best driver to load. Driver rank is established based on whether the driver is signed and how closely the driver's Plug and Play ID matches the device's Plug and Play ID.

If no driver is available on the computer, and the computer is connected to the Internet, then the operating system will use the Windows Update service to search for available drivers on the Windows Update site. If an appropriate driver is found, Windows Server 2003 copies it and installs it on the computer.

Device drivers included with or installed under Windows Server 2003 must meet the standards of the Windows Logo program. Device drivers that have passed the Windows Hardware Quality Lab (WHQL) compatibility tests are electronically signed, and Windows Server 2003 detects the digital signature. For system stability, it is recommended that you use only signed device drivers with Windows Server 2003. A message notifies the user if an unsigned driver is being installed.

The extent of Plug and Play support depends on both the hardware device and the device driver. If a driver does not support Plug and Play, its devices behave as devices that are not Plug and Play. This might result in the loss of some operating system functionality. For example, power management features such as hibernation might not work. Some bus types, such as Peripheral Component Interconnect (PCI) and Universal Serial Bus (USB) take full advantage of Plug and Play. Older bus types, such as Industry Standard Architecture (ISA), do not and require more user interaction to ensure that devices are correctly installed.

For more information about device ranking and device signing, see "Device Driver Technical Reference."

**Plug and Play in Managed Environments**
When you install a Plug and Play device and are connected to the Internet, Windows Server 2003 automatically accesses Windows Update to search for a device driver.

In order to install devices using the hardware wizards, you must be logged on as an administrator or a member of the Administrators group. You then use hardware wizards to search the Windows Update site for device drivers.

You can control how administrators update and install hardware devices by using Group Policy. For example, an IT administrator in a highly managed network environment can control the ability of administrators to install new hardware and to access the Internet automatically when Windows Server 2003 searches for device drivers. You can set a policy to disallow any access to Windows Update, preventing certain administrators from automatically accessing Windows Update. You can also set a policy for downloading the updates from the Windows Update Catalog and distributing through a corporate network as needed.

In a managed environment, you can control how Plug and Play operates by changing any or all of the following, including downloading updates from Windows Update:

**Specific information sent or received**
The Code Download Manager (CDM) calls Windows Update to find and download device drivers. The CDM also calls Help and Support Center, which logs Plug and Play IDs for devices for which Microsoft does not have drivers. Neither of these communications is under the direct control of Plug and Play. The CDM handles all of the communication between the computer and Windows Update. None of this communication between the computer and the Internet uniquely identifies the user.

**Default and recommended settings**
Plug and Play is enabled by default.

**Triggers**
A user controls whether to run Windows Update. If Automatic Updates is enabled following setup, it is triggered about once per day when there is an Internet connection.

**User notification**
When searching for a device driver, Windows Update sends a list of available drivers to the user's computer.

**Logging**
If you use a Plug and Play driver with a device that is not Plug and Play, any associated issues or problems are recorded in the event log.

**Encryption**
Data transfer is based on interaction with Windows Update. The data is transferred using HTTPS.

**Transmission protocol and ports**
The protocol for non-secure transmission is HTTP and the port used is 80; the protocol for secure transmission is HTTPS and the port used is 443.

**Ability to disable**
Plug and Play cannot be disabled, since system instability would result. You can disable access to Windows Update using Group Policy.

**Plug and Play Services**
Plug and Play in Windows Server 2003 provides the following services.

- Detection of a Plug and Play device
- Determination of hardware resource requirements
- Determination of device identification number
- Allocation of hardware resources
- Dynamic loading, initialization, and unloading of drivers
- Notification to other drivers and applications when a new device is available
- Coordination with power management components to install and remove devices

**How Plug and Play Detection Occurs on ACPI Systems**

1. The function driver for the bus detects a new device on the bus.
2. The bus driver notifies Plug and Play Manager that a set of devices has changed.
3. Plug and Play Manager queries the driver for the current list of devices on the bus.
4. When Plug and Play Manager obtains the current list of devices, it determines whether any devices have been added or removed.
5. Plug and Play Manager gathers information about the new device and begins configuring it.
6. Plug and Play Manager checks the registry to determine whether the device has been installed on the computer before; if not, Plug and Play Manager creates a registry key and any associated subkeys about the device in the registry.
7. Plug and Play Manager attempts to find and load the function drivers and filter drivers for the device.
8. Plug and Play Manager assigns resources to the device if needed and issues an I/O request packet (IRP) to start the device.

**Plug and Play Device States**
The following diagram shows how a device transitions through Plug and Play states as it is configured, started, and stopped.

### Plug and Play Device States




The process illustrated in the previous diagram begins when a device is first recognized as physically present. This is indicated in the diagram above by the box in the upper left labeled "Device is physically present". Before the a Plug and Play device is recognized as physically present, the device has not yet been identified by the system. This could be because the device has just been inserted into the system; however, the device could have been present at boot time and was not detected then for some reason.

Plug and Play Manager begins configuration of the device by enumerating it, then identifying and initializing the drivers for the device, and then creating and attaching the function driver objects and any optional filter driver objects.

Plug and Play Manager calls a routine to load each driver if the driver is not yet loaded. Once the driver is initialized, it must be ready to initialize the device. When a device receives a request from Plug and Play Manager, the driver starts the device and is ready to process I/O requests for the device.

Upon any changes in the device status, such as the device being physically removed from the system, Plug and Play Manager will reconfigure the hardware resources of the device. If a device is reconfigured, Plug and Play Manager directs the drivers to restart the device.

**Removing Devices While the System is Running**

Some Plug and Play devices can be installed or removed while the system is running. For example, USB, IEEE 1394, and PC Card devices can be added to and removed from a fully powered system. When such hardware is added or removed, the operating system automatically detects insertion or removal of the device and manages system and device configuration as required. If the device is not designed to be removed while the system is running, it is recommended that you notify the operating system in advance to avoid problems. The Safely Remove Hardware application notifies the operating system that a device will be removed.

The following table shows the different types of Plug and Play devices and whether or not they can be removed while the system is turned on.

**Plug and Play Device Connections and Installation Guidelines**

| Devices on These Buses or Connectors | Can Be Added to or Removed from a Running System? | System Must be Turned Off Before Device is Added to or Removed from System? |
|---|---|---|
| USB, IEEE 1394, PC Card devices, CardBus devices | Yes. Remove hardware by using the Safely Remove Hardware application if it appears in the notification area. | No. |
| PCI, ISA, EISA | No. | Yes. |
| Docking station | Varies among computer manufacturers; most support docking and undocking while the computer is running. | Varies among computer manufacturers; most support docking and undocking while the computer is running. |

Even when removing a device from a bus that supports the removal of devices while the system is running, it is best to use the Safely Remove Hardware application to ensure a safe removal of hardware from the system. The Safely Remove Hardware application informs Windows that the user intends to remove a device. This gives Windows an opportunity to prepare for the removal by taking steps such as halting data transfers to the device and unloading device drivers. If a device supports safe removal, a Safe Removal icon displays in the notification area next to the task bar. The icon displays a list of the Plug and Play

devices supporting safe removal currently attached to the system. If you do not see the icon, the device does not support this feature, and you cannot unplug or eject your device using Safe Removal.

When hardware is removed from a running system without using the Safely Remove Hardware application to notify the operating system in advance of the removal, it is referred to as surprise removal. Surprise removal is particularly a concern for storage devices for which write caching is enabled; when such devices are surprise-removed, data loss or data corruption might occur. To reduce the likelihood of data loss or data corruption due to surprise removal of consumer-oriented storage devices, Windows Server 2003 disables write caching by default for these devices (such as cameras that include IEEE 1394 or USB storage and small-form-factor storage devices such as compact flash). While write caching policy addresses this particular issue, it is recommended that you continue to use the Safely Remove Hardware application when it appears in the notification area. Also, disabling write caching might slow the performance of consumer storage devices.

Write caching is enabled by default for high-performance external storage devices such as IEEE 1394 hard drives and SCSI hard drives and for storage devices inside the computer that cannot be surprise-removed.

**How Plug and Play Interoperates with an ACPI BIOS**
In Windows Server 2003, Plug and Play support is optimized for computers that include an ACPI BIOS. ACPI devices are defined by the ACPI Specification, a hardware and software interface specification that combines and enhances the Plug and Play and Advanced Power Management (APM) standards. ACPI devices include low-level system devices such as batteries, thermal zones, and other devices defined in a system's *ACPI namespace,* a hierarchical namespace that an ACPI BIOS uses to reference objects. The ACPI BIOS is responsible for tasks such as describing hardware that is not visible to Plug and Play because the hardware is connected to a bus that does not directly support Plug and Play. For example, the ACPI BIOS describes and helps in configuring devices such as system timers and programmable interrupt controllers on the motherboard, which is not on a bus that supports Plug and Play.

For all Plug and Play features to work on a given system, it must include an ACPI BIOS and hardware devices and drivers that are compliant with Plug and Play. An Advanced Power Management (APM) BIOS or a Plug and Play BIOS does not enable all Plug and Play features and is not as robust as an ACPI BIOS.

**Network Ports Used by Plug and Play**
The following table describes the port used by Plug and Play when searching Windows Update for the most current device drivers.

### Port Assignments for Plug and Play

| Service Name | UDP | TCP |
|---|---|---|
| HTTP, IIS | None | 80 |
| HTTPS | None | 443 |

**Related Information**
The following resources contain additional information that is relevant to this section.

- Device Driver Technical Reference
- “Microsoft Windows Driver Development Kits” on the Microsoft Web site

# Plug and Play Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Plug and Play Tools and Settings**
**In this section**


- Plug and Play Tools
- Plug and Play Registry Entries
- Plug and Play Group Policy Settings
- Plug and Play WMI Classes
- Network Ports Used by Plug and Play
- Related Information


The following tools, registry entries, Group Policy settings, network ports, and Windows Management Instrumentation (WMI) classes are associated with Plug and Play.

**Plug and Play Tools**
The following tools are associated with Plug and Play.

**Devcon.exe: Device Console**
**Category**
Devcon ships with the Windows Driver Development Kit.

**Version compatibility**
Windows 2000, Windows XP, and Windows Server 2003

Device Console (Devcon) is a command-line tool that displays detailed information about devices and lets you search for and manipulate devices from the command line. Devcon enables, disables, installs, configures, and removes devices on the local computer and displays detailed information about devices on local and remote computers.

For more information about Devcon.exe and the syntax used to carry out commands, see the Windows Driver Development Kit.

**GenINF.exe: INF File Creation Wizard**
**Category**
GenINF ships with the Windows Driver Development Kit.

**Version compatibility**
GenINF generates .inf files for Windows 2000, Windows XP, and Windows Server 2003.

GenINF is a GUI-based tool that guides users through the process of creating device driver .inf files. GenINF does not generate fully valid .inf files. However, the skeleton .inf files it generates should help developers complete their device driver .inf files.

GenINF is located in the \Tools subdirectory of the DDK installation directory. Because GenINF is GUI-based, it can be executed from within the Windows Server 2003 User Interface or from the command prompt. GenINF takes no arguments. GenINF supports all device classes, but includes specific support for only certain device classes. For a list of these classes, see the Geninf.htm file that is located in the Tools\geninf subdirectory of the DDK.

For more information about GenINF.exe see the Windows Driver Development Kit.

**Net.exe: Net Start or Net Stop**
**Category**
Command-line tool that ships with Windows Server 2003

**Version compatibility**
Windows XP, Windows 200, and Windows Server 2003

Net.exe allows you to start and stop the Plug and Play service.

For more information about the syntax used with Net.exe, see the Windows Server 2003 help.

**SigVerif.exe: File Signature Verification**
**Category**
Operating system tool

**Version compatibility**
Windows XP, Windows 200, and Windows Server 2003

File Signature Verification checks to see which system files and device driver files are digitally signed and displays its findings. If you have enabled logging, the search results are also written to a log file.

**Svchost.exe: Service Host Process**
**Category**
Operating system tool

**Version compatibility**
Windows 2000, Windows XP, and Windows Server 2003

Generic host process name for services that are run from dynamic-link libraries (DLLs). At startup, Svchost.exe checks the services portion of the registry to construct a list of services that it needs to load. There can be multiple instances of Svchost.exe running at the same time. Each Svchost.exe session can contain a grouping of services, so that separate services can be run depending on how and where Svchost.exe is started.

**Plug and Play Registry Entries**
The following registry keys are associated with Plug and Play.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

For more information about registry entries, see "Registry Reference" in the Tools and Settings Collection.

**PnPManager**
**Registry path**
HKEY_LOCAL_MACHINE\HARDWARE\RESOURCEMAP\PnPManager

**Version**
Windows XP, Windows Server 2003

This key contains registry entries for each device installed on the system. Each of the subkeys and entries for this key are dependent upon the specific device installed on a system. Because of this, these subkeys and entries are not listed here.

**PlugPlay**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services

**Version**
Windows Server 2003

The Plug and Play registry key enables a computer to recognize and adapt to hardware changes with little or no user input. Stopping or disabling this service will result in system instability.

**Description**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay

**Version**
Windows XP, Windows 2000, Windows Server 2003

Text field that displays the description of the Plug and Play service.

**DisplayName**
**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay

**Version**
Windows XP, Windows 2000, Windows Server 2003

Text field that displays the default service name, Plug and Play.

**ErrorControl**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay

**Version**
Windows XP, Windows 2000, Windows Server 2003

Tells the system how to boot if a Plug and Play driver fails. The following table lists the accepted values for this registry entry.

### ErrorControl Values

| Value | Description |
|---|---|
| 00 | Ignore. If the driver fails to load or initialize, then startup continues as normal and no warning message appears. |
| 01 | Normal. If the driver fails to load or initialize, then startup proceeds, but a warning message appears. |
| 02 | Severe. If the driver fails to load or initialize, then startup is declared to have failed and restarts by using the LastKnownGood control set. If startup is already using the LastKnownGood control set, then startup continues |
| 03 | Critical. If the driver fails to load or initialize, then startup is declared to have failed and restarts by using the LastKnownGood control set. If startup is already using the LastKnownGood control set, then startup stops and enters Debugging Mode. |

**Group**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay

**Version**
Windows XP, Windows 2000, Windows Server 2003

Specifies which group the service belongs to when Plug and Play services are started

**ImagePath**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay

**Version**
Windows XP, Windows 2000, Windows Server 2003

Location of Services Control Manager, which uses the information located in the subkeys to initialize Plug and Play services. The default path is %SystemRoot%\system32\services.exe.

**ObjectName**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay

**Version**
Windows XP, Windows 2000, Windows Server 2003

Name of the account under which the Plug and Play service runs. The default value is LocalSystem, which sets the greatest privileges. In order to run as any other user, the service must provide a password for the user. If the service is a kernel driver, this specifies which kernel object loads it.

**Start**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay

**Version**
Windows XP, Windows 2000, Windows Server 2003

Determines when Plug and Play services should be started on a system. The following table lists the accepted values for this

registry entry.

**Start Values**

| Value | Description |
|---|---|
| 00 | Boot. Starts the service during the boot sequence. |
| 01 | System. Starts the service when loaded by the I/O subsystem in the kernel. |
| 02 | Automatic. Starts the service automatically. |
| 03 | Manual. Starts the service when a user initiates it from the control panel or a command line. |
| 04 | Disable. Loads the service, but neither the system nor a user can start it. |

**Type**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay

**Version**
Windows XP, Windows 2000, Windows Server 2003

Identifies the type of service. The following table lists the accepted values for this registry entry.

**Type Values**

| Value | Description |
|---|---|
| 01 | Kernel-mode device driver |
| 02 | File system driver |
| 04 | Network adapter |
| 10 | Win32 service that can run as a standalone process |
| 20 | Win32 service that shares address space with other services at the same time |

**Security**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay\Security.

**Security**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PlugPlay

**Version**
Windows Server 2003

Holds the account policy information for the workstation or member server, as well as the security rights users have with the local Security Accounts Manager (SAM).

**Plug and Play Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with Plug and Play.

**Group Policy Settings Associated with Plug and Play**

| Group Policy Setting | Description |
|---|---|

| | |
|---|---|
| Remove Computer from Docking Station | Determines whether a user can undock a portable computer from its docking station without logging on. If this policy is set, the user must log on before removing the portable computer from its docking station. If this policy is not set, the user can remove the portable computer from its docking station without logging on. |
| Allow undock without having to log on | Determines whether a user must log on to request that a portable computer be removed from a docking station. If this setting is not set, the user must log on to request an undock, at which time the user must have been granted the Remove Computer from Docking Station privilege. If this setting is set, a user can request an undock by pressing the portable computer's physical eject button. |
| Allowed to format and eject removable media | Determines who is allowed to format and eject removable NTFS media. This capability can be given to Administrators, Administrators and Power Users, or Administrators and Interactive Users. |
| Prevent users from installing printer drivers | Determines who is allowed to install a printer driver as part of adding a network printer. For a computer to print to a network printer, the driver for that network printer must be installed on the local computer. If this setting is set, only Administrators and Power Users can install a printer driver as part of adding a network printer. If this setting is not set, any user can install a printer driver as part of adding a network printer. This setting can be used to prevent unprivileged users from downloading and installing an untrusted printer driver. |
| Restrict CD-ROM access to locally logged-on user | Determines whether a CD-ROM is accessible to both local and remote users simultaneously. If set, this policy allows only the interactively logged-on user to access removable CD-ROM media. If no one is logged on interactively, the CD-ROM can be shared over the network. If this policy is not set, then the local user and remote users can access the CD-ROM simultaneously. |
| Restrict floppy access to locally logged-on user only | Determines whether removable floppy media are accessible to both local and remote users simultaneously. If set, this policy allows only the interactively logged-on user to access removable floppy media. If no one is logged on interactively, the floppy media can be shared over the network. If this policy is not set, then the local user and remote users can access the floppy media simultaneously. |
| Unsigned driver installation behavior | Determines what happens when an attempt is made to install a device driver (by means of Setup API) that has not been certified by the Windows Hardware Quality Lab (WHQL). |

For more information about Group Policy settings, see "Group Policy Settings Reference" in the Tools and Settings Collection.

**Plug and Play WMI Classes**

The following table lists and describes the WMI (Windows Management Instrumentation) classes that are associated with Plug and Play.

### WMI Classes Associated with Plug and Play

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| WMIBinaryMofResource | \root\wmi\ | Windows XP, Windows 2000 Professional, Windows Server 2000, Windows NT Server 4.0 Service Pack 4 and later, Windows Server 2003 |
| Win32_DeviceSettings | \root\cimv2 | Windows XP, Windows 2000 Professional, Windows Server 2000, Windows NT Server 4.0 Service Pack 4 and later, Windows Server 2003 |
| Win32_PnP | \root\cimv2 | Windows XP, Windows 2000 Professional, Windows Server 2000, Windows NT Server 4.0 Service Pack 4 and later, Windows Server 2003 |

| | | |
|---|---|---|
| Win32_PnPAllocatedResource | \root\cimv2 | Windows XP, Windows 2000 Professional, Windows Server 2000, Windows NT Server 4.0 Service Pack 4 and later, Windows Server 2003 |
| Win32_PnPDevice | \root\cimv2 | Windows XP, Windows 2000 Professional, Windows Server 2000, Windows NT Server 4.0 Service Pack 4 and later, Windows Server 2003 |
| Win32_PnPEntity | \root\cimv2 | Windows XP, Windows 2000 Professional, Windows Server 2000, Windows NT Server 4.0 Service Pack 4 and later, Windows Server 2003 |
| Win32_PnPSignedDriver | \root\cimv2 | Windows XP, Windows 2000 Professional, Windows Server 2000, Windows NT Server 4.0 Service Pack 4 and later, Windows Server 2003 |
| Win32_PnPSignedDriverCIMDataFile | \root\cimv2 | Windows XP, Windows 2000 Professional, Windows Server 2000, Windows NT Server 4.0 Service Pack 4 and later, Windows Server 2003 |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by Plug and Play**

The following table describes the port used by Plug and Play when searching Windows Update for the most current device drivers.

### Port Assignments for Plug and Play

| Service Name | UDP | TCP |
|---|---|---|
| HTTP, IIS | None | 80 |
| HTTPS | None | 443 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Device Driver Technical Reference
- "Microsoft Windows Driver Development Kits" on the Microsoft Web site
- "Registry Reference" in the Tools and Settings Collection

# Installation and Setup Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Microsoft Windows Server 2003 operating system provides several installation and setup technologies. You can use these technologies to perform an interactive setup, which prompts the user for configuration information, or an automated installation, which requires no user input.

Interactive setup technologies are used to install or upgrade an operating system on a single computer. Interactive setup technologies are typically used by small organizations that have few servers, or medium and large organizations that need to install the operating system on servers that have one-of-a-kind configuration settings.

Automated installation technologies are used to install or upgrade operating systems on numerous computers. Automated installation technologies include Remote Installation Services (RIS), the System Preparation (Sysprep) tool, and unattended installation answer files. Automated installation technologies are typically used by medium and large organizations during large-scale rollouts when it would be too slow and costly to have administrators or technicians interactively install the operating system on individual computers.

**Installation and Setup Components**
The installation and setup technologies include the following components.

**Interactive Setup**
Interactive setup relies on the Setup program in Windows Server 2003 and provides a method for carrying out an interactive installation or upgrade on a single computer. Interactive setup can be started by running Setup.exe, but also by running Winnt32.exe or Winnt.exe, both of which offer command-line options. The Setup program is also the basis for the automated installation technologies, including RIS, Sysprep, and unattended installation.

Running Setup interactively is a simple way to accomplish a server installation or upgrade. The installation or upgrade can be carried out from the CD drive or across the network (after you copy the installation files that are on the product CD to a shared folder). In addition, interactive setup requires little preparatory work compared to automated installation, and does not require supplemental tools or programs. Also, interactive setup can be used to install an operating system on a computer that is not connected to a network.

**Unattended Installation**
Unattended installation is an automated installation technology in Windows Server 2003 that you can use to install or upgrade an operating system with minimal user intervention. Unattended installation relies on the Setup program and an answer file to automate the setup process. An answer file is a simple text file that provides details about how to install and configure the operating system. To start an unattended installation, you run Winnt32.exe or Winnt.exe with a command-line parameter that specifies the answer file to use. Then, during the installation process, necessary information is obtained from instructions contained in the answer file instead of your responses to prompts.

**Sysprep**
Sysprep is an automated installation technology that you can use to install Windows Server 2003 and Microsoft Windows XP operating systems. The Sysprep tool is typically used in conjunction with a non-Microsoft disk imaging tool or Microsoft Windows Server 2003 Automated Deployment Services (ADS) to perform image-based installations. Image-based installation is a method of copying or cloning preconfigured operating systems (and, optionally, software applications) onto destination computers. After you set up a master installation — an installation with the operating system, software applications, and configuration settings that you want to install onto the destination computers in your organization, Sysprep prepares the master installation so that you can create a disk image; that is, a functionally identical replica of the disk containing the master installation, that can be copied onto multiple computers. Next, you use a disk-imaging program to create the disk image of the master installation. After you copy the disk image onto a destination computer and start the destination computer, a shortened version of the Setup program runs. The shortened version of Setup configures only user-specific and computer-specific settings, such as computer name, domain membership, and regional options. You can automate this last part of the setup process by using an answer file, a simple text file that instructs the Setup program how to configure the various operating system settings.

**Remote Installation Services**
RIS is an automated installation technology that can be used to install Windows Server 2003, Windows XP, and Microsoft Windows 2000 operating systems. RIS uses Pre-Boot Execution Environment (PXE) technology to enable client computers without an operating system to start up and connect remotely to a RIS server, which installs a supported operating system.

You can use RIS to perform either CD-based installations or image-based installations. To perform CD-based installations you

use the Remote Installation Services Setup (Risetup.exe) tool, which installs an operating system on a destination computer by using an answer file and the installation files that are on the product CD. To perform image-based installations you use the Remote Installation Preparation Wizard (Riprep.exe) tool, which is similar to the Sysprep tool and prepares a master installation for disk imaging.

**Installation and Setup Scenarios**
Although you can use the installation and setup technologies to create a wide range of interactive and automated deployment solutions, each technology has certain requirements and limitations that make it better suited to a specific deployment solution. The following describes optimum scenarios for each of the installation and setup technologies.

**Interactive Setup Scenarios**
Interactive setup is the most versatile and robust setup and installation technology. Interactive setup does not require network connectivity and it does not require any supplemental tools or programs. Because interactive setup has no limitations or restrictions, you can use it to install an operating system on any type of server or create any type of server configuration. However, the following installation scenarios are particularly well-suited for interactive setup.

**Installation in an organization that requires only a small number of servers**
Interactive setup is useful when you are installing a small number of servers, because you do not need to do any preparatory work related to creating answer files or images.

**One-of-a-kind server configuration**
Interactive setup is useful if you are installing an operating system on a server that requires unique configuration settings or contains unique hardware. For example, you might use interactive setup to create a file server that contains special redundant array of independent disks (RAID) hardware and therefore requires you to install special device drivers, or a bastion server or proxy server that requires special security settings.

**Master installations for image-based deployment technologies**
You can use interactive setup to create master installations for image-based installation technologies, such as image-based installation with Sysprep, or RIS when it is being used for image-based installation. For example, to use Sysprep, you need to create a master installation from which you create a disk image. You can create the master installation manually by running Setup.exe, and then manually configuring operating system settings and installing software.

**Upgrades**
You can use interactive setup to upgrade a Windows server operating system to Microsoft Windows Server 2003, Standard Edition; Microsoft Windows Server 2003, Enterprise Edition; or Microsoft Windows Server 2003, Web Edition operating systems. An upgrade is useful if you want to retain existing server settings (for example, user rights, groups, and permissions settings) and you do not want to reinstall applications and files.

**Unattended Installation Scenarios**
Unattended installation is the most versatile and robust automated installation technology, and can be used to roll out any type of server or any server configuration. However, the following installation scenarios are particularly well suited for unattended installation.

**Heterogeneous hardware**
Unattended installation is ideal if you are installing the same server configuration on computers that have dissimilar or heterogeneous hardware. In addition, unattended installation is equally adept at handling other hardware differences, including differences among network adapters, media devices (CD drives), and video adapter cards. In large part, this is because unattended installation is just an automated process that uses the standard Setup program. Therefore, whatever device Setup can detect or install, unattended installation can also detect or install.

**Upgrades**
You can use unattended installation to automate the upgrade of a Windows server operating system to Microsoft Windows Server 2003, Standard Edition; Microsoft Windows Server 2003, Enterprise Edition; or Microsoft Windows Server 2003, Web Edition operating systems. No other automated installation technology can perform upgrades. Unattended upgrade installations are useful if you want to retain existing server settings (for example, user rights, groups, and permissions settings) and you do not want to reinstall applications and files.

**Special server configurations**
Unattended installation is an ideal automated installation technology for deploying domain controllers, and member servers that run Certificate Services and Cluster service. Although you can use other automated installation technologies to deploy these types of servers, those technologies require you to use scripts or supplemental answer files after the automated installation is finished. With unattended installation, the answer file contains all of the necessary settings for installing these types of servers.

**Master installations for image-based deployment technologies**
You can use unattended installation to create master installations for image-based installation technologies, such as the

Sysprep tool, or RIS when it is being used for image-based installation. For example, to use Sysprep, you need to create a master installation from which you create a disk image. You can create the master installation by performing an unattended installation on the master computer.

**Sysprep Scenarios**

Image-based installation with Sysprep can only be used to perform clean installations of an operating system. Image-based installation with Sysprep also requires a disk imaging program and usually requires high-speed network connectivity among the destination computers you are deploying. Given these requirements, several installation scenarios are particularly well suited for image-based installations with Sysprep.

**Homogeneous hardware configurations**

If you are installing the same server configuration on computers that have similar or homogeneous hardware, you can customize hardware configurations on the disk image, thereby avoiding many post-installation configuration and installation tasks. For example, you can use image-based installation with Sysprep to rapidly roll out the same file server configuration to a group of computers that have similar processor configurations and similar storage configurations.

**Rapid rollouts**

If you need to install the operating system as rapidly as possible, you can copy a preconfigured disk image onto a destination computer much faster than you can perform an unattended installation or a Remote Installation Services (RIS) installation. For example, you can use image-based installation with Sysprep to quickly deploy critical servers in your organization or quickly reinstall an operating system on member servers after a hard disk failure or some other catastrophic event.

**Application deployment**

If you want to roll out applications at the same time that you roll out the operating system, you can include the applications on the disk image. For example, if your organization's standard computer configuration includes an antivirus program, an e-mail program, and an office suite, you can install and configure these programs on the master installation so that the disk image includes these programs.

**RIS Scenarios**

Like Sysprep, RIS can be used only to perform clean installations of an operating system. RIS also requires PXE-enabled network adapters and high-speed network connectivity among the destination computers you are deploying. In addition, RIS can be used as a CD-based or image-based deployment solution. Given these requirements and options, several installation scenarios are particularly well suited for RIS installation.

**Minimize user interaction**

You can configure RIS so that minimal user interaction is necessary to configure and install an operating system. In this scenario users can initiate a RIS installation on a destination computer by pressing the F12 key. After pressing the F12 key, all configuration and installation tasks are automated and cannot be altered.

**Application deployment**

Because you can use RIS as an image-based deployment solution, you can deploy preinstalled and preconfigured applications with the operating system. For example, if your organization's standard computer configuration includes an antivirus program, an e-mail program, and an office suite, you can install and configure these programs on the master installation so that the disk image includes these programs.

# Setup Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setup Technical Reference**

The Setup program in the Microsoft Windows Server 2003 operating systems provides a method for carrying out an interactive installation or upgrade on a single computer. The Setup program is also the basis for the following installation or upgrade methods that are not interactive: unattended installation, Remote Installation Services (RIS), and image-based installation with the System Preparation tool (Sysprep).

Interactive setup involves a simple repeating process in which Setup alternates between prompting for information and performing appropriate tasks, such as copying files, saving settings, and so on. In contrast, for installation methods that are not interactive, the input does not come from a person responding to prompts, but from another source, such as an answer file (for unattended installation or RIS) or an operating system image (for Sysprep or RIS).

**In this subject**


- What Is Setup?
- How Setup Works
- Setup Tools and Settings

# What Is Setup?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Setup?**

**In this section**

- Related Information

The Setup program in Windows Server 2003 provides a method for carrying out an interactive installation or upgrade on a single computer. The Setup program is also the basis for the following installation or upgrade methods that are not interactive: unattended installation, Remote Installation Services (RIS), and image-based installation with the System Preparation tool (Sysprep).

Running Setup interactively is a simple way to accomplish a server installation or upgrade. The installation or upgrade can be carried out from the CD drive or across the network (after you share the installation files that are on the product CD). No preparatory work related to creating answer files or images is necessary. In contrast, installation methods that are not interactive require preparatory work, but are more efficient in the context of a large-scale rollout. Installation methods that are not interactive are also useful for achieving consistency and precision in the configuration on each server.

**How Interactive Setup Differs from Installation Methods That Are Not Interactive**

Interactive Setup involves a simple repeating process in which Setup alternates between prompting for information and performing the appropriate tasks, such as copying files, saving settings, and so on. The following figure shows how the process repeats until the installation or upgrade is complete.

### Interactive Installation or Upgrade




In contrast, for installation methods that are not interactive, the input does not come from a person responding to prompts, but from another source, such as an answer file (for unattended installation or RIS) or an operating system image (for Sysprep or RIS).

**Related Information**

The following resources contain additional information that is relevant to this section.

- Remote Installation Services Technical Reference
- Sysprep Technical Reference
- Unattended Installation Technical Reference

# How Setup Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setup Architecture**

The Setup program contains components that start or restart the computer, gather information, copy files, and create or adjust settings. The components involved and the sequence of actions they perform are slightly different when Setup is started in different ways. The following figures and tables show the components and briefly describe their relationships for the two most common ways of performing an interactive setup:

- An upgrade when you insert the product CD while the computer is running the existing operating system and then wait for a display to appear.
- A clean installation when you insert the product CD and then restart the computer from the CD drive.

### Note

- Multiple methods are available for running Setup. This Technical Reference focuses on two methods, but it is not intended to provide guidance for choosing a method. For information about choosing a method, see "Related Information" later in the section.

**Setup Components Used During an Upgrade**

The following figure and corresponding table show the components that are involved when you perform an upgrade by starting the computer from the existing operating system, inserting the product CD in the CD drive, and then waiting for a display to appear.

The figure and table list both the components and the related information files for each stage of the upgrade. Setup components use information files to store information that is later retrieved. Note that information might be retrieved by a component different from the one that stored it. Information might also be retrieved after restarting the computer.

For more information about the operating systems from which you can start an upgrade, see "Setup Processes" later in this section.

### Setup Components Used During an Upgrade




The following table describes the Setup components used for each stage of an upgrade.

### Setup Components Used During an Upgrade

| Components and Descriptions Grouped by Stage | Corresponding Information Files |
|---|---|
| **Initial stage, also known as data-collection stage** | |
| **Setup.exe (autorun)**<br><br>Setup.exe runs automatically when you insert the product CD in the CD drive on a computer that is turned on and running an operating system that can be upgraded to Windows Server 2003. When this happens, Setup.exe offers options for viewing information about the product CD or for beginning the setup process. | None |
| **Winnt32.exe** and **Winnt32u.dll**<br><br>If Setup.exe is allowed to start as described in the previous line of this table, and you choose to begin the setup process, Setup.exe runs Winnt32.exe. | **For the compatibility report**: Ntcompat.inf in the \I386\compdata folder on the product CD<br><br>**For other actions**: Winnt.sif in $win_nt$.~bt folder<br><br>*.inf files in \I386\Winntupg folder and subfolders on product CD |
| **Intermediate stage, also known as text-mode stage** | |
| **$LDR$ file (copied and renamed from Setupldr.bin in \I386 folder on product CD)**<br><br>As part of the previous stage, Winnt32.exe causes the computer to restart with the Setup loader, instead of NTLDR, the loader that would otherwise be used when starting a server running Windows. | Winnt.sif in $win_nt$.~bt folder<br><br>Txtsetup.sif in $win_nt$.~bt folder and in root directory (of the boot disk on the computer)<br><br>Migrate.inf in $win_nt$.~bt folder (stores drive letter assignments)<br><br>For some computers that have a mass storage controller, Unsupdrv.inf will be created. For more information, see the paragraphs after this table. |
| **Setupdd.sys**<br><br>Setupdd.sys carries out most of the actions during text-mode stage. | Winnt.sif in $win_nt$.~bt folder<br><br>Txtsetup.sif in $win_nt$.~bt folder and in root directory (of the boot disk on the computer) |
| **Final stage, also known as GUI-mode stage** | |

| | |
|---|---|
| **Syssetup.dll, Setupapi.dll, Netshell.dll, and other DLLs**<br><br>For the final stage of an upgrade, Syssetup.dll coordinates most of the actions. The operating system running during this stage is very similar to the completely upgraded operating system, although only a small number of services are started. Syssetup.dll uses a variety of other DLLs for configuration tasks, for example, iesetup.dll for configuring Internet Explorer and Ocgen.dll for configuring certain operating system components such as Windows Management Instrumentation (WMI). | Syssetup.dll uses:<br><br>$Winnt$.inf in *systemroot*\system32<br><br>Syssetup.inf and Layout.inf in *systemroot*\inf<br><br>At this stage, the operating system is very similar to the completely upgraded operating system and uses many of the same .inf files (not listed here). |

The Unsupdrv.inf file mentioned in the preceding table is created to store settings for a mass storage driver in certain situations. If a mass storage driver is installed on the computer, Setup will copy the driver into the $win_nt$.~bt folder unless the product CD has an equivalent, updated driver or the existing driver is on a list of those that must be disabled for Windows Server 2003. When Setup copies the driver, it also adds the driver file to its list of files and stores the driver settings retrieved from the existing registry in Unsupdrv.inf.

When the computer first restarts after interactive setup, a few final user settings are applied, much like the process that happens at the beginning of any restart. The upgrade is then complete.

**Setup Components Used During a Clean Installation When You Start the Computer from the Product CD**

The following figure and corresponding table show the components that are involved in a clean installation when you start the computer from the product CD.

The figure and table list both the components and the related information files for each stage of Setup. Setup components use information files to store information that is later retrieved. Note that information might be retrieved by a component different from the one that stored it. Information might also be retrieved after restarting the computer.

### Note

- The following figure and table describe the contents of the I386 folder on the product CD. This folder contains the Setup components for installations on an x86-based computer. The Setup components on an Itanium architecture-based computer are located in the IA64 folder on the product CD.

### Setup Components Used During a Clean Installation When You Start the Computer from the Product CD




The following table describes the Setup components used for each stage of a clean installation.

**Setup Components Used During a Clean Installation When You Start the Computer from the Product CD**

| Components and Descriptions Grouped by Stage | Information File |
|---|---|
| **Initial stage, also known as data-collection stage and text-mode stage** | |

| | |
|---|---|
| **Setupldr.bin in the I386 folder on the product CD**<br><br>When the computer is started from the product CD, the boot sector on the CD initializes the setup loader instead of NTLDR, the loader normally used when starting a server running Windows. | Txtsetup.sif in the I386 folder on the product CD<br><br>Optionally, you can use the information file Winnt.sif. For more information, see the note after this table. |
| **Setupdd.sys**<br><br>Setupdd.sys carries out most of the actions during text-mode stage. | Txtsetup.sif in the I386 folder on the product CD<br><br>Optionally, you can use the information file Winnt.sif. For more information, see the note after this table. |
| **Final stage, also known as GUI-mode stage** | |
| **Syssetup.dll, Setupapi.dll, Netshell.dll, and other DLLs**<br><br>For the final stage of a clean installation, Syssetup.dll coordinates most of the actions. The operating system running during this stage is very similar to the complete Windows Server 2003 operating system, although only a small number of services are started. Syssetup.dll uses a variety of other DLLs for configuration tasks, for example, Iesetup.dll for configuring Internet Explorer and Ocgen.dll for configuring certain operating system components such as Windows Management Instrumentation (WMI). | Syssetup.dll uses:<br><br>$Winnt$.inf in *systemroot*\system32<br><br>Syssetup.inf and Layout.inf in *systemroot*\inf<br><br>At this stage, the operating system is very similar to the complete Windows Server 2003 operating system and uses many of the same .inf files (not listed here). |

When the computer first restarts after interactive setup, a few final user settings are applied, much like the process that happens at the beginning of any restart. The installation is then complete.

**Note**

- If you start the computer from the product CD and also insert a floppy disk containing an information file called Winnt.sif, Setup reads the copy of Winnt.sif on the floppy disk. For more information about Winnt.sif, see "Unattended Installation Technical Reference."

**Additional Setup Components Used in Other Installation or Upgrade Scenarios**

Some Setup components are called only if you start or restart Setup in a specific way. The following table lists and describes those components.

**Additional Setup Components Used in Other Installation or Upgrade Scenarios**

| Component | Installation or Upgrade Scenario |
|---|---|
| $Win_nt$.~ls folder<br><br>Dosnet.inf in the \I386 folder on the product CD | $Win_nt$.~ls is created when you perform an installation or upgrade by sharing Setup files in a network folder. During the initial stage of setup, all the folders and files needed for the new operating system are copied to the $Win_nt$.~ls folder on the local computer. This folder is also created when you run Winnt32.exe with the /makelocalsource option.<br><br>Dosnet.inf contains the list of files to copy to $Win_nt$.~ls. |
| Migrate.inf in $win_nt$.~bt folder | Migrate.inf stores the computer's drive letter assignments when you start either an upgrade or a clean installation on a computer running an earlier operating system such as Windows NT 4.0 or Windows 2000. |

| | |
|---|---|
| Sprestrt.exe and Undo_Guimode.txt | Sprestrt.exe, which runs at the start of GUI mode, determines if GUI mode ran previously and failed. If so, Undo_Guimode.txt is used to bring files back to the state appropriate for the beginning of GUI mode. |

**Log Files Created by Setup**
A number of log files are created during an installation or upgrade and are saved in the *systemroot* folder, for example, C:\Windows. The names and descriptions of the log files are as follows.

**Commonly Used Log Files**
**Setupapi.log**
Log file for Setupapi.dll, which processes .inf files during hardware installation and registration. This file can include entries that begin with #E, which indicates problems or errors.

**Setuperr.log**
Logs errors that occur during the setup process.

**Setuplog.txt**
Logs the overall setup process. This log includes information about the process of setup, the process of applying user settings after the computer's first restart into the new operating system, and the process of activation. The last line in this file shows the last action Setup was able to complete.

**Upgrade.txt**
Default name for the file that stores the compatibility report. (You can also save the report with a different name.) The compatibility report is created when you carry out any of the following:

- Insert the product CD in the CD drive, wait for the display to appear, and then click the appropriate options for checking system compatibility.
- Begin an upgrade of the existing operating system.
- Run Winnt32.exe with the /checkupgradeonly option.

**Other Log Files**
**Setupact.log**
Records some file copying and many actions related to hardware, for example, the detection of devices in the system. For upgrades, this file can contain some information about whether the computer hardware is compatible with Windows Server 2003.

**Winnt32.log**
Logs details about the initial stage of Setup, sometimes called the data-collection stage. If you attempted to upgrade to a product version that is not on the upgrade path from the existing operating system (for example, an upgrade to an English version of Windows Server 2003 from a non-English version of Windows 2000 Server), details will be logged in this file. This file is only created when you start Setup while the previous operating system is running, not when you start the computer from the product CD.

**Wsdu.log**
Logs Dynamic Update actions, if Dynamic Update is run.

**Logs created by other components**
During the last stage of Setup (also known as GUI-mode stage), the Windows Server 2003 operating system installation or upgrade is nearly complete, and therefore operating system components begin to create logs. The file name often indicates the component that created the log.

**Setup Processes**
An interactive setup can be started using a variety of methods. This Setup Technical Reference provides detailed descriptions of two of the most common methods for starting Setup:

- An upgrade when the computer is started from the operating system to be upgraded.
- A clean installation when the computer is started from the product CD.

The following information defines these two methods as well as other methods used for starting Setup. You can use almost all methods either for a clean installation or an upgrade. The exception is when you start the computer from the product CD, which is used for clean installations only.

**Note**

- Multiple methods are available for running Setup. This Technical Reference focuses on two methods, but it is not intended to provide guidance for choosing a method. For information about choosing a method, see "Related Information" later in this section.

**Methods for Starting Setup When the Computer is Started from the Product CD**

If you insert the product CD in the computer's CD drive and then start the computer from that drive, Setup starts automatically. This method for starting Setup differs from the other methods in the following ways:

- This is usually the most straightforward way to begin a clean installation. With this method, you can change the partitions on the hard disk and format partitions as needed. With other methods, you can make changes to the partitions or format the hard disk during installation only if those changes do not affect files that Setup must store on the hard disk while it is running.

- You must use the basic input/output system (BIOS) or Extensible Firmware Interface (EFI) to configure the computer so that you can start it from the CD drive.

- This method is used only for a clean installation (not an upgrade).

- You cannot use Dynamic Update if you start the computer from the product CD because the computer will not have network or Internet connectivity.

**Methods for Starting Setup When the Computer is Running MS-DOS**

If you start Setup from a computer that is running MS-DOS, you must use Winnt.exe; you cannot run Setup.exe or Winnt32.exe. Starting Setup from MS-DOS is the least preferred method; it does not offer all the options available with the other methods.

**Methods for Starting Setup When the Computer is Running Any 32-bit or 64-bit Windows Operating System**

The following table describes different methods for starting Setup from a computer that is running any 32-bit or 64-bit Windows operating system.

| **Medium and Method for Starting Setup** | **Notes** |
|---|---|
| **Medium:** Windows Server 2003 product CD<br><br>**Method:** Insert CD in local CD drive and wait for display to appear. | If the operating system already on the computer is one from which you can upgrade, this is the most straightforward way to begin an upgrade.<br><br>Allowing the display to appear after the CD is inserted in the drive is the same as running Setup.exe (in the root directory on the CD).<br><br>You can also start an installation or upgrade by changing to the I386 folder on the product CD and running Winnt32.exe. (For Itanium architecture-based computers, the folder is the IA64 folder on the product CD.) Winnt32.exe is described later in this table. |
| **Medium:** Windows Server 2003 product CD<br><br>**Method:** Insert CD in remote CD drive that has been shared on the network. On the local computer, run Setup.exe from the remote CD drive. | With this method, you do not have to keep track of a product CD or carry it from computer to computer.<br><br>Running Setup across the network requires slightly more free space on the local hard disk during Setup than if it is run locally. |
| **Medium:** A shared folder in which you place files from the Windows Server 2003 product CD.<br><br>**Method:** Run Setup.exe across the network. | With this method, you do not have to keep track of a product CD or carry it from computer to computer.<br><br>Running Setup across the network requires slightly more free space on the local hard disk during Setup than if it is run locally. |

| | |
|---|---|
| **Medium:** A shared folder in which you place files from the Windows Server 2003 product CD as well as other files. **Method:** From the I386 subfolder in the shared folder, run Winnt32.exe across the network, using appropriate options. (For Itanium architecture-based computers, run Winnt32.exe from the IA64 folder on the product CD.) | With this method and with appropriate preparation of the shared installation folder or subfolders, you can specify options with Winnt32.exe to change the installation process slightly. For example, you can create a shared subfolder with additional drivers, and have the subfolder installed at the same time as the operating system. As another example, you can have Dynamic Update use updated Setup files that you have downloaded to your shared installation folder. You download these files from the Windows Update Web site. For more information about Winnt32.exe, see "Command line reference A-Z" on the Microsoft TechNet Web site.<br><br>Running Winnt32.exe across the network requires slightly more free space on the local hard disk during installation than if Winnt32.exe is run locally.<br><br>Winnt32.exe can also be used for unattended installation. For more information, see "Unattended Installation Technical Reference." |

**Note**

- When you run Setup from an existing operating system, you have the option of installing a new operating system while leaving the old one in place. This results in a computer with multiple operating systems installed. For information about using the NTFS file system on such a computer, see "NTFS Technical Reference."

**Supported Upgrade Paths**

The operating systems from which you can upgrade depend on the specific product in the Windows Server 2003 family that you have purchased:

**Supported Upgrade Paths in Windows Server 2003**

| | **Standard Edition** | **Enterprise Edition** | **Datacenter Edition** | **Web Edition** | **Windows Small Business Server 2003** |
|---|---|---|---|---|---|
| Windows NT 4.0* Server | ● | ● | | | |
| Windows NT 4.0* Terminal Server Edition | ● | ● | | | |
| Windows NT 4.0* Enterprise Edition | | ● | | | |
| Windows 2000 Server | ● | ● | | | ● |
| Windows 2000 Advanced Server | | ● | | | |
| Windows 2000 Datacenter Server | | | ● | | |
| Windows Server 2003 Standard Edition | ● | ● | | | ● |
| Windows Server 2003 Enterprise Edition | | ● | | | |
| Windows Server 2003 Datacenter Edition | | | ● | | |
| Windows Server 2003 Web Edition | | | | ● | |
| Windows Server 2003** Beta3/RC1/RC2 | ● | ● | ● | ● | *** |
| Small Business Server 2000 | | | | | ● |

| | | | | | |
|---|---|---|---|---|---|
| Windows Small Business Server 2003 | | | | | • |

* Windows NT 4.0 upgrade is supported by Service Pack 5 (SP5) or later. If an earlier version of the service pack is installed, the upgrade is not possible.

** Interim releases of Windows Server 2003 will upgrade to the release to manufacturer (RTM) code of the same edition. For example, RC1 Standard Edition upgrades to RTM Standard Edition.

*** Release candidate (RC) to RTM code for Windows Small Business Server 2003 is supported.

**How Setup Performs an Upgrade**

In the process of performing an upgrade, Setup goes through several stages, carrying out specific actions in an orderly way during each stage. The Setup stages and actions are defined as follows:

- **Initial stage, also known as data-collection stage:** You start Setup from the existing operating system. Setup prompts you for items such as the product key and your decision about whether to run Dynamic Update. Setup also performs a compatibility check (with a report) during this stage. Finally, Setup copies the files necessary for starting the computer in the next stage and then restarts.

- **Intermediate stage, also known as text-mode stage:** Setup copies most of the files that will become the new operating system and then restarts.

- **Final stage, also known as GUI-mode stage:** Setup loads appropriate device drivers, applies network settings, copies additional files necessary for the operating system, installs Start menu items, and makes changes to the registry. At the end of this stage, Setup removes temporary files it used and then restarts the computer into the newly upgraded operating system.

The following numbered list shows the sequence of actions in each stage of an upgrade.

**Upgrade: Initial Stage (Also Known as Data-Collection Stage)**

1. You start Setup, typically by inserting the product CD in the CD drive and waiting for a display to appear. Setup examines the computer to determine which operating system exists on it, and offers the following options:

   - Installing the operating system

   - Performing additional tasks, that is, setting up a Remote Desktop Connection, browsing the CD, or viewing setup instructions and release notes

   - Checking system compatibility

   If the operating system that exists on the computer is in the Windows Server 2003 family, the display also offers the option to install optional Windows components (such as Certificate Service or Terminal Server).

   You can start Setup from an existing operating system in other ways. For more information, see "Setup Processes" earlier in this section. For some of these options, the display described in step 2 is the first display you see.

2. If you choose the first option in step 1, Setup asks you whether you want to perform an upgrade or a new installation.

3. Setup confirms that you agree with the product license.

4. Setup prompts you to type the product key.

   The 25-character product key is printed on a sticker on the back of your Windows Server 2003 product CD case.

5. Setup checks that the amount of memory and the space on the hard disk are adequate for the version of the product to which you are upgrading.

6. Setup asks whether to use Dynamic Update to get updated Setup files from the Microsoft Windows Update Web site. If you agree, Setup performs Dynamic Update (actions are recorded in Wsdu.log), and then checks that the space on the hard disk is still adequate for the upgrade.

7. Setup checks your system for compatibility, and displays a report if there are incompatibilities. For many types of incompatibilities, you can continue the upgrade. If, however, an incompatibility exists that is known to prevent Setup from completing, you cannot continue the upgrade. The report is saved with a default file name of Upgrade.txt, and actions are recorded in Winnt32.log.

8. Setup copies files to the hard disk. These files are necessary for starting the computer in the next stage.

9. Setup restarts the computer.

**Upgrade: Intermediate Stage (Also Known as Text-Mode Stage)**

1. Setup rechecks that the space on the hard disk is adequate, and then copies most of the files that will become the new operating system, replacing previous versions of the files where they exist.

2. Setup restarts the computer again.

**Upgrade: Final Stage (Also Known as GUI-Mode Stage)**

1. Using Plug and Play technology, Setup loads appropriate device drivers for the attached devices.

2. Setup configures settings using the existing registry as a baseline.

3. Setup copies files to the hard disk for operating system components that have not yet been installed.

4. Setup installs Start menu items and makes changes to the registry.

5. Setup sets file and folder permissions to protect the new operating system, for example, for files in the root directory, the system root folder, and the registry.

6. Setup runs Windows File Protection to ensure that all upgraded files are valid.

7. Setup removes any temporary files used during the installation process.

8. Setup restarts the computer. A few final user settings are applied after this restart, much like the process that happens at the beginning of any restart. The upgrade is then complete.

**How Setup Performs a Clean Installation when the Computer is Started from the Product CD**

In the process of performing a clean installation, Setup goes through several stages, carrying out specific actions in an orderly way during each stage. This Setup stages and actions are defined as follows:

- **Initial stage, also known as data-collection stage and text-mode stage:** You start Setup, typically by starting the computer from the product CD. Setup displays a character-oriented window during this stage. It collects information, mainly about what partitions to create or delete and what file system to use on an unformatted partition. Setup carries out your instructions for creating, deleting, and formatting partitions, then copies files and restarts.

- **Final stage, also known as GUI-mode stage:** Setup displays a graphical user interface (GUI) during this stage. It loads appropriate device drivers, then prompts for the following information:

    - Regional and language preferences

    - Your name and the name of the organization

- Product key
- Licensing mode
- Computer name and Administrator account password
- Date and time settings
- Network settings

Also during the final stage, Setup applies the network settings you specified, copies the files, installs Start menu items, and makes changes to the registry. At the end of this stage, Setup removes temporary files it used and then restarts the computer into the newly installed operating system.

The following numbered list shows the sequence of actions in each stage of a clean installation where the computer has been started from the product CD.

**Clean Installation: Initial Stage (Also Known as Data-Collection Stage and Text-Mode Stage)**

1. You turn the computer on and start it from the product CD.

   In order to start the computer from a CD, either you must press a key at the appropriate prompt, or you must prepare the computer ahead of time by configuring the BIOS or EFI to specify a boot sequence that will cause the computer to start from the CD drive. For more information, see the manufacturer's instructions for your computer.

2. Setup gives you the opportunity to load a mass storage driver (press F6). At this time, you can also press F5 to load a Hardware Abstraction Layer (HAL) file.

3. Setup gives you the opportunity to run Automated System Recovery (press F2).

4. If you continue with installation (not Automated System Recovery), Setup loads a set of generic device drivers.

5. Setup displays a welcome screen and gives you the opportunity to start the Recovery Console (press R).

6. Setup confirms that you agree with the product license.

7. Setup searches the hard disks for an existing installation of the product that you are installing. If an existing installation is found, Setup gives you the opportunity to repair that installation (press R).

8. Setup checks that the amount of memory and the hard disk size are adequate for the version of the product that you are installing.

9. Setup gives you the opportunity to delete or create partitions on the hard disk. If you create a partition and specify it as the installation partition, you must specify a file system (NTFS or FAT) and specify whether to use quick format (integrity of sectors is not verified) or full format (bad sectors are identified and tracked).

10. After carrying out the specified partition and file system changes, Setup checks that the partition you selected contains enough space to hold the installation, and then copies files to the partition. The files are most of those that will become the new operating system.

11. Setup restarts the computer.

**Clean Installation: Final Stage (Also Known as GUI-Mode Stage)**

1. Using Plug and Play technology, Setup loads appropriate device drivers for the attached devices.

2. Setup requests information about the following items:

    - Regional and language options
    - Your name and the name of the organization (the name of the organization can be left blank)
    - Product key

      The 25-character product key is printed on a sticker on the back of your Windows Server 2003 product CD case.
    - The licensing mode
    - Computer name and administrator password
    - Date and time settings

3. Setup installs a basic set of network components to begin the process of establishing network connectivity.
4. Setup requests information about the following items:

    - Networking settings
    - Workgroup or computer domain

5. Setup configures the settings that were specified.
6. Setup copies files to the hard disk for operating system components that have not yet been installed.
7. Setup installs Start menu items and makes changes to the registry.
8. Setup sets file and folder permissions to protect the new operating system, for example, for files in the root directory, the systemroot folder, and the registry.
9. Setup runs Windows File Protection to ensure that all installed files are valid.
10. Setup removes any temporary files used during the installation process.
11. Setup restarts the computer. A few final user settings are applied after this restart, much like the process that happens at the beginning of any restart. The installation is then complete.

**Configuration Elements That You Can Change After Setup**

In Windows Server 2003, most of what you specify during Setup can be changed at a later time. The following table specifies what you can change after Setup.

**Configuration Elements That You Can Change After Setup**

| Configuration Elements Specified During Setup | Configuration Elements You Can Change After Setup |
|---|---|
| **Partitions:** Arrangement, order, or format of partitions on the hard disk | You can use Disk Manager to change various aspects of the partitions. However, unless you reinstall, you cannot perform actions that involve deleting, moving, or changing the letter of the boot partition (also called the boot volume) or the system partition (also called the system volume). |

| | |
|---|---|
| **Components:** Installed components of the operating system | You can use Add or Remove Programs in Control Panel to add or remove Windows components after installation, for example, the Remote Storage component. To add components, you will need access to the product CD or a network share containing files from the product CD. |
| **Regional and language options Support for East Asian languages:** These languages include Chinese, Japanese, or Korean, or the complex script and right-to-left languages, which include Arabic, Armenian, Georgian, Hebrew, the Indic languages, Thai, or Vietnamese. | You can use Regional and Language Options in Control Panel to change most of these options after installation. However, to add support for East Asian languages after installation, you will need access to the product CD or a shared network folder containing files from the product CD. |
| **Name** entered by the person who installed the computer (in Personalize Your Software) | You can change this setting by changing the following registry key: \HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft \Windows NT\CurrentVersion\RegisteredOwner |
| **Other options** | You can change other settings after installation, usually through Control Panel or through Administrative Tools (in Control Panel). |

**Note**

- Incorrectly editing the registry might severely damage your system. Before making changes to the registry, you should back up any valued data on the computer

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Automating and Customizing Installations" in the Windows Server 2003 Deployment Kit on the Microsoft Web site.
- "Global Development and Computing Portal" on the Microsoft Web site for information about supporting multiple user languages.
- "Dynamic Update" on the Microsoft Windows Update Web site for information about creating a shared installation folder in which you can download updated Setup files that you can use with Dynamic Update.
- Remote Installation Services Technical Reference
- Sysprep Technical Reference
- Unattended Installation Technical Reference

# Setup Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Setup Tools
- Related Information


This section describes the Winnt32.exe and Winnt.exe commands.

**Setup Tools**
The following tools are associated with Setup.

**Winnt32.exe**
Winnt32.exe is on the Windows Server 2003 product CD in the I386 folder. For Itanium architecture-based computers, Winnt32.exe is located in the IA64 folder on the product CD.

**Version Compatibility**
Winnt32.exe can be run on a computer running Windows 95, Windows 98, Windows Millennium Edition, Windows NT, Windows 2000, Windows XP, or a product in the Windows Server 2003 family. If you run Winnt32.exe on a computer running Windows NT 4.0, you must first apply Service Pack 5 or later.

To install Windows Server 2003, use the version of Winnt32.exe on your Windows Server 2003 product CD, not the version of Winnt32.exe that exists on a CD for a different product.

**Description**
Winnt32.exe provides some ways to run Setup that contrast with starting the computer from the product CD or starting Setup by inserting the product CD and waiting for a display to appear automatically. With options available with Winnt32.exe, you can make slight changes to the interactive installation process, or you can run unattended installation (by using an answer file).

The typical way you would use Winnt32.exe with interactive installations is by creating a shared network folder containing files from the Windows Server 2003 product CD. You would then prepare the folder to work with specific Winnt32.exe options and run Winnt32.exe across the network using those options. For example, you can create a shared subfolder with additional drivers and have the subfolder installed at the same time as the operating system. As another example, you can perform an installation in which you use Dynamic Update files that you have downloaded and prepared in a special shared folder. (When you use the correct options with Winnt32.exe, the files are downloaded from the Windows Update Web site.)

For more information about Winnt32.exe, see “Command line reference A-Z” on the Microsoft TechNet Web site. For more information about using Winnt32.exe for unattended installation, see “Unattended Installation Technical Reference.”

**Winnt.exe**
**Category**
Winnt.exe is on the Windows Server 2003 product CD in the I386 folder.

**Version Compatibility**
Winnt.exe can be run on a computer running MS-DOS. To install Windows Server 2003, use the version of Winnt.exe on your Windows Server 2003 product CD, not the version of Winnt.exe that exists on the CD for a different product.

**Description**
Winnt.exe is a command-line interface for starting an installation or upgrade on a computer running MS-DOS.

Winnt.exe provides some ways to run Setup that contrast with starting the computer from the product CD or starting Setup by inserting the product CD and waiting for a display to appear automatically. With options available with Winnt.exe, you can make slight changes to the interactive installation process, or you can run unattended installation (by using an answer file).

The typical way you would use Winnt.exe with interactive installations is by creating a shared network folder containing files from the Windows Server 2003 product CD. You would then prepare the folder to work with specific Winnt.exe options and run Winnt.exe across the network using those options. For example, you can create a shared subfolder with additional drivers and have the subfolder installed at the same time as the operating system.

For more information about Winnt.exe, see “Command line reference A-Z” on the Microsoft TechNet Web site. For more

information about using Winnt.exe for unattended installation, see "Unattended Installation Technical Reference."

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Automating and Customizing Installations" in the Windows Server 2003 Deployment Kit on the Microsoft Web site.
- "Windows Installer" on MSDN for information about supporting efficient installation and configuration of applications.
- Remote Installation Services Technical Reference
- Sysprep Technical Reference
- Unattended Installation Technical Reference

# Unattended Installation Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Unattended Installation Technical Reference**
Unattended installation is an automated installation technology in Microsoft Windows Server 2003 that you can use to install or upgrade an operating system with minimal user intervention. Unattended installation uses an answer file to automatically answer the questions that the Windows Setup program asks during the Windows setup process. Unattended installation is typically used during large-scale rollouts when it would be too slow and costly to have administrators or technicians interactively install the operating system on individual computers.

**In this subject**


- What Is Unattended Installation?
- How Unattended Installation Works
- Unattended Installation Tools and Settings

# What Is Unattended Installation?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Unattended Installation?**

**In this section**


- Common Unattended Installation Scenarios
- Unattended Installation Dependencies
- Technologies Related to Unattended Installation
- Related Information


Unattended installation is an automated installation technology in Microsoft Windows Server 2003 that you can use to install or upgrade an operating system with minimal user intervention. Unattended installation is typically used during large-scale rollouts when it would be too slow and costly to have administrators or technicians interactively install the operating system on individual computers.

One of the essential components of an unattended installation is an answer file. An answer file is a simple text file that instructs the Windows Setup program how to install and configure the operating system. To perform an unattended installation, you run Setup with a special command-line parameter that instructs Setup to read the answer file. Then, during the setup process, instead of prompting you for installation and configuration information interactively, Setup follows the instructions contained in the answer file.

The following figure shows how an unattended installation differs from an interactive setup.

**Comparison of Unattended Installation and Interactive Setup**




Like all automated installation technologies, unattended installation has many benefits that standard interactive setup does not.

**Fewer errors during installation**

Because unattended installation uses an answer file to install and configure the operating system, there is minimal interaction by administrators or technicians during the setup process. This reduces the number of errors that are introduced during the setup process.

**Greater consistency during a rollout**
By using the same answer file to install and configure the operating system, you can ensure that all of the computers in your organization are set up exactly the same way.

**Shorter installation times**
Unattended installation is faster than interactive setup because Setup does not have to prompt administrators or technicians for configuration information and wait for them to enter it; instead, Setup reads configuration settings from an answer file.

**Lower support costs**
By minimizing errors during the setup process, increasing the consistency of the computers in your organization, and reducing the amount of time a technician needs to spend setting up a computer, you can reduce the overall support costs in your organization.

Unattended installation is a suitable automated installation technology if you need to:

- Upgrade a Windows operating system to Microsoft Windows XP Professional, Microsoft Windows Server 2003, Standard Edition; or Microsoft Windows Server 2003, Enterprise Edition operating systems.
- Perform clean, automated installations on computers that have heterogeneous hardware configurations.
- Perform clean, automated installations on specific types of servers, such as domain controllers, remote access servers, and servers that run Certificate Services or the Cluster service.
- Configure a wide range of operating system settings during an automated installation without using batch files and scripts.

**Common Unattended Installation Scenarios**
Although unattended installation is a robust installation technology, and can be used to roll out many types of server configurations, some installation scenarios are particularly well suited for unattended installation.

**Heterogeneous Hardware**
Unattended installation is ideal if you are installing the same server configuration on computers that have dissimilar or heterogeneous hardware. For example, you can use unattended installation to roll out the same file server configuration to a group of computers that have different processor configurations and different storage configurations because an unattended installation can detect and install the appropriate hardware abstraction layer (HAL) and the appropriate hard disk controller during the installation process. Other automated installation technologies — particularly image-based technologies — have a limited ability to detect and install different HALs and hard disk controllers during the installation process.

In addition, unattended installation is equally adept at handling other hardware differences, including differences among network adapters, media devices (CD drives), and video adapter cards. In large part, this is because unattended installation is just an automated process that uses the standard Setup program. Therefore, whatever device Setup can detect or install, unattended installation can also detect or install.

**Upgrades**
You can use unattended installation to upgrade a Windows server operating system to Microsoft Windows Server 2003, Standard Edition or Microsoft Windows Server 2003, Enterprise Edition operating systems. You cannot use unattended installation to upgrade a Windows operating system to Windows Server 2003, Web Edition. No other automated installation technology can perform upgrades.

Unattended upgrade installations are commonly performed by small organizations that have only a limited number of servers and a limited number of IT professionals to perform installations. An unattended upgrade installation is preferred in this situation because existing server settings are retained during an upgrade (for example, user rights, groups, and permissions settings), and applications and files do not need to be reinstalled.

The following table shows the types of upgrades that you can perform by using unattended installation.

**Types of Upgrades by Using Unattended Installation of Windows Server 2003**

| Existing Operating System | Upgrade to Windows Server 2003, Standard Edition | Upgrade to Windows Server 2003, Enterprise Edition |
|---|---|---|
| Microsoft Windows NT Server 3.51 | | |
| Microsoft Windows NT Server version 4.0 (with Service Pack 5) | ● | ● |
| Microsoft Windows NT version 4.0, Terminal Server Edition | ● | ● |
| Microsoft Windows NT version 4.0, Enterprise Edition | | ● |
| Microsoft Windows 2000 Server | ● | ● |
| Microsoft Windows 2000 Advanced Server | | ● |
| Microsoft Windows Server 2003, Standard Edition | | ● |

There are three exceptions to the upgrades listed in the previous table:

- You cannot upgrade from one localized version of a Windows operating system to another localized version of a Windows operating system. For example, you cannot upgrade from the Japanese version of Windows 2000 Server to the English version of Windows Server 2003, Standard Edition.

- You cannot upgrade the operating system of a computer that has the Multilingual User Interface Pack (MUI Pack) except to the English version of the operating system.

- You cannot upgrade the operating system from Windows NT 4.0 unless you have Service Pack 5 or later installed. It is not possible to upgrade from Windows NT 4.0 with Service Pack 4 or earlier.

**Note**

- You cannot use unattended installation to install or upgrade to Windows Server 2003, Datacenter Edition.

**Domain Controllers and Servers Running Certificate Services and Cluster Service**

Unattended installation is an ideal automated installation technology for deploying domain controllers, and member servers that run Certificate Services and Cluster service. Although you can use other automated installation technologies to deploy these types of servers, those technologies require you to use scripts or supplemental answer files after the automated installation is finished. With unattended installation, the answer file contains all of the necessary settings for installing these types of servers.

**Master Installations for Image-based Deployment Technologies**

You can use unattended installation to create master installations for image-based installation technologies, such as the System Preparation (Sysprep) tool, or Remote Installation Services (RIS) when it is being used for image-based installation. For example, to use Sysprep, you need to create a master installation from which you create a disk image. You can create the master installation manually by running Setup, and then manually configuring operating system settings and installing software, or you can automate the creation of the master installation by performing an unattended installation on the master computer.

By using unattended installation to create master installations you:

- Introduce fewer installation errors because unattended installation relies on an answer file for configuration settings.

- Enhance consistency among all of your master installations because you can use the same answer file as a basis for each of your master installations.

- Decrease the amount of time it takes to modify and create a master installation because unattended installation is much simpler and quicker than manual installation.

- Maintain a record of each of your master installations because the settings are all recorded in an answer file (with the exception of password and product key information).

**Unattended Installation Dependencies**
To perform an unattended installation, you need the following:

- Answer file
- Operating system files, including the Windows Setup program

Because an unattended installation is an automated way of using Setup, you can perform an unattended installation only by running one of the two Windows Setup programs, Winnt.exe or Winnt32.exe. Winnt.exe is a 16-bit Setup program that you can use to perform an unattended installation by starting a destination computer with a 16-bit operating systems such as MS-DOS. Winnt32.exe is a 32-bit Setup program that you can use to perform an unattended installation by starting a destination computer with a 32-bit operating system such as Windows 95, Windows 98, Windows NT, Windows 2000, Windows XP, or Windows Server 2003. You can also perform an unattended installation on an Itanium architecture-based computer by running Winnt32.exe from the command prompt on the following: Windows XP 64-Bit Edition; the 64-bit version of Windows Server 2003, Enterprise Edition; or the 64-bit version of Windows Server 2003, Datacenter Edition.

An answer file can be created in two ways: by using a text editor such as Notepad, or by using Setup Manager (Setupmgr.exe). Setup Manager is a wizard-like tool that prompts you for a series of configuration parameters, and then writes your responses to an answer file.

You also need a Windows Server 2003 product CD, or a shared distribution folder ("distribution share") that contains the Windows Server 2003 installation files. If you store the installation files on a distribution share, each of your destination computers needs a high-speed network connection and appropriate permissions to access the server on which the distribution share resides.

**Technologies Related to Unattended Installation**
Unattended installation is similar to two other automated installation technologies: image-based installation with the Sysprep tool and Remote Installation Services (RIS) installation. Like unattended installation, these technologies are designed specifically as mass roll-out or deployment solutions. Although you can use these technologies to automate the installation of Windows Server 2003, each of these technologies is based on a different installation process, and relies on a different set of programs and tools.

**Image-based installation with Sysprep**
Image-based installation with the Sysprep tool is a method of copying, also known as *cloning*, preconfigured operating systems and software applications to servers.

Image-based installation with Sysprep is a suitable automated installation technology if you need to:

- Install identical operating systems and software configurations on multiple computers.
- Install an operating system and software configuration as quickly as possible.
- Eliminate administrator and technician interaction during the setup process.
- Minimize post-installation tasks.
- Perform clean installations of an operating system, rather than upgrade existing installations.
- Install operating systems on computers that have similar hardware and a compatible HAL.

You need two tools to perform an image-based installation with Sysprep: Sysprep, which can be found on any Windows Server 2003 product CD; and a disk-imaging program. Usually, you use a non-Microsoft disk-imaging program, which you must purchase from a vendor. However, you can instead perform an image-based installation with Sysprep by using Windows Server 2003 Automated Deployment Services (ADS). ADS is part of the Dynamic Systems Initiative (DSI), and can be used to rapidly deploy Windows server operating systems.

Image-based installation with Sysprep requires substantial planning and design, but it is the fastest method of installing

operating systems and applications. You cannot use it to perform operating system upgrades.

**RIS installation**
RIS installation is a method of cleanly installing an operating system and applications over a network with minimal administrator or technician interaction. RIS uses Pre-Boot Execution Environment (PXE) technology to enable client computers without an operating system to start up from a PXE-enabled network adapter and connect remotely to a RIS server, which installs a supported operating system.

RIS is a suitable automated installation technology if you need to:

- Eliminate administrator and technician interaction during the setup process.
- Minimize post-installation tasks.
- Perform clean installations of an operating system, rather than upgrade existing installations.

RIS installation is more complex and requires more upfront configuration and planning than image-based installation and unattended installation. However, RIS installation is an efficient deployment technology if you have a high-speed network and you need to perform clean installations of operating systems and preconfigured applications on a large number of computers.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Setup Technical Reference
- Sysprep Technical Reference
- Remote Installation Services Technical Reference

# How Unattended Installation Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Unattended Installation Architecture
- Unattended Installation Physical Structures
- Unattended Installation Processes and Interactions
- Related Information


Unattended installation is an automated installation technology that you can use to install an operating system with minimal administrator or technician intervention. Unattended installation is typically used during large-scale rollouts when it would be too slow and costly to have administrators or technicians interactively install the operating system on individual computers. It can also be used to upgrade operating systems.

One of the essential components of an unattended installation is an answer file that you create. An answer file is a simple text file that instructs the Windows Setup program how to install and configure the operating system. To perform an unattended installation, you run Setup by using a special command-line parameter that instructs Setup to read the answer file, or you start the destination computer with the product CD and then insert a floppy disk that contains an answer file. Then, during the setup process, instead of prompting you for installation and configuration information interactively, Setup follows the instructions contained in the answer file.

The following sections provide an in-depth view of how unattended installation works in an optimal environment. An optimal environment for unattended installation contains the following elements:

- Plug and Play hardware that is compatible with Windows Server 2003.
- High-bandwidth network connectivity between the computer on which the unattended installation is being performed (destination computer) and any installation file that resides in a shared folder (distribution share).
- Physical access to the destination computer.
- BIOS and hardware configured so that you can start the computer from the product CD.
- BIOS and hardware configured so that you can start the computer from the floppy disk.

**Unattended Installation Architecture**

Unattended installation relies on the same setup components as an interactive setup. These components are used in three distinct setup stages, including the initial stage (also known as data-collection stage), the intermediate stage (also known as text-mode stage), and the final stage (also known as GUI-mode stage). These three stages culminate with a final restart, after which a user logs onto the computer for the first time. In addition, unattended installation relies on several answer files and optional components, such as a distribution share, a product CD, and a renaming answer file to rename long file and folder names. How these components interact is illustrated in the following figure.

**Unattended Installation Architecture**




The following are descriptions of the primary unattended installation components.

**$ldr$**

The Setup loader, similar to Ntldr, which starts the destination computer in a special setup state during which the intermediate stage of unattended installation occurs. This file is named Setupldr.bin on the product CD. Setupldr.bin is renamed $ldr$ and copied to the destination computer's hard disk when you do not start the destination computer from the product CD.

**$win_nt$.~ls**

Temporary installation folder created on the hard disk of the destination computer that contains the files and folders that are copied from the I386 folder in the distribution share.

**$win_nt$.~bt**

Temporary installation folder created on the hard disk of the destination computer that contains temporary startup files. The startup files are used to restart the computer into a special setup state that allows the intermediate stage of unattended installation to occur.

**$Winnt$.inf**

Final version of the unattended installation answer file that Setup creates and saves in the *systemroot*\System32 folder on the destination computer. Setup makes this a hidden file and names it $winnt$.inf, regardless of what the name of the original answer file is.

**$$Rename.txt**

A file that lists all files contained in a particular folder that need to be renamed from an 8.3 format name to a long file name during the setup process. All file names must be in 8.3 format if you run Setup by using Winnt.exe, or if you copy a set of files from one computer to another using MS-DOS.

**Cmdlines.txt**

A configurable text file that you use to customize an unattended installation. Cmdlines.txt contains a list of commands that run synchronously after Setup finishes, but before a computer restarts.

**Destination computer**

Computer on which you are installing an operating system. During a large-scale rollout, there can be typically hundreds or thousands of destination computers.

**Distribution share (I386\)**

Hierarchical folder structure that is stored in a shared folder on a remote computer (usually a server). The distribution share contains all of the Windows Server 2003 installation files that are necessary to perform an unattended installation. You can use Setupmgr.exe to create a distribution folder, or you can create one manually by using Windows Explorer.

**[GUIRunOnce]**

A section in an unattended installation answer file that contains a list of commands, programs, scripts, or batch files that run the first time a user logs on to the computer after the operating system is installed.

**Setupapi.dll**

Dynamic-link library that is responsible for installing hardware.

**Setupdd.sys**
System file responsible for performing tasks during the intermediate stage of unattended installation, such as creating or modifying the system folder (usually, C:\Windows or C:\Winnt), copying system files and folders from the temporary folder or the product CD to the system folder, and building the registry. It is also responsible for creating partitions and formatting disks.

**Setupldr.bin**
The Setup loader, similar to Ntldr, which starts the destination computer in a special setup state during which the intermediate stage of unattended installation occurs. Setupldr.bin is renamed $ldr$ and copied to the destination computer's hard disk if you do not start the destination computer from the product CD.

**Syssetup.dll**
Dynamic-link library that performs installation and configuration tasks during the final stage of unattended installation. This component might have to access the CD in order to install optional components.

**Technician computer**
Computer used by an administrator or technician to use Setup Manager or Notepad to create and modify answer files, such as Unattend.txt or Cmdlines.txt.

**Txtsetup.sif**
System information file that Setup uses during the setup process. Txtsetup.sif defines the names of installation files and folders, where files and folders are copied, and how to process existing files and folders (for example, delete or overwrite). Txtsetup.sif is saved in the I386 folder of the distribution share and resides in the I386 folder of every product CD.

**Unattend.txt**
Default name for an unattended installation answer file. You can name an answer file anything you want, but Unattend.txt is the name that Setupmgr.exe uses when you first create an answer file.

**Winlogon.exe**
A component of the Windows operating system that provides interactive logon support. Winlogon is designed around an interactive logon model that consists of three components: the Winlogon executable, a Graphical Identification and Authentication dynamic-link library (DLL) referred to as the GINA, and many network providers.

**Winnt32.exe**
The 32-bit and 64-bit Setup program. You can run the 32-bit version of this program on Windows 95, Windows 98, Windows Millennium Edition, Windows NT, Windows 2000, Windows XP, or a product in the Windows Server 2003 family. The 32-bit version of this program can be found in the I386 folder on any 32-bit Windows operating system product CD. You can run the 64-bit version of this program on Windows XP 64-bit Edition, the 64-bit version of Windows Server 2003, Enterprise Edition, and the 64-bit version of Windows Server 2003, Datacenter Edition. The 64-bit version of this program can be found in the IA64 folder on any 64-bit operating system product CD.

**Winnt32u.dll**
Dynamic-link libraries that perform installation tasks during the initial stage of unattended installation. These include assessing the current operating system that is installed on the destination computer, copying installation files and folders from the distribution share to the temporary folders, parsing the answer file, and creating intermediate information files that are used by the intermediate stage of unattended installation.

**Winnt.exe**
The 16-bit Setup program. You can use this program to start an unattended installation from any 16-bit operating system, such as MS-DOS.

**Winnt.sif**
Name you must use for the answer file if you are performing an unattended installation by starting a destination computer from the Windows Server 2003 product CD. Winnt.sif is also the name of an intermediate information file that Setup creates after the initial stage of unattended installation. This intermediate file contains all of the information that is in the answer file (regardless what the name of the answer file is), as well as information about the destination computer and the type of installation you are performing.

**Unattended Installation Physical Structures**
Two physical structures are essential to an unattended installation: the distribution share and the answer file. You can create these structures by using Setup Manager (Setupmgr.exe) or you can create them manually by using Windows Explorer and a text editor, such as Notepad. These structures are illustrated in the following figure and are described in the sections that follow.

### Unattended Installation Physical Structures




**Distribution Share Structure**

A distribution share consists of a single, top-level folder named I386 and several subfolders. The following figure shows the structure of a distribution share.

**Distribution Share Structure**




**Distribution Share Folders**

The following sections describe the folders that can exist in a distribution folder.

**I386**

This is the distribution share. You create it at the root of the distribution server (the server on which the distribution share is located). This folder must be named I386 if it contains installation files for a 32-bit operating system; it must be called IA64 if it contains installation files for a 64-bit operating system. You can create this folder manually by copying the I386 folder that is on a Windows Server 2003 or Windows XP Professional product CD to your distribution server. You can also use Setup Manager to create this folder.

**$OEM$**

You create the $OEM$ subfolder directly beneath the I386 folder. During the setup process, you can automatically copy folders, standard 8.3 format files, and any tools needed for your automated installation process to the $OEM$ subfolder. However, the contents of this folder are copied only when you include the following setting in your answer file:

```
[Unattended]
OemPreInstall = Yes
```

If you want to create the $OEM$ subfolder in a different location than the distribution share, you can use the OemFilesPath parameter in the answer file to specify a different location.

The $OEM$ subfolder can contain the optional file Cmdlines.txt, which contains a list of commands to be run at the end of the final stage of unattended installation. These commands can be used to install additional services or applications that you want to include with your installation.

As long as Setup finds the $OEM$ subfolder in the root of the distribution share, Setup copies all the files found in this folder

to a temporary folder that is created during the initial stage of unattended installation on the destination computer.

**$OEM$\Textmode**
The $OEM$\Textmode subfolder contains new or updated files for installing mass storage device drivers and hardware abstraction layers (HALs). These files can include an OEM HAL, drivers for a small computer system interface (SCSI) device, and a Txtsetup.oem file, which directs the loading and installing of these components. Be sure to include the Txtsetup.oem file in this folder: most device drivers cannot be installed without a Txtsetup.oem file. All files placed in the $OEM$\Textmode subfolder (HALs, drivers, and Txtsetup.oem) must be listed in the [OEMBootFiles] section of the answer file.

**$OEM$\$$**
The $OEM$\$$ subfolder is equivalent to the systemroot or windir environment variable. The subfolder contains additional files that need to be copied to the various subfolders of the Windows Server 2003 or Windows XP Professional installation directory. The structure of this subfolder must match the structure of a standard Windows Server 2003 or Windows XP Professional installation, where $OEM$\$$ matches systemroot or windir (for example, C:\Windows), $OEM$\$$\Fonts matches *systemroot*\Fonts, and so on. Each subfolder must contain the files that need to be copied to the corresponding system folder on the destination computer.

**$OEM$\$$\Help**
A subfolder that contains the OEM Help files to be copied to the *systemroot*\Help folder during the setup process.

**$OEM$\$$\System32**
A subfolder that contains files to be copied to the *systemroot*\System32 folder during the setup process.

**$OEM$\$1**
This folder points to the drive on which Windows Server 2003 or Windows XP Professional is installed; $1 is equivalent to the systemdrive environment variable. For example, if you are installing Windows Server 2003 on drive D, $OEM$\$1 points to drive D.

**$OEM$\$1\Pnpdrvrs**
You can use the $OEM$\$1\Pnpdrvrs subfolder to place new or updated Plug and Play device drivers in your distribution shares. These folders are copied to the *systemdrive*\Pnpdrvrs location on the destination computer. Adding the OemPnPDriversPath parameter to your answer file directs Windows Server 2003 or Windows XP Professional to look (both during and after the setup process) for new or updated Plug and Play drivers in the folders that you created, in addition to those originally included with the system. Note that you can replace the file name Pnpdrvrs with any name that is eight or fewer characters long.

**$OEM$\$1\Sysprep**
The $OEM$\$1\Sysprep subfolder is optional. This subfolder contains the files that you need to run the System Preparation (Sysprep.exe) tool.

**$OEM$\drive_letter**
During the intermediate stage of unattended installation, the structure and contents of each $OEM$*drive_letter* subfolder are copied to the root of the corresponding drive on the destination computer. For example, files that you place in the $OEM$\D subfolder are copied to the root of drive D. You can also create subfolders within these subfolders. For example, $OEM$\E\Misc causes Setup to create a subfolder called Misc on drive E.

**Converting Short File Names to Long File Names by Using $$Rename.txt**
If you are using Winnt.exe to perform an unattended installation, your distribution share can contain only files and folders that have short file names. This is because Winnt.exe runs only on MS-DOS, and MS-DOS can process only files and folders that use the 8.3 naming convention. The 8.3 naming convention allows only eight characters to the left of the decimal point and three characters to the right of the decimal point. Because of the short file name limitation, you need to shorten the names of files and folders so that you can put them into your distribution share. You can convert these short file and folder names back to long file and folder names during the setup process by using the $$Rename.txt file.

$$Rename.txt lists all of the files and folders that need to be renamed. You need only one $$Rename.txt file to rename files in an entire distribution share. If you are using Winn32.exe to perform an unattended installation, short file names are automatically converted to long file names during the setup process.

To use $$Rename.txt, put the file in the $1 folder in your distribution share. Saving the $$Rename.txt file in this location will ensure that the $$Rename.txt file is copied to the systemdrive on the destination computer. Setup automatically looks for $$Rename.txt; if it finds a $$Rename.txt file, Setup renames the files and folders that are in your distribution share. The syntax for $$Rename.txt is:

```
[
section_name_1
]
```

```
short_name_2 = long_name_2short_name_1 = long_name_1short_name_x= "long_name_x"
.
.
.
[section_name_2]
short_name_1 = long_name_1short_name_2 = long_name_2short_name_x= "long_name_x"
.
.
.
```

Where:

- *section_name_x* is the path to the folder that contains the files and folders. A section does not need to be named; it can have a backslash (\) as a name, which indicates that the section contains the names of the files or subfolders that are in the root of the drive.

- *short_name_x* is the name of the file or folder within this named folder that needs to be renamed. The name *must not* be enclosed in quotation marks.

- *long_name_x* is the new name of the file or folder. If the name contains spaces or commas, it *must* be enclosed in quotation marks.

**Answer File Structure**

The answer file for unattended installations of Windows Server 2003 is usually named Unattend.txt, but you can name it anything you want. However, if you start a destination computer from a product CD and the answer file is on a floppy disk, you must name the answer file Winnt.sif. Setup will then detect the answer file on the disk without user input.

The answer file supplies Setup with answers to all the questions that an administrator or technician would be asked during a standard, interactive setup. The answer file also contains information about your installation and configuration requirements. In addition, an answer file tells Setup how to interact with the distribution shares and files that you have created (or the installation files on the product CD).

An answer file contains multiple sections — some mandatory and some optional — that you can modify. Section names are always enclosed in square brackets (for example, [Unattended]).

Every section has one or more entries that contain installation and configuration information. Entries are usually composed of an entry name, an equal sign, and a value (for example, `ComputerName = Computer1`). The entry name represents a specific computer setting or action; the value represents the unique way you want the setting configured or the action performed. In some cases, entries do not have an entry name. In these cases, a value or a list of values (usually commands) appear in the answer file below a section name. The following example shows a section that contains only values and does not have entry names.

```
[GUIRunOnce]
dcpromo /answer:a:\installdc.txt
%systemroot%\cluster\cluster.exe
```

Use the following guidelines when you create or modify an answer file:

- Sections do not have to be any particular order.

- Sections and entries are not case sensitive.

- Entries must have valid values. For example, “on” is not equivalent to “enable,” and “off” is not equivalent to “disable.” Invalid sections and entries can generate errors, or cause Setup to fail.

- Each entry must have a value.

- Each section and entry must appear on a separate line.

- Comments must begin with a semi-colon (;). Do not use “REM” to specify a comment.

**Note**

- Do not modify undocumented sections or entries in an answer file; doing so can cause unpredictable results.

The following table lists all of the valid sections that can appear in an answer file for an unattended installation.

**Answer File Sections for Unattended Installation**

| Section | Description |
|---|---|
| [*adapter_specific_protocol_section*] | Contains an entry for specifying a network component. Must be used in conjunction with the [Networking] section. |
| [Branding] | Contains entries for branding Microsoft Internet Explorer during an unattended installation. If you do not specify this section, Setup uses the default for Windows Server 2003. |
| [CertSrv_Client] | Contains entries for the Web client component of Certificate Services. You do not need to include this section in the answer file to install the Web client component of Certificate Services on the computer. Enabling Certificate Services in the [Components] section of the answer file installs the necessary files. |
| [CertSrv_Server] | Contains entries for specifying the server components of Certificate Services. The entries in this section are supported only for the Windows Server 2003. |
| [Components] | Contains entries for installing the optional components of Windows XP Professional and Windows Server 2003. |
| [Data] | Contains entries for starting directly from the Windows Server 2003 product CD when performing an unattended installation. |
| [DCInstall] | Contains entries for installing a domain controller after the initial setup of the operating system finishes. |
| [Display] | Contains entries for specifying display settings for graphics devices. |
| [FavoritesEx] | Contains the default Favorites settings for the Windows browser. If not present, Setup uses the default Uniform Resource Locator (URL) for Windows XP or Windows Server 2003. |
| [Fax] | Contains entries for specifying custom installation information for the Fax service. |
| [GuiRunOnce] | Contains commands that run the first time a user logs on to the computer after final (GUI-mode) stage of Setup completes. |
| [GuiUnattended] | Contains entries for preparing the graphical user interface (GUI) for unattended installation. |
| [Homenet] | Contains entries for installing Home Networking settings for network adapters, Internet Connection Sharing (ICS), and Internet Connection Firewall (ICF). |
| [Identification] | Contains entries for specifying the network identification of a computer. If these entries are not present, Setup adds the computer to the default workgroup. If there is not enough information for this entry, Setup prompts the user to provide this information. |
| [IEHardening] | Contains entries for configuring Web content zones. |
| [InternetServer] | Contains entries for configuring the installation folder for Internet Information Services (IIS) components. |
| [LicenseFilePrintData] | Contains entries for installing licensing information for Windows Server 2003. |

| | |
|---|---|
| [MassStorageDrivers] | Contains an entry for loading retail or original equipment manufacturer (OEM) small computer system interface (SCSI) drivers in the Setup Loader during the intermediate (text-mode) stage of Setup |
| [*MS_AppleTalk parameters*] | Contains entries for configuring AppleTalk. |
| [*MS_ATMArps parameters*] | Contains no entries. |
| [*MS_ATMLANE parameters*] | Contains no entries. |
| [*MS_ATMUni parameters*] | Contains no entries. |
| [*MS_L2TP parameters*] | Contains entries for the Layer 2 Tunneling Protocol (L2TP). |
| [*MS_MSClient parameters*] | Contains entries for Client for Microsoft Networks. |
| [*MS_NetMon parameters*] | Contains no entries. |
| [*MS_NWClient parameters*] | Contains entries for Client Service for NetWare. |
| [*MS_NWIPX parameters*] | Contains entries for the Internetwork Packet Exchange (IPX) protocol. Its parameters are network adapter–independent. |
| [*MS_NwSapAgent parameters*] | Contains no entries. |
| [*MS_PPTP parameters*] | Contains entries for the Point-to-Point Tunneling Protocol (PPTP). |
| [*MS_Psched parameters*] | Contains no entries. |
| [*MS_RAS parameters*] | Contains entries for the Routing and Remote Access service. |
| [*MS_RasSrv parameters*] | Contains entries for the Routing and Remote Access service for remote access servers. |
| [*MS_Server parameters*] | Contains entries for the File and Print Service. |
| [*MS_TCPIP parameters*] | Contains entries for configuring the TCP/IP protocol. |
| [*MS_WLBS parameters*] | Contains entries for configuring Network Load Balancing (formerly, Windows Load Balancing Service, or WLBS) and binding Network Load Balancing to multiple network adapters. |
| [NetAdapters] | Contains entries for network adapters that you want to install. |
| [NetBindings] | Contains entries for disabling, enabling, and demoting binding rules for network components such as services, protocols, and adapters. |
| [NetClients] | Contains an entry for specifying the network clients to install. |
| [NetOptionalComponents] | Contains entries for installing optional network components. |
| [NetProtocols] | Contains an entry for installing the network protocols. |
| [NetServices] | Contains entries for installing network services. |
| [Networking] | Contains no entries. However, to configure network protocols in unattended installation, you must include the [Networking] section name in your answer file. |
| [OEMBootFiles] | Contains entries for specifying OEM-supplied boot files. |
| [*params.adapter_instance*] | Contains entries for specifying an adapter's instance on a computer by using either PCI location information or a Plug and Play device ID. |
| [PCHealth] | Contains entries for configuring error reporting, Remote Assistance, and Help and Support Center. |

| | |
|---|---|
| [*protocol_parameters_section*] | Contains an entry for specifying network protocol parameters that are also specified in the [NetProtocols] section. |
| [Proxy] | Contains entries for specifying proxy server settings for the Windows XP or Windows Server 2003 browser. |
| [RegionalSettings] | Contains entries for installing regional options during the setup process. |
| [SetupMgr] | Contains entries for specifying the server and folder in which the distribution share resides and the name of the distribution share. This section is used only by Setup Manager when an existing distribution share is edited. It is not used by Setup during an unattended installation. |
| [SetupParams] | Contains an entry for running an additional post-Setup command. |
| [Shell] | Contains entries for specifying the visual style of the user interface. |
| [SNMP] | Contains entries for specifying the Simple Network Management Protocol (SNMP) service. |
| [SystemFileProtection] | Contains entries for specifying Windows File Protection (WFP) settings. |
| [SystemRestore] | Contains entries for specifying the System Restore feature. |
| [TapiLocation] | Contains entries for specifying the telephony API (TAPI). It is valid only if a modem is present on the computer. |
| [TerminalServices] | Contains entries for configuring Terminal Services features (Terminal Server and Remote Desktop). |
| [Unattended] | Contains entries for running Setup. |
| [Uninstall] | Contains entries for allowing users who upgrade to Microsoft Windows XP to remove the operating system if they encounter problems. |
| [URL] | Contains entries for specifying default Uniform Resource Locator (URL) settings for the browser available in Windows XP or Windows Server 2003. |
| [UserData] | Contains entries for specifying user settings during the setup process. |
| [WindowsFirewall] | This section must be included in your answer file to configure firewall sections and entries in unattended installations. To configure Windows Firewall settings, you must include the [WindowsFirewall] section along with one or more of the other Windows Firewall settings sections. |
| [WindowsFirewall.*profile_name*] | Contains entries for specifying a domain type of profile, a standard type, or a profile that applies to both domain and standard. |
| [WindowsFirewall.*program_name*] | Contains entries for adding programs to the Windows Firewall Exceptions list. |
| [WindowsFirewall.*service_name*] | Contains entries for adding services to the Windows Firewall Exceptions list. |
| [WindowsFirewall.*portopening_name*] | Contains entries for adding static ports to the Windows Firewall Exceptions list. |
| [WindowsFirewall.*icmpsetting_name*] | The default configuration for Windows Firewall blocks all Internet Control Message Protocol (ICMP) message types, the entries in this section allows you to modify this behavior. |
| [Win9xUpg] | Contains entries for specifying an unattended upgrade from Windows 98, Windows 98 Second Edition, or Windows Millennium Edition to either Windows XP Home Edition or Windows XP Professional. |

You do not need to add every possible section and entry to an answer file. Most sections and entries are optional, and Setup ignores optional sections and entries that are not present in an answer file.

**Entries for Fully-Automated Unattended Installation**
If you want to perform a fully-automated unattended installation, you must configure the following sections and entries in your answer file. If you do not provide values for these entries in your answer file, Setup will display a prompt asking for the values.

**[GuiUnattended]**
You must specify values for the AdminPassword and TimeZone entries. The value for AdminPassword cannot begin with an asterisk (*). Using a password that begins with an asterisk can cause the password to be set to a null value, which can be a security risk.

**[Identification]**
You must specify values for the JoinDomain, DomainAdmin, and DomainAdminPassword entries. However, if you are concerned about security you might want to leave the DomainAdminPassword blank because this entry cannot be encrypted in the answer file. If you leave the DomainAdminPassword entry blank, you will be prompted for the password during the unattended installation.

**[LicenseFilePrintData]**
You must specify values for the AutoMode and AutoUsers entries if you are installing Windows Server 2003.

**[Networking]**
If your destination computer requires network connectivity, you must specify values for various network protocol entries.

**[Unattended]**
You must specify values for the UnattendMode and TargetPath entries.

**[UserData]**
You must specify values for the FullName and ComputerName entries.

**Entries for Unattended Upgrade**
If you are performing an unattended upgrade, you can use only the following answer file sections and entries. All other sections and entries are ignored during an unattended upgrade.

**[Unattended]**
You must specify values for either the NtUpgrade or the Win9xUpgrade entry. You must use the entry `OEMPreinstall = No`, which will prevent the files and folders in the $OEM$ folder in your distribution share from being copied to the destination computer. You can also set values for the AutoActivate, DUDisable, DUShare, and DUStopOnError entries.

**[UserData]**
You can set a value for the ProductKey entry. If you do not, the administrator or technician will be prompted for a product key.

In addition, your answer file must reside on a floppy disk and you must use the product CD instead of a distribution share.

For more information about answer file sections, including descriptions and examples of entries and values, see "Unattended Installation Tools and Settings."

**Unattended Installation Processes and Interactions**
Because there are numerous different ways to perform an unattended installation, it is impractical to describe all of the processes and interactions that occur for each method and variation. Therefore, this section describes the processes and interactions that occur during the two most commonly used methods of performing an unattended installation: a fully-automated unattended installation when you start the destination computer from a 32-bit operating system, and a fully-automated unattended installation when you start the destination computer from the product CD.

In addition, this section describes log file interactions that occur during an unattended installation as well as the processes that occur when you perform post-installation tasks by using a Cmdlines.txt file or the [GUIRunOnce] section of an answer file.

**Unattended Installation When You Start the Computer from a 32-bit Operating System**
The following figure shows the processes and interactions that occur during a fully-automated unattended installation of Windows Server 2003 when you start the destination computer from a 32-bit operating system and you copy the installation files from a distribution share. The 32-bit operating system can be any of the following: Windows XP, Windows NT, Windows 2000, and Windows Server 2003. The distribution share can be saved on any accessible server that you can secure with file and folder permissions.

### Note

- You can also copy the installation files from the product CD; however, it is more common to use a distribution share when you start the destination computer from a 32-bit operating system.

**Unattended Installation from a 32-bit Operating System Using a Distribution Share**




This figure is based on the following requirements:

- The destination computer is running Windows 2000 Server (or an equivalent 32-bit Windows operating system).
- The distribution share is named Win2003_Std and is saved in the OS_Shares folder on the distribution server.
- The answer file is named Unattend.txt and is saved in the distribution share.
- All of the installation files, including Winnt32.exe, are saved in the I386 folder in the distribution share.
- The unattended installation is a clean installation of the operating system (not an upgrade).
- The unattended installation is started from the existing Windows 2000 Server operating system by running Winnt32.exe at the command line.

In addition, the following answer file is used to perform this installation:

```
 [Unattended]
    UnattendMode=FullUnattended
    OemSkipEula=Yes
    OemPreinstall=Yes
    FileSystem = LeaveAlone
    TargetPath= *

[GuiUnattended]
    AdminPassword=ca8d43d0eae56d1019ac6cc7a3608820ae698f3d8ce553674c94fb185
    EncryptedAdminPassword=Yes
    OEMSkipRegional=1
    TimeZone=004
    OemSkipWelcome=1

[UserData]
    ProductKey=#####-#####-#####-#####-#####
    FullName="My Name"
    OrgName="Microsoft Corporation"
    ComputerName=Server2

[Display]
    BitsPerPel=24
    Xresolution=800
    YResolution=600
    Vrefresh=70

[LicenseFilePrintData]
    AutoMode=PerServer
    AutoUsers=5

[TapiLocation]
    CountryCode=1
```

```
    Dialing=Tone
    AreaCode=425
    LongDistanceAccess="9"

[SetupMgr]
    DistFolder=\\Server1\OS_Shares
    DistShare=Win2003_Std

[Identification]
    JoinDomain=MyDomain
    DomainAdmin=MyAdmin
    DomainAdminPassword=MyPassword

[Networking]
    InstallDefaultComponents=Yes
```

**Starting the Unattended Installation**
The processes and interactions described in this section are valid only when you start the installation by mapping a network drive to the distribution share from the destination computer, navigating to the distribution share, and running Winnt32.exe with the `/unattend` parameter. You need to map a network drive to the distribution share because all of the installation files, including Winnt32.exe, and the answer file, Unattend.txt, are saved in the distribution share.

**Initial Stage**
After you start an unattended installation from a 32-bit operating system, the initial stage (also known as data-collection stage) of unattended installation begins. The following actions occur during the initial stage of an unattended installation:

**Assess current operating system**
The current operating system configuration is assessed, including operating system version, available disk space, hardware, and installed components.

**Process answer file**
The answer file is parsed and the results are saved in a new file named Winnt.sif. This file contains all of the settings that are in the answer file as well as other information that was gathered during the system assessment.

**Begin Welcome to Windows Setup**
The Welcome to Windows Setup dialog boxes appear briefly and then close. During an interactive setup this sequence of dialog boxes prompts you for several configuration settings, including installation type (upgrade or clean installation), product key, acceptance of the End-User License Agreement (EULA), installation folder, location of installation files, language options, accessibility options, and dynamic update options. Because this information is specified in the answer file, the introductory dialog box stays open for only a few seconds and then closes; other dialog boxes are not displayed.

**Copy files to temporary folders**
During the initial stage, Setup creates two temporary folders on the destination computer's hard disk — $win_nt$.~ls and $win_nt$.~bt. The $win_nt$.~ls folder (known as the local source folder) contains the installation files, which are copied from the $OEM$ folder in the distribution share.

The $win_nt.~bt folder (known as the boot folder) contains all of the files that are necessary for starting the computer and initiating the setup process after all of the installation files are copied to the destination computer's hard disk. These files are copied from the distribution share and include basic device drivers, such as mouse, keyboard, mass storage, and video drivers. The $win_nt.~bt folder also contains Txtsetup.sif and Winnt.sif. Txtsetup.sif contains definitions and processing instructions for the files and folders that are installed during the setup process. Winnt.sif contains all of the answer file settings as well as other information that was gathered during system assessment.

A file known as $ldr$ is also created during the initial stage. This file is similar to Ntldr and is used to start the computer in a special setup state during which the intermediate stage of unattended installation occurs (this stage is also known as the text-mode stage).

**Load boot sector and change boot.ini**
The partition boot sector is changed so that it contains the master boot code for starting the computer with $ldr$. The Boot.ini file is also modified so that the computer uses the startup files in the $Win_nt$.~bt folder, which start the computer in a special setup state during which the intermediate stage of Setup occurs.

At the end of initial stage, when all of the installation files have been copied to the destination computer's hard disk, the destination computer restarts and the intermediate stage of the unattended installation begins.

**Intermediate Stage**
The following actions occur during the intermediate stage (also known as the text-mode stage) of an unattended installation from a 32-bit operating system:

**Start the Computer in a Special Setup State**
The computer starts by running $ldr$ and the files that are in $Win_nt$.~bt. During this stage of an interactive setup you would be prompted for information about partitions, file system formats, and system folder name and location. Because it is not possible to configure partitions or file formats when you start an installation from an existing operating system, and all other configuration information is contained in the answer file, these prompts do not appear.

**Build directory structure**
Setup creates the system folder structure and copies files from the temporary installation folder ($win_nt$.~ls) to the system folders. For example, for a default installation, Setup creates the Windows folder and all subfolders (Inf, Help, System, System32, and so forth). The Txtsetup.sif file defines which folders are created and which files are copied to which folders. In addition, Winnt.sif, the file that contains the original answer file settings and other operating system information, is renamed $winnt$.inf and saved in the *systemroot*\System32 folder. Finally, the default registry hive files are created in the *systemroot*\System32\Config folder.

### Note

- Not all files are copied from the temporary installation folders to the system folders during the intermediate stage of the installation. Many of the files for optional components are installed during the final stage of Setup instead. This makes the installation process more efficient.

Load boot sector and change boot.ini

The partition boot sector is changed again so that it contains the master boot code for starting the computer with Ntldr instead of $ldr$. The Boot.ini file is also modified so that the computer uses the Windows Server 2003 startup files instead of the startup files in the $Win_nt$.~bt folder.

At the end of the intermediate stage, when all of the installation files have been copied from the temporary folders to the system folders, the destination computer restarts and the final stage of the unattended installation begins.

**Final Stage**
The following actions occur during the final stage (also known as the GUI-mode stage) of an unattended installation from a 32-bit operating system:

**Start Windows Server 2003**
The destination computer starts by using the Windows Server 2003 startup files and device drivers instead of the temporary startup files in $Win_nt$.~bt. For example, the computer uses Ntldr instead of $ldr$, and the computer uses the startup files and device drivers in the systemroot folder instead of the startup files and device drivers in the $Win_nt$.~bt folder. As with an interactive setup, the hardware detection and installation process also occurs.

**Install and configure optional components and settings**
Optional components are installed (or not installed) and optional settings are configured based on the settings in the answer file. Examples of optional components and settings include regional and language options, text services and input languages, user name, company name, licensing mode, computer name, administrator account password, time zone settings, network adapter settings, protocol settings, and file and printer sharing. During an interactive setup you would be prompted for the installation information and the configuration settings, but these dialog boxes do not appear during a fully-automated unattended installation.

**Copy files from temporary folders**
Based on the answer file settings for the final stage of the installation process, additional files might need to be copied from the temporary installation folders to the system folders. As with an interactive setup, a dialog box shows the progress of this file-copy process; however, no administrator or technician interaction is necessary or possible.

**Complete the installation**
Windows performs the final configuration tasks. This includes configuring start menu items, registering components, deleting temporary installation folders, and running any commands specified in the Cmdlines.txt file.

At the end of the final stage, all of the required installation tasks have been completed and the destination computer is ready to run Windows Server 2003. If the answer file contains any instructions for performing post-installation tasks, Windows then restarts the destination computer, logs on with the Administrator account, and performs the post-installation tasks that are specified in the answer file.

**Unattended Installation When You Start the Computer from a Product CD**
The following figure shows the processes and interactions that occur during a fully-automated unattended installation when you start the destination computer from a product CD.

**Unattended Installation from a Product CD**




This figure is based on the following requirements:

- The destination computer is running Windows 2000 Server.
- The answer file is named Winnt.sif and is saved on a floppy disk.
- The unattended installation is a clean installation of the operating system (not an upgrade).
- The primary partition on the hard disk is reformatted during the unattended installation.
- The unattended installation is started by starting the destination computer from the product CD.
- The unattended installation is fully automated (no administrator or technician interaction required).

In addition, the following answer file is used to perform this installation:

```
 [Data]
    AutoPartition=1
    MsDosInitiated="0"
    UnattendedInstall="Yes"

[Unattended]
    UnattendMode=FullUnattended
    OemSkipEula=Yes
    OemPreinstall=No
    TargetPath=\WINDOWS
    Repartition = Yes

[GuiUnattended]
    AdminPassword=ca8d43d0eae56d10197a3608820ae698fa3254ef3d8ce553674c94fb185
    EncryptedAdminPassword=Yes
    OEMSkipRegional=1
    TimeZone=4
    OemSkipWelcome=1

[UserData]
    ProductKey=#####-#####-#####-#####-#####
    FullName="MyName"
    OrgName="MyCompany"
    ComputerName=MyComputer

[Display]
    BitsPerPel=24
    Xresolution=800
    YResolution=600
    Vrefresh=70

[LicenseFilePrintData]
    AutoMode=PerServer
    AutoUsers=5
```

```
[TapiLocation]
    CountryCode=1
    Dialing=Tone
    AreaCode=425
    LongDistanceAccess="9"

[Identification]
    JoinDomain=MyDomain
    DomainAdmin=MyAdmin
    DomainAdminPassword=MyPassword

[Networking]
    InstallDefaultComponents=Yes
```

**Starting the Unattended Installation from the Product CD**
The processes and interactions described in this section are valid only when you start the destination computer from the product CD and quickly insert the floppy disk on which you saved your answer file into the floppy disk drive. You must insert the floppy disk immediately after the destination computer starts from the product CD because the Setup program immediately searches the floppy disk drive for a Winnt.sif file. If the Setup program does not find a floppy disk in the floppy disk drive, it assumes you are performing an interactive setup. When Setup finds a Winnt.sif file, Setup reads the Winnt.sif and initiates an unattended installation.

**Initial and Intermediate Stages**
After you start an unattended installation from the product CD, the initial and intermediate stages of unattended installation begin (these stages are also known as the data-collection and text-mode stages). Unlike in most installations, the initial and intermediate stages are not separate processes during an unattended installation from the product CD. Instead, these stages are combined into a single process and the destination computer does not restart between them.

The following actions occur during the initial and intermediate stages of an unattended installation from a product CD:

**Begin the setup process**
Setup loads the files that are necessary to start the computer from the product CD. During this process a text-based Setup dialog box provides information about the Setup files that are loading. This is the same dialog box that appears at the beginning of an interactive setup. Usually, during an interactive setup from a product CD, several text-based dialog boxes appear after this initial Setup dialog box. For example, these dialog boxes ask you whether you are repairing an existing installation, whether you want to overwrite an existing operating system, and whether you want to use an existing partition or create a new partition. Because all of this information is contained in the answer file, these dialog boxes do not appear during an unattended installation from the product CD.

### Note

- The Dynamic Update feature cannot be used during an unattended installation from the product CD because there is no network connectivity when you start a computer from the product CD.

**Process answer file**
After the Setup files are loaded, the answer file is parsed and the results are saved in a new file named Winnt.sif. This file contains all of the settings that are in the answer file. A dialog box notifies you when the answer file is being processed, but it requires no administrator or technician interaction because you are performing a fully-automated unattended installation.

**Format and repartition**
The primary partition is deleted and a new partition is created and formatted based on the information in the answer file. A dialog box notifies you when the partition is being created and the status of the formatting process, but the dialog box requires no administrator or technician interaction because you are performing a fully-automated unattended installation.

**Load boot sector**
The partition boot sector is changed so that it contains the master boot code for Windows Server 2003. This small executable code initiates the startup process and instructs the computer where to find Ntldr.

Build directory structure

Setup creates the system folder structure on the hard disk. For example, for a default installation, Setup creates the Windows folder and all subfolders (Inf, Help, System, System32, and so on). The Txtsetup.sif file defines which folders are created.

**Copy files to system folders**
During the file-copy process Setup copies files from the I386 folder on the product CD to the system folders on the hard disk. The Txtsetup.sif file contains definitions and processing instructions for the files and folders that are copied. In addition,

Winnt.sif, the file that contains the original answer file settings, is renamed $winnt$.inf and saved in the *systemroot*\System32 folder.

At the end of initial and intermediate stage, when all of the installation files have been copied to the destination computer's hard disk, the destination computer restarts and the final stage of the unattended installation begins.

**Final Stage**
The following actions occur during the final stage (also known as the GUI-mode stage) of an unattended installation from a product CD:

**Start Windows Server 2003**
Windows Server 2003 starts by using the Ntldr file (not Setupldr.bin) and all of the necessary startup files and device drivers that are in the system folder. As with an interactive setup, the hardware detection and installation process occurs.

**Install and configure optional components and settings**
Optional components are installed (or not installed) and optional settings are configured based on the settings in the answer file. Examples of optional components and settings include regional and language options, text services and input languages, user name, company name, licensing mode, computer name, administrator account password, time zone settings, network adapter settings, protocol settings, and file and printer sharing. During an interactive setup you would be prompted for the installation information and the configuration settings, but these dialog boxes do not appear during a fully-automated unattended installation.

**Copy files from the product CD**
Based on the answer file settings for the final stage of the installation process, additional files might need to be copied from the product CD to the system folders. As with an interactive setup, a dialog box shows the progress of this file-copy process; however, no administrator or technician interaction is necessary or possible.

**Complete the installation**
Setup performs the final configuration tasks. This includes configuring start menu items, registering components, deleting temporary files, and running any commands specified in the Cmdlines.txt file.

At the end of the final stage, all of the required installation tasks have been completed and the destination computer is ready to run Windows Server 2003. If the answer file contains any instructions for performing post-installation tasks, Setup then restarts the destination computer, logs on with the Administrator account, and performs the post-installation tasks that are specified in the answer file.

**Post Installation Processes**
Several unattended installation processes can occur after an operating system is installed on a destination computer. These processes occur when you automate post-installation configuration tasks. Examples of post-installation configuration tasks include running command-line programs to configure system services, running scripts to configure Group Policy settings, and installing software programs. The most commonly used mechanisms for automating post-installation tasks is the Cmdlines.txt file, the `winnt32` command with the `/cmd` parameter, and the [GUIRunOnce] and [SetupParams] sections of the unattended installation answer file.

### Note

- Several answer file sections contain entries that are not configured during the intermediate or final stages of an unattended installation. These settings are configured after the operating system is installed on the destination computer and the destination computer has restarted. The answer file sections include [DCInstall], [CertSrv_Client], and [CertSrv_Server].

**Cmdlines.txt Processes and Interactions**
A Cmdlines.txt file contains a list of commands, programs, scripts, or batch files that run just after the final stage of Setup finishes but before the destination computer restarts. By default, Cmdlines.txt does not exist, so you must manually create the file with a text editor, such as Notepad.

The syntax for Cmdlines.txt is as follows:

```
              [Commands]
command_1command_2
     .
     .
command x
```

Where *command_1*, *command_2*, and *command_x* refer to the commands, programs, scripts, or batch files that you want to

run when final stage of Setup is complete. Note that all commands must be enclosed in quotation marks. Also, if you are using the command window (Cmd.exe) to run commands, programs, scripts, or batch files, then you need to use the `/c` parameter with the `cmd` command. For example, to create a new folder named Test on drive C, you type:

```
cmd /c mkdir c:\Test
```

In addition, to use a Cmdlines.txt file, you need to:

- Add the following section, entry, and value to your answer file:

```
[Unattended]
OemPreinstall = Yes
```

- Copy all of the programs, scripts, and batch files that are listed in your Cmdlines.txt file to the $OEM$ folder in your distribution share.

- Save the Cmdlines.txt file in the $OEM$ folder in your distribution share.

When you use the entry `OemPreinstall = Yes` in the [Unattended] section of your answer file, Setup searches the $OEM$ folder in your distribution share for a Cmdlines.txt file. If Setup finds a Cmdlines.txt file, then the current path is changed to c:\$win_nt$.~ls\$OEM$. This is why you need to put all programs, scripts, or batch files that are listed in Cmdlines.txt into the $OEM$ folder in your distribution share.

You can only have one Cmdlines.txt file, but you can use it to run multiple commands, programs, scripts, or batch files. When these commands, programs, scripts, or batch files run, the following processes and interactions occur:

- The commands, programs, scripts, and batch files run synchronously, which means each command, program, script, or batch file runs to completion before the next one starts running.

- All commands, programs, scripts, and batch files specified in a Cmdlines.txt run under the Local System security account.

- Network connectivity is not available when the commands, programs, scripts, and batch files listed in a Cmdlines.txt file run.

- Windows Installer packages (.msi files) cannot be run from a Cmdlines.txt file.

**Note**

- You cannot use Cmdlines.txt if you are using a product CD and a Winnt.sif file to perform an unattended installation. You can use Cmdlines.txt only if you are installing from a distribution share.

**[GuiRunOnce] Processes and Interactions**

The [GuiRunOnce] section of an answer file contains a list of commands, programs, scripts, or batch files that run the first time an administrator or technician logs on to the computer after the operating system is installed. Because the [GuiRunOnce] section does not exist in a default answer file built by Setup Manager, you must add it manually by editing the answer file with a text editor, such as Notepad.

The syntax for [GuiRunOnce] is:

```
[GuiRunOnce]
command_1command_2
      .
      .
command_x
```

Where *command_1*, *command_2*, and *command_x* refer to the commands, programs, scripts, or batch files that you want to run after the operating system is installed and the administrator or technician logs on. Note that all commands must be

enclosed in quotation marks.

The commands that are listed in the [GuiRunOnce] section are stored in the following registry location:

HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Runonce

When an administrator or technician logs onto the destination computer for the first time after the unattended installation is complete, Windows Server 2003 reads and then runs the commands that are listed in this registry location. When these commands, programs, scripts, or batch files run, the following processes and interactions occur:

- The commands, programs, scripts, and batch files run synchronously, which means each command, program, script, or batch file runs to completion before the next one starts running.
- The commands, programs, scripts, and batch files listed in the [GuiRunOnce] section run under the security account of the user who logged onto the destination computer (usually an administrator or a technician).
- Network connectivity is available when the commands, programs, scripts, and batch files run.
- You can use Windows Installer packages (.msi files) to install applications (unlike Cmdlines.txt).
- You can run commands, programs, scripts, and batch files from hard disk drives, CD drives, shared folders on the network, or other storage devices (unlike Cmdlines.txt, which runs everything from the $OEM$ folder in your distribution share).

In addition, there are some interactions and processes that you might need to control if you run certain commands, programs, scripts, and batch files with the [GuiRunOnce] section.

**Prevent restarts**
You can run commands, programs, scripts, and batch files only if you can prevent the command, program, script, or batch file from restarting the computer. When a computer restarts, all remaining entries in the [GuiRunOnce] section are lost because the Runonce registry entry is deleted. If the system restarts before completing entries listed in the [GuiRunOnce] section, the remaining items will not run. Therefore, you need to suppress restarts. If you cannot suppress a restart within a command, program, script, or batch file, you can try to repackage the command, program, script, or batch file into a Windows Installer package.

**Do not rely on the Windows Explorer shell**
The Windows Explorer shell is not loaded when the operating system starts running commands, programs, scripts, or batch files that are listed in the [GuiRunOnce] section of an answer file. In some cases, you can get an updated installer program from the application vendor that does not rely on the Windows Explorer shell. If you cannot, you might be able to repackage the command, program, script, or batch file as a Windows Installer package (.msi file).

**Control the processes by using the /wait parameter**
Commands, programs, scripts, and batch files often start and stop several different processes. In some cases, when you are running multiple programs — especially installation programs — this can inadvertently cause the next program listed in the [GuiRunOnce] section to start before the previous program is finished running. When this occurs, the second program usually fails. To prevent this, you can run the programs from batch files by using the `start` command with the `/wait` parameter. This forces each program to run to completion before the next command listed in [GuiRunOnce] runs.

**Other Post-Installation Processes and Interactions**
In addition to the Cmdlines.txt file and the [GuiRunOnce] section of an answer file, two other mechanisms can create post-installation processes and interactions during an unattended installation. These mechanisms are the [SetupParams] section of an answer file, and the `winnt32` command with the `/cmd` parameter.

**[SetupParams]**
The [SetupParams] section of an answer file contains a single command, program, script, or batch file that runs immediately after the commands, programs, scripts, or batch files listed in the Cmdlines.txt file run. You typically use this answer file section and entry to perform a post-installation task on only a few computers without changing the Cmdlines.txt file. Because the [SetupParams] section does not exist in a default answer file built by Setup Manager, you must add it manually by editing the answer file with a text editor, such as Notepad.

The syntax for the [SetupParams] section is:

```
[SetupParams]
```

```
UserExecute =path_and_file_name_or_command
```

Where *path_and_file_name_or_command* refers to the command, program, script, or batch file that you want to run. Note that the *path_and_file_name_or_command* must be enclosed in quotation marks. Also, you can omit the path portion of *path_and_file_name_or_command* only if the application is in the systemroot or *systemroot*\System32 folder or search path.

**Winnt32 /cmd**

The `winnt32` command with the `/cmd` parameter specifies a single command, program, script, or batch file that runs immediately after the commands, programs, scripts, or batch files listed in the Cmdlines.txt file run. You typically use this answer file section and entry to perform a post-installation task on only a few computers without changing the Cmdlines.txt file.

The syntax for the `winnt32` command with the `/cmd` parameter is:

```
Winnt32 /cmd:
path_and_file_name_or_command
```

Where *path_and_file_name_or_command* refers to the command, program, script, or batch file that you want to run.

**Log File Interactions During an Unattended Installation**

Several log files are created during an unattended installation and saved in the systemroot folder. The following table provides the names, locations, and descriptions of the Setup log files. The descriptions provide a summary of the contents of each log file when the logging level is set to its default state. You can increase the logging level (increase log file details) for the Winnt32.log and Setupact.log files by using the `winnt32` command with the `/debug` parameter. By default, the logging level is set to 2, which means severe errors, common errors, and warnings are included in the log files. By increasing the logging level you can include basic setup information and detailed debugging information in the Winnt32.log and Setupact.log log files.

**Setup Log Files for an Unattended Installation**

| Log File Name | Location | Description |
|---|---|---|
| Setupact.log | Systemroot | Log file for the initial (data-collection) and final ( GUI-mode) stage of unattended installation. Contains information about detected hardware, installation requirements, and a list of device driver files that are copied from the product CD or $Win_nt$.~ls to various system folders. |
| Setupapi.log | Systemroot | Log file for Setupapi.dll, which processes .inf files during hardware installation and registration. |
| Setuperr.lot | Systemroot | Log file for errors that occur during the initial (data -collection) and final (GUI-mode) stages of unattended installation. This log file consolidates Setup errors; the errors that appear in this log file also appear in other log files. |
| Setuplog.txt | Systemroot | Log file for the initial (data-collection) and final (GUI-mode) stages of unattended installation, as well as some post-installation tasks that occur after Setup completes and the destination computer restarts. Contains information about the files that are copied and deleted, the settings that are configured, and the security settings that are applied to the destination computer. |
| Winnt32.log | Systemroot | Log file for the initial (data-collection) stage of the setup process. Contains information about the existing operating system, hard disk size, Boot.ini configuration, and a list of the files that are copied to the temporary installation folders $Win_nt$.~bt and $Win_nt$.~ls. |

Setupapi.log and Setuplog.txt are not overwritten if you restart the unattended installation process: that is, information is appended to the log file each time you prematurely interrupt and restart the unattended installation process. All of the other log files are overwritten every time the unattended installation process is prematurely interrupted and restarted.

**Note**

- No log files are generated during the intermediate (text-mode) stage of Setup.

In addition to these log files, three other log files contain relevant information about the unattended installation process. The following table provides the names, locations, and descriptions of these log files.

**Supplemental Log Files for an Unattended Installation**

| Log File Name | Location | Description |
|---|---|---|
| Wsdu.log | Systemroot | Log file for Dynamic Update. This is generated only if you use Dynamic Update during an unattended installation, which you can only do if you are performing an unattended installation from an existing 32-bit operating system such as Windows NT, Windows 2000, or Windows Server 2003. |
| Dberr.txt | *systemroot*\System32\CatRoot2 | Log file for catalog file registrations. Contains information about catalog files that are not registered, and exists only if any nonregistered catalogs are identified during unattended installation. If a catalog is not registered, then the files in the catalog file are considered unsigned and will generate errors in Setuperr.log. |
| Scesetup.log | *systemroot*\security\logs | Log file for security settings. |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Setup Technical Reference
- Sysprep Technical Reference
- Remote Installation Services Technical Reference

# Unattended Installation Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Unattended Installation Tools
- Unattended Installation Answer File Settings


The following tools and answer file settings are associated with unattended installations. You use the tools and answer file settings to perform unattended installations.

**Unattended Installation Tools**
The following tools are associated with unattended installations.

**Winnt.exe: Windows Setup program for 16-bit operating systems**
**Category**
Windows setup and installation tool available on every Windows Server 2003 product CD.

**Version compatibility**
Runs on MS-DOS and Windows for Workgroups.

You can use this tool to perform an unattended installation from the 16-bit MS-DOS and Windows for Workgroups operating systems. Typically, you use Winnt.exe to perform an unattended installation on a computer that does not have an operating system and you start the computer with an MS-DOS startup floppy disk. You cannot perform an upgrade with Winnt.exe; you can only perform clean installations with Winnt.exe.

For more information about Winnt.exe, see "Command Line References" in the Tools and Settings Collection.

**Winnt32exe: Windows Setup program for 32-bit and 64-bit operating systems**
**Category**
Windows setup and installation tool available on every Windows Server 2003 product CD.

**Version compatibility**
Runs on Windows Server 2003, Windows XP, Windows 2000, Windows NT, Windows Millennium Edition (Me), Windows 98, and Windows 95.

You can use this tool to perform an unattended installation or upgrade from a Windows 32-bit operating system by using the 32-bit version of Winnt32.exe and from 64-bit operating systems by using the 64-bit version of Winnt32.exe. Typically, you use Winnt32.exe to perform an unattended installation or upgrade on a computer that has an operating system; you start the computer from the existing operating system and then run Winnt32.exe from the command line.

For more information about Winnt32.exe, see "Command Line References" in the Tools and Settings Collection.

**Setupmgr.exe: Setup Manager**
**Category**
Windows deployment tool available on the Windows Server 2003 product CD.

**Version compatibility**
Runs on the 32-bit versions of Windows Server 2003 and Windows XP.

Setup Manager is a wizard-like program that helps you configure distribution shares and answer files that can be used to perform automated installations, including unattended installation, image-based installation with Sysprep, and Remote Installation Services (RIS) installations. Setup Manager prompts you for information in a series of dialog boxes. Setup Manager configures answer file settings and distribution shares based on the information you provide in these dialog boxes.

Usually, you run Setup Manager on the administrator or technician computer. Setup Manager does not affect the system files of the technician computer on which the tool runs and does not actually run Setup. Also, you do not need to specify all of your customizations with Setup Manager. Instead, you can create an answer file in Setup Manager and then manually edit the file with Notepad. If you later open the answer file in Setup Manager, the tool displays your changed settings.

Setup Manager is not installed when you install Windows XP or Windows Server 2003. You must manually install Setup Manager on the administrator or technician computer.

**To install Setup Manager**

1. Insert a Windows XP or Windows Server 2003 product CD into your CD drive and open the \Support\Tools folder.

2. Double-click **Deploy.cab**.

3. Right-click the file named **Setupmgr.exe** and click **Extract**.

4. In the **Select a Destination** dialog box, select a location on your hard drive where you want to save Setup Manager, and then click **Extract**.

You can run Setup Manager from the graphical user interface by double-clicking the Setupmgr.exe file that you saved to your hard disk, or from the command line by typing `setupmgr`.

**Note**

- If you use Setup Manager to create or modify a distribution share that is not on the local computer, be sure you run Setup Manager under a user account that has the minimum user rights to connect to the remote computer. This is because Setup Manager sends the user name and the password across the network in plaintext when it connects to a remote distribution share.

The following sections provide information about each Setup Manager dialog box. The section names correspond to the dialog box name that you see in Setup Manager.

**New or Existing Answer File**
Enables you to create a new answer file or edit an existing one.

Answer file sections and entries configured:

- None

**Type of Setup**
Provides three Setup options: Unattended, Sysprep, or Remote Installation Services (RIS).

Answer file sections and entries configured:

- None

**Product**
Prompts you to select the Windows product that you want to install.

Answer file sections and entries configured:

- None

**User Interaction**
Prompts you to select the level of user interaction you want. This dialog box is not used if you are configuring a Sysprep or RIS installation.

Answer file sections and entries configured:

- [Unattended]

- UnattendMode

**Distribution Share**
Enables you to create a new distribution share, modify an existing one, or set up Windows from a product CD.

Answer file sections and entries configured:

- None

**Location of Setup Files**
Prompts you for the location of the Setup files you want to copy.

Answer file sections and entries configured:

- None

**Distribution Share**
Prompts you for the location of the distribution share you want to create or modify. This section is used only by Setup Manager when an existing distribution share is edited. It is not used by Setup during an unattended installation.

Answer file sections and entries configured:

- [SetupMgr]
- DistFolder
- DistShare

**Name and Organization**
Enables you to provide a default user name and organization name.

Answer file sections and entries configured:

- [UserData]
- FullName
- OrgName

**Display Settings**
Enables you to change the change settings for colors, screen area, and refresh frequency.

Answer file sections and entries configured:

- [Display]
- BitsPerPel
- XResolution
- YResolution
- VRefresh

**Time Zone**
Prompts you to select a default time zone.

Answer file sections and entries configured:

- [GuiUnattended]
- TimeZone

**Product Key**
Prompts you to provide a Product Key to identify your copy of Windows.

Answer file sections and entries configured:

- [UserData]
- ProductKey

**Computer Names**
Prompts you to provide names for the destination computers.

Answer file sections and entries configured:

- [UserData]
- ComputerName

**Administrator Password**
Enables you to specify a password for the Administrator account on all destination computers.

Answer file sections and entries configured:

- [GuiUnattended]
- AdminPassword
- EncryptedAdminPassword

**Networking Computers**
Enables you to choose typical or custom network settings.

Answer file sections and entries configured:

- [Networking]
- [NetAdapters] (if custom settings are configured)
- [NetClients] (if custom settings are configured)
- [NetProtocols] (if custom settings are configured)

**Workgroup or Domain**
Prompts you to specify how the destination computers will participate in a network.

Answer file sections and entries configured:

- [Identification]

**Windows Components**
Enables you to install commonly used Windows components.

Answer file sections and entries configured:

- [Components]

**Telephony**
Enables you to provide telephony information for the user.

Answer file sections and entries configured:

- [TapiLocation]

**Regional Settings**
Enables you to customize Windows for different regions and languages.

Answer file sections and entries configured:

- [RegionalSettings]

**Languages**
Prompts you to select the language groups you want to use.

Answer file sections and entries configured:

- [RegionalSettings]
- Language

**Browser and Shell Settings**
Enables you to customize the browser and shell settings for Windows.

Answer file sections and entries configured:

- [FavoritesEx]

**Installation Folder**
Enables you to specify the name of the Windows folder.

Answer file sections and entries configured:

- [Unattended]
- TargetPath

**Install Printers**
Enables you to install network printers on the destination computers.

Answer file sections and entries configured:

- [GuiRunOnce]

**Run Once**
Enables you to automatically run a command the first time the user logs on.

Answer file sections and entries configured:

- [GuiRunOnce]

**Additional Commands**
Enables you to add commands to run at the end of unattended installation. Settings from this page are saved only if you choose to create a new distribution share or modify an existing one. If you choose to set up from a product CD, anything entered on this page is ignored.

Answer file sections and entries configured:

- [Commands] (in the Cmdlines.txt file)

**Unattended Installation Answer File Settings**

This section describes the settings that you can configure in an unattended installation answer file. For more information about how to structure an answer file, see "Unattended Installation Physical Structures" in How Unattended Installation Works.

The following table lists all of the valid sections that can appear in an answer file for an unattended installation.

**Answer File Sections for Unattended Installation**

| Section | Description |
|---|---|
| [adapter_specific_protocol_section] | Contains an entry for specifying a network component. Must be used in conjunction with the [Networking] section. |
| [Branding] | Contains entries for branding Microsoft Internet Explorer during an unattended installation. If you do not specify this section, Setup uses the default for Windows Server 2003. |
| [CertSrv_Client] | Contains entries for the Web client component of Certificate Services. You do not need to include this section in the answer file to install the Web client component of Certificate Services on the computer. Enabling Certificate Services in the [Components] section of the answer file installs the necessary files. |
| [CertSrv_Server] | Contains entries for specifying the server components of Certificate Services. The entries in this section are supported only for the Windows Server 2003. |
| [Components] | Contains entries for installing the optional components of Windows XP Professional and Windows Server 2003. |
| [Data] | Contains entries for starting directly from the Windows Server 2003 product CD when performing an unattended installation. |
| [DCInstall] | Contains entries for installing a domain controller after the initial setup of the operating system finishes. |
| [Display] | Contains entries for specifying display settings for graphics devices. |
| [FavoritesEx] | Contains the default Favorites settings for the Windows browser. If not present, Setup uses the default Uniform Resource Locator (URL) for Windows XP or Windows Server 2003. |
| [Fax] | Contains entries for specifying custom installation information for the Fax service. |
| [GuiRunOnce] | Contains commands that run the first time a user logs on to the computer after final (GUI-mode) stage of Setup completes. |
| [GuiUnattended] | Contains entries for preparing the graphical user interface (GUI) for unattended installation. |
| [Homenet] | Contains entries for installing Home Networking settings for network adapters, Internet Connection Sharing (ICS), and Internet Connection Firewall (ICF). |
| [Identification] | Contains entries for specifying the network identification of a computer. If these entries are not present, Setup adds the computer to the default workgroup. If there is not enough information for this entry, Setup prompts the user to provide this information. |
| [IEHardening] | Contains entries for configuring Web content zones. |
| [InternetServer] | Contains entries for configuring the installation folder for Internet Information Services (IIS) components. |
| [LicenseFilePrintData] | Contains entries for installing licensing information for Windows Server 2003. |

| | |
|---|---|
| [MassStorageDrivers] | Contains an entry for loading retail or original equipment manufacturer (OEM) small computer system interface (SCSI) drivers in the Setup Loader during the intermediate (text-mode) stage of Setup |
| [MS_AppleTalk parameters] | Contains entries for configuring AppleTalk. |
| [MS_ATMArps parameters] | Contains no entries. |
| [MS_ATMLANE parameters] | Contains no entries. |
| [MS_ATMUni parameters] | Contains no entries. |
| [MS_L2TP parameters] | Contains entries for the Layer 2 Tunneling Protocol (L2TP). |
| [MS_MSClient parameters] | Contains entries for Client for Microsoft Networks. |
| [MS_NetMon parameters] | Contains no entries. |
| [MS_NWClient parameters] | Contains entries for Client Service for NetWare. |
| [MS_NWIPX parameters] | Contains entries for the Internetwork Packet Exchange (IPX) protocol. Its parameters are network adapter–independent. |
| [MS_NwSapAgent parameters] | Contains no entries. |
| [MS_PPTP parameters] | Contains entries for the Point-to-Point Tunneling Protocol (PPTP). |
| [MS_Psched parameters] | Contains no entries. |
| [MS_RAS parameters] | Contains entries for the Routing and Remote Access service. |
| [MS_RasSrv parameters] | Contains entries for the Routing and Remote Access service for remote access servers. |
| [MS_Server parameters] | Contains entries for the File and Print Service. |
| [MS_TCPIP parameters] | Contains entries for configuring the TCP/IP protocol. |
| [MS_WLBS parameters] | Contains entries for configuring Network Load Balancing (formerly, Windows Load Balancing Service, or WLBS) and binding Network Load Balancing to multiple network adapters. |
| [NetAdapters] | Contains entries for network adapters that you want to install. |
| [NetBindings] | Contains entries for disabling, enabling, and demoting binding rules for network components such as services, protocols, and adapters. |
| [NetClients] | Contains an entry for specifying the network clients to install. |
| [NetOptionalComponents] | Contains entries for installing optional network components. |
| [NetProtocols] | Contains an entry for installing the network protocols. |
| [NetServices] | Contains entries for installing network services. |
| [Networking] | Contains no entries. However, to configure network protocols in unattended installation, you must include the [Networking] section name in your answer file. |
| [OEMBootFiles] | Contains entries for specifying OEM-supplied boot files. |
| [params.adapter_instance] | Contains entries for specifying an adapter's instance on a computer by using either PCI location information or a Plug and Play device ID. |
| [PCHealth] | Contains entries for configuring error reporting, Remote Assistance, and Help and Support Center. |

| | |
|---|---|
| [protocol_parameters_section] | Contains an entry for specifying network protocol parameters that are also specified in the [NetProtocols] section. |
| [Proxy] | Contains entries for specifying proxy server settings for the Windows XP or Windows Server 2003 browser. |
| [RegionalSettings] | Contains entries for installing regional options during the setup process. |
| [SetupData] | Contains the OSLoadOptionsVar entry, which is used to modify the /NoExecute option in Boot.ini used by Data Execution Prevention (DEP). |
| [SetupMgr] | Contains entries for specifying the server and folder in which the distribution share resides and the name of the distribution share. This section is used only by Setup Manager when an existing distribution share is edited. It is not used by Setup during an unattended installation. |
| [SetupParams] | Contains an entry for running an additional post-Setup command. |
| [Shell] | Contains entries for specifying the visual style of the user interface. |
| [SNMP] | Contains entries for specifying the Simple Network Management Protocol (SNMP) service. |
| [SystemFileProtection] | Contains entries for specifying Windows File Protection (WFP) settings. |
| [SystemRestore] | Contains entries for specifying the System Restore feature. |
| [TapiLocation] | Contains entries for specifying the telephony API (TAPI). It is valid only if a modem is present on the computer. |
| [TerminalServices] | Contains entries for configuring Terminal Services features (Terminal Server and Remote Desktop). |
| [UDDIDatabase] | Contains entries for configuring the optional Universal Description, Discovery and Integration (UDDI) Database feature. |
| [UDDIWeb] | Contains entries for configuring the optional Universal Description, Discovery and Integration (UDDI) Web services feature. |
| [Unattended] | Contains entries for running Setup. |
| [Uninstall] | Contains entries for allowing users who upgrade to Microsoft Windows XP to remove the operating system if they encounter problems. |
| [URL] | Contains entries for specifying default Uniform Resource Locator (URL) settings for the browser available in Windows XP or Windows Server 2003. |
| [UserData] | Contains entries for specifying user settings during the setup process. |
| [WindowsFirewall] | This section must be included in your answer file to configure firewall sections and entries in unattended installations. To configure Windows Firewall settings, you must include the [WindowsFirewall] section along with one or more of the other Windows Firewall settings sections. |
| [WindowsFirewall.profile_name] | Contains entries for specifying a domain type of profile, a standard type, or a profile that applies to both domain and standard. |
| [WindowsFirewall.program_name] | Contains entries for adding programs to the Windows Firewall Exceptions list. |
| [WindowsFirewall.service_name] | Contains entries for adding services to the Windows Firewall Exceptions list. |
| [WindowsFirewall.portopening_name] | Contains entries for adding static ports to the Windows Firewall Exceptions list. |

| | |
|---|---|
| [WindowsFirewall.icmpsetting_name] | The default configuration for Windows Firewall blocks all Internet Control Message Protocol (ICMP) message types, the entries in this section allows you to modify this behavior. |
| [Win9xUpg] | Contains entries for specifying an unattended upgrade from Windows 98, Windows 98 Second Edition, or Windows Millennium Edition to either Windows XP Home Edition or Windows XP Professional. |

# Sysprep Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sysprep Technical Reference**
The System Preparation tool (Sysprep) is a technology that you can use with other deployment tools to install Microsoft Windows operating systems with minimal administrator or technician intervention. You typically use the Sysprep tool in conjunction with a non-Microsoft disk imaging tool or Microsoft Windows Server 2003 Automated Deployment Services (ADS) to perform image-based installations. Image-based installations are ideal for large-scale rollouts when it would be too slow and costly to have administrators or technicians interactively install the operating system on individual computers.

**In this subject**


- What Is Sysprep?
- How Sysprep Works
- Sysprep Tools and Settings

# What Is Sysprep?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Sysprep?**

**In this section**


- Common Sysprep Scenarios
- Sysprep Limitations
- Sysprep Dependencies
- Technologies Related to Sysprep
- Related Information


The System Preparation tool (Sysprep) is a technology that you can use with other deployment tools to install Microsoft Windows operating systems with minimal intervention by an administrator or technician. Sysprep is typically used during large-scale rollouts when it would be too slow and costly to have administrators or technicians interactively install the operating system on individual computers.

You typically use the Sysprep tool in conjunction with a non-Microsoft disk imaging tool or Microsoft Windows Server 2003 Automated Deployment Services (ADS) to perform image-based installations. Image-based installation is a method of copying or cloning preconfigured operating systems (and, optionally, software applications) onto destination computers. After you set up a master installation — an installation with the operating system, software applications, and configuration settings that you want to install onto the destination computers in your organization, Sysprep prepares the master installation so that you can create a disk image; that is, a functionally identical replica of the disk containing the master installation, that can be copied onto multiple computers. The disk-imaging program creates the disk image of the master installation. After the disk image is copied onto a destination computer, and you start the destination computer, a shortened version of the Windows Setup program runs. The shortened version of Setup configures only user-specific and computer-specific settings, such as computer name, domain membership, and regional options. You can automate this last part of the setup process by using an answer file, a simple text file that instructs the Setup program how to configure the various operating system settings.

The following figure shows how an image-based installation differs from an interactive setup.

**Comparison of Image-based Installation and Interactive Setup**




Like all automated installation technologies, imaged-based installation with Sysprep has many benefits that standard interactive setup does not.

**Fewer errors during installation**

Because image-based installation with Sysprep uses preconfigured disk images and answer files to install and configure the operating system, there is minimal administrator or technician interaction during the setup process. This reduces the number of errors that are introduced during the setup process.

**Greater consistency**

By using the same disk image and answer file to install and configure operating systems, you can ensure that all of the computers in your organization are set up exactly the same way.

**Shortest installation times**

Image-based installation with Sysprep is faster than interactive setup because the operating system and applications are preinstalled and preconfigured on the disk image and require only minimal configuration after the disk image is copied to a destination computer. In addition, you can speed up the installation process by using an answer file. Instead of prompting administrators or technicians for configuration information, Setup reads configuration settings from an answer file.

**Lower support costs**

By minimizing errors during the setup process, increasing the consistency of the computers in your organization, and reducing the amount of time a technician needs to spend setting up a computer, you can reduce the overall support costs in your organization.

**Common Sysprep Scenarios**

Several installation scenarios are particularly well suited for image-based installations with Sysprep.

**Homogeneous hardware configurations**

If you are installing the same server configuration on computers that have similar or homogeneous hardware, you can customize hardware configurations on the disk image, thereby avoiding many post-installation configuration and installation tasks. For example, you can use image-based installation with Sysprep to rapidly roll out the same file server configuration to a group of computers that have similar processor configurations and similar storage configurations.

**Rapid rollout scenarios**

If you need to install the operating system as rapidly as possible, you can copy a preconfigured disk image onto a destination computer much faster than you can perform an unattended installation or a Remote Installation Services (RIS) installation. For example, you can use image-based installation with Sysprep to quickly deploy critical servers in your organization or quickly reinstall an operating system on member servers after a hard disk failure or some other catastrophic event. .

**Application deployment**

If you want to roll out applications at the same time that you roll out the operating system, you can include the applications on

the disk image. For example, if your organization's standard computer configuration includes an antivirus program, an e-mail program, and an office suite, you could install and configure these programs on the master installation so that the disk image includes these programs.

**Sysprep Limitations**
You can use Sysprep and image-based installation to deploy operating systems and software applications onto many types of servers. However, image-based installation has several significant limitations and can be used only when certain conditions are met.

**Clean installation only**
You can use image-based installation with Sysprep only to install a clean version of the operating system and clean versions of software applications. You cannot use image-based installation to upgrade an operating system or software configuration.

**Limited server configuration**
Some server components must be installed and configured after an image-based installation with Sysprep is complete. These components include Certificate Services, Cluster service, and any software that is dependent on the Active Directory directory service. They also include any application or service that stores the computer name or the computer SID and cannot recover if the computer name or SID changes.

**HAL compatibility**
You can perform an image-based installation with Sysprep only if the hardware abstraction layer (HAL) on the disk image is compatible with the hardware on the destination computer. In some cases, Windows Server 2003 automatically upgrades the HAL that is on a disk image to suit the HAL requirements of a destination computer, but this upgrade can occur only if the master installation is built on an APIC uniprocessor-system (UP) computer and the image of that master installation is copied to a compatible multiprocessor-system destination computer.

**Special domain controller installation process**
You cannot deploy preconfigured domain controllers by using image-based installation with Sysprep. However, you can configure a domain controller by first deploying a member server and then automatically running a script that runs Dcpromo.exe, the Active Directory Installation Wizard.

**Limited configuration of some security settings**
You cannot use image-based installation with Sysprep to deploy computers that contain any files that are encrypted by using Encrypting File System (EFS). In addition, you cannot use image-based installation to deploy systems that have already been configured with NTFS security settings, such as file and folder permissions, unless the disk-imaging program supports the NTFS file system. However, you can use a script to configure these settings after the image-based installation is complete.

### Note

- You can use Sysprep only to prepare a master installation that is installed on the master computer's drive C.

**Sysprep Dependencies**
You need two tools to perform an image-based installation: Sysprep, which can be found on any Microsoft Windows XP Professional or Windows Server 2003 product CD; and a disk-imaging program. Setup Manager can also be a useful tool for performing image-based installations with Sysprep, but it is not required.

**Sysprep**
The Sysprep tool consists of three separate programs: Sysprep.exe, Setupcl.exe, and Factory.exe. However, you run only Sysprep.exe; Setupcl.exe and Factory.exe are secondary programs that Sysprep.exe runs as needed. Sysprep is found in the Support\Tools folder on any Windows XP Professional or Windows Server 2003 product CD, under Deploy.cab.

You run Sysprep on the master computer before you create an image of the master computer's hard disk. Sysprep configures various operating system settings on the master computer to ensure that every copy of the master computer's disk image is unique when you distribute it to a destination computer. Specifically, Sysprep configures a master installation so that unique security identifiers (SIDs) are generated on each destination computer. Sysprep also configures the master computer's disk image so that every destination computer starts in a special setup mode known as Mini-Setup. After you copy a disk image onto a destination computer, Mini-Setup runs the first time you start the destination computer.

**Disk-imaging program**
You can use either a non-Microsoft disk-imaging program or ADS to create an image of the master computer's hard disk. You also use the disk-imaging program to copy the disk image from the master computer onto a shared folder or a CD, and from the shared folder or CD onto a destination computer.

### Note

- You can perform an image-based installation with Sysprep without purchasing a non-Microsoft disk-imaging program by using Windows Server 2003 Automated Deployment Services (ADS). ADS is an early deliverable of the Dynamic Systems Initiative (DSI), and can be used to rapidly deploy Windows server operating systems.

If you are not using ADS to create disk images, you need to purchase a non-Microsoft disk imaging program.

**Setup Manager**
In addition to Sysprep and a disk imaging program, there is an optional tool called Setup Manager that you can use to create an answer file for image-based installations with Sysprep. You can also create an answer file manually by using a text editor such as Notepad. Usually, you use Setup Manager to create a basic answer file, and then you use Notepad to customize the answer file.

**Technologies Related to Sysprep**
Image-based installation with Sysprep is similar to two other automated installation technologies: unattended installation and Remote Installation Services (RIS) installation. Like image-based installation with Sysprep, these technologies are designed specifically as mass roll-out or deployment solutions. Although you can use all of these technologies to automate the installation of Windows Server 2003, each of these technologies is based on a different set of installation processes, and relies on a different set of programs and tools.

**Unattended Installation**
Unattended installation is a method of automating clean installations and upgrades with minimal administrator or technician interaction. It relies on an answer file that you create and can be performed from the product CD or from a shared distribution folder.

Unattended installations are particularly useful if you need to:

- Perform automated installations on computers that have heterogeneous hardware configurations.
- Perform automated installations on specific types of servers, such as domain controllers, remote access servers, and servers that run Certificate Services or the Cluster service.
- Upgrade Windows-based servers to Microsoft Windows Server 2003, Standard Edition or Microsoft Windows Server 2003, Enterprise Edition operating systems.
- Configure a wide range of operating system settings during an automated installation without using batch files and scripts.

In addition to these deployment solutions, unattended installation is a useful method of creating master installations for image-based and RIS installations.

Unattended installation requires some upfront planning and design, but it is the most flexible and robust automated installation technology.

**RIS Installation**
RIS installation is a method of cleanly installing an operating system and applications over a network with minimal administrator or technician interaction. RIS uses Pre-Boot Execution Environment (PXE) technology to enable client computers without an operating system to start up and connect remotely to a RIS server, which installs a supported operating system.

You can use RIS to perform either CD-based installations or image-based installations. To perform CD-based installations you use the Remote Installation Services Setup (Risetup.exe) tool, which installs an operating system on a destination computer by using an answer file and the installation files that are on the product CD. To perform image-based installations you use the Remote Installation Preparation Wizard (Riprep.exe) tool, which is similar to the Sysprep tool and prepares a master installation for disk imaging.

RIS is a suitable automated installation technology if you need to:

- Install operating systems on computers that have similar hardware and a compatible HAL.
- Eliminate administrator or technician interaction during the setup process.
- Minimize post-installation tasks.

- Perform clean installations of an operating system, rather than upgrade existing installations.
- Install software applications with the operating system.

RIS installation is more complex and requires more upfront configuration and planning than image-based installation and unattended installation. However, RIS installation is an efficient deployment technology if you have a high-speed network and you need to perform clean installations of operating systems and preconfigured applications on a large number of computers.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Setup Technical Reference
- Unattended Installation Technical Reference
- Remote Installation Services Technical Reference

# How Sysprep Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Sysprep Works**

**In this section**


- Sysprep Architecture
- Sysprep Physical Structures
- Sysprep Processes and Interactions
- Related Information


The System Preparation tool (Sysprep) is a technology that you can use when performing an image-based installation. Image-based installation is a method of copying or cloning preconfigured operating systems and software applications onto destination computers. Image-based installation is the fastest method of installing an operating system and applications, and is typically used during large-scale rollouts when it would be too slow and costly to have administrators or technicians interactively install the operating system on individual computers.

These sections provide an in-depth view of how Sysprep works in an optimal environment. An optimal environment for Syprep contains the following:

- Plug and Play hardware that is compatible with Windows Server 2003.
- High-bandwidth network connectivity between the destination computers and the computer on which the disk images are stored.
- Physical access to the destination computers.
- BIOS and hardware that support starting from the floppy drive and the CD drive.
- Disk-imaging software, such as a non-Microsoft disk-imaging program that is compatible with Microsoft Windows Server 2003 or Windows Server 2003 Automated Deployment Services (ADS).

**Sysprep Architecture**

Sysprep relies on several components, including three executable programs and several answer files. How these components interact with each other at the various stages of an image-based installation is illustrated in the following figure.

**Sysprep Architecture**




The Sysprep tool consists of three components: Sysprep.exe, Setupcl.exe, and Factory.exe. These components run on both the master computer — the computer on which you are creating your master installation — and the destination computer — the computer on which you are installing a disk image. The answer files (Cmdlines.txt, Sysprep.inf, and Winbom.ini) are created on a technician or administrator computer, and are used on the master computer and the destination computer when Sysprep and its associated components run. The disk imaging tool runs on the master computer, where it is used to create a disk image, and on the destination computer, where it is used to copy a disk image from a CD or a remote server onto the destination computer's hard disk.

The following are components used in an image-based installation with Sysprep.

**Cmdlines.txt**
A configurable text file that you use to perform installation and configuration tasks at the end of the Mini-Setup stage. Cmdlines.txt contains a list of commands that run synchronously after the Mini-Setup stage finishes, but before a computer restarts.

**Destination computer**
Computer on which you are installing an operating system. For image-based installations, this is the computer to which you are copying a disk image. During a large-scale rollout, there can be typically hundreds or thousands of destination computers.

**Disk image**
Functionally identical replica of a hard disk that can be copied onto multiple computers. You need a disk imaging tool to create a disk image.

**Disk-imaging tool**
Program that creates and manages disk images. You use the disk-imaging tool on the master computer to create a disk image and save the disk image to a CD, DVD, or a remote server. You use the disk-imaging tool on the destination computer to copy the disk image from a CD, DVD, or remote server to the destination computer's hard drive. You can purchase a non-Microsoft disk-imaging program to perform an image-based installation with Sysprep, but you can also use ADS.

**Distribution server**
Computer that contains the disk images you want to copy onto destination computers.

**Factory.exe**
Sysprep component that controls the Factory-mode stage of an image-based installation with Sysprep. You can configure Sysprep so that a destination computer starts in Factory mode. Factory mode puts a destination computer into a network-enabled state that allows you to perform installation and configuration tasks before you prepare the destination computer for final delivery to a user. Factory.exe must exist in the same folder on the master computer as Setupcl.exe and Sysprep.exe, which is typically the *systemdrive*\Sysprep folder.

**Master computer**
Computer from which a disk image is derived. A master computer contains a master installation, which is the system disk on which you have installed and configured the operating system and software applications that you want to install onto the

destination computers in your organization.

**Mini-Setup**
A wizard that is a subset of Setup. Mini-Setup runs automatically after a disk image is copied to a destination computer and the destination computer is started. Mini-Setup provides prompts for user-specific information, configures operating system settings, and detects new hardware. You can automate the Mini-Setup stage by using Sysprep.inf.

**Registry**
Repository where Sysprep stores configuration and setup information prior to the disk imaging process. Sysprep reads the configuration and setup information from this repository the first time a destination computer starts after delivery to a user.

**Setupcl.exe**
Sysprep component that resets security identifiers (SIDs) on the master computer prior to running a disk imaging tool and creating a disk image of the master computer's system disk. Setupcl.exe can also run at the end of the Factory-mode stage if you run Sysprep with the Factory mode option. Setupcl.exe must exist in the same folder on the master computer as Factory.exe and Sysprep.exe, which is typically the *systemdrive*\Sysprep folder.

**Setup Manager**
Deployment tool that creates a basic Sysprep.inf answer file. You use Setup Manager on the administrator or technician computer, and then transfer the answer file to the master computer prior to creating the disk image.

**Sysprep.exe**
Primary Sysprep component. You run Sysprep.exe on the master computer prior to running a disk-imaging tool to create a disk image. Sysprep.exe also runs automatically after you copy a disk image to a destination computer and then start the destination computer for the first time. Sysprep.exe can be run from the command line or it can be run in the Windows graphical user interface (GUI). Sysprep.exe must exist in the same folder on the master computer as Setupcl.exe and Factory.exe, which is typically the *systemdrive*\Sysprep folder.

**Sysprep.inf**
An answer file that can be used to automate the Mini-Setup stage.

**Syssetup.dll**
Dynamic-link library that performs installation and configuration tasks during the Factory mode stage and the Mini-Setup stage of an image-based installation with Sysprep.

**Technician computer**
Computer used by an administrator or technician to use Setup Manager or Notepad to create and modify answer files, such as Sysprep.inf and Winbom.ini.

**Winbom.ini**
An answer file that you can use to automate installation tasks when a computer is started in Factory mode. Factory mode is useful for updating drivers, running Plug and Play enumeration, installing applications, testing, configuring the computer with customer data, or making other configuration changes before you deliver a computer to an end user.

**Winlogon.exe**
A component of the Windows operating system that provides interactive logon support. Winlogon is designed around an interactive logon model that consists of three components: the Winlogon executable, a Graphical Identification and Authentication dynamic-link library (DLL) referred to as the GINA, and many network providers.

**Sysprep Physical Structures**
The structures of two files are essential to an image-based installation with Sysprep: the Sysprep.inf answer file, which automates the Mini-Setup stage of an image-based installation, and the Winbom.ini answer file, which automates the Factory-mode stage of an image-based installation. You can create the Sysprep.inf answer file by using Setup Manager (Setupmgr.exe) or you can create it manually by using a text editor, such as Notepad. You must create the Winbom.ini answer file manually by using a text editor, such as Notepad.

As with most answer files, Sysprep.inf and Winbom.ini contain multiple predefined sections — some mandatory and some optional — that you can modify. Section names are always enclosed in square brackets (for example, [Unattended]).

Every section has one or more entries that contain installation and configuration information. Entries are usually composed of an entry name, an equal sign, and a value (for example, `ComputerName = Computer1`). The entry name represents a specific computer setting or action; the value represents the unique way you want the setting configured or the action performed. In some cases, entries do not have an entry name. In these cases, a value or a list of values (usually commands) appear in the answer file below a section name. The following example shows a section that contains only values and does not have entry names.

```
[GUIRunOnce]
```

```
dcpromo /answer:a:\installdc.txt
%systemroot%\cluster\cluster.exe
```

Use the following guidelines when you create or modify an answer file:

- Sections do not have to be any particular order.
- Sections and entries are not case sensitive.
- Entries must have valid values. For example, “on” is not equivalent to “enable,” and “off” is not equivalent to “disable.” Invalid sections and entries can generate errors, or cause Setup to fail.
- Each entry must have a value.
- Each section and entry must appear on a separate line.
- Comments must begin with a semi-colon (;).Do not use “REM” to specify a comment.

**Note**

- Do not modify undocumented sections or entries in an answer file; doing so can cause unpredictable results.

**Sysprep.inf Structure**

You can use Sysprep.inf to automate the Mini-Setup stage of an image-based installation. This corresponds to the final stage, or GUI-mode stage, of an interactive setup. The following table lists all of the valid sections that can appear in a Sysprep.inf file.

**Answer File Sections for Sysprep.inf**

| Section | Description |
|---|---|
| [*adapter_specific_protocol_section*] | Contains an entry for specifying a network component. Used in conjunction with the [Networking] section. |
| [Display] | Contains entries for specifying display settings for graphics devices. |
| [GuiRunOnce] | Contains commands to run the first time a user logs on to the computer after the Mini-Setup stage completes and the destination computer restarts. |
| [GuiUnattended] | Contains entries for configuring the way the Mini-Setup stage runs as well as time and date settings. . |
| [Identification] | Contains entries for specifying the network identification of a computer. If these entries are not present, Setup adds the computer to the default workgroup. If there is not enough information for this entry, Setup prompts the user to provide this information. |
| [IEHardening] | Contains entries for configuring Web content zones. |
| [LicenseFilePrintData] | Contains entries for installing licensing information for Windows Server 2003. |
| [*MS_AppleTalk parameters*] | Contains entries for configuring AppleTalk. |
| [*MS_ATMArps parameters*] | Contains no entries. |
| [*MS_ATMLANE parameters*] | Contains no entries. |
| [*MS_ATMUni parameters*] | Contains no entries. |
| [*MS_L2TP parameters*] | Contains entries for the Layer 2 Tunneling Protocol (L2TP). |
| [*MS_MSClient parameters*] | Contains entries for Client for Microsoft Networks. |

| | |
|---|---|
| [*MS_NetMon parameters*] | Contains no entries. |
| [*MS_NWClient parameters*] | Contains entries for Client Service for NetWare. |
| [*MS_NWIPX parameters*] | Contains entries for the Internetwork Packet Exchange (IPX) protocol. Its parameters are network adapter–independent. |
| [*MS_NwSapAgent parameters*] | Contains no entries. |
| [*MS_PPTP parameters*] | Contains entries for the Point-to-Point Tunneling Protocol (PPTP). |
| [*MS_Psched parameters*] | Contains no entries. |
| [*MS_RAS parameters*] | Contains entries for the Routing and Remote Access service. |
| [*MS_RasSrv parameters*] | Contains entries for the Routing and Remote Access service for remote access servers. |
| [*MS_Server parameters*] | Contains entries for the File and Print Service. |
| [*MS_TCPIP parameters*] | Contains entries for configuring the TCP/IP protocol. |
| [*MS_WLBS parameters*] | Contains parameters for configuring Network Load Balancing (formerly, Windows Load Balancing Service, or WLBS) and binding Network Load Balancing to multiple network adapters. |
| [NetAdapters] | Contains entries that specify network adapters. |
| [NetBindings] | Contains entries for disabling, enabling, and demoting binding rules for network components such as services, protocols, and adapters. |
| [NetClients] | Contains an entry for specifying the network clients to install. |
| [NetOptionalComponents] | Contains entries for installing optional network components. |
| [NetProtocols] | Contains an entry for installing the network protocols. |
| [NetServices] | Contains entries for installing network services. |
| [Networking] | Contains no entries. However, to configure network protocols, you must include the [Networking] section name in your answer file. |
| [*params.adapter_instance*] | Contains entries for specifying an adapter's instance on a computer. |
| [*protocol_parameters_section*] | Contains an entry for specifying network protocol parameters that are also specified in the [NetProtocols] section. |
| [RegionalSettings] | Contains entries for installing regional options during the setup process. |
| [Sysprep] | Contains an entry for automatically generating the entries in the pre-existing [SysprepMassStorage] section and then installing those mass-storage controllers. |
| [SysprepMassStorage] | Contains entries for identifying the different mass-storage controllers that you want to install on each destination computer. |
| [TapiLocation] | Contains entries for specifying the telephony API (TAPI). It is valid only if a modem is present on the computer. |
| [Unattended] | Contains entries that control the Mini-Setup stage. |
| [UserData] | Contains entries for specifying user settings during the setup process. |

You do not need to add every possible section and entry to the Sysprep.inf file. Most sections and entries are optional, and Mini-Setup ignores optional sections and entries that are not present in the answer file.

If you want to fully automate the Mini-Setup stage, you must configure the following sections and entries in Sysprep.inf. If you

do not provide values for these entries in Sysprep.inf, a prompt will appear during the Mini-Setup stage asking for the values.

**[GuiUnattended]**
You must provide values for the AdminPassword, EncryptedAdminPassword, and TimeZone entries. The value for AdminPassword cannot begin with an asterisk (*). Using a password that begins with an asterisk can cause the password to be set to a null value, which can be a security risk. You must also include the following entries and values: `OEMSkipWelcome = 1` and `OEMSkipRegional = 1`.

**[Identification]**
You must specify values for either the JoinWorkgroup entry or the JoinDomain entry. If you choose the latter, you must also provide values for the DomainAdmin and DomainAdminPassword entries. However, if you are concerned about security you might want to leave the DomainAdminPassword blank because this entry cannot be encrypted in the answer file. If you leave the DomainAdminPassword entry blank, you will be prompted for the password during the Mini-Setup stage.

**[LicenseFilePrintData]**
You must specify values for the AutoMode and AutoUsers entries if you are installing Windows Server 2003. This section does not apply to Windows XP.

**[Networking]**
At a minimum you must set `InstallDefaultComponents=Yes`.

**[Unattended]**
You must accept the End User License Agreement (EULA) by setting `OEMSkipEula = Yes`. You must also provide a value for the InstallFilesPath entry, which specifies the location of any files that are necessary during the Mini-Setup stage, such as language files.

**[UserData]**
You must specify values for the ComputerName, FullName, and OrgName entries. You must also specify a value for the ProductKey entry.

**Winbom.ini Structure**
You can use Winbom.ini to automate the Factory-mode stage of an image-based installation with Sysprep. The following table lists all of the valid sections that can appear in a Winbom.ini file.

### Answer File Sections for Winbom.ini

| Section | Description |
|---|---|
| [Components] | Contains entries for installing the components of Windows XP Professional and Windows Server 2003. |
| [ComputerSettings] | Contains entries for configuring user settings on the destination computer. |
| [DesktopShortcutsFolder] | Contains an entry for specifying a shortcuts folder on the desktop. |
| [DiskConfig] | Contains entries for configuring and partitioning one or more physical hard disks when Factory.exe runs in the Windows Preinstallation Environment (Windows PE) |
| [*description*.config] | Contains entries for configuring partitions located on a physical disk. |
| [Factory] | Contains entries for managing the actions of Factory.exe. |
| [NetCards] | Contains an entry for configuring the network adapter on the destination computer, if Setup requires a network connection. |
| [OEMLink] | Contains entries for adding graphics and text to the Start menu so that users have easy access to an OEM's customized folder, .htm file, or .exe file for a tour. |
| [OEMRun] | Enables you to run external applications and command-line scripts when auditing the computer in Factory mode, by using the `Sysprep -factory` command. |
| [OEMRunOnce] | Enables you to control how and when external applications and command shell scripts run when auditing the computer in Factory mode, by using the `Sysprep -factory` command. |
| [*section_name*] | Contains entries for configuring the preinstallation of applications during the Factory-mode stage of Sysprep. This section can be used only with the [OEMRunOnce] section. |

| | |
|---|---|
| [PnPDrivers] | Lists all the updated drivers to copy onto the destination computer. |
| [PnPDriverUpdate] | Contains entries for specifying optional user credentials and a source location of drivers to place on an installed image. |
| [SetupHomenet] | Contains an entry for setting up a firewall on a computer with a preconfigured Internet connection. |
| [Shell] | Contains entries for specifying the visual style of the user interface. |
| [StartMenuMFUList] | Contains entries for installing shortcut links in the Start menu. |
| [UpdateSystem] | Contains entries, called directives, for performing registry or file system updates. |
| [UserAccounts] | Contains entries for controlling the creation of user accounts on the destination computer. |
| [UserDefined] | Contains branding information or any other data that other processes might need to access. |
| [WinPE] | Contains entries for identifying the configuration set and Windows operating system that you install. |
| [WinPE.Net] | Contains entries for controlling the settings that Windows Preinstallation Environment (Windows PE) uses to connect to a network. |

You do not need to add every possible section and entry to the Winbom.ini file. All sections and entries are optional, and Factory.exe ignores optional sections and entries that are not present in the answer file.

For more information about answer file sections, including descriptions and examples of entries and values, see "Sysprep Tools and Settings".

**Sysprep Processes and Interactions**

Sysprep configures various operating system settings on the master computer to ensure that every copy of the master computer's disk image is unique when you distribute it to a destination computer. Specifically, Sysprep configures a master installation (and hence, the disk image of that master installation) so that a unique computer *security identifier* (SID) is generated on each destination computer. A SID is a variable length data structure that identifies a user, a group, or a computer. Internal processes in Windows refer to the SID instead of the computer name or user name or group name. Sysprep resets SIDs by running *Setupcl.exe*, which is a Sysprep component.

You can also use Sysprep to configure the master computer's disk image so that every destination computer starts in a special setup mode known as *Mini-Setup*. Mini-Setup is the first-run experience on Windows Server 2003 and Windows XP, however, you can also configure Sysprep so that Windows Welcome is the first-run experience on Windows XP.

After you copy a disk image onto a destination computer, Mini-Setup runs the first time you start the destination computer. Mini-Setup provides the administrator, technician, or user the opportunity to configure operating system settings that are unique to the user or the computer. For example, computer name, regional options, time zone, and display settings. The Mini-Setup stage is a shortened version of the final stage of an interactive setup.

You can also use Sysprep to configure the master computer's disk image so that every destination computer starts in a special intermediate setup mode known as *Factory mode*. Factory mode is controlled by *Factory.exe* and is commonly used in a manufacturing environment where every computer requires some customization prior to final delivery; however, it is also useful for some corporate deployments. You can use Factory mode if you want to customize individual computers or groups of computers after you have copied a disk image onto a destination computer but before Mini-Setup runs. For example, Factory mode gives you the ability to update drivers, run Plug and Play enumeration, install applications, test hardware and software, and configure computer settings. Factory mode is an extension of Mini-Setup and encompasses many of the same installation and configuration tasks as the final (GUI-mode) stage of an interactive setup.

In addition to preparing a master computer for disk duplication, you can configure Sysprep so that the Mini-Setup and Factory-mode stages are automated. You use an answer file known as *Sysprep.inf* to automate the Mini-Setup stage of an image-based installation; you use an answer file known as *Winbom.ini* to automate the Factory-mode stage of an image-based installation. You can also use an answer file known as *Cmdlines.txt* to automate installation and configuration tasks at the end of the Mini-Setup stage. Cmdlines.txt does not contain any predefined, structured entries like Sysprep.inf and Winbom.ini. Instead, Cmdlines.txt contains command-line instructions similar to a batch file. You can use Cmdlines.txt to perform customized installation and configuration tasks such as installing software programs and running command-line programs.

This section describes the primary processes and interactions that occur when you:

- Run the Sysprep tool to prepare a master installation for disk imaging.
- Run a destination computer in Factory mode.
- Run Mini-Setup to perform final installation and configuration tasks.
- Perform post-installation tasks with Cmdlines.txt and the [GuiRunOnce] section.

This section does not discuss all of the processes and interactions that occur during an image-based installation: this section discusses only the processes and interactions that occur when you run the Sysprep tool. For example, the disk-imaging process and the image-copying process are not discussed in this section because these processes vary according to the disk-imaging program that you use.

**Tasks Performed by Sysprep**
You can use the Sysprep tool to perform four tasks.

**Prepare a master installation for final delivery**
This is the primary purpose of the Sysprep tool. To do this you run the `Sysprep -reseal` command, or click the **Reseal** button in the Sysprep interface.

**Configure a destination computer for Factory mode**
This task is performed primarily in manufacturing environments where every computer requires some customization prior to final delivery. However, it is also useful for some corporate deployments. To do this you run the `Sysprep -factory` command, or click the **Factory** button in the Sysprep interface.

**Configure a destination computer for auditing**
This task is performed to audit a destination computer before Mini-Setup runs. To do this you run the `Sysprep -audit` command, or click the **Audit** button in the Sysprep interface.

**Remove hidden mass storage devices**
This task is performed to remove references to devices that were present on the master computer but are not present on the destination computer. To do this you run the `Sysprep -clean`, `Sysprep -factory`, or `Sysprep -reseal` command, or click **Factory** or **Reseal** in the Sysprep interface.

The following sections describe the processes and interactions that occur when you use the four main parameters of Sysprep. The processes and interactions are presented in the approximate order in which they occur. It is not possible to describe these processes and interactions in precise order because the order might vary slightly depending on the way you perform the image-based installation and the options you choose when you run the Sysprep tool.

**Running Sysprep -reseal**
There are three types of processes and interactions that occur when you run the `Sysprep -reseal` command: verification and data gathering, configuration, and cleanup.

**Verification and Data Gathering**
Before any configuration settings are changed on the master installation, Sysprep checks various configuration settings and gathers information about the master installation. The main verification and data gathering processes include:

- Verifying that no other instance of Sysprep is running on the master installation. If there is another version of Sysprep running, Sysprep generates an error message.
- Verifying command-line syntax. For example, several command-line parameters must be used exclusively, such as `-bmsd`, `-clean`, `-audit`, `-factory`, and `-reseal`. If these parameters are used with other parameters, Sysprep generates an error.
- Verifying that the user is a member of the local Administrators group. If not, Sysprep generates an error and exits.
- Verifying that the user has sufficient privileges to run Sysprep. Users must have the following privileges: back up files and directories, modify firmware environment variables, restore files and directories, and shut down the system.
- Verifying that Setupcl.exe is present in the current directory (usually, *systemroot*\Sysprep). If the file is present, Sysprep copies it to *systemroot*\System32. If the file is not present, Sysprep checks for it in *systemroot*\System32. If it is not present there, Sysprep generates an error message and exits. This process is skipped if you use the `-nosidgen` parameter.

- Displaying a dialog box that describes the Sysprep tool. Users must verify that they have read this message. This process is skipped if you use the `-quiet` parameter.

- Verifying that that master computer is not joined to a domain. If it is, Sysprep removes the computer from the domain.

- Verifying that the screensaver is not enabled. If it is, Sysprep disables it and saves the settings so it can be enabled and reconfigured during the Mini-Setup stage.

**Configuration**
After the verification and data gathering processes are complete, Sysprep configures the master installation. The main configuration processes and interactions include:

- Checking the operating system version to be sure it is either Windows XP or Windows Server 2003. If you are not running Sysprep on either operating system, Sysprep generates an error and exits.

- Finding Sysprep.inf and saving the path information.

- Checking whether the master computer is a member of a domain, and if it is, disjoining it from the domain.

- Adding the information string to the registry if you specified a value for the OEMDuplicatorString entry in the [GuiUnattended] section of Sysprep.inf.

- Configuring the registry so that Setupcl.exe either runs or does not run the first time a destination computer is started. This is depends on whether you use the `-nosidgen` parameter. Setupcl.exe is responsible for resetting SIDs.

- Removing mass storage devices from the critical device database that were included in the [SysprepMassStorage] section but are not present on the destination computer.

- Saving network configuration settings so that they can be reapplied the first time the destination computer is started, and uninstalling all network adapters so that they will be detected the first time a destination computer is started.

**Cleanup**
After the configuration processes complete, Sysprep cleans up user-specific and computer-specific settings and data that you don't want copied to a destination computer. The main cleanup processes include:

- Deleting Systems Management Server settings.

- Deleting the event logs.

- Deleting per-user settings, such as user settings for Windows Media Player and Internet Connection Wizard.

- Deleting the list of most recently used documents, applications, and Web sites.

- Modifying device paths in the registry if you specified a path in the InstallFilesPath entry or the OemPnPDriversPath in the [Unattended] section of Sysprep.inf.

After these processes run, the master computer either shuts down or restarts depending on whether you use the `-reboot` or `-noreboot` parameters with Sysprep.

**Running Sysprep -factory**
There are two types of processes and interactions that occur when you run the `Sysprep -factory` command: configuration and cleanup.

**Configuration**
The main configuration processes and interactions that occur when you run the `Sysprep -factory` command include:

- Renaming Winbom.ini so that the file is not used again in the event that the computer is restarted during the Factory-mode stage.

- Checking to see if Factory.exe exists and saving the path information.
- Verifying that Factory.exe and Sysprep.exe are on the same drive and in the same location on the drive.
- Making various registry modifications so that the destination computer starts in factory mode the first time it is started.
- Making sure the user has privileges to update the registry.
- Checking whether the master computer is a member of a domain, and if it is, disjoining it from the domain.
- Configuring the registry so that Setupcl.exe runs the first time a destination computer is started.
- Building the [SysprepMassStorage] section in Sysprep.inf if the BuildMassStorageSection entry is present in the [Sysprep] section of Sysprep.inf or you used the `-bmsd` parameter with Sysprep.
- Saving network configuraton settings so they can be reapplied the first time the destination computer is started, and uninstalling all network adapters (except legacy network adapters) so they will be detected the first time a destination computer is started.

**Cleanup Processes**
After the configuration processes complete, Sysprep cleans up user-specific and computer-specific settings and data that you don't want to appear on a destination computer. The main cleanup processes include:

- Deleting per-user settings, such as user settings for Windows Media Player and Internet Connection Wizard.
- Deleting the list of most recently used documents, applications, and Web sites.
- Modifying device paths in the registry if you specified a path in the InstallFilesPath entry or the OemPnPDriversPath in the [Unattended] section of Sysprep.inf.
- Deleting the event logs.

After these processes run, the master computer shuts down.

**Running Sysprep -audit**
When you use Sysprep to configure a destination computer for auditing, Sysprep changes the registry so that the destination computer starts in an audit mode the next time it is started. When the technician or administrator is finished performing auditing tasks, Sysprep must be run again with the `-factory` parameter or the `-reseal` parameter, or by clicking **Factory** or **Reseal** in the Sysprep interface.

**Running Sysprep -clean**
When you use Sysprep to remove hidden mass storage devices, Sysprep does two things:

- Finds Sysprep.inf and removes the BuildMassStorageSection entry, if it exists.
- Removes mass storage devices from the critical device database that were included in the [SysprepMassStorage] section but are not present on the destination computer.

After the command `Sysprep -clean` is run, the computer is typically shut down and is ready for delivery to its final destination.

**Factory Mode Processes and Interactions**
You can use Factory mode to perform configuration and installation tasks on a destination computer. It is most often used to perform configuration and installation tasks that you could not perform on the master installation. This includes such tasks as configuring operating system settings, installing additional device drivers, and installing applications. When you use Factory mode, every destination computer starts in a network-enabled state without starting Mini-Setup. In this state, Factory.exe processes its answer file, Winbom.ini, and performs these actions:

- Copies drivers from a network source to the computer.
- Starts Plug and Play enumeration.
- Stages, installs, and uninstalls applications on the computer from source files located on either the computer or a network source.
- Performs configuration tasks specified in the Winbom.ini file.

**Locating a Winbom.ini File**
Sysprep searches these locations in this order for a Winbom.ini file:

- The path and file name specified by the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Factory\Winbom.
- The root of all removable media drives that are not CD drives, such as a floppy disk drive.
- The root of all CD drives.
- The location of Factory.exe, usually the %systemdrive%\Sysprep folder.
- The root of %systemdrive%.

After locating a Winbom.ini file, Sysprep reads the value of the WinbomType entry in the [Factory] section. If the value of WinbomType is not appropriate (for example, you use the entry `WinbomType = WinPE` but the computer is not running WinPE), the Winbom.ini file is ignored and Sysprep continues searching for a Winbom.ini file. If the value of the WinbomType entry is appropriate, Sysprep reads the value of the NewWinbom entry in the [Factory] section. If a value is specified for NewWinbom, and a Winbom.ini file is located at that location, Sysprep examines that Winbom.ini file for a NewWinbom entry.

This cycle continues until Sysprep locates a Winbom.ini file of the correct type that does not contain a NewWinbom entry or for a maximum of ten times, whichever occurs first. Sysprep then continues to run in Factory mode, performing the tasks that are specified in the Winbom.ini file.

When you run Sysprep in Factory mode, the NewWinbom entry is processed only once at each start.

**Processing Winbom.ini**
Factory.exe processes Winbom.ini entries in different ways, depending on the value of the WinbomType entry in the [Factory] section of the Winbom.ini file.

- If you use the entry `WinbomType = WinPE`, you can start a computer by using the Windows Preinstallation Environment (Windows PE) as the operating system, and run the command `factory -winpe` to preinstall Windows from a configuration set specified in the Winbom.ini file.
- If you use the entry `WinbomType = Factory`, the Winbom.ini file can be used to control Sysprep in factory mode, which you run by using the command `Sysprep -factory`.
- If you use the entry `WinbomType = OOBE`, you can use the Winbom.ini file to configure Windows XP during Windows Welcome. After you run the command `Sysprep -reseal`, Windows Welcome runs the next time you start the computer.

A state is the smallest group of Winbom.ini sections and entries that is logged as a single unit. For each value of the WinbomType entry, the states are listed in the order in which Factory.exe processes them. States always restart the computer, meaning that if the computer restarts either programmatically or unexpectedly, Factory.exe resumes by processing its last state, along with any states that are processed on every restart of the computer for that value of WinbomType.

Specific states can apply to more than one value of the WinbomType entry, and Factory.exe might process the same state multiple times during a particular instance of running the tool.

The following sections contain information about the states that occur during each type of Winbom.ini processing. The information for each state is listed in a table formatted as follows:

| | |
|---|---|
| State | Name of the state as it appears in the user interface and/or the log file. This name is localized and can be different in each language build of Factory.exe. |
| Winbom.ini | List of the Winbom.ini sections and entries processed in this state.<br>• If only a section is listed, then all entries in the section are processed.<br>• If the value for an entry is not specified, the entry can have any value that is syntactically correct. |
| Default | Specifies the default value if the entry is omitted or blank.<br>• "Nothing" means that the state does nothing if the entry is not specified.<br>• Otherwise, the values listed are used as specified, even when the entries are not present |
| Executes | Specifies if this state is processed once or on every boot of the computer.<br>• "Once" means that after the state completes, Factory.exe does not process that state again until you run either `Sysprep -factory` or `Sysprep -reseal`. Any entries read in those states are not read again from that state if you run `Sysprep -audit`.<br>• "Always" means that the state is processed on every start. If the computer restarts in the middle of processing the Winbom.ini file or if you manually restart the computer while auditing the installation, all of the "Always" states are processed again. However, each state is processed only once per start, unless otherwise noted. |
| Variables | Specifies if the entries processed in this state support environment variables.<br>• If No, then none of the entries support them.<br>• If Yes, then all the entries listed (or the whole section) support them.<br>• Otherwise, the individual Winbom.ini entries that support environment variables are listed.<br><br>If the state occurs before a user account logs on to the system or during Windows Welcome, then that state supports only system environment variables, not per user-defined environment variables, because the state is processed in the system context. |
| Comments | Additional comments, if any. |

**Winbom.ini Processing When Preinstalling Windows from Windows PE**

The states in this section contain the Winbom.ini entries that are processed when you start a computer by using Windows PE as the operating system, and run the command `factory -winpe` to preinstall Windows from a configuration set specified in the Winbom.ini file. WinbomType, located in the [Factory] section of the Winbom.ini file, must include the value WinPE.

The following tables describe the states in the order in which they are processed.

| State | Setting display resolution |
|---|---|
| Winbom.ini | [ComputerSettings]<br>DisplayResolution<br>DisplayRefresh |

| | |
|---|---|
| Default | [ComputerSettings]<br><br>DisplayResolution = "800x600x16" |
| Executes | Always |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Partitioning disks and formatting drives |
| Winbom.ini | [DiskConfig]<br><br>DiskN<br><br>[description.Config]<br><br>WipeDisk<br><br>ForceFormat<br><br>FileSystemN<br><br>PartitionTypeN<br><br>QuickFormatN<br><br>SetActiveN<br><br>SizeN<br><br>N is the (1-based) index of the physical disk. |

| | |
|---|---|
| Default | On 32-bit versions of Windows:<br>• [description.Config]<br>• WipeDisk = No<br>• ForceFormat = No<br>• FileSystemN = NTFS<br>• PartitionTypeN = Primary<br>• QuickFormatN = Yes<br>• SetActiveN = No<br><br>On 64-bit versions of Windows:<br>• [description.Config]<br>• WipeDisk = Yes<br>• ForceFormat = No<br>• FileSystemN = NTFS<br>• PartitionTypeN = Primary<br>• QuickFormatN = Yes<br>• SetActiveN = No |
| Executes | Always |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Creating Page File |
| Winbom.ini | [WinPE]<br>PageFileSize |
| Default | If the computer has less than 64 MB of physical memory:<br>• [WinPE]<br>• PageFileSize = 64 |
| Executes | Always |
| Variables | No |
| Comments | If the machine has more than 64 MB of physical memory, the page file size is 0 by default. |

| | |
|---|---|
| State | Starting Networking |
| Winbom.ini | [WinPE.Net]<br>Startnet<br>IPConfig<br>SubnetMask<br>Gateway |
| Default | [WinPE.Net]<br>Startnet = Yes<br>IPConfig = DHCP |
| Executes | Always |
| Variables | No |
| Comments | If StartNet = No, then subsequent values are ignored.<br>If IPConfig = DHCP, then SubnetMask and Gateway are ignored. |

| | |
|---|---|
| State | Installing applications |
| Winbom.ini | [OEMRunOnce] |
| Default | Nothing |
| Executes | Always |
| Variables | No |
| Comments | Credentials come from the UserName and Password entries in the [Factory] section. |

| | |
|---|---|
| State | Launching [OEMRun] programs |
| Winbom.ini | [OEMRun] |
| Default | None |
| Executes | Always |
| Variables | No |
| Comments | Credentials come from the UserName and Password entries in the [Factory] section. |

| | |
|---|---|
| State | Installing the Configuration Set |
| Winbom.ini | [WinPE]<br>ConfigSet<br>Lang<br>Quiet<br>SKU<br>SourceRoot |

| | |
|---|---|
| Default | [WinPE]<br><br>Quiet = No |
| Executes | Always |
| Variables | No |
| Comments | Credentials come from [WinPE] section or from the UserName and Password entries in the [Factory] section. |

| | |
|---|---|
| State | Processing INF settings |
| Winbom.ini | [UpdateSystem] |
| Default | None |
| Executes | Always |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Finishing Windows PE |
| Winbom.ini | [WinPE]<br><br>Restart |
| Default | [WinPE]<br><br>Restart = Prompt |
| Executes | Always |
| Variables | No |
| Comments | None |

**Winbom.ini Processing During Sysprep in Factory Mode**

The states in this section are the Winbom.ini entries processed when you run the command `Sysprep -factory`. The WinbomType entry, located in the [Factory] section of the Winbom.ini file, must include the value Factory.

The following tables describe the states in the order in which they are processed.

**Note**

- The `Sysprep -factory` command starts very early in the operating system boot sequence to enable operations that must occur before other system processes start. For example, during Plug and Play device enumeration, other system services and boot processes run, the `Sysprep -factory` command performs other tasks with the full Windows operating system running in the background.

| | |
|---|---|
| State | Extending system drive |
| Winbom.ini | [ComputerSettings]<br><br>ExtendPartition |
| Default | [ComputerSettings]<br><br>ExtendPartition = 0 |
| Executes | Once |

| | |
|---|---|
| Variables | No |
| Comments | You can extend only NTFS file system partitions. The partition to extend must have unpartitioned space available on its disk following the partition. |

| | |
|---|---|
| State | Setting source path |
| Winbom.ini | [ComputerSettings]<br><br>SourcePath |
| Default | Nothing |
| Executes | Once |
| Variables | Yes |
| Comments | This should point to the folder that contains the I386 folder. If the files exist in the root, and this entry is present, and the source path that it points to does not exist, Factory.exe moves the files from the root to where the source path points. |

| | |
|---|---|
| State | Installing test certificate |
| Winbom.ini | [ComputerSettings]<br><br>TestCert |
| Default | Nothing |
| Executes | Once |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Setting factory computer name |
| Winbom.ini | [Factory]<br><br>FactoryComputerName<br><br>RebootAfterComputerName |
| Default | [Factory]<br><br>RebootAfterComputerName = No |
| Executes | Once |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Copying updated drivers |
| Winbom.ini | [PNPDrivers]<br><br>[PNPDriverUpdate]<br><br>DevicePath<br><br>TargetRoot |

| | |
|---|---|
| Default | [PNPDrivers]<br>[PNPDriverUpdate]<br>TargetRoot = %SYSTEMROOT%\drivers |
| Executes | Once |
| Variables | Yes |
| Comments | None |

| | |
|---|---|
| State | Starting Plug and Play |
| Winbom.ini | Nothing |
| Default | Nothing |
| Executes | Always |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Setting display resolution |
| Winbom.ini | [ComputerSettings]<br>DisplayResolution<br>DisplayRefresh |
| Default | [ComputerSettings]<br>DisplayResolution = "800x600x16" |
| Executes | Always |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Customizing shell settings |
| Winbom.ini | [Shell]<br>MSNExplorer<br>DefaultStartPanelOff<br>DefaultThemesOff<br>CustomDefaultThemeFile<br>DoDesktopCleanup |
| Default | Nothing |
| Executes | Once |
| Variables | CustomDefaultThemeFile |
| Comments | DefaultStartPanelOff, DefaultThemesOff, and CustomDefaultThemeFile change the relevant setting only when new users are created. These settings are not changed on the administrator/owner account. |

| | |
|---|---|
| State | Preparing for automatic logon |
| Winbom.ini | [ComputerSettings]<br><br>AuditAdminAutoLogon |
| Default | [ComputerSettings]<br><br>AuditAdminAutoLogon = No |
| Executes | Once |
| Variables | No |
| Comments | If AuditAdminAutoLogon = Yes, then Factory.exe logs on to the desktop by using the administrator account. On Windows XP Home Edition, the owner account is used instead of the administrator account. |

| | |
|---|---|
| State | Waiting for Plug and Play to finish |
| Winbom.ini | [PNPDriverUpdate]<br><br>WaitForPnP |
| Default | [PNPDriverUpdate]<br><br>WaitForPnP = Yes |
| Executes | Always |
| Variables | No |
| Comments | Runs only if WaitForPnP = Yes. |

| | |
|---|---|
| State | Updating installed drivers |
| Winbom.ini | [PnPDriverUpdate]<br><br>UpdateInstalledDrivers |
| Default | If the first-run experience is Windows Welcome:<br>• [PNPDriverUpdate]<br>• UpdateInstalledDrivers = Yes<br><br>If the first-run experience is Mini-Setup:<br>• [PNPDriverUpdate]<br>• UpdateInstalledDrivers = No |
| Executes | Once |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Setting display resolution |

| | |
|---|---|
| Winbom.ini | [ComputerSettings]<br>DisplayResolution<br>DisplayRefresh |
| Default | [ComputerSettings]<br>DisplayResolution = "800x600x16" |
| Executes | Always |
| Variables | No |
| Comments | This state is run a second time, in case Plug and Play installed a new display driver. |

| | |
|---|---|
| State | Optimizing Shell |
| Winbom.ini | [Factory]<br>OptimizeShell |
| Default | [Factory]<br>OptimizeShell = Yes |
| Executes | Always |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Customizing shell font settings |
| Winbom.ini | [ComputerSettings]<br>FontSmoothing |
| Default | Nothing |
| Executes | Once |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Customizing shell settings |
| Winbom.ini | [Shell]<br>DefaultStartPanelOff<br>DefaultThemesOff<br>CustomDefaultThemeFile |
| Default | Nothing |
| Executes | Once |
| Variables | CustomDefaultThemeFile |
| Comments | DefaultStartPanelOff, DefaultThemesOff, and CustomDefaultThemeFile change the relevant setting only when new users are created. These settings are not changed on the administrator/owner account. |

| | |
|---|---|
| State | Customizing power settings |
| Winbom.ini | [ComputerSettings]<br>Hibernation<br>PowerScheme |
| Default | Nothing |
| Executes | Once |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Customizing home network settings |
| Winbom.ini | [SetupHomeNet]<br>EnableFirewall |
| Default | Nothing |
| Executes | Once |
| Variables | No |
| Comments | If this section exists, Factory.exe uses the Home Networking feature to configure settings. |

| | |
|---|---|
| State | Adding user identities |
| Winbom.ini | [UserAccounts] |
| Default | Nothing |
| Executes | Once |
| Variables | No |
| Comments | User accounts are not configured on the Windows Server 2003 family. |

| | |
|---|---|
| State | Processing INF settings |
| Winbom.ini | [UpdateSystem] |
| Default | Nothing |
| Executes | Always |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Configuring Optional Components |
| Winbom.ini | [Components] |
| Default | Nothing |
| Executes | Once |

| | |
|---|---|
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Installing applications |
| Winbom.ini | [OEMRunOnce] |
| Default | Nothing |
| Executes | Always |
| Variables | Yes |
| Comments | Credentials come from the UserName and Password entries in the [Factory] section. |

| | |
|---|---|
| State | Adding Start Menu frequently used links |
| Winbom.ini | [StartMenuMFUlist] |
| Default | Nothing |
| Executes | Once |
| Variables | Yes |
| Comments | None |

| | |
|---|---|
| State | Recording OEM application preferences |
| Winbom.ini | [Shell]<br>DefaultClientStartMenuInternet<br>DefaultClientMail<br>DefaultClientMedia<br>DefaultClientIM<br>DefaultClientJavaVM |
| Default | Nothing |
| Executes | Once |
| Variables | Yes |
| Comments | None |

| | |
|---|---|
| State | Applying OEM Start Menu branding |
| Winbom.ini | [OEMLink]<br>[DesktopShortcutsFolder] |
| Default | Nothing |
| Executes | Once |
| Variables | Yes |

| | |
|---|---|
| Comments | None |

| | |
|---|---|
| State | Launching [OEMRun] programs |
| Winbom.ini | [OEMRun] |
| Default | Nothing |
| Executes | Always |
| Variables | Yes |
| Comments | Credentials come from the UserName and Password entries in the [Factory] section |

| | |
|---|---|
| State | Waiting for Plug and Play to Finish |
| Winbom.ini | [PnPDriverUpdate]<br>WaitForPnP |
| Default | [PNPDriverUpdate]<br>WaitForPnP = Yes |
| Executes | Always |
| Variables | No |
| Comments | Runs only if WaitForPnP = Yes. |

| | |
|---|---|
| State | Launching Sysprep |
| Winbom.ini | [Factory]<br>Reseal<br>ResealFlags<br>ResealMode |
| Default | On Windows XP Home Edition or Windows XP Professional:<br>• [Factory]<br>• Reseal = No<br>• ResealoMode = OOBE<br><br>On Windows XP 64-Bit Edition or the Windows Server 2003 family:<br>• [Factory]<br>• ResealMode = Mini |
| Executes | Always |
| Variables | No |
| Comments | On Windows XP 64-Bit Edition or the Windows Server 2003 family, the default behavior is for Sysprep to run and present its user interface. |

**Winbom.ini Processing During Windows Welcome**

The states in this section are the Winbom.ini entries processed at the beginning of Windows Welcome on Windows XP. These states are not processed if the first-run experience is Mini-Setup. The WinbomType entry, located in the [Factory] section of the Winbom.ini file, must include the value OOBE.

The following tables describe the states in the order in which they are processed.

| | |
|---|---|
| State | Extending system drive |
| Winbom.ini | [ComputerSettings]<br><br>ExtendPartition |
| Default | Nothing |
| Executes | Once |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Setting source path |
| Winbom.ini | [ComputerSettings]<br><br>SourcePath |
| Default | Nothing |
| Executes | Once |
| Variables | Yes |
| Comments | The value of SourcePath the folder that contains the I386 folder. If the files exist in the root, this entry is present, and the source path that it points to does not exist, then Factory.exe moves the files from the root to the location specified by SourcePath. |

| | |
|---|---|
| State | Installing test certificate |
| Winbom.ini | [ComputerSettings]<br><br>TestCert |
| Default | Nothing |
| Executes | Once |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Updating installed drivers |
| Winbom.ini | [PnPDriverUpdate]<br><br>UpdateInstalledDrivers |
| Default | [PNPDriverUpdate]<br><br>UpdateInstalledDrivers = Yes |
| Executes | Once |

| | |
|---|---|
| Variables | No |
| Comments | Only functions if UpdateInstalledDrivers = Yes. |

| | |
|---|---|
| State | Setting display resolution |
| Winbom.ini | [ComputerSettings]<br>DisplayResolution<br>DisplayRefresh |
| Default | [ComputerSettings]<br>DisplayResolution = "800x600x16" |
| Executes | Once |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Optimizing Shell |
| Winbom.ini | [Factory]<br>OptimizeShell |
| Default | [Factory]<br>OptimizeShell = Yes |
| Executes | Always |
| Variables | No |
| Comments | Does not function if OptimizeShell = No. If OptimizeShell = Yes, the Windows shell settings are configured every time Windows Welcome runs, even if that causes the computer to restart. |

| | |
|---|---|
| State | Customizing shell settings |
| Winbom.ini | [Shell]<br>MSNExplorer<br>DefaultStartPanelOff<br>DefaultThemesOff<br>CustomDefaultThemeFile<br>DoDesktopCleanup |
| Default | Nothing |
| Executes | Once |
| Variables | CustomDefaultThemeFile |
| Comments | DefaultStartPanelOff, DefaultThemesOff, and CustomDefaultThemeFile change the relevant setting only when new users are created. These settings are not changed on the administrator/owner account. |

| | |
|---|---|
| State | Customizing power settings |
| Winbom.ini | [ComputerSettings]<br>Hibernation<br>PowerScheme |
| Default | Nothing |
| Executes | Once |
| Variables | No |
| Comments | None |

| | |
|---|---|
| State | Configuring Optional Components |
| Winbom.ini | [Components] |
| Default | Nothing |
| Executes | Once |
| Variables | No |
| Comments | None |

**Mini-Setup Processes and Interactions**

Mini-Setup is a shortened version of an interactive setup. Syssetup.dll controls Mini-Setup, which is the same DLL that controls the final stage (GUI mode stage) of an interactive setup. However, not all of the interactive setup steps are performed during Mini-Setup because the operating system was already installed when the master installation was created.

The primary processes and interactions that occur during the Mini-Setup stage are described below in the approximate order in which they occur.

- Search for a Sysprep.inf file. If a Sysprep.inf is found, read it and save the settings.
- Begin Plug and Play enumeration and installation of devices.
- Determine whether any network components need to be reconfigured because the SIDs changed, and perform the necessary configuration tasks.
- Extend the primary partition, if specified in the answer file.
- Start the Mini-Setup wizard, which displays the following dialog boxes (if Sysprep.inf contains the configuration information required for a dialog box, the box is not displayed or it is displayed briefly and then closed).
    - Welcome to the Windows Setup Wizard (click Next)
    - End-User License Agreement (accept or not accept)
    - Regional and Language Options
    - Personalize Your Software (user name and organization name)
    - Your Product Key
    - Licensing Modes (per server or per device or per user)
    - Computer Name and Administrator Password

    - Date and Time Settings
    - Networking Settings (typical or custom)
    - Workgroup or Computer Domain
    - Completing the Windows Setup Wizard (click Finish)
- Execute the commands in Cmdlines.txt.
- Write any commands listed in the [GUIRunOnce] section to the registry so they can run the first time the computer starts and a user logs on.
- Update the HAL, if specified in the answer file.
- Delete the Sysprep folder, which is in the systemroot root.

When the Mini-Setup stage finishes, the computer restarts and the entries specified in the [GUIRunOnce] section are run.

**Post Installation Processes and Interactions**
Several Sysprep installation processes occur after an operating system is installed on a destination computer. These processes occur when you automate post-installation configuration tasks. Examples of post-installation configuration tasks include running command-line programs to configure system services, running scripts to configure Group Policy settings, and installing software programs. The most commonly used mechanisms for automating post-installation tasks is the Cmdlines.txt file and the [GUIRunOnce] in the Sysprep.inf answer file.

**Cmdlines.txt Processes and Interactions**
A Cmdlines.txt file contains a list of commands, programs, scripts, or batch files that run just after the Mini-Setup stage finishes but before the destination computer restarts. By default, Cmdlines.txt does not exist, so you must manually create the file with a text editor, such as Notepad.

The syntax for Cmdlines.txt is as follows:

```
              [Commands]
"command_1"
"command_2"
     .
     .
"command_x"
```

Where *command_1*, *command_2*, and *command_x* refer to the commands, programs, scripts, or batch files that you want to run when the Mini-Setup stage is complete. Note that all commands must be enclosed in quotation marks. Also, if you are using the command window (Cmd.exe) to run commands, programs, scripts, or batch files, then you need to use the `/c` parameter with the `cmd` command.

For example, to create a new folder named Test on drive C, you type:

```
              "cmd /c mkdir c:\Test"
```

In addition, to use a Cmdlines.txt file, you need to:

- Add the following section, entry, and value to your Sysprep.inf answer file:

  ```
  [Unattended]
  InstallFilesPath = C:\Sysprep\i386
  ```

- In the distribution share that you used to create the master installation, create the following folder:

*systemdrive*\Sysprep\i386\$OEM$.

- Copy Cmdlines.txt to the *systemdrive*\Sysprep\i386\$OEM$ folder.

You can only have one Cmdlines.txt file, but you can use it to run multiple commands, programs, scripts, or batch files. When these commands, programs, scripts, or batch files run, the following processes and interactions occur:

- The commands, programs, scripts, and batch files run synchronously, which means each command, program, script, or batch file runs to completion before the next one starts running.
- All commands, programs, scripts, and batch files specified in a Cmdlines.txt run under the Local System security account.
- Network connectivity is not available when the commands, programs, scripts, and batch files listed in a Cmdlines.txt file run.
- Windows Installer packages (.msi files) cannot be run from a Cmdlines.txt file.

**[GuiRunOnce] Processes and Interactions**

The [GuiRunOnce] section of an answer file contains a list of commands, programs, scripts, or batch files that run the first time an administrator or technician logs on to the destination computer after the operating system is installed.

The syntax for [GuiRunOnce] is:

```
              [GuiRunOnce]
"command_1"
"command_2"
     .
     .
"command_x"
```

Where *command_1*, *command_2*, and *command_x* refer to the commands, programs, scripts, or batch files that you want to run after the operating system is installed and the administrator or technician logs on. Note that all commands must be enclosed in quotation marks.

The commands that are listed in the [GuiRunOnce] section are stored in the following registry location:

HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Runonce

When an administrator or technician logs onto the destination computer for the first time after the Mini-Setup stage is complete, Windows Server 2003 reads and then runs the commands that are listed in this registry location. When these commands, programs, scripts, or batch files run, the following processes and interactions occur:

- The commands, programs, scripts, and batch files run synchronously, which means each command, program, script, or batch file runs to completion before the next one starts running.
- The commands, programs, scripts, and batch files listed in the [GuiRunOnce] section run under the security account of the user who logged onto the destination computer (usually an administrator or a technician).
- Network connectivity is available when the commands, programs, scripts, and batch files run.
- You can use Windows Installer packages (.msi files) to install applications (unlike Cmdlines.txt).
- You can run commands, programs, scripts, and batch files from hard disk drives, CD drives, shared folders on the network, or other storage devices (unlike Cmdlines.txt, which runs everything from the $OEM$ folder in your distribution share).

In addition, there are some interactions and processes that you might need to control if you run certain commands, programs, scripts, and batch files with the [GuiRunOnce] section.

**Prevent restarts**

You can run commands, programs, scripts, and batch files only if you can prevent the command, program, script, or batch file from restarting the computer. When a computer restarts, all remaining entries in the [GuiRunOnce] section are lost because the Runonce registry entry is deleted. If the system restarts before completing entries listed in the [GuiRunOnce] section, the remaining items will not run. Therefore, you need to suppress restarts. If you cannot suppress a restart within a command, program, script, or batch file, you can try to repackage the command, program, script, or batch file into a Windows Installer package.

**Do not rely on the Windows Explorer shell**

The Windows Explorer shell is not loaded when the operating system starts running commands, programs, scripts, or batch files that are listed in the [GuiRunOnce] section of an answer file. In some cases, you can get an updated installer program from the application vendor that does not rely on the Windows Explorer shell. If you cannot, you might be able to repackage the command, program, script, or batch file as a Windows Installer package (.msi file).

**Control the processes by using the /wait parameter**

Commands, programs, scripts, and batch files often start and stop several different processes. In some cases, when you are running multiple programs — especially installation programs — this can inadvertently cause the next program listed in the [GuiRunOnce] section to start before the previous program is finished running. When this occurs, the second program usually fails. To prevent this, you can run the programs from batch files by using the `start` command with the `/wait` parameter. This forces each program to complete before the next command listed in [GuiRunOnce] runs.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Setup Technical Reference
- Unattended Installation Technical Reference
- Remote Installation Services Technical Reference

# Sysprep Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sysprep Tools and Settings**

**In this section**


- Sysprep Tools
- Sysprep.inf Answer File Settings
- Winbom.ini Answer File Settings
- Sysprep Command-Line Parameters


The following tools, answer files, and command-line parameters are associated with image-based installations with the System Preparation tool (Sysprep). You use the tools and answer files to perform image-based installations with Sysprep.

**Sysprep Tools**

The following tools are associated with image-based installations with Sysprep.

**Winnt.exe: Windows Setup program for 16-bit operating systems**

**Category**

Windows setup and installation tool available on the Microsoft Windows Server 2003 product CD.

**Version compatibility**

Runs on Microsoft MS-DOS and Windows for Workgroups.

You can use this tool to initiate an unattended installation from the 16-bit operating systems MS-DOS and Windows for Workgroups. Typically, you use Winnt.exe to perform an unattended installation on a master computer prior to running the Sysprep tool. You cannot perform an upgrade with Winnt.exe; you can only perform clean installations with Winnt.exe.

For more information about Winnt.exe, see “Command Line References” in the Tools and Settings Collection.

**Winnt32exe: Windows Setup program for 32-bit and 64-bit operating systems**

**Category**

Windows setup and installation tool available on the Windows Server 2003 product CD.

**Version compatibility**

Runs on Microsoft Windows Server 2003, Windows XP, Windows 2000, Windows NT, Windows Millennium Edition, Windows 98, and Windows 95.

You can use this tool to perform an interactive setup or an unattended installation on a computer that has an existing operating system. To perform an interactive setup or an unattended installation from a 32-bit Windows operating system you must use the 32-bit version of Winnt32.exe. To perform an interactive setup or an unattended installation from a 64-bit Windows operating system you must use the 64-bit version of Winnt32.exe. Typically, you use Winnt32.exe to create a master installation on a master computer prior to running the Sysprep tool.

For more information about Winnt32.exe, see “Command Line References” in the Tools and Settings Collection.

**DiskPart.exe: DiskPart**

**Category**

Windows Server 2003 operating system tool.

**Version compatibility**

DiskPart is identical for all versions of Windows Server 2003. DiskPart includes a few 64-bit parameters that are only available on Itanium-based computers.

DiskPart is a text-mode command interpreter that enables you to manage objects (disks, partitions, or volumes) by using scripts or direct input from a command prompt.

For more information about DiskPart, see “Command Line References” in the Tools and Settings Collection.

**Setupmgr.exe: Setup Manager**
**Category**
Windows deployment tool available on the Windows Server 2003 product CD.

**Version compatibility**
Runs on the 32-bit versions of Windows Server 2003 and Windows XP.

Setup Manager is a wizard-like program that helps you configure distribution shares and answer files that can be used to perform automated installations, including unattended installation, image-based installation with Sysprep, and Remote Installation Services (RIS) installations. Setup Manager prompts you for information in a series of dialog boxes. Setup Manager configures answer file settings and distribution shares based on the information you provide in these dialog boxes.

Usually, you run Setup Manager on the administrator or technician computer. Setup Manager does not affect the system files of the technician computer on which the tool runs and does not actually run Setup. Also, you do not need to specify all of your customizations with Setup Manager. Instead, you can create an answer file in Setup Manager and then manually edit the file with Notepad. If you later open the answer file in Setup Manager, the tool displays your changed settings.

Setup Manager is not installed when you install Windows XP or Windows Server 2003. You must manually install Setup Manager on the administrator or technician computer.

**To install Setup Manager**

1. Insert a Windows XP or Windows Server 2003 product CD into your CD drive and open the \Support\Tools folder.
2. Double-click **Deploy.cab**.
3. Right-click the file named **Setupmgr.exe** and click **Extract**.
4. In the **Select a Destination** dialog box, select a location on your hard drive where you want to save Setup Manager, and then click **Extract**.

You can run Setup Manager from the graphical user interface by double-clicking the Setupmgr.exe file that you saved to your hard disk, or from the command line by typing `setupmgr`.

### Note

- If you use Setup Manager to create or modify a distribution share that is not on the local computer, be sure you run Setup Manager under a user account that has the minimum user rights to connect to the remote computer. This is because Setup Manager sends the user name and the password across the network in plaintext when it connects to a remote distribution share.

The following sections provide information about each Setup Manager dialog box. The section names correspond to the dialog box name that you see in Setup Manager.

**New or Existing Answer File**
Enables you to create a new answer file or edit an existing one.

Answer file sections and entries configured:

- None

**Type of Setup**
Provides three Setup options: Unattended, Sysprep, or Remote Installation Services (RIS).

Answer file sections and entries configured:

- None

**Product**
Prompts you to select the Windows product that you want to install.

Answer file sections and entries configured:

- None

**User Interaction**
Prompts you to select the level of user interaction you want. This dialog box is not used if you are configuring a Sysprep or RIS installation.

Answer file sections and entries configured:

- [Unattended]
- UnattendMode

**Distribution Share**
Enables you to create a new distribution share, modify an existing one, or set up Windows from a product CD.

Answer file sections and entries configured:

- None

**Location of Setup Files**
Prompts you for the location of the Setup files you want to copy.

Answer file sections and entries configured:

- None

**Distribution Share**
Prompts you for the location of the distribution share you want to create or modify. This section is used only by Setup Manager when an existing distribution share is edited. It is not used by Setup during an unattended installation.

Answer file sections and entries configured:

- [SetupMgr]
- DistFolder
- DistShare

**Name and Organization**
Enables you to provide a default user name and organization name.

Answer file sections and entries configured:

- [UserData]
- FullName
- OrgName

**Display Settings**
Enables you to change the change settings for colors, screen area, and refresh frequency.

Answer file sections and entries configured:

- [Display]
- BitsPerPel
- XResolution

- YResolution
- VRefresh

**Time Zone**
Prompts you to select a default time zone.

Answer file sections and entries configured:

- [GuiUnattended]
- TimeZone

**Product Key**
Prompts you to provide a Product Key to identify your copy of Windows.

Answer file sections and entries configured:

- [UserData]
- ProductKey

**Computer Names**
Prompts you to provide names for the destination computers.

Answer file sections and entries configured:

- [UserData]
- ComputerName

**Administrator Password**
Enables you to specify a password for the Administrator account on all destination computers.

Answer file sections and entries configured:

- [GuiUnattended]
- AdminPassword
- EncryptedAdminPassword

**Networking Computers**
Enables you to choose typical or custom network settings.

Answer file sections and entries configured:

- [Networking]
- [NetAdapters] (if custom settings are configured)
- [NetClients] (if custom settings are configured)
- [NetProtocols] (if custom settings are configured)

**Workgroup or Domain**
Prompts you to specify how the destination computers will participate in a network.

Answer file sections and entries configured:

- [Identification]

**Windows Components**
Enables you to install commonly used Windows components.

Answer file sections and entries configured:

- [Components]

**Telephony**
Enables you to provide telephony information for the user.

Answer file sections and entries configured:

- [TapiLocation]

**Regional Settings**
Enables you to customize Windows for different regions and languages.

Answer file sections and entries configured:

- [RegionalSettings]

**Languages**
Prompts you to select the language groups you want to use.

Answer file sections and entries configured:

- [RegionalSettings]

- Language

**Browser and Shell Settings**
Enables you to customize the browser and shell settings for Windows.

Answer file sections and entries configured:

- [FavoritesEx]

**Installation Folder**
Enables you to specify the name of the Windows folder.

Answer file sections and entries configured:

- [Unattended]

- TargetPath

**Install Printers**
Enables you to install network printers on the destination computers.

Answer file sections and entries configured:

- [GuiRunOnce]

**Run Once**
Enables you to automatically run a command the first time the user logs on.

Answer file sections and entries configured:

- [GuiRunOnce]

**Additional Commands**
Enables you to add commands to run at the end of the Mini-Setup stage.

Answer file sections and entries configured:

- [Commands] (in the Cmdlines.txt file)

**Sysprep.inf Answer File Settings**
This section describes the settings that you can configure in the Sysprep.inf answer file, which controls the Mini-Setup stage of an image-based installation with Sysprep. For more information about how to structure the Sysprep.inf answer file, see "Sysprep Physical Structures" in How Sysprep Works.

The following table lists all of the valid sections that can appear in a Sysprep.inf file.

**Answer File Sections for Sysprep.inf**

| Section | Description |
|---|---|
| [adapter_specific_protocol_section] | Contains an entry for specifying a network component. Used in conjunction with the [Networking] section. |
| [Display] | Contains entries for specifying display settings for graphics devices. |
| [GuiRunOnce] | Contains commands to run the first time a user logs on to the computer after the Mini-Setup stage completes and the destination computer restarts. |
| [GuiUnattended] | Contains entries for configuring the way the Mini-Setup stage runs as well as time and date settings. . |
| [Identification] | Contains entries for specifying the network identification of a computer. If these entries are not present, Setup adds the computer to the default workgroup. If there is not enough information for this entry, Setup prompts the user to provide this information. |
| [IEHardening] | Contains entries for configuring Web content zones. |
| [LicenseFilePrintData] | Contains entries for installing licensing information for Windows Server 2003. |
| [MS_AppleTalk parameters] | Contains entries for configuring AppleTalk. |
| [MS_ATMArps parameters] | Contains no entries. |
| [MS_ATMLANE parameters] | Contains no entries. |
| [MS_ATMUni parameters] | Contains no entries. |
| [MS_L2TP parameters] | Contains entries for the Layer 2 Tunneling Protocol (L2TP). |
| [MS_MSClient parameters] | Contains entries for Client for Microsoft Networks. |
| [MS_NetMon parameters] | Contains no entries. |
| [MS_NWClient parameters] | Contains entries for Client Service for NetWare. |
| [MS_NWIPX parameters] | Contains entries for the Internetwork Packet Exchange (IPX) protocol. Its parameters are network adapter–independent. |
| [MS_NwSapAgent parameters] | Contains no entries. |
| [MS_PPTP parameters] | Contains entries for the Point-to-Point Tunneling Protocol (PPTP). |

| | |
|---|---|
| [MS_Psched parameters] | Contains no entries. |
| [MS_RAS parameters] | Contains entries for the Routing and Remote Access service. |
| [MS_RasSrv parameters] | Contains entries for the Routing and Remote Access service for remote access servers. |
| [MS_Server parameters] | Contains entries for the File and Print Service. |
| [MS_TCPIP parameters] | Contains entries for configuring the TCP/IP protocol. |
| [MS_WLBS parameters] | Contains parameters for configuring Network Load Balancing (formerly, Windows Load Balancing Service, or WLBS) and binding Network Load Balancing to multiple network adapters. |
| [NetAdapters] | Contains entries that specify network adapters. |
| [NetBindings] | Contains entries for disabling, enabling, and demoting binding rules for network components such as services, protocols, and adapters. |
| [NetClients] | Contains an entry for specifying the network clients to install. |
| [NetOptionalComponents] | Contains entries for installing optional network components. |
| [NetProtocols] | Contains an entry for installing the network protocols. |
| [NetServices] | Contains entries for installing network services. |
| [Networking] | Contains no entries. However, to configure network protocols, you must include the [Networking] section name in your answer file. |
| [params.adapter_instance] | Contains entries for specifying an adapter's instance on a computer. |
| [protocol_parameters_section] | Contains an entry for specifying network protocol parameters that are also specified in the [NetProtocols] section. |
| [RegionalSettings] | Contains entries for installing regional options during the setup process. |
| [Sysprep] | Contains an entry for automatically generating the entries in the pre-existing [SysprepMassStorage] section and then installing those mass-storage controllers. |
| [SysprepMassStorage] | Contains entries for identifying the different mass-storage controllers that you want to install on each destination computer. |
| [TapiLocation] | Contains entries for specifying the telephony API (TAPI). It is valid only if a modem is present on the computer. |
| [Unattended] | Contains entries that control the Mini-Setup stage. |
| [UserData] | Contains entries for specifying user settings during the setup process. |

**Winbom.ini Answer File Settings**

This section describes the settings that you can configure in the Winbom.ini answer file, which controls the Factory mode stage of an image-based installation with Sysprep. For more information about how to structure the Winbom.ini answer file, see "Sysprep Physical Structures" in How Sysprep Works.

The following table lists all of the valid sections that can appear in a Winbom.ini file.

**Answer File Sections for Winbom.ini**

| Section | Description |
|---|---|
| [Components] | Contains entries for installing the components of Windows XP Professional and Windows Server 2003. |

| | |
|---|---|
| [ComputerSettings] | Contains entries for configuring user settings on the destination computer. |
| [DesktopShortcutsFolder] | Contains an entry for specifying a shortcuts folder on the desktop. |
| [DiskConfig] | Contains entries for configuring and partitioning one or more physical hard disks when Factory.exe runs in the Windows Preinstallation Environment (Windows PE) |
| [description.config] | Contains entries for configuring partitions located on a physical disk. |
| [Factory] | Contains entries for managing the actions of Factory.exe. |
| [NetCards] | Contains an entry for configuring the network adapter on the destination computer, if Setup requires a network connection. |
| [OEMLink] | Contains entries for adding graphics and text to the Start menu so that users have easy access to an OEM's customized folder, .htm file, or .exe file for a tour. |
| [OEMRun] | Enables you to run external applications and command-line scripts when auditing the computer in Factory mode, by using the `Sysprep -factory` command. |
| [OEMRunOnce] | Enables you to control how and when external applications and command shell scripts run when auditing the computer in Factory mode, by using the `Sysprep -factory` command. |
| [section_name] | Contains entries for configuring the preinstallation of applications during the Factory-mode stage of Sysprep. This section can be used only with the [OEMRunOnce] section. |
| [PnPDrivers] | Lists all the updated drivers to copy onto the destination computer. |
| [PnPDriverUpdate] | Contains entries for specifying optional user credentials and a source location of drivers to place on an installed image. |
| [SetupHomenet] | Contains an entry for setting up a firewall on a computer with a preconfigured Internet connection. |
| [Shell] | Contains entries for specifying the visual style of the user interface. |
| [StartMenuMFUList] | Contains entries for installing shortcut links in the Start menu. |
| [UpdateSystem] | Contains entries, called directives, for performing registry or file system updates. |
| [UserAccounts] | Contains entries for controlling the creation of user accounts on the destination computer. |
| [UserDefined] | Contains branding information or any other data that other processes might need to access. |
| [WinPE] | Contains entries for identifying the configuration set and Windows operating system that you install. |
| [WinPE.Net] | Contains entries for controlling the settings that Windows Preinstallation Environment (Windows PE) uses to connect to a network. |

**Sysprep Command-Line Parameters**

Sysprep.exe can be run from the command line or it can be run in the Windows graphical user interface (GUI). Use the following syntax when running Sysprep.exe from the command line.

```
sysprep [-bmsd] | {{[-factory] | [-reseal]} {[-clean] | [-activated] [-
audit] [-forceshutdown] [-mini] [-noreboot] [-nosidgen] [-pnp] [-quiet] [-reboot]}}
```

**Parameters**

Sysprep.exe uses the following parameters.

**-bmsd**

If the [SysprepMassStorage] section header exists in Sysprep.inf, and you run the `Sysprep -bmsd` command, then Sysprep populates the [SysprepMassStorage] section with the `Plug_and_Play_ID = path_to_device_inf_file` entries corresponding to the Plug and Play IDs of mass-storage devices specified in the Machine.inf, Scsi.inf, Pnpscsi.inf, and Mshdc.inf files. Sysprep builds the list of mass-storage devices; it does not install these devices in the critical device database or do any other processing.

After using the `Sysprep -bmsd` command to generate entries in the [SysprepMassStorage] section, you can delete items from this section before running `Sysprep -reseal` or `Sysprep -factory`. Installing a smaller number of items in the critical device database reduces the time required for this image to reboot into the operating system.

Do not use the `-bmsd` parameter with any other command-line parameters.

**-factory**
Restarts the destination computer in a network-enabled state without running Mini-Setup. This parameter is useful for updating drivers, running Plug and Play enumeration, installing applications, testing, configuring the computer with customer data, or making other configuration changes. For companies that use disk imaging (or cloning) software, Factory mode can reduce the number of images you need to create.

When you have finished your Factory mode tasks, run the `Sysprep -reseal` command to prepare the computer for delivery.

**-reseal**
Prepares the computer for final delivery and configures the computer so that Mini-Setup runs the next time the computer is started. If you run the `Sysprep -factory` command, you must seal the installation as the last step in your preinstallation process, either by running the `Sysprep -reseal` command or by clicking the Reseal button in the Sysprep dialog box.

**-clean**
Clears unused mass-storage drivers added by the [SysprepMassStorage] section of Sysprep.inf, and removes previously installed but nonexistent devices created by Plug and Play.

**-activated**
Does not reset the grace period for Windows Product Activation.

Use this parameter only if you have activated the Windows operating system on the destination computer prior to running Mini-Setup. The Product Key that you use to activate the Windows installation must match the Product Key located on the COA sticker attached to that computer.

**-audit**
Restarts the computer into Factory mode without generating new security IDs (SIDs) or processing any items in the [OEMRunOnce] section of Winbom.ini.

Use this parameter only if the computer is in Factory mode.

**-forceshutdown**
Shuts down the computer after Sysprep finishes.

Use this option if your computers have an ACPI BIOS that does not shut down properly with the Sysprep default settings.

**-mini**
Configures Windows XP Professional to use Mini-Setup instead of Windows Welcome. This parameter has no effect on Windows XP Home Edition, where the first-run experience is always Windows Welcome. Likewise, this parameter is not needed on Windows Server 2003 because Mini-Setup is the default first-run experience.

**-noreboot**
Modifies registry keys without the system restarting or preparing for duplication. This option is used mainly for testing, specifically to see if the registry is modified properly. This option is not recommended because making changes to a computer after Sysprep has run may invalidate the preparation done by Sysprep. Do not use this option in a production environment.

**-nosidgen**
Runs Sysprep without generating new SIDs. You must use this option if you do not duplicate the computer, or if the computer is a domain controller.

**-pnp**
Runs the full Plug and Play device enumeration and installation during the Mini-Setup stage. This command-line option has no effect if the first-run experience is Windows Welcome.

Use the `-pnp` parameter only to detect and install legacy, non-Plug and Play devices. Do not use the `-pnp` parameter on computers that use only Plug and Play devices. Otherwise, you will increase the time required for the first-run experience without providing any additional benefit to the user.

**-quiet**
Runs Sysprep without displaying onscreen confirmation messages. This is useful if you automate Sysprep. For example, if you plan to run Sysprep immediately after an unattended installation, add the `Sysprep -quiet` command to the [GuiRunOnce] section of the Unattend.txt file.

**-reboot**

Forces the computer to restart automatically and then start the Mini-Setup stage or the Factory-mode stage, as specified. This is useful when you want to audit the system and verify that the first-run experience operates correctly.

# Remote Installation Services Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services Technical Reference**
Remote Installation Services (RIS) is an automated installation technology that you can use to create installation images of operating systems or of complete computer configurations, including desktop settings and applications. These installation images can then be made available to users at client computers. RIS is typically used during large-scale deployments when it would be too slow and costly to have administrators or end users interactively install the operating system on individual computers.

When you use RIS, a CD-ROM is not required for installing an operating system on a client computer. The client computers must support remote startup with the Pre-Boot eXecution Environment (PXE) ROM, or they must be started with a remote-startup floppy disk and have a network adapter that is supported by that disk.

The Setup program in Microsoft Windows Server 2003 provides the method for carrying out an automated installation using RIS.

**In this subject**


- What Is Remote Installation Services?
- How Remote Installation Services Work
- Remote Installation Services Tools and Settings

# What Is Remote Installation Services?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Remote Installation Services?**

**In this section**


- Remote Installation Services Dependencies and Requirements
- Remote Installation Services Limitations
- Technologies Related to Remote Installation Services


Remote Installation Services (RIS) is an automated installation technology that you can use to create installation images of operating systems or of complete computer configurations, including desktop settings and applications. These installation images can then be made available to users at client computers. RIS is typically used during large-scale deployments when it would be too slow and costly to have administrators or end users interactively install the operating system on individual computers.

Compared to a standard interactive installation, RIS decreases installation time, eliminates the need to use CDs during installation, and minimizes end user and administrator interaction during installations. To accomplish this, RIS relies on Pre-Boot eXecution Environment (PXE) technology. PXE technology enables computers without an operating system to start up and connect remotely over a TCP/IP network connection to a RIS server, which then installs an operating system. The client computers must support remote booting with the PXE ROM, or they must be started with a remote-startup floppy disk and have a network adapter that is supported by that disk.

To perform a RIS installation, an administrator or end user typically boots a client computer from a PXE-capable network adapter by pressing F12 early in the startup process. This instructs the computer to acquire an IP address and connect to a RIS server. The client computer then downloads a small version of the Windows Setup program from the RIS server. You can configure a RIS installation so that the Windows Setup program interactively steps the administrator or end user through a shortened setup process, or you can configure it so that the entire setup process is automated after the end user or administrator presses F12. Also, you can configure a RIS installation so the installation is image-based or setup-based. By using image-based installations, you can deploy applications with an operating system. By using setup-based installations, you have more flexibility in terms of handling hardware differences among the computers in your organization.

The following figure shows how a RIS installation differs from an interactive installation.

**Comparison of RIS and Interactive Installations**




Like all automated installation technologies, RIS has many benefits that a standard interactive installation does not. These benefits include:

**Fewer errors during installation**
Because RIS installation uses either preconfigured disk images or answer files to install and configure the operating system, there is minimal user interaction during the setup process. This reduces the number of errors that are introduced during setup.

**Greater consistency**
By using the same disk image or answer file to install and configure the operating system, you can ensure that all of the computers in your organization are set up exactly the same way.

**Shorter installation times**
RIS installation is faster than interactive installation because the operating system and applications are preinstalled and preconfigured on the disk image and require only minimal configuration after the disk image is copied to a destination computer. In addition, you can speed up the installation process by using the fully automated installation features of RIS. Instead of prompting end users or administrators for configuration information, the Setup program reads configuration settings from an answer file.

**Lower support costs**
By minimizing errors during the setup process, increasing the consistency of the computers in your organization, and reducing the amount of time an administrator needs to spend setting up a computer, you can reduce the overall support costs in your organization.

RIS is more complex and requires more upfront configuration and planning than other automated installation technologies. This is because RIS relies on several operating system services and protocols, as well as specialized hardware that supports the PXE specification. However, RIS installation is an efficient deployment technology if you have a high-speed network and you need to install clean operating systems and preconfigured applications onto a large number of computers with minimal end user or administrator interaction.

**RIS Dependencies and Requirements**
RIS relies on specialized hardware as well as several Windows Server 2003 features and technologies. This includes the following:

**High speed network connectivity**
The computers on which you are installing an operating system must all be connected to a reliable, high speed network. Wide area networks (WANs) that are connected with analog modems or Integrated Services Digital Network (ISDN) do not have

sufficient bandwidth to support RIS installation.

**PXE-enabled hardware**
To take full advantage of RIS, each of the computers on which you are installing an operating system must have a Peripheral Component Interconnect (PCI), Mini-PCI, or CardBus type of network adapter that supports the PXE specification. PXE-enabled network adapters allow client computers to be remotely booted. In addition, the basic input/output system (BIOS) for each computer must be capable of using the network adapter as a boot device (also known as remote-boot capability). A ROM BIOS that is at least version .99n satisfies this requirement. Also, most computers that conform to the Net PC or PC98 specifications have a PXE remote boot–enabled network adapter and remote boot–enabled BIOS.

The published PXE specification defines the remote boot process and also establishes the PXE compliance standards for hardware manufacturers and other vendors. RIS uses PXE environment extensions to Dynamic Host Configuration Protocol (DHCP), an industry-supported technology, to allow workstations to do the following:

- Boot remotely using their network adapters to access boot code from a network location.
- Install an operating system from a remote source to a client's local hard disk.

The PXE environment is built upon Internet protocols and services that are widely used in the computer industry. This includes TCP/IP, DHCP, and TFTP. The PXE extensions to the DHCP protocol enable information to be sent to network-bootable systems and also allow these systems to locate remote installation services.

**TCP/IP network**
Your network must support TCP/IP to perform RIS installations.

**Dedicated RIS server**
You need to install Remote Installation Services on a dedicated server running Windows Server 2003. The RIS server must have at least two NTFS disk partitions — one for booting the server operating system and another to contain the directory structure for the client operating system images or distribution shares. The partition containing the images or distribution shares must be large enough to store one or more operating system images, depending on your requirements. In addition, the RIS server must run the following services: Remote Installation Services, Single Instance Store service (SIS), and the Trivial File Transfer Protocol Daemon (TFTPD).

**DNS**
RIS servers rely on the Domain Name System (DNS) to locate the required Active Directory domain controllers to facilitate domain operations. If you use Windows Server 2003 DNS, you have the benefit of dynamic updates for your DNS server. However, it is not a requirement to use Windows Server 2003 DNS for RIS to function. Whichever DNS server software you use, it must support the SRV RR record type and the dynamic update protocol specified in RFCs 2052 and 2136, respectively.

**DHCP**
RIS servers require a DHCP server on the network which is authorized and has an activated scope. Remote boot–enabled computers must receive an IP address from a DHCP server before they can contact a RIS server to request an operating system installation. You can install Windows Server 2003 DHCP or you can use the DHCP services provided with Windows 2000 Server. In addition, you can use a non-Microsoft DHCP server software.

**Active Directory**
You must install RIS on a computer running Windows Server 2003 in an Active Directory domain. For best results, configure this computer as a member server. Although you can install RIS on a domain controller, the heavy traffic load generated by RIS can impact the performance of the domain controller.

RIS uses Active Directory to locate RIS clients and other RIS servers. You can administer the RIS server from the Active Directory Users and Computers snap-in (Dsa.msc) located on the RIS server. For more information about Active Directory, see the *Directory Services Guide* of the *Windows Server 2003 Resource Kit*.

**RIS Limitations**
Although RIS is a fast, efficient automated installation technology, it has the following limitations:

**Clean installation only**
You can only use RIS to install a clean version of an operating system, with or without software applications. You cannot use RIS to upgrade an operating system or software configuration.

**Limited configuration of server components**
You might not be able to include all the server components that you want to provide with RIS installation. For example, some server components require that you install and configure them only after the RIS-based installation is complete. This can include components such as Certificate Services, the Cluster service, or software that is dependent on Active Directory.

**No support for installing domain controllers**
You cannot install a preconfigured domain controller using a RIS image. However, you can use RIS to install a stand-alone server and then configure the server as a domain controller by running the Active Directory Installation Wizard.

**No support for configuring encryption and security settings**
You cannot use RIS to deploy files that are encrypted with a system such as the Encrypting File System (EFS). Also, you cannot use RIS to deploy systems with preconfigured user-level security settings such as file and folder permissions. To configure these settings, you can run a script after completing your RIS-based installation.

**No support for wireless networks**
You cannot use RIS to deploy images over a wireless network. Wireless networks do not support remote-booting computers that use PXE technology.

**Limited support for multihomed computers**
Multihomed RIS servers are supported if the network adapters use multiple separate subnets or if all network adapters use the same subnet. In both cases, the RIS server must also be the DHCP server. The DHCP server must have active scopes for each subnet serviced and must be authorized for each IP address on the network adapters being serviced.

**Technologies Related to RIS**
RIS installation is similar to other automated installation technologies: unattended installation, image-based installation with Sysprep and Windows Server 2003, and Automated Deployment Services (ADS). Like RIS installation, these technologies are designed specifically as mass deployment solutions. Although you can use all of these technologies to automate the installation of Windows Server 2003, each of these technologies is based on a different set of installation processes, and rely on a different set of programs and tools. If you want to automate installations with minimal user interaction, use an unattended installation. To copy preconfigured operating systems and software applications, you should use an image-based installation with the Sysprep tool. RIS is typically used during large-scale deployments with minimal administrator or end user interaction.

**Unattended installation**
Unattended installation is a method of automating clean installations and upgrades with minimal user interaction. Unattended installations are particularly useful if you are:

- Upgrading a Windows server to Windows Server 2003, Standard Edition or Windows Server 2003, Enterprise Edition operating systems.
- Performing automated installations on computers that have various hardware configurations.
- Performing automated installations on specific types of servers, such as domain controllers, remote access servers, and servers that run Certificate Services or the Cluster service.
- Configuring a wide range of operating system settings during an automated installation without using batch files and scripts.

In addition to these deployment solutions, unattended installation is a useful method of creating master installations for image-based and RIS installations.

To perform an unattended installation, you first create an answer file, which is a text file that contains answers to the questions that Windows Setup normally prompts you for during an installation. After you configure your answer file, you typically create a distribution share (a folder that contains the Windows Server 2003 installation files) as well as any device drivers or other files that are required to customize the installation.

### Note

- You do not need to use a distribution share to perform an unattended installation. You can use an operating system CD instead of a distribution share.

After you have created an answer file and a distribution share, you are ready to start the unattended installation on a destination computer. To do this, you run the Windows Setup program (either Winnt.exe or Winnt32.exe), and specify the name of the answer file that you want Setup to use and the location of the distribution share that contains the installation files. Setup then runs and carries out all of the instructions that are specified in the answer file.

Unattended installation requires some upfront planning and design, but it is the most flexible and robust automated installation technology.

**Image-based installation with Sysprep**
Image-based installation with the System Preparation (Sysprep) tool is a method of copying, also known as *cloning*, preconfigured operating systems and software applications onto servers.

Image-based installation with Sysprep is a suitable automated installation technology if you need to:

- Install identical operating systems and software configurations on multiple computers.
- Install an operating system and software configuration as quickly as possible.
- Perform clean installations of an operating system, rather than upgrade an existing installation.
- Minimize end user interaction and post-installation tasks.
- Install operating systems on computers that have similar hardware and hardware abstraction layers (HALs).

To perform an image-based installation with Sysprep, you first set up a master installation — a computer with the operating system, software applications, and configuration settings that you want to install onto the destination computers in your organization. Then you run Sysprep, which prepares the master installation so that you can create a disk image (that is, a functionally identical replica of its disk) that can be copied onto multiple computers. Next, you use a non-Microsoft disk-imaging program to create the disk image of the master installation. Finally, you copy the disk image onto your destination computers.

You need two tools to perform an image-based installation with Sysprep: Sysprep, which can be found on any Windows Server 2003 operating system CD, and a non-Microsoft disk-imaging program, which you must purchase from a non-Microsoft vendor. You run Sysprep on the master computer before you create an image of the master computer's hard disk. Sysprep configures various operating system settings on the master computer to ensure that every copy of the master computer's disk image is unique when you distribute it to a destination computer. Specifically, Sysprep configures a master installation so that unique security identifiers (SIDs) are generated on each destination computer. Sysprep also configures the master computer's disk image so that every destination computer starts in a special setup mode known as *Mini-Setup*. After you copy a disk image onto a destination computer, Mini-Setup runs the first time you start the destination computer.

You use the non-Microsoft disk-imaging program to create an image of the master computer's hard disk. You also use the disk-imaging program to copy the disk image from the master computer onto a shared folder or a CD, and from the shared folder or CD onto a destination computer.

Image-based installation with Sysprep requires substantial planning and design, but it is the fastest method of installing operating systems and applications. You cannot use it to perform operating system upgrades.

**Automated Deployment Services**
You can use Automated Deployment Services (ADS) to remotely and automatically deploy Windows 2000 Server and Windows Server 2003 family operating system images onto servers without an operating system. You can also use ADS to configure, maintain, and manage servers. By using ADS, you can reduce administration costs of your data center. ADS is designed for organizations that are running data centers like service providers and corporations.

ADS reduces the number and complexity of the steps required to deploy and configure a server by using task sequences. The processes required to deploy and configure a server have been captured by a set of scripts or tasks that are easily extensible by the data center to follow their business and technical processes. Placing the configuration steps within scripts makes it easy to automatically perform the deployment or configuration of a server or group of servers based on external events (such as the arrival of a customer order, or the need to add additional servers because the existing servers are overloaded).

With ADS, you use a single server, called a "Controller," to perform large-scale server administration in your data center. The Controller, together with the other ADS services, enables you to deploy operating system images onto servers or to repurpose existing devices with new operating system images. In the data-center environment, you can use ADS to:

- Mount an image as a file and edit the image.
- Remotely purpose a device that has no operating system to a useful state or repurpose a device from one state to another state.
- Run extensible and configurable operations, such as scripts, on one or more systems from a single administration point.

- Leverage the ADS WMI Object Model to build custom applications.

# How Remote Installation Services Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Remote Installation Services Works**

**In this section**


- Remote Installation Services Architecture
- Remote Installation Services Physical Structures
- Remote Installation Services Dependencies
- Remote Boot and Installation Setup Processes
- Network Ports Used by Remote Installation Services
- Related Information


Remote Installation Services (RIS) is a Windows component that you can install with Windows Server 2003 or add at any time after the operating system is installed. RIS is an automated installation technology that you can use to create installation images of operating systems or of complete computer configurations, including desktop settings and applications. These installation images can then be made available to users at client computers. RIS is typically used during large-scale deployments when it would be too slow and costly to have administrators or end users interactively install the operating system on individual computers.

These sections provide a detailed view of how RIS works in an optimal environment. You can create an optimal environment for RIS by doing the following:

- RIS requires a significant amount of disk space. You should dedicate an entire partition or preferably an entire physical disk to store the RIS installation images.
- Use the appropriate number of Remote Installation Services (RIS) servers on your network.
- Use RIS with computers that use the Pre-Boot eXecution Environment (PXE) architecture. PXE is a remote-boot technology that allows the client computer to begin a boot sequence from the network adapter. When a PXE-enabled client computer starts for the first time, it uses the Dynamic Host Configuration Protocol (DHCP) to request an Internet Protocol (IP) address and the IP address of an active RIS server. As part of the initial request, the client computer sends out its globally unique identifier (GUID), which identifies the client computer within Active Directory.
- Dedicate an entire partition to the RIS folder tree. RIS requires a significant amount of disk space.

**Remote Installation Services Architecture**

RIS consists of several components that facilitate the remote installation of client operating systems. The Remote Installation service (Binlsvc) is the boot service that interacts with directory services to remotely boot a client computer. Trivial File Transfer Protocol Daemon (TFTPD) is a service protocol used to transfer the files needed to remotely boot, maintain, and troubleshoot a client computer. These and other RIS components are described in detail in this section.

The Setup program in Windows Server 2003 provides the basis for creating installation images using RIS. For background information about how Setup works, see Setup Technical Reference.

The following figure shows how a client that requests an installation image interacts with the RIS server to verify the client account and server settings, and then finally receives the downloaded Client Installation Wizard.

**Remote Installation Services Architecture**




**Remote Installation Service (Binlsvc)**
This service is part of the Remote Installation Services component. It detects PXE-initiated DHCP requests from RIS clients and facilitates a response to those requests. The Remote Installation service also directs clients to files on the RIS server that initiate the installation process and then responds to Client Installation Wizard requests. In addition, the Remote Installation service checks Active Directory to verify client credentials, determines if a client can be serviced, and confirms whether to create a new computer account object or reset an existing account on behalf of the client. Also, if a client that is prestaged in Active Directory has settings specifying that a particular RIS server must answer the client, then the Remote Installation service facilitates the response to that client from the specified RIS server.

The Remote Installation service was formerly known as the Boot Information Negotiation Layer (BINL) service in Windows 2000 and Windows XP Professional.

**Trivial File Transfer Protocol Daemon (TFTPD)**
This service is used by a RIS server to download the Client Installation Wizard and the initial files that are required to start the remote installation process on the client computer.

On a RIS server, TFTP is called a *daemon* or service (TFTPD), while on the client, it is referred to as a *protocol* (TFTP).

The first file that downloads is Startrom.com, which is a small startup program that displays the **Press F12 for Network Boot** prompt to the client. If the user presses F12 within 3 seconds, the Client Installation Wizard downloads to the client so the installation process can begin. The file Startrom.com is located on your RIS server in the directory path *ServerName*\RemoteInstall\OSChooser\i386\.

### Note

- For installations of Windows XP 64-Bit Edition Version 2003, the first file downloaded is Oschoice.efi. It is not necessary to press F12 for these installations.

**Single Instance Store Service**
This service consists of an NTFS file-system filter driver and a groveler agent that interacts with RIS images. The Single Instance Store (SIS) service reduces the hard-disk storage requirements for RIS images. The SIS service does this by monitoring the RIS server partition for duplicate files. To reduce the amount of disk space used by the installation folders, SIS "grovels" through the partition containing the RIS installation directory using the groveler agent. Whenever the groveler agent finds a duplicate file, the SIS service copies the original file into a directory and copies an NTFS reparse point that contains the current location, size, and attributes of the original file. (Reparse points are NTFS file system objects that have a definable attribute containing user-controlled data and that are used to extend functionality in the input/output (I/O) subsystem.) By doing this, the SIS service retains only a single instance of the file, while replacing duplicate files with links to the single instance. The SIS service then can store the duplicate files it finds in RIS images and reduce the amount of disk space that is used on your RIS server.

**RIS Components**
With the following RIS components, you can install, configure, and implement RIS in your organization.

**Remote Installation Services**
RIS is a Windows component that you can install with Windows Server 2003 or add at any time after the operating system is installed. Services that install with RIS include the Remote Installation service, TFTPD, and the SIS service.

**Remote Installation Services Setup (Risetup.exe)**
You can use Risetup.exe to initially set up the RIS server and create at least one CD-based operating system image. You can initiate the setup process from the **Start** menu of your RIS server. If you click **Remote Installation Services Setup** from the **Administrative Tools** menu, the Remote Installation Services Setup Wizard starts. The wizard performs the following actions:

- Requests preliminary information, including the installation folder name and the path to the operating system installation files.

- Copies Windows installation files.
- Updates the Client Installation Wizard screens.
- Creates a default answer file (Ristndrd.sif).
- Starts the RIS services, which include the Remote Installation service, TFTPD, and the SIS service.
- Authorizes the DHCP server.
- Risetup is also used to create any additional CD-based operating system images after the initial installation is created.

**Remote Installation Preparation Wizard (Riprep.exe)**
You can use Riprep.exe to create a customized image of an operating system such as Windows XP Professional. To create an image means that you create a replica of a hard disk that you can install on other computers in your organization. You use Riprep to create an image of an existing operating system installation on a master computer and then replicate that image onto an available RIS server on your network. The image can include the operating system with default parameters applied by the administrator, or the operating system with a preconfigured desktop, locally-installed applications, and drivers.

A 64-bit version of the Remote Installation Preparation (RIPrep) Wizard is included with the 64-bit versions of Windows Server 2003. For more information about using the RIPrep Wizard to create an installation image of a 64-bit operating system, see article 891128, "Updated Remote Installation Service (RIS) functionality in Windows Server 2003 Service Pack 1," in the Microsoft Knowledge Base.

**Remote Boot Floppy Generator (Rbfg.exe)**
You can use Rbfg.exe to create remote-boot floppy disks for RIS client computers that are not PXE-enabled, so that these clients can emulate the remote boot process and install an operating system over the network using RIS. However, for RIS clients that are not PXE-enabled to use the remote-boot floppy disk, they must each have a supported Peripheral Component Interconnect (PCI)–based network adapter supported by the Rbfg.exe utility. PC Card adapters, CardBus adapters, and adapters that are not PCI-based are not supported. You can view the list of supported network adapters by starting the Rbfg.exe utility and clicking **Adapter List**. The client computer's basic input/output system (BIOS) must be configured to use the network adapter as its primary boot device.

**Client Installation Wizard (OSChooser)**
OSChooser is a service of the Client Installation Wizard that is run by the client computer. It is a text-based program downloaded by the RIS server to the client that allows the client to communicate with the RIS server during the installation process. Remote Installation service is the server component that sends a default set of Client Installation Wizard screens to guide the client through the remote installation process. Clients that can remote boot use this wizard to log on and select from operating system installation options. You can customize these setup screens to meet the needs of your organization.

**Active Directory Users and Computers Extension for RIS (Dsa.msc)**
When you create the RIS server, the Active Directory Users and Computers extension installs on the RIS server. The extension provides a **Remote Install** tab within the computer account **Properties** dialog box of each RIS server that you can use to manage the RIS server. You can start this extension by specifying the Microsoft Management Console (MMC) snap-in Dsa.msc in the **Run** dialog box or you can start it from the command line.

You can use the Active Directory Users and Computers extension to manage RIS locally or through a Terminal Services session on another network computer. You can also administer RIS from a computer running Windows XP Professional if you install the Adminpak.msi on that computer.

**Remote Installation Services Physical Structures**
The physical structure of RIS depends on whether you are using a CD-based or an image-based installation and whether you are also using an answer file.

**CD-based RIS installations**
Using Risetup.exe, you can create a CD-type image, also called a flat image, to be used with RIS installations on client computers. Using a flat image is similar to setting up a client directly from the installation CD, except that the operating system image source files reside on an available RIS server.

Use an installation image created by Risetup if you want to distribute the network equivalent of CD-based installation functionality. A Risetup image is a replica of an operating system CD file structure, located across the network on a remote

server using RIS.

You create Risetup images by running the Risetup Wizard on a RIS server, while using an operating system CD to create the image. When using Risetup images, you cannot provide a fully-configured clone of an operating system with applications and desktop customizations, as you can with Riprep images. However, you can add applications and drivers to the distribution folder where the Risetup images are located and use answer files to install the applications and specify the location of drivers.

Installing a Risetup image is similar to setting up a workstation directly from a CD. However, the source files are located across the network on a RIS server. The following figure shows the directory structure of your RIS server under the RemoteInstall folder where Risetup images are stored. You can define the name of the folder <imagename> where the images are located.

**Directory Structure of a RIS Server**




When you install remote installation services on your RIS server, it automatically creates one Risetup image of the server operating system and stores it under the Images folder within the RIS directory structure, as shown in the figure. This image is available to remote boot–enabled clients. Clients that request installation of an operating system can access Risetup images on a remote RIS server, if you configure them to do so.

You can make additional Risetup images using operating system CDs for Windows 2000 Professional, Windows 2000 Server, and Windows 2000 Advanced Server, in addition to Windows XP Professional and Windows Server 2003. To create a Risetup image, place the CD in the CD drive of your RIS server and run the Risetup Wizard from the **Images** tab of RIS server **Properties**. You can also specify the following command string at the command-line interface to start the Risetup Wizard:

```
risetup -add
```

You can create and associate multiple answer files with Risetup images, which allows you to customize the applications and drivers that you want to install with each image. However, you cannot include preconfigured application or desktop configurations with a Risetup image. Also, Risetup images take longer to install than an equivalent size Riprep image.

**Image-based RIS installations**

Use an image created by Riprep if you want to distribute an image of a fully-configured workstation complete with applications. A Riprep image is essentially a file system image that is located on a remote RIS server. It is similar to the hard disk-images you create using a non-Microsoft disk-imaging tool and the Windows System Preparation tool (Sysprep).

You create Riprep images by running the Riprep Wizard (Riprep.exe) on a master computer which has the operating system configuration, applications and settings, and desktop customizations that you want to deploy to client computers in your organization.

Riprep images are most useful for cloning a standard operating system configuration to clients. Riprep images generally require more disk space on your RIS server than Risetup images because they usually include preconfigured applications and tools. However, they install faster than equivalent size Risetup images.

To create a Riprep-based installation, you must first set up a *master installation*. This is the reference computer that contains the operating system, software applications, and configuration settings you plan to install on destination computers in your organization. After you configure the master installation, you run Riprep.exe, which is on the server at the following location: *servername*\reminst\admin\i386\riprep. This converts the master installation into a remote installation image, which is a functionally identical replica of the master computer disk that you can install on multiple destination computers. Riprep also replicates the image to a RIS server where it is available for installation on remote boot–enabled client computers. Clients who request installation of an operating system can access Riprep-based images on a remote RIS server if you configure them to do so.

The best way to install the operating system on your master computer is to use RIS with an unattended installation. However, you can also install the operating system locally using the appropriate operating system CD. If you do this, use the disk-partitioning utility on the Windows Server 2003 installation CD and use the text-mode Setup to clear the disk partition and ensure a clean installation.

Riprep configures various operating system settings on the master computer to ensure that every copy of the master computer's disk image is unique when you install it on destination computers. This includes resetting the security identifiers (SIDs) and access control lists (ACLs). Riprep also configures the master installation image so that, after the initial installation of the image, every destination computer starts in a special setup mode known as Mini-Setup.

**RIS Client HAL Types**

If you want to create an image-based RIS installation using the tool Riprep.exe, you first need to determine if your RIS clients have a Hardware Abstraction Layer (HAL) that is compatible with the master computer from which you create your image. For example, if the master computer where you run Riprep.exe has an Advanced Configuration and Power Interface (ACPI) HAL, then the client computers you designate to receive operating system images that are generated from that master computer must also have an ACPI HAL. The HAL type is indicated by the original file name of the file Hal.dll.

You need to verify how many different HAL types exist in your organization. This determines how many different master installations you will need with HALs that are compatible with the client HALs in your organization. To verify the type of HAL on your client computers, you can do one of the following:

- Use a management tool such as Systems Management Server (SMS) to obtain your client inventory, from which you can determine the HAL types.
- View the properties of Hal.dll to determine the HAL types.

You can install RIS-based operating system images if any of the following conditions are true regarding HALs:

- The master and destination computer HALs are identical.
- The master and destination computers both have either uniprocessor or multiprocessor Advanced Programmable Interrupt Controller (APIC) HALs.
- The master and destination computers both have either uniprocessor or multiprocessor Advanced Configuration and Power Interface (ACPI) HALs.

**Using answer files with a RIS installation**

Using answer files with a RIS installation is similar to using an unattended answer file with RIS-specific sections and values. For a CD-based installation, the default answer file is Ristndrd.sif, and for an image-based installation, the default answer file is Riprep.sif. Multiple answer files can be associated with an installation. You can create the answer files by using Setup Manager (Setupmgr.exe) or you can create it manually by using a text editor, such as Notepad.

As with most answer files, Ristndrd.sif and Riprep.sif contain multiple predefined sections that you can modify. Section names are always enclosed in square brackets (for example, [Unattended]).

### RIS-Specific Sections in an Answer File

| Section | Description |
|---|---|
| [RemoteInstall] | Contains entries Repartition and UseWholeDisk. |
| [Oschooser] | Contains entries for values which are displayed after the operating system is selected. For example, the image type and start file. |

For additional information about unattended installation, see Unattended Installation Technical Reference.

**Remote Installation Services Dependencies**

Remote Installation Services (RIS) uses or depends upon a variety of technologies and components.

**Network infrastructure**

RIS depends on a network infrastructure that can identify computers on the network. RIS therefore requires that DHCP, DNS, and Active Directory be available on the network.

**Remote-boot technology**
RIS uses remote-boot technology to enable a computer with no operating system to begin a boot sequence from the network adapter, after which an operating system can be installed across the network. This remote-boot technology is called the Pre-Boot eXecution Environment (PXE). For client computers that are not PXE-enabled—that is, pre–Net PC/PC98 computers—you can use a special remote-boot disk, Rbfg.exe, if the client computer has a supported network adapter. Most computers that conform to the Net PC or PC98 specifications have a PXE remote boot–enabled network adapter and remote boot–enabled BIOS. A computer on which you use the remote boot disk must have a Peripheral Component Interconnect (PCI)–based network adapter supported by the Rbfg.exe utility. PC Card adapters, CardBus adapters, and adapters that are not PCI–based are not supported. You can view the list of supported network adapters by starting the Rbfg.exe utility and clicking Adapter List.

**Technology to create an installation image**
RIS includes a technology that you can use to create an image of an operating system that will be installed on client computers. You can create the image in one of two formats: a flat image or a Remote Installation Preparation Wizard (RIPrep) image format. Use the flat image option to create an image directly from a set of operating system files, for example, the files on a CD. Use the RIPrep image format to create an image that includes an operating system with specific settings and applications, such as an image that complies with a corporate desktop standard.

**Technologies for communication between RIS client and server**
RIS includes technologies that establish communication between a client computer that has begun a PXE boot sequence and the RIS server that contains installation images that are available to that client. For this process, RIS uses DHCP to provide the client with an Internet Protocol (IP) address. RIS then downloads the Client Installation Wizard, which prompts the user to log on and provides a menu of installation options that are customized for that user. You can control these images through Group Policy.

**PXE architecture**
Remote Installation Services (RIS) uses the Dynamic Host Configuration Protocol (DHCP) that follows the Pre-Boot eXecution Environment (PXE) architecture to start a client computer.

When a new PXE remote boot–enabled client computer starts for the first time, it uses DHCP to request an Internet Protocol (IP) address and the IP address of an active RIS server. As part of the initial request, the client computer sends out its globally unique identifier (GUID), which identifies the client computer within Active Directory.

After the client computer receives an IP address from the DHCP server, it then requests service from a RIS server. The request to the RIS server is a broadcast, which means all of the RIS servers on the network will receive the request. If the RIS server can answer, it queries its preferred domain controller for the client computer's GUID. If the server fails to find the GUID, it queries the global catalog for the forest. If the GUID is found in either location, the client computer is considered "known". If the GUID is not found, it is considered "unknown". In this case —that is, the query by the RIS server found no computer account object in Active Directory with this computer's GUID —only RIS servers that are configured to answer unknown clients will do so.

If the client is known, all the available RIS servers query the domain to determine whether the client computer account has been prestaged with settings that specify that the client computer is to be answered by a particular server. If so, only the designated server answers the service request. The other servers respond by telling the client which server will answer it. This is called a referral. If the client computer account settings do not require the client computer to be answered by a particular server, all RIS servers answer it if they are aware of this client and configured to answer it. Each RIS server then offers its own files to download.

After the client computer receives a reply to its service request, it initiates a Trivial File Transfer Protocol (TFTP) download of the boot program. The Remote Installation Services implementation of this program is called Startrom.com (located at \\RIS_server_name\REMINST\oschooser\i386\startrom.com).

After Startrom.com has been downloaded, the client computer runs it. The default version of Startrom.com then prompts the user to press F12 to initiate a network installation. If the user fails to press F12 within three seconds, the network boot is stopped, and the client computer attempts to start up from the next available boot device. When the user presses F12, the client computer uses TFTP to download the Client Installation Wizard. The wizard then requests the user to log on to the network.

**Remote Installation Services directory service planning**
The Remote Installation Services (RIS) environment relies on a well-designed and well-planned Active Directory architecture. Normally, determining the physical location of a server can be a challenge. This can become even more difficult when you must find and configure several servers that are located in different buildings, different offices, or on different floors. Using Active Directory, however, can simplify this process.

In a given domain, Active Directory provides organizational units or containers that you can use to organize users and resources into logical administrative groups. This makes it easier to locate and configure servers at multiple locations.

**Remote Boot and Installation Setup Processes**

RIS uses PXE technology to enable RIS client computers without an operating system to initiate the boot sequence from their network adapters, thus facilitating operating system installations from remote network locations. To initiate the remote boot process and set up a RIS-based operating system installation, PXE interacts with the Dynamic Host Configuration Protocol (DHCP), the Remote Installation services, and TFTPD, as shown in the following figure:

**How RIS Works**




When you start a new PXE-enabled RIS client computer, the following sequence of events occurs:

1. The client computer initiates the communication by sending a DHCP Discover broadcast on its subnet. A DHCP server with an active scope for that subnet will issue an IP address to the client.

2. All Remote Installation servers that receive the client's DHCP Discover broadcast extract (from the PXE data portion of the packet) the universally unique identifier (UUID) of the client that is requesting service. The RIS server then queries its preferred domain controller to search for this UUID in all prestaged computer accounts in Active Directory.

   If the domain controller does not find the UUID in the local domain, the RIS server queries the global catalog to locate the client computer account. If the UUID is found in either location, the client computer is recognized as a known client; otherwise, it is considered an unknown client. If the client is unknown, it will only receive an answer from a RIS server that is configured to answer unknown clients, provided that one exists on the network.

3. If the client is known, all available RIS servers query the domain to determine whether the prestaged client computer account has a setting that specifies that only a particular RIS server can answer the client.

   If this is the case, then only the designated RIS server answers the service request, and other RIS servers simply notify the client of the particular RIS server that is configured to answer it. If the client computer account does not have a setting that requires it to be answered by a particular server, any RIS server can answer the request. However, the client only receives service from the first RIS server it contacts.

4. The user receives a prompt to press the F12 key to initiate a network service boot request from the RIS server.

5. Using TFTPD, the contacted RIS server downloads the Client Installation Wizard to the RIS client, along with all client dialog boxes contained within the wizard.

6. The Client Installation Wizard prompts the user to log on with a valid user name, password, and domain name.

7. The user receives a selection of operating system images that are hosted on the RIS server for installation on the client computer.

   The list of operating system images that are offered to the user is based on the user's credentials or security group membership.

**PXE Specification**

The published PXE specification defines the remote-boot process and also establishes the PXE compliance standards for hardware manufacturers and other vendors. RIS uses PXE environment extensions to DHCP, an industry-supported technology, to allow workstations to do the following:

- Boot remotely using their network adapters to access boot code from a network location.
- Install an operating system from a remote source to a client's local hard disk.

The PXE environment is built upon Internet protocols and services that are widely used in the computer industry. This includes TCP/IP, DHCP, and TFTP. The PXE extensions to the DHCP protocol enable information to be sent to systems that support remote network booting and also allow these systems to locate remote installation services.

**Note**

- Network adapters that meet the PXE .99n specification will work correctly with RIS.

**RIS Technology Limitations**

You can use RIS technology to install operating systems, with or without software applications, to portable and desktop computers in your organization, which include member servers, stand-alone servers, and domain controllers. However, limitations to the scope of RIS-based operating system installations include the following:

**Clean Installs**

You can only use RIS to provide a clean version of an operating system, with or without software applications. You cannot use RIS to upgrade an operating system or software configuration.

**Server Components**

If you use RIS to install a server operating system, you might not be able to include all the server components that you want to provide with the RIS image. For example, some server components require that you install and configure them only after the RIS-based installation is complete. This can include components such as Certificate Services, Cluster service, or software that is dependent on Active Directory.

**Domain controllers**

You cannot install a preconfigured domain controller using a RIS image. However, you can use RIS to install a stand-alone server and then configure the server as a domain controller by running the Active Directory Installation Wizard.

**Encryption and security settings**

You cannot use RIS to deploy files that are encrypted with a system such as the Encrypting File System (EFS). Also, you cannot use RIS to deploy systems with preconfigured user-level security settings such as file and folder permissions. To configure these settings, you can run a script after completing your RIS-based installation.

**Wireless networks**

Wireless networks do not support remotely booting computers using PXE technology.

**Multihomed computers**

Multihomed RIS servers are supported if the network adapters use multiple separate subnets or if all network adapters service the same subnet. In both cases the RIS server must also be the DHCP server. The DHCP server must have active scopes for each subnet serviced and must be authorized for each IP address on the network adapters being serviced.

**Supported operating systems**

RIS has certain limitations depending on the operating system that you are installing. For more information about operating systems supported by RIS, see "Operating Systems Supported by Remote Installation Services" in Help and Support Center for Windows Server 2003.

**Network Ports Used by Remote Installation Services**

The following table lists the ports that are used by RIS.

### Port Assignments for Remote Installation Services

| Service Name | UDP | TCP |
|---|---|---|
| Dynamic Host Configuration Protocol (DHCP) | 67 | N/A |
| Boot Information Negotiation Layer (BINL) | 4011 | N/A |

| | | |
|---|---|---|
| Trivial File Transfer Protocol Daemon (TFTPD) | 69 | N/A |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Setup Technical Reference
- Unattended Installation Technical Reference
- Sysprep Technical Reference

# Remote Installation Services Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services Tools and Settings**

### In this section


- Remote Installation Services Tools
- Remote Installation Services Group Policy Settings
- Remote Installation Services Answer-File Settings
- Network Ports Used by Remote Installation Services
- Related Information


**Remote Installation Services Tools**

The following tools are associated with Remote Installation Services.

**Risetup.exe: Remote Installation Services Setup program**

**Category**

Tool included in the Windows Server 2003 family.

**Version compatibility**

Available on computers running the Windows Server 2003 family of products or the Windows 2000 Server family.

This component is used to initially set up the RIS server and create at least one CD-based operating system image. You can initiate the setup process from the Start menu of your RIS server. By clicking **Remote Installation Services Setup** from the Administrative Tools group, a wizard starts and does the following:

- Requests preliminary information, including the installation folder name and the path to the operating system installation files.
- Copies Windows installation files.
- Updates the Client Installation Wizard screens.
- Creates a default answer file (Ristndrd.sif).
- Starts RIS services.
- Authorizes the DHCP server.

For more information about Risetup.exe, see “Command-Line References” in the Tools and Settings Collection.

### Note

- Risetup is also used to create any additional CD-based operating system images after the initial installation is created.

**Riprep.exe: Remote Installation Preparation Wizard**

**Category**

Tool included in the Windows Server 2003 family.

**Version compatibility**

Available on computers running the Windows Server 2003 family of products or the Windows 2000 Server family.

Riprep.exe is used to create a customized image of an operating system such as Windows XP Professional. To create an image means that you create a replica of a hard disk that you can install on other computers in your organization. You use Riprep to image an existing Windows XP Professional operating system installation on a master computer and replicate that image to an available RIS server on your network. The image can include the operating system with default parameters applied, or the operating system with a preconfigured desktop, locally installed applications, and drivers.

For more information about Riprep.exe, see Windows Server 2003 Help.

**Rbfg.exe: Remote Boot Floppy Generator**
**Category**
Tool included in the Windows Server 2003 family.

**Version compatibility**
Available on computers running the Windows Server 2003 family of products or the Windows 2000 Server family.

Rbfg.exe is used to create remote-boot floppy disks for RIS client computers that are not PXE-enabled, so that these clients can emulate the remote boot process and install an operating system over the network using RIS. However, for RIS clients that are not PXE-enabled to use the remote-boot floppy disk, they must each have a supported Peripheral Component Interconnect (PCI) network adapter–based network adapter supported by the Rbfg.exe utility. PC Card adapters, CardBus adapters, and adapters that are not PCI–based are not supported. You can view the list of supported network adapters by starting the Rbfg.exe utility and clicking Adapter List. The client computer's basic input/output system (BIOS) must be configured to use the network adapter as its primary boot device.

For more information about Riprep.exe, see Windows Server 2003 Help.

**OSChooser.exe: Client Installation Wizard**
**Category**
Tool included in the Windows Server 2003 family.

**Version compatibility**
Available on computers running the Windows Server 2003 family of products or the Windows 2000 Server family.

OSChooser.exe is a service of the Client Installation Wizard that is run by the client computer. It is a text-based program downloaded by the RIS server to the client that allows the client to communicate with the RIS server during the installation process. Remote Installation service is the server component that sends a default set of Client Installation Wizard screens to guide the client through the remote installation process. Clients that can remote boot use this wizard to log on and select from operating system installation options. You can customize these setup screens to meet the needs of your organization.

For more information about Riprep.exe, see Windows Server 2003 Help.

**DSA.msc: Active Directory Users and Computers Extension for RIS**
**Category**
Tool included in the Windows Server 2003 family.

**Version compatibility**
Available on computers running the Windows Server 2003 family of products or the Windows 2000 Server family.

When you create the RIS server, the Active Directory Users and Computers extension installs on the RIS server. The extension provides a **Remote Install** tab within the computer account **Properties** dialog box of each RIS server that you can use to manage the RIS server. You can start this extension by specifying the Microsoft Management Console (MMC) snap-in Dsa.msc in the **Run** dialog box or you can start it from the command line.

You can use the Active Directory Users and Computers extension to manage RIS locally or through a Terminal Services session on another network computer. You can also administer RIS from a computer running Windows XP Professional if you install the Adminpak.msi on that computer.

For more information about Riprep.exe, see Windows Server 2003 Help.

**Remote Installation Services Group Policy Settings**
You can use Group Policy to configure certain installation options presented by the Client Installation Wizard. With Group Policy, you can create different Group Policy objects to apply to specific users. The following table lists and describes the Group Policy settings that are associated with RIS. Each of these settings has three different settings associated with it: Enabled, Disabled, and Not Configured.

### Group Policy Settings Associated with RIS

| Group Policy Setting | Description |
|---|---|
| Automatic Setup | Predefines the computer name and a location within Active Directory for the client computer accounts. (Another alternative is prestaging client computer accounts.) |
| Custom Setup | Allows users to define a unique name for this computer and specify where the computer account will be created within Active Directory. |
| Restart Setup | Restarts an operating system installation attempt if it fails during the text-mode portion of the setup process. |
| Tools | Gives users the ability to access tools, including the Recovery Console, from the Client Installation Wizard. |

For more information about Group Policy settings, see "Group Policy Settings Reference" in the Tools and Settings Collection.

**Remote Installation Services Answer-File Settings**

To simplify the process of setting up an operating system on multiple computers, you can run Setup unattended. To run Setup unattended, you create and use an answer file, a customized script that answers the Setup questions automatically. The unattended installation Setup answer files (.sif) are text files that describe the installation process for an installation image. An unattended installation template, which you can customize, is provided by RIS. The .sif file must be associated with one of the available installation images.

[OsChooser] and [RemoteInstall] are specific answer file entries for RIS. These are used in addition to the other answer file entries in an unattended installation. For information about unattended installation, see "Unattended Installation Technical Reference."

The section below describes the settings that you can configure in an unattended installation answer file. For more information about how to structure an answer file, see "Unattended Installation Physical Structures" in How Unattended Installation Works.

The following table lists all of the valid sections that can appear in an answer file for a RIS installation.

**Answer File Sections for RIS Installation**

| Section | Description |
|---|---|
| [adapter_specific_protocol_section] | Contains an entry for specifying a network component. Must be used in conjunction with the [Networking] section. |
| [Branding] | Contains entries for branding Microsoft Internet Explorer during an unattended installation. If you do not specify this section, Setup uses the default for Windows Server 2003. |
| [CertSrv_Client] | Contains entries for the Web client component of Certificate Services. You do not need to include this section in the answer file to install the Web client component of Certificate Services on the computer. Enabling Certificate Services in the [Components] section of the answer file installs the necessary files. |
| [CertSrv_Server] | Contains entries for specifying the server components of Certificate Services. The entries in this section are supported only for the Windows Server 2003. |
| [Components] | Contains entries for installing the optional components of Windows XP Professional and Windows Server 2003. |
| [Data] | Contains entries for starting directly from the Windows Server 2003 product CD when performing an unattended installation. |
| [DCInstall] | Contains entries for installing a domain controller after the initial setup of the operating system finishes. |
| [Display] | Contains entries for specifying display settings for graphics devices. |
| [FavoritesEx] | Contains the default Favorites settings for the Windows browser. If not present, Setup uses the default Uniform Resource Locator (URL) for Windows XP or Windows Server 2003. |

| | |
|---|---|
| [Fax] | Contains entries for specifying custom installation information for the Fax service. |
| [GuiRunOnce] | Contains commands that run the first time a user logs on to the computer after final (GUI-mode) stage of Setup completes. |
| [GuiUnattended] | Contains entries for preparing the graphical user interface (GUI) for unattended installation. |
| [Homenet] | Contains entries for installing Home Networking settings for network adapters, Internet Connection Sharing (ICS), and Internet Connection Firewall (ICF). |
| [Identification] | Contains entries for specifying the network identification of a computer. If these entries are not present, Setup adds the computer to the default workgroup. If there is not enough information for this entry, Setup prompts the user to provide this information. |
| [IEHardening] | Contains entries for configuring Web content zones. |
| [InternetServer] | Contains entries for configuring the installation folder for Internet Information Services (IIS) components. |
| [LicenseFilePrintData] | Contains entries for installing licensing information for Windows Server 2003. |
| [MassStorageDrivers] | Contains an entry for loading retail or original equipment manufacturer (OEM) small computer system interface (SCSI) drivers in the Setup Loader during the intermediate (text-mode) stage of Setup |
| [MS_AppleTalk parameters] | Contains entries for configuring AppleTalk. |
| [MS_ATMArps parameters] | Contains no entries. |
| [MS_ATMLANE parameters] | Contains no entries. |
| [MS_ATMUni parameters] | Contains no entries. |
| [MS_L2TP parameters] | Contains entries for the Layer 2 Tunneling Protocol (L2TP). |
| [MS_MSClient parameters] | Contains entries for Client for Microsoft Networks. |
| [MS_NetMon parameters] | Contains no entries. |
| [MS_NWClient parameters] | Contains entries for Client Service for NetWare. |
| [MS_NWIPX parameters] | Contains entries for the Internetwork Packet Exchange (IPX) protocol. Its parameters are network adapter–independent. |
| [MS_NwSapAgent parameters] | Contains no entries. |
| [MS_PPTP parameters] | Contains entries for the Point-to-Point Tunneling Protocol (PPTP). |
| [MS_Psched parameters] | Contains no entries. |
| [MS_RAS parameters] | Contains entries for the Routing and Remote Access service. |
| [MS_RasSrv parameters] | Contains entries for the Routing and Remote Access service for remote access servers. |
| [MS_Server parameters] | Contains entries for the File and Print Service. |
| [MS_TCPIP parameters] | Contains entries for configuring the TCP/IP protocol. |
| [MS_WLBS parameters] | Contains entries for configuring Network Load Balancing (formerly, Windows Load Balancing Service, or WLBS) and binding Network Load Balancing to multiple network adapters. |
| [NetAdapters] | Contains entries for network adapters that you want to install. |

| | |
|---|---|
| [NetBindings] | Contains entries for disabling, enabling, and demoting binding rules for network components such as services, protocols, and adapters. |
| [NetClients] | Contains an entry for specifying the network clients to install. |
| [NetOptionalComponents] | Contains entries for installing optional network components. |
| [NetProtocols] | Contains an entry for installing the network protocols. |
| [NetServices] | Contains entries for installing network services. |
| [Networking] | Contains no entries. However, to configure network protocols in unattended installation, you must include the [Networking] section name in your answer file. |
| [OEMBootFiles] | Contains entries for specifying OEM-supplied boot files. |
| [OsChooser] | Contains entries for selecting the operating system. Only clients installed with access to Remote Installation Services (RIS) can use these entries. |
| [params.adapter_instance] | Contains entries for specifying an adapter's instance on a computer by using either PCI location information or a Plug and Play device ID. |
| [PCHealth] | Contains entries for configuring error reporting, Remote Assistance, and Help and Support Center. |
| [protocol_parameters_section] | Contains an entry for specifying network protocol parameters that are also specified in the [NetProtocols] section. |
| [Proxy] | Contains entries for specifying proxy server settings for the Windows XP or Windows Server 2003 browser. |
| [RegionalSettings] | Contains entries for installing regional options during the setup process. |
| [RemoteInstall] | Contains entries for Remote Installation Services (RIS). You can only use this section for client operating systems that are installed by RIS. |
| [SetupData] | Contains the OSLoadOptionsVar entry, which is used to modify the /NoExecute option in Boot.ini used by Data Execution Prevention (DEP). |
| [SetupMgr] | Contains entries for specifying the server and folder in which the distribution share resides and the name of the distribution share. This section is used only by Setup Manager when an existing distribution share is edited. It is not used by Setup during an unattended installation. |
| [SetupParams] | Contains an entry for running an additional post-Setup command. |
| [Shell] | Contains entries for specifying the visual style of the user interface. |
| [SNMP] | Contains entries for specifying the Simple Network Management Protocol (SNMP) service. |
| [SystemFileProtection] | Contains entries for specifying Windows File Protection (WFP) settings. |
| [SystemRestore] | Contains entries for specifying the System Restore feature. |
| [TapiLocation] | Contains entries for specifying the telephony API (TAPI). It is valid only if a modem is present on the computer. |
| [TerminalServices] | Contains entries for configuring Terminal Services features (Terminal Server and Remote Desktop). |
| [UDDIDatabase] | Contains entries for configuring the optional Universal Description, Discovery and Integration (UDDI) Database feature. |

| | |
|---|---|
| [UDDIWeb] | Contains entries for configuring the optional Universal Description, Discovery and Integration (UDDI) Web services feature. |
| [Unattended] | Contains entries for running Setup. |
| [Uninstall] | Contains entries for allowing users who upgrade to Microsoft Windows XP to remove the operating system if they encounter problems. |
| [URL] | Contains entries for specifying default Uniform Resource Locator (URL) settings for the browser available in Windows XP or Windows Server 2003. |
| [UserData] | Contains entries for specifying user settings during the setup process. |
| [WindowsFirewall | This section must be included in your answer file to configure firewall sections and entries in unattended installations. To configure Windows Firewall settings, you must include the [WindowsFirewall] section along with one or more of the other Windows Firewall settings sections. |
| [WindowsFirewall.profile_name] | Contains entries for specifying a domain type of profile, a standard type, or a profile that applies to both domain and standard. |
| [WindowsFirewall.program_name] | Contains entries for adding programs to the Windows Firewall Exceptions list. |
| [WindowsFirewall.service_name] | Contains entries for adding services to the Windows Firewall Exceptions list. |
| [WindowsFirewall.portopening_name] | Contains entries for adding static ports to the Windows Firewall Exceptions list. |
| [WindowsFirewall.icmpsetting_name] | The default configuration for Windows Firewall blocks all Internet Control Message Protocol (ICMP) message types, the entries in this section allows you to modify this behavior. |
| [Win9xUpg] | Contains entries for specifying an unattended upgrade from Windows 98, Windows 98 Second Edition, or Windows Millennium Edition to either Windows XP Home Edition or Windows XP Professional. |

**Network Ports Used by Remote Installation Services**

The following table lists the ports are used by RIS.

### Port Assignments for Remote Installation Services

| Service Name | UDP | TCP |
|---|---|---|
| Dynamic Host Configuration Protocol (DHCP) | 67 | N/A |
| Boot Information Negotiation Layer (BINL) | 4011 | N/A |
| Trivial File Transfer Protocol Daemon (TFTPD) | 69 | N/A |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Setup Technical Reference
- Unattended Installation Technical Reference
- Sysprep Technical Reference

# Group Policy Collection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Collection**
Group Policy is an infrastructure that allows you to implement specific configurations for users and computers. Group Policy settings are contained in Group Policy objects (GPOs), which are linked to the following Active Directory directory service containers: sites, domains, or organizational units (OUs). The settings within GPOs are then evaluated by the affected targets, using the hierarchical nature of Active Directory. Consequently, Group Policy is one of the top reasons to deploy Active Directory because it allows you to manage user and computer objects.

Group Policy is one of a group of management technologies, collectively known as IntelliMirror management technologies, which provide users with consistent access to their applications, application settings, roaming user profiles, and user data, from any managed computer—even when they are disconnected from the network. IntelliMirror is implemented through a set of Microsoft Windows features, including Active Directory, Group Policy, Software Installation, Windows Installer, Folder Redirection, Offline Folders, and Roaming User Profiles.

This collection includes detailed information about each of the following areas of Group Policy:

- Core Group Policy
- Group Policy Components
- Group Policy Administrative Tools

This page introduces Group Policy management concepts and architecture, summarizes the areas included in the Group Policy collection, and describes Group Policy scenarios.

**Group Policy Management**
Administrators face increasingly complex challenges in managing the IT infrastructure. You must deliver and maintain customized desktop configurations for more types of workers such as mobile users, information workers, or others assigned to strictly defined tasks, such as data entry. Security settings and updates must be delivered efficiently to all the computers and devices in the organization. New users need to be productive quickly without costly training. In the event of a computer breakdown or disaster, service must be restored with a minimum of data loss and interruption. All of these tasks, known collectively as Change and Configuration Management, must be achieved at the lowest possible cost. You need to be able to implement change quickly and affect large numbers of users and computers. Group Policy is the infrastructure that allows you to implement change on the object level in Active Directory.

You need to be able to define configurations once and rely on the operating system to enforce that state. With Active Directory, GPOs can be linked to sites, domains, and OUs, allowing Group Policy settings to be applied to users and computers. In addition, GPOs can be used to help manage server computers, through many server-specific operational and security settings. This infrastructure provides a high degree of flexibility, allowing you to customize configurations, such as delivering a specific piece of software to specialized users based on their membership in an OU. In addition, the Group Policy Management Console (GPMC) simplifies implementation and management of Group Policy.

**Group Policy Architecture**
Group Policy uses a document-centric approach to creating, storing, and associating Group Policy settings. Similar to the way in which Microsoft Word stores information in .doc files, Group Policy settings are contained in GPOs. A GPO is a virtual object; policy-setting information is stored in two locations: the Active Directory container to which the GPO is linked, and the Sysvol on the domain controller.

Group Policy is configured primarily through the use of two tools: Group Policy Object Editor, (previously known as the Group Policy snap-in, Group Policy Editor, or Gpedit) and Group Policy Management Console (GPMC), available for download from the Microsoft Web site. Whereas Group Policy Object Editor is used to configure and modify settings within GPOs, GPMC is used to create, view, and manage GPOs. Group Policy architecture is shown in the following diagram, which shows how the primary components interact through read or write access. Components are described in the figure below.

**Group Policy Architecture**




**Group Policy Components**

| Component | Description |
|---|---|
| Server (Domain Controller) | In an Active Directory forest, the domain controller is a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. |
| Active Directory | Active Directory, the Windows-based directory service, stores information about objects in a network and makes this information available to users and network administrators. Administrators link GPOs to Active Directory containers such as sites, domains, and OUs that include user and computer objects. In this way, Group Policy settings can be targeted to users and computers throughout the organization. |
| Group Policy object (GPO) | A GPO is a collection of Group Policy settings, stored at the domain level as a virtual object consisting of a Group Policy container (GPC) and a Group Policy template (GPT). The GPC, which contains information on the properties of a GPO, is stored in Active Directory on each domain controller in the domain. The GPT contains the data in a GPO and is stored in the Sysvol in the **/Policies** sub-directory. GPOs affect users and computers that are contained in sites, domains, and OUs. |
| Sysvol | Sysvol is a shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain. The Sysvol contains the data in a GPO: the GPT, which includes Administrative Template-based Group Policy settings, security settings, script files, and information regarding applications that are available for software installation. It is replicated using the File Replication Service (FRS). |
| Local Group Policy object | The local Group Policy object (local GPO) is stored on each individual computer, in the hidden **%systemroot%\System32\GroupPolicy** directory. Each computer running Windows 2000, Windows XP Professional, Windows XP 64-Bit Edition, Windows XP Media Center Edition, or Windows Server 2003 has exactly one local GPO, regardless of whether the computers are part of an Active Directory environment.<br><br>Local GPOs do not support certain extensions, such as Folder Redirection or Group Policy Software Installation. Local GPOs do support many security settings, but the Security Settings extension of Group Policy Object Editor does not support remote management of local GPOs. Local GPOs are always processed, but are the least influential GPOs in an Active Directory environment, because Active Directory-based GPOs have precedence.<br><br>Although you can configure local GPOs on individual computers, the full power of Group Policy can only be realized in a Windows Server network with Active Directory installed. In addition, some features and Group Policy settings require client computers running Windows XP. |
| Group Policy Object Editor | Group Policy Object Editor is a Microsoft Management Console (MMC) snap-in that is used to edit GPOs. It was previously known as the Group Policy snap-in, Group Policy Editor, or Gpedit. |

| | |
|---|---|
| Server-Side Snap-Ins | The MMC snap-in is loaded, by default, in Group Policy Object Editor. Server-side snap-in extensions provide the user interface to allow you to configure various policy settings while client-side extensions implement the actual policy settings on target client computers.<br><br>Snap-in extensions include Administrative Templates, Scripts, Security Settings, Software Installation, Folder Redirection, Remote Installation Services, Internet Explorer Maintenance, Disk Quotas, Wireless Network Policy, and QoS Packet Scheduler. Snap-ins may in turn be extended. For example, the Security Settings snap-in includes several extension snap-ins. Developers can also create their own MMC extension snap-ins to Group Policy Object Editor to provide additional Group Policy settings. |
| Client-Side Extensions | Client-side extensions (CSEs) run within dynamic-link libraries (DLLs) and are responsible for implementing Group Policy at the client computer. The following CSEs are loaded, by default, in Windows Server 2003:<br><br>Administrative Templates, Wireless Network Policies, Folder Redirection, Disk Quotas, QoS Packet Scheduler, Scripts, Security, Internet Explorer Maintenance, EFS Recovery, Software Installation, and IP Security. |
| Group Policy Management Console (GPMC) | GPMC is a new tool designed to simplify implementation and management of Group Policy. It consists of a new MMC snap-in and a set of scriptable interfaces for managing Group Policy. The Group Policy Management Console provides:<br>• A user interface based on how customers use and manage Group Policy, rather than on how the technology is built.<br>• Import/Export, Copy/Paste, and searching of GPOs.<br>• Simplified management of Group Policy-related security.<br>• Reporting (printing, saving, read-only access to GPOs) for GPO and Resultant Set of Policy (RSoP) data.<br>• Backup/Restore of GPOs.<br>• Scripting of GPO operations that are exposed within this tool (but NOT scripting of settings within a GPO). |
| Resultant Set of Policy (RSoP) snap-in | The Resultant Set of Policy (RSoP) snap-in is an MMC snap-in that that simplifies Group Policy implementation and troubleshooting. RSoP uses Windows Management Instrumentation (WMI) to determine how Group Policy settings are applied to users and computers. For RSoP functionality, it is recommended to use the reporting features in GPMC. |
| Winlogon | A component of the Windows operating system that provides interactive logon support, Winlogon is the service in which the Group Policy engine runs. |
| Group Policy engine | The Group Policy engine is the framework that handles common functionalities across client-side extensions including scheduling of Group Policy application, obtaining GPOs from relevant configuration locations, and filtering and ordering of GPOs. |
| File System | The NTFS file system on client computers. |

| | |
|---|---|
| Registry | A database repository for information about a computer's configuration, the registry contains information that Windows continually references during operation, such as:<br>1. Profiles for each user.<br>2. The programs installed on the computer and the types of documents that each can create.<br>3. Property settings for folders and program icons.<br>4. The hardware on the system.<br>5. Which ports are being used.<br><br>The registry is organized hierarchically as a tree, and it is made up of keys and their subkeys, hives, and entries. The Group Policy engine has read and write access to the Registry.<br><br>Registry settings can be controlled via the Group Policy Administrative Templates extension. |
| Event Log | The Event log is a service, located in Event Viewer, which records events in the system, security, and application logs. The Group Policy engine has write access to the Event Log on client computers and domain controllers. The Help and Support Center on each computer has read access to the Event Log. |
| Help and Support Center | The Help and Support Center is a component on each computer that provides HTML reports on the Group Policy settings currently in effect on the computer. |
| Resultant Set of Policy (RSoP) infrastructure | All Group Policy processing information is collected and stored in a Common Information Model Object Management (CIMOM) database on the local computer. This information, such as the list, content and logging of processing details for each GPO, can then be accessed by tools using WMI.<br><br>In logging mode (Group Policy Results), RSoP queries the CIMOM database on the target computer, receives information about the policies and displays it in GPMC. In planning mode (Group Policy Modeling), RSoP simulates the application of policy using the Group Policy Directory Access Service (GPDAS) on a domain controller. GPDAS simulates the application of GPOs and passes them to virtual client-side extensions on the domain controller. The results of this simulation are stored to a local CIMOM database on the domain controller before the information is passed back and displayed in GPMC. |
| WMI | WMI is a management infrastructure that supports monitoring and controlling of system resources through a common set of interfaces and provides a logically organized, consistent model of Windows operation, configuration, and status.<br><br>WMI makes data about a target computer available for administrative use. Such data can include hardware and software inventory, settings, and configuration information. For example, WMI exposes hardware configuration data such as CPU, memory, disk space, and manufacturer, as well as software configuration data from the registry, drivers, file system, Active Directory, the Windows Installer service, networking configuration, and application data. WMI Filtering in Windows Server 2003 allows you to create queries based on this data. These queries (also called WMI filters) determine which users and computers receive all of the policy configured in the GPO where you create the filter. |

**Core Group Policy**
This subject explains Group Policy infrastructure including how the Group Policy engine controls policy processing, including retrieval of GPOs, invocation of individual extensions, and other infrastructure functionality.

**Group Policy Components**
The Group Policy Components subcollection describes the role of extensions including server-side snap-in extensions and client-side extensions. These extensions include: Administrative Templates, Software Installation, Security Settings, Scripts, Remote Installation Services, Internet Explorer Maintenance, Folder Redirection, QoS Packet Scheduler, Disk Quotas, and Wireless Network Policies.

**Group Policy Administrative Tools**
This subcollection explains administrative tools including the Group Policy Object Editor, Group Policy Management Console,

and the Resultant Set of Policy (RSoP) snap-in.

**Group Policy Scenarios**
Group Policy is used to define configurations for groups of users and computers. With Group Policy, you can specify specific configurations for a wide range of areas including Administrative Templates (registry-based policies), security, software installation, scripts, folder redirection, remote installation services, and Internet Explorer maintenance. Group Policy settings are contained in a GPO. By associating a GPO with selected Active Directory system containers—sites, domains, and organizational units—the GPO's Group Policy settings are applied to the users and computers in those Active Directory containers. This section provides an overview of what you can do with Group Policy.

**Managing Desktops, Applications, and Components with Registry-Based Policies**
Administrative Templates (or .adm files) enable you to control registry settings using Group Policy, providing the means to configure the behavior and appearance of the desktop, including the operating system, components, and applications. Windows comes with a predefined set of Administrative template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a GPO. These .adm files are stored in two locations by default: inside GPOs in the **Sysvol** folder and in the **%windir%\inf** directory on the local computer.

**Managing Security**
Group Policy is used to manage the following types of securityoptions for users, clients, servers, and domain controllers:

- **Security settings**. These Group Policy settings are used to define values for various security-relevant operating system parameters, such as password policy, user rights assignment, audit policy, registry values, file and registry ACLs, and service startup modes.

- **IPSec policies**. These Group Policy settings are used to configure IPSec services for authenticating or encrypting network traffic. An IPSec policy consists of a set of security rules, and each security rule consists of an IP filter with an action.

- **Software restriction policies**. These Group Policy settings are used to help protect computers from code that is not trusted by identifying and specifying which applications are permitted to run.

- **Wireless network policies**. These Group Policy settings are used to configure settings for the Wireless Configuration Service, a user-mode service that operates on each of the IEEE 802.11 wireless network adapters that are installed on a computer.

- **Public Key Policies**. These Group Policy settings are used to:

    - Specify that computers automatically submit a certificate request to an enterprise certification authority and install the issued certificate.

    - Create and distribute a certificate trust list.

    - Establish common trusted root certification authorities.

    - Add encrypted data recovery agents and change the encrypted data recovery policy settings.

**Implementing Group Policy–based Software Installation**
The Software Installation snap-in is used to centrally manage software. Software can be assigned or published to users and assigned to computers. Group Policy-based software installation can be used to install software applications when a computer is started, when the user logs on, or on demand. Software installation Group Policy settings can be applied to users or computers in an Active Directory structure.

Group Policy-based software installation can also be used to upgrade deployed applications or remove earlier applications that are no longer required. Users can be restricted from installing any software from local media, such as a CD-ROM, or disk, or other unapproved applications.

Medium and large organizations may wish to consider using Systems Management Server (SMS). SMS provides advanced capabilities such as inventory-based targeting, status reporting, server- and client-side scheduling, multisite facilities, complex targeting, centralized hardware and software inventory, remote diagnostic tools, software metering, software distribution-point population and maintenance, support for Windows 95, Windows 98, Windows NT 4.0, Windows 2000, and Windows XP clients, and enhanced software deployment features. SMS does not require Active Directory.

**Managing Remote Operating System Installations**
Remote Installation Services (RIS) is used to control the behavior of the Remote Operating System Installation feature as displayed to client computers. Remote Installation enables administrators to perform a new installation of Windows on Preboot eXecution Environment (PXE) remote boot-enabled client computers throughout an organization. Using a customized, fully automated installation process from a remote source, an administrator does not have to visit the new computer to install a new operating system and core applications.

**Managing with Scripts**
Scripts are used to automate tasks at computer startup and shutdown, and at user logon and logoff. Scripts can be written in any language supported by Windows Script Host including the Microsoft Visual Basic development system, Scripting Edition (VBScript), JavaScript, PERL, and MS DOS-style batch files (.bat and .cmd).

**Managing Internet Explorer**
Internet Explorer Maintenance is used to manage and customize Internet Explorer on computers running Windows 2000 or later. You can set options for the Browser UI, connections, URLs, proxy settings, security zones, Favorites, and Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening). .

**Managing Folder Redirection**
You can use folder redirection to redirect special directories on Windows 2000 or Windows Server 2003 from their default user profile location to an alternate location on the network. These special folders include **My Documents**, **Application Data**, **Desktop**, and the **Start** menu.

# Core Group Policy Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Core Group Policy Technical Reference**
Group Policy is an infrastructure that allows you to specify managed configurations for users and computers in an Active Directory directory service environment. Group Policy settings are contained in Group Policy objects (GPOs). GPOs exist in a domain, and can be linked to the following Active Directory containers: sites, domains, or organizational units (OUs). A GPO consists of two parts — one part is located in Active Directory, the other is located in the Sysvol folder — both of which are replicated to all domain controllers in the domain. The settings within GPOs are then evaluated by the affected targets, using the hierarchical nature of Active Directory. Core Group Policy or the Group Policy engine is the framework that handles common functionalities across all client side extensions.

**In this subject**


- What Is Core Group Policy?
- How Core Group Policy Works
- Core Group Policy Tools and Settings

# What Is Core Group Policy?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Change and Configuration Management
- Core Group Policy Infrastructure
- Core Group Policy Scenarios
- Core Group Policy Dependencies
- Related Information


Group Policy is an infrastructure used to deliver and apply one or more desired configurations or policy settings to a set of targeted users and computers within an Active Directory environment. This infrastructure consists of a Group Policy engine and multiple client-side extensions (CSEs) responsible for writing specific policy settings on target client computers.

Group Policy settings are contained in Group Policy objects (GPOs), which live in the domain and can be linked to the following Active Directory containers: sites, domains, or organizational units (OUs). The settings within GPOs are then evaluated by the affected targets, using the hierarchical nature of Active Directory. Consequently, Group Policy is one of the top reasons to deploy Active Directory.

Group Policy is one of a group of management technologies, collectively known as IntelliMirror, that provides users with consistent access to their applications, application settings, roaming user profiles, and user data, from any managed computer — even when they are disconnected from the network. IntelliMirror is implemented through a set of Windows features, including Active Directory, Group Policy, Group Policy-based Software Installation, Windows Installer, Folder Redirection, Offline Folders, and Roaming User Profiles.

Core Group Policy or the Group Policy engine is the framework that handles common functionalities across Administrative Template settings and other client-side extensions. The following figure shows how the Group Policy engine interacts with other components as part of processing policy settings. You use Group Policy Management Console (GPMC) to create, view, and manage GPOs and use Group Policy Object Editor to set and configure the policy settings in GPOs.

**Group Policy Components**




**Change and Configuration Management**

Administrators face increasingly complex challenges in managing their IT infrastructures. You must deliver and maintain customized desktop configurations for many types of workers, including mobile users, information workers, or others assigned

to strictly defined tasks, such as data entry. Changes to standard operating system images might be required on an ongoing basis. Security settings and updates must be delivered efficiently to all the computers and devices in the organization. New users need to be productive quickly without costly training. In the event of a computer failure or disaster, service must be restored with a minimum of data loss and interruption.

Specifically, an IT department must respond to various factors that require change in an IT environment including:

- New operating systems and applications.
- Updates to operating systems and applications.
- New hardware.
- New business requirements that require configuration changes.
- Security influences that require configuration changes.
- New users.

Managing this change can be viewed as a continuous cycle, in which new business requirements demand changes that must first be tested before they can be deployed as a standard configuration. This cycle is shown in the following figure.

**Change and Configuration Management Process**




These roles, known collectively as Change and Configuration Management, enable administrators to implement change quickly and affect large numbers of users and computers at the lowest possible cost. You can use Group Policy to maintain standard operating environments for specific groups of users, such as developers or information workers. As software changes and policies change over time, Group Policy can be used to update the already-deployed standard operating environment until the image can be updated. Group Policy can also enforce rules, if necessary, by restricting the programs that can be run on company computers. For example, it can prevent access to games or other programs unrelated to the workplace.

Group Policy is a key enabling technology that allows you to implement Change and Configuration Management along with other technologies in IntelliMirror. For example, you can deploy new operating systems with Remote Installation Services or other imaging technology. You can deliver updates to computers throughout the network using Software Update Services (SUS). Although you can deploy software using Group Policy, larger organizations might want to use Microsoft Systems Management Server (SMS) to take advantage of the scalability that SMS provides.

In summary, Group Policy is the delivery mechanism that allows you to implement change and configuration for users and computers on the object level in Active Directory. Because you can target Group Policy settings to individual objects throughout the Active Directory hierarchy, Group Policy is the central enabling technology that allows organizations to effectively use Active Directory as a management tool. In addition, the Group Policy Management Console simplifies implementation and management of Group Policy.

**Core Group Policy Infrastructure**

Group policy is an infrastructure with pluggable extensions. Extensions that exist on client computers include Administrative Templates (also known as registry-based policy), Security Settings, Software Installation, Folder Redirection, Scripts, and Wireless Network Policies.

The policy settings exist in a GPO. A GPO is a virtual object that lives in the domain; part of the GPO is located in Active Directory and is called the Group Policy container. The other part of a GPO is located in the Sysvol and is called the Group Policy template. When policy settings need to be applied, the framework calls each extension and the extension then applies the necessary settings.

Each Group Policy extension has two extensions — a client extension that is called by the Group Policy engine to apply policy, and a server-side extension that plugs into Group Policy Object Editor to define and set the policy settings that need to be applied to client computers.

Although you can configure Local Group Policy objects (Local GPOs) on individual computers, the full power of Group Policy can only be realized in a Windows 2000 or Windows Server 2003-based network with Active Directory installed. In addition, some features and policy settings require client computers running Windows XP.

**Group Policy and Active Directory**
Active Directory organizes objects by sites, domains, and OUs. Domains and OUs are organized hierarchically, making the containers and the objects within them easy to manage. The settings defined in a GPO can only be applied when the GPO is linked to one or more of these containers.

By linking GPOs to sites, domains, and OUs, you can implement Group Policy settings for as broad or as narrow a portion of the organization as you want. GPO links affect users and computers in the following ways:

- A GPO linked to a site applies to all users and computers in the site.
- A GPO linked to a domain applies directly to all users and computers in the domain and by inheritance to all users and computers in child OUs. Note that policy is not inherited across domains.
- A GPO linked to an OU applies directly to all users and computers in the OU and by inheritance to all users and computers in child OUs.

When a GPO is created, it is stored in the domain. When the GPO is linked to an Active Directory container, such as an OU, the link is a component of that container, not a component of the GPO.

An example of how GPOs can be linked to sites, domains, and OUs is shown in the following figure.

**Sample Active Directory Organizational Structure**




In this configuration, the Servers OUs have the following GPOs applied: A1, A2, A3, A4, A6. The Marketing OUs have the following GPOs applied: A1, A2, A3, A5.

**Loopback processing with merge or replace**
Loopback is an advanced Group Policy setting that is useful on computers in certain closely managed environments, such as servers, kiosks, laboratories, classrooms, and reception areas. Setting loopback causes the User Configuration settings in GPOs that apply to the computer to be applied to every user logging on to that computer, instead of, or in addition to, the User Configuration settings of the user. This allows you to ensure that a consistent set of policies is applied to any user logging on to a particular computer, regardless of their location in Active Directory. Loopback is controlled by the setting, User Group Policy loopback processing mode, which is located in **Computer Configuration\Administrative Templates\System\Group Policy**. Loopback only works when both the user account and the computer account are in a Windows 2000 or later domain. Loopback does not work for computers joined to a workgroup. Loopback is not enabled if the computer or user is not in an Active Directory domain.

**Filtering the Scope of the Group Policy Object**
Group Policy is a powerful tool for managing the Windows Server 2003 environment. The value of Group Policy can only be realized through properly applying the GPOs to the Active Directory containers you want to manage. Determining which users and computers will receive the settings in a GPO is referred to as "scoping the GPO." Scoping a GPO is based on three factors:

- The site(s), domain(s), or organization unit(s) where the GPO is linked.
- The security filtering on the GPO.
- The WMI filter on the GPO.

You can use the site, domain, and OU links from a GPO as the primary targeting principle for defining which computers and users should receive a GPO. You can then use security filtering and WMI filtering to further reduce the set of computers and users to which the GPO will apply.

Scoping or targeting of a GPO allows you to apply or deny an entire GPO; you cannot choose to filter settings within a GPO.

**Group Policy Inheritance**
In addition to the ability to filter the scope of GPOs, you can change the way GPOs are applied by managing Group Policy inheritance. In most environments, the actual settings for a given user and computer are the result of the combination of GPOs that are applied at a site, domain, or OU. When multiple GPOs apply to these users and computers, the settings in the GPOs are aggregated. The settings deployed by GPOs linked to higher containers (parent containers) in Active Directory are inherited by default to child containers and combine with any settings deployed in GPOs linked to child containers. If multiple GPOs attempt to set a setting to conflicting values, the GPO with the highest precedence sets the setting. GPO processing is based on a "last writer wins" model, and GPOs that are processed later have precedence over GPOs that are processed earlier.

**Viewing and Reporting of Policy Settings**
In order to properly implement, troubleshoot, and plan Group Policy, administrators need to be able to quickly view the settings in a GPO. When multiple GPOs apply to a given user or computer, they can contain conflicting policy settings. For most policy settings, the final value of the policy setting is set only by the highest precedent GPO that contains that setting. Resultant Set of Policy (RSoP) helps you understand and identify the final set of policy that is applied as well as settings that did not apply as a result of policy inheritance.

Specifically, Resultant Set of Policy helps you determine:

- The final value of the setting that is applied as a result of all the GPOs.
- The final GPO that set the value of this setting (also known as the winning GPO).
- Precedence details that show any other GPOs that attempted to set this setting and the value that each GPO attempted to set for that policy setting.

Group Policy Management Console, available as a separate download from the Microsoft Web site, addresses some common reporting requirements including the ability to document all the settings in a GPO to a file for printing or viewing. Users can either print the reports, or save them to a file as either HTML or XML.

**Delegating Administration of Group Policy**
Organizations need to be able to delegate administration of Group Policy to other administrators who can take responsibility for a given OU, domain, or other container. Active Directory is designed to allow you to delegate control of portions of the directory service in managing aspects of Group Policy. The following areas can be delegated:

- **GPO delegation.** This includes permission to create GPOs in a domain or permission to edit an existing GPO. Note that having permission to edit a GPO does not include any delegated rights on the GPO links.
- **Link delegation.** This includes permission to add, delete, or change links to GPOs. Note that having link delegation does not include any delegated rights on the GPO itself.
- **RSOP delegation.** This includes permission to run RSoP (in either planning or logging mode) on objects under a container.
- **WMI filter delegation.** This includes permission to create WMI filters or permission to edit an existing filter.

In GPMC, delegation is simplified because it manages the various Access Control Entries (ACEs) required for a task as a single bundle of permissions for the task. You can also use the Access Control List (ACL) editor to view or manage these permissions manually.

The underlying mechanism for achieving delegation is the application of the appropriate DACLs to GPOs and other objects in Active Directory. This mechanism is identical to using security groups to filter the application of GPOs to various users. You can also specify Group Policy to control who can use MMC snap-ins. For example, you can use Group Policy to manage the rights to create, configure, and use MMC consoles, and to control access to individual snap-ins.

**Core Group Policy Scenarios**

The Group Policy engine is designed to apply policy configurations to individual computers and users through Group Policy objects. Settings within a GPO are configured by individual client-side extensions and are applied to individual computers and users by the client-side extension. This section summarizes the core scenarios for the Group Policy engine.

- **Scheduling of Group Policy application.** Group Policy starts each time the computer starts or a user logs in, a process called foreground policy application. Group Policy is also applied in the background at regular refresh intervals. In addition, it can be forced to apply through command line tools such as Gpupdate.

- **Obtaining GPOs from the relevant configuration locations.** The Group Policy engine obtains GPOs from the appropriate site, domain, and OU containers, known collectively as Scope of Management (SOM).

- **Handling special cases affecting all CSEs.** The Group Policy engine implements additional changes specified by the administrator such as changing the link order in which GPOs should be applied. In addition, the Group Policy engine handles any loopback processing that has been set by an administrator. Loopback processing, typically used for public workstations or kiosks, specifies that the user settings defined in the computer's GPOs replace or are merged with the user settings normally applied to the user.

- **Filtering and ordering of GPOs.** The Group Policy engine checks for any conditions set by the administrator to filter GPOs or specify the order in which GPOs should be applied.

- **Configuring CSEs.** CSEs can be configured to run only in some conditions, as specified in the registry.

- **Maintaining version numbers and histories for all CSEs.** A history list is maintained for each CSE in the Registry, showing when the CSE last applied policy. A status is also maintained to figure out whether the client-side extension applied policy successfully last time the policy was applied by the CSE.

- **Calling into the CSE.** When the Group Policy engine has determined that a particular CSE needs to be executed, it loads the dynamic link library (DLL) associated with the CSE and loads up the entry point.

- **Notifying various components of any changes made by Group Policy.** After all the extensions have been called, the Group Policy engine updates the registry information to specify what the next policy foreground application should be, as well as schedule the next background refresh time.

- **Processing RSoP data.** The Group Policy engine periodically refreshes the RSoP data, ensuring that actual policy application is updated and stored in the WMI namespace on each computer.

**Group Policy Dependencies**

Group Policy has several key dependencies. Domain-based Group Policy requires an Active Directory environment with DNS properly configured.

**Active Directory**

Active Directory is the Windows 2000 Server and Windows Server 2003 directory service that stores information about all objects on the computer network and makes this information easy for administrators and users to find and apply. With Active Directory, users can gain access to resources anywhere on the network with a single logon. Similarly, administrators have a single point of administration for all objects on the network, which can be viewed in a hierarchical structure. In a network environment, Group Policy depends on Active Directory as the targeting framework that allows you to link GPOs to specific Active Directory containers such as sites, domains, or OUs.

In a stand-alone environment without Active Directory, you can use Local Group Policy objects to configure settings on individual computers.

**Domain Name System (DNS)**

DNS is a hierarchical, distributed database that contains mappings of DNS domain names to various types of data, such as IP addresses. DNS enables the location of computers and services by user-friendly names, and it also enables the discovery of other information stored in the database.

Group Policy application requires clients to access specified servers, including domain controllers and other servers such as share points and install points. Group Policy management also requires access to domain controllers. DNS is used to locate and identify these servers. In Windows 2000 Server and later Active Directory requires DNS support. If the network is

functioning, but clients or consoles such as the Group Policy Object Editor or GPMC are unable to locate the servers, there might be a problem with your network's DNS system.

**Replication**

Group Policy depends on other technologies in order to properly replicate between domain controllers in a network environment. A GPO is a virtual object stored in both Active Directory and the Sysvol of a domain controller. Property settings, stored in the Group Policy container, are replicated through Active Directory replication. Replication automatically copies the changes that originate on a writable directory partition replica to all other domain controllers that hold the same directory partition replica. More specifically, a destination domain controller pulls these changes from the source domain controller. Data settings, stored in the Sysvol as the Group Policy template, are replicated through the File Replication Service (FRS), which provides multi-master file replication for designated directory trees between designated servers running Windows Server 2003. The Group Policy container stores GPO properties, including information about version, GPO status, and a list of components that have settings in the GPO. The Group Policy template is a directory structure within the file system that stores Administrative Template-based policy settings, security settings, script files, and information regarding applications that are available for software installation. The Group Policy template is located in Sysvol in the \Policies sub-directory for its domain. GPOs are identified by their globally unique identifiers (GUIDs) and stored at the domain level. The settings from a GPO are only applied when the Group Policy container and Group Policy template are synchronized.

**DFS publishing**

The Sysvol folder is shared on each domain controller and is accessible through the UNC path \\dcname.domainname\sysvol.

The Sysvol is also published as a domain-based Distributed File System (DFS) share. This allows clients to access the Sysvol by using the generic path *domainname*\sysvol. A request for a DFS referral for *domainname*\sysvol will always return a replica in the same Active Directory site as the client if one is available. This is the mechanism that the Group Policy client-side extensions use to retrieve a local copy of the Group Policy template information.

**Related Information**

The following contains additional information that is relevant to this section.

Group Policy Settings Reference for Windows Server 2003

# How Core Group Policy Works

Updated: September 2, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Core Group Policy Works**

**In this section**


- Core Group Policy Architecture
- Core Group Policy Physical Structure
- Core Group Policy Processes and Interactions
- Network Ports Used by Group Policy
- Related Information


Core Group Policy or the Group Policy engine is the infrastructure that processes Group Policy components including server-side snap-in extensions and client-side extensions. You use administrative tools such as Group Policy Object Editor and Group Policy Management Console to configure and manage policy settings.

At a minimum, Group Policy requires Windows 2000 Server with Active Directory installed and Windows 2000 clients. Fully implementing Group Policy to take advantage of all available functionality and the latest policy settings depends on a number of factors including:

- Windows Server 2003 with Active Directory installed and with DNS properly configured.
- Windows XP client computers.
- Group Policy Management Console (GPMC) for administration.

**Core Group Policy Architecture**

The Group Policy engine is a framework that handles client-side extension (CSE) processing and interacts with other elements of Group Policy, as shown in the following figure:

**Core Group Policy Architecture**




The following table describes the components that interact with the Group Policy engine.

**Core Group Policy Components**

| Component | Description |
| --- | --- |
| Server (domain controller) | In an Active Directory forest, the domain controller is a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. |
| Active Directory | Active Directory, the Windows-based directory service, stores information about objects on a network and makes this information available to users and network administrators. Administrators link Group Policy objects (GPOs) to Active Directory containers such as sites, domain, and organizational units (OUs) that include user and computer objects. In this way, policy settings can be targeted to users and computers throughout the organization. |
| Sysvol | The Sysvol is a set of folders containing important domain information that is stored in the file system rather than in the directory. The Sysvol folder is, by default, stored in a subfolder of systemroot folder (%\systemroot\sysvol\sysvol) and is automatically created when a server is promoted to a domain controller. The Sysvol contains the largest part of a GPO: the Group Policy template, which includes Administrative Template-based policy settings, security settings, script files, and information regarding applications that are available for software installation. It is replicated through the File Replication Service (FRS) between all domain controllers in a domain. |
| Group Policy object (GPO) | A GPO is a collection of Group Policy settings, stored at the domain level as a virtual object consisting of a Group Policy container and a Group Policy template. The Group Policy container, which contains information about the properties of a GPO, is stored in Active Directory on each domain controller in the domain. The Group Policy template contains the data in a GPO and is stored in the Sysvol in the /Policies subdirectory. GPOs affect users and computers that are contained in sites, domains, and OUs. |
| Local Group Policy object | The Local Group Policy object (Local GPO) is stored on each individual computer, in the hidden **Windows\System32\GroupPolicy** directory. Each computer running Windows 2000, Windows XP Professional, Windows XP 64-Bit Edition, Windows XP Media Center Edition, or Windows Server 2003 has exactly one Local GPO, regardless of whether the computers are part of an Active Directory environment.<br><br>Local GPOs are always processed, but are the least influential GPOs in an Active Directory environment, because Active Directory-based GPOs have precedence. |
| Winlogon | A component of the Windows operating system that provides interactive logon support, Winlogon is the service in which the Group Policy engine runs. |
| Group Policy engine | The Group Policy engine is the framework that handles common functionalities across registry-based settings and client-side extensions (CSEs). |
| Client-side extensions | CSEs run within dynamic-link libraries (DLLs) and are responsible for implementing Group Policy at the client computer.<br><br>The CSEs are loaded on an as-needed basis when a client computer is processing policy |
| File system | The NTFS file system on client computers. |

| | |
|---|---|
| Registry | A database repository for information about a computer's configuration, the registry contains information that Windows continually references during operation, such as:<br>• Profiles for each user.<br>• The programs installed on the computer and the types of documents that each can create.<br>• Property settings for folders and program icons.<br>• The hardware on the system.<br>• Which ports are being used.<br><br>The registry is organized hierarchically as a tree, and it is made up of keys and their subkeys, hives, and entries.<br><br>Registry settings can be controlled through Group Policy, specifically, Administrative Templates (.adm files). Windows Server 2003 comes with a predefined set of Administrative Template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a GPO. These .adm files are stored in two locations by default: inside GPOs in the Sysvol folder and in the **Windows\inf** directory on the local computer. |
| Event log | The Event log is a service, located in Event Viewer, which records events in the system, security, and application logs. |
| Help and Support Center | The Help and Support Center is a component on each computer that provides HTML reports on the policy settings currently in effect on the computer. |
| Resultant Set of Policy (RSoP) infrastructure | All Group Policy processing information is collected and stored in a Common Information Model Object Management (CIMOM) database on the local computer. This information, such as the list, content, and logging of processing details for each GPO, can then be accessed by tools using Windows Management Instrumentation (WMI). |
| WMI | WMI is a management infrastructure that supports monitoring and controlling of system resources through a common set of interfaces and provides a logically organized, consistent model of Windows operation, configuration, and status.<br><br>WMI makes data about a target computer available for administrative use. Such data can include hardware and software inventory, settings, and configuration information. For example, WMI exposes hardware configuration data such as CPU, memory, disk space, and manufacturer, as well as software configuration data from the registry, drivers, file system, Active Directory, the Windows Installer service, networking configuration, and application data. WMI filtering in Windows Server 2003 allows you to create queries based on this data. These queries (WMI filters) determine which users and computers receive all of the policy configured in the GPO where you create the filter. |

**Group Policy Engine Architecture**

The primary purpose of Group Policy is to apply policy settings to computers and users in an Active Directory domain. GPOs can be targeted through Active Directory containers, such as sites, domains, and OUs, containing user or computer objects. The Group Policy engine is in userenv.dll, which runs inside the Winlogon service. This is shown in the following figure.

**Group Policy Engine Architecture and CSE Components**




**Group Policy Engine Architecture and CSE Components**

| Component | Description |
|---|---|
| Group Policy Engine | The framework that handles functionalities across CSEs, the Group Policy engine runs inside userenv.dll. |
| Winlogon.exe | A component of the Windows operating system that provides interactive logon support, Winlogon is the service in which the Group Policy engine runs. Winlogon is the only system component that actively interacts with the Group Policy engine. |
| Userenv.dll | Userenv.dll runs inside Winlogon and contains the Group Policy engine and the Administrative Templates extension. |
| gptext.dll | Used to configure Scripts, IP Security, QoS Packet Scheduler, and Wireless settings. |
| fdeploy.dll | Used to configure folder redirection |
| scecli.dll | Used to configure security settings. |
| iedkcs32.dll | Used to manage various Internet Explorer settings. |
| appmgmts.dll | Used to configure software installation settings. |
| dskquota.dll | Used for setting disk quotas. |

**RSoP Architecture**

Resultant Set of Policy (RSoP) uses WMI to determine how policy settings are applied to users and computers. RSoP has two modes: logging mode and planning mode. Logging mode determines the resultant effect of policy settings that have been applied to an existing user and computer based on a site, domain, and OU. Logging mode is available on Windows XP and later operating systems. Planning mode simulates the resultant effect of policy settings that are applied to a user and computer. Planning mode requires a Windows Server 2003 computer as a domain controller. For RSoP functionality, using GPMC is recommended, which includes RSoP features integrated with the rest of GPMC. In GPMC, RSoP logging mode is referred to as Group Policy Results; planning mode is referred to as Group Policy Modeling.

The following figure shows the high-level architecture of RSoP for Group Policy Results and Group Policy Modeling:

**RSoP Architecture**




Windows Server 2003 collects Group Policy processing information and stores it in a WMI database on the local computer. (The WMI database is also known as the CIMOM database.)This information, such as the list, content and logging of processing details for each GPO, can then be accessed by tools using WMI.

In Group Policy Results, RSoP queries the WMI database on the target computer, receives information about the policies and displays it in GPMC. In Group Policy Modeling, RSoP simulates the application of policy using the Resultant Set of Policy Provider on a domain controller. Resultant Set of Policy Provider simulates the application of GPOs and passes them to virtual CSEs on the domain controller. The results of this simulation are stored to a local WMI database on the domain controller before the information is passed back and displayed in GPMC (or the RSoP snap-in). This is explained in greater detail in the following section.

**Windows Management instrumentation (WMI) and Common Information Model Object Management (CIMOM)**
WMI provides a common scriptable interface to retrieve, and in some cases set, a wide variety of system and application information. WMI is implemented through the winmgmt.exe service. The WMI information hierarchy is modeled as a hierarchy of objects following the Common Information Model (CIM) standards.

This information hierarchy is extensible, which allows different applications and services to expose configuration information by supplying a WMI provider. WMI providers are the interface between the WMI service and the application's data in its native format.

WMI data can be dynamic (generated on demand when required by a management application) or static. Static data is stored in the CIMOM database. This data can be accessed at any time (security controls permitting) by management applications. RSoP uses WMI and the CIMOM to write, store and query RSoP settings information.

**Resultant Set of Policy Provider**
RSoP in planning mode has the special requirement that no settings information is actually applied to the client system during the RSoP data generation. In fact, in many planning scenarios, there might not be a computer or user object to apply the settings to. To meet this requirement, the RSoP provider runs on domain controllers and performs some of the functions of a client system for GPO application.

The RSoP provider is actually a WMI provider that performs the role of Winlogon in invoking CSEs to log RSoP information to the CIM repository. It takes parameters supplied by the RSoP wizard to select GPOs from the directory. It uses the following parameters:

- **Scope of management (SOM).** This is the combination of user object OU and computer object OU and Site (although the latter is optional). Either the User or Computer SOM can be omitted but one must be specified. This can be specified as either an existing computer, user, or both, or as existing OUs for the computer, user, or both.

- **Security group memberships for the computer and user objects.** By default, these are the existing security group memberships if actual user or computer objects are chosen as the SOM. Security filtering can be ignored entirely or a new or a modified set of groups can be chosen.

- **WMI filter.** New or modified WMI filters can be applied to the GPOs during the RSoP generation.

RSoP provider is a service on a domain controller and runs in system context. There are two ramifications to this design. The service manually evaluates the security descriptor of each GPO in the SOM against the user object and computer object security identifiers (SIDs) and their security group membership. If Active Directory is locked down, some security group membership analyses might fail until the user is provided the correct access.

In addition, because the RSoP provider has Domain Admin-equivalent access rights, some control needs to be placed on who can generate RSoP information in the directory. This control is achieved by an extended access right — GenerateRSoPData. To execute an RSoP session for a particular container, the user must have the Generate RSoP (planning) access right for that container.

**Planning Mode (Group Policy Modeling)**
In planning mode, RSoP provider performs the RSoP data generation in the following steps:

1. RSoP tool gets the user, computer and domain controller name from the wizard.

2. RSoP connects to the WMI database on the specified domain controller.

3. WMI service in turn calls the WMI Provider (an out-of-process service provider) on the same computer to create the RSoP data.

4. RSoP provider gets the list GPOs for the user and computer from Active Directory.

5. RSoP provider populates the WMI database with instances of the GPOs for user and computer.

6. The list of registered CSEs is retrieved. Each of the policy extensions is dynamically loaded in succession by the RSoP provider and the list of computer and user GPOs is passed to each extension. Each policy extension takes the list of GPOs and instead of applying the policy, it populates the WMI database with instances of policy objects that describe the effective policy.

7. After the WMI database population is over, RSoP provider returns the namespace under which the RSoP data was created to the WMI service. This returns the namespace to the RSoP tool.

8. The RSoP tool connects to the namespace on the WMI database on the domain controller. RSoP navigates or iterates through the populated data in the WMI database using WMI enumeration APIs to retrieve the policy data.

9. The RSoP data is displayed to the user. When the user is done looking at the RSoP data, the RSoP tool calls RsopDeleteSession on the WMI service to delete the data in WMI's database that was previously created.

**Logging Mode (Group Policy Results)**
In logging mode, the RSoP data generation is controlled by Winlogon and is part of the normal GPO processing operation.

The Winlogon process retrieves the list of GPOs from the Active Directory using the security context of the user or computer.

- Winlogon process populates the WMI database with instances of the GPOs.

- The list of registered policy extensions is retrieved. Each of the policy extensions is dynamically loaded in succession by the Winlogon process and the list of GPOs is passed to each policy extension. Each extension takes the list of GPOs and, in addition to applying the policy, it populates the WMI database with the policies set and which GPOs applied them.

**Core Group Policy Physical Structure**
Understanding where GPOs are stored and how they are structured can help you troubleshoot problems you might encounter when you implement Group Policy. Although GPOs can be linked to sites, domains, and OUs, they are stored only in the domain. As explained earlier, a GPO is a virtual object that stores its data in two locations: a Group Policy container and a Group Policy template.

**Group Policy Container**
A Group Policy container is a location in Active Directory that stores GPOs and their properties. The properties of a GPO include both computer and user Group Policy information. The Policies container is the default location of GPOs. The path to the Policies container, in Lightweight Directory Access Protocol (LDAP) syntax, is **CN=Policies,CN=System,DC=Domain_Name**, where the Domain_Name values specify a fully qualified domain name (FQDN).

The Active Directory store contains the Group Policy container of each GPO in the domain. The Group Policy container contains attributes that are used to deploy GPOs to the domain, to OUs, and to sites within the domain. The Group Policy container also

contains a link to the file system component of a GPO — the Group Policy template. Some of the information in a Group Policy container includes:

- **Version information.** Ensures that the information is synchronized with the Group Policy template information.
- **Status information.** Indicates whether the user or computer portion of the GPO is enabled or disabled.
- **List of components.** Lists (extensions) that have settings in the GPO. These attributes are **gPCMachineExtensionNames** and **gPCUserExtensionNames**.
- **File system path.** Specifies the Universal Naming Convention (UNC) path to the Sysvol folder. This attribute is **gPCFileSysPath**.
- **Functionality version.** Gives the version of the tool that created the GPO. Currently, this is version 2. This attribute is **gPCFunctionalityVersion**.
- **WMI filter.** Contains the distinguished name of the WMI filter. This attribute is **gPCWQLFilter**.

**System Container**

Each Windows Server 2003 domain contains a System container. The System container stores per-domain configuration settings, including GPO property settings, Group Policy container settings, IP Security settings, and WMI policy settings. IP Security and WMI policy are deployed to client computers through the GPO infrastructure.

The following subcontainers of the System container hold GPO-related settings:

- **Policies.** This object contains groupPolicyContainer objects listed by their unique name. Each groupPolicyContainer object holds subcontainers for selected computer and user policy settings.
- **Domain, OUs and Sites.** These objects contain two GPO property settings, **gPLink** and **gPOptions**.
- **Default Domain Policy.** This object contains the **AppCategories** container, which is part of the Group Policy Software installation extension.
- **IP Security.** This object contains IP Security policy settings that are linked to a GPO. The linked IP Security policy is applied to the recipients (user or computer) of the GPO.
- **WMIPolicy.** This object contains WMI filters that can be applied to GPOs. WMI filters contain one or more Windows Query Language (WQL) statements.

**System\Policies Container**

The System container is a top level container found in each domain naming context. It is normally hidden from view in the Active Directory Users and Computers snap-in but can be made visible by selecting “Advanced Features” from the snap-in View menu inside MMC. (Objects appear hidden in the Active Directory Users and Computers snap-in when they have the property showInAdvancedViewOnly = TRUE.) Group Policy information is stored in the Policies subcontainer of this container. Each GPO is identified by a GroupPolicyContainer object stored within the Policies container.

The Group Policy container is located in the **Domain_Name/System/Policies** container. Each Group Policy container is given a common name (CN) and this name is also assigned as the container name. For example, the name attribute of a Group Policy container, might be: {923B9E2F-9757-4DCF-B88A-1136720B5AF2}, which is also assigned to the Group Policy container's CN attribute.

The default GPOs are assigned the same Group Policy container CN on all domains. All other GPOs are assigned a unique CN. The default GPOs and their Group Policy container common names are:

- Default Domain Policy: {31B2F340-016D-11D2-945F-00C04FB984F9}.
- Default Domain Controllers Policy: {6AC1786C-016F-11D2-945F-00C04fB984F9}.

Knowing the common names of the default GPOs will help you distinguish them from non-default GPOs.

**How a Group Policy Container is Named**

Group Policy containers are named automatically when they are created. The CN of each Group Policy container is a GUID (Globally Unique Identifier). This is distinct from and unrelated to the Object GUID given to each Active Directory object. The CN is the name of the Group Policy container used to ensure uniqueness of Group Policy container names within the Policies container. There is no requirement for these GUIDs to be unique between domains (the Default Domain Policy and the Default Domain Controllers Policy GPOs each have identical GUIDs in all Active Directory installations). However, an Object GUID is always unique across all installations of the Active Directory store.

The following table shows permissions on Group Policy container:

**Default Group Policy Container Permissions**

| Trustee | Access |
|---|---|
| Authenticated Users | Read, Apply Group Policy |
| Domain Admins | Read, Write |
| System | Read, Write |

**GPO Attributes in the Policies CN**

GPOs are created by instantiating the groupPolicyContainer class in the Active Directory schema and storing the resulting GPO in the System/Policies container of the Active Directory store. After creating a GPO, you can review its CN from the Active Directory Users and Computers snap-in by enabling the Advanced view and then expanding the Policies CN. You can review all GPO attributes and their values from the Active Directory Services Interface Editor snap-in, ADSI Edit.

Object attributes are either mandatory or optional, as defined in the Active Directory schema. The CN attribute is mandatory for the class **Container**, Group Policy container's parent class. Three attributes — instanceType, objectCategory, and objectClass — are mandatory for the class **Top**, CN's parent class. Thus the Group Policy container class inherits all four mandatory attributes. The following table describes the mandatory attributes:

**Mandatory Attributes of the groupPolicyContainer Class**

| Name | Description |
|---|---|
| CN | The common name of the GPO. This is in the form of a GUID to avoid GPO naming conflicts within the Policies container. |
| instanceType | An attribute that dictates how an object is instantiated from its class on a particular server. In this case, it describes how the groupPolicyContainer class is created into a GPO in the Active Directory. A GPO is assigned the instanceType value of 4. |
| objectCategory | An object class name, including the object's path, used to group object's of the instantiated class. For example, the objectCategory of a GPO in the contoso.com domain is: CN=Group-Policy-Container,CN=Schema,CN=Configuration,DC=contoso,DC=com. |
| objectClass | The list of classes from which this class is derived. For a GPO, the objectClass is Container, groupPolicyContainer, and top. |

There are also a number of optional attributes inherited from the top class, and others that are assigned directly to the Group Policy container. Many optional attributes are required in order for the Group Policy container to function properly. For example, the GPCFileSysPath optional attribute must be present or the Group Policy container will not be linked to its corresponding Group Policy template.

**GroupPolicyContainer Subcontainers**

Within the GroupPolicyContainer there are a series of subcontainers. The first level of subcontainers — User and Machine — belong to the class Container. These two containers are used to separate some User-specific and Computer-specific Group Policy components.

**Group Policy Container-Related Attributes of Domain, Site, and OU Containers**

Windows Server 2003 uses domain, DNS, site, and organizational unit classes to create domain, site, and OU container objects respectively. These objects contain two optional Group Policy container-related attributes, **gPLink** and **gPOptions**. The **gPLink**

property contains the prioritized list of GPOs and the **gPOptions** property contains the Block Policy Inheritance setting.

The **gPLink** attribute holds a list of all Group Policy containers linked to the container and a number for each listed Group Policy container, that represents the **Enforced** (previously known as No Override) and **Disabled** option settings. The list appears in priority order from lowest to highest priority GPO.

The **gPOptions** attribute holds an integer value that indicates whether the Block Policy Inheritance option of a domain or OU is enabled (1) or disabled (0).

**Managing Group Policy Links for a Site, Domain, or OU**
To manage GPO links to a site, domain, or OU, you must have read and write access to the **gPLink** and **gPOptions** properties. By default, Domain Admins have this permission for domains and organizational unit, and only Enterprise Admins and Domain Admins of the forest root domain can manage links to sites. Active Directory supports security settings on a per-property basis. This means that a non-administrator can be delegated read and write access to specific properties. In this case, if non-administrators have read and write access to the gPLink and gPOptions properties, they can manage the list of GPOs linked to that site, domain, or OU.

**How WMIPolicy Objects are Stored and Associated with Group Policy Container Objects**
A single WMI filter can be assigned to a Group Policy container. The Group Policy container stores the distinguished name of the filter in **gPCWQLFilter** attribute. The Group Policy container locates the assigned filter in the System/WMIPolicy/SOM container. Each Windows Server 2003 domain stores its WMI filters in this Active Directory container. Each WMI filter stored in the SOM container lists the rules that define the WMI filter. Each rule is listed separately. For example, consider a WMI filter containing the following three WQL queries:

```
SELECT * FROM Win32_Product WHERE IdentifyingNumber = "{5E076CF2-EFED-43A2-A623-
13E0D62EC7E0}"
SELECT * FROM Win32_Product WHERE IdentifyingNumber = "{242365CD-80F2-11D2-989A-
00C04F7978A9}"
SELECT * FROM Win32_Product WHERE IdentifyingNumber = "{00000409-78E1-11D2-B60F-
006097C998E7}"
```

Three WMI rules are defined in the details of the filter. Each rule contains a number of attributes, including the query language (WQL) and the WMI namespace queried by the rule.

**Group Policy Template**
The majority of Group Policy settings are stored in the file system of the domain controllers. This part of each GPO is known as the Group Policy template. The GroupPolicyContainer object for each GPO has a property, GPCFileSysPath, which contains the UNC path to its related Group Policy template.

All Group Policy templates in a domain are stored in the ***domain_name*\Sysvol*domain_name*\Policies** folder, where domain_name is the FQDN of the domain. The Group Policy template for the most part stores the actual data for the policy extensions, for example Security Settings inf file, Administrative Template-based policy settings .adm and .pol files, applications available for the Group Policy Software installation extension, and potentially scripts.

**The Gpt.ini File**
The Gpt.ini file is located at the root of each Group Policy template. Each Gpt.ini file contains GPO version information. Except for the Gpt.ini files created for the default GPOs, a display name value is also written to the file.

Each Gpt.ini file contains the GPO version number of the Group Policy template.

```
[General]
Version=65539
```

Normally, this is identical to the version-number property of the corresponding GroupPolicyContainer object. It is encoded in the same way — as a decimal representation of a 4 byte hexadecimal number, the upper two bytes of which contain the GPO user settings version and the lower two bytes contain the computer settings version. In this example the version is equal to 10003 hexadecimal giving a user settings version of 1 and a computer settings version of 3.

Storing this version number in the Gpt.ini allows the CSEs to check if the client is out of date to the last processing of policy settings or if the currently applied policy settings (cached policies) are up-to-date. If the cached version is different from the version in the Group Policy template or Group Policy container, then policy settings will be reprocessed.

**Group Policy Template Subfolders**
The Group Policy template folder contains the following subfolders:

- **Machine.** Includes a Registry.pol file that contains the registry settings to be applied to computers. When a computer initializes, this Registry.pol file is downloaded and applied to the **HKEY_LOCAL_MACHINE** portion of the registry. The **Machine** folder can contain the following subfolders (depending on the contents of the GPO):
    - **Scripts\Startup.** Contains the scripts that are to run when the computer starts up.
    - **Scripts\Shutdown.** Contains the scripts that are to run when the computer shuts down.
    - **Applications.** Contains the advertisement files (.aas files) used by the Windows installer. These are applied to computers.
    - **Microsoft\Windows NT\Secedit.** Contains the Gpttmpl.inf file, which includes the default security configuration settings for a Windows Server 2003 domain controller.
    - **Adm.** Contains all of the .adm files for the GPO.
- User. Includes a Registry.pol file that contains the registry settings to be applied to users. When a user logs on to a computer, this Registry.pol file is downloaded and applied to the HKEY_CURRENT_USER portion of the registry. The User folder can contain the following subfolders (depending on the contents of the GPO):
    - **Applications.** Contains the advertisement files (.aas files) used by the Windows installer. These are applied to users.
    - **Documents and Settings.** Contains the Fdeploy.ini file, which includes status information about the Folder Redirection options for the current user's special folders.
    - **Microsoft\RemoteInstall.** Contains the OSCfilter.ini file, which holds user options for operating system installation through Remote Installation Services.
    - **Microsoft\IEAK.** Contains settings for the Internet Explorer Maintenance snap-in.
    - **Scripts\Logon.** Contains all the user logon scripts and related files for this GPO.
    - **Scripts\Logoff.** Contains all the user logoff scripts and related files for this GPO.

The User and Machine folders are created at install time, and the other folders are created as needed when policy is set.

The permissions of each Group Policy template reflect the read and write permissions applied to the GroupPolicyContainer through the Group Policy Object Editor. These permissions are automatically maintained and are shown in the following table.

**Default Group Policy Template Permissions**

| Trustee | Access |
|---|---|
| Authenticated Users | Read and Execute |
| Administrators | Full Control |
| Group Policy Creator Owners | Read and Execute |
| Creator Owner | Full Control (Subfolders and Files only) |
| System | Full Control |

**Group Policy Object Editor use of Sysvol**

Each policy setting changed in a GPO causes at least two files to be rewritten — the GPT.ini and the file holding the changed setting. Making many changes to a GPO can cause a lot of network traffic as Sysvol replicates these changes. This congestion should only occur on a local area network where Sysvol replication occurs frequently. Across wide area network links, the

inter-site replication schedule will cause these changes to be amalgamated into a smaller amount of traffic (for example, four changes to the Registry.pol file will result in only a single file replication).

#### The Local Group Policy Object

The Local GPO has no Active Directory component. Information stored in the Group Policy container of an Active Directory GPO is instead stored in the Group Policy template of a Local GPO. The Group Policy template of a Local GPO is located in the **Windows\system32\GroupPolicy** folder. The Gpt.ini file in this GroupPolicy folder must hold more management information than its counterpart in a domain-based GPO because there is no Active Directory component to hold this information. The following table shows the attributes for the Group Policy template.ini file.

**Local GPO GPT.INI Attributes**

| Attribute | Description |
|---|---|
| gPCUserExtensionNames | Includes a list of GUIDs that tells the client-side engine which CSEs have User data in the GPO.<br><br>The format is: [{GUID of CSE}{GUID of MMC extension}{GUID of second MMC extension if appropriate}][repeat first section as appropriate]. |
| GPCMachineExtensionNames | Includes a list of GUIDs that tells the client-side engine which CSEs have computer data in the GPO. |
| Options | Refers to GPO options such as User portion disabled or Computer portion disabled. |

The following extensions are disabled in a Local GPO:

- Group Policy Software installation extension
- Folder Redirection

The following extensions have reduced functionality in a Local GPO:

- Public Key policies; EFS only; there are not any options for trust lists or auto enrollment.
- Security Settings; there are not any options for Restricted Groups or File System, Registry or Service Access Control Lists (ACLs).

#### Core Group Policy Processes and Interactions

Application of GPOs to targeted users and computers relies on many interactive processes. This section explains how GPOs are applied and filtered to Active Directory containers such as sites, domains, and OUs. It includes information about how the Group Policy engine processes GPOs in conjunction with CSEs. In addition, it explains how Group Policy is replicated among domain controllers.

#### Group Policy Processing Rules

GPOs that apply to a user or computer do not all have the same precedence. Settings that are applied later can override settings that are applied earlier. Group Policy settings are processed in the following order:

- **Local Group Policy object.** Each computer has exactly one Group Policy object that is stored locally. This processes for both computer and user Group Policy processing.
- **Site.** Any GPOs that have been linked to the site that the computer belongs to are processed next. Processing is in the order that is specified by the administrator, on the Linked Group Policy Objects tab for the site in GPMC. The GPO with the lowest link order is processed last, and therefore has the highest precedence.
- **Domain.** Processing of multiple domain-linked GPOs is in the order specified by the administrator, on the Linked Group Policy Objects tab for the domain in GPMC. The GPO with the lowest link order is processed last, and therefore has the highest precedence.
- **Organizational units.** GPOs that are linked to the organizational unit that is highest in the Active Directory hierarchy are processed first, then GPOs that are linked to its child organizational unit are processed, and so on. Finally, the GPOs that are linked to the organizational unit that contains the user or computer are processed.

To summarize, the Local GPO is processed first, and the organizational unit to which the computer or user belongs (the one that it is a direct member of) is processed last. All of this processing is subject to the following conditions:

- WMI or security filtering that has been applied to GPOs.

- Any domain-based GPO (not Local GPO) can be enforced by using the **Enforce** option so that its policies cannot be overwritten. Because an Enforced GPO is processed last, no other settings can write over the settings in that GPO. If you have more than one Enforced GPO, it's possible to set the same setting in each GPO to a different value, in which case, the link order of the GPOs determines which one contains the final settings.

- At any domain or organizational unit, Group Policy inheritance can be selectively designated as **Block Inheritance.** However, because enforced GPOs are always applied, and cannot be blocked, blocking inheritance does not prevent policy from **Enforced** GPOs from applying.

Every computer has a single Local GPO that is always processed regardless of whether the computer is part of a domain or is a stand-alone computer. The Local GPO can't be blocked by domain-based GPOs. However, settings in domain GPOs always take precedence since they are processed after the Local GPO.

**Targeting GPOs**

The site, domain, and OU links from a GPO are used as the primary targeting principle for defining which computers and users should receive a GPO. Security filtering and WMI filtering can be used to further reduce the set of computers and users to which the GPO will apply. The Group Policy engine uses the following logic in processing GPOs: If a GPO is linked to a domain, site, or OU that applies to the user or computer, the Group Policy engine must then determine whether the GPO should be added to its GPO list for processing. A GPO is blocked from processing in the following circumstances:

- **The GPO is disabled.** You disable either or both the computer or user components of a GPO from its Policy Properties dialog box.

- **The computer or user does not have permission to read and apply the GPO.** You control permission to a GPO through security filtering, as explained in the following section.

- **A WMI filter applied to a GPO evaluates to false on the client computer.** A WMI filter must evaluate to true before the Group Policy engine will allow it to be processed, as explained in the following section.

**Security Filtering**

Security filtering is a way of refining which users and computers will receive and apply the settings in a GPO. By using security filtering to specify that only certain security principals within a container where the GPO is linked apply the GPO, you can narrow the scope of a GPO so that it applies only to a single group, user, or computer. Security filtering determines whether the GPO as a whole applies to groups, users, or computers; it cannot be used selectively on different settings within a GPO.

In order for the GPO to apply to a given user or computer, that user or computer must have both **Read** and **Apply Group Policy (AGP)** permissions on the GPO, either explicitly, or effectively though group membership.

By default, all GPOs have **Read** and **AGP** both **Allowed** for the Authenticated Users group. The Authenticated Users group includes both users and computers. This is how all authenticated users receive the settings of a new GPO when it is applied to an organizational unit, domain or site. Therefore, the default behavior is for every GPO to apply to every Authenticated User. By default, Domain Admins, Enterprise Admins, and the local system have full control permissions, without the Apply Group Policy access-control entry (ACE). However, administrators are members of Authenticated Users, which means that they will receive the settings in the GPO by default.

These permissions can be changed to limit the scope to a specific set of users, groups, or computers within the organizational unit, domain, or site. The Group Policy Management Console manages these permissions as a single unit, and displays the security filtering for the GPO on the GPO Scope tab. In GPMC, groups, users, and computers can be added or removed as security filters for each GPO.

**How Security Filtering is Processed**

Before processing a GPO, the Group Policy engine checks the Access Control List ACL associated with the GPO. If an ACE on a GPO denies a security principal to which the computer or user belongs, either the Apply Group Policy or Read permission, the Group Policy engine does not add the GPO to its list of GPOs to process. Additionally, an ACE on a GPO must allow the

appropriate security principal both Apply Group Policy and Read permissions in order for the Group Policy engine to add the GPO to the GPO processing list.

If appropriate permissions are granted to the GPO, it is added to the list of GPOs to download.

In general, **Deny** ACEs should be avoided because you can achieve the same results by granting or not granting **Allow** permissions.

**Windows Management Instrumentation (WMI) Filtering**
WMI makes data about a target computer available for administrative use. Such data can include hardware and software inventory, settings, and configuration information. For example, WMI exposes hardware configuration data such as CPU, memory, disk space, and manufacturer, as well as software configuration data from the registry, drivers, file system, Active Directory, the Windows Installer service, networking configuration, and application data.

WMI filtering allows you to filter the application of a GPO by attaching a Windows Query Language query to a GPO. The queries can be written to query WMI for multiple items. If the query returns true for all queried items, then the GPO will be applied to the target user or computer.

When a GPO that is linked to a WMI filter is applied on the target computer, the filter is evaluated on the target computer. If the WMI filter evaluates to false, the GPO is not applied (except if the client computer is running Windows 2000, in which case the filter is ignored and the GPO is always applied). If the WMI filter evaluates to true, the GPO is applied.

The WMI filter is a separate object from the GPO in the directory. A WMI filter must be linked to a GPO in order to apply. Each GPO can have only one WMI filter; however the same WMI filter can be linked to multiple GPOs. WMI filters, like GPOs, are stored only in domains. A WMI filter and the GPO it is linked to must be in the same domain.

**How Windows Management Instrumentation (WMI) Filtering is Processed**
If, after security filtering, appropriate permissions are granted to the GPO, it is added to the list of GPOs to download. Upon download, the Group Policy engine reads the **gPCWQLFilter** attribute in the Group Policy container to determine if a WMI filter is applied to the GPO. If so, the WMI filter, which contains one or more WQL statements, is evaluated. If the statement evaluates to true, then the GPO is processed. There are tradeoffs in using WMI filters because they can increase the amount of time it takes to process policy especially if the filter to be evaluated takes a long time to process.

**Windows Management Instrumentation (WMI)Filtering Scenarios**
Sample uses of WMI filters include:

- Services. Computers where DHCP is turned on.
- Registry. Computers that have this registry key populated.
- Hardware inventory. Computers with a Pentium III processor.
- Software inventory. Computers with Visual Studio .NET installed.
- Hardware configuration. Computers with network interface cards (NICs) on interrupt level 3.
- Software configuration. Computers with multi-casting turned on.
- Associations. Computers that have any services dependent on Systems Network Architecture (SNA) service.

Client support for WMI filters exists only on Windows XP, Windows Server 2003, and later operating systems. Windows 2000 clients will ignore any WMI filter and the GPO is always applied, regardless of the WMI filter. WMI filters are only available in domains that have at least one Windows Server 2003 domain controller.

**Application of Group Policy**
Application of Group Policy involves a series of processes, beginning with user and computer logon.

**Initial Processing of Group Policy**
Group Policy for computers is applied at computer startup. For users, Group Policy is applied when they log on. In Windows 2000, the processing of Group Policy is synchronous, which means that computer Group Policy is completed before the logon dialog box is presented, and user Group Policy is completed before the shell is active and available for the user to interact with it. As explained in the following section, Windows XP with Fast Logon-enabled (which is the default setting) allows users to log

on while Group Policy is processed in the background.)

**Synchronous and Asynchronous Processing**
Synchronous processes can be described as a series of processes where one process must finish running before the next one begins. Asynchronous processes, on the other hand, can run on different threads simultaneously because their outcome is independent of other processes.

You can change the default processing behavior by using a policy setting for each GPO so that processing is asynchronous instead of synchronous. For example, if the policy has been set to remove the **Run** command from the **Start** menu, it is possible under asynchronous processing that a user could logon prior to this policy taking effect, so the user would initially have access to this functionality.

**Fast Logon in Windows XP Professional**
By default in Windows XP Professional, the Fast Logon Optimization feature is enabled for both domain and workgroup members. This means that policy settings apply asynchronously when the computer starts and when the user logs on. This process of applying policies is similar to a background refresh process. As a result, users can logon and begin using the Windows shell faster than they would with synchronous processing. Fast Logon Optimization is always off during logon under the following conditions:

- When a user first logs on to a computer.
- When a user has a roaming user profile or a home directory for logon purposes.
- When a user has synchronous logon scripts.

Note that under the preceding conditions, computer startup can still be asynchronous. However, because logon is synchronous under these conditions, logon does not exhibit optimization. The following table compares policy processing of Windows 2000 and Windows XP client computers.

**Default Policy Processing for Client Computers**

| Client | Application at startup/log on | Application at refresh |
|---|---|---|
| Windows 2000 | Synchronous | Asynchronous |
| Windows XP Professional | Asynchronous | Asynchronous |

Windows XP clients support Fast Logon Optimization in any domain environment. Fast Logon Optimization can be disabled with the following policy setting:

**Computer Configuration\Administrative Templates\System\Logon\ Always wait for the network at computer startup and logon.**

Note that Fast Logon Optimization is not a feature of Windows Server 2003.

**Folder Redirection and Software Installation Policies**
Note that when Fast Logon Optimization is on, a user might need to log on to a computer twice before folder redirection policies and software installation policies are applied. This occurs because the application of these types of policies requires the synchronous policy application. During a policy refresh (which is asynchronous), the system sets a flag indicating that the application of folder redirection or a software installation policy is required. The flag forces synchronous application of the policy at the user's next logon.

**Time Limit for Processing of Group Policy**
Under synchronous processing, there is a time limit of 60 minutes for all of Group Policy to finish processing on the client computer. Any CSEs that are not finished after 60 minutes are signaled to stop, in which case the associated policy settings might not be fully applied.

**Background Refresh of Group Policy**
In addition to the initial processing of Group Policy at startup and logon, Group Policy is applied subsequently in the background on a periodic basis. During a background refresh, a CSE will only reapply the settings if it detects that a change was made on the server in any of its GPOs or its list of GPOs.

In addition, software installation and folder redirection processing occurs only during computer startup or user logon. This is because background processing could cause undesirable results. For example, in software installation, if an application is no

longer assigned, it is removed. If a user is using the application while Group Policy tries to uninstall it or if an assigned application upgrade takes place while someone is using it, errors would occur. Although the Scripts CSE is processed during background refresh, the scripts themselves only run at startup, shutdown, logon, and logoff, as appropriate.

**Periodic Refresh Processing**
By default, Group Policy is processed every 90 minutes with a randomized delay of up to 30 minutes — for a total maximum refresh interval of up to 120 minutes.

Group Policy can be configured on a per-extension basis so that a particular extension is always processed during processing of policy even if the GPOs haven't changed. Policy settings for each extension are located in **Computer Configuration\Administrative Templates\System\Group Policy**.

**On-Demand Processing**
You can also trigger a background refresh of Group Policy on demand from the client. However, the application of Group Policy cannot be pushed to clients on demand from the server.

**Messages and Events**
When Group Policy is applied, a **WM_SETTINGCHANGE** message is sent, and an event is signaled. Applications that can receive window messages can use them to respond to a Group Policy change. Those applications that do not have a window to receive the message (as with most services) can wait for the event.

**Refreshing Policy from the Command Line**
You can update or refresh Group Policy settings manually through a command line tool. On Windows 2000, you can use Secedit with the /refreshpolicy option; on Windows XP and Windows Server 2003, you can use Gpupdate.

**Group Policy and Slow Links**
When Group Policy detects a slow link, it sets a flag to indicate to CSEs that a policy setting is being applied across a slow link. Individual CSEs can determine whether or not to apply a policy setting over the slow link. The default settings are as follows:

### Default Slow Link Settings

| Extension | Default Setting |
|---|---|
| Security Settings | On (and cannot be turned off) |
| Administrative Templates | On (and cannot be turned off) |
| Software Installation | Off |
| Scripts | Off |
| Folder Redirection | Off |

**Group Policy Loopback Support**
Group Policy is applied to the user or computer, based on where the user or computer object is located in Active Directory. However, in some cases, users might need policy applied to them, based on the location of the computer object, not the location of the user object. The Group Policy loopback feature gives you the ability to apply User Group Policy, based on the computer that the user is logging onto. The following figure shows a sample site, domain, and OU structure and is followed by a description of the changes that can occur with loopback processing.

### Sample Active Directory Structure



Normal user Group Policy processing specifies that computers located in the Servers organizational unit have the GPOs A3, A1, A2, A4, and A6 applied (in that order) during computer startup. Users of the Marketing organizational unit have GPOs A3, A1, A2, and A5 applied (in that order), regardless of which computer they log on to.

In some cases this processing order might not be what you want. An example is when you do not want applications that have been assigned or published to the users of the Marketing organizational unit to be installed while they are logged on to the computers in the Servers organizational unit. With the Group Policy loopback feature, you can specify two other ways to retrieve the list of GPOs for any user of the computers in the Servers organizational unit:

- Merge mode. In this mode, the computer's GPOs have higher precedence than the user's GPOs. In this example, the list of GPOs for the computer is A3, A1, A2, A4, and A6, which is added to the user's list of A3, A1, A2, A5, resulting in A3, A1, A2, A5, A3, A1, A2, A4, and 9A6 (listed in lowest to highest priority).

- Replace mode. In this mode, the user's list of GPOs is not gathered. Only the list of GPOs based upon the computer object is used. In this example, the list is A3, A1, A2, A4, and A6.

The loopback feature can be enabled by using the **User Group Policy loopback processing mode** policy under **Computer Settings\Administrative settings\System\Group Policy.**

The processing of the loopback feature is implemented in the Group Policy engine. When the Group Policy engine is about to apply user policy, it looks in the registry for a computer policy, which specifies which mode user policy should be applied in.

**How the Group Policy Engine Processes Client-Side Extensions**
Client-side extensions are the components running on the client system that process and apply the Group Policy settings to that system. There are a number of extensions that are pre-installed in Windows Server 2003. Other Microsoft applications and third party application vendors can also write and install additional extensions to implement Group Policy management of these applications.

The default Windows Server 2003 CSEs are listed in the following table:

Default Windows Server 2003 CSEs

| Client-Side Extension | Active Directory Component | Sysvol Component |
|---|---|---|
| Software Installation | PackageRegistration objects | .aas files |
| Security Settings | | Gptmpl.inf |
| Folder Redirection | | fdeploy.ini |
| Scripts | | Scripts.ini |
| IP Security | IPSec Policy objects | |
| Internet Explorer Maintenance | | .ins + branding .inf files. |
| Administrative Templates | | Registry.pol and .adm files |
| Disk Quota | | Registry.pol |
| EFS Recovery | | Registry.pol |
| Remote Installation | | Oscfilter.ini |
| Wireless Network Policies | | Registry.pol and .adm files |
| QoS Packet Scheduler | | Registry.pol and .adm files |

**Client-Side Extension Operation**
CSEs are called by the Winlogon process at computer startup, user logon and at the Group Policy refresh interval. CSEs are registered with Winlogon in the registry. This registration information includes a DLL and a DLL entry point (function call) by which the CSE processing can be initiated. The Winlogon process uses these to trigger Group Policy processing.

Each extension can opt not to perform processing at any of these points (for example, avoid processing during background

refresh).

**Client-Side Extensions Registered with WinLogon**
Each of the CSEs is registered under the following key:

**HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Winlogon\GPExtensions**

Each extension is identified by a key named after the GUID of the extension. These extensions are shown in the following table.

**CSE Extensions**

| Extension GUID | Extension Name |
|---|---|
| 25537BA6-77A8-11D2-9B6C-0000F8080861 | Folder Redirection |
| 35378EAC-683F-11D2-A89A-00C04FBBCFA2 | Administrative Templates Extension |
| 3610EDA5-77EF-11D2-8DC5-00C04FA31A66 | Disk Quotas |
| 426031c0-0b47-4852-b0ca-ac3d37bfcb39 | QoS Packet Scheduler |
| 42B5FAAE-6536-11D2-AE5A-0000F87571E3 | Scripts |
| 827D319E-6EAC-11D2-A4EA-00C04F79F83A | Security |
| A2E30F80-D7DE-11d2-BBDE-00C04F86AE3B | Internet Explorer Maintenance |
| B1BE8D72-6EAC-11D2-A4EA-00C04F79F83A | EFS Recovery |
| C6DC5466-785A-11D2-84D0-00C04FB169F7 | Software Installation |
| E437BC1C-AA7D-11D2-A382-00C04F991E27 | IP Security |

**How Group Policy Processing History Is Maintained on the Client Computer**
Each time GPOs are processed, a record of all of the GPOs applied to the user or computer is written to the registry. GPOs applied to the local computer are stored in the following registry path:

**HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Group Policy\History**

GPOs applied to the currently logged on user are stored in the following registry path:

**HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Group Policy\History**

**Preferences and Policy Configuration**
Manipulating these registry values directly is not recommended. Most of the items in which you might need to change the behavior of an extension (such as forcing a CSE to run over a slow link), are available as Group Policy settings. These can be found in the Group Policy Object Editor in the following location:

**Computer Settings\Administrative Templates\System\Group Policy**

The behavior can be changed for the following CSEs:

- Administrative Templates (Registry-based policy)
- Internet Explorer Maintenance
- Software Installation
- Folder Redirection
- Scripts
- Security

- IP Security
- EFS recovery
- Disk Quotas

**Order of Extension Processing**
Administrative Templates policy settings are always processed first. Other extensions are processed in an indeterminate order.

**Policy Application Processes**
There are two primary milestones that the Group Policy engine uses for GPO processing:

- Creating the list of GPOs targeted at the user or computer.
- Invoking the relevant CSEs to process the policy settings relevant to them within the GPO list.

The following figure shows the steps required to reach the first milestone in GPO processing, GPO list creation.

**GPO List Creation**




Creating the GPO list involves the following steps:

1. Query the Active Directory for the **gPLink** and **gPOptions** properties in the Site and Domain hierarchies to which the user or computer object belongs.
2. Query the Active Directory for the GroupPolicyContainer objects referenced in the **gPLink** properties.
3. Evaluate security filtering to determine if the user or computer have the Apply Group Policy access permission to the GPO.
4. Evaluate the WMI query against the WMI repository on the client computer to determine if the computer meets the query requirements.

Once the GPO list is created, the Group Policy engine and the CSEs work together to process Group Policy template components. The following figure shows the steps required to determine which CSEs to call.

**Determining CSEs to Call**




Determining which CSEs to call involves the following steps:

1. Retrieve the list of CSEs registered with Winlogon.

2. Check to see whether it is appropriate to run a particular CSE (for example, whether background processing or slow link processing is enabled for the extension).

3. Check the CSE history against list of Applied GPOs. GPOs with new version numbers and GPOs that have settings relevant to the CSE (that is, they have the CSE extension GUID in the Group Policy container gpcUserExtension or gpcMachineExtension properties) are added to the Changed GPO List. GPOs no longer in the Applied GPO List are added to the Deleted GPO List.

4. Check to see whether the appropriate CSE should be processing policy settings for the user or the computer.

5. Check the version number listed in the GPO against its recorded version history in the registry to determine whether the GPO needs reprocessing.

If all of the version numbers are unchanged the MaxNoGPOListChanges interval might have expired; if so, the CSE processes policy settings without regard to an unchanged version number.

Steps 3 through 5 are repeated by each CSE for all GPOs in the GPO list. After one CSE is done, the next CSE that needs to run repeats the entire process.

Group Policy updates are dynamic and occur at specific intervals. If there have been no changes to Group Policy, the client computer still refreshes the security policy settings at regular intervals for the GPO.

If no changes are discovered, GPOs are not processed. Security policies have a periodic force apply every 16 hours. For security policies, there is a value that sets a maximum limit of how long a client can function without reapplying non-changed GPOs. By default, this setting is every 16 hours plus a randomized delay of up to 30 minutes. Even when GPOs that contain security policy settings do not change, the policy is reapplied every 16 hours

**Group Policy Replication**
In a domain that contains more than one domain controller, Group Policy information takes time to propagate, or replicate, from one domain controller to another. Low bandwidth network connections between domain controllers slow replication. The Group Policy infrastructure has mechanisms to manage these issues.

Each GPO is stored partly in the Sysvol on the domain controller and partly in Active Directory. GPMC and Group Policy Object Editor present and manage the GPO as a single unit. For example, when you set permissions on a GPO in GPMC, GPMC is actually setting permissions on objects in both Active Directory and the Sysvol. It is not recommended that you manipulate these separate objects independently outside of GPMC and the Group Policy Object Editor. As shown in the following figure, it is important to understand that these two separate components of a GPO rely on different replication mechanisms. The file system portion is replicated through FRS, independently of the replication handled by Active Directory. Only the Sysvol subfolder (%systemroot%\SYSVOL\sysvol) is shared and replicated. Sysvol was designed to allow multiple domain's Sysvols to be replicated in the same tree — each domain's Sysvol is contained under a subfolder of the Sysvol share. For the current domain, a copy of the domain's Sysvol subtree is also stored directly under the **%systemroot%\SYSVOL\domain** folder.

**Group Policy Replication**




FRS is a multi-master replication service that synchronizes folders between two or more Windows Server 2003 or Windows 2000 systems. Modified files are queued for replication at the point the file is closed. In the case of conflicting modifications between two copies of an FRS replica, the file with the latest modification time will overwrite any other copies. This is referred to as a "last-writer-wins" model.

FRS replication topology configuration is stored as a combination of child objects of each FRS replica partner (in the FRS Subscriptions subcontainer) and objects within another hidden subcontainer of the domain System container. Replication links between systems are maintained as FRS subscription objects. These objects specify the replica partner and the replication schedule. It is possible to view the schedule by browsing to an FRS subscription object and viewing the properties. The replica partner is stored as the object GUID of the computer account of that partner.

The Sysvol folder is a special case of FRS replication. Active Directory automatically maintains the subscription objects and their schedules as the directory replication is built and maintained. It is possible, but not recommended, to modify the properties (for example, the schedule) of the Sysvol subscription objects manually.

The FRS replication schedule only approximates to the directory replication schedule so it is possible for the directory-based Group Policy information and the file-based information to get temporarily out of synch. Since GPO version information is stored in both the Group Policy container object and in the Group Policy template, any discrepancy can be viewed with tools such as Gpotool.exe and Repadmin.exe.

For those Group Policy extensions that store data in only one data store (either Active Directory or Sysvol), this is not an issue, and Group Policy is applied as it can be read. Such extensions include Administrative Templates, Scripts, Folder Redirection, and most of the Security Settings.

For any Group Policy extension that stores data in both storage places (Active Directory and Sysvol), the extension must properly handle the possibility that the data is unsynchronized. This is also true for extensions that need multiple objects in a single store to be atomic in nature, since neither storage location handles transactions.

An example of an extension that stores data in Active Directory and Sysvol is Group Policy Software installation extension. The .aas files are stored on Sysvol and the Windows Installer package definition is in Active Directory. If the .aas file exists, but the corresponding Active Directory components are not present, the software is not installed. If the .aas file is missing, but the package is known in Active Directory, application installation fails gracefully and will be retried on the next processing of Group Policy.

The tools used to manage Active Directory and Group Policy, such as GPMC, the Group Policy Object Editor, and Active Directory Users and Computers all communicate with domain controllers. If there are several domain controllers available, changes made to objects like users, computers, organizational units, and GPOs might take time to appear on other domain controllers. The administrator might see different data depending on the last domain controller on which changes were made and which domain controller they are currently viewing the data from.

If multiple administrators manage a common GPO, it is recommended that all administrators use the same domain controller when editing a particular GPO, to avoid collisions in FRS. Domain Admins can use a policy to specify how Group Policy chooses a domain controller — that is, they can specify which domain controller option should be used. The Group Policy domain controller selection policy setting is available in the Administrative Templates node for User Configuration, in the System\Group Policy subcontainer.

**Network Ports Used by Group Policy**

### Port Assignments for Group Policy

| Service Name | UDP | TCP |
|---|---|---|
| Lightweight Directory Access Protocol | n/a | 389 |
| SMB | n/a | 445 |
| DCOM | Dynamicallly assigned | Dynamically assigned |
| RPC | Dynamically assigned | Dynamically assigned |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Group Policy Settings Reference for Windows Server
- Group Policy Administrative Tools

# Core Group Policy Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Group Policy Tools
- Group Policy Settings
- Group Policy WMI Classes
- Related Information


This section summarizes the tools and settings associated with Group Policy.

**Group Policy Tools**
The following is a list of administrative tools associated with Group Policy. For more information about Group Policy administrative tools, see the following topics in this collection:

- Group Policy Administrative Tools
- Group Policy Object Editor Tools and Settings
- Group Policy Management Console Technical Reference
- What Is Resultant Set of Policy?

For information about tools specific to Group Policy extensions, see the appropriate Group Policy Components topics in this collection. For more information about other Resource Kit tools, see the Windows Server 2003 Deployment Kit Tools page.

**GPUpdate.exe**
This tool is used for refreshing local and Active Directory policy settings on the computer from which you run the **GPUpdate** command.

**Category**
This command-line tool is included in Windows XP and Windows Server 2003.

**Version compatibility**
You can use GPUpdate locally on Windows XP and higher computers to refresh policy immediately. On computers running Windows 2000, this behavior is provided by the using the secedit.exe command line tool, with a specific parameter.

GPUpdate refreshes local Group Policy settings and Group Policy settings that are stored in Active Directory, including security settings, on the computer from which it is run. This command supersedes the now obsolete **/refreshpolicy** option for the secedit command line tool. For more information about GPUpdate, type **GPUpdate /?** at the command line.

**GPResult.exe**
GPResult.exe is a Group Policy tool for examining the settings applied during Group Policy refresh.

**Category**
There are two versions of GPResult: One shipped with the Windows 2000 Resource Kit; the other is included with Windows XP and Windows Server 2003. The Windows 2000 version runs only locally on Windows 2000. The windows Server 2003 version runs locally or remotely on Windows XP or Windows Server 2003.

The different versions are not compatible.

**Version compatibility**
GPResult utilizes Resultant Set of Policy (RSoP) data. You can use GPResult that shipped with Windows Server 2003 family on Windows XP and higher computers. You can use GPResult that shipped with Windows 2000 Resource Kit for Windows 2000.

GPResult for Windows Server 2003 displays Group Policy settings and Resultant Set of Policy (RSoP) for a user or a computer. Because you can apply overlapping levels of policies to any computer or user, the Group Policy feature generates a resulting set of policies at logon. GPResult displays the resulting set of policies that were enforced on the computer for the specified user at logon.

GPResult for Windows 2000 estimates the Group Policy settings that would be applied at a specific computer. Full documentation for this version of GPResult is available in the readme file distributed with the tool.

**Dcgpofix.exe: Dcgpofix**
**Category**
Dcgpofix ships with Windows Server 2003.

**Version compatibility**
You can run Dcgpofix only on servers running Windows Server 2003 family. This tool can restore default domain policy and default domain controllers policy to their original state after installation, except for some security-related settings that are impossible to return to their exact original state. When you run Dcgpofix, you will lose any changes made to these Group Policy objects. For more information about Dcgpofix, type **Dcgpofix /?** at the command line.

This tool should be used as a last-resort disaster-recovery tool. A better solution is to use GPMC to back up and restore these GPOs.

**GPMonitor.exe: Group Policy Monitor Tool**
**Category**
Group Policy Monitor tool is included in the Windows Server 2003 Deployment Kit.

**Version compatibility**
The Group Policy Monitor tool works on Windows XP and higher computers. Group Policy Monitor tool collects Group Policy information at every Group Policy refresh and sends that information to a centralized location that you specify. You can then use the Group Policy Monitor user interface (UI) to view the data. The Group Policy Monitor UI can provide a historical view of policy changes. The UI is also designed to make it easy to navigate through historical snapshots of data and trace changes. For more information about the Group Policy Monitor tool, type **GPMonitor /?** at the command line. You can find full documentation for the Group Policy Monitor tool in the Windows Server 2003 Deployment Kit Tools.

**GPOTool.exe: Group Policy Verification Tool**
**Category**
Group Policy Verification tool is included in the Windows Server 2003 Deployment Kit.

**Version compatibility**
The Group Policy Verification tool works on Windows 2000 and higher computers. You use Group Policy Verification tool to check the health of the Group Policy objects on domain controllers. The tool checks GPOs for consistency on each domain controller in your domain. The tool also determines whether the policies are valid and displays detailed information about replicated Group Policy objects (GPOs).

GPOTool.exe ships with the Microsoft Windows 2003 Server Resource Kit and is also available as a free download at the Gpotool.exe: Group Policy Verification Tool page.

For more information about the Group Policy Verification tool, type **GPOTool /?** at the command line. You can find full documentation for Group Policy Verification tool in the Windows Server 2003 Deployment Kit Tools.

**AdmX.exe: ADM File Parser**
**Category**
The ADM File Parser (AdmX) is a command-line tool that enables an administrator to export Group Policy settings to a tab-delimited text file. The administrator can then use the text produced by ADM File Parser (AdmX) to find changes for the policy settings between different versions of the operating systems. AdmX is for use only with policies based on administrative templates.

**Version compatibility**
The AdmX.exe tool runs on Windows 2000, Windows Server 2003, and Windows XP Professional. AdmX.exe also requires the Microsoft .NET Framework 1.0.

Administrators use AdmX.exe to extract policy setting information from .adm files and to list differences between two .adm files.

For more information about the ADM File Parser, type **AdmX.exe /?** at the command line.

**Other Tools**
For information about tools specific to Group Policy extensions, see the appropriate Group Policy Components Tools topics in

this collection.

**Group Policy Settings**
Many Group Policy settings are configurable within the Administrative Template (.adm) files supplied with Windows 2000 and Windows Server 2003. Administrators can access these files from the Group Policy Object Editor for a specific GPO, and can configure Group Policy settings contained in the template files.

A registry setting exists for a Group Policy template setting only when the setting state has been changed from **Not Configured** to either **Enabled** or **Disabled**. You can change the setting state from the appropriate Administrative Template node in the Group Policy Object Editor.

To download a complete reference to the Group Policy settings in the .adm files supplied with Windows Server 2003, see the Group Policy Settings Reference for Windows Server 2003. This reference is in Microsoft Excel spreadsheet format.

For a reference to settings specific to Group Policy behavior, and for a list of registry locations that are associated with Group Policy on the domain controller, target, or both, see the Group Policy Object Editor Tools and Settings topic in this collection.

For a reference to settings specific to Group Policy extensions, see the appropriate Group Policy Components topics in this collection.

**Group Policy WMI Classes**
For complete information about WMI classes associated with Group Policy, see the WMI SDK documentation on MSDN.

**Related Information**
The following resources contain additional information that is relevant to this section.

- The Group Policy Administrative Tools topic in this collection.
- The Group Policy Object Editor Tools and Settings topic in this collection.
- The What Is Group Policy Management Console? topic in this collection.
- The What Is Resultant Set of Policy? topic in this collection.
- The appropriate Group Policy Components Tools topics in this collection.
- The Group Policy Management Console page.
- The Group Policy Settings Reference for Windows Server 2003.
- The Microsoft Platform SDK page.
- The Windows Server 2003 Deployment Kit Tools page.
- The Tools and Settings Collection.

# Group Policy Components

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Components**

Depending on the size and complexity of the organization, administrators must set, maintain, and manage various standard configurations for groups of users and computers. Users may require specialized software or have unique needs, such as the ability to always work remotely. Client computers must be set up for specific purposes with configurations for kiosks, labs, general office use, and other roles.

Group Policy is the infrastructure that enables you to deliver and apply these and other configurations to targeted users and computers within an Active Directory directory service environment.

The Group Policy infrastructure consists of a Group Policy engine and various individual components known as client-side extensions that "plug" into the infrastructure. These client-side extensions are used to configure Group Policy settings either by modifying the registry (through the Administrative Templates extension) or setting Group Policy settings for security settings, software installation, folder redirection, wireless network settings, and other areas. Group Policy settings are contained in Group Policy objects (GPOs) that are applied to users and computers.

Group Policy components provide you with the means of enforcing and maintaining a wide variety of configurations, thereby reducing the risk of disruption in your computing environment. In addition, the Group Policy infrastructure is designed to be flexible and extensible, allowing you to develop your own extensions, if desired.

You administer Group Policy components using two primary tools. The first tool is Group Policy Management Console (GPMC), which lets you create, view, and manage GPOs. The second tool is Group Policy Object Editor, which which lets you configure and modify settings within GPOs.

**Group Policy Component Architecture**

Group Policy settings are implemented by a variety of components. There are two basic types of components:

- Extensions that allow you to administer and configure Group Policy settings using Group Policy Object Editor. These are server-side snap-ins that are extensions to the Microsoft Management Console (MMC) snap-in.
- Extensions that interpret Group Policy settings and apply them on the client computer. These are extensions to the logon process – Winlogon – and are known as client-side extensions.

**Group Policy Component Architecture**




**Group Policy Components**

**Server Side MMC Snap-in Extensions**

You configure Group Policy settings using Microsoft Management Console (MMC) snap-ins. The snap-in provides the basic organization of the Group Policy Object Editor and the top-level namespace (Software Settings, Windows Settings, and Administrative Templates). There are server-side snap-in extensions for most of the client-side extensions shown in the figure earlier. However, some client-side extensions, such as Disk Quotas, do not require separate server side snap-ins. In many cases client-side extensions only require registry settings to be transferred to the client system and these can be configured using the Administrative Templates snap-in.

Server-side snap-in extensions include:

- Administrative Templates.
- Security Settings.
- IP Security.
- Scripts.
- Software Installation.
- Internet Explorer Maintenance.

All MMC snap-ins and extensions are registered under the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MMC\SnapIns
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MMC\Nodetypes
```

Each extension may have further child extensions. In this way, Group Policy Object Editor consists of an extensible tree of administrative components. This allows components to be updated or new components to be added without affecting other components.

**Client-Side Extensions**

Client-side extensions are the components running on the client system that process and apply the Group Settings to that system. There are a number of extensions that are preinstalled in Windows Server 2003. Other Microsoft applications and third-party application vendors can also write and install additional extensions to implement Group Policy management of these applications. The default CSEs are listed in the following table.

**Default Client-Side Extensions**

| Client-Side Extension | Active Directory Component | Sysvol Component |
|---|---|---|
| Administrative Templates | | Registry.pol (+ .adm) |
| Software Installation | PackageRegistration objects | .aas files |
| Security Settings | | Gptmpl.inf |
| Folder Redirection | | fdeploy.ini |
| Scripts | | Scripts |
| IP Security | IPSec Policy objects | |
| Internet Explorer Maintenance | | .ins and branding .inf files. |
| Disk Quotas | | Registry.pol |
| EFS Recovery | | Registry.pol |
| Remote Installation | | Oscfilter.ini |
| QoS Packet Scheduler | | Registry.pol and .adm files. |

Client-side extensions are implemented within DLLs that are installed and registered on the client computer during operating

system installation. To trigger Group Policy processing, the Group Policy engine running within Winlogon calls each CSE DLL using the CSE registration settings contained below the Winlogon key in the registry.

Each of the CSEs is registered under the following registry key: `HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtensions`

**Group Policy Component Scenarios**
Group Policy is used to define configurations for groups of users and computers. With Group Policy, administrators can specify Group Policy settings for registry-based policies, security, software installation, scripts, folder redirection, remote installation services, and Internet Explorer maintenance. Group Policy settings are contained in a GPO. By linking a GPO with selected Active Directory system containers—sites, domains, and organizational units—the GPO's Group Policy settings are applied to the users and computers in those Active Directory containers.

**Administrative Templates**
Administrative Templates enable administrators to control registry settings using Group Policy, providing the means to configure the behavior and appearance of the desktop, including the operating system, components, and applications. Windows comes with a predefined set of Administrative Template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a GPO. These .adm files are stored in two locations by default: inside GPOs in the **Sysvol** folder and in the **%windir%\inf** directory on the local computer. Administrative Templates Group Policy settings are also referred to as .adm settings or registry-based settings.

As new versions of Windows are released, new Administrative Templates policy settings are added. In addition to supporting these new settings, each successive version of Windows supports all registry policy settings that were available in earlier versions of Windows. For example, the Windows Server 2003 family of operating systems supports all Administrative Templates policy settings available in Windows 2000 and Windows XP.

**Security Settings**
You can define security settings that are enforced on any number of computers, by using the Security Settings extension of Group Policy. By specifying the security configuration of a Group Policy object (GPO), you can affect the computers in a site, domain, or organizational unit to which the GPO is linked.

You can configure the following types of security settings:

- Account Policies. This includes settings for Password Policy, Account Lockout Policy, and Kerberos Policy.
- Local Policies. This includes settings for Audit Policy, User Rights Assignment, and Security Options.
- Event Log. This includes settings for application, system, and security Event Log settings.
- Restricted Groups. This specifies policies for membership of security-sensitive groups.
- System Services. This specifies policies for startup and permissions for system services.
- Registry. This includes policies that specify permissions for registry keys.
- File System. This includes policies that specify permissions for folders and files.

**IP Security Policy**
Internet Protocol security (IPSec) is a framework of open standards designed to ensure private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services. IPSec supports network-level data integrity, data confidentiality, data origin authentication, and replay protection.

In an Active Directory environment, you can use Group Policy to centralize IPSec policy distribution and management. You can use the **IP Security Policies** snap-in to configure IPSec policies to meet the security requirements of a user, group, application, domain, site, or global organization. You can also specify IPSec policies on a local computer that is not a member of a domain.

**Software Restriction Policies**
The Windows Server 2003 family of operating systems and Windows XP include software restriction policies, a new policy-based feature that identifies software programs running on computers in a domain, and controls the ability of those programs to run. Software restriction policies are a part of Microsoft's security and management strategy to assist enterprises in increasing the reliability, integrity, and manageability of their computer networks. You can specify software restriction policies in an Active Directory environment or on a local computer.

You can use **Software Restriction Policies** to:

- Fight viruses.
- Regulate which ActiveX controls can be downloaded.
- Run only digitally signed scripts.
- Enforce that only approved software is installed on system computers.
- Highly restrict (lock down) a computer.

**Wireless Network Policies**
A new Wireless Network (IEEE 802.11) Policies Group Policy extension allows you to configure wireless network settings that are part of Group Policy for Computer Configuration. Wireless network settings include the list of preferred networks, Wired Equivalent Privacy (WEP) settings, and IEEE 802.1X settings. These settings are downloaded to targeted domain members, making it much easier to deploy a specific configuration for secure wireless connections to wireless client computers.

**Public Key Policies**
These Group Policy settings are used to:

- Specify that computers automatically submit a certificate request to an enterprise certification authority and install the issued certificate.
- Create and distribute a certificate trust list.
- Establish common trusted root certification authorities.
- Add encrypted data recovery agents and change the encrypted data recovery policy settings.

**Software Installation**
The Software Installation snap-in is used to centrally manage software. Software can be assigned or published to users and assigned to computers. Group Policy-based Software Installation can be used to install software applications when a computer is started, when the user logs on, or on demand. Software installation Group Policy settings can be applied to users or computers in an Active Directory structure.

Group Policy-based Software Installation can also be used to upgrade deployed applications or remove earlier applications that are no longer required. Users can be restricted from installing any software from local media, such as a CD-ROM, or disk, or other unapproved applications.

Medium and large organizations may wish to consider using Microsoft Systems Management Server (SMS). SMS provides advanced capabilities such as inventory-based targeting, status reporting, server- and client-side scheduling, multisite facilities, complex targeting, centralized hardware and software inventory, remote diagnostic tools, software metering, software distribution-point population and maintenance, support for Windows 95, Windows 98, Windows NT 4.0, Windows 2000, and Windows XP clients, and enhanced software deployment features. SMS does not require Active Directory.

**Remote Installation Services**
Remote Installation Services (RIS) is an optional component that is included in the Windows Server operating system and works with other Windows Server 2003 technologies to implement the Remote Operating System Installation feature. Administrators use Remote Operating System Installation to remotely install a copy of the Windows XP Professional operating system on supported computers. Administrators use the RIS extension of Group Policy to specify which options are presented to users by the Client Installation Wizard, for example, Automatic Setup, Custom Setup, and Restart Setup.

**Scripts**
Scripts are used to automate tasks at computer startup and shutdown, and at user logon and logoff. Scripts can be written in any language supported by Windows Script Host including the Microsoft Visual Basic development system, Scripting Edition (VBScript), JavaScript, PERL, and MS DOS-style batch files (.bat and .cmd).

**Internet Explorer Maintenance**
Internet Explorer Maintenance is used to manage and customize Internet Explorer on computers running Windows 2000 or later. Administrators can set options for the Browser UI, connections, URLs, proxy settings, security zones, Favorites, and

Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening).

**Folder Redirection**
Folder redirection is used to redirect special directories on Windows 2000 or Windows Server 2003 from their default user profile location to an alternate location on the network. These special folders include **My Documents**, **Application Data**, **Desktop**, and the **Start** menu.

**Disk Quotas**
Disk quotas provide a way to limit each users utilization of disk space on a volume. Disk quotas are implemented as part of the NTFS File System. Disk quota settings are read from the registry at startup and during GPO processing. Disk Quotas are configured through the Administrative Templates snap-in and are stored in the registry.pol GPO registry settings file – there is no GPO snap-in for disk quotas. The Disk Quota CSE retrieves the resultant setting from the client system's registry and configures the disk quota parameters accordingly. Prioritization of multiple GPO settings is handled by the Administrative Templates CSE.

**QoS Packet Scheduler**
Quality of service (QoS) allows you to use your existing resources efficiently and guarantee that critical applications receive high-quality service, without having to expand as quickly or even over-provision your networks. Quality of Service parameters are configured as registry settings using the Administrative Templates snap-in. The QoS Packet Scheduler CSE reads these settings from the registry and configures the Packet Scheduler service on the client system with these parameters.

# Administrative Templates Extension Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Administrative Templates extension consists of a Microsoft Management Console (MMC) server-side snap-in used to configure policy settings and a client-side extension that sets registry keys on target computers. Administrative Templates policy is also known as registry-based policy or registry policy. Although other Group Policy extensions manipulate the registry to implement policy settings, the Administrative Templates Extension exists specifically for registry-based policy.

**In this subject**


- What Is Administrative Templates Extension?
- How Administrative Templates Extension Works
- Administrative Templates Extension Tools and Settings

# What Is Administrative Templates Extension?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Core Scenarios for Administrative Templates Extension
- Related Information


The Administrative Templates Extension is the largest of all available Group Policy extensions and includes more than 700 policy settings for applications and operating system components. These policy settings are applied by modifying the registry on target clients. Administrative Templates policy settings is also referred to as registry-based policy or simply registry policy.

Administrators need a simple way to configure policy settings and apply those changes to many users and computers throughout the network. You need to be able to modify policy settings quickly and be able to delete policy settings and remove them from all target computers without the risk of old policy settings remaining in the registry. In addition, developers need a way to integrate policy management into new applications. Administrative Templates provides dynamic management capabilities to administrators and an infrastructure for developers to policy-enable their applications.

The Administrative Templates Extension consists of a server-side snap-in that is loaded by default in Group Policy Object Editor and a client-side extension that writes policy settings that manipulate registry keys on target client computers. The server-side snap-in loads a predefined set of Administrative Template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a Group Policy object (GPO).

.Adm files are Unicode files which consist of a hierarchy of categories and subcategories that define how the options are displayed through the Group Policy Object Editor and GPMC. They also indicate the registry locations where changes should be made if a particular selection is made, specify any options or restrictions (in values) that are associated with the selection, and in some cases, indicate a default value to use if a selection is activated. It's important to note that the functionality of .adm files is limited. The only purpose of .adm files is to enable a user interface to configure policy settings. .Adm files do not contain actual policy settings; these are contained in registry.pol files located in the Sysvol on domain controllers.

**Core Scenarios for Administrative Templates Extension**
You use Administrative Templates Extension to modify and remove registry-based policy settings. Developers can create .adm files to integrate policy settings management into their applications.

**Manipulating Administrative Template Extension Policy Settings**
The Administrative Templates server-side snap-in provides an Administrative Templates node that appears twice in Group Policy Object Editor; once under the Computer Configuration node, and again under the User Configuration node. The settings under the computer node manipulate registry settings for the computer, while settings under the user node manipulate registry settings for users. Although some policy settings require simple UI elements, such as text boxes to enter values, most settings contain only the following options:

- Enabled
- Disabled
- Not configured

The default setting is not configured, so you only need to decide whether to turn the setting on or to turn it off.

Administrative Templates provide administrators with a Group Policy interface for the DSQ and QoS CSEs, among others.

**Removing Policy Settings Implemented by Administrative Template Extension**
In addition to being easy to set, administrators can also easily remove default registry-based policy settings implemented with Administrative Templates Extension. This is accomplished by having all default registry-based policy settings from Windows 2000 or later stored in one of four specific registry keys:

- HKEY_LOCAL_MACHINE\SOFTWARE\policies (preferred location)

- HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\policies
- HKEY_CURRENT_USER\SOFTWARE\policies (preferred location)
- HKEY_ CURRENT_USER\SOFTWARE\Microsoft\Windows\CurrentVersion\policies

Policy settings that are stored in these locations are known as true policies. Registry keys used for true policy are secured so they cannot be modified by a non-administrator. More importantly, when Group Policy changes, for any reason, all settings from these registry keys are removed, and new registry keys are written based on the policy settings in the GPO. This prevents the behavior that was often present in Windows NT 4.0, whereby System Policies resulted in persistent settings in the user and computer registry. The policy remained in effect until the value was reversed, either by a counteracting policy or by editing the registry. Persistent policy settings persist because they are stored outside the approved registry locations. These types of settings are known as preferences.

All the policy settings in the default administrative template files manipulate settings under the approved registry keys—resulting in true policy settings. This means that they will not cause persistent settings in the registry when the GPO that applies them is no longer in effect. By default, Group Policy Object Editor only displays true policy settings.

**Default Administrative Template Files**
Windows Server 2003 includes the following administrative template files: System.adm, Inetres.adm, Conf.adm, Wmplayer.adm, and Wuau.adm, which contain all the settings initially displayed in the Administrative Templates node. These administrative template files are stored in two locations by default: inside GPOs in the Sysvol folder and in the Windows\inf directory on the local computer.

As new versions of Windows are released, new policy settings are added. In addition to supporting these new settings, each successive version of Windows supports all registry policy settings that were available in earlier versions of Windows. For example, the Windows Server 2003 family supports all registry policy settings available in Windows 2000 and Windows XP.

On Windows XP and Windows Server 2003, each registry setting contains a "Supported on" tag that indicates which operating system versions support that policy setting. If a setting is specified and deployed to a client operating system that does not support that setting, the settings are ignored.

Because all successive iterations of administrative template files include settings from earlier versions, and because there is no harm if a new setting is applied inadvertently to a computer running an earlier operating system that does not support that setting, it is recommended to always create and edit GPOs from a computer that has the latest administrative template files available.

**Extending Registry-Based Policy with Administrative Template Files**
If a developer needs to provide policy settings specific to an application, they can extend registry-based policy by using administrative template files. The Administrative Templates Extension to Group Policy provides this capability.

Developers can specify that that Administrative Templates Extension writes settings under the secure registry keys to create true policy settings; these policy settings are designed to be manipulated by an administrator using Group Policy Object Editor. Developers can also specify that the application write policy settings outside the secure registry keys to create preferences; these settings are designed to be manipulated by end-users working from within the developer's application. The Group Policy settings set by administrators take precedence over user preferences. Because Group Policy stores the policy settings set by administrators under the approved Group Policy keys, users cannot change or disable these policy settings.

**Related Information**
The following contains additional information that is relevant to this section.

- Group Policy Settings Reference in the Tools and Settings Collection

# How Administrative Templates Extension Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Administrative Templates are the primary means of configuring the client computer's registry settings through Group Policy. Organizations benefit the most by using Administrative Templates policy settings in a Windows Server 2003 environment with Active Directory installed and Windows XP client computers.

**In this section**


- Administrative Templates Extension Architecture
- Administrative Templates Extension Physical Structure
- Administrative Templates Extension Processes and Interactions
- Network Ports Used by Administrative Templates Extension
- Related Information


**Administrative Templates Extension Architecture**

The following diagram shows how components interact in implementing Administrative Templates policy settings.

**Administrative Templates Extension Architecture**




Components and protocols significant to Administrative Templates Extension are included in the following table.

**Administrative Templates Components**

| Component | Description |
|---|---|
| Administrator | The computer that you use to configure Administrative Templates policy settings in Local Group Policy Editor or Group Policy Management Console (GPMC). |

| | |
|---|---|
| Domain controller | The server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. Each domain controller contains:<br>• The Group Policy container (GPC), which stores information about GPO properties in Active Directory.<br>• The Group Policy template (GPT), which stores GPO data in the Sysvol. Data includes the Registry.pol file that stores Administrative Templates policy settings. |
| Target client | The computer(s) to which you intend to apply Administrative Templates policy settings. |
| Registry | A database repository for information about a computer's configuration, the registry is organized hierarchically as a tree, and is made up of keys and their subkeys, hives, and entries. Administrative Templates Extension directly modifies registry keys in order to configure Machine and User policy settings. |
| Administrative Templates Snap-in Extension | The MMC server-side snap-in extension used to configure and modify Administrative Templates-based policy settings. The snap-in is contained within userenv.dll located in Windows\System32\. It appears in Local Group Policy Editor or GPMC as a node under Computer Configuration and User Configuration. |
| Administrative Templates client-side extension | Administrative Templates client-side extension runs inside userenv.dll and is responsible for modifying the registry according to the Administrative Template policy settings that you configure in Local Group Policy Editor or GPMC.<br>Userenv.dll is registered at the following location:<br>HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtensions.<br>This is the location of all Group Policy CSE registration. |
| Active Directory | A requirement for implementing domain-based Group Policy, Active Directory provides the containers in which you link Group Policy objects (GPOs). |
| Group Policy engine | The framework that handles functionalities across client-side extensions (CSEs). The Group Policy engine invokes the Administrative Templates client-side extension with a list of GPOs to be applied. |
| Event log | The Event log is a service, located in Event Viewer, which records events in the system, security, and application logs. |
| File system | The NTFS file system on client computers. |
| Group Policy object | Administrative Templates policy settings are contained in GPOs, which can then be linked to Active Directory containers such as sites, domains, and OUs. |
| Local Group Policy object | The GPO stored on each individual computer in the hidden *%systemroot%*\System32\GroupPolicy directory. Although you can configure Administrative Templates policy settings in local GPOs, they are the least influential GPOs in an Active Directory environment, because Active Directory-based GPOs have precedence. |
| Kerberos | Any time the registry extensions (client or server) access any information about the Sysvol, or any other domain resource, authentication traffic is generated to the Active Directory. Kerberos is an authentication mechanism used to verify user or host identity. The Kerberos V5 authentication protocol is the default authentication service. Administrative Templates client-side extension communicates with Active Directory using Kerberos. |
| NTLM | A security package that provides authentication between clients and servers. NTLM is also used by Administrative Templates client-side extension to communicate with Active Directory. |

| | |
|---|---|
| WMI | Windows Management Infrastructure (WMI) is a management infrastructure that supports monitoring and controlling of system resources through a common set of interfaces and provides a logically organized, consistent model of Windows operation, configuration, and status.<br><br>Each computer contains a WMI database that stores information about policy settings. WMI makes data about a target computer available for administrative use.<br><br>The actual Administrative Templates policy settings that are in effect on a client computer (known as Resultant Set of Policy or RSoP) is periodically updated. During processing, the Administrative Templates CSE writes the processing data to the RSoP namespace in WMI.<br><br>The namespaces in WMI to which the extension writes are:<br><br>● Root/RSoP/Computer.<br><br>● Root/RSoP/User/<user SID>. |
| LDAP | The Lightweight Directory Access Protocol used to communicate with Active Directory. |
| SMB | Server Message Block (SMB) protocol is the primary method of file and print sharing. Both the Administrative Templates client-side extension and the Administrative Templates snap-in use SMB to access the Sysvol as well as back up and retrieve files to a remote file system. The client computer also uses SMB to read the Sysvol on the domain controller.<br><br>When you edit policy settings, the Administrative Templates snap-in writes changes to the Registry.pol file in the GPO located on the Sysvol. (Or the local GPO on the local computer). The ADM files in the ADM directory can also be updated at administration time.<br><br>During processing the Administrative Templates CSE only reads the Registry.pol files from the GPO. No other files are read. No files are written to the Sysvol during processing. Communication to the Sysvol takes place through standard Win32 file system API calls. |
| Distributed Component Object Model (DCOM) | DCOM is used by the Administrative Templates CSE to communicate with the WMI database. |
| dskquoui.dll | The dskquoui.dll file provides the Quota tab user interface on NTFS volumes. |
| dskquota.dll | The dskquota.dll file provides the Disk Quotas client-side extension and server-side extension. |
| gptext.dll | The gptext.dll file provides the QoS Packet Scheduler client-side extension and server-side extension |

**Administrative Templates Extension Physical Structure**

The following figure shows the files used by the Administrative Templates Extension.




Descriptions of these files are included in the following table.

**Physical Structure Components**

| Component | Description |
|---|---|
| Sysvol | The Sysvol is a set of folders containing important domain information that is stored in the file system rather than in Active Directory. The Sysvol folder is stored in a subfolder of systemroot folder (*%\systemroot*\sysvol\sysvol) and is automatically created when a server is promoted to a domain controller. |
| . adm files | .Adm files are Unicode text files that enable a user interface to allow you to modify Registry-based policy settings using Local Group Policy Editor or GPMC. .Adm files do not contain any policy setting information. Actual policy setting information is contained in Registry.pol files.<br><br>.Adm files are stored in two places:<br><br>• Local .adm files stored on the computer where you run Local Group Policy Editor or GPMC.<br><br>• Domain-based .adm files stored on the Sysvol. These files are copies of the local .adm files and are updated whenever you modify Administrative Templates policy settings. |
| System.adm | Provides policy settings to configure the operating system. System.adm is loaded by default in Windows Server 2000, Windows XP, and Windows Server 2003. |
| Inetres.adm | Provides policy settings to configure Internet Explorer. Ineteres.adm is loaded by default in Windows 2000 Server, Windows XP, and Windows Server 2003. |
| Conf.adm | Provides policy settings to configure NetMeeting. Conf.adm is loaded by default in Windows 2000 Server, Windows XP, and Windows Server 2003.<br><br>**Note:** Conf.adm is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. |
| Wmplayer.adm | Provides policy settings to configure Windows Media Player. Wmplayer.adm is loaded by default in Windows XP and Windows Server 2003. Wmplayer.adm is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. |

| | |
|---|---|
| Wuau.adm | Provides policy settings to configure Windows Update. Wuau.adm is loaded by default in Windows 2000 Service Pack 3 (SP3), Windows XP Service Pack 1 (SP1), and Windows Server 2003. |
| Registry.pol | Unicode files that contain the Administrative Templates policy settings that are to be applied to the computer or user portion of the registry.<br><br>In contrast to .adm files, the Registry.pol files contain the actual Group Policy settings used by the Group Policy engine during processing. Registry.pol files contain instructions to add or delete registry keys, corresponding to the policy settings you specify in Local Group Policy Editor or GPMC. |
| Ntuser.pol | An archive file on client computers that is updated each time the Administrative Templates CSE sets a policy setting. As processing completes, an Ntuser.pol containing the history of applied registry based policies in the Group Policy managed policies tree is written to the root of the user's profile containing user policy settings and to the "all users" profile containing computer policy settings.(%Allusersprofile%\ntuser.pol for computer policy and %userprofile%\ntuser.pol for user policy.) |
| Userenv.dll | The dynamic link library in which the Administrative Templates CSE runs. Userenv.dll is part of the code of the Group Policy engine and is always called whenever policy settings are processed. |
| Userenv.log | A log file that records warnings and events as a result of Administrative Templates processing. The Userenv.log contains more verbose logging than the event log and is not enabled by default. It is located on client target computers in Windows\debug\usermode. UserEnv.log also contains entries for profiles. |
| Gptext.log | A log file located on client computers in Windows\debug\usermode. You can use gptext.log to view errors in the processing of Administrative Templates policy settings. |

**Administrative Templates Extension Processes and Interactions**
The Administrative Templates CSE is invoked by the Group Policy engine with a list of GPOs to be applied. Each GPO contains a Registry.pol file. In order to modify settings in the registry, the Administrative Templates CSE parses the ntuser.pol files to remove previously set registry settings, returning the registry settings to a "default" state. The Administrative Templates Extension then retrieves the Registry.pol files from each GPO in the list and applies the Registry.pol changes to the local registry in order of precedence for the GPOs. Data is also written to WMI by the Administrative Templates Extension for RSoP to retrieve later. As the Registry.pol files are processed, the ntuser.pol file is created to centrally manage the registry keys that have been set. Events are written to the event log during this process. If debugging is enabled, data is written to the userenv.log as well.

Unlike other CSEs, the Administrative Templates CSE is coded within the Group Policy engine as a piece of userenv.dll. As a result, the extension is always called whenever Group Policy is processed. In addition, the Administrative Templates CSE is always called first by the Group Policy engine, in advance of any other CSEs that might be processed. Administrative Templates CSE will run only if there are any changes in any the GPOs that the Administrative Templates CSE has enabled.

**True Policies and Preferences**
The Administrative Templates CSE has control over a part of the registry for both user and computer registry hives and treats these specially. These parts are for the computer and user hives respectively:

- HKEY_LOCAL_MACHINE\SOFTWARE\policies (preferred location)
- HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\policies
- HKEY_CURRENT_USER\SOFTWARE\policies (preferred location)
- HKEY_ CURRENT_USER\SOFTWARE\Microsoft\Windows\CurrentVersion\policies

These trees cannot be modified by a non-administrator. Because all keys and values beneath these paths are erased before applying the resultant registry policy settings, the registry policies applied in these subtrees will only persist as long as a valid Group Policy setting exists. Policy settings that are stored in these specific locations of the registry are known as true policies.

All the policy settings in the standard Administrative Template files that shipped with Windows 2000 Server and Windows Server 2003 use true policies. This prevents the behavior that was often present in Windows NT 4.0, whereby System Policies resulted in persistent settings in the user and computer registry. The policy remained in effect until the value was reversed, either by a counteracting policy or by editing the registry. These settings are stored outside the approved registry locations listed and are known as preferences.

Although Group Policy settings take priority over preferences, they do not overwrite or touch the registry key used by the preference. If both a policy and preference are present, the preference will be successfully restored if the policy is removed or disabled. Preference settings persist in the registry until they are reversed by a counteracting policy setting or by editing the registry.

The configuration of the wallpaper on the Windows desktop illustrates an example of simultaneous policy and preference settings. In the Windows shell, it is possible for a user to configure their desktop wallpaper using the Display icon in Control Panel. An administrator can also configure desktop wallpaper using a default policy setting called **Active Desktop Wallpaper**, which can be found under **User Configuration\Administrative Templates\Desktop\Active Desktop** node in Local Group Policy Editor or GPMC.

The following table lists the resultant behavior for Group Policy settings and preferences.

**Results of Group Policy Settings and Preferences**

| Scenario | Policy present | Preference present | Resultant behavior |
|---|---|---|---|
| 1 | No | No | Default |
| 2 | No | Yes | Preference configures behavior |
| 3 | Yes | No | Policy configures behavior |
| 4 | Yes | Yes | Policy configures behavior; preference ignored |

It is common practice to offer both a preference and a policy setting for most applications. The application reads the registry keys and uses them accordingly. Registry-based data is appropriate for many types of policy settings and is also the least complex way to create custom policy settings. In addition, Registry-based policy managed through administrative template files automatically supports Resultant Set of Policy (RSoP) capabilities.

**Registry.pol Files**

Registry.pol files are formed by the following process:

- When you start Local Group Policy Editor or GPMC, a temporary registry tree is created that consists of two nodes: **USER** and **MACHINE**.
- As you navigate the Administrative Templates node of the Local Group Policy Editor or GPMC, .adm file nodes are displayed. The .adm files within Local Group Policy Editor or GPMC nodes are loaded dynamically when a particular node is selected, and the .adm file is then cached.
- When a policy is selected in the details pane (the right side of the Local Group Policy Editor or GPMC), the temporary registry is queried to determine whether the selected policy already has registry values assigned to it; if it does, those values are displayed in the **Policy** dialog box. If the selected policy does not have a registry value assigned to it, the default value from the .adm file or from the associated MMC snap-in extension is used.
- After you modify a policy, the registry values that you specify are written to the appropriate portion of the temporary registry (either **MACHINE** or **USER**).
- When you close Local Group Policy Editor or GPMC, the temporary registry hives are exported to the Registry.pol files in the appropriate folders of the Group Policy template: GPT\Machine and GPT\User.
- The next time you start Local Group Policy Editor or GPMC for the same Group Policy object for which you have previously set Group Policy settings, the registry information from the corresponding Registry.pol files is imported into the temporary registry tree. Therefore, when you view the policy settings, they reflect the current state.

**Registry.pol in the Local GPO**

During processing, the local Registry.pol files in the Machine and User directories of the Local GPO are read. There is a Registry.pol file in both the Machine and User directories of the Local GPO directory structure for machine- and user-based policies respectively. The Local GPO is distinct from domain policy in that it updates the Registry.pol file as each policy setting is set and refreshes policy settings on the computer at the time each policy setting is configured.

**.Adm Files**

Administrative Template files describe where Registry-based policy settings are stored in the registry, by associating a description and explain text with a registry key and value. Local Group Policy Editor or GPMC displays only the descriptive text and provides various dialog boxes that you can use to modify the setting. A section of the user's hive is mapped to the registry policy setting. .Adm files, unlike Registry.pol files, do not affect the actual policy processing by the Administrative Templates CSE. .Adm files only affect the display of the policy settings in the Local Group Policy Editor or GPMC snap-in. If an .adm file is removed, the settings corresponding to the .adm file will not appear in Local Group Policy Editor or GPMC. However, the policy settings that are configured from the .adm file will remain in the Registry.pol file and continue to apply to the appropriate target client or user.

Each administrative workstation that is used to run Local Group Policy Editor or GPMC stores .adm files in the Windows\Inf folder. When GPOs are created and first edited, the .adm files from this folder are copied to the \adm subfolder in the Group Policy template (GPT).

By default, when GPOs are edited, Local Group Policy Editor or GPMC compares the time stamps of the .adm files in the workstation's Windows\Inf folder with those that are stored in the GPT \adm folder. If the workstation's files are newer, Local Group Policy Editor or GPMC copies these files to the GPT \adm folder, overwriting any existing files of the same name. This comparison occurs when the Administrative Templates node (computer or user configuration) is selected in Local Group Policy Editor or GPMC, regardless of whether you actually edit the GPO. The .adm files stored in the Group Policy template can be updated by viewing a GPO in Local Group Policy Editor or GPMC. The process is simplified for local GPOs where all adm files are stored locally in a single adm folder.

Because of the importance of time stamps on .adm file management, the editing of system-supplied .adm files is not recommended. If a new policy setting is required, Microsoft recommends that you create a custom .adm file. This prevents the replacement of system-supplied .adm files when service packs are released.

**Using the latest .adm files**
As a general rule, each operating system or service pack release includes a superset of the .adm files provided by earlier releases, including policy settings that are specific to operating systems that are different to those of the new release. For example, the .adm files that are provided with Windows Server 2003 include all policy settings for all operating systems, including those that are only relevant to Windows 2000 or Windows XP Professional. This means that only viewing a GPO from a computer with the new release of an operating system or service pack effectively upgrades the .adm files.

**How .adm files are handled by Local Group Policy Editor or GPMC**
By default Local Group Policy Editor or GPMC attempts to read .adm files from the GPO (from the Sysvol on the domain controller). Alternatively, the .adm file can be read from the local workstation computer. This behavior can be controlled by a policy setting.

By default, if the version of the .adm file found on the local computer is newer (based on the time stamp of the file) than the version on the Sysvol, the local version is copied to the Sysvol and is then used to display the settings. This behavior can be controlled by a policy setting.

If the GPO contains registry settings for which there is no corresponding .adm file, these settings cannot be seen in Local Group Policy Editor or GPMC. However, the policy settings are still active and will be applied to users or computers targeted by the GPO.

**How .adm files are handled by Group Policy Management Console**
GPMC uses .adm files to display the friendly names of policy settings when generating HTML reports for GPOs, Group Policy Modeling, and Group Policy Results.

By default, GPMC uses the local .adm file, regardless of time stamp. If the file is not found, GPMC will look in the GPO's directory on Sysvol.

You can specify an alternate path for where to find .adm files. If specified, this takes precedence over the previous locations. GPMC never copies the .adm file to the Sysvol.

**Extensions that use Administrative Templates**
There are additional extensions that are located within the Administrative Templates for Computer Configuration in the Local Group Policy Editor or GPMC. These are:

- Disk Quotas
- QoS Packet Scheduler.

**Disk Quotas**
Disk Quotas are used to manage NTFS file system disk space. Administrators use the Disk Quotas Extension to configure

Group Policy for Disk Quotas on target computers. The Disk Quotas Extension includes a server-side extension and a client-side extension.

Administrators manage Disk Quotas policy settings under the following node in the Local Group Policy Editor or GPMC: Computer Configuration\Administrative Templates\System\Disk Quotas. The Disk Quotas node is the user interface for the server-side component of the Disk Quotas extension. There is no user interface for the client-side component, although you can view changes made by the CSE on the Quota Property tab for NTFS volumes. The Group Policy engine, using the Disk Quotas client-side extension component, applies settings to the target computer.

The Disk Quotas CSE is registered with Winlogon in the registry at the following path: {HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\WindowsNT\CurrentVersion\Winlogon\GPExtensions\{3610eda5-77ef-11d2-8dc5-00c04fa31a66}

**QoS Packet Scheduler**
QoS Packet Scheduler Extension is an extension to Local Group Policy Editor or GPMC. Administrators use QoS Packet Scheduler Extension to set QoS Packet Scheduler Group Policy.

QoS Packet Scheduler Extension is included in the same binary (gptext.dll) as the Scripts, IP Security, and Wireless Group Policy extensions. Administrators manage QoS Packet Scheduler policy settings under the following node in the Local Group Policy Editor or GPMC: Computer Configuration\Administrative Templates\Network\QoS Packet Scheduler.

The QoS Packet Scheduler CSE is registered with Winlogon in the registry at the following path:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\ GPExtensions\{426031c0-0b47-4852-b0ca-ac3d37bfcb39}

**Network Ports Used by Administrative Templates Extension**
Administrative Templates CSE communicates with Active Directory using LDAP to complete the following tasks:

- GPO list retrieval.
- Group Policy container retrieval.

It uses SMB to complete the following tasks:

- Request DFS referral for \\domainname\sysvol.
- SysvolDFS replica location \\dcname.domainname\sysvol.
- Open and read GPT.INI.
- Return GPT.INI.
- Open and read GPT settings files.
- Return GPT file.

### Port Assignments for Group Policy Administrative Templates

| Service Name | UDP | TCP |
|---|---|---|
| Lightweight Directory Access Protocol | n/a | 389 |
| SMB | n/a | 445 |

**Related Information**
The following contains additional information that is relevant to this section.

- “Group Policy Settings Reference” in the Tools and Settings Collection

# Administrative Templates Extension Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Administrative Templates Extension Tools
- Administrative Templates Extension Registry Locations
- Administrative Templates Extension Group Policy Settings
- Administrative Templates Extension WMI Classes
- Network Ports Used by Administrative Templates Extension
- Related Information


This section describes the tools and templates available for the Group Policy Administrative Templates extension.

**Administrative Templates Extension Tools**
The following tools are associated with the Group Policy Administrative Templates extension. For more information about Group Policy administrative tools, see the following topics in this collection:

- Group Policy Administrative Tools
- Group Policy Object Editor Tools and Settings
- Group Policy Management Console Technical Reference
- What Is Resultant Set of Policy?

For information about tools specific to other Group Policy extensions, see the appropriate Group Policy Components Tools topics in this collection. For more information about other Resource Kit tools, see the Windows Server 2003 Resource Kit Tools page.

**AdmX.exe: ADM File Parser**
**Category**
The ADM File Parser (AdmX) is a command-line tool that enables an administrator to export Group Policy settings to a tab-delimited text file. The administrator can then use the text produced by ADM File Parser (AdmX) to find changes for the policy settings between different versions of the operating systems. AdmX is for use only with policies based on administrative templates.

**Version compatibility**
The AdmX.exe tool runs on Windows 2000, Windows Server 2003, and Windows XP Professional. AdmX.exe also requires the Microsoft .NET Framework 1.0.

Administrators use AdmX.exe to perform the following two tasks:

- Extract policy setting information from .adm files.
- Compare two .adm files. This can be used to determine what has changed between two like files, such as the system.adm file from two different operating systems.

For more information about the ADM File Parser, type **AdmX.exe /?** at the command line.

**GPUpdate.exe**

This tool is used for refreshing local and Active Directory policy settings on the computer from which you run the **GPUpdate** command.

**Category**
This command-line tool is included in Windows XP and Windows Server 2003.

**Version compatibility**
You can use GPUpdate locally on computers running Windows XP and higher to refresh policy immediately. On computers running Windows 2000, this behavior is provided by the using the secedit.exe command line tool, with a specific parameter.

GPUpdate refreshes changed local Group Policy settings and Group Policy settings that are stored in Active Directory, including security settings, on the computer from which it is run. This command supersedes the now obsolete **/refreshpolicy** option for the secedit command line tool. Group Policy can be forced to reprocess all policy settings whether there has been a change or not by using the **/force** switch.

For more information about GPUpdate, see "Core Group Policy Tools and Settings" in this collection.

**GPResult.exe: Group Policy Results**
This command line tool displays Resultant Set of Policy (RSoP) for a user or computer. You can use GPResult.exe to see what policy is in effect and to troubleshoot problems on computers running Windows XP or later.

**Category**
There are two versions of GPResult: One shipped with the Windows 2000 Resource Kit; the other is included with Windows XP and Windows Server 2003. The Windows 2000 version of GPResult runs only locally on Windows 2000. The Windows Server 2003 version runs locally or remotely on Windows XP or Windows Server 2003.

The different versions are not compatible.

**Version compatibility**
The GPResult.exe that ships with Windows XP and Windows Server 2003 family is completely different than the original version of GPResult.exe that shipped in the Windows 2000 Resource Kit. The newer version cannot be used to view policy information for computers running Windows 2000.

For more information about GPResult, see "Core Group Policy Tools and Settings" in this collection.

**Dcgpofix.exe: Dcgpofix**
**Category**
Dcgpofix ships with Windows Server 2003.

**Version compatibility**
You can run Dcgpofix only on servers in a Windows Server 2003 domain. This tool can restore default domain policy and default domain controllers policy to their original state after installation. When you run Dcgpofix, you will lose any changes made to these Group Policy objects. This tool should be used as a last-resort disaster-recovery tool. A better solution is to use Group Policy Management Console (GPMC) to back up and restore these GPOs.

For more information about Dcgpofix, type **Dcgpofix /?** at the command line.

**GPMonitor.exe: Group Policy Monitor Tool**
**Category**
Group Policy Monitor tool is included in the Windows Server 2003 Deployment Kit.

**Version compatibility**
The Group Policy Monitor tool works on Windows XP and higher computers. Group Policy Monitor tool collects Group Policy information at every Group Policy refresh and sends the information to a centralized location that you specify. You can then use the Group Policy Monitor user interface (UI) to view the data. The Group Policy Monitor UI can provide a historical view of policy changes. The UI is also designed to make it easy to navigate through historical snapshots of data and trace changes. For more information about the Group Policy Monitor tool, type **GPMonitor /?** at the command line. You can find full documentation for the Group Policy Monitor tool in the Windows Server 2003 Resource Kit Tools.

**GPOTool.exe: Group Policy Verification Tool**
**Category**
Group Policy Verification Tool is included in the Windows Server 2003 Deployment Kit.

**Version compatibility**
The Group Policy Verification tool works on Windows 2000 and higher computers. You can use Group Policy Verification Tool to check the health of the Group Policy objects on domain controllers. The tool checks GPOs for consistency on each domain controller in your environment. The tool also determines whether the policies are valid and displays detailed information about replicated Group Policy objects (GPOs).

GPOTool.exe ships with the Microsoft Windows 2003 Server Resource Kit and is also available as a free download at the Gpotool.exe: Group Policy Verification Tool page.

For more information about the Group Policy Verification tool, type **GPOTool /?** at the command line. You can find full documentation for Group Policy Verification Tool in the Windows Server 2003 Resource Kit Tools.

**Regview.exe: Registry.pol Viewer Tool**
**Category**
Regview.exe is a Group Policy tool for viewing the contents of any registry.pol file, which is created by the Group Policy Object Editor for containing the administrative settings. It is included in the Windows Server 2003 Deployment Kit.

**Version compatibility**
The Group Policy Viewer tool works on Windows 2000 and higher computers. Because the registry.pol file is not an ASCII file, you cannot easily view the contents. If an administrator needs to troubleshoot or just wants to document the contents of the file, this tool supplies the contents in a delimited format.

**Other Tools**
For information about tools specific to Group Policy extensions, see the appropriate Group Policy Components Tools topics in this collection.

**Administrative Templates Extension Registry Locations**
The following registry locations are associated with the Administrative Templates extension. All default registry-based policy settings from Windows 2000 or later are stored in one of four specific registry keys:

- HKEY_LOCAL_MACHINE\SOFTWARE\Policies (preferred location)
- HKEY_CURRENT_USER\SOFTWARE\Policies (preferred location)
- HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\policies
- HKEY_ CURRENT_USER\SOFTWARE\Microsoft\Windows\CurrentVersion\policies

Policy settings that are stored in these locations are known as true policies.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. You should not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**Administrative Templates Extension Group Policy Settings**
Administrative templates (.adm files) enable administrators to configure registry settings by using Group Policy. These settings appear under the Administrative Templates item for both User Configuration and Computer Configuration in the console tree of the Group Policy Object Editor (GPEdit.dll), and in HTML reports produced by the Group Policy Management Console (GPMC).

For a full list of Administrative Templates policy settings, see Group Policy Settings Reference for Windows Server 2003.

**Group Policy Settings**
For more information about Group Policy settings, see the Group Policy Settings Reference for Windows Server 2003.

**Administrative Templates Extension WMI Classes**
Like other Administrative Templates policy settings, RSoP data is updated whenever policies applied to a computer changes. The extension writes this data into WMI during processing so that the data maintained in WMI is consistent with what is applied to the computer. The administrative template files themselves used in editing the GPO are stored in the Windows/inf folder of the administrative computer. The extension writes to the following WMI namespaces:

- Root/RSoP/Computer
- Root/RSoP/User/<user SID>

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by Administrative Templates Extension**
The Group Policy Administrative Templates extension communicates with Active Directory using Lightweight Directory Access Protocol (LDAP) to complete the following tasks:

- GPO list retrieval
- Group Policy container retrieval

The Administrative Templates extension uses Server Message Block (SMB) to complete the following tasks:

- Request Distributed File System (DFS) referral for \\domainname\sysvol
- SysvolDFS replica location \\dcname.domainname\sysvol
- Open and read GPT.INI
- Return GPT.INI
- Open and read Group Policy template (GPT) settings files
- Return GPT file

The following table lists the network ports used by the Administrative Templates extension.

**Port Assignments for Group Policy Administrative Templates**

| Service Name | UDP | TCP |
|---|---|---|
| Lightweight Directory Access Protocol | n/a | 389 |
| SMB | n/a | 445 |

**Related Information**
The following resources contain additional information that is relevant to this section.

- The Group Policy Administrative Tools topic in this collection.
- The Group Policy Object Editor Tools and Settings topic in this collection.
- The What Is Group Policy Management Console? topic in this collection.
- The appropriate Group Policy Components Tools topics in this collection.
- The Implementing Registry-Based Group Policy page.
- The Group Policy Management Console page.
- The Group Policy Settings Reference for Windows Server 2003.
- The Microsoft Platform SDK page.
- The Windows Server 2003 Resource Kit Tools page.
- The Tools and Settings Collection

# Group Policy Software Installation Extension Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By using the Group Policy Software installation extension of Microsoft Windows Server 2003, an administrator can centrally deploy, install, and manage applications throughout an organization. From a central location, the administrator can also perform routine maintenance tasks such as upgrading, patching, and removing applications without going to individual workstations, which gives users reliable access to the applications that they need to perform their jobs, no matter which computer they are using.

The Group Policy Software installation extension relies on both the Active Directory directory service included with Windows Server 2003 and Group Policy.

**In this subject**


- What Is Group Policy Software Installation Extension?
- How Group Policy Software Installation Extension Works
- Group Policy Software Installation Extension Tools and Settings

# What Is Group Policy Software Installation Extension?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Group Policy Software Installation Extension?**

**In this section**


- Group Policy Software Installation Extension Scenarios
- Group Policy Software Installation Extension Dependencies
- Comparing Microsoft Software Management Solutions


This section provides a high-level overview of the Group Policy Software installation extension and compares it with other Microsoft software-management technologies.

Microsoft offers several key technologies to aid organizations with deploying new software and software upgrades or updates. One of the key reasons for implementing the Group Policy Software installation extension, a Group Policy-based technology, is to lower the company's total cost of ownership of its computers. Because installing, upgrading, and updating applications are critical business processes, companies can leverage the built-in Group Policy software distribution technology to increase users' productivity as changes occur.

The following table describes the capabilities of the Group Policy Software installation extension.

**Group Policy-Based Software Installation Capabilities**

| Capability | Description |
|---|---|
| Central location to create distribution instructions | Windows Server 2003 administrators create the packages and configure the details of the packages that are published to Active Directory. |
| Central location to initiate delivery of software | By using Group Policy configurations, software is made available to computers and users that are part of the Active Directory hierarchy. |
| Scheduling | Software is made available to users immediately in **Add or Remove Programs** in Control Panel, or is automatically installed at computer startup or user logon. |
| Distribution targeting | Software deployment uses the Active Directory and Group Policy infrastructure services that are built into Windows Server 2003 to perform its targeting. |
| Installation and configuration status | Group Policy Results (formerly known as RSoP) provides detailed information on the client about what has been installed and why. |
| Distribution reporting | Group Policy provides interfaces for collecting data on individual software installations, but provides no such network-wide data collection mechanisms. |
| Disaster recovery for applications | With Group Policy, assigned applications automatically reinstall after system loss, and published applications are available for users to reinstall on-demand or by choosing the software in **Add or Remove Programs**. |

**Group Policy Software Installation Extension Scenarios**

The Group Policy Software installation extension enables the administrator to create a controlled environment, providing on-demand software installation and automatic repair of applications. Users benefit from reliable access to the applications that they need to perform their jobs on any computer they use on their network.

Scenarios where organizations might use Group Policy-based software distribution include the following:

- Publish applications that are nonessential for the users. When software is published for a user, it does not initially appear

to be installed on the computer. There is no Windows Installer advertisement information about the software on the computer in the registry, on the desktop, or on **Start** menu as a shortcut. On an as-needed basis, the user installs the published software by using **Add or Remove Programs** in **Control Panel**. Users can also install the published application by selecting a file that has a file name extension for an application.

- Assign software to users or computers for either of the following reasons:
  - To make a particular application available to all users of a computer, assign the application to the computer.
  - To make mission-critical software available to users or computers at all times, assign the application to the users.

**Group Policy Software Installation Extension Dependencies**
The Microsoft Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems provide software distribution by using Group Policy, which is built on the management infrastructure services in the operating system. Group Policy requires Active Directory directory service. Windows Server 2003 provides a robust and feature-laden software distribution mechanism for organizations by using Active Directory. The Group Policy Software installation extension also depends on Windows Installer.

For detailed information about Group Policy dependencies, see "Core Group Policy Technical Reference."

**Comparing Microsoft Software Management Solutions**
Microsoft offers several software solutions for networked users. An organization might already have objectives and requirements for a software installation and management product. This section compares Microsoft technologies for software deployment.

Group Policy, which is built-in to the Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems, offers a convenient method for distributing software in an Active Directory environment, especially if it already uses Group Policy for other purposes, such as securing client and server computers. However, a Group Policy-based software installation has some basic limitations, including difficulties with scheduling installation, consistently managing network bandwidth, and providing feedback on the status of the installation.

To carefully schedule installations, manage network use, perform hardware and software inventory, or monitor installation status, consider using Microsoft Systems Management Server (SMS). For more information about SMS, see the Microsoft Systems Management Server Product Information page.

Using the right solutions can benefit an organization by providing a centralized, efficient means to perform routine tasks such as updating software. The following table compares the various software management technologies.

**Comparing Software Management Technologies**

| Management Function | Group Policy | Microsoft Systems Management Server (SMS) | Terminal Services | Microsoft Software Update Services (SUS) |
|---|---|---|---|---|
| Patch and upgrade Windows XP, Windows Server 2003, and Windows 2000 | N/A | Yes<br><br>When using SMS for software management, also use it to patch Windows-based computers instead of SUS. | Although Terminal Services does not automate patching, it can be used it to remotely log on and apply patches. | Windows patches only (no upgrade) |
| Create consistent user environment (persistence of data, software, and settings) | Yes | Software only | Yes | N/A |
| Perform disaster recovery for applications in Windows 2000, Windows XP, and Windows Server 2003 | Yes | Yes | N/A | N/A |

| | | | | |
|---|---|---|---|---|
| Perform inventory, advanced deployment, troubleshooting, and diagnostic tools | Limited | Yes | Limited | None |
| Manage environments that are not Active Directory-based | No | Yes | Yes | Yes (Windows patches only) |

Although all these Microsoft management technologies provide important software distribution capabilities, SMS is the preferred Microsoft software distribution solution for medium-sized, and especially for enterprise-sized, organizations. SMS provides advanced features for deploying and managing software, Windows patches, and critical updates. For organizations that use SMS as a software management solution, the Microsoft Software Update Services (SUS) SMS Feature Pack for SMS 2.0 provides the features of SUS for users of SMS 2.0 for distributing patches and critical updates to clients. However, SUS used with the Automatic Updates client is the recommended solution for distributing Windows patches in conjunction with Group Policy–based software distribution.

Although there are specific instances where an administrator might choose one software deployment method over another, many of these Microsoft technologies can be used together, depending on an organization's needs.

The following sections describe each of these solutions.

**Group Policy Software Installation Extension**

The Group Policy Software installation extension is well-suited to deploy and manage software if the organization is small or medium in size, and the following conditions exist:

- The administrator has deployed Active Directory.
- The administrator has determined that Group Policy provides the management features the organization requires.
- The administrator has a base of client computers running Windows 2000 Professional or Windows XP Professional and member servers running Windows 2000 Server and Windows Server 2003. Note that the servers do not have to run Windows Server 2003 to use Group Policy–based software deployment.

Group Policy can also serve the needs of large enterprises that use other software installation solutions, such as SMS, across the organization. In conjunction with SMS or other solutions, Group Policy can be useful for distributing software within various groups, such as individual divisions, where the advanced capabilities of SMS might not be needed.

The following figure shows the Group Policy Object Editor interface after a new package has been added to a GPO by using the Group Policy Software installation extension.

**Group Policy Software Installation Extension**




This topic details the Group Policy Software installation extension for providing application management.

**Systems Management Server**

SMS is appropriate for organizations if any of the following conditions exist:

- The organization is medium or large in size.

- Its users are running operating systems earlier than Windows 2000 Professional.
- The organization requires more advanced capabilities for planning, scheduling, distributing, and tracking software.

The advanced capabilities of SMS include such features as inventory-based targeting, status reporting, server-side and client-side scheduling, support for multisite facilities, centralized hardware and software inventorying, remote diagnostic tools, software metering, software distribution-point population and maintenance, and other enhanced software deployment features. SMS also provides support for Microsoft Windows 95, Windows 98, Windows NT 4.0, Windows 2000, and Windows XP clients. Additionally, SMS does not require Active Directory. For more information about SMS, see the Microsoft Systems Management Server Product Information page.

**Terminal Services**

Terminal Services can be very useful if an organization has Windows-based desktop applications that require frequent updates, and the users who require those applications are in remote locations and have low bandwidth. When used as a terminal server, a server becomes a Windows application server. This allows the user to run Windows-based applications remotely on the server while only the mouse, keyboard, and display data are transmitted to the local computer. Terminal Services allows an administrator to offer users software as a remote service instead of as a local installation package, as with Group Policy-based software distribution. For more information about Microsoft Terminal Services, see the Terminal Services topic.

**Microsoft Software Update Services**

Microsoft Software Update Services (SUS) can be used to quickly acquire and distribute critical Windows patches to computers in an organization. By using SUS, an administrator can choose which of the latest critical or security patches to download, test them in a company-standard operating environment, and then efficiently deploy the patches to the appropriate computers running the Automatic Updates client.

# How Group Policy Software Installation Extension Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Group Policy Software Installation Extension Works**

**In this section**


- Terms and Definitions
- Group Policy Software Installation Extension Architecture
- Group Policy Software Installation Extension Components
- Group Policy Software Installation Extension Processes and Interactions


The Group Policy Software installation extension allows administrators to use the Group Policy Object Editor to centrally manage the installation of software on all client computers in an organization. This is accomplished either by assigning applications to users or computers, or by publishing applications for users.

Software can be assigned on a per-user or per-computer basis when an organization does not want to give users the choice to install or remove the software. For example, if a user removes a user-assigned application by using Add or Remove Programs in Control Panel, the Group Policy Software installation extension automatically reapplies the advertisement information after the user logs on or the computer restarts, and the software is reinstalled the next time a user selects it or tries to open a file with an associated file name extension. It is not possible for a user to delete a computer-assigned application. In most cases, packages that are assigned to users or computers include applications that are essential but do not create congestion between the clients and the software distribution points.

Group Policy-based software deployment allows administrators to publish software for users only; in other words, published software is not available for computers. Publishing software for users allows them to decide if and when they want to install it. They can install the software from a list of published applications in Add or Remove Programs. For example, not everyone in the organization requires software for project management. Therefore, a software administrator is likely to publish a project management package for only those users who require it. Managers who require the software can then choose to install it.

Users can always see both assigned and published applications in Add or Remove Programs. Add or Remove Programs includes an active Web link for each application, which provides users with the support information they need to install certain applications. However, administrators can overwrite the default link by using the Group Policy Software installation extension. The support link then corresponds to the internal product support resources.

This Web link can also point to a support page that includes information such as an FAQ about a specified application, a help desk article about using the application, or instructions for requesting support.

**Group Policy–based Software Installation in the Ideal Environment**

The following components are necessary to deploy software using the Group Policy Software installation extension:

- Group Policy.
- Windows 2000 Server or Windows Server 2003 domain controllers.
- A network that uses Active Directory directory service–based domains.
- Target computers that run Windows 2000 Server, Windows 2000 Professional, Windows XP Professional, or Windows Server 2003.

The Group Policy Management Console (GPMC) is not required but necessary for simple, powerful, and efficient administration. See the Group Policy topic for more information about GPMC.

**Terms and Definitions**

The following are terms used in this document to describe aspects of Group Policy–based software deployment.

**Advertisement**
The process of making an application that supports Windows Installer appear installed on the desktop, by installing the shortcuts for the application on the Start menu or desktop, and by populating the registry with the application information.

**Advertisement script**
The output of the Group Policy Software installation extension Microsoft Management Console (MMC) snap-in, which advertises an application that supports Windows Installer. The advertisement script is stored in Active Directory.

**Assigned**
One of the modes of targeting supported by Group Policy-based software deployment. Administrators can assign the software that users require to perform their jobs. Administrators can assign software to users and to computers. User-assigned applications are advertised on the desktop; they appear to be installed on the user's desktop, but they get installed the first time that the user starts the application, typically from the Start menu. If a user deletes an assigned application, it will be re-advertised. Computer-assigned applications are installed the next time that the computer restarts.

**Package**
A Windows Installer .msi file. At its core is a database that contains the instructions to manage the software on a computer.

**Published**
One of the modes of targeting supported by Group Policy-based software deployment. Administrators can publish software to users. Published applications are not advertised on the desktop; there is no physical appearance of published applications on the user's desktop. Users install published applications from the Add or Remove Programs in Control Panel or by opening a file type associated with an application.

**Repackaging**
The process of preparing software for distribution by making an image of a correctly configured computer, then installing the software and making a post-installation image of the computer. The repackaging software takes the difference of the two images and creates the necessary installation instructions to reproduce the installation.

**Group Policy Software installation extension snap-in**
The MMC component of Group Policy Object Editor that an administrator uses to manage software.

**Transform**
(Verb) The process of modifying or customizing a Windows Installer-based .msi package.

**Transform**
(Noun) A Windows Installer customized package that, if associated with another Windows Installer-based .msi package at deployment time, results in a customized installation.

**Windows Installer**
Windows Installer is a Windows operating system–based service that can install, modify, repair, and remove software. Windows Installer includes the operating system–based service, a package format, and an application-programming interface (API) that allows both the operating system and applications to interact with the service to install, modify, or repair the software.

**Group Policy Software Installation Extension Architecture**
This section details the architecture behind Group Policy–based software distribution. The Group Policy Software installation extension performs two processes: application publishing and application assignment. The following figure shows the server-side perspective of Group Policy–based software deployment.

### Server Architecture and Components of the Group Policy Software Installation Extension




**Application Publishing**

*Publishing* is simply advertising that a package is available to be installed – it does not mandate that an application be installed. A published package is created by using the Group Policy Software installation extension. The Group Policy Software installation extension extracts advertising information about the package and stores this, in modified form, in both the directory and the Sysvol folder. Advertising information includes the name and installation path of the application; it also includes information about the file associations, Program IDs (ProgID), and Component Object Model (COM) Class IDs (CLSID) that the application supports or implements.

Published applications can be either Windows Installer .msi files or text-based .zap files. Each is treated differently: The Group Policy Software installation extension uses the Windows Installer subsystem to extract advertising information from the .msi package. This information is programmed into the .msi package by the author of the package, normally the application vendor. A .zap package is a simple text wrapper around a setup command. The author of the .zap file can include information such as file name extension and COM support in the .zap file. This information is extracted directly by the Group Policy Software installation extension.

A subset of the advertising information – file name extensions, CLSIDs, and ProgIDs – is used to create an Active Directory package publication object (PackageRegistration) in the ClassStore container. This is located in the *GPO GUID*\User\Class Store\Packages subcontainer of the Group Policy object (GPO) in which the package is published and given a newly generated globally unique identifier (GUID) for a common name (CN). This published package then becomes available to all recipients of the advertised application.

The Group Policy Software installation client-side extension (CSE) also creates an Application Advertising script in the Sysvol folder that is derived from the same advertising information. This has the same name as the related PackageRegistration object with a file name extension of .aas. In this case, more of the advertising information – such as Start menu shortcuts created by the application, which are omitted from the PackageRegistration object – is used in the creation of the .aas file. This file is stored in *GPO GUID*\User\Applications.

Published applications do not cause any change on the client computer during normal GPO processing, so no information is downloaded from Active Directory or Sysvol. However, the applications published to the user can be viewed in Add or Remove Programs. The user must perform some action to cause the published application to be installed. The simplest case is that the user chooses to install the application from Add or Remove Programs.

The Group Policy Software installation extension also supports a more sophisticated method for installing applications. If a user attempts to open a document with a file name extension for which there is no application currently installed, or tries to call a COM object for which there is no local definition in the registry, the Group Policy Software installation extension service searches Active Directory for PackageRegistration objects that support the file name extension or COM CLSID and that have been published to the user by using a GPO. If one is found, the .aas script is downloaded to the client computer and runs to create advertisements on the computer. When the user again attempts to open a document or a COM operation is retried, the advertisement triggers the installation of the application (or at least that part of it required to support the document or COM operation attempted). If more than one suitable PackageRegistration object is found, the one in the highest-priority GPO is selected. Where there is more than one in the same GPO, the one with the latest creation date is selected.

The following figure shows the Group Policy–based software publishing architecture.

**Group Policy–based Software Deployment Processing Architecture and Components (Published Software)**




**Application Assignment**

Assigning is a superset of publishing. The same description for published-application processing applies to assignment except that the application is either fully or partly installed on the client computer during GPO processing.

Assignment differs from publishing in that the Group Policy Software installation extension either advertises the application on the client computer (in the case of a user assignment) or fully installs the application (in the case of a computer assignment). Assigning to a user creates an advertisement of the application on the client computer – this can include application shortcuts on the Start menu and registry entries defining COM and program registrations supported by the application. Unlike publishing, applications can also be assigned to computers. The Group Policy Software installation CSE only processes during startup and logon, never during periodic refresh. By default, it also does not process over a slow link (this behavior can be changed).

The creation of an assigned package is similar to that of a published package. This produces an almost identical PackageRegistration object and .aas file. The only difference is that the PackageFlags property of the PackageRegistration object identifies this as an Assigned rather than a Published application.

Advertising is a function of Windows Installer and works by placing entry points to the application on the client system. Examples of these entry points are Start menu shortcuts, COM Class IDs, Program IDs, and Multipurpose Internet Mail Extensions (MIME) types. These entry points are configured so that, if any of them are called, Windows Installer installs the application. Windows Installer checks to see whether the application or application component that implements the call has been installed. If not, it installs the application. To the user, the advertisements appear to be already installed.

During GPO processing, the Group Policy Software installation CSE queries the directory for available assigned PackageRegistration objects for applied GPOs that have changed or are new. Any packages found are compared to those already downloaded onto the client computer – the client computer records received packages in the registry and caches .aas files. Any new or changed packages are processed. The corresponding .aas file is read from the Sysvol folder and then cached in %systemroot%\appmgmt. This folder also contains AppMgmt.ini, which stores configuration settings.

The following figure shows Group Policy–based software assignment architecture.

**Group Policy–based Software Deployment Processing Architecture and Components (Assigned Software)**




**Software Distribution Point Servers**

*Deploying applications* is the process of setting up and managing distribution points or server shares where users have access to the packaged software and can install it on their computers. These servers act as file shares for the .msi packages identified in software installation GPOs.

Two methods can be used for deploying and managing software distribution point servers:

- Universal naming convention (UNC) path to a server share.
- Distributed File System (DFS).

**UNC Paths to Server Shares**

By using UNC paths, a user or application can specify the physical server and share names to gain access to file information. For example: *Server**Share**Path**File_name*. A UNC path can be used to allow direct access to a shared file by mapping to a network drive, where the drive letter denotes *Server**Share*. Administrators can also perform a **net use** command to navigate beyond the redirected drive.

However, as networks grow and as organizations begin using existing storage for new purposes, mapping a single drive letter to individual shares becomes inefficient. Also, despite the fact that users and applications can refer to UNC names directly, the increasing number of places users must go to retrieve data can be overwhelming.

**DFS and Software Distribution Point Servers**

DFS provides fault tolerance for software distribution points by mapping a specific logical name to shared folders on multiple file servers. This way, software remains available for installation, regardless of whether one of the physical servers where the software deployment files reside becomes unavailable. DFS also improves storage scalability because administrators can deploy additional or higher-performance file servers and present the storage on the new computers as new directories in an existing namespace.

When using DFS in combination with Group Policy–based software deployment, organizations benefit from its load-sharing abilities and location-independence. These features simplify management and optimize the installation for users. Instead of allowing all users to install software from a single server, and taxing the server, a DFS namespace can distribute network traffic across multiple servers.

**Group Policy Software Installation Extension Components**

Group Policy–based software deployment uses the Windows Server 2003 features and technologies described in the following table.

**Essential Tools and Components for Group Policy–based Software Deployment**

| Component or Tool | General Description | Combined with Group Policy Software Installation Extension |
|---|---|---|

| | | |
|---|---|---|
| Active Directory | A hierarchical collection of objects including domains, sites, organization units (OUs), users, computers, and printers that allow an organization to manage these resources. | Provides the scope of management mechanism to locate users and computers.<br><br>Stores software deployment information through Group Policy. |
| Group Policy | An administrative tool for defining and controlling the way programs, network resources, and the operating system work for users and computers in an organization. In an Active Directory environment, Group Policy is applied to users or computers on the basis of their membership within sites, domains, or OUs. | Enables administrators to deploy applications in a GPO associated with one or more Active Directory containers, such as sites, domains, or OUs. Use the Group Policy Software Installation extension snap-in to deploy applications. |
| Windows Installer | A service based on an operating system, which provides software installation services using a standard package format. Administrators can use Windows Installer to manage the installation, modification, upgrade, and removal of software applications. | Installs, modifies, upgrades, and removes software applications. |
| Group Policy Software installation extension | An extension of the Group Policy Object Editor snap-in that includes a user interface that allows administrators to deploy and manage software. | Communicates with Active Directory, GPOs, and Windows Installer to assign or publish applications as follows:<br><br>● Assigns software to users. Installs user-assigned applications entirely the first time the user logs on after deployment, or allows users to install certain components or features of an application as needed.<br><br>● Assigns software to computers. Installs an application the next time the computer starts. The application is available for all the users on that computer.<br><br>● Publishes applications for users only: Users can choose to install the software from a list of published applications located in Add or Remove Programs in Control Panel. |
| Appmgmts.dll | The client-side extension for the software installation component of Group Policy. | This is the extension file on the client that configures software installation settings. |
| Group Policy Management Console (GPMC) | A new tool that consists of an MMC snap-in and command-line tools. This tool unifies management of all aspects of Group Policy across an enterprise. GPMC allows administrators to manage all GPOs, Windows Management Instrumentation (WMI) filters, and permissions on a network. | Group Policy Modeling (formerly known as RSoP planning) allows administrators to run hypothetical scenarios to verify software configurations under various sites, domains, and OUs. Provides printable HTML-based reports.<br><br>Group Policy Results(formerly known as RSoP logging)verifies which software applications are properly installed for a specific group of users or computers. It also pinpoints the causes of unintended removal or damage to software. Provides printable HTML-based reports. |
| Add or Remove Programs | A user interface in Control Panel of Windows XP Professional and Windows 2000 Professional. Add or Remove Programs lets users manage software on their own computers. | Lists both published and assigned applications so that users can install, modify, and remove software from their desktop computers. |

| | | |
|---|---|---|
| Command-line and Graphical User Interface (GUI) tools or scripts | These include GPResult.exe, GPOTool.exe, GPUpdate.exe, ReplMon.exe, NetDiag.exe, InstallShield, and the new Group Policy Management Console (GPMC) MMC snap-in. Some are installed by default; others must be installed manually. | Helps manage, optimize, or troubleshoot Group Policy–based software deployment. |

**How Windows Installer Works**

Many software authors develop applications to include native Windows Installer packages. A native Windows Installer package contains a single product, such as Microsoft Office 2000, which can be made up of several features, such as Microsoft Word, Microsoft Excel, and Microsoft PowerPoint. However, software can be configured to install features individually. When the user selects an uninstalled feature, the feature is installed. Each feature (Word, Excel, PowerPoint, and so on) contains components, such as a thesaurus, spelling checker, or an additional user-interface language. When the user selects a feature or component that is not installed, the feature or component is automatically installed. Automatic (or "on-demand") installation of selected features saves network bandwidth during the initial installation. It also gives users only the features and components that they need to do their jobs. Automatic installation can also save space on users' local disk drives. However, this type of installation delays the availability of a feature when the user first selects it.

Windows Installer packages ensure that an accidentally deleted file, such as Winword.exe, is reinstalled the next time the user tries to start Word because Windows Installer detects and reinstalls missing files.

Windows Installer technology uses the following two components to help an organization install and manage software:

- A Windows Installer package (an .msi file), which is a database containing information that describes the installed state of an application. The Windows Installer package performs the installation, modification, and removal of the software.
- An application programming interface (API) that allows applications and management tools to interact with Windows Installer to install or remove additional features of the application after the initial installation is completed.

The managed state of an application includes installation, modification, upgrade, or removal. Windows Installer provides administrators with consistent and reliable methods to customize installations, update and upgrade applications, and resolve configuration problems. The following are some of its advantages:

- Helps manage the state of software during and after installation.
- Defines a standard format for application setup and tracks components such as groups of files, registry entries, shortcuts, and other aspects of the application that must be managed together.
- Detects whether software is installed correctly or whether a program file is missing. It can immediately reinstall the damaged or missing files.
- Repairs applications and ensures that they are installed or removed without overwriting or deleting files that are required by another application.
- Enables administrators to modify an installation by adding or removing features after the installation.
- Enables deployment of 32-bit and 64-bit Windows applications using Windows Installer version 2.0.

**Note**

- Versions of Windows Installer earlier than version 2.0 can install and manage 32-bit Windows Installer packages only on 32-bit operating systems. Windows Installer version 2.0 supports all of the setup functionality that is available in earlier Windows Installer versions.

Software installed by using the Windows Installer uses the following parameters:

```
\>msiexec /I \\servername\share\<software.msi> TRANSFORMS = \\servername\share\
<software.mst> /qb
```

In this example, *Software.msi* is the Windows Installer package, and *Software.mst* is the transform applied at deployment time.

**Transforms for Customizing Windows Installer Packages**
A transform (.mst file) is a collection of specified changes that are applied to a base Windows Installer package file at the time of deployment. Windows Installer packages are modular to simplify deployment. When transforms are applied to an .msi file, Windows Installer can dynamically add or modify data in the installation database to customize the installation of the application.

For example, Office 2000 provides the Office 2000 Custom Installation wizard, which can be used to build transforms. Administrators can create a transform for managing the configuration of Office 2000 that is deployed to users in an organization. Other tools in the Windows Installer software development kit (SDK), or other non-Microsoft tools, can also help to create transforms for Windows Installer packages that do not include their own custom installation tools.

Transforms tailor the installation of an application. Although they are optional, transforms can be used for a variety of purposes including customizing and encapsulating.

**Customizing**
Customizing can involve configuring installations so that a particular set of features from a specified software application, or suite of software applications, is installed locally on the computer.

Transforms can also add new features to an existing application's package. For example, they can add custom Excel templates for an organization's Finance group.

**Encapsulating**
Administrators can encapsulate numerous customizations of a base package that are required by different groups. For example, in organizations where the Finance and Marketing departments require different installations of an application, the base package of the application can be made available to everyone at one software distribution point. Then the appropriate customizations can be distributed separately as transforms to each group of users.

**Reauthoring Applications for Windows Installer**
When administrators reauthor an application, they create applications that adhere to the Windows Installer format. Reauthoring is essentially redeveloping the setup portion of the application to take full advantage of the advanced capabilities of Windows Installer.

There are some authoring tools available to help developers create new Windows Installer packages, but the procedures can be resource-intensive and costly. If an organization determines that the application will play a key role in its future, it is important to weigh the benefits of reauthoring with the costs.

**.Zap Files**
Applications that do not use the .msi file format for the Windows Installer Service can be set up for distribution by creating a text file that has a .zap file name extension. When organizations have many applications that do not contain native Windows Installer packages and plan to discontinue these applications, they often create software installation settings (.zap) files for the installation executable files, such as Setup.exe or Install.exe. Additionally, if an organization uses custom applications that do not have Windows Installer support, but plan to use them in the long term, .zap files might be the only choice. Users of .zap files do not benefit from the capabilities of Windows Installer. The Group Policy Software installation extension recognizes .zap files that wrap 32-bit or 64-bit Setup.exe files into a .zap file format. Applications published in this format are available for users to install by using Add or Remove Programs.

Because these applications do not use Windows Installer setup programs, they do not do the following:

- Use elevated permissions for installation. While applications that are installed by using .zap files run their original setup programs, they do not run with the elevated permissions that Windows Installer packages have. If installing the application requires administrative permissions, only users who have those permissions can install it.
- Install a feature on the first use of the feature.
- Roll back an unsuccessful operation, such as install, modify, repair, or removal.
- Detect a broken state and automatically repair it.
- Implement customized installations (also known as transforms).

**64-Bit Application Packaging**
In Windows Server 2003, the Group Policy Software installation extension and Windows Installer version 2.0 continue to

support and protect the investment that organizations have made in 32-bit applications. Additionally, Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, introduce support for 64-bit application installation.

Windows Installer version 2.0 can install three types of Windows Installer packages on a computer running a 64-bit operating system:

- 32-bit packages that contain only 32-bit components
- 64-bit packages containing some 64-bit components and some 32-bit components
- 64-bit packages containing only 64-bit components

**Note**

- 64-bit applications cannot be published or installed on 32-bit operating systems. Group Policy will attempt to install an incorrectly configured 64-bit package for 32-bit clients, but then removes the package during logon. It then installs the package again the next time a user logs on, and so on.

The Group Policy Software installation extension can be configured to allow or disallow installation of 32-bit applications to 64-bit clients. More .zap applications fail on 64-bit clients than on 32-bit clients. Unless the default behavior is changed, 32-bit .zap applications are deployed so that they are not listed in Add or Remove Programs on 64-bit clients.

**Repackaging Applications for Windows Installer**
Administrators repackage an .msi package for Windows Installer when they cannot reauthor it. Repackaging an application for Windows Installer involves making an image of a fresh computer (including the registry settings, files, and system settings), installing the software, and then taking a post-installation image of the computer. The repackaging software detects the difference between the two images, and then creates the necessary instructions to reproduce the installation. If any registry changes, files changes, or system setting changes occur during the capture process, they are included in the installation. Repackaging is useful when administrators do not have control over dynamic-link library (DLL) files, source code, and registry entries, or for applications about which they do not have in-depth knowledge.

Repackaging requires a thorough knowledge of the application's installation program and of the Windows Installer configuration on the Windows platform. Success with repackaging is affected by the state of the computer on which an administrator performs the repackaging. The computer should have only the operating system and operating system service packs installed before running the repackaging software. Because of this limitation, and other issues, repackaging is not recommended.

Repackaging is not a function or a feature of Windows Installer. As with reauthoring applications, several vendors provide tools to enable administrators to repackage applications for a variety of needs. The same vendors who provide tools to reauthor applications can also help repackage them.

During repackaging, an application, the existing components, such as DLLs, .ini files, registry settings, and shortcuts, are replaced, and then the administrator creates a path for Windows Installer to find these items at installation time.

For more information about repackaging applications in the .msi format, see the Windows Installer SDK for detailed documentation about the Windows Installer package format.

**Group Policy Software Installation Extension Processes and Interactions**
Because software can be published for users, assigned to users, or assigned to computers, administrators can establish a workable combination of those three options to meet an organization's software management goals.

**Publishing Software for Users**
The benefit of publishing software, instead of assigning it, is that it requires less management when change occurs in the Active Directory structure. Typically, applications that are nonessential for the users are published. When software is published for a user, it does not initially appear to be installed on the computer. There is no Windows Installer advertisement information about the software on the computer in the registry, on the desktop, or on Start menu as a shortcut. On an as-needed basis, the user installs the published software by using Add or Remove Programs in Control Panel. An exception is when an application is published in a new GPO, when the GPO must be simultaneously linked to the users in a site, domain, or OU. If a GPO is linked and the software is deployed at the same time, the GPO must be refreshed by an administrator or by a user logging on before the application appears in Add or Remove Programs.

Users can also install a published application by opening an associated document if the application is deployed to do that (if Auto-Installis selected).

The user can remove published software, and then later choose to reinstall it, by using Add or Remove Programs.

The following figure shows how Group Policy–based software publishing works.

**Group Policy–based Software Publishing for Users**




**Assigning Software**

There are three methods for assigning software: assign to users on-demand, assign to users, or assign to computers.

**Note**

- By default, Group Policy allows administrators to configure a user-assigned application that has a staggered, on-demand installation. Windows Server 2003 allows administrators to turn off the default installation and install the entire application at once. This is the same behavior as computer-assigned application installation.

**Assigning Software to be Available on Demand**

After a software package is assigned to users in a site, domain, or OU, the software is advertised on user desktops. The application becomes available to the user the next time the user logs on (if the application's GPO applies to that user). The application is fully installed by the user from the Start menu, from Add or Remove Programs, from a desktop shortcut, or by opening a document (on demand) that has a file name extension that is associated with the application.

The user can remove the software, and then later choose to reinstall it as they did previously. Group Policy ensures that assigned applications that are available on-demand are available, regardless of whether users remove them, and that the applications are available again the next time the user logs on or starts the computer.

**Assigning Software to Users**

After a software package is assigned to users in a site, domain, or OU, administrators can use the **Install this application at logon** option to install the whole application the next time the computer starts, or after the user logs off and then logs on again. Alternatively, for users of Windows XP Professional, Logon Optimization can delay the installation of software to the second logon to prevent installation during the first logon. The application is also immediately available in Add or Remove Programs.

The user can remove the software, and then later choose to reinstall it.

**Note**

- Some published applications might not appear in Add or Remove Programsin a domain that has multiple domain controllers until the changes have replicated to all domain controllers in the domain.

The following figure shows the client-side perspective of assigning Group Policy–based software deployment.

**Group Policy–based Software Assignment to Users**




**Assigning Software to Computers**

After assigning a software package to computers in a site, domain, or OU, the software is installed the next time the computer restarts or the user logs on. For users of Windows XP Professional, Logon Optimization can delay the installation of software to the second logon to prevent installation during the first logon.

Only the local or network administrator can remove the software, though a user can repair the software.

The following figure shows the client-side perspective of publishing Group Policy–based software deployment.

**Group Policy–based Software Assignment to Computers**




**Patches for Installed Applications**

Installing a patch updates a previously installed application or applications. A patch package (.msp file) is obtained from the software manufacturer or from the internal developers of the original program. Installing patches allows existing applications to be updated without removing the product. This preserves the customizations of the installation and can lower the cost of making the change. The patch might change only a few bytes of a single application file. It is more efficient to distribute those few bytes than to remove and redeploy the whole product.

**Note**

- A patch might change all of the files and registry keys in a product.

After a patch is applied, it is cached on the user's computer, this allows the user to:

- Perform any installation on demand.
- Reinstall the application.
- Repair the application.
- Remove the application.

A patch package (.msp file) does not include a database as a regular installation package does. Instead, it contains, at a minimum, one database transform that adds patching information to the database of its target Windows Installer .msi package. For an application to perform maintenance operations, such as adding, removing, or repairing the installation, the package codes for both the installed application and the source must match.

**Software Upgrades**

During an upgrade where the ProductVersion or ProductCode property of the package has changed, Windows Installer searches for upgradeable products by querying the Upgrade table of the upgrade package. The newer version of the product is installed. If Windows Installer finds an older version of the product, it removes the old version. The author of the application's setup can choose to remove the old version, and then install the new version. The maintenance mode and removal do not trigger these actions because "remove existing versions of an application" is now automatic.

The Windows Installer package for a repackaged application does not have declared upgrade relationships. Administrators must manually create upgrade relationships. The Group Policy Software installation extension can be used to create upgrade relationships between the new package and the packages that the application replaces. This includes making a formal upgrade relationship between two similar products from completely different vendor's. A vendors application can replace another's, or an administrator can upgrade a repackaged application.

# Group Policy Software Installation Extension Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Software Installation Extension Tools and Settings**

**In this section**


- Group Policy Software Installation Extension Tools
- Group Policy Software Installation Extension Group Policy Settings and Registry Entries
- Related Information


This section summarizes the tools and settings associated with the Group Policy Software installation extension.

**Group Policy Software Installation Extension Tools**

The following tools are associated with the Group Policy Software installation extension:

- InstallShield
- Application Experience Lookup Service
- Windows Installer

**InstallShield**

**Category**

InstallShield products do not ship with Windows.

**Version compatibility**

InstallShield Corporation creates several tools for building Windows Installer packages that work on all versions of Windows managed by Group Policy. Make sure you use the correct version for the systems that you manage.

InstallShield creates popular tools for developing Windows Installer .msi packages.

**Application Experience Lookup Service**

**Category**

The Application Experience Lookup Service is a new service included in Windows Server 2003 with Service Pack 1 (SP1).

**Version compatibility**

This service is part of an infrastructure that provides a way to apply fixes to applications in order to ensure that they run on newly released Windows operating systems or service packs.

This service needs to be running for the application fixes to work. There are no entry points to this service for customizations and it is for operating system internal use only. There is no out-of-the box communication in the service. This service does not use any Active Directory, network, or internet resources.

The functionality of the service can be disabled though Group Policy settings for application compatibility. When this setting is disabled, the service will continue to run, but there will be no calls made to the service. The service itself cannot be stopped or disabled.

**Windows Installer**

**Category**

Windows Installer ships with Microsoft Windows Server 2003 family, Windows XP, Windows 2000, and Windows Millennium Edition (Windows Me). The installer is also provided as a service pack for Microsoft Windows NT version 4.0, Windows 98, and Windows 95.

**Version compatibility**

Windows Installer version 2.0 adds advanced features and requires Windows NT 4.0 with Service Pack 6 or later,

Windows 2000, Windows Me, or Windows XP. Earlier Windows Installer versions require Windows NT 4.0 with Service Pack 3 or later, Windows 2000, or Windows Me.

Windows Installer supports advertisement of applications and features according to operating system. The following table outlines Windows Installer advertisement support on different operating systems.

**Group Policy Software Installation Advertisement Support on Different Operating Systems**

| Operating System | Advertisement Support |
|---|---|
| Windows 2000<br><br>Windows XP | • Shortcuts and their icons.<br><br>**NOTE**: AppId and Typelib information is only written when an advertised component is installed.<br><br>• Extensions and their icons specified in the ProgId table.<br><br>• Shell and command Verbs registered underneath the ProgId key.<br><br>• CLSID contexts and InProcHandler.<br><br>• Install-On-Demand through OLE is only available programmatically through CoCreateInstance (C or C++) and CreateObject or GetObject (Visual Basic). |
| Windows 98<br><br>Windows Me | All of the above except CLSID, which is only written when installing an advertised component.<br><br>Shell and MIME support. |
| Microsoft Windows 95<br><br>with IE4.01 Service Pack 1 installed<br><br>with Windows Desktop Update installed<br><br>(shell32.dll of 4.72.3110.0 or newer)<br><br>Windows NT 4.0<br><br>with IE4.01 Service Pack 1 installed<br><br>with Windows Desktop Update installed<br><br>(shell32.dll of 4.72.3110.0 or newer) | All of the above except CLSID, which is only written when installing an advertised component.<br><br>Shell and MIME support. |
| Windows 95<br><br>Windows NT 4.0<br><br>(shell32.dll older than 4.72.3110.0) | Advertisement is not supported by these platforms. |

On Windows 98 or Windows 95 with the updated shell32.dll, advertised shortcuts do not work until the computer is restarted. This only affects the first product that installs the package for Windows Installer. The installation of the product might not require a restart, but any advertised shortcuts do not work until the computer has been restarted. Advertised shortcuts of subsequent installations work without a restart. Conditional statements can check the **ShellAdvtSupport** property and **Version9X** property.

Windows Installer is a Windows operating system-based service that reduces the total cost of ownership by allowing

administrators to manage the installation, modification, upgrade, and removal of software applications using a standard package format.

Windows Installer includes the operating system-based service, a package format, and an application-programming interface (API) that allows both the operating system and applications to interact with the service to install, modify, or repair the software.

**Group Policy Software Installation Extension Group Policy Settings and Registry Entries**
In addition to setting configuration options for the application in **Properties**, you can use several Group Policy settings to control the behavior of Windows Installer and the Add or Remove Programs feature of Windows.

The following tables list the Group Policy settings and associated registry keys that control Windows Installer and Add or Remove Programs. The settings are all part of the System.adm file.

The following table lists the Group Policy Machine settings and associated registry keys that control Windows Installer. These settings are found in these locations:

- Group Policy Location: MACHINE\Administrative Templates\Windows Components\Windows Installer
- Registry Location: HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows

**Group Policy Settings for Windows Installer (Machine)**

| Setting | Description |
|---|---|
| Disable Windows Installer | Disables or restricts the use of Windows Installer. This setting can prevent users from installing software on their systems or permit users to install only those programs offered by a system administrator. |
| Always install with elevated privileges | Directs Windows Installer to use system permissions when it installs any program on the system. This setting extends elevated privileges to all programs. |
| Prohibit rollback | Prohibits Windows Installer from generating and saving the files it needs to reverse an interrupted or unsuccessful installation. This setting prevents Windows Installer from recording the original state of the system and sequence of changes it makes during installation. It also prevents Windows Installer from retaining files it intends to delete later. As a result, Windows Installer cannot restore the computer to its original state if the installation does not complete. |
| Remove browse dialog box for new source | Prevents users from searching for installation files when they add features or components to an installed program. This setting disables the **Browse** button beside the Use feature from list in the Windows Installer dialog box. |
| Prohibit patching | Prevents users from using Windows Installer to install patches. |
| Disable IE security prompt for Windows Installer scripts | Allows Web-based programs to install software on the computer without notifying the user. |
| Enable user control over installs | Permits users to change installation options that typically are available only to system administrators. This setting bypasses some of the security features of Windows Installer. It permits installations to complete that otherwise would be halted due to a security violation. |

| | |
|---|---|
| Enable user to browse for source while elevated | Allows users to search for installation files during privileged installations. This setting enables the **Browse** button in the Use feature from dialog box. As a result, users can search for installation files, even when the installation program is running with elevated system privileges. |
| Enable user to use media source while elevated | Allows users to install programs from removable media, such as floppy disks and CD-ROMs, during privileged installations. This setting permits all users to install programs from removable media, even when the installation program is running with elevated system privileges. |
| Enable user to patch elevated products | Allows users to upgrade programs during privileged installations. This setting permits all users to install patches, even when the installation program is running with elevated system privileges. |
| Allow admin to install from Terminal Services session | Allows Terminal Services administrators to install and configure programs remotely. |
| Cache transforms in secure location on workstation | Saves copies of transform files in a secure location on the local computer. |
| Logging | Specifies the types of events that Windows Installer records in its transaction log for each installation. The log, **Msi.log**, appears in the **Temp** directory of the system volume. |
| Prohibit User Installs | Allows you to configure user installs. This setting is useful in environments where the administrator only wants per-computer applications installed, such as on a kiosk or a Windows Terminal Server. |
| Turn off creation of System Restore Checkpoints | If you disable this setting or do not configure it, the Windows Installer automatically creates a System Restore checkpoint each time an application is installed. |

The following table lists the Group Policy User settings and associated registry keys that control Windows Installer. These settings are found in these locations:

- Group Policy Location: USER\Administrative Templates\Windows Components\Windows Installer
- Registry Location: HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows

**Group Policy Settings for Windows Installer (User)**

| Setting | Description |
|---|---|
| Always install with elevated privileges | Directs Windows Installer to use system permissions when it installs any program on the system. This setting extends elevated privileges to all programs. |
| Search order | Specifies the order in which Windows Installer searches for installation files. |

| | |
|---|---|
| Prohibit rollback | Prohibits Windows Installer from generating and saving the files it needs to reverse an interrupted or unsuccessful installation. This setting prevents Windows Installer from recording the original state of the system and sequence of changes it makes during installation. It also prevents Windows Installer from retaining files it intends to delete later. As a result, Windows Installer cannot restore the computer to its original state if the installation does not complete. |
| Prevent removable media source for any install | Prevents users from installing programs from removable media. |

The following table lists the Group Policy User settings and associated registry keys that control Add or Remove Programs. These settings are found in these locations:

- Group Policy Location: USER\Administrative Templates\Control Panel\Add or Remove Programs
- Registry Location: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies\

**Group Policy Settings for Add or Remove Programs (User)**

| Setting | Description |
|---|---|
| Remove **Add or Remove Programs** | Prevents users from using **Add or Remove Programs**. This setting removes **Add or Remove Programs** from Control Panel and removes the **Add or Remove Programs** item from menus. |
| Hide Change or Remove Programs page | Removes the **Change or Remove Programs** button from the **Add or Remove Programs** bar. |
| Hide Add New Programs page | Removes the **Add New Programs** button from the **Add or Remove Programs** bar. |
| Hide Add/Remove Windows Components page | Removes the **Add/Remove Windows Components** button from the **Add or Remove Programs** bar. |
| Hide the Set Program Access and Defaults page | Removes the **Set Program Access and Defaults** button from the **Add or Remove Programs** bar. |
| Hide the Add a program from CD-ROM or floppy disk option | Removes the **Add a program from CD-ROM or floppy disk section** from the **Add New Programs** page. |
| Hide the Add programs from Microsoft option | Removes the **Add programs from Microsoft** section from the **Add New Programs** page. |
| Hide the Add programs from your network option | Prevents users from viewing or installing published programs. This setting removes the **Add programs from your network** section from **the Add New Programs** page. |
| Go directly to Components Wizard | Prevents users from using **Add or Remove Programs** to configure installed services. This setting removes the **Set up services** section of the **Add/Remove Windows Components** page. |
| Remove Support Information | Removes links to the **Support Info** dialog box from programs on the **Change or Remove Programs** page. |
| Specify default category for **Add New Programs** | Specifies the category of programs that appears when users open the **Add New Programs** page. If you enable this setting, only the programs in the category you specify are displayed when the **Add New Programs** page opens. |

The following table lists the Group Policy Machine settings and associated registry keys for application compatibility. These settings are found in these locations:

- Group Policy Location: MACHINE\Administrative Templates\Windows Components\Application Compatibility

- Registry Location: HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows

**Group Policy Settings for Application Compatibility (Machine)**

| Setting | Description |
|---|---|
| Turn Off Application Compatibility Engine | Controls the state of the application compatibility engine in the system. |
| Turn Off Program Compatibility Wizard | Controls the state of the Program Compatibility Wizard. When enabled, this setting disables the start page of the wizard in Help and Support, and in the Start menu. |
| Remove Program Compatibility Property Page | Controls the visibility of the Program Compatibility property page shell extension. |
| Turn On Application Help Log Events | Blocks known incompatible applications and displays a dialog to the end-user regarding the problem. |
| Prevent access to 16-bit applications | Specifies whether to prevent the MS-DOS subsystem (ntvdm.exe) from running on this computer. This setting affects the launching of 16-bit applications in the operating system. |

The following table lists the Group Policy User settings and associated registry keys for application compatibility. These settings are found in these locations:

- Group Policy Location: USER\Administrative Templates\Windows Components\Application Compatibility

- Registry Location: HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows

**Group Policy Setting for Application Compatibility (User)**

| Setting | Description |
|---|---|
| Prevent access to 16-bit applications | Specifies whether to prevent the MS-DOS subsystem (ntvdm.exe) from running for all users. This setting affects the launching of 16-bit applications in the operating system. |

**Related Information**

The following resources contain additional information that is relevant to this section.

- How Group Policy Software Installation Extension Works

- Core Group Policy Technical Reference

- Group Policy Administrative Tools

- Group Policy Object Editor Tools and Settings

- Group Policy Management Console Technical Reference

- "Group Policy Settings Reference" in the Tools and Settings Collection

- "Resource Kit Tools" in the Tools and Settings Collection

- "Microsoft Platform SDK" on MSDN for more information about Windows Installer.

# Security Settings Extension Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In an Active Directory directory service environment, you can use the Security Settings snap-in to modify the security policy assigned to a Group Policy object that is linked to an organizational unit, domain, or site. The Security Settings snap-in is an extension of the Group Policy Object Editor Microsoft Management Console (MMC) snap-in. You can use the Security Settings snap-in from any computer that is joined to a domain. On a local computer, you can use the Local Security Policy snap-in to modify account and local policies, public key policies, and IP security policies.

**In this subject**


- What Is Security Settings Extension?
- How Security Settings Extension Works
- Security Settings Extension Tools and Settings

# What Is Security Settings Extension?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Security Settings Extension Overview
- Securing Distributed Systems
- Related Information


Security settings policies are rules that administrators configure on a computer or multiple computers for the purpose of protecting resources on a computer or network. The Security Settings extension of the Group Policy Object Editor snap-in allows you to define security configurations as part of a Group Policy object (GPO). The GPOs are linked to Active Directory containers such as sites, domains, or organizational units, and enable administrators to manage security settings for multiple computers from any computer joined to the domain. Security settings policies are used as part of your overall security implementation to help secure domain controllers, servers, clients and other resources in your organization.

Security settings can control:

- User authentication to a network or computer.
- The resources that users are permitted to access.
- Whether to record a user's or group's actions in the Event log.
- Membership in a group.

**Security Settings Extension Overview**

Administrators can use the Security Settings extension to specify the security configuration of a Group Policy object. This in turn affects all the computers in the Active Directory container to which the Group Policy object is linked. All security policies are computer-based policies.

To manage security configurations for multiple computers, you can use one of the following options:

- Edit specific security settings in a GPO.
- Use the Security Templates snap-in to create a security template that contains the security policies you want to apply, and then import the security template to a Group Policy object. A security template is a file that represents a security configuration, and can be imported to a GPO, or applied to a local computer, or can be used to analyze security.

The Security Settings extension of Group Policy Object Editor includes the following types of security policies:

- **Accounts Policies.** These polices are defined on computers; they affect how user accounts can interact with the computer or domain. Accounts policies include the following types of policies:
    - **Password Policy.** These policies determine settings for passwords, such as enforcement and lifetimes. Password policies are used for domain accounts.
    - **Account Lockout Policy.** These policies determine the conditions and length of time that an account will be locked out of the system. Account lockout policies are used for domain or local user accounts.
    - **Kerberos Policy.** These policies are used for domain user accounts; they determine Kerberos-related settings, such as ticket lifetimes and enforcement.

- **Local Policies.** These policies apply to a computer and include the following types of policy settings:
    - **Audit Policy.** Specifies whether to log security events into the Security log on the computer, and specifies what types of security events to log (success, failure, or both).
    - **User Rights Assignment.** Specify the users or groups that have logon rights or privileges on a computer.
    - **Security Options.** Specify security settings for the computer such as Administrator and Guest Account names, access to floppy disk drive and CD-ROM drive, installation of drivers, logon prompts and so on.
- **Event Log.** Specifies settings for the application, security, and system logs.
- **Restricted Groups.** Controls membership in a security-sensitive (restricted) group, and specifies the groups to which a restricted group should belong. Restricted Groups policy states that only the members that you have added can belong to that group.
- **System Services.** Specifies startup options for system services, and define access permissions.
- **Registry.** Specifies access permissions (on discretionary access control lists (DACLs) and audit settings (on system access control lists (SACLs)) for registry keys.
- **File System.** Specifies access permissions and audit settings for file system objects.

**Securing Distributed Systems**

Organizations are faced with the complex, costly, and time-consuming task of securing and managing network systems, and keeping their desktops and servers protected with up-to-date security updates.

Administrators' tasks include:

- Keeping desktops and servers up-to-date with the latest security patches.
- Ensuring that the corporate security policies are enforced across desktops and servers.
- Monitoring systems for potential security compromises.

Organizations need efficient ways to maintain network security and manage updates, while reducing total costs for security management.

**Policy-Based Security Settings Management**

Windows Server 2003, Windows 2000, and Windows XP provide an integrated policy-based management infrastructure to help administrators manage and enforce their security policies.

Windows Server 2003 and Windows 2000, through Group Policy and Active Directory, enable IT administrators to define and apply security settings policies to users, groups, and network servers and clients. A group of servers with the same functionality can be created (for example, a Microsoft Internet Information Services (IIS) server), and then Group Policy objects can be used to apply common security settings to the group. If more servers are added to this group later, many of the common security settings are automatically applied, reducing deployment and administrative labor.

Using security settings policies simplifies and centralizes security configuration and management for computers running Windows Server 2003. Security policies can help reduce administrative tasks by automating some processes for applying security to servers. Computers running Windows Server 2003 that are members of a domain periodically access Active Directory; if they detect that a new policy exists or that an existing one has been changed, they automatically download the policy and apply it locally.

**Common Scenarios for Using Security Settings Policies**

Security Settings policies are used to manage the following aspects of security: accounts policy, local policy, user rights assignment, registry values, file and registry Access Control Lists (ACLs), and service startup modes.

As part of your security strategy, you can create GPOs with security settings policies configured specifically for the various computer roles in your organization, such as domain controllers, file servers, member servers, clients, and so on.

You can create an organizational unit (OU) structure that groups computers according to their roles. Using OUs is the best method for separating specific security requirements for the different computer roles in your network. This approach also allows you to apply customized security templates to each class of server or computer. After creating the security templates, you create a new Group Policy object (GPO) for each of the OUs, and then import the security template (.inf file) into the new GPO.

Importing a security template to a Group Policy object ensures that any accounts to which the GPO is applied automatically receive the template's security settings when the Group Policy settings are refreshed. On a workstation or server, the security settings are refreshed every 90 minutes, and on a domain controller, this process occurs every 5 minutes if changes have occurred in any of the GPO settings that apply. The settings are also refreshed every 16 hours, whether or not any changes have occurred.

By using Group Policy-based security configurations in conjunction with the delegation of administration, you can ensure that specific security settings, rights, and behavior are applied to all servers and computers within an OU. This approach makes it simple to update a number of servers with any additional changes required in the future.

**Dependencies on Other Operating System Components**
For computers that are members of a Windows Server 2003 or Windows 2000 Server domain, Security Settings policies depend on the following components:

**Active Directory.**
The Windows-based directory service, Active Directory, stores information about objects on a network and makes this information available to administrators and users. By using Active Directory, you can view and manage network objects on the network from a single location, and users can access permitted network resources by using a single logon.

**Group Policy.**
The infrastructure within Active Directory that enables directory-based configuration management of user and computer settings on computers running Windows Server 2003, Windows 2000, and Windows XP Professional operating systems. By using Group Policy, you can define configurations for groups of users and computers, including policy settings for Windows Server 2003 registry-based policies, software installation, scripts, folder redirection, Remote Installation Services, Internet Explorer maintenance, and security.

**Domain Name System (DNS).**
A hierarchical naming system used for locating domain names on the Internet and on private TCP/IP networks. DNS provides a service for mapping DNS domain names to IP addresses, and IP addresses to domain names. This allows users, computers, and applications to query the DNS to specify remote systems by fully qualified domain names rather than by IP addresses.

**Winlogon.**
A component of the Windows operating system that provides interactive logon support. Winlogon is designed around an interactive logon model that consists of three components: the Winlogon executable, a Graphical Identification and Authentication dynamic-link library (DLL), and any number of network providers.

**Setup.**
Security configuration interacts with the operating system setup process during a clean installation or upgrade of Windows Server 2003 or Windows 2000 Server.

**Security Accounts Manager (SAM).**
A Windows service used during the logon process. SAM maintains user account information, including groups to which a user belongs.

**Local Security Authority (LSA).**
A protected subsystem that authenticates and logs users onto the local system. LSA also maintains information about all aspects of local security on a system, collectively known as the Local Security Policy of the system.

**Windows Management Instrumentation (WMI).**
A component of the Microsoft Windows operating system, WMI is the Microsoft implementation of Web-Based Enterprise Management (WBEM), which is an industry initiative to develop a standard technology for accessing management information in an enterprise environment. WMI provides access to information about objects in a managed environment. Through WMI and the WMI application programming interface (API), applications can query for and make changes to static information in the Common Information Model (CIM) repository and dynamic information maintained by the various types of providers.

**Resultant Set of Policy (RSoP).**
An enhanced Group Policy infrastructure that uses WMI in order to make it easier to plan and debug policy settings. RSoP provides public methods that expose what an extension to Group Policy would do in a what-if situation, and what the extension has done in an actual situation. This allows administrators to determine easily the combination of policy settings that apply to, or will apply to, a user or computer.

**Service Control Manager (SCM).**
Used for configuration of service startup modes and security.

**Registry.**
Used for configuration of registry values and security.

**File system.**
Used for configuration of security.

**File system conversions.**
Security is set when an administrator converts a file system from FAT to NTFS.

**Microsoft Management Console (MMC).**
The user interface for the Security Settings tool is an extension of the Group Policy Object Editor MMC snap-in.

**Security Settings Policies and Group Policy**
The Security Settings extension of Group Policy Object Editor is part of the Windows Server 2003 Security Configuration Manager tool set. The following components are associated with Security Settings: a configuration engine; an analysis engine; a template and database interface layer; setup integration logic; and the secedit.exe command line tool. The security configuration engine runs on computers running Windows Server 2003, Windows 2000 and Windows XP and is responsible for handling security configuration editor-related security requests for the system on which it runs. The analysis engine analyzes system security for a given configuration and saves the result. The template and database interface layer handles reading and writing requests from and to the template or database (for internal storage). The Security Settings extension of Group Policy Object Editor handles Group Policy from a domain-based or local computer. The security configuration logic integrates with Windows Server 2003 and Windows 2000 setup and manages system security for a clean installation or upgrade of Windows Server 2003 and Windows 2000 systems. Security information is stored in templates (.inf files) or in the Secedit.sdb database.

The following figure shows Security Settings and related components, including:

- Scesrv.dll. Provides the core security engine functionality.
- Scecli.dll. Provides the client-side interfaces to the security configuration engine and provides data to Resultant Set of Policies (RsoP).
- Wsecedit.dll. The Security Settings extension of Group Policy Object Editor. Scecli.dll is loaded into Wsecedit.dll to support the Security Settings user interface.
- Gpedit.dll. The Group Policy Object Editor MMC snap-in.

**Security Settings Policies and Related Components**




**Related Information**

The following resources contain additional information that is relevant to this section.

- "What Is Core Group Policy?" in the Group Policy Technical Reference collection
- "Security Policy Settings" in the Security collection.

# How Security Settings Extension Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Security Settings Extension Architecture
- Security Settings Policy Processes and Interactions
- Related Information


The Security Settings extension of the Group Policy Object Editor MMC snap-in allows you to specify security configurations as part of Group Policy objects, which can be linked to Active Directory sites, domains, or organizational units.

An Active Directory-based deployment of security settings policy requires deployment of at least Windows 2000 Server, and Windows 2000 Professional or Windows XP clients. However, to gain full advantage of the latest security settings policies, you need:

- Windows Server 2003 with Active Directory installed and DNS dynamic update protocol properly configured.
- Windows XP client computers.
- Group Policy Object Editor, which you can access by using Group Policy Management Console (GPMC).
- The Security Settings extension of Group Policy Object Editor.

For computers that are not participating in a domain, you can use the Local Security Policy snap-in.

This section discusses how the Security Settings extension of Group Policy Object Editor works. It presents information about how security settings are stored and how security settings policy is processed.

**Security Settings Extension Architecture**

The Security Settings extension of Group Policy Object Editor is part of the Security Configuration Manager tools, as shown in the following diagram:

**Security Settings Architecture**




The security settings configuration and analysis tools include a security configuration engine, which provides local computer (non-domain member) and Group Policy-based configuration and analysis of security settings policies. The security configuration engine also supports the creation of security policy files. The primary components of the security configuration engine are:

**Scecli.dll**
Provides client side interfaces to the security configuration engine and does Resultant Set of Policies (RsoP) logging during policy propagation.

**Scesrv.dll**
Provides core security engine functionality including support for import, configure, analyze, and policy propagation operations.

Communications between components of the Security Settings extension occurs by using the following methods:

- Component Object Model (COM) calls
- Local remote procedure call (LRPC)
- Lightweight Directory Access Protocol (LDAP)
- Active Directory Service Interfaces (ADSI)
- Server Message Block (SMB)
- Win32 APIs
- Windows Management Instrumentation (WMI) calls

The following table describes the Security Settings-related components:

**Components and Description**

| Component | Description |
|---|---|

| | |
|---|---|
| Scesrv.dll | This .dll is hosted in Services.exe and runs under local system context. Scesrv.dll provides core Security Configuration Manager functionality, such as import, configure, analyze, and policy propagation.<br><br>Scesrv.dll performs configuration and analysis of various security-related system parameters by calling corresponding system APIs, including LSA, SAM, and the registry.<br><br>Scesrv.dll exposes APIs such as import, export, configure, and analyze. It checks that the request is made over LRPC (Windows XP) and fails the call if it is not.<br><br>On domain controllers, Scesrv.dll receives notifications of changes made to SAM and the LSA that need to be synchronized across domain controllers. Scesrv.dll incorporates those changes into the Default Domain Controller Policy GPO by using in-process Scecli.dll template modification APIs.<br><br>Scesrv.dll also performs configuration and analysis operations. |
| Scecli.dll | This is the client-side interface or wrapper to scesrv.dll. Scecli.dll is loaded into Wsecedit.dll to support MMC snap-in UIs. It is used by setup to configure default system security and security of files, registry keys, and services installed by the Setup API .inf files.<br><br>The command line version of the security configuration and analysis user interfaces, Secedit.exe, uses Scecli.dll.<br><br>Scecli.dll implements the client side extension for Group Policy.<br><br>Scesrv.dll uses Scecli.dll to download applicable Group Policy files from Sysvol in order to apply Group Policy security settings to the local computer.<br><br>Scecli.dll logs application of security policy into WMI (RSoP).<br><br>Scesrv.dll policy filter uses Scecli.dll to update Default Domain Controller Policy GPO when changes are made to SAM and LSA. |
| Wsecedit.dll | The Security Settings extension of the Group Policy Object Editor snap-in. You use this tool to configure security settings in a Group Policy object for a site, domain, or organizational unit. You can also use SecuritySettings to import security templates to a GPO. |
| Secedit.sdb | This is a permanent system database used for policy propagation including a table of persistent settings (sometimes referred to as a *tattoo* table) for rollback purposes. |
| User databases | A user database is any database other than the system database created by administrators for the purposes of configuration or analysis of security. |
| .Inf Templates | These are text files that contain declarative security settings. They are loaded into a database before configuration or analysis. Group Policy security policies are stored in .Inf files on the Sysvol folder of domain controllers, where they are downloaded (by using file copy) and merged into the system database during policy propagation. |

**Security Settings Policy Processes and Interactions**
Before learning about how Security Settings policy is processed, it is useful to review some basic concepts related to Group Policy processing.

**Group Policy Processing**
When a computer starts and a user logs on, computer policy and user policy are applied according to the following sequence:

- The network starts. Remote Procedure Call System Service (RPCSS) and Multiple Universal Naming Convention Provider (MUP) start.
- An ordered list of Group Policy objects is obtained for the computer. The list might depend on these factors:
    - Whether the computer is part of a domain and, therefore, subject to Group Policy through Active Directory.
    - The location of the computer in Active Directory.
    - Whether the list of Group Policy objects has changed. If the list of Group Policy objects has not changed, no

processing is done.

- Computer policy is applied. These are the settings under Computer Configuration from the gathered list. This is a synchronous process by default and occurs in the following order: local, site, domain, organizational unit, child organizational unit, and so on. No user interface appears while computer policies are processed.

- Startup scripts run. This is hidden and synchronous by default; each script must complete or time out before the next one starts. The default time-out is 600 seconds. You can use several policy settings to modify this behavior.

- The user presses CTRL+ALT+DEL to log on.

- After the user is validated, the user profile loads; it is governed by the policy settings that are in effect.

- An ordered list of Group Policy objects is obtained for the user. The list might depend on these factors:

    - Whether the user is part of a domain and, therefore, subject to Group Policy through Active Directory.

    - Whether loopback policy processing is enabled, and if so, the state (Merge or Replace) of the loopback policy setting.

    - The location of the user in Active Directory.

    - Whether the list of Group Policy objects has changed. If the list of Group Policy objects has not changed, no processing is done.

- User policy is applied. These are the settings under User Configuration from the gathered list. This is synchronous by default and in the following order: local, site, domain, organizational unit, child organizational unit, and so on. No user interface appears while user policies are processed.

- Logon scripts run. Unlike Windows NT 4.0 scripts, Group Policy-based logon scripts are hidden and asynchronous by default. The user object script (which, as in Windows NT 4.0, is run in a normal window) runs last.

- The operating system user interface that is prescribed by Group Policy appears.

**Note**

- In mixed environments, there are three special cases you need to consider:

- If the computer account object is in a Windows NT 4.0 domain and the user account object is in Active Directory, only computer (not user) System Policy is processed when the user logs on. Then, user (not computer) Group Policy is processed.

- If the computer account object is in Active Directory and the user account object is in a Windows NT 4.0 domain, computer (not user) Group Policy is processed during computer startup. When the user logs on, user (not computer) System Policy is processed.

- If the Windows XP computer and user accounts are members of a Windows NT 4.0 domain, only System Policy (not Group Policy) for the computer and user is applied when the user logs on.

**Group Policy Objects (GPOs) Storage**

A Group Policy object is a virtual object that is identified by a Globally Unique Identifier (GUID) and stored at the domain level. The policy setting information of a GPO is stored in the following two locations:

- Group Policy containers in Active Directory. The Group Policy container is an Active Directory container that contains GPO properties, such as version information, GPO status, plus a list of other component settings.

- Group Policy templates in a domain's system volume folder(Sysvol). The Group Policy template is a file system folder that includes policy data specified by .adm files, security settings, script files, and information about applications that are available for installation. The Group Policy template is located in the Sysvol folder in the domain \Policies subfolder.

The **GROUP_POLICY_OBJECT** structure provides information about a GPO in a GPO list, including the version number of the GPO, a pointer to a string that indicates the Active Directory portion of the GPO, and a pointer to a string that specifies the path to the file system portion of the GPO.

**Group Policy Processing Order**

Group Policy settings are processed in the following order:

1. **Local Group Policy object:** Each Windows XP computer has exactly one Group Policy object that is stored locally.

2. **Site:** Any Group Policy objects that have been linked to the site are processed next. Processing is synchronous and in an order that is specified by the administrator.

3. **Domain:** Processing of multiple domain-linked Group Policy objects is synchronous and in an order specified by the administrator.

4. **Organizational units:** Group Policy objects that are linked to the organizational unit that is highest in the Active Directory hierarchy are processed first, then Group Policy objects that are linked to its child organizational unit, and so on. Finally, the Group Policy objects that are linked to the organizational unit that contains the user or computer are processed.

At the level of each organizational unit in the Active Directory hierarchy, one, many, or no Group Policy objects can be linked. If several Group Policy objects are linked to an organizational unit, their processing is synchronous and in an order that is specified by the administrator.

This order means that the local Group Policy object is processed first, and Group Policy objects that are linked to the organizational unit of which the computer or user is a direct member are processed last, which overwrites the earlier Group Policy objects.

This is the default processing order and administrators can specify exceptions to this order. A Group Policy object that is linked to a site, domain, or organizational unit (not a local Group Policy object) can be set to **Enforced** with respect to that site, domain, or organizational unit, so that none of its policy settings can be overridden. At any site, domain, or organizational unit, you can mark Group Policy inheritance selectively as **Block Inheritance**. Group Policy object links that are set to **Enforced** are always applied, however, and they cannot be blocked.

**Security Settings Policy Processing**

Security Settings policy is processed in the context of Group Policy processing.

1. During Group Policy processing, the Group Policy engine determines which security settings policies to apply.

2. If security settings policies exist in a GPO, Group Policy invokes the Security Settings client-side extension.

3. The Security Settings extension downloads the policy from the appropriate location such as a specific domain controller.

4. The Security Settings extension merges all security settings policies according to precedence rules. The processing is according to the Group Policy processing order of local, site, domain, and organizational unit (OU), as described earlier in the Group Policy Processing Order section.If multiple GPOs are in effect for a given computer and there are no conflicting policies, then the policies are cumulative and are merged.

   This example uses the Active Directory structure shown in the following figure (Multiple GPOs and Merging of Security Policy). A given computer is a member of OU2, to which the **GroupMembershipPolGPO** GPO is linked. This computer is also subject to the **UserRightsPolGPO** GPO, which is linked to OU1, higher in the hierarchy. In this case, no conflicting policies exist so the computer receives all of the policies contained in both the **UserRightsPolGPO** and the **GroupMembershipPolGPO** GPOs.

   **Multiple GPOs and Merging of Security Policy**




5. The resultant security policies are stored in secedit.sdb, the security settings database. The security engine gets the security template files and imports them to secedit.sdb.

6. The security settings policies are applied to computers.

The following figure illustrates the security settings policy processing.

**Security Settings Policy Processing**




**Merging of Security Policies on Domain Controllers**

Password policies, Kerberos, and some security options are only merged from GPOs that are linked at the root level on the domain. This is done to keep those settings synchronized across all domain controllers in the domain. The following security options are merged:

- **Network Security: Force logoff when logon hours expire**
- **Accounts: Administrator account status**
- **Accounts: Guest account status**
- **Accounts: Rename administrator account**
- **Accounts: Rename guest account**

Another mechanism exists that allows security policy changes made by administrators using net accounts to be merged into the **Default Domain Policy** GPO. User rights changes that are made by using Local Security Authority (LSA) APIs are filtered into the **Default Domain Controllers Policy** GPO.

**Special Considerations for Domain Controllers**

If an application is installed on a primary domain controller (PDC) with operations master role (also known as flexible single master operations or FSMO) and the application makes changes to user rights or password policy, these changes must be communicated to ensure that synchronization across domain controllers occurs. Scesrv.dll receives a notification of any changes made to the security account manager (SAM) and LSA that need to be synchronized across domain controllers and incorporates the changes into the **Default Domain Controller Policy** GPO by using Scecli.dll template modification APIs.

**When Security Settings are Applied**
There are several instances after you have edited the security settings policies, when the settings are refreshed on the computers in the organizational unit linked to your Group Policy object:

- When a computer is restarted.
- Every 90 minutes on a workstation or server and every 5 minutes on a domain controller.
- By default, Security Policy settings delivered by Group Policy are also applied every 16 hours (960 minutes) even if a GPO has not changed.

  It is possible to change this default period by using the registry entry **MaxNoGPOListChangesInterval** in the following subkey:

  **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon \GPExtentions\ {827D319E-6AC-11D2-A4EA-00C04F79F83A},**

  The data type of this entry is **REG_DWORD** and the value is number of minutes.

### Note

- Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference on the *Microsoft Windows Server 2003 Deployment Kit* companion CD.

**Persistence of Security Settings Policy**
Security settings can persist even if a setting is no longer defined in the policy that originally applied it.

In Windows Server 2003 and Windows XP, security settings might persist in the following cases:

- The setting has not been previously defined for the computer.
- The setting is for a registry security object.
- The settings are for a file system security object.

In Windows 2000, security settings might persist even if the setting is no longer defined in the GPO that originally applied it. All settings applied through local policy or through a Group Policy object are stored in a local database on your computer. Whenever a security setting is modified, the computer saves the security setting value to the local database, which retains a history of all the settings that have been applied to the computer. If a policy first defines a security setting and then no longer defines that setting, then the setting takes on the previous value in the database. If a previous value does not exist in the database then the setting does not revert to anything and remains defined as is. This behavior is sometimes referred to as "tattooing."

Registry and file security settings will maintain the values applied through Group Policy until that setting is set to other values.

On domain controllers running Windows Server 2003 or Windows 2000, all security settings persist.

**Permissions Required for Policy to Apply**
Both Apply Group Policy and Read permissions are required to have the settings from a Group Policy object apply to users or groups, and computers.

**Filtering Security Policy**
By default, all GPOs have Read and Apply Group Policy both Allowed for the AuthenticatedUsers group. The Authenticated Users group includes both users and computers. Security settings policies are computer-based. To specify which client computers will or will not have a Group Policy object applied to them, you can deny them either the Apply Group Policy or Read permission on that Group Policy object. Changing these permissions allows you to limit the scope of the GPO to a specific set of computers within a site, domain, or OU.

### Note

- Do not use security policy filtering on a domain controller as this would prevent security policy from applying to it.

**Migration of GPOs Containing Security Settings**
In some situations, a policy administrator might want to migrate GPOs from one domain environment to another environment. The two most common scenarios are test-to-production migration, and production-to-production migration. The GPO copying process has implications for some types of security settings.

Data for a single GPO is stored in multiple locations and in various formats; some data is contained in Active Directory and other data is stored on the Sysvol share on the domain controllers. Certain policy data might be valid in one domain but might be invalid in the domain to which the GPO is being copied. For example, Security Identifiers (SIDs) stored in security policy settings are often domain-specific. So copying GPOs is not as simple as taking a folder and copying it from one computer to another.

The following security policies can contain security principals and might require some additional work to successfully move them from one domain to another.

- User rights assignment
- Restricted groups
- Services
- File system
- Registry
- The GPO DACL, if you choose to preserve it during a copy operation

To ensure that data is copied correctly, administrators can use Group Policy Management Console (GPMC). When migrating a GPO from one domain to another, GPMC ensures that all relevant data is properly copied. GPMC also introduces migration tables, which can be used to update domain-specific data to new values as part of the migration process. GPMC hides much of the complexity involved in the migrating GPO operations, and provides simple and reliable mechanisms for performing operations such as copy and backup of GPOs.

For more information about GPMC, see "How Group Policy Management Console Works" in this collection.

**Security Settings Policies**
Security Settings includes the following types of policies:

**Account Policies**
Defined for computers, account policies affect how user accounts can interact with the computer or domain. Account policies are implemented at the domain level. For domain controllers, Windows Server 2003 and Windows 2000 allow only one domain account policy: the account policy applied to the root domain of the domain tree. The domain account policy becomes the default account policy of any Windows-based computer that is a member of the domain; this account policy affects all domain users. If an account policy is defined for an organizational unit (OU), it will affect only local accounts on member servers. The account policy settings for the OU affect the local policy on any computers contained in the OU.

For domain accounts, only one account policy is permitted per domain. This account policy must be specified in the **Default Domain Policy** GPO, or in a new GPO that is linked to the root of the domain and has precedence over the **Default Domain Policy** GPO. The **Default Domain Policy** GPO is created when a computer running Windows Server 2003 or Windows 2000 Server is first promoted to domain controller. A domain controller always gets the account policy from a GPO linked to the domain, by default from the **Default Domain Policy** GPO.

Account policies include the following types of policies:

**Password Policy**
Used for domain or local user accounts. Determines settings for passwords, such as enforcement and lifetimes.

**Account Lockout Policy**
Used for domain or local user accounts. Determines the circumstances and length of time that an account will be locked out of the system.

**Kerberos Policy**

Used for domain user accounts. Determines Kerberos-related settings, such as ticket lifetimes and enforcement. Kerberos policies do not exist in local computer policy.

**Local Policies**
These policies apply to the local computer and include the following settings:

**Audit Policy**
Determine whether security events are logged into the Security log on the computer. Also determines whether to log successful attempts, failed attempts or both. (The Security log is part of Event Viewer.)

**User Rights Assignment**
Determine which users or groups have logon rights or privileges on the computer.

**Security Options**
Enable or disable security settings for the computer, such as digital signing of data, Administrator and Guest account names, floppy drive and CD-ROM access, driver installation, and logon prompts.

**Event Log**
Event Log security policies specify attributes for events logs such as log size maximum, access rights, and retention methods. There are settings for application, security, and system logs.

**Restricted Groups**
You can define two properties for restricted groups, Members, and Members Of. The Members list defines who belongs and who does not belong to the restricted group. The Member Of list specifies which other groups the restricted group belongs to.

**System Services**
You can specify security policy for starting system services in automatic or manual mode, or to disable them. You can also set security properties on the objects to specify access permissions, inheritance or auditing settings, and ownership.

**Registry**
You can configure security descriptors for registry keys.

**File System**
You can use Security Settings to configure security descriptors for files, folders for file system objects.

**Predefined Security Templates**
Windows 2000 and Windows Server 2003 provide predefined security templates that you can use as a starting point for creating security policies that are customized to meet different organizational requirements. The templates can be customized using the Security Templates snap-in. After you customize the predefined security templates, they can be used to configure an individual computer or multiple computers. To configure individual computers, you can use the Security Configuration and Analysis snap-in, the secedit.exe command line tool, or you can import the template into Local Security Settings. You can configure multiple computers by importing a template into a Group Policy object.

You can also use a security template as a baseline for analyzing a system for potential security gaps or policy violations using the Security Configuration and Analysis snap-in. By default, the predefined security templates are stored in: *Systemroot*\Security\Templates.

**Compatible (compatws.inf)**
Default permissions for workstations and servers are primarily granted to three local groups: Administrators, Power Users, and Users. Administrators have the highest privilege, while Users have the lowest. One way to significantly improve the security, reliability, and total cost of system ownership is to do the following:

- Make sure that end-users are members of the Users group

- Deploy applications that can run successfully under a User context.

  Applications that comply with the *Certified for Windows Program* meet this criterion. However, applications that are not certified for Windows 2000 or Windows Server 2003 might not run under a User context. *The Certified for Windows Program* is designed to ensure that applications that run on Windows Server 2003 or Windows 2000 Server can meet the highest standards for reliability, availability, security, and supportability.

  If you must support non-certified applications, you have two options:

  - Allow end-users to be members of the Power Users group.

- Relax the default permissions granted to the Users group.

Since Power Users have inherent capabilities, such as creating users, groups, printers and shares, some customers would rather relax the default User permissions than allow end-users to be members of the Power Users group. This is precisely what the Compatible template is for. The Compatible template loosens the default file and registry permissions granted to Users in a manner that is consistent with the requirements of most non-certified applications. Additionally, since it is assumed that the administrator applying the compatible template does not want end users to be Power Users, the compatible template also removes all members of the Power Users group.

**Note**

- The Compatible template should *not* be applied to domain controller computers. For example, do not import the compatible template to the **Default Domain Policy** GPO or the **Default Domain Controller Policy** GPO.

Information about the Microsoft Certified for Windows Program is available on the Windows Server 2003 Web site.

**Secure (secure*.inf)**
The Secure templates (secure*.inf) define enhanced security settings that are least likely to impact application compatibility. For example, the Secure templates define stronger password, lockout, and audit settings.

The Secure templates also limit the use of the LAN Manager and NTLM authentication protocols by configuring clients to send only NTLMv2 responses and Servers to refuse LAN Manager responses:

- In order to apply securews.inf to a member computer, all of the domain controllers that contain the accounts of all users that will logon to the client must be running Windows NT 4.0 Service Pack 4 or higher.

- In order to apply securews.inf to a member computer that is joined to a domain containing Windows NT 4.0 domain controllers, the clocks between the NT 4.0 domain controllers and the member computers must be within 30 minutes of each other.

- If a client is configured with securews.inf, it will not be able to connect to servers that only use the LAN Manager authentication protocol or Windows NT 4.0 servers prior to SP 4 using a local account defined on the target server.

- If a client is configured with securews.inf, it will not be able to connect to Windows 2000 or Windows NT 4.0 servers using a local account defined on the target server unless the clock on the target server is within 30 minutes of the clock on the client.

- If a client is configured with securews.inf, it will not be able to connect to a computer running Windows XP by using a local account that is defined on the target computer unless the clock on the target computer is within 20 hours of the clock on the client.

- If a client is configured with securews.inf, then it will not be able to connect to servers running LAN Manager that are operating in share-level security mode.

- If a server is configured with securews.inf, then a user with a local account on that server will not be able to connect to the server by using that local account from a client computer that is running only LAN Manager.

- If a Windows 2000 server is configured with securews.inf, then a client using a local account on that server that is also configured to use NTLMv2 authentication will not be able to connect unless the clocks on the two computers are within 30 minutes of each other.

- If a computer running Windows 2000 or Windows Server 2003 is configured with securews.inf, then a client with a local account on that server that is also configured to use NTLMv2 authentication will not be able to connect unless the clocks on the two computers are within 20 hours of each other.

- If a domain controller is configured with securedc.inf, then a user with an account in that domain will not be able to connect to any member server from a client computer that is running only LAN Manager using their domain account.

LAN Manager clients include Windows for Workgroups, Windows 95, and Windows 98 platforms that do not have the Active Directory Client Extensions Pack installed. If the Active Directory Client Extensions Pack is installed on Windows 95 or Windows 98 platforms, those clients can use NTLMv2. Microsoft Windows Millennium Edition supports NTLMv2.

Windows NT-based computers (SP 4 and higher) can be configured to send only NTLMv2 responses by setting **HKEY_LOCAL_MACHINE\System CurrentControlSet\Control\Lsa\LMCompatibilityLevel** =3 or higher. The **LMCompatibilityLevel** registry value is exposed in the **Local Policies\Security Options** node of the **Security Configuration Manager** tools (**Local Security Policy** and **Security Settings** extension of **Group Policy Object Editor**) as **Network Security: Lan Manager authentication level** policy setting. Clients that have the secure or high secure template applied to them will have **LMCompatibility** set to 3 or higher.

The Secure templates also provide further restrictions for anonymous users by preventing anonymous users (for example, users from untrusted domains) from doing the following:

- Enumerating account names and shares
- Performing SID to Name or Name to SID translations

Finally, the Secure templates enable server-side SMB packet signing, which is disabled by default for workstations and servers. Since client-side SMB packet signing is enabled by default, SMB packet signing is always negotiated when workstations and servers are operating at the secure level.

**High Secure (hisec*.inf)**
The High Secure templates are supersets of the Secure templates that impose further restrictions on the levels of encryption and signing required for authentication and for the data that flows over Secure Channels and between SMB clients and servers. For example, while the Secure templates cause servers to refuse LAN Manager responses, the High Secure templates cause servers to refuse both LAN Manager and NTLM responses. The secure template enables server-side SMB packet signing, but the High Secure template requires it. The High Secure Template also requires strong (128 bit) encryption and signing for the Secure Channel data that constitutes domain-member and domain-domain trust relationships.

- To apply hisecws.inf to a member computer, the following requirements must be met:
    - All the domain controllers that contain the accounts of all users that will logon to the client must be running Windows NT 4.0 Service Pack 4 or higher, Windows 2000, or Windows Server 2003.
    - All the domain controllers for the domain to which the client is joined must be running at least Windows 2000.
- To apply hisecdc.inf to a domain controller, all of the domain controllers in all trusted or trusting domains must be running Windows 2000 or later.
- If a client is configured with hisecws.inf, it will not be able to connect to servers running only LAN Manager or Windows NT 4.0 servers prior to SP 4 using a local account defined on the target server.
- If a client is configured with hisecws.inf, it will not be able to connect to Windows 2000 or Windows NT 4.0 servers using a local account defined on the target server unless the clock on the target server is within 30 minutes of the clock on the client.
- If a client is configured with hisecws.inf, it will not be able to connect to a computer running Windows 2000 server or Windows Server 2003 using a local account defined on the target server unless the clock on the target server is within 20 hours of the clock on the client.
- If a client is configured with hisecws.inf, then it will not be able to connect to servers running LAN Manager operating in share level security mode.
- If a server is configured with hisecws.inf, then a client with a local account on that server will not be able to connect to it from a client that does not support NTLMv2.
- If a server is configured with hisecws.inf, then a Windows NT client with a local account on that server will not be able to connect to the server unless the Windows NT client is configured to send NTLMv2 responses.

- If a server is configured with hisecws.inf, then all clients that want to use SMB to connect to that server must have client-side SMB packet signing-enabled. Client-side SMB packet-signing is enabled by default for all Windows 2000 computers.

- If a domain controller is configured with hisecdc.inf, then a user with an account in that domain will not be able to connect to member servers by using the domain account if the connection is being made from a computer using only LAN Manager authentication protocol.

- If a domain controller is configured with hisecdc.inf, then a user with an account in that domain will not be able to connect to member servers using that domain account unless the following conditions apply:

    - Both client and target server are Windows 2000 or Windows Server 2003 and, thus, can use Kerberos rather than LAN Manager-based authentication.

    - The client is configured to send NTLMv2 responses.

- If a domain controller is configured with hisecdc.inf, then LDAP clients will not be able to bind with the Active Directory LDAP server unless data signing is negotiated. Thus, bind requests using **ldap_simple_bind** or **ldap_simple_bind_s** are rejected. By default, all Microsoft LDAP clients that ship with Windows XP or Windows Server 2003 request data signing if Transport Layer Security\Secure Sockets Layer (TLS\SSL) protocol is not already being used. If TLS\SSL is being used, then data signing is considered to be negotiated.

Clients that do not support NTLMv2 include Windows for Workgroups, Windows NT clients prior to SP 4, and Windows 95 and Windows 98 platforms that do not have the DS Client Pack installed. Windows NT-based computers (SP 4 and higher) can be configured to send only NTLMv2 responses by setting **HKEY_LOCAL_MACHINE\System CurrentControlSet\Control\Lsa\LMCompatibilityLevel** = 3 or higher. The **LMCompatibilityLevel** registry value is exposed in the **Local Policies\Security Options** node of the security configuration manager tools (Local Security Settings and Security Settings extension of Group Policy Object Editor) as the **Network Security: Lan Manager authentication level** policy setting. Clients that have the secure or high secure template applied to them will have **LMCompatibility** set to 3 or higher.

In addition to further restrictions on the use of computers running LAN Manager protocols and the requirements for encryption and signing of secure channel and SMB traffic, the High Secure templates also restrict clients from caching domain credentials for subsequent logons when the domain controllers aren't available. For example, this would prohibit traveling users who don't have a network connection available for domain authentication from being able to login to their laptop.

Finally, Hisecws.inf uses restricted group settings to:

- Remove all members of the Power Users group.

- Ensure that only Domain Admins and the local Administrator account are members of the local Administrators group.

Hisecws defines these group restrictions under the assumption that only applications that are certified for Windows 2000 or Windows Server 2003 are deployed. With certified applications in place, neither the unsecured Compatible template nor the unsecured Power Users group is needed. Instead, end users can successfully run certified applications under the secure context of a normal User, which is defined by the default file system and registry ACLs.

**Root Directory Permissions (rootsec.inf)**
Specifies the new root permissions introduced with Windows XP. By default, rootsec.inf defines these permissions for the root of the system drive. This template can be used to reapply the root directory permissions if they are inadvertently changed or the template can be modified to apply the same root permissions to other volumes. As specified, the template does not overwrite explicit ACEs defined on child objects, it propagates only the permissions that are inherited by child objects.

**No Terminal Server SID (notssid.inf)**
The default file system and registry ACLs on servers grant permissions to a Terminal Server SID. The Terminal Server SID is used only when Terminal Server is running in application compatibility mode. If Terminal Server is not being used, this template can be applied to remove the unnecessary Terminal Server SID from these file system and registry locations. Removing the ACE for the Terminal Server SID from these default file system and registry locations does not increase the security of the system. Instead of removing the Terminal Server SID, it is recommended to simply run Terminal Server in Full Security mode rather than Application Compatibility mode. When running in Full Security mode, the Terminal Server SID is not

used.

**Setup security.inf**
Setup security is a computer-specific template that represents the default security settings applied during installation of the operating system including the file permissions for the root of the system drive. This template, or portions thereof, can be used for disaster recovery purposes. Setup security.inf should *not* be applied by using Group Policy.

### Note

- Do not apply the predefined security templates to production systems without performing comprehensive testing first to ensure that the security and functionality requirements for your specific organization are met.

You can view Security template settings by using the Security Templates snap-in. You can also examine the changes that a template would make if it were applied by using the Security Configuration and Analysis snap-in. To examine the proposed changes, import the template into a database then perform an analysis against that database. Performing an analysis does not change any system settings, but will highlight the differences between the settings specified in the template and the settings as they currently exist on the system.

**Related Information**
The following resources contain additional information that is relevant to this section.

- "How Core Group Policy Works" in the Group Policy Collection.

- For information about the Microsoft Certified for Windows Program, see the Windows Server 2003 Web site.

# Security Settings Extension Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this subject**


- Resultant Set of Policies
- Command Line Tools
- Security Settings Policies
- Related Information


This section presents an overview of Resultant Set of Policies (RSoP), which you use to determine which policy settings are currently in effect for a computer or user, and to assess how policy settings would affect computers or users if a specific Group Policy object were applied to them. It also describes the Windows Server 2003 command-line tools for configuring and analyzing security settings.

**Resultant Set of Policies**

Windows XP and the Windows Server 2003 family of operating systems support an enhanced Group Policy infrastructure that utilizes Windows Management Instrumentation (WMI) to collect Group Policy-related data for planning and troubleshooting Group Policy. This structure is the Resultant Set of Policy (RSoP), a query engine that polls existing policy settings and planned policy settings, and then reports the results of those queries. RSoP polls existing policies based on site, domain, domain controller, and organizational unit. RSoP gathers this information from the Common Information Management Object Model (CIMOM) database (also known as the CIM-compliant object repository) by using WMI. Administrators can use RSoP in one of two modes. To determine which policy settings are in effect for a particular computer or user, administrators use Resultant Set of Policy – Logging Mode, and to evaluate how policy settings would affect a computer or user if a specific Group Policy object were applied to users or computers, they use Resultant Set of Policy – Planning Mode.

In Windows Server 2003, you can use the Group Policy Results node Group Policy Management Console (GPMC) to access the Resultant Set of Policy – Logging Mode capabilities. Group Policy Results represents the actual resultant set of policy that was applied to a given user and computer. This information is obtained by directly querying the target user/computer. Each sub-node represents a different RSoP query for a given user/computer combination. The Group Policy Modeling node in GPMC allows administrators to access the Resultant Set of Policy – Planning Mode capabilities of Windows Server 2003.

**Note**

- As with other Group Policy settings, you must fully test your implementation in a test domain before you deploy your security settings to your production environment.

Some of the security settings extensions of Group Policy provide RSoP classes to represent data pertaining to security policy settings. The *Security Policy RSoP Classes* section, later in this document, lists the RSoP classes for security policy settings.

On Windows 2000 computers, you can use the Gpresults.exe tool to display information about how Group Policy affects both the currently logged-on user and the computer. For information about the Gpresults syntax, see "IPSec Policy Extension Tools and Settings" in this collection.

The Gpresults command-line tool is available in the Windows 2000 Resource Kit. To download this tool, see the download site for Windows 2000 Server Resource Kit tools.

For more detailed information about RSoP and WMI and to download SDKs, see the Microsoft Platform SDK link on the Web Resources page.

For more information about GPMC, see the Group Policy Management Console link on the Web Resources page.

**Security Policy RSoP Classes**

The RSoP Windows Management Instrumentation (WMI) Method Provider supports the following security policy classes, as listed in the following table.

**Security Policy RSoP Classes**

| Class | Description |
|---|---|
| RSOP_AuditPolicy | This class represents the security setting for a local Group Policy that relates to the auditing of an event type. Events can include, among others, system events and account management events. |
| RSOP_File | This class represents a security policy setting that defines the access permissions and audit settings for a securable file system object. |
| RSOP_RegistryKey | This class represents a security policy setting that defines the access permissions and audit settings for a particular registry key. |
| RSOP_RegistryValue | This class represents specific security-related registry values. |
| RSOP_RestrictedGroup | This class represents a security policy setting that defines the members of a restricted (security-sensitive) group. |
| RSOP_SecurityEventLogSettingBoolean | This class represents a security policy setting that determines whether or not guests can access the system, application and security event logs. |
| RSOP_SecurityEventLogSettingNumeric | This class represents a security policy setting that determines numeric properties related to the system, application and security event logs. Properties include the number of days to retain entries and maximum log size. |
| RSOP_SecuritySettingBoolean | This class represents the Boolean security setting for an account policy. Account policies include password policies and account lockout policies. |
| RSOP_SecuritySettingNumeric | This class represents the numeric security setting for an account policy. Account policies include password policies, account lockout policies, and Kerberos-related policies. |
| RSOP_SecuritySettings | This is the abstract class from which other RSoP security classes derive. Instances of this class are not logged. RSOP_SecuritySettings derives from the RSOP_PolicySetting class. |
| RSOP_SecuritySettingString | This class represents the string security setting for an account policy. |
| RSOP_SystemService | This class represents the security policy setting that defines the start-up mode and access permissions for a particular system service. |
| RSOP_UserPrivilegeRight | This class represents the security setting for a local Group Policy that relates to the assignment of a particular user privilege. |

**RSOP_PolicySetting**

The **RSOP_PolicySetting** WMI class is the class from which policy objects for client-side extensions are inherited. An instance of this class corresponds to a specific policy setting. This class was added for Windows XP.

Requirements for this class are as follows:

- **Client:** Included in Windows XP Professional.
- **Server:** Included in Windows Server 2003.

**Command Line Tools**

This section describes the Windows Server 2003 command line tools for configuring, analyzing, and updating security settings:

- Secedit.exe
- Gpupdate.exe

**Secedit.exe**

You can use the secedit.exe command to configure and analyze system security by comparing your current configuration to at least one template.

**Note**

- **Secedit /refreshpolicy** has been replaced with **gpupdate**. To refresh local Group Policy settings and Group Policy settings that are stored in Active Directory, including security settings, use **gpudate**. See "Gpupdate," later in this section.

Secedit supports the following commands:

- analyze
- configure
- export
- import
- validate
- generateRollback

**Secedit /analyze**
This command allows you to analyze the security settings on a computer by comparing them against the baseline settings in a database. You can view the results of the analysis in the Security Configuration and Analysissnap-in.

**Syntax**

```
              secedit /analyze /dbFileName.sdb [/cfgFileName] [/overwrite]
[/logFileName] [/quiet]
```

**Parameters**
**/db** *FileName* **.sdb**

This parameter specifies the database used to perform the analysis.

**/cfg** *FileName*

This parameter specifies a security template to import into the database prior to performing the analysis. Security templates are created using the Security Templates snap-in.

**/log** *FileName*

This parameter specifies a file in which to log the status of the configuration process. If not specified, configuration data is logged in the scesrv.log file which is located in the %windir%\security\logs directory.

**/quiet**

This parameter specifies that the analysis process should take place without further comments.

**Examples**
The following is an example of how you can use this command:

**secedit /analyze /db** hisecws.sdb

**Secedit /configure**
You can use secedit /configure to configure local computer security by applying the settings stored in a database.

**Syntax**

```
              secedit
              /configure
              /db
              FileName [/cfg FileName ] [/overwrite][/areasArea1 Area2 ...]
```

```
[/logFileName] [/quiet]
```

**Parameters**

**/db** *FileName*

This parameter specifies the database used to perform the security configuration.

**/cfg** *FileName*

This parameter specifies a security template to import into the database prior to configuring the computer. Security templates are created using the **Security Templates** snap-in.

**/overwrite**

This parameter specifies that the database should be emptied prior to importing the security template. If this parameter is not specified, the settings in the security template are accumulated in the database. If this parameter is not specified and there are conflicting settings in the database and the template being imported, the template settings take precedence.

**/areas** *Area1 Area2*

This parameter specifies the security areas to be applied to the system. If this parameter is not specified, all security settings defined in the database are applied to the system. To configure multiple areas, separate each area by a space. The following security areas are supported:

| Area name | Description |
|---|---|
| SECURITYPOLICY | This includes Account Policies, Audit Policy, Event Log settings, and Security Options. |
| GROUP_MGMT | This includes Restricted Groups settings. |
| USER_RIGHTS | This includes User Rights Assignment. |
| REGKEYS | This includes Registry permissions. |
| FILESTORE | This includes File System permissions. |
| SERVICES | This includes System Services settings. |

**/log** *FileName*

This parameter specifies a file in which to log the status of the configuration process. If not specified, configuration data is logged in the scesrv.log file, which is located in the %windir%\security\logs directory.

**/quiet**

This parameter specifies that the configuration process should take place without prompting the user.

**Examples**

The following are examples of how you can use this command:

**secedit /configure /db** *hisecws.sdb***/cfg***FileName*

**hisecws.inf /overwrite /log** *hisecws.log*

If **db***FileName* (database filename) doesn's exist, **secedit** creates a new database using the settings in **cfg***FileName* (template filename) and applies the configuration. If **db***FileName* exists, then **secedit** merges the settings into the database before applying the newly merged configuration. If you omit **cfg***FileName*, **secedit** applies the configuration using the settings already stored in the database.

**Secedit /export**

Running **secedit /export** allows you to export the security settings stored in the database.

**Syntax**

```
              secedit
              /export [/DBFileName] [/mergedpolicy] [/CFG FileName] [/areasArea1 Area2
...] [/logFileName] [/quiet]
```

**Parameters**
**/db** *FileName*

This parameter specifies the database used to configure security.

**/mergedpolicy**

This parameter merges and exports domain and local policy security settings.

**/CFG** *FileName*

This parameter specifies the template the settings will be exported to.

**/areas** *Area1 Area2*

This parameter specifies the security areas to be exported to a template. If an area is not specified, all areas are exported. Each area should be separated by a space. The following table lists the security areas that can be exported.

**Security Areas and Descriptions**

| Area name | Description |
|---|---|
| SECURITYPOLICY | This includes Account Policies, Audit Policy, Event Log settings, and Security Options. |
| GROUP_MGMT | This includes Restricted Groups settings. |
| USER_RIGHTS | This includes User Rights Assignment. |
| REGKEYS | This includes Registry permissions. |
| FILESTORE | This includes File System permissions. |
| SERVICES | This includes System Services settings. |

**/log** *FileName*

This parameter specifies a file in which to log the status of the export process. If not specified, the default is %windir%\security\logs\scesrv.log.

**/quiet**

This parameter specifies that the configuration process should take place without prompting the user.

**Examples**
The following is an example of how you can use this command:

**secedit /export /db *hisecws* .inf /log *hisecws* .log**

**Secedit /import**
Running **secedit /import** allows you to import a security template into a database so that the settings specified in the template can be applied to a system or analyzed against a system.

**Syntax**

```
secedit
/import
/db
FileName
.sdb
/cfg
FileName.inf [/overwrite] [/areasArea1 Area2 ...] [/logFileName] [/quiet]
```

**Parameters**
**/db** *FileName* **.sdb**

This parameter specifies the database to which the security template settings will be imported.

**/CFG** *FileName*

This parameter specifies a security template to import into the database. Security templates are created using the Security Templates snap-in.

**/overwrite** *FileName*

This parameter specifies that the database contents should be cleared prior to importing the security template. If this parameter is not specified, the settings in the security template are accumulated in the database. If this parameter is not specified and there are conflicting settings in the database and the template being imported, the template settings take precedence.

**/areas** *Area1 Area2*

This parameter specifies the security areas to be exported to a template, as listed in the following table. If an area is not specified, all areas are exported. Each area should be separated by a space.

| Area name | Description |
|---|---|
| SECURITYPOLICY | This includes Account Policies, Audit Policy, Event Log settings, and Security Options. |
| GROUP_MGMT | This includes Restricted Groups settings. |
| USER_RIGHTS | This includes User Rights Assignment. |
| REGKEYS | This includes Registry permissions. |
| FILESTORE | This includes File System permissions. |
| SERVICES | This includes System Services settings. |

**/log** *FileName*

This parameter specifies a file in which to log the status of the export process. If not specified, the default is %windir%\security\logs\scesrv.log.

**/quiet**

This parameter specifies that the configuration process should take place without prompting the user.

**Examples**
The following is an example of how you can use this command:

**secedit /import /db hisecws.sdb /cfg hisecws.inf /overwrite**

**Secedit /validate**
You can use **secedit /validate** to validate the syntax of a security template to be imported into a database for analysis or application to a system.

**Syntax**

```
secedit
/validate FileName
```

**Parameters**
*FileName*

This parameter specifies the file name of the security template you have created with Security Templates.

**Examples**
The following is an example of how you can use this command:

**secedit /validate /cfg filename**

**Secedit /GenerateRollback**
You can run **secedit / GenerateRollback** to generate a rollback template with respect to a configuration template. When applying a configuration template to a computer you have the option of creating rollback template which, when applied, resets the security settings to the values before the configuration template was applied.

**Syntax**

```
secedit /GenerateRollback
/CFG
FileName.inf /RBK SecurityTemplatefilename.inf [/logRollbackFileName.inf]
[/quiet]
```

**Parameters**
**/CFG** *FileName*

This parameter specifies the file name of the security template for which you want to create a rollback template of.

**/RBK** *FileName*

This parameter specifies the file name of the security template that will be created as the rollback template.

**Gpupdate**
You can use **gpupdate** to refresh local Group Policy settings and Group Policy settings that are stored in Active Directory, including security settings. This command supersedes the now obsolete **/refreshpolicy** option for the **secedit** command.

**Syntax**

```
gpupdate [/target:{computer | user}] [/force] [/wait:Value] [/logoff]
[/boot]
```

**Parameters**
**/target:{computer | user}**
This parameter specifies to processes only the *Computer* settings or the current *User* settings. By default, both the computer settings and the user settings are processed.

**/force**
Using this parameter ignores all processing optimizations and reapplies all settings.

**/wait:Value**
Indicates the number of seconds that policy processing must wait to finish. The default is 90 minutes with a randomized delay of up to 30 minutes — for a total maximum refresh interval of up to 120 minutes. 0 equals no wait, and -1 equals wait indefinitely.

**/logoff**
This parameter specifies that a user log off occurs after the refresh has completed. This is required for those Group Policy client-side extensions that do not process on a background refresh cycle but that do process when the user logs on, such as user Group Policy Software Installation and Folder Redirection. This option has no effect if there are no extensions called that require the user to log off.

**/boot**
Using this parameter restarts the computer after the refresh has completed. This is required for those Group Policy client-side extensions that do not process on a background refresh cycle but that do process when the computer starts up, such as computer Group Policy Software Installation. This option has no effect if there are no extensions called that require the computer to be restarted.

**/?**
Using this parameter displays help at the command prompt.

**Examples**
The following examples show how you can use the **gpupdate** command:

- **gpupdate**

  This command triggers a Group Policy refresh.

- **gpupdate /target:computer**

  This command triggers a Group Policy refresh of the Computer Settings policies only.

- **gpupdate /force /wait:100**

  This command triggers a Group Policy refresh, reapplies all policy settings, and indicates to wait 100 seconds for policy

to finish processing.

- **gpupdate /boot**

  This command triggers a Group Policy refresh and then causes the computer to restart.

**Security Settings Policies**

Security Settings include policy settings to control the following aspects of security:

- Account Policies
- Local Policies
- Event Log
- Restricted Groups
- System Services
- Registry
- File System

For a description of all security settings policies, see"Security Policy Settings" in the Security Collection.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Security Policy Settings" in the Security Collection.
- "How Core Group Policy Works" in this collection.
- "What Is Resultant Set of Policy?" in this collection.
- "IPSec Policy Extension Tools and Settings" in this collection.
- For information about RSoP and WMI and to download SDKs, see the Microsoft Platform SDK link on the Web Resources page.
- For information about GPMC, see the Group Policy Management Console link on the Web Resources page.
- The Tools and Settings Collection.

# IPSec Policy Extension Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Internet Protocol Security (IPSec) consists of a set of open standards designed to ensure private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services.

IPSec uses policy-based management to define IPSec configuration options and security settings to enable two systems to agree on how to secure traffic between them. To configure IPSec security services, you create IPSec policies. An IPSec policy is a collection of settings and rules that determine IPSec behavior, including the types of network traffic IPSec examines, whether to permit or block traffic, whether to negotiate security, and how to authenticate an IPSec peer.

IPSec is included in the Microsoft Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; Windows Server 2003, Web Edition operating systems; the 64-bit editions of Windows Server 2003, and Windows 2000 Server and Windows 2000 Professional, Windows XP Professional, and Windows XP Home Edition operating systems (Kerberos authentication is not available in Windows XP Home edition).

**In this subject**


- What Is IPSec Policy Extension?
- How IPSec Policy Extension Works
- IPSec Policy Extension Tools and Settings

# What Is IPSec Policy Extension?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is IP Security Policy Extension?**

**In this section**


- IPSec Policy Settings and Rules
- IPSec Policy Core Scenarios
- IPSec Policy Advantages
- Related Information


Internet Protocol Security (IPSec) is an architecture defined by the Internet Engineering Task Force (IETF) RFC 2401. This architecture involves several protocols that perform various functions in the architecture.

A network is not secure until servers can identify the computers communicating with them. IPSec enables secure, trusted communications between IP addresses. The system behind the IP address has an identity that is verified by using an authentication process. The only computers that must be aware of IPSec are the sending and receiving computers. Each computer handles security at its respective end, and assumes that the medium over which the communication takes place is not secure. Any computers that route data between the source and destination computer are not required to support IPSec.

You can deploy IPSec in several scenarios:

- A local area network (LAN) for client/server and server-to-server.
- A wide area network (WAN) for router-to-router and gateway-to-gateway.
- A remote network's dial-up clients and for Internet access from private networks.

Both the sending and the receiving computer require IPSec configuration (an IPSec policy) to specify options and security settings to enable the two systems to agree on how to secure traffic between them. The IP Security Policies extension of the Group Policy Object Editor snap-in is used to configure IPSec policies based on Active Directory. You can also configure IPSec policies on a computer that is not a domain member; however, it is more difficult to assign and manage IPSec policy configurations and trust relationships for computers that are not members of a trusted domain.

**IPSec Policy Settings and Rules**

An IPSec policy contains general IPSec policy settings that apply regardless of the rules that are configured, and one or more IPSec policy rules that determine IPSec behavior.

**General IPSec Policy Settings**

The general IPSec policy settings specify the name of the policy, its description for administrative purposes, polling interval for policy change, key exchange settings, and key exchange methods. The following table lists these general settings.

**General IPSec Policy Settings**

| Settings | Description |
|---|---|
| Name | The name for the policy. |
| Description | The optional text description used to describe the purpose of the IPSec policy. It is recommended that you complete and update this description to provide a summary of the settings and rules for the policy. |

| | |
|---|---|
| Policy change poll interval | Specifies the number of minutes between consecutive polls for changes in IPSec policies based on Active Directory. This polling does not detect a change in domain or organizational unit membership, or the assignment or removal of assignment of a new policy. These events are detected when the Winlogon service polls for changes in Group Policy, which occurs every 90 minutes by default. |
| Key exchange settings | Define the way in which new keys are derived and how often they are renewed. |
| Key exchange methods | Determine the ways in which identities are protected during the key exchange.<br><br>The default key exchange settings and methods are configured to work for most IPSec deployments. Unless there are special security requirements, default settings should not have to be changed. |

**IPSec Policy Rules**
IPSec policy rules specify settings such as the type of traffic IPSec examines, whether to permit or block traffic, whether to negotiate security, and how to authenticate an IPSec peer.

IPSec rules include a list of IP filters that specify a particular subset of inbound or outbound network traffic that should be secured. A filter is required to cover any traffic to which an associated IPSec rule applies. IPSec filters are inserted into the IP layer of the TCP/IP networking protocol stack on the computer so that they can filter all inbound or outbound IP packets.

**IP Filter Settings**
An IP filter contains the following settings:

**The source and destination address of the IP packet.**
You can specify any IP address assigned to the IPSec peer, a single IP address, IP addresses by DNS name, or groups of addresses to specify IP subnets.

**The protocol over which the packet is being transferred.**
By default, all protocols in the TCP/IP protocol suite are selected. However, you can specify an individual protocol for this filter to meet special requirements, including custom protocols.

**The source and destination port for TCP and UDP.**
By default, all TCP and UDP ports are selected, but you can select a specific TCP or UDP port.

**IPSec Rule Entries**
An IPSec rule contains the following configurable entries:

**Filter list.**
This list contains one or multiple predefined filters that specify IP addresses and the types of traffic to which an action (permit or block traffic, or negotiate security) is applied.

**Filter action.**
This entry specifies the security requirements for data transmission.

You can configure one of the following options for a filter action:

- **Permit.** Select this option to permit traffic. IPSec passes this traffic without modification or the requirement for security. This is appropriate for traffic from specific computers that are not IPSec-capable. Be sure to limit the IP filter list to a minimal scope when using this type of filter action, so you do not let traffic through which should be secured.

- **Block.** Select this option to block traffic. IPSec silently discards this traffic. Be sure to use an IP filter list that appropriately defines the scope of traffic when using a blocking filter action so that you do not block valid computers from communicating.

- **Negotiate security.** Select this option to negotiate IPSec parameters. IPSec requires the negotiation of security associations and the sending or receiving of IPSec-secured traffic. If you choose this option, you can also configure: security methods; allowance of initial incoming unsecured traffic (**Accept unsecured communication, but always respond using IPSec**); enabling of communication with non-IPSec-enabled computers (**Allow unsecured communication with non-IPSec-aware computer**); and generation of session keys from new keying material (**Use session key perfect forward secrecy (PFS)**).

**Authentication.**

This entry contains one or more authentication methods in order of preference that are used for protection during Internet Key Exchange (IKE) negotiations. The available authentication methods are the Kerberos v5 protocol, the use of a certificate issued from a specified certification authority, or a pre-shared key. The IKE protocol securely establishes a trust relationship between each computer, negotiates security options, and dynamically generates shared, secret cryptographic keying material.

**Tunnel endpoint.**
This entry contains settings that determine whether traffic is tunneled and, if it is, the tunnel endpoint. The tunnel endpoint is configured on the **Tunnel Setting** tab in the properties of an IPSec rule within an IPSec policy.

**Connection type.**
This entry contains a setting that specifies whether the rule applies to only local area network (LAN) connections, to only Point-to-Point Protocol (PPP)-based connections, or to both types of connections. The interface applicability is configured on the **Connection Type** tab in the properties of an IPSec rule within an IPSec policy.

For more detailed information about IPSec, see "What Is IPSec?" in the Networking collection.

**IPSec Policy Core Scenarios**
Corporate networks face various types of external and internal threats from untrusted computers, such as denial-of-service attacks, data corruption, and data theft. Network-based attacks on any application or service running on an internal corporate network server can result in:

- Loss of service through denial-of-service of the application, the service, or the network.
- Data corruption.
- Theft of information such as user credentials or data theft.
- Administrative control of the server or other network computers.

To protect network servers against network-based attacks, administrators can use IPSec. To provide enhanced security for networks against untrusted network attacks, you can require IPSec-authenticated, signed, and encrypted communication between computers.

You can use IPSec as part or your organization's security strategy to provide protection against network-based attacks from untrusted computers. IPSec is intended for use in environments where untrusted network access and attacks on network traffic are a realistic threat.

The features of IPSec (such as strong, cryptographic-based authentication and encryption) make it particularly useful for securing traffic that must traverse untrusted network paths such as the corporate intranet or the Internet. IPSec is also appropriate for securing traffic that uses protocols and applications that do not provide sufficient security for communications.

By using policy based on Active Directory, you can secure most communications for a group of servers. You can also use command-line tools to create, modify, and assign IPSec policies. The tools you use vary, depending on the operating system running on the computers. For more information about command-line tools, see the "IPSec Policy Extension Tools and Settings" section in this collection.

**IPSec Policy Advantages**
IPSec provides a key line of defense against private network and Internet attacks. IPSec keeps data intact as it is transmitted across the network and protects communications between private network computers, domains, sites, remote sites, extranets, and dial-up clients.

By using cryptographic protection services, security protocols, and dynamic key management, IPSec:

- Protects the content of IP packets.
- Provides a defense against active and passive network attacks by using packet filtering and the enforcement of trusted communication.

IPSec provides protection against the following types of attacks:

**Sniffers (lack of confidentiality).**
The Encapsulating Security Payload (ESP) protocol in IPSec encrypts the payload of IP packets to provide data confidentiality.

**Data modification.**
IPSec uses cryptography-based keys, shared only by the sending and receiving computers, to create a cryptographic checksum for each IP packet. Any modification to the packet data alters the checksum, indicating to the receiving computer that the packet was altered in transit.

**Identity spoofing, password-based, and application-layer attacks.**
IPSec permits the exchange and verification of identities without exposing that information to interpretation by an attacker. To establish trust between the communicating systems, and to ensure that only trusted systems can communicate with each other, IPSec uses mutual authentication. After identities are established, IPSec uses cryptography-based keys to create a cryptographic checksum for each IP packet. The checksum ensures that only the computers that have knowledge of the keys could have sent each packet.

**Man-in-the-middle attacks.**
IPSec combines mutual authentication with shared, cryptography-based keys to provide end-to-end data integrity and information hiding.

**Denial-of-service attacks.**
IPSec uses IP packet filtering methodology to determine whether communication is allowed, secured, or blocked, according to the IP address ranges, IP protocols, or specific TCP and UDP ports.

**IPSec Management with Group Policy**
To provide strong, cryptography-based security to protect data, and to reduce administrative costs, IPSec uses policy-based administration. You can configure IPSec policies to meet the security requirements of a computer, application, organizational unit, domain, site, or global enterprise. You can use the **IP Security Policies Management** snap-in provided in Windows XP, Windows 2000, and Windows Server 2003 to define IPSec policies for computers through Active Directory (for domain members), or on the local computer (for computers that are not members of a domain).

**Related Information**
The following resources contain additional information that is relevant to this section:

- "How Core Group Policy Works" in the Group Policy collection.
- "What Is IPSec?" in the Networking collection.

# How IPSec Policy Extension Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How IP Security Policy Extension Works**

**In this section**


- IPSec Policy Extension Architecture
- IPSec Policy Extension Processes and Interactions
- Related Information


An Active Directory deployment is recommended for central management and distribution of Internet Protocol Security (IPSec) policies. Deployment of IPSec policy in an Active Directory environment requires Windows Server 2003, Windows 2000 Server, and Windows 2000 Professional or Windows XP Professional clients. The IP Security Policies extension of Group Policy Object Editor is used for managing IPSec in Active Directory environments. Policy based on Active Directory is not supported for computers running Windows XP Home Edition.

For computers that are not members of a domain, you can use the IP Security Policy Management snap-in to manage the local computer.

To successfully deploy IPSec, the following conditions must also be met:

- Each computer that must establish IPSec-secured communications has an assigned IPSec policy. The assigned policy must be compatible with the IPSec policy that is assigned to other computers with which it must communicate.
- Authentication is configured correctly and an appropriate authentication method is selected in the IPSec policy, so that mutual authentication can occur between IPSec peers.
- Routers, firewalls, or other filtering devices are configured correctly on all parts of the corporate network to permit IPSec protocol traffic, if IPSec-secured traffic must pass through these devices.
- IPSec policy compatibility issues must be considered if computers are running different versions of the Windows operating system (Windows Server 2003, Windows XP, and Windows 2000).
- For client computers that need to establish IPSec-secured connections with servers, servers must be adequately sized to support those connections. If necessary, IPSec hardware offload network adapters can be used.

This section discusses how IPSec Policy Extensions work. It presents information about IPSec architecture and how IPSec policies are stored.

**IPSec Policy Extension Architecture**

The basic components of IPSec in Windows Server 2003 include the Policy Agent, the Internet Key Exchange (IKE) module, the IPSec driver, and the TCP/IP driver, as shown in the following figure.

**IPSec Policy Extension Architecture**




The following table describes the IPSec policy-related components.

**IPSec Policy-Related Components and Descriptions**

| Component | Description |
|---|---|
| IP Security Policy Management Snap-in | This is the administrative tool for configuring IPSec policies in Active Directory or on the local computer. When the IP Security Policy Management snap-in loads as part of a GPO, it appears in the **Computer Configuration\Security Settings** node of Group Policy Object Editor(displayed as IP Security Policies). IP Security policies set from a GPO that is linked to a site, domain, organizational unit (OU), or local computer follow GPO inheritance and filtering rules. |

| | |
|---|---|
| Policy Agent | The Policy Agent controls the retrieval and distribution of IPSec policy and maintains the most current configured policy data for the IPSec Driver and IKE module. The Service Control Manager starts the Policy Agent, which runs in the context of the local security authority (LSA). The Policy Agent performs the following tasks when it is started (when the Local Security Authority is initialized):<br>1. Starts the IKE module.<br>2. Retrieves IPSec policies.<br>3. Determines the filter list order.<br>4. Delivers IP filters to the IPSec driver.<br>5. Delivers both Phase 1 and Phase 2 settings to IKE.<br>6. Polls for policy updates.<br>For more information about the Policy Agent tasks, see *Process Flow*, later in this section. |
| Internet Key Exchange (IKE) Module | The IKE module is used to establish a combination of mutually agreeable policy settings and keys that define the security services, protection mechanisms, and cryptographic keys between communicating peers. This combination is called a security association (SA). The IPSec driver in Windows 2000 and Windows Server 2003 uses the IKE module to protect all corresponding network traffic.<br>To create an SA between the two computers, the IETF has established a standard method of SA and key exchange resolution, which combines the Internet Security Association and Key Management Protocol (RFC 2408) and the Oakley Key Determination Protocol (RFC 2412). This standard method is IKE and is described in RFC 2409. |
| IPSec Driver | The IPSec driver receives the active IP filter list from the IPSec Policy Agent, and then attempts to match every inbound and outbound packet against the filters in the list. If a packet matches a filter, the IPSec driver applies the filter action. If a packet does not match any filters, the packet is passed back without modification to the TCP/IP driver to be received or transmitted.<br>If the filter action permits transmission, the packet is received or sent with no modifications. If the action blocks transmission, the packet is discarded. If the action requires the negotiation of security, main mode and quick mode SAs are negotiated. |
| TCP/IP Driver | The TCP/IP driver is the Windows 2000 implementation of the TCP/IP protocol. It starts the IPSec driver. |
| Policy Store | This maintains both IPSec policy descriptions and interfaces to applications and other tools that provide for policy data management. The policy store accesses IPSec policy data stored in either the local registry or Active Directory. |
| Local Registry | The local registry stores the locally configured IPSec policies, the local cache, and other IPSec settings. |
| Local Cache | The local cache stores IPSec policies after they are pulled from an Active Directory domain controller. The local cache is part of the registry. |
| Windows Management Instrumentation (WMI) | WMI is a component of the Microsoft Windows operating system and is the Microsoft implementation of Web-Based Enterprise Management (WBEM), which is an industry initiative to develop a standard technology for accessing management information in an enterprise environment. WMI uses the Common Information Model (CIM) industry standard to represent systems, applications, networks, devices, and other managed components. You can use WMI to automate administrative tasks in an enterprise environment. |
| IP Security Container | Each domain OU has an internal system container in Active Directory. IPSec policies created with IP Security Policiesare stored in this location. |

**IPSec Policy Extension Processes and Interactions**

This section discusses how IPSec Policies are processed, assigned, and stored.

The sequence of events is as follows:

1. During computer startup, the TCP/IP driver starts the IPSec driver.

2. The Service Control Manager starts the IPSec Policy Agent. The Policy Agent service performs the following tasks when it is initialized:

   - In preparation for loading the IPSec policy from either Active Directory or the local registry, the Policy Agent initializes the policy state and polling interval.

   - The policy agent loads the IPSec policy from the appropriate store. If boot time security is enabled, the IPSec driver mode is changed from the boot mode (block or stateful) to normal operational mode.

   - The policy agent retrieves the appropriate IPSec policy (if one has been assigned) from the Active Directory directory store or from the registry.

     To determine if IPSec policies exist, the Policy Agent checks the local registry and the Active Directory domain or OU to which the computer belongs. The Policy Agent checks the registry for the value that indicates the location of the computer's IPSec policy in Active Directory. If the computer is a member of a domain and an IPSec policy exists, this value was set during computer initialization. The registry key checked is:

     **HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\IPSec \GPTIPSECPolicy\DSIPSECPolicyPath**.

     If the computer is not a member of a domain, or if there is no IPSec policy in Active Directory, a value does not exist.

     If IPSec policies are found in Active Directory, in a Group Policy object, the Policy Agent retrieves the policies and stores them in the local cache. Note that the local Active Directory cache is not used for locally assigned policies. If IPSec policies are defined both locally and in Active Directory, the policies based on Active Directory override the local policies.

     If there are no IPSec policies in Active Directory or the registry, or if the IPSec Policy Agent cannot connect to Active Directory, the IPSec Policy Agent waits for policy to be assigned or activated.

3. The IKE module, which runs in the context of the local security authority (LSA) is started.

4. After the policy load process is completed, the Policy Agent begins a service loop.

   In Windows Server 2003, if an unexpected event occurs, the IPSec driver is put into Block mode and all network traffic is blocked. This is the network security component failing to a secure mode. This applies only to Windows Server 2003; Windows 2000 and Windows XP do not block traffic.

For more information about Windows Server 2003 IPSec, see "What Is IPSec?" in the *Networking Collection*.

**Assignment of IPSec Policy**

You can assign IPSec policy to a local computer and to computers in Active Directory sites, domains or organizational units. For computers that are not members of a domain, you can assign local IPSec policy.

The main method for building an IPSec policy is to use the **IP Security Policy Management** snap-in to create, modify, and activate IPSec policies, and then assign them to computers in either Active Directory (for domain members) or to the local computer (for computers that are not members of a domain).

When you assign an IPSec policy in a GPO, a Lightweight Directory Access Protocol (LDAP) pointer to the IPSec policy is recorded in the **IPSecOwnersReference** attribute of the GPO. When the GPO is applied to the computer, the IP Security Policies extension of Group Policy Object Editor reads the LDAP pointer from the GPO and writes it to the registry. When the IPSec service starts on the computer, it reads the LDAP pointer and queries Active Directory to obtain the IPSec policy settings. The IPSec service maintains a current local cache of the IPSec Policy by periodically polling for changes. The polling frequency

is based on a polling interval that is specified on the IPSec policy.

You can also create and assign persistent IPSec policy to secure a computer even if a local IPSec policy or an Active Directory IPSec policy cannot be applied. This policy adds to or overrides the local or Active Directory policy, and remains in effect regardless of whether other policies are applied. Persistent IPSec policies enhance security by providing a secure transition from computer startup to IPSec policy enforcement. Persistent policy also provides backup security in the event of an IPSec policy corruption, or if errors occur during the application of local or domain-based IPSec policy. To configure persistent policies, you must use the **netsh ipsec static set store location=persistent** command.

**Administering IPSec Policies with Group Policy Management Console**
IPSec policies are stored separately from the GPOs that have policies applied; therefore, administrators need to be aware of the following differences. Domain-based IPSec policies are stored in the IP Security Policiescontainer in the System container.

**Back-up and Restore of IPSec Policies**
Group Policy Management Console (GPMC) provides back-up and restore operations for Group Policy objects (GPOs). These operations help administrators maintain their Group Policy deployment in the event of a disaster. It is a recommended practice to regularly back up all GPOs. To provide consistency, when you use these tools, make sure that you include IPSec policies. However; GPMC tools do not store the IPSec policies themselves; therefore, you cannot use GPMC tools to back up and restore the actual IPSec policies. You can only back-up and restore information about whether the IPSec policies are assigned to Group Policy objects and, if so, to which Group Policy objects.

The IPSec policy objects in the IP Security Policies container in Active Directory can be backed up and restored by using the Export Policies and Import Policies commands in the IP Security Policiesconsole. The Export Policies command exports all IPSec policy objects from the policy store into one .IPSec file. The Import Policies command imports all IPSec policy objects in the .IPSec file into the destination policy store.

### Note

- Importing IPSec policy into Active Directory overwrites existing IPSec policy objects.

For information about GPMC, see "What Is Group Policy Management Console?" in this collection.

**Delegation of Rights to Manage IPSec Policies**
GPMC delegation of rights and security filtering permissions only apply to the GPO; they do not apply to the IPSec policy that is assigned in the GPO. Thus, delegation of edit rights in GPMC only enables a user to assign or unassign an existing IPSec policy in the specific GPO, and only if the user also has read access rights to the IPSec policy. To delegate rights to create, edit, or delete IPSec policies, administrators must perform the delegation on the IPSec policies container.

**Applied IPSec Policies**
You can use the GPMC Group Policy Results wizard to show which GPOs are applied to a computer, including the IPSec policy assignments. **Advanced View** in the computer node shows which IPSec policy is assigned to a specific computer.

On computers running Windows XP, IPSec does not provide Resultant Set of Policy (RSoP) information. GPMC Group Policy Results only shows the GPO being processed, but does not show which IPSec policy is assigned. To view the assigned IPSec policy, you can use **netdiag /test:ipsec.**

For information about netdiag, see "IPSec Policy Extension Tools and Settings."

**IPSec Order of Precedence**
Although most Group Policy settings are cumulative, you can assign only one IPSec policy to a computer at a time. If multiple IPSec policies have been assigned at different Active Directory levels, the last one applied is the one that takes effect. IPSec policy settings differ from other types of Group Policy settings in that IPSec policy settings from different OUs are not merged.

IPSec policy uses the same precedence sequence as other Group Policy settings, from lowest to highest. The following list outlines the default order in which policies are applied. You can override this order by using Group Policy options, including **Enforced**, **Block Policy Inheritance**, and the loopback policy processing option (the **User Group Policy loopback processing mode** policy setting):

- **Local GPO.** Each computer has one local GPO. If a computer is not a member of a domain, this is the only place where you can assign IPSec policy. You can assign an IPSec policy by editing the local GPO, but you can also assign a local policy directly in the IP Security Policy Management snap-in, outside of the Group Policy context, or by using the Netsh IPSec context. When IPSec policy assignments are made outside of the GPO context, the GPO cannot display the local IPSec policy that is assigned.

- **Site.** IPSec policies are infrequently assigned at the site level because this would require all computers within a site to

have the same security needs, which is unlikely. Furthermore, if the computer moves to another subnet — such as when a user travels to another location with a laptop that uses DHCP — different policies might be in effect, resulting in different security behaviors.

- **Domain.** Simple IPSec policies are often assigned at the domain level and then superseded by more specific IPSec policies, as required on various OUs.

- **OU.** Specific IPSec policies are assigned to groups of computers. This is the last policy applied under normal conditions and, therefore, the policy takes precedence. If an OU is nested within another OU, the IPSec policy assigned to the nested OU takes precedence.

**Persistence of IPSec Policy**
An IPSec policy that has been assigned might remain active (persist) even after a GPO to which it is assigned has been deleted. When you assign an IPSec policy in a GPO, a Lightweight Directory Access Protocol (LDAP) pointer to the IPSec policy is recorded in the IPSecOwnersReference attribute of the GPO. The IP Security Policies extension reads the LDAP pointer from the GPO and writes it to the registry when the GPO is applied to the computer. This registry entry can persist if you delete a GPO that contains an assigned IPSec policy. Although the IPSecOwnersReference LDAP pointer is updated when the GPO is deleted, it points to an unreadable object (the location of the deleted objects container), which IPSec services cannot access. As a result, IPSec policy fails and an error is logged.

To prevent an IPSec policy from persisting, you must unassign the IPSec policy in the GPO beforeyou delete the GPO. Unassigning the IPSec policy deletes the LDAP pointer from the GPO.

**Storage of IPSec Policy Settings**
Domain-based IPSec policy objects are stored in the IP Security Policies container in Active Directory, which is separate from the GPOs to which the IPSec policies are applied. The domain administrator must grant permissions to the IP Security Policies container for other delegated administrators to administer IPSec policies. Standard delegation tools cannot be used to delegate permissions to administer IPSec policies. Instead, domain administrators must use Active Directory Service Interfaces (ADSI) Edit for this purpose.

ADSIEdit is a Microsoft Management Console (MMC) snap-in that domain administrators can use to edit objects in the Active Directory database. When domain administrators delegate permissions to others to administer IPSec policies, the delegated administrators must have Full Control permissions to all IPSec policy objects in the IP Security Policies container. After an IPSec administrator creates an IPSec policy, a member of the Group Policy Creator Owners group or other delegated owner of the GPO can assign the IPSec policy to the appropriate GPOs.

### Note

- To modify ACLs on the IP Security Policies container, domain administrators should use the tools that are provided with the Windows 2000 or Windows Server 2003 operating system CD. When domain administrators modify ACLs on the IP Security Policies container, they must specify that all modifications are also propagated to all child objects. Incorrect modification of the ACLs on the IP Security Policies container, failure to propagate modifications to the child objects of this container, or both, can result in the failure of IPSec policy settings to be properly applied.

On standalone computers, IPSec policy is stored in the local registry.

ADSIEdit is one of the Windows Server 2003 Support Tools which are included on the Windows Server 2003 operating system CD. You install these tools separately by running the Support Tools Setup program, **Suptools.msi**, from the **\Support\Tools** folder.

For more information about permissions for the IPSec Policies container, see article 329194, “Permissions on the IPSec Policy Store” in the Microsoft Knowledge Base. To find this article, see the Microsoft Knowledge Baselink on the Web Resources page.

**IP Security Container Default Permissions**
In Windows Server 2003, the default discretionary access control lists (DACLs) for the IP Security container are described in the following tables. The IPSec container owner is the Domain Admin, and group membership is Domain Admin.

The following table lists access rights related to the IPSec container and describes their values.

### IPSec Container Access Rights Values and Descriptions

| Access Right Value | Description |
|---|---|

| | |
|---|---|
| **ADS_RIGHT_DS_READ_PROP** | The right to *read* properties of the object. The **ObjectType** member of an access control entry (ACE) can contain a GUID that identifies a property set or property. If **ObjectType** does not contain a GUID, the ACE controls the right to read all of the object properties. |
| **ADS_RIGHT_DS_WRITE_PROP** | The right to *write* properties of the object. The **ObjectType** member of an ACE can contain a GUID that identifies a property set or property. If **ObjectType** does not contain a GUID, the ACE controls the right to write all of the object properties. |
| **ADS_RIGHT_DS_CONTROL_ACCESS** | The right to perform an operation controlled by an extended access right. The **ObjectType** member of an ACE can contain a GUID that identifies the extended right. If **ObjectType** does not contain a GUID, the ACE controls the right to perform all extended right operations associated with the object. |
| **ADS_RIGHT_DS_LIST** | The right to list a particular object. If the user is not granted such a right, and the user does not have ADS_RIGHT_ACTRL_DS_LIST set on the object parent, the object is hidden from the user. This right is ignored if the third character of the **dSHeuristics** property is '0' or not set. |
| **ADS_RIGHT_DS_LIST_OBJECT** | The right to list a particular object. If the user is not granted such a right, and the user does not have ADS_RIGHT_ACTRL_DS_LIST set on the object parent, the object is hidden from the user. This right is ignored if the third character of the **dSHeuristics** property is '0' or not set. |
| **ADS_RIGHT_DS_CREATE_CHILD** | The right to create children of the object. The **ObjectType** member of an ACE can contain a GUID that identifies the type of child object whose creation is controlled. If **ObjectType** does not contain a GUID, the ACE controls the creation of all child object types. |
| **ADS_RIGHT_DS_DELETE_CHILD** | The right to delete children of the object. The **ObjectType** member of an ACE can contain a GUID that identifies a type of child object whose deletion is controlled. If **ObjectType** does not contain a GUID, the ACE controls the deletion of all child object types. |
| **ADS_RIGHT_READ_CONTROL** | The right to read data from the security descriptor of the object, not including the data in the SACL. |
| **ADS_RIGHT_WRITE_DAC** | The right to modify the discretionary access-control list (DACL) in the object security descriptor. |
| **ADS_RIGHT_WRITE_OWNER** | The right to assume ownership of the object. The user must be an object trustee. The user cannot transfer the ownership to other users. |
| **ADS_RIGHT_DELETE** | The right to delete the object. |
| **ADS_RIGHT_DS_DELETE_TREE** | The right to delete all children of this object, regardless of the permissions of the children. |
| **ADS_RIGHT_DS_SELF** | The right to perform an operation controlled by a validated write access right. The **ObjectType** member of an ACE can contain a GUID that identifies the validated write. If **ObjectType** does not contain a GUID, the ACE controls the rights to perform all valid write operations associated with the object. |

The following table lists the access control rights for the IPSec container.

**DACLs for IPSec Container**

| DACLs | Access Rights |
|---|---|

| | |
|---|---|
| Allow domain computers | **ADS_RIGHT_DS_READ_PROP**<br>**ADS_RIGHT_DS_LIST**<br>**ADS_RIGHT_DS_LIST_OBJECT**<br>**ADS_RIGHT_READ_CONTROL** |
| Allow Group Policy Creator Owner | **ADS_RIGHT_DS_READ_PROP**<br>**ADS_RIGHT_DS_LIST**<br>**ADS_RIGHT_DS_LIST_OBJECT**<br>**ADS_RIGHT_READ_CONTROL** |
| Allow Domain Admin | **ADS_RIGHT_DS_READ_PROP**<br>**ADS_RIGHT_DS_WRITE_PROP**<br>**ADS_RIGHT_DS_CONTROL_ACCESS**<br>**ADS_RIGHT_DS_LIST**<br>**ADS_RIGHT_DS_LIST_OBJECT**<br>**ADS_RIGHT_DS_CREATE_CHILD**<br>**ADS_RIGHT_DS_DELETE_CHILD**<br>**ADS_RIGHT_READ_CONTROL**<br>**ADS_RIGHT_WRITE_DAC**<br>**ADS_RIGHT_WRITE_OWNER**<br>**ADS_RIGHT_DS_SELF** |
| Allow Local SYSTEM | **ADS_RIGHT_DS_READ_PROP**<br>**ADS_RIGHT_DS_WRITE_PROP**<br>**ADS_RIGHT_DS_CONTROL_ACCESS**<br>**ADS_RIGHT_DS_LIST**<br>**ADS_RIGHT_DS_LIST_OBJECT**<br>**ADS_RIGHT_DS_CREATE_CHILD**<br>**ADS_RIGHT_DS_DELETE_CHILD**<br>**ADS_RIGHT_READ_CONTROL**<br>**ADS_RIGHT_WRITE_DAC**<br>**ADS_RIGHT_WRITE_OWNER**<br>**ADS_RIGHT_DELETE**<br>**ADS_RIGHT_DS_DELETE_TREE**<br>**ADS_RIGHT_DS_SELF** |

The following changes are implemented in Windows Server 2003, and they apply to both new and upgrade installations of Windows Server 2003:

- An inheritable DACL is included with the IP Security container.
- An empty DACL is included with the built-in IPSec policies. Because the DACL is empty, objects can inherit permissions

from the IP Security container and other parent containers in Active Directory.

- The default DACL for all IPSec objects is changed so that all new objects that you create have an empty DACL and can inherit permissions. As a result, IPSec policies that are created from Windows 2000, Windows XP or Windows Server 2003 management clients have the new permissions.

To view these ACLs, you can use the ADSIEdit.exe tool (described earlier in *Storage of IPSec Policy Settings*) or Ldp.exe. The Ldp.exe tool provides an interface to perform Lightweight Directory Access Protocol (LDAP) operations against Active Directory. Ldp.exe and ADSIEdit.exe are two of the Windows Server 2003 Support Tools, which are included in the Windows Server 2003 operating system CD. You can install these tools by running the Support Tools Setup program, Suptools.msi, from the **Support\Tools** folder.

For more information about delegating IPSec permissions, see "Delegating Permissions to Modify Active Directory Based IPSec Policy" in the *Using Microsoft Windows IPSec to Help Secure an Internal Corporate Network Server* article available at the Microsoft Download Center.

For more information about control access rights, search for "Control Access Rights" in the *Directory Services SDK Documentation* in the Microsoft Platform SDK on MSDN.

**IPSec Policy Polling Interval**

The general IPSec policy settings include a polling interval, **Policy change poll interval**, which specifies the number of minutes between consecutive polls for changes in IPSec policies based on Active Directory. The default value for the **Policy change poll interval** is 180 minutes. This polling does not detect a change in domain or organizational unit membership, or the assigning or unassigning of a new policy; these events are detected when the Winlogon service polls for changes in Group Policy, which occurs every 90 minutes by default.

To generate an immediate update of changes in IPSec policy settings, instead of waiting for the next polling period, you can use the **Gpupdate** command, by typing the following at the command prompt: **gpupdate /target:computer**.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "How Core Group Policy Works" in the Group Policy Collection.
- "What Is IPSec?" in the Networking Collection.
- "Deploying IPSec" in the *Deploying Network Services* guide of the Microsoft Windows Server 2003 Deployment Kit.
- "Control Access Rights" in the Directory Services SDK Documentation in the Microsoft Platform SDK on MSDN.

# IPSec Policy Extension Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP Security Policy Extension Tools and Settings**

**In this subject**


- Resultant Set of Policies
- Tools for Viewing IPSec Policy Assignment on Computers Running Windows Server 2003
- Tools for Viewing IPSec Policy Assignment on Computers Running Windows 2000
- Tools for Viewing IPSec Policy Assignment on Computers Running Windows XP
- Tools for Managing and Monitoring IPSec Policy
- Network Ports Used by IPSec
- Related Information


This section is an overview of IPSec tools, Resultant Set of Policies (RSoP) in particular, which are used to determine which IPSec policy settings are currently in effect for a computer or a user. RSoP is also used to assess how policy settings would affect computers or users if a specific Group Policy object were applied to them. This section also describes the Windows Server 2003 command-line tools for configuring and analyzing IPSec policy settings.

**Resultant Set of Policies**

Windows XP and the Windows Server 2003 family support an enhanced Group Policy infrastructure that utilizes Windows Management Instrumentation (WMI) to collect Group Policy-related data for planning and troubleshooting Group Policy. This structure is the Resultant Set of Policy (RSoP), a query engine that polls existing policy settings and planned policy settings, and then reports the results of those queries. RSoP polls existing policies based on site, domain, domain controller, and organizational unit. RSoP gathers this information from the Common Information Management Object Model (CIMOM) database (also known as CIM-compliant object repository) by using WMI.

Administrators can use RSoP in one of two modes. To determine which policy settings are in effect for a particular computer or user, administrators use Resultant Set of Policy (RSoP) Logging Mode; to evaluate how policy settings would affect a computer or user if a specific Group Policy object were applied to users or computers, they use Resultant Set of Policy (RSoP) Planning Mode.

In Windows Server 2003, you can use the **Group Policy Modeling** node in the Group Policy Management Console (GPMC) to access the RSoP Planning Mode capabilities of Windows Server 2003. By using this feature, administrators can simulate policy settings applied to users and computers with Group Policy before actually applying the policy settings. You can also use the **Group Policy Results** node of GPMC to access the RSoP Logging Mode capabilities. Group Policy Results represents the actual resultant set of policy that was applied to a given user and computer. This information is obtained by directly querying the target user or computer. Each subnode represents a different RSoP query for a given user/computer combination.

**RSOP_IPSECPolicySetting**

The **RSOP_IPSECPolicySetting** WMI class represents the policy data for the IPSec extension. The **RSOP_IPSECPolicySetting** class is derived from **RSOP_PolicySetting**, and is included in Windows Server 2003. This class was available starting with Windows Server 2003 family and might not be available in future versions of Windows.

For more detailed information about RSoP and WMI and to download SDKs, see the Microsoft Platform SDK page.

For more information about GPMC, see the Group Policy Management Console page.

**Tools for Viewing IPSec Policy Assignment on Computers Running Windows Server 2003**

IPSec provides support for RSoP to enhance deployment and troubleshooting of IPSec Policy in Windows Server 2003. You can run an RSoP Logging-Mode query or an RSoP Planning-Mode query to view detailed settings (the filter rules, filter actions, authentication methods, tunnel endpoints, and connection types) for the IPSec policy that is currently being applied, or would

be applied if you deployed that IPSec policy.

The RSoP console displays the following information for each Group Policy object (GPO) that contains an IPSec policy assignment: the name of the IPSec policy, the name of the GPO to which the IPSec policy is assigned, the IPSec policy precedence, and the name of the site, domain, and organizational unit to which the GPO containing the IPSec policy applies (that is, the scope of management for the GPO). The lower the number in the **Precedence** column, the higher the precedence of the IPSec policy. A precedence of 1 indicates the IPSec policy that is being applied.

To provide support for RSoP in Windows Server 2003, the **RSOP_IPSECPolicySetting** WMI class was created. For more information, see *RSOP_IPSECPolicySetting* earlier in this document.

**RSoP and the Group Policy Management Console**
If you have installed Group Policy Management Console (GPMC), you can find out which policies are currently applied to a specific computer, including the IPSec policy assignments, by using the **Group Policy Results** node in GPMC. This runs the GPMC **Group Policy Results** wizard, which replaces the Resultant Set of Policy (RSoP) Logging Mode capabilities. Running the wizard shows which IPSec policy is assigned to a specific computer by right-clicking on the computer node, and then selecting **Advanced View.**

GPMC is available as a download from the Microsoft Web site. For more information about GPMC, see the Group Policy Management Console page.

**Netsh Commands for IPSec**
**Netsh** is a command-line scripting utility that you can run on a local or remote computer to display or modify the network configuration of a computer that is currently running. **Netsh** also provides a scripting feature so you can run a group of commands in batch mode against a specified computer, and it can also save a configuration script in a text file for archival purposes or to help you configure other servers.

**Netsh** interacts with other operating system components using dynamic-link library (DLL) files. **Netsh** helper DLLs provide an extensive set of features called a context, which is a group of commands specific to a networking component. These contexts extend the functionality of **netsh** by providing configuration and monitoring support for one or more services, utilities, or protocols. For example, **Nshipsec.dll** provides **netsh** the context and set of commands necessary to configure and manage IPSec policies and state.

To run a **netsh** command, you must start **netsh** from the **Cmd.exe** prompt and change to the context that contains the command you want to use. The contexts that are available to you depend on which networking components you have installed.

For computers running members of the Windows Server 2003 family, you can use the **netsh** commands for Internet Protocol security (**netsh ipsec** context) as an alternative to the console-based management and diagnostic capabilities provided by the IP Security Policy Management and the IP Security Monitor snap-ins. By using the **netsh** commands for IPSec, you can script IPSec configuration, and change the IPSec configuration for troubleshooting. You can also use this command to extend the security and manageability of IPSec. For example, you can use the **netsh** commands for IPSec to enable IPSec driver event logging, set default traffic exemptions, and configure computer startup security.

**Note**

- For computers running Windows 2000, use **Ipsecpol.exe**, which is provided with the Windows 2000 Server Resource Kit.
- For more information about **Ipsecpol.exe**, see "*Ipsecpol.exe: Internet Protocol Security Policies Tool*" on the download site for Windows 2000 Server Resource Kit tools.

The **Netsh ipsec** command supports both static and dynamic mode:

**Netsh ipsec static mode commands.**
You can use the **netshipsec** static commands to perform the same management tasks that you can perform by using the IP Security Policy Management console. By using these commands, you can create, modify, and assign IPSec policies without immediately affecting the configuration of the active IPSec policy.

**Netsh ipsec dynamic mode commands.**
You can use the **netsh ipsec** dynamic commands to display the active state of IPSec and to immediately affect the configuration of the active IPsec policy. These commands directly configure the security policy database (SPD). Changes that you make to an IPSec policy while using these commands take effect only while the IPSec service is running. If the IPSec service is stopped, the dynamic policy settings are discarded. You can use the **netsh ipsec** dynamic mode commands to perform the monitoring tasks of the IPSec Monitor snap-in. The **netsh ipsec** dynamic command allows the active state of IPSec to be

displayed, and it also allows modification of the state (unlike the IPSec Monitor tool, which does not allow modification of the state).

You can run these commands from the Windows Server 2003 command prompt or from the command prompt for the **netsh ipsec** context. For these commands to work at the Windows Server 2003 command prompt, you must type **netsh ipsec** before typing ipsec commands (such as **add filter** or **add rule**) and parameters.

To find more information about **netsh ipsec**, see "Command Line References" in Tools and Settings.

**Gpupdate**
This command-line tool refreshes local Group Policy settings and Group Policy settings that are stored in Active Directory, including security settings. This command supersedes the now obsolete **/refreshpolicy** option for the **secedit** command.

**Syntax**

```
gpupdate [/target:{computer | user}] [/force] [/wait:Value] [/logoff]
[/boot]
```

**Parameters**
The following table lists the **Gpupdate** parameters.

**Gpupdate parameters**

| Parameter | Description |
|---|---|
| **/target:**{**computer** \| **user**} | Processes only the *Computer* settings or the current *User* settings. By default, both the computer settings and the user settings are processed. |
| **/force** | Ignores all processing optimizations and reapplies all settings. |
| **/wait:** *Value* | Number of seconds that policy processing waits to finish. The default is 600 seconds. 0 equals no wait, and -1 equals wait indefinitely. |
| **/logoff** | Logs off after the refresh has completed. This is required for those Group Policy client-side extensions that do not process on a background refresh cycle but that do process when the user logs on, such as user Group Policy Software Installation and Folder Redirection. This option has no effect if there are no extensions called that require the user to log off. |
| **/boot** | Restarts the computer after the refresh has completed. This is required for those Group Policy client-side extensions that do not process on a background refresh cycle but that do process when the computer starts up, such as computer Group Policy Software Installation. This option has no effect if there are no extensions called that require the computer to be restarted. |
| **/?** | Displays help at the command prompt. |

**Examples**
The following examples show how you can use the **gpupdate** command:

**gpupdate**

**gpupdate /target:computer**

**gpupdate /force /wait:**100

**gpupdate /boot**

**Tools for Viewing IPSec Policy Assignment on Computers Running Windows 2000**
On computers running Windows 2000, you can use the Group Policy Results tool, **Gpresult.exe,** to display information about how Group Policy affects both the currently logged-on user and the computer.

**Gpresult** displays Group Policy information regarding:

- The last time that Group Policy was applied and the domain controller that applied the Group Policy setting for the user and the computer.

- A list of all applied GPOs and their details, including a summary of the extensions that each GPO contains.

- Registry settings that are applied and the details.
- Folders that are redirected and their details.
- Disk quota information.
- Internet Protocol (IP) Security settings.
- Scripts.

**Gpresult.exe Syntax**

Gpresult.exe uses the following syntax:

```
gpresult [/v] [/s] [/c] [/u]
```

The following table lists the **Gpresult** parameters.

**Gpresult.exe parameters**

| Parameter | Explanation |
|---|---|
| **/v** | Use this parameter to run Gpresult.exe in verbose mode. When you use this parameter, the following information is displayed (in additional to the information that is typically displayed):<br>• A list of the user's security rights.<br>• GPO details including globally unique identifier (GUID), friendly name, version, and source.<br>• Details for the following Group Policy extensions:<br>1. Administrative Templates (registry-based policy settings)<br>2. Application management<br>3. Disk quotas<br>4. Folder redirection<br>5. IP Security<br>6. Scripts |
| **/s** | Use this parameter to run Gpresult.exe in super-verbose mode. When you use this parameter, the following information is displayed (in addition to the information that is typically displayed):<br>• Binary values of binary registry settings (when applicable)<br>• A detailed list of the programs that are displayed in the Add or Remove Programs tool in Control Panel<br>• The Group Policy container and Group Policy template version numbers of the GPO |
| **/c** | Use this parameter to display information about computer settings only. |

| | |
|---|---|
| **/u** | Use this parameter to display information about user settings only. |

The **Gpresults** command-line tool is available in the Windows 2000 Resource Kit. To download this tool, see Windows 2000 Server Resource Kit tools Web site.

**Gpotool.exe**

You can use Gpotool.exe on a domain controller to check Group Policy object integrity and monitor policy replication. GpoTool performs the following tasks:

- Checks Group Policy object consistency. The tool reads mandatory and optional directory services properties such as version, friendly name, extension globally unique identifiers (GUIDs), and SYSVOL data (Gpt.ini). GpoTool also compares directory services and SYSVOL version numbers and performs other consistency checks. If the extensions property contains any GUID, the functionality version must be 2 and the user/computer version must be greater than 0.

- Checks Group Policy object replication. The tool compares the Group Policy object instances from each domain controller.

- Displays information about a particular Group Policy object, including properties that cannot be accessed by means of the Group Policy snap-in, such as functionality version and extension GUIDs.

- Browses Group Policy objects. A command-line option can search based on either friendly name or GUID. A partial match is also supported for both name and GUID.

- Designates preferred domain controllers. By default, all available domain controllers in the domain are used, but this can be overwritten with a list of domain controllers supplied from the command line.

- Provides cross-domain support. You can use the command-line option to check Group Policy objects in different domains.

- Runs in verbose mode. If all the Group Policy objects are acceptable, the tool displays a validation message. If errors occur, the tool displays information about the corrupted Group Policy objects. A command-line option can turn on verbose information about each Group Policy object that is being processed.

**GpoTool Syntax**

```
gpotool [/gpo:GPO[,GPO]] [/domain:DNSname] [/dc:DC[,DC]] [/checkacl]
[/verbose] [/new:GPO[,GPO...]] [/del:GPO[,GPO...]] [/?] [/help]
```

Where:

**/gpo:** *GPO*[,*GPO*] processes preferred Group Policy objects *GPO*[,*GPO*]. Partial GUID and friendly name matches are accepted for *GPO*. If *GPO* is not specified, the tool processes all Group Policy objects in the domain.

**/domain:** *DNSname* specifies the DNS name for the domain hosting the Group Policy objects. If a name is not specified, the user's domain is used.

**/dc:** *DC*[,*DC*] finds the preferred list of domain controllers *DC*[,*DC*]. If not specified, finds all controllers in the domain.

**/checkacl** verifies the SYSVOL ACL. For faster processing, this step is skipped by default.

**/verbose** displays detailed information.

**/new:** *GPO*[,*GPO*...] create new Group Policy objects with the specified friendly names *GPO*[,*GPO*...].

**/del:** *GPO*[,*GPO*...] deletes Group Policy objects with the specified friendly names *GPO*[,*GPO*...].

**/?** or **/help** displays GpoTool syntax.

**Tools for Viewing IPSec Policy Assignment on Computers Running Windows XP**

On computers running Windows XP, IPSec does not provide RSoP information. GPMC **Group Policy Results** shows the GPO being processed, but it does not show which IPSec policy is assigned. To view the assigned IPSec policy, you must use the **netdiag /test:ipsec** command. **Netdiag.exe** is a command-line tool that you can use to display information about IPSec

policies and statistics, reports network configuration, and tests basic networking capabilities and domain-based functionality.

Specifically, **Netdiag.exe** displays information about IPSec policy assignments, including the name of the active IPSec policy, the name of the Group Policy object that assigned the policy, and the policy path.

For Windows XP, **Netdiag.exe** is available for installation from the Windows XP CD. You install it by running the **Setup.exe** file from the **\Support\Tools** folder, and choosing the **Complete** setup option.

**Netdiag.exe**
On computers running Windows XP, you can use **Netdiag.exe** to view information about IPSec policies and statistics, report network configuration, and test basic networking capabilities and domain-based functionality.

**Syntax**
**Netdiag** uses the following syntax:

```
netdiag [/q] [/v] [/l] [/debug] [/d:domain_name] [/fix] [/dcaccountenum]
[/test:test_name] [/skip:test_name]
```

**Parameters**
You can use the following parameters with **Netdiag**:

| Parameter | Explanation |
|---|---|
| **/q** | Use this parameter to specify quiet output and display errors only. |
| **/v** | Use this parameter to run Netdiag in verbose mode and display information about the actions that are performed. |
| **/l** | Use this parameter to write output to the Netdiag.log file. The Netdiag.log file is created in the same folder in which you run Netdiag. |
| **/debug** | Use this parameter to run Netdiag in debug mode. This parameter specifies a more verbose output than when you use the **/v** parameter. |
| **/d:domain_name** | Use this parameter to locate a domain controller in the specified domain. |
| **/fix** | Use this parameter to correct minor issues. |
| **/dcaccountenum** | Use this parameter to enumerate domain controller computer accounts. |

| | |
|---|---|
| **/test:** *test_name* | Use this parameter to specify the test or tests that you want to run, where *test_name* can be any of the following values:<br>• **Autonet**: Automatic Private IP Addressing (APIPA) address test<br>• **Bindings**: Bindings test<br>• **Browser**: Redir and Browser test<br>• **DcList**: Domain controller list test<br>• **DefGw**: Default gateway test<br>• **DNS**: Domain Name Service (DNS) test<br>• **DsGetDc**: Domain controller discovery test<br>• **IpConfig**: Internet Protocol (IP) address configuration test<br>• **IpLoopBk** : IP address loopback ping test<br>• **IPSec**: IP Security Protocol (IPSec) security test<br>• **IPX**: Internetwork Packet Exchange (IPX) test<br>• **Kerberos**: Kerberos Test<br>• **Ldap**: Lightweight Directory Access Protocol (LDAP) test<br>• **Member**: Domain membership test<br>• **Modem**: Modem diagnostics test<br>• **NbtNm**: NetBIOS over TCP/IP (NetBT) name test<br>• **Ndis**: Netcard queries test<br>• **NetBTTransports**: NetBT transports test<br>• **Netstat**: Netstat information test<br>• **NetWare**: NetWare test<br>• **Route**: Routing table test<br>• **Trust**: Trust relationship test<br>• **WAN**: Wide Area Network (WAN) configuration test<br>• **WINS**: Windows Internet Name Service (WINS) service test<br>• **Winsock**: Winsock test |

**Example**

To use **Netdiag** to display the currently active Internet Protocol security (IPSec) policy, type the following line:

**netdiag /test:ipsec /debug**

For more information about **Netdiag**, see the Microsoft Knowledge Base page, and search for Knowledge Base article number *321708 How to: Use the Network Diagnostics Tool (Netdiag.exe) in Windows 2000*.

**Tools for Managing and Monitoring IPSec**

This section describes tools for administering and monitoring IPSec policies:

- IP Security Policy Managementsnap-in. You can apply IPSec policies to domains, sites, organizational units, or any Group Policy object in Active Directory, as well as to the local computer. To create, modify, and assign IPSec policies, you can use the IP Security Policy Management snap-in.

- IP Security Monitorsnap-in. You can use this tool to view details about IPSec policy settings as they are applied to a computer.

**IP Security Policy Management**

Microsoft Windows XP and the Windows Server 2003 family provide the IP Security Policy Management snap-in, which you can use to define IPSec policies for computers through Active Directory (for domain members) or on the local computer (for computers that do not belong to domains). By using IP Security Policy Management, you can create IPSec policies to meet the security requirements of a computer, application, organizational unit, domain, site, or global corporation.

**IP Security Monitor**

On computers running Windows XP and the Windows Server 2003 family of operating systems, you can use the **IP Security Monitor** snap-in to view IPSec policy settings as they are applied to the computer. This tool monitors IPSec SAs, rekeys, negotiation errors, and other IPSec statistics.

### Note

- On computers running Windows 2000, you use the **Ipsecmon.exe** program to view IPSec policy settings.

You use IP Security Monitor to perform the following tasks:

- Monitor IPSec information for local computer and remote computers.

- View details about active IPSec policies, including the name, description, date last modified, store, path, organizational unit, and Group Policy object name.

- View the following IPSec details in main mode and quick mode:

  - Generic filters and specific filters

  - Statistics

  - Security associations

- Customize refresh rates, and use DNS name resolution for filter and security association output.

- Search for specific main mode or quick mode filters that match any source or destination IP address, a source or destination IP address on your local computer, or a specific source or destination IP address.

**Network Ports Used by IPSec**

IPSec uses the following ports:

| IPSec protocol | Uses |
|---|---|
| IPSec ESP | Protocol 50. |
| IPSec AH | Protocol 51. |

| | |
|---|---|
| ISAKMP | UDP port 500. |
| IPSec NAT-T | Source port of **Any** or UDP port 4500 (a source port of **Any** might be used because the network address translator might translate source port UDP 4500 to a different source port).<br><br>Destination port of UDP 4500. |

**Related Information**

The following resources contain additional information that is relevant to this section.

- "What Is IPSec?" in the Networking collection.
- For more information about RSoP and WMI and to download SDKs, see the Microsoft Platform SDK page.
- For more information about GPMC, see the Group Policy Management Console page.
- The Tools and Settings Collection.

# Software Restriction Policies Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft Windows XP and the Windows Server 2003 family include software restriction policies, a new Group Policy-based feature that identifies software programs running on computers in a domain, and controls the ability of those programs to run. Software restriction policies are part of the Microsoft security and management strategy to assist enterprises in increasing the reliability, integrity, and manageability of their computers.

**In this subject**


- What Are Software Restriction Policies?
- How Software Restriction Policies Work
- Software Restriction Policies Tools and Settings

# What Are Software Restriction Policies?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Are Software Restriction Policies?**

**In this section**


- Protecting Computer System Integrity
- Dependencies on Other Operating System Components
- Related Information


Software restriction policies provide administrators with a Group Policy-driven mechanism to identify software and control its ability to run on the local computer. These policies can be used to protect computers running Microsoft Windows XP Professional against known conflicts and safeguard the computers against security threats such as malicious viruses and Trojan horse programs. You can also use software restriction policies to create a highly restricted configuration for computers, in which you allow only specifically identified applications to run. Software restriction policies are integrated with Microsoft Active Directory and Group Policy. You can also create software restriction policies on stand-alone computers.

Software restriction policies are trust policies, which are regulations set by an administrator to restrict scripts and other code that is not fully trusted from running.

**Protecting Computer System Integrity**

Business users collaborate by using e-mail, instant messaging, and peer-to-peer applications. As these collaborations increase, especially with the use of the Internet in business computing, so do the threats from malicious code, such as worms, viruses, and malicious user or attacker threats.

Users might receive hostile code in many forms, ranging from native Windows executable files (.exe files), to macros in documents (such as .doc files), to scripts (such as .vbs files). Malicious users or attackers often use social engineering methods to get users to run code containing viruses and worms. (Social engineering is a term for tricking people into revealing their password or some form of security information.) If such code is activated, it can generate denial-of-service attacks on the network, send sensitive or private data to the Internet, put the security of the computer at risk, or damage the contents of the hard disk drive.

IT organizations and users must be able to determine which software is safe to run and which is not. With the large numbers and forms that hostile code can take, this becomes a difficult task.

To help protect their network computers from both hostile code and unknown or unsupported software, organizations can implement software restriction policies as part of their overall security strategy.

**Group Policy-Based Software Restrictions**

Administrators can use software restriction policies to define which applications are allowed or not allowed to run on a target computer. The configuration is deployed in Group Policy objects.

To help organizations address the problem of unknown code, administrators can use software restriction policies to perform the following tasks:

- Fight computer viruses
- Regulate which ActiveX controls can run
- Run only digitally signed scripts
- Enforce that only approved software is run on system computers
- Create a highly restricted configuration for a computer (for example, stipulate that only specific programs are allowed to run)

**Scenarios for Using Software Restriction Policies**
Administrators can use software restriction policies for the following tasks:

- Define what is trusted code
- Design a flexible Group Policy for regulating scripts, executable files, and ActiveX controls

Software restriction policies are enforced by the operating system and by applications (such as scripting applications) that comply with software restriction policies.

Specifically, administrators can use software restriction policies for the following purposes:

- Specify which software (executable files) can run on clients
- Prevent users from running specific programs on shared computers
- Specify who can add trusted publishers to clients
- Set the scope of the software restriction policies (specify whether policies affect all users or a subset of users on clients)
- Prevent executable files from running on the local computer, organizational unit (OU), site, or domain. This would be appropriate in cases when you are not using software restriction policies to address potential issues with malicious users.

**Advantages of Software Restriction Policies**
Using software restriction policies provides the following advantages:

**Administrators can enforce software restriction policies in domains or on the local computer.**
The software restriction policies can be implemented in a large enterprise with various types of computers and applications, and can also be applied in a stand-alone environment for computers that are not members of a domain. Software restriction policies leverage Active Directory and Group Policy for manageability. The software restriction policies are stored in a GPO. Administrators create the software restriction policy, and then define which applications are trusted and which are not. The software restriction policy is enforced at run time and users do not receive prompts allowing them to choose whether to run executable files.

**Software restriction policies apply to multiple types of files.**
Software restriction policies provide control over Microsoft Visual Basic Scripting Edition (VBScript), JScript, and other scripting languages. Software restriction policies also integrate with the Windows Installer feature to provide control over the types of software packages (.msi files) that can be installed on clients. This feature includes an application programming interface (API), which you can use to coordinate the software restriction policy run time with other run times.

**Software restriction policies provide flexibility.**
Administrators can prohibit unauthorized scripts from running, regulate Microsoft ActiveX controls, or lockdown client computers.

**Software restriction policies use strong cryptography to identify software.**
Software restriction policies can identify software by using a hash or a certificate.

**Dependencies on Other Operating System Components**
Software restriction policies rely on the following operating system components:

**Active Directory.**
The Windows-based directory service, Active Directory, stores information about objects on a network and makes this information available to administrators and users. By using Active Directory, you can view and manage network objects on the network from a single location, and users can access permitted network resources by using a single logon.

**Group Policy.**
The infrastructure within Active Directory, Group Policy, enables directory-based configuration management of user and computer settings on computers running Windows Server 2003, the Windows 2000 family, and the Microsoft Windows XP Professional operating systems. By using Group Policy, you can define configurations for groups of users and computers, including policy settings for Windows Server 2003 registry-based policies, software installation, scripts, folder redirection, Remote Installation Services, Microsoft Internet Explorer maintenance, and security.

**Windows Installer.**

Windows Installer is an installation and configuration service that reduces the total cost of ownership. Windows installer is included in Microsoft Windows Server 2003 family, Windows 2000, Windows XP, and Windows Millennium Edition (Windows Me). It is also provided as a service pack to Microsoft Windows NT version 4.0, Windows 98, and Windows 95. Windows Installer version 2.0 is integrated with Software Restriction Policies in Microsoft Windows XP and Windows Server 2003. The Windows Installer version 2.0 only installs packages allowed to run at the unrestricted level.

Patches or transforms must also be allowed to run at the unrestricted level. If a package, patch, or transform is not configured to run fully trusted (unrestricted), the Windows Installer displays an error message and logs an entry in the application event log. Software restriction policy is evaluated the first time an application is installed, when a new patch is applied, and when the installation package is re-cached.

Msiexec.exe is the executable program that interprets packages and installs products.

**Authenticode and WinVerifyTrust.**
These components are used to process signed executable files. Microsoft Authenticode, which is based on industry standards, allows developers to include information about themselves and their code with their programs through the use of digital signatures. The WinVerifyTrust function performs a trust verification action on a specified object.

The following figure illustrates the interactions of software restriction policies, Active Directory, and Group Policy.

**Software Restriction Policies and Related Components**




**Related Information**
The following resource contains additional information that is relevant to this section.

- "What Is Core Group Policy?" in the Group Policy Collection.

# How Software Restriction Policies Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Software Restriction Policies Work**

**In this section**


- Software Restriction Policies Architecture
- Software Restriction Policies Overview
- Software Restriction Policies Processes and Interactions
- Related Information


Software restriction policies provide administrators with a policy-driven mechanism for identifying the software programs running on computers in a domain and for controlling the ability of those programs to run.

Deployment based on Active Directory requires Windows Server 2003, and Windows XP Professional clients, specifically:

- Windows Server 2003 with Active Directory installed and dynamic DNS properly configured.
- Windows XP client computers.
- The **Group Policy Object Editor** and the **Software Restriction Policies** extension of **Group Policy Object Editor**.

For Windows XP computers that are not participating in a domain, you can use the **Local Security Settings** snap-in to access **Software Restriction Policies**.

**Software Restriction Policies Architecture**

Software restriction policies provide a mechanism for the operating system and applications compliant with software restriction policies to restrict the runtime execution of software programs.

At a high level, software restriction policies consist of the following components:

- Software restriction policies (sometimes referred to as Safer) Application Programming Interfaces (APIs), which are used to create and configure the rules that constitute the software restriction policy. There also are software restriction policies APIs for querying, processing, and enforcing software restriction policies.
- A software restriction policies management tool. This consists of the **Software Restriction Policies** extension of the **Group Policy Object Editor** snap-in, which administrators use to create and edit the software restriction policies.
- A set of operating system APIs and applications that call the software restriction policies APIs to provide enforcement of the software restriction policies at runtime.
- Active Directory and Group Policy. Software restriction policies depend on the Group Policy infrastructure to propagate the software restriction policies from the Active Directory to the appropriate clients, and for scoping and filtering the application of these policies to the appropriate target computers.
- Authenticode and WinVerify Trust APIs which are used to process signed executable files.
- Event Viewer. Software restriction policies functions log events to the Event Viewer logs.
- Resultant Set of Policies (RSoP), which can aid in the diagnosing of the effective policy that will be applied to a client.

The following figure shows an example of the Group Policy Object Editor targeted on one of three different Group Policy objects (GPOs) that apply to a workstation. The Group Policy Object Editor is viewing a GPO containing software restriction policies, A1GPO, which is associated with the University domain. When the computer starts or the user logs on, the policies from A1GPO, A2GPO, and A3GPO are downloaded and applied.

**Software Restriction Policies, Group Policy and Active Directory**




After Group Policy downloads the software restriction policy, the policy is stored in the registry. This policy is processed during the calls to the software restriction policies (Safer) enforcement APIs. For example, when users double-click an executable file (.exe) in Windows Explorer, the enforcement API, **SaferIdentifyLevel**, is called to determine the rule details that are associated with that executable file. In this example, if the software restriction policy does not indicate that the program is fully trusted, the executable file does not run. **ShellExecute** calls **SaferIdentifyLevel** when a non-natively executable file is double-clicked (this is covered only by **Designated File Types**). For natively-executable files (a portable executable file or PE), **CreateProcess** calls **SaferIdentifyLevel**. See "Default Designated File Types," later in this section for more information.

Software restriction policies consist of the following components.

**Software Restriction Policies Components**

| Component | Description |
|---|---|
| Software Restriction Policies (Safer) API | This component is an API for creating and manipulating software restriction policies. The functions in this API are used to get and set the rules that constitute a software restriction policy. Safer APIs provide applications that launch programs from external sources the ability to query security policy for approval before an executable file is launched. |
| Management tools | These components include the **Group Policy Object Editor** snap-in, the **Software Restriction Policies** extension, and command line tools. The management tools use the Software Restriction Policies (Safer) APIs to allow the computer administrator to author and edit software restriction policies. |

| | |
|---|---|
| Safer Enforcement APIs | In Windows Server 2003, these components include functions (**SaferIdentifyLevel** and **SaferComputeTokenFromLevel**) that the operating system and applications compliant with software restriction policies use to enforce the software restriction policy. These functions can process the set of software restriction policies rules contained in the policy and return the most appropriate rule for a given program. |
| Operating system APIs and applications | A set of operating system APIs and applications, compliant with software restriction policies, which call into the Safer enforcement functions to provide enforcement of the policy at runtime. These include Microsoft Office Macros, a .DLL loader, Component Object Model (COM), and these APIs:<br>1. **ShellExecute.** Performs an operation on a specified file.<br>2. **CreateProcess.** Creates a new process and its primary thread. The new process runs the specified executable file in the security context of the calling process. |
| Microsoft Windows Script Host | Windows Script Host is not a component of Software Restriction Policies; however, it interacts with Software Restriction Policies. Windows Script Host calls into Software Restriction Policies to apply software restriction policies to scripts that are being run. Note that other script engines do not do this. If you want to perform script checking, you can only run script hosts that call software restriction policies.<br>Windows Script Host is a language-independent scripting host for Windows Script compatible scripting engines. It brings simple, powerful, and flexible scripting to the Windows 32-bit platform, enabling you to run scripts from both the Windows desktop and the command prompt. |
| Windows Installer | Windows Installer (Msiexec) is not a component of Software Restriction Policies, but it also interacts with Software Restriction Policies. Windows Installer calls into Software Restriction Policies to apply software restriction policies to scripts that are being installed. Windows Installer implements the Software Restriction Policies API. |
| Group Policy | Group Policy is a feature included in Windows XP, Windows 2000, and Windows Server 2003. Software restriction policies use the Group Policy's *Registry.pol* files.<br>The Administrative Templates extension of Group Policy saves information in Registry.pol files. These files contain the customized registry settings that you specify (by using Group Policy) to be applied to the computer or user portion of the registry. One Registry.pol file contains registry settings that are specific to the **HKEY_LOCAL_MACHINE** key; it is stored in the Group Policy Template folder, in the \Machine subfolder. The other Registry.pol file contains registry settings that are specific to the **HKEY_CURRENT_USER** key; it is stored in the Group Policy Template folder, in the \User subfolder.<br>Software restriction policies do not use a client side extension. Software restriction policies depend on Group Policy to propagate the policy from the Active Directory to the appropriate clients and use the appropriate filtering and scoping. |
| Event Viewer | Software restriction policies functions log events in the Event log. The **SaferRecordEventLogEntry** function saves messages to an event log. |
| Authenticode and WinVerifyTrust | These components are used to process signed executable files. Microsoft Authenticode, which is based on industry standards, enables developers to include information about themselves and their code with their programs through the use of digital signatures. The **WinVerifyTrust** function performs two actions: signature checking on a specified object and trust verification action. |
| Windows Management Instrumentation (WMI) | The Software Restriction Policies extension queries the registry policy by using WMI. WMI is a component of the Microsoft Windows operating system and is the Microsoft implementation of Web-Based Enterprise Management (WBEM), which is an industry initiative to develop a standard technology for accessing management information in an enterprise environment. WMI uses the Common Information Model (CIM) industry standard to represent systems, applications, networks, devices, and other managed components. You can use WMI to automate administrative tasks in an enterprise environment. |
| Resultant Set of Policy (RSoP) | RSoP is a Group Policy feature that aids in diagnosing effective policy settings that will be applied to a client. Software restriction policies provide support for RSoP. |

| | |
|---|---|
| Registry | Software restriction policies are stored in the registry. In **HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\Safer** or in **HKEY_CURRENT_USER\SOFTWARE\Policies\Microsoft\Windows**. |

For information about the Software Restriction Policies (Safer) functions, see the topic "Safer" in the Security section of the Microsoft Platform SDK on MSDN.

**Software Restriction Policies Overview**

A software restriction policy consists of a default rule that defines which programs can run, and a set of exceptions to the default rule.

You can set the default rule, which specifies whether software can run, to either **Unrestricted** (program can run) or **Disallowed** (program cannot run). If you set the default rule to **Unrestricted**, you can also define exceptions to specify a set of programs that are not allowed to run. The **Unrestricted** default setting is appropriate in an environment with lightly managed clients. For example, you can prevent users from installing a program that conflicts with existing programs by creating a rule to block that program.

For environments that require a more secure method of software control, you should set the default rule to **Disallowed** and then allow only a specific set of programs to run. If you use the **Disallowed** default setting, you have to identify all the software that users are permitted to run and create appropriate rules for them. The **Disallowed** default setting is the preferred default for securing Windows XP clients.

**Rules in Software Restriction Policies**

Software restriction policies use rules to identify one or more software applications, and to specify whether the software is permitted to run. By creating rules, administrators can identify software that is an exception to the default rule. Each rule can include descriptive text to indicate why the rule was created.

To identify software, software restriction policies support the following rules:

- **Hash.** A cryptographic fingerprint of the file
- **Certificate.** A software publisher certificate used to digitally sign a file
- **Path.** The local or universal naming convention (UNC) path of where the file is stored
- **Zone.** Internet Zone

**Hash Rule**

A hash is a cryptographic fingerprint that uniquely identifies a file regardless of its name or where it is accessed. With a hash rule, software can be renamed or moved into another location on a disk, but it will still match the hash rule because the rule is based on a cryptographic calculation involving file contents. Hash rules work effectively in a static environment. If the software in your clients is upgraded frequently, hash rules can become difficult to manage. Whenever a program executable is updated, the hash rule needs to be updated to support the new executable version.

For secure environments using a default level of **Disallowed**, hash rules are a very secure way to identify applications that you want to allow to run as **Unrestricted**. For environments using a default level of **Unrestricted**, hash rules allow you to specifically identify one particular version of an application to disallow.

A hash rule consists of three pieces of data, separated by colons:

- MD5 or SHA-1 hash value
- File length
- Hash algorithm ID

A hash rule uses the following format:

```
[MD5 or SHA1 hash value]:[file length]:[hash algorithm id]
```

Files that are digitally signed use the hash value contained in the signature, either SHA-1 or MD5. Files that are not digitally

signed use an MD5 hash.

For example, the following hash rule matches a file with a length of 126 bytes and with contents that match the MD5 (denoted by the hash algorithm identifier of 32771) hash of 7bc04acc0d6480af862d22d724c3b049:

```
7bc04acc0d6480af862d22d724c3b049:126:32771
```

**Certificate Rule**
A certificate rule specifies a code-signed software publisher's certificate. For example, a company can require that all scripts and ActiveX controls be signed with a particular set of publisher certificates. Certificates used in a certificate rule can be issued from a commercial certificate authority (CA) such as VeriSign, a Windows 2000 or Windows Server 2003 Public Key Infrastructure (PKI), or a self-signed certificate.

A certificate rule uses signed hashes contained in the signature of the signed file to match files regardless of name or location. If you want to make exceptions to a certificate rule, you can use a hash rule to identify the exceptions. Certificate rules provide more flexibility than hash rules. By using a certificate rule, you can allow (or deny) multiple applications at once. Certificate rules also allow low maintenance of the software restriction policies rule-set, but they can provide the security of hash rules.

**Path Rule**
A path rule can specify a folder or fully qualified path to a program. When a path rule specifies a folder, it matches any program contained in that folder and any programs contained in subfolders. Software restriction policies support local and Uniform Naming Convention (UNC) paths.

The administrator must define all directories for launching a specific application in the path rule. For example, if the administrator creates a shortcut on the desktop to launch an application, then in the path rule, the administrator must also grant the user Read access rights to both the executable file and the shortcut paths to run the application. If all the path information necessary for launching the application in the path rule is not defined, it can trigger the **Software Restricted** warning when the user attempts to run the application.

**Environment variables in path rules**
You can use environment variables in a path rule. Since path rules are evaluated in the client environment, the ability to use environment variables (for example, %Windir%) allows a rule to adapt to a particular user's environment.

### Note

- Environment variables are *not* protected by Access Control Lists (ACLs). There are two types of environment variables: **Users** and **System**. Users that are able to start a command prompt can redefine the **Users** environment variable to a different path. The reason for this is that if an administrator uses %Windir% in a path rule, then a user can run Cmd.exe, change the value of %Windir% to a directory where the user's .exe file is, run his or her .exe file, and then exit that shell or change %Windir% back. It is also possible for a user to make changes programmatically or by using the **System Properties** page in Control Panel. Only users in the Administrators group have permissions to change the **System** environment variable.

- If you need to provide enhanced security in your environment, use the registry-based path rules, which is the preferred method. See "Registry Path Rules," later in this section.

The following examples show instances of applying environment variables to a path rule:

- "%UserProfile%" matches C:\Documents and Settings\User and all subfolders under this directory.

- "%ProgramFiles%\Application" matches C:\Program Files\Application and all subfolders under this directory.

**Wildcards in path rules**
A path rule can incorporate the '?' and '*' wildcards, allowing rules such as "*.vbs" to match all Visual Basic Script files. The following examples illustrate the use of wildcards:

- "\\DC-??\login$" matches \\DC-01\login$, \\DC-02\login$

- "*\Windows" matches C:\Windows, D:\Windows, E:\Windows

- "c:\win*" matches c:\winnt, c:\windows, c:\windir

**Registry path rules**

Many applications store paths to their installation folders or application directories in the Windows registry. You can create a path rule to look up these registry keys. For example, some applications can be installed anywhere on the file system. These locations might not be easily identifiable by using specific folder paths, such as C:\Program Files\Microsoft Platform SDK, or environment variables, such as %ProgramFiles%\Microsoft Platform SDK. If the program stores its application directories in the registry, you can create a path rule that uses the value stored in the registry, such as **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\PlatformSDK\Directories\Install Dir%**. This type of path rule is called a registry path rule. Using registry path rules is a recommended practice because users cannot bypass the path rule by undermining the environment variable expansion.

The registry path format is:

**%[Registry Hive]\[Registry Key Name]\[Value Name]%**

**Note**

- A registry path rule suffix must not contain a backslash (\) character immediately after the last percent sign (%) in the rule. The registry hive name must be written out; abbreviations will not work.

The following conditions apply to registry paths rules:

- The registry path must be enclosed in percent signs (%).
- The registry value must be a **REG_SZ** or **REG_EXPAND_SZ**. You cannot use HKLM as an abbreviation for **HKEY_LOCAL_MACHINE**, or HKCU as an abbreviation for **HKEY_CURRENT_USER**.
- If the registry value contains environment variables, these will be expanded when the policy is evaluated.
- A registry path rule can also contain a suffix path such as **%HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\Shell Folders\Cache%OLK*.** This registry path rule identifies the folder that Microsoft Outlook XP uses to store attachments before launching them. The attachment folder always starts with the letters "OLK" so the rule uses wildcard matching. For example, this rule matches the following path: C:\Documents and Settings\username\Local Settings\Temporary Internet Files\OLK4.

**Note**

- When you set a path rule, you should check the access control list (ACL) entries on the path. If users have write access to a path, they can modify its contents. For example, if you allow C:\Program Files, any power user on the computer can copy software into the Program Files folder.

When the default rule is set to **Disallowed**, there are four registry paths that are setup so the operating system has access to system files for normal operation. These registry path rules are created as a safeguard against locking yourself and all other users out of the system. These registry rules are set to **Unrestricted**. Only advanced users should consider modifying or deleting these rules. The registry path rule settings are:

- **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SystemRoot%**
- **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SystemRoot%*.exe**
- **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\ CurrentVersion\SystemRoot%\System32*.exe**
- **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\ProgramFilesDir%**

**Path rule precedence**

When there are multiple matching path rules, the most specific matching rule takes precedence.

The following is a set of paths, from highest precedence (more specific match) to lowest precedence (more general match):

- Drive:\Folder1\Folder2\FileName.Extension
- Drive:\Folder1\Folder2*.Extension
- *.Extension
- Drive:\Folder1\Folder2\
- Drive:\Folder1\

**Zone Rule**
A rule can identify software from the Internet Explorer zone from which it is downloaded. These zones are:

- Internet
- Intranet
- Restricted Sites
- Trusted Sites
- My Computer

The current version of the Internet zone rule applies to only Windows Installer (*.msi) packages. This rule does not apply to software downloaded by using Internet Explorer. All other file types affected by zone rules are listed in the table on designated file types that appears under the "Default Designated File Types" heading, later in this section. There is one list of designated file types that is shared by all zone rules.

### Note

- Each rule is associated with a globally unique identifier (GUID). An example GUID is {f8c2c158-e1af-4695-bc93-07cbefbdc594}. Two identical rules, for example two **Disallowed**hash rules on the same program, have a different GUID associated with each. This GUID is stored in the registry along with the rule data.

**Rules Precedence**
Rules for software restriction policies are evaluated according to a specific order. The rules that more specifically match a program take precedence over rules that more generally match the same program. If two identical rules with differing security levels are established for the same software, the rule with the highest security level takes precedence. For example, if two hash rules, one with the security level **Disallowed** and one with the security level **Unrestricted**, are applied to the same software program, the rule with the security level **Disallowed** takes precedence, and the program will not run. The following ordered list defines the precedence order for the rules from the most specific (hash) to the least specific rule (default):

1. Hash rule
2. Certificate rule
3. Path rule
4. Zone rule
5. Default rule

**Policy Enforcement Options**
There are two policy enforcement options that influence the behavior of a software restriction policy. The following options define how software restriction policies are applied for all files, including signed files:

- **Apply software restriction policies to the following: All files**. This option includes dynamic-link libraries (DLLs).

- **Apply software restriction policies to the following users**: **All users except local administrators.** This option prevents the software restriction policies from applying to local administrators.

#### Including DLLs

Most programs consist of an executable file and many supporting DLLs. By default, software restriction policy rules are not enforced on DLLs. Not checking DLLs is the recommended option for most customers for the following reasons:

- Including DLLs degrades system performance because it has to check all libraries linked to the application. For example, if a user runs 10 programs during a logon session, the software restriction policy evaluates each program. When the **Apply software restriction policies to the following: All files** is selected, the software restriction policy evaluates each DLL load within each program. If each program uses 20 DLLs, this results in 10 executable program checks plus 200 DLL checks, so the software restriction policy must perform 210 evaluations. A program such as Internet Explorer consists of an executable file, iexplore.exe, and many supporting DLLs.

- Setting the default security level to **Disallowed** forces the system to identify both the main executable file andits supporting DLLs before the .exe file can run, which places added burden on the system.

### Note

- To provide the highest assurance possible for the programs running in environments that want the highest assurance possible when running programs, DLL Checking is the recommended option. DLL checking is recommended because while viruses primarily target executable files, some specifically target DLLs. To ensure that a program does not contain a virus, you can use a set of hash rules that identify the executable file and its constituent DLLs.

#### Excluding Administrators

An administrator might want to disallow programs from running for most users, but allow administrators to run all of them. For example, an administrator might have a shared computer that multiple user's connect to using terminal server. The administrator might want users to run only specific applications on the computer, but allow members in the local Administrators group to run any program. You can use the **Apply software restriction policies to the following users: All users except local administrators** enforcement option to do this.

If the software restriction policy is created in a GPO linked to an object in Active Directory, instead of using the **Skip Administrators** option, consider denying the **Apply Group Policy** permission on the GPO to the Administrators group. This approach consumes less network traffic because GPO settings that do not apply to administrators are not downloaded.

### Note

- Software restriction policy defined in local security policy objects cannot filter user groups. In this case, use the Skip Administrators option.

#### Default Designated File Types

Software restriction policy rules apply automatically to some types of executable files, regardless of their extension type. This includes the following types of files:

- Windows operating system executable files (.exe, .com, .dll)

- Windows scripting files (when processed by Windows Script Host, such as .vbs, and .wsh files)

- Command batch files (when processed by Windows CMD, such as .bat, .cmd)

- Installation packages when processed by Windows Installer (.msi)

You can apply software restriction policies rules to additional file types, even when they are handled by non-Microsoft programs. However, if the other program does not enforce software restriction policies itself, the software restriction policies restrictions will only apply when either Windows Explorer or Internet Explorer is used to open the file. Opening the other program and then opening the executable script from within that program circumvents software restriction policies. To resolve this issue, contact the software provider and request that they enforce software restriction policies when opening executable

scripts, or disallow execution of the other scripting host.

The file types are listed in the **Designated File Types Properties** dialog box. This dialog is accessible in the Group Policy Object Editor, in **Computer Configuration\Windows Settings\Security Settings\Software Restriction Policies**, under **Designated File Types** in the details pane. Additional file types to which you want to apply rules can be added to this list. For example, for Perl scripting files you might choose to add .pl and other file types associated with the Perl engine to the **Designated file types** list under the **General** tab of the **Designated File Types Properties** dialog box. Note that you can check only the initial script this way; if a Perl script calls another Perl script, there will be no software restriction policies check.

The following table lists what is considered to be executable code. This is in addition to standard program files such as .exe, .dll, and .vbs.

**Default Designated File Types**

| File Extension | File Description |
|---|---|
| .ade | Microsoft Access Project Extension |
| .adp | Microsoft Access Project |
| .bas | Visual Basic Class Module |
| .bat | Windows Batch File |
| .chm | Compiled HTML Help File |
| .cmd | Windows NT Command Script |
| .com | Application |
| .cpl | Control Panel extension |
| .crt | Security Certificate |
| .exe | Application |
| .hlp | Windows Help File |
| .hta | HTML Applications |
| .inf | Setup Information File |
| .ins | Internet Communication Settings |
| .isp | Internet Communication Settings |
| .lnk1 | Shortcut |
| .mdb | MDB File |
| .mde | Microsoft Access MDE Database |
| .msc | Microsoft Common Console Document |
| .msi | Windows Installer Package |
| .msp | Windows Installer Patch |
| .ocx | ActiveX Controls |
| .pcd | Photo CD Image |
| .pif | Shortcut to Program |
| .reg | Registration Entries |

| | |
|---|---|
| .scr | Screen Saver |
| .shs | Shell Scrap Object |
| .url | Internet Shortcut (Uniform Resource Locator) |
| .vb | VB File |
| .wsc | Windows Script Component |

[1] If you click a .lnk file (shortcut) in the UI, ShellExecute applies software restriction policies to it if the .lnk extension is included in Designated File Types. Otherwise, ShellExecute runs the target, where software restriction policies get applied in CreateProcess.

If the file type of the shortcut's (.lnk) target is not an executable file directly, such as a .doc file, and if .doc is included in **Designated File Types,** software restriction policies rules are applied to the .doc file. If .doc is not included in **Designated File Types,ShellExecute** runs the handler with specified options (such as running "WinWord filename.doc"). At this point, **CreateProcess** applies software restriction policies to the application, WinWord, and not to the filename.doc, which was the actual target of the .lnk.

#### Trusted Publishers

You can use the **Trusted Publishers Properties** dialog box to configure which users can select trusted publishers. You can also determine which, if any, certificate revocation checks are performed before trusting a publisher. With certificate rules enabled, software restriction policy will check a certificate revocation list (CRL) to ensure the software's certificate and signature are valid. This can decrease performance when starting signed programs. The options under the **General** tab of the **Trusted Publishers Properties** dialog box. You access this dialog in the **Group Policy Object Editor** snap-in, by navigating to **Computer Configuration\Windows Settings\Security Settings\Software Restriction Policies**, anddouble-clicking **Trusted Publishers** in the details pane.

### Note

- If you enable the Trusted Publishers option, you must also explicitly add the certificates of allowed ActiveX publishers to the Trusted Publisher certificate store. Adding the certificates is required to ensure that installation and other operations work properly and to maintain the appropriate system functions. If you don't add the certificates, ActiveX controls will not install, and some applications which perform signature validation might exhibit unexpected behavior. All trust validations using the default Authenticode policy fail unless the signing certificate is in the Trusted Publishers store. For example, Windows Update automatically validates its own signature and the signatures of downloaded files. The Trusted Publishers restrictions are only recommended for the most restrictive environments.

#### Default Settings for a Software Restriction Policy

The default settings for a software restriction policy include the following:

- Default **Security Level**: **Unrestricted**
- **Enforcement** options:
  - Apply to Files: All software files except libraries (such as DLLs)
  - Apply to Users: All users
- **Additional Rules**: none
- **Designated File Types**: These are defined in the "Default Designated File Types" section, earlier in this document.
- **Trusted Publishers**:
  - **Select Trusted Publishers**: End users
  - **Publisher** certificate revocation checking: Not selected

- **Timestamp** certificate revocation checking: Not selected

**Example: Applying Software Restriction Policies in a Highly Restricted Environment**
In some cases, you might want to manage all of the software that runs on a computer. For example, even when users have insufficient rights to replace system files or files in shared folders such as Program Files, if the user's have a place on the file system to which they can write, they could copy a program there and start it from that location.

Administrators can create a highly-restricted (locked down) configuration to prevent users from modifying the contents of the Program Files or Windows folders, and allow them to run only software that is installed by the administrator. In this case, you would also need to ensure that users are not administrators on their local computers.

The following policy can help prevent users from running malicious code:

- Default Security Level: Disallowed
- Apply software restriction policies to the following users: All users except administrators
- Path Rules:
    - %WINDIR% Unrestricted
    - %PROGRAMFILES% Unrestricted

This policy disallows all software on the users computer, except software that is installed in the Windows directory, Program Files directory, or their respective subfolders. The policy does not apply to administrators.

If all the programs a user needs are not installed in %WINDIR% or in %PROGRAMFILES%, or there are programs in those folders that the administrator does not want the user to run, the administrator can make additional exceptions as follows:

- Path Rules:
    - %WINDIR%\regedit.exe: Disallowed
    - %WINDIR%\system32\cmd.exe: Disallowed
    - \\CORP_DC_??\scripts: Unrestricted
    - %HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates \InoculateIT\6.0\Path\HOME%: Unrestricted

As a result of these exceptions, the following occurs:

- Both the command prompt (cmd.exe) and the registry editor (regedit.exe) are disallowed.
- An exception is created to allow login scripts to run on the users computer.
- The use of the question mark "?" wildcard allows the rule to match \\CORP_DC_01, \\CORP_DC_02, and others.
- A registry path rule is added that allows the anti-virus software on the computer to run.

**Software Restriction Policies Processes and Interactions**
This section discusses how software restriction policies are processed and stored, and also explains how software restriction policies are merged.

**Software Restriction Policies Processing**
Software restriction policy is distributed by using the normal Group Policy process. That is, software restriction policies are saved in Group Policy objects (GPOs), and the management tools use the Group Policy Interfaces (**IGroupPolicyObject**) to save and load the software restriction policies. GPOs are linked to sites, domains, or organizational units (OUs) in Active Directory. An administrator can use the standard Group Policy mechanisms to apply additional filtering options to a GPO to

have it apply to a select set of users or computers within an OU.

When a computer starts up or when a user logs on, the Group Policy engine queries Active Directory for any GPOs that apply to it. It downloads and applies these GPOs in a specific order (local, site, domain, and OU), where the closest OU to the user or computer object is the last GPO applied. If Group Policy fails to download any of the GPOs, the entire process is aborted and no GPOs are applied.

In regard to software restriction policies (also referred to as Safer), Group Policy processing essentially involves downloading a number of Registry.pol files from Active Directory. The Registry.pol files contain the registry keys and values set through the Safer policy APIs. The keys and values of each policy are then copied into the local user's registry in the **HKEY_LOCAL_MACHINE** or **HKEY_CURRENT_USER** registry hive as appropriate. This process is repeated for each GPO that contains a software restriction policy. Group Policy is periodically refreshed: on clients, security policy reapplies every 16 hours regardless of whether policy changes have occurred. If there is a policy change, it applies within 90 minutes of the change. All the policy distribution is handled by the Group Policy infrastructure and not by the software restriction policies extension directly. The Safer Enforcement APIs base their decisions on the effective policy, which is the result of combining the software restriction policies in the **HKEY_LOCAL_MACHINE** and **HKEY_CURRENT_USER** registry hives.

Software restriction policies do not implement a client side extension to Group Policy.

**Group Policy Processing Order**

By default, Group Policy is inherited and cumulative, and it affects all computers and users in an Active Directory container. Group Policy objects are processed according to the following order:

- **Local Group Policy object.** Each computer has one Group Policy object that is stored locally. This processes for both computer and user Group Policy processing. The local GPO can't be blocked by domain-based GPOs. However, settings in domain GPOs always take precedence because they are processed after the local GPO.

- **Site.** Any GPOs that have been linked to the site that the computer belongs to are processed next. They are processed in the order specified by the administrator, on the **Linked Group Policy Objects** tab for the site in the **Group Policy Management Console** (GPMC). The GPO with the lowest link order is processed last, and therefore has the highest precedence.

- **Domain.** Multiple domain-linked GPOs are processed in the order specified by the administrator, on the **Linked Group Policy Objects** tab for the domain in GPMC. The GPO with the lowest link order is processed last, and therefore has the highest precedence.

- **Organizational units.** GPOs that are linked to the OU that is highest in the Active Directory hierarchy are processed first. Then GPOs that are linked to its child OU are processed, and then any lower child OUs. Finally, the GPOs that are linked to the OU that contains the user or computer are processed.

This processing is subject to the following conditions:

- WMI or security filtering that has been applied to GPOs.

- A domain-based GPO can be enforced by using the **Enforce** option to ensure that it's policies cannot be overwritten. Because an Enforced GPO is processed last, no other settings can overwrite the settings in that GPO. If more than one Enforced GPO exist, its possible to set the same setting in each GPO to a different value, in which case, the link order of the GPOs determines which one contains the final settings.

- At any domain or OU, Group Policy inheritance can be selectively designated as **Block Inheritance.** However, because enforced GPOs are always applied and cannot be blocked, blocking inheritance does not prevent policy from **Enforced** GPOs from being applied.

**Group Policy Refresh Period**

By default, Group Policy is processed every 90 minutes plus a randomized offset of up to 30 additional minutes for a total maximum refresh interval of up to 120 minutes. For domain controllers, the default period is every 5 minutes. You can change the default values for computers by using the **Group Policy Refresh Interval for Computers** Group Policy setting under **Administrative Templates** in the **Default Domain Controllers** GPO. For domain controllers, the **Group Policy Refresh Interval for Domain Controllers** setting controls this default; it is available under **Computer Configuration\Administrative Templates\System\Group Policy**.

For computers, the **Gpupdate** command can also be used to trigger a background refresh of Group Policy on demand from the client.

**Merging for Multiple Software Restriction Policies**
Whenever two or more Group Policy objects apply to a user or computer, the policies are merged. When two or more software restriction policies are merged, the following occurs:

- The GPO with the highest precedence sets the following values:
    - Default Security Level
    - Designated File Types
    - Skip Administrators
    - DLL Checking
- The rules from multiple GPOs are always merged. Thus, all additional rules from all GPOs are preserved.

A software restriction policy can be set for user scope and computer scope. The following semantics are observed when merging user and computer scope:

- The more restrictive default security level is chosen.
- The list of designated file types in the computer policy, if present, is used. If not present, the list of designated file types in the user policy is used.
- The Skip Administrators value is always chosen from the computer policy.
- If DLL checking is enabled in either policy, then it is enabled.
- All the rules between user and computer policies are merged.

**Storage of Software Restriction Policies Settings**
After a software restriction policy is applied, the software restriction policy configuration is stored in the system registry. The security access control list (ACL) protecting these registry keys allows only members of the Administrators group and the SYSTEM account to change them.

Software restriction policies use these registry keys:
**HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\Safer** and
**HKEY_CURRENT_USER\Software\Policies\Microsoft\Windows.**

- **CodeIdentifiers**.
- **Default Level** (DWORD). 40000 for **Unrestricted**, 0 for **Disallowed**.
- **ExecutableTypes** (REG_MULTI_SZ). List of extensions for designated file types.
- **TransparentEnabled** (DWORD). Defines which files to include during rule evaluation. 0 means no enforcement, 1 indicates to exclude DLLs in evaluation, and 2 indicates to include all files in evaluation.
- **PolicyScope** (DWORD). Defines the scope of users to which this policy applies. 0 applies to all users, and 1 applies to all users except administrators. Note that this value is only valid for HKEY_LOCAL_MACHINE.
- **AuthenticodeEnabled** (DWORD). Defines if certificate rules should be applied, 0 means disabled, 1 indicates that certificate rules will be applied. Note that this value is only valid for HKEY_LOCAL_MACHINE.
- **LogFileName** (REG_SZ). An optional registry value that needs to be set manually. This value defines a path to the log file

that captures advanced logging. If this value is set, advanced logging is enabled and written to this log file. Note that this value is only valid for HKEY_LOCAL_MACHINE.

- The zero (**0**) key**.** Indicates a **SAFER_LEVELID_DISALLOWED** security level. Entries under this key are **Disallowed** rules.

    - **Hashes.** All hash rules are defined under this key.

      **{GUID}**. This is uniquely generated for each Rule defined.

      **Description** (REG_SZ). This is the administrator-defined description of the rule.

      **FriendlyName** (REG_SZ). This is the name of the hash.

      **HashAlg** (REG_DWORD). This is the type of hash that is stored in **ItemData**. For files that are not digitally signed or are signed with MD5, this value is defined as 0x8003 or CALG_MD5 or (ALG_CLASS_HASH | ALG_TYPE_ANY | ALG_SID_MD5).

      **ItemData** (REG_BINARY). The actual hash to the file. This value should always be 16 bytes and is generated with a call to **CodeAuthzpComputeImageHash()**.

      **ItemSize** (REG_QWORD). This is the actual file size of the target file. To evaluate this value, reverse the bytes. For example, if the value data was: 00 0A 01 00 00 00 00 00, the file size would actually be 010a00 or 68,096 bytes.

      **LastModified** (REG_QWORD). This is the date and time down to seconds of when this entry was last updated. Several utilities exits to extract this name into a readable format however just as the ItemSize data was reversed so should this value.

      **SaferFlags** (REG_DWORD). Not used and will always be set to zero.

- **Path.** All Path rules are defined under this key.

    - **{GUID}.** This is uniquely generated for each Rule defined.

    - **Description** (REG_SZ). This is an administrator defined description of the rule.

    - **ItemData** (REG_EXPAND_SZ Or REG_SZ). This is the path to be used for this rule. Note that if the data provided for this value includes any environment variables that must be expanded the value type will be REG_EXPAND_SZ otherwise it will be REG_SZ.

    - **LastModified** (REG_QWORD). This is the date and time down to seconds of when this entry was last updated. Several utilities exits to extract this name into a readable format however just as the ItemSize data was reversed so should this value.

    - **SaferFlags** (REG_DWORD). This is not used and will always be set to zero.

- **UrlZones**. All zone rules are defined under this key.

    - **{GUID}.** This is uniquely generated for each Rule defined.

    - **ItemData** (REG_BINARY). Defines the URLZone to which this rule is defined. URL Zones are defined in urlmon.h and are of type enum tagURLZONE. For example if a value of 0x3 appears in this value that evaluates to URLZONE_INTERNET.

    - **LastModified** (REG_QWORD). This is the date and time down to seconds of when this entry was last updated. Several utilities exits to extract this name into a readable format however just as the ItemSize data was reversed so should this value.

- **SaferFlags** (REG_DWORD). This is not used and will always be set to zero.

- **262144**. Indicates a **SAFER_LEVELID_FULLYTRUSTED** security level. Entries under this key are **Unrestricted** rules.

**Certificate Rules**
Certificate rules are stored in a separate key in the registry.

Certificate rules for user software restriction policies are stored in this registry key:

- **HKEY_CURRENT_USER\SOFTWARE\Policies\Microsoft\SystemCertificates**

Certificate rules for computer software restriction policies are stored in this registry key:

- **HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\SystemCertificates**

**Resultant Set of Policy**
To obtain information about policy stored in the registry, the Software Restriction Policies extension queries the registry policy by using WMI.

The **RSOP_RegistryPolicySetting** WMI namespace represents policy data for registry or administrative template extension. Policy settings in an administrative template file (.adm) involve setting values on registry keys that are abstracted by this namespace. Software Restriction Policies also store information in the registry.

To retrieve this registry data, the **IwbemServices** WMI interface's **ExecQuery** method is used to perform a query to retrieve objects from the **RSOP_RegistryPolicySetting** namespace.

**Related Information**
The following resources contain additional information that is relevant to this section.

- "How Core Group Policy Works" in the Group Policy Collection.

- "How Group Policy Software Installation Extension Works" in the Group Policy Collection.

# Software Restriction Policies Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Software Restriction Policies Tools and Settings**

**In this section**


- Software Restriction Policies Tools
- Software Restriction Policies Group Policy Settings
- Related Information


Administrators can use command line tools to refresh Group Policy settings, including software restriction policies, and to display the resulting set of policies that were enforced on the computer for a specified user at logon. To assess the policy settings that are in effect for a computer or user, administrators use the Resultant Set of Policy (RSoP) snap-in.

**Software Restriction Policies Tools**

The following tools are associated with Software Restriction Policies:

**Gpupdate.exe**

This tool is used for refreshing local and Active Directory policy settings on the computer from which you run the **gpupdate** command. For more information about this tool, see "Core Group Policy Tools and Settings" in this collection.

**Gpresult.exe**

This tool enables you to examine the Group Policy settings applied during a policy refresh. For more information about this tool, see "Core Group Policy Tools and Settings" in this collection.

**Resultant Set of Policy (RSoP)**

This tool polls for existing Group Policy settings and planned policy settings, and then reports the results of those queries. For more information about RSoP, see "What Is Resultant Set of Policy?" in this collection.

**Software Restriction Policies Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with software restriction policies.

**Software Restriction Policies Security Levels and Additional Rules**

| Security Levels | Description |
|---|---|
| Disallowed | Does not allow the specified software to run. |
| Unrestricted | Allows the specified software to run on the computer with the full rights of the currently logged on user. |

**Additional Rules**

| Security Levels | Description |
|---|---|
| Hash Rule | A series of bytes with a fixed length that uniquely identifies a software program or file. A hash (also called a message digest) is obtained by applying a one-way mathematical function (sometimes called a hash algorithm) to an arbitrary amount of data. If the input data changes, the hash changes. The hash can be used in many operations, including authentication and digital signing.<br><br>If you create a hash rule for a software program, software restriction policies calculate a hash of the program. When a user tries to open a software program, a hash of the program is compared to existing hash rules for software restriction policies. The hash of a software program is always the same, regardless of where the program is located on the computer. However, if any changes are made to the software program, its hash also changes, and it no longer matches the hash in the hash rule for software restriction policies. |

| | |
|---|---|
| Path Rule | Identifies a program according to a folder or its fully qualified path. Both URL and UNC paths are permitted. You can use the following in path rules: environment variables, wildcards (question mark "?" and asterisk "*"), and registry path rules. |
| Certificate Rule | Identifies software based on a signed certificate. You create a certificate rule that identifies software and then specify a security level to either allow or not allow the software to run. |
| Internet Zone Rule | Identifies software from a zone that is specified through Internet Explorer. The zones are Internet, Intranet, Restricted sites, Trusted sites, and My Computer. These rules apply only to Windows Installer packages (.msi files). |

The following table lists the software restriction policy options.

**Software Restriction Policy Options**

| Options | Description |
|---|---|
| Enforcement | Enforcement enables you to specify whether to turn on dynamic-link library (DLL) checking and skip administrators to prevent the software restriction policies from applying to local administrators.<br><br>Use the Apply software restriction policies to the following option to select one of the following:<br><br>• **All software files except libraries (such as DLLs).** This option specifies that the rules will not affect DLLs.<br><br>• **All software files.** This option applies software restriction policies to all files, including DLLs. This option turns on DLL checking.<br><br>Use the Apply software restriction policies to the following users option to select one of the following:<br><br>• **All users.** This option specifies that all the rules you define apply to all users.<br><br>**All users except local administrators.** This option prevents software restriction policies from applying to local administrators. This is used when administrators want to prevent most users from running certain programs, but allow local administrators to run any program. |
| Designated File Types | The Designated File Types dialog box lists the file types to which the software restriction policy applies. The list represents the types of files that are considered executable files. The rules you specify in a software restriction policy apply only to the file types listed in the Designated File Types list. If you want to be able to set rules on additional file types, add them to the Designated File Types list. |
| Trusted Publishers | The Trusted Publishers options enable you to configure settings related to ActiveX controls and other signed content. Trusted Publishers includes the following options:<br><br>• **Enterprise Administrators:** Use this option to allow only domain administrators to make decisions regarding signed active content.<br><br>• **Local computer Administrators:** Use this option to allow local computer administrators to make all decisions regarding signed active content.<br><br>• **End Users:** Use this option to allow any user to make decisions regarding signed active content.<br><br>• **Publisher:** Use this option to ensure that the certificate used by the software publisher has not been revoked.<br><br>• **Timestamp:** Use this option to ensure that the certificate used by the organization that time-stamped the active content has not been revoked. |

For more information about software restriction policies rules and enforcement options, see "How Software Restriction Policies Work" in this Collection.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "What Is Resultant Set of Policy?" in this collection.
- "What Is Group Policy Management Console?" in this collection.
- For more information about the command-line tools listed in this document, see "Core Group Policy Tools and Settings" in this collection.
- For GPMC download information, see the Group Policy Management Console page.
- Tools and Settings Collection

# Scripts Extension Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft Windows 2000 Server and Windows Server 2003 use Group Policy settings to provide directory-based desktop configuration management. These Group Policy settings are stored in a Group Policy object (GPO). You can create and manage GPOs by using the Group Policy Object Editor.

General classes of Group Policy settings are implemented as extensions to the operating system. These extensions are packaged as dynamic-link libraries (DLLs), and exist as two types of extension:

- Server-side extension — A Microsoft Management Console (MMC) snap-in that is used for administration and configuration.
- Client-side extension (CSE) — An extension that runs on the client computer that interprets and applies the MMC type settings to the client computer, also known as the target computer.

The GPO settings are configured via the server side extension, and applied to individual computers and users, by the client-side extension (CSE). The Group Policy engine is the framework used to manage all client side extensions, including the Scripts client-side extension.

**In this subject**


- What Is Scripts Extension?
- How Scripts Extension Works
- Scripts Extension Tools and Settings

# What Is Scripts Extension?

Updated: December 13, 2010

Applies To: Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Vista

**What Is Scripts Extension?**

Active Directory domains use the Scripts extension to run scripts on client computers. The Scripts extension consists of two components; an MMC server-side snap-in and a client-side extension (CSE). A Group Policy CSE is a dynamic-link library (DLL) on the client (target) that interacts with the Group Policy infrastructure and that implements Group Policy scripts on the target. You can control how the Scripts CSE implements scripts by using the Group Policy Object Editor on the server. The following figure shows Group Policy Object Editor, a Microsoft Management Console (MMC) snap-in with the scripts node.

**Group Policy Object Editor**




Before the introduction of the GPO Scripts extension, user logon scripts were added to the user account object. Only one script was added to the object, and the script was a user script; no computer scripts were allowed.

The Scripts CSE now allows the following four script types, corresponding to four of the target events triggering the Scripts extension:

- User logon scripts
- User logoff scripts
- Computer startup scripts
- Computer shutdown scripts

Other Scripts extension trigger events are:

- User policy refresh event
- User policy force refresh event
- Computer policy refresh event
- Computer policy force refresh event

- Secedit tool or gpupdate

These events tell the computer which scripts need to be run at each script execution opportunity; these events do not run the scripts.

Scripts are added to the Group Policy object (GPO), and multiple scripts for each script type can be added to the GPO. A script can be written in any language supported by the client computer, with Windows Script Host (WSH) supported languages and command files being the most common. WSH languages include VBScript and JScript, which are the languages used for the sample scripts included with Group Policy Management Console (GPMC).

Script names and their command-line arguments are stored in the registry when script policy processes. You can specify Administrative Templates Group Policy settings that modify how Group Policy scripts behave by using the options shown in the following table. These settings are located under Computer Configuration|Administrative Templates|System|Scripts.

**Computer Configuration Script Options**

| Option | Description |
|---|---|
| Allow logon scripts when NetBIOS or WINS is disabled<br><br>This setting requires Windows Vista or later. | Enable this setting to allow user logon scripts to run if NetBIOS or WINS is disabled during logons across forests without the DNS suffixes being configured. |
| Run logon scripts synchronously | Enable this option to force the system to run the scripts synchronously, one after another. This option exists for computer and user configuration, which can have a different value. In case of conflict, the computer configuration setting prevails. |
| Run startup scripts asynchronously | The default setting is for scripts to run synchronously (and hidden). Use this option to optimize the startup/logon processes so users can logon before startup scripts have finished. |
| Run startup scripts visible | Enable this option to run startup scripts in a visible window. |
| Run shutdown scripts visible | Enable this option to run shutdown scripts in a visible window. |
| Maximum wait time for Group Policy scripts | Use this option to set the script timeout interval. The default interval is 600 seconds (10 minutes), and valid intervals range from 0 to 32000 seconds.<br><br>If the value of this setting is greater than 900 seconds (15 minutes), for computers running Windows 7 you need to raise the value under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\gpsvc for PreshutdownTimeout to a corresponding value that is slightly greater than the value of the "Maximum wait time for Group Policy scripts" value. For example, in order for "Maximum wait time for Group Policy scripts" to work with a value of 30 minutes, you should set the **gpsvc for PreshutdownTimeout** registry entry to a value of 1800000 (The registry setting value is in milliseconds). |

| | |
|---|---|
| Run Windows PowerShell scripts first at computer startup, shutdown<br><br>This setting requires Windows 7 or Windows Server 2008 R2. | Enable this option to run PowerShell scripts before non-PowerShell scripts at computer startup and shutdown. |
| Run Windows PowerShell scripts first at user logon, logoff<br><br>This setting requires Windows 7 or Windows Server 2008 R2. | Enable this option to run PowerShell scripts before non-PowerShell scripts at logon and logoff. |

You can specify user configuration Group Policy settings that modify how Group Policy scripts behave by using the options shown in the following table. These settings are located under User Configuration|Administrative Templates|System|Scripts

**User Configuration Script Options**

| Option | Description |
|---|---|
| Run logon scripts synchronously | Enable this option to force the system to run the script synchronously, one after another. This option also exists for Computer configuration, which can have a different value. In case of conflict, the computer configuration setting prevails. On computers running Windows XP, the explorer starts before scripts finish running. |
| Run legacy logon scripts hidden | Enable this option to run legacy, Windows NT version 4–type logon scripts hidden. By default these run visible and so can be canceled by users. |
| Run logon scripts visible | Enable this option to run logon scripts in a visible window. |
| Run logoff scripts visible | Enable this option to run logoff scripts in a visible window. |
| Run Windows PowerShell scripts first at user logon, logoff<br><br>This setting requires Windows 7 or Windows Server 2008 R2. | Enable this option to run PowerShell scripts before non-PowerShell scripts at logon and logoff. |

**Change History**

| Date | Revision |
|---|---|

| | |
|---|---|
| December 13, 2010 | Updated the topic so it applies to Windows Vista, Windows Server 2008, Windows 7, and Windows Server 2008 R2.<br><br>Added information about configuring the value of the **gpsvc for PreshutdownTimeout** registry entry for computers running Windows 7 when the Maximum wait time for the Group Policy scripts setting is greater than 900 seconds.<br><br>Added settings that require Windows Vista, Windows Server 2008, Windows 7, or Windows Server 2008 R2.<br><br>Deleted row that incorrectly listed Run logoff scripts synchronously. |

# How Scripts Extension Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Scripts Extension Works**

The Scripts extension is one of the Group Policy extensions that automate configuration of the Group Policy infrastructure. The Scripts extension consists of two parts. The server-side extension, a Microsoft Management Console (MMC) snap-in, enables administrators to make enterprise-wide script Group Policy settings. The Scripts client side extension (CSE) implements those settings on the client computer. A separate process, Userinit, actually runs the scripts.

### In this section


- Scripts Extension Architecture
- Scripts Extension Protocols
- Scripts Extension Physical Structure
- Scripts Extension Processes and Interactions
- Network Ports Used by Scripts Extension


**Scripts Extension Architecture**

The Group Policy Scripts client-side extension (CSE) is managed by the Group Policy core, which stores a list of Group Policy objects (GPOs) received from Active Directory. The GPO list is then passed into the Scripts CSE where the scripts.ini file of each GPO is read from the Sysvol folder. This information required for script processing is stored in the client registry. When Group Policy is finished processing, the WinLogon process retrieves the registry information, and then creates a Userinit process that actually runs the scripts. During logoff and shutdown, WinLogon again creates a Userinit process to run the relevant scripts. The Userinit process uses the ShellExecute command to run each script. The following figure shows the Scripts extension architecture and the client-server interaction of the two extension components.

### Scripts Extension Architecture




Components important to the Scripts extension, seen in the preceding figure, are described in the table below. Components not seen in the figure, but important to the process are also described.

### Scripts Logical Architecture Components

| Component | Description |
|---|---|

| | |
|---|---|
| Group Policy engine | The Framework that manages and implements the Group Policy settings and configurations, made by the admin, across all client-side extensions (CSE). Userinit actually executes the script; not the CSE. Userenv.dll is the Group Policy engine module. |
| Scripts client-side extension (CSE) | The Scripts CSE is the component that is called by the Group Policy engine, and that applies the scripts setting. The Scripts CSE writes the relevant script information into the registry. It does not run the scripts. |
| WinLogon | WinLogon is the service that contains the Group Policy engine. |
| Resultant Set of Policy (RSoP) snap-in | Displays the results of Group Policy, including what scripts have run and when they were last run. For more information about RSoP, see "What Is Resultant Set of Policy?." |
| Userinit | Called by Winlogon to run Group Policy scripts by calling ShellExecute. |
| Local GPO | Contains Group Policy settings for the local computer, including potential scripts policies. |

The CSE registration information is written at setup to the HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Winlogon\ GPExtensions registry key. This registry key structure exists on both the target and on the domain controller systems.

#### Scripts Extension Protocols

The various protocols used within the Scripts extension architecture are summarized in the following table.

**Scripts Extension Protocols**

| Protocol | Description |
|---|---|
| Win32 | Windows API function call. |
| NTLM or Kerberos | Authentication protocol used by the client to authenticate with Active Directory. |
| Server Message Block (SMB) | SMB is a file access protocol used by Userinit to access the Sysvol folder. |
| Distributed Component Object Model (DCOM) | DCOM is used by CSEs to communicate with Windows Management Instrumentation (WMI). |

#### Scripts Extension Physical Structure

Most Group Policy settings are stored in the Group Policy template, part of the domain controller's file system. Group Policy templates are stored in the %systemroot%\SYSVOL\sysvol*domainname*\policiessubfolder. The root of each Group Policy template contains a Gpt.ini file. Each Group Policy template contains a Machine and User subfolder, both of which contain a Scripts subfolder that contains the Scripts.ini file. The Machine Scripts subfolder contains a Shutdown and a Startup subfolder. The User Scripts subfolder contains a Logoff and a Logon subfolder. These four subfolders contain their respective Group Policy scripts to be implemented on the target by the Scripts CSE.

Components important to the Scripts CSE are described in the following table.

**Scripts Physical Architecture Components**

| Component | Description |
|---|---|
| Sysvol folder | The Sysvol folder contains the set of folders shared on each domain controller that stores *file-system* domain information, as compared to *registry* domain information. The Sysvol folder is one of the locations that the Scripts.ini file is stored. The other location is the Local GPO. |
| Group Policy template | The Group Policy template is part of the GPO, and is the portion of Group Policy settings stored in the file system and shared between domain controllers. The Group Policy template makes up the majority of the Group Policy settings. Each Group Policy template folder is a subfolder of the Policies folder and has a globally unique identifier (GUID). |

| | |
|---|---|
| Gpt.ini | Gpt.ini is a file in the Group Policy template root folder that stores the GPO version number of the Group Policy template. The Group Policy CSEs use this version number to confirm that directory and file components are synchronized correctly. |
| ADM folder | The ADM folder stores the System.adm file, among others. |
| Machine folder | The Machine folder stores GPO Computer settings files. One of the subfolders is the Scripts folder. |
| User folder | The User folder stores GPO User settings files. One of the folders is the Scripts folder. |
| Scripts folder | The Scripts folder, a subfolder of both the Machine and User folders, stores the Scripts.ini file and references scripts to be run. |
| Scripts.ini file | The Scripts.ini file is stored on the Sysvol folder in the Scripts folder and stores the path to each script. |
| External Scripts server | A server, other than the domain controller, on which the actual scripts can be stored. |
| Scripts Server-side extension Snap-in | A snap-in that is a server side component that loads with Group Policy Object Editor. The snap-in displays a Scripts node under Windows Settings of both Computer Configuration and User Configuration. |
| Logon folder | A subfolder of the Sysvol folder containing any user logon scripts to be implemented by the Scripts CSE. |
| Logoff folder | A subfolder of the Sysvol folder containing any user logoff scripts to be implemented by the Scripts CSE. |
| Startup folder | A subfolder of the Sysvol folder containing any computer startup scripts to be implemented by the Scripts CSE. |
| Shutdown folder | A subfolder of the Sysvol folder containing any computer shutdown scripts to be implemented by the Scripts CSE. |

**Scripts Extension Processes and Interactions**

The Scripts CSE processes all applied GPOs. The Scripts CSE reads the command-line options for the scripts from the Scripts.ini file. If a full path is specified, the full path is stored in the registry. If only a name is specified, the CSE creates the full path to the file in the Sysvol folder, prior to storing it in the registry. On Windows XP and better, The Scripts CSE stores this information in the client's registry, in the following keys:

- HKEY_LOCAL_MACHINE\Software\Policies\Windows\System\Scripts
- HKEY_CURRENT_USER\Software\Policies\Windows\System\Scripts

The script type is stored as one of the following registry values:

- Logon
- Logoff
- Startup
- Shutdown

The Scripts CSE sorts scripts by using the following rules:

- By GPO priority
- By priority order within each GPO

Scripts are usually stored in the Startup and Shutdown subfolders of the Machine\Scripts folder or in the Logon and Logoff subfolders of the User\Scripts folder. Scripts can be stored in a different location, or on a different server. However, for the script to run, the script file must be available and accessible (read access) for it to be run. Note that the user must have read access for logon and logoff scripts and the computer must have access for startup and shutdown scripts. Both are members of the Authenticated Users group, so by default both have access to run their respective scripts. The script is run from the file path location stored in Scripts.ini and the script is not cached on the client computer. If the network is unavailable when scripts are set to run, they will not run.

**Network Ports Used by Scripts Extension**

The following table summarizes the network protocols and port numbers used by the Scripts CSE. The Scripts extension impersonates the user and computer for network communication.

**Port Assignments for Scripts CSE**

| Task | Port (Protocol) |
|---|---|
| Retrieve GPO list | TCP 398 (LDAP) |
| Retrieve Group Policy container | TCP 398 (LDAP) |
| Request Distributed File System (DFS) referral for the Sysvol folder on the domain controller | TCP 445 (SMB) |
| Determine Sysvol DFS replica location for the Sysvol folder on the domain controller | TCP 445 (SMB) |
| Open and read Gpt.ini | TCP 445 (SMB) |
| Open and read Group Policy template settings files | TCP 445 (SMB) |
| Return Group Policy template file | TCP 445 (SMB) |

# Scripts Extension Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scripts Extension Tools and Settings**

This section contains references to tools associated with the Scripts Extension. This section also contains settings, Registry entries, WMI Classes, and Network information associated with the Scripts extension.

**In this section**

- Scripts Extension Group Policy Settings
- Related Information

**Scripts Extension Group Policy Settings**

You can use the Windows Settings\Scripts node under both Computer Configuration and User Configuration to specify scripts to be run. Note that the scripts server-side extension permits you to specify the path to the scripts that need to be run. It does not create the script. The following tables describe the Computer and User Configuration Windows Settings\Scripts node settings.

**Computer Configuration\Windows Settings\Scripts**

| Option | Description |
|---|---|
| Scripts (Startup\Shutdown) | Use this node to specify the script files to be run when the computer starts up and shuts down. You can add or delete script files by selecting the button to open the explorer window. Add or delete the script files using this window. |

**User Configuration\Windows Settings\Scripts**

| Option | Description |
|---|---|
| Scripts (Logon\Logoff) | Use this node to specify the script files to be run when the user logs on and logs off. You can add or delete script files by selecting the button to open the explorer window. Add or delete the script files using this window. |

You can also specify Administrative Templates Group Policy settings that modify how Group Policy scripts behave. There is an Administrative Templates node for both Computer Configuration and User Configuration in Group Policy Object Editor. The following tables summarize the Group Policy settings that are associated with Scripts Extension in a Windows 2000 or greater domain.

**Computer Configuration\Administrative Templates\System\Scripts**

| Option | Description |
|---|---|
| Run logon scripts synchronously | Enable this option to force the system to run the scripts synchronously, one after another. This option exists for computer and user configuration. Computer setting has precedence. |
| Run startup scripts asynchronously | Use this option to optimize the startup/logon processes so users can logon before startup scripts have finished. By default startup scripts run synchronously. |
| Run startup scripts visible | Enable this option to run startup scripts in a visible window. By default scripts run invisible, users are not able to interact or cancel scripts that run invisible. |
| Run shutdown scripts visible | Enable this option to run shutdown scripts in a visible window. By default scripts run invisible, users are not able to interact or cancel scripts that run invisible. |
| Maximum wait time for Group Policy scripts | Use this option to set the script timeout interval. The default interval is 600 seconds (10 minutes), and valid intervals range from 0 to 32000 seconds. This affects both computer and user scripts that run synchronously. |

**Computer Configuration\Administrative Templates\System\Group Policy**

| Option | Description |
|---|---|
| Scripts policy processing | When enabled, this changes how the scripts client-side extension processes scripts policy. (Note that this affects only the processing of scripts policy—it does not affect how or when scripts are run).<br><br>When enabled, options include:<br><br>• Processing scripts policy when on a slow link. (Note that this will not stop existing scripts from running scripts policy; it stops scripts from processing scripts policy).<br><br>• Processing scripts policy during periodic background refresh.<br><br>• Always process scripts policy, even if there were no changes detected. |

**User Configuration\Administrative Templates\System\Scripts**

| Option | Description |
|---|---|
| Run logon scripts synchronously | Enable this option to force the system to run the script synchronously, one after another. This option also exists for Computer configuration, which can have a different value. In case of conflict, the computer configuration setting prevails. On computers running Windows XP, the Windows Explorer starts before scripts finish running. |
| Run legacy logon scripts hidden | Enable this option to run legacy, Windows NT 4–type logon scripts hidden. By default these run visible and so can be canceled by users. |
| Run logon scripts visible | Enable this option to run logon scripts in a visible window. By default scripts run invisible, users are not able to interact or cancel scripts that run invisible. |
| Run logoff scripts visible | Enable this option to run logoff scripts in a visible window. By default scripts run invisible, users are not able to interact or cancel scripts that run invisible. |

There are policies that control which Microsoft Management Console (MMC) snap-ins can and cannot be run. By setting these policies, a user can be restricted from administering scripts policies by restricting the scripts MMC snap-in. These policies are shown in the following table.

**User Configuration\Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions**

| Option | Description |
|---|---|
| Scripts (Logon/Logoff) | The default setting is Not Configured, which causes the Restrict users to the explicitly permitted list of snap-ins setting to apply. Enabling Scripts (Logon/Logoff) permits the snap-in. Disabling the setting prohibits the snap-in. |
| Scripts (Startup\Shutdown) | The default setting is Not Configured, which causes the Restrict users to the explicitly permitted list of snap-ins setting to apply. Enabling Scripts (Startup\Shutdown) permits the snap-in. Disabling the setting prohibits the snap-in. |
| Restrict users to the explicitly permitted list of snap-ins | The default setting of Not Configured and Disabled permits any snap-in that is not explicitly prohibited. If this setting is Enabled, only explicitly Enabled snap-ins are permitted. |

For more information about Group Policy settings, see the Group Policy Settings Reference for Windows Server 2003.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Group Policy Settings Reference for Windows Server 2003
- Tools and Settings Collection

# Wireless Network Policies Extension Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Wireless Network Policies Extension is a Group Policy extension that you can use to automate configuration of Wireless Network Group Policy settings. Configuration of Wireless Group Policy settings requires, at minimum, one domain controller running a Microsoft Windows Server 2003 operating system. The Wireless Network Policies Extension is supported on client computers running Windows XP Professional (Service Pack 1 and later) and using IEEE 802.11 access with IEEE 802.1x authentication in a wireless network.

**In this subject**


- What Is Wireless Network Policies Extension?
- How Wireless Network Policies Extension Works
- Wireless Network Policies Extension Tools and Settings

# What Is Wireless Network Policies Extension?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Wireless Network Policies Extension?**
This section describes the Wireless Network Policies Extension. On computers running Microsoft Windows Server 2003 Active Directory, domain administrators use the Wireless Network Policies Extension to automate wireless Group Policy management. The Wireless Network Policies Extension is one of the Group Policy client-side extensions that "plug" into the Group Policy infrastructure.

Wireless Network Policies Extension is one of the general classes of Group Policy settings that are implemented as extensions to the operating system. These extensions are packaged as DLLs, and exist as two types of extensions:

- MMC snap-in—An administration and configuration extension that runs on the server.
- Client-Side Extension—An extension that interprets and applies the MMC type settings to the target system (client), and which runs on the client.

The following figure shows an overview of a wireless network. The MMC snap-in fits in the Directory box on the server, and the CSE fits in the Wireless Client box.

**Wireless Network Overview**




Organizations need a way to centrally manage wireless Group Policy configuration for multiple users. The administrator needs to make specific wireless Group Policy decisions, including:

- Group Policy object (GPO) application criteria.
- Number of GPOs required.
- Location of GPOs.
- Number of Wireless LAN (WLAN) profiles required.

The Wireless Network Policies Extension provides central wireless Group Policy management that includes:

- Active Directory-based group filtering to provide Wireless Network Policy to a single computer global group.

- A single GPO Wireless Network Policy.
- A single GPO created and applied from the ForestRootDomain object.
- A single WLAN profile configured for 802.1x-compliant organizations. For multiple profiles, you have the option to support a phased migration from a legacy production WLAN.

This solution implementation is seen in the following figure. The figure shows the Group Policy Object Editor for a GPO, named gpo_Name. The Wireless Network (IEEE 802.11) Policies node contains a single wireless policy, named GPfield. It is from this central location that you implement the above wireless Group Policy solution.

**Group Policy Object Editor for the gpo_Name GPO**




You can use the Wireless Network Policies Extension to configure Active Directory-based Group Policy configurations for client computers. Specifically, the Wireless Network Policies Extension enables you to:

- Make wireless network Group Policy settings to protect a wireless network from unauthorized access by client computers with a compatible WLAN adapter.
- Protect data transferred over the wireless network, based on Group Policy settings.
- Make wireless network Group Policy settings to implement either certificate-based or password-based authentication for client computers accessing the wireless network.
- Configure Group Policy to add User authentication to computer authentication, ensuring round-the-clock network availability.

This Group Policy affects wireless client interaction between wireless Access Points (AP), the RADIUS server, and other wireless networks.

# How Wireless Network Policies Extension Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Wireless Network Policies Extension Works**
This section discusses how the Wireless Network Policies Extension works.

**In this section**

- Wireless Network Policies Extension Architecture
- Wireless Network Policies Extension Protocols
- Wireless Network Policies Extension Physical Structure
- Network Protocols Used by Wireless Network Policies Extension

**Wireless Network Policies Extension Architecture**
The Wireless Network Policies Extension consists of two subsystems:

- Wireless MMC snap-in
- Wireless Client-Side Extension (CSE)

**Wireless MMC Snap-in**
The Wireless MMC snap-in is the server-side tool used by the administrator to make the wireless Group Policy settings into Group Policy objects (GPOs). The snap-in functionality for the GPO is accessible from **Edit** on the pop-up menu for the GPO within the Group Policy Management Console (GPMC). The following figure shows the snap-in for the GPO named **gpo_Name**, and shows the **Properties** dialog box for the policy named **GPfield**.

**Group Policy Object Editor MMC Snap-in**




**Wireless Client-Side Extension (CSE)**

Gptext.dll, located in the **systemroot%\system32** folder on the target, is the Client-Side Extension (CSE) component, which implements the settings made by the administrator on the server using the snap-in. The Wireless Extension downloads these settings from the domain controller to the target and stored in the registry. This must be done over a wired (cable) network initially, due to a certificate services requirement. This initial download cannot be done over a wireless connection. Subsequent Computer Configuration Group Policy settings can be downloaded to a wireless device after the initial download. There is no UI for the CSE on the target.

Wireless CSE settings are stored in a subkey under the following registry key on the target:

**HKLM\Software\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtensions**

The wireless CSE subkey name is a GUID and the wireless subkey and its values are seen in the figure below. The other GUID subkeys correspond to the other CSEs.

**Target Wireless CSE Registry Key and Values**




**Wireless Network Policies Extension Protocols**

The figure below shows the conceptual view of wireless network access. The Wireless Network Extension MMC snap-in sits, conceptually, in the Management "box" of infrastructure services. The CSE sits in the Wireless Client "box" of the WLAN Component.

**Wireless Network Overview**




Components and protocols significant to the Wireless Network Policies Extension are included in the following table:

**Wireless Network Policies Extension Components and Protocols**

| Component | Description |
|---|---|
| IEEE 802.11 | A shared, wireless local area network standard. |
| IEEE 802.1x | A Port-based Network Access Control standard. |
| IEEE 802.11i | A draft standard that specifies improvements to WLAN networking security. |
| Wi-Fi Protected Access (WPA) | An interim standard, agreed to by wireless vendors, that is used until IEEE 802.11i is ratified. |
| wireless Access Point (wireless AP) | The more general term is Network Access Service (NAS). This is a device defined by the 802.11 standard that bridges the wireless and wired networks. |
| Remote Authentication Dial-In User Service (RADIUS) | Performs authentication and authorization for many types of networks, including wireless. |
| Internet Authentication Service (IAS) | A Microsoft implementation of RADIUS. |
| Extensible Authentication Protocol (EAP) | An authentication protocol used by, and required by the IEEE 802.1x standard. |
| Protected Extensible Authentication Protocol (PEAP) | One of two EAP authentication protocols available in the Wireless extension. Both are supported by Windows XP Service Pack 1. |
| EAP-Transport Layer Security (EAP-TLS) | One of two EAP authentication protocols (PKI-based) supported by Windows XP Service Pack 1. EAP-TLS is the EAP type recommended for certificate-based client authentication. |
| Public key infrastructure (PKI) | PKI is an implementation of public-key cryptography. Public-key cryptography provides data privacy, robust user and computer authentication, and documentation of action taken; all in a distributed security model. |

| | |
|---|---|
| certificate | A credential used for authentication. |
| Transport Level Security (TLS) | TLS provides mutual authentication, negotiation, and endpoint key determination. |
| Microsoft Challenge-Handshake Authentication Protocol version 2 (MS-CHAP v2) | An authentication protocol. |
| PEAP-EAP-MSCHAP-v2 | The recommended EAP type for password-based client authentication. |
| Wired Equivalent Privacy (WEP) algorithm | An algorithm on which the IEEE 802.11 WLAN Standard authentication and encryption services are based. |

**Wireless Network Policies Extension Physical Structure**

The physical structure of the Wireless Network Policies Extension, and how it fits into a wireless infrastructure, is seen in the Wireless Network Overview figure in the Protocols section above. Settings made by the administrator in the server-side snap-in are transferred to the CSE during initial wired connection between target and domain. These settings are transferred based on target (Computer) identity. Many of the settings are stored in the CSE registry key seen in the Target Wireless CSE Registry Key and Values figure, above.

One of the settings stored in a different registry location is the EAP type. This setting is stored in the following registry key on the client, and on the domain controller:

- HKLM\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\[#]
    - [#] = 13 for Smart Card or other certificate
    - [#] = 25 for Protected EAP (PEAP)
    - [#] = 26 for Secured password (EAP-MSCHAP v2)
    - [#] = 4 for MD-5 Challenge (not configurable via Extension)

On the wireless client (target) at machine startup, all CSEs are managed by the Group Policy Engine (GPE). The GPE stores a list of all GPOs received from the Active Directory. The GPE passes the GPO list into the Wireless CSE.

Group Policy templates (GPTs), the file system portion of Group Policy Settings, are stored in the **[%systemroot%\SYSVOL]\sysvol\domainname\policies** subfolder of the sysvol tree. The root of each GPT contains a GPT.INI file. Each GPT contains a MACHINE and USER subfolder. Wireless policy is located in the MACHINE subfolder.

The GPT.ini file, stored on the domain controller SYSVOL, contains the GPO version number of the Group Policy template (GPT). The GptTmpl.inf file, stored in the MACHINE\Microsoft\Windows NT\SecEdit folder for each GPO, contains machine GPO settings including Wireless Network Policy Extension settings.

At client machine startup, when Group Policy is finished processing, Winlogon gets the CSE registry information, creates a Userinit process, and passes Userinit the Wireless CSE information. The GPE, running in the same process as Winlogon, calls Userinit; Winlogon doesn't call Userinit. The Wireless CSE is registered with Winlogon in the registry, as seen in the Target Wireless CSE Registry figure above. Other Wireless Network Group Policy settings are also stored in the registry, and are retrieved and applied at this time.

Windows XP and Windows Server 2003 stores computer and user authentication behavior settings in the AuthMode value in the following registry key:

HKLM\Software\Microsoft\EAPOL\Parameters\General\Global

Valid AuthMode registry values are:

- 0 for With user authentication
- 1 for With user re-authentication
- 2 for Computer only

The following table summarizes the components that interact during target startup.

**Physical Structure Components**

| Component | Description |
|---|---|
| Group Policy Engine | The Framework that manages and implements the Group Policy settings and configurations, made by the admin, across all Client Side Extensions (CSE). Userenv.dll is the GPE module. The GPE runs in the same process as Winlogon. It is from that process that Userinit is called. The component actually calling Userinit is the GPE, not Winlogon. |
| Wireless Client-Side Extension (CSE) | The Wireless CSE is the central communication point for the GPE, and is contained within gptext.dll located in the %systemroot%\system32 folder. This DLL runs on the client system that is responsible for interpreting and implementing the Group Policy Management Console wireless settings made by the administrator through the wireless snap-in. It also runs on the domain controller to provide the Group Policy Object Editor snap-in. |
| WinLogon | WinLogon is the component that calls the Wireless CSE at machine startup and timed intervals. Winlogon is the only system component that actively interacts with the Group Policy Engine. All CSEs are registered with WinLogon in the registry under the following registry key on both the domain controller and the client machine:<br><br>HKLM\Software\Microsoft\WindowsNT\CurrentVersion\Winlogon\GPExtensions<br><br>**Note**<br><br>Direct registry manipulation of CSE behavior is not recommended. Most behavior settings are available for modification in the Group Policy Object Editor. |
| Userinit | Called by the process in which both Winlogon and the GPE are running. |
| SYSVOL | The SYSVOL is the set of folders shared on each domain controller, which store *file-system* domain information, as compared to *registry* domain information. Sysvol is one of the locations in which the GPT.INI file is stored. |
| Group Policy template (GPT) | The GPT is part of the GPO, and is the portion of Group Policy Settings stored in the file system and shared between domain controllers. The GPT makes up the majority of the Group Policy settings. Each GPT folder is a subfolder of the Policies folder and has a GUID for a name. |
| GPT.INI | GPT.INI is a file in the GPT root folder that stores the GPO version number of the GPT. The Group Policy Client Side Extensions use this version number to confirm that directory and file components are synchronized correctly. |
| ADM (Adm) folder | The ADM folder stores the System.adm file, among others. |
| MACHINE folder | The Machine folder stores GPO Computer settings files. One of the subfolders is the SecEdit folder. |
| SecEdit folder | The SecEdit folder, a subfolder of the MACHINE folder, stores the GptTmpl.inf file. |

**Network Ports Used by Wireless Network Policies Extension**

The following table summarizes the network protocols and port numbers used by the Wireless Network Policies Extension.

**Port Assignments for Wireless Network Policies Extension**

| Task | Port (Protocol) |
|---|---|
| Retrieve GPO list | TCP 398 (LDAP) |
| Retrieve Group Policy container | TCP 398 (LDAP) |
| Request Distributed File System (DFS) referral for the Sysvol folder on the domain controller | TCP 445 (SMB) |

| | |
|---|---|
| Determine Sysvol DFS replica location for the Sysvol folder on the domain controller | TCP 445 (SMB) |
| Open and read Gpt.ini | TCP 445 (SMB) |
| Open and read Group Policy template settings files | TCP 445 (SMB) |
| Return Group Policy template file | TCP 445 (SMB) |
| RADIUS Messages | UDP 1812, 1813 |

# Wireless Network Policies Extension Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless Network Policies Extension Tools and Settings**

**In this section**


- Wireless Network Policies Extension Tools
- Wireless Network Policies Extension Registry Entries
- Wireless Network Policies Extension Group Policy Settings
- Related Information


This section contains references to tools associated with the Wireless Extension. This section also contains settings, Registry entries, and Network information associated with the Wireless extension.

**Wireless Network Policies Extension Tools**

The following tools are associated with Wireless Network Policies Extension.

**Wlanmon.dll: Wireless Monitor**

**Category**

Wlanmon.dll is a Microsoft Management Console (MMC) snap-in.

**Version compatibility**

Wlanmon.dll runs only on Windows Server 2003.

Wlanmon.dll is used to monitor Access Point information and Wireless Client information.

**Wlsnp.dll: Wireless Group Policy Snap-In**

**Category**

Wlsnp.dll is an MMC snap-in.

**Version compatibility**

Wlsnp.dll runs only on Windows Server 2003.

Wlsnp.dll is used to create and edit Wireless policy on Windows Server 2003 operating systems.

**Wireless Network Policies Extension Registry Entries**

The following registry entries are associated with Wireless Network Policies Extension.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. You should not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**GPExtensions\ {0ACDD40C-75AC-47ab-BAA0-BF6DE7E7FE63} (Domain Controller and Target)**

**Registry path**

HKEY_CURRENT_USER\Software\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The Registry settings in this section are for a domain controller running Windows Server 2003 and a target running Windows XP SP1 or later.

This key contains Wireless Group Policy extension information.

For more information about this registry entry, see the Registry Reference for Windows Server 2003.

**GPExtensions\{0ACDD40C-75AC-47ab-BAA0-BF6DE7E7FE63} (Domain Controller and Target)**

**Registry path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**
The Registry settings in this section are for a domain controller running Windows Server 2003 and a target running Windows XP SP1 or later.

This key contains specific Wireless Group Policy extension information.

For more information about this registry entry, see the Registry Reference for Windows Server 2003.

**Wireless (Domain Controller and Target)**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\windows\

Note that Windows in the Registry path is "windows" for Windows Server 2003 and "Windows" for Windows XP.

**Version**
The Registry settings in this section are for a domain controller running Windows Server 2003 and a target running Windows XP SP1 or later.

This registry key stores specific wireless Group Policy settings.

For more information about this registry entry, see the Registry Reference for Windows Server 2003.

**Wireless Network Policies Extension Group Policy Settings**
You can specify Administrative Templates Group Policy settings that modify wireless Group Policy. You can also modify settings in the Wireless Group Policy Properties dialog box. The following two tables summarize Administrative Template settings.

**Computer Configuration\Administrative Templates\System\Group Policy**

| Group Policy Setting | Description |
|---|---|
| Wireless Policy Processing | Options are:<br>• Not Configured (default)<br>• Enabled<br>• Disabled |

**Computer Configuration\Administrative Templates\Network\Network Connections**

| Group Policy Setting | Description |
|---|---|
| Prohibit use of Internet Connection Sharing on your Domain Name System (DNS) domain network | Options are:<br>• Not Configured (default)<br>• Enabled<br>• Disabled |
| Prohibit use of Internet Connection Firewall on your DNS domain network | Options are:<br>• Not Configured (default)<br>• Enabled<br>• Disabled |

| | |
|---|---|
| Prohibit installation and configuration of Network Bridge on your DNS domain network | Options are:<br>• Not Configured (default)<br>• Enabled<br>• Disabled |
| IEEE 802.1x Certificate Authority for Machine Authentication | Options are:<br>• Not Configured (default)<br>• Enabled<br>• Disabled |

The remaining tables summarize settings from the various tabs on the Wireless Group Policy Properties dialog box.

**General Tab**

| Option | Description |
|---|---|
| Name | String of your choice |
| Description | String of your choice |
| Check for policy Changes every *x* minutes | 180 is the default |
| Networks to access | Drop-down list box choices are:<br>• Any available network (access point preferred) (default)<br>• Access point (infrastructure) networks only<br>• Computer-to-computer (ad hoc) networks only |
| Use Windows to configure wireless network settings for clients (check box) | Default is check box selected. Clearing the check box disables the option. |
| Automatically connect to non-preferred networks (check box) | Default is check box cleared. Selecting the check box enables the option. |

**Preferred Networks Tab**

| Option | Description |
|---|---|
| Networks: (list) | The list, by default, is empty. Adding a network to the list enables the Move Up and Move Down buttons. |
| Add, Edit and Remove buttons | The Add and Edit button dialog boxes are the same. |

**Network Properties Tab**

| Option | Description |
|---|---|

| | |
|---|---|
| Network Name (SSID) | String for network name. |
| Description | String for description |
| Wireless Network Key (WEP) (3 check boxes) | This network requires a key for the following:<br>• Data encryption (WEP-enabled) (default is selected)<br>• Network authentication (Shared mode) (default is cleared)<br>• The key is provided automatically (default is selected) |
| This is a computer-to-computer (ad hoc) network; wireless access points are not used. | Default is check box cleared. Selecting the check box enables the option. |

**IEEE 802.1x Tab**

| **Option** | **Description** |
|---|---|
| Enable network access control using IEEE 802.1x | Default is check box selected. Clearing the check box disables the option. |
| EAPOL-Start message | Drop-down list box choices are:<br>• Do not transmit<br>• Transmit (default)<br>• Transmit per IEEE 802.1x |
| Parameters (seconds) | Text Box settings are:<br>• Max start (default is 3)<br>• Start period (default is 60)<br>• Held period (default is 60)<br>• Authentication period (default is 30) |
| EAP type | Drop-down list box choices are:<br>• Smart Card or other certificate (default)<br>• Protected EAP (PEAP) |
| Authenticate as guest when user or computer information is unavailable | Default is check box cleared. Selecting the check box enables the option. |
| Authenticate as computer when computer information is available | Default is check box selected. Clearing the check box disables the option. |

| | |
|---|---|
| Computer authentication | Drop-down list box choices are:<br>• With user authentication<br>• With user re-authentication (default)<br>• Computer only |

**EAP type (Smart Card or other certificate) Settings Tab**

| Option | Description |
|---|---|
| When connecting | Option button choices are:<br>• Use my smart card<br>• Use a certificate on this computer (default)<br>  • Use simple certificate selection (Recommended) [(check box) (default is selected)] |
| Validate server certificate (check box) | Default is check box selected. Clearing the check box disables the option. |
| Connect to these servers (check box) | Default is check box cleared and list box disabled. Selecting the check box enables the text box, in which you enter the servers. |
| Trusted Root Certificate Authorities (list box) | Default is none of the check boxes in the list are selected. Selecting the check box enables the option. |
| Use a different user name for the connection (check box) | Default is check box is cleared. Selecting the check box enables the option. |

**EAP type (Protected EAP (PEAP)) Settings Tab**

| Option | Description |
|---|---|
| Validate server certificate (check box) | Default is check box selected. Clearing the check box disables the option. |
| Connect to these servers (check box) | Default is check box cleared and list box disabled. Selecting the check box enables the text box, in which you enter the servers. |
| Trusted Root Certificate Authorities (list box) | Default is none of the check boxes in the list are selected. Selecting the check box enables the option. |
| Select Authentication Method | Drop-down list box options are:<br>• Secured password (EAP-MSCHAP v2) [(default)]<br>• Smart Card or other certificate |
| Enable Fast Reconnect (check box) | Default is check box cleared. Selecting the check box enables the option. |

**Select Authentication Method (Secured Password (EAP-MSCHAP v2)) Properties Tab**

| Option | Description |
|---|---|

| | |
|---|---|
| When connecting: (check box) | Automatically use my Windows logon name and password (and domain if any) [(default is check box is selected)] |

**Select Authentication Method (Smart Card or other certificate) Properties Tab**

| Option | Description |
|---|---|
| When connecting | Option button choices are:<br>• Use my smart card<br>• Use a certificate on this computer (default)<br>  • Use simple certificate selection (Recommended) [(check box) (default is selected)] |
| Validate server certificate (check box) | Default is check box selected. Clearing the check box disables the option. |
| Connect to these servers (check box) | Default is check box cleared and list box disabled. Selecting the check box enables the text box, in which you enter the servers. |
| Trusted Root Certificate Authorities (list box) | Default is none of the check boxes in the list are selected. Selecting the check box enables the option. |
| Use a different user name for the connection (check box) | Default is check box cleared. Selecting the check box enables the option. |

For more information about Group Policy settings, see the Group Policy Settings Reference for Windows Server 2003.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Registry Reference for Windows Server 2003
- Group Policy Settings Reference for Windows Server 2003
- Tools and Settings Collection

# Folder Redirection Extension Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Folder Redirection is a feature of IntelliMirror that allows users and administrators to redirect the path of a folder to a new location. The new location can be a folder on the local computer or a shared directory on a network. Users have the ability to work with documents on a server as if the documents were located on the local hard disk. For example, if you redirect the My Documents folder, which is usually stored on the computer's local hard disk drive, to a network location, the documents in the folder will be available to the user from any computer on the network.

**In this subject**


- What Is Folder Redirection Extension?
- How Folder Redirection Extension Works
- Folder Redirection Extension Tools and Settings

# What Is Folder Redirection Extension?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Benefits of Folder Redirection
- Common Folder Redirection Scenarios
- Folder Redirection Dependencies on or Interactions with Other Technologies
- Related Information


The Folder Redirection extension, a feature of IntelliMirror, enables an administrator to redirect the location of certain folders in the user profile to different path, such as a shared network location. When these redirected folders are accessed, either by the operating system or by applications, the operating system automatically redirects to the location on a network share specified by the administrator. From a user perspective, this feature is similar to a roaming profile because users have the same folder settings regardless of which computers they use. These settings remain on the network share and enable users to access their files from any computer they use to logon on to the network. The following figure illustrates the relationships and interactions of the Folder Redirection extension with other networking components and Group Policy features.

**Folder Redirection Extension Relationships and Interactions**




The following special folders can be redirected:

- Application Data
- Desktop
- My Documents
    - My Pictures
- Start Menu

**Note**

- My Pictures folder is no longer shown in the Folder Redirection Node. To simplify the user interface and to help support the best practice that the My Pictures folder should always follow the My Documents folder, the My Pictures folder is not shown in the Folder Redirection node for new GPO's. If you have previously redirected the My Pictures folder separately, the My Pictures node will still appear.

**Benefits of Folder Redirection**

When fully deployed, IntelliMirror uses Active Directory and Group Policy for policy-based management of user desktops. A Windows XP or Windows 2000 Professional desktop can be automatically configured to meet specific requirements of a user's business roles, group memberships, and location. Group Policy and the Active Directory are not necessary for every IntelliMirror feature. Some of the features can be set on the local level or through local polices. An organization can tailor IntelliMirror to meet its needs. When planning to use IntelliMirror, an organization should assess which features it needs and then implement the technology required.

By using IntelliMirror on both the server and client, administrators can protect and manage user data and settings. Non-recoverable data from local workstations can be copied to servers, where it can be easily backed up and centrally managed. Personalized data, applications, and settings can follow each user to different computers throughout the network. Administrators can easily replace faulty computers and restore all user data and settings on a new computer.

Previously, administrators who wanted to redirect folders to the network had to use logon scripts to change registry values. In Windows Server 2003 and Windows XP, the same task can be accomplished by using the Folder Redirection extension of Group Policy. For example, you could redirect a user's My Documents directory to \\Server\Share\%username%.

Folder Redirection provides the following benefits:

- **Improved Roaming User Profile performance.** You can use Folder Redirection to reduce the time it takes to log on to and log off from the network. The My Documents folder is part of the Roaming User Profile (RUP), which means that the My Documents folder and its contents are copied back and forth between the client computer and the server each time users log on and log off. Redirecting the My Documents folder to a location outside of the user profile can significantly decrease the amount of time it takes to log on and log off because not all of the data in the user profile is transferred to the desktop each time the user logs on — only the data that user requires is transferred.

- **User documents available from any computer.** To ensure that users' documents are available when they roam from one computer to another, you can use Folder Redirection to redirect user documents to a server location accessible from any computer connected to the domain.

- **User documents available when not connected to the corporate network.** Folder Redirection can be used with Offline Files technologies to make users' network-based My Documents folder available when they are disconnected from the corporate network. By default in Windows XP and Windows Server 2003, any redirected shell folders such as My Documents, Desktop, Start Menu, and Application Data are automatically made available Offline.

  **Note**

  - Offline Files are disabled by default in Windows Server 2003.

- **Increased security and availability of user data.** When you use Folder Redirection to redirect user data to a network location, the data stored on a shared network server can be backed up as part of routine system administration. This protects user data and requires no action on the part of the user.

- **Reduced support costs and user down-time.** If there are multiple hard disk drives in a user's computer, user data and files can be redirected to a hard disk on the user's local computer other than the one on which the operating system is installed. This protects the user's data if the operating system needs to be reinstalled or if computer hardware needs to be replaced.

- **Control allocation of disk space.** Administrators can use Group Policy to set disk quotas, which limit the amount of space taken up by users' folders.

**Common Folder Redirection Scenarios**

This section presents sample scenarios that illustrate some of the practical uses of Folder Redirection. Each of the scenarios fits into an entire picture or can be seen as a separate event. Each scenario also shows how the components of IntelliMirror benefit

the entire organization by reducing the time and effort associated with maintaining the computing environment.

**The new hire**
One of the most critical and time consuming IT tasks is setting up a computer for a new hire. In an organization that uses IntelliMirror, the new hire logs on to a new computer and finds documents and shortcuts already on the desktop. There are shortcuts to common files, URLs, and folders that are useful to all employees (for example, the employee handbook, shortcuts to the departmental shared documents store, and the user's departmental guidelines and procedures). Desktop options, application configurations, Internet settings, and so on, are all configured to the corporate standard, ensuring that if the user needs to call the help desk, the support staff knows what configuration the user started with.

In this example, the user gets a pre-configured user profile that was set up for all new users, and was configured before the new hire logged on to the network. The administrator configured a computer to look and behave according to the corporate standard, and then, using the User Profile utility built into the System Control Panel application, copied the user profile to a Default User folder on the domain controller's Netlogon share. When the new hire logged onto the network for the first time, Windows copied this default profile to the local computer and used this profile as the basis for the new hire's profile. In addition to configuring the default profile the user received, the administrator also used Group Policy to redirect the user's My Documents folder to a network location, so that the user's documents are safely stored on a network server and can be backed up regularly.

**The laptop user**
In this scenario, a laptop user working at the office creates several documents and saves them to his or her My Documents folder. After saving documents, the user logs off, unplugs the laptop computer from the network and takes it home. While at home and off the network, the user continues to edit the documents saved earlier in the My Documents folder.

The user returns to the office and logs on to the network. Since the user has worked offline, a dialog box appears advising the user that data in My Documents has changed and is being synchronized with the network copy.

In this scenario, the user's My documents folder has been redirected to a network server, the documents are transparently saved to the network location and also saved in the local computer's cache (because the network folder is setup to be available offline), so that they are available when the computer is disconnected from the network.

The whole process can be transparent to the user; the experience is no different than saving documents to the local hard disk.

As soon as the user reconnects to the network, IntelliMirror attempts to reconnect to the network location of the redirected folders. When IntelliMirror reconnects, it determines if there are differences in the data between the local copy of the folder and the network copy. In this scenario, the user has made modifications to a document on the local computer. IntelliMirror identifies this change and prompts the user to update the version stored on the network.

**Computer replacement**
In this scenario the computer that a user is working on suddenly stops working because of a complete hardware failure. The user calls the support line, and receives a new computer, loaded with only the Windows XP Professional operating system. Without waiting for technical assistance, the user plugs in the new computer, connects it to the network, and starts it. When the user logs on to the corporate network, he or she will find that the desktop has the same appearance as the original computer that it replaced. It will have the same color scheme, the user's preferred background picture on the screensaver, and all the application icons, shortcuts, and favorites. More importantly, all the user's data files will have been restored.

In a disaster recovery scenario, IntelliMirror assists in getting the user's computer replaced and running quickly with the minimum of support. In this example, because the user was configured to use roaming user profiles, a copy of the user's working environment was safely stored on a network server. When the new computer arrived, the user was able to log on and the server copy of the user's profile was downloaded to the new computer. An administrator could also have used Folder Redirection to redirect the user's key folders such as My Documents and Application Data, to ensure that the user's documents were safely stored on the server.

**A shared computer environment**
In this scenario, a user works in a department where the computer he or she uses might change from day to day—a call center or IT support environment, for example. The user is working on an important document late one night when the shift ends. The user saves the document and logs off the computer. When the user returns to work the next day, he or she logs onto the first available computer—a different computer from the one used the previous night. The user logs onto the network, and sees that the desktop has the same look and feel as the original computer. The user opens the My Documents folder on the desktop and finds the document exactly where he or she saved it and continues the work started the previous night.

In this example, the user was configured to use roaming user profiles, so that a copy of the user's working environment was stored on a network server. When the user logged onto the computer, the user's existing preferences, shortcuts and documents were copied to the local computer, so that the user could continue working as if using the original computer. A variation on this scenario is using roaming profiles in conjunction with Folder Redirection. Users can have the same work

environment and have access to the same documents on any computer. Changes made on one computer are synchronized with the other computer the next time the user logs on.

**Folder Redirection Dependencies**
The Folder Redirection extension is used most effectively when combined with other related Group Policy and Windows Server technologies, such as Roaming User Profiles, Offline Files, and Synchronization Manager.

**Folder Redirection with Roaming User Profiles**
You can also combine Folder Redirection and roaming user profiles to decrease logon and logoff times for roaming and mobile users. A common scenario is to redirect the My Documents and My Pictures folders, and allow the Application Data, Desktop, and Start Menu folders to roam with the profile. In addition to improved availability and administrative benefits from storing the data on the network, users also realize performance gains when using low-speed network connections and in subsequent logon sessions. Not all of the data in the user profile is transferred to the desktop each time the user logs on — only the data that user accesses during a session is transferred. Because only some documents are copied, performance is improved when the user's' profiles are copied from the server.

When you combine the use of Folder Redirection and roaming user profiles, you can also provide fast computer replacement. If a user's computer needs to be replaced, the user's data can quickly be copied from the server locations to a replacement computer. Combining Folder Redirection with roaming profiles gives the benefits of roaming profiles, while minimizing network traffic caused by synchronization of the profile.

### Note

- To decrease initial logon time to a new computer, you should redirect the location of the My Documents folder outside of the user's roaming profile.

**Local profiles**
Using Folder Redirection with local profiles can provide some of the benefits of roaming profiles (such as having a user's data available at any computer or maintaining data on the server) without the need to implement roaming profiles. Remember though, using Folder Redirection with a local profile would only result in the user's documents and files being available from all computers. To have settings and configurations move with the user, you would need to use roaming profiles.

**Folder Redirection with Offline Files**
By using Offline Files, users can continue to work with a cached copy of files stored on a network location, even when they are not connected to a network. If your organization has mobile users with portable computers, Offline Files gives them access to their files when they are not connected to the network, and ensures that they are always working with the current version of network files. Offline Files technology stores the data in the computer's cache to make network files available offline. The view of shared network items that you have made available offline remains as it is when connected, even if users lose a connection to or manually disconnect from the network. Users can continue to work with the Offline Files as they normally do when online. Users have the same access permissions to those files and folders as when they are connected to the network. When the network connection is restored, any changes they made while working offline are updated to the network so that both locations are synchronized.

Offline Files is a stand-alone technology, which means that you do not need to use it with Folder Redirection. However, using Offline Files with Folder Redirection allows for greater flexibility and availability of user files. For example, if a shortcut to a file is available offline, that file is made available offline, but if a shortcut to a folder is available offline, the contents of that folder are not available offline. If you pair the two technologies, Offline Files and Folder Redirection, both the shortcut and the folder are available offline.

Shared files or folders on a Windows Server 2003 or Windows XP-based network can be available offline. By default in Windows XP and Windows Server 2003, any redirected shell folders such as My Documents, Desktop, Start Menu, and Application Data are automatically made available offline. You can also make files available for offline use from any computer that is sharing files using server message block-based file and printer sharing, including Windows 2000, Windows 95, Windows 98, and Windows NT 4.0.

### Note

- Make sure that you turn off Offline Folders for shares where roaming user profiles are stored. If you do not turn off Offline Folders for a user's profile, you might experience synchronization problems since both Offline Folders and Roaming Profiles will try to synchronize the files in a user's profile.

**Synchronization Manager**
When using Offline Files and folders, users can synchronize all network resources by using the Synchronization Manager. The

Synchronization Manager can be set to automatically synchronize some or all resources. For example, users can set certain files and folders to be synchronized every time they log on or off the network. The Synchronization Manager quickly scans the system for any changes, and if it detects changes, the resources are automatically updated. Only resources that have changed are updated—vastly speeding up the synchronization process.

**Related Information**

The following resources contain additional information that is relevant to this section.

- The Group Policy Settings Reference for Windows Server 2003.
- The Group Policy Administrative Tools topic in this collection.

# How Folder Redirection Extension Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Folder Redirection Architecture
- Folder Interactions Processes and Interactions
- Related Information


Folder Redirection, a client-side extension (CSE) of Group Policy, enables an administrator of an organization to set policies that redirect users' folders to a different location, such as a shared network folder. When used in combination with Roaming User profiles and offline files, Folder Redirection enables administrators to centrally manage and protect user data and settings including policy-based management of user desktops.

Folder Redirection provides the greatest benefit when used with other components of IntelliMirror in an Active Directory environment. Folders can be redirected only from client computers running Windows 2000 or Windows XP.

**Folder Redirection Architecture**

The Folder Redirection client-side extension consists of a dynamic-linked library (DLL), fdeploy.dll, a redirection information file, {25537BA6-77A8-11D2-9B6C-0000F8080861}.ini, which includes information about Folder Redirection options specific to each user profile on the computer, and {1C08E84D-F112-4252-978B-EC82A225CC20}.ini which stores information about previous locations used for Folder Redirection.. Fdeploy.dll is loaded during the logon process by the Group Policy engine. The Folder Redirection policies are passed to fdeploy.dll, which then examines the policy settings and redirects user folders based on those settings.

The first step in Folder Redirection is to get a list of all Folder Redirection policies and evaluate them. The policies are handed down to the fdeploy.dll by the Group Policy engine through an API call into fdeploy.dll. The policies come in the form of two linked lists. The first list is a list of added policies. This refers to a list of Folder Redirection policies that are still applicable to the user. The second list is a list of deleted policies that refers to policies that used to be applicable to the user but are no longer applicable. The changes that result from Folder Redirection are then updated with the shell. The client-side cache, also referred to as offline files, is also updated to reflect the files moved by Folder Redirection.

**Folder Redirection Client-side Extension Architecture**




The following table describes the components that interact with Folder Redirection extension.

**Folder Redirection Components**

| Component | Description |
|---|---|

| | |
|---|---|
| Fdeploy.dll | Client-side extension DLL that processes policy settings for Folder Redirection. |
| {25537BA6-77A8-11D2-9B6C-0000F8080861}.ini | Locally cached redirection information file specific to each user profile. |
| {1C08E84D-F112-4252-978B-EC82A225CC20}.ini | A file indicating the status of the previous locations used for redirected folders. This file is used to detect server location changes and user name changes. |
| Server (domain controller) | In an Active Directory forest, the domain controller is a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. |
| Active Directory | Active Directory, the Windows-based directory service, stores information about objects on a network and makes this information available to users and network administrators. Administrators link Group Policy objects (GPOs) to Active Directory containers such as sites, domain, and organizational units (OUs) that include user and computer objects. In this way, policy settings can be targeted to users and computers throughout the organization. |
| Sysvol | The Sysvol is a set of folders containing important domain information that is stored in the file system rather than in the directory. The Sysvol folder is, by default, stored in a subfolder of systemroot folder (%\systemroot\sysvol\sysvol) and is automatically created when a server is promoted to a domain controller. The Sysvol contains the largest part of a GPO: the Group Policy template, which includes Administrative Template-based policy settings, security settings, script files, and information regarding applications that are available for software installation. It is replicated through the File Replication Service (FRS) between all domain controllers in a domain. |
| Group Policy object (GPO) | A GPO is a collection of Group Policy settings, stored at the domain level as a virtual object consisting of a Group Policy container and a Group Policy template. The Group Policy container, which contains information about the properties of a GPO, is stored in Active Directory on each domain controller in the domain. The Group Policy template contains the data in a GPO and is stored in the Sysvol in the /Policies subdirectory. GPOs affect users and computers that are contained in sites, domains, and OUs. |
| Local Group Policy object | The Local Group Policy object (Local GPO) is stored on each individual computer, in the hidden **Windows\System32\GroupPolicy** directory. Each computer running Windows 2000, Windows XP Professional, Windows XP 64-Bit Edition, Windows XP Media Center Edition, or Windows Server 2003 has exactly one Local GPO, regardless of whether the computers are part of an Active Directory environment.<br><br>Local GPOs are always processed, but are the least influential GPOs in an Active Directory environment, because GPOs based on Active Directory have precedence. |
| Winlogon | A component of the Windows operating system that provides interactive logon support, Winlogon is the service in which the Group Policy engine runs. |
| Group Policy engine | The Group Policy engine is the framework that handles common functionalities across registry-based settings and client-side extensions (CSEs). |
| Client-side extensions | CSEs run within dynamic-link libraries (DLLs) and are responsible for implementing Group Policy at the client computer.<br><br>The CSEs are loaded on an as-needed basis when a client computer is processing policy. |
| File system | The NTFS file system on client computers. |

| | |
|---|---|
| Registry | A database repository for information about a computer's configuration, the registry contains information that Windows continually references during operation, such as:<br>• Profiles for each user.<br>• The programs installed on the computer and the types of documents that each can create.<br>• Property settings for folders and program icons.<br>• The hardware on the system.<br>• Which ports are being used.<br><br>The registry is organized hierarchically as a tree, and it is made up of keys and their subkeys, hives, and entries.<br><br>Registry settings can be controlled through Group Policy, specifically, Administrative Templates (.adm files). Windows Server 2003 comes with a predefined set of Administrative Template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a GPO. These .adm files are stored in two locations by default: inside GPOs in the Sysvol folder and in the **Windows\inf** directory on the local computer. |
| Event log | The Event log is a service, located in Event Viewer, that records events in the system, security, and application logs. |
| Help and Support Center | The Help and Support Center is a component on each computer that provides HTML reports on the policy settings currently in effect on the computer. |
| Resultant Set of Policy (RSoP) infrastructure | All Group Policy processing information is collected and stored in a Common Information Model Object Management (CIMOM) database on the local computer. This information, such as the list, content, and logging of processing details for each GPO, can then be accessed by tools using Windows Management Instrumentation (WMI). |
| WMI | WMI is a management infrastructure that supports monitoring and controlling of system resources through a common set of interfaces and provides a logically organized, consistent model of Windows operation, configuration, and status.<br><br>WMI makes data about a target computer available for administrative use. Such data can include hardware and software inventory, settings, and configuration information. For example, WMI exposes hardware configuration data such as CPU, memory, disk space, and manufacturer, as well as software configuration data from the registry, drivers, file system, Active Directory, the Windows Installer service, networking configuration, and application data. WMI filtering in Windows Server 2003 enables you to create queries based on this data. These queries (WMI filters) determine which users and computers receive all of the policy configured in the GPO where you create the filter. |
| SMB | Server Message Block (SMB) protocol is the primary method of file and print sharing. Folder Redirection client-side extension use SMB to access the Sysvol as well as back up and retrieve files to a remote file system. The client computer also uses SMB to read the Sysvol on the domain controller. |

The following table shows the location of the files used by Folder Redirection.

**Folder Redirection Component Locations**

| **Component** | **Location** |
|---|---|
| Fdeploy.dll | %windir%\system32 |
| {25537BA6-77A8-11D2-9B6C-0000F8080861}.ini | %USERPROFILE%\Local Settings\Application Data\Microsoft\Windows\File Deployment\ |
| Fdeploy.log | %windir%\debug\usermode\ |

**Locally Cached Redirection Information File**
Whenever a folder is redirected, fdeploy.dll saves the details of the redirection in a local file. This data is used later by fdeploy.dll to determine how to process redirected folders if a user's personal name (UPN) or security group membership changes. The file is also used with roaming user profiles.

The redirection information file is stored in the %USERPROFILE%\Local Settings\Application Data\Microsoft\Windows\File Deployment\ folder.

The file has one section per folder currently redirected through Folder Redirection policy. If there is no policy applicable for a particular folder, then that section is removed from this file. For each folder, the following information is maintained:

- **User Name.** The user name of the user when the redirection was performed. This information is used to handle UPN changes.
- **Path.** The path to which the folder is redirected.
- **Flags.** The flags describing the Folder Redirection settings.
- **Group.** The string representation of the SID for the security group that was responsible for that particular path to get chosen as the redirection destination.
- **GPO.** The GUID of the GPO that was responsible for redirection.

### An example cached redirection file:

```
[Desktop]
Username=username
Path=%USERPROFILE%\Desktop
Flags=31
Group=S-1-5-21-397955417-626881126-188441444-3038017
GPO={AFC36721-F1D5-46BA-981C-FC33BDE293D7}
```

**Folder Redirection Processes and Interactions**
Folder Redirection client-side extension interacts with the Group Policy engine as well as the shell.

Folder Redirection processing contains five steps:

1. Determine which user folders to redirect based on changes to Group Policy settings at time of logon.
2. Determine the target location specified for redirection and confirm the user has access rights to that location.
3. If the target folder does not exist, the folder is created and the appropriate access control list (ACL) rights are set.
4. If the folder exists, access rights and folder ownership are checked.
5. If desired, the files contained within specified folders are moved to the new location, which also deletes them from the source folder if the source folders are local.

If the policy for Folder Redirection does not specify that the user be granted exclusive access to the redirected location, fdeploy.dll confirms only that the destination folder exists. If the target folder does not exist, it is created. However, if the policy requires that the user be given exclusive access to the destination folder, fdeploy.dll does one of the following, depending on whether the destination folder already exists or not.

### Note

- With Windows XP SP1 the user must be the owner of the folder on the server. If not, the ownership check fails, and redirection fails unless a policy is set to ignore this. This was added in SP1.
- If the destination already exists, fdeploy.dll checks whether the user is the owner of the folder. If the folder is owned by the user, fdeploy.dll proceeds; otherwise it fails and aborts the redirection of that folder.

- If the destination does not exist, fdeploy.dll creates the folder and then sets ACLs on it so that only the user and the local system account have full access to it, but no other users have any access to it. Fdeploy.dll also ensures that the folder does not inherit any Access Control Entries (ACE) from its parent. Fdeploy.dll then sets the user as the owner of the folder. This is done because if the user belongs to the local administrators group on the destination server, the local administrators group becomes the owner of the group by default. By setting the user as the owner of the folder you can assure that any quota accounting for that user is done correctly. This also ensures that any subsequent redirections to the same location from other computers do not fail due to the ownership check performed on pre-existing folders described previously. If configuration of either ownership or Access Control Lists (ACL) fails, the process is aborted and redirection fails for the folder specified. If not, redirection proceeds and the files are copied to the redirected location.

Once this is done, fdeploy.dll deletes the files from the source. However, there is an exception to this. If the source is a network location and the destination is a local location, the files are not deleted from the source. This ensures that any subsequent redirection operations from other computers also get a copy of the files on the local computer as required by the policy.

**Folder Redirection Interaction with the Shell**
There are two main points of interaction with the shell. The first point is when fdeploy.dll determines the current location of a folder. For this, it uses the **SHGetFolderPath** API from shell32.dll. The second point of interaction is when fdeploy.dll needs to notify the shell about the new location of a folder after it applies redirection settings. For this, it uses the **SHSetFolderPath** API from shell32.dll.

In addition, there is one other point of interaction which is specific to redirection of the My Documents folder. In the Windows shell, the properties dialog box for the My Documents folder has a **Target** tab. Settings on the **Target** page enable a user to change the location of this folder. If the folder is redirected through policy, this functionality is not desirable and should be disabled. Therefore, fdeploy.dll needs to set the policy which disables this functionality. This is done by creating a DWORD value called **DisablePersonalDirChange** under the key HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer.

**Folder Redirection and Environment Variables**
Environment variables can be used to define the path to the redirected location. Supported environment variables include %USERNAME% and %USERPROFILE%.

These are the only environment variables supported because other variables are not defined when the Folder Redirection extension is loaded by Winlogon. Note that fdeploy.dll does not do any sort of validation to ensure that only certain allowed variables are present in the paths. It merely expands the paths and if any variables are not defined, they stay as they are. Homedir redirection is an exception. Redirection to the home directory of the user is only allowed for the My Documents folder; fdeploy.dll does perform validation to ensure that other folders cannot be redirected to that location.

**Folder Redirection and Mapped Drives**
Because Folder Redirection is processed early in the logon process, drives mapped with logon scripts (including the homedrive for folders other than My Documents), the Folder Redirection client side extension is not able to redirect to these locations. At the time that redirection takes place, the drives do not exist hence redirection fails.

**Using Offline Files Settings**
Using offline files settings on a shared network folder where user data is stored is especially useful for users of portable computers. It is recommended that you use Folder Redirection in conjunction with offline files.

The offline files feature of Windows, also referred to as client-side caching, is designed to make network access of files more efficient by keeping a local copy of the files in the offline files cache. It also facilitates access to those files when the user is not connected to the network. Since Folder Redirection moves a user's files to and from network locations, it is important for it to move files within the local offline files cache to provide a seamless and transparent experience to the user. In most cases when you use Folder Redirection, you will combine it with offline files so that users can access cached copies of the redirected folders when disconnected from the network. The following table provides some recommended settings for redirecting folders.

### Note

- All redirected folders are pinned by default.

### Recommended Configuration for Offline Files

| Redirected Folder | Recommended Offline Folder Settings |
|---|---|

| | |
|---|---|
| My Documents | Auto-caching for documents or manual caching for documents (if you want users to have to manually make files and folders available offline). |
| My Pictures | Auto-caching for documents or manual caching for documents (if you want users to have to manually make files and folders available offline). |
| Application Data | Auto-caching for programs. |
| Desktop | Auto-caching for programs if the desktop is read-only. |

**Folder Redirection and Software Installation Policies**

When logon optimization is enabled, a user might need to log on to a computer twice before Folder Redirection policies and software installation policies are applied. This is because application of these types of policies requires the synchronous policy application. During a policy refresh (which is asynchronous), the system sets a flag that indicates that the application of Folder Redirection or a software installation policy is required. The flag forces synchronous application of the policy at the user's next logon.

Because background refresh is the default behavior in Windows XP, Folder Redirection and Software Installation might require as many as three logons to apply changes.

This behavior exists because Folder Redirection and Software Installation cannot apply during an asynchronous or background application of policy. Folder Redirection can be applied only when processed synchronously.

Here is a sample scenario showing how polices are applied:

- An administrator deploys a software package to User A.
- User A logs on fast and receives a background (asynchronous) application of policy.
- Because the policy application was asynchronous, the software that was set to be installed cannot be installed at this time. Instead the computer is tagged, indicating that software needs to be installed.
- The next time the user logs on, the computer instead logs on the user synchronously to allow the software package to be installed. (This is the same behavior as Windows 2000). This results in one extra logon for the software to be installed.
- In the case of Advanced Folder Redirection, because policy is evaluated based on security group membership three logons will be required: the first logon to update the cached user object (and security group membership), the second logon for policy to detect the change in security group membership and require a foreground policy application, and the third logon to actually apply Folder Redirection policy in the foreground.

  **Note**

  - When a client running Windows XP logs onto a Windows 2000 or Windows Server 2003 Active Directory, all Software Installation policy settings for Windows 2000 clients will be applied and work successfully on the Windows XP client.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Core Group Policy Technical Reference in the Group Policy collection.
- Core Group Policy Tools and Settings in the Group Policy collection.

# Folder Redirection Extension Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Folder Redirection Tools
- Folder Redirection Settings
- Folder Redirection Registry Keys
- Related Information


**Folder Redirection Tools**
The following is a list of tools available for supporting and configuring Folder Redirection.

**Folder Redirection Logging**
Errors and failures of Folder Redirection processing are logged as events in the Application Event log. In addition, Folder Redirection can provide a detailed log to aid troubleshooting. To create a detailed log file for Folder Redirection, add a value to the following registry key:

HKLM\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics

Set: FdeployDebugLevel = Reg_DWORD 0x0f

A log file, Fdeploy.log, is created and stored in the *%windir%*\debug\usermode\ folder.

**Group Policy Administrative Tools**
The following Group Policy administrative tools are used to administer Group Policy components:

- Group Policy Management Console
- Group Policy Object Editor
- Resultant Set of Policy Snap-in

For more information about these tools, please see the "Group Policy Administrative Tools" section.

**Folder Redirection Settings**
There are two types of Folder Redirection. The type you choose to use depends on your network configuration and security considerations.

- **Basic.** This option redirects everyone's folder to the same location. All folders affected by this Group Policy object will be stored on the same network share.
- **Advanced.** This option specifies locations for various user groups. Folders are redirected to different network shares based on security group membership. For example, folders belonging to users in the Accounting group can be redirected to the Finance server, while folders belonging to users in the Sales group can be redirected to the Marketing server.

The following list describes the settings available for redirecting folders.

- **Grant the user exclusive rights to My Documents.** If selected, this sets the NTFS security descriptor for the %username% folder to Full Control for the user and local system only; this means that administrators and other users do not have access rights to the folder. This option is enabled by default. Rights to this folder can later be modified to enable administrators to access these folders.

- **Move the contents of My Documents to the new location.** This setting moves any file the user has in the local My Documents folder to the server share. This option is enabled by default.

- **Leave the folder in the new location when policy is removed.** This setting specifies that files remain in the new location when the Group Policy object no longer applies. This option is enabled by default.

- **Redirect the folder back to the local user profile location when policy is removed.** If enabled, this setting specifies for the folder to be copied back to the local profile location if the Group Policy object no longer applies.

There are two additional settings available only for My Pictures folder.

- **Make My Pictures a subfolder of My Documents.** If selected, when the My Documents folder is redirected, My Pictures remains a subfolder of My Documents. By default, My Pictures automatically follows the My Documents folder.

- **Do not specify administrative policy for My Pictures.** If selected, Group Policy does not control the location of My Pictures; this is determined by the user profile.

The following table outlines the behavior of redirected folders, and the files contained in them, when the policy for Folder Redirection is removed.

**Summary of Folder Redirection Policy Removal Settings**

| Move the contents of special folder to the new location setting | Setting selected for the Policy Removal option | When policy is removed |
|---|---|---|
| Enabled | Redirect the folder back to the user profile location when policy is removed. | The special folder returns to its user profile location.<br><br>The folder contents are copied back to the user profile location.<br><br>The contents are not deleted from the redirected location.<br><br>The user can continue to access the contents, but only on the local computer. |
| Disabled | Redirect the folder back to the user profile location when policy is removed. | The special folder returns to its user profile location.<br><br>Note**:** In this case, the folder contents are not copied or moved to the user profile location. As a result, the user can no longer see the contents. |
| Either Enabled or Disabled | Leave the folder in the new location when policy is removed. | The special folder remains at its redirected location.<br><br>The contents remain at the redirected location.<br><br>The user continues to have access to the contents at the redirected folder. |

**Note**

- Changing the redirection option to "Not Configured" does not redirect the folder to the local profile. This option means that Folder Redirection is not configured—if a folder was previously redirected it will continue to be redirected to the previous location. If administrators wish to return the folder to the local user profile they should use the **Redirect to the local user profile** setting.

**Folder Redirection Registry Keys**

The following registry entries are associated with Folder Redirection.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. You should not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in

unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**Note**

- Folder Redirection extension does not own the User Shell Folder keys, but simply manipulates them.

The following tables describe the registry settings associated with Folder Redirection.

**HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders**

| Value Name | Type | Purpose | Platforms |
|---|---|---|---|
| Personal | REG_EXPAND_SZ | Populated with redirected path from policy. | W2K, WinXP |
| My Pictures | REG_EXPAND_SZ | Populated with redirected path from policy. | W2K, WinXP |
| Desktop | REG_EXPAND_SZ | Populated with redirected path from policy. | W2K, WinXP |
| AppData | REG_EXPAND_SZ | Populated with redirected path from policy. | W2K, WinXP |
| Start Menu | REG_EXPAND_SZ | Populated with redirected path from policy. | W2K, WinXP |
| Startup | REG_EXPAND_SZ | Populated with redirected path from policy. | W2K, WinXP |
| Programs | REG_EXPAND_SZ | Populated with redirected path from policy. | W2K, WinXP |

**HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\Shell Folders**

| Value Name | Type | Purpose | Platforms |
|---|---|---|---|
| Personal | REG_SZ | Populated with expanded redirected path from policy. | W2K, WinXP |
| My Pictures | REG_SZ | Populated with expanded redirected path from policy. | W2K, WinXP |
| Desktop | REG_SZ | Populated with expanded redirected path from policy. | W2K, WinXP |
| AppData | REG_SZ | Populated with expanded redirected path from policy. | W2K, WinXP |
| Start Menu | REG_SZ | Populated with expanded redirected path from policy. | W2K, WinXP |
| Startup | REG_SZ | Populated with expanded redirected path from policy. | W2K, WinXP |
| Programs | REG_SZ | Populated with expanded redirected path from policy. | W2K, WinXP |

**HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Diagnostics**

| Value Name | Type | Purpose | Platforms |
|---|---|---|---|
| FDeployDebugLevel | REG_DWORD | Determines the logging level of Folder Redirection. | W2K, WinXP |
| FDeployDebugBreak | REG_DWORD | Makes the Folder Redirection client break into the debugger at startup. | W2K, WinXP |

**HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Explorer**

| Value Name | Type | Purpose | Platforms |
|---|---|---|---|
| DisablePersonalDirChange | REG_DWORD | Prevent users from changing the location of the My Documents folder through the UI on the My Documents properties sheet if the folder is redirected through policy. | W2K, WinXP |

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtensions\{25537BA6-77A8-11D2-9B6C-0000F8080861}**

(This is registration information from fdeploy.dll)

| Value Name | Type | Purpose | Platforms |
|---|---|---|---|
| DllName | REG_EXPAND_SZ | Name of the dll. | W2K, WinXP |
| EventSources | REG_MULTI_SZ | For event logging. | W2K, WinXP |
| GenerateGroupPolicy | REG_SZ | Entry point for planning mode RSoP. | WinXP |
| ProcessGroupPolicyEx | REG_SZ | Entry point for diagnostic mode RSoP and regular policy application. | WinXP |
| NoBackgroundPolicy | REG_DWORD | Disposition during background policy application. | W2K, WinXP |
| NoGPOListChanges | REG_DWORD | Disposition when the GPO list has not changed. | W2K, WinXP |
| NoMachinePolicy | REG_DWORD | Disposition when processing machine policies. | W2K, WinXP |
| NoSlowLink | REG_DWORD | Disposition when connected over slow link. | W2K, WinXP |
| PerUserLocalSettings | REG_DWORD | Determine if history should be maintained per-user per computer-. | W2K, WinXP |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Core Group Policy Tools and Settings
- The Group Policy Settings Reference for Windows Server 2003
- The Tools and Settings Collection

# Internet Explorer Maintenance Extension Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Internet Explorer Maintenance Extension is used to manage Internet Explorer settings in a computing environment that uses Group Policy. The settings are grouped into the Browser User Interface, Connection Settings, Custom URLs, Security, and Program nodes.

Administrators can customize the appearance of the browser, preset and manage browser connection settings, set the default URLs displayed by the browser, and set the default programs used for each Internet service. Administrators can also preset the security zone, content rating, certification authority, and Authenticode settings.

The Internet Explorer Maintenance Extension of Group Policy is used to configure Internet Explorer only in domains using the Windows 2000 or Windows Server 2003 operating systems. It is possible to customize Internet Explorer before deployment and to manage Internet Explorer in other domain environments by using the Microsoft Internet Explorer Administration Kit (IEAK).

**In this subject**


- What Is Internet Explorer Maintenance Extension?
- How Internet Explorer Maintenance Extension Works
- Internet Explorer Maintenance Extension Tools and Settings

# What Is Internet Explorer Maintenance Extension?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Internet Explorer Maintenance Extension?**

**In this subject**


- Common Internet Explorer Maintenance Extension Scenarios
- Internet Explorer Maintenance Extension Dependencies on or Interactions with Other Technologies
- Related Information


The Internet Explorer Maintenance Extension is part of the Group Policy Object Editor, and it enables you to define an Internet Explorer configuration as part of a Group Policy Object (GPO). The GPOs are linked to Active Directory containers such as sites, domains, or organizational units, and enable administrators to manage the Internet Explorer configuration for multiple users on any computer joined to the domain.

The Internet Explorer Maintenance Extension uses two sets of extensions to accomplish its purpose: a snap-in extension to the Group Policy Object Editor, and a Client-Side Extension (CSE). The snap-in extension (ieaksie.dll) is used to configure Internet Explorer settings in a GPO. The Client-Side Extension (iedkcs32.dll) is a dynamic-link library (DLL) on the client computer that implements the Internet Explorer Maintenance Extension settings contained in the GPO.

The security risks inherent in exposure to the Internet for large computer networks require organizations to adopt a computer security policy. Users can accidentally compromise security for the entire network by misconfiguring their own, or a coworker's, computer. Providing users with information useful in solving problems or in accomplishing their jobs is not very effective if users are not familiar with the information, or if the information is difficult to access. The Internet Explorer Maintenance Extension enables administrators to stipulate custom favorites, links, security, interface, and other settings to specify an Internet Explorer configuration as part of a GPO. Administrators are able to enforce Internet-related security standards, and provide a common browser interface within the organization. Specifying custom links provides users easy access to useful information, and can aid in reducing end-user help calls.

**Common Internet Explorer Maintenance Extension Scenarios**

The Internet Explorer Maintenance Extension can be used to configure either mandatory or default settings for Internet Explorer.

**Mandatory Internet Explorer Settings**

Mandatory settings are used to enforce security, interface and other Internet Explorer settings by making users unable to change those settings. To configure mandatory settings, the settings are created while Internet Explorer Maintenance Extension is in Policy mode. When a GPO containing Policy mode settings is applied to a user configuration, the settings are reapplied only when Group Policy is forcefully reapplied to the target computer or changes are made to the GPO. Without further configuration using Group Policy Administrative Templates Extension settings, users are able to change their browser settings. There are two additional steps that must be completed to impose the mandatory status of the GPO:

- The `Computer Configuration\Administrative Templates\System\Group Policy\Internet Explorer Maintenance policy processing` setting must be configured to process settings even if they have not been changed.
- All of the settings in the `User Configuration\Administrative Templates\Windows Components\Internet Explorer\Internet Control Panel` node must be enabled.

**Default Internet Explorer Settings**

Internet Explorer Maintenance Extension can also be used to configure default settings for users the first time the GPO is applied to their computers. This provides users with the same starting configuration for their browsers, but enables them to personalize the configuration. To configure default settings, the settings are created while Internet Explorer Maintenance Extension is in Preference mode. Once a GPO containing Preference mode settings is applied to a user configuration, the settings are not reapplied unless there are changes to the GPO settings, even when Group Policy is forcefully reapplied to the target computer. Preference Mode also enables two additional groups of settings:

- **Corporate Settings.** These are used to configure temporary internet file settings, and download locations for ActiveX controls and Java code.

- **Internet Settings.** These are used to configure Internet Explorer link and text colors, AutoComplete settings, how often Internet Explorer checks for updates, and other advanced settings.

**Internet Explorer Maintenance Extension Settings**

The mode setting for Internet Explorer Maintenance Extension settings is exclusive within a GPO. Policy and Preference mode settings cannot coexist in the same GPO. The following table lists the available Internet Explorer Maintenance Extension Group Policy settings and their descriptions. It also notes if a setting is available in only Preference mode or Policy mode.

**Internet Explorer Maintenance Extension Configuration Settings**

| **Browser User Interface** | **Description** |
|---|---|
| Browser Title (Policy mode only) | This setting customizes the text that appears in the title bar of the Internet Explorer browser. The text is appended to the string "Microsoft Internet Explorer provided by." |
| Custom Logo (Policy mode only) | This setting replaces the static Internet Explorer logo in the upper-right corner of the Internet Explorer window with a custom static logo. It also replaces the animated Internet Explorer logo in the upper-right corner of the Internet Explorer window with a properly formatted custom animated logo. |
| Browser Toolbar Customizations | This setting customizes the background and buttons of the Internet Explorer toolbar. |
| **Connection** | **Description** |
| Connection Settings | This setting enables you to import connection settings from the Internet Control Panel Connections tab of an already-configured computer. It restricts how users can interact with connection settings through the System Policies and Restrictions page. It is not necessary to import current settings in order to set theses restrictions. Connections Settings is also used to remove old dial-up connection settings from users' computers. |
| Automatic Browser Configuration | Automatic Configuration (auto-config) is used to update a user's computer after deployment by specifying a URL to an .ins file, an auto-proxy URL, or both. You can set the interval in minutes for when auto-config will happen. If the interval value is left blank, or at zero, auto-config will happen only when the browser has been started and navigates to a page. |
| Proxy Settings | This setting specifies to which proxy servers, if any, users connect. |
| User Agent String (Policy mode only) | The user agent string is what the browser sends to visited servers to identify itself. It is often used to keep Internet traffic statistics. You can enter custom text that will be appended to the default Internet Explorer string. The default string is different for each platform. |
| **URLs** | **Description** |
| Favorites and Links | You can customize Favorites and Links by specifying the URLs. You can also specify the order of each folder, add an icon for each Favorite and Link, and import an existing folder structure. |
| Important URLs | You can specify a custom home page, search bar URL, and online support page. |
| **Security** | **Description** |
| Security Zones and Content Ratings | You can customize the settings of each security zone as well as customize the privacy settings. Content ratings enable you to prevent users from viewing sites with risky content. |

| | |
|---|---|
| Authenticode Settings | Authenticode enables you to designate software publishers as trustworthy. You can prevent users from adding new trusted publishers while using the browser. Enabling this lockdown does not prevent access to the Content control panel, but doing so is possible using the Administrative Templates Extension of Group Policy. |
| **Programs** | **Description** |
| Programs | You can import the current default programs settings. The programs selected specify which program Windows automatically uses for each Internet service. |
| **Advanced** | **Description** |
| Corporate Settings (Preference mode only) | These settings specify the location of a file containing settings used to make the browser work best for your corporation. These settings can be useful in reducing the cost of supporting applications in your corporation. Advanced features, as well as details of other features, can all be configured through these options. |
| Internet Settings (Preference mode only) | These settings specify the location of a file used to preset defaults for Internet Explorer settings that are not set through the IEAK Wizard. Most of these settings apply to defaults in the Internet Control Panel. |

**Internet Explorer Maintenance Extension Dependencies on or Interactions with Other Technologies**
Deployment of Internet Explorer Maintenance Extension settings requires Group Policy in a Windows 2000 or Windows 2003 Active Directory environment, and Windows 2000 Professional or Windows XP clients running Internet Explorer.

The Internet Explorer Maintenance Extension leverages the Internet Explorer Administration Kit (IEAK) management infrastructure, a part of Internet Explorer, in order to configure Internet Explorer. IEAK is an application used to deploy and manage customized Internet Explorer software packages. The packages can be used to configure Internet Explorer, just like the Internet Explorer Maintenance Extension, or to install completely customized versions.

**Internet Explorer Maintenance Extension and Internet Explorer Enhanced Security Configuration**
Another Technology the Internet Explorer Maintenance Extension interacts with is the Internet Explorer Enhanced Security Configuration component of Windows Server 2003. Internet Explorer Enhanced Security Configuration, also known as Microsoft Internet Explorer Hardening, reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings. You can deploy Internet Explorer Enhanced Security Configuration settings to target computers through Internet Explorer Maintenance only if Internet Explorer Enhanced Security Configuration is enabled on the target computer. This is because the Internet Explorer Maintenance CSE stores these settings on the target computer in the equivalent registry location from where they were imported. Those registry keys are read only by Windows Server 2003 because it is currently the only operating system that supports Internet Explorer Hardening.

If Internet Explorer Enhanced Security Configuration is enabled on your computer, the sites list is read from and written to **HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Internet Settings\ZoneMap\ESCDomains**.

If Internet Explorer Enhanced Security Configuration is disabled on your computer, or if your computer is running an earlier version of Windows, the sites list is read from and written to **HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Internet Settings\ZoneMap\Domains**.

Internet Explorer Enhanced Security Configuration settings (as imported by the computer when editing the GPO) can be applied only to computers on which Internet Explorer Enhanced Security Configuration is already enabled. Likewise, if Internet Explorer Enhanced Security Configuration is disabled on your computer, you can deploy security settings and trusted sites to other servers running with Internet Explorer Enhanced Security Configuration disabled. These variables are shown in the following table.

**Expected Behaviors on Target Computers for HKEY_CURRENT_USER Preferences with Internet Explorer Maintenance**

| Computer Importing Settings to the GPO (Your administrative computer) | Windows Server 2003 with Internet Explorer Enhanced Security Configuration enabled | Windows Server 2003 with Internet Explorer Enhanced Security Configuration disabled | Windows 2000 Server |
|---|---|---|---|

| | | | |
|---|---|---|---|
| | (writes to HKCU\...\ESCDomain) | (writes to HKCU\...\Domain | (reads HKCU\...\Domain) |
| Windows Server 2003 with Internet Explorer Enhanced Security Configuration enabled<br><br>(writes to HKCU\...\ESCDomain) | Security settings and trusted sites | No security settings or trusted sites | No security settings or trusted sites |
| Windows Server 2003 with Internet Explorer Enhanced Security Configuration disabled (writes to HKCU\...\Domain) | No security settings or trusted sites | Security settings and trusted sites | Security settings and trusted sites |
| Windows 2000 Server<br><br>(writes to HKCU\...\Domain) | No security settings or trusted sites | Security settings and trusted sites | Security settings and trusted sites |

**Related Information**

The following contains additional information that is relevant to this section.

- Group Policy Settings Reference in the Tools and Settings Collection.
- "What Is Core Group Policy?" in the Group Policy collection.

# How Internet Explorer Maintenance Extension Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Internet Explorer Maintenance Extension Works**

**In this section**


- Internet Explorer Maintenance Extension Architecture
- Internet Explorer Maintenance Extension Physical Structure
- Internet Explorer Maintenance Extension Processes and Interactions
- Related Information


The Internet Explorer Maintenance Extension of the Group Policy Object Editor enables administrators to define an Internet Explorer configuration as part of a Group Policy Object (GPO). The GPO is linked to Active Directory containers such as sites, domains, or organizational units (OUs), and enables management of the Internet Explorer configuration for multiple users on any computer joined to the domain that is capable of using Group Policy.

Deployment of Internet Explorer Maintenance Extension settings requires Group Policy in a Windows 2000 or Windows 2003 Active Directory environment, and Windows 2000 Professional or Windows XP clients.

**Internet Explorer Maintenance Extension Architecture**

The following figure illustrates the components important to the Internet Explorer Maintenance Extension.

**Internet Explorer Maintenance Extension Architecture**




These components are described in the following table. Components not seen in the figure, but important to the process, are also described.

**Internet Explorer Maintenance Extension Logical Architecture Components**

| Component | Description |
|---|---|
| Group Policy engine | This component is the framework that manages and implements the Group Policy settings and configurations, made by the admin, across all client-side extensions (CSE). Userenv.dll is the Group Policy engine module. |
| Internet Explorer Maintenance Client-Side extension (CSE) | The Internet Explorer Maintenance CSE is the component that is called by the Group Policy engine, and that applies the Internet Explorer Maintenance settings. The Internet Explorer Maintenance CSE writes the relevant information into the registry. |
| WinLogon | WinLogon is the service that contains the Group Policy engine. |

| | |
|---|---|
| Resultant Set of Policy (RSoP) snap-in | This component displays the results of Group Policy, including what Group Policy settings have been applied and when they were last applied. For more information about RSoP, see "What Is Resultant Set of Policy?." |
| Local GPO | Contains Group Policy settings for the local computer, including potential Internet Explorer Maintenance policies. |

The CSE registration information is written at setup to the HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Winlogon\ **GPExtensions** registry key. This registry key structure exists on both the target and domain controller systems.

**Internet Explorer Maintenance Extension Physical Structure**

Understanding where Internet Explorer Maintenance Extension policy settings are stored and how they are structured can help you troubleshoot problems you might encounter when you implement Internet Explorer Maintenance. Although GPOs can be linked to sites, domains, and OUs, they are stored only in the domain. See the "How Core Group Policy Works" topic in this collection for more information about how Group Policy stores its data.

The following table lists the setting types and the locations where Internet Explorer Maintenance Extension policy configuration files are stored on both the local computer and the domain.

**Physical Structure Components**

| Setting Type | Policy File Name |
|---|---|
| Browser Title | install.ins |
| Custom Bitmaps | Install.ins<br>\Branding\Logo\<<small logo file name>><br>\Branding\Logo\<<big logo file name>><br>\Branding\Animbmp (empty folder created) |
| Toolbar Customization | \install.ins<br>\Branding\Btoolbar\<<color logo file name>><br>\Branding\Btoolbar\<<grayscale logo file name>><br>\Branding\Toolbmp\<<toolbar bmp file name >> |
| Connection Settings | \install.ins<br>\Branding\cs\connect.set<br>\Branding\cs\cs.dat |
| Automatic Browser Configuration | \install.ins |
| Proxy Settings | \install.ins |
| User Agent String | \install.ins |
| Favorites and Links | \install.ins |
| Important URLs | \install.ins |
| Security Zones | \install.ins<br>\Branding\Zones\seczones.inf<br>\Branding\Zones\seczrsop.inf |

| | |
|---|---|
| Content Ratings | \install.ins<br>\Branding\Ratings\ratings.inf<br>\Branding\Ratings\ratrsop.inf |
| Authenticode Settings | \install.ins<br>\Branding\Authcode\authcode.inf |
| Programs | \install.ins<br>\Branding\Programs\programs.inf |
| Corporate Settings | \Branding\Adm\inetcorp.adm<br>\Branding\Adm\inetcorp.inf |
| Internet Settings | \Branding\Adm\inetset.adm<br>\Branding\Adm\inetset.inf |

Domain policy settings use the Fully Qualified Domain Name (FQDN) to reference GPOs. There are two main paths where the configuration files are stored:

- Domain policy files are stored in the folder `\\FQDN\Sysvol\FQDN\Policies\<GPOGUID>\User\Microsoft\IEAK`
- Local Machine policy files are stored in the folder `%windir%\System32\GroupPolicy\User\Microsoft\IEAK`

The following figure shows the files used by the Internet Explorer Maintenance Extension and where they are stored on both the domain controller and client computers.

**Internet Explorer Maintenance Extension File Storage**




**Internet Explorer Maintenance Extension Processes and Interactions**

When working with Internet Explorer Maintenance settings, you can use one of two interfaces. To configure Internet Explorer Maintenance Extension settings, use the Group Policy Object Editor. Use the Group Policy Management Console (GPMC) to view the Internet Explorer Maintenance Extension settings contained within a GPO.

**Using Group Policy Object Editor with Internet Explorer Maintenance**

To configure Internet Explorer Maintenance settings, an Administrator sets up Internet Explorer on a client computer with the settings to be included in the GPO. The Administrator then uses the Group Policy Object Editor to import the settings for the Security Zones, Content Ratings, Authenticode Settings, Programs, and Connection Settings, areas of the Internet Explorer Maintenance Extension and saves them as part of a GPO. The following figure shows the Internet Explorer Maintenance Extension interface used to import Connection Settings into a GPO.

**Importing Internet Explorer Settings**




**Configuring and Importing Internet Explorer Maintenance Settings to a GPO**

Administrators import settings from the appropriate settings dialog boxes in the Internet Explorer Maintenance extension of Group Policy Object Editor. The following things occur when the settings are imported:

- The IEAK Engine (ieakeng.dll) hosts the Internet Options Control Panel (inetcpl.cpl), which then reads the current settings from the registry.
- The Administrator then modifies the settings using the user interface of inetcpl.cpl.
- When the settings are saved, they are written back to the registry by inetcpl.cpl. Ieakeng.dll then imports them to the appropriate GPO files.

The following figure illustrates the process of importing Internet Explorer Maintenance Extension settings.

**Importing Internet Explorer Settings in XP**




**Note**

- If an administrator tries to view the settings in a GPO by clicking **Modify Settings**, the current settings from the registry, instead of the GPO, are immediately imported. Clicking **OK** then overwrites the settings stored in the GPO with the settings in effect on the client, deleting the settings previously contained in the GPO. In this event, the administrator

cannot view the GPO to find out what the previous settings were. It then becomes difficult to reconfigure the settings.

**Using the Group Policy Management Console to View a GPO**

To avoid overwriting the Internet Explorer Maintenance settings in a GPO, use GPMC to view the Internet Explorer Maintenance settings. GPMC runs on Windows XP Professional SP1 and Windows Server 2003 computers, and can manage Group Policy in either Windows 2000 or Windows Server 2003 domains. To see the settings contained in a GPO using GPMC, an Administrator views the **Settings** tab of the GPO as shown in the following figure.

**Viewing Internet Explorer Maintenance Settings in GPMC**




**Applying GPO Settings to a Client Computer**

The Internet Explorer Maintenance Extension uses the Internet Explorer Administration Kit (IEAK) infrastructure for both storage of settings and application to the client system.

When Group Policy is applied, Client-Side Extensions process the GPO. Internet Explorer Maintenance settings are handled by the Internet Explorer branding DLL (iedkcs32.dll). The Group Policy CSE invokes iedkcs32.dll, and two things happen:

1. The Group Policy CSE copies all IEAK settings files created using Internet Explorer Maintenance, listed in the previous **Physical Structure Components** table, to the following locations:

   Documents and Settings\<<username>>\Application Data\Microsoft\Internet Explorer\Custom Settings\Custom0\

   And

   Documents and Settings\<<username>>\Application Data\Microsoft\Network\Connections\pbk\Rasphone.pbk (for connection settings)

   Note that the policy's directory structure shown in the previous Physical Structure Components table is not replicated.

2. The Branding DLL then applies the settings from the downloaded files to the registry on the client system. There are four possible locations for the registry settings:

   - HKLM\Software\Policies (preferred)

- HKLM\Software\Microsoft\Windows\CurrentVersion\Policies
- HKCU\Software\Policies (preferred)
- HKCU\Software\Microsoft\Windows\CurrentVersion\Policies

These locations have security permissions that a standard user cannot modify in order to change applied policy settings. These keys are created the first time a GPO `configures` them.

Because the IEAK settings files are stored in the user's profile, the user has full-access permissions and can modify their contents. When the GPO is updated, the files are copied from the policy's directory structure back to the user profile and any changes the user might have made are overwritten. Although users can modify files in their own profile, attempting to execute the .inf file will give them an Access is denied error if they attempt to write settings to a key located in the secure registry branches previously specified.

If the user has a roaming profile, the IEAK settings files in the profile can be applied when roaming. This will happen if a roaming profile user logs on to the network from a computer that can't use Group Policy, or from one that isn't linked to a GPO containing Internet Explorer Maintenance Extension settings. If a user has manually changed the Internet Explorer Maintenance Extension settings located in their user profile, the user's settings will be applied to the computer. This has the potential of circumventing browser and security settings configured by the administrator. However, any settings appropriately locked-down in the registry (such as security and connection settings) will not have this problem.

**Related Information**
The following contains additional information that is relevant to this section.

- The Microsoft Internet Explorer Administration Kit (IEAK) Web site.
- How Core Group Policy Works in the Group Policy collection.

# Internet Explorer Maintenance Extension Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Explorer Maintenance Extension Tools and Settings**

**In this section**


- Internet Explorer Maintenance Extension Tools
- Internet Explorer Maintenance Registry Entries
- Internet Explorer Maintenance Group Policy Settings
- Related Information


This section summarizes the tools and settings associated with the Group Policy Internet Explorer Maintenance Extension.

**Internet Explorer Maintenance Extension Tools**

The following tools are associated with the Internet Explorer Maintenance Extension:

- The Microsoft Internet Explorer Administration Kit (IEAK)
- Group Policy Management MMC snap-in (GPMC)
- Resultant Set of Policy (RSoP)
- Group Policy Object Editor

This topic discusses the Microsoft Internet Explorer Administration Kit (IEAK) and the settings associated with the Internet Explorer Maintenance Extension.

GPMC, the Group Policy Object Editor, and RSoP are described in the "Group Policy Administrative Tools" topic in this collection. GPMC is documented in full at the Group Policy Management Console page, and in the "What Is Group Policy Management Console?" topic in this collection.

Additional command-line tools for managing Group Policy are described in the "Core Group Policy Tools and Settings" topic in this collection.

**Microsoft Internet Explorer Administration Kit**

The Microsoft Internet Explorer Administration Kit (IEAK) is a set of tools that provide administrators with powerful, easy-to-use options designed to save time and money in deploying and managing Web solutions for organizations and internet service providers (ISPs).

**Category**

To download a tool kit for customizing and distributing Microsoft Internet Explorer, see the Microsoft Internet Explorer Administration Kit (IEAK) Web site for Internet Explorer 6 Service Pack 1. To download the IEAK for previous versions of Internet Explorer, use the site's navigation to find the **Downloads** page.

The programs and tools in the IEAK 6 Service Pack 1 (SP1) simplify deploying and managing custom browser software packages. The features and resources include:

- **IEAK Toolkit.** The IEAK Toolkit contains tools, programs, and sample files, such as the animated bitmap tools, a sample sign-up and add-ons. The toolkit also includes the IExpress Wizard; administrators can use it to control setup options and to integrate command-line switches, which control how custom browser packages are installed.
- **IEAK Profile Manager.** After deploying Internet Explorer, the IEAK Profile Manager enables administrators to change browser settings and restrictions automatically.

- **Internet Explorer Customization Wizard.** The Internet Explorer Customization Wizard breaks browser customization into five stages. Step-by-step screens guide the administrator through the process of creating custom browser packages. When these packages are installed on clients' desktops, clients receive customized versions of Internet Explorer with the settings and options that the administrator selected.

- **IEAK Help.** The IEAK Help includes many conceptual and procedural topics viewable through **Index**, **Contents**, and **Search** tabs. Topics from the IEAK Help can also be printed.

**Version compatibility**
The IEAK is designed to manage Internet Explorer on any Windows operating system. Internet Explorer 6 is a core component of Windows XP and Windows Server 2003.

**IEdkcs32.dll**
**Category**
IEdkcs32.dll is an extension that provides settings for Group Policy Internet Explorer Maintenance.

**Version compatibility**
IEdkcs32.dll is compatible with Windows XP1 and Windows Server 2003.

**Internet Explorer Maintenance Extension Registry Entries**
The following registry entries are associated with the Internet Explorer Maintenance Extension.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. You should not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**ConnectionSettings**
**Registry path**
HKCU\Software\Microsoft\IEAK\BrandedFeatures\ConnectionSettings

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key stores the names of all network connections assigned to the computer using the Internet Explorer Maintenance Extension. The ConnectionSettings key contains a REG_MULTI_SZ entry named **NamesList**, which holds a list of the connection names.

**Button GUID**
**Registry path**
HKCU\Software\Microsoft\Internet Explorer\Extensions\<*Button GUID*>

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key stores configuration data for each button added to the Internet Explorer toolbar using Internet Explorer Maintenance Extension. The key name is variable, and is assigned by the Internet Explorer Maintenance Client-Side Extension (CSE) when the button is created. The **Button GUID** registry key contains the following Internet Explorer Maintenance Extension configured entries:

- **ButtonText** – The toolbar caption for the button.

- **Exec** – The program or action associated with the button.

- **HotIcon** – The toolbar button shown when the button is selected or the mouse pointer is over it.

- **Icon** – The standard toolbar button.

- **Default Visible** – Specifies whether or not the button is visible on the toolbar.

**Help_Menu_URLs**
**Registry path**
HKCU\Software\Microsoft\Internet Explorer\Help_Menu_URLs

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key stores the URL associated with the Online Support choice located on Internet Explorer's Help menu. The **Online_Support** entry contains the configured URL.

**Main**
**Registry path**
HKCU\Software\Microsoft\Internet Explorer\Main

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key stores configuration data for the home page, search bar, and the browser title. The **Main** registry key contains the following Internet Explorer Maintenance Extension configured entries:

- **Window Title** – This entry contains the test string displayed in the Internet Explorer title bar. Text entered in the Internet Explorer Maintenance Extension interface is appended to the string "Microsoft Internet Explorer provided by" and is then stored in this registry key entry.
- **Start Page** – This entry stores the URL for the assigned browser home page.
- **Search Bar** – This entry stores the URL displayed in the search toolbar pane when a user clicks the search button.
- **Use Custom Search URL** – This entry specifies whether or not the browser uses the custom search page located in the Search Bar entry.

**Toolbar**
**Registry path**
HKCU\Software\Microsoft\Internet Explorer\Toolbar

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key stores configuration data used to customize the Internet Explorer toolbar background, and the animated and static logos. The **Toolbar** registry key contains the following Internet Explorer Maintenance Extension configured entries:

- **BackBitmapIE5** – This entry specifies the bitmap used to customize the background for the Internet Explorer toolbar.
- **SmallBitmap** – This entry specifies the bitmap used for the small static logo.
- **BigBitmap** – This entry specifies the bitmap used for the large static logo.
- **SmBrandBitmap** – This entry specifies the bitmap used for the small animated logo.
- **BramdBitmap** – This entry specifies the bitmap used for the large animated logo.

**Default**
**Registry path**
HKCU\Software\Microsoft\RAS\AutoDial\Default

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key is used to specify the default internet connection. The **DefaultInternet** entry contains the name of the default

connection used to access the Internet.

**Internet Settings**
The following registry entries are located under HKCU\Software\Microsoft\Windows\CurrentVersion\Internet Settings.

**Automatic Configuration script**
**Registry path**
HKCU\Software\Microsoft\Windows\CurrentVersion\Internet Settings

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

The **AutoConfigURL** entry is used to specify where Internet Explorer checks for a configuration script. The configuration script is used to update the Internet Explorer setup.

**Proxy Server Location**
**Registry path**
HKCU\Software\Microsoft\Windows\CurrentVersion\Internet Settings

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key stores configuration data used to specify the proxy servers Internet Explorer uses. The following registry entries are configured for the Proxy Server Location by the Internet Explorer Maintenance Extension:

- **ProxyEnable** – This entry specifies whether or not the proxy server is currently enabled.
- **ProxyOverride** – This entry specifies the Web sites Internet Explorer bypasses the proxy server to access.
- **ProxyServer** – This entry specifies the URL and port number for the proxy server.

**Dial -up and VPN settings**
**Registry path**
HKCU\Software\Microsoft\Windows\CurrentVersion\Internet Settings\Connections

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key stores configuration data used to specify the configuration of dial-up and VPN connections. The following registry entries are configured for Dial-up and VPN Settings by the Internet Explorer Maintenance Extension:

- **<Connection Name> –** This entry is dependent on the name of the connection, and specifies the connection name displayed on the Internet Properties Connections tab.
- **DefaultConnectionSettings** – This entry specifies the configuration of the default connection used by Internet Explorer.
- **SavedLegacySettings** – This entry specifies the configuration used by network connections other than the default connection.

**Zone Map Domains**
**Registry path**
HKCU\Software\Microsoft\Windows\CurrentVersion\Internet Settings\ZoneMap\Domains\<*Domain or UNC*>

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

These registry keys specify which domain names have been configured within a security zone. Each domain or UNC that has an entry within a security zone has its own registry key in this path. The <*protocol*> entry takes its name from the corresponding protocol type (http, ftp, and so on). The contained entry specifies which zone the domain with protocol is part of.

**Protocol**
**Registry path**
HKCU\Software\Microsoft\Windows\CurrentVersion\Internet Settings\ZoneMap\Domains\<*Domain or UNC*>\<*protocol*>

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key specifies which protocols within a domain name have been placed in a security zone. There are no entries associated with this registry key.

**User Agent String**
**Registry path**
HKCU\Software\Microsoft\Windows\CurrentVersion\Internet Settings\User Agent\PostPlatform

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key contains the Internet Explorer Maintenance Extension settings specifying the user agent string. The name of the entry used to specify the string is the same as the string itself.

**Ratings**
The following registry entries are located under HKLM\Software\Microsoft\Windows\CurrentVersion\policies\Ratings.

**Password Hint and Encrypted Password**
**Registry path**
HKLM\Software\Microsoft\Windows\CurrentVersion\policies\Ratings

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key contains the settings for Internet Explorer ratings. The Hint entry contains the clue used to help someone remember the Ratings password, and the **Key** entry contains the encrypted Ratings password.

**Rating Policy**
**Registry path**
HKLM\Software\Microsoft\Windows\CurrentVersion\policies\Ratings\PICSRules\.Default\0\PRPolicy\<#>

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

The <#> registry key contains information about the Ratings of a domain or UNC. The registry key is a variable number generated in the order the **Ratings** entry was created. The **PRPolicyAttribute** entry specifies whether the entry is approved or restricted.

**Rating Policy Configuration**
**Registry path**
HKLM\Software\Microsoft\Windows\CurrentVersion\policies\Ratings\PICSRules\.Default\0\PRPolicy\<#>\PRPolicySub\0

**Version**
The registry settings in this section are for Windows 2000 Professional or Windows XP Professional clients, or for Windows Server 2003 without the Internet Explorer Enhanced Security Configuration enabled.

This registry key contains the Ratings configuration of a Domain or UNC. The following entries specify the Ratings configuration:

- **PRBUHost** – Specifies the rating's configured Domain or UNC name.
- **PRBUPort** – Specifies which port is used.
- **PRBUScheme** – Specifies which rating level the domain or UNC is configured to use.
- **PRBUUrl** – Specifies the URL for the specified Ratings Service's Web site.

**Internet Explorer Maintenance Extension Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with the Internet Explorer Maintenance Extension.

**Internet Explorer Maintenance Extension Configuration Settings**

| **Browser User Interface** | **Description** |
|---|---|
| Browser Title (Policy mode only) | This setting customizes the text that appears in the title bar of the Internet Explorer browser. The text is appended to the string "Microsoft Internet Explorer provided by." |
| Custom Logo (Policy mode only) | This setting replaces the static Internet Explorer logo in the upper-right corner of the Internet Explorer window with a custom static logo. It also replaces the animated Internet Explorer logo in the upper-right corner of the Internet Explorer window with a properly formatted custom animated logo. |
| Browser Toolbar Customizations | This setting customizes the background and buttons of the Internet Explorer toolbar. |
| **Connection** | **Description** |
| Connection Settings | These setting allow importing connection settings from the Internet Control Panel Connections tab of an already configured computer. They restrict how users can interact with connection settings through the System Policies and Restrictions page. It is not necessary to import current settings in order to set theses restrictions. They are also used to remove old dial-up connection settings from users' computers. |
| Automatic Browser Configuration | Automatic Configuration (auto-config) is used to update a user's computer after deployment by specifying a URL to an .ins file, an auto-proxy URL, or both. You can set the interval in minutes for when auto-config will happen. If the interval value is left blank, or at zero, auto-config will only happen when the browser has been started and navigates to a page. |
| Proxy Settings | These settings specify to which proxy servers, if any, user's connect. |
| User Agent String (Policy mode only) | The user agent string is what the browser sends to visited servers to identify itself. It is often used to keep Internet traffic statistics. You can enter custom text that will be appended to the default Internet Explorer string. The default string is different for each platform. |
| **URLs** | **Description** |
| Favorites and Links | You can customize Favorites and Links by specifying the URLs. You can also specify the order of each folder, add an icon for each Favorite and Link, and import an existing folder structure. |
| Important URLs | You can specify a custom home page, search bar URL, and online support page. |
| **Security** | **Description** |
| Security Zones and Content Ratings | You can customize the settings of each security zone as well as customizing the privacy settings. Content ratings let you prevent users from viewing sites with risky content. |
| Authenticode Settings | Authenticode enables you to designate software publishers as trustworthy. You can prevent users from adding new trusted publishers while using the browser. Enabling this lockdown does not prevent access to the Content control panel, but doing so is possible using the Administrative Templates Extension of Group Policy. |
| **Programs** | **Description** |
| Programs | You can import the current default programs settings. The programs selected specify which program Windows automatically uses for each Internet service. |
| **Advanced** | **Description** |

| | |
|---|---|
| Corporate Settings (Preference mode only) | These setting specify the location of a file containing settings used to make the browser work best for your corporation. These settings can be useful in reducing the cost of supporting applications in your corporation. Advanced features, as well as details of other features, can all be configured through these options. |
| Internet Settings (Preference mode only) | These settings specify the location of a file used to preset defaults for Internet Explorer settings which are not set through the IEAK Wizard. Most of these settings apply to defaults in the Internet Control Panel. |

- For more information about Group Policy settings, see the Group Policy Settings Reference in the Tools and Settings Collection.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Tools and Settings Collection in the Tools and Settings Collection.
- The Microsoft Internet Explorer Administration Kit (IEAK) Web site.
- The Core Group Policy Tools and Settings topic in this collection.
- The What Is Group Policy Management Console? topic in this collection.
- The Group Policy Administrative Tools topic in this collection.
- The Group Policy Management Console page.
- The Tools and Settings Collection.

# Remote Installation Services Extension Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft Windows 2000 Server and Windows Server 2003 use Group Policy settings to provide directory-based desktop configuration management. These Group Policy settings are stored in a Group Policy object (GPO). You can create and manage GPOs by using the Group Policy Object Editor.

General classes of Group Policy settings are implemented as extensions to the operating system, and can exist as two types of extensions:

- **Server-side extension (SSE).** A Microsoft Management Console (MMC) snap-in that is used for administration and configuration.
- **Client-side extension (CSE).** An extension that runs on the client computer and interprets and applies the MMC-type settings to the client computer. The client computer is also known as the target computer.

The GPO settings are configured through the server-side extension, and are applied to individual computers and users by the Group Policy engine and the client-side extension (CSE). The Group Policy engine is the framework used to manage client side-extensions.

Remote Installation Services (RIS) is different from most other Group Policy extensions. RIS has no CSE or dynamic-link library (DLL) providing a CSE because RIS is used by the client before the client has an operating system installed.

**In this subject**


- What Is Remote Installation Services Extension?
- How Remote Installation Services Extension Works
- Remote Installation Services Extension Tools and Settings

# What Is Remote Installation Services Extension?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Remote Installation Services Extension?**

**In this section**


- Remote Installation Services Extension Overview
- Remote Installation Services Extension Core Scenarios
- Remote Installation Services Extension Dependencies
- Related Information


Remote Installation Services (RIS), originally introduced as an optional component with Windows 2000 Server, is an optional component of Windows Server 2003 family. RIS enables administrators or users to remotely install a copy of the Windows Server 2003, Windows XP Professional, or Windows 2000 operating system. Administrators use Remote Installation Services Extension to specify which options are presented to users by the Client Installation Wizard.

**Remote Installation Services Extension Overview**

Administrators who install multiple client operating systems on “bare-metal” computers—computers that do not have an operating system—one at a time from the installation CD, can spend a lot of time at each computer. This takes administrators away from their other responsibilities. Customizing and configuring each newly installed operating system to meet organization needs is susceptible to errors and takes even more time. Having an inexperienced end user install the operating system from the CD can be frustrating for the end user, result in installation errors, and increase the number of technical support service calls.

With Windows Server 2003, Microsoft provides two features to help administrators deploy Windows clients and servers over a network. These two features are RIS and Automated Deployment Services (ADS). Administrators can use ADS for deploying server farms. For more information about ADS, see AutomatedDeploymentServicesSupportResources under ManagementServices on Windows Server 2003 Support Center.

You can use RIS to:

Provide an operating system to users on demand.

Provide an operating system image that includes specific settings and applications.

Create automated installation images of products in the Windows Server 2003 family, Windows XP, and Windows 2000.

You can also combine RIS with IntelliMirror features such as user documents and settings, Software Installation, and Group Policy. This combination can improve the efficiency of computer management in your organization, and reduce the number of technical support service calls.

The following figure shows the Remote Installation Services node of the Group Policy Object Editor, which is used to configure RIS policy settings. This figure shows the RIS server-side extension user interface, provided by rigpsnap.dll. There is no RIS client-side extension.

**Group Policy Object Editor**




The settings you make in the RIS SSE determine which of the four options the user sees in the Client Installation Wizard of the target computer, as seen in the following figure.

**Client Installation Wizard**




This figure shows each of the options that are presented to the user on the target computer during a remote installation of an operating system. Because all four settings are enabled in the SSE, all four settings appear in the wizard. Maintenance and Troubleshooting in the wizard corresponds to Tools in the SSE.

**Remote Installation Services Extension Core Scenarios**

The core scenario for Remote Installation Services is operating system installation with no involvement of the administrator during client installation. By using the Remote Installation Services Extension node in the Group Policy Object Editor to configure RIS policy settings, the administrator can predetermine the options presented to users during installation. This saves the administrator time, and ensures consistent deployment of client computers across the organization.

**Remote Installation Services Extension Dependencies**
You can edit Remote Installation Group Policy settings on any computer that has the Group Policy Object Editor with the Remote Installation Services node (rigpsnap.dll). However, note the following about Remote Installation Services:

The Microsoft version of Active Directory is required.

The client computer using RIS to install an operating system must have a network card that:

- Supports Pre-boot Execution Environment (PXE), or
- Is supported by the RIS remote boot floppy disk.

Windows XP Professional must have Windows Server 2003 Administration Tools Pack installed to show RIS settings in Group Policy Object Editor.

RIS Help is not available by default in the Group Policy Management Console (GPMC) running on Windows XP Professional. Install Windows Help from the Windows Server 2003 CD onto the computer running Windows XP Professional.

RIS is not included in the Windows Server 2003, Web Edition operating system.

RIS installation image files must be located on a Windows NT File System (NTFS) partition not containing the system or boot files.

**Related Information**
The following contains additional information that is relevant to this section.

- Automated Deployment Services Support Resources under Management Services on Windows Server 2003 Support Center.

# How Remote Installation Services Extension Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Remote Installation Services Extension Works**

**In this section**


- Remote Installation Services Extension Architecture
- Remote Installation Services Extension Protocols
- Remote Installation Services Extension Physical Structure
- Remote Installation Services Extension Processes and Interactions
- Network Ports Used by Remote Installation Services Extension


Administrators configure Remote Installation Services Extension using the Remote Installation Services (RIS) server-side extension (SSE), accessed from the Group Policy Object Editor. There is no RIS client-side extension (CSE). This section describes how RIS Policy is configured in the SSE, and how those settings affect Remote Installation Services.

**Remote Installation Services Extension Architecture**

The following figure shows the high-level architecture of Remote Installation Services.

**Remote Installation Services Extension Architecture**




In the figure, RIS SSE is the Remote Installation Services (RIS) server-side extension (SSE). This is the Remote Installation Services node in the Group Policy Object Editor for the GPO.

Components significant to RIS are summarized in the following table.

**Remote Installation Services Extension Architecture Components**

| Component | Description |
|---|---|

| | |
|---|---|
| Group Policy Object Editor | The Group Policy Object Editor is an MMC server-side extension packaged as the dynamic-link library (DLL) named Gpedit.dll. Each of the nodes in the Group Policy Object Editor, including the Remote Installation Services node, is itself an extension. |
| Remote Installation Services (RIS) server-side extension (SSE) | The RIS SSE is a Group Policy Object Editor extension, packaged as the DLL named rigpsnap.dll. Administrators use the RIS extension to configure the Remote operating system Installation options presented to target users. |
| Rigpsnap.dll | Rigpsnap.dll is the RIS SSE component. |
| Lightweight Directory Access Protocol (LDAP) | LDAP is a communication protocol used to query Active Directory directory services. |
| Server Message Block (SMB) | SMB is a communication protocol used for file access between the RIS extension and Active Directory. |
| Group Policy Container | A Group Policy Container is a location in Active Directory where GPOs and their properties are stored. |
| Group Policy Template (GPT) | The GPT is the part of the GPO that consists of the domain controller file system, and in which most Group Policy settings are stored. |
| Sysvol | The Sysvol folder contains the subset of folders shared on each domain controller. These folders store *file-system* domain information, as compared to *registry* domain information. |
| Pre-Boot Execution Environment (PXE) | PXE is a combination of standardized protocols and services that enables remote startup of computers over a network. Remote Installation Services uses PXE. |
| Boot Information Negotiation Layer (BINL) service | BINL is a service that runs on the RIS Server. BINL handles client requests, authenticates clients with Active Directory, and confirms that correct policy and settings are applied to the client during remote installation. |
| Trivial File Transfer Protocol (TFTP) service (daemon) | TFTP is a RIS Server service that downloads the OSChooser, Client Installation Wizard (CIW) files, and other boot files during the RIS process. |
| Client Installation Wizard (CIW) | CIW is the graphical user interface that guides the user of a client computer through the remote installation of an operating system. CIW consists of screens and a program called OSChooser that displays the screens. |
| Single Instance Store (SIS) service | SIS is a pair of components used to save disk and main memory file cache space when two or more files share content. Other attributes of the files might be different, such as path, owner, access control lists, and so on. Writes to one file don't necessarily affect the other files. SIS saves space by providing links that are semantically identical to independent files. |
| SIS filter | The SIS filter is one of the SIS components. The SIS filter is a kernel-level file system filter driver that implements links for identical files. |
| SIS Groveler service | The SIS groveler service is one of the SIS components that detects identical files and reports them to the SIS filter for link creation. |
| DHCP | Dynamic Host Configuration Protocol is the communication protocol used between the client and the RIS server (and the DHCP server). RIS uses DHCP, which follows the PXE architecture, to bootstrap the client computer. |
| Bootp | Bootp is the RIS Server. |

**Remote Installation Services Extension Protocols**

Protocols used by Remote Installation Services are shown in the previous RIS Architecture figure, and are described in the following table.

**Remote Installation Services Protocols**

| Protocols | Description |
|---|---|
| Trivial File Transfer Protocol (TFTP) | TFTP is the service that hosts file download from the RIS Server to the client during remote operating system installation. |
| Lightweight Directory Access Protocol (LDAP) | LDAP is a communication protocol used by the RIS SSE on the Administrator computer to query Active Directory directory services. |
| Server Message Block (SMB) | SMB is a protocol used by the Administrative computer to access the Sysvol on the Domain Controller. SMB is the primary file and print sharing protocol. |
| NTLM v1 | Windows 2000 uses NTLM v1 for encryption. |
| NTLM v2 | Windows Server 2003 uses NTLM v2 for encryption. |
| DHCP | The client uses Dynamic Host Configuration Protocol (DHCP) to send out a DHCP discover asking for an IP address and a PXE boot server. The RIS Server and DHCP server respond using DHCP. |

**Remote Installation Services Extension Physical Structure**
The following figure shows the RIS file structure.

**RIS Policy Files**




The following table summarizes the files used by Remote Installation Services and the RIS extension.

**Remote Installation Services Policy Physical Structure Components**

| Component | Description |
|---|---|

| | |
|---|---|
| Setup Installation Files (.sif) | The RIS Server uses Setup Installation Files to communicate the installation options to the RIS client. These files have the .sif extension. |
| .osc files | The RISetup.exe tool, used to setup the RIS Server, installs a default set of Client Installation Wizard files. These files have an .osc extension, and are stored in the \OSChooser\English folder on the RIS Server. These files are text files formatted in OSCML, and can be modified. |
| Client Installation Wizard (CIW) files | CIW files are .osc files that guide the user through a remote operating system installation on the client. |
| Winnt.sif | Winnt.sif is a Setup Installation File downloaded from the RIS Server to the client and used to detect the client hardware abstraction layer (HAL) type. |
| Ntldr | Ntldr is the executable that starts the CIW. Ntldr is on the client, and is the same as the oschoice.exe file on the RIS Server, copied and saved as Ntldr on the client. |
| Startrom.com file | Startrom.com is a file located in the \\RIS_server_name\REMINST\oschooser\i386 folder on the RIS server. Startrom.com is the first file downloaded by Ntldr to start the client boot process. |
| Welcome.osc file | The Welcome.osc file provides the first Client Installation Wizard screen presented to the client during the Client Installation Wizard. |
| Choice.osc file | The Choice.osc file stores all possible remote operating system installation choices that can be presented to the client user. |
| Oscfilter.ini | The Oscfilter.ini file stores settings made in the RIS SSE, and it determines which of the possible operating system installation options in Choice.osc, are presented to the client. |
| Oschoice.osc | The Oschoice.osc file stores the remote operating system installation list (if more than one) presented to the client user as determined by the oscfilter.ini file and choice.osc file combination. |

**Remote Installation Services Extension Processes and Interactions**
A high-level overview of the remote installation of an operating system follows.

**How a Remote Operating System Installation Works**
The RIS Server and the RIS Client both must be connected to the wired network. Wireless connection to the network is not supported.

When the client user presses the F12 key, the client's BIOS initiates a PXE boot by sending a DHCP discover and implementing PXE option tags. The BinL service on the RIS Server, which monitors port 4011, receives the DHCP discover from the client. The server downloads the startrom.com file image to the client using TFTP. The server loads the network driver on the client using the Uniform Network Device Interface (UNDI). The client transmits the client's GUID/MAC to the server, which records the GUID/MAC. The server sends the Ntldr (oschoice.exe saved as Ntldr on the client) to the client using TFTP, and the client starts Ntldr/OSChoice.exe.

The server downloads the winnt.sif file to the client to detect the hardware abstraction layer (HAL) type. The client and server negotiate an OSC connection on port 401. Ntldr presents Client Installation Wizard (CIW) screens for user entry. Settings in Active Directory, the Registry, and files on the RIS Server determine the content of these screens.

One of the CIW screens is the Installation Choices screen. The RIS server queries the Active Directory Sysvol to determine the remote operating system installation options to display in this screen. The Active Directory Sysvol stores operating system choice settings in the oscfilter.ini file. You can determine the settings in the oscfilter.ini file using the Remote Installation Services node in the GPO editor. The Remote Installation Services node provides an interface for you to control the remote operating system installation options displayed on this one screen in the CIW. You can manually edit the oscfilter.ini file in Notepad to obtain the same result.

### Note

- You can also control the client installation process by manually editing the .osc and .sif files on the RIS server. For more information, see the Remote Installation Services help, which is available from the Windows Server 2003 family Group Policy Object Editor.

Depending on which of the remote operating system installation choices the client user makes, OSChoice receives the name of the next application to start, and sets the fully qualified path to the RIS server. The name of the loader is set, and a soft TFTP restart occurs. OSChoice passes the TFTP restart block, which is a local packet of information, to the setup loader. The restart block is an in-memory location for X86 computers and is passed using launch arguments in EFI architecture.

**RIS SSE Impact on Remote O/S Installation**

The settings you make in the RIS SSE are saved in the oscfilter.ini file, which is located in the \WINDOWS\SYSVOL\sysvol\<domain>\Policies<gpo_guid>\User\Microsoft\RemoteInstall folder on the domain controller. The oscfilter.ini file has the following structure:

```
[Choice]
OSAUTO=0
CUSTOM=1
RESTART=0
TOOLS=0
```

Your three options for each of the four settings in the RIS SSE, and their impact on the settings in the oscfilter.ini file are:

- **Not Configured.** This removes the choice and value from the oscfilter.ini file.
- **Enabled.** This sets the value for the choice to 1.
- **Disabled.** This sets the value for the choice to 0.

The choice.osc file on the RIS server is not edited by your settings in the RIS SSE; it remains unchanged. Only the oscfilter.ini file is edited. The Choice.osc file is formatted in OSChooser Markup Language (OSCML) which is HTML 2.0. You can manually edit this file to customize the login options presented to the client user. The following code shows the default Choice.osc file.

```
<OSCML>
<META KEY=F3 ACTION="REBOOT">
<META KEY=F1 href="choichlp">
<META SERVER ACTION="DNRESET">
<META SERVER ACTION="FILTER CHOICE">
<TITLE>  Client Installation Wizard                          Main Menu</TITLE>
<FOOTER>  [ENTER] continue             [F1] help          [F3] restart
computer</FOOTER>
<BODY left=5 right=75>
<br>
<br>
Use the arrow keys to select one of the following options:<br>
<P left=8>
<FORM>
<SELECT SIZE=10>
<OPTION VALUE="OSAUTO" TIP="This is the easiest way to install an operating system on
your
computer. Most installation options are already configured by your network
administrator.">
Automatic Setup
<OPTION VALUE="CUSTOM" TIP="With this option, you can define a unique name for this
computer
and specify where the computer account will be created within the directory service.
Select this option if you are setting up this computer for someone else within your
company.">
Custom Setup
<OPTION VALUE="RESTART" TIP="A previous remote installation attempt has been detected
on this
computer. Select this option to restart a previously started installation.">
Restart a Previous Setup Attempt
<OPTION VALUE="TOOLS" TIP="This option gives you access to tools for keeping your
computer
up-to-date and for troubleshooting problems.">
Maintenance and Troubleshooting
</SELECT>
</FORM>
</P>
```

<BR>
<BOLD>Description:</BOLD>  
<TIPAREA>
</BODY>
</OSCML>

**Network Ports Used by Remote Installation Services Extension**

The following table lists port assignments for Remote Installation Services extension.

**Port Assignments for Remote Installation Services**

| Service Name | UDP | TCP |
|---|---|---|
| DHCP | n/a | 67, 68 |
| RPC | Dynamically assigned | Dynamically assigned |
| LDAP | n/a | 398 |
| SMB | n/a | 445 |
| COM | Dynamically assigned | Dynamically assigned |
| DCOM | Dynamically assigned | Dynamically assigned |

# Remote Installation Services Extension Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services Extension Tools and Settings**

**In this section**


- Remote Installation Services Extension Tools
- Remote Installation Services Extension Registry Entries
- Remote Installation Services Extension Group Policy Settings
- Related Information


This section includes Tools and Settings associated with Remote Installation Services Extension.

**Remote Installation Services Extension Tools**

The following tools are associated with Remote Installation Services.

**DSA.msc: Active Directory Users and Computers**

**Category**

DSA.msc is the Active Directory Users and Computers snap-in used by Administrators to manage Active Directory in Windows 2000 and Windows Server 2003 domains.

**Version compatibility**

DSA.msc runs on Windows 2000 Server, Windows XP SP1 with the Windows Server 2003 Administration Tools Pack installed, and Windows Server 2003.

DSA.msc ships with Windows 2000 and Windows Server 2003.

**GPMC.msc: Group Policy Management Console**

**Category**

GPMC.msc is a snap-in used by Administrators to manage Group Policy in Windows 2000 and Windows Server 2003 domains. Remote Installation Services policy can be configured from GPMC.

**Version compatibility**

GPMC.msc runs on Windows XP SP1, and Windows Server 2003.

For GPMC download information, see the Group Policy Management Console link on the Web Resources page.

**Remote Installation Services Extension Registry Entries**

The following registry entries are associated with Remote Installation Services Extension.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. You should not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**BINLSVC**

**Registry path**

HKEY_LOCAL_MACHINE \System\CurrentControlSet\Services\

**Version**

Windows Server 2003.

This key and its sub-keys store Remote Installation Service settings for the BinL service used by Remote Installation Services (RIS).

For more information about this registry entry, see the Registry Reference for Windows Server 2003.

**Groveler**
**Registry path**
HKEY_LOCAL_MACHINE \SYSTEM\CurrentControlSet\Service\

**Version**
Windows Server 2003.

This key and its sub-keys store Remote Installation Service settings for the Groveler service used by RIS.

For more information about this registry entry, see the Registry Reference for Windows Server 2003.

**SIS**
**Registry path**
HKEY_LOCAL_MACHINE \SYSTEM\CurrentControlSet\Service\

Version

**Windows Server 2003.**
This key and its sub-keys store Remote Installation Service settings for the SIS service used by RIS.

For more information about this registry entry, see the Registry Reference for Windows Server 2003.

**TFTPD**
**Registry path**
HKEY_LOCAL_MACHINE \SYSTEM\CurrentControlSet\Service\

**Version**
Windows Server 2003.

This key and its sub-keys store Remote Installation Service settings for the TFTPD service used by RIS.

For more information about this registry entry, see the Registry Reference for Windows Server 2003.

**Remote Installation Services Extension Group Policy Settings**
You can use the following four general settings to configure Remote Installation Services options:

- Group Policy Settings
    - Remote Installation Services node
    - Administrative Templates
- RIS Server Authorization (DHCP MMC)
- Authorized RIS Server Properties (ADUC MMC)
- Discretionary Access Control Lists (DACLs)

**Group Policy Settings**
This section summarizes the two types of Group Policy settings.

**Remote Installation Services Node Settings**
The following tables list and describe Group Policy settings that are associated with Remote Installation Services Extension.

### Remote Installation Services Node Policy Settings

| Group Policy Setting | Description |
|---|---|

| | |
|---|---|
| Automatic Setup | Automatic Setup is one of the four installation options. When Automatic Setup is the only installation choice, users on the client are not offered installation choices. Automatic Setup settings are:<br>• Not Configured (Default)<br>• Enabled<br>• Disabled |
| Custom Setup | Custom Setup is one of the four installation options, with settings of:<br>• Not Configured (Default)<br>• Enabled<br>• Disabled |
| Restart Setup | Restart Setup is one of the four installation options, with settings of:<br>• Not Configured (Default)<br>• Enabled<br>• Disabled |
| Tools (Maintenance and Troubleshooting) | Tools is one of the four installation options, with settings of:<br>• Not Configured (Default)<br>• Enabled<br>• Disabled |

**Administrative Templates Settings**

RIS template settings are found at the following node in Group Policy Object Editor:

- **User Configuration\Administrative Templates\Windows Components\MMC\Restricted/Permitted snap-ins\Group Policy\Group Policy snap-in extensions**

The setting is **Remote Installation Services** and the setting options are:

- Not Configured (Default)
- Enabled
- Disabled

**RIS Server Authorization (DHCP MMC) Settings**

RIS settings are also available from the Dynamic Host Configuration Protocol MMC, and the setting is **RIS Authorization**. You use this option to specify which RIS Servers are allowed to run on your network. RIS Servers are disabled by default, and you must specifically authorize a RIS Server.

**Authorized RIS Server Properties (ADUC MMC) Settings**

Once RIS servers are authorized in the DHCP MMC, the RIS server has settings accessible from the Active Directory Users and Computers (ADUC) MMC. These settings are located on tabs in the **Properties** dialog box for the RIS Server. The RIS server also has a New Answer File wizard.

The following table summarizes the Remote Install Tab settings.

**Remote Install Tab Settings**

| Setting | Description |
|---|---|
| Respond to client computers requesting service | You use this option to pre-stage a client. You pre-stage a client by adding the computer name to Active Directory, and then adding that name to this setting list. This setting is disabled by default. |
| Do not respond to unknown client computers | You use this option to prevent/allow clients that are not pre-staged to install a remote operating system. This setting is disabled by default. |

The next table shows the settings for the **New Clients Tab**.

**New Clients Tab**

| Setting | Description |
|---|---|
| Generate client computer names using | Options for this setting are:<br>• First initial, Last name<br>• Last name, First initial<br>• First name, Last initial<br>• Last initial, First name<br>• User name (default)<br>• NP plus MAC<br>• Custom |
| Where client account is created | Option button choices are:<br>• Default directory service location (Computers)<br>• Same location as that of the user setting up the client computer<br>• The following directory service location (Browse) |

The next table shows the settings for the **Images Tab**.

**Images Tab**

| Setting | Description |
|---|---|

| | |
|---|---|
| Images | This option lists installation images installed on the RIS Server, with no images listed as the default. The following actions are available for this option:<br>• Add (starts Add Wizard)<br>• Remove<br>• Properties<br>• Refresh |

The next table shows the settings for the **Tools Tab**.

**Tools Tab**

| Setting | Description |
|---|---|
| Tools | This option provides maintenance and troubleshooting tools that are installed on the RIS Server, with no tools listed as the default. The following actions are available for this option:<br>• Remove<br>• Properties<br>• Refresh |

The next table shows the settings for the **Security Tab**.

**Security Tab**

| Setting | Description |
|---|---|
| Group or user names | You use this option to add or remove users and groups to the list. |
| Permissions for *<user or group selected>* | You use this option to allow or deny the permission for the selected user or group, for the following permissions:<br>• Full Control<br>• Read<br>• Write<br>• Create All Child Objects<br>• Delete All Child Objects<br>• Special Permissions<br>• Advanced button |

The next table shows the settings for the **New Answer File Wizard**.

**New Answer File Wizard**

| Setting | Description |
|---|---|

| | |
|---|---|
| Associate a new answer file to an existing image (option button) | When you select this option and click the Next button, your options are:<br>• Windows image sample files<br>• Another remote installation server<br>• An alternate location |
| Add a new installation image (option button) | When you select this option and click the Next button, the Add Installation Image Wizard displays, which you can use to set:<br>• Path to image<br>• Folder name<br>• Friendly Description<br>• Help Text<br>• CIW screen options |

**Discretionary Access Control List Setting**

The following table shows the Discretionary Access Control List setting for the RIS server.

**Discretionary Access Control List (DACL) Setting**

| Setting | Description |
|---|---|
| OS image availability to users in CIW | Available to all users by default. |

For more information about Group Policy settings, see the Group Policy Settings Reference for Windows Server 2003.

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Group Policy Settings Reference for Windows Server 2003.
- MSDN.
- For GPMC download information, see the Group Policy Management Console link on the Web Resources page.
- Tools and Settings Collection

# Group Policy Administrative Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Administrative Tools**

There are three primary tools used to administer Group Policy: Microsoft Group Policy Management Console (GPMC), Group Policy Object Editor, and Resultant Set of Policy (RSoP) snap-in. Each of these tools is a Microsoft Management Console (MMC). Administrators use GPMC for the bulk of Group Policy management tasks. Group Policy Object Editor is for editing Group Policy objects. Although administrators can still use the RSoP snap-in for reporting and planning the effects of Group Policy, much of its functionality has been subsumed into GPMC.

**Group Policy Administrative Tools Architecture**

The diagram below illustrates the relationship between the domain controller, client, and the three primary tools used to administer Group Policy.

**Group Policy Administrative Tools Architectural Diagram**




**Components of Group Policy Administrative Tools Architectural Diagram**

| Component | Description |
|---|---|
| Group Policy Object Editor | The Group Policy Object Editor is used to edit GPOs. It was previously known as the Group Policy snap-in, Group Policy Editor, or Gpedit. A notable feature of the Group Policy Object Editor is its extensibility. Developers can extend the server-side snap-ins that ship with Group Policy Object Editor or they can develop completely new extensions for implementing Group Policy.<br><br>The Group Policy Object Editor is capable of read and write access to Active Directory, Sysvol, and the Local GPO. |
| Server-Side Snap-Ins | The nodes of the Group Policy Object Editor are MMC snap-ins. These snap-ins include Administrative Templates, Scripts, Security Settings, Software Installation, Folder Redirection, Remote Installation Services, and Internet Explorer Maintenance. Snap-ins may in turn be extended. For example, the Security Settings snap-in includes several extension snap-ins. Developers can also create their own MMC extension snap-ins to the Group Policy Object Editor to provide additional policy settings. Extensions are capable of read and write access to the Local GPO. |

| | |
|---|---|
| Group Policy Management Console (GPMC) | GPMC makes Group Policy much easier to manage by providing a view of GPOs, sites, domains, and organizational units (OU) across an enterprise. GPMC can be used to manage either Windows Server 2003 or Windows 2000 domains.<br><br>GPMC simplifies the management of Group Policy by providing a single place for managing core aspects of Group Policy, such as scoping, delegating, filtering, and manipulating inheritance of GPOs. You can also back up GPOs to the file system as well as restore GPOs from backups. GPMC includes features that enable an administrator to predict how GPOs are going to affect the network as well as to determine how GPOs have actually changed settings on any particular computer or user.<br><br>GPMC is capable of read and write access to the Sysvol using the SMB protocol. It is also capable of read and write access to Active Directory via the LDAP protocol. In addition, GPMC is capable of read access to the event log and RSoP infrastructure. |
| Resultant Set of Policy (RSoP) snap-in | The Resultant Set of Policy snap-in is an MMC used to determine which policy settings are in effect for a given User or Computer, or to predict the effect of applied policy.<br><br>The snap-in itself is contained within the same binary as the Group Policy Object Editor snap-in (gpedit.dll). The user interface is mostly a read-only view of the same information available in the Group Policy Object Editor. However there is one important difference: While the Group Policy Object Editor can show only a single GPO setting at a time, the RSoP snap-in shows the cumulative effect of many GPOs.<br><br>For RsoP functionality it is recommended to use GPMC, which includes its own integrated RsoP features.<br><br>The RSoP snap-in is capable of read access to the Active Directory, Sysvol, Event Log, RSoP infrastructure, and Local GPO. Although the RSoP snap-in is capable of read only access to the Active Directory and Sysvol, most of the work of predicting or reporting Group Policy is done using RPC/COM communication with the RSoP provider, either on the client or the domain controller. |
| Domain Controller (Server) | In an Active Directory forest, the domain controller is a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. GPOs are stored in two parts of domain controllers: The Active Directory database and the Sysvol share. |
| Client | In an Active Directory forest, settings from GPOs are applied to clients. GPMC and the RSoP snap-in query the client to determine how policy has been applied to a particular user or computer. |
| Active Directory | Active Directory, the Windows-based directory service, stores information about objects on a network and makes this information available to users and network administrators. Administrators link GPOs to Active Directory containers such as sites, domains, and OUs that include user and computer objects. In this way, policy settings can be targeted to users and computers throughout the organization. |
| Sysvol | Sysvol is a shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain. The Sysvol contains the largest part of a GPO: the Group Policy template (GPT), which includes Administrative Template-based policy settings, security settings, and script files. File Replication Service (FRS) replicates this information throughout the network. |
| RsoP infrastructure | All Group Policy processing information is collected and stored in a Common Information Model Object Management (CIMOM) database on the local computer. This information, such as the list of GPOs that have been processed, as well as content and logging of processing details for each GPO, can then be accessed by tools using Windows Management Instrumentation (WMI).<br><br>With Group Policy Results in GPMC, or logging mode for the RSoP snap-in, the RSoP service is used to query the CIMOM database on the target computer; it receives information about the policies that were applied and displays the resulting information in GPMC or the RSoP snap-in.<br><br>With Group Policy Modeling in GPMC, or planning mode for the RSoP snap-in, the RSoP service simulates the application of policy using the Group Policy Directory Access Service (GPDAS) on a Domain Controller. GPDAS simulates the application of GPOs and passes them to virtual client-side extensions on the Domain Controller. The results of this simulation are stored in a local CIMOM database on the domain controller before the information is passed back and displayed in either GPMC or the RSoP snap-in. |

**Group Policy Administrative Tools Components**

The following is a list of Group Policy administrative tools.

**Group Policy Management Console**
In the past, administrators have been required to use several Microsoft tools to manage Group Policy, such as the Active Directory Users and Computers, Active Directory Sites and Services, Group Policy Object Editor and Resultant Set of Policy snap-ins. With the introduction of Group Policy Management Console (GPMC), most administrative tasks have been integrated into a single, unified console that also offers several new capabilities.

**Group Policy Object Editor**
Group Policy Object Editor is an MMC snap-in used to edit the policy settings in Group Policy objects (GPOs). Like all MMC snap-ins, its functionality can be customized or extended by means of MMC snap-in extensions. When you use Group Policy Object Editor, various extensions are included by default, including Administrative Templates, Scripts, Security Settings, Software Installation, Folder Redirection, Remote Installation Services, and Internet Explorer Maintenance. Snap-ins may in turn be extended. For example, the Security Settings snap-in includes several extension snap-ins. Developers can also create their own MMC extension snap-ins to the Group Policy Object Editor to provide additional policy settings.

**Resultant Set of Policy Snap-in**
Resultant Set of Policy (RSoP) snap-in is an MMC used to predict the effect of GPOs on the network as a whole, or to determine the effect Group Policy has had on a specific user or computer. Although administrators can use the RSoP snap-in, much of its functionality has been subsumed into GPMC.

**Group Policy Administrative Tools Deployment Scenarios**
The following is a list of common scenarios supported by each of the Group Policy Administrative Tools.

**Group Policy Management Console**
An administrator uses GPMC to manage Group Policy in an Active Directory environment. Although the majority of computers on the network might be running Windows 2000 server or Windows 2000 professional, the administrator might download and install GPMC on a machine running Windows XP SP1 or later. GPMC requires Windows XP or later.

GPMC offers the administrator a persistent view of the Group Policy environment on the network, including icons that represent GPOs, GPO links, sites, domains, and organizational units (OU) in the selected forest. With GPMC, an administrator can do any of the administrative tasks previously only available from the Group Policy tab of the Active Directory administrative tools.

GPMC can also be used to generate RSoP data that either predicts the cumulative effect of GPOs on the network, or reports the cumulative effect of GPOs on a particular user or computer. In addition, the administrator can use GPMC to perform GPO operations never possible before, like backing up and restoring a GPO, copying a GPO, or even migrating a GPO to another forest. Reading or generating HTML or XML reports of GPO settings is also possible.

**Group Policy Object Editor**
An administrator uses Group Policy Object Editor to manipulate settings in a GPO. Typically an administrator accesses Group Policy Object Editor by electing to edit a GPO from within GPMC. The Group Policy Object Editor opens, allowing the administrator to change settings for that GPO. If the administrator had two GPOs linked to an OU and wanted to manage settings in both, he would have to open them one at a time. This is because Group Policy Object Editor can only display settings for GPOs one at a time. To see how settings for multiple GPOs might affect an OU, the administrator would use a different tool — either GPMC or the RSoP snap-in.

**Resultant Set of Policy snap-in**
An administrator uses Resultant Set of Policy snap-in to predict the cumulative effect of GPOs on the network, or report the cumulative effect of GPOs on a particular user or computer. With the advent of GPMC, the administrator can generate reports with most of the information formerly only available through the RSoP snap-in. The RSoP snap-in does provide some information that GPMC does not. For example, when multiple GPOs attempt to set the same Group Policy setting differently, both GPMC and RSoP snap-in report how the setting is ultimately set and which GPO is responsible for the setting, but only RSoP snap-in can report all the GPO(s) that attempted and failed to manipulate the setting.

# Group Policy Management Console Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Management Console Technical Reference**
In the past, administrators have been required to use several Microsoft tools to manage Group Policy, such as the Active Directory Users and Computers, Active Directory Sites and Services, Group Policy Object Editor and Resultant Set of Policy snap-ins. With the introduction of Group Policy Management Console (GPMC), most administrative tasks have been integrated into a single, unified console that also offers several new capabilities.

**In this subject**


- What Is Group Policy Management Console?
- How Group Policy Management Console Works
- Group Policy Management Console Tools and Settings

# What Is Group Policy Management Console?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Group Policy Management Console?**

**In this section**


- Group Policy Management Console Core Scenarios
- Group Policy Management Console Dependencies
- Related Topics


The Group Policy Management Console (GPMC) is a new and comprehensive administrative tool for Group Policy management.

Prior to GPMC, administrators used property pages in various Active Directory administrative tools to manage Group Policy. For example, an administrator who wanted to implement policy for users might open the Active Directory Users and Computers snap-in, find an appropriate Organizational Unit (OU) and open its property page to access the Group Policy tab. On the Group Policy tab, the administrator might do any of a dozen or so administrative tasks, like creating Group Policy object links or manipulating their order to achieve the desired results. Whatever the tasks, when the administrator leaves the Group Policy tab, access to a visual representation of Group Policy ends and a view that focuses on Active Directory's user and computer objects appears.

GPMC integrates the existing Group Policy functionality of the property pages on the Active Directory administrative tools into a single, unified console dedicated to Group Policy management tasks; GPMC also expands management capabilities with new features. Administrators still use Active Directory administrative tools to manage Active Directory, but GPMC replaces the Group Policy management functionality of those tools with its own.

**Group Policy Management Console Core Scenarios**

There are many core scenarios for GPMC. Administrators use GPMC to perform all Group Policy management tasks, with the exception of configuring individual policy settings in Group Policy Objects themselves, which is done with Group Policy Object Editor. The scenarios below describe how an administrator uses GPMC to manage Group Policy.

**Creating and Editing GPOs**

Administrators use GPMC to create a GPO with no initial settings. An administrator can also create a GPO and linking it to an Active Directory container at the same time. To configure individual settings within a GPO, an administrator edits the GPO from within GPMC and Group Policy Object Editor appears with the GPO loaded. An administrator can use either GPMC or Group Policy Object Editor to disable or enable computer, user, or both computer and user nodes within a GPO.

**Scoping GPOs**

An administrator can use GPMC to link GPOs to sites, domains, or OUs in the Active Directory. Administrators must link GPOs to apply settings to users and computers in Active Directory Containers. Linking GPOs is the primary mechanism by which administrators apply Group Policy settings.

In addition to linking, an administrator can manipulate permissions on GPOs to manage how Group Policy applies. Prior to GPMC, an administrator would have to manually manipulate access control entries (ACE) to modify the scope of the GPO. For example, the administrator might remove **Read** and **Apply Group Policy** from the Authenticated Users group for GPO1. This effectively disables GPO1, since users in the Authenticated Users group require both **Read** and **Apply Group Policy** permissions to process Group Policy. To apply the settings in GPO1 to select network users or computers, the administrator would add a new security principal (typically a security group containing the target users or computers) to the ACL on the GPO and set **Read** and **Apply Group Policy** permissions. This is known as security filtering.

With GPMC, security filtering has been simplified. The administrator adds the security principal to the GPO, and GPMC automatically sets the **Read** and **Apply Group Policy** permissions.

Administrators can also use GPMC to link WMI Filters to GPOs. WMI Filters allow an administrator to dynamically determine the scope of GPOs based on attributes (available through WMI) of the target computer. A WMI filter consists of one or more queries that are evaluated to be either true or false against the WMI repository of the target computer. The WMI filter is a separate object from the GPO in the directory.

**Manipulating Inheritance**
Group Policy can be applied to users and computers at the site, domain, or OU level. GPOs from parent containers are inherited by default. When multiple GPOs apply to these users and computers, the settings in the GPOs are aggregated. For most policy settings, the final value of a given policy setting is set only by the highest precedent GPO that contains that setting. (However, the final value for a few settings will actually be the combination of values across GPOs.)

Group Policy determines precedence of GPOs by the order of processing for the GPOs. GPOs processed last have highest precedence. GPOs follow the SDOU rule for processing; site first, then domain, and then followed by OU, including nested OUs. A nested OU is one that has another OU as its parent. In the case of nested OUs, GPOs associated with parent OUs are processed prior to GPOs associated with child OUs. In this processing order, sites are applied first but have the least precedence. OUs are processed last and have the highest precedence.

When a container has multiple GPO links, administrators can use GPMC to manipulate the link order for every container. GPMC assigns each link a Link Order number; the GPO link with Link Order of 1 has highest precedence on that container.

Administrators can use GPMC to Block Inheritance. This is the ability to prevent an OU or domain from inheriting GPOs from any of its parent container. Note that Enforced GPO links (see below) will always be inherited.

Administrators can use GPMC to set GPO links to Enforced (previously known as No Override). This is the ability to specify that a GPO should take precedence over any GPOs that are linked to child containers. Enforcing a GPO link works by moving that GPO to the end of the processing order.

An administrator can also use GPMC to enable or disable a GPO link. If an administrator enables a GPO link, Group Policy processes the linked GPO. If the link is not enabled, Group Policy does not process the linked GPO.

**GPO Operations**
GPO operations refer to the ability to backup (export), restore, import, and copy GPOs. Backing up a GPO consists of making a copy of GPO data to the file system. Note that the Backup function also serves as the export function for GPOs. Backed up GPOs can be used either in conjunction with Restore or Import operations.

Restoring a GPO takes an existing GPO backup and re-instantiates it back in the domain. The purpose of a restore is to reset a specific GPO back to the identical state it was in when it was backed up. This restoration does not include GPO links. This is because the links are a property of the container the GPO is linked to, not the GPO itself. Since a restore operation is specific to a particular GPO, it is based on the GUID and domain of the GPO. Therefore, a restore operation cannot be used to transfer GPOs across domains.

Importing a GPO allows you to transfer settings from a backed up GPO to an existing GPO. You can perform this operation within the same domain, across domains, or across forests. This allows for many interesting capabilities such as staging of a test GPO environment in a lab before importing into a production environment.

Restoring and Importing a GPO will remove any existing settings already in the target GPO. Only the settings in the backup will be in the GPO when these operations are complete.

Copying a GPO is similar to an export/import operation only the GPO is not saved to a file system location first. In addition, a copy operation creates a new GPO as part of the operation, whereas an import uses an existing GPO as its destination.

**Reporting of GPO Settings**
GPMC can display a report of the defined settings in a given GPO. This report can be generated by any user with read access to the GPO. Without GPMC, users that did not have write access to a GPO could not read and view the settings in that GPO. This is because the Group Policy Object Editor requires the user to have read and write permissions to the GPO to open it. Some examples of users that might need to read and view but not edit a GPO include security audit teams that need to read but not edit GPO settings, helpdesk personnel that are troubleshooting a Group Policy issue, and OU administrators that may need to read and view the settings from inherited GPOs. With GPMC these users now have read access to the settings.

The HTML reports also make it easy for the administrator to view settings that are contained in a GPO at a glance. Alternatively, administrators can expand and contract individual sections within the report by clicking the heading for each section.

GPMC also solves some common reporting requirements including the ability to document all the settings in a GPO to a file for printing or viewing. Using a context menu, users can either print the reports, or save them to a file in either HTML or XML format.

**Search for GPO**
GPMC provides extensive capabilities to search for GPOs within a domain or across all domains in a forest. This search feature allows an administrator to search for GPOs based on the following criteria:

- Display name of the GPO.

- Whether or not a specific domain has containers that link to the GPO.
- The permissions set on the GPO.
- The WMI filter that is linked to the GPO.
- The type of policy settings that have been set in the User Configuration or Computer Configuration in the GPO, such as folder redirection or security settings. Note that you cannot search based on the individual settings configured in a GPO.
- GUID of the GPO.

**Group Policy Modeling**
Windows Server 2003 has a powerful new Group Policy management feature that allows the user to simulate a policy deployment that would be applied to users and computers before actually applying the policies. This feature, known in Windows Server 2003 as Resultant Set of Policy (RSoP) Planning Mode, is integrated into GPMC as Group Policy Modeling. This feature requires a domain controller that is running Windows Server 2003 in the forest, because the simulation is performed by a service that is only present on Windows Server 2003 domain controllers. However, with this feature, you can simulate the resultant set of policy for any computer in the forest, including those running Windows 2000.

**Group Policy Results**
This feature allows administrators to determine the resultant set of policy that was applied to a given computer and (optionally) the user that logged on to that computer. The data that is presented is similar to Group Policy Modeling data, however, unlike Group Policy Modeling, this data is not a simulation. It is the actual resultant set of policy data obtained from the target computer. Unlike Group Policy Modeling, the data from Group Policy Results is obtained from the client, and is not simulated on the DC. The client must be running Windows XP, Windows Server 2003 or later. It is not possible to get Group Policy Results data for a Windows 2000 computer. (However, with Group Policy Modeling, you can simulate the RSoP data).

**GPMC Scripting**
The GPMC user interface is based on a set of COM interfaces that accomplish all of the operations performed by GPMC. These interfaces are available to Windows scripting technologies like JScript and VBScript, as well as programming languages such as Visual Basic and VC++. An administrator can use these interfaces to automate many Group Policy management tasks.

These interfaces are discussed in detail in the GPMC software development kit (SDK) located in the **%programfiles%\gpmc\scripts\gpmc.chm** Help file on systems where GPMC has been installed. The contents of the GPMC SDK are also available in the Platform SDK.

For more information about GPMC interfaces, see MSDN.

**Group Policy Management Console Dependencies**
Group Policy Management Console requires the MMC infrastructure to function. In addition, GPMC has its own set of system installation requirements, as well as requirements for certain GPMC features to function.

**GPMC System Installation Requirements**
Although GPMC can manage both Windows 2000 and Windows Server 2003 domains with Active Directory, the tool itself must be installed on a computer running Windows Server 2003 or Windows XP Professional (with Windows XP Service Pack 1 (or later) and the Microsoft .NET Framework). Note that when installing GPMC on Windows XP Professional SP1, a hotfix is required. This hotfix (Q326469) is included with GPMC. GPMC Setup prompts you to install Windows XP QFE Q326469 if it is not already present.

**GPMC Feature Requirements**
GPMC exposes features that are available in the underlying operating system. Because new features have been added to Group Policy since Windows 2000, certain features will only be available in GPMC depending on the operating system that has been deployed on the domain controllers and clients. This section describes these dependencies. In general, there are four key issues that determine whether a feature is available in GPMC:

- Windows Server 2003 Active Directory schema must be available to delegate Group Policy Modeling or Group Policy Results.
- Windows Server 2003 domain controller must be available to run Group Policy Modeling.
- Windows Server 2003 domain configuration (ADPrep /DomainPrep) must be available to use WMI Filters.

- Clients must be running Windows XP or Windows Server 2003 in order to generate Group Policy Results data.

GPMC dependencies upon Windows and Active Directory platform appear below:

**GPMC Dependencies on Windows and Active Directory**

| Dependency | Feature | Reason |
|---|---|---|
| Windows Server 2003 Active Directory Schema | Delegation of Group Policy Modeling and Group Policy Results | The Generate Resultant Set of Policy (Logging) and Generate Resultant Set of Policy (Planning) permissions needed for this operation are only available with the Windows Server 2003 Active Directory schema. |
| Windows Server 2003 Domain Controller in the forest | Group Policy Modeling | The simulation is performed by the Resultant Set of Policy Service which is only available on domain controllers running Windows Server 2003. |
| Windows Server 2003 domain configuration (DomainPrep) | WMI Filters | ADPREP /DomainPrep configures the domain for Windows 2003 Active Directory including configuration for WMI Filters. |
| Clients must be running Windows XP or Windows Server 2003 | Group Policy Results | Clients must be instrumented to log Group Policy Results data when policy is processed. This capability is only available on the listed systems. |

There is no dependency from the Group Policy perspective on whether a domain is in native mode or mixed mode.

**Related Information**

The following resource contains additional information that is relevant to this section:

- GPMC Interfaces on MSDN.

# How Group Policy Management Console Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Group Policy Management Console Works**

Administrators use Group Policy Management Console (GPMC) to manage Group Policy. An administrator must have an Active Directory environment to use GPMC to manage Group Policy. In an ideal environment, the clients are running Windows XP while the servers are running Windows 2003. However, GPMC can also be used to manage Windows 2000 domains.

**In this section**


- Group Policy Management Console Architecture
- Group Policy Management Console Interfaces
- Group Policy Management Console Processes and Interactions
- Related Information


**Group Policy Management Console Architecture**

GPMC is one of three administrative tools used to manage Group Policy. The following diagram shows all three of those tools, as well as the domain controller and a client. In addition, the diagram describes the different communication protocols being used by each tool (LDAP, SMB, RPC/COM, DNS), the interactions between the tool, the domain controller and the client, and whether those interactions are READ (lines with a single arrow head) or READ/WRITE (lines with double arrow heads).

**Group Policy Management Console Architectural Diagram**




**Components of Group Policy Management Console Architectural Diagram**

| Component | Description |
|---|---|

| | |
|---|---|
| Group Policy Management Console (GPMC) | GPMC makes Group Policy much easier to manage by providing a view of GPOs, sites, domains, and OUs across an enterprise. GPMC can be used to manage either Windows Server 2003 or Windows 2000 domains.<br><br>GPMC simplifies the management of Group Policy by providing a single place for managing core aspects of Group Policy, such as scoping, delegating, filtering, and manipulating inheritance of GPOs. You can also back up GPOs to the file system as well as restore GPOs from backups. GPMC includes features that enable an administrator to predict how GPOs are going to affect the network as well as to determine how GPOs have actually changed settings on any particular computer or user.<br><br>GPMC is capable of read and write access to the Sysvol using the SMB protocol. It is also capable of read and write access to Active Directory via the LDAP protocol. In addition, GPMC is capable of read access to the event log and RSoP infrastructure. |
| Server (Domain Controller) | In an Active Directory forest, the domain controller is a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. GPOs are stored in two parts of domain controllers: The Active Directory database and the Sysvol.<br><br>In addition to storing the GPOs used in normal policy processing, GPMC and the RSoP snap-in use the domain controller and the RSoP infrastructure to simulate the application of policy across the network. |
| Active Directory | Active Directory, the Windows-based directory service, stores information about objects on a network and makes this information available to users and network administrators. Administrators link GPOs to Active Directory containers such as sites, domains, and OUs that include user and computer objects. In this way, policy settings can be targeted to users and computers throughout the organization. Prior to GPMC, administrators used property pages in various Active Directory administrative tools to manage Group Policy. |
| Sysvol | The Sysvol is a shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain. The Sysvol contains the largest part of a GPO: the Group Policy template (GPT), which includes Administrative Template-based policy settings, security settings, script files, and information regarding applications that are available for software installation. File Replication Service (FRS) replicates this information throughout the network. |
| DNS/WINS | Domain Name Service (DNS) and Windows Naming Service (WINS) are network name services. Network name services are used to translate IP addresses to domain names. GPMC can use either DNS or WINS to locate domain controllers. |
| LDAP Protocol | Lightweight Directory Access Protocol (LDAP) is the protocol used by the Active Directory directory service. GPMC uses LDAP to access the directory store on the domain controller. The client computer on which Group Policy is processed also uses LDAP to read the directory store on the domain controller. |
| SMB Protocol | The Server Message Block (SMB) protocol is the primary protocol used for file and print sharing. SMB can also be used for named pipes and mail slots. GPMC uses SMB to access the Sysvol as well as back up and retrieve files to a remote file system. The client also uses SMB to read the sysvol on the domain controller. |
| DNS protocol /WINS protocol | Domain Name Service (DNS) protocol and Windows Internet Name Service (WINS) protocol are used for locating computers on a network by using hostnames instead of IP addresses. GPMC can use either to locate domain controllers. |
| RPC/COM | RPC (Remote Procedure Call), DCOM (Distributed Component Object Model) and COM (Component Object Model) enable data exchange between different processes. The different process can be on the same machine, on the local area network, or across the Internet.<br><br>The various binary files that make up the GPMC components primarily use COM calls to communicate. The GPMC COM interfaces are also all scriptable so that an administrator can automate Group Policy Management.<br><br>DCOM and RPC are used by GPMC for communication with the RSoP provider on the client or domain controller for Group Policy Results and Group Policy Modeling reports. |

| | |
|---|---|
| WMI | WMI is a management infrastructure that supports monitoring and controlling of system resources through a common set of interfaces and provides a logically organized, consistent model of Windows operation, configuration, and status.<br><br>WMI makes data about a target computer available for administrative use. Such data can include hardware and software inventory, settings, and configuration information. For example, WMI exposes hardware configuration data such as CPU, memory, disk space, and manufacturer, as well as software configuration data from the registry, drivers, file system, Active Directory, the Windows Installer service, networking configuration, and application data. WMI Filtering in Windows Server 2003 allows you to create queries based on this data. These queries (also called WMI filters) determine which users and computers receive the policy configured in the GPO. |
| RsoP infrastructure | Group Policy processing information is collected and stored in a Common Information Model Object Management (CIMOM) database on the local computer. This information, such as the list of GPOs, as well as the content and logging of processing details for each GPO, can then be accessed by tools using WMI.<br><br>With Group Policy Results in GPMC, or logging mode for the RSoP snap-in, the RSoP service is used to query the CIMOM database on the target computer; it receives information about the policies and displays it in GPMC or the RSoP snap-in.<br><br>With Group Policy Modeling in GPMC, or planning mode for the RSoP snap-in, the RSoP service simulates the application of policy using the Group Policy Directory Access Service (GPDAS) on a Domain Controller. GPDAS simulates the application of GPOs and passes them to virtual client-side extensions on the Domain Controller. The results of this simulation are stored to a local CIMOM database on the domain controller before the information is passed back and displayed in either GPMC or the RSoP snap-in. |
| Group Policy Engine | The Group Policy Engine is a framework that handles common functionality across client-side extensions. GPMC does not communicate directly with the Group Policy Engine. For more information about the Group Policy Engine, see Core Group Policy Technical Reference. |
| Client-Side Extensions | Client-side extensions (CSEs) consist of one or more dynamic-link libraries (DLLs) that are responsible for implementing Group Policy at the client computer. CSEs typically correspond to snap-ins: Administrative Templates, Scripts, Security Settings, Software Installation, Folder Redirection, Remote Installation Services, and Internet Explorer Maintenance. GPMC does not communicate directly with the Client-side extensions. For more information about client-side extensions, see Group Policy Components. |
| File System | GPMC can write to the file system of the local machine or any remote machine. GPMC writes to the file system for GPO operations, such as backups or copy, and for saving HTML/XML reports. |
| Event Log | The Event log is a service that records events in the various logs on the computer. GPMC is capable of read and write access to the Event Log on client computers and domain controllers. |
| Local Group Policy object | The local Group Policy object (local GPO) is stored on each individual computer, in the hidden **%systemroot%\System32\GroupPolicy** directory. Each computer running Windows 2000, Windows XP Professional, Windows XP 64-Bit Edition, or Windows Server 2003 has exactly one local GPO, regardless of whether the computers are part of an Active Directory environment.<br><br>GPMC does not offer access to the local GPO. For more information about the local Group Policy Object, see Core Group Policy Technical Reference. |

**Group Policy Management Console Interfaces**

GPMC includes a set of programmable interfaces designed for use by administrators writing scripts, as well as programmers using languages such as Visual Basic or C++.

The following is a comprehensive list of all the Group Policy Management Console (GPMC) interfaces.

**Comprehensive List of GPMC Interfaces**

| Interface | Description |
|---|---|
| **IGPM** | Accesses GPMC interfaces to retrieve GPMC information and to create objects. Manages and queries Group Policy objects (GPOs) and GPO links. Queries scope of management (SOM) objects. |

| | |
|---|---|
| **IGPMAsyncCancel** | Cancels an asynchronous GPMC operation, such as a backup, restore, import, copy, or report generation. The GPMC function called by the client will asynchronously return a pointer to the **IGPMAsyncCancel** interface. **IGPMAsyncCancel** is not available through scripting. |
| **IGPMAsyncProgress** | Enables client notification about the progress of an operation. Implemented by the client and passed as a parameter to GPMC methods that can run asynchronously. **IGPMAsyncProgress** is not available through scripting. |
| **IGPMBackup** | Retrieves properties of **GPMBackup** objects. Deletes **GPMBackup** objects |
| **IGPMBackupCollection** | Accesses a collection of GPO backups. |
| **IGPMBackupDir** | Queries GPO backups and GPO backup collections. |
| **IGPMClientSideExtension** | Queries client-side extension properties such as ID and display name. |
| **IGPMConstants** | Retrieves the values of GPMC constants. |
| **IGPMCSECollection** | Accesses a collection of client-side extension objects. |
| **IGPMDomain** | Queries SOM objects. Creates, restores, and queries GPOs. Creates and queries WMI filters. WMI filter queries use WMI Query Language (WQL). |
| **IGPMGPO** | Manages an individual GPO. Deletes, copies, imports, or backs up a GPO. Enables or disables user and computer configuration options. Sets security information. Retrieves multiple GPO properties for a given GPO and generates reports of the settings in the GPO. |
| **IGPMGPOCollection** | Accesses a collection of GPOs. |
| **IGPMGPOLink** | Manages a GPO link from a SOM. Sets and retrieves properties of GPO links, such as the link order or whether a link is enabled or enforced. |
| **IGPMGPOLinksCollection** | Accesses a collection of GPO links. |
| **IGPMPermission** | Retrieves permission properties. |
| **IGPMResult** | Retrieves status message information while performing GPMC operations on GPOs, such as restore, import, backup and copy. |
| **IGPMRSOP** | Queries Resultant Set of Policy (RSoP) in logging or planning mode. Generates RSoP reports. |
| **IGPMSearchCriteria** | Defines the criteria for search operations. |
| **IGPMSecurityInfo** | Defines the set of permissions that exist on a SOM, GPO or WMI filter. Removes or adds permissions. |
| **IGPMSitesContainer** | Queries SOM objects for particular sites in a forest. |
| **IGPMSOM** | Creates and retrieves GPO links for a SOM. Sets and retrieves security attributes and properties for a SOM. |
| **IGPMSOMCollection** | Accesses a collection of SOMs. |
| **IGPMStatusMessage** | Retrieves properties of status messages of GPO operations. |
| **IGPMStatusMsgCollection** | Accesses a collection of status messages. |
| **IGPMTrustee** | Retrieves information about a trustee that is a user, group, or computer in the domain. |
| **IGPMWMIFilter** | Sets or retrieves security attributes and properties of a WMI filter. |
| **IGPMWMIFilterCollection** | Accesses a collection of WMI filters. |
| **IGPMMapEntry** | Retrieves map entries. |

| | |
|---|---|
| **IGPMMapEntryCollection** | Accesses a collection of map entries. |
| **IGPMigrationTable** | Accesses a migration table. |

For more information about GPMC interfaces, see MSDN.

**Group Policy Management Console Processes and Interactions**

This section defines how GPO operations are handled by GPMC, as well as important differences between how GPMC and Group Policy Object Editor handle Administrative template files.

**GPO Operations in Group Policy Management Console**

With Group Policy Management Console (GPMC) administrators can back up, restore, import, or copy Group Policy objects (GPOs). When copying or importing GPOs across domains, an administrator can use migration tables to facilitate the copy or import operation.

**Backups**

Backing up a Group Policy object (GPO) copies the data in the GPO to the file system. The backup function also serves as the export capability for GPOs. A GPO backup can be used to restore the GPO to the backed-up state, or to import the settings in the backup to another GPO. Because each backup is identified by a unique backup ID, multiple backups of the same or different GPO can be stored in the same file system location. The collection of backups in a given file system location can be managed using the GPMC or through the scriptable interfaces. When you open **Manage Backups** from the **Group Policy Objects** node in GPMC, the view is automatically filtered to show only backups of GPOs from that domain. When opened from the **Domains** node, all backups are displayed, regardless of which domain they are from.

**Information saved in a backup**

Backing up a GPO saves all information that is stored inside the GPO to the file system. This includes the following information:

- GPO globally unique identifier (GUID) and domain.
- GPO settings.
- Discretionary access control list (DACL) on the GPO.
- WMI filter link, if there is one, but not the filter itself.
- Links to IP Security Policies, if any.
- XML report of the GPO settings, which can be viewed as HTML from within GPMC.
- Date and time stamp of when the backup was taken.
- User-supplied description of the backup.

**Information not saved in a backup**

Backing up a GPO only saves data that is stored inside the GPO. Data that is stored outside the GPO is not available when the backup is restored to the original GPO or imported into a new one. This data that becomes unavailable includes the following information:

- Links to a site, domain, or organizational unit.
- WMI filter.
- IP Security policy.

**Restore**

Restoring a Group Policy object (GPO) re-creates the GPO from the data in the backup. A restore operation can be used in both of the following cases: the GPO was backed up but has since been deleted, or the GPO is live and you want to roll back to a known previous state. The end effect is the same in either case, except as noted below in some cases for GPOs that contain software installation settings.

GPMC identifies the GPO by its domain and globally unique identifier (GUID). The purpose of a restore operation is to return the GPO to its original state, so the restore operation retains the original GUID even if it is recreating a deleted GPO. This is a key difference between the restore operation and the import or copy operations. You cannot use restore to transfer GPOs to different domains or forests. That capability is provided by import and copy.

**Information replaced in a restore**
A restore operation replaces the following components of a GPO:

- GPO settings.
- The Discretionary Access Control List (DACL) on the GPO.
- WMI filter links (but not the filters themselves).

**Information not replaced in a restore**
The restore operation does not restore objects that are not part of the GPO. This includes the following:

- Links to a site, domain, or organizational unit. Links are an attribute of the site, domain, or organizational unit, not the GPO. Any existing links in these containers will continue to be used, for example, when restoring an existing GPO to a previous state. However, if you have deleted a GPO and all links to the GPO, you must add these links back after restoring the GPO.
- WMI filters. A GPO only contains a link to the WMI filter. If the WMI filter does not exist at the time of restore, the link will be dropped, otherwise the link will be restored.
- IPSec Policies. A GPO only contains a link to the IPSec policy. If the IPSec policy does not exist at the time of restore, the link will be dropped, otherwise the link will be restored.

**Permissions required to restore a GPO from backup**
The permissions necessary to perform a restore of a GPO vary, depending on whether you are restoring an existing GPO or if you are restoring a GPO that has been deleted since it was backed up. Version numbers for the GPO are handled differently as well. The following table summarizes the situation for existing and deleted GPOs:

**Permissions Required to Restore GPOs from Backup**

| GPO state | Permission needed to restore GPO from backup | GPO version number |
|---|---|---|
| Existing GPO | The user must have Edit settings, delete, and modify permissions on the GPO, as well as read access to the file system location where the backup is stored. This does NOT require GPO creation rights. | Incremented by 1, which will trigger client refreshes of settings from the GPO. |
| Deleted GPO | The user must have the right to create GPOs in the domain, as well as read access to the file system location where the backup is stored. | Retained unchanged from the backed-up GPO. |

**Restoring GPOs with software installation settings**
When restoring a deleted GPO that contains Software Installation settings, some side effects are possible depending on the circumstances under which the GPO is restored.

When restoring a GPO that contains software installation settings, it is possible that:

- Cross-GPO upgrade relationships that upgrade applications in the GPO being restored, if any, are not preserved after restore. A cross-GPO upgrade is one in which the administrator has specified that an application should upgrade another application, and the two applications are not in the same GPO. Note that upgrade relationships are preserved when applications — in the GPO being restored — upgrade applications in other GPOs.
- If the client computer has not yet detected that the GPO has been deleted (either because the user has not logged on again or rebooted since the deletion of the GPO), and the application was deployed with the option to **Uninstall this application when it falls out of scope of management** then the next time the user logs on:
    - Published applications that the user has previously installed will be removed.

- Assigned applications will be uninstalled before re-installation.

This issue can be avoided if all of the following conditions are met:

- You perform the restore on a Windows Server 2003 domain controller instead of a Windows 2000 domain controller.
- The user performing the restore has permissions to re-animate tombstone objects in the domain.
- The time elapsed between deletion and restoration of the GPO does not exceed the tombstone lifetime specified in Active Directory.

Tombstone re-animation is a new feature of Windows Server 2003 Active Directory. By default, only Domain Admins and Enterprise Admins have this permission but you can delegate this right to additional users at the domain level using the ACL editor.

As a general rule, if you deploy software using Group Policy, it is recommended that you perform the restoration of GPOs that contain application deployments using a domain controller running Windows Server 2003 and that you grant the tombstone re-animation right to the users who will be performing restoration of those GPOs.

Finally, when restoring a GPO that contains software installation settings, if you are using categories to classify applications, the application in the restored GPO will appear in its original category only if the category exists at the time of restoration. Note that the category definition is not part of the GPO.

**Import**
The Import operation transfers settings into an existing GPO in Active Directory using a backed-up GPO in the file system location as its source. Import operations can be used to transfer settings from one GPO to another GPO within the same domain, to a GPO in another domain in the same forest, or to a GPO in a domain in a different forest. The import operation always places the backed-up settings into an existing GPO. It erases any pre-existing settings in the destination GPO. Import does not require trust between the source domain and destination domain. Therefore it is useful for transferring settings across forests and domains that don't have trust. Importing settings into a GPO does not affect its discretionary access control list (DACL), links on sites domains or organizational units to that GPO, or a link to a WMI filter.

When using import to transfer GPO settings to a GPO in a different domain or different forest, you might want to use a migration table in conjunction with the import operation. A migration table allows you to facilitate the transfer of references to security groups, users, computers, and UNC paths in the source GPO to new values in the destination GPO.

**Copy**
A copy operation allows you to transfer settings from an existing Group Policy object (GPO) in Active Directory directly into a new GPO. The new GPO created during the copy operation is given a new globally unique identifier (GUID) and is unlinked. You can use a copy operation to transfer settings to a new GPO in the same domain, another domain in the same forest, or a domain in another forest. Because a copy operation uses an existing GPO in Active Directory as its source, trust is required between the source and destination domains. Copy operations are suited for moving Group Policy between production environments, and for migrating Group Policy that has been tested in a test domain or forest to a production environment, as long as there is trust between the source and destination domains.

Copying is similar to backing up followed by importing, but there is no intermediate file system step, and a new GPO is created as part of the copy operation. The import operation, in contrast with the copy operation, does not require trust.

**Specifying the discretionary access control list (DACL) on the new GPO**
You have two options for specifying the DACL to use on the new GPO:

- Use the default permissions that are used when creating new GPOs.
- Preserve the DACL from the source GPO. For this option, you can specify a migration table, used to facilitate the transfer of references to security groups, users, computers, and UNC paths in the source GPO to new values in the destination GPO. If you specify a migration table for the copy operation, and you choose the option to preserve the DACL from the source GPO, the migration table will apply to any security principals in the DACL.

**Copying within a domain compared with copying to another domain**

Copying a GPO to another domain is slightly different from copying it to the same domain.

When copying a GPO within the same domain, you have a simple choice of two options, just described, for choosing the DACL. However, for copy operations to another domain, GPMC presents a wizard to facilitate the operation. The wizard guides you through the following choices:

- Choice of DACL for the new GPO, as described earlier.
- Specification of migration table, if applicable. A migration table allows you to facilitate the transfer of references to security groups, users, computers, and UNC paths in the source GPO to new values in the destination GPO.

When copying a GPO within the same domain, any link to a WMI filter is preserved. However, when copying a GPO to a new domain, the link is dropped because WMI filters can only be linked to GPOs within the same domain.

**Migration Tables**
When you copy or import a Group Policy object (GPO) from one domain to another, you can use a migration table to tell Group Policy Management Console (GPMC) how domain-specific data should be treated.

The key challenge when migrating GPOs from one domain to another is that some information in the GPO is actually specific to the domain where the GPO is defined. When transferring the GPO to a new domain, it may not always be desirable, or even possible, to use the exact same settings. Items that can be specific to a domain include references to security principals such as users, groups, and computers as well as UNC paths. The solution is to modify these references in the GPO that are domain-specific, during the import or copy operation, so that the settings in the destination GPO are written with the appropriate information for the destination domain. GPMC supports this capability using migration tables.

A migration table is a file that maps references to users, groups, computers, and UNC paths in the source GPO to new values in the destination GPO. A migration table consists of one or more mapping entries. Each mapping entry consists of a type, source reference, and destination reference. If you specify a migration table when performing an import or copy, each reference to the source entry will be replaced with the destination entry when writing the settings into the destination GPO.

The migration table will apply to any references in the settings within a GPO, whether you are performing an import or copy operation. In addition, during a copy operation, if you choose the option to preserve the discretionary access control list (DACL) on the GPO, the migration table will also apply to both the DACL on the GPO and the DACLs on any software installation settings in the GPO.

Migration tables store the mapping information as XML, and have their own file name extension, .migtable. You can create migration tables using the Migration Table Editor (MTE). The MTE is a convenient tool for viewing and editing migration tables without having to work in, or be familiar with, XML. The MTE is associated with the .migtable extension so that when you double-click a migration table, it opens in the MTE. The MTE is installed with GPMC. You can also create and edit migration tables using any XML editor or using the GPMC scripting interfaces.

**Settings impacted by migration tables**
There are two types of settings that may not transfer well across domains:

- Users, groups, and computers (security principals) that are referenced in the following places: the settings of the GPO, in the ACL for the GPO itself, and the ACL for any software installation settings in the GPO. Security principals do not transfer well for several reasons. For instance, domain local groups are invalid in external domains but other groups are valid if there is trust in place. Even with a trust between domains, it may not always be appropriate to use the exact same group in the new domain. If there is no trust, then none of the security principals in the source domain will be available in the destination domain.
- UNC paths. When you are migrating GPOs across domains that do not have trust, such as from test to production environments, users in the destination domain may not have access to paths in the source domain.

The following items can contain security principals and can be modified using a migration table:

- Security policy settings of the following types:
    - User rights assignment
    - Restricted groups

    - Services
    - File system
    - Registry
- Advanced folder redirection policies.
- The GPO discretionary access control list (DACL), if you choose to preserve it during a copy operation.
- The DACL on software installation objects. This is only preserved if the option to copy the GPO DACL is specified. Otherwise the default DACL is used.

  **Note**

    - Security Principals referenced in Administrative Templates settings cannot be migrated using a migration table.

The following items can contain UNC paths and can be modified using a migration table:

- Folder redirection policies.
- Software installation policies (for software distribution points).
- Pointers to scripts deployed through Group Policy (such as startup and shutdown scripts) that are stored outside the GPO. Note that the script itself is not copied as part of the GPO copy operation, unless the script is stored inside the source GPO.

**Options for specifying migration tables**
Migration tables are specified when performing import and copy operations. There are three options for using migration tables with import and copy:

- Do not use a migration table. This option copies the GPO exactly as it is. All references to security principals and UNC paths are copied identically.
- Use a migration table. This option maps any references in the GPO that are in the migration table. References that are not contained in the migration table are copied as is.
- Use a migration table exclusively. This option requires that all references to security principals and UNC paths that are referenced in the source GPO be specified in the migration table. If a security principal or UNC path is referenced in the source GPO and is not included in the migration table, the import or copy operation fails.

In addition, cross-domain copy operations will apply the migration table to the DACL on the GPO (and any software installation settings) if you choose the option **Preserve or migrate the existing permissions**.

When performing a copy or import, the wizard scans the source GPO to determine if there are any references to security principals or UNC paths in the GPO. If there are, you have the opportunity to specify a migration table. During cross-domain copy operations, if you specify the option **Preserve or migrate the permissions on the GPO**, the wizard will always present the opportunity to specify a migration table, because a DACL by definition contains security principals.

**Contents of Migration tables**
A migration table consists of one or more mapping entries. Each mapping entry has three pieces of information: type, source and destination. The following sections describe these categories in more detail.

**Source Types**
Source types describe the type of domain-specific information for the source GPO. Migration tables support the following types of source types:

### XML Elements of Source Types

| Source Type | How This Type Appears in XML Editor |
|---|---|
| User | <Type>User</Type> |
| Computer | <Type>Computer</Type> |
| Domain Local Group | <Type>LocalGroup</Type> |
| Domain Global Group | <Type>GlobalGroup</Type> |
| Universal Group | <Type>UniversalGroup</Type> |
| UNC Path | <Type>UNCPath</Type> |
| Free Text or security identifier (SID). This category is only for use with security principals that are specified as free text and raw SIDs. | <Type>Unknown</Type> |

**Note**

- Built-in groups such as Administrators and Account Operators have the same SID in all domains. If references to built-in groups are stored in the GPO using their resolved format (based on the underlying SID), they cannot be mapped using migration tables. However, if references to built-in groups are stored as free text, you can map them using a migration table, and in this case, you must specify source type=Free Text or SID.

**Source Reference**

A source reference is the specific name of the User, Computer, Group or UNC Path referenced in the source GPO. For example, **\\server1\publicshare** is a specific example of a UNC path. The type of the source reference must match whatever source type has been specified in the migration table.

**Source Reference Syntax**

| Source Reference | Syntax |
|---|---|
| UPN | someone@example.com |
| SAM | example\someone |
| DNS | example.com\someone |
| Free text | someone (must be specified as the unknown type) |
| SID | S-1-11-111111111-111111111-1111111111-1112 (must be specified as the unknown type.) |

**Destination Name**

The destination name specifies how the name of the User, Computer, Group or UNC Path in the source GPO should be treated upon transfer to the destination GPO.

**Destination Name Options**

| Destination Name | Description |
|---|---|
| Same as source | Copy without changes. Equivalent to not putting the source value in the migration table at all.<br><br>MTE value: <Same As Source><br><br>XML tag:<DestinationSameAsSource /> |
| None | Removes the User, Computer or Group from the GPO. This option cannot be used with UNC paths<br><br>MTE value: <None><br><br>XML tag: <DestinationNone /> |

| | |
|---|---|
| Map by relative name | For example, map SourceDomain\Group1 to TargetDomain\Group1. This option cannot be used with UNC paths.<br><br>MTE value: <Map by Relative name><br><br>XML tag: <DestinationByRelativeName /> |
| Explicitly specify value | In the destination GPO, replace the source value with the exact literal value you specify.<br><br>MTE value: <exact name to use in target><br><br>XML tag: <Destination>Exact Name to use in Destination</Destination> |

**Note**

- Administrators can specify security principals for destination names using any of the syntactical formats described in source references, with the exception of using a raw SID. A raw SIDs can never be used for the destination name.

#### Default Entries

Valid migtable files must contain the following entries, which identifies the namespace used by migtables. If you are creating migtables by hand, you need to add this. Otherwise MTE does this for you.

```
<?xml version="1.0" encoding="utf-16"? >
<MigrationTable
xmlns="http://www.microsoft.com/GroupPolicy/GPOOperations/MigrationTable">
```

Each mapping entry appears in the following format:

```
<Mapping>
<Type>Type</Type>
<Source>Source</Source>
<Destination>Destination</Destination>
</Mapping>
```

#### The Migration Table Editor

The Migration Table Editor (MTE) is provided with Group Policy Management Console (GPMC) to facilitate the editing of migration tables. Migration tables are used for copying or importing Group Policy objects (GPOs) from one domain to another, in cases where the GPOs include domain-specific information that must be updated during copy or import.

MTE displays three columns of information:

### Migration Table Editor columns

| Source Name | Source Type | Destination Name |
|---|---|---|
| Name of the user, computer, group, or UNC path, referenced in the source GPO. | The type of the reference named in the Source Name column. For example, Domain Global Group. | The new value, after copy or import to the new GPO; or the method used to calculate the new value. |

Note that you do not specify Destination Type. The Destination Type is determined during the import or copy operation by checking the actual reference identified in the Destination Name.

#### Migration Table Editor features

- You can edit text fields using Cut, Copy, and Paste options, for example when you right-click an item in **the Source Name** column.
- You can add computers, users, and groups, for source and destination names, by using a **Browse** dialog box.
- You can fill in **Source Type** fields by using a drop-down menu.
- MTE provides automatic syntax checking to make sure the XML file is properly populated and also ensures that only compatible entries are entered in the table. For example, \\server1\share1 is not a valid security group name name, and

server1 is not a valid UNC path so MTE prompts the user to correct those entries.

- You can validate the overall migration table using MTE, in addition to basic syntax checking using the **Validate Table** option. The **Validate Table** option checks to ensure:

  - The existence of destination security principals and UNC paths. This is important to know before you copy or import a GPO, because if there are unresolvable entries in the migration table, then the copy or import operation might fail.

  - Source entries with UNC paths do not have destinations of MapByRelative name or NoDestination.

  - Type of the destination entry in the table matches the real type in the destination source.

- You can auto-populate a migration table by scanning one or more GPOs or backups to extract all references to security principals and UNC paths, and then enter these items into the table as source name entries. This capability is provided by the **Populate from GPO** and **Populate from Backup** options available on the **Tools** menu. To complete the migration table, you only need to adjust the destination values. By default, the destination name for each entry will be set to Same as source when you use either of the Populate options.

  **Note**

  - With either auto-populate option, when the source GPO or backup contains references to built-in groups such as Administrators or Backup Operators, these references will not be entered into the migration table if the references are stored in the source GPO or backup in their resolved format [based on security identifier (SID)]. This is because built-in groups have the same SID in all domains and therefore cannot be mapped to new values using a migration table. However, if references to built-in groups are stored as free text, they will be captured by the auto-populate options. In this case, you can map them as long as you specify the source type as Free Text or SID.

**Administrative Templates in GPMC and Group Policy Object Editor**

Administrators using Group Policy Management Console need to understand that Administrative template files are handled differently, depending upon whether they are working with Group Policy Management Console or Group Policy Object Editor. How each of the tools handles Administrative template files can be manipulated using Group Policy, but it is important to understand first the differences between the two.

**Administrative Templates**

Administrative templates, or .adm files, enable administrators to control registry settings using Group Policy. Windows comes with a predefined set of Administrative template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a GPO. These .adm files are stored in two locations by default: inside GPOs in the **Sysvol** folder and in the **%windir%\inf** directory on the local computer.

Windows Server 2003 includes the following .adm files: System.adm, Inetres.adm, Conf.adm, Wmplayer.adm, and Wuau.adm, which contain all the settings initially displayed in the **Administrative Templates** node.

### Default Administrative Templates

| Administrative Template File | Contains | For Use on | Description |
|---|---|---|---|
| System.adm | Settings to configure the Operating System | Windows 2000 or Windows Server 2003 | Loaded by default. |
| Inetres.adm | Settings to configure Internet Explorer | Windows 2000 or Windows Server 2003 | Loaded by default |
| Conf.adm | Settings to configure NetMeeting v3 | Windows 2000 or Windows Server 2003. Note: This tool is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. | Loaded by default. |

| | | | |
|---|---|---|---|
| Wmplayer.adm | Settings to configure Windows Media Player | Windows XP, Windows Server 2003.<br><br>**Note:** This tool is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. | Loaded by default. |
| Wuau.adm | Settings to configure Windows Update | Windows 2000 SP3, Windows XP SP1, Windows Server 2003 | Loaded by default. |

**Handling Administrative Template Files in GPMC**

GPMC uses .adm files to display the friendly names of policy settings when generating HTML reports for GPOs, Group Policy Modeling, and Group Policy Results.

By default, GPMC uses the local .adm file, regardless of time stamp. If the file is not found, then GPMC will look in the GPO's directory on Sysvol.

The user can specify an alternate path for where to find .adm files. If specified, this takes precedence over the previous locations.

GPMC never copies the .adm file to the Sysvol.

**Handling Administrative Template Files in Group Policy Object Editor**

Group Policy Object Editor uses .adm files to display available registry-based policy settings in the Administrative Templates section of a GPO. This includes Group Policy for the Windows Server 2003 operating system and its components and for applications.

By default, it attempts to read .adm files from the GPO (from the Sysvol on the domain controller). Alternatively, the .adm file can be read from the local workstation computer. This behavior can be controlled by a policy setting.

By default, if the version of the .adm file found on the local computer is newer (based on the time stamp of the file) than the version on the Sysvol, the local version is copied to the Sysvol and is then used to display the settings. This behavior can be controlled by a policy setting.

If the GPO contains registry settings for which there is no corresponding .adm file, these settings cannot be seen in Group Policy Object Editor. However, the policy settings are still active and will be applied to users or computers targeted by the GPO.

Policy settings pertaining to a user who logs on to a given workstation or server are written to the **User** portion of the registry database under **HKEY_CURRENT_USER**. Computer-specific settings are written to the **Local Machine** portion of the registry under **HKEY_LOCAL_MACHINE**.

**Related Information**

The following resources contain additional information that is relevant to this section:

- "GPMC Interfaces" on MSDN.
- Core Group Policy Technical Reference
- Group Policy Components

# Group Policy Management Console Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Management Console Tools and Settings**

**In this section**


- Group Policy Management Console Tools
- Group Policy Management Console Group Policy Settings
- Related Information


This section describes tools and settings that are associated with Group Policy Management Console (GPMC) Service Pack 1 (SP1).

**Group Policy Management Console Tools**

The following tools are associated with Group Policy Management Console (GPMC) SP1.

**Group Policy Management Console**

**Category**

The GPMC SP1 is a Microsoft Management Console (MMC) snap-in used to administer Group Policy for Windows 2000 and Windows Server 2003-based domains.

**Version compatibility**

GPMC runs on both Windows Server 2003 and on Windows XP SP1 with the Microsoft .NET Framework v1 and later, and with QFE Q326469 applied. QFE Q326469 is automatically installed by gpmc.msi on computers with English as the default language. GPMC can be used to manage both Windows 2000 and Windows Server 2003 domain Group Policy.

For more information about GPMC, see the What Is Group Policy Management Console? topic in this collection.

**Migration Tool Editor**

**Category**

The Migration Table Editor (MTE) is provided with GPMC to facilitate the editing of migration tables. Migration tables are used for copying or importing Group Policy objects (GPOs) from one domain to another, in cases where the GPOs include domain-specific information that must be updated during copy or import. When transferring the GPO to a new domain, it might not always be desirable, or even possible, to use the exact same settings. Items that can be specific to a domain include references to security principals such as users, groups, and computers as well as Universal Naming Convention (UNC) paths.

**Version compatibility**

The Migration Table Editor is provided with GPMC.

**Group Policy Management Console Group Policy Settings**

You can permit or prohibit use of GPMC through an Administrative Templates policy setting, as shown in the following table.

| Administrative Templates Policy Setting | Description |
|---|---|
| Group Policy Management<br><br>USER\Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy | This setting permits or prohibits the use of this snap-in. If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited. If this setting is not configured, the setting of the Restrict users to the explicitly permitted list of snap-ins setting determines whether this snap-in is permitted or prohibited. |

For a complete reference of all Administrative Templates policy settings, see Group Policy Settings Reference for Windows Server 2003.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Group Policy Settings Reference for Windows Server 2003.
- Tools and Settings Collection

# Group Policy Object Editor Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Object Editor Technical Reference**
Group Policy Object Editor is a Microsoft Management Console (MMC) snap-in used to edit the policy settings in Group Policy objects (GPOs). Like all MMC snap-ins, its functionality can be customized, or extended, by means of MMC snap-in extensions. When you use Group Policy Object Editor, various extensions are included by default.

**In this subject**


- What Is Group Policy Object Editor?
- How Group Policy Object Editor Works
- Group Policy Object Editor Tools and Settings

# What Is Group Policy Object Editor?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Group Policy Object Editor Core Scenarios
- Group Policy Object Editor Dependencies


Group Policy Object Editor is a Microsoft Management Console (MMC) snap-in used for configuring and modifying Group Policy settings within Group Policy objects (GPOs).

Administrators need to be able to quickly modify Group Policy settings for multiple users and computers throughout a network environment. The Group Policy Object Editor provides administrators with a hierarchical tree structure for configuring Group Policy settings in GPOs. These GPOs can then be linked to sites, domains, and organizational units (OU) containing computer or user objects.

Group Policy Object Editor consists of two main sections: **User Configuration**, which holds settings that are applied to users (at logon and periodic background refresh), and **Computer Configuration**, which holds settings that are applied to computers (at startup and periodic background refresh). These sections are further divided into the different types of policies that can be set, such as Administrative Templates, Security, or Folder Redirection.

To work efficiently, administrators need to have immediate access to information about the function and purpose of individual policy settings. For Administrative Templates policy settings, Group Policy Object Editor provides information about each policy setting directly in the Web view of the console. This information is called explain text. Explain text shows operating system requirements, defines the policy setting, and includes any specific details about the effect of enabling or disabling the policy setting.

In addition, developers should be able to quickly and easily add Group Policy support to their software products. The Group Policy Object Editor is designed to be extensible. The easiest way for developers to extend Group Policy Object Editor for their applications is to write custom Administrative Template files that "plug in" to Group Policy Object Editor.

**Group Policy Object Editor Core Scenarios**
There are two core scenarios for Group Policy Object Editor: editing GPOs, and extending the user interface (UI) to accommodate new applications or features. Both of these scenarios are described in detail in the following section.

**Editing Group Policy Objects**
Group Policy Object Editor is the primary tool used for configuring policy settings within a GPO. Group Policy Object Editor operates as an extension to Group Policy Management Console (GPMC). When an administrator elects to edit a GPO from within GPMC, Group Policy Object Editor appears, displaying the settings for that particular GPO. If GPMC is not available, Group Policy Object Editor operates as an extension to Active Directory management tools, such as the Active Directory Users and Computers snap-in or the Active Directory Sites and Services snap-in. Regardless of the tool an administrator uses to call Group Policy Object Editor, the primary function of Group Policy Object Editor is to edit settings within GPOs.

In addition to editing Active Directory-based GPOs, Group Policy Object Editor can also edit the local Group Policy object (local GPO). There is a local GPO stored on each computer running Windows 2000, Windows XP Professional, Windows XP 64-Bit Edition, or Windows Server 2003, regardless of whether the computers are part of an Active Directory environment. Local GPOs are always processed, but are the least influential GPOs in an Active Directory environment, because Active Directory-based GPOs have precedence.

Local GPOs do not support certain extensions, such as Folder Redirection or Group Policy Software Installation. Local GPOs do support many security settings, but the Security Settings extension of Group Policy Object Editor does not support remote management of local GPOs.

**Extending Group Policy Object Editor**
The Group Policy Object Editor is designed so that it can be extended. This architecture enables developers to create applications and features that can be administered with Group Policy. All of the nodes in the Group Policy Object Editor are themselves MMC snap-in extensions. These extensions include Administrative Templates, Scripts, Security Settings, Software Installation, Folder Redirection, Remote Installation Services, and Internet Explorer Maintenance. Extension snap-ins can in turn be extended. For example, the Security Settings snap-in includes several snap-in extensions.

Writing Group Policy extensions allows developers to use their own data store for policy settings, create a richer UI through the Group Policy Object Editor, or implement their own client-side processing for Group Policy. Developers can provide any one or all of the previous elements, or combine them with existing parts of the Group Policy infrastructure. For example, a developer could write a server-side extension snap-in to the Group Policy Object Editor that provides a UI on the server and a corresponding client-side extension that processes your data on the client. Alternately, a developer could write only a client-side extension and allow an existing Group Policy Object Editor extension to provide the UI, or vice versa.

**Extending Registry-Based Policy**

Registry-based policy is the least complex method to implement policy, and for administrators, registry-based policy settings are easy to configure and deploy. In addition, registry-based policy settings managed through administrative template (.adm) files automatically support Resultant Set of Policy (RSoP) capabilities.

Registry-based policy settings work by modifying the registry on the client computer. They utilize the Administrative Templates node of Group Policy Object Editor for a UI, and the default client side extension for registry-based policy to implement settings on the client. In addition to these components, registry-based policy settings also utilize an administrative template file. Written by the developer, an .adm file is a text file that specifies the registry-based policy that users modify through the Group Policy Object Editor. This file is then loaded into the Group Policy Object Editor. The Group Policy Object Editor displays the information contained in loaded .adm files as administrators navigate through the Administrative Templates node.

**Group Policy Object Editor Dependencies**

You must have edit rights on a GPO in order to open it in Group Policy Object Editor.

# How Group Policy Object Editor Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**

- Group Policy Object Editor Architecture
- Group Policy Object Editor Processes and Interactions

Administrators use Group Policy Object Editor to manipulate the settings in Group Policy objects (GPOs). The ideal environment contains:

- Windows Server 2003 with Active Directory installed
- Windows XP client computers

In this environment, an administrator can work with GPOs stored on the server. For these types of GPOs, all of the extensions to Group Policy Object Editor are functional. If Active Directory is not installed on your network, Group Policy Object Editor can be used only to edit the local GPO on individual client computers. For the local GPO, some of the extensions to Group Policy Object Editor are not functional.

**Group Policy Object Editor Architecture**

The Group Policy Object Editor is one of three administrative tools used to manage Group Policy. The following diagram shows all three of the tools as well as the domain controller and a client. In addition, the diagram describes the different communication protocols being used by each tool (LDAP, SMB, RPC/COM, DNS), the interactions between the tool, the domain controller and the client, and whether those interactions are READ or READ/WRITE.

**Group Policy Object Editor Architecture**




**Component Descriptions of the Group Policy Object Editor Architectural Diagram**

| Component | Description |
|---|---|

| | |
|---|---|
| Group Policy Object Editor | The Group Policy Object Editor is the Microsoft Management Console (MMC) snap-in used to edit GPOs. It was previously known as the Group Policy snap-in, Group Policy Editor, or Gpedit. The most notable feature of the Group Policy Object Editor is its extensibility. Administrators can extend the server-side snap-ins that ship with Group Policy Object Editor or they can develop completely new extensions for implementing Group Policy.<br><br>The Group Policy Object Editor is capable of read and write access to Active Directory, Sysvol, and the Local GPO. |
| Server Side Snap-ins | The nodes of the Group Policy Object Editor are also MMC server-side snap-ins. These snap-ins include Administrative Templates, Scripts, Security Settings, Software Installation, Folder Redirection, Remote Installation Services, and Internet Explorer Maintenance. Snap-ins can, in turn, be extended. For example, the Security Settings snap-in includes several extension snap-ins. Developers can also create their own MMC extension snap-ins to the Group Policy Object Editor to provide additional policy settings. Extensions are capable of read and write access to the Local GPO. |
| Server (Domain Controller) | In an Active Directory forest, the domain controller is a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. GPOs are stored in two parts of domain controllers: The Active Directory database and the Sysvol. |
| Active Directory | Active Directory, the Windows-based directory service, stores information about objects on a network and makes this information available to users and network administrators. Administrators link GPOs to Active Directory containers such as sites, domain, and OUs that include user and computer objects. In this way, policy settings can be targeted to users and computers throughout the organization. |
| Sysvol | Sysvol is a shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain. The Sysvol contains the largest part of a GPO, the Group Policy template (GPT), which includes Administrative Template-based policy settings, security settings, script files, and information regarding applications that are available for software installation. File Replication Service (FRS) replicates this information throughout the network.<br><br>Group Policy Object Editor uses .adm files to display available registry-based policy settings in the Administrative Templates section of a GPO. By default it attempts to read .adm files from the Sysvol on the domain controller. Alternatively, the .adm file can be read from the local workstation computer. This behavior can be controlled by a policy setting.<br><br>By default, if the time stamp of the .adm file found on the local computer is newer than the version on the Sysvol, the local version is copied to the Sysvol and is then used to display the settings. This behavior can be controlled by a policy setting.<br><br>If the GPO contains registry settings for which there is no corresponding .adm file, these settings cannot be seen in the Group Policy Object Editor. |
| LDAP Protocol | LDAP (Lightweight Directory Access Protocol) is the protocol used by the Active Directory directory service. Group Policy Object Editor uses LDAP to access the directory store on the domain controller. The client also uses LDAP to read the directory store on the domain controller. |
| SMB Protocol | SMB (Server Message Block) protocol is the primary method of file and print sharing. SMB can also be used for abstractions such as named pipes and mail slots. Group Policy Object Editor uses SMB to access the Sysvol as well as back up and retrieve files to a remote file system. The client also uses SMB to read the sysvol on the domain controller. |
| Group Policy Engine | The Group Policy Engine is a framework that handles common functionalities across client side extensions. |
| Client Side Extensions | Client-side extensions (CSEs) are dynamic-link libraries (DLLs) that are responsible for implementing Group Policy at the client computer. CSEs correspond to server-side snap-ins: Administrative Templates, Scripts, Security Settings, Software Installation, Folder Redirection, Remote Installation Services, and Internet Explorer Maintenance.<br><br>The client-side extensions are loaded on an as-needed basis when a client computer is processing policy. The client computer first gets a list of Group Policy objects. Next, it loops through all the client-side extensions and determines whether each client-side extension has any data in any of the GPOs. If a client-side extension has data in a GPO, the client-side extension is called with the list of Group Policy objects that it should process. If the client-side extension does not have any settings in any of the GPOs, it is not called. |

| | |
|---|---|
| Local Group Policy object | The local Group Policy object (local GPO) is stored on each individual computer, in the hidden **%systemroot%\System32\GroupPolicy** directory. Each computer running Windows 2000, Windows XP Professional, Windows XP 64-Bit Edition, or Windows Server 2003 has exactly one local GPO, regardless of whether the computers are part of an Active Directory environment.<br><br>Local GPOs do not support certain extensions, such as Folder Redirection or Group Policy Software Installation. Local GPOs do support many security settings, but the Security Settings extension of the Group Policy Object Editor does not support remote management of local GPOs. Local GPOs are always processed, but are the least influential GPOs in an Active Directory environment, because Active Directory-based GPOs have precedence. |

**Group Policy Object Editor Processes and Interactions**

Administrators using Group Policy Object Editor need to understand that Administrative template files are handled differently, depending upon whether they are working with Group Policy Object Editor or Group Policy Management Console. The way each of the tools handles Administrative template files can be manipulated using Group Policy, but it is important to understand the differences between the two.

**Administrative Templates in Group Policy Object Editor and GPMC**

Administrative templates, or .adm files, enable administrators to control registry settings using Group Policy. Windows comes with a predefined set of Administrative template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a GPO. These .adm files are stored in two locations by default: inside GPOs in the Sysvol folder and in the %windir%\inf directory on the local computer.

Windows Server 2003 includes the following .adm files: System.adm, Inetres.adm, Conf.adm, Wmplayer.adm, and Wuau.adm. These files contain all the settings initially displayed in the Administrative Templates node.

### Default Administrative Templates

| .Adm file | Contains | For Use on | Description |
|---|---|---|---|
| System.adm | Settings to configure the Operating System. | Windows 2000 or Windows Server 2003. | Loaded by default. |
| Inetres.adm | Settings to configure Internet Explorer. | Windows 2000 or Windows Server 2003. | Loaded by default. |
| Conf.adm | Settings to configure NetMeeting v3. | Windows 2000 or Windows Server 2003. Note: This tool is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. | Loaded by default. |
| Wmplayer.adm | Settings to configure Windows Media Player. | Windows XP, Windows Server 2003. Note: This tool is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. | Loaded by default. |
| Wuau.adm | Settings to configure Windows Update. | Windows 2000 SP3, Windows XP SP1, Windows Server 2003. | Loaded by default. |

**Handling .Administrative Template Files in Group Policy Object Editor**

Group Policy Object Editor uses .adm files to display available registry-based policy settings in the Administrative Templates section of a GPO. These settings include registry-based Group Policy for the Windows Server 2003 operating system and its components and for applications.

By default Group Policy Object Editor attempts to read .adm files from the GPO (from the Sysvol on the domain controller). Alternatively, the .adm file can be read from the local workstation computer. This behavior can be controlled by a policy setting.

By default, if the version of the .adm file found on the local computer is newer (based on the time stamp of the file) than the version on the Sysvol, the local version is copied to the Sysvol and is then used to display the settings. This behavior can be controlled by a policy setting.

If the GPO contains registry settings for which there is no corresponding .adm file, these settings cannot be seen in the Group Policy Object Editor. However, the policy settings are still active and will be applied to users or computers targeted by the GPO.

Policy settings pertaining to a user who logs on to a given workstation or server are written to the User portion of the registry database under **HKEY_CURRENT_USER**. Computer-specific settings are written to the Local Machine portion of the registry

under **HKEY_LOCAL_MACHINE**.

**Handling Administrative Template Files in GPMC**

GPMC uses .adm files to display the friendly names of policy settings when generating HTML reports for GPOs, Group Policy Modeling, and Group Policy Results.

By default, GPMC uses the local .adm file, regardless of time stamp. If the file is not found, then GPMC will look in the GPO's directory on Sysvol.

The user can specify an alternate path for where to find .adm files. If specified, this takes precedence over the previous locations.

GPMC never copies the .adm file to the Sysvol.

# Group Policy Object Editor Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Group Policy Object Editor Tools
- Group Policy Object Editor Group Policy Settings
- Related Information


This section describes the tools and settings that are associated with the Group Policy Object Editor server-side extension (SSE) when accessed from the Group Policy Management Console (GPMC) with Service Pack (SP) 1 applied. The Group Policy Object Editor SSE is a Microsoft Management Console (MMC) snap-in that administrators use to edit policy settings for Group Policy Objects (GPOs).

**Group Policy Object Editor Tools**
The following tools are associated with the Group Policy Object Editor.

**Gpedit.dll: Group Policy Object Editor**
**Category**
The Group Policy Object Editor is contained in gpedit.dll. Group Policy Object Editor is an MMC snap-in installed with Windows 2000, Windows XP, and Windows Server 2003. Gpedit.dll also contains the Resultant Set of Policy (RSoP) snap-in.

**Version compatibility**
You can access the Group Policy Object Editor from the Active Directory Users and Computers on Windows 2000, on Windows XP with the Windows Server 2003 Administration Tools Pack installed, and on Windows Server 2003. Administrators use the Group Policy Object Editor to edit policy settings in Group Policy objects (GPOs).

For a discussion of RSoP features provided by Gpedit.dll, see the What Is Resultant Set of Policy? topic in this collection. For more information about Group Policy Object Editor, see the What Is Group Policy Object Editor? topic in this collection.

**MMC.exe**
**Category**
MMC.exe provides a common host environment for snap-ins such as the Group Policy Object Editor and the GPMC.

**Version compatibility**
See MMC.exe documentation for version compatibility.

MMC.exe supplies programmatic interfaces that integrate snap-ins with the console. MMC.exe itself provides no management behavior. All management behavior is supplied by the snap-in. MMC interfaces deal only with user interface extensions.

**Group Policy Object Editor MMC Extensions**
The remaining tools in this section are Microsoft Management Console (MMC) snap-ins that are Group Policy server-side extensions. They extend the Group Policy Object Editor (gpedit.dll).

**Gptext.dll**
**Category**
Gptext.dll is an MMC extension to gpedit.dll that provides settings for Scripts, IP Security, and Wireless Group Policy.

**Version compatibility**
Gptext.dll is compatible with Windows XP SP1 and Windows Server 2003.

For more information about the Scripts, IP Security, and Wireless extensions, see the What Is Scripts Extension?, What Is Security Settings Extension?, and What Is Wireless Network Policies Extension? topics in this collection.

**Fdeploy.dll**
**Category**
Fdeploy.dll is an MMC extension to gpedit.dll that provides settings for Folder Redirection Group Policy.

**Version compatibility**
Fdeploy.dll is compatible with Windows XP SP1 and Windows Server 2003.

For more information about the Folder Redirection extension, see the What Is Folder Redirection Extension? topic in this collection.

**Scecli.dll**
**Category**
Scecli.dll is an MMC extension to gpedit.dll that provides settings for Security Settings Group Policy.

**Version compatibility**
Scecli.dll is compatible with Windows XP SP1 and Windows Server 2003.

For more information about the Security Settings extension, see the What Is Security Settings Extension? topic in this collection.

**Dskquota.dll**
**Category**
Dskquota.dll is an MMC extension to gpedit.dll that provides settings for Disk Quota Group Policy.

**Version compatibility**
Dskquota.dll is compatible with Windows XP1 and Windows Server 2003.

**Iedkcs32.dll**
**Category**
Iedkcs32.dll is an MMC extension to gpedit.dll that provides settings for Internet Explorer Maintenance Group Policy.

**Version compatibility**
Iedkcs32.dll is compatible with Windows XP SP1 and Windows Server 2003.

For more information about the Internet Explorer Maintenance extension, see the What Is Internet Explorer Maintenance Extension? topic in this collection.

**Appmgmts.dll**
**Category**
Appmgmts.dll is an MMC extension to gpedit.dll that provides settings for Software Installation Group Policy.

**Version compatibility**
Appmgmts.dll is compatible with Windows XP SP1 and Windows Server 2003.

For more information about the Software Installation extension, see the What Is Group Policy Software Installation Extension? topic in this collection.

**Group Policy Object Editor Group Policy Settings**
You can permit or prohibit use of Group Policy Object Editor through an Administrative Templates policy setting, as shown in the following table. MMC has policies that will prevent snap-ins from running. The list of snap-ins, including the Group Policy server side extension snap-ins, can be viewed in the results pane by selecting the tree-view node.

| Administrative Templates Policy Setting | Description |
|---|---|
| Group Policy Object Editor<br><br>USER Configuration\Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins\Group Policy | Permits or prohibits use of this snap-in. If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited. If this setting is not configured, the setting of the Restrict users to the explicitly permitted list of snap-ins setting determines whether this snap-in is permitted or prohibited. |

For a complete reference of all Administrative Templates policy settings, see Group Policy Settings Reference for Windows Server 2003.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Registry Reference for Windows Server 2003.
- Group Policy Settings Reference for Windows Server 2003.
- MSDN.

- Tools and Settings Collection.

# Resultant Set of Policy Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resultant Set of Policy Technical Reference**
Resultant Set of Policy (RSoP) snap-in is a Microsoft Management Console (MMC). Administrators use RSoP snap-in to report and plan the cumulative effects of Group Policy objects. Although administrators can use the RSoP snap-in for reporting and planning the effects of Group Policy, much of its functionality has been subsumed into Group Policy Management Console (GPMC), which provides a much better experience for the network administrator.

**In this subject**


- What Is Resultant Set of Policy?
- How Resultant Set of Policy Works
- Resultant Set of Policy Tools and Settings

# What Is Resultant Set of Policy?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Resultant Set of Policy?**

**In this section**


- Resultant Set of Policy Snap-in Core Scenarios
- Similar Technologies for Viewing Resultant Set of Policy Data
- Resultant Set of Policy Snap-in Dependencies


One challenge of Group Policy administration is to understand the cumulative effect of a number of Group Policy objects (GPOs) on any given computer or user, or how changes to Group Policy, such as reordering the precedence of GPOs or moving a computer or user to a different organizational unit (OU) in the directory, might affect the network. The Resultant Set of Policy (RSoP) snap-in offers administrators one solution. Administrators use the RSoP snap-in to see how multiple Group Policy objects affect various combinations of users and computers, or to predict the effect of Group Policy settings on the network.

**Resultant Set of Policy Snap-in Core Scenarios**

There are two core scenarios for the Resultant Set of Policy (RSoP) snap-in: reporting the effect of policy on various combinations of users and computers, and predicting the affect of policy. The reporting functionality of RSoP snap-in is known as logging mode. The predictive functionality is known as planning mode. Both of these scenarios are described in detail in the following respective sections.

Whether in logging or planning mode, an administrator typically accesses the RSoP snap-in by opening an empty Microsoft Management Console (MMC), adding the RSoP snap-in to the console, and then using the Resultant Set of Policy Wizard to collect data from various computers on the network for either logging or planning mode.

**Logging Mode**

An administrator uses logging mode to report on the current state of Group Policy settings. The scope of reports can include Group Policy settings for various targets, including a computer, a user, or both. The administrator selects combinations of targets in the Resultant Set of Policy Wizard.

When the administrator finishes with the wizard, the RSoP snap-in queries the Windows Management Instrumentation (WMI) repository on the client computer for information about the Group Policy settings for the targets. RSoP snap-in displays these settings.

Logging mode in the RSoP snap-in is useful for troubleshooting Group Policy. The RSoP snap-in lists each policy setting, and from which GPO the displayed setting came. GPO precedence is also available, as is any error information that was logged by either the Group Policy engine or any of the Group Policy client side extensions. Using this information, an administrator can determine which GPOs are applying a Group Policy setting and which GPOs are not.

In addition to using an empty MMC to access RSoP snap-in, there is another scenario an administrator is just as likely to use for troubleshooting Group Policy settings. An administrator can run RSOP.msc from the command prompt on a client computer to report on the current computer with the current user logged on. In this manner, the administrator avoids having to select targets in the Resultant Set of Policy Wizard.

**Planning Mode**

An administrator uses planning mode to perform “what if” scenarios with Group Policy. Like logging mode, the end result is a report of Group Policy settings. However, unlike logging mode, this report contains simulated data.

The administrator uses RSoP snap-in and a Windows Server 2003 domain controller. First, the administrator must use the Resultant Set of Policy Wizard to define the scope of the report. This is done by selecting various targets. In planning mode these targets must include both user information and computer information; however, the administrator can select either a particular user or computer or an Active Directory container for users or computers.

After specifying the scope of the plan, the administrator uses the wizard to specify a simulated Group Policy settings environment. Because these Group Policy environment settings are simulated, an actual client computer might behave differently in live tests. It is not possible for a domain controller to accurately determine how a particular client computer will actually behave.

The following simulated environment settings are available:

- A slow network connection
- Loopback processing (merge mode and replace mode)
- Site-based policy implementation
- Security group memberships for the user, the computer, or both
- WMI Filters that are linked to the user, the computer, or both

The administrator can elect to simulate any, all, or none of the environment variables. Whatever the choices, when the wizard finishes, the RSoP snap-in requests a simulation of policy from the domain controller. The domain controller uses Resultant Set of Policy Provider to simulate the application of GPOs. This service passes the GPO settings to virtual client-side extensions on the domain controller. The results of the simulation are stored in the WMI repository on the domain controller before the information is passed back to the RSoP snap-in for analysis. It is important to remember that the results displayed in the RSoP snap-in are not actual Group Policy settings, but simulated Group Policy based on the settings created using the wizard. If a custom client side extension exists on a client but does not exist on the domain controller, then any Group Policy settings this custom client side extension might create would not appear in the simulation results.

Planning mode in the RSoP snap-in is useful for planning Group Policy. The RSoP snap-in lists each GPO from which the displayed setting came as well as any other lower priority GPOs that attempted to configure settings. Using this information, an administrator can determine which GPOs are applying a policy setting and which GPOs are not.

**Similar Technologies for Viewing Resultant Set of Policy Data**
Although administrators can use the RSoP snap-in for reporting and planning the effects of Group Policy, much of its functionality has been subsumed into Group Policy Management Console (GPMC), which provides a much better experience for the network administrator.

The following table gives the equivalent names of GPMC and RSoP snap-in features that utilize the RSoP infrastructure.

**Equivalent Names and Descriptions of GPMC and RSoP Snap-in Features**

| GPMC feature | RSoP snap-in feature | Description |
|---|---|---|
| Group Policy Results | Logging Mode | Reports the effect of policy on a computer or user |
| Group Policy Modeling | Planning Mode | Predicts the affect of policy on the network |

Although GPMC provides functionality that subsumes most of the reporting features of RSoP snap-in, there is some Group Policy information that can only be reported on using RSoP snap-in. For example, RSoP snap-in lists each GPO from which the displayed setting came as well as any other lower priority GPOs that attempted to configure settings. Using this information, an administrator can determine which GPOs are applying a policy setting and which GPOs are not. In these cases, an administrator can use GPMC to open the RSoP snap-in by electing to view advanced information about a Group Policy Results or Group Policy Modeling report.

**Resultant Set of Policy Snap-in Dependencies**
Resultant Set of Policy snap-in has the following dependencies.

- Resultant Set of Policy snap-in is an MMC console; as such, it requires MMC infrastructure to function.
- Resultant Set of Policy infrastructure was introduced in Windows XP. Therefore, you cannot run gather RSoP data for a computer running Windows 2000 or lower.
- To run RSoP on a remote computer, you must be logged on as a member of the local Administrators group, or be delegated **Generate Resultant Set of Policy (logging)** rights for logging mode, or **Generate Resultant Set of Policy (planning)** rights for planning mode.
- To run RSoP planning mode, you must have a Windows Server 2003 domain controller.

- For planning mode involving security group membership, the administrator needs access to see membership of security groups.

# How Resultant Set of Policy Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Resultant Set of Policy Works**

**In this section**

- Resultant Set of Policy Snap-in Architecture
- Related Topics

Administrators can use the Resultant Set of Policy (RSoP) snap-in for two purposes: to predict the cumulative effect of Group Policy objects (GPOs), or to determine the actual result of Group Policy settings on a particular computer, user, or user on a computer.

Although administrators can use the RSoP snap-in for reporting and planning the effects of Group Policy, much of its functionality has been subsumed into Group Policy Management Console (GPMC), which provides a much better experience for the network administrator.

In an ideal environment, administrators are encouraged to use the GPMC features for simulating Group Policy or determining the effect of Group Policy on a particular user or computer.

**Resultant Set of Policy Snap-in Architecture**

The RSoP snap-in is one of three administrative tools used to manage Group Policy. The following diagram shows all three of the tools, as well as the domain controller and a client computer. In addition, the diagram describes the different communication protocols being used by each tool (LDAP, SMB, RPC/COM); the interactions between RSoP, the domain controller, and the client; and whether those interactions are READ or READ/WRITE.

**Resultant Set of Policy Snap-in Architectural Diagram**




**Component Description of Resultant Set of Policy Snap-in Architectural Diagram**

| Component | Description |
|---|---|

| | |
|---|---|
| Resultant Set of Policy snap-in | The RSoP snap-in is an MMC used to determine which policy settings are in effect for a given computer, user, or user on a computer, or to predict the effect of applied policy.<br><br>The RSoP snap-in itself is contained within the same binary as the Group Policy Object Editor. Thus, the user interface is a read-only view of the same information available in Group Policy Object Editor. However, there is one important difference: while Group Policy Object Editor can show all settings from a single GPO at a time, the RSoP snap-in can show the cumulative effect of many GPOs.<br><br>For RSoP snap-in functionality, administrators can use GPMC, which includes its own integrated RSoP infrastructure reporting features.<br><br>RSoP snap-in is capable of read access to the Active Directory, Sysvol, Event Log, RSoP infrastructure, and Local GPO on the target computer. Although RSoP is capable of read-only access to the Active Directory and Sysvol, most of the work of predicting or reporting Group Policy is done using RPC/COM communication with the RSoP provider, either on the client or the domain controller. |
| Domain Controller (Server) | In an Active Directory forest, the domain controller is a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. GPOs are stored in two parts of domain controllers: The Active Directory database (sometimes called Group Policy Container) and the Sysvol (known as the Group Policy template). |
| Active Directory | Active Directory, the Windows-based directory service, stores information about objects on a network and makes this information available to users and network administrators. Administrators link GPOs to Active Directory containers such as sites, domain, and OUs that include user and computer objects. In this way, policy settings can be targeted to users and computers throughout the organization. |
| Sysvol | Sysvol is a shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain. The Sysvol contains the largest part of a GPO: the Group Policy template, which includes Administrative Template-based policy settings, security settings, script files, and information regarding applications that are available for software installation. File Replication Service (FRS) replicates this information throughout the network. |
| LDAP Protocol | LDAP (Lightweight Directory Access Protocol) is the protocol used by the Active Directory directory service. RSoP snap-in uses LDAP for authentication and delegation checks. The client also uses LDAP to read the directory store on the domain controller. |
| SMB Protocol | SMB (Server Message Block) protocol is the primary method of file and print sharing. SMB can also be used for abstractions such as named pipes and mail slots. RSoP snap-in and the client use SMB to access the Sysvol on the domain controller. |
| RPC/COM | RPC (Remote Procedure Call), DCOM (Distributed Component Object Model) and COM (Component Object Model) enable data exchange between different processes. The different process can be on the same computer, on the local area network, or across the Internet.<br><br>COM and RPC are used by the RSoP snap-in for communication with the RSoP provider on the client or domain controller. |
| WMI | WMI is a management infrastructure that supports the monitoring and controlling of system resources through a common set of interfaces and provides a logically organized, consistent model of Windows operation, configuration, and status.<br><br>WMI makes available data about a target computer for administrative use. Such data can include hardware and software inventory, settings, and configuration information. For example, WMI exposes hardware configuration data such as CPU, memory, disk space, and manufacturer, as well as software configuration data from the registry, drivers, file system, Active Directory, the Windows Installer service, networking configuration, and application data. WMI Filtering in Windows Server 2003 allows you to create queries based on this data. These queries (also called WMI filters) determine which users and computers receive all of the policy configured in the GPO where you create the filter. |

| | |
|---|---|
| RSoP infrastructure | All Group Policy processing information is collected and stored in a namespace in WMI. This information, such as the list, content and logging of processing details for each GPO, can then be accessed by tools using WMI.<br><br>In logging mode, RSoP snap-in queries the database on the target computer, receives information about the policies and displays it in the RSoP snap-in.<br><br>In planning mode, RSoP snap-in simulates the application of policy using the Resultant Set of Policy Provider on a Domain Controller. This simulates the application of GPOs and passes them to Group Policy client-side extensions on the Domain Controller. The results of this simulation are stored to a local WMI database on the domain controller before the information is passed back and displayed in the RSoP snap-in. |
| Event Log | The Event log is a service that records events in various logs. The RSoP snap-in reads the Event Log on client computers and domain controllers in order to provide information about error events. |
| Local Group Policy object | The local Group Policy object (local GPO) is stored on each individual computer, in the hidden **%systemroot%\System32\GroupPolicy** directory. Each computer running Windows 2000, Windows XP Professional, Windows XP 64-Bit Edition, or Windows Server 2003 has exactly one local GPO, regardless of whether the computers are part of an Active Directory environment.<br><br>Local GPOs do not support certain extensions, such as Folder Redirection or Group Policy Software Installation. Local GPOs do support many security settings, but the Security Settings extension of the Group Policy Object Editor does not support remote management of local GPOs. Local GPOs are always processed, but can be overridden by domain policy in an Active Directory environment, because GPOs linked to Active Directory containers have precedence. |

**Related Information**

- Group Policy Management Console Technical Reference
- Group Policy Object Editor Technical Reference
- Core Group Policy Technical Reference

# Resultant Set of Policy Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resultant Set of Policy Tools and Settings**

**In this section**


- Resultant Set of Policy Tools
- Resultant Set of Policy Group Policy Settings
- Resultant Set of Policy WMI Classes
- Related Information


This section describes the tools and settings that are associated with Resultant Set of Policy (RSoP), a Microsoft Management Console (MMC) snap-in that administrators use to report and plan cumulative effects of Group Policy.

**Resultant Set of Policy Tools**

The following tools are associated with Resultant Set of Policy.

**RSoP Extension**

**Category**

The Resultant Set of Policy (RSoP) extension is an MMC snap-in included in the same binary as the Group Policy Object Editor (gpedit.dll). The RSoP extension consists of a planning mode service, and a logging mode service. These two services correspond to the Group Policy Modeling and Group Policy Results nodes, respectively, in Group Policy Management Console (GPMC).

**Version compatibility**

The RSoP feature, which consists of two services, requires a Windows Server 2003 domain controller for the planning mode, but requires only a Windows Server 2003 schema for remote logging delegation. A Windows Management Instrumentation (WMI) namespace, Web Based Enterprise Management (WBEM), on which each computer policy has processed, is required for logging mode. This namespace is present in Windows XP and later.

The RSoP snap-in displays RSoP data provided by one or both of its two services, depending on domain controller operating system and forest schema. The RSoP snap-in requires Windows XP or later. RSoP logging is always processed on the targeted client computer. For remote logging, the RSoP data is generated on the remote computer and displayed on the client computer running the RSoP snap-in.

For more information about Resultant Set of Policy, see the What Is Resultant Set of Policy? topic in this collection.

**Gpresult.exe: Group Policy Results**

**Category**

This command line tool displays RSoP for a user or computer. You can use gpresult to see what policy is in effect and to troubleshoot problems on computers running Windows XP and later.

**Version compatibility**

The gpresult.exe that ships with Windows XP and Windows Server 2003 family is completely different than the original version of gpresult.exe that shipped in the Windows 2000 Resource Kit. The newer version cannot be used to view policy information for computers running Windows 2000.

For more information about Group Policy Tools and Settings, see the Core Group Policy Tools and Settings topic in this collection.

**Resultant Set of Policy Group Policy Settings**

This section describes Resultant Set of Policy specific settings. Using these settings, the administrator can determine which MMC snap-ins and RSoP services are permitted for the GPO. For a complete reference of these settings, see Group Policy Settings Reference for Windows Server 2003.

### Group Policy Settings Associated with RSoP

| Group Policy Setting | Description |
| --- | --- |
| Resultant Set of Policy Provider (Computer Configuration/Windows Settings/Security Settings/System Services) | This setting controls enabling/disabling the RSoP service. |
| Turn off RSoP logging (Computer Configuration/Administrative Templates/System/Group Policy) | This setting allows you to enable or disable Resultant Set of Policy (RSoP) logging on a client computer. |
| Disallow Interactive Users from generating RSoP data (Computer and User Configuration/Administrative Templates/System/Group Policy) | This setting controls the ability of users to view their Resultant Set of Policy (RSoP) data.<br><br>If this setting is enabled, interactive users cannot generate RSoP data.<br><br>If this setting is not configured or disabled, interactive Users can generate RSoP. |
| Resultant Set of Policy snap-in (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the “Restrict users to the explicitly permitted list of snap-ins” setting determines whether this snap-in is permitted or prohibited. |
| Administrative Templates (Computers) (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy/RSoP snap-in extensions) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the “Restrict users to the explicitly permitted list of snap-ins” setting determines whether this snap-in is permitted or prohibited. |
| Administrative Templates (Users) (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy/RSoP snap-in extensions) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the “Restrict users to the explicitly permitted list of snap-ins” setting determines whether this snap-in is permitted or prohibited. |
| Folder Redirection (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy/RSoP snap-in extensions) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the “Restrict users to the explicitly permitted list of snap-ins” setting determines whether this snap-in is permitted or prohibited. |
| Internet Explorer Maintenance (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy/RSoP snap-in extensions) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the “Restrict users to the explicitly permitted list of snap-ins” setting determines whether this snap-in is permitted or prohibited. |

| | |
|---|---|
| Scripts (Logon/Logoff) (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy/RSoP snap-in extensions) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the "Restrict users to the explicitly permitted list of snap-ins" setting determines whether this snap-in is permitted or prohibited. |
| Scripts (Startup/Shutdown) (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy/RSoP snap-in extensions) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the "Restrict users to the explicitly permitted list of snap-ins" setting determines whether this snap-in is permitted or prohibited. |
| Security Settings (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy/RSoP snap-in extensions) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the "Restrict users to the explicitly permitted list of snap-ins" setting determines whether this snap-in is permitted or prohibited. |
| Software Installation (Computers) (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy/RSoP snap-in extensions) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the "Restrict users to the explicitly permitted list of snap-ins" setting determines whether this snap-in is permitted or prohibited. |
| Software Installation (Users) (User Configuration/Administrative Templates/Windows Components/MMC/Restricted/Permitted snap-ins/Group Policy/RSoP snap-in extensions) | This setting permits or prohibits use of this snap-in.<br><br>If you enable this setting, the snap-in is permitted. If you disable the setting, the snap-in is prohibited.<br><br>If this setting is not configured, the setting of the "Restrict users to the explicitly permitted list of snap-ins" setting determines whether this snap-in is permitted or prohibited. |

For more information about Group Policy settings, see the Group Policy Settings Reference for Windows Server 2003.

**Resultant Set of Policy WMI Classes**

There are two categories of WMI namespaces: computer and user. Each user has its own RSoP namespace based on the user's Security Identifier (SID). Administrators can automate RSoP using scripting and WMI classes. The following table lists the WMI classes associated with RSoP.

**RSoP WMI Classes**

| Class Name | Namespaces | Version Compatibility |
|---|---|---|
| RsopLoggingModeProvider | ROOT\RSOP | Windows XP<br><br>Windows Server 2003 |
| RsopPlanningModeProvider | ROOT\RSOP | Windows XP<br><br>Windows Server 2003<br><br>(Windows Server 2003 Domain Controller Required) |

In addition, administrators can automate RSoP using the **IGPMRSOP** interface included with GPMC interfaces. The **IGPMRSOP** interface provides methods that support making RSoP queries in both logging and planning mode.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Core Group Policy Tools and Settings.
- What Is Resultant Set of Policy?.
- Registry Reference for Windows Server 2003.
- Group Policy Settings Reference for Windows Server 2003.
- MSDN.
- Tools and Settings Collection.

# High Availability and Scalability Collection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**High Availability and Scalability Collection**
In a competitive environment, enterprises are required to have computer systems that are both highly available and easily scalable. High availability is essential to both customer satisfaction and employee productivity. Scalability is vital in responding to spikes in business and supporting the rapid rollout of new products and services.

The following technologies provide for high availability and scalability in a Microsoft Windows Server 2003 environment:

- Clustering Technologies
- Resource and Memory Management Technologies

High availability is the ability to provide user access to a service or application for a high percentage of scheduled time by attempting to reduce unscheduled outages and mitigate the impact of scheduled downtime for particular servers. A solution is highly available if it meets the organization's scheduled uptime goals. Included in the Windows Server 2003 family of operating systems are two technologies used to improve availability, server clusters and Network Load Balancing.

Scalability is the ability to increase or decrease computing capacity almost instantly, while delivering on value. Scalability is delivered through two methods: scale up and scale out. Typically, both scale up and scale out technologies are employed by organizations to address the capacity needs of Web applications, databases, and business applications. The scaling method used, scale out or scale up, depends on the types of applications being used and on the type of functionality needed. The scale up technologies included in the 32-bit versions of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, are: 4-gigabyte tuning (4GT), Physical Address Extension (PAE) X86, and Windows System Resource Manager. Available in all versions of the Windows Server 2003 family of products is Network Load Balancing, which can be used to not only improve availability, but also as a scale out solution.

High availability and scale up and scale out technologies are designed to respond to the changing requirements demanded of today's enterprise, regardless of where changes occur in the architecture. Whether it is handling Web site traffic, query processing, or large-scale database applications, high availability and scaling flexibility provide support for the most demanding Information Technology (IT) environments. By using these technologies, administrators can provide a highly available IT infrastructure that can incrementally scale as needed, without having to over-buy computing power.

### Notes

- Although not currently described in this Technical Reference, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, also support Cache Coherent Non-Uniform Memory Architecture (ccNUMA) and Large Symmetric Multiprocessing (Large SMP). These two scale up technologies enhance and optimize processing in large multi-CPU servers.
- Also not discussed in this collection of technologies, the 64-bit versions of Windows Server 2003 provide scale up capabilities as well. Because of the inherent capabilities of the 64-bit operating system, the scale up technologies described in this collection are not required on the 64-bit versions of the Windows Server 2003 family.

**Clustering Technologies**
Availability is the percentage of time that a system is available for users. Availability is increased by reducing the amount of time that a system is down for any reason, such as maintenance, upgrades, or recovery from a failure. Server clusters and Network Load Balancing can improve availability by removing the physical server as a single point of failure. A server cluster is a group of independent computer systems that work together as a single system. If one of the servers in the cluster becomes unavailable, another server in the cluster can automatically replace the unavailable server and continue to provide users access to the system.

Network Load Balancing combines up to 32 servers running any product in the Windows Server 2003 family into a single cluster. Each server can run an identical set of processes and applications. If one of the servers becomes unavailable, network requests can automatically be redirected to other servers in the Network Load Balancing cluster thus eliminating any loss of system availability.

### Note

- Maintaining highly available systems requires not only good choices in technology, but careful attention to design and operational practices. Implementing server clusters and Network Load Balancing will not ensure high availability. Many other factors, not covered in this technical reference, must also be considered and addressed.

**Resource and Memory Management Technologies**
Scaling up refers to incrementally adding system hardware resources (processors and memory) to increase overall system performance. Two of the resource and memory management technologies described in this technical reference used for scaling up, 4-gigabyte tuning (4GT) and Physical Address Extension (PAE) X86, are key components of the scale up capabilities. Further, Windows System Resource Manager provides for the effective management of the powerful servers used in scale up scenarios. Windows System Resource Manager makes it possible to assign, manage, and guarantee system resources for applications. Windows System Resource Manager also provides resource utilization statistics.

### Note

- Not only must the hardware and operating system be able to effectively scale up, but to gain the benefits of scale up capability, applications must be designed to scale up as well.

**Clustering Technologies**
The two clustering technologies described here are server clusters and Network Load Balancing. Both of these technologies provide a means for two or more computers to work together to provide a higher level of availability than can be obtained by using a single computer.

A brief description of these technologies is provided in the following sections.

For more information about clustering technologies, see "Clustering Technologies."

**Server Clusters**
A server cluster is a group of independent computer systems running Microsoft Windows Server 2003, Enterprise Edition, or Microsoft Windows Server 2003, Datacenter Edition, and working together as a single system to ensure that critical applications and resources remain available to clients. The servers in a cluster remain in constant communication through the exchange of periodic messages. If one of the servers becomes unavailable as a result of failure or maintenance, another server immediately begins providing service (a process known as failover). Server clusters also allow users and administrators to access and manage the servers as a single system rather than as separate computers.

**Network Load Balancing**
Network Load Balancing improves availability by clustering multiple servers running identical services and applications. The servers in the cluster can concurrently respond to different client requests, even multiple requests from the same client. Similar to server clusters, the servers in a Network Load Balancing cluster remain in constant communication through the exchange of periodic messages. If one of the servers becomes unavailable, the other servers in the cluster can continue to respond to client requests with no loss of availability. For more information about Network Load Balancing, see "Network Load Balancing Technical Reference."

**Resource and Memory Management Technologies**
4GT and PAE X86 increase the amount of RAM available to applications. Windows System Resource Manager makes it possible to assign, manage, and guarantee system resources for applications. Windows System Resource Manager also provides resource utilization statistics.

A brief description of these technologies is provided in the following sections.

For more information about scale up technologies, see "Resource and Memory Management Technologies."

### Notes

- Although not currently described in this Technical Reference, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, also support Cache Coherent Non-Uniform Memory Architecture (ccNUMA) and Large Symmetric Multiprocessing (Large SMP). These two scale up technologies enhance and optimize processing in large multi-CPU servers.

- Also not discussed in this collection of technologies, the 64-bit versions of Windows Server 2003 provide scale up capabilities as well. Because of the inherent capabilities of the 64-bit operating system, the scale up technologies described in this collection are not required on the 64-bit versions of the Windows Server 2003 operating system.

**4-gigabyte Tuning (4GT)**
Additional RAM is often added to servers as they are scaled up to meet increasing demands. Applications can benefit from accessing this additional memory. 4GT can make more of the computer's RAM available to applications by making less virtual memory available to the operating system. 4GT allows applications to address 3 gigabytes (GB) of virtual memory instead of the 2 GB normally allocated for user mode processes. This is a 50 percent increase in virtual memory for the applications, allowing more data to be cached and potentially increasing performance. For more information about 4GT, see "4GT Technical Reference."

**Physical Address Extension (PAE) X86**
As even more physical memory is added to servers, the operating system must be specifically enabled to access any memory beyond 4 GB. PAE X86 provides access to this additional memory. PAE X86 allows software that uses the Address Windowing Extensions (AWE) API set and runs on computers with more than 4 GB of physical memory to map more physical memory into the application's virtual address space. Applications not using the AWE API set can also benefit from PAE X86 because the operating system uses the larger physical memory to reduce paging and thus increase performance. For more information about PAE X86, see PAE X86 Technical Reference."

**Windows System Resource Manager**
Windows System Resource Manager (WSRM) is a tool used to manage server resources such as processor cycles and memory utilization. Using policies, Windows System Resource Manager allocates specific amounts of resources to applications or processes. These policies allow system administrators to set usage limits based on a percentage of available CPU resources as well as set affinities between managed applications and certain CPUs in the case of multi-processor servers. For more information about Windows System Resource Manager, see "Windows System Resource Manager Technical Reference."

**High Availability and Scalability Scenarios**
Each of the technologies described in this collection can be used in multiple ways for multiple scenarios. However, certain technologies are better suited for particular situations, as described in the following sections.

Additional deployment scenarios for 4GT, PAE X86, Windows System Resource Manager, and Network Load Balancing are described in the Technical Reference section for each technology.

**High availability**
Applications servicing customers as well as line-of-business applications such as databases, messaging servers, enterprise resource planning (ERP) applications and core file and print services are central to a company's operations and profitability. It is essential that these applications remain highly available. While both server clusters and Network Load Balancing can be used to improve the availability of many of these applications, each technology is better suited for certain scenarios. Server clusters are particularly valuable for messaging applications, database applications, and file and print services. Network Load Balancing is most commonly used with Web and File Transfer Protocol (FTP) servers, proxy servers and firewall services, such as computers running Microsoft Internet Security and Acceleration (ISA) Server, virtual private network (VPN) servers, servers running Windows Media technologies player, and terminal servers.

**Scale up**
Many of the same applications requiring high availability also present the need for scalability. Applications such as transaction processing applications and databases are architected to perform best on a single, monolithic server requiring scale up capabilities. The scale up technologies 4GT, PAE X86, and Windows System Resource Manager, are typically used in server consolidation scenarios and other enterprise systems that make use of these large servers with large amounts of memory and multiple CPUs. Some examples of scale up usage models include:

- Back-end database systems that drive ERP
- Customer relationship management (CRM) applications
- Decision support and data mining

**Scale out**
Scale out scenarios are ones in which applications are architected for distributed processing or that might need incremental hardware additions. Scale out scenarios also require the ability for centralized management of multiple servers. The same applications that can benefit from improved availability using Network Load Balancing can also often be scaled out with Network Load Balancing as well. As described earlier, these are Web and File Transfer Protocol (FTP) servers, proxy servers and firewall services, such as computers running Internet Security and Acceleration (ISA) Server, virtual private network (VPN) servers, servers running Windows Media technologies player, and terminal servers.

# Clustering Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Clustering Technologies**

The clustering technologies in products in the Microsoft Windows Server 2003 operating system are designed to help you achieve high availability and scalability for applications that are critically important to your business. These applications include corporate databases, e-mail, and Web-based services such as retail Web sites. By using appropriate clustering technologies and carefully implementing good design and operational practices (for example, configuration management and capacity management), you can scale your installation appropriately and ensure that your applications and services are available whenever customers and employees need them.

High availability is the ability to provide user access to a service or application for a high percentage of scheduled time by attempting to reduce unscheduled outages and mitigate the impact of scheduled downtime for particular servers. Scalability is the ability to easily increase or decrease computing capacity. A cluster consists of two or more computers working together to provide a higher level of availability, scalability, or both than can be obtained by using a single computer. Availability is increased in a cluster because a failure in one computer results in the workload being redistributed to another computer. Scalability tends to be increased, because in many situations it is easy to change the number of computers in the cluster.

Windows Server 2003 provides two clustering technologies: server clusters and Network Load Balancing (NLB). Server clusters primarily provide high availability; Network Load Balancing provides scalability and at the same time helps increase availability of Web-based services.

Your choice of cluster technologies (server clusters or Network Load Balancing) depends primarily on whether the applications you run have long-running in-memory state:

- **Server clusters** are designed for applications that have long-running in-memory state or frequently updated data. These are called stateful applications. Examples of stateful applications include database applications such as Microsoft SQL Server 2000 and messaging applications such as Microsoft Exchange Server 2003.

  Server clusters can combine up to eight servers.

- **Network Load Balancing** is intended for applications that do not have long-running in-memory state. These are called stateless applications. A stateless application treats each client request as an independent operation, and therefore it can load-balance each request independently. Stateless applications often have read-only data or data that changes infrequently. Web front-end servers, virtual private networks (VPNs), and File Transfer Protocol (FTP) servers typically use Network Load Balancing. Network Load Balancing clusters can also support other TCP- or UDP-based services and applications.

  Network Load Balancing can combine up to 32 servers.

In addition, with Microsoft Application Center 2000 Service Pack 2, you can create another type of cluster, a Component Load Balancing cluster. Component Load Balancing clusters balance the load between Web-based applications distributed across multiple servers and simplify the management of those applications. Application Center 2000 Service Pack 2 can be used with Web applications built on either the Microsoft Windows 2000 or Windows Server 2003 operating systems.

**Multitiered Approach for Deployment of Multiple Clustering Technologies**

Microsoft does not support the configuration of server clusters and Network Load Balancing clusters on the same server. Instead, use these technologies in a multitiered approach.

**Clustering Technologies Architecture**

A cluster consists of two or more computers (servers) working together. For server clusters, the individual servers are called nodes. For Network Load Balancing clusters, the individual servers are called hosts.

**Basic Architecture for Server Clusters**

The following diagram shows a four-node server cluster of the most common type, called a single quorum device cluster. In this type of server cluster, there are multiple nodes with one or more cluster disk arrays (often called the cluster storage) and a connection device (bus). Each of the disks in the disk array are owned and managed by only one node at a time. The quorum resource on the cluster disk array provides node-independent storage for cluster configuration and state data, so that each

node can obtain that data even if one or more other nodes are down.

**Four-Node Server Cluster Using a Single Quorum Device**




**Basic Architecture for Network Load Balancing Clusters**

The following diagram shows a Network Load Balancing cluster with eight hosts. Incoming client requests are distributed across the hosts. Each host runs a separate copy of the desired server application, for example, Internet Information Services. If a host failed, incoming client requests would be directed to other hosts in the cluster. If the load increased and additional hosts were needed, you could add them dynamically to the cluster.

**Network Load Balancing Cluster with Eight Hosts**




**Clustering Technologies Scenarios**

This section describes the most common scenarios for using server clusters and Network Load Balancing.

**Scenarios for Server Clusters**

This section provides brief descriptions of some of the scenarios for server cluster deployment. The scenarios cover three different aspects of server cluster deployment:

- The applications or services on the server cluster.
- The type of storage option: SCSI, Fibre Channel arbitrated loops, or Fibre Channel switched fabric.
- The number of nodes and the ways that the nodes can fail over to each other.

**Applications or Services on a Server Cluster**

Server clusters are usually used for services, applications, or other resources that need high availability. Some of the most common resources deployed on a server cluster include:

- **Printing**
- **File sharing**
- **Network infrastructure services**. These include the DHCP service and the WINS service.

- **Services that support transaction processing and distributed applications**. These services include the Distributed Transaction Coordinator (DTC) and Message Queuing.

- **Messaging applications**. An example of a messaging application is Microsoft Exchange Server 2003.

- **Database applications**. An example of a database application is Microsoft SQL Server 2000.

#### Types of Storage Options

A variety of storage solutions are currently available for use with server clusters. As with all hardware that you use in a cluster, be sure to choose solutions that are listed as compatible with Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. Also be sure to follow the vendor's instructions closely.

The following table provides an overview of the three types of storage options available for server clusters running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition:

**Storage Options for Server Clusters**

| Storage Option | Maximum Number of Supported Nodes |
|---|---|
| SCSI | Two |
| Fibre Channel arbitrated loop | Two |
| Fibre Channel switched fabric | Eight |

#### Number of Nodes and Failover Plan

Another aspect of server cluster design is the number of nodes used and the plan for application failover:

- **N-node Failover Pairs**. In this mode of operation, each application is set to fail over between two specified nodes.

- **Hot-Standby Server /N+I**. Hot-standby server operation mode reduces the overhead of failover pairs by consolidating the "spare" (idle) node for each pair into a single node, providing a server that is capable of running the applications from each node pair in the event of a failure. This mode of operation is also referred to as active/passive.

  For larger clusters, N+I mode provides an extension of the hot-standby server mode where N cluster nodes host applications and I cluster nodes are spare nodes.

- **Failover Ring**. In this mode of operation, each node in the cluster runs an application instance. In the event of a failure, the application on the failed node is moved to the next node in sequence.

- **Random**. For large clusters running multiple applications, the best policy in some cases is to allow the server cluster to choose the fail over node at random.

#### Scenarios for Network Load Balancing

This section provides brief descriptions of some of the scenarios for deployment of Network Load Balancing. The scenarios cover three different aspects of Network Load Balancing deployment:

- The types of servers or services in Network Load Balancing clusters.

- The number and mode of network adapters on each host.

#### Types of Servers or Services in Network Load Balancing Clusters

In Network Load Balancing clusters, some of the most common types of servers or services are as follows:

- Web and File Transfer Protocol (FTP) servers.

- ISA servers (for proxy servers and firewall services).

- Virtual private network (VPN) servers.
- Windows Media servers.
- Terminal servers.

**Number and Mode of Network Adapters on Each Network Load Balancing Host**

Another aspect of the design of a Network Load Balancing cluster is the number and mode of the network adapter or adapters on each of the hosts:

| **Number and Mode of Network Adapters on Each Host** | **Use** |
|---|---|
| Single network adapter in unicast mode | A cluster in which ordinary network communication among cluster hosts is not required and in which there is limited dedicated traffic from outside the cluster subnet to specific cluster hosts. |
| Multiple network adapters in unicast mode | A cluster in which ordinary network communication among cluster hosts is necessary or desirable. It is also appropriate when you want to separate the traffic used to manage the cluster from the traffic occurring between the cluster and client computers. |
| Single network adapter in multicast mode | A cluster in which ordinary network communication among cluster hosts is necessary or desirable but in which there is limited dedicated traffic from outside the cluster subnet to specific cluster hosts. |
| Multiple network adapters in multicast mode | A cluster in which ordinary network communication among cluster hosts is necessary and in which there is heavy dedicated traffic from outside the cluster subnet to specific cluster hosts. |

# Server Clusters Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server Clusters Technical Reference**
With the Microsoft Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, operating systems you can use server clusters to ensure that users have constant access to important server-based resources. With clustering, you create several cluster nodes that appear to users as one server. If one of the nodes in the cluster fails, another node begins to provide service (a process known as failover). In this way, server clusters can increase the availability and scalability of critical applications and resources.

Server clusters are based on one of the two clustering technologies in Windows Server 2003. The other clustering technology is Network Load Balancing. Server clusters are designed for applications that have long-running in-memory state or frequently updated data. Typical uses for server clusters include file servers, print servers, database servers, and messaging servers. Network Load Balancing is intended for applications that do not have long-running in-memory state. Typical uses for Network Load Balancing include Web front-end servers, virtual private networks (VPNs), and File Transfer Protocol (FTP) servers.

**In this subject**


- What Is a Server Cluster?
- How a Server Cluster Works
- Server Cluster Tools and Settings

# What Is a Server Cluster?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is a Server Cluster?**

**In this section**


- Introduction to Server Clusters
- Types of Server Clusters
- Related Information


A server cluster is a group of independent servers running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, and working together as a single system to provide high availability of services for clients. When a failure occurs on one computer in a cluster, resources are redirected and the workload is redistributed to another computer in the cluster. You can use server clusters to ensure that users have constant access to important server-based resources.

Server clusters are designed for applications that have long-running in-memory state or frequently updated data. Typical uses for server clusters include file servers, print servers, database servers, and messaging servers.

**Introduction to Server Clusters**

A cluster consists of two or more computers working together to provide a higher level of availability, reliability, and scalability than can be obtained by using a single computer. Microsoft cluster technologies guard against three specific types of failure:

- **Application and service failures,** which affect application software and essential services.
- **System and hardware failures,** which affect hardware components such as CPUs, drives, memory, network adapters, and power supplies.
- **Site failures in multisite organizations,** which can be caused by natural disasters, power outages, or connectivity outages.

The ability to handle failure allows server clusters to meet requirements for high availability, which is the ability to provide users with access to a service for a high percentage of time while reducing unscheduled outages.

In a server cluster, each server owns and manages its local devices and has a copy of the operating system and the applications or services that the cluster is managing. Devices common to the cluster, such as disks in common disk arrays and the connection media for accessing those disks, are owned and managed by only one server at a time. For most server clusters, the application data is stored on disks in one of the common disk arrays, and this data is accessible only to the server that currently owns the corresponding application or service.

Server clusters are designed so that the servers in the cluster work together to protect data, keep applications and services running after failure on one of the servers, and maintain consistency of the cluster configuration over time.

**Dependencies on Other Technologies**

Server clusters require network technologies that use IP-based protocols and depend on the following basic elements of network infrastructure:

- The Active Directory directory service (although server clusters can run on Windows NT, which does not use Active Directory).
- A name resolution service, that is, Windows Internet Name Service (WINS), the Domain Name System (DNS), or both. You can also use IP broadcast name resolution. However, because IP broadcast name resolution increases network traffic, and is ineffective in routed networks, within this Technical Reference we assume that you are using WINS or DNS.

In addition, for IP addressing for clients, Dynamic Host Configuration Protocol (DHCP) is often used.

**Note**

- The Cluster service does not support the use of IP addresses assigned from a DHCP server for the cluster administration address (which is an IP address resource associated with the cluster name) or any other IP address resources. If possible, we recommend that you use static IP addresses for all cluster systems.

**Types of Server Clusters**

There are three types of server clusters, based on how the cluster systems, called nodes, are connected to the devices that store the cluster configuration and state data. This data must be stored in a way that allows each active node to obtain the data even if one or more nodes are down. The data is stored on a resource called the quorum resource. The data on the quorum resource includes a set of cluster configuration information plus records (sometimes called checkpoints) of the most recent changes made to that configuration. A node coming online after an outage can use the quorum resource as the definitive source for recent changes in the configuration.

The sections that follow describe the three different types of server clusters:

- Single quorum device cluster, also called a standard quorum cluster
- Majority node set cluster
- Local quorum cluster, also called a single node cluster

**Single Quorum Device Cluster**

The most widely used cluster type is the single quorum device cluster, also called the standard quorum cluster. In this type of cluster there are multiple nodes with one or more cluster disk arrays, also called the cluster storage, and a connection device, that is, a bus. Each disk in the array is owned and managed by only one server at a time. The disk array also contains the quorum resource. The following figure illustrates a single quorum device cluster with one cluster disk array.

**Single Quorum Device Cluster**




Because single quorum device clusters are the most widely used cluster, this Technical Reference focuses on this type of cluster.

**Majority Node Set Cluster**

Windows Server 2003 supports another type of cluster, the majority node set cluster. In a majority node set cluster, each node maintains its own copy of the cluster configuration data. The quorum resource keeps configuration data consistent across the nodes. For this reason, majority node set clusters can be used for geographically dispersed clusters. Another advantage of majority node set clusters is that a quorum disk can be taken offline for maintenance and the cluster as a whole will continue to operate.

The major difference between majority node set clusters and single quorum device clusters is that single quorum device clusters can operate with just one node, but majority node set clusters need to have a majority of the cluster nodes available for the server cluster to operate. The following figure illustrates a majority node set cluster. For the cluster in the figure to continue to operate, two of the three cluster nodes (a majority) must be available.

**Majority Node Set Cluster**




This Technical Reference focuses on the single quorum device cluster.

**Local Quorum Cluster**

A local quorum cluster, also called a single node cluster, has a single node and is often used for testing. The following figure illustrates a local quorum cluster.

**Local Quorum Cluster**




This Technical Reference focuses on the single quorum device cluster, which is explained earlier in this section.

**Related Information**

The following sources provide information that is relevant to this section.

- "Planning Server Deployments" in the Windows Server 2003 Deployment Kit on the Microsoft Web site for information about planning for high availability and scalability, choosing among the different types of clustering technology, and designing and deploying server clusters.

- Network Load Balancing Technical Reference

# How a Server Cluster Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How a Server Cluster Works**

**In this section**


- Server Cluster Architecture
- Server Cluster API
- Server Cluster Processes
- Related Information


A server cluster is a collection of servers, called nodes that communicate with each other to make a set of services highly available to clients. Server clusters are based on one of the two clustering technologies in the Microsoft Windows Server 2003 operating systems. The other clustering technology is Network Load Balancing. Server clusters are designed for applications that have long-running in-memory state or frequently updated data. Typical uses for server clusters include file servers, print servers, database servers, and messaging servers.

This section provides technical background information about how the components within a server cluster work.

**Server Cluster Architecture**

The most basic type of cluster is a two-node cluster with a single quorum device. For a definition of a single quorum device, see "What Is a Server Cluster?." The following figure illustrates the basic elements of a server cluster, including nodes, resource groups, and the single quorum device, that is, the cluster storage.

**Basic Elements of a Two-Node Cluster with Single Quorum Device**




Applications and services are configured as resources on the cluster and are grouped into resource groups. Resources in a resource group work together and fail over together when failover is necessary. When you configure each resource group to include not only the elements needed for the application or service but also the associated network name and IP address, then that collection of resources runs as if it were a separate server on the network. When a resource group is configured this way,

clients can consistently get access to the application using the same network name, regardless of which node the application is running on.

The preceding figure showed one resource group per node. However, each node can have multiple resource groups. Within each resource group, resources can have specific dependencies. Dependencies are relationships between resources that indicate which resources need to come online before another resource can come online. When dependencies are configured, the Cluster service can bring resources online or take them offline in the correct order during failover.

The following figure shows two nodes with several resource groups in which some typical dependencies have been configured between resources. The figure shows that resource groups (not resources) are the unit of failover.

**Resource Dependencies Configured Within Resource Groups**




**Cluster Service Component Diagrams and Descriptions**

The Cluster service runs on each node of a server cluster and controls all aspects of server cluster operation. The Cluster service includes multiple software components that work together. These components perform monitoring, maintain consistency, and smoothly transfer resources from one node to another.

Diagrams and descriptions of the following components are grouped together because the components work so closely together:

- Database Manager (for the cluster configuration database)
- Node Manager (working with Membership Manager)
- Failover Manager
- Global Update Manager

Separate diagrams and descriptions are provided of the following components, which are used in specific situations or for specific types of applications:

- Checkpoint Manager
- Log Manager (quorum logging)
- Event Log Replication Manager
- Backup and Restore capabilities in Failover Manager

**Diagrams of Database Manager, Node Manager, Failover Manager, Global Update Manager, and Resource Monitors**

The following figure focuses on the information that is communicated between Database Manager, Node Manager, and Failover Manager. The figure also shows Global Update Manager, which supports the other three managers by coordinating updates on other nodes in the cluster. These four components work together to make sure that all nodes maintain a consistent view of the cluster (with each node of the cluster maintaining the same view of the state of the member nodes as the others) and that resource groups can be failed over smoothly when needed.

**Basic Cluster Components: Database Manager, Node Manager, and Failover Manager**




The following figure shows a Resource Monitor and resource dynamic-link library (DLL) working with Database Manager, Node Manager, and Failover Manager. Resource Monitors and resource DLLs support applications that are cluster-aware, that is, applications designed to work in a coordinated way with cluster components. The resource DLL for each such application is responsible for monitoring and controlling that application. For example, the resource DLL saves and retrieves application properties in the cluster database, brings the resource online and takes it offline, and checks the health of the resource. When failover is necessary, the resource DLL works with a Resource Monitor and Failover Manager to ensure that the failover happens smoothly.

**Resource Monitor and Resource DLL with a Cluster-Aware Application**




**Descriptions of Database Manager, Node Manager, Failover Manager, Global Update Manager, and Resource Monitors**
The following descriptions provide details about the components shown in the preceding diagrams.

**Database Manager**
Database Manager runs on each node and maintains a local copy of the cluster configuration database, which contains information about all of the physical and logical items in a cluster. These items include the cluster itself, cluster node membership, resource groups, resource types, and descriptions of specific resources, such as disks and IP addresses. Database Manager uses the Global Update Manager to replicate all changes to the other nodes in the cluster. In this way, consistent configuration information is maintained across the cluster, even if conditions are changing such as if a node fails and the administrator changes the cluster configuration before that node returns to service.

Database Manager also provides an interface through which other Cluster service components, such as Failover Manager and Node Manager, can store changes in the cluster configuration database. The interface for making such changes is similar to the interface for making changes to the registry through the Windows application programming interface (API). The key difference is that changes received by Database Manager are replicated through Global Update Manager to all nodes in the cluster.

**Database Manager functions used by other components**
Some Database Manager functions are exposed through the cluster API. The primary purpose for exposing Database Manager functions is to allow custom resource DLLs to save private properties to the cluster database when this is useful for a particular clustered application. (A private property for a resource is a property that applies to that resource type but not other resource types; for example, the SubnetMask property applies for an IP Address resource but not for other resource types.) Database Manager functions are also used to query the cluster database.

**Node Manager**
Node Manager runs on each node and maintains a local list of nodes, networks, and network interfaces in the cluster. Through regular communication between nodes, Node Manager ensures that all nodes in the cluster have the same list of functional nodes.

Node Manager uses the information in the cluster configuration database to determine which nodes have been added to the cluster or evicted from the cluster. Each instance of Node Manager also monitors the other nodes to detect node failure. It does this by sending and receiving messages, called heartbeats, to each node on every available network. If one node detects a communication failure with another node, it broadcasts a message to the entire cluster, causing all nodes that receive the message to verify their list of functional nodes in the cluster. This is called a regroup event.

Node Manager also contributes to the process of a node joining a cluster. At that time, on the node that is joining, Node Manager establishes authenticated communication (authenticated RPC bindings) between itself and the Node Manager component on each of the currently active nodes.

### Note

- A down node is different from a node that has been evicted from the cluster. When you evict a node from the cluster, it is removed from Node Manager's list of potential cluster nodes. A down node remains on the list of potential cluster nodes even while it is down; when the node and the network it requires are functioning again, the node joins the cluster. An evicted node, however, can become part of the cluster only after you use Cluster Administrator or Cluster.exe to add the node back to the cluster.

**Membership Manager**
Membership Manager (also called the Regroup Engine) causes a regroup event whenever another node's heartbeat is interrupted (indicating a possible node failure). During a node failure and regroup event, Membership Manager and Node Manager work together to ensure that all functioning nodes agree on which nodes are functioning and which are not.

**Cluster Network Driver**
Node Manager and other components make use of the Cluster Network Driver, which supports specific types of network communication needed in a cluster. The Cluster Network Driver runs in kernel mode and provides support for a variety of functions, especially heartbeats and fault-tolerant communication between nodes.

**Failover Manager and Resource Monitors**
Failover Manager manages resources and resource groups. For example, Failover Manager stops and starts resources, manages resource dependencies, and initiates failover of resource groups. To perform these actions, it receives resource and system state information from cluster components on the node and from Resource Monitors. Resource Monitors provide the execution environment for resource DLLs and support communication between resources DLLs and Failover Manager.

Failover Manager determines which node in the cluster should own each resource group. If it is necessary to fail over a resource group, the instances of Failover Manager on each node in the cluster work together to reassign ownership of the resource group.

Depending on how the resource group is configured, Failover Manager can restart a failing resource locally or can take the failing resource offline along with its dependent resources, and then initiate failover.

**Global Update Manager**
Global Update Manager makes sure that when changes are copied to each of the nodes, the following takes place:

- Changes are made atomically, that is, either all healthy nodes are updated, or none are updated.

- Changes are made in the order they occurred, regardless of the origin of the change. The process of making changes is coordinated between nodes so that even if two different changes are made at the same time on different nodes, when the changes are replicated they are put in a particular order and made in that order on all nodes.

Global Update Manager is used by internal cluster components, such as Failover Manager, Node Manager, or Database Manager, to carry out the replication of changes to each node. Global updates are typically initiated as a result of a Cluster API call. When an update is initiated by a node, another node is designated to monitor the update and make sure that it happens on all nodes. If that node cannot make the update locally, it notifies the node that tried to initiate the update, and changes are not made anywhere (unless the operation is attempted again). If the node that is designated to monitor the update can make the update locally, but then another node cannot be updated, the node that cannot be updated is removed from the list of functional nodes, and the change is made on available nodes. If this happens, quorum logging is enabled at the same time, which ensures that the failed node receives all necessary configuration information when it is functioning again, even if the original set of nodes is down at that time.

**Diagram and Description of Checkpoint Manager**
Some applications store configuration information locally instead of or in addition to storing information in the cluster configuration database. Applications might store information locally in two ways. One way is to store configuration information in the registry on the local server; another way is to use cryptographic keys on the local server. If an application requires that locally-stored information be available on failover, Checkpoint Manager provides support by maintaining a current copy of the local information on the quorum resource.

The following figure shows the Checkpoint Manager process.

### Checkpoint Manager




Checkpoint Manager handles application-specific configuration data that is stored in the registry on the local server somewhat differently from configuration data stored using cryptographic keys on the local server. The difference is as follows:

- For applications that store configuration data in the registry on the local server, Checkpoint Manager monitors the data while the application is online. When changes occur, Checkpoint Manager updates the quorum resource with the current configuration data.

- For applications that use cryptographic keys on the local server, Checkpoint Manager copies the cryptographic container to the quorum resource only once, when you configure the checkpoint. If changes are made to the cryptographic container, the checkpoint must be removed and re-associated with the resource.

Before a resource configured to use checkpointing is brought online (for example, for failover), Checkpoint Manager brings the locally-stored application data up-to-date from the quorum resource. This helps make sure that the Cluster service can recreate the appropriate application environment before bringing the application online on any node.

**Note**

- When configuring a Generic Application resource or Generic Service resource, you specify the application-specific configuration data that Checkpoint Manager monitors and copies. When determining which configuration information must be marked for checkpointing, focus on the information that must be available when the application starts.

Checkpoint Manager also supports resources that have application-specific registry trees (not just individual keys) that exist on the cluster node where the resource comes online. Checkpoint Manager watches for changes made to these registry trees when the resource is online (not when it is offline). When the resource is online and Checkpoint Manager detects that changes have been made, it creates a copy of the registry tree on the owner node of the resource and then sends a message to the owner node of the quorum resource, telling it to copy the file to the quorum resource. Checkpoint Manager performs this function in batches so that frequent changes to registry trees do not place too heavy a load on the Cluster service.

**Diagram and Description of Log Manager (for Quorum Logging)**

The following figure shows how Log Manager works with other components when quorum logging is enabled (when a node is down).

**Log Manager and Other Components Supporting Quorum Logging**




When a node is down, quorum logging is enabled, which means Log Manager receives configuration changes collected by other components (such as Database Manager) and logs the changes to the quorum resource. The configuration changes logged on the quorum resource are then available if the entire cluster goes down and must be formed again. On the first node coming online after the entire cluster goes down, Log Manager works with Database Manager to make sure that the local copy of the configuration database is updated with information from the quorum resource. This is also true in a cluster forming for the first time — on the first node, Log Manager works with Database Manager to make sure that the local copy of the configuration database is the same as the information from the quorum resource.

**Diagram and Description of Event Log Replication Manager**

Event Log Replication Manager, part of the Cluster service, works with the operating system's Event Log service to copy event log entries to all cluster nodes. These events are marked to show which node the event occurred on.

The following figure shows how Event Log Replication Manager copies event log entries to other cluster nodes.

**How Event Log Entries Are Copied from One Node to Another**




The following interfaces and protocols are used together to queue, send, and receive events at the nodes:

- The Cluster API
- Local remote procedure calls (LRPC)
- Remote procedure calls (RPC)
- A private API in the Event Log service

Events that are logged on one node are queued, consolidated, and sent through Event Log Replication Manager, which broadcasts them to the other active nodes. If few events are logged over a period of time, each event might be broadcast individually, but if many are logged in a short period of time, they are batched together before broadcast. Events are labeled to show which node they occurred on. Each of the other nodes receives the events and records them in the local log. Replication of events is not guaranteed by Event Log Replication Manager — if a problem prevents an event from being copied, Event Log Replication Manager does not obtain notification of the problem and does not copy the event again.

**Diagram and Description of Backup and Restore Capabilities in Failover Manager**

The Backup and Restore capabilities in Failover Manager coordinate with other Cluster service components when a cluster node is backed up or restored, so that cluster configuration information from the quorum resource, and not just information from the local node, is included in the backup. The following figure shows how the Backup and Restore capabilities in Failover Manager work to ensure that important cluster configuration information is captured during a backup.

**Backup Request on a Node That Does Not Own the Quorum Resource**




**DLLs Used by Core Resource Types**

A number of DLLs that are used by core resource types are included with server clusters in Windows Server 2003. The resource DLL defines and manages the resource. The extension DLL (where applicable) defines the resource's interaction with Cluster Administrator.

**Core Resource Types and Their Associated DLLs**

| Resource Types | DLLs |
|---|---|
| Physical Disk<br><br>Internet Protocol (IP) Address<br><br>Network Name<br><br>Print Spooler<br><br>File Share<br><br>Generic Application<br><br>Generic Script<br><br>Generic Service<br><br>Local Quorum<br><br>Majority Node Set | Resource DLL: Clusres.dll<br><br>Extension DLL: Cluadmex.dll |
| Volume Shadow Copy Service Task | Resource DLL: VSSTask.dll<br><br>Extension DLL: VSSTskEx.dll |

| | |
|---|---|
| DHCP Service<br><br>WINS Service | Resource DLL: Clnetres.dll<br><br>Extension DLL: Clnetrex.dll |
| Distributed Transaction Coordinator (DTC) | Resource DLL: Mtxclu.dll<br><br>Extension DLL: not applicable |
| Message Queuing | Resource DLL: Mqclus.dll<br><br>Extension DLL: not applicable |

**Cluster Service Files**

The following table lists files that are in the cluster directory (*systemroot*\cluster, where *systemroot* is the root directory of the server's operating system).

**Cluster Service Files in *Systemroot* \Cluster**

| File | Description |
|---|---|
| Cladmwiz.dll | Cluster Administrator Wizard |
| Clcfgsrv.dll | DLL file for Add Nodes Wizard and New Server Cluster Wizard |
| Clcfgsrv.inf | Setup information file for Add Nodes Wizard and New Server Cluster Wizard |
| Clnetres.dll | Resource DLL for the DHCP and WINS services |
| Clnetrex.dll | Extension DLL for the DHCP and WINS services |
| Cluadmex.dll | Extension DLL for core resource types |
| Cluadmin.exe | Cluster Administrator |
| Cluadmmc.dll | Cluster Administrator MMC extension |
| Clusres.dll | Cluster resource DLL for core resource types |
| Clussvc.exe | Cluster service |
| Debugex.dll | Cluster Administrator debug extension |
| Mqclus.dll | Resource DLL for Message Queuing |
| Resrcmon.exe | Cluster Resource Monitor |
| Vsstask.dll | Resource DLL for Volume Shadow Copy Service Task |
| Vsstskex.dll | Extension DLL for Volume Shadow Copy Service Task |
| Wshclus.dll | Winsock helper for the Cluster Network Driver |

The following table lists log files for server clusters.

**Log Files for Server Clusters**

| Log File | Folder Location | Description |
|---|---|---|
| cluster.log (default name) | *systemroot*\Cluster | Records the activity of the Cluster service, Resource Monitor, and resource DLLs on that node. The default name of this log can be changed by changing the System environment variable called **ClusterLog**. |
| cluster.oml | *systemroot*\Cluster | Records the creation and deletion of cluster objects and other activities of the Object Manager of the cluster; useful for a developer writing a tool for analyzing the translation of GUIDs to friendly names in the cluster. |

| | | |
|---|---|---|
| clcfgsrv.log | *systemroot*\system32\LogFiles\Cluster | Records activity of Cluster configuration wizards; useful for troubleshooting problems during cluster setup. |
| clusocm.log | *systemroot*\system32\LogFiles\Cluster | Records cluster-related activity that occurs during an operating system upgrade. |
| cluscomp.log | *systemroot*\system32\LogFiles\Cluster | Records the activity that occurs during the compatibility check at the start of an operating system upgrade on a cluster node. |

The following table lists files that are in *systemroot*\system32, *systemroot*\inf, or subfolders in *systemroot*\system32.

**Additional Cluster Service Files**

| File | Folder | Description |
|---|---|---|
| clusapi.dll | *systemroot*\system32 | Server Cluster API |
| clusocm.dll | *systemroot*\system32\Setup | Cluster extension for the Optional Component Manager |
| clusocm.inf | *systemroot*\inf | Cluster INF file for the Optional Component Manager |
| clussprt.dll | *systemroot*\system32 | A DLL that enables the Cluster service on one node to send notice of local cluster events to the Event Log service on other nodes |
| cluster.exe | *systemroot*\system32 | Cluster command-line interface |
| msclus.dll | *systemroot*\system32 | Cluster Automation Server |
| Resutils.dll | *systemroot*\system32 | Utility routines used by resource DLLs |
| Clusnet.sys | *systemroot*\system32\drivers | Cluster Network Driver |
| Clusdisk.sys | *systemroot*\system32\drivers | Cluster Disk Driver |

The following table lists files that have to do with the quorum resource and (for a single quorum device cluster, the most common type of cluster) are usually in the directory *q*:*mscs*, where *q* is the quorum disk drive letter and *mscs* is the name of the directory.

**Files Related to the Quorum Resource**

| File | Description |
|---|---|
| Quolog.log | The quorum log, which contains records of cluster actions that involve changes to the cluster configuration database. |
| Chk*.tmp | Copies of the cluster configuration database (also known as checkpoints). Only the latest one is needed. |
| {GUID} | Directory for each resource that requires checkpointing; the resource GUID is the name of the directory. |
| {GUID}*.cpt | Resource registry subkey checkpoint files. |
| {GUID}*.cpr | Resource cryptographic key checkpoint files. |

**Server Cluster API**

With the Server Cluster application programming interface (API), developers can write applications and resource DLLs for server clusters. The following table lists Server Cluster API subsets.

**Subsets in the Server Cluster API**

| API Subset | Description |
|---|---|
| Cluster API | Works directly with cluster objects and interacts with the Cluster service. |
| Resource API | Manages resources through a Resource Monitor and a resource DLL. |

| | |
|---|---|
| Cluster Administrator Extension API | Enables a custom resource type to be administered by Cluster Administrator. |

For more information, see Server Cluster APIson MSDN.

**Server Cluster Processes**

The following sections provide information about the following processes:

- How nodes form, join, and leave a cluster.
- How heartbeats, regroup events, and quorum arbitration work in a cluster. These processes help the cluster to keep a consistent internal state and maintain availability of resources even when failures occur.
- How resource groups are brought online, taken offline, failed over, and failed back.

**How Nodes Form, Join, and Leave a Cluster**

Nodes must form, join, and leave a cluster in a coordinated way so that the following are always true:

- Only one node owns the quorum resource at any given time.
- All nodes maintain the same list of functioning nodes in the cluster.
- All nodes can maintain consistent copies of the cluster configuration database.

**Forming a Cluster**

The first server that comes online in a cluster, either after installation or after the entire cluster has been shut down for some reason, forms the cluster. To succeed at forming a cluster, a server must:

- Be running the Cluster service.
- Be unable to locate any other nodes in the cluster (in other words, no other nodes can be running).
- Acquire exclusive ownership of the quorum resource.

If a node attempts to form a cluster and is unable to read the quorum log, the Cluster service will not start, because it cannot guarantee that it has the latest copy of the cluster configuration. In other words, the quorum log ensures that when a cluster forms, it uses the same configuration it was using when it last stopped.

The sequence in which a node forms a cluster is as follows:

1. The node confirms that it can start the Cluster service.
2. The node reviews the information stored in the local copy of the cluster configuration database.
3. Using information from the local copy of the cluster configuration database, the node confirms that no other nodes are running.

   If another node is running, then the node that started most recently joins the cluster instead of forming it.
4. Using information from the local copy of the cluster configuration database, the node locates the quorum resource.
5. The node confirms that it can acquire exclusive ownership of the quorum resource and that it can read from the quorum resource. If it can, the node takes ownership.
6. The node compares the sequence numbers on the copy of the cluster configuration database on the quorum resource and the sequence numbers on the quorum log against the sequence numbers on the node's local copy of the cluster configuration database.

7. The node updates its local copy of the cluster configuration database with any newer information that might be stored on the quorum resource.

8. The node begins to bring resources and resource groups online.

**Joining a Cluster**
The sequence in which a server joins an existing cluster is as follows:

1. The node confirms that it can start the Cluster service.

2. The node reviews the information stored in the local copy of the cluster configuration database.

3. The node that is joining the cluster tries to locate another node (sometimes called a sponsor node) running in the cluster. The node goes through the list of other nodes in its local configuration database, trying one or more until one responds.

   If no other nodes respond, the node tries to form the cluster, starting by locating the quorum resource.

4. Node Manager on the sponsor node authenticates the new node. If the joining node is not recognized as a cluster member, the request to join the cluster is refused.

5. Node Manager on the joining node establishes authenticated communication (authenticated RPC bindings) between itself and the Node Manager component on each of the currently active nodes.

6. Database Manager on the joining node checks the local copy of the configuration database. If it is outdated, Database Manager obtains an updated copy from the sponsor node.

**Leaving a Cluster**
A node can leave a cluster when the node shuts down or when the Cluster service is stopped. When a node leaves a cluster during a planned shutdown, it attempts to perform a smooth transfer of resource groups to other nodes. The node leaving the cluster then initiates a regroup event.

Functioning nodes in a cluster can also force another node to leave the cluster if the node cannot perform cluster operations, for example, if it fails to commit an update to the cluster configuration database.

**Heartbeats, Regroup Events, and Quorum Arbitration**
When server clusters encounter changing circumstances and possible failures, the following processes help the cluster to keep a consistent internal state and maintain availability of resources:

- Heartbeats

- Regroup events

- Quorum arbitration

**Heartbeats**
Heartbeats are single User Datagram Protocol (UDP) packets exchanged between nodes once every 1.2 seconds to confirm that each node is still available. If a node is absent for five consecutive heartbeats, the node that detected the absence initiates a regroup event to make sure that all nodes reach agreement on the list of nodes that remain available.

Server cluster networks can be private (node-to-node communication only), public (client-to-node communication), or mixed (both node-to-node and client-to-node communication). Heartbeats are communicated across all networks; however, the monitoring of heartbeats and the way the cluster interprets missed heartbeats depends on the type of network:

- On private or mixed networks, which both carry node-to-node communication, heartbeats are monitored to determine whether the node is functioning in the cluster.

  A series of missed heartbeats can either mean that the node is offline or that all private and mixed networks are offline; in either case, a node has lost its ability to function in the cluster.

- On public networks, which carry only client-to-node communication, heartbeats are monitored only to determine whether a node's network adapter is functioning.

**Regroup Events**
If a node is absent for five consecutive heartbeats, a regroup event occurs. (Membership Manager, described earlier, starts the regroup event.)

If an individual node remains unresponsive, the node is removed from the list of functional nodes. If the unresponsive node was the owner of the quorum resource, the remaining nodes also begin the quorum arbitration process. After this, failover begins.

**Quorum Arbitration**
Quorum arbitration is the process that occurs when the node that owned the quorum resource fails or is unavailable, and the remaining nodes determine which node will take ownership. When a regroup event occurs and the unresponsive node owned the quorum resource, another node is designated to initiate quorum arbitration. A basic goal for quorum arbitration is to make sure that only one node owns the quorum resource at any given time.

It is important that only one node owns the quorum resource because if all network communication between two or more cluster nodes fails, it is possible for the cluster to split into two or more partial clusters that will try to keep functioning (sometimes called a "split brain" scenario). Server clusters prevent this by allowing only the partial cluster with a node that owns the quorum resource to continue as the cluster. Any nodes that cannot communicate with the node that owns the quorum resource stop working as cluster nodes.

**How Clusters Keep Resource Groups Available**
This section describes how clusters keep resource groups available by monitoring the health of resources (polling), bringing resource groups online, and carrying out failover. Failover means transferring ownership of the resources within a group from one node to another. This section also describes how resource groups are taken offline as well as how they are failed back, that is, how resource groups are transferred back to a preferred node after that node has come back online.

Transferring ownership can mean somewhat different things depending on which of the group's resources is being transferred. For an application or service, the application or service is stopped on one node and started on another. For an external device, such as a Physical Disk resource, the right to access the device is transferred. Similarly, the right to use an IP address or a network name can be transferred from one node to another.

Resource-related activities in server clusters include:

- Monitoring the health of resources (polling).
- Bringing a resource group online.
- Taking a resource group offline.
- Failing a resource group over.
- Failing a resource group back.

The administrator of the cluster initiates resource group moves, usually for maintenance or other administrative tasks. Group moves initiated by an administrator are similar to failovers in that the Cluster service initiates resource transitions by issuing commands to Resource Monitor through Failover Manager.

**Resource Health Monitoring (Polling)**
Resource Monitor conducts two kinds of polling on each resource that it monitors: Looks Alive (resource appears to be online) and Is Alive (a more thorough check indicates the resource is online and functioning properly).

When setting polling intervals, it can be useful to understand the following:

- If a Generic Application resource has a long startup time, you can lengthen the polling interval to allow the resource to finish starting up. In other words, you might not require a custom resource DLL to ensure that the resource is given the necessary startup time.
- If you lengthen the polling intervals, you reduce the chance that polls will interfere with each other (the chance for lock contention).

- You can bypass Looks Alive polling by setting the interval to 0.

**How a Resource Group Comes Online**
The following sequence is used when Failover Manager and Resource Monitor bring a resource group online.

1. Failover Manager uses the dependency list (in the cluster configuration) to determine the appropriate order in which to bring resources online.

2. Failover Manager works with Resource Monitor to begin bringing resources online. The first resource or resources started are ones that do not depend on other resources.

3. Resource Monitor calls the **Online** entry point of the first resource DLL and returns the result to Failover Manager.

    - If the entry point returns ERROR_IO_PENDING, the resource state changes to **OnlinePending**. Resource Monitor starts a timer that waits for the resource either to go online or to fail. If the amount of time specified for the pending timeout passes and the resource is still pending (has not entered either the Online or Failed state), the resource is treated as a failed resource and Failover Manager is notified.

    - If the Online call fails or the **Online** entry point does not move the resource into the Online state within the time specified in the resource DLL, the resource enters the Failed state, and Failover Manager uses Resource Monitor to try to restart the resource, according to the policies defined for the resource in its DLL.

    - When the resource enters the Online state, Resource Monitor adds the resource to its list of resources and starts monitoring the state of the resource.

4. The sequence is repeated as Failover Manager brings the next resource online. Failover Manager uses the dependency list to determine the correct order for bringing resources online.

After resources have been brought online, Failover Manager works with Resource Monitor to determine if and when failover is necessary and to coordinate failover.

**How a Resource Group Goes Offline**
Failover Manager takes a resource group offline as part of the failover process or when an administrator moves the group for maintenance purposes. The following sequence is used when Failover Manager takes a resource group offline:

1. Failover Manager uses the dependency list (in the cluster configuration) to determine the appropriate order in which to bring resources offline.

2. Failover Manager works with Resource Monitor to begin taking resources offline. The first resource or resources stopped are ones on which other resources do not depend.

3. Resource Monitor calls the **Offline** entry point of the resource DLL and returns the result to Failover Manager.

    - If the entry point returns ERROR_IO_PENDING, the resource state changes to **OfflinePending**. Resource Monitor starts a timer that waits for the resource either to go offline or to fail. If the amount of time specified for the pending timeout passes and the resource is still pending (has not entered either the Offline or Failed state), the resource is treated as a failed resource and Failover Manager is notified.

    - If the Offline call fails or the **Offline** entry point does not move the resource into the Offline state within the time specified in the resource DLL, Failover Manager uses Resource Monitor to terminate the resource and the resource enters the Failed state.

    - If the Offline call succeeds, Resource Monitor takes the resource off its list and stops monitoring the resource.

4. The sequence is repeated as Failover Manager brings the next resource offline. Failover Manager uses the dependency list to determine the correct order for taking resources offline.

**How Failover Occurs**

Group failover happens when the group or the node that owns the group fails. Individual resource failure causes the group to fail if you configure the **Affect the group** property for the resource.

Failover takes two forms, as described in the sections that follow:

- Resource failure and group failure (without node failure)
- Node failure or loss of communication between nodes

**Resource Failure and Group Failure (Without Node Failure)**

When a resource fails, the following process occurs:

1. Resource Monitor detects the failure, either through Looks Alive or Is Alive polling or through an event signaled by the resource. Resource Monitor calls the **IsAlive** entry point of the resource DLL to confirm that the resource has failed.

2. If **IsAlive** fails, the state of the resource changes to Failed.

   If you configured the resource to be restarted on failure, Failover Manager attempts to restart the resource by trying to bring it online. If the attempts to bring the resource online fail more than is allowed by the restart **Threshold** and **Period** properties, Resource Monitor stops polling the resource.

3. Through Resource Monitor, Failover Manager calls the **Terminate** entry point of the resource DLL.

   The rest of this process concerns how the group fails over.

4. If the resource is set to **Affect the group**, the sequence continues. Otherwise, it ends without further action. If the resource is set to **Affect the group**, Failover Managers on the cluster nodes work together to reassign ownership of the group.

5. On the node on which the resource failed, Failover Manager terminates the resource that failed and the resources that depend on it, and then Failover Manager takes the remaining resources in the dependency tree offline in order of dependencies.

6. Failover Manager on the node on which the resource failed notifies Failover Manager on the node that will take ownership of the resource (and also notifies Failover Manager on other nodes about the changes that are happening).

7. If any of the resources have been configured so that application-specific configuration information (registry subkeys) for that resource is checkpointed, Checkpoint Manager restores the saved registry subkeys for those resources from the quorum resource.

8. The destination node Failover Manager brings the resources online one by one, using the dependency list to determine the correct order.

9. The node that now owns the group turns control of the group's resources over to their respective Resource Monitors.

**Node Failure or Loss of Communication Between Nodes**

Failover that occurs when a node fails is different from failover that occurs when a resource fails. For the purposes of clustering, a node is considered to have failed if it loses communication with other nodes.

As described in previous sections, if a node misses five heartbeats, this indicates that it has failed, and a regroup event (and possibly quorum arbitration) occurs. After node failure, surviving nodes negotiate for ownership of the various resource groups. On two-node clusters the result is obvious, but on clusters with more than two nodes, Failover Manager on the surviving nodes determines group ownership based on the following:

- The nodes you have specified as possible owners of the affected resource groups.
- The order in which you specified the nodes in the group's Preferred Owners list.

**Note**

- When setting up a preferred owners list for a resource group, we recommend that you list *all* nodes in your server cluster and put them in priority order.

**How Failback Occurs**

Failback is the process by which the Cluster service moves resource groups back to their preferred node after the node has failed and come back online. You can configure both *whether* and *when* failback occurs. By default, groups are not set to fail back.

The node to which the group will fail back initiates the failback. Failover Manager on that node contacts Failover Manager on the node where the group is currently online and negotiates for ownership. The instances of Failover Manager on the two nodes work together to smoothly transfer ownership of the resource group back to the preferred node.

You can test failback configuration settings by following procedures in Help and Support Center.

**Related Information**

The following sources provide information that is relevant to this section.

- "Planning Server Deployments" in the Windows Server 2003 Deployment Kit on the Microsoft Web site for more information about failover policies, choices for cluster storage, and ways that applications can operate within a server cluster.
- What Is a Server Cluster?

# Server Cluster Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server Cluster Tools and Settings**

**In this section**


- Server Cluster Tools
- Server Cluster WMI Classes
- Related Information


The information in this section is helpful for learning about tools to administer, diagnose, and recover server clusters. It also provides information about tools for configuring disks and for migrating a clustered print server to Windows Server 2003 operating systems. In addition, this section provides information about WMI classes associated with server clusters.

**Server Cluster Tools**

The following tools are associated with server clusters.

**Cluadmin.exe: Cluster Administrator**

**Category**

Tool included in Windows Server 2003, Standard Edition, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, operating systems. The tool is also included in the Windows Server 2003 Administration Tools Pack.

**Version Compatibility**

Cluster Administrator can run on computers running Windows Server 2003, Standard Edition, Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, and Windows XP Professional with Windows Server 2003 Administration Tools Pack installed.

Cluster Administrator can target server cluster nodes running Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, and Windows NT Server 4.0, Enterprise Edition.

Cluster Administrator is the graphical user interface (GUI) for server clusters. It displays information about the cluster and its nodes, resources, and groups. With Cluster Administrator, you can do a variety of tasks, including joining nodes to a cluster, managing cluster objects, establishing groups, initiating failover, monitoring cluster activity, and other tasks.

**Cluster.exe**

**Category**

Tool included in Windows Server 2003, Standard Edition, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, operating systems. The tool is also included in the Windows Server 2003 Administration Tools Pack.

**Version Compatibility**

Cluster.exe can run on computers running Windows Server 2003, Standard Edition, Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, and Windows XP Professional with Windows Server 2003 Administration Tools Pack installed.

Cluster.exe can target server cluster nodes that are running Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, and Windows NT Server 4.0, Enterprise Edition.

- Cluster.exe is the command-line interface for server clusters. Cluster.exe provides all the functionality of Cluster Administrator, the graphical user interface (GUI), plus several additional functions:
- The **cluster resource** command includes the **/private** option, which can be used to view or change a resource's private properties in the cluster configuration database.
- The **cluster** command includes the **/changepass** option, which can be used to change the Cluster service account password.

- The **cluster resource** command (with no options) can provide information about a cluster resource when Cluster Administrator does not display information quickly, which sometimes happens if a resource is in a pending state.

- Commands can be put together into scripts that can be run routinely, possibly by staff members less highly trained than those who wrote the scripts.

For more information about Cluster.exe, see "Command line reference A-Z" on the Microsoft TechNet Web site.

**Clusdiag.msi: Cluster Diagnostics and Verification Tool**
**Category**
Tool included in the Windows Server 2003 Resource Kit. An updated version of this tool is available.

**Version Compatibility**
The Cluster Diagnostics tool can run on computers running Windows Server 2003, Windows 2000, and Windows XP Professional.

The updated version of the tool can target server clusters running Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

The original version of the tool can target server clusters running Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

The Cluster Diagnostics tool provides a graphical user interface (GUI) for performing diagnostic tests on a preproduction cluster, as well as an interface for viewing multiple log files to aid in troubleshooting cluster failures. The tool can generate various cluster-related reports from the information gathered during testing. An example of a report that the tool can generate is a diagram showing resource dependencies.

To find more information about Clusdiag.msi, see Resource Kit Tool Updates in the Tools and Settings Collection.

**Clusterrecovery.exe: Server Cluster Recovery Utility**
**Category**
Tool included in the Windows Server 2003 Resource Kit.

**Version Compatibility**
The Server Cluster Recovery Utility can be installed on computers running Windows Server 2003 and Windows XP Professional.

The Server Cluster Recovery Utility can target server clusters running Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

The Server Cluster Recovery Utility is primarily intended for:

- Restoring resource checkpoint files.

- Replacing a failed disk or recovering from disk signature changes.

- Migrating data to a different disk on the shared bus.

To find more information about Clusterrecovery.exe, see Resource Kit Tools Help in the Tools and Settings Collection.

**Confdisk.exe: Disk Configuration Tool**
**Category**
Tool included in the Windows Server 2003 Resource Kit.

**Version Compatibility**
The Disk Configuration Tool can be installed on computers running Windows Server 2003 and Windows XP Professional.

The Disk Configuration Tool can target server clusters running Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

The Disk Configuration Tool restores a disk signature from an Automated System Recovery (ASR) set.

To find more information about Confdisk.exe, see Resource Kit Tools Help in the Tools and Settings Collection.

**Printmig.exe: Print Migrator 3.1 Tool**
**Category**
Tool is available at the Download Center at the Microsoft Web site.

**Version Compatibility**
The tool supports moving printers, including clustered print servers, from Windows NT 4.0 to Windows 2000 or Windows Server 2003 operating systems. The tool also supports the conversion of LPR ports to standard TCP/IP ports.

For more information about print servers and Print Migrator 3.1, see the Microsoft Web site. To download this tool, see Print Migrator 3.1 at the Microsoft Web site.

**Server Cluster WMI Classes**
With Windows Server 2003, you can take advantage of new support for managing server clusters through Windows Management Instrumentation (WMI). You can also continue to use Cluster Automation Server, the cluster management interface that was first available in Windows NT Server 4.0, Enterprise Edition. However, Cluster Automation Server is an interface specific to server clusters. In contrast, the WMI interface has the advantage of being a standardized, current interface in Windows operating systems.

The following tables list and describe the WMI classes that are associated with server clusters. The WMI classes are grouped as follows:

- Classes for the cluster and its elements
- Association classes
- Event classes

**WMI Classes for Server Clusters—Classes for the Cluster and Its Elements**

| Class Name | Name-space | Version Compatibility |
|---|---|---|
| MSCluster_Cluster | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_Network | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_NetworkInterface | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_Node | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_Property | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_Resource | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ResourceGroup | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ResourceType | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_Service | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |

**WMI Classes for Server Clusters—Association Classes**

| Class Name | Name-space | Version Compatibility |
|---|---|---|

| | | |
|---|---|---|
| MSCluster_ClusterToNetwork | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ClusterToNetworkInterface | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ClusterToNode | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ClusterToQuorumResource | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ClusterToResource | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ClusterToResourceGroup | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ClusterToResourceType | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_NetworkToNetworkInterface | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_NodeToActiveGroup | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_NodeToActiveResource | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_NodeToHostedService | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_NodeToNetworkInterface | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ResourceGroupToPreferredNode | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ResourceGroupToResource | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ResourceToDependentResource | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ResourceToPossibleOwner | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_ResourceTypeToResource | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |

**WMI Classes for Server Clusters—Event Classes**

| **Class Name** | **Name-space** | **Version Compatibility** |
|---|---|---|
| MSCluster_Event | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_EventGroupStateChange | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_EventObjectAdd | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |

| | | |
|---|---|---|
| MSCluster_EventObjectRemove | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_EventPropertyChange | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_EventResourceStateChange | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| MSCluster_EventStateChange | MSCluster | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |

For information about Cluster Automation Server and descriptions of many WMI classes, see the SDK documentation on MSDN.

**Related Information**

The following source provides information that is relevant to this section.

For information about command syntax for Cluster.exe and about backing up server clusters, see Help and Support Center in Windows Server 2003, Standard Edition, Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition operating systems.

# Network Load Balancing Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Load Balancing Technical Reference**
Network Load Balancing is a network driver that distributes the load for networked client/server applications across multiple cluster servers. It is part of the Windows scale out functionality and is one of three Windows Clustering technologies. Server clusters, ideal for applications requiring stateful connections, and Component Load Balancing, ideal for stateless COM+ applications, are the other two technologies. Component Load Balancing is a feature of Microsoft Application Center 2000. It is not a feature of the Microsoft Windows Server 2003 family.

Network Load Balancing distributes client requests across a set of servers. It is particularly useful for ensuring that stateless applications, such as a Web server running Internet Information Services (IIS), can be scaled out by adding additional servers as the load increases. Network Load Balancing allows you to easily replace a malfunctioning server or add a new server, which provides scalability.

Network Load Balancing clusters provide scalability for services and applications based on Transmission Control Protocol (TCP) and User Data Protocol (UDP) by combining up to 32 servers running the following operating systems into a single cluster:

- Windows Server 2003, Standard Edition
- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition
- Windows Server 2003, Web Edition

By using Network Load Balancing to build a group of identical, clustered computers, you can enhance the scalability of the following servers over your corporate local area network (LAN):

- Web and File Transfer Protocol (FTP) servers
- Proxy servers and firewall services, such as computers running Internet Security and Acceleration (ISA) Server
- Virtual private network (VPN) servers
- Servers running Windows Media technologies player
- Terminal servers

The sections in this subject include an overview of Network Load Balancing, an in-depth description of how it works, and information about tools and settings.

**In this subject**


- What Is Network Load Balancing?
- How Network Load Balancing Technology Works
- Network Load Balancing Tools and Settings

# What Is Network Load Balancing?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Network Load Balancing?**

**In this section**


- Common Scenarios for Network Load Balancing
- Technologies Related to Network Load Balancing
- Interacting with Other Technologies
- Related Information


For most Information Technology (IT) departments, Internet servers must support applications and services that run 24 hours a day, 7 days a week, such as financial transactions, database access, and corporate intranets. In addition, network applications and servers need the ability to scale performance to handle large volumes of client requests without creating unwanted delays.

Network Load Balancing clusters enable you to manage a group of independent servers as a single system for greater scalability, increased availability, and easier manageability. You can use Network Load Balancing to implement enterprise-wide scalable solutions for the delivery of Transmission Control Protocol/Internet Protocol (TCP/IP) based services and applications.

Network Load Balancing has many advantages over other load balancing solutions that can introduce single points of failure or performance bottlenecks. Because there are no special hardware requirements for Network Load Balancing, you can use any industry standard compatible computer in a Network Load Balancing cluster.

Network Load Balancing works by distributing client requests across a set of servers. It is particularly useful for ensuring that stateless applications, such as Web pages from a server running Internet Information Services (IIS), are highly available and can be scaled out by adding additional servers as the load increases. The ease with which Network Load Balancing allows you to replace a malfunctioning server or add a new server provides scalability.

Network Load Balancing offers the following benefits:

**Scalability**
Network Load Balancing supports up to 32 computers in a single cluster. Hosts can be added and removed without interrupting cluster availability.

**Load balancing**
Load-balances multiple server requests, from either the same client, or from several clients, across multiple hosts in the cluster.

**Increased availability**
Automatically detects and recovers from a failed or offline computer and balances the network load when hosts are added or removed. Recovers and redistributes the workload within seconds.

Many enterprise solutions must address client access to services and applications that are based on connections to selected TCP/IP addresses, protocols, and port numbers. For example, IIS provides service to clients on IP (TCP, 80). If this single IP host were to fail or become overloaded, client access to the service or application may be prevented or fall below a designated performance level.

Configuring multiple hosts to increase scalability and availability for applications and services is one solution. However, this solution may involve specialized network hardware, complex network configuration, and management of individual hosts. For example, multiple hosts functioning as Web servers, each with an individual IP address, could be resolved by multiple entries in round robin DNS. Each server is independent and if a server fails, the static load balancing provided by round robin DNS may prevent clients from accessing their Web application.

To resolve client connection problems, Network Load Balancing allows multiple computers or hosts, configured in a logical group called a network load balancing cluster, to respond to client connection requests made to a single virtual IP address.

**Network Load Balancing Driver**
Network Load Balancing is a driver, Wlbs.sys, which you must load on each host in the cluster. Wlbs.sys includes a statistical

mapping algorithm that the hosts of the cluster collectively use to determine which host handles each incoming request.

You install the driver on each of the cluster hosts, and you configure the cluster to present a virtual IP address to client requests. The client requests go to all of the hosts in the cluster, but only the mapped host accepts and handles the request. All of the other hosts in the cluster drop the request.

**Network Load Balancing Cluster Configuration**
After the driver is installed, it must be configured before the host can join a cluster. You must configure three groups of information about each host: cluster parameters, host parameters, and port rules, before it is possible to create or join a cluster. Configuring the driver allows you to:

- Select the cluster virtual IP address option.
- Customize the cluster according to the various hosts' capacities and sources of client requests.
- Specify that one host handles all of the client requests with the others serving as failover alternatives.
- Divide the load of incoming client requests among the hosts evenly or according to a specified load partitioning weight.

**Common Scenarios for Network Load Balancing**
This section includes scenarios representing common implementations of Network Load Balancing.

**IIS Server (Web Farm)**
An IIS Server Web farm is the most common scenario for Network Load Balancing. Below are two examples of how Network Load Balancing can be used to service a Web farm.

**Servicing Multiple Web Sites**
A Network Load Balancing cluster can be used to host multiple Web sites with different IP addresses. To do this, enter the additional IP addresses as you create the cluster. Additional IP addresses can also be added to an existing cluster.

The virtual cluster feature for Network Load Balancing (available in Windows Server 2003) makes it simpler to service multiple Web sites. Using this feature, you can define port rules that apply to only one of the specific IP addresses that you add to the cluster or to all of the IP addresses.

**Servicing a Web Site with Active Server Pages**
Web sites that use Active Server Pages (ASP) can maintain session state across client connections. Network Load Balancing helps preserve client access to session information by ensuring that all TCP/IP connections from a single client are directed to the same cluster host.

There are, however, situations in which a client can connect with one cluster host, and then have subsequent connections load-balanced to different hosts. Such situations include the following:

- A host is added to the cluster, and Network Load Balancing load-balances new connections from this client to the host.
- Multiple client-side proxy servers cause multiple connections from the same client to originate from different IP addresses.

If either of the preceding situations arises, ASP applications must provide a means to retrieve and manage session state even if a client connects to multiple cluster hosts as part of a single session. The following are two strategies for addressing this issue:

- Use a means at the ASP level, such as a cookie, to retrieve the ASP client state across the Network Load Balancing cluster hosts.
- Encapsulate in a client-side cookie the state from a specific client request. The cookie gives the server the context for subsequent client requests. This solution works only if there is a relatively small amount of state associated with each client transaction. As state grows larger, it becomes increasingly difficult to have the client forward a large cookie to the server with every request.

**Virtual Private Network**
Network Load Balancing can be used with virtual private network (VPN) servers to load-balance PPTP clients. To ensure compatibility with clients running earlier versions of Windows, such as Windows 98 and Windows NT 4.0, it is important to

configure the TCP/IP properties correctly. To do this, assign only a single virtual IP address (the cluster's primary IP address) to the network adapter used by Network Load Balancing. Do not assign a dedicated IP address on any network adapter on this subnet. This restriction does not apply for Windows 2000 or Windows Server 2003 clients. Assigning only a single virtual IP address to the network adapter used by Network Load Balancing ensures that network traffic returning from the host to the client originates from the virtual IP address to which the client sent the request.

**Single-Server Failover Support**
Although you can use Network Load Balancing to provide failover support for applications, using a server cluster is the preferred solution for this scenario. However, if you choose to achieve failover support with Network Load Balancing, this section describes how to do this.

You must first start the application on every host in the cluster. Network Load Balancing does not restart an application on failover. It assumes that an instance of the application is running on each host in the cluster.

In order for Network Load Balancing to provide failover support for a specific application, the files that the application uses must be simultaneously accessible to all hosts that run the application. To achieve this, these files should reside on a back-end file server.

Some applications open shared files only on client request. For these applications, using Network Load Balancing for failover is possible. However, other applications require that these files be continuously open exclusively by one instance of the application. Because this is not possible with Network Load Balancing, you must instead use server clusters to provide failover support for these types of applications.

**Technologies Related to Network Load Balancing**
Network Load Balancing is one of three Windows Clustering technologies. In addition to Network Load Balancing, there are server clusters and Component Load Balancing.

**Server clusters**
Server clusters are used to ensure that stateful applications such as a database, for example, SQL Server, can be kept highly available by failing over the application from one server to another in the event of a failure. Multiple servers (nodes) in a cluster remain in constant communication. If one of the nodes in a cluster is unavailable as a result of failure or maintenance, another node immediately begins providing service (a process known as failover). Users who access the cluster are constantly connected to server-based resources.

With Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure server clusters in one of three available cluster models, depending on your requirements.

Network Load Balancing load-balances requests between front-end Web servers and server clusters does the same thing for back-end database access. Network Load Balancing and server clusters cannot both be active on the same computer since both technologies control and configure network adapters and they may interfere with one another. By joining Network Load Balancing and server clusters together, an overall clustering scheme is created, as shown in the following diagram:

**Integrated Clustering Scheme Using Network Load Balancing and Server Clusters**




**Component Load Balancing**
Component Load Balancing is the dynamic load balancing feature of COM+. It enables the creation of 8-node application clusters behind a Component Load Balancing router, enabling the COM+ applications to scale while also ensuring their constant availability. Component Load Balancing is ideal for stateless COM+ applications. Component Load Balancing is a feature of Microsoft Application Center 2000. It is not a feature of the Windows Server 2003 family.

**Interacting with Other Technologies**
Network Load Balancing interacts with the following technologies:

**Terminal server clusters**
Network Load Balancing can be used to load-balance terminal server clusters. Load balancing pools the processing resources of several servers using the TCP/IP networking protocol. You can use this service with a cluster of terminal servers to scale the performance of a single terminal server by distributing sessions across multiple servers. Session Directory (included in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition) keeps track of disconnected sessions on the cluster, and ensures that users are reconnected to those sessions.

**DNS resolution**
Network Load Balancing does not support dynamic Domain Name System (DNS) resolution, where the name of the cluster is automatically registered by the host when the host starts. This functionality must be disabled on the Network Load Balancing interface for both DNS and Windows Internet Name Service (WINS); otherwise, each host's computer name will be registered with the cluster IP address. When using Network Load Balancing with DNS, you will need to directly configure the DNS server to register the name. Also, Dynamic Host Configuration Protocol (DHCP) should not be used on the network interface to which Network Load Balancing is bound; however, DHCP can be used on other interfaces.

**L2TP/IPSec**
Network Load Balancing in Windows Server 2003 supports both Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) virtual private network (VPN) sessions. However, Network Load Balancing must be configured in single affinity mode.

**Kerberos authentication**
Network Load Balancing supports Kerberos with applications it load-balances.

**.NET Framework remoting**
Network Load Balancing supports .NET Framework remoting, which uses method invocations from client to server over a single TCP connection. This means that once a connection is established, it is reused for subsequent method invocations and is closed only after the connection remains idle for a pre-configured amount of time. Network Load Balancing can load-balance these connections, but load balancing will likely be uneven because it is the TCP connection that is load-balanced and not the method invocation. Since one client gets "pinned" to a specific server, the load will be well distributed only if you have many clients connected to the cluster at the same time. Each client will get load-balanced, but the connection will stay open for a long period of time.

**Related Information**
The following resource contains additional information that is relevant to this section.

- Clustering Technologies

# How Network Load Balancing Technology Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Network Load Balancing Terms and Definitions
- Network Load Balancing Architecture
- Network Load Balancing Protocols
- Application Compatibility with Network Load Balancing
- Network Load Balancing for Scalability
- Scaling Network Load Balancing Clusters
- Performance Impact of Network Load Balancing


When Network Load Balancing is installed as a network driver on each of the member servers, or hosts, in a cluster, the cluster presents a virtual Internet Protocol (IP) address to client requests. The client requests go to all the hosts in the cluster, but only the host to which a given client request is mapped accepts and handles the request. All the other hosts drop the request. Depending on the configuration of each host in the cluster, the statistical mapping algorithm, which is present on all the cluster hosts, maps the client requests to particular hosts for processing.

**Basic Diagram for Network Load Balancing Clusters**

The following figure shows two connected Network Load Balancing clusters. The first cluster is a firewall cluster with two hosts and the second cluster is a Web server cluster with four hosts.

**Two Connected Network Load Balancing Clusters**




**Network Load Balancing Terms and Definitions**

Before you review the Network Load Balancing components and processes, it is helpful to understand specific terminology. The following terms are used to describe the components and processes of Network Load Balancing.

**affinity**

For Network Load Balancing, the method used to associate client requests to cluster hosts. When no affinity is specified, all network requests are load-balanced across the cluster without respect to their source. Affinity is implemented by directing all client requests from the same IP address to the same cluster host.

**convergence**

The process of stabilizing a system after changes occur in the network. For routing, if a route becomes unavailable, routers send update messages throughout the network, reestablishing information about preferred routes.

For Network Load Balancing, a process by which hosts exchange messages to determine a new, consistent state of the cluster and to elect the default host. During convergence, a new load distribution is determined for hosts that share the handling of

network traffic for specific Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports.

**dedicated IP address**
The IP address of a Network Load Balancing host used for network traffic that is not associated with the Network Load Balancing cluster (for example, Telnet access to a specific host within the cluster). This IP address is used to individually address each host in the cluster and therefore is unique for each host.

**default host**
The host with the highest host priority for which a **drainstop** command is not in progress. After convergence, the default host handles all of the network traffic for Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports that are not otherwise covered by port rules.

**failover**
In server clusters, the process of taking resource groups offline on one node and bringing them online on another node. When failover occurs, all resources within a resource group fail over in a predefined order; resources that depend on other resources are taken offline before, and are brought back online after, the resources on which they depend.

**heartbeat**
A message that is sent at regular intervals by one computer on a Network Load Balancing cluster or server cluster to another computer within the cluster to detect communication failures.

Network Load Balancing initiates convergence when it fails to receive heartbeat messages from another host or when it receives a heartbeat message from a new host.

**multicast media access control (MAC) address**
A type of media access control address used by multiple, networked computers to receive the same incoming network frames concurrently. Network Load Balancing uses multicast MAC addresses to efficiently distribute incoming network traffic to cluster hosts.

**multihomed computer**
A computer that has multiple network adapters or that has been configured with multiple IP addresses for a single network adapter.

Network Load Balancing supports multihomed servers by allowing multiple virtual IP addresses to be assigned to the cluster adapter.

**throughput**
A performance measure defined as the number of client requests processed by a Network Load Balancing cluster per unit of time.

**virtual cluster**
A Network Load Balancing cluster that you create by assigning specific port rules to specific virtual IP addresses. With virtual clusters, you can use different port rules for different Web sites or applications hosted on the cluster, provided each Web site or application has a different virtual IP address.

**virtual IP address**
An IP address that is shared among the hosts of a Network Load Balancing cluster. A Network Load Balancing cluster might also use multiple virtual IP addresses, for example, in a cluster of multihomed Web servers.

**Network Load Balancing Architecture**
Network Load Balancing runs as a network driver logically beneath higher-level application protocols, such as HTTP and FTP. On each cluster host, the driver acts as a filter between the network adapter's driver and the TCP/IP stack, allowing a portion of the incoming network traffic to be received by the host. This is how incoming client requests are partitioned and load-balanced among the cluster hosts. To maximize throughput and availability, Network Load Balancing uses a fully distributed software architecture, and an identical copy of the Network Load Balancing driver that runs in parallel on each cluster host. The figure below shows the implementation of Network Load Balancing as an intermediate driver in the Windows Server 2003 network stack.

### Typical Configuration of a Network Load Balancing Host




The following table describes the principal components related to the Network Load Balancing architecture.

**Components of the Network Load Balancing Architecture**

| Component | Description |
|---|---|
| Nlb.exe | The Network Load Balancing control program. You can use Nlb.exe from the command line to start, stop, and administer Network Load Balancing, as well as to enable and disable ports and to query cluster status. |
| Nlbmgr.exe | The Network Load Balancing Manager control program. Use this command to start Network Load Balancing Manager. |
| Wlbs.exe | The former Network Load Balancing control program. This has been replaced by Nlb.exe. However, you can still use Wlbs.exe rather than Nlb.exe if necessary, for example, if you have existing scripts that reference Wlbs.exe. |
| Wlbsprov.dll | The Network Load Balancing WMI provider. |
| Nlbmprov.dll | The Network Load Balancing Manager WMI provider. |
| Wlbsctrl.dll | The Network Load Balancing API DLL. |
| Wlbs.sys | The Network Load Balancing device driver. Wlbs.sys is loaded onto each host in the cluster and includes the statistical mapping algorithm that the cluster hosts collectively use to determine which host handles each incoming request. |

**Application and Service Environment**

When a Web server application maintains state information about a client session across multiple TCP connections, it is important that all TCP connections for the client are directed to the same cluster host.

Network Load Balancing can load-balance any application or service that uses TCP/IP as its network protocol and is associated with a specific TCP or UDP port. The distributed algorithm, which is used to determine which host responds to a TCP connection request or incoming UDP packet, can include the port rule in the decision. Including the port rule in the decision means that for any client, different members of the Network Load Balancing cluster may service connection requests or packets addressed to different port rule on the virtual IP address.

**Note**

- While configuring a Network Load Balancing cluster, consider the type of application or service that the server is providing, and then select the appropriate configuration for Network Load Balancing hosts.

**Client State**

A discussion of Network Load Balancing clusters requires clarification of two kinds of client states, application data state and session state, because they are essential to configuring Network Load Balancing.

**Application data state**

In terms of application data, you must consider whether the server application makes changes to the data store and whether the changes are synchronized across instances of the application (the instances that are running on the Network Load Balancing cluster hosts). An example of an application that does not make changes to the data store is a static Web page supported by an IIS server.

Means must be provided to synchronize updates to data state that need to be shared across servers. One such means is use of

a back-end database server that is shared by all instances of the application. An example would be an Active Server Pages (ASP) page that is supported by an IIS server and that can access a shared back-end database server, such as a SQL Server.

**Session state**
Session state (or intraclient state) refers to client data that is visible to a particular client for the duration of a session. Session state can span multiple TCP connections, which can be either simultaneous or sequential. Network Load Balancing assists in preserving session state through client affinity settings. These settings direct all TCP connections from a given client address or class of client addresses to the same cluster host. This allows session state to be maintained by the server application in the host memory.

Client/server applications that embed session state within "cookies" or push it to a back-end database do not need client affinity to be maintained.

An example of an application that requires maintaining session state is an e-commerce application that maintains a shopping cart for each client.

**Network Load Balancing Parameters**
By setting port rules, cluster parameters, and host parameters, you gain flexibility in configuring the cluster, which enables you to customize the cluster according to the various hosts' capacities and sources of client requests. Cluster parameters apply to the entire cluster, while the host parameters apply to a specific host.

**Port rules**
The Network Load Balancing driver uses port rules that describe which traffic to load-balance and which traffic to ignore. By default, the Network Load Balancing driver configures all ports for load balancing. You can modify the configuration of the Network Load Balancing driver that determines how incoming network traffic is load-balanced on a per-port basis by creating port rules for each group of ports or individual ports as required. Each port rule configures load balancing for client requests that use the port or ports covered by the port range parameter. How you load-balance your applications is mostly defined by how you add or modify port rules, which you create on each host for any particular port range.

**Affinity**
Affinity is the method used to associate client requests to cluster hosts. Network Load Balancing assists in preserving session state through client affinity settings for each port rule that Network Load Balancing creates. These settings direct all TCP connections from a given client address or class of client addresses to the same cluster host. Directing the connections to the same cluster host allows the server applications in the designated host memory to correctly maintain the session state.

**Running Network Load Balancing in an Optimal Environment**
The sections in this document provide in-depth information about how Network Load Balancing works in an optimal environment. An optimal environment for Network Load Balancing is defined as follows:

- Two or more network adapters in each cluster host are used.
- The Transmission Control Protocol/Internet Protocol (TCP/IP) network protocol is the only protocol used on the cluster adapter. Other protocols, such as Internetwork Packet Exchange (IPX), should not be added to this adapter.
- Within a given cluster, all cluster hosts must operate in either unicast or multicast mode, but not both.
- Network Load Balancing should not be enabled on a computer that is part of a server cluster.
- All hosts in a cluster must belong to the same subnet and the cluster's clients must be able to access this subnet.
- The appropriate hardware resources should be used. The goal in tuning Network Load Balancing and the applications it load-balances is to determine which hardware resource will experience the greatest demand, and then to adjust the configuration to relieve that demand and maximize total throughput.
- Cluster parameters, port rules, and host parameters are correctly configured:
    - Cluster parameters and port rules for each unique virtual IP address are identical across all hosts. Each unique virtual IP address must be configured with the same port rules across all hosts that service that virtual IP address. However, if you have multiple virtual IP addresses configured on a host, each of those virtual IP addresses can have a different set of port rules.
    - Port rules are set for all ports used by the load-balanced application. For example, FTP uses port 20, port 21, and

ports 1024–65535.

    - The dedicated IP address is unique and the cluster IP address is added to each cluster host.

- Network performance should be optimized by limiting switch port flooding.

- Windows Internet Name Service (WINS), Dynamic Host Configuration Protocol (DHCP), and Domain Name System (DNS) services can be run on Network Load Balancing cluster hosts; however, the interface to which Network Load Balancing is bound cannot use a DHCP address.

**Network Load Balancing Driver**
The Network Load Balancing driver is installed on a computer configured with Transmission Control Protocol/Internet Protocol (TCP/IP) and is bound to a single network interface called the cluster adapter. The driver is configured with a single IP address, the cluster primary IP address, on a single subnet for all of the member servers (or hosts) on the cluster. Each host has an identical media access control (MAC) address that allows the hosts to concurrently receive incoming network traffic for the cluster's primary IP address (and for additional IP addresses on multihomed hosts).

Incoming client requests are partitioned and load-balanced among the cluster hosts by the Network Load Balancing driver, which acts as a rule-based filter between the network adapter's driver and the TCP/IP stack. Each host receives a designated portion of the incoming network traffic.

**Distributed Architecture**
Network Load Balancing is a distributed architecture, with an instance of the driver installed on each cluster host. Throughput is maximized to all cluster hosts by eliminating the need to route incoming packets to individual cluster hosts, through a process called filtering. Filtering out unwanted packets in each host improves throughput; this process is faster than routing packets, which involves receiving, examining, rewriting, and resending the packets.

Another key advantage to the fully distributed architecture of Network Load Balancing is the enhanced availability resulting from (n-1) way failover in a cluster with n hosts. In contrast, dispatcher-based solutions create an inherent single point of failure that you must eliminate by using a redundant dispatcher that provides only one-way failover. Dispatcher-based solutions offer a less robust failover solution than does a fully distributed architecture.

**Load Balancing Algorithm**
The Network Load Balancing driver uses a fully distributed filtering algorithm to statistically map incoming client requests to the cluster hosts, based upon their IP address, port, and other information.

When receiving an incoming packet, all hosts within the cluster simultaneously perform this mapping to determine which host should handle the packet. Those hosts not required to service the packet simply discard it. The mapping remains constant unless the number of cluster hosts changes or the filter processing rules change.

The filtering algorithm is much more efficient in its packet handling than centralized load balancers, which must modify and retransmit packets. Efficient packet handling allows for a much higher aggregate bandwidth to be achieved on industry standard hardware.

The distribution of client requests that the statistical mapping function effects is influenced by the following:

- Host priorities

- Multicast or unicast mode

- Port rules

- Affinity

- Load percentage distribution

- Client IP address

- Client port number

- Other internal load information

The statistical mapping function does not change the existing distribution of requests unless the membership of the cluster changes or you adjust the load percentage.

**Unicast and Multicast Modes**
Network Load Balancing requires at least one network adapter; and different hosts in a cluster can have a different number of adapters, but all must use the same network IP transmission mode, either unicast or multicast.

Unicast mode is the default mode, but you can configure the Network Load Balancing driver to operate in either mode.

When you enable unicast support, the unicast mode changes the cluster adapter's MAC address to the cluster MAC address. This cluster address is the same MAC address that is used on all cluster hosts. When this change is made, clients can no longer address the cluster adapters by their original MAC addresses.

When you enable multicast support, Network Load Balancing adds a multicast MAC access to the cluster adapters on all of the cluster hosts. At the same time, the cluster adapters retain their original MAC addresses.

### Note

- The Network Load Balancing driver does not support a mixed unicast and multicast environment. All cluster hosts must be either multicast or unicast; otherwise, the cluster will not function properly.

If clients are accessing a Network Load Balancing cluster through a router when the cluster has been configured to operate in multicast mode, be sure that the router meets the following requirements:

- Accepts an Address Resolution Protocol (ARP) reply that has one MAC address in the payload of the ARP structure but appears to arrive from a station with another MAC address, as judged by the Ethernet header.

- Accepts an ARP reply that has a multicast MAC address in the payload of the ARP structure.

If your router does not meet these requirements, you can create a static ARP entry in the router. For example, some routers require a static ARP entry because they do not support the resolution of unicast IP addresses to multicast MAC addresses.

**Subnet and Network Considerations**
The Network Load Balancing architecture with a single MAC address for all cluster hosts maximizes use of the subnet's hub and/or switch architecture to simultaneously deliver incoming network traffic to all cluster hosts.

Your network configuration will typically include routers, but may also include layer 2 switches (collapsed backbone) rather than the simpler hubs or repeaters that are available. Cluster configuration, when using hubs, is predictable because the hubs distribute IP traffic to all ports.

### Note

- If client-side network connections at the switch are significantly faster than server-side connections, incoming traffic can occupy a prohibitively large portion of server-side port bandwidth.

**Network Adapters**
Network Load Balancing requires only a single network adapter, but for optimum cluster performance, you should install a second network adapter on each Network Load Balancing host. In this configuration, one network adapter handles the network traffic that is addressed to the server as part of the cluster. The other network adapter carries all of the network traffic that is destined to the server as an individual computer on the network, including cluster communication between hosts.

### Note

- Network Load Balancing with a single network adapter can provide full functionality if you enable multicast support for this adapter.

**Selecting an IP Transmission Mode**
When you are implementing a Network Load Balancing solution, the Internet Protocol transmission mode that is selected and the number of network adapters that are required are dependent upon the following network requirements:

- Layer 2 switches or hubs

- Peer-to-peer communication between hosts
- Maximized communication performance

For example, a cluster supporting a static Hypertext Markup Language (HTML) Web application can have a requirement to synchronize the Web site copies of a large number of cluster hosts. This scenario requires interhost peer-to-peer communications. You select the number of network adapters and the IP communications mode to meet this requirement.

There is no restriction on the number of network adapters, and different hosts can have a different number of adapters. You can configure Network Load Balancing to use one of four different models.

**Single Network Adapter in Unicast Mode**
The single network adapter in unicast mode is suitable for a cluster in which you do not require ordinary network communication among cluster hosts, and in which there is limited dedicated traffic from outside the cluster subnet to specific cluster hosts. In this model, the computer can also handle traffic from inside the subnet if the IP datagrams do not carry the same MAC address as the cluster adapter.

**Single Network Adapter in Multicast Mode**
This model is suitable for a cluster in which ordinary network communication among cluster hosts is necessary or desirable, but in which there is limited dedicated traffic from outside the cluster subnet to specific cluster hosts.

**Multiple Network Adapter in Unicast Mode**
This model is suitable for a cluster in which ordinary network communication among cluster hosts is necessary or desirable, and in which there is comparatively heavy dedicated traffic from outside the cluster subnet to specific cluster hosts.

This mode is the preferred configuration used by most sites because a second network adapter may enhance overall network performance.

**Multiple Network Adapter in Multicast Mode**
This model is suitable for a cluster in which ordinary network communication among cluster hosts is necessary, and in which there is heavy dedicated traffic from outside the cluster subnet to specific cluster hosts.

**Comparison of Modes**
The advantages and disadvantages of each model are listed in the following table.

### Components of the Network Load Balancing Architecture

| Adapter | Mode | Advantages | Disadvantages |
|---|---|---|---|
| Single | Unicast | Simple configuration | Poor overall performance |
| Single | Multicast | Medium performance | Complex configuration |
| Multiple | Unicast | Best balance | None |
| Multiple | Multicast | Best balance | Complex configuration |

**Network Load Balancing Addressing**
The Network Load Balancing cluster is assigned a primary Internet Protocol (IP) address. This IP address represents a virtual IP address to which all of the cluster hosts respond, and the remote control program that is provided with Network Load Balancing uses this IP address to identify a target cluster.

**Primary IP address**
The primary IP address is the virtual IP address of the cluster and must be set identically for all hosts in the cluster. You can use the virtual IP address to address the cluster as a whole. The virtual IP address is also associated with the Internet name that you specify for the cluster.

**Dedicated IP address**
You can also assign each cluster host a dedicated IP address for network traffic that is designated for that particular host only. Network Load Balancing never load-balances the traffic for the dedicated IP addresses, it only load-balances incoming traffic from all IP addresses other than the dedicated IP address.

The following figure shows how IP addresses are used to respond to client requests.

**Network Load Balancing Cluster**




**Distribution of Cluster Traffic**

When the virtual IP address is resolved to the station address (MAC address), this MAC address is common for all hosts in the cluster. You can enable client connections to only the required cluster host when more packets are sent. The responding host then substitutes a different MAC address for the inbound MAC address in the reply traffic. The substitute MAC address is referred to as the Source MAC address. The following table shows the MAC addresses that will be generated for a cluster adapter.

| IP Mode | MAC Address | Explanation |
|---|---|---|
| Unicast inbound | 02-BF-W-X-Y-Z | W-X-Y-Z = IP address<br><br>Onboard MAC disabled |
| Multicast inbound | 03-BF-W-X-Y-Z | W-X-Y-Z = IP address<br><br>Onboard MAC enabled |
| Source outbound | 02-P-W-X-Y-A | W-X-Y-Z = IP address<br><br>P = Host priority |

In the unicast mode of operation, the Network Load Balancing driver disables the onboard MAC address for the cluster adapter. You cannot use the dedicated IP address for interhost communications because all of the hosts have the same MAC address.

In multicast mode of operation, the Network Load Balancing driver supports both the onboard and the multicast address. If your cluster configuration will require connections from one cluster host to another, for example, when making a NetBIOS connection to copy files, use multicast mode or install a second network adapter.

If the cluster hosts were attached to a switch instead of a hub, the use of a common MAC address would create a conflict because Layer 2 switches expect to see unique source MAC addresses on all switch ports. To avoid this problem, Network Load Balancing uniquely modifies the source MAC address for outgoing packets, for example, a cluster MAC address of 02-BF-1-2-3-4 is set to 02-p-1-2-3-4, where p is the host's priority within the cluster.

This technique prevents the switch from learning the cluster's inbound MAC address, and as a result, incoming packets for the cluster are delivered to all of the switch ports. If the cluster hosts are connected to a hub instead of to a switch, you can disable masking of the source MAC address in unicast mode to avoid flooding upstream switches. You disable Network Load Balancing by setting the Network Load Balancing registry parameter MaskSourceMAC to 0. The use of an upstream level three switch will also limit switch flooding.

The unicast mode of Network Load Balancing induces switch flooding to simultaneously deliver incoming network traffic to all of the cluster hosts. Also, when Network Load Balancing uses multicast mode, switches often flood all of the ports by default to deliver multicast traffic. However, the multicast mode of Network Load Balancing gives the system administrator the opportunity to limit switch flooding by configuring a virtual LAN within the switch for the ports corresponding to the cluster hosts.

**Port Rules**

Port rules are created for individual ports and for groups of ports that Network Load Balancing requires for particular applications and services. The filter setting then defines whether the Network Load Balancing driver will pass or block the traffic.

The Network Load Balancing driver controls the distribution and partitioning of Transmission Control Protocol (TCP) and User

Datagram Protocol (UDP) traffic from connecting clients to selected hosts within a cluster by passing or blocking the incoming data stream for each host. Network Load Balancing does not control any incoming IP traffic other than TCP and UDP for ports that a port rule specifies.

You can add port rules or update parameters by taking each host out of the cluster in turn, updating its parameters, and then returning it to the cluster. The host joining the cluster handles no traffic until convergence is complete. The cluster does not converge to a consistent state until all of the hosts have the same number of rules. For example, if a rule is added, it does not take effect until you have updated all of the hosts have been updated and they have rejoined the cluster.

### Note

- Internet Group Membership Protocol (IGMP), ARP, or other IP protocols are passed unchanged to the TCP/IP protocol software on all of the hosts within the cluster.

Port rules define individual ports or groups of ports for which the driver has a defined action. You need to consider certain parameters when creating the port rules, such as the:

- TCP or UDP port range for which you should apply this rule.
- Protocols for which this rule should apply (TCP, UDP, or both).
- Filtering mode chosen: multiple hosts, single host, or disabled.

When defining the port rules, it is important that the rules be exactly the same for each host in the cluster because if a host attempts to join the cluster with a different number of rules from the other hosts, the cluster will not converge. The rules that you enter for each host in the cluster must have matching port ranges, protocol types, and filtering modes.

**Filtering Modes**
The filter defines for each port rule whether the incoming traffic is discarded, handled by only one host, or distributed across multiple hosts. The three possible filtering modes that you can apply to a port rule are:

**Multiple hosts**
With this filtering mode, clusters can equally distribute the load among the hosts or each host can handle a specified load weight.

**Single host**
This filtering mode provides fault tolerance for the handling of network traffic with the target host defined by its priority.

**Disabled**
This filtering mode lets you build a firewall against unwanted network access to a specific range of ports; the driver discards the unwanted packets.

**Load Weighting**
When the filter mode for a port rule is set to multiple hosts, the Load Weight parameter specifies the percentage of load-balanced network traffic that this host should handle for the associated rule. Allowed value ranges are from 0 to 100.

### Note

- To prevent a host from handling any network traffic for a port rule, set the load weight to 0.

Because hosts can dynamically enter or leave the cluster, the sum of the load weights for all cluster hosts does not have to equal 100. The percentage of host traffic is computed as the local load percentage value divided by the load weight sum across the cluster.

If you balance the load evenly across all of the hosts with this port rule, you can specify an equal load distribution parameter instead of specifying a load weight parameter.

**Priority**
When the filter mode for a port rule is set to single, the priority parameter specifies the local host's network traffic for the associated port rule. The host with the highest handling priority for this rule among the current cluster members will handle all of the traffic.

The allowed values range from 1, the highest priority, to the maximum number of hosts allowed, 32. This value must be unique

for all hosts in the cluster.

**Supporting Multiple Client Connections**
In a load-balanced multiserver environment, managing and resolving client, application, and session state for individual clients can be complex. By default, in a Network Load Balancing solution, different hosts in the cluster can service multiple client connections.

When a client creates an initial connection to a host in the cluster, the application running on this host holds the client state. If the same host does not service subsequent connections from the client, errors can occur if the application instances do not share the client state between hosts.

For example, application development for an ASP-based Web site can be more difficult if the application must share the client state among the multiple hosts in the cluster. If all of the client connections can be guaranteed to go to the same server, you can solve the difficulties with the application that is not sharing the client state among host instances.

By using a Network Load Balancing feature called affinity, you can ensure that the same cluster host handles all of the TCP connections from one client IP address. Affinity allows you to scale applications that manage session state spanning multiple client connections. In a Network Load Balancing cluster, with affinity enabled, initial client connection requests are distributed according to the cluster configuration, but after you have established the initial client request the same host will service all of the subsequent requests from that client.

**Affinity**
Clients can have many TCP connections to a Network Load Balancing cluster; the load-balancing algorithm will potentially distribute these connections across multiple hosts in the cluster.

If server applications have client or connection state information, this state information must be made available on all of the cluster hosts to prevent errors. If you cannot make state information available on all of the cluster hosts, you cannot use client affinity to direct all of the TCP connections from one client IP address to the same cluster host. Directing TCP connections from the IP address to the same host allows an application to maintain state information in the host memory.

For example, if a server application (such as a Web server) maintains state information about a client's site navigation status that spans multiple TCP connections, it is critical that all of the TCP connections for this client state information be directed to the same cluster host to prevent errors.

Affinity defines a relationship between client requests from a single client address or from a Class C network of clients (where IP addresses range from 192.0.0.1 to 223.255.255.254) and one of the cluster hosts. Affinity ensures that requests from the specified clients are always handled by the same host. The relationship lasts until convergence occurs (namely, until the membership of the cluster changes) or until you change the affinity setting. There is no time-out — the relationship is based only on the client IP address.

You can distribute incoming client connections based on the algorithm as determined by the following client affinity settings:

**No Affinity**
This setting distributes client requests more evenly, when maintaining session state is not an issue, you can use this setting to speed up response time to requests. For example, because multiple requests from a particular client can go to more than one cluster host, clients that access Web pages can get different parts of a page or different pages from different hosts. This setting is used for most applications.

With this setting, the Network Load Balancing statistical mapping algorithm uses both the port number and entire IP address of the client to influence the distribution of client requests.

**Single Affinity**
When single affinity is used, the entire source IP address (but not the port number) is used to determine the distribution of client requests.

You typically set single affinity for intranet sites that need to maintain session state. This setting always returns each client's traffic to the same server, thus assisting the application in maintaining client sessions and their associated session state.

Note that client sessions that span multiple TCP connections (such as ASP sessions) are maintained as long as the Network Load Balancing cluster membership does not change. If the membership changes by adding a new host, the distribution of client requests is recomputed, and you cannot depend on new TCP connections from existing client sessions ending up at the same server. If a host leaves the cluster, its clients are partitioned among the remaining cluster hosts when convergence completes, and other clients are unaffected.

**Class C Affinity**
When Class C affinity is used, only the upper 24 bits of the client's IP address are used by the statistical-mapping algorithm. A Class C unicast IP address ranges from 192.0.0.1 to 223.255.255.254. The first three octets indicate the network, and the last

octet indicates the host on the network. Network Load Balancing provides optional session support for Class C IP addresses (in addition to support for single IP addresses) to accommodate clients that make use of multiple proxy servers at the client site. Class-based IP addressing has been superceded by Classless Interdomain Routing (CIDR). This option is appropriate for server farms that serve the Internet. Client requests coming over the Internet might come from clients sitting behind proxy farms. In this case, during a single client session, client requests can come into the Network Load Balancing cluster from several source IP addresses during a session.

Class C affinity addresses this issue by directing all the client requests from a particular Class C network to a single Network Load Balancing host.

There is no guarantee, however, that all of the servers in a proxy farm are on the same Class C network. If the client's proxy servers are on different Class C networks, then the affinity relationship between a client and the server ends when the client sends successive requests from different Class C network addresses.

**Heartbeats and Convergence**

Network Load Balancing cluster hosts exchange heartbeat messages to maintain consistent data about the cluster's membership. By default, when a host fails to send out heartbeat messages within five seconds, it is deemed to have failed. Once a host has failed, the remaining hosts in the cluster perform convergence and do the following:

- Establish which hosts are still active members of the cluster.
- Elect the host with the highest priority as the new default host.
- Ensure that all new client requests are handled by the surviving hosts.

In convergence, surviving hosts look for consistent heartbeats. If the host that failed to send heartbeats once again provides heartbeats consistently, it rejoins the cluster in the course of convergence. When a new host attempts to join the cluster, it sends heartbeat messages that also trigger convergence. After all cluster hosts agree on the current cluster membership, the client load is redistributed to the remaining hosts, and convergence completes.

The following figure shows how the client load is evenly distributed among four cluster hosts before convergence takes place:

**Network Load Balancing Cluster Before Convergence**




The following figure shows a failed host and how the client load is redistributed among the three remaining hosts after convergence.

**Network Load Balancing Cluster After Convergence**




Convergence generally only takes a few seconds, so interruption in client service by the cluster is minimal. During

convergence, hosts that are still active continue handling client requests without affecting existing connections. Convergence ends when all hosts report a consistent view of the cluster membership and distribution map for several heartbeat periods.

By editing the registry, you can change both the number of missed messages required to start convergence and the period between heartbeats. However, be aware that making the period between heartbeats too short increases network overhead on the system. Also be aware that reducing the number of missed messages increases the risk of erroneous host evictions from the cluster.

**Note**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Network Load Balancing Protocols**

Network Load Balancing can load-balance any application or service that uses TCP/IP as its network protocol and is associated with a specific Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port. For applications to work with Network Load Balancing, they must use TCP connections or UDP data streams.

You can also enable Internet Group Management Protocol (IGMP) support on the cluster hosts to control switch flooding when operating in multicast mode. IGMP is a protocol used by Internet Protocol version 4 (IPv4) hosts to report their multicast group memberships to any immediately neighboring multicast routers.

**Application Compatibility with Network Load Balancing**

In general, Network Load Balancing can load-balance any application or service that uses TCP/IP as its network protocol and is associated with a specific TCP or UDP port.

**Compatible Applications**

Compatible applications must have the following requirements in order to work with Network Load Balancing:

- TCP connections or UDP data streams.

- Synchronization of data. If client data that resides on one of the hosts changes as a result of a client transaction, applications must provide a means of synchronizing data updates if that data is shared on multiple instances across the cluster.

- Single or Class C affinity. If session state is an issue, applications must use single or Class C affinity or provide a means (such as a client cookie or reference to a back-end database) of maintaining session state in order to be uniformly accessible across the cluster.

As highlighted in the previous two bullets, before using Network Load Balancing, you must consider an application's requirements regarding client state. Network Load balancing can accommodate both stateless and stateful connections. However, certain factors must be considered for a stateful connection.

Application servers maintain two kinds of stateful connections:

- Interclient state. A state whose data updates must be synchronized with transactions that are performed for other clients, such as merchandise inventory at an e-commerce site. Because any host within the cluster can service client traffic, information that must be persistent across multiple requests or that must be shared among clients needs to be shared in a location that is accessible from all cluster hosts. Updates to information that is shared among the hosts needs to be synchronized, for example by using a back-end database server.

- Intraclient state. A state that must be maintained for a given client throughout a session (that can span multiple connections), such as a shopping cart process at an e-commerce site.

Network Load Balancing can be used with applications that maintain either type of state, although applications requiring interclient state are typically best accommodated with server clusters rather than Network Load Balancing. While intraclient state can be accommodated through the Network Load Balancing affinity settings, interclient state, on the other hand, requires additional components outside of Network Load Balancing.

**Interclient State**

Some applications require interclient state, that is, the applications make changes to data that must be synchronized across all

instances of the application. For example, any type of application that uses Microsoft SQL Server for data storage, such as a user registration database, requires interclient state. When applications do share and change data (require interclient state), the changes need to be properly synchronized. Each host can use local, independent copies of databases that are merged offline as necessary. Alternatively, the clustered hosts can share access to a separate, networked database server.

A combination of these approaches can also be used. For example, static Web pages can be replicated among all clustered servers to ensure fast access and complete fault tolerance. However, database requests would be forwarded to a common database server that handles updates for the multiple cluster hosts.

Network Load Balancing is a viable solution for applications that require interclient state.

**Intraclient State**
A second kind of state intraclient state (also known as session state), refers to client data that is visible to a particular client for the duration of a session. Session state can span multiple TCP connections, which can be either simultaneous or sequential. Network Load Balancing assists in preserving session state through client affinity settings. These settings direct all TCP connections from a given client address or group of client addresses to the same cluster host. This allows session state to be maintained by the server application in the host memory.

**Incompatible Applications**
Applications that are not compatible with Network Load Balancing have one or more of the following characteristics:

- They bind to actual computer names (examples of such applications are Exchange Server and Distributed File System).
- They have files that must be continuously open for writing (for example, Exchange Server). In a Network Load Balancing cluster, multiple instances of an application (on different cluster hosts) should not have a file simultaneously opened for writing unless the application was designed to synchronize file updates. This is generally not the case.

**Network Load Balancing for Scalability**
Network Load Balancing scales the performance of a server-based program, such as a Web server, by distributing its client requests across multiple identical servers within the cluster; you can add more servers to the cluster as traffic increases. Up to 32 servers are possible in any one cluster. The following figure represents how additional servers can support more users.

**How Servers Scale Out Across a Cluster**




You can improve the performance of each individual host in a cluster by adding more or faster CPUs, network adapters and disks, and in some cases by adding more memory. These additions to the Network Load Balancing cluster is called scaling up, and requires more intervention and careful planning than scaling out. Limitations of applications or the operating system configuration could mean that scaling up by adding more memory may not be as appropriate as scaling out.

You can handle additional IP traffic by simply adding computers to the Network Load Balancing cluster as necessary. Load balancing, in conjunction with the use of server clusters, is part of a scaling approach referred to as scaling out. The greater the number of computers involved in the load-balancing scenario, the higher the throughput of the overall server cluster.

**Scaling Network Load Balancing Clusters**
Network Load Balancing clusters have a maximum of 32 hosts, and all of the hosts must be on the same subnet. If a cluster

cannot meet the performance requirements of a clustered application, such as a Web site, because of a host count or subnet throughput limitation, then you can use multiple clusters to scale out further.

Combining round robin DNS and Network Load Balancing results in a very scalable and highly available configuration. Configuring multiple Network Load Balancing clusters on different subnets and configuring DNS to sequentially distribute requests across multiple Network Load Balancing clusters can evenly distribute the client load that is distributed across several clusters. When multiple Network Load Balancing Web server clusters are configured with round robin DNS, the Web servers are made resilient to networking infrastructure failures.

When you use round robin DNS in conjunction with Network Load Balancing clusters, each cluster is identified in DNS by the cluster virtual IP. Because each cluster is automatically capable of both load balancing and fault tolerance, each DNS-issued IP address will function until all hosts in that particular cluster fail. Round robin DNS enables only a limited form of TCP/IP load balancing for IP-based servers when used without Network Load Balancing. When used with multiple individual hosts, such as Web servers, round robin DNS does not function effectively as a solution. If a host fails, round robin DNS continues to route requests to the failed server until the server is removed from DNS.

**Performance Impact of Network Load Balancing**
The performance impact of Network Load Balancing can be measured in four key areas:

- CPU overhead on the cluster hosts, which is the CPU percentage required to analyze and filter network packets (lower is better).
- Response time to clients, which increases with the non-overlapped portion of CPU overhead, called latency (lower is better).
- Throughput to clients, which increases with additional client traffic that the cluster can handle prior to saturating the cluster hosts (higher is better).
- Switch occupancy, which increases with additional client traffic (lower is better) and must not adversely impact port bandwidth.

In addition, scalability determines how its performance improves as hosts are added to the cluster. Scalable performance requires that CPU overhead and latency not grow faster than the number of hosts.

**CPU Overhead**
All load-balancing solutions require system resources to examine incoming packets and make load-balancing decisions, and thus impose an overhead on network performance. As previously noted, dispatcher-based solutions examine, modify, and retransmit packets to particular cluster hosts. (They usually modify IP addresses to retarget packets from a virtual IP address to a particular host's IP address.) In contrast, Network Load Balancing independently delivers incoming packets to all cluster hosts and applies a filtering algorithm that discards packets on all but the desired host. Filtering imposes less overhead on packet delivery than re-routing, which results in lower response time and higher overall throughput.

**Throughput and Response Time**
Network Load Balancing scales performance by increasing throughput and minimizing response time to clients. When the capacity of a cluster host is reached, it cannot deliver additional throughput, and response time grows non-linearly as clients awaiting service encounter queuing delays. Adding another cluster host enables throughput to continue to climb and reduces queuing delays, which minimizes response time. As customer demand for throughput continues to increase, more hosts are added until the network's subnet becomes saturated. At that point, throughput can be further scaled by using multiple Network Load Balancing clusters and distributing traffic to them using Round Robin DNS.

**Switch Occupancy**
The filtering architecture for Network Load Balancing relies on the broadcast subnet of the LAN to deliver client requests to all cluster hosts simultaneously. In small clusters, this can be achieved using a hub to interconnect cluster hosts. Each incoming client packet is automatically presented to all cluster hosts. Larger clusters use a switch to interconnect cluster hosts, and, by default, Network Load Balancing induces switch-flooding to deliver client requests to all hosts independently. It is important to ensure that switch-flooding does not use an excessive amount of switch capacity, especially when the switch is shared with computers outside the cluster. The percentage of switch bandwidth consumed by flooding of client requests is called its switch occupancy.

# Network Load Balancing Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Network Load Balancing Tools
- Network Load Balancing Registry Entries
- Network Load Balancing WMI Classes
- Network Port Used by Network Load Balancing


This section provides summary information about the tools, registry entries, WMI classes, and network ports associated with Network Load Balancing. This information is helpful when evaluating Network Load Balancing to determine what tools are available to assist in deployment and operations tasks.

**Network Load Balancing Tools**
The following tools are associated with Network Load Balancing.

**NLB.exe command: Network Load Balancing control program**
**Category**
Tool included in the Windows Server 2003 family, Windows 2000 Standard Server SP 3 and later, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

**Version compatibility**
Available on computers running the Windows Server 2003 family of products, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

The Nlb.exe command replaces Wlbs.exe. Windows NT Load Balancing Service (WLBS) is the former name of Network Load Balancing in Windows NT Server 4.0. For reasons of backward compatibility, WLBS continues to be used in certain instances.

You can perform most cluster operations by using the nlb.exe command. Nlb.exe performs the following functions:

- Suspend and resume all cluster operations. The reason you would suspend operations is to override any remote control commands that might be issued. Resuming cluster operations would not restart the operations, but would enable use of cluster-control commands, including remote control commands.
- Start and stop cluster operations on the specified hosts.
- Disable all new traffic handling on the specified hosts.
- Enable and disable traffic handling for the rule whose port range contains the specified port.
- Display the current cluster state and the list of host priorities for the current members of the cluster.
- Display information about a given port rule.
- Reload the Network Load Balancing driver's current parameters from the registry.
- Display extensive information about the current Network Load Balancing parameters, cluster state, and past cluster activity. The last several event log records produced by Network Load Balancing are shown, including the binary data attached to those records. This assists in technical support and debugging.
- Display information about the current Network Load Balancing configuration.

- Display the media access control (MAC) address corresponding to the specified cluster name or IP address.

**Nlbmgr.exe: Network Load Balancing Manager**
**Category**
Tool included in the Windows Server 2003 family.

**Version compatibility**
Available on computers running the Windows Server 2003 family of products, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server. This tool is also part of the Windows Server 2003 Administration Tools Pack, which you can install on computers running Windows XP Professional. You can use Network Load Balancing Manager on Windows XP Professional only to manage Network Load Balancing clusters on remote computers running the Windows Server 2003 family of products. You cannot install the Network Load Balancing service itself on Windows XP Professional.

Network Load Balancing Manager is used to create and manage Network Load Balancing clusters and all cluster hosts from a single computer, and you can also replicate the cluster configuration to other hosts. By centralizing administration tasks, Network Load Balancing Manager helps eliminate many common configuration errors.

**Network Load Balancing Registry Entries**
The following registry entries are associated with Network Load Balancing.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

For more information about the following registry entries, see the Registry Reference in Tools and Settings Collection.

**network adapter ID subkey**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey.

**AliveMsgPeriod**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

Modifying this registry entry allows you to control both the message exchange period and the number of missed messages required to initiate convergence. Modifying the **AliveMsgTolerance** registry entry will also allow you to do this. The **AliveMsgPeriod** value holds the message exchange period in milliseconds, and the **AliveMsgTolerance** value specifies how many exchanged messages from a host can be missed before the cluster initiates convergence.

The values you enter should reflect your installation's requirements. A longer message exchange period reduces the heartbeat traffic needed to maintain fault tolerance, but it increases the time that it takes for Network Load Balancing to stop sending network messages to an offline host. Likewise, increasing the number of message exchanges prior to convergence reduces the number of unnecessary convergence initiations due to network congestion, but it too increases the time for an offline host to stop receiving network traffic.

Using the default values, 5 seconds are needed to discover a missing host, and another 5 seconds are needed for the cluster to redistribute the load. A total of 10 seconds to stop sending network traffic to an offline host should be acceptable for most TCP/IP applications. This configuration incurs very low networking overhead.

Default value: 1000 (1 second)

Possible range: 100–10000

**AliveMsgTolerance**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

Modifying this registry entry allows you to control both the message exchange period and the number of missed messages required to initiate convergence. Modifying the **AliveMsgPeriod** registry entry will also allow you to do this. For more information, see the previous description for **AliveMsgPeriod**.

Default value: 5

Possible range: 5–100

**DescriptorsPerAlloc**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

Determines the number of connection descriptors allocated at a time. Connection descriptors are used to track TCP connections.

Default value: 512

Possible range: 16–1024

**MaskSourceMAC**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

Enables masking of the media access control (MAC) address. The MAC address is used for communication between network adapters on the same subnet. Each network adapter has an associated MAC address.

If the host is connected to a switch when Network Load Balancing is running in unicast mode, set the value of **MaskSourceMAC** to 1 (the default). If the Network Load Balancing host is running in unicast mode and is attached to a hub that is then connected to a switch, set the value of this entry to 0. If Network Load Balancing is running in multicast mode, this setting has no effect.

Default value: 1

Possible range: 0–1

**MaxDescriptorAllocs**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

Determines the maximum number of times the connection descriptors can be allocated. This value limits the maximum memory footprint of Network Load Balancing.

Default value: 512

Possible range: 1–1024

**NetmonAliveMsgs**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

Determines whether Network Monitor (NetMon) captures Network Load Balancing heartbeat messages on the local host.

To allow NetMon to capture Network Load Balancing heartbeat messages on the local host, set the value of this entry to 1. To get the best performance, leave the value of this entry at its default value of 0.

**NumActions**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

This key is an internal Network Load Balancing entry. Increase the value of this entry only if you encounter an event log message that advises you to do so.

Default value: 100

Possible range: 5–500

**NumAliveMsgs**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

This is an internal Network Load Balancing entry. Increase the value of this entry only if you encounter an event log message that advises you to do so.

Default value: 66

Possible range: 5–500

**NumPackets**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

This is an internal Network Load Balancing entry. Increase the value of this entry only if you encounter an event log message that advises you to do so.

Default value: 200

Possible range: 5–500

**RemoteControlUDPPort**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WLBS\Parameters\Interface*network adapter ID* subkey

**Version**
Available in the Windows Server 2003 family, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server.

Determines the UDP port that is used by Network Load Balancing to receive remote control messages. Note that for backwards compatibility, Network Load Balancing automatically listens to port 1717. If you decide to firewall the remote control port to block remote control messages, you also always need to firewall port 1717.

Do not use ports that are commonly used for other purposes (such as those used for FTP or DNS).

Default value: 2504

### Note

- DecriptorsPerAlloc, MaxDescriptorAllocs, NumActions, NumPackets, and NumAliveMsgs only take effect when the Network Load Balancing driver is binding to the network adapter (upon restarting the hosts, or through disabling and then re-enabling the network adapter). The remaining settings will take effect when the Network Load Balancing driver binds to the network adapter, or when Network Load Balancing is reloaded (using Nlb.exe reload command).

**Network Load Balancing WMI Classes**
The following table lists and describes the WMI classes that are associated with Network Load Balancing.

### WMI Classes Associated with Network Load Balancing

| Class Name | Namespace | Version Compatibility |
|---|---|---|

| | | |
|---|---|---|
| MicrosoftNLB_Cluster | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_ClusterClusterSetting | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_ClusterSetting | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_ExtendedStatus | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_Node | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_NodeNodeSetting | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_NodeSetting | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_NodeSettingPortRule | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_ParticipatingNode | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_PortRule | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_PortRuleDisabled | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_PortRuleFailover | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |
| MicrosoftNLB_PortRuleLoadBalanced | root\cimv2 | Included in Windows Server 2003, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Port Used by Network Load Balancing**

Network Load Balancing uses the following network port.

### Network Ports Used by Network Load Balancing

| Service Name | UDP |
|---|---|
| NLBS | 2504 |

Port numbers in a range of 0 to 65,535 are currently supported by Network Load Balancing for TCP/UDP. The default port range is 0 to 65,535. To load-balance all client requests with a single port rule, use the default port range. By using the default port range, you do not have to worry about which port or ports are associated with the application whose client requests you are load-balancing.

# Resource and Memory Management Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource and Memory Management Components**

4GT and PAE X86 are available only in the 32-bit editions of Microsoft Windows Server 2003, Enterprise Edition, or Microsoft Windows Server 2003, Datacenter Edition, of the Windows Server 2003 family. Both technologies can provide user applications with access to additional RAM. The default virtual address space and flat addressing used on the 64-bit versions of these products negates the need to use these technologies on the 64-bit platforms.

4GT makes more of the computer's virtual memory available to applications by making less virtual memory available to the operating system. 4GT allows applications to address 3 gigabytes (GB) of virtual memory instead of the 2 GB normally allocated for user mode processes. This is a 50 percent increase in virtual memory for the applications, allowing more data to be cached by the application, which can significantly increase performance.

PAE X86 provides the operating system access to memory beyond the default memory address limit of 4 GB. PAE X86 allows software that uses the Address Windowing Extensions (AWE) API set and runs on computers with more than 4 GB of physical memory to map additional physical memory into the application's virtual address space. Applications not using the AWE API set can also benefit from PAE X86 because the operating system uses the larger physical memory to reduce paging and thus increase performance.

Windows System Resource Manager (WSRM) is a tool used to programmatically manage server resources such as processor utilization and memory allocation. Using policies, Windows System Resource Manager allocates specific amounts of resources to applications or processes. These policies allow system administrators to set usage limits based on a percentage of available CPU resources as well as set affinities between managed applications and certain CPUs in the case of multi-processor servers. Windows System Resource Manager consists of two components: the service and the client. The service component can be installed only on the 32-bit and 64-bit versions of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. The client component can be installed only on 32-bit versions of the Windows Server 2003 family. The client component does not support the 64-bit versions of these operating systems.

**Resource and Memory Management Scenarios**

The resource and memory management technologies described in this collection are typically used in server consolidation scenarios and other enterprise systems that make use of large servers with large amounts of memory and multiple CPUs.

4GT is most commonly used when servers are hosting memory intensive applications, such as database management systems, whose performance is directly affected by the availability of physical memory. Servers can take advantage of 4GT provided they have at least 2 GB of RAM available.

PAE X86 is used on servers with more than 4 GB of physical memory. PAE X86 allows the Windows executive software (also known as the kernel) to take full advantage of all available physical memory, including memory above 4 GB, up to 64 GB. Certain data-intensive applications, such as database management systems and scientific and engineering software that need access to very large caches of data, can use the Application Windowing Extensions (AWE) API set to access the memory that PAE can provide beyond 4 GB.

Windows System Resource Manager is most effective on systems where the CPU is in high demand. In these situations, Windows System Resource Manager manages competing processes by raising and lowering process priority. Windows System Resource Manager also provides a mechanism to store and retrieve records of the behavior of managed processes. With this information, Windows System Resource Manager can be used to generate reports of process-resource usage on a per-user, per-application, or other relevant basis. This accounting data can also be used to provide charge-back accounting.

# 4GT Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**4GT Technical Reference**
The technology called 4-gigabyte tuning (4GT), also known as application memory tuning, or the **/3GB** switch, is one of two technologies that increase the amount of physical memory available to user mode applications. The other technology is Physical Address Extension (PAE) X86. For more information about PAE X86, see PAE X86 Technical Reference.

4GT allows memory-intensive applications running on 32-bit versions of the Microsoft Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, operating systems to use 50 percent more virtual memory by making less virtual memory available to the operating system. This subject will provide an overview of 4GT along with an in-depth description of how 4GT works.

### Note

- 4GT is not required on the 64-bit versions of the Windows Server 2003 family.

### In this subject


- What Is 4GT?
- How 4GT Works
- 4GT Tools and Settings

# What Is 4GT?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is 4GT?**

**In this section**


- Common 4GT Scenarios
- Technologies Related to 4GT
- Related Information


The technology called 4-gigabyte tuning (4GT), also known as application memory tuning, or the **/3GB** switch, is one of two technologies that increase the amount of physical memory available to user mode applications. The other technology is Physical Address Extension (PAE) X86. For more information about PAE X86, see "PAE X86 Technical Reference."

4GT allows memory-intensive applications running on Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, 32-bit, to use 50 percent more virtual memory by making less virtual memory available to the operating system. In addition, 4GT can be used on Windows Server 2003, Standard Edition, but only in non-production environments.

Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, provide applications with a flat 32-bit virtual address space that can describe up to 4 gigabytes (GB) of virtual memory. The address space is usually split so that 2 GB of virtual address space is directly accessible to the applications and the other 2 GB is only accessible to the Windows executive software (also known as the kernel). Because of this 2 GB virtual memory limit, applications that are memory intensive and that manage their memory directly through their own caching methods, such as database management systems (DBMS), can experience reduced performance.

4GT makes more of the computer's virtual memory available to applications by making less virtual memory available to the operating system. By enabling 4GT, applications are able to access 3 GB of virtual memory instead of the 2 GB normally allocated for user mode processes. This is a 50 percent increase in virtual memory, allowing more data to be cached and potentially significantly increasing performance.

**Common 4GT Scenarios**

4GT is most commonly used when servers are hosting memory intensive applications, such as database management systems, the performance of which is directly affected by the availability of physical memory. However, it would be ineffective for every application to automatically be provided with the 3 GB address space that 4GT allows. Thus, to provide a selective use of 4GT, 4GT must first be explicitly enabled in the operating system and the application must also be explicitly configured to make use of the extended memory. For more information about 4GT, see "How 4GT Works."

4GT should not be used if either of the following is true:

- Greater than 16 GB of physical memory are available and PAE X86 is enabled on the server.
- Important performance characteristics of the application are adversely affected by limiting the kernel to 1 GB of RAM, for example, more than 1 GB of RAM might be needed to support the maximum number of concurrent connections.

**Note**

- 4GT is not required on the 64-bit versions of the Windows Server 2003 family.

**Technologies Related to 4GT**

The following technologies and components are integral in understanding 4GT.

**Memory Manager**

The 32-bit versions of Windows are able to address $2^{32}$ bytes (4 GB) of memory by default. Windows splits that memory, allocating 2 GB of memory to the kernel and up to 2 GB of memory to each application's private address space. This is done in order to provide protection for both the operating system and the applications. The kernel mode functions get 2 GB so that

they will have enough room to perform mandatory processes without running out of space. User mode applications therefore cannot address memory in the kernel mode space and accidentally corrupt kernel mode memory.

Each application can also utilize up to 2 GB of memory. This is true regardless of the number of applications on the server and despite the operating system's ability to natively recognize only 4 GB of total memory. If the total memory requirements of all applications exceed 4 GB, the memory manager uses a paging file to substitute for RAM. This process is described in the following paragraphs.

The memory manager is responsible for allocating and tracking memory assignments for both the kernel and applications. The memory manager translates virtual memory addresses used by the operating system and applications to actual physical memory locations. The translation of virtual memory to physical memory is transparent to the applications. User mode processes are never able to directly write to real memory and never actually know where their data resides. A user mode process can request a block of memory and write to it. The data written to the memory location might be written to real memory, or might be written to a paging file. A paging file (also known as a swap file) is a file on the hard disk that the memory manager uses to hold data that does not fit in memory. The memory manager moves data from the paging file to memory as needed and moves data from memory to the paging file to make room for new data.

The paging process allows the operating system to overcome its native 2 GB limit. However, it also has a direct impact on performance because of the time necessary to write or retrieve data from disk. 4GT modifies the way the memory manager works, allowing user mode applications that can effectively use greater virtual address space to achieve higher performance.

The kernel address space can also use up to 2 GB of address virtual memory, but all processes in the kernel space share the same 2 GB. Each process does not get its own unique memory space use without affecting the other kernel mode processes. Much of the memory in the kernel mode space also uses a virtual address and depending upon the function, some memory can be paged out to disk as well. However, there is a special type of memory that cannot be paged out to disk. One of the key items that reside in this space is the portion of the memory manager that handles virtual address translation and its related items and drivers for disk access. These items must remain resident in physical memory at all times.

**Page Table Entries**

As described in the previous section, the memory manager translates, tracks, and organizes real and virtual memory for applications and for portions of the kernel. Because of this, the memory manager must index where in real memory or where in a page file the data resides. When an application asks for the data, the memory manager can refer to the index to determine where that data actually resides and then retrieve it for the application. The following diagram shows a simplified explanation of the translation process.

**Use of Page Table Entries**




Page Table with Page Table Entries (PTEs)

The Page Table, one of the internal tables used by the memory manager in translating virtual memory addresses to physical addresses, resides in the kernel's memory space. The kernel's allocated memory space (either 2 GB by default, or 1 GB with 4GT enabled) is divided into two sections, the Paged Pool area and the Page Table Area, as shown in the following figure.

**Paged Pool and Page Table Areas**




On a computer in the Windows Server 2003 family, the Page Table area, which contains the page table entries (PTEs), is 64 MB by default when 4GT is enabled. The Paged Pool space is approximately 256 MB. While these different allocations share the same memory area, the partition between them is fixed at startup. If the operating system runs out of space in either the PTE area or the Paged Pool area, the other area cannot donate space to it, and programs may begin to encounter unexpected errors. This can happen in certain scenarios when using 4GT. For more information about PTEs and 4GT, see "How 4GT Works." By editing the registry, it is possible to limit the size of the Paged Pool space, which will then free additional memory for system PTEs.

### Note

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Physical Address Extension (PAE) X86**

Physical Address Extension (PAE) X86 is another memory management technology that can, in certain scenarios, but used in conjunction with 4GT.

As described earlier in this section, the operating system, by default, can only address up to a total of 4 GB of physical memory. PAE X86 allows the operating system to access additional memory beyond 4 GB. This offers a direct benefit to the applications because more physical memory can be mapped into the application's virtual address space.

Applications that are specifically designed and written to use additional physical memory can benefit from increased performance using PAE X86. Through the use of the Address Windowing Extensions (AWE) API set, the operating system can make the additional physical memory directly available to the application. This is particularly beneficial for large, data intensive application such as database management applications.

Applications not specifically designed to use the additional memory can also benefit from PAE X86 because the operating system uses the larger physical memory to reduce paging and thus increase performance. This is particularly beneficial on consolidation servers hosting multiple applications.

In many situations, it is valuable to use both 4GT and PAE X86. The combination of these two technologies maximizes the amount of physical memory available to applications. However, 4GT cannot be used with PAE X86 when the server has more than 16 GB of physical memory. In this situation, 4GT should be disabled and only PAE X86 should be used.

**Related Information**

The following resource contains additional information that is relevant to this section.

- PAE X86 Technical Reference

# How 4GT Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How 4GT Works**

**In this section**


- 4GT Architecture
- 4GT Processes and Interactions


The technology called 4-gigabyte tuning (4GT) is a scale up component in the Windows Server 2003 family. The Microsoft operating environment delivers scalability via two methods: scale up and scale out. Scaling up refers to running a single application or image on a single server and having the ability to incrementally add system hardware resources (processors and memory) to increase overall system performance. Scaling out refers to distributing the computing workload among multiple servers with the ability to add or subtract servers to increase or decrease capacity.

Enterprises are required to scale flexibly to respond to spikes in business and to support the rapid roll out of new products and services. It is therefore increasingly important that an organization's computing infrastructure provides the ability to increase or decrease computing capacity almost instantly, while continuing to deliver value. 4GT helps provide this scalability.

There are few restrictions regarding the environment in which 4GT can be enabled. 4GT has the following hardware and software requirements:

- x86-based processor
- 2 gigabytes (GB) or more of RAM
- Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition
- Windows Server 2003, Standard Edition, but only in non-production environments.

4GT should not be used if either of the following is true:

- Greater than 16 GB of physical memory are available and Physical Address Extension (PAE) X86 is enabled on the server.
- Important performance characteristics of the application are adversely affected by limiting the kernel to 1 GB of RAM, for example, more than 1 GB of RAM might be needed to support the maximum number of concurrent connections.

**Note**

- 4GT is not required on the 64-bit versions of the Windows Server 2003 family.

**4GT Architecture**

4GT alters the way the memory manager allocates the virtual address space available on the computer. The following diagram and corresponding table illustrate this process.

**4GT Architecture**




When virtual memory space is requested by the user mode applications or the kernel, the memory manager determines where that memory will physically reside. The virtual address to physical address translation information is kept in the page table, which is part of the kernel's virtual address space. Using this translation data, the memory manager accesses the data

The following table describes the components related to the 4GT architecture.

**Components of the 4GT Architecture**

| Component | Description |
|---|---|
| Kernel or User Application | The user mode application or kernel process that requests data at a specific virtual address space. |
| Memory Manager | Operating system component that allocates, organizes, and tracks memory assignments. |
| 4 GB Virtual Address Space | Virtual addresses representing the 4 GB of RAM that is recognized by default by the operating system. Using 4GT, this 4 GB is divided to give user mode applications 3 GB of virtual address space and the kernel 1 GB of virtual address space. The kernel's 1 GB is further allocated into many other specific sections, including the page table area. |
| Page Table | A portion of the kernel's virtual address space that indexes virtual memory addresses to physical addresses. |
| Paging File | The files on a hard disk where data is stored instead of storing that data in RAM. |
| RAM | The physical RAM where data is stored. |

**4GT Processes and Interactions**

4GT is not enabled by default. In order to enable 4GT, the **/3GB** switch must be added to the server's Boot.ini file. The addition of this option causes the operating system to change the default virtual address allocations for applications and the kernel. Upon starting an application, the operating system examines the header file for that application. If the application header includes the "IMAGE_FILE_LARGE_ADDRESS_AWARE" flag and the **/3GB** switch has been used to enable 4GT, then the application has access to the greater virtual address space.

Instead of dividing the available 4 GB of virtual memory evenly between the kernel and any specific user mode application, 4GT provides less virtual address space to the kernel. This results in 3 GB of virtual memory to be allocated to the application. With 4GT, the virtual addresses available to the application range from 0x00000000 through 0XBFFFFFF. This is an increase of the default range of 0x00000000 through 0x7FFFFFFF. The virtual memory address range for the kernel shrinks from 0x80000000 through 0xFFFFFFFF to 0xC0000000 through 0xFFFFFFFF, as shown in the following diagram.

**Virtual Memory Allocation**




4GT can be used with Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. In addition, 4GT can be used with Windows Server 2003, Standard Edition, but only in non-production environments.

**The /Userva Parameter**

In some cases, restricting the kernel to only 1 GB of memory might not be desirable, but the kernel might not need 2 GB either. In order to further refine memory management options, a new parameter, **/Userva**, has been introduced in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. This parameter can only be used in conjunction with the **/3GB** switch. The **/Userva** parameter allows an administrator to determine the specific size of the user mode virtual address space between 2 GB and 3 GB. The value specified by the **/Userva** parameter in the Boot.ini files (expressed in MB) determines the amount of memory the memory manager will return to the kernel. The amount of memory returned is 3072 MB less the value specified by the **/Userva** parameter. For example, to increase the kernel's memory space by 172 MB, the following parameter combination would be used in the Boot.ini file:

**3GB /Userva=2900**

This additional kernel address space is held in reserve and used as additional address space for Page Table Entries (PTEs) in the event that the system runs out of free PTE space. This address space is not allocated to PTEs until the system runs low on PTE space. For more information about PTEs, see "Page Table Entries" later in this section.

**Page File**

Page files store data that cannot be kept in RAM. The default page file size is equal to 1.5 times the amount of physical memory on the system. On a 4 GB computer, this can result in a page file larger than 4 GB. Under such circumstances, the effectiveness of this large of page file could be minimal, based on the usage of the system, such as when hosting database management applications. However, there are also times when even larger page files are needed, such as when multiple applications are consolidated on to one system. Therefore, on servers with 4 GB or more of memory, it is often preferable to change the default size of the page file based on the usage of the server. Regardless, on x86-based computers with 4 GB or more of physical memory there should be a page file sufficiently large to capture a kernel memory dump.

A page file size of at least 2,050 MB is the minimum size necessary for computers with 4 GB or more of physical memory to capture a complete kernel memory dump. The page file can be much larger (up to three times the physical RAM is typical) if the system is being used for application consolidation.

**Page Table Entries**

The page table is used by the memory manager to track the physical location of the page files. When using 4GT, because of the smaller amount of virtual memory allocated to the kernel, the number of Page Table Entries (PTEs) that can be recorded by the memory manager in the page table is also significantly reduced. In a standard configuration (without 4GT), there are typically 80,000 to 140,000 PTEs available. This is reduced to only 40,000 PTEs when using 4GT. This can lead to problems when the amount of physical memory increases and the kernel is required to keep track of larger and larger amounts of PTEs.

**4GT and PAE X86**

PAE X86 is another technology that allows access to additional physical memory. 4GT and PAE X86 can be combined to provide applications large amounts of virtual memory and provide a significant performance boost. However, using 4GT reduces the amount of PTE space available to the kernel, while using PAE X86 dramatically increases the amount of memory that must be indexed and translated by the memory manager. Thus this combination will exhaust system kernel space much earlier than normal. Because of this, the memory manager imposes a virtual memory limit of 16 GB on a system with both 4GT and PAE X86 enabled. Even if a system has 32 GB or more of physical memory, if both options are enabled, only 16 GB of memory will be recognized.

### Note

- Even though the memory manager imposes a hard limit of 16 GB when both 4GT and PAE X86 are enabled, it is possible to encounter problems with lesser amount of memory, such as 8 GB or 12 GB. Therefore, the kernel should be given as much memory as possible.

# 4GT Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**4GT Tools and Settings**

### In this section

- 4GT Registry Entry
- Related Information

The following section describes one registry entry associated with 4-gigabyte tuning (4GT).

### Note

- 4GT is not required on the 64-bit versions of the Windows Server 2003 family.

**4GT Registry Entry**

The following registry entry is associated with 4GT.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**PagedPoolSize**

**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\

**Version**

Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition

This registry key determines the number of bytes allocated to the paged pool. It can be reduced in order to provide additional space for Paged Table Entries (PTEs). Because some computers require a certain amount of paged pool area size, reducing this value will not be an option in all cases. You should not typically need to change this registry value. The operating system calculates an optimal value for this entry based on the amount of physical memory on the computer.

### Note

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Related Information**

The following resources contain additional information that is relevant to this section.

- How 4GT Works
- Registry Reference in Tools and Settings Collection

# PAE X86 Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PAE X86 Technical Reference**
Physical Address Extension (PAE) X86 is one of two technologies that increase the amount of physical or virtual memory available to applications. The other technology is 4-gigabyte tuning (4GT). For more information about 4GT, see "4GT Technical Reference."

PAE X86 allows applications using the Address Windowing Extensions (AWE) API set and running on a computer with more than 4 gigabytes (GB) of physical memory to map additional physical memory into the application's virtual address space. Applications not using the AWE API set can also benefit from PAE X86 because the operating system uses the larger physical memory to reduce paging and thus increase performance. This is of particular benefit to consolidation servers hosting multiple applications.

### Note

- PAE X86 is not required on the 64-bit versions of the Microsoft Windows Server 2003 family.

### In this subject


- What Is PAE X86?
- How PAE X86 Works
- PAE X86 Tools and Settings

# What Is PAE X86?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is PAE X86?**

### In this section


- Common PAE X86 Scenarios
- Technologies Related to PAE X86
- Related Information


Physical Address Extension (PAE) X86 is one of two technologies that increase the amount of physical or virtual memory available to user mode applications. The other technology is 4-gigabyte tuning (4GT). For more information about 4GT, see "4GT Technical Reference."

By default, servers running the 32-bit editions of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, support only 32-bit memory addressing and thus are able to access only 4 gigabytes (GB) of physical memory. However, hardware that offers much larger amounts of physical memory is available. Both the operating system and many applications can benefit from using this additional memory.

PAE X86 allows servers running Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, on a 32-bit hardware platform to access physical memory beyond 4 GB. To do this, PAE changes the memory addressing from 32-bit addressing mode to 64-bit addressing mode, allowing the operating system and high performance drivers and applications access to the additional physical memory. When using the larger amount of physical memory made available by PAE X86, operating system performance can be increased, particularly when the server is hosting multiple applications. Applications can also have direct access to the physical memory beyond 4 GB provided they use the Address Windowing Extensions (AWE) API set.

When using PAE X86, all of the physical memory in the system is treated as general purpose memory. The operating system is able to use this memory for caching and virtual memory management without major changes.

Applications can also use this memory without major changes. The most noticeable difference is that more applications can be running before paging occurs because there is more memory for each process and its associated virtual address space. This also means that more processes can be resident in memory. Because applications not needing larger physical memory do not need to be modified or use the AWE API set, virtually all applications can benefit from PAE X86.

**Common PAE X86 Scenarios**

PAE X86 allows the Windows executive software (also known as the kernel) to take full advantage of all available physical memory, including memory above 4 GB, up to 64 GB. This can result in a significant reduction in paging operations and can improve the performance of the server when multiple applications are hosted on a single computer, such as in application consolidation scenarios.

Additionally, certain data-intensive applications, such as database management systems and scientific and engineering software that need access to very large caches of data can also use PAE to address memory beyond 4 GB. These applications can use the AWE API to access up to 64 GB of memory using a memory window inside their 4 GB process. Just four system calls are used to declare the AWE window and to map additional physical memory inside that window.

### Note

- PAE X86 is not required on the 64-bit versions of the Windows Server 2003 family.

**Technologies Related to PAE X86**

The following technologies and components are integral in understanding PAE X86.

**Memory Manager**

The memory manager translates virtual memory addresses used by the operating system and applications to actual physical memory locations. The translation of virtual memory to physical memory is transparent to the application. User mode processes are never able to directly write to real memory and never actually know where their data resides. A user mode

process can request a block of memory and write to it. The data written to the memory location might be written to real memory, or might be written to a paging file. A paging file (also known as a swap file) is a file on the hard disk that the memory manager uses to hold data that does not fit in memory. The memory manager moves data from the paging file to memory as needed and moves data from memory to the paging file to make room for new data.

The 32-bit versions of Windows are able to address 232 bytes (4 GB) of memory by default. PAE X86 changes that addressing from 32-bit to 64-bit addressing mode. Even with the additional memory made available through PAE X86, if the total memory requirements of all applications exceeds the physical memory available, the memory manager uses the paging file to substitute for real memory; overcoming the physical memory limitation.

The paging process allows the operating system to overcome the real physical memory limits. However, it also has a direct impact on performance because of the time necessary to write or retrieve data from disk. PAE X86 modifies the way the memory manager translates the 32-bit virtual memory address to a physical address, allowing the address to be translated into $2^{36}$ potential physical memory addresses. Access to the additional memory then reduces the amount of paging operations needed.

**Page Table Entries**

As described in the previous section, the memory manager translates, tracks, and organizes real and virtual memory for applications and for portions of the kernel. Because of this, the memory manager must index where in real memory or where in a page file the data resides. When an application asks for the data, the memory manager can refer to the index to determine where that data actually resides and then retrieve it for the application. The following diagram shows a simplified explanation of the translation process.

### Use of Page Table Entries




Page Table with Page Table Entries (PTEs)

The Page Table, one of the internal tables used by the memory manager in translating virtual memory addresses to physical addresses, resides in the kernel's memory space. The kernel's allocated memory space (either 2 GB by default, or 1 GB with 4GT enabled) is divided by function into numerous sections, two of which are the Paged Pool area and the Page Table area, as shown in the following figure.

### Page Pool and Page Table Areas




On a computer running a member of the Windows Server 2003 family, the Page Table area, which contains the Page Table Entries (PTEs), is 64 megabytes (MB) by default when 4GT is enabled. The Paged Pool space is approximately 256 MB. While these different allocations share the same memory area, the partition between them is fixed at startup. This becomes important

when PAE X86 is enabled because the access to additional physical memory that PAE X86 allows requires additional PTEs for tracking and translation. If the operating system runs out of space in either the PTE area or the Paged Pool area, the other area cannot donate space to it, and programs may not receive requested resources, which can lead to unexpected errors for the application. By editing the registry, it is possible to limit the size of the Paged Pool space, which will then free additional memory for system PTEs.

**Note**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**4-gigabyte tuning (4GT)**

Another memory management technology is 4-gigabyte tuning (4GT) that can, in certain scenarios, but used in conjunction with PAE X86.

By default, the 32-bit editions of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, provide applications with a flat 32-bit virtual address space that can address up to 4 GB of virtual memory. The address space is usually split so that 2 GB of virtual address space is directly accessible to the applications and the other 2 GB is only accessible to the kernel. Because of this 2 GB virtual memory limit, applications that are memory intensive and that manage their memory directly through their own caching methods, such as database management systems (DBMS), can experience reduced performance.

4GT can make more of the computer's virtual memory available to applications by making less virtual memory available to the operating system. With 4GT enabled, applications that are designed to be aware of this larger address space are able to address 3 GB of virtual memory instead of the 2 GB normally allocated for user mode processes. This is a 50 percent increase in virtual memory, allowing more data to be cached and potentially significantly increasing performance.

In many situations, it is valuable to use both 4GT and PAE X86. The combination of these two technologies maximizes the amount of physical memory available to applications. However, 4GT cannot be used with PAE X86 when the server has more than 16 GB of physical memory. In this situation, 4GT should be disabled and only PAE X86 should be used.

**Application Windowing Extensions (AWE)**

Because an application's virtual address space can only be extended to a maximum of 3 GB (using 4GT) in a 32-bit operating environment, in order to increase performance, applications (and database applications in particular because they use large datasets and do extensive swapping of data from disk to memory) need a method to map large portions of data into memory and keep that data in real memory at all times. This ability should also allow an application to map and remap data into its virtual memory space very quickly.

Application Windowing Extensions (AWE) is the solution to this need. AWE is an application programming interface (API) set that complements PAE X86. AWE is a set of APIs that allows software developers to create applications that use up to 64 GB of physical non-paged memory in a 32-bit virtual address space on 32-bit platforms. This technology allows for windowed views to this physical memory from within the application's virtual address space.

AWE allows an application to reserve sections of real memory that cannot be paged out to the paging file or otherwise manipulated, except by the reserving application. AWE will instead keep the data in real memory at all times. Because the memory manager does not manage this memory, it will never be swapped out to the paging file.

Unlike PAE, AWE is not an option set at startup. Applications must directly invoke the AWE APIs in order to use them, and the applications must be developed using the AWE APIs in order for the applications to benefit from the AWE APIs.

**Related Information**

The following resource contains additional information that is relevant to this section.

- "4GT Technical Reference"

# How PAE X86 Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How PAE-X86 Works**

**In this section**


- PAE X86 Architecture
- PAE X86 Processes and Interactions


Physical Address Extension (PAE) X86 is one of the scale up technologies in the Windows Server 2003 family. The Microsoft operating environment delivers scalability through two methods: scale up and scale out. Scaling up refers to the ability to incrementally add system hardware resources (processors and memory) to increase overall system performance. Scaling out refers to distributing the computing workload among multiple servers with the ability to add or subtract servers to increase or decrease capacity.

Enterprises are required to scale flexibly to respond to spikes in business and to support the rapid rollout of new products and services. It is therefore increasingly important that an organization's computing infrastructure provides the ability to increase or decrease computing capacity almost instantly, while continuing to deliver on value. PAE X86 helps provide this scalability.

**Note**

- PAE X86 is not required on the 64-bit versions of the Windows Server 2003 family.

**PAE X86 Architecture**

PAE X86 alters the way the memory manager translates memory addresses. The memory manager moves from a two-level linear address translation to a three-level address translation. The extra layer of translation is what provides access to physical memory beyond 4 gigabytes (GB). The components used in the PAE X86 method of translation are shown in the following diagram.




The following table describes the components related to the PAE X86 architecture.

**Components of the PAE X86 Architecture**

| Component | Description |
|---|---|
| Virtual Address | A virtual address that is translated into a specific physical address. This is the address that an application or the kernel uses when accessing data in memory. |
| Page Directory Pointer Index (PDPI) | A table listing pointers to specific page directories. |
| Page Directory | A table with 512 entries, each pointing to a specific page table. |
| Page Table | A table with 512 entries, each pointing to a specific 4 kilobyte (KB) portion of memory. |

| | |
|---|---|
| 4 KB Page | A 4 KB portion of physical memory. |

**PAE X86 Processes and Interactions**
PAE X68 is not automatically enabled. It must first be enabled by adding the **/PAE** switch to the Boot.ini file. By using this switch, the operating system loads a different kernel from \i386\Driver.cab file. The kernel file is Ntkrnlpa.exe for uniprocessor systems and Ntkrpamp.exe for multiprocessor systems. The kernel is loaded based on only the switch and independently from the existing physical memory.

Physical memory is treated as general purpose memory, so no new APIs are needed for the kernel to access the memory above 4 GB. Direct I/O can be done to addresses above 4 GB. This requires Dual Address Cycle (DAC) capable or 64-bit Peripheral Component Interconnect (PCI) devices. Devices and drivers offering this capability are called Large Memory Enabled (LME).

Because the 32-bit editions of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, do not have a kernel PAE or LME API interface, the PAE X86 kernel ensures that many items are identical to the standard kernel, including:

- Kernel memory space organization is unchanged.
- Registry flags work the same.
- Non-paged pool size remains the same.
- The 4-gigabyte tuning (4GT) feature is supported for up to 16 GB of RAM.
- The application header IMAGE_FILE_LARGE_ADDRESS_AWARE, which allows applications to make use of 4GT, will continue to work.
- "Well-known" kernel addresses remain in the same locations.

Applications do not need to be aware of PAE X86. Once enabled, the kernel will automatically make use of the additional memory for file caching, potentially improving the performance of applications running on the server.

**Virtual Address Translation without PAE X86**
When an application or the kernel requests access to physical memory, the memory manager must first translate the request, from a virtual address to a physical address. The normal implementation of virtual address translation uses a two-level page table structure. When PAE X86 is not enabled, the 32-bit virtual address is interpreted as three separate components: a 10-bit Page Directory Index, a 10-bit Page Table Index, and a 12-bit Byte Index. The following figure shows the translation of the linear virtual memory address to a physical memory location when PAE X86 is not enabled.




The following actions are involved in translating a virtual address when PAE X86 is not enabled:

- The memory manager locates the Page Directory for the current process. The memory manager is informed of the

location of the Page Directory for the process by a special control register.

- The 10-bit Page Directory Index in the virtual address is used to locate a Page Directory Entry (PDE) that defines the location of the Page Table needed to translate the virtual memory address.
- The 10-bit Page Table Index in the virtual address is used to locate the correct Page Table Entry (PTE).
- The PTE is used to locate the correct 4 KB page of memory.
- After accessing this page, the 12-bit Byte Index in the virtual address is used to locate the physical address of the desired data.

**Virtual Address Translation with PAE X86**

With PAE X86 enabled, the operating system moves from a two-level linear address translation to a three-level address translation. The memory manager still implements Page Directories and Page Tables, but a third level of translation, the Page Directory Pointer Table, is also added. The extra layer of translation is what provides access to physical memory beyond 4 GB. The following figure shows the translation of the virtual address to a physical memory location when PAE X86 is enabled.




Instead of splitting the virtual memory address into three separate components for indexing into memory tables, it is split into four separate components: a 2-bit Page Directory Pointer Index (PDPI), two 9-bit components (a Page Directory Index and a Page Table Index), and a 12-bit Byte Index.

The control register now points the memory manager to the Page Directory Pointer Index (PDPI) that is 2-bits wide. The first two bits of the virtual address are used as an index into the table, with the resulting value pointing to a Page Directory. The first 9-bit field in the virtual address is then used to index the Page Directory. The indexed value then points to a Page Table. The second 9-bit field is an index into the Page Table. The resulting Page Table Entry (PTE) points to the actual 4 KB page in memory where the desired byte is located. The remaining 12 bits of data in the virtual address point to the specific byte of data on that page.

**Page Table Entries**

As described earlier in this section, the Page Table is used by the memory manager as part of the translation process between virtual memory addresses and physical memory addresses. When using the large amounts of memory made available through PAE X86, the number of PTEs needed to track that physical memory can increase dramatically. This can lead to problems when the amount of physical memory allocated to the kernel for storing PTEs is reduced, such as when using 4GT. With 4GT, the smaller amount of virtual memory allocated to the kernel means that the number of PTEs that can be recorded by the memory manager is much smaller. In a standard configuration (without 4GT), there are typically 80,000 to 140,000 PTEs available. This is reduced to only 40,000 PTEs when using 4GT. The next section describes the implications of this change.

**PAE X86 and 4GT**

4GT is another technology that makes more physical memory available to user applications. PAE X86 and 4GT can be combined to provide applications large amounts of virtual memory and provide a significant performance boost. However, using 4GT reduces the amount of PTEs available to the kernel, while using PAE X86 dramatically increases the amount of memory that must be indexed and translated by the memory manager. This combination will exhaust system kernel space much earlier than normal. Because of this, the memory manager imposes a virtual memory limit of 16 GB on a system with

both 4GT and PAE X86 enabled. Even if a system has 32 GB or more of physical memory, if both options are enabled, only 16 GB of memory will be recognized. Thus, if PAE X86 is enabled on a system with more that 16 GB of physical memory, 4GT should not be used.

**Note**

- Even though the memory manager imposes a hard limit of 16 GB when both 4GT and PAE X86 are enabled, it is possible to encounter problems with lesser amount of memory, such as 8 GB or 12 GB. Therefore, the kernel should be given as much memory as possible.

**Application Windowing Extensions (AWE)**

Virtually all applications can benefit from the increased performance resulting from the kernel's use of PAE X86. However, in order to directly access the additional physical memory enabled by PAE X86, applications must use the Application Windowing Extensions (AWE) API set specifically designed for this purpose.

AWE is an API set that complements PAE X86. Unlike PAE X68, AWE is not an option set at startup and applications must directly invoke the AWE APIs in order to use them. Applications must be developed using the AWE API set in order to benefit from it.

Because an application's virtual address space can only be extended to 3 GB (using 4GT) on a 32-bit operating environment, applications (and database applications in particular, since they use large data sets and do a lot of swapping of data from disk to memory) need a method to map large portions of data into memory and keep that data in real memory at all times in order to increase performance.

AWE provides applications with this ability. AWE allows an application to reserve portions of real memory that cannot be paged out to the paging file or otherwise manipulated except by the reserving application. AWE will keep this data in real memory at all times. Because the memory manager does not manage this memory, it will never be swapped out to the paging file. The application is completely responsible for handling the memory.

AWE places the following restrictions on how this additional memory may be used, primarily because these restrictions allow extremely fast mapping, remapping, and freeing. Fast memory management is important for these potentially enormous address spaces and for attaining the highest possible performance.

- Virtual address ranges allocated for AWE are not sharable with other processes (and therefore not inheritable). Similarly, two different AWE virtual addresses within the same process are not allowed to map the same physical page. These restrictions provide fast remapping and cleanup when memory is freed.

- Since this memory is never paged, the physical pages that can be allocated for an AWE region are limited by the number of physical pages present on the computer. The memory is reserved until the application explicitly frees it or exits. The physical pages allocated for a given process can be mapped into any AWE virtual region within the same process. Applications that use AWE must be careful not to take so much physical memory that they cause other applications to page excessively or prevent creation of new processes or threads due to lack of resources.

- AWE virtual addresses are always read/write and cannot be protected. Read-only memory, no access memory, guard pages, or other similar calls cannot be specified.

- AWE address ranges cannot be used to buffer data for graphics or video calls.

- An AWE memory range cannot be split, nor can pieces of it be deleted. Instead, the entire virtual address range must be deleted as a unit when deletion is required.

- Applications can map multiple regions simultaneously, provided they do not overlap.

- Applications that use AWE are not supported in emulation mode. That is, an x86-based application that uses AWE functions must be recompiled to run on another processor, whereas most applications can run without recompiling under an emulator on other operating environments.

**Paging File**

Page files store data that cannot be kept in RAM. The default page file size is equal to 1.5 times the amount of physical memory on the system. On a 4 GB computer, this can result in a page file larger than 4 GB. Under such circumstances, the effectiveness

of this large of page file could be minimal, based on the usage of the system, such as when hosting database management applications. However, there are also times when even larger page files are needed, such as when multiple applications are consolidated on to one system. Therefore, on servers with 4 GB or more of memory, it is often preferable to change the default size of the page file based on the usage of the server. Regardless, on x86-based computers with 4 GB or more of physical memory there should be a page file sufficiently large to capture a kernel memory dump.

A page file size of at least 2,050 megabytes (MB) is the minimum size necessary for computers with 4 GB or more of physical memory to capture a complete kernel memory dump. The page file can be much larger (up to three times the physical RAM is typical) if the system is being used for application consolidation.

# PAE X86 Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PAE X86 Tools and Settings**

### In this section


- PAE X86 Registry Entry
- Related Information


The following section describes one registry entry associated with Physical Address Extension (PAE) X86.

### Note

- PAE X86 is not required on the 64-bit versions of the Windows Server 2003 family.

**PAE X86 Registry Entry**

The following registry entry is associated with PAE X86.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**PagedPoolSize**

**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\

**Version**

Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition

This registry key determines the number of bytes allocated to the paged pool. It can be reduced in order to provide additional space for Paged Table Entries (PTEs). This is sometimes necessary when using PAE X86 in conjunction with 4-gigabyte tuning (4GT). However, because some computers require a certain amount of paged pool area size, reducing this value will not be an option in all cases. You should not typically need to change this registry value. The operating system calculates an optimal value for this entry based on the amount of physical memory on the computer.

### Note

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Related Information**

The following resources contain additional information that is relevant to this section.

- How PAE X86 Works
- Registry Reference in Tools and Settings Collection

# Windows System Resource Manager Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows System Resource Manager Technical Reference**
Windows System Resource Manager is a feature of Microsoft Windows Server 2003, Enterprise Edition, and Microsoft Windows Server 2003, Datacenter Edition.

With Windows System Resource Manager, administrators can control how CPU and memory resources are dynamically allocated to applications, services, and processes. Managing resources in this way improves system performance and reduces the chance that applications, services, or processes will interfere with the rest of the system. It also creates a more consistent and predictable experience for users of applications and services running on the computer.

You can use Windows System Resource Manager to manage multiple applications on a single computer or users on a computer on which Terminal Services is installed.

**In this subject**


- What is Windows System Resource Manager?
- How Windows System Resource Manager Works

# What is Windows System Resource Manager?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Windows System Resource Manager?**

**In this section**

- Common Windows System Resource Manager Scenarios

With Windows System Resource Manager, administrators can control how CPU and memory resources are dynamically allocated to applications, services, and processes. Managing resources in this way improves system performance and reduces the chance that applications, services, or processes will interfere with the rest of the system. It also creates a more consistent and predictable experience for users of applications and services running on the computer.

**Common Windows System Resource Manager Scenarios**

Windows System Resource Manager is designed for scenarios that involve running multiple applications on a single physical computer. Using Windows System Resource Manager, you can create resource-allocation policies, which determine how CPU and memory resources are allocated to running processes, and process-matching criteria, which use matching logic to group processes for easy management. Windows System Resource Manager supports the following scenarios:

- Allocating resources through server consolidation, to reduce the ability of applications to interfere with each other, such as when running multiple instances of a single application, or combining line-of-business applications with other applications or services.
- Managing users, either individually or by security group, in large terminal server systems by application, which limits the effect heavily used resources have on other users.
- Administering resource usage of multiple SQL Server instances sharing a server.
- Administering resource usage of Internet Information Services (IIS) 6.0 application pools on a server, such as when one server hosts multiple Web sites.
- Reports on memory usage and CPU time to support service level agreement (SLA) metrics.

Windows System Resource Manager is included with only Microsoft Windows Server 2003, Enterprise Edition, and Microsoft Windows Server 2003, Datacenter Edition.

# How Windows System Resource Manager Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Windows System Resource Manager Works**

**In this section**


- Terms and Definitions
- Windows System Resource Manager Architecture
- Windows System Resource Manager Processes and Interactions


Windows System Resource Manager (WSRM) is designed to manage multiple applications on a single computer or multiple users on a computer on which Terminal Services is in use. This supports a variety of scenarios, including consolidation, and can increase the efficiency with which physical hardware on a server is used by running applications.

Resource usage is controlled by using a dynamic process priority management algorithm that determines how to allocate system resources among multiple processes. You can administer resource policies and monitor resource usage with Windows System Resource Manager using Microsoft Management Console (MMC) or a command-line interface.

Windows System Resource Manager is available only on Microsoft Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. The following condition increases the effectiveness of Windows System Resource Manager:

- CPU consumption targets are more likely to be met if processes are contesting for CPU bandwidth. This is because there is more contention among processes for scarce CPU resources when the system is experiencing heavy load. Windows System Resource Manager does not enforce resource allocation policies until CPU consumption reaches 75 percent. Before this point, all running processes are already getting all the CPU resources they can use. Windows System Resource Manager is needed only if one or more running processes are unable to use CPU resources because of other running processes using too much CPU bandwidth.

The following restrictions must be considered before using Windows System Resource Manager:

- Do not use Windows System Resource Manager in conjunction with another resource manager because neither resource manager is aware that the other is attempting to control system resources.
- Do not use Windows System Resource Manager to manage applications or processes that modify their process priority, memory limits, or processor affinity dynamically because this can have an adverse effect on the application being managed. Instead, use the application's own interface for controlling resources.
- Processes running in job objects cannot be managed because job objects include a management layer. Instead, use the application that creates the job objects to control the resources allocated to processes in the job object.
- The service component of Windows System Resource Manager can be installed only on the 32-bit and 64-bit versions of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.
- The client component of Windows System Resource Manager can be installed only on all versions of Windows 2000 and Windows XP Professional, as well as the 32-bit versions of the Windows Server 2003 operating systems. The client component does not support the 64-bit versions of these operating systems.

**Terms and Definitions**

The following terms are associated with Windows System Resource Manager.

**committed memory**

Memory that either resides in the working set of a process or maps to a disk file, such as a page file or other memory-mapped

file.

**CPU consumption target or CPU allocation**
Specifies a target CPU bandwidth for the resource allocation. Windows System Resource Manager uses a dynamic process priority management algorithm to try to limit the CPU usage of a process to the target specified in the matching resource-allocation policy of that process.

**default group**
Includes all running processes not matched by a specific process-matching criterion. Processes in the default group are allocated all of the managed CPU resources that are not allocated to processes matching the active resource-allocation policy.

**managed CPU**
The total CPU resources being used by managed processes.

**managed process**
All running processes, except those on the system-defined exclusion list or the user-defined exclusion list. They are managed by Windows System Resource Manager.

**process**
An object created by Windows to represent an executing program. This includes the virtual address space and the control information necessary for the program to run.

**process-matching criterion**
The mechanism Windows System Resource Manager uses to match running processes to a resource-allocation policy.

**process priority**
A value for each process that helps determine how CPU resources are used. Windows System Resource Manager dynamically changes the priority of a process so the CPU usage matches the targeted CPU allocation for that process.

**processor affinity**
On systems with more than one processor, processor affinity specifies the set of processors on which the process is exclusively allowed to run. Also known as CPU affinity.

**resource allocation**
The part of a resource-allocation policy that specifies the set of resources to be allocated to the processes matching to a single process-matching criterion.

**resource-allocation policy**
An ordered list of resource allocations that specifies the management behavior of Windows System Resource Manager at a given time.

**sub-allocation**
The process by which a resource allocation is divided by multiple child resource allocations.

**system-defined exclusion list**
The list of system processes that are excluded from management by Windows System Resource Manager. The system-defined exclusion list is built into Windows System Resource Manager and cannot be modified by users.

**unmanaged process**
A process on either the system-defined exclusion list or the user-defined exclusion list. Windows System Resource Manager never applies a resource-allocation policy to an unmanaged process, even if a resource-allocation policy would otherwise match.

**user-defined exclusion list**
A modifiable list of processes to be excluded from management by Windows System Resource Manager.

**working set memory**
For a process, the physical memory that is consumed during the time when the process is running and can be paged to disk.

**Windows System Resource Manager Architecture**
Windows System Resource Manager consists of two components: the service and the client. For the client, two administrative interfaces are provided: the command-line administration tool (Wsrmc.exe) and the WSRM snap-in available in Microsoft Management Console (MMC). The administrator creates resource-allocation policies and process-matching criteria to control how CPU and memory resources are allocated to applications, services, and processes. These policies are managed by the Windows System Resource Manager service, which then uses the memory manager and kernel scheduler to implement the policies.

The following figure illustrates the different components of Windows System Resource Manager and the interactions between

the components.

**Windows System Resource Manager Architecture**

The following table describes the principal components related to Windows System Resource Manager architecture.

**Components of Windows System Resource Manager Architecture**

| Component | Description |
|---|---|
| Command-line Interface (Wsrmc.exe) | Provides administrative control of the Windows System Resource Manager service using the command-line interface. The command-line interface provides equivalent administrative control to the WSRM snap-in. |
| WSRM Snap-in (MMC interface) | Provides administrative control of the Windows System Resource Manager service using the MMC interface. |
| Distributed Common Object Model (DCOM) Interface | The remote application programming interface (API) through which the client and Windows System Resource Manager service communicate. |
| WSRM Service | Performs resource management by using logic components that interact with the operating system and running processes.<br><br>The service continually monitors the resource usage of each managed process in the group. If a process exceeds its resource allocation, the service attempts to bring the resource usage of the process back to its specified target. |
| Accounting Database | Stores information about managed processes on a per-process basis. Provides a mechanism through which data on the behavior of managed processes can be stored, grouped, and retrieved. |
| Policy Store | Stores all resource-allocation policies, process-matching criteria, and other related information. |
| WSRM Settings | Stores the current management settings. |
| Calendar | Stores all calendar-related information, such as calendar events and schedules. |
| Memory Manager | Allocates memory to running processes. |
| Kernel Scheduler | Determines when and for how long each process in the system is allowed to use the processors in the computer. |

**Windows System Resource Manager Processes and Interactions**

This section describes the different processes and interactions of Windows System Resource Manager.

**Windows System Resource Manager Service**

The Windows System Resource Manager (WSRM) service monitors the state of managed processes and takes action when those processes do not meet their target CPU allocations. The service does the actual resource management work by using logic components that interact with the Windows scheduler and running processes.

Logic components are defined in the following table.

**Logic Components**

| Logic Components | Description |
|---|---|
| Exclusion logic | Matches all processes on either the user-defined or system-defined exclusion lists and marks them to be excluded from management by Windows System Resource Manager. |
| Matching logic | Examines all processes not on the exclusion lists and groups them by matching process-matching criterion for management. |

| | |
|---|---|
| Process priority logic | Implements the dynamic process priority management algorithm to raise and lower the priority of a running process. Changing the priority of a process affects how CPU resources are allocated among running processes. |
| Memory logic | Sets the memory limits specified in the associated resource allocation, if the limits are defined. |
| Affinity logic | When processor affinity is specified in the managing resource allocation policy, configures the process to run on the set of processors specified. |
| Job accounting logic | Periodically writes resource consumption statistics to a private job accounting database. Also, provides an interface for querying and exporting stored resource consumption statistics from the job accounting database. Job accounting logic is not included in the figures that appear later in this document because it does not apply to the interactions illustrated in those figures. |

While management is active, the Windows System Resource Manager service continually monitors the resource usage of each managed process in the group. If a process momentarily exceeds its resource allocation, the service attempts to bring the resource usage of the process back to its specified target.

**How the Service Matches Processes**

A running process is matched when the Windows System Resource Manager service discovers it, or when the active resource-allocation policy is changed, which causes the service to re-poll all processes again. The system-defined exclusion list and the user-defined exclusion list are checked first and all matching processes are placed in the excluded processes group. Processes that do not match any process-matching criterion are automatically assigned to the default group. After a process is assigned to a group, the service enforces the associated resource allocation for the process by setting resource limits or modifying the priority of the process. No changes are ever made to processes in the system-defined exclusion list or the user-defined exclusion list.

The service makes no other changes to the process or the process-matching criterion unless the process-matching criterion or the resource-allocation policy is edited. Changes made to the user-defined exclusion list do not take effect until a new resource-allocation policy is set to manage, or the current resource-allocation policy is stopped and restarted. If either the process-matching criterion or the resource-allocation policy is edited, the service re-polls all of the processes running on the system following the standard allocation algorithm.

**The Dynamic Process Priority Management Algorithm**

During each poll, the service checks to determine whether the processes match their target allocations and whether any new processes have been created since the last poll.

If a process exceeds its target allocation, the service uses the dynamic process priority management algorithm to try to make the process conform to its target allocation. If changing the priority of the process does not make the CPU usage of the process conform to its target, the algorithm adjusts the target allocation of the process. It does so by using a process called re-allocation, which is described below.

If a new process is identified during the polling interval, the Windows System Resource Manager service takes the following steps to determine whether the new process should be managed:

- The Windows System Resource Manager service matches the properties of the new process to each process-matching criterion in the managing policy. It does so in the order of the process-matching criteria that is specified in the resource-allocation policy.

- When the properties of the new process are found to match a process-matching criterion, the new process is then grouped with other processes that have properties matching the same process-matching criterion. The new process is then subject to the resource allocations specified in the resource-allocation policy.

- If the process does not match any of the process-matching criteria, it is added to the default group. Processes matching in the default group are given CPU bandwidth not allocated to other running processes. No memory limits are applied to processes in the default group.

**Process-Matching Criterion**

Process matching is the mechanism that Windows System Resource Manager uses to match processes and then aggregate them into groups to be managed.

A process-matching criterion creates an association between a group of one or more matched processes and a resource-allocation policy. Windows System Resource Manager then manages all of the matched processes using the resource allocations specified in the resource-allocation policy. If a CPU allocation is specified, the CPU resources are distributed evenly

among all the processes in the group. If a memory limit is specified, the limit is applied on a per-process basis to all the processes in the group.

Windows System Resource Manager maintains the specified resource allocations until the managing policy is changed.

The following table defines the entries that each process-matching criterion can include.

**Events for Process-Matching Criterion**

| Event | Description |
|---|---|
| Included files or command lines | Lists wildcard or regular expression match strings. |
| Excluded files or command lines | Lists wildcard or regular expression match strings. |
| Included users or groups | Lists user account names or group account names. |
| Excluded users or groups | Lists user account names or group account names. |

Windows System Resource Manager uses the following process to decide if a process matches a given process-matching criterion.

- If the process-matching criterion has entries for included files or commands lines, Windows System Resource Manager attempts to match each match string against the process properties in the following order, stopping on the first match:

  1. Process file name
  2. Fully-qualified process path and file name
  3. Full process command line

- If the process matches the included files or command-line entries, then Windows System Resource Manager attempts to match the entries for excluded files or command lines, if they are defined. The process for matching the excluded entries is identical to the process for included entries. If the process matches against the excluded entries, then the process is treated as one that does not match the process-matching criterion.

- If the process-matching criterion has entries for included users or groups, Windows System Resource Manager matches each user or group name against the user account used to create the process, stopping on the first match. A user account matches if it is the same as the user account in the process. A group matches if the user account used to create the process is a member of the group. If the process matches against the excluded entries, then the process is treated as one that does not match the process-matching criterion.

- If the process matches the included users or groups entries, then Windows System Resource Manager attempts to match the entries for excluded users or groups if they are present. The process for matching the excluded entries is identical to that of the included entries. If the process matches against the excluded users or groups entries, then the process is treated as one that does not match the process-matching criterion.

- If both a files or command-line entry and user or group entry are provided, then both must match in order for the process to match the process-matching criterion.

**Resource-Allocation Policies**

A resource-allocation policy is the mechanism that specifies the resource usage for a managed process.

After a process-matching criterion is defined, it can be used by any number of resource-allocation policies. Only one resource-allocation policy can be in effect and managing at a time. The managing resource-allocation policy determines the resource usage for all managed processes, in other words, those processes not included in the system-defined or user-defined exclusion lists. A resource-allocation policy is an ordered list of resource allocations. Each resource allocation consists of a process-matching criterion and one or more allocations, such as a CPU consumption target, memory limit, or processor affinity. The allocations specified apply to all processes that match the process-matching criterion associated with the resource allocation.

**The Default Group**
All running processes that are not matched by any of the process-matching criteria in the managing resource-allocation policy are grouped in the default group.

The default group is allocated any CPU bandwidth that is unaccounted for by the managing policy. This bandwidth is distributed on an equal-shares basis. This ensures that all processes in the default group get an equal proportion of the total CPU bandwidth that is available to the default group. Windows System Resource Manager enforces this allocation, which prevents one or more applications in the default group from using more than their share of CPU bandwidth.

**Exclusion Lists**
While Windows System Resource Manager can effectively manage most applications, it is best not to use Windows System Resource Manager to manage processes that many applications depend on, such as Windows system services. This is likely to have adverse effects, so Windows System Resource Manager maintains a fixed list of Windows processes that should not be managed. This list is called the system-defined exclusion list. In addition, there is a user exclusion list intended for any other services or processes that might not behave properly under management. A process that is included on either of these lists is known as an unmanaged process.

**The System-Defined Exclusion List**
The system-defined exclusion list cannot be modified; however, you can exclude additional processes that you do not want Windows System Resource Manager to manage by manually adding them to the user-defined exclusion list.

To ensure that only designated system processes are excluded from being managed, the processes on the system-defined exclusion list are matched using the exact process path. This prevents other processes with similar names from being matched by mistake.

The following system processes are included by default on the system-defined exclusion list:

- Csrss.exe
- Dfssvc.exe
- Dumprep.exe
- Dwwin.exe
- Llssrv.exe
- Lsass.exe
- Msdtc.exe
- Smss.exe
- Services.exe
- Spoolsv.exe
- System.exe
- Taskmgr.exe
- Winlogon.exe
- Wmiprvse.exe
- Wsrm.exe
- Wsrmc.exe

**The User-Defined Exclusion List**
Processes on the user-defined exclusion list are either added manually or the system adds several processes by default. The processes included on the user-defined exclusion list by default are several Svchost.exe processes. Important operating system services run inside the Svchost.exe processes. To ensure that these processes have access to their required CPU bandwidth, they are excluded from management. It is not usually necessary to manage these processes; however, if one or more of these processes start to consume excessive CPU bandwidth, you can manage them by removing them from the user-defined exclusion list. This allows Windows System Resource Manager to control the CPU bandwidth usage of these processes.

Although it is rarely necessary, you can include processes on the user-defined exclusion list that do not respond well to management and any other processes that you do not want Windows System Resource Manager to manage. Processes that Windows System Resource Manager does not control are unmanaged and can, therefore, consume resources freely.

Because the behavior of unmanaged applications cannot be controlled, the effectiveness of Windows System Resource Manager can decrease if you include on the user-defined exclusion list multiple processes or processes that use resources intensively.

**Calendar**
The Windows System Resource Manager calendar schedules the automatic activation of resource-allocation policies. Policies can be activated on a one-time or recurring basis. Schedules provide flexibility to manage multiple resource-allocation policies and to activate various resource-allocation policies at different, predefined times over a 24-hour period.

When the calendar is enabled, and no policy is scheduled for the current time, the default calendar resource-allocation policy is in effect. The default calendar policy is configurable and can be a user-defined resource-allocation policy or the default calendar resource-allocation policy. The default calendar policy is configurable and can be set to any policy.

The following table describes the three types of calendar events you can use to schedule resource-allocation policies.

### Types of Calendar Events

| Type of Calendar Event | Description |
|---|---|
| One time | Create a one-time calendar event to schedule one resource-allocation policy or schedule to run once. Requires a starting date and time and duration (or ending date and time), and a specified resource-allocation policy to run. |
| Recurring event | Create a recurring event to schedule one resource-allocation policy or schedule to run more than once. Requires a starting date and time and duration (or ending date and time), a recurrence pattern, a recurrence length, and a specified resource-allocation policy to run. |
| Schedule (day schedule applied to calendar dates) | Create a schedule to activate multiple resource-allocation policies to run over a 24-hour period. Policy schedules can be reused and applied to one or more dates on the calendar. You can change the date on which to apply the schedule without affecting the entire schedule. |

**Accounting**
Windows System Resource Manager provides a mechanism to store and retrieve records of the behavior of managed processes. Accounting information is stored in an accounting database for easy administration and retrieval. This information is logged in the accounting database at a default interval of 10 minutes, which is configurable. The interval determines the frequency with which the data for all managed processes on the computer managed by Windows System Resource Manager is logged to the accounting database. With this information, process-resource usage reports can be generated on a per-user, per-application, or other relevant basis.

**How CPU Resources Are Managed**
CPU resources are represented as both a percentage of total CPU bandwidth and as a set of physical processors. These resources are controlled through the use of CPU consumption targets and processor affinity, respectively. Windows System Resource Manager uses the dynamic process priority management algorithm to determine how CPU bandwidth is allocated.

**Managed CPU and Total CPU**
CPU bandwidth targets in resource allocation policies are expressed as a percentage from 0 to 100. This represents the total managed CPU resources of the managed computer. The total bandwidth for both a one-processor system and an eight-processor system is 100 percent. The maximum CPU bandwidth percentage you can specify in a resource allocation, however, is 99 percent because a minimum of 1 percent is always reserved for the default group.

Managed CPU corresponds to the CPU bandwidth that is available to the processes Windows System Resource Manager is managing.

Total CPU corresponds to the total CPU bandwidth on the system. The total CPU resources available for Windows System Resource Manager to manage might not match the total CPU resources on the computer.

The difference between the managed CPU and the total CPU in a system is determined by how CPU is used by unmanaged and managed processes. Unmanaged processes can consume as much of the available CPU resources as needed. Windows System Resource Manager can manage only those CPU resources that are not being consumed by unmanaged processes. This means that the managed available CPU might be less than the total CPU available in the system. As the percentage of total CPU resources consumed by unmanaged processes increases, the ability of Windows System Resource Manager to meet target CPU allocations for managed processes decreases. Thus, if you limit the number of unmanaged processes in the system, you can maximize the accuracy of overall system resource management.

**CPU Targets**
To control the allocation of managed CPU, you can set CPU consumption targets for each resource allocation. The CPU consumption target is expressed as a percentage that corresponds to the total managed CPU available on the system when the resource allocation is created. After you create a CPU allocation, the consumption target percentage is deducted from the total available managed CPU resources. For example, if you create a resource allocation with 30 percent as the CPU consumption target, subsequent resource allocations will have only 69 percent available.

CPU consumption targets are applied to the group of processes matched by the resource allocation's associated process-matching criterion.

For example, if you create a resource allocation with a CPU consumption target of 10 percent CPU and a process-matching criterion that matches 10 processes, the 10 percent CPU is applied to the group of processes rather than to each process individually.

The distribution of CPU bandwidth among multiple processes in a group is controlled by the Management Rule property of each resource allocation. The Management Rule can take one of three possible values:

- **Standard.** This is the default setting. It means that Windows System Resource Manager does not control the distribution of CPU bandwidth among processes in the group.

- **Equal per process.** Windows System Resource Manager will manage individual processes in the group according to equal targets, each an equal share of the group's target.

- **Equal per user.** Windows System Resource Manager will manage sub-groups of processes in the group, assigning equal shares according to the user that started the process.

**Allocation**
In order to manage processes toward their CPU allocation targets, Windows System Resource Manager samples the CPU consumption of all managed processes in the system and computes a moving average of each over the last few samples. The priorities of processes consuming above their targets are lowered to prevent the consumption of CPU bandwidth that is allocated to other processes. Similarly, processes consuming below their targets have their priorities raised to give them better access to CPUs in the system. Processes that are unable to consume their allocated CPU bandwidth, even after appropriate priority adjustments have been made to the priorities of other processes, will have their unconsumed allocation re-allocated to other processes. This occurs according to the proportions specified in the resource allocations.

**Sub-allocation**
A sub-allocation is a subdivision of a resource allocation. Each subdivided resource allocation is referred to as the parent allocation, and each of its sub-allocations is a child resource allocation. Sub-allocations can be further subdivided, so there is no effective limit to the size of the resource hierarchy. A sub-allocation behaves similarly to a resource allocation; however, its CPU percentage targets are interpreted as a percentage of its parent's CPU allocation, rather than a percent of the managed CPU of the whole machine.

In a sub-allocation, CPU resources are allocated first to the child resource-allocation policies, starting with the child that is at the bottom of the hierarchy. This last child in the hierarchy gets as close to its specified CPU allocation as possible. The remaining CPU resources that this last child does not use become available to its parent. The parent consumes bandwidth up to its specified CPU allocation. Any remaining CPU bandwidth then becomes available to its parent, if one exists. If a parent does not exist, the remaining CPU resources are allocated back to the default group.

The only relationship between an allocation and its parent or sub-allocation concerns CPU bandwidth. That is, although an allocation and its parent or sub-allocation share the same CPU bandwidth, the processes matched to these allocations are different.

**Priority-order Chain**

You can use resource sub-allocation to create a specific type of resource-allocation hierarchy called a priority-order chain.

You can configure a single resource allocation with multiple child resource allocations, each with its own relative allocation. Doing so, you can create a hierarchy.

To create the priority-order chain, you order the sub-allocations in increasing priority as you move down the hierarchy. That is, the lower priority resource allocation should be matched at the higher levels, and the higher-priority resource allocation should be matched at the lower levels of the hierarchy.

This is illustrated in the following example:

- A resource-allocation named *ra0* is created with a process-matching criterion named *pmc0* and allocated 30 percent of CPU.
- A sub-allocation named *sub_ra1* is created for a different process-matching criterion, such as *pmc1*; it is allocated 20 percent of CPU.
- In turn, *sub_ra1* has a sub-allocation named *sub_ra2* with a different process-matching criterion, such as *pmc2*, and is allocated 30 percent of CPU.
- *sub_ra2* has a sub-allocation named *sub_ra3* with yet another process-matching criterion, such as *pmc3*, and is allocated 40 percent of CPU.
- The resource sub-allocations are numbered in an increasing order that corresponds to their preferred relative priority. Each sub-allocation and associated process-matching criterion have a higher precedence than their associated parent resource allocation, which is described as follows:
    - *sub_ra3* has the highest relative priority and will consume up to 40 percent of the total CPU bandwidth allocated in the policy.
    - *sub_ra2* can consume 30 percent of the remaining total bandwidth allocated that was not consumed by *sub_ra3*.
    - *sub_ra1* can consume 20 percent of the remaining total bandwidth allocated in the policy that was not consumed by *sub_ra2* and *sub_ra3*.
    - *ra0* has the lowest relative priority and will be able to consume only 10 percent of the total bandwidth allocated in the policy, assuming there is any that is unused by its sub-allocations, because the remaining 90 percent of the total bandwidth available to the entire resource allocation is allocated to its sub-allocations.

If any of the sub-allocations cannot use all of their allocation, the unused percentage is returned to the parent allocation. This process continues until either an allocation can use the unused allocation, or the parent allocation is the default group. In the latter case, the unused allocation is redistributed in the usual procedure as defined by the standard allocation algorithm.

**Re-allocation**

The amount of CPU bandwidth a process consumes does not always match its allocation. If a process cannot consume any part of its CPU bandwidth allocation, the unconsumed portion is re-allocated to other processes that can consume it. This is known as resource re-allocation.

The resource re-allocation sequence occurs as follows:

- The unconsumed allocation is passed from the process to its parent process-matching criterion and any sibling process-matching criteria that might exist. Other processes matched by either the parent or sibling process-matching criteria can then try to use the unconsumed allocation. The unconsumed allocation is redistributed to the other processes in proportion to their CPU consumption targets, as specified in the resource-allocation policy. For example, if a policy has resource allocations with process-matching criterion A allocated 20 percent and B allocated 40 percent, it would also implicitly have the default group at 40 percent. If the default group is only able to consume 20 percent, but A and B could each consume more than their allocation, then the other 20 percent not used by default would be reallocated to A and B proportionally. Process-matching criterion A would get one-third of the 20 percent, and B would get the other two-thirds.
- If the processes matched by either the parent or sibling process-matching are unable to use all of the unconsumed re-

allocation, the remaining unused portion is passed to the parent of the parent process-matching criterion.

- The re-allocation sequence continues until the re-allocation is consumed or the process-matching criterion has no parent. In the latter case, the unconsumed allocation is distributed among all top-level process-matching criteria, including the default group.

When a process attempts to use more than the allocated CPU bandwidth, Windows System Resource Manager lowers the priority of that process. This allows competing processes to use the CPU bandwidth equal to the difference between all remaining CPU bandwidth and the CPU bandwidth allocated to the process.

If a process is using less bandwidth than its allocated amount, the extra bandwidth is re-allocated to other managed processes. This occurs according to the default bandwidth allocation rules.

Because Windows System Resource Manager attempts to meet the CPU needs of running processes, it does not prevent an application from using unconsumed CPU bandwidth that no other process is using. Instead, a process is limited to its allocation only if there is not enough CPU bandwidth to satisfy the CPU demands of every running process.

**CPU Consumption Target Interactions**

In the following figure, a resource allocation policy is running and includes a CPU consumption target. To achieve the target CPU consumption, the Windows System Resource Manager service sends information to the kernel scheduler to control how managed processes are scheduled for CPU resources.

**Interactions that Occur When a CPU consumption Target Is Set**




**How Memory Resources Are Managed**

In Windows System Resource Manager, memory-allocation policies are configurable properties in a resource allocation. Memory-allocation policies limit the amount of working set memory and committed memory processes are allowed to consume.

Limits are applied on a per-process basis. For example, if you create a resource allocation that specifies a working set limit of 10 megabytes (MB) and a process-matching criterion that matches six processes, the limit of 10 MB is applied individually to each of the six processes.

**Working Set Memory**

You can set an upper limit on the process working set of the resource-allocation policy's process groups. Through the memory manager, Windows System Resource Manager prevents the working set of a process from exceeding the limit defined in the resource allocation. If the limit is reached, subsequent memory allocations do not fail, but existing pages in the working set of the process are replaced. This meets the memory limit without inducing any failures in the running application. Applications are generally unaware that the memory manager is reusing physical memory because applications are running in a virtual memory space.

**Committed Memory**
You can set an upper limit on the committed memory that a process consumes. Typically, if there is a constant increase in the committed memory a process consumes, it is due to a memory leak. By setting a limit on the amount of committed memory a process consumes, you can have Windows System Resource Manager intervene if a memory leak occurs. When the limit is reached, Windows System Resource Manager can either terminate the process or log an error in the event log.

The Windows System Resource Manager service maintains the committed memory limits in the current active resource allocation policy and actively monitors the committed memory usage of each matching process to determine when the limit is reached.

**Memory Limit Interactions**
Windows System Resource Manager enforces limits on committed memory and working set memory for processes. Windows System Resource Manager is not involved in the actual allocation of memory to a process, other than to signal the memory manager when there are memory limits in the currently active policy to be enforced. If the memory manager receives a signal from Windows System Resource Manager that a process is exceeding its memory limit, it will not allocate additional memory to a process.

In the following figure, a resource allocation policy is running and includes a memory limit.

**Interactions That Occur When a Memory Limit Is Set**




**Network Ports Used by Windows System Resource Manager**
Windows System Resource Manager uses only the ports already used by the Distributed Common Object Model (DCOM) service running on the computer. For more information about the ports used by DCOM, see "Using Distributed COM with Firewalls" at the Microsoft Web site.

# Networking Collection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Networking Collection**
Microsoft Windows Server 2003 offers a wide variety of technologies to meet the complex needs of today's connected environments. The networking technologies in Windows Server 2003 are designed to support all types of businesses, from small-inter-office network configurations to the largest enterprise solutions. The Windows Server 2003 networking technologies that are described in the Networking Collection are grouped into the following sub-collections:

- Core Networking Technologies
- Routing Technologies
- Network Monitoring Technologies
- Name Resolution Technologies
- Network Configuration Technologies
- Network Access Technologies
- Network Interoperability Technologies
- Remote Access Technologies
- Data Streaming and E-mail Technologies
- Network Security Technologies

**Core Networking Technologies**
The sub-collection "Core Networking Technologies" contains information about the TCP/IP protocol suite that is installed by default with Windows Server 2003. TCP/IP is an industry-standard suite of protocols designed to enable enterprise-wide network communications. There are two versions of TCP/IP that are supported by Windows Server 2003:

- TCP/IP with Internet Protocol version 4 (IPv4)
- TCP/IP with Internet Protocol version 6 (IPv6)

**IPv4**
IPv4 is a suite of protocols and standards based on the original IP specification described in RTF 791 in the IETF RFC Database and is in widespread use today on the Internet and on private networks. IPv4 has a relatively small address space that is quickly being depleted as Internet use expands. The need for more IP addresses and support for newer networking technologies are motivating factors in the adoption of IPv6.

**IPv6**
IPv6 is a suite of protocols and standards that supports a much larger address space than IPv4. IPv6 has 128-bit (16-byte) source and destination IP addresses. In contrast, IPv4 has 32-bit (4 bytes) source and destination IP addresses. IPv6 has many other security and efficiency improvements.

**Routing Technologies**
The sub-collection "Routing Technologies" contains information about the routing technologies that are supported by Windows Server 2003. Routing technologies manage the flow of data between network segments, also known as subnets. These routing technologies include:

- Unicast routing

- Multicast routing
- Network address translation (NAT)

**Unicast Routing**
Unicast routing forwards packets from one host to another host using the unicast destination IP address. This topic describes the IPv4 unicast routing technologies included in the Routing and Remote Access service

**Multicast Routing**
Multicast IP routing forwards packets from one host to multiple hosts using the multicast destination IP address. This topic describes the IPv4 multicast forwarding technologies included in the TCP/IP (IPv4) protocol and the Routing and Remote Access service.

**NAT**
Network address translation (NAT) functionality is part of the Routing and Remote Access service. A server that has been configured as a NAT-enabled router, with a private IP address and at least one public IP address, translates the private addresses (and TCP or UDP port numbers) in outgoing packets. The outgoing packets can then be forwarded to a resource on a public network, such as the Internet. The NAT-enabled router also translates incoming traffic and forwards the incoming packets to the appropriate address on the private network.

**Network Monitoring Technologies**
The sub-collection "Network Monitoring Technologies" contains information about network data-gathering services that are supported by Windows Server 2003. These services include:

- Computer Browser service
- Simple Network Management Protocol (SNMP)

**Computer Browser Service**
Computer Browser service is used by Windows-based computers to collect, distribute, and obtain the browser list, a collection of information about workgroups, domains, and the computers within them. The browser list is available to programs that request it, such as My Network Places.

**SNMP**
Simple Network Management Protocol (SNMP) is a network management protocol and infrastructure widely used in TCP/IP networks to remotely monitor, configure, and troubleshoot network resources from a centrally located SNMP management system.

**Name Resolution Technologies**
The sub-collection "Name Resolution Technologies" contains information about Windows Server 2003 support for name resolution services that include:

- Domain Name System (DNS)
- Windows Internet Name Service (WINS)

**DNS**
Domain Name System (DNS) is a domain name resolution service that allows client computers to register their domain names and IP addresses in a dynamic, distributed database and to resolve the domain names of network resources to their IPv4 or IPv6 addresses.

**WINS**
WINS is a NetBIOS name resolution service that allows client computers to register their NetBIOS names and IP addresses in a dynamic, distributed database and to resolve the NetBIOS names of network resources to their IPv4 addresses.

**Network Configuration Technologies**
The sub-collection "Network Configuration Technologies" contains information about Windows Server 2003 support for network configuration services that include:

- Dynamic Host Configuration Protocol (DHCP)

- Windows Time Service

**DHCP**
Dynamic Host Configuration Protocol (DHCP) is used to automate the configuration of hosts on a network. DHCP clients request and DHCP servers allocate IP addresses and network configuration settings. DHCP servers maintain a database of IP address ranges, configuration settings, and the current set of allocated IP addresses.

**Windows Time Service**
Windows Time service uses the Network Time Protocol (NTP) to enable date and time synchronization among network resources on a Windows-based network. The Windows Time service can be used to automatically synchronize the time of a local computer with other computers on a network.

**Network Access Technologies**
The sub-collection "Network Access Technologies" contains information about Windows Server 2003 support for secure network access solutions including:

- 802.11 Wireless
- Virtual Private Network (VPN)
- Internet Authentication Service (IAS)
- Connection Manager

**802.11 Wireless**
The IEEE 802.11 is a set of standards for wireless LAN connectivity that allows the extension of a wired LAN to include wireless, mobile clients. For secure wireless connections, 802.11-based wireless networks use authentication (the IEEE 802.1X standard) and encryption (either Wired Equivalent Privacy or Wi-Fi Protected Access).

**VPN**
Virtual Private Network (VPN) technologies allow the creation of secure point-to-point logical links across shared or public networks such as the Internet. Remote access VPN connections allow VPN clients to securely connect to a private network over the Internet. Site-to-site VPN connections allow routers to connect the sites of an organization together across the Internet.

**IAS**
Internet Authentication Service (IAS) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy. IAS RADIUS servers perform centralized connection authentication and authorization. IAS RADIUS servers also provide network access accounting services. IAS RADIUS proxies forward authentication and accounting messages from RADIUS clients to RADIUS servers, and from RADIUS servers to RADIUS clients.

**Connection Manager**
Connection Manager is a set of components that is used to create and maintain customized remote access client service profiles that enable users to connect to a network using administrator-defined connection parameters.

**Network Interoperability Technologies**
The sub-collection "Network Interoperability Technologies" contains information about technologies that support the core functionality that enables remote administration and directory service synchronization including:

- Remote Procedure Call (RPC)
- Services for NetWare (SFN)

**RPC**
Remote Procedure Call (RPC) is an interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality residing in a different process. That process can be on the same computer, on the local area network (LAN), or across the Internet.

**Services for NetWare**
Services for Netware (SFN) provides directory object migration and synchronization between computers that are running Windows Server 2003 with Active Directory and computers that are running Novell's NetWare operating system with NDS, eDirectory, and Bindery directory services.

**Remote Access Technologies**
The sub-collection “Remote Access Technologies” contains information about Windows Server 2003 support for remote access solutions including:

- Dial-up Remote Access
- Telnet
- Terminal Services

**Dial-up Remote Access**
Dial-up Remote Access enables remote access clients to connect to a network. Remote access clients use the available telecommunications infrastructure to create a temporary physical or virtual circuit to a port on a remote access server that is connected to a network. After the connection between the remote access client and the remote access server is established, the remote access server forwards packets between the remote access client and the network.

**Telnet**
Telnet is a protocol that enables remote connections from a remote access client to a host. A local command prompt on a remote access client can be used to run command-line programs, shell commands, and scripts in a remote command console session.

**Terminal Services**
Terminal Services delivers the Windows Server 2003 desktop and the latest Windows-based applications to a wide variety of desktops, including those that cannot normally run Windows. Terminal Services uses terminal emulation to host multiple, simultaneous client sessions. Terminal Services enables process scheduling, application-sharing, and multi-user functionality.

**Data Streaming and E-mail Technologies**
The sub-collection “Data Streaming and E-mail Technologies” contains information about Windows Server 2003 support for technologies that provide a baseline network infrastructure for streaming multimedia data. This sub-collection also contains information about Windows Server 2003 support for the transfer and retrieval of e-mail. The technologies in this sub-collection include:

- Asynchronous Transfer Mode (ATM)
- Quality of Service (QoS)
- POP3 Mail Server

**ATM**
Asynchronous Transfer Mode (ATM) is connection-oriented networking technology that establishes a defined network path, called a virtual channel, between two endpoints for the purpose of sending isochronous (time-dependent) network traffic. When the virtual channel is established, ATM also negotiates a QoS contract for the virtual channel. The QoS contract between the endpoints defines the bandwidth, the maximum transit delay, and an acceptable variance in the transit delay.

**QoS**
Quality of Service (QoS) is a set of service requirements that a network must meet to ensure an adequate service level for data transmission. QoS allows real-time programs to make the most efficient use of network bandwidth.

**POP3 Mail Server**
The POP3 service is a standard messaging protocol that is supported in Windows Server 2003. POP3 provides e-mail transfer and retrieval services.

**Network Security Technologies**
The sub-collection "Network Security Technologies" contains information about Windows Server 2003 network security technologies, which can help protect your network from intrusions, malicious software (malware), viruses, worms, and attacks that rely on unsolicited incoming traffic.

**Windows Firewall**
Windows Firewall is a stateful host-firewall technology that inspects and filters all IP version 4 (IPv4) and IP version 6 (IPv6) network traffic. As a stateful firewall, Windows Firewall tracks the state of each network connection and determines whether unsolicited incoming traffic is allowed or dropped. Windows Firewall blocks unsolicited incoming traffic unless the traffic is a response to a request by the host (solicited traffic) or it is specifically allowed (in which case, it has been added to the Windows

Firewall exceptions list). Exceptions can be specified by program name, UDP or TCP port number, or system service name. Aside from a few Internet Control Message Protocol (ICMP) messages, Windows Firewall allows all outgoing traffic.

# Core Networking Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In Microsoft Windows Server 2003, core networking tasks are accomplished by using TCP/IP. TCP/IP consists of a suite of protocols, of which TCP and IP are two. This suite of protocols was originally designed to solve a communications problem among the branches of the United States military. In the 1960s, each of the military branches obtained bids from different vendors to provide computer systems for their branch. The Army chose Digital Equipment Corporation (DEC), the Air Force chose International Business Machines (IBM), and the Navy chose Unisys. Soon after, the military branches discovered that they needed their computer systems to communicate with each other to facilitate coordination between the military branches. The Department of Defense (DoD) launched a research project in 1969 to connect the systems of various vendors together to form a network of networks. The DoD developed TCP/IP with IP version 4 (IPv4) to connect this network of networks — the collection of networks now known as the Internet. TCP/IP is still used to connect business networks across the world.

**Note**

- The word internetwork refers to multiple TCP/IP networks connected with routers. The Internet is a worldwide public IP internetwork. An intranet is a private IP internetwork.

IPv4 proved to be robust, easily implemented and interoperable. It has also stood the test of scalability from an internetwork to today's global Internet. However, the initial design did not fully allow for the following:

- The exponential growth of the Internet and the impending depletion of the IPv4 address space.

  The scarcity of IPv4 addresses has forced some organizations to use a network address translator to map multiple private addresses to a single public IP address. Network address translators promote reuse of the private address space, but they do not support standards-based network layer security or the correct mapping of all higher layer protocols. Moreover, the translators can create problems when connecting two organizations that use the private address space.

  Additionally, the rising prominence of Internet-connected devices and appliances ensures that the public IPv4 address space will eventually be depleted.

- The requirement for security at the IP level.

- Private communication over a public medium, such as the Internet, requires encryption services that protect the data being sent from being viewed or modified in transit. Although a standard now exists for providing security for IPv4 packets, known as Internet Protocol Security, or IPSec, this standard is optional.

- The growth of the Internet and the ability of Internet backbone routers to maintain large routing tables.

  Because of the way that IPv4 network IDs have been and are currently allocated, there are often over 85,000 routes in the routing tables of Internet backbone routers. The current IPv4 Internet routing infrastructure is a combination of both flat and hierarchical routing.

- The need for better support for real-time delivery of data, or Quality of Service (QoS).

  QoS standards exist for IPv4, but real-time traffic support relies on the IPv4 Type of Service (ToS) field and the identification of the payload, typically using either a User Datagram Protocol (UDP) or TCP port. Unfortunately, the IPv4 ToS field has limited functionality, and over time there have been various local interpretations. In addition, payload identification using a TCP and User Datagram Protocol (UDP) port is not possible when the IPv4 packet payload is encrypted.

To address these and other concerns, the Internet Engineering Task Force (IETF) developed IP version 6 (IPv6). IPv6, previously called IP next generation, incorporates the concepts of many proposed methods for updating the IPv4 protocol. The design of IPv6 is intentionally targeted for minimal impact on upper and lower layer protocols by avoiding the random addition of new features.

Windows Server 2003 supports TCP/IP with IPv4 and IPv6.

**Core Networking Components**
The core networking protocol, TCP/IP, which is installed by default on computers that run Windows Server 2003 is an industry-standard suite of protocols designed for large internetworks spanning wide area network (WAN) links. TCP/IP in Windows Server 2003 was designed to make it easy to integrate Microsoft systems into large-scale corporate, government and public networks, and to provide the ability to operate over those networks in a more secure manner. There are two versions of TCP/IP available for use with Windows Server 2003: TCP/IP version 4 (with IPv4) and TCP/IP version 6 (with IPv6).

**TCP/IP (with IPv4)**
TCP/IP in Windows Server 2003 enables enterprise networking and connectivity that have the following features:

- A standard, routable enterprise networking protocol that is the most complete and accepted protocol available. All modern, network operating systems offer TCP/IP support, and most large networks rely on TCP/IP for much of their network traffic.

- A technology for connecting dissimilar systems. Many standard connectivity tools are available to access and transfer data between dissimilar systems, including File Transfer Protocol (FTP) and Telnet, a terminal emulation protocol. Several of these standard tools are included with Windows Server 2003.

- A robust, scalable, cross-platform client/server framework. TCP/IP in Windows Server 2003 offers the Windows Sockets (Winsock) interface, which is ideal for developing client/server applications that can run on Winsock-compliant TCP/IP protocol implementations from other vendors.

- A method of gaining access to the Internet. The Internet consists of thousands of networks worldwide, connecting research facilities, universities, libraries, and private companies.

For more information about Windows Server 2003 TCP/IP, see TCP/IP Technical Reference.

**IPv6**
The following features of the IPv6 protocol overcome the limitations of IPv4:

- New header format

- Large address space

- Efficient and hierarchical addressing and routing infrastructure

- Stateless and stateful address configuration

- Built-in security measures

- Better support for QoS

- New protocol for neighboring node interaction

- Extensibility

IPv6 includes new capabilities such as scoped addresses, stateless autoconfiguration, lowering the complexity and management burden, and mandatory IPSec, permitting end-to-end data authentication, data integrity, and privacy of connections. In addition to the new capabilities, IPv6 brings back the capability of end-to-end communications; making networking applications simpler as the network again becomes transparent

### Note

- In IPv4, end-to-end communications have been largely replaced by a client/server model because network address translators are better behaved for client/server communications. The translators also extend the lifetime of IPv4 by enabling multiple operating systems to share a single, public IPv4 address.

IPv6 supports always-on technologies such as Digital Subscriber Line (DSL) and cable modems, personal data assistants (PDAs), and cellular Mobile Information Appliances. Windows Server 2003 supports the capability to send IPv6 traffic over existing IPv4 networks.

For more detailed information about these features, see IPv6 Technical Reference.

**Core Networking Scenarios**

Computers that run Windows operating systems use TCP/IP to communicate with other computers in corporate intranets, and across the Internet. As such, TCP/IP is an essential component of any Windows network configuration. Windows Server 2003 is the first release of the Windows operating system that includes IPv6. Ideally, IPv6 is used in a pure environment, meaning an environment where IPv6 is the exclusive Internet protocol used between computers. Currently, however, pure IPv6 environments are attainable only when you use computers running Windows operating systems that support IPv6 and routers that support IPv6 routing. As IPv6 replace IPv4, pure IPv6 environments will eventually replace IPv4. Until that occurs, the transition technologies described in this technical reference can be used to facilitate coexistence and provide a migration path from IPv4 to IPv6.

# TCP/IP Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Transmission Control Protocol/Internet Protocol, or TCP/IP, is a protocol suite developed for communications over, often dissimilar, networks. Microsoft Windows Server 2003 provides extensive support for the TCP/IP suite for IP version 4 (IPv4), as both a protocol and a set of services for connectivity and management of IP internetworks. Knowledge of the basic concepts of TCP/IP is an absolute requirement for the proper understanding of the configuration, deployment, and troubleshooting of Windows Server 2003.

This subject describes how TCP/IP with IPv4 relates to other networking protocols, the functions it performs, how addresses are structured and assigned, and how packets are structured and routed.

**In this subject**


- What Is TCP/IP?
- How TCP/IP Works
- TCP/IP Tools and Settings

# What Is TCP/IP?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is TCP/IP?**

**In this section**


- Windows Server 2003 TCP/IP
- TCP/IP Core Protocols
- TCP/IP Services
- Related Information


Transmission Control Protocol/Internet Protocol, or TCP/IP, is an industry-standard suite of protocols designed for large internetworks. TCP/IP, which was developed in 1969 by the U.S. Department of Defense (DoD) Advanced Research Projects Agency (DARPA), is the result of a resource-sharing experiment called Advanced Research Projects Agency Network (ARPANET). The TCP/IP protocol was developed to provide high-speed communication network links. Since 1969, ARPANET has grown into a worldwide community of networks known as the Internet.

Before TCP/IP, there was no way for computers to communicate easily and securely on public networks. Windows Server 2003 TCP/IP was designed to make it easy to integrate Microsoft systems into large-scale corporate, government, and public networks, and to provide the ability to operate over those networks in a secure manner. The Windows Server 2003 TCP/IP protocol is installed by default and, unlike previous versions of Windows, cannot be uninstalled. However, you can reset the TCP/IP configuration to a default state with the **netsh interface ip reset** command.

The Windows TCP/IP suite contains *core protocol elements*, *services*, and the *interfaces* between them. The Transport Driver Interface (TDI) and the Network Device Interface Specification (NDIS) are public, and their specifications are available from Microsoft. In addition, there are a number of higher-level interfaces available to user-mode applications. The most commonly used are Windows Sockets and NetBIOS.

**Windows Server 2003 TCP/IP**

Windows Server 2003 TCP/IP enables enterprise networking and connectivity. Adding TCP/IP to a Windows Server 2003 configuration offers the following advantages:

- A standard, routable enterprise networking protocol that is the most complete and accepted protocol available. All modern network operating systems offer TCP/IP support, and most large networks rely on TCP/IP for much of their network traffic.
- A technology for connecting dissimilar systems. Many standard connectivity utilities are available to access and transfer data between dissimilar systems, including File Transfer Protocol (FTP) and Telnet, a terminal emulation protocol. Several of these standard utilities are included with Windows Server 2003.
- A robust, scalable, cross-platform client/server framework. Windows Server 2003 TCP/IP offers the Windows Sockets interface, which is ideal for developing client/server applications that can run on Windows Sockets–compliant TCP/IP protocol implementations from other vendors.
- A method of gaining access to the Internet. The Internet consists of thousands of networks worldwide, connecting research facilities, universities, libraries, private companies, and individuals.

**Note**

- The term Internet refers to the worldwide public Internet. An intranet refers to a private IP-based internetwork.

**Support for Standard Features**

Windows Server 2003 TCP/IP supports the following standard features:

- Ability to bind to multiple network adapters with different media types.
- Logical and physical multihoming.
- Internal IP routing capability.
- Internet Group Management Protocol (IGMP) (IP multicasting).
- Duplicate IP address detection.
- Multiple configurable default gateways.
- Dead gateway detection.
- Automatic Path Maximum Transmission Unit (PMTU) discovery.
- Internet Protocol security (IPSec).
- Quality of Service (QoS).
- Asynchronous Transfer Mode Tutorial (ATM) services.
- Virtual Private Networks (VPNs) with the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol with IPSec (L2TP/IPSec).

**Windows Server 2003 TCP/IP Features**

The features and improvements of TCP/IP for Windows Server 2003 include the following:

- Windows Server 2003 and Windows XP (Service Pack 1 or later) now include a production-quality IPv6 protocol stack. For more information about IPv6, see IPv6 Technical Reference.
- Auto-negotiation of RFC 1323 options (window scaling and TCP timestamps).
- Default support of network interface cards providing large send offload (LSO) and checksum offload.
- Internet Group Management Protocol (IGMP) version 3.
- Reliable multicast with Pragmatic General Multicast (PGM).
- Alternate configuration of a static IP address configuration.
- Automatic determination of the interface-related and default route metrics.

**TCP/IP Standards**

The standards for TCP/IP are published in a series of documents that are called Requests for Comments (RFCs). RFCs describe the internal workings of the Internet. Some RFCs describe network services or protocols and their implementations, whereas others summarize policies. TCP/IP standards are always published as RFCs, although not all RFCs specify standards.

TCP/IP standards are not developed by a committee, but rather by consensus. Anyone can submit a document for publication as an RFC. Documents are reviewed by a technical expert, a task force, or the RFC editor, and then assigned a status. The status specifies whether a document is being considered for becoming a standard.

There are five RFC status assignments, shown below.

**RFC Status Assignments**

| Status | Description |
|---|---|
| Required | Must be implemented on all TCP/IP-based hosts and gateways. |
| Recommended | Encouraged that all TCP/IP-based hosts and gateways implement the RFC specifications. Recommended RFCs are usually implemented. |
| Elective | Implementation is optional. Application has been agreed to but is not a requirement. |
| Limited Use | Not intended for general use. |
| Not Recommended | Not recommended for implementation. |

If a document is being considered for becoming a standard, it goes through stages of development, testing, and acceptance known as the Internet Standards Process. These stages are formally labeled maturity levels. The following table lists the three maturity levels for Internet Standards.

**Maturity Levels for Internet Standards**

| Maturity Level | Description |
|---|---|
| Proposed Standard | This specification is generally stable, has resolved known design issues, is believed to be well understood, has received significant community review, and appears to hold enough community interest to be considered valuable. |
| Draft Standard | This must be well understood and known to be quite stable, both in its semantics and as a basis for developing an implementation. |
| Internet Standard | This specification is characterized by a high degree of technical maturity, and it is generally agreed that the specified protocol or service provides significant benefit to the Internet community. |

When a document is published, it is assigned an RFC number. If changes are required, a new RFC is published with a new number. The original RFC is never updated. Therefore, it is important to verify that you have the most recent RFC on a particular topic.

RFCs can be obtained in several ways. To obtain any RFC, or a full and current indexed listing of all RFCs published to date, see the IETF RFC Database.

#### TCP/IP Core Protocols

The TCP/IP protocol component that is installed in your network operating system is a series of interconnected protocols called the TCP/IP core protocols. All other applications and other protocols in the TCP/IP protocol suite rely on the basic services provided by the following protocols: IP, ARP, ICMP, IGMP, TCP, and UDP. For more information about these protocols, see How TCP/IP Works.

#### TCP/IP Services

The Windows Server 2003 operating system provides the following TCP/IP-related services:

- Dynamic Host Configuration Protocol (DHCP) client and server and DHCP Relay Agent (with the Routing and Remote Access service).
- In the absence of a DHCP server, Automatic Private IP Addressing (APIPA) or an alternate static IP address configuration is used.
- Windows Internet Name Service (WINS), a NetBIOS name client and server.
- Domain Name System (DNS) client and server, including support for DNS dynamic updates.
- Dial-up support using the Point-to-Point Protocol (client and server) and Serial Line Internet Protocol (client only).
- PPTP and L2TP/IPSec, used for remote access and site-to-site VPN connections.
- TCP/IP network printing (client only with the Lpr.exe and Lpq.exe tools).

- SNMP agent.
- NetBIOS interface.
- Network Location Service.
- Windows Sockets version 2 (Winsock2) interface.
- Remote Procedure Call (RPC) support.
- Network Dynamic Data Exchange (NetDDE).
- Computer browsing (My Network Places) across IP routers.
- Reliable multicast with the Pragmatic General Multicast (PGM) protocol.
- Basic TCP/IP connectivity utilities, including: **finger, ftp, rcp, rexec, rsh, telnet,** and **tftp.**
- Server and client software for simple network protocols, including: Character Generator, Daytime, Discard, Echo, and Quote of the Day.
- Routing Information Protocol (RIP) listener (for Windows XP Professional) and RIP and Open Shortest Path First (OSPF) (with the Routing and Remote Access service).
- Network Address Translator (NAT) capabilities using either the Internet Connection Service (ICS) or the NAT/Basic Firewall routing protocol component of the Routing and Remote Access Service.
- Stateful firewalling capabilities using either the Internet Connection Firewall (ICF) or the NAT/Basic Firewall routing protocol component of the Routing and Remote Access service.
- Multicast forwarding and IGMP router and proxy capabilities with the Routing and Remote Access Service.
- TCP/IP management and diagnostic tools, including: **arp, ipconfig, nbtstat, netsh, netstat, ping, pathping, route, nslookup,** and **tracert**.

**Related Information**
For more information about RFCs, see the IETF RFC Database.

# How TCP/IP Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How TCP/IP Works**

**In this section**


- TCP/IP Protocol Architecture
- IPv4 Addressing
- Name Resolution
- IPv4 Routing
- Physical Address Resolution
- Related Information


TCP/IP for IP version 4 (IPv4) is a networking protocol suite that Microsoft Windows uses to communicate over the internet with other computers. It interacts with Windows naming services like DNS and security technologies, such as IPsec primarily, as these help facilitate the successful and secure transfer of IP packets between machines.

Ideally, TCP/IP is used whenever Windows-based computers communicate over networks.

This subject describes the components of the TCP/IP Protocol Suite, the protocol architecture, which functions TCP/IP performs, how addresses are structured and assigned, and how packets are structured and routed.

Microsoft Windows Server 2003 provides extensive support for the Transmission Control Protocol/Internet Protocol (TCP/IP) suite, as both a protocol and a set of services for connectivity and management of IP internetworks. Knowledge of the basic concepts of TCP/IP is an absolute requirement for the proper understanding of the configuration, deployment, and troubleshooting of IP-based Windows Server 2003 and Microsoft Windows 2000 intranets.

**TCP/IP Protocol Architecture**

TCP/IP protocols map to a four-layer conceptual model known as the DARPA model, named after the U.S. government agency that initially developed TCP/IP. The four layers of the DARPA model are: Application, Transport, Internet, and Network Interface. Each layer in the DARPA model corresponds to one or more layers of the seven-layer Open Systems Interconnection (OSI) model.

The following figure shows the TCP/IP protocol architecture.

**TCP/IP Protocol Architecture**




**Note**

- The architectural diagram above corresponds to the Internet Protocol TCP/IP and does not reflect IP version 6. To see a

TCP/IP architectural diagrm that includes IPv6, see *How IPv6 Works* in this technical reference.

**Network Interface Layer**
The Network Interface layer (also called the Network Access layer) handles placing TCP/IP packets on the network medium and receiving TCP/IP packets off the network medium. TCP/IP was designed to be independent of the network access method, frame format, and medium. In this way, TCP/IP can be used to connect differing network types. These include local area network (LAN) media such as Ethernet and Token Ring and WAN technologies such as X.25 and Frame Relay. Independence from any specific network media allows TCP/IP to be adapted to new media such as asynchronous transfer mode (ATM).

The Network Interface layer encompasses the Data Link and Physical layers of the OSI model. Note that the Internet layer does not take advantage of sequencing and acknowledgment services that might be present in the Network Interface layer. An unreliable Network Interface layer is assumed, and reliable communication through session establishment and the sequencing and acknowledgment of packets is the function of the Transport layer.

**Internet Layer**
The Internet layer handles addressing, packaging, and routing functions. The core protocols of the Internet layer are IP, ARP, ICMP, and IGMP.

- The Internet Protocol (IP) is a routable protocol that handles IP addressing, routing, and the fragmentation and reassembly of packets.
- The Address Resolution Protocol (ARP) handles resolution of an Internet layer address to a Network Interface layer address, such as a hardware address.
- The Internet Control Message Protocol (ICMP) handles providing diagnostic functions and reporting errors due to the unsuccessful delivery of IP packets.
- The Internet Group Management Protocol (IGMP) handles management of IP multicast group membership.

The Internet layer is analogous to the Network layer of the OSI model.

**Transport Layer**
The Transport layer (also known as the Host-to-Host Transport layer) handles providing the Application layer with session and datagram communication services. The core protocols of the Transport layer are Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP).

- TCP provides a one-to-one, connection-oriented, reliable communications service. TCP handles the establishment of a TCP connection, the sequencing and acknowledgment of packets sent, and the recovery of packets lost during transmission.
- UDP provides a one-to-one or one-to-many, connectionless, unreliable communications service. UDP is used when the amount of data to be transferred is small (such as data that fits into a single packet), when you do not want the overhead of establishing a TCP connection, or when the applications or upper layer protocols provide reliable delivery.

The TCP/IP Transport layer encompasses the responsibilities of the OSI Transport layer.

**Application Layer**
The Application layer lets applications access the services of the other layers and defines the protocols that applications use to exchange data. There are many Application layer protocols and new protocols are always being developed.

The most widely known Application layer protocols are those used for the exchange of user information:

- The Hypertext Transfer Protocol (HTTP) is used to transfer files that make up the Web pages of the World Wide Web.
- The File Transfer Protocol (FTP) is used for interactive file transfer.
- The Simple Mail Transfer Protocol (SMTP) is used for the transfer of mail messages and attachments.
- Telnet, a terminal emulation protocol, is used for logging on remotely to network hosts.

Additionally, the following Application layer protocols help facilitate the use and management of TCP/IP networks:

- The Domain Name System (DNS) is used to resolve a host name to an IP address.
- The Routing Information Protocol (RIP) is a routing protocol that routers use to exchange routing information on an IP internetwork.
- The Simple Network Management Protocol (SNMP) is used between a network management console and network devices (routers, bridges, intelligent hubs) to collect and exchange network management information.

Examples of Application layer interfaces for TCP/IP applications are Windows Sockets and NetBIOS. Windows Sockets provides a standard application programming interface (API) under Windows Server 2003. NetBIOS is an industry-standard interface for accessing protocol services such as sessions, datagrams, and name resolution. More information on Windows Sockets and NetBIOS is provided later in this chapter.

The TCP/IP Application layer encompasses the responsibilities of the OSI Session, Presentation, and Application layers.

**TCP/IP Core Protocols**

The TCP/IP protocol component that is installed in your network operating system is a series of interconnected protocols called the core protocols of TCP/IP. All other applications and other protocols in the TCP/IP protocol suite rely on the basic services provided by the following protocols: IP, ARP, ICMP, IGMP, TCP, and UDP.

**IP**

IP is a connectionless, unreliable datagram protocol primarily responsible for addressing and routing packets between hosts. Connectionless means that a session is not established before exchanging data. Unreliable means that delivery is not guaranteed. IP always makes a "best effort" attempt to deliver a packet. An IP packet might be lost, delivered out of sequence, duplicated, or delayed. IP does not attempt to recover from these types of errors. The acknowledgment of packets delivered and the recovery of lost packets is the responsibility of a higher-layer protocol, such as TCP. IP is defined in RFC 791.

An IP packet consists of an IP header and an IP payload. The following table describes the key fields in the IP header.

**Key Fields in the IP Header**

| IP Header Field | Function |
|---|---|
| Source Address | The IP address of the original source of the IP datagram. |
| Destination Address | The IP address of the final destination of the IP datagram. |
| Identification | Used to identify a specific IP datagram and to identify all fragments of a specific IP datagram if fragmentation occurs. |
| Protocol | Informs IP at the destination host whether to pass the packet up to TCP, UDP, ICMP, or other protocols. |
| Checksum | A simple mathematical computation used to verify the bit-level integrity of the IP header. |
| Time to Live (TTL) | Designates the number of network segments on which the datagram is allowed to travel before being discarded by a router. The TTL is set by the sending host and is used to prevent packets from endlessly circulating on an IP internetwork. When forwarding an IP packet, routers are required to decrease the TTL by at least one. |

**Fragmentation and reassembly**

If a router receives an IP packet that is too large for the network to which the packet is being forwarded, IP fragments the original packet into smaller packets that fit on the downstream network. When the packets arrive at their final destination, IP on the destination host reassembles the fragments into the original payload. This process is referred to as fragmentation and reassembly. Fragmentation can occur in environments that have a mix of networking media, such as Ethernet and Token Ring.

The fragmentation and reassembly works as follows:

- When an IP packet is sent by the source, it places a unique value in the **Identification** field.
- The IP packet is received at the router. The IP router notes that the maximum transmission unit (MTU) of the network

onto which the packet is to be forwarded is smaller than the size of the IP packet.

- IP divides the original IP payload into fragments that fit on the next network. Each fragment is sent with its own IP header that contains:
    - The original **Identification** field identifying all fragments that belong together.
    - The **More Fragments Flag** indicating that other fragments follow. The **More Fragments Flag** is not set on the last fragment, because no other fragments follow it.
    - The **Fragment Offset** field indicating the position of the fragment relative to the original IP payload.

When the fragments are received by IP at the remote host, they are identified by the **Identification** field as belonging together. The **Fragment Offset** field is then used to reassemble the fragments into the original IP payload.

**ARP**
When IP packets are sent on shared access, broadcast-based networking media — such as Ethernet or Token Ring — the media access control (MAC) address corresponding to a forwarding IP address must be resolved. ARP uses MAC-level broadcasts to resolve a known forwarding or next-hop IP address to its MAC address. ARP is defined in RFC 826.

**ICMP**
Internet Control Message Protocol (ICMP) provides troubleshooting facilities and error reporting for packets that are undeliverable. For example, if IP is unable to deliver a packet to the destination host, ICMP sends a Destination Unreachable message to the source host. The following table shows the most common ICMP messages.

**Common ICMP Messages**

| ICMP Message | Function |
|---|---|
| Echo Request | Troubleshooting message used to check IP connectivity to a desired host. The ping utility sends ICMP Echo Request messages. |
| Echo Reply | Response to an ICMP Echo Request. |
| Redirect | Sent by a router to inform a sending host of a better route to a destination IP address. |
| Source Quench | Sent by a router to inform a sending host that its IP datagrams are being dropped due to congestion at the router. The sending host then lowers its transmission rate. Source Quench is an elective ICMP message and is not commonly implemented. |
| Destination Unreachable | Sent by a router or the destination host to inform the sending host that the datagram cannot be delivered. |

The following table describes the most common ICMP Destination Unreachable ICMP messages.

**Common ICMP Destination Unreachable Messages**

| Destination Unreachable Message | Description |
|---|---|
| Host Unreachable | Sent by an IP router when a route to the destination IP address cannot be found. |
| Protocol Unreachable | Sent by the destination IP node when the **Protocol** field in the IP header cannot be matched with an IP client protocol currently loaded. |
| Port Unreachable | Sent by the destination IP node when the Destination Port in the UDP header cannot be matched with a process using that port. |
| Fragmentation Needed and DF Set | Sent by an IP router when fragmentation must occur but is not allowed due to the source node setting the **Don't Fragment (DF)** flag in the IP header. |

| | |
|---|---|
| Source Route Failed | Sent by an IP router when delivery of the IP packet using source route information (stored as source route option headers) fails. |

ICMP does not make IP a reliable protocol. ICMP attempts to report errors and provide feedback on specific conditions. ICMP messages are carried as unacknowledged IP datagrams and are themselves unreliable. ICMP is defined in RFC 792.

**IGMP**

Internet Group Management Protocol (IGMP) is a protocol that manages host membership in IP multicast groups on a network segment. An IP multicast group, also known as a host group, is a set of hosts that listen for IP traffic destined for a specific IP multicast address. IP multicast traffic is sent to a single MAC address but processed by multiple IP hosts. A specific host listens on a specific IP multicast address and receives all packets to that IP address.

The following are some of the additional aspects of IP multicasting:

- Host group membership is dynamic, hosts can join and leave the group at any time.
- A host group can be of any size.
- Members of a host group can span IP routers across multiple networks. This situation requires IP multicast support on the IP routers and the ability for hosts to register their group membership with local routers. Host registration is accomplished using IGMP.
- A host can send traffic to an IP multicast address without belonging to the corresponding host group.

For a host to receive IP multicasts, an application must inform IP that it will receive multicasts at a specified IP multicast address. If the network technology supports hardware-based multicasting, the network interface is told to pass up packets for a specific IP multicast address. In the case of Ethernet, the network adapter is programmed to respond to a multicast MAC address corresponding to the specified IP multicast address.

A host supports IP multicast at one of the following levels:

- Level 0: No support to send or receive IP multicast traffic.
- Level 1: Support exists to send but not receive IP multicast traffic.
- Level 2: Support exists to both send and receive IP multicast traffic. Windows Server 2003, Windows 2000, Microsoft Windows NT version 3.5 and later, and TCP/IP support level 2 IP multicasting.

The protocol to register host group information is IGMP, which is required on all hosts that support level 2 IP multicasting. IGMP packets are sent using an IP header.

IGMP messages take three forms.

- **Host Membership Report**. When a host joins a host group, it sends an IGMP Host Membership Report message to the all-hosts IP multicast address (224.0.0.1) or to the specified IP multicast address declaring its membership in a specific host group by referencing the IP multicast address. A host can also specify the specific sources from which multicast traffic is needed.
- **Host Membership Query**. When a router polls a network to ensure that there are members of a specific host group, it sends an IGMP Host Membership Query message to the all-hosts IP multicast address. If no responses to the poll are received after several polls, the router assumes no membership in that group for that network and stops advertising that multicast group information to other routers.
- **Group Leave**. When a host is no longer interested in receiving multicast traffic sent to a specific IP multicast address and it sent the last IGMP Host Membership Report message in response to an IGMP Host Membership Query, it sends an IGMP Group Leave message to the specific IP multicast address. Local routers verify that the host sending the IGMP Group Leave message is the last group member for that multicast address on that subnet. If no responses to the poll are received after several polls, the router assumes no membership in that group for that subnet and stops advertising that multicast group information to other routers.

For IP multicasting to span routers across an internetwork, multicast routing protocols are used by routers to communicate host group information so that each router supporting multicast forwarding is aware of which networks contain members of which host groups. IGMP is defined in RFCs 1112 and 2236.

**TCP**

TCP is a reliable, connection-oriented delivery service. The data is transmitted in segments. Connection-oriented means that a connection must be established before hosts can exchange data. Reliability is achieved by assigning a sequence number to each segment transmitted. An acknowledgment is used to verify that the data is received. For each segment sent, the receiving host must return an acknowledgment (ACK) within a specified period for bytes received. If an ACK is not received, the data is retransmitted. TCP is defined in RFC 793.

TCP uses byte-stream communications, wherein data within the TCP segment is treated as a sequence of bytes with no record or field boundaries. The following table describes the key fields in the TCP header.

**Key Fields in the TCP Header**

| Field | Function |
|---|---|
| Source Port | TCP port of sending host. |
| Destination Port | TCP port of destination host. |
| Sequence Number | Sequence number of the first byte of data in the TCP segment. |
| Acknowledgment Number | Sequence number of the byte the sender expects to receive next from the other side of the connection. |
| Window | Current size of a TCP buffer on the host sending this TCP segment to store incoming segments. |
| TCP Checksum | Verifies the bit-level integrity of the TCP header and the TCP data. |

**TCP ports**

A TCP port provides a specific location for delivery of TCP segments. Port numbers below 1024 are well-known ports and are assigned by the Internet Assigned Numbers Authority (IANA). The following table lists a few well-known TCP ports.

**Well-Known TCP Ports**

| TCP Port Number | Description |
|---|---|
| 20 | FTP (Data Channel) |
| 21 | FTP (Control Channel) |
| 23 | Telnet |
| 80 | HTTP used for the World Wide Web |
| 139 | NetBIOS session service |

**TCP three-way handshake**

A TCP connection is initialized through a three-way handshake. The purpose of the three-way handshake is to synchronize the sequence number and acknowledgment numbers of both sides of the connection and exchange TCP window sizes or the use of large window sizes or TCP timestamps. The following steps outline the process:

1. The initiator of the TCP connection, typically a client, sends a TCP segment to the server with an initial Sequence Number for the connection and a window size indicating the size of a buffer on the client to store incoming segments from the server.

2. The responder of the TCP connection, typically a server, sends back a TCP segment containing its chosen initial Sequence Number, an acknowledgment of the client's Sequence Number, and a window size indicating the size of a buffer on the server to store incoming segments from the client.

3. The initiator sends a TCP segment to the server containing an acknowledgment of the server's Sequence Number.

TCP uses a similar handshake process to end a connection. This guarantees that both hosts have finished transmitting and that all data was received.

**UDP**

UDP provides a connectionless datagram service that offers unreliable, best-effort delivery of data transmitted in messages. This means that neither the arrival of datagrams nor the correct sequencing of delivered packets is guaranteed. UDP does not recover from lost data through retransmission. UDP is defined in RFC 768.

UDP is used by applications that do not require an acknowledgment of receipt of data and that typically transmit small amounts of data at one time. NetBIOS name service, NetBIOS datagram service, and SNMP are examples of services and applications that use UDP. The following table describes the key fields in the UDP header.

**Key Fields in the UDP Header**

| Field | Function |
|---|---|
| Source Port | UDP port of sending host. |
| Destination Port | UDP port of destination host. |
| UDP Checksum | Verifies the bit-level integrity of the UDP header and the UDP data. |

**UDP ports**

To use UDP, an application must supply the IP address and UDP port number of the destination application. A port provides a location for sending messages. A port functions as a multiplexed message queue, meaning that it can receive multiple messages at a time. Each port is identified by a unique number. It is important to note that UDP ports are distinct and separate from TCP ports even though some of them use the same number. The following table lists a few well-known UDP ports.

**Well-Known UDP Ports**

| UDP Port Number | Description |
|---|---|
| 53 | Domain Name System (DNS) name queries |
| 69 | Trivial File Transfer Protocol (TFTP) |
| 137 | NetBIOS name service |
| 138 | NetBIOS datagram service |
| 161 | SNMP |

**TCP/IP Application Interfaces**

For applications to access the services offered by the core TCP/IP protocols in a standard way, network operating systems like Windows Server 2003 make industry-standard application programming interfaces (APIs) available. APIs are sets of functions and commands that are programmatically called by application code to perform network functions. For example, a Web browser application connecting to a Web site needs access to TCP's connection establishment service.

The following figure shows two common TCP/IP APIs, Windows Sockets and NetBIOS, and their relation to the core protocols.

**APIs for TCP/IP**




**Windows Sockets Interface**

The Windows Sockets API is a standard API under Windows Server 2003 for applications that use TCP and UDP. Applications written to the Windows Sockets API run on many versions of TCP/IP. TCP/IP utilities and the SNMP service are examples of applications written to the Windows Sockets interface.

Windows Sockets provides services that allow applications to bind to a particular port and IP address on a host, initiate and accept a connection, send and receive data, and close a connection. There are two types of sockets:

- A stream socket provides a two-way, reliable, sequenced, and unduplicated flow of data using TCP.
- A datagram socket provides a one-way or two-way flow of data using UDP.

A socket is defined by a protocol and an address on the host. The format of the address is specific to each protocol. In TCP/IP, the address is the combination of the IP address and port. Two sockets, one for each end of the connection, form a bi-directional communications path.

To communicate, an application specifies the protocol, the IP address of the destination host, and the port of the destination application. After the application is connected, information can be sent and received.

**NetBIOS Interface**

NetBIOS allows applications to communicate over a network. NetBIOS defines two entities, a session-level interfaceand a session management and data transport protocol.

The NetBIOS interface is a standard API for user applications to submit network input/output (I/O) and control directives to underlying network protocol software. An application program that uses the NetBIOS interface API for network communication can be run on any protocol software that supports the NetBIOS interface.

NetBIOS also defines a protocolthat functions at thesession/transport level. This is implemented by the underlying protocol software (such as the NetBIOS Frames Protocol NBFP — a component of NetBEUI or NetBIOS over TCP/IP (NetBT)), which performs the network I/O required to accommodate the NetBIOS interface command set. NetBIOS over TCP/IP is defined in RFCs 1001 and 1002. NetBT is enabled by default, however Windows Server 2003 allows you to disable NetBT for an environment that contains no NetBIOS-based network clients or applications.

NetBIOS provides commands and support for NetBIOS Name Management, NetBIOS Datagrams, and NetBIOS Sessions.

**NetBIOS name management**

NetBIOS name management services provide the following functions:

- **Name registration and release**. When a TCP/IP host initializes, it registers its NetBIOS names by broadcasting or directing a NetBIOS name registration request to a NetBIOS Name Server such as a WINS server. If another host has registered the same NetBIOS name, either the host or a NetBIOS Name Server responds with a negative name registration response. The initiating host receives an initialization error as a result. When the workstation service on a host is stopped, the host discontinues broadcasting a negative name registration response when someone else tries to use the name and sends a name release to a NetBIOS Name Server. The NetBIOS name is said to be released and available for use by another host.
- **Name Resolution**. When a NetBIOS application wants to communicate with another NetBIOS application, the IP address of the NetBIOS application must be resolved. NetBT performs this function by either broadcasting a NetBIOS name query on the local network or sending a NetBIOS name query to a NetBIOS Name Server.

The NetBIOS name service uses UDP port 137.

**NetBIOS datagrams**
The NetBIOS datagram service provides delivery of datagrams that are connectionless, unsequenced, and unreliable. Datagrams can be directed to a specific NetBIOS name or broadcast to a group of names. Delivery is unreliable in that only the users who are logged on to the network receive the message. The datagram service can initiate and receive both broadcast and directed messages. The NetBIOS datagram service uses UDP port 138.

**NetBIOS sessions**
The NetBIOS session service provides delivery of NetBIOS messages that are connection-oriented, sequenced, and reliable. NetBIOS sessions use TCP connections and provide session establishment, keepalive, and termination. The NetBIOS session service allows concurrent data transfers in both directions using TCP port 139.

**IPv4 Addressing**
For IP version 4, each TCP/IP host is identified by a logical IP address. The IP address is a Network layer address and has no dependence on the Data-Link layer address (such as a MAC address of a network adapter). A unique IP address is required for each host and network component that communicates using TCP/IP and can be assigned manually or by using Dynamic Host Configuration Protocol (DHCP).

The IP address identifies a system's location on the network in the same way a street address identifies a house on a city block. Just as a street address must identify a unique residence, an IP address must be globally unique to the internetwork and have a uniform format.

Each IP address includes a network ID and a host ID.

- The network ID (also known as a network address) identifies the systems that are located on the same physical network bounded by IP routers. All systems on the same physical network must have the same network ID. The network ID must be unique to the internetwork.

- The host ID (also known as a host address) identifies a workstation, server, router, or other TCP/IP host within a network. The host address must be unique to the network ID.

**IPv4 Address Syntax**
An IP address consists of 32 bits. Instead of expressing IPv4 addresses 32 bits at a time using binary notation ($Base_2$), it is standard practice to segment the 32 bits of an IPv4 address into four 8-bit fields called *octets*. Each octet is converted to a decimal number (base 10) from 0–255 and separated by a period (a dot). This format is called *dotted decimal notation*. The following table provides an example of an IP address in binary and dotted decimal formats.

### An IP Address in Binary and Dotted Decimal Formats

| Binary Format | Dotted Decimal Notation |
|---|---|
| 11000000 10101000 00000011 00011000 | 192.168.3.24 |

For example, the IPv4 address of 11000000101010000000001100011000 is:

- Segmented into 8-bit blocks: 11000000 10101000 00000011 00011000.

- Each block is converted to decimal: 192 168 3 24

- The adjacent octets are separated by a period: 192.168.3.24.

The notation *w.x.y.z* is used when referring to a generalized IP address, and is shown the following figure.

### IP Address




**Types of IPv4 Addresses**

The Internet standards define the following types of IPv4 addresses:

- *Unicast*. Assigned to a single network interface located on a specific subnet on the network and used for one-to-one communications.

- *Multicast*. Assigned to one or more network interfaces located on various subnets on the network and used for one-to-many communications.

- *Broadcast*. Assigned to all network interfaces located on a subnet on the network and used for one-to-everyone-on-a-subnet communications.

The following sections describe these types of addresses in detail.

**IPv4 Unicast Addresses**

The IPv4 unicast address identifies an interface's location on the network in the same way a street address identifies a house on a city block. Just as a street address must identify a unique residence, an IPv4 unicast address must be globally unique to the network and have a uniform format.

Each IPv4 unicast address includes a network ID and a host ID.

- The network ID (also known as a network address) is the fixed portion of an IPv4 unicast address that identifies the set of interfaces that are located on the same physical or logical network segment as bounded by IPv4 routers. A network segment on TCP/IP networks is also known as a subnet. All systems on the same physical or logical subnet must use the same network ID and the network ID must be unique to the entire TCP/IP network.

- The host ID (also known as a host address) is the variable portion of an IPv4 unicast address that is used to identify a network node's interface on a subnet. The host ID must be unique to the network ID.

If the network ID is unique to the TCP/IP network and the host ID is unique to the network ID, then the entire IPv4 unicast address consisting of the network ID and host ID is unique to the entire TCP/IP network.

**IPv4 Multicast Addresses**

IPv4 multicast addresses are used for single-packet one-to-many delivery. On an IPv4 multicast-enabled intranet, an IPv4 packet addressed to an IPv4 multicast address is forwarded by routers to the subnets on which there are hosts listening to the traffic sent to the IPv4 multicast address. IPv4 multicast provides an efficient one-to-many delivery service for many types of communication.

IPv4 multicast addresses are defined by the class D Internet address class: 224.0.0.0/4. IPv4 multicast addresses range from 224.0.0.0 through 239.255.255.255. IPv4 multicast addresses for the 224.0.0.0/24 address prefix (224.0.0.0 through 224.0.0.255) are reserved for local subnet multicast traffic.

**IPv4 Broadcast Addresses**

IPv4 uses a set of broadcast addresses to provide a one-to-everyone on the subnet delivery service. Packets sent to IPv4 broadcast addresses are processed by all the interfaces on the subnet. The following are the different types of IPv4 broadcast addresses:

- **Network broadcast**. Formed by setting all the host bits to 1 for a classful address prefix. An example of a network broadcast address for the classful network ID 131.107.0.0/16 is 131.107.255.255. Network broadcasts are used to send packets to all interfaces of a classful network. IPv4 routers do not forward network broadcast packets.

- **Subnet broadcast**. Formed by setting all the host bits to 1 for a classless address prefix. An example of a network broadcast address for the classless network ID 131.107.26.0/24 is 131.107.26.255. Subnet broadcasts are used to send packets to all hosts of a classless network. IPv4 routers do not forward subnet broadcast packets. For a classful address prefix, there is no subnet broadcast address, only a network broadcast address. For a classless address prefix, there is no network broadcast address, only a subnet broadcast address.

- **All-subnets-directed broadcast.** Formed by setting all the original classful network ID host bits to 1 for a classless address prefix. A packet addressed to the all-subnets-directed broadcast was defined to reach all hosts on all of the subnets of a subnetted class-based network ID. An example of an all-subnets-directed broadcast address for the subnetted network ID 131.107.26.0/24 is 131.107.255.255. The all-subnets-directed broadcast is the network broadcast address of the original classful network ID. IPv4 routers can forward all-subnets directed broadcast packets, however the use of the all-subnets-directed broadcast address is deprecated in RFC 1812.

- **Limited broadcast.** Formed by setting all 32 bits of the IPv4 address to 1 (255.255.255.255). The limited broadcast address is used for one-to-everyone delivery on the local subnet when the local network ID is unknown. IPv4 nodes typically only use the limited broadcast address during an automated configuration process such as Boot Protocol (BOOTP) or DHCP. For example, with DHCP, a DHCP client must use the limited broadcast address for all traffic sent until the DHCP server acknowledges the use of the offered IPv4 address configuration. IPv4 routers do not forward limited broadcast packets.

**Internet Address Classes**

The Internet community originally defined address classes to accommodate different types of addresses and networks of varying sizes. The class of address defined which bits were used for the network ID and which bits were used for the host ID. It also defined the possible number of networks and the number of hosts per network. Of five address classes, class A, B, and C addresses were defined for IPv4 unicast addresses. Class D addresses were defined for IPv4 multicast addresses and class E addresses were defined for experimental uses.

**Class A**

Class A network IDs were assigned to networks with a very large number of hosts. The high-order bit in a class A address is always set to zero, which makes the address prefix for all class A networks and addresses 0.0.0.0/1 (or 0.0.0.0, 128.0.0.0). The next seven bits (completing the first octet) are used to enumerate class A network IDs. Therefore, address prefixes for class A network IDs have an 8-bit prefix length (/8 or 255.0.0.0). The remaining 24 bits (the last three octets) are used for the host ID. The address prefix 0.0.0.0/0 (or 0.0.0.0, 0.0.0.0) is a reserved network ID and 127.0.0.0/8 (or 127.0.0.0, 255.0.0.0) is reserved for loopback addresses. Out of a total of 128 possible class A networks, there are 126 networks and 16,777,214 hosts per network.

### Note

- All-Zeros and All-Ones Host IDs are Reserved

- When enumerating host IDs for a given network ID, the two host IDs in which all the bits in the host ID are set to 0 (the all-zeros host ID) and all the bits in the host ID is set to 1 (the all-ones host ID) are reserved and cannot be assigned to network node interfaces. Hence, in the calculation above in which there are 24 bits for class A host IDs, the total number of possible host IDs is 16,777,216 (224). When you subtract the two reserved host IDs, the total number of usable host IDs is 16,777,214.

The following figure illustrates the structure of class A addresses.

**Structure of class A addresses**




**Class B**

Class B network IDs were assigned to medium to large-sized networks. The two high-order bits in a class B address are always set to 10, which makes the address prefix for all class B networks and addresses 128.0.0.0/2 (or 128.0.0.0, 192.0.0.0). The next 14 bits (completing the first two octets) are used to enumerate class B network IDs. Therefore, address prefixes for class B network IDs have a 16-bit prefix length (/16 or 255.255.0.0). The remaining 16 bits (last two octets) are used for the host ID. With 14 bits to express class B network IDs and 16 bits to express host IDs, this allows for 16,384 networks and 65,534 hosts per network.

The following figure illustrates the structure of class B addresses.

**Structure of class B addresses**




**Class C**

Class C addresses were assigned to small networks. The three high-order bits in a class C address are always set to 110, which makes the address prefix for all class C networks and addresses 192.0.0.0/3 (or 192.0.0.0, 224.0.0.0). The next 21 bits (completing the first three octets) are used to enumerate class C network IDs. Therefore, address prefixes for class C network IDs have a 24-bit prefix length (/24 or 255.255.255.0). The remaining 8 bits (the last octet) are used for the host ID. With 21 bits to express class C network IDs and 8 bits to express host IDs, this allows for 2,097,152 networks and 254 hosts per network.

The following figure illustrates the structure of class C addresses.

**Structure of class C addresses**




**Class D**

Class D addresses are reserved for IPv4 multicast addresses. The four high-order bits in a class D address are always set to 1110, which makes the address prefix for all class D addresses 224.0.0.0/4 (or 224.0.0.0, 240.0.0.0).

**Class E**

Class E addresses are reserved for experimental use. The high-order bits in a class E address are set to 1111, which makes the address prefix for all class E addresses 240.0.0.0/4 (or 240.0.0.0, 240.0.0.0)

The following table is a summary of the Internet address classes A, B, and C that can be used for IPv4 unicast addresses.

**Internet Address Class Summary**

| Class | Value for *w* | Network ID Portion | Host ID Portion | Network IDs | Host IDs per Network |
|---|---|---|---|---|---|
| A | 1-126 | *w* | *x.y.z* | 126 | 16,777,214 |
| B | 128-191 | *w.x* | *y.z* | 16,384 | 65,534 |
| C | 192-223 | *w.x.y* | *z* | 2,097,152 | 254 |

**Modern Internet Addresses**

The Internet address classes are an obsolete unicast address allocation method that proved to be an inefficient way to assign network IDs and addresses to organizations connected to the Internet. For example, a large organization with a class A network ID can have up to 16,777,214 hosts. However, if the organization only uses 70,000 host IDs, then 16,707,214 potential IPv4 unicast addresses for the Internet are wasted.

On the modern-day Internet, IPv4 address prefixes are handed out to organization's based on the organization's actual need for Internet-accessible IPv4 unicast addresses using a method known as Classless Inter-Domain Routing (CIDR). For example, an organization determines that it needs 2,000 Internet-accessible IPv4 unicast addresses. The Internet Corporation for Assigned Names and Numbers (ICANN) or an Internet service provider (ISP) allocates an IPv4 address prefix in which 21 bits are fixed, leaving 11 bits for host IDs. From the 11 bits for host IDs, the organization can create 2,032 possible IPv4 unicast addresses.

CIDR-based address allocations typically start at 8 bits. The following table lists the required number of host IDs and the corresponding prefix length for CIDR-based address allocations.

**Host IDs Needed and CIDR-based Prefix Lengths**

| Number of Host IDs | Prefix Length | Dotted Decimal |
|---|---|---|
| 2–254 | /24 | 255.255.255.0 |
| 255–510 | /23 | 255.255.254.0 |

| | | |
|---|---|---|
| 511–1,022 | /22 | 255.255.252.0 |
| 1,021–2,046 | /21 | 255.255.248.0 |
| 2,047–4,094 | /20 | 255.255.240.0 |
| 4,095–8,190 | /19 | 255.255.224.0 |
| 8,191–16,382 | /18 | 255.255.192.0 |
| 16,383–32,766 | /17 | 255.255.128.0 |
| 32,767–65,534 | /16 | 255.255.0.0 |

**Public and Private Addresses**

If you want direct (routed) connectivity to the Internet, then you must use public addresses. If you want indirect (proxied or translated) connectivity to the Internet, you can use either public or private addresses. If your intranet is not connected to the Internet in any way, you can use any unicast IPv4 addresses that you want. However, you should use private addresses to avoid network renumbering when your intranet is eventually connected to the Internet.

**Public addresses**

Public addresses are assigned by ICANN and consist of either historically allocated class-based network IDs or, more recently, CIDR-based address prefixes that are guaranteed to be globally unique on the Internet. For CIDR-based address prefixes, the value of *w* (the first octet) is in the ranges of 1 through 126 and 128 through 223, with the exception of the private address prefixes described in "Private Addresses."

When the public addresses are assigned, routes are added to the routers of the Internet so that traffic sent to an address that matches the assigned public address prefix can reach the assigned organization. For example, when an organization is assigned an address prefix in the form of a network ID and prefix length, that address prefix also exists as a route in the routers of the Internet. IPv4 packets destined to an address within the assigned address prefix are routed to the proper destination.

**Private addresses**

Each IPv4 interface requires an IPv4 address that is globally unique to the IPv4 network. In the case of the Internet, each IPv4 interface on a subnet connected to the Internet requires an IPv4 address that is globally unique to the Internet. As the Internet grew, organizations connecting to the Internet required a public address for each interface on their intranets. This requirement placed a huge demand on the pool of available public addresses.

When analyzing the addressing needs of organizations, the designers of the Internet noted that for many organizations, most of the hosts on an organization's intranet did not require direct connectivity to the Internet. Those hosts that did require a specific set of Internet services, such as Web access and e-mail, typically access the Internet services through Application layer gateways such as proxy servers and e-mail servers. The result is that most organizations only required a small amount of public addresses for those nodes (such as proxies, servers, routers, firewalls, and translators) that were directly connected to the Internet.

For the hosts within the organization that do not require direct access to the Internet, IPv4 addresses that do not duplicate already-assigned public addresses are required. To solve this addressing problem, the Internet designers reserved a portion of the IPv4 address space and named this space the private address space. An IPv4 address in the private address space is never assigned as a public address. IPv4 addresses within the private address space are known as private addresses. Because the public and private address spaces do not overlap, private addresses never duplicate public addresses.

The private address space specified in RFC 1918 is defined by the following address prefixes:

- 10.0.0.0/8 (10.0.0.0, 255.0.0.0)

  Allows the following range of valid IPv4 unicast addresses: 10.0.0.1 to 10.255.255.254. The 10.0.0.0/8 address prefix has 24 host bits that can be used for any addressing scheme within the private organization.

- 172.16.0.0/12 (172.16.0.0, 255.240.0.0)

  Allows the following range of valid IPv4 unicast addresses: 172.16.0.1 to 172.31.255.254. The 172.16.0.0/12 address prefix has 20 host bits that can be used for any addressing scheme within the private organization.

- 192.168.0.0/16 (192.168.0.0, 255.255.0.0)

Allows the following range of valid IPv4 unicast addresses: 192.168.0.1 to 192.168.255.254. The 192.168.0.0/16 address prefix has 16 host bits that can be used for any addressing scheme within the private organization.

Because the IPv4 addresses in the private address space will never be assigned by ICANN to an organization connected to the Internet, there will never be routes for the private address prefixes in Internet routers. You cannot connect to a private address over the Internet. Therefore, a host that has a private address must send its Internet traffic requests to an Application layer gateway (such as a proxy server) that has a valid public address or through a network address translator (NAT) that translates the private address into a valid public address.

**Illegal addresses**
Private organization intranets that do not need an Internet connection can choose any address scheme they want, even using public address prefixes that have been assigned by ICANN. If that organization later decides to connect to the Internet, its current address scheme might include addresses already assigned by ICANN to other organizations. These addresses conflict with existing public addresses assigned by ICANN and are known as illegal addresses. Connectivity from illegal addresses to Internet locations is not possible because the routers of the Internet send traffic destined to ICANN-allocated address prefixes to the assigned organizations, not to the organizations using illegal addresses.

For example, a private organization chooses to use the 206.73.118.0/24 address prefix for its intranet. The public address prefix 206.73.118.0/24 has been assigned by ICANN to the Microsoft Corporation and routes exist on the Internet routers to send all packets for IPv4 addresses on 206.73.118.0/24 to Microsoft routers. As long as the private organization does not connect to the Internet, there is no problem because the two address prefixes are on separate IPv4 networks; therefore they are unique to each separate network. If the private organization later connects directly to the Internet and continues to use the 206.73.118.0/24 address prefix, any Internet response traffic to locations matching the 206.73.118.0/24 address prefix is sent to Microsoft routers, not to the routers of the private organization.

**Automatic Private IP Addressing**
An interface on a computer running Windows Server 2003 and Windows XP that is configured to obtain an IPv4 address configuration automatically that does not successfully contact a Dynamic Host Configuration Protocol (DHCP) server uses its alternate configuration, as specified on the Alternate Configuration tab.

If the Automatic Private IP Address option is selected on the Alternate Configuration tab and a DHCP server cannot be found, Windows TCP/IP uses Automatic Private IP Addressing (APIPA). Windows TCP/IP randomly selects an IPv4 address from the 169.254.0.0/16 address prefix and assigns the subnet mask of 255.255.0.0. This address prefix has been reserved by the ICANN and is not reachable on the Internet. APIPA allows single-subnet Small Office/Home Office (SOHO) networks to use TCP/IP without static configuration or the administration of a DHCP server. APIPA does not configure a default gateway. Therefore, only local subnet traffic is possible.

**Special IPv4 Addresses**
The following are special IPv4 addresses:

- 0.0.0.0

  Known as the unspecified IPv4 address, it is used to indicate the absence of an address. The unspecified address is used only as a source address when the IPv4 node is not configured with an IPv4 address configuration and is attempting to obtain an address through a configuration protocol such as Dynamic Host Configuration Protocol (DHCP).

- 127.0.0.1

  Known as the IPv4 loopback address, it is assigned to an internal loopback interface, enabling a node to send packets to itself.

**Unicast IPv4 Addressing Guidelines**
When assigning network IDs to the subnets of an organization, use the following guidelines:

- The network ID must be unique on the IPv4 network.

  If the network ID is for a subnet on which there are hosts that are directly accessible from the Internet, you must use a public IPv4 address prefix assigned by ICANN or an Internet service provider. If the network ID is for a subnet that is not directly accessible by the Internet, use either a legal public address prefix or a private address prefix that is unique on your private intranet.

- The network ID cannot begin with the numbers 0 or 127.

  Both of these values for the first octet are reserved and cannot be used for IPv4 unicast addresses.

When assigning host IDs to the interfaces of nodes on an IPv4 subnet, use the following guidelines:

- The host ID must be unique on the subnet.
- You cannot use the all-zeros or all-ones host IDs.

When defining the range of valid IPv4 unicast addresses for a given address prefix, use the following standard practice:

- For the first IPv4 unicast address in the range, set all the host bits in the address to 0, except for the low-order bit, which is set to 1.
- For the last IPv4 unicast address in the range, set all the host bits in the address to 1, except for the low-order bit, which is set to 0.

For example, to express the range of addresses for the address prefix 192.168.16.0/20:

- The first IPv4 unicast address in the range is 11000000 10101000 0001**000000000001** (host bits are bold), or 192.168.16.1.
- The last IPv4 unicast address in the range is 11000000 10101000 0001**111111111110** (host bits are bold), or 192.168.31.254.

**Name Resolution**

While IP is designed to work with the 32-bit IP addresses of the source and the destination hosts, computers users are much better at using and remembering names than IP addresses.

When a name is used as an alias for an IP address, a mechanism must exist for assigning that name to the appropriate IP node — to ensure its uniqueness and resolution to its IP address.

In this section, the mechanisms used for assigning and resolving host names (which are used by Windows Sockets applications), and NetBIOS names (which are used by NetBIOS applications) are discussed.

**Host Name Resolution**

A host name is an alias assigned to an IP node to identify it as a TCP/IP host. The host name can be up to 255 characters long and can contain alphabetic and numeric characters and the "-" and "." characters"." Multiple host names can be assigned to the same host. For Windows Server 2003–based computers, the host name does not have to match the Windows Server 2003 computer name.

Windows Sockets applications, such as Microsoft Internet Explorer, can use one of two values to connect to the destination: the IP address or a host name. When the IP address is specified, name resolution is not needed. When a host name is specified, the host name must be resolved to an IP address before IP-based communication with the desired resource can begin.

Host names most commonly take the form of a domain name with a structure that follows Internet conventions. Name resolution, and domain names work the same whether they are used for IPv4 or IPv6 addresses.

**Domain Names**

To meet the need for a scaleable, customizable naming scheme for a wide variety of organizations, InterNIC has created and maintains a hierarchical namespace called the Domain Name System (DNS). The DNS naming scheme looks like the directory structure for files on a disk. Usually, you trace a file path from the root directory through subdirectories to its final location and its file name. However, a host name is traced from its final location back through its parent domains up to the root. The unique name of the host, representing its position in the hierarchy, is its Fully Qualified Domain Name (FQDN). The top-level domain namespace with second-level and subdomains is shown in the following figure.

**Domain Name System**




The domain namespace includes the following categories:

- The root domain, which is indicated by "" (null), represents the root of the namespace.
- Top-level domains, directly below the root, represent types of organizations. InterNIC is responsible for the maintenance of top-level domain names on the Internet. The following table has a partial list of the Internet's top-level domain names.

**Internet Top-Level Domain Names**

| **Domain Name** | **Meaning** |
|---|---|
| com | Commercial organization |
| edu | Educational institution |
| gov | Government institution |
| mil | Military group |
| net | Major network support center |
| org | Organization other than those above |
| int | International organization |
| *<country/region code>* | Each country/region (geographic scheme) |

- Second-level domains, below the top-level domains, represent specific organizations within the top-level domains. InterNIC is responsible for maintaining and ensuring uniqueness of second-level domain names on the Internet.
- Subdomains are below the second-level domain. Individual organizations are responsible for the creation and maintenance of subdomains.

For example, for the FQDN **websrv.wcoast.reskit.com**:

- The trailing period (**.**) denotes that this is an FQDN with the name relative to the root of the domain namespace. The trailing period is usually not required for FQDNs and if it is missing it is assumed to be present.
- **com** is the top-level domain, indicating a commercial organization.
- **reskit** is the second-level domain, indicating the Resource Kit Corporation.
- **wcoast** is a subdomain of **reskit.com** indicating the West Coast division of the Resource Kit Corporation.
- **websrv** is the name of the Web server in the West Coast division.

Domain names are not case-sensitive.

Organizations not connected to the Internet can implement whatever top and second-level domain names they want. However, typical implementations follow InterNIC specifications so that eventual participation in the Internet will not require a renaming process.

**Host Name Resolution Using a Hosts File**

One common way to resolve a host name to an IP address is to use a locally stored database file that contains IP-address-to-host-name mappings. On most UNIX systems, this file is /etc/hosts. On Windows Server 2003 systems, it is the Hosts file in the **%systemroot%\System32\Drivers\Etc** directory.

The following is an example of the contents of the Hosts file:

```
#
Table of IP addresses and host names
#
127.0.0.1    localhost
131.107.34.1    router
172.30.45.121    server1.central.reskit.com s1
```

Within the Hosts file:

- Multiple host names can be assigned to the same IP address. Note that the server at the IP address 172.30.45.121 can be referred to by its FQDN (server1.central.reskit.com) or a nickname (s1). This allows the user at this computer to refer to this server using the nickname **s1** instead of typing the entire FQDN.

- Entries can be case sensitive depending on the platform. Entries in the Hosts file for UNIX computers are case-sensitive. Entries in the Hosts file for Windows Server 2003, Windows XP, and Windows 2000–based computers are not case sensitive.

For computers running Windows Server 2003, Windows XP, and Windows 2000, the entries in the Hosts file are loaded into the DNS client resolver cache. When resolving host names, the DNS client resolver cache is always checked.

The advantage of using a Hosts file is that it is customizable for the user. Users can create whatever entries they want, including easy-to-remember nicknames for frequently accessed resources. However, the individual maintenance of the Hosts file does not scale well to storing large numbers of FQDN mappings.

**Host Name Resolution Using a DNS Server**

To make host name resolution scalable and centrally manageable, IP address mappings for FQDNs are stored on DNS servers. To enable the querying of a DNS server by a host computer, a component called the DNS resolver is enabled and configured with the IP address of the DNS server. The DNS resolver is a built-in component of TCP/IP protocol stacks supplied with most network operating systems, including Windows Server 2003.

When a Windows Sockets application is given an FQDN as the destination location, the application calls a Windows Sockets function to resolve the name to an IP address. The request is passed to the DNS resolver component in the TCP/IP protocol. The DNS resolver packages the FQDN request as a DNS Name Query packet and sends it to the DNS server.

DNS is a distributed naming system. Instead of storing all the records for the entire namespace on each DNS server, each DNS server stores only the records for a specific portion of the namespace. The DNS server is authoritative for the portion of the namespace that corresponds to records stored on that DNS server. In the case of the Internet, hundreds of DNS servers store various portions of the Internet namespace. To facilitate the resolution of any valid domain name by any DNS server, DNS servers are also configured with pointer records to other DNS servers.

The following process outlines what happens when the DNS resolver component on a host sends a DNS query to a DNS server. This process is shown in the following figure and is simplified so that you can gain a basic understanding of the DNS resolution process.

1. The DNS resolver component of the DNS client formats a DNS Name Query Request message containing the FQDN and sends it to the configured DNS server.

2. The DNS server checks the FQDN in the DNS Name Query Request message against locally stored address records. If a record is found, the IP address corresponding to the requested FQDN is sent back to the client.

3. If the FQDN is not found, the DNS server forwards the request to a DNS server that is authoritative for the FQDN.

4. The authoritative DNS server returns the reply, which contains the resolved IP address, back to the original DNS server.

5. The original DNS server sends the IP address mapping information to the client.

**Resolving an FQDN by using DNS servers**




To obtain the IP address of a server that is authoritative for the FQDN, DNS servers on the Internet go through an iterative process of querying multiple DNS servers until the authoritative server is found. For more information about DNS name-resolution processes, see the DNS Technical Reference.

**Combining a Local Database File with DNS**

TCP/IP implementations, including Windows Server 2003, allow the use of both a local database file and a DNS server to resolve host names. When a user specifies a host name in a Windows Sockets–based TCP/IP application:

- TCP/IP checks the DNS client resolver cache (loaded with entries from the Hosts file and other previously resolved host names) for a matching name. If a matching name is not found in the local database file, the host name is packaged as a DNS Name Query Request message and sent to the configured DNS server.

Combining methods allows the user to have a local database file for resolving personalized nicknames and to use the globally distributed DNS database to resolve FQDNs.

**NetBIOS Name Resolution**

NetBIOS name resolution is the process of successfully mapping a NetBIOS name to an IP address. A NetBIOS name is a 16-byte address used to identify a NetBIOS resource on the network. A NetBIOS name is either a unique (exclusive) or group (nonexclusive) name. When a NetBIOS process communicates with a specific process on a specific computer, a unique name is used. When a NetBIOS process communicates with multiple processes on multiple computers, a group name is used.

The NetBIOS name acts as a Session layer application identifier. For example, the NetBIOS session service operates over TCP port 139. All NetBT session requests are addressed to TCP destination port 139. When identifying a NetBIOS application with which to establish a NetBIOS session, the NetBIOS name is used.

An example of a process using a NetBIOS name is the File and Printer Sharing for Microsoft Networks component (the Server service) on a Windows Server 2003–based computer. When you start your computer, the Server service registers a unique NetBIOS name based on your computer's name. The exact name used by the Server service is the 15-character computer name plus a sixteenth character of 0x20. If the computer name is not 15 characters long, it is padded with spaces up to 15 characters long. Other network services, such as the Workstation or Messenger service, also use the computer name to build their NetBIOS names. The sixteenth character is used to uniquely identify each service.

**Note**

- The Messenger service refered to here is not Windows Messenger. Windows Messenger is a Microsoft application included in Windows Server 2003 that allows real-time messaging and collaboration.

The Server service on the file server you specify corresponds to a specific NetBIOS name. For example, when you attempt to connect to the computer called CORPSERVER, the NetBIOS name corresponding to the Server service is "CORPSERVER <20>" (note the padding using the space character). Before a file and print sharing connection can be established, a TCP connection must be created. In order for a TCP connection to be established, the NetBIOS name "CORPSERVER <20>" must be resolved to an IP address.

To view the NetBIOS names registered by NetBIOS processes running on a Windows Server 2003 computer, type **nbtstat -n** at the Windows Server 2003 command prompt.

**NetBIOS Node Types**

The exact mechanism by which NetBIOS names are resolved to IP addresses depends on the node's configured NetBIOS Node Type. RFC 1001 defines the NetBIOS Node Types listed in the following table.

**NetBIOS Node Types**

| Node Type | Description |
|---|---|
| B-node (broadcast) | B-node uses broadcasted NetBIOS name queries for name registration and resolution. B-node has two major problems: (1) In a large internetwork, broadcasts can increase the network load, and (2) Routers typically do not forward broadcasts, so only NetBIOS names on the local network can be resolved. |
| P-node (peer-peer) | P-node uses a NetBIOS name server (NBNS), such as Windows Internet Name Service (WINS), to resolve NetBIOS names. P-node does not use broadcasts; instead, it queries the name server directly. The most significant problem with P-node is that all computers must be configured with the IP address of the NBNS, and if the NBNS is down, computers are not able to communicate even on the local network. |
| M-node (mixed) | M-node is a combination of B-node and P-node. By default, an M-node functions as a B-node. If it is unable to resolve a name by broadcast, it uses the NBNS of P-node. |
| H-node (hybrid) | H-node is a combination of P-node and B-node. By default, an H-node functions as a P-node. If it is unable to resolve a name through the NetBIOS name server, it uses a broadcast to resolve the name. |

When NetBT is enabled, Windows Server 2003–based computers are B-node by default and become H-node when configured for a WINS server. Windows Server 2003 also uses a local database file called Lmhosts to resolve remote NetBIOS names.

**IPv4 Routing**

After the host name or NetBIOS name is resolved to an IP address, the IP packet must be sent by the sending host to the resolved IP address. Routing is the process of forwarding a packet based on the destination IP address. Routing involves both the TCP/IP host and an IP router. A router is a device that forwards the packets from one network to another. Routers are also commonly referred to as gateways. Both the sending host and router need to make a determination about how the packet is forwarded.

To make these determinations, the IP layer consults a routing table stored in memory. Routing table entries are created by default when TCP/IP initializes and additional entries are added either manually by a system administrator or automatically through communication with routers.

**Direct and Indirect Delivery**

IP packets use at least one of two types of delivery based on whether the final destination is located on a directly attached network. These two types of delivery are known as direct and indirect delivery.

- Direct delivery occurs when the IP node (either the sending node or an IP router) forwards a packet to the final destination on a directly attached network. The IP node encapsulates the IP packet in a frame format for the Network Interface layer (such as Ethernet or Token Ring) addressed to the destination's MAC address.

- Indirect delivery occurs when the IP node (either the sending node or an IP router) forwards a packet to an intermediate node (an IP router) because the final destination is not on a directly attached network. The IP node encapsulates the IP packet in a frame format for the Network Interface layer (such as Ethernet or Token Ring) addressed to the IP router's MAC address.

IP routing is a combination of direct and indirect deliveries.

In the following figure, when sending packets to node B, node A performs a direct delivery. When sending packets to node C, node A performs an indirect delivery to Router 1, and Router 1 performs an indirect delivery to Router 2, and then Router 2 performs a direct delivery to node C.

**Direct and Indirect Deliveries**




**IP Routing Table**

A routing table is present on all IP nodes. The routing table stores information about IP networks and how they can be reached (either directly or indirectly). Because all IP nodes perform some form of IP routing, routing tables are not exclusive to IP routers. Any node loading the TCP/IP protocol has a routing table. There are a series of default entries according to the configuration of the node and additional entries can be entered either manually through TCP/IP utilities or dynamically through interaction with routers.

When an IP packet is to be forwarded, the routing table is used to determine:

- **The next-hop IP address.** For a direct delivery, the next-hop IP address is the destination IP address in the IP packet. For an indirect delivery, the next-hop IP address is the IP address of a router.

- **The next-hop interface.** The next-hop interface identifies the physical or logical interface, such as a network adapter, that is used to forward the packet to either its destination or the next router.

**IP Routing Table Entry Types**

Entries in the IP routing table contain the following information:

- **Network ID.** The network ID or destination corresponding to the route. The network ID can identify a specific subnet, be a summarized route, or an IP address for a host route. In the Windows Server 2003 IP routing table, this is the Network Destination column.

- **Network mask.** The mask that is used to match a destination IP address to the network ID. In the Windows Server 2003 IP routing table, this is the Netmask column.

- **Next hop.** The IP address of the next hop. In the Windows Server 2003 IP routing table, this is the Gateway column.

- **Interface.** An indication of which network interface is used to forward the IP packet.

- **Metric.** A number used to indicate the cost of the route so the best route among possible multiple routes to the same destination can be selected. A common use of the metric is to indicate the number of hops (routers crossed) to the network ID.

Entries in the routing table can be used to store the following types of routes:

- **Directly attached network ID.** Aroute for network IDs that are directly attached. For directly attached networks, the Next Hop field can be blank or contain the IP address of the interface on that network.

- **Remote network ID.** A route for network IDs that are not directly attached but are available across other routers. For remote networks, the Next Hop field is the IP address of a local router.

- **Host route.** A route to a specific IP address. Host routes allow routing to occur on a per-IP address basis. For host routes, the network ID is the IP address of the specified host and the network mask is 255.255.255.255.

- **Default route.** The default route is designed to be used when a more specific network ID or host route is not found. The default route network ID is 0.0.0.0 with a network mask of 0.0.0.0.

**Route Determination Process**
To determine which routing table entry is used to find the next-hop address and interface, IP uses the following process:

- For each entry in a routing table, IP performs a bit-wise logical AND operation between the destination IP address and the network mask. It compares the result with the network ID of the entry for a match.

- A list of matching routes is compiled. The route that has the longest match (the route with the largest number of bits that match the destination IP address) is chosen. The longest matching route is the most direct route to the destination IP address. If multiple matching entries are found (for example, multiple routes to the same network ID), the router uses the lowest metric to select the best route. If multiple entries have the longest match and the lowest metric, the router designates one of them as the routing table entry. For Windows Server 2003 TCP/IP, the route chosen corresponds to the route associated with the interface that is first in the network binding order.

The end result of the route-determination process is a single route in the routing table that yields a next-hop IP address and interface. If the route-determination process fails to find a route, IP indicates a routing error. For the sending host, an IP routing error message is sent to the upper layer protocol, such as TCP or UDP. For a router, an ICMP Destination Unreachable–Host Unreachable message is sent to the sending host.

**Routing Table for Windows Server 2003**
The following table shows the default routing table for a Windows Server 2003–based host (not a router). The host has a single network adapter and has the IP address 157.60.27.90, subnet mask 255.255.240.0, and a default gateway of 157.60.16.1.

### Windows Server 2003 Routing Table

| Network Destination | Netmask | Gateway | Interface | Metric | Purpose |
|---|---|---|---|---|---|
| 0.0.0.0 | 0.0.0.0 | 157.60.16.1 | 157.60.27.90 | 1 | Default Route |
| 127.0.0.0 | 255.0.0.0 | 127.0.0.1 | 127.0.0.1 | 1 | Loopback Network |
| 157.60.16.0 | 255.255.240.0 | 157.60.27.90 | 157.60.27.90 | 1 | Directly Attached Network |
| 157.60.27.90 | 255.255.255.255 | 127.0.0.1 | 127.0.0.1 | 1 | Local Host |
| 157.60.255.255 | 255.255.255.255 | 157.60.27.90 | 157.60.27.90 | 1 | Network Broadcast |
| 224.0.0.0 | 240.0.0.0 | 157.60.27.90 | 157.60.27.90 | 1 | Multicast |
| 255.255.255.255 | 255.255.255.255 | 157.60.27.90 | 157.60.27.90 | 1 | Limited Broadcast |

**Default Route**
The entry corresponding to the default gateway configuration is a network destination of 0.0.0.0 with a network mask (netmask) of 0.0.0.0. Any destination IP address that is logically ANDed with 0.0.0.0 results in 0.0.0.0. Therefore, for any IP address, the default route produces a match. If the default route is chosen because no better routes were found, the IP packet is forwarded to the IP address in the Gateway column (the default gateway of 157.60.16.1), by using the interface assigned the IP address in the Interface column.

**Loopback Network**
The loopback network entry is designed to take any IP address of the form 127.x.y.z and forward it to the special loopback address of 127.0.0.1.

**Directly Attached Network**
The local network entry corresponds to the directly attached network. IP packets destined for the directly attached network are not forwarded to a router but sent directly to the destination. Note that the Gateway and Interface columns match the IP address of the node. This indicates that the packet is sent from the network adapter corresponding to the node's IP address.

**Local Host**
The local host entry is a host route (network mask of 255.255.255.255) corresponding to the IP address of the host. All IP packets sent to the IP address of the host are forwarded to the loopback address.

**Network Broadcast**
The network broadcast entry is a host route (network mask of 255.255.255.255) corresponding to the all-subnets directed broadcast address (all subnets of class B network ID 157.60.0.0). Packets addressed to the all-subnets directed broadcast are

sent from the network adapter corresponding to the node's IP address.

**Multicast**
The multicast addresses route is used to send any multicast IP packets from the network adapter corresponding to the node's IP address.

**Limited Broadcast**
The limited broadcast address is a host route (network mask of 255.255.255.255). Packets addressed to the limited broadcast are sent from the network adapter corresponding to the node's IP address.

**Viewing the IP Routing Table**
To view the IP routing table on a Windows Server 2003-based computer, type **route print** at a Windows Server 2003 command prompt.

When determining the next-hop IP address and interface from a route in the routing table:

- If the gateway address is the same as the interface address, the next-hop IP address is set to the destination IP address of the IP packet.
- If the gateway address is not the same as the interface address, the next-hop IP address is set to the gateway address.

For example, when traffic is sent to 157.60.16.48, the most specific matching route is the route for the directly attached network (157.60.16.0/20). The next-hop IP address is set to the destination IP address (157.60.16.48) and the interface is the network adapter that has been assigned the IP address 157.60.27.90.

When sending traffic to 192.168.0.79, the most specific matching route is the default route (0.0.0.0/0). The next-hop IP address is set to the gateway address (157.60.16.1) and the interface is the network adapter that has been assigned the IP address 157.60.27.90.

**Maintenance of Routing Table Entries**
For IP routing to occur efficiently between routers in the IP internetwork, routers must be configured with remote network IDs or a default route. On large IP internetworks, one of the challenges faced by network administrators is how to maintain the routing tables on their IP routers so that IP traffic flow is traveling the best path and is fault tolerant.

There are two methods of maintaining routing table entries on IP routers.

**Manual**
Static IP routers have routing tables that do not change unless manually changed by a network administrator.

Static routing relies on the manual administration of the routing table. Remote network IDs are not discovered by static routers and must be manually configured. Static routers are not fault-tolerant. If a static router goes down, neighboring routers do not sense the fault and inform other routers.

**Automatic**
Dynamic IP routers have routing tables that change automatically based on the exchange of routing information with other routers.

Dynamic routing employs the use of routing protocols, such as Routing Information Protocol (RIP) and Open Shortest Path First (OSPF), to dynamically update the routing table through the exchange of routing information between routers. Remote network IDs are discovered by dynamic routers and automatically entered into the routing table. Dynamic routers are fault-tolerant. If a dynamic router goes down, the fault is detected by neighboring routers, which send the changed routing information to the other routers in the internetwork.

**Physical Address Resolution**
Based on the destination IP address and the route determination process, IP determines the next-hop IP address and interface. IP then sends the IP packet, the next-hop IP address, and the interface to ARP.

If the next-hop IP address is the same as the destination IP address, then ARP performs a direct delivery. In a direct delivery, the MAC address corresponding to the destination IP address must be resolved.

If the next-hop IP address is not the same as the destination IP address, then ARP performs an indirect delivery. The next-hop IP address is the IP address of a router between the current IP node and the final destination. In an indirect delivery, the MAC address corresponding to the IP address of the router must be resolved.

To resolve a next-hop IP address to its MAC address, ARP uses broadcast traffic on shared access networking media (such as Ethernet or Token Ring) to send out a broadcasted ARP Request frame. An ARP Reply, containing the MAC address

corresponding to the requested next-hop IP address, is sent back to the sender of the ARP Request.

**ARP Cache**
To keep the number of broadcasted ARP Request frames to a minimum, many TCP/IP protocol stacks incorporate an ARP cache, which is a table of recently resolved IP addresses and their corresponding MAC addresses. TCP/IP checks the ARP cache before sending an ARP Request frame. Each interface has its own ARP cache.

Depending on the vendor implementation, the ARP cache can have the following qualities:

- ARP cache entries can be dynamic (based on ARP Replies) or static. Static ARP entries are permanent and are manually added by using a TCP/IP utility such as the ARP tool provided with Windows Server 2003. Static ARP cache entries are used to prevent ARP Requests for commonly used local IP addresses, such as routers and servers. The problem with static ARP entries is that they have to be manually updated when network interface equipment changes.

- Dynamic ARP cache entries have a time-out value associated with them to remove entries in the cache after a specified period of time. Dynamic ARP cache entries for Windows Server 2003 TCP/IP are given a maximum time of 10 minutes before being removed.

To view the ARP cache on a Windows Server 2003–based computer, type **arp -a** at a Windows Server 2003 command prompt.

**ARP Process**
IP sends ARP the IP packet, the next-hop IP address, and the next-hop interface. Whether performing a direct or indirect delivery, ARP carries out the following process, as shown in the following figure.

**ARP process**




1. Based on the next-hop address and interface, ARP consults the appropriate ARP cache for an entry for the next-hop IP address. If an entry is found, ARP skips to step 6.

2. If an entry is not found, ARP builds an ARP Request frame containing the MAC address of the interface sending the ARP Request, the IP address of the interface sending the ARP Request, and the next-hop IP address. ARP then broadcasts the ARP Request using the appropriate interface.

3. All hosts receive the broadcasted frame and the ARP Request is processed. If the receiving host's IP address matches the requested IP address (the next-hop IP address), its ARP cache is updated with the address mapping of the sender of the ARP Request.

4. If the receiving host's IP address does not match the requested IP address, the ARP Request is silently discarded.

5. The receiving host formulates an ARP Reply containing the requested MAC address and sends it directly to the sender of the ARP Request.

6. When the ARP Reply is received by the sender of the ARP Request, it updates its ARP cache with the address mapping.

7. The ARP Request host and the ARP Reply host have each other's address mappings in their ARP caches.

8. ARP sends the IP packet to the next-hop node by addressing it to the resolved MAC address.

**End-to-End Delivery**
The IP routing processes for all nodes involved in the delivery of an IP packet include the sending host, the intermediate routers, and the destination host.

**IP on the Sending Host**
When a host sends a packet, the packet is transmitted from an upper layer protocol (TCP, UDP, or ICMP) to IP, and then IP on the sending host does the following:

1. Sets the Time-to-Live (TTL) value to either a default or application-specified value.

2. Checks its routing table for the best route to the destination IP address.

   If no route is found, IP sends a routing error message to the upper layer protocol (TCP, UDP, or ICMP).

3. Determines the next-hop IP address and the interface, based on the most specific matching route.

4. Sends the packet, the next-hop IP address, and the next-hop interface to Address Resolution Protocol (ARP), and then ARP resolves the next-hop IP address to its media access control (MAC) address and forwards the packet.

**IP on the Router**
When a packet is received at a router, the packet is passed to IP, and IP on the router does the following:

1. Verifies the IP header checksum.

   If the IP header checksum fails, the IP packet is discarded without notification to the user. This is known as a silent discard.

2. Verifies whether the destination IP address in the IP packet corresponds to an IP address assigned to a router interface.

   If so, the router processes the IP packet as the destination host (see step 3 in the following "IP on the Destination Host" section).

3. If the destination IP address is not the router, IP decrements the Time-to-Live (TTL).

   If the TTL is 0, the router discards the packet and sends an ICMP Time Expired–TTL Expired in Transit message to the sender.

4. If the TTL is 1 or greater, IP updates the TTL field and calculates a new IP header checksum.

5. IP checks its routing table for the best route to the destination IP address in the IP packet.

   If no route is found, the router discards the packet and sends an ICMP Destination Unreachable–Host Unreachable message to the sender.

6. Based on the best route found, IP determines the next-hop IP address and interface.

7. IP sends the packet, the next-hop IP address, and the interface to ARP, and then ARP forwards the packet to the appropriate MAC address.

This entire process is repeated at each router in the path between the source and destination host.

**IP on the Destination Host**
When a packet is received at the destination host, it is passed up to IP, and IP on the destination host does the following:

1. Verifies the IP header checksum.

   If the IP header checksum fails, the IP packet is silently discarded.

2. Verifies that the destination IP address in the IP packet corresponds to an IP address assigned to the host.

   If the destination IP address is not assigned to the host, the IP packet is silently discarded.

3. Passes the IP packet without the IP header to the appropriate upper-level protocol, based on the IP protocol field.

   If the protocol does not exist, ICMP sends a Destination Unreachable–Protocol Unreachable message back to the sender.

4. For TCP and UDP packets, IP checks the destination port and processes the TCP segment or UDP header.

   If no application exists for the UDP port number, ICMP sends a Destination Unreachable–Port Unreachable message back to the sender. If no application exists for the TCP port number, TCP sends a Connection Reset segment back to the sender.

**Related Information**

- For more information about RFCs, see the IETF RFC Database.

# TCP/IP Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP Tools and Settings**

**In this section**

TCP/IP Tools

Microsoft Windows Server 2003 and Microsoft Windows XP Professional support many TCP/IP diagnostic and remote networking tools that help manage and troubleshoot your network. Windows Server 2003 adds new IPv6 tools and updates several standard IPv4 tools to include IPv6 functionality.

**TCP/IP Tools**

The following sections "TCP/IP Diagnostic Tools" and "TCP/IP Remote Networking Tools" describe the diagnostic and remote networking tools associated with TCP/IP.

**Other Tools**

In addition to the Microsoft TCP/IP tools described in this document, Windows Server 2003 supports the tools shown in the following table. The Ipsec6 and Ttcp tools are new IPv6 tools.

**Other TCP/IP-Related Tools**

| Tool | Description |
|---|---|
| Ipsec6 | You can use **Ipsec6** to experiment with managing Internet Protocol security (IPSec) in an IPv6 test environment by configuring IPSec policies and security associations. |
| Ttcp | You can use this tool to troubleshoot TCP and UDP traffic. You can find the **Ttcp** tool in the **Valueadd\Msft\Net\Tools** folder of the Windows Server 2003 CD-ROM. |
| Event Viewer | Tracks errors and events recorded in Application, Security, and System logs. |
| Network Monitor | Captures and displays network traffic. The full version, which can capture all frames sent on a network segment, is part of the Microsoft Systems Management Server (SMS) product. A limited version, which can capture frames that are sent to or from the network adapter of the computer on which Network Monitor is installed, is included with Windows Server 2003. |
| Registry Editor | The registry editor Regedit.exe allows viewing and editing of registry parameters. |
| Simple Network Management Protocol (SNMP) service | Provides statistical information about network devices to SNMP management systems, such as HP OpenView, IBM NetView, or Sun Net Manager. |
| System Monitor | Analyzes TCP/IP network performance. System Monitor is a component of in the Performance Logs and Alerts snap-in. |

**TCP/IP Diagnostic Tools**

TCP/IP supports several diagnostic tools in Windows Server 2003 that include new support for IPv6:

- Netstat
- Pathping
- Ping
- Tracert

In addition, Windows Server 2003 introduces the **Netsh** commands for both IPv6 and for Ipv6 6to4. You can use the first set of

tools to query and configure IPv6 interfaces, addresses, caches, and routes. You can use the second set of commands to query or configure the 6to4 service on either a 6to4 host or a 6to4 router.

The following table lists TCP/IP diagnostic tools that you can use to identify and resolve TCP/IP networking problems.

**TCP/IP Diagnostic Tools**

| Tool | Description |
|---|---|
| Arp | View and manage the Address Resolution Protocol (ARP) cache on the interfaces of the local computer. |
| Hostname | Display the name of the computer. |
| Ipconfig | Display current TCP/IP network configuration values, update or release Dynamic Host Configuration Protocol (DHCP) allocated leases, and display, register, or flush Domain Name System (DNS) names. |
| Nbtstat | Check the state of current NetBIOS over TCP/IP (NetBT) connections, view and update the NetBIOS name cache, and determine the names registered with Windows Internet Name Service (WINS). |
| Netdiag | Check the state of a network client to help isolate TCP/IP-related connectivity problems, including verifying that DNS is available and functioning correctly. |
| Netsh | Provides thirteen sets of commands (called contexts) for performing a wide range of network configuration tasks. Windows Server 2003 adds a new context for managing IPv6 to the **netsh** command set. |
| Netstat | Display statistics for current TCP/IP connections. Windows Server 2003 adds IPv6 parameters to the **netstat** command. |
| Nslookup | Check records, domain host aliases, domain host services, and operating system information by querying DNS servers. |
| Pathping | Trace a path to a remote system and report packet losses at each router along the way. Windows Server 2003 adds IPv6 parameters to the **pathping** command. |
| Ping | Send Internet Control Message Protocol (ICMP) Echo messages to verify IP connectivity. Windows Server 2003 adds IPv6 parameters to the **ping** command. |
| Route | Display the IP routing table, and add, edit, or delete IPv4 routes. Route for Windows Server 2003 also displays IPv6 routes. |
| Tracert | Trace a path to a destination. Windows Server 2003 adds IPv6 parameters to the **tracert** command. |

All of the tools in the previous table, except **Netdiag**, are installed with the Windows Server 2003 operating system. **Netdiag** is one of the support tools delivered with the second Windows Server 2003 operating system CD, which you install separately from the operating system itself.

**Arp.exe: Arp**
**Category**
The **Arp** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Arp** command-line tool to display and modify entries in the local Address Resolution Protocol (ARP) cache. The ARP cache, which is a memory-resident list, contains one or more tables that store IP addresses and the corresponding Ethernet, Token Ring, or wireless LAN physical addresses that have been resolved from other computers on the same subnet. Typically, a physical address is the Media Access Control (MAC) address. A separate table exists for each network adapter installed on the computer.

**Hostname.exe: Hostname**
**Category**
The **Hostname** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Hostname** command-line tool to display the host name of the computer on which you run the command.

**Hostname** does not display the fully qualified domain name (FQDN) of the computer. You can find the computer name and its FQDN through Control Panel on the computer.

**Ipconfig.exe: Ipconfig**
**Category**
The **IPconfig** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Ipconfig** command-line tool to display the current configuration of the installed IP stack on a networked computer and to refresh DCHP and DNS settings. The **Ipconfig** command is often one of the first commands you use to check the status of the connection when you experience communication problems on a TCP/IP network. Ipconfig is most useful for managing computers that obtain an IP address automatically, such as by using DHCP or through alternate configuration.

When used without a parameter, **Ipconfig** displays the IPv4 address, subnet mask, and default gateway for all adapters on a computer. If the computer has the IPv6 protocol installed, **Ipconfig** also displays the IPv6 address information.

**Nbstat.exe: Nbtstat**
**Category**
The **Nbtstat** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Nbtstat** command-line tool to troubleshoot NetBIOS name-resolution problems.

When a network is functioning correctly, NetBIOS over TCP/IP (NetBT) resolves NetBIOS names to IP addresses. NetBT uses several options for NetBIOS name resolution, including local NetBIOS name cache lookup, WINS server query, broadcast, and LMHOSTS lookup. After NetBT methods are exhausted, Windows TCP/IP converts the NetBIOS name to a host name and attempts host name resolution, including checking the local host name, checking the DNS client resolver cache, and querying DNS servers.

Use **Nbtstat** to display a variety of information, including:

NetBT protocol statistics.

NetBIOS name tables for both the local computer and for remote computers. The NetBIOS name table is the list of NetBIOS names that corresponds to NetBIOS applications running on that computer.

NetBIOS name cache. The NetBIOS name cache is the table that contains NetBIOS name-to-IP address mappings.

Also use **Nbtstat** to refresh the NetBIOS name cache and the names registered with WINS.

**Netdiag.exe: Netdiag**
**Category**
**Netdiag** is a command-line tool delivered with the support tools on the second Windows Server 2003 operating system CD, which you install separately from the operating system itself.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Netdiag** command-line tool, which is one of the support tools delivered with the second Windows Server 2003 operating system CD, to help isolate TCP/IP-related connectivity problems. After you install the Support Tools, Netdiag.exe appears in the Program Files\Support Tools folder on the drive on which the Support Tools are installed.

**Netdiag** provides tests that report information about a computer and its network configuration, which can help diagnose network problems. Use the **Netdiag** logging parameter (**-l**) to capture and store output from **Netdiag** tests. One use of **Netdiag** logging is to add the output of Netdiag tests to your network's baseline documentation and to create updated test logs each time important changes to a computer's configuration are made.

In addition to troubleshooting TCP/IP issues, you can also use Netdiag to examine a computer's Internetwork Packet Exchange (IPX) and NetWare configurations.

### Note

- The IPX protocol is supported only on the 32-bit platforms in Windows Server 2003.

Despite the name similarity and some overlap in function, the Netdiag command-line tool that you install from the Windows Server 2003 operating system CD as part of the set of Support Tools is not the same as the following two versions of a different network diagnostic tool:

**Network Diagnostics**, which you access under Help and Support Center Tools in Help and Support Center for Windows Server 2003.

**Netsh diagnostic** (**diag**)commands, which you access by typing **netsh -c diag** at a command prompt.

**Network Diagnostics** and the **Netshdiag** commands use the same Dynamic Link Library (DLL) and run the same set of tests. However, **Netsh** lets you test user-defined destinations, whereas **Network Diagnostics** does not.

Although there is some overlap between these tools and the **Netdiag** support tool described here, for the most part **Netdiag** provides a different set of tests.

**Netsh.exe: Netsh**
**Category**
The **Netsh** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Netsh** command-line tool to locally or remotely display or modify the configuration of services or protocols on computers running Windows. In addition, the **Netsh** command-line interface is scriptable, which lets you perform batch configurations or administration from a centralized location.

The **Netsh** commands, which were first introduced in Windows 2000, help you manage a standard TCP/IP network, that is, an IPv4 network. For organizations that have begun to explore IPv6 by introducing IPv6 into their network, Windows Server 2003 and Windows XP add a new set of **Netsh** commands for IPv6.

For administrative convenience, the **Netsh** commands are grouped into sets, called contexts. Each context provides commands appropriate for a specific area of networking functionality. A context is implemented through an associated **Netsh** helper, which is a dynamic link library (DLL) file that provides the capabilities for that context. **Netsh** directs the context command that you enter to the appropriate helper, and the helper then carries out the command. The **Netsh** helper DLLs interact with other operating system components, such as WINS or TCP/IP. For example, Winsmon.dll provides the set of commands for managing the WINS service, and Ifmon.dll provides the set of commands for managing IPv4 interfaces.

A context can contain one or more additional contexts, called a subcontext. For example, the IPv6 context contains the 6to4 subcontext, which provides a separate set of commands for managing the 6to4 service (which encapsulates IPv6 traffic with an IPv4 header before it is sent over an IPv4 network) on either a 6to4 host or a 6to4 router.

**Netsh** is extensible. Developers can create additional contexts to manage networking services in addition to the contexts provided by Windows.

**Netstat.exe: Netstat**
**Category**
The **Netstat** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Netstat** command-line tool to display active TCP/IP connections, ports on which the computer is listening, Ethernet statistics, the IP routing table, IPv4 statistics (for the IP, ICMP, TCP, and UDP protocols), and IPv6 statistics (for the IPv6, ICMPv6, TCP over IPv6, and UDP over IPv6 protocols).

**Nslookup.exe: Nslookup**
**Category**
The **Nslookup** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Nslookup** command-line tool for querying and troubleshooting the DNS infrastructure. For example, **Nslookup** can provide host-name resolution.

**Pathping.exe: Pathping**
**Category**
The **Pathping** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use **Pathping**, a command-line tool that combines features of **Ping** and **Tracert**, to obtain additional information that neither of those tools provides. Specifically, you can use **Pathping** to identify the route to a remote host, then ping the remote host for a period of time to collect and report statistics. Pathping information includes information about the intermediate routers visited on the path, the Round-Trip Time (RTT) value, and link-loss information.

**Ping.exe: Ping**
**Category**
The **Ping** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Ping** command-line tool as your primary tool for troubleshooting IP-level connectivity between two TCP/IP computers. **Ping** sends ICMP Echo or ICMPv6 Echo Request messages to perform network diagnostics and to test ability to reach a specific destination. **Ping** can use IPv4 or IPv6 addresses. If a name is specified, **Ping** uses the address that is resolved.

**Ping** lets you specify the size of packets to use (the default is 32 bytes), how many to send, whether to record the route used, what Time-To-Live (TTL) value to use, whether or not to set the **Don't Fragment** flag, and so on.

**Route.exe: Route**
**Category**
The **Route** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Route** command-line tool to view and modify the local IP routing table.

For two hosts to exchange IP datagrams, they must both have a route to each other, or they must use a default gateway that knows a route between the two. Typically, routers exchange information using a protocol such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). The RIP Listening service is available for Microsoft Windows XP Professional, and full routing protocols are supported by Windows Server 2003 in the Routing and Remote Access service.

All symbolic names used for the specified network destination of the route are looked up in the network database file NETWORKS. The symbolic names for the gateway are looked up in the host name database file HOSTS. If the command is **print** or **delete**, the destination value can be a wildcard value specified by an asterisk ("*****"). If the destination specified contains a ***** or **?**, it is treated as a shell pattern and only matching destination routes are printed. The asterisk matches any string, and the question mark matches any one character. For example, **157.*.1**, **157.***, **127.***, and ***224*** are all valid uses of the wildcard asterisk.

Using an invalid combination of a destination and netmask value generates a "The route addition failed: The specified mask parameter is invalid. (Destination & Mask) != Destination." error. This sort of error message appears, for example, when a bitwise logical AND between the destination and mask does not equal the destination value.

**Tracert.exe: Tracert**
**Category**
The **Tracert** command-line tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Tracert** command-line route-tracing tool to display the path between the sending host and a destination.

The path that **Tracert** displays is a list of near-side router interfaces of the routers in the path between the source host and destination. **Tracert** uses the IP Time-to-Live (TTL) field in Internet Control Message Protocol (ICMP) Echo Requests and ICMP Time Exceeded-TTL Exceeded in Transit messages to determine the path from a source to a destination through an IP internetwork.

Note that some routers silently drop packets with expired TTLs. These routers do not appear in the **Tracert** display.

**Tracert** works by incrementing the TTL value by one for each ICMP Echo Request it sends, and then waiting for an ICMP Time Exceeded-TTL Exceeded in Transit message. The TTL values of the **Tracert** packets start with an initial value of one; the TTL of each trace after the first is incremented by one. A packet sent out by **Tracert** travels one hop further on each successive trip.

### Note

- The UNIX version of **Tracert** performs the same function as the Windows version, except that the IP payload is a UDP packet addressed to a (presumably) unknown destination UDP port. Intermediate routers send back ICMP Time Expired-TTL Exceeded in Transit messages recording the route taken, and the final destination sends back an ICMP Destination Unreachable-Port Unreachable message.

- The UDP payload from the UNIX **Tracert** tool can cross routers and some firewalls, whereas the ICMP Echo Request messages might not, due to ICMP filtering. To avoid this problem in Windows Server 2003, turn off packet filtering and then try using **Tracert** again.

**TCP/IP Remote Networking Tools**

Microsoft TCP/IP includes several remote networking tools. The following table lists the tools included with Microsoft TCP/IP that you can use to communicate with remote computers.

**TCP/IP Remote Tools**

| Tool | Description |
|---|---|
| Finger | Displays information about a user or users on a specified remote computer (typically, a computer running UNIX) that is running the Finger service. |
| Ftp | Transfers files over the Internet to, and from, a computer running a File Transfer Protocol (FTP) server service, such as the FTP component of Microsoft Internet Information Services (IIS). |
| Rcp | Copies files between a computer running Windows Server 2003 and a computer running Rshd, the UNIX remote shell service. |
| Rexec | Runs commands on a remote computer (typically, a computer running UNIX). |
| Rsh | Runs commands on remote hosts using the Rsh service, the UNIX remote shell service. |
| Telnet | Starts terminal emulation with a remote host running a Telnet server service. |
| Tftp | Transfers files to and from a remote computer (typically, a computer running UNIX) that is running the Trivial File Transfer Protocol (TFTP) service. |

**Note**

- All passwords used by Windows networking services are encrypted. However, the **Ftp**, **Rexec**, and **Telnet** connectivity tools rely on plaintext password authentication by the remote computer. Plaintext passwords are not encrypted before being sent over the network. This enables another user equipped with a network capture tool such as Network Monitor on the same network to obtain a user's remote account password. For this reason, choose different passwords from those used for computers running Windows Server 2003 or domains when connecting to non-Microsoft remote computers with the **Ftp**, **Rexec**, or **Telnet** tools. Note that the protocols themselves prohibit encryption; the use of plaintext passwords is not recommended by Microsoft.

**Finger.exe: Finger**
**Category**
The **Finger** connectivity tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Finger** connectivity tool to display information about a user or users on a specified remote computer (typically, a host computer running UNIX) that is running the Finger service. Output varies based on how the remote host specifies **Finger** output format.

**Ftp.exe: Ftp**
**Category**
The **Ftp** connectivity tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

FTP is a protocol that defines how to transfer files from one computer to another over a TCP/IP network, such as the Internet or a company intranet. Use the **Ftp** connectivity tool to transfer files to and from a host running an FTP server service, such as the FTP component of Microsoft Internet Information Services (IIS). You can use **Ftp** interactively or in batch mode to process ASCII text files.

### Note

- The FTP service is a component of Microsoft Internet Information Services (IIS). However, when you install IIS on a server, FTP is not installed unless you explicitly specify that it be installed. If you install IIS without FTP, you can use **Add or Remove Windows Components** in **Add or Remove Programs** in **Control Panel** to install FTP later.

Typically, users use **Ftp** to download files from a location on the Internet or to upload files to a location on the Internet. Before a user can log on to an FTP server, the server administrator must grant the user permission to access the **inetput\ftproot** folder, which is created along with the Default FTP site by default when you install IIS FTP, or a folder for the use of a particular user or group under the ftproot folder. Alternatively, you can create a new FTP site, and you can use folders other than ftproot and its children for FTP files.

Similarly, users can upload a file from the local computer to an FTP server using the **put** command. However, for users to successfully use the **put** command, the administrator must grant them write permissions for the FTP folder to which the file will be uploaded.

**Rcp.exe: Rcp**
**Category**
The **Rcp** connectivity tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Rcp** connectivity tool to copy files between a computer running Windows Server 2003 and a computer running **Rshd**, the UNIX remote shell server service or daemon. **Rshd** stands for **Rsh d**aemon; daemon is the UNIX term for service. In addition to running **Rsh**, the remote computer must also support the **Rcp** tool.

Also use **Rcp** for third-party transfer to copy files between two computers running **Rsh** when the command is issued from computers running Windows Server 2003. The UNIX **Rsh** server service is not available on computers running Windows Server 2003, but the computer running Windows Server 2003 can participate as the computer from which the commands are issued.

**Permitting Network Access**
The .rhosts file typically permits network access on UNIX systems. The .rhosts file lists computer names and associated logon names that have access to remote computers. When you run **Rcp**, **Rexec**, or **Rsh** tools remotely with a correctly configured .rhosts file, you do not need to provide logon and password information for the remote computer. The .rhosts file must be in the user's home directory on the remote computer.

The .rhosts file specifies which remote computers or users can access a local account using the **Rcp** or **Rsh** commands (**Rsh** is described later in this chapter). For you to access the remote computer using **Rcp**, the .rhosts file (or a file called Hosts.equiv) must exist on the remote computer. **Rcp** transmits the local user name to the remote computer. The remote computer uses this name and the IP address (usually resolved to a computer name) of the requesting computer to determine whether to grant access. No provision exists for specifying a password to access an account using **Rcp**.

If the user is logged on to a Windows Server 2003 domain, the domain controller must be available to resolve the currently logged-on name because the logged-on name is not cached on the local computer. Because the user name is required as part of the Rcp protocol, the command fails if it cannot obtain the user name.

The .rhosts file is a text file in which each line is an entry. An entry consists of the local host name, the local user name, and any comments about the entry. Each entry is separated by a tab or space, and comments begin with the number sign (**#**).

For more information about the implementation of the .rhosts file on a remote computer, see the remote computer's documentation.

**Specifying Computers (Hosts)**
Use Host.user to specify a user name other than the current user name. If you use Host.user with **Source** (which specifies which files to copy), the .rhosts file on the remote computer must contain an entry for user. For example:

```
Rcp host99.user7:file1 corp7.admin:file2
```

In this example, the .rhosts file on Corp7 must have an entry for User7 on Host99.

If you type a computer name as a fully qualified domain name (FQDN) containing periods, you must append a user name to the host name. This prevents the last element of the domain name from being interpreted as a user name. For example:

```
Rcp domain-name1.user:johns domain-name2.user:buddyg
```

**Performing Remote Processing**
Remote processing is performed by a command run from the user's logon shell on most UNIX computers. The user's **.profile** or **.cshrc** file is run before file names are parsed. You can specify that exported shell variables be used in remote file names by using one of the escape characters. The escape characters are the backslash (\), quotation mark ("), and apostrophe (').

**Copying Files**
If you try to copy several files to a file rather than to a directory, only the last file is copied. Also, the **rcp** command cannot copy a file onto itself (*Source* and *Path/Destination* cannot be the same).

**Rexec.exe: Rexec**
**Category**
The **Rexec** connectivity tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Rexec** connectivity tool to run commands on a remote computer that is running the Rexecd service, the UNIX remote execute service. Rexecd stands for Rexec daemon; daemon is the UNIX term for service. Before carrying out the specified command, **Rexec** authenticates the user name on the remote host by prompting for a password. Windows Server 2003 does not provide the Rexec service.

**Rexec** copies standard input to the remote command, standard output of the remote command to its standard output, and the standard error of the remote command to its standard error. Interrupt, Quit, and Terminate signals are propagated to the remote command. **Rexec** generally terminates when the remote command completes.

You cannot use **Rexec** to run most interactive commands. For example, you cannot use **Rexec** to run **vi** or **emacs** commands. Use Telnet to run interactive commands.

**Rsh.exe: Rsh**
**Category**
The **Rsh** connectivity tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Rsh** connectivity tool to run commands on remote computers that are running the Rsh service, the UNIX remote shell service. For information about the .rhosts file used to enable this tool, see the description of the **Rcp** tool earlier in this section.

If the user is logged on to a Windows Server 2003 domain, the domain controller must be available to resolve the user name, because it is not cached on the local computer. Because the user name is required as part of the Rsh protocol, the command fails if the user name cannot be obtained.

**Rsh** copies standard input to the remote command, standard output of the remote command to its standard output, and the standard error of the remote command to its standard error. **Rsh** generally terminates when the remote command completes.

**Telnet.exe: Telnet**
**Category**
The **Telnet** connectivity tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

The **Telnet** connectivity tool starts terminal emulation with a remote host running a Telnet server service. The remote computer must also be using TCP/IP. Telnet provides DEC VT 100, DEC VT 52, or ANSI emulation, using the connection-based services of TCP.

To provide terminal emulation from a computer running Windows, the remote host must be running TCP/IP and a Telnet Server service. The Windows-based Telnet user must also have a user account on the remote Telnet Server.

**Note**

- Windows Server 2003 and Microsoft Windows XP Professional provide the Telnet Client and Telnet Server components. These are built in, but you must use the Services snap-in to start the **Telnet** service before it can serve Telnet clients.

**Tftp.exe: Tftp**
**Category**
The **Tftp** connectivity tool is installed with the operating system.

**Version compatibility**
This tool is included with all versions of Windows that include TCP/IP.

Use the **Tftp** connectivity tool to transfer files over the Internet to and from a remote computer (typically, a computer running UNIX) that is running the Trivial File Transfer Protocol (TFTP) service. This tool is similar to **Ftp**, but it does not provide user authentication. However, the files you transfer using **Tftp** do require UNIX read and write permissions. You can use **Tftp** only for unidirectional transfer of files.

# IPv6 Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The current Internet has fueled significant innovation and growth in network computing. To support this growth and the demand for more collaborative communication experiences, the current version of the Internet Protocol (IP), called IP Version 4 (IPv4), is being replaced with a new standard.

IPv4 is an agreed-upon set of protocols, or rules, that allow computers to communicate with each other by specifying the format of packets and the addressing scheme.

The new version of IP, called IP version 6 (IPv6), performs the same functions as IPv4 but resolves unanticipated IPv4 design issues.

As a core networking protocol in the Microsoft Windows Server 2003 operating systems, IPv6 serves as one of two Internet protocols that enable computers that are running Microsoft Windows operating systems to communicate on intranets and over the Internet.

This subject describes how IPv6 relates to other networking protocols, the functions that IPv6 performs, how IPv6 addresses are structured and assigned, how IPv6 packets are structured and routed, and how IPv6 can interoperate with IPv4.

**In this subject**


- What Is IPv6?
- How IPv6 Works
- IPv6 Tools and Settings

# What Is IPv6?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is IPv6?**

**In this section**


- Introduction to IPv6
- IPv6 in Windows Server 2003
- Transitioning from IPv4 to IPv6
- Limitations of IPv4
- IPv6 Features that Fix IPv4 Limitations
- Differences Between IPv4 and IPv6
- Related Information


**Introduction to IPv6**

Microsoft is delivering support for the emerging update to the Internet Protocol, commonly referred to as IP version 6 — or simply IPv6 (RFC 2460). This protocol suite is based on a standard from the Internet Engineering Task Force (IETF), and it is designed to significantly increase the size of the address space used to identify communication endpoints in the Internet, thereby allowing it to continue its tremendous growth rate. As a core networking protocol in Windows Server 2003, IPv6 serves as one of two Internet protocols that enable computers running Windows to communicate on intranets and over the Internet.

The recent broad adoption of always-on technologies such as Digital Subscriber Line (DSL) and cable modems, coupled with the pending integration of personal data assistants (PDAs) and cellular phones into always-addressable Mobile Information Appliances, significantly elevates the urgency to expand the address space that Internet-connected systems use to communicate. The address space currently used is defined as part of the Internet Protocol, or IP (the network layer of the TCP/IP protocol suite). The version of IP commonly used today, Version 4 (IPv4), has not been substantially changed since RFC 791 was published in 1981. Since then, IPv4 has proven to be robust, easily implemented, and interoperable. It has withstood the test of scaling an internetwork (a network of networks) to a global utility the size of today's Internet.

IPv6 will continue the tradition of the IPv4 protocol, which gained much of its acceptance by defining mechanisms to tie systems together over a wide variety of disparate networking technologies. Already defined link-layer mappings for transporting IPv6 include Ethernet, Point-to-Point Protocol (PPP), Fiber Distributed Data Interface (FDDI), Token Ring, Asynchronous Transfer Mode (ATM), Frame Relay, and IEEE 1394. From an architectural perspective, an IPv4-based infrastructure appears to systems that are enabled for IPv6 as a single segment non-broadcast multi-access (NBMA) network. The capability to send IPv6 traffic over existing IPv4 networks will provide an initial reach as broad as the current Internet, limited only by the endpoints' ability and readiness to make use of it.

New capabilities that are expected to drive rapid adoption include scoped addresses; stateless autoconfiguration, which lowers complexity and management burden; and mandatory IP security (IPSec), which permits end-to-end data authentication and integrity and increases privacy of connections. In addition, technologies that extend the lifetime of IPv4 (such as network address translation, or NATs) frequently do not work with existing applications, and those technologies are already restricting the flexibility to deploy new applications. NATs are popular today because they allow multiple systems to share a single public IPv4 address. However, they tend to enforce a client/server usage model where the client uses private address space with only the server existing in public address space. IPv6 brings back the capability of "end-to-end control of communications," making networking applications simpler as the network again becomes transparent.

**IPv6 in Windows Server 2003**

Wireless technologies are emerging in ways that make ad-hoc networks between personal devices more feasible. Setting up systems to work in an ad-hoc mode is challenging enough, but many of these personal devices will also need to perform in the managed environment of the workplace. Switching between these modes is frequently frustrating and is significantly more

involved than using either mode on its own. To reduce the complexity, IPv6 has defined an architectural principle that systems are required to simultaneously support multiple addresses. Coupling this capability with scoped addresses results in the ability to move easily and automatically between ad-hoc and managed environments. Tthe IPv6 implementation will automatically adapt itself to the current needs, be it ad-hoc, home, or business connections.

To address concerns about security and privacy, the Microsoft IPv6 implementation includes IPSec, which provides data authenticity, data integrity, and data confidentiality across the array of protocols used by the various applications. Providing the capability at the network layer frees developers from having to add specific security capabilities to every application.

In addition, Microsoft helped standardize temporary addresses. To make stateless autoconfiguration work well and to ensure global uniqueness, the standards community chose the underlying hardware address (the MAC address) for use as part of an IPv6 address. The side effect of this approach is that all communications are traceable to the specific hardware device.

Although it is technically necessary to have a published (over some scope), globally unique address to receive incoming connections, the address of an originator requires only current global uniqueness (not publication). To alleviate this potential privacy concern, Microsoft has authored RFC 3041 to define a locally generated address mechanism where the result is valid only for a period that the local system or application determines. The ability of IPv6 systems to simultaneously support multiple addresses allows each application to use an independent address, an application to use a different address for each service to which it connects, or both.

Peer-to-peer applications that are made easier using IPv6 include IP telephony and video tele-conferencing. These and similar applications are likely to take advantage of the Quality of Service (QoS) features defined for IPv6. Although many QoS features have also been defined as add-ons for IPv4, the mechanism selected was to redefine the meaning for the Type of Service field of the IP header, which caused collisions with historical implementations. The effort to provide QoS for IPv4 has been a struggle due to differing models of deployment. This effort is not wasted though, because it is forcing many details to be worked through — from hardware capabilities to business practices. Systems that are enabled for IPv6 will be able to leverage this effort to provide an array of service levels that are consistent from end to end.

**Transitioning from IPv4 to IPv6**

The conversion from IPv4 to IPv6 will be a larger task for the industry than was the preparation for year 2000. This protocol change will affect nearly all networked applications, end systems, infrastructure systems, and network architectures. This change must be approached with responsibility to prevent the costly, unproductive missteps that often result from broad, premature availability of technologies. Unlike the year 2000 issue, the conversion to IPv6 has no specific timeline. However, as noted earlier, the rate of IPv4 address consumption is rapidly increasing. Simplicity of deployment will be the key to rapid adoption.

Like IPv4 (where early deployments frequently transited X.25 networks), IPv6 deployment will start at the edge of the network, taking advantage of framing within any available network technology. Internet service providers (ISPs) will deploy native IPv6 routing based on customer demand. However, this conversion may be slow because ISPs will need several years to replace network equipment. Microsoft is taking the approach that encapsulating IPv6 packets within IPv4 will allow incremental deployments of end systems that will, in turn, demonstrate the demand to the ISPs.

To stay on the high performance path of the existing routers, computers that are running Windows and that are enabled for IPv6 will default to tunneling over IPv4 unless the ISP provides a specific indication to do otherwise and a native IPv6 path exists end to end. The only requirement is that systems that are directly connected to an ISP must receive at least one public IPv4 address. (The address ranges specified in RFC 1918 are not public.) Other systems in a home or business will receive 6to4 (RFC 3056) prefix Router Advertisement messages from the directly connected system.

In the presence of NATs that are not enabled for IPv6 where only private addresses are available, a supplementary technology will be used. This technology will tunnel IPv6 traffic over NATs by including a User Datagram Protocol (UDP) header that can be used to provide a mechanism for 6to4-type tunneling across the IPv4 Internet. In enterprise environments, an incremental upgrade to IPv6 is possible using the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP). ISATAP allows IPv6-only hosts and subnets to fully co-exist and interoperate with IPv4 hosts and subnets in an intranet. In partnership with 6to4 technology, a comprehensive incremental migration solution is available to businesses that are taking their corporate networks through this transition.

Despite these approaches, the transition will not be easy. Most manufacturers will produce systems that support both IPv4 and IPv6 so that, if connections are not possible using IPv6, the systems can fall back and succeed using IPv4 (if IPv4 connectivity existed before the introduction of IPv6). The overall goal is to ensure a smooth transition and deployments where updated applications can take advantage of the new protocol without breaking existing functionality. To this end, new Windows APIs have been defined to specifically isolate the legacy applications from unintentional exposure to protocol differences, including the larger IPv6 addresses.

Microsoft has taken four key steps to deliver IPv6:

- In 1998, Microsoft Research released an IPv6 implementation to help the community that was developing standards for IPv6 understand and test the protocol during its definition.

- In March 2000, a technology preview was released for computers that were running Windows 2000. This release helped developers become familiar with the concepts and capabilities they would encounter when they enabled their applications to use IPv6.

- In October 2001, Windows XP was released with a developer preview IPv6 stack and key components of the system enabled for IPv6 so developers could begin the task of enabling their applications for IPv6. It is also expected that early adoption customers will start using IPv6 in test labs. This testing will allow those customers to have better visibility into managing their eventual rollout, and it will help identify any issues that need to be addressed in networked products.

- In November 2001, beta 3 for Windows Server 2003 was released with the first edition of the Microsoft production stack and components enabled for IPv6.

Current releases, including both Windows Server 2003 and Windows XP with Service Pack 1, provide general availability of this stack. Customers can begin production rollout of systems and applications that take advantage of this protocol.

The Microsoft implementation of IPv6 is easy to deploy because it includes stateless address autoconfiguration (including temporary addresses), automatic tunneling over existing IPv4 networks, and appropriate use of scoped addresses.

Because IP is a fundamental and pervasive technology within the operating system, it is not feasible to retrofit versions of Windows prior to Windows Server 2003 and Windows XP. However, to maintain backwards compatibility, versions of Windows that are enabled for IPv6 will also provide the capability to natively communicate using IPv4 for the foreseeable future. Although translation between IPv4 and IPv6 will be necessary in some cases (such as late in a transition when new IPv6-only devices need to access yet-to-be-retired IPv4-only systems), it is not expected to be the norm for early deployments. Whenever the IPv6-only devices arrive, the issues that surround address translation are typically specific to a given application. Thus as they arise, these scenarios will require targeted development on a case-by-case basis.

**Limitations of IPv4**
IPv4's initial design did not anticipate the following:

- The recent exponential growth of the Internet and the impending exhaustion of the IPv4 address space.

  Because IPv4 addresses have become relatively scarce, some organizations have been forced to use NATs to map multiple private addresses to a single public IP address. NATs promote reuse of the private address space, but they do not support standards-based network-layer security or the correct mapping of all upper layer protocols. NATs can also create problems when they connect two organizations that use the private address space.

  Additionally, the rising prominence of Internet-connected devices and appliances ensures that the public IPv4 address space will eventually be depleted.

- The requirement for security at the Internet layer.

  Private communication over a public medium such as the Internet requires encryption services that protect the data being sent from being viewed or modified in transit. IPSec provides security for IPv4 packets, but this standard is optional, and proprietary solutions prevail.

- 7The growth of the Internet and the ability of Internet backbone routers to maintain large routing tables.

  Because of the way that IPv4 network IDs have been and are currently allocated, the routing tables of Internet backbone routers routinely contain more than 85,000 routes. The routing infrastructure of the IPv4 Internet combines both flat and hierarchical routing.

- The need to better support real-time delivery of data — also called quality of service (QoS).

  QoS standards exist for IPv4, but real-time traffic support relies on the IPv4 Type of Service (TOS) field and the identification of the payload, which is typically done using a UDP or TCP port. Unfortunately, the IPv4 TOS field has limited functionality, and various local interpretations developed over time. In addition, payload identification using a TCP

and UDP port is not possible when the payload is encrypted.

**IPv6 Features that Fix IPv4 Limitations**
IPv6 includes the following features:

- New header format
- Larger address space
- Efficient and hierarchical addressing and routing infrastructure
- Stateless and stateful address configuration
- Built-in security
- Better support for QoS
- New protocol for neighboring node interaction
- Extensibility

The following subsections discuss each of these new features in detail.

**New Header Format**
The IPv6 header has a new format that is designed to minimize header overhead. This optimization is achieved by moving both non-essential fields and optional fields to extension headers that appear after the IPv6 header. Intermediate routes can process the streamlined IPv6 header more efficiently.

IPv4 headers and IPv6 headers do not interoperate. IPv6 is not a superset of functionality that is backward compatible with IPv4. A host or router must use an implementation of both IPv4 and IPv6 to recognize and process both header formats. The IPv6 header is only twice as large as the IPv4 header, even though IPv6 addresses are four times as large as IPv4 addresses.

**Larger Address Space**
IPv6 has 128-bit (16-byte) source and destination IP addresses. Although 128 bits can express over 3.4×1038 possible combinations, the large address space of IPv6 has been designed for multiple levels of subnetting and address allocation from the Internet backbone to the individual subnets within an organization.

Even though only a small number of the possible addresses are currently allocated for use by hosts, plenty of addresses are available for future use. With a much larger number of available addresses, address-conservation techniques, such as the deployment of NATs, are no longer necessary.

**Efficient and Hierarchical Addressing and Routing Infrastructure**
IPv6 global addresses that are used on the IPv6 portion of the Internet are designed to create an efficient, hierarchical, and summarizable routing infrastructure that is based on the common occurrence of multiple levels of Internet service providers.

**Stateless and Stateful Address Configuration**
To simplify host configuration, IPv6 supports both stateful address configuration (as in the presence of a DHCP server) and stateless address configuration (as in the absence of a DHCP server). With stateless address configuration, hosts on a link automatically configure themselves with IPv6 addresses for the link (called link-local addresses) and with addresses that they derive from prefixes that local routers advertise. Even in the absence of a router, hosts on the same link can configure themselves with link-local addresses and communicate without manual configuration.

**Built-in Security**
The IPv6 protocol suite requires support for IPSec. This requirement provides a standards-based solution for network security needs and promotes interoperability between different IPv6 implementations.

**Better Support for QoS**
New fields in the IPv6 header define how traffic is handled and identified. Traffic identification (using a Flow Label field in the IPv6 header) allows routers to identify and provide special handling for packets belonging to a flow, which is a series of packets between a source and a destination. Because the IPv6 header identifies the traffic, QoS can be supported even when the packet payload is encrypted through IPSec.

**New Protocol for Neighboring Node Interaction**
The Neighbor Discovery protocol for IPv6 is a series of Internet Control Message Protocol for IPv6 (ICMPv6) messages that manage the interaction of nodes on the same link (known as neighboring nodes). Neighbor Discovery replaces the broadcast-based Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and ICMPv4 Redirect messages with efficient multicast and unicast Neighbor Discovery messages.

**Extensibility**
IPv6 can easily be extended by adding extension headers after the IPv6 header. Unlike options in the IPv4 header, which can support only 40 bytes of options, the size of IPv6 extension headers is constrained only by the size of the IPv6 packet.

**Differences Between IPv4 and IPv6**
The following table highlights some of the key differences between IPv4 and IPv6.

### Differences between IPv4 and IPv6

| IPv4 | IPv6 |
|---|---|
| Source and destination addresses are 32 bits (4 bytes) in length. | Source and destination addresses are 128 bits (16 bytes) in length. For more information, see "IPv6 Addressing" in "How IPv6 Works." |
| IPSec support is optional. | IPSec support is required. |
| IPv4 header does not identify packet flow for QoS handling by routers. | IPv6 header contains Flow Label field, which identifies packet flow for QoS handling by router. |
| Both routers and the sending host fragment packets. | Only the sending host fragments packets; routers do not. |
| Header includes a checksum. | Header does not include a checksum. |
| Header includes options. | All optional data is moved to IPv6 extension headers. |
| Address Resolution Protocol (ARP) uses broadcast ARP Request frames to resolve an IP address to a link-layer address. | Multicast Neighbor Solicitation messages resolve IP addresses to link-layer addresses. For more information, see "Neighbor Discovery" in "How IPv6 Works." |
| Internet Group Management Protocol (IGMP) manages membership in local subnet groups. | Multicast Listener Discovery (MLD) messages manage membership in local subnet groups. |
| ICMP Router Discovery is used to determine the IPv4 address of the best default gateway, and it is optional. | ICMPv6 Router Solicitation and Router Advertisement messages are used to determine the IP address of the best default gateway, and they are required. For more information, see "Neighbor Discovery" in "How IPv6 Works." |
| Broadcast addresses are used to send traffic to all nodes on a subnet. | IPv6 uses a link-local scope all-nodes multicast address. |
| Must be configured either manually or through DHCP. | Does not require manual configuration or DHCP. |
| Uses host address (A) resource records in Domain Name System (DNS) to map host names to IPv4 addresses. | Uses host address (AAAA) resource records in DNS to map host names to IPv6 addresses. |
| Uses pointer (PTR) resource records in the IN-ADDR.ARPA DNS domain to map IPv4 addresses to host names. | Uses pointer (PTR) resource records in the IP6.ARPA DNS domain to map IPv6 addresses to host names. |
| Must support a 576-byte packet size (possibly fragmented). | Must support a 1280-byte packet size (without fragmentation). |

**Related Information**
For more information about TCP/IP, see "TCP/IP Technical Reference."

# How IPv6 Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How IPv6 Works**
**In this section**


- IPv6 Architecture
- IPv6 Addressing
- Neighbor Discovery
- IPv6 Routing
- End-to-End IPv6 Delivery Process
- IPv6 Transition Technologies
- Related Information


Internet Protocol version 6 (IPv6) is a networking protocol that allows Windows users to communicate with other users over the Internet. It interacts with Windows naming services such as Domain Name System (DNS) and uses security technologies such as Internet Protocol security (IPSec), because they help facilitate the successful and secure transfer of IP packets between computers.

Ideally, IPv6 is used in a pure environment, that is, an environment where IPv6 is the exclusive Internet protocol used between computers. Currently, however, pure IPv6 transmissions are attainable only with routers that support IPv6 and computers that are running Windows and that support IPv6. As IPv6 supplants IPv4, pure IPv6 across the Internet will become more prevalent and will eventually replace IPv4. Until that occurs, the transition technologies described in this reference can be used to bridge the technological gap between IPv4 and IPv6.

In addition to describing the transition technologies between IPv4 and IPv6, this subject describes how IPv6 relates to other networking protocols, which functions IPv6 performs, how IPv6 addresses are structured and assigned, and how IPv6 packets are structured and routed.

**IPv6 Architecture**
The IPv6 protocol component that is installed in Windows operating systems is a series of interconnected protocols that include Internet Control Message Protocol version 6 (ICMPv6), Multicast Listener Discovery (MLD), and Neighbor Discovery. These core protocols replace the Internet layer protocols in the Defense Advanced Research Projects Agency (DARPA) model. All protocols above the Internet layer rely on the basic services that IPv6 provides. Protocols at the Host-to-Host Transport and Application layers are largely unchanged, except when addresses are part of the payload or part of the data structures that the protocol maintains. For example, both Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) must be updated to perform new checksum calculations that include IPv6 addresses. TCP must be updated to store IPv6 addresses in its internal Transmission Control Block (TCB). Routing Information Protocol (RIP) must be updated to send and receive IPv6 route prefixes.

The following figure shows the architecture of the IPv6 core protocols in relation to the Open Systems Interconnection (OSI) model, the TCP/IP protocol architecture, and the other protocols in the TCP/IP suite.

**IPv6 Architecture**




The following table explains each of the IPv6 core protocols in more detail.

**IPv6 Core Protocols**

| Protocol | Function |
|---|---|
| IPv6 | IPv6 is a routable protocol that is responsible for the addressing, routing, and fragmenting of packets by the sending host. IPv6 replaces Internet Protocol version 4 (IPv4). |
| ICMPv6 | ICMPv6 is responsible for providing diagnostic functions and reporting errors due to the unsuccessful delivery of IPv6 packets. ICMPv6 replaces ICMPv4. |
| Neighbor Discovery | Neighbor Discovery is responsible for the interaction of neighboring nodes and includes message exchanges for address resolution, duplicate address detection, router discovery, and router redirects. Neighbor Discovery replaces Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and the ICMPv4 Redirect message. |
| Multicast Listener Discovery | Multicast Listener Discovery is a series of three ICMPv6 messages that replace version 2 of the Internet Group Management Protocol (IGMP) for IPv4 to manage subnet multicast membership. |

**IPv6**

Like IPv4, IPv6 is a connectionless, unreliable datagram protocol that is primarily responsible for addressing and routing packets between hosts. Connectionless means that a session is not established before data is exchanged. Unreliable means that delivery is not guaranteed. IPv6 always makes a best effort attempt to deliver a packet. An IPv6 packet might be lost, delivered out of sequence, duplicated, or delayed. IPv6 does not attempt to recover from these types of errors. The acknowledgment of packets delivered and the recovery of lost packets is the responsibility of an upper layer protocol, such as TCP.

An IPv6 packet consists of an IPv6 header and an IPv6 payload. The IPv6 payload consists of zero or more IPv6 extension headers and an upper layer protocol data unit, such as an ICMPv6 message, a TCP segment, or a UDP message. The following figure shows the structure of an IPv6 packet.

**Structure of an IPv6 Packet**




The following table describes the key fields in the IPv6 header.

**Key Fields in the IPv6 Header**

| IP Header Field | Function |
|---|---|
| Source Address | Identifies the IPv6 address of the original source of the packet. |
| Destination Address | Identifies the IPv6 address of the intermediate or final destination of the packet. |

| | |
|---|---|
| Next Header | Identifies either the extension header immediately following the IPv6 header or an upper layer protocol, such as ICMPv6, TCP, or UDP. |
| Hop Limit | Designates the number of subnets, or network segments, on which the packet is allowed to travel before being discarded by a router. The sending host sets the Hop Limit, which prevents packets from endlessly circulating on an IPv6 internetwork. When forwarding an IPv6 packet, routers decrement the Hop Limit by one. |

**IPv6 Extension Headers**

Zero or more extension headers of varying lengths can be present. If any extension headers are present, a Next Header field in the IPv6 header indicates the first extension header. In a typical IPv6 packet, no extension headers are present. The sending host adds one or more extension headers only if either an intermediate router or the destination requires special handling.

Within each extension header is another Next Header field that indicates the next extension header. The last extension header indicates the upper layer protocol (such as TCP, UDP, or ICMPv6) contained within the upper layer protocol data unit.

The IPv6 header and extension headers replace the existing IPv4 header with options. The new extension header format allows IPv6 to be augmented to support future needs and capabilities. Unlike options in the IPv4 header, IPv6 extension headers have no maximum size and can expand to accommodate all the extension data needed for IPv6 communication.

RFC 2460 defines the following extension headers, which all IPv6 nodes must support:

- Hop-by-Hop Options header
- Destination Options header
- Routing header
- Fragment header
- Authentication header
- Encapsulating Security Payload header

**Fragmentation in IPv6**

If an IPv4 router receives a packet that is too large for the network segment to which the packet is being forwarded and fragmentation of the packet is allowed, IPv4 fragments the original packet into smaller packets that fit on the downstream network segment. In IPv6, only the sending host performs fragmentation. If an IPv6 router cannot forward a packet because it is too large, the router sends an ICMPv6 Packet Too Big message to the sending host and discards the packet.

Use of the Fragment extension header facilitates sending host fragmentation and destination host reassembly.

**ICMPv6**

Like IPv4, IPv6 does not provide facilities for reporting errors. Instead, IPv6 uses an updated version of the Internet Control Message Protocol (ICMP) named ICMPv6. ICMPv6 has the common ICMPv4 functions of reporting delivery or forwarding errors and providing a simple echo service for troubleshooting.

The ICMPv6 protocol also provides a framework for the following:

- Neighbor Discovery

  Neighbor Discovery is a series of five ICMPv6 messages that manage node-to-node communication on a link. Neighbor Discovery replaces Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and the ICMPv4 Redirect message. For more information about Neighbor Discovery, see "Neighbor Discovery" later in this section.

- Multicast Listener Discovery (MLD)

  Multicast Listener Discovery is a series of three ICMP messages that manage subnet multicast membership. Multicast Listener Discovery replaces version 2 of the Internet Group Management Protocol (IGMP) for IPv4. For more information about Multicast Listener Discovery, see "Multicast Listener Discovery" later in this section.

ICMPv6 is required for an IPv6 implementation. The following table lists and describes the ICMPv6 messages that are defined in RFC 2463.

**Common ICMPv6 Messages**

| ICMPv6 Message | Function |
|---|---|
| Echo Request | Sent to check IPv6 connectivity to a particular host. |
| Echo Reply | Sent in response to an ICMPv6 Echo Request. |
| Destination Unreachable | Sent by a router or the destination host to inform the sending host that the packet or payload cannot be delivered. |
| Packet Too Big | Sent by a router to inform a sending host that a packet is too large to forward. |
| Time Exceeded | Sent by a router to inform a sending host that the Hop Limit of an IPv6 packet has expired. |
| Parameter Problem | Sent by a router to inform a sending host that an error was encountered in processing the IPv6 header or an IPv6 extension header. |

A series of defined Destination Unreachable messages has been defined. The following table lists and describes the most common of these messages.

**Common Destination Unreachable Messages**

| Destination Unreachable Message | Description |
|---|---|
| No Route Found | Sent by an IPv6 router when a route to the destination IPv6 address cannot be found. |
| Communication Prohibited by Administrative Policy | Sent when administrative policy prohibits communication with the destination host. This message is typically sent when a firewall discards a packet. |
| Destination Address Unreachable | Sent when the destination address is unreachable. This message is typically sent when the destination's link layer address cannot be resolved. |
| Destination Port Unreachable | Sent when the destination port is unreachable. This message is typically sent when an IPv6 packet containing a UDP message arrived at the destination but no applications were listening on the destination UDP port. |

ICMPv6 does not make IPv6 a reliable protocol. ICMPv6 attempts to report errors and to provide feedback on specific conditions. ICMPv6 messages are carried as unacknowledged IPv6 datagrams and are themselves unreliable.

**Path MTU Discovery**

The maximum transmission unit (MTU) for a path is the minimum link MTU of all links on a path between a source and a destination. IPv6 packets that are smaller than the path MTU do not require fragmentation by the host and are successfully forwarded by all routers on the path. To discover the path MTU, the sending host uses the receipt of ICMPv6 Packet Too Big messages.

Sending hosts discover the path MTU through the following process:

1. The sending host assumes that the path MTU is the link MTU of the interface on which the traffic is being forwarded.

2. The sending host sends IPv6 datagrams at the path MTU size.

3. If a router on the path is unable to forward the packet over a link because the packet is larger than the link MTU, the router sends an ICMPv6 Packet Too Big message back to the sending host and discards the packet. The Packet Too Big message contains the MTU of the link on which the forwarding failed.

4. The sending host sets the path MTU for packets being sent to the destination to the value of the MTU field in the Packet Too Big message.

The sending host starts again at step 2 and repeats steps 2 through 4 as many times as necessary to discover the path MTU. The sending host determines the path MTU when the node receives either an acknowledgement from the destination or no more Packet Too Big messages.

RFC 1981 recommends that IPv6 nodes support path MTU discovery. Nodes that do not provide this support must use the minimum link MTU of 1280 bytes as the path MTU.

**Changes in Path MTU**
Changes in routing topology can change the path between source and destination. If a path MTU becomes smaller, sending hosts receive Packet Too Big messages, and the nodes must repeat steps 2 through 4 to discover the new path MTU. If a path MTU becomes larger, sending hosts can detect this change (as described in RFC 1981) by attempting to send a larger IPv6 packet after a minimum of 5 minutes (10 minutes are recommended) upon receiving a Packet Too Big message.

**Neighbor Discovery**
Neighbor Discovery is a set of ICMPv6 messages and processes that determine relationships between neighboring nodes. Neighbor Discovery replaces ARP, ICMP Router Discovery, and ICMP Redirect used in IPv4, and Neighbor Discovery provides additional functionality.

Hosts use Neighbor Discovery to:

- Discover neighboring routers.
- Discover addresses, address prefixes, and other configuration parameters.

Routers use Neighbor Discovery to:

- Advertise their presence, host configuration parameters, and on-link prefixes.
- Inform hosts of a better next-hop address to forward packets for a specific destination.

Nodes use Neighbor Discovery to:

- Resolve the link-layer address of a neighboring node to which an IPv6 packet is being forwarded.
- Determine when the link-layer address of a neighboring node has changed.
- Determine whether a neighbor is still reachable.

The following table lists and describes the Neighbor Discovery processes that RFC 2461 describes.

**IPv6 Neighbor Discovery Processes**

| Neighbor Discovery Process | Description |
|---|---|
| Router discovery | The process by which hosts discover the local routers on attached links. Equivalent to ICMPv4 Router Discovery. For more information, see “Router Discovery” later in this section. |
| Prefix discovery | The process by which hosts discover the network prefixes for local link destinations. Similar to the ICMPv4 Address Mask Request/Reply. For more information, see “Router Discovery” later in this section. |
| Parameter discovery | The process by which hosts discover additional operating parameters, including the link MTU and the default hop limit for outgoing packets. For more information, see “Router Discovery” later in this section. |
| Address autoconfiguration | The process for configuring IP addresses for interfaces in either the presence or absence of a server that provides stateful address configuration using a protocol such as Dynamic Host Configuration Protocol version 6 (DHCPv6). For more information, see “Address Autoconfiguration” later in this section. |
| Address resolution | The process by which nodes resolve a neighbor’s IPv6 address to its link-layer address. Equivalent to ARP in IPv4. For more information, see “Address Resolution” later in this section. |

| | |
|---|---|
| Next-hop determination | The process by which nodes determine the IPv6 address of the neighbor to which a packet is being forwarded based on the destination address. The forwarding or next-hop address is either the destination address or the address of an on-link default router. For more information, see "End-to-End IPv6 Delivery Process" later in this section. |
| Neighbor unreachability detection | The process by which nodes determine that a neighbor is no longer receiving packets. For more information, see "Neighbor Unreachability Detection" later in this section. |
| Duplicate address detection | The process by which nodes determine whether an address considered for use is already in use by a neighboring node. Equivalent to using gratuitous ARP frames in IPv4. For more information, see "Duplicate Address Detection" later in this section. |
| Redirect function | The process of informing a host of a better first-hop IPv6 address to reach a destination. Equivalent to the use of the IPv4 ICMP Redirect message. For more information, see "Redirect Function" later in this section. |

**Multicast Listener Discovery**

Multicast Listener Discovery (MLD) is the IPv6 equivalent of Internet Group Management Protocol version 2 (IGMPv2) for IPv4. MLD is a set of ICMPv6 messages that routers and nodes exchange to enable routers to discover which multicast addresses have listening nodes for each attached interface. As with IGMPv2, MLD discovers only those multicast addresses for which at least one listener exists, not the list of individual multicast listeners for each multicast address. MLD is described in RFC 2710.

In contrast to IGMPv2, MLD uses ICMPv6 messages instead of defining its own message structure. The three types of MLD messages are:

- Multicast Listener Query

  Routers send Multicast Listener Query messages to query a link for multicast listeners. Multicast Listener Query messages are either General Query messages or Multicast-Address-Specific Query messages. Routers send General Query messages to query for multicast listeners of all multicast addresses. Routers send Multicast-Address-Specific Query messages to query for multicast listeners of a specific multicast address. The two message types are distinguished by the multicast destination address in the IPv6 header and a multicast address within the Multicast Listener Query message.

- Multicast Listener Report

  Multicast listeners send Multicast Listener Report messages to either report interest in receiving multicast traffic for a specific multicast address or to respond to a Multicast Listener Query message.

- Multicast Listener Done

  Multicast listeners send Multicast Listener Done messages to report that they are no longer interested in receiving multicast traffic for a specific multicast address.

An MLD message packet consists of an IPv6 header, a Hop-by-Hop Options extension header, and the MLD message. The Hop-by-Hop Options extension header contains the IPv6 Router Alert Option that RFC 2711 describes. This option ensures that routers process MLD messages sent to multicast addresses on which the routers are not listening.

**IPv6 Application Interfaces**

As in IPv4, applications can use the services of IPv6 by using the Windows Sockets application programming interface (API). A socket is defined by a protocol and an address on the host. The format of the address is specific to each protocol. For TCP and UDP traffic over IPv6, the address is the combination of the IPv6 address and port. Two sockets, one for each end of the connection, form a bi-directional communications path.

To communicate with IPv6, an application can use Windows Sockets to specify the protocol (IPv6), the IPv6 address of the destination host, and the port of the destination application. After the application is connected, information can be sent and received.

Until IPv6 is broadly adopted, most hosts will be IPv6/IPv4 hosts, using implementations of both the IPv6 and IPv4 protocols. Developers could modify current applications to test for and then selectively use either IPv4 or IPv6 as their transport for data. However, it makes more sense to modify the applications to use a single Windows Sockets function to perform network operations regardless of whether IPv4, IPv6, or both are installed.

A good example of making applications IP version-independent is to replace all instances of the Windows Sockets function **gethostbyname( )** with the Windows Sockets function **getaddrinfo( )**. The **gethostbyname( )** function is commonly used to resolve a specified name to an IPv4 address and never returns an IPv6 address. The function **getaddrinfo( )** resolves names to both IPv4 and IPv6 addresses, depending on whether IPv4 and IPv6 is installed and which types of addresses the Domain Name System (DNS) name query returns. The application need not test for the existence of both protocols or both types of addresses in the response. By replacing **gethostbyname( )** with **getaddrinfo( )**, the developer makes the application IP version-independent for name resolution.

Windows Sockets has been extended to support IPv6 either directly, using IPv6-specific Windows Sockets functions and structures, or indirectly, using IP version-independent functions. The following figure shows the Windows Sockets architecture for Windows Server 2003.

**Windows Sockets for Both IPv4 and IPv6**




**IPv6 Addressing**

The most obvious distinguishing feature of IPv6 is its use of much larger addresses. The size of an address in IPv6 is 128 bits, which is four times larger than an IPv4 address. A 32-bit address space includes 232 or 4,294,967,296 possible addresses. A 128-bit address space includes $2^{128}$ or 340,282,366,920,938,463,463,374,607,431,768,211,456 (or $3.4\times10^{38}$) possible addresses.

The IPv4 address space was designed in the late 1970s, and it seemed impossible to exhaust. However, addresses were not allocated in a way that anticipated changes in technology and an explosion in the number of hosts on the Internet. The IPv4 address space was consumed to the point that, by 1992, it clearly needed a replacement.

It is even harder to conceive that the IPv6 address space will be consumed. A 128-bit address space provides 655,570,793,348,866,943,898,599 ($6.5\times10^{23}$) addresses for every square meter of the Earth's surface.

It is important to remember that the decision to make the IPv6 address 128 bits long was not so that every square meter of the Earth could have $6.5\times10^{23}$ addresses. Rather, the relatively large size of the IPv6 address is designed to be subdivided into hierarchical routing domains that reflect the topology of the modern Internet. The use of 128 bits allows multiple levels of hierarchy and flexibility in designing hierarchical addressing and routing that is currently lacking in the IPv4-based Internet.

RFC 3513 describes the IPv6 addressing architecture.

**IPv6 Address Syntax**

IPv4 addresses are represented in dotted-decimal format. These 32-bit addresses are divided along 8-bit boundaries. Each set of 8 bits is converted to its decimal equivalent and separated from the other sets by periods. For IPv6, the 128-bit address is divided along 16-bit boundaries. Each 16-bit block is converted to a 4-digit hexadecimal number and separated by colons. The resulting representation is known as colon-hexadecimal.

The following is an IPv6 address in binary form:

```
0010000111011010000000001101001100000000000000000010111100111011
0000001010101010000000001111111111111100010100010011100010110 10
```

The 128-bit address is divided along 16-bit boundaries:

```
0010000111011010   0000000011010011   0000000000000000   0010111100111011
0000001010101010   0000000011111111   1111111000101000   1001110001011010
```

Each 16-bit block is converted to hexadecimal and delimited with colons. The result is:

```
21DA:00D3:0000:2F3B:02AA:00FF:FE28:9C5A
```

IPv6 representation can be further simplified by removing the leading zeros within each 16-bit block. However, each block must have at least a single digit. With leading zero suppression, the address representation becomes:

```
21DA:D3:0:2F3B:2AA:FF:FE28:9C5A
```

**Compressing Zeros**

Some types of addresses contain long sequences of zeros. To further simplify the representation of IPv6 addresses, a contiguous sequence of 16-bit blocks set to 0 in the colon-hexadecimal format can be compressed to "::," known as double-colon.

For example, the link-local address of FE80:0:0:0:2AA:FF:FE9A:4CA2 can be compressed to FE80::2AA:FF:FE9A:4CA2. The multicast address FF02:0:0:0:0:0:0:2 can be compressed to FF02::2.

Zero compression can be used to compress only a single contiguous series of 16-bit blocks expressed in colon-hexadecimal notation. You cannot use zero compression to include part of a 16-bit block. For example, you cannot express FF02:30:0:0:0:0:0:5 as FF02:3::5.

To determine how many 0 bits are represented by the double colon, you can count the number of blocks in the compressed address, subtract this number from 8, and then multiply the result by 16. For example, the address FF02::2 has two blocks (the "FF02" block and the "2" block.) The number of 0 bits expressed by the double colon is 96 (96 = (8 – 2)×16).

Zero compression can be used only once in a given address. Otherwise, you could not determine the number of 0 bits represented by each double colon.

**IPv6 Prefixes**

The prefix is the part of the address that indicates which bits have fixed values or reflect the subnet identifier. Prefixes for IPv6 subnet identifiers and routes are expressed in the same way as Classless Inter-Domain Routing (CIDR) notation for IPv4, that is, in address/prefix-length notation. For example, 21DA:D3::/48 is a route prefix and 21DA:D3:0:2F3B::/64 is a subnet prefix.

### Note

- IPv4 implementations commonly use a dotted decimal representation of the network prefix known as the subnet mask. IPv6 does not support subnet masks. IPv6 supports only the prefix-length notation.

**Types of IPv6 Addresses**

IPv6 supports three types of addresses:

- Unicast

  A unicast address identifies a single interface within the scope of the type of unicast address. With the appropriate unicast routing topology, packets addressed to a unicast address are delivered to a single interface. To accommodate load-balancing systems, RFC 3513 allows multiple interfaces to use the same address as long as they appear as a single interface to the IPv6 implementation on the host.

- Multicast

  A multicast address identifies multiple interfaces. With the appropriate multicast routing topology, packets addressed to a multicast address are delivered to all interfaces that are identified by the address. A multicast address is used for one-to-many communication, with delivery to multiple interfaces.

- Anycast

  An anycast address identifies multiple interfaces. With the appropriate routing topology, packets addressed to an anycast address are delivered to a single interface, the nearest interface that is identified by the address. The nearest interface is defined as being closest in terms of routing distance. An anycast address is used for one-to-one-of-many communication, with delivery to a single interface.

IPv6 addresses always identify interfaces, not nodes. A node is identified by any unicast address that is assigned to one of its interfaces.

**Note**

- RFC 3513 does not define a broadcast address. All types of IPv4 broadcast addressing are performed in IPv6 using multicast addresses. For example, the subnet and limited broadcast addresses from IPv4 are replaced with the link-local scope all-nodes multicast address of FF02::1.

**Unicast IPv6 Addresses**

Unicast IPv6 addresses fall into one of five types:

- Global unicast addresses
- Link-local addresses
- Site-local addresses
- Special addresses
- Compatibility addresses

**Global unicast addresses**

Global unicast addresses are equivalent to public IPv4 addresses. They are globally routable and reachable on the IPv6 Internet.

Unlike the current IPv4-based Internet, which is a mixture of both flat and hierarchical routing, the IPv6-based Internet has been designed from its foundation to support efficient, hierarchical addressing and routing. The scope (that is, the region of the IPv6 internetwork over which the address is unique) of a global unicast address is the entire IPv6 Internet.

The following figure shows the structure of a global unicast address as defined in RFC 3587.

**The Global Unicast Address**




Global unicast addresses contain four fields.

- The three high-order bits are set to 001. The address prefix for currently assigned global addresses is 2000::/3.
- The Global Routing Prefix indicates the global routing prefix for a specific organization's site. The combination of the three fixed bits and the 45-bit Global Routing Prefix creates a 48-bit site prefix, which is assigned to an individual site of an organization. After this prefix is assigned, routers on the IPv6 Internet forward IPv6 traffic matching the 48-bit prefix to the routers of the organization's site.
- The Subnet ID is used within an organization's site to identify subnets. This field is 16 bits long. The organization's site can use these 16 bits within its site to create 65,536 subnets or multiple levels of addressing hierarchy and an efficient routing infrastructure.
- The Interface ID indicates the interface on a specific subnet within the site. This field is 64 bits long.

The following figure shows how fields within the global unicast address create a three-level structure.

**The Three-level Structure of the Global Unicast Address**




The public topology is the collection of larger and smaller ISPs that provide access to the IPv6 Internet. The site topology is the collection of subnets within an organization's site. The interface identifier identifies a specific interface on a subnet within an organization's site. For more information about global unicast addresses, see RFC 3587 in the IETF RFC Database.

**Local-use unicast addresses**
Local-use unicast addresses fall into one of two types:

- Link-local addresses are used between on-link neighbors and for Neighbor Discovery processes.
- Site-local addresses are used between nodes in the same site.

**Link-local addresses**
Nodes use link-local addresses when communicating with neighboring nodes on the same link. For example, on a single-link IPv6 network with no router, hosts use link-local addresses to communicate with other hosts on the link. Link-local addresses are equivalent to Automatic Private IP Addressing (APIPA) IPv4 addresses autoconfigured on computers that are running Windows. APIPA addresses use the 169.254.0.0/16 prefix. The scope of a link-local address is the local link.

A link-local address is required for Neighbor Discovery processes and is always automatically configured, even in the absence of all other unicast addresses. For more information about how link-local addresses are configured, see "Address Autoconfiguration" later in this section.

The following figure shows the structure of the link-local address.

### The Link-Local Address




Link-local addresses always begin with FE80. With the 64-bit interface identifier, the prefix for link-local addresses is always FE80::/64. An IPv6 router never forwards link-local traffic beyond the link.

**Site-local addresses**
Site-local addresses are equivalent to the IPv4 private address space (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16). For example, private intranets that do not have a direct, routed connection to the IPv6 Internet can use site-local addresses without conflicting with global addresses. Site-local addresses are not reachable from other sites, and routers must not forward site-local traffic outside the site. Site-local addresses can be used in addition to global addresses. The scope of a site-local address is the site (the organization internetwork).

Unlike link-local addresses, site-local addresses are not automatically configured and must be assigned through either stateless or stateful address configuration.

The following figure shows the structure of the site-local address.

### The Site-Local Address




The first 10 bits are always fixed for site-local addresses, beginning with FEC0::/10. After the 10 fixed bits is a 54-bit subnet identifier (Subnet ID field) that provides 54 bits with which you can create a hierarchical and summarizable routing infrastructure within the site. After the Subnet ID field is a 64-bit Interface ID field that identifies a specific interface on a subnet.

**Special IPv6 addresses**
The following are special IPv6 addresses:

- Unspecified address

  The unspecified address (0:0:0:0:0:0:0:0 or ::) indicates the absence of an address. It is equivalent to the IPv4 unspecified address of 0.0.0.0. The unspecified address is typically used as a source address for packets attempting to verify the uniqueness of a tentative address. The unspecified address is never assigned to an interface or used as a destination address.

- Loopback address

  The loopback address (0:0:0:0:0:0:0:1 or ::1) is used to identify a loopback interface, enabling a node to send packets to itself. It is equivalent to the IPv4 loopback address of 127.0.0.1. Packets addressed to the loopback address must never be sent on a link or forwarded by a router.

**Compatibility addresses**
To aid in the migration from IPv4 to IPv6 and the coexistence of both types of hosts, the following addresses are defined:

- IPv4-compatible address

  The IPv4-compatible address, 0:0:0:0:0:0:*w.x.y.z* or ::*w.x.y.z* (where *w.x.y.z* is the dotted decimal representation of a public IPv4 address), is used by IPv6/IPv4 nodes that are communicating using IPv6. IPv6/IPv4 nodes support both IPv4 and IPv6 protocols. When an IPv4-compatible address is used as an IPv6 destination, the IPv6 traffic is automatically encapsulated with an IPv4 header and sent to the destination using the IPv4 infrastructure.

- IPv4-mapped address

  The IPv4-mapped address, 0:0:0:0:0:FFFF:*w.x.y.z* or ::FFFF:*w.x.y.z*, is used to represent an IPv4-only node to an IPv6 node. It is used only for internal representation. The IPv4-mapped address is never used as a source or destination address of an IPv6 packet.

- 6to4 address

  The 6to4 address is used for communicating between two nodes running both IPv4 and IPv6 over an IPv4 routing infrastructure. The 6to4 address is formed by combining the prefix 2002::/16 with the 32 bits of a public IPv4 address of the node, forming a 48-bit prefix. 6to4 is a tunneling technique described in RFC 3056.

- ISATAP address

  The Internet draft titled "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)" defines ISATAP addresses used between two nodes running both IPv4 and IPv6 over an IPv4 routing infrastructure. ISATAP addresses use the locally administered interface ID ::0:5EFE:*w.x.y.z* where *w.x.y.z* is any unicast IPv4 address, which includes both public and private addresses.

  The ISATAP interface ID can be combined with any 64-bit prefix that is valid for IPv6 unicast addresses. This includes the link-local address prefix (FE80::/64), site-local prefixes, and global prefixes.

- Teredo address

  Teredo addresses use the prefix 3FFE:831F::/32. Beyond the first 32 bits, Teredo addresses are used to encode the IPv4 address of a Teredo server, flags, and the encoded version of the external address and port of a Teredo client. An example of a Teredo address is 3FFE:831F:CE49:7601:8000:EFFF:62C3:FFFE. Teredo addresses are used to represent a host when using the automatic tunneling mechanism defined in the Internet draft titled "Teredo: Tunneling IPv6 over UDP through NATs." For more information, see "Teredo" later in this section.

**Multicast IPv6 Addresses**
In IPv6, multicast traffic operates in the same way that it does in IPv4. Arbitrarily located IPv6 nodes can listen for multicast traffic on arbitrary IPv6 multicast addresses. IPv6 nodes can listen to multiple multicast addresses at the same time. Nodes can join or leave a multicast group at any time.

IPv6 multicast addresses have the first 8 bits set to 1111 1111. An IPv6 address is easy to classify as multicast because it always begins with "FF." Multicast addresses cannot be used as source addresses or as intermediate destinations in a Routing header.

Beyond the first 8 bits, multicast addresses include additional structure to identify their flags, scope, and multicast group. The following figure shows the structure of the IPv6 multicast address.

**The IPv6 Multicast Address**




The fields in the multicast address are:

- Flags

Indicates flags set on the multicast address. The size of this field is 4 bits. As of RFC 3513, the only flag defined is the Transient (T) flag. The T flag uses the low-order bit of the Flags field. When set to 0, the T flag indicates that the multicast address is a well-known multicast address that has been permanently assigned by the Internet Assigned Numbers Authority (IANA). When set to 1, the T flag indicates that the multicast address is a transient multicast address that IANA has not permanently assigned.

- Scope

  Indicates the scope of the IPv6 internetwork for which the multicast traffic is intended. The size of this field is 4 bits. In addition to using information provided by multicast routing protocols, routers use the multicast scope to determine whether they can forward multicast traffic. The most prevalent values for the Scope field are 1 (interface-local scope), 2 (link-local scope), and 5 (site-local scope).

  For example, traffic with the multicast address of FF02::2 has a link-local scope. An IPv6 router never forwards this traffic beyond the local link.

- Group ID

  Identifies the multicast group and is unique within the scope. The size of this field is 112 bits. Permanently assigned group IDs are independent of the scope. Transient group IDs are relevant only to a specific scope. Multicast addresses from FF01:: through FF0F:: are reserved, well-known addresses.

To identify all nodes for the interface-local and link-local scopes, the following addresses are defined:

- FF01::1 (interface-local scope all-nodes multicast address)
- FF02::1 (link-local scope all-nodes multicast address)

To identify all routers for the interface-local, link-local, and site-local scopes, the following addresses are defined:

- FF01::2 (interface-local scope all-routers multicast address)
- FF02::2 (link-local scope all-routers multicast address)
- FF05::2 (site-local scope all-routers multicast address)

With 112 bits for the group ID, it is possible to have $2^{112}$ group IDs. However, because of the way in which IPv6 multicast addresses are mapped to Ethernet multicast media access control (MAC) addresses, RFC 3513 recommends assigning the group ID from the low-order 32 bits of the IPv6 multicast address and setting the remaining original group ID bits to 0. By using only the low-order 32 bits, each group ID maps to a unique Ethernet multicast MAC address. The following figure shows the recommended IPv6 multicast address.

**The Recommended IPv6 Multicast Address Using a 32-bit Group ID**

**Solicited-node Address**

The solicited-node address facilitates the efficient querying of network nodes during address resolution. In IPv4, the ARP Request frame is sent to the MAC-level broadcast, disturbing all nodes on the network segment, including those that are not running IPv4. IPv6 uses the Neighbor Solicitation message to perform address resolution. However, instead of using the local-link scope all-nodes multicast address as the Neighbor Solicitation message destination, which would disturb all IPv6 nodes on the local link, the solicited-node multicast address is used. The following figure shows how the solicited-node multicast address comprises the prefix FF02::1:FF00:0/104 and the last 24-bits of the IPv6 address that is being resolved.

**Mapping Unicast IPv6 Addresses to Solicited Node IPv6 Addresses**




For example, for the node with the link-local IPv6 address of FE80::2AA:FF:FE28:9C5A, the corresponding solicited-node address is FF02::1:FF28:9C5A. This node is listening for multicast traffic at the solicited-node address of FF02::1:FF28:9C5A and, for interfaces that correspond to a physical network adapter, has registered the corresponding multicast address with the network adapter. To resolve the address of FE80::2AA:FF:FE28:9C5A to its link-layer address, a neighboring node sends a Neighbor Solicitation to the solicited-node address of FF02::1:FF28:9C5A.

The result of using the solicited-node multicast address is that address resolutions, a common occurrence on a link, are not required to use a mechanism that disturbs all network nodes. By using the solicited-node address, address resolution disturbs very few nodes. In practice, due to the relationship between the Ethernet MAC address, the IPv6 interface ID, and the solicited-node address, the solicited-node address acts as a pseudo-unicast address for very efficient address resolution.

**Anycast IPv6 Addresses**

An anycast address is assigned to multiple interfaces. The routing infrastructure forwards packets that are addressed to an anycast address to the nearest interface to which the anycast address is assigned. To facilitate delivery, the routing infrastructure must track the interfaces that have been assigned anycast addresses and their distance in terms of routing metrics. At present, anycast addresses are used only as destination addresses and are assigned only to routers. Anycast addresses are assigned out of the unicast address space, and the scope of an anycast address is the scope of the type of unicast address from which the anycast address is assigned.

The Subnet-Router anycast address is predefined and required. It is created from the subnet prefix for a given interface. To construct the Subnet-Router anycast address, the bits in the subnet prefix are fixed at their appropriate values, and the remaining bits are set to 0. All router interfaces attached to a subnet are assigned the Subnet-Router anycast address for that subnet. The Subnet-Router anycast address is used to communicate with one of multiple routers attached to a remote subnet.

**IPv6 Addresses for a Host**

An IPv4 host with a single network adapter typically has a single IPv4 address assigned to that adapter. An IPv6 host, however, usually has multiple IPv6 addresses — even with a single interface. An IPv6 host is assigned the following unicast addresses:

- A link-local address for each interface.
- Unicast addresses for each interface (which could be a site-local address and one or multiple global unicast addresses).
- The loopback address (::1) for the loopback interface.

Typical IPv6 hosts are logically multi-homed because they have at least two addresses with which they can receive packets — a link-local address for local link traffic and a routable site-local or global address.

Additionally, each host is listening for traffic on the following multicast addresses:

- The interface-local scope all-nodes multicast address (FF01::1).
- The link-local scope all-nodes multicast address (FF02::1).
- The solicited-node address for each unicast address on each interface.
- The multicast addresses of joined groups on each interface.

**IPv6 Addresses for a Router**

An IPv6 router is assigned the following unicast addresses:

- A link-local address for each interface.
- Unicast addresses for each interface (which could be a site-local address and one or multiple global unicast addresses).

- A Subnet-Router anycast address.
- Additional anycast addresses (optional).
- The loopback address (::1) for the loopback interface.

Additionally, each router is listening for traffic on the following multicast addresses:

- The interface-local scope all-nodes multicast address (FF01::1).
- The interface-local scope all-routers multicast address (FF01::2).
- The link-local scope all-nodes multicast address (FF02::1).
- The link-local scope all-routers multicast address (FF02::2).
- The site-local scope all-routers multicast address (FF05::2).
- The solicited-node address for each unicast address on each interface.
- The multicast addresses of joined groups on each interface.

**IPv6 Interface Identifiers**
The last 64 bits of an IPv6 address are the interface identifier that is unique to the 64-bit prefix of the IPv6 address. IPv6 interface identifiers are determined as follows:

- Derived from the Extended Unique Identifier (EUI)-64 address.
- Generated randomly and changed over time to provide a level of anonymity.
- Assigned during stateful address autoconfiguration (for example, through DHCPv6).

**EUI-64 Address-based Interface Identifiers**
RFC 3513 states that all unicast addresses that use the prefixes 001 through 111 must also use a 64-bit interface identifier derived from the EUI-64 address. EUI-64 addresses have 64-bits and are defined by the Institute of Electrical and Electronic Engineers (IEEE). EUI-64 addresses are either assigned to network adapters or derived from IEEE 802 addresses.

**IEEE 802 addresses**
Traditional interface identifiers for network adapters use a 48-bit address called an IEEE 802 address. It consists of a 24-bit company ID (also called the manufacturer ID) and a 24-bit extension ID (also called the board ID). The combination of the company ID, which is uniquely assigned to each manufacturer of network adapters, and the board ID, which is uniquely assigned to each network adapter at the time of assembly, produces a globally unique 48-bit address. This 48-bit address is also called the physical, hardware, or media access control (MAC) address.

The following figure shows the structure of the 48-bit IEEE 802 address.

### 48-bit IEEE 802 Address




The IEEE 802 address includes the following defined bits:

- Universal/Local (U/L)

  The next-to-the-low order bit in the first byte indicates whether the address is universally or locally administered. If the U/L bit is set to 0, the IEEE (through the designation of a unique company ID) has administered the address. If the U/L bit

is set to 1, the address is locally administered, which means that the network administrator has overridden the manufactured address and specified a different address.

- Individual/Group (I/G)

  The low-order bit of the first byte indicates whether the address is an individual address (unicast) or a group address (multicast). If the I/G bit is set to 0, the address is a unicast address. If the I/G bit is set to 1, the address is a multicast address.

For a typical IEEE 802 network adapter address, both the U/L and I/G bits are set to 0, indicating a universally administered, unicast MAC address.

**IEEE EUI-64 addresses**
The IEEE Extended Unique Identifier (EUI)-64 address represents a newer standard for network interface addressing. The company ID is still 24 bits, but the extension ID is 40 bits, creating a much larger address space for each network adapter manufacturer. The EUI-64 address uses the U/L and I/G bits in the same way as the IEEE 802 address.

The following figure shows the structure of the EUI-64 address.

**EUI-64 Address**




**Mapping IEEE 802 addresses to EUI-64 addresses**
To create an EUI-64 address from an IEEE 802 address, the 16 bits of 11111111 11111110 (0xFFFE) are inserted into the IEEE 802 address between the company ID and the extension ID. The following figure shows an example of such a conversion.

**Conversion of an IEEE 802 Address to an EUI-64 Address**




**Mapping EUI-64 addresses to IPv6 interface identifiers**
To obtain the 64-bit interface identifier for an IPv6 unicast address, the U/L bit in the EUI-64 address is complemented. (If it is set to 1, it is changed to 0, and if it is set to 0, it is changed to 1.) The following figure shows the conversion for a universally administered, unicast EUI-64 address.

**Conversion of a Universally Administered, Unicast EUI-64 Address to an IPv6 Interface Identifier**

To obtain an IPv6 interface identifier from an IEEE 802 address, you must first map the IEEE 802 address to an EUI-64 address, and then complement the U/L bit. The following figure shows this conversion process for a universally administered, unicast IEEE 802 address.

**Conversion of a Universally Administered, Unicast IEEE 802 Address to an IPv6 Interface Identifier**




**IEEE 802 address conversion example**

Host A has the Ethernet MAC address of 00-AA-00-3F-2A-1C. First, it is converted to EUI-64 format by inserting FF-FE between the third and fourth bytes, yielding 00-AA-00-FF-FE-3F-2A-1C. The U/L bit, which is the seventh bit in the first byte, is then complemented. Before the conversion, the first byte in binary form is 00000000. When the seventh bit is complemented, it becomes 00000010 (0x02). The final result is 02-AA-00-FF-FE-3F-2A-1C which, when converted to colon-hexadecimal notation, becomes the interface identifier 2AA:FF:FE3F:2A1C. As a result, the link-local address that corresponds to the network adapter with the MAC address of 00-AA-00-2A-1C is FE80::2AA:FF:FE3F:2A1C.

### Note

- When complementing the U/L bit, add 0x2 to the first byte if the address is universally administered, and subtract 0x2 from the first byte if the address is locally administered.

**Temporary Address Interface Identifiers**

In today's Internet, a typical Internet user connects to an Internet service provider (ISP) and obtains an IPv4 address using the Point-to-Point Protocol (PPP) and the Internet Protocol Control Protocol (IPCP). Each time the user connects, a different IPv4 address might be obtained. Because of this, it is difficult to track a dial-up user's traffic on the Internet on the basis of IP address.

For dial-up connections that are based on IPv6, the user is assigned a 64-bit prefix through router discovery and stateless address autoconfiguration after the connection is made. If the interface identifier is always based on the EUI-64 address (as derived from the static IEEE 802 address), it is possible to identify the traffic of a specific node regardless of the prefix, making it easy to track a specific user and their use of the Internet. To address this concern and provide a level of anonymity, RFC 3041 describes an alternative IPv6 interface identifier that is randomly generated and changes over time.

The initial interface identifier is generated by using random numbers. For computers that cannot store any historical information for generating future interface identifier values, a different random interface identifier is generated each time the IPv6 protocol is initialized. For IPv6 systems that have storage capabilities, a history value is stored. When the IPv6 protocol is initialized, a new interface identifier is generated through the following process:

1. Retrieve the history value from storage, and append the interface identifier based on the EUI-64 address of the adapter.

2. Compute a Message Digest-5 (MD5) one-way encryption hash over the quantity in step 1.

3. Save the last 64 bits of the MD5 hash computed in step 2 as the history value for the next interface identifier computation.

4. Take the first 64 bits of the MD5 hash computed in Step 2, and set the seventh bit to 0. The seventh bit corresponds to the U/L bit which, when set to 0, indicates a locally administered IPv6 interface identifier. The result is the IPv6 interface identifier.

The resulting IPv6 address, based on this random interface identifier, is known as a temporary or anonymous address. Temporary addresses are generated for public address prefixes that use stateless address autoconfiguration. Temporary addresses are used for the shorter of the valid and preferred lifetimes:

- The lifetimes included in the Prefix Information option in the received Router Advertisement message.

- Local default values of one week for valid lifetime and one day for preferred lifetime.

After the valid lifetime of the temporary address expires, a new interface identifier and temporary address is generated.

**Subnetting Global Address Prefixes**
For most users of IPv6, subnetting the IPv6 address space consists of using subnetting techniques to divide the Subnet ID portion of the global address. This approach allows for route summarization and delegation of the remaining address space to different portions of an IPv6 intranet.

For the global address, the first 48 bits are fixed and allocated by an ISP.

Subnetting the subnet ID portion of a global address space requires two steps:

1. Determine the number of bits to be used for the subnetting.

2. Enumerate the new subnetted network prefixes.

You are not required to subnet in any specific fashion. The subnetting technique described here assumes that subnetting is done by dividing the 16-bit address space of the subnet ID using the high-order bits in the subnet ID. Although this method promotes hierarchical addressing and routing, it is not required. For example, in a small organization with a small number of subnets, you can also create a flat addressing space for the subnet ID by numbering the subnets from 0 to 65,535.

**Step 1: Determining the Number of Subnetting Bits**
The number of bits being used for subnetting determines the possible number of new subnetted network prefixes that can be allocated to portions of your network based on geographical or departmental divisions. In a hierarchical routing infrastructure, you need to determine how many network prefixes, and therefore how many bits, are needed at each level in the hierarchy. The more bits you choose for the various levels of the hierarchy, the fewer bits you have to enumerate individual subnets in the last level of the hierarchy.

For example, a network administrator can decide to implement a two-level hierarchy reflecting a geographical/departmental structure, use 4 bits for the geographical level, and use 6 bits for the departmental level. With this approach, each department in each geographical location has only 6 bits of subnetting space left (16 - 6 - 4) or only 64 (=$2^6$) subnets per department.

On any given level in the hierarchy, some number of bits is already fixed by the next level up in the hierarchy (*f*), some number of bits is used for subnetting at the current level in the hierarchy (*s*), and some number of bits remain for the next level down in the hierarchy (*r*). At all times, $f+s+r = 16$. The following figure shows this relationship.

**Subnetting a Subnet ID**




**Step 2: Enumerating Subnetted Network Prefixes**
Based on the number of bits used for subnetting, you must list the new subnetted network prefixes. There are two main approaches:

- Hexadecimal — Enumerate new subnetted network prefixes by using hexadecimal representations of the subnet ID and increment.

- Decimal — Enumerate new subnetted network prefixes by using decimal representations of the subnet ID and increment.

Either method produces the same result: an enumerated list of subnetted network prefixes.

**To create the enumerated list of subnetted network prefixes using the hexadecimal method**

1. Based on *s*, the number of bits chosen for subnetting, and *m*, the length of the network prefix being subnetted, calculate the following:

$f = m - 48$

$f$ is the number of bits within the global address prefix that are already fixed.

$n = 2^s$

$n$ is the number of network prefixes that are obtained.

$i = 2^{16-(f+s)}$

$i$ is the incremental value between each successive subnet ID expressed in hexadecimal.

$l = 48+f+s$

$l$ is the length of the new subnetted network prefixes.

2. Create a two-column table with *n* entries. The first column is the network prefix number (starting with 1), and the second column is the new subnetted network prefix.

3. In the first table entry, the subnetted network prefix is the original network prefix with the new prefix length. For example, based on *F*, the hexadecimal value of the subnet ID being subnetted, the subnetted network prefix is [*48-bit prefix*]:*F*::/*l*.

4. In the next table entry, increase the value within the subnet ID portion of the site-local or global address by *i*. For example, in the second table entry, the subnetted prefix is [*48-bit prefix*]:*F*+*i*::/*l*.

5. Repeat step 4 until the table is complete.

For example, to perform a 3-bit subnetting of the site-local network prefix 3FFE:FFFF:0:C000::/51, first calculate the values of the number of prefixes, the increment, and the new prefix length. The starting values are $F$=0xC000, $s$=3, and $f$=51-48=3. The number of prefixes is 8 ($n=2^3$). The increment is 0x400 ($i=2^{16-(3+3)}$=1024=0x400). The new prefix length is 54 ($l$=48+3+3).

Next, construct a table with 8 entries. The entry for the network prefix 1 is 3FFE:FFFF:0:C000::/54. Additional entries in the table are successive increments of *i* in the subnet ID portion of the network prefix, as shown in the following table.

**Hexadecimal Subnetting Technique for Network Prefix 3FFE:FFFF:0:C000::/51**

| Network Prefix | Subnetted Network Prefix |
|---|---|
| 1 | 3FFE:FFFF:0:C000::/54 |
| 2 | 3FFE:FFFF:0:C400::/54 |
| 3 | 3FFE:FFFF:0:C800::/54 |
| 4 | 3FFE:FFFF:0:CC00::/54 |
| 5 | 3FFE:FFFF:0:D000::/54 |
| 6 | 3FFE:FFFF:0:D400::/54 |
| 7 | 3FFE:FFFF:0:D800::/54 |
| 8 | 3FFE:FFFF:0:DC00::/54 |

**To create the enumerated list of subnetted network prefixes using the decimal method**

1. Based on *s*, the number of bits chosen for subnetting, *m*, the length of the network prefix being subnetted, and *F*, the hexadecimal value of the subnet ID being subnetted, calculate the following:

$f = m - 48$

*f* is the number of bits within the prefix that are already fixed.

$n = 2^s$

*n* is the number of network prefixes that are obtained.

$i = 2^{16-(f+s)}$

*i* is the incremental value between each successive subnet ID.

$l = 48+f+s$

*l* is the length of the new subnetted network prefixes.

*D* = decimal representation of *F*

2. Create a three-column table with *n* entries. The first column is the network prefix number (starting with 1), the second column is the decimal representation of the subnet ID portion of the new subnetted network prefix, and the third column is the new subnetted network prefix.

3. In the first column of the first row, the decimal representation of the subnet ID is *D* and the subnetted prefix is [*48-bit prefix*]:*F*::/*l*.

4. In the second column of the first row, increase the value of the decimal representation of the subnet ID by *i*. For example, in the second table entry, the decimal representation of the subnet ID is *D*+*i*.

5. In the third column of the first row, convert the decimal representation of the subnet ID to hexadecimal, and construct the prefix from [*48-bit prefix*]:[*subnet ID*]::/*l*. For example, in the second table entry, the subnetted network prefix is [*48-bit prefix*]:[*D*+*i (converted to hexadecimal)*]::/*l*.

6. Repeat steps 4 and 5 until the table is complete.

For example, to perform a 3-bit subnetting of the site-local network prefix 3FFE:FFFF:0:C000::/51, first calculate the values of the number of prefixes, the increment, the new prefix length, and the decimal representation of the starting subnet ID. The starting values are *F*=0xC000, *s*=3, and *f*=51-48=3. The number of prefixes is 8 ($n=2^3$). The increment is 1024 ($i=^{216-(3+3)}$). The new prefix length is 54 (*l*=48+3+3). The decimal representation of the starting subnet ID is 49152 (*D*=0xC000=49152).

Next, construct a table with 8 entries. The entry for the network prefix 1 is 49192 and 3FFE:FFFF:0:C000::/54. Additional entries in the table are successive increments of *i* in the subnet ID portion of the network prefix, as shown in the following table.

**Decimal Subnetting for Network Prefix 3FFE:FFFF:0:C000::/51**

| Network Prefix | Decimal Representation of Subnet ID | Subnetted Network Prefix |
|---|---|---|
| 1 | 49192 | 3FFE:FFFF:0:C000::/54 |
| 2 | 50176 | 3FFE:FFFF:0:C400::/54 |
| 3 | 51200 | 3FFE:FFFF:0:C800::/54 |
| 4 | 52224 | 3FFE:FFFF:0:CC00::/54 |
| 5 | 53248 | 3FFE:FFFF:0:D000::/54 |
| 6 | 54272 | 3FFE:FFFF:0:D400::/54 |
| 7 | 55296 | 3FFE:FFFF:0:D800::/54 |

| | | |
|---|---|---|
| 8 | 56320 | 3FFE:FFFF:0:DC00::/54 |

**Note**

- RFC 3513 allows the use of subnetted network prefixes where the bits being used for subnetting are set to all 0s (the all-zeros subnetted network prefix) and all 1s (the all-ones subnetted network prefix).

#### IPv6 and Name Resolution

Resolving names to IPv6 addresses is supported if either the local Hosts file (stored in *systemroot*\System32\Drivers\Etc) contains entries for IPv6 addresses or the Domain Name System (DNS) infrastructure contains resource records for IPv6 addresses. RFC 1886 describes enhancements to DNS for IPv6, and the DNS name resolver and DNS Server service in Windows Server 2003 support them.

**Note**

- NetBIOS over IPv6 is not supported. Therefore, NetBIOS name resolution techniques cannot be used to resolve names to IPv6 addresses.

In RFC 1886, a new DNS resource record type, AAAA (called quad A), is used for resolving a fully qualified domain name to an IPv6 address. It is comparable to the host address (A) resource record used with IPv4. The resource record type is named AAAA (Type value of 28) because IPv6 addresses are four times as large as IPv4 addresses. The following is an example of a AAAA resource record:

```
host1.microsoft.com    IN    AAAA   FEC0::2AA:FF:FE3F:2A1C
```

A host must specify either a AAAA query or a general query for a specific host name to receive IPv6 address resolution data in the DNS query answer sections.

RFC 1886 also describes the IP6.INT domain created for IPv6 reverse queries. However, according to RFC 3152, Internet Engineering Task Force (IETF) consensus has been reached that the IP6.ARPA domain be used instead of IP6.INT. Also called pointer queries, reverse queries determine a host name based on the IP address. To create the namespace for reverse queries, each hexadecimal digit in the fully expressed 32-digit IPv6 address becomes a separate level in inverse order in the reverse domain hierarchy.

For example, the reverse lookup domain name for the address FEC0::2AA:FF:FE3F:2A1C (fully expressed as FEC0:0000:0000:0000:02AA: 00FF:FE3F:2A1C) is:

C.1.A.2.F.3.E.F.F.F.0.0.A.A.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.C.E.F.IP6.ARPA.

The DNS support that RFC 1886 describes represents a simple way to both map host names to IPv6 addresses and to provide reverse name resolution.

#### IPv4 Addresses and IPv6 Equivalents

The following table lists IPv4 addresses and concepts next to their IPv6 equivalents.

**IPv4 Addresses and their IPv6 Equivalents**

| IPv4 Address | IPv6 Address |
|---|---|
| Internet address classes | Not applicable in IPv6 |
| Multicast addresses (224.0.0.0/4) | IPv6 multicast addresses (FF00::/8) |
| Broadcast addresses | Not applicable in IPv6 |
| Unspecified address is 0.0.0.0 | Unspecified address is :: |
| Loopback address is 127.0.0.1 | Loopback address is ::1 |
| Public IP addresses | Global unicast addresses |
| Private IP addresses (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16) | Site-local addresses (FEC0::/10) |

| | |
|---|---|
| Autoconfigured addresses (169.254.0.0/16) | Link-local addresses (FE80::/64) |
| Text representation: Dotted decimal notation | Text representation: Colon-hexadecimal format with suppression of leading zeros and zero compression. IPv4-compatible addresses are expressed in dotted decimal notation. |
| Network bits representation: Subnet mask in dotted decimal notation or prefix length notation | Network bits representation: Prefix length notation only |
| DNS name resolution: IPv4 host address (A) resource record | DNS name resolution: IPv6 host address AAAA resource records (RFC 1886) or A6 records (RFC 2874) |
| DNS reverse resolution: IN-ADDR.ARPA domain | DNS reverse resolution: IP6.INT domain (RFC 1886) or IP6.ARPA domain (RFC 2874) |

**Neighbor Discovery**
Neighbor Discovery uses ICMPv6 messages to manage neighboring node interaction. Neighbor Discovery replaces ARP, ICMP Router Discovery, and ICMP Redirect and provides additional functionality.

**Neighbor Discovery Messages**
All functions of Neighbor Discovery are performed with the following messages:

- Router Solicitation

  IPv6 hosts send Router Solicitation messages to discover IPv6 routers present on the link. To prompt IPv6 routers to respond immediately, hosts send multicast Router Solicitation messages rather than waiting for a periodic Router Advertisement message.

- Router Advertisement

  IPv6 routers send Router Advertisement messages either periodically or in response to the receipt of Router Solicitation messages. Router Advertisement messages contain the information required by hosts to determine what the link prefixes are, what the link MTU is, whether or not to use address autoconfiguration, and the duration for which addresses created through address autoconfiguration are both valid and preferred.

- Neighbor Solicitation

  IPv6 hosts send Neighbor Solicitation messages to discover the link-layer addresses of on-link IPv6 nodes. Neighbor Solicitation messages include the link-layer address of the sender. Hosts send Neighbor Solicitation messages to multicast addresses to resolve addresses and to unicast addresses to verify the reachability of a neighboring node.

- Neighbor Advertisement

  IPv6 nodes send Neighbor Advertisement messages in response to the receipt of Neighbor Solicitation messages. Nodes also send unsolicited Neighbor Advertisements to inform neighboring nodes of changes in link-layer addresses. Neighbor Advertisement messages contain information that nodes require to determine the type of Neighbor Advertisement message, the link-layer address of the sender, and the sender's role on the network.

- Redirect

  IPv6 routers send Redirect messages to inform an originating host of a better first-hop address for a specific destination. Routers send Redirect messages for unicast traffic only and only to originating hosts. Only hosts process Redirect messages.

To ensure that all Neighbor Discovery messages were sent by a node on the local link, all Neighbor Discovery messages are sent with a hop limit of 255. When a Neighbor Discovery message is received, the Hop Limit field in the IPv6 header is checked. If the field is not set to 255, the message is silently discarded. Verifying that the Neighbor Discovery message has a hop limit of 255 provides protection from network attacks that are based on Neighbor Discovery and launched from off-link nodes. If a

message has a hop limit of 255, a router could not have forwarded the message from an off-link node.

**Neighbor Discovery Options**
RFC 2461 defines the following Neighbor Discovery options:

- Source Link-Layer Address Option

  The Source Link-Layer Address option indicates the link-layer address of the Neighbor Discovery message sender. Neighbor Solicitation, Router Solicitation, and Router Advertisement messages include this option. This option is not included when the source address of the Neighbor Discovery message is the unspecified address (::).

- Target Link-Layer Address Option

  The Target Link-Layer Address option indicates the link-layer address of the neighboring node to which IPv6 packets should be directed. Neighbor Advertisement and Redirect messages include this option.

- Prefix Information Option

  The Prefix Information option is sent in Router Advertisement messages to indicate both address prefixes and information about address autoconfiguration. Each Router Advertisement message can include multiple Prefix Information options, indicating multiple address prefixes.

- Redirected Header Option

  The Redirected Header option is sent in Redirect messages to specify the IPv6 packet that caused the router to send a Redirect message. This option can contain all or part of the redirected IPv6 packet, depending on the size of the IPv6 packet that was initially sent.

- MTU Option

  The MTU option is sent in Router Advertisement messages to indicate the IPv6 MTU of the link. This option can be used when the IPv6 MTU for a link is not well known or must reflect a translational or mixed-media bridging configuration. The MTU option overrides the IPv6 MTU that the interface hardware reports.

  In bridged or Layer-2 switched environments, different link-layer technologies with different link-layer MTUs can exist on the same network segment. In this case, differences in IPv6 MTUs between nodes on the same network are not discovered through Path MTU Discovery. The MTU option indicates the highest IPv6 MTU supported by all link-layer technologies on the network segment.

**Host Data Structures**
To facilitate interactions between neighboring nodes, RFC 2461 defines the following host data structures as examples of how to store information for Neighbor Discovery processes:

- Neighbor cache

  Stores the on-link address of a neighbor, its corresponding link-layer address, and an indication of the neighbor's reachability state. The neighbor cache is equivalent to the ARP cache in IPv4.

- Destination cache

  Stores information about forwarding or next-hop addresses for destinations to which traffic has recently been sent. Entries in the destination cache contain the destination IP address (either local or remote), the previously resolved next-hop address, and the Path MTU for the destination.

- Prefix list

  Lists on-link prefixes. Each entry defines a range of IP addresses for destinations that are directly reachable (neighbors). This list is populated from prefixes that routers advertise in Router Advertisement messages.

- Default router list

  Lists IP addresses corresponding to on-link routers that send Router Advertisement messages and that are eligible to be default routers.

RFC 2461 defines these structures as an example of an IPv6 host conceptual model. An IPv6 implementation is not required to create these exact data structures as long as the external behavior of the host is consistent with RFC 2461. For example, the IPv6 protocol for Windows Server 2003 uses a routing table rather than a prefix list and default router list. The routing table contains sufficient information to determine how to forward IPv6 packets. For more information, see "IPv6 Routing" later in this section.

**Neighbor Discovery Message Exchanges**
The Neighbor Discovery protocol provides message exchanges for the following processes:

- Address resolution
- Duplicate address detection
- Router discovery
- Redirect function

**Address Resolution**
To resolve addresses, IPv6 nodes exchange Neighbor Solicitation and Neighbor Advertisement messages to resolve the link-layer address of the on-link next-hop address for a given destination. The sending host sends a multicast Neighbor Solicitation message on the appropriate interface. The multicast address of the Neighbor Solicitation message is the solicited-node multicast address derived from the target IP address. The Neighbor Solicitation message includes the link-layer address of the sending host in the Source Link-Layer Address option.

When the target host receives the Neighbor Solicitation message, it updates its own neighbor cache based on the source address of the Neighbor Solicitation message and the link-layer address in the Source Link-Layer Address option. Next, the target host sends a unicast Neighbor Advertisement to the Neighbor Solicitation sender. The Neighbor Advertisement includes the Target Link-Layer Address option.

After receiving the Neighbor Advertisement from the target host, the sending host updates its neighbor cache with an entry for the target host based on the information in the Target Link-Layer Address option. At this point, the sending host and the target host can send unicast IPv6 traffic to each other.

As an example of this process, Host A has an Ethernet MAC address of 00-AA-00-11-11-11 and a corresponding link-local address of FE80::2AA:FF:FE11:1111. Host B has an Ethernet MAC address of 00-AA-00-22-22-22 and a corresponding link-local address of FE80::2AA:FF:FE22:2222. To send a packet to Host B, Host A must resolve Host B's link-layer address.

The following figure shows Host A sending a solicited-node multicast Neighbor Solicitation message (based on Host B's IP address) to the address of FF02::1:FF22:2222.

**Multicast Neighbor Solicitation Message for Address Resolution**




Host B, having registered the solicited-node multicast address of 33-33-FF-22-22-22 with its Ethernet adapter, receives and

processes the Neighbor Solicitation. The following figure shows Host B responding with a unicast Neighbor Advertisement message.

**Unicast Neighbor Advertisement Message for Address Resolution**




**Duplicate Address Detection**

IPv4 nodes use ARP Request messages and a method called gratuitous ARP to detect duplicate IP addresses on their local links. Similarly, IPv6 nodes use the Neighbor Solicitation message to detect duplicate address use on their local links.

With gratuitous ARP, the Source Protocol Address and Target Protocol Address fields in the ARP Request message header are set to the IPv4 address for which duplication is being detected. In IPv6 duplicate address detection, the Target Address field in the Neighbor Solicitation message is set to the IPv6 address for which duplication is being detected.

Duplicate address detection differs from address resolution in these ways:

- In a Neighbor Solicitation message for duplicate address detection, the Source Address field in the IPv6 header is set to the unspecified address (::). The address being queried for duplication cannot be used until it is determined that the address is not already in use.

- In the Neighbor Advertisement reply to a Neighbor Solicitation message for duplicate address detection, the Destination Address in the IPv6 header is set to the link-local scope all-nodes multicast address (FF02::1). Because the sender of the Neighbor Solicitation message for duplicate address detection is not using the desired address, it cannot receive unicast Neighbor Advertisement messages. Therefore, the Neighbor Advertisement message is multicast.

Upon receipt of the multicast Neighbor Advertisement message sent in response to the Neighbor Solicitation message, the node disables the use of the duplicate address on the interface. If the node does not receive a Neighbor Advertisement message that defends the use of the IPv6 address, it initializes the address on the interface.

**Router Discovery**

Router discovery is the process through which nodes attempt to discover the set of routers on the local link. Router discovery in IPv6 is similar to ICMP Router Discovery for IPv4 described in RFC 1256.

An important difference between ICMPv4 Router Discovery and IPv6 router discovery is the mechanism through which the default router is changed when the current one becomes unavailable. In ICMPv4 Router Discovery, the Router Advertisement message includes an Advertisement Lifetime field. Advertisement Lifetime is the time after which the router, upon sending its last Router Advertisement message, can be considered unavailable. In the worst case, a router can become unavailable and hosts will not attempt to discover a new default router until the Advertisement Lifetime has elapsed.

IPv6 has a Router Lifetime field in the Router Advertisement message. This field indicates the length of time that the router can be considered a default router. However, if the current default router becomes unavailable, the condition is detected through neighbor unreachability detection instead of the Router Lifetime field in the Router Advertisement message. Because neighbor unreachability detection determines that the router is no longer reachable, a new router is chosen immediately from the default router list. For more information, see "Neighbor Unreachability Detection" later in this section.

In addition to configuring a default router, IPv6 router discovery also configures the following:

- The default setting for the Hop Limit field in the IPv6 header.

- A determination of whether the node should use a stateful address protocol, such as Dynamic Host Configuration Protocol for IPv6 (DHCPv6), for addresses and other configuration parameters.

- The timers used in reachability detection and the retransmission of Neighbor Solicitation messages.

- The list of network prefixes defined for the link. Each network prefix contains both the IPv6 network prefix and its valid and preferred lifetimes. If indicated, a network prefix combined with the interface identifier creates a stateless IP address configuration for the receiving interface. A network prefix also defines the range of addresses for nodes on the local link.

- The MTU of the local link.

The IPv6 router discovery processes are the following:

- IPv6 routers periodically send Router Advertisement messages on the local link advertising their existence as routers. They also provide configuration parameters such as default hop limit, MTU, and prefixes.

- Active IPv6 hosts on the local link receive the Router Advertisement messages and use the contents to maintain their default router lists, prefix lists, and other configuration parameters.

- A host that is starting up sends a Router Solicitation message to the link-local scope all-routers multicast address (FF02::2). Upon receipt of a Router Solicitation message, each router on the local link send a unicast Router Advertisement message to the node that sent the Router Solicitation message. The node receives the Router Advertisement messages and uses their contents to build the default router and prefix lists and to set other configuration parameters.

**Redirect Function**

Routers use the redirect function to inform originating hosts of a better first-hop neighbor to which traffic should be forwarded for a specific destination. Routers use the redirect function for two purposes:

- A router informs an originating host of the IP address of a router available on the local link that is closer to the destination. The term closer is a routing metric function used to reach the destination network segment. This condition can occur when multiple routers are on a network segment, the originating host chooses a default router, and it is not the best one to use to reach the destination.

- A router informs an originating host that the destination is a neighbor (it is on the same link as the originating host). This condition can occur when the prefix list of a host does not include the prefix of the destination. Because the destination does not match a prefix in the list, the originating host forwards the packet to its default router.

The following steps occur in the IPv6 redirect process:

1. The originating host sends a unicast packet to its default router.

2. The router processes the packet and notes that the address of the originating host is a neighbor. Additionally, the router notes that both the originating host and the next-hop are on the same link.

3. The router forwards the packet to the appropriate next-hop address.

4. The router sends the originating host a Redirect message. In the Target Address field of the Redirect message is the next-hop address of the node to which the originating host should send packets addressed to the destination.

   For packets redirected to a router, the Target Address field is set to the link-local address of the router. For packets redirected to a host, the Target Address field is set to the destination address of the packet originally sent.

   The Redirect message includes the Redirected Header option. The message might also include the Target Link-Layer Address option.

5. Upon receiving the Redirect message, the originating host updates the destination address entry in the destination cache with the address in the Target Address field. If the Redirect message includes the Target Link-Layer Address option, its contents are used to create or update the corresponding entry in the neighbor cache.

Only the first router in the path between the originating host and the destination sends redirect messages, and (like ICMPv6 error messages) they are rate limited. Hosts never send Redirect messages, and routers never update routing tables based on the receipt of a Redirect message.

**Neighbor Unreachability Detection**

Reachability is defined as the ability to send an IPv6 packet to a neighboring node and have the IPv6 layer of the neighbor receive and process that packet. When a node sends a packet to a router, the packet is delivered to the router's IPv6 layer and then forwarded to the next hop. When a node sends a packet to a neighboring node, the packet is delivered to the node's IPv6 layer. It is important to note that the definition of reachability does not require delivery to a remote node across a router — only to the neighboring router.

When a neighbor becomes unreachable, IPv6 detects this condition and attempts to correct it. To determine whether a neighbor is reachable, IPv6 relies on either upper layer protocols that indicate communication progress or receipt of a Neighbor Advertisement message that has been sent in response to a unicast Neighbor Solicitation message.

For TCP traffic, communication progress is indicated when new data or acknowledgement segments for sent data are received. For UDP traffic, a progress indication might not be present. In this case, the node sends unicast Neighbor Solicitation messages to the next-hop neighbor to monitor its ongoing reachability.

Only the receipt of a solicited Neighbor Advertisement message is considered proof of reachability. A solicited Neighbor Advertisement message, which has its Solicited flag set to 1, is sent only in response to a Neighbor Solicitation message. Unsolicited Neighbor Advertisement or Router Advertisement messages are not considered proof of reachability.

Neighbor unreachability detection detects symmetric reachability. In this instance, packets must be able to travel to and from the desired neighboring node. When a Neighbor Solicitation message is sent and a solicited Neighbor Advertisement message is received, the path between the nodes is confirmed. For an unsolicited Neighbor Advertisement or Router Advertisement message, only the path from the node sending the message is confirmed. This is called asymmetric reachability.

For a specific local node, reachability is confirmed only by the node that sends the Neighbor Solicitation message and receives the Neighbor Advertisement message. The node sending the Neighbor Advertisement message receives no confirmation that the message reached the intended node. For two neighboring nodes to both determine reachability, each node must exchange Neighbor Solicitation and Neighbor Advertisement messages with the other.

The reachability of a neighboring node is determined by monitoring the state of the neighboring node's entry in the neighbor cache. RFC 2461 defines the following states for a neighbor cache entry:

- INCOMPLETE

  IPv6 address resolution, which is using a solicited-node multicast Neighbor Solicitation message, is in progress. The INCOMPLETE state is entered when a neighbor cache entry is created but does not yet have the node's corresponding link-layer address. RFC 2461 recommends that three successive multicast Neighbor Solicitation messages be sent.

- REACHABLE

  Reachability has been confirmed by receipt of a solicited unicast Neighbor Advertisement message. The neighbor cache entry stays in the REACHABLE state until the number of milliseconds indicated in the Reachable Time field in the Router Advertisement message elapses. The entry stays in the REACHABLE state as long as upper layer protocols such as TCP indicate that communication is progressing. Each time that progress is indicated, the reachable time for the entry is refreshed.

- STALE

  Reachable time (the duration since the last reachability confirmation was received) has elapsed. The neighbor cache entry enters the STALE state after the number of milliseconds in the Reachable Time field in the Router Advertisement message (or a host default value) elapses, and the entry remains in this state until a packet is sent to the neighbor. The entry also enters the STALE state when the host receives an unsolicited Neighbor Advertisement message that is advertising the link-layer address.

- DELAY

  To allow time for upper layer protocols to provide reachability confirmation before sending Neighbor Solicitation

messages, the neighbor cache entry enters the DELAY state and waits a configurable period of time after sending a packet. RFC 2461 recommends a value of 5 seconds. If reachability is not confirmed by the delay time, then the entry enters the PROBE state, and a unicast Neighbor Solicitation message is sent.

- PROBE

  Reachability confirmation is in progress for a neighbor cache entry that was in the STALE or DELAY state. Unicast Neighbor Solicitation messages are sent at intervals corresponding to a retransmission timer field in the Router Advertisement message that this host received. A configurable variable determines the number of Neighbor Solicitation messages sent before the reachability detection process is abandoned and the neighbor cache entry is removed. RFC 2461 recommends sending three successive unicast Neighbor Solicitation messages.

The following figure shows the state diagram of an entry in the neighbor cache.

**States of a Neighbor Cache Entry**




If the unreachable neighbor is a router, the host chooses another router from the default router list and performs both address resolution and unreachability detection on it.

**Address Autoconfiguration**

A highly useful aspect of IPv6 is its ability to automatically configure itself without the use of a stateful configuration protocol, such as Dynamic Host Configuration Protocol for IPv6 (DHCPv6). By default, an IPv6 host can configure a link-local address for each interface. By using router discovery, a host can also determine the addresses of routers, additional addresses, and other configuration parameters. The Router Advertisement message indicates whether a stateful address configuration protocol should be used.

Address autoconfiguration can be performed only on multicast-capable interfaces. RFC 2462 describes address autoconfiguration.

**Autoconfigured Address States**

Autoconfigured addresses are in one or more of the following states:

- Tentative

  The address is in the process of being verified as unique. Verification occurs through duplicate address detection.

- Preferred

  An address for which uniqueness has been verified. A node can send and receive unicast traffic to and from a preferred address. Router Advertisement messages include the period of time that an address can remain in the tentative and preferred states.

- Deprecated

  An address that is still valid but whose use is discouraged for new communication. Existing communication sessions can continue to use a deprecated address. Nodes can send and receive unicast traffic to and from deprecated addresses.

- Valid

  An address from which unicast traffic can be sent and received. The valid state covers both the preferred and deprecated states. Router Advertisement messages include the amount of time that an address remains in the valid state. The valid lifetime must be longer than or equal to the preferred lifetime.

- Invalid

  An address for which a node can no longer send or receive unicast traffic. An address enters the invalid state after the valid lifetime expires.

The following figure shows the relationship between the states of an autoconfigured address and the preferred and valid lifetimes.

**States of an Autoconfigured Address**




With the exception of link-local addresses, address autoconfiguration is specified only for hosts. Routers must obtain address and configuration parameters through another means (for example, manual configuration).

**Types of Autoconfiguration**

Autoconfiguration falls into three types:

1. Stateless

   Configuration is based on Router Advertisement messages. These messages include stateless address prefixes and require that hosts not use a stateful address configuration protocol.

2. Stateful

   Configuration is based on a stateful address configuration protocol, such as DHCPv6, to obtain addresses and other configuration options. Hosts use stateful address configuration when they receive Router Advertisement messages that do not include address prefixes and that require the hosts to use a stateful address configuration protocol. A host will also use a stateful address configuration protocol when no routers are present on the local link.

3. Both

   Configuration is based on Router Advertisement messages. These messages include stateless address prefixes but require hosts to use a stateful address configuration protocol.

For all autoconfiguration types, a link-local address is always configured.

**Autoconfiguration Process**

Address autoconfiguration for an IPv6 node occurs as follows:

1. A tentative link-local address is derived, based on the link-local prefix of FE80::/64 and the 64-bit interface identifier.

2. Duplicate address detection is performed to verify the uniqueness of the tentative link-local address.

   If the address is already in use, the node must be configured manually.

   If the address is not already in use, the tentative link-local address is assumed to be unique and valid. The link-local address is initialized for the interface. The corresponding solicited-node multicast link-layer address is registered with the

network adapter.

3. The host sends a Router Solicitation message.

4. If the host receives no Router Advertisement messages, then it uses a stateful address configuration protocol to obtain addresses and other configuration parameters. Windows Server 2003 does not support the use of a stateful address configuration protocol for IPv6.

5. If the host receives a Router Advertisement message, the host is configured based on the information in the message

6. For each stateless address prefix that the message includes:

   A tentative address is derived from the address prefix and the appropriate 64-bit interface identifier.

   The uniqueness of the tentative address is verified.

   If the tentative address is in use, the address is not initialized for the interface.

   If the tentative address is not in use, the address is initialized. Initialization includes setting the valid and preferred lifetimes based on information in the Router Advertisement message. Initialization also includes registering the corresponding solicited-node multicast link-layer address with the network adapter.

7. If specified in the Router Advertisement message, the host uses a stateful address configuration protocol to obtain additional addresses or configuration parameters.

**IPv6 Routing**

Similar to IPv4, IPv6 nodes use local routing tables to determine how to forward packets. IPv6 routing table entries are created by default when IPv6 initializes, and either a system administrator adds entries manually or communication with routers adds them automatically.

A routing table is present on each IPv6 node that is running Windows Server 2003. The routing table stores information about IPv6 network prefixes and how they can be reached (either directly or indirectly). Routing tables are not exclusive to IPv6 routers. IPv6 hosts that are running Windows Server 2003 also have IPv6 routing tables. Before a routing table is used, the destination cache is checked for an entry matching the destination address in the packet being forwarded. If the destination cache does not contain an entry for the destination address, the routing table is used to determine:

1. The next-hop address

   For a direct delivery (in which the destination is on a local link), the next-hop address is the destination address in the packet. For an indirect delivery (in which the destination is not on a local link), the next-hop address is the address of a router.

2. The next-hop interface

   The interface identifies the physical or logical interface that is used to forward the packet either to its destination or to the next router.

After the next-hop address and interface are determined, the destination cache is updated. Subsequent packets forwarded to the destination use the destination cache entry, rather than the routing table.

**Entry Types in IPv6 Routing Tables**

Entries in IPv6 routing tables store the following types of routes:

- Directly attached network routes

  Routes for network prefixes that are directly attached and typically have prefixes that are 64 bits long.

- Remote network routes

Routes for network prefixes that are not directly attached but are available across other routers. Remote network routes can be subnet network prefixes or prefixes for address spaces.

- Host routes

  Routes to specific IPv6 addresses. Host routes allow routing to occur for each IPv6 address. For host routes, the route prefix is a specific IPv6 address with a prefix length of 128 bits. Prefixes for both types of network routes are shorter than 128 bits.

- Default route

  The default route is used when a more specific network route or host route is not found. The default route prefix is ::/0.

**Route Determination Process**

IPv6 uses the following process to determine which routing table entry is used for the forwarding decision:

- For each entry in a routing table, IPv6 compares the bits in the network prefix to the same bits in the destination IPv6 address for the number of bits in the prefix length in the route. If all the bits in the network prefix match all the bits in the destination IPv6 address for the number of bits in the prefix length for the route, the route is a match for the destination.

- The list of matching routes is compiled. The route that has the largest prefix length (the route that matched the most high-order bits with the destination IPv6 address) is chosen. The longest matching route is the most specific route to the destination. If multiple entries with the longest match are found (multiple routes to the same network prefix, for example), the router uses the lowest metric to identify the best route. If multiple entries exist that are the longest match and the lowest metric, IPv6 can choose which routing table entry to use.

The match for any given destination finds routes in the following order:

1. A route that matches the entire destination IPv6 address (a host route with a 128-bit prefix length).

2. A route with the longest prefix length that matches the destination.

3. The default route (the network prefix ::/0).

The end result of the route determination process is the selection of a single route in the routing table. The selected route yields a next-hop address and an interface. If the route determination process on the sending host fails to find a route, IPv6 assumes that the destination is locally reachable. If the route determination process on a router fails to find a route, IPv6 sends an ICMPv6 Destination Unreachable-No Route Found message to the sending host and discards the packet.

**End-to-End IPv6 Delivery Process**

The following sections describe the process of forwarding an IPv6 packet from the sending host across one or more IPv6 routers to the final destination. These sections assume that the Hop-by-Hop Options, Destination Options, and Routing extension headers are not present.

**IPv6 on the Sending Host**

The process by which hosts send IPv6 packets combines the local host's data structures with the Neighbor Discovery protocol. IPv6 hosts use the following algorithm when sending a packet to an arbitrary destination:

1. Set the Hop Limit value to either a default or application-specified value.

2. Check the destination cache for an entry that matches the destination address.

3. If the destination cache contains an entry that matches the destination address, obtain the next-hop address and interface index from the destination cache entry. If the destination is a mobile IPv6 node and the sending host is a mobile correspondent node, the destination cache entry might contain a pointer to care-of destination cache. If so, obtain the next-hop address from the corresponding entry in the care-of destination cache. For more information about IPv6 mobility support, see the Internet draft titled "Mobility Support in IPv6." Go to step 7.

4. If the destination cache does not contain an entry that matches the destination address, check the local IPv6 routing table for the longest matching route to the destination address.

5. Based on the longest matching route, determine the next-hop address and the interface index used for forwarding the packet.

   If no route is found, IPv6 assumes that the destination is directly reachable. The next-hop address is set to the destination address, and an interface index is chosen.

6. Update the destination cache.

7. Check the neighbor cache for an entry that matches the next-hop address.

8. If the neighbor cache contains an entry that matches the next-hop address, obtain the link-layer address.

9. If the neighbor cache does not contain an entry that matches the next-hop address, use address resolution to obtain the link-layer address for the next-hop address.

   If address resolution is not successful, indicate an error.

10. Send the packet using the link-layer address of the entry in the neighbor cache.

The following figure shows this sending host process.

**Sending Host Process**




**IPv6 on the Router**

IPv6 routers use the following algorithm when they receive and forward a packet to an arbitrary destination:

1. Verify whether the destination address in the IPv6 packet corresponds to an address assigned to a router interface.

   If so, the router processes the IPv6 packet as the destination host. (For more information, see step 2 in the following subsection, "IPv6 on the Destination Host.")

2. Decrement the value of the Hop Limit field by 1.

   If the value of the Hop Limit field is 0, send an ICMPv6 Time Exceeded message to the sender, and discard the packet.

3. If the value of the Hop Limit field is 1 or greater, update the Hop Limit field in the IPv6 header of the packet.

4. Check the destination cache for an entry that matches the destination address.

5. If the destination cache contains an entry that matches the destination address, obtain the next-hop address and interface index in the entry, and go to step 9.

6. Check the local IPv6 routing table for the longest matching route to the destination address.

7. Based on the longest matching route, determine the next-hop IPv6 address and the interface index used for forwarding the packet.

   If no route is found, send an ICMPv6 Destination Unreachable-No Route Found message to the sending host, and discard the packet.

8. Update the destination cache.

9. If the interface on which the packet was received is the same as the interface on which the packet is being forwarded, the interface is a point-to-point link, and the Destination Address field matches a prefix assigned to the interface, send an ICMPv6 Destination Unreachable-Address Unreachable message to the sending host, and discard the packet. This prevents the needless "ping-pong" forwarding of IPv6 packets between the two interfaces on a point-to-point link for a packet whose destination matches the prefix of the point-to-point link but does not match the address of either interface. This condition and its solution is described in the Internet draft titled "Avoiding ping-pong packets on point-to-point links."

10. If the interface on which the packet was received is the same as the interface on which the packet is being forwarded and the Source Address field matches a prefix assigned to the interface, send a Redirect message to the sending host (subject to rate limiting).

11. Compare the link MTU of the next-hop interface to the size of the IPv6 packet being forwarded.

12. If the link MTU is smaller than the packet, send an ICMPv6 Packet Too Big message.

13. Check the neighbor cache of the next-hop interface for an entry that matches the next-hop address.

14. If the neighbor cache of the next-hop interface contains an entry that matches the next-hop address, obtain the link-layer address.

15. If the neighbor cache of the next-hop interface does not contain an entry that matches the next-hop address, use address resolution to obtain the link-layer address for the next-hop address.

    If address resolution is not successful, send an ICMPv6 Destination Unreachable-Destination Address Unreachable message, and discard the packet.

16. Send the packet using the link-layer address of the neighbor cache entry.

The following figures show this router forwarding process.

**Router Forwarding Process**




**Router Forwarding Process**




This entire process is repeated at each router in the path between the source and destination host.

**IPv6 on the Destination Host**

IPv6 hosts use the following algorithm when they receive IPv6 packets:

1. Verify whether the destination address in the packet corresponds to an IPv6 address that is assigned to a host interface.

   If the destination address is not assigned to a host interface, silently discard the packet.

2. Based on the Next Header field, process extension headers (if present), and pass the upper-layer PDU to the appropriate upper-layer protocol.

   If the protocol does not exist, send an ICMPv6 Parameter Problem message back to the sender, and discard the packet.

3. For TCP and UDP packets, check the destination port.

   If no application exists for the UDP port number, send an ICMPv6 Destination Unreachable-Destination Port Unreachable message back to the sender, and discard the packet. If no application exists for the TCP port number, send a TCP Connection Reset segment back to the sender, and discard the packet.

4. For TCP and UDP packets, pass the TCP segment or UDP message to the appropriate application.

**IPv6 Internet Connection Firewall**

A firewall provides security by creating a protective boundary between a computer or network and the Internet. When you install the Advanced Networking Pack for Windows XP, IPv6 Internet Connection Firewall (ICF) is provided and enabled by default to dynamically restrict traffic from the Internet. IPv6 ICF is different from the existing ICF in Windows XP, which restricts IPv4 traffic.

IPv6 ICF:

- Runs automatically and provides filters for all network connections on which IPv6 is enabled.
- Monitors all outbound traffic and dynamically creates inbound packet filters for response traffic. This is known as stateful filtering. IPv6 ICF silently discards all unsolicited inbound traffic.
- Logs IPv6 traffic events to a log file (separate from IPv4 ICF). By default, this log file is located at: *systemroot*\Pfirewall-v6.log.

IPv6 ICF protects your computer by silently dropping all packets that are initiated from a source outside the IPv6 ICF computer, such as the Internet. IPv6 ICF stops common attempts to illegally gain access to your computer or network by malicious Internet users using techniques such as port scanning. Instead of notifying you of inbound discarded traffic, IPv6 ICF creates entries in a security log from which you can view the activity that the firewall tracks.

IPv6 ICF is configured using commands in the **netsh firewall** context. You can use netsh commands to configure IPv6 ICF to allow specific types of ICMPv6 traffic or traffic to specific TCP or UDP ports.

**Note**

- Windows XP shows only the IPv4 ICF configuration in Network Connections. IPv6 ICF might appear disabled, but it is actually enabled and filtering IPv6 traffic.

**IPv6 Transition Technologies**

Protocol transitions are not easy, and the transition from IPv4 to IPv6 is no exception. Protocol transitions are typically deployed by installing and configuring the new protocol on all nodes within the network and verifying that all node and router operations work. Although this might be possible in a small- or medium-sized organization, the challenge of making a rapid protocol transition in a large organization is very difficult. Additionally, given the scope of the Internet, rapid protocol transition becomes an impossible task.

The transition from IPv4 to IPv6 will take years, and organizations or hosts within organizations might continue to use IPv4 indefinitely. Therefore, although migration is the long-term goal, equal consideration must be given to the interim coexistence of IPv4 and IPv6 nodes.

RFC 1752 defines the following transition criteria:

- Existing IPv4 hosts can be upgraded at any time, independent of the upgrade of other hosts or routers.
- Hosts that use only IPv6 can be added at any time, without dependencies on other hosts or routing infrastructure.
- IPv4 hosts on which IPv6 is installed can continue to use their IPv4 addresses and do not need additional addresses.
- Little preparation is required to either upgrade IPv4 nodes to IPv6 or deploy new IPv6 nodes.

The inherent lack of dependencies between IPv4 and IPv6 hosts, IPv4 routing infrastructure, and IPv6 routing infrastructure requires several mechanisms that allow seamless coexistence.

**Node Types**
RFC 2893 defines the following node types:

**IPv4-only Node**
A node that implements only IPv4 (and has only IPv4 addresses). This node does not support IPv6. Most hosts and routers installed today are IPv4-only nodes.

**IPv6-only Node**
A node that implements only IPv6 (and has only IPv6 addresses). This node is able to communicate only with IPv6 nodes and applications. This type of node is not common today, but it might become more prevalent as smaller devices such as cellular phones and handheld computing devices include the IPv6 protocol.

**IPv6/IPv4 Node**
A node that implements both IPv4 and IPv6. This node is IPv6-enabled if it has an IPv6 interface configured.

**IPv4 Node**
A node that implements IPv4 (it can send and receive IPv4 packets). An IPv4 node can be an IPv4-only node or an IPv6/IPv4 node.

**IPv6 Node**
A node that implements IPv6 (it can send and receive IPv6 packets). An IPv6 node can be an IPv6-only node or an IPv6/IPv4 node.

For coexistence to occur, the largest number of nodes (IPv4 or IPv6 nodes) can communicate using an IPv4 infrastructure, an IPv6 infrastructure, or an infrastructure that is a combination of IPv4 and IPv6. True migration is achieved when all IPv4 nodes are converted to IPv6-only nodes. However, for the foreseeable future, practical migration is achieved when as many IPv4-only nodes as possible are converted to IPv6/IPv4 nodes. IPv4-only nodes can communicate with IPv6-only nodes only through an IPv4-to-IPv6 proxy or translation gateway.

**Address Compatibility**
The following addresses are defined to help IPv4 and IPv6 nodes coexist:

**IPv4-compatible Addresses**
IPv6/IPv4 nodes that are communicating with IPv6 over an IPv4 infrastructure use IPv4-compatible addresses, 0:0:0:0:0:0:w.x.y.z or ::w.x.y.z (where w.x.y.z is the dotted decimal representation of a public IPv4 address). When an IPv4-compatible address is used as an IPv6 destination, IPv6 traffic is automatically encapsulated with IPv4 headers and sent to their destinations using the IPv4 infrastructure.

**IPv4-mapped Addresses**
The IPv4-mapped address, 0:0:0:0:0:FFFF:w.x.y.z or ::FFFF:w.x.y.z, is used to represent an IPv4-only node to an IPv6 node. It is used only for internal representation. The IPv4-mapped address is never used as a source or destination address of an IPv6 packet. The IPv6 protocol for Windows Server 2003 does not support IPv4-mapped addresses. Some IPv6 implementations use IPv4-mapped addresses when translating traffic between IPv4-only and IPv6-only nodes.

**6over4 Addresses**
Each 6over4 address comprises a valid 64-bit unicast address prefix and the interface identifier ::WWXX:YYZZ (where WWXX:YYZZ is the colon-hexadecimal representation of w.x.y.z, a unicast IPv4 address assigned to an interface). An example of a link-local 6over4 address based on the IPv4 address of 131.107.4.92 is FE80::836B:45C. 6over4 addresses represent a host that use the automatic tunneling mechanism defined in RFC 2529.

**6to4 Addresses**
6to4 addresses are based on the prefix 2002:WWXX:YYZZ::/48 (where WWXX:YYZZ is the colon-hexadecimal representation of w.x.y.z, a public IPv4 address assigned to an interface). 6to4 addresses represent sites that use the automatic tunneling mechanism defined in RFC 3056, also known as 6to4. For more information, see “6to4” later in this section.

**ISATAP Addresses**
Each Intra-site Automatic Tunnel Addressing Protocol (ISATAP) addresses comprise a valid 64-bit unicast address prefix and the interface identifier ::0:5EFE:w.x.y.z (where w.x.y.z is a unicast IPv4 address assigned to an interface). An example of a link-local ISATAP address is FE80::5EFE:131.107.4.92. ISATAP addresses represent hosts that use the automatic tunneling mechanism defined in the Internet draft titled “Intra-Site Automatic Tunnel Addressing Protocol (ISATAP).” For more information, see “ISATAP” later in this section.

**Teredo Addresses**
Teredo addresses use the prefix 3FFE:831F::/32. An example of a Teredo address is 3FFE:831F:CE49:7601:8000:EFFF:62C3:FFFE.

Beyond the first 32 bits, Teredo addresses encode the IPv4 address of a Teredo server, flags, and the encoded version of the external address and port of a Teredo client. Teredo addresses represent hosts that use the automatic tunneling mechanism defined in the Internet draft titled "Teredo: Tunneling IPv6 over UDP through NATs." For more information, see "Teredo" later in this section.

**Transition Mechanisms**

The following mechanisms are used to help IPv6 coexist with an IPv4 infrastructure and to provide an eventual transition to an IPv6-only infrastructure:

- Dual IP layer
- IPv6 over IPv4 tunneling
- DNS infrastructure

**Dual IP Layer**

The dual IP layer is an implementation of the TCP/IP suite of protocols that includes both an IPv4 Internet layer and an IPv6 Internet layer. IPv6/IPv4 nodes use this mechanism to communicate with both IPv4 and IPv6 nodes. A dual IP layer contains a single implementation of Host-to-Host Transport layer protocols such as TCP and UDP. All upper layer protocols in an implementation of dual IP layer can communicate over IPv4, IPv6, or IPv6 tunneled in IPv4.

The following figure shows a dual IP layer architecture.

**A Dual IP Layer Architecture**




The IPv6 protocol for Windows Server 2003 is not a dual IP layer. The IPv6 protocol driver, Tcpip6.sys, contains a separate implementation of TCP and UDP and is sometimes referred to as a dual-stack implementation. The following figure shows the dual stack architecture of the IPv6 protocol for Windows Server 2003.

**The Dual Stack Architecture of the IPv6 Protocol for Windows Server 2003**




Although the IPv6 protocol for Windows Server 2003 is not a dual IP layer, it functions in the same way as a dual IP layer in terms of providing functionality for IPv6 transition.

**IPv6 over IPv4 Tunneling**

IPv6 over IPv4 tunneling is the encapsulation of IPv6 packets with an IPv4 header so that IPv6 packets can be sent over an IPv4 infrastructure. Within the IPv4 header:

- The IPv4 Protocol field is set to 41 to indicate an encapsulated IPv6 packet.
- The Source and Destination fields are set to IPv4 addresses of the tunnel endpoints. The tunnel endpoints are either manually configured as part of the tunnel interface or are automatically derived from the sending interface, the next-hop address of the matching route, or the source and destination IPv6 addresses in the IPv6 header.

The following figure shows IPv6 over IPv4 tunneling.

**IPv6 over IPv4 Tunneling**




For IPv6 over IPv4 tunneling, the IPv6 path MTU is typically 20 less than the IPv4 path MTU However, if the IPv4 path MTU is not stored for each tunnel, an intermediate IPv4 router might need to fragment the IPv4 packet. In this case, the Don't Fragment flag in the IPv4 header must be set to 0 when IPv6 over IPv4 tunneled packets are sent.

**DNS Infrastructure**

Because most users refer to network resources by name rather than by address, successful coexistence requires upgrading the DNS infrastructure. Upgrading the DNS infrastructure consists of populating the DNS servers with records to support IPv6 name-to-address and address-to-name resolution. After the sending host obtains addresses using a DNS name query, the node must select which addresses to use for communication.

**Address Records**

The DNS infrastructure must contain the following resource records (populated either manually or dynamically) for the successful resolution of domain names to addresses:

- A records for IPv4-only and IPv6/IPv4 nodes
- AAAA records for IPv6-only and IPv6/IPv4 nodes

**Pointer Records**

The DNS infrastructure must contain the following resource records (populated either manually or dynamically) for the successful resolution of address to domain names (reverse queries):

- PTR records in the IN-ADDR.ARPA domain for IPv4-only and IPv6/IPv4 nodes
- PTR records in the IP6.ARPA domain for IPv6-only and IPv6/IPv4 nodes (optional)

**Address Selection Rules**

For name-to-address resolution, after the querying node obtains the set of addresses corresponding to the name, the node must determine the set of addresses to choose as source and destination for outbound packets.

This choice is not an issue in today's environment, in which IPv4-only nodes prevail. However, in an environment in which IPv4 and IPv6 coexist, the set of addresses returned in a DNS query might contain multiple IPv4 and IPv6 addresses. The querying host is configured with at least one IPv4 address and (typically) multiple IPv6 addresses. It is not an easy task to decide which type of address (IPv4 versus IPv6) and then which scope (public versus private for IPv4 and link-local versus global versus compatible for IPv6) to use for both the source and the destination addresses.

Default address selection rules are currently under discussion and are described in RFC 3484.

You can view the default address selection rules for the IPv6 protocol for Windows Server 2003 using the **netsh interface ipv6 show prefixpolicy** command to display the prefix policy table. You can modify the entries in the prefix policy table using the **netsh interface ipv6 add|set|delete prefixpolicy** commands. By default, IPv6 addresses in DNS query responses are preferred over IPv4 addresses.

**Tunneling Configurations**

RFC 2893 defines the following configurations with which to tunnel IPv6 traffic between IPv6/IPv4 nodes over an IPv4 infrastructure:

- Router-to-Router
- Host-to-Router or Router-to-Host
- Host-to-Host

**Note**

- IPv6 over IPv4 tunneling describes only an encapsulation of IPv6 packets with an IPv4 header so that IPv6 nodes are reachable across an IPv4 infrastructure. Unlike tunneling for Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP), no messages are exchanged for tunnel setup, maintenance, or termination. Additionally, IPv6 over IPv4 tunneling does not provide security for tunneled IPv6 packets.

#### Router-to-Router

In the tunneling configuration between routers, two IPv6/IPv4 routers connect two IPv4 or IPv6 infrastructures over an IPv4 infrastructure. The tunnel spans a logical link in the path between the source and destination. The IPv6 over IPv4 tunnel between the two routers acts as a single hop. Routes within each IPv4 or IPv6 infrastructure point to the IPv6/IPv4 router on the edge. For each IPv6/IPv4 router, an interface represents the IPv6 over IPv4 tunnel and routes that use the tunnel interface.

The following figure shows tunneling between two routers.

**Router-to-Router Tunneling**




Examples of this tunneling configuration include:

- An IPv6-only test lab that tunnels across an organization's IPv4 infrastructure to reach the IPv6 Internet.
- Two IPv6-only routing domains that tunnel across the IPv4 Internet.
- A 6to4 router that tunnels across the IPv4 Internet to reach another 6to4 router or a 6to4 relay router. For more information about "6to4", see "6to4" later in this section.

#### Host-to-Router and Router-to-Host

In the host-to-router tunneling configuration, an IPv6/IPv4 node that resides within an IPv4 infrastructure creates an IPv6 over IPv4 tunnel to reach an IPv6/IPv4 router. The tunnel spans the first segment of the path between the source and destination nodes. The IPv6 over IPv4 tunnel between the IPv6/IPv4 node and the IPv6/IPv4 router acts as a single hop.

On the IPv6/IPv4 node, a tunnel interface representing the IPv6 over IPv4 tunnel is created, and a route (typically a default route) is added using the tunnel interface. The IPv6/IPv4 node tunnels the IPv6 packet based on the matching route, the tunnel interface, and the next-hop address of the IPv6/IPv4 router.

In the router-to-host tunneling configuration, an IPv6/IPv4 router creates an IPv6 over IPv4 tunnel across an IPv4 infrastructure to reach an IPv6/IPv4 node. The tunnel endpoints span the last segment of the path between the source node and destination node. The IPv6 over IPv4 tunnel between the IPv6/IPv4 router and the IPv6/IPv4 node acts as a single hop.

On the IPv6/IPv4 router, a tunnel interface representing the IPv6 over IPv4 tunnel is created, and a route (typically a subnet route) is added using the tunnel interface. The IPv6/IPv4 router tunnels the IPv6 packet based on the matching subnet route, the tunnel interface, and the destination address of the IPv6/IPv4 node.

The following figure shows host-to-router (for traffic traveling from Node A to Node B) and router-to-host (for traffic traveling from Node B to Node A) tunneling.

**Host-to-Router and Router-to-Host Tunneling**




Examples of host-to-router and router-to-host tunneling include:

- An IPv6/IPv4 host that tunnels across an organization's IPv4 infrastructure to reach the IPv6 Internet.

- An ISATAP host that tunnels across an IPv4 network to an ISATAP router to reach the IPv4 Internet, another IPv4 network, or an IPv6 network. For more information about "ISATAP", see "ISATAP" later in this section.

- An ISATAP router that tunnels across an IPv4 network to reach an ISATAP host.

**Host-to-Host**
In the tunneling configuration between hosts, an IPv6/IPv4 node that resides within an IPv4 infrastructure creates an IPv6 over IPv4 tunnel to reach another IPv6/IPv4 node that resides within the same IPv4 infrastructure. The tunnel spans the entire path between the source and destination nodes. The IPv6 over IPv4 tunnel between the IPv6/IPv4 nodes acts as a single hop.

On each IPv6/IPv4 node, an interface representing the IPv6 over IPv4 tunnel is created. Routes might indicate that the destination node is on the same logical subnet defined by the IPv4 infrastructure. Based on the sending interface, the optional route, and the destination address, the sending host tunnels the IPv6 traffic to the destination.

The following figure shows tunneling between hosts.

**Host-to-Host Tunneling**




Examples of this tunneling configuration include:

- IPv6/IPv4 hosts that use ISATAP addresses to tunnel across an organization's IPv4 infrastructure.

- IPv6/IPv4 hosts that use IPv4-compatible addresses to tunnel across an organization's IPv4 infrastructure.

**Types of Tunnels**
RFC 2893 defines the following types of tunnels:

- Configured

- Automatic

**Configured Tunnels**
A configured tunnel is one whose endpoints must be configured manually. In a configured tunnel, the endpoints have IPv4 addresses that are not derived from the IPv6 source or destination addresses or the next-hop address of the matching route.

Tunnels between routers are typically configured manually. To configure the tunnel interface, the IPv4 addresses of the tunnel endpoints must be manually specified along with static routes that use the tunnel interface.

To manually configure tunnels for the IPv6 protocol for Windows Server 2003, use the **netsh interface ipv6 add v6v4tunnel** command.

**Automatic Tunnels**
An automatic tunnel is one whose endpoints do not need to be configured manually. Tunnel endpoints are determined from logical tunnel interfaces, routes, and source and destination IPv6 addresses.

The IPv6 protocol for Windows Server 2003 supports the following automatic tunneling technologies:

- 6to4, which is enabled by default.

- ISATAP, which is enabled by default. For more information, see "ISATAP" later in this section.

- IPv6 Automatic Tunneling, which is disabled by default.

- 6over4, which is disabled by default.

Additionally, the IPv6 protocol for Windows XP with Service Pack 1 also supports Teredo when the Advanced Networking Pack

for Windows XP is installed. Teredo is enabled by default. For more information, see "Teredo" later in this section.

**6to4**

6to4 is a technology that assigns addresses and automatically configures tunnels between routers to provide unicast IPv6 connectivity between IPv6 sites and hosts across the IPv4 Internet. 6to4 uses the global address prefix:

```
2002:WWXX:YYZZ::/48
```

in which *WWXX*:*YYZZ* is the colon-hexadecimal representation of a public IPv4 address (*w.x.y.z*) assigned to a site or host. The full 6to4 address is:

```
2002:WWXX:YYZZ:[Subnet ID]:[Interface ID]
```

RFC 3056 describes 6to4 in the following terms:

- 6to4 host

Any IPv6 host that is configured with at least one 6to4 address (a global address with the 2002::/16 prefix). 6to4 hosts do not require any manual configuration, and they create 6to4 addresses using standard address autoconfiguration mechanisms.

- 6to4 router

An IPv6/IPv4 router that supports the use of a 6to4 tunnel interface and that is typically used to forward 6to4-addressed traffic between the 6to4 hosts within a site and other 6to4 routers or 6to4 relay routers on an IPv4 internetwork, such as the Internet. 6to4 routers require additional processing logic for proper encapsulation and decapsulation, and they might require manual configuration.

- 6to4 relay router

An IPv6/IPv4 router that forwards 6to4-addressed traffic between 6to4 routers on the Internet and hosts on the IPv6 Internet.

The following figure shows 6to4 components.

**6to4 Components**




Within a site, IPv6 routers advertise 2002:*WWXX*:*YYZZ*:[*Subnet ID*]::/64 prefixes so that hosts can create an autoconfigured 6to4 address and 64-bit prefix routes are used to deliver traffic between 6to4 hosts within the site. Hosts on individual subnets are automatically configured with a 64-bit subnet route for direct delivery to neighbors and a default route with the next-hop address of the advertising router. All IPv6 traffic that does not match a 64-bit prefix used by one of the subnets within the site is forwarded to a 6to4 router on the site border.

The 6to4 router on the site border has a 2002::/16 route that is used to forward traffic to other 6to4 sites and a default route (::/0) that is used to forward traffic to a 6to4 relay router.

In the example network shown in the previous figure, Host A and Host B can communicate with each other because of a default route using the next-hop address of the 6to4 router in Site 1. When Host A communicates with Host C in another site, Host A sends the traffic to the 6to4 router in Site 1 as IPv6 packets. The 6to4 router in Site 1, using the 6to4 tunnel interface and the 2002::/16 route in its routing table, encapsulates the packet with an IPv4 header and tunnels the packet to the 6to4 router in

Site 2. When the 6to4 router in Site 2 receives the packet, the router removes the IPv4 header and, using the 64-bit prefix route in its routing table, forwards the IPv6 packet to Host C.

In this example, Host A (with the interface ID ID_A) resides on subnet 1 within Site 1, which uses the public IPv4 address of 157.60.91.123. Host C (with the interface ID ID_C) resides on subnet 2 within Site 2, which uses the public IPv4 address of 131.107.210.49. When the 6to4 router in Site 1 sends the IPv4-encapsulated IPv6 packet to the 6to4 router in Site 2, the addresses in the IPv4 and IPv6 headers are as listed in the following table.

**Example 6to4 Addresses**

| Field | Value |
|---|---|
| IPv6 Source Address | 2002:9D3C:5B7B:1:[ID_A] |
| IPv6 Destination Address | 2002:836B:D231:2:[ID_C] |
| IPv4 Source Address | 157.60.91.123 |
| IPv4 Destination Address | 131.107.210.49 |

For a more detailed example of 6to4 traffic using ISATAP-derived interface identifiers, see "ISATAP" later in this section.

The following types of communication are possible when 6to4 hosts, an IPv6 routing infrastructure within a site, a 6to4 router at the site boundary, and a 6to4 relay router are used:

- A 6to4 host can communicate with another 6to4 host within the same site.

  This type of communication is available with the IPv6 routing infrastructure, which provides reachability to all hosts within the site. The previous figure shows this type of communication between Host A and Host B.

- A 6to4 host can communicate with 6to4 hosts in other sites across the IPv4 Internet.

  This type of communication occurs when a 6to4 host forwards IPv6 traffic that is destined to a 6to4 host in another site to the 6to4 router for the local site. The 6to4 router for the local site tunnels the IPv6 traffic through the IPv4 Internet to the 6to4 router at the destination site. The 6to4 router at the destination site removes the IPv4 header and forwards the IPv6 packet to the appropriate 6to4 host by using the IPv6 routing infrastructure of the destination site. The previous figure shows this type of communication between Host A and Host C.

- A 6to4 host can communicate with hosts on the IPv6 Internet.

  This type of communication occurs when a 6to4 host forwards IPv6 traffic that is destined for an IPv6 Internet host to the 6to4 router for the local site. The 6to4 router for the local site tunnels the IPv6 traffic to a 6to4 relay router that is connected to both the IPv4 Internet and the IPv6 Internet. The 6to4 relay router removes the IPv4 header and forwards the IPv6 packet to the appropriate IPv6 Internet host by using the routing infrastructure of the IPv6 Internet. The previous figure shows this type of communication between Host A and Host D.

All of these types of communication use IPv6 traffic without having to obtain either a direct connection to the IPv6 Internet or an IPv6 global address prefix from an ISP.

**Communicating Using a 6to4 Address**

The IPv6 protocol for Windows Server 2003 contains a 6to4 component that supports 6to4 hosts and 6to4 routers. If a public IPv4 address is assigned to an interface on the host and a global prefix is not received in a router advertisement, the 6to4 component:

- Configures 6to4 addresses on the 6to4 Tunneling Pseudo-Interface for all public IPv4 addresses that are assigned to interfaces on the computer.

- Creates a 2002::/16 route that forwards all 6to4 traffic with the 6to4 Tunneling Pseudo-Interface (interface index 3). All traffic forwarded by this host to 6to4 destinations is encapsulated with an IPv4 header.

- Performs a DNS query to obtain the IPv4 address of a 6to4 relay router on the Internet. You can also use the **netsh interface ipv6 6to4 set relay** command to specify the DNS name to query. If the query is successful, a default route is

added using the 6to4 Tunneling Pseudo-Interface, and the next-hop address is set to the 6to4 address of the 6to4 relay router.

The results of 6to4 component autoconfiguration vary depending on the configuration of the host. The following figure shows how 6to4 is configured for different types of hosts running Windows Server 2003 (except IPv6 host D).

**6to4 for Hosts Running Windows**




For a host that is assigned a private IPv4 address or receives a router advertisement for a global prefix, no 6to4 addresses are assigned to the 6to4 Tunneling Pseudo-Interface. Addresses are autoconfigured based on the global prefix, and the routing table contains a 64-bit global prefix route and default route. This configuration corresponds to Host A, Host B, and Host C in the previous figure.

For a host that is assigned a public IPv4 address and does not receive a router advertisement for a global prefix, a 6to4 address of the form 2002:*WWXX*:*YYZZ*::*WWXX*:*YYZZ* is automatically assigned to the 6to4 Tunneling Pseudo-Interface. A 2002::/16 route using the 6to4 Tunneling Pseudo-Interface is added, and, if the DNS query for the 6to4 relay router is successful, a default route using the 6to4 address of the 6to4 relay router as the next hop is added. This configuration corresponds to Host E in the previous figure, a host that is directly connected to the IPv4 Internet. In this case, the host is acting as its own site and its own 6to4 router.

The 6to4 component can also enable a computer running Windows Server 2003 as a 6to4 router by utilizing the configuration of the Internet Connection Sharing (ICS) feature. This configuration corresponds to 6to4 router 1 and 6to4 router 2 in the previous figure.

If ICS is enabled on an interface that is assigned a public IPv4 address, the 6to4 component automatically:

- Enables IPv6 forwarding on both the public and private interfaces.

  The public interface is connected to the Internet. The private interface is connected to a single-subnet intranet and uses private IPv4 addresses from the 192.168.0.0/24 prefix.

- Sends Router Advertisement messages on the private interface.

  The router advertisements advertise the ICS computer as a default router and contain a global 6to4 address prefix that is based on the public IPv4 address assigned to the public interface. The Subnet ID in the 6to4 address prefix is set to the interface index of the interface on which the advertisements are sent.

For example, for an ICS computer using the public IPv4 address of 131.107.23.89 and interface 5 as the private interface, the advertised prefix would be 2002:836B:1759:5::/64. Private hosts that receive this router advertisement would create global addresses through normal address autoconfiguration. The hosts would also add a 2002:836B:1759:5::/64 route for the local subnet and a default route with a next-hop address of the link-local address of the ICS computer's private interface. Private hosts can communicate with each other on the same subnet using the 2002:836B:1759:5::/64 route. For all other 6to4 sites or the IPv6 Internet, the IPv6 packets are forwarded to the ICS computer using the default route.

For traffic to other 6to4 sites, the ICS computer uses its 2002::/16 route and encapsulates the IPv6 traffic with IPv4 headers and

sends the packets across the IPv4 Internet to another 6to4 router. For all other IPv6 traffic, the ICS computer uses its default route and encapsulates the IPv6 traffic with IPv4 headers and sends the packets across the IPv4 Internet to a 6to4 relay router.

**Note**

- The 6to4 component does not perform network address translation on the IPv6 packets that it forwards. ICS translates network addresses for IPv4 packets being forwarded to and from private hosts. The 6to4 component uses the ICS configuration to determine the public IPv4 address and public interface.

**ISATAP**

ISATAP is a technology that assigns addresses, configures tunnels between hosts and between routers and hosts, and provides unicast IPv6 connectivity between IPv6 hosts across an IPv4 intranet. ISATAP is described in the Internet draft titled "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)." ISATAP hosts do not require any manual configuration, and they create ISATAP addresses using standard mechanisms for address autoconfiguration.

IPv6/IPv4 nodes can use ISATAP to communicate on an IPv4 network. ISATAP addresses use the locally administered interface identifier ::0:5EFE:*w.x.y.z* where the *w.x.y.z* portion is any public or private unicast IPv4 address.

The ISATAP interface identifier can be combined with any 64-bit prefix that is valid for IPv6 unicast addresses. This includes the link-local address prefix (FE80::/64) and global prefixes (including 6to4 prefixes).

Like IPv4-compatible addresses, 6over4 addresses, and 6to4 addresses, ISATAP addresses contain embedded IPv4 addresses that are used to determine either the source or destination IPv4 addresses within the IPv4 header when ISATAP-addressed IPv6 traffic is tunneled across an IPv4 network.

By default, the IPv6 protocol for Windows Server 2003 automatically configures the link-local ISATAP address of FE80::5EFE:*w.x.y.z* on the Automatic Tunneling Pseudo-Interface (interface index 2) for each IPv4 address that is assigned to the node. This link-local ISATAP address allows two hosts to communicate over an IPv4 network by using each other's link-local ISATAP addresses.

For example, Host A is configured with the IPv4 address of 10.40.1.29, and Host B is configured with the IPv4 address of 192.168.41.30. When the IPv6 protocol for Windows Server 2003 is started, Host A is automatically configured with the ISATAP address of FE80::5EFE:10.40.1.29, and Host B is automatically configured with the ISATAP address of FE80::5EFE:192.168.41.30. The following figure shows this configuration.

**An Example ISATAP Configuration**




When Host A sends IPv6 traffic to Host B by using Host B's link-local ISATAP address, the source and destination addresses for the IPv6 and IPv4 headers are as listed in the following table.

**Example Link-Local ISATAP Addresses**

| Field | Value |
|---|---|
| IPv6 Source Address | FE80::5EFE:10.40.1.29 |
| IPv6 Destination Address | FE80::5EFE:192.168.41.30 |
| IPv4 Source Address | 10.40.1.29 |
| IPv4 Destination Address | 192.168.41.30 |

To test connectivity, use the ping command. For example, Host A would use the following command to ping Host B by using its link-local ISATAP address:

```
ping FE80::5EFE:192.168.41.30%2
```

Because the destination of the ping command is a link-local address, the *%ZoneID* portion of the command specifies the interface index of the interface from which traffic is sent. In this case, %2 specifies interface 2, which is the interface index assigned to the Automatic Tunneling Pseudo-Interface on Host A. The Automatic Tunneling Pseudo-Interface uses the link-local ISATAP address assigned to the interface as a source. The Automatic Tunneling Pseudo-Interface also uses the last 32 bits in the source and destination IPv6 addresses (corresponding to the embedded IPv4 addresses) as the source and destination IPv4 addresses.

**ISATAP Router**

The use of link-local ISATAP addresses allows IPv6/IPv4 hosts on the same logical subnet (an IPv4 network) to communicate with each other, but not with other IPv6 addresses on other subnets. To communicate outside the logical subnet using ISATAP-derived global addresses, IPv6 hosts using ISATAP addresses must tunnel their packets to an ISATAP router. The following figure shows this configuration.

**An ISATAP Router**




An ISATAP router is an IPv6 router that:

- Forwards packets between ISATAP hosts on a logical subnet (an IPv4 network) and hosts on other subnets.

  The other subnets can be other IPv4 networks (such as a portion of an organization network or the IPv4 Internet) or subnets in a native IPv6 routing domain (such as an organization's IPv6 network or the IPv6 Internet).

- Acts as a default router for ISATAP hosts.

- Advertises address prefixes to identify the logical subnet on which ISATAP hosts are located. ISATAP hosts use the advertised address prefixes to configure global ISATAP addresses.

When an ISATAP host receives a Router Advertisement message from an ISATAP router, a default route (::/0) is added using the Automatic Tunneling Pseudo-Interface with next-hop address set to the link-local ISATAP address that corresponds to the logical subnet interface of the ISATAP router. When packets destined to locations outside the logical subnet are sent, they are tunneled to the IPv4 address that corresponds to the ISATAP router's interface on the logical subnet defined by the IPv4 network containing the ISATAP router and ISATAP host. The ISATAP router then forwards the IPv6 packet.

For the IPv6 protocol for Windows Server 2003, the configuration of the intranet IPv4 address of the ISATAP router is obtained through one of the following:

- The resolution of the name ISATAP to an IPv4 address.

- The **netsh interface ipv6 isatap set router** command.

**Resolving the ISATAP Name**

When the IPv6 protocol for Windows Server 2003 starts, it attempts to resolve the name "ISATAP" to an IPv4 address using normal TCP/IP name resolution techniques. These techniques include the following:

1. Checking the local host name.

2. Checking the DNS client resolver cache, which includes the entries in the Hosts file. This file is located in the *systemroot*\system32\drivers\etc folder.

3. Forming a fully qualified domain name and sending a DNS name query. For example, if the computer running Windows Server 2003 is a member of the example.microsoft.com domain (and example.microsoft.com is the only domain name in the search list), the computer sends a DNS query to resolve the name isatap.example.microsoft.com.

4. Converting the ISATAP name into the NetBIOS name "ISATAP <00>" and checking the NetBIOS name cache.

5. Sending a NetBIOS name query to the configured Windows Internet Name Service (WINS) servers.

6. Sending NetBIOS broadcasts.

7. Checking the Lmhosts file. This file is located in the *systemroot*\system32\drivers\etc folder.

If the name is resolved, the host sends an IPv4-encapsulated Router Solicitation message to the ISATAP router. The ISATAP router responds with an IPv4-encapsulated unicast Router Advertisement message that contains prefixes to use for autoconfiguring ISATAP-based addresses and, optionally, advertising itself as a default router.

To ensure that at least one of these attempts is successful, you can do one of the following:

- If the ISATAP router is a computer running Windows Server 2003, name the computer ISATAP, and it will automatically register the appropriate records in DNS and WINS.

- Manually create an ISATAP address (A) record in the appropriate domain in DNS. For example, create an A record for isatap.example.microsoft.com.

- Manually create a static WINS record in WINS for the NetBIOS name "ISATAP <00>."

- Add the following entry to the Hosts file of the computers that need to resolve the name ISATAP:

  IPv4Address ISATAP

- Add the following entry to the Lmhosts file of the computers that need to resolve the name ISATAP:

  IPv4Address ISATAP

**Resolving the _ISATAP Name for Windows XP (without SP1)**
When the IPv6 protocol for Windows XP (without SP1) starts, it attempts to resolve the name "_ISATAP," rather than "ISATAP." To ensure that a computer running Windows XP (without SP1) can resolve the name _ISATAP, you can do one of the following:

- Manually create _ISATAP canonical name (CNAME) records in the appropriate domains in DNS. A CNAME record maps a name that is an alias to another name. For example, assuming that an A record already exists for the name isatap.example.microsoft.com, create a CNAME record that maps _isatap.example.microsoft.com to isatap.example.microsoft.com.

- Manually create a static WINS record for the NetBIOS name "_ISATAP <00>."

- Add the following entry to the Hosts file of the computers running Windows XP:

  *IPv4Address* _ISATAP

- Add the following entry to the Lmhosts file of the computers running Windows XP:

  *IPv4Address* _ISATAP

### Note

- Windows XP with SP1 attempts to resolve the name "ISATAP" to determine the IPv4 address of the ISATAP router. The methods described here are not needed if all your computers are running either Windows Server 2003 or Windows XP with SP1.

**Using the Netsh Interface Ipv6 Isatap Set Router Command**
Although the automatic resolution of the ISATAP name is the recommended method for configuring the IPv4 address of the ISATAP router, you can configure this address manually by using the **netsh interface ipv6 isatap set router** command. The syntax of this command is:

**netsh interface ipv6 isatap set router** *AddressOrName*

where *AddressOrName* is the name or IPv4 address of the ISATAP router's intranet interface. For example, if the ISATAP router's IPv4 address is 192.168.39.1, the command is:

**netsh interface ipv6 isatap set router 192.168.39.1**

After a host has been configured, it sends an IPv4-encapsulated Router Solicitation message to the ISATAP router. The ISATAP router responds with an IPv4-encapsulated unicast Router Advertisement message that contains prefixes to use for autoconfiguring ISATAP-based addresses. This additional configuration is needed only when no IPv6 router is on the host's subnet.

**ISATAP and 6to4 Example**
The following figure shows two ISATAP hosts that are using 6to4 prefixes to communicate across the Internet even though each site is using the 192.168.0.0/16 private address space internally.

**Communication Between ISATAP Hosts in Different 6to4 Sites**




In this configuration:

- ISATAP Host A automatically configures a link-local ISATAP address of FE80::5EFE:192.168.12.9 on its Automatic Tunneling Pseudo-Interface.
- 6to4 Router A automatically configures a link-local ISATAP address of FE80::5EFE:192.168.204.1 on its Automatic Tunneling Pseudo-Interface.
- 6to4 Router B automatically configures a link-local ISATAP address of FE80::5EFE:192.168.39.1 on its Automatic Tunneling Pseudo-Interface.
- ISATAP Host B automatically configures a link-local ISATAP address of FE80::5EFE:192.168.141.30 on its Automatic Tunneling Pseudo-Interface.

ISATAP Host A can reach 6to4 Router A and all other hosts within Site A using link-local ISATAP addresses. However, ISATAP Host A cannot reach any addresses outside Site A. 6to4 Router A constructs the global prefix 2002:9D36:1:5::/64. (9D36:1 is the colon-hexadecimal notation for 157.54.0.1 and 5 is the interface index of 6to4 Router A's intranet interface.) 6to4 Router A also advertises the prefix using a Router Advertisement message on its intranet interface. However, ISATAP Host A is not on 6to4

Router A's subnet and will never create a global address based on this 6to4 prefix.

To configure ISATAP Host A to receive the router advertisement from 6to4 Router A, the network administrator for Site A has configured 6to4 Router A as an ISATAP router. The administrator has also added an A record to Site A's DNS infrastructure so that the name ISATAP is resolved to the IPv4 address of 192.168.204.1. Upon startup, the IPv6 protocol on Host A resolves the ISATAP name and sends a Router Solicitation message to the addresses listed in the following table.

**Addresses in the Router Solicitation Message**

| Field | Value |
|---|---|
| IPv6 Source Address | FE80::5EFE:192.168.12.9 |
| IPv6 Destination Address | FF02::2 |
| IPv4 Source Address | 192.168.12.9 |
| IPv4 Destination Address | 192.168.204.1 |

6to4 Router A receives the Router Solicitation message from ISATAP Host A and sends back a unicast Router Advertisement message advertising itself as a default router with a Prefix Information option to automatically configure IPv6 addresses using the prefix 2002:9D36:1:2::/64. (9D36:1 is the colon-hexadecimal notation for 157.54.0.1, and 2 is the interface index of 6to4 Router A's Automatic Tunneling Pseudo-Interface.)

The Router Advertisement message is sent to the addresses listed in the following table.

**Addresses in the Router Advertisement Message**

| Field | Value |
|---|---|
| IPv6 Source Address | FE80::5EFE:192.168.204.1 |
| IPv6 Destination Address | FE80::5EFE:192.168.12.9 |
| IPv4 Source Address | 192.168.204.1 |
| IPv4 Destination Address | 192.168.12.9 |

ISATAP Host A receives the Router Advertisement message and autoconfigures the address 2002:9D36:1:2:0:5EFE:192.168.12.9. The host also configures a default route (::/0) using the Automatic Tunneling Pseudo-Interface (interface index 2) with the next-hop address of FE80::5EFE:192.168.204.1 and a 2002:9D36:1:2::/64 route using the Automatic Tunneling Pseudo-Interface.

Similarly, 6to4 Router B is configured as an ISATAP router, and Site B has an appropriate A record in its DNS infrastructure so that ISATAP Host B autoconfigures the address 2002:836B:1:2:0:5EFE:192.168.141.30. (836B:1 is the colon-hexadecimal notation for 131.107.0.1.) ISATAP Host B also configures a default route (::/0) using the Automatic Tunneling Pseudo-Interface (interface index 2) with the next-hop address of FE80::5EFE:192.168.39.1 and a 2002:836B:1:2::/64 route using the Automatic Tunneling Pseudo-Interface.

ISATAP Host A can now send a packet to ISATAP B. Packet addressing comprises three parts (as shown in the previous figure during the packet's trip from ISATAP Host A to ISATAP Host B).

**Part 1: from ISATAP Host A to 6to4 Router A**

ISATAP Host A sends the IPv6 packet with the ::/0 route that uses the Automatic Tunneling Pseudo-Interface to the next-hop address of FE80::5EFE:192.168.204.1. By using this route, the next-hop address for this packet is set to the link-local ISATAP address of 6to4 Router A (FE80::5EFE:192.168.204.1).

Using the Automatic Tunneling Pseudo-Interface, the packet is tunneled using IPv4 from the address assigned to its intranet interface (192.168.12.9) to the embedded address in the ISATAP interface ID of the next-hop address (192.168.204.1). The resulting addresses are listed in the following table.

**Addresses in Part 1**

| Field | Value |
|---|---|

| | |
|---|---|
| IPv6 Source Address | 2002:9D36:1:2:0:5EFE:192.168.12.9 |
| IPv6 Destination Address | 2002:836B:1:2:0:5EFE:192.168.141.30 |
| IPv4 Source Address | 192.168.12.9 |
| IPv4 Destination Address | 192.168.204.1 |

**Part 2: from 6to4 Router A to 6to4 Router B**

6to4 Router A receives the IPv4 packet and removes the IPv4 header. 6to4 Router A then forwards the IPv6 packet with the 2002::/16 route that uses the 6to4 Tunneling Pseudo-Interface. The router, by using this route, ensures that the next-hop address for this packet is set to the destination address (2002:836B:1:2:0:5EFE:192.168.141.30).

The packet is tunneled using IPv4 and the 6to4 Tunneling Pseudo-Interface from the address assigned to its Internet interface (157.54.0.1) to the embedded address in the 6to4 prefix (836B:1) of the next-hop address (131.107.0.1). The resulting addresses are listed in the following table.

**Addresses in Part 2**

| Field | Value |
|---|---|
| IPv6 Source Address | 2002:9D36:1:2:0:5EFE:192.168.12.9 |
| IPv6 Destination Address | 2002:836B:1:2:0:5EFE:192.168.141.30 |
| IPv4 Source Address | 157.54.0.1 |
| IPv4 Destination Address | 131.107.0.1 |

**Part 3: from 6to4 Router B to ISATAP Host B**

6to4 Router B receives the IPv4 packet and removes the IPv4 header. 6to4 Router B then forwards the IPv6 packet with the 2002:836B:1:2::/64 route that uses the ISATAP Automatic Tunneling Pseudo-Interface. The router, by using this route, ensures that the next-hop address for this packet is set to the destination address (2002:836B:1:2:0:5EFE:192.168.141.30).

Because the packet is forwarded using Automatic Tunneling Pseudo-Interface, the packet is tunneled using IPv4 from the address assigned to its intranet interface (192.168.39.1) to the embedded address in the ISATAP interface ID of the next-hop IPv6 address (192.168.141.30). 6to4 Router B sets the addresses in the forwarded packet as listed in the following table.

**Addresses in Part 3**

| Field | Value |
|---|---|
| IPv6 Source Address | 2002:9D36:1:2:0:5EFE:192.168.12.9 |
| IPv6 Destination Address | 2002:836B:1:2:0:5EFE:192.168.141.30 |
| IPv4 Source Address | 192.168.39.1 |
| IPv4 Destination Address | 192.168.141.30 |

**PortProxy**

To facilitate communication between nodes or applications that cannot connect using a common Internet layer protocol (IPv4 or IPv6), the IPv6 protocol for Windows Server 2003 provides PortProxy, a component that allows the following traffic to be proxied:

- IPv4 to IPv4

  TCP traffic to an IPv4 address is proxied to TCP traffic to another IPv4 address.

- IPv4 to IPv6

  TCP traffic to an IPv4 address is proxied to TCP traffic to an IPv6 address.

- IPv6 to IPv6

TCP traffic to an IPv6 address is proxied to TCP traffic to another IPv6 address.

- IPv6 to IPv4

  TCP traffic to an IPv6 address is proxied to TCP traffic to an IPv4 address.

  The most interesting and useful proxying for IPv6/IPv4 coexistence and migration is from IPv4 to IPv6 and from IPv6 to IPv4. For coexistence and migration, PortProxy enables the following scenarios:

- An IPv4-only node can access an IPv6-only node.

  In the IPv4 DNS infrastructure of the IPv4-only node, the name of the IPv6-only node resolves to an IPv4 address assigned to an interface of the PortProxy computer. (Resolving this name might require an A record to be configured manually in DNS.) The PortProxy computer is configured to proxy IPv4 to IPv6. All TCP traffic sent by the IPv4-only node is proxied in a manner similar to Internet proxy servers: the IPv4-only node establishes a TCP connection with the PortProxy computer, and the PortProxy computer establishes a separate TCP connection with the IPv6-only node. The TCP connection data is transferred between the IPv4-only node and the IPv6-only node by the PortProxy component.

- An IPv6-only node can access an IPv4-only node.

  In the IPv6 DNS infrastructure of the IPv6-only node, the name of the IPv4-only node resolves to an IPv6 address assigned to an interface of the PortProxy computer. (Resolving this name might require AAAA records to be manually configured in DNS.) The PortProxy computer is configured to proxy IPv6 to IPv4. TCP traffic sent by the IPv6-only node to the PortProxy computer is proxied to the IPv4-only node.

- An IPv6 node can access an IPv4-only service running on a PortProxy computer.

  In the IPv6 DNS infrastructure of the IPv6-only node, the name of the IPv6/IPv4 node resolves to an IPv6 address assigned to an interface of the PortProxy computer. The PortProxy computer is configured to proxy from IPv6 to IPv4 on itself. TCP traffic sent by the IPv6 node to the PortProxy computer is proxied to an IPv4-only service or application running on the PortProxy computer.

The last scenario allows IPv6 nodes to access services that are running on a server that does not support IPv6. For example, Windows Server 2003 includes Telnet Client, which supports IPv6, and Telnet Server, which does not. However, you can allow IPv6 nodes to access the Telnet service by running PortProxy on the computer that is running Telnet Server.

To configure the PortProxy component, use the **netsh interface portproxy add|set|delete v4tov4|v4tov6|v6tov4|v6tov6** commands. For example, use the **netsh interface portproxy add v6tov4 23** command to enable IPv6 on the Telnet service (using TCP port 23) on a computer that is running Windows Server 2003.

By default, the IPv6/IPv4 node that is running Telnet Server dynamically registers both its IPv6 and IPv4 addresses in DNS. By default, a computer that is running Windows Server 2003 queries DNS for all record types, preferring IPv6 addresses to IPv4 addresses. When the Telnet client is a computer that is running Windows Server 2003, it attempts to connect using IPv6 first. With PortProxy properly configured, the first attempt to connect using an IPv6 address of the computer that is running Telnet Server should succeed without manual configuration of DNS records.

**Note**

- The PortProxy component is provided only with the IPv6 protocol for Windows Server 2003.

- The PortProxy component works only for TCP traffic and for application-layer protocols that do not embed address or port information inside the upper-layer PDU. PortProxy has no facilities to check for and change embedded address or port information in upper layer PDUs that are being proxied. For example, PortProxy cannot enable the FTP server service for IPv6 because the FTP PORT command embeds IPv4 address information inside the FTP PDU.

**IPv6 Automatic Tunneling**

RFC 2893 defines IPv6 Automatic Tunneling as the tunneling that occurs when IPv4-compatible addresses (::*w.x.y.z* where *w.x.y.z* is a public IPv4 address) are used. IPv6 Automatic Tunneling is a host-to-host tunnel between two IPv6/IPv4 hosts using IPv4-compatible addresses.

For example, when Host1 (with the public IPv4 address of 157.60.91.123 and corresponding IPv4-compatible address of ::157.60.91.123) sends traffic to Host2 (with the IPv4 address of 131.107.210.49 and corresponding IPv4-compatible address of ::131.107.210.49), the addresses in the IPv4 and IPv6 headers are as listed in the following table.

**Example IPv6 Automatic Tunneling Addresses**

| Field | Value |
|---|---|
| IPv6 Source Address | ::157.60.91.123 |
| IPv6 Destination Address | ::131.107.210.49 |
| IPv4 Source Address | 157.60.91.123 |
| IPv4 Destination Address | 131.107.210.49 |

To test connectivity, use the ping command. For example, Host A would use the following command to ping Host B by using its IPv4-compatible address:

```
ping ::131.107.210.49
```

The IPv6 protocol for Windows Server 2003 does not use IPv4-compatible addresses by default. To enable IPv4-compatible addresses, use the **netsh interface ipv6 set state v4compat=enabled** command. When the IPv6 protocol for Windows Server 2003 is enabled, communication to IPv4-compatible addresses is facilitated by a ::/96 route in the IPv6 routing table that uses the Automatic Tunneling Pseudo-Interface (interface index 2). This route indicates that all packets with destination addresses in which the first 96 bits are set to 0 are forwarded to their destination addresses with the Automatic Tunneling Pseudo-Interface. The Automatic Tunneling Pseudo-Interface uses the last 32 bits in the source and destination IPv6 addresses (corresponding to the embedded IPv4 addresses) as the source and destination IPv4 addresses for the outgoing IPv4 packet.

**Note**

- In this section, the term "IPv6 Automatic Tunneling" refers to the use of IPv4-compatible addresses. The term "automatic tunneling" is tunneling that occurs without manual configuration, independent of the type of addressing being used.

- IPv4-compatible addresses are not widely used because they are defined only for public IPv4 addresses, and their functionality has been replaced with ISATAP for IPv4 intranets and 6to4 for the IPv4 Internet. For more information, see "ISATAP" and "6to4" earlier in this section.

#### 6over4

6over4, also known as IPv4 multicast tunneling, is a technology that automatically configures tunnels between hosts, between routers, and between hosts and routers. 6over4 provides unicast and multicast connectivity between IPv6 nodes across an IPv4 intranet. RFC 2529 describes 6over4. 6over4 hosts use a valid 64-bit prefix for unicast addresses and the interface identifier ::*WWXX*:*YYZZ*, where *WWXX*:*YYZZ* is the colon-hexadecimal representation of the IPv4 address (*w.x.y.z*) that is assigned to the host. By default, 6over4 hosts automatically configure the link-local address FE80::*WWXX*:*YYZZ* on each 6over4 interface.

6over4 treats an IPv4 infrastructure as a single link with multicast capabilities. Neighbor Discovery processes (such as address resolution and router discovery) work over a physical link with multicast capabilities. To emulate a multicast-capable link, the IPv4 infrastructure must be enabled for IPv4 multicast traffic. The following figure shows a 6over4 configuration.

**A 6over4 Configuration**




To facilitate IPv6 multicast communication over an infrastructure that is enabled for IPv4 multicast traffic, RFC 2529 defines the following mapping to translate an IPv6 multicast address to an IPv4 multicast address:

239.192.[*second to last byte of IPv6 address*].[*last byte of IPv6 address*]

The following are example mappings for IPv6 multicast addresses:

- FF02::1 (link-local scope all-hosts multicast address) is mapped to 239.192.0.1.
- FF02::2 (link-local scope all-routers multicast address) is mapped to 239.192.0.2.
- FF02::1:FF28:9C5A (example solicited-node multicast address) is mapped to 239.192.156.90.

When 6over4 is enabled, the IPv4 layer uses Internet Group Membership Protocol (IGMP) messages to inform local IPv4 routers that it will receive IPv4 multicast traffic that is sent to the mapped IPv4 multicast addresses. Hosts that are enabled for 6over4 also register additional multicast MAC addresses with their network adapters that correspond to the mapped IPv4 multicast addresses. For example, for an Ethernet adapter:

- The corresponding multicast MAC address for 239.192.0.1 is 01-00-5E-40-00-01.
- The corresponding multicast MAC address for 239.192.0.2 is 01-00-5E-40-00-02.
- The corresponding multicast MAC address for 239.192.156.90 is 01-00-5E-40-9C-5A.

Because the IPv4 infrastructure acts as a multicast-capable link, hosts can use Neighbor Solicitation and Neighbor Advertisement messages to resolve each other's link-layer addresses. The 6over4 link-layer addresses are the tunnel endpoints. Hosts and routers can use Router Solicitation and Router Advertisement messages for router, prefix, and parameter discovery. The following figure shows the format for the Source and Target Link-Layer Address options that, as defined in RFC 2529, facilitate Neighbor Discovery.

**Source and Target Link-Layer Address Options for 6over4**




For example, Host1 has the public IPv4 address of 157.60.91.123 and corresponding link-local 6over4 address of FE80::9D3C:5B7B. Host2 has the public IPv4 address of 131.107.210.49 and corresponding link-local 6over4 address of FE80::836B:D231. When Host 1 sends traffic to Host 2, the addresses in the IPv4 and IPv6 headers are as listed in the following table.

**Example 6over4 Addresses**

| Field | Value |
|---|---|

| | |
|---|---|
| IPv6 Source Address | FE80::9D3C:5B7B |
| IPv6 Destination Address | FE80::836B:D231 |
| IPv4 Source Address | 157.60.91.123 |
| IPv4 Destination Address | 131.107.210.49 |

The use of 6over4 for the IPv6 protocol for Windows Server 2003 is disabled by default. To enable the use of 6over4, use the **netsh interface ipv6 set state 6over4=enabled** command. This command creates a 6over4 tunneling interface for each IPv4 address assigned to the computer. If the computer receives Router Advertisement messages over any of these interfaces (through the multicast mapping mechanism described earlier), appropriate addresses for this interface and routes that use this interface are automatically configured.

Routes and the 6over4 tunneling interface facilitate communication to 6over4 addresses. For example, the interface index for the 6over4 tunneling interface of a host is 5. (The actual interface index for the 6over4 tunneling interface varies depending on the configuration of the computer.) A Router Advertisement message from a router with the link-local 6over4 address of FE80::C0A8:1501 is received. The router is advertised as a default router, and the message contains the auto-configuration prefix FEC0:0:0:21A8::/64. The host configures a default route with the next-hop address of FE80::C0A8:1501 and a subnet route for prefix FEC0:0:0:21A8::/64 that uses interface index 5.

When packets are sent using the default or FEC0:0:0:21A8::/64 routes, the sending host uses the appropriate locally assigned 6over4 address as a source. The node also uses the last 32 bits in the source and destination IPv6 addresses (corresponding to the embedded IPv4 addresses) as the source and destination IPv4 addresses for the outgoing IPv4 packet.

A packet to a destination that matches the prefix FEC0:0:0:21A8::/64 is sent to the next-hop address of the destination using the 6over4 tunneling interface. The 6over4 tunneling interface uses address resolution for the destination address to determine the source and destination link-layer addresses (and corresponding IPv4 addresses) to use when sending the IPv4-encapsulated IPv6 packet.

A packet to a destination that matches the default route is sent to the next-hop address of FE80::C0A8:1501 using the 6over4 tunneling interface. The 6over4 tunneling interface uses address resolution for the next-hop address to determine the source and destination link-layer addresses (and corresponding IPv4 addresses) to use when sending the IPv4-encapsulated IPv6 packet.

**Note**

- 6over4 is not widely used because it requires an IPv4 multicast infrastructure.
- The difference between using ISATAP versus 6over4 on an IPv4 intranet is that 6over4 supports IPv6 multicast, and ISATAP currently does not.
- When you enable 6over4 with the **netsh interface ipv6 set state 6over4=enabled** command, it creates a 6over4 tunneling interface with a name that is based on a globally unique identifier (GUID). To create a 6over4 tunneling interface with a friendlier name, use the **netsh interface ipv6 add 6over4tunnel** command instead.

**Teredo**

Teredo, also known as IPv4 network address translation (NAT) traversal for IPv6, is an IPv6/IPv4 transition technology. Teredo provides address assignment and host-to-host automatic tunneling for unicast IPv6 connectivity when IPv6/IPv4 hosts are located behind one or multiple IPv4 NATs. To traverse IPv4 NATs, IPv6 packets are sent as IPv4-based User Datagram Protocol (UDP) messages. This section provides an overview of Teredo — including Teredo addresses and packet structures — and detailed explanations of how communication is initiated between Teredo clients, Teredo host-specific relays, and IPv6-only hosts that use the IPv4 Internet, the IPv6 Internet, Teredo servers, and Teredo relays.

Teredo is an address assignment and automatic tunneling technology that provides unicast IPv6 connectivity across the IPv4 Internet. 6to4 is a well-defined automatic tunneling technology that provides unicast IPv6 connectivity across the IPv4 Internet. However 6to4 works best when a 6to4 router exists at the edge of the site. The 6to4 router uses a public IPv4 address to construct the 6to4 prefix and acts as an IPv6 advertising and forwarding router. The 6to4 router encapsulates and decapsulates IPv6 traffic sent to and from site nodes.

6to4 relies on the configuration of a public IPv4 address and the implementation of 6to4 routing functionality in the edge device. Many small office/home office (SOHO) configurations include an IPv4 network address translator (NAT). For more

information about how network address translation works, see "Overview of Network Address Translators (NATs)" later in this section. In most NAT configurations, the device providing NAT functionality is not capable of being a 6to4 router. Even if all NAT devices could support 6to4, some configurations contain multiple levels of NATs. In these configurations, a 6to4-capable NAT device will not work because it does not have a public IPv4 address.

Teredo fills the need for 6to4 functionality in modern-day NATs and multi-layered NAT configurations by tunneling IPv6 packets between the hosts within the sites. In contrast, 6to4 uses tunneling from the edge device. Tunneling between hosts presents another issue for NATs: IPv4-encapsulated IPv6 packets are sent with the Protocol field in the IPv4 header set to 41. Most NATs translate only TCP or UDP traffic, and they must either be manually configured to translate other protocols or have installed NAT editors that handle the translation. Because Protocol 41 translation is not a common feature of NATs, IPv4-encapsulated IPv6 traffic will not traverse typical NATs. Therefore, to allow IPv6 traffic to traverse one or multiple NATs, the IPv6 packet is encapsulated as an IPv4 UDP message, containing both an IPv4 and a UDP header. UDP messages can be translated universally by NATs and can traverse multiple layers of NATs.

To summarize, Teredo is an IPv6/IPv4 transition technology that allows automatic IPv6 tunneling between hosts that are located across one or more IPv4 NATs. IPv6 traffic from Teredo hosts can traverse NATs because it is sent as IPv4 UDP messages. If a NAT supports UDP port translation, then the NAT supports Teredo. The exception is a symmetric NAT. For more information, see "Types of NATs" later in this section.

Teredo is designed as a last-resort transition technology for IPv6 connectivity. If native IPv6, 6to4, or ISATAP connectivity is present, the host does not act as a Teredo client. As more IPv4 NATs are upgraded to support 6to4 and IPv6 connectivity becomes ubiquitous, Teredo will be used less and less and eventually eliminated.

**Network Address Translators**

As defined in RFC 1631, a network address translator (NAT) is an IPv4 router that can translate the IP addresses and TCP/UDP port numbers of packets as it forwards them. For example, consider a small business network with multiple computers that connect to the Internet. Without a NAT, this business would need to obtain a public IP address for each computer on the network. With a NAT, however, the small business can use private addressing (as described in RFC 1918) and have the NAT map its private addresses to a single or to multiple public IP addresses.

NAT is a common solution for the following combination of requirements:

- The administrator wants to leverage a single connection over multiple computers, rather than connecting each one to the Internet.
- The administrator wants to use private addressing.
- The administrator wants to allow access to Internet resources without having to deploy a proxy server.

**How network address translation works**

When a private user on the small business intranet connects to an Internet resource, the user's TCP/IP protocol creates an IP packet with the following values set in the IP and TCP or UDP headers (bold text indicates the fields that are affected by the NAT):

- Destination IP Address: IP address of the Internet resource
- Source IP Address: **Private IP address**
- Destination Port: TCP or UDP port of the Internet resource
- Source Port: **Source application TCP or UDP port**

The source host or another router forwards this IP packet to the NAT, which translates the addresses of the outgoing packet as follows:

- Destination IP Address: IP address of the Internet resource
- Source IP Address: **ISP-allocated public address**
- Destination Port: TCP or UDP port of the Internet resource

- Source Port: **Remapped source application TCP or UDP port**

The NAT sends the remapped IP packet over the Internet. The responding computer sends back a response to the NAT. The response contains the following addressing information:

- Destination IP Address: ISP-allocated public address
- Source IP Address: IP address of the Internet resource
- Destination Port: Remapped source application TCP or UDP port
- Source Port: TCP or UDP port of the Internet resource

When the NAT maps and translates the addresses and forwards the packet to the intranet client, the packet contains the following addressing information:

- Destination IP Address: **Private IP address**
- Source IP Address: IP address of the Internet resource
- Destination Port: **Source application TCP or UDP port**
- Source Port: TCP or UDP port of the Internet resource

For outgoing packets, the source IP address and TCP/UDP port numbers are mapped to a public source IP address and a possibly changed TCP/UDP port number. For incoming packets, the destination IP address and TCP/UDP port numbers are mapped to the private IP address and original TCP/UDP port numbers.

For example, a small business is using the 192.168.0.0/24 private network ID for its intranet and has been allocated a single public IP address of 131.107.0.1 by its ISP. When a user with the private address 192.168.0.99 on the small business intranet connects to a Web server at the IP address 157.60.0.1, the user's TCP/IP protocol creates an IP packet with the following values set in the IP and TCP or UDP headers:

- Destination IP Address: 157.60.0.1
- Source IP Address: **192.168.0.99**
- Destination Port: 80
- Source Port: **1025**

The source host forwards this packet to the NAT device, which translates the addresses of the outgoing packet as follows:

- Destination IP Address: 157.60.0.1
- Source IP Address: **131.107.0.1**
- Destination Port: 80
- Source Port: **5000**

The NAT sends the remapped packet over the Internet. The Web server sends back a response to the NAT. The response contains the following addressing information:

- Destination IP Address: **131.107.0.1**

- Source IP Address: 157.60.0.1
- Destination Port: **5000**
- Source Port: 80

When the NAT maps and translates the addresses and forwards the packet to the intranet client, the packet contains the following addressing information:

- Destination IP Address: **192.168.0.99**
- Source IP Address: 157.60.0.1
- Destination Port: **1025**
- Source Port: 80

The following figure shows the configuration for this example.

**NAT example**




The mappings for private to public traffic are stored in a NAT translation table, which can contain two types of entries:

- Dynamic mappings

  Created when private network clients initiate communication. Dynamic mappings are removed from the NAT translation table after a specified amount of time, unless traffic that corresponds to the mapping refreshes it.

- Static mappings

  Configured manually so that communication that Internet clients initiate can be mapped to specific private network addresses and ports. Static mappings are needed when you want to make servers (for example, Web servers) or applications (for example, games) on the private network available to computers that are connected to the Internet. Static mappings are not automatically removed from the NAT translation table.

The NAT forwards traffic from the Internet to the private network only if the NAT translation table contains an appropriate mapping. In this way, the NAT provides a level of protection for computers that are connected to private network segments. However, a NAT should not be used in place of a fully featured firewall when Internet security is a concern.

**Types of NATs**

NATs fall into the following types:

- Cone NATs

  A NAT in which the NAT translation table entry stores a mapping between an internal address and port number and an external address and port number. When the NAT translation table entry is in place, inbound traffic to the external address and port number from any source address and port number is translated.

- Restricted NATs

  A NAT in which the NAT translation table entry stores a mapping between an internal address and port number and an external address and port number, for either specific source addresses or specific source address and port numbers. An inbound packet from an unknown external address or port number is silently discarded.

- Symmetric NATs

  A NAT that maps the same internal address and port number to different external addresses and ports, depending on the external destination address (for outbound traffic).

Teredo can work over only cone and restricted NATs. Teredo cannot work over symmetric NATs.

**Teredo Components**
The following figure shows the components of the Teredo infrastructure:

- Teredo clients
- Teredo servers
- Teredo relays
- Teredo host-specific relays

**Components of the Teredo infrastructure**




**Teredo client**
A Teredo client is an IPv6/IPv4 node that supports a Teredo tunneling interface through which packets are tunneled to either other Teredo clients or nodes on the IPv6 Internet (through a Teredo relay). A Teredo client communicates with a Teredo server to obtain an address prefix from which a Teredo-based IPv6 address is configured or to help initiate communication with other Teredo clients or hosts on the IPv6 Internet.

The Advanced Networking Pack for Windows XP includes a Teredo client.

**Teredo server**
A Teredo server is an IPv6/IPv4 node that is connected to both the IPv4 Internet and the IPv6 Internet, supports a Teredo tunneling interface over which packets are received. The general role of the Teredo server is to help configure addresses for Teredo clients and to facilitate the initial communication between Teredo clients and other Teredo clients or between Teredo clients and IPv6-only hosts. The Teredo server listens on UDP port 3544 for Teredo traffic.

The Advanced Networking Pack for Windows XP does not include Teredo server functionality. To facilitate communication between computers using the Advanced Networking Pack for Windows XP, Microsoft has deployed Teredo servers on the IPv4 Internet.

**Teredo relay**
A Teredo relay is an IPv6/IPv4 router that can forward packets between Teredo clients on the IPv4 Internet (using a Teredo tunneling interface) and IPv6-only hosts. In some cases, the Teredo relay interacts with a Teredo server to help it facilitate initial communication between Teredo clients and IPv6-only hosts. The Teredo relay listens on UDP port 3544 for Teredo traffic.

The Advanced Networking Pack for Windows XP does not include Teredo relay functionality. Microsoft does not plan to deploy

any Teredo relays on the IPv4 Internet. Individual Internet service providers (ISPs) could deploy their own Teredo relays. The implementation of the Teredo client in the Advanced Networking Pack for Windows XP will work with a Teredo relay when sending traffic to an IPv6-only host on the IPv6 Internet. Teredo relays are not needed to communicate with Teredo host-specific relays.

**Teredo host-specific relay**

Communication between Teredo clients and IPv6 hosts that are configured with global addresses must go through a Teredo relay. This is required for IPv6-only hosts connected to the IPv6 Internet. However, when the IPv6 host enabled for both IPv6 and IPv4 and connected to both the IPv4 Internet and IPv6 Internet, then communication should occur between the Teredo client and the IPv6 host over the IPv4 Internet, rather than having to traverse the IPv6 Internet and go through a Teredo relay.

A Teredo host-specific relay is an IPv6/IPv4 node that has an interface and connectivity to both the IPv4 Internet and the IPv6 Internet and that can communicate directly with Teredo clients over the IPv4 Internet, without the need for an intermediate Teredo relay. The connectivity to the IPv4 Internet can be through a public IPv4 address or through a private IPv4 address and a neighboring NAT. The connectivity to the IPv6 Internet can be through a direct connection to the IPv6 Internet or through an IPv6 transition technology such as 6to4, where IPv6 packets are tunneled across the IPv4 Internet. The Teredo host-specific relay listens on UDP port 3544 for Teredo traffic.

The implementation of the Teredo client in the Advanced Networking Pack for Windows XP includes Teredo host-specific relay functionality. When the Advanced Networking Pack for Windows XP is installed, the Teredo host-specific relay functionality is automatically enabled if the computer running Windows XP SP1 has a global address assigned. A global address can be assigned from a Router Advertisement message from a native IPv6 router, an ISATAP router, or a 6to4 router. If the computer running Windows XP SP1 does not have a global address, Teredo client functionality is enabled.

The Teredo host-specific relay allows Teredo clients to efficiently communicate with 6to4 hosts, IPv6 hosts with a non-6to4 global prefix, or ISATAP or 6over4 hosts within organizations that use a global prefix for their addresses, provided both hosts are using the Advanced Networking Pack for Windows XP.

**Teredo Addresses**

The following figure shows the format of Teredo addresses.

### Teredo Address Format




A Teredo address consists of the following:

- Teredo prefix (Prefix)

  The first 32 bits are for the Teredo prefix, which is the same for all Teredo addresses. The Internet Assigned Numbers Authority (IANA) has not yet defined this prefix, although the prefix 3FFE:831F::/32 is used for initial deployment.

- Teredo server IPv4 address

  The next 32 bits contain the IPv4 public address of the Teredo server that helped configure this Teredo address.

- Flags

  The next 16 bits for are reserved for Teredo flags. The only defined flag is the high-order bit known as the Cone flag. The Cone flag is set when the NAT that is connected to the Internet is a cone NAT. The determination of whether the NAT that is connected to the Internet is a cone NAT occurs during the initial configuration of a Teredo client.

- Obscured external port

  The next 16 bits store an obscured version of the external UDP port that corresponds to all Teredo traffic for this Teredo client. When the Teredo client sends its initial packet to a Teredo server, the source UDP port of the packet is mapped by the NAT to a different, external UDP port. The Teredo client maintains this port mapping so that it remains in the NAT's translation table. Therefore, all Teredo traffic for the host uses the same external, mapped UDP port. The Teredo server determines the external UDP port from the source UDP port of the incoming initial packet that the Teredo client sent, and the Teredo server sends the port information back to the Teredo client.

  The external port is obscured by XORing the external port with 0xFFFF. For example, the obscured version of external port

5000 in hexadecimal format is EC77 (5000 = 0x1388, 0x1388 XOR 0xFFFF = 0xEC77). Obscuring the external port prevents NATs from translating the external port within the payload of the packets that they are forwarding.

- Obscured external address

  The last 32 bits store an obscured version of the external IPv4 address that corresponds to all Teredo traffic for this Teredo client. Just like the external port, when the Teredo client sends its initial packet to a Teredo server, the source IP address of the packet is mapped by the NAT to a different, external (public) address. The Teredo client maintains this address mapping so that it remains in the NAT's translation table. Therefore, all Teredo traffic for the host uses the same external, mapped, public IPv4 address. The Teredo server determines the external IPv4 address from the source IPv4 address of the incoming initial packet that the Teredo client sent, and the Teredo server sends the address information back to the Teredo client.

  The external address is obscured by XORing the external address with 0xFFFFFFFF. For example, the obscured version of the public IPv4 address 131.107.0.1 in colon-hexadecimal format is 7C94:FFFE (131.107.0.1 = 0x836B0001, 0x836B0001 XOR 0xFFFFFFFF = 0x7C94FFFE). Obscuring the external address prevents NATs from translating the external address within the payload of the packets that they are forwarding.

The following figure shows an example of two Teredo clients and their addressing.

**Teredo Addressing Example**




For Teredo client A, the following are used to construct its Teredo address:

- Its external address and port for its Teredo traffic is 157.60.0.1, UDP port 4096.
- Its Teredo server is at the public IPv4 address of 206.73.118.1.
- It has determined that it is behind a cone NAT.

Therefore, using the Teredo address format of Prefix:*ServerAddr*:Flags:*ObscExtPort*:*ObscExtAddr*, the address for Teredo client A is 3FFE:831F:CE49:7601:8000:EFFF:62C3:FFFE. This address is based on the following:

- CE49:7601 is the colon-hexadecimal version of 206.73.118.1.
- 8000 is the Flags field in which the Cone flag is set to 1, indicating that Teredo client A's

NAT is a cone NAT.

- EFFF is the obscured version of 4096 (0x1000).
- 62C3:FFFE is the obscured version of 157.60.0.1.

For Teredo client B, the following are used to construct its Teredo address:

- Its external address and port for its Teredo traffic is 131.107.0.1, UDP port 8192.
- Its Teredo server is at the public IPv4 address of 206.73.118.1.
- It has determined that it is behind a restricted NAT.

Therefore, using the Teredo address format of Prefix:*ServerAddr*:Flags:*ObscExtPort*:*ObscExtAddr*, the address for Teredo client B is 3FFE:831F:CE49:7601:0:DFFF:7C94:FFFE. This address is based on the following:

- CE49:7601 is the colon-hexadecimal version of 206.73.118.1.
- 0 is the Flags field in which the Cone flag is set to 0, indicating that Teredo client B's NAT is a restricted NAT.
- DFFF is the obscured version of 8192 (0x2000).
- 7C94:FFFE is the obscured version of 131.107.0.1.

Teredo addresses are assigned to Teredo clients only. Teredo servers, relays, and host-specific relays are not assigned Teredo addresses.

**Communicating Using a Teredo Address**
Initial configuration for Teredo clients is accomplished by sending a series of Router Solicitation messages to Teredo servers. The clients use the responses to derive a Teredo address and determine whether they are behind cone, restricted, or symmetric NATs. If a Teredo client is behind a symmetric NAT, then it cannot function. You can see what type of NAT a Teredo client has discovered from the display of the **netsh interface ipv6 show teredo** command.

Based on the Router Advertisement messages that the Teredo client receives, the Teredo client constructs its Teredo address from the following:

- The first 64 bits are set to the value that the Prefix Information option of the Router Advertisement message includes. The 64-bit prefix that the Teredo server advertises consists of the Teredo prefix (32 bits) and the public IPv4 address of the Teredo server (32 bits).
- The next 16 bits are either 0x8000 (cone NAT) or 0x0 (restricted NAT).
- The next 16 bits are set to the obscured external UDP port number that a special Teredo header in the Router Advertisement message includes.
- The last 32 bits are set to the obscured external IP address that a special Teredo header in the Router Advertisement message includes.

The success of initial communication between Teredo clients that are located in different sites depends on whether those sites are using cone NATs or restricted NATs.

When both Teredo clients are located behind cone NATs, the NAT translation table entries for Teredo traffic for each Teredo client allow traffic from any source IP address or source UDP port. Therefore, a Teredo client in one site can send packets directly to a Teredo client in another site without the use of additional packets to establish NAT translation table entries.

When both Teredo clients are located behind restricted NATs, the clients must establish additional NAT translation table entries before they can exchange unicast packets. The following figure shows the initial communication process between Teredo clients that are located in different sites when both sites are using restricted NATs.

**Initial Communication Process Between Teredo Clients that Are Located in Different Sites when Both Sites Are Using Restricted NATs**




To send an initial communication packet from Teredo Client A to Teredo Client B, the following process is used:

1. Teredo Client A sends a bubble packet directly to Teredo Client B. A bubble packet contains no data and is used to create or maintain NAT mappings. Because Teredo Client B is behind a restricted NAT, Teredo traffic from an arbitrary source IPv4 address and UDP port number is not allowed unless the NAT translation table contains a source-specific entry. Assuming that no such entry exists, the restricted NAT silently discards the bubble packet. However, when the restricted NAT for Teredo Client A forwarded the bubble packet, it created a source-specific NAT translation table entry that will allow future packets sent from Teredo Client B to be forwarded to Teredo Client A.

2. Teredo Client A sends a bubble packet to Teredo Client B through Teredo Server 2 (Teredo Client B's Teredo server). The IPv4 destination address in the bubble packet is set to the IPv4 address of Teredo Server 2, which Teredo Client A determines from the third and fourth blocks of Teredo Client B's Teredo address.

3. Teredo Server 2 forwards the bubble packet to Teredo Client B. The restricted NAT for Teredo Client B forwards the packet because a source-specific mapping exists for Teredo traffic from Teredo Server 2 (established by the initial configuration of Teredo Client B).

4. Teredo Client B responds to the bubble packet received from Teredo Client A with its own bubble packet, which is sent directly to Teredo Client A. Because Teredo Client A's restricted NAT has a source-specific mapping for Teredo traffic from Teredo Client B (as established by the initial bubble packet sent from Teredo Client A in step 1), it forwards the bubble packet to Teredo Client A.

5. When Teredo Client A receives the bubble packet from Teredo Client B, Teredo Client A determines that source-specific NAT mappings exist for both NATs. Teredo Client A sends an initial communication packet directly to Teredo Client B.

This process occurs transparently to the user at Teredo Client A.

**Related Information**

For more information about TCP/IP, see "TCP/IP Technical Reference."

The standards for IPv6 are published in a series of documents called Requests for Comments (RFCs). RFCs are an evolving series of reports, proposals for protocols, and protocol standards that describe the internal workings of IPv6 and the Internet.

Although IPv6 standards are always published as RFCs, not all RFCs specify standards. RFCs are authored by individuals who voluntarily write and submit draft proposals for new protocols or specifications to the Internet Engineering Task Force (IETF) and other working groups. Submitted Internet drafts are first reviewed by a technical expert, a task force, or an RFC editor, and then the drafts are assigned a status.

If an Internet draft passes this initial review stage, the IETF circulates it to the larger Internet community for a period of time for further comment and review, and the IETF assigns an RFC number to the draft. This RFC number remains constant.

If the proposed specification changes, the IETF assigns a higher RFC number to the updated draft and circulates it. Higher RFC numbers identify more recent documents.

For more information, see the IETF RFC Database.

**Table of RFCs and Internet Drafts**

The following table lists the RFCs and Internet drafts that relate to the implementation of the IPv6 protocol in Windows Server 2003 and Windows XP.

**Table of IPv6-Related RFCs and Internet Drafts**

| RFC Number or Internet Draft | Title |
|---|---|
| 1752 | The recommendation for the IP next generation protocol |
| 1828 | IP authentication using keyed MD5 |
| 1886 | DNS extensions to support IP version 6 |
| 1993 | Transition mechanisms for IPv6 hosts and routers |
| 1981 | Path MTU discovery for IP version 6 |
| 2185 | Routing aspects of IPv6 transition |
| 2373 | IP version 6 addressing architecture |
| 2374 | An IPv6 aggregatable global unicast address format |
| 2401 | Security architecture for the Internet protocol |
| 2402 | IP authentication header |
| 2403 | The use of HMAC-MD5-96 within ESP and AH (implemented for AH only) |
| 2404 | The use of HMAC-SHA-1-96 within ESP and AH (implemented for AH only) |
| 2406 | IP encapsulating security payload (ESP) |
| 2428 | FTP extensions for IPv6 and NATs |
| 2460 | Internet protocol, version 6 (IPv6) specification |
| 2461 | Neighbor discovery for IP version 6 (IPv6) |
| 2462 | IPv6 stateless address autoconfiguration |
| 2463 | Internet Control Message Protocol (ICMPv6) for the Internet protocol version 6 (IPv6) specification |
| 2464 | Transmission of IPv6 packets over Ethernet networks |
| 2465 | Management Information Base for IP Version 6: Textual Conventions and General Group (for Windows Server 2003 only) |
| 2467 | Transmission of IPv6 packets over FDDI networks |
| 2526 | Reserved IPv6 subnet anycast addresses |
| 2529 | Transmission of IPv6 over IPv4 domains without explicit tunnels |
| 2553 | Basic socket interface extensions for IPv6 |
| 2710 | Multicast listener discovery (MLD) for IPv6 (implemented for host only) |
| 2711 | IPv6 router alert option (implemented for host only) |
| 3041 | Privacy extensions for stateless address autoconfiguration in IPv6 |

| | |
|---|---|
| 3056 | Connection of IPv6 domains via IPv4 clouds |
| 3484 | Default Address Selection for Internet Protocol version 6 (IPv6) |
| Internet draft | An extension of format for IPv6 scoped addresses |
| Internet draft | Intra-site Automatic Tunnel Addressing Protocol (ISATAP) |
| 3513 | IP version 6 addressing architecture |
| Internet draft | IP version 6 scoped address architecture |
| Internet draft | Mobility support in IPv6 (host only) |
| Internet draft | Routing of scoped addresses in the Internet protocol version 6 (IPv6) |
| Internet draft | Site prefixes in neighbor discovery |
| Internet draft | The UDP lite protocol |
| Internet draft | IPv6 Near-Unique Site-Local Addresses |
| Internet draft | Management Information Base for the User Datagram Protocol (UDP) |
| Internet draft | Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) |
| Internet draft | Default Router Preferences and More-Specific Routes |

# IPv6 Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 Tools and Settings**

**In this section**


- IPv6 Tools
- Related Information


**IPv6 Tools**

The following tools are associated with Internet Protocol version 6 (IPv6).

**IPv6-aware TCP/IP Tools**

| Tool | Description |
|---|---|
| Ftp | Use to communicate with remote computers. |
| Ipconfig | Display current TCP/IP and IPv6 network configuration values, update or release Dynamic Host Configuration Protocol (DHCP) allocated leases, and display, register, or flush Domain Name System (DNS) names. |
| Ipsec6.exe | Use to configure Internet Protocol security (IPSec) policies and security associations between two IPv6 hosts. |
| Netsh | Provides 13 sets of commands for performing network configuration tasks. Windows Server 2003 adds a new context for managing IPv6 to the **netsh** command set. |
| Netstat | Display statistics for current TCP/IP connections. Windows Server 2003 adds IPv6 parameters to the **netstat** command. |
| Pathping | Trace a path to a remote system and report packet losses at each router along the way. Windows Server 2003 adds IPv6 parameters to the **pathping** command. |
| Ping | Send Internet Control Message Protocol (ICMP) Echo Requests to verify that TCP/IP is configured correctly and that a remote TCP/IP computer is available. Windows Server 2003 adds IPv6 parameters to the **ping** command. |
| Route | Display the IP routing table, and add or delete IPv6 routes. |
| Tracert | Trace a path to a remote system. Windows Server 2003 adds IPv6 parameters to the **tracert** command. |

**FTP**

The FTP connectivity tool included with Microsoft TCP/IP can be used to communicate with remote computers.

**Category**

The IPv6-aware version of this tool is included on the Windows Server 2003 product CD.

**Version compatibility**

You can run this command on computers running Windows XP or Windows Server 2003.

You can use the FTP tool to transfer files to and from a host running an FTP server service, such as the FTP component of Internet Information Services (IIS). The File Transfer Protocol (FTP) is a protocol that defines how to transfer files from one computer to another over a TCP/IP network, such as the Internet or a company intranet.

**Note**

- The FTP service is a component of IIS. However, when you enable IIS on a server, FTP is not enabled unless you explicitly enable it. If you install IIS without FTP, you can use **Add or Remove Windows Components** in **Add or Remove Programs** in the **Control Panel** to install FTP later.

For more information about TCP/IP, see "TCP/IP Technical Reference."

**Ipconfig**
**Category**
The IPv6-aware version of this tool is included on the Windows Server 2003 product CD.

**Version compatibility**
You can run this command on computers running Windows XP or Windows Server 2003.

You can use the Ipconfig command-line tool to display the current configuration of the installed IP stack on a networked computer and to refresh DCHP and DNS settings. The **ipconfig** command is often one of the first commands you use to check the status of the connection when you experience communication problems on a TCP/IP network. Ipconfig is most useful for managing computers that obtain an IP address automatically, such as by using DHCP or Automatic Private IP Addressing (APIPA).

When called with no parameters, Ipconfig displays the Internet Protocol version 4 (IPv4) address, subnet mask, and default gateway for all adapters on a computer. If the computer has an IPv6 address, Ipconfig also displays the IPv6 address information.

**Ipsec6.exe**
**Category**
The Ipsec6.exe tool is included on the Windows Server 2003 product CD.

**Version compatibility**
You can run Ipsec6.exe on computers running Windows Server 2003 and use Ipsec6.exe to target computers running versions of Windows that support IPv6 (Windows XP Service Pack 1 [SP1] and later and Windows Server 2003).

You can use Ipsec6.exe to configure IPSec policies and security associations between two IPv6 hosts. This configuration creates an IPSec security association (SA) between two hosts on the same subnet. The SA performs authentication by using the Authentication Header (AH) and either the Message Digest 5 (MD5) or Secure Hash Algorithm 1 (SHA1) hashed message authentication code (HMAC) algorithms. This configuration provides data origin authentication and data integrity for all traffic between two IPv6 hosts. Ipsec6.exe has multiple commands, each with its own set of parameters:

**ipsec6 sp** [*Interface*]

Displays the active security policies. Alternatively, displays the active security policies for a specific interface.

**ipsec6 sa**

Displays the active security associations.

**ipsec6 |** *FileNameWithNoExtension*

Loads the security policies from FileName.spd and the security associations from FileName.sad.

**ipsec6 s |** *FileNameWithNoExtension*

Saves the current security policies to FileName.spd and the current security associations to FileName.sad. You can use this command to create files that are used to configure security policy and security associations. When there are no security policies or security associations, this command creates FileName.spd for security policies and FileName.sad for security associations. You can use these files as templates to configure the desired security policies or security associations by modifying them with a text editor.

**ipsec6 d** [{**sp** | **sa**}] [*Index*]

Deletes the security policies (using the **sp** parameter) or security associations (using the **sa** parameter) from the list of active security policies and security associations, as specified by index number. You can use **ipsec6 sp** or **ipsec6 sa** to display the index number.

**ipsec6 m** [{**on** | **off**}]

Specifies whether binding updates that are used for mobile IPv6 are protected by IPSec. This is enabled by default.

**Note**

- This implementation of IPSec for IPv6 is not recommended for use in a production environment because it relies on static keying and has no provisions for updating keys upon sequence number reuse.

- When you manually configure Security Parameters Indexes (SPIs), always use random numbers. Do not use sequential numbers for SPIs, or you will compromise the security of your IPSec for IPv6 policies.

- The IPv6 protocol for the Windows Server 2003 family supports the use of IPSec Encapsulating Security Payload (ESP) with NULL encryption.

**Netsh Commands for Interface IPv6**
**Category**
Netsh commands for Interface IPv6 are included on the Windows Server 2003 product CD.

**Version compatibility**
You can run Netsh commands for Interface IPv6 on computers running Windows Server 2003 and use Netsh commands for Interface IPv6 to target computers running versions of Windows that support IPv6 (Windows XP SP1 and later and Windows Server 2003).

The Netsh commands for Interface IPv6 provide a command-line tool that you can use to query and configure IPv6 interfaces, address, caches, and routes.

The Interface IPv6 context of **netsh** has a subcontext for 6to4, a transition technology described in the *HowIPv6 Works* section of this Technical Reference, that allows communication between IPv6/IPv4 nodes across the IPv4 Internet. You can use the commands in the **netsh interface IPv6 6to4** context to configure or display the configuration of the IPv6 Helper service on either a 6to4 host or a 6to4 router. In addition, the Interface IPv6 context of **netsh** has a subcontext for Intra-Site Automatic Tunnel Addressing Protocol (ISATAP). ISATAP is an address assignment and tunneling mechanism for communication between IPv6/IPv4 nodes within an IPv4 intranet. You can use the commands in the **nesth interface ipv6 isatap** context to configure the IPv4 address of the ISATAP router.

You can run these commands from the command prompt in the Windows Server 2003 family of products, or from the command prompt for the **netsh interface IPv6** context. For these commands to work at the Windows Server 2003 family command prompt, you must type **netsh interface ipv6** before typing commands and parameters as they appear in the following reference.

**6to4**
Specifies that the **6to4** context of **netsh interface IPv6 6to4** is used.

**Syntax**

```
6to4
```

**add6over4tunnel**
Creates a 6over4 interface using the specified IPv4 address.

**Syntax**

```
add 6over4tunnel [[interface=]String] [localaddress=]IPv4Address
[[store=]{active | persistent}]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[**localaddress=**]*IPv4Address*

Required. Specifies the IPv4 address that is encapsulated.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command creates a 6over4 interface using the IPv4 address 10.1.1.1 on the interface named “Private.”

**add 6over4tunnel “Private” 10.1.1.1**

**add address**
Adds an IPv6 address to a specified interface. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s).

For example, 2d represents two days.

**Syntax**

```
add address [[interface=]String] [address=]IPv6Address [[type=]{unicast | anycast}] [[validlifetime=]{Integer | infinite}] [[preferredlifetime=]{Integer | infinite}] [[store=]{active | persistent}]
```

**Parameters**

[[**interface=**]*String*]

Specifies an interface name or index.

[**address=**]*IPv6Address*

Required. Specifies the IPv6 address to add.

[[**type=**]{**unicast** | **anycast**}]

Specifies whether a unicast address (**unicast**) or an anycast address (**anycast**) is added. The default selection is **unicast**.

[[**validlifetime=**]{*Integer* | **infinite**}]

Specifies the lifetime over which the address is valid. The default value is **infinite**.

[[**preferredlifetime=**]{*Integer* | **infinite**}]

Specifies the lifetime over which the address is preferred. The default value is **infinite**.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**

This example command adds the IPv6 address FE80::2 to the interface named "Private."

**add address "Private" FE80::2**

**add dns**

Adds a new DNS server IP address to the statically-configured list of DNS servers for the specified interface.

**Syntax**

```
add dns [interface=]String [address=]IPAddress [[index=]Integer]
```

**Parameters**

[**interface=**]*String*

Required. Specifies, by name, which interface will have a DNS server IP address added to its list of DNS server IP addresses.

[**address=**]*IPAddress*

Required. Specifies the IPv6 address of the DNS server to add to the list.

[[**index=**]*Integer*]

Specifies the position on the statically-configured list in which to place the DNS server IP address specified in **address**. By default, the DNS server IP address is added to the end of the list.

**Remarks**

If an index is specified, the DNS server is placed in that position in the list.

**Examples**

In the first example command, a DNS server with the IPv6 address FEC0:0:0:FFFF::1 is added to the list of DNS server IP addresses for the interface named "Local Area Connection." In the second example, a DNS server with the IPv6 address FEC0:0:0:FFFF::2 is added at index 2 as the second server on the list of servers for the interface named "Local Area Connection."

**add dns "Local Area Connection" FEC0:0:0:FFFF::1**

**add dns "Local Area Connection" FEC0:0:0:FFFF::2 index=2**

**add prefixpolicy**
Adds a source and destination address selection policy for a specified prefix.

**Syntax**

```
                add prefixpolicy [prefix=]IPv6Address/Integer [precedence=]Integer
[label=]Integer [[store=]{active | persistent}]
```

**Parameters**
[**prefix=**]*IPv6Address***/***Integer*

Required. Specifies the prefix for which to add a policy in the policy table. *Integer* specifies the prefix length.

[**precedence=**]*Integer*

Required. Specifies the precedence value used for sorting destination addresses in the policy table.

[**label=**]*Integer*

Required. Specifies the label value that allows for policies that require a specific source address prefix for use with a destination address prefix.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command adds a prefix policy for prefix ::/96, with a precedence value of 3 and a label value of 4.

**add prefixpolicy ::/96 3 4**

**add route**
Adds a route for a specified prefix. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s). For example, 2d represents two days.

**Syntax**

```
                add route [prefix=]IPv6Address/Integer [[interface=]String]
[[nexthop=]IPv6Address] [[siteprefixlength=]Integer] [[metric=]Integer] [[publish=]{no
| yes | immortal}] [[validlifetime=]{Integer | infinite}] [[preferredlifetime=]{Integer
| infinite}] [[store=]{active | persistent}]
```

**Parameters**
[**prefix=**]*IPv6Address***/***Integer*

Required. Specifies the prefix for which to add a route. *Integer* specifies the prefix length.

[[**interface=**]*String*]

Specifies an interface name or index.

[[**nexthop=**]*IPv6Address*]

Specifies the gateway address, if the prefix is not on-link.

[[**siteprefixlength=**]*Integer*]

Specifies the prefix length for the entire site, if the prefix is not on-link.

[[**metric=**]*Integer*]

Specifies the route metric.

[[**publish=**]{**no** | **yes** | **immortal**}]

Specifies whether routes are advertised (**yes**), advertised with an infinite lifetime (**immortal**), or not advertised (**no**) in Route

Advertisements. The default selection is **no**.

[[**validlifetime=**]{*Integer* | **infinite**}]

Specifies the lifetime over which the route is valid. The default value is **infinite**.

[[**preferredlifetime=**]{*Integer* | **infinite**}]

Specifies the lifetime over which the route is preferred. The default value is **infinite**.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command adds a route on the interface named "Internet" with a prefix of 3FFE:: and a prefix length of 16 bits (3FFE::/16). The **nexthop** value is FE80::1.

**add route 3FFE::/16 "Internet" FE80::1**

**add v6v4tunnel**
Creates an IPv6-in-IPv4 tunnel.

**Syntax**

```
                add v6v4tunnel [[interface=]String] [localaddress=]IPv4Address
[remoteaddress=]IPv4Address [[neighbordiscovery=]{enabled | disabled}] [[store=]{active
| persistent}]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[**localaddress=**]*IPv4Address*

Required. Specifies the IPv4 address of the local tunnel endpoint.

[**remoteaddress=**]*IPv4Address*

Required. Specifies the IPv4 address of the remote tunnel endpoint.

[[**neighbordiscovery=**]{**enabled** | **disabled**}]

Specifies whether Neighbor Discovery is enabled (**enabled**) or disabled (**disabled**) on the interface. The default selection is **disabled**.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command creates an IPv6-in-IPv4 tunnel between the local address 10.0.0.1 and the remote address 192.168.1.1 on the interface "Private."

**add v6v4tunnel "Private" 10.0.0.1 192.168.1.1**

**delete address**
Modifies an IPv6 address on a specified interface.

**Syntax**

```
                delete address [[interface=]String] [address=]IPv6Address [[store=]
{active | persistent}]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[**address=**]*IPv6Address*

Required. Specifies the IPv6 address to delete.

[[**store=**]{**active** | **persistent**}]

Specifies whether the deletion lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command deletes the address FE80::2 from the interface named "Private."

**delete address "Private" FE80::2**

**delete destinationcache**
Clears the destination cache. If an interface is specified, clears the cache only on that interface. If an address is also specified, deletes only that destination cache entry.

**Syntax**

```
delete destinationcache [[interface=]String] [[address=]IPv6Address]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[[**address=**]*IPv6Address*]

Specifies the address of the destination.

**Remarks**
When no parameters are specified, all entries in the destination caches for all interfaces are deleted.

**Example**
This example command deletes the destination cache for the interface named "Private."

**delete destinationcache "Private"**

**delete dns**
Deletes statically-configured DNS server IPv6 addresses for a specific interface.

**Syntax**

```
delete dns [interface=]String [[address=]{IPv6Address | all}]
```

**Parameters**
[**interface=**]*String*

Required. Specifies the interface, by name, for which you want to remove a DNS server from the list of DNS servers.

[[**address=**]{*IPv6Address* | **all**}]

Specifies the DNS server IPv6 address to delete. If **all** is specified, all DNS server IPv6 addresses on the list for the interface are deleted.

**Examples**
In the first example command, the DNS server IPv6 address FEC0:0:0:FFFF::1 is deleted from the list of addresses for the connection named "Local Area Connection." In the second example command, all DNS server IPv6 addresses are deleted for the connection named "Local Area Connection."

**delete dns "Local Area Connection" FEC0:0:0:FFFF::1**

**delete dns "Local Area Connection" all**

**delete interface**
Deletes a specified interface from the IPv6 stack.

**Syntax**

```
delete interface [[interface=]String] [[store=]{active | persistent}]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[[**store=**]{**active** | **persistent**}]

Specifies whether the deletion lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
This example command deletes the interface named "Private."

**delete interface "Private"**

**delete neighbors**
Specifies that all entries in the neighbor cache are deleted. If an interface is specified, clears the cache only on that interface. If an address is also specified, deletes only that neighbor cache entry.

**Syntax**

```
delete neighbors [[interface=]String] [[address=]IPv6Address]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[[**address=**]*IPv6Address*]

Specifies the address of the neighbor.

**Example**
This example command removes all entries from the neighbor cache on the interface named "Private."

**delete neighbors "Private"**

**delete prefixpolicy**
Deletes the source and destination address selection policy for a specified prefix.

**Syntax**

```
delete prefixpolicy [prefix=]IPv6Address/Integer [[store=]{active |
persistent}]
```

**Parameters**
[**prefix=**]*IPv6Address***/***Integer*

Required. Specifies the prefix (*IPv6Address*) and prefix length (*Integer*) to delete from the policy table.

[[**store=**]{**active** | **persistent**}]

Specifies whether the deletion lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command deletes the prefix ::/96 from the policy table.

**delete prefixpolicy ::/96**

**delete route**
Deletes an IPv6 route.

**Syntax**

```
delete route [prefix=]IPv6Address/Integer [[interface=]String]
[[nexthop=]IPv6Address] [[store=]{active | persistent}]
```

**Parameters**

[**prefix=**]*IPv6Address***/***Integer*

Required. Specifies the prefix of the route to delete.

[[**interface=**]*String*]

Specifies an interface name or index.

[[**nexthop=**]*IPv6Address*]

Specifies the gateway address, if the prefix is not on-link.

[[**store=**]{**active** | **persistent**}]

Specifies whether the deletion lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**

This example command deletes the route with the prefix 3FFE::/16 and the gateway FE80::1 from the interface named "Internet."

**delete route 3FFE::/16 "Internet" FE80::1**

**dump**

Dumps the network adapter IPv6 configuration to the command prompt window when run within the **netsh** context. When used in a batch file or script, output can be saved in a text file.

**Syntax**

```
netsh interface ipv6 dump > [
PathAndFileName
]
```

**Parameters**

[*PathAndFileName*]

Specifies both the location to which to save the file and the name of the destination file to which the configuration is saved.

**Remarks**

After file output is obtained, you can use the **netsh exec** command to either configure another computer with the same IPv6 configuration or to restore the original configuration on the same computer.

All IPv6 configuration information is saved with the **dump** command. For example, if an ISATAP or a 6to4 configuration is defined on an interface, the **dump** command saves these settings in the text file.

**Examples**

In the first example, the command is run manually at the **netsh interface ipv6** context of a command prompt. The IPv6 configuration is displayed in the command prompt window, and can be copied and pasted into a text file. In the second example, the **dump** command is run in a batch file, and the configuration is saved to a text file named Ipv6_conf.txt at the location C:\Temp.

**dump**

**netsh interface ipv6 dump > C:\temp\ipv6_conf.txt**

**install**

Installs IPv6.

**Syntax**

```
install
```

**isatap**
Specifies that the **isatap** context of **netsh interface IPv6 isatap** is used.

**Syntax**

```
isatap
```

**Remarks**
ISATAP is used for communication between IPv6 and IPv4 nodes within an IPv4 site. It is described in the Internet draft titled "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)."

**renew**
Restarts IPv6 interfaces.

**Syntax**

```
renew [[interface=]String]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

**Example**
**renew "Private"**

**reset**
Resets the IPv6 configuration state.

**Syntax**

```
reset
```

**set address**
Modifies an IPv6 address on a specified interface. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s). For example, 2d represents two days.

**Syntax**

```
set address [[interface=]String] [address=]IPv6Address [[type=]{unicast
| anycast}] [[validlifetime=]{Integer | infinite}] [[preferredlifetime=]{Integer |
infinite}] [[store=]{active | persistent}]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[**address=**]*IPv6Address*

Required. Specifies the IPv6 address to modify.

[[**type=**]{**unicast** | **anycast**}]

Specifies whether the address is marked as a unicast address (**unicast**) or as an anycast address (**anycast**). The default selection is **unicast**.

[[**validlifetime=**]{*Integer* | **infinite**}]

Specifies the lifetime over which the address is valid. The default value is **infinite**.

[[**preferredlifetime=**]{*Integer* | **infinite**}]

Specifies the lifetime over which the address is preferred. The default value is **infinite**.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command sets the address FE80::2 on the interface named “Private” as an anycast address.

**set address "Private" FE80::2 anycast**

**set global**
Modifies global configuration parameters.

**Syntax**

```
                set global [[defaultcurhoplimit=]Integer] [neighborcachelimit=]Integer
[[routecachelimit=]Integer] [[reassemblylimit=]Integer] [[store=]{active | persistent}]
```

**Parameters**
[[**defaultcurhoplimit=**]*Integer*]

Specifies the default hop limit of packets sent.

[**neighborcachelimit=**]*Integer*

Required. Specifies the maximum number of neighbor cache entries.

[[**routecachelimit=**]*Integer*]

Specifies the maximum number of route cache entries.

[[**reassemblylimit=**]*Integer*]

Specifies the maximum size of the reassembly buffer.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command sets global parameters for all IPv6-enabled interfaces on the computer. The default hop limit is set to 32, the maximum number of neighbor cache entries is set to 100, and the maximum number of route cache entries is 100,000.

**set global 32 100 100000**

**set interface**
Modifies interface configuration parameters.

**Syntax**

```
                set interface [[interface=]String] [[forwarding=]{enabled | disabled}]
[[advertise=]{enabled | disabled}] [[mtu=]Integer] [[siteid=]Integer]
[[metric=]Integer] [[firewall=]{enabled | disabled}] [[siteprefixlength=]Integer]
[[store=]{active | persistent}]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[[**forwarding=**]{**enabled** | **disabled**}]

Specifies whether packets arriving on this interface can be forwarded to other interfaces. The default selection is **disabled**.

[[**advertise=**]{**enabled** | **disabled**}]

Specifies whether Router Advertisements are sent on this interface. The default selection is **disabled**.

[[**mtu=**]*Integer*]

Specifies the MTU of this interface. The default MTU is the natural MTU of the link.

[[**siteid=**]*Integer*]

Specifies the site scope zone identifier.

[[**metric=**]*Integer*]

Specifies the interface metric, which is added to route metrics for all routes over the interface.

[[**firewall=**]{**enabled** | **disabled**}]

Specifies whether to operate in firewall mode.

[[**siteprefixlength=**]*Integer*]

Specifies the default length of the global prefix for the entire site.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command sets the interface with the name "Private," with a **siteid** of 2 and a **metric** of 2. All other parameter values are left at the default values.

**set interface "Private" siteid=2 metric=2**

**set mobility**
Modifies mobility configuration parameters.

**Syntax**

```
                set mobility [[security=]{enabled | disabled}]
[[bindingcachelimit=]Integer] [[correspondentnode=]enabled | disabled] [[store=]{active
| persistent}]
```

**Parameters**
[[**security=**]{**enabled** | **disabled**}]

Specifies whether binding updates must be secured.

[[**bindingcachelimit=**]*Integer*]

Specifies the maximum number of binding cache entries.

[[**correspondentnode=**]**enabled** | **disabled**]

Specifies whether Correspondent Node functionality is set to **enabled** or the default of **disabled**.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
**set mobility security=disabled bindingcachelimit=1000 corr=enabled**

**set prefixpolicy**
Modifies a source and destination address selection policy for a specified prefix.

**Syntax**

```
                set prefixpolicy [prefix=]IPv6Address/Integer [precedence=]Integer
[label=]Integer [[store=]{active | persistent}]
```

**Parameters**
[**prefix=**]*IPv6Address***/***Integer*

Required. Specifies the prefix for which to add a policy in the policy table. *Integer* specifies the prefix length.

[**precedence=**]*Integer*

Required. Specifies the precedence value used for sorting destination addresses in the policy table.

[**label=**]*Integer*

Required. Specifies the label value that allows for policies that require a specific source address prefix for use with a destination address prefix.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**
This example command sets a policy in the policy table for the prefix ::/96, with a precedence value of 3 and a label value of 4.

**set prefixpolicy ::/96 3 4**

**set privacy**
Modifies parameters related to temporary address generation. If **randomtime=** is specified, **maxrandomtime=** is not used. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s). For example, 2d represents two days.

**Syntax**

```
                set privacy [[state=]{enabled | disabled}] [[maxdadattempts=]Integer]
[[maxvalidlifetime=]Integer] [[maxpreferredlifetime=]Integer]
[[regeneratetime=]Integer] [[maxrandomtime=]Integer] [[randomtime=]Integer] [[store=]
{active | persistent}]
```

**Parameters**
[[**state=**]{**enabled** | **disabled**}]

Specifies whether temporary addresses are enabled.

[[**maxdadattempts=**]*Integer*]

Specifies the number of duplicate address detection attempts made. The default value is 5.

[[**maxvalidlifetime=**]*Integer*]

Specifies the maximum lifetime over which a temporary address is valid. The default value is 7d (seven days).

[[**maxpreferredlifetime=**]*Integer*]

Specifies the maximum lifetime over which an anonymous address is preferred. The default value is 1d (one day).

[[**regeneratetime=**]*Integer*]

Specifies the duration of time that elapses when a new address is generated prior to deprecating a temporary address. The default value is 5s (five seconds).

[[**maxrandomtime=**]*Integer*]

Specifies the upper limit to use when computing a random delay at boot. The default value is 10m (ten minutes).

[[**randomtime=**]*Integer*]

Specifies a time value to use, instead of a value generated at boot.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**set route**
Modifies route parameters. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s). For example, 2d represents two days.

**Syntax**

```
set route [prefix=]IPv6Address/Integer [[interface=]String]
[[nexthop=]IPv6Address] [[siteprefixlength=]Integer] [[metric=]Integer] [publish=]{no |
yes | immortal}] [[validlifetime=]{Integer | infinite}] [[preferredlifetime=]{Integer |
infinite}] [[store=]{active | persistent}]
```

**Parameters**

[**prefix=**]*IPv6Address***/***Integer*

Required. Specifies the prefix (*IPv6Address*) and prefix length (*Integer*) of the route to modify.

[[**interface=**]*String*]

Specifies an interface name or index.

[[**nexthop=**]*IPv6Address*]

Specifies the gateway address, if the prefix is not on-link.

[[**siteprefixlength=**]*Integer*]

Specifies the prefix length for the entire site, if the prefix is not on-link.

[[**metric=**]*Integer*]

Specifies the route metric.

[[**publish=**]{**no** | **yes** | **immortal**}]

Specifies whether routes are advertised (**yes**), advertised with an infinite lifetime (**immortal**), or not advertised (**no**) in Route Advertisements. The default selection is **no**.

[[**validlifetime=**]{*Integer* | **infinite**}]

Specifies the lifetime over which the route is valid. The default value is **infinite**.

[[**preferredlifetime=**]{*Integer* | **infinite**}]

Specifies the lifetime over which the route is preferred. The default value is **infinite**.

[[**store=**]{**active** | **persistent**}]

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Example**

This example command sets a route on the interface named “Internet.” The route prefix is 3FFE::, and has a length of 16 bits. The gateway address, defined by the **nexthop** parameter, is FE80::1.

**set route 3FFE::/16 "Internet" FE80::1**

**set state**

Enables or disables IPv4 compatibility. The default value for all parameters is **disabled**.

**Syntax**

```
set state [[v4compat=](enabled | disabled | default)]
```

**Parameters**

[[**v4compat=**](**enabled** | **disabled** | **default**)]

Specifies whether IPv4 compatible interfaces are created. To both disable and delete IPv4 compatible interfaces, specify **default**. To disable IPv4 compatible interfaces without deleting them, specify **disabled**.

**Examples**

In the first example command, IPv4-compatible addresses are disabled, and any previously existing interfaces are deleted. In the second example command, IPv4-compatible addresses are enabled.

**set state default**

**set state v4compat=enabled**

**show address**
Displays all IPv6 addresses, or all addresses on a specified interface.

**Syntax**

```
show address [[interface=]String] [[level=]{normal | verbose}]
[[store=]{active | persistent}]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[[**level=**]{**normal** | **verbose**}]

Specifies whether one line per interface is displayed (**normal**) or additional information is displayed for each interface (**verbose**). When no interface is specified, the default selection is **normal**. When an interface is specified, the default selection is **verbose**.

[[**store=**]{**active** | **persistent**}]

Specifies whether active (**active**) or persistent (**persistent**) addresses are displayed. The default selection is **active**.

**show bindingcacheentries**
Displays all binding cache entries.

**Syntax**

```
show bindingcacheentries
```

**show destinationcache**
Displays destination cache entries. If an interface is specified, displays the cache only on that interface. If an address is also specified, displays only that destination cache entry.

**Syntax**

```
show destinationcache [[interface=]String] [[address=]IPv6Address]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[[**address=**]*IPv6Address*]

Specifies the address of the destination.

**show dns**
Displays the DNS server configuration for a specific interface or interfaces.

**Syntax**

```
show dns [[interface=]String]
```

**Parameters**
[[**interface=**]*String*]

Specifies the interface, by name, for which you want to display configured DNS server IPv6 addresses. If no interface is specified, servers for all interfaces are displayed.

**Example**
In this example command, DNS server IPv6 addresses configured on the “Local Area Connection” interface are displayed.

**show dns "Local Area Connection"**

**show global**
Displays global configuration parameters.

**Syntax**

```
show global [[store=]{active | persistent}]
```

**Parameters**
[[**store=**]{**active** | **persistent**}]

Specifies whether active (**active**) or persistent (**persistent**) information is displayed. The default selection is **active**.

**show interface**
Displays information about all interfaces, or about a specified interface.

**Syntax**

```
show interface [[interface=]String] [[level=]{normal | verbose}]
[[store=]{active | persistent}]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[[**level=**]{**normal** | **verbose**}]

Specifies whether one line per interface is displayed (**normal**) or additional information is displayed for each interface (**verbose**). When no interface is specified, the default selection is **normal**. When an interface is specified, the default selection is **verbose**.

[[**store=**]{**active** | **persistent**}]

Specifies whether active (**active**) or persistent (**persistent**) interfaces are displayed. The default selection is **active**.

**show joins**
Displays all IPv6 multicast addresses, or all multicast addresses on a specified interface.

**Syntax**

```
show joins [[interface=]String] [[level=]{normal | verbose}]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[[**level=**]{**normal** | **verbose**}]

Specifies whether one line per interface is displayed (**normal**) or additional information is displayed for each interface (**verbose**). When no interface is specified, the default selection is **normal**. When an interface is specified, the default selection is **verbose**.

**show mobility**
Displays mobility configuration parameters.

**Syntax**

```
show mobility [[store=]{active | persistent}]
```

**Parameters**
[[**store=**]{**active** | **persistent**}]

Specifies whether active (**active**) or persistent (**persistent**) information is displayed. The default selection is **active**.

**show neighbors**
Displays neighbor cache entries. If an interface is specified, displays only the cache on that interface. If an address is also

specified, displays only that neighbor cache entry.

**Syntax**

```
show neighbors [[interface=]String] [[address=]IPv6Address]
```

**Parameters**
[[**interface=**]*String*]

Specifies an interface name or index.

[[**address=**]*IPv6Address*]

Specifies the address of the neighbor.

**show prefixpolicy**
Displays prefix policy table entries used in source and destination address selection.

**Syntax**

```
show prefixpolicy [[store=]{active | persistent}]
```

**Parameters**
[[**store=**]{**active** | **persistent**}]

Specifies whether active (**active**) or persistent (**persistent**) information is displayed. The default selection is **active**.

**show privacy**
Displays privacy configuration parameters.

**Syntax**

```
show privacy [[store=]{active | persistent}]
```

**Parameters**
[[**store=**]{**active** | **persistent**}]

Specifies whether active (**active**) or persistent (**persistent**) information is displayed. The default selection is **active**.

**show routes**
Displays route table entries.

**Syntax**

```
show routes [[level=]{normal | verbose}] [[store=]{active |
persistent}]
```

**Parameters**
[[**level=**]{**normal** | **verbose**}]

Specifies whether only normal routes (**normal**) or routes used for loopback (**verbose**) are displayed. The default selection is **normal**.

[[store=]{active | persistent}]

Specifies whether active (**active**) or persistent (**persistent**) routes are displayed. The default selection is **active**.

**show siteprefixes**
Displays the site prefix table.

**Syntax**

```
show siteprefixes
```

**uninstall**
Uninstalls IPv6.

**Syntax**

```
uninstall
```

**Netsh Interface IPv6 6to4**

You can use the following commands in the **netsh interface IPv6 6to4** context to display the configuration of or configure the 6to4 service on either a 6to4 host or a 6to4 router.

**set interface**

Configures the 6to4 service on an interface.

**Syntax**

```
set interface [name=] InterfaceName [[routing=] {enabled | disabled | default}]
```

**Parameters**

[**name=**] *InterfaceName*

Required. Specifies the name of the interface for which you want to set 6to4 service configuration. *InterfaceName* must match the name of the interface specified in Network Connections. If *InterfaceName* contains any spaces, it must be enclosed in quotes.

[[**routing=**] {**enabled** | **disabled** | **default**}]

Specifies whether the forwarding of 6to4 packets received on the interface is enabled, disabled, or set to its default value.

**Remarks**

This command enables, disables, or sets to default the routing behavior of the 6to4 service on a specified interface.

The default setting for the **routing** parameter is **enabled**, which enables routing on private interfaces if Internet Connection Sharing (ICS) is used.

**show interface**

Displays the 6to4 service routing configuration on all interfaces, or on a specified interface.

**Syntax**

```
show interface [[name=] InterfaceName]
```

**Parameters**

[[**name=**] *InterfaceName*]

Specifies the name of the interface for which you want to display the 6to4 service configuration. *InterfaceName* must match the name of the interface specified in Network Connections. If *InterfaceName* contains any spaces, it must be enclosed in quotes.

**Remarks**

If an interface name is not specified, the 6to4 configuration for all interfaces is displayed.

**set relay**

Configures the name of the 6to4 relay router for the 6to4 service. Additionally, specifies how often the name is resolved and the state of the relay component for the 6to4 service.

**Syntax**

```
set relay [[name=] {RelayDNSName | default}] [[state=] {enabled | disabled | automatic | default}] [[interval=] {ResInterval | default}]
```

**Parameters**

[[**name=**] {*RelayDNSName* | **default**}]

Specifies either the fully qualified domain name (FQDN) of a 6to4 relay router on the IPv4 Internet (*RelayDNSName*), or sets the relay name to its default value of 6to4.ipv6.microsoft.com (**default**).

[[**state=**] {**enabled** | **disabled** | **automatic** | **default**}]

Specifies whether the state of the relay component for the 6to4 service is enabled, disabled, automatically enabled if a public IPv4 address is configured, or set to its default value.

[[**interval=**] {*ResInterval* | **default**}]

Specifies how often the name of the relay router is resolved in minutes (*ResInterval*) or sets the resolution interval to its default value of 1440 minutes (**default**).

**Remarks**
The 6to4 relay router is a router that provides an access point between the IPv4 Internet and the 6bone (the native IPv6 portion of the Internet). To access 6bone resources from a 6to4 router, the 6to4 router encapsulates 6to4 traffic with an IPv4 header and sends it to the IPv4 address of the 6to4 relay router. The 6to4 relay router removes the IPv4 header and forwards the traffic to the 6bone. For return traffic, the 6to4 relay router encapsulates IPv6 traffic and sends it to the 6to4 router at the 6to4 host's site.

The default name of the 6to4 relay router is 6to4.ipv6.microsoft.com.

The default state is **automatic**, which enables the forwarding of native IPv6 traffic to a relay router when a public IPv4 address is assigned to any interface.

The default resolution interval is 1440 minutes (once each day).

**show relay**
Displays the relay router configuration for the 6to4 service.

**Syntax**

```
show relay
```

**set routing**
Sets both the state of routing and the inclusion of site-local address prefixes in Router Advertisements that are sent by the 6to4 router.

**Syntax**

```
set routing [[routing=] {enabled | disabled | automatic | default}]
[[sitelocals=] {enabled | disabled | default}]
```

**Parameters**
[[**routing=**] {**enabled** | **disabled** | **automatic** | **default**}]

Specifies whether the state of routing on a 6to4 router is enabled, disabled, automatically enabled if ICS is enabled, or set to its default value.

[[**sitelocals=**] {**enabled** | **disabled** | **default**}]

Specifies whether the advertising of site-local address prefixes, in addition to 6to4 address prefixes, is enabled, disabled, or set to its default value.

**Remarks**
The default setting for the **routing** parameter is **automatic**, which enables routing on private interfaces when ICS is used.

The default setting for the **sitelocals** parameter is **enabled**, which enables the advertising of site-local prefixes when site-local addresses are configured on private interfaces.

**show routing**
Displays the routing configuration of the 6to4 service.

**Syntax**

```
show routing
```

**set state**
Configures the state of the 6to4 service.

**Syntax**

```
set state [[state=] {enabled | disabled | default}] [[undoonstop=] {enabled | disabled | default}] [[6over4=] {enabled | disabled | default}]
```

**Parameters**

[[**state=**] {**enabled** | **disabled** | **default**}]

Specifies whether the state of the 6to4 service is enabled, disabled, or set to its default value.

[[**undoonstop=**] {**enabled** | **disabled** | **default**}]

Specifies whether the reversal of all automatic configuration that has been performed by the 6to4 service occurs when the service stops is enabled, disabled, or set to its default value.

**Remarks**

The default setting for the **state** parameter is **enabled**, which enables the 6to4 service.

The default setting for the **undoonstop** parameter is **enabled**, which reverses all automatic configuration performed by the 6to4 service when the service is stopped.

**show state**

Displays the state of the 6to4 service.

**Syntax**

```
show state
```

**reset**

Resets the 6to4 service.

**Syntax**

```
reset
```

**Netsh Interface IPv6 ISATAP**

You can use the following commands to configure the ISATAP router.

**set router**

Specifies the ISATAP router information, including router name, state, and resolution interval.

**Syntax**

```
set router [[name=]{String | default}] [[state=]{Enabled | Disabled | Default}] [[interval]=Integer]
```

**Parameters**

[[**name=**]{*String* | **default**}]

Specifies whether the router is named with a string. If **default** is specified, the system reverts to using the default name.

[[**state=**]{**Enabled** | **Disabled** | **Default**}]

Specifies whether the ISATAP router relays packets between subnets.

[[**interval**]=*Integer*]

Specifies the router resolution interval, in minutes. The default interval is 1440 (24 hours).

**Example**

The following example command sets the router name to isatap, enables the router, and sets the resolution interval to 120 minutes:

**set router isatap enabled 120**

**show router**

Displays configuration information for the ISATAP router.

**Syntax**

```
show router
```

**Remarks**
This command displays the router name, the relay state, and the resolution interval.

**Netstat**
**Category**
The IPv6-aware version of this tool is included on the Windows Server 2003 product CD.

**Version compatibility**
You can run this command on computers running Windows XP or Windows Server 2003.

You can use the Netstat command-line tool to display active TCP/IP connections, ports on which the computer is listening, Ethernet statistics, the IP routing table, IPv4 statistics (for the IP, ICMP, TCP, and User Datagram Protocol [UDP]), and IPv6 statistics (for the IPv6, ICMPv6, TCP over IPv6, and UDP over IPv6 protocols).

**Pathping**
**Category**
The IPv6-aware version of this tool is included on the Windows Server 2003 product CD.

**Version compatibility**
You can run this command on computers running Windows XP or Windows Server 2003.

You can use Pathping, an IP packet route-tracing command-line tool that combines features of Ping and Tracert, to obtain additional information that neither of those tools provides. Specifically, you can use Pathping to discover the route to a remote host; it then pings the remote host for a period of time and collects and reports statistics. Pathping path information includes information about the intermediate routers visited on the path, the Round-Trip Time (RTT) value, and link-loss information.

**Ping**
**Category**
The IPv6-aware version of this tool is included on the Windows Server 2003 product CD.

**Version compatibility**
You can run this command on computers running Windows XP or Windows Server 2003.

You can use the Ping command-line tool as your primary tool for troubleshooting IP-level connectivity between two TCP/IP computers. Ping sends ICMP Echo Request or ICMPv6 messages to perform network diagnostics and to test reachability for a specific destination. By default, Ping queries for both IPv4 and IPv6 addresses and uses the addresses returned by the operating system.

Ping lets you specify the size of packets to use (the default is 32 bytes), how many to send, whether to record the route used, which Time-To-Live (TTL) value to use, and so on.

**Route**
**Category**
The IPv6-aware version of this tool is included on the Windows Server 2003 product CD.

**Version compatibility**
You can run this command on computers running Windows XP or Windows Server 2003.

You can use the Route command-line tool to view and modify the local IP routing table.

For two hosts to exchange IP datagrams, they must both have a route to each other, or they must use a default gateway that knows a route between the two. Typically, routers exchange information using a protocol such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). RIP Listening service is available for Microsoft Windows XP Professional, and full routing protocols are supported by Windows Server 2003 in the Routing and Remote Access service.

**Tracert**
**Category**
The IPv6-aware version of this tool is included on the Windows Server 2003 product CD.

**Version compatibility**
You can run this command on computers running Windows XP or Windows Server 2003.

You can use the Tracert command-line route-tracing tool to display the path between the sending host and a destination.

The path that Tracert displays is a list of near-side router interfaces of the routers in the path between the source host and destination. Tracert uses the IP TTL field in ICMP Echo Requests and ICMP Time Exceeded messages to determine the path from a source to a destination through an IP internetwork.

Some routers silently drop packets with expired TTLs. These routers do not appear in the Tracert display.

Tracert works by incrementing the TTL value by one for each ICMP Echo Request it sends, and then waiting for an ICMP Time Exceeded message. The TTL values of the Tracert packets start with an initial value of one; the TTL of each trace after the first is incremented by one. A packet sent out by Tracert travels one hop further on each successive trip.

**Note**

- The UNIX version of Tracert performs the same function as the Windows version, except that the IP payload is a UDP packet addressed to a (presumably) unknown destination UDP port. Intermediate routers send back ICMP Time Expired messages recording the route taken, and the final destination sends back an ICMP Destination Unreachable-Port Unreachable message.

- The UDP payload from the UNIX Tracert tool can cross routers and firewalls, whereas the ICMP Echo Request messages might not due to ICMP filtering. To avoid this problem in Windows Sever 2003, turn off packet filtering and then try using Tracert again.

**Related Information**

The following resources contain additional information that is relevant to this section:

- "TCP/IP Technical Reference."

- Internet Protocol Version 6 on the Microsoft Web site.

# Routing Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing Technologies**
Routing technologies manage the flow of data between network segments, which are also known as *subnets*. Windows Server 2003 includes several routing technologies, including demanddial routing, unicast Internet Protocol (IP) routing, IP multicasting, and network address translation (NAT) functionality, which are discussed in this subcollection.

**Routing Components**
This subcollection contains information about the following routing technologies:

- Demand-dial routing
- Unicast IP routing
- IP multicasting
- NAT

**Demand-Dial Routing**
Demand-dial routing (also known as dial-on-demand routing) is the forwarding of packets over Point-to-Point Protocol (PPP) links within a wide-area network (WAN) infrastructure of dial-up telecommunications technologies, including Public Switched Telephone Network (PSTN), Integrated Services Digital Network (ISDN), X.25, and Asynchronous Transfer Mode (ATM) over asymmetric digital subscriber line (ADSL). Demand-dial routing allows you to connect branch offices through the Internet and to implement dial-up site-to-site virtual private network (VPN) connections.

Because demand-dial routing is an implementation of Routing and Remote Access technologies, including unicast routing and multicasting, it is integrated within the Routing and Remote Access architecture. For information about the architecture of Routing and Remote Access, see Unicast IPv4 Routing Technical Reference and IPv4 Multicasting Technical Reference.

**Unicast IP Routing**
Unicast routing forwards packets from one host to another host using the unicast destination IP address.

The Unicast IP Routing subject describes Internet Protocol version 4 (IPv4) unicast routing technologies included in the Routing and Remote Access service. The IP networking protocol is fundamental to the unicast routing that occurs on any Windows-based IPv4 network. This includes networks that use hardware routers, software-based routers, or a combination of both.

In addition to supporting IP routing, Windows Server 2003 also supports Appletalk routing; however, Windows Server 2003 does not support Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) routing. Although, at some time in the future, IPv4 might be replaced by the next generation of IP, IP version 6 (IPv6), at present most unicast IP routing takes place over IPv4 networks. Windows Server 2003 supports both IPv4 and IPv6. This subject describes unicast IP routing in an IPv4 environment. For more information about Windows Server 2003 and IPv6, see "IPv6 Technical Reference."

**IP Multicasting**
Internet Protocol (IP) multicasting is the delivery of a single IP datagram to a group of hosts on a TCP/IP-based network. IP multicasting is the most practical method for point-to-multipoint data transmission because unicast IP routing requires multiple copies of the datagrams, and broadcast does not work across routers.

The IP Multicasting subject describes the IPv4 multicast forwarding functionalities included in the TCP/IP (IPv4) protocol and the Routing and Remote access service.

**NAT**
Windows Server 2003 provides network address translation (NAT) functionality as part of the Routing and Remote Access service. NAT provides a method for translating the IPv4 addresses of computers on one network into IPv4 addresses of computers on a different network. A NAT-enabled IP router deployed at the boundary where a private network meets a public network allows computers on the private network to access computers on the public network by providing this translation service.

**Routing Scenarios**

The ability for people to communicate by sending messages from one computer to another, whether the computers are located on the same network or on networks at opposite sides of the globe, has become commonplace. These communications are supported by routing methods described in the following scenarios.

**Demand-dial routing scenario**
Organizations looking to reduce the telephone costs incurred by their dial-in users may want to consider implementing demand-dial routing. Demand-dial routing reduces dial-in costs because the connection becomes active only when data is sent to the remote site. When no data has been sent over the link for a specified amount of time, the link is disconnected. By making a demand-dial connection, you can use existing dial-up telephone lines instead of leased lines for low-traffic situations. Demand-dial routing also provides organizations the ability to specify the hours that a router is allowed to dial out to make demand-dial connections. Demand-dial routing can significantly reduce connection costs.

**Unicast routing scenario**
Unicast routing over IP internetworks is a major part of the technology that makes such communication possible.

Today, the majority of internetwork traffic worldwide is over IPv4 networks, and most user-initiated traffic across IPv4 internetworks is unicast traffic. Any Windows-based network supports unicast IP routing, including networks that use only hardware routers, networks that use software-based routers such as that provided by the Windows Server 2003 Routing and Remote Access service, or a combination of hardware and software routers.

**IP multicasting scenario**
Recent trends in multimedia transmission across intranets and the Internet, such as video conferencing, live event broadcasting, and distance learning, require a bandwidth-conserving method of delivering data to multiple interested hosts. Of the delivery methods supported by IP, multicasting is the method that best addresses this need.

**NAT scenario**
The impetus initiating the development of NAT technology was the need to provide a temporary solution to the IPv4 address-depletion problem faced by the Internet community. The number of available globally unique (public) IPv4 addresses is far too few to accommodate the rapidly increasing number of computers that need access to the Internet. Although the long-term solution — the development of Internet Protocol version 6 (IPv6) addresses — exists, IPv6 is not yet widely adopted. NAT technology lets computers on any network use reusable private addresses to connect to computers with globally unique public addresses on the Internet.

NAT enables computers on small- to medium-sized organizations with private networks to access resources on the Internet or other public network. The computers on a private network are configured with reusable private IPv4 addresses; the computers on a public network are configured with globally unique IPv4 (or rarely, at present, IPv6) addresses. A typical deployment is a small office or home office (SOHO) or a medium-sized business that uses Routing and Remote Access NAT technology to enable computers on the internal corporate network to connect to resources on the Internet without having to deploy a proxy server.

# Demand Dial Routing Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Demand-Dial Routing Technical Reference**
Demand-dial routing is the forwarding of packets between networks and over Point-to-Point Protocol (PPP) links, such as analog phone lines and Integrated Services Digital Network (ISDN). Microsoft Windows Server 2003 provides extensive support for demand-dial routing, which is available as part of Routing and Remote Access.

**In this subject**


- What Is Demand Dial Routing?
- How Demand Dial Routing Works
- Demand Dial Tools and Settings

# What Is Demand Dial Routing?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Types of Demand-Dial Connections
- Components of a Demand-Dial Connection
- Demand-Dial Routing vs. Remote Access
- Components of Secure Demand-Dial Connections
- Related Information


Demand-dial routing is available as part of Routing and Remote Access that is included in Microsoft Windows Server 2003.

Demand-dial routing is the forwarding of packets over Point-to-Point Protocol (PPP) links within a wide-area network (WAN) infrastructure of dial-up telecommunications technologies, including Public Switched Telephone Network (PSTN), Integrated Services Digital Network (ISDN), X.25, and Asynchronous Transfer Mode (ATM) over asymmetric digital subscriber line (ADSL). Demand-dial routing allows you to connect branch offices through the Internet and to implement dial-up site-to-site virtual private network (VPN) connections.

For more information about dial-up telecommunications technologies, see "Dial-up Remote Access Technical Reference."

**Types of Demand-Dial Connections**
Demand-dial connections are characterized as either on-demand or persistent and as either two-way initiated or one-way initiated.

These characteristics determine the configuration of the demand-dial interface.

**On-Demand and Persistent Connections**
Demand-dial connections are either on-demand or persistent.

**On-Demand Connections**
On-demand connections are used when the cost of using the communications link is time-sensitive. For example, the charges for long distance analog phone calls are on a per-minute basis. On-demand connections make the connection when traffic is forwarded and terminate the connection after a configured amount of idle time.

Idle disconnect behavior is configured on the calling router and the answering router.

- On the calling router, the idle disconnect time is set on the **General** tab of the properties of the demand-dial interface.
- On the answering router, the idle disconnect time is set on the **Dial-In Constraints** tab of the profile properties of the remote access policy being used by the demand-dial connection.

**Persistent Connections**
Persistent connections use a dial-up WAN technology when the cost of the link is fixed and the connection can be active 24 hours a day. Examples of WAN technologies for persistent demand-dial connections include local calls that use analog phone lines, leased analog lines, and flat-rate ISDN. If a persistent connection is lost, the calling router immediately attempts to reestablish the connection.

Persistent connection behavior must be configured on the calling router and the answering router.

**Two-Way and One-Way Initiated Connections**
Demand-dial connections are either two-way initiated or one-way initiated.

**Two-Way Initiated Connections**
With two-way initiated connections, either router can be the answering router or the calling router, depending on which router

initiates the connection. Both routers must be configured to initiate and accept a demand-dial connection. You use two-way initiated connections when traffic from either router can create the demand-dial connection. Two-way initiated demand-dial connections require that:

- Both routers are configured as local area network (LAN) and WAN routers.
- User accounts are added for both routers so that the authentication credentials of the calling router are accessed and validated by the answering router.
- Demand-dial interfaces are fully configured on both routers and include the phone number of the answering router and user account credentials to authenticate the calling router.
- Static routes are configured on both routers.

For two-way initiated demand-dial routing to work properly, the user account names of the calling routers on both sides of the connection must match the name of a demand-dial interface. The following table shows an example of this configuration.

**Example of Two-Way Initiated Connection Configuration**

| Router | User Account Name | Demand-Dial Interface Name |
|---|---|---|
| Corporate office router | CorpHub | NewYorkRouter |
| Branch office router | NewYorkRouter | CorpHub |

For a description of the two-way connection process, see "How Demand Dial Routing Works."

**One-Way Initiated Connections**

With one-way initiated connections, one router is always the answering router and the other router is always the calling router. In one-way initiated connections, the routing configuration is simplified because user accounts, demand-dial interfaces, and static IP routes do not need to be fully configured on both sides of the connection. Instead of configuring a demand-dial interface and static routes on the answering router, static routes are added to the dial-in properties of the user account of the calling router.

If your answering router is in a Windows Server 2003 or Windows 2000 Server mixed-mode domain or in a Microsoft Windows NT version 4.0 domain, static routes on the user account are not available. In this case, one-way initiated connections require that:

- Both routers are configured as LAN and WAN routers.
- A user account is added for the authentication credentials of the calling router.
- A demand-dial interface is configured at the calling router with the user credentials of the user account. A demand-dial interface is configured at the answering router with the same name as the user account that is used by the calling router. Because the demand-dial interface of the answering router is not used to dial out, it is not configured with the phone number of the calling router or with valid user credentials.

For a description of the one-way connection process, see "How Demand Dial Routing Works."

**Components of a Demand-Dial Connection**

A demand-dial connection contains the following components:

- Calling router, which initiates the demand-dial connection.
- Answering router, which accepts the demand-dial connection initiated by the calling router.
- Connection medium, which is either a physical medium or a tunnel medium. For more information about connection media, see "Connection Medium" later in this section.

**Components of a Demand-Dial Connection**




**Common Components for Routers**

The following components are common to both the calling router and the answering router:

- Routing and Remote Access
- Port

**Routing and Remote Access**

Routing and Remote Access on the calling router must be configured as a LAN and WAN router and configured for IP address allocation and authentication methods. IP addresses can be allocated either by using Dynamic Host Configuration Protocol (DHCP) or a static address pool.

**Port**

A port is a logical or physical communications channel that can support a single PPP connection. Physical ports are based on equipment installed in the calling router. VPN ports are logical ports.

**Calling Router Components**

In addition to Routing and Remote Access and a port, the calling router contains the following components:

- Demand-dial interface
- Route

**Demand-Dial Interface**

A demand-dial interface configured on the calling router represents the PPP connection and contains configuration information such as the port to use, the addressing used to create the connection (such as a phone number), authentication and encryption methods, and authentication credentials.

**Route**

An IP route in the routing tables of the calling router is configured to use a demand-dial interface to forward traffic.

**Answering Router Components**

In addition to Routing and Remote Access and a port, the answering router contains the following components:

- User account
- Demand-dial interface
- Route
- Remote access policy

### Note

- Two-way initiated and one-way initiated connections require different configurations for the answering router. For more information about two-way initiated and one-way initiated connections, see "Types of Demand-Dial Connections" in this section.

**User Account**

To authenticate the calling router, the credentials of the calling router must be verified by the properties of a corresponding user account. A user account for the calling router must be either locally present or available through Windows Server 2003

security. If the answering router is configured for RADIUS authentication, then the RADIUS server must have access to the user account of the calling router.

The user account must have the following settings:

- On the **Dial-in** tab, remote access permission is set to either **Allow access** or **Control access through Remote Access Policy**.
- On the **General** or **Account** tab, **User must change password at next logon** is disabled and **Password never expires** is enabled.

For a one-way initiated connection, configure static IP routes that are added to the answering router's routing table when the demand-dial connection is made.

**Demand-Dial Interface**
For two-way initiated connections, a demand-dial interface configured on the answering router represents the PPP connection to the calling router. For a one-way initiated connection using static routes on the user account of the calling router, a demand-dial interface on the answering router does not need to be configured.

**Route**
For two-way initiated connections, an IP route in the routing tables of the calling router is configured to use a demand-dial interface to forward traffic.

For one-way initiated connections, you can configure the user account of the calling router with static IP routes.

**Remote Access Policy**
To specify connection parameters that are specific to demand-dial connections, create a separate remote access policy that uses the **Windows-Groups** attribute set to the group, which has all of the user accounts for calling routers as members. A separate remote access policy for demand-dial connections is not required.

**Connection Medium**
The PPP link is established over either a physical medium or a tunnel medium. Physical media includes PSTN, ISDN, X.25, and ATM over ADSL. Tunnel media includes Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP). For more information about PPP, see "Dial-up Remote Access Technical Reference."

**PPPoE**
PPPoE is a method of encapsulating PPP frames so that they can be sent over an Ethernet network. Using PPPoE and a broadband modem, LAN clients can gain individual authenticated access to high-speed data networks. For more information about PPPoE, see "Dial-up Remote Access Technical Reference."

**Demand-Dial Routing vs. Remote Access**
Remote access is not the same as demand-dial routing. Remote access connects a single user to a network. Demand-dial routing connects networks together. However, both remote access and demand-dial routing use PPP to negotiate and authenticate the connection and encapsulate data sent on the connection. As implemented in Routing and Remote Access for Windows Server 2003, both remote access and demand-dial connections can be enabled separately but share the same:

- Behavior of the dial-in properties of user accounts.
- Security, including authentication protocols and encryption.
- Use of remote access policies.
- Use of Windows or Remote Authentication Dial-In User Service (RADIUS) as authentication providers.
- IP address allocation configuration.
- Use of PPP features, such as Microsoft Point-to-Point Compression (MPPC), Multilink Protocol (MP), and Bandwidth Allocation Protocol (BAP).
- Troubleshooting facilities, including event logging, Windows or RADIUS authentication and accounting, logging, and tracing.

**Remote Access Clients and Routers**
Because the routing service and the remote access service coexist on a server running Routing and Remote Access, both routers and remote access clients can call the same phone number. The server running Routing and Remote Access that answers the call must be able to distinguish a remote access client from a router that is calling to create a demand-dial connection. To differentiate a remote access client from a demand-dial router, the user name in the authentication credentials sent by the calling router must exactly match the name of a demand-dial interface on the answering router. Otherwise, the incoming connection is assumed to be a remote access connection.

**Components of Secure Demand-Dial Connections**
Security for demand-dial connections uses the following security methods:

- Remote access permission
- User-level and computer-level authentication
- One-way and mutual authentication
- Encryption
- Packet filtering

For information about callback and caller ID, see "Dial-up Remote Access Technical Reference." For information about remote access account lockout, see "IAS Technical Reference."

**Remote Access Permission**
The user account of the calling router must be a valid account in the security database of the answering router or RADIUS server (if RADIUS authentication is being used); it must be granted remote access permission either explicitly in the user account (remote access permission of the dial-in properties of the user account is set to **Allow access**) or implicitly through the remote access permission setting on a remote access policy (the remote access permission of the dial-in properties of the user account is set to **Control access through Remote Access Policy** and a matching remote access policy remote access permission is set to **Grant remote access**).

**User-Level and Computer-Level Authentication**
The calling router can be authenticated at the user level and the computer level.

**User-Level Authentication**
As part of the PPP connection establishment process, the calling router's credentials must be authenticated. User-level authentication occurs through one of the following PPP authentication methods:

- Password Authentication Protocol (PAP)
- Shiva Password Authentication Protocol (SPAP)
- Challenge Handshake Authentication Protocol (CHAP)
- Microsoft Challenge Handshake Authentication Protocol version 1 (MS-CHAP)
- Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2)
- Extensible Authentication Protocol-Message Digest 5 Challenge (EAP-MD5)
- Extensible Authentication Protocol-Transport Layer Security (EAP-TLS)

For all of the authentication methods except EAP-TLS, the calling router's credentials consist of a user name, a domain, and a password. For all of the authentication methods except PAP, the password is sent over the connection in an encrypted or hashed form.

In the case of EAP-TLS, the calling router's credentials consist of a user certificate that is validated by the answering router. EAP-TLS requires a public key infrastructure (PKI) to issue and validate certificates.

For more information about PPP authentication protocols, see "Dial-up Remote Access Technical Reference."

**Computer-Level Authentication**
Computer-level authentication for demand-dial routing occurs in the following cases:

- When Internet Protocol security (IPSec) is used for a L2TP over IPSec demand-dial connection, computer-level authentication is performed through the exchange of computer certificates (also known as machine certificates) during the establishment of the IPSec security association.

- When EAP-TLS is used for user-level authentication, the answering router authenticates itself to the calling router by sending its computer certificate.

Computer certificates require a PKI to issue and validate certificates.

**One-Way and Mutual Authentication**
Authentication of the demand-dial routing connection can be one-way or mutual.

**One-Way Authentication**
With one-way authentication, the calling router authenticates itself to the answering router. PAP, SPAP, CHAP, MS-CHAP v1, and EAP-MD5 authentication methods provide for the passing of credentials from the calling router to the answering router only. With one-way authentication, the calling router does not receive any verification that the answering router is the authorized router. One-way authentication does not provide protection from unauthorized or masquerading answering routers.

**Mutual Authentication**
With mutual authentication, the calling router authenticates itself to the answering router and the answering router authenticates itself to the calling router. Both ends of the connection verify the identity of the other end of the connection. MS-CHAP v2 and EAP-TLS authentication methods provide mutual authentication.

With MS-CHAP v2, both sides of the connection send a hash of a challenge string and the user password. If successful, both ends of the connection receive verification that the other end of the connection has access to the user account's password.

With EAP-TLS, the calling router sends a user certificate that is validated by the answering router and the answering router sends a computer certificate that is validated by the calling router. EAP-TLS is the most secure form of mutual authentication; however, it requires a PKI.

Windows NT 4.0 with the Routing and Remote Access Service (RRAS) supports a feature called two-way authentication. Two-way authentication uses one-way authentication methods to perform mutual authentication. When two-way authentication is enabled on a demand-dial interface, the calling router forces the answering router to authenticate itself after the calling router authenticates itself to the answering router. A Windows 2000 calling router never requests to authenticate a Windows NT 4.0 RRAS answering router. However, a Windows 2000 answering router authenticates itself when requested by a Windows NT 4.0 RRAS calling router.

**Encryption**
There are two forms of encryption available for demand-dial connections: Microsoft Point-to-Point Encryption (MPPE) and IPSec.

**MPPE**
All PPP connections, including PPTP but not including L2TP, can use MPPE. MPPE uses the Rivest-Shamir-Adleman (RSA) RC4 stream cipher and is used only when the EAP-TLS, MS-CHAP, or MS-CHAP v2 authentication methods are used.

MPPE can use 40-bit, 56-bit, or 128-bit encryption keys — the higher the number of bits, the higher the key strength. The 40-bit key is designed for backward compatibility and international use. The 56-bit key is designed for international use and adheres to United States encryption export laws. The 128-bit key is designed for North American use. By default, the highest key strength supported by the calling router and answering router is negotiated during the connection-establishment process. If the answering router requires a higher key strength than is supported by the calling router, the connection attempt is rejected.

**IPSec**
For demand-dial connections using L2TP over IPSec, encryption is determined by the establishment of the IPSec security association (SA). The available encryption algorithms include:

- Data Encryption Standard (DES) with a 56-bit key.

- Triple DES (3DES), which uses three 56-bit keys and is designed for high-security environments.

The initial encryption keys are derived from the IPSec authentication process. For more information, see "VPN Technical Reference."

**Packet Filtering**
Routing and Remote Access supports IP packet filtering, which prevents certain types of network packets from either being sent or received. IP packet filtering specifies the type of traffic that is allowed into (input filters) and out of (output filters) the router.

**Demand-Dial Interface Packet Filtering**
You can set packet filters for each demand-dial interface and configure the filters to do one of the following:

- Pass through all traffic except packets prohibited by the filters.
- Discard all traffic except packets allowed by the filters.

IP packet filtering only occurs when the demand-dial interface is in a connected state.

Packet filtering is especially useful for an extranet, a portion of your private intranet that is accessible to business partners over demand-dial connections. For example, when a business partner makes a demand-dial connection, packet filters on the demand-dial interface can restrict the TCP/IP traffic to specific network segments or specific resources, as identified by IP address and TCP or UDP port number.

**Remote Access Policy Profile Packet Filtering**
In addition to demand-dial interface packet filtering, IP packet filters can be configured on the profile of the remote access policy configured for calling routers. While primarily designed to restrict the traffic of remote access connections, remote access policy profile-based IP packet filters can be used for demand-dial routing. Rather than configure the same IP packet filters on many demand-dial interfaces, if all the demand-dial connections share the same IP packet filters and remote access policy, then remote access policy profile packet filters allow you to configure the IP packet filters once for all of the demand-dial connections.

**Related Information**
The following Technical References contain additional information that is relevant to this section.

- VPN Technical Reference
- IAS Technical Reference
- Dial-up Remote Access Technical Reference
- IPv4 Multicasting Technical Reference
- Unicast IPv4 Routing Technical Reference
- Connection Manager Technical Reference

# How Demand Dial Routing Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Demand-Dial Routing and Routing Protocols
- Demand-Dial Routing Processes and Interactions
- Related Information


Microsoft Windows Server 2003 provides extensive support for demand-dial routing, which is available as part of Routing and Remote Access.

With demand-dial routing, packets are forwarded between networks and over Point-to-Point Protocol (PPP) links. The PPP link is represented inside the Windows Server 2003 router as a demand-dial interface. Demand-dial interfaces are used to connect networks (for example, between branch offices and the corporate office) through the Internet or through dial-up site-to-site virtual private networks (VPNs).

For more information about dial-up site-to-site VPNs, see "Dial-up Site-to-Site VPN" later in this section.

With local area network (LAN) and permanent wide area network (WAN) links, the interface that is used to forward the packets is always in an active or connected state. The packets can be forwarded without having to create the physical or logical connection. However, the demand-dial interface can be either in a connected state or a disconnected state, depending on the type of connection. The demand-dial interface must be changed to a connected state before the packet can be forwarded.

For information about the types of demand-dial connections, see "What Is Demand Dial Routing?."

**Note**

- Because demand-dial routing is an implementation of Routing and Remote Access technologies, including unicast routing and multicasting, it is integrated within the Routing and Remote Access architecture. For information about the architecture of Routing and Remote Access, see "Unicast IPv4 Routing Technical Reference" and "IPv4 Multicasting Technical Reference."

**Demand-Dial Routing and Routing Protocols**

The method for updating the routing tables of demand-dial routers depends on the type of demand-dial connection. For on-demand connections, use static routing. For persistent connections, use dynamic routing with routing protocols.

For more information about on-demand and persistent connections, see "What Is Demand Dial Routing?"

**Routing with On-Demand Connections**

Static routing is recommended for on-demand connections for the following reasons:

- The routing protocols that are provided with Routing and Remote Access—Open Shortest Path First (OSPF) and Routing Information Protocol (RIP) for Internet Protocol (IP)—can either cause the connection to be made each time routing tables are advertised or keep the connection open permanently if the advertising interval is less than the idle time-out.

- To trigger an on-demand connection, it is necessary to add static routes to the route table. RIP and OSPF do not require static routes to operate. Rather, RIP and OSPF are designed to prevent an administrator from having to add static routes to a route table. RIP and OSPF simply listen for dynamic route updates and add new routes to the route table as appropriate.

- RIP for IP and OSPF detect unsuccessful network connections and, through time-out mechanisms, report these problems to neighboring routers. This means that these protocols assume that network connections are permanent unless they are disconnected. This assumption is inaccurate for on-demand connections, which are connected only when needed and are disconnected when no traffic is being routed.

Because of the time, distance, and cost-sensitive nature of typical dial-up WAN links, the use of routing protocols for on-demand connections is not recommended.

For more information about OSPF and RIP for IP, see "Unicast IPv4 Routing Technical Reference."

**Manual Configuration of Static Routes**
Manual configuration of static routes is the adding of static routes (for TCP/IP traffic, static IP routes) that are available across the demand-dial interface through the use of Routing and Remote Access.

**Using a Default IP Route for an On-Demand Connection**
The default IP route (0.0.0.0/0) summarizes all IP destinations and becomes the route used to forward IP packets when a more specific route is not found. Although the default IP route can be used to simplify static routing for on-demand connections, be sure to use the default IP route carefully.

Using the default IP route assumes that all other destinations are in the direction of the default route. The assumption is valid if you are using a demand-dial interface to connect to the Internet. However, if you are using a demand-dial interface to connect a branch office to a corporate office in a private intranet, using the default IP route can produce undesirable results for unreachable traffic.

If a default IP route is configured to use a demand-dial interface, then it is possible that a demand-dial connection can be initiated for IP traffic that is unreachable. For example, if an organization is using the private IP network 10.0.0.0/8 for its address space and a branch office uses 10.1.1.0/24 for the hosts of the branch office, then the static routing of the branch office router can be configured as follows:

**Example 1: Configure a static route for 10.0.0.0/8 using the on-demand demand-dial interface.**
If someone in the branch office attempts to send traffic to the destination IP address of 192.168.0.1, the router at the branch office does not have a route in its routing table for the packet. The packet is discarded and an ICMP (Internet Control Message Protocol) Destination Unreachable-Host Unreachable message is sent to the sending host.

**Example 2: Configure a default route for 0.0.0.0/0 using the on-demand demand-dial interface.**
If someone in the branch office attempts to send traffic to the destination IP address of 192.168.0.1, the router at the branch office initiates the connection and forwards the packet across the demand-dial connection to the corporate router. Neither the corporate router nor another router on the corporate intranet has a route in its routing table for the packet. The packet is discarded and an ICMP Destination Unreachable-Host Unreachable message is sent to the sending host.

In Example 1, the branch office immediately determines that the IP traffic destination is unreachable and does not initiate a connection to the corporate office.

In Example 2, the IP traffic destination is determined to be unreachable only after the branch office initiates the connection and the IP traffic reaches the corporate office.

For more information about default IP routes and routing tables, see "Unicast IPv4 Routing Technical Reference."

**Auto-Static Updates**
When the number of routes is large or when routes change, manual configuration of static routes is no longer a viable administrative option. For this reason, Routing and Remote Access supports auto-static updates, the process of automatically adding static routes to the routing table.

When you configure an interface to use auto-static update mode, the router sends a request to other routers and inherits routes. The routes are saved in the routing table as auto-static routes and are kept even if the router is restarted or the interface becomes unavailable. Auto-static updates are supported in RIP for IP, but are not available for use with OSPF.

### Note

- The "auto" in "auto"-static refers to the automatic adding of the requested routes as static routes in the routing table. The sending of the request for routes is performed through an explicit action while the demand-dial interface is in a connected state. Auto-static updates are not automatically performed every time a demand-dial connection is made.

An auto-static update requests all known routes or services from the router on the other side of the connection and adds them to the routing tables of the requesting router. An auto-static update is a one-time, one-way exchange of routing information. After the routes are sent, the two routers do not periodically advertise even though the connection remains in a connected state.

To use auto static updates for IP routes, the demand-dial interface must be added to the RIP routing protocol. The default operation mode for demand-dial interfaces for RIP is Autostatic update mode. The default outgoing packet protocol is RIP version 2 multicast. The default settings are correct when initiating a connection with another Windows Server 2003 router.

**Note**

- When an auto-static update is requested, the existing routes that were obtained through a previous auto-static update are deleted before the request for routes is sent. If there is no response to the request, then the router cannot replace the routes it has deleted. This might lead to a loss of connectivity to remote networks.

Auto-static updates can be scheduled or made manually.

**Scheduled Auto-Static Updates**
You can schedule auto-static updates to occur periodically by using a combination of Netsh utility scripts and Task Scheduler. For more information, see "Demand Dial Tools and Settings."

**Manual Auto-Static Updates**
In the Routing and Remote Access snap-in, you can manually update static IP routes.

The transfer of IP routing information occurs through RIP for IP. The router on which the update was initiated sends a General RIP Request. The router on the other end of the connection responds with a RIP Response containing all of the appropriate routes in its IP routing table. The RIP routes received by the requesting router are automatically added as static routes to the requesting router's IP routing table. For more information about RIP messages, see "Unicast IPv4 Routing Technical Reference."

**Routing with Persistent Connections**
For persistent demand-dial connections, you can enable routing protocols to operate the same way as LAN interfaces to provide dynamic updates of routing tables. The following table describes the configuration of the routing protocol for a persistent demand-dial interface.

**Routing Protocol Configuration for a Persistent Demand-Dial Connection**

| Routing Protocol | Configuration Change |
|---|---|
| RIP for IP | Default operation mode is Periodic update mode and triggered updates are enabled. |
| OSPF | On the **General** tab on the properties of an OSPF interface, select the Point-to-Point network type. This is the default network type for demand-dial interfaces. For a persistent site-to-site VPN connection, on the **Advanced** tab on the properties of an OSPF interface, you might want to increase the values of the transit delay, the re-transmit interval, the hello interval, and the dead interval to account for the delay of forwarding OSPF packets across the Internet. |

Routing protocols provided by Windows Server 2003 Routing and Remote Access must run over numbered connections. Although unnumbered connections are supported by Routing and Remote Access, routing protocols cannot operate over them. For unnumbered connections, you must use static routing.

For more information about numbered and unnumbered connections, see "Address Allocation for TCP/IP-Based Connections" later in this section.

**Using MP and BAP**
Multilink Protocol (MP) and Bandwidth Allocation Protocol (BAP) can be used with demand-dial routing to automatically add physical connections to a Multilink PPP connection when network traffic increases and remove physical connections when the network traffic decreases.

**Note**

- If one of the physical links of the Multilink logical link is dropped, and BAP is not being used for the connection, then the dropped link will not be redialed.

For more information about Multilink PPP, see "Dial-up Remote Access Technical Reference."

**Demand-Dial Routing Processes and Interactions**
This section describes the following processes that take place when establishing demand-dial-routing connections.

- Connection establishment delay
- Address allocation for TCP/IP-based connections
- Two-way initiated and one-way initiated connections

- Dial-up site-to-site VPN connection

**Connection Establishment Delay**
The connection establishment process, which consists of creating a physical or a logical connection and a PPP connection, introduces a delay in the forwarding of the packet called the connection establishment delay. The length of the connection establishment delay varies, depending on the type of physical or logical connection being established. For example, the connection establishment delay for analog phone lines or X.25 dialing in to a packet assembler-disassembler (PAD) can be 10 to 20 seconds or more. For Integrated Services Digital Network (ISDN) lines, the connection establishment delay can be as short as 3 to 5 seconds.

Because the connection establishment delay affects applications that are being used across a demand-dial connection, the following factors must be considered:

- How long it takes for the application to abandon the attempt to establish network communications, also known as application time-out. If the application time-out is longer than the connection establishment delay, then the application fails to establish communications and presents an error message to the user.

- How many times the application attempts to establish network communications. On the first attempt, network traffic is forwarded to the demand-dial router, which begins the connection establishment process. Due to the limited size of a buffer in the router, additional packets that arrive during the connection establishment process to be forwarded across the demand-dial connection might overwrite the initial application connection attempt packet. If the application tries to establish communications multiple times, then there is a better chance of forwarding an application connection attempt packet after the connection is established.

Applications that have long time-outs or multiple retries might not fail while waiting for the link to become available. Interactive applications, such as Web browsers and Telnet, might fail when first connecting. However, when the user retries, the connection attempt succeeds because the first connection attempt created the demand-dial connection.

After the connection is established, packets are forwarded across the demand-dial connection. Because the costs of demand-dial connections are typically time-sensitive, after a configured amount of idle time, the demand-dial link is terminated.

**Address Allocation for TCP/IP-Based Connections**
For a TCP/IP-based demand-dial connection, TCP/IP connection settings are negotiated. An important element of a TCP/IP-based PPP connection is the allocation of an IP address.

**Numbered Connections**
When a call is initiated between demand-dial routers (either servers running Windows Server 2003 Routing and Remote Access or servers running Windows NT 4.0 Routing and Remote Access Service [RRAS]), the calling router requests an IP address from the answering router; the answering router requests an IP address from the calling router. If both routers have an IP address to give each other, the logical interface on the PPP connection for each router is assigned an IP address from the other router. This is known as a numbered connection.

**Windows NT 4.0 Connections**
Windows NT 4.0 RRAS demand-dial connections require a numbered connection. If either the calling or the answering router does not have an IP address to assign to the other router, the connection establishment process fails. For both routers to assign IP addresses to each other, you must configure both routers for the proper IP address assignment behavior through the properties of RRAS.

### Note

- With RRAS, the default configuration for IP address assignment is to obtain an IP address through DHCP and no DHCP server is available on the network. Because no IP address is allocated to the other router, the connection establishment process fails.

- This common configuration problem is resolved by the following features in Windows Server 2003 Routing and Remote Access:

- Use of Automatic Private IP Addressing (APIPA) addresses.

- Support for unnumbered connections.

**APIPA Addresses**
In Routing and Remote Access, the default IP address assignment configuration is still to obtain an IP address from a DHCP server. However, if a DHCP server is not available, APIPA addresses in the range 169.254.0.0 through 169.254.255.255 with subnet mask 255.255.0.0 are used. Therefore, the default configuration does not prevent two routers from establishing a numbered connection. For more information about APIPA, see "DHCP Technical Reference."

**Unnumbered Connections**
By default, demand-dial connections do not require a numbered connection. The calling and answering routers request an IP address from each other during the connection establishment process. However, if one of the routers rejects the request, both routers continue the connection establishment process. The logical interface on the point-to-point connection does not have an assigned IP address. This is known as an unnumbered connection.

With unnumbered connections, servers running Routing and Remote Access can connect to other routers that do not assign an IP address during the connection establishment process. For example, some Internet service providers (ISPs) prefer to use an unnumbered connection to conserve IP addresses. The Internet connection can be unnumbered because you would typically configure a default static IP route rather than run a routing protocol.

**Two-Way and One-Way Initiated Connections**
This section illustrates the complexity and elegance of demand-dial routing through examples of the two-way initiated and the one-way initiated connections.

The following figure shows the basic configuration between two offices that use demand-dial IP routing.

### Demand-Dial IPv4 Routing Between Two Offices




The Seattle office has a server running Routing and Remote Access that acts as a remote access server and demand-dial router. All computers in the Seattle office are connected to the 172.16.1.0 network (subnet mask 255.255.255.0). The Seattle router (hereafter referred to as Router 1) has a modem connected to its COM1 port; the phone number of the modem is 555-0111.

The New York office has a server running Routing and Remote Access that acts as a remote access server and demand-dial router. All computers in the New York office are connected to the 172.16.2.0 network (subnet mask 255.255.255.0). The New York router (hereafter referred to as Router 2) has a modem connected to its COM2 port; the phone number of the modem is 555-0122.

The user on the computer with the IP address of 172.16.1.10 must be able to connect to the user on the computer with the IP address of 172.16.2.20, and vice versa.

This section describes the process by which Router 1 and Router 2 are configured; it also provides examples of two-way and one way initiated connections that use those router configurations.

**Router Configuration**
The following describes the configuration of Router 1 and Router 2.

**Configuring Router 1**
The configuration for demand-dial routing on Router 1 consists of the following three steps:

1. Create a demand-dial interface.

2. Create a static route.

3. Create an account that Router 2 uses when calling Router 1.

**Create a demand-dial interface**
By using Routing and Remote Access with the Demand-Dial Interface Wizard, the administrator at Router 1 creates a demand-

dial interface called DD_NewYork with the configuration shown in the following table.

**Demand-Dial Interface Configuration for Router 1**

| Configuration Information | Description |
|---|---|
| Equipment | Modem on COM1 |
| Phone number | 555-0122 |
| Protocols | TCP/IP |
| Authentication credentials | DD_Seattle with a password |

**Create a static route**

By using Routing and Remote Access, the Demand-Dial Interface Wizard will already have the name and destination of the static route with the configuration shown in the following table.

**Static Route Configuration for Router 1**

| Static Route Information | Description |
|---|---|
| Interface | DD_NewYork |
| Destination | 172.16.2.0 |
| Network mask | 255.255.255.0 |
| Metric | 1 |

**Note**

- Because the demand-dial connection is a point-to-point connection, the gateway IP address cannot be configured.

**Create a user account with dial-in permissions**

By using Active Directory Users and Computers or Local Users and Groups, the administrator at Router 1 creates a user account with the settings in the following table.

**User Account with Dial-In Permissions for Router 1**

| Account Information | Description |
|---|---|
| Account name | DD_NewYork with a password |
| Account settings | Clear the **User must change password at next logon** check box and select the **Password never expires** check box |

The DD_NewYork account is granted dial-in permissions either through the dial-in properties of the user account or through remote access policies. For more information about remote access policies, see "Dial-up Remote Access Technical Reference" and "IAS Technical Reference."

**Configuring Router 2**

The configuration for demand-dial routing on Router 2 consists of the following three steps:

1. Create a demand-dial interface.
2. Create a static route.
3. Create an account that Router 2 uses when calling Router 1.

**Create a demand-dial interface**

By using Routing and Remote Access with the Demand-Dial Interface Wizard, the administrator at Router 2 creates a demand-dial interface called DD_Seattle with the configuration shown in the following table.

**Demand-Dial Interface Configuration for Router 2**

| Configuration Information | Description |
|---|---|
| Equipment | Modem on COM2 |
| Phone number | 555-0111 |
| Protocols | TCP/IP |
| Authentication credentials | DD_NewYork with a password |

**Create a static route**

By using Routing and Remote Access, the Demand-Dial Interface Wizard will already have the name and destination of the static route with the configuration shown in the following table.

**Static Route Configuration for Router 2**

| Static Route Information | Description |
|---|---|
| Interface | DD_Seattle |
| Destination | 172.16.1.0 |
| Network mask | 255.255.255.0 |
| Metric | 1 |

**Note**

- Because the demand-dial connection is a point-to-point connection, the gateway IP address cannot be configured.

**Create a user account with dial-in permissions**

By using Active Directory Users and Computers or Local Users and Groups, the administrator at Router 2 creates a user account with the settings in the following table.

**User Account with Dial-In Permissions for Router 2**

| Account Information | Description |
|---|---|
| Account name | DD_Seattle with a password |
| Account settings | Clear the **User must change password at next logon** check box and select the **Password never expires** check box |

The DD_Seattle account is granted dial-in permissions either through the dial-in properties of the user account or through remote access policies. For more information about remote access policies, see "Dial-up Remote Access Technical Reference" and "IAS Technical Reference."

**Resulting Configuration**

The following figure shows the demand-dial routing configuration in terms of the demand-dial interfaces, static routes, and user accounts for the routers for the Seattle and New York offices.

**Resulting Configuration for the Seattle and New York Offices**




**Example of a Two-Way Initiated Demand-Dial Connection**

For two-way initiated demand-dial routing to work properly, authentication credentials (such as user account name) for the demand-dial interface must be created on both servers running Routing and Remote Access. This example shows a proper configuration and is summarized in the following table.

**Resulting Demand-Dial Routing Configuration**

| Router | Demand-Dial Interface Name | User Account Name |
|---|---|---|
| Router 1 | DD_NewYork | DD_Seattle |
| Router 2 | DD_Seattle | DD_NewYork |

The following is an example of the process by which a two-way initiated connection takes place. For more information about two-way initiated connections, see "What Is Demand Dial Routing?."

When the user at 172.16.1.10 tries to connect to a resource at 172.16.2.20, the following events occur:

1. Packets from 172.16.1.10 destined for 172.16.2.20 are forwarded to Router 1.
2. Router 1 receives the packet from 172.16.1.10 and checks its routing table. It finds a route to 172.16.2.20 that uses the DD_NewYork interface.
3. Router 1 checks the state of the DD_NewYork interface and finds it is in a disconnected state.
4. Router 1 retrieves the configuration of the DD_NewYork demand-dial interface.
5. Based on the DD_NewYork configuration, Router 1 uses the modem on COM1 to dial the number 555-0122.
6. Router 2 answers the incoming call.
7. Router 2 requests authentication credentials from the incoming caller.
8. Router 1 sends the user name, DD_Seattle, with its associated password.
9. Upon receiving the authentication credentials, Router 2 checks the user name and password against the security features of Windows Server 2003 and verifies that Router 1 has dial-in permission through the dial-in properties of the DD_Seattle user account and the configured remote access policies.
10. Router 2 must now determine whether the incoming caller is a dial-up networking client or a router that is creating a demand-dial connection. Router 2 looks in its list of demand-dial interfaces to find one that matches the user name sent by Router 1 as part of the authentication credentials. Router 2 finds a demand-dial interface, DD_Seattle, that matches the user name credential.
11. Router 2 changes the state of the DD_Seattle demand-dial interface to a connected state.

12. Router 1 forwards the packet from the computer at 172.16.1.10 across the demand-dial connection to Router 2.

13. Router 2 receives the packet and forwards it to the computer at 172.16.2.20.

14. The response to the connection request by the computer at 172.16.1.10 is forwarded to Router 2 by the computer at 172.16.2.20.

15. Router 2 receives the packet destined for 172.16.1.10 and checks its routing table. It finds a route to 172.16.1.10 by using the DD_Seattle interface.

16. Router 2 checks the state of the DD_Seattle interface and finds it is in a connected state.

17. Router 2 forwards the packet to Router 1.

18. Router 1 forwards the packet to the computer at 172.16.1.10.

**Router vs. Remote Access Client**

If the user name credential does not match the name of an appropriate demand-dial interface, the calling entity is identified as a remote access client, which can result in routing problems.

For example, if Router 1 uses DialUpRouter1 as its user name credential, then Router 1 is identified as a remote access client rather than a router (assuming DialUpRouter1 is a valid account with dial-in permissions). Packets are routed from the user at 172.16.1.10 to the user at 172.16.2.20, as described earlier.

However, response packets from 172.16.2.20 to 172.16.1.10 are forwarded to Router 2 which, upon inspecting its routing table, determines that the interface to use is DD_Seattle. DD_Seattle is in a disconnected state. Based on the configuration for DD_Seattle, COM2 should be used. However, COM2 is currently being used for a remote access client (Router 2). Router 2 then tries to locate a modem that is not being used. If found, Router 2 dials Router 1 and forwards the packet after the connection has been established. If another modem cannot be found, the response packets from 172.16.2.20 to 172.16.1.10 are dropped.

**Example of a One-Way Initiated Demand-Dial Connection**

The following is an example of the process by which a one-way initiated connection takes place. For more information about one-way initiated connections, see "What Is Demand Dial Routing?"

When the user at 172.16.1.10 tries to connect to a resource at 172.16.2.20, the following events occur:

1. Packets from 172.16.1.10 destined for 172.16.2.20 are forwarded to Router 1.

2. Router 1 receives the packet from 172.16.1.10 and checks its routing table. It finds a route to 172.16.2.20 by using the DD_NewYork interface.

3. Router 1 checks the state of the DD_NewYork interface and finds it is in a disconnected state.

4. Router 1 retrieves the configuration of the DD_NewYork demand-dial interface.

5. Based on the DD_NewYork configuration, Router 1 uses the modem on COM1 to dial the number 555-0122.

6. Router 2 answers the incoming call.

7. Router 2 requests authentication credentials from the incoming caller.

8. Router 1 sends the user name, DD_Seattle, with its associated password.

9. Upon receiving the authentication credentials, Router 2 checks the user name and password against the security features of Windows Server 2003 and verifies that Router 1 has dial-in permission through the dial-in properties of the DD_Seattle user account and the configured remote access policies.

10. Router 2 retrieves the static route (172.16.1.0 with the subnet mask of 255.255.255.0) that is configured on the

DD_Seattle user account and creates a corresponding static route in its routing table. If Router 2 is configured with routing protocols, Router 2 uses routing protocols to communicate with neighboring routers so that the route to the Seattle network is propagated to all of the routers in the New York office.

11. Router 2 must now determine whether the incoming caller is a dial-up networking client or a router creating a demand-dial connection. Router 2 looks in its list of demand-dial interfaces and does not find one called DD_Seattle. Therefore, Router 2 considers the connection to the Seattle office to be a remote access connection.

12. Router 1 forwards the packet from the computer at 172.16.1.10 across the demand-dial connection to Router 2.

13. Router 2 receives the packet and forwards it to the computer at 172.16.2.20.

14. The response to the connection request by the computer at 172.16.1.10 is forwarded to Router 2 by the computer at 172.16.2.20.

15. Router 2 receives the packet destined for 172.16.1.10 and checks its routing table. A route to 172.16.1.10 is found by using the connection to Router 1.

16. Router 2 forwards the packet to Router 1.

17. Router 1 forwards the packet to the computer at 172.16.1.10.

**Note**

- When the connection is made, the static routes on the user account of the calling router are added to the routing table of the answering router. If routing protocols are used to propagate the new static route, then there is a delay between the time the connection is made and the time when all of the routers on the intranet of the answering router are aware of the new route. Therefore, hosts on the intranet of the calling router might experience a delay between the time that the connection is made and the time when they begin to receive traffic from hosts on the intranet of the answering router.

**Dial-up Site-to-Site VPN**

A site-to-site VPN connection (also known as a router-to-router VPN connection) is typically used to connect remote offices together when both routers are connected to the Internet through permanent WAN links, such as T1 or Frame Relay. In this configuration, the VPN connection is always available. However, when a permanent WAN link is not possible or practical, you can configure a dial-up site-to-site VPN connection.

The following figure shows a dial-up site-to-site VPN connection.

**Dial-up Site-to-Site VPN Connection**




A dial-up site-to-site VPN connection consists of two demand-dial interfaces and two static routes that are configured on the VPN client (the calling router) as follows. In this discussion, the VPN client is the branch office router.

- A demand-dial interface to dial in to a local ISP.

- A demand-dial interface for the site-to-site VPN connection.

- A static host route to dynamically connect to the Internet.

- A static route to reach the locations of the corporate office.

A dial-up site-to-site VPN connection is automatically established when you route traffic. For example, in a branch office configuration, upon receipt of a packet to be routed to the corporate office, the branch office router uses a dial-up link to connect to a local ISP and then creates a site-to-site VPN connection with the corporate office router located on the Internet.

**Note**

- This discussion assumes that the corporate office router (the answering router) is connected to the Internet by using a permanent WAN link. It is possible to have both routers connected to the Internet by using a dial-up WAN link. However, this is feasible only if the ISP supports demand-dialing routing to customers; the ISP calls the customer router when an IP datagram is to be delivered to the customer. Demand-dial routing to customers is not widely supported by ISPs.

**Router Configuration**

The branch office router and the corporate office router must be configured for the site-to-site VPN connection.

A dial-up site-to-site VPN connection at the branch office router is configured as follows:

1. A demand-dial interface for the Internet connection is configured for the appropriate equipment (a modem or ISDN device), the phone number of the local ISP, and the user name and password to gain Internet access.

2. A demand-dial interface for the VPN connection with the corporate office router is configured for a VPN port (a Point-to-Point Tunneling Protocol [PPTP] or Layer Two Tunneling Protocol [L2TP] port), the IP address of the interface on the Internet for the corporate office router, and a user name and password that can be verified by the VPN server. The user name must match the name of a demand-dial interface on the corporate office router.

3. A static host route to the corporate office router is created that uses the ISP demand-dial interface.

4. A static route (or routes) for the IP network IDs of the corporate intranet is created that uses the VPN demand-dial interface.

The corporate office router is configured as follows:

1. A demand-dial interface for the VPN connection with the branch office router is configured for a VPN port (a PPTP or L2TP port). The demand-dial interface must have the same name as the user name in the authentication credential that is used by the branch office router to create the VPN connection.

2. A static route (or routes) for the IP network IDs of the branch office is created that uses the VPN demand-dial interface.

**Dial-up Site-to-Site VPN Connection**

The dial-up site-to-site VPN connection is automatically initiated by the branch office router through the following process:

1. Packets sent to a corporate hub network location from a computer in the branch office are forwarded to the branch office router.

2. The branch office router checks its routing table and finds a route to the corporate intranet network ID, which uses the VPN demand-dial interface.

3. The branch office router checks the state of the VPN demand-dial interface and finds that the demand-dial interface is in a disconnected state.

4. The branch office router retrieves the configuration of the VPN demand-dial interface.

5. Based on the VPN demand-dial interface configuration, the branch office router attempts to initialize a site-to-site VPN connection at the IP address of the corporate office router on the Internet.

6. To establish a VPN, either a Transmission Control Protocol (TCP) (by using PPTP) or User Datagram Protocol (UDP) (by using L2TP/Internet Protocol security [IPSec]) packet must be sent to the corporate office router that acts as the VPN

server. The VPN establishment packet is created.

7. To forward the VPN establishment packet to the corporate office router, the branch office router checks its routing table and finds the host route that is using the ISP demand-dial interface.

8. The branch office router checks the state of the ISP demand-dial interface and finds that the demand-dial interface is in a disconnected state.

9. The branch office router retrieves the configuration of the ISP demand-dial interface.

10. Based on the ISP demand-dial interface configuration, the branch office router uses its modem to dial and establish a connection with its local ISP.

11. Immediately after the ISP connection is made, the VPN establishment packet is sent to the corporate office router.

12. A site-to-site VPN connection is negotiated between the branch office router and the corporate office router. As part of the negotiation, the branch office router sends authentication credentials that are verified by the corporate office router.

13. The corporate office router checks its demand-dial interfaces and finds one that matches the user name sent during authentication and changes the interface to a connected state.

14. The branch office router forwards the routed packet across the VPN and the corporate office router forwards the packet to the appropriate intranet location.

15. When the intranet location responds to the packet sent to it by the user in the branch office, the packet is forwarded to the corporate office router.

16. The corporate office router checks its routing table and finds a route to the branch office network that uses the VPN demand-dial interface.

17. The corporate office router checks the state of the VPN demand-dial interface and finds that the interface it is in a connected state.

18. The response packet is forwarded across the Internet by using the VPN connection.

19. The response packet is received by the branch office router and is forwarded to the original user.

**Related Information**

The following resources contain additional information that is relevant to this section.

- VPN Technical Reference
- IAS Technical Reference
- Unicast IPv4 Routing Technical Reference
- IPv4 Multicasting Technical Reference
- Dial-up Remote Access Technical Reference
- DHCP Technical Reference
- Connection Manager Technical Reference
- IETFRFC Database

# Demand Dial Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Demand-Dial Routing Tools
- Demand-Dial Routing Registry Entries


The "Demand-Dial Routing Tools" section describes tools that are useful for auditing, customizing, and troubleshooting demand-dial routing connections.

The "Demand-Dial Routing Registry Entries" section describes registry entries that:

- Enable negotiation of the Point-to-Point Protocol (PPP) connection with the endpoint.
- Indicate the router managers installed in the system.
- Indicate the types of router interfaces in the system.

**Demand-Dial Routing Tools**
The following tools are associated with demand-dial routing.

**Iprtrmgr.dll**
**Category**
Iprtrmgr.dll, also known as the IP router manager, is a dynamic-link library (DLL) available as part of the operating system.

**Version compatibility**
Iprtrmgr.dll is available on all platforms that support Routing and Remote Access, including Windows Server 2003, Windows 2000 Server, and Windows XP.

Iprtrmgr.dll manages the routing protocols. Acting as liaison between the routing protocols and the dynamic interface manager (DIM), Iprtrmgr.dll does the following:

- Obtains configuration information from the DIM.
- Communicates IP packet-filtering configuration to the IP filtering driver.
- Communicates IP routing configuration information to the IP forwarder in the TCP/IP protocol.
- Maintains an interface database of all IP routing interfaces.
- Loads and communicates configuration information to IP routing protocols that are supplied with Windows Server 2003, such as Routing Information Protocol (RIP) for IP and Open Shortest Path First (OSPF).
- Initiates demand-dial connections for routing protocols by communicating with the DIM.

In addition, Iprtrmgr.dll updates the route information for the Route Table Manager (RTM) and keeps the RTM synchronized with the stack forwarding table.

For more information about Iprtrmgr.dll, see "Windows Support Tools" in the Tools and Settings Collection.

**Mprdim.dll**
**Category**
Mprdim.dll, also known as the DIM, is a dynamic-link library available as part of the operating system.

**Version compatibility**

Mprdim.dll is available in all platforms that support Routing and Remote Access, including Windows Server 2003, Windows 2000 Server, and Windows XP.

The primary responsibility of Mprdim.dll is the administration of the Routing and Remote Access server computer and to relay information to other routing components, such as the IP router manager.

Mprdim.dll does the following:

- Supports a remote procedure call (RPC) interface for Simple Network Management Protocol (SNMP)-based management functions used by management utilities, such as the Routing and Remote Access snap-in.
- Loads configuration information from the Windows registry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess.
- Reads the router phone book file: *systemroot*\system32\RAS\Router.pbk.
- Communicates with the Connection Manager for demand-dial connections.
- Communicates configuration information to the router managers (for example, the IP router manager).
- Passes information about static routes and filters to the appropriate router manager.
- Manages all routing interfaces, including local area network (LAN), persistent demand-dial, and IP-in-IP interfaces.
- Sends information to the Active Directory Service if the computer is configured as a member of a domain.

For more information about Mprdim.dll, see "Windows Support Tools" in the Tools and Settings Collection.

**Rtm.dll**
**Category**
Rtm.dll, also known as the RTM, is a dynamic-link library available as part of the operating system.

**Version compatibility**
Rtm.dll is available in all platforms that support Routing and Remote Access, including Windows Server 2003, Windows 2000 Server, and Windows XP.

Rtm.dll is the central repository of routing information for all routing protocols and components, such as router managers.

Rtm.dll does the following:

- Maintains a user-mode route table for all routes from all possible route sources.
- Provides application programming interfaces (APIs) to accomplish the following tasks:
    - Adding, deleting, and enumerating routes that are used by the routing protocols.
    - Querying and modifying routing information (other than adding, deleting, and enumerating routes) stored in the RTM.
    - Registering routes.
    - Registering to receive notifications of changes to RTM information.
- Ages learned routes. The age of a route is the time, in seconds, since the route was created or last updated.
- Provides routing information to routing protocols, management programs, and monitoring programs.
- Communicates only the best routes to the appropriate IP forwarder driver. For IP routes, the best route is the route with the lowest cost based on specified criteria, such as the number of networks crossed, the type of network crossed (for example, private or public), or a monetary or bandwidth limit. The best routes become the routes in the IP forwarding

table.

For more information about Rtm.dll, see "Windows Support Tools" in the Tools and Settings Collection.

**Netsh**
**Category**
Netsh is a command-line scripting tool used to configure multiple networking components and is available as part of the operating system.

**Version compatibility**
Netsh is available in Windows Server 2003, Windows 2000 Server, and all versions of Windows that include TCP/IP.

The Netsh command-line syntax is as follows:

```
netsh [-a AliasFile] [-c Context] [-r RemoteMachine]
       [Command | -f ScriptFile]
```

Netsh enables you to display or modify the network configuration of a local or remote computer. You can schedule auto-static updates to occur periodically by using a combination of Netsh utility scripts and Task Scheduler. The transfer of routing information occurs by using RIP for IP. To perform a scheduled automated auto-static update, use the following Netsh commands:

```
netsh interface set interface name=DemandDialInterfaceName connect=CONNECTED
netsh routing ip rip update DemandDialInterfaceName
netsh interface set interface name=DemandDialInterfaceName connect=DISCONNECTED
```

For example, to automatically update the RIP-for-IP routes by using a demand-dial connection called CorpHub, you type the following Netsh commands:

```
netsh interface set interface name=CorpHub connect=CONNECTED
netsh routing ip rip update CorpHub
netsh interface set interface name=CorpHub connect=DISCONNECTED
```

You can run these commands from a batch file or you can place them in a Netsh script file. For example, the script file, Corphub.scp, runs the following commands for CorpHub:

```
interface set interface name=CorpHub connect=CONNECTED
routing ip rip update CorpHub
interface set interface name=CorpHub connect=DISCONNECTED
```

To run the Corphub.scp script, type the following at a command prompt:

```
netsh -f corphub.scp
```

After the batch file or Netsh script file is created, you can execute the batch file or Netsh script on a scheduled basis by using Task Scheduler.

For more information about Netsh, see "Command-Line References" in the Tools and Settings Collection.

**Demand-Dial Routing Registry Settings**
The following registry settings are associated with demand-dial routing.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_LOCAL_MACHINE\SYSTEM**
The following registry key is located under
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn

**BuiltIn**

**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Rasman\PPP\ControlProtocols\

**Version**
The BuiltIn registry setting is available in Windows Server 2003, Windows 2000 Server, and Windows XP.

After a demand-dial connection (physical or logical) is made, a PPP connection is negotiated with the endpoint. By default, the calling router does not register itself with the DNS or WINS servers of the answering router unless the value of **RegisterRoutersWithNameServers** is set to 1.

**HKEY_LOCAL_MACHINE\SYSTEM**
The following registry key is located under
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Interfaces

**Interfaces**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\

**Version**
The Interfaces registry setting is available in Windows Server 2003, Windows 2000 Server, and Windows XP.

The Interfaces key is populated with the interfaces that have been added to the local system from the Router configuration. Each interface has an associated Type (Internal, Dedicated, or Dynamic) and subkeys, such as Ip, for each router manager.

**HKEY_LOCAL_MACHINE\SOFTWARE**
Every router manager installed in the system has a registry key created under the Router registry key. For the Ip subkey, the DllPath name specifies the location of the DLL that corresponds to the router manager; the ProtocolID name specifies the protocol family identifier for the router manager.

The following registry key is located under
HKEY_LOCAL_MACHINE\Software\Microsoft\Router\CurrentVersion\RouterManagers\Ip\

**OSPF**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Microsoft\Router\CurrentVersion\RouterManagers\Ip\

**Version**
The OSPF registry setting is available in Windows Server 2003 and Windows 2000 Server.

Values for the OSPF registry key are set as shown in the following table.

### OSPF Registry Key Values

| Name | Type | Data |
|---|---|---|
| ConfigDll | REG_SZ | ipadmin.dll |
| DllName | REG_SZ | ospf.dll |
| ProtocolID | REG_DWORD | 0x0000000d |
| Title | REG_SZ | Open Shortest Path First |

The following registry key is located under
HKEY_LOCAL_MACHINE\Software\Microsoft\Router\CurrentVersion\RouterManagers\Ip\

**IPBOOTP**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Microsoft\Router\CurrentVersion\RouterManagers\Ip\

**Version**
The IPBOOTP registry setting is available in Windows Server 2003 and Windows 2000 Server.

Values for the IPBOOTP registry key are set as shown in the following table.

### IPBOOTP Registry Key Values

| Name | Type | Data |
|---|---|---|

| | | |
|---|---|---|
| ConfigDll | REG_SZ | ipadmin.dll |
| DllName | REG_SZ | ipbootp.dll |
| ProtocolID | REG_DWORD | 0x0000270f |
| Title | REG_SZ | DHCP Relay Agent |

The following registry key is located under HKEY_LOCAL_MACHINE\Software\Microsoft\Router\CurrentVersion\RouterManagers\Ip\

**IPRIP**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Microsoft\Router\CurrentVersion\RouterManagers\Ip\

**Version**
The IPRIP registry setting is available in Windows Server 2003 and Windows 2000 Server.

Values for the IPRIP registry key are set as shown in the following table.

### IPRIP Registry Key Values

| Name | Type | Data |
|---|---|---|
| ConfigDll | REG_SZ | ipadmin.dll |
| DllName | REG_SZ | iprip2.dll |
| ProtocolID | REG_DWORD | 0x00000008 |
| Title | REG_SZ | RIP Version 2 for Internet Protocol |

**HKEY_LOCAL_MACHINE\SOFTWARE**
The following registry key is located under HKEY_LOCAL_MACHINE\Software\Microsoft\Router\CurrentVersion\UiConfigDlls\

**UiConfigDlls**
Under the Router registry key, in addition to registry keys for the router manager, there is the UiConfigDlls registry key.

**Registry path**
HKEY_LOCAL_MACHINE\Software\Microsoft\Router\CurrentVersion\UiConfigDlls\

**Version**
The UiConfigDlls registry setting is available in Windows Server 2003, Windows 2000 Server, and Windows XP.

The UiConfigDlls registry key requires that globally unique identifier (GUID) values are set as shown in the following table.

### UiConfigDlls Registry Key Values

| Name | Type | Data |
|---|---|---|
| <guid1> | REG_SZ | ifadmin.dll |
| <guid2> | REG_SZ | ipadmin.dll |
| <guid3> | REG_SZ | ddadmin.dll |

# IPv4 Multicasting Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv4 Multicasting Technical Reference**
Internet Protocol (IP) multicasting is the delivery of a single IP datagram to a group of hosts, regardless of their location, on a TCP/IP-based network. Of the three delivery methods supported by IP — multicasting, unicasting, and broadcasting — multicasting is the most practical method for point-to-multipoint data transmission, because unicasting requires multiple copies of datagrams and broadcasting does not work across routers.

Multicasting in an IP version 4 (IPv4) network is different from multicasting in an IP version 6 (IPv6) network. This document describes multicasting with IPv4. IPv4 is frequently referred to as IP and that convention is used in this document. For information about multicasting with IPv6, see "IPv6 Technical Reference."

Although Windows Server 2003 does not provide any multicast routing protocols, the Routing and Remote Access service provides a platform that supports these multicast routing protocols. In addition, you can configure Routing and Remote Access as a multicast forwarder for limited network configurations.

**In this subject**


- What Is IPv4 Multicasting?
- How IPv4 Multicasting Works
- IPv4 Multicasting Tools and Settings

# What Is IPv4 Multicasting?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP Multicasting Overview**

The concept of group membership is central to IP multicasting. IP multicast datagrams are sent to a group, and only members of the group receive the datagrams. A group is identified by a single IP multicast address, which is an IP address in the Class D range of 224.0.0.0 to 239.255.255.255 (designated as 224.0.0.0/4 in classless interdomain routing (CIDR) notation). These Class D addresses are known as group addresses. A source host sends multicast datagrams to a group address. Destination hosts inform a local router that they need to join the group.

In an IP multicast-enabled intranet, any host can send IP multicast datagrams to any group address, and any host can receive IP multicast datagrams from any group address regardless of its location. To facilitate this capability, the hosts and routers on the intranet must support IP multicasting. Hosts use the Internet Group Management Protocol (IGMP) for establishing group membership. Routers use multicast routing protocols for forwarding multicast data.

The following figure illustrates a multicast-enabled intranet.

**Multicast-Enabled Network**




In this illustration, the hosts and routers are multicast-enabled so that the following can occur:

- The sending host sends multicast datagrams to a designated group address.
- The routers forward the multicast datagrams to any network segments that include group members. Routers can forward multicast traffic across a network, between networks, and across the Internet.
- The receiving hosts inform a local router to join the group, and then they receive all subsequent datagrams sent to the group address.
- If a receiving host leaves the group and detects that it might be the last group member on the subnet, it can contact the local router to leave the group, informing the router to stop forwarding the multicast datagrams to that subnet.

**Benefits of IP Multicasting**

Multicasting provides an efficient way to support high-bandwidth, one-to-many applications on a network:

- Multicasting can dramatically reduce network traffic by sending a single copy of the data.
- Hosts can be configured for multicasting without hardware upgrades.
- Because newer routers already support multicast forwarding and multicast routing protocols, enabling multicasting on a network is practical and cost-effective.

Multicasting is useful for many types of one-to-many applications, such as the following:

- Multimedia, such as video conferencing and collaborative computing.
- Automatic discovery of resources in an internetwork (in Windows Server 2003 for example, TCP/IP router discovery uses multicasting by default, and WINS uses multicasting during automatic discovery of replication partners).
- Datacasting such as file distribution or database synchronization.
- Mobile computer support such as remote address book updating.
- Distribution of organizational publications.

**IP Multicasting with Routing and Remote Access**

Windows Server 2003 does not provide multicast routing protocols, such as Distance Vector Multicast Routing Protocol (DVMRP), Multicast Extensions to Open Shortest Path First (MOSPF), and Protocol Independent Multicast (PIM), although Routing and Remote Access does support multicast routing protocols developed by independent software vendors (ISVs).

As an alternative, you can use the Routing and Remote Access service to forward multicast traffic. In this case, the Routing and Remote Access service uses IGMP as an IP routing protocol component. Router interfaces are configured in one of two operating modes: IGMP router mode or IGMP proxy mode. The purpose of IGMP router mode is to forward multicast traffic in a single-router intranet. The purpose of IGMP proxy mode is to connect a single-router intranet to a multicast-capable intranet or the Internet.

Although Routing and Remote Access uses IGMP in a limited way to enable multicast forwarding on an intranet, it is not the equivalent of a true multicast routing protocol. The Routing and Remote Access IGMP routing protocol component supports multicast forwarding for several specific internetwork topologies. For more information about the supported topologies, see "How IPv4 Multicasting Works."

**Related Information**

The following resources contain additional information that is relevant to this section.

- IPv6 Technical Reference
- TCP/IP Technical Reference
- Unicast IPv4 Routing Technical Reference

# How IPv4 Multicasting Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- IP Multicasting Architecture
- IP Multicasting Protocols
- IP Multicasting Processes and Interactions
- IP Multicasting with Routing and Remote Access
- IP Multicasting Interfaces
- Related Information


Internet Protocol (IP) multicasting is the sending of a single datagram to multiple hosts on a network or internetwork. Of the three delivery methods supported by IP, multicasting is the method that is most practical for one-to-many delivery. Unlike IP multicasting, IP unicasting sends a separate datagram to each recipient host. IP broadcasting sends a single datagram to all hosts on a single network segment (also known as subnet), even to those not interested in receiving it. Recent trends toward multimedia applications such as video conferencing necessitate the use of multicasting to efficiently send traffic to multiple hosts.

Before sending or receiving IP multicast data, a network must be enabled for multicasting, as follows:

- Hosts must be configured to send and receive multicast data..
- Routers must support the Internet Group Membership Protocol (IGMP), multicast forwarding, and multicast routing protocols.

Windows Server 2003 supports IP multicasting in the following ways:

- All computers running Windows Server 2003 can both send and receive multicast data.
- Windows Server 2003 TCP/IP (IPv4) can forward multicast IP traffic.
- Windows Server 2003 Routing and Remote Access supports IGMP for forwarding multicast traffic in a single-router intranet and for passing group membership information from a single-router intranet to a multicast-enabled intranet or the Internet.
- Routing and Remote Access provides a platform for true IP multicast routing protocols for forwarding multicast traffic across an intranet.

This section covers the following:

- The components of IP multicasting and the role of each
- The protocols that support multicasting
- How multicasting works
- How Routing and Remote Access supports multicasting

- The interfaces available for developing multicast-enabled applications and multicast routing protocols

This document describes how multicasting works in an IPv4 network. In addition to IPv4, the Windows Server 2003 family includes the latest version of IP, known as IP version 6 (IPv6). For information about how multicasting works with IPv6, see "IPv6 Technical Reference."

**IP Multicasting Architecture**

To support multicasting in an internetwork, the hosts and routers must be multicast-enabled. In an IP multicast-enabled intranet, any host can send IP multicast datagrams, and any host can receive IP multicast datagrams, including sending and receiving across the Internet.

The source host sends multicast datagrams to a single Class D IP address, known as the group address. Any host that is interested in receiving the datagrams contacts a local router to join the multicast group and then receives all subsequent datagrams sent to that address.

Routers use a multicast routing protocol to determine which subnets include at least one interested multicast group member and to forward multicast datagrams only to those subnets that have group members or a router that has downstream group members. The IP header of a multicast datagram includes a Time-to-Live (TTL) value that determines how far routers can forward a multicast datagram.

The following figure shows how multicasting components fit in an internetwork.

**IP Multicasting Architecture**




In this illustration, the following is occurring:

1. Host A in Subnet 1 is the multicasting source and is sending multicast data to the multicast group address.

2. Host B in Subnet 1 has requested group membership from its local router. Because Host B has joined the multicast group, its network adapter is listening for datagrams sent to the multicast group address. The remaining host in Subnet 1 has not requested group membership, so its network adapter is filtering out traffic sent to the multicast group address.

3. Routers forward multicast data to any subnet that has group members or a router that has downstream group members. In this example, the routers are forwarding multicast data from Subnet 1 to Subnet 3. The data is also sent across Subnet 2, although it has no group members, because there are downstream group members.

4. Host C in Subnet 3 has requested group membership and is also receiving the multicast data.

5. Host D in Subnet 3 is sending a request to its local router to join the group. After Host D joins the group, its network adapter will also listen for traffic sent to the multicast group address.

Alternatively, multicast data might come from the multicast-enabled portion of the Internet, known as the MBone.

The following table introduces the multicasting components, which are described later in this section.

**IP Multicasting Components**

| Component | Description |
|---|---|
| Host (source or receiver) | A host is any client or server on the network. A multicast-enabled host is configured to send and receive (or only send) multicast data. |
| Router | A multicast router is capable of handling host requests to join or leave a group and of forwarding multicast data to subnets that contain group members. A multicast router can be either a non-Microsoft router that uses a multicast routing protocol or a Windows Server 2003-based server running the Routing and Remote Access service. |
| Multicast address | A Class D IP address used for sending IP multicast data. An IP multicast source sends the data to a single multicast address, as described later in this section. A specific IP multicast address is also known as a group address. |
| Multicast group | A multicast group is the set of hosts that listen for a specific IP multicast address. A multicast group is also known as a host group. |
| MBone | The Internet multicast backbone, or the portion of the Internet that supports multicast routing. |

**Host**

On a multicast-enabled network, a host that is configured to send multicast data can send datagrams to a single designated Class D IP address so that multiple hosts can receive the data. A host that is configured to receive multicast data can use IGMP to join a multicast group and then listen for datagrams sent to the multicast address. Hosts can send and receive multicast data from anywhere on an intranet or the Internet.

A host supports IP multicasting at one of the following levels:

- Level 0 - No support for sending or receiving IP multicast datagrams.
- Level 1 - Supports sending but not receiving IP multicast datagrams.
- Level 2 - Supports both sending and receiving IP multicast datagrams.

TCP/IP for Windows Server 2003 (as well as Windows XP) supports all levels of IP multicasting and, by default, is configured for Level 2 support. You can modify the level of support by changing the **IGMPLevel** subkey under HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\ in the registry.

**Note**

- It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**Router**

The role of a multicast router is to:

- Manage multicast group membership by processing IGMP requests to join or leave groups.
- Forward multicast traffic to subnets of the internetwork that contain multicast group members.

Multicast routing is more complex than unicast routing. The following table compares the characteristics of unicast routing and multicast routing.

**Unicast Routing and Multicast Routing Comparison**

| Unicast Routing Characteristics | Multicast Routing Characteristics |
|---|---|

| | |
|---|---|
| Unicast traffic is sent to a globally unique destination. | Multicast traffic is sent to an ambiguous group destination. |
| Unicast routes in the routing table summarize ranges of globally unique destinations. | Because group addresses represent different groups with different members, group addresses generally cannot be summarized in the IP multicast forwarding table. |
| Unicast routes are comparatively consistent, so the routing table needs to be updated only when the topology of the internetwork changes. Unicast routing protocols update the unicast routing table. | The location of group members is not consistent, so the IP multicast forwarding table might need to be updated whenever a group member joins or leaves a multicast group. Multicast routing protocols update the IP multicast forwarding table. |

Multicast routers can be hardware or software routers. Non-Microsoft routers that are multicast-enabled run a multicast routing protocol, such as one of the following:

- Distance Vector Multicast Routing Protocol (DVMRP)
- Multicast Extensions to Open Shortest Path First (MOSPF)
- Protocol Independent Multicast (PIM)

These protocols are briefly described later in this document. These protocols are not included with Windows Server 2003.

Windows Server 2003 Routing and Remote Access can be used for multicast forwarding. The Routing and Remote Access service includes an IGMP routing protocol component that uses two modes, router mode and proxy mode, to add entries to the TCP/IP multicast forwarding table. Because Routing and Remote Access is not designed to replace a true multicast routing protocol, it supports several specific configurations, which are described later in this document.

**Multicast Addresses**

IP addresses in the Class D range are reserved for IP multicasting. Class D addresses are in the range of 224.0.0.0 to 239.255.255.255, with the four high-order bits always set to 1110. In network prefix or classless interdomain routing (CIDR) notation, IP multicast addresses are summarized as 224.0.0.0/4.

Within the Class D range of addresses, certain ranges of multicast addresses have special meaning, as follows:

- Addresses in the range of 224.0.0.0 to 224.0.0.255 (224.0.0.0/24) are reserved for local subnet multicast traffic. Datagrams sent to addresses in this range are not forwarded by IP routers.
- Addresses in the range of 224.0.1.0 to 238.255.255.255 are known as globally scoped addresses, which means they can be used for multicasting across an intranet or the Internet. Some addresses within this range are reserved by the Internet Assigned Numbers Authority (IANA) for special purposes.
- Addresses in the range of 239.0.0.0 to 239.255.255.255 (239.0.0.0/8) are reserved for administratively scoped addresses. Administratively scoped addresses as defined by RFC 2365 are used to prevent the forwarding of multicast traffic across boundaries configured for the address. Multicast boundaries are described in more detail later in this document.

The following are examples of reserved IP multicast addresses:

- 224.0.0.1 - All hosts on this subnet
- 224.0.0.2 - All routers on this subnet
- 224.0.0.5 - Open Shortest Path First (OSPF) version 2, designed to reach all OSPF routers
- 224.0.0.6 - OSPF version 2, designed to reach all OSPF-designated routers and designated backup routers
- 224.0.0.9 - Routing Information Protocol (RIP) version 2
- 224.0.1.1 - Network Time Protocol

For the latest list of reserved multicast addresses, see the IANA link.

#### MAC-Level Multicast Addresses

As described previously, a multicast datagram is sent to a single group address, and only members of the group receive the datagram. For a network adapter to receive a multicast datagram, the IP multicast address must be mapped to a hardware, or media access control (MAC), multicast address. The network adapter then listens for any datagrams that are addressed to this MAC-level address and passes the datagrams up to the operating system.

The Internet Assigned Numbers Authority (IANA) has reserved a range of multicast addresses for Ethernet and Fiber Distributed Data Interface (FDDI) networks. The following table lists the MAC-level multicast addresses used for different types of networks.

### MAC-Level Multicast Addresses

| Network Type | Multicast Addresses | Related RFC |
|---|---|---|
| Ethernet and FDDI | 0x01-00-5E-00-00-00 to 0x01-00-5E-7F-FF-FF | 1112 |
| Token Ring | 0xC0-00-00-04-00-00 | 1469 |

Token Ring can use the same method of MAC-level multicast addressing as Ethernet and FDDI, but many Token Ring network adapters do not support it. By default, therefore, all IP multicast addresses on Token Ring networks are mapped to the single functional address specified in the preceding table. You can modify the use of this functional address by editing the **TrFunctionalMcastAddress** subkey under HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\ in the registry. IP multicasting on Token Ring network adapters is supported for Windows 2000, Windows XP, and the Windows Server 2003 family.

### Note

- It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

#### IP Multicast to MAC-Level Multicast Address Mapping

To map an IP multicast address to a MAC-level multicast address, the 23 low-order bits of the IP multicast address are mapped directly to the 23 low-order bits of the MAC-level multicast address. The 25 high-order bits in the MAC-level multicast address are fixed. The following figure illustrates how the IP multicast address is mapped to the MAC-level multicast address.

### IP Multicast Address Mapped to Ethernet or FDDI MAC Address




This figure also illustrates the following:

- The 4 high-order bits of the IP multicast address are fixed according to the Class D address definition.

- The 5 remaining bits (bits 5 to 9) in the IP multicast address are not mapped to the MAC-level multicast address, so the resulting MAC address is not unique. This means that it is possible for a host to receive MAC-level multicast datagrams for up to 32 groups to which it does not belong. However, IP drops datagrams that do not match a multicast address that is registered by an application or another protocol.

- The low-order bit of the first octet in a MAC-level destination address is the Individual/Group bit. This bit is set to 0 for a unicast or individual address and to 1 for a multicast address or a broadcast address.

**Example of MAC-level address mapping**
If the IP multicast address is 224.192.16.1, the MAC-level multicast address becomes 0x01-00-5E-40-10-01, as illustrated in the following figure.

### Sample Address Mapping




The 23 low-order bits in the IP multicast address map directly to the MAC-level multicast address. Because the entire third and fourth octets are included in the mapping, they can be converted directly to hexadecimal. The high-order bit of the second octet is dropped from the mapping, so that octet must be converted to binary first. In this example, dropping the high-order bit results in a value of 1000000, which converts to decimal 64 and hexadecimal 0x40. The shading in the figure shows that the 25 high-order bits of the MAC-level multicast address are fixed values.

**Multicast Groups**
The set of hosts listening for a specific IP multicast address is called a multicast group or a host group. Multicast groups have the following attributes:

- Group membership is dynamic; hosts can join and leave the group at any time.
- Members of a group can be located anywhere on a multicast-enabled network.
- A host need not be a member of a multicast group to send data to the group's IP multicast address.
- Group members can receive both unicast and multicast traffic.
- A host can be a member of any number of multicast groups.
- A multicast group has no size limits.
- Multicast groups can be either transient or permanent. Permanent groups are assigned a well-known multicast address. For example, multicast address 224.0.0.1 is a permanent address for the all-hosts multicast group. Membership in a permanent group is transient; only the address is permanent.

**MBone**
The MBone, or Internet multicast backbone, is the portion of the Internet that supports the forwarding of Internet-based IP multicast traffic. The MBone, used in conjunction with tunneling, provides a way for multicast traffic to travel across portions of the Internet that do not support multicasting.

Tunneling is the encapsulation of a datagram within another datagram. For multicast traffic to traverse the Internet, each IP multicast datagram must be encapsulated within an IP unicast datagram, which is known as IP-in-IP tunneling. An IP-in-IP tunnel is used to forward information between endpoints in an IP internetwork that have differing capabilities, thus acting as a logical link between the endpoints. The encapsulating IP header includes the addresses of the endpoints as the source and destination addresses. The following figure illustrates an encapsulated IP datagram.

### IP-in-IP Datagram Structure




The MBone consists of a series of multicast-enabled islands. These islands are collections of contiguous networks that use

routers that support multicast routing. IP-in-IP tunneling allows the multicast traffic to travel from one island to another. The IP header of the unicast datagram is addressed from a router in one multicast-enabled island to a router in another multicast-enabled island.

The MBone is used for audio and video multicasts of Internet Engineering Task Force (IETF), the National Aeronautics and Space Administration, and the United States House of Representatives and Senate meetings, among other uses. Support for the MBone can vary among Internet service providers (ISPs).

**IP Multicasting Protocols**

A variety of protocols are used for various aspects of IP multicasting. The following table briefly identifies the protocols described in this section.

**Protocols Used for IP Multicasting**

| Protocol | Description | RFC |
|---|---|---|
| IGMP | Internet Group Management Protocol. Hosts use this protocol to make group membership requests. Routers use this protocol to query for active group members in a subnet. (Note: IPv6 uses Multicast Listener Discovery (MLD) instead of IGMP to manage group membership.) | 1112 |
| DVMRP | Distance Vector Multicast Routing Protocol. This is a multicast routing protocol. | 1075 |
| MOSPF | Multicast Extensions to Open Shortest Path First. This is a multicast routing protocol. | 1584 |
| PIM | Protocol-Independent Multicast. This is a multicast routing protocol. | 2117 |
| MADCAP | Multicast Address Dynamic Client Allocation Protocol. This protocol enables a host to request an IP multicast address from a multicast address allocation server. | 2730 |
| PGM | Pragmatic General Multicast. PGM is a reliable multicast transport protocol. Typically, multicast datagrams are delivered by using UDP, which is an inherently unreliable transport protocol. Hosts and routers that are PGM-enabled can recover lost multicast datagrams. | 3208 |

**IGMP**

Internet Group Management Protocol (IGMP) maintains host group membership on a local subnet. Hosts use IGMP to communicate multicast group membership requests with their local multicast router. Routers receive the group membership requests and periodically send out queries to determine which host groups are active or inactive on the local subnet. This protocol is required to support Level 2 multicasting.

IGMP messages are encapsulated in IP datagrams and use IP protocol number 0x02. The following illustrates the resulting datagram.

**Encapsulated IGMP Message**




There are three versions of IGMP:

- IGMP version 1 is provided by TCP/IP for Microsoft Windows NT Server 4.0 Service Pack 3 and earlier.
- IGMP version 2 is provided by TCP/IP for Microsoft Windows NT Server 4.0 Service Pack 4 and later and Windows 2000. Routers using this version are backward-compatible with hosts using IGMP version 1.
- IGMP version 3 is provided by TCP/IP for Windows XP and Windows Server 2003. Routers using this version are backward-compatible with hosts using IGMP version 2 or IGMP version 1. Hosts running Windows XP or Windows Server 2003 use IGMP version 3 by default. To configure a host running Windows XP SP2 or later to use IGMP version 1 or IGMP version 2, use the **IGMPVersion** subkey under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\ in the registry.

  **Note**

- It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

This section describes the differences between the three versions of IGMP.

**IGMP Version 1**
IGMP version 1 consists of the following two types of messages:

- Host Membership Report
- Host Membership Query

The following table describes the fields in IGMP version 1 messages.

**Fields in IGMP Version 1 Messages**

| IGMP Version 1 Field | Length | Description |
|---|---|---|
| Version | 4 bits | Identifies the IGMP version. It is set to 1. |
| Type | 4 bits | Identifies the type of message:<br>• 0x1 for Host Membership Query<br>• 0x2 for Host Membership Report |
| Unused | 1 octet | Set to 0 by sender, ignored by receiver. |
| Checksum | 2 octets | Used to test for data transmission errors. |
| Group Address | 4 octets | Specifies the group address for the type of message:<br>• 0.0.0.0 for Host Membership Query<br>• The multicast group address to be joined for Host Membership Report |

**Host Membership Report**
To join a multicast group, a host sends an IGMP Host Membership Report message to the specific group address, regardless of whether other hosts on the subnet are already multicast group members. Unlike a multicast router, a host does not keep track of the multicast group membership of other hosts on its subnet. A multicast router uses a special listening mode called multicast promiscuous mode so that it can receive all IP multicast traffic. Because it listens in multicast promiscuous mode, a multicast router receives and processes IGMP Host Membership Report messages that are sent to any multicast address.

**Host Membership Query**
An IGMP version 1 multicast router (known as a multicast query router) periodically sends an IGMP Host Membership Query message to IP address 224.0.0.1 (the all-hosts group) to obtain updated information about multicast group membership on the subnet. The multicast routing protocols determine which router is to send the query messages.

Multicast routers do not need to track the specific hosts that belong to a group, but only that at least one host in a subnet belongs to the group. For each multicast group that has members on the subnet, one multicast group member responds with an IGMP Host Membership Report message. To keep multiple hosts on a particular subnet from sending IGMP Host Membership Report messages for the same group, a random response timer on the hosts delays the transmission of the IGMP Host Membership Report message. If another host on the same subnet sends a message before the response timer expires, a host does not send its own message.

**Addresses and TTLs used for IGMP version 1 messages**
The following table summarizes the values of the source and destination IP addresses in the IP header, the TTL in the IP header,

and the Group Address in the IGMP header for the two types of IGMP version 1 messages.

**Addresses and TTLs Used for IGMP Version 1 Messages**

| Address | IGMP Host Membership Report | IGMP Host Membership Query |
|---|---|---|
| Source IP address (IP header) | *<HostInterfaceIPAddress>* | *<RouterInterfaceIPAddress>* |
| Destination IP address (IP header) | *<GroupAddress>* | 224.0.0.1 |
| TTL | 1 | 1 |
| Group address | *<GroupAddress>* | 0.0.0.0 |

IGMP version 1 has no explicit mechanism for hosts to notify a router that it is ready to leave a multicast group.

For more information about IGMP version 1, see Appendix I of RFC 1112 in the IETF RFC Database.

**IGMP Version 2**

IGMP version 2 extends the functionality of IGMP version 1 by including the following:

- A procedure for determining which router in a subnet is to be the multicast query router
- A new message for when hosts leave a multicast group
- A new group-specific query message
- A new version of the Host Membership Report

The following table describes the fields in IGMP version 2 messages.

**Fields in IGMP Version 2 Messages**

| IGMP Version 2 Field | Length | Description |
|---|---|---|
| Type | 1 octet | Identifies the type of message:<br>• 0x11 for Host Membership Query<br>• 0x12 for IGMP version 1 Host Membership Report<br>• 0x16 for IGMP version 2 Host Membership Report<br>• 0x17 for Leave Group Message |
| Maximum Response Time | 1 octet | Specifies the maximum allowable time (in tenths of a second) for a host to respond to a router query. Used only in query messages.<br>For Routing and Remote Access routers, this value is configured with the **Query response interval** from either of the following sources:<br>• Routing and Remote Access Microsoft Management Console (MMC) snap-in **Router** tab for the IGMP interface properties<br>• **Netsh routing ip igmp set interface** command |
| Checksum | 2 octets | Used to test for data transmission errors. |

| | | |
|---|---|---|
| Group Address | 4 octets | Specifies the group address for the type of message:<br>• 0.0.0.0 for general Host Membership Query message<br>• The specific multicast group address for Group-Specific Host Membership Query, Host Membership Report, and Leave Group messages |

**Note**

- The Version field is eliminated and the Type field is 8 bits in IGMP version 2. The hexadecimal value of the Type field in version 2 messages is the result of concatenating the version 1 hexadecimal values of the Version field (0x1) and the Type field (0x1 for Host Membership Query and 0x2 for Host Membership Report).

**Query router election**

In contrast to IGMP version 1, in which the multicast routing protocols determine the multicast query router, IGMP version 2 uses a simple election process to choose one router on each subnet to send the periodic IGMP Host Membership Query messages. The elected router is the one with the numerically lowest IP address. During the election process, each router listens for IGMP queries from other routers. If a router receives a query with a lower source IP address, that router does not send multicast queries. If a router does not receive a query with a lower source IP address, it becomes the multicast query router for the subnet.

**Leave Group message**

The Leave Group message reduces leave latency, which is the time it takes for a multicast router to stop forwarding multicast traffic to a subnet that no longer contains multicast group members. If a host responds to the last IGMP query, it might be the last or only member of the multicast group. When this host leaves the group, it sends an IGMP Leave Group message to the all-routers multicast group (224.0.0.2). Upon receipt of the Leave Group message, the router sends a series of group-specific queries to the multicast group. If no host responds to the group-specific queries, the router determines that the multicast group on that subnet contains no members.

**IGMP group-specific query**

An IGMP Host Membership Query is sent to 224.0.0.1 (the all-hosts group) to find out which multicast groups have members on the subnet. IGMP version 2 routers can also send a group-specific query to a specific group address to find out if that group has any members on the subnet.

**IGMP version 2 Host Membership Report**

The IGMP version 2 Host Membership Report has the same function as the IGMP version 1 Host Membership Report, except that it is used by IGMP version 2 routers.

**Addresses and TTLs used for IGMP version 2 messages**

The following table summarizes the values of the source IP address, the destination IP address and the TTL in the IP header, as well as the value of the Group Address in the IGMP header for the two additional types of IGMP version 2 messages.

### Addresses and TTLs Used for IGMP Version 2 Messages

| Address | IGMP Leave Group Message | IGMP Group-Specific Query |
|---|---|---|
| Source IP address (IP header) | *<HostInterfaceIPAddress>* | *<RouterInterfaceIPAddress>* |
| Destination IP address (IP header) | 224.0.0.2 | *<GroupBeingQueried>* |
| TTL | 1 | 1 |
| Group address | *<GroupBeingLeft>* | *<GroupBeingQueried>* |

For more information about IGMP version 2, see RFC 2236 in the IETF RFC Database.

**IGMP Version 3**

IGMP version 3 extends the functionality of IGMP version 2 by adding support for source filtering, also known as source-specific multicasting. With IGMP version 3, it is possible to have multiple sources for a multicast data stream so that hosts can join the group and receive the data from the nearest source. For example, you might use this feature for an organization-wide video conference. You can have multiple sources multicast the video conference and place the sources in different regions so that hosts can join the group and specify the source topologically nearest them.

IGMP version 3 has a modified version of the Host Membership Query message and includes a new version of the Host

Membership Report.

**Source filtering**
Source filtering is the ability of a host to specify that it receive multicast traffic only from specific source addresses. With IGMP version 1 and IGMP version 2, a host reports group membership regardless of the source of the multicast traffic sent to its group. If multicast traffic is sent to a group from multiple sources, a subnet with a group member receives one multicast datagram from each source.

IGMP version 3 allows a host to specify either of the following for a particular multicast group:

- The sources from which it is to receive messages pertaining to the group (a list of included sources).
- The sources from which it is not to receive messages pertaining to the group (a list of excluded sources).

Multicast routers and multicast routing protocols use the source-specific information in the IGMP version 3 Host Membership Report message to prevent forwarding and delivering multicast messages from specific sources to subnets that have no interested group members.

**IGMP version 3 Host Membership Query**
The IGMP version 3 Host Membership Query message has the same Type value and the same format as the IGMP version 2 Host Membership Query message, except that it includes additional fields after the Group Address field. The additional fields provide query parameters for routers and specify the included or excluded sources for the multicast group being queried. The list of included or excluded sources is used only for queries to a specific group that uses source filtering.

The first 8 octets of the IGMP version 3 Host Membership Query message are the same as IGMP version 2. The Type value for the modified Host Membership Query remains 0x11. The following table describes the fields in IGMP version 3 Host Membership Query messages.

**Fields in IGMP Version 3 Host Membership Query Messages**

| IGMP Version 3 Field | Length | Description |
|---|---|---|
| Type | 1 octet | Identifies the type of message:<br>• 0x11 for Host Membership Query |
| Maximum Response Time | 1 octet | Specifies the maximum allowable time (in tenths of a second) for a host to respond to a general (not group-specific) router query.<br>For Routing and Remote Access routers, this value is configured with the **Query response interval** from either of the following sources:<br>• Routing and Remote Access snap-in **Router** tab for the IGMP interface properties<br>• **Netsh routing ip igmp set interface** command |
| Checksum | 2 octets | Used to test for data transmission errors. |
| Group Address | 4 octets | Specifies the group address for the query message:<br>• 0.0.0.0 for general Host Membership Query<br>• The specific multicast group address for Group-Specific Host Membership Query, Host Membership Report, and Leave Group |
| Reserved | 4 bits | Set to 0 by sender, ignored by receiver. |
| Suppress Router-Side Processing | 1 bit | Set to 1 to specify that receiving routers are to suppress normal timer updates when they receive a query. |

| | | |
|---|---|---|
| Querier's Robustness Variable (QRV) | 3 bits | Allows tuning for expected datagram loss on the network. IGMP can recover QRV-1 lost datagrams.<br><br>For Routing and Remote Access routers, this value is configured with the **Robustness variable** from either of the following sources:<br><br>• Routing and Remote Access snap-in **Router** tab for the IGMP interface properties<br><br>• **Netsh routing ip igmp set interface** command |
| Querier's Query Interval Code | 1 octet | Specifies how many seconds the router waits between general queries.<br><br>For Routing and Remote Access routers, this value is configured with the **Query interval** from either of the following sources:<br><br>• Routing and Remote Access snap-in **Router** tab for the IGMP interface properties<br><br>• **Netsh routing ip igmp set interface** command |
| Number of Sources | 2 octets | Specifies the number of source addresses included in the query message. |
| Source Addresses | 4 octets per address | Specifies the IP addresses of multicast sources. |

**IGMP version 3 Host Membership Report**

A host uses the IGMP version 3 Host Membership Report message to specify all the multicast groups it wants to join, and for each group address, either the sources from which it is to receive multicast data (the list of included sources) or the sources from which it is not to receive multicast data (the list of excluded sources).

The IGMP version 3 Host Membership Report message contains a series of group records. Each record contains the multicast group address and the associated list of sources. The Host Membership Report message is sent to the multicast address 224.0.0.22, which is a reserved multicast address for all IGMP version 3-capable multicast routers.

The following table describes the fields in IGMP version 3 Host Membership Report messages.

**Fields in IGMP Version 3 Host Membership Report Messages**

| IGMP Version 3 Field | Length | Description |
|---|---|---|
| Type | 1 octet | Identifies the type of message:<br><br>• 0x12 for IGMP version 1 Host Membership Report<br><br>• 0x16 for IGMP version 2 Host Membership Report<br><br>• 0x22 for IGMP version 3 Host Membership Report |
| Reserved | 1 octet | Set to 0 by sender, ignored by receiver. |
| Checksum | 2 octets | Used to test for data transmission errors. |
| Reserved | 2 octets | Set to 0 by sender, ignored by receiver. |
| Number of Records | 2 octets | Specifies the number of group records in the message. |
| Group Records | Variable | Each record specifies the IP address for a multicast group being joined or left and a list of included or excluded sources. |

**Addresses and TTLs used for IGMP version 3 messages**

The following table summarizes the values of the source and destination IP addresses, the TTL in the IP header, and the value of the Group Address in the IGMP header for the two additional types of IGMP version 3 messages.

**Addresses and TTLs Used for IGMP Version 3 Messages**

| Address | IGMP Host Membership Query for Specific Groups and Sources | IGMP Host Membership Report |
|---|---|---|
| Source IP address (IP header) | <*RouterInterfaceIPAddress*> | <*HostInterfaceIPAddress*> |
| Destination IP address (IP header) | <*GroupBeingQueried*> | 224.0.0.22 |
| TTL (IP header) | 1 | 1 |
| Group address (IGMP header) | <*GroupBeingQueried*> | <*GroupAddress*> in individual group address records |

For more information about IGMP version 3, see RFC 3376 in the IETF RFC Database.

**Multicast Routing Protocols: DVMRP, MOSPF, and PIM**

The goals of multicast routing protocols include the following:

- Forward traffic away from the source to prevent loops.
- Minimize or eliminate multicast traffic to subnets that do not need to receive it.
- Minimize CPU and memory load on the router for scalability.
- Minimize the traffic overhead of a routing protocol.
- Minimize the join latency, which is the time it takes for the first group member on a subnet to begin receiving group traffic.

Multicast routing protocols use a variety of algorithms to create efficient multicast delivery paths through the internetwork. Unicast routers use the IP destination address to forward traffic toward the destination. Because multicast traffic is sent to a group address, multicast routers use the source address to forward traffic away from the source to avoid looping the traffic back to the source.

This section provides a brief introduction to the most common multicasting routing protocols. These protocols are not provided with Windows Server 2003.

**Distance Vector Multicast Routing Protocol (DVMRP)**

DVMRP is a multicast routing protocol that grew from Routing Information Protocol (RIP), a unicast routing protocol that forwards datagrams based on information about the next hop toward the destination. In contrast to RIP, DVMRP forwards datagrams based on information about the previous hop back toward the source.

DVMRP builds a delivery tree for determining where to forward datagrams. It floods the first multicast datagrams to a multicast group across the internetwork and then prunes the branches of the tree that lead to subnets that do not include group members.

DVMRP is most efficient in environments where multicast group members are densely distributed over the internetwork. DVMRP runs on internetworks with multicast routers and also supports tunneling (encapsulating multicast datagrams within a unicast datagram for forwarding through non-multicast routers). DVMRP does not scale well.

**Multicast Extensions to OSPF (MOSPF)**

MOSPF is a multicast routing protocol that is an extension to Open Shortest Path First (OSPF). This protocol builds a delivery tree for determining where to forward datagrams by using group membership information from IGMP and the OSPF link-state database. The trees are built on demand for each source-group pair.

MOSPF is intended for use on a single organization’s network and does not scale well. MOSPF requires OSPF as its accompanying unicast routing protocol. It can sometimes put a heavy load on router CPU bandwidth.

MOSPF is most efficient in environments where multicast group members are densely distributed over the internetwork. It does not support tunneling but can work in multicast environments with non-MOSPF routers.

**Protocol-Independent Multicast (PIM)**
PIM consists of two protocols: PIM Dense Mode (PIM-DM) and PIM Sparse Mode (PIM-SM). PIM-DM is most efficient in environments where multicast group members are densely distributed over the internetwork. PIM-SM is most efficient in environments where group members are sparsely distributed. A multicast group that uses PIM can declare itself sparse or dense.

The PIM protocol establishes routes to multicast groups whose members span wide-area and interdomain internetworks. PIM functions independently of any specific unicast routing protocol, although it does use the existing unicast routing table.

**PIM-DM**
This protocol is designed for multicast groups whose members are distributed densely over an area where bandwidth is plentiful. Like DVMRP, PIM-DM first floods multicast traffic across the internetwork and then prunes the subnets that have no group members.

PIM-DM is interoperable with PIM-SM. PIM-DM does not scale well.

**PIM-SM**
PIM-SM is the most widely used multicast routing protocol. It is designed for multicast groups with members distributed sparsely across a large region and is most efficient in a WAN environment. The flood-and-prune method used by PIM-DM can cause unnecessary traffic. In contrast, PIM-SM forwards multicast traffic only to routers that explicitly request it. With PIM-SM, routers must explicitly join and leave multicast groups to receive or stop receiving multicast traffic.

**MADCAP**
Multicast Address Dynamic Client Allocation Protocol (MADCAP) is a method for hosts to request a multicast address from a multicast address allocation server. If multiple hosts use the same IP multicast address for different applications, multicast traffic can be forwarded to the incorrect group members. MADCAP prevents this problem by allocating unique multicast addresses. It supports dynamic assignment and configuration of IP multicast addresses on TCP/IP-based networks.

MADCAP is an extension of the Dynamic Host Configuration Protocol (DHCP) standard, but is separate from DHCP. The Windows Server 2003 DHCP service supports both DHCP and MADCAP, but the functions are independent.

Multicast administrative scopes configured on the DHCP server define ranges of IP multicast addresses. Unique IP multicast addresses are allocated to a single DHCP client, similar to how unique unicast IP addresses are assigned to clients.

For more information about MADCAP and its support in Windows Server 2003, see "DHCP Technical Reference." For more information about the IETF definition of MADCAP, see RFC 2730 in the IETF RFC Database.

**PGM**
Pragmatic General Multicast (PGM) is a reliable multicast transport protocol for multicast-enabled applications that require properly sequenced, duplicate-free delivery of datagrams. Typically, multicast datagrams are sent as UDP datagrams, but UDP is inherently unreliable because it does not guarantee packet delivery or retransmit lost packets. PGM ensures that a host that is a member of a multicast group either receives all the transmitted datagrams or detects unrecoverable lost datagrams. It is implemented on both source and destination hosts.

PGM is a transport-layer protocol that runs directly over IP and uses IP protocol number 113. It is the type of reliable multicast protocol known as router assist. With this type of protocol, PGM-enabled receiving hosts use negative acknowledgement (NAK) packets to request the retransmission of lost multicast datagrams, and PGM-enabled routers on the network can retransmit lost packets on behalf of the sender. The PGM-enabled routers define a logical PGM topology that overlays the physical IP internetwork topology. This topology defines the series of hops used for PGM packets between the sending host and the receiving hosts. Routers are not required for PGM: the PGM topology can consist of a single logical hop between the sending host and its receiving hosts.

PGM is currently the only reliable multicast protocol supported by Windows Server 2003. To use PGM with Windows Server 2003, you must add the Reliable Multicast Protocol component in Network Connections to sending and receiving hosts and create PGM-enabled applications. For more information about creating PGM-enabled applications, see "Reliable Multicast Programming (PGM)" on MSDN.

With PGM, a session is established during multicast datagram transmission. The source host caches transmission for retransmission in the event of an error. If a receiving host detects an error, it sends a unicast NAK to the next upstream PGM router. The PGM router confirms receipt of the NAK packet and forwards the NAK packet upstream until it reaches the source host. The source host then retransmits the datagram and routers forward it only to the subnets that require it.

**IP Multicasting Processes and Interactions**

In a multicast-enabled intranet, hosts send and receive multicast datagrams and routers manage group membership and forward multicast datagrams. The Windows Server 2003 implementation of TCP/IP supports all levels of multicasting.

**Host Processes**

In a Level 2 multicast-enabled intranet, any host can send IP multicast datagrams to any group address, and any host can receive IP multicast datagrams from any group address. In a Level 1 multicast-enabled intranet, any host can send multicast datagrams but cannot receive them. This section describes what happens when hosts send and receive multicast datagrams.

**Sending IP Multicast Datagrams**

To send IP multicast datagrams, a host determines the IP multicast address to use, constructs the multicast datagram, and then places the multicast datagram on the medium.

An IP multicast address can be obtained in any of the following ways:

- Assigned by the application
- Selected from well-known (reserved) IP multicast addresses defined by the IANA
- Allocated by an address allocation protocol, such as MADCAP

The sending host constructs the IP datagram by using its own IP address as the source address, the obtained IP multicast address as the destination address, and an appropriate TTL value.

The TTL value for an IP multicast datagram determines how many subnets the datagram can travel. As specified earlier, IP multicast addresses in the range 224.0.0.0 to 224.0.0.255 (224.0.0.0/24) are reserved for local subnet multicast traffic. By convention, the TTL value is 1 for a datagram sent to an address in this range, but routers do not forward multicast traffic sent to an address in this range even if the TTL is greater than 1. For multicast datagrams that are to be forwarded across an intranet or the Internet, the TTL value must be high enough to reach all members of the multicast group. The following table describes the conventional definitions for TTL values.

**TTL Value Conventional Definitions**

| TTL Value | Definition |
|---|---|
| 0 | Restricted to the same host |
| 1 | Restricted to the local subnet, no router hops |
| 32 | Restricted to the site |
| 64 | Restricted to the region |
| 128 | Restricted to the continent |
| 255 | Worldwide (unrestricted) |

After the host constructs the IP multicast datagram, the datagram passes through the IP routing process and is sent to Address Resolution Protocol (ARP), which maps the IP address to the hardware address. An ARP extension processes IP multicast addresses. If the IP address is in the multicast range, ARP does not check the ARP cache or broadcast a request datagram to locate the corresponding MAC address. Instead, for Ethernet and FDDI networks, ARP maps the IP destination address to the MAC-level address, as described in "IP Multicast Architecture." For Token Ring networks, ARP assigns a single functional address, as described in "IP Multicasting Architecture" earlier in this document. For more information about how ARP works, see "How TCP/IP Works." For more information about the IP routing process, see "Unicast IPv4 Routing Technical Reference."

**Receiving IP Multicast Datagrams**

To receive IP multicast datagrams, a host notifies IP to process multicast datagrams for a specific group address, registers the appropriate multicast MAC address with the network adapter, and notifies the local router to join the multicast group.

To determine the specific group address to use, the host can obtain an address from the application or derive the address from a universal resource locator (URL). On the MBone, multicasts are typically announced with the Session Description Protocol (SDP). SDP provides all the information needed to receive a multicast, including the multicast name and description, the times it is active, the type of media it uses, as well as the IP addresses, ports, and protocol it uses. The information is sent to a well-known IP address and port where a host running the session directory tool receives it.

After determining the group address, the host informs IP to process multicast datagrams with that destination address. For shared access technologies such as Ethernet, FDDI, and Token Ring (in which all attached nodes use a common wire), IP contacts the Network Interface layer to have the MAC-level multicast address that corresponds to the group address added to the table of interesting destination MAC addresses, which identifies the datagrams to be received and passed up to the operating system. To accomplish this task in the Windows Server 2003 family, IP uses the **NDISRequest()** function.

If the group address of interest is outside the range of 224.0.0.0 through 224.0.0.255 or 224.0.0.0/24 (the group address is non-local), the host sends an IGMP Host Membership Report message to the local router to inform it to forward datagrams for that group address to the host's subnet.

If the host has multiple applications that use the same IP multicast address, IP keeps track of group membership for each application and passes a copy of the multicast datagram to each application. If a host is multihomed, IP keeps track of group membership for each subnet.

A host stops receiving datagrams for a specific multicast group address when the corresponding MAC-level address is removed from the table of interesting destination MAC addresses. This occurs when a multicast-enabled application notifies Windows Sockets to stop receiving datagrams for a specific group address, such as at application shutdown.

**Router Processes**
Multicast routers handle multicast group membership and forward IP multicast datagrams to the subnets that have at least one group member. Multicast routers do the following:

- Send and receive IGMP group membership messages.
- Communicate group membership to other multicast routers.
- Receive all IP multicast traffic.
- Forward IP multicast traffic.

**Send and Receive IGMP Group Membership Messages**
**Receive and process IGMP Host Membership Report and Leave Group messages**
Multicast routers receive IGMP Host Membership Report messages from any hosts on local subnets and process these messages by placing entries that track group membership in the IP multicast forwarding table. Because multicast routers listen in multicast promiscuous mode, they receive all IGMP Host Membership Report messages sent to any group address.

An IGMP version 2 host that might be the last group member on a subnet sends a Leave Group message. Multicast routers process these messages by sending group-specific queries to the multicast group. If no host responds to the queries, the multicast router removes the entry for the subnet from the IP multicast forwarding table.

**Query local subnets for host membership status**
Any single subnet can support any combination of IGMP version 1, IGMP version 2, and IGMP version 3 hosts. Because IGMP version 1 does not support Leave Group messages, an IGMP version 1 host cannot notify the local router when it leaves a multicast group. In addition, because UDP is used to send multicast datagrams, it is possible for IGMP Host Membership Report and Leave Group messages to be lost. In cases such as these, multicast routers can continue to forward multicast traffic to a subnet that no longer has any group members.

To improve leave latency, multicast routers periodically send IGMP Host Membership Query messages to the local subnet to obtain updated host membership information. For each multicast group that has members on the subnet, one multicast group member responds to the query with an IGMP Host Membership Report message. Multicast routers then process these messages to update the IP multicast forwarding table.

**Communicate Group Membership to Other Multicast Routers**
In a multicast-enabled internetwork that contains multiple routers, each router communicates group membership information with its neighboring routers so that group members located anywhere on the internetwork receive the group traffic.

Multicast routers communicate group membership information either explicitly, by exchanging group address and subnet information, or implicitly by informing upstream routers whether group members exist downstream from the source host. The routers use a multicast routing protocol such as DVMRP, MOSPF, or PIM to exchange the group membership information.

**Receive All IP Multicast Traffic**
To receive all multicast traffic, a multicast router must put its network adapter in a special listening mode. A listening mode is the way a network adapter analyzes the destination MAC address of incoming datagrams to determine whether to process them further.

For shared access technologies such as Ethernet, FDDI, and Token Ring, the normal listening mode for network adapters is unicast listening mode. In unicast listening mode, the only datagrams that are considered for further processing are in a table of interesting destination MAC addresses on the network adapter. Typically, the only interesting addresses are the broadcast address (0xFF-FF-FF-FF-FF-FF) and the MAC address of the network adapter.

The special listening mode needed to receive all multicast traffic is called multicast promiscuous mode. Multicast promiscuous mode analyzes the Individual/Group bit in the destination MAC address to determine whether the datagram requires further processing. The Individual/Group bit for Ethernet and FDDI addresses is the low-order bit of the first octet in the destination MAC address. This bit is set to 1 for multicast addresses. When the network adapter is placed in multicast promiscuous mode, any datagrams with the Individual/Group bit set to 1 are passed up to the operating system for further processing.

Multicast promiscuous mode is different from promiscuous mode. In promiscuous mode, all datagrams — regardless of the destination MAC address — are passed up to the operating system for processing. Promiscuous mode is used by protocol analyzers, also known as network sniffers, such as the full version of Network Monitor that is part of the Microsoft Systems Management Server.

A network adapter that supports promiscuous mode might not support multicast promiscuous mode. However, multicast promiscuous mode is supported by most network adapters.

**Forward IP Multicast Traffic**

Multicast routers forward multicast traffic downstream, away from the source host and toward subnets that have at least one group member. The Windows Server 2003 implementation of TCP/IP supports multicast forwarding.

At each hop in the multicast forwarding process, the router decrements the TTL in the IP header. If the resulting TTL is greater than zero, the router compares the source address and destination group address to entries in the local IP multicast forwarding table and then forwards the datagram over the appropriate interfaces. If the TTL equals 0 after decrementing, the router silently discards the multicast datagram. Destination addresses in the range of 224.0.0.0 to 224.0.0.255 (224.0.0.0/24) designate multicast traffic on the local subnet. Multicast routers do not forward datagrams sent to addresses in this range, regardless of the TTL value.

Entries in an IP multicast forwarding table do not identify which, or how many, hosts are group members, but only that at least one host is a member. An entry in an IP multicast forwarding table identifies a subnet that has either a group member or a router that has downstream group members. A multicast router, therefore, can forward multicast datagrams to a subnet that has no group members because the subnet is on the path to other subnets that do have group members.

**Multicast boundaries**

Multicast boundaries are configurable administrative barriers that limit how far routers can forward multicast traffic across an internetwork. Without boundaries, a multicast router forwards all appropriate multicast traffic throughout the internetwork. With boundaries, a multicast router forwards appropriate multicast traffic only over a limited area of the internetwork. There are two types of multicast boundaries, as follows:

- **Scope-based boundaries.** A scope is a range of IP multicast addresses and is defined by specifying the address range as an IP address and mask. Multicast addresses in the range of 239.0.0.0 to 239.255.255.255 (239.0.0.0/8) are used to specify administrative scopes, as defined in RFC 2365. However, the range of 239.255.0.0/16 is defined as a local scope and cannot be divided further.

  An administrative scope becomes a boundary, beyond which multicast datagrams are not forwarded. Routers that support administratively scoped multicasts can have each interface configured with a scoped boundary. Such a router is called a boundary router. Datagrams that match the interface boundary definition of a boundary router do not get forwarded in either direction.

- **TTL-based boundaries.** TTL-based boundaries prevent the forwarding of multicast datagrams beyond a specified number of router hops, as described earlier in this section. TTL-based boundaries apply to all multicast datagrams regardless of the group address. TTL-based boundaries are less precise than scope-based boundaries because of interactions with multicast routing protocols.

**IP Multicasting with Routing and Remote Access**

Routing and Remote Access can forward multicast traffic in limited network configurations. The primary use of Routing and Remote Access as a multicast router is to connect a subnet to a multicast-enabled intranet that contains routers running multicast routing protocols. To fully support efficient multicast forwarding on a multiple-router intranet, it is necessary to install multicast routers that run one or more multicast routing protocols.

Windows Server 2003 Routing and Remote Access supports IGMP version 1, version 2, and version 3 and also supports

multicast boundaries and multicast heartbeat. This section describes how Routing and Remote Access supports multicast traffic.

**Routing and Remote Access Architecture**

The following figure illustrates the components of Windows Server 2003 Routing and Remote Access. The components in clear boxes are unique to multicasting and are discussed in this document. For information about components that pertain to unicasting or to both unicasting and multicasting, see "Unicast IPv4 Routing Technical Reference."

### IP Multicast Components in Routing and Remote Access

The following table describes the multicast components of Routing and Remote Access.

**Multicast Components in Routing and Remote Access**

| Component | Description |
|---|---|
| IGMP | In Routing and Remote Access, IGMP is implemented as a routing protocol component. This implementation of IGMP has two modes: router mode and proxy mode. This component does the following:<br>• Communicates multicast group membership information to the Multicast Group Manager.<br>• Uses Windows Sockets to send and receive IGMP traffic. |
| Multicast Group Manager | The Multicast Group Manager component does the following:<br>• Maintains multicast group membership.<br>• Coordinates operations between all multicast routing protocols when multiple multicast routing protocols are enabled on a router.<br>• Communicates IP multicast forwarding table entries from multicast routing protocols to the IP Multicast Forwarder. |
| IP Multicast Forwarder | The IP Multicast Forwarder component does the following:<br>• Stores the multicast forwarding entries received from Multicast Group Manager.<br>• Communicates source and group information from received multicast traffic to Multicast Group Manager.<br>• Forwards multicast datagrams. |

**Routing and Remote Access Support for Multicast Forwarding**

In Routing and Remote Access, IGMP is implemented as a routing protocol component. The IGMP routing protocol component has two modes: IGMP router mode and IGMP proxy mode.

The Routing and Remote Access Server Setup Wizard adds the IGMP routing protocol component. You can then use the Routing and Remote Access snap-in to accomplish the following:

- Add router interfaces to the IGMP routing protocol.
- Configure each interface for one of the two operating modes and for the IGMP version.
- Remove IGMP as a routing protocol component.

You can configure IGMP router interfaces by using the **Router** tab on the properties of an IGMP interface in the Routing and Remote Access snap-in or by using the **netsh routing ip igmp set interface** command.

**IGMP Router Mode**

An interface that is configured for IGMP router mode supports the following functions:

- **Listens in multicast promiscuous mode.** If the network technology supports it, such as with Ethernet and FDDI, the router-mode interface is placed in multicast promiscuous mode. If the network technology does not support multicast promiscuous mode, an IP Router Manager event is logged in the system event log.

- **Receives and processes IGMP group membership messages.** The IGMP router-mode interface receives and processes IGMP Host Membership Report messages and Leave Group messages that are sent by hosts on the subnet as they join and leave groups.

- **Sends IGMP Host Membership queries.** The IGMP router-mode interface sends periodic general (all-hosts) queries and group-specific queries to obtain updated group membership information from the subnet.

- **Elects an IGMP query router.** Only one router on a subnet sends the IGMP queries. As an IGMP multicast router, the IGMP router-mode interface elects an IGMP query router for the subnet.

- **Maintains the IP multicast forwarding table.** The IP multicast forwarding table for the TCP/IP protocol is updated based on IGMP group membership messages and multicast datagrams received on the router-mode interface. An IP multicast forwarding table entry contains the multicast group address, the source IP address, a list of interfaces that receive forwarded traffic (next-hop interfaces), and the single interface on which traffic must be received to be forwarded (the previous-hop interface). The multicast group and source IP addresses are obtained from the multicast datagram. The next-hop interfaces are determined by the registration of multicast group members using IGMP (and any multicast routing protocols, if present). The previous-hop interface is the interface that is closest — in terms of routing metrics — to the source of the IP traffic. The previous-hop interface is determined by checking the multicast routing table and identifying a single interface that is the best route back to the source of the IP multicast datagram. The best route is the closest matching multicast route based on the routing metrics.

You can configure IGMP router mode on multiple interfaces. For each interface, any version of IGMP can be configured. The default version is IGMP version 3.

**IGMP Proxy Mode**

The purpose of IGMP proxy mode is to act as a proxy for multicast hosts on IGMP router-mode interfaces. As a proxy, the IGMP proxy-mode interface forwards IGMP membership messages for the multicast hosts to a router that is running a multicast routing protocol.

IGMP proxy mode is designed to connect a Windows Server 2003-based router to a multicast-enabled IP internetwork that uses a multicast routing protocol, such as DVMRP, MOSPF, or PIM. The IGMP proxy-mode interface acts like a host and joins multicast groups on behalf of hosts on the IGMP router-mode interfaces. Multicast traffic that is sent to host members on IGMP router-mode interfaces is received on the IGMP proxy-mode interface and is forwarded by the IP multicast forwarding process. Multicast traffic that is sent by hosts on IGMP router-mode interfaces is flooded on the IGMP proxy-mode interface where a downstream IP multicast router can either forward the traffic or ignore it.

An interface configured for IGMP proxy mode performs the following functions:

- 2**Forwards IGMP Membership Messages.** The IGMP proxy-mode interface forwards all IGMP Host Membership Report messages and IGMP Leave Group messages received on the IGMP router-mode interfaces. The IGMP proxy-mode interface does not use the IP multicast forwarding table to forward these messages. It forwards the messages with a TTL value of 1.

- **Registers multicast MAC addresses.** For shared-access technologies such as Ethernet, the network adapter remains in unicast listening mode. For each multicast group that is registered by an IGMP Host Membership Report and forwarded on the IGMP proxy-mode interface, the IGMP router-mode interface adds the MAC address mapped from the multicast group address to the table of interesting destination MAC addresses on the network adapter for the proxy-mode interface. The network adapter then listens for any datagrams that are addressed to these addresses and passes them up to the operating system for further processing. The network adapter remains in unicast listening mode, unless there are too many MAC addresses to fit in the network adapter's table of interesting MAC addresses. In this case, the IGMP routing protocol component enables multicast promiscuous listening mode on the network adapter.

- **Maintains the IP multicast forwarding table.** The IP multicast forwarding table for the TCP/IP protocol is updated based on multicast datagrams received on both the IGMP router-mode interface and the IGMP proxy-mode interface. When non-local multicast traffic is received on an IGMP router-mode interface, entries are added to the IP multicast

forwarding table to forward the traffic across the IGMP proxy-mode interface. These entries cause all non-local multicast traffic that is received on an IGMP router-mode interface to be forwarded across the IGMP proxy-mode interface, regardless of whether there are downstream group members.

You can enable IGMP proxy mode only on one IGMP routing protocol component interface. Enable IGMP proxy mode on the interface that is attached to a subnet where there is a multicast router running multicast routing protocols.

**IGMP Router Mode vs. IGMP Proxy Mode**
The following table compares the functions of IGMP router mode and IGMP proxy mode.

### IGMP Router Mode and IGMP Proxy Mode

| Function | IGMP Router Mode | IGMP Proxy Mode |
|---|---|---|
| Listening mode | Multicast promiscuous mode. | Unicast listening mode. |
| IGMP router or host | Acts as an IGMP-based multicast router and listens for IGMP Host Membership Report messages. | Acts as an IGMP-based host by forwarding IGMP Host Membership Reports and responding to IGMP queries. Listens for IGMP Host Membership Report messages as a host, not as a router. |
| Updating the IP multicast forwarding table | Updates the IP multicast forwarding table based on IGMP traffic. | Updates the IP multicast forwarding table to flood (forward throughout the intranet) non-local multicast traffic that is received on IGMP router-mode interfaces. |
| IGMP queries | Sends IGMP queries to maintain an up-to-date forwarding table. | Sends no IGMP queries. |

**IGMP Router-Mode Processes and Interactions**
A Routing and Remote Access router running the IGMP routing protocol component with an interface configured for IGMP router mode maintains group membership information and the IP multicast forwarding table.

**Maintaining group membership information**
The process of managing group membership information for hosts on an IGMP router-mode interface is as follows:

1. When a host joins a multicast group, it sends an IGMP Host Membership Report message to the router on the local subnet.

2. The router updates the IP multicast forwarding table with the group address and the IGMP router-mode interface that received the message.

3. The IGMP routing protocol adds the multicast MAC address (mapped from the multicast group address, as described in "IP Multicasting Architecture") to the table of interesting destination MAC addresses on the network adapter for IGMP proxy-mode interface.

   **Note**

   - The **NDISRequest( )** function adds the multicast MAC address to the table of interesting addresses on the network adapter.

4. The router forwards the IGMP Host Membership Report message using the IGMP proxy-mode interface.

5. The neighboring multicast router receives the IGMP Host Membership Report message, updates its IP multicast forwarding table, and informs downstream multicast routers that a host member exists on the IGMP proxy-mode interface subnet.

6. Periodically (based on the Querier's Query Interval Code), the router elected as querier sends an IGMP Host Membership Query message. If the router receives a response from a multicast group member, the router updates the expiration time for the group and retains the group entry in the IP multicast forwarding table. If the router does not receive a response to the query, the router sets the expiration time for the group to 0 and removes the group entry from the IP multicast

forwarding table.

**Forwarding IP multicast traffic**
The process of forwarding multicasts from hosts on a router-mode interface works as follows:

1. A multicast host sends multicast traffic to a specific group address, and if the address is not local, the IGMP router-mode interface receives the traffic.

2. When the router receives the first datagram for a specific multicast group, the IGMP routing protocol component adds an entry to the IP multicast forwarding table to indicate that group members exist on the IGMP proxy-mode interface.

3. The multicast traffic passes to the multicast forwarding process, and, based on entries in the IP multicast forwarding table, the traffic is forwarded using the IGMP proxy-mode interface.

4. The neighboring multicast router receives the multicast traffic and passes it to the multicast forwarding process. The multicast router either forwards the traffic to group members or silently discards it based on the entries in its IP multicast forwarding table.

**Routing and Remote Access Support for Multicast Boundaries**
Routing and Remote Access supports multicast boundaries, as described earlier in this document. In addition to supporting scope-based boundaries and TTL-based boundaries, Windows Server 2003-based routers support boundaries defined by the rate of multicast traffic. With multicast rate limiting, multicast traffic forwarding can be restricted for traffic beyond a specified rate in kilobytes per second. You can configure multicast boundaries by using the General node under the IP Routing node in the console tree of the Routing and Remote Access snap-in.

**Multicast Heartbeat**
Multicast heartbeat is the ability of a server running Routing and Remote Access to listen for a regular multicast notification to a specified group address to verify that IP multicast connectivity is available on the network. If the heartbeat is not received within a configured amount of time, the server running Routing and Remote Access sets the multicast heartbeat status of the configured interface to inactive.

A polling mechanism can periodically check the multicast heartbeat status to detect when the multicast heartbeat is missing. For example, Simple Network Management Protocol (SNMP) can be used to send a message when the status becomes inactive.

Simple Network Time Protocol (SNTP) is frequently used for the multicast heartbeat. SNTP performs time synchronization and uses the reserved IP multicast address 224.0.1.1. If the SNTP server (the source of the heartbeat traffic) is strategically placed, the loss of the heartbeat indicates a problem with the IP multicast routing infrastructure. Windows Server 2003 includes an SNTP server called the Windows Time service (W32Time) and an SNTP client. For more information about SNTP, see RFC 1769.

You can configure the multicast heartbeat by using the **Multicast Heartbeat** tab of the properties for the General node under the IP Routing node in the console tree of the Routing and Remote Access snap-in.

**Supported Multicast Configurations**
The IGMP routing protocol component of Windows Server 2003 Routing and Remote Access is not designed as a substitute for multicast routing protocols such as DVMRP, MOSPF, or PIM. Therefore, this protocol works best for specific internetwork configurations of limited size. The recommended and supported configurations for a Windows Server 2003-based router using the IGMP routing protocol component are described in the following sections.

**Single-Router Intranet**
A Windows Server 2003-based router can provide full multicasting capabilities in a single-router intranet. The following figure illustrates a single-router intranet.

### Single-Router Intranet




In this configuration, each router interface is added to the IGMP routing protocol component and is configured for IGMP router mode.

Any host on any subnet can send multicast traffic to, and receive multicast traffic from, any other host. The Windows Server 2003-based router forwards all multicast traffic to subnets that contain group members.

**Single-Router Intranet Connected to the Internet**

A Windows Server 2003-based router supports multicasting in a single-router intranet that is connected to the MBone. The following figure illustrates this configuration.

**Single-Router Intranet Connected to the Internet**




In this configuration, all router interfaces are added to the IGMP routing protocol component. All private intranet interfaces are configured for IGMP router mode, and the Internet interface is configured for IGMP proxy mode.

As hosts join and leave multicast groups, they send IGMP membership messages to the local intranet router interface. These messages are then copied to the Internet interface. Multicast traffic from the Internet is sent to the Internet interface. When the Internet interface receives the multicast traffic, it forwards the traffic to the host on the appropriate subnet. Multicast traffic sent by a host on an intranet subnet is copied to the Internet interface. The multicast router at the ISP either ignores or forwards the multicast traffic, depending on whether there are any group members downstream.

If multicast traffic is sent between two hosts on the private intranet, the traffic is nevertheless copied to the Internet interface and is received by the multicast router at the ISP. You can prevent intranet multicast traffic from being copied to the Internet interface if you configure applications and MADCAP servers to use IP multicast addresses from the administratively scoped range (239.0.0.0 to 239.254.255.255) and configure the Internet interface with the appropriate scope-based boundary.

**Peripheral Router in a Multicast-Enabled Intranet**

A peripheral router is a router that is attached to multiple networks, one of which contains another router. In this case, the other router is a multicast router running a multicast routing protocol that is used in a multicast-enabled intranet. A peripheral router in a multicast-enabled intranet is a configuration that is very similar to the single-router intranet connected to the Internet. In this configuration, a Windows Server 2003-based router can act as the peripheral router.

The following figure illustrates a Windows Server 2003-based router acting as a peripheral router in a multicast-enabled intranet.

**Peripheral Router in a Multicast-Enabled Intranet**




In this configuration, all the peripheral router interfaces are added to the IGMP routing protocol component. All interfaces

except the one with the multicast router are configured for IGMP router mode. The interface connected to the subnet that contains the multicast router is configured for IGMP proxy mode.

When a host joins or leaves a multicast group, it sends an IGMP membership message. The message is copied to the interface that is configured for IGMP proxy mode. Multicast traffic from the multicast-enabled intranet is forwarded to the subnet attached to the IGMP proxy-mode interface. When the Windows Server 2003-based router receives the multicast traffic, it forwards the traffic to the host on the appropriate subnet. Multicast traffic that is sent by a host on an attached subnet is copied to the IGMP proxy-mode interface. The multicast router either ignores the multicast traffic or forwards it to downstream group members.

**Multicast Support for Remote Access Clients**
The Windows Server 2003 IGMP routing protocol component is commonly used to provide multicast services for dial-up or virtual private network (VPN) remote access clients. In this situation, the remote access or VPN server acts as a peripheral router to a multicast-enabled IP internetwork.

Two configurations are possible, depending on which of the following types of connectivity is provided to the remote access clients:

- Access to the MBone is provided by an ISP for dial-up clients.
- Access to a private multicast-enabled intranet is provided for dial-up or VPN clients.

**MBone access for ISP dial-up clients**
The following figure illustrates a configuration in which the Windows Server 2003 Routing and Remote Access service acts as an ISP to provide access to the Internet for dial-up clients.

### MBone Access for ISP Dial-Up Clients




In this configuration, the interface named **Internal**, which represents a logical adapter to which all the ISP clients connect, and the interface that is connected to the Internet are added to the IGMP routing protocol component. The **Internal** interface is configured for IGMP router mode, and the Internet interface is configured for IGMP proxy mode.

When a remote access client joins or leaves a multicast group, it sends an IGMP membership message. This message is copied to the interface that is connected to the Internet. Multicast traffic from the Internet is sent to the IGMP proxy-mode interface. When this interface receives the multicast traffic, it forwards the traffic to the connected remote access client. When a connected remote access client sends multicast traffic, the traffic is forwarded to other connected group members and is copied to the IGMP proxy-mode interface. The multicast router on the Internet either ignores the multicast traffic or forwards it to downstream group members.

**Intranet access for dial-up or VPN clients**
The following figure illustrates a configuration in which the Windows Server 2003 Routing and Remote Access service provides access to a private intranet for dial-up or VPN remote access clients.

### Intranet Access for Dial-Up or VPN Clients



In this configuration, the interface named **Internal** and the interface that is connected to the multicast-enabled intranet are added to the IGMP routing protocol component. The **Internal** interface represents all remote access clients. The intranet interface is attached to a subnet that contains a router running a multicast routing protocol, such as DVMRP, MOSPF, or PIM. The **Internal** interface is configured for IGMP router mode, and the intranet interface is configured for IGMP proxy mode.

When a remote access client joins or leaves a multicast group, it sends an IGMP membership message. This message is copied to the interface that is connected to the multicast-enabled intranet. Multicast traffic from the intranet is sent to the IGMP proxy-mode interface. When this interface receives the multicast traffic, it forwards the traffic to connected remote group members. When a remote access client sends multicast traffic, the traffic is forwarded to other connected group members and is copied to the IGMP proxy-mode interface. The multicast router on the intranet either discards the multicast traffic or forwards it to downstream group members.

**Multicast Support for Branch Office Networks**

A Windows Server 2003-based router can provide multicast forwarding capabilities for single-router branch offices connected to a hub office that has a multicast-enabled intranet. This configuration requires proper configuration of both the branch office and hub office router.

The following figure illustrates the use of a Windows Server 2003-based router to support multicasting for a branch office network.

**Multicast Support for Branch Office Networks**




For the branch office router, all the interfaces are added to the IGMP routing protocol component. The interfaces for the branch office subnets are configured for IGMP router mode, and the interface that connects to the hub office router is configured for IGMP proxy mode. The interface that connects to the hub office router can be a LAN interface (such as one that uses a T-carrier or frame relay connection) or a demand-dial interface (such as one that uses dial-up analog phone lines, ISDN, or a router-to-router VPN connection). A demand-dial interface can be either on-demand or persistent.

For the hub office router, all the interfaces are added to the IGMP routing protocol component. The interface that connects to the branch office router is configured for IGMP router mode. This interface can be a LAN interface (such as one that uses a T-carrier or frame relay connection) or a demand-dial interface (such as one that uses dial-up analog phone lines, ISDN, or a router-to-router VPN connection). The interface that connects to the multicast-enabled intranet is configured for IGMP proxy mode.

When a host joins or leaves a multicast group, it sends an IGMP membership message. The message is copied across the branch office link to the multicast-enabled intranet. Multicast traffic from branch office hosts is copied from the branch office subnets across the branch office link to the multicast-enabled intranet.

If multicast traffic is sent between two hosts on the branch office intranet, the traffic is copied to the branch office link, which results in inefficient use of bandwidth over the link to the hub office. You can configure applications and MADCAP servers to use IP multicast addresses from the administratively scoped range (239.0.0.0 to 239.254.255.255) to prevent local branch office multicast traffic from being copied to the branch office link. In this case, the interface to the hub office is configured with the appropriate scope-based boundaries.

**IP Multicasting Interfaces**

Windows Server 2003 provides application programming interfaces (APIs) that support the following types of development for multicasting:

- Multicast applications

- Reliable multicasting for multicast applications
- Multicast routing protocols

**Multicast Applications**
You can write multicast applications by using Windows Sockets, which is an API for developing applications that communicate by using TCP/IP. Windows Sockets 2 supports IGMP versions 2 and 3. For more information about how the Windows Sockets interface relates to TCP/IP, see "How TCP/IP Works."

For more information about developing multicast applications, see the Platform SDK Windows Sockets 2 documentation about Multicast Programming on MSDN.

**Reliable Multicasting for Multicast Applications**
PGM is an API for developing multicast applications that use a reliable multicast transport protocol rather than UDP. PGM, which is also known as reliable multicast, is implemented in Windows through Windows Sockets. You can use PGM to write multicast applications that have reliability requirements.

For more information about how to use reliable multicast in Windows, see the Microsoft Windows Server 2003 Family Platform SDK for Windows Sockets 2 documentation about Reliable Multicast Programming (PGM) on MSDN.

**Multicast Routing Protocols**
Routing and Remote Access is an extensible platform that supports non-Microsoft routing protocols. The Multicast Group Manager (MGM) API enables developers to write multicast routing protocols that work with routers running the MGM component of Routing and Remote Access.

For more information about developing multicast routing protocols with the MGM API, see the Routing and Remote Access Services SDK documentation about MGM Programming Considerations on MSDN.

**Related Information**
The following resources contain additional information that is relevant to this section.

**Related Technical Reference Subjects**

- IPv6 Technical Reference
- TCP/IP Technical Reference
- Unicast IPv4 Routing Technical Reference

**Related RFCs**
For more information about IP multicasting, see the following RFCs in the IETF RFC Database.

- RFC 2365, "Administratively Scoped IP Multicast."
- RFC 1075, "Distance Vector Multicast Routing Protocol."
- RFC 1112, "Host extensions for IP multicasting."
- RFC 2236, "Internet Group Management Protocol, Version 2."
- RFC 3376, "Internet Group Management Protocol, Version 3."
- RFC 1853, "IP in IP Tunneling."
- RFC 1469, "IP Multicast over Token-Ring Local Area Networks."
- RFC 2730, "Multicast Address Dynamic Client Allocation Protocol (MADCAP)."
- RFC 1584, "Multicast Extensions to OSPF."

- RFC 3208, “PGM Reliable Transport Protocol Specification.”
- RFC 2117, “Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification.”
- RFC 1769, “Simple Network Time Protocol (SNTP).”

# IPv4 Multicasting Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- IP Multicasting Tools
- IP Multicasting Registry Entries


This section describes the tools and registry entries used to configure Routing and Remote Access routers for multicasting and to troubleshoot multicast forwarding.

**IP Multicasting Tools**
The following tools are associated with IPv4 multicasting.

**Mcast.exe: Multicast Packet Tool**
**Category**
The Multicast Packet Tool is included with Windows Server 2003 Resource Kit tools.

**Version compatibility**
The Multicast Packet Tool runs on, and targets, Windows XP Professional and the Windows Server 2003 family.

The Multicast Packet Tool is a command-line tool that you can use to send multicast packets or to listen for packets being sent to a multicast group address. This tool is useful for testing multicast connectivity between computers on an internetwork.

To find more information about the Multicast Packet Tool, see "Resource Kit Tools Help" in the Tools and Settings Collection.

**Mrinfo.exe**
**Category**
The **Mrinfo.exe** tool is included with Windows 2000, Windows XP Professional, and Windows Server 2003.

**Version compatibility**
The **Mrinfo.exe** tool runs on Windows 2000, Windows XP Professional, and the Windows Server 2003 family. You can use this tool for troubleshooting multicasting on Windows 2000-based and Windows Server 2003-based routers.

This tool displays the configuration of a multicast router to help troubleshoot multicast forwarding problems. It sends an Internet Group Management Protocol (IGMP) query message to the specified router, which responds to the query with the version number, list of interfaces, list of neighbors on each interface, metrics, Time to Live (TTL) thresholds, and other data.

For more information about **Mrinfo.exe**, see "Command Line References" in the Tools and Settings Collection.

**Mtrace**
**Category**
Like Windows 2000, the Windows Server 2003 family does not provide a version of Mtrace, a multicast tracing tool.

**Version compatibility**
Although the Windows Server 2003 family does not provide a version of Mtrace, the server running Routing and Remote Access does respond to command queries from non-Microsoft Mtrace tools.

You can use an Mtrace tool for troubleshooting multicast forwarding problems. It traces a multicast path from the source to a receiver.

**Netsh Commands for Multicast Routing**
**Category**
The Netsh commands for multicast routing are included with the Windows Server 2003 operating system.

**Version compatibility**
The Netshcommands for multicast routing run on, and can target, the Windows Server 2003 family and Windows 2000 Server.

The Netsh tool provides command-line support for many of the multicast display and configuration options provided by the Routing and Remote Access Microsoft Management Console (MMC) snap-in.

The Netsh commands that support Routing and Remote Access server administration for multicasting are available from the

command prompt for the **netshrouting ip** context.

The command prompt for the **Netsh routing ip show** context provides support for displaying multicast forwarding table entries and statistics for multicast forwarding table entries.

The command prompt for the **Netsh routing ip igmp** context provides support for configuring specific interfaces and for displaying group members on a specific interface or on all-router interfaces.

For more information about **Netsh.exe**, see "Command Line References" in the Tools and Settings Collection.

**Rrasmgmt.msc: Routing and Remote Access Snap-In**
**Category**
The Routing and Remote Access snap-in is included with Windows 2000 Server and the Windows Server 2003 family.

**Version compatibility**
The Routing and Remote Access snap-in runs on Windows 2000, Windows XP Professional, and on the Windows Server 2003 family. You can use this tool to manage Windows 2000-based and Windows Server 2003-based routers.

The Routing and Remote Access snap-in configures and troubleshoots Routing and Remote Access multicast routers. For example, it supports the following:

- Adding interfaces to the router, including demand-dial interfaces
- Adding routing protocols
- Configuring for IGMP router mode or IGMP proxy mode
- Configuring IGMP router-mode settings
- Setting administrative scopes
- Configuring multicast heartbeats
- Setting IGMP event-logging levels

In addition, the snap-in displays the following:

- Multicast statistics (statistics compiled by IP for each multicast group)
- Multicast forwarding table (used by IP to forward IP multicast traffic)
- IGMP group table (group membership information for all groups registered on all IGMP router-mode interfaces)
- IGMP interface group table (group membership information for all groups registered on a specific IGMP router-mode interface)

**IP Multicasting Registry Entries**
The following registry entries are associated with multicasting.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_LOCAL_MACHINE\SYSTEM**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\.

**EnableMulticastForwarding**
**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\

**Version**
The **EnableMulticastForwarding** entry is supported by Windows 2000, Windows XP Professional, and the Windows Server 2003 family.

The **EnableMulticastForwarding** entry enables or disables the forwarding of IP multicast traffic.

**IGMPLevel**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\

**Version**
The **IGMPLevel** entry is supported by Windows 2000, Windows XP Professional, and the Windows Server 2003 family.

The **IGMPLevel** entry specifies the level of IP multicasting support for hosts. The possible levels are no support for the host, the host can only send multicast traffic, or the host can both send and receive multicast traffic.

To find more information about **IGMPLevel**, see "Registry Reference" in the Tools and Settings Collection.

**TrFunctionalMcastAddress**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\

**Version**
The **TrFunctionalMcastAddress** entry is supported by Windows 2000, Windows XP Professional, and the Windows Server 2003 family. IP multicast on Token Ring network adapters is not supported by Microsoft Windows NT Server 4.0 and earlier versions of Microsoft Windows.

The **TrFunctionalMcastAddress** entry controls whether multicast traffic on a Token Ring network is sent to the functional address 0xC0-00-00-04-00-00 or is sent as a MAC-level broadcast.

**HKEY_LOCAL_MACHINE\SOFTWARE**
The following registry entry is located under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing\IGMPV2\.

**EnableFileTracing**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing\IGMPV2\

**Version**
The **EnableFileTracing** entry is supported by Windows 2000, Windows XP Professional, and the Windows Server 2003 family.

The **EnableFileTracing** entry records detailed information about IGMP routing protocol component processes in the *systemroot*\Tracing\IGMPv2.log file. The **netsh ras tracing igmp enabled** command also enables IGMP tracing.

# NAT Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NAT Technical Reference**
Network address translation (NAT) provides a method for translating the Internet Protocol version 4 (IPv4) addresses of computers on one network into IPv4 addresses of computers on a different network. A NAT-enabled IP router deployed at the boundary where a private network, such as a corporate network, meets a public network, such as the Internet, allows computers on the private network to access computers on the public network by providing this translation service.

The impetus initiating the development of NAT technology was the need to provide a temporary solution to the IPv4 address-depletion problem faced by the Internet community. The number of available globally unique (public) IPv4 addresses is far too few to accommodate the rapidly increasing number of computers that need access to the Internet. Although the long-term solution — the development of Internet Protocol version 6 (IPv6) addresses — exists, IPv6 is not yet widely adopted. NAT technology lets computers on any network use reusable private addresses to connect to computers with globally unique public addresses on the Internet.

In the Microsoft Windows Server 2003 and Microsoft Windows 2000 Server operating systems, the Routing and Remote Access service supports a NAT solution by allowing the optional configuration of a routing protocol component that provides network address translation. Computers on a private network can access a public network by means of a NAT-enabled router that runs the Routing and Remote Access service. All traffic leaving, or entering, the private network must travel by way of the NAT-enabled router. The Windows Server 2003 Routing and Remote Access service improves on the NAT routing protocol component provided by Windows 2000 Server by adding support for a basic firewall, which helps protect the public interface of a computer that is configured to provide NAT, and by enabling connected remote access clients to access the Internet.

**In this subject**


- What Is NAT?
- How NAT Works
- NAT Tools and Settings

# What Is NAT?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section:**


- NAT Solutions
- NAT Connection Scenarios
- How NAT Differs from ICS
- Related Information


Windows Server 2003 provides network address translation (NAT) functionality as a part of the Routing and Remote Access service. NAT enables computers on small- to medium-sized organizations with private networks to access resources on the Internet or other public network. The computers on a private network are configured with reusable private Internet Protocol version 4 (IPv4) addresses; the computers on a public network are configured with globally unique IPv4 (or, rarely at present, Internet Protocol version 6 [IPv6]) addresses. A typical deployment is a small office or home office (SOHO), or a medium-sized business, that uses Routing and Remote Access NAT technology to enable computers on the internal corporate network to connect to resources on the Internet without having to deploy a proxy server.

The following figure is a simplified depiction of the Routing and Remote Access NAT service. When a client on the private network sends a request for information to a computer on the public network, the NAT-enabled router, located at the border between the private and the public networks, translates the outgoing packets and then sends them to the destination computer on the public network. The router also translates the response from the public network resource and returns the response to the client on the local network.

**NAT-Enabled Router Translating Request and Response Packets**




**NAT Solutions**

The Windows Server 2003 Routing and Remote Access implementation of NAT provides the following solutions to problems faced by small- or medium-sized networks:

- Reusing IP addresses
- Bundling NAT, DHCP, and DNS functionality for small networks
- Providing firewall protection to the NAT-enabled router
- Providing Internet access to remote access clients
- Providing PPPoE broadband Internet access

**Reusing IP Addresses**

Any computer with a direct (routed) connection to the Internet requires a public IP address from the Internet Assigned Numbers Authority (IANA). IANA allocates public addresses and guarantees them to be globally unique on the Internet. The limited and ever-decreasing availability of public IPv4 addresses is one of the most compelling problems facing the IP Internet. The major long-term solution to the problem of address depletion is the development of IPv6, which introduces a new

addressing model that uses 16-byte addresses expressed in colon-hexadecimal notation rather than the familiar four-byte IPv4 addresses typically expressed in dotted-decimal notation. Until the long-term IPv6 solution is widely in use, however, other methods to ensure that IPv4 addresses remain available are urgently required. One temporary solution that reduces the demand for IPv4 addresses is address reuse.

NAT, originally defined in RFC 1631 and extended in RFC 3022, was developed explicitly to provide a method that enabled an unlimited number of organizations to reuse private IPv4 addresses on their networks, thus substantially decreasing the demand for new public IPv4 addresses. This need to decrease the demand for public IPv4 addresses was the initial impetus for the creation of NAT technology.

Private networks can use any range of IP addresses for computers on their internal network. Prior to RFC 1918, published in February 1996, the typical practice was to assign public addresses to all hosts that use TCP/IP whether or not they needed access to the Internet. RFC 1918 describes the three ranges of IPv4 addresses that it recommends that organizations use for private networks. RFC 3022, published in January 2001, recommends that private networks that implement NAT in order to access the Internet use the private address space described in RFC 1918.

The private address space described in RFC 1918 consists of three sets of addresses reserved by IANA for use by private networks. These are:

- 10.0.0.0 through 10.255.255.255 (or 10.0.0.0/8 in Classless Inter-domain Routing [CDIR] notation)
- 172.16.0.0 through 172.31.255.255 (or 172.16.0.0/12)
- 192.168.0.0 through 192.168.255.255 (or 192.168.0.0/16)

IANA has designated these addresses as nonroutable, that is, networks that use these addresses cannot directly connect to the Internet (or other public network) through an Internet router. Instead, they need to access a router that supports NAT so that these nonroutable addresses can be translated into public addresses for routing over the Internet.

Thus, all interfaces connected to a private network with a Routing and Remote Access NAT-enabled router are assigned private IP addresses, including the NAT-enabled router itself, which has a private address on its private interface (LAN card) and one or more public addresses on at least one public interface (demand-dial connection or LAN card). The NAT-enabled router is the conduit through which a private network computer sends a request out to a public network and through which a response is received.

Network Address Port Translation (NAPT), described in RFC 3022, extends address translation by using many-to-one mapping when, as is typically the case, the number of internal addresses is greater than the number of public addresses available on the NAT-enabled router. The NAT-enabled router might have only one public address configured, or it might have a pool of public addresses. Using NAPT, it is possible to map many connections through a single public address by assigning each connection a different port number.

In addition to the conservation of IPv4 addresses, NAT technology provides the following additional incentives to choose reusable private addresses:

- Using IPv4 private addresses eliminates the need for an organization to obtain public IPv4 addresses from its ISP. This is an advantage because any organization that uses public addresses that then changes to a different ISP must obtain new public addresses for each computer from the new provider and then reconfigure the network accordingly. With NAT, an organization can use the same private addresses regardless of which ISP provides its Internet access. If such an organization decides to change ISPs, the only reconfiguration required is on the public interface of the NAT-enabled router, which will need a new public address from the new ISP.
- Using private IPv4 addresses, for which you incur no cost, is less expensive than buying a block of public IPv4 addresses from an ISP.
- Using private IPv4 addresses lets organizations continue to use IPv4 networks rather than migrate their networks to IPv6, a process that can be costly, especially for larger organizations, and that might require equipment upgrades or replacement as well as retraining for network administrators.

For a discussion of public versus private addressing in TCP/IP networks, see “Choosing Public or Private Addresses” in the chapter Designing a TCP/IP Network of the Windows Server 2003 Deployment Kit. For information about IPv6, including IPv6 addressing, see “Introducing IPv6 on Your Network” in the same chapter.

**Bundling NAT, DHCP, and DNS for Small Networks**
Organizations with small networks often do not need or want to expend resources setting up a full-fledged network infrastructure. For these organizations, the NAT/Basic Firewall component provided by the Windows Server 2003 Routing and Remote Access service supports NAT in conjunction with simplified versions of the Dynamic Host Configuration Protocol (DHCP), and Domain Name System (DNS) services.

| **Note** |
| --- |
| The name "NAT/Basic Firewall" indicates only that NAT is configured; it does not necessarily indicate that the optional Basic Firewall feature is also configured. You can configure Basic Firewall at the same time that you install NAT, or you can configure Basic Firewall later. You can deploy NAT with both Basic Firewall and another firewall technology, or you can use NAT with another firewall technology without configuring Basic Firewall. For more information, see "Providing Firewall Protection to the NAT-enabled Router" later in this section. |

In addition to the translation component, which translates IP addresses and TCP or UDP port numbers for packets sent between a private network and the Internet, the NAT/Basic Firewall component of the Routing and Remote Access service also provides (optionally) automatic addressing and name-resolution proxying:

- **DHCP Allocator.** The NAT/Basic Firewall automatic addressing capability. If enabled, the DHCP allocator feature allows the NAT-enabled router to act as a simplified DHCP server that automatically assigns to the computers on the private network an IP address, a subnet mask, a default gateway, and the IP address of a DNS server. If multiple routed subnets are configured, you must use a DHCP server rather than the DHCP allocator.

- **DNS Proxy.** The NAT/Basic Firewall automatic name-resolution capability. If enabled, the DNS proxy feature allows the NAT-enabled router to act as a DNS server for the computers on the private network. When the DNS proxy receives a DNS name resolution request, the DNS proxy sends a request to the Internet-based DNS server for which the NAT-enabled router is configured and then provides the results of the response to the private network computer from which the request originated.

For more information about the DHCP allocator and DNS proxy, see "Optional NAT Subsystems" in "How NAT Works".

**Providing Firewall Protection to the NAT-Enabled Router**
To computers on a public network, such as the Internet, all requests from a private network that has deployed NAT appear to come from a public IPv4 address on the NAT-enabled router. The private network IPv4 address scheme remains hidden. By using a NAT, the computers on the private network gain some measure of protection because the NAT-enabled router does not forward traffic from the Internet into the private network unless a private network client had requested it or unless specific traffic is explicitly allowed. Although NAT technology lets organizations that implement it hide their private network topology from the Internet behind the NAT-enabled router, the Internet interface (also known as the public interface) of the NAT-enabled router is not hidden from the Internet. Without firewall protection, the Internet interface of the NAT-enabled router is vulnerable to attack.

The Windows Server 2003 Routing and Remote Access NAT component includes a simple stateful firewall called Basic Firewall, which uses technology similar to that used by the Internet Connection Firewall (ICF) feature available in Windows XP, Windows XP SP1, and Windows Server 2003. A stateful firewall uses dynamic packet filtering to examine incoming and outgoing packets. This new integration of Routing and Remote Access NAT with Basic Firewall lets network administrators protect the Internet interface of any computer running a member of the Windows Server 2003 family that is acting as a NAT. By enabling the optional Basic Firewall feature on the Internet interface (or multiple public interfaces) of the NAT-enabled router, an administrator ensures that all packets that are received on the public interface that do not correspond to traffic requested by the NAT computer (either for itself or for private network clients), or that are explicitly allowed, are discarded.

For more information about Basic Firewall, see the "Optional NAT Subsystems" section in "How NAT Works".

**Providing Internet Access to Remote Access Clients**
A computer running a member of the Windows 2000 Server family might be configured to provide remote access to its private network — for traveling employees or employees working from home, for example. The same server might also be configured as a NAT-enabled router to provide access to the Internet for computers within the private network.

Although you cannot configure a computer running Windows 2000 Server to provide Internet access to connected remote access clients by using the Routing and Remote Access snap-in, you can do so by using the command-line Netsh tool. You can use a Netsh command to add the **Internal** interface, which is a logical interface that represents the connection to all remote access clients, as a private interface to the NAT/Basic Firewall component of the Routing and Remote Access service. This configuration allows connected remote access clients to access the Internet.

Computers running a member of the Windows Server 2003 family now allow you to use the Routing and Remote Access snap-in to add the **Internal** interface as a private interface to the NAT/Basic Firewall component to enable connected remote access clients to access the Internet. With Windows Server 2003, you can also still use Netsh to accomplish the same task.

| **Note** |
| --- |
| When you enable the Routing and Remote Access service on a computer running Windows Server 2003, an interface named **Internal** appears under **Routing Interfaces** in the Routing and Remote Access snap-in. Do not delete the **Internal** interface. |

**Providing PPPoE Broadband Internet Access**

The number of multiple-computer households is increasing at a rapid rate. One efficient option for such households, as well as small businesses, is to use Point-to-Point Protocol over Ethernet (PPPoE) to establish a single broadband connection (such as through a DSL line or cable modem) to an ISP. PPPoE makes high-speed Internet access available to all the computers in the home or small office.

The Windows Server 2003 Routing and Remote Access service introduces the use of PPPoE as a connection type for demand-dial connections. This functionality enables a home or small business to use a NAT-enabled router and the PPPoE broadband Internet connection to connect their private network to the Internet.

**NAT Connection Scenarios**

Scenarios in which you might deploy Windows Server 2003 Routing and Remote Access NAT include connecting a business or home to the Internet, or using NAT on one side of two connected geographically remote offices.

**Connecting a Business or Home to the Internet**

In a typical deployment that uses Windows Server 2003 Routing and Remote Access NAT, a small- or medium-sized business (or home office) assigns private IPv4 addresses to its computers and installs a Routing and Remote Access NAT-enabled router. The computers on the private network cannot gain access to the Internet directly because they have reusable private addresses. Instead, they gain access to the Internet through the NAT-enabled router.

The NAT-enabled router has both a private interface and a public interface, the public IPv4 address of the latter provided by a local ISP. When a client computer sends a request out to a resource on the Internet, the request is sent initially to the NAT-enabled router, which translates the request packets, forwards them to the Internet resource, accepts responses from the Internet resource, re-translates the packets, and then returns the response to the client.

Another example in which NAT technology can provide access to the Internet is the case of a small ISP that serves home users who have dial-up connections. In this case, the NAT-enabled router is located at the ISP. Whenever the customer dials in, the ISP assigns a private IPv4 address to the customer's computer. When the customer requests access to a server on the Internet, the NAT-enabled router at the ISP translates the outgoing request and, later, the incoming response.

**Using NAT on One Side of Two Remote Offices**

Another deployment option is to use NAT on one side (typically, the smaller office) of a connection that links offices in different geographical locations. The Windows Server 2003 Routing and Remote Access service provides two types of virtual private network (VPN) site-to-site connections. The following table describes the circumstances in which you can use a NAT in conjunction with a VPN connection.

### Locating a Network Using a VPN Behind a NAT

| Type of VPN Site-to-Site Connection | Can You Use NAT? | Description |
| --- | --- | --- |
| PPTP VPN | Yes | In most cases, you can locate PPTP–based calling routers behind a NAT-enabled router (or configure one computer as both the calling router and the NAT-enabled router) in order to allow computers with private addresses in a SOHO network to share a single connection to the Internet. With a VPN connection, the site-to-site connection from the small office to the main office is "tunneled" through the Internet. Windows Server 2003 Routing and Remote Access NAT includes a NAT editor that can accurately translate PPTP-tunneled data. |

| | | |
|---|---|---|
| L2TP/IPSec VPN | Yes, but only if you use the IPSec NAT Traversal (NAT-T) feature | With Windows Server 2003–based calling or answering routers, you can use the IPSec feature called NAT traversal (NAT-T) to create L2TP/IPSec connections across NATs. Using NAT-T requires running Windows Server 2003 on both the calling and answering routers (or appropriately configured Cisco routers). With NAT-T, computers with private addresses that are hidden behind a NAT can use IPSec to connect to a remote site if these computers have the NAT-T update installed (for computers running Windows XP Professional SP1). No NAT editor exists for L2TP/IPSec, so the only way to use NAT is by implementing IPSec NAT-T. |

For more information about site-to-site VPN connections, see the chapter "Connecting Remote Sites" in the Networking collection of the "Windows Server 2003 Deployment Kit".

For more information about IPSec NAT-T, see "End-to-End Traffic Through an ISA-Secured Network Address Translator" in "What Is IPSec?" in the Security collection of this technical reference.

For more information about NAT editors provided by the Routing and Remote Access service, see "How NAT Works" in the Networking collection of this technical reference.

**How NAT Differs from ICS**

Internet Connection Sharing (ICS) provides a network translation capability that is an alternative option to that provided by Routing and Remote Access NAT. ICS, which is typically used by networks that have two to 10 computers, is provided by Network Connections on Windows XP as well as on Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition. Because ICS and Routing and Remote Access NAT share common drivers, they cannot coexist on the same network.

With ICS, as with Routing and Remote Access NAT, you can connect computers on a SOHO network to the Internet by using just one connection. Although configuring Routing and Remote Access NAT is simple, configuring ICS is simpler — a single configuration step (selecting a single check box) provides a translated connection to the Internet for all of the computers on a single-subnet private network.

ICS and Routing and Remote Access NAT share the following features:

- **Address translation.** Both ICS and Routing and Remote Access NAT translate inbound and outbound IPv4 traffic by modifying IPv4 addresses and ports.

- **Public and private interfaces on the computer that connects to the network.** Both ICS and Routing and Remote Access NAT require two interfaces on the computer that provides address translation in order to work: one public and one private. The private interface, typically a LAN adapter or the **Internal** interface, connects the ICS or Routing and Remote Access NAT computer to the computers on the private network. The public interface, typically a DSL, cable, or dial-up modem, connects the private network to the Internet.

- **Firewall protecting the Internet interface.** Windows XP, Windows XP SP1, and Windows Server 2003 include ICF and Windows Server 2003 Routing and Remote Access NAT includes Basic Firewall, to protect the public interface of the workstation, server, or NAT router that provides Internet access to the other computers on the private network.

- **DHCP Allocator.** Both ICS and (optionally) Routing and Remote Access NAT can use this simplified DHCP service, which assigns the IP address, subnet mask, default gateway, and DNS server on the private network.

- **DNS Proxy.** Both ICS and (optionally) Routing Remote Access NAT can use this simplified DNS service, which resolves DNS names on behalf of private network clients and forwards queries.

The most significant differences between Routing and Remote Access NAT and ICS NAT are the following:

- **SOHO or larger network.** ICS is designed specifically for SOHO networks. Routing and Remote Access NAT scales from homes or small organizations to medium-sized networks.

- **Computer connection to the Internet.** With ICS, the connection to the Internet is managed most often by a Windows XP workstation, although you can also configure ICS on a computer running Windows 2000 Server or Windows Server 2003. With Routing and Remote Access NAT, the connection to the Internet or other public network is managed by a NAT-enabled router running the Windows Server 2003 or Windows 2000 Server Routing and Remote Access service.

- **Configuration limits.** ICS allocates addresses for internal hosts automatically from an address range that is not configurable (192.168.0.0/24) and supports only one public interface obtained from an ISP through which network computers can access the Internet. Routing and Remote Access NAT lets you access the Internet (or other public network) through multiple interfaces and lets you choose your own addressing scheme.

- **DNS and DHCP configuration.** DNS and DHCP configuration are similar, but not identical, for ICS and Routing and Remote access NAT:

    - **ICS NAT.** Typically, ICS is used in a network with two to 10 computers. Because a network this small might not have DNS servers or DHCP servers, ICS provides the DHCP allocator and the DNS proxy to provide automatic address configuration and name resolution for private network hosts. If, however, a DHCP server is deployed on a network that uses ICS, the ICS DHCP allocator yields to the DHCP server. Similarly, if the Microsoft DNS service is installed on the same computer as the DNS proxy, the ICS DNS proxy yields to the DNS server. On a private network that uses ICS, if a DNS server is deployed on a separate computer than the computer on which ICS is installed, either a DHCP server must also be installed on a different computer than the computer on which ICS is installed, or all clients on the private network must be manually configured with the IP address of the DNS server.

    - **Routing and Remote Access NAT.** Typically, Routing and Remote Access NAT is used in any network from small to medium in size. Routing and Remote Access NAT also provides the DHCP allocator and the DNS proxy for use in smaller networks that do not install a DNS server or a DHCP server. If a network that uses Routing and Remote Access NAT has a DHCP server in place, do not enable (or disable, if necessary) the DHCP allocator on the Routing and Remote Access NAT-enabled router. The DNS proxy can work in conjunction with a DNS server. For more information about how the DHCP allocator and the DNS proxy work in Routing and Remote Access NAT, see "Optional NAT Subsystems" in "How NAT Works."

- **Number of network segments.** ICS supports only a single private subnet. Routing and Remote Access NAT can support multiple private subnets. A network with multiple subnets must use a DHCP server rather than the DHCP allocator.

For more information about ICS and NAT, see Internet Connection Sharing and network address translation in Help and Support Center for Windows Server 2003 and Internet Connection Sharing Overview.

**Related Information**

For more information about NAT, see the following RFCs in the IETF RFC Database:

- RFC 1631, "The IP Network Address Translator."

- RFC 3022, "Traditional IP Network Address Translator (Traditional NAT)."

- RFC 1918, "Address Allocation for Private Internets."

- RFC 2663, "IP Network Address Translator (NAT) Terminology and Considerations."

The following resources contain additional information that is relevant to this section:

- "Choosing Public or Private Addresses" and "Introducing IPv6 on Your Network" in the chapter Designing a TCP/IP Network of the Windows Server 2003 Deployment Kit.

- "Connecting Remote Sites" in the Windows Server 2003 Deployment Kit.

- "End-to-End Traffic Through an ISA-Secured Network Address Translator" in "What Is IPSec?"

- "Network Address Translation (NAT)" in *Internetworking with TCP/IP, Vol. 1: Principles, Protocols, and Architecture, Fourth Edition*, by Douglas E. Comer, Prentice Hall, New Jersey, 2000.

# How NAT Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- NAT Architecture
- NAT Network Structure
- Optional NAT Subsystems
- NAT Processes and Interactions
- Related Information


The network address translation (NAT) functionality provided by the Routing and Remote Access service enables computers on a private network to access computers on a public network, such as the Internet. In Windows Server 2003, the Routing and Remote Access NAT/Basic Firewall component is typically referred to as a routing protocol component. Strictly speaking, NAT is not a routing protocol; it does not send and receive traffic nor does it instigate changes to the routing table. However, the Routing and Remote Access service interfaces with the NAT/Basic Firewall component through the routing protocol interface, and, in this sense, NAT functions just like a routing protocol. Although NAT is not a true routing protocol, it is installed as if it were by using the Routing and Remote Access snap-in.

This document explains how Routing and Remote Access NAT works. It covers the components of Routing and Remote Access NAT architecture, a typical network behind a NAT-enabled router, optional NAT subsystems designed to simplify configuration for smaller networks, and NAT processes and interactions. The latter section explains how NAT works with TCP/IP, how address and port translation work, how dynamic and static mappings and IP reservations work, how NAT editors work, and includes a flowchart showing how inbound and outbound packets are processed.

For an introductory description of the problems Routing and Remote Access NAT is designed to solve and a quick sketch of typical NAT scenarios, see "What Is NAT?." For information about tools used to install, configure, and manage Routing and Remote Access NAT, see "NAT Tools and Settings."

**NAT Architecture**

Understanding the architecture of Routing and Remote Access NAT requires recognizing the relationship of the NAT routing protocol component to other routing components, to the NAT editors and proxies, to TCP/IP, and to the TCP/IP Address Resolution Protocol (ARP) module. The following figure illustrates Routing and Remote Access NAT architecture:

**Routing and Remote Access NAT Architecture**




The following table describes each component depicted in the figure.

**Routing and Remote Access NAT Components**

| Component | Description |
|---|---|
| IP Router Manager | A component of the Routing and Remote Access service. The IP Router Manager uses input/output (I/O) controls to do the following:<br>• Add interfaces to and remove interfaces from NAT.<br>• Provide NAT with its configuration, which is stored as part of the IP global and per-interface information in the registry. |
| NAT | A routing protocol component that, when enabled on a server running the Routing and Remote Access service, translates IPv4 addresses and TCP or UDP port numbers of request packets originating from a client on a private network, forwards the translated packets to a destination computer on the Internet (or other public network), and then performs reverse translation for response packets sent by the destination computer back to the client. NAT can also translate packets for specifically allowed incoming traffic originating from a computer on the Internet to a client computer on the private network.<br>For more information, see "NAT Network Structure" and "NAT Processes and Interactions" later in this document. |
| NAT Editors | Programs that perform translation for the data of certain network protocols that NAT on its own cannot translate. Routing and Remote Access NAT has the following built-in editors:<br>• **Internet Control Message Protocol (ICMP)**, which supports the transmission of error messages across a NAT.<br>• **Point-to-Point Tunneling Protocol (PPTP)**, which supports PPTP VPN traffic across a NAT.<br>For more information, see "How NAT Editors Work" in "NAT Processes and Interactions" later in this document. |
| NAT Proxies | Programs that redirect packets to a module that builds a new packet modified so that the newly modified packet can cross a NAT.<br>Routing and Remote Access NAT includes the File Transfer Protocol (FTP) proxy, which supports FTP traffic across a NAT, and the H.323 proxy, which supports NetMeeting calls across a NAT. |

| | |
|---|---|
| TCP/IP | The industry standard suite of protocols for transmitting data between computers over networks. Routing and Remote Access NAT is implemented as a routing protocol component that the TCP/IP protocol component in Windows calls when a data packet is sent or received. TCP/IP allows NAT to make modifications to a packet's data (translating IP addresses and TCP or UDP port numbers) and to ensure that checksums within a modified packet are correct.<br><br>For more information, see "How NAT Works with TCP/IP" and "How Address and Port Translation Work" in "NAT Processes and Interactions" later in this document. |
| ARP | The TCP/IP network layer protocol that uses broadcast traffic on the local network to resolve a logically assigned IPv4 address to its physical hardware or media access control (MAC) layer address. When mapping between public and private address pools when a pool of public addresses exists, NAT uses TCP/IP's proxy ARP mechanism to respond to ARP queries for addresses in a public address pool. (Proxy ARP refers to the use of a particular machine, such as a NAT-enabled IP router, to respond to ARP requests for computers other than itself.) |

**NAT Network Structure**

In a typical scenario, a Routing and Remote Access NAT-enabled router provides access to the Internet for a small- or medium-sized organization that uses private IPv4 addresses. To support this type of network, in addition to a server running the Routing and Remote Access service with the NAT/Basic Firewall component enabled, the typical network configuration includes, at minimum, client computers running a Web browser, such as Internet Explorer, and a firewall, such as the Basic Firewall option available with Routing and Remote Access NAT. The NAT-enabled router lets the clients on the private network access a computer on the Internet, such as a Web server running Internet Information Services (IIS).

Unless the network contains multiple routed segments, neither a DHCP nor a DNS server is required because NAT functionality in Windows Server 2003 includes the DHCP allocator and the DNS proxy capabilities, which provide automatic addressing and name resolution, respectively. For more information about the DHCP allocator and the DNS proxy, see "DHCP Allocator" and "DNS Proxy" in "Optional NAT Subsystems" later in this document.

If a DHCP server or a DNS server does exist on the network, Routing and Remote Access NAT is also designed to work with those servers. In addition, Routing and Remote Access NAT is designed to work in networks with or without the Active Directory directory service.

The following figure shows an example of a single-subnet private network that uses a Routing and Remote Access NAT-enabled router to provide access to the Internet. Computers on the private network have a private IPv4 address, except for the router, which has both a private and a public address. The Web server, like any computer on the Internet, has a public IPv4 address.

**Client on a Private Network Behind a Routing and Remote Access NAT Accessing a Resource over the Internet**




The following table describes each component shown in the preceding figure.

**Components of a Private Network Behind a NAT-enabled Router**

| Component | Description |
|---|---|
| Client PCs | Each client has a private IPv4 address configured on its network adapter. Clients need only a Web browser and access to a NAT-enabled router to be able to access the Web server across the Internet. |

| | |
|---|---|
| Routing and Remote Access NAT-enabled router | The NAT-enabled router has a private IP v4 address configured on its private interface, and a public IPv4 address configured on its Internet interface (in this example, a dial-up modem). The router has the optional Basic Firewall, DHCP allocator, and DNS proxy NAT components enabled. For more information, see "Optional NAT Subsystems."<br><br>For a larger network, an administrator might also configure a more sophisticated firewall in addition to using Basic Firewall to protect the Internet interface of the NAT-enabled router and might use a DHCP server (required if the network has more than one segment) and a DNS server. |
| Web server | A Web server on the Internet provides resources needed by the client computers on the private network. Like all computers on the Internet, it has a public IPv4 address. |

For information about why NAT was developed to mediate between private and public networks, see "What Is NAT?" in this collection.

**Optional NAT Subsystems**

One major type of organization that deploys Routing and Remote Access NAT is the small office or home office (SOHO) customer. Routing and Remote Access NAT has three optional sub-components — Basic Firewall, DHCP allocator, and DNS proxy — that are designed specifically to provide vital services while simplifying configuration for smaller-sized networks.

**Basic Firewall**

In Windows Server 2003, the Routing and Remote Access NAT routing protocol component is now integrated with Basic Firewall in order to provide protection to the NAT-enabled router's Internet interface. Basic Firewall is a simple stateful packet filter that prevents unsolicited traffic originating from a public network, such as the Internet, from reaching a private network through the public interface of the NAT-enabled router. Although enabling the Basic Firewall feature is optional, enabling it is a recommended practice even if your network has other firewall software installed. Enabling Basic Firewall helps secure the NAT-enabled router and strengthens overall network security by protecting the private network from unwanted Internet traffic.

Basic Firewall, which combines stateful packet filtering of network traffic with a set of static packet filters, monitors traffic that travels through the interface for which Basic Firewall is enabled. For an interface configured for private network traffic and to provide NAT, each packet's source and destination addresses are recorded in a table and all incoming traffic from the public network is compared to the entries in that table. Traffic from a public network can reach the private network only if the table contains an entry that shows that the communication exchange is allowed.

Windows Server 2003 provides two methods to configure the Basic Firewall that is now integrated with Routing and Remote Access NAT:

- If you use the Routing and Remote Access Server Setup Wizard to enable NAT, you can enable the Basic Firewall option during the same setup procedure.

- Alternatively, if the Routing and Remote Access service is already installed with NAT enabled, you can use the Routing and Remote Access snap-in to enable Basic Firewall at any time.

  **Note**

  - In addition to the Basic Firewall included in the Windows Server 2003 Routing and Remote Access service to protect the public interface of a NAT router, Windows XP, Windows XP SP1, and Windows Server 2003 include Internet Connection Firewall (ICF) to protect the public interface of any server or workstation that provides Internet access to computers on a private network. ICF, which provides a firewall technology similar to Basic Firewall, can be configured by using Network Connections.

For more information about installing and configuring Routing and Remote Access NAT and its optional components, see "NAT Tools and Settings" in this collection.

**TCP or UDP Unicast Packets**

If Routing and Remote Access NAT has Basic Firewall configured, an incoming TCP or UDP unicast packet is dropped unless it meets one of the following conditions:

- The packet belongs to a previously established mapping, or meets the UDP loose source port criteria (see "How Dynamic Mapping, Static Mapping, and IP Reservations Work" in "NAT Processes and Interactions" later in this document).

- The packet is a DHCP response packet, that is, UDP with source port 67 and destination port 68.

- The packet matches a static port mapping (see "How Dynamic Mapping, Static Mapping, and IP Reservations Work" in "NAT Processes and Interactions" later in this document).

- The packet matches a ticket (dynamic or editor-created). A ticket, in this context, is a statically defined port exemption that allows inbound packets that match this ticket.

- The packet matches a redirect. A redirect changes the destination address and port of packets that match the redirect to the address or port specified in the redirect. In Routing and Remote Access NAT, a redirect is used either by the NAT FTP editor to enable FTP traffic across a NAT, or it is used by NAT deployed in combination with the Microsoft Internet Security and Acceleration (ISA) Server to provide high-security Internet access.

**Broadcast and Multicast Packets**
If Routing and Remote Access NAT has Basic Firewall configured, the firewall always accepts broadcast and multicast packets and passes them to the NAT component. However, on a computer running the Windows Server 2003, Windows XP SP1, or later operating system, NAT is set by default to drop all inbound broadcast and multicast packets. To change this default behavior, an administrator can add the following registry key and set it to **1** to allow inbound non-unicast traffic, that is, to allow broadcast and multicast packets across Basic Firewall:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IpNat\Parameters\AllowInboundNonUnicastTraffic**

By default, **AllowInboundNonUnicastTraffic** is set to **0**, which blocks inbound non-unicast traffic.

**Specially Defined Packets**
If Routing and Remote Access NAT has Basic Firewall configured, the firewall always accepts specially defined packets for the following IP protocols, which are already protected by IPSec, and passes them to the NAT component:

- IP protocol 50 (or 0x32 in hexadecimal) for Internet Protocol security Encapsulating Security Payload (IPSec ESP), an IPSec protocol that provides confidentiality in addition to authentication, integrity, and replay.

- IP protocol 51 (or 0x33 in hexadecimal) for IPSec Authentication Header (IPSec AH), a header that provides authentication integrity, and anti-replay.

**ICMP Traffic**
The ICMP protocol, an extension to the Internet Protocol (IP), enables the generation of error messages, test packets, and informational messages related to IP. ICMP is described in RFC 792.

If Routing and Remote Access NAT has Basic Firewall configured, the following firewall behavior for ICMP traffic is defined:

- Inbound echo, timestamp, mask, and router request packets are dropped.

- Inbound echo, timestamp, mask, and router replies are dropped unless there exists a mapping automatically created by a corresponding outbound request.

- Inbound time exceeded, parameter problem, destination unreachable, and source quench are always forwarded.

- Outbound time exceeded, parameter problem, destination unreachable, and source quench are always dropped.

- Redirects are always dropped.

An administrator can configure the ICMP behavior to allow message types that are dropped by default (for example, to allow an echo request).

**DHCP Allocator**
In Windows Server 2003, the DHCP allocator feature of the Routing and Remote Access NAT routing protocol component provides a simplified form of automatic addressing. Windows Server 2003 (as well as Windows 2000 and Windows NT) supports automatic address acquisition by using the DHCP service. On starting up, a DHCP client obtains its addressing information from a DHCP server automatically, including its default gateway and DNS server. However, DHCP server configuration is complex and might be inappropriate for many SOHO networks. In response to this need, Windows, beginning

with Windows 98, also supports Automatic Private IP Addressing (APIPA), in which client computers, when no DHCP server is available, automatically configure themselves with private addresses in the range 169.254.0.1 to 169.254.255.254 with a subnet mask of 255.255.0.0 (or 169.254.0.0/16 in Classless Inter-domain Routing [CDIR] notation). Computers using APIPA addresses can communicate only with other computers that are also using APIPA addresses within a single subnet but cannot connect to computers outside the subnet. This approach works, therefore, only for strictly local networks, but does not cover the case where a local network has a link to the Internet.

The DHCP allocator feature provided by Routing and Remote Access NAT provides a solution that is a combination of the two approaches just described. The DHCP allocator is disabled by default. If you enable it, the Routing and Remote Access NAT-enabled router will function as a simplified DHCP server, that is, it will automatically assign IPv4 addresses to local clients and advertise itself to its DHCP clients as both the default gateway and the DNS server to the Internet. Specifically, the DHCP allocator assigns the following to each client on the private network:

- An IPv4 address and subnet mask.
- The IPv4 address of a default gateway (the private address of the NAT-enabled router).
- The IPv4 address of a DNS server (the private address of the NAT-enabled router; the router will forward a name resolution request that it receives from a client on its private network to the Internet–based DNS server for which the router is configured).

You must configure computers on the private network as DHCP clients in order to automatically receive the IP configuration.

Because the DHCP allocator provided by Routing and Remote Access NAT maintains no database, it does not require additional configuration beyond enabling the DHCP allocator feature. In Routing and Remote Access NAT, the DHCP allocator randomly allocates addresses from the range of configured addresses. For each address generated for a new client, the DHCP allocator uses the TCP/IP ARP module to ascertain whether that address is in use by another client. If a client replies to the ARP message indicating that it is using that address, a new address is generated. Each address assignment has a lease time of 10 minutes to avoid duplicate addresses that might be used by unavailable computers.

**Note**

- Configuring the NAT DHCP allocator feature works well when all computers on the local network are connected to the same subnet. On the same subnet, computers share each other's broadcasts and, therefore, can resolve each other's names by using broadcast NetBIOS name queries. If the network has multiple subnets, you must use a DHCP server rather than the DHCP allocator capability of the NAT/Basic Firewall component.

The DHCP allocator allocates addresses only to clients on permanently connected interfaces (such as LAN media) that are configured as private. You can configure this feature at either of two points:

- If you use the Routing and Remote Access Server Setup Wizard to enable NAT and you choose a network interface with a static IPv4 address as the private interface, you are given the option toenable the DHCP allocator feature for addressing services and the DNS Proxy feature for name resolution services. (For more information about "DNS Proxy," see "DNS Proxy," described next.)
- Alternatively, if the Routing and Remote Access service is already installed with NAT enabled, you can use the Routing and Remote Access snap-in to enable the DHCP allocator at any time.

For more information about installing and configuring Routing and Remote Access NAT and its optional components, see "NAT Tools and Settings" in this collection.

**DNS Proxy**

In Windows Server 2003, the DNS proxy feature of the Routing and Remote Access NAT routing protocol component provides a simplified name resolution service. Windows Server 2003 (as well as Windows 2000 and Windows NT) supports the assignment of DNS server addresses to DHCP clients during DHCP address assignment. This scheme fails in the case of a SOHO where the DNS server resides either with an ISP or on a corporate network because the correct DNS server address cannot be known until after a connection is established. Moreover, the typical SOHO will not have a DHCP server configured.

To solve this problem, a Routing and Remote Access NAT-enabled router that has both the DHCP allocator (described earlier) and the DNS proxy features configured provides its own address as the DNS server address to DHCP clients on the local

network. Then, when the connection to the Internet is established, the NAT-enabled router proxies DNS name resolution requests to the Internet DNS server address assigned to the Internet interface of the router. The DNS proxy makes the name resolution process transparent from the point of view of local computers. Like the DHCP allocator, the DNS proxy is disabled by default.

The following list explains how DNS proxy works:

1. The private computer sends a DNS name query request to the DNS proxy (source: private computer IP address, destination: DNS proxy private IP address).

2. DNS proxy saves the DNS name query request in memory and sends a new DNS name query request (source: DNS proxy public IP address, destination: Internet DNS server IP address).

3. The Internet DNS server resolves the name and sends a DNS name query response (source: Internet DNS server IP address, destination: DNS proxy public IP address).

4. DNS proxy matches the received DNS name query response with the saved DNS name query request and sends a new DNS name query response to the private computer (source: DNS proxy private IP address, destination: private computer IP address).

To the private computer, the DNS proxy is a DNS server. To the Internet DNS server, the DNS proxy is a DNS client.

Another possible scenario in a small private network in which no DHCP server exists is to enable the DNS proxy and disable the DHCP allocator. In this case, for each computer on the private network, a static IP address, subnet mask, and default gateway (the IP address of the NAT-enabled router) is configured by using the Network Connections folder to access the TCP/IP properties page for the network adapter on the private network. In addition, on the NAT-enabled router, the address of a DNS server on the Internet is configured as the DNS server for the router by using the TCP/IP properties page on the public interface of the NAT-enabled router. This configuration enables the NAT-enabled router to pass DNS name resolution requests to the Internet DNS server for which it is configured. For more information, see "Using Network Connections to Configure NAT Clients on the Private Network" in "NAT Tools and Settings."

In a setting where the local network already has a DNS server configured, the DNS proxy can still be useful. In all likelihood, an existing DNS server is configured to forward Internet name resolution requests to a static Internet DNS server address. This static configuration will fail whenever the address of the Internet DNS server changes. With automatic name resolution provided by the DNS proxy feature, you can configure the existing DNS server to forward Internet DNS requests to the local address of the DNS proxy, and the DNS proxy will then handle the task of dynamically finding the Internet DNS server address when the Internet connection is established.

Using DNS proxy provides an adequate solution even if multiple routed segments exist. In that event, you must disable the DHCP allocator (if it is enabled), configure a full DHCP server, and then configure the DHCP server to give the NAT-enabled router's local address on each subnet as the DNS server address for clients on that subnet. However, if multiple subnets are configured, the recommended option is to automatically install a full DNS server in default mode, which provides caching and name registration in addition to the name resolution.

**NAT Processes and Interactions**

This section covers the following topics:

- How NAT works with TCP/IP
- How address and port translation work
- How dynamic and static mappings and IP reservations work
- How NAT editors work
- How mapping and translation work together

**How NAT Works with TCP/IP**

Routing and Remote Access NAT is implemented as a routing protocol component that TCP/IP calls when a packet is sent or received. For the benefit of modules such as firewalls, the TCP/IP driver supports a pointer that sets the callout that TCP/IP

invokes for incoming and outgoing data packets. The parameters of the TCP/IP callout include the IP header and data and the context for the incoming or outgoing interface. The function to which TCP/IP calls out indicates whether the data packet should be forwarded, dropped, or processed as local host traffic. TCP/IP also allows the function to provide a replacement for the data packet, to accommodate any modifications that the NAT driver might make to the packet.

When mapping between public and private addresses when a pool of public addresses exists, NAT uses TCP/IP's Proxy ARP mechanism to respond to ARP queries for public addresses in the pool.

In the translation process, the NAT routing protocol component makes routing decisions in order to determine the interface over which to forward any given packet. To avoid consulting TCP/IP for each packet, a cache is implemented in the NAT. To keep the cache consistent with the TCP/IP forwarding database, the NAT registers for notification of route changes with TCP/IP.

**How Address and Port Translation Work**

The Routing and Remote Access NAT-enabled router, located, logically, at the edge point between the private and public networks, mediates between its private network and the resources on the Internet that its users want to access. To do so, the router uses both address and port translation, as described in this section.

The NAT driver modifies the source address and port in the header of TCP or UDP packets sent from a client on the private network to the public network, while also modifying the destination address and port in packets returned from the public network to the private network. Specifically, for an application for which the IPv4 address and port information is located in the IP and TCP or UDP headers (the type of application for which NAT was originally designed), the NAT-enabled router can translate the following:

- IPv4 address in the IP header
- Port number (either a TCP port number in the TCP header, or a UDP port number in the UDP header)

For example, traffic between a private network and the Internet — Hypertext Transfer Protocol (HTTP) traffic used to access Web servers — requires the translation of the IPv4 address in the IP header and, possibly, the TCP port in the TCP header. Some applications, however, embed address or port information inside the data portion of the packet or do not use TCP or UDP ports to identify the data stream. Translation for these applications is more complicated, and for such translation to take place, an appropriate NAT editor must exist. For more information about NAT editors, see "How NAT Editors Work" later in this section.

**Outgoing and Incoming Packet Translation**

When a computer user on the private network connects to a resource on the Internet (or other public network), the computer's TCP/IP protocol creates an IP packet with values for the source address and port as well as for the destination address and port. This information is set in the IP header and in the TCP header (or UDP header), as shown in the center column of the following table. The source computer (or an intermediate router) forwards this packet to the NAT-enabled router. The router then translates the source address and port of the outgoing packet, as shown in the right column of the same table. Information modified by this translation is shown in the table in italic text:

**Example of Translation of Address and Port of Outgoing Packet**

| Address and Port | Outgoing Packet Header When Client Initiates Request | Outgoing Packet Header after NAT Translates Source Address and Port |
|---|---|---|
| Source IPv4 Address | *Private IPv4 address* | *ISP-allocated public IPv4 address* [1] |
| Source Port | *Source application TCP or UDP port* | *Remapped source application TCP or UDP port* |
| Destination IPv4 Address | Internet resource IPv4 address | Internet resource IPv4 address |
| Destination Port | Internet resource TCP or UDP port | Internet resource TCP or UDP port |

[1] This example is a case in which the public address on the Internet interface of the NAT-enabled router is allocated by an organization's ISP

The NAT-enabled router then sends the remapped IP packet over the Internet. When the Internet resource sends back a response to the NAT, the incoming packet contains the information shown in the center column of the next table, "Example of Translation of Address and Port of Incoming Packet." After the NAT maps and translates the destination address and port, it contains the information as shown in the right column of that table. The NAT-enabled router forwards the translated packet to

the intranet client that made the initial request. Information modified by this translation is shown in the following table in italic text.

**Example of Translation of Address and Port of Incoming Packet**

| Address and Port | Incoming Packet Header When Internet Resource Sends Response | Incoming Packet Header After NAT Translates Destination Address and Port |
|---|---|---|
| Source IPv4 Address | Internet resource IPv4 address | Internet resource IPv4 address |
| Source Port | Internet resource TCP or UDP port | Internet resource TCP or UDP port |
| Destination IPv4 Address | *ISP-allocated public IPv4 address* [1] | *Private IPv4 address* |
| Destination Port | *Remapped source application TCP or UDP port* | *Source application TCP or UDP port* |

[1] This example is a case in which the public address on the Internet interface of the NAT-enabled router is one allocated by an organization's ISP

For outgoing packets, the source IPv4 address and TCP or UDP port numbers are mapped to a public source IPv4 address and a possibly changed TCP or UDP port number. For incoming packets, the destination IPv4 address and TCP or UDP port numbers are mapped to the private IPv4 address and original TCP or UDP port number.

**Network Address Port Translation**

All computers on the private network behind the NAT-enabled IP router, including the router itself, have a unique private IPv4 address and a private port number. The router also has at least one globally unique (public) IPv4 address. Optionally, an administrator can configure the router with a pool of public addresses so that multiple clients can request public network resources at the same time. The router can handle traffic for multiple clients on its private network even when it does not have a sufficient number of addresses in its public address pool to accommodate all of the clients requesting access to the public network. It can accommodate additional clients by using Network Address Port Translation (NAPT) to modify port information in the TCP or UDP packets being transmitted. RFC 3022 describes NAPT.

NAPT extends address translation by using many-to-one mapping — multiple private addresses are mapped to a single public address on the Internet interface of the NAT-enabled router. This is helpful because the number of clients with private addresses is ordinarily far larger than the available number of public addresses on the NAT-enabled router, even if the router is configured with a relatively large number of addresses in a public address pool. When the public addresses in the pool have all been translated, NAT starts to perform port translation instead of address translation.

For TCP or UDP communications, the address and port translation functionality made available by NAPT permits a client on a private network to access multiple computers on the Internet, and it enables multiple private clients to access the same Internet computer at the same time. Because both TCP and UDP ports have 16 bits, NAPT allows up to 65,535 ($2^{16}$-1) communications to take place at the same time. One is subtracted from the total number because port 0 for both TCP and UPD is reserved.

**Address and Port Translation Example**

How NAPT works is best understood through an example. Consider the simple case of a private network with three computers — Client A, Client B, and Client C — that use the private network IPv4 range of 192.168.0.0 with a subnet mask of 255.255.255.0. The NAT-enabled router is configured with a public pool containing the following IPv4 addresses:

157.54.35.38

157.54.35.39

The user on Client A initiates a request to look at a Web page on the Internet, so Client A sends a TCP packet to the NAT-enabled router. The NAT driver translates the packet and adds an entry for it to the NAT Mapping Table. The following figure depicts the packet translation and the entry made into the mapping table for Client A:

**Packet Translation and NAT Mapping Table Entry for Client A**

| NAT Mapping Table |
|---|
| **192.168.0.10: 3576 ← TCP → 157.54.35.38: 5000 (Dynamic Address Translation)** |

Next, the user on Client B initiates another TCP session with a host on the public network. The following figure depicts the packet translation for Client B and the new entry for Client B, which is added to the mapping table that already contains the entry for Client B.

**Packet Translation and NAT Mapping Table Entry for Client B**




| NAT Mapping Table |
|---|
| 192.168.0.10: 3576 ← TCP → 157.54.35.38: 5000 (Dynamic Address Translation) |
| **192.168.0.11: 2258 ← TCP → 157.54.35.39: 5000 (Dynamic Address Translation)** |

Finally, the user on Client C initiates traffic from the private to the public network and thus sends a packet to the NAT-enabled router. However, the public address pool configured on the NAT-enabled router contains only two addresses, which have already been used for Clients A and B. In this case, the NAT driver translates Client C's private address:port pair to a unique public address:port pair by changing the public port number. That is, when the pool of public addresses, which is fewer in number than the number of clients on the private network, is exhausted, the NAT driver switches from address translation to port translation. The following figure depicts the packet translation for Client C and the new entry for Client C in the mapping table.

**Packet Translation and NAT Mapping Table Entry for Client C**

| NAT Mapping Table |
|---|
| 192.168.0.10: 3576 ← TCP → 157.54.35.38: 5000 (Dynamic Address Translation) |
| 192.168.0.11: 2258 ← TCP → 157.54.35.39: 5000 (Dynamic Address Translation) |
| **192.168.0.12: 1944 ← TCP → 157.54.35.39: 5001 (Dynamic Port & Address Translation)** |

When the resource on the public network to which each client sent a request sends a response back to the client on the private network, the same process occurs in reverse: The NAT driver intercepts the response packet, looks up the packet's destination address in the NAT Mapping Table, finds the corresponding address of a private network client, modifies the destination address and destination port in the response packets — this time translating the packet's public destination address and port number back to the client's private address and port number — recomputes the checksum (which it must do any time it replaces an address in a packet header), and then sends the response across the private network to the client.

**How Dynamic and Static Mappings and IP Reservations Work**

In addition to the dynamic mapping just introduced in "How Address and Port Translation Work," Routing and Remote Access NAT technology provides for two additional types of entries in the NAT Mapping Table: static mapping and IP reservations. Each of the three mapping types is designed for a different purpose.

**Dynamic Mapping**

Whenever a NAT-enabled router receives an outgoing TCP or UDP packet from a local client that is initiating communications with a computer on the public network, the NAT driver creates an entry in the NAT Mapping Table, called a 5-tuple entry, that contains the following five pieces of information:

{protocol (TCP or UDP), private address, private port, public address, public port}

An alternative notation representing this 5-tuple entry is the following:

{protocol (TCP or UPD), source address, source port, destination address, destination port}

The private address and port pair are automatically (dynamically) mapped to the public address and port pair. This is the type of mapping illustrated earlier in "How Address and Port Translation Work." This entry in the NAT Mapping Table enables the NAT-enabled router to direct a response packet from a computer on the public network back to the client on the private

network that sent initial the request.

Because the number of mappings that can be established is limited by the number of available 16-bit TCP and UDP ports, the NAT driver must eventually delete the dynamic mappings that it creates in order to free up port numbers for use in new mappings. A dynamic mapping entry remains in the mapping table for the length of time that the administrator specifies on the **Translation** tab for the properties of the **NAT/Basic Firewall** component in the Routing and Remote Access snap-in. Routing and Remote Access NAT, by default, uses the RFC 1631 recommended timeouts of 24 hours for idle TCP mappings and 1 minute for idle UDP mappings.

**Static Mapping**
With dynamic mapping, in order for a TCP or UDP packet to pass through the NAT from the public network to the private network, a mapping must have been established previously by a packet sent from the private to the public network. However, if a client on the public network attempts to establish a TCP or UDP session with a computer on the private network (such as a Web server on the private network or a computer on the private network that provides a game application), no dynamic mapping will be found and the incoming packet will be discarded.

If an organization wants to allow such incoming traffic to a specific computer on the private network, Routing and Remote Access NAT allows an administrator to use the **Services and Ports** tab on the properties page of the public interface in the Routing and Remote Access snap-in to configure a static mapping for this traffic. Thus, the way that the NAT-enabled router forwards Internet traffic into its private network is either in response to traffic initiated by a user on the private network, which creates a dynamic mapping, or because an administrator has configured a static mapping to enable Internet users to access specific resources on the private network. (For an alternative to static mapping, see "IP Reservations" later in this section).

Static mappings have limited usefulness for connections between a private network and the Internet because of the large number of possible connections. Too many connections can make the NAT Mapping Table grow excessively and thus slow router performance.

A static mapping consists of a 5-tuple entry identical in content to a dynamic mapping entry:

{protocol (TCP or UDP), private address, private port, public address, public port}

Unlike a dynamic mapping, a static mapping explicitly matches a given TCP or UDP port number to both the private and the public address in the static entry in the mapping table. For example, to set up a Web server on a computer on a private network, an administrator could create a static mapping that maps [*Public IPv4 Address*, TCP Port 80] to [*Private IPv4 Address*, TCP Port 80]. Port 80 is, by convention, assigned to the World Wide Web (WWW) and used by Web servers for HTTP traffic. When a packet arrives for *Public IPv4 Address* and Port 80, the NAT-enabled router directs the packet to the Web server on the private network.

Any server running on a well-known port on the private network (such as HTTP or FTP servers) can have only one instance accessible to clients running on the public network.

Multiple interfaces can have translation enabled at the same time, as in the case of a home where a user connects both to the Internet and to a corporate network. Therefore, static mappings in Routing and Remote Access NAT can be configured on a per-interface basis.

A static mapping entry remains in the mapping table until an administrator deletes it.

**UDP Source Port Allocation and Loose Source Matching**
To better support various types of peer-to-peer applications, the NAT mapping behavior for UDP differs from that of TCP in the following two ways:

**How NAT chooses the source port for outbound dynamic mappings.**
When creating a new TCP mapping for an outbound packet, the NAT driver chooses a source port without regard for already existing mappings as long as such a choice does not result in a conflict. In contrast, when choosing a source port for a UDP mapping for an outbound packet, the NAT driver determines if a mapping exists that has the same private address and port. If such a mapping exists, the NAT driver will use the same public port for the new mapping. For example:

- If a client on the private network makes a TCP connection to two different computers on the public network from the same source port, the NAT driver will choose different source ports for those mappings.

- If a client on the private network sends UDP packets to two different computers on the public network from the same source port, the NAT driver will use the same source port for both mappings.

**How NAT determines whether an inbound packet matches an existing dynamic or static mapping.**
For TCP, an inbound packet must exactly match the 5-tuple for a mapping (that is, protocol, source address, source port,

destination address, and destination port). For UDP, however, an inbound packet must match only the protocol, destination address, and destination port of a mapping — the source address and source port of the packet are effectively ignored. This "loose matching behavior" applies only if the private port is greater than 1024. Allowing this behavior for ports below 1024 would introduce a security risk because it might allow unfettered access to such sensitive TCP and UDP ports as 137 (NetBIOS Name service) and 445 (Microsoft Common Internet File System [CIFS]).

**IP Reservation**

If you want to allow a computer on the Internet to initiate a connection to a computer on the private network behind the NAT-enabled router, an alternative to using static mapping is to configure an IP reservation to handle this traffic. If you establish a public IPv4 address reservation to one of the private IPv4 addresses on the private network, all incoming traffic addressed to the specified public address is sent to the private address reserved for it. Thus, instead of putting a computer directly on the public network, the traffic is transmitted through the NAT-enabled router. For both static mapping and IP reservations, this is an advantage because, typically, the router is configured with a firewall and other filters.

For example, you might want to use an IP reservation to allow traffic to a Web server on the private network. You can allow such incoming traffic by first using the **Address Pool** tab on the properties page of the public interface in the Routing and Remote Access snap-in to configure an IP address pool (you cannot configure an IP reservation until after you configure at least one IP address pool). You can then configure the IP reservation itself by using the **Reservations** button on the **Address Pool** tab.

IP reservations are particularly useful when the number of public IP addresses on the NAT-enabled router is large — reservations provide an easy way to map public to private addresses. In addition, IP reservations are useful for IP protocols that pick port numbers randomly, which means that the administrator does not need to know the port numbers in advance. IP reservations can also be used for IP protocols that do not use ports because this process is independent of the protocol (whether TCP, UDP, or ICMP). For an example showing how IP reservations work, see "Static Mapping and IP Reservation Examples" later in this document.

For outgoing traffic from a computer on the private network for which an IP reservation exists, the NAT-enabled router does not use the public IP address reserved for that private computer. Outgoing traffic from a computer on the private network to any destination on the Internet, including traffic to a Web server for which an IP reservation exists, is handled by the NAT-enabled router by using standard address and port translation.

When a network administrator uses the Routing and Remote Access snap-in to create an IP reservation, a 2-tuple entry is entered in the NAT Mapping Table:

{private address, public address}

An IP reservation remains in the mapping table until an administrator deletes it.

**Dynamic and Static Mappings More Restrictive than IP Reservations**

The behavior of dynamic and static mapping entries and IP reservations is determined by the amount of information each stores:

- 5-tuple dynamic and static mappings are more restrictive than IP reservations because incoming or outgoing traffic must match all five pieces of information in the NAT Mapping Table entry before the packet will be forwarded; or, in the case of UDP for ports greater than 1024, an inbound packet must match three pieces of information — the protocol, destination address, and destination port.

- 2-tuple IP reservations are less restrictive than either dynamic or static mappings because incoming or outgoing traffic that has a destination IPv4 address that matches the public IPv4 address in the NAT Mapping Table entry will be forwarded.

**Static Mapping and IP Reservation Examples**

To continue the example presented earlier in "How Address and Port Translation Work," which illustrates dynamic mapping, an administrator next decides to use Client B as a Web server and therefore uses the Routing and Remote Access snap-in to add a static mapping for Client B. This mapping means that all inbound TCP traffic directed to 157.54.35.38 and port 80 on the NAT-enabled router will match that address:port pair to Client B's own address and port 80. Thus, the NAT Mapping Table shown in the earlier figure, "Packet Translation and NAT Mapping Table Entry for Client C," changes to include a second entry for Client B (the new fourth row), as shown in the following table.

**Static Mapping of Client B's Private Address and Port to a Public Address and Port**

| Client | Private Address:Port | Protocol | Public Address:Port | Type |
|---|---|---|---|---|

| | | | | |
|---|---|---|---|---|
| A | 192.168.0.10:3576 | TCP | 157.54.35.38:5000 | Dynamic address translation |
| B | 192.168.0.11:2258 | TCP | 157.54.35.39:5000 | Dynamic address translation |
| C | 192.168.0.12:1944 | TCP | 157.54.35.39:5001 | Dynamic port and address translation |
| B | 192.168.0.11:80 | TCP | 157.54.35.38:80 | Static address translation |

If the administrator then uses the Routing and Remote Access snap-in to add an IP reservation to the NAT Mapping Table, reserving public IPv4 address 157.54.35.39 for Client A, any existing dynamic mappings to 157.54.35.39 are preempted by this change. That is, the earlier dynamic mappings to the public address 157.54.35.39 no longer exist. Only the IP reservation for 157.54.35.39 remains in the NAT Mapping Table, as shown in the following table.

**IP Reservation of Client A's Private Address to a Public Address**

| Client | Private Address:Port | Protocol | Public Address:Port | Type |
|---|---|---|---|---|
| A | 192.168.0.10:3576 | TCP | 157.54.35.38:5000 | Dynamic address translation |
| B | 192.168.0.11:80 | TCP | 157.54.35.38:80 | Static address translation |
| A | 192.168.0.10:* | * | 157.54.35.39.* | IP reservation address translation |

With an IP reservation, as the asterisks in the last row of the preceding table indicate, it is irrelevant which protocol (TCP or UDP) is used or what the destination port number is on an inbound packet. All that matters is that the destination address is matched to the public IPv4 address in the reservation. Therefore, in the example in the preceding table, any traffic with a destination IPv4 address of 157.54.35.39 is forwarded to client A.

**Summarizing Packet Translation**

To summarize outbound and inbound packet translation:

- When a TCP or UDP packet outbound to the Internet from a client on the private network arrives at the private interface of the NAT-enabled router, the NAT driver looks in the NAT Mapping Table for an existing dynamic mapping. If none is found, the NAT driver creates a new mapping, updates the checksum, and sends the packet out to the Internet over a dial-up, broadband, or other connection.

- When a TCP or UDP packet inbound from the Internet arrives at the Internet interface of the NAT-enabled router, the NAT driver looks in the NAT Mapping Table for an existing dynamic or static mapping or for an IP reservation. If any of the three types of entries is found, the NAT driver translates the incoming packet and delivers it over the private interface of the NAT-enabled router to the client. Otherwise, the packet is discarded.

**How NAT Editors Work**

In Routing and Remote Access NAT, the NAT driver is designed to translate IPv4 address and TCP or UDP port information for traffic between a private and a public network. If the IPv4 address and port information is located only in the IP and TCP or UDP headers — as is the case, for example, for HTTP (World Wide Web) traffic — the NAT driver can translate the application protocol transparently. However, the NAT driver, on its own, cannot translate packets in the following two types of cases:

**An IPv4 address, TCP port, or UDP port is stored in the payload.**

Certain applications embed address or port information inside the data portion of the packet rather than in the IP and TCP or UDP headers. One such application is the File Transfer Protocol (FTP) — for the FTP **port** command, FTP stores the IPv4 address in the FTP header. In addition, FTP stores the IPv4 address in dotted decimal format, which means that the translated IPv4 address in the FTP header can be a different size, and, therefore, the TCP sequence numbers must be modified to ensure that no data is lost. Other examples of applications that the NAT driver cannot translate transparently include streaming applications, such as RealAudio and CUSeeMe, which use a TCP connection to dynamically select a UDP port over which data is then transmitted.

**TCP or UDP is not being used to identify the data stream.**

Data that is tunneled from one geographic site to another through the Internet by using the Point-to-Point Tunneling Protocol (PPTP) does not use a TCP or UDP header. Instead, PPTP data uses a Generic Routing Encapsulation (GRE) header, and the Call ID, which is stored in the GRE header, identifies the data stream. If the NAT driver does not correctly translate the Call ID within the GRE header, connectivity problems can occur.

NAT editors are application level gateways (ALGs) that have been developed to translate packets in these two cases. A NAT editor examines the packets sent by such applications and modifies the data portion of the packet as appropriate. Because such

modifications might change the size of packets, a NAT editor must also adjust sequence numbers for TCP connections after any translation has changed the number of bytes in the TCP data stream.

The Windows Server 2003 Routing and Remote Access service installs two built-in NAT editors along with the NAT routing protocol component — Internet Control Message Protocol (ICMP) to support error messages, and Point-to-Point Tunneling Protocol (PPTP) to support PPTP VPN traffic. Both editors are always on. The NAT driver consults the editors whenever the payload of the packet being translated matches one of the installed editors. The editor modifies the payload and returns the result to the NAT driver.

### Note

- In addition to the NAT editors, Routing and Remote Access NAT includes two NAT proxies, which redirect packets to a module that builds a new packet modified so that it can cross a NAT. The NAT proxies are the File Transfer Protocol (FTP) proxy, which supports FTP traffic across a NAT, and the H.323 proxy, which supports NetMeeting calls across a NAT. Unlike NAT editors, which are always on, NAT proxies are disabled by default. Both proxies can be turned on or off by using the appropriate Netsh command in the **netsh routing ip nat context**: you can enable the NAT proxies by using **add ftp** or **add h323**; you can disable the NAT proxies by using **delete ftp** or **delete h323**.

#### How Mapping and Translation Work Together

The two flowcharts in this section, "Processing for Outbound and Inbound NAT Traffic" and "Packet Translation Process," outline the overall process that the Routing and Remote Access NAT-enabled router uses to manage outbound and inbound packets.

First, the router determines whether the packet is outbound or inbound. Depending on the outcome of that initial determination, the following occurs:

#### Outbound traffic

For traffic from the private network that is outbound to the Internet (or other public network), the NAT driver assesses whether or not an address/port mapping (whether a static or dynamic mapping, or an IP reservation) exists for the packet. If such a mapping exists, the NAT driver uses the existing mapping; if none exists, the NAT driver creates a new dynamic mapping. Whether a single IPv4 address or a pool of addresses is available determines how that mapping is made (see "How Address and Port Translation Work" earlier in this section). Next, the NAT driver checks for editors and invokes one if necessary (assuming that one exists). If no editor is needed, the NAT driver performs standard network address translation of the TCP or UDP header and the IP header (for an outbound packet, translating the source TCP or UDP port and the source IP address), and then refreshes the NAT Mapping Table. Finally, the NAT-enabled router uses its Internet interface to forward the outgoing packet to the resource computer on the Internet.

#### Inbound traffic

For a response or other specifically allowed incoming traffic from the Internet (or other public network) that is inbound to the private network, the NAT driver also assesses whether or not an address/port mapping (whether static or dynamic mapping, or an IP reservation) exists for the packet. If a mapping exists, the NAT driver uses the existing mapping; if no mapping exists, the incoming packet is discarded. Next, the NAT driver checks for editors and invokes one if necessary (assuming that one exists). If no editor is needed, the NAT driver performs standard network address translation of the TCP or UDP header and the IP header (for an inbound packet, translating the destination TCP or UDP port and the destination IP address), and then refreshes the NAT Mapping Table. Finally, the NAT-enabled router uses its private interface to forward the incoming packet to the appropriate computer on the private network.

Of the following two flowcharts, the second is a continuation of the first. In the first flowchart, "Processing for Outbound and Inbound NAT Traffic," the two process boxes (one on the left and another on the right) labeled "Performs packet translation" both point to the second flowchart, "Packet Translation Process."

### Processing for Outbound and Inbound NAT Traffic




**Packet Translation Process**




**Related Information**

For more information about NAT, see the following RFCs in the IETF RFC Database:

- RFC 1631, "The IP Network Address Translator (NAT)."
- RFC 3022, "Traditional IP Network Address Translator (Traditional NAT)."
- RFC 1918, "Address Allocation for Private Internets."
- RFC 2663, "IP Network Address Translator (NAT) Terminology and Considerations."

The following resource contains additional information that is relevant to this section:

- "Network Address Translation (NAT)" in *Internetworking with TCP/IP, Vol. 1: Principles, Protocols, and Architecture, Fourth Edition*, by Douglas E. Comer, Prentice Hall, New Jersey, 2000.

# NAT Tools and Settings

Updated: July 31, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- NAT Tools
- NAT Registry Entries
- Related Information


You can use the network address translation (NAT) tools and registry settings described here to enable, configure, and manage Routing and Remote Access NAT on a computer running Windows Server 2003.

**NAT Tools**
The following tools are associated with the NAT routing protocol component provided by the Microsoft Windows Server 2003 Routing and Remote Access service:

- Routing and Remote Access snap-in
- Network Connections
- Netsh command-line tools for Routing and Remote Access NAT

**Graphical User Interface Tools**
The graphical user interface tools used to install and configure Routing and Remote Access NAT include Network Connections, used to configure TCP/IP properties for NAT clients to provide the client computers access to the Internet, and the Routing and Remote Access snap-in, used to install and configure Routing and Remote Access NAT on a server.

**Network Connections Category**
Network Connections is included with the versions of the Windows operating systems as described in the next paragraph, "Version compatibility."

**Version compatibility**
You can run Network Connections on any computer running Windows Server 2003, Windows XP, Windows 2000, or Windows NT version 4.0. However, Windows NT 4.0 does not support Routing and Remote Access NAT.

**Network Connections options for NAT clients**
You use the Network Connections tool to configure TCP/IP properties for NAT clients so that the client computers can use a NAT-enabled router to gain access to the Internet (or other public network).

The properties page for TCP/IP on client computers can be used for a variety of purposes. The following table lists which TCP/IP options are used to enable NAT clients to interact with a NAT-enabled router on a private network.

**TCP/IP Properties Options Used for a NAT Client**

| Page | Tab | Option |
|---|---|---|

| | | |
|---|---|---|
| TCP/IP Properties | General | The following options on the adapter of the NAT client are used to enable the NAT client to access the NAT-enabled router:<br>1. Under **Use the following IP address**, the values needed for following options are specific to a NAT client:<br>• **IP address.** The private IPv4 address and subnet mask of the client.<br>• **Subnet mask.** The private subnet mask of the client.<br>• **Default gateway.** The private IPv4 address of the NAT-enabled router.<br>2. Under **Use the following DNS server addresses**, the value needed for the following option is specific to a NAT client:<br>**Preferred DNS server.** The private IPv4 address of the NAT-enabled router, which is acting as a DNS proxy. |
| TCP/IP Properties | Alternate Configuration | No NAT-related configuration is needed on this tab. |
| Advanced TCP/IP Settings | IP Settings | **Default gateways.** A gateway using the private address of the NAT-enabled router (if appropriate).<br>**Caution:** This option is appropriate to use only if the client computer receives its IP address from a DHCP server (that is, this option is not appropriate if the client obtains its IP address from the NAT DHCP allocator). |

**Routing and Remote Access Snap-in**
**Category**
The Routing and Remote Access snap-in is used with the Routing and Remote Access service, which is included with Windows Server 2003 and Windows 2000 Server. Routing and Remote Access is disabled by default. You can use the Routing and Remote Access snap-in, under **Administrative Tools**, to enable and configure the Routing and Remote Access service, including the optional NAT routing protocol component.

**Version compatibility**
The Routing and Remote Access snap-in is provided by the Routing and Remote Access service on computers running Windows Server 2003 and Windows 2000 Server. For Windows NT 4.0, the Routing and Remote Access Service (RRAS) and its snap-in are available as a separate download from the Microsoft Windows NT Server Routing and Remote Access Service Download page at http://go.microsoft.com/fwlink/?LinkId=22441. However, Windows NT 4.0 RRAS does not include NAT.

The following sections briefly summarize NAT-related tasks for which you can use the Routing and Remote Access snap-in. (In addition to NAT, the Routing and Remote Access snap-in is also used to configure LAN routing, dial-up or VPN remote access connections, and site-to-site connections between geographically remote networks.)

- Enabling NAT while running the Routing and Remote Access wizard
- Enabling both VPN and NAT while running the Routing and Remote Access wizard

**Enabling NAT while running the Routing and Remote Access Wizard**
If the Routing and Remote Access service is not yet enabled on a computer running Windows Server 2003, you can enable NAT when you run the Routing and Remote Access Wizard by selecting the **Network Address Translation (NAT)** option. The wizard also lets you choose to enable the Basic Firewall feature.

Choosing **Network Address Translation (NAT)** when you run the wizard establishes the following:

- Configures the IP address of the private network interface (the LAN card that connects to the private network segment).
- Configures the public interface. If the connection is a non-permanent connection (such as a dial-up modem), the wizard creates a demand-dial interface to the ISP and creates a default static route that uses the Internet interface. (If the

connection is permanent, such as DDS, T-Carrier, Frame Relay, permanent ISDN, xDSL, or cable modem, the wizard does not create a demand-dial interface or static route for the interface.)

- Adds the NAT routing protocol component.
- Adds Internet and private network interfaces to the NAT routing protocol component.
- If you chose the option to enable **Basic Firewall** while running the wizard, the wizard configures a basic stateful firewall on the public interface connected to the Internet.

**Note**

- If the network already has a firewall and you do not select the **Basic Firewall** option while running the wizard, the Routing and Remote Access snap-in entry for the NAT routing protocol component (under **IP Routing** in the console tree) displays as **NAT/Basic Firewall**. The name "NAT/Basic Firewall" does not indicate whether Basic Firewall is configured.
- You can confirm whether Basic Firewall is configured by using the **NAT/Basic Firewall** tab on the properties page of the public (Internet-connected) interface.

**Enabling both VPN and NAT while running the Routing and Remote Access Wizard**

If the Routing and Remote Access service is not yet enabled on a computer running Windows Server 2003, you can configure the server both to provide NAT for the private network and also to accept VPN connections. You can do so when you run the Routing and Remote Access Wizard by selecting the **Virtual Private Network (VPN) access and NAT** option when the wizard begins.

Choosing **Virtual Private Network (VPN) access and NAT** specifies that computers on the Internet cannot determine the IP addresses of any computer on the private network, yet allows VPN clients to connect to computers on the private network.

**Routing and Remote Access snap-in options for NAT-enabled routers**

If the Routing and Remote Access service is already enabled on a server, or if you installed NAT by using the Routing and Remote Access Wizard and want to modify the NAT configuration, you can use the tools provided by the Routing and Remote Access snap-in to enable and configure, or modify, Routing and Remote Access NAT.

The Routing and Remote Access snap-in can be used for a variety of purposes unrelated to Routing and Remote Access NAT. The following table lists which options under the **General** and **NAT/Basic Firewall** nodes in the Routing and Remote Access snap-in are used for NAT-related tasks and describes the location in the Routing and Remote Access snap-in used for each task.

**NAT-related Options in the Routing and Remote Access Console Tree**

| Node | Task |
|---|---|
| General | Adding network address translation:<br>• Under the server name for the server to be configured as the NAT-enabled router, expand **IP Routing**, right-click **General**, select **New Routing Protocol**, and then choose **NAT/Basic Firewall**. |

| | |
|---|---|
| NAT/Basic Firewall | Adding and configuring public or private interfaces for the NAT routing protocol component:<br>• Right-click **NAT/Basic Firewall**, and then click **New Interface** to add an internal interface to connect to the private network or to add a public interface to connect to the Internet.<br><br>**Note:** You do not need to manually configure public or private interfaces for the NAT component if you used the Routing and Remote Access Setup wizard to configure NAT.<br><br>Viewing the NAT mapping table:<br>• Click **NAT/Basic Firewall**, right-click the public interface in the details pane, and then click **Show Mappings**.<br><br>Viewing DHCP allocator information:<br>• Right-click **NAT/Basic Firewall**, and then select **Show DHCP Allocator Information** to display the number of instances for each of the following:<br>  • Messages ignored<br>  • DECLINE messages received<br>  • DISCOVER messages received<br>  • INFORM messages received<br>  • RELEASE messages received<br>  • REQUEST messages received<br>  • ACK messages sent<br>  • BOOTP replies sent<br>  • NAK messages sent<br>  • OFFER messages sent<br><br>Viewing DNS proxy information:<br>• Right-click **NAT/Basic Firewall**, and then select **Show DNS Proxy Information** to display the number of instances for each of the following:<br>  • Messages ignored<br>  • Queries received<br>  • Responses received<br>  • Queries sent<br>  • Responses sent |

The following table describes how each tab on the **NAT/Basic Firewall Properties** page in the Routing and Remote Access snap-in is used for NAT-related tasks.

**NAT-related Options on the NAT/Basic Firewall Properties Page**

| Tab | Task |
| --- | --- |
| General | Specifying the level of errors and warnings to be logged in the System Log in Event Viewer:<br>• **Log errors only.** Specifies that only errors are logged in the System Log in Event Viewer.<br>• **Log errors and messages.** Specifies that both errors and warnings are logged in the System Log in Event Viewer.<br>• **Log the maximum amount of information.** Specifies that the maximum amount of information is logged in the System Log in Event Viewer.<br>• **Disable event logging.** Specifies that no events are logged in the System Log in Event Viewer. |
| Translation | Specifying the number of minutes that a dynamic mapping for a TCP session or for a UDP message remains in the NAT Mapping Table. |
| Address Assignment | Configuring the DHCP allocator feature:<br>1. Specify whether the NAT-enabled router will provide DHCP-based address assignment to DHCP clients on the private network.<br>2. Specify both the private address range and any exclusions; that is, specify any addresses within the specified range of addresses that should not be assigned to DHCP clients on the private network because they are already in use.<br>If multiple routed subnets are configured, you must use a DHCP server rather than the DHCP allocator. |
| Name Resolution | Configure the DNS proxy feature:<br>1. Specify whether the NAT-enabled router relays DNS name resolution requests from hosts on the private network to the configured DNS server for the NAT-enabled router.<br>2. Specify whether a connection is attempted by using the selected demand-dial interface when a DNS name resolution request is received by a host on the private network. |

The following table describes how each tab on the properties page of the public (Internet-connected) interface in the details pane of the Routing and Remote Access snap-in is used for NAT-related tasks.

**NAT-related Options on the Public Interface Properties Page**

| Tab | Task |
| --- | --- |

| | |
|---|---|
| NAT/Basic Firewall | Configuring NAT:<br>● Select **Enable NAT on this interface** to enable the router to send data to and receive data from the Internet over this interface.<br><br>Configuring Basic Firewall:<br>● Select **Enable a basic firewall on this interface** to protect computers on the private network from unsolicited Internet traffic.<br><br>Configuring static packet filters:<br>● Under **Static packet filters**, select **Inbound Filters** or **Outbound** Filters to establish inbound and outbound packet filters on the public interface to restrict traffic based on packet attributes such as IP address or protocol. For example, you can use this option to configure filters for PPTP or L2TP/IPSec VPN connections, as described in the Windows Server 2003 Deployment Guide section about "Configuring Packet Filters for a VPN Server" in Deploying a VPN Remote Access Server Solution. |
| Address Pool | Configuring one or more IP address ranges:<br>● Select **Add** to configure an IP address pool on the public interface. You cannot configure an IP reservation until you configure at least one IP address pool.<br><br>Configuring an IP reservation to allow incoming traffic to a computer (such as a Web server) on the private network:<br>● Under **Reserve public addresses**, click **Reservations** to configure an IP reservation by specifying a public IP address from one of the configured ranges. This ensures that the reserved address cannot be used for address translation. |
| Services and Ports | Configuring a static mapping for the services on your network to which you want to provide access for Internet users:<br>● Select one of more from the list of services provided (such as **FTP Server**, **Web Server (HTTP)**, or **Remote Desktop**).<br><br>-or-<br><br>● Click **Add** to configure a service that is not on the list to which you want to provide access for Internet users.<br><br>**Note:** Configuring a service on this tab creates a static entry in the NAT Mapping Table and creates exceptions in the Basic Firewall that allow the specified incoming traffic. |
| ICMP | Configuring Internet Control Message Protocol (ICMP) options:<br>● Select the requests for error and status information (ICMP messages) from the Internet to which this computer will respond. |

The following table describes how the tab on the properties page of the **Internal** interface in the details pane of the Routing and Remote Access snap-in is used for NAT-related tasks.

**NAT-related Options on the Internal Properties Page**

| Tab | Task |
|---|---|

| | |
|---|---|
| NAT/Basic Firewall | Configuring static packet filters:<br>• Under **Static packet filters**, select **Inbound Filters** or **OutboundFilters** to establish inbound and outbound packet filters on the private interface to restrict traffic based on packet attributes such as IP address or protocol<br><br>**Note:** If the NAT/Basic Firewall component is enabled, the option **Private interface connected to private network** is selected on this tab by default and cannot be unselected. |

**Netsh Command-Line Tools for Routing and Remote Access NAT**

Netsh provides several sets of commands (also known as *contexts*) for performing a wide range of network configuration tasks. The Netsh Routing IP NAT commandsprovide the Netsh context for Routing and Remote Access NAT.

**Netsh.exe: Netsh Routing IP NAT Commands**

**Category**

The Netsh Routing IP NAT commands, a subset of the Netsh command-line toolset, are included with the Windows Server 2003 operating system.

**Version compatibility**

The Netsh Routing IP NAT commands are compatible with Windows Server 2003. Netsh commands were first introduced in Windows 2000 Server and were expanded to include additional commands, including commands to manage NAT, in Windows Server 2003.

The Netsh commands are designed to help network administrators manage a TCP/IP network. You can use the Netsh command-line set of tools to locally or remotely display or modify the configuration of services or protocols on Windows–based computers. The Netsh command-line interface is scriptable, which lets you perform batch configurations or network administration from a centralized location. In addition to the Netsh Routing IP NAT commands that are designed specifically for Routing and Remote Access NAT, NAT also inherits commands from the Netsh Routing context and the Netsh Routing IP context.

The following table contains a brief description of the commands available in the Netsh Routing IP NAT context.

**Commands Available in the Netsh Routing IP NAT Context**

| NAT Context Command | Description |
|---|---|
| ? or help | When typed at a **netsh routing ip nat>** prompt, either **?** or **help** displays a complete list of all commands in the Netsh Routing IP NAT context, including all commands inherited from the global Netsh context as well as commands inherited from the Netsh Routing and Netsh Routing IP subcontexts.<br><br>When typed at a **netsh routing ip nat>** prompt, a command name followed by **?** (such as **show ?**) displays information about that command. |
| add addressmapping | Adds an IP address mapping to the NAT address pool for the specified interface. |
| add addressrange | Adds an address range to the NAT address pool for the specified interface. |
| add ftp | Enables the NAT proxy for FTP (supports FTP traffic across a NAT). |
| add h323 | Enables the NAT proxy for H.323 (supports NetMeeting calls across a NAT). |
| add interface | Configures NAT on the specified interface. |
| add portmapping | Adds a protocol port mapping for either the TCP or the UDP protocol type on the NAT interface. |
| delete addressmapping | Deletes an address mapping from the NAT address pool for the specified interface. |

| | |
|---|---|
| delete addressrange | Deletes an address range from the NAT address pool for the specified interface. |
| delete ftp | Disables the NAT proxy for FTP. |
| delete h323 | Disables the NAT proxy for H323. |
| delete interface | Removes NAT from the specified interface. |
| delete portmapping | Deletes a protocol port mapping for either the TCP or the UDP protocol type from the specified NAT-enabled interface. |
| set global | Sets the following global parameters for NAT:<br>• The timeout value, in minutes, for TCP mappings.<br>• The timeout value, in minutes, for UDP mappings.<br>• Which events should be logged. The **none** parameter specifies that no events related to NAT should be logged. The **error** parameter specifies that only errors related to NAT should be logged. The **warn** parameter specifies that only warnings related to NAT should be logged. The **info** parameter specifies that all events related to NAT should be logged. |
| set interface | Configures NAT parameters for the specified interface. |
| show global | Displays NAT global configuration. That is, it displays the current defaults for the following:<br>• TCP timeout (in minutes)<br>• UDP timeout (in minutes)<br>• Logging level (such as errors only) |
| show interface | Displays NAT configuration for the specified interface. |

For more information about Netsh, see "Command-Line Reference for Windows Server 2003, Standard Edition" in the Tools and Settings Collection.

**NAT Registry Entries**

The following registry entry, associated with Routing and Remote Access NAT, is the only registry entry that an administrator might want to modify by using the registry editor.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**AllowInboundNonUnicastTraffic**
**Registry path**
**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IpNat\Parameters**

**Version**

Windows Server 2003, Windows XP SP1, or later.

If Routing and Remote Access NAT has Basic Firewall configured, the firewall always accepts broadcast and multicast packets and passes them to the NAT component. However, on a computer running the Windows Server 2003, Windows XP SP1, or later operating system, the following registry key for NAT is set by default to drop all inbound broadcast and multicast packets. If you need to change this default behavior, add the following registry key and set it to **1**. Setting the key to **1** allows broadcast and multicast packets to cross Basic Firewall:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IpNat\Parameters\AllowInboundNonUnicastTraffic**

By default, **AllowInboundNonUnicastTraffic** is set to **0**, which blocks inbound unicast traffic.

For more information about this registry entry, see the Registry Reference for Windows Server 2003.

**Related Information**
The following resources contain additional information that is relevant to this section:

- "Command-Line Reference for Windows Server 2003, Standard Edition" in the Tools and Settings Collection
- "Registry Reference for Windows Server 2003" in the Tools and Settings Collection

# Unicast IPv4 Routing Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Unicast IPv4 Routing Technical Reference**
This technical reference describes unicast Internet Protocol version 4 (IPv4) routing and explains how it works on an internetwork using servers, including routers, that run the Microsoft Windows Server 2003 operating system. IPv4 (often referred to as IP) is a major component of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite, which is the standard networking protocol used in most networking environments today as well as on the Internet. A TCP/IP-based internetwork, typically, is called an IP internetwork, although the IP network protocol is only a part of TCP/IP.

IP is the routable network protocol that makes communication possible across IP internetworks. Unicast IP routing refers to the process of transmitting IP packets from a single node (a host or a router) to a single node on an IP internetwork. Although, as the name Internet Protocol indicates, IP packet transmission was originally intended for traffic over the Internet, the transmission of IP packets is now also the primary type of traffic used on most intranets.

The information provided here is relevant to the IP unicast routing that occurs on any Windows-based network, including networks that use hardware routers, software-based routers such as those that use the Windows Server 2003 Routing and Remote Access service, or a combination of hardware and software routers. Included is a description of the multiple features of the Routing and Remote Access service that support unicast IP routing.

In addition to supporting IP routing, Windows Server 2003 also supports AppleTalk routing. The Windows Server 2003 family no longer supports Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) routing. Microsoft Windows 2000 Server supports unicast routing for IP, AppleTalk, and IPX networks.

Although IPv4 will likely be replaced in the future by the next generation of IP, IP version 6 (IPv6), at present most unicast IP routing takes place over IPv4 networks. Windows Server 2003 supports both IPv4 and IPv6. For more information about Windows Server 2003 support for IPv6, including support for static routing in an IPv6 environment, see "IPv6 Technical Reference."

**In this subject**


- What Is Unicast IPv4 Routing?
- How Unicast IPv4 Routing Protocols and Services Work
- How Unicast IPv4 Routing Works
- Unicast IPv4 Routing Tools and Settings

# What Is Unicast IPv4 Routing?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Unicast IPv4 Routing?**

**In this section**


- IPv4 Routing Enables Most Communications Worldwide
- Windows Server 2003 Supports Multiple Unicast IPv4 Routing Scenarios
- How Unicast IPv4 Routing Differs from Related or Similar Technologies
- Related Information


Internet Protocol (IP), a part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite, is the routable network protocol that enables the routing of network traffic across any type of IP internetwork, including Windows internetworks, UNIX internetworks, and mixed network environments. The Windows Server 2003 family of operating systems uses IP to route communications, including unicast traffic, between computers or other network devices on private intranets, ranging in size from very small to intercontinental enterprise internetworks. IP also enables communication across the public Internet, which is itself an IP-based internetwork.

Currently, most unicast IP routing takes place over IP version 4 (IPv4) internetworks. Windows Server 2003 also supports the next generation of IP, IP version 6 (IPv6). This technical reference describes unicast IPv4 routing on IPv4 internetworks. Frequently, IPv4 internetworks are referred to simply as IP internetworks, and this technical reference follows that convention. For more information about Windows Server 2003 support for IPv6, including support for static routing in an IPv6 environment, see "IPv6 Technical Reference."

An IP internetwork, like any computer internetwork, consists of smaller networks joined together by the use of interconnecting devices known as routers. IP routing is the process of forwarding IP packets from a network device on one part of an IP internetwork to a network device on another network segment (subnet). An IP packet or datagram is a unit of information sent over an IP network that includes data intended for the recipient as well as a header containing routing information (the source and destination addresses and error-control data). IP routers forward packets between network segments.

Unicast IP routing is the process that enables unicast IP packets to be forwarded across an IP internetwork from a sending node to a destination node through one or more intermediate routers. A node is any network device that is running the TCP/IP protocol. A host is a node that does not perform routing, such as a user workstation or a non-router server. A router is a node that performs routing. That is, a router forwards packets that are not destined for the router itself either directly to the destination or to another router on the route to the destination.

**IPv4 Routing Enables Most Communications Worldwide**

Over the past few decades, the ability for people to communicate by sending messages from one computer to another, whether the computers are located on the same network in an office building or on networks at opposite sides of the globe, has become so commonplace that it is difficult to imagine a world where such communication is not possible. IP routing provides the infrastructure that enables all other IP-based technologies, such as Domain Name System (DNS), to function. The routing of unicast IP packets over IP internetworks is a major part of the technology that makes such communication possible.

Today, the majority of internetwork traffic worldwide is over IPv4 networks, and most user-initiated traffic across IPv4 internetworks is unicast traffic. Unicast IP routing occurs on every IP internetwork, including:

- An IP intranet not connected to the Internet
- The global Internet
- Intranets that connect to the Internet or to each other through the Internet

The major operating systems for which TCP/IP is the primary network protocol are Windows and UNIX.

Any Windows network supports unicast IP routing. These include networks that use only hardware routers, networks that use software-based routers such as the Routing and Remote Access service included with Windows Server 2003, or networks that use a combination of hardware and software routers.

Although all modern networking operating systems support the TCP/IP protocol suite, Windows Server 2003 TCP/IP provides the best platform for connecting Windows-based systems to earlier Windows-based systems and to non-Windows-based systems. In Windows Server 2003, TCP/IP supports enterprise networking and connectivity on computers running the Windows Server 2003, Microsoft Windows XP, Windows 2000, Windows NT, Windows Millennium Edition, and Windows 98 operating systems and on computers running the UNIX operating system.

For more information about Windows Server 2003 TCP/IP, see "TCP/IP Technical Reference."

**Windows Server 2003 Supports Multiple Unicast IPv4 Routing Scenarios**
A computer running the Windows Server 2003 Routing and Remote Access service that has two or more network adapters (each configured with an appropriate IP address and subnet mask) is a software router that can provide a wide range of services that support IP routing. A major part of this support is support for static and dynamic IP routing:

- **Static IP routing.** An administrator manually configures routing information, and routing information does not change unless the administrator manually updates or deletes it. Using static routing exclusively is appropriate only for a small internetwork.

- **Dynamic IP routing.** An administrator configures a router to use a dynamic routing protocol. A router that uses a dynamic routing protocol automatically generates routing information, shares that information with other routers, and updates the information when routing changes occur. A large internetwork that primarily uses dynamic routing typically also uses some manually configured static routes.

A computer running Routing and Remote Access that has only one interface is, by definition, a non-router computer and can provide some routing-related networking services, such as IP packet filtering or acting as a DHCP relay agent.

For a Windows Server 2003–based IP internetwork that uses hardware routers exclusively and does not install any Routing and Remote Access routers, IP routing is still the mechanism used for routing packets on the internetwork. This is because the primary networking protocol used by Windows Server 2003 is TCP/IP, and the Windows Server 2003 TCP/IP protocol is installed by default on any computer running Windows Server 2003. TCP/IP can be uninstalled on earlier versions of the Windows server operating systems, but it cannot be uninstalled on Windows Server 2003.

The Windows Server 2003 Routing and Remote Access service supports the following unicast IP routing scenarios:

**Static route management**
Support for static route management. A static IP route is an IP route that a network administrator adds to the IP routing table manually. Static IP routing is an inherent function of the TCP/IP protocol for Windows Server 2003. You can manage static routes by using the Routing and Remote Access snap-in, by using the **route** command-line tool, or by using the set of Netsh routing IP commands. A server running the Routing and Remote Access service can act exclusively as a static IP router, or it can act primarily as a dynamic IP router that also includes manually configured static routes.

**Routing Information Protocol (RIP) for IP**
Support for versions 1 and 2 of Routing Information Protocol (RIP) for IP (RIP v1 and RIP v2). A distance-vector routing protocol, RIP for IP is the primary dynamic routing protocol used in small or medium-sized IP internetworks. A server running the Routing and Remote Access service configured with RIP v2 can function as a dynamic IP router. RIP v1 is obsolete but is included to support a mixed environment that contains both RIP v1 and RIP v2 routers.

**Open Shortest Path First (OSPF)**
Support for the industry standard OSPF dynamic routing protocol. A link state routing protocol commonly used in large or very large IP internetworks, the OSPF routing protocol is available only on 32-bit versions (not 64-bit versions) of the Windows Server 2003 family. A server running the Routing and Remote Access service configured with OSPF can function as a dynamic IP router.

**DHCP relay agent**
Support for the RFC 1542–compliant Dynamic Host Configuration Protocol (DHCP) Relay Agent service, also known as Bootstrap Protocol (BOOTP) Relay Agent. DHCP Relay Agent transfers messages between DHCP clients and DHCP servers located on separate network segments. Any IP subnet that contains DHCP clients requires either a DHCP server or a DHCP relay agent to provide address leases to those DHCP clients. Using one or more DHCP relay agents makes it unnecessary to install a separate DHCP server on each subnet in an internetwork.

**IPv4 packet filtering**

Support for separately configured input and output filters for each IPv4 interface. IPv4 packet filtering specifies which traffic is allowed into and out of each interface, based on filters defined by source and destination IP addresses, Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) port numbers, Internet Control Message Protocol (ICMP) types and codes, and IP protocol numbers. In addition to IP packet filtering, the Windows Server 2003 Routing and Remote Access service also includes support for Basic Firewall, which uses stateful firewall technology similar to that used by the Internet Connection Firewall (ICF) feature available in Windows XP, Windows XP SP1, and Windows Server 2003. Basic Firewall can be used in conjunction with Routing and Remote Access NAT or on any other Internet-connected interface.

**ICMP router discovery**
Support for ICMP Router Solicitation and Router Advertisement messages to allow the automated discovery of default routers by hosts. If both routers and hosts are configured to use ICMP router discovery, this feature simplifies how IP hosts are configured with the IP addresses of local routers and it provides a way for hosts to discover routers that are down.

**Network address translation (NAT)**
Support for network address translation, which translates private and public addresses to allow small office or home office (SOHO) networks or medium-sized networks that use private addresses to connect to the Internet. The Routing and Remote Access NAT routing protocol component also includes a DHCP allocator and a Domain Name System (DNS) proxy to simplify the configuration of SOHO hosts.

Each of these unicast IP routing scenarios, except NAT, is described "How Unicast IPv4 Routing Protocols and Services Work" in this technical reference. For more information about Routing and Remote Access NAT, see "NAT Technical Reference."

For an introductory overview of protocol-independent unicast routing principles, see "Unicast Routing Overview" in the "*Microsoft Windows 2000 Resource Kit*." Windows 2000 Server supports both IP and IPX, and "Unicast Routing Overview" includes a discussion of both.

**How Unicast IPv4 Routing Differs from Related or Similar Technologies**
The following sections briefly summarize how routable and routing protocols differ from each other, how the IP routable protocol differs from the AppleTalk and IPX routable protocols, how IPv4 differs from IPv6, and how unicasting differs from broadcasting and multicasting.

**Routable Protocols and Routing Protocols**
To understand the types of protocols associated with routing, it is helpful to distinguish between routable protocols and routing protocols:

- **Routable protocol.** A network protocol that uses layer 3 (Network layer) addresses to forward packets to their destination across an internetwork. For example, IP is the routable protocol in the TCP/IP protocol suite that is responsible for IP addressing, routing, and the fragmentation and reassembly of IP packets. For information about routable protocols other than IP that provide similar services for other types of internetworks, see "IP, AppleTalk, and IPX/SPX" described later in this section.

- **Routing protocol.** A protocol that enables routers to exchange routing information with each other automatically across an internetwork by exchanging a series of periodic or on-demand messages. The major unicast IP routing protocols that route traffic over private IP internetworks are RIP v2 and OSPF. RIP v1 is now obsolete for enterprise-size internetworks.

For more information about RIP and OSPF, see "How Unicast IPv4 Routing Protocols and Services Work."

**IP, AppleTalk, and IPX/SPX Routable Protocols**
As just described, a routable protocol is a network protocol that enables the routing of packets across an internetwork. Three major routable protocols include:

- **IP internetworks use IP**. The IP portion of the TCP/IP protocol suite constitutes the industry standard for routing IP packets between nodes on an IP internetwork and enables communication across any set of interconnected IP networks. IP, the routable protocol described in this document, is the primary routable protocol for the Windows and UNIX operating systems.

- **AppleTalk internetworks use AppleTalk.** AppleTalk is the protocol suite developed by Apple Computer for AppleTalk Macintosh networks. Routing Table Maintenance Protocol (RTMP) is the routing protocol used in the AppleTalk environment. AppleTalk internetworks are also known as AppleTalk internets, which are small, physical networks connected by routers.

- **NetWare internetworks use IPX/SPX.** Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) is the

routable protocol developed by Novell for NetWare networks. A typical IPX internetwork might contain Novell NetWare clients and servers running IPX located in multiple networks connected by IPX-enabled routers.

Windows Server 2003 and Windows 2000 Server support the AppleTalk protocol stack and AppleTalk routing software. Windows 2000 Server supports IPX routing, but Windows Server 2003 does not.

**IPv4 and IPv6 Routing**

In addition to IPv4, Windows Server 2003 TCP/IP supports the latest version of IP, IPv6. IPv6 includes unicast and multicast address types that are more complex than IPv4 unicast and multicast address types. IPv6 also includes a new type of address called anycast addresses. IPv6 addressing is a major departure from IPv4 addressing. The most obvious difference is that IPv4 uses 4-byte source and destination IP addresses, typically expressed in the familiar dotted-decimal notation, whereas IPv6 uses 16-byte addresses, expressed in colon-hexadecimal notation. Colon-hexadecimal notation uses eight 4-digit hexadecimal numbers, with colons separating the 16-bit blocks (the 4-digit numbers).

In Windows Server 2003, the Routing and Remote Access service does not support the dynamic routing protocols RIPng (the version of RIP for IPv6) or OSPF for IPv6. However, a computer running Windows Server 2003 can act as a static IPv6 router, and it is possible to route IPv6 traffic over an IPv4 internetwork using an IPv6 transition technology. For more information about routing in an IPv6 environment, see "How IPv6 Works" in "IPv6 Technical Reference."

**Unicasting, Broadcasting, and Multicasting**

In addition to forwarding unicast IP packets, IP routers can also forward broadcast and multicast IP packets from a device on one network to multiple devices on another network. Although most user-initiated traffic on an IP internetwork is unicast traffic, to understand the role of unicasting on IP networks, you must also understand how its function differs from broadcasting and multicasting:

- **Unicasting** is one-to-one communication of IP packets from a sending node to a receiving node on an IP internetwork. Unicast packets might pass through other devices as they are transmitted across the internetwork; however, unicast packets, by definition, are always sent from a single sending node to a single destination node. Any technology that involves sending packets from one node to another over an IP-based internetwork uses unicast routing.

- **Broadcasting** is one-to-everyone communication of IP packets between one node and all other accessible nodes on the same subnet. If routers are configured to forward internetwork-level broadcast traffic (highly discouraged for all IP routers), internetwork-level broadcasting forwards packets to all other network segments except the network segment on which the packet was received. The uses of broadcast communication on a subnet include announcing the availability of network services, resolving names into addresses, and resolving IP addresses into Media Access Control (MAC) addresses. DHCP and the NetBIOS name service are IP services that use broadcasting.

- **Multicasting** is one-to-many communication of IP packets between one node and multiple nodes that elect to participate in a specific multicast group. Multicast communication is used primarily for multiple-user multimedia applications such as video conferencing, distance learning, and collaborative computing. For more information about multicast routing on an IPv4 internetwork, see "IPv4 Multicasting Technical Reference."

Unicast routing performs broadcasting and multicasting for specific purposes. For example, a RIP v1 announcement (in which the router advertises the contents of its routing table to other routers) is always broadcasted. A RIP v2 router can be configured either to broadcast or to multicast its announcements.

**Related Information**

The following resources contain additional information that is relevant to this section.

- How Unicast IPv4 Routing Protocols and Services Work

- TCP/IP Technical Reference

- IPv6 Technical Reference

- IPv4 Multicasting Technical Reference

- "Unicast Routing Overview" in the "*Microsoft Windows 2000 Resource Kit*."

- "*Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference*" by Joseph Davies and Thomas Lee,

2003, Microsoft Press.

# How Unicast IPv4 Routing Protocols and Services Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Unicast IPv4 Routing Protocols and Services Work**

**In this section**


- Unicast IP Routing Architecture in Windows Server 2003
- How RIP Works
- How OSPF Works
- How the DHCP Relay Agent Works
- How IP Packet Filtering Works
- How ICMP Router Discovery Works
- IPv4 and Routing and Remote Access APIs
- Related Information


The Microsoft Windows Server 2003 family of operating systems provides extensive support for unicast Internet Protocol version 4 (IPv4) routing — that is, routing IPv4 packets from a source node to a destination node identified by its unicast destination IPv4 address. IP, a part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite, is the routable network protocol that enables the routing of network traffic across any type of IP internetwork, including Windows internetworks, UNIX internetworks, and mixed network environments. Windows Server 2003 includes support for unicast IPv4 dynamic routing protocols as well as for additional unicast routing features provided by the Windows Server 2003 Routing and Remote Access service. IPv4 is frequently referred to as IP, and that convention is used in this document.

**Note**

- Windows Server 2003 also supports Internet Protocol version 6 (IPv6). For information about Windows Server 2003 support for IPv6, including support for static routing in an IPv6 environment, see "How IPv6 Works" in IPv6 Technical Reference.

An organization's implementation of unicast IP routing can be simple or complex, depending on factors such as the size of its internetwork, the use of Dynamic Host Configuration Protocol (DHCP) to allocate IP addresses, connectivity to the Internet, and the presence of non-Windows hosts on the internetwork. To determine how best to implement or modify the unicast IP routing infrastructure for your internetwork, the first step is to understand how unicast IP routing works.

The Windows Server 2003 Routing and Remote Access service does not include support for the Internetwork Packet Exchange (IPX) protocol, such as Routing Information Protocol (RIP) for IPX, Service Advertising Protocol (SAP) for IPX, Network basic input/output system (NetBIOS) over IPX, or IPX packet filtering.

**Unicast IP Routing Architecture in Windows Server 2003**

Unicast IP routing occurs on any TCP/IP-based internetwork by using IP routers to forward IP packets from one network segment to another. Unicast IP routing occurs on a Windows Server 2003 IP internetwork whether or not any computers run the Routing and Remote Access service. You can use hardware routers on your internetwork, software routers, or a combination of hardware and software routers. The Windows Server 2003 Routing and Remote Access service can play a role in implementing and supporting unicast IP routing features on a Windows network even when most, or all, routers are hardware routers.

To see an architectural diagram showing where IP fits within the TCP/IP protocol suite, see "TCP/IP Protocol Architecture" in "How TCP/IP Works" in TCP/IP Technical Reference. The following figure shows the architecture of the Routing and Remote Access service in Windows Server 2003. The components are differentiated in the figure as follows:

- Clear boxes represent components that are used for unicast IP routing and multicast IP routing, as well as for other Routing and Remote Access–supported features.
- Dark blue boxes represent components that are unique to unicast routing.
- Light blue boxes represent components that are unique to multicast routing.
- The orange box represents the only Routing and Remote Access component — Simple Network Management Protocol (SNMP) — that is used neither for unicast nor for multicast routing.

**Unicast IP Components of the Windows Server 2003 Routing and Remote Access Service**




The following table describes the unicast IP routing components depicted in the figure, including components that are shared by unicast routing and multicast routing as well as components that are unique to unicast routing. Not included in the table are the non-unicast components: multicast-only components (light blue boxes) or SNMP (orange box). For information about SNMP, see "SNMP Technical Reference." For information about multicast IP routing, see "IPv4 Multicasting Technical Reference."

**Unicast Routing Components in the Routing and Remote Access Service**

| Component | Description |
|---|---|
| Routing and Remote Access snap-in | Management application in the Administrative Tools folder for administering unicast IP routing and other routing and remote access features. |
| Netsh | Command-line management application that provides sets of commands (known as *contexts*) used to perform a wide range of network configuration tasks, including unicast routing tasks. The Netsh contexts specific to unicast routing include **netsh routing** and **netsh routing ip**. |
| AAA | A set of components that provide authentication, authorization, and accounting (AAA) for the Routing and Remote Access service when it is configured to act as the Windows authentication provider or the Windows accounting provider. Alternatively, the computer running the Routing and Remote Access service can be configured with the Remote Authentication Dial-In User Service (RADIUS) authentication or accounting provider in the Internet Authentication Service (IAS), in which case, the local AAA components are not used. |

| | |
|---|---|
| Dynamic Interface Manager<br><br>(Mprdim.dll) | Component that provides the following services for unicast IP routing and other Routing and Remote Access–supported features:<br><br>• Communicates configuration information to the IP Router Manager.<br><br>• Manages all routing interfaces, including local area network (LAN), dial-up interfaces, non-persistent demand-dial interfaces, and persistent demand-dial interfaces. Unlike Windows 2000 Server, Windows Server 2003 Routing and Remote Access does not support IP-in-IP tunnels (which encapsulate an IP packet within another IP packet and forward the packet across the tunnel) and thus does not support IP-in-IP interfaces (logical interfaces used to send IP packets across an IP-in-IP tunnel).<br><br>• Loads configuration information from the Windows Server 2003 registry.<br><br>• Supports a remote procedure call (RPC) interface for SNMP-based management functions used by management tools such as the Routing and Remote Access snap-in.<br><br>• Communicates with the Connection Manager for demand-dial connections. |
| Connection Manager | A set of components that provides the following services:<br><br>• Manages wide area network (WAN) devices.<br><br>• Establishes connections by using the Telephony application programming interface (TAPI).<br><br>• Negotiates Point-to-Point Protocol (PPP) control protocols, including Extensible Authentication Protocol (EAP). PPP is an industry standard suite of protocols for the use of point-to-point links to transport multiprotocol datagrams. EAP is an extension to PPP that allows for arbitrary authentication mechanisms to be employed for the validation of a PPP connection.<br><br>• Implements multilink protocol (MP) and Bandwidth Allocation Protocol (BAP). MP, also known as MLP or Multilink Point-to-Point Protocol (MPPP), is an extension to PPP that is used to aggregate multiple physical links into a single logical link. BAP is a PPP control protocol that is used on a multilink connection to add and remove links dynamically.<br><br>**Note:** Routing and Remote Access Connection Manager, described here, is not the same as the Connection Manager client dialer. |
| TAPI | An API used by communications programs to work with telephony and network services. Telephony API, or TAPI, provides services to create, monitor, and terminate connections in a hardware-independent manner. Connection Manager uses TAPI to create or receive demand-dial connections. |

| | |
|---|---|
| IP Router Manager<br><br>(Iprtrmgr.dll) | Component that starts and manages other routing components, including routing protocols, the routing table manager, and the multicast group manager. The IP Router Manager provides the following services:<br><br>• Obtains configuration information from the dynamic interface manager.<br><br>• Communicates IP packet filtering configuration to the IP filtering driver.<br><br>• Communicates IP routing configuration information to the IP forwarder in the TCP/IP protocol.<br><br>• Maintains an interface database of all IP routing interfaces.<br><br>• Loads and communicates configuration information to IP routing protocols that are supplied with Windows Server 2003, such as Routing Information Protocol (RIP) and Open Shortest Path First (OSPF).<br><br>• Initiates demand-dial connections for routing protocols by communicating with Dynamic Interface Manager. |
| IP RIP<br><br>(Iprip2.dll) | A unicast distance vector dynamic routing protocol. The RIP for IP routing protocol (also known simply as RIP) provides the following services:<br><br>• Communicates RIP learned routes by using the Route Table Manager.<br><br>• Uses Windows Sockets (Winsock) to send and receive RIP traffic. Winsock is an implementation of the industry-standard Sockets API for the Windows operating system.<br><br>• Exports management APIs to support SNMP management information bases (MIBs) and other management applications by using the IP Router Manager. A MIB is a set of objects, which represent various types of information about a device that are used by SNMP to manage the device.<br><br>The Windows Server 2003 Routing and Remote Access service supports (for IPv4 only) both RIP version 1 and version 2 (RIP v1 and RIP v2). However, RIP v1 is considered outdated. For more information about RIP v1 and RIP v2, see “How RIP” later in this document. |
| OSPF<br><br>(Ospf.dll) | A unicast link state dynamic routing protocol. The OSPF routing protocol performs the following services:<br><br>• Communicates OSPF learned routes by using the Route Table Manager.<br><br>• Uses Winsock to send and receive OSPF traffic.<br><br>• Exports management APIs to support SNMP MIBs and other management applications by using the IP Router Manager.<br><br>The OSPF routing protocol is available only on 32-bit versions of Windows Server 2003. For more information about OSPF, see “How OSPF Works” later in this document. |

| | |
|---|---|
| Route Table Manager<br><br>(Rtm.dll)<br><br>(Also called Routing Table Manager) | The central repository for routing information for all routing protocols that operate under the Routing and Remote Access service and for other components such as the IP Router Manager. The Route Table Manager provides the following services:<br>• Maintains a user-mode routing table for all routes from all possible route sources.<br>• Exposes APIs for adding, deleting, and enumerating routes that are used by the routing protocols.<br>• Timestamps learned routes.<br>• Provides routing information to routing protocols, management programs, and monitoring programs.<br>• Communicates only the best routes to the appropriate IP forwarder driver. For IP routes, the best route is the route with the lowest cost based on specified criteria, such as the number of networks crossed, the type of network crossed (for example, private or public), or a monetary or bandwidth limit. The best routes become the routes in the IP forwarding table.<br>• Supports multiple instances of the same protocol, such as the Microsoft implementation of OSPF and other implementations of OSPF. |
| IP Filtering<br><br>(Ipfltdrv.sys) | Component that performs the following services:<br>• Obtains configuration information from the IP Router Manager.<br>• Applies IP filters after the IP forwarder has found a route. |
| IP Unicast Forwarder | A component of the TCP/IP protocol (Tcpip.sys). The IP Unicast Forwarder performs the following services:<br>• Obtains configuration information from the IP Router Manager.<br>• Stores the IP forwarding table, a table of the best routes obtained from the Route Table Manager.<br>• Determines the next-hop address and interface for IP packets being forwarded.<br>• Forwards unicast IP traffic from one router interface to another router interface.<br>• Initiates a demand-dial connection. |
| NDIS Wrapper<br><br>(Ndis.sys) | Code that surrounds all Network Driver Interface Specification (NDIS) device drivers. Windows Server 2003 NDIS (and Windows 2000 Server NDIS) is implemented by a file called Ndis.sys, referred to as the NDIS Wrapper. NDIS enables transport protocols such as TCP/IP to communicate with an underlying network adapter or other hardware device so that the network adapter can send or receive data over the network. The NDIS Wrapper provides a uniform interface between protocol drivers and NDIS device drivers. |

**Note**

- Network address translation (NAT) is a feature of the Routing and Remote Access service and is part of unicast routing. However, NAT is not shown in the preceding figure or table because NAT is not a routing protocol. For information about the NAT routing protocol component and Routing and Remote Access NAT architecture, see "How NAT Works" in NAT Technical Reference.

**Routing and Remote Access Unicast IP Routing Processes**

Typical unicast IP routing processes for the components of the Routing and Remote Access service described in the preceding section include the following types of communications.

**Incoming and outgoing packet (transit traffic)**

An incoming packet is passed by the TCP/IP driver to the IP forwarder, which finds a route and then passes the packet to the IP

filtering driver to check for input and output filters. If the packet is approved for acceptance by the input filters and for forwarding by the output filters, it is passed back to the IP forwarder driver, which uses NDIS to forward the packet over the appropriate interface. If the input or output filters do not permit the packet to be received or forwarded, the packet is silently discarded. If a route is not found, an Internet Control Message Protocol (ICMP) Destination Unreachable-Host Unreachable message is sent back to the source of the packet.

**Incoming packet (local host traffic)**
An incoming packet is passed by the TCP/IP driver first to the IP forwarder, which notes that the packet is not to be routed (because the destination IP address is an address assigned to the router or it is a broadcast address). The IP forwarder then hands the packet to the IP filtering driver to check for input filters. If the packet is accepted by the input filters, it is passed up to the TCP/IP driver, which processes the packet. If the packet is not accepted by the input filters, the packet is silently discarded.

**Outgoing packet (local host traffic)**
An outgoing packet is passed by the TCP/IP driver to the IP filtering driver, which checks for output filters. If the packet is approved for sending by the output filters, the packet is passed to the IP forwarder, which sends the packet by using the best route over an appropriate interface that uses NDIS. If the packet is not approved by the output filters, the packet is silently discarded. If a route is not found, an IP routing error is passed to the source application of the packet.

**Routing protocol network communication**
The RIP for IP and OSPF routing protocols operate like any other Winsock application that sends or receives IP packets. For a brief description of the Windows Sockets 2 (Winsock2) API, see "IPv4 and Routing and Remote Access APIsW2K3TR" later in this document. For more information about the Windows Sockets API, see "Microsoft Platform SDK Windows Sockets 2."

**Routing table updates**
The RIP and OSPF routing protocols update routes in the Route Table Manager. Based on the best route and route source ranking, the table of best routes (the IP forwarding table) is updated in the IP forwarder.

**Routing and Remote Access Features by Windows Server Operating System**
In addition to support for unicast IP routing features, the Routing and Remote Access service for Windows Server 2003 also includes a variety of other routing and remote access features, most of which interact with or support unicast IP routing — for example, LAN and WAN support, both of which are integral to unicast IP routing, and the Routing and Remote Access snap-in, which you can use to configure unicast IP routing features. The following table lists the routing and remote access features available in Windows operating systems from Windows NT 3.51 to Windows Server 2003.

**Comparison of Routing and Remote Access Features in Windows Operating Systems**

| Feature | Windows NT Server 3.51 SP2 | Windows NT Server 4.0 SP3 | Windows 2000 Server | Windows Server 2003 |
|---|---|---|---|---|
| Static IP routing | ● | ● | ● | ● |
| RIP for IP version 1 (RIP v1) | ● | ● | ● | ● |
| RIP for IP version 2 (RIP v2) | | ● | ● | ● |
| RIP for IPX | ● | ● | ● | |
| SAP for IPX | ● | ● | ● | |
| NetBIOS over IPX | ● | ● | ● | |
| Routing and remote access combined into one service | | ● | ● | ● |
| Inclusion of routing and remote access with Windows | | ● | ● | ● |
| OSPF | | ● | ● | ● |
| IP packet filtering | | ● | ● | ● |
| Basic stateful firewall | | | | ● |
| IPX packet filtering | | ● | ● | |
| ICMP router discovery | | ● | ● | ● |

| | | | | |
|---|---|---|---|---|
| DHCP relay agent | ● | ● | ● | ● |
| Network address translation (NAT) | | | ● | ● |
| Remote Authentication Dial-In User Service (RADIUS) client | | ● | ● | ● |
| AppleTalk routing | | | ● | ● |
| Remote access routing (dial-up) | | ● | ● | ● |
| Remote access routing using Point-to-Point Tunneling Protocol (PPTP) virtual private networking | | ● | ● | ● |
| Remote access routing using Layer Two Tunneling Protocol over Internet Protocol Security L2TP/IPSec virtual private networking | | | ● | ● |
| IP-in-IP tunneling | | | ● | |
| SNMP MIB | | | ● | ● |
| LAN and WAN support1 | | ● | ● | ● |
| Multicast IP forwarding | | | ● | ● |
| Graphical administrative tool:<br><br>Routing and RAS Admin | | ● | | |
| Graphical administrative tool:<br><br>Routing and Remote Access snap-in (MMC) | | | ● | ● |
| Command-line tool: Routemon.exe | | ● | | |
| Command-line toolset: Netsh.exe (Netsh routing IP commands) | | | ● | ● |

1. The Routing and Remote Access service works with a wide variety of hardware platforms and can run over any of the many LAN and WAN network adapters supported by Windows. For information about supported network adapters, see "Windows Server Catalog."

IPX routing is available only on Windows NT Server 3.51 Service Pack 2 (SP2), Windows NT Server 4.0 SP3, and the 32-bit versions of Windows 2000 Server. IPX is not available on any version of Windows Server 2003. OSPF is available only on Windows NT Server 3.51 SP2, Windows NT Server 4.0 SP3, and on 32-bit versions of Windows 2000 Server and Windows Server 2003.

For more information about the following topics, see the following references:

- IPv4 Multicasting Technical Reference
- VPN Technical Reference
- SNMP Technical Reference
- NAT Technical Reference

**Unicast IP Routing in a Windows-based Internetwork**

A typical IPv4 internetwork might contain a mix of computers running Windows Server 2003, Windows XP, Windows 2000 Server, Windows 2000 Professional, or UNIX operating systems. These computers might be located in multiple subnets connected by hardware routers from Cisco Systems and software routers running the Windows Server 2003 Routing and Remote Access service. Such an internetwork can easily communicate with computers on the global Internet because the Internet is also an IP internetwork.

A medium-size or enterprise-size Windows IP internetwork typically deploys the Active Directory directory service, DNS, and

DHCP, and the following routing-related services:

- A routing protocol, such as RIP v2 or OSPF, to enable routing information exchange between routers on an IP internetwork
- DHCP relay agents to enable DHCP clients on a subnet with no DHCP server to request IP addresses from a DHCP server located on a different subnet
- IP packet filtering, such as Web traffic filtering or L2TP/IPSec traffic filtering, to allow only specific types of traffic
- ICMP router discovery to enable IP hosts to discover the best default gateway router on a subnet

**How RIP Works**

Routing Information Protocol (RIP) for IP facilitates the dynamic exchange of routing information between RIP routers over IP internetworks. Routers can add or remove routes automatically as networks are added or removed from the internetwork. RIP, the best known and currently most widely used of the distance vector dynamic routing protocols for IP internetworks, is an open standard developed by the Internet Engineering Task Force (IETF).

RIP version 1 (RIP v1), which is now outmoded, was the first routing protocol accepted as a standard for TCP/IP. The updated RIP version 2 (RIP v2) supports simple password authentication (a form of router identification, not a security option) and, more important, provides improved support for classless networks. The Windows Server 2003 Routing and Remote Access service supports both RIP v1 and RIP v2 (for IPv4 only).

### Note

- The Windows Server 2003 Routing and Remote Access service does not support RIPng, the version of RIP for IPv6, or any other IPv6 routing protocol. For information about Windows Server 2003 support for IPv6, including support for static routing in an IPv6 environment, see "How IPv6 Works" in IPv6 Technical Reference.

RIP has its origins in the Xerox Network Systems (XNS) version of RIP. RIP for IP became a popular routing protocol due to its inclusion in the Berkeley Software Distribution (BSD) UNIX (BSD 4.2 and later) as the routed server daemon (a UNIX daemon is similar to a Windows Server 2003 service). The BSD 4.2 version of UNIX, introduced in 1983, incorporated the TCP/IP protocol suite and initiated the explosive growth of TCP/IP and what later became known as the Internet. RIP v1 is defined in RFC 1058 and RIP v2 is defined in RFC 2453. In this document, when neither RIP v1 nor RIP v2 is explicitly cited, information about RIP applies to both versions.

This section describes:

- RIP features appropriate for small- to medium-sized internetworks
- How RIP convergence works
- Basic processes common to RIP v1 and RIP v2
- How RIP v1 works
- How RIP v2 works

**RIP Features Adppropriate for Small- to Medium-Sized Internetworks**

RIP is a useful solution for a small- to medium-sized IP internetwork (10 to 50 networks). As a dynamic routing protocol, it can update topology changes as networks are added or removed and as links fail and are restored.

Although RIP is simple to administer and well supported in the industry, it has some of the limitations inherent to its original LAN-based design. These limitations, in combination, make RIP an inappropriate solution for enterprise-sized IP internetworks.

**RIP Hop Counts**

RIP uses a hop count as the metric to determine the best route among multiple possible routes stored in an IP routing table. The hop count is the number of routers that must be crossed to reach a specific subnet. A computer on the local subnet is one hop, and each router crossed after that is an additional hop.

RIP has a maximum hop count of 15. Therefore, any two hosts can be separated by no more than 15 routers. Subnets that are 16 or more hops away are unreachable and are thus said to be an "infinite" distance away. Any attempt to forward packets to hosts on that subnet result in an ICMP Destination Unreachable–Host Unreachable message from the RIP router, indicating that the packet cannot be delivered. Although the accumulated hop count between any two subnets must not exceed 15 hops, this restriction is not as limiting as it might seem because, if routers are placed in a hierarchical structure, 15 hops can cover a large number of destinations.

The RIP hop count is independent of the Time-to-Live (TTL) field in the IP header. On an IP internetwork, a subnet 16 hops away would normally be reachable for an IP packet with an adequate TTL. However, to the RIP router, that subnet is unreachable.

Using the number of routers to cross as the basis for selecting the best route can produce an unwanted result. For example, if a T1 link is the primary connection between two sites and a slower satellite link is used as a backup in case the T1 link fails, RIP assigns both links the same metric. When a router chooses between these two routes, it might select either one. If the satellite link is selected, traffic is slower. To avoid this type of problem, the cost of each interface is configurable. Typically, the hop count value for slow links is set to multiple hops to ensure that fast links, which are assigned a lower hop count value, take precedence over slow links. In this example, you can assign a higher cost to the satellite interface to ensure that the T1 link is selected as the best route whenever both routes are available. Only if the T1 link fails will the satellite link be selected.

Although the maximum diameter of a RIP internetwork is defined by RFCs 1058 and 2453 to be 15 routers, for a server running the Routing and Remote Access service, the effective diameter is only 14 routers. This is because, for a Routing and Remote Access RIP router, all non-RIP learned routes — such as static routes (including static routes for directly connected networks), OSPF routes, and SNMP routes — have a fixed hop count of 2. Therefore, when the Routing and Remote Access RIP router advertises its directly connected networks, it advertises them at a hop count of 2 even though there is only one physical router to cross.

### Note

- In a networking environment that includes Routing and Remote Access RIP routers, make sure that the accumulated costs (hop counts) between any two endpoints on the internetwork does not exceed 14.

**RIP Routing Table Entries**

When a RIP router first starts up, its routing table contains only the networks to which it is physically connected. This router, as it initiates, sends periodic announcements that contain its routing table entries to inform neighboring RIP routers about the networks that it can reach. In turn, it receives route advertisements from other RIP routers that it adds to its own routing table.

For a subnet that has multiple paths, RIP allows storage of multiple entries in a routing table. For RIP routers that do store a complete list of all network IDs and all possible routes to reach each subnet, the routing table can have hundreds or even thousands of entries in a large multipath IP internetwork. The IP routing process chooses the route with the lowest metric (lowest hop count) as the best route. Because only 25 routes can be sent in a single RIP packet, large routing tables must be sent as multiple RIP messages.

Although in theory RIP can store multiple entries for a multipath subnet in a routing table, typical RIP router implementations — including Windows Server 2003 — store only a single lowest-metric route for any subnet in the routing table. If multiple lowest hop count routes are received by RIP, the first lowest-metric route received is stored in the routing table.

**"RIP Route Advertising"**

Typically, RIP either broadcasts or multicasts announcements. However, RIP can also unicast RIP announcements to specific RIP neighbors.

**RIP broadcasting of announcements for mixed environment**

Typically, a RIP router uses an IP subnet and MAC-level broadcast to advertise the entire contents of its routing table on all attached networks. The announcement interval is every 30 seconds by default. RIP v1 announcements are always broadcasted. RIP v2 announcements can be broadcasted or multicasted.

You must use broadcasting if your network is a mixed environment that includes both RIP v1 and RIP v2 routers. Both RIP v1 routers and RIP v2 routers can process a RIP v2 broadcast announcement.

**RIP multicasting of announcements for the RIP v2 environment**

RIP v2 routers are configured, by default, to broadcast RIP announcements, but they have the capability to multicast RIP announcements to the IP multicast address of 224.0.0.9 instead. The router does not register this multicast address by using Internet Group Management Protocol (IGMP) because the address is in the local network control block range, as defined by the Internet Assigned Numbers Authority (IANA).

Use RIP v2 multicast announcements only if all neighboring RIP routers on the network connected to this interface are

configured for RIP v2. A RIP v1 router cannot process a RIP v2 multicast announcement.

RIP v2 announcements must be multicast if you are performing auto-static updates across demand-dial links. An auto-static update is a one-time, one-way transfer of routing information. A router that has a demand-dial interface configured for auto-static updates sends a message across an active connection requesting all of the routes of the router on the other side of the connection. This request is *not* automatic; the request is sent only when an administrator issues a command or a scheduled task initiates auto-static updates. When the remote router responds, the routes of the remote router are automatically entered in the routing table of the requesting router. The routes are saved in the routing table even if the router is restarted or the interface goes down.

For RIP v2 to support auto-static updates, you must configure a demand-dial interface on each router to use RIP v2 multicast announcements and to accept RIP v2 announcements. Otherwise, the RIP request for routes does not receive a response from the remote router. (OSPF does not support auto-static updates.)

For more information about RIP v2 support for multicast RIP announcements, see "RIP v2 Updates RIP for Modern Internetworks" later in this document.

**RIP unicasting of announcements to specific RIP neighbors**
Although RIP typically broadcasts or multicasts announcements, it can also unicast RIP announcements to specific RIP routers known as RIP neighbors.

Unicasting RIP announcements to RIP neighbors was originally developed specifically to support nonbroadcast multiple access (NBMA) network technologies. NBMA networks, such as Frame Relay, X.25, or Asynchronous Transfer Mode (ATM), can connect more than two routers but cannot broadcast a single frame to all routers attached to that network. However, the configuration of RIP neighbors is also useful for the wider purpose of ensuring that RIP announcements are directed to specific neighboring RIP routers. Configuring a RIP router to unicast RIP announcement to specified neighboring RIP routers helps prevent RIP traffic from being received by any other node.

Configuring RIP routers to unicast to neighboring RIP routers can interfere with the ability of silent RIP hosts to receive RIP traffic. Silent RIP is a feature that lets an IP host listen to routing traffic exchanged by RIP routers and update its routing table accordingly. If you use silent RIP hosts, you can solve this interference problem by adding silent RIP hosts as neighbors. Use a static IP address for a silent RIP host. Otherwise, if you let DHCP assign an IP address to a silent RIP host, that address might change, and this host will no longer receive the unicast RIP messages sent by the RIP router.

How you configure RIP for an NBMA network depends on how Frame Relay virtual circuits appear as network interfaces on a RIP server running Routing and Remote Access. Frame Relay virtual circuits appear in the Routing and Remote Access snap-in either as a single adapter or as multiple adapters. A virtual circuit is a logical path or connection between nodes in a network. After the nodes of a virtual circuit establish a connection, the nodes can communicate as if they were directly connected by physical wires.

The following descriptions assume that you have a Frame Relay network, but other NBMA technologies, such as X.25 and ATM, function in a similar manner:

- **Single-adapter model unicasts RIP announcements.** All virtual circuits on a RIP router function as a single adapter. In the single-adapter model, also known as the NBMA model, the network of the Frame Relay service provider (also called the Frame Relay cloud) is treated as an IP network, and the endpoints on the cloud are assigned IP addresses from a designated IP network ID. To ensure that RIP traffic is received by all of the appropriate endpoints on the cloud, you must configure the Frame Relay interface to unicast its RIP announcements to each appropriate endpoint. On a Routing and Remote Access RIP router, you do this by designating the interface as an NBMA network and by configuring RIP to unicast announcements to the specified neighbor routers.

  In a hub-and-spoke Frame Relay topology in a RIP-based IP internetwork, the Frame Relay interface for the hub router must have split-horizon processing disabled. Otherwise, the spoke routers never receive each other's routes. For more information, see "Split Horizon Processing" later in this section.

- **Multiple-adapter model broadcasts or multicasts RIP announcements.** Each virtual circuit functions as a separate adapter. In the multiple-adapter model, each Frame Relay virtual circuit functions as a point-to-point link with its own network ID, and the endpoints on the cloud are assigned IP addresses from a designated IP network ID. Because each virtual circuit is its own point-to-point connection, you can either broadcast (assuming both endpoints are on the same IP network ID) or multicast RIP announcements.

On a RIP router running Windows Server 2003 Routing and Remote Access, broadcasting or multicasting RIP announcements is configurable in the Routing and Remote Access snap-in by navigating to the **IP Routing\RIP** node, and then using the

**General** tab of the ***InterfaceName*Properties** page. Unicasting RIP announcements to specific RIP neighbors is configurable on the **Neighbors** tab of the same ***InterfaceName*Properties** page.

**RIP Convergence**

Convergence is among the most important factors when considering features that make RIP appropriate for small- to medium-sized internetworks but unsuitable for large enterprise internetworks. Convergence is the process by which routers update routing tables after a change in network topology occurs — the change is replicated to all routers that need to know about it. When convergence is achieved, the internetwork is in a stable state and all routing occurs along optimal paths. If, after convergence is achieved, a link or router fails, the internetwork must reconfigure itself to reflect the new topology.

By default, on a RIP internetwork, each routing table entry that is learned through RIP is given a time-out value of three minutes after the last time the entry was received in a RIP announcement from a neighboring RIP router. Therefore, when a router stops responding because of a hardware or software failure, several minutes can elapse before RIP propagates the topology change throughout the entire internetwork. The delay that occurs before the successful propagation of routes is achieved is known as slow convergence, which is a characteristic of RIP internetworks.

For more information about RIP convergence, including how slow convergence occurs and how to improve convergence time on a RIP internetwork, see "How RIP Convergence Works," described next.

**How RIP Convergence Works**

RIP, like most distance vector routing protocols, announces its routes in an unsynchronized and unacknowledged manner. Although this imprecise announcement scheme can lead to convergence problems, you can modify the announcement algorithms to reduce convergence time on a RIP network in most situations.

**How Count-to-Infinity Occurs**

The classic distance vector convergence problem, known as count-to-infinity, results from the asynchronous announcement scheme. When a RIP router updates its routing table based on RIP announcements from other routers, the RIP router stores only the best route to each network and it updates a lower-cost route with a higher-cost route only if the higher-cost route is announced by the same source as the current lower-cost route. In some situations, as illustrated in the following series of five figures, this causes the count-to-infinity problem.

Assume that the IP internetwork in the following figure has converged. For simplicity, only announcements sent by Router 1 on Network 2 and Router 2 on Network 2 are depicted; announcements sent by Router 1 on Network 1 and Router 2 on Network 3 are not included in the discussion.

**Count-to-Infinity Phase 1: Converged RIP Internetwork**




Now, assume that the link from Router 2 to Network 3 fails, and the failure is sensed by Router 2. As the next figure shows, Router 2 changes the hop count for the route to Network 3 to indicate that it is unreachable, an infinite distance away. For RIP, infinity is 16.

**Count-to-Infinity Phase 2: Link to Network 3 Fails**




However, before Router 2 can advertise the new hop count to Network 3 in a scheduled announcement, Router 2 receives an

announcement from Router 1. The Router 1 announcement contains a route to Network 3, which is two hops away. Because two hops away is a better route than 16 hops, Router 2 updates its routing table entry for Network 3, changing it from 16 hops to 3 hops, as shown in the following figure.

**Count-to-Infinity Phase 3: Router 2 Updates from Router 1**




When Router 2 announces its new routes, Router 1 notes that Network 3 is available three hops away through Router 2. Because the route to Network 3 on Router 1 was originally learned from Router 2, Router 1 updates its route to Network 3 to four hops, as shown in the following figure.

**Count-to-Infinity Phase 4: Router 1 Updates from Router 2**




Now, when Router 1 announces its new routes, Router 2 notes that Network 3 is available four hops away through Router 1. Because the route to Network 3 on Router 2 was originally learned from Router 1, Router 2 updates its route to Network 3 to five hops, as shown in the following figure.

**Count-to-Infinity Phase 5: Router 2 Updates Again from Router 1**




The two routers continue to announce routes to Network 3 with higher and higher hop counts until infinity (16) is reached. At that point, Network 3 is considered unreachable, and the route to Network 3 is eventually timed out of the routing table. This routing loop is known as the count-to-infinity problem.

The count-to-infinity problem is one of the reasons why the maximum hop count of RIP for IP internetworks is set to 15 (16 for unreachable). A higher maximum hop count would make the convergence time longer when count-to-infinity occurs.

Note also that during the count-to-infinity process in the preceding example, the route from Router 1 to Network 3 is through Router 2, and the route from Router 2 to Network 3 is through Router 1. A routing loop exists between Router 1 and Router 2 for Network 3 for the duration of the count-to-infinity problem.

**Reducing Convergence Time**

To speed recovery of RIP internetworks when topology changes occur and to help avoid the count-to-infinity problem and routing loops that persist for the duration of the count-to-infinity process, an administrator can modify the RIP announcement process by using the following algorithms:

- Split horizon
- Split horizon with poison reverse
- Triggered updates

On a Routing and Remote Access RIP router, these options are configurable in the Routing and Remote Access snap-in by navigating to the **IP Routing\RIP** node in the console tree, and then accessing the **Advanced** tab of the appropriate ***InterfaceName*Properties** page. By default, all three options are enabled to improve convergence on a RIP internetwork.

**Split Horizon**
Split horizon is a routing technique that reduces convergence time by preventing or reducing the occurrence of the count-to-infinity problem and the associated routing loops that can occur in distance vector routing protocols such as RIP. When split-horizon processing is enabled, information about routes is not announced over the router interface through which that information was received. Thus, a router configured to use split horizon processing omits routes learned from a specific neighboring router in updates that it returns to that neighbor. If split-horizon processing is not configured, routes learned on a network are announced on the same network. The Windows Server 2003 Routing and Remote Access service enables split-horizon processing by default.

Split horizon eliminates the count-to-infinity problem and routing loops during convergence in single-path internetworks and reduces the chances of encountering count-to-infinity problems in multipath internetworks. The following figure shows how split horizon processing keeps the RIP router from advertising routes in the direction from which they were learned.

**Split Horizon Announcement Process**




**Split Horizon with Poison Reverse**
Split horizon with poison reverse processing refers to routing updates that explicitly indicate that a network or a subnet is unreachable, rather than merely implying that a network is unreachable by not including it in updates. Whereas the simple split horizon method omits routes learned from one neighbor in updates sent to that neighbor, the split horizon with poison reverse method includes such routes in updates, but announces them with a hop count of 16, indicating that the subnet for these routes is unreachable. Split horizon with poison reverse processing can be enabled only if split horizon processing is enabled.

In a single-path internetwork, split horizon with poison reverse has no greater benefit than split horizon. However, in a multipath internetwork, split horizon with poison reverse greatly reduces count-to-infinity problems and routing loops. Count-to-infinity problems can still occur in a multipath internetwork because routes to networks can be learned from multiple sources.

In the following figure split horizon with poison reverse advertises learned routes as unreachable in the direction from which they are learned. Split horizon with poison reverse, however, has the disadvantage of requiring additional RIP message overhead because all networks are advertised.

**Split Horizon with Poison Reverse Announcement Process**




On a Windows Server 2003 Routing and Remote Access RIP router, split-horizon processing and poison-reverse processing (both enabled by default) are configurable in the Routing and Remote Access snap-in by navigating to the **IP Routing\RIP** node in the console tree, and then selecting or clearing the **Enable split-horizon processing** and **Enable poison-reverse processing** options on the **Advanced** tab of the ***InterfaceName*Properties** page. Alternatively, if you do not want both options enabled, you can select **Enable split-horizon processing** and clear **Enable poison-reverse processing**. However, you cannot select **Enable poison-reverse processing** unless you first select **Enable split-horizon processing**.

**Triggered Updates**

In addition to making periodic announcements to other local RIP routers, a RIP router can also communicate routing information through triggered updates. A triggered update is an update that is announced as soon as a network topology change occurs; only the change is announced. Triggered updates improve the convergence time of RIP internetworks but do so at the expense of additional network traffic.

The trigger for an update to the routing table is one of the following:

- The router starts advertising the network ID for a subnet attached to a newly operational interface on the router.
- The router starts advertising the network ID for a subnet attached to a newly inoperational interface on the router as 16 hops (unreachable).
- The router modifies its routing table in response to an update from a neighboring router.
- The router removes a routing table entry because it has timed out.

For example, if a network becomes unavailable due to a problem such as a link or router failure, a RIP router updates its own routing table and can use a triggered update to announce the unavailable network with a hop count of 16. An interval to wait between sending triggered updates is typically specified on the router (on a Routing and Remote Access RIP router, the default is 5 seconds). Otherwise, if all routers immediately send triggered updates, these updates could cause a cascade of traffic across the IP internetwork.

On a Windows Server 2003 Routing and Remote Access RIP router, triggered updates processing (enabled by default) is configurable in the Routing and Remote Access snap-in by navigating to the **IP Routing\RIP** node in the console tree, and then selecting the **Enable triggered updates** option on the **Advanced** tab of the ***InterfaceName*Properties** page. The maximum time (in seconds) that elapses between triggered updates is configurable by using the **General** tab on the **RIP Properties** page.

**Basic Processes Common to RIP v1 and RIP v2**

The standard operation of a RIP for IP router consists of three processes: initialization, when the router learns the routes of the internetwork from neighboring routers; ongoing periodic advertisements; and advertisement of unreachable routes when a router is brought down intentionally or fails.

**Initialization**

During startup, the RIP router announces its locally attached networks on all of its interfaces. Neighboring RIP routers process the RIP announcement and add routes to their routing tables, as appropriate.

The initializing RIP router also sends a general RIP request to all locally attached networks. The general RIP request is a special RIP message that requests all routes. Neighboring RIP routers receive the general RIP request and send a unicast reply to the requesting router. The initializing RIP router uses the replies to build its routing table.

**Ongoing Maintenance**

By default, a RIP router announces its routes on all of its interfaces every 30 seconds. The type of routing announcement

depends on whether the RIP router is configured for split horizon or split horizon with poison reverse processing. The RIP router also continuously listens for RIP announcements from neighboring routers in order to add or update the routes in its own routing table.

**Advertisement of Unreachable Routes**
RIP routers send announcements for routers that are shut down by administrators for reasons such as routine maintenance and for links or routers that fail unexpectedly.

**Administrative Router Shutdown**
If a RIP router is turned off correctly through an administrative action, it sends a triggered update to neighboring routers on all locally attached networks as it shuts down, to inform neighboring routers that networks available through it have become unavailable. The triggered update announces the networks with a hop count of 16 (unreachable). The neighboring RIP routers then use triggered updates to propagate this topology change throughout the IP internetwork.

**Response to Topology Changes**
As dynamic routers, functioning RIP routers react to changes in the internetwork topology from downed links or downed routers:

**Downed link**
If a link that corresponds to one of the router's interfaces fails, and if the failure is detected by the interface hardware and is indicated to the router, the router sends this change out to neighboring routers as a triggered update.

**Downed router**
If a router stops responding due to a power outage or other hardware or software failure, it cannot inform neighboring routers that the networks available through it have become unavailable. To prevent the lingering existence of unreachable networks in routing tables, each route that is learned by RIP has a maximum lifetime of 3 minutes by default. If the entry is not refreshed by the receipt of another announcement in 3 minutes, the entry's hop count is changed to 16 (unreachable) and is, therefore, eventually removed from the routing table.

Thus, if a RIP router fails unexpectedly, it takes up to 3 minutes for the neighboring routers to mark as unreachable the routes that are learned from the failed router. Only then do they propagate the topology change throughout the internetwork by using triggered updates.

**How RIP v1 Works**
RIP v1, defined in RFC 1058, is still deployed in many small- to medium-sized intranets. However, RIP v1 is now considered outdated because it does not fully support classless addressing — including variable length subnet masking (VLSM), which enables configuring subnets of variable size, and classless interdomain routing (CIDR), which enables route summarization.

RIP v2 and OSPF do support CIDR and VLSM, so there is no reason to deploy RIP v1 as the primary routing protocol even for a small network. However, if routers that do not support RIP v2 exist on an IP internetwork, it is possible to deploy a mixed environment that includes both RIP v1 and RIP v2 routers. For more information about mixed environments, see "Mixed RIP v1 and RIP v2 Environments" later in this section.

**RIP v1 Message Format**
RIP messages are sent as User Datagram Protocol (UDP) messages from the router interface IP address and UDP port 520 to the subnet broadcast IP address and UDP port 520. The RIP v1 message consists of a 4-byte RIP header and up to 25 RIP routes. The maximum size of the RIP message is 504 bytes. With the 8-byte UDP header, the maximum size of the RIP message is a 512-byte IP payload. The following figure illustrates the format of a RIP v1 message.

### Message Format for RIP v1




The following table describes the fields depicted in the figure. The four fields named **Must Be Zero** were reserved for future use; as will be described later, RIP v2 uses three of these fields.

### RIP v1 Message Fields

| Field | Description |
|---|---|

| | |
|---|---|
| Command | A 1-byte field that contains a value of either 1 or 2, indicating request or response, respectively.<br>• **Request.** A value of **1** indicates a RIP request to neighboring routers asking for all (a general RIP request) or part of their routing tables.<br>• **Response.** A value of **2** indicates a RIP response from a neighboring router, consisting of all or part of the neighbor's routing table. A RIP response can be sent in response to a RIP request from a requesting router, or as a periodic or triggered update message generated by the sending router. |
| Version | A 1-byte field that is set to **1** to indicate RIP v1. |
| Family Identifier | A 2-byte field that identifies the routable protocol family. In an IP internetwork, this field is set to **2** to indicate the IP routable protocol family. |
| IP Address | A 4-byte field that is set to the IP network ID, which can be any of the following:<br>• A class-based network ID<br>• A subnetted network ID (advertised only within the subnetted network)<br>• An IP address (for a host route)<br>• 0.0.0.0 (for the default route)<br>For a general RIP request, the **IP Address** field is set to **0.0.0.0**. |
| Metric | A 4-byte field that indicates the number of hops to the IP subnet. This field must be a value from 1 to 16. The metric is set to **16** in a general RIP request or to indicate that the subnet is unreachable in a RIP response. |

**Why RIP v1 Is Outdated**

The following sections explain the primary historical and technical reasons that make RIP v1 out-of-date in a modern IP internetwork.

**MAC-level broadcasting of RIP announcements designed for LANs**

RIP v1 was designed in 1988 to satisfy the dynamic routing needs of LAN-based IP internetworks. Shared-access LAN technologies, such as Ethernet or Token Ring, support MAC-level broadcasting, in which a single packet can be received and processed by nodes on a subnet. All RIP v1 route announcements are addressed to the IP subnet (all host bits are set to 1) and MAC-level broadcast addresses. Non-RIP nodes also receive RIP v1 announcements.

In modern IP internetworks, the use of MAC-level broadcasts is unsatisfactory because all nodes must process all broadcasts. For large or very large RIP internetworks, the amount of broadcast traffic on each subnet can become substantial enough to affect node performance.

**Subnet mask not announced with the route**

RIP v1 originated in an era when the Internet was still using network IDs based on the standard Internet address classes. Today, however, the use of CIDR (for route summarization) and VLSM (for configuring subnets of varying size) is widespread, making the standard address classes obsolete.

RIP v1 was designed for class-based IP internetworks where the network ID can be determined from the values of the first 3 bits of the IP address in the RIP route. Because the subnet mask is not included with the route with RIP v1, the RIP router must determine the network ID based on a limited set of information. For each route in a RIP v1 message, the RIP v1 router performs the following process:

- If the network ID fits one of the standard address classes (Class A, Class B, or Class C), the default class-based subnet mask is assumed.
- If the network ID does not fit any address class, one of the following occurs:
    - If the network ID fits the subnet mask of the interface on which it is received, the subnet mask of the interface on which it was received is assumed.

- If the network ID does not fit the subnet mask of the interface on which it is received, the network ID is assumed to be a host route with the subnet mask 255.255.255.255.

As a result of the preceding assumptions, RIP v1 can misinterpret summarized routes as a single class-based network ID rather than as a range of class-based network IDs. In addition, with noncontiguous subnets, RIP v1 can misinterpret subnet routes that are advertised outside the network ID being subnetted as host routes.

RIP v1 can support only physically contiguous, fixed-size subnets. RIP v1 routers can support subnetted environments to a limited extent because they do not advertise the subnets of a subnetted class-based network ID outside the subnetted region of the IP internetwork. Subnets of a network ID in a RIP v1 environment must be contiguous, because only the class-based network ID is advertised outside the subnetted environment. Noncontiguous subnets, also known as discontiguous subnets, occur when another network with a different network ID separates subnets of a classful network. If the subnets of a network ID are noncontiguous, the class-based network ID is announced by separate RIP v1 routers in different parts of the internetwork. As a result, IP traffic can be forwarded to the wrong subnet.

**No protection from rogue RIP routers**
RIP v1 does not provide any protection from a rogue RIP router that starts up on a subnet and announces false or inaccurate routes. RIP v1 announcements are processed regardless of their source. A malicious user can use this lack of protection to overwhelm RIP routers with hundreds or thousands of false or inaccurate routes. However, on a RIP router running Windows Server 2003 Routing and Remote Access, you can configure RIP to accept only RIP announcements from specified IP addresses.

**How RIP v2 Works**
RIP v2, which is defined in RFC 2453, reduces or eliminates some of the problems associated with RIP v1. The decision to refine RIP was controversial in the context of newer, smarter routing protocols, such as OSPF. However, RIP was extended in the new RIP v2 implementation because RIP has the following advantages over OSPF:

- **Easy to implement.** In its simplest form (when defaults are accepted), implementing RIP for IP is as easy as configuring IP addresses and subnet masks for each router interface and then turning on the router.

- **Widely in use.** RIP is installed in a large number of small-sized and medium-sized IP internetworks.

**RIP v2 Updates RIP for Modern Internetworks**
RIP v2 adds several key features to RIP v1 that help today's IP internetworks minimize broadcast traffic, conserve IP addresses by using VLSM and CIDR, and help secure the routing environment from incorrectly or maliciously configured routers.

**Support for multicast RIP announcements**
As explained earlier, RIP v2 supports sending RIP announcements to the IP multicast address of 224.0.0.9. In addition, in contrast to RIP v1, non-RIP nodes are not disturbed by RIP v2 router announcement traffic when multicasting is used.

A disadvantage of the new multicast announcement feature, however, is that silent RIP nodes, if they exist, must also listen for the multicast traffic that is sent to 224.0.0.9. If your computers use silent RIP, verify that the silent RIP nodes can listen for multicasted RIP v2 announcements before you deploy multicasted RIP v2. The Windows Server 2003 (or Windows 2000 Server) Routing and Remote Access service can be configured to send out RIP v1, RIP v2, or both types of packets.

For more information about multicast RIP announcements, see "RIP Route Advertising" earlier in this section.

**Support for subnet masks**
RIP v2 announcements send the subnet mask (also known as a network mask) along with the network ID. Because RIP v2 has this capability, it supports VLSM and CIDR, which means that RIP v2 can operate in variable-length subnetted internetwork environments and can use route summarization. It also means that RIP v2 can manage noncontiguous subnets without experiencing the routing errors that occur with RIP v1.

**Support for simple password authentication**
RIP v2 supports the use of simple password authentication to verify the origin of incoming RIP announcements. Simple password authentication for RIP v2 is defined in RFC 1723. Enabling simple password authentication helps prevent an unauthorized or incorrectly configured RIP router from corrupting RIP routes on a subnet.

Newer authentication mechanisms, such as Message Digest 5 (MD5), are also available in some RIP v2 environments. However, Windows Server 2003 Routing and Remote Access supports only simple password authentication for RIP v2.

If you use RIP v2 simple password authentication, you must configure all RIP v2 interfaces on the same network with the same case-sensitive password. A RIP v2 router discards RIP announcements received from any RIP v2 interface whose password does not match the password configured on the interface of this RIP v2 router. You can use the same password for all of the

networks on your IP internetwork, or you can vary the password for each network.

**Note**

- Simple password authentication is a form of router identification. It does not provide security for RIP messages. Because the password is sent in plaintext, any user with a protocol analyzer (also known as a network sniffer), such as the Network Monitor tool, can capture the RIP v2 announcements and view the authentication password.

**Support for RIP v1 and RIP v2 compatibility**

RIP v1 was designed to be compatible with future versions of RIP. If a RIP v1 router receives a message and if the RIP version indicated in the RIP header is not 1, the RIP v1 router does not discard the RIP announcement but instead processes only the RIP v1 defined fields.

In addition, RIP v2 routers send a RIP v1 response to a RIP v1 request, unless they are configured to send only RIP v2 announcements.

**RIP v2 Message Format**

RIP v1 routers can process RIP v2 announcements in a mixed environment that contains both types of routers, because RIP v2 does not modify the structure of the RIP message format. Instead, RIP v2 makes use of fields that were defined in RIP v1 as **Must be Zero**.

The **Command**, **Family Identifier**, **IP Address**, and **Metric** fields in RIP v2 are used as they are for RIP v1. For RIP v2, the **Version** field is set to **2** to indicate that the message is a RIP v2 message. The following figure illustrates the RIP v2 message format.

**Message Format for RIP v2**




RIP v2 uses three of the four fields that were named **Must Be Zero** in RIP v1. The following table describes the new fields for RIP v2.

**Additional Message Fields Defined in RIP v2**

| Field | Description |
|---|---|
| Route Tag | A 2-byte field used to distinguish routes announced by RIP routers from non-RIP routes. The **Route Tag** field is defined by RFC 1723.<br><br>Routers that can support multiple routing protocols typically allow the **Route Tag** to be configured to indicate routes imported from differing sources. For example, an administrator typically sets the **Route Tag** for routes imported from Exterior Gateway Protocol (EGP) or OSPF to an arbitrary value to differentiate them from RIP routes.<br><br>The Windows Server 2003 Routing and Remote Access snap-in supports the configuration of the **Route Tag** on RIP v2 interfaces on the **General** tab of the ***InterfaceName* Properties** page. Routing and Remote Access does not allow you to specify a **Route Tag** by route, only by interface. For example, a route tag can be used on a router that has OSPF enabled on all of its interfaces except one, on which RIP is used. The route tag for routes on the RIP interface can be configured to tag all of the OSPF routes learned on all the other interfaces. |
| Subnet Mask | A 4-byte field that contains the subnet mask (also known as a network mask) of the network ID specified in the **IP Address** field. |

| | |
|---|---|
| Next Hop | A 4-byte field that contains the forwarding IP address (also known as the gateway address) for the network ID specified in the **IP Address** field. If the next hop is set to **0.0.0.0**, the next-hop IP address for the route is assumed to be the IP address of the source of the route announcement.<br><br>The **Next Hop** field is designed to prevent the occurrence of non-optimal routing situations. The usefulness of the **Next Hop** field is illustrated by the following example:<br><br>• If the **Next Hop** field is not used, when a router announces a host route for a host that resides on the same subnet as the router interface that is advertising the route, the next-hop IP address for the host route is the IP address of the router's interface, not the IP address of the host. Other routers on that subnet that receive the announcement forward packets destined for the host's IP address to the announcing router's IP address rather than to the host. This routing issue, which can occur with RIP v1, is avoided by RIP v2 because it uses the **Next Hop** field.<br><br>• If the **Next Hop** field is used, the router announces the host route with the host's IP address in the **Next Hop** field. Other routers on that subnet that receive the announcement correctly forward packets destined for the host to the host's IP address, rather than incorrectly forwarding the packets to the announcing router.<br><br>Because the **Next Hop** field becomes the **Gateway** address field in the routing table, the IP address in the **Next Hop** field must be directly reachable through a router interface. |

**RIP v2 Simple Password Authentication Message Format**

The simple password authentication process for RIP v2 announcements uses the first route entry in the RIP message to store authentication information. The first route entry must be used for this authentication mechanism, which leaves a maximum of 24 routes available in a RIP v2 authenticated announcement (because the RIP message format allows up to 25 routes per RIP message). To indicate authentication, the **Family Identifier** field is set to **0xFFFF**. The **Authentication Type** field (which is usually used as the **Route Tag** field) indicates the type of authentication being used. A value of **1** in the **Authentication Type** field indicates simple password authentication.

The 16-byte **Authentication Value** field, which follows the **Authentication Type** field, stores the left-justified, null-padded, case-sensitive, plaintext password used for simple password authentication. The following figure illustrates the RIP v2 authentication message format for plaintext passwords.

**RIP v2 Message Format Using Authentication**




RIP v1 routers disregard the first route in a RIP v2 authenticated announcement because RIP v1 does not recognize the **Family Identifier** (that is, **0xFFFF**) for the route.

**Mixed RIP v1 and RIP v2 Environments**

To take advantage of all RIP v2 features, use only RIP v2 in a RIP internetwork. For internetworks that contain one or more routers that do not support RIP v2, it is possible to establish a mixed environment that contains both RIP v1 and RIP v2 routers. However, you must exercise caution if you establish a mixed environment to avoid experiencing routing problems. Follow these guidelines to avoid routing problems in a mixed environment:

- Because RIP v1 routers cannot interpret the **Subnet Mask** field in the route, RIP v2 routers must not announce routes that can be misinterpreted by a RIP v1 router. Therefore, do not use VLSMs or noncontiguous subnets in a mixed environment.

- Because RIP v1 routers will misinterpret a summarized route as a single network ID rather than as a range of network IDs, do not use summarized routes in a mixed environment.

- For a RIP v2 router interface to make announcements in a way that allows RIP v1 routers to process the announced routes, the RIP v2 routers must summarize subnet routes when announcing outside a subnetted environment. A specific subnet route announced to a RIP v1 router can be misinterpreted as a host route. Therefore, in a mixed environment, use

only contiguous, fixed-size subnets or clear the **Disable subnet summarization** setting on the **Advanced** tab of each interface.

- If RIP v2 routers are on the same subnet as RIP v1 routers, you must configure the interface for each RIP v2 router to broadcast its announcements on that subnet. RIP v1 routers cannot process multicasted RIP v2 announcements.

**How OSPF Works**

Open Shortest Path First (OSPF) for IP enables OSPF routers to dynamically exchange routing information with each other over complex IP internetworks. Routers can add or remove routes automatically as networks are added or removed from the internetwork, dynamically building and synchronizing a database of the OSPF network topology. As its name implies, OSPF is designed to calculate the shortest path to any destination within an OSPF autonomous system (AS). OSPF, the best known and most widely used link state routing protocol, is an open standard developed by the Internet Engineering Task Force (IETF) as an alternative to RIP. OSPF is defined in RFC 2328. The Windows Server 2003 Routing and Remote Access service supports OSPF for IPv4 only.

### Notes

- The OSPF routing protocol is available only on 32-bit versions of Windows Server 2003
- The Windows Server 2003 Routing and Remote Access service does not support OSPF (or any other dynamic routing protocol) for IPv6. For information about Windows Server 2003 support for IPv6, including support for static routing in an IPv6 environment, see "IPv6 Routing" in "How IPv6 Works" in IPv6 Technical Reference.

This section describes:

- OSPF features appropriate for large internetworks
- How OSPF convergence works
- How OSPF forms adjacencies to synchronize the LSDB
- How OSPF multiple-area autonomous systems work

**OSPF Features Appropriate for Large Internetworks**

OSPF, an Interior Gateway Protocol (IGP) used for routing IP traffic within a single autonomous system, is commonly used on a multipath, dynamic IP internetwork that consists of more than 50 (but fewer than 100) networks.

Typically, OSPF is deployed in large private IP internetworks, including internetworks that span multiple geographical regions, nations, or continents. OSPF functions efficiently in a WAN environment by using summarized routes to reduce the size of routing tables.

You can also easily take advantage of the advanced features of OSPF in a much smaller network by accepting the default OSPF settings and the default settings for OSPF interfaces.

Both the Internetwork Operating System (IOS) used by many hardware routers and the Windows Server 2003 Routing and Remote service support OSPF.

**OSPF Database, Areas, and Routers**

In an OSPF routed internetwork, an autonomous system is made up of one or more areas. An OSPF area is made up of a logical collection of networks, routers, and links, and each area includes special OSPF routers that enable communication between the constituent areas. Each OSPF area uses a database to store its routing information. OSPF updates this database dynamically as network topology changes.

**OSPF database**

An OSPF area has a database that is shared by the routers in that area, a method to keep the database updated, and a method to calculate the routing table from the database.

**All OSPF routers in an area share the same database**

An OSPF router stores information about the state of each link between itself and other routers in an OSPF-specific database called a link state database (LSDB). The LSDB is a topological map of all routers in the same OSPF area. The LSDB maps either all routers within the autonomous system (if only one area exists) or it maps all routers within the local OSPF area (if multiple

areas exist). If multiple areas exist, each OSPF router stores a separate LSDB for each connected area.

**OSPF routers exchange advertisements to create and update the database**
OSPF routers compile the shared LSDB by flooding link state advertisements (LSAs) within the autonomous system (if only one OSPF area exists) or within the local area (if multiple OSPF areas exist). Flooding means sending traffic received on one interface out all of the other interfaces on that device. An LSA is a multicast packet that an OSPF router uses to advertise information about network neighbors and path costs. Each OSPF router in an area must receive a valid LSA from each other router.

**OSPF routers calculate the routing table from the database**
After an OSPF router compiles its LSDB from the LSAs that it receives from other routers, each router, using itself as the root, calculates a tree of shortest paths to each other router and subnet, and then creates IP routing table entries from this tree, storing a single entry for each subnet.

For detailed information about how OSPF manages its database, advertisements, and routing table, see "How OSPF Convergence Works" later in this section.

**OSPF areas**
Each OSPF area is identified by an area identification number called an area ID. The 32-bit OSPF area ID is expressed in dotted decimal notation. The area ID is an administrative identifier that is unrelated to an IP address or an IP network ID. However, if all networks within one OSPF area correspond to a single subnetted network ID, for administrative convenience you can set the area ID to match the network ID. For example, if an area contains all of the subnets of the IP network 10.1.0.0/16, you can set the area ID to 10.1.0.0.

The following table describes the possible types of OSPF area. The only required area type is a backbone area.

### Types of OSPF Areas

| Area Type | Description |
|---|---|
| Standard | An OSPF area that propagates intra-area, inter-area, and external routes:<br>• Intra-area routes are communicated within the local OSPF area.<br>• Inter-area routes are communicated between OSPF areas (connected by a backbone area) that are located in the same autonomous system. OSPF uses summarized routes for inter-area routes.<br>• External routes are communicated from an external route source into the OSPF autonomous system by an Autonomous System Border Router (ASBR). An external resource can be either from another OSPF autonomous system or from a non-OSPF source, such as another routing protocol.<br>The LSDB in an area stores OSPF router LSAs for intra-area routes, summary LSAs for inter-area routes, and external route LSAs for external routes. |
| Backbone | A standard OSPF area that is either the sole area in the OSPF autonomous system or the central area to which all other OSPF areas of an OSPF autonomous system connect. The OSPF backbone area requires the reserved area ID of 0.0.0.0; an alternative notation, the area ID of 0, is sometimes used to denote the backbone area.<br>When more than one area exists in an OSPF autonomous system, the OSPF backbone area (typically, a single high-bandwidth network) distributes routing information among all attached areas. Traffic between areas is first routed to the backbone, then routed to the destination area, and then finally routed to the destination host within the destination area. |
| Transit | A non-backbone standard OSPF area that connects another OSPF area (that does not have a physical link to the backbone) to the backbone by using a logical point-to-point connection called a virtual link. |

| | |
|---|---|
| Stub | An OSPF area that, like a standard OSPF area, accepts and advertises inter-area and intra-area routes but, unlike a standard area, does not accept or distribute external routes from external OSPF autonomous systems or from non-OSPF sources.<br><br>Internal OSPF routers in a stub area use the default route (destination 0.0.0.0, network mask 0.0.0.0) to reach all destinations outside the autonomous system. A stub area, typically, has only one router (an Area Border Router, or ABR) that connects the stub area to the rest of the autonomous system. |
| Totally Stub | A type of stub area that propagates only intra-area routes and the default route within the area; it does not accept or advertise either inter-area routes or external routes.<br><br>Internal OSPF routers in a totally stub area use the default route to send traffic destined outside the area to the local ABR, which forwards the traffic towards its destination in another OSPF area or outside the autonomous system. |

On a Windows Server 2003 Routing and Remote Access OSPF router, OSPF areas are viewable in the Routing and Remote Access snap-in by right-clicking **OSPF** under **IP Routing** in the console tree, and then clicking **Show Areas**.

For more information about OSPF areas, see "How OSPF Multiple-Area Autonomous Systems Work" later in this section.

**OSPF Routers**

When an OSPF autonomous system contains two or more areas and connects to other destinations outside the OSPF autonomous system, the routers it contains can operate in one of four roles. The following table describes these OSPF routers.

**Types of OSPF Routers**

| Router Type | Description |
|---|---|
| Internal router | An OSPF router that has two or more interfaces that all connect within the OSPF area in which the router resides. All internal routers within an area share the same LSDB. |
| Area border router (ABR) | An OSPF router that has two or more interfaces that connect directly to networks in two or more OSPF areas. One interface must connect to the backbone area; thus, an ABR is always a backbone router but the converse is not always true — a backbone router is not necessarily an ABR. An ABR stores multiple LSDBs, one for each attached area, including one for the backbone area.<br><br>An ABR routes traffic between the backbone area and all other areas connected to the backbone by sending routes for its attached area to the backbone, which, in turn, distributes this information to other OSPF areas not attached to this router within the autonomous system. |
| Backbone router | An OSPF router that has at least one interface connected to the backbone area and that routes traffic within the backbone area. All internal routers of the backbone area and all ABRs are backbone routers. |
| Autonomous system boundary router (ASBR) | An OSPF router that imports routing information from external sources and advertises external destinations throughout its own autonomous system. An ASBR in an OSPF autonomous system can connect to and exchange routes with:<br>• An ASBR in another OSPF autonomous system, by connecting the two autonomous systems.<br>• Non-OSPF sources, by connecting its own autonomous system to the external source. For example, an ASBR can connect to an external network running a different interior gateway protocol, such as RIP for IP.<br><br>Only a standard OSPF area can contain an ASBR. Each router in an OSPF autonomous system can reach the ASBRs within its autonomous system. |

**OSPF Unitless Metrics**

In contrast to the RIP hop count cost metric, OSPF uses a unitless metric to indicate the preference for one link over another. Factors in the metric include the following:

- **Delay.** The time that a packet requires to reach the destination network. Delay indicates the speed (LAN links have a low delay; WAN links have a high delay) or congestion level of the path.

- **Throughput (link speed).** The amount of data that can be sent along the path per second. Throughput does not necessarily indicate the bit rate of the link because, for example, a very busy Ethernet link might have a lower throughput

than an unused 64-kilobits per second (Kbps) WAN link.

- **Reliability.** The path constancy. Some types of links are more likely to fail than others. For example, leased line WAN links are more reliable than dial-up lines.

The cost value is advertised as the link cost for this interface in the LSA of the router. The lower the cost assigned to an interface, the more often the OSFP router uses the interface to forward packets. Therefore, typically, you assign lower costs to the interfaces of preferred paths for IP traffic. The accumulated cost for network segments in an OSPF internetwork can be as great as 65,535, which enables OSPF to accommodate enterprise-size networks.

**OSPF Route Advertising**
OSPF route advertising includes the characteristics described in the following sections.

**Advertises subnet mask**
OSPF advertises the subnet mask of a route with its network ID. As a result, OSPF supports VLSM and CIDR and, therefore, supports subnetting, noncontiguous subnets, and route summarization.

**Advertises external routes**
In OSPF, an external route is a route that is learned from a non-OSPF source, such as a static route manually configured by a network administrator or a RIP route learned from a network that uses the RIP routing protocol. Routes outside an OSPF autonomous system are advertised within the autonomous system, which enables OSPF routers to calculate the least-cost route to external networks. OSPF can use the default route (destination 0.0.0.0, network mask 0.0.0.0) to summarize routes imported from external sources.

For more information about external routes, see “Propagating External Source Routes to the OSPF Autonomous System” later in this section.

**Advertises summarized routes**
In an OSPF autonomous system with multiple areas, OSPF can be configured to send summary route advertisements in the format {*Destination, Network Mask*} to advertise networks located inside one area to other OSPF areas. The destination and network mask specify a range of addresses to be summarized. Summarizing — consolidating multiple routes into a single advertisement — is configured on an ABR at the boundary between an OSPF area and the backbone area. After receiving the summarized route from one area, the ABR forwards the summarized route through the backbone area and then through another ABR into another area.

Because OSPF can include multiple areas and can use summarized routes (which reduces the size of OSPF routing tables), OSPF can scale to large and very large internetworks. In contrast, RIP internetworks cannot be subdivided and RIP v1 performs no route summarization beyond summarizing for all subnets of a class-based network ID.

For more information about using summarized routes for inter-area routing, see “Using Route Summarization Between OSPF Areas” later in this section.

**Advertises based on network type**
The address type, multicast or unicast, used for an OSPF announcement is determined by the type of network — broadcast, point-to-point, or nonbroadcast multiple access (NBMA) — to which the OSPF interface is connected.

**OSPF multicasts announcements on broadcast networks (default)**
A broadcast network is a multiple access network, that is, more than two routers are connected to the network, which can broadcast a single frame to all routers attached to that network. Ethernet, Token Ring, and Fiber Distributed Data Interface (FDDI) are broadcast networks.

On broadcast networks, OSPF routers use the following two reserved multicast IP addresses to send OSPF messages:

- **224.0.0.5 - AllSPFRouters.** Used to send OSPF messages to all OSPF routers on the same subnet. OSPF routers use the **AllSPFRouters** multicast address to send OSPF Hello messages to each other in order to discover each other. In addition, the OSPF designated router (DR) and backup designated router (BDR) use this address to send Link State Update and Link State Acknowledgment messages. (For more information about OSPF designated routers, see “How OSPF Forms Adjacencies to Synchronize the LSDB” later in this section.)

- **224.0.0.6 - AllDRouters.** Used to send OSPF messages to all OSPF designated routers and BDRs on the same subnet. All OSPF routers, except designated routers themselves, use this multicast address to send Link State Update and Link State Acknowledgment messages to designated routers and BDRs.

The multicast protocol, Internet Group Management Protocol (IGMP), is not needed because these addresses are in the link-local range. The link-local range refers to addresses in the range of 224.0.0.0 to 224.0.0.255 (224.0.0.0/24). This multicast range is reserved for local subnet multicast traffic. IP routers do not forward packets sent to addresses in this range.

**OSPF multicasts announcements on point-to-point networks**

A point-to-point network is a network that connects only two routers. T1/E1, T3/E3, ISDN, and other types of dial-up WAN links as well as leased-line WAN links are point-to-point networks.

On point-to-point networks, OSPF routers use the **AllSPFRouters** multicast IP address (**224.0.0.5**) to send all OSPF messages.

**OSPF unicasts announcements to specific OSPF neighbors on NBMA networks**

A nonbroadcast multiple access (NBMA) network is a network that can connect more than two routers but cannot broadcast a single frame to all routers attached to that network. X.25, Frame Relay, and Asynchronous Transfer Mode (ATM) are NBMA networks. In addition to being unable to broadcast, NBMA networks also cannot multicast.

How you configure OSPF for an NBMA network depends on how Frame Relay virtual circuits appear as network interfaces on an OSPF server running Routing and Remote Access. Frame Relay virtual circuits appear in the Routing and Remote Access snap-in either as a single adapter or as multiple adapters. The following descriptions assume that you have a Frame Relay network, but other NBMA technologies, such as X.25 and ATM, function in a similar manner.

- **Single-adapter model unicasts OSPF announcements.** All virtual circuits on an OSPF router function as a single adapter. In the single-adapter model, also known as the *NBMA model*, the network of the Frame Relay service provider (also called the Frame Relay cloud) is treated as an IP network, and the endpoints on the cloud are assigned IP addresses from a designated IP network ID. To ensure that OSPF traffic is received by all of the appropriate endpoints on the cloud, you must configure the Frame Relay interface to unicast its OSPF announcements to each appropriate endpoint. On a Routing and Remote Access OSPF router, you do this by designating the interface as an NBMA network and by configuring OSPF to unicast announcements to the specified neighbor routers. (Up to this point, the single-adapter model functions the same with both RIP and OSPF.)

  In a hub-and-spoke Frame Relay topology in an OSPF-based IP internetwork, the Frame Relay interface for the hub router must have a router priority set to **1** or greater and the Frame Relay interfaces for the spoke routers must have a router priority set to **0**. Otherwise, the hub router, which is the only router that can communicate with all of the spoke routers, cannot become the designated router and adjacencies cannot form across the Frame Relay network. For more information about the router priority setting, see “Using Designated Routers” later in this section.

- **Multiple-adapter model multicasts OSPF announcements.** Each virtual circuit functions as a separate adapter. In the multiple-adapter model, each Frame Relay virtual circuit functions as a point-to-point link with its own network ID, and the endpoints on the cloud are assigned IP addresses from a designated IP network ID. Because each virtual circuit is its own point-to-point connection, you can configure the interface as the point-to-point network type rather than as the NBMA network type.

On an OSPF router running Windows Server 2003 Routing and Remote Access, the network type (broadcast, point-to-point, or NBMA) is configurable in the Routing and Remote Access snap-in by navigating to the **IP Routing\OSPF** node, and then using the **General** tab of the ***InterfaceName*Properties** page. In addition, for the NBMA network type, neighbor routers on the NBMA network can be specified on the **NBMA Neighbors** tab of the same ***InterfaceName*Properties** page.

### Note

- The use of OSPF to provide remote access over nonpermanent, non-persistent, dial-up WAN links, such as analog phone lines or dial-up ISDN, is not recommended.

**OSPF Simple Password Authentication**

OSPF supports the use of simple password authentication to verify the origin of incoming OSPF announcements. Simple password authentication, which is enabled by default on Windows Server 2003 Routing and Remote Access OSPF interfaces, helps prevent an unauthorized or incorrectly configured OSPF router from corrupting OSPF messages on a subnet. By default, Windows Server 2003 Routing and Remote Access OSPF routers send the default plaintext password "12345678" in their OSPF Hello messages.

If you specify that plaintext passwords are required for an OSPF area:

- All interfaces on the same network segment must use an identical password.

- Interfaces in the same area that are on different networks can have different passwords.

**Note**

- Simple password authentication is a form of router identification. It is not a security option. Because the password is sent in plaintext, any user with a protocol analyzer (also known as a network sniffer), such as the Network Monitor tool, can capture the OSPF messages and view the authentication password.

**OSPF Convergence**

Of the features listed in this section that make it possible for OSPF to support large internetworks, fast convergence time is one of the most important. OSPF can detect and propagate topology changes faster than RIP. Unlike RIP, the count-to-infinity problem does not occur on an OSPF network and OSPF-calculated routes are always loop-free.

For more information about OSPF convergence, see "How OSPF Convergence Works," described next.

**How OSPF Convergence Works**

OSPF-enabled routers go through the following stages to achieve convergence of an IP internetwork:

1. Compiling the LSDB.
2. Calculating the Shortest Path First tree.
3. Creating routing table entries.

The following three sections describe this process in detail. The result of this process, sometimes referred to as fast convergence, is that OSFP can quickly recalculate routing information if routers fail and can avoid the occurrence of routing loops.

**Compiling the LSDB**

OSPF compiles the LSDB by means of an ongoing exchange of LSAs between neighboring routers, leading to a state in which each router is synchronized with its neighbors. After the autonomous system converges (if only one OSPF area exists) or after the local area converges (if multiple OSPF areas exist), each router has the appropriate entries in its LSDB.

To create the LSDB, each OSPF router must receive a valid LSA from each router in the autonomous system (if only one OSPF area exists) or from each router in the local area (if multiple OSPF areas exist). The process through which OSPF routers send and receive the LSAs is called flooding (sending traffic received on one interface out all other interfaces). Each router initially sends an LSA that contains its own configuration. As it receives LSAs from other routers, it propagates those LSAs to its neighbor routers. In this way, an LSA from a specific router is flooded across the autonomous system so that each other router learns that router's LSA. Although the flooding of LSAs might appear to cause a large amount of network traffic, OSPF is very efficient in the propagation of LSA information.

The following figure shows a simple, single-area OSPF autonomous system, the flooding of LSAs between neighboring routers, and the resulting LSDB compiled on one router. After the process completes, all routers contain an identical LSDB.

**Flooding LSAs to Create the LSDB in a Single-Area Autonomous System**




On a Windows Server 2003 Routing and Remote Access OSPF router, the OSPF LSDB is viewable in the Routing and Remote

Access snap-in by right-clicking the **OSPF** routing protocol in the console tree, and then clicking **Show Link State Database**.

To keep track of LSAs in the LSDB, each router is assigned a router ID, which identifies the router in the autonomous system. The router ID is a 32-bit dotted decimal number that is unique to the autonomous system. The router ID is not the IP address of one of the router's interfaces, and it is not used as a destination IP address for sending information to the router. Typically, the largest or smallest IP address assigned to the router is the router ID. Because IP addresses are unique, this convention ensures that OSPF router IDs are also unique.

**Calculating the Shortest Path First Tree**
After the LSDB is compiled, each OSPF router uses the Dijkstra algorithm to perform a least-cost path calculation on the information in the database. The Dijkstra algorithm, from a branch of mathematics called graph theory, is an efficient method of calculating a set of least-cost paths relative to a source node. The router, using itself as the root, performs this calculation to create a tree of shortest paths to each other router and subnet in the OSPF autonomous system (if only one OSPF area exists) or to each other router in the local area (if multiple OSPF areas exist). This tree, which contains a single least-cost path to each router and subnet, is known as the Shortest Path First (SPF) tree. A least-cost path is the path with the lowest accumulated cost. Because the least-cost path calculation is performed by each router with itself as the root of the tree, the SPF tree is different for each router.

OSPF uses the Dijkstra algorithm to calculate the SPF tree. The result of the Dijkstra algorithm is the set *SPF{}*, which is a cost-sorted list of least-cost paths that contain the path (the series of nodes and links) and its accumulated cost from the source node S.

**Creating Routing Table Entries**
OSPF creates entries in the IP routing table of each router from the SPF tree. If only one OSPF area exists, the routing table stores a single route for each subnet in the autonomous system. For multi-area OSPF internetworks, typically you configure route summarization at each ABR so that only summary routes that represent the destinations of each area are added as routes to the routing table. As explained earlier, the metric for the routing table entry is the OSPF-calculated cost, not a hop count (as is the case for RIP).

To calculate the IP routing table entries from the SPF tree, OSPF analyzes the set SPF{} that results from calculating the SPF tree. The analysis produces a series of OSPF routes that are added to the routing table. Each route contains:

- An IP destination (the network ID) and its associated network mask.
- The next-hop IP address of the appropriate neighboring router.
- The next-hop interface over which the neighboring router is reachable.
- The OSPF-calculated cost to the subnet.

**Example of OSPF Convergence**
The example described in this section illustrates how an OSPF internetwork compiles the LSDB, performs the least-cost analysis, and creates routing table entries. This example is deliberately simplified to help provide an understanding of the basic principles of OSPF convergence.

**Example: Compiling the LSDB in a single-area autonomous system**
At each router interface, the network administrator assigns a unitless cost metric (reflecting bandwidth, delay, or reliability), indicating a preference for using that interface, as shown in the following figure.

### Cost Assigned for Each OSPF Router Interface in an OSPF Area



After the LSA flooding process completes convergence, each router in the single-area autonomous system has received an LSA from all other routers, and each router in the system now contains an identical LSDB. The following table illustrates the common LSDB now shared by all routers in the area.

**After Convergence Occurs, All Routers Contain the Same LSDB Information**

| Router | Attached Networks and Costs |
|---|---|
| R1 | Net 1-Cost 2, Net 3-Cost 5, Net 4-Cost 2 |
| R2 | Net 1-Cost 1, Net 2-Cost 4, Net 6-Cost 2 |
| R3 | Net 2-Cost 4, Net 3-Cost 2, Net 5-Cost 3, Net 7-Cost 2 |
| R4 | Net 4-Cost 3, Net 5-Cost 2 |
| R5 | Net 6-Cost 2, Net 7-Cost 3 |

**Example: Calculating the SPF Tree**

In the next step, OSPF applies the Dijkstra algorithm to the example single-area OSPF autonomous system. The following figure illustrates the resulting SPF tree calculation, as performed by router R4. With R4 acting as the root, the SPF tree calculation determines the series of connected routers and networks that represent the least-cost path to each router and to each subnet.

**SPF Tree Calculation Performed by Router R4**




**Example: Creating Routing Table Entries**

Next, OSPF creates the routing table from the results of the SPF tree calculation, as illustrated in the following figure.

**OSPF Derives Routing Table Entries from the SPF Tree Calculation**




The exchange of LSAs by OSPF routers on a large IP internetwork might require a brief, but large, burst of bandwidth. After the exchange of advertisements and before convergence, a large amount of memory is required to process the LSDB. In addition, running the Dijkstra algorithm requires high CPU use. Therefore, networks with links that frequently appear and disappear can cause performance issues on routers because of the bandwidth needed to generate LSAs, to store the database, and to run the Dijkstra algorithm.

Instead of using one OSPF area for the entire autonomous system, you can minimize OSPF overhead and performance issues by dividing the OSPF autonomous system into multiple areas. For more information about OSPF areas, see "How OSPF Multiple-Area Autonomous Systems Work" later in this section.

**How OSPF Forms Adjacencies to Synchronize the LSDB**

Link state routing algorithms rely on the synchronization of LSDB information among routers in an autonomous system.

However, if every router exchanges routing information with every other router, the exchange of redundant information can occur, which increases network traffic and enlarges the size of the LSDB. To avoid this unnecessary overhead, each router is required to synchronize only with specific neighboring routers known as adjacent routers. When the LSDBs of two neighboring OSPF routers are synchronized, the routers are said to be adjacent. OSPF enables the exchange of routing information only between those neighboring OSPF routers that share an adjacency relationship. Not every pair of neighboring routers becomes adjacent.

OSPF requires that adjacencies be established in order to compile the correct entries in the LSDB, and only then can OSPF calculate the SPF tree and create entries in the routing table. Failure to establish adjacencies is one of the main problems that can prevent convergence in an OSPF internetwork.

**Establishing an OSPF Adjacency**

When an OSPF router initializes, it multicasts periodic OSPF Hello messages onto all directly connected subnets in order to discover neighboring OSPF routers and to communicate its own existence to them. The OSPF Hello message contains configuration information, such as the router's router ID and the list of neighboring routers from which the router has received a Hello message. Initially, however, the neighbor list in the OSPF Hello message does not contain any neighbors.

The initializing OSPF router, while multicasting its own OSPF Hello messages, also listens for Hello messages sent from neighboring routers. From incoming Hello messages, the initializing router determines the specific router or routers with which to form an adjacency. On a multiple-access network (broadcast or NBMA), adjacencies are formed with a designated router and a backup designated router (BDR) that are identified as such in the incoming Hello messages. The designated router forms adjacencies with other routers on its subnet, avoiding redundant adjacencies by establishing only the minimum required number of adjacencies. For more information about how designated routers avoid establishing redundant adjacencies, see "Using Designated Routers" later in this section.

To form an adjacency:

1. Two routers describe the contents of their LSDBs in an exchange of database description packets known as the database exchange process. It is during the database exchange process that the two neighboring routers form a master and subordinate relationship if they become adjacent router's. The content of each router's LSDB is acknowledged by its neighboring router.

2. Each router compares its LSAs with the LSAs of its neighbor and notes which LSAs it needs to request from the neighbor in order to synchronize the LSDB. The missing or more recent LSAs are then requested through Link State Request messages. Link State Update messages are sent in response to Link State Request messages and their receipt is acknowledged. When all Link State Request messages of both routers have received a response, the LSDBs of the router and its neighboring router are fully synchronized and an adjacency is formed.

3. After the adjacency has formed, each adjacent router sends a periodic Hello message to inform the adjacent router that the router is still active on the subnet. OSPF uses the lack of Hello messages from a neighbor to detect a downed router.

4. If an event, such as a downed link or router or the addition of a new subnet that changes the LSDB of one router occurs, the databases of adjacent routers are no longer synchronized. The changed router sends Link State Update messages to its adjacent neighbors, who acknowledge the messages. After this exchange, the LSDBs of the adjacent routers are again synchronized.

**OSPF neighbor states as an adjacency forms**

The neighboring OSPF routers go through a series of states (as defined in RFC 2328) during the establishment of an adjacency. The following table lists these states in the adjacency relationship in progressive order.

### States of OSPF Neighbors as an Adjacency Forms

| Neighbor State | Description |
|---|---|
| Down<br><br>(initial state) | Router A has received no OSPF Hello message from neighbor router B. However, on NBMA networks, router A can send a Hello message to router B even though router B is in a **Down** state. |

| | |
|---|---|
| Attempt<br><br>(for NBMA networks only) | Router A has received no OSPF Hello message from neighbor router B recently, despite attempts to contact router B by sending OSPF Hello messages to it. |
| Init | Router A has received an OSPF Hello message from neighbor router B, but the router ID of router A does not appear in the neighbor list of the Hello message sent by router B. This indicates that bidirectional communication has not yet been established between router A and router B. |
| 2-way | Router A has received an OSPF Hello message from neighbor router B, and the router ID of router A now appears in the neighbor list of the Hello message sent by router B. The same process occurs in reverse, establishing two-way communications between the two routers.<br><br>After the **2-way** state is achieved, router A determines whether to become adjacent with router B. On a broadcast network (the typical case) or on an NBMA network, an OSPF router remains in a **2-way** state with all other OSPF routers except the designated router and the BDR; it will establish a **Full** state (that is, a fully adjacent state) at a later stage only with the designated router and the BDR. |
| ExStart | Master and subordinate roles for the database exchange process are negotiated between router A and neighbor router B in this stage. This negotiation is the first phase of the adjacency relationship. The router with the higher router ID becomes the master and starts the exchange. |
| Exchange | Router A and neighbor router B exchange database description packets (also known as database descriptor packets). The contents of each database description packet is compared to the information stored in each router's LSDB to determine whether new or more current link state information is available in the database description packet sent by the other router. |
| Loading | Based on information in the exchanged database description packets, router A and neighbor router B send each other Link State Request messages asking for missing or more recent LSAs. Each responds by sending the requested link state information in Link State Update messages (which contain the entire LSA) to the other router. |
| Full | The LSDBs of router A and neighbor router B are fully synchronized, and the two routers are now said to be fully adjacent with each other. This adjacency relationship is now advertised in LSAs.<br><br>On a broadcast network (the typical case) or on an NBMA network, after completing this series of states to establish an adjacency relationship, a **Full** state occurs only between an OSPF router and its designated router and BDR; the **2-way** state is the normal state for most OSPF routers in a broadcast network.<br><br>On a point-to-point network, after completing this series of states to establish an adjacency relationship, the **Full** state is the normal state for the two OSPF routers on either side of the link. |

On a Windows Server 2003 Routing and Remote Access OSPF router, the neighbor state of neighboring routers is viewable in the Routing and Remote Access snap-in by right-clicking **OSPF** in the **IP Routing** container in the console tree, and then clicking **Show Neighbors**. The window that opens, **OSPF Neighbors**, displays all neighboring routers and each neighbor's current state. In a broadcast network:

- For neighbor routers where an adjacency is established (the designated router and BDR), **Full** appears in the **State** column.
- For neighbor routers where no adjacency is established, nor will be established because the router is neither a designated router nor a BDR, **2-way** appears in the **State** column.

For more information about designated routers and BDRs, see "Using Designated Routers" later in this section.

**OSPF adjacency configuration requirements**

A common OSPF problem is that a required adjacency between two neighboring routers does not occur. To successfully establish an adjacency, the following requirements must be met:

- The router IDs of the two neighboring routers *must not match*. Router IDs must be globally unique within an autonomous system. Duplicate router IDs prevent an adjacency. A router ID is different from an IP address.

- Settings on the two neighboring routers *must match* in all of the following parameters:

    - **Authentication type.** If authentication is used, the neighboring routers must use the same authentication type.

    - **Simple password authentication password.** If simple password authentication is enabled, the neighboring routers must use the same case-sensitive password.

    - **Hello interval.** The periodic interval at which Hello messages are sent must be the same on both routers. The default is 10 seconds.

    - **Dead interval.** The amount of time after which an adjacent router is considered down if the neighboring router does not receive a Hello message must be the same on both routers. In an Internet-Draft, titled "OSPF Benchmarking Terminology and Concepts," IETF recommends that the dead interval value be four times the Hello interval. The default is 40 seconds.

    - **Area ID.** The area ID, which is configured on the interface of each router, must be the same on both routers. This is a 32-bit dotted decimal number that identifies the area of the autonomous system to which the router is attached. An OSPF area ID is different from an IP address.

    - **Location in an OSPF stub area.** The two neighbor routers must agree whether or not they are in a stub area.

**Adding a Router to a Converged OSPF Internetwork**

When a new OSPF router initializes on an existing and converged OSPF internetwork, the LSA for the new OSPF router must be propagated to all existing OSPF routers on the internetwork through the process of flooding. After each existing router receives the new LSA, it performs the Dijkstra algorithm, recalculates the SPF tree, and creates new routing table entries. After all routing tables on all routers are updated, the internetwork has reconverged. For multi-area OSPF internetworks on which route summarization is used, the set of routes that summarize the networks of the new router must be configured on each ABR of the area in which the new router exists. This ensures that the networks attached to the new router are propagated across the entire OSPF autonomous system.

The following figure illustrates the convergence process for a new router, R1, and the new adjacency propagation process in an example single-area OSPF internetwork. In this example, the link between R1 and R2 is a dedicated WAN link, which is one type of point-to-point link.

**Adjacency Propagation of a New Router**




The preceding figure shows the convergence process for the new OSPF router, R1. This process works as follows:

1. After router R1 is turned on, it begins sending periodic Hello messages across the point-to-point WAN link to router R2. Router R2 also sends periodic Hello messages across the link to R1. R1 and R2 form an adjacency.

2. R1 and R2 exchange database description packets. R1's packet contains information only about itself. R2's packet contains the latest LSAs for all routers on the internetwork (except R1).

3. R1 sends a Link State Request message to R2, requesting the advertisements of all routers on the internetwork. R2 sends the requested advertisements to R1 as Link State Update messages.

4. R2 sends a Link State Request message to R1, requesting its advertisements. R1 sends its advertisements to R2 as a Link State Update message. R1 and R2 now have synchronized their respective LSDBs. When R1 and R2 receive the advertisements, they calculate their respective SPF trees and routing tables.

5. When R2 is synchronized with R1, R2 sends a Link State Update message to all other OSPF routers to which it is adjacent (routers R3 and R4). The update message contains the advertisements learned from R1. When R3 and R4 receive the advertisements from R2, they recalculate their respective SPF trees and routing tables.

6. R3 and R4 each floods the information in a separate Link State Update message to their respective adjacent routers (routers R5 and R6). When R5 and R6 receive the flooded advertisement about R1, they recalculate their respective SPF trees and routing tables.

The example single-area OSPF internetwork has reconverged after adding R1 and its associated subnet.

**Using Designated Routers**

On a point-to-point link, such as a dedicated WAN link, the OSPF adjacency relationship must be established between the two routers on either side of the link. However, on multiple access networks — broadcast or NBMA networks that contain multiple routers — the adjacencies must be controlled. On a broadcast network with *n* routers, a total of *n* (*n*–1)/2 adjacencies are formed if no control mechanism is present. For example, a broadcast network with six OSPF routers where each router establishes an adjacency with each other router would form a total of 15 adjacency relationships. In addition, unneeded flooding traffic would occur as each router attempts to synchronize with all of its adjacent routers.

**Election of designated routers**

To solve this scaling problem, every multiple access network (both broadcast and NBMA) elects a designated router, which forms adjacencies with all other routers on the subnet. On a multiple access network, the designated router establishes only the needed number of adjacencies (eliminating redundant adjacencies), which in turn reduces both network traffic and the size of the LSDB. On a broadcast network with *n* routers, a total of (*n* -1) adjacencies must be formed; any additional adjacencies are redundant. For example, a broadcast network with six OSPF routers will form a total of five adjacency relationships. Because the designated router is adjacent to all other routers, it acts as a hub for distributing link state information and for synchronizing LSDBs.

The designated router is elected through the exchange of OSPF Hello messages. Each Hello message specifies the current designated router (if elected), and the router ID and router priority of the sending router. The router priority is an interface-specific OSPF configuration parameter that is used to elect the designated router — the router with the highest router priority is elected. The default router priority is set to **1**. Setting a router priority of **0** prevents a router from becoming a designated router. If multiple routers have the same highest router priority, the router with the highest router ID is elected designated router. The recommended practice is to assign router priorities to ensure that the least busy routers are the designated router and BDR.

On a Windows Server 2003 Routing and Remote Access OSPF router, the router priority setting is configurable in the Routing and Remote Access snap-in by navigating to the **IP Routing\OSPF** node, and then specifying the desired priority for the **Router priority** option on the **General** tab of the ***InterfaceName*Properties** page.

If a designated router is already elected for a subnet, a new router on the subnet does not become the designated router even if it has a higher router priority than the current designated router.

### Note

- At least one router on the multiple access network (broadcast or NBMA) must be configured with a router priority of **1** or greater. If all routers on a multiple access network are configured with a router priority of **0**, no router becomes the designated router and no adjacencies are established. Without adjacencies, the LSDB cannot be synchronized and no transit traffic (traffic across that subnet) can be forwarded.

**Designated routers on broadcast networks**

The following figure illustrates why a broadcast network requires a designated router. Without a designated router, six separate adjacencies (illustrated in the top half of the figure) are formed on a broadcast network that contains four routers. The redundant adjacencies consume local and network resources in order to maintain the adjacencies and to propagate changes to the LSDB.

When the same network uses a designated router, only three adjacencies (illustrated in the bottom half of the figure) are needed to synchronize the LSDB.

**Designated Router Eliminates Redundant Adjacencies on a Broadcast Network**




**Designated routers on NBMA networks**

For NBMA networks, such as a Frame Relay network in a hub-and-spoke configuration, the designated router must be the hub router because only the hub router can communicate with all of the other routers. To ensure that the hub router is the designated router, set its router priority to **1** (or greater). To ensure that no spoke router becomes a designated router, set the router priority on spoke routers to **0**.

The following figure illustrates a designated router on a Frame Relay network that uses a hub-and-spoke configuration.

**Designated Routers on a Frame Relay Network**




**Backup designated routers**

The designated router forms adjacencies with other routers on its subnet and acts as a central distribution point for topological changes on a multiple access network. If the designated router becomes unavailable, adjacencies must be reestablished with a new designated router. Until the adjacencies re-form and the internetwork reconverges, a temporary loss of connectivity for transit traffic can result.

To prevent connectivity loss due to the loss of a designated router, a backup — the BDR — is also elected for each multiple access network. Like the designated router, the BDR is adjacent to all routers on the network. When the designated router fails, the backup immediately becomes the designated router by sending LSAs that announce its new role to all adjacent routers. For a brief time only, transit traffic might be impaired while the BDR takes over the role of the designated router.

Like the designated router, the BDR is elected by the exchange of Hello messages. Each Hello message contains a field for the BDR of the subnet. If a backup is not specified, the router with the highest router priority that is not already the designated router becomes the BDR. If multiple routers have the highest router priority, the router with the highest router ID is elected the BDR.

If a BDR is already elected for a subnet, a new router on the subnet does not become the BDR even if it has a higher router priority than the current BDR. After the initial election, the designated router and the BDR retain their status, even if a higher priority router joins the subnet. Only after both the designated router and the BDR fail does the current highest priority router become the new designated router, and the current router with the second highest router priority becomes the new BDR.

**Interface states for adjacent routers**

Each OSPF interface can be in one of several states after routers form an adjacency. The following table lists these interface states.

**Interface States for Adjacent Routers**

| Interface State | Description |
|---|---|
| Down | The initial interface state. No Hello messages have been sent or received. |
| Loopback | The interface to the network is looped back (internally configured so that outgoing packets are routed back to the source computer) through hardware or software. |
| Waiting | The interface is sending and receiving Hello messages to determine the designated router and the BDR for the subnet. |
| Point-to-point | The interface is adjacent to its neighbor on a point-to-point network or through a virtual link and is not the designated router or the BDR. |
| Other | The interface is on a multiple access network (broadcast or NBMA) and is not the designated router or the BDR. |
| Designated router | The interface is the designated router for the attached subnet. |
| Backup designated router | The interface is the BDR for the attached subnet. |

On a Windows Server 2003 Routing and Remote Access OSPF router, the interface state for an OSPF interface is viewable in the Routing and Remote Access snap-in by clicking the **IP Routing\OSPF** node. In the details pane, which displays the OSPF interfaces, the **State** column indicates the interface's current state. By viewing the interface state for each OSPF interface on a specific subnet, you can determine the designated router and the BDR for the subnet.

**How OSPF Multiple-Area Autonomous Systems Work**

In an OSPF autonomous system that includes a large set of contiguous networks in a single OSPF area, each OSPF router must maintain in its database the LSAs of every other router in that area. Therefore, each router in a large single-area OSPF autonomous system has a large LSDB. The SPF calculation for a large LSDB requires a substantial amount of processing. In addition, the resulting routing table is large because it contains a route for each subnet in the autonomous system. A large OSPF autonomous system is typically configured to contain more than one OSPF area because each area then has its own smaller LSDB. Using multiple areas in an OSPF autonomous system makes processing throughout the autonomous system more efficient.

OSPF is designed to support dividing the autonomous system into multiple OSPF areas in order to decrease the size of the LSDB and to reduce the processing overhead for calculating the SPF tree and the routing table. Dividing a large IP internetwork into several OSPF areas, each made up of groups of contiguous subnets, helps keep the size of the LSDB to a minimum. The size of the LSDB is more manageable and processing is faster because LSAs for networks and routers inside one area are flooded within that area but are not flooded to routers outside the area. Each area becomes its own link state domain with its own LSDB.

For a description of the possible types of OSPF areas and the types of OSPF routers needed inside and between OSPF areas, see "OSPF Database, Areas, and Routers" earlier in this section.

**Routing Between OSPF Areas**

Routing between OSPF areas in an autonomous system takes the following course:

1. A host in one area sends a packet destined for a host in another area to an internal router within its area. The sending host has no knowledge of either its own area or the area of the destination.

2. The internal router forwards the packet along the least-cost path to the nearest ABR, which (by definition) has an

interface that connects to the backbone area. The ABR is also a backbone router.

3. The ABR forwards the packet along the least-cost path to the nearest ABR that is connected to the area containing the IP address of the destination host.

4. The ABR connected to the area containing the destination host forwards the packet along the least-cost path, including through intermediate internal routers, if any, to the destination host.

The following figure shows the route of a packet as it is forwarded across R1, an internal router in area 0.0.0.1, to ABR R2, across the backbone area (area 0.0.0.0) to router R3, and then to the destination host in area 0.0.0.2.

**Inter-Area Routing in OSPF**




OSPF routers do not make routing decisions based on area IDs. All routing decisions are based on entries in the IP routing table. For example, in the inter-area routing shown in this figure, the backbone routers do not explicitly forward the packet to area 0.0.0.2. Rather, the routers forward the packet along the least-cost path to the router that is the best match for the destination IP address in the packet.

**Using Route Summarization Between OSPF Areas**

OSPF routers within an OSPF area perform routing *within* that area by using the least-cost path to the destination subnet. Each router in that area has a route to each subnet within the area. When multiple OSPF areas exist, OSPF can be configured to manage routing *between* areas by advertising networks located inside one area to other areas by using summary route advertisements in the format {*Destination, Network Mask*}. These destination/mask pairs, known as inter-area routes, summarize a range of addresses within this OSPF area to other areas in the same autonomous system. Using route summarization for inter-area routes reduces the number of entries in the routing table of OSPF routers within each area.

Using multiple OSPF areas in combination with route summarization makes very large OSPF internetworks feasible. When designing areas for an OSPF autonomous system, assign network IDs that can be expressed as a small number of summary routes. For example, assign network IDs for the private IP address space 10.17.0.0/16 and 10.29.0.0/16 to the subnets of an OSPF area so that all of the locations within the area can be summarized with two routes. If an area can be summarized with a single route, you can set the area ID of that area to the IP network ID that summarizes the entire area. If an area has multiple ABRs, ensure that they all are configured to summarize the same routes.

Configured ABRs on the OSPF backbone advertise summarized routes for all areas to which they are connected to all other routers on the backbone. Therefore, the routing table on each router in an OSPF area contains both the individual routes available within the local area and the routes corresponding to the summary advertisements received from configured ABRs connected to other areas in the autonomous system.

In the following figure, router R1 (the ABR on the border of area 0.0.0.1 and area 0.0.0.0) uses a summary advertisement to announce all routes (the list of address ranges) inside area 0.0.0.1 to all other backbone routers (including R2 in 0.0.0.2 and R3 in 0.0.0.3). In turn, R1 receives summary advertisements from R2 about area 0.0.0.2 and from R3 about area 0.0.0.3. In addition, R1 is configured with summary advertisement information for the backbone area, area 0.0.0.0. By flooding, R1 propagates incoming summary routing information to all routers within area 0.0.0.1. Each router within area 0.0.0.1 incorporates summary routing information from areas 0.0.0.0, 0.0.0.2, and 0.0.0.3 into its routing table.

**Summary Routes Reduce Routing Table Size**




Using route summarization hides the topology — the networks and their path costs — of one area from the rest of the autonomous system. The network topology for each area is discoverable only by routers within that area. Conversely, routers within each area cannot discover detailed topological information about other areas.

Hiding the topology of an OSPF area from other areas by using route summarization protects the rest of the autonomous system from revising its routing tables when events cause a network in another area to alternately fail and reconnect. Whenever a subnet reconnects, OSPF propagates the event as a Link State Update message and floods the update through adjacencies to routers within the area. This condition, known as route flapping, causes routing table instability because numerous routes are withdrawn and re-advertised in rapid succession. However, because all subnets within the area are advertised outside the area by using summary routes, OSPF does not flood the Link State Update outside the area and route flapping is avoided.

**Using Virtual Links to Areas Not Connected to the OSPF Backbone**

OSPF areas in an IP internetwork are connected to a backbone area that links the areas in the autonomous system together. Typically, an ABR that connects an OSPF area to the backbone area has a physical connection to the backbone area. However, when it is not possible or practical to use a physical connection to the backbone, you can configure a virtual link between an OSPF area and the backbone. A virtual linkis a logical point-to-point connection between an ABR of an OSPF area that is not physically connected to the backbone area and an ABR that is physically connected to the backbone.

You can configure a virtual link between any two routers that have an interface to a single common non-backbone area, known as the transit area. The transit area must have an ABR that is connected to the backbone. You cannot configure virtual links across multiple transit areas. OSPF forms a virtual adjacency across the virtual link to enable the exchange of routing information between the two routers connected by the virtual link. Such routers are known as virtual link neighbors.

In the following figure, area 0.0.0.3 does not have a router physically connected to the backbone (area 0.0.0.0). Therefore, a virtual link between routers R2 and R3 is configured across the transit area (area 0.0.0.2). R2 and R3 are virtual link neighbors.

**OSPF Virtual Link**




A virtual link requires values for the following settings on both virtual link neighbors:

- The area ID of the transit area for the virtual link.
- The OSPF router ID of the virtual link neighbor.
- The transit delay interval, which specifies the estimated number of seconds that it takes to transmit OSPF packets to the virtual link neighbor router.

- The retransmit interval, which specifies how long the OSPF router waits before retransmitting Link State messages. This number should be significantly larger than the expected round-trip delay between the two routers at either end of the transit area. Because this can be difficult to estimate for a virtual link, it is better to make this number too large rather than too small.

- Adjacency settings, including the Hello interval, dead interval, and plaintext password. Just as in physical adjacencies, the adjacency settings of the two virtual link neighbors must match before OSPF can successfully establish an adjacency.

On a Windows Server 2003 Routing and Remote Access OSPF router, a virtual link is configurable in the Routing and Remote Access snap-in by navigating to the **IP Routing\OSPF** node, and then using the **Virtual Interfaces** tab on the **OSPF Properties** page. Existing virtual links are viewable by right-clicking **OSPF** in the **IP Routing** container, and then clicking **Show Virtual Interfaces**. The **Virtual Interfaces** window displays all configured virtual interfaces and their states.

**Propagating External Source Routes to the OSPF Autonomous System**

An *external route* is defined as any route whose source is external to the OSPF autonomous system. By default, the set of OSPF routers that define an organization's OSPF autonomous system propagate only OSPF routes that correspond to directly connected network segments within the OSPF autonomous system. OSPF can also import external routes by defining an OSPF router as an autonomous system boundary router (ASBR). External routes can come from the following sources:

- RIP v2

- Local routes

- Static routes

- Static (non-demand-dial) routes

- Auto-static routes

- Routes that SNMP sets on the router

OSPF uses one or more ASBRs to learn and propagate external routes throughout an OSPF autonomous system. The ASBR uses a series of external route LSAs to advertise the availability of external routes, flooding the external route LSAs throughout the autonomous system (except in stub areas). The external route LSAs become part of the SPF tree and are included in the routing table calculation. Traffic to external networks is routed within the autonomous system by using the least-cost path to the ASBR.

The following figure shows an autonomous system that uses an ASBR to manage external route traffic.

**An ASBR Imports Routes from an External Source (RIP) into an OSPF Autonomous System**




**ASBRs and External Route Filters**

By default, an OSPF router configured to act as an ASBR imports and advertises all external routes from all external route sources. You can choose to block routes from one or more route sources, and you can filter out specific external routes. For

example, on a Windows Server 2003 Routing and Remote Access OSPF router configured as an ASBR, you can set either or both of the following types of restrictions:

- **Exclude specific external route sources.** You can configure the ASBR to accept or discard routes originating from one or more external sources, such as the RIP v2 routing protocol, static routes, or SNMP-generated routes. Accepting routes from all listed external route sources is the default option.

- **Exclude individual routes.** You can configure the ASBR to accept or discard specific routes for all external route sources or for a particular external route source by specifying one or more {*Destination, Network Mask*} pairs.

A combination of these filters, configured at the ASBR, can prevent the ASBR from advertising inappropriate routes and help ensure that the OSPF autonomous system is protected from incorrect or malicious routing information.

**ASBRs and Default Routes**

A router configured to act an ASBR advertises, by default, all external routes, including its own default route. This default route is used by all OSPF routers in the autonomous system. A valid default route has a next-hop IP address that points to a router that can reach the external network outside the OSPF autonomous system. An invalid default route has a next-hop IP address that points to a router within the OSPF autonomous system.

If the default gateway configured on the ASBR has a next-hop IP address that points to itself (its own IP address), the default route is not valid and packets forwarded by using that default route are dropped. To avoid this problem, use either of the following solutions:

- Do not configure an IP address for the default gateway on the TCP/IP properties of the ASBR.

- If the ASBR must have a default gateway, create an OSPF external route filter to filter out its own default route (destination 0.0.0.0, network mask 0.0.0.0). The default gateway configured on the ASBR must point to an IP router that is external to the OSPF autonomous system. Then, all routers within the autonomous system can correctly use that external router as their default gateway.

**ASBRs and Remote Access Servers**

A Windows Server 2003 Routing and Remote Access OSPF router in an OSPF environment might also be configured to function as a remote access server in order to support remote or mobile users. Typically, when you create a static IP address pool for remote access connections, address ranges in the address pool can represent separate subnets (known as off-subnet address ranges).

For the router to advertise correctly the routes for off-subnet address ranges, you must:

- Configure the router as an ASBR in the Routing and Remote Access snap-in by navigating to the **IP Routing\OSPF** node, and then selecting the **Enable autonomous system boundary router** option on the **General** tab of the **OSPFProperties** page.

- Use the **External Routing** tab on the **OSPF Properties** page to select **Route Filters**, specify **Accept listed routes**, and then add the routes that correspond to the off-subnet address ranges of the static IP address pool used for remote access connections.

### Note

- The use of OSPF to provide remote access over nonpermanent, non-persistent, dial-up WAN links, such as analog phone lines or dial-up ISDN, is not recommended.

**Using OSPF Stub Areas**

OSPF provides two types of area designed to improve performance in an OSPF autonomous system: standard stub areas and an extension to stub areas known as totally stub areas. The Windows Server 2003 Routing and Remote Access service supports both. OSPF stub areas are defined in RFC 2328.

**Stub areas summarize external routes**

Using an OSPF stub area decreases the amount of routing information that floods into it and therefore reduces routing overhead on the stub area's internal routers. OSPF does not flood routes external to the OSPF autonomous system into or through a stub area. To deliver packets from a stub area to any destination IP address outside of the OSPF autonomous

system, OSPF uses only a single route, the default route (destination 0.0.0.0, network mask 0.0.0.0). In OSPF, any destination not reachable through an intra-area or inter-area route is reachable through the default route; thus, OSPF uses the default route to route packets from a stub area to any destination not reachable within the autonomous system.

A stub area:

- **Reduces memory load.** Because a stub area does not advertise individual external networks, less memory is required on internal stub area routers than is required on internal routers in a standard non-stub OSPF area.
- **Reduces routing table size.** Using the default route for routing to external networks reduces the size of the routing table.

All routers in a stub area must be configured as stub routers. Configuring a router as a stub router automatically sets the E-bit flag in the OSPF Hello message on its interface to **0**. The **0** setting identifies this router as a stub router when it attempts to establish an adjacency with another router. Otherwise, by default, for an area that is not a stub area, OSPF sets the E-bit to **1**, which allows the router to accept and flood external routes into the OSPF area.

Stub areas have the following constraints:

- You cannot configure the backbone area as a stub area.
- You cannot configure virtual links that use a stub area as a transit area.
- You cannot place an ASBR inside a stub area.

**Example Stub Area**

Area 0.0.0.3 in the following figure is configured as a stub area. All traffic to external route locations must travel through its single ABR, R3. Router R3 advertises a default route for distribution inside area 0.0.0.3 instead of flooding external routes into the area.

**OSPF Stub Areas**




A stub area can contain a single ABR (a single entry and exit point) as in the example figure, or it can contain multiple ABRs. If the stub area contains multiple ABRs, each ABR must be configured to advertise a default route instead of external routes. When configured, any of the ABRs can be used to reach external route destinations even if the path to the external destination through the closest ABR is not the optimal path.

**Totally stub areas summarize external and inter-area routes**

A totally stub area uses the default route for route summarization of routes that are external to the totally stub area's OSPF autonomous system as well as for inter-area routes within its OSPF autonomous system. The following table compares stub areas and totally stub areas.

**Contrasting an OSPF Stub Area with a Totally Stub Area**

| Stub Area | Totally Stub Area |
|---|---|
| Does not advertise external routes. | Does not advertise external routes or OSPF inter-area routes. |

| Routing table contains: | Routing table contains: |
| --- | --- |
| • Intra-area routes<br>• Inter-area routes<br>• A default route (to reach destinations external to its autonomous system) | • Intra-area routes<br>• A default route (to reach destinations in other OSPF areas within its autonomous system and to reach destinations external to its autonomous system) |

One appropriate reason to deploy a totally stub area is to configure a branch office that does not have high traffic requirements to other branch offices. The branch office uses the default route to reach the central office and other branch offices.

On a Windows Server 2003 Routing and Remote Access OSPF router, a stub area or a totally stub area is configurable in the Routing and Remote Access snap-in by navigating to the **IP Routing\OSPF** node, selecting the **Add** option on the **Areas** tab of the **OSPFProperties** page, and then selecting **Stub area** on the **General** tab of the **OSPF Area Configuration** dialog box. A stub area is differentiated from a totally stub area by selecting the **Import summary advertisements** option for a stub area or ensuring that **Import summary advertisements** is not selected (which is the default) for a totally stub area.

**How the DHCP Relay Agent Works**

Any IP subnet that contains DHCP clients requires either a DHCP server or a DHCP relay agent to provide address leases to those DHCP clients. A computer acting as a DHCP relay agent, also known as a Bootstrap Protocol (BOOTP) relay agent, relays DHCP messages between DHCP clients on an IP subnet with no DHCP server and a DHCP server located on a different subnet.

Using one or more DHCP relay agents makes it unnecessary to install a separate DHCP server on each subnet in an IP internetwork because the DHCP relay agent enables a DHCP server to lease addresses to clients on a different subnet than the subnet on which the DHCP server is located. With DHCP relay agents, two DHCP servers (the second one is a backup) on one subnet can service an entire internetwork if the subnets with no DHCP server have an IP router configured to act as a DHCP relay agent.

The Windows Server 2003 Routing and Remote Access service provides a DHCP relay agent when the **DCHP Relay Agent** routing protocol component is added and configured with the IP address of one or more DHCP servers. The Routing and Remote Access DHCP relay agent is compliant with RFC 1542.

A DHCP relay agent mediates between a DHCP client and DHCP server in three steps:

1. A client on a subnet that has no DHCP server but that does contain a DHCP relay agent broadcasts a request for a DHCP lease.

2. The DHCP relay agent accepts the request and forwards it as a unicast message to the IP address of a DHCP server on another subnet.

3. The DHCP server sends a response (a lease offer) to the unicast IP address of the DHCP relay agent, which then forwards the response as either a broadcast or unicast transmission to the DHCP client.

**Note**

- Configuring the DHCP relay agent on a computer running any of the following will impair the proper functioning of DHCP:

- The DHCP Server service

- The Routing and Remote Access NAT routing protocol component with automatic addressing enabled

- Internet Connection Sharing (ICS)

A DHCP relay agent helps direct DHCP messages between a DHCP server and a DHCP client located on different subnets and takes an active role in recording information required for DHCP configuration.

**Initial DHCP Configuration**
Initial DHCP configuration is initiated by a DHCP client that has never leased an IP address, has released its IP address, or has received a DHCPNack message (indicating a negative response to a DHCP request) when attempting to lease a previous IP address. As described next, the initial DHCP configuration process uses four DHCP messages: DHCPDiscover, DHCPOffer, DHCPRequest, and DHCPAck.

**DHCPDiscover — DHCP client broadcasts initial request to discover a DHCP server**
When a Windows DHCP client computer starts, it broadcasts a message to obtain or renew a lease from a DHCP server. The DHCP client sends a DHCPDiscover message, which contains the MAC address of the DHCP client, to the limited broadcast IP address (255.255.255.255) and the MAC-level broadcast address. If no DHCP server exists on the client's subnet, the DHCP relay agent (if one is configured) receives the DHCPDiscover message and forwards it to a DHCP server on another subnet.

As defined in RFC 1542, the DHCP relay agent can forward the packet to an IP broadcast, multicast, or unicast address. In practice, DHCP relay agents forward a DHCPDiscover message to unicast IP addresses that correspond to configured DHCP servers. Before forwarding the original DHCPDiscover message, the DHCP relay agent makes the following changes:

1. Increments the Hop Count field in the DHCP header. The DHCP Hop Count field (which is different from the Time to Live [TTL] field in the IP header) is used to indicate how many DHCP relay agents have handled this DHCP message. When the configured maximum Hop Count is exceeded, the DHCPDiscover message is silently discarded. The default maximum hop count for the Windows Server 2003 DHCP relay agent is 4.

2. If needed, updates the Relay IP Address field (also known as the Gateway IP Address field) in the DHCP header. When the DHCP client sends the DHCPDiscover message, the Relay IP Address field is set to 0.0.0.0. If the relay IP address is 0.0.0.0, the DHCP relay agent records the IP address of the interface on which the DHCPDiscover message was received. If the relay IP address is not 0.0.0.0, the DHCP relay agent does not modify it. The Relay IP Address field records the first router interface encountered by the DHCPDiscover message.

3. Changes the source IP address of the DHCPDiscover message to the IP address of the interface on which the broadcasted DHCPDiscover message was received.

4. Changes the destination IP address of the DHCPDiscover message to the unicast address of the DHCP server configured on the DHCP relay agent.

The DHCP relay agent sends the DHCPDiscover message to the DHCP server (located on another subnet) as a unicast IP packet rather than as an IP and MAC-level broadcast. If the DHCP relay agent is configured with multiple DHCP servers, it sends each DHCP server a copy of the DHCPDiscover message.

**DHCPOffer — DHCP server offers IP address to the DHCP client**
When the DHCP server responds to the DHCP client request for an IP address, the DHCP server uses the Relay IP Address field in the following ways:

1. The relay IP address and the subnet masks of the server's configured scopes are compared by using a logical **AND** operation. When the comparison finds a scope whose network ID matches the network ID of the relay IP address, the DHCP server allocates an IP address from that scope.

2. When the DHCP server sends the lease offer back to the client, the server sends the DHCPOffer message to the source address of the DHCPDiscover message.

After the DHCP relay agent receives the relay IP address, it uses the address to determine which interface over which to send the DHCPOffer message. The DHCP relay agent then looks at the value of the Broadcast flag in the DHCPOffer message. The Broadcast flag is set by the original DHCP client to indicate, when set to **1**, that the client cannot receive unicast IP traffic until the DHCP lease process for the unicast IP address is concluded. Based on the value of the Broadcast flag, the DHCP Relay Agent either readdresses and sends the DCHPOffer message to the limited broadcast IP address and MAC-level broadcast address (if the Broadcast flag is set to **1**) or readdresses and sends the DHCPOffer message to the client using the offered IP address as the destination IP address and the client's MAC address as the destination MAC address (if the Broadcast flag is set to **0**). Windows Server 2003 and Windows XP DHCP clients set the Broadcast flag to **0**.

**DHCPRequest — Client requests to lease the IP address offered by the DHCP server**
As it does with the DHCPDiscover message, the DHCP client sends the DHCPRequest message, containing the MAC address of the client, to the limited IP broadcast address 255.255.255.255 and to the MAC-level broadcast address. The DHCP relay agent receives this packet and readdresses and sends it as a unicast IP packet to the configured DHCP server or servers.

**DHCPAck — DHCP server acknowledges the DHCP client request and approves its lease**
The DHCP server initially unicasts the DHCPAck message to the relay IP address, as it did with the DHCPOffer message. When the DHCP relay agent receives the DHCPAck message, it uses the Broadcast flag to forward it as a broadcast or unicast transmission to the DHCP client. The DHCPAck message is the signal from the DHCP server to the DHCP client that the server approves the lease that the client will now use, and the server updates its DHCP database.

When the client receives the acknowledgement, it configures its TCP/IP properties using any DHCP option information (such as default gateway and DNS server addresses) contained in the DHCPAck message and completes the initialization of TCP/IP.

**Restarting a DHCP Client and Obtaining an IP Address**
By default, when a Windows Server 2003, Windows XP, or Windows 2000 DHCP client shuts down, it does not send a DHCPRelease message to the DHCP server. Instead, when the DHCP client restarts, it attempts to obtain the IP address it last used by exchanging DHCPRequest and DHCPAck messages.

### Note

- For Windows Server 2003, Windows XP, and Windows 2000 DHCP clients, you can override this default by configuring the DHCP server to send the **Release DHCP lease on shutdown** Microsoft-vendor-specific DHCP option.

**DHCPRequest**
When a Windows Server 2003, Windows XP, or Windows 2000 DHCP client restarts, it attempts to lease its previously allocated IP address by using a broadcasted DHCPRequest message. The DHCPRequest, which is sent to the limited IP broadcast address 255.255.255.255 and to the MAC-level broadcast address, contains the MAC address and the previously allocated IP address of the client. The DHCP relay agent receives this packet and treats the message in a way similar to the way that it treats a DHCPDiscover message (described earlier). Before forwarding the packet, the DHCP relay agent does the following:

1. Increments the Hop Count field in the DHCP header.

2. In the Relay IP Address field, records the IP address of the interface on which the DHCPRequest message was received.

3. Changes the source IP address to the IP address of the interface on which the broadcasted DHCPDiscover message was received.

4. Changes the destination IP address to the unicast address of the DHCP server that is recorded in the DHCPRequest, and then readdresses and sends the packet as a unicast IP packet to the DHCP server.

**DHCPAck and DHCPNack**
When the DHCP server receives the DHCPRequest, the server compares the network ID of the client's allocated IP address to the network ID of the relay IP address. Based on the result of that comparison, the server does one of the following:

- If the two network IDs are the same and if the IP address is available and thus can be reallocated to the DHCP client, the DHCP server sends a DHCPAck to the IP address of the DHCP relay agent. When the DHCP relay agent receives the DHCPAck, it readdresses and sends the message to the DHCP client.

- If the two network IDs are the same and if the IP address is not available and thus cannot be reallocated to the DHCP client, the DHCP server sends a DHCPNack to the DHCP relay agent. When the DHCP relay agent receives the DHCPNack, it readdresses and sends the message to the DHCP client. The DHCP client then restarts the IP address allocation process with a DHCPDiscover.

- If the two network IDs are not the same, the DHCP client has moved to a different subnet, and the DHCP server sends a DHCPNack to the DHCP relay agent. When the DHCP relay agent receives the DHCPNack, it readdresses and sends the message to the DHCP client. The DHCP client then restarts the IP address allocation process by sending a DHCPDiscover message.

**How IP Packet Filtering Works**
Windows Server 2003 supports several types of IP packet filtering.

**Routing and Remote Access IP Packet Filters**
Windows Server 2003 supports the following types of IP packet filtering to protect interfaces of a Routing and Remote Access

router:

- Basic Firewall, provided by the NAT/Basic Firewall component of the Routing and Remote Access service. Basic Firewall is intended for use on Internet-connected interfaces (public interfaces). For more information about Basic Firewall packet filtering, see "Basic Firewall" later in this section.

- IP packet filtering, enabled on the **General** tab of the **Properties** page of any IP interface. Although input and output filters configured in this way are typically referred to as "IP packet filters," all of the other types of packet filters listed in this section introduction are also IP packet filters. For more information about IP packet filtering provided by Routing and Remote Access, see "IP Packet Filtering" later in this section.

**Other Windows IP Packet Filters**

In addition to the IP packet filter capabilities provided by the Routing and Remote Access service, Windows Server 2003 supports the following types of IP packet filtering:

- Internet Connection Firewall (ICF) — a Network Connections component available in Windows XP, Windows XP SP1, and Windows Server 2003 — is a simple stateful firewall similar to Basic Firewall.

- TCP/IP packet filtering, configured by using the **Options** tab of the **Advanced TCP/IP Settings** page of TCP/IP properties, is used only for inbound local host traffic.

- Internet Protocol security (IPSec) packet filtering, provided as part of IPSec, is included with Windows Server 2003. IPSec can perform host-based packet filtering to provide limited firewall capabilities that help defend end systems against network-based attacks. An end system is a network device that cannot forward packets between portions of a network. For more information about IPSec packet filtering, see "IPSec Technical Reference."

**Basic Firewall**

Basic Firewall, provided by the NAT/Basic Firewall component of the Routing and Remote Access service, is a simple stateful firewall that can be enabled to protect a public interface on a computer running the Windows Server 2003 Routing and Remote Access service. A stateful firewall uses dynamic packet filtering to examine incoming and outgoing packets. Basic Firewall, which combines stateful packet filtering of network traffic with a set of static packet filters, examines network traffic in order to prevent unsolicited traffic originating on the Internet from being processed by an Internet-connected interface. Basic Firewall uses technology similar to that used by ICF.

The network address translation capability provided by the Windows Server 2003 Routing and Remote Access service can be configured with or without the Basic Firewall feature enabled. Conversely, you can enable Basic Firewall on any Internet-connected interface even when you have not deployed NAT. For a Basic Firewall with NAT, all packets that are received on the Internet interface of a NAT router that do not correspond to traffic requested by the NAT router (either for itself or for private intranet clients) are discarded. For a Basic Firewall without NAT, all packets that are received on the Internet interface of a Routing and Remote Access server that does not perform network address translation that do not correspond to traffic requested by this server are discarded.

Basic Firewall in conjunction with Routing and Remote Access NAT is primarily for use by small office or home office (SOHO) customers. For more information about using NAT with Basic Firewall, including information about how Basic Firewall filters TCP and UDP unicast packets, broadcast and multicast packets, specially defined packets, and ICMP traffic, see "How NAT Works" in NAT Technical Reference.

You can also use the NAT/Basic Firewall component of Routing and Remote Access in conjunction with the IP packet filtering capability provided by IPSec. Alternatively, you can use IPSec with ICF, which also provides stateful packet filtering.

**IP Packet Filtering**

To increase security on an IP network, an IP router configured with packet filters can allow or prevent the flow of specific types of IP traffic. IP packet filtering, by providing a way for the network administrator to define precisely which IP traffic is received and sent by the router, is an important component of connecting corporate intranets to public networks like the Internet.

When packet filtering is configured, a packet arriving at a router might be forwarded to its destination, rejected with an error message returned to notify the sender, or dropped with no error message sent.

**Filtering Inbound and Outbound Traffic**

IP packet filtering consists of creating a series of definitions, called filters, which define for the router what types of traffic are allowed or blocked on each interface. Filters can be set for incoming and outgoing traffic. Input filters define what inbound traffic the router is allowed to process locally (as the destination) or to forward (as a router) on a specific interface. Output

filters define what outbound traffic the router is allowed to send from that interface.

To filter incoming and outgoing traffic to allow a specific subset of traffic, you can implement packet filtering on a router or on a non-router computer. A computer running Routing and Remote Access that has only one interface is a non-router computer.

**Planning Carefully When Configuring Packet Filters**
Because you can configure both input and output filters for each interface, it is possible to create contradictory filters. For example, the input filter on one interface allows certain inbound traffic, but the output filter on another interface does not allow the same traffic to be sent. The result is that the traffic is not passed across the router or exchanged with a non-router computer on which these contradictory packet filters are configured.

When planning how to implement packet filters, ensure that the filters are not too restrictive. Excessively restrictive filters can impair the functionality of other protocols that might be operating on the computer. For example, if a computer that is running Windows Server 2003 is also running Internet Information Services (IIS) as a Web server and if packet filters are defined so that only Web-based traffic is allowed, you cannot use the **ping** tool to perform basic IP troubleshooting. Likewise, the Web server cannot communicate with a domain controller. If the Web server is a silent RIP host, the filters prevent the silent RIP process from receiving the RIP announcements.

Therefore, when troubleshooting connectivity or IP-based network problems on a computer running Windows Server 2003, begin by determining whether that computer is using IP packet filtering. If it is, verify whether packet filtering on that computer prevents outgoing or incoming packets for the protocol experiencing the problem.

**Windows Server 2003 IP Packet Filtering**
You can configure a computer running the Windows Server 2003 Routing and Remote Access service either to pass all traffic except that which is blocked by filters, or to discard all traffic except that which is allowed by filters. For example, you might want to configure a filter to allow all traffic except Telnet traffic (TCP port 23), or you might want to set up filters on a dedicated Web server to process only Web-based TCP traffic (TCP port 80).

**Routing and Remote Access Input and Output Filters**
The following table shows that input and output filters work in a similar manner, differing only in that they apply either to inbound or to outbound traffic.

### Similar Functionality of Input and Output Filters

| Function | Input Filters | Output Filters |
|---|---|---|
| Filtering by exception | Configured on an exception basis. You can configure the filter action to either allow all traffic except that which is specified, or to block all traffic except that which is specified. | Configured on an exception basis. You can configure the filter action to either transmit all traffic except that which is specified, or to drop all traffic except that which is specified. |
| Using multiple filters | When multiple filters are configured, the separate filters applied to the inbound packet are compared by using a logical OR operation. If the packet matches at least one of the configured filters, it is allowed or blocked depending on the filter action setting. | When multiple filters are configured, the separate filters applied to the outbound packet are compared by using a logical OR operation. If the packet matches at least one of the configured filters, it is transmitted or dropped depending on the filter action setting. |
| Troubleshooting | You can quickly and temporarily disable input packet filtering while preserving the list of configured packet filters. Do this by adding a filter that allows all traffic and set the input filter action to block all traffic except that which is specified. After troubleshooting, remove the all-traffic filter and reset the input filter action to its original value. | You can quickly and temporarily disable output packet filtering while preserving the list of configured packet filters. Do this by adding a filter that allows all traffic and set the output filter action to block all traffic except that which is specified. After troubleshooting, remove the all-traffic filter and reset the output filter action to its original value. |

**Filterable Fields in IP, TCP, UDP, and ICMP Headers**
Windows Server 2003 Routing and Remote Access can filter incoming and outgoing IP packets based on specified fields of the IP, TCP, UDP, and ICMP headers.

The Routing and Remote Access service does not support either of the following:

- Creating a filter on any protocols other than IP, TCP, UDP, and ICMP.
- Creating a filter based on *all* of the fields of the IP, TCP, UDP, or ICMP header. The Routing and Remote Access service does support creating a filter on specific fields of each of these headers.

For more information about the IP, TCP, UDP, and ICMP protocols, including the complete set of fields for each protocol's header, see *Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference* by Joseph Davies and Thomas Lee, 2003, Microsoft Press.

**IP Header**

The IP header has 13 fields, including source and destination IP addresses and a field that identifies the upper layer protocol (such as TCP, UDP, or ICMP) contained within the IP packet payload. Because the IP header contains the packet's source and destination addresses, the TCP, UDP, and ICMP headers do not need to contain these addresses.

To add or edit an IP packet filter, in the Routing and Remote Access snap-in select **Inbound Filters** or **Outbound Filters** on the **General** tab of the ***InterfaceName*Properties** page in the **IP Routing\General** container. You can define filters for the following IP header fields:

- **Source IP Address.** The IP address of the source network, which you can configure with a subnet mask to allow or block a single source or an entire range of IP addresses (corresponding to an IP subnet or a summarized route) to be specified with a single filter entry.

- **Destination IP Address.** The IP address of the destination network, which you can configure with a subnet mask to allow or block a single destination or an entire range of IP addresses (corresponding to an IP subnet or a summarized route) to be specified with a single filter entry.

- **Protocol.** An identifier used to indicate the upper layer protocol contained within the IP packet payload. For example, 6 indicates the TCP protocol, 17 indicates the UDP protocol, and 1 indicates the ICMP protocol.

**TCP Header**

TCP is one of two transport layer protocols within the TCP/IP protocol suite used for transporting data (the other is UDP). TCP, which supports only one-to-one communication between two hosts, guarantees delivery of data, and guarantees that packets are delivered in the same order in which they were sent. The TCP header has 11 fields, including fields for source and destination port.

If you select TCP as the protocol when adding or editing an IP packet filter, you can define filters for the following two TCP header fields:

- Source Port

- Destination Port

### Note

- Routing and Remote Access does not support the configuration of a range of TCP ports in one step. To configure a range of TCP ports, you must configure a separate filter for each port in the range.

**UDP Header**

UDP is one of two transport layer protocols within the TCP/IP protocol suite used for transporting data (the other is TCP). UDP is simpler than TCP, supports one-to-one and one-to-many communications, and, unlike TCP, sends packets without first negotiating a connection between the sending and receiving hosts, and provides only a few error recovery services. UDP thus supports a "best-effort" delivery in contrast to TCP's guaranteed delivery. The UDP header has 4 fields, including fields for source and destination port.

If you select UDP as the protocol when adding or editing an IP packet filter, you can define filters for the following two UDP fields:

- Source Port

- Destination Port

### Note

- Routing and Remote Access does not support the configuration of a range of UDP ports in one step. To configure a range of UDP ports, you must configure a separate filter for each port in the range.

**ICMP Header**

ICMP enables the generation of delivery error messages, diagnostic messages, and informational messages related to IP (such as the ICMP redirect message used to dynamically update the IP routing table).

If you select ICMP as the protocol when adding or editing an IP packet filter, you can define filters for the following two ICMP header fields:

- **Type.** Indicates the type of ICMP packet, such as an Echo or Echo Reply message.

- **Code.** Indicates one of the possible multiple functions within a specified type. If a type has only one function, the **Code** field is set to **0**.

The combination of ICMP Type and ICMP Code determines which ICMP message is sent. The following table lists commonly used ICMP types and codes.

**Common ICMP Types and Codes and Associated Messages**

| ICMP Type | ICMP Code | Use |
|---|---|---|
| 0 | 0 | Echo reply |
| 8 | 0 | Echo |
| 3 | 0 | Destination unreachable - Network unreachable |
| 3 | 1 | Destination unreachable - Host unreachable |
| 3 | 2 | Destination unreachable - Protocol unreachable |
| 3 | 3 | Destination unreachable - Port unreachable |
| 3 | 4 | Destination unreachable - Fragmentation needed and the **Don't Fragment** flag is set |
| 4 | 0 | Source quench |
| 5 | 1 | Redirect - Redirected datagrams for the host |
| 9 | 0 | Router Advertisement |
| 10 | 0 | Router Solicitation |
| 11 | 0 | Time exceeded - TTL expired in transit |
| 11 | 1 | Time exceeded - Fragmentation reassembly expiration |
| 12 | 0 | Parameter problem |

**Note**

- For a complete list of ICMP message types, see "ICMP TYPE NUMBERS" at the Internet Assigned Numbers Authority Web site.

**Routing and Remote Access Filters for IP Header Fields**

The following table shows the IP header fields for which it is possible to create a filter and the corresponding fields of the dialog boxes in the Routing and Remote Access snap-in in which the filter is created. These settings are configurable in the Windows Server 2003 Routing and Remote Access snap-in by selecting **Inbound Filters** or **Outbound Filters** on the **General** tab of the ***InterfaceName*Properties** page for an interface in the **IP Routing\General** container, and then selecting **New** or **Edit** to access the **Add IP Filter** dialog box or the **Edit IP Filter** dialog box, respectively. Typically, filters are configured in pairs, either {*Destination Network/Protocol*}, or {*Source Network/Protocol*}.

**IP Header Fields and Corresponding Routing and Remote Access Name**

| IP Header Field | Name in Routing and Remote Access Add IP Filter or Edit IP Filter Dialog Box |
|---|---|
| Source IP Address | Source Network |
| Destination IP Address | Destination Network |
| Protocol | Protocol |

**Filtering Scenarios**

This section illustrates filter configurations for commonly implemented filtering scenarios.

**Note**

- If you combine any sets of filters, ensure that the filters do not cancel each other out. Note that AND is used within a filter; OR is used between filters.

**Local Host Filtering**

Local host filtering configured on the Internet interface of an IP host permits the interface to receive only traffic that a non-router host can receive. This includes traffic to the unicast address of the host, broadcast addresses, and multicast addresses. This type of filtering allows the receipt of packets but disables the forwarding of packets on the interface on which local host filtering is enabled. Use local host filtering when an intranet is connected to the Internet and you do not want direct routing of packets between the intranet and the Internet.

Local host filtering on an interface effectively disables unicast routing on that interface because the only unicast traffic allowed through the interface is that which is destined specifically for that host. Transit unicast traffic is dropped. Local host filtering permits the interface to receive only the unicast traffic destined for this host, the broadcast traffic intended for all host's on the host's subnet, or multicast traffic.

As shown in the example in the following table, for the local host filtering scenario, five specific input filters are used on the Internet interface. In the example, the local host computer has the IP address 172.16.5.98 with the subnet mask of 255.255.255.0 (a subnet of the private IP network ID 172.16.0.0/12). The Internet interface is viewable in the Routing and Remote Access snap-in in the details pane of the **IP Routing\General** container. By adding the filters shown in the table on the **General** tab of the **Properties** page of the Internet interface, and then selecting the **Drop all packets except those that meet the criteria below** filter action, you can configure the computer to block all packets except those explicitly allowed by the configured filters. For the local host filters, the default protocol **Any** is used and no source network address is needed.

**Example: Five Input Filters for Local Host Filtering**

| Filter | Filter Characteristics |
|---|---|
| Host IP address input filter (allows packets unicast to this computer) | Destination network:<br>• IP address of this host: 172.16.5.98<br>• Subnet mask: 255.255.255.255<br><br>**Note:** An IP address for a particular computer in conjunction with a subnet mask of 255.255.255.255 represents a host route (a route to a specific IP address). |

| | |
|---|---|
| Subnet broadcast address input filter (allows packets broadcast to hosts on the local subnet) | Destination network:<br>• IP address of this host's subnet broadcast: 172.16.5.255<br>• Subnet mask: 255.255.255.255<br><br>**Note:** On a classless network (as is the case in this example), the IP subnet broadcast address is formed by setting all host bits to 1 (255 in decimal notation represents eight ones in binary notation).<br><br>Only a classless network has a subnet broadcast address; it does not have a network broadcast address. (In contrast, a classful network does not have a subnet broadcast address; it has only a network broadcast address.) |
| All-subnets-directed broadcast address input filter (allows packets broadcast to hosts on all subnets of a subnetted class-based network ID) | Destination network:<br>• IP address of this host's all-subnets-directed broadcast: 172.16.255.255<br>• Subnet mask: 255.255.255.255<br><br>**Note:** On a classless network (as is the case in this example), the all-subnets-directed broadcast address is formed by assuming that the network ID is the original classful network ID, and then setting all host bits to **1** (255 in decimal notation). If 172.16.5.98 were a class-based address, it would be a Class B address. Therefore, in the IP address, the last two octets are set to 255 for the all-subnets-directed broadcast address filter. |
| Limited broadcast address input filter (allows packets sent to the limited broadcast address) | Destination network:<br>• IP address of the limited broadcast: 255.255.255.255<br>• Subnet mask: 255.255.255.255<br><br>**Note:** The IP limited broadcast address is formed by setting all 32 bits in the IP address to 1 (255.255.255.255). All hosts, classful or classless, process packets addressed to the limited broadcast address, typically during an automated configuration process, such as the DHCP automated configuration process. |
| All possible multicast traffic input filter (allows all IP multicast packets) | Destination network:<br>• IP address for multicast traffic: 224.0.0.0<br>• Subnet mask: 240.0.0.0<br><br>**Note:** The address 224.0.0.0 with the subnet mask 240.0.0.0 is the entry in the routing table that informs IP to forward the multicast packet to the destination (in this example, to 172.16.5.98). |

**Web Traffic Filtering**

Web traffic filtering is configured on Web servers that run the Routing and Remote Access service in addition to Microsoft Internet Information Services (IIS). Web traffic filtering ensures that only Web-based traffic to and from IIS is processed to protect against malicious attacks on other services running on the Web server.

For a Web server connected to the Internet, enabling Web traffic filtering requires adding one input filter and one output filter on the Internet interface. The Internet interface is viewable in the Routing and Remote Access snap-in in the details pane of the **IP Routing\General** container. For example, if the default port of your Web server is port 80, you can add filters on the **General** tab of the **Properties** page of the Internet interface that block all packets (by using the **Drop all packets except those that meet the criteria below** filter action) except those allowed to and from the Web server by the filters described in the following table.

**Example: Input and Output Filters for Web Traffic Filtering**

| Filter | Filter Characteristics |
|---|---|
| Input filter (allows packets sent specifically to this Web server) | Destination network:<br>• IP address: *IP address of this Web server*<br>• Subnet mask: 255.255.255.255<br><br>Protocol: TCP<br>• Source port: *None*<br>• Destination port: 80 |
| Output filter (allows packets sent from this Web server) | Source network:<br>• IP address: *IP address of this Web server*<br>• Subnet mask: 255.255.255.255<br><br>Protocol: TCP<br>• Source port: 80<br>• Destination port: *None* |

If these filters are the only filters configured, the only traffic allowed through the interface is TCP traffic to and from IIS through TCP port 80 on the Windows Server 2003–based computer.

**FTP Traffic Filtering**

File Transfer Protocol (FTP) traffic filtering is done on hosts that run IIS with the FTP service enabled and run the Routing and Remote Access service. FTP traffic filtering ensures that only FTP-based traffic to and from FTP is processed to protect against malicious attacks on other services running on the FTP server.

For an FTP server connected to the Internet, enabling FTP traffic filtering requires adding two input filters and two output filters on the Internet interface. The Internet interface is viewable in the Routing and Remote Access snap-in in the details pane of the **IP Routing\General** container. For example, you can add filters on the **General** tab of the **Properties** page of the Internet interface that block all packets (by using the **Drop all packets except those that meet the criteria below** filter action) except those allowed to and from the FTP server by the filters described in the following table.

**Note**

- The example assumes that you are using the default ports of the FTP server, 20 and 21. If you use ports other than 20 and 21, substitute the appropriate ports for ports 20 and 21 in these filters.

**Example: Input and Output Filters for FTP Traffic Filtering**

| Filter | Filter Characteristics |
|---|---|

| | |
|---|---|
| Input filter (allows packets sent specifically to the FTP control port on this FTP server) | Destination network:<br>• IP address: *IP address of this FTP server*<br>• Subnet mask: 255.255.255.255<br><br>Protocol: TCP<br>• Source port: *None*<br>• Destination port: 21 (the FTP control port) |
| Input filter (allows packets sent specifically to the FTP data port on this FTP server) | Destination network:<br>• IP address: *IP address of this FTP server*<br>• Subnet mask: 255.255.255.255<br><br>Protocol: TCP<br>• Source port: *None*<br>• Destination port: 20 (the FTP data port) |
| Output filter (allows packets sent from the FTP control port on this FTP server) | Source network:<br>• IP address:. *IP address of this FTP server*<br>• Subnet mask: 255.255.255.255<br><br>Protocol: TCP<br>• Source port: 21 (the FTP control port)<br>• Destination port: *None* |
| Output filter (allows packets sent from the FTP data port on this FTP server) | Source network:<br>• IP address: *IP address of this FTP server*<br>• Subnet mask: 255.255.255.255<br><br>Protocol: TCP<br>• Source port: 20 (the FTP data port)<br>• Destination port: *None* |

If these filters are the only filters configured, the only traffic allowed through the interface is TCP traffic to and from FTP on the Windows Server 2003–based computer.

**PPTP Traffic Filtering**

Point-to-Point Tunneling Protocol (PPTP) traffic filtering on hosts that are acting as PPTP VPN servers ensures that only PPTP-based traffic to and from the PPTP VPN service on the host is processed. PPTP traffic filtering helps secure the PPTP VPN server

from malicious attacks on other services running on the PPTP VPN server. For a PPTP VPN server connected to the Internet, PPTP traffic filtering is configured on the Internet interface.

For information about IP packet filtering for PPTP traffic, see "How VPN Works" in VPN Technical Reference.

**L2TP/IPSec Traffic Filtering**
Layer Two Tunneling Protocol over IPSec (L2TP/IPSec) traffic filtering on hosts that are acting as L2TP/IPSec VPN servers ensures that only L2TP-based traffic to and from the L2TP/IPSec VPN service on the host is processed. L2TP traffic filtering helps secure the L2TP/IPSec VPN server from malicious attacks on other services running on the L2TP/IPSec VPN server. For an L2TP/IPSec VPN server connected to the Internet, L2TP traffic filtering is configured on the Internet interface.

For information about IP packet filtering for L2TP/IPSec traffic, see "How VPN Works" in VPN Technical Reference.

**Denying Spoofed Packets from Private IP Addresses**
One way to create denial-of-service attacks is to flood servers with packets, such as TCP connection request packets, from addresses to which there can be no reply. In these cases, the malicious users spoof, or substitute, the source IP address of the packets with something other than the IP address of the interface on which the packets originated. An easy address to spoof is a private address because a response sent to a private address on the Internet results in an ICMP Destination Unreachable-Host Unreachable message.

To drop Internet traffic from spoofed private IP addresses, you can use the Windows Server 2003 Routing and Remote Access snap-in to add three input filters on the Internet interface. The Internet interface is viewable in the Routing and Remote Access snap-in in the details pane of the **IP Routing\General** container. You can add filters on the **General** tab of the **Properties** page of the Internet interface that accept all packets (by using the **Receive all packets except those that meet the criteria below** filter action) except the following:

- A source IP address of 10.0.0.0 with a subnet mask 255.0.0.0.
- A source IP address of 172.16.0.0 with a subnet mask 255.240.0.0.
- A source IP address of 192.168.0.0 with a subnet mask 255.255.0.0.

**Filtering Fragmented IP Datagrams**
The Windows Server 2003 Routing and Remote Access service also supports the filtering of fragmented IP datagrams, which are datagrams that contain a fragment of an IP payload. Source hosts or routers fragment IP payloads so that the resulting IP datagram is small enough to be sent on the subnet of the next hop. The filtering of fragmented IP datagrams by the Routing and Remote Access service applies only to incoming traffic.

All interfaces are viewable in the Routing and Remote Access snap-in in the details pane of the **IP Routing\General** container. You can enable filtering of fragmented IP datagrams on the **General** tab on the properties of each interface by selecting **Enable fragmentation checking**. This setting specifies that the router discard all fragmented IP packets that do not correspond with traffic allowed by the filters (if any) configured for this interface. To prevent the router from forwarding all fragmented IP packets for transit traffic on any interface, enable fragmentation checking on all interfaces of the router. This setting does not prevent the forwarding of fragmented packets sent from the router.

You can use filtering of fragmented IP datagrams to prevent the "ping of death," which is a denial-of-service attack where malicious users send multiple 64-KB ICMP Echo messages. The 64-KB messages are fragmented and must be reassembled at the destination host. For each separate 64-KB message, the TCP/IP protocol must allocate memory, tables, timers, and other resources. If a server, including a Windows Server 2003–based computer, receives a substantial number of fragmented messages, it can become overloaded, with the result that the servicing of valid information requests is impaired. Enabling the filtering of fragmented IP datagrams causes incoming fragmented IP datagrams to be discarded immediately.

**How ICMP Router Discovery Works**
Windows Server 2003 TCP/IP and the Windows Server 2003 Routing and Remote Access service include support for the router advertisement and discovery capability that is documented in RFC 1256. ICMP itself is described in RFC 792. Enabling router discovery, an alternative to using manually or DHCP-configured default gateways, provides an improved method of configuring and detecting the default gateway. For IP nodes, the default gateway is the IP router (on the same subnet) that forwards traffic for the node to destinations not located on the local subnet.

### Note

- In IPv6, the Neighbor Discovery protocol replaces the IPv4 features ICMP Router Discovery, ARP, and ICMP Redirect. For information about IPv6, see "How IPv6 Works" in IPv6 Technical Reference.

ICMP Router Solicitation messages and Router Advertisement messages simplify how IP hosts are configured with the IP addresses of local routers and provide a way for hosts to determine whether routers are available. If enabled on both router and host, ICMP router discovery automates the discovery and configuration of the default gateway on the client. Client computers can dynamically discover the best default gateway to use on a subnet and can automatically switch to another default gateway if the first default gateway fails or if the network administrator changes router preferences.

ICMP router discovery governs the interaction between a host and a router in the following way:

- A router periodically sends a Router Advertisement message to all hosts on the subnet to notify them that the router is available. A router also sends a Router Advertisement message in response to a Router Solicitation message sent by a host.

- A host listens for Router Advertisement messages from routers. A host also sends a Router Solicitation message to discover active routers on its subnet and to discover whether known routers are currently unavailable.

On a computer running Windows Server 2003, ICMP Router Advertisement messages are disabled by default. TCP/IP for Windows Server 2003, Windows 2000, and Windows 98 supports ICMP Router Solicitation messages to discover the default gateway as well as the receipt of ICMP Router Advertisements; this capability is also disabled by default.

**Router Advertisements Sent by Routers**
Router Advertisements are explicit notifications sent by IP routers to IP hosts on the subnet that the router is available. A router sends out a pseudo-periodic Router Advertisement by using an ICMP message (Type 9, Code 0). The Router Advertisement message can be sent to the all-systems IP multicast address of 224.0.0.1, to the local IP broadcast address, or to the limited broadcast address (255.255.255.255). By default, a router sends a Router Advertisement to the all-systems IP multicast address.

**Router Solicitations Sent by Hosts**
When an IP host that supports ICMP router discovery needs to be configured with a default gateway (either upon initialization or because its default gateway is unavailable), the host sends out a Router Solicitation message by using an ICMP message (Type 10, Code 0). The Router Solicitation message can be sent to the all-routers IP multicast address of 224.0.0.2, to the local IP broadcast address, or to the limited broadcast address (255.255.255.255). By default, hosts send Router Solicitation messages to the all-routers IP multicast address when an interface initializes. Routers on the host's subnet that support ICMP router discovery immediately respond with a Router Advertisement message, and the host chooses the router for which an administrator has configured the highest preference level as its default gateway. A host sends a maximum of three solicitations at intervals of approximately 600 milliseconds.

### Note

- ICMP router discovery is not a routing protocol. Routers advertise only their existence; they do not advertise the best way to reach a specific destination subnet. If a host uses a non-optimal route, an ICMP Redirect message redirects the host to the preferred route.

**ICMP Router Discovery Settings**
Settings for ICMP router discovery for a router are configurable in the Windows Server 2003 Routing and Remote Access snap-in on the **General** tab of the properties for an interface in the **IP Routing\General** container. If you select the **Enable router discovery advertisements** option on the **General** tab, you can configure the time interval after which a router is considered down, the level of preference for this router to act as a default gateway for hosts, and the maximum and minimum rate at which the router periodically sends ICMP router advertisements.

On a client computer running Windows XP Professional, you can use the **netsh routing ip routerdiscoveryadd interface** command to configure router discovery for the specified interface from a command prompt. On a client computer running Windows 2000 Professional, you can enable router discovery by using the Registry Editor to configure the **PerformRouterDiscovery** and **SolicitationAddressBcast** registry entries.

**IPv4 and Routing and Remote Access APIs**
Microsoft platform software development kits (SDKs) provide application programming interfaces (APIs) useful for third-party developers who create software that interacts with Windows operating systems. The following are important APIs for developers working in a unicast IPv4 routing environment that includes Windows Server 2003.

**Internet Protocol Helper (IP Helper) API**
This API assists network administration of the local computer by enabling applications to retrieve information about network configuration settings of the local computer and to modify that configuration. In addition, IP Helper notification mechanisms ensure that an application is informed when specific aspects of the local computer network configuration change. The IP Helper

API is applicable in any computing environment where programmatically manipulating TCP/IP configuration is useful. Typical usage includes retrieving information about network configuration, managing network adapters, managing interfaces, managing IP addresses, using the address resolution protocol (ARP), retrieving information about IP and ICMP, managing routing, receiving notification of network events, and retrieving information about TCP and UDP. In Windows Server 2003, IP Helper has been extended to allow the retrieval of information for IPv6 and its components. For more information about the IP Helper API, see "Microsoft Platform SDK IP Helper."

**Routing and Remote Access Service API**
This API makes the Routing and Remote Access service extensible. Third-party routing protocol developers can write additional routing protocols that interact with the routing table manager of the Routing and Remote Access service. Third-party software vendors can also create applications to administer the routing and remote access capabilities of Windows Server 2003 (as well as Windows 2000 Server). For example, to augment the built-in support for RIP and OSPF, a third-party vendor can develop support for Border Gateway Protocol (BGP) for IP that works with the Routing and Remote Access service. For more information, see "Microsoft Platform SDK Routing and Remote Access Service."

**Windows Sockets 2 (Winsock2) API**
This API enables programmers to create advanced Internet, intranet, and other network-capable applications to transmit application data across the wire, independent of the network protocol being used. In Windows Server 2003, new Windows Sockets programming elements extend the capability of Winsock to simplify programming and to provide IPv6 compatibility. For more information about the Windows Sockets API, see "Microsoft Platform SDK Windows Sockets 2."

For comprehensive information about APIs supported by Microsoft Windows, including Windows Server 2003; Windows XP; Windows Advanced Server Limited Edition; Windows Millennium Edition; Windows 2000; Windows NT versions 4.0 and 3.51; Windows 98 and Windows 95; as well as Microsoft Internet Explorer, see "Microsoft Platform SDK Windows API."

To install the SDK for Windows Server 2003, see "Microsoft Platform SDK," and then click "Windows SDK" in the left column.

**Related Information**
The following resources contain additional information that is relevant to this section.

- For more information about each of the RFCs cited in this document, see IETF RFC Database.
- TCP/IP Technical Reference.
- Routing and Remote Access.
- IPv4 Multicasting Technical Reference.
- *Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference*, by Joseph Davies and Thomas Lee, 2003, Redmond, WA: Microsoft Press.
- *Routing in the Internet, Second Edition*, by Christian Huitema, 2000, Englewood Cliffs, NJ: Prentice Hall.
- *Interconnections: Bridges and Routers, Second Edition*, by Radia Perlman, 2000, Reading, MA: Addison-Wesley.
- *Internetworking with TCP/IP: Volume 1 Principles, Protocols, and Architectures, Fourth Edition*, by Douglas E. Comer, 2000, Englewood Cliffs, New Jersey: Prentice Hall.
- *OSPF: Anatomy of an Internet Routing Protocol*, by John T. Moy, 1998, Reading, MA: Addison-Wesley.

# How Unicast IPv4 Routing Works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Unicast IPv4 Routing Internetwork Structure
- How Host Routing and Router Routing Work
- How IPv4 Addressing and IPv4 Routing Interact
- How the IPv4 Routing Table Works
- How Static and Dynamic Routing Work
- Related Information


Internet Protocol version 4 (IPv4) is a routable network protocol that enables the routing of network traffic across any type of IPv4 internetwork, including Windows or UNIX internetworks or mixed network environments. A part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite, IPv4 is used by Windows Server 2003 to route unicast, multicast, and broadcast traffic between computers or other network devices on private intranets and across the public Internet. IPv4 is frequently referred to as IP, and that convention is used in this document.

| **Note** |
| --- |
| Windows Server 2003 also supports Internet Protocol version 6 (IPv6). For information about IPv6, including static routing in an IPv6 environment, see "How IPv6 Works" in the IPv6 Technical Reference. |

Unicast IP routing is the process of transferring data in the form of unicast IP packets across an IP internetwork from a single source node (a computer or other network device) to a single destination node. An IP router is a node that forwards IP packets between interfaces (which can be physical network adapters or logical interfaces such as tunnels) based on information that it stores in a database called a routing table. On IP internetworks, routers configured with dynamic routing protocols can exchange information with each other automatically in order to keep routing table information up-to-date. Dynamic routing contrasts with static routing, in which an administrator must manually configure routing information.

Familiarity with the protocol-independent unicast IPv4 routing principles described here provides a foundation for the in-depth discussion of dynamic routing protocols and other routing-related services provided by the Windows Server 2003 Routing and Remote Access service. For more information about Routing and Remote Access components used for unicast IPv4 routing, see How Unicast IPv4 Routing Protocols and Services Work.

**Unicast IPv4 Routing Internetwork Structure**

In unicast routing, data is transferred across an internetwork from a source node to a destination node. Routers, which connect different networks on the internetwork, accept a packet sent from one network to another and forward the packet toward its destination. This traffic requires the following types of protocols:

- **Routable protocol.** A routable protocol, such as IP or Internetwork Packet Exchange (IPX), is a network protocol that a router uses to forward packets toward their destination across the internetwork. IP, which uses network layer addresses to forward packets, is the routable protocol for an IP internetwork.

- **Dynamic routing protocol.** A dynamic routing protocol is used by routers to exchange routing information with each other dynamically. The exchange process allows routers to determine the most efficient route over which to forward a packet. The most common dynamic routing protocols used by IP are Routing Information Protocol for IP (RIP for IP) and Open Shortest Path First (OSPF).

In addition to routable and routing protocols, unicast IP routing requires the routing components and structures described in this section:

- Unicast IPv4 routing internetwork components
- IPv4 routing infrastructure

**Unicast IPv4 Routing Internetwork Components**

The following figure depicts a simple IP internetwork in which unicast IP routing takes place.




The following table describes each of the routing components, as defined by the International Organization for Standardization (ISO), depicted in the figure.

**Components of a Simple IP Internetwork**

| Internetwork Component | Description |
|---|---|
| End system | A network device that cannot forward IP packets. End systems are also known as *hosts*. A computer that is not acting as a router is an example of an end system. (Collectively, hosts and routers are known as *nodes*.) |
| Intermediate system | A network device that can forward IP packets that are not addressed to it. Bridges, switches, and routers are examples of intermediate systems. |
| Router | An intermediate system used to connect network segments (subnets) based on a common network layer protocol, such as IP. In the preceding figure, the intermediate system depicted is a router. Routers include:<br>• **Hardware routers.** A device that performs routing as a dedicated function and that is specifically designed and optimized for routing. For example, hardware routers made by Cisco Systems, Inc., are widely deployed in IP internetworks.<br>• **Software routers.** A computer that performs routing as one of multiple processes running on the computer. For example, a computer running Windows Server 2003 with the Routing and Remote Access service enabled is a software router.<br>A router can be configured to perform only static routing, only dynamic routing, or, more typically, primarily dynamic routing supplemented by some manually configured static routes. |
| Network | A portion of the networking infrastructure (encompassing repeaters, hubs, bridges, or Layer 2 switches) that is bounded by an intermediate system and whose computers and other network devices share a common network address called a network ID. Also referred to as a *network segment* or a *subnet*. |
| Internetwork | Two or more networks (subnets) connected by routers. In the preceding figure, all of the depicted components combined constitute an internetwork. |

When a node on an IP network sends a unicast IP packet to a node on another network, it typically first forwards the packet to a router. The router examines the packet and uses the destination internetwork address (defined in the following table) in the packet header to determine which of the router's interfaces to use to forward the packet towards its destination. The following table lists the generic internetworking term and the corresponding IP term used for each part of an internetwork address.

**Parts of IPv4 Unicast Internetwork Addresses**

| Generic Internetwork Address Term | IP Internetwork Address Term | Description |
|---|---|---|

| | | |
|---|---|---|
| Internetwork address<br><br>(Network address + host address) | IP address<br><br>(Network ID + host ID) | In IPv4, a 32-bit logical address for a node on an IPv4 network. By combining the network address portion and the host address portion, a unicast IPv4 address uniquely identifies a node on an internetwork.<br><br>Hosts and routers use IPv4 addresses to route packets over an IPv4 network. Unicast IPv4 addresses enable a packet to be sent from a single source node to a single destination node. |
| Network address | Network ID | In IPv4, the portion of an IP address that identifies a single network segment (also known as a subnet) bounded by routers on an IP internetwork. All systems on the same subnet share the same network ID. |
| Host address | Host ID or node ID | In IPv4, the portion of an IP address that uniquely identifies an individual node on an IP subnet. |

The network layer header of any IPv4 packet sent from a source node to a destination node on an IPv4 internetwork includes the following:

- A source IPv4 address, which contains a source network ID and a source host ID.
- A destination IPv4 address, which contains a destination network ID and a destination host ID.
- A Time-To-Live (TTL) value, which is used to prevent the packet from endlessly circulating on the internetwork. The TTL starts at a maximum value and decreases by one for each link crossed; if zero is reached, the packet is discarded and a message is returned to the sending node from which the packet originated.

| **Note** |
|---|
| The TTL link count (the number of Network Interface Layer links crossed to reach a given destination) is not the same as the hop count (the number of routers crossed to reach a specific destination). The link count, which equals the number of network segments, is the hop count plus one. For example, if two nodes are separated by five routers, the hop count is 5 but the TTL link count is 6. |

The following resources provide additional information about the structure of unicast IP routing components on an IPv4 internetwork:

- For an architectural diagram showing where IP fits within the TCP/IP protocol suite, see "TCP/IP Protocol Architecture" in "How TCP/IP Works" in TCP/IP Technical Reference.
- For information about the architecture of unicast IPv4 routing specific to the Windows Server 2003 Routing and Remote Access service, see "Unicast IP Routing Architecture in Windows Server 2003" in How Unicast IPv4 Routing Protocols and Services Work.
- For information about the architecture of the NAT unicast routing protocol component of the Windows Server 2003 Routing and Remote Access service, see "How NAT Works" in NAT Technical Reference.

**IPv4 Routing Infrastructure**

The routing infrastructure refers to the entire structure of a routed internetwork over which routing takes place. A routed internetwork is designed to route packets between its individual networks either over one path or over one of multiple paths; its structure is either flat or built into a network, subnet, and sub-subnet hierarchy; and it either groups networks under one administrative authority (autonomous system) that shares a common routing protocol, or it is divided into multiple areas within one autonomous system.

This section briefly summarizes the following background information for understanding unicast IPv4 routing on the IPv4 internetwork:

- Single-path vs. multipath routing infrastructures
- Flat vs. hierarchical routing infrastructures
- Autonomous systems

A detailed discussion of these structural issues is not included in this technical reference. However, familiarity with these topics is important in understanding how unicast IP routing is implemented in a Windows internetwork and in considering which dynamic routing protocols and other routing-related services to deploy on your IP internetwork.

**Single-Path vs. Multipath Routing Infrastructures**

A single-path routing infrastructure is one in which a unicast IP packet can take only a single route between any source node and any destination node. Although single-path routing simplifies routing tables and packet flow paths, single-path internetworks are not fault tolerant. On an internetwork in which all routing information is manually configured, a router cannot detect a failure, and if a failure occurs, networks beyond the failure are unreachable until the failure is resolved. By contrast, a router running a dynamic routing protocol can detect a failure, but, because only a single path exists, networks beyond the failure are also unreachable for a dynamic router for the duration of the failure. A downed link or router must be brought back up before packets can again be delivered successfully to all locations on the internetwork. Single-path routing is thus not feasible for a large internetwork that depends on consistent network availability.

By contrast, a multipath routing infrastructure is one in which a unicast IP packet can take one of several paths between any source node and any destination node. A multipath IP internetwork that uses dynamic routers is fault tolerant because more than one path exists, which means that routers can not only detect a failure but can also route packets to networks beyond the downed link or router by using an alternative path.

Typically, multipath internetworks provide more reliable network availability than single-path internetworks but are more complex to configure. In addition, although multipath networks that use distance vector–based routing protocols provide better availability than single-path networks, they can experience routing loops if configuration errors occur or during convergence (the process by which routers update routing tables after a change in network topology and reestablish a stable state). A multipath network that uses a link state routing protocol, such as OSPF, provides more efficient convergence and thus better stability than a distance vector routing protocol, such as RIP, but requires more effort to plan and deploy.

For more information about static and dynamic routing, see How Static and Dynamic Routing Work later in this document.

**Flat vs. Hierarchical Routing Infrastructures**

An internetwork that uses a flat routing infrastructure consists of multiple network segments connected to each other without a parent-child structure. All routers are peers, and each network segment is stored as a separate route in the routing table. The network IDs have no network or subnet structure and routes cannot be summarized, which means it is not possible to group a set of network segments as a single entry in the routing table. Therefore, although flat routing can work well for small- or medium-sized IP internetworks, it is not well-suited to large internetworks because it requires that its routers store an ever-increasing number of routes in their routing tables.

By contrast, an IP internetwork that uses a hierarchical routing infrastructure consists of multiple network segments sharing the same network ID prefix grouped into a hierarchy that consists of a network, subnet, and sub-subnet structure. Each group of network segments in the hierarchy can be represented as a single network ID through route summarization and can therefore be represented in the IP routing table as a single route. The routing table entry for the highest level (the network) is also the route used for regions of the network and for individual subnets.

In a hierarchical routing infrastructure, a collection of contiguous networks connected by routers that share routing information is referred to as a routing region (also known as a *routing domain* or *routing area*). One part of the internetwork that is divided into routing regions can be connected by a common routing region, called a *backbone*, to another part of the internetwork that is also divided into routing regions. Routers within each routing region perform intradomain routing. Routers connected to the backbone perform interdomain routing. Only a relatively few routes are required on the backbone of a well-planned hierarchical internetwork.

Some IP internetworks use both flat and hierarchical routing. For example, the Internet, originally designed as a flat routing infrastructure until the Internet community helped pioneer the introduction of classless addressing and summarized routing, now incorporates both flat and hierarchical routing. One of the factors that results in the storing of large numbers of routes in the routing tables of Internet backbone routers is the continued inclusion of flat routing within the Internet. These large Internet backbone routing tables impede the speed at which Internet traffic can be forwarded and is one of several factors driving the development of IPv6. The hierarchical routing infrastructure and new address types provided by IPv6 are designed to require relatively few routing entries in the routing tables of Internet backbone routers.

For more information about route summarization on a hierarchical IPv4 internetwork, see "CIDR Blocks and Hierarchical Routing Enabled by Route Summarization" in How IP Addressing and IP Routing Interact later in this document.

**Autonomous Systems**

Very large IP internetworks are divided into separate entities called autonomous systems. An autonomous system (AS) is a portion of the internetwork — either a single network or a collection of networks — under a common administrative authority,

such as an enterprise, a business division, or other organization, that shares a common routing protocol. The autonomous system is described in RFC 1930, "Guidelines for creation, selection, and registration of an Autonomous System (AS)" in the IETF RFC Database.

An autonomous system is sometimes said to be defined by the use of the shared routing protocol that it uses. For example, a contiguous portion of an IP internetwork that uses OSPF to exchange routing information among a group of routers is under OSPF administrative authority and is, therefore, an OSPF autonomous system.

In a large organization, one autonomous system might be subdivided into multiple routing regions (routing domains or areas) that define a hierarchy within that autonomous system. Regions within an autonomous system can use summarized routes for traffic between the regions to decrease the number of routes exchanged and stored in routing tables.

**Types of autonomous system**

The three types of autonomous system are:

- **Stub AS.** An autonomous system that is connected to only one other autonomous system. Any data sent to, or received from, a destination outside the autonomous system must travel over that connection. A network that is connected to only one other autonomous system has the same identifying number (autonomous system number, described in the next subsection) as the autonomous system to which it is connected. A small campus network is an example of a stub autonomous system.

- **Transit AS.** An autonomous system that is connected to two or more other autonomous systems and that allows traffic between other autonomous systems (that is not destined for a node within that autonomous system) to cross through it. The network of an ISP is an example of a transit autonomous system.

- **Multihomed AS.** An autonomous system that is connected to two or more other autonomous systems but does not allow transit traffic between other autonomous systems to cross through it. A multihomed autonomous system differs from a stub autonomous system in that the entrance and exit points for data traveling to or from the autonomous system can vary depending on which connection offers the shortest route to the destination. An example of a multihomed autonomous system is a corporate network that uses two or more Internet connections to different ISPs.

The terms stub AS and transit AS are not to be confused with the OSPF stub area or OSPF transit area (described in "How OSPF Works" in How Unicast IPv4 Routing Protocols and Services Work) because the former refer to types of autonomous system whereas the latter refer to types areas within an OSPF autonomous system.

Many enterprise IP networks and most ISPs consist of one autonomous system. Only very large internetworks, such as the Internet, are divided into multiple autonomous systems.

**Autonomous system numbers**

In most cases, an autonomous system is assigned a globally unique 16-bit number called an autonomous system number (ASN). Autonomous systems use ASNs to identify themselves to each other so that they can exchange routing information. ASNs, which are required for autonomous systems connected to the Internet, are obtained from Internet numbers registries, such as:

- **American Registry for Internet Numbers (ARIN)**, for the Americas, Caribbean, and sub-Saharan Africa regions.

- **Asian-Pacific Network Information Center (APNIC)**, for the Asia Pacific region.

- **Réseaux IP Européens (RIPE)**, for the Europe, Middle East, and northern Africa regions.

In nearly all cases, unless an organization is multihomed to more than one ISP or its routing policy is different from that of its ISP, it does not need its own ASN because it can use the ISP's ASN.

For information about routing protocols used by autonomous systems, see "Routing Protocols Used Between and Within Autonomous Systems" in the section How Static and Dynamic Routing Work later in this document.

**How Host Routing and Router Routing Work**

A route is the path that a packet travels across an internetwork from a sending node to a destination node. Routing, which takes place on a packet-by-packet basis, is the process of determining the best route to use to forward a packet as it crosses the internetwork. On an IP internetwork, the routing of a packet across the IP internetwork is based on the packet's destination IP address.

Routing can be either host routing or router routing:

- **Host routing.** The sending host forwards a packet. Based on the destination network ID, the sending host determines whether to forward the packet to a destination host on its subnet or forward the packet to a router for further action.

- **Router routing.** A router receives a packet (from a sending host or another router) that is not addressed to the router and that must be forwarded from one network to another, which it then sends it on its way. The packet is forwarded between routers (when the destination network is not directly attached to the router) or between the router and the destination host (when the destination network is directly attached to the router).

In the following figure, the source host forwards to Router 1 a packet destined for the destination host on a different subnet. In this example, the initial step in the routing process — from the source host to Router 1 — is host routing. Router 1 then forwards the packet to Router 2, which in turn forwards the packet to the destination host. In this example, this second step in the routing process — from Router 1 to Router 2 to the destination host — is router routing.




The following subsections describe how host routing and router routing work on an IP internetwork.

**Host Routing**

When an IP host sends data to another host on an IP internetwork, the sending host might first obtain the IP address of the destination host by referencing the logical name of the destination host. Hosts on an IP internetwork typically use Domain Name System (DNS) name resolution to resolve a DNS domain name to an IP address.

In host routing, as depicted in the following figure, two types of delivery are possible: direct and indirect delivery.




- **Direct delivery in host routing.** After obtaining the destination IP address, the source host determines whether the destination is on the locally attached network segment. When the source and destination hosts are on the same network segment, the source host sends a packet directly to the destination host and does not use a router (labeled "Direct Delivery" in the above figure). In this case, the destination host is called a *local host* or *neighboring host*. The source host uses Address Resolution Protocol (ARP) to obtain the destination link-layer address (typically, the Ethernet media access control, or MAC, address) to which it sends the packet. In a direct delivery, the destination IP address and the destination link-layer address refer to the same node.

- **Indirect delivery in host routing.** When the destination host is not on the locally attached network segment, the source host cannot deliver a packet directly to the destination. In this case, the destination host is called a *remote host*. Therefore, the source delivers the packet to an intermediate router by using ARP to obtain the MAC address of the near-side interface of the router, and then sends the packet to that interface, as shown in the preceding figure. In an indirect delivery, the destination IP address and the destination MAC address refer to different nodes.

During an indirect delivery, the source host forwards the packet to a router on its network segment by determining the router corresponding to the first hop. If multiple routers exist in the path between the source host and the destination host on the internetwork, as the packet is forwarded from router to router, the packet's network layer destination address (IP address) remains the same, but its MAC address keeps changing to that of the next router along the path.

**Host Determination of the First Hop**

An IP source host determines the next-hop IP address of a neighboring router by using one of the following methods:

**IP address currently in the host routing table**

A routing table on the host provides the IP address of a router that can forward the packet towards its destination. For more information about routing tables, see How the IPv4 Routing Table Works later in this document.

### ICMP dynamic updates of the host routing table

If a sending host forwards a packet to a neighboring router but another neighboring router is closer to the destination, the router that initially receives the packet can send an ICMP Redirect message to the sending host, informing it of the better router to use as the next-hop for the destination. When the sending host receives the Redirect message, it adds a host route for the destination to the local routing table. TCP/IP for Windows Server 2003 supports the dynamic update of the IP routing table based on the receipt of ICMP Redirect messages.

### Default route

A default route, which summarizes destinations that are not located on the local subnet, is the route on which IP packets are sent when no other route to the destination network is found in the routing table. To simplify the configuration of IP nodes and to reduce the overhead that exists if each host stored routes for all network IDs on the internetwork, a source host is configured with a single default route. In IPv4, the default route is {Network ID: 0.0.0.0, Subnet Mask: 0.0.0.0}, also expressed as 0.0.0.0/0. For more information about the default route, see How the IPv4 Routing Table Works later in this document.

### Eavesdropping

IP hosts can listen to routing protocol traffic if the capability known as *eavesdropping* or *wiretapping* is enabled. Eavesdropping hosts maintain the same routing information as do routers. An example of eavesdropping is silent RIP, which is the ability of an IP host to listen to RIP routing traffic exchanged by RIP routers and to update the routing table on the host. A computer that uses silent RIP processes RIP announcements but does not announce its own routes. You can use silent RIP on a non-router host (such as a workstation or a server running Routing and Remote Access that has only one interface and is therefore not a router) to produce a routing table that contains as much detail as the routing table on a RIP router. By including detailed routes in its routing table, a silent RIP host can perform routing more efficiently.

For example, you can configure computers running Windows XP as silent RIP hosts in an environment that includes RIP routers by installing the RIP Listener networking component. Instead of a user or administrator manually adding routes to the routing table, the computer learns other routes on the network by listening to RIP messages and then adds advertised routes in those messages to its routing table. You do not need to configure a silent RIP host with a default gateway (IP address of a directly reachable IP router).

The eavesdropping provided by silent RIP can be useful in some remote access environments, such as when a remote access computer connects to a corporate network over a dial-up connection. The additional routes added by RIP Listener can improve network access and performance over the dial-up connection by reducing the number of packets that are sent to the incorrect router. Windows Server 2003, Microsoft® Windows® XP, Windows® 2000, Microsoft Windows NT® Workstation version 4.0 Service Pack 4 (SP4) and later, and Windows NT Server version 3.51 Service Pack 2 (SP2) and later support Silent RIP.

#### Host Determination of the Entire Path: Source Routing

Some routable protocols, including IP, can enable the source host to do more than simply determine the first hop. In this scenario, IP enables the source host to determine the path between the source host and the destination, and then includes the list of routers in the network layer header of the packet. The network layer header is used by each router to forward the packet along the indicated path. This process is known generically as *source routing*.

Source routing is *not* typically implemented on IP internetworks because in source routing the path must be known. Typically, instead of using source routing, IP routing decisions are made by source hosts and IP routers based on their respective local routing tables. However, in network testing and debugging situations, you might sometimes want to specify an exact route that overrides the route as determined by routing tables as the packet travels towards its destination. Specifying an exact route through the IP internetwork is known as *IP source routing*.

In IP source routing, the source host specifies the entire route through successive IP routers between the source and destination. Each IP router addresses the IP packet to the next router by using the **Destination IP Address** field in the IP header.

IP supports two types of source routing:

- **Loose source routing.** The IP address of the next router can be one or more routers away (multiple hops).
- **Strict source routing.** The next router must be a neighboring router (single hop).

**Note**

Token Ring source routing is a MAC–sublayer routing scheme; it is not related to the internetwork-based source routing discussed here.

**Router Routing**

When a router receives a packet, the router must either deliver the packet to the destination host or to another router. As in host routing, router routing has two types of delivery, direct and indirect, as depicted in the following figure.




- **Direct delivery in router routing**. If a router is directly attached to the destination network, it forwards the packet to the local destination host by addressing the packet to the link-layer address of the destination host. The router thus performs a direct delivery to the destination.

- **Indirect delivery in router routing**. If a router is not directly attached to the network destination, it forwards the packet to an intermediate router. The first router chooses the intermediate router based on the next-hop address of the closest matching route in the routing table. The router forwards the packet by addressing the packet to the link-layer address of the intermediate router. The router thus performs an indirect delivery to the next router in the path to the destination.

**How IPv4 Addressing and IPv4 Routing Interact**

Although an IPv4 routing infrastructure can use either class-based or classless routing, the development of classless routing protocols has made classful routing obsolete for medium-sized or enterprise networks. When routing protocols are referred to as classful or classless, these designations refer to routing protocols that can handle classful or classless unicast addressing as follows:

- **Classful addressing.** Classful routing protocols recognize only classful addresses — that is, the standard IPv4 unicast address classes. Classful routing protocols determine the network ID based on the IP address class alone, do not advertise a subnet mask when advertising routes, and therefore cannot recognize a summarized route (a single route that represents a range of network IDs) in routing advertisements. Routing for multiple subnets created from one network ID is possible when classful routing protocols are used, but all subnets of a class-based network ID are of equal size and must be contiguous or else routing errors can occur. RIP version 1 (RIP v1) is an example of a classful routing protocol.

- **Classless addressing.** Classless routing protocols recognize classless addresses — that is, IPv4 unicast addresses used for variable-length subnetting or route summarization as well as IP addresses defined as standard classful IPv4 unicast addresses. Classless routing protocols include the subnet mask along with the IP address when advertising routing information and can therefore recognize and announce subnets of varying sizes as well as summary routes (each of which groups several subnets as a single entry in a routing table). Classless routing protocols let you create noncontiguous subnets from the same network ID; routers can correctly recognize these noncontiguous subnets. RIP version 2 (RIP v2), OSPF, and Border Gateway Protocol version 4 (BGP v4) are examples of classless routing protocols.

Understanding IPv4 unicast addressing is essential to understanding unicast IP routing because IP routing uses unicast IP addresses to direct packets to their destination. Medium-sized and large organizations today typically implement classless routing solutions, thus taking advantage of the flexibility that classless addressing enables. The largest IP internetwork in the world, the Internet, supports classless routing, although it was designed originally as a class-based routing infrastructure.

**Note**

The summary of IPv4 unicast addressing described in the following sections is provided as a basis for understanding the role that addressing plays in routing.

For a more complete discussion of IPv4 addressing, including unicast, multicast, and broadcast addressing, as well as an in-depth discussion of addressing, subnetting, and route summarization, see "Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference" by Joseph Davies and Thomas Lee, 2003, Microsoft Press.

This section describes:

- How classful addressing limits classful routing
- How classless addressing enables classless routing

**How Classful Addressing Limits Classful Routing**

Although the development of classless IP addressing has made classful IP addressing obsolete for enterprise networks, configuring classless IP addresses requires knowledge of the classful IP addressing structure out of which classless addressing later developed.

The only flexibility that classful addressing offers is the configuration of fixed-size subnets. Because classful routing protocols, such as RIP v1, recognize only the original standard IP address classes and do not include the subnet mask in routing advertisements, neither variable-length subnetting nor route summarization are possible.

This section includes:

- Classful addressing summary
- Classful subnetting enables only fixed-size subnets

**Classful Addressing Summary**

Of the five standard address classes, Class A, B, and C addresses, collectively known as IPv4 unicast addresses, are assigned to specific devices on an IPv4 network. Class D addresses, known as multicast addresses, are used for IP multicasting, simultaneously sending a message from one source to more than one network destination (for more information about multicasting, see IPv4 Multicasting Technical Reference). Class E addresses are reserved for experimental purposes.

To be able to configure fixed-size subnetting on your network (as well as variable length subnetting or route summarization, described later), you must first understand the default formats of the unicast IPv4 addresses. All 32-bit IPv4 addresses contain four 8-bit bytes called *octets* separated for readability by spaces (in binary notation) or represented as four decimal numbers separated by dots (in dotted decimal notation). If all or some of the actual numbers are not known, all or part of the IP address in dotted decimal notation is typically represented by *w.x.y.z*. The following table shows a single IP address in both decimal and binary notation.

**IP Address in Decimal and Binary Notation**

| IP Address in Decimal Notation | IP Address in Binary Notation |
|---|---|
| 131.107.65.37 | 10000011 01101011 01000001 00100101 |

**Network ID bits and host ID bits**

Like an IP address, a network ID in an IP internetwork is typically expressed in dotted decimal or binary notation. It is important to recall, however, that a network ID is not an IP address. A network ID is the leftmost portion of an IP address and identifies the set of nodes that are located on the same subnet — the network ID is the network address shared by all network interfaces on one subnet. An IP address, made up of the network ID portion and the host ID portion combined, is the unique address assigned to one network interface on a node.

Classful unicast IPv4 addresses have the following formats:

- **Class A.** The first byte, or octet, represents the network ID, and the three remaining bytes are used for the host ID.
- **Class B.** The first 2 bytes represent the network ID, and the last 2 bytes are used for the host ID.
- **Class C.** The first 3 bytes are used for the network ID, and the final byte is used for the host ID.

**Unicast IPv4 address ranges**

The following table shows the address ranges and the number of possible networks and hosts for each type of IPv4 unicast address class. Keep in mind that a network ID is not an IP address, even though both are expressed by using dotted decimal notation.

**Classful Unicast IPv4 Address Classes**

| Class | Leading Bit Pattern | First Network ID in Range | Last Network ID in Range | Available Networks | Hosts per Network |
|---|---|---|---|---|---|
| A | 0 | 1.0.0.0 | 126.0.0.0 | 126 | 16,777,214 |
| B | 10 | 128.0.0.0 | 191.255.0.0 | 16,384 | 65,534 |
| C | 110 | 192.0.0.0 | 223.255.255.0 | 2,097,152 | 254 |

| Note |
|---|
| The network ID 127.0.0.0 is reserved for loopback addresses, which are used for testing TCP/IP operation. |

The range in the first octet is used to identify the address class. For example, an address in the format 10.*x.y.z* is immediately recognizable as a class A address because 10 falls within the range 1–126. Similarly, an address in the format 131.107.*y.z* is identifiable as a class B address, and an address in the format 206.73.118.*z* is a class C address.

**Classful subnet masks**

A subnet mask, paired with an IP address, is a 32-bit value that enables a host or router to distinguish the network ID and host ID portions of its associated IPv4 address and thus determine the network ID (shared by all nodes on that subnet) for a specific IP address. When binary notation is used, the network ID bits in the subnet mask are identifiable because they are set to 1 and the host ID bits are identifiable because they are set to 0. When decimal notation is used, the network ID bits for a classful subnet mask are set to 255 (which is eight ones when translated into binary), and the host ID bits are set to 0. The following table illustrates the structure of the subnet masks used for classful unicast IPv4 addresses.

**Subnet Masks Used to Distinguish Network and Subnet Bits**

| Class | Decimal Subnet Mask: Network ID Bits 255 and Host ID Bits 0 | Binary Subnet Mask: Network ID Bits 1 and Host ID Bits 0 |
|---|---|---|
| A | 255.0.0.0 | 11111111 00000000 00000000 00000000 |
| B | 255.255.0.0 | 11111111 11111111 00000000 00000000 |
| C | 255.255.255.0 | 11111111 11111111 11111111 00000000 |

Although the table lists the default subnets masks now used for each classful address type, subnet masks were actually not defined until RFC 950, which defines subnetting procedures for fixed-size subnets, was published in August 1985.

For a technique for using an {IP Address, Subnet Mask} pair to determine the network ID for any type of IP address, including classless addresses, see “Addressing Mechanisms Required by Classless Routing” later in this section.

**Classful Subnetting Enables Only Fixed-Size Subnets**

Without some means of subdividing networks, all available IPv4 addresses would have been depleted long ago. Classful subnetting, dividing a network into subnets of equal size rather than simply using the default classful formats, is the initial strategy that was defined in RFC 950 to handle this problem and is still often used in smaller organizations.

The RIP v1 routing protocol, which does not advertise a subnet mask with the route, can reliably route only fixed-size subnets. RIP v1 advertises only the class-based network ID and therefore cannot perform routing in an environment that includes variable-size subnets.

**Using host ID bits to configure multiple subnets**

Subnet addressing lets you use one or more bits from the host ID part of the IPv4 address to configure the division of the network into multiple subnets. When subnetting is used, the subnet mask consists of the octets assigned to the network ID (set to 1) plus one or more bits for the subnet (now also set to 1). Therefore, when subnetting is implemented, the number of bits in the network ID and the host ID, respectively, vary and cannot be inferred by looking at the IP address alone. However, when each IP address on a network divided into subnets is paired with a subnet mask, the subnet mask is used to distinguish the network and subnetting bits (set to 1) from the host ID bits (set to 0).

The following table shows three alternative ways to use 131.107.65.37.

**Three Alternative Fixed-Size Subnet Masks for One IP Address**

| Bits Used for Subnets | Decimal Subnet Mask | Binary Subnet Mask | Fixed-Size Subnets | Hosts per Subnet |
|---|---|---|---|---|
| 0 | 255.255.0.0 | 11111111 11111111 00000000 00000000 | 1 | 65,534 |
| 8 | 255.255.255.0 | 11111111 11111111 11111111 00000000 | 256 | 254 |
| 4 | 255.255.240.0 | 11111111 11111111 11110000 00000000 | 16 | 4,094 |

Using the class B address 131.107.65.37 with its default subnet mask of 255.255.0.0 (first row in the example table) provides one subnet that can have up to 65,534 hosts. Because, typically, you do not need that many hosts on a single subnet, you might decide to create multiple subnets. Of the many possible choices you can make, the second and third rows in the example illustrate two:

- If you use a subnet mask that allocates an additional 8 host ID bits for subnets (shown in the second row), you obtain 256 (that is, $2^8$) subnetted network IDs (256 subnets of equal size), each of which can support up to 254 hosts.

  **Note**

  The number of hosts per subnet is listed as 254 rather than 256 because although 8 bits (28) are reserved for the host ID, you must subtract 2 (that is 28–2) because many vendors, including Microsoft, support only subnetting that excludes host IDs that consist of all ones or all zeros.

- If you use a subnet mask that allocates 4 host ID bits for subnets (shown in the third row), you obtain 16 (that is $2^4$) subnets of equal size, each of which can support up to 4,094 hosts ($2^{12}$–2).

**Using succinct network prefix notation to indicate subnetting**

Network prefix length notation, an alternative type of notation, lets you represent an {IP Address, Subnet Mask} pair in a different, more succinct form. The format is *IP address/p*, where *p* defines the number of network ID bits in the subnet mask. The remaining bits (calculated by subtracting the number of bits, *p*, from 32) represent the number of host ID bits in the subnet mask. The network ID bits (the bits set to 1) are considered the prefix because, when viewed in binary format, they always precede the host ID bits (the bits set to 0).

Network prefix length notation is sometimes referred to as Classless Interdomain Routing (CIDR) notation. CIDR, as its name implies, also encompasses the entire area of classless routing, including the use of subnet masks to enable variable-length subnets and route summarization, both described later.

The following table again shows the same address, 131.107.65.37, configured in three alternative ways, this time contrasting each address in network prefix length notation (left column) with the more cumbersome {IP Address, Subnet Mask} pair in both decimal and binary notation. Note that the number of bits set to 1 in the binary subnet mask is the same as the number after the forward slash in the network prefix length notation version of the same address, and the number of bits set to 0 is equal to 32 minus *p*.

**Network Prefix Length Notation Compared to {IP Address, Subnet Mask} Notation**

| Network Prefix Length Notation | Address / Subnet Mask (Decimal Notation) | Address / Subnet Mask (Binary Notation) |
|---|---|---|
| 131.107.65.37/16 | 131.107.65.37/255.255.0.0 | 10000011 01101011 01000001 00100101/<br>11111111 11111111 00000000 00000000 |
| 131.107.65.37/24 | 131.107.65.37/255.255.255.0 | 10000011 01101011 01000001 00100101/<br>11111111 11111111 11111111 00000000 |
| 131.107.65.37/20 | 131.107.65.37/255.255.240.0 | 10000011 01101011 01000001 00100101/<br>11111111 11111111 11110000 00000000 |

**How Classless Addressing Enables Classless Routing**

An IP internetwork that uses classless dynamic routing protocols can take advantage of classless addressing to produce the following types of network configuration:

- Subnets, each of the same size

- Subnets of different sizes
- A set of network IDs grouped into a block represented by a summary route; or an extension of this technique creating a network structure that maps to a region, department, or building; and subnet structure with a summary route representing each level of the hierarchy.

To understand the development and purpose of classless routing, a quick summary of the historical background is helpful. In the early 1990s, the Internet community recognized two interrelated problems:

- **Inadequate address space.** The existing public IPv4 address space was inadequate for handling the exponential growth of the Internet.
- **Overloaded routing tables.** Existing routers and existing routing protocols were inadequate for handling the related explosive growth of routing tables.

The long-term solution to both of these problems is the development and deployment of IPv6 (for information about Windows Server 2003 support for IPv6, see IPv6 Technical Reference). The short-term solution, described here, is to make IPv4 address allocation more efficient and to enable the aggregation of routing information. The classless interdomain routing (CIDR) techniques described in the early 1990s in RFCs 1518 and 1519 promoted conservation of the IPv4 address space and reduced the size of routing tables.

Initially, the new address allocation scheme focused on combining a set of class C addresses into one logical network (as an alternative to using class B addresses, whose availability is limited). A related goal was to create a hierarchical rather than a flat network topology, enabling a medium-size or large organization to tailor the structure of its network to match its current needs while allowing for future growth. One CIDR technique is route summarization, in which a single summary route is used to represent each level of the hierarchy, thus reducing the size of routing tables. Another CIDR technique provides a method for subdividing networks into variable-size subnets.

A major benefit of the new address allocation scheme used in conjunction with routing protocols capable of handling classless routing (such as RIP v2 or OSPF) was to free the IPv4 address space from the limitations of the original class-based IPv4 addressing system.

This section describes:

- Addressing mechanisms required by classless routing
- Subnets of varying sizes enabled by VLSM
- CIDR blocks and hierarchical routing enabled by route summarization

**Addressing Mechanisms Required by Classless Routing**
Classless routing requires the following addressing mechanisms:

**Subnet mask explicitly advertised by routing protocol**
A classless routing protocol, by definition, is one that advertises the subnet mask with the route. Routers that support a classless routing protocol can successfully route IP packets among subnets of variable size and can advertise summary routes.

**{IP Address, Subnet Mask} pair used to derive network ID**
In a routing environment that uses classless routing, the bitwise logical AND operation performed on an IP address and its subnet mask yields the network ID.

For example, {131.107.0.0, 255.255.0.0} denotes the network ID. In the subnet mask, the 16 leftmost contiguous bits are set to 1 (recall that 255 in decimal notation is 11111111 in binary). If you perform the bitwise logical AND operation on the IP address 131.107.20.30 with the subnet mask 255.255.0.0, the result is the network ID 131.107.0.0, sometimes written simply as 131.107. The following table illustrates the logical AND operation performed on each bit in the IP address with each corresponding bit in the subnet mask — the result is 1 only if both bits are 1; otherwise, the result is 0.

**Logical AND Operation to Find the Network ID**

| Description | Decimal Notation | Binary Notation |
|---|---|---|

| | | |
|---|---|---|
| IP address | 131.107.20.30 | 10000011 01101011 00010100 00011110 |
| Subnet Mask | 255.255.0.0 | 11111111 11111111 00000000 00000000 |
| Result of AND operation = network ID | 131.107.0.0 | 10000011 01101011 00000000 00000000 |

If you perform the AND operation on any other class B IP address with the class B default subnet mask, 255.255.0.0, the result is also the network ID 131.107.0.0.

For class A, B, and C addresses using the default subnet mask, the AND operation is not needed because you already know, by definition, the subnet mask and network ID for any classful IP address. However, you can use the AND operation to determine the network ID for subnetted classful IP addresses and for classless IP addresses for which no default subnet mask or definition of the network ID exists.

**{Network ID, Subnet Mask} pair used to express a summary route**

Classless routing can use a {Network ID, Subnet Mask} pair to represent a summary route. As explained earlier, a summary route is a single entry in a routing table that represents a set of entries, that is, the summary route represents a range of network IDs. RFC 1519 introduced the use of a summary route to express a range of class C network IDs to be consolidated into a single block, but the RFC also suggested that an arbitrary network number and mask can be configured for all routing destinations.

**Longest match routing**

With classless routing, routing to all destinations is done on a longest-match basis. This means that, for a specific destination that matches multiple {Network ID, Subnet Mask} pairs in the routing table, the router forwarding the IP packet uses the match with the longest mask (the mask with the most number of bits set to 1) to forward the packet.

For more information about these mechanisms, see How the IPv4 Routing Table Works later in this document.

**Subnets of Varying Sizes Enabled by VLSM**

As described earlier, standard subnetting creates subnets of equal size in a networking environment that supports only classful routing protocols. Subnets of equal size can be inefficient if each subnet does not need to support the same number of hosts. Therefore, a more useful method of subnetting, known as *variable length subnetting*, can be configured by using variable length subnet masks (VLSMs). Classless routing protocols support variable length subnetting, which is defined in RFC 1878.

Using VLSMs, you can perform subnetting on the same address space (network ID) multiple times to produce subnets of different sizes on the same network. All subnetted network IDs are unique and can be distinguished from each other by their corresponding subnet mask.

For example, by starting with a classful class B network ID of 131.107.0.0/16, you can configure one subnet that supports 1,022 hosts, a second that supports 510 hosts, and a third that supports only two hosts (for a point-to-point WAN link between two routers). The following table shows how.

**Example Using Three VLSMs to Create Three Various-Sized Subnets from 131.107.0.0**

| Network Prefix Length Notation | Decimal Mask | Binary Mask | Host |
|---|---|---|---|
| Subnet 1:<br><br>131.107.0.0/22 | 255.255.252.0 | 11111111 11111111 11111100 00000000 | 1,022<br><br>($2^{10}$–2) |
| Subnet 2:<br><br>131.107.208.0/23 | 255.255.254.0 | 11111111 11111111 11111110 00000000 | 510<br><br>($2^{9}$–2) |
| Subnet 3:<br><br>131.107.208.84/30 | 255.255.255.252 | 11111111 11111111 11111111 11111100 | 2<br><br>($2^{2}$–2) |

In this example, creating three subnets, each supporting a variable number of hosts, is the hypothetical goal. However, this simple example does not exhaust the possibilities. Instead of simply creating three subnets with 1,022, 510, and 2 hosts, respectively, you could create a more complex result by taking advantage of the number of possible subnets as well as the number of possible hosts. Because 131.107.0.0/16 is a class B network ID, the number of possible subnets is determined by the number of bits set to 1 in the third and fourth octets:

- **22-bit prefix length (Subnet 1):** 64 ($2^{6}$) possible subnets, each supporting up to 1,022 hosts.
- **23-bit prefix length (Subnet 2):** 128 ($2^{7}$) possible subnets, each supporting up to 510 hosts.

- **30-bit prefix length (Subnet 3):** 16,384 ($2^{14}$) possible subnets, each supporting up to 2 hosts.

When using VLSM, make sure that you do not overlap blocks of addresses. A good technique is to start with equal-size subnets and then subdivide them.

If variable-length subnets are also contiguous, the routes for all of the subnets can be summarized: a router can be configured to advertise the set of subnets to another router by advertising only the original network ID. Route summarization is described in the next section.

**CIDR Blocks and Hierarchical Routing Enabled by Route Summarization**

This section describes how classless routing uses a summary route to represent a range of network IDs (also known as a *CIDR block*) and how a multi-level hierarchy uses summary routes at each level of the hierarchy to make routing more efficient.

**How a summary route represents a range of network IDs**

As explained earlier, route summarization was introduced by RFC 1519 as a way to consolidate several class C network IDs into one logical network ID, sometimes referred to as a *supernet*. Because not enough class B addresses were available to meet the escalating demand, this technique enables two or more class C network IDs to be allocated to an organization in order to provide an adequate number of IP addresses.

Whereas subnetting enables the division of networks into smaller subnets of equal or variable size, route summarization enables the grouping of contiguous subnets. The use of a subnet mask is required for route summarization just as it is for subnetting. However, these subnets masks differ as follows:

- **Subnet mask used in subnetting.** Uses one or more host ID bits to create a longer subnet mask (prefix length) and uses these bits to create subnets. Subnet masks for subnetted network IDs are always more specific (that is, have more bits set to 1) than the network ID that is being subnetted.

- **Subnet mask used in route summarization.** Uses one or more network ID bits to create a shorter subnet mask (prefix length) in order to represent a set of network IDs as a single route in the routing table. This group of network IDs is the CIDR block. Subnet masks for summarized network IDs are always less specific (that is, have fewer bits set to 1) than the network IDs that are being summarized.

The following table lists a series of eight class C network IDs that, in this example, are combined into one block. The first network ID is 208.147.168.0. The first 21 bits are the same for each of the eight network IDs grouped into the single CIDR block. The last 3 bits of the third octet, which are taken from the class C–based network ID, range from 000 through 111. In decimal notation, the range is 0 through 7, or a total of 8 contiguous subnets, combined to form one block.

**Eight Class C Network IDs to be Combined into One Block**

| Range | Network ID (Decimal) | Network ID (Binary) |
|---|---|---|
| 0 (first network ID) | 208.147.168.0 | 11010000 10010011 10101000 00000000 |
| 1 (second network ID) | 208.147.169.0 | 11010000 10010011 10101001 00000000 |
| 2 (third network ID) | 208.147.170.0 | 11010000 10010011 10101010 00000000 |
| 3 (fourth network ID) | 208.147.171.0 | 11010000 10010011 10101011 00000000 |
| 4 (fifth network ID) | 208.147.172.0 | 11010000 10010011 10101100 00000000 |
| 5 (sixth network ID) | 208.147.173.0 | 11010000 10010011 10101101 00000000 |
| 6 (seventh network ID) | 208.147.174.0 | 11010000 10010011 10101110 00000000 |
| 7 (eighth network ID) | 208.147.175.0 | 11010000 10010011 10101111 00000000 |

This CIDR block is expressed as {208.147.168.0, 255.255.248.0} or, in network prefix length notation, as 208.147.168.0/21. A single route in the IP routing table represents this block. Without a list such as that shown in the preceding table, you can determine the number of network IDs in the block by raising the number of summarized bits to the power of 2. In this class C example, the subnet mask indicates that the number of bits used for summarization is three: 24–21=3, that is, 24 bits set to 1 in a standard class C subnet mask minus 21 actual bits set to 1 leaves 3 bits for summarization. Therefore, there are $2^3$=8 network IDs in this block.

Looking at this CIDR block (which, in this simple example, does not use variable length subnetting), you can double-check that it summarizes 8 class C networks by confirming that both of the following calculations produce the same number:

- Possible host IDs supported by a classful class C network ID (256) × network IDs in the block (8) = the number of possible host IDs (2048).

- Two (2) raised to the number of bits set to 0 in the subnet mask ($2^{11}$) = the number of possible host IDs (2048).

The assignable bits can be used either simply for host IDs or to create subnets of variable size, as the needs of the organization dictate. Typically, using VLSM to divide the 2048 possible host IDs into variable size subnets is the more practical option.

In this example, as the two calculations demonstrate, the CIDR block can be viewed from two perspectives:

- As a block of 8 contiguous class C network IDs grouped into one logical network ID.

- As a classless address space that has 21 fixed bits and 11 assignable bits.

**How a multi-level hierarchy using summary routes streamlines routing**

Hierarchical routing requires a method of address allocation that can reflect a hierarchical network topology and that can reduce the size of routing tables. Hierarchical routing accomplishes both requirements by using a single route to advertise a set of routes at each level of the hierarchy.

ISPs, organizations that subscribe to ISPs, and subnets within organizations are some of the typical levels in the routing hierarchy of the public Internet. An element lower in the hierarchy reports summary routing information to the level above it. For example, the allocation of large blocks of the public IPv4 address space to network service providers (such as ISPs), which then allocate subsets of these blocks to smaller organizations (such as organizations that subscribe to an ISP) is a practice recommended in RFCs 1518 and 1519.

Thus, a router within an organization does not need to advertise routing information about every subnet in that organization. Instead, the router can be configured to maintain routing information on a per region basis, in which each routing region is summarized with a single route except for the regions to which the router is connected. Likewise, an ISP router does not need to advertise routing information about individual subnets within each of the organizations that make up the set of its subscribers. Instead, the router can maintain routing information on a per subscriber basis.

Classless routing protocols, which recognize a summarized route represented by a {Network ID, Subnet Mask} pair that represents a set of reachable IP addresses, are designed to handle routing in a hierarchical infrastructure. By expressing a range of network IDs as a single route in the routing table at each level of the hierarchy, route summarization decreases the number of routes on routers. With fewer routes to check, routers can forward packets faster. One major use of route summarization is to reduce the burden for routers on the Internet backbone.

The example shown in the following figure depicts a simple hierarchical routing structure divided into routing regions based on the private network ID 10.0.0.0/8. Route summarization is used at the border of each region. Each arrow is labeled with the summarized route that is advertised outside its region by the router or routers at the border of that region. (This example figure is used by permission from the authors of "Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference" by Joseph Davies and Thomas Lee, Redmond, WA, Microsoft Press, 2003.)




The structure shown in this example illustrates both the advantage of hierarchical routing and its cost:

- **Advantage.** The efficiency provided by hierarchical routing reduces the size of routing tables, reduces the amount of communication required between routers in order for them to update their routing tables, and therefore reduces network traffic. With smaller routing tables, routers forward packets faster.

- **Cost.** The initial planning required to establish such a structure takes time. In a large organization, this initial investment in time is well worth the effort because of the long-term efficiency gained.

#### How the IPv4 Routing Table Works

On an IPv4 internetwork, IP makes routing decisions by consulting a database of routes known as the *IP routing table*. Each IP node uses its own routing table to store information about routes usable for forwarding unicast IP packets to any reachable location. A node running Windows Server 2003 (or Windows XP) builds its routing table automatically, based on its current TCP/IP configuration, static routing table entries, information provided by dynamic routing protocols, and ICMP Redirect messages.

This section describes:

- IPv4 routing table
- Structure of the IPv4 routing table
- How IP consults the routing table
- How to manage the IPv4 routing table

#### IPv4 Routing Table

Each entry in the routing table corresponds to one route. When an IP node needs to forward an IP packet, IP searches the routing table on the node for a route that most closely matches the destination IP address in the packet. Typically, the process is as follows:

- **For an IP host.** A host can either send a packet directly to its destination, or it can send the packet en route to the destination through a default route {Network ID: 0.0.0.0, Subnet Mask: 0.0.0.0} that points to its default gateway. A *default gateway* is a router that connects separate IP network segments.
- **For an IP router.** A router forwards a packet either by using an explicit static route for a specific network segment, a summarized route, a host route, or a default route.

#### Routing Table Example

An administrator or user can view the current state of the IP routing table on an IP node at any time. For example, consider a node configured as follows:

- **IP address:** 157.54.27.90
- **Subnet mask:** 255.255.240.0
- **Default gateway:** 157.54.16.1

In network prefix length notation, the IP address and subnet mask pair {157.54.27.90, 255.255.240.0} for this node is expressed as 157.54.27.90/20. Typing **route print** at a command prompt on this node produces the output shown in the following table.

**Example Routing Table**

```
Network Destination        Netmask          Gateway       Interface   Metric
        0.0.0.0            0.0.0.0     157.54.16.1   157.54.27.90       20
      127.0.0.0          255.0.0.0       127.0.0.1      127.0.0.1        1
    157.54.16.0      255.255.240.0   157.54.27.90   157.54.27.90       20
   157.54.27.90    255.255.255.255       127.0.0.1      127.0.0.1       20
 157.54.255.255    255.255.255.255   157.54.27.90   157.54.27.90       20
      224.0.0.0          240.0.0.0   157.54.27.90   157.54.27.90       20
255.255.255.255    255.255.255.255   157.54.27.90   157.54.27.90        1
```

This routing table, displayed on a computer running Windows Server 2003, differs from the routing table on a computer running Windows 2000 Server in the following ways:

- **New netmask for the multicast address route.** In the routing table on a server running Windows 2000 Server, the

Network Destination field and the Netmask field for the route used for IPv4 multicast addresses (defined as those in the range from 224.0.0.0 through 239.255.255.255) are both 224.0.0.0. However, the subnet mask used for IPv4 multicast addresses is 240.0.0.0. Therefore, the new Netmask of 240.0.0.0 that is used in the Windows Server 2003 routing table is more appropriate than the 224.0.0.0 Netmask used in Windows 2000, based on the class D definition of IPv4 multicast addresses.

- **Automatic metric feature.** On a server running Windows 2000 Server, the default routing metric for each route (each row) in the routing table is set to 1. On a computer running Window Server 2003 or Windows XP, the routing metric for each route is set to a value determined by the link speed of the interface. (This automatic metric feature does not apply to the default route metric when the default gateway is manually configured with a metric.)

In the output shown in "Example Routing Table," the routing metric that is set to 20 for several rows is the metric that Windows Server 2003 or Windows XP TCP/IP uses for a 100-Mbps Ethernet interface, for an 802.11g adapter, or for an 802.11a adapter. The following table lists the criteria that Windows Server 2003 or Windows XP uses to assign metrics for routes that are bound to network interfaces of varying speeds.

**Metrics Assigned by the Automatic Metric Feature**

| Link Speed | Metric |
|---|---|
| Greater than 200 Mbps | 10 |
| Greater than 20 Mbps, and less than or equal to 200 Mbps | 20 |
| Greater than 4 Mbps, and less than or equal to 20 Mbps | 30 |
| Greater than 500 Kbps, and less than or equal to 4 Mbps | 40 |
| Less than or equal to 500 Kbps | 50 |

The automatic metric feature is enabled by default through the **Automatic metric** check box on the **IP Settings** tab on **Advanced TCP/IP Settings** of the TCP/IP protocol. For DHCP-assigned default gateways, you can override the default behavior of automatically calculating a metric for the default route based on adapter speed by using the Microsoft-specific DHCP option called **Default Router Metric Base.**

**Routing Table vs. Forwarding Table**

Many router implementations, including the Windows Server 2003 Routing and Remote Access service, use both a routing table and a forwarding table:

- **Routing table.** Stores all routes from all possible sources. The command **netsh ro ip sh rtmr** (that is, **netsh routing ip show rtmroutes**) shows the routing table.

- **Forwarding table.** Used by IP to forward packets. The commands **route print** or **netsh ro ip sh rtmd** (that is, **netsh routing ip show rtmdestinations**) show the IP forwarding table.

The Routing and Remote Access service on a Windows Server 2003 router maintains the IP routing table by using a component called *Route Table Manager* (abbreviated as **rtm** in **rtmroutes** and **rtmdestinations** in the commands shown earlier). Route Table Manager updates the IP forwarding table (which is contained within the TCP/IP protocol) based on incoming route information from multiple sources.

The contents of the routing table do not necessarily match the contents of the forwarding table. However, in this introductory discussion, the routing table and the forwarding table are treated as if they are identical.

**Routing Table Stored in RAM on Local Node**

On an IPv4 internetwork, all routing decisions made by a host or a router are based on information stored in the IP routing table on that node. Because the routing table is a local database that physically resides in the remote access memory (RAM) of the node making the routing decision, the IP routing table is rebuilt whenever the node restarts. Thus, there is no single, unified view of the internetwork that is gathered by a server and downloaded to each host and router — all nodes do not have the same view of the internetwork, and all traffic does not flow along predictable pathways.

Each router in a path between a source and a destination makes a local routing decision based on its local routing table. Therefore, the path that packets take from the source to the destination might not be the same as the path taken by response packets from the destination back to the source.

If the information in the local routing table of a node is incorrect due to configuration errors or changing network conditions, routing problems can result. Troubleshooting routing problems might involve the analysis of the routing tables of the hosts (source and destination) and analysis of all of the routers that forward packets between them.

**Structure of the IPv4 Routing Table**

Understanding the IPv4 routing table requires familiarity with:

- The types of routes stored in the routing table.
- Available methods for configuring a default gateway, any of which creates a default route in the IP routing table.
- The meaning of the columns and rows that make up the routing table database.

**Types of Routes**

The following table describes the types of routes stored in an IP routing table. For a description of each route in an actual routing table, see the table "Routes in an IP Routing Table" later in this section.

**Types of Routes Stored in an IP Routing Table**

| Route Type | Description |
|---|---|
| Local network route<br><br>(Directly attached network route) | A route to a specific local network ID. This route identifies a network segment that is directly attached to the node.<br><br>For a local network route, the Gateway column (sometimes called *Next Hop*) might be blank or might contain the IP address of the interface on that network segment. |
| Remote network route | A route to a specific remote network ID. This route identifies a network segment that is not directly attached to the node but is available across one or more routers.<br><br>For a remote network route, the Gateway (Next Hop) column is the IP address of a local router located between the node and the remote network. |
| Host route | A route to a specific IP address (network ID + host ID) on the internetwork. Instead of making a routing decision based on just the network ID, as is the case for either a local or a remote network route, the routing decision for a host route is based on the combination of network ID and host ID. For a host route, the Network Destination column is the IP address of the specified host, and the Netmask column is 255.255.255.255.<br><br>Typically, a host route is used to create a custom route to a specific host in order to control or optimize specific types of internetwork traffic. |
| Default route | A route that is used when there are no other closer matching routes for the destination in the routing table. Including a default route in the routing table means that the routing table does not need to store routes for every network ID on the internetwork. Using a default route, therefore, simplifies the configuration of hosts or routers. |

**Default Gateways and Default Routes**

Configuring a default gateway creates a default route in the IP routing table. For IP nodes, the default gateway (also called a *default router*) is a neighboring IP router that forwards unicast traffic for the node by providing a next-hop IP address (the Gateway column in the routing table) and interface (the Interface column in the routing table) for all destinations that are not located on the local subnet. The default gateway address is the IP address of a directly reachable IP router.

The default gateway on a computer running Windows Server 2003 or Windows XP is configured by using one of the following methods:

- **TCP/IP properties General tab**. If you want the node to obtain its IP address configuration by using manual configuration, the default gateway is the IP address located in the **Default gateway** field on the **General** tab on the TCP/IP properties page. Multiple default gateways are configurable by adding them on the **IP Settings** tab on the **Advanced TCP/IP Settings** page of TCP/IP properties.
- **TCP/IP properties Alternate Configuration tab.** If you want the node to obtain its IP address configuration by using the user-configured alternate configuration option, the default gateway is the IP address located in the **Default gateway**

field on the **Alternate Configuration** tab on the TCP/IP properties page. You can specify only a single default gateway. The TCP/IP alternate configuration feature lets a computer function in two or more networks — one configured with static IP addresses and another configured with DHCP — without reconfiguration of network adapter parameters. When no DHCP server is available, the connection uses the configuration specified on the **Alternate Configuration** tab.

- **DHCP.** If you want the node to obtain its IP address configuration by using DHCP, the default gateway is the value of the first IP address in the Router DHCP option. The Router DHCP option specifies an ordered list of one or more default gateways. If you are using a Windows Server 2003–based DHCP server, the Router option is located in the **Details** pane in the DHCP snap-in console tree under *ServerName**ScopeName***Scope Options**.

- **ICMP Router Discovery.** If you want the node to automatically discover the best default gateway router available on a subnet, you can configure nodes on that subnet to use ICMP router discovery. ICMP Router Solicitation messages and Router Advertisement messages exchanged between routers and hosts enable hosts to dynamically discover which local routers are available, which routers are down, and which router is currently the best default gateway to use on a subnet. By using ICMP router discovery, a host can automatically switch to another default gateway if its current default gateway becomes unavailable. For information about support for ICMP router discovery in the Windows Server 2003 Routing and Remote Access service, see "How ICMP Router Discovery Works" in How Unicast IPv4 Routing Protocols and Services Work.

| **Note** |
|---|
| If the host obtains its IP address configuration by using Automatic Private IP Addressing (APIPA), no default gateway is configured and no default route is created in the routing table. APIPA is useful only for a single subnet. |

**Fields and Records in the IP Routing Table Database**

As with any database, understanding the IP routing table requires understanding the relationship between the records (rows) and the fields (columns) that make up the IP routing table. This relationship is best understood by using an example. The following table shows the same information displayed in "Example Routing Table" earlier in this section. It labels each row with the route type for that entry.

**Example Information Stored in an IP Routing Table**

| Type | Network Destination | Netmask | Gateway | Interface | Metric |
|---|---|---|---|---|---|
| Default route | 0.0.0.0 | 0.0.0.0 | 157.54.16.1 | 157.54.27.90 | 20 |
| Loopback network route | 127.0.0.0 | 255.0.0.0 | 127.0.0.1 | 127.0.0.1 | 1 |
| Local network route<br><br>(Directly attached network route) | 157.54.16.0 | 255.255.240.0 | 157.54.27.90 | 157.54.27.90 | 20 |
| Local host route<br><br>(Local IP address route) | 157.54.27.90 | 255.255.255.255 | 127.0.0.1 | 127.0.0.1 | 20 |
| Network broadcast route<br><br>(All-subnets-directed broadcast route) | 157.54.255.255 | 255.255.255.255 | 157.54.27.90 | 157.54.27.90 | 20 |
| Multicast address route | 224.0.0.0 | 240.0.0.0 | 157.54.27.90 | 157.54.27.90 | 20 |
| Limited broadcast route | 255.255.255.255 | 255.255.255.255 | 157.54.27.90 | 157.54.27.90 | 1 |

**Routing table columns**

The following table describes each field, or column, in the IP routing table.

**Columns for Each Routing Table Entry**

| Column | Description |
|---|---|

| | |
|---|---|
| Network Destination | Contains one of the following:<br>• **IP address.** An IP address for a host route<br>• **Network ID.** A class-based, subnetted, or summarized network ID |
| Netmask<br><br>(Subnet Mask) | Contains the subnet mask that IP applies to the destination IP address. If the destination AND the subnet mask equals the network destination, the route is a match for the destination. The process is as follows:<br>1. IP performs a bit-wise logical AND operation between the packet's destination IP address and the value in the Netmask column.<br>2. The result is compared to the value in the Network Destination column.<br>3. When IP compares each bit in the destination IP address to the corresponding bit in the subnet mask, the bit-wise AND operation works as follows:<br>  • If both bits are 1, the resulting bit is 1.<br>  • If both bits are not 1, the resulting bit is 0.<br><br>**Examples:**<br>• **Default route.** The netmask for the default route is 0.0.0.0, which translates to the binary value 00000000 00000000 00000000 00000000. When IP performs the AND operation between any destination IP address and this subnet mask, the result is all zeros. Therefore, no bits need to match, and any destination matches the default route.<br>• **Host route.** The netmask for a host route (a route that matches a specific IP address) is 255.255.255.255, which translates to the binary value 11111111 11111111 11111111 11111111. When IP performs the AND operation between a host route and this subnet mask, only a destination IP address that is equal to the IP address of the host route matches this route.<br>• **Packet with IP address 157.54.16.48.** When IP performs the AND operation between 157.54.16.48 and the netmask for the directly attached network route (in this example, 255.255.240.0), the result is 157.54.16.0. The subnet mask 255.255.240.0 requires that all eight bits in the first two octets and the first four bits in the third octet must match (240=11110000). Therefore, the closest match for this packet is the directly attached network (157.54.16.0/20). |
| Gateway<br><br>(Next Hop or Forwarding Address) | Contains either the IP address of a local network adapter for direct deliveries or the IP address of an IP router (such as a default gateway router) on the local network segment for indirect deliveries. This is the IP address that the local node uses to forward the IP packet. |
| Interface | Indicates the network interface (either a physical network adapter or a logical interface) that will be used to forward an IP packet. |

| | |
|---|---|
| Metric | Indicates the relative cost of routes so that the best route among possible multiple closest matching routes to the same destination can be selected. If there are multiple routes to the same destination with different metrics, the route with the lowest metric is selected.<br><br>Metrics can express the route preference in different ways:<br><br>• **Hop Count.** Indicates the number of routers to cross to reach a destination network. A computer on the local subnet is one hop, and each router crossed after that is an additional hop.<br><br>• **Delay.** Indicates the amount of time that the packet takes to reach the destination network. Delay refers to the speed (LAN links have a lower delay; WAN links have a higher delay) or congestion level of the path.<br><br>• **Throughput.** Indicates the amount of data that can be sent along the path per second. Throughput does not necessarily indicate the bit rate of the link because a very busy Ethernet link might have a lower throughput than an unused 64 kilobits per second (Kbps) WAN link.<br><br>• **Reliability.** Indicates a measure of the path constancy. Some types of links are more likely to fail than others. For example, leased-line WAN links are more reliable than dial-up lines.<br><br>RIP uses a hop count metric. OSPF can use hop count, delay, throughput, or reliability factors for its cost metric. |

**Routing table rows**

The following table describes each record, or row, in the IP routing table. Each row represents a route. In the Description column, all references are to the output "Example Routing Table" shown earlier.

**Routes in an IP Routing Table**

| Route | Destination / Netmask (Network Prefix Length Notation) | Description |
|---|---|---|
| Default route | 0.0.0.0/0.0.0.0<br><br>(0.0.0.0/0, sometimes written as 0/0) | The route used if there are no other closer matching routes for the destination address in the IP packet. In the example, the IP packet is forwarded to 157.54.16.1 (the default gateway) by using interface 157.54.27.90 (the IP address of the node). |
| Loopback network route | 127.0.0.0/255.0.0.0<br><br>(127.0.0.0/8) | The route used for any IP address of the form 127.*x.y.z* (127.0.0.0 through 127.255.255.255). In the example, as for any IP address in this range, the IP packet is forwarded to the special loopback address 127.0.0.1 by using the loopback interface. |
| Directly attached network route (Local network route) | *varies/varies* | The route used for an IP address on the locally attached subnet. The IP packet is forwarded directly to the destination IP address (not to an intermediate router). In the example, the IP packet is forwarded to its destination by using the interface assigned to 157.54.27.90 — that is, the Gateway and Interface columns contain an IP address of the node, indicating that the packet will be sent directly to its destination out a network adapter that is assigned an IP address on this node. |
| Local host route<br><br>(Local IP address route) | *varies*/255.255.255.255<br><br>(*varies*/32) | The route used if the IP address in the packet is the same as the IP address of the local host. The subnet mask 255.255.255.255 indicates that this is a host route. In the example, as for any IP address that matches the local host, the IP packet is forwarded to the loopback address 127.0.0.1 by using the loopback interface. |

| | | |
|---|---|---|
| Network broadcast route<br><br>(All-subnets-directed broadcast route) | *varies*/255.255.255.255<br><br>(*varies*/32) | The route used for an IP address that matches the all-subnets-directed broadcast address. The subnet mask 255.255.255.255 indicates that this is a host route. The IP packet is forwarded as a MAC-level broadcast to hosts on all subnets of a subnetted IP network ID by using an interface that is assigned an IP address on the local node.<br><br>A network broadcast route exists only if the local network ID is subnetted from a class-based network ID. |
| Multicast address route | 224.0.0.0/240.0.0.0<br><br>(224.0.0.0/4) | The route used for any of the class D IP addresses reserved for multicast traffic. The IP multicast packet is forwarded as a MAC-level multicast to the multicast group by using an interface that is assigned an IP address on the local node. |
| Limited broadcast route | 255.255.255.255/255.255.255.255<br><br>(255.255.255.255/32) | The route used for an IP address that matches the limited broadcast address. The subnet mask 255.255.255.255 indicates that this is a host route. The IP packet is forwarded as a MAC-level broadcast to all hosts on the same network segment by using an interface that is assigned an IP address on the local node.<br><br>The limited broadcast route is used during the DHCP configuration process, such as when a DHCP client uses the limited broadcast address to send all traffic while it waits for the DHCP server to acknowledge its lease of an IP address. |

**How IP Consults the Routing Table**

When an IP host or router needs to forward a packet to its destination, IP consults the routing table on that node to determine where to route the packet. This section provides a summary of how IP consults each row and column, and then provides a detailed description of the route determination process.

**Synopsis of How IP Consults the Routing table**

The following table summarizes how IP makes use of the IP routing table.

**Overview of How IP Uses Routing Table Rows and Columns**

| Action | Summary Steps |
|---|---|
| IP searches each row | IP checks each route in the table to find the closest matching route. Routes that match the destination vary from the most specific type of route to the least specific type of route. The closest matching routes have the following order:<br><br>1. Host route<br><br>2. Route representing the network ID for a local or remote subnet<br><br>3. Route representing a summarized network ID<br><br>4. Default route |
| Simultaneously, IP uses each column to assess each row | IP uses each column in every row as follows:<br><br>1. Network Destination and Netmask. Used together to determine if the destination address in the IP packet matches the route.<br><br>2. Metric. Used to determine which matching entry to select if two or more closest matching routes are found.<br><br>3. Gateway and Interface. Used together to determine the next-hop address to which to send the packet and the next-hop interface over which to forward the packet. |

| | |
|---|---|
| Result | Three alternative routing decisions are possible:<br>• Pass the packet up to the loopback interface (if the destination of the packet is this computer or of the form 124.*x.y.z*).<br>• Forward the packet by using one of the locally attached network interface cards (if a closest matching route was found).<br>• Discard the packet (if no matching route was found), and then (on a host) send an error message internally or (on a router) send an error message to the source host. |

The next section describes in detail the process by which IP determines which routing decision to make.

**Details of How IP Makes the Routing Decision**

IP determines the best route to a destination by comparing the destination IP address in the packet to each route currently present in the routing table. The detailed process is as follows:

1. **Determine whether the packet's destination IP address matches one or more routes.** For each entry in the routing table, IP on the node performs a bitwise logical AND operation between the destination IP address of the packet and the subnet mask listed in the Netmask column. IP compares the resulting value to the value in the Network Destination column:

   - **Match.** If one or more matching routes are found, IP compiles a list of the matching routes. IP determines if a route produces a match as follows:

     **Host route.** All 32 bits match the destination IP address. The route is to a single destination address.

     **Route representing the network ID for a local or remote subnet.** All of the bits in the network ID match the destination IP address. This route is either a route to a destination on the local subnet, or it is a route to a destination on a non-local network through a router.

     **Route representing a summarized network ID.** All of the bits in the summarized network ID match the destination IP address. The route is to a destination on the set of subnets summarized by this summary route.

     **Default route.** All destination IP addresses match the default route. This is the route used when there is no more specific match.

   - **No match.** If no route is found (including no default route), IP indicates an error condition. A "no match" condition cannot occur if a default route exists. If the node is a host, an IP routing error is sent internally to an upper layer protocol (such as TCP, UDP, or ICMP). If the node is a router, an ICMP Destination Unreachable-Host Unreachable message is sent to the source host.

   For examples illustrating the AND operation between a destination IP address and the value in the Netmask column, see the Netmask row in the table "Columns for Each Routing Table Entry" in "Structure of the IPv4 Routing Table" earlier in this section.

2. **Determine the single route to use to send or forward the packet.** The result of the route determination process is the choice of a single route in the routing table:

   - If one and only one route has the longest match — the *longest match* is the route with the highest number of bits set to 1 in the Netmask column (the longest prefix length) — that route is selected.

   - If multiple longest match routes are found (for example, multiple routes to the same network ID), IP selects the route with the lowest metric.

   - If multiple longest match routes with the same lowest metric exist, IP selects the route associated with the interface

that is first in the binding order.

| Note |
| --- |
| If a computer has more than one interface (network adapter), the binding order is the order in which the interfaces are accessed by network services. This order reflects the order in which TCP/IP is bound to each of the interfaces. To change the relative binding order of interfaces on a computer running Windows Server 2003 or Windows XP, open Network Connections, select the **Advanced** menu, select **Advanced Settings**, click the interface whose binding order you want to change, and then click the up arrow or the down arrow, as appropriate. |

3. **Determine the next-hop address and interface.** After a route is selected, IP determines from the routing table entry what the next-hop IP address is and which interface (physical network adapter or logical port) to use to forward the packet:

   - **Direct delivery (to the destination node).** If the destination is on a subnet directly connected to the host or router, IP delivers the packet to the destination node. In this case, the address in the Gateway column is the same as the address in the Interface column (or the Gateway column is blank), and the next-hop IP address is set to the destination IP address in the IP packet. The interface used is the one specified in the Interface column of the selected route.

     Again using the example shown in "Example Routing Table" described earlier, if traffic is sent to 157.54.16.48, the most specific route is the route for the directly attached network (157.54.16.0/20). The next-hop IP address is set to the destination IP address (157.54.16.48) and the interface used is the adapter card that has been assigned the IP address 157.54.27.90 (the interface on the local node).

   - **Indirect delivery (to the next router).** If the destination is not on a subnet directly connected to the node, IP delivers the packet to a neighboring router on a directly connected subnet for further routing. In this case, the address in the Gateway column is not the same as the address in the Interface column, and the next-hop IP address is set to the address in the Gateway column. The interface used is the one specified in the Interface column of the selected route.

     For example, if traffic is sent to 157.60.0.79, the most specific route is the default route (0.0.0.0/0). The next-hop IP address is set to the gateway address (157.54.16.1) and the interface used is, in this case again, the adapter card that has been assigned the IP address 157.54.27.90 (the interface on the local node).

4. **Hand off the packet to ARP.** IP hands the packet, the next-hop IP address, and the next-hop interface to ARP. ARP resolves the next-hop IP address to its MAC address and forwards the packet, as appropriate, to the next hop, which is either the destination or the router that will forward it towards its destination.

For more information about the IP routing process, including details of the host forwarding process, the router forwarding process, and the destination host receiving process, see "IP Routing from Sending Host to Destination" in "Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference" by Joseph Davies and Thomas Lee, Redmond, WA, Microsoft Press, 2003.

**How to Manage the IPv4 Routing Table**

On a computer running the Windows Server 2003 Routing and Remote Access service, several methods for managing the IP routing table are available:

- The **route** command-line tool, which is designed specifically to display and manage the IP routing table.

- The Netsh routing IP context, which includes several commands for managing IP routes. In Windows Server 2003, Netsh contains more commands than it did in Windows 2000 Server.

- The Routing and Remote Access snap-in, which displays the routing table with an additional column that lists the source from which the route was obtained. View the routing table by right-clicking the **IP Routing\Static Routes** node in the console tree and then selecting **Show IP Routing Table**.

In addition to using the **route print** command to view the routing table, you can also use **route** to add, change, or delete routes. Because the IP routing table is stored in RAM and is rebuilt whenever a computer is restarted, one common practice is to use **route add** with the **-p** option to add a static route that is persistent, that is, a route that persists even after the computer is restarted. Persistent routes are stored in the registry.

The IP routing table that appears in the Routing and Remote Access snap-in includes one additional column that the **route print** command does not display. This is the **Protocol** column, which identifies how the router learned each route. For example, **Local** identifies a route on the local node; **Static** identifies a static route that the administrator configured manually; **RIP** or **OSPF** identifies dynamic routes learned from routers configured to use those routing protocols. The following screenshot displays the IP routing table on a Windows Server 2003 Routing and Remote Access router.

RRAS-1 - IP Routing Table

| Destination | Network mask | Gateway | Interface | Metric | Protocol |
|---|---|---|---|---|---|
| 192.168.10.0 | 255.255.255.0 | 192.168.0.1 | 192.168 | 1 | Static (non demand-dial) |
| 192.168.0.0 | 255.255.255.0 | 192.168.0.1 | 192.168 | 2 | OSPF |
| 10.199.17.0 | 255.255.255.0 | 10.199.17.80 | Corpnet | 2 | OSPF |
| 0.0.0.0 | 0.0.0.0 | 10.199.17.1 | Corpnet | 20 | Network management |
| 255.255.255.255 | 255.255.255.255 | 192.168.0.1 | 192.168 | 1 | Local |
| 255.255.255.255 | 255.255.255.255 | 172.16.0.1 | 172.16 | 1 | Local |

If the **Protocol** column lists **RIP** or **OSPF** entries, or any other entries in addition to Local entries, the existence of these entries confirms that the router is receiving routes.

For more information about tools used for unicast IP routing, see Unicast IPv4 Routing Tools and Settings.

**How Static and Dynamic Routing Work**

For information exchange between routers to work efficiently on an IP internetwork, routers must have knowledge of other network IDs or be configured with a default route. On medium-sized or large internetworks, the routing tables must be maintained in an up-to-date state so that network traffic travels along optimal paths. How the routing tables are maintained defines the distinction between static and dynamic routing:

- **Static routing.** An administrator manually enters routing information. Routing information does not change unless the administrator manually updates or deletes it. The term static router refers to a router in an internetwork that uses only static routing. Using static routing exclusively is appropriate only for a small internetwork.

- **Dynamic routing.** An administrator configures a router to automatically generate routing information, share it with other routers, and update it when routing changes occur. The term dynamic router refers to a router in an internetwork that uses only dynamic routing or, more typically, that uses primarily dynamic routing in conjunction with some manually configured static routes. For example, each interface on a dynamic router is typically manually configured.

**Static Routing**

Static routing is based on routing table entries that are manually configured on a static router and that do not change when internetwork topology changes. On a static router, a network administrator with knowledge of the internetwork topology builds and updates the routing table. The routing table on a static router either contains a route for every subnet on the internetwork, or the router is configured to use route summarization and thus the routing table contains only summary routes. Static routers do not exchange routing information with each other.

On an internetwork based on static routing, an IP node is manually configured so that its default gateway address matches the IP address of the local router interface. When a node sends a packet to another network, it forwards the packet to the local router, which checks its own routing table and determines which route to use to forward the packet.

Static routing can work well for small internetworks. Static routing can also be appropriate for connecting a branch office to a main office across a WAN link. However, static routing does not scale well to large or frequently changing internetworks because of the manual administration it requires.

Static routers are not fault tolerant — a static router cannot sense that a neighboring router is no longer available over an operational link or that a remote network segment is no longer reachable if a link goes down. Because a static router is unaware of these changes, it cannot make adjustments to its routing table and requires an administrator to update the router.

A computer running Windows Server 2003 that has two or more network adapters (each configured with an IP address and subnet mask defining routes for local subnets) can function as a static IP router. Typically, to avoid conflicts in case the router is connected to two or more disjoint subnets, only one of the network adapters is configured with a default gateway. For information about disjoint networks, including how to configure an Internet Security and Acceleration (ISA) or virtual private network (VPN) server to prevent disjoint network unreachability, see Default Gateway Behavior for Windows TCP/IP.

On a server running Windows Server 2003 Routing and Remote Access, you can add static routes in three ways:

- In the console tree of the Routing and Remote Access snap-in, by adding a route on the **IP Routing\Static Routes** node
- By using the **route add** command-line utility (with the **-p** switch to add a persistent static route)
- By accessing the **netsh routing ip** context in a command prompt window and then using the **add persistentroute** command

**Static routing entries**

An administrator configuring a static routing entry manually specifies the following information:

- **Destination.** A destination IP address for this route, which can be a host address, a subnet address, a summarized network address, or the destination for the default route (0.0.0.0).
- **Network mask.** The subnet mask for the static route. This number is used in conjunction with the destination IP address. For example, a mask of 255.255.255.255 means that only an exact match of the network mask with the destination can use this route. A mask of 0.0.0.0 means that any destination can use this route. For more information, see "Details of How IP Makes the Routing Decision" earlier in this document.
- **Gateway.** The next-hop IP address for this route. For a LAN interface, the gateway address is the IP address of a neighboring router. For a demand-dial interface, no gateway address is configured.
- **Metric.** The cost (such as the number of hops to the destination, or a unitless metric, such as the speed or reliability of the link) associated with this route. When multiple routes to the same destination exist, IP selects the route with the lowest metric.

For more information about how IP handles entries in the routing table, see How the IPv4 Routing Table Works earlier in this section.

**Using static routing with dynamic routing**

Using static routing exclusively is appropriate only for a small internetwork that is rarely reconfigured. However, a medium-size or large internetwork that primarily uses dynamic routing typically also uses some manually configured static routes. For example, you might use a static route with a higher metric as the redundant backup for a dynamically configured route. In addition, you might use dynamic routing for most paths but configure a few static paths where you want to ensure that network traffic follows a particular route, such as configuring routers to force traffic over a specific path to a high-bandwidth link.

**Autostatic updates across temporary WAN links**

Using dynamic routing protocols across temporary dial-up WAN links is not feasible because dynamic routing protocols typically advertise the contents of the routing table at periodic intervals. Such automatic periodic advertisements are appropriate only for permanently connected LAN or WAN links. If a dial-up WAN link is configured to use dynamic routing, it might incur long-distance charges each time one router calls another router to establish a connection and update the routing table. A more cost-effective way to update the routing table over a dial-up WAN link is to configure autostatic updates. RIP supports autostatic updates; OSPF does not.

An autostatic route is a static route that is dynamically obtained through an exchange of RIP messages. An administrator can enable autostatic updates on an interface that is used for a demand-dial connection so that — whenever the demand-dial link is in a connected state — a request for updated routes can be sent over that interface. Autostatic updates are automatic only in the sense that, when the administrator gives the command to update routes, the router sends an update request across the link to the router on the other side of the connection, and then the router automatically adds the requested routes as static routes to its routing table. Although the term autostatic update might seem to imply an automatic operation, autostatic updates are not automatically performed when the demand-dial connection is made but only if the command to perform the update is given.

As an alternative to issuing a command to initiate a manual autostatic update, if you use Routing and Remote Access RIP routers, you can use the Routing and Remote Access snap-in, Netsh commands, and Task Scheduler to schedule autostatic updates:

1. In the Routing and Remote Access snap-in, enable autostatic update mode by navigating to the **IP Routing\RIP** node in the console tree, and then selecting the **Auto-static update mode** option on the **General** tab of the *InterfaceName*

**Properties** page.

2. Create a batch file or Netsh script file that contains the following commands:

   **netsh interface set interface name**=*DemandDialInterfaceName* **connect=CONNECTED**

   **netsh routing ip rip update** *DemandDialInterfaceName*

   **netsh interface set interface name**=*DemandDialInterfaceName* **connect=DISCONNECTED**

3. Open **Scheduled Tasks** in **Control Panel**, double-click **Add Scheduled Task**, and then follow the wizard instructions to establish the schedule for the batch file or Netsh script to run.

**Static routing in IPv6**

A computer running the Windows Server 2003 family can act as a static IPv6 router that forwards IPv6 packets between interfaces based on the contents of the IPv6 routing table. You can configure static routes by using the **netsh interface ipv6 add route** command.

For information about Windows Server 2003 support for static routing in an IPv6 environment, see "How IPv6 Works" in IPv6 Technical Reference.

**Dynamic Routing**

Dynamic routing is based on routing table entries that are automatically configured and automatically updated by routers whenever the topology of the internetwork changes. Dynamic routing protocols enable this ongoing maintenance of routing tables by using a series of periodic or on-demand messages containing routing information about the internetwork environment that is exchanged among routers. Any dynamic routing protocol enables routers to perform the following basic functions:

- Detect other routers (called *neighbors*) with which it will exchange routing information.
- Store the routing data that it collects from other routers as entries in its routing table.

A router with dynamically configured and updated routing tables is known as a dynamic router. Except for initial configuration, dynamic routers require little ongoing maintenance and can therefore scale to large internetworks. For dynamic routing to work, the same routing protocol must be installed on each router on the internetwork.

Dynamic routing in a multipath IP internetwork is fault tolerant (for a brief description of multipath routing, see "Single-Path vs. Multipath Routing Infrastructures" earlier in this document). Dynamic routes learned from other routers have a finite lifetime: If a router or link fails, dynamic routers detect the change in the internetwork topology when the lifetime of the learned route expires in the routing table. This change then propagates to other routers so that all routers on the internetwork become aware of the new topology.

The ability to scale and the ability to recover from internetwork faults make dynamic routing the better choice for medium, large, and very large internetworks. Routing protocols do, however, increase network traffic on the internetwork. This additional traffic can be an important factor in planning WAN link usage.

Two or more dynamic routing protocols cannot interoperate on the same internetwork because each protocol collects different types of data and uses different techniques to converge after topological changes occur. However, you can configure two or more routing protocols on a single router to connect internetworks that use different routing protocols. For example, you can configure RIP on an interface that connects to a RIP-based internetwork and configure OSPF on another interface that connects to an OSPF-based internetwork. A router with these interfaces can be configured to exchange routing information between the two networks. The exchange of routing information between networks that use different dynamic routing protocols is sometimes referred to as *route redistribution*.

A computer running the Windows Server 2003 Routing and Remote Access service can act as a dynamic router that also supports static routing. An administrator can use the Routing and Remote Access snap-in to add RIP, OSPF, or both, and then enable the protocol on the appropriate interface by adding the interface to the RIP or OSPF protocol displayed in the console tree. On a Routing and Remote Access router, you can use the snap-in to add a static route to an internetwork configured to use dynamic routing by using the **IP Routing\Static Routes** node in the console tree.

**Convergence**

An important element of a routing protocol implementation is its ability to sense and recover from internetwork faults. How

quickly the routing protocol can enable recovery from failures on the internetwork is determined by the type of failure, how it is sensed, and how the routing information is propagated throughout the internetwork.

When all routers on an internetwork have the correct routing information in their routing tables, the internetwork is said to have converged. *Convergence* is the process by which routers update routing tables after a change in network topology occurs — the change is replicated to all routers that need to know about it. When convergence is achieved, the internetwork is in a stable state and all routing occurs along optimal paths.

When a link or router fails, the internetwork must reconfigure itself to reflect the new topology. Information in routing tables must be updated. Until the internetwork reconverges, it is in an unstable state in which routing loops and black holes (both described later) can occur. The time it takes for the internetwork to reconverge is known as the *convergence time*. The convergence time varies based on the routing protocol and the type of failure — downed link or downed router.

For more information about convergence, see "Impact of Routing Loops and Black Holes During Convergence" later in this section, and see "How RIP Works" and "How OSPF Works" in How Unicast IPv4 Routing Protocols and Services Work.

**Routing Protocols Used Between and Within Autonomous Systems**

Autonomous systems (described earlier in "IPv4 Routing Infrastructure") use two types of routing protocols to update routing information:

- The protocols used to distribute routing information between two or more autonomous systems are known as Exterior Gateway Protocols (EGPs).

- The protocols used to distribute routing information within a single autonomous system are known as Interior Gateway Protocols (IGPs).

The following figure depicts two autonomous systems using interior and exterior protocols to communicate.




EGPs are inter–autonomous system routing protocols. EGPs define the way that all routes within the autonomous system are advertised outside of the autonomous system. Routers that connect autonomous systems to the Internet backbone use an EGP to advertise routing information to each other. Advertising can include a list of routes in a flat routing infrastructure or a list of summarized routes in a hierarchical routing infrastructure.

EGPs are independent of the IGPs used within an autonomous system and can enable the exchange of routes between autonomous systems that use different IGPs.

The EGPs for IP internetworks include:

- **Exterior Gateway Protocol (EGP).** An obsolete EGP that was developed to communicate information between autonomous systems on the Internet. Although EGP is the eponym for the generic protocol type for inter-autonomous system routing protocols, it is no longer used on the Internet because it cannot support multipath environments or classless routing. EGP was defined in the now obsolete RFC 904, "Exterior Gateway Protocol."

- **Border Gateway Protocol (BGP).** The EGP that is currently used to communicate information between autonomous systems on the Internet. The Internet is a large IP internetwork divided into several autonomous systems that are connected by the Internet's core routers, which use BGP for communication among themselves. BGP uses autonomous system numbers (ASNs) to avoid routing loops and to implement policy-based routing on the Internet backbone. Unlike its obsolete predecessor EGP, BGP supports complex multipath networks and classless routing. BGP is defined in RFC 1771, "A Border Gateway Protocol 4 (BGP-4)" and RFC 1772, "Application of the Border Gateway Protocol in the Internet" in the IETF RFC Database.

The Windows Server 2003 Routing and Remote Access service does not include EGP or BGP.

IGPs are intra–autonomous system routing protocols. Routers use an IGP to forward routing information to the other routers within the autonomous system. IGPs can distribute routes within the autonomous system in either a flat or hierarchical manner.

The IGPs for IP internetworks include:

- **RIP for IP.** An RFC-based distance vector IGP.
- **OSPF.** An RFC-based link state IGP.
- **IGRP.** A distance vector IGP developed by Cisco Systems, Inc.

The Windows Server 2003 Routing and Remote Access service includes RIP and OSPF, but it does not include a version of IGRP. For detailed information about RIP and OSPF, see the next section, "Common IGP Dynamic Routing Protocols," and see the companion document How Unicast IPv4 Routing Protocols and Services Work.

A third-party software vendor can develop a version of IGRP that works with the Routing and Remote Access service. For information about Microsoft platform software development kits (SDKs) that provide application programming interfaces (APIs) useful for third-party developers who create software that interacts with Windows Server 2003, including the Routing and Remote Access service, see "IPv4 and Routing and Remote Access APIs" in How Unicast IPv4 Routing Protocols and Services Work.

**Common IGP Dynamic Routing Protocols**

The most common types of dynamic routing protocols used within a single autonomous system are:

- Distance vector routing protocols
- Link state routing protocols

A distance vector routing protocol, such as RIP v2, is more appropriate for a relatively small, simple network that is not expected to grow rapidly. For a large, complex internetwork, a link state routing protocol, such as OSPF, is more appropriate. You must use RIP v2 or OSPF to support VLSM and route summarization. Although the outdated RIP v1 is still widely used in private networks, it does not support either VLSM or route summarization and thus is not well suited for enterprise networks.

| **Note** |
| --- |
| The Windows Server 2003 Routing and Remote Access service does not support the dynamic routing protocols RIPng (the version of RIP for IPv6) or OSPF for IPv6. However, a computer running Windows Server 2003 can act as a static IPv6 router, and it is possible to route IPv6 traffic over an IPv4 internetwork using an IPv6 transition technology. For information about Windows Server 2003 support for static IP routing in an IPv6 environment, see "How IPv6 Works" in IPv6 Technical Reference. |

Understanding how distance vector and link state routing protocols work is essential to choosing the type of dynamic routing that best suits your network needs. The primary differences between distance vector and link state routing protocols include:

- What routing information is exchanged
- How the information is exchanged
- How quickly the internetwork can recover from a downed link or a downed router

**Distance Vector Dynamic Routing Protocols**

A router that uses distance vector routing protocols, the earliest type of dynamic routing protocol, advertises the routes in its routing table to other routers at regular intervals. The routing information that is periodically exchanged between routers by using distance vector routing protocols is typically unsynchronized and unacknowledged.

A distance vector routing protocol advertises the number of hops to a network destination (the distance) and the direction in which a packet can reach a network destination (the vector). The distance vector algorithm, also known as the *Bellman-Ford*

*algorithm*, enables a router to pass route updates to its neighbors at set intervals. Each neighbor then adds its own distance value and forwards the routing information on to its immediate neighbors. The result of this process is a table containing the cumulative distance to each network destination.

The following table lists several common distance vector routing protocols based on three major routable network protocols.

**Common Distance Vector Dynamic Routing Protocols**

| Routable Network Protocol | Distance Vector Dynamic Routing Protocol |
|---|---|
| IP | Routing Information Protocol (RIP) for IP versions 1 and 2 (RIP v1 and RIP v2)<br><br>Interior Gateway Routing Protocol (IGRP) |
| IPX | Routing Information Protocol (RIP) for IPX |
| AppleTalk | Routing Table Maintenance Protocol (RTMP) |

The Windows Server 2003 Routing and Remote Access service includes RIP v1 and RIP v2 for IP and RTMP for AppleTalk. Windows Server 2003 Routing and Remote Access does not include IGRP or RIP for IPX.

The advantages of distance vector dynamic routing protocols include:

- **Simpler.** Distance vector routing protocols use simple router advertisement processes that are easy to understand.
- **Easy to configure.** In its most basic form, configuring a distance vector routing protocol is as easy as enabling the protocol on the router interfaces.

The disadvantages of distance vector dynamic routing protocols include:

- **Large routing tables.** Multiple routes to a specific network ID appear as multiple entries in the routing table. In a large internetwork with multiple paths, the IP routing table can have hundreds or thousands of entries.

  This disadvantage is eliminated in some RIP router implementations, including Windows Server 2003, which store only one route — the route with the lowest metric — for any subnet in the routing table.
- **High network traffic overhead.** Route advertising is done at periodic intervals, even (unnecessarily) after the internetwork has converged.
- **Does not scale.** Because of the size of the routing table and the high network traffic overhead, distance vector routing protocols do not scale well to large or very large internetworks.
- **High convergence time.** Because of the unsynchronized and unacknowledged way that distance vector information is exchanged, convergence of the internetwork can take several minutes. During convergence, routing loops and black holes can occur and result in lost or undeliverable data.

For more information about the RIP distance vector routing protocol, see the section "How RIP Works" in How Unicast IPv4 Routing Protocols and Services Work.

**Link State Dynamic Routing Protocols**

A router that uses a link state dynamic routing protocol, such as OSPF, stores information about the state of each link between itself and other routers. Link state routing protocols address some of the limitations of distance vector routing protocols. The most important of these improvements is that link state dynamic routing protocols enable faster convergence than distance vector routing protocols.

Unlike distance vector routing protocols, which broadcast updates to all routers at regularly scheduled intervals, link state routing protocols provide updates only when a network link changes state. When such an event occurs, a notification called a *link state advertisement (LSA)* is sent throughout the network to update routing tables. LSAs, which include the router network IDs, are sent at startup and again any time that changes in the internetwork topology are detected. Link state updates are sent by using unicast or multicast traffic rather than by broadcasting. Link state routers build a database of LSAs and use the database to calculate the routing table. Routing information exchanged between link state dynamic routers is synchronized and acknowledged.

Link state routing protocols are more reliable and require less bandwidth than do distance vector routing protocols. However,

link state routing protocols are also more complex, more memory-intensive, and place a greater load on the CPU of a router.

The following table lists two major link state routing protocols used on internetworks based on two major routable network protocols.

**Common Link State Dynamic Routing Protocols**

| Routable Network Protocol | Link State Dynamic Routing Protocol |
|---|---|
| IP | Open Shortest Path First (OSPF) |
| IPX | NetWare Link Services Protocol (NLSP) |

The Windows Server 2003 Routing and Remote Access service includes OSPF; it does not include NLSP.

The advantages of link state dynamic routing protocols include:

- **Smaller routing tables.** Only a single optimal route for each network ID is stored in the routing table.
- **Low network traffic overhead.** Routers that use a link state dynamic routing protocol exchange less routing information after the internetwork has converged than do distance vector dynamic routing protocols. For example, after convergence OSPF sends only Hello packets at regular intervals. Unlike distance vector routing protocols, which broadcast updates to all routers at regularly scheduled intervals, link state routing protocols provide updates (in the form of an LSA) only when a network link changes state.
- **Ability to scale.** Because of smaller routing tables and low network traffic overhead, link state routing protocols scale well to large and very large internetworks.
- **Lower convergence time.** Link state routing protocols have a much lower convergence time than distance vector routing protocols. An internetwork that uses link state routing protocols converges without producing routing loops.

The disadvantages of link state dynamic routing protocols include:

- **Complexity.** Link state routing protocols are more complex and difficult to understand than distance vector routing protocols.
- **More time needed to configure.** A link state routing protocol implementation requires additional planning and configuration.
- **Resource intensive.** For very large internetworks, the database of LSAs and the calculation of routing table entries can be memory and processor intensive.

For more information about the OSPF link state dynamic routing protocol, see the section "How OSPF Works" in How Unicast IPv4 Routing Protocols and Services Work.

**Impact of Routing Loops and Black Holes During Convergence**

Routing problems can occur when either a host routing table or a router routing table contains information that does not reflect the correct topology of the internetwork. On an internetwork that uses dynamic routing, the routing internetwork must reconfigure itself to reflect the new topology whenever a link or router fails. The internetwork is in an unstable state during the period of time it takes for convergence to take place — during this period of time, routing loops and black holes can occur.

**Routing loops**

During the routing process, packets are forwarded by the sending host and then by one or more routers in the optimal direction as determined by information in each node's local routing table. If the routing table entries on all routers are correct, a unicast packet takes the best path from the source node to the destination node. However, if any routing table entry is incorrect, either through configuration errors or through learned routes that do not accurately reflect the topology of the internetwork, routing loops can occur.

A *routing loop* is a path through the internetwork that runs in a circle instead of reaching the intended destination. A routing loop occurs when routers forward traffic to each other in a loop that does not include the network segment of the destination.

The following figure illustrates an example of a routing loop that occurs when inaccurate information is stored in local routing tables.




In this example:

- According to the routing table on Router 1, the optimal route to Network 10 is through Router 2.
- According to the routing table on Router 2, the optimal route to Network 10 is through Router 3.
- According to the routing table on Router 3, the optimal route to Network 10 is through Router 1.

The result would be an infinite loop — except that routable protocols use a counter in the network layer header of the packet to prevent the packet from perpetually looping. Each time a router passes the packet from one network segment to another, the router either increases or decreases the counter. If the count reaches its maximum value (when increasing) or reaches 0 (when decreasing), the packet is discarded by the router. For example, when an IP node sends an IP packet, IP sets a maximum link count in the Time to Live (TTL) field in the IP header. Each IP router crossed by the packet decreases the TTL value in the packet by one. When the TTL value is 0, the IP router discards the packet and sends an ICMP Time Exceeded-TTL Exceeded in Transit message back to the sending node from which the packet originated.

One common type of routing loop can be avoided by ensuring that neighboring routers are not configured with default routes that point to each other: that is, on each router, the default route (0.0.0.0, 0.0.0.0) must not have the **Default gateway** field configured with the IP address of the other router. Because a default route passes all traffic that is not on a directly connected network to the configured router, routers that have default routes pointing to each other can produce routing loops for traffic with an unreachable destination.

By typing **ping -i 255** *IP_address* at a command prompt, you can initiate the detection of a possible routing loop. The **-i** parameter sets the TTL value in the ICMP Echo message. If the command results in the message **TTL Expired In Transit**, this indicates a possible routing loop. By typing **tracert** *IP_address* at a command prompt, you can then confirm that a routing loop exists if the output displays a set of repeating router IP addresses.

**Black holes**

The IP protocol is a connectionless, datagram-based protocol, which, by definition, does not guarantee a successful delivery. IP attempts a best effort, unacknowledged delivery to the next hop or to the final destination, which can lead to conditions on the internetwork in which data is lost.

If a downstream router fails and the failure is not detected by an upstream router, the upstream router continues to forward packets to the failed router. Because the failed downstream router does not receive the packets, the packets forwarded by the upstream router are dropped from the internetwork. The upstream router is said to be sending packets to a *black hole*, which is defined as a case in which packets on an internetwork are lost without any indication of an error.

In the following figure, Router 1 is not informed that Router 2 has failed and, therefore, continues to forward packets to Router 2. Router 2 is now a black hole because the packets are dropped and no error message is returned to Router 1.




A black hole can form when a link or router fails, and the failure is not yet detected. In a dynamic routing environment, routers detect failed links or failed routers when the lifetime of learned routes expires in their routing tables. In a static routing environment, a black hole persists until the functionality of the link or router is restored or until the static routers are reconfigured by the network administrator.

Black holes can also occur when an active router discards packets without indicating why the packets are discarded. For example, a Path Maximum Transmission Unit (PMTU) black hole router discards IP packets that need to be fragmented without returning a message to the sender indicating the error. A PMTU black hole router silently discards IP packets that require

fragmentation if the Don't Fragment (DF) flag in the IP header of the received packet is set to 1. PMTU black hole routers can be difficult to detect because packets of smaller sizes are forwarded.

**Link state protocols more robust than distance vector protocols**
One of the disadvantages of distance vector dynamic routing protocols, such as RIP, is that, on a large internetwork, several minutes might be required for convergence to occur. This delay can allow routing loops or black holes to occur during convergence.

Link state dynamic routing protocols provide more robust protection against black holes and routing loops. For example, OSPF prevents routing loops due to unsynchronized databases by not advertising links until the routers at each end of a link have synchronized their databases. OSPF also reduces the occurrence of black holes caused by one-way links by not including links in its database until the links are known to be bidirectional.

**Ranking Route Sources**
On a network that uses multiple sources of routing information, route sources are ranked in order of preference. If two or more route sources provide a route to the same destination, the route learned by the route source with a lower rank number is the preferred route. One reason that route sources must be ranked — either by accepting the default ranks or by reordering them — is because metric definitions differ for different route sources, including the RIP and OSPF routing protocols. Rather than trying to reconcile these dissimilar metrics, the router uses the route learned from the most preferred route source.

For example, if you configure a Routing and Remote Access router to use both RIP and OSPF, the Routing and Remote Access service adds both RIP-learned routes and OSPF-learned routes to the Route Table Manager (RTM) IP routing table. If you specify that the metric of an OSPF-learned route is 5 and that the metric of the corresponding RIP-learned route is 3, and if OSPF is the preferred routing protocol, RTM adds only the OSPF route to the IP forwarding table.

| ⚠**Caution** |
| --- |
| If you enable both RIP and OSPF on the same router, make sure that you configure each routing protocol on a separate interface, each connecting to a different network. |

If you use Windows Server 2003 Routing and Remote Access routers, you can modify the default preference levels for the following route sources in the Routing and Remote Access snap-in by navigating to the **IP Routing\General** node in the console tree and selecting the **Preference Levels** tab on the **General Properties** page:

- Local.
- Static, which refers to a static route assigned to a demand-dial interface that is added by using the Routing and Remote Access snap-in.
- Static (non demand-dial), which refers to a static route assigned to a LAN interface that is added by using the Routing and Remote Access snap-in.
- Auto-static.
- Network Management, such as SNMP, which refers to a static route assigned to a LAN interface that is added by using the **route add** command-line tool. You cannot use **route add** to add a route for a demand-dial interface.
- OSPF.
- RIP v2.

Alternatively, you can specify the relative preference level of any of these route sources by using the **netsh routing ip set preferenceforprotocol** command.

**Related Information**
The following resources contain additional information that is relevant to this section.

- For more information about each of the RFCs cited in this document, see the IETF RFC Database.
- What Is Unicast IPv4 Routing?

- How Unicast IPv4 Routing Protocols and Services Work.
- Unicast IPv4 Routing Tools and Settings.
- TCP/IP Technical Reference.
- IPv6 Technical Reference.
- IPv4 Multicasting Technical Reference.
- "Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference" by Joseph Davies and Thomas Lee, Redmond, WA, Microsoft Press, 2003.

# Unicast IPv4 Routing Tools and Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Unicast IPv4 Routing Tools
- Unicast IPv4 Routing Registry Entries
- Related Information


You can use the unicast IPv4 routing tools and registry settings described here to enable, configure, and manage the unicast IPv4 routing features provided by the Microsoft® Windows® Server™ 2003 Routing and Remote Access service.

**Unicast IPv4 Routing Tools**

The sets of tools listed in the following table are associated with unicast IPv4 routing as supported by the Windows Server 2003 Routing and Remote Access service.

**Available Tools for Unicast IPv4 Routing**

| Category | Tools in This Set |
|---|---|
| Command-line tools used for unicast IPv4 routing | • Netsh.exe, including:<br>  • Netsh ras commands for file tracing<br>  • Netsh routing IP commands for IP packet filtering and for configuring interfaces, persistent and non-persistent static routes, and preference levels for routing protocols (route sources)<br>  • Netsh routing IP OSPF commands<br>  • Netsh routing IP relay commands (for the DHCP relay agent)<br>  • Netsh IP RIP commands<br>  • Netsh routing IP routerdiscovery commands (for ICMP router discovery)<br>• Pathping.exe<br>• Ping.exe<br>• Route.exe<br>• Tracert.exe |

| | |
|---|---|
| Graphical user interface tools used for unicast IPv4 routing | • Network Connections (default gateway options only)<br>• Rrasmgt.msc — Routing and Remote Access snap-in, including:<br>  • Options for enabling, disabling, or refreshing the Routing and Remote Access service<br>  • General options (event logging and route source preference options)<br>  • RIP for IP options<br>  • OSPF options<br>  • DHCP Relay Agent options<br>  • IP packet filtering options<br>  • ICMP router discovery options |

Each of these tools is described in the following sections.

You can specify commands for the command-line tools for one or more remote servers by establishing a remote access connection to each server from a workstation running Microsoft Windows XP or from another server, and then opening a command prompt window in the remote access window for each server. You can also specify Network Connections and Routing and Remote Access snap-in options by establishing a remote access connection. In addition, you can use the Remote Access snap-in to manage multiple servers running the Routing and Remote Access service by accessing the node for each server.

**Netsh.exe: Netsh Commands for Managing IPv4 Routing**

**Category** The Netsh commands for IPv4 routing are a subset of the Netsh command-line tools that are included with the Windows Server 2003 operating system. In many cases, a Netsh command can be used to manage the Routing and Remote Access service as an alternative to using an option available in the Routing and Remote Access snap-in. Netsh also includes commands for IPv6.

**Version compatibility** The Netsh commands for IP routing are compatible with the Windows Server 2003 operating system. Netsh commands were first introduced in Microsoft Windows 2000 Server and were expanded to include additional commands in Windows Server 2003.

The Netsh commands help network administrators manage a TCP/IP network. You can use Netsh locally or remotely to display or modify the configuration of services or protocols on Windows-based computers. The Netsh command-line interface is scriptable, which lets you perform batch configurations or remote network administration from a single location.

Netsh provides sets of commands (also known as contexts) for performing a wide range of network configuration tasks. In addition to a few commands from the Netsh ras and Netsh interface contexts, this document describes several Netsh Routing IP contexts used in conjunction with the unicast IPv4 routing features of the Routing and Remote Access service.

For more information about Netsh, see Command-Line Reference for Windows Server 2003, Standard Edition. In particular, for information about IP routing-related commands not included in this technical reference, see the following Netsh contexts:

- **Netsh ras context,** for remote access configuration.
- **Netsh interface context,** for configuration of demand-dial interfaces.
- **Netsh routing IP NAT context,** for network address translation (NAT).
- **Netsh interface IPv6 context,** for manually configuring IPv6 settings.
- **Netsh routing IP IGMP context,** for multicast routing.

Each of the listed Netsh contexts contains unicast routing–related commands not included in this technical reference, except for the IGMP context, which contains multicast routing–related commands.

**How to access Netsh commands**

You can access a particular command in any Netsh context in one of the following two ways:

**Type the entire path for the command at the command prompt.** For example, to view IPv4 router discovery configuration data, type the following command at a command prompt:

**C:\>netsh routing ip routerdiscovery show interface name="Local Area Connection"**

The preceding command produces the following output:

```
Router Discovery Information for Local Area Connection 3

---------------------------------------------------------------------

Router Discovery is: Disabled

Minimum Advertisement Interval: 7 minutes

Maximum Advertisement Interval: 10 minutes

Advertisement Lifetime: 30 minutes

Preference Level: 0
```

**Access the context containing the command that you want to use, and then type only the specific command.** The following series of commands displays the same output as the preceding example:

**C:\>netsh**

**netsh>routing**

**netsh routing>ip**

**netsh routing ip>routerdiscovery**

**netsh routing ip routerdiscovery>show interface name="Local Area Connection"**

The preceding series of commands produce the following output:

```
Router Discovery Information for Local Area Connection 3

---------------------------------------------------------------------

Router Discovery is: Disabled

Minimum Advertisement Interval: 7 minutes

Maximum Advertisement Interval: 10 minutes

Advertisement Lifetime: 30 minutes

Preference Level: 0
```

The following sections contain a brief description of unicast IPv4 routing commands that are related to the Routing and Remote Access service and that are available for the following Netsh contexts:

- Netsh ras context (includes only commands for configuring file tracing)
- Netsh routing IP context (includes only commands for configuring IP packet filtering, interfaces, persistent and non-persistent static routes, and preference levels for route sources)
- Netsh routing IP OSPF context
- Netsh routing IP relay context
- Netsh routing IP RIP context
- Netsh routing IP routerdiscovery context

**Netsh ras context**

The Routing and Remote Access service provides extensive tracing capability to help troubleshoot complex network problems. Separate routing components can be independently enabled so that tracing information can be logged to files (known as *file tracing*). Logged data includes information about internal component variables, function calls, and interactions. You can enable and disable tracing for routing protocols while the router is running.

| **Note** |
| --- |
| You can also enable file tracing for Routing and Remote Access components, such as the RIP and OSPF routing protocols, by editing the registry (see Unicast IPv4 Routing Registry Entries later in this document). However, using Netsh commands is the recommended method. |

Tracing consumes system resources and must be used sparingly. After you obtain the trace or identify the problem, immediately disable tracing. In particular, do not leave tracing enabled on multiprocessor computers.

The Netsh ras context includes primarily remote access–related commands, which are not included here. The following table lists only the file tracing commands for the RIP v2 and OSPF routing protocols that are available in the Netsh ras context.

**Unicast IPv4 Routing Commands in the Netsh Ras Context**

| Command | Task |
| --- | --- |
| Set | You can use the following commands:<br>• **Set tracing** *component* **enable**. Enables file tracing for the specified component. For example, **set tracing iprip2 enable** or **set tracing ospf enable**.<br>• **Set tracing** *component* **disable**. Disables file tracing for the specified component. For example, **set tracing iprip2 disable** or **set tracing ospf disable**.<br>• **Set tracing * enable**. Enables file tracing on all components. |
| Show | You can use the following commands:<br>• **Show tracing.** Lists components for which it is possible to enable file tracing and shows whether file tracing for each one is currently enabled or disabled.<br>• **Show tracing** *component*. Shows whether file tracing for the specific component is enabled. For example, **show tracing iprip2** or **show tracing ospf**. |

Most of the commands available in the Netsh ras context are used for remote access networking, such as for demand-dial or virtual private network (VPN) connections. These commands are also unicast IPv4 commands, but they are not covered in this technical reference. For more information about tools for remote access networking, see Demand Dial Tools and Settings and VPN Tools and Settings.

**Netsh routing IP context**

The partial set of commands included in this document for this context are the Netsh commands used to configure IP packet filtering, which can help increase security on a private IP internetwork connected to the Internet by specifying which IP packet traffic is accepted or rejected; the commands used to configure interfaces, persistent and non-persistent static routes; and preference levels for route sources. Also included, in separate sections later in this document, are several relevant contexts accessible under the Netsh routing IP context: the OSPF, RIP, relay (for the DHCP relay agent), and IP router discovery contexts.

Excluded from this technical reference are Netsh commands used for IPv4 multicast routing, network address translation (NAT), DHCP, DNS, and remote access.

The following table describes commands available in the Netsh routing IP context that are relevant to the unicast IPv4 routing services provided by the Routing and Remote Access service.

**Unicast IPv4 Routing Commands in the Netsh Routing IP Context**

| Command | Task |
| --- | --- |
| Add | • **Add filter**. Adds an IP packet filter to the specified interface and configures one or more of the following parameters: |

- The name of the interface on which to add the filter.
- The type of filter (input filter, output filter, or a filter for a dial-up interface).
- The source IP address and subnet mask of the packet to be filtered.
- The destination IP address and subnet mask of the packet to be filtered.
- The protocol type (TCP, TCP-EST, or UDP) and the source and destination ports for the packet to be filtered.
- TCP-EST refers to TCP [established] and is used, for example, in a site-to-site VPN or dial-up connection when a calling router connects to an answering router. If you specify **TCP [established]**, traffic is accepted only if the calling router has already initiated the TCP connection.
- The ICMP type and ICMP code of the packet to be filtered.

- **Add interface**. Enables IP forwarding on an interface as it adds the interface (**add interface name="***Interface Name***" state=enable**) or disables IP forwarding on an interface as it adds the interface (**add interface name= "***Interface Name***" state=disable**).

- **Add persistentroute**. Adds a persistent static route to the specified interface, including:
  - The destination IP address and subnet mask for the route.
  - The name of the interface on which the route is to be added.
  - The next hop (gateway) IP address for the route (omit for point-to-point interfaces).
  - The connection type for this route, that is, **static** (which triggers a demand-dial connection) or **nondod** (which does not trigger a demand-dial connection).
  - The preference for this route, and its metric. The metric is used to break ties among routes with equal preference.
  - The view (unicast, multicast, or both) for this route.

- **Add preferenceforprotocol**. Sets the preference level for any of the following route sources:
  - Autostatic
  - Local
  - NetMgmt (network management, such as SNMP)
  - Nondod (a static route assigned to a LAN interface)
  - OSPF
  - RIP
  - Static (a static route assigned to a demand-dial interface)

**Note**

| | |
|---|---|
| | If two routes exist to the same destination, the route learned by the route source with a higher preference is preferred over the route learned by the route source with a lower preference regardless of the route metric. For example, if the preference level for OSPF is higher than the preference level for RIP v2, an OSPF route with a metric of 5 is entered into the routing table rather than the same route discovered by RIP with a metric of 3.<br><br>• **Add rtmroute**. Adds a non-persistent route used for network management to the specified interface name, including:<br>  ◦ The destination IPaddress and subnet mask for the route.<br>  ◦ The name of the interface on which the route is to be added.<br>  ◦ The next hop (gateway) IP address for the route (omit for point-to-point interfaces).<br>  ◦ The preference for this route, and its metric. The metric is used to break ties among routes with equal preference).<br>  ◦ The view (unicast, multicast, or both) for this route. |
| Delete | • **Delete filter**. Deletes an IP packet filter from the specified interface.<br>• **Delete interface**. Deletes the specified interface, which also removes IP forwarding on that interface.<br>• **Delete persistentroute**. Deletes a persistent static route from the specified interface.<br>• **Delete preferenceforprotocol**. Deletes the preference level for a specified routing protocol type (also known as *route source*).<br>• **Delete rtmroute**. Deletes a non-persistent route used for network management when either the interface name or the next hop count is specified. |
| Dump | Creates a script that contains the current configuration for this server. If you save the script to a file, you can use it to restore altered configuration settings or to configure a new server.<br><br>"Current configuration" refers to all configuration settings in and below the specific node in the Routing and Remote Access snap-in at which the command "dump" is given. Because this table refers to the Netsh routing IP context, in this case, the dump command includes all configuration settings in the **IP Routing node** and any configured child nodes, including the **General, Static Routes**, **DHCP Relay Agent**, **OSPF**, **NAT/Basic Firewall**, and **IGMP** nodes.<br><br>If you plan to make extensive changes to the configuration of this server, a recommended practice is to begin the configuration session by using the **dump** command before you start, in case you need to restore the configuration later. |
| Set | • **Set filter**. Changes the default action (drop or forward) for a filter type (input, output, or dial-up) or changes the fragmentation check setting (enable or disable) on the specified interface.<br>• **Set interface**. Enables or disables the specified interface.<br>• **Set loglevel**. Sets the global logging level to one of the following: none; errors only; errors and warnings; or errors, warnings, and informational messages.<br>• **Set persistentroute**. Modifies a persistent static route on the specified interface by changing one or more of the following: |

- The destination IP address and subnet mask for the specified route.
- The name of the interface on which the route is to be modified.
- The next hop for the route (omit for point-to-point interfaces).
- The connection type for this route that is **static** (which triggers a demand-dial connection) or **nondod** (which does not trigger a demand-dial connection).
- The preference for the route.
- The metric for the route. The metric is used to break ties among routes with equal preference.
- Whether the route is valid only for unicast traffic, only for multicast traffic, or for both types of traffic.

- **Set preferenceforprotocol**. Sets the preference level for any of the following protocol types (also known as *route sources*):
  - Autostatic
  - Local
  - NetMgmt (network management, such as SNMP)
  - Nondod (a static route assigned to a LAN interface)
  - OSPF
  - RIP
  - Static (a static route assigned to a demand-dial interface)
- **Set rtmroute**. Modifies a non-persistent route used for network management on the specified interface:
  - The destination IP address and subnet mask for the route.
  - The interface on which the route is to be added.
  - The next hop (gateway) IP address for the route (omit for point-to-point interfaces).
  - The preference for this route.
  - The metric for this route. The metric is used to break ties among routes with equal preference.
  - The view (unicast, multicast, or both) for this route.

| | |
|---|---|
| Show | • **Show filter**. For all interfaces or for the specified interface, indicates whether input, output, or demand-dial IP packet filters are configured and shows whether fragmentation checking is enabled. If a filter exists for a given interface, displays filter information for that interface.<br><br>• **Show interface**. For all interfaces or for the specified interface, displays information including interface type, vendor, routing protocols running on a given interface, and whether or not an interface is enabled.<br><br>• **Show loglevel**. Displays the global logging level, which is one of the following: none; errors only; errors and warnings; or errors, warnings, and informational messages.<br><br>• **Show persistentroutes**. For all interfaces or for the specified interface, displays persistent static routes, including:<br><br>  ◦ Each route and its address prefix . The address prefix is shown in classless interdomain routing (CIDR) notation, for example, 10.0.0.1/8.<br><br>  ◦ The protocol (route source) that added this route.<br><br>  ◦ The preference for this route, and its metric.<br><br>  ◦ The next hop (gateway) IP address for the route, and the interface that this route uses to forward packets.<br><br>  ◦ The view (unicast, multicast, or both) for this route.<br><br>• **Show preferenceforprotocol**. Lists the current preference level (relative priority) for each IP routing protocol type (each type of route source).<br><br>• **Show protocol**. Displays all configured IP protocols, including unicast IP protocols.<br><br>• **Show rtmdestinations** [1]. Displays the IP forwarding table. IP uses the forwarding table to forward packets.<br><br>• **Show rtmroutes** [1]. Displays the IP routing table. IP uses the routing table to store all routes from all possible sources. |

1. The Routing and Remote Access service on a router running Windows Server 2003 maintains the IP routing table by using a component called the Route Table Manager, abbreviated as "rtm" in the commands **rtmroutes** and **rtmdestinations**. The Route Table Manager updates the IP forwarding table based on incoming route information from multiple sources.

**Netsh routing IP OSPF context**

Open Shortest Path First (OSPF), the most common routing protocol used on large and very large IPv4 internetworks, enables the dynamic exchange of IPv4 routing information between OSPF routers across complex internetworks. Typically, OSPF is deployed on private IP internetworks connected to the public Internet, including internetworks that span multiple geographical regions, nations, or continents.

OSPF is structured as an autonomous system made up of one or more OSPF areas. Each area is made up of a logical collection of networks, routers, and links, and each area includes special OSPF routers that enable communication among the constituent areas. OSPF routers in each area dynamically build and synchronize a database of the OSPF network topology, called a *link state database* (LSDB). The LSDB is a topological map of all routers in the same OSPF area, and OSPF routers can update the LSDB in response to topology changes.

The following table describes the commands available in the Netsh routing IP OSPF context. The OSPF routing protocol is provided by the Windows Server 2003 Routing and Remote Access service.

**Commands in the Netsh Routing IP OSPF Context**

| Command | Task |
|---|---|

| | |
|---|---|
| Add | • **Add area**. Adds an OSPF area. Specify the area ID, which is a 32-bit number expressed in dotted decimal notation. The only required OSPF area is the OSPF backbone area, which requires the reserved area ID of 0.0.0.0. An OSPF backbone area is either the sole area in the OSPF autonomous system, or it is the central area to which all other OSPF areas in this OSPF autonomous system connect and through which packets are routed from one non-backbone area to another non-backbone area.<br><br>Although an OSPF area ID is neither an IP address nor an IP network ID, if this is not the backbone area and if the entire area can be summarized with a single network ID, you can set the area ID to match the single network ID. For example, if an area contains all of the subnets for the IP network 10.1.0.0/16, you can set the area ID to 10.1.0.0.<br><br>• **Add interface**. Enables OSPF on the specified interface.<br><br>• **Add neighbor**. Adds an OSPF neighbor router on a non-broadcast multiple access (NBMA) network. Specify the IP address of the OSPF router at the opposite end of the NBMA link from this router, and then specify the router priority of the neighbor router.<br><br>• **Add protofilter**. Adds an OSPF protocol filter (autostatic, local, RIP, SNMP, non-demand-dial, or static).<br><br>• **Add range**. Adds the network ID and subnet mask pair that define a range of addresses for the specified area ID. This {network ID, subnet mask} pair is known as a *summary route* and is used to summarize routes within this OSPF area.<br><br>• **Add routefilter**. Adds an OSPF route filter. Specify the IP address and subnet mask pair of the OSPF route to be filtered.<br><br>• **Add virtif**. Adds a virtual interface (also called a *virtual link*). Specify the OSPF transit area ID and the OSPF router ID of the virtual OSPF neighbor router. A virtual link is a logical connection that enables the exchange of routing information between a backbone area border router (ABR) and an ABR that is not physically connected to the backbone area across a transit area. You cannot configure a virtual link across the backbone area. |
| Delete | • **Delete area**. Deletes the specified OSPF area by deleting its area ID.<br><br>• **Delete interface**. Disables OSPF on the specified interface.<br><br>• **Delete neighbor**. Deletes the specified OSPF neighbor router on an NBMA network.<br><br>• **Delete protofilter**. Deletes the specified OSPF protocol filter (autostatic, local, RIP, SNMP, non-demand-dial, or static).<br><br>• **Delete range**. Deletes the IP address and subnet mask pair that define a range of addresses (a summary route) for the specified OSPF area ID.<br><br>• **Delete routefilter**. Deletes the specified OSPF route filter.<br><br>• **Delete virtif**. Deletes the specified virtual interface on this OSPF router. |
| Dump | Creates a script that contains the current configuration for OSPF. If you save the script to a file, you can use it to restore altered configuration settings or to configure a new OSPF router.<br><br>If you plan to make extensive changes to the configuration of an OSPF router, a recommended practice is to begin the configuration session by using the **dump** command before you start, in case you need to restore the configuration later. |

| | |
|---|---|
| Install | Installs the OSPF routing protocol under the IP routable protocol. |
| Set | • **Set area**. For the specified OSPF area, sets the following:<br>  ◦ Whether plaintext passwords are required. Plaintext passwords are used by routers to identify each other, not for security. Passwords are enabled by default, and all interfaces in the same area must use the same password.<br>  ◦ Whether this area is a stub area, and, if so, the cost of the summary default route that an OSPF router in this area uses to reach all destinations external to its OSPF autonomous system.<br>  ◦ Whether the OSPF area imports summary advertisements.<br>• **Set global**. For the specified OSPF router, sets the following:<br>  ◦ Whether the router acts as an autonomous system boundary router (ASBR).<br>  ◦ Whether the amount of information logged is one of the following: none; errors only; errors and warnings; or errors, warnings, and informational messages.<br>• **Set interface**. For the specified OSPF interface, sets the following:<br>  ◦ The name, IP address, and subnet mask of the interface; whether the interface is enabled or disabled; and the OSPF area ID.<br>  ◦ Whether the interface is on an NBMA, point-to-point, or broadcast network, and the interface priority.<br>  ◦ The estimated number of seconds to transmit a link state update packet over this interface; the interval in seconds between link state advertisement retransmissions for adjacencies; the interval in seconds between transmissions of hello packets; the number of seconds before a neighboring router considers this router to be down; and (for an NBMA interface only) the interval in seconds between OSPF network polls.<br>  ◦ The cost of sending a packet from this interface.<br>  ◦ The password (if simple password authentication is enabled for this OSPF area). Passwords are enabled by default, and all interfaces in the same area must use the same password.<br>  ◦ The maximum size, in bytes, of IP packets carrying OSPF information that can be sent without fragmentation.<br>• **Set protofilter**. For the OSPF protocol filter (autostatic, local, RIP, SNMP, non-demand-dial, or static), specifies whether the action for IP packets is **drop** or **accept**.<br>• **Set routefilter**. For the OSPF route filter, specifies whether IP packets are accepted or dropped.<br>• **Set virtif**. For the specified OSPF virtual interface (virtual link), sets the following:<br>  ◦ The IP address of the transit area, and the OSPF router ID of the virtual neighbor router.<br>  ◦ The estimated number of seconds to transmit a link state update packet over this interface; the interval in seconds between link state advertisement retransmissions for adjacencies; the interval in seconds between transmissions of hello packets; and the number of seconds before a neighboring router considers this router to be down. |

| | |
|---|---|
| | • The password (if simple password authentication is enabled). |
| Show | • **Show area**. Shows configuration information about this OSPF area, including:<br>  • Whether the area is created and enabled as an OSPF area, or created but not enabled.<br>  • Authentication type (such as simple password authentication), if any. Plaintext passwords for this type of authentication are used by routers to identify each other, not for security.<br>  • Whether the area is a stub area, and if so, its metric.<br>  • Whether the area imports summary advertisements.<br>• **Show areastats**. For each OSPF area, or for the specified area, shows the state (**Up** or **Down**) and the shortest path first (SPF) count. The SPF count is the number of least-cost paths that this router has calculated from itself to each other router and subnet in this OSPF autonomous system (if only one OSPF area exists) or to each other router in the local area (if multiple OSPF areas exist).<br>• **Show global**. Shows global configuration information for this OSPF router:<br>  • Whether the router is created and enabled as an OSPF router, or created but not enabled.<br>  • The OSPF router ID.<br>  • Whether the router is an area border router (ABR).<br>  • The specified level of logging for this router.<br>• **Show interface**. Shows OSPF configuration information for the specified OSPF interface, including:<br>  • Whether the interface is created and enabled as an OSPF interface, or created but not enabled.<br>  • The IP address, subnet mask, area ID, interface type, and router priority of the OSPF interface.<br>  • The current values for transit delay, hello interval, dead interval, poll interval, metric, and maximum transmission unit (MTU) size.<br>• **Show lsdb**. Displays information about the OSPF link state database (LSDB):<br>  • **Area**. The identification number of the OSPF area for each link state advertisement (LSA) in the table.<br>  • **Type**. The type of advertisement, including Router, Net-Summary, Stub, Network, ASBR-Summary, or AS External.<br>  • **Link State ID**. The identification number of each advertisement shown in the output. The advertisements vary by advertisement type.<br>  • **Adv Router**. The OSPF router ID of the advertising router.<br>  • **Age**. The age (in seconds) since the advertisement originated.<br>  • **Sequence**. A sequence number that identifies old or duplicate advertisements.<br>• **Show neighbor**. Lists OSPF neighbor routers and information about each neighbor: |

| | |
|---|---|
| | • **Neighbor**. The IP address of each neighbor router.<br>• **Neighbor ID**. Each neighbor router's OSPF router ID.<br>• **State**. Displays one of the following: Down; Attempt; Init; 2-way; ExStart; Exchange; Loading; or Full. For a definition of each of these states, see "How OSPF Works" in How Unicast IPv4 Routing Protocols and Services Work.<br>• **Priority**. The router priority of each neighboring router.<br>• **State-changes and errors**. The total number of changes of state and errors.<br>• **Queue length**. The length of the retransmit queue.<br>• **Show protofilter**. For the OSPF protocol filter (autostatic, local, RIP, SNMP, non-demand-dial, or static), shows whether the action for IP packets is **drop** or **accept**.<br>• **Show routefilter**. Displays OSPF route filter information.<br>• **Show virtif**. Displays information about the specified virtual link.<br>• **Show virtifstats**. Displays information about all OSPF virtual links:<br>• **Area ID**. The area ID of the transit area for the virtual link.<br>• **Neighbor ID**. The neighboring router's OSPF router ID.<br>• **State**. Displays one of the following: Down; Loopback; Waiting; Point-to-point; Designated-router; Backup designated-router; Other; or Unknown.<br>• **Sent packets**. The number of packets sent over this virtual link.<br>• **Received packets**. The number of packets received over this virtual link.<br>• **Dropped packets**. The number of packets that failed to be transmitted over this virtual link. |
| Uninstall | Removes the OSPF routing protocol. |

**Netsh routing IP relay context**

Any IPv4 subnet that contains DHCP clients requires either a DHCP server or a properly configured DHCP relay agent to provide address leases to those DHCP clients. A computer acting as a DHCP relay agent forwards DHCP messages between DHCP clients on an IP subnet that has no DHCP server and a DHCP server located on a different subnet. Using DHCP relay agents make it unnecessary to install a separate DHCP server on each subnet.

The following table describes the commands related to DHCP and DHCP relay agent in the Netsh routing IP relay context. The capability to enable a router to act as a DHCP relay agent is provided by the Windows Server 2003 Routing and Remote Access service.

**Commands in the Netsh Routing IP Relay Context**

| Command | Task |
|---|---|
| Add | • **Add dhcpserver**. Adds a DHCP server to the global list of DHCP servers.<br>• **Add interface**. Enables the DHCP Relay Agent service on the specified interface. |

| | |
|---|---|
| Delete | • **Delete dhcpserver**. Deletes a DHCP server from the global list of DHCP servers.<br><br>• **Delete interface**. Disables the DHCP Relay Agent service on the specified interface. |
| Dump | Creates a script that contains the current configuration for the DHCP relay agent. If you save the script to a file, you can use it to restore altered configuration information.<br><br>If you plan to make extensive changes to the configuration of a DHCP relay agent, a recommended practice is to begin the configuration session by using the **dump** command before you start, in case you need to restore the configuration later. |
| Install | Installs the DHCP Relay Agent service. |
| Set | • **Set global**. Sets the logging level for a DHCP relay agent to one of the following: none; errors only; errors and warnings; or errors, warnings, and informational messages.<br><br>• **Set interface**. On the specified interface, updates the following DHCP relay agent information:<br><br>  • Whether the DHCP Relay Agent service is enabled or disabled.<br><br>  • The maximum hop count before a DHCP packet is dropped.<br><br>  • The minimum number of seconds that the DHCP relay agent waits before forwarding a DHCP packet (from a DHCP client computer initiating an IP address renewal request) to a listed DHCP server.<br><br>  • The delay allows a DHCP server on the local subnet (if one is available) time in which to respond to the DHCP client's request. If no DHCP server responds after the configured interval, the DHCP relay agent then processes the request. |

| | |
|---|---|
| Show | • **Show global**. Displays the following information about the DHCP relay agent configured on this router:<br>  • The logging level: none; errors only; errors and warnings; or errors, warnings, and informational messages.<br>  • The maximum number of DHCP packets in the router's queue (for all interfaces).<br>  • The number of DHCP servers in the global list.<br>• **Show ifbinding**. For the specified interface, displays IP address bindings. If no interface is specified, displays IP address bindings for each interface on which the DHCP relay agent is enabled.<br>• **Show ifconfig**. For the specified interface, displays:<br>  • Whether the DHCP relay agent is bound and enabled.<br>  • The maximum hop count before a DHCP packet is dropped.<br>  • The minimum number of seconds from startup before a DHCP packet is sent to a listed DHCP server.<br>  If no interface is specified, displays this information for each interface on which the DHCP relay agent is enabled.<br>• **Show ifstats**. For the specified interface, displays:<br>  • The number of send and receive failures.<br>  • Address Resolution Protocol (ARP) update failures.<br>  • Requests and replies received or discarded.<br>  If no interface is specified, displays statistics for each interface on which the DHCP relay agent is enabled.<br>• **Show interface**. For the specified interface, displays:<br>  • The state (enabled or disabled, and bound or unbound).<br>  • The relay mode (enabled or disabled).<br>  • The maximum hop count.<br>  • The minimum number of seconds that a DHCP relay agent waits after a DHCP client initiates an IP address renewal request before the DHCP relay agent forwards the DHCP client request to a remote DHCP server. |
| Uninstall | Removes the DHCP Relay Agent service. |

**Netsh routing IP RIP context**

Routing Information Protocol (RIP), the primary dynamic routing protocol used on small or medium-sized IPv4 internetworks (10 to 50 networks), enables the dynamic exchange of IPv4 routing information between RIP routers. RIP routers can update IP internetwork topology changes as networks are added or removed and as links fail and are restored.

The following table describes the commands available in the Netsh routing IP RIP context. The RIP routing protocol is provided by the Windows Server 2003 Routing and Remote Access service, which supports both the updated version of RIP, RIP version 2 (RIP v2) and the earlier version of RIP, RIP version 1 (RIP v1).

**Commands Available in the Netsh Routing IP RIP Context**

| Command | Task |
|---|---|
| Add | • **Add acceptfilter**. Adds an acceptance filter for routes received on an interface by specifying the IP address range.<br>• **Add announcefilter**. Adds a filter for routes announced on an interface by specifying the IP address range.<br>• **Add interface**. Enables RIP on the specified interface, and configures the following:<br>  ◦ A metric value for routes based on the interface.<br>  ◦ Whether to perform on-demand updates or periodic updates.<br>  ◦ Whether to announce no routes, RIP v1 routes only, RIP v1 and v2 routes, or RIP v2 routes only.<br>  ◦ Whether to accept no routes, RIP v1 routes only, RIP v1 and v2 routes, or RIP v2 routes only.<br>  ◦ An interval for route expiration, route removal, and full update.<br>  ◦ Whether simple password authorization is required. This type of authorization uses plaintext passwords, which are used by routers to identify each other, not for security.<br>  ◦ The route tag.<br>  ◦ Whether to use unicast mode in addition to the default mode, use only unicast mode, or to disable unicast mode. The default mode is to use broadcast or multicast only (but not unicast), based on the RIP version used (that is, RIP v1 broadcast, RIP v 2 broadcast, or RIP v 2 multicast).<br>  ◦ Whether to include, exclude, or disable acceptance filters.<br>  ◦ Whether to include, exclude, or disable announcement filters.<br>  ◦ A plaintext password (1–16 characters).<br>• **Add neighbor**. Adds a RIP neighbor router on an interface. Specify the IP address for the RIP neighbor router.<br>• **Add peerfilter**. Adds a filter for another RIP router from which this router will accept announcements. Specify the IP address of the peer server. By default, RIP announcements from all sources are accepted. By configuring filters accepting announcements from one or more RIP peers, RIP announcements from unauthorized RIP routers are discarded. |
| Delete | • **Delete acceptfilter**. Deletes an acceptance filter for routes received on an interface.<br>• **Delete announcefilter**. Deletes an announcement filter set for an interface.<br>• **Delete interface**. Disables RIP from the specified interface.<br>• **Delete neighbor**. Deletes a RIP neighbor from an interface.<br>• **Delete peerfilter**. Deletes a filter that specifies another RIP router from which this router will accept announcements. |

| | |
|---|---|
| Dump | Creates a script that contains the current configuration for RIP. If you save the script to a file, you can use it to restore altered configuration information or to configure a new RIP router.<br><br>If you plan to make extensive changes to the configuration of a RIP router, a recommended practice is to begin the configuration session by using the **dump** command before you start, in case you need to restore the configuration later. |
| Install | Installs the RIP routing protocol under the IP routable protocol. |
| Set | • **Set flags**. Sets RIP-related flags for the specified interface by specifying one or more of the following values:<br><br>  ◦ **Clear**. Clear all currently set RIP-related flags.<br><br>  ◦ **Splithorizon**. Enable RIP for split horizon.<br><br>  ◦ **Poisonreverse**. Enable RIP for poison reverse.<br><br>  ◦ **Triggeredupdates**. Enable RIP triggered updates.<br><br>  ◦ **Cleanupupdates**. Enable cleanup updates for RIP.<br><br>  ◦ **Accepthostroutes**. Enable acceptance of host routes.<br><br>  ◦ **Acceptdefaultroutes**. Enable acceptance of default routes.<br><br>  ◦ **Senddefaultroutes**. Enable sending of default routes.<br><br>  ◦ **Nosubnetsummary**. Disable subnet summarization.<br><br>• **Set global**. Sets the following global RIP parameters:<br><br>  ◦ Whether to set the level of logging for RIP events to one of the following: none; errors only; errors and warnings; or errors, warnings, and informational messages.<br><br>  ◦ Minimum interval in minutes between triggered updates.<br><br>  ◦ Whether to include peers, exclude peers, or disable peer mode.<br><br>• **Set interface**. Modifies one or more of the following on the specified interface:<br><br>  ◦ A metric value for routes based on this interface.<br><br>  ◦ Whether to perform on-demand updates or periodic updates.<br><br>  ◦ Whether to announce no routes, RIP v1 routes only, RIP v1 and v2 routes, or RIP v2 routes only.<br><br>  ◦ Whether to accept no routes, RIP v1 routes only, RIP v1 and v2 routes, or RIP v2 routes only.<br><br>  ◦ An interval for route expiration, route removal, and full update.<br><br>  ◦ Whether simple password authorization is required. This type of authorization uses plaintext passwords, which are used by routers to identify each other, not for security.<br><br>  ◦ The route tag.<br><br>  ◦ Whether to also use unicast mode in addition to default mode, only use unicast mode, or to disable unicast mode. The default mode is to use broadcast or multicast only (but not unicast), based on the |

| | |
|---|---|
| | RIP version used (that is, RIP v1 broadcast, RIP v 2 broadcast, or RIP v 2 multicast).<br>● Whether to include, exclude, or disable acceptance filters.<br>● Whether to include, exclude, or disable announcement filters.<br>● A plaintext password (16 characters maximum). |

| | |
|---|---|
| Show | • **Show flags**. Displays RIP flags set on the specified interface.<br><br>• **Show global**. Displays global parameters for this RIP router:<br><br>  • The specified level of logging: none; errors only; errors and warnings; or errors, warnings, and informational messages.<br><br>  • The maximum size for the receive queue and for the send queue.<br><br>  • The minimum interval in minutes between triggered updates.<br><br>  • The status (include, exclude, or disable) for peer filter mode.<br><br>  • The number of peer filters.<br><br>• **Show globalstats rr=** *n*. Displays the number of route changes and the number of responses sent every *n* seconds. In this command, rr stands for refresh rate, that is, the interval, in seconds, that you specify after which the statistics for route changes and responses sent are refreshed.<br><br>• **Show ifbinding**. Displays IP address bindings for the specified RIP interface.<br><br>• **Show ifstats**. Displays RIP statistics for the specified RIP interface, including:<br><br>  • Send and receive failures.<br><br>  • Requests sent and received.<br><br>  • Responses sent and received.<br><br>  • Bad response packets and bad response entries received.<br><br>  • Triggered updates sent.<br><br>• **Show interface**. Displays RIP configuration for the specified interface, including:<br><br>  • Metric.<br><br>  • Update, accept, and announce modes.<br><br>  • Interval for route expiration, route removal, and full update.<br><br>  • Authentication type and password.<br><br>  • Route tag.<br><br>  • Unicast peer mode, accept filter mode, and announce filter mode.<br><br>  • Unicast peer count, accept filter count, and announce filter count.<br><br>• **Show neighbor**. Shows RIP neighbor statistics. |
| Uninstall | Removes the RIP routing protocol |

**Netsh routing IP routerdiscovery context**

In unicast IPv4 routing, ICMP router discovery is an alternative to using manually configured or DHCP-configured default

gateways.

| Note |
| --- |
| In IPv6 unicast routing, IPv6 Neighbor Discovery (ND) replaces and improves the functionality of Address Resolution Protocol (ARP), ICMP Router Discovery, and ICMP Redirect, which together are the IPv4 counterparts for IPv6 ND. ND, which is a set of messages and processes that determine relationships between neighboring nodes, is described in RFC 2461, "Neighbor Discovery for IP Version 6 (IPv6)," in the IETF RFC Database. |

ICMP router discovery automates the discovery and configuration of the default gateway so that IP hosts can dynamically discover the best default gateway to use on a subnet and can automatically switch to another default gateway if the first default gateway fails or if the network administrator changes router preferences.

The following table describes the commands for ICMP router discovery in the Netsh routing IP routerdiscovery context. ICMP router discovery is provided by the Windows Server 2003 Routing and Remote Access service.

**Commands in the Netsh Routing IP Routerdiscovery Context**

| Command | Task |
| --- | --- |
| Add | **Add interface**. On the specified interface, specifies the following:<br>• Whether ICMP router discovery is enabled or disabled.<br>• A minimum and a maximum interval in minutes between router discovery packets.<br>• The time in minutes after which a host considers the interface to be failed (which is after it has received its last Router Advertisement message from this router).<br>• The preference level for using this interface as the default gateway. A higher number indicates a higher level of preference. |
| Delete | **Delete interface**. Disables ICMP router discovery configuration for the specified interface. |
| Dump | Creates a script that contains the current configuration of the ICMP router discovery component. If you save the script to a file, you can use it to restore altered configuration information.<br>If you plan to make extensive changes to the configuration of the ICMP router discovery component, a recommended practice is to begin the configuration session by using the **dump** command before you start, in case you need to restore the configuration later |
| Set | **Set interface**. On the specified interface, updates the following:<br>• Whether ICMP router discovery is enabled or disabled.<br>• A minimum and a maximum interval in minutes between router discovery packets.<br>• The time in minutes after which a host considers the interface to be failed (which is after it has received its last Router Advertisement message from this router).<br>• The preference level for using this interface as the default gateway. A higher number indicates a higher level of preference. |
| Show | **Show interface**. For the specified interface, displays:<br>• Whether ICMP router discovery is enabled or disabled.<br>• The minimum and maximum advertisement intervals.<br>• The preference level for using this interface as the default gateway. |

| | |
|---|---|
| Uninstall | Removes ICMP router discovery from all interfaces on this server. If, later, you want to reenable ICMP router discovery by using the **add interface** command, you must add ICMP router discovery one interface at a time. |

**Management and Diagnostic Command-Line Tools for Unicast IPv4 Routing**

In addition to the Netsh command-line tool, Windows Server 2003 also supports the Pathping, Ping, Route, and Tracert command-line tools for unicast IPv4 and IPv6 network traffic. This section describes each tool as used for unicast IPv4 routing–related tasks in an IPv4 internetworking environment. For information about these tools in an IPv6 internetworking environment, see IPv6 Tools and Settings.

**Pathping.exe**

**Category** The Pathping.exe diagnostic tool for IPv4 internetworks is included with the Windows Server 2003 operating system. The Windows Server 2003 version of Pathping.exe also supports IPv6.

**Version compatibility** The Pathping.exe tool runs on Windows Server 2003, Windows XP, and Windows 2000.

Pathping combines information provided by Tracert (identifying which routers exist along the path) and Ping (repeatedly sending messages to each router) with additional information that neither of those tools provides.

The Pathping tool uses ICMP Echo messages to test router latency, link latency, and packet losses on an IPv4 internetwork. Pathping tests each hop by pinging multiple times over a period of time, and then displaying delay and packet loss information per hop. Specifically, Pathping records the following:

- The average round-trip time.
- The packet loss when sending ICMP Echo messages to each router.
- The packet loss when sending ICMP Echo messages across the links between each router.

The output of the Pathping command can help you determine if there is a high-loss link or router in the path by showing the degree of packet loss at each router or link.

For more information about Pathping.exe, see Command-Line Reference for Windows Server 2003, Standard Edition.

**Ping.exe**

**Category** The Ping.exe diagnostic tool for IPv4 internetworks is included with the Windows Server 2003 operating system. The Windows Server 2003 version of Ping.exe also supports IPv6.

**Version compatibility** The Ping.exe tool runs on Windows Server 2003, Windows XP, Windows 2000, and Windows NT® 4.0.

The Ping tool tests reachability on an IPv4 internetwork by sending an ICMP Echo message from a source host to a destination IP address and recording the results. Ping indicates whether the destination responds by displaying an ICMP Echo Response message and how much time elapses (in milliseconds) before a reply is received. If Ping receives no response, it displays an error message instead.

In addition to testing connectivity to a destination computer, successfully pinging from the local computer by IP address, not by name, verifies that TCP/IP is functioning properly on the local computer. Likewise, successfully pinging the local router verifies that the router is functioning.

For more information about Ping.exe, see Command-Line Reference for Windows Server 2003, Standard Edition.

**Route.exe**

**Category** The Route.exe management tool for IPv4 internetworks is included with the Windows Server 2003 operating system. The Windows Server 2003 version of Route.exe also supports IPv6.

**Version compatibility** The Route.exe tool runs on Windows Server 2003, Windows XP, Windows 2000, and Windows NT 4.0.

The Route tool displays and modifies entries in the IPv4 routing table on the local computer. You can use Route to add a temporary or persistent route, to change an existing route, or to remove a route from the IP routing table. The Route tool has four subcommands, listed in the following table.

**Subcommands for Route.exe**

| Subcommand | Task |
|---|---|

| | |
|---|---|
| Route print | Prints the IPv4 routing table and (if IPv6 exists) the IPv6 routing table.<br><br>You can also use **route print** to display only part of the routing table. For example, if you use route print to display the entire routing table and then decide that you want to look only at the set of routes that begin with, say, 157, you can use **route print 157*** to print only those routes that match 157*. |
| Route add | Adds a static route.<br><br>In Windows Server 2003, this subcommand does not support IPv6. |
| Route delete | Deletes a static route.<br><br>In Windows Server 2003, this subcommand does not support IPv6. |
| Route change | Modifies an existing static route. This command is used only to modify gateway or metric values.<br><br>In Windows Server 2003, this subcommand does not support IPv6 |

The IP routing table is stored in RAM, which means that a non-persistent static route disappears whenever a computer restarts. Typically, therefore, you use **route add** with the **-p** option to add a persistent static route that continues to exist even after a computer restarts.

Persistent routes are stored in the registry in the **PersistentRoutes** entry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters. Because you can add, modify, or delete a persistent static route by using the Route tool, you do not need to configure this registry entry directly.

For more information about Route.exe, see Command-Line Reference for Windows Server 2003, Standard Edition.

**Tracert.exe**

**Category** The Tracert.exe diagnostic tool for IPv4 internetworks is included with the Windows Server 2003 operating system. The Windows Server 2003 version of Tracert.exe also supports IPv6.

**Version compatibility** The Tracert.exe tool runs on Windows Server 2003, Windows XP, Windows 2000, and Windows NT 4.0.

The Tracert route tracking tool tests reachability, network latency, routing loops, and other issues on an IPv4 internetwork by determining the path (the series of routers) that unicast IPv4 traffic takes between a source host and destination IP address. Tracert sends ICMP Echo messages with incrementally higher values of the Time to Live (TTL) field in the IP header and then returns (lists) the series of near-side router interfaces on each router between the source and destination hosts. If the intended destination host does not receive the packets sent, Tracert displays the last router that successfully forwarded the packets.

For more information about Tracert.exe, see Command-Line Reference for Windows Server 2003, Standard Edition.

**Network Connections**

For unicast IPv4 routing, Network Connections is used to configure default gateway options and other options for TCP/IP.

**Category** Network Connections is included with the Windows Server 2003 operating system.

**Version compatibility** Network Connections runs on Windows Server 2003 and Windows XP. Network Connections is known as Network and Dial-up Connection in Windows 2000. Network Connections is known as the Network component of Control Panel in Windows NT 4.0.

The Network Connections management tool enables the configuration of TCP/IP properties for a variety of purposes. The following table describes only the options used for manually configuring the default gateway and the related automatic metric option.

**TCP/IP Properties Options in Network Connections**

| Page | Tab | Task |
|---|---|---|

| | | |
|---|---|---|
| TCP/IP Properties | General | Manually configuring a network adapter to point to a default gateway:<br>• Under **Use the following IP address**, specify:<br>  • **IP address** and **Subnet mask**. Together, the IP address and subnet mask identify this computer and the subnet on which it is located.<br>  • **Default gateway**. The IP address for the default gateway identifies a local IP router (on the same subnet as this computer) that forwards traffic for this computer to destinations on a different subnet. |
| Advanced TCP/IP Settings | IP Settings | Adding, editing, or removing additional default gateways for this network connection:<br>• Under **Default gateways**, click **Add:**<br>  Opens the **TCP/IP Gateway Address** dialog box, which you can use to configure an IP address and an automatic or manually configured metric. If the default gateway specified on the **General** tab fails, the local computer can switch to an alternative default gateway listed here if dead gateway detection is enabled.<br>  See the registry entry **EnableDeadGWDetect** later in this document.<br>Enabling the option to automatically calculate the interface cost based on the speed of the interface:<br>• Select **Automatic metric**. An interface metric is enabled that is calculated automatically by TCP/IP based on the speed of the interface. If multiple closest matching routes to a destination exist in the IP routing table, IP uses the route with the lowest cost so that the fastest interface is used to forward traffic to the default gateway.<br>Disabling the option to automatically calculate the interface cost based on the speed of the interface:<br>• Clear **Automatic metric**, and then specify an **Interface metric**. If you disable **Automatic metric**, the **Interface metric** field becomes available to manually configure a cost for this interface. |

For related information about tools for configuring IP routers and hosts to enable hosts to dynamically discover the best default gateway to use on a subnet, see "Netsh routing IP routerdiscovery context" earlier in this document and see "ICMP router discovery options" later in this document.

**Rrasmgmt.msc: Routing and Remote Access Snap-in**

For unicast IPv4 routing, the Routing and Remote Access snap-in is used to configure options for a variety of functions provided by the Routing and Remote Access service. These functions include event logging, route source preference levels, the RIP for IP and OSPF routing protocols, the DHCP relay agent, IP packet filtering, and ICMP router discovery.

**Category** The Routing and Remote Access snap-in of Microsoft Management Console (MMC) is included with Windows Server 2003 operating system. Routing and Remote Access is disabled by default.

**Version compatibility** The Routing and Remote Access snap-in is provided by the Routing and Remote Access service on computers running Windows Server 2003 and Windows 2000 Server. For Windows NT 4.0, the Routing and Remote Access Service (RRAS) and its snap-in are available as a separate download from the Microsoft Windows NT Server Routing and Remote Access Service Download page.

You can use the Routing and Remote Access snap-in to enable, configure, or modify, unicast IPv4 routing options for RIP, OSPF, DHCP relay agent, IP packet filtering, and ICMP router discovery. The Routing and Remote Access snap-in can also be used to configure options for a variety of other routing or remote access services unrelated to these unicast IPv4 routing topics, such as configuring remote access connections, multicast IPv4 routing options, or network address translation (NAT).

**Options for enabling, disabling, or refreshing the Routing and Remote Access service**

The Routing and Remote Access snap-in is available by default in Administrative Tools in Control Panel. However, the Routing and Remote Access service managed by means of this snap-in is disabled by default.

| Note |
| --- |
| You do not use the Add or Remove Programs icon in Control Panel to install or uninstall the Routing and Remote Access service. |

The following table describes the options to enable, disable, or refresh the Windows Server 2003 Routing and Remote Access service. These settings are available in the console tree of the Routing and Remote Access snap-in by right-clicking the *ServerName* node.

**Options to Enable, Disable, or Refresh the Routing and Remote Access Service**

| Node | Task |
| --- | --- |
| *ServerName* | Enabling the Routing and Remote Access service:<br>• **Configure and Enable Routing and Remote Access**. Starts Routing and Remote Access Server Setup Wizard. To enable LAN routing, select **Custom configuration** on the wizard's **Configuration** page, and then select **LAN routing** on the **Custom configuration** page.<br>**Note**<br>Alternatively, you can use the wizard to configure and enable connections for remote access dial-up or VPN, NAT, or a site-to-site VPN connection between two private networks from the **Configuration** page. If you select the **Custom configuration** option from the **Configuration** page, you can enable any of these connections in addition to LAN routing, or you can enable LAN routing only.<br>Disabling the Routing and Remote Access service:<br>• **Disable Routing and Remote Access**. Disables the Routing and Remote Access service, which removes all Routing and Remote Access registry settings, but it does not uninstall the service or snap-in.<br>Refreshing the Routing and Remote Access service:<br>• **Disable Routing and Remote Access**, and then **Configure and Enable Routing and Remote Access**. Refreshes the Routing and Remote Access service by first disabling it, which removes all current configuration, and then reenabling it, which starts the wizard so that you can re-enable the service and configure new settings. You do not need to restart the server. |

**General options**

Options available in the **IP Routing\General** container in the Routing and Remote Access snap-in take effect at the level of the server rather than at the level of an individual component of the Routing and Remote Access service. General options include those for setting global logging levels and for ranking route sources. You can also set logging levels at the level of the individual component.

This section includes unicast IPv4 routing options available in the console tree on the **IP Routing\General** node of the Routing and Remote Access snap-in and in the details pane on the *InterfaceName* **Properties** page. Multicast routing options are not included in this section. For more information, see IPv4 Multicasting Tools and Settings.

**General options in the console tree**

The following table describes the options for adding a new unicast IPv4 routing protocol to the Routing and Remote Access snap-in console tree and for adding an interface to the details pane. These options are available by right-clicking the **IP Routing\General** node.

**Options on the IP Routing\General Node in the Console Tree**

| Node | Task |
| --- | --- |

| | |
|---|---|
| General | Adding a new routing protocol to the console tree:<br>• **New Routing Protocol**. Provides a list from which to select one or more of the following protocols to add to the Routing and Remote Access console tree under the **IP Routing** node:<br>  • DHCP Relay Agent<br>  • IGMP Router and Proxy<br>  • RIP<br>  • OSPF<br>  • NAT/Basic Firewall<br>  This document describes the management of the DHCP relay agent, RIP, and OSPF components in the Routing and Remote Access snap-in. For information about IGMP, see IPv4 Multicasting Tools and Settings. For information about NAT/Basic Firewall, see NAT Tools and Settings.<br><br>Adding a new interface to the details pane:<br>• **New Interface.** Provides a list of available interfaces on this server that you can add to the details pane of the **IP Routing\General** node. |

The following table describes Routing and Remote Access options for configuring global IP logging and for ranking route source preference levels. These settings are available in the console tree of the Routing and Remote Access snap-in on the properties page of the **IP Routing\General** node. The properties page on the **General** node also includes a tab for **Multicast Scopes**, which is not covered in this technical reference.

**Options on the IP Routing\General Properties Page**

| Tab | Task |
|---|---|
| Logging | Configuring events to be logged in the System Log in Event Viewer for this server:<br>• **Log errors only**<br>• **Log errors and warnings**<br>• **Log the maximum amount of information** (that is, errors, warnings, and informational messages)<br>• **Disable event logging**<br><br>**Note**<br>Use information about routing events stored in the System Log to troubleshoot routing processes. Logging puts a considerable load on system resources and should be enabled only when needed. After the event is logged or the problem is identified, immediately reset logging to its default value, **Log errors only**. |

| | |
|---|---|
| Preference Levels | Configuring the rank of route sources in order of preference by moving one or more of the listed route sources up or down (route sources are listed here in default order):<br>• **Local**<br>• **Static** (a static route assigned to a demand-dial interface)<br>• **Static (non demand-dial)** (a static route assigned to a LAN interface)<br>• **Auto-static**<br>• **Network management** (such as SNMP)<br>• **Open Shortest Path First (OSPF)**<br>• **RIP Version 2 for Internet Protocol**<br><br>**Note**<br>If two routes exist to the same destination, the route learned by the route source with a higher preference is preferred regardless of the route metric. For example, if the preference level for OSPF is higher than the preference level for RIP v2, an OSPF route with a metric of 5 is entered into the routing table rather than the same route discovered by RIP with a metric of 3. |

**General options on the InterfaceName Properties page**

For a description of unicast IPv4 routing configuration options available on the **General** tab of the *InterfaceName* **Properties** page in the **IP Routing\General** container, see the separate sections for "IP Packet Filtering Options" and "ICMP router discovery options" later in this document.

This *InterfaceName* **Properties** page also provides the **Configuration** tab to specify whether the IP address for this interface is statically configured or obtained dynamically from a DHCP server, and it provides the **Multicast Boundaries** and **Multicast Heartbeat** tabs to configure options related to IPv4 multicast routing. None of these options is described here because they are outside the scope of this technical reference.

**RIP Options**

Routing Information Protocol (RIP), the primary dynamic routing protocol used on small or medium-sized IPv4 internetworks (10 to 50 networks), enables the dynamic exchange of IPv4 routing information between RIP routers. By adding RIP to the Routing and Remote Access snap-in console tree, adding RIP interfaces to the details pane, and then configuring RIP options, you can enable RIP routers to update topology changes as networks are added or removed and as links fail and are restored.

The Windows Server 2003 Routing and Remote Access service supports both the updated version of RIP, RIP version 2 (RIP v2) and the earlier version of RIP, RIP version 1 (RIP v1). The version of RIP installed when you add the RIP protocol to the Routing and Remote Access snap-in console tree is RIP v2, which includes options that make RIP v1 optionally available.

This section includes unicast IPv4 routing options available in the console tree of the Routing and Remote Access snap-in on the **IP Routing\RIP** node and in the details pane on the *InterfaceName* **Properties** page.

**RIP options in the console tree**

The following table describes options for adding one or more interfaces to the details pane and for displaying information about neighboring RIP routers. These options are available by right-clicking the **IP Routing\RIP** node.

**Options on the IP Routing\RIP Node in the Console Tree**

| Node | Task |
|---|---|

| | |
|---|---|
| RIP | Adding a new interface to the details pane:<br>• **New Interface**. Provides a list of available interfaces on this server that you can add to the details pane of the **IP Routing\RIP** node.<br><br>Displaying neighboring RIP routers:<br>• **Show Neighbors**. Displays a table of RIP routers configured as neighbor routers to this RIP router. For each RIP neighbor, the table lists its IP address, RIP version, and any bad packets or bad routes. |

The following table describes Routing and Remote Access options for the RIP routing protocol. These settings are available in console tree of the Routing and Remote Access snap-in on properties page of the **IP Routing\RIP** node.

**Options on the RIP Properties Page**

| Tab | Task |
|---|---|
| General | Configuring the delay before this RIP router sends triggered updates:<br>• **Maximum delay (seconds)**. Specifies the maximum amount of time that this router waits before it sends triggered updates (updates that are announced as soon as a network topology change occurs, rather than at periodic intervals). The default value is 5 seconds.<br><br>Configuring events to be logged in the System Log in Event Viewer for this RIP router:<br>• **Log errors only**<br>• **Log errors and warnings**<br>• **Log the maximum amount of information** (that is, errors, warnings, and informational messages)<br>• **Disable event logging**<br><br>**Note**<br>Use information about routing events stored in the System Log to troubleshoot routing processes. Logging puts a considerable load on system resources and should be enabled only when needed. After the event is logged or the problem is identified, immediately reset logging to its default value, **Log errors only**. |
| Security | Configuring the RIP routers from which this RIP router accepts announcements:<br>• **Accept announcements from all routers (default)**. Specifies that this RIP router accepts announcements from all RIP routers.<br>• **Accept announcements from listed routers only**. Specifies that this RIP router accepts only announcements from the listed RIP routers.<br>• **Ignore announcements from all listed routers**. Specifies that this RIP router ignores announcements from the listed RIP routers but accepts announcements from all other RIP routers.<br>• **Router IP address**. Specifies the list (if any) of routers whose announcements you want to accept or ignore. |

**RIP for IP options on the InterfaceName Properties page**

The following table describes Routing and Remote Access options for a RIP interface. These settings are available in the Routing and Remote Access snap-in on the specified tab of the *InterfaceName* **Properties** page for an interface in the **IP Routing\RIP** container.

**RIP Options on the InterfaceName Properties Page**

| Tab | Task |
|---|---|

| | |
|---|---|
| General | Configuring **Operation mode** options for this interface:<br><br>• **Auto-static update mode (default for demand-dial interfaces)**. Sends RIP announcements only when other routers request an update. Routes learned are marked as static routes in the IP routing table and remain in the routing table until they are manually deleted.<br><br>• **Periodic update mode (default for LAN interfaces)**. Sends RIP announcements periodically every *n* seconds as specified in **Periodic announcement interval** on the **Advanced** tab. Routes learned are marked as RIP routes in the IP routing table and are removed if the router restarts.<br><br>Configuring **Outgoing packet protocol** options for sending outgoing RIP route announcements over this interface:<br><br>• **RIP version 1 broadcast**. Broadcasts RIP v1 announcements. Specify this protocol only if all neighboring RIP routers are also RIP v1 routers.<br><br>• **RIP version 2 broadcast**. Broadcasts RIP v2 announcements. Specify this protocol for a mixed environment that uses both RIP v1 and RIP v2.<br><br>• **RIP version 2 multicast**. Multicasts RIP v2 announcements. Specify this protocol only if all neighboring RIP routers are also RIP v2 routers.<br><br>• **Silent RIP**. Disables outgoing RIP announcements on this interface. Specify this option to configure this non-router computer to listen for incoming RIP announcements and update its routing table but not to announce its own routes.<br><br>Configuring **Incoming packet protocol** options for receiving incoming RIP announcements over this interface:<br><br>• **Ignore incoming packets**. Configures this router only to announce routes (but not to receive incoming route announcements).<br><br>• **RIP version 1 and 2**. Configures this router to accept both RIP v1 and RIP v2 announcements.<br><br>• **RIP version 1 only**. Configures this router to accept only RIP v1 announcements.<br><br>• **RIP version 2 only**. Configures this router to accept only RIP v2 announcements.<br><br>Configuring metric and route tag options for this interface:<br><br>• **Added cost for routes**. Configures the cost associated with sending packets over this interface. For example, set a higher cost on a heavily trafficked route to reduce its use. Increasing the cost for all RIP routes reduces your maximum RIP network diameter.<br><br>• **Tag for announced routes**. Specifies a tag number for the routes advertised on this interface in order to distinguish these routes from non-RIP routes. Used on a router that can support multiple routing protocols.<br><br>Configuring the simple password authentication option for this interface:<br><br>• **Activate authentication**. Enables a form of authentication that uses a plaintext password for RIP v2 announcements over this interface. All incoming and outgoing RIP v2 packets from neighboring RIP v2 routers connected to this interface must contain the same plaintext password. This type of authentication uses plaintext passwords, which are used by routers to identify each other, not for security.<br><br>• **Password**. Specifies a plaintext password (1–16 characters). |

| | |
|---|---|
| Security | Configuring RIP security options for incoming routes on this interface:<br>• **Accept all routes**. Configures this RIP router to look at each route entry in an incoming RIP announcement to determine whether to update its routing table.<br>• **Accept all routes in the ranges listed**. Configures this RIP router to look at each route entry in an incoming RIP announcement but processes the route only if it falls within one of the ranges listed.<br>• **Ignore all routes in the ranges listed**. Configures this RIP router to look at each route entry in an incoming RIP announcement and discards the route if it falls into one of the ranges listed.<br><br>Configuring RIP security options for outgoing routes on this interface:<br>• **Accept all routes**. Configures this RIP router to include all appropriate route entries in outgoing RIP announcements.<br>• **Accept all routes in the ranges listed**. Configures this RIP router to include a route in the outgoing RIP announcement only if the route falls into one of the ranges listed.<br>• **Ignore all routes in the ranges listed**. Configures this RIP router to exclude a route from the outgoing RIP announcement if the route falls into one of the ranges listed. |
| Neighbors | Configuring how this RIP router interacts with neighboring RIP routers through this interface:<br>• **Use broadcast or multicast only**. Specifies that all RIP announcements are broadcasted or multicasted as specified in **Outgoing packet protocol** on the **General** tab.<br>• **Use neighbors in addition to broadcast or multicast**. Specifies that RIP announcements are unicasted to specific neighboring routers in addition to broadcasted or multicasted as specified in **Outgoing packet protocol** on the **General** tab.<br>• **Use neighbors instead of broadcast or multicast**. Specifies that RIP announcements are unicasted only to specific neighboring routers. For nonbroadcast networks, such as Frame Relay, select this option to enable RIP traffic between specific routers.<br>• **IP address**. Specifies the IP address of one or more neighboring RIP routers. Valid only if one of the preceding options to send RIP announcements only to specific neighboring routers is selected. |
| Advanced | Configuring RIP announcement and route interval options for this interface:<br>• **Periodic announcement interval (seconds)**. Specifies the interval — from 15 seconds through 86,400 seconds (24 hours) — between periodic RIP announcements. The default value is 30 seconds. Valid only if **Periodic update mode** is selected as the **Operation mode** on the **General** tab.<br>• **Time before routes expire (seconds)**. Specifies the lifetime — from 15 seconds through 259,200 seconds (72 hours) — of a route that is learned through RIP. The default value is 180 seconds. If a route is not updated with another RIP announcement within the specified time, the route expires. Valid only if **Periodic update mode** is selected as the **Operation mode** on the **General** tab.<br>• **Time before route is removed (seconds)**. Specifies the amount of time — from 15 seconds through 259,200 seconds (72 hours) — after which a RIP-learned route that has expired is removed from the routing table. The default value is 120 seconds. Valid only if **Periodic update mode** is selected as the **Operation mode** on the **General** tab.<br><br>Configuring split-horizon processing and triggered updates options for this interface:<br>• **Enable split-horizon processing**. Enabled by default. Specifies that routes learned from a neighboring RIP |

| | |
|---|---|
| | router are not announced in RIP announcements that this router sends to that neighbor (unless poison-reverse processing is also enabled, in which case such routes are announced with a metric of 16 to indicate explicitly that those routes are unavailable). Helps reduce convergence time and avoid routing loops.<br><br>• **Enable poison-reverse processing**. Enabled by default. Specifies that routes learned from a neighboring RIP router are announced with a metric of 16 (unreachable) in RIP announcements that this router sends to that neighbor. Available only if split horizon is enabled (if you disable split-horizon, poison-reverse processing is also disabled).<br><br>• **Enable triggered updates**. Specifies that a change in network topology, such as a new route or a change in metric, triggers an immediate update announcement that includes only the change. Configure **Maximum delay (seconds)** between triggered updates on the **General** tab on the **RIP Properties** page.<br><br>• **Send clean-up updates when stopping**. Enabled by default. Specifies that, if this RIP router stops, the router advertises as unreachable all of the routes that it makes available on all of its interfaces as it shuts down.<br><br>Configuring host and default route options for this interface:<br><br>• **Process host routes in received announcements**. Specifies that host routes in received RIP announcements are accepted. By default, host routes are ignored.<br><br>• **Include host routes in sent announcements**. Specifies that host routes are included in RIP announcements. By default, host routes are not included.<br><br>• **Process default routes in received announcements**. Specifies that default routes in received RIP announcements are accepted. By default, default routes are ignored.<br><br>• **Include default routes in sent announcements**. Specifies that default routes are included in RIP announcements. By default, default routes are not included.<br><br>Configuring subnet summarization option for this interface:<br><br>• **Disable subnet summarization**. Selected by default. Specifies that subnet routes are not summarized in the form of the class-based network ID when announced on a network that is not a subnet of the class-based network ID. Available only if **Outgoing packet protocol** on the **General** tab is set to either **RIP version 2 broadcast** or **RIP version 2 multicast**. In a mixed environment that contains both RIP v1 and RIP v2 routers, use only contiguous, fixed-size subnets, or clear this option. |

**OSPF options**

Open Shortest Path First (OSPF), the most common routing protocol used on large and very large IPv4 internetworks, enables the dynamic exchange of IPv4 routing information between OSPF routers across complex internetworks. Adding OSPF to the Routing and Remote Access snap-in console tree, adding OSPF interfaces to the details pane, and then configuring OSPF options enables OSPF routers to update topology changes as networks are added or removed and as links fail and are restored.

This section includes unicast IPv4 routing options available in the console tree of the Routing and Remote Access snap-in on the **IP Routing\OSPF** node and in the details pane on the *InterfaceName* **Properties** page.

**OSPF options in the console tree**

The following table describes options for adding one or more interfaces to the details pane and for displaying information about the OSPF autonomous system. These options are available by right-clicking the **IP Routing\OSPF** node.

**Options on the IP Routing\OSPF Node in the Console Tree**

| Node | Task |
|---|---|
| OSPF | Adding a new interface to the details pane:<br><br>• **New Interface**. Provides a list of available interfaces on this server that you can add to the details pane of the **IP Routing\OSPF** node. |

Displaying information about the OSPF autonomous system:

- **Show Areas**. Displays a table of OSPF areas, including:
    - **Area ID**. The identification number of each OSPF area.
    - **State**. Up or down.
    - **SPF computations**. The number of Shortest Path First (SPF) computations. An OSPF router creates a tree (the SPF tree) of shortest paths to each other router and subnet in the OSPF autonomous system.
- **Show Link-state Database**. Displays a table of information about the OSPF LSDB, including:
    - **Area ID**. The identification number of the OSPF area for each link state advertisement (LSA) in the table.
    - **Type**. The type of advertisement, such as Router, Summary (Network), Stub, Network, Summary (ASBR), or AS External.
    - **Link state ID**. The identification number of each advertisement in the table. The LSAs vary by advertisement type.
    - **Advertising router**. The OSPF router ID of the advertising router
    - **Age**. The age (in seconds) since the advertisement originated.
    - **Sequence**. A number that identifies old or duplicate advertisements.
- **Show Neighbors**. Lists information about neighboring OSPF routers, including:
    - **Neighbor**. The IP address of each neighboring router.
    - **Neighbor ID**. Each neighboring router's OSPF router ID.
    - **Type**. The network type of each neighboring router, such as NBMA.
    - **State**. The state of each neighbor router:Down, Attempt, Init, 2-way, ExStart, Exchange, Loading, or Full. For a definition of each of these states, see "How OSPF Works" in How Unicast IPv4 Routing Protocols and Services Work.
    - **Priority**. The router priority of each neighboring router.
    - **State-changes and errors**. The total number of changes of state and errors.
    - **Retransmit-queue length**. The length of the retransmit queue.
- **Show Virtual Interfaces**. Displays a table of virtual interfaces (virtual links). A virtual link is a logical connection that enables the exchange of routing information between a backbone area border router (ABR) and an ABR that is not physically connected to the backbone area across a transit area. You cannot configure a virtual link across the backbone area. For each virtual link, the table includes:
    - **Area ID**. The area ID of the transit area for the virtual link.
    - **Neighbor ID**. The neighboring router's OSPF router ID.

| | |
|---|---|
| | • **State**. The state of each virtual interface:Down, Loopback, Waiting, Point-to-point, Designated-router, Backup designated-router, Other, or Unknown.<br>• **Sent packets**. The number of packets sent over this virtual link.<br>• **Received packets**. The number of packets received over this virtual link.<br>• **Dropped packets**. The number of packets that failed to be transmitted over this virtual link. |

The following table describes Routing and Remote Access options for OSPF properties. These settings are available in the console tree of the Routing and Remote Access snap-in on the properties page of the IP Routing\OSPF node.

**Options on the OSPF Properties Page**

| Tab | Task |
|---|---|
| General | Configuring the OSPF router ID:<br>• **Router identification**. Specifies the 32-bit number (in dotted decimal notation) that uniquely identifies this OSPF router on the network. The OSPF router ID is not an IP address. Ensure that the number is unique, for example, by selecting the highest or the lowest IP address assigned to one of the interfaces on this computer.<br><br>Configuring this OSPF router as an autonomous system boundary router (ASBR):<br>• **Enable autonomous system boundary router**. Specifies that the OSPF router is an ASBR and therefore advertises routing information into the OSPF autonomous system that is from route sources external to the OSPF autonomous system. If you select this option, you can also use the **External Routing** tab to specify the route sources — such as static routes or the Routing Information Protocol (RIP) — from which this router will accept routes.<br><br>Configuring events to be logged in System Log in Event Viewer for this OSPF router:<br>• **Log errors only**<br>• **Log errors and warnings**<br>• **Log the maximum amount of information** (that is, errors, warnings, and informational messages)<br>• **Disable event logging**<br><br>**Note**<br>Use information about routing events stored in the System Log to troubleshoot routing processes. Logging puts a considerable load on system resources and should be enabled only when needed. After the event is logged or the problem is identified, immediately reset logging to its default value, **Log errors only**. |

| | |
|---|---|
| Areas | Configuring OSPF areas (each OSPF interface on this OSPF router must belong to at least one OSPF area):<br><br>• **Add**. Opens the **OSPF Area Configuration** dialog box, which you can use to configure the following options:<br><br>On the **General** tab:<br><br>• **Area ID**. Specifies an ID (a 32-bit number expressed in dotted decimal notation) that identifies the OSPF area. The only required OSPF area is the OSPF backbone area, which requires the reserved area ID of 0.0.0.0. An OSPF backbone area is either the sole area in the OSPF autonomous system, or it is the central area to which all other OSPF areas in this OSPF autonomous system connect and through which packets are routed from one non-backbone area to another non-backbone area.<br><br>Although an OSPF area ID is neither an IP address nor an IP network ID, if this is not the backbone area and if this area can be summarized with a single network ID, you can set the area ID to match the single network ID. For example, if an area contains all of the subnets for the IP network 10.1.0.0/16, you can set the area ID to 10.1.0.0.<br><br>• **Enable plaintext password**. Enabled by default. Specifies that a plaintext password is required for the OSPF area so that OSPF routers can identify each other (this is not a security option). If this password option is enabled, all interfaces in this area that are on the same subnet must use the same password. By default, the password is 12345678.<br><br>• **Stub area**. Specifies that this area is a stub area and therefore does not accept external routes from outside this OSPF autonomous system. Routing from a stub area to destinations outside this OSPF area, which includes destinations in other areas within its own autonomous system as well as to destinations outside its autonomous system, uses a summary default route. You cannot configure the backbone as a stub area, nor can you configure a virtual link through a stub area.<br><br>• **Stub metric**. Specifies the cost of the summary default route that this OSPF router uses to reach all destinations external to this OSPF autonomous system.<br><br>• **Import summary advertisements**. Specifies that inter-area routes for the OSPF autonomous system are imported into this stub area. By default, this check box is cleared and all non-intra-area destinations within the OSPF autonomous system are routed based on a single default route.<br><br>On the **Ranges** tab:<br><br>• **Destination** and **Network mask**. Specifies the network ID and subnet mask pair (that is, a summary route) that defines a range of addresses that belongs to this OSPF area. OSPF uses OSPF ranges to summarize the routes within this OSPF area.<br><br>• **Edit**. Opens the **OSPF Area Configuration** dialog box, which you can use to modify settings for the selected OSPF area. These settings are the same as those listed just above for **Add** on the **General** tab and on the **Ranges** tab.<br><br>• **Remove**. Removes the selected OSPF area. |

| | |
|---|---|
| Virtual Interfaces | Configuring a virtual interface for this OSPF router:<br><br>**Note**<br>A virtual interface, also called a virtual link, is a logical connection that enables the exchange of routing information between a backbone area border router (ABR) and an ABR that is not physically connected to the backbone area. You cannot configure a virtual link across the backbone area or a stub area.<br><br>• **Add**. Opens the **OSPF Virtual Interface Configuration** dialog box, which you can use to create a virtual link by configuring the following options:<br><br>On the **General** tab:<br><br>**Note**<br>The following settings must be the same on both routers at each end of the virtual link.<br><br>• **Transit area ID**. The OSPF area ID for the transit area for this virtual link. Select the one that you want to use from the provided list of non-backbone area IDs configured on this OSPF router.<br><br>• **Virtual neighbor router ID**. The OSPF router ID of the router at the other end of the virtual link.<br><br>• **Transit delay (seconds)**. The estimated number of seconds that it takes to transmit OSPF packets to the virtual link neighbor router over this interface. The default is 1 second.<br><br>• **Retransmit interval (seconds)**. The interval that the OSPF router waits before retransmitting link state messages. This value should be substantially larger than the expected round-trip delay between the two routers at either end of the transit area. The default is 10 seconds.<br><br>• **Hello interval (seconds)**. The interval between hello packets that are sent out over this interface. The default is 10 seconds.<br><br>• **Dead interval (seconds)**. The interval that a router's neighbors wait after they cease to hear the router's hello packets before they declare that the router is down. The default is 60 seconds.<br><br>• **Plaintext password**. The plaintext password to be shared by both routers at each end of the virtual link. This password is not the same as the plaintext password shared by all routers in an OSPF area so that OSPF routers in that area can identify each other.<br><br>• **Edit**. Opens the **OSPF Virtual Interface Configuration** dialog box, which you can use to modify the selected virtual interface. These settings are the same as those listed for the **Add** option.<br><br>• **Remove**. Removes the selected virtual interface. |

| | |
|---|---|
| External Routing | Configuring external route source options:<br>• **Accept routes from all route sources except those selected**. Specifies that routes from all selected route sources are ignored.<br>• **Ignore routes from all route sources except those selected**. Specifies that only routes from selected route sources are accepted.<br>• **Route sources**. Specifies non-OSPF route sources from which this OPSF router accepts or ignores routes:<br>  • Auto-static routes<br>  • Local routes<br>  • RIP version 2 for Internet Protocol<br>  • SNMP routes<br>  • Static (non-demand dial) Routes (refers to a static route assigned to a LAN interface)<br>  • Static routes (refers to a static route assigned to a demand-dial interface)<br>• **Route filters**. Opens the **OSPF External Route Filters** dialog box, which you can use to specify filters that accept or ignore listed external routes, and to specify the list of routes (by destination and subnet mask) to be accepted or ignored. |

**OSPF options on the InterfaceName Properties page**

The following table describes Routing and Remote Access options for an OSPF interface. These settings are available in the Windows Server 2003 Routing and Remote Access snap-in on the specified tab of the *InterfaceName* **Properties** page for an interface in the **IP Routing\OSPF** container.

**OSPF Options on the InterfaceName Properties Page**

| Tab | Task |
|---|---|

| | |
|---|---|
| General | Configuring OSPF on this interface:<br><br>• **Enable OSPF for this address**. Displays the IP address assigned to this interface, and specifies whether OSPF is enabled on this interface. If two or more IP addresses are assigned to the interface (by using the **Advanced TCP/IP Settings** page in **TCP/IP Properties**), select the IP address that you want to configure from the list provided. If more than one IP address is configured for this interface, you must configure OSPF settings on the **General**, **NBMA Neighbors**, and **Advanced** tabs for each IP address. Specify the following:<br><br>  • **Area ID**. The OSPF area ID for the interface. Select the one that you want to use from the provided list of available areas.<br><br>  • **Router priority**. The OSPF router priority for the interface. When two or more OSPF routers attached to a subnet that does not have an existing designated router attempt to become the designated router of the subnet, the router with the highest router priority setting takes precedence. If a tie exists, the router with the higher router ID takes precedence. An interface with a priority of 0 cannot become the designated router. The priority setting is ignored for point-to-point links.<br><br>  • **Cost**. The cost, advertised in the router's link state advertisement (LSA), of sending a packet out over this interface. Because an interface with a lower cost is used more frequently, typically, you configure a faster interface with a lower cost.<br><br>  • **Password**. The plaintext password for this OSPF area, which is enabled by default. All interfaces in the same OSPF area that are on the same subnet must use the same password.<br><br>Configuring the network type on this interface by selecting one of the following:<br><br>• **Broadcast**. Specifies that the interface is a broadcast OSPF interface on a multiple access broadcast network in which multiple routers can be connected to the network. Ethernet, Token Ring, and FDDI are examples of broadcast networks.<br><br>• **Point-to-point**. Specifies that this interface is a point-to-point OSPF interface that connects a single pair of routers. T1/E1, T3/E3, ISDN, and dial-up links are examples of point-to-point networks.<br><br>• **Non-broadcast multiple access (NBMA)**. Specifies that this interface is an NBMA OSPF interface. If you specify this type of interface, you can then configure neighboring routers by using the **NBMA Neighbors** tab. X.25, Frame Relay, and Asynchronous Transfer Mode (ATM) are examples of NBMA networks. |
| NBMA Neighbors | Configuring NBMA neighbors (available if **Non-broadcast multiple access (NBMA)** is selected as the network type on this interface's **General** tab) for this interface:<br><br>• **IP address**. Displays the IP address assigned to the interface. If two or more IP addresses are assigned to the interface, select the IP address that you want to configure from the list provided. If multiple IP addresses are configured for the interface, you must configure settings for each IP address.<br><br>• **Neighboring IP address**. Specifies the IP address (not the OSPF router ID) of a neighbor router on the NBMA network that you want to add.<br><br>• **Router priority**. Specifies the OSPF router priority of the neighbor router. A neighbor is eligible to become the designated router if its router priority is 1 or greater. |

| | |
|---|---|
| Advanced | Configuring advanced OSPF options for this interface:<br><br>• **IP address**. Displays the IP address assigned to the interface. If two or more IP addresses are assigned to the interface, select the IP address that you want to configure from the list provided. If multiple IP addresses are configured for the interface, you must configure settings for each IP address.<br><br>• **Transmit delay (seconds)**. The estimated number of seconds it takes for this OSPF router to transmit an OSPF link state update packet over this interface. This value should take into account the transmission and propagation delays of the interface and network media. The default setting is 1 second.<br><br>• **Retransmit interval (seconds)**. The number of seconds between link state advertisement (LSA) retransmissions on this interface. This value should exceed the expected round-trip delay between any two routers on the attached subnet; if this value is too low, needless retransmissions result. The value needs to be higher on low-speed serial lines. A sample value for a local area network (LAN) subnet is 5 seconds.<br><br>• **Hello interval (seconds)**. The interval in seconds between transmissions of OSPF Hello packets on this interface. This setting must be the same for all routers on the same subnet. The shorter the hello interval, the faster topological changes are detected. However, a shorter interval also results in more OSPF traffic. A sample value for an X.25 network is 30 seconds. A sample value for a LAN is 10 seconds.<br><br>• **Dead interval (seconds)**. The number of seconds before a neighbor router determines that this router is unavailable. The router is determined to be unavailable if a neighboring router does not receive a Hello packet sent by this router within the specified interval. For example, if the hello interval setting is 15 seconds and the dead interval is 60 seconds, after four unreceived Hello packets, neighbor routers declare that this router is unavailable. Typically, the value for this setting is four times the value in the hello interval. This value must be the same for all OSPF router interfaces attached to the same subnet.<br><br>• **Poll interval (seconds)**. For NBMA interfaces only, the number of seconds between OSPF network polls. If a neighboring router is unavailable, it might still be necessary to send Hello packets to the unavailable neighbor in order to determine if the connection has been restored. This is called the *poll interval*. Typically, the poll interval is at least twice as long as the dead interval. A sample value for a poll interval on an X.25 network is 2 minutes.<br><br>• **Maximum transmission unit (MTU) size (bytes)**. The maximum size of an IP datagram carrying OSPF information that can be sent without fragmentation. The default value of 1,500 bytes is the default IP MTU for an Ethernet network. |

**DHCP Relay Agent options**

Any IPv4 subnet that contains DHCP clients requires either a DHCP server or a DHCP relay agent to provide address leases to those DHCP clients. A computer acting as a DHCP relay agent, also known as a Bootstrap Protocol (BOOTP) relay agent, relays DHCP messages between DHCP clients on an IP subnet that has no DHCP server and a DHCP server located on a different subnet. By adding the DHCP relay agent to the Routing and Remote Access snap-in console tree, adding one or more DHCP relay agent interfaces to the details pane, and then configuring DHCP relay agent options, you can enable a router to act as a DCHP relay agent. Installation of a DHCP relay agent makes it unnecessary to install a separate DHCP server on each subnet in an IP internetwork.

This section includes unicast IPv4 routing options available in the console tree in the Routing and Remote Access snap-in on the **IP Routing\DHCP Relay Agent** node and in the details pane on the *InterfaceName* **Properties** page.

**DHCP relay agent options in the console tree**

The following table describes the option for adding one or more interfaces to the details pane. This option is available by right-clicking the **IP Routing\DHCP Relay Agent** node.

**Options on the IP Routing\DHCP Relay Agent Node in the Console Tree**

| Node | Task |
|---|---|

| | |
|---|---|
| DHCP Relay Agent | Adding a new interface to the details pane:<br>• **New Interface**. Provides a list of available interfaces on this server that you can add to the details pane of the **IP Routing\DHCP Relay Agent** node. |

The following table describes Routing and Remote Access options for DHCP relay agent properties. These settings are available in the console tree of the Routing and Remote Access snap-in on the properties page of the **IP Routing\DHCP Relay Agent** node.

**Options on the DHCP Relay Agent Properties Page**

| Tab | Task |
|---|---|
| General | Configuring one or more DHCP servers for local clients to contact through the DHCP relay agent:<br>• **Server address**. Adds the IP address of one or more DHCP servers to create a list of available DHCP servers on the network. This router acts as a DHCP relay agent by relaying DHCP messages between DHCP clients on the local subnet and the DHCP server or servers specified in this list. |

**DHCP relay agent options on the InterfaceName Properties page**

The following table describes Routing and Remote Access options for a DHCP relay agent interface. These settings are available in the Windows Server 2003 Routing and Remote Access snap-in on the **General** tab of the *InterfaceName* **Properties** page for an interface in the **IP Routing\DHCP Relay Agent** container.

**DHCP Relay Agent Options on the InterfaceName Properties Page**

| Tab | Task |
|---|---|
| General | Configuring a DHCP relay agent interface over which to route packets exchanged between DHCP clients on the local subnet and a DHCP server on another subnet:<br>• **Relay DHCP packets**. Enables DHCP messages to be relayed between DHCP clients on the subnet to which this interface is attached and a DHCP server on a different subnet. Enabled by default when you add an interface to the **IP Routing\DHCP Relay Agent** node in the console tree.<br>• **Hop-count threshold**. Specifies the maximum number of DHCP relay agents that can handle DHCP relayed traffic. The default value is 4 hops; the maximum value is 16 hops.<br>• **Boot threshold (seconds)**. Specifies the number of seconds the DHCP relay agent waits before forwarding DHCP messages. The default value is 4 seconds. This option is useful when you want a local DHCP server to respond first to a DHCP client's request for an IP address, but, if no local DHCP server responds, you want the DHCP relay agent to forward the message to a remote DHCP server. |

**IP packet filtering options**

Configuring IP packet filters on an IP router helps increase security on a private IP internetwork that connects to a public network such as the Internet. When IP packet filtering is configured, a packet arriving at a router might be forwarded to its destination, rejected with an error message returned to notify the sender, or dropped with no error message sent.

The following table describes options for IP packet filtering available in the Windows Server 2003 Routing and Remote Access snap-in by selecting **Inbound Filters** or **Outbound Filters** on the **General** tab of the *InterfaceName* **Properties** page for an interface in the **IP Routing\General** container, and then selecting **New** or **Edit**. Note that, although the options for adding an inbound or an outbound filter shown in the table appear together in the same dialog box, typically filters are configured in pairs, either {**Destination Network/Protocol**}, or {**Source Network/Protocol**}. After you add a filter, you select one of the following filter actions:

- Receive all packets except those that meet the criteria below.
- Drop all packets except those that meet the criteria below.

If you configure multiple parameters on a particular filter, when the filter is applied to an incoming packet, the parameters of the filter are compared by using a logical **AND** operator. The fields in the packet must match all of the criteria of the filter to have the filter action applied.

**IP Packet Filtering Options on the InterfaceName Properties Page**

| Tab | Task |
|---|---|
| General | Adding inbound filters or outbound filters (in the **Add IP Filter** or **Edit IP Filter** dialog box):<br><br>• **Source Network** field:<br><br>  • **IP address**. The source network ID or IP address.<br><br>  • **Subnet mask**. The subnet mask that corresponds to the source network ID, or 255.255.255.255 for a source IP address. The subnet mask bits must encompass all of the bits used in the **IP address** field. The IP address cannot be more specific than the subnet mask, that is, the IP address (in binary notation) cannot have a bit set to 1 when the corresponding bit in the subnet mask is set to 0.<br><br>• **Destination Network** field:<br><br>  • **IP address**. The destination network ID or IP address.<br><br>  • **Subnet mask**. The subnet mask that corresponds to the destination network ID, or 255.255.255.255 for a destination IP address. The subnet mask bits must encompass all of the bits being used in the **IP address** field. The IP address cannot be more specific than the subnet mask, that is, the IP address (in binary notation) cannot have a bit set to 1 when the corresponding bit in the subnet mask is set to 0.<br><br>• **Protocol** field:<br><br>  • **TCP**. Displays text boxes in which to specify a source TCP port and a destination TCP port. One or both can be specified. If nothing is specified in these text boxes, they default to **0**, which signifies any port.<br><br>  • **TCP [established]**. Displays text boxes in which to specify a source TCP port and a destination TCP port. Use this option when you want to define TCP traffic (source TCP port and destination TCP port) for TCP connections already established by the router.<br><br>  **TCP [established]** is used, for example, in a site-to-site VPN or dial-up connection when a calling router connects to an answering router. If you specify **TCP [established]**, TCP traffic is accepted only if the calling router initiates the TCP connection.<br><br>  • **UDP**. Displays text boxes in which to specify a source UDP port and a destination UDP port. One or both can be specified. If nothing is specified in these text boxes, they default to **0**, which signifies any port.<br><br>  • **ICMP**. Displays text boxes in which to specify an ICMP code and an ICMP type. One or both can be specified. If nothing is specified in these text boxes, they default to **255**, which signifies any code and any type.<br><br>  • **Any**. Configures a filter for any protocol (**Any** is the default).<br><br>  • **Other**. Lets you specify the IP protocol identifier number for the protocol you want (rather than select a protocol by name). For example, TCP is IP protocol 6, UDP is IP protocol 17, and ICMP is IP protocol 1.<br><br>**Note**<br>For a complete list of IP protocol numbers, which are used to configure firewalls, routers, and proxy servers, see the Internet Assigned Numbers Authority (IANA) list of <A HREF="http://go.microsoft.com/fwlink/?LinkID=3858" TARGET="_blank">IP protocol numbers</A> at http://go.microsoft.com/fwlink/?LinkID=3858. (The IPv4 “Protocol” field is called the “Next Header” field in IPv6.)<br><br>Configuring fragmentation checking:<br><br>• **Enable fragmentation checking**. If selected, specifies that this router discards all fragmented IP packets that it |

| | |
|---|---|
| | receives on this interface. This option applies only to incoming traffic. |

For examples of filter configurations for commonly implemented IP filtering scenarios on a Routing and Remote Access router, see "Filtering Scenarios" in How Unicast IPv4 Routing Protocols and Services Work.

**ICMP router discovery options**

Configuring ICMP router discovery, an alternative to using manually configured or DHCP-configured default gateways, provides an improved method of configuring and detecting the default gateway. For IP hosts, the default gateway is the IP router on the local subnet that forwards traffic to destinations that are not located on the local subnet.

The ICMP router discovery feature automates the discovery and configuration of the default gateway so that IP hosts can dynamically discover the best default gateway to use on a subnet and can automatically switch to another default gateway if the first default gateway fails or if the network administrator changes router preferences.

The following table describes Routing and Remote Access options for configuring ICMP router discovery on an IP router. These settings are available in the Windows Server 2003 Routing and Remote Access snap-in on the **General** tab of the *InterfaceName* **Properties** page for an interface in the **IP Routing\General** container.

**ICMP Router Discovery Options on the InterfaceName Properties Page**

| Tab | Task |
|---|---|
| General | Configuring options for ICMP router discovery on this interface:<br>• **Enable router discovery advertisements**. Specifies whether ICMP router discovery advertisements are enabled on the interface. If selected, you must also specify settings for the following (or accept the defaults):<br>  • **Advertisement lifetime (in minutes)**. The interval after which a host determines that a router is unavailable (since its last Router Advertisement message). The default is 30 minutes.<br>  • **Level of preference**. The preference level of this router to be the default gateway. The router with the highest preference level becomes the default gateway. The default is 0, which means that any other router configured with 1 or a higher number is preferred over this router by router discovery clients.<br>  • **Advertisement interval minimum time (in minutes)**. The minimum amount of time between Router Advertisement messages sent by this router. The default is 7 minutes.<br>  • **Advertisement interval maximum time (in minutes)**. The maximum amount of time between Router Advertisement messages sent by this router. The default is 10 minutes. Router Advertisements are sent at a random interval between the minimum and maximum times. |

In addition to enabling ICMP router discovery on the IP router, you must also configure router discovery on IP client computers on that router's subnet:

- For information about the tool used to configure ICMP router discovery on a client running Windows XP Professional, see "Netsh routing IP routerdiscovery context" in the section "Command-Line Tools for Unicast IPv4 Routing" earlier in this document.

- For information about the tool used to configure ICMP router discovery on a client running Windows 2000 Professional, see "PerformRouterDiscovery" and "SolicitationAddressBcast" in the section Unicast IPv4 Routing Registry Entries later in this document.

**Unicast IPv4 Routing Registry Entries**

When the Routing and Remote Access service is enabled, it creates and maintains its settings in the Windows Server 2003 registry. For performance reasons, most Routing and Remote Access service configuration information is stored in binary format in large configuration blocks, not as separate registry entries that you can easily view or modify by using the Registry Editor (Regedit.exe) tool. This section includes only unicast IPv4 routing-related Routing and Remote Access registry entries that are configurable.

For more information about these and other registry entries, see the Windows Server 2003 Resource Kit Registry Reference.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It

is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to change configuration settings rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing**

The Routing and Remote Access service provides an extensive tracing capability to help troubleshoot complex network problems. Separate routing components can be independently enabled to log tracing information to files — this capability is called file tracing. File tracing includes information about internal component variables, function calls, and interactions. You can enable and disable tracing for routing protocols while the router is running. This section includes only file tracing registry keys for the RIP v2 and OSPF routing protocols.

The following registry keys are located under the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing.

**IPRIP2**

**Registry path** HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing

**Version** Windows Server 2003 and Windows 2000 Server

Configuring entries under the IPRIP2 key enables file tracing, which logs tracing information about the RIP v2 routing protocol to files so that you can troubleshoot problems on a RIP internetwork.

| **Caution** |
|---|
| The recommended method for enabling file tracing for RIP is to use the Netsh tool rather than by changing the registry settings shown here. For more information, see “Netsh ras context” earlier in this document. |

Tracing consumes system resources and must be used sparingly. After you obtain the trace or identify the problem, immediately disable tracing. In particular, do not leave tracing enabled on multiprocessor computers.

The following table lists the registry entries for configuring file tracing for RIP v2.

**Entries for Enabling Tracing Under the IPRIP2 Registry Key**

| Entry | Task |
|---|---|
| EnableFileTracing | Configuring file tracing:<br>• 1 enables file tracing for the RIP v2 routing protocol.<br>• 0 disables file tracing for the RIP v2 routing protocol. File tracing is disabled by default. |
| FileDirectory | Configuring location of the file tracing log file:<br>• By default, the location of the tracing file is set to %windir%\Tracing\IPRIP2.LOG file.<br>• To set a new default location, you can right-click this entry and then specify an alternative for %windir%\Tracing. You cannot change the log file name. |
| FileTracingMask | Configuring how much tracing information is saved in the log file:<br>• By default, the level for file tracing is set to 0xFFFF0000, which is the maximum level.<br>• To set a new default level for file tracing, you can right-click this entry and then specify an alternative level in the range from 0 to 0xFFFF0000. |

| | |
|---|---|
| MaxFileSize | Configuring maximum log file size:<br>• By default, the maximum size of the log file is set to 0x100000, which is 1 MB.<br>• To set a new maximum size for the log file, you can right-click this entry and then specify an alternative maximum size. |

**OSPF**

**Registry path** HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing

**Version** Windows Server 2003 and Windows 2000 Server

Configuring entries under the OSPF key enables file tracing, which logs tracing information about the OSPF routing protocol to files so that you can troubleshoot problems on an OSPF internetwork.

| **Caution** |
|---|
| The recommended method for enabling file tracing for OSPF is to use the Netsh tool rather than by changing the registry settings shown here. For more information, see "Netsh ras context" earlier in this document. |

Tracing consumes system resources and must be used sparingly. After you obtain the trace or identify the problem, immediately disable tracing. In particular, do not leave tracing enabled on multiprocessor computers.

The following table lists the registry entries for configuring file tracing for OSPF.

**Entries for Enabling Tracing Under the OSPF Registry Key**

| Entry | Task |
|---|---|
| EnableFileTracing | Configuring file tracing:<br>• 1 enables file tracing for the OSPF routing protocol.<br>• 0 disables file tracing for the OSPF routing protocol. |
| FileDirectory | Configuring location of the file tracing log file:<br>• By default, the location of the tracing file is set to %windir%\Tracing\OSPF.LOG file.<br>• To set a new default location, you can right-click this entry and then specify an alternative for %windir%\Tracing. You cannot change the log file name. |
| FileTracingMask | Configuring how much tracing information is saved in the log file:<br>• By default, the level for file tracing is set to 0xFFFF0000, which is the maximum level.<br>• To set a new default level for file tracing, you can right-click this entry and then specify an alternative level in the range from 0 to 0xFFFF0000. |
| MaxFileSize | Configuring maximum log file size:<br>• By default, the maximum size of the log file is set to 0x100000, which is 1 MB.<br>• To set a new maximum size for the log file, you can right-click this entry and then specify an alternative maximum size. |

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Tcpip\Parameters**

The Routing and Remote Access service uses several TCP/IP registry keys. The entries included in this section are those that are specific to unicast IPv4 routing. For additional entries for TCP/IP, see TCP/IP Tools and Settings.

The following registry entries are located under the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters.

**DisableIPSourceRouting**

**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version** Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0, or Windows 98

This entry determines whether IP source routing is enabled. IP source routing lets the sender use a tool, such as Tracert.exe or Ping.exe, to specify the exact route that an IP packet takes through successive IP routers across the IP internetwork.

Source routing is not typically implemented on IP internetworks. Instead, IP routing is typically handled through routing decisions made by source hosts and IP routers based on entries in their respective local routing tables. Although you might want to use IP source routing in network testing and debugging situations, disable IP source routing as soon as you are finished testing or troubleshooting.

**Values for DisableIPSourceRouting**

| Value | Setting |
|---|---|
| 0 | Forwards all packets. |
| 1 (Default) | Does not forward source routed packets. |
| 2 | Drops all incoming source routed packets. |

**EnableDeadGWDetect**

**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version** Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0, or Windows 98

This entry determines whether dead gateway detection is enabled. With **EnableDeadGWDetect** enabled, TCP can detect a failure of the default gateway and inform IP to switch to a different default gateway if one is available. A default gateway is an IP router on the local subnet that can forward traffic to destinations located beyond the local subnet.

For optimal security, setting **EnableDeadGWDetect** to **0** is recommended in order to prevent an attacker from forcing this computer to switch to an unintended default gateway.

**Values for EnableDeadGWDetect**

| Value | Setting |
|---|---|
| 0 | Disables dead gateway detection. |
| 1 (default) | Enables dead gateway detection. |

**EnableFastRouteLookup**

**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version** Windows Server 2003, Windows 2000 Server

This entry, which is created only if the Routing and Remote Access Service is enabled, determines whether the fast route lookup feature is enabled. Fast route lookups make route searches faster but use a significant amount of non-paged pool memory. Fast route lookups are used only if the computer runs Windows Server 2003 or Windows 2000 Server and contains at least 64 MB of memory.

**Values for EnableFastRouteLookup**

| Value | Setting |
|---|---|
| 0 (default) | Disables fast route lookup. |
| 1 | Enables fast route lookup. |

**EnableICMPRedirects**

**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version** Windows Server 2003, Windows XP, Windows 2000

This entry determines whether a computer running Windows Server 2003 TCP/IP, or another version of Windows TCP/IP that supports Internet Control Message Protocol (ICMP) Redirect messages, updates its routing table in response to ICMP Redirect messages sent to it by a network device such as a router. An ICMP Redirect message instructs the computer to direct IP packets for the recipient along a different route. A computer running Windows Server 2003 accepts ICMP redirection messages from

any node in the route between it and the destination computer.

You must restart the computer for changes to this entry to take effect.

**Values for EnableICMPRedirects**

| Value | Setting |
|---|---|
| 0 | Disables updates of the IP routing table in response to ICMP Redirect messages. |
| 1 (default) | Enables updates of the IP routing table in response to ICMP Redirect messages. |

**EnablePMTUBHDetect**

**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version** Windows Server 2003, Windows XP, Windows 2000

This entry determines whether TCP tries to detect black hole routers during the path maximum transmission unit (PMTU) discovery process. For performance reasons, PMTU black hole router detection is disabled by default for TCP/IP in Windows Server 2003, Windows XP, and Windows 2000.

When PMTU black hole router detection is enabled, if several retransmissions of a TCP segment are not acknowledged, TCP tries sending segments with the Don't Fragment flag in the IP header of the IP packet set to **0** (which enables an IP router to fragment the TCP segment). If, as a result, a segment with the Don't Fragment flag set to **0** is acknowledged, the TCP maximum segment size (MSS) is decreased and the Don't Fragment flag is set to **1** (which prevents an IP router from fragmenting the TCP segment) in subsequent segments on the connection.

Enabling PMTU black hole detection increases the maximum number of retransmissions that are performed for a specific segment and can therefore decrease overall performance.

This registry entry is not present by default and must be added by using the Registry Editor tool, and then restarting Windows.

**Values for EnablePMTUBHDetect**

| Value | Setting |
|---|---|
| 0 (default) | Disables PMTU black hole detection. |
| 1 | Enables PMTU black hole detection. |

**EnablePMTUDiscovery**

**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version** Windows Server 2003, Windows XP, Windows 2000

This entry determines whether TCP uses a fixed, default maximum transmission unit (MTU) of 576 bytes for all connections that are not to destinations on a locally attached subnet, or alternatively, whether TCP attempts to detect the actual path maximum transmission unit (PMTU).

When PMTU Discovery is enabled, TCP attempts to discover the MTU, or largest packet size, allowed by the path to a remote host. By discovering the PMTU and limiting TCP segments to this size, TCP can eliminate fragmentation at IP routers along the path that connect networks with different MTUs. Fragmentation reduces TCP throughput and increases network congestion.

By default, this entry applies to all interfaces. However, you can use the Registry Editor to change the default value of this entry for one or more interfaces on this computer.

For optimal security, set **EnablePMTUDiscovery** to **0** in order to prevent an attacker from forcing the MTU value to a very small value and thereby overworking the TCP/IP stack.

This registry entry is not present by default and must be added by using the Registry Editor tool, and then restarting Windows.

**Values for EnablePMTUDiscovery**

| Value | Setting |
|---|---|
| 0 | Causes TCP to use an MTU of 576 bytes for all connections to nodes outside of the local subnet. |
| 1 (default) | Enables TCP to attempt to discover the actual MTU of the path to a remote node. |

**IPEnableRouter**

**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version** Windows Server 2003, Windows XP, Windows 2000

This entry determines whether IP forwarding is enabled on a computer on which the Routing and Remote Access service is not enabled. As the name **IPEnableRouter** indicates, manually enabling this entry enables the computer to act as a router. That is, when IP forwarding is enabled, the computer can route IP packets that it receives to nodes on all networks to which it is connected.

**Values for IPEnableRouter**

| Value | Setting |
|---|---|
| 0 (default) | Disables IP forwarding. The host does not route IP packets. |
| 1 | Enables IP forwarding. The host routes IP packets. |

A server running the Routing and Remote Access service does not use the **IPEnableRouter** registry entry; this entry is therefore, correctly, set to 0.

If you want to enable IP forwarding on a workstation or on a server on which you do not want to enable the Routing and Remote Access service, you must use the Registry Editor to set the **IPEnableRouter** registry entry to **1**. Typically, a workstation or multi-homed server that is not a router should not be configured to forward packets between networks, unless it is acting as a firewall.

**IPEnableRouterBackup**

**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version** Windows Server 2003, Windows XP, Windows 2000

This entry stores a backup copy of the value that was stored in the **IPEnableRouter** entry when the operating system of the computer was most recently installed or upgraded. The computer uses this entry to:

- Determine whether the value of **IPEnableRouter** has changed
- Restore the previous value if the value of **IPEnableRouter** has changed

Do not delete this entry or change its value. If you do, you might disable the IP packet routing mechanism.

There is no default value for this entry. The computer enters a value when you install or upgrade Windows.

**Values for IPEnableRouterBackup**

| Value | Setting |
|---|---|
| 0 | Disables storing a backup copy of the value stored in **IPEnableRouter**. |
| 1 | Enables storing a backup copy of the value stored in **IPEnableRouter**. |

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\interfaceGUID**

The Routing and Remote Access service uses several TCP/IP registry keys for interfaces. The entries included in this section are specific to unicast IPv4 routing features supported by the Routing and Remote Access service. For additional entries for TCP/IP interfaces, see TCP/IP Tools and Settings.

The following registry entries are located under the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interfaceGUID*.

Before you can use the registry editor to configure a registry entry located under this key, you must first determine the globally unique identifier (GUID) for the TCP/IP interface that you want to configure.

**To determine the GUID for a specific TCP/IP interface**

1. Open Network Connections, and then note the name of the LAN connection, such as "Local Area Connection."
2. Open the Registry Editor (click **Start**, click **Run**, type **regedit.exe**, and then click **OK**).
3. In the tree view (the left pane), expand the following key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Network\{4D36E972-E325-11CE-BFC1-08002BE10318}
4. Under this key are one or more keys for the globally unique identifiers (GUIDs) corresponding to the installed LAN connections on this computer. Each of these GUID keys has a Connection subkey. Open each of the *GUID*\Connection keys, and then look for the Name setting whose value matches the name of your LAN connection from step 1.

5. When you have found the *GUID*\Connection key that contains the Name setting that matches the name of your LAN connection, write down, take a screenshot, or otherwise note the GUID value.

6. Use the tree view to open the following key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interfaceGUID*.

7. Right-click the *interfaceGUID* key you want in the tree view, and then view or configure that entry.

**PerformRouterDiscovery**
**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interfaceGUID*

**Version** Windows Server 2003, Windows XP, Windows 2000, Windows 98

| **Note** |
|---|
| The recommended method for configuring router discovery on a host running Windows Server 2003 or Windows XP is to set DHCP options, instead of modifying the registry. |

The **PerformRouterDiscovery** entry determines whether Windows TCP/IP attempts to perform ICMP router discovery as a host, as defined by RFC 1256, on this interface. Enabling router discovery on a client or non-router server enables that computer to listen for Router Advertisement messages from an IP router, to send Router Solicitation messages to discover active routers on its subnet, and to switch to another router if the first router fails or the network administrator changes router preferences.

Windows does not add the **PerformRouterDiscovery** entry to the registry. To enable ICMP router discovery, you must add this entry (as type REG_DWORD) and set its value to one of the options listed in the following table.

**Values for Enabling PerformRouterDiscovery**

| Value | Setting |
|---|---|
| 0 | Disables router discovery. |
| 1 | Enables router discovery. |
| 2 (default) | Enables DHCP-controlled router discovery.<br><br>The value 2 is the default, but it is used only if the DHCP options **Perform Router Discovery** and **Router Solicitation Address** are configured. The DHCP **Perform Router Discovery** option is DHCP scope option 31 and the DHCP **Router Solicitation Address** option is DHCP scope option 32. Both are configured by using the DHCP snap-in. For more information, see "DHCP Options" in DHCP Tools and Settings. |

For ICMP router discovery to take place, in addition to configuring **PerformRouterDiscovery** on a client or non-router server, the following requirements must also be met:

- A Routing and Remote Access router on the same subnet must be configured for ICMP router discovery, to enable it to send Router Advertisement messages to all hosts on the subnet. For more information, see "Netsh routing IP routerdiscovery context" or "ICMP router discovery options" earlier in this technical reference.

- The entry **SolicitationAddressBcast** must be added to the registry, and its value configured. For more information, see "SolicitationAddressBcast" later in this technical reference.

**PerformRouterDiscoveryBackup**
**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interfaceGUID*

**Version** Windows Server 2003, Windows XP, Windows 2000

This entry stores a backup copy of the value stored in the **PerformRouterDiscovery** entry (if an entry for **PerformRouterDiscovery** has been configured). The computer uses this entry only when IP routing is enabled and the value of **PerformRouterDiscovery** is **0**.

| **Caution** |
|---|
| This entry is used internally by the Windows operating system. Do not delete this entry or change its value. If you do, you might disable the router discovery mechanism. |

There is no default value for this entry.

**Values for PerformRouterDiscoveryBackup**

| Value | Setting |
|---|---|
| 0 | Disables storing a backup copy of the value stored in **PerformRouterDiscovery**. |
| 1 | Enables storing a backup copy of the value stored in **PerformRouterDiscovery**. |

**SolicitationAddressBcast**

**Registry path** HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces*interfaceGUID*

**Version** Windows Server 2003, Windows XP, Windows 2000, Windows 98

| **Note** |
|---|
| The recommended method for configuring router discovery on a host running Windows Server 2003 or Windows XP is to set DHCP options, instead of modifying the registry. |

The **SolicitationAddressBcast** entry, as defined by RFC 1256, determines whether Windows sends router discovery messages as broadcasts rather than as multicasts on this interface. Enabling router discovery on a client or non-router server computer lets that computer listen for **Router Advertisement** messages from a router and lets that computer send **Router Solicitation** messages to discover active routers on its subnet and to discover the unavailability of previously available routers that are currently down.

Windows does not add the **SolicitationAddressBcast** entry to the registry. To specify how router solicitation messages are sent, you must add this entry (as type REG_DWORD) and set its value to one of the options listed in the following table.

**Values for Enabling SolicitationAddressBcast**

| Value | Setting |
|---|---|
| 0 (Default) | Enables the sending of router solicitation messages to the all-routers multicast group (224.0.0.2). This is the recommended option. |
| 1 | Enables the sending of router solicitation messages as limited broadcast messages by using the IP address 255.255.255.255.<br><br>The value 1 is not recommended because it causes all clients to broadcast router solicitation messages, creating a large volume of unnecessary broadcast traffic that must be processed by other clients. |

For ICMP router discovery to take place, in addition to configuring the value for this entry, the following requirements must also be met:

- A Routing and Remote Access router on the same subnet must be configured for ICMP router discovery, to enable it to send Router Advertisement messages to all hosts on the subnet. For more information, see "Netsh routing IP routerdiscovery context" or "ICMP router discovery options" earlier in this technical reference.

- The entry **PerformRouterDiscovery** must be added to the registry, and its value configured. For more information, see "PerformRouterDiscovery" earlier in this technical reference.

Note that **PerformRouterDiscovery** option 2 enables DHCP-controlled router discovery, but only if the DHCP options **Perform Router Discovery** and **Router Solicitation Address** are configured.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Command-Line Reference for Windows Server 2003, Standard Edition

- Windows Server 2003 Resource Kit Registry Reference

- "Parameters Configurable Using the Registry Editor" in Microsoft Windows Server 2003 TCP/IP Implementation Details

- IETF RFC Database

- How Unicast IPv4 Routing Protocols and Services Work
- TCP/IP Tools and Settings
- IPv6 Tools and Settings
- Demand Dial Tools and Settings
- IPv4 Multicasting Tools and Settings
- NAT Tools and Settings
- VPN Tools and Settings
- DHCP Tools and Settings

# Network Monitoring Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Monitoring Technologies**
In Microsoft Windows Server 2003 network monitoring technologies help network administrators track and manage the functions of devices on the network.

**Network Monitoring Components**
There are two network monitoring technologies that help network administrators track and manage the functions of devices on the network:

- Simple Network Management Protocol (SNMP)

  SNMP is a network management protocol that can monitor and alert administrators regarding the condition of hardware components such as workstations and servers.

- Microsoft Computer Browser Service

  The Computer Browser Service in computers running Windows enables network administrators and users to browse and gain access to the resources of the network. The browser service provides a convenient listing of devices such as workstations and printers, and allows network users access to them with the proper authentication.

**SNMP**
SNMP management software is used to monitor any network device on which SNMP agent software is installed. The SNMP agent, an optional component of Windows Server 2003, interacts with other SNMP management software to enable the transport of information about network status between monitored devices and applications and the management systems that monitor them.

**SNMP Management Systems and Agents**
SNMP provides a method of managing network hosts such as workstation or server computers, routers, bridges, and hubs from a centrally-located computer that runs network management software. SNMP performs management services by using a distributed architecture of management systems and agents.

**SNMP management system**
The management system, also called management console, sends information and update requests to an SNMP agent. Any computer running SNMP management software is an SNMP management system. The management software application does not need to run on the same host as the SNMP agent.

The SNMP management system requests information from a managed computer, (SNMP agent), such as the amount of hard disk space available or the number of active sessions. The management system can also initiate a change to an agent's configuration. However, this is rare because most clients have read-only access.

**SNMP agent**
The SNMP agent responds to management system requests for information. Any computer running SNMP agent software is an SNMP agent. The Windows Server 2003 SNMP service, which is agent software, responds to information requests from one or multiple management systems. The SNMP service can be configured to determine which statistics are tracked and which management systems are authorized to request information.

Typically, agents do not originate messages, but only respond to them. A trap message is the only agent-initiated SNMP communication. A trap is an alarm-triggering event on an agent, such as a system reboot or illegal access, which provides enhanced security.

Management hosts and agents belong to an SNMP community, which is a collection of hosts grouped together for administrative purposes. Defining communities provides security by allowing only management systems and agents within the same community to communicate.

**Computer Browser Service**
The Computer Browser service is the mechanism that collects and distributes the list of workgroups and domains and the servers within them. The list is displayed in the window of Microsoft Windows Network and related windows in My Network

Places.

The architecture of the browser service is complete when several computers are brought together on a network to perform Browser Service roles. Individually, these roles are just parts of the Browser Service, but when combined on the network, they form a browser system that enables efficient sharing of network resource information. The computers that perform these roles make up the components of the browser system.

**Browser System Roles**
There are four Browser Service roles that computers can play in the browser system:

- Potential browser
- Backup browse server
- Master browse server
- Domain master browse server

**Network Monitoring Scenarios**
The need for network monitoring corresponds to the size of the network. Small networks with few devices attached and with a small geographical size might be easily managed without SNMP or a large reliance on the Computer Browser Service. A large network with many devices is more difficult to monitor. Unlike a small network where administrators might be able to visually inspect network devices regularly, large networks might require the use of SNMP and the Computer Browser Service to make it possible to track network devices from a single computer. Thus, these technologies become more useful and more essential as networks expand.

# SNMP Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SNMP Technical Reference**
Simple Network Management Protocol (SNMP) lets you monitor network devices configured as SNMP agents. Microsoft Windows Server 2003 provides SNMP agent software that works with third-party SNMP management software to monitor the status of managed devices and applications. Use SNMP to monitor network performance, audit network usage, detect network faults or inappropriate access, and, in some cases, configure remote devices.

**In this subject**


- What Is SNMP?
- How SNMP Works

# What Is SNMP?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is SNMP?**

**In this section**


- Automating Network Management
- Related Information


Simple Network Management Protocol (SNMP) is a popular protocol for network management. It is used for collecting information from, and configuring, network devices, such as servers, printers, hubs, switches, and routers on an Internet Protocol (IP) network. Microsoft Windows Server 2003 provides SNMP agent software that works with third-party SNMP management software to monitor the status of managed devices and applications.

**Automating Network Management**

Large networks with hundreds or thousands of nodes are difficult to manage without a large staff to monitor every computer. SNMP, which is widely used in local area networks (LANs), lets you monitor network nodes from a management host. You can monitor network devices such as servers, workstations, printers, routers, bridges, and hubs, as well as services such as Dynamic Host Configuration Protocol (DHCP) or Windows Internet Name Service (WINS).

Use SNMP management software to monitor any network device on which you install SNMP agent software. The SNMP agent, which is an optional component of Windows Server 2003, interacts with third-party SNMP management software to enable the sharing of network status information between monitored devices and applications and the SNMP management system that monitors them. The following figure shows this process.

**Computer Running SNMP Manager Requests Status Information from Devices Running SNMP Agent**




Using SNMP, you can monitor network performance, audit network usage, detect network faults or inappropriate access, and in some cases configure remote devices. SNMP is designed to be deployed on the largest possible number of network devices, to have minimal impact on the managed nodes, to have minimal transport requirements, and to continue working when most other network applications fail.

**SNMP Background**

Developed in 1988 to provide network-device-monitoring capability for TCP/IP-based networks, SNMP was approved as an Internet standard in 1990 by the Internet Architecture Board (IAB) and has been in wide use since that time. More recently, Internetwork Packet Exchange (IPX)-based networks have added support for SNMP. Currently, most network equipment vendors provide SNMP support in their products.

Windows Server 2003 supports SNMP versions 1 and 2c, which are based on the industry standards described in the following table.

**Industry Standards for SNMP Versions 1 and 2c**

| Standard | RFCs |
|---|---|

| | |
|---|---|
| SNMP | RFC 1157, "A Simple Network Management Protocol (SNMP)," a standard that defines how communication occurs between SNMP-capable devices and defines the SNMP message types.<br><br>Windows Server 2003 also supports SNMPv2c, which is defined in RFC 1901, "Introduction to Community-Based SNMPv2," RFC 1905, "Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)," and RFC 1906, "Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)." SNMPv2c updates protocol operations by introducing the **GetBulk** request and authentication based on community names.<br><br>RFC 1420, "SNMP over IPX" describes SNMP support now provided by the Internetwork Packet Exchange (IPX) protocol used in Novell NetWare environments. SNMP typically uses the User Datagram Protocol (UDP). SNMP over IPX is intended for use in environments in which the UDP transport is not available. IPX provides a datagram service similar to that provided by IP/UDP. |
| MIB | RFC 1213, "Management Information Base for Network Management of TCP/IP based Internets: MIB II." A standard that defines a set of objects that represent information about IP and IPX components on your network, such as the list of network interfaces, the routing table, the Address Resolution Protocol (ARP) table, the list of open TCP connections, or Internet Control Message Protocol (ICMP) statistics.<br><br>RFC 1213 is updated (but not made obsolete) by RFC 2011, "SNMPv2 Management Information Base for the Internet Protocol using SMIv2," RFC 2012, "SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2," and RFC 2013, "SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2." |
| SMI | RFC 2578, "Structure of Management Information Version 2 (SMIv2)." A standard that describes the object syntax for specifying how MIB data is referenced and stored. |

For a list of additional SNMP-related RFCs, see "Related Information" below.

**Implementing Windows Server 2003 and SNMP**

SNMP is a service that you can install on computers that are running TCP/IP or the IPX protocol. Windows Server 2003 supports SNMPv1 and SNMPv2c, which define how network management information is structured, stored, and communicated between SNMP agents and SNMP managers on IP-based or IPX-based networks. Windows Server 2003 adds support for the IPv6 Management Information Base (MIB) for SNMP.

**SNMP agent**

The Windows Server 2003 SNMP service provides an SNMP agent that enables the remote monitoring of computers running Windows Server 2003, Windows XP, or Windows 2000. The Windows XP 64-Bit Edition operating system and the 64-bit version of Windows Server 2003 do not support the Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) protocol.

The Windows Server 2003 SNMP agent lets you manage the services listed in the following table.

**Services That Can Be Monitored or Configured Using SNMP Agent**

| Service | Monitor | Configure | Windows Operating Systems That Provide SNMP Support for This Service Version |
|---|---|---|---|
| WINS | • | • | Windows Server 2003, Windows XP, and Windows 2000. |
| DHCP | • | | Windows Server 2003, Windows XP, Windows 2000, and Windows NT-based Dynamic Host Configuration Protocol (DHCP). |
| IIS | • | | Windows Server 2003, Windows XP, Windows 2000, and Windows NT-based Microsoft Internet Information Services (IIS). |
| LAN Manager | • | • | Microsoft LAN Manager. |
| Routing and Remote Access | • | | Windows Server 2003, and Windows 2000 Routing and Remote Access service. |
| IAS | • | • | Windows Server 2003, and Windows 2000 Internet Authentication Service (IAS). |

| | | | |
|---|---|---|---|
| System Monitor counters | • | | Windows Server 2003, Windows XP, and Windows 2000 System Monitor counters.<br><br>You can use SNMP to monitor all performance counters related to System Monitor TCP/IP, including ICMP, IP, Network Interface, TCP, UDP, DHCP, File Transfer Protocol (FTP), WINS, and IIS performance counters. |

In addition to retrieving status information, SNMP can send a **Set** request to configure any object to which the SNMP manager has read/write permission.

**SNMP manager**

The Windows Server 2003 SNMP service provides only the SNMP agent. Unless you develop your own SNMP management application, you must install third-party SNMP management software, such as HP Openview, Novell NMS, IBM NetView, or Sun Net Manager to work with the SNMP agent. You can install third-party SNMP management software on one or more hosts.

Alternatively, you can develop your own SNMP management software application by using the two application programming interfaces (APIs) that Windows Server 2003 provides:

WinSNMP API (Wsnmp32.dll) provides a set of functions for encoding, decoding, sending, and receiving SNMP messages.

Management API (Mgmtapi.dll) provides a basic set of functions for developing fast and simple SNMP management systems.

**Related Information**

For more information about SNMP, see the following RFCs in the IETF RFC Database.

The following RFCs relate to SNMP version 1:

RFC 1157, “Simple Network Management Protocol (SNMP).”

RFC 1155, “Structure and Identification of Management Information for TCP/IP-based Internets.”

RFC 1213, “Management Information Base for Network Management of TCP/IP-based Internets: MIB-II.”

RFC 1573, “Evolution of the Interfaces Group of MIB-II.”

The following RFCs relate to SNMP version 2:

RFC 1901, “Introduction to Community-Based SNMPv2.”

RFC 1902, “Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2).”

RFC 1903, “Textual Conventions for SNMPv2.”

RFC 1904, “Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2).”

RFC 1905, “Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2).”

RFC 1906, “Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2).”

RFC 1907, “Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2).”

RFC 1908, “Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework.”

RFC 2089, “V2ToV1 Mapping SNMPv2 onto SNMPv1 Within a Bi-Lingual SNMP Agent.”

RFC 2465, “Management Information Base for IP Version 6: Textual Conventions and General Group.”

RFC 2466, “Management Information Base for IP Version 6: ICMPv6 Group.”

# How SNMP Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How SNMP Works**
**In this section**


- SNMP Architecture
- SNMP Manager/Agent Communication
- SNMP Management Systems and Agents
- SNMP Messages
- SNMP Communities
- Related Information


Use Simple Network Management Protocol (SNMP) management software to monitor any network device configured with SNMP agent software. The SNMP agent, which is an optional component of Windows Server 2003, interacts with third-party SNMP management software to enable the flow of network status information between monitored devices and applications and the management systems that monitor them.

You can use SNMP in environments that include large networks with hundreds or thousands of nodes that would otherwise be difficult and costly to monitor. SNMP allows monitoring of network devices such as servers, workstations, printers, routers, bridges, and hubs, as well as services such as Dynamic Host Configuration Protocol (DHCP) or Windows Internet Name Service (WINS).

The following sections describe the architecture, components, and processes used by SNMP.

**SNMP Architecture**
To perform its monitoring services, SNMP uses a distributed architecture of management systems and agents and several related components. Windows Server 2003 provides an SNMP agent that is designed to be capable of interacting with any SNMP manager. The following components are the building blocks of SNMP and the Windows Server 2003 SNMP agent:

- SNMP management systems and agents
- Management Information Base (MIB)
- SNMP Messages
- SNMP Communities
- The communication process between SNMP managers and agents

The internal architecture of the Windows Server 2003 implementation of SNMP is divided into management and agent functions, which, in some cases, overlap. The following figure illustrates how the Windows Server 2003 SNMP structure fits into the layers of the underlying TCP/IP protocol architecture.

**Windows Server 2003 SNMP Components at Each TCP/IP Layer**




The following table lists the Windows Server 2003 SNMP internal components.

**Windows Server 2003 SNMP Components**

| Component Name | Associated Programs | Component Type | Description |
|---|---|---|---|
| Microsoft SNMP Service<br><br>(Also called master agent or extendible agent) | Snmp.exe | Agent | Receives SNMP requests and delivers them to the appropriate SNMP subagent DLL for processing. The service is also responsible for intercepting events (traps) from the SNMP subagents and forwarding trap messages to the appropriate management systems. |
| SNMP Subagents<br><br>(Also called extension agents) | Inetmib1.dll, Hostmib.dll, Lmmib2.dll, and others | Agent | Provides a set of entry points. When an SNMP request is received, the SNMP service delivers it to the appropriate subagent by calling one of these entry points. After the subagent processes the message, it passes the information back to the SNMP service, which then forwards the message to the SNMP manager. |
| SNMP Utility API | Snmpapi.dll | Both Agent and Manager | Provides utilities that the SNMP service uses for memory management operations, address-decoding routines, object identifier handling routines, and so forth.<br><br>Provides a set of routines that SNMP subagents use to handle and order SNMP objects.<br><br>Although use of the Snmpapi.dll is not required, the framework defined by this tool greatly facilitates the development of new SNMP subagents. |
| WinSNMP API<br><br>-and-<br><br>Management API | Wsnmp32.dll<br><br>-and-<br><br>Mgmtapi.dll | Manager | Helps develop SNMP management software applications:<br><br>• The WinSNMP API provides a set of functions for encoding, decoding, sending, and receiving SNMP messages.<br><br>• The Management API is a simple API that resides on top of the WinSNMP and SNMP Utility APIs. It provides a limited set of functions that you can use to develop basic SNMP management applications quickly. |
| SNMP Trap Service | Snmptrap.exe | Manager | Uses the WinSNMP API to forward a trap message sent by an SNMP agent to the appropriate SNMP manager application. |

| | | | |
|---|---|---|---|
| SNMP Manager Application | Snmputil.exe | Manager | Provides a basic command-line utility with which to retrieve information from any SNMP agent in your network. This example of a management application was developed using the Management API. |

**SNMP Manager/Agent Communication**
The following figure shows an example of the communication process between SNMP managers and agents.

**Manager Requesting Number of Sessions from an SNMP Agent**




The communication between the SNMP manager and SNMP agent in the example proceeds in the following fashion:

1. The SNMP manager, Host A, forms an SNMP message that contains an information request (Get) for the number of active sessions, the name of the community to which the SNMP manager belongs, and the destination of the message — the IP address (131.107.3.24) of the SNMP agent, Host B.

   The SNMP manager can use either the Microsoft SNMP Management API library (Mgmtapi.dll) or the Microsoft WinSNMP API library (Wsnmp32.dll) to perform this step.

2. The SNMP manager sends the information request to Host B by using the SNMP service libraries.

3. When Host B receives the message, it verifies that the community name (MonitorInfo) contained in the packet is on its list of acceptable community names, evaluates the request against the agent's list of access permissions for that community, and verifies the source IP address.

   If the community name or access permission is incorrect, and the SNMP service has been configured to send an authentication trap, the agent sends an "authentication failure" trap to the specified trap destination, Host C. Hosts B and C belong to the TrapAlarm community.

4. The master agent component of the SNMP agent calls the appropriate extension agent to retrieve the requested session information from the MIB.

5. Using the session information that it retrieved from the extension agent, the SNMP service forms a return SNMP message that contains the number of active sessions and the destination — the IP address (131.107.7.29) of the SNMP manager, Host A.

6. Host B sends the response to Host A.

**SNMP Management Systems and Agents**
SNMP locates the network management component on one or more computers and locates the managed component on multiple managed devices:

- **SNMP manager.** An SNMP manager, also known as an SNMP management system or a management console, is any computer that sends queries for IP-related information to a managed computer, known as an SNMP agent. In some

cases, the SNMP manager can send a request to an SNMP agent to change a configuration value.

- **SNMP agent.** An SNMP agent is any computer or other network device that monitors and responds to queries from SNMP managers. The agent can also send a trap message to the manager when specified events, such as a system reboot or illegal access, occur.

A computer on which you install SNMP management software is an SNMP manager, and a computer on which you install agent software, such as the Microsoft SNMP agent included with Windows Server 2003, is an SNMP agent. The SNMP manager displays the information it receives in a user-friendly graphical user interface. You configure SNMP options, including traps, on the SNMP agent, but the SNMP agent does not display the managed information that it sends to an SNMP manager. For more information about SNMP requests and trap messages, see "SNMP Messages" later in this section.

To enable SNMP communications between an SNMP manager and SNMP agents, you configure the SNMP manager and the SNMP agents that it manages as members of an SNMP community. The community name functions like a password to authenticate communications between the SNMP manager and agent. The SNMP community is an SNMP-defined group, not a group defined in the Active Directory directory service. For more information about SNMP communities, see "SNMP Communities" later in this section.

An SNMP manager can request the following types of information from the SNMP agents that it monitors:

- Network protocol identification and statistics
- Dynamic identification (discovery) of devices attached to the network
- Hardware and software configuration data
- Device performance and usage statistics
- Device error and event messages
- Program and application usage statistics

If you assign the SNMP manager write permission for the SNMP agent, the SNMP manager can also send a configuration request to the agent (using a **Set** message) to change a local parameter. However, **Set** requests are limited to a small set of client parameters that have read-write access defined. Most client parameters allow only read-only access.

**Understanding the Management Information Base (MIB)**

When an SNMP manager requests information from an SNMP agent, the SNMP agent retrieves the current value of the requested information from the Management Information Base (MIB). The MIB defines the managed objects that an SNMP manager monitors (or sometimes configures) on an SNMP agent.

Each system in a network (workstation, server, router, bridge, and so forth) maintains a MIB that reflects the status of the managed resources on that system, such as the version of the software running on the device, the IP address assigned to a port or interface, the amount of free hard drive space, or the number of open files. The MIB does not contain static data, but is instead an object-oriented, dynamic database that provides a logical collection of managed object definitions. The MIB defines the data type of each managed object and describes the object.

**The MIB Tree**

The SNMP-related branches of the MIB tree are located in the internet branch, which contains two main types of branches:

- *Public branches (mgmt=2)*, which are defined by the Internet Engineering Task Force (IETF) RFCs, are the same for all SNMP-managed devices.
- *Private branches (private=4)*, which are assigned by the Internet Assigned Numbers Authority (IANA), are defined by the companies and organizations to which these branches are assigned.

The following figure shows the structure of the SNMP MIB tree. There are no limits on the width and depth of the MIB tree.

**The Structure of the SNMP MIB Tree**




Immediately beneath the root of the MIB tree, International Organization for Standardization (iso), is the Organization (org) branch, followed by Department of Defense (dod), and then Internet (internet). Management (mgmt), the main public branch, defines network management parameters common to devices from all vendors. Underneath the Management branch is MIB-II (mib-2), and beneath this are branches for common management functions such as system management, printers, host resources, and interfaces.

The private branch of the MIB tree contains branches for large organizations, organized under the enterprises branch. Each organization has a root branch node under this object. Each organization creates its own subset of MIB branches and objects, which must comply with a common definition of SNMP information known as Structure of Management Information (SMI). SMI defines the allowed data types for MIB objects.

**MIB Objects**

At the programmatic level, the definition of each MIB object that an SNMP agent manages includes the following elements:

- The object name and object identifier (also known as an OID).
- A text description of the object.
- The object's data-type definition (such as counter, string, gauge, or address).
- The index for objects that are assigned complex data types. The index specifies the key field for the table — that is, the field that can be used to identify a row.

  The only complex SNMP data type that is allowed is a table, and tables cannot be nested. Examples include the list of a system's network interfaces, a routing table, or the Address Resolution Protocol (ARP) table.
- The level of access to the object (such as read or read/write) that is allowed.
- Size restrictions.
- Range information.

SNMP references each MIB variable by using its unique object identifier, which identifies the location of a given managed object within the MIB namespace. The object identifier reflects the object's position within the hierarchy of the MIB tree, containing a sequence of subidentifiers that begin at the root of the MIB tree and end at the object (leaf node). Subidentifiers are separated with a period.

To reference a MIB object, you can use either numeric or text subidentifiers. For example, the following text-based object identifier is interchangeable with its numeric counterpart, shown beneath it. The value of this object identifier, in either format, identifies the current operational state of a network adapter.

.iso.org.dod.internet.mgmt.mib-2.interfaces.ifTable.ifEntry.ifOperStatus

.1.3.6.1.2.1.2.2.1.8

Corresponding to each object identifier is a value that represents the current state of the object. SNMP, which accesses only the

leaf nodes in the MIB tree, references a MIB variable by the dotted numeric string that represents its object identifier in order to retrieve the current value of the variable.

Microsoft has the authority to assign names to objects within the 311 (microsoft) namespaces. For example, Microsoft (itself represented in MIB notation as .1.3.6.1.4.1.311) defines various MIB branches under its root node, such as:

.1.3.6.1.4.1.311.1.3 — Dynamic Host Configuration Protocol (DHCP)

.1.3.6.1.4.1.311.1.2 — Windows Internet Name Service (WINS)

.1.3.6.1.4.1.311.1.7.2 — File Transfer Protocol (FTP)

.1.3.5.1.4.1.311.1.7.3 — Hypertext Transfer Protocol (HTTP)

**Interacting with MIB Objects**
To understand how SNMP interacts with MIB objects, you must first understand that the SNMP agent consists of a master agent and one or more subagents, also known as extension agents, which are the dynamic link libraries (DLLs) that SNMP uses to process a request received from an SNMP manager.

When the SNMP service starts, it loads each SNMP extension agent that is listed in the registry. When the master agent receives a request from the SNMP manager, it passes the request to an extension agent, which accesses the appropriate MIB and returns the current value of the relevant MIB object to the SNMP master agent. The SNMP agent then returns the information to the SNMP manager. An SNMP agent can also send unsolicited MIB information to the SNMP manager in response to an event that has been defined as an SNMP trap (unexpected event). In some cases, the SNMP agent changes (or sets) the value of the MIB object.

For example, in order to assess the network load at a particular point in your network, you can use SNMP to retrieve the value of a counter that tracks the number of packets sent over a link to the SNMP agent. Or you might use SNMP to retrieve the value that represents the amount of disk space available for storage. Although SNMP is used less commonly to set the value of a MIB object, to ensure that a link is not used, you might use SNMP to set the link state value of a MIB object to a disabled state.

**Compiling a New or Updated MIB File by Using Mibcc.exe**
As explained earlier, the SNMP-related branches of the MIB tree are located in the internet branch of the tree. The internet branch contains public branches that are defined by the IETF and private branches that are defined by large organizations. When an organization creates its own subset of MIB branches and objects, or updates an existing MIB file, the new or updated MIB file must be created in compliance with SMI-prescribed data types.

If your organization adds or updates a new MIB file, use the Mibcc.exe tool to compile the MIB file so that the SNMP Management API (Mgmtapi.dll) can use the MIB objects in the new or updated MIB file. After you compile the MIB file, you can reference objects by their text object identifiers instead of their numeric object identifiers. The ASN.1 language is used to define the formats of the protocol data units (PDUs) that are exchanged by SNMP entities and to define the objects that are managed through SNMP. Mibcc.exe converts the ASN.1 MIB description into the binary Mib.bin file, which the Management API then uses to map text-based object names to numeric object identifiers.

You can find Mibcc.exe in the C:\Program Files\Resource Kit folder when you install the Windows Server 2003 Resource Kit companion CD. The Mib.bin file is located in *systemroot* on Windows Server 2003.

**SNMP Messages**
SNMP sends operation requests and responses as SNMP messages. An SNMP message consists of an SNMP protocol data unit (PDU) plus additional message header elements defined by the relevant RFC. An SNMP agent sends information in two situations:

- When it responds to a request from an SNMP manager
- When a trap event occurs

**SNMP Message Types**
The following table describes each type of message that the SNMP manager and agent programs use to communicate with each other.

### SNMP Message Types

| SNMP Message | From / To | Message Description |
|---|---|---|

| | | |
|---|---|---|
| Get | Manager / agent | Accesses and retrieves the current value of one or more MIB objects on an SNMP agent . |
| GetResponse | Agent / manager | Replies to a **Get**, **GetNext**, or **Set** operation. |
| GetNext | Manager / agent | Browses the entire tree of MIB objects, reading the values of variables in the MIB sequentially. Typically, you use **GetNext** to obtain information from selected columns from one or more rows of a table. **GetNext** is especially useful for browsing dynamic tables, such as an internal IP route table or an ARP table, reading through the table one row at a time. |
| GetBulk<br><br>(SNMPv2c only) | Manager / agent | Retrieves data in units as large as possible within the given constraints on the message size. **GetBulk**, which accesses multiple values at one time without using a **GetNext** message, minimizes the number of protocol exchanges required to retrieve a large amount of information.<br><br>To avoid fragmentation, restrict the maximum message size to a size smaller than the path maximum transmission unit (MTU), the largest frame size allowed for a single frame on your network. Typically, when it is not known how many rows are in a table, **GetBulk** is used (rather than **GetNext**) to browse all rows in the table. |
| Set | Manager / agent | Changes the current value of a MIB object. In order to update a MIB value on the SNMP agent, the SNMP manager must have write access to the object. **Set** is used infrequently, because most MIB objects are read-only by default, so that unauthorized changes cannot be made. |
| Trap | Agent / manager | Notifies the specified SNMP manager (the trap destination) when an unexpected event occurs locally on the managed host. You can use traps for limited security checking (such as notifying the trap destination if the agent receives an information request from an SNMP manager that it does not recognize) or for troubleshooting (such as notifying the trap destination if the WINS service fails). |

SNMP uses the connectionless User Datagram Protocol (UDP) service to transmit SNMP messages. SNMP uses the simple UDP transport service, which guarantees neither delivery nor correct sequencing of delivered packets, so that SNMP can continue functioning after many other network services have failed. By default, UDP port 161 is used to listen for SNMP messages and port 162 is used to listen for SNMP traps. If necessary — for example, because your organization already uses ports 161 and 162 for some other protocol or service — you can change these port settings by configuring the local Services file (this Services file is different from the Windows Services snap-in). The following figure shows messages moving between an SNMP manager and several SNMP agents.

**Messages Sent Between an SNMP Manager and its Managed Devices**




**Get Message Requires Instance Identifier**

When typing an SNMP command (such as Snmputil.exe), you can specify a MIB object by using either numeric or text subidentifiers. For example, as explained earlier, both of the following notations refer to the object identifier whose value is the current operational status of an interface:

.iso.org.dod.internet.mgmt.mib-2.interfaces.ifTable.ifEntry.ifOperStatus

.1.3.6.1.2.1.2.2.1.8

The object identifier of a MIB object identifies the type of object. The instance identifier of a MIB object, used only in SNMP **Get** operations, indicates which value to retrieve. That is, a **Get** operation must specify *ObjectIdentifier.SpecificValueToRetrieve*. Two types of instance identifiers exist:

- **Object is not in a table.** If an object can return only one value (such as the description of a system), you append the suffix **.0** to the end of the object identifier.

- **Object is in a table.** If an object can return different values (such as a router that has more than one interface), you append an index at the end of the object identifier to indicate which instance to return. For example, if a router has four interfaces and you want to use a **Get** command with Snmputil.exe to return the current status of an interface that has an interface index (ifIndex) of 3, you type **.3** at the end of the object identifier.

**SNMP Trap Types**

The SNMP standard defines six types of traps and provides a seventh trap type that an organization can use to implement enterprise-specific traps. The following table lists the seven trap types.

**SNMP Trap Types**

| Trap Type | Trap Name | Description |
|---|---|---|
| 0 | ColdStart | The SNMP agent initialized its configuration tables. |
| 1 | WarmStart | The SNMP agent re-initialized its configuration tables. |
| 2 | LinkDown | The state of a network adapter on the SNMP agent changed from up to down. |
| 3 | Linkup | The state of a network adapter on the SNMP agent changed from down to up. |
| 4 | authenticationFailure | The SNMP agent received a message from an SNMP manager, but the message contained an invalid community name. |
| 5 | egpNeighborLoss | The SNMP agent could not communicate with its Exterior Gateway Protocol (EGP) peer. |
| 6 | enterpriseSpecific | Reserved for vendor-defined error conditions and error codes. |

The Windows Server 2003 SNMP service can send enterprise-specific traps on behalf of the Event To Trap Translator.

**SNMP Communities**

SNMP version 2c introduced the use of shared community names for authentication of SNMP traffic. By default, in Windows Server 2003, no community names are configured and the SNMP service accepts SNMP requests only from the host name, localhost. To enable successful communication between an SNMP agent and an SNMP manager, you must configure at least one community name in SNMP. SNMP community names are case-sensitive.

A community name acts as a password that is shared, typically, by multiple SNMP agents and one or more SNMP managers. You configure the SNMP manager and the computers or devices that it manages as members of a single SNMP community. An SNMP agent only accepts requests from SNMP managers that are on the agent's list of acceptable community names.

When an SNMP manager sends a query or set request to the SNMP agent, the SNMP service compares the community name of the requestor with the community name of the agent and its access permissions. If the names match, the SNMP manager is successfully authenticated, and the agent replies to the query or performs the set request. If the community names do not match, the SNMP agent considers the request a failed access attempt and — if configured to do so — can send an SNMP trap message notifying the trap destination that an improper access has been attempted.

One computer or device can belong to several communities. SNMP community names are specific to SNMP and are not related to Active Directory or Workgroup names. In the past, the SNMP default community name, Public, was used widely for read-only access and the name, Private, was used for read-write access. To increase security, specify a unique community name and change it frequently, as you would a password.

By defining an SNMP community, you can group computers together for administrative convenience and provide limited security by restricting communication to only those computers that are members of the community. Two common bases for grouping computers into an SNMP community are by physical proximity or by function. The recommended practice is to create SNMP communities based on function rather than location.

In addition to specifying that the SNMP agent can communicate only with a specified set of SNMP managers, you can specify the access permission (such as **Read Only** or **Read Write**) for each community and specify that a trap be sent, alerting a trap destination to the authentication failure if either the community name or access permission is incorrect. For the community name, specify the most restrictive permission possible. For example, do not choose the **Read Write** permission if the **Read Only** permission is sufficient. Also check the permission level of the MIB access variable. It interacts with the community security level. For example, if your community access level is set to **Read Only** and the MIB access variable is set to **Read Write**, the following operations are valid: **Get**, **GetNext**, **GetResponse**, and **Trap**. However, the **Set** request is blocked by the more restrictive community access level.

You might want to specify one community name for the SNMP agent and its trap destination and another community name for the SNMP agent and any SNMP manager that sends the agent query and set requests. Using separate community names for requests and traps ensures that the agent has an available destination to which it can send an alert if an unauthorized SNMP manager tries to request information.

The security that the SNMP service provides through the use of community names used in conjunction with access permissions and authentication traps is rudimentary. Because community names are sent across the network as clear text (which is unencrypted), even secret and frequently changed community names with restricted permissions represent a potential security risk. Clear text messages are easily intercepted and decoded by network analyzers such as Microsoft Network Monitor. Community names sent in clear text can be captured and used by unauthorized personnel to access valuable information about network resources.

To help prevent the interception of SNMP messages, configure Internet Protocol Security (IPSec) policies on all SNMP managers and SNMP agents.

**Related Information**

- For more information about SNMP, see the following RFCs in the IETF RFC Database.

The following RFCs relate to SNMP version 1:

- RFC 1157, “Simple Network Management Protocol (SNMP).”
- RFC 1155, “Structure and Identification of Management Information for TCP/IP-based Internets.”
- RFC 1213, “Management Information Base for Network Management of TCP/IP-based Internets: MIB-II.”
- RFC 1573, “Evolution of the Interfaces Group of MIB-II.”

The following RFCs relate to SNMP version 2:

- RFC 1901, “Introduction to Community–Based SNMPv2.”
- RFC 1902, “Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2).”
- RFC 1903, “Textual Conventions for SNMPv2.”
- RFC 1904, “Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2).”
- RFC 1905, “Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2).”
- RFC 1906, “Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2).”
- RFC 1907, “Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2).”
- RFC 1908, “Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework.”
- RFC 2089, “V2ToV1 Mapping SNMPv2 onto SNMPv1 Within a Bi-Lingual SNMP Agent.”

# Computer Browser Service Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computer Browser Service Technical Reference**
Users often rely on the list of workgroups, domains, and computers visible from the My Network Places folder to determine which resources are accessible from their local computers. Computer Browser service, which operates on all Microsoft Windows–based computers, collects and distributes this list. Computer Browser service also provides interoperability with domains and computers that do not use the Active Directory directory service.

The Computer Browser Service Technical Reference documentation includes information about Computer Browser service mechanisms, time intervals, announcements and elections, and interoperability with various network configurations.

**In this subject**


- What Is Computer Browser Service?
- How Computer Browser Service Works

# What Is Computer Browser Service?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Computer Browser Service?**
Computer Browser service is the mechanism that collects and distributes the list of workgroups and domains and the servers within them. The list displays in the Microsoft Windows Network window and related windows in My Network Places.

Microsoft Active Directory services in Windows Server 2003 and Windows XP create a searchable list of computers and resources that is separate from the list of domains, workgroups, and servers in My Network Places. Computer Browser service provides backward compatibility with computers running earlier versions of Windows that must use Network Basic Input/Output System (NetBIOS) over TCP/IP (NetBT) and are not Active Directory–capable.

**In this section**

- Computer Browser Service Concepts

In Windows Server 2003, Computer Browser service is included with Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and the 64-bit editions of the Windows Server 2003 family. It is also included in other Windows operating systems as follows:

- Microsoft Windows XP
- Microsoft Windows 2000
- Microsoft Windows Millennium Edition
- Microsoft Windows 98

Computer Browser service (and NetBT, which it uses), is enabled by default when any of the above operating systems are installed.

Users often need to know which resources are accessible from their local computers. Computer Browser service enables users to easily browse for available domains, workgroups, servers, and resources in My Network Places. Computer Browser service also provides interoperability with domains and computers that do not use the Active Directory directory service.

In pure Windows Server 2003, Windows XP, and Windows 2000 environments that are configured to use Active Directory, Active Directory itself can display information about available network resources. In Windows 98 and Windows Millennium Edition, however, Computer Browser service must be used for this purpose. Computer Browser service is also still required in Windows Server 2003, Windows XP, and Windows 2000 for using applications such as My Network Places and Windows Explorer, or the **net view** command.

**Computer Browser Service Concepts**
Networked computers can play a variety of roles as different types of browse servers. The different types of browse servers and browse clients together are called the browser system.

In versions of Windows that support Computer Browser service, the operating system assigns tasks to specific computers on the network to provide browse services. Any networked computer that can collect, maintain, and distribute a browse list can be a browse server. Computers that are designated as browse servers work together to provide a centralized browse list, which contains a list of all known domains, workgroups, and the set of file servers in the domain to which the computer belongs.

Browse clients on the network access the browse list when users want to view the list of workgroups and domains or the list of servers in a workgroup or domain. For example, when a user at a computer running Windows 98 opens My Network Places, Computer Browser service computer obtains a copy of the browse list from a browse server on the same subnet in order to display the requested list of domains, workgroups, and computers. By working together, browse servers reduce the CPU cycles and network traffic needed to build and maintain the browse list.

When an individual server starts, it announces its presence by sending a broadcast datagram called a host announcement on the subnet. The announcement is received by a master browse server for the workgroup or domain. When the master browse

server receives a host announcement from a computer, it adds that computer to the browse list.

When a domain spans more than one subnet, the master browse server:

- Maintains the browse list for the portion of the domain on its subnet.
- Provides lists of backup browse servers on local subnets of a TCP/IP-based network.

If a TCP/IP-based subnet encompasses more than one domain, each domain has its own master browse server and backup browse servers. On networks using the NWLink IPX/SPX–Compatible protocol, routers forward the NetBIOS over IPX broadcast host announcement packets, a mechanism which allows the network to use only one master browse server.

Whenever a computer starts, if the value of the **MaintainServerList** entry in its registry is set to **Auto**, the master browse server must instruct the newly started computer whether to become a backup browse server.

As the master browse server for the network, a computer running Windows Server 2003 contacts the primary domain controller (PDC) every 12 minutes to provide it with the local network's browse list and to obtain the domain-wide browse list.

For a computer running Windows 98 to be a master browse server, the computer and the PDC for the domain must both be WINS clients. A Windows 98 master browse servermust also:

- Use TCP/IP.
- Use WINS for name resolution.
- Be in a workgroup that has the same name as the domain.

**Domain Master Browse Server**
The domain master browse server:

- Collects announcements for the entire domain.
- Collects browse lists from master browse servers on other subnets.
- Provides a browse list to other master browse servers.

A PDC typically functions as the domain master browse server on its subnet. When Computer Browser service runs on a PDC, the PDC is always the domain master browse server. To ensure that the PDC becomes the domain master browse server, you must give the PDC priority in browse server elections. For information about browse server elections, see "Browser Elections" in How Computer Browser Service Works.

In a domain that uses TCP/IP and that spans more than one subnet, each subnet functions as an independent browsing entity, with each subnet having its own master browse server and backup browse servers. On the other hand, NWLink and NetBEUI Frame (NBF) transports do not use the domain master browse server role; these transports have only a single master browse server for the entire network. When a domain spans multiple subnets, the master browse server of each subnet announces itself as the master browse server to the domain master browse server, using a directed datagram called a MasterBrowserAnnouncement. The domain master browse server then sends a remote **NetServerEnum** API call to each master browse server, to collect the list of servers from each subnet. The domain master browse server merges the server list from each subnet master browse server with its own server list, forming the server-list portion of the browse list for the domain. This process is repeated every 12 minutes to ensure that the domain master browse server has an updated and complete browse list of all the servers in the domain.

**Note**

- The domain master browse server must be able to resolve the server name of each master browse server on a TCP/IP network (using WINS, for example). Each master browse server must be able to resolve the name of the domain, as well as the computer name of the primary domain controller.

For additional information about Computer Browser service, see Help and Support Center in Windows Server 2003.

# How Computer Browser Service Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Computer Browser Service Works**

**In this section**


- Computer Browser Service Architecture
- Registration and Propagation
- Computer Browser Service Announcements
- Browser Elections
- Number of Browsers in a Domain
- Multiple Subnets
- Computer Browser Service Across an IP Router
- Name Resolution


Computer Browser service is a networking technology used by Windows-based computers to collect, distribute, and obtain browse lists, which are collections of information about workgroups, domains, and the computers within them. The service is essential for computers running Windows 98 or Windows Millennium Edition to browse for resources and it is required in Windows Server 2003, Windows XP, and Windows 2000 for using applications such as My Network Places, Windows Explorer, and the **net view** command.

Computer Browser service typically uses connectionless server broadcasts to communicate between nodes. Registration, announcements, and browser elections are performed by the service to provide the network with a list of available resources.

Computer Browser service operates ideally in a Windows network where several computers play various browser roles. Together, these browser roles form the browser system.

**Computer Browser Service Architecture**

The architecture of Computer Browser service is complete when several computers are brought together on a network to perform browse service roles. Individually, these roles are just parts of the browse service, but when combined on the network, they form a browser system that enables efficient sharing of network resource information.

**Browser System Roles**

There are five browse service roles that computers can play in the browser system:

- Non-browser
- Potential browser
- Backup browse server
- Master browse server
- Domain master browse server

The figure below shows an example of browser system roles played by computers on a single subnet.

**Browser System Containing Browsers and Non-Browsers on a Subnet**




The table below lists each of the browse service roles and describes its function in the architecture of the browser system.

**Browser Roles and Functions**

| Role | Function |
|---|---|
| Non-browser | A non-browser is a computer that has been configured not to maintain a browse list. |
| Potential browser | A potential browser is a computer that is configured to be capable of maintaining a browse list. Such a computer can act as a backup browse server if instructed to do so by the master browse server. Or, it can be elected as a master browse server. |
| Backup browse server | The backup browse server receives a copy of the browse list from the master browse server and distributes the list, upon request, to computers in the domain or workgroup. All Windows Server 2003 domain controllers are configured as either master browse servers or backup browse servers. |
| Master browse server | The master browse servercollects the information necessary to create and maintain the browse list. The browse list includes the names of all servers in the same domain or workgroup as the master browse server and a list of all domains on the network. |
| Domain master browse server | The domain master browse server collects browse lists from master browse servers on other subnets and also serves as the master browse server for its subnet. This role is typically performed by the primary domain controller (PDC) for a domain. |

**Registration and Propagation**

When a server starts, it immediately broadcasts a Host Announcement message to announce its presence to the master browse server. Host Announcement messages are transmitted again after four, eight, and 12 minutes, and then every 12 minutes thereafter.

Allowing for the loss of a few messages, the server acting as the master browse server adds the new server to its browse list within 12 minutes after the new server starts. The process of adding a new server entry to the master browse server's browse list is called registration.

After the initial host announcements are sent, new entries are transmitted through unicast sessions:

- To the browse list from the master browse server and to the domain master browse server.
- From the domain master browse server to master browse servers when master browse servers and domain master browse servers are located on different subnets.

Unicast sessions provide reliable data delivery of browse list updates.

Within the second 12 minutes after the server starts, the master browse server connects to the domain master browse server to obtain the domain-wide list. Concurrently, the domain master browse server connects to the master browse server and obtains a browse list containing the new server. Master browse servers on remote subnets connect to the domain master browse server at 12-minute intervals, so they are quickly informed of the new server's existence. Immediately after these

remote master browse servers are informed, all of the backup browse servers will connect to their respective master browse servers, and so will also be informed of the new server. The figure below shows the propagation sequence for a host announcement.

**Host Announcement Propagation**




Removing computers from the browse list when they are no longer connected to the network can take longer than adding new ones. To allow for messages that might be lost during transmission, the master browse server does not remove a server from its list until three announcement periods (36 minutes) have passed.

If the server is shut down gracefully, it sends a single Host Announcement message to the master browse server, indicating that it is no longer acting as a server. Upon receipt of this message, the master browse server immediately removes the server from its local list. On a healthy network, the removal of a server's name takes 36 minutes or less to propagate to all the master browse servers in the domain.

If the server is not shut down gracefully, or if network connectivity is lost, the server remains in the master browse server's list for up to 36 minutes. After that time, the domain master browse server is notified by the master browse server to remove the server name from its list. Within another 12 minutes, a master browse server on a remote subnet obtains the domain-wide list from the domain master browse server, and within an additional 12 minutes, each backup browse server on the remote subnet is notified to remove the server name from its list. Thus, removing a computer from the browse list can take as long as 72 minutes.

Unlike WINS or DNS server roles, browse server roles are not statically defined. Instead, they are dynamically defined with periodic browser elections. (As a result, it can be difficult to determine the flow of communication that is used by the servers in a specific instance to provide the browse list to a particular browse client.)

If a master browse server is gracefully shut down, it forces an election and a new master browse server is established promptly. This conversion occurs most quickly when a backup browse server, which already has a fully populated list, wins the election.

If a server acting as the master browse server for a subnet is not gracefully shut down, or if the Request Election message is lost, there might be a delay of several minutes until browsing is active again on that subnet. If a client cannot find a master browse server by sending the Get Backup List Request message, the client forces an election. If no client requests a browse list, it can take up to 12 minutes until a backup browse server discovers that there is no master browse server. Whenever a backup browse server discovers that there is no master browse server, it forces an election, and within another 12 minutes, browsing is active again.

**Browser Requests**

The purpose of Computer Browser service is to make the browse list — a list of workgroups, domains, and the computers within them — available to browse clients. To obtain the list of network resources, the browse client must first select an appropriate computer from which to request a copy of the list. The client selects a computer by sending a broadcast Get Backup List Request message. The master browse server on the subnet responds with a unicast message containing a list of up to three backup browse servers. The browse client then chooses one of the backup browse servers and sends it a message requesting the browse list. The selected browse server sends the contents of its browse list. The figure below shows the flow of the messages.

**Browser Request Flow**




If the master browse server for the workgroup or domain being queried cannot be found after three attempts, the client forces the election of a new master browse server in the domain.

**Computer Browser Service Announcements**

Every computer running Windows that has the file sharing component installed and enabled on a network interface broadcasts a periodic (every 12 minutes) Host Announcement message on the local subnet. On the master browse server for the subnet, these messages are used to maintain the list of available servers in the browse list

### Browser Announcements




Master browse servers receive announcements from, and return lists of backup browse servers to, computers running Windows 98, Windows Millennium Edition, Windows 2000, Windows XP, or Windows Server 2003.

When a computer starts, if the value of the **MaintainServerList** entry in its registry is **Auto**, the master browse server is responsible for telling the computer whether or not to become a backup browse server.

When a computer becomes the master browse server by winning an election, its browse list is empty, so it broadcasts an Announcement Request message. All computers that receive an Announcement Request message must reply with an announcement after a random delay of up to 30 seconds. This 30-second range for responses prevents the master browse server from becoming overloaded and losing replies, and it protects the network from being flooded with responses.

A master browse server cannot be directly forced to rebuild the browse list for a workgroup or domain. However, shutting down and restarting a computer that is configured as the preferred master browse server, or stopping and restarting Computer Browser service on that computer, forces the building of a new browse list. This happens because when a preferred master browse server starts, it forces an election, which it wins. The preferred master browse server has no browse list at startup, so it then forces all members of the domain or workgroup to announce themselves.

If a master browse server receives an announcement from another computer that claims to be the master browse server, the master browse server receiving the announcement demotes itself from its role as a master browse server and forces an election. This demotion process ensures that there is never more than one master browse server in each workgroup or domain.

**Non-Browser Announcements**

After a non-browser computer starts, it periodically broadcasts a Host Announcement message on the subnet to announce its availability. Announcements are sent 1 minute, 2 minutes, 4 minutes, 8 minutes, and 12 minutes after the non-browser starts, and every 12 minutes thereafter. If the master browse server has not received an announcement from the non-browser for three consecutive announcement periods (36 minutes), the master browse server removes the non-browser from the browse list.

**Potential Browser Announcements**

Most computers are potential browsers — that is, they are capable of becoming either backup browse servers or master browse servers. These computers announce themselves in the same manner as non-browsers.

**Backup Browse Server Announcements**

Backup browse servers announce themselves in the same manner as non-browsers. However, backup browse servers do participate in browser elections. Backup browse servers connect to the master browse server every 12 minutes to obtain

updates to the browse list. The backup browse server stores these lists and returns the browse list to any client that sends it a browse request. If the backup browse server cannot find the master browse server, it forces an election.

**Browser Announcement Time Configuration**

To change how often a computer running Computer Browser service announces itself, add a registry entry called **Announce** with a value that has the REG_DWORD data type. Set the value of the **Announce** entry to the number of seconds that the browser should wait between announcements. Add **Announce** to the following registry subkey:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\parameters**

For example, the default value of the **Announce** entry is 720, which defines an announcement interval of 12 minutes.

The value of the **Announce** entry must be changed on all computers in the workgroup or domain before all computers can use the new value. Decreasing this value has the effect of decreasing the length of time that an unavailable computer appears on the browse list, but it also increases announcement traffic. On the other hand, increasing the value of **Announce** reduces the amount of announcement traffic, but it increases the length of time that an unavailable computer appears on the browse list.

### Note

- Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editors bypass the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Windows Server 2003. To configure or customize Windows Server 2003, use the programs in Microsoft Management Console (MMC) or Control Panel whenever possible.

**Browser Elections**

Browser elections occur to select a new master browse server under the following circumstances:

- When a computer cannot locate a master browse server.
- When a preferred master browse server comes online.
- When a Windows-based domain controller starts.
- When a back-up browse server cannot contact the master browse server to obtain updates to the browse list.

A computer initiates an election by broadcasting a Request Election message. When an election occurs, Computer Browser service on the computer forcing the election logs an event in the system log, indicating that it forced the election.

All browsers can receive Request Election messages. When a browser receives a Request Election message, it examines the election criteria of that message. If the browser has a higher ranking value of the same election criteria than the sender of the Request Election message, the browser sends its own Request Election message and enters an election in progress state. If the browser does not have higher ranking value of the same election criteria than the sender of the Request Election message, the browser attempts to determine which computer is the new master browse server. The following figure shows computers performing a browser election.

### Browser Election




RE =Request Election message

The election criteria for a browser are based on the current role of the browser in the domain and its current state, as shown in

the following table.

**Browser Election Criteria**

| Election Criteria Value | Windows Election Field |
|---|---|
| Windows 98 | 0x01000000 |
| Windows XP Professional, Windows 2000 | 0x10000000 |
| Windows Server 2003, Windows 2000 Server | 0x20000000 |
| Election version | 0x00FFFF00 |
| Version-specific criteria | 0x000000FF |
| PDC | 0x00000080 |
| WINS system | 0x00000020 |
| Preferred master | 0x00000008 |
| Running master | 0x00000004 |
| MaintainServerList = Yes | 0x00000002 |
| Running backup browse server | 0x00000001 |

The browser uses all of the appropriate election criteria to determine the election ranking of the sending computer.

The following rules determine whether a browser has won an election:

- If the election version of the browser is greater than the election version of the sender, the browser wins. If not, the browser uses the next election criterion. (The election version is a constant value that identifies the version of the browser-election protocol; it is not related to the operating-system version.)

- If the value of the election criteria for the browser is greater than that for the sender, the browser wins. If not, the browser uses the next election criterion.

- If the browser has been running longer than the sender, the browser wins. If not, the browser uses the next election criterion.

- If none of the rules above have determined the election, then the server with the lexically (alphabetically, including numbers and symbols) lowest name wins. For example, a server named “A” will win an election to become master browse server over a server named “B.”

**Running Election State**

When a browser receives a Request Election message indicating that it wins an election, the browser enters the running election state. While in this state, the browser sends out a Request Election message after a delay. The delay is based on the current role of the browser in the domain:

- Master browse servers and the primary domain controllers delay for 100 microseconds (ms).

- Backup browse servers and backup domain controllers randomly delay for 200 ms and 600 ms.

- All other browsers randomly delay between 800 ms and 3000 ms.

The browser sends up to four Request Election messages. If no other browser responds with a Request Election message that wins the election, the browser that sent the Request Election messages is promoted to master browse server. If the browser receives a Request Election message indicating that another computer wins the election, and the computer is currently the master browse server, the computer demotes itself from master browse server and becomes a backup browse server.

**Number of Browsers in a Domain**
Computer Browser service uses the following rules to set the number of browsers in a domain:

- If there is currently a primary domain controller (PDC) in the domain, it is designated as the domain master browse server.

- Every backup domain controller (BDC) in the domain is a backup browse server for the domain. The only exception to this rule occurs when the PDC has failed. In that case, the BDC becomes the domain master browse server.

**Note**

- PDC failure can be a problem because the DOMAIN[1B] entry, used to locate the Domain Master browse server, is only updated if the backup domain controller is promoted.

If the registry of a computer has the value of the **MaintainServerList** entry set to **Yes**, the computer is a backup browse server for the domain or IP subnet.

If no backup browse servers are selected for the domain based on the preceding rules, the master browse server sets the number of backup browse servers to be used for the domain, and then selects — from among those computers that have the value of the **MaintainServerList** entry set to **Auto —** that number of specific computers to act as browsers.The table below shows the number of backup browse servers that are selected, based on the number of computers in the domain.

**Number of Master or Backup Browse Servers in a Domain**

| Number of Computers* | Number of Backup Browse Servers | Number of Master Browse Servers |
|---|---|---|
| 1 | 0 | 1 |
| 2 to 31 | 1 | 1 |
| 32 to 63 | 2 | 1 |

* For each additional 32 computers added to a domain, another backup browse server is added.

In a TCP/IP network, each subnet independently enforces the preceding rules.

**Multiple Subnets**
Users might need to query multiple workgroups and domains to retrieve a list of servers within their workgroup or domain and create lists of other workgroups and domains. The figure below shows Computer Browser service operating with multiple domains and workgroups on the same subnet.

**Computer Browser Service Across Multiple Domains**




Upon becoming a master browse server, the master browse server for each subnet broadcasts a Local Master Announcement message every minute for the first five minutes. After the first five minutes, the master browse server broadcasts a Local

Master Announcement message once every 12 minutes. If a subnet has not announced itself for three consecutive announcement periods, the subnet is removed from the list. Thus, a subnet might remain in the browse list for up to 45 minutes after failing or being shut down.

A Local Master Announcement message contains the following information:

- The name of the domain or workgroup
- The name of a master browse server (which might or might not be a PDC) for that domain
- The version number of the operating system

If the browser computer is running Windows Server 2003 or Windows 2000 Server, the Local Master Announcement message also specifies whether or not the browser computer is the domain PDC.

**Computer Browser Service Across an IP Router**

When you use domains that are split across routers, the TCP/IP subnets function as independent browsing entities, each with its own master browse server and backup browse servers. Therefore, browser elections occur within each subnet. In contrast, workgroups do not span across an IP router.

Domain master browse servers collect the domain browse list, which includes the set of computers in the domain, and the set of other domains and workgroups. The master browse server on each subnet collects the set of computers in the domain and other domains and workgroups from received announcements; then the master browse server forwards this information to the domain master browse server. The domain master browse server periodically downloads the domain browse list to each master browse server, which in turn updates its backup browse servers.

The figure below shows Computer Browser service functioning across an IP router.

**Computer Browser Service Across an IP Router**




When a domain spans multiple subnets, the master browse servers for each subnet use a unicast Master Announcement message to announce themselves to the domain master browse server. This message notifies the domain master browse server that the sending computer is a master browse server in the same domain. When the domain master browse server receives a Master Browse Server Announcement message, it returns to the "announcing" master browse server a request for a list of the server's in that master browse server's subnet. When that list is received, the domain master browse server merges it with its own server list.

This process, repeated every 12 minutes, guarantees that the domain master browse server has a complete browse list of all the servers in the domain. Thus, when a client sends a browse request to a backup browse server, the backup browse server can return a list of all the servers in the domain, regardless of the subnet on which those servers are located.

**Note**

- It is possible on some routers to enable the forwarding of NetBIOS broadcasts. This causes Computer Browser service to work as if all computers were connected to the same subnet. However, this practice propagates broadcast traffic and can create problems with browser elections. Microsoft recommends that you do not enable the forwarding of NetBIOS (UDP port 137 and UDP port 138) broadcast packets across routers.

**Name Resolution**

Name resolution is critical for browsing to operate properly across IP routers. For most networks, this means that the Windows

Internet Name Service (WINS) infrastructure must be working properly:

- All computers that have the potential to become master browse servers on an IP internetwork must be able to resolve the *DomainName*<1b> entry for the domain master browse server.

- The master browse servers must also be able to resolve the *ComputerName*<00> of the domain master browse server.

The unicast Master Announcement messages sent by the master browse servers on each subnet contain the name of the master browse server computer. The domain master browse server resolves the NetBIOS name *ComputerName*<00> to obtain the IP address of the master browse server and to request the master browse server's locally collected browse list.

A browse client must be able to resolve the names of the backup browse servers sent to it by the master browse server on the subnet. When a client views the shared resources of a specific computer, it must be able to resolve the NetBIOS name *ComputerName*<20>. Therefore, all clients in the domain must be able to resolve the IP address of any server in the domain.

To resolve NetBIOS names so that browsing can work across IP routers, you can use WINS (recommended) or Lmhosts file entries.

**Windows Internet Name Service**
The Windows Internet Name Service (WINS) resolves IP addresses from NetBIOS names so that messages can be sent to the targeted computer. Implementing WINS eliminates the need to configure the Lmhosts file. Using WINS requires the following configuration:

- A WINS server is configured on a computer running the Windows Server 2003 family or Windows 2000 Server.

- Clients are WINS-enabled.

WINS clients can be computers running Windows Server 2003, Windows XP, Windows 2000, Windows Me, and Windows 98.

It is recommended that you implement WINS for name resolution and browsing support.

Windows Me and Windows 98 clients will participate in domain browsing only if they are configured to use a workgroup name that is the same as the domain name.

**Lmhosts File**
NetBIOS name resolution is typically performed through broadcasts, which will resolve names only on the local subnet. To resolve names of computers located on another subnet, the Lmhosts file (located in the directory *systemroot*\System32\drivers\etc) must be configured. The Lmhosts file must contain a NetBIOS name–to–IP address mapping for all computers that are not on the local subnet.

To implement communication between subnets and the domain master browse server, the administrator must configure the Lmhosts file with the NetBIOS names and IP addresses of all browsers. To ensure that the master browse server for each subnet can access the domain's PDC, the PDC for each domain must have an entry in the Lmhosts file for each master browse server. Also, each entry must have the tag #DOM, which designates the named computer as a domain controller.

To guarantee that the domain master browse server can request the local browse list from the master browse server of each subnet, the Lmhosts file on the domain master browse server must contain entries for all the potential master browse servers. Thus, the Lmhosts file on the master browse server of each subnet needs to contain the following information:

- IP address and NetBIOS name of the domain master browse server

- Domain name, preceded by the tags #PRE and #DOM

The following is an example of such an Lmhosts file entry:

```
131.107.7.80 MasterBrowserName #PRE #DOM:domain_name
```

# Name Resolution Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Name Resolution Technologies**
Microsoft Windows Server 2003 operating systems use name resolution to translate the numerical IP addresses that are used for TCP/IP communications to computer names, which are easier for users to remember. With name resolution, computer names are assigned to the IP addresses of the source and destination hosts, and are then used to contact these hosts, rather than using the 32-bit IP addresses.

In Windows Server 2003, there are two types of TCP/IP names to resolve:

- Host names
- NetBIOS names

**Name Resolution Components**
Windows Server 2003 provides Domain Name System (DNS) for host name resolution and Windows Internet Name Service (WINS) for NetBIOS name resolution.

**DNS Name Resolution**
DNS name resolution means successfully mapping a DNS domain or host name to an IP address. A host name is an alias that is assigned to an IP node to identify it as a TCP/IP host. The host name can be up to 255 characters long and can contain alphabetic and numeric characters, hyphens, and periods. Multiple host names can be assigned to the same host.

Windows Sockets (Winsock) programs, such as Internet Explorer and the FTP utility, can use one of two values for the destination host: the IP address or a host name. When the IP address is specified, DNS name resolution is not needed. When a host name is specified, the host name must be resolved to an IP address before IP-based communication with the desired resource can begin.

Host names can take various forms. The two most common forms are a nickname and a domain name. A nickname is an alias to an IP address that individuals can assign and use. A domain name is a structured name in a hierarchical namespace called DNS. An example of a domain name is www.microsoft.com.

Nicknames are resolved through entries in the Hosts file, which is stored in the systemroot\System32\Drivers\Etc folder. For more information, see TCP/IP Technical Reference.

Domain names are resolved by sending DNS name queries to a DNS server. The DNS server is a computer that stores domain name-to-IP address mapping records or has knowledge of other DNS servers. The DNS server resolves the queried domain name to an IP address and returns the name-to-IP address mapping records in response to a query.

Windows Server 2003 client computers must be configured with the IP address of a DNS server in order to resolve domain names. Active Directory-based computers running Microsoft Windows XP Professional or Windows Server 2003 operating systems must be configured with the IP address of a DNS server. For more information, see DNS Technical Reference.

**WINS Name Resolution**
WINS name resolution means successfully mapping a NetBIOS name to an IP address. A NetBIOS name is a 16-byte address that is used to identify a NetBIOS resource on the network. A NetBIOS name is either a unique (exclusive) or group (nonexclusive) name. When a NetBIOS process is communicating with a specific process on a specific computer, a unique name is used. When a NetBIOS process is communicating with multiple processes on multiple computers, a group name is used.

The exact mechanism by which NetBIOS names are resolved to IP addresses depends on the NetBIOS node type that is configured for the node. RFC 1001, "Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods," defines the NetBIOS node types, as listed in the following table.

**NetBIOS Node Types**

| Node Type | Description |
|---|---|

| | |
|---|---|
| B-node(broadcast) | B-node uses broadcast NetBIOS name queries for name registration and resolution. B-node has two major limitations: (1) Broadcasts disturb every node on the network, and (2) Routers typically do not forward broadcasts, so only NetBIOS names on the local network can be resolved. |
| P-node (peer-peer) | P-node uses a NetBIOS name server (NBNS), such as a WINS server, to resolve NetBIOS names. P-node does not use broadcasts; instead, it queries the name server directly. |
| M-node (mixed) | M-node is a combination of B-node and P-node. By default, an M-node functions as a B-node. If an M-node is unable to resolve a name by broadcast, it queries a NBNS using P-node. |
| H-node(hybrid) | H-node is a combination of P-node and B-node. By default, an H-node functions as a P-node. If an H-node is unable to resolve a name through the NBNS, it uses a B-node to resolve the name. |

Computers running Windows Server 2003 operating systems are B-node by default and become H-node when they are configured with a WINS server. Those computers can also use a local database file called Lmhosts to resolve remote NetBIOS names. The Lmhosts file is stored in the *systemroot*\System32\Drivers\Etc folder. For more information, see TCP/IP Technical Reference.

Typically, Windows-based computers are configured with the IP address of a WINS server so remote NetBIOS names can be resolved. Active Directory-based computers, such as Windows XP Professional, Microsoft Windows 2000 and Windows Server 2003 operating systems, must be configured with the IP address of a WINS server if they are to communicate with computers running Microsoft Windows NT, Windows 95, Windows 98, or Windows Millennium Edition that are not Active Directory-based. For more information about WINS, see WINS Technical Reference.

**Name Resolution Scenarios**

DNS is the name resolution service of Windows 2000, Windows XP, and Windows Server 2003. Windows clients use DNS for name resolution and service location, including locating Active Directory domain controllers for logon. Pre-Active Directory clients (Windows NT 3.5 and 3.51, Windows NT 4.0, Windows 95, and Windows 98), however, rely on NetBIOS, which can use WINS, broadcast, or flat Lmhosts file. In particular, WINS is used by Pre-Active Directory clients for domain controller location. Windows NT 4.0-based clients can register themselves in Windows Server 2003 WINS. Clients running Windows 2000, Windows XP, and Windows Server 2003 can register in Windows NT 4.0 WINS.

The DNS implementation in Windows 2000, Windows XP, and Windows Server 2003 is WINS-aware, and, as a result, a combination of both DNS and WINS is typically used to achieve maximum efficiency in locating various network resources in a mixed environment.

# DNS Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS Technical Reference**

**In this subject**


- What Is DNS?
- How DNS Works
- DNS Tools and Settings


Domain Name System (DNS) is the default name resolution service used in a Microsoft Windows Server 2003 network. DNS is part of the Windows Server 2003 TCP/IP protocol suite and all TCP/IP network connections are, by default, configured with the IP address of at least one DNS server in order to perform name resolution on the network. Windows Server 2003 components that require name resolution will attempt to use this DNS server before attempting to use the previous default Windows name resolution service, Windows Internet Name Service (WINS).

Windows Server 2003 includes an IETF standard-based DNS server as a service. The Windows Server 2003 DNS Server service is RFC-compliant and therefore it is fully compatible with any other RFC-compliant DNS server. A DNS Client resolver is included as a service in all client and server versions of the Windows operating system.

The Technical Reference DNS subject describes how the DNS Server and Client services form a Windows Server 2003 DNS infrastructure.

# What Is DNS?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is DNS?**

**In this section**


- Technologies That Use DNS
- Related Information


Domain Name System (DNS) is one of the industry-standard suite of protocols that comprise TCP/IP. Microsoft Windows Server 2003. DNS is implemented using two software components: the DNS server and the DNS client (or resolver). Both components are run as background service applications.

Network resources are identified by numeric IP addresses, but these IP addresses are difficult for network users to remember. The DNS database contains records that map user-friendly alphanumeric names for network resources to the IP address used by those resources for communication. In this way, DNS acts as a mnemonic device, making network resources easier to remember for network users.

The Windows Server 2003 DNS Server and Client services use the DNS protocol that is included in the TCP/IP protocol suite. DNS is part of the application layer of the TCP/IP reference model.

**DNS in TCP/IP**




For more information and to view logical diagrams illustrating how DNS fits with other Windows Server 2003 technologies, see "How DNS Works" in this collection.

By default, Windows Server 2003 DNS is used for all name resolution in a Windows Server 2003 network. In the most typical scenario, when a Windows Server 2003 network user specifies the name of a network host or an internet DNS domain name, the DNS Client service running on the Windows Server 2003 computer of the user contacts a DNS server to resolve the name to an IP address.

**Technologies That Use DNS**

**DNS and Active Directory**

Windows Server 2003 Active Directory directory service uses DNS as its domain controller location mechanism. When any of the principal Active Directory operations is performed, such as authentication, updating, or searching, Windows Server 2003 computers use DNS to locate Active Directory domain controllers and these domain controllers use DNS to locate each other. For example, when a network user with an Active Directory user account logs in to an Active Directory domain, the user's computer uses DNS to locate a domain controller for the Active Directory domain to which the user wants to log in. For more information about integrating DNS and Active Directory, see "How DNS Works" in this collection.

**DNS and WINS**

The earlier method of name resolution for a Windows network was Windows Internet Name Service (WINS). DNS is different than WINS in that DNS is a hierarchical namespace and WINS is a flat namespace. Down-level clients and applications that rely on NetBIOS names continue to use WINS for name resolution. Since Windows Server 2003 DNS is WINS-aware, a combination of both DNS and WINS can be used in a mixed environment to achieve maximum efficiency in locating various network services and resources. For more information about using DNS in a mixed environment, see "How DNS Works" in this collection.

**DNS and DHCP**

For Windows Server 2003 DNS, the DHCP service provides default support to register and update information for legacy DHCP clients in DNS zones. Legacy clients typically include other Microsoft TCP/IP client computers that were released prior to Windows 2000. The Windows Server 2003 DNS-DHCP integration enables a DHCP client that is unable to dynamically update DNS resource records directly to have this information updated in DNS forward and reverse lookup zones by the DHCP server.

**Related Information**

The following resources contain additional information that is relevant to this section:

- How DNS Works

# How DNS Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How DNS Works**
**In this section**


- DNS Architecture
- DNS Protocol
- DNS Physical Structure
- DNS Processes and Interactions
- Network Ports Used By DNS
- Related Information


Domain Name System (DNS) is the default name resolution service used in a Microsoft Windows Server 2003 network. DNS is part of the Windows Server 2003 TCP/IP protocol suite and all TCP/IP network connections are, by default, configured with the IP address of at least one DNS server in order to perform name resolution on the network. Windows Server 2003 components that require name resolution will attempt to use this DNS server before attempting to use the previous default Windows name resolution service, Windows Internet Name Service (WINS).

Typically, Windows Server 2003 DNS is deployed in support of Active Directory directory service. In this environment, DNS namespaces mirror the Active Directory forests and domains used by an organization. Network hosts and services are configured with DNS names so that they can be located in the network, and they are also configured with DNS servers that resolve the names of Active Directory domain controllers.

Windows Server 2003 DNS is also commonly deployed as a non-Active Directory, or standard, Domain Name System solution, for the purposes of hosting the Internet presence of an organization, for example.

**DNS Architecture**
DNS architecture is a hierarchical distributed database and an associated set of protocols that define:

- A mechanism for querying and updating the database.
- A mechanism for replicating the information in the database among servers.
- A schema of the database.

DNS originated in the early days of the Internet when the Internet was a small network established by the United States Department of Defense for research purposes. The host names of the computers in this network were managed through the use of a single HOSTS file located on a centrally administered server. Each site that needed to resolve host names on the network downloaded this file. As the number of hosts on the Internet grew, the traffic generated by the update process increased, as well as the size of the HOSTS file. The need for a new system, which would offer features such as scalability, decentralized administration, support for various data types, became more and more obvious.

The Domain Name System introduced in 1984 became this new system. With DNS, the host names reside in a database that can be distributed among multiple servers, decreasing the load on any one server and providing the ability to administer this naming system on a per-partition basis. DNS supports hierarchical names and allows registration of various data types in addition to host name to IP address mapping used in HOSTS files. Because the DNS database is distributed, its potential size is unlimited and performance is not degraded when more servers are added.

The original DNS was based on Request for Comment (RFC) 882 (“Domain Names: Concepts and Facilities”) and RFC 883 (Domain Names–Implementation and Specification), which were superseded by RFC 1034 (“Domain Names–Concepts and

Facilities"), and RFC 1035 ("Domain Names–Implementation and Specification"). Additional RFCs that describe DNS security, implementation, and administrative issues later augmented the original design specifications.

The implementation of DNS — Berkeley Internet Name Domain (BIND) — was originally developed for the 4.3 BSD UNIX Operating System. The Microsoft implementation of DNS became a part of the operating system in Microsoft Windows NT Server 4.0. The Windows NT 4.0 DNS server, like most DNS implementations, has its roots in RFCs 1034 and 1035.

The RFCs used in Microsoft Windows 2000 and Windows Server 2003 operating systems are 1034, 1035, 1886, 1996, 1995, 2136, 2308, and 2052.

**DNS Domain Names**

The Domain Name System is implemented as a hierarchical and distributed database containing various types of data, including host names and domain names. The names in a DNS database form a hierarchical tree structure called the domain namespace. Domain names consist of individual labels separated by dots, for example: mydomain.microsoft.com.

A Fully Qualified Domain Name (FQDN) uniquely identifies the hosts position within the DNS hierarchical tree by specifying a list of names separated by dots in the path from the referenced host to the root. The next figure shows an example of a DNS tree with a host called mydomain within the microsoft.com. domain. The FQDN for the host would be mydomain.microsoft.com.

**Understanding the DNS Domain Namespace**

The DNS domain namespace, as shown in the following figure, is based on the concept of a tree of named domains. Each level of the tree can represent either a branch or a leaf of the tree. A branch is a level where more than one name is used to identify a collection of named resources. A leaf represents a single name used once at that level to indicate a specific resource.

### DNS Domain Name Hierarchy




The previous figure shows how Microsoft is assigned authority by the Internet root servers for its own part of the DNS domain namespace tree on the Internet. DNS clients and servers use queries as the fundamental method of resolving names in the tree to specific types of resource information. This information is provided by DNS servers in query responses to DNS clients, who then extract the information and pass it to a requesting program for resolving the queried name. In the process of resolving a name, keep in mind that DNS servers often function as DNS clients, querying other servers in order to fully resolve a queried name.

**How the DNS Domain Namespace Is Organized**

Any DNS domain name used in the tree is technically a domain. Most DNS discussions, however, identify names in one of five ways, based on the level and the way a name is commonly used. For example, the DNS domain name registered to Microsoft (microsoft.com.) is known as a second-level domain. This is because the name has two parts (known as labels) that indicate it is located two levels below the root or top of the tree. Most DNS domain names have two or more labels, each of which indicates a new level in the tree. Periods are used in names to separate labels.

The five categories used to describe DNS domain names by their function in the namespace are described in the following table, along with an example of each name type.

### Types of DNS Domain Names

| Name Type | Description | Example |
|---|---|---|
| Root domain | This is the top of the tree, representing an unnamed level; it is sometimes shown as two empty quotation marks (""), indicating a null value. When used in a DNS domain name, it is stated by a trailing period (.) to designate that the name is located at the root or highest level of the domain hierarchy. In this instance, the DNS domain name is considered to be complete and points to an exact location in the tree of names. Names stated this way are called fully qualified domain names (FQDNs). | A single period (.) or a period used at the end of a name, such as "example.microsoft.com." |
| Top level domain | A name used to indicate a country/region or the type of organization using a name. | ""."com", which indicates a name registered to a business for commercial use on the Internet. |
| Second level domain | Variable-length names registered to an individual or organization for use on the Internet. These names are always based upon an appropriate top-level domain, depending on the type of organization or geographic location where a name is used. | ""microsoft.com. ", which is the second-level domain name registered to Microsoft by the Internet DNS domain name registrar. |
| Subdomain | Additional names that an organization can create that are derived from the registered second-level domain name. These include names added to grow the DNS tree of names in an organization and divide it into departments or geographic locations. | ""example.microsoft.com. ", which is a fictitious subdomain assigned by Microsoft for use in documentation example names. |
| Host or resource name | Names that represent a leaf in the DNS tree of names and identify a specific resource. Typically, the leftmost label of a DNS domain name identifies a specific computer on the network. For example, if a name at this level is used in a host (A) RR, it is used to look up the IP address of computer based on its host name. | ""host-a.example.microsoft.com.", where the first label ("host-a") is the DNS host name for a specific computer on the network. |

**DNS and Internet Domains**

The Internet Domain Name System is managed by a Name Registration Authority on the Internet, responsible for maintaining top-level domains that are assigned by organization and by country/region. These domain names follow the International Standard 3166. Some of the many existing abbreviations, reserved for use by organizations, as well as two-letter and three-letter abbreviations used for countries/regions are shown in the following table:

**Some DNS Top-level Domain Names (TLDs)**

| DNS Domain Name | Type of Organization |
|---|---|
| com | Commercial organizations |
| edu | Educational institutions |
| org | Non-profit organizations |
| net | Networks (the backbone of the Internet) |
| gov | Non-military government organizations |
| mil | Military government organizations |
| arpa | Reverse DNS |
| "xx" | Two-letter country code (i.e. us, au, ca, fr) |

**Resource Records**

A DNS database consists of resource records (RRs). Each RR identifies a particular resource within the database. There are various types of RRs in DNS. This section provides information about the common structure of resource records. RRs are

discussed in greater detail in “Resource Records in DNS” later in this document.

The following table provides detailed information about structure of common RRs.

**Common DNS Resource Records**

| Description | Class | Time To Live (TTL) | Type | Data |
|---|---|---|---|---|
| Start of Authority | Internet (IN) | Default TTL is 60 minutes | SOA | Owner Name<br>Primary Name Server DNS Name, Serial Number<br>Refresh Interval<br>Retry Interval<br>Expire Time<br>Minimum TTL |
| Host | Internet (IN) | Record-specific TTL if present, or else zone (SOA) TTL | A | Owner Name (Host DNS Name)<br>Host IP Address |
| Name Server | Internet (IN) | Record-specific TTL if present, or else zone (SOA) TTL | NS | Owner Name<br>Name Server DNS Name |
| Mail Exchanger | Internet (IN) | Record-specific TTL if present, or else zone (SOA) TTL | MX | Owner Name<br>Mail Exchange Server DNS Name, Preference Number |
| Canonical Name<br>(an alias) | Internet (IN) | Record-specific TTL if present, or else zone (SOA) TTL | CNAME | Owner Name (Alias Name)<br>Host DNS Name |

**Distributing the DNS Database: Zone Files and Delegation**

A DNS database can be partitioned into multiple zones. A zone is a portion of the DNS database that contains the resource records with the owner names that belong to the contiguous portion of the DNS namespace. Zone files are maintained on DNS servers. A single DNS server can be configured to host zero, one or multiple zones.

Each zone is anchored at a specific domain name referred to as the zone’s root domain. A zone contains information about all names that end with the zone’s root domain name. A DNS server is considered authoritative for a name if it loads the zone containing that name. The first record in any zone file is a Start of Authority (SOA) RR. The SOA RR identifies a primary DNS name server for the zone as the best source of information for the data within that zone and as an entity processing the updates for the zone.

A name within a zone can also be delegated to a different zone that is hosted on a different DNS server. Delegation is a process of assigning responsibility for a portion of a DNS namespace to a DNS server owned by a separate entity. This separate entity could be another organization, department or workgroup within your company. Such delegation is represented by the NS resource record that specifies the delegated zone and the DNS name of the server authoritative for that zone. Delegating across multiple zones was part of the original design goal of DNS.

The primary reasons to delegate a DNS namespace include:

- A need to delegate management of a DNS domain to a number of organizations or departments within an organization.
- A need to distribute the load of maintaining one large DNS database among multiple DNS servers to improve the name resolution performance as well as create a DNS fault tolerant environment.
- A need to allow for a host’s organizational affiliation by including them in appropriate domains.

The name server (NS) RRs facilitate delegation by identifying DNS servers for each zone and the NS RRs appear in all zones.

Whenever a DNS server needs to cross a delegation in order to resolve a name, it will refer to the NS RRs for DNS servers in the target zone.

In the figure below, the management of the microsoft.com. domain is delegated across two zones, microsoft.com. and mydomain.microsoft.com.

**DNS Delegation**




**Note**

- If multiple NS records exist for a delegated zone identifying multiple DNS servers available for querying, the Windows Server 2003 DNS Server service will be able to select the closest DNS server based on the round trip intervals measured over time for every DNS server.

**Replicating the DNS Database**

There could be multiple zones representing the same portion of the namespace. Among these zones there are three types:

- Primary
- Secondary
- Stub

Primary is a zone to which all updates for the records that belong to that zone are made. A secondary zone is a read-only copy of the primary zone. A stub zone is a read-only copy of the primary zone that contains only the resource records that identify the DNS servers that are authoritative for a DNS domain name. Any changes made to the primary zone file are replicated to the secondary zone file. DNS servers hosting a primary, secondary or stub zone are said to be authoritative for the DNS names in the zone.

As mentioned above, a DNS server can host multiple zones. A DNS server can therefore host both a primary zone (which has the writeable copy of a zone file) and a separate secondary zone (which obtains a read-only copy of a zone file). A DNS server hosting a primary zone is said to be the primary DNS server for that zone, and a DNS server hosting a secondary zone is said to be the secondary DNS server for that zone.

**Note**

- A secondary or stub zone cannot be hosted on a DNS server that hosts a primary zone for the same domain name.

**Zone Transfer**

The process of replicating a zone file to multiple DNS servers is called zone transfer.Zone transfer is achieved by copying the zone file from one DNS server to a second DNS server. Zone transfers can be made from both primary and secondary DNS servers.

A master DNS server is the source of the zone information during a transfer. The master DNS server can be a primary or

secondary DNS server. If the master DNS server is a primary DNS server, then the zone transfer comes directly from the DNS server hosting the primary zone. If the master server is a secondary DNS server, then the zone file received from the master DNS server by means of a zone transfer is a copy of the read-only secondary zone file.

The zone transfer is initiated in one of the following ways:

- The master DNS server sends a notification (RFC 1996) to one or more secondary DNS servers of a change in the zone file.

- When the DNS Server service on the secondary DNS server starts, or the refresh interval of the zone has expired (by default it is set to 15 minutes in the SOA RR of the zone), the secondary DNS server will query the master DNS server for the changes.

**Types of Zone File Replication**
There are two types of zone file replication. The first, a full zone transfer (AXFR), replicates the entire zone file. The second, an incremental zone transfer (IXFR), replicates only records that have been modified. Zone transfer is discussed in detail later in this document.

BIND 4.9.3 and earlier DNS server software, as well as Windows NT 4.0 DNS, support full zone transfer (AXFR) only. There are two types of the AXFR: one requires single record per packet, the other allows multiple records per packet. The Windows 2000 and Windows Server 2003 DNS Server service supports both types of zone transfer, but by default uses multiple records per packet. It can be configured differently for compatibility with servers that do not allow multiple records per packet, such as BIND servers versions 4.9.4 and earlier.

**Querying the Database**
DNS queries can be sent from a DNS client (resolver) to a DNS server, or between two DNS servers.

A DNS query is merely a request for DNS resource records of a specified resource record type with a specified DNS name. For example, a DNS query can request all resource records of type A (host) with a specified DNS name.

There are two types of DNS queries that may be sent to a DNS server:

- Recursive

- Iterative

A recursivequery forces a DNS server to respond to a request with either a failure or a successful response. DNS clients (resolvers) typically make recursive queries. With a recursive query, the DNS server must contact any other DNS servers it needs to resolve the request. When it receives a successful response from the other DNS server(s), it then sends a response to the DNS client. The recursive query is the typical query type used by a resolver querying a DNS server and by a DNS server querying its forwarder, which is another DNS server configured to handle requests forwarded to it. For more information about forwarders, see “Forwarding” later in this document.

When a DNS server processes a recursive query and the query cannot be resolved from local data (local zone files or cache of previous queries), the recursive query must be escalated to a root DNS server. Each standards-based implementation of DNS includes a cache file (or root server hints) that contains entries for the root DNS servers of the Internet domains. (If the DNS server is configured with a forwarder, the forwarder is used before a root server is used.)

An iterative query is one in which the DNS server is expected to respond with the best local information it has, based on what the DNS server knows from local zone files or from caching. This response is also known as a referral if the DNS server is not authoritative for the name. If a DNS server does not have any local information that can answer the query, it simply sends a negative response. A DNS server makes this type of query as it tries to find names outside of its local domain(s) (when it is not configured with a forwarder). It may have to query a number of outside DNS servers in an attempt to resolve the name.

The following figure shows an example of both types of queries.

**DNS Query Types**




As shown in the graphic above, a number of queries were used to determine the IP address for www.whitehouse.gov. The query sequence is described below:

1. Recursive query for www.whitehouse.gov (A resource record)

2. Iterative query for www.whitehouse.gov (A resource record)

3. Referral to the .gov name server (NS resource records, for .gov); for simplicity, iterative A queries by the DNS server (on the left) to resolve the IP addresses of the Host names of the name server's returned by other DNS servers have been omitted.

4. Iterative query for www.whitehouse.gov (A resource record)

5. Referral to the whitehouse.gov name server (NS resource record, for whitehouse.gov)

6. Iterative query for www.whitehouse.gov (A resource record)

7. Answer to the interative query from whitehouse.gov server (www.whitehouse.gov's IP address)

8. Answer to the original recursive query from local DNS server to Resolver (www.whitehouse.gov's IP address)

**Time to Live for Resource Records**

The Time to Live (TTL) value in a resource record indicates a length of time used by other DNS servers to determine how long to cache information for a record before expiring and discarding it. For example, most resource records created by the DNS Server service inherit the minimum (default) TTL of one hour from the start of authority (SOA) resource record, which prevents extended caching by other DNS servers.

A DNS client resolver caches the responses it receives when it resolves DNS queries. These cached responses can then be used to answer later queries for the same information. The cached data, however, has a limited lifetime specified in the TTL parameter returned with the response data. TTL ensures that the DNS server does not keep information for so long that it becomes out of date. TTL for the cache can be set on the DNS database (for each individual resource record, by specifying the TTL field of the record and per zone through the minimum TTL field of the SOA record) as well as on the DNS client resolver side by specifying the maximum TTL the resolver allows to cache the resource records.

There are two competing factors to consider when setting the TTL. The first is the accuracy of the cached information, and the second is the utilization of the DNS servers and the amount of network traffic. If the TTL is short, then the likelihood of having old information is reduced considerably, but it increases utilization of DNS servers and network traffic, because the DNS client must query DNS servers for the expired data the next time it is requested. If the TTL is long, the cached responses could become outdated, meaning the resolver could give false answers to queries. At the same time, a long TTL decreases utilization of DNS servers and reduces network traffic because the DNS client answers queries using its cached data.

If a query is answered with an entry from cache, the TTL of the entry is also passed with the response. This way the resolvers that receive the response know how long the entry is valid. The resolvers honor the TTL from the responding server; they do not reset it based on their own TTL. Consequently, entries truly expire rather than live in perpetuity as they move from DNS server to DNS server with an updated TTL.

**Note**

- In general, never configure the TTL to zero. The different between a setting of 0 or 60 is minimal to the accuracy of the record, but when the TTL is set to 0 there is a major impact on DNS server performance because the DNS server is constantly querying for the expired data.

**Updating the DNS Database**

Since the resource records in the zone files are subjected to changes, they must be updated. The implementation of DNS in Windows 2000 and Windows Server 2003 supports both static and dynamic updates of the DNS database. The details of the dynamic update are discussed later in this document.

**DNS Architecture Diagrams**

The following diagrams illustrate how the DNS Client and Server services work and provide additional information regarding name resolution, update, and administration operations.

The first diagram illustrates the DNS Client service architecture in its name resolution and update operations. In this diagram, name resolution architecture is demonstrated using a Web browser and Microsoft Outlook and updates are represented by the DHCP client.

**DNS Client Service Architecture**




The following diagram illustrates the DNS Server service architecture with its administration tools and the Windows Management Instrumentation (WMI) interface.

**DNS Server Service Architecture**




**DNS Protocol**

The DNS protocol consists of DNS different types of DNS messages that are processed according to the information in their message fields. This section discusses the different types of DNS messages and the different fields in each message type.

In this section, the following DNS message topics are discussed:

- Message types
- DNS query message format
- DNS query message header
- DNS query question entries
- DNS resource records
- Name query message
- Name query response
- Reverse name query message
- DNS update message format
- DNS update message flags
- Dynamic update response message

**Message Types**

There are three types of DNS messages:

- Queries
- Responses
- Updates

Queries and responses are defined in the original DNS standard, and updates are defined in RFC 2136. All three types follow a common message format.

**DNS Query Message Format**
The common DNS message format has a fixed-length, 12-byte header and a variable position reserved for question, answer, authority, and additional DNS resource records. The common message format can be illustrated as follows:

**Standard DNS Query Message Format**

| **DNS Message Format** |
|---|
| DNS Header (fixed length) |
| Question Entries (variable length) |
| Answer Resource Records (variable length) |
| Authority Resource Records (variable length) |
| Additional Resource Records(variable length) |

**DNS Query Message Header**
The DNS message header contains the following fields, in the following order:

**DNS Query Message Header Fields**

| **Field Name** | **Description** |
|---|---|
| **Transaction ID** | A 16-bit field identifying a specific DNS transaction. The transaction ID is created by the message originator and is copied by the responder into its response message. Using the transaction ID, the DNS client can match responses to its requests. |
| **Flags:** | A 16-bit field containing various service flags that are communicated between the DNS client and the DNS server, including: |
| Request/response | 1-bit field set to 0 to represent a name service request or set to 1 to represent a name service response. |
| Operation code | 4-bit field represents the name service operation of the packet: 0x0 is a query. |
| Authoritative answer | 1-bit field represents that the responder is authoritative for the domain name in the query message. |
| Truncation | 1-bit field that is set to 1 if the total number of responses exceeded the User Datagram Protocol (UDP) datagram. Unless UDP datagrams larger than 512 bytes or EDNS0 are enabled, only the first 512 bytes of the UDP reply are returned. |
| Recursion desired | 1-bit field set to 1 to indicate a recursive query and 0 for iterative queries. If a DNS server receives a query message with this field set to 0 it returns a list of other DNS servers that the client can choose to contact. This list is populated from local cache data. |
| Recursion available | 1-bit field set by a DNS server to 1 to represent that the DNS server can handle recursive queries. If recursion is disabled, the DNS server sets the field appropriately. |
| Reserved | 3-bit field that is reserved and set to 0. |

| | |
|---|---|
| Return code | 4-bit field holding the return code:<br>• 0 is a successful response (query answer is in the query response).<br>• 0x3 is a name error, indicating that an authoritative DNS server responded that the domain name in the query message does not exist. For more information about return codes, see "Related Information" at the end of this document. |
| **Question Resource Record count** | A 16-bit field representing the number of entries in the question section of the DNS message. |
| **Answer Resource Record count** | A 16-bit field representing the number of entries in the answer section of the DNS message. |
| **Authority Resource Record count** | A 16-bit field representing the number of authority resource records in the DNS message. |
| **Additional Resource Record count** | A 16-bit field representing the number of additional resource records in the DNS message. |

**DNS Query Question Entries**

The DNS message's Question Entries section contains the domain name that is being queried and has the following three fields:

**DNS Query Question Entry Fields**

| **Field Name** | **Description** |
|---|---|
| **Question Name** | The domain name that is being queried. DNS domain names are expressed as a series of labels, such as microsoft.com, but in the Question Name field the domain name is encoded as a series of length-value pairs consisting of a 1-byte file that indicates the length of the value, followed by the value (the label). For example, the domain microsoft.com is expressed as *0x09microsoft0x03com0x00*, where the hexadecimal digits represent the length of each label, the ASCII characters indicate the individual labels, and the final 0 indicates the end of the name. |
| **Question Type** | Uses a 16-bit integer to represents the resource record type that should be returned, as expressed below: |
| Type value | Record(s) Returned |
| 0x01 | Host (A) record |
| 0x02 | Name server (NS) record |
| 0x05 | Alias (CNAME) record |
| 0x0C (12) | Reverse-lookup (PTR) record |
| 0x0F (15) | Mail exchange (MX) record |
| 0x21 (33) | Service (SRV) record |
| 0xFB (251) | Incremental zone transfer (IXFR) record |
| 0xFC (252) | Standard zone transfer (AXFR) record |

| | |
|---|---|
| 0xFF (255) | All records |
| **Question Class** | Represents the IN (Internet) question class and is normally set to 0x0001. |

**DNS Resource Records**

The answer, authority, and additional information sections of a DNS response message can contain resource records that answer the query message question section. Resource records are formatted as follows:

**DNS Resource Record Message Fields**

| Field Name | Description |
|---|---|
| **Resource record name** | The DNS domain name recorded as a variable-length field following the same formatting as the Question Name field. |
| **Resource record type** | The resource record type value. |
| **Resource record class** | The resource record class code, the Internet class, 0x0001. |
| **Time-to-live** | The TTL expressed in seconds as a 32-bit unsigned field. |
| **Resource data length** | 2-byte field indicating the length of the resource data. |
| **Resource data** | Variable-length data corresponding to the resource record type. |

The Resource Record Name field is encoded in the same way as the Question Name field unless the name is already present elsewhere in the DNS message, in which case a 2-byte field is used in place of a length-value encoded name and acts as a pointer to the name that is already present.

**Name Query Message**

A Name Query message format is the same as the DNS message format described above. In a typical Name Query message, the DNS message fields would be set as follows:

**DNS Name Query Message Fields**

| Field Name | Description |
|---|---|
| **Query identifier (Transaction ID)** | Set to a unique number to enable the DNS client resolver to match the response to the query. The query response transaction ID always matches the query request transaction ID. |
| **Flags** | Set to indicate a standard query with recursion enabled. |
| **Question count** | Set to 1. |
| **Question entry** | Set to the domain name queried and the resource record type to return. |

**Name Query Response**

A Name Query Response message format is the same as the DNS message format described above. In a typical Name Query message, the DNS message fields would be set as follows:

**DNS Name Query Response Fields**

| Field Name | Description |
|---|---|
| **Query identifier (Transaction ID)** | Set to a unique number to enable the DNS client resolver to match the response to the query. |
| **Flags** | Set to indicate a standard query with recursion enabled. |
| **Question count** | Set to 1. |

| | |
|---|---|
| **Question entry** | Set to the domain name queried and the resource record type to return. |

**Reverse Name Query Message**

Reverse name query messages use the common message format with the following differences:

- The DNS client resolver constructs the domain name in the in-addr.arpa domain based on the IP address that is queried.
- A Pointer (PTR) resource record is queried rather than a host (A) resource record.

**DNS Update Message Format**

The DNS update message format uses a header defining the update operation to be performed and a resource record set that contains the update. The DNS update message format has the following fields:

- **Identification**. A 16-bit identifier assigned by the DNS client requestor. This identifier is copied in the corresponding reply and can be used by the requestor to match replies to outstanding requests, or by the server to detect duplicated requests from some requestor.
- **Flags**. A 16-bit DNS update message flags field. For a description of each flag, see "DNS Update Message Flags" below.
- **Number of zone entries**. The number of resource records in the Zone entry section.
- **Number of prerequisite resource records**. The number of resource records in the Prerequisite resource records section.
- **Number of update resource records**. The number of resource records in the Update resource records section.
- **Number of additional resource records**. The number of resource records in the Additional resource records section.
- **Zone entry**. Denotes the zone of the records being updated. All records to be updated must be in the same zone, and therefore the Zone Section is allowed to contain exactly one record. It has three values: ZNAME is the zone name, the ZTYPE must be SOA, and the ZCLASS is the zone's class.
- **Prerequisite resource records**. Contains a set of resource record prerequisites which must be satisfied at the time the update message is received by the master DNS server. There are five possible sets of values that can be expressed:
    - Resource record set exists (value independent). At least one resource record with a specified name and type (in the zone and class specified by the Zone Section) must exist.
    - Resource record set exists (value dependent). A set of resource records with a specified name and type exists and has the same members with the same data as the resource record set specified in this section.
    - Resource record set does not exist. No resource records with a specified name and type (in the zone and class denoted by the Zone section) exist.
    - Name is in use. At least one resource record with a specified name (in the zone and class specified by the Zone section) exists. This prerequisite is not satisfied by empty nonterminals.
    - Name is not in use. No resource record of any type is owned by a specified name. This prerequisite is satisfied by empty nonterminals.
- **Update resource records**. Contains the resource records that are to be added or deleted from the zone. One of four operations are performed during the update:
    - Add resource records to an resource records set.
    - Delete an resource records set

- Delete all resource records sets from a name.
- Delete a resource record from an resource records set.

- **Additional resource records**. Contains resource records which are related to the update, or to new resource records being added by the update.

**DNS Update Message Flags**

The DNS update message flags field uses the following flags:

- **Request/response**. 1-bit field set to 0 to represent an update request and 1 to represent an update response.
- **Operation code**. 4-bit field set to 0x5 for DNS updates.
- **Reserved**. 7-bit reserved field set to 0.
- **Return code**. 4-bit field containing codes to represent the result of the update query. The codes are as follows:

  **DNS Update Message Flag Field Return Code Values**

| Result Code Value | Description |
|---|---|
| 0 (NOERROR) | No error; successful update. |
| 1 (FORMERR) | Format error; DNS server did not understand the update request. |
| 0x2 (SERVFAIL) | DNS server encountered an internal error, such as a forwarding timeout |
| 0x3 (NXDOMAIN) | A name that should exist does not exist. |
| 0x4 (NOTIMP) | DNS server does not support the specified Operation code. |
| 0x5 (REFUSED) | DNS server refuses to perform the update because |
| 0x6 (YXDOMAIN) | A name that should not exist does exist. |
| 0x7 (YXRRSET) | A resource record set that should not exist does exist. |
| 0x8 (NXRRSET) | A resource record set that should exist does not exist. |
| 0x9 (NOTAUTH) | DNS server is not authoritative for the zone named in the Zone section. |
| 0xA (NOTZONE) | A name used in the Prerequisite or Update sections is not within the zone specified by the Zone section. |

**Dynamic Update Response Message**

The dynamic update response message follows the same format as the DNS update message, with the exception of the DNS flags. The dynamic update response message header flags indicate whether or not the update is successful by including the successful response code or one of the error codes described in DNS update message flags above.

**DNS Physical Structure**

The logical structure of Windows Server 2003 DNS involves DNS namespace partitioning, which extends the DNS domain name hierarchy into multiple subdomains. The physical structure of DNS involves distributing the DNS database using DNS servers to host DNS zones for the subdomains of the DNS domain name hierarchy. Both the DNS Client and Server service applications manage the physical DNS data in the DNS database.

**DNS Client Service**

The Windows Server 2003 operating system, as well as Microsoft Windows XP and Windows 2000 include a DNS Client service. This service performs all necessary DNS lookups and provides a local cache for DNS queries that reduces DNS network traffic and speeds name resolution.

This service can be stopped and started using the **Services** console. Computers running Windows 2000, Windows XP and Windows Server 2003 enable the DNS Client service by default.

The Windows Server 2003 DNS Client service performs the following tasks:

- Registers its names in DNS.
- Name resolution.
- Caching responses to name resolution queries.
- Removes previously resolved names from the cache when it receives a negative response for the name.
- Negative caching.
- Keeps track of transitory (Plug and Play) network connections and the DNS server lists based on their IP configurations.
- Maintains connection-specific domain name suffixes.
- Prioritizes which DNS servers it uses according to whether they respond to a query if multiple DNS server are configured on the client.
- Prioritizes the multiple A resource records it receives from a DNS server based on their IP address.
- Initiates a network failure timeout when all DNS Client service queries time out, and does not submit any queries for 30 seconds. This feature applies to every adapter separately.

Windows XP, Windows 2000 and Windows Server 2003 DNS client configuration involves the following settings in the TCP/IP properties for each computer:

- **Domain Names**. Domain names are to form the fully qualified domain name (FQDN) for DNS clients.
- **Host names**. A DNS computer or host name for each computer. For example, in the fully qualified domain name (FQDN) wkstn1.example.microsoft.com., the DNS computer name is the leftmost label client1.
- **Primary DNS suffixes**. A primary DNS suffix for the computer, which is placed after the computer or host name to form the FQDN. Using the previous example, the primary DNS suffix would be example.microsoft.com.
- **Connection-specific names.** Each network connections of a multihomed computer can be configured with a connection-specific DNS domain name
- **NetBIOS names**. NetBIOS names are used to support legacy Microsoft networking technology.
- **DNS servers list.** A list of DNS servers for clients to use when resolving DNS names, such as a preferred DNS server, and any alternate DNS servers to use if the preferred server is not available.
- **DNS suffix search list**. The DNS suffix search list or search method to be used by the client when it performs DNS query searches for short, unqualified domain names.

**Domain Names**

The domain name is used with the client computer name to form the fully qualified domain name (FQDN), known also as the full computer name. In general, the DNS domain name is the remainder of the FQDN that is not used as the unique host name for the computer.

For example, the DNS domain name used for a client computer could be the following: If the FQDN, or **Full computer name**, is wkstn1.example.microsoft.com, the domain name is the example.microsoft.com portion of this name.

DNS domain names have two variations — a DNS name and a NetBIOS name. The full computer name (a fully qualified DNS

name) is used during querying and location of named resources on your network. For earlier version clients, the NetBIOS name is used to locate various types of NetBIOS services that are shared on your network.

An example that shows the need for both NetBIOS and DNS names is the Net Logon service. In Windows Server 2003 DNS, the Net Logon service on a domain controller registers its service (SRV) resource records on a DNS server. For Windows NT Server 4.0 and earlier versions, domain controllers register a DomainName entry in Windows Internet Name Service (WINS) to perform the same registration and to advertise their availability for providing authentication service to the network.

When a client computer is started on the network, it uses the DNS resolver to query a DNS server for SRV records for its configured domain name. This query is used to locate domain controllers and provide logon authentication for accessing network resources. A client or a domain controller on the network optionally uses the NetBIOS resolver service to query WINS servers, attempting to locate DomainName [1C] entries to complete the logon process.

Your DNS domain names should follow the same standards and recommended practices that apply to DNS computer naming described in the previous section. In general, acceptable naming conventions for domain names include the use of letters A through Z, numerals 0 through 9, and the hyphen (-). The use of the period (.) in a domain name is always used to separate the discrete parts of a domain name, commonly known as labels. Each label corresponds to an additional level defined in the DNS namespace tree.

For most computers, the primary DNS suffix configured for the computer can be the same as its Active Directory domain name, although the two values can be different.

**Host Names**
Computers using the underlying TCP/IP protocol of a Windows-based network use an IP address, a 32-bit numeric value (in the case of IPv4) or a 128-bit numeric value (in the case of IPv6), to identify the computer network connection of network hosts. However, network users prefer to use memorable, alphanumeric names. To support this need, network resources in a Windows-based network are identified by both alphanumeric names and IP addresses. DNS and WINS are two name resolution mechanisms that enable the use of alphanumeric names, and convert these names into their respective IP addresses.

**NetBIOS vs. DNS Computer Names**
In networks running Windows NT 4.0 and earlier, users typically locate and access a computer on the network using a NetBIOS (Network Basic Input Output System) name. In Windows 2000, Windows XP, and Windows Server 2003 operating systems, users locate and access a computer using DNS. In this implementation of DNS, a computer is identified by its full computer name, which is a DNS fully qualified domain name (FQDN).

**Primary DNS Suffixes**
The full computer name is a concatenation of the single-label host name, such as hostcomputer, and a multilabel primary DNS suffix name, such as corp.example.com, which is the DNS name of the Active Directory domain to which the computer is joined. Using the host and primary DNS suffix examples, the full computer name is hostcomputer.corp.example.com.

The host name is the same as the computer name specified during the installation of Windows Server 2003 and is listed in System Properties. The primary DNS suffix name is the same as the domain name specified during installation of Windows Server 2003 and is listed in System Properties. The full computer name is also listed in System Properties.

In addition, connection-specific DNS suffixes can be applied to the separate network adapter connections used by a multihomed computer. Connection-specific DNS suffixes identify the host when it is connected to separate networks that use different domain names. When using connection-specific DNS suffixes, a full computer name is also a concatenation of the host name and a connection-specific DNS suffix.

Using its host name and DNS suffixes, a single computer can have its full computer name configured using two possible methods:

- A primary full computer name, which applies as the default full computer name for the computer and all of its configured network connections.

- A connection-specific full computer name, which can be configured as an alternate DNS domain name that applies only for a single network adapter installed and configured on the computer.

Note that when using Active Directory, by default, the primary DNS suffix portion of a computer's full computer name must be the same as the name of the Active Directory domain where the computer is located. To allow different primary DNS suffixes, a domain administrator may create a restricted list of allowed suffixes by creating the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is created and managed by the domain administrator using Active Directory Service Interfaces (ADSI) or the Lightweight Directory Access Protocol (LDAP).

**Connection-specific Names**

As shown in the following figure, a multihomed server computer named "host-a" can be named according to both its primary and connection-specific DNS domain names.

**Connection-specific DNS Names**




In this example, the server computer host-a attaches to two separate subnets — Subnet 1 and Subnet 2 — which are also linked at redundant points using two routers for additional paths between each subnet. Given this configuration, host-a provides access as follows through its separately named local area network (LAN) connections:

- The name "host-a.public.example.microsoft.com" provides access using LAN connection 1 over Subnet 1, a lower-speed (10 megabit) Ethernet LAN, for normal access to users who have typical file and print service needs.

- The name "host-a.backup.example.microsoft.com" provides access using LAN connection 2 over Subnet 2, a higher-speed (100 megabit) Ethernet LAN, for reserved access by server applications and administrators who have special needs, such as troubleshooting server networking problems, performing network-based backup, or replicating zone data between servers.

In addition to the connection-specific DNS names, the computer can also be accessible using either of the two LAN connections by specifying its primary DNS domain name, "host-a.example.microsoft.com".

When configured as shown, a computer can register resource records in DNS according to its three distinct names and sets of IP addresses, as shown in the following table:

**DNS Names**

| DNS Name | IP Addresses | Description |
|---|---|---|
| host-a.example.microsoft.com | 10.1.1.11, 10.2.2.22 | Primary DNS name for computer. The computer registers A and PTR resource records for all configured IP addresses under this name in the "example.microsoft.com" zone. |
| host-a.public.example.microsoft.com | 10.1.1.11 | Connection-specific DNS name for LAN connection 1, which registers A and PTR resource records for IP address 10.1.1.11 in the "public.example.microsoft.com" zone. |
| host-a.backup.example.microsoft.com | 10.2.2.22 | Connection-specific DNS name for LAN connection 2, which registers A and PTR resource records for IP address 10.2.2.22 in the "backup.example.microsoft.com" zone. |

When a computer changes between connections to different networks hosting different DNS domains, the host name does not need to be changed unless there is a host in the new DNS domain with the same host name. The primary DNS suffix for the computer can be changed from the old domain name to the new domain and the computer will register the new full computer name in DNS.

**Note**

- If you have any multihomed dynamic update clients and at least one adapter is using DHCP, configure the DHCP server to update resource records according to the request of the client. If the DHCP server is configured to register both A and PTR resource records, the DHCP server may replace all A resource records for the name it attempts to register. As a result, A resource records that correspond to the computer's other IP addresses might be deleted.

In Windows XP and Windows Server 2003 operating systems, the full computer name is viewed and set in the **Computer Name** tab of **System Properties**. Connection-specific DNS suffixes are configured in the **Advanced TCP/IP Settings** dialog box of the **Internet Protocol (TCP/IP) Properties** for a network connection.

**NetBIOS Names**

The NetBIOS name is created using the first 16 bytes of the host name, which is the maximum number of characters for NetBIOS names. The NetBIOS name is a 16-byte string that uniquely identifies a computer or service for network communication. If the DNS host name is 15 or fewer bytes, the NetBIOS name is the host name plus enough spaces to form a 15-byte name, followed by a 1-byte unique identifier. The sixteenth byte specifies the network service associated with the computer. When the computer name exceeds the maximum length for NetBIOS, the NetBIOS computer name is created by truncating the host name to form a 15-byte name, followed by a 1-byte unique identifier.

**Note**

- Because host names are encoded in UTF-8 format, they do not necessarily have only 1 byte per character. ASCII characters are 1 byte each, but the size of extended characters is more than 1 byte. For details of naming restrictions and a fuller description of UTF8, see “Naming Restrictions for Hosts and Domains” later in this chapter.

If you have an investment in using NetBIOS names to support legacy Microsoft networking technology, it is recommended that you revise NetBIOS computer names used on your network to prepare for migration to a standard DNS-only environment. This prepares your network well for long term growth and compatibility with future naming requirements. For example, if you use the same computer name for both NetBIOS and DNS resolution, consider converting any special characters in your current NetBIOS names that do not comply with DNS naming standards, such as the underscore (_). While these characters are permitted in NetBIOS names, they are more often incompatible with traditional DNS host naming requirements and most existing DNS resolver client software.

**Note**

- If WINS lookup is enabled for zones hosted by your Windows DNS servers, you need to use the same name for both NetBIOS and DNS computer naming. Otherwise, the results of clients attempting to query and resolve the names of these computers will be inconsistent.

The following table summarizes the differences between DNS and NetBIOS names using the example FQDN client1.example.com.

**DNS and NetBIOS Names**

| Name Type | Description |
|---|---|
| NetBIOS name | The NetBIOS name is used to uniquely identify the computer’s NetBIOS services. This unique NetBIOS name is resolved to the IP address of the computer through broadcast, WINS, or the LMHosts file. By default, the computer’s NetBIOS name is the same as the host name up to 15 bytes, plus any spaces necessary to make the name 15 bytes long, plus the service identifier.<br><br>For example, a NetBIOS name might be Client1. |
| Host name | The host name is the first label of a FQDN.<br><br>For example, the first label of the FQDN client1.example.com is client1. |
| Primary DNS suffix | Every Windows XP and Windows Server 2003 computer can be assigned a primary DNS suffix to be used in name resolution and name registration. You can view the primary DNS suffix for your computer from the **Computer Name** tab of **SystemProperties**.<br><br>The primary DNS suffix is also known as the primary domain name.<br><br>For example, the FQDN client1.example.com has the primary DNS suffix example.com. |

| | |
|---|---|
| Connection-specific DNS suffix | The connection-specific DNS suffix is a DNS suffix that is assigned to a network connection. The connection-specific DNS suffix is also known as an adapter-specific DNS suffix. For example, a connection-specific DNS suffix might be acquired01-ext.com. |
| Fully Qualified Domain Name (FQDN) | The FQDN is a DNS name that uniquely identifies the computer in the DNS namespace. By default, it is a concatenation of the host name, the primary DNS suffix, and a period. For example, an FQDN might be client1.example.com. |
| Full computer name | The full computer name is the FQDN for a Windows XP, Windows 2000 or Windows Server 2003 computer. It is the concatenation of the host name and primary DNS suffix (or host name and connection-specific DNS suffix). |

**DNS Servers List**

For DNS clients to operate effectively, a prioritized list of DNS name servers must be configured for each computer to use when processing queries and resolving DNS names. In most cases, the client computer contacts and uses its preferred DNS server, which is the first DNS server on its locally configured list. Listed alternate DNS servers are contacted and used when the preferred server is not available. For this reason, it is important that the preferred DNS server be appropriate for continuous client use under normal conditions.

For computers running Windows XP, the DNS server list is used by clients only to resolve DNS names. When clients send dynamic updates, such as when they change their DNS domain name or a configured IP address, they might contact these servers or other DNS servers as needed to update their DNS resource records.

By default, the DNS client on Windows XP does not attempt dynamic update over a Remote Access Service (RAS) or virtual private network (VPN) connection. By default, the Windows XP and Windows Server 2003 DNS Client service does not attempt dynamic update of top-level domain (TLD) zones. Any zone named with a single-label name is considered a TLD zone, for example, com, edu, blank, my-company. When DNS clients are configured dynamically using a DHCP server, it is possible to have a larger list of provided DNS servers. To effectively share the load when multiple DNS servers are provided in a DHCP options-specified list, you can configure a separate DHCP scope that rotates the listed order of DNS and WINS servers provided to clients.

**DNS Suffix Search List**

For DNS clients, you can configure a DNS domain suffix search list that extends or revises their DNS search capabilities. By adding additional suffixes to the list, you can search for short, unqualified computer names in more than one specified DNS domain. Then, if a DNS query fails, the DNS Client service can use this list to append other name suffix endings to your original name and repeat DNS queries to the DNS server for these alternate FQDNs.

For computers and servers, the following default DNS search behavior is predetermined and used when completing and resolving short, unqualified names.

When the suffix search list is empty or unspecified, the primary DNS suffix of the computer is appended to short unqualified names, and a DNS query is used to resolve the resultant FQDN. If this query fails, the computer can try additional queries for alternate FQDNs by appending any connection-specific DNS suffix configured for network connections.

If no connection-specific suffixes are configured or queries for these resultant connection-specific FQDNs fail, then the client can then begin to retry queries based on systematic reduction of the primary suffix (also known as devolution).

For example, if the primary suffix were “example.microsoft.com”, the devolution process would be able to retry queries for the short name by searching for it in the microsoft.“com” and “com” domains.

When the suffix search list is not empty and has at least one DNS suffix specified, attempts to qualify and resolve short DNS names are limited to searching only those FQDNs made available by the specified suffix list. If queries for all FQDNs that are formed as a result of appending and trying each suffix in the list are not resolved, the query process fails, producing a “name not found” result.

If the domain suffix list is used, clients continue to send additional alternate queries based on different DNS domain names when a query is not answered or resolved. Once a name is resolved using an entry in the suffix list, unused list entries are not tried. For this reason, it is most efficient to order the list with the most commonly used domain suffixes first.

Domain name suffix searches are used only when a DNS name entry is not fully qualified. To fully qualify a DNS name, a trailing period (**.**) is entered at the end of the name.

**Name Restrictions for Hosts and Domains**

Different DNS implementations allow different characters and lengths, and differ from NetBIOS naming restrictions. The

following table shows the restrictions for standard DNS names, DNS names in Windows Server 2003 and NetBIOS Names.

| Restriction | Standard DNS (Including Windows NT 4.0) | DNS in Windows Server 2003 | NetBIOS |
|---|---|---|---|
| Characters | Supports RFC 1123, which permits "A" to Z, "a" to "z", "0" to "9", and the hyphen (-) | Several different configurations are possible, as described at the end of this section | Permits Unicode characters, numbers, white space, symbols: ! @ $ % ^ & ) ( . - _ { } ~ |
| Fully qualified domain name length | Permits 63 bytes per label and 255 bytes for an FQDN | Permits 63 bytes per label and 255 bytes for an FQDN; the FQDN for an Active Directory domain name is limited to 64 bytes | Permits 16 bytes for a host name |

**Note**

- Although you can create long, complex DNS names, it is recommended that you create shorter, user-friendly names.

Adherence to DNS RFC 1123 can present a problem during an upgrade from Windows NT 4 to Windows 2000 or Windows Server 2003 because the host name of a Windows 2000 or Windows Server 2003 computer is set to the NetBIOS name that was set prior to the upgrade. NetBIOS names can use characters that violate the RFC 1123 DNS name restrictions, and converting NetBIOS names to strict DNS names can be time consuming.

Windows 2000 and Windows Server 2003 support migration by supporting a wider character set than RFC 1123. RFC 2181 enlarges the character set allowed in DNS names. It states that a DNS label can be any binary string, and it does not necessarily need to be interpreted as ASCII. Based on this definition, Microsoft has proposed that the DNS name specification be readjusted to accommodate a larger character set: UTF-8 character encoding, as described in RFC 2044.

UTF-8 character encoding is a superset of ASCII and a translation of the UCS-2 (also known as Unicode) character encoding. The UTF-8 character set includes characters from most of the world's written languages; this enables a far greater range of possible names. The Windows Server 2003 DNS service includes support for UTF-8 character encoding. For more information about UTF-8, see "Unicode Character Support" later in this document.

However, before using additional characters in DNS names, consider the following issues:

- Some third-party DNS client software supports only the characters listed in RFC 1123. Third-party DNS client software may not able to resolve the DNS names of computers with names that use characters outside the set supported by RFC 1123.

- A DNS server that does not support UTF-8 encoding might accept a zone transfer of a zone containing UTF-8 names, but it cannot write back those names to a zone file or reload those names from a zone file. Therefore, you must not transfer a zone that contains UTF-8 characters to a DNS server that does not support them.

- If you attempt to register a DNS name in Active Directory that contains an extended character that cannot be rendered in an LDAP distinguished name, Active Directory will respond with an invalid syntax error.

  You can configure the Windows Server 2003 DNS server to allow or reject the use of UTF-8 characters in DNS names. You can do this for each DNS server administered using the DNS console.

  **Note**

  - When you are modifying a host name or DNS suffix, or creating an Active Directory domain, if you enter a DNS name that includes UTF-8 or underscore characters not listed in RFC 1123, a warning message appears explaining that some DNS server implementations might not support these characters.

**Subnet Prioritization**

DNS subnet prioritization returns local IP addresses to the DNS Client service in preference to IP addresses on different subnets. This feature reduces network traffic by encouraging client computers to connect to network resources closer to them.

For example, suppose the three Web servers that host www.example.com are located on different subnets. On the DNS server, you can create the following resource records:

```
www.example.com.  IN A  172.16.64.11
www.example.com.  IN A  172.17.64.22
www.example.com.  IN A  172.18.64.33
```

When a user queries for www.example.com, all three resource records are returned. With subnet prioritization, the DNS Client service reorders the list of records returned so that it begins with the IP addresses from networks to which the computer is directly connected. For example, if a user with the IP address 172.17.64.93 queries for www.example.com, the DNS Client service returns the resource records in the following order:

```
www.example.com.  IN A  172.17.64.22
www.example.com.  IN A  172.16.64.11
www.example.com.  IN A  172.18.64.33
```

**Note**

- Subnet prioritization prevents the DNS Client service from using the Round Robin feature on the DNS Server service. For more information, see “Round Robin” later in this document.

Although subnet prioritization can reduce network traffic across subnets, you might prefer to use round robin as described in RFC 1794.

**DNS Server Service**

The DNS Server service is the component that provides the server implementation of DNS. The settings discussed in this section include:

- Disabling the use of recursion.
- Round robin use of resource records.
- Subnet prioritization.
- Advanced parameters.

**Disabling Recursion**

By default, recursion is enabled for the DNS Server service, and clients typically request that the server use recursion to resolve a name when sending a query. If recursion is disabled, the DNS Server service always uses referral, regardless of the client request.

In general, DNS servers can answer queries for names outside of their authoritative zones in two ways:

- Servers can send referral answers, which are an immediate response to the requesting client with a list of resource records for other DNS servers that it knows about that appear to be closer or more likely to be of help in resolving the queried name.
- Servers can use recursion to query other servers on behalf of the requesting client, attempting to fully resolve the name. Recursive lookups continue until the server receives an authoritative answer for the queried name. The server then forwards this answer in response to the original query from the requesting client.

In most cases, recursion is disabled on a DNS server when DNS clients are limited to resolving names authoritatively managed on a specific server. For example, this is the case when a DNS server has only DNS names data for an internal network or when the DNS server is incapable of resolving external DNS names (such as Internet DNS names) and clients are expected to retry another DNS server to resolve these names.

**Note**

- If you disable recursion on the DNS server, you will not be able to use forwarders on the same server. For more information about forwarders, see “Forwarding” later in this document.

**Round Robin**

Round robin DNS is a method of managing server congestion by distributing connection loads across multiple servers (containing identical content). Round robin works on a rotating basis in that one server IP address is handed out, then moves to the back of the list; the next server IP address is handed out, then it moves to the end of the list; and so on, depending on the number of servers being used. This works in a looping fashion.

This local balancing mechanism is used by DNS servers to share and distribute network resource loads. You can use round robin to rotate all resource record types contained in a query answer if multiple RRs are found.

By default, DNS uses round robin to rotate the order of RR data returned in query answers where multiple RRs of the same type exist for a queried DNS domain name. This feature provides a simple method for load balancing client use of Web servers and other frequently queried multihomed computers.

If round robin is disabled for a DNS server, the order of the response for these queries is based on a static ordering of RRs in the answer list as they are stored in the zone (either its zone file or Active Directory).

**Example: Round-robin rotation**
A forward lookup-type query (for all Host Address [A] RRs that match a DNS domain name) is made for a multihomed computer (multihomed.example.microsoft.com) that has three IP addresses. Separate A RRs are used to map the host's name to each of these IP addresses in the zone. In the stored example.microsoft.com zone, the RRs appear in this fixed order:

```
multihomed   IN  A  10.0.0.1
multihomed   IN  A  10.0.0.2
multihomed   IN  A  10.0.0.3
```

The first DNS client that queries the server to resolve this host's name receives the list in default order. When a second client sends a subsequent query to resolve this name, the list is rotated as follows:

```
multihomed   IN  A  10.0.0.2
multihomed   IN  A  10.0.0.3
multihomed   IN  A  10.0.0.1
```

**Restricting round robin rotation for selected RR types**
By default, DNS will perform round robin rotation for all RR types. You can now specify that certain RR types are not to be round-robin rotated in the registry. These modifications can be made in the registry.

**Restricting round-robin rotation for all RR types**
By default, all RR types are rotated, except those that have been specified as excluded from rotation in the registry.

**Subnet Prioritization**
By default, the DNS Server service uses local subnet prioritizing as the method for giving preference to IP addresses on the same network when a client query resolves to a host name that is mapped to more than one IP address. This feature requires that the client application attempt to connect to the host using its closest (and typically fastest) IP address available for connection.

The DNS Server service uses local subnet priority as follows:

1. The DNS Server service determines if local subnet prioritization is needed to order the query response.

   If more than one A resource record (RR) matches the queried host name, the DNS Server service can reorder the records by their subnet location. If the queried host name only matches a single A resource record, or if the IP network address of the client does not match an IP network address for any of the mapped addresses in an answer list of multiple RRs, no prioritizing is necessary.

2. For each RR in the matched answer list, the DNS Server service determines which records (if any) match the subnet location of the requesting client.

3. The DNS Server service reorders the answer list so that A RRs which match the local subnet of the requesting client are placed first in the answer list.

4. Prioritized by subnet order, the answer list is returned to the requesting client.

**Simple example: Local network prioritizing**
A multihomed computer, multihomed.example.microsoft.com, has three A RRs for its three separate host IP addresses in the

example.microsoft.com zone. A separate A RR is used for each of the host's addresses, which appear in this order in the zone:

```
multihomed   IN  A  192.168.1.27
multihomed   IN  A  10.0.0.14
multihomed   IN  A  172.16.20.4
```

If a DNS client resolver at IP address 10.4.3.2 queries the server for the IP addresses of host multihomed.example.microsoft.com, the DNS Server service notes that the originating IP network address (10.0.0.0) of the client matches the network (class A) portion of the 10.0.0.4 address in the answer list of RRs. The DNS Server service then reorders the addresses in the response as follows:

```
multihomed   IN  A  10.0.0.14
multihomed   IN  A  192.168.1.27
multihomed   IN  A  172.16.20.4
```

If the IP address of the requesting client has no local network match with any of the RRs in the answer list, then the list is not prioritized.

**Complex example: Local subnet prioritizing**
In Windows Server 2003, addresses are prioritized by matching the class C subnet by default, regardless of the subnet mask or address class of the target address.

For example, a multihomed computer, multihomed.example.microsoft.com, has four A RRs for four separate host IP addresses in the example.microsoft.com zone. Two of these IP addresses are for nonlocal networks. The other two IP addresses share a common IP network address but, because IP subnetting is used, represent different physical subnetted network connections based on their custom (nondefault) subnet mask value of 255.255.248.0. These example RRs appear in the following order in the zone:

```
multihomed   IN  A  192.168.1.27
multihomed   IN  A  172.16.22.4
multihomed   IN  A  10.0.0.14
multihomed   IN  A  172.16.31.5
```

If the IP address of the requesting client is 172.16.22.8, both of the IP addresses that match the same IP network as the client, the 172.16.0.0 network, are returned at the top of the answer list to the client. However, in this example, the 172.16.22.4 address is placed ahead of the 172.16.31.5 address because it matches the client IP address down through the 172.16.20.0 subnet address.

The reordered answer list returned by the DNS Server service would be:

```
multihomed   IN  A  172.16.22.4
multihomed   IN  A  172.16.31.5
multihomed   IN  A  192.168.1.27
multihomed   IN  A  10.0.0.14
```

### Note

- IP subnetting is imposed by using a custom or nondefault subnet mask value with all of the IP addresses on a network. Local subnet priority supersedes the use of round robin rotation for multihomed names. When round robin is enabled, however, RRs continue to be rotated using round robin as the secondary method of sorting the response list.

**Advanced DNS Server Service Parameters**
When initialized for service, DNS servers use server configuration settings taken from the parameters stated in a boot information file, the registry, and possibly zone information provided through Active Directory integration.

In most situations, the installation defaults are acceptable and should not require modification. However, when needed, you can use the DNS console to tune the following advanced parameters, accommodating special deployment needs and situations.

### DNS Server Service Advanced Parameters

| Value | Description |
|---|---|

| | |
|---|---|
| Disable recursion | Determines whether or not the DNS server uses recursion. By default, the DNS Server service is enabled to use recursion. |
| BIND secondaries | Determines whether to use fast transfer format when transferring a zone to DNS servers running legacy Berkeley Internet Name Domain (BIND) implementations.<br><br>By default, all Windows-based DNS servers use a fast zone transfer format, which uses compression and can include multiple records per TCP message during a connected transfer. This format is also compatible with more recent BIND-based DNS servers that run versions 4.9.4 and later. |
| Fail on load if bad zone data | Sets the DNS server to parse files strictly.<br><br>By default, the DNS Server service logs data errors, ignore any erred data in zone files, and continue to load a zone. This option can be reconfigured using the DNS console so that the DNS Server service logs errors and fails to load a zone file containing records data that is determined to have errors. |
| Enable round robin | Determines whether the DNS server uses round robin to rotate and reorder a list of resource records (RRs) if multiple RRs exist of the same type exisit for a query answer. |
| Enable netmask ordering | Determines whether the DNS server reorders A resource records within the same resource record set in its response to a query based on the IP address of the source of the query. |
| Secure cache against pollution | Determines whether the server attempts to clean up responses to avoid cache pollution. This setting is enabled by default.<br><br>By default, DNS servers use a secure response option that eliminates adding unrelated resource records included in a referral answer to their cache. In most cases, any names added in referral answers are typically cached and help expedite the speed of resolving subsequent DNS queries.<br><br>With this feature, however, the server can determine that referred names are potentially polluting or insecure and discard them. The server determines whether to cache the name offered in a referral on the basis of whether or not it is part of the exact related DNS domain name tree for which the original queried name was made.<br><br>For example, if a query was originally made for “example.microsoft.com” and a referral answer provided a record for a name outside of the “microsoft.com” domain name tree, such as msn.com, then that name would not be cached where this feature is enabled for use. |

**Resource Records in DNS**

DNS resource records are the data that is associated with DNS names in the DNS namespace. Each domain name of the DNS namespace tree contains a set of resource records, and each resource record in the set contains different types of information relating to the domain name. A DNS query includes the DNS domain name that is to be resolved and the type of information desired (the resource records that are requested). Queries for the IP addresses of DNS hosts return A resource records, and queries for the DNS servers authoritative for a DNS domain name return name server (NS) resource records.

Resource records are typically discussed in two categories: authority records and other records. Authority records identify the DNS servers that are authoritative for the domain names in the DNS namespace and how their zones should be managed, and all other DNS records contain information about a domain name that is unrelated to authority.

**Authority Records**

Zones are based on a concept of server authority. When a DNS server is configured to load a zone, it uses two types of resource records to determine the authoritative properties of the zone:

- First, the start of authority (SOA) resource record indicates the name of origin for the zone and contains the name of the server that is the primary source for information about the zone. It also indicates other basic properties of the zone.

- Next, the name server (NS) resource record is used to notate which DNS servers are designated as authoritative for the zone. By listing a server in the NS RR, it becomes known to others as an authoritative server for the zone. This means that any server specified in the NS RR is to be considered an authoritative source by others, and is able to answer with certainty any queries made for names included in the zone.

The SOA and NS resource records occupy a special role in zone configuration. They are required records for any zone and are

typically the first resource records listed in files.

**The SOA resource record**

The start of authority (SOA) resource record is always first in any standard zone. It indicates the DNS server that either originally created it or is now the primary server for the zone. It is also used to store other properties such as version information and timings that affect zone renewal or expiration. These properties affect how often transfers of the zone are done between servers authoritative for the zone.

The SOA resource record contains the following information:

**SOA Resource Record Fields**

| Field | Description |
|---|---|
| Primary server (owner) | The host name for the primary DNS server for the zone. |
| Responsible person | The e-mail address of the person responsible for administering the zone. A period (**.**) is used instead of an at sign (**@**) in this e-mail name. |
| Serial number | The revision number of the zone file. This number increases each time a resource record in the zone changes. It is important that this value increase each time the zone is changed, so that either partial zone changes or the fully revised zone can be replicated to other secondary servers during subsequent transfers. |
| Refresh interval | The time, in seconds, that a secondary DNS server waits before querying its source for the zone to attempt renewal of the zone. When the refresh interval expires, the secondary DNS server requests a copy of the current SOA record for the zone from its source, which answers this request. The secondary DNS server then compares the serial number of the source server's current SOA record (as indicated in the response) with the serial number in its own local SOA record. If they are different, the secondary DNS server requests a zone transfer from the primary DNS server. The default for this field is 900 seconds (15 minutes). |
| Retry interval | The time, in seconds, a secondary server waits before retrying a failed zone transfer. Normally, this time is less than the refresh interval. The default value is 600 seconds (10 minutes). |
| Expire interval | The time, in seconds, before a secondary server stops responding to queries after a lapsed refresh interval where the zone was not refreshed or updated. Expiration occurs because at this point in time, the secondary server must consider its local data unreliable. The default value is 86,400 seconds (24 hours). |
| Minimum (default) TTL | The default Time-To-Live (TTL) of the zone and the maximum interval for caching negative answers to name queries. The default value is 3,600 seconds (1 hour). |

The following is an example of a default SOA resource record:

```
@   IN  SOA     nameserver.example.microsoft.com.  postmaster.example.microsoft.com. (
                                1            ; serial number
                                3600         ; refresh   [1h]
                                600          ; retry     [10m]
                                86400        ; expire    [1d]
                                3600 )       ; min TTL   [1h]
```

In the example SOA record shown above, the primary or originating server for the zone is shown as nameserver.example.microsoft.com. The e-mail address for the person to contact regarding questions about this zone is postmaster.example.microsoft.com.

Periods are used to represent e-mail addresses when writing and storing DNS domain names in a zone. In an e-mail application, the previous example address would instead likely appear as postmaster@example.microsoft.com. The parentheses used in the SOA resource record as it appears in a zone file are used to enable wrapping of the record over multiple lines of text. If an individual TTL value is assigned and applied to a specified resource record used in the zone, it overrides the minimum (default) TTL set in the SOA record.

**The NS resource record**

Name server (NS) resource records can be used to assign authority to specified servers for a DNS domain name in two ways:

- By establishing a list of authoritative servers for the domain so that those servers can be made known to others that request information about this domain (zone).

- By indicating authoritative DNS servers for any subdomains that are delegated away from the zone.

In the case of assigning servers with host names in the same zone, corresponding address (A) resource records are normally used in the zone to resolve the names of specified servers to their IP addresses. For servers that are specified using this RR as part of a zone delegation to a subdomain, the NS resource record usually contains out-of-zone names. For the out-of-zone names to be resolved, A resource records for the specified out-of-zone server's might be needed. When these out-of-zone NS and A records are needed to provide delegation, they are known as glue records.

The following table shows the basic syntax of how a NS RR is used.

**Basic Syntax of a Name Server Resource Record**

| |
|---|
| **Description:** Used to map a DNS domain name as specified in **owner** to the name of hosts operating DNS servers specified in the *name_server_domain_name* field. |
| **Syntax:** *owner ttl* **IN NS** *name_server_domain_name*. |
| **Example:**<br>`example.microsoft.com.  IN NS nameserver1.example.microsoft.com` |

**Other Important Records**

After a zone is created, additional resource records need to be added to it. The following table lists the most common resource records (RRs) to be added.

**Common DNS Resource Records**

| Resource Record | Description |
|---|---|
| Host (A) | For mapping a DNS domain name to an IP address used by a computer. |
| Alias (CNAME) | For mapping an alias DNS domain name to another primary or canonical name. |
| Mail Exchanger (MX) | For mapping a DNS domain name to the name of a computer that exchanges or forwards mail. |
| Pointer (PTR) | For mapping a reverse DNS domain name based on the IP address of a computer that points to the forward DNS domain name of that computer. |
| Service location (SRV) | For mapping a DNS domain name to a specified list of DNS host computers that offer a specific type of service, such as Active Directory domain controllers. |

**Host (A) resource records**

Host (A) resource records are used in a zone to associate DNS domain names of computers (or hosts) to their IP addresses, and can be added manually, Windows clients and servers can also use the DHCP Client service to dynamically register and update their own A resource records in DNS when an IP configuration change occurs. DHCP-enabled client computers running earlier versions of Microsoft operating systems can have their A resource records registered and updated by proxy if they obtain their IP lease from a qualified DHCP server (only the Windows 2000 and Windows Server 2003 DHCP Server service currently supports this feature).

The host (A) resource record is not required for all computers, but is needed by computers that share resources on a network. Any computer that shares resources and needs to be identified by its DNS domain name, needs to use A resource records to provide DNS name resolution to the IP address for the computer.

Most A RRs that are required in a zone can include other workstations or servers that share resources, other DNS servers, mail servers, and Web servers. These resource records comprise the majority of resource records in a zone database.

**Alias (CNAME) resource records**

Alias (CNAME) resource records are also sometimes called canonical names. These records allow you to use more than one name to point to a single host, making it easy to do such things as host both an FTP server and a Web server on the same

computer. For example, the well-known server names (ftp, www) are registered using CNAME RRs that map to the DNS host name, such as "server-1" for the server computer that hosts these services.

CNAME RRs are recommended for use in the following scenarios:

- When a host specified in an A RR in the same zone needs to be renamed.
- When a generic name for a well-known server such as www needs to resolve to a group of individual computers (each with individual A RRs) that provide the same service. An example would be a group of redundant Web servers.

When renaming a computer with an existing A RR in the zone, you can use a CNAME RR temporarily, to allow a grace period for users and programs to switch from specifying the old computer name to using the new one. To do this, you need the following:

- For the new DNS domain name of the computer, a new A RR is added to the zone.
- For the old DNS domain name, a CNAME RR is added that points to the new A RR.
- The original A RR for the old DNS domain name (and its associated PTR RR if applicable) is removed from the zone.

When using a CNAME RR for aliasing or renaming a computer, set a temporary limit on how long the record is used in the zone before removing it from DNS. If you forget to delete the CNAME RR and later its associated A RR is deleted, the CNAME RR can waste server resources by trying to resolve queries for a name no longer used on the network.

The most common or popular use of a CNAME RR is to provide a permanent DNS aliased domain name for generic name resolution of a service-based name, such as www.example.microsoft.com to more than one computer or one IP address used in a Web server. For example, the following shows the basic syntax of how a CNAME RR is used.

```
alias_name IN CNAME primary_canonical_name
```

In this example, a computer named host-a.example.microsoft.com needs to function as both a Web server named "www.example.microsoft.com."; and an FTP server named "ftp.example.microsoft.com."; To achieve the intended use for naming this computer, you can add and use the following CNAME entries in the example.microsoft.com zone:

```
                host-a    IN  A      10.0.0.20
ftp       IN  CNAME  host-a
www       IN  CNAME  host-a
```

If you later decide to move the FTP server to another computer, separate from the Web server on "host-a", simply change the CNAME RR in the zone for ftp.example."microsoft.com" and add an additional A RR to the zone for the new computer hosting the FTP server.

Based on the earlier example, if the new computer were named "host-b.example.microsoft.com", the new and revised A and CNAME RRs would be as follows:

```
                host-a    IN  A      10.0.0.20
host-b    IN  A      10.0.0.21
ftp       IN  CNAME  host-b
www       IN  CNAME  host-a
```

**Mail exchanger (MX) resource records**

The mail exchanger (MX) RR is used by e-mail applications to locate a mail server based on a DNS domain name used in the destination address for the e-mail recipient of a message. For example, a DNS query for the name example.microsoft.com; could be used to find an MX RR, enabling an e-mail application to forward or exchange mail to a user with the e-mail address user@microsoft.com.

The MX RR shows the DNS domain name for the computer or computers that process mail for a domain. If multiple MX RRs exist, the DNS Client service attempts to contact mail servers in the order of preference from lowest value (highest priority) to highest value (lowest priority). The following shows the basic syntax for use of an MX RR.

```
mail_domain_name IN MX preference mailserver_host
```

By using the MX RRs shown below in the example.microsoft.com zone, mail addressed to user@example.microsoft.com is delivered to user@mailserver0.example.microsoft.com first if possible. If this server is unavailable, the resolver client can then use user@mailserver1.example.microsoft.com instead.

```
@          IN  MX   1    mailserver0
@          IN  MX   2    mailserver1
```

Note that the use of the at sign (**@**) in the records indicates that the mailer DNS domain name is the same as the name of origin (example.microsoft.com) for the zone.

**Pointer (PTR) resource records**
Pointer (PTR) RRs are used to support the reverse lookup process, based on zones created and rooted in the in-addr.arpa domain. These records are used to locate a computer by its IP address and resolve this information to the DNS domain name for that computer.

PTR RRs can be added to a zone in several ways:

- You can manually create a PTR RR for a static TCP/IP client computer using the DNS , either as a separate procedure or as part of the procedure for creating an A RR.
- Computers use the DHCP Client service to dynamically register and update their PTR RR in DNS when an IP configuration change occurs.
- All other DHCP-enabled client computers can have their PTR RRs registered and updated by the DHCP server if they obtain their IP lease from a qualified server. The Windows 2000 and Windows Server 2003 DHCP Server service provides this capability.

The pointer (PTR) resource record is used only in reverse lookup zones to support reverse lookup.

**Service location (SRV) resource records**
To locate Active Directory domain controllers, service location (SRV) RRs are required. Typically, you can avoid manual administration of the SRV RR when installing Active Directory.

By default, the Active Directory installation wizard attempts to locate a DNS server based on the list of preferred or alternate DNS servers, configured in any of its TCP/IP client properties, for any of its active network connections. If a DNS server that can accept dynamic update of the SRV RR (and other RRs related to registering Active Directory as a service in DNS) is contacted, the configuration process is complete.

If, during the installation, a DNS server that can accept updates for the DNS domain name used to name your Active Directory is not found, the wizard can install a DNS server locally and automatically configure it with a zone to support the Active Directory domain.

For example, if the Active Directory domain that you chose for your first domain in the forest was example.microsoft.com, a zone rooted at the DNS domain name of example.microsoft.com would be added and configured to use with the DNS server running on the new domain controller.

Whether or not you install the DNS Server service locally, a file (Netlogon.dns) is written and created during the Active Directory installation process that contains the SRV RRs and other RRs needed to support the use of Active Directory. This file is created in the systemroot\System32\Config folder.

If you are using a DNS server that fits one of the following descriptions, you should use the records in Netlogon.dns to manually configure the primary zone on that server to support Active Directory.

- The computer operating your DNS server is running on another platform, such as UNIX, and cannot accept or recognize dynamic updates.
- A DNS server at this computer that is not the DNS Server service provided with the Windows Server 2003 operating system is authoritative for the primary zone corresponding to the DNS domain name for your Active Directory domain.

- The DNS server supports the SRV RR, as defined in the Internet draft, "A DNS RR specifying the location of services (DNS SRV)", but does not support dynamic updates.

- For example, the DNS Server service provided with Windows NT Server 4.0, when updated to Service Pack 4 or later, fits this description.

In the future, the SRV RR might also be used to register and lookup other well-known TCP/IP services on your network if applications implement and support DNS name queries that specify this record type.

Other additional resource records are supported by Windows Server 2003 DNS and are used less frequently in most zones. These additional types of resource records can be added as needed using the DNS console. For more information about supported resource records, see "Related Information" later in this document.

**DNS-related Files**
The following files relate to using and configuring DNS servers and clients.

| File | Description |
|---|---|
| Boot | BIND boot configuration file. This file is not created by the DNS console. However, as an optional configuration for the DNS Server service, it can be copied from another DNS server running the Berkeley Internet Name Domain (BIND) server implementation of DNS. To use this file with the DNS Server service, you need to click From file in Server properties. On BIND servers, this file is often called the "named.boot" file. |
| Cache.dns | Used to preload resource records into the DNS server names cache. DNS servers use this file to help locate root servers on either your network or the Internet.<br><br>By default, this file contains DNS resource records that prime the local cache of the server with the addresses of authoritative root servers for the Internet. If you are setting up a DNS server to resolve Internet DNS names, the information in this file is required unless you enable the use of another DNS server as a forwarder to resolve these names.<br><br>Traffic to the Internet root servers is heavy, but because host names are not usually resolved at this level, the load can be reasonably handled. Instead, the root hints file provides referral information that can be useful during DNS name resolution to redirect a query to other servers that are authoritative for names located beneath the root.<br><br>For DNS servers operating privately on your internal network, the DNS console can learn and replace the contents of this file with internal root servers on your network, provided they are reachable through the network when you are setting up and configuring new DNS servers. It can be updated using the DNS console from the Root Hints tab located under the applicable server properties.<br><br>This file preloads the server names cache when it is started. |
| Root.dns | Root zone file. This file can appear at a DNS server if it is configured as a root server for your network. |
| *zone_name*.dns | Used when a standard zone (either primary or secondary) is added and configured for the server. Files of this type are not created or used for primary type zones that are directory-integrated, which are stored in the Active Directory database. |

These files can be found in the systemroot\System32\Dns folder on the server computer.

**Zones and Zone Transfer**
DNS distributes the DNS namespace database using DNS zones, which store name information about one or more DNS domains. There are three types of DNS zones supported in Windows Server 2003:

- Primary zone. Original copy of a zone where all resource records are added, modified, and deleted.

- Secondary zone. Read-only copy of the primary zone that is created and updated by transferring zone data from the primary zone.

- Stub zone. Read-only copy of the primary zone containing only the DNS resource records for the DNS servers listed in the zone (SOA, NS, and glue A resource records).

**Difference Between Zones and Domains**
A zone starts as a storage database for a single DNS domain name. If other domains are added below the domain used to create the zone, these domains can either be part of the same zone or belong to another zone. Once a subdomain is added, it can then either be:

- Managed and included as part of the original zone records, or
- Delegated away to another zone created to support the subdomain

For example, the following figure shows the microsoft.com domain, which contains domain names for Microsoft. When the microsoft.com domain is first created at a single server, it is configured as a single zone for all of the Microsoft DNS namespace. If, however, the microsoft.com domain needs to use subdomains, those subdomains must be included in the zone or delegated away to another zone.

**DNS Domain Names and Subdomain Names**




In this example, the example.microsoft.com domain shows a new subdomain — the example.microsoft.com domain — delegated away from the microsoft.com zone and managed in its own zone. However, the microsoft.com zone needs to contain a few resource records to provide the delegation information that references the DNS servers that are authoritative for the delegated example.microsoft.com subdomain.

If the microsoft.com zone does not use delegation for a subdomain, any data for the subdomain remains part of the microsoft.com zone. For example, the subdomain dev.microsoft.com is not delegated away but is managed by the microsoft.com zone.

**Why Zone Replication and Zone Transfers Are Needed**
Because of the important role that zones play in DNS, it is intended that they be available from more than one DNS server on the network, to provide availability and fault tolerance when resolving name queries. Otherwise, if a single server is used and that server is not responding, queries for names in the zone can fail. For additional servers to host a zone, zone transfers are required to replicate and synchronize all copies of the zone used at each server configured to host the zone.

When a new DNS server is added to the network and is configured as a new secondary server for an existing zone, it performs a full initial transfer of the zone to obtain and replicate a full copy of resource records for the zone. For most earlier DNS server implementations, this same method of full transfer for a zone is also used when the zone requires updating after changes are made to the zone. For DNS servers running Windows Server 2003, the DNS service supports incremental zone transfer, a revised DNS zone transfer process for intermediate changes.

**Domain Delegation**
A subordinate domain name, or subdomain, can be delegated from the DNS zone where its parent name is stored to a zone on another DNS server. When deciding whether to delegate your DNS namespace to make additional zones, consider the following reasons to use additional zones:

- A need to delegate management of part of your DNS namespace to another location or department within your organization.
- A need to divide one large zone into smaller zones for distributing traffic loads among multiple servers
- A need to extend the namespace by adding numerous subdomains at once, such as to accommodate the opening of a new branch or site.

If, for any of these reasons, you could benefit from delegating subordinate domain names to additional zones, it might make sense to restructure your namespace by adding additional zones. When choosing how to structure zones, you should use a plan that reflects the structure of your organization.

When delegating zones within your namespace, be aware that for each new zone you create, you will need delegation records in other zones that point to the authoritative DNS servers for the new zone. This is necessary both to transfer authority and to provide correct referral to other DNS servers and clients of the new servers being made authoritative for the new zone.

When a standard primary zone is first created, it is stored as a text file containing all resource record information about a single DNS server. This server acts as the primary master for the zone. Zone information can be replicated to other DNS servers to improve fault tolerance and server performance.

When structuring your zones, there are several good reasons to use additional DNS servers for zone replication:

- Additional DNS servers provide zone redundancy, enabling DNS names in the zone to be resolved for clients if a primary server for the zone stops responding.
- Additional DNS servers can be placed so as to reduce DNS network traffic. For example, adding a DNS server to the opposing side of a low-speed wide area network (WAN) link can be useful in managing and reducing network traffic.
- Additional DNS secondary servers can be used to reduce loads on a primary server for a zone.

**Example: Delegating a subdomain to a new zone**
As shown in the following figure, when a new zone for a subdomain (example.microsoft.com) is created, delegation from the parent zone (microsoft.com) is needed.

**Delegating a Subdomain**




In this example, an authoritative DNS server computer for the newly delegated example.microsoft.com subdomain is named based on a derivative subdomain included in the new zone (ns1.na.example.microsoft.com). To make this server known to others outside of the new delegated zone, two RRs are needed in the microsoft.com zone to complete delegation to the new zone.

These RRs include:

- An NS RR to effect the delegation. This RR is used to advertise that the server named ns1.na.example.microsoft.com is an authoritative server for the delegated subdomain.
- An A RR (also known as a glue record) is needed to resolve the name of the server specified in the NS RR to its IP address. The process of resolving the host name in this RR to the delegated DNS server in the NS RR is sometimes referred to as glue chasing.

  **Note**

  - When zone delegations are correctly configured, normal zone referral behavior can sometimes be circumvented

if you are using forwarders in your DNS server configuration. For more information, see "Forwarding" later in this document.

**Incremental Zone Transfers**

Incremental zone transfers are described in RFC 1995 as an additional DNS standard for replicating DNS zones. When incremental transfers are supported by both a DNS server acting as the source for a zone and any servers that copy the zone from it, the incremental transfer provides a more efficient method of propagating zone changes and updates.

In earlier DNS implementations, any request for an update of zone data required a full transfer of the entire zone database using an AXFR query. With incremental transfer, an alternate query type (IXFR) can be used instead. This allows the secondary server to pull only those zone changes it needs to synchronize its copy of the zone with its source, either a primary or secondary copy of the zone maintained by another DNS server.

With IXFR zone transfers, differences between the source and replicated versions of the zone are first determined. If the zones are identified to be the same version — as indicated by the serial number field in the SOA resource record of each zone — no transfer is made.

If the serial number for the zone at the source is greater than at the requesting secondary server, a transfer is made of only those changes to RRs for each incremental version of the zone. For an IXFR query to succeed and changes to be sent, the source DNS server for the zone must keep a history of incremental zone changes to use when answering these queries. The incremental transfer process requires substantially less traffic on a network and zone transfers are completed much faster.

**Example: Zone Transfer**

A zone transfer might occur during any of the following scenarios:

- When the refresh interval expires for the zone
- When a secondary server is notified of zone changes by its master server
- When the DNS Server service is started at a secondary server for the zone
- When the DNS console is used at a secondary server for the zone to manually initiate a transfer from its master server

Zone transfer requests are always initiated at the secondary server for a zone, and then sent to their configured master servers which act as their source for the zone. Master servers can be any other DNS server that loads the zone, such as either the primary server for the zone or another secondary server. When the master server receives the request for the zone, it can reply with either a partial or full transfer of the zone to the secondary server.

**Zone transfer process**

As shown in the following figure, zone transfers between servers follow an ordered process. This process varies depending on whether a zone has been previously replicated, or if initial replication of a new zone is being performed.

### Zone Transfer Process




1. During new configuration, the destination server sends an initial "all zone" transfer (AXFR) request to the master DNS server configured as its source for the zone.
2. The master (source) server responds and fully transfers the zone to the secondary (destination) server.

   The zone is delivered to the destination server requesting the transfer with its version established by use of a Serial number field in the properties for the start of authority RR. The SOA RR also contains a stated refresh interval in seconds

(by default, 900 seconds or 15 minutes) to indicate when the destination server should next request to renew the zone with the source server.

3. When the refresh interval expires, an SOA query is used by the destination server to request renewal of the zone from the source server.

4. The source server answers the query for its SOA record.

   This response contains the serial number for the zone in its current state at the source server.

   The destination server checks the serial number of the SOA record in the response and determines how to renew the zone.

   If the value of the serial number in the SOA response is equal to its current local serial number, it concludes that the zone is the same at both servers and that a zone transfer is not needed. The destination server then renews the zone by resetting its refresh interval based on the value of this field in the SOA response from its source server.

   If the value of the serial number in the SOA response is higher than its current local serial number, it concludes that the zone has been updated and that a transfer is needed.

5. If the destination server concludes that the zone has changed, it sends an IXFR query to the source server, containing its current local value for the serial number in the SOA record for the zone.

6. The source server responds with either an incremental or full transfer of the zone.

   If the source server supports incremental transfer by maintaining a history of recent incremental zone changes for modified resource records, it can answer with an incremental zone transfer (IXFR) of the zone.

   If the source server does not support incremental transfer, or does not have a history of zone changes, it can answer with a full (AXFR) transfer of the zone instead.

   **Note**

   - For servers running Windows 2000 and Windows Server 2003, incremental zone transfer through IXFR query is supported. For earlier versions of the DNS service and for many other DNS server implementations, incremental zone transfer is not available and only full-zone (AXFR) queries and transfers are used to replicate zones.

**DNS Notify**

Windows-based DNS servers support DNS Notify, an update to the original DNS protocol specification that permits a means of initiating notification to secondary servers when zone changes occur (RFC 1996). DNS notification implements a push mechanism for notifying a select set of secondary servers for a zone when the zone is updated. Servers that are notified can then initiate a zone transfer as described above, to pull zone changes from their master servers and update their local replicas of the zone.

For secondaries to be notified by the DNS server acting as their configured source for a zone, each secondary server must first have its IP address in the notify list of the source server. When using the DNS console, this list is maintained in the Notify dialog box, which is accessible from the Zone Transfer tab located in zone Properties.

In addition to notifying the listed servers, the DNS console permits you to use the contents of the notify list as a means to restrict or limit zone transfer access to only those secondary servers specified in the list. This can help prevent an undesired attempt by an unknown or unapproved DNS server to pull, or request, zone updates.

**DNS notification process**

The following is a brief summary of the typical DNS notification process for zone updates:

- The local zone at a DNS server acting as a master server, a source for the zone to other servers, is updated. When the zone is updated at the master or source server, the serial number field in the SOA RR is also updated, indicating a new local version of the zone.

- The master server sends a DNS notify message to other servers that are part of its configured notify list.
- All secondary servers that receive the notify message can then respond by initiating a zone transfer request back to the notifying master server.

The normal zone transfer process can then continue as described in the previous section.

You cannot configure a notify list for a stub zone.

Use DNS notification only to notify servers operating as secondary servers for a zone. For replication of directory-integrated zones, DNS notification is not needed. This is because any DNS servers that load a zone from Active Directory automatically poll the directory (as specified by the SOA resource record's Refresh Interval) to update and refresh the zone. In these cases, configuring a notify list can actually degrade system performance by causing unnecessary additional transfer requests for the updated zone.

**Note**

- By default, the Windows Server 2003 DNS Server service will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.

**Stub Zones**

In Windows Server 2003, a new zone type named stub zone has been introduced. A stub zone is a copy of a zone that contains only those resource records necessary to identify the authoritative DNS servers for that zone. A stub zone is used to resolve names between separate DNS namespaces. The need for this type of resolution can occur when a corporate merger requires that the DNS servers for two separate DNS namespaces resolve names for clients in both namespaces.

A stub zone consists of:

- The start of authority (SOA) resource record, name server (NS) resource records, and the glue A resource records for the delegated zone.
- The IP address of one or more master servers that can be used to update the stub zone.

The master servers for a stub zone are one or more DNS servers authoritative for the child zone, usually the DNS server hosting the primary zone for the delegated domain name.

**Stub zone resolution**

When a DNS client performs a recursive query operation on a DNS server hosting a stub zone, the DNS server uses the resource records in the stub zone to resolve the query. The DNS server sends an iterative query to the authoritative DNS servers specified in the NS resource records of the stub zone as if it were using NS resource records in its cache. If the DNS server cannot find the authoritative DNS servers in its stub zone, the DNS server hosting the stub zone attempts standard recursion using its root hints.

The DNS server will store the resource records it receives from the authoritative DNS servers listed in a stub zone in its cache, but it will not store these resource records in the stub zone itself; only the SOA, NS, and glue A resource records returned in response to the query are stored in the stub zone. The resource records stored in the cache are cached according to the Time-to-Live (TTL) value in each resource record. The SOA, NS, and glue A resource records, which are not written to cache, expire according to the expire interval specified in the stub zone's SOA record, which is created during the creation of the stub zone and updated during transfers to the stub zone from the original, primary zone.

If the query was an iterative query, the DNS server returns a referral containing the servers specified in the stub zone.

**Stub zone updates**

Stub zone updates involve the following conditions:

- When a DNS server loads a stub zone, it queries the zone's master server for the SOA resource record, NS resource records at the zone's root, and glue A resource records.
- During updates to the stub zone, the master server is queried by the DNS server hosting the stub zone for the same resource record types requested during the loading of the stub zone.

- The Refresh interval of the SOA resource record determines when the DNS server hosting the stub zone will attempt a zone transfer (update).
- If an update fails, the Retry interval of the SOA resource record determines when the update is retried.
- Once the Retry interval has expired without a successful update, the expiration time as specified in the Expires field of the SOA resource record determines when the DNS server stops using the stub zone data.

**Root Hints**

Root hints are used to prepare servers authoritative for non-root zones so that they can learn and discover authoritative servers that manage domains located at a higher level or in other subtrees of the DNS domain namespace. These hints are essential for servers authoritative at lower levels of the namespace when locating and finding servers under these conditions.

For example, suppose a DNS server (Server A) has a zone called sub.example.microsoft.com. In the process of answering a query for a higher-level domain, such as the example.microsoft.com domain, Server A needs some assistance to locate an authoritative server (such as Server B) for this domain.

In order for Server A to find Server B, or any other servers that are authoritative for the microsoft.com domain, it needs to be able to query the root servers for the DNS namespace. The root servers can then refer Server A to the authoritative servers for the com domain. The servers for the com domain can, in turn, offer referral to Server B or other servers that are authoritative for the microsoft.com domain.

The root hints used by Server A must have helpful hints to the root servers for this process to locate Server B (or another authoritative server) as intended.

To configure and use root hints correctly, first determine how the following applies to your DNS servers:

- Are you using DNS on the Internet or on a private network?
- Is the server used as a root server?

By default, the DNS Server service implements root hints using a file, Cache.dns, stored in the systemroot\System32\Dns folder on the server computer. This file normally contains the NS and A resource records for the Internet root servers. If, however, you are using the DNS Server service on a private network, you can edit or replace this file with similar records that point to your own internal root DNS servers.

Another server configuration in which root hints are treated differently is one in which a DNS server is configured to be used by other DNS servers in an internal namespace as a forwarder for any DNS queries of names managed externally (the Internet, for example). Even though the DNS server used as a forwarder can be located internally on the same network as servers using it as a forwarder, it needs hints for the Internet root servers to work properly and resolve external names.

**Note**

- If you are operating internal root servers, do not use root hints. Instead, delete the Cache.dns file entirely for any of your root servers.

**EDNS0**

Extension Mechanisms for DNS (EDNS0 as defined in RFC 2671) allow DNS requestors to advertise the size of their UDP packets and facilitate the transfer of packets larger than 512 octets, the original DNS restriction for UDP packet size (RFC 1035). When a DNS server receives a request over the UDP Transport Layer, it identifies the requestor's UDP packet size from the option (OPT) RR and scales its response to contain as many resource records as are allowed in the maximum UDP packet size specified by the requestor.

Windows Server 2003 DNS support for EDNS0 is enabled by default. It can be disabled using the registry. Locate the following registry subkey:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

Add the entry **EnableEDNSProbes** to the subkey. Give the entry a DWORD value and set it to 0x0 to disable EDNS0.

Use extreme caution when editing the registry Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the

system..

For more information about resource records, see "Related Information" at the end of this document.

**ENDS0 UDP responses**
Before a DNS server assumes that the requestor supports EDNS0, the DNS server must receive a query containing an OPT resource record. An OPT record contains no actual DNS data and its contents relate to the UDP Transport Layer message only. The OPT record stores the sender's UDP payload size in its CLASS field and lists the number of octets in the largest UDP payload that the requestor can deliver in the requestor's network.

When the DNS server receives a query containing an OPT record advertising the maximum UDP packet size, it will truncate any UDP response's size larger than the limit specified in the OPT record.

By default, the DNS server includes OPT resource records indicating its UDP maximum in responses to queries containing OPT resource records.

If the DNS server receives a query that does not contain an OPT resource record, it assumes the requestor's server does not support EDNS0 and will respond to the requestor assuming that the sender does not accept UDP packets larger than 512 octets. In this case, the DNS server will truncate its UDP response size to a maximum of 512 octets.

**EDNS0 UDP queries**
Before the requesting DNS server sends a query, it checks its cache to identify if the responding DNS server supports EDNS0. If the responding DNS server supports EDNS0, the requesting DNS server attaches an OPT resource record to the additional section of the query it sends. (All queries have five sections: header, question, answer, authority, and additional.) If, according to the requesting DNS server's cache, the responding DNS server does not support EDNS0, the requesting DNS server will not attach an OPT resource record to the query before it is sent.

**Identifying and caching EDNS0 support**
When the DNS server receives a request or response from a host containing an OPT record, the DNS server caches the EDNS version supported by the host (such as EDNS0). If there is no OPT record in a request or response from a host, the DNS server's cache will indicate that the host does not support EDNS0. If the cache already indicates that the host does support ENDS0, then cache will not be changed.

The default value for how long a host's EDNS0 support information is cached is 86400 (one day specified in seconds). This value can be modified in the registry.

The DNS server decides that a host does not support EDNS0 when it requests an OPT resource record and receives a response containing one of the following RCODE values in the header, as shown in the following table.

### EDNS0 Failure Code Values

| Name | Value | Description |
|---|---|---|
| FORMERR | 1 | Format Error. The name server did not interpret the OPT resource record. |
| SERVFAIL | 2 | Server Failure. The name server did not process the query because of a problem with the name server. |
| NOTIMPL | 4 | Not Implemented. The name server does not support the kind of query requested. |

(The RCODE field, or response code field, is a 4-bit field set in the header section as part of responses.) In this situation (as a requester), the DNS server identifies that the server does not support EDNS0 and caches this information.

### Note

- When considering packet sizes, you should take account of the network transmission path's discovered Maximum Transmission Unit (MTU), if this information is available. When configuring the UDP packet size to be larger than 512 octets, remember that the UDP packets must travel through devices other than UDP hosts and these devices, such as routers, may not support UDP packets larger than 512 octets. The maximum UDP packet size should always be compared with the MTU, which in some cases may be smaller. It is recommended that you establish the maximum UDP packet length support for all devices, and configure your UDP hosts for this maximum.

**DNSSEC**
Windows Server 2003 DNS provides basic support of the DNS Security Extensions (DNSSEC) protocol as defined in RFC 2535. The current feature support allows DNS servers to perform as secondary DNS servers for existing DNSSEC–compliant, secure zones. DNS supports the storing and loading of the DNSSEC-specific resource records. Currently, a DNS server is not capable

of signing zones and resource records (creating cryptographic digital signatures) or validating the SIG RRs. The DNSSEC resource records are KEY, SIG, and NXT. For more information about resource records, see the "DNS Resource Records Reference" later in this document. For more information about related RFCs, see "Related Information" at the end of this document.

**Server support for DNSSEC**
When loading a zone containing DNSSEC resource records, the DNS server loads these records along with all other types of resource records contained in the zone. When receiving a zone transfer containing DNSSEC resource records (SIG, KEY, NXT), the DNS server writes these records to the zone storage (zone data file or Active Directory) along with all other resource records.

When a DNS server receives a request or response containing DNSSEC resource records, it does not verify the digital signatures but caches the response and uses it for ensuing queries. When a DNS server receives a request for a resource record in a zone also containing DNSSEC resource records, it attaches the appropriate DNSSEC records to the response.

When a signed zone contains resource records for an owner name, including a CNAME resource record for that name, the DNS server will return the DNSSEC resource records associated with the owner name and the CNAME resource record's alias name. The DNS server will not suppress the retrieval of the CNAME resource record, and it will not return a SIG resource record for the canonical name. Rather, it will return the SIG resource record for the alias name.

**Client support for DNSSEC**
The DNS client does not read and store a key for the trusted zone and, consequently, it does not perform any cryptography, authentication, or verification. When a resolver initiates a DNS query and the response contains DNSSEC resource records, programs running on the DNS client will return these records and cache them in the same manner as any other resource records. This is the extent to which Windows XP DNS clients support DNSSEC. When the DNS client receives the SIG RR relating to the RRset, it will not perform an additional query to obtain the associated KEY record or any other DNSSEC records.

Resolvers do not authenticate resource records by verifying the signature information contained in the SIG resource record. The DNS client does not contain any information to indicate which resource records have been authenticated or to what extent they have been authenticated.

When a resolver receives a response or performs a query operation, it does not recognize the checking disabled (CD) query header bit, which in DNSSEC indicates that the data is authenticated by the server according to its policies, nor does it set the authentic data query header bit, which in DNSSEC indicates that unauthenticated data is acceptable to the resolver.

### Note

- If there is a signing agent running on the DNS server running Windows Server 2003 that signs the zone resource records, this DNS server may also be used as a primary server for DNSSEC-compliant secure zones.

**Active Directory Integration**
The DNS Server service is integrated into the design and implementation of Active Directory. Active Directory provides an enterprise-level tool for organizing, managing, and locating resources in a network.

In addition to supporting a conventional way of maintaining and replicating DNS zone files, the implementation of DNS in Windows Server 2003 has the option of using Active Directory as the data storage and replication engine for DNS. DNS is the domain controller location mechanism for Active Directory.

**Benefits of Integrating DNS with Active Directory**
Using Active Directory as the data storage and replication engine provides the following benefits:

- DNS replication is performed by Active Directory, so there is no need to support a separate replication topology for DNS servers.
- Active Directory service replication provides per-property replication.
- Active Directory service replication is secure.
- A primary DNS server is eliminated as a single point of failure. Original DNS replication is single-master and relies on a primary DNS server to update all the secondary servers. Unlike original DNS replication, Active Directory service replication is multimaster; an update can be made to any domain controller in it, and the change will be propagated to other domain controllers. In this way if DNS is integrated into Active Directory service the replication engine will always synchronize the DNS zone information.

Thus Active Directory integration significantly simplifies the administration of a DNS namespace. At the same time, standard zone transfers to other DNS servers (DNS servers other than Windows 2000 Server or Windows Server 2003 DNS servers, as well as earlier versions of Microsoft DNS servers) is still supported.

**How DNS Integrates with Active Directory**

When you install Active Directory on a server, you promote the server to the role of a domain controller for a specified domain. When completing this process, you are prompted to specify a DNS domain name for the Active Directory domain for which you are joining and promoting the server.

### Note

- When specifying the DNS domain name for the first Active Directory domain in the first Active Directory forest, known as the forest root domain, Windows Server 2003 does not support single-label domain names.

If during this process, a DNS server authoritative for the domain that you specified either cannot be located on the network or does not support the DNS dynamic update protocol, you are prompted with the option to install a DNS server. This option is provided because a DNS server is required to locate this server or other domain controllers for members of an Active Directory domain.

Once you have installed Active Directory, you have two options for storing and replicating your zones when operating the DNS server at the new domain controller:

- Standard zone storage, using a text-based file.

  Zones stored this way are located in .dns files that are stored in the systemroot\System32\Dns folder on each computer operating a DNS server. Zone file names correspond to the name you choose for the zone when creating it, such as "example.microsoft.com.dns" if the zone name was "example.microsoft.com."

- Directory-integrated zone storage, using the Active Directory database.

  Zones stored this way are located in the Active Directory tree under the domain or application directory partition. Each directory-integrated zone is stored in a dnsZone container object identified by the name you choose for the zone when creating it.

Zones stored in text files are typically referred to as file-backed zonesand zones stored in Active Directory are referred to as Active Directory-integrated zones.

In Active Directory-integrated DNS, each DNS zone becomes an Active Directory container object (DnsZone). The DnsZone object contains a DnsNode leaf object for every unique name within that zone. The DnsNode object has a DnsRecord multivalued attribute with an instance of a value for every record associated with the object's name.

### Note

- Active Directory-integrated zones can only be loaded onto a domain controller when the Windows Server 2003 DNS Server service is running on the domain controller.

**Replication Model**

When DNS zone information is stored in Active Directory and an update is made to a DNS server, it writes the update data to Active Directory. Active Directory is now responsible for replicating the data to other domain controllers. The DNS servers running on other domain controllers will poll the updates from the Active Directory.

Because Active Directory service uses the multimaster replication model, DNS updates can be written to any Active Directory-integrated DNS server, and the data will automatically be replicated across all the domain controllers.

The ability to write to Active Directory from multiple domain controllers at the same time can create a conflicting situation where the changes are made to the same object on two different DNS servers. The conflict will eventually be resolved in favor of the last update made to the object based on the timestamps of the updates. The same rule is applied in the case where two or more nodes with the same name are created on two or more DNS servers. Until the conflict is resolved and the DNS server, containing invalid update, polls the valid data from the Active Directory, it is possible that requests for the same object made to two different DNS servers will be resolved differently. This is why the Active Directory database is called *loosely consistent*.

Active Directory-integration provides an advantage over file-backed zones. With file-backed zones, only the primary authoritative DNS server for a zone can modify the zone. With Active Directory-integration, all domain controllers for the Active Directory domain where the zone is stored can modify the zone and then replicate the changes to other domain controllers. This replication process is called multimaster replication because multiple domain controllers can update the zone.

Windows Server 2003 domain controllers replicate resource records contained in Active Directory–integrated zones using Active Directory replication. Zones stored in Active Directory can also be transferred to secondary servers to create secondary zones in the same way that file-backed zones are transferred.

Active Directory–integration provides numerous benefits, including fault tolerance, security, simplified management, and more efficient replication of zones.

**Active Directory-integrated Zones**
When you configure a primary zone to be Active Directory–integrated, the zone is stored in Active Directory.

The following figure shows this configuration.

**Active Directory–integrated Zone**




The DNS server component contains only a copy of the zone. When DNS starts up, it reads a copy of the zone from Active Directory (step 1). Then, when the DNS server receives a change, it writes the change to Active Directory (step 2).

Through Active Directory replication, the zone is replicated to other domain controllers in the same domain. Also, through standard zone transfer, the DNS server can send its copy of the zone to any secondary DNS servers that request it. The DNS server can perform both incremental and full zone transfers. The following figure shows how the same zone can be replicated by using both Active Directory replication and standard zone transfer.

**Active Directory Replication and Zone Transfer**




By default, when the DNS Server service starts on a domain controller, it checks whether Active Directory is available and if it contains any DNS zones. If Active Directory does have zones, the DNS Server service loads them from Active Directory. The DNS Server service automatically writes back to the boot file at regular intervals. The boot file can be updated manually by using the DNS console.

The DNS Server service also loads the root hints and server and zone parameters from different locations depending on its settings. The table below shows the locations from which the DNS Server service loads and to which it writes zones, root hints, and server and zone parameters.

**Note**

- Changing the startup type is not recommended and could result in DNS infrastructure errors.

**How the DNS Server Loads Zones, Root Hints, and Parameters**

| Task | Load Data On Startup Set To: From File | Load Data On Startup Set To: From Registry | Load Data On Startup Set To: From Active Directory and Registry |
|---|---|---|---|
| **Read root hints from:** | Root hints file | If available, the root hints file. Otherwise, if the Directory is available and contains root hints, the directory | If the directory is available and contains root hints, from the directory. Otherwise, from the root hints file |
| **Write root hints to:** | Root hints file | Root hints file | If the directory is available, the directory |
| **Read zones from:** | Boot file, to get list of zones, then from zone files | Registry | The directory (for Active Directory–integrated zones) and the registry |
| **Write zones to:** | Boot file and the registry | Registry and, if the zone is Active Directory–integrated, the directory | Registry and, if the zone is Active Directory–integrated, the directory |
| **Read server and zones parameters from:** | Boot file and the registry | Registry and (for Active Directory–integrated zones) the directory | The directory (for Active Directory–integrated zones) and the registry |
| **Write server and zones parameters to:** | Boot file and the registry | Registry (for all zones) and (for Active Directory–integrated zones) the directory | The directory (for Active Directory–integrated zones) and the registry |

If you change the setting of the DNS Server service, it first writes the root hints file, zones, and parameters to the locations specified in the default setting, and then the DNS Server service reads them from the new setting.

**Directory-integrated Zone Storage Location**

Active Directory is an object-oriented X.500-compliant database that organizes resources available on your network in a hierarchical tree-like structure. This database is managed by a set of domain controllers. The portion of the Active Directory database for which a specific domain controller is authoritative is physically located on the same computer as the domain controller is. Every resource in Active Directory is represented by an object. There are two distinct types of objects supported by Active Directory:

- Containers–objects that can contain other container and leaf objects.
- Leafs–objects representing a specific resource within the Active Directory service tree.

Each Active Directory object has attributes associated with it that define particular characteristics of the object.

The classes of objects in the Active Directory database, as well as the attributes of each object, are defined in the Active Directory schema. In other words, the schema contains definitions for each class object available in Active Directory. The following are examples of Active Directory service class objects:

- User
- Group
- Organizational Unit
- DnsZone
- DnsNode

In Windows 2000 Server, Active Directory-integrated zones were stored in the domain partition of the directory. Zones stored in the domain partition are replicated to all domain controllers in the domain. This replication scope is not necessary for some

applications, such as DNS. Windows Server 2003 Active Directory provides a new type of partition, called application directory partition, to enable different replication scopes. Application directory partitions provide storage for application-specific data that can be replicated to any arbitrary set of domain controllers in the forest, as few or as many as needed by the application that uses the data.

**Note**

- At the time a DNS application directory partition is created (manually by an administrator or programmatically by an application), the domain naming flexible single master operations (FSMO) role for the forest must be held by a domain controller running Windows Server 2003. Following application directory partition creation, you can move the domain naming role to a domain controller that is running Windows 2000, if needed.

There are two default Windows Server 2003 DNS application directory partitions created to allow for different DNS zone replication scopes: the forest-wide DNS application directory partition and the domain-wide application directory partition. Active Directory-integrated zones can be stored in the domain or application directory partitions. The following table describes the Windows Server 2003 Active Directory storage options available for DNS zones.

**Windows Server 2003 Active Directory Storage Options**

| **Active Directory Storage Option** | **Description** |
|---|---|
| Domain partition | Active Directory domain partition for each domain in the forest. DNS zones stored in this partition are replicated to all domain controllers in the domain. This is the only Active Directory storage option for DNS zones that are replicated to domain controllers running Windows 2000 Server. |
| Forest-wide DNS application directory partition | DNS application directory partition for the entire forest. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controllers in the forest.<br><br>This DNS application directory partition is created when you install the DNS Server service on the first Windows Server 2003<br><br>domain controller in the forest. |
| Domain-wide DNS application directory partition | DNS application directory partition for each domain in the forest. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controllers in the domain. For the forest root domain, this DNS application directory partition is created when you first install the DNS Server service on a Windows Server 2003 domain controller in the forest.<br><br>For each new domain in the forest (child domain), this DNS application directory partition is created when you first install the DNS Server service on a Windows Server 2003 domain controller for the new domain. |
| Custom DNS application directory partition | DNS application directory partition for any domain controller that is enlisted in its replication scope. This type of DNS application directory partition does not exist by default and must be created. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controller that enlist in the partition. |

**Note**

- You can specify the DNS application directory partition where you want to store the zone by using the DNS console or the Dnscmd support tool, or you can change the zone storage option once the zone is created using DNS console or Dnscmd.

**Active Directory DNS Objects**

By default, Windows Server 2003 Active Directory-integrated zones are stored in the domain-wide application directory partition of the directory. The zone information is stored with the container whose distinguished name is: CN=MicrosoftDNS,DC=DNS domain name. The zone information for zones stored on domain controllers in the example.com domain would be stored in CN=MicrosoftDNS,DC=DomainDnsZones,DC=example,DC=com.

The DNS application directory partitions are not displayed by all Active Directory administrative tools. To see these directory partitions, you can use dnscmd (command-line tool) or ADSI Edit (adsiedit.msc) in Support Tools.

The replication scope of a custom DNS application directory partition is defined by the number of domain controllers enlisted in the partition. By default, only members of the Enterprise Admins group can enlist a DNS server in a DNS application directory partition.

Once a DNS server is enlisted in a DNS application directory partition, you can store DNS zones in that application directory partition using the DNS console. The required FQDN value specifies the name of the new DNS application directory partition. You must use a DNS fully qualified domain name.

The MicrosoftDNS container holds other objects that represent the zones and the individual DNS resource records in those zones. The following table lists the DNS object types that are stored in the Active Directory.

**DNS Objects Stored in Active Directory**

| Object | Description |
|---|---|
| DnsZone | Container created when a zone is stored in Active Directory. |
| DnsNode | Leaf object used to map and associate a name in the zone to resource data. |
| DnsRecord | Multivalued attribute of a dnsNode object used to store the resource records associated with the named node object. |
| DnsProperty | Multivalued attribute of a dnsZone object used to store zone configuration information. |

The MicrosoftDNS container object contains one or more dnsZone container objects. Each dnsZone object represents one zone.

MicrosoftDNS contains the following dnsZone objects:

- The reverse lookup zone, 72.16.172.in-addr.arpa
- The forward lookup zone, example.com
- The root hints, RootDNSServers

You can view DNS objects from within the Active Directory Users and Computers console. The following graphic shows the dnsZone objects in Active Directory.

**DNS Objects in Active Directory**




The dnsZone container object contains a dnsNode leaf object for every unique name within that zone. This figure shows the following dnsNode objects within the dnsZone container object for example.com:

- **@**, which signifies that the node has the same name as the dnsZone object.
- **delegated**, a delegated subdomain.
- **host.notdelegated**, a host in the domain not delegated.example.com, a domain that is controlled by the zone on example.com.
- **host1**, a host in the domain example.com.
- **mailserver**, the mail server in the domain example.com.
- **nameserver**, the name server in example.com.
- **notdelegated**, the domain notdelegated.example.com, which is controlled by the zone on example.com.

The dnsNode leaf object has a multivalued attribute called **dnsRecord** with an instance of a value for every record associated with the object's name. In this example, the dnsNode leaf object mailserver.example.com has an "A" attribute containing the IP address.

Although you can view the zone objects from within the Active Directory Users and Computers component, the Active Directory Users and Computers component cannot interpret the values of the **dnsRecord** attribute. If you want to view the DNS domain hierarchy and associated records, use the DNS console. Alternatively, if you want to view the zones, you can retrieve them by using Nslookup.

**Note**

- Windows 2000 domain controllers do not support application directory partitions and therefore you cannot view zones stored in DNS application directory partitions using Windows 2000 administration tools.

**DNS Processes and Interactions**

DNS processes and interactions involve the communications between DNS clients and DNS servers during the resolution of DNS queries and dynamic update, and between DNS servers during name resolution and zone administration. Secondary processes and interactions depend on the support for technologies such as Unicode and WINS.

For information about TCP/IP DNS messages, see "DNS Protocol" in this document.

**How DNS Queries Work**

When a DNS client needs to look up a name used in a program, it queries DNS servers to resolve the name. Each query message the client sends contains three pieces of information, specifying a question for the server to answer:

1. A specified DNS domain name, stated as a fully qualified domain name.
2. A specified query type, which can either specify a resource record by type or a specialized type of query operation.
3. A specified class for the DNS domain name. For Windows DNS servers, this should always be specified as the Internet (IN) class.

For example, the name specified could be the FQDN for a computer, such as ""host-a.example.microsoft.com."", and the query type specified to look for an address (A) resource record by that name. Think of a DNS query as a client asking a server a two-part question, such as "Do you have any A resource records for a computer named 'hostname.example.microsoft.com.' "? When the client receives an answer from the server, it reads and interprets the answered A resource record, learning the IP address for the computer it asked for by name.

DNS queries resolve in a number of different ways. A client can sometimes answer a query locally using cached information obtained from a previous query. The DNS server can use its own cache of resource record information to answer a query. A DNS server can also query or contact other DNS servers on behalf of the requesting client to fully resolve the name, then send an answer back to the client. This process is known as recursion.

In addition, the client itself can attempt to contact additional DNS servers to resolve a name. When a client does so, it uses separate and additional queries based on referral answers from servers. This process is known as iteration.

In general, the DNS query process occurs in two parts:

- A name query begins at a client computer and is passed to a resolver, the DNS Client service, for resolution.
- When the query cannot be resolved locally, DNS servers can be queried as needed to resolve the name.

Both of these processes are explained in more detail in the following sections.

**Part 1: DNS Client Service Resolver**
The following figure shows an overview of the complete DNS query process.

**Overview of DNS Query Process**




As shown in the initial steps of the query process, a DNS domain name is used in a program on the local computer. The request is then passed to the DNS Client service for resolution using locally cached information. If the queried name can be resolved, the query is answered and the process is completed.

The local resolver cache can include name information obtained from two possible sources:

- If a Hosts file is configured locally, any host name-to-address mappings from that file are preloaded into the cache when the DNS Client service is started.
- Resource records obtained in answered responses from previous DNS queries are added to the cache and kept for a period of time.

If the query does not match an entry in the cache, the resolution process continues with the client querying a DNS server to resolve the name.

**Part 2: Querying a DNS Server**
As indicated in the previous figure, the client queries a preferred DNS server. The actual server used during the initial client/server query is selected from a global list.

When the DNS server receives a query, it first checks to see if it can answer the query authoritatively based on resource record information contained in a locally configured zone on the server. If the queried name matches a corresponding resource record in local zone information, the server answers authoritatively, using this information to resolve the queried name.

If no zone information exists for the queried name, the server then checks to see if it can resolve the name using locally cached information from previous queries. If a match is found here, the server answers with this information. Again, if the preferred server can answer with a positive matched response from its cache to the requesting client, the query is completed.

If the queried name does not find a matched answer at its preferred server — either from its cache or zone information — the query process can continue, using recursion to fully resolve the name. This involves assistance from other DNS servers to help resolve the name. By default, the DNS Client service asks the server to use a process of recursion to fully resolve names on behalf of the client before returning an answer.

In order for the DNS server to do recursion properly, it first needs some helpful contact information about other DNS servers in the DNS domain namespace. This information is provided in the form of root hints, a list of preliminary resource records that can be used by the DNS service to locate other DNS servers that are authoritative for the root of the DNS domain

namespace tree. Root servers are authoritative for the domain root and top-level domains in the DNS domain namespace tree.

By using root hints to find root servers, a DNS server is able to complete the use of recursion. In theory, this process enables any DNS server to locate the servers that are authoritative for any other DNS domain name used at any level in the namespace tree.

For example, consider the use of the recursion process to locate the name “host-b.example.microsoft.com.” when the client queries a single DNS server. The process occurs when a DNS server and client are first started and have no locally cached information available to help resolve a name query. It assumes that the name queried by the client is for a domain name of which the server has no local knowledge, based on its configured zones.

First, the preferred server parses the full name and determines that it needs the location of the server that is authoritative for the top-level domain, “com”. It then uses an iterative query to the “com” DNS server to obtain a referral to the “microsoft.com” server. Next, a referral answer comes from the “microsoft.com” server to the DNS server for “example.microsoft.com”.

Finally, the “example.microsoft.com.” server is contacted. Because this server contains the queried name as part of its configured zones, it responds authoritatively back to the original server that initiated recursion. When the original server receives the response indicating that an authoritative answer was obtained to the requested query, it forwards this answer back to the requesting client and the recursive query process is completed.

Although the recursive query process can be resource-intensive when performed as described above, it has some performance advantages for the DNS server. For example, during the recursion process, the DNS server performing the recursive lookup obtains information about the DNS domain namespace. This information is cached by the server and can be used again to help speed the answering of subsequent queries that use or match it. Over time, this cached information can grow to occupy a significant portion of server memory resources, although it is cleared whenever the DNS service is cycled on and off.

The following three figures illustrate the process by which the DNS client queries the servers on each adapter.

**Querying the DNS Server, Part 1**




**Querying the DNS Server, Part 2**




**Querying the DNS Server, Part 3**




The DNS Client service queries the DNS servers in the following order:

1. The DNS Client service sends the name query to the first DNS server on the preferred adapter's list of DNS servers and waits one second for a response.

2. If the DNS Client service does not receive a response from the first DNS server within one second, it sends the name query to the first DNS servers on all adapters that are still under consideration and waits two seconds for a response.

3. If the DNS Client service does not receive a response from any DNS server within two seconds, the DNS Client service sends the query to all DNS servers on all adapters that are still under consideration and waits another two seconds for a response.

4. If the DNS Client service still does not receive a response from any DNS server, it sends the name query to all DNS servers on all adapters that are still under consideration and waits four seconds for a response.

5. If it the DNS Client service does not receive a response from any DNS server, the DNS client sends the query to all DNS servers on all adapters that are still under consideration and waits eight seconds for a response.

If the DNS Client service receives a positive response, it stops querying for the name, adds the response to the cache and returns the response to the client.

If the DNS Client service has not received a response from any server within eight seconds, the DNS Client service responds with a time-out. Also, if it has not received a response from any DNS server on a specified adapter, then for the next 30 seconds, the DNS Client service responds to all queries destined for servers on that adapter with a time-out and does not query those servers. Only computers running Windows 2000 or Windows Server 2003 return this time-out.

If at any point the DNS Client service receives a negative response from a server, it removes every server on that adapter from consideration during this search. For example, if in step 2, the first server on Alternate Adapter A gave a negative response, the DNS Client service would not send the query to any other server on the list for Alternate Adapter A.

The DNS Client service keeps track of which servers answer name queries more quickly, and it moves servers up or down on

the list based on how quickly they reply to name queries.

The following figure shows how the DNS client queries each server on each adapter.

**Multihomed Name Resolution**




**Alternate Query Responses**

The previous discussion of DNS queries assumes that the process ends with a positive response returned to the client. However, queries can return other answers as well. These are the most common query answers:

- An authoritative answer
- A positive answer
- A referral answer (used by the Windows Server 2003 DNS Server service only)
- A negative answer

An authoritative answer is a positive answer returned to the client and delivered with the authority bit set in the DNS message to indicate the answer was obtained from a server with direct authority for the queried name.

A positive response can consist of the queried RR or a list of RRs (also known as an RRset) that fits the queried DNS domain name and record type specified in the query message.

A referral answer contains additional RRs not specified by name or type in the query. This type of answer is returned to the client if the recursion process is not supported. The records are meant to act as helpful reference answers that the client can use to continue the query using iteration. A referral answer contains additional data such as RRs that are other than the type queried. For example, if the queried host name was "www" and no A RRs for this name were found in this zone but a CNAME RR for "www" was found instead, the DNS server can include that information when responding to the client. If the client is able to use iteration, it can make additional queries using the referral information in an attempt to fully resolve the name for itself.

A negative response from the server can indicate that one of two possible results was encountered while the server attempted to process and recursively resolve the query fully and authoritatively:

- An authoritative server reported that the queried name does not exist in the DNS namespace.
- An authoritative server reported that the queried name exists but no records of the specified type exist for that name.

The resolver passes the results of the query, in the form of either a positive or negative response, back to the requesting program and caches the response.

If the resultant answer to a query is too long to be sent and resolved in a single UDP message packet, the DNS server can initiate a failover response over TCP port 53 to answer the client fully in a TCP connected session.

Disabling the use of recursion on a DNS server is generally done when DNS clients are being limited to resolving names to a specific DNS server, such as one located on your intranet. Recursion might also be disabled when the DNS server is incapable

of resolving external DNS names, and clients are expected to fail over to another DNS server for resolution of these names. If you disable recursion on the DNS server, you will not be able to use forwarders on the same server.

By default, DNS servers use several default timings when performing a recursive query and contacting other DNS servers. These defaults include:

- A recursion retry interval of 3 seconds. This is the length of time the DNS service waits before retrying a query made during a recursive lookup.

- A recursion time-out interval of 15 seconds. This is the length of time the DNS service waits before failing a recursive lookup that has been retried.

Under most circumstances, these parameters do not need adjustment. However, if you are using recursive lookups over a slow-speed WAN link, you might be able to improve server performance and query completion by making slight adjustments to the settings.

**How Iteration Works**
Iteration is the type of name resolution used between DNS clients and servers when the following conditions are in effect:

- The client requests the use of recursion, but recursion is disabled on the DNS server.

- The client does not request the use of recursion when querying the DNS server.

An iterative request from a client tells the DNS server that the client expects the best answer the DNS server can provide immediately, without contacting other DNS servers.

When iteration is used, a DNS server answers a client based on its own specific knowledge about the namespace with regard to the names data being queried. For example, if a DNS server on your intranet receives a query from a local client for "www.microsoft.com", it might return an answer from its names cache. If the queried name is not currently stored in the names cache of the server, the server might respond by providing a referral — that is, a list of NS and A resource records for other DNS servers that are closer to the name queried by the client.

When iteration is used, a DNS server can further assist in a name query resolution beyond giving its own best answer back to the client. For most iterative queries, a client uses its locally configured list of DNS servers to contact other name servers throughout the DNS namespace if its primary DNS server cannot resolve the query.

The Windows Server 2003 DNS Client service does not perform recursion.

**How Caching Works**
As DNS servers process client queries using recursion or iteration, they discover and acquire a significant store of information about the DNS namespace. This information is then cached by the server.

Caching provides a way to speed the performance of DNS resolution for subsequent queries of popular names, while substantially reducing DNS-related query traffic on the network.

As DNS servers make recursive queries on behalf of clients, they temporarily cache resource records (RRs). Cached RRs contain information obtained from DNS servers that are authoritative for DNS domain names learned while making iterative queries to search and fully answer a recursive query performed on behalf of a client. Later, when other clients place new queries that request RR information matching cached RRs, the DNS server can use the cached RR information to answer them.

When information is cached, a Time-To-Live (TTL) value applies to all cached RRs. As long as the TTL for a cached RR does not expire, a DNS server can continue to cache and use the RR again when answering queries by its clients that match these RRs. Caching TTL values used by RRs in most zone configurations are assigned the Minimum (default) TTL which is set in the zone's start of authority (SOA) resource record. By default, the minimum TTL is 3,600 seconds (one hour) but can be adjusted or, if needed, individual caching TTLs can be set at each RR.

### Note

- By default, the DNS Server service uses a root hints file, cache.dns, that is stored in the systemroot\System32\Dns folder on the server computer. This file contains the NS and A resource records for the root servers of the DNS namespace (the Internet root servers or intranet root servers). When the DNS Server service is started, the root server list is queried for a current list of all the root servers. The results of the query are used to update the root hints file. This operation is also performed periodically while the service is running. When changes are made to the root hints by an administrator, these

changes are written back to the root hints file.

**Reverse Lookup**

In most DNS lookups, clients typically perform a forward lookup, which is a search based on the DNS name of another computer as stored in an address (A) resource record. This type of query expects an IP address as the resource data for the answered response.

DNS also provides a reverse lookup process, enabling clients to use a known IP address during a name query and to look up a computer name based on its address. A reverse lookup takes the form of a question, such as "Can you tell me the DNS name of the computer that uses the IP address 192.168.1.20?"

DNS was not originally designed to support this type of query. One problem for supporting the reverse query process is the difference in how the DNS namespace organizes and indexes names and how IP addresses are assigned. If the only method available to answer the previous question was to search all domains in the DNS namespace, a reverse query would take too long and require too much processing to be useful.

To solve this problem, a special domain called the in-addr.arpa domain was defined in the DNS standards and reserved in the Internet DNS namespace to provide a practical and reliable way to perform reverse queries. To create the reverse namespace, subdomains within the in-addr.arpa domain are formed using the reverse ordering of the numbers in the dotted-decimal notation of IP addresses.

This reversed ordering of the domains for each octet value is needed because, unlike DNS names, when IP addresses are read from left to right, they are interpreted in the opposite manner. When an IP address is read from left to right, it is viewed from its most generalized information (an IP network address) in the first part of the address to the more specific information (an IP host address) contained in the last octets.

For this reason, the order of IP address octets must be reversed when building the in-addr.arpa domain tree. The IP addresses of the DNS in-addr.arpa tree can be delegated to companies as they are assigned a specific or limited set of IP addresses within the Internet-defined address classes.

Finally, the in-addr.arpa domain tree, as built into DNS, requires that an additional RR type — the pointer (PTR) RR — be defined. This RR is used to create a mapping in the reverse lookup zone that typically corresponds to a host (A) named RR for the DNS computer name of a host in its forward lookup zone.

The in-addr.arpa domain applies for use in all TCP/IP networks that are based on Internet Protocol version 4 (IPv4) addressing. The New Zone Wizard automatically assumes that you are using this domain when creating a new reverse lookup zone.

If you are installing DNS and configuring reverse lookup zones for an Internet Protocol version 6 (IPv6) network, you can specify an exact name in the New Zone Wizard. This will permit you to create reverse lookup zones in the DNS console that can be used to support IPv6 networks, which use a different special domain name, the ip6.arpa domain.

For information about IPv6 and DNS, including examples of how to create and use ip6.arpa domain names as described in RFC 1886 ("DNS Extensions to support IP version 6"), see "Related Information" at the end of this section.

### Note

- The configuration of PTR resource records and reverse lookup zones for identifying hosts by reverse query is strictly an optional part of the DNS standard implementation. You are not required to use reverse lookup zones, although for some networked applications, they are used to perform security checks.

**Example: Reverse Query (for IPv4 networks)**

The following figure shows an example of a reverse query initiated by a DNS client (host-b) to learn the name of another host (host-a) based on its IP address, 192.168.1.20.

### Reverse Query




The reverse query process as shown in this figure occurs in the following steps:

1. The client, “host-b”, queries the DNS server for a pointer (PTR) RR that maps to the IP address of 192.168.1.20 for “host-a”.

   Because the query is for PTR records, the resolver reverses the address and appends the in-addr.arpa domain to the end of the reverse address. This forms the fully qualified domain name (“20.1.168.192.in-addr.arpa.”) for which to be searched in a reverse lookup zone.

2. Once located, the authoritative DNS server for “20.1.168.192.in-addr.arpa” can respond with the PTR record information. This includes the DNS domain name for “host-a”, completing the reverse lookup

Keep in mind that if the queried reverse name is not answerable from the DNS server, normal DNS resolution (either recursion or iteration) can be used to locate a DNS server that is authoritative for the reverse lookup zone and that contains the queried name. In this sense, the name resolution process used in a reverse lookup is identical to that of a forward lookup.

**Note**

- The DNS console provides a means for you to configure a subnetted reverse lookup “classless” zone when the **Advanced** view is selected. This allows you to configure a zone in the in-addr.arpa domain for a limited set of assigned IP addresses where a nondefault IP subnet mask is used with those addresses.

**Forwarding**

A forwarder is a Domain Name System (DNS) server on a network used to forward DNS queries for external DNS names to DNS servers outside of that network. You can also forward queries according to specific domain names using conditional forwarders.

A DNS server on a network is designated as a forwarder by having the other DNS servers in the network forward the queries they cannot resolve locally to that DNS server. By using a forwarder, you can manage name resolution for names outside of your network, such as names on the Internet, and improve the efficiency of name resolution for the computers in your network.

**Directing Name Queries Using Forwarders**

The following figure illustrates how external name queries are directed using forwarders.

**External Name Queries Directed Using Forwarders**




Without having a specific DNS server designated as a forwarder, all DNS servers can send queries outside of a network using their root hints. As a result, a lot of internal, and possibly critical, DNS information can be exposed on the Internet. In addition to this security and privacy issue, this method of resolution can result in a large volume of external traffic that is costly and inefficient for a network with a slow Internet connection or a company with high Internet service costs.

When you designate a DNS server as a forwarder, you make that forwarder responsible for handling external traffic, thereby limiting DNS server exposure to the Internet. A forwarder will build up a large cache of external DNS information because all of the external DNS queries in the network are resolved through it. In a small amount of time, a forwarder will resolve a good portion of external DNS queries using this cached data and thereby decrease the Internet traffic over the network and the response time for DNS clients.

**Behavior of a DNS Server Configured to Use Forwarding**

A DNS server configured to use a forwarder will behave differently than a DNS server that is not configured to use a forwarder. A DNS server configured to use a forwarder behaves as follows:

1. When the DNS server receives a query, it attempts to resolve this query using the primary and secondary zones that it hosts and its cache.

2. If the query cannot be resolved using this local data, then it will forward the query to the DNS server designated as a forwarder.

3. The DNS server will wait briefly for an answer from the forwarder before attempting to contact the DNS servers specified in its root hints.

When a DNS server forwards a query to a forwarder it sends a recursive query to the forwarder. This is different than the iterative query that a DNS server will send to an other DNS server during standard name resolution (name resolution that does not involve a forwarder).

**Forwarding Sequence**
The sequence in which the forwarders configured on a DNS server are used is determined by the order in which the IP addresses are listed on the DNS server. After the DNS server forwards the query to the forwarder with the first IP address, it waits a short period for an answer from that forwarder (according to the DNS server's time out setting) before resuming the forwarding operation with the next IP address. It continues this process until it receives an affirmative answer from a forwarder.

Unlike conventional resolution, where a roundtrip time (RTT) is associated with each server, the IP addresses in the forwarders list are not ordered according to roundtrip time and must be reordered manually to change preference.

**Forwarders and Delegation**
A DNS server configured with a forwarder and hosting a parent zone will use its delegation information before forwarding queries. If no delegation record exists for the DNS name in the query, then the DNS server will use its forwarders to resolve the query.

**Forwarders and Root Servers**
A common error when configuring forwarding is to attempt to configure forwarding on the root servers of a private DNS namespace. The goal of attempting to configure forwarding on root servers for a private DNS namespace is to forward all offsite queries to Internet DNS servers. Root servers cannot be configured with standard forwarding. If a root server is queried about any domain name, then it will refer to a DNS server that can answer the question (from its local zones, cache), or it will respond with a failure (NXDOMAIN), but it cannot be configured to forward to specific servers.

### Note

- A root server can be configured with a conditional forwarder. Conditional forwarding can be used to forward queries between root servers in separate DNS namespaces, although the DNS servers for the top-level domains in the namespace are better suited for this method of resolution.

**Conditional Forwarders**
A conditional forwarder is a DNS server on a network that is used to forward DNS queries according to the DNS domain name in the query. For example, a DNS server can be configured to forward all the queries it receives for names ending with widgets.example.com to the IP address of a specific DNS server or to the IP addresses of multiple DNS servers.

**Intranet Name Resolution**
A conditional forwarder can be used to improve name resolution for domains within your intranet. Intranet name resolution can be improved by configuring DNS servers with forwarders for specific internal domain names. For example, all DNS servers in the domain widgets.example.com could be configured to forward queries for names that end with test.example.com to the authoritative DNS servers for merged.widgets.example.com, thereby removing the step of querying the root servers of example.com, or removing the step of configuring DNS servers in the widgets.example.com zone with secondary zones for test.example.com.

**Internet Name Resolution**
DNS servers can use conditional forwarders to resolve queries between the DNS domain names of companies that share information. For example, two companies, Widgets Toys and TailspinToys, want to improve how the DNS clients of Widgets Toys resolve the names of the DNS clients of Tailspin Toys. The administrators from Tailspin Toys inform the administrators of Widgets Toys about the set of DNS servers in the Tailspin Toys network where Widgets can send queries for the domain dolls.tailspintoys.com. The DNS servers within the Widgets Toys network are configured to forward all queries for names ending with dolls.tailspintoys.com to the designated DNS servers in the network for Tailspin Toys. Consequently, the DNS servers in the Widgets Toys network do not need to query their internal root servers, or the Internet root servers, to resolve queries for names ending with dolls.tailspintoys.com.

**Dynamic Update**

Dynamic update enables DNS client computers to register and dynamically update their resource records with a DNS server whenever changes occur. This reduces the need for manual administration of zone records, especially for clients that frequently move or change locations and use DHCP to obtain an IP address.

The DNS Client and Server services support the use of dynamic updates, as described in RFC 2136, "Dynamic Updates in the Domain Name System." The DNS Server service allows dynamic update to be enabled or disabled on a per-zone basis at each server configured to load either a standard primary or directory-integrated zone. By default, the Windows Server 2003 DNS Client service will dynamically update host (A) resource records in DNS when configured for TCP/IP. The Windows Server 2003 DNS Server service is configured, by default, to allow only secure dynamic update. You must change this configuration if you will be using dynamic update only.

**Protocol Description**
RFC 2136 introduces a new opcode or message format called UPDATE. The update message can add and delete resource records from a specified zone as well as test for prerequisite conditions. Update is atomic, that is, all prerequisites must be satisfied or no update operation will take place.

As in any conventional DNS implementation, the zone update must be committed on a primary DNS server for that zone. If an update is received by a secondary DNS server, it will be forwarded up the replication topology until it reaches the primary DNS server. Note that in the case of an Active Directory integrated zone, an update for a resource record in a zone may be sent to any DNS server running on an Active Directory domain controller whose data store contains the zone.

A zone transfer process will always lock a zone so that a secondary DNS server receives a consistent zone view while transferring the zone data. When the zone is locked it can no longer accept dynamic updates. If the zone is large and is locked very often for zone transfer purposes, it will starve dynamic update clients, and the system can become unstable. The Windows Server 2003 DNS Server service queues the update requests that arrived during the zone transfer and processes them after the zone transfer is completed.

**How client and server computers update their DNS names**
By default, computers that are statically configured for TCP/IP attempt to dynamically register host (A) and pointer (PTR) resource records for IP addresses configured and used by their installed network connections. All computers register records based on their fully qualified domain name (FQDN).

The following defaults also apply to how computers update their DNS names:

- By default, the DNS client on Windows XP does not attempt dynamic update over a Remote Access Service or virtual private network connection. To modify this configuration, you can modify the advanced TCP/IP settings of the particular network connection or modify the registry.

- By default, the DNS client does not attempt dynamic update of top-level domain (TLD) zones. Any zone named with a single-label name is considered a TLD zone, for example, com, edu, blank, my-company.

- By default, the primary DNS suffix portion of a computer's FQDN is the same as the name of the Active Directory domain to which the computer is joined. To allow different primary DNS suffixes, a domain administrator may create a restricted list of allowed suffixes by modifying the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is managed by the domain administrator using Active Directory Service Interfaces (ADSI) or the Lightweight Directory Access Protocol (LDAP).

Dynamic updates can be sent for any of the following reasons or events:

- An IP address is added, removed, or modified in the TCP/IP properties configuration for any one of the installed network connections.

- An IP address lease changes or renews with the DHCP server any one of the installed network connections. For example, when the computer is started or if the ipconfig /renew command is used.

- The ipconfig /registerdns command is used to manually force a refresh of the client name registration in DNS.

- At startup time, when the computer is turned on.

- A member server is promoted to a domain controller.

When one of the previous events triggers a dynamic update, the DHCP Client service (not the DNS Client service) sends updates. This is designed so that if a change to the IP address information occurs because of DHCP, corresponding updates in DNS are performed to synchronize name-to-address mappings for the computer. The DHCP Client service performs this function for all network connections used on the system, including connections not configured to use DHCP.

The update process described above assumes that installation defaults are in effect for computers running Windows 2000, Windows XP, or servers running Windows Server 2003. Specific names and update behavior is tunable where advanced TCP/IP properties are configured to use non-default DNS settings.

In addition to the full computer name (or primary name) of the computer, additional connection-specific DNS names can be configured and optionally registered or updated in DNS.

**Example: How dynamic update works**
Dynamic updates are typically requested when either a DNS name or IP address changes on the computer. For example, suppose a client named "oldhost" is first configured in System properties with the following names:

| | |
|---|---|
| Computer name | oldhost |
| DNS domain name of computer | example.microsoft.com |
| Full computer name | oldhost.example.microsoft.com |

In this example, no connection-specific DNS domain names are configured for the computer. Later, the computer is renamed from "oldhost" to "newhost", resulting in the following name changes on the system:

| | |
|---|---|
| Computer name | newhost |
| DNS domain name of computer | example.microsoft.com |
| Full computer name | newhost.example.microsoft.com |

Once the name change is applied in **System properties**, you are prompted to restart the computer. When the computer restarts Windows, the DHCP Client service performs the following sequence to update DNS:

1. The DHCP Client service sends a start of authority (SOA) type query using the DNS domain name of the computer.

   The client computer uses the currently configured FQDN of the computer (such as "newhost.example.microsoft.com") as the name specified in this query.

2. The authoritative DNS server for the zone containing the client FQDN responds to the SOA-type query.

   For standard primary zones, the primary server (owner) returned in the SOA query response is fixed and static. It always matches the exact DNS name as it appears in the SOA RR stored with the zone. If, however, the zone being updated is directory-integrated, any DNS server that is running on a domain controller for the Active Directory domain in the FQDN can respond and dynamically insert its own name as the primary server (owner) of the zone in the SOA query response.

3. The DHCP Client service then attempts to contact the primary DNS server.

   The client processes the SOA query response for its name to determine the IP address of the DNS server authorized as the primary server for accepting its name. It then proceeds to perform the following sequence of steps as needed to contact and dynamically update its primary server:

   - It sends a dynamic update request to the primary server determined in the SOA query response.
   - If the update succeeds, no further action is taken.
   - If this update fails, the client next sends an NS-type query for the zone name specified in the SOA record.
   - When it receives a response to this query, it sends an SOA query to the first DNS server listed in the response.

- After the SOA query is resolved, the client sends a dynamic update to the server specified in the returned SOA record.
- If the update succeeds, no further action is taken.
- If this update fails, then the client repeats the SOA query process by sending to the next DNS server listed in the response.

4. Once the primary DNS server that can perform the update is contacted, the client sends the update request and the DNS server processes it.

   The contents of the update request include instructions to add A (and possibly PTR) resource records for "newhost.example.microsoft.com" and remove these same record types for "oldhost.example.microsoft.com", the name that was previously registered.

   The DNS server also checks to ensure that updates are permitted for the client request. For standard primary zones, dynamic updates are not secured, so any client attempt to update succeeds. For Active Directory–integrated zones, updates are secured and performed using directory-based security settings. For more information, see "Secure Dynamic Updates" later in this document.

Dynamic updates are sent or refreshed periodically. By default, computers send a refresh once every seven days. If the update results in no changes to zone data, the zone remains at its current version and no changes are written. Updates result in actual zone changes or increased zone transfer only if names or addresses actually change.

Note that names are not removed from DNS zones if they become inactive or are not updated within the refresh interval (seven days). DNS does not use a mechanism to release or tombstone names, although DNS clients do attempt to delete or update old name records when a new name or address change is applied.

When the DHCP Client service registers A and PTR resource records for a computer, it uses a default caching Time to Live (TTL) of 15 minutes for host records. This determines how long other DNS servers and clients cache a computer's records when they are included in a query response.

**DNS and DHCP Clients and Servers**
Windows 2000, Windows XP, and Windows Server 2003 DHCP clients are dynamic update-aware and can initiate the dynamic update process. A DHCP client negotiates the process of dynamic update with the DHCP server when the client leases an IP address or renews the lease, determining which computer updates the A and PTR resource records of the client. Depending on the negotiation process, the DHCP client, the DHCP server, or both, update the records by sending the dynamic update requests to the primary DNS servers that are authoritative for the names that are to be updated.

Clients and servers that are running versions of Windows earlier than Windows 2000 do not support dynamic update. The Windows 2000 and Windows Server 2003 DHCP Server service can perform dynamic updates on behalf of clients that do not support the DHCP Client service FQDN option (which is described in the following section). For example, clients that are running Microsoft Windows 95, Windows 98, and Windows NT do not support the FQDN option. However, this functionality can be enabled in the **DNS** tab of the server properties for the DHCP console. The DHCP server first obtains the name of legacy clients from the DHCP REQUEST packet. It then appends the domain name given for that scope and registers the A and PTR resource records.

In some cases, stale PTR or A resource records can appear on DNS servers when the lease of a DHCP client expires. For example, when a Windows 2000, Windows XP or Windows Server 2003 DHCP client tries to negotiate a dynamic update procedure with a Windows NT 4.0 DHCP server, the DHCP client must register both A and PTR resource records itself. Later, if the Windows 2000, Windows XP or Windows Server 2003 DHCP client is improperly removed from the network, the client cannot deregister its A and PTR resource records and they become stale.

If a stale A resource record appears in a zone that allows only secure dynamic updates, no computer is able to register any other resource record for the name in that A resource record. To prevent problems with stale PTR and A resource records, you can enable the aging and scavenging feature. For more information about the aging and scavenging feature, see "Understanding Aging and Scavenging" in this document.

To provide fault tolerance for dynamic updates, consider Active Directory integration for those zones that accept dynamic updates from Windows Server 2003 network-based clients. To speed up the discovery of authoritative DNS servers, you can configure each client with a list of preferred and alternate DNS servers that are primary for that directory-integrated zone. If a

client fails to update the zone with its preferred DNS server because the DNS server is unavailable, the client can try an alternate server. When the preferred DNS server becomes available, it loads the updated, directory-integrated zone that includes the update from the client.

**Dynamic Update Process for Network Connections Configured by DHCP**

To negotiate the dynamic update process, the DHCP client sends its fully qualified domain name (FQDN) to the DHCP server in the DHCPREQUEST packet by using the DHCP Client service FQDN option. The DHCP server then replies to the DHCP client by sending a DHCP acknowledgment (DHCPACK) message that includes the FQDN option (option code 81).

The table below lists the fields of the FQDN option of the DHCPREQUEST packet.

**Fields in the FQDN Option of the DHCPREQUEST Packet**

| Field | Explanation |
|---|---|
| Code | Specifies the code for this option (81). |
| Len | Specifies the length, in octets, of this option (minimum of 4). |
| Flags | Can be one of the following values:<br>0. Client wants to register the A resource record and requests that the server update the PTR resource record.<br>1. Client wants server to register the A and PTR resource records.<br>3. DHCP server registers the A and PTR resource records regardless of the request of the client. |
| RCODE1 and RCODE 2 | The DHCP server uses these fields to specify the response code from the A and PTR resource records registrations performed on the client's behalf and to indicate whether it attempted the update before sending DHCPACK. |
| Domain Name | Specifies the FQDN of the client. |

The conditions under which DHCP clients send the FQDN option and the actions taken by DHCP servers depend on the operating system that the client and server are running and how the client and server are configured.

The client requests a dynamic update depending on whether or not it is running Windows Server 2003 operating system, Windows 2000, or earlier. It also depends on the client configuration. Clients can take any of the following actions:

- By default, the Windows 2000, Windows XP and Windows Server 2003 DHCP Client service sends the FQDN option with the Flags field set to 0 to request that the client update the A resource record, and the DHCP Server service updates the PTR resource record. After the client sends the FQDN option, it waits for a response from the DHCP server. Unless the DHCP server sets the Flags field to 3, the client then initiates an update for the A resource record. If the DHCP server does not support or is not configured to perform registration of the DNS record, then no FQDN is included in the DHCP servers response and the client attempts registration of the A and PTR resource records.

- If the DHCP client is running a Windows operating system earlier than Windows 2000, or if the client is Windows 2000 and it is configured not to register DNS resource records, then the client does not send the FQDN option. In this case, the client does not update either record.

Depending on what the DHCP client requests, the DHCP server can take different actions. If the DHCP client sends a DHCPREQUEST message without the FQDN option, behavior depends on the type of DHCP server and how it is configured. The DHCP server can update both records if it is configured to update records on behalf of DHCP clients that do not support the FQDN option.

In the following cases, the DHCP server does not perform any action:

- The DHCP server does not support dynamic update (for example, a Windows NT 4.0 server).

- The DHCP server is running Windows 2000 or Windows Server 2003 operating system and is configured not to do dynamic updates for clients that do not support the FQDN option.

- The DHCP server is running Windows 2000 or Windows Server 2003 operating system and configured not to register

DNS resource records.

If the Windows 2000, Windows XP or Windows Server 2003 network–based DHCP client requests that the server updates the PTR resource record but not the A resource record, behavior depends on the type of DHCP server and how it is configured. The server can perform any of the following actions:

- If the DHCP server is running Windows 2000 or Windows Server 2003 operating system and is configured not to perform dynamic updates, its response does not contain the FQDN option and does not update either resource record. In this case, the DHCP client attempts to update both the A and PTR resource records, if it capable.

- If the DHCP server is running Windows 2000 or Windows Server 2003 operating system and is configured to update according to the request of the DHCP client, the server attempts to update the PTR resource record. The DHCP server DHCPACK message to the DHCP client contains the FQDN option with the Flags set to 0, confirming that the DHCP server updates the PTR record. The DHCP client then attempts to update the A resource record, if it is capable.

If the DHCP server is running Windows 2000 or Windows Server 2003 operating system and is configured to always update A and PTR both records, the DHCP server attempts to update both resource records. The DHCP server DHCPACK message to the DHCP client contains the FQDN option with the Flags set to 3, notifying the DHCP client that the DHCP server updates A and PTR records. In this case, the DHCP client does not attempt to update either resource record.

**Dynamic Update Process for Statically Configured and Remote Access Clients**
Statically configured clients and remote access clients do not rely on the DHCP server for DNS registration. Statically configured clients dynamically update their A and PTR resource records every time they start and then every 24 hours in case the records become corrupted or need to be refreshed in the DNS database.

Remote access clients can dynamically update A and PTR resource records when a dial-up connection is made. They can also attempt to withdraw, or deregister, the A and PTR resource records when the user closes down the connection explicitly. Computers running Windows 2000 or Windows Server 2003 operating system with a remote access network connection attempt the dynamic registration of the A and PTR records corresponding to the IP address of this connection. By default, the DNS Client service on Windows XP does not attempt dynamic update over a Remote Access Service (RAS) or virtual private network (VPN) connection. To modify this configuration, you can modify the advanced TCP/IP settings of the particular network connection or modify the registry.

In all operating systems, if a remote access client does not receive a successful response from the attempt to deregister a DNS resource record, or if for any other reasons fails to deregister a resource record within four seconds, the DNS client closes the connection. In such cases, the DNS database might contain a stale record.

If the remote access client fails to deregister a DNS resource record, it adds a message to the event log, which you can view by using Event Viewer. The remote access client never deletes stale records, but the remote access server attempts to deregister the PTR resource record when the client is disconnected.

Windows 2000 dial-up networking clients attempt to register A and PTR records for the dial-up connection. By default, the Windows XP and Windows Server 2003 DNS Client service dial-up networking clients do not attempt to update A and PTR records automatically. Due to the nature of their business, it is common that ISPs do not enable dynamic updating of DNS information by their customers. If you use an ISP that does not support dynamic update, configure the connection properties to prevent the computer from performing dynamic updates.

**Dynamic Update Process for Multihomed Clients**
If a dynamic update client is multihomed (has more than one network connection and associated IP address), it registers the first IP address for each network connection by default. If you do not want it to register these IP addresses, you can configure the network connection to not register IP addresses.

The dynamic update client does not register all IP addresses with the DNS servers in all namespaces that the computer is connected to. For example, a multihomed computer, client1.noam.example.com, is connected to both the Internet and the corporate intranet. Client1 is connected to the intranet by adapter A, a DHCP adapter with the IP address 172.16.8.7. Client1 is also connected to the Internet by adapter B, a remote access adapter with the IP address 10.3.3.9. Client1 resolves intranet names by using a name server on the intranet, NoamDC1, and resolves Internet names by using a name server on the Internet, ISPNameServer.

**Time to Live**
Whenever a dynamic update client registers in DNS, the associated A and PTR resource records include the Time to Live (TTL), which by default is set to 10 minutes for records registered by the Net Logon service, and 15 minutes for records registered by the DHCP Client service. If the DNS Server service dynamically registers records for its own zones, the default TTL is 20

minutes. You can change the default setting in the registry. A small value causes cached entries to expire sooner, which increases DNS traffic but decreases the risk of cached records becoming outdated. Expiring entries quickly is useful for computers that frequently renew their DHCP leases. Long retention times are useful for computers that renew their DHCP leases infrequently.

**Resolving Name Conflicts**
When the DNS Client service attempts to register an A record and it discovers that the authoritative DNS zone already contains an A record for the same name but with a different IP address, by default, the DNS Client service attempts to replace the existing A record (s) with the new A record containing the IP address of the DNS client. As a result, any computer on the network can modify the existing A record unless secure dynamic update is used. Zones that are configured for secure dynamic update allow only authorized users to modify the resource record.

You can change the default setting so that the DNS Client service aborts the registration process and logs the error in Event Viewer, instead of replacing the existing A record.

**Secure Dynamic Update**
DNS update security is available only for zones that are integrated into Active Directory. Once you directory-integrate a zone, access control list (ACL) editing features are available in the DNS console so you can add or remove users or groups from the ACL for a specified zone or resource record. ACLs are for DNS administration access control only, and do not influence DNS query resolution.

By default, dynamic update security for DNS servers and clients are handled as follows:

- DNS clients attempt to use unsecured dynamic update first. If an unsecured update is refused, clients try to use secure update.

  Also, clients use a default update policy that permits them to attempt to overwrite a previously registered resource record, unless they are specifically blocked by update security.

- Once a zone becomes Active Directory–integrated, DNS servers running Windows Server 2003 default to allowing only secure dynamic updates.

  When using standard zone storage, the default for the DNS Server service is to not allow dynamic updates on its zones. For zones that are either directory-integrated or use standard file-based storage, you can change the zone to allow all dynamic updates which permits all updates to be accepted.

  Dynamic update is a recent additional DNS standard specification, defined in RFC 2136. For more information about RFCs, see "Related Information" at the end of this document.

  The dynamic registration of DNS resource records can be restricted with the use of registry entries.

**How secure dynamic update works**
The secure dynamic update process is described below:

- To initiate a secure dynamic update, the DNS client first initiates the security context negotiation process, during which the tokens are passed between client and server using TKEY resource records. At the end of the negotiation process the security context is established.

- Next, the DNS client sends the dynamic update request (containing resource records for the purpose of adding, deleting, or modifying data) to the DNS server, signed using the previously established security context and passing the signature in the TSIG resource record, included in the dynamic update packet.

- The server attempts to update Active Directory using the clients credentials and sends the result of the update to the client. These results are signed using the security context and pass the signature in the TSIG resource record included in the response.

**Secure dynamic update process**
The secure dynamic update process is described below:

1. The DNS client queries the preferred DNS server to determine which DNS server is authoritative for the domain name it is attempting to update. The preferred DNS server responds with the name of the zone and the primary DNS server that

is authoritative for the zone.

2. The DNS client attempts a standard dynamic update, and if the zone is configured to allow only secure dynamic updates (the default configuration for Active Directory-integrated zones), the DNS server refuses the non-secure update. Had the zone been configured for standard dynamic update rather than secure dynamic update, the DNS server would have accepted the DNS client's attempt to add, delete, or modify resource records in that zone.

3. The DNS client and DNS server begin TKEY negotiation.

4. First, the DNS client and DNS server negotiate an underlying security mechanism. Windows dynamic update clients and DNS servers can only use the Kerberos protocol.

5. Next, by using the security mechanism, the DNS client and DNS server verify their respective identities and establish the security context.

6. The DNS client sends the dynamic update request to the DNS server, signed using the established security context . The signature is included in the signature field of the TSIG resource record that is included in the dynamic update request packet. The DNS server verifies the origin of the dynamic update packet by using the security context and the TSIG signature.

7. The DNS server attempts to add, delete, or modify resource records in Active Directory. Whether or not it can make the update depends on whether the DNS client has the proper permissions to make the update and whether the prerequisites have been satisfied.

8. The DNS server sends a reply to the DNS client stating whether or not it was able to make the update, signed using the established security context. The signature is included in the signature field of the TSIG resource record that is included in the dynamic update response packet. If the DNS client receives a spoofed reply, it ignores it and waits for a signed response.

   **Note**

   - DHCP clients running Windows 2000, Windows XP, or Windows Server 2003 explicitly request that the DHCP server update only pointer (PTR) resource records used in DNS for the reverse lookup and resolution of the client's IP address to its name. These clients update their address (A) resource records for themselves.

   - Clients running earlier versions of Windows cannot make an explicit request for DNS dynamic update protocol preference. For these clients, the DHCP service updates both the PTR and the A resource records.

**Security for DHCP Clients That Do Not Support the FQDN Option**

Windows DHCP clients that do not support the FQDN option (option 81) are not capable of dynamic updates. If you want the A and PTR resource records for these clients dynamically registered in DNS, you must configure the DHCP server to perform dynamic updates on their behalf.

However, having the DHCP server to perform secure dynamic updates on behalf of DHCP clients that do not support the FQDN option is undesirable because when a DHCP server performs a secure dynamic update on a name, that DHCP server becomes the owner of that name, and only that DHCP server can update any record for that name. This can cause problems in a few different circumstances.

For example, suppose that the DHCP server DHCP1 created an object for the name nt4host1.example.com and then stopped responding, and that later the backup DHCP server, DHCP2, tried to update a record for the same name, nt4host1.example.com. In this situation, DHCP2 is not able to update the name because it does not own the name. In another example, suppose DHCP1 added an object for the name nt4host1.example.com, and then the administrator upgraded nt4host1.example.com to a Windows 2000-based computer. Because the Windows 2000-based computer did not own the name, it would not be able to update DNS records for the name.

To solve this problem, the built-in security group called DnsUpdateProxy is provided. If all DHCP servers are added as members of the DnsUpdateProxy group, one servers records can be updated by another server if the first server fails. Also, because all objects created by the members of the DnsUpdateProxy group are not secured, the first user (that is not a member

of the DnsUpdateProxy group) to modify the set of records associated with a DNS name becomes its owner. When legacy clients are upgraded, they can therefore take ownership of their name records at the DNS server. If every DHCP server registering resource records for older clients is a member of the DnsUpdateProxy group, the problems discussed earlier are eliminated.

**Securing records when using the DnsUpdateProxy group**
DNS domain names that are registered by the DHCP server are not secure when the DHCP server is a member of the DnsUpdateProxy group. As a result, do not use this group in an Active Directory integrated-zone that allows only secure dynamic updates without taking additional steps to allow records created by members of the group to be secured.

To protect against unsecured records, or to allow members of the DnsUpdateProxy group to register records in zones that allow only secured dynamic updates, Windows Server 2003 DHCP and DNS allow you to create a dedicated user account and configure DHCP servers to perform DNS dynamic updates with the user account credentials (user name, password, and domain). The credentials of one dedicated user account can be used by multiple DHCP servers.

The dedicated user account is a standard user account used only is supplying DHCP servers with credentials for DNS dynamic update registrations. Each DHCP server supplies these credentials when registering names on behalf of DHCP clients using DNS dynamic update. The dedicated user account is created in the same forest where the primary DNS server for the zone to be updated resides. The dedicated user account can also be located in another forest as long as the forest it resides in has a forest trust established with the forest containing the primary DNS server for the zone to be updated.

When installed on a domain controller, the DHCP Server service inherits the security permissions of the domain controller and has the authority to update or delete any DNS record that is registered in a secure Active Directory-integrated zone (this includes records that were securely registered by other computers running Windows 2000 or Windows Server 2003, including domain controllers). When installed on a domain controller, configure the DHCP server with the credentials of the dedicated user account to prevent the server from inheriting, and possibly misusing, the power of the domain controller.

Configure a dedicated user account and configure the DHCP Server service with the account credentials under the following circumstances:

- A domain controller is configured to function as a DHCP server.
- The DHCP server is configured to perform DNS dynamic updates on behalf of DHCP clients.
- The DNS zones to be updated by the DHCP server are configured to allow only secure dynamic updates.

Once you have created a dedicated user account, you can configure DHCP servers with the user account credentials by using the DHCP console or by using the Netsh command (**netsh dhcp server set dnscredentials**).

**Note**

- If the supplied credentials belong to an object (such as a computer) that is a member of the DnsUpdateProxy security group, the next object to register the same name record in DNS will become the record owner.
- If you have specified credentials (user name, domain, and password) that the DHCP server uses when registering DHCP client computers in DNS, these credentials are not backed up with either synchronous or asynchronous backup. After a DHCP database is restored, new credentials must be configured.

**Controlling Update Access to Zones and Names**
Access to the DNS zones and resource records stored in Active Directory is controlled access control lists (ACLs). ACLs can be specified for the DNS Server service, an entire zone or for specific DNS names. By default, any authenticated Active Directory user can create the A or PTR resource records in any zone. Once an owner name has been created for a zone (regardless of the type of resource record), only the users or groups specified in the ACL for that name that have write permission are enabled to modify records corresponding to that name. While this approach is desirable in most scenarios, some special situations need to be considered separately.

**DNSAdmins group**
By default, the DNSAdmins group has full control of all zones and records in the Windows Server 2003 domain in which it is specified. In order for a user to be able to enumerate zones in a specific Windows Server 2003 domain, the user (or a group the user belongs to) must be enlisted in the DNSAdmin group.

It is possible that a domain administrator may not want to grant full control to all users listed in the DNSAdmins group.

Typically, this would be the result if a domain administrator wanted to grant full control for a specific zone and read only control for other zones in the domain to a set of users. To accomplish this, the domain administrator can create a separate group for each of the zones, and add specific users to each group. Then the ACL for each zone will contain a group with full control for that zone only. At the same time, all of the groups will be included in the DNSAdmins group, which can be configured to have read permissions only. As a result of the fact that a zone's ACL always contains the DNSAdmins group, all users enlisted in the zone-specific groups will have read permission for all the zones in the domain.

**Reserving names**
The default DNS Server service configuration of allowing any authenticated user to create a new name in a zone may not be sufficient for environments that require a high level of security. In such cases, the default ACL can be changed to allow for the creation of objects in a zone by certain groups or users only. Per-name administration of ACLs provides another solution to this problem. An administrator may reserve a name in a zone leaving the rest of the zone open for the creation of any new objects by all authenticated users. To accomplish this, an administrator creates a record for the reserved name and sets the appropriate list of groups or users in the ACL. As a result, only the users listed in the ACL will be able to register another record under the reserved name.

**Understanding Aging and Scavenging**
DNS servers running Windows Server 2003 support aging and scavenging features. These features are provided as a mechanism for performing cleanup and removal of stale resource records (RRs), which can accumulate in zone data over time.

With dynamic update, RRs are automatically added to zones when computers start on the network. However, in some cases, they are not automatically removed when computers leave the network. For example, if a computer registers its own host (A) RR at startup and is later improperly disconnected from the network, its host (A) RR might not be deleted. If your network has mobile users and computers, this situation can occur frequently.

If left unmanaged, the presence of stale RRs in zone data might cause some problems. The following are examples:

- If a large number of stale RRs remain in server zones, they can eventually take up server disk space and cause unnecessarily long zone transfers.
- DNS servers loading zones with stale RRs might use outdated information to answer client queries, potentially causing the client's to experience name resolution problems on the network.
- The accumulation of stale RRs at the DNS server can degrade its performance and responsiveness.
- In some cases, the presence of a stale RR in a zone could prevent a DNS domain name from being used by another computer or host device.

To solve these problems, the DNS Server service has the following features:

- Time stamping, based on the current date and time set at the server computer, for any RRs added dynamically to primary-type zones. In addition, time stamps are recorded in standard primary zones where aging/scavenging is enabled.
- For RRs that you add manually, a time stamp value of zero is used, indicating that they are not affected by the aging process and can remain without limitation in zone data unless you otherwise change their time stamp or delete them.
- Aging of RRs in local data, based on a specified refresh time period, for any eligible zones. Only primary type zones that are loaded by the DNS Server service are eligible to participate in this process.
- Scavenging for any RRs that persist beyond the specified refresh period. When a DNS server performs a scavenging operation, it can determine that RRs have aged to the point of becoming stale and remove them from zone data. Servers can be configured to perform recurring scavenging operations automatically, or you can initiate an immediate scavenging operation at the server.

  **Note**

  - By default, the aging and scavenging mechanism for the DNS Server service is disabled. It should only be enabled when all parameters are fully understood. Otherwise, the server could be accidentally configured to delete records that should not be deleted. If a record is accidentally deleted, not only will users fail to resolve queries for that record, but any user can create the record and take ownership of it, even on zones configured for secure

dynamic update.

The server uses the contents of each RR-specific time stamp, along with other aging/scavenging properties that you can adjust or configure, to determine when it scavenges records.

**Prerequisites for Aging and Scavenging**
Before the aging and scavenging features of DNS can be used, several conditions must be met:

1. Scavenging and aging must be enabled both at the DNS server and on the zone.

   By default, aging and scavenging of resource records is disabled.

2. Resource records must either be dynamically added to zones or manually modified to be used in aging and scavenging operations.

Typically, only those resource records added dynamically using the DNS dynamic update protocol are subject to aging and scavenging.

You can, however, enable scavenging for other resource records added through non-dynamic means. For records added to zones in this way, either by loading a text-based zone file from another DNS server or by manually adding them to a zone, a time stamp of zero is set. This makes these records ineligible for use in aging/scavenging operations.

In order to change this default, you can administer these records individually, to reset and permit them to use a current (non-zero) time stamp value. This enables these records to become aged and scavenged.

**Note**

- In the case of changing a zone from standard primary to Active Directory–integrated, you might want to enable scavenging of all existing resource records in the zone. To enable aging for all existing resource records in a zone, you can use the AgeAllRecords command, which is available through the dnscmd command-line tool.

**Aging and Scavenging Terminology**
The following list indicates new or revised terms that have been introduced to help specifically when discussing aging and scavenging.

**Current server time** The current date and time on the DNS server. This number can be expressed as an exact numeric value at any point in time.

**No-refresh interval** An interval of time, determined for each zone, as bounded by the following two events:

- The date and time when the record was last refreshed and its time stamp set.

- The date and time when the record next becomes eligible to be refreshed and have its time stamp reset.

This value is needed to decrease the number of write operations to the Active Directory database. By default, this interval is set to seven days. It should not be increased to an unreasonably high level, because the benefits of the aging and scavenging feature might either be lost or diminished.

**Record refresh** When a DNS dynamic update is processed for a resource record when only the resource record time stamp, and no other characteristics of the record, are revised. Refreshes generally occur for the following reasons:

- When a computer is restarted on the network and, if at startup, its name and IP address information are consistent with the same name and address information it used prior to being shut down, it sends a refresh to renew its associated resource records for this information.

- A periodic refresh is sent by the computer while it is running.

- The Windows XP and Windows Server 2003 DNS Client service renews DNS registration of client resource records every 24 hours. When this dynamic update occurs, if the dynamic update request does not cause modification to the DNS database, then it is considered to be a refresh and not a resource record update.

- Other network services make refresh attempts, such as: DHCP servers which renew client address leases, cluster servers which register and update records for a cluster, and the Net Logon service, which can register and update resource records used by Active Directory domain controllers.

**Record update** When a DNS dynamic update is processed for a resource record where other characteristics of the record in addition to its time stamp are revised. Updates generally occur for the following reasons:

- When a new computer is added to the network and, at startup, it sends an update to register its resource records for the first time with its configured zone.

- When a computer with existing records in the zone has a change in IP address, causing updates to be sent for its revised name-to-address mappings in DNS zone data.

- When the Net Logon service registers a new Active Directory domain controller.

**Refresh interval** An interval of time, determined for each zone, as bounded by the following two distinct events:

- The earliest date and time when the record becomes eligible to be refreshed and have its time stamp reset.

- The earliest date and time when the record becomes eligible to be scavenged and removed from the zone database.

This value should be large enough to allow all clients to refresh their records. By default, this interval is set to seven days. It should not be increased to an unreasonably high level, because the benefits of the aging and scavenging feature might either be lost or diminished.

**Resource record (RR) time stamp** A date and time value used by the DNS server to determine removal of the resource record when it performs aging and scavenging operations.

**Scavenging period** When automatic scavenging is enabled at the server, this period represents the time between repetitions of the automated scavenging process. The default value for this is seven days. To prevent deterioration of DNS server performance, the minimum allowed value for this is one hour.

**Scavenging servers** An optional advanced zone parameter that enables you to specify a restricted list of IP addresses for DNS servers that are enabled to perform scavenging of the zone. By default, if this parameter is not specified, all DNS servers that load a directory-integrated zone (also enabled for scavenging) attempt to perform scavenging of the zone. In some cases, this parameter can be useful if it is preferable that scavenging only be performed at some servers loading the directory-integrated zone. To set this parameter, you must specify the list of IP addresses for the servers enabled to scavenge the zone in the ScavengingServers parameter for the zone. This can be done using the dnscmd command, a command-line based tool for administering Windows DNS servers.

**Start scavenging time** A specific time, expressed as a number. This time is used by the server to determine when a zone becomes available for scavenging.

#### When Scavenging Can Start

Once all prerequisites for enabling the use of scavenging are met, scavenging can start for a server zone when the current server time is greater than the value of the start scavenging time for the zone.

The server sets the time value to start scavenging on a per-zone basis whenever any one of the following events occurs:

- Dynamic updates are enabled for the zone.

- A change in the state of the **Scavenge stale resource records** check box is applied. You can use the DNS console to modify this setting at either an applicable DNS server or one of its primary zones.

- The DNS server loads a primary zone enabled to use scavenging. This can occur when the server computer is started or when the DNS Server service is started.

- When a zone resumes service after having been paused.

When any of the previous events occur, the DNS server sets the value of start scavenging time by calculating the following sum:

Current server time + Refresh interval = Start scavenging time

This value is used as a basis of comparison during scavenging operations.

**Example of the aging and scavenging process for a sample record**
To understand the process of aging and scavenging at the server, consider the life span and successive stages of a single resource record, as it is added to a server and zone where this process is in effect and then aged and removed from the database.

1. A sample DNS host, "host-a.example.microsoft.com", registers its host (A) resource record at the DNS server for a zone where aging/scavenging is enabled for use.

2. When registering the record, the DNS server places a time stamp on this record based on current server time.

   After the record time stamp is written, the DNS server does not accept refreshes for this record for the duration of the zone no-refresh interval. It can, however, accept updates prior to that time. For example, if the IP address for "host-a.example.microsoft.com" changes, the DNS server can accept the update. In this case, the server also updates (resets) the record time stamp.

3. Upon expiration of the no-refresh period, the server begins to accept attempts to refresh this record.

   Once the initial no-refresh period ends, the refresh period immediately begins for the record. During this time, the server does not suppress attempts to refresh the record for its remaining life span.

4. During and after the refresh period, if the server receives a refresh for the record, it processes it.

   This resets the time stamp for the record based on the method described in step 2.

5. When subsequent scavenging is performed by the server for the "example.microsoft.com" zone, the record (and all other zone records) are examined by the server.

   Each record is compared to current server time on the basis of the following sum to determine whether the record should be removed:

   Record time stamp + **No-refresh interval** for zone + **Refresh interval** for zone

   - If the value of this sum is greater than current server time, no action is taken and the record continues to age in the zone.

   - If the value of this sum is less than current server time, the record is deleted both from any zone data currently loaded in server memory and also from the applicable DnsZone object store in Active Directory for the directory-integrated "example.microsoft.com" zone.

**Unicode Character Support**
Originally, Internet host names were restricted to the character set specified in RFCs 952 and 1123. These restrictions include limiting names to using uppercase and lowercase letters (A-"Z", a-z), numbers (0-9) and hyphens (-). In addition, the first character of the DNS name can be a number and names must be encoded and represented using US-ASCII-based characters.

These requirements were maintained when DNS was introduced as part of RFC 1035, one of the core DNS standards specifications. For use of DNS in international settings, this requirement has significant limitations where extended character sets are used for local naming standards.

To remove these limitations, Microsoft expands DNS character support beyond the RFC 1035 specification. The DNS service now provides enhanced default support for UTF-8, a Unicode transformation format.

**What is UTF-8?**
UTF-8 is the recommended character set for protocols evolving beyond the use of ASCII. The UTF-8 protocol provides for

support of extended ASCII characters and translation of UCS-2, a 16-bit Unicode character set that encompasses most of the world's writing systems. UTF-8 enables a far greater range of names than can be achieved using ASCII or extended ASCII encoding for character data.

Computers running Windows 2000, Windows XP, and Windows Server 2003 operating systems are UTF-8 aware. This means that when UTF-8 encoded characters are received or used as data by the server, the server can load and store this data in its zones. Although Windows-based DNS servers are UTF-8 aware, they remain compatible with other DNS servers that use traditional US-ASCII data encoding and current DNS standards.

**How the DNS Service Implements UTF-8**
To provide standards compatibility and interoperability with other DNS implementations, the DNS service uses uniform downcasing of any received character data. In this process, the DNS service converts all uppercase characters used in standard US-ASCII data to lowercase equivalent data for the following reasons:

- To maintain compatibility with current and existing DNS standards.
- To provide interoperability with DNS server implementations that do not recognize or support UTF-8 encoding.

To understand why uniform downcasing was chosen, several related points must first be considered from the current revised Internet standards for DNS. Several key points in the standards pertain directly to how character data is to be handled between DNS servers and other servers and clients. These include the following:

- Any binary string can be used in a DNS name. (RFC 2181)
- DNS servers must be able to compare names in a case-insensitive way. (RFC 1035)
- The original case for character data should be preserved whenever possible as the data is entered into the system. (RFC 1035)

Because case insensitivity is a required part of the core DNS standard and case preservation is an optional recommendation, uniform downcasing was chosen to provide an effective standards-compliant solution. By downcasing UTF-8 encoded names before transmission, other DNS servers (which are not UTF-8 aware) are able to receive and perform successful binary comparisons of the data and obtain the desired results.

**Considerations for Interoperability With UTF-8**
The DNS Server service can be configured to allow or disallow the use of UTF-8 characters on a per-server basis. Although other DNS server implementations that are not UTF-8 aware might be able to accept the transfer of a zone containing UTF-8 encoded names, these servers might not be able to write back those names to a zone file or reload those names from a zone file. Administrators should exercise caution when transferring a zone containing UTF-8 names to a DNS server that is not UTF-8-aware.

Some protocols place restrictions on the characters allowed in a name. In addition, names that are intended to be globally visible (RFC 1958) should contain ASCII-only characters, as recommended in RFC 1123.

The use of UTF-8 for transformation of Unicode characters is not noticeable for general users. Only where Network Monitor or another similar tool is used to analyze DNS-related traffic over the physical network are UTF-8 encoded characters observable.

In addition to DNS server support for the UTF-8 encoding format, the client resolver defaults to using the UTF-8 character encoding format.

Names encoded in UTF-8 format must not exceed the size limits clarified in RFC 2181, which specifies a maximum of 63 octets per label and 255 octets per name. Character count is insufficient to determine size because some UTF-8 characters exceed one octet in length.

The UTF-8 encoding protocol adapts to use with existing DNS protocol implementations that expect US-ASCII characters because representation of US-ASCII characters in UTF-8 is identical, byte for byte, to the US-ASCII representation. DNS client or server implementations that do not recognize UTF-8 characters always encode names in the US-ASCII format. Those names are correctly interpreted by the DNS Server service.

The DNS service provides the ability to configure name checking to allow or restrict the use of UTF-8 characters in DNS data.

By default, multibyte UTF-8 name checking is used, allowing the greatest tolerance when the DNS service processes characters. This is the preferred name-checking method for most privately operated DNS servers that are not providing name service for

Internet hosts.

**WINS Lookup Integration**
Support for using Windows Internet Name Service (WINS) is provided to look up DNS names that cannot be resolved by querying the DNS domain namespace. To accomplish WINS lookup, two specific resource record types are used and can be enabled for any zones loaded by the DNS service:

- The WINS resource record, which can be enabled to integrate WINS lookup into forward lookup zones
- The WINS-R resource record, which can be enabled to integrate node adapter status request for reverse lookup zones

**WINS Resource Record**
The WINS and DNS services are used to provide name resolution for the NetBIOS namespace and the DNS domain namespace, respectively. Although both DNS and WINS can provide a separate and useful name service to clients, WINS is mainly needed to provide support for older clients and programs that require support for NetBIOS naming.

However, the DNS service can work with WINS to provide combined name searches in both namespaces when resolving a DNS domain name not found in zone information. To provide this interoperability, a new record (the WINS record) was defined as part of the zone database file.

The WINS resource record is specific to computers running Windows NT 4.0 and earlier, Windows 2000, and Windows Server 2003 operating systems and can be attached only to the domain of origin for a zone. The presence of a WINS resource record can instruct the DNS service to use WINS to look up any forward queries for host names or names that are not found in the zone database. This functionality is particularly useful for name resolution required by clients that are not WINS-aware (for example, UNIX) for the names of computers not registered with DNS, such as Windows 95 or Windows 98 computers.

**How WINS Lookup Works**
The following is an example of a DNS client (host-b)querying its DNS server in an attempt to look up the address for another computer named "host-a.example.microsoft.com."

**WINS Lookup**




In Step 1, the client queries its preferred DNS server. In Steps 2 through 8, the normal process of recursion proceeds as the preferred DNS server queries other DNS servers in succession on behalf of the client. This process concludes at Step 8, when the DNS server for the example.microsoft.com zone is located through the previous chain of referral answers. At this point in the process, the server contacted is a DNS server either running Windows NT Server 4.0, Windows 2000, or Windows Server 2003.

When the DNS server for the example.microsoft.com zone receives the query for host-a, it looks in its configured zone to see if a matching address (A) resource record (RR) can be found. If no A record is found and the zone is enabled to use WINS lookup, the server does the following:

- The DNS server separates the host part of the name (host-a) from the fully qualified domain name contained in the DNS query.

  The host part of the name is the first label in the queried DNS domain name before a period is used in the name.

- The server then sends a NetBIOS name request to the WINS server using the host name, host-a.

- If the WINS server can resolve the name, it returns the IP address to the DNS server.
- The DNS server then compiles an A resource record using the IP address resolved through the WINS server and returns this record to the original preferred DNS server that was queried by the requesting client, host-b.
- The preferred DNS server then passes the query answer back to the requesting client.

**How WINS Reverse Lookup Works**

There is also a WINS-R record or WINS reverse lookup entry that can be enabled and added to reverse lookup zones. However, because the WINS database is not indexed by IP address, the DNS service cannot send a reverse name lookup to WINS to get the name of a computer given its IP address.

Because WINS does not provide reverse lookup capability, the DNS service instead sends a node adapter status request directly to the IP address implied in the DNS reverse query. When the DNS server gets the NetBIOS name from the node status response, it appends the DNS domain name back onto the NetBIOS name provided in the node status response and forwards the result to the requesting client.

**Note**

- WINS and WINS-R resource records are proprietary to the DNS Server service provided by Windows. You can prevent these resource records from being included in zone transfers to other DNS server implementations. For more information, see "How WINS Lookup Works" in this section.

**Network Ports Used By DNS**

During DNS resolution, DNS messages are sent from DNS clients to DNS servers or between DNS servers. Messages are sent over UDP and DNS servers bind to UDP port 53. When the message length exceeds the default message size for a User Datagram Protocol (UDP) datagram (512 octets), the first response to the message is sent with as much data as the UDP datagram will allow, and then the DNS server sets a flag indicating a truncated response. The message sender can then choose to reissue the request to DNS server using TCP (over TCP port 53). The benefit of this approach is that it takes advantage of the performance of UDP but also has a backup failover solution for longer queries.

In general, all DNS queries are sent from a high-numbered source port (above 1023) to destination port 53, and responses are sent from source port 53 to a high-numbered destination port. The following table lists the UDP and TCP ports used for different DNS message types.

**UDP and TCP Port Assignments for DNS Servers**

| Traffic Type | Source of Transmission | Source Port | Destination of Transmission | Destination Port |
|---|---|---|---|---|
| Queries from local DNS server | Local DNS server | Any port number above 1023 | Any remote DNS server | 53 |
| Responses to local DNS server | Any remote DNS server | 53 | Local DNS server | Any port number above 1023 |
| Queries from remote DNS server | Any remote DNS server | Any port number above 1023 | Local DNS server | 53 |
| Responses to remote DNS server | Local DNS server | 53 | Any remote DNS server | Any port number above 1023 |

**Note**

- The Windows Server 2003 DNS Server service supports Extension Mechanisms for DNS (EDNS0 as defined in RFC 2671), which allow DNS requestors to advertise the size of their UDP packets and facilitate the transfer of packets larger than 512 bytes. When a DNS server receives a request over UDP, it identifies the requestor's UDP packet size from the option (OPT) resource record and scales its response to contain as many resource records as are allowed in the maximum UDP packet size specified by the requestor.

- Windows Server 2003 DNS support for EDNS0 is enabled by default. It can be disabled using the registry. Locate the following registry subkey:
- HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters
- Add the entry **EnableEDNSProbes** to the subkey. Give the entry a DWORD value and set it to 0x0 to disable EDNS0.
- Use extreme caution when editing the registry Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system.

**Related Information**
The following resources contain additional information that is relevant to this section.

**RFCs for the DNS Server Service**
Request for Comments (RFCs) are an evolving series of reports, proposals for protocols, and protocol standards used by the Internet community. Domain Name System (DNS) specifications are based on approved RFCs published by the Internet Engineering Task Force (IETF) and other working groups.

The following RFCs can be found in the IETF RFC Database. These documents contain specifications used to design and implement the DNS Server service and DNS Client service.

- RFC 1034 -- Domain Names — Concepts and Facilities
- RFC 1035 -- Domain Names — Implementation and Specification
- RFC 1123 -- Requirements for Internet Hosts — Application and Support
- RFC 1886 -- DNS Extensions to Support IP Version 6
- RFC 1995 -- Incremental Zone Transfer in DNS
- RFC 1996 -- A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)
- RFC 2136 -- Dynamic Updates in the Domain Name System (DNS UPDATE)
- RFC 2181 -- Clarifications to the DNS Specification
- RFC 2308 -- Negative Caching of DNS Queries (DNS NCACHE)
- RFC 2535 -- Domain Name System Security Extensions (DNSSEC)
- RFC 2671 -- Extension Mechanisms for DNS (EDNS0)
- RFC 2782 -- A DNS RR for specifying the location of services (DNS SRV)

**Internet Drafts for DNS**
The following Internet drafts contain specifications used to design and implement the DNS Server service and DNS Client service:

- Draft-skwan-utf8-dns-02.txt -- Using the UTF-8 Character Set in the Domain Name System
- Draft-ietf-dhc-dhcp-dns-08.txt -- Interaction between DHCP and DNS
- Draft-ietf-dnsind-tsig-11.txt -- Secret Key Transaction Signatures for DNS (TSIG)
- Draft-ietf-dnsind-tkey-00.txt -- Secret Key Establishment for DNS (TKEY RR)

- Draft-skwan-gss-tsig-04.txt -- GSS Algorithm for TSIG (GSS-TSIG)

**Other Specifications for DNS**

The following additional specification is used to design and implement the DNS Server service and DNS Client service. This specification is published by the ATM Forum. You can obtain this specification by downloading it from the ATM Forum FTP site.

- af-saa-0069.000.doc -- ATM Name System Specification version 1.0

**Obtaining DNS RFCs and Other Internet Draft Specifications**

You can obtain RFCs and other Internet drafts from the RFC Editor Web site. RFCs are classified as one of the following: approved Internet standards, proposed Internet standards (circulated in draft form for review), Internet best practices, or For Your Information (FYI) documents. Internet drafts are not assigned an RFC number but are specifications currently in the proposal stage and generally published by one of the IETF working groups, of which several groups focus on discussion of DNS-related topics.

### Note

- RFCs 1034 and 1035 define the original DNS standard protocol to support domain name services in a TCP/IP environment. These RFCs describe the protocols in a detailed manner, emphasizing the underlying ideas and techniques used in all DNS implementations.

**DNS Resource Records Reference**

A DNS database consists of one or more zone files used by the DNS server. Each zone holds a collection of structured resource records, the following of which are supported by the DNS Server service.

**Format of DNS Resource Records**

All resource records have a defined format that uses the same top-level fields, as described in the following table.

### Format of DNS Resource Records

| Field | Description |
|---|---|
| Owner | Indicates the DNS domain name that owns a resource record. This name is the same as that of the console tree node where a resource record is located. |
| Time to Live (TTL) | For most resource records, this field is optional. It indicates a length of time used by other DNS servers to determine how long to cache information for a record before expiring and discarding it.<br><br>For example, most resource records created by the DNS Server service inherit the minimum (default) TTL of 1 hour from the start of authority (SOA) resource record which prevents overlong caching by other DNS servers.<br><br>For an individual resource record, you can specify a record-specific TTL that overrides the minimum (default) TTL inherited from the SOA RR. TTL values of zero are not recommended as they increase DNS server utilization unnecessarily.<br><br>In the resource record examples below, the TTL field is omitted wherever it is optional. The TTL field is included in the syntax for each record to indicate where it may be added. |
| Class | Contains standard mnemonic text indicating the class of the resource record. For example, a setting of “IN” indicates that the resource record belongs to the Internet class, which is the only class supported by Windows Server 2003 DNS. This field is required. |
| Type | Contains standard mnemonic text indicating the type of resource record. For example, a mnemonic of “A” indicates that the resource record stores host address information. This field is required. |
| Record-specific data | A required, variable-length field that contains information describing the resource. The format of this information varies according to the type and class of the resource record. |

**DNS Resource Records (Alphabetical List)**

**A**

**Description:** Host address (A) resource record. Maps a DNS domain name to an Internet Protocol (IP) version 4 32-bit address. For more information, see RFC 1035.

**Syntax::**

*owner class ttl* **A** *IP_v4_address*

**Example:**

```
host1.example.microsoft.com.      IN A 127.0.0.1
```

**AAAA**

**Description:** IPv6 host address (AAAA) resource record. Maps a DNS domain name to an Internet Protocol (IP) version 6 128-bit address. For more information, see RFC 1886.

**Syntax:** *owner class ttl* **AAAA** *IP_v6_address*

**Example:**

```
ipv6_host1.example.microsoft.com.      IN AAAA 4321:0:1:2:3:4:567:89ab
```

**AFSDB**

**Description:** Andrew File System Database (AFSDB) resource record. Maps a DNS domain name in the *server_host_name* field to the host name for a server computer of a server subtype. The **subtype** field can have either of the following recognized and supported values:

- A value of **1** to indicate that the server is an AFS version 3.0 volume location server for the named AFS cell.
- A value of **2** to indicate that the server is an authenticated name server holding the cell-root directory node for the server that uses either Open Software Foundation's (OSF) DCE authenticated cell-naming system or HP/Apollo's Network Computing Architecture (NCA).

For more information, see RFC 1183.

**Syntax::**

*owner ttl class* **AFSDB** *subtype server_host_name*

**Example:**

```
example.microsoft.com.   AFSDB   1 afs-server1.example.microsoft.com.
```

**ATMA**

**Description:** Asynchronous Transfer Mode address (ATMA) resource record. Maps a DNS domain name in the *owner* field to an ATM address referenced in the *atm_address* field.

For more information, see "ATM Name System Specification Version 1.0" at the ATM Forum FTP site. At the ATM Forum FTP site, download **af-saa-0069.000.doc** or **af-saa-0069.000.rtf**.

**Syntax:** *owner ttl class* **ATMA** *atm_address*

**Example:**

```
atm-host   ATMA   47.0079.00010200000000000000.00a03e000002.00
```

**CNAME**

**Description:** Canonical name (CNAME) resource record. Maps an aliased or alternate DNS domain name in the **owner** field to a canonical or primary DNS domain name specified in the *canonical_name* field. The canonical or primary DNS domain name used in the data is required and must resolve to a valid DNS domain name in the namespace.

| |
|---|
| **Syntax:** *owner ttl class* **CNAME** *canonical_name* |
| **Example:**<br>`aliasname.example.microsoft.com.  CNAME  truename.example.microsoft.com.` |

**HINFO**

| |
|---|
| **Description:** Host information (HINFO) resource record. Specifies the type of CPU and operating system in the *cpu_type* and *os_type* fields, respectively, for the host DNS domain name in the **owner** field. Well-known CPU and operating system types that are most often used are noted in RFC 1700. This information can be used by application protocols such as FTP, which use special procedures when communicating with computers of a known CPU and operating system type. |
| **Syntax:** *owner ttl class* **HINFO** *cpu_type os_type* |
| **Example:**<br>`my-computer-name.example.microsoft.com.  HINFO  INTEL-386 WIN32` |

**ISDN**

| |
|---|
| **Description:** Integrated Services Digital Network (ISDN) resource record. Maps a DNS domain name to an ISDN telephone number. Telephone numbers used with this record should follow ITU-T E.163/E.164 international telephone numbering standards, which are compatible with current international phone numbering plans already in use. For more information, see RFC 1183. |
| **Syntax:** *owner ttl class* **ISDN** *isdn_address sub_address* |
| **Example:**<br>`my-isdn-host.example.microsoft.com.  ISDN  141555555539699 002` |

**KEY**

| |
|---|
| **Description:** Public key resource record. Contains a public key that is associated with a zone. In full DNSSEC implementation, resolvers and servers use KEY resource records to authenticate SIG resource records received from a signed zones. KEY resource records are signed by the parent zone, allowing a server that knows a parent zone's public key to discover and verify the child zone's key. Name servers or resolvers receiving resource records from a signed zone obtain the corresponding SIG record then retrieve the zone's KEY record. For more information, see RFC 2535. |
| **Syntax:** *owner class* **KEY** *protocol digital_signature_algorithm (DSA) public_key* |
| **Example:**<br>`widgets.microsoft.com IN KEY 0x0000 3 0` |

**MINFO**

| |
|---|
| **Description:** Mailbox mail list information (MINFO) resource record. Specifies (in *responsible_mailbox*) a domain mailbox name for a responsible person who maintains a mailing list or mailbox specified in the **owner** field. The *error_mailbox* field can also be used to specify a domain mailbox that receives error messages related to this mailing list or mailbox. Mailboxes specified for responsible contacts and error forwarding must be the same as valid mailbox (MB) records that already exist in the current zone. For more information, see RFC 1035. |
| **Syntax:** *owner ttl class* **MINFO** *responsible_mailbox error_mailbox* |
| **Example:**<br>`administrator.example.microsoft.com.  MINFO resp-mbox.example.microsoft.com err-mbox.example.microsoft.com` |

**MX**

| |
|---|
| **Description:** Mail exchanger (MX) resource record. Provides message routing to a mail exchanger host, as specified in *mail_exchanger_host*, for mail sent to the domain name specified in the **owner** field. A 2-digit preference value indicates preferred ordering if multiple exchanger hosts are specified. Each exchanger host must have a corresponding host (A) address resource record in a valid zone. For more information, see RFC 1035. |
| **Syntax:** *owner ttl class* **MX** *preference mail_exchanger_host* |
| **Example:**<br>`example.microsoft.com. MX 10 mailserver1.example.microsoft.com` |

**NS**

| |
|---|
| **Description:** Used to map a DNS domain name as specified in **owner** to the name of hosts operating DNS servers specified in the *name_server_domain_name* field. |
| **Syntax:** *owner ttl* IN **NS** *name_server_domain_name* |
| **Example:**<br>`example.microsoft.com.  IN NS nameserver1.example.microsoft.com` |

**NXT**

| |
|---|
| **Description:** Next resource record. NXT resource records indicate the nonexistence of a name in a zone by creating a chain of all of the literal owner names in that zone. They also indicate what resource record types are present for an existing name. For more information, see RFC 2535. |
| **Syntax:** *owner class* **NXT** *next_domain_name last_record_type NXT* |
| **Example:**<br>`east.widgets.microsoft.com. IN NXT ftp.widgets.microsoft.com. A NXT` |

**OPT**

| |
|---|
| **Description:** Option resource record. One OPT resource record can be added to the additional data section of either a DNS request or response. An OPT resource record belongs to a particular transport level message, such as UDP, and not to actual DNS data. Only one OPT resource record is allowed, but not required, per message. For more information, see RFC 2671. |
| **Syntax:** *name* **OPT** *class ttl rdlen rdata* |
| **Example:**<br>`null OPT IN 1280 0 0` |

**PTR**

| |
|---|
| **Description:** Pointer (PTR) resource record. Points from the name in **owner** to another location in the DNS namespace as specified by *targeted_domain_name*. Often used in special domains such as the in-addr.arpa domain tree to provide reverse lookups of address-to-name mappings. In most cases, each record provides information that points to another DNS domain name location, such as a corresponding host (A) address resource record in a forward lookup zone. For more information, see RFC 1035. |
| **Syntax:** *owner ttl class* **PTR** *targeted_domain_name* |
| **Example:**<br>`1.0.0.10.in-addr.arpa.  PTR host.example.microsoft.com.` |

**RP**

**Description:** Responsible Person (RP) resource record. Specifies the domain mailbox name for a responsible person in *mailbox_name*. This name is then mapped to a domain name in *text_record_name* for which (TXT) resource records exist in the same zone. When RP records are used in DNS queries, subsequent queries are used to retrieve associated text (TXT) resource record information. For more information, see RFC 1183.

**Syntax:** *owner ttl class* **RP** *mailbox_name text_record_name*

**Example:**

```
example.microsoft.com. RP admin.example.microsoft.com.
admin-info.example.microsoft.com.
admin-info.example.microsoft.com. TXT "Joe Administrator,
(555) 555-0110"
```

**SIG**

**Description:** Signature resource record. Encrypts an RRset to a signer's (RRset's zone owner) domain name and a validity interval. For more information, see RFC 2535.

**Syntax:**

*owner class* **SIG** *ttl signature_expiration signature_inception key_identifier signer_name{digital_signature}*

**Example:**

```
widgets.microsoft.com IN SIG 86400 19700101000000 19700101000000 49292

widgets.microsoft.com.{

AIYADP8d3zYNyQwW2EM4wXVFdslEJcUx/fxkfBeH1El4ixPFhpfHF
ElxbvKoWmvjDTCmfiYy2X+8XpFjwICHc398kzWsTMKlxovpz2FnCTM=

}
```

**SOA**

**Description:** Start of authority (SOA) resource record. Indicates the name of origin for the zone and contains the name of the server that is the primary source for information about the zone. It also indicates other basic properties of the zone. The SOA resource record is always first in any standard zone. It indicates the DNS server that either originally created it or is now the primary server for the zone. It is also used to store other properties such as version information and timings that affect zone renewal or expiration. These properties affect how often transfers of the zone are done between servers authoritative for the zone.

In the example below, the owner (primary DNS server) is specified as "@" because the domain name is the same as the origin of all data in the zone (example.microsoft.com.). This is a standard notation convention for resource records and is most often seen in the SOA record.

**Syntax:**

*owner class* **SOA** *name_server responsible_person (serial_number refresh_interval retry_interval expiration minimum_time_to_live)*

**Example:**

```
@   IN SOA     nameserver.example.microsoft.com.
postmaster.example.microsoft.com. (
1            ; serial number
3600         ; refresh   [1h]
600          ; retry     [10m]
86400        ; expire    [1d]
3600 )       ; min TTL   [1h]
```

**SRV**

**Description:** Service locator (SRV) resource record. Allows multiple servers providing a similar TCP/IP-based service to be located using a single DNS query operation. This record enables you to maintain a list of servers for a well-known server port and transport protocol type ordered by preference for a DNS domain name. For example, in Windows Server 2003 DNS, it provides the means to locate domain controllers that use Lightweight Directory Access Protocol (LDAP) service over TCP port 389.

The purposes of each of the specialized fields used in an SRV resource record are as follows:

**service** A symbolic name for the desired service. For well-known services, a reserved universal symbolic name such as "_telnet" or "_smtp" is defined in RFC 1700. If a well-known service name is not defined in RFC 1700, a local or user-preferred name can be used instead. Some widely used TCP/IP services, notably the Post Office Protocol (POP), do not have a single universal symbolic name. If RFC 1700 assigns a name for a service indicated in this field, the RFC-defined name is the only name that is legal to use. Only locally defined services can be named locally.

**protocol** Indicates the transport protocol type. Typically, this is either TCP or UDP, although any transport protocol named in RFC 1700 can be used.

**name** The DNS domain name referred to by this resource record. The SRV resource record is unique among other DNS record types in that it is not used to perform the search or query.

**priority** Sets the preference for a host specified in the **target** field. DNS clients that query for SRV resource records attempt to contact the first reachable host of the lowest numbered preference listed here. Although target hosts have the same stated preference value, they can be tried in random order. The range of preference values is 0 to 65535.

**weight** Can be used in addition to *preference* to provide a load-balancing mechanism where multiple servers are specified in the **target** field and are all set to the same level of preference. When selecting a target server host among those of equal preference, this value can be used to set an added level of preference that can be used to determine the exact order or balancing of selection for the target hosts used in an answered SRV query. When a non-zero value is used, servers of equal preference are tried in proportion to the weight of this value. The range of values is 1 to 65535. If load balancing is not needed, use a value of **0** in this field to make the record easier to read.

**port** The server port on the **target** host that provides the service indicated in the **service** field. The range of port numbers is 0 to 65535, although the number is often a well-known assigned service port number, as specified in RFC 1700. Unassigned ports can be used as needed.

**target** Specifies the DNS domain name of the host that provides the type of service being requested. For each host name used, a corresponding host address (A) resource record in the DNS namespace is required. A single period (**.**) can be used in this field to indicate authoritatively that the requested service specified in this SRV resource record is not available at this DNS domain name.

For more information, see the Internet draft "A DNS RR for specifying the location of services (DNS SRV)."

**Syntax:** service.protocol.name ttl class **SRV** preference weight port target

**Example:**

```
_ldap._tcp._msdcs    SRV 0 0   389 dc1.example.microsoft.com
SRV 10 0   389 dc2.example.microsoft.com
```

**TXT**

**Description:** Text (TXT) resource record. Maps a DNS domain name specified in the **owner** field to a string of characters in *text_string* serving as descriptive text. For more information, see RFC 1035.

**Syntax:**

*owner ttl class* **TXT** *text_string*

**Example:**

```
example.microsoft.com.    TXT  "This is an example of additional domain name
information."
```

**WKS**

| |
|---|
| **Description:** Well-known service (WKS) resource record. Describes the well-known TCP/IP services supported by a particular protocol on a specific IP address. WKS records provide TCP and UDP availability information for TCP/IP servers. If a server either supports both TCP and UDP for a well-known service or has multiple IP addresses that support a service, then multiple WKS records are used. For more information, see RFC 1035. |
| **Syntax:**<br><br>*owner ttl class* **WKS** *address protocol service_list* |
| **Example:**<br><br>`example.microsoft.com.  WKS 10.0.0.1 TCP ( telnet smtp ftp )` |

**X25**

| |
|---|
| **Description:** X.25 (X25) resource record. Maps a DNS domain name in the **owner** field to a Public Switched Data Network (PSDN) address number specified in *psdn_number*. PSDN numbers used with this record should follow the X.121 international numbering plan. For more information, see RFC 1183. |
| **Syntax:**<br><br>*owner ttl class* **X25** *psdn_number* |
| **Example:**<br><br>`example.microsoft.com.  X25 52204455506` |

**DNS Server Log Reference**

The following is a partial list of events recorded in the DNS server log. Event messages logged by the DNS Client service are logged in System log, and can also be viewed using Event Viewer at the applicable client computer. If you need to troubleshoot persistent problems, you can also enable for use advanced debugging options at the DNS server for obtaining a more detailed log of server activity. You can view this log by using the DNS console after the DNS Server service is installed on the server computer.

**Some Events Recorded in the DNS Server Log**

| Event ID | Description |
|---|---|
| 2 | The DNS server has started.<br><br>This message generally appears at startup when either the server computer is started or the DNS Server service is manually started. |
| 3 | The DNS server has shut down.<br><br>This message generally appears when either the server computer is shut down or the DNS Server service is stopped manually. |
| 408 | The DNS server could not open socket for address *IPaddress*. Verify that this is a valid IP address for the server computer.<br><br>If this is a valid IP address for the server computer, verify that no other application that would attempt to use the same DNS server port (such as another DNS server application) is running. By default, DNS uses TCP port 53. |

| | |
|---|---|
| 413 | The DNS server will send requests to other DNS servers on a port other than its default port (TCP port 53).<br><br>This DNS server is multihomed and has been configured to restrict DNS Server service to only some of its configured IP addresses. For this reason, there is no assurance that DNS queries made by this server to other remote DNS servers will be sent using one of the IP addresses enabled for the DNS server.<br><br>This might prevent query answer responses returned by these servers from being received on the DNS port that the server is currently configured to use. To avoid this problem, The DNS server sends queries to other DNS servers using an arbitrary non-DNS port, and the response is received regardless of the IP address used.<br><br>Typically, limiting the DNS server to using only its configured DNS port for sending queries to other DNS servers is the response to this event. This configuration is performed using the DNS console's server properties **Interfaces** tab as follows:<br><br>1. Either select **All IP addresses** to enable the DNS server to listen on all configured server IP addresses<br><br>2. Or, if you continue to select and use **Only the following IP addresses**, limit the **IP address** list to a single server IP address. |
| 414 | The server computer currently has no primary DNS suffix configured. Its DNS name is currently a single label host name. For example, its currently configured name is "host" rather than "host.example.microsoft".com or another fully qualified name.<br><br>While the DNS server has only a single label name, default resource records created for its configured zones use only this single label name when mapping the host name for this DNS server. This can lead to incorrect and failed referrals when clients and other DNS servers use these records to locate this server by name.<br><br>In general, the DNS server should be reconfigured with a full DNS computer name appropriate for its domain or workgroup use on your network. |
| 708 | The DNS server did not detect any zones of either primary or secondary type. It will run as a caching-only server but will not be authoritative for any zones. |
| 3150 | The DNS server wrote a new version of zone *zonename* to file *filename*. You can view the new version number in the DNS manager Zone Properties dialog box, Serial Number field.<br><br>This event should only appear if the DNS server is configured to operate as a root server. |
| 6527 | Zone *zonename* expired before it could obtain a successful zone transfer or update from a master server acting as its source for the zone. The zone has been shut down.<br><br>This event ID might appear when the DNS server is configured to host a secondary copy of the zone from another DNS server acting as its source or master server. It is always a good practice to verify that this server has network connectivity to its configured master server.<br><br>If the problem were to continue, one or more of the following options is avaialble:<br><br>1. Delete the zone and recreate it, specifying either a different master server, or an updated and corrected IP address for the same master server.<br><br>2. If zone expiration continues, consider adjusting the expire interval. |

For more information about all other messages not described in this topic, refer directly to the message text itself as it appears in the DNS server event log.

# DNS Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS Tools and Settings**

**In this section**


- DNS Tools
- DNS Registry Entries
- DNS Group Policy Settings
- DNS WMI Classes
- Related Information


**DNS Tools**

The following tools are associated with the Domain Name System (DNS).

**Dnscmd.exe: Dnscmd**

**Category**

This tool is included in the Microsoft Windows Server 2003 Support Tools.

**Version compatibility**

This tool runs on the Microsoft Windows 2000 Server and Windows Server 2003 operating systems.

Dnscmd is used to view the properties of DNS servers, zones, and resource records. In addition, Dnscmd is used to modify all aspects of the DNS Server service, including creating and deleting zones and resource records. Dnscmd can also be useful for developing scripts for configuring a DNS server.

**Dnslint.exe: DNSLint**

**Category**

This tool is a free download from Microsoft.

**Version compatibility**

This tool runs on Microsoft Windows XP, and the Windows 2000 Server and Windows Server 2003 operating systems.

DNSLint is a Microsoft Windows utility that can be used to help diagnose common DNS name resolution issues. It can be targeted to look for specific DNS record sets and ensure that they are consistent across multiple DNS servers. It can also be used to verify that DNS records used specifically for Active Directory directory service replication are correct.

**Dnsmgmt.msc: DNS Console**

**Category**

This tool is included in the Windows 2000 Server and Windows Server 2003 operating systems, and is installed when the DNS Server service is installed. This tool is also installed with either the Windows 2000 Server or Windows Server 2003 Administration Tools Pack (Adminpak.msi).

**Version compatibility**

This tool runs on the Windows 2000 Server and Windows Server 2003 operating systems. When installed from one of the administration tools packs, this tool can also run on Microsoft Windows 2000 Professional and Windows XP.

The DNS console is used to administer the DNS Server service. It can be used to modify all aspects of the DNS Server service, including creating and deleting zones and resource records, and forcing replication events between DNS server physical memory and DNS databases. The DNS console can also be used to perform diagnostics on the DNS infrastructure of a network.

**Eventvwr.exe: Event Viewer**

**Category**

This tool is included in all Windows server and client operating systems.

**Version compatibility**

This tool runs on Windows XP, and the Windows 2000 Server and Windows Server 2003 operating systems.

You can use Event Viewer to monitor events recorded in event logs. Typically, a computer stores the Application, Security, and System logs. It could also contain other logs, depending on the computer's role and the applications that are installed on it. For example, DNS servers write DNS-related events, such as any errors that occur while attempting to start the DNS Server service, to log files which can be read using Event Viewer.

**Ipconfig.exe: Ipconfig**
**Category**
This tool is included in all Windows server and client operating systems.

**Version compatibility**
This tool runs on all Windows server and client operating systems.

Ipconfig displays all current TCP/IP network configuration values and refreshes Dynamic Host Configuration Protocol (DHCP) and DNS settings. Used without parameters, Ipconfig displays IPv6 addresses or the IPv4 address, subnet mask, and default gateway for all adapters.

**Netmon.exe: Network Monitor**
**Category**
This tool is installed along with Windows Server 2003 and the Windows Server 2003 Administration Tools Pack.

**Version compatibility**
This tool runs on all Windows operating systems.

Network Monitor captures data about the packets on a network and logs them for subsequent analysis. The monitored data can be filtered many different ways including protocol, ports, physical addresses, and logical addresses. Network Monitor can be useful in many situations, such as when you are troubleshooting an environment that has a firewall between a DNS server and a client, or between two DNS servers.

Network Monitor comes in two versions: the Network Monitor that is provided as part of the Windows Server 2003 operating system, and the Network Monitor that is part of Microsoft Systems Management Server (SMS). The version of Network Monitor that is included with the Windows Server 2003 operating system only captures data about network packets being sent to or from the server on which you run Network Monitor; it also captures data about network broadcasts that are received. The Network Monitor that is included with SMS enables you to monitor all network packets on a network segment.

**Nslookup.exe: Nslookup**
**Category**
This tool is included in all Microsoft Windows server and client operating systems.

**Version compatibility**
This tool runs on all Windows server and client operating systems.

Nslookup is used to query DNS servers and to obtain detailed responses. The information obtained using Nslookup can be used to diagnose and solve name resolution problems, verify that resource records are added or updated correctly in a zone, and debug other server-related problems.

**DNS Registry Entries**
The following registry entries are associated with DNS.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, modify settings though the use of policy settings or through dedicated management tools, such as the DNS snap-in for the Microsoft Management Console (MMC), to accomplish tasks, rather than editing the registry directly. If you must edit the registry, use extreme caution.

### Note

- In Windows Server 2003, the location in the registry of the settings relating to DNS zones is different from the location used in Windows 2000 Server.
- In Windows 2000 Server, the location is:
- HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Zones

- In Windows Server 2003, the location is:
- HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\DNS Server\Zones

**DNS\Parameters**
The following registry entries are located under HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters.

**EDNSCacheTimeout**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

**Version**
Windows Server 2003

By default, the DNS server includes option (OPT) resource records indicating its User Datagram Protocol (UDP) maximum in responses to queries containing OPT resource records. The value of the registry key **EDNSCacheTimeout** determines how long the DNS server will keep information about the Extension Mechanisms for DNS (EDNS) versions supported by other DNS servers that have responded to a query with a OPT resource record. The value for the registry key is specified in seconds between 3,600 (1 hour) and 15,724,800 (182 days).

**EnableDirectoryPartitions**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

**Version**
Windows Server 2003

If you do not want the default DNS application directory partitions to be created automatically, you must disable the **EnableDirectoryPartitions** registry key. The values for this key are **0x0** (disable) and **0x1** (enable).

For more information about this registry entry, see the "Registry Reference" in the "Tools and Settings Collection."

**EnableEDNSProbes**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

**Version**
Windows Server 2003

The registry entry **EnableEDNSProbes** enables or disables the EDNS0 response for the DNS Server service. To configure the DNS server to include an OPT resource record only in response to EDNS0 requests containing OPT resource records, set the registry entry to a value of **0x1** (DWORD). In Microsoft Windows Server 2003 this setting is enabled by default. Setting the value to **0x0** (DWORD) will disable the feature.

**MaximumUdpPacketSize**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

**Version**
Windows Server 2003

The registry entry **MaximumUdpPacketSize** specifies the maximum UDP packet size advertised by the DNS server. The default value is 1,280 bytes. The value must be between 512 and 16,384 in decimal format (200 and 4,000 in hexadecimal format). UDP packets must travel through devices other than UDP hosts, such as routers, and some of these devices might not support UDP packets larger than 512 bytes. The maximum UDP packet length that is supported by all network devices should be determined, along with the path's Maximum Transmission Unit (MTU), before configuring this registry setting.

**RoundRobin**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

**Version**
Windows Server 2003

By default, DNS will perform round-robin rotation for all resource record types. The default setting for round-robin rotation is contained in the registry entry **RoundRobin** (REG_DWORD). By default, this entry's value is **1**, rotating all RR types except

those listed in the **DoNotRoundRobinTypes** registry entry. If the value of **RoundRobin** is set to **0**, then no resource record types will be round-robin rotated.

**DoNotRoundRobinTypes**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

**Version**
Windows Server 2003

By default, DNS will perform round-robin rotation for all resource record types. You can specify that certain resource record types are not to be round-robin rotated in the registry. The registry entry **DoNotRoundRobinTypes** (REG_SZ) allows a string value to be entered containing a list of resource record types. By modifying this entry, you turn off round-robin rotation for specific resource record types. For example, to prevent round-robin rotation for A, PTR, SRV, and NS record types, you would enter the following value for the registry entry: **a ptr srv ns**.

**DNSCache\Parameters**
The following registry entries are located under
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dnscache\Parameters.

**AlternateComputerNames**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dnscache\Parameters

**Version**
Windows Server 2003

The **AlternateComputerNames** registry entry specifies a list of alternate Fully Qualified Domain Names (FQDNs) for the computer. When the DNS Client service registers the host address (A) and pointer (PTR) resource records for the computer, it will also register the FQDNs specified in this key. This key is used by the Netdom.exe utility and is not typically modified by administrators.

**RegistrationRefreshInterval**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dnscache\Parameters

**Version**
Windows Server 2003

The **RegistrationRefreshInterval** registry entry specifies how often the DNS Client service will register the A DNS resource records for its host computer. The value for this registry key is specified in seconds.

**UpdateTopLevelDomainZones**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dnscache\Parameters

**Version**
Windows Server 2003

Windows XP and Windows 2000 with Service Pack 4 (SP4) and Windows Server 2003-based computers never try to dynamically update the root," "., zone.

The domain members and domain controllers in a domain with a single-label DNS name usually need to dynamically register DNS records in a single-label DNS zone that matches the DNS name of the domain. If an Active Directory forest root has a single-label DNS name, all domain controllers in the forest usually need to dynamically register DNS records in a single-label DNS zone that matches the DNS name of the forest root.

By default, Windows XP and Windows 2000 Service Pack 4 (SP4) and Windows Server 2003-based DNS clients do not attempt dynamic updates of any single-label DNS zone (.com, .net, .org, corp, for example). To enable such clients to attempt dynamic updates of a single-label DNS zones, set the **UpdateTopLevelDomainZones** (REG_DWORD) registry value to **0x1**.

These configuration changes should be applied to all domain controllers and members of an Active Directory domain with single-label DNS names. If a domain with a single-level domain name is a forest root, these configuration changes should be applied to all of the domain controllers in the forest, unless the separate zones _msdcs.*forest name*, _sites.*forest name*, _tcp.*forest name* and _udp.*forest name* are delegated from the *forest name* zone. You must restart the computer for the **UpdateTopLevelDomainZones** registry edit to take effect.

**Tcpip\Parameters**

The following registry entries are located under
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters.

**RegistrationEnabled (This registry entry is a global setting.)**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version**
Windows Server 2003

The **RegistrationEnabled** registry entry specifies that the DNS Client service should register all of the network connections of a computer in DNS. This is a global setting that is applied to all interfaces on a computer. If the need exists to apply this setting to an individual interface, see the registry entry for
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters \interfaces\[*interfaceGUID*}.

**SearchList**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Version**
Windows Server 2003

The **SearchList** registry entry specifies the DNS suffix search list for the network connection used by the computer. For DNS clients, the DNS suffix search list extends or revises their DNS search capabilities. By adding additional suffixes to the list, the DNS client searches for short, unqualified computer names in more than one specified DNS domain. Then, if a DNS query fails, the DNS Client service can use this list to append other name suffix endings to the original name query and then repeat DNS queries to the DNS server for these alternate FQDNs. When the suffix search list is empty or unspecified, the primary DNS suffix of the computer is appended to short unqualified names, and a DNS query is used to resolve the resultant FQDN. If this query fails, the computer can try additional queries for alternate FQDNs by appending any connection-specific DNS suffix configured for network connections.

**Interfaces\{interfaceGUID}**
The following registry entries are located under
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\{*interfaceGUID*}.

**NameServer**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\{*interfaceGUID*}

**Version**
Windows Server 2003

The **NameServer** registry entry specifies the IP address of the DNS server that is used by the network connection of the computer.

**RegistrationEnabled (This is an interface-specific registry setting.)**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\{*interfaceGUID*}

**Version**
Windows Server 2003

The **RegistrationEnabled** registry entry specifies that the DNS Client service should register the computer's network connection in DNS. This entry is interface specific and can be set on any interfaces on a computer. The {*interfaceGUID*} portion of the registry path determines which interface the setting is applied to.

There is a global version of this setting that gets applied to all interfaces. Information about this setting can be found under the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\RegistrationEnabled.

For more information about this registry entry, see the "Registry Reference" in the Tools and Settings Collection.

**DNS Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with the DNS Client service.

### Group Policy Settings Associated With DNS Client Service

| Group Policy Name | Purpose |
|---|---|

| | |
|---|---|
| Primary DNS Suffix | Specifies the primary DNS suffix. |
| Dynamic Update | Specifies whether or not the DNS client should perform dynamic updates. |
| DNS Suffix Search List | Specifies the DNS server search list. |
| Primary DNS Suffix Devolution | Specifies whether DNS name devolution should be used in resolving names. |
| Register PTR Records | Specifies whether reverse lookup name registrations should be performed. |
| Registration Refresh Interval | Specifies the default refresh interval. |
| Replace Addresses In Conflicts | Specifies whether addresses in conflict should be replaced. |
| Register DNS records With Connection-specific DNS Suffix | Specifies whether DNS registration should be attempted for each network adapter. |
| TTL Set in the A and PTR Records | Specifies Time to Live (TTL) for dynamic registrations. |
| Update Security Level | Specifies the security level for updates. |
| Update Top Level Domain Zones | Specifies whether the computer sends dynamic updates to top level domain zones. |

For more information about Group Policy settings, see the "Group Policy Settings Reference" in the "Tools and Settings Collection."

**DNS WMI Classes**

The following table lists and describes the WMI classes that are associated with the DNS Server service.

**WMI Classes Associated With DNS Server Service**

| Microsoft DNS WMI Class | Description |
|---|---|
| MicrosoftDNS_Server | Describes a DNS server. Every instance of this class might be associated with one instance of class MicrosoftDNS_Cache, one instance of class MicrosoftDNS_RootHints, and multiple instances of class MicrosoftDNS_Zone. |
| MicrosoftDNS_Domain | Represents a domain in a DNS hierarchy tree. |
| MicrosoftDNS_Zone | Describes a DNS zone. Every instance of the class MicrosoftDNS_Zone must be assigned to exactly one DNS Server. Zones might be associated with multiple instances of the classes MicrosoftDNS_Domain and MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_Cache | Describes a cache existing on a DNS server (do not confuse this with a cache file that contains root hints). This class simplifies visualizing the containment of DNS objects, rather than representing a real object. The class, MicrosoftDNS_Cache, is a container for the resource records cached by the DNS server.<br><br>Every instance of the class MicrosoftDNS_Cache must be assigned to exactly one DNS server. It might be associated with multiple instances of MicrosoftDNS_Domain and MicrosoftDNS_ResourceRecord. |

| | |
|---|---|
| MicrosoftDNS_RootHints | Describes the RootHints stored in a cache file on a DNS server. This class simplifies visualizing the containment of DNS objects, rather than representing a real object. Class MicrosoftDNS_RootHints is a container for the resource records stored by the DNS server in a cache file.<br><br>Every instance of the class MicrosoftDNS_RootHints must be assigned to exactly one DNS server. It might be associated with multiple instances of the MicrosoftDNS_ResourceRecord class. |
| MicrosoftDNS_Statistic | Represents a single DNS server statistic. |
| MicrosoftDNS_ServerDomainContainment | Every instance of the class MicrosoftDNS_ServerDomainContainment might contain multiple instances of the class MicrosoftDNS_Domain. |
| MicrosoftDNS_DomainDomainContainment | Every instance of the MicrosoftDNS_DomainDomainContainment class might contain multiple other instances of MicrosoftDNS_Domain. |
| MicrosoftDNS_DomainResourceRecordContainment | Every instance of the class MicrosoftDNS_DomainResourceRecordComtainment might contain multiple instances of the MicrosoftDNS_ResourceRecord class. |
| MicrosoftDNS_ResourceRecord | Represents the general properties of a DNS RR. |
| MicrosoftDNS_AAAAType | Represents an IPv6 Address (AAAA), often pronounced quad-A, RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_AFSDBType | Represents an Andrew File System Database Server (AFSDB) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_ATMAType | Represents an ATM Address-to-Name (ATMA) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_AType | Represents an Address (A) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_CNAMEType | Represents a Canonical Name (CNAME) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_HINFOType | Represents a Host Information (HINFO) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_ISDNType | Represents an ISDN RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_KEYType | Represents a KEY RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MBType | Represents a Mailbox (MB) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MDType | Represents a Mail Agent for Domain (MD) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MFType | Represents a Mail Forwarding Agent for Domain (MF) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MGType | Represents an MG RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MINFOType | Represents an Mail Information (MINFO) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MRType | Represents a Mailbox Rename (MR) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_MXType | Represents a Mail Exchanger (MX) RR. Subclass of MicrosoftDNS_ResourceRecord. |

| | |
|---|---|
| MicrosoftDNS_NSType | Represents a Name Server (NS) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_NXTType | Represents a Next (NXT) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_PTRType | Represents a Pointer (PTR) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_RPType | Represents a Responsible Person (RP) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_RTType | Represents a Route Through (RT) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_SIGType | Represents a Signature (SIG) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_SOAType | Represents a Start Of Authority (SOA) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_SRVType | Represents a Service (SRV) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_TXTType | Represents a Text (TXT) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_WINSRType | Represents a WINS-Reverse (WINSR) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_WINSType | Represents a WINS RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_WKSType | Represents a Well-Known Service (WKS) RR. Subclass of MicrosoftDNS_ResourceRecord. |
| MicrosoftDNS_X25Type | Represents an X.25 (X25) RR. Subclass of MicrosoftDNS_ResourceRecord. |

For more information about many WMI classes, see the WMI SDK documentation on MSDN.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Microsoft Platform SDK on MSDN for more information about WMI classes that are associated with the DNS Server service.
- "Group Policy Settings Reference" in the "Tools and Settings Collection" for information about Group Policy settings that are associated with the DNS Client service.
- "Registry Reference" in the "Tools and Settings Collection" for information about registry entries that are associated with DNS.

# WINS Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WINS Technical Reference**
Windows Internet Name Service (WINS) is a NetBIOS name resolution service that allows client computers to register their NetBIOS names and IP addresses in a dynamic, distributed database and to resolve the NetBIOS names of network resources to their IP addresses.

WINS is needed in any Microsoft Windows network environment that includes operating systems or applications that rely on NetBIOS.

**In this subject**


- What Is WINS?
- How WINS Technology Works
- WINS Tools and Settings

# What Is WINS?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Role of WINS in the Network**

Although NetBIOS and NetBIOS names can be used with network protocols other than TCP/IP, WINS was designed specifically to support NetBIOS over TCP/IP (NetBT). WINS is required for any environment in which users access resources that have NetBIOS names. If you do not use WINS in such a network, you cannot connect to a remote network resource by using its NetBIOS name unless you use Lmhosts files, and you might be unable to establish file and print sharing connections.

The following figure illustrates the role of WINS for computers that use NetBIOS names. Typically, DHCP is used to assign IP addresses automatically.

**WINS Name Registration and Resolution**




In a typical scenario, the following occurs:

1. ClientA, which uses NetBIOS and is a WINS client, sends a name registration request to its configured primary WINS server (WINSA) when it starts up and joins the network. WINSA adds ClientA's NetBIOS name and IP address to the WINS database.

2. When ClientB needs to connect to ClientA by its name, it requests the IP address from the WINS server.

3. The WINS server locates the corresponding entry in its database and replies with ClientA's IP address.

**Summary of WINS Benefits**

WINS provides the following benefits over other NetBIOS name resolution methods:

- WINS name resolution reduces NetBIOS name query broadcast traffic because clients can query a WINS server directly instead of broadcasting queries.

- WINS enables the Computer Browser service to collect and distribute browse lists across IP routers.

- The WINS dynamic name-to-address database supports NetBIOS name registration and resolution in environments where DHCP-enabled clients are configured for dynamic TCP/IP address allocation.

- The WINS database also supports centralized management and replicates name-to-address mappings to other WINS servers.

- WINS and DNS can be used in the same environment to provide combined name searches in both namespaces.

**WINS and DNS**

WINS and DNS are both name resolution services for TCP/IP networks. While WINS resolves names in the NetBIOS namespace, DNS resolves names in the DNS domain namespace. WINS primarily supports clients that run older versions of Windows and applications that use NetBIOS. Windows 2000, Windows XP, and Windows Server 2003 use DNS names in addition to NetBIOS names. Environments that include some computers that use NetBIOS names and other computers that use domain names must include both WINS servers and DNS servers.

**Related Information**

The following resources contain additional information that is relevant to this section.

The following resources contain additional information that is relevant to this section.

- DHCP Technical Reference
- DNS Technical Reference
- TCP/IP Technical Reference

# How WINS Technology Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- NetBIOS Overview
- WINS Architecture
- WINS Protocol
- WINS Processes and Interactions
- Related Information


Windows Internet Name Service (WINS) is a name resolution service that maps NetBIOS names to IP addresses in networks that use NetBIOS over TCP/IP (NetBT). As such, the primary purpose of WINS is to support clients that run older versions of Windows and applications that use NetBIOS.

In earlier versions of Microsoft operating systems, NetBIOS names were required for locating network resources. In Microsoft Windows NT 4.0, for example, clients used NetBIOS to locate Windows NT 4.0 domain controllers, as well as for file and print sharing. Therefore, WINS is required for name resolution in network environments that include computers running Windows versions earlier than Windows 2000.

WINS is not required in a network that consists entirely of computers running Windows 2000, Windows XP, Windows Server 2003, or other TCP/IP-based systems, such as most versions of UNIX, that fully support the use of Domain Name System (DNS) names. If NetBIOS and, consequently, WINS are eliminated from such an environment, however, any applications and services that depend on NetBIOS, such as the Computer Browser service, will not function. Organizations that need to continue to support older versions of Windows still rely on NetBIOS and WINS.

WINS can be used in conjunction with DNS, which resolves domain names, to connect to any network resource by referencing the resource's user-friendly name. Used in conjunction with Dynamic Host Configuration Protocol (DHCP), WINS supports dynamic IP address assignments, thus eliminating the administrative burden of maintaining static IP addresses.

This section describes how WINS works in a properly configured environment, which includes the following:

- NetBT is enabled.
- DNS is configured correctly.
- DHCP is configured correctly.
- The network, including routers, is configured to handle name resolution traffic.
- WINS-configured servers are designed with sufficient CPU, memory, and hard disk space to handle WINS activity.

The following sections describe WINS components, how they fit in a network environment, and how NetBIOS name resolution processes work.

**NetBIOS Overview**

An understanding of NetBIOS is essential to understanding WINS. In the OSI reference model, NetBIOS is a session-layer protocol. Applications use this protocol to communicate over NetBIOS-compatible transports. NetBIOS and applications that use NetBIOS can run over both TCP/IP and Internetwork Packet Exchange (IPX). When it runs over TCP/IP, NetBIOS is commonly referred to as NetBT.

RFCs 1001 and 1002 define the functions and features of three NetBIOS services in a TCP/IP environment: name service, session service, and datagram service. The name service encompasses registering NetBIOS names, renewing and releasing

names, and resolving NetBIOS names to IP addresses. WINS is the Microsoft implementation of a NetBIOS name server (NBNS), as defined by the RFCs. The other two services are beyond the scope of this section.

**NetBIOS Names**

RFCs 1001 and 1002 provide for 16-byte NetBIOS names. The first 15 characters of a NetBIOS name are user-definable. Windows reserves the sixteenth character for a name suffix.

The user-definable portion of a NetBIOS name cannot begin with an asterisk (*). Use of extended characters, the underscore (_) and period (.) in particular, can result in name resolution errors in environments that also use DNS name resolution. NetBIOS names are not case-sensitive.

The name suffix, the final byte of a NetBIOS name reserved by Windows, specifies the resource type. A WINS client can use the suffix to identify the specific service it needs to access. That is, a WINS client can construct the NetBIOS name based on a computer, user, or domain name and append the suffix to locate a specific resource on that computer. For example, the Messenger service has a suffix of 0x03. To send a message to COMPUTER42, the client sends a message to the Messenger service on COMPUTER42 by specifying the following NetBIOS name:

```
COMPUTER42      [03]
```

In the full 16-character name notation, the first 15 characters are padded with spaces, if necessary. A NetBIOS name can be a unique name, which means it maps to a single IP address, or a group name, which means it maps to multiple IP addresses. For example, each computer that runs Windows Server 2003 has a Server service for sharing files and a Workstation service for accessing the files shared by the Server service. These services each have a unique NetBIOS name, so client computers can locate the appropriate server and use the file-sharing service. On the other hand, the domain controllers within a single domain share a group name.

The following table describes some of the NetBIOS name suffixes used by Windows Server 2003. The suffixes are listed in hexadecimal format. Other applications, such as Microsoft Exchange and Microsoft Systems Management Server (SMS), register NetBIOS names with different suffixes.

**Common Windows NetBIOS Name Suffixes**

| 15-Character Name | Suffix (Hexadecimal) | Type | Usage |
|---|---|---|---|
| *ComputerName* | 00 | Unique | Registered by the Workstation service on a WINS client. Typically, this name is called the NetBIOS computer name. |
| *DomainName* | 00 | Group | Registered by the Workstation service for receiving browser broadcasts from LAN Manager-based computers. |
| *MSBROWSE* | 01 | Group | Registered by the master browse server for each subnet. |
| *ComputerName* | 03 | Unique | Registered by the Messenger service on a WINS client. This name is usually appended to the NetBIOS computer name and to the name of the logged-on user. (The Messenger service is not related to Windows Messenger.) |
| *UserName* | 03 | Unique | Identifies the Messenger service for the logged-on user. User names are registered in the WINS database by the Messenger service. |
| *ComputerName* | 1B | Unique | Registered by each Windows Server 2003 domain controller running as the domain master browse server. This name record is used to allow remote browsing of domains. When a WINS server is queried for this name, a WINS server returns the IP address of the computer that registered this name. |
| *DomainName* | 1B | Unique | Registered by each Primary Domain Controller (PDC) that is running as the domain master browser. This name is used to support remote browsing of domains. When a WINS server is queried for this name, it returns the IP address of the computer that registered the name. |
| *DomainName* | 1C | Group | Registered for use by domain controllers within the domain. This NetBIOS name can include up to 25 IP addresses: one for the PDC and 24 for backup domain controllers (BDCs). |

| | | | |
|---|---|---|---|
| *DomainName* | 1D | Unique | Registered only by a master browse server. There is only one master browse server per subnet. Backup browse servers use this name to communicate with the master browse server to retrieve the list of available servers from the master browse server. WINS servers always return a positive registration response for *DomainNname*[1D], even though they do not add this name to the WINS database. Therefore, when a WINS server is queried for D*omainName*[1D], the server always responds with a broadcast address, which forces the client to broadcast to resolve the name. |
| *DomainName* | 1E | Group | Registered for any computers configured to be network browse servers. This is a normal group name. Any computers configured to be network browse servers can broadcast to this name and listen for broadcasts to this name to elect a master browse server. A statically-mapped group name uses this name to register itself on the network. When a WINS server receives a name query for a name ending with [1E], the WINS server always returns the network broadcast address for the local network of the requesting client. The client can then use this address to broadcast to the group members. |
| *ComputerName* | 1F | Unique | Registered by the Network Dynamic Data Exchange (NetDDE) services; appears only if the NetDDE services are started on the computer. |
| *ComputerName* | 20 | Unique | Registered by the Server service on the WINS client. The WINS client uses this service to share folders and printers on the network. |
| *GroupName* | 20 | Group | A special group (also called Internet group) name is registered with WINS servers to identify groups of computers for administrative purposes. For example, “printers” could be a registered group name used to identify an administrative group of print servers. |
| *ComputerName* | BE | Unique | Registered by the Network Monitoring Agent service and appears only if the service is started on the WINS client computer. If the computer name has fewer than 15 characters, plus symbols (**+**) are added to expand the name to 15 characters. |
| *ComputerName* | BF | Unique | Registered by the Network Monitoring Utility (included with Microsoft Systems Management Server). If the computer name has fewer than 15 characters, plus symbols (**+**) are added to expand the name to 15 characters. |

The NetBIOS namespace is flat, which means that names can be used only once within a network. By contrast, DNS uses fully qualified domain names (FQDNs), which combine a host name with the hierarchical name of its domain. For example, a NetBIOS name of WINSserver01 might be an FQDN of WINSserver01.reskit.com.

Being able to use a NetBIOS name only once within a network can be a restriction with imposing implications for large networks. Using a NetBIOS scope, as described in RFC 1001, is one solution to this situation. A character string, known as the scope identifier, works something like a domain name in DNS and defines the population of computers across which a registered NetBIOS name is known. The scope identifier is appended to NetBIOS names during name services processes. However, the meaning of scope varies based on the NetBIOS node. For example, the scope for B-node includes all the NetBIOS nodes in the local subnet, the scope for P-node includes all the NetBIOS nodes in the routed network that use the same NetBIOS name service, and the scope for M-node and H-node is an intersection of the B-node and P-node scopes. (NetBIOS nodes are defined later in this section.) This situation causes name resolution problems in environments that use more than one type of NetBIOS node. Because only NetBIOS nodes that use the same scope ID can resolve each other’s names, the use of different scope IDs can create communication problems. Therefore, use of the scope identifier is discouraged in Windows Server 2003.

**NetBIOS Nodes**

Client computers use NetBIOS name services to register, renew, and release names and to query and resolve IP addresses for network resources. NetBIOS session and datagram services use these name services to locate resources. For example, if a user wants to map a network drive to a file share on a remote computer by using Windows Explorer, the user types the remote computer name and share name in the **Map Network Drive** window. Windows Explorer uses NetBIOS name services to obtain the remote computer's IP address and then establishes a NetBIOS session over which a file-sharing session is established to access the shared resource.

NetBIOS name services are carried out either by using IP broadcasts or by using a NetBIOS name server. Broadcasts are appropriate in single-subnet local area networks (LANs), but rapidly cause administrative problems in multiple-subnet networks. A NetBIOS name server is the desirable alternative for larger networks. To allow for various sizes of networks, RFCs

1001 and 1002 use the concept of node type to determine the specific method a computer uses to perform name services. The following table describes the four NetBIOS node types.

**NetBIOS Node Types**

| Name Resolution Node | Description |
|---|---|
| B-node | Broadcast node. The client computer sends IP broadcast messages to all the computers on the subnet. The computer that has the requested name sends back a confirmation with its IP address. This is the default mode for non-WINS clients. |
| P-node | Point-to-point node. The client computer only communicates with a NetBIOS name server (in Windows Server 2003-based networks, this is the WINS server) to register names and resolve names to IP addresses. |
| M-node | Mixed node. The client computer uses a mix of B-node and P-node communications to register and resolve NetBIOS names: it first attempts to resolve a name by using broadcasts, and if that is unsuccessful it queries a WINS server. |
| H-node | Hybrid node. The client computer uses a hybrid of B-node and P-node: it first queries a WINS server, and if that is unsuccessful it sends broadcasts. |

The default node type for Windows-based WINS clients is B-node. When a computer is configured to use a WINS server, whether manually or through DHCP, it becomes H-node. The other NetBIOS node types are configured through DHCP by configuring the **WINS/NBT Node Type** option for a DHCP client.

Because B-node is based on broadcasts, it does not scale beyond a single subnet. Configuring routers to forward NetBIOS broadcasts also does not scale well and is not recommended. To address these limitations of RFC-compliant B-node, Microsoft extended B-node for name service operations. Microsoft's extension, known as modified B-node, adds the following two name lookup facilities:

**NetBIOS name cache**
An in-memory cache of recently resolved NetBIOS names that can be searched before attempting other name resolution methods.

**LMHosts file**
A flat file, located in *systemroot*\System32\Drivers\Etc, that contains a static list of NetBIOS names and their IP addresses. This file does not exist by default.

**WINS Architecture**
The WINS service is Microsoft's implementation of a NetBIOS name server. WINS consists of two primary components: WINS servers and WINS clients. Each WINS server includes a dynamic database of NetBIOS name-to-IP address mappings. WINS clients communicate with a WINS server to register their NetBIOS names in the WINS database and to query the WINS database to resolve the IP addresses associated with other NetBIOS names. A special kind of WINS client is the WINS proxy. WINS proxies support NetBIOS name services in a limited way for client computers that are not WINS-enabled. The following figure illustrates how WINS components fit in a network environment.

**WINS Component Configuration**




This example shows two subnets (Subnet 1 and Subnet 2) that include WINS clients. The third subnet (Subnet 3) includes B-node clients and a WINS proxy to act on the behalf of the B-node clients during name service requests. There are two WINS servers, each of which contains a WINS database. The WINS components illustrated in the figure are described in the following table.

**WINS Components**

| Component | Description |
|---|---|
| WINS server | A computer that provides the WINS name registration and resolution service. A WINS server replicates the WINS database to other WINS servers so that complete name-to-address resolution information is available, regardless of which WINS server a WINS client uses. |
| WINS client | Any WINS-enabled computer that communicates with a WINS server to register, refresh, or release its NetBIOS name or to resolve the NetBIOS name of another computer. |
| WINS proxy | A WINS client computer that is configured to act on the behalf of other, non-WINS-enabled computers, allowing these computers limited participation in name registration, release, and resolution processes. |
| WINS database | A dynamically updated list of NetBIOS names and their associated IP addresses, including IP addresses assigned by DHCP. In networks that have multiple WINS servers, the servers exchange database updates through replication. |

**WINS Servers**

A WINS server handles name registration requests from WINS clients by registering their names and IP addresses, provided the requested unique name is not in active use. It also responds to NetBIOS name queries submitted by clients by returning the IP address of a queried name if it is listed in the WINS database. Either local clients (on the same subnet as the WINS server) or remote clients (across a router from the server) can register with a WINS server.

When a WINS-enabled client starts on the network, it sends its name and IP address in a registration request directly to its configured primary WINS server. The WINS server that registers the NetBIOS name of the client in its database is said to be the owner of the WINS name entry.

WINS servers can replicate the contents of their databases to other WINS servers.

**Primary and Secondary WINS Servers**

WINS servers can act as either a primary WINS server or a secondary WINS server to a client. The difference between primary and secondary WINS servers is simply the priority in which clients contact them. A primary WINS server is the first server a client contacts to perform its NetBIOS name service operations. A client contacts a secondary WINS server only when a primary WINS server is unable to fulfill the request, for example if it is unavailable when the client makes the request or unable to resolve a name for the client.

If a primary WINS server fails to fulfill a request, the client makes the same request of its secondary WINS server. If more than two WINS servers are configured for the client, the client tries the additional secondary WINS servers until the list is exhausted or one of the WINS servers successfully responds to the request. After a client uses a secondary WINS server, it periodically tries to switch back to its primary WINS server for future name service requests.

**WINS Clients**
WINS clients are computers that are configured to make direct use of a WINS server. WINS clients attempt to register their names with a WINS server when they start or join the network. Thereafter, they query the WINS server to resolve NetBIOS names as needed. WINS clients communicate with WINS servers to:

- Register client names in the WINS database.
- Refresh client names in the WINS database.
- Release client names from the WINS database.
- Resolve names by obtaining name-to-IP address mappings from the WINS database.

Most WINS clients typically have more than one NetBIOS name that they must register for use with the network. These names are used to publish various types of network services, such as the Messenger or Workstation services, that each computer can use to communicate with other computers on the network.

WINS client computers are configured with either static or dynamic IP addresses. Clients that use static IP addresses are manually configured with a set of specific WINS servers. Clients that use dynamic IP addresses use the **WINS/NBNS Servers** option in DHCP to specify the WINS servers.

For Windows 2000, Windows XP, and Windows Server 2003, WINS clients can be configured with up to 12 secondary WINS servers. This feature is useful in an environment in which there are many mobile clients, many NetBIOS-based resources, and services are used often. In such an environment, updates to the WINS database might not be replicated throughout the network of WINS servers at the time of a request, and it can be helpful for clients to be able to query more than two WINS servers. Clients use the additional server addresses only if the primary and secondary servers fail to respond.

WINS clients running any of the following versions of Windows can communicate with a WINS server running Windows Server 2003:

- 32-bit versions of the Windows Server 2003 family
- 64-bit versions of Windows Server 2003, Datacenter Edition and Windows Server 2003, Enterprise Edition
- 32-bit versions of Windows XP Professional and Windows XP Home Edition
- 64-bit version of Windows XP Professional
- Windows Millennium Edition
- Windows 2000 Professional; Windows 2000 Server; Windows 2000 Advanced Server; and Windows 2000 Datacenter Server
- Windows NT Workstation 4.0 and Windows NT Server 4.0
- Windows 98
- Linux and UNIX clients (with Samba installed)

Clients that are not WINS-enabled can participate in name registration and resolution processes in a limited way by using WINS proxies.

**WINS Proxies**
A WINS proxy is a WINS client that can perform limited NetBIOS name services, such as registration, release, and resolution, on behalf of other non-WINS clients that are on a routed TCP/IP network.

Typically, older NetBIOS clients that are not, or cannot be, configured to use WINS are B-node clients. These computers use broadcasts to register their NetBIOS names on the network and to resolve NetBIOS name queries. Computers configured for B-node cannot use WINS directly because WINS servers do not respond to broadcasts. A WINS proxy can listen on the local subnet for B-node name query broadcasts and respond for the names that are not on the local subnet. WINS proxies

communicate with a WINS server by using unicast datagrams.

RCF 1001 recommends against using B-node name resolution in a routed network, but in practice, the B-node NetBIOS node type is sometimes useful and sometimes cannot be removed or updated. Microsoft designed WINS proxies to support B-node clients in a routed network.

**WINS Database**

The WINS database resides on the WINS server. It stores NetBIOS name-to-IP address mappings for the network. WINS clients use this mapping information to connect with other network resources. NetBIOS names are registered in the WINS database whenever a network resource becomes available for use, for example when a computer starts up. Therefore, the contents of the WINS database change over time as computers join and leave the network.

The size of the WINS database depends on the number of WINS clients in the network and the number of NetBIOS names each uses. Unique name entries typically are 42 bytes. Internet group entries can include up to 25 addresses and, therefore, are longer entries. The databases on all WINS servers in a network are about equal in size. All the databases have the same number of NetBIOS name entries, except for those pending replication.

The actual size of the database, however, can be much larger than the number of entries, because space is not automatically reclaimed when unused records are deleted. WINS database compaction recovers the unused space. Windows Server 2003 provides both dynamic and manual database compaction. The manual database compaction is most efficient, but requires the database to be offline. The dynamic database compaction runs during idle time after a database update and allows the database to remain online.

The use of multiple WINS servers in a NetBIOS environment provides redundancy and load balancing. When multiple WINS servers are used, updates to a WINS database replicate among the WINS servers designated as replication partners until the updates are in all the WINS databases in the network. Although each name entry is eventually replicated to all other WINS databases, a name entry is said to be owned by the WINS server that originated the entry.

WINS servers running Windows Server 2003 can replicate with WINS servers running earlier versions of Windows.

**WINS Database Files**

The Windows Server 2003 WINS database uses the Extensible Storage Engine (ESE). This database imposes no limit on the number of records that a WINS server can store or replicate.

WINS uses the Jet database format for storing its data. Jet produces the J*n*.log and other files by default in the *systemroot*\System32\Wins folder. The files that are created and used by the database on each WINS server are described in the following table.

**WINS Server Database Files**

| File | Description |
|---|---|
| J50.log and J50*xxxxx*.log | A log of all WINS database transactions. This file is used by WINS to recover data if necessary.<br><br>Log files maintain a specific size of 1024 KB. When WINS tries to write more transactions to the log than the file size can accommodate, it renames the current log file and creates a new one. The new log file has a file name in the format J*n*.log, where *n* is a decimal number, such as J50.log. The format of the renamed log file is Jet*xxxxx*.log, where each *x* denotes a hexadecimal number from 0 to F. Renamed log files are maintained in the same folder as the current log file.<br><br>All entries in the log file are written to the database and the log is deleted every three hours, when the WINS database is successfully backed up, and when the WINS Server service is shut down properly. |
| J50.chk | A checkpoint file that indicates the location of the last information successfully written from the transaction logs to the database. During data recovery, the checkpoint file indicates where recovery begins. The checkpoint file is updated every time data is written to the database file (Wins.mdb). |

| | |
|---|---|
| Wins.mdb | The WINS database, which contains two tables: the IP address-to-owner ID mapping table and the name-to-IP address mapping table.<br><br>The name-to-IP address mapping table stores NetBIOS names and the IP addresses assigned to them. These entries are created during name registration requests and during replication. This table includes a clustered index (logical order of the key values is the same as the physical order of the rows in the table) on the name field for quick record retrieval during name queries. The primary index is built by concatenating owner ID and version ID in ascending order for quick retrieval of records that fall within a range of version IDs.<br><br>The IP address-to-owner ID mapping table contains a row for each WINS server that has entries in the name-to-IP address mapping table. A row contains the IP address of the WINS server and its identifier, which is the owner ID field for the name entries it owns. This table is used to build the owner-to-version table that determines which replication partners have the latest updates. |
| Winstmp.mdb | A temporary file that is created by the WINS service. This file functions as a swap file during index maintenance operations and can remain in the *systemroot*\System32\Wins folder after a system failure. |
| Res#.log | These are reserved log files, which function in emergencies when the server runs out of disk space. If a server attempts to create another transaction log file and there is insufficient disk space, the server flushes any outstanding transactions into these reserved log files. The service then shuts down and logs an event to Event Viewer. |

To increase speed and efficiency of data storage, current transactions are written to log files rather than directly to the database. Therefore, the most current view of the data includes both the database and any transactions in the log files. Both the database and the log files are used for recovery if the WINS service abruptly or unexpectedly stops. If the service stops in an unexpected manner, the log files are automatically used to re-create the correct state of the WINS database.

**Name Data**

Name entries in the name-to-IP address mapping table include data used during name query requests, as well as data used to age and replicate the entries. Data used for name service requests or for aging and replicating entries is described in the following table.

**Name Data**

| Data | Description |
|---|---|
| Name | The NetBIOS name. |
| Address | The IP address associated with the name.<br><br>A unique name has a single IP address. A special group entry can have multiple IP addresses. |
| Name type | The type of name.<br><br>A name is classified by how it is added to the WINS database:<br><br>• Static (the name is manually added to the WINS database)<br><br>• Dynamic (the name is automatically added to the WINS database during NetBIOS name registration)<br><br>A name is also classified as one of the following:<br><br>• Unique name<br><br>• Normal group<br><br>• Special group, also known as an Internet group<br><br>• Multihomed (a computer that uses multiple IP addresses) |

| | |
|---|---|
| State | The status of the NetBIOS name entry, which can be one of the following:<br>• Active (name is in active use)<br>• Released (network node has shut down, name is not in use)<br>• Extinct, also known as tombstone (name was not reregistered, is ready for deletion) |
| Owner ID | The ID of the WINS server that owns the entry. The owner is the WINS server that first added the entry to the WINS database.<br>WINS automatically assigns and increments an integer value to represent owner IDs at the local server. These owner IDs are mapped to the WINS server IP addresses. Owner IDs are used internally at the local server, and their values differ from WINS server to WINS server. |
| Version ID | A counter that indicates which name entry updates need to be replicated. |
| Time stamp | A date and time stamp that is updated when a NetBIOS name entry is updated. WINS uses specific time intervals to determine when to renew or delete name entries from the database. WINS calculates the time stamp from the current date and time plus a time interval. |

**WINS Time Intervals**

Name entries in the WINS database have time stamps that determine when specific aging activities are to take place. WINS uses four configurable time intervals to calculate the time stamps. The following table describes the four time intervals.

**WINS Time Intervals**

| Time Interval | Description |
|---|---|
| Renew interval (also known as the Time-to-Live [TTL]) | Specifies when a client must refresh its name. The default is six days.<br>During a name registration request, the entry is time-stamped with the sum of the current time and this interval value. If a name entry is not refreshed before this time stamp expires, WINS releases the name entry. |
| Extinction interval (also known as tombstone interval) | Specifies the interval between when an entry is marked as released and when it is marked as extinct. The default depends on the renew interval and, if the WINS server has replication partners, on the maximum replication time interval, which is the configurable amount of time between pull replications. The maximum allowable value is six days.<br>When an entry is marked as extinct, the entry is time-stamped with the sum of the current time and the extinction time-out on the WINS server that is configured as the pulling replication partner (the partner that requests replication). |
| Extinction time-out (also known as tombstone time-out) | Specifies the interval between when an entry is marked extinct and when the entry is finally deleted (scavenged) from the database. The default depends on the renew interval and, if the WINS server has replication partners, on the maximum replication time interval. The default is six days. |
| Verification interval | Specifies the interval after which a WINS server must verify that old names it does not own are still active. The default value is 24 days.<br>When an active entry is replicated, it is time-stamped with the sum of the current time and the verification interval on the WINS server requesting replication (pull partner). |

The Windows Server 2003 defaults provided for these values take the following into consideration:

- Minimized level of network traffic
- Minimized load on WINS servers
- Various configurations in which WINS servers might be deployed

- Minimized window between replications
- Exceptional conditions, such as long weekends and large volumes in worst-case scenarios

**Server Clocks**
Replication and scavenging (database cleanup) algorithms rely on a reasonably consistent system clock. Setting the system clock forward or backward affects these algorithms. However, because time stamps are always entered locally, the WINS servers do not need to be time-synchronized.

**WINS Protocol**
The WINS protocol is based on, and is compatible with, the protocols defined for NetBIOS name services specified in RFCs 1001 and 1002. Therefore, WINS is compatible with any other implementation of these RFCs. That is, another RFC-compliant implementation of the client can communicate with a WINS server. Similarly, a Microsoft TCP/IP client can communicate with other implementations of an NBNS. To provide scalability for larger networks, however, the WINS implementation of NBNS supports replication of the WINS database between WINS servers. WINS replication is Microsoft proprietary technology, so WINS replication is not compatible with other NBNS implementations. This section briefly describes the format of NetBIOS packets and the ports used by NetBIOS.

**NetBIOS name services packets**
NetBIOS name service packets consist of messages sent between a WINS client and a WINS server. These message packets are similar in structure to DNS packets, but with additional types and codes for NetBIOS. The high-level format of name services packets is described in the following table.

**Name Services Packet Format**

| Message Section | Description |
|---|---|
| Name service header | Fixed length, 12 octets. This section contains information about the type of message and counts for the other records included in the message. |
| Question entries<br><br>or<br><br>Answer resource records | Variable length, optional. Length is defined in the header. This section contains the NetBIOS name for name registration, refresh, and release messages.<br><br>Variable length, optional. Length is defined in the header. This section contains the resource records returned in response to a request. |
| Authority resource records | This section is not used in Windows Server 2003. |
| Additional resource records | Variable length, optional. Length is defined in the header. This section contains other resource records that are not provided as a direct reply to a request. |

Each name service packet contains a header and one or more additional entries. The sections included for a packet depend on the message. The format of name service packets used by Windows Server 2003 is compatible with the specification provided in RFC 1002. For a detailed description of packet fields, see RFC 1002.

**WINS ports**
The ports listed in the following table are used by specific WINS services.

**WINS Port Assignments**

| Service Name | UDP | TCP |
|---|---|---|
| NetBIOS name service | 137 (client and server) | |
| WINS replication | | 42 |
| WINS replication partner automatic discovery | 42 | |

The WINS client and WINS Server services listen on User Datagram Protocol (UDP) port 137. UDP is an unreliable transport protocol in that it does not establish a session or connection between the sender and the receiver. Because of this, it is possible for more requests than responses to be sent. That is, a client might need to send a request multiple times before it receives a

response. However, a client never receives more than one response to a request.

**WINS Processes and Interactions**

WINS clients communicate with WINS servers for NetBIOS name services. WINS servers communicate with other WINS servers for WINS database replication. WINS servers also periodically run a process to identify and clean up extinct names in the WINS database.

NetBIOS name services include four types of processes between a WINS client and a WINS server. The following table briefly describes these processes, which are discussed in more detail in the following sections.

**NetBIOS Name Service Processes**

| Name Service Process | Description |
|---|---|
| Name registration | Name registration is a WINS client requesting the use of a NetBIOS name on the network. WINS validates that the NetBIOS name is not in active use and then adds it to the WINS database.<br><br>When a NetBIOS name is registered in the WINS database, other computers can resolve that name and use the resulting IP address to communicate with the network resource.<br><br>Names typically are registered when a NetBIOS application starts, such as when a computer or service starts up, or when a user logs on to the network. |
| Name renewal | Each NetBIOS name registered in the WINS database has a specified TTL. Before this time expires, the name is refreshed. |
| Name release | When a NetBIOS application shuts down, such as when a computer or service is stopped or a user logs off the network, the name is released by the WINS client and is marked as released in the WINS database. |
| Name resolution | When an application needs to communicate with a NetBIOS application on another NetBIOS node, the NetBIOS name for the other node is resolved to its IP address. |

Clients that are not configured to use WINS can also participate in these processes to a limited extent by using a WINS proxy. The remainder of this section describes what happens during typical and WINS proxy name service processes, the scavenge process, and the replication process.

**Name Registration**

Before a NetBIOS name can be resolved in a WINS environment, it must be registered with a WINS server. For example, when a WINS client starts up, it sends a name registration request to the WINS server for each NetBIOS name it uses. A name registration request can be for a unique name (maps to a single IP address) or for a group name (maps to multiple IP addresses).

As illustrated in the following figure, a WINS client (ClientC) sends a name registration request directly to its configured WINS server, WINSA.

**Name Registration**




In this example, ClientC sends a name registration request to its configured primary WINS server, WINSA. WINSA adds an entry for ClientC, with its IP address, in the WINS database and sends a positive reply back to ClientC.

Registering a unique NetBIOS name by using a WINS server works as follows:

1. When a network resource starts up, the WINS client sends a name registration request, which includes the NetBIOS name

and IP address, to its configured WINS server.

2. The WINS server checks to see if the name has already been registered, and if so, whether it is in active use. The WINS server sends a Wait for Acknowledgment (WACK) message to the WINS client while it attempts to determine whether the name can be registered.

3. If the name is not in use, the server adds the NetBIOS name and IP address to the WINS database and sends the client a positive name registration reply, which includes a TTL value that determines when the name will expire.

4. If the name is in active use, the WINS server does not add the name and address to the WINS database and sends the WINS client a negative name registration reply.

The precise process for registering NetBIOS names depends on the client's node type. B-node clients use broadcasts instead of WINS servers to register names. M-node and H-node clients combine the use of broadcasts and WINS P-node clients use only WINS.

Because WINS servers do not respond to broadcasts, the actions described in this section do not apply to broadcast registration. The following table describes the results of WINS server replies.

**WINS Server Replies**

| Server Reply | Result |
|---|---|
| No reply | If the WINS client has one or more secondary servers configured, it tries each server. If the WINS client does not receive an immediate response from a WINS server, it can try the server up to three times. If the WINS client does not receive a reply from any of its configured WINS servers, name registration fails and the application can display an error to the user. |
| Positive reply | The name can be resolved by other computers for communication and no other network resource can use the same name (if the name is unique). |
| Negative reply | The name registration fails. Windows Server 2003 processes, such as the Server service, write an error event to the system event log. |

**Name Conflicts**

During a NetBIOS name registration request, a WINS server verifies that the requested name is not a duplicate of a NetBIOS name already in use in the network. The WINS server first determines whether the name exists in the WINS database. If the name is already registered, the WINS server determines whether the name is in active use or whether it is available. A variety of factors determine whether the name is available.

In some cases, a WINS server can determine from the database entry whether a NetBIOS name is active. If the WINS server cannot determine in this way whether the name is active, it might need to challenge the name. To challenge a name, the WINS server sends a name query request to the WINS client that registered the name. If the name is active, the WINS client that registered the name sends a name query response and is said to defend the name.

The following factors influence the action taken by a WINS server when resolving a name conflict:

- Whether the name is a unique name or a group name.

- The state of the registered name. A registered name can be active, released, or extinct.

- Whether the name is owned by the WINS server handling the registration request or another WINS server.

- Whether the IP address is the same or different from the IP address specified in the name registration request.

- Whether the name is statically or dynamically added to the WINS database.

**Names with different IP addresses**

If a WINS client sends a name registration request for a unique name that already exists in the WINS database, but with a different IP address, the WINS server first checks the status of the existing name. If the name is in a released or extinct state, the

WINS server treats the request as a new registration and sends a positive name registration reply to the WINS client requesting the name.

If the database status of the existing name is active, the WINS server determines whether the unique name is still in use by sending a challenge to the WINS client that registered the name, as illustrated in the following figure.

**Name Challenge**




In this example, the following steps take place:

1. A WINS client (ClientA at IP address 192.168.0.17) sends a name registration request to its configured primary WINS server (WINSA).

2. WINSA finds the unique name "ClientA" with a different IP address in the WINS database. WINSA sends a Wait for Acknowledgement (WACK) message to the requesting WINS client. The WACK message includes the estimated time of the wait in the TTL field.

   WINSA then sends a name query request to the WINS client at IP address 192.168.0.20. WINSA waits 500 milliseconds between challenges, and if ClientA is multihomed, WINSA tries each IP address. If WINSA receives no response from the first challenge, it makes two additional attempts.

3. The name ClientA is still in active use, so the WINS client sends a name query response back to WINSA.

4. WINSA then sends a negative name registration response back to the requesting WINS client. If the registered name "ClientA" was no longer in active use, no name query response would be sent, and WINSA could send a positive name registration response to the requesting WINS client.

**Names with the same IP address**

If a WINS client sends a name registration request for a unique name with the same IP address as one that already exists in the WINS database, the action a WINS server takes depends on the status and ownership of the existing name, as follows:

- If the registered name has a status of active and is owned by the WINS server handling the new registration request, the WINS server sends a positive name registration reply to the WINS client making the request.

- If the registered name has a released or extinct status, or if it is owned by a different WINS server, the WINS server treats it as a new registration: it sends a positive name registration reply to the WINS client making the request and takes ownership of the name.

**Group names**

Group names are handled differently than unique names. A group name can be a normal group name or a special (also known as Internet) group name. Normal group names can have many associated IP addresses, but the member addresses are not stored in the WINS database. WINS responds to a name query for a normal group name with the limited broadcast address 255.255.255.255. Routers do not forward packets addressed to the limited broadcast address, so special groups were

developed to support communications between subnets.

Special groups are used for special, user-defined administrative groups. For example, a special group can be used to group resources such as file servers or printers. Special groups are associated with multiple IP addresses.

A WINS server handles name conflicts for a group name as follows:

- If a WINS client sends a unique name registration request that conflicts with a registered normal group name, a WINS server always sends a negative name registration reply.

- If a WINS client sends a name registration request for a special group name, a WINS server always sends a positive name registration reply, adding the requested name as an additional member of the group. In this situation, if the existing name has a state of released or extinct, the name registration request is treated as a new registration.

- If a WINS client sends a name registration request for a *DomainName*[1D] name, a WINS server always returns a positive registration response, even though the WINS server does not add this name to the WINS database. When a WINS client queries a WINS server for this type of name, the server responds with a broadcast address, which forces the WINS client to broadcast to resolve the name.

**Static names**

In addition to supporting dynamic name registration, WINS supports static name registration for computers that do not directly communicate with WINS. Static names are added manually to the WINS database with the WINS Microsoft Management Console (MMC) snap-in or with the **netsh wins server add name** command. Static names do not automatically age and drop from the WINS database like dynamically-added names do. Static names remain in the database indefinitely unless there is manual intervention.

If a WINS client sends a name registration request for an existing static name, a WINS server by default always sends a negative name registration reply. This behavior can be overridden to always send a positive name registration reply.

**Positive Name Registration Replies**

When a WINS server sends a WINS client a positive name registration reply for a name that does not yet exist, it adds the name entry to the WINS database as follows:

- Calculates time stamp by adding current time to the renew interval

- Takes ownership

- Sets state to active

- Assigns a new version ID

When a WINS server sends a WINS client a positive name registration reply for a name and IP address that already exist in the WINS database, it updates the name entry in the WINS database. The following table describes the database updates for unique, dynamic name registrations.

### Note

- A WINS server does not update existing database records for normal group entries or static entries because the server returns negative name registration replies for these types of name conflicts (unless the default behavior for group entries is overridden).

### Name Registration Database Updates

| State of Name | Server Action |
|---|---|
| Owned and active | Updates time stamp |
| Replica and active | Updates time stamp, takes ownership, increments version ID |
| Released | Updates time stamp, makes active, increments version ID |

| | |
|---|---|
| Owned and extinct | Updates time stamp, makes active, increments version ID |
| Replica and extinct | Updates time stamp, takes ownership, makes active, increments version ID |

The time stamp is calculated as the sum of the current time and the renew interval. Changes that increment the version ID are replicated to all other WINS servers.

**Burst Handling**

High-volume (burst) server loads occur when a large number of WINS clients simultaneously attempt to register their names with a WINS server, such as when power is restored after an outage. A WINS server can use burst mode to respond positively to client requests before it processes and updates the WINS database.

Burst mode uses a configurable burst-queue size as a threshold value. This value specifies how many name registration and name renewal (also known as refresh) requests a WINS server processes normally before it starts burst-mode handling. By default, the value is 500. When the number of name registration and refresh requests exceeds the threshold value, a WINS server initiates burst mode.

During burst mode, a WINS server immediately responds to registration and refresh requests with a positive response that includes a very small TTL. The TTL value included in the response varies for different clients to distribute the actual registration load over time. This process slows the refresh and retry rate for new WINS clients and regulates the amount of WINS client network traffic.

For example, if the burst queue size is 500 entries and more than 500 requests are active, the WINS server replies immediately to the next 100 WINS registration and refresh requests by sending immediate positive responses, with a starting TTL value of 5 minutes. For each additional round of 100 client requests, the WINS server adds 5 minutes to the TTL, up to a maximum of 50 minutes. If WINS client traffic still arrives at burst levels, the WINS server responds to the next round of 100 client requests starting again with a TTL value of 5 minutes. The WINS server continues to increment the TTL until it reaches its maximum intake of 25,000 requests. At 25,000 requests, the WINS server begins dropping requests.

During burst handling, name release requests are queued along with name registration and refresh requests. When the queue becomes full, however, release requests are silently dropped. In this case, the names are released when their TTL expires.

Name query requests are queued indefinitely during burst handling.

**Name Renewal**

Dynamic NetBIOS names are not owned permanently. When a WINS server successfully registers a name, it sends the WINS client a positive name registration reply with a TTL value that specifies the life span of the name. The default TTL value in Windows Server 2003 is six days. It is the responsibility of the WINS client to renew the name before the TTL expires.

If the WINS client does not renew the NetBIOS name before the TTL expires, the state of the name is changed to released, and the WINS server eventually deletes the name from the WINS database, as described later. Static names do not expire, so they do not need to be renewed.

To renew a name, a WINS client sends a name refresh request to the WINS server. The WINS server treats name refresh requests similarly to name registration requests, by sending a name refresh reply that includes a new TTL. The following figure illustrates a renewal request.

### Name Renewal




In this example, the WINS client, ClientC, sends a name refresh request to its configured primary server, WINSA. WINSA updates the database entry for ClientC to reflect an active state and sends a positive name refresh reply to ClientC with a new TTL.

WINS clients attempt to renew their NetBIOS names when one-half of the TTL value is reached or when the computer or the service restarts.

If the WINS server does not respond to the name refresh request, the WINS client continues to try the primary WINS server every 10 minutes for one hour, and then tries the secondary WINS server every 10 minutes for one hour. The process of trying first the primary server for an hour and then the secondary server for an hour continues until the name expires or the WINS server responds and renews the name.

**Name Release**

A NetBIOS name is released when it is no longer in active use by the WINS client that registered it. The owning WINS server does not respond to or resolve name queries for these names, unless the WINS client again registers the name.

Names can be released actively or passively. Expired names that are not renewed are released passively. WINS clients actively release their NetBIOS names when NetBIOS applications terminate properly, such as when a computer is shut down or a service is stopped, as shown in the following illustration.

**Name Release**




In this illustration, the following occurs:

1. When the computer, ClientC, shuts down properly, it sends a name release request to the WINS server, WINSA.

2. WINSA changes the state of ClientC's name entry in the WINS database to released, time-stamps the entry with the sum of the current time and the extinction interval, and takes ownership of the of name entry.

   **Note**

   - For active releases in Windows 98, the WINS server makes itself the owner of the name entry, if it is not already. In Windows NT 4.0 and later, if the WINS client sends a name release request to a different WINS server than it registered with, the WINS server changes the name entry to extinct instead of released, as described later in this section.

3. WINSA sends a positive name release response back to ClientC.

When a name entry is changed to the released state, the version ID is not updated and the change is not replicated to other WINS servers.

Released name entries are handled differently if the WINS server releasing the name is not the owner of the name entry, as described in the following table.

**Released Name Handling**

| Server Is Owner of Name Entry | Server Is Not Owner of Name Entry |
|---|---|
| State set to release | State set to extinct |
| Time stamp calculated from current time plus extinction interval | Time stamp calculated from current time plus extinction interval plus extinction time-out |
| Release status information is not replicated to other WINS servers | Extinct status information is replicated to other WINS servers |

The different approach for non-owners prevents inconsistencies between primary and secondary WINS servers, because released name entries are not replicated to other servers while extinct name entries are replicated.

If a network resource does not shut down properly, or the WINS client is unable to contact the WINS server during shutdown, the name might not be automatically released.

The released state simplifies WINS name registration for WINS clients that shut down and then restart. If a name is released, the WINS server does not challenge the name when the network resource restarts. If a proper shut down did not occur and the name is not released, the WINS server challenges the name on the next registration request if the IP address is different, as can happen when a DHCP-enabled laptop changes subnets. In such a case, the challenge fails (because the client cannot defend the old IP address) and the registration succeeds, but the registration requires extra processing.

**Name Extinction and Deletion**

When a name entry in the WINS database is marked released and the name is not reregistered, its state progresses to extinct. The name entry is marked with a series of time intervals that determine when the entry changes from released to extinct, and when it can be deleted from the database. The scavenge process checks the time intervals and makes these changes. This process runs automatically, based on the relationship between the configurable renewal and extinction intervals.

**Name extinction**

In general, name entries pass through the states described in the following table.

**Progression of Inactive States**

| Current State of Name Entry | Resulting Status of Name Entry |
|---|---|
| Active state, TTL interval elapsed, WINS server is the owner of the name entry | Released |
| Active state, TTL interval elapsed, WINS server is not the owner of the name entry | Extinct |
| Released state, extinction interval is not elapsed | Released |
| Released state, extinction interval is elapsed | Extinct |
| Extinct state, extinction time-out is not elapsed | Extinct, not ready for deletion |
| Extinct state, extinction time-out is elapsed | Extinct, deleted during scavenging |

Because changing a name entry to the extinct state results in replication, rapid synchronization of WINS databases occurs. The extinct state prevents inconsistencies from lingering. For example, if the primary WINS server of a WINS client is unavailable when the client shuts down, the name release request is directed to the secondary WINS server. If the primary WINS server is available again when the client restarts, the client registers and continues to renew with the primary WINS server, which has not recorded any change in the status of the client, while the secondary WINS server still reflects the released state of the name entry.

**Name deletion**

Extinct name entries that have expired extinction time-out intervals are automatically deleted from the WINS database by the scavenge process. Because changing the status to extinct causes name entries to be replicated, replication typically ensures that the name entries are extinct in all copies of the WINS database before any of the entries are deleted, thus ensuring that the name entries are deleted from all copies of the database.

Under certain abnormal conditions, however, names no longer in use can remain in the database. For example, if an extinct entry is deleted before it is replicated, the active state of the record in the replica databases never changes. This might happen if a replication copy of the database is not reachable during the extinction time-out period.

When an active entry is replicated, it is time-stamped on the pulling server (the WINS server requesting the replication) with the sum of the current time and the verification interval. The default verification interval is 24 days. If the scavenge process determines that there are entries with an expired verification interval, the WINS server sends a name query to the WINS server that owns the name entry to determine whether the version ID is still valid. The following table describes the possible courses of action based on the owning WINS servers response.

**Verification Interval Evaluation**

| Owning WINS Server Response | Action Taken |
|---|---|
| Negative response (name entry is invalid) | Name entry is deleted |

| | |
|---|---|
| Positive response (name entry is valid) | Time stamp is updated |
| No response | Name entry left intact until next verification interval or until an administrator manually scavenges it |

The scavenge process determines which name entries need to be updated or deleted based on ownership, state, and time stamp. This process releases name entries when appropriate, deletes extinct name entries, and verifies whether names are still in use after the verification interval expires. The scavenge process runs periodically based on a scavenging time interval. The scavenging time interval starts at WINS server startup and is half the renew interval (by default, 72 hours). If a WINS server is stopped and restarted before 72 hours have elapsed, the scavenging time interval is reset. If a WINS server is stopped on a daily basis, scavenging does not occur. Scavenging can also be initiated manually.

The following table summarizes the changes the scavenge process makes to name entries in the WINS database.

**Changes to Name Entries During Scavenging**

| Pre-Scavenge State | Post-Scavenge State |
|---|---|
| Owned, active state, renew interval not expired | Unchanged |
| Owned, active state, renew interval expired | Marked released |
| Owned, released state, extinction interval not expired | Unchanged |
| Owned, released state, extinction interval expired | Marked extinct |
| Owned, extinct state, extinction time-out not expired | Unchanged |
| Owned, extinct state, extinction time-out expired | Deleted |
| Replica, extinct state, extinction time-out expired | Deleted |
| Replica, active state, verification interval not expired | Unchanged |
| Replica, active state, verification interval expired, owner challenged, positive response (name entry exists in owner's database) | Time stamp updated (current time plus verification interval) |
| Replica, active state, verification interval expired, owner challenged, negative response (name does not exist in owner's database) | Deleted |
| Replica, active state, verification interval expired, owner challenged, no response received | Unchanged |

During the first scavenging run, all the actions described in the previous table are performed except for deleting name entries. Name entries are not deleted until at least three days after WINS server startup. This delay allows sufficient time for replicating the entries that are ready to be deleted.

In addition to the database cleanup that occurs during the scavenge process, database consistency checks verify name entries in the local database and update name entries as necessary to match the owner entry. Version ID consistency checks verify that owner version IDs are higher than replica version IDs at other WINS servers.

**Name Resolution**

When a computer needs to communicate with another network resource that uses a NetBIOS name, the computer must first acquire the network resource's IP address before communications can be established. Name resolution is the process of obtaining the IP address for a NetBIOS name.

As the following figure illustrates, the WINS client sends the NetBIOS name to a WINS server, the WINS server sends the IP address back to the client, and then the client uses the IP address to connect to the remote computer.

**Name Resolution**




In this figure, ClientA, a WINS client, needs to connect to ClientB. ClientA sends a name query request for the NetBIOS name ClientB to its primary WINS server, WINSA. WINSA finds ClientB in the WINS database and returns the IP address 192.168.1.17 to ClientA, and ClientA then connects to ClientB.

The precise steps for resolving a NetBIOS name depend on the client's NetBIOS node type. B-node clients broadcast the name request to the local subnet. P-node clients use a WINS server only. M-node clients first try broadcasting, and if that is unsuccessful, they use a WINS server. H-node clients first try a WINS server, and if that is unsuccessful, they try broadcasting. This section focuses on the WINS server part of the process.

WINS clients running Windows 2000, Windows XP, or Windows Server 2003 are configured by default to use DNS first to resolve names that are longer than 15 characters or that have periods (.) within the name. If the client is configured to use a DNS server, DNS is used again as a final option if a WINS query is unsuccessful for names that are less than 15 characters and that do not have periods.

For NetBIOS name resolution, a WINS client typically performs the following progression of steps until either the name is resolved or all methods have failed to return an IP address:

1. Client checks its local NetBIOS name cache of remote names. Recently resolved NetBIOS names (up to 16, by default) are held in this cache for 10 minutes.

2. Client sends the NetBIOS name query request to its configured primary WINS server. If the primary WINS server fails to respond, either because it is not available or because it does not have an entry for the name, the client tries to contact all other configured WINS servers in the order they are listed in its configuration.

3. If the client uses H-node, it broadcasts the NetBIOS name query to the local subnet up to three times. (If the client uses M-node, it performs this step before contacting the WINS server (step 2). If the client uses P-node, it does not perform this step.)

4. Client checks the Lmhosts file, if Lmhosts is enabled in the **WINS** tab of the advanced TCP/IP properties.

5. Client converts the NetBIOS name to a host name by removing the last byte and any spaces at the end of the NetBIOS name.

6. Client checks to see if the host name corresponding to the converted NetBIOS name matches the local host name.

7. Client checks the DNS client resolver cache.

8. Client tries a DNS server, if it is configured for one.

If the WINS server resolves the name query request, the WINS server sends a positive name query reply that includes the requested IP address to the client. When the client receives the reply, it uses the address to establish a connection to the desired network resource.

If the WINS server cannot resolve the name query request or if the request is for a unique name that is in a released or extinct state, it sends a negative name query reply to the client. If none of these steps can resolve the name, Windows returns an error.

**Name Queries for Normal Groups**
A normal group name does not have an IP address associated with it, so when a WINS server receives a name query request for a registered normal group name, it cannot return a specific address. Instead, the WINS server returns the limited broadcast address (255.255.255.255). The WINS client then issues a broadcast to its local subnet to resolve the name.

If the request is for a normal group name that is in a released or extinct state, a WINS server sends a positive name query reply and also sends the limited broadcast address (255.255.255.255).

**Name Queries for Special Groups**
A special group can have multiple IP addresses (up to 25) associated with it. When a WINS server receives a name query request for a special group, it returns all the IP addresses for which the TTL has not expired.

### Note

- The WINS server sends as many addresses as it can fit into a UDP message. Because special groups are limited to 25 addresses, all the addresses fit into the UDP message.

**WINS Proxy Name Services**
When a WINS proxy is located on a subnet that includes B-node clients, the WINS proxy can mediate B-node name services requests for registration, release, and resolution by taking the following actions:

- Perform a name lookup during name registration and send a negative name registration reply if it finds the name in its local cache of remote names or in the WINS database.
- Delete the B-node client name from its cache of remote names, when a client releases its name.
- Attempt to resolve a B-node name query request by using information from either its local cache of remote names or from the WINS server.

**Proxy Support for Name Registration Requests**
When a B-node client broadcasts a name registration request, the WINS proxy receives the broadcast and then does a name lookup for names that do not exist on the local subnet. First it uses its local cache of remote NetBIOS names, and if the name is not in the cache, the WINS proxy sends a name query request to the WINS server. If the unique name is found in the WINS database, the proxy sends a negative name registration reply to the B-node client.

**Proxy Support for Name Release**
When a B-node network resource shuts down, WINS proxies simply delete the name from the local name cache in response to the name release request.

**Proxy Support for Name Resolution**
WINS proxies support NetBIOS name resolution for B-node clients. They receive B-node name query broadcasts and respond for names that are not found on the local subnet.

The following figure illustrates how a WINS proxy resolves names for non-WINS clients.

### WINS Proxy Name Resolution




In this example, the WINS proxy (ClientB) performs the following steps to resolve names for the B-node client (ClientA):

1. ClientA broadcasts a NetBIOS name query request to the local subnet (Subnet 1).

2. ClientB receives the broadcast and checks its cache of remote names for the requested NetBIOS name. The WINS proxy always differentiates name queries for local names from remote names. It compares the address of names it resolves to its own address by using the subnet mask, and if the two match, the WINS proxy does not respond to the name query.

3. If ClientB locates the NetBIOS name in its cache, ClientB sends a positive name query reply to ClientA with the IP address.

4. In this example, ClientB cannot locate the NetBIOS name in its cache, so ClientB sends a name query request to the WINS server, WINSA, and enters the requested name in its cache of remote names in a “resolving” state. If ClientB receives another query for the same name before the WINSA responds, ClientB does not query WINSA again.

5. If WINSA cannot resolve the name, it sends a negative name query reply to ClientB, which in turn sends a negative name query reply to ClientA.

6. In this example, WINSA can resolve the name, so it sends a positive name query reply with the IP address to ClientB. ClientB updates the cache entry with the IP address and changes the state to resolved.

   By default, WINS proxies cache remote name mappings resolved by the WINS server for 6 minutes.

7. ClientB sends a positive name query response to ClientA and uses the name mapping in its cache to respond to subsequent NetBIOS name queries broadcast from any B-node clients on the subnet. A WINS proxy sends reply messages only if the response is already in its cache.

### Note

- When a WINS proxy responds to a query for a multihomed client or for a group record that contains a list of IP addresses, it returns only the first address to the B-node client.

**Replication**

Typically, large environments include multiple WINS servers. When multiple WINS servers are used, they can be configured to replicate, or copy, records from one WINS database to other WINS databases. Replication keeps all the WINS databases synchronized and ensures that the names registered on one WINS server are available on all WINS servers. Replication supports the following:

- Reduced name services traffic over WANs

- Redundancy

- Load balancing

- Scalability

For a WINS server to replicate data, it must be configured with at least one other WINS server as its replication partner. Replication partners are WINS servers that are configured to have a replication relationship with each other.

There are three types of replication partners, as described in the following table.

**Types of Replication Partners**

| Type of Partner | Description |
|---|---|
| Push | A WINS server that is configured to notify its pull partner when there are updated entries in the WINS database that are available for replication. Push partners initiate replication. |
| Pull | A WINS server that is configured to request the updated WINS database entries that are ready for replication. |
| Push/pull | A WINS server that is configured to act both as a push partner and a pull partner. This is the default configuration. |

A WINS server can be configured to discover its replication partners automatically (known as autodiscovery) or it can be manually configured for specific replication partners. WINS servers periodically announce their presence on the network. The WINS announcements are sent on a multicast address that is reserved for WINS (224.0.1.24). WINS servers that are configured for automatic discovery listen for the announcements and find other WINS servers on the network. Any WINS server discovered this way is automatically added to the replication partners list as a push/pull partner.

Regardless of whether servers are configured for one-way replication (push only or pull only) or for two-way replication (push/pull), the primary and secondary WINS servers for a client must have a push and pull relationship to each other to ensure that WINS functions properly if the primary server fails.

The following figure illustrates how replication works. Two WINS servers, WINSA and WINSB, are configured to be replication partners.

**WINS Database Replication**




In this figure, a WINS client (ClientA) on Subnet 1 registers its name with its primary WINS server, WINSA, and its name is updated in the WINSA database. Another WINS client (ClientB) on Subnet 3 registers its name with its primary WINS server, WINSB, and its name is updated in the WINSB database. If, after replication, ClientA sends a name query request to WINSA to find an IP address for ClientB, ClientB will be in the WINSA database.

WINS replication is always incremental, which means that only changed records, rather than all records, are replicated.

**Persistent Connections**

By default WINS uses persistent connections between replication partners. When they use persistent connections, WINS servers do not disconnect from their replication partners each time replication is completed and, therefore, avoid the overhead and delays of establishing a new connection each time replication begins. Typically when WINS servers are interconnected through high-speed LAN links, it is preferable to keep connections open.

Persistent connections increase the speed of replication because a server can immediately send updates to its partners, making replication more consistent. The bandwidth used by persistent connections is minimal because the connection is usually idle.

Persistent connections can be used with push partners and pull partners, either with a specific partner or with all the replication partners of a WINS server.

**Version ID**
Every name entry in a WINS database has a version ID field. This field is incremented whenever a change that is to be replicated takes place. That is, the version ID for a name is incremented whenever that name is registered, changed from a released to an active state, or changed to an extinct state.

The value in version ID reflects the total number of replicable changes that occurred to owned entries for that particular WINS server. The value is incremented by hexadecimal 1 every time a replicable change occurs. For example, if WINSA is the owner for ClientA, which has a version ID of 4B3, and for ClientB, which has a version ID of 4C2, ClientB was the fifteenth replicated change in the WINSA database after the ClientA change.

The number of changes in version ID is one way that WINS determines when to trigger replication between two servers.

**Push Partners**
A push partner is a WINS server that is configured to push or notify other WINS servers when it has database updates that are ready to be replicated. Push partners initiate replication. They notify replication partners when it is time for the partners to pull updates, but they do not send any updates.

A push partner can be configured to use any of the following events to trigger the push notification:

- WINS server startup.
- An IP address change in a name-to-address mapping in the database.
- A specified number of version ID changes occur. This number is the number of local changes and does not include changes pulled from other partners. WINS can use this number as a threshold for determining how many changes must be made to the database before it starts push replication to its partners. This number can apply globally to all partners to which it pushes or individually to specific partners.

When persistent connections are used for push replication and the number of version ID changes for triggering replication is set to 0 (the default), the push partner sends a push trigger notifying its replication partners every time an update occurs. In this situation, push replication is not dependent on the number of times version ID changes but occurs at every update. (An update is defined here as a change that increments the highest version ID in the local WINS database for records the server owns.) If the number of version ID changes is greater than 0, the frequency of push notifications is reduced.

If persistent connections are not used for push replication, push replication is dependent on the number of times the version ID changes. If this number is set to zero, push replication never occurs, except if it is set to occur at service startup or when IP addresses change. If the number of version ID changes is greater than 0 when persistent connections are not used, the frequency of push notifications is increased.

For example, suppose that the highest version ID for server WINSA is 1C4. WINSA is also set to push to its replication partners when there have been two updates in its database. Suppose also that the following four sets of database changes are then made to the WINS database at WINSA:

1. ClientA refreshes its previously registered name but keeps the same IP address.
2. ClientB registers its name, creating a new dynamic name-to-address mapping in the database.
3. ClientC reregisters a previously released name and changes its IP address.
4. A pull replication occurs at a specified interval with another WINS server.

When these changes are entered into the database, the highest version ID for WINSA is 1C6 because two of the changes — update 2 (registering a new name) and update 3 (changing an IP address) — cause the version ID for that server to be incremented.

After these changes occur, WINSA then sends a push notification to its replication partners.

**Pull Partners**
A pull partner is a WINS server that is configured to pull, or request replication of, updated WINS database entries from other

WINS servers (those configured as its push partners) at a configured interval. Push replication can occur in response to a specified event, in response to a manual request by an administrator, or in response to a push trigger.

A pull partner pulls or requests entries that have a higher version ID than the last entry it received from its configured replication partner.

A pull partner can be configured to use either of the following events to trigger the request for replication:

- WINS server startup
- A user-specified amount of elapsed time, which can apply globally to all the partners from which it pulls or individually to specific partners

When the specified time has elapsed, the pull partner sends a replication request to its configured partners. The partners respond with incremental updates, if any have occurred since the last replication.

Database updates are not physically replicated until a pull partner sends a replication request in response to the notification sent by a push partner. Therefore, pull replication can be viewed as the single process used to replicate data between WINS databases.

The IP address-to-owner ID table in the WINS database is used during the replication process. In addition, all WINS servers maintain an internal push partner owner-to-version mapping table. During WINS initialization, a WINS server scans the name-to-IP address mapping table to find the highest version ID for each owner in the IP address-to-owner ID table. The WINS server uses the information to build the push partner owner-to-version table, which maps the owner IDs to their highest version IDs. The WINS server maintains this table in memory and never commits it to the database. WINS uses the push partner owner-to-version table to determine the state of the database at any point before or after replication.

During replication, the WINS server updates the push partner owner-to-version table with the highest version IDs sent by its push partners. The WINS server then scans the version ID field in the name-to-IP address table to find the highest version ID for each owner and compares the results with the numbers in the push partner owner-to-version table to determine which push partner has the latest data for each owner. The WINS server also updates its IP address-to-owner ID table with additional owners obtained from its push partners.

The following figure illustrates how WINS manages replication. In this example, four WINS servers (WINSA, WINSB, WINSC, and WINSD) are all configured as push/pull partners for a central server, WINSH.

**Example of Pull Replication**




Before replication begins in this example, the owner-to-version table at WINSA contains the highest version IDs for all the owners in its database. Assume that these values are the values shown in the following table.

**WINSA Owner-to-Version Table**

| Owner | Name | Highest Version ID |
|---|---|---|

| | | |
|---|---|---|
| 192.168.1.5 | WINSA | 1F |
| 192.168.2.5 | WINSH | 125 |
| 192.168.3.5 | WINSB | 31 |
| 192.168.4.5 | WINSC | 100 |
| 192.168.5.5 | WINSD | 200 |

In reality, pull replication can occur dynamically between WINSH and WINSA, WINSB, WINSC, and WINSD at various times, so the table might not remain static for long. To simplify the details of an individual pull replication cycle, however, the example below focuses on what might happen when WINSA replicates its database with WINSH.

In this example, the following changes occur in the replicated WINS network:

1. Assume that first WINSB, WINSC, and WINSD each replicate with WINSH, based on the highest version ID.

    a. WINSB has added or updated one new record since its last replication cycle, which is replicated at WINSH.

    b. WINSC has the same version ID and does not have any records to be replicated at WINSH.

    c. WINSD has updated 20 records since its last replication cycle, which are replicated at WINSH.

2. Assume that after WINSB, WINSC, and WINSD replicate with WINSH, WINSH sends a push trigger to WINSA, initiating replication. WINSA then pulls the updated owner-to-version information from WINSH.

3. WINSA responds by pulling entries from WINSH.

    WINSA determines which updates it needs to pull based on the difference between its locally-held highest version ID for each owner and the corresponding highest version IDs provided by WINSH. WINSA specifies the range of records to be pulled and replicated by comparing version IDs. In this example, WINSA pulls updates that occurred on WINSH with a version ID greater than 125, updates that occurred on WINSB with a version ID greater than 31, and updates that occurred on WINSD with a version ID greater than 200.

If another pull replication to the same partner is already queued, WINS treats the later request as further down in the queue and discards it in favor of keeping the one that is higher in the queue.

**Accepting and Blocking Replication Partners**
WINS servers can be configured to accept updates only from specific replication partners (known as persona grata) or to block updates from specific replication partners (known as persona non grata) during replication. A WINS server can be configured for only one of these options. These options apply either during push-pull replication or during pull replication.

**Accepting replication partners**
When a WINS server is configured to accept specific replication partners, only updates from WINS servers that are included in the list are replicated. For example, accepting replication from specific partners might be useful when WINS servers are configured in a hub-and-spokes pattern, where a central server is the hub and other distributed servers are the spokes. By including only the IP addresses for the hubs in the list, you can replicate updates only between the hubs, without including the spokes.

If a WINS server is configured to accept updates from specific partners and no list of accepted WINS servers is provided, no updates are replicated. Accepting updates only from specific partners is a new WINS feature in Windows Server 2003.

**Blocking replication partners**
When a WINS server is configured to block specific replication partners, only updates from WINS servers that are not in the list are replicated. For example, blocking replication partners is useful when a WINS server is removed from service. After a WINS server is removed from service, name entries it owns can continue to be replicated to other servers. Replication of these records can occur for one of the following reasons:

- Static entries that were originally added to the now inactive WINS server replicate indefinitely to the other active WINS servers in the network until the entries are deleted manually.

- Dynamic entries that were originally added to the now inactive WINS server are not immediately deleted from the WINS database. WINS must check with the owner server before it deletes dynamic entries. If a database entry cannot be verified as expired because the owner server is no longer in use, other WINS server's retain the entry until it can be verified.

The blocking replication partners option can prevent replication of name entries from inactive WINS servers.

If a WINS server is configured to block updates and no list of blocked WINS servers is provided, updates are replicated for all WINS servers.

**Owner ID and Time Stamp During Replication**
WINS replicates only name entries in the active and extinct states. During replication, the fields for these database entries are copied, with the exception of owner ID and time stamp.

The owner ID is obtained from the local owner ID-to-version ID mapping table because the value that is used locally to represent a specific WINS server differs from server to server. For example, WINSD might be represented as owner ID 2 on WINSB and as owner ID 3 on WINSA.

When a name entry is replicated, it is given a new time stamp. Active name entries are time-stamped with the sum of the local current time and the verification interval. Extinct name entries are time-stamped with the sum of the local current time and the extinction time-out.

**Name Conflicts During Replication**
Although name conflicts are usually handled at the time of name registration, it is possible for the same name to be registered at two different WINS servers. This situation might happen if a WINS client registers the same name at a second WINS server before the updates from the first WINS server replicate to the second. In this case, WINS resolves the conflict at replication time.

Conflicts at replication can be between two unique name entries, between a unique name entry and a group name entry, or between two group name entries.

**Conflicts between unique name entries**
When two unique names conflict, WINS resolves the conflict based on three factors:

- State of the entries (active, released, or extinct)
- Ownership of the entries
- Same or different IP addresses

This section describes how WINS resolves conflicts between unique names.

**Conflicts between two replicas**
When two replicas conflict, the new replica overwrites the existing replica, regardless of whether the addresses match. The only exception to this rule is if the existing replica is active and the new replica is extinct. If the existing replica is active, the new replica is extinct, and different WINS servers own the replicas, then the new replica does not overwrite the existing replica. If the existing replica is active, the new replica is extinct, and the same WINS server owns both, the new extinct replica overwrites the existing active replica.

**Conflicts between owned entries and replicas with the same IP address**
The replica replaces the owned entry unless the owned entry is active and the replica is extinct. In this case, WINS increments the version ID of the owned entry so that it replicates during the next iteration.

**Conflicts between owned entries and replicas with different IP addresses**
The replica replaces the owned entry unless the owned entry is active. If the owned entry is active and the replica is extinct, WINS increments the version ID of the owned entry so that it replicates during the next iteration. If both the owned entry and the replica are active, the server receiving the replica challenges the client that registered the owned entry to determine whether the client still uses the name. If it does still use the name, WINS sends the client designated in the replica a name conflict demand. This message puts the client in a conflict state, which forces the client to put the name in a conflict state. A name in a conflict state is marked and the name is no longer used.

**Conflicts between unique entries and group entries**
When a unique entry and a group entry conflict, WINS keeps the group entry. If the WINS server owns the unique entry and the entry is not in a released or extinct state, the WINS server requests the client named in the unique entry to release the

name.

**Conflicts between two special group entries**
When there is a conflict between two special groups, the new replica replaces the existing entry unless the existing entry is active. If the existing entry is active, WINS increments its version ID so that it replicates at the next iteration. If the new replica is also active, the WINS server updates the member list of the existing entry with any new members from the replica. If the list of active members grows to more than 25, the extra members are not added but are dropped silently.

**Conflicts involving multihomed entries**
If one of the name entries involved in a name conflict is for a multihomed node, the conflict is resolved in one of the following ways:

- If a multihomed replica conflicts with an extinct or released entry in the database, WINS replaces the existing entry in the database with the replica, unless the existing entry is a normal group and is in a released state. This is the same as the other scenarios in which a single-address entry conflicts with a released normal group entry.

- If a multihomed replica that is extinct conflicts with an active existing entry and both entries are owned by the same server, WINS replaces the active existing entry. If the active existing entry is a replica owned by a different server, WINS does not replace the existing entry. If the active existing entry is owned by the local WINS server and is a unique entry, WINS increments the version ID of the existing record to force replication.

- If an active multihomed replica conflicts with an active, unique, multihomed replica in the local database with the same owner, WINS replaces the existing entry. If the owner is different, WINS does not replace it. If the entry in the database is owned by the local WINS server, and if the members of the record (a single member in the case of a unique record) are a subset of the members of the replica, WINS changes the time stamp of the existing entry and increments its version ID to force replication. If the members of the replica are not a subset, the WINS server challenges the addresses in the existing entry. If there are no responses to the challenges, WINS replaces the existing entry. If there is at least one response to the challenges, WINS notifies the client to release the name from all addresses before it replaces the existing entry with the replica.

- If a multihomed replica conflicts with an existing active group entry, WINS increments the version ID of the existing entry to force replication.

- If a single-address replica conflicts with an existing inactive multihomed entry, WINS replaces the existing entry with the replica. If the replica conflicts with an existing active multihomed entry and both entries are owned by the same owner, WINS replaces the existing entry with the replica. If the existing multihomed entry is a replica owned by a different server, WINS does not replace it. If, however, the existing multihomed entry is owned by the local WINS server and the new replica is a unique record, then WINS challenges the clients that use the addresses in the multihomed record. If there are no responses to the challenges, WINS replaces the existing entry. If there is at least one response to the challenges, WINS sends requests for the name to the addresses in the existing entry and then updates the existing entry. The address in the unique replica, if present in the member list, is ignored in this situation.

**WINS Convergence Time**
There is a latency period between the time a database entry is updated at one WINS server and the time the entry is replicated to all other WINS servers in the network. This latency period is known as the convergence time for the WINS system. Convergence time can be computed as the sum of the configured replication periods from one server to the next over the longest path. The following figure illustrates a replication configuration and indicates the replication periods between nodes. This figure is not intended to be a realistic configuration, but merely to illustrate a longest replication path.

**WINS Convergence**



In this figure, the longest replication path is from WINSB to WINSA to WINSC. A name registration request from one of the WINS clients causes the WINSB database to be updated. Changes to the WINSB database are replicated to WINSA every hour, and changes are replicated from WINSA to WINSC every two hours. Therefore, the convergence time for the name registration request in this example is three hours.

The convergence time can vary among different types of name services. For example, name registration requests that occur at one server replicate to the server's replication partner at the next replication event. Name release requests, however, are not replicated. NetBIOS names are released when a network resource shuts down. Because it is common for these names to be reused with the same IP address when the network resource is restarted, replicating the name releases unnecessarily increases the network load of replication. Since these name service requests are not replicated, however, a NetBIOS name can be in the release state in one WINS database and still be in the active state in all other WINS databases.

In addition, some changes can take place and not be reflected in any of the WINS databases. For example, if a WINS client shuts down improperly, as in the case of an unexpected power outage, it is unable to send a name release request to its primary WINS server, and the registered names are not marked as released in the WINS database.

**Replication and Name Services Summary**

The following table describes typical events that might occur in a NetBIOS network, the name service operations they invoke, the resulting WINS database updates, and whether or not the changed name entry is replicated. These scenarios assume that the WINS server handling the name service request is the owner for the name entry and that each name service request results in a positive response.

**Typical WINS Scenarios**

| Event | Name Service Request / WINS Process | WINS Database Updates | Replication |
|---|---|---|---|
| Computer starts up | Name registration | • Name added to database<br>• State set to active<br>• Owner ID set<br>• Version ID incremented<br>• Renew interval set | Yes |
| Computer shuts down | Release | • State set to released<br>• Owner ID set<br>• Extinction interval set | No |
| Computer restarts, same name and IP address | Name registration | • State set to active<br>• Owner ID set<br>• Version ID incremented<br>• Renew interval set | Yes |

| | | | |
|---|---|---|---|
| Computer stays up, renew interval nears | Name refresh | • Name added to database<br>• State set to active<br>• Owner ID set<br>• Version ID incremented<br>• Renew interval set | Yes |
| Computer shuts down and restarts on different subnet with different primary WINS server, same name but different IP address | Name registration, which is challenged | • State remains active (release from shutdown not replicated)<br>• Owner ID set<br>• Version ID incremented<br>• Renew interval set | Yes |
| Computer shuts down and stays down, extinction interval expires | Database scavenge | • State set to extinct<br>• Version ID incremented<br>• Extinction time-out set | Yes |
| Computer remains shut down, extinction time-out expires | Database scavenge | • Entry deleted during scavenging | No |

**Related Information**

The following resources contain additional information that is relevant to this section.

- "How DHCP Technology Works"
- "How DNS Works"

For more information about WINS, see the following RFCs in the IETF RFC Database.

- RFC 1001, "Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods."
- RFC 1002, "Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Detailed Specifications."

# WINS Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- WINS Tools
- WINS Registry Settings
- DHCP Options for WINS
- Network Ports Used by WINS


The "WINS Tools" section describes tools that are useful for configuring, managing, and monitoring WINS servers and databases. Some of the tools are specific to WINS, while others are not.

The "WINS Registry Settings" section describes registry entries that are useful in specific circumstances and need to be set directly by using Regedit.exe. Registry entries that are not described in this section are modified by the tools described in the "WINS Tools" section.

The "DHCP Options for WINS" section describes the DHCP options that relate specifically to WINS.

The "Network Ports Used by WINS" section lists the ports that WINS uses for specific purposes.

**WINS Tools**
The following tools are associated with WINS.

**Jetpack.exe: Jetpack**
**Category**
The Jetpack tool is included as part of the operating system.

**Version compatibility**
Jetpack runs on Windows NT 4.0, Windows 2000, and the Windows Server 2003 family.

This tool compacts the WINS database while the database is offline. (Database compaction also runs as an automatic background process during idle time after a database update while the server is online. Dynamic compaction, however, does not reclaim as much space as offline compaction.)

Jetpack copies database information to a temporary database, deletes the original WINS database, and then renames the temporary database to the original name.

**Nbtstat.exe: Nbtstat**
**Category**
The **Nbtstat** tool is included as part of the operating system.

**Version compatibility**
**Nbtstat** runs on, and can target, computers running Windows NT 4.0, Windows 2000, Windows XP, and the Windows Server 2003 family.

This tool displays NetBIOS over TCP/IP (NetBT) protocol statistics, NetBIOS name tables for both the local computer and remote computers, and the NetBIOS name cache. It can be used to refresh the NetBIOS name cache and to release and reregister the names registered with WINS during troubleshooting.

To find more information about Nbtstat.exe, see "Command Line References" in the "Tools and Settings Collection."

**Netsh.exe: Netsh**
**Category**
The **Netsh** commands for WINS are included as part of the operating system.

**Version compatibility**
The **Netsh** commands for WINS run on, and can target, computers running the Windows Server 2003 family and Windows 2000 Server. There are some functional differences between the commands on Windows Server 2003 and Windows 2000

Server.

This tool provides a command-line alternative to the WINS Microsoft Management Console (MMC) snap-in for administering WINS servers. Command-line management can be advantageous in the following ways:

- Commands run in interactive mode can more efficiently manage WINS servers across slow wide area network (WAN) links.
- Commands run in batch mode can automate recurring tasks for managing many servers.

The **Netsh** extension that supports WINS server administration is available from the **wins>** prompt, which is available from the **netsh>** prompt.

The **server** command is available at the **netsh wins>** prompt for switching to a specific server context. The following table describes the commands available at the **netsh wins server>** prompt for administering a WINS server.

To find more information about **Netsh** commands for WINS, see "Command Line References" in the "Tools and Settings Collection."

**Network Connections**
**Category**
Network Connections is included as part of the operating system.

**Version compatibility**
Network Connections runs on Windows 2000, Windows XP, and the Windows Server 2003 family.

Network Connections is used to configure WINS clients for the following:

- Manually configure primary and secondary WINS servers (if DHCP option 44 is not used to configure the IP addresses of WINS servers)
- Enable Lmhosts lookup
- Enable or disable NetBT

These configuration options are available from the **WINS** tab in the advanced Internet Protocol (TCP/IP) properties for a connection.

**Wins.dll: WINS Replication Network Monitor Parser**
**Category**
WINS Replication Network Monitor Parser is included with Windows Server 2003 Resource Kit tools.

**Version compatibility**
WINS Replication Network Monitor Parser runs on Windows 2000 Server and the Windows Server 2003 family.

The WINS Replication Network Monitor Parser tool is a DLL that parses WINS data from within Network Monitor, a feature of Microsoft Windows. This tool is useful as a way to view and troubleshoot network issues related to WINS replication.

To find more information about WINS Replication Network Monitor Parser, see "Resource Kit Tools" in the "Tools and Settings Collection."

**Winsmgmt.msc: WINS Microsoft Management Console Snap-in**
**Category**
The WINS MMC snap-in is included with the Windows Server 2003 family and Windows 2000 Server, and it is installed when the WINS Server service is installed. It is also included in the Windows Server 2003 Administration Tools Pack.

**Version compatibility**
The WINS snap-in runs on, and can target, computers running Windows 2000, Windows XP, and the Windows Server 2003 family.

The WINS snap-in configures and manages WINS servers and databases by supporting the functions described in the following table.

| Function | Options |
|---|---|

| | |
|---|---|
| Configures server properties | • Frequency for updating server statistics<br>• Database backup settings<br>• Time intervals for aging and deleting name entries in the database<br>• Database consistency-checking settings<br>• Logging settings<br>• Burst handling settings<br>• Database path and name<br>• Starting version ID value<br>• LAN Manager compatibility |
| Configures database replication properties | • General properties, such as which WINS servers are used for replication and how static mappings are handled<br>• Push partner properties, such as when push replication is initiated and whether persistent connections are used<br>• Pull partner properties, such as when pull replication is initiated, how many attempts are made, and whether persistent connections are used<br>• Blocking or accepting replication only from specific partners<br>• Automatic discovery of replication partners |
| Manages WINS servers | • Add and delete WINS servers<br>• Add, remove, and configure replication partners<br>• Start and stop WINS service<br>• Display server statistics |

| | |
|---|---|
| Manages WINS databases | • Manually initiate replication<br>• Display database records<br>• Manually initiate the scavenge (database cleanup) process<br>• Initiate backup or restore<br>• Manually delete records<br>• Delete owners<br>• Initiate database consistency checks, which verify local database records against records in the owner databases<br>• Initiate version ID consistency checks, which verify that owner version IDs are higher than replica version IDs at other WINS servers<br>• Create static name mappings<br>• Import Lmhosts files<br>• Verify name records |

**WINS Registry Entries**
The following registry entries are associated with WINS and might need to be modified to customize a WINS client or WINS server configuration. Additional registry entries associated with WINS are modified by options in the WINS snap-in and **Netsh** commands. These tools should be used to update registry entries whenever possible.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The registry entries associated with WINS are under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters. These registry entries are used by NetBT and are used to configure WINS clients.

**BcastNameQueryCount**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\

**Version**
The BcastNameQueryCount entry is included in Windows NT 4.0, Windows 2000, Windows XP, and the Windows Server 2003 family.

The BcastNameQueryCount entry specifies the number of times that NetBT broadcasts a unique name query if no response is received.

For more information about BcastNameQueryCount, see "Registry Reference" in the "Tools and Settings Collection."

**BcastQueryTimeout**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\

**Version**
The BcastQueryTimeout entry is included in Windows NT 4.0, Windows 2000, Windows XP, and the Windows Server 2003 family.

The BcastQueryTimeout entry specifies the time in milliseconds between successive broadcast queries for a unique name.

For more information about BcastQueryTimeout, see "Registry Reference" in the "Tools and Settings Collection."

**CacheTimeout**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\

**Version**
The CacheTimeout entry is included in Windows NT 4.0, Windows 2000, Windows XP, and the Windows Server 2003 family.

The CacheTimeout entry specifies the time in milliseconds that NetBT caches NetBIOS names.

For more information about CacheTimeout, see "Registry Reference" in the "Tools and Settings Collection."

**EnableProxy**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\

**Version**
The EnableProxy entry is included in Windows NT 4.0, Windows 2000, Windows XP, and the Windows Server 2003 family.

The EnableProxy entry configures a computer to act as a WINS proxy for B-node clients on the local subnet.

**NodeType**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\

**Version**
The NodeType entry is included in Windows NT 4.0, Windows 2000, Windows XP, and the Windows Server 2003 family.

The NodeType entry specifies the method NetBT uses to register and resolve NetBIOS names. The default NodeType is B-node. Configuring a client to use a WINS server sets this entry to H-node. The registry entry can be modified with Regedit.exe for M-node or P-node.

For more information about NodeType, see "Registry Reference" in the "Tools and Settings Collection."

**Size/Small/Medium/Large**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\

**Version**
The Size/Small/Medium/Large entry is included in Windows NT 4.0, Windows 2000, Windows XP, and the Windows Server 2003 family.

The Size/Small/Medium/Large entry specifies the maximum number of entries that can be held in the NetBIOS name cache of recently resolved NetBIOS names.

**ScopeID**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\

**Version**
The ScopeID entry is included in Windows NT 4.0, Windows 2000, Windows XP, and the Windows Server 2003 family.

The ScopeID entry sets the NetBIOS scope, which is an optional, additional qualifier for NetBIOS names. Use of the scope identifier is discouraged.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WINS**
The registry entries under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WINS\Parameters are used to configure WINS servers. Most of these registry entries are set by using the WINS snap-in or the WINS **Netsh** commands. The registry entries described in this section need to be modified by using Regedit.exe.

**DataFiles**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WINS\Parameters\

**Version**
The DataFiles entry is included in Windows NT 4.0, Windows 2000, and the Windows Server 2003 family.

The DataFiles subkey contains a list of files that contain static name records. WINS uses these files to initialize and reinitialize the local database. This subkey does not exist in the registry by default. This subkey and its entries are needed if static name records are used.

**InitTimePause**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WINS\Parameters\

**Version**
The InitTimePause entry is included in Windows NT 4.0, Windows 2000, and the Windows Server 2003 family.

The InitTimePause registry entry determines whether WINS starts in a paused state and remains in that state until the first replication is complete. While in a paused state, WINS does not accept any name service requests.

This registry entry can be used in conjunction with the WINS snap-in option for triggering push or pull replication at system startup to restore a WINS database by using replication.

**DHCP Options for WINS**
When Dynamic Host Configuration Protocol (DHCP) is used to dynamically assign IP addresses in a WINS environment, you can use DHCP options to define the primary and secondary WINS servers and the name resolution method used by DHCP clients. The following table lists the DHCP options that are associated with WINS.

### DHCP Options Associated with WINS

| DHCP Option | Description |
|---|---|
| 44 | Identifies the WINS servers to be used by DHCP clients. |
| 46 | Defines the NetBIOS Node Type: B-node, P-node, M-node, or H-node. This option typically does not need to be set because WINS clients are automatically set to H-node when they are configured with WINS servers. |
| 47 | Identifies the NetBIOS scope ID to be used by the DHCP client. Using NetBIOS scope IDs is not recommended. |

For more information about DHCP options, see "DHCP Tools and Settings."

**Network Ports Used by WINS**
The ports listed in the following table are used by specific WINS services.

### WINS Port Assignments

| Service Name | UDP | TCP |
|---|---|---|
| NetBIOS Name Service | 137 (client and server) | |
| WINS replication | | 42 |
| WINS replication partner automatic discovery | 42 | |

# Network Configuration Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Configuration**
Network configuration is vital to the proper operation of any network infrastructure. Administrators must be able to configure IP addresses as well as other configuration information, such as options configured by using Dynamic Host Configuration Protocol (DHCP) or the Windows Time service. These services offer administrators scalable solutions to configure client computers on the network.

**Network Configuration Components**
Microsoft Windows Server 2003 provides DHCP and the Windows Time service for network configuration.

**DHCP**
DHCP is a client/server protocol that automatically provides an IP host with its IP address and other related configuration information such as the subnet mask and default gateway. RFC 2131 and RFC 2132 define DHCP as an Internet Engineering Task Force (IETF) standard based on the Bootstrap Protocol (BOOTP), with which it shares many implementation details. DHCP allows hosts to obtain all necessary TCP/IP configuration information from a DHCP server.

All Microsoft Windows Server 2003 products (including Standard Edition, Enterprise Edition, Web Edition, and Datacenter Edition) include a DHCP Server service, which is an optional networking component. All Windows clients include the DHCP Client as part of TCP/IP, including Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0, Windows Millennium Edition (Windows Me), and Windows 98.

DHCP simplifies the administrative management of IP address configuration by automating this process for network clients. The DHCP servers, which are computers running the DHCP Server service, allocate IP addresses and related TCP/IP configuration settings to DHCP-enabled clients on the network. DHCP relay agents help forward DHCP messages between DHCP clients and DHCP servers.

**Windows Time Service**
The Windows Time service (W32Time) in Windows Server 2003 synchronizes the date and time for all computers running on a Windows Server 2003 network. Time synchronization is critical for the proper operation of many Windows services and line-of-business applications. The Windows Time service uses the Network Time Protocol (NTP) to synchronize computer clocks on the network so that an accurate clock value, or time stamp, can be assigned to network validation and resource access requests. The service integrates NTP and time providers, making it a reliable and scalable time service for enterprise administrators.

**Network Configuration Scenarios**
Every device on a TCP/IP-based network must have a unique unicast IP address to access the network and its resources. Without DHCP, IP addresses must be configured manually on each computer on the network, changes must be tracked so that unused addresses can be reused by other computers, and computers must be manually reconfigured each time they move between subnets.

DHCP allows this entire process to be automated and centrally managed. The DHCP server maintains a pool of IP addresses and leases an address to any DHCP-enabled client when it starts on the network. Because the IP addresses are dynamic (leased) rather than static (permanently assigned), addresses no longer in use are automatically returned to the pool for reallocation.

DHCP also allows the network administrator to centrally manage additional configuration settings by specifying DHCP options, which pass additional information to the DHCP client when it obtains a lease.

The Windows Time service is essential to the successful operation of Kerberos authentication. It is therefore essential for authentication based on Active Directory. Any application that uses Kerberos, including most security services, relies on time synchronization between the computers that are participating in the authentication request. Active Directory domain controllers must also have synchronized clocks to help ensure accurate data replication.

Although Windows Time service was originally designed to provide clock synchronization for Kerberos authentication, many current applications use timestamps to ensure transactional consistency, to record the time of important events, and other time-sensitive information.

# DHCP Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP Technical Reference**
Microsoft Windows Server 2003 includes Dynamic Host Configuration Protocol (DHCP), a TCP/IP standard that enables the centralized management of IP addresses and other optional configuration information for client computers on a network.

DHCP simplifies IP address configuration by automating the process for clients on a network. DHCP servers, allocate IP addresses and related TCP/IP configuration settings to DHCP-enabled clients on the network. DHCP relay agents are used to forward initial DHCP messages between DHCP clients and DHCP servers if the clients and servers are not on the same IP broadcast subnet.

The DHCP Server service for Windows Server 2003 is based on Internet Engineering Task Force (IETF) standards. DHCP specifications are defined in Requests for Comments (RFCs) published by the Dynamic Host Configuration working group of the IETF. The following RFCs specify the core DHCP standards that Windows Server 2003 DHCP supports:

- RFC 2131: Dynamic Host Configuration Protocol
- RFC 2132: DHCP Options and BOOTP Vendor Extensions

Before you can understand DHCP you must have a working knowledge of TCP/IP networks. Specifically, you must understand how subnets work within a TCP/IP network. You should also review and understand the terms listed in "What Is DHCP?."

**In this subject**


- What Is DHCP?
- How DHCP Technology Works
- DHCP Tools and Settings
- DNS registration changes for Windows Server 2003 based DHCP Servers

# What Is DHCP?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What is DHCP?**

Dynamic Host Configuration Protocol (DHCP) is a client/server protocol that automatically provides an Internet Protocol (IP) host with its IP address and other related configuration information such as the subnet mask and default gateway. RFCs 2131 and 2132 define DHCP as an Internet Engineering Task Force (IETF) standard based on Bootstrap Protocol (BOOTP), a protocol with which DHCP shares many implementation details. DHCP allows hosts to obtain necessary TCP/IP configuration information from a DHCP server.

The Microsoft Windows Server 2003 operating system includes a DHCP Server service, which is an optional networking component. All Windows-based clients include the DHCP client as part of TCP/IP, including Windows Server 2003, Microsoft Windows XP, Windows 2000, Windows NT 4.0, Windows Millennium Edition (Windows Me), and Windows 98.

**Note**

- It is necessary to have an understanding of basic TCP/IP concepts, including a working knowledge of subnets before you can fully understand DHCP. For more information about TCP/IP, see "TCP/IP Technical Reference."

**Benefits of DHCP**

In Windows Server 2003, the DHCP Server service provides the following benefits:

- **Reliable IP address configuration.** DHCP minimizes configuration errors caused by manual IP address configuration, such as typographical errors, or address conflicts caused by the assignment of an IP address to more than one computer at the same time.

- **Reduced network administration.** DHCP includes the following features to reduce network administration:

    - Centralized and automated TCP/IP configuration.

    - The ability to define TCP/IP configurations from a central location.

    - The ability to assign a full range of additional TCP/IP configuration values by means of DHCP options.

    - The efficient handling of IP address changes for clients that must be updated frequently, such as those for portable computers that move to different locations on a wireless network.

    - The forwarding of initial DHCP messages by using a DHCP relay agent, thus eliminating the need to have a DHCP server on every subnet.

**Why use DHCP**

Every device on a TCP/IP-based network must have a unique unicast IP address to access the network and its resources. Without DHCP, IP addresses must be configured manually for new computers or computers that are moved from one subnet to another, and manually reclaimed for computers that are removed from the network.

DHCP enables this entire process to be automated and managed centrally. The DHCP server maintains a pool of IP addresses and leases an address to any DHCP-enabled client when it starts up on the network. Because the IP addresses are dynamic (leased) rather than static (permanently assigned), addresses no longer in use are automatically returned to the pool for reallocation.

The network administrator establishes DHCP servers that maintain TCP/IP configuration information and provide address configuration to DHCP-enabled clients in the form of a lease offer. The DHCP server stores the configuration information in a database, which includes:

- Valid TCP/IP configuration parameters for all clients on the network.

- Valid IP addresses, maintained in a pool for assignment to clients, as well as excluded addresses.
- Reserved IP addresses associated with particular DHCP clients. This allows consistent assignment of a single IP address to a single DHCP client.
- The lease duration, or the length of time for which the IP address can be used before a lease renewal is required.

A DHCP-enabled client, upon accepting a lease offer, receives:

- A valid IP address for the subnet to which it is connecting.
- Requested DHCP options, which are additional parameters that a DHCP server is configured to assign to clients. Some examples of DHCP options are Router (default gateway), DNS Servers, and DNS Domain Name. For a full list of DHCP options, see "DHCP Tools and Settings."

**Terms and Definitions**

The following table lists common terms associated with DHCP.

**DHCP Terms and Definitions**

| Term | Definition |
|---|---|
| DHCP server | A computer running the DHCP Server service that holds information about available IP addresses and related configuration information as defined by the DHCP administrator and responds to requests from DHCP clients. |
| DHCP client | A computer that gets its IP configuration information by using DHCP. |
| Scope | A range of IP addresses that are available to be leased to DHCP clients by the DHCP Server service. |
| Subnetting | The process of partitioning a single TCP/IP network into a number of separate network segments called subnets. |
| DHCP option | Configuration parameters that a DHCP server assigns to clients. Most DHCP options are predefined, based on optional parameters defined in Request for Comments (RFC) 2132, although extended options can be added by vendors or users. |
| Option class | An additional set of options that can be provided to a DHCP client based on its computer class membership. The administrator can use option classes to submanage option values provided to DHCP clients. There are two types of options classes supported by a DHCP server running Windows Server 2003: vendor classes and user classes. |
| Lease | The length of time for which a DHCP client can use a DHCP-assigned IP address configuration. |
| Reservation | A specific IP address within a scope permanently set aside for leased use by a specific DHCP client. Client reservations are made in the DHCP database using the DHCP snap-in and are based on a unique client device identifier for each reserved entry. |
| Exclusion/exclusion range | One or more IP addresses within a DHCP scope that are not allocated by the DHCP Server service. Exclusions ensure that the specified IP addresses will not be offered to clients by the DHCP server as part of the general address pool. |
| DHCP relay agent | Either a host or an IP router that listens for DHCP client messages being broadcast on a subnet and then forwards those DHCP messages directly to a configured DHCP server. The DHCP server sends DHCP response messages directly back to the DHCP relay agent, which then forwards them to the DHCP client. The DHCP administrator uses DHCP relay agents to centralize DHCP servers, avoiding the need for a DHCP server on each subnet. Also referred to as a *BOOTP relay agent*. |

| | |
|---|---|
| Unauthorized DHCP server | A DHCP server that has not explicitly been authorized. Sometimes referred to as a *rogue DHCP server*.<br><br>In a Windows Server 2003 domain environment, the DHCP Server service on an unauthorized server running Windows Server 2003 fails to initialize. The administrator must explicitly authorize all DHCP servers running Windows Server 2003 that operate in an Active Directory service domain environment. At initialization time, the DHCP Server service in Windows Server 2003 checks for authorization and stops itself if the server detects that it is in a domain environment and the server has not been explicitly authorized. |
| Automatic Private IP Addressing (APIPA) | A TCP/IP feature in Windows XP and Windows Server 2003 that automatically configures a unique IP address from the range 169.254.0.1 through 169.254.255.254 with a subnet mask of 255.255.0.0 when the TCP/IP protocol is configured for automatic addressing, the **Automatic private IP address** alternate configuration setting is selected, and a DHCP server is not available. The APIPA range of IP addresses is reserved by the Internet Assigned Numbers Authority (IANA) for use on a single subnet, and IP addresses within this range are not used on the Internet. |
| Superscope | A configuration that allows a DHCP server to provide leases from more than one scope to clients on a single physical network segment. |
| Multicast IP addresses | Multicast IP addresses allow multiple clients to receive data that is sent to a single IP address, enabling point-to-multipoint communication. This type of transmission is often used for streaming media transmissions, such as video conferencing. |
| Multicast Scope | A range of multicast IP addresses that can be assigned to DHCP clients. A multicast scope allows dynamic allocation of multicast IP addresses for use on the network by using the MADCAP protocol, as defined in RFC 2730. |
| BOOTP | An older protocol with similar functionality; DHCP is based on BOOTP. BOOTP is an established protocol standard used for configuring IP hosts. BOOTP was originally designed to enable boot configuration for diskless workstations. Most DHCP servers, including those running Windows Server 2003, can be configured to respond to both BOOTP requests and DHCP requests. |

# How DHCP Technology Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How DHCP Works**

DHCP provides an automated way to distribute and update IP addresses and other configuration information on a network. A DHCP server provides this information to a DHCP client through the exchange of a series of messages, known as the DHCP conversation or the DHCP transaction. If the DHCP server and DHCP clients are located on different subnets, a DHCP relay agent is used to facilitate the conversation.

### Note

- It is necessary to have an understanding of basic TCP/IP concepts, including working knowledge of subnets before you can have a full understanding of DHCP. For more information about TCP/IP, see "TCP/IP Technical Reference."

### In this section

- DHCP Architecture
- DHCP Protocols
- DHCP Processes and Interactions

**DHCP Architecture**

The DHCP architecture consists of DHCP clients, DHCP servers, and DHCP relay agents on a network. The clients interact with servers using DHCP messages in a DHCP conversation to obtain and renew IP address leases.

**DHCP Client Functionality**

A *DHCP client* is any network-enabled device that supports the ability to communicate with a DHCP server in compliance with RFC 2131, for the purpose of obtaining dynamic leased IP configuration and related optional information.

DHCP provides support for client computers running any of the following Microsoft operating systems:

- Windows NT version 4.0
- Windows 2000
- Windows XP
- Windows Server 2003
- Windows 98
- Windows Millennium Edition

**Automatic IP Configuration**

DHCP supports Automatic Private IP Addressing (APIPA), which enables computers running Windows 2000, Windows XP, and Windows Server 2003 to configure an IP address and subnet mask if a DHCP server is unavailable at system startup and the **Automatic private IP address** Alternate Configuration setting is selected. This feature is useful for clients on small private networks, such as a small-business office or a home office.

The DHCP Client service on a computer running Windows XP and Windows Server 2003 uses the following process to auto-configure the client:

1. The DHCP client attempts to locate a DHCP server and obtain an IP address and configuration.
2. If a DHCP server cannot be found or does not respond after one minute, the DHCP client checks the settings on the

**Alternate Configuration** tab of the properties of the TCP/IP protocol.

If **Automatic private IP address** is selected, the DHCP client auto-configures its IP address and subnet mask by using a selected address from the Microsoft-reserved Class B network, 169.254.0.0, with the subnet mask 255.255.0.0. The DHCP client tests for an address conflict to ensure that the IP address is not in use on the network. If a conflict is found, the client selects another IP address. The client retries auto-configuration up to 10 times.

If **User Configured** is selected, the DHCP client configures a static IP address configuration. The DHCP client tests for an address conflict to ensure that the IP address is not already in use on the network. If a conflict is found, the DHCP client indicates the error condition to the user.

3. When the DHCP client succeeds in self-selecting an address, it configures its network interface with the IP address. The client then continues to check for a DHCP server in the background every five minutes. If a DHCP server responds, the DHCP client abandons its self-selected IP address and uses the address offered by the DHCP server (and any other DHCP option information that the server provides) to update its IP configuration settings.

If the DHCP client obtained a lease from a DHCP server on a previous occasion, and the lease is still valid (not expired) at system startup, the client tries to renew its lease. If, during the renewal attempt, the client fails to locate any DHCP server, it attempts to ping the default gateway listed in the lease, and proceeds in one of the following ways:

- If the ping is successful, the DHCP client assumes that it is still located on the same network where it obtained its current lease, and continues to use the lease as long as the lease is still valid. By default the client then attempts, in the background, to renew its lease when 50 percent of its assigned lease time has expired.

- If the ping fails, the DHCP client assumes that it has been moved to a network where a DHCP server is not available. The client then auto-configures its IP address by using the settings on the **Alternate Configuration** tab. When the client is auto-configured, it attempts to locate a DHCP server and obtain a lease every five minutes.

**Local Storage**
Windows Server 2003 DHCP supports local storage, which allows clients to store DHCP information on their own hard disks. Local storage is useful because it enables the client to store its last leased IP address, so that when the client starts it first attempts to renew the lease of its previous IP address. Local storage also enables a client to be shut down and restarted and it will use its previously leased address and configuration, even if the DHCP server is unreachable or offline at the time that the client computer is restarted.

**DHCP Server Responsibilities**
The DHCP servers maintain scopes, reservations, and options as set by the administrator.

**Scopes**
A scope must be properly defined and activated before DHCP clients can use the DHCP server for automatic TCP/IP configuration. A DHCP scope is an administrative collection of IP addresses and TCP/IP configuration parameters that are available for lease to DHCP clients of a specific subnet. The network administrator creates a scope for each subnet.

A scope has the following properties:

- A scope name, assigned when the scope is created.

- A range of possible IP addresses from which to include or exclude addresses used in DHCP lease offers.

- A unique subnet mask, which determines the network ID for an IP address in the scope.

- Lease duration values.

Each DHCP scope can have a single continuous range of IP addresses. To use several address ranges within a single scope you must first define the entire address range for the scope, and then set exclusion ranges.

**Lease Durations**
When a scope is created, the lease duration is set to eight days by default. However there are situations when the administrator might want to change the lease duration. The following are examples of adjusting the lease duration due to individual network

consideration:

- An organization has a large number of IP addresses available and configurations that rarely change. The administrator increases the lease duration to reduce the frequency of lease renewal exchanges between clients and the DHCP server. Because the DHCP clients are renewing their leases less frequently, DHCP-related network traffic is reduced.

- A limited number of IP addresses are available and client configurations change frequently or clients move often in or out of the network. The administrator reduces the lease duration. This increases the rate at which unused addresses are returned to the available address pool for reassignment.

For example, consider the ratio between connected computers and available IP addresses. If 40 computers share 254 available addresses, the demand for reusing addresses is low. A long lease time, such as a few months, might be appropriate in such a situation. However, if 230 computers must share the same address pool, demand for available addresses is greater, and a shorter lease time, for example a few days, is more appropriate.

**Note**

- Although it is possible to configure a client with infinite lease duration, use infinite lease durations with caution. Even relatively stable environments have a certain amount of client turnover. At a minimum, computers might be added and removed, moved from one office to another, or network adapters might be replaced. If a client with an infinite lease is removed from the network without releasing its lease, the DHCP server is not notified, and the IP address is not automatically reused. Also, when using an infinite lease, options set on the DHCP server are not automatically updated on the DHCP client, because the client is never required to renew its lease and obtain the new options. It is recommended that reservations be used rather than infinite lease durations.

**Exclusion Ranges**

When you create a new scope, immediately exclude the addresses of existing statically configured computers from the scope. By using exclusion ranges, you can exclude specific IP address ranges within a scope so that those addresses are not offered to clients. Assign IP addresses within exclusion ranges to computers or devices that must have a static IP address, such as servers, firewalls, or routers.

You can use excluded IP addresses on your network by manually configuring these addresses at computers that do not use DHCP to obtain an address, or by configuring reservations for these addresses.

**Reservations**

You can reserve IP addresses for assignment to specified computers or devices on the network. Reservations ensure that a specified hardware device on a subnet always receives the same IP address lease. Use reservations for DHCP-enabled devices that must always have the same IP address on your network, such as servers that do not support Domain Name System (DNS) dynamic update.

**Note**

- If multiple DHCP servers are each configured with scopes that cover addresses that must be reserved, the reservations must be specified on each DHCP server. Otherwise, the client might receive an IP address from one of the DHCP servers that does not contain the reservation, and therefore might not receive the IP address reserved for the client.

**Superscopes**

A superscope allows a DHCP server to provide leases from more than one scope to clients on a single physical subnet. Before you can create a superscope, you must use the DHCP Microsoft Management Console (MMC) snap-in to define at least one of the scopes to be included in the superscope. Scopes added to a superscope are called *member scopes*. Superscopes can resolve DHCP Server service issues in several different ways; these issues include situations in which:

- Support is needed for DHCP clients on a single physical network segment — such as a single Ethernet LAN segment — where multiple logical IP networks are used. When more than one logical IP network is used on a physical network, these configurations are also known as *multinets*. In a situation where multinets are used, clients might not be able to communicate directly with each other, because the clients might be on different logical subnets, even if they are on the same physical network segment. In this case, routing must be enabled to allow the clients to communicate with each other. Also, a router or BOOTP/DHCP relay agent must be configured on the subnet to allow DHCP messages to travel between the logical subnets.

- Support is needed for DHCP clients that are in a multinet located on the other side of BOOTP relay agents.
- Clients need to be migrated to a new scope.

**Interactions between Client and Server**
DHCP servers and DHCP clients communicate through a series of DHCP messages. To obtain a lease, the DHCP client initiates a conversation with a DHCP server using a series of these DHCP messages.

**DHCP Messages**
The following list includes the eight types of messages that can be sent between DHCP clients and servers. For more information about the structure and specifics of each of these packets, see "DHCP Message Format" later in this section.

**DHCPDiscover**
Broadcast by a DHCP client when it first attempts to connect to the network. The DHCPDiscover message requests IP address information from a DHCP server.

**DHCPOffer**
Broadcast by each DHCP server that receives the client DHCPDiscover message and has an IP address configuration to offer to the client. The DHCPOffer message contains an unleased IP address and additional TCP/IP configuration information, such as the subnet mask and default gateway. More than one DHCP server can respond with a DHCPOffer message. The client accepts the best offer, which for a Windows DHCP client is the first DHCPOffer message that it receives.

**DHCPRequest**
Broadcast by a DHCP client after it selects a DHCPOffer. The DHCPRequest message contains the IP address from the DHCPOffer that it selected. If the client is renewing or rebinding to a previous lease, this packet might be unicast directly to the server.

**DHCPAck**
Broadcast by a DHCP server to a DHCP client acknowledging the DHCPRequest message. At this time, the server also forwards any options. Upon receipt of the DHCPAck, the client can use the leased IP address to participate in the TCP/IP network and complete its system startup. This message is typically broadcast, because the DHCP client does not officially have an IP address that it can use at this point. If the DHCPAck is in response to a DHCPInform, then the message is unicast directly to the host that sent the DHCPInform message.

**DHCPNack**
Broadcast by a DHCP server to a DHCP client denying the client's DHCPRequest message. This might occur if the requested address is incorrect because the client moved to a new subnet or because the DHCP client's lease has expired and cannot be renewed.

**DHCPDecline**
Broadcast by a DHCP client to a DHCP server, informing the server that the offered IP address is declined because it appears to be in use by another computer.

**DHCPRelease**
Sent by a DHCP client to a DHCP server, relinquishing an IP address and canceling the remaining lease. This is unicast to the server that provided the lease.

**DHCPInform**
Sent from a DHCP client to a DHCP server, asking only for additional local configuration parameters; the client already has a configured IP address. This message type is also used by DHCP servers running Windows Server 2003 to detect unauthorized DHCP servers.

**DHCP Lease Process**
A DHCP-enabled client obtains a lease for an IP address from a DHCP server. Before the lease expires, the DHCP client must renew the lease or obtain a new lease. Leases are retained in the DHCP server database for a period of time after expiration. By default, this grace period is four hours and cleanup occurs once an hour for a DHCP server running Windows Server 2003. This protects a clients lease in case the client and server are in different time zones, the internal clocks of the client and server computers are not synchronized, or the client is off the network when the lease expires.

**Obtaining a New Lease**
A DHCP client initiates a conversation with a DHCP server when it is seeking a new lease, renewing a lease, rebinding, or restarting. The DHCP conversation consists of a series of DHCP messages passed between the DHCP client and DHCP servers. The following figure shows an overview of this process when the DHCP server and DHCP client are on the same subnet.

### DHCP Lease Process Overview




1. The DHCP client requests an IP address by broadcasting a DHCPDiscover message to the local subnet.

2. The client is offered an address when a DHCP server responds with a DHCPOffer message containing an IP address and configuration information for lease to the client. If no DHCP server responds to the client request, the client sends DHCPDiscover messages at intervals of 0, 4, 8, 16, and 32 seconds, plus a random interval of between -1 second and 1 second. If there is no response from a DHCP server after one minute, the client can proceed in one of two ways:

   - If the client is using the Automatic Private IP Addressing (APIPA) alternate configuration, the client self-configures an IP address for its interface.

   - If the client does not support alternate configuration, such as APIPA, or if IP auto-configuration has been disabled, the client network initialization fails.

   In both cases, the client begins a new cycle of DHCPDiscover messages in the background every five minutes, using the same intervals as before (0, 4, 8, 16, and 32 seconds), until it receives a DHCPOffer message from a DHCP server.

3. The client indicates acceptance of the offer by selecting the offered address and broadcasting a DHCPRequest message in response.

4. The client is assigned the address and the DHCP server broadcasts a DHCPAck message in response, finalizing the terms of the lease.

When the client receives acknowledgment, it configures its TCP/IP properties by using the DHCP option information in the reply, and completes its initialization of TCP/IP.

In rare cases, a DHCP server might return a negative acknowledgment to the client. This can happen if a client requests an invalid or duplicate address. If a client receives a negative acknowledgment (DHCPNack), the client must begin the entire lease process again.

When the DHCP client and the DHCP server are on the same IP broadcast subnet, the DHCPDiscover, DHCPOffer, DHCPRequest, and DHCPAck messages are sent to identify clients by means of IP-level broadcasts sent to the limited broadcast address and the media access control (MAC) broadcast address.

When the DHCP server and DHCP client are not on the same subnet either a router or a host on the DHCP client's subnet must act as a DHCP relay agent to support the forwarding of DHCP messages between the DHCP client and the DHCP server.

**Renewing a Lease**

The DHCP client first attempts to renew its lease when 50 percent of the original lease time, known as *T1*, has passed. At this point the DHCP client sends a unicast DHCPRequest message to the DHCP server that originally granted its lease. If the server is available, and the lease is still available, the server responds with a unicast DHCPAck message and the lease is renewed.

If the original DHCP server is available, but the client's current lease is no longer available, the DHCP server responds with a DHCPNack message, and the client immediately starts the process to obtain a new lease. This can happen if the client has changed subnets or if the DHCP server cannot fulfill the lease request for some other reason.

If there is no response from the DHCP server, the client waits until 87.5 percent of the lease time has passed (known as *T2*). At T2, the client enters the rebinding state, and broadcasts a DHCPRequest message to attempt to renew the lease from any available DHCP server. If no DHCP server is available by the time the lease expires, the client immediately unbinds itself from the existing lease and starts the process to obtain a new lease, beginning with a DHCPDiscover message.

**Preventing Address Conflicts**
Windows Server 2003 DHCP has both server-side and client-side conflict detection to prevent duplicate IP addresses on your network.

**Client Conflict Detection**
Client computers running Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0, Windows Millennium Edition, and Windows 98 automatically check to determine if an IP address is already in use before using it.

After the DHCP client receives a lease from the DHCP server, the client sends an Address Resolution Protocol (ARP) request to the address that it has been assigned. If a reply to the ARP request is received, the client has detected a conflict and sends a DHCPDecline message to the DHCP server. The DHCP server attaches a BAD_ADDRESS value to the IP address in the scope for the length of the lease. The client then begins the lease process again, and is offered the next available address in the scope.

### Note

- ARP requests do not traverse routers. Clients use ARP requests rather than pings (ICMP Echo messages) because pings require the sender to have an IP address.

**Server Conflict Detection**
If your network includes older DHCP clients that do not perform conflict detection themselves, you can enable conflict detection on the DHCP server. By default, the Windows Server 2003 DHCP Server service does not perform any conflict detection.

To detect conflicts, the DHCP server pings (sends an ICMP Echo message to) an IP address before offering that address to clients in a new lease. The DHCP server only pings addresses that have not been successfully and previously leased. If a client requests a lease on an IP address that it already had or is requesting a renewal, the DHCP server does not ping the IP address.

If conflict detection is enabled, an administrator-defined number of pings are sent. The server waits 1 second for a reply. Because the time required for a client to obtain a lease is equal to the number of pings used, choose this value carefully because it directly impacts the overall performance of the server. In general, one ping is sufficient.

If a response to the ping is received, a conflict is registered and that address is not offered to clients requesting a lease from the server. The DHCP server then attaches a BAD_ADDRESS value to that IP address in the scope. The DHCP server then tries to lease the next available address. If the duplicate address is removed from the network, the BAD_ADDRESS value attached to the IP address can be deleted from the scope's list of active leases, and then the address returns to the pool. Addresses are marked as BAD_ADDRESS for the length of the lease for which the scope is configured. If the BAD_ADDRESS entry is not manually removed, it will automatically be removed after a period of time equal to the lease time for the scope.

### Note

- In general, use server conflict detection only as a troubleshooting aid when you suspect that duplicate IP addresses are in use on your network. Each additional conflict detection attempt adds to the time needed to negotiate leases for DHCP clients.

**DHCP Options**
*DHCP options* are additional configuration parameters that a DHCP server assigns to clients. Options can also be used for DHCP communication between the server computer and client computers.

The most specific options take precedence over the least specific options. This simplifies DHCP management and allows a flexible administration that can range from per-server default settings to common settings for a specific subnet and individualized client settings when needed for special circumstances. In most cases, the option values are specified in the **Options** dialog box on the DHCP server, scope, or reservation.

DHCP options can be configured for specific values and enabled for assignment and distribution to DHCP clients based on:

- **Server options.** These options apply globally for all scopes and classes defined at each DHCP server and any clients that it services. Configured server option values always apply unless they are overridden by options assigned to other scope, class, or client reservation.

- **Scope options.** These options apply to any clients that obtain a lease within that particular scope. Configured scope option values always apply to all computers obtaining a lease in a given scope unless they are overridden by options assigned to class or client reservation.

- **Class options.** These options apply to any clients that specify that particular DHCP Class ID value when obtaining a scope lease. Configured class option values always apply to all computers configured as members in a specified DHCP option class unless they are overridden by options assigned to a client reservation.

- **Reserved client options.** These options apply only to the client corresponding to the reservation. Reserved client option values override all other server, scope, or class assigned option values.

Options are typically applied at each DHCP server at the server or scope level. To precisely manage or customize option settings for a group or class of computers, specify either a user or vendor class assignment that overrides the broader server or scope option defaults.

For special requirements, such as clients with special functions, assign options for specific reserved clients.

Options can also be used to separate and distribute appropriate options for clients with similar or special configuration needs. For example, DHCP clients on the same floor of a building can be configured with the same DHCP Class ID value to assign them membership in the same option class. You can then distribute additional or varied option data to that class during the lease process, overriding any scope or globally provided default options.

**Note**

- Statically configured values on a client override any DHCP options of any type or level.

Many options are predefined on a DHCP server running Windows Server 2003. Other standard DHCP options can be added as needed to support any other DHCP client software that recognizes or requires the use of these additional options. The DHCP Server service running on Windows Server 2003 supports all options defined in RFC 2132, although most DHCP clients use or support only a small subset of the available RFC-specified options.

The following table contains a list of default DHCP options requested by DHCP clients running Windows Server 2003 and Windows XP. For a complete reference of DHCP options, see "DHCP Tools and Settings."

**Default DHCP Options**

| Option Code | Option Name |
|---|---|
| 1 | Subnet mask |
| 3 | Router |
| 6 | DNS servers |
| 15 | DNS domain name |
| 44 | WINS/NBNS servers |
| 46 | WINS/NetBT node type |
| 47 | NetBIOS scope ID |
| 51 | Lease time |
| 58 | Renewal (T1) time value |
| 59 | Rebinding (T2) time value |
| 31 | Perform router discovery |
| 33 | Static route |
| 43 | Vendor-specific information |
| 249 | Classless static routes |

**DHCP Option Parameters**

DHCP servers can be configured to provide optional data that fully configures TCP/IP on a client. Some of the most common

DHCP option types configured and distributed by the DHCP server during the lease process include parameters for the default gateway, DNS, and WINS.

Clients can be configured with:

- **Information options.** You can explicitly configure these options and any associated values provided to clients.
- **Protocol options.** You can implicitly configure these options used by the DHCP Server service based on server and scope property settings.

You can use the DHCP snap-in to configure these properties and set them for an entire scope or for a single, reserved client scope.

**Information Options**

The following table lists the most common types of DHCP information options that can be configured for DHCP clients. These options can be enabled and configured for each scope that you configure on a DHCP server. Depending on your network infrastructure, some of these options can be configured as server options, such as DNS domain name.

**Common Information Options**

| Code | Description |
|---|---|
| 3 | Router |
| 6 | DNS server |
| 15 | DNS domain name |
| 44 | WINS/NBNS servers |

Clients can request these DHCP options, and can use the values to set their TCP/IP configurations for the duration of the lease.

**Protocol Options**

The following table shows protocol options that DHCP clients can be configured to use when communicating with a DHCP server to obtain or renew a lease.

**Common Protocol Options**

| Code | Description |
|---|---|
| 51 | Lease time |
| 53 | DHCP message type |
| 55 | Special option type used to communicate a parameter request list to the DHCP server |
| 58 | Renewal time value (T1) |
| 59 | Rebind time value (T2) |

The values provided to clients for lease time, T1, and T2 are taken from the scope settings on the DHCP server. The value provided for DHCP message type is automatically set depending on which packet of the DHCP conversation is being sent.

**Option Classes**

Option classes allow quick introduction of custom applications for enterprise networks. DHCP option classes provide a way to easily configure network clients with the parameters necessary to meet the special requirements of custom applications. Equipment from multiple vendors on a network can also use different option code numbers for different functions. The options used to support vendor classes — the vendor class identifier and the vendor-specific option — are defined in the Internet DHCP options standard reference, RFC 2132.

Windows Server 2003 includes two types of option classes: vendor-defined and user-defined. These classes can be configured on your servers to offer specialized client support in the following ways:

- Add and configure vendor-defined classes for managing DHCP options assigned to clients identified by vendor type.

- Add and configure user-defined classes for managing DHCP options assigned to clients that need a similar DHCP option configuration.

After options classes are defined on a DHCP server, scopes on the server can be configured to assign options for specific user-defined and vendor-defined option classes.

**Vendor Classes**
Vendor-defined option classes can be used by DHCP clients to identify the client's vendor type and configuration when obtaining a lease from the DHCP server. The client can include the vendor class ID option (option code 60) when it requests or selects a lease from a DHCP server to identify its vendor class during the lease process.

The vendor class identifier information is a string of character data interpreted by the DHCP servers. Vendors can choose to define specific vendor class identifiers to convey particular configuration or other identification information about a client. For example, the identifier might encode the client's hardware or software configuration. Most vendor types are derived from standard reserved hardware and operating system-type abbreviation codes listed in RFC 1700.

When vendor options are specified, the server performs the following additional steps to provide a lease to the client:

- The server verifies that the vendor class identified by the client request is a recognized class defined on the server.

  If the vendor class is recognized, the server checks to see if any additional DHCP options are configured for this class in the active scope.

  If the vendor class is not recognized, the server ignores the vendor class identified in the client request, and returns options allocated to the default vendor class (which includes all DHCP Standard options).

- If the scope contains options configured specifically for use with clients in this vendor-defined class, the server returns those options using the vendor-specific option type (option code 43) as part of its acknowledgment message.

In most cases, the default vendor class — the DHCP Standard option class — provides a default vendor class for any Windows DHCP clients or other DHCP clients that do not specify a vendor class ID. In some cases, you might define additional vendor classes for other DHCP clients, such as printers or some types of UNIX clients. When you add other vendor classes for these purposes, make sure that the vendor class identifier you use to configure the class at the server matches the identifier used by clients for your third-party vendor.

**User Classes**
User classes allow DHCP clients to differentiate themselves by specifying what type of client they are, such as desktop or server computer. For computers running Windows Server 2003, Windows XP, and Windows 2000, you can define specific user class identifiers to convey information about a client's software configuration, its physical location in a building, or about its user preferences. For example, an identifier can specify that DHCP clients are members of a user class called "2nd floor, West", which has need for a specific set of router, DNS, and WINS server settings. An administrator can then configure the DHCP server to include different option values depending on the user class of client receiving the lease.

Windows Server 2003 user classes can be used as follows:

- DHCP client computers can include the DHCP user class option when sending DHCP request messages to the DHCP server. This can specifically identify the client as part of a user class on the server.

- DHCP servers running the Windows 2000 Server or Windows Server 2003 DHCP Server service can recognize and interpret the DHCP user class option from clients and provide additional options (or a modified set of DHCP options) based on the client's user class identity.

For example, shorter leases can be assigned to wireless clients. Or perhaps a particular set of clients might need a specific set of routes, a specific DNS server, or a specific default gateway.

**Note**

- If user classes are not specified, default settings, such as server options or scope options, are assigned.

A user class can be either a default or custom user class. Microsoft provides three default user classes, as described in the

following table.

**Default User Classes Provided by Windows DHCP**

| Class Type | Class ID String | Description |
|---|---|---|
| Default User Class | (Unspecified) | This class is typically used by most DHCP clients. Clients that are included in this class:<br>• DHCP clients that cannot be configured with a user class or a user class ID. This is true for most Windows-based DHCP clients prior to Windows 2000.<br>• Clients running Windows Server 2003, Windows XP, or Windows 2000 configured with a class ID unknown to the DHCP server.<br>• Clients that do not otherwise specify a user class ID. |
| Default Routing and Remote Access class | RRAS.Microsoft | This class is used by the Windows 2000 Server or Windows Server 2003 DHCP Server service to classify clients making a PPP-type connection through a remote access server. Typically, this class includes most dial-up networking clients that use DHCP to obtain a lease, including remote access clients that cannot be configured with a Routing and Remote Access user class or a Routing and Remote Access user class ID.<br>See "DHCP and Routing and Remote Access" later in this topic for details about the interaction between a Routing and Remote Access server and a DHCP server and how DHCP servers identify remote access clients. |
| Default BOOTP class | BOOTP | This class is used by the Windows 2000 Server or Windows Server 2003 DHCP Server service to classify any clients recognized as BOOTP clients. |

Use the Microsoft default user classes to isolate specific configuration details for clients with special needs, such as older clients or clients that use BOOTP or Routing and Remote Access. For example, you might want to include and assign special BOOTP option types (such as option codes 66 and 67) for clients that are BOOTP type, or shorten the lease time for remote access clients.

You can also add and configure custom user classes for use by DHCP clients running Windows 2000, Windows XP, and Windows Server 2003. For a custom user class to work properly, the client must use the same custom identifier when requesting options as was used when the class was defined on the DHCP server

The user class option field permits only one ASCII text string to be used for identifying clients. This means each client computer can be identified only as a member of a single user class by the DHCP server. You can use additional user classes to make new hybrids from your other user classes to accommodate clients that need configuration for multiple user classes. For example, if you have two user classes, one called "mobile" with short lease times assigned and another called "engineer" with an option assigned to configure a high-performance server for its clients, you can make a new hybrid user class called "mobile-engineer" that contains both special option value settings.

**MADCAP and Multicast DHCP**

Multicast Address Dynamic Client Allocation Protocol (MADCAP) is modeled after the DHCP standard. MADCAP assists in simplifying and automating configuration of multicast groups on your network, but it is not required for the operation of multicast groups or for the DHCP Server service. Multicast scopes provide only multicast address configuration and do not support or use other DHCP-assignable options.

Multicast scopes configured on the DHCP server define ranges of IP multicast addresses. Similar to allocating unicast IP addresses, IP multicast addresses are allocated to MADCAP clients. A MADCAP address is configured separately from a primary IP address. Computers that use either static or dynamic IP configuration through a DHCP server can be MADCAP clients.

In Windows Server 2003, the DHCP Server service supports both DHCP and MADCAP, although these services function separately. Clients of one do not depend on the use or configuration of the other.

Clients that do not support the MADCAP service or are unable to contact and obtain multicast configuration from a MADCAP server can be configured in other ways so that they participate in either permanent or temporary multicast groups on the network.

In all TCP/IP networks, each computer requires a unique primary unicast IP address for each network interface. You must assign this required primary unicast IP address before you can configure a computer to support and use secondary IP addresses such as multicast IP addresses.

**DHCP Protocols**

In Windows Server 2003, the DHCP Server service includes support for the Dynamic Host Configuration Protocol (DHCP), the Multicast Address Dynamic Client Allocation Protocol (MADCAP), and the Bootstrap Protocol (BOOTP).

**DHCP**

DHCP servers communicate with DHCP clients by using a series of DHCP messages. The format of DHCP messages is based on the message format used with the BOOTP protocol.

RFC 2131 defines the format for each message sent between a DHCP client and a DHCP server. The following table shows the possible fields in the DHCP messages.

**DHCP Message Fields**

| Field Name | Friendly Name | Field Length (Octets) | Description |
|---|---|---|---|
| op | Message Type | 1 | Message type |
| htype | Hardware Address Type | 1 | Hardware address type. Defined at http://www.iana.org/assignments/arp-parameters |
| hlen | Hardware Address Length | 1 | Hardware address length in octets |
| hops | Hops | 1 | Value is set to zero by DHCP clients. Optionally used to count the number of relay agents that forwarded the message. |
| xid | Transaction ID | 4 | A random number used to associate messages and responses between a client and a server. |
| secs | Seconds | 2 | Seconds elapsed since client began address acquisition or renewal process. |
| flags | Flags | 2 | Flags set by client. The Broadcast flag is set if the client cannot receive unicast IP datagrams (for example, before it is configured with an IP address). |
| ciaddr | Client IP Address | 4 | This field is only filled in if the client has an IP address and can respond to ARP requests. |
| yiaddr | Your IP Address | 4 | Address given to the DHCP client by the DHCP server |
| siaddr | DHCP Server IP Address | 4 | IP address of the server that is offering a lease |
| giaddr | Gateway IP Address | 4 | DHCP relay agent IP address |
| chaddr | Client Hardware Address | 16 | Client hardware address |
| sname | Server Host Name | 64 | Optional server host name. Not used in Windows Server 2003 |

| | | | |
|---|---|---|---|
| file | Boot File Name | 128 | The name of the file containing the boot image for a BOOTP client |
| options | Options | variable | Optional parameters field. In the DHCP protocol packet, each option begins with a single octet tag, which holds the option code, and a second octet, which describes the option data length, in bytes. For a complete list of the DHCP options available by default on a DHCP server running on Windows Server 2003, see "DHCP Tools and Settings." |

For a complete view of how these fields are used in each DHCP message, see RFC 2131 or use a network monitoring tool, such as Netmon, to view the DHCP messages.

**MADCAP**

Windows Server 2003 includes a Multicast Address Dynamic Client Allocation Protocol (MADCAP) Server service to support dynamic assignment and configuration of IP multicast addresses on TCP/IP-based networks.

Whereas DHCP unicast scopes provide client configurations by allocating ranges of IP addresses for point-to-point communication between two networked computers, multicast scopes provide ranges for multicast IP addresses. These addresses are reserved for multicast operation using directed transmission from one point to multiple points.

A multicast address is shared by many computers. A group of TCP/IP computers can use a single multicast IP address to send directed communication to all computers with which they share the use of the group address. An IP datagram that is sent to the multicast address is forwarded to all members of that multicast group.

**Dynamic Membership**

Multicast addresses support dynamic membership, allowing individual computers to join or leave the multicast group at any time. The size of the group is not limited, and computers can be members of multiple groups. In addition, any computer that uses TCP/IP can send datagrams to any multicast group.

**Multicast Address Ranges**

You can permanently reserve multicast group addresses or temporarily assign and use them. A permanent group is made by permanently reserving a multicast IP address (224.0.0.0 to 239.255.255.255) with the Internet Assigned Numbers Authority (IANA). The reserved address then becomes a well-known address, indicating a specific multicast group that exists regardless of whether group member computers are present on the network. Any multicast IP address that is not permanently reserved with the IANA can then be used dynamically to assign and form temporary multicast groups. These temporary groups can exist as long as one or more computers on the network are configured with the group's address and actively share in its use.

**BOOTP**

Bootstrap Protocol (BOOTP) is a computer configuration protocol developed before DHCP. DHCP improves on BOOTP and resolves specific limitations that BOOTP had as a computer configuration service. RFC 951 defines BOOTP.

Whereas BOOTP configures diskless workstations with limited boot capabilities, DHCP configures networked computers, that have local hard drives and full boot capabilities.

Likewise, although both BOOTP and DHCP allocate IP addresses to clients during startup, they use different methods of allocation. BOOTP typically provides fixed allocation of a single IP address for each client, permanently reserving this address in the BOOTP server database. DHCP typically provides dynamic, leased allocation of available IP addresses, reserving each DHCP client address temporarily in the DHCP database.

Because of the relationship between BOOTP and DHCP, both protocols share some defining characteristics. BOOTP and DHCP use nearly identical request messages and reply messages. Both protocols enclose each protocol message in a single User Datagram Protocol (UDP) datagram of 576 bytes. Message headers are the same for both BOOTP and DHCP, except for the final message header field that carries optional data. For BOOTP, this optional field is called the *vendor-specific area* and is limited to 64 bytes. For DHCP, this optional field is called the *options field* and is at least 312 bytes long.

Both BOOTP and DHCP use the same reserved protocol ports for sending and receiving messages between servers and clients. Both BOOTP and DHCP servers use UDP port 67 to listen for and receive client request messages. BOOTP and DHCP clients typically reserve UDP port 68 for accepting message replies from either a BOOTP server or DHCP server.

Because DHCP and BOOTP messages use nearly identical format types and packet structures, and use the same well-known service ports, BOOTP or DHCP relay agent programs usually treat BOOTP and DHCP messages as the same message type and do not differentiate between them.

BOOTP clients do not rebind or renew configuration with the BOOTP server except when the system restarts, whereas DHCP clients do not require a system restart to rebind or renew configuration with the DHCP server. Instead, clients automatically enter the rebinding state at defined intervals to renew their leased address allocation with the DHCP server. This process occurs in the background and is transparent to the user.

BOOTP uses a two-phase bootstrap configuration process in which clients contact BOOTP servers to perform address determination and boot file name selection, and clients also contact Trivial File Transfer Protocol (TFTP) servers to perform file transfer of their boot image. DHCP uses a single-phase boot configuration process whereby a DHCP client negotiates with a DHCP server to determine its IP address and obtain any other initial configuration details it needs for network operation.

Because BOOTP clients contact TFTP servers to perform file transfer of their boot image and Windows Server 2003 does not provide a TFTP file service, you need a third-party TFTP server to support BOOTP clients that must boot from an image file (usually diskless workstations). You also need to configure your DHCP server to provide supported BOOTP/DHCP options.

**DHCP Options Supported for BOOTP Clients**

To obtain other options, BOOTP clients must specify DHCP option code 55 (the Options Request List parameter) in the BOOTP request. BOOTP clients that do not specify option 55 can still retrieve the options listed in the following table from DHCP servers running Windows NT Server 4.0 or later, if they are configured on the server.

**DHCP Options for BOOTP Clients**

| Code | Option Name |
|---|---|
| 1 | Subnet Mask |
| 3 | Router |
| 4 | Time Server |
| 5 | Name Server |
| 9 | LPR Server |
| 12 | Computer Name |
| 15 | Domain Name |
| 17 | Root Path |
| 42 | NTP Servers |
| 44 | WINS Server |
| 45 | NetBIOS over TCP/IP Datagram Distribution Server |
| 46 | NetBIOS over TCP/IP Node Type |
| 47 | NetBIOS over TCP/IP Scope |
| 48 | X Window System Font Server |
| 49 | X Window System Display Manager |
| 69 | SMTP Server |
| 70 | POP3 Server |

DHCP servers running Windows Server 2003 return the options in the order listed above and return as many options as can fit in a single datagram response. For more information about individual DHCP options, see "DHCP Tools and Settings."

**Note**

- When configuring client reservations for use with BOOTP clients, remember that DHCP options can apply equally to DHCP and BOOTP clients.

**BOOTP Table**

Each record in the BOOTP table has three fields of information that is returned to the BOOTP client:

- **Boot Image**. Identifies the generic file name (such as "unix") of the requested boot file, based on the BOOTP client's hardware type.

- **File Name**. Identifies the full path of the boot file (such as "/etc/vmunix") that the BOOTP server returns to the client by using TFTP.

- **File Server**. Identifies the name of the TFTP server used to store the boot file.

To add entries in the BOOTP table, use the DHCP snap-in.

**DHCP Processes and Interactions**
In Windows Server 2003, the DHCP Server service interacts with several other services, including the Active Directory directory service, DNS, and the Routing and Remote Access service.

**Detecting Unauthorized DHCP Servers**
An unauthorized DHCP server on a network can cause a variety of problems, such as the leasing of incorrect IP addresses and options. To protect against this type of problem, when a DHCP server running Windows 2000 or Windows Server 2003 starts on the network, it first attempts to determine if it is authorized to service clients. There are different methods used depending on how the network is configured.

**Unauthorized Domain Member DHCP Servers**
A domain member DHCP server queries Active Directory. The DHCP server compares its IP address and server name to the list of authorized DHCP servers. If either the server name or IP address is found on the list of authorized DHCP servers, the server is authorized as a DHCP server. If no match is found, the server is not authorized in Active Directory, the server does not respond to DHCP traffic, and a system event is logged.

### Note

- This process of authorizing DHCP servers is useful for only DHCP servers running Windows 2000 or Windows Server 2003. This process cannot be used for DHCP servers running Windows NT Server 4.0, or servers running non-Windows-based DHCP Server services. Only a member of the Enterprise Admins group can authorize or unauthorize a DHCP server in Active Directory.

**Unauthorized Workgroup DHCP Servers**
A Windows Server 2003 workgroup member DHCP server uses the following process to detect other DHCP servers currently running on the reachable network and to determine if it is authorized to provide service.

1. When the DHCP Server service starts, it sends a DHCPInform request message to the reachable network, using the local limited broadcast address (255.255.255.255), to locate other DHCP servers on the network.

   This message includes several vendor-specific option types that are known and supported by other DHCP servers running Windows Server 2003. These other DHCP servers will respond with a DHCPAck containing information indicating if they are authorized domain member or workgroup member servers.

2. When queried, other DHCP servers running Windows 2000 and Windows Server 2003 reply with DHCPAck messages to acknowledge and answer with workgroup or domain membership information.

3. If an Active Directory domain member DHCP server is found, then the workgroup member server determines that it is not authorized and does not service clients. If other workgroup servers are found, the workgroup member server determines that it is authorized to service clients, and begins service. It then performs the check again at one-hour intervals.

**DHCP and DNS**
Domain Name System (DNS) servers provide name resolution for network clients. DNS resolves a fully qualified domain name (FQDN) to its corresponding IP address.

Although DHCP provides a powerful mechanism for automatically configuring client IP addresses, prior to Windows 2000, the DHCP Server service did not notify DNS to update the DNS records on behalf of the client. Specifically, DHCP did not map the client name to an IP address and did not update IP address-to-name mappings using DNS dynamic update.

Without a way for DHCP to interact with DNS, the information maintained by DNS for a DHCP client might be incorrect. For example, a client can acquire its IP address from a DHCP server, but the DNS records might not reflect the IP address acquired

nor provide a mapping from the new IP address to the FQDN.

**DNS Dynamic Updates**

In Windows 2000 and Windows Server 2003, DHCP servers and clients can register record updates if the DNS server supports DNS dynamic updates. In Windows 2000 Server and Windows Server 2003, the DNS service supports DNS dynamic updates.

A DHCP server running Windows Server 2003 can register with a DNS server and update pointer (PTR) and address (A) resource records on behalf of its DHCP-enabled clients by using the DNS dynamic update protocol.

The ability to register A and PTR resource records lets a DHCP server act as a DNS registration proxy for clients using Windows NT 4.0, Windows 98, or Windows Millennium Edition, and possibly other clients that are not able to register the updates on their own, as shown in the following figure.

**DHCP Server Performing DNS Dynamic Update on Behalf of DHCP Client**




DHCP clients running Windows 2000, Windows XP, and Windows Server 2003 interact with DNS differently than DHCP clients running earlier versions of Windows. DHCP clients running Windows XP, Windows 2000, or Windows Server 2003 typically update their own dynamic forward lookup names, as shown in the following figure.

**DHCP Client and DHCP Server Performing DNS Dynamic Update**




An additional DHCP option code (option code 81) enables the return of a client's FQDN to the DHCP server. If implemented, the DHCP server can dynamically update an individual computer's resource records on a DNS server by using the DNS dynamic update protocol.

For more information about DNS dynamic updates, see "DNS Technical Reference."

**Secure DNS Dynamic Updates**

By itself, DNS dynamic update is not secure; any client can modify DNS records. When secure DNS dynamic update is configured, the authoritative name server accepts updates only from clients and servers that are authorized to make DNS dynamic updates to the appropriate objects in Active Directory. Secure DNS dynamic update is available only on Active Directory–integrated zones.

Secure DNS dynamic update protects zones and resource records from being modified by unauthorized users by allowing you to specify the users and groups that can modify zones and resource records. By default, Windows Server 2003, Windows XP Professional, and Windows 2000 clients attempt unsecured DNS dynamic updates first. If that request fails, they attempt secure updates.

When using multiple DHCP servers and secure DNS dynamic updates, add each of the DHCP servers as members of the DnsUpdateProxy global security group so that any DHCP server can perform a secure DNS dynamic update for any record. Otherwise, when a DHCP server performs a secure DNS dynamic update for a record, that DHCP server is the only computer that can update the record.

**DHCP and Routing and Remote Access**

The Windows Server 2003 DHCP Server service interacts with the Windows Server 2003 Routing and Remote Access Server service in two specific ways. When Routing and Remote Access is used to provide remote access to PPP clients, the remote access server obtains IP addresses from the DHCP server, which it then assigns to the PPP clients.

DHCP also interacts with routers when DHCP clients and DHCP servers are on different subnets from each other. In this

situation, a router that can act as a DHCP relay agent must be present on the subnet of the DHCP client. You can use the Windows Server 2003 Routing and Remote Access service to act as a DHCP relay agent.

**Configuration of PPP Clients**

When the Routing and Remote Access service is configured to use DHCP to obtain IP addresses for TCP/IP based clients, the Routing and Remote Access service instructs the DHCP Client service to obtain 10 IP addresses from a DHCP server when the first PPP client connects. The Routing and Remote Access service uses the first IP address obtained from DHCP for the Internal interface, which is a logical interface that represents the connections to all PPP-based clients. Subsequent addresses are assigned to TCP/IP-based PPP clients as they connect. After the PPP client disconnects, the now-unassigned IP address is reused for a future PPP connection.

The remote access server uses the IP addresses from these leases to configure PPP clients, but discards all options contained in the leases.

When all 10 IP addresses are used, the Routing and Remote Access service obtains another block of 10 IP addresses from the DHCP server.

With a Windows NT 4.0–based remote access server, DHCP-allocated addresses are recorded and reused when the remote access service is restarted. In Windows Server 2003 and Windows 2000 Server, the Routing and Remote Access service releases all DHCP-allocated IP addresses by using DHCPRelease messages each time the service is stopped.

If the DHCP server becomes unavailable, the DHCP Client service on the Routing and Remote Access server assigns APIPA addresses to TCP/IP-based PPP clients. APIPA addresses for remote access connectivity work only if the network to which the remote access client is attached is also using APIPA addresses (which is not a recommended configuration). If the local network is not using APIPA addresses, remote access clients can only obtain point-to-point remote access connectivity.

The Routing and Remote Access service uses a specific LAN interface to obtain DHCP-allocated IP addresses for remote access clients. You can select which LAN interface to use on the **IP** tab of the **Properties** dialog box of a server in the **Routing and Remote Access** snap-in. If the Routing and Remote Access server has more than one LAN interface installed, the Routing and Remote Access Server Setup Wizard prompts you to select the LAN interface.

**Options for PPP Clients**

Although the remote access server running Windows Server 2003 discards all options from the leases it obtains from the DHCP server, PPP clients do receive specific configuration information, such as WINS server and DNS server assignments, from the settings of the remote access server as part of the negotiation of the PPP connection. However, clients running Windows 2000, Windows XP, or Windows Server 2003 can receive additional configuration information from the DHCP server, by using a DHCPInform message after the connection has been established. These options are available only if the VPN server has the DHCP Relay Agent routing protocol component configured with the IP address of the DHCP server.

The following three figures show the three steps of the remote access server obtaining leases from the DHCP server.

First, when the first PPP client connects to the remote access server, the remote access server obtains 10 IP addresses from the DHCP server as shown in the following figure.

**Remote Access Server Obtains IP Addresses**




Next, the remote access server uses Internet Protocol Control Protocol (IPCP) to configure the IP address of the client, as well as assign the DNS server and WINS server settings that are configured on the selected interface of the remote access server, as shown in the following figure.

**Remote Access Server Configures PPP Client**




After the remote access client has an IP address, it sends a unicast DHCPInform message to request options from the DHCP server to the remote access server. The remote access server must also be configured with the DHCP Relay Agent routing

protocol component. The remote access server, acting in its role as a DHCP relay agent, then sends the DHCPInform message to the DHCP server. The DHCP server responds with the options in a DHCPAck message, which is sent back to the remote access server. Finally, the DHCP relay agent on the remote access server sends the DHCPAck message to the remote access client, as shown in the following figure.

**PPP Client Obtains DHCP Options**




The following table lists the DHCP options that are requested by the client in the DHCPInform message.

**DHCP Options Requested in DHCPInform Message**

| Code | Description |
|---|---|
| 1 | Subnet mask |
| 6 | DNS servers |
| 15 | DNS domain name |
| 43 | Vendor-specific information |
| 44 | WINS/NBNS servers |
| 249 | Classless static routes |

**DHCP Relay Agents**

A DHCP relay agent is a hardware device or software program that can pass DHCP or BOOTP messages between DHCP clients and servers, according to the RFC 2131 specification for DHCP. DHCP relay agents act as proxies, forwarding messages from a subnet to one or more DHCP servers. Some DHCP messages are sent as broadcasts, so without relay agents and the ability to pass DHCP and BOOTP messages across routers, every subnet on a network must have its own DHCP server.

Most routers support acting as a DHCP relay agent. Alternatively, if a router cannot function as a DHCP relay agent, a computer that can function as a DHCP relay agent must be configured on each subnet to which the router is connected.

In cases where it is impractical or impossible to configure routers to act as a DHCP relay agent, you can configure a computer running Windows Server 2003 or Windows 2000 Server, to act as a relay agent by enabling the Routing and Remote Access service and installing the DHCP Relay Agent routing protocol component.

**How Relay Agents Work**

The following figure shows how Client C on Subnet 2 obtains a DHCP address lease from DHCP Server 1 on Subnet 1.

**Using a Relay Agent**




1. DHCP Client C broadcasts a DHCPDiscover message on Subnet 2 as a UDP datagram over well-known UDP port 67, which is the port reserved and shared for BOOTP and DHCP server communication.

2. The relay agent, in this case a DHCP relay-enabled router, examines the Gateway IP Address field (also known as the giaddr field) in the DHCP message header. If the field has an IP address of 0.0.0.0, the agent fills the Gateway IP Address field with the IP address assigned to the interface on which the DHCPDiscover was received, and forwards the

DHCPDiscover message as a unicast message to the DHCP server on Subnet 1.

3. When DHCP Server 1 receives the DHCPDiscover message, it examines the Gateway IP Address field to determine if the packet was relayed. The DHCP server then determines whether it can supply an IP address lease to clients on the subnet indicated by the address in the Gateway IP Address field.

   For example, if the Gateway IP Address field has an IP address of 192.168.45.2, the DHCP server checks its DHCP scopes for a scope range that includes 192.168.45.2 (the gateway IP address in the packet). To find a match, the DHCP server determines whether the IP address 192.168.45.2 matches the network ID of each scope. It determines the network ID of each scope by ANDing any IP address in the scope with its corresponding subnet mask. In this case, the DHCP server checks to see which scope includes addresses for Subnet 2. If a scope exists that matches this criterion, the DHCP server selects an available address from the matched scope to use in an IP address lease offer response to the client.

4. DHCP Server 1 sends a DHCPOffer message directly to the relay agent identified in the Gateway IP Address field.

5. The relay agent relays the DHCPOffer message to the DHCP client.

   Depending on the value of the Broadcast flag in the DHCPOffer message (which was copied from the DHCPDiscover message), the DHCP relay agent either unicasts or broadcasts the DHCPOffer message.

6. By using a similar process, a DHCPRequest message is relayed from client to server, and a DHCPAck message is relayed from server to client.

For more information about the DHCP relay agent, see "How the DHCP Relay Agent Works" in How Unicast IPv4 Routing Protocols and Services Work.

# DHCP Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP Tools and Settings**

**In this section**


- DHCP Tools
- DHCP Options
- Related information


**DHCP Tools**

The following tools are associated with DHCP.

**DHCP Snap-in**

The DHCP snap-in allows you to perform a variety of administrative tasks for your DHCP servers:

- Create and manage scopes, including superscopes and multicast scopes.
- Create and manage properties for scopes, such as options, reservations, and exclusion ranges.
- Review active leases for each scope.

**Category**

The DHCP snap-in Microsoft Management Console (MMC) appears as an administrative tool after you install DHCP from Add/Remove Windows Components in Control Panel. The DHCP snap-in can also be added to Windows Server 2003 or Windows XP by installing the Windows Server 2003 Administrative Tools Pack. This allows remote administration of DHCP servers running Windows 2000 Server or Windows Server 2003 from a Windows XP-based workstation.

**Version Compatibility**

The Windows Server 2003 DHCP snap-in is compatible with DHCP servers running Windows Server 2003 and Windows 2000.

**Netsh**

Netsh is a command-line scripting tool that allows you to display or modify the network configuration of a computer. Netsh also provides a scripting feature that allows you to run a group of commands in batch mode against a specified computer. Netsh can also save a configuration script in a text file for archival purposes or for reuse in configuring other servers.

Commands in the **netsh dhcp** context provide a command-line method to help with the administration of DHCP servers and provides an equivalent alternative to console-based management. All commands in **netshdhcp** context can also be executed against a specified remote server.

**Category**

Netsh is a command-line tool.

**Version compatibility**

You can use netsh commands with Windows 2000 Server and Windows Server 2003.

**Network Monitor**

You can use the Network Monitor tool or a commercial packet analyzer (also known as a network sniffer), to capture and view packets such as DHCP messages.

In Windows 2000 Server and Windows Server 2003, Network Monitor is installed as an optional management and monitoring component by using Add or Remove Programs in Control Panel. After it is installed, you can run Network Monitor from the Administrative Tools folder.

**Category**

Network Monitor is available in Microsoft Systems Management Server or with Windows 2000 Server and Windows Server 2003.

**Version compatibility**
You can use Network Monitor to capture and view packets in Windows XP, Windows 2000, or Windows Server 2003.

**DHCP Options**
This section lists the predefined options available for use with the Microsoft Windows Server 2003 DHCP service. These options are defined according to the updated standards reference for DHCP options in RFC 2132, "DHCP Options and BOOTP Vendor Extensions."

Use the DHCP Microsoft Management Console (MMC) snap-in to specifically configure each option value, and enable the option for assignment and distribution to DHCP clients based on server, scope, class, or client-specific levels of preference.

In this section:

- **Basic options.** These options were originally defined in RFC 1497 and relisted in RFC 2123.
- **IP host options.** These options affect the operation of the IP layer, on a per-host basis.
- **IP interface options.** These options affect operation of the IP layer on a per-interface basis.
- **Link layer options.** These options affect operation of the data-link layer on a per-interface basis.
- **TCP options.** These options affect operation of the TCP layer on a per-interface basis.
- **Application layer options.** These options affect the application layer operations on a per-interface basis.
- **NetBIOS over TCP/IP options.** These options are used to support NetBIOS over TCP/IP.
- **Vendor-specific options.** These options are specified for vendor class use..
- **User class options.** These options are specified for user class use.
- **DHCP extensions.** These options are used to implement default protocol interaction and system behavior between servers and clients.
- **Administrator-defined options.** These options are not preconfigured in Windows Server 2003, but can be defined by an administrator.
- **Microsoft options.** These options are only available for use with supported Microsoft DHCP clients, such as computers running Windows Server 2003 or Windows XP.

**Note**

- For all DHCP options that use a list of IP addresses as the value data, the IP addresses are always used in order of preference by the DHCP client. For example, the first address in the list is used first.

**Basic Options**
The following sections list the basic DHCP options originally defined in RFC 1497 and updated in RFC 2132, for use with DHCP and the Boot Protocol (BOOTP) service. The BOOTP service refers to options as vendor extensions.

The DHCP service supports configuration and distribution of any options assigned using the DHCP Manager snap-in. By default, Microsoft DHCP-enabled clients require and provide storage and interpretation for options 1 (Subnet Mask), 3 (Router), 6 (DNS Servers), and 15 (DNS Domain Name).

**Pad Option**
A single octet of zero ("00") used for padding. This option differs from most DHCP options in that it does not use a length or value field. This option used to ensure that subsequent DHCP options are aligned on word boundaries the same as they appear in the DHCP packet. This option does not require configuration.

**Code**
0

**Length**
Not used.

**Value**
Not used.

### Structure of Pad Option

| Code |
| --- |
| 0 |

**End Option**
A single octet of decimal 255 (“FF”) used to indicate the end of a DHCP options area in DHCP messages. This option differs from most DHCP options because it does not use a length or value field. Typically, it is used at the end of the options field to indicate that there is no more option data in a DHCP message. It can also be used within the message, in connection with vendor-specific information (option 43), to indicate the end of an encapsulated vendor-specific options subfield. This option does not require configuration.

**Code**
255

**Length**
Not used.

**Value**
Not used.

### Structure of End Option

| Code |
| --- |
| 255 |

**Subnet Mask**
This option specifies the subnet mask of the client subnet, as described in RFC 950. The value for this option is taken from the **Subnet Mask** field, as defined in the **DHCP ScopeProperties** dialog box in DHCP Manager. A DHCP client running Windows XP or Windows Server 2003 requests this option.

**Code**
1

**Length**
Fixed, 4 octets.

**Value**
Unsigned 32-bit (4 octet) integer representing the subnet mask for an IP address provided in a DHCP message.

### Structure of Subnet Mask

| Code | Length | Subnet Mask |
| --- | --- | --- |
| 1 | 4 | *subnet mask in binary format* |

**Time Offset**
This option specifies an offset value (in seconds) from the Coordinated Universal Time (UTC) that applies to the client subnet. This value is configurable as a signed 32-bit integer. Positive offset values indicate a subnet location east of the zero meridian. Negative offset values indicate a subnet location west of the zero meridian.

**Code**
2

**Length**
Fixed, 4 octets.

**Value**
Signed 32-bit integer used for offset of UTC.

**Structure of Time Offset**

| Code | Length | Time Offset |
|---|---|---|
| 2 | 4 | *time* |

**Router**
This option specifies a list of IP addresses for routers on the client subnet. When more than one router is assigned, the client interprets and uses the addresses in the specified order. This option is normally used to assign a default gateway to DHCP clients on a subnet. A DHCP client running Windows XP or Windows Server 2003 requests this option.

**Code**
3

**Length**
Variable. Minimum length of 4 octets; octet length increases in multiples of 4 for each router address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each assigned router.

**Structure of Router**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 3 | *n* | *IP address in binary format* | *IP address in binary format* |

**Time Server**
This option specifies a list of IP addresses for time servers, as defined in RFC 868, that are available to the client. When more than one time server is assigned, the client interprets and uses the addresses in the specified order.

**Code**
4

**Length**
Variable. Minimum length of 4 octets; octet length increases in multiples of 4 for each time server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each assigned time server.

**Structure of Time Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 4 | *n* | *IP address in binary format* | *IP address in binary format* |

**IEN Name Server**
This option specifies a list of IP addresses for Internet Engineering Note (IEN) name servers available to the client. When more than one server is assigned, the client interprets and uses the addresses in the specified order.

**Code**
5

**Length**
Variable. Minimum length of 4 octets; octet length increases in multiples of 4 for each IEN name server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each assigned IEN name server.

**Structure of IEN Name Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 5 | *n* | *IP address in binary format* | *IP address in binary format* |

**DNS Servers**
This option specifies a list of IP addresses for DNS name servers available to the client. When more than one server is assigned, the client interprets and uses the addresses in the specified order. A DHCP client running Windows XP and Windows Server 2003 requests this option.

**Code**
6

**Length**
Variable. Minimum length of 4 octets; octet length increases in multiples of 4 for each Domain Name System (DNS) server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each assigned DNS server.

**Structure of DNS Servers**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 6 | *n* | *IP address in binary format* | *IP address in binary format* |

**Log Server**
This option specifies a list of IP addresses for Massachusetts Institute of Technology Lab for Computer Science (MIT-LCS) User Datagram Protocol (UDP) log servers available to the client. When more than one server is assigned, the client interprets and uses the addresses in the specified order.

**Code**
7

**Length**
Variable. Minimum length of 4 octets; octet length increases in multiples of 4 for each log server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each assigned log server.

**Structure of Log Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 7 | *n* | *IP address in binary format* | *IP address in binary format* |

**Cookie Server**
This option specifies a list of IP addresses for cookie servers, as defined in RFC 865, available to the client. When more than one server is assigned, the client interprets and uses the addresses in the specified order.

**Code**
8

**Length**
Variable. Minimum length of 4 octets; octet length increases in multiples of 4 for each cookie server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each assigned cookie server.

**Structure of Cookie Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 8 | *n* | *IP address in binary format* | *IP address in binary format* |

**LPR Server**
This option specifies a list of IP addresses for line printer (LPR) servers, as defined in RFC 1179, available to the client. When more than one server is assigned, the client interprets and uses the addresses in the specified order.

**Code**
9

**Length**
Variable. Minimum length of 4 octets; octet length increases in multiples of 4 for each LPR server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each assigned LPR server.

**Structure of LPR Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 9 | *n* | *IP address in binary format* | *IP address in binary format* |

**Impress Server**
This option specifies a list of IP addresses for Imagen Impress servers available to the client. When more than one server is assigned, the client interprets and uses the addresses in the specified order.

**Code**
10

**Length**
Variable. Minimum length of 4 octets; octet length increases in multiples of 4 for each Impress server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each assigned Impress server.

**Structure of Impress Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 10 | *n* | *IP address in binary format* | *IP address in binary format* |

**Resource Location Server**
This option specifies a list of IP addresses for resource location servers, as defined in RFC 887, available to the client. When more than one server is assigned, the client interprets and uses the addresses in the specified order.

**Code**
11

**Length**
Variable. Minimum length of 4 octets; octet length increases in multiples of 4 for each resource location server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each assigned resource location server.

**Structure of Resource Location Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 11 | *n* | *IP address in binary format* | *IP address in binary format* |

**Host Name**
This option specifies a host name for the client. In some cases, this name can also be fully qualified by appending the name value provided here with the DNS domain name, as specified in DHCP option 15. For Windows clients, this option is not supported for use when configuring the client's host name, which is set for computers running Windows XP or Windows Server 2003 on the **Computer Name** tab in the **System Properties** dialog box on the client computer.

**Code**
12

**Length**
Length varies depending on data in value. Minimum length is 1 octet. Maximum length is limited to 63 characters, or one octet for each character used in the host name configured for use with this option.

**Value**
ASCII character text.

**Structure of Host Name**

| Code | Length | Host Name |
|---|---|---|
| 12 | *n* | *name* |

**Boot File Size**
This option specifies the size of the default boot image file for the client.

**Code**

13

**Length**
Fixed, 2 octets.

**Value**
Unsigned 16-bit integer to indicate the number of 512-octet blocks needed to make up the boot file.

**Structure of Boot File Size**

| Code | Length | File Size |
|---|---|---|
| 13 | 02 | *16-bit integer* |

**Merit Dump File**
This option specifies the path name of a file to which the client's core memory image should be dumped in the event the client terminates abnormally. Data used for a value is in ASCII character text format. The length of the value field depends on the number of characters used in the path specified. For example, if the path entered has 20 characters, the value field for this option should also be 20 octets in length.

**Code**
14

**Length**
Length varies depending on data in value. Minimum length is 1 octet.

**Value**
ASCII character text.

**Structure of Merit Dump File**

| Code | Length | Dump File Path |
|---|---|---|
| 14 | *n* | *path name* |

**DNS Domain Name**
This option specifies the domain name that the DHCP client should use when resolving unqualified domain names with the Domain Name System (DNS). For DHCP clients running Windows 2000, Windows XP, and Windows Server 2003, the DNS Domain Name option becomes the connection-specific DNS name assigned to the DHCP-configured interface. The connection-specific DNS name is used to construct fully qualified domain names (FQDNs) that are registered using DNS dynamic update. The length of the value field depends on the number of characters used in the DNS domain name specified. For example, if the domain name has 20 characters, the value field for this option is 20 octets in length. A DHCP client running Windows XP or Windows Server 2003 requests this option.

**Code**
15

**Length**
Length varies depending on data in value. Minimum length is 1 octet.

**Value**
ASCII character text.

**Structure of DNS Domain Name**

| Code | Length | Domain Name |
|---|---|---|
| 15 | *n* | *domain name* |

**Swap Server**
This option specifies the IP address of the client's swap server.

**Code**
16

**Length**
Length is fixed at 4 octets.

**Value**
A single IP address for the client's swap server (unsigned 32-bit integer).

### Structure of Swap Server

| Code | Length | Swap Server Address |
|---|---|---|
| 16 | 4 | *IP address in binary format* |

**Root Path**
This option specifies the path of the client's root disk. The path value is formatted as ASCII text. The length of the value field depends on the number of characters used in the root path specified. For example, if the root path entered has 20 characters, the value field for this option is 20 octets in length.

**Code**
17

**Length**
Length varies depending on data in value. Minimum length is 1 octet.

**Value**
ASCII character text.

### Structure of Root Path

| Code | Length | Root Disk Path |
|---|---|---|
| 17 | *n* | *path name* |

**Extensions Path**
This option specifies an extension path file that can be retrieved using Trivial File Transfer Protocol (TFTP). The file contains information to be interpreted as the 64-octet vendor-extension field within a BOOTP response. To allow more than 64 octets of BOOTP vendor extension information, this option can be enabled. When enabled, the length of the specified extension path file is not constrained in size and all references in the extensions file to tag 18 (such as instances of the BOOTP Extensions Path field) are ignored.

**Code**
18

**Length**
Length varies depending on data in value. Minimum length is 1 octet.

**Value**
ASCII character text.

### Structure of Extension Path

| Code | Length | Extensions Path |
|---|---|---|
| 18 | *n* | *file name* |

**IP Host Options**
The following sections describe the DHCP options that affect the operation of the IP layer, on a per-host basis.

**IP Forwarding Enable/Disable**
This option specifies whether the DHCP client should enable or disable forwarding of datagrams at the IP layer.

**Code**
19

**Length**
Length is fixed at 1 octet.

**Value**
Accepted values for this option are:

- 1 = Enable IP forwarding

- 0 = Disable IP forwarding

**Structure of IP Forwarding Enable/Disable**

| Code | Length | Value |
|---|---|---|
| 19 | 1 | 0 \| 1 |

**Non-Local Source Routing Enable/Disable**
This option specifies whether the DHCP client enables or disables the forwarding at the IP layer of datagrams that contain source routing information and were sent by a non-local host.

**Code**
20

**Length**
Length is fixed at 1 octet.

**Value**
Accepted values for this option are:

- 1 = Enable forwarding of datagrams from non-local sources
- 0 = Disable forwarding of datagrams from non-local sources

**Structure of Non-Local Source Routing Enable/Disable**

| Code | Length | Value |
|---|---|---|
| 20 | 1 | 0 \| 1 |

**Policy Filter**
This option specifies policy filters for non-local source routing on the client. The filters consist of a list of IP address and mask pairs specifying destination and mask pairs for which incoming datagrams should be source-route filtered. The client discards any source routed datagram with a next-hop address that does not match one of the filters. For further information about policy filtering as it applies to this option, see RFC 1122 in the IETF RFC Database.

**Code**
21

**Length**
Variable. Minimum length is 8 octets for a single destination and mask pair. Length increases in multiples of 8 octets for each additional pair used.

**Value**
Two consecutive, unsigned 32-bit integers indicating a paired value, consisting of an IP address followed by a subnet mask.

**Structure of Policy Filter**

| Code | Length | Address 1 | Subnet Mask 1 | Address N | Subnet Mask N |
|---|---|---|---|---|---|
| 21 | *n* | *IP address in binary format* | *subnet mask in binary format* | *IP address in binary format* | *subnet mask in binary format* |

**Maximum Datagram Reassembly Size**
This option specifies the maximum size datagram that the client needs to be prepared to reassemble.

**Code**
22

**Length**
Fixed, 2 octets.

**Value**
Unsigned 16-bit integer specifying the maximum datagram size for reassembly. The minimum size for a datagram is 576

octets.

**Structure of Maximum Datagram Reassembly Size**

| Code | Length | File Size |
|---|---|---|
| 22 | 2 | *16-bit integer* |

**Default IP Time-To-Live**
This option specifies the default Time-To-Live (TTL) that the client uses for the datagrams it sends.

**Code**
23

**Length**
Fixed, 1 octet.

**Value**
A number (in seconds) between 1 and 255.

**Structure of Default IP Time-To-Live**

| Code | Length | TTL |
|---|---|---|
| 23 | 1 | *TTL value in seconds* |

**Path MTU Aging Time-out**
This option specifies the time-out for aging Path Maximum Transmission Unit (MTU) values. Values are found by the Path MTU discovery process, as defined in RFC 1191.

**Code**
24

**Length**
Fixed, 4 octets.

**Value**
A number (in seconds) that specifies a time-out value.

**Structure of Path MTU Aging Time-out**

| Code | Length | Time-out |
|---|---|---|
| 24 | 4 | *time-out value in seconds* |

**Path MTU Plateau Table**
This option specifies a table of MTU sizes to use when performing Path MTU discovery, as defined in RFC 1191.

**Code**
25

**Length**
Variable. Minimum length is 2 octets and increases in multiples of 2.

**Value**
A table formatted as a list of 16-bit unsigned integers, ordered from smallest to largest. The minimum tabled MTU value cannot be smaller than **68**.

**Structure of Path MTU Plateau Table**

| Code | Length | MTU Size 1 | MTU Size N |
|---|---|---|---|
| 25 | *n* | *MTU size in binary format* | *MTU size in binary format* |

**IP Interface Options**
The following sections describe the DHCP options that affect operation of the IP layer on a per-interface basis.

**Interface MTU**

This option specifies the MTU size that can be used on a specified host adapter interface.

**Code**
26

**Length**
Fixed, 2 octets.

**Value**
A 16-bit unsigned integer specifying the interface MTU. The minimum value for the MTU is **68**.

**Structure of Interface MTU**

| Code | Length | MTU |
|---|---|---|
| 26 | 2 | *interface MTU* |

**All Subnets Are Local**
This option specifies whether the client assumes that all subnets within the client's internetwork use the same MTU size as the local subnet on which the client is connected.

**Code**
27

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option are:

- 1 = Clients assume all subnets are local and share the same MTU size
- 0 = Clients assume some subnets are not local and that smaller MTU sizes might be in use on remote subnets

**Structure of All Subnets Are Local**

| Code | Length | Value |
|---|---|---|
| 27 | 1 | 0 \| 1 |

**Broadcast Address**
This option specifies the broadcast address used on the client's subnet.

**Code**
28

**Length**
Fixed, 4 octets.

**Value**
Typically, the limited broadcast IP address (255.255.255.255), but can be modified using legal values for broadcast addresses, as specified in section 3.2.1.3 of RFC 1122.

**Structure of Broadcast Address**

| Code | Length | Broadcast Address |
|---|---|---|
| 28 | 4 | *broadcast address in binary format* |

**Perform Mask Discovery**
This option specifies whether the client uses Internet Control Message Protocol (ICMP) for subnet mask discovery.

**Code**
29

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option are:

- 1 = Client performs subnet mask discovery
- 0 = Client does not perform subnet mask discovery

### Structure of Perform Mask Discovery

| Code | Length | Value |
|---|---|---|
| 29 | 1 | 0 \| 1 |

**Mask Supplier**
This option specifies whether the client responds to subnet mask requests using ICMP.

**Code**
30

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option are:

- 1 = Client responds to subnet mask requests
- 0 = Client does not respond to subnet mask requests

### Structure of Mask Supplier

| Code | Length | Value |
|---|---|---|
| 30 | 1 | 0 \| 1 |

**Perform Router Discovery**
This option specifies whether the client solicits routers using the router discovery method in RFC 1256. A DHCP client running Windows XP or Windows Server 2003 requests this option.

**Code**
31

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option are:

- 1 = Client performs router discovery
- 0 = Client does not perform router discovery

### Structure of Perform Router Discovery

| Code | Length | Value |
|---|---|---|
| 31 | 1 | 0 \| 1 |

**Router Solicitation Address**
This option specifies the IP address to which the client submits router solicitation requests.

**Code**
32

**Length**
Fixed, 4 octets.

**Value**
Unsigned 32-bit integer representing an IP address.

### Structure of Router Solicitation Address

| Code | Length | Address |
|---|---|---|
| 32 | 4 | *IP address in binary format* |

**Static Routes**
This option specifies a list of classful static routes that the DHCP client automatically adds to its IP routing table. Multiple routes to the same destination are listed in descending order of priority. The default route of 0.0.0.0 is an illegal destination for a static route. To configure the default route, use the Router DHCP option to assign a default gateway. A DHCP client running Windows XP or Windows Server 2003 requests this option.

**Code**
33

**Length**
Variable. Minimum length of 8 octets; octet length increases in multiples of 8 for each additional static route provided with this option.

**Value**
A list of IP address pairs. Each set of 8 octets provides two consecutive IP addresses pairing the destination (as a classful network ID) and a router address (the IP address of the router interface on the subnet to which this scope-specific option is configured) used for each route. The first 4 octets specify the destination classful network ID, and the second 4 octets specify the router IP address.

### Structure of Static Routes

| Code | Length | Destination 1 | Router 1 | Destination N | Router N |
|---|---|---|---|---|---|
| 33 | *n* | *IP address in binary format* | *IP address in binary format* | *IP address in binary format* | *IP address in binary format* |

**Classless Static Route**
This option specifies a list of classful static routes that the client automatically adds to its IP routing table. A DHCP client running Windows XP or Windows Server 2003 requests this option. The Classless Static Route can be used to configure split tunneling for remote access virtual private network (VPN) clients running Windows XP and Windows Server 2003. For more information, see "VPN Technical Reference."

**Code**
249

### Note

- This is the same as option 121, as defined in RFC 3442.

**Length**
Variable. Minimum length of 5 octets; maximum octet length depends on the nature and number of the classless static routes. Each route entry includes a destination descriptor and a router and can vary from a minimum of 5 octets in length to maximum 9 octets in length. For more information about how classless static routes are constructed using a destination descriptor, see RFC 3442 in the IETF RFC Database.

**Value**
The destination descriptor and the router IP address. The value is encoded. For the encoding scheme, see RFC 3442.

### Structure of Classless Static Route

| Code | Length | Destination 1 | Router 1 | Destination N | Router N |
|---|---|---|---|---|---|
| 249 | *n* | *Destination descriptor, as defined in RFC 3442* | *IP address in binary format* | *Destination descriptor, as defined in RFC 3442* | *IP address in binary format* |

**Link Layer Options**
The following sections describe the DHCP options that affect operation of the data link layer on a per-interface basis.

**Trailer Encapsulation**
This option specifies whether the client negotiates the use of trailers, as described in RFC 893, when using the Address Resolution Protocol (ARP).

**Code**
34

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option are:

- 1 = Client attempts to use trailer
- 0 = Client does not attempt to use trailers

### Structure of Trailer Encapsulation

| Code | Length | Value |
|---|---|---|
| 34 | 1 | 0 \| 1 |

**ARP Cache Time-Out**
This option specifies the time-out for ARP cache entries.

**Code**
35

**Length**
Fixed, 4 octets.

**Value**
An unsigned 32-bit integer specifying a time-out value, in seconds.

### Structure of ARP Cache Time-Out

| Code | Length | Time |
|---|---|---|
| 35 | 4 | *time-out value in seconds* |

**Ethernet Encapsulation**
This option specifies whether the client uses Ethernet II (RFC 894) or IEEE 802.3 (RFC 1042) encapsulation (if the interface is Ethernet).

**Code**
36

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option are:

- 1 = Client uses RFC 1042 encapsulation
- 0 = Client uses RFC 894 encapsulation

### Structure of Ethernet Encapsulation

| Code | Length | Value |
|---|---|---|

| | | |
|---|---|---|
| 36 | 1 | 0 \| 1 |

**TCP Options**
The following sections describe the DHCP options that affect operation of the TCP layer on a per-interface basis.

**TCP Default TTL**
This option specifies the default TTL that the client uses when sending TCP segments.

**Code**
37

**Length**
Fixed, 1 octet.

**Value**
An unsigned 8-bit integer specifying a Time-To-Live (TTL) value in seconds. The minimum TTL value is **1**.

### Structure of TCP Default TTL

| Code | Length | TTL |
|---|---|---|
| 37 | 1 | *TTL value in seconds* |

**TCP Keep-Alive Interval**
This option specifies the interval the client waits before sending a keep-alive message on a TCP connection. A value of **0** indicates that the client should not send keep-alive messages on connections unless specifically requested by an application.

**Code**
38

**Length**
Fixed, 4 octets.

**Value**
An unsigned 32-bit integer that specifies a keep-alive interval, in seconds.

### Structure of TCP Keep-Alive Interval

| Code | Length | Time |
|---|---|---|
| 38 | 4 | *keep-alive interval in seconds* |

**TCP Keep-Alive Garbage**
This option specifies whether or not the client sends TCP keep-alive messages with an octet of garbage data for compatibility with older implementations.

**Code**
39

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option are:

- 1 = Client sends keep-alive garbage octet
- 0 = Client does not send keep-alive garbage octet

### Structure of TCP Keep-Alive Garbage

| Code | Length | Value |
|---|---|---|
| 39 | 1 | 0 \| 1 |

**Application Layer Options**

The following sections describe the DHCP options that affect the application layer operations on a per-interface basis. These are miscellaneous options used to configure programs and services.

**NIS Domain Name**
This option specifies the Network Information Service (NIS) domain name as an ASCII string.

**Code**
40

**Length**
Length varies depending on data in value. Minimum length is 1 octet.

**Value**
ASCII character text.

### Structure of NIS Domain Name

| Code | Length | NIS Domain Name |
|---|---|---|
| 40 | *n* | *NIS domain name* |

**NIS Servers**
This option lists the IP addresses in order of preference for Network Information Service (NIS) servers available to the client.

**Code**
41

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each NIS server.

### Structure of NIS Servers

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 41 | *n* | *IP address in binary format* | *IP address in binary format* |

**NTP Servers**
This option lists the IP addresses in the order of preference for Network Time Protocol (NTP) servers available to the client.

**Code**
42

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server address listed.

**Value**
Unsigned 32-bit integer representing the IP address of each NTP server.

### Structure of NTP Servers

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 42 | *n* | *IP address in binary format* | *IP address in binary format* |

**X Window System Font Servers**
This option lists the IP addresses in the order of preference for X Window System font servers available to the client.

**Code**
48

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server address listed.

**Value**
Unsigned 32-bit integer representing each server IP address.

**Structure of X Window System Font Servers**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 48 | *n* | *IP address in binary format* | *IP address in binary format* |

**X Window System Display Manager Servers**
This option lists the IP addresses in the order of preference for X Window System display manager servers available to the client.

**Code**
49

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server address listed.

**Value**
Unsigned 32-bit integer representing each server IP address.

**Structure of X Window System Display Manager Servers**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 49 | *n* | *IP address in binary format* | *IP address in binary format* |

**NIS+ Domain Name**
This option specifies the name of the client's Network Information Service Plus (NIS+) domain name as an ASCII string.

**Code**
64

**Length**
Length varies depending on the data in its value. Minimum length is 1 octet.

**Value**
ASCII character text.

**Structure of NIS+ Domain Name**

| Code | Length | NIS+ Domain Name |
|---|---|---|
| 64 | *n* | *NIS+ domain name* |

**NIS+ Servers**
This option lists the IP addresses in the order of preference for Network Information Service Plus (NIS+) servers available to the client.

**Code**
65

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server address listed.

**Value**
Unsigned 32-bit integer representing IP address of NIS+ servers.

**Structure of NIS+ Servers**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 65 | *n* | *IP address in binary format* | *IP address in binary format* |

**Mobile IP Home Agents**
This option lists the IP addresses in the order of preference for mobile IP home agents available to the client.

**Code**
68

**Length**

Variable. Minimum length is 0 octets. A length of 0 octets signifies that no mobile IP home agents are available. Octet length increases in multiples of 4 for each mobile IP home agent address listed.

**Value**
Unsigned 32-bit integer representing IP address of a mobile IP home agent.

**Structure of Mobile IP Home Agents**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 68 | *n* | *IP address in binary format* | *IP address in binary format* |

**Simple Mail Transfer Protocol (SMTP) Server**
This option lists the IP addresses in order of preference for SMTP servers available to the client.

**Code**
69

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server IP address listed.

**Value**
Unsigned 32-bit integers representing IP addresses of servers.

**Structure of Simple Mail Transfer Protocol (SMTP) Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 69 | *n* | *IP address in binary format* | *IP address in binary format* |

**Post Office Protocol 3 (POP3) Server**
This option lists the IP addresses in order of preference for POP3 servers available to the client.

**Code**
70

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server IP address listed.

**Value**
Unsigned 32-bit integers representing IP addresses of servers.

**Structure of Post Office Protocol 3 (POP3) Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 70 | *n* | *IP address in binary format* | *IP address in binary format* |

**Network News Transport Protocol Server**
This option lists the IP addresses in order of preference for NNTP servers available to the client.

**Code**
71

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server IP address listed.

**Value**
Unsigned 32-bit integers representing IP addresses of servers.

**Structure of Network News Transport Protocol Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 71 | *n* | *IP address in binary format* | *IP address in binary format* |

**Default World Wide Web Server**
This option lists the IP addresses in order of preference for default Web servers available to the client.

**Code**
72

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server IP address listed.

**Value**
Unsigned 32-bit integers representing IP addresses of servers.

**Structure of Default World Wide Web Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 72 | *n* | *IP address in binary format* | *IP address in binary format* |

**Default FFinger Server**
This option lists the IP addresses in order of preference for default Finger servers available to the client.

**Code**
73

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server IP address listed.

**Value**
Unsigned 32-bit integers representing IP addresses of servers.

**Structure of Default Finger Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 73 | *n* | *IP address in binary format* | *IP address in binary format* |

**Default Internet Relay Chat Server**
This option lists the IP addresses in order of preference for default IRC servers available to the client.

**Code**
74

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server IP address listed.

**Value**
Unsigned 32-bit integers representing IP addresses of servers.

**Structure of Default Internet Relay Chat Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 74 | *n* | *IP address in binary format* | *IP address in binary format* |

**StreetTalk Server**
This option lists the IP addresses in order of preference for StreetTalk servers available to the client.

**Code**
75

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server IP address listed.

**Value**
Unsigned 32-bit integers representing IP addresses of servers.

**Structure of StreetTalk Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|

| | | | |
|---|---|---|---|
| 75 | *n* | *IP address in binary format* | *IP address in binary format* |

**StreetTalk Directory Assistance Server**
This option lists the IP addresses in order of preference for StreetTalk Directory Assistance (STDA) servers available to the client.

**Code**
76

**Length**
Variable. Minimum length is 4 octets; octet length increases in multiples of 4 for each server IP address listed.

**Value**
Unsigned 32-bit integers representing IP addresses of servers.

**Structure of StreetTalk Directory Assistance Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 76 | *n* | *IP address in binary format* | *IP address in binary format* |

**NetBIOS over TCP/IP Options**
The following options are used to support NetBIOS over TCP/IP. All Microsoft-based DHCP clients and DHCP servers can recognize and support the use of these options.

**WINS/NBNS Servers**
This option lists the IP addresses for NetBIOS name servers (NBNSes) on the network. In Windows, a Windows Internet Name Service (WINS) server is an NBNS. A DHCP client running Windows XP or Windows Server 2003 requests this option.

**Code**
44

**Length**
Variable. Minimum length is 4 octets; length can be increased by multiples of 4 for each address listed.

**Value**
Each 4 octets in this field contains an NBNS server IP address, specified as an unsigned 32-bit integer.

**Structure of WINS/NBNS Servers**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 44 | *n* | *IP address in binary format* | *IP address in binary format* |

**NetBIOS over TCP/IP Datagram Distribution Server**
This option lists the IP addresses for NetBIOS datagram distribution (NBDD) servers.

**Code**
45

**Length**
Variable. Minimum length is 4 octets; length can be increased only by multiples of 4.

**Value**
Each 4 octets in this field contains an NBDD server IP address, specified as an unsigned 32-bit integer.

**Structure of NetBIOS over TCP/IP Datagram Distribution Server**

| Code | Length | Address 1 | Address N |
|---|---|---|---|
| 45 | *n* | *IP address in binary format* | *IP address in binary format* |

**NetBIOS over TCP/IP Node Type**
Configures the client node type for NetBIOS over TCP/IP (NetBT) clients, as described in RFCs 1001and 1002. On multihomed computers, the node type is assigned for the computer, not to individual network adapters. A DHCP client running Windows XP or Windows Server 2003 requests this option.

**Code**

46

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option are:

- 1 = b-node
- 2 = p-node
- 4 = m-node
- 8 = h-node

**Structure of NetBIOS over TCP/IP Node Type**

| Code | Length | Value |
|---|---|---|
| 46 | 1 | 1 \| 2 \| 4 \| 8 |

**NetBIOS over TCP/IP Scope ID**
This option specifies a string representing the NetBIOS over TCP/IP scope ID for the client, as specified in RFCs 1001 and 1002. On multihomed computers, the node type is assigned for the computer, not to individual network adapters. A DHCP client running Windows XP or Windows Server 2003 requests this option. The use of NetBIOS scope IDs is not recommended.

**Code**
47

**Length**
Variable. Minimum length is 1 octet. Octet length is equal to the number of characters used in NetBIOS scope ID.

**Value**
This option specifies the NetBIOS over TCP/IP scope identifier used by the client. The format used for these scope IDs is described in RFCs 1001 and 1002. For character-set restrictions, see the RFCs.

**Structure of NetBIOS over TCP/IP Scope ID**

| Code | Length | NetBIOS Scope |
|---|---|---|
| 47 | *n* | *scope identifier* |

**Vendor-Specific Options**
This section describes reserved DHCP options specified for vendor class use. The vendor-specific options are specified in RFC 2132. Vendor classes can be used by the DHCP service and DHCP clients running Windows 2000, Windows XP, and Windows Server 2003. For any other DHCP clients, default classes provided by the DHCP service can be used to group and classify non-identifying clients at the DHCP server.

The Windows Server 2003 DHCP snap-in provides a default vendor class called *DHCP Standard*. This class can be used to group and classify clients that do not identify a vendor class to the DHCP service.

**Vendor-Specific Information**
This option is used by clients and servers to exchange vendor-specific information. Servers not equipped to interpret the information ignore it. Clients that expect but do not receive the information attempt to operate without it.

In some cases, a vendor uses this option to send more than one information item; therefore, this option can serve as a subfield for encapsulating vendor-specific options. When encapsulating options, DHCP servers maintain the same syntax (the same sequence of code, length, and value fields) for each encapsulated option as it would normally appear in the full standard options field. The following exceptions are for the encapsulated, vendor-specific subfield:

- A “Magic cookie” field cannot be used.
- All standard option codes — other than the padding option (0) or the end option (255) — can be redefined, but should

conform to the code, length, value syntax sequence described in RFC 2131.

- If present, the end option (255) signifies the end of the encapsulated vendor options, but not the end of the encapsulated vendor-specific subfield. If no end option is present, the end for the encapsulated vendor-specific subfield is taken from its stated length. For more information, see RFC 2132 in the IETF RFC Database.

**Code**
43

**Length**
Variable. Minimum length is 1 octet.

**Value**
An object of *n* octets (where *n* is equal to the length specified with this option). The definition of values stored for this option is vendor specific, and values provided here are presumed to be interpreted by vendor-specific code on DHCP clients and the DHCP server.

**Structure of Vendor-Specific Information**

| Code | Length | Value |
|---|---|---|
| 43 | *n* | *vendor-specific information* (which can include subfield *bytes 1-n*) |

When this option uses an encapsulated vendor-specific subfield, the information bytes 1–*n* have the following format.

| Code | Length | Data Item |
|---|---|---|
| T1 | *n* | *d1*, *d2*, ...*dn* |
| T2 | *n* | *d1*, *d2*, *dn* |

**Vendor Class Identifier**
Can be used by DHCP clients to identify their vendor type and configuration when communicating with DHCP servers. Vendors can define their own specific identifier values, like conveying a particular hardware or operating system configuration.

For Windows Server 2003, all computers that function either as DHCP servers or DHCP clients can use and support this option. When vendor classes are used, the DHCP server responds to identifying clients by using option 43, the reserved option for returning vendor-specific information to the client.

DHCP servers that do not automatically interpret this option are expected to ignore it. For earlier, Windows-based clients and other clients that do not support this option, the DHCP service classifies these clients as part of the default vendor class — the DHCP Standard option class — predefined for Windows-based DHCP servers.

**Code**
60

**Length**
Variable. Minimum is 1 octet; Length varies according to *n* (the number of octets used as an identifier).

**Value**
A value of *n* octets, which can be interpreted by DHCP servers that support vendor-specific classification of clients.

**Structure of Vendor Class Identifier**

| Code | Length | Value |
|---|---|---|
| 60 | *n* | *vendor class identifier* |

**User Class Options**
This section describes reserved DHCP options specified for user classes. User classes can be used by the DHCP service and DHCP-enabled client computers running Windows Server 2003. For other DHCP clients, default classes provided by the DHCP service can be used to group and classify non-identifying clients at the DHCP server.

**User Class Information**
A DHCP client can use this option to identify the user class of which it is a member when communicating with the DHCP server. The information contained in this option is a Network Virtual Terminal (NVT) ASCII text object that represents the user class ID.

You can use the DHCP snap-in to define specific user classes. When user classes are created, each class sets an identifying string of information to be used by the DHCP service to classify identifying clients. Also, a default user class is created for classifying clients that are unable to support a user class ID.

User classes can be helpful for separating client computers that have a shared or common need for similar software configuration or user preferences. For example, an identifier can specify that a particular DHCP client be a member of the class "accounting auditors," who have special service needs, such as a particular database server.

Computers running Windows Server 2003 support sending or using this option. Legacy DHCP clients do not send a class ID and cannot recognize DHCP user class IDs. Such client's are assigned to the Default User Class, a user class predefined for immediate use in the DHCP snap-in. Other user classes must be manually created.

**Code**
77

**Length**
Variable. Minimum is 2 octets.

**Value**
ASCII character text.

**Structure of User Class Information**

| Code | Length | User Class Information |
|---|---|---|
| 77 | *n* | *c1*, *c2*, *c3*, *c4 ...cn* |

**DHCP Extensions**
The following options are specific to DHCP and are used to implement default protocol interaction and system behavior between servers and clients. Some of these options are implicitly set when you configure server and scope properties using the DHCP snap-in.

**Requested IP Address**
This option can be used by clients when sending a DHCPDiscover message to request a specific IP address from the DHCP server.

**Code**
50

**Length**
Fixed, 4 octets.

**Value**
Single, unsigned 32-bit integer representing a requested IP address.

**Structure of Requested IP Address**

| Code | Length | Requested IP Address |
|---|---|---|
| 50 | *4* | *IP address in binary format* |

**IP Address Lease Time**
This option is used to negotiate and exchange lease-time information between DHCP clients and servers in two possible ways. First, the option can be used in a DHCPDiscover or DHCPRequest message sent by a client to request a lease time for its IP address. Second, the option can be used in a DHCPOffer message sent by a server to specify a lease time for the client. This option is configured on the **Scope Properties** dialog box.

**Code**
51

**Length**
Fixed, 4 octets.

**Value**
Single, unsigned 32-bit integer representing a clients lease time (in seconds).

**Structure of IP Address Lease Time**

| Code | Length | Lease Time |
|---|---|---|
| 51 | 4 | *lease time in seconds* |

**Option Overload**
Used in messages sent by a DHCP server to indicate that either of the standard message fields in a DHCP packet for *server_host_name* (**sname**) and *boot_file_name* (**file**) can be used to hold options (a condition also known as *overloaded*).

When this option is used, it extends the options area in each packet by indicating that unused space for one or both of these two standard fields should be allocated to the area used to carry DHCP options.

**Code**
52

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option include:

- 1 = **File** field is overloaded
- 2 = **Sname** field is overloaded
- 3 = Both **file** and **sname** fields are overloaded

### Structure of Option Overload

| Code | Length | Value |
|---|---|---|
| 52 | 1 | 1 \| 2 \| 3 |

**TFTP Server Name**
This option specifies the host name of the Trivial File Transfer Protocol (TFTP) server when the *server_host_name* (**sname**) field in a DHCP message is overloaded and used for carrying additional DHCP options.

**Code**
66

**Length**
Variable depending on data in value. Minimum length is 1 octet.

**Value**
ASCII character text.

### Structure of TFTP Server Name

| Code | Length | TFTP Server |
|---|---|---|
| 66 | *n* | *TFTP server host name* |

**Boot File Name**
This option specifies the name of a boot image file on the TFTP server when the *boot_file_name* (**file**) field in a DHCP message is overloaded and used for carrying additional DHCP options.

**Code**
67

**Length**
Variable depending on data in value. Minimum length is 1 octet.

**Value**
ASCII character text.

### Structure of Boot File Name

| Code | Length | Boot File Name |
|---|---|---|
| 67 | *n* | *boot image file name* |

**DHCP Message Type**
This option is required for use in all DHCP messages to convey the type of message being sent.

**Code**
53

**Length**
Fixed, 1 octet.

**Value**
Accepted values for this option are:

- 1 = DHCP Discover message (DHCPDiscover)
- 2 = DHCP Offer message (DHCPOffer)
- 3 = DHCP Request message (DHCPRequest)
- 4 = DHCP Decline message (DHCPDecline)
- 5 = DHCP Acknowledgment message (DHCPAck)
- 6 = DHCP Negative Acknowledgment message (DHCPNak)
- 7 = DHCP Release message (DHCPRelease)
- 8 = DHCP Informational message (DHCPInform)

### Structure of DHCP Message Type

| Code | Length | Value |
|---|---|---|
| 53 | 1 | 1 \| 2 \| 3 \| 4 \| 5 \| 6 \| 7 \| 8 |

**Server Identifier**
This option is used in DHCPOffer and DHCPRequest messages, and might appear in DHCP Acknowledgment messages (DHCPAck, DHCPNak). The server identifier is the IP address of a selected DHCP server.

This option is used in three possible ways. First, servers include this option in DHCPOffer messages so that clients can distinguish between multiple lease offers. Second, clients include this option in DHCPRequest messages to select a lease and indicate which offer is accepted from multiple lease offers. Third, clients can use the contents of this option for unicast transmission of DHCPRequest messages to specific DHCP servers to renew a current lease.

**Code**
54

**Length**
Fixed, 4 octets.

**Value**
A single, unsigned 32-bit integer representing the IP address that identifies the DHCP server.

### Structure of Server Identifier

| Code | Length | Address |
|---|---|---|
| 54 | 4 | *IP address in binary format* |

**Parameter Request List**

Used by a DHCP client to request specific option values from the DHCP server. Each option is requested and listed by a single octet value containing a valid or recognized DHCP option Code for the server.

For clients that use this option, the list can be ordered by preference. Although the DHCP server is not required to return options in the order they are requested it does attempt to insert the requested options in the requested order.

**Code**
55

**Length**
Variable. Minimum of 1 octet; length increases by 1 octet for each option code included in the request list.

**Value**
List of 8-bit values, each representing an option code between 0 and 255.

**Structure of Parameter Request List**

| Code | Length | Option Codes |
|---|---|---|
| 55 | *n* | *c1*, *c2*...*cn* |

**Optional Message**
Can be used by both the DHCP server and DHCP clients in the following ways:

- A server can use this option to provide and embed an error message in a DHCP Negative Acknowledgment (DHCPNak) message in the event of a failure.
- A client can use this option in a DHCPDecline message to indicate why it declined offered parameters.

The message consists of a variable-length ASCII text string, which the receiving computer can then either log or display.

**Code**
56

**Length**
Minimum of 1 octet. Length depends on the length of the sent message.

**Value**
ASCII character text.

**Structure of Optional Message**

| Code | Length | Text |
|---|---|---|
| 56 | *n* | *c1*, *c2*...*cn* |

**Maximum Message Size**
Used by client to specify the maximum length for a DHCP message it can accept. A client can include this option in DHCPDiscover or DHCPRequest messages; however, it does not include this option in DHCPDecline messages.

**Code**
57

**Length**
Fixed, 2 octets.

**Value**
A 16-bit integer indicating the maximum size, in octets, for a DHCP message. The minimum value for this option is **576**.

**Structure of Maximum Message Size**

| Code | Length | Value |
|---|---|---|
| 57 | 2 | *maximum size* |

**Renewal Time Value (T1)**

The time value of this option is typically 50 percent of the client's lease time. To adjust the value, change the length of the client lease in the client scope properties or the per-user class on the DHCP server. You can also change the value by using the Netsh tool.

**Code**
58

**Length**
Fixed, 4 octets.

**Value**
A 32-bit unsigned integer indicating the number of seconds before the client begins to renew its address lease with the DHCP server.

**Structure of Renewal Time Value (T1)**

| Code | Length | T1 Interval |
|---|---|---|
| 58 | 4 | *renewal interval in seconds* |

**Rebinding Time Value (T2)**
The time value of this option is typically 87.5 percent of the full configured duration (lease time) for a client's lease. To adjust the value, change the length of the client lease in the properties for the client's scope or per-user class on the DHCP server. You can also change the value by using the Netsh tool.

**Code**
59

**Length**
Fixed, 4 octets.

**Value**
A 32-bit unsigned integer indicating the number of seconds before the client enters the rebinding state (if it has not renewed its current address lease with the DHCP server).

**Structure of Rebinding Time Value (T2)**

| Code | Length | T2 Interval |
|---|---|---|
| 59 | 4 | *rebinding interval in seconds* |

**Client Unique Identifier**
This option is used by clients to specify their unique identifier to the server. This option is most useful for reserved clients. When a reserved client contacts the server, the DHCP service checks and matches the client identifier value to a corresponding identifier used to configure an address reservation in the server's database. When a matching reservation is found, the DHCP server returns the reserved address and its related parameters to the correct client.

Each clients identifier must be unique among all other client identifiers used on the DHCP client's local subnet and any remote subnets reachable using DHCP relay. Vendors and system administrators are responsible for choosing client identifiers that meet this requirement for uniqueness. One way to ensure unique values is to use the client's media access control (MAC) address as the client identifier value. Media access control addresses are encoded in the client's network adapter hardware, and are assigned to hardware manufacturers in such a way as to ensure that they are unique for each device.

**Code**
61

**Length**
Variable length; minimum length is 2 octets.

**Value**
A series of 2 or more octets treated as a single value by the DHCP server. Servers can interpret and use this value to uniquely identify clients. The client identifier can consist of type-value pairs similar to the DHCP Header fields "htype" and "chaddr," which are defined by the DHCP protocol.

**Structure of Client Unique Identifier**

| Code | Length | Type | Client Identifier |
|---|---|---|---|
| 61 | *n* | *t1* | *i1*, *i2*...*in* |

**Administrator-Defined Options**
This section describes DHCP options that are reserved and specified for use by RFC 2132, but are not predefined for use in the DHCP snap-in. Administrators can add these options to support DHCP clients that recognize these options.

**Proxy Autodiscovery for Internet Explorer 5**
This option points to the configuration file that the client uses for automatic configuration of Internet Explorer 5.*x* and Internet Explorer 6. The configuration file can be a .pac, .jvs, .js, or .ins file created by a system or Web administrator when deploying Internet Explorer 5.*x* on an intranet. It might include settings for other Internet Explorer configurable options, such as which home page to use, or settings for locating and using a proxy server.

The option is communicated between Internet Explorer client computers and the DHCP server using the DHCPInform message. DHCPInform is currently supported for DHCP servers running Windows Server 2003 and DHCP clients running Windows 2000, Windows XP, and Windows Server 2003.

The use of additional DHCP configuration is supported by Internet Explorer 5.*x* and Internet Explorer 6, but not earlier versions, because those use different methods for automatic detection and configuration of proxy server settings.

You can add and configure an alias (CNAME) resource record at the DNS server to support Internet Explorer proxy server auto-discovery and configuration features.

**Code**
252

**Length**
Variable

**Value**
A URL that points to the configuration file that the client should use for automatic configuration of Internet Explorer 5.*x* and Internet Explorer 6.

**Structure of Proxy Autodiscovery for Internet Explorer 5**

| Code | Length | Value |
|---|---|---|
| 252 | *n* | *URL name* |

**Microsoft Options**
This section describes reserved DHCP options defined by Microsoft. These options are only available for use with supported Windows-based DHCP clients, such as computers running Windows Server 2003.

The Microsoft options are provided as encapsulated vendor-specific data fields within the vendor-specific information option.

Currently, administrators can assign these options are by using the DHCP snap-in through the following vendor classes: Microsoft options and Microsoft Windows 2000 options.

**Disable NetBIOS over TCP/IP (NetBT)**
This option can be used to selectively enable or disable NetBT for DHCP-enabled computers running Windows Server 2003, Windows XP, and Windows 2000. By default, if this option is not present, Windows Server 2003, Windows XP, and Windows 2000 enable the use of NetBT for network connections that are configured to use TCP/IP. Earlier Windows-based clients require NetBT and do not support this option.

**Code**
1

**Length**
4 octets

**Value**
Accepted values for this option are:

- 1= NetBT remains enabled
- 2=Disable NetBT for DHCP clients

**Structure of Disable NetBIOS over TCP/IP (NetBT)**

| Code | Length | NetBT |
|---|---|---|
| 001 | 4 | 1 \| 2 |

**Release DHCP Lease on Shutdown**
This option can be used to control whether DHCP-enabled computers running Windows 2000, Windows XP, or Windows Server 2003 sends a DHCPRelease message to the DHCP server when a shutdown occurs. It is actually implemented and interpreted as a bit-masked value by the DHCP Client service. By default, these clients do not send DHCPRelease messages on proper shutdown.

**Code**
2

**Length**
4 octets

**Value**
Accepted values for this option are:

- 1 = DHCP clients send a DHCPRelease message on proper shutdown
- 0 = DHCP clients do not send a DHCPRelease message on proper shutdown

**Structure of Release DHCP Lease on Shutdown**

| Code | Length | Release |
|---|---|---|
| 002 | 4 | 0 \| 1 |

**Default Router Metric Base**
This option can be used to set the default base metric for DHCP clients running Windows Server 2003. When this option is set, the DHCP Client service uses the value configured here as the base metric for its default gateways.

**Code**
3

**Length**
4 octets

**Value**
This value represents a specified router metric base to be used for all default gateway routes used by DHCP-enabled clients running Windows Server 2003. This value can be assigned as an integer representing a cost metric ranging from 1 through 9,999. It is used in calculating the fastest, most reliable, and least expensive routes. If a value is not specified, a default of either one (1) or the currently set interface-specific metric is used.

**Structure of Default Router Metric Base**

| Code | Length | Route Metric |
|---|---|---|
| 003 | *4* | *Router metric base* |

**Related Information**
The following resources contain additional information that is relevant to this section.

- The IETF RFC Database
- What Is DHCP?
- How DHCP Technology Works

# DNS registration changes for Windows Server 2003 based DHCP Servers

Updated: August 19, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This article describes the DNS registration changes for DHCP Servers that are running Windows Server® 2003.

**Overview**

When the DHCP Server role is installed on a domain controller, the DHCP Server inherits the security permissions of the domain controller. To prevent possible misuse of the domain controller's elevated permissions, DHCP Servers that are installed on Windows Server 2003 do not register DNS records on behalf of the clients that are associated with the DHCP Server unless the DHCP Server is explicitly configured with DNS credentials. This behavior prevents a potential escalation of privilege to clients of the DHCP Server.

**Configuring DNS credentials on a DHCP Server that is running Windows Server 2003**

Follow these steps to configure the DNS credentials of a DHCP Server that is installed on a domain controller running the Windows Server 2003 operating system:

1. Start the DHCP management console. Click **Start**, then click **Run** and type **dhcpmgmt.msc**.

2. In the DHCP navigation pane, expand the server node and right-click **IPv4**, and then click **Properties**.

3. In the resulting dialog box, click the **Advanced** tab, and then click **Credentials**.

4. In the **DNS dynamic update credentials** dialog box, enter the DNS credentials (User name, Domain, and Password) for the account that is used by the DHCP Server for DNS registrations.

| **Note** |
| --- |
| These credentials should belong to a member of the DnsAdmins or the DnsUpdateProxy group. |

**Related information**

For more information about creating and adding a user account as a member of a group by using the Active Directory® Users and Groups console, see Changing group memberships.

# Windows Time Service Technical Reference

Updated: August 17, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 Foundation, Windows Server 2008 R2, Windows Server 2008 R2 Foundation

**Windows Time Service Technical Reference**

**In this guide**


- Where to Find Windows Time Service Configuration Information
- What is the Windows Time Service?
- Importance of Time Protocols
- How the Windows Time Service Works
- Windows Time Service Tools and Settings


| **Note** |
| --- |
| In Windows Server® 2003 and Microsoft Windows 2000 Server, the directory service is named Active Directory® directory service. In Windows Server 2008 R2 and Windows Server 2008, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to AD DS, but the information is also applicable to Active Directory. |

The Windows® Time service, also known as W32Time, synchronizes the date and time for all computers running in an AD DS domain. Time synchronization is critical for the proper operation of many Windows services and line-of-business applications. The Windows Time service uses the Network Time Protocol (NTP) to synchronize computer clocks on the network so that an accurate clock value, or time stamp, can be assigned to network validation and resource access requests. The service integrates NTP and time providers, making it a reliable and scalable time service for enterprise administrators.

| **Important** |
| --- |
| The W32Time service is not a full-featured NTP solution that meets time-sensitive application needs and is not supported by Microsoft as such. For more information, see Microsoft Knowledge Base article 939322, Support boundary to configure the Windows Time service for high-accuracy environments (http://go.microsoft.com/fwlink/?LinkID=179459). |

**Where to Find Windows Time Service Configuration Information**

This guide does **not** discuss configuring the Windows Time service. There are several different topics on Microsoft TechNet and in the Microsoft Knowledge Base that do explain procedures for configuring the Windows Time service. If you require configuration information, the following topics should help you locate the appropriate information.

- To configure the Windows Time service for the forest root primary domain controller (PDC) emulator, see:
    - Configure the Windows Time service on the PDC emulator in the forest root domain (http://go.microsoft.com/fwlink/?LinkId=91969)
    - Configuring a time source for the forest (http://go.microsoft.com/fwlink/?LinkId=111146)
    - Microsoft Knowledge Base article 816042, How to configure an authoritative time server in Windows Server (http://go.microsoft.com/fwlink/?LinkID=60402), which describes configuration settings for computers running Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, and Windows Server 2003 R2.
    - Microsoft Knowledge Base article 216734, How to configure an authoritative time server in Windows 2000 (http://go.microsoft.com/fwlink/?LinkID=23042), which describes configuration settings for computers running

Windows 2000 Server.

- To configure the Windows Time service on any domain member client or server, or even domain controllers that are not configured as the forest root PDC emulator, see Configure a client computer for automatic domain time synchronization (http://go.microsoft.com/fwlink/?LinkId=91376).

| **Warning** |
| --- |
| Some applications may require their computers to have high-accuracy time services. If that is the case, you may choose to configure a manual time source, but be aware that the Windows Time service was not designed to function as a highly accurate time source. Ensure that you are aware of the support limitations for high-accuracy time environments as described in Microsoft Knowledge Base article 939322, Support boundary to configure the Windows Time service for high-accuracy environments (http://go.microsoft.com/fwlink/?LinkID=179459). |

- To configure the Windows Time service on any Windows-based client or server computers that are configured as workgroup members instead of domain members see Configure a manual time source for a selected client computer (http://go.microsoft.com/fwlink/?LinkId=91377).

- To configure the Windows Time service on a host computer that runs a virtual environment, see Microsoft Knowledge Base article 816042, How to configure an authoritative time server in Windows Server (http://go.microsoft.com/fwlink/?LinkID=60402). If you are working with a non-Microsoft virtualization product, be sure to consult the documentation of the vendor for that product.

- To configure the Windows Time service on a domain controller that is running in a virtual machine, it is recommended that you partially disable time synchronization between the host system and guest operating system acting as a domain controller. This enables your guest domain controller to synchronize time for the domain hierarchy, but protects it from having a time skew if it is restored from a Saved state. For more information, see Microsoft Knowledge Base article 976924, You receive Windows Time Service event IDs 24, 29, and 38 on a virtualized domain controller that is running on a Windows Server 2008–based host server with Hyper-V (http://go.microsoft.com/fwlink/?LinkID=192236) and Deployment Considerations for Virtualized Domain Controllers (http://go.microsoft.com/fwlink/?LinkID=192235).

- To configure the Windows Time service on a domain controller acting as the forest root PDC emulator that is also running in a virtual computer, follow the same instructions for a physical computer as described in Configure the Windows Time service on the PDC emulator in the Forest Root Domain (http://go.microsoft.com/fwlink/?LinkId=91969).

- To configure the Windows Time service on a member server running as a virtual computer, use the domain time hierarchy as described in (Configure a client computer for automatic domain time synchronization (http://go.microsoft.com/fwlink/?LinkId=91376).

**What is the Windows Time Service?**

The Windows Time service (W32Time) provides network clock synchronization for computers without the need for extensive configuration.

The Windows Time service is essential to the successful operation of Kerberos version 5 authentication and, therefore, to AD DS-based authentication. Any Kerberos-aware application, including most security services, relies on time synchronization between the computers that are participating in the authentication request. AD DS domain controllers must also have synchronized clocks to help to ensure accurate data replication.

The Windows Time service is implemented in a dynamic link library called W32Time.dll. W32Time.dll is installed by default in the **%Systemroot%\System32** folder during operating system setup and installation.

W32Time.dll was originally developed for Windows 2000 Server to support a specification by the Kerberos V5 authentication protocol that required clocks on a network to be synchronized. Starting with Windows Server 2003, W32Time.dll provided increased accuracy in network clock synchronization over the Windows 2000 Server operating system and, in addition, supported a variety of hardware devices and network time protocols by means of time providers. Although originally designed to provide clock synchronization for Kerberos authentication, many current applications use timestamps to ensure transactional consistency, to record the time of important events, and other business-critical, time-sensitive information. These applications benefit from time synchronization between computers that is provided by the Windows Time service.

**Importance of Time Protocols**

Time protocols communicate between two computers to exchange time information and then use that information to synchronize their clocks. With the Windows Time service time protocol, a client requests time information from a server and synchronizes its clock based on the information that is received.

The Windows Time service uses NTP to help synchronize time across a network. NTP is an Internet time protocol that includes the discipline algorithms necessary for synchronizing clocks. NTP is a more accurate time protocol than the Simple Network Time Protocol (SNTP) that is used in some versions of Windows; however, W32Time continues to support SNTP to enable backward compatibility with computers running SNTP-based time services such as Windows 2000.

**See Also**

**Concepts**

How the Windows Time Service Works

Windows Time Service Tools and Settings

**Other Resources**

Microsoft Knowledge Base article 902229

Videos about the Windows Time Service

# How the Windows Time Service Works

Updated: March 12, 2015

Applies To: Windows SBS 2003, Windows SBS 2008, Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 Foundation, Windows Server 2008 R2, Windows Server 2008 R2 Foundation, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Windows Time Service Architecture
- Windows Time Service Time Protocols
- Windows Time Service Processes and Interactions
- Network Ports Used by Windows Time Service


**Note**

This topic explains only how the Windows Time service (W32Time) works. For information about how to configure Windows Time service, see the list of topics in the section Where to Find Windows Time Service Configuration Information.

**Note**

In Windows Server® 2003 and Microsoft Windows 2000 Server, the directory service is named Active Directory® directory service. In Windows Server 2008 and later versions, the directory service is named Active Directory Domain Services (AD DS). The rest of this topic refers to AD DS, but the information is also applicable to Active Directory.

Although the Windows Time service is not an exact implementation of the Network Time Protocol (NTP), it uses the complex suite of algorithms that is defined in the NTP specifications to ensure that clocks on computers throughout a network are as accurate as possible. Ideally, all computer clocks in an AD DS domain are synchronized with the time of an authoritative computer. Many factors can affect time synchronization on a network. The following factors often affect the accuracy of synchronization in AD DS:

- Network conditions
- The accuracy of the computer's hardware clock
- The amount of CPU and network resources available to the Windows Time service

**Important**

The W32Time service is not a full-featured NTP solution that meets time-sensitive application needs and is not supported by Microsoft as such. For more information, see Microsoft Knowledge Base article 939322, Support boundary to configure the Windows Time service for high-accuracy environments (http://go.microsoft.com/fwlink/?LinkID=179459).

Computers that synchronize their time less frequently, such as computers running Windows XP Home Edition, computers with intermittent network connections, or computers that are not joined to a domain, are configured by default to synchronize with time.windows.com. Because they do not synchronize their clock frequently and because the factors that affect time accuracy may not be known, it is impossible to guarantee time accuracy on computers that have intermittent or no network connections.

An AD DS forest has a predetermined time synchronization hierarchy. The Windows Time service synchronizes time between computers within the hierarchy, with the most accurate reference clocks at the top. If more than one time source is configured on a computer, Windows Time uses NTP algorithms to select the best time source from the configured sources based on the computer's ability to synchronize with that time source. The Windows Time service does not support network synchronization

from broadcast or multicast peers. For more information about these NTP features, see RFC 1305 in the IETF RFC Database.

Every computer that is running the Windows Time service uses the service to maintain the most accurate time. In most cases, it is not necessary to configure the Windows Time service. Computers that are members of a domain act as a time client by default. In addition, the Windows Time service can be configured to request time from a designated reference time source, and can also be configured to provide time to clients.

The degree to which a computer's time is accurate is called a stratum. The most accurate time source on a network (such as a hardware clock) occupies the lowest stratum level, or stratum one. This accurate time source is called a reference clock. An NTP server that acquires its time directly from a reference clock occupies a stratum that is one level higher than that of the reference clock. Resources that acquire time from the NTP server are two steps away from the reference clock, and therefore occupy a stratum that is two higher than the most accurate time source, and so on. As a computer's stratum number increases, the time on its system clock may become less accurate. Therefore, the stratum level of any computer is an indicator of how closely that computer is synchronized with the most accurate time source.

When the W32Time Manager receives time samples, it uses special algorithms in NTP to determine which of the time samples is the most appropriate for use. The time service also uses another set of algorithms to determine which of the configured time sources is the most accurate. When the time service has determined which time sample is best, based on the above criteria, it adjusts the local clock rate to allow it to converge toward the correct time. If the time difference between the local clock and the selected accurate time sample (also called the time skew) is too large to correct by adjusting the local clock rate, the time service sets the local clock to the correct time. This adjustment of clock rate or direct clock time change is known as clock discipline.

**Windows Time Service Architecture**

The Windows Time service consists of the following components:

- Service Control Manager
- Windows Time Service Manager
- Clock Discipline
- Time providers

The following figure shows the architecture of the Windows Time service.

**Windows Time Service Architecture**




The Service Control Manager is responsible for starting and stopping the Windows Time service. The Windows Time Service Manager is responsible for initiating the action of the NTP time providers included with the operating system. The Windows Time Service Manager controls all functions of the Windows Time service and the coalescing of all time samples. In addition to providing information about the current system state, such as the current time source or the last time the system clock was updated, the Windows Time Service Manager is also responsible for creating events in the event log.

The time synchronization process involves the following steps:

- Input providers request and receive time samples from configured NTP time sources.
- These time samples are then passed to the Windows Time Service Manager, which collects all the samples and passes them to the clock discipline subcomponent.

- The clock discipline subcomponent applies the NTP algorithms which results in the selection of the best time sample.
- The clock discipline subcomponent adjusts the time of the system clock to the most accurate time by either adjusting the clock rate or directly changing the time.

If a computer has been designated as a time server, it can send the time on to any computer requesting time synchronization at any point in this process.

**Windows Time Service Time Protocols**
Time protocols determine how closely two computers' clocks are synchronized. A time protocol is responsible for determining the best available time information and converging the clocks to ensure that a consistent time is maintained on separate systems.

The Windows Time service uses the Network Time Protocol (NTP) to help synchronize time across a network. NTP is an Internet time protocol that includes the discipline algorithms necessary for synchronizing clocks. NTP is a more accurate time protocol than the Simple Network Time Protocol (SNTP) that is used in some versions of Windows; however W32Time continues to support SNTP to enable backward compatibility with computers running SNTP-based time services, such as Windows 2000.

**Network Time Protocol**
Network Time Protocol (NTP) is the default time synchronization protocol used by the Windows Time service in the operating system. NTP is a fault-tolerant, highly scalable time protocol and is the protocol used most often for synchronizing computer clocks by using a designated time reference.

NTP time synchronization takes place over a period of time and involves the transfer of NTP packets over a network. NTP packets contain time stamps that include a time sample from both the client and the server participating in time synchronization.

NTP relies on a reference clock to define the most accurate time to be used and synchronizes all clocks on a network to that reference clock. NTP uses Coordinated Universal Time (UTC) as the universal standard for current time. UTC is independent of time zones and enables NTP to be used anywhere in the world regardless of time zone settings.

**NTP Algorithms**
NTP includes two algorithms, a clock-filtering algorithm and a clock-selection algorithm, to assist the Windows Time service in determining the best time sample. The clock-filtering algorithm is designed to sift through time samples that are received from queried time sources and determine the best time samples from each source. The clock-selection algorithm then determines the most accurate time server on the network. This information is then passed to the clock discipline algorithm, which uses the information gathered to correct the local clock of the computer, while compensating for errors due to network latency and computer clock inaccuracy.

The NTP algorithms are most accurate under conditions of light-to-moderate network and server loads. As with any algorithm that takes network transit time into account, NTP algorithms might perform poorly under conditions of extreme network congestion. For more information about the NTP algorithms, see RFC 1305 in the IETF RFC Database.

**NTP Time Provider**
The Windows Time service is a complete time synchronization package that can support a variety of hardware devices and time protocols. To enable this support, the service uses pluggable time providers. A time provider is responsible for either obtaining accurate time stamps (from the network or from hardware) or for providing those time stamps to other computers over the network.

The NTP provider is the standard time provider included with the operating system. The NTP provider follows the standards specified by NTP version 3 for a client and server, and can interact with SNTP clients and servers for backward compatibility with Windows 2000 and other SNTP clients. The NTP provider in the Windows Time service consists of the following two parts:

- **NtpServer output provider.** This is a time server that responds to client time requests on the network.
- **NtpClient input provider.** This is a time client that obtains time information from another source, either a hardware device or an NTP server, and can return time samples that are useful for synchronizing the local clock.

Although the actual operations of these two providers are closely related, they appear independent to the time service. Starting with Windows 2000 Server, when a Windows computer is connected to a network, it is configured as an NTP client. Also, computers running the Windows Time service only attempt to synchronize time with a domain controller or a manually specified time source by default. These are the preferred time providers because they are automatically available, secure

sources of time.

**NTP Security**
Within an AD DS forest, the Windows Time service relies on standard domain security features to enforce the authentication of time data. The security of NTP packets that are sent between a domain member computer and a local domain controller that is acting as a time server is based on shared key authentication. The Windows Time service uses the computer's Kerberos session key to create authenticated signatures on NTP packets that are sent across the network. NTP packets are not transmitted inside the Net Logon secure channel. Instead, when a computer requests the time from a domain controller in the domain hierarchy, the Windows Time service requires that the time be authenticated. The domain controller then returns the required information in the form of a 64-bit value that has been authenticated with the session key from the Net Logon service. If the returned NTP packet is not signed with the computer's session key or is signed incorrectly, the time is rejected. All such authentication failures are logged in the Event Log. In this way, the Windows Time service provides security for NTP data in an AD DS forest.

Generally, Windows time clients automatically obtain accurate time for synchronization from domain controllers in the same domain. In a forest, the domain controllers of a child domain synchronize time with domain controllers in their parent domains. When a time server returns an authenticated NTP packet to a client that requests the time, the packet is signed by means of a Kerberos session key defined by an interdomain trust account. The interdomain trust account is created when a new AD DS domain joins a forest, and the Net Logon service manages the session key. In this way, the domain controller that is configured as reliable in the forest root domain becomes the authenticated time source for all of the domain controllers in both the parent and child domains, and indirectly for all computers located in the domain tree.

The Windows Time service can be configured to work between forests, but it is important to note that this configuration is not secure. For example, an NTP server might be available in a different forest. However, because that computer is in a different forest, there is no Kerberos session key with which to sign and authenticate NTP packets. To obtain accurate time synchronization from a computer in a different forest, the client needs network access to that computer and the time service must be configured to use a specific time source located in the other forest. If a client is manually configured to access time from an NTP server outside of its own domain hierarchy, the NTP packets sent between the client and the time server are not authenticated, and therefore are not secure. Even with the implementation of forest trusts, the Windows Time service is not secure across forests. Although the Net Logon secure channel is the authentication mechanism for the Windows Time service, authentication across forests is not supported.

**Hardware Devices That Are Supported by the Windows Time Service**
Hardware-based clocks such as GPS or radio clocks are often used as highly accurate reference clock devices. By default, the Windows Time service NTP time provider does not support the direct connection of a hardware device to a computer, although it is possible to create a software-based independent time provider that supports this type of connection. This type of provider, in conjunction with the Windows Time service, can provide a reliable, stable time reference.

Hardware devices, such as a cesium clock or a Global Positioning System (GPS) receiver, provide accurate current time by following a standard to obtain an accurate definition of time. Cesium clocks are extremely stable and are unaffected by factors such as temperature, pressure, or humidity, but are also very expensive. A GPS receiver is much less expensive to operate and is also an accurate reference clock. GPS receivers obtain their time from satellites that obtain their time from a cesium clock. Without the use of an independent time provider, Windows time servers can acquire their time by connecting to an external NTP server, which is connected to a hardware device by means of a telephone or the Internet. Organizations such as the United States Naval Observatory provide NTP servers that are connected to extremely reliable reference clocks.

Many GPS receivers and other time devices can function as NTP servers on a network. You can configure your AD DS forest to synchronize time from these external hardware devices only if they are also acting as NTP servers on your network. To do so, configure the domain controller functioning as the primary domain controller (PDC) emulator in your forest root to synchronize with the NTP server provided by the GPS device. To do so, see Configure the Windows Time service on the PDC emulator in the Forest Root Domain (http://go.microsoft.com/fwlink/?LinkId=91969).

**Simple Network Time Protocol**
The Simple Network Time Protocol (SNTP) is a simplified time protocol that is intended for servers and clients that do not require the degree of accuracy that NTP provides. SNTP, a more rudimentary version of NTP, is the primary time protocol that is used in Windows 2000. Because the network packet formats of SNTP and NTP are identical, the two protocols are interoperable. The primary difference between the two is that SNTP does not have the error management and complex filtering systems that NTP provides. For more information about the Simple Network Time Protocol, see RFC 1769 in the IETF RFC Database.

**Time Protocol Interoperability**
The Windows Time service can operate in a mixed environment of computers running Windows 2000, Windows XP, and Windows Server 2003, because the SNTP protocol used in Windows 2000 is interoperable with the NTP protocol in Windows XP and Windows Server 2003.

The time service in Windows NT Server 4.0, called TimeServ, synchronizes time across a Windows NT 4.0 network. TimeServ is an add-on feature available as part of the *Microsoft Windows NT 4.0 Resource Kit* and does not provide the degree of reliability of time synchronization that is required by Windows Server 2003.

The Windows Time service can interoperate with computers running Windows NT 4.0 because they can synchronize time with computers running Windows 2000 or Windows Server 2003; however, a computer running Windows 2000 or Windows Server 2003 does not automatically discover Windows NT 4.0 time servers. For example, if your domain is configured to synchronize time by using the domain hierarchy–based method of synchronization and you want computers in the domain hierarchy to synchronize time with a Windows NT 4.0 domain controller, you have to configure those computers manually to synchronize with the Windows NT 4.0 domain controllers.

Windows NT 4.0 uses a simpler mechanism for time synchronization than the Windows Time service uses. Therefore, to ensure accurate time synchronization across your network, it is recommended that you upgrade any Windows NT 4.0 domain controllers to Windows 2000 or Windows Server 2003.

**Windows Time Service Processes and Interactions**

The Windows Time service is designed to synchronize the clocks of computers on a network. The network time synchronization process, also called time convergence, occurs throughout a network as each computer accesses time from a more accurate time server. Time convergence involves a process by which an authoritative server provides the current time to client computers in the form of NTP packets. The information provided within a packet indicates whether an adjustment needs to be made to the computer's current clock time so that it is synchronized with the more accurate server.

As part of the time convergence process, domain members attempt to synchronize time with any domain controller located in the same domain. If the computer is a domain controller, it attempts to synchronize with a more authoritative domain controller.

Computers running Windows XP Home Edition or computers that are not joined to a domain do not attempt to synchronize with the domain hierarchy, but are configured by default to obtain time from time.windows.com.

To establish a computer running Windows Server 2003 as authoritative, the computer must be configured to be a reliable time source. By default, the first domain controller that is installed on a Windows Server 2003 domain is automatically configured to be a reliable time source. Because it is the authoritative computer for the domain, it must be configured to synchronize with an external time source rather than with the domain hierarchy. Also by default, all other Windows Server 2003 domain members are configured to synchronize with the domain hierarchy.

After you have established a Windows Server 2003 network, you can configure the Windows Time service to use one of the following options for synchronization:

- Domain hierarchy-based synchronization
- A manually-specified synchronization source
- All available synchronization mechanisms
- No synchronization.

Each of these synchronization types is discussed in the following section.

**Domain Hierarchy–Based Synchronization**

Synchronization that is based on a domain hierarchy uses the AD DS domain hierarchy to find a reliable source with which to synchronize time. Based on domain hierarchy, the Windows Time service determines the accuracy of each time server. In a Windows Server 2003 forest, the computer that holds the primary domain controller (PDC) emulator operations master role, located in the forest root domain, holds the position of best time source, unless another reliable time source has been configured. The following figure illustrates a path of time synchronization between computers in a domain hierarchy.

**Time Synchronization in an AD DS Hierarchy**




**Reliable Time Source Configuration**

A computer that is configured to be a reliable time source is identified as the root of the time service. The root of the time service is the authoritative server for the domain and typically is configured to retrieve time from an external NTP server or hardware device. A time server can be configured as a reliable time source to optimize how time is transferred throughout the domain hierarchy. If a domain controller is configured to be a reliable time source, Net Logon service announces that domain controller as a reliable time source when it logs on to the network. When other domain controllers look for a time source to synchronize with, they choose a reliable source first if one is available.

**Time Source Selection**

The time source selection process can create two problems on a network:

- Additional synchronization cycles.
- Increased volume in network traffic.

A cycle in the synchronization network occurs when time remains consistent between a group of domain controllers and the same time is shared between them continuously without a resynchronization with another reliable time source. The Windows Time service's time source selection algorithm is designed to protect against these types of problems.

A computer uses one of the following methods to identify a time source to synchronize with:

- If the computer is not a member of a domain, it must be configured to synchronize with a specified time source.
- If the computer is a member server or workstation within a domain, by default, it follows the AD DS hierarchy and synchronizes its time with a domain controller in its local domain that is currently running the Windows Time service.

If the computer is a domain controller, it makes up to six queries to locate another domain controller to synchronize with. Each query is designed to identify a time source with certain attributes, such as a type of domain controller, a particular location, and whether or not it is a reliable time source. The time source must also adhere to the following constraints:

- A reliable time source can only synchronize with a domain controller in the parent domain.
- A PDC emulator can synchronize with a reliable time source in its own domain or any domain controller in the parent

domain.

If the domain controller is not able to synchronize with the type of domain controller that it is querying, the query is not made. The domain controller knows which type of computer it can obtain time from before it makes the query. For example, a local PDC emulator does not attempt to query numbers three or six because a domain controller does not attempt to synchronize with itself.

The following table lists the queries that a domain controller makes to find a time source and the order in which the queries are made.

**Domain Controller Time Source Queries**

| Query Number | Domain Controller | Location | Reliability of Time Source |
|---|---|---|---|
| 1 | Parent domain controller | In-site | Prefers a reliable time source but it can synchronize with a non-reliable time source if that is all that is available. |
| 2 | Local domain controller | In-site | Only synchronizes with a reliable time source. |
| 3 | Local PDC emulator | In-site | Does not apply.<br><br>A domain controller does not attempt to synchronize with itself. |
| 4 | Parent domain controller | Out-of-site | Prefers a reliable time source but it can synchronize with a non-reliable time source if that is all that is available. |
| 5 | Local domain controller | Out-of-site | Only synchronizes with a reliable time source. |
| 6 | Local PDC emulator | Out-of-site | Does not apply.<br><br>A domain controller does not attempt to synchronize with itself. |

**Note**

- A computer never synchronizes with itself. If the computer attempting synchronization is the local PDC emulator, it does not attempt Queries 3 or 6.

Each query returns a list of domain controllers that can be used as a time source. Windows Time assigns each domain controller that is queried a score based on the reliability and location of the domain controller. The following table lists the scores assigned by Windows Time to each type of domain controller.

**Score Determination**

| Domain Controller Status | Score |
|---|---|
| Domain controller located in same site | 8 |
| Domain controller marked as a reliable time source | 4 |
| Domain controller located in the parent domain | 2 |
| Domain controller that is a PDC emulator | 1 |

When the Windows Time service determines that it has identified the domain controller with the best possible score, no more queries are made. The scores assigned by the time service are cumulative, which means that a PDC emulator located in the same site receives a score of nine.

If the root of the time service is not configured to synchronize with an external source, the internal hardware clock of the computer governs the time.

**Manually-Specified Synchronization**

Manually-specified synchronization enables you to designate a single peer or list of peers from which a computer obtains time. If the computer is not a member of a domain, it must be manually configured to synchronize with a specified time source. A computer that is a member of a domain is configured by default to synchronize from the domain hierarchy, manually-specified synchronization is most useful for the forest root of the domain or for computers that are not joined to a domain. Manually specifying an external NTP server to synchronize with the authoritative computer for your domain provides reliable time. However, configuring the authoritative computer for your domain to synchronize with a hardware clock is actually a better solution for providing the most accurate, secure time to your domain.

Manually-specified time sources are not authenticated unless a specific time provider is written for them, and they are therefore vulnerable to attackers. Also, if a computer synchronizes with a manually-specified source rather than its authenticating domain controller, the two computers might be out of synchronization, causing Kerberos authentication to fail. This might cause other actions requiring network authentication to fail, such as printing or file sharing. If only the forest root is configured to synchronize with an external source, all other computers within the forest remain synchronized with each other, making replay attacks difficult.

**All Available Synchronization Mechanisms**
The "all available synchronization mechanisms" option is the most valuable synchronization method for users on a network. This method allows synchronization with the domain hierarchy and may also provide an alternate time source if the domain hierarchy becomes unavailable, depending on the configuration. If the client is unable to synchronize time with the domain hierarchy, the time source automatically falls back to the time source specified by the **NtpServer** setting. This method of synchronization is most likely to provide accurate time to clients.

**Stopping Time Synchronization**
There are certain situations in which you will want to stop a computer from synchronizing its time. For example, if a computer attempts to synchronize from a time source on the Internet or from another site over a WAN by means of a dial-up connection, it can incur costly telephone charges. When you disable synchronization on that computer, you prevent the computer from attempting to access a time source over a dial-up connection.

You can also disable synchronization to prevent the generation of errors in the event log. Each time a computer attempts to synchronize with a time source that is unavailable, it generates an error in the Event Log. If a time source is taken off of the network for scheduled maintenance and you do not intend to reconfigure the client to synchronize from another source, you can disable synchronization on the client to prevent it from attempting synchronization while the time server is unavailable.

It is useful to disable synchronization on the computer that is designated as the root of the synchronization network. This indicates that the root computer trusts its local clock. If the root of the synchronization hierarchy is not set to **NoSync** and if it is unable to synchronize with another time source, clients do not accept the packet that this computer sends out because its time cannot be trusted.

The only time servers that are trusted by clients even if they have not synchronized with another time source are those that have been identified by the client as reliable time servers.

**Disabling the Windows Time Service**
The Windows Time service (W32Time) can be completely disabled. If you choose to implement a third-party time synchronization product that uses NTP, you must disable the Windows Time service. This is because all NTP servers need access to User Datagram Protocol (UDP) port 123, and as long as the Windows Time service is running on the Windows Server 2003 operating system, port 123 remains reserved by Windows Time.

**Network Ports Used by Windows Time Service**
The Windows Time service communicates on a network to identify reliable time sources, obtain time information, and provide time information to other computers. It performs this communication as defined by the NTP and SNTP RFCs.

### Port Assignments for the Windows Time Service

| Service name | UDP | TCP |
|---|---|---|
| NTP | 123 | N/A |
| SNTP | 123 | N/A |

**See Also**
**Reference**
Windows Time Service Technical Reference
**Concepts**
Windows Time Service Tools and Settings
**Other Resources**

# Windows Time Service Tools and Settings

Updated: May 23, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 Foundation, Windows Server 2008 R2, Windows Server 2008 R2 Foundation

**In this section**


- Windows Time Service Tools
- Windows Time Service Registry Entries
- Windows Time Service Group Policy Settings
- Network Ports Used by the Windows Time Service
- Related Information


**Note**

This topic contains information only about tools and settings for Windows Time service (W32Time). If you only want to synchronize time for a domain-joined client computer, see Configure a client computer for automatic domain time synchronization. For additional topics about how to configure Windows Time service, see the list of topics in the section Where to Find Windows Time Service Configuration Information.

**Caution**

You should not use the **Net time** command to configure or set time when the Windows Time service is running.

Also, on older computers that run Windows XP or earlier, the command **Net time /querysntp** displays the name of a Network Time Protocol (NTP) server with which a computer is configured to synchronize, but that NTP server is used only when the computer's time client is configured as NTP or AllSync. That command has since been deprecated.

Most domain member computers have a time client type of NT5DS, which means that they synchronize time from the domain hierarchy. The only typical exception to this is the domain controller that functions as the primary domain controller (PDC) emulator operations master of the forest root domain, which is usually configured to synchronize time with an external time source. To view the time client configuration of a computer, run the **W32tm /query /configuration** command from an elevated Command Prompt in starting in Windows Server 2008, and Windows Vista®, and read the **Type** line in the command output. For more information, see How Windows Time Service Works (http://go.microsoft.com/fwlink/?LinkId=117753). You can run the command **reg query HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Parameters** and read the value of **NtpServer** in the command output.

**Important**

The W32Time service is not a full-featured NTP solution that meets time-sensitive application needs and is not supported by Microsoft as such. For more information, see Microsoft Knowledge Base article 939322, Support boundary to configure the Windows Time service for high-accuracy environments (http://go.microsoft.com/fwlink/?LinkID=179459).

**Windows Time Service Tools**

The following tools are associated with the Windows Time service.

**W32tm.exe: Windows Time**

**Category**

This tool is installed as part of Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server® 2008 R2 default installations.

**Version compatibility**

This tool works on Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2 default installations.

W32tm.exe is used to configure Windows Time service settings. It can also be used to diagnose problems with the time service. W32tm.exe is the preferred command line tool for configuring, monitoring, or troubleshooting the Windows Time service.

The following tables describe the parameters that are used with W32tm.exe.

**W32tm.exe Primary Parameters**

| **Parameter** | **Description** |
|---|---|
| W32tm /? | W32tm command line help |
| W32tm /register | Registers the time service to run as a service and adds default configuration to the registry. |
| W32tm /unregister | Unregisters the time service and removes all configuration information from the registry. |
| w32tm /monitor<br><br>[/domain:<domain name>]<br>[/computers:<name>[,<name>[,<name>...]]]<br>[/threads:<num>] | domain – specifies which domain to monitor. If no domain name is given, or neither the domain nor computers option is specified, the default domain is used. This option might be used more than once.<br><br>computers – monitors the given list of computers. Computer names are separated by commas, with no spaces. If a name is prefixed with a '*', it is treated as a PDC. This option might be used more than once.<br><br>threads – specifies the number of computers to analyze simultaneously. The default value is 3. Allowed range is 1-50. |
| w32tm /ntte <NT time epoch> | Convert an NT system time, in (10^-7)s intervals from 0h 1-Jan 1601, into a readable format. |
| w32tm /ntpte <NTP time epoch> | Convert an NTP time, in (2^-32)s intervals from 0h 1-Jan 1900, into a readable format. |
| w32tm /resync<br><br>[/computer:<computer>]<br><br>[/nowait]<br><br>[/rediscover]<br><br>[/soft] | Tells a computer that it should resynchronize its clock as soon as possible, throwing out all accumulated error statistics.<br><br>computer:<computer> – Specifies the computer that should resynchronize. If not specified, the local computer will resynchronize.<br><br>nowait – do not wait for the resynchronize to occur; return immediately. Otherwise, wait for the resynchronize to complete before returning.<br><br>rediscover – Redetect the network configuration and rediscover network sources, then resynchronize.<br><br>soft – resynchronize using existing error statistics. Not useful, provided for compatibility. |
| w32tm /stripchart<br><br>/computer:<target><br><br>[/period:<refresh>]<br><br>[/dataonly]<br><br>[/samples:<count>] | Display a strip chart of the offset between this computer and another computer.<br><br>computer:<target> – the computer to measure the offset against.<br><br>period:<refresh> – the time between samples, in seconds. The default is 2s.<br><br>dataonly – display only the data without graphics.<br><br>samples:<count> – collect <count> samples, then stop. If not specified, samples will be collected until **Ctrl+C** is pressed. |

| | |
|---|---|
| w32tm /config<br><br>[/computer:<target>]<br><br>[/update]<br><br>[/manualpeerlist:<peers>]<br><br>[/syncfromflags:<source>]<br><br>[/LocalClockDispersion: <seconds>]<br><br>[/reliable:(YES\|NO)]<br><br>[/largephaseoffset: <milliseconds>] | computer:<target> – adjusts the configuration of <target>. If not specified, the default is the local computer.<br><br>update – notifies the time service that the configuration has changed, causing the changes to take effect.<br><br>manualpeerlist:<peers> – sets the manual peer list to <peers>, which is a space-delimited list of DNS and/or IP addresses. When specifying multiple peers, this option must be enclosed in quotes.<br><br>syncfromflags:<source> – sets what sources the NTP client should synchronize from. <source> should be a comma separated list of these keywords (not case sensitive):<br><br>MANUAL – include peers from the manual peer list.<br><br>DOMHIER – synchronize from a domain controller (DC) in the domain hierarchy.<br><br>LocalClockDispersion:<seconds> – configures the accuracy of the internal clock that W32Time will assume when it can't acquire time from its configured sources.<br><br>reliable:(YES\|NO) – set whether this computer is a reliable time source.<br><br>This setting is only meaningful on domain controllers.<br><br>YES – this computer is a reliable time service.<br><br>NO – this computer is not a reliable time service.<br><br>largephaseoffset:<milliseconds> – sets the time difference between local and network time which W32Time will consider a spike. |
| w32tm /tz | Display the current time zone settings. |
| w32tm /dumpreg<br><br>[/subkey:<key>]<br><br>[/computer:<target>] | Display the values associated with a given registry key.<br><br>The default key is HKLM\System\CurrentControlSet\Services\W32Time<br><br>(the root key for the time service).<br><br>subkey:<key> – displays the values associated with subkey <key> of the default key.<br><br>computer:<target> – queries registry settings for computer <target> |
| w32tm /query [/computer: <target>] {/source \| /configuration \| /peers \| /status} [/verbose] | This parameter was first made available in the Windows Time client versions of Windows Vista, and Windows Server 2008.<br><br>Display a computer's Windows Time service information.<br><br>**computer:<target>** – Query the information of **<target>**. If not specified, the default value is the local computer.<br><br>**Source** – Display the time source.<br><br>**Configuration** – Display the configuration of run time and where the setting comes from. In verbose mode, display the undefined or unused setting too.<br><br>**peers** – Display a list of peers and their status.<br><br>**status** – Display Windows Time service status.<br><br>**verbose** – Set the verbose mode to display more information. |

| | |
|---|---|
| w32tm /debug {/disable \| {/enable /file:<name> /size: <bytes> /entries:<value> [/truncate]}} | This parameter was first made available in the Windows Time client versions of Windows Vista, and Windows Server 2008.<br><br>Enable or disable the local computer Windows Time service private log.<br><br>**disable** – Disable the private log.<br><br>**enable** – Enable the private log.<br><br>• **file:<name>** – Specify the absolute file name.<br><br>• **size:<bytes>** – Specify the maximum size for circular logging.<br><br>• **entries:<value>** – Contains a list of flags, specified by number and separated by commas, that specify the types of information that should be logged. Valid numbers are 0 to 300. A range of numbers is valid, in addition to single numbers, such as 0-100,103,106. Value 0-300 is for logging all information.<br><br>**truncate** – Truncate the file if it exists. |

For more information about **W32tm.exe**, see Help and Support Center in Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2.

**Windows Time Service Registry Entries**

The following registry entries are associated with the Windows Time service.

This information is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system.

When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**Warning**

Some of the preset values that are configured in the System Administrative template file (System.adm) for the Group Policy object (GPO) settings are different from the corresponding default registry entries. If you plan to use a GPO to configure any Windows Time setting, be sure that you review Preset values for the Windows Time service Group Policy settings are different from the corresponding Windows Time service registry entries in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=186066). This issue applies to Windows Server 2008 R2, Windows Server 2008, Windows Server 2003 R2, and Windows Server 2003.

Many registry entries for the Windows Time service are the same as the Group Policy setting of the same name. The Group Policy settings correspond to the registry entries of the same name located in:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time**.

There are several registry keys at this registry location. The Windows Time settings are stored in values across all of these keys.

**Note**

Many of the values in the W32Time section of the registry are used internally by W32Time to store information. These values should not be manually changed at any time. Do not modify any of the settings in this section unless you are familiar with the setting and are certain that the new value will work as expected. The following registry entries are located under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time**

Some of the parameters are stored in clock ticks in the registry and some are in seconds. To convert the time from clock ticks to seconds:

- 1 minute = 60 sec
- 1 sec = 1000 ms

- 1 ms = 10,000 clock ticks on a Windows system, as described at DateTime.Ticks Property.

For example, 5 minutes would become 5*60*1000*10000 = 3000000000 clock ticks.

**AllowNonstandardModeCombinations**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry indicates that non-standard mode combinations are allowed in synchronization between peers. The default value for domain members is 1. The default value for stand-alone clients and servers is 1.

**AllowNonstandardModeCombinations**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry indicates that non-standard mode combinations are allowed in synchronization between clients and servers. The default value for domain members is 1. The default value for stand-alone clients and servers is 1.

**AnnounceFlags**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls whether this computer is marked as a reliable time server. A computer is not marked as reliable unless it is also marked as a time server.

- 0x00 Not a time server
- 0x01 Always time server
- 0x02 Automatic time server
- 0x04 Always reliable time server
- 0x08 Automatic reliable time server

The default value for domain members is 10. The default value for stand-alone clients and servers is 10.

**CompatibilityFlags**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the following compatibility flags and values:

- DispersionInvalid: 0x00000001
- IgnoreFutureRefTimeStamp: 0x00000002
- AutodetectWin2K: 0x80000000
- AutodetectWin2KStage2: 0x40000000

The default value for domain members is 0x80000000. The default value for stand-alone clients and servers is 0x80000000.

**CrossSiteSyncFlags**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry determines whether the service chooses synchronization partners outside the domain of the computer. The options and values are:

- None: 0
- PdcOnly: 1
- All: 2

This value is ignored if the NT5DS value is not set. The default value for domain members is 2. The default value for stand-alone clients and servers is 2.

**DllName**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the location of the DLL for the time provider.

The default location for this DLL on both domain members and stand-alone clients and servers is %windir%\System32\W32Time.dll.

**DllName**
Registry path

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer**

Version

Windows Server 2003 Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the location of the DLL for the time provider.

The default location for this DLL on both domain members and stand-alone clients and servers is %windir%\System32\W32Time.dll.

**Enabled**

**Registry path**

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry indicates if the NtpClient provider is enabled in the current Time Service.

- Yes 1
- No 0

The default value on domain members is 1. The default value on stand-alone clients and servers is 1.

**Enabled**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry indicates if the NtpServer provider is enabled in the current Time Service.

- Yes 1
- No 0

The default value on domain members is 1. The default value on stand-alone clients and servers is 1.

**EventLogFlags**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls the events that the time service logs.

- Time Jump: 0x1
- Source Change: 0x2

The default value on domain members is 2. The default value on stand-alone clients and servers is 2.

**EventLogFlags**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the events logged by the Windows Time service.

- 0x1 reachability changes
- 0x2 large sample skew (This is applicable to Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2 only)

The default value on domain members is 0x1. The default value on stand-alone clients and servers is 0x1.

**FrequencyCorrectRate**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls the rate at which the clock is corrected. If this value is too small, the clock is unstable and overcorrects. If the value is too large, the clock takes a long time to synchronize. The default value on domain members is 4. The default value on stand-alone clients and servers is 4.

| **Note** |
| --- |
| 0 is an invalid value for the FrequencyCorrectRate registry entry. On Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2computers, if the value is set to 0 the Windows Time service will automatically change it to 1. |

**HoldPeriod**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls the period of time for which spike detection is disabled in order to bring the local clock into synchronization quickly. A spike is a time sample indicating that time is off a number of seconds, and is usually received after good time samples have been returned consistently. The default value on domain members is 5. The default value on stand-alone clients and servers is 5.

**InputProvider**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry indicates if the NtpClient provider is enabled.

- Yes 1
- No 0

The default value on domain members is 1. The default value on stand-alone clients and servers is 1.

**InputProvider**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry indicates if the NtpServer provider is enabled.

- Yes 1

- No 0

The default value on domain members is 1. The default value on stand-alone clients and servers is 1.

**LargePhaseOffset**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies that a time offset greater than or equal to this value in $10^{-7}$ seconds is considered a spike. A network disruption such as a large amount of traffic might cause a spike. A spike will be ignored unless it persists for a long period of time. The default value on domain members is 50000000. The default value on stand-alone clients and servers is 50000000.

**LargeSampleSkew**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows Server 2003 and Windows Server 2008

This entry specifies the large sample skew for logging in seconds. To comply with Security and Exchange Commission (SEC) specifications, this should be set to three seconds. Events will be logged for this setting only when EventLogFlags is explicitly configured for 0x2 large sample skew. The default value on domain members is 3. The default value on stand-alone clients and servers is 3.

**LastClockRate**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry is maintained by W32Time. It contains reserved data that is used by the Windows operating system, and any changes to this setting can cause unpredictable results. The default value on domain members is 156250. The default value on stand-alone clients and servers is 156250.

**LocalClockDispersion**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls the dispersion (in seconds) that you must assume when the only time source is the built-in CMOS clock. The default value on domain members is 10. The default value on stand-alone clients and servers is 10.

**MaxAllowedPhaseOffset**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the maximum offset (in seconds) for which W32Time attempts to adjust the computer clock by using the

clock rate. When the offset exceeds this rate, W32Time sets the computer clock directly. The default value for domain members is 300. The default value for stand-alone clients and servers is 1.

In order for W32Time to set the computer clock gradually, the offset must be less than the MaxAllowedPhaseOffset value and satisfy the following equation at the same time:

|CurrentTimeOffset| / (PhaseCorrectRate*UpdateInterval) < SystemClockRate / 2

The CurrentTimeOffset is measured in clock ticks, where 1ms = 10,000 clock ticks on a Windows system.

SystemClockRate and PhaseCorrectRate are also measured in clock ticks. To get the SystemClockRate, you can use the following command and convert it from seconds to clock ticks using the formula of seconds*1000*10000:

```
W32tm /query /status /verbose
ClockRate: 0.0156000s
```

SystemclockRate is the rate of the clock on the system. Using 156000 seconds as an example, the SystemclockRate would be = 0.0156000 * 1000 * 10000 = 156000 clock ticks.

MaxAllowedPhaseOffset is also in seconds. To convert it to clock ticks, multiply MaxAllowedPhaseOffset*1000*10000.

The following two examples show how to apply

Example 1: Time differs by 4 minutes (For example, your time is 11:05 AM and the time sample received from a peer and believed to be correct is 11:09 AM).

phasecorrectRate = 1

UpdateInterval = 30000 (clock ticks)

systemclockRate = 156000 (clock ticks)

MaxAllowedPhaseOffset = 10min = 600 seconds = 600*1000*10000=6000000000 clock ticks

|currentTimeOffset| = 4mins = 4*60*1000*10000 = 2400000000 ticks

Is CurrentTimeOffset < MaxAllowedPhaseOffset?

2400000000 < 6000000000 = TRUE

AND does it satisfy the above equation? (|CurrentTimeOffset| / (PhaseCorrectRate*UpdateInterval) < SystemClockRate / 2)

Is 2,400,000,000 / (30000*1) < 156000/2

Is 100,000 < 78,000

NO/FALSE

Therefore W32tm would set the clock back immediately.

| **Note** |
| --- |
| In this case, if you want to set the clock back slowly, you would need to adjust the values of PhaseCorrectRate or updateInterval in the registry as well to ensure the equation results in TRUE. |

Example 2: Time differs by 3 minutes.

phasecorrectRate = 1

UpdateInterval = 30000 (clock ticks)

systemclockRate = 156000 (clock ticks)

MaxAllowedPhaseOffset = 10min = 600 seconds = 600*1000*10000=6000000000 clock ticks

currentTimeOffset = 3mins = 3*60*1000*10000 = 1800000000 clock ticks

Is CurrentTimeOffset < MaxAllowedPhaseOffset?

1800000000 < 6000000000 = TRUE

AND does it satisfy the above equation? (|CurrentTimeOffset| / (PhaseCorrectRate*UpdateInterval) < SystemClockRate / 2)

Is 3 mins (1,800,000,000) / (30000*1) < 156000/2

Is 60,000 < 78,000

YES/TRUE

In this case the clock will be set back slowly.

**MaxClockRate**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry is maintained by W32Time. It contains reserved data that is used by the Windows operating system, and any changes to this setting can cause unpredictable results. The default value for domain members is 155860. The default value for stand-alone clients and servers is 155860.

**MaxNegPhaseCorrection**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the largest negative time correction in seconds that the service makes. If the service determines that a change larger than this is required, it logs an event instead. Special case: 0xFFFFFFFF means always make time correction. The default value for domain members is 0xFFFFFFFF. The default value for stand-alone clients and servers is 54,000 (15 hrs).

**MaxPollInterval**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the largest interval, in $\log_2$ seconds, allowed for the system polling interval. Note that while a system must poll according to the scheduled interval, a provider can refuse to produce samples when requested to do so. The default value for domain controllers is 10. The default value for domain members is 15. The default value for stand-alone clients and servers is 15.

**MaxPosPhaseCorrection**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2.

This entry specifies the largest positive time correction in seconds that the service makes. If the service determines that a change larger than this is required, it logs an event instead. Special case: 0xFFFFFFFF means always make time correction. The default value for domain members is 0xFFFFFFFF. The default value for stand-alone clients and servers is 54,000 (15 hrs).

**MinClockRate**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry is maintained by W32Time. It contains reserved data that is used by the Windows operating system, and any changes to this setting can cause unpredictable results. The default value for domain members is 155860. The default value for stand-alone clients and servers is 155860.

**MinPollInterval**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the smallest interval, in $\log_2$ seconds, allowed for the system polling interval. Note that while a system does not request samples more frequently than this, a provider can produce samples at times other than the scheduled interval. The default value for domain controllers is 6. The default value for domain members is 10. The default value for stand-alone clients and servers is 10.

**NtpServer**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Parameters**

**Version**
Windows Server 2003 and Windows Server 2008

This entry specifies a space-delimited list of peers from which a computer obtains time stamps, consisting of one or more DNS names or IP addresses per line. Each DNS name or IP address listed must be unique. Computers connected to a domain must synchronize with a more reliable time source, such as the official U.S. time clock.

- 0x01 SpecialInterval
- 0x02 UseAsFallbackOnly
- 0x04 SymmetricActive

  For more information about this mode, see Windows Time Server: 3.3 Modes of Operation (http://go.microsoft.com/fwlink/?LinkId=208012).

- 0x08 Client

There is no default value for this registry entry on domain members. The default value on stand-alone clients and servers is time.windows.com,0x1.

| **Note** |
| --- |
| For more information on available NTP Servers, see Microsoft Knowledge Base article 262680 (http://go.microsoft.com/fwlink/?LinkId=186067) |

**PhaseCorrectRate**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls the rate at which the phase error is corrected. Specifying a small value corrects the phase error quickly, but might cause the clock to become unstable. If the value is too large, it takes a longer time to correct the phase error.

The default value on domain members is 1. The default value on stand-alone clients and servers is 7.

| Note |
| --- |
| 0 is an invalid value for the PhaseCorrectRate registry entry. On Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2computers, if the value is set to 0, the Windows Time service automatically changes it to 1. |

**PollAdjustFactor**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls the decision to increase or decrease the poll interval for the system. The larger the value, the smaller the amount of error that causes the poll interval to be decreased. The default value on domain members is 5. The default value on stand-alone clients and servers is 5.

**ResolvePeerBackOffMaxTimes**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the maximum number of times to double the wait interval when repeated attempts to locate a peer to synchronize with fail. A value of zero means that the wait interval is always the minimum. The default value on domain members is 7. The default value on stand-alone clients and servers is 7.

**ResolvePeerBackOffMinutes**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the initial interval to wait, in minutes, before attempting to locate a peer to synchronize with. The default value on domain members is 15. The default value on stand-alone clients and servers is 15.

**ServiceDll**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Parameters**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry is maintained by W32Time. It contains reserved data that is used by the Windows operating system, and any changes to this setting can cause unpredictable results. The default location for this DLL on both domain members and stand-alone clients and servers is %windir%\System32\W32Time.dll.

**ServiceMain**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Parameters**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry is maintained by W32Time. It contains reserved data that is used by the Windows operating system, and any changes to this setting can cause unpredictable results. The default value on domain members is SvchostEntry_W32Time. The default value on stand-alone clients and servers is SvchostEntry_W32Time.

**SpecialPollInterval**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the special poll interval in seconds for manual peers. When the SpecialInterval 0x1 flag is enabled, W32Time uses this poll interval instead of a poll interval determine by the operating system. The default value on domain members is 3,600. The default value on stand-alone clients and servers is 604,800.

**SpecialPollTimeRemaining**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient**

**Version**
Windows Server 2003 and Windows Server 2008

This entry is maintained by W32Time. It contains reserved data that is used by the Windows operating system. It specifies the time in seconds before W32Time will resynchronize after the computer has restarted. Any changes to this setting can cause unpredictable results. The default value on both domain members and on stand-alone clients and servers is left blank.

**SpikeWatchPeriod**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry specifies the amount of time that a suspicious offset must persist before it is accepted as correct (in seconds). The default value on domain members is 900. The default value on stand-alone clients and workstations is 900.

**Type**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Parameters**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry Indicates which peers to accept synchronization from:

- **NoSync.** The time service does not synchronize with other sources.
- **NTP.** The time service synchronizes from the servers specified in the **NtpServer.** registry entry.
- **NT5DS.** The time service synchronizes from the domain hierarchy.
- **AllSync.** The time service uses all the available synchronization mechanisms.

The default value on domain members is **NT5DS.** The default value on stand-alone clients and servers is **NTP.**

**UpdateInterval**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows XP, Windows Vista, Windows 7, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2.

This entry specifies the number of clock ticks between phase correction adjustments. The default value for domain controllers is 100. The default value for domain members is 30,000. The default value for stand-alone clients and servers is 360,000.

| **Note** |
| --- |
| 0 is an invalid value for the UpdateInterval registry entry. On computers running Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2, if the value is set to 0 the Windows Time service automatically changes it to 1. |

The following three registry entries are not a part of the W32Time default configuration but can be added to the registry to obtain increased logging capabilities. The information logged to the System Event log can be modified by changing value for the EventLogFlags setting in the Group Policy Object Editor. By default, the time service creates a log in Event Viewer every time that it switches to a new time source.

| **Warning** |
| --- |
| Some of the preset values that are configured in the System Administrative template file (System.adm) for the Group Policy object (GPO) settings are different from the corresponding default registry entries. If you plan to use a GPO to configure any Windows Time setting, be sure that you review Preset values for the Windows Time service Group Policy settings are different from the corresponding Windows Time service registry entries in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=186066). This issue applies to Windows Server 2008 R2, Windows Server 2008, Windows Server 2003 R2, and Windows Server 2003. |

The following registry entries must be added in order to enable W32Time logging:

**FileLogEntries**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls the amount of entries created in the Windows Time log file. The default value is none, which does not log any Windows Time activity. Valid values are 0 to 300. This value does not affect the event log entries normally created by Windows Time.

**FileLogName**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls the location and file name of the Windows Time log. The default value is blank, and should not be changed unless **FileLogEntries** is changed. A valid value is a full path and file name that Windows Time will use to create the log file. This value does not affect the event log entries normally created by Windows Time.

**FileLogSize**

**Registry path**
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config**

**Version**
Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2

This entry controls the circular logging behavior of Windows Time log files. When **FileLogEntries** and **FileLogName** are defined, this entry defines the size, in bytes, to allow the log file to reach before overwriting the oldest log entries with new entries. Any positive number is valid, and 3000000 is recommended. This value does not affect the event log entries normally created by Windows Time.

**Windows Time Service Group Policy Settings**

You can configure most W32Time parameters by using the Group Policy Object Editor. This includes configuring a computer to be an NTPServer or NTPClient, configuring the time synchronization mechanism, and configuring a computer to be a reliable time source.

| **Note** |
| --- |
| Group Policy settings for the Windows Time service can be configured on Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2 domain controllers and can be applied only to computers running Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2. |

You can find the Group Policy settings used to configure W32Time in the Group Policy Object Editor snap-in in the following locations:

- Computer Configuration\Administrative Templates\System\Windows Time Service

  Configure **Global Configuration Settings** here.

- Computer Configuration\Administrative Templates\System\Windows Time Service\Time Providers

  Configure **Windows NTP Client** settings here.

  Enable **Windows NTP Client** here.

  Enable **Windows NTP Server** here.

| **Warning** |
| --- |
| Some of the preset values that are configured in the System Administrative template file (System.adm) for the Group Policy object (GPO) settings are different from the corresponding default registry entries. If you plan to use a GPO to configure any Windows Time setting, be sure that you review Preset values for the Windows Time service Group Policy settings are different from the corresponding Windows Time service registry entries in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=186066). This issue applies to Windows Server 2008 R2, Windows Server 2008, Windows Server 2003 R2, and Windows Server 2003. |

The following table lists the global Group Policy settings that are associated with the Windows Time service and the pre-set value associated with each setting. For more information about each setting, see the corresponding registry entries in "Windows Time Service Registry Entries" earlier in this subject. The following settings are contained in a single GPO called **Global Configuration Settings**.

**Global Group Policy Settings Associated with Windows Time**

| Group Policy Setting | Pre-Set Value |
| --- | --- |
| AnnounceFlags | 10 |
| EventLogFlags | 2 |
| FrequencyCorrectRate | 4 |
| HoldPeriod | 5 |
| LargePhaseOffset | 1280000 |
| LocalClockDispersion | 10 |
| MaxAllowedPhaseOffset | 300 |
| MaxNegPhaseCorrection | 54,000 (15 hours) |
| MaxPollInterval | 15 |

| | |
|---|---|
| MaxPosPhaseCorrection | 54,000 (15 hours) |
| MinPollInterval | 10 |
| PhaseCorrectRate | 7 |
| PollAdjustFactor | 5 |
| SpikeWatchPeriod | 90 |
| UpdateInterval | 100 |

The following table lists the available settings for the **Configure Windows NTP Client** GPO and the pre-set values that are associated with the Windows Time service. For more information about each setting, see the corresponding registry entries in "Windows Time Service Registry Entries" earlier in this subject.

**NTP Client Group Policy Settings Associated with Windows Time**

| Group Policy Setting | Default Value |
|---|---|
| NtpServer | time.windows.com,0x1 |
| Type | Default options:<br>• **NTP.** Use on computers that are not joined to a domain.<br>• **NT5DS.** Use on computers that are joined to a domain. |
| CrossSiteSyncFlags | 2 |
| ResolvePeerBackoffMinutes | 15 |
| ResolvePeerBackoffMaxTimes | 7 |
| SpecialPollInterval | 3600 |
| EventLogFlags | 0 |

**Network Ports Used by the Windows Time Service**

Windows Time follows the NTP specification, which requires the use of UDP port 123 for all time synchronization communication. This port is reserved by Windows Time and remains reserved at all times. Whenever the computer synchronizes its clock or provides time to another computer, that communication is performed on UDP port 123.

| **Note** |
|---|
| If you have a computer with multiple network adapters (also called a multihomed computer), you cannot selectively enable the Windows Time service based on the network adapter. |

**Related Information**

The following resources contain additional information that is relevant to this section.

- RFC *1305* in the IETF RFC Database

**See Also**

**Other Resources**

Videos about the Windows Time Service

# Network Access Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Access Technologies**

Network access is a critical component of any organization's network infrastructure. Administrators have to dynamically enable and control network access for employees, vendors, business partners, and entire branch offices that are in a variety of physical locations and that are using many different types of devices. Enabling network access, while maintaining security, is an enormous challenge for administrators. You can use the Microsoft Windows Server 2003 operating system to implement a comprehensive network access solution that can include the following variety of technologies:

- Secure access to wireless local area networks (WLANs)
- Secure remote access to intranets over the Internet
- Services that provide centralized network connection authentication, authorization, and accounting
- Tools and services that support the creation and distribution of components to manage secure client-side network access

**Network Access Technologies Architecture**

In Windows Server 2003, there are four network access technology components:

- 802.11 wireless
- Virtual private networks (VPNs)
- Internet Authentication Service (IAS)
- Connection Manager

The following figure shows the components of a network access technology infrastructure. In this example, the Connection Manager service profile initiates remote access client connections to dial-up or VPN-based remote access servers. The remote access servers act as Remote Authentication Dial-In User Service (RADIUS) clients and send the access requests to an IAS server for authentication and authorization. Similarly, the 802.11 wireless client requests access from the wireless access point, which sends the access requests to an IAS server.

The IAS server receives and processes the access requests and queries a user account database to verify access client credentials.

**Network Access Technology Architecture**




**Network Access Technologies Components**
The following network access technologies provide a range of options for secure client access to an organization's network.

**802.11 Wireless**
The WLAN protocol, 802.11, and associated technologies, such as the 802.1X protocol, are IEEE standards supported by Microsoft to provide secure wireless networking that allows you to extend your core network infrastructure to roaming wireless clients.

The IEEE 802.11 wireless standard defines the specifications for the physical layer and the media access control (MAC) sublayer for wireless communications. The 802.1X standard defines port-based network access control used to provide authenticated network access for Ethernet networks. Although this standard was originally designed for wired Ethernet networks, it has been adapted for use on 802.11 wireless LANs.

Until the recent development and wide adoption of WiFi, also known as 802.11b, network clients needed to be physically connected to a LAN to obtain high-speed network access to the LAN.

For more information, see 802.11 Wireless Technical Reference.

**VPN**
A VPN is the extension of a private network that encompasses secure links across shared or public networks, such as the Internet, so that remote clients can securely connect to an organization's network resources. VPN technology also enables remote offices to securely connect to each other so that they can share resources and information over the Internet. A VPN connection uses secure authentication and data encryption to preserve the privacy and integrity of data as it traverses the Internet.

For more information, see VPN Technical Reference.

**IAS**
IAS is the Microsoft implementation of the RADIUS standard defined by the Internet Engineering Task Force (IETF) in RFCs 2138 and 2139. IAS performs centralized authentication, authorization, auditing, and accounting of users who are connecting to a wireless or wired local area network (LAN), using a remote access connection to connect to a private intranet. A server running IAS can be configured to perform a RADIUS server role or a RADIUS proxy role.

Configured as a RADIUS server, an IAS server accepts connection requests from access clients and performs connection authentication and authorization for many types of network access, including wireless, wired, and remote access dial-up and VPN connections. You can use an IAS server to authenticate users in databases on domain controllers running Windows NT Server version 4.0, Windows 2000 Server, or Windows Server 2003. Configured as a RADIUS proxy, an IAS proxy server forwards authentication and accounting messages to a RADIUS server. When used as a RADIUS proxy, an IAS proxy server is a central switching or routing point through which RADIUS access and accounting messages flow.

**Connection Manager**
You can use Connection Manager to create and distribute customized remote access client connections, called service profiles. Service profiles are used to automatically configure remote access clients so they can connect to servers that are running the Routing and Remote Access service or other dial-up or VPN servers by using administrator-defined connection settings. By customizing remote access connections and controlling how users remotely connect to a network, administrators can simplify

their remote access solution, reduce the organizational resources dedicated to assisting remote users, and increase the security of remote connections to the private intranet.

For more information, see Connection Manager Technical Reference.

**Network Access Technologies Scenarios**

The four network access technologies in this section can be deployed in a variety of configurations to meet the specific needs of your environment. When an access client attempts to connect to a wireless access point or a remote access server, the wireless access point or remote access server acts as a RADIUS client and communicates with an IAS server for authentication and authorization to complete the network connection. Connection Manager profiles can be used to automate client-side remote access client configuration

# Connection Manager Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Manager Technical Reference**
Connection Manager is a suite of components that provides administrators with the ability to create and distribute customized remote access connections, called service profiles, and to create, distribute, and automatically update customized phone books. Connection Manager service profiles appear as network connections on client computers, and profiles can be used to connect to remote networks through servers that are running Routing and Remote Access, Internet Authentication Service (IAS), or remote access and virtual private networking technologies from companies other than Microsoft.

Administrators can manually customize and distribute remote access connections to a network, allow users to manually create remote access connections, or both. However, these methods are not easy to document or maintain, and user interfaces and connection properties vary widely. In addition, access requirements are more difficult, if not impossible, to enforce. With the Connection Manager suite, administrators can create and maintain one or more remote access connections customized specifically for access to their networks. Administrators can provide their users with a branded, customized connection experience that is the same on all supported operating systems and for which the administrator has already configured all of the connection properties. By customizing remote access connections and controlling how users remotely connect to a network, administrators who use the Connection Manager suite can simplify their remote access solutions, reduce the organizational resources dedicated to assisting remote users, and decrease the risk that remote connections pose to the network.

**In this subject**


- What Is Connection Manager?
- How Connection Manager Works
- Connection Manager Tools and Settings

# What Is Connection Manager?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Connection Manager?**

**In this section**


- Connection Manager Components
- Common Connection Manager Scenarios
- Connection Manager Dependencies
- Related Information


The Connection Manager suite of components provides administrators with the ability to create and maintain customized remote access connections and customized phone books of access numbers so that users can connect to a network using administrator-defined connection features. Administrators can distribute these customized remote access connections, called service profiles, by using a variety of methods, and they can customize the user interface with text and graphics that are unique to their organizations. Connection Manager service profiles appear in the connections folder on client computers, have their own distinct and brandable user interface, and can be used to connect to remote networks through servers that are running Routing and Remote Access, Internet Authentication Service (IAS), or remote access and virtual private networking technologies from companies other than Microsoft.

By customizing remote access connections and controlling how users remotely connect to a network, administrators who use the Connection Manager suite can simplify their remote access solutions, reduce the organizational resources dedicated to assisting remote users, and decrease the risk that remote connections pose to the network. The Connection Manager suite includes phone book services, so access numbers can be updated every time a user connects to the network without any action required by the user, regardless of the location of the user. Using the Connection Manager suite also helps organizations that have users in multiple geographic locations because administrators can ensure that local access numbers and virtual private network (VPN) servers are listed first. Administrators can create service profiles in multiple languages, providing a similar connection experience for all users regardless of language.

Although administrators can allow users to manually create remote access connections to a network, this method has major disadvantages for large networks with many remote access users. These disadvantages include:

- The procedure for manually configuring remote access connections varies between Microsoft Windows operating systems. Administrators and support personnel would require separate procedures for supporting users on each operating system.
- The procedure for troubleshooting manually configured remote access connections varies between Windows operating systems. Administrators and support personnel would require multiple troubleshooting procedures.
- If any dial-up access numbers change, either the administrator or the user must reconfigure the connection by manually typing the new number. Administrators must manually inform users of access number changes.
- Administrators do not have access to user-configured remote access connections, so enforcing network access requirements (such as the use of antivirus software) on remote computers is difficult.

The Connection Manager suite helps administrators solve these and other issues. Administrators can create connections that have a consistent user interface across all Windows operating systems, use specific authentication protocols, check for required programs, verify registry settings, update files and phone books, or perform any combination of these tasks. Using the Connection Manager suite provides numerous benefits that manually creating connections does not. These benefits include:

**Branding**

Administrators can customize the user interface to identify their organizations. Customizable elements include graphics, icons, and messages. Administrators can also customize online Help and phone books, and they can include support resources that

are unique to their network needs.

**Connection support**
Administrators can customize individual aspects of the user connection, such as by providing routing table updates, by configuring user proxy settings for use during the connection, and by specifying what authentication protocols are used or what VPN protocol to try first when a connection is attempted. Administrators can also provide choices for users, such as a list of VPN servers that users can choose from when they make a connection. Users can save their favorite connection settings. Connection settings are saved based on user logon credentials, so multiple users on the same computer can share a service profile without having to share or re-create their settings.

**Customizable program and application functionality**
Administrators can incorporate custom functionality, including installing and running programs to enhance the connection experience for users. These programs can run automatically at various points during the connection process, such as when users disconnect or when an error occurs.

**Integrated support for access numbers and phone books**
Administrators can specify what phone books and access numbers are available to users. Businesses can contract access number support with multiple large Internet service providers (ISPs) and distribute these access numbers in their phone books. Administrators can configure service profiles to update phone books automatically when users connect to the organization network, eliminating the need for manually maintaining and updating lists of access numbers.

**Logging**
Administrators can configure whether connection attempts are logged by default on user computers and whether users can turn logging on or off. Log files can be used to troubleshoot connection problems.

**Platform support**
Administrators can create service profiles that provide a similar connection experience to users whose computers are running different Windows operating systems. Alternatively, administrators can create service profiles that require a specific Windows operating system to connect to the organization network.

**Connection Manager Components**
The Connection Manager suite consists of Connection Manager, the Connection Manager Administration Kit, and Connection Point Services. Connection Point Services itself consists of Phone Book Administrator and Phone Book Service.

**Connection Manager**
Connection Manager is client connection software that administrators can customize and distribute to users. The customized and distributed Connection Manager software, including all additional programs and applications that administrators include with it, composes a service profile. A service profile consists of all of the files that Connection Manager needs to install itself and then enable users to connect to the network. Administrators can maximize or minimize the identification of the service or organization, depending on what they include in a service profile.

**Connection Manager Administration Kit**
The Connection Manager Administration Kit (CMAK) is a wizard that guides administrators as they create and customize a Connection Manager service profile and then builds the service profile as a compressed, self-installing executable that administrators can distribute to users. Although using the CMAK wizard to build custom service profiles is simple, creating the custom elements requires careful planning and development.

**Phone Book Administrator**
Phone Book Administrator (PBA) is a tool to create and maintain phone book files for use with Connection Manager service profiles. Each phone book is a collection of points of presence (POPs) provided by either the administrator or one or more ISPs with which the organization has contracted to provide Internet access to its users. Each POP provides a local access number and connection settings for a specific region within a country or dependency. PBA compresses and publishes these phone books to a server running Phone Book Service. Administrators who do not intend to include phone books in their service profiles do not need to install PBA.

**Phone Book Service**
Phone Book Service (PBS) is an Internet Information Services (IIS) extension. When a service profile that is configured to check for phone book updates establishes a connection, Connection Manager queries the PBS server for a phone book update. PBS compares the version of the phone book that the service profile is using with the most recent files in the phone book database and then sends any updates to the service profile. Administrators who do not intend to include phone books in their service profiles do not need to install PBS.

**Common Connection Manager Scenarios**
Connection Manager is designed so that administrators can customize it to meet a wide variety of business needs. The following are some of the most common Connection Manager scenarios. Although the scenarios are fictional, administrators who set up and maintain remote access solutions commonly encounter the problems that the scenarios detail.

**Businesses with many remote or traveling users**
A sales company has many representatives who work in the field and who frequently need to make remote access connections to the corporate network. Long-distance charges are prohibitive, and the company does not want the expense of maintaining a large dial-up remote access solution. VPN connections across the Internet are more cost-effective but present more difficulties for inexperienced users. The best solution would be for users to connect to the Internet by dialing a local access number and then make VPN connections to the corporate network. This type of connection is known as a double-dial connection.

The company contracts with several large ISPs to provide local numbers for dial-up access to the Internet across the sales area of the company. The company imports the POP data from the ISPs into a single phone book using PBA and publishes the phone book to a PBS server on the perimeter network. Then the company configures a service profile that allows sales representatives to connect to the Internet by picking a local access number wherever they might be and then seamlessly makes a VPN connection to the corporate network. The company installs the service profile on all corporate laptops and also makes the service profile available for download on the corporate intranet.

**Businesses with specific requirements for network access**
A technology company is increasingly worried about security but wants to continue to allow remote access connections for traveling users and for employees who work from home. The company wants to use Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec) VPN connections, but it is concerned about the ability of their users who are not joined to the corporate domain to get the appropriate certificates that these connections require. Additionally, the company has issued multiple requests for its employees to upgrade the operating systems on their home computers and to run current antivirus software on those computers, but evidence strongly suggests that not all employees have complied. The best solution would be to grant full access to the corporate network only to those remote access connections that comply with the company's requirements, restricting other connections to a quarantined network that has access only to the software that compliance requires.

The company configures a certification authority on their corporate intranet and configures Routing and Remote Access on a server that is running Microsoft Windows Server 2003. Remote access policies are configured for Point-to-Point Tunneling Protocol (PPTP) connections, L2TP connections, and for quarantining all remote access connections that do not meet company requirements for operating system, connection software, and antivirus software. An administrator writes a custom action script and uses the Connection Manager suite and the Windows Server 2003 Resource Kit Tools to create a Connection Manager service profile. This profile checks for required programs, registry settings, and files; quarantines the remote access connection until compliance has been verified; and requests and installs a certificate for an L2TP/IPSec connection. The administrator makes the service profile available for download on the corporate intranet and the quarantine network, and the administrator burns the profile onto CD-ROMs for distribution to new employees.

**Businesses that want to brand their connections**
An ISP has access points and customers all over the world. The customer base uses a diverse set of operating systems, including many versions of Windows operating systems. The company wants to provide a branded, distinctive service to all of its customers, including its own icons, graphics, customized online Help, and troubleshooting steps. The company also wants to reduce the amount of training necessary for its support personnel to troubleshoot customer connections. The best solution would be to provide distinctively branded connections that resemble each other as much as possible in usage and appearance, regardless of operating system and language.

For its Windows users, the company creates phone books in each language that it supports and then publishes them to PBS servers. The company then creates a dial-up service profile in each language that it supports, makes the profiles available on its public Web site, and mails floppy disks with the appropriate service profile to its existing customers.

**Connection Manager Dependencies**
Because Connection Manager is designed to be customizable, the actual dependencies and requirements of individual service profiles vary widely. However, each component of the Connection Manager suite has some consistent dependencies.

**Connection Manager**
Connection Manager requires one of the following operating systems:

- Microsoft Windows 98
- Microsoft Windows 2000
- Microsoft Windows Millennium Edition
- Microsoft Windows XP
- Microsoft Windows Server 2003

Connection Manager also requires a supported version of Internet Explorer to be installed on the same computer. Internet Explorer does not have to be set as the default browser, and the user does not have to use Internet Explorer. Supported versions of Internet Explorer include:

- Internet Explorer 4.01
- Internet Explorer 5
- Internet Explorer 5.5
- Internet Explorer 6

Connection Manager requires the following protocols and interfaces:

- Point-to-Point Protocol (PPP)
- TCP/IP
- Telephony Application Programming Interface (TAPI) for dial-up connections
- PPTP or L2TP for VPN connections

**Connection Manager Administration Kit**
The Connection Manager Administration Kit (CMAK) must be installed on a supported operating system. Although the operating system is the only specific requirement, creating a service profile that meets particular needs might require additional programs or files to be available on the computer on which CMAK is installed. These programs and files include those that the administrator wants to use in custom actions or as additional files.

CMAK must be installed on one of the following operating systems:

- Windows Server 2003
- Windows XP Professional with the Windows Server 2003 Administration Tools Pack

**Phone Book Service**
Phone Book Service (PBS) is an Internet Information Services (IIS) extension. It depends on the following components:

- Internet Information Services
- World Wide Web (WWW)
- File Transfer Protocol (FTP)
- TCP/IP

PBS must be installed on a server running Windows Server 2003.

**Phone Book Administrator**
Phone Book Administrator (PBA) depends on the following components:

- FTP
- TCP/IP

PBA can be installed on either Windows Server 2003 or Windows XP Professional.

**Connection Manager Logical Diagram**

The following figure illustrates how a simple Connection Manager deployment works and how phone books are updated.

**Sample Connection Manager Deployment**




**Related Information**

The following resources contain additional information that is relevant to this section.

- VPN Technical Reference
- "Resource Kit Tools" in Tools and Settings Collection
- Connection Manager Administration Kit
- Routing and Remote Access
- Microsoft Systems Architecture
- Network Access Quarantine Control in Windows Server 2003

# How Connection Manager Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Connection Manager Works**

**In this section**


- Connection Manager Terminology
- Connection Manager Architecture
- Connection Manager Protocols
- Connection Manager Interfaces
- Connection Manager Physical Structure
- Connection Manager Processes and Interactions
- Network Ports Used by Connection Manager
- Related Information


Connection Manager is a suite of components that provides administrators with the ability to create and distribute customized remote access connections and to create, distribute, and automatically update customized phone books. Connection Manager service profiles appear as network connections on client computers, and profiles can be used to connect to remote networks through servers running Routing and Remote Access, Internet Authentication Service (IAS), or remote access and virtual private networking technologies from companies other than Microsoft.

This section provides an in-depth view of how Connection Manager works in an optimal environment. Connection Manager can perform well in many environments; it is designed to be customized for individual network needs. However, for the purpose of this section, an optimal environment is defined as follows:

- Telephone and network infrastructure is in place, address space has been leased, domain names have been registered, and an Internet presence has been established.
- The Active Directory directory service and Group Policy are configured correctly on the network, and appropriate Group Policy settings and permissions are applied.
- Domain Name System (DNS) servers are configured correctly on the network.
- Dynamic Host Configuration Protocol (DHCP) servers are configured correctly on the network.
- Remote access servers are configured correctly and deployed properly on the network.
- Routers are configured correctly on the network.
- Remote Authentication Dial-in User Service (RADIUS) authentication is being used, and it is configured correctly on the network.
- Firewalls and filters are configured correctly on all parts of the network, including between the intranet, the perimeter network, and the Internet.
- Internet Information Services (IIS) is configured correctly.

- A local account (not a domain account) with minimal permissions has been created for posting phone books. This account is disabled except when phone books are being posted.
- File Transfer Protocol (FTP) anonymous access is disabled. FTP is disabled for all accounts except the one that is used to post phone books.
- The FTP service is started immediately before a phone book is posted, and the service is stopped immediately after the phone book has been posted.
- Permissions for the Phone Book Service (PBS) folder and for the PBSData folder in the FTP virtual root have been set appropriately.
- All remote access computers are running the Microsoft Windows XP Professional operating system.
- The Internet service providers (ISPs) with which the organization has contracted service provide phone books or point of presence (POP) data in a database form that administrators can easily import into phone books.

**Connection Manager Terminology**
Before you review Connection Manager components and processes, it is helpful to understand the terminology. The following subsections provide a brief introduction and illustration of the Connection Manager suite and a glossary of terms.

**Connection Manager Components**
Connection Manager refers to both a component and a suite of components. The Connection Manager component is customizable remote access connection software. The Connection Manager suite supports the creation, distribution, and maintenance of customized remote access connections and phone books. The Connection Manager suite consists of the Connection Manager component, the Connection Manager Administration Kit component, and Connection Point Services. Connection Point Services itself consists of Phone Book Administrator and Phone Book Service. Unlike Connection Manager, there is no program named Connection Point Services.

**Connection Manager**
Connection Manager is client connection software that administrators can customize and distribute to users. Administrators can customize many aspects of a Connection Manager service profile, including interface elements, authentication protocols, and programs that run at specific points during the connection. A Connection Manager service profile is compressed into a self-installing, self-extracting executable that can be easily distributed. Most Connection Manager service profiles fit on 3.5-inch disks.

Connection Manager 1.3 is the most recent version; older versions do not have all of the functionality discussed in the Windows Server 2003 Technical Reference. Connection Manager 1.3 is included with the Microsoft Windows Server 2003 and Windows XP operating systems. Service profiles that are created for these operating systems do not need to include Connection Manager 1.3. Service profiles that are created for the Microsoft Windows 2000, Windows Millennium Edition, and Windows 98 operating systems should include Connection Manager 1.3 with the service profile.

**Connection Manager Administration Kit**
To create a Connection Manager service profile, administrators need the Connection Manager Administration Kit (CMAK), which includes:

- a wizard for creating and customizing a Connection Manager service profile and building the service profile as a compressed, self-installing executable
- customizable templates for online Help
- default graphics
- pre-configured custom actions
- customizable templates for service-profile files

Administrators can perform additional customization using a plain-text editor (such as Notepad) to edit service-profile files and then using the CMAK wizard to re-build the service profile. Although using the CMAK wizard is simple, creating a service

profile that exactly meets the needs of a particular network or business environment requires careful planning and development.

CMAK 1.3 is the most recent version, and it is included with Windows Server 2003. Earlier versions of CMAK do not have all of the functionality discussed in the Windows Server 2003 Technical Reference.

**Phone Book Service**
Phone Book Service (PBS) is an Internet Information Services (IIS) extension. For service profiles that are configured to use and check for updated phone books, Connection Manager queries the PBS server after it connects to the Internet. PBS compares the phone book version reported by Connection Manager with the most recent version file. When appropriate, it passes the appropriate update file to the Connection Manager service profile. Administrators who do not intend to include phone books in their service profiles do not need to install PBS.

The most recent version of PBS is included with Windows Server 2003. Earlier versions of PBS do not have all of the functionality discussed in the Windows Server 2003 Technical Reference.

**Phone Book Administrator**
Phone Book Administrator (PBA) is a tool to create, maintain, and post phone book files for use with Connection Manager service profiles. Each phone book is a collection of POPs. Each POP provides a local access number and connection settings for a specific region within a country or dependency. PBA compresses phone books into .cab files, which administrators can post to the PBS server using FTP.

The most recent version of PBA is included with Windows Server 2003 and Windows XP Professional. Earlier versions of PBA do not have all of the functionality discussed in the Windows Server 2003 Technical Reference. Although administrators can use the Windows Server 2003 version of PBS with earlier versions of PBA, the Windows Server 2003 version of PBA is not compatible with earlier versions of PBS.

**Connection Manager Glossary of Terms**
The following terms describe the components and elements of Connection Manager:

**.cab file**
The file name extension for a cabinet file. A cabinet file is a compressed data file that contains phone book information or installation information for Connection Manager. Phone books have two types of .cab files, Full.cab and Delta.cab. Full.cab files contain full phone books. Delta.cab files contain only changes to the phone book. PBA creates a Delta.cab file the first five times that an administrator changes information in a phone book. If the administrator makes a sixth set of changes, PBA creates a Full.cab file.

**.cmp file**
The file name extension for a connection profile file. Every service profile has at least one .cmp file, named *ServiceProfileName*.cmp. The .cmp files contain user-related information. By editing this file, administrators can provide a first-time-only population of user information. Because users can overwrite this information, any settings that administrators specify are available the first time the service profile is used.

**.cms file**
The file name extension for a service provider file. Every service profile has at least one .cms file, named *ServiceProfileName*.cms. The service provider file specifies the configuration of the phone book and most of the other functions of a service profile. Most advanced customization for a service profile is done by editing the .cms file for a particular service profile, by using either the **Advanced Customization** page of the CMAK wizard or a plain-text editor.

**.inf file**
The file name extension for an information file. Every service profile has one .inf file. The .inf file specifies installation information for service profiles. Administrators can configure some setup and uninstallation information in an .inf file, but they should thoroughly test the installation after making any changes. Information files cannot be edited from the **Advanced Customization** page of the CMAK wizard; they must be edited with a plain-text editor.

**.pbk file**
The file name extension for a phone book file. A phone book file is a text file that contains lists of POP information. Phone book files are compressed before they are transferred to the server.

**.pbr file**
The file name extension for a region file. A region file is a text file that helps categorize POPs. A POP can reference one of many geographical regions listed in the region file.

**.sed file**
The file name extension for a connection extraction file. Every service profile has one .sed file. The .sed file contains the instructions for building a self-extracting executable (.exe) file for service profiles. Administrators should never edit any .sed file.

**.ver file**
The file name extension for a version file. Every time that a phone book is updated, the version file is incremented.

**advanced customization**
The process of manually editing service-profile files to achieve specific results. Administrators can use the CMAK wizard to customize most features of Connection Manager service profiles. However, some features require administrators to edit the service profile files, changing how Connection Manager handles certain functions. Administrators can edit these files either by using the **Advanced Customization** page of the CMAK wizard or by building the service profile, editing the service-profile files using a plain-text editor, and then rebuilding the service profile. This process is called advanced customization.

**component profile**
A service profile that is merged into another service profile. Administrators can merge much of the information in existing service profiles into a new service profile by using the CMAK wizard.

**custom action**
An additional program that starts seamlessly at a specified point during the remote access connection. A custom action can be a dynamic-link library (DLL) file; an executable file such as a .bat, .exe, or .cmd file; or a shell-executable file, such as a .txt or .doc file. Custom actions can run at any of nine points during the connection.

**dial-up connection**
A remote access connection that uses a modem to connect to a network.

**direct connection**
A remote access connection that uses a technology such as a digital subscriber line (DSL) or a cable modem to connect to a network.

**disconnect action**
A custom action that runs immediately before the connection ends. Disconnect actions run even if Connection Manager did not initialize the disconnection. For example, if a disruption in telephone service terminates a connection, Connection Manager will attempt to run the disconnect actions specified in the service profile after the unexpected termination.

**double-dial connection**
A remote access connection that first uses a modem to connect to an ISP and then makes a VPN connection to a specific network.

**key**
The generic term for an entry in a Connection Manager service-profile file. All valid entries in service-profile files have key names. Each key must have an appropriate value to be valid.

**monitored action**
A custom action that runs after a connection is established and, for a VPN connection, after the tunnel is established. Each monitored action runs every time the user connects to the service, whether through a dial-up connection, a direct connection, or a double-dial connection. All monitored actions must be .exe files because monitored actions run asynchronously. Connection Manager monitors the status of all monitored actions and starts the disconnect sequence when the last monitored action closes.

**on-cancel action**
A custom action that runs as soon as users click **Cancel** during a connection attempt. On-cancel actions do not run when users click the **Cancel** button to close Connection Manager.

**on-error action**
A custom action that runs whenever an error occurs during a connection.

**phone book**
A collection of one or more POP entries, with each POP supplying a telephone number that provides dial-up access to an intranet or an ISP. Phone books give users complete POP information, so when they travel they can connect to different Internet access points rather than being restricted to a single POP.

**post-connect action**
A custom action that runs after a connection is established and, for a VPN connection, after the tunnel is established. Each post-connect action runs every time the user connects, whether through a dial-up connection or a direct connection.

**pre-connect action**
A custom action that runs as soon as users click **Connect.** These actions run before Connection Manager establishes a connection to the service.

**pre-dial action**
A custom action that runs after users click **Connect** but before the computer starts to dial the connection to the service. Pre-

connect actions run before pre-dial actions.

**pre-init action**
A custom action that runs as soon as users start Connection Manager. These actions run before the Connection Manager logon screen appears.

**pre-tunnel action**
A custom action that runs after a connection with the ISP is established but before a tunnel to the VPN server is established. This type of action is available only if the service profile is configured for VPN connections, and it will run only when users are using the VPN connection option.

**section**
A header within service-profile files. Section names are always contained within brackets. Some section names are preset; others can be added by administrators.

**service profile**
A customized Connection Manager remote access connection used to connect to an ISP, a corporate network, or other network. Service profiles are occasionally referred to as Connection Manager profiles.

**service-profile files**
The collective term for all files needed to build a Connection Manager profile, including but not limited to the .inf, .sed, .cms, and .cmp files.

**top-level profile**
A profile that contains information from other service profiles. Administrators can merge much of the information in existing service profiles into a new service profile by using the CMAK wizard.

**value**
The data required to configure a key in service-profile files.

**Connection Manager Architecture**
The following figure illustrates how the components of the Connection Manager suite work together. This figure illustrates a Connection Manager service profile that is configured to use phone books. VPN-only service profiles, which do not use phone books, require neither PBA nor PBS.

**Overall Architecture for Connection Manager**




The following table describes the components of the Connection Manager architecture.

**Connection Manager Architecture Components**

| Component | Description |
|---|---|
| Phone Book Administrator (PBA) | Used to create and maintain phone books that Connection Manager uses. |

| | |
|---|---|
| Connection Manager Administration Kit (CMAK) | Used to create Connection Manager service profiles. |
| Phone Book Service (PBS) | Used to distribute phone books to Connection Manager. |
| Connection Manager | Used to connect to a remote network. In the architecture in the previous figure, the service profile first connects to an ISP using a POP from its phone book. Then the service profile makes a VPN connection to the remote network. |
| Remote access server | Used to provide remote access. For example, a remote access server might be a computer that is running Windows Server 2003 and Routing and Remote Access. |

**Connection Manager Protocols**

Administrators can customize Connection Manager service profiles to use specific protocols. For example, administrators can specify the use of data-link protocols such as Point-to-Point Tunneling Protocol (PPTP) and transport protocols such as NetBEUI. Additionally, programs that administrators include for custom actions could use a variety of protocols. The following table shows the protocols that the Connection Manager suite uses by default.

### Connection Manager Protocols

| Component | Description |
|---|---|
| Point-to-Point Protocol (PPP) | A data-link layer protocol for transmitting data across point-to-point links. PPP allows remote access technologies (including Connection Manager) and devices to interoperate. |
| TCP/IP | A protocol suite that provides communication across interconnected networks such as the Internet. |
| Hypertext Transfer Protocol (HTTP) | An application-layer protocol that specifies the client/server interaction between Web browsers and Web servers. Connection Manager service profiles use HTTP to communicate with PBS servers. |
| File Transfer Protocol (FTP) | An application-layer protocol that is used to transfer files between hosts on a TCP/IP network. FTP is used to post phone books to PBS servers. |

**Connection Manager Interfaces**

By default, Connection Manager service profiles use some remote access application programming interfaces (APIs) and Telephony Application Programming Interface (TAPI). Administrators can customize service profiles to use any of the remote access APIs documented on Microsoft Developer Network (MSDN). In addition to using remote access APIs, Connection Manager has its own set of macros, DLL parameters, and registry key values for use with custom actions.

**Macros**

Administrators can use some command-line macros in custom actions to pass arguments. When these macros are used, Connection Manager replaces them with the actual run-time information for the parameter. The following table describes the command-line macros that custom actions support.

### Connection Manager command-line macros

| Macro | Description |
|---|---|
| %ServiceName% | The service name of the profile. |
| %UserPrefix% | The user-name prefix used for this connection. |
| %UserSuffix% | The user-name suffix used for this connection. |
| %UserName% | The user name without any realm user-name prefix or suffix. |
| %Profile% | The location and file name of the active connection profile (.cmp) file. |
| %ServiceDir% | The path to the profile directory. |

| | |
|---|---|
| %Domain% | The Active Directory domain for the connection. |
| %InetUserName% | The user name for the Internet connection. |
| %ConnectionType% | A value that identifies the connection type: 0, 1, or 2.<br>0 = dial-up<br>1 = direct<br>2 = double-dial |
| %DialRasPhoneBook% | The full path to the phone book file. |
| %TunnelRasPhoneBook% | The full path to the phone book file used for the VPN portion of this connection. |
| %DialRasEntry% | The service name or remote access entry name for the dial-up connection. |
| %TunnelRasEntry% | The service name or remote access entry name for the tunnel connection. |
| %AutoRedial% | A Boolean value that is 1 if this dial attempt is an automatic redial or 0 if not an automatic redial. |
| %PopName% | The description for the phone number in the phone book. |
| %ErrorCode% | The Win32 error code. |
| %LastErrorSource% | The origin of the last error. |
| %CurrentFavorite% | The connection settings that the user saved in the **Settings Saved As** box. |
| %TunnelServerAddress% | The IP or DNS address of the VPN server, if any. |
| %ClientIPAddress% | The IP address of the computer on which the Connection Manager profile is installed. |
| %ServerIPAddress% | The IP address of the remote access server, if any. |
| %Interactive% | A Boolean value that is used in a custom action to determine whether to display a user interface. Administrators can incorporate this macro in a custom action to display an interactive user interface (such as an error message). Administrators should use this macro only with programs that can either complete (if Connection Manager is running in an interactive state) or that are able to take other action, such as failing gracefully without an error message (if Connection Manager is running in a non-interactive state). This macro was designed to be used in conjunction with the **Program interacts with the user** check box. For example, if the **Program interacts with the user** check box is cleared, the %interactive% macro should be added to the custom action parameters so that the custom action can behave according to the state in which Connection Manager is running. Using this macro might require modification to the custom action itself. |

When parameters are specified, Connection Manager reads the parameter string and passes it to the **pszCommandLine** parameter of the DLL or sends the string as part of the command line to an executable (available through the Windows API **GetCommandLine**).

If the custom action is a DLL, the first token of the parameter string is the function entry point to call. This token is removed from the parameters before the string is passed to the DLL (so the **pszCommandLine** string does not start with the DLL entry point name as the first parameter). When creating the DLL, administrators must ensure that the exported name of the DLL matches the name specified as the first parameter, because name decoration can result in different names. (For example, in C++, name decoration can cause the exported name to be different from the actual function name.)

Connection Manager cannot denote a null pointer in a string, so it passes the string NULL (which it passes for all undefined parameters). Use quotes only when a parameter contains spaces. When using an executable other than a DLL, the function name is not required.

**DLL Parameters**

The following table describes the DLL parameters that Connection Manager supports. DLLs only run synchronously; Connection Manager starts the action and then waits for the function to return before continuing. The first argument is the function name within the DLL to call. Administrators specify the argument in the Parameters box in the **Add/Edit Custom Actions** page of the CMAK wizard.

**Connection Manager DLL Parameters**

| Parameter | Description |
|---|---|
| hWndParent | Handle to the Connection Manager logon dialog box or NULL; used as the parent window for any user interface that the custom action displays. |
| hinstDll | Handle to the instance of this DLL. |
| pszCommandLine | String pointer to the command-line arguments. |
| dwReserved | Reserved for future use. |

The syntax for each of these custom actions is:

```
HRESULT WINAPI function
    [IN] HWND hWndParent
    [IN] HINSTANCE hinstDll
    [IN] LPCSTR pszCommandLine
    [IN] DWORD dwReserved
```

Custom action DLLs often have a comment that appears while they run. These comments have the form **Running Custom Action** *Description*, and they are specified in the **Description** box on the **Add/EditCustom Actions** page of the CMAK wizard. The Connection Manager interface is frozen and does not accept input while this type of action runs.

Custom action DLLs require a return value. If the return value is less than 0, the DLL call fails. (SUCCEEDED macro returns False.) In this case, an error message (including the return value) appears in the form "**Custom Action** *Description* **failed** *ReturnValue*." If a pre-init, pre-dial, pre-connect, pre-tunnel, or post-connect DLL action fails, the connection attempt is ended.

The DLL is unloaded after the function call. For DLLs that require extended times for the program to run, administrators should implement WM_PAINT messages in the DLL to ensure that user actions do not disrupt graphics.

Administrators should provide the exact name of exported functions when they build custom action DLLs. When using a language that changes the exported name (such as C++), administrators should use a .def file to preserve the exported function names.

**Registry key values**

For the most part, administrators do not need to set registry keys for use with Connection Manager. The exception is if the administrator wants custom actions to run on user computers before users have logged on. For security reasons, custom actions are disabled by default during logon; they will not run if users select the **Log on using dial-up networking** check box at the logon screen. For a custom action to run during the logon process, values must be specified for the **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Connection Manager***ProfileName***\WinLogon Actions** registry key. The following table describes the fields for the registry key.

**Logon Registry Key Fields**

| Field | Description |
|---|---|
| Name | The name of the executable file that will run. |
| Type | REG_DWORD |
| Data | A value indicating the location of the executable. Supported values for the Data field are:<br><br>**Value**: 0x00000000(0)<br><br>**Location of executable:** %windir%\system32<br><br>**Value**: 0x00000001(1)<br><br>**Location of executable:** Profile directory |

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use

extreme caution.

For more information about remote access APIs, see the Microsoft Platform SDK on MSDN.

**Connection Manager Physical Structure**
Each component of the Connection Manager suite requires physical structures to be stored on the computer on which the component is installed. The following figure illustrates components in the Connection Manager suite and where they exist in relation to each other.

**Physical Placement of Connection Manager Suite Components in a Network Architecture**




At a minimum, Connection Manager requires at least 2 MB to install. The amount of disk space actually required to install a service profile varies depending on how the service profile has been configured. Depending on the service-profile configuration, Connection Manager has the following requirements:

- One of the following operating systems:
    - Microsoft Windows 98
    - Microsoft Windows 2000
    - Microsoft Windows Millennium Edition
    - Microsoft Windows XP
    - Microsoft Windows Server 2003
- A supported version of Internet Explorer. Supported versions of Internet Explorer include Internet Explorer 4.01, Internet Explorer 5, Internet Explorer 5.5, or Internet Explorer 6. Users do not have to set Internet Explorer as their default browser or use the software.
- For dial-up connections, a 28.8 Kbps modem or faster connection. Connection Manager can automatically configure the modem.

After installation, Connection Manager creates directories on the system drive and stores information needed to connect in these directories as follows:

- **\Documents and Settings\ *User* \Application Data\Microsoft\Network\Connections\Cm** This directory contains the .cmp file for each service profile. Connection Manager also creates a subdirectory for each service profile, using the eight-character name for the service profile. This subdirectory contains the rest of the service-profile files, including the .cms file, the .pbk file, all custom icons and graphics, and any custom action files.
- **\Documents and Settings\ *User* \Local Settings\Temp** This directory is the default location for the Connection Manager log files. Administrators can specify where the log files are stored through advanced customization.

- **\Software\Microsoft\Connection Manager** Depending on the service profile and how it is installed, Connection Manager creates the necessary registry keys for the service profile in **HKEY_CURRENT_USER**, **HKEY_LOCAL_MACHINE**, or both.

CMAK requires Windows Server 2003 and 5 MB of free space in order to be installed. After installation, CMAK creates directories on the system drive and stores information needed to create and build service profiles as follows:

- **\Program Files\CMAK\Profiles\Support** This directory contains all the template service-profile files, DLL files for pre-configured custom actions such as automatic proxy configuration, DLL files for installation on specific Windows operating systems, and information files. It also contains a subdirectory, \CMHelp, which contains the template Help files for administrators to customize.

- **\Program Files\CMAK\Profiles\ *ServiceProfileName*** When a service profile is created, CMAK creates a subdirectory in the Profiles directory using the eight-character service profile name and stores the service-profile files in that subdirectory.

PBA requires Windows Server 2003 or Windows XP Professional and 1 MB of free space in order to be installed. After installation, PBA creates directories on the system drive and stores information needed to create and post phone books as follows:

- **\Program Files\PBA** This directory contains the program file, as well as the template database files, the country.txt file, and the context-sensitive Help file for PBA.

- **\Program Files\PBA\ *PhoneBookName*** When a phone book is created, PBA creates a subdirectory in its program directory using the eight-character phone book name and stores all the phone book files in that subdirectory.

- **\Program Files\PBA\ *PhoneBookName* \ *PostNumber*** When a phone book is posted, PBA creates an enumerated folder in the phone book subdirectory and stores the cabinet file and the phone book database file in it.

PBS requires Windows Server 2003 and less than 1 MB to install. However, it also requires IIS to run, so the total installation requires at least 11 MB. After installation, PBS creates directories on the system drive and stores information needed to distribute phone book updates as follows:

- **\Program Files\Phone Book Service\Bin** This directory contains the DLL files for PBS. A virtual directory, named PBServer, for the \Bin directory is created in IIS under the default Web site.

- **\Program Files\Phone Book Service\Data** The Data directory contains subdirectories for each phone book published to the PBS server. These subdirectories contain the cabinet files and a database subdirectory. The Data directory and its subdirectories have virtual directories created for them in IIS. The virtual directory for the Data directory is called PBSData.

**Performance Limitations for Connection Manager**

Phone Book Administrator and Phone Book Server have some performance limitations and requirements.

**Phone Book Administrator**

- A phone book name must contain no more than eight characters, cannot consist of all digits, and must not contain a space or any symbols. Symbols include but are not limited to: **! , ; * = / \ : ? < > | . & % { } [ ] @ ( ) ` ~**

- Phone books can be updated no more than 32,676 times.

- Phone books can contain no more than 65,000 POPs.

- When administrators import POPs to a phone book that has not yet been posted, the import file must contain no more than 32,000 POPs.

- When administrators import POPs to a phone book that has already been posted, the import file must contain no more

than 6,000 POPs (combined adds, edits, and deletes) for each subsequent post.

- Phone Book Administrator is available only in English, French, German, Japanese, and Spanish. Administrators can create phone books in another language by running PBA on an operating system that has been optimized for that language.

- The same phone book cannot be published to more than one phone book server. Publishing the same phone book to more than one phone book server causes version conflicts that will prevent users from obtaining the most recent version of the phone book. However, administrators can use the Distributed File System (DFS) console to automate replication of phone books to multiple servers from the server to which the phone book was published. Administrators can also replicate phone books manually by copying the folder to the other servers.

- Deleting a phone book removes the indexed references. However, released .cab files remain in the file system of the computer on which PBA is running. To remove the .cab files, administrators must manually delete the phone book directory and all of its contents.

**Phone Book Servers**

When a user logs on using a service profile that automatically requests phone book updates, the Phone Book Service (PBS) server receives an HTTP query, initiating the phone-book update process. The maximum number of hits per second depends on the size of the update (.cab) file. Larger update files can slow performance by as much as 25 percent.

The following table details how many hits the listed processors are capable of handling on a dedicated server with 128 megabytes (MB) of RAM. These estimates are based on small update files of about 5 kilobytes (KB).

**Performance Estimates for PBS Servers**

| Processor | Hits/sec | Hits/hour | Hits/day |
|---|---|---|---|
| Intel Pentium III 600 MHz | 275 | 990,000 | 23,760,000 |
| Intel Pentium III 500 MHz (dual processors) | 250 | 900,000 | 21,600,000 |
| Intel Celeron 400 MHz | 125 | 450,000 | 10,800,000 |
| Intel Pentium II 300 MHz | 125 | 450,000 | 10,800,000 |

Deleting a phone book removes the indexed references. However, released .cab files remain on the file system of the PBS server. To remove the .cab files, administrators must manually delete the phone book directory and all of its contents.

**Connection Manager Processes and Interactions**

Connection Manager service profiles are customizable, so the exact processes and interactions between a Connection Manager service profile and other components vary. The following is a brief description of what happens during the creation, distribution, and usage of a double-dial Connection Manager service profile that includes a phone book and automatically checks for phone book updates as its first post-connect custom action.

1. An administrator creates a phone book using PBA.

    a. PBA creates a subdirectory for the phone book files.

    b. PBA creates the version file, the phone book files, and the cabinet files in the phone book subdirectory.

2. The administrator posts the phone book to the PBS server.

    a. PBA creates a post subdirectory in the phone book file subdirectory.

    b. PBA creates a database file in the post subdirectory and copies the cabinet files from the phone book file directory into that subdirectory.

    c. An FTP control session is opened between the computer on which PBA is running and the PBS server on TCP port 21. PBA sends the user name and password of the posting account to the PBS server. The PBS server authenticates the account credentials and sends a response. PBA opens an FTP data session on TCP port 20 and uploads the

contents of the post subdirectory.

d. PBS creates a subdirectory for the phone book under its Data subdirectory. The cabinet file is copied into this subdirectory, and the phone book database is copied into the Database directory.

e. Both FTP sessions are closed.

3. The administrator creates a service profile using CMAK.

   a. CMAK creates a subdirectory for the service profile.

   b. CMAK creates the service-profile files and copies all custom files (including but not limited to graphics, online Help files, and programs for custom actions) and additional files into the service profile subdirectory.

   c. CMAK creates the self-installing executable and saves it to the service profile subdirectory.

4. The administrator distributes the Connection Manager service profile by copying the service profile onto floppy disks and distributing the disks to users.

5. The user installs the service profile on a home computer that is running Windows XP.

   a. Connection Manager creates a subdirectory for the service-profile files in the appropriate user context and unpacks the service-profile files.

   b. Connection Manager creates registry keys in the appropriate user context as the service profile configuration requires (for example, for custom actions).

   c. An icon for the service profile is created in the Network Connections folder.

6. The user opens the Network Connections folder and opens the Connection Manager service profile.

   a. Any pre-init custom actions are run.

   b. The Connection Manager logon screen appears.

7. The user provides the information that this connection requires, as the administrator determined when creating the service profile, and clicks **Connect**.

   a. Any pre-connect custom actions are run.

   b. Any pre-dial custom actions are run.

   c. Connection Manager dials the connection.

      — The ISP answers the call and negotiates a connection speed.

      — The ISP authenticates the user name and password and establishes the connection to the Internet.

   d. Any pre-tunnel custom actions are run.

   e. Connection Manager makes the VPN connection to the corporate network.

      — The remote access server on the corporate network answers the connection request.

      — The user name and password are authenticated, and any Group Policy settings are applied.

— The connection to the corporate intranet is authorized.

f. The Automatically download phone-book updates post-connect action is run.

— Connection Manager sends an HTTP GET request on port 80 to the PBS server using the following format: http://*PhoneBookServerName*/pbserver/pbserver.dll?osarch=0&ostype=0&osver=1&cmver=1&lcid=1033&pbver=*Version*&pb=*PhoneBookName*

— PBS checks the version number of the phone book that is enumerated in pbver= against the version number of the phone book that PBS has. If the version number in the request is the same or higher than the version that PBS has, PBS provides no phone book update. If the number is lower but less than five versions lower, PBS sends a delta phone book. If the number is more than five versions lower, PBS sends a full phone book. For example, assume that the version number of the phone book that PBS has is 10. If the version number in the request is also 10, PBS does nothing. If the version number in the request is 7, PBS sends a delta phone book. If the version number in the request is 3, PBS sends a full phone book.

— If PBS has sent a phone book update, Connection Manager installs the phone book.

g. Any other post-connect custom actions are run.

h. Any monitored action custom actions are run.

8. The user disconnects from the corporate network.

   a. Before the connection is closed, any disconnect custom actions are run.

   b. Connection Manager closes the connection.

**Network Ports Used by Connection Manager**

The following table describes the network ports that Connection Manager service profiles commonly use. Connection Manager service profiles can include custom actions that run scripts or tools that might use additional protocols, so the exact network ports that service profiles use will vary.

**Port Assignments for Connection Manager**

| Service Name | UDP | TCP |
|---|---|---|
| HTTP | | 80 |
| FTP | | 20/21 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- VPN Technical Reference
- Microsoft Platform SDK on MSDN
- Resource Kit Tools in Tools and Settings Collection
- Connection Manager Administration Kit
- Routing and Remote Access
- Microsoft Systems Architecture
- Network Access Quarantine Control in Windows Server 2003

# Connection Manager Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Manager Tools and Settings**

**In this section**


- Connection Manager Tools
- Connection Manager Predefined Custom Actions
- Connection Manager Keys and Values
- Connection Manager Registry Entries
- Network Ports Used by Connection Manager
- Related Information


This section contains information about the tools, predefined custom actions, keys and values, registry entries, and ports that are associated with Connection Manager. One of the tools for Connection Manager, the Connection Manager Administration Kit, includes several predefined custom actions. Predefined custom actions are executables that require minimal configuration from the administrator to perform common connection tasks, such as automatically downloading phone book updates. In addition to configuring predefined custom actions, administrators can configure hundreds of individual settings, called keys, to precisely control individual aspects of a Connection Manager service profile.

**Connection Manager Tools**

The following tools are associated with Connection Manager.

**Cmak.exe: Connection Manager Administration Kit**

**Category**

Operating system (optional component)

**Version compatibility**

Connection Manager Administration Kit 1.3 is available as part of the Management and Monitoring Tools in Microsoft Windows Server 2003. It is compatible with Windows Server 2003 and with Microsoft Windows XP Professional if the Windows Server 2003 Administration Tools Pack is installed.

The Connection Manager Administration Kit wizard guides administrators through building service profiles that meet their networking needs. Administrators can configure predefined custom actions, define and include their own, or both. Administrators can customize most elements of a service profile, from graphics to authentication protocols. At the end of the wizard, the service profile is created and compressed into a self-installing, self-extracting executable that administrators can easily distribute.

For more information about the Connection Manager Administration Kit, see Connection Manager Administration Kit.

**Pbserver.dll: Phone Book Service**

**Category**

Operating system (optional component)

**Version compatibility**

Phone Book Service is available as part of the Management and Monitoring Tools in Windows Server 2003. It is compatible with Windows Server 2003.

Phone Book Service (PBS) is an Internet Information Services (IIS) extension. When a service profile that is configured to check for phone book updates establishes a connection, the service profile queries the PBS server for a phone book update. PBS compares the version of the phone book that the service profile is using with the most recent files in the phone book database and then sends any updates to the service profile.

**Pbadmin.exe: Phone Book Administrator**

**Category**
Valueadd tool

**Version compatibility**
Phone Book Administrator 1.2 is a Valueadd tool for Windows Server 2003 and Windows XP Professional. It is compatible with Windows Server 2003 and with Windows XP Professional.

Phone Book Administrator is a tool to create and maintain phone book files for use with Connection Manager service profiles. Each phone book is a collection of points of presence (POPs). Each POP provides a local access number, including connection settings, for a specific region within a country or dependency. PBA compresses phone books into .cab files and posts the files to the PBS server using File Transfer Protocol (FTP).

**Connection Manager Predefined Custom Actions**
The following predefined custom actions are included with Connection Manager Administration Kit.

**Cmdl32.exe: Automatic Phone Book Download**
**Custom Action Type**
Post-connect

**Purpose**
Automatically check for phone book updates; download and install updates as necessary.

**Configuration**
In the CMAK wizard, on the **Phone Book** page, select the **Automatically download phone book updates** check box, specify the URL for the PBS server in **Connection Point Services server**, and specify the name of phone book in **Phone Book Name**.

Associated keys:

- PBFile
- RegionFile
- PBURL

For more information about automatic phone book downloads, see Providing phone book support. For more information about Connection Manager keys, see "Connection Manager Keys and Values" later in this section.

**Cmproxy.dll: Automatic Proxy Configuration**
**Custom Action Type**
Post-connect (configure user proxy settings); Disconnect (restore user proxy settings)

**Purpose**
Ensure that the user has appropriate access to internal and external resources by reconfiguring user proxy settings while the remote computer is connected to the organizational network.

**Configuration**
In the CMAK wizard, click **Automatically configure proxy settings**, specify the proxy file in **Proxy settings file**, and optionally select the **Restore the users' previous proxy settings after disconnecting** check box.

Associated keys: none

To use automatic proxy configuration, administrators must include proxy settings files in plain text with the service profiles. The following table describes the syntax elements of a proxy setting file:

| Key | Value |
|---|---|
| [Automatic Proxy] | No value. Section header for automatic proxy settings. |
| [Manual Proxy] | No value. Section header for manual proxy settings. |

| | |
|---|---|
| AutoProxyEnable | 0 = Automatic proxy detection is disabled.<br><br>1 = Automatic proxy detection is enabled.<br><br>**Note**: This feature is not supported for Internet Explorer 4 and earlier. |
| AutoConfigScriptEnable | 0 = Automatic proxy script will not be executed.<br><br>1 = Automatic proxy script will be executed. |
| AutoConfigScript | The URL that points to the automatic proxy script. |
| ProxyEnable | 0 = Manual proxy detection is disabled.<br><br>1 = Manual proxy detection is enabled. |
| ProxyServer | The name and the port of the proxy server. |
| ProxyOverride | The Domain Name System (DNS) addresses, separated by semi-colons, that should not be passed on to the proxy server. Use the <local> value to bypass the proxy server for local addresses. If the <local> value is included, it must be the final value in the string. |

For more information about automatic proxy configuration, see Using automatic proxy configuration.

**Cmroute.dll: Routing Table Updates**
**Custom Action Type**
Post-connect

**Purpose**
To manage network traffic and security by reconfiguring user routing tables while the remote computer is connected to the organization network.

**Configuration**
In the CMAK wizard, on the **Routing Table Update** page, click **Define a routing table update**, specify the route file in **Route file to include**, optionally specify the URL for an updated route file in **URL to route file**, and optionally select the **If this URL is unavailable, disconnect the client** check box.

Associated keys: none

To use routing table updates, administrators must include routing table update files in plain text with the service profiles. Each line in those files adds or deletes a route. The syntax of each line is as follows:

```
Command Destination
MASK
Netmask Gateway
METRIC
Metric
IF
Interface
```

Certain parameters can contain the value of **default**. In those cases, the appropriate information from the client computer is used. The following table provides more information on each parameter.

| Parameter | Description |
|---|---|
| *Command* | Specifies whether to add or delete a route.<br><br>**ADD** adds a route.<br><br>**DELETE** deletes a route. |
| *Destination* | Specifies the destination IP address on the target network. |
| **MASK** | Specifies that the next parameter is the netmask value. |

| | |
|---|---|
| *Netmask* | Specifies the netmask. |
| *Gateway* | Specifies the gateway. Should contain the value of **default**. |
| **METRIC** | Specifies that the next parameter is the cost for the destination. The **METRIC** parameter is optional. If **METRIC** is not included, do not include *Metric*. |
| *Metric* | Specifies the cost for the destination. Should contain the value of **default**; its default value is 1. |
| **IF** | Specifies that the next parameter is the interface number. |
| *Interface* | Specifies the interface number. Should contain almost always contain the value of **default**. |

To remove the default gateway, you can add the following line to the update file:

```
REMOVE_GATEWAY
```

Unlike the **ADD** and **DELETE** commands, the **REMOVE_GATEWAY** command requires no further syntax.

**Routing table precedence**
During the connection process, update files are applied in a specific sequence.

1. All update files are retrieved.

2. If a service profile contains an update file (static file), routing information from that file is applied to the client computer first.

3. If a service profile contains an update URL, routing information from that file (dynamic file) is applied to the client computer after the static file. The downloaded dynamic file is then immediately deleted.

4. If either the static file or the dynamic file contains the REMOVE_GATEWAY command, this command is run after all other routes have been applied.

Because of this order of precedence, a dynamic file could add routes, replace routes specified in the static file, or both.

If you use the default interface value, the routing changes that you specify are reversed when the user disconnects.

For more information about routing table updates, see Including routing table updates.

**Cmdl32.exe /VPN: Download Update for VPN File**
**Custom Action Type**
Post-connect

**Purpose**
Allows the automatic downloading of an updated VPN file from a specified server. For service profiles that support VPN files, this file is used to update the list of servers available on the **VPN** tab in the properties dialog box for the service profile. Users can choose a VPN server from the list to use for the connection. This tab also contains a customizable text message.

**Configuration**
In the CMAK wizard, on the **VPN Support** page, select the **Phone book from this profile** check box, click **Allow the user to choose a VPN server before connecting**, and specify a VPN file.

Associated keys:

- TunnelFile

- TunnelDUN

To enable users to choose a VPN server from a list when they connect to the network, administrators must include a VPN file in plain text with the service profile. The following table identifies sections, keys, and values that compose a VPN file.

| [Section] or key name | Value |
|---|---|
| [Settings] | Section header for the keys that contain VPN settings. |
| Default | The friendly name of the default VPN server for this service profile. If you create this key but do not give it a value or if you do not create this key, the user must select a VPN server the first time the user connects using the service profile. |
| UpdateURL | The URL of the Web server that contains updates for this VPN file. If this key is added, a post-connect action is added automatically to the service profile. This post-connect action updates the VPN file. |
| Message | The text that appears on the **VPN** tab in the properties dialog box for the service profile. This message cannot exceed 256 characters, and it must be a single paragraph. If this key is created but left blank, default text is used. |
| [VPN Servers] | Section header for the keys that identify VPN servers. |
| *FriendlyName* | The name of this key is the friendly name of one of the VPN servers available to your users. The value of this key is the DNS or IP address of the VPN server for which you named the key. If you want this VPN server to use a specific set of network and security settings, you must follow the DNS address with a comma and the name of a VPN entry that you will edit in the **VPN Entries** page of the CMAK wizard. If you do not specify a VPN entry, the VPN server will use the settings for the default VPN entry. |

For more information about user choice of VPN servers, see Implementing VPN support.

**Connection Manager Keys and Values**

The following keys and values are associated with Connection Manager. The online Help for the Connection Manager Administration Kit includes tables of all keys, settings, and values for Connection Manager, sorted by section header types as they are found in .cms files. The tables below are sorted by usage type. For more information about keys and values, see Advanced Customization.

**Connection Manager Section Headings**

**These section headers define the keys that are contained within the sections.**

| [Section] | Value | Set in |
|---|---|---|
| [Animated Actions] | Specifies the keys that control when to play segments of the animation. | CMAK<br>.cms |
| [Animated Logo] | Specifies the animation for the Connection Manager logon dialog box.<br>The [Animated Logo] section specifies the location of the animation (.spl or .swf) file. This function is provided for compatibility with Connection Manager 1.0, which supports Macromedia Flash animation files.<br>The AnimatedLogo key in the [Connection Manager] section determines whether an animation is used. | .cms |
| [Auto Applications] | Specifies monitored actions that run during a connection. Monitored actions are similar to custom actions, but they cause the disconnect sequence to start as soon as the last monitored application closes. This section uses the following keys: #, #&Flags, and #&Description. Unlike other custom actions, monitored actions do not support DLLs. | CMAK<br>.cms |
| [Connect Actions] | Specifies the post-connect actions that run after a connection is established and, for VPN connections, after the tunnel is established. This section uses the following keys: #, #&Flags, and #&Description. | CMAK<br>.cms |
| [Connection Manager] | Specifies miscellaneous service profile keys and values. | CMAK<br>.cms<br>.cmp |

| | | |
|---|---|---|
| [CustomButton Actions] | Specifies the keys that control what programs start when the custom button is clicked. The format of the value is # = *Program*, where # is an ID number from 0 to *n* and *Program* is the complete name of the program, including the extension and the path to the program. Programs start in sequence: 0 first, 1 next, and so on. Custom button programs cannot be DLLs. | CMAK .cms |
| [Disconnect Actions] | Specifies the disconnect actions that run immediately before disconnecting. If a connection is prematurely dropped, these actions are run immediately after the disconnect. This section uses the following keys: #, #&Flags, and #&Description. | CMAK .cms |
| [Logging] | Specifies the logging keys. | .cms |
| [Menu Options] | Specifies the keys that control the shortcut menu that appears when a user right-clicks the notification area shortcut for your service profile. | CMAK .cms |
| [Networking&*TunnelDUN*] | Specifies the configuration to be used for all VPN connections. In this section title, *TunnelDUN* is the value of the TunnelDUN key in the [Connection Manager] section. | CMAK |
| [On-Cancel Actions] | Specifies the on-cancel actions that run when the user cancels a session after starting the connection attempt (any time that the user clicks **Connect** and then clicks **Cancel**). This section uses the following keys: #, #&Flags, and #&Description. | CMAK .cms |
| [On-Error Actions] | Specifies the on-error actions that run when an error occurs. This section uses the following keys: #, #&Flags, and #&Description. | CMAK .cms |
| [Pre-Connect Actions] | Specifies the pre-connect actions that run before a connection is established. This section uses the following keys: #, #&Flags, and #&Description. | CMAK .cms |
| [Pre-Dial Actions] | Specifies the pre-dial actions that run immediately before dialing. This section uses the following keys: #, #&Flags, and #&Description. | CMAK .cms |
| [Pre-Init Actions] | Specifies the pre-init actions that run as soon as the user accesses the profile. This section uses the following keys: #, #&Flags, and #&Description. | CMAK .cms |
| [Pre-Tunnel Actions] | Specifies the pre-tunnel actions that run after a connection is established but before a VPN connection is established. This section uses the following keys: #, #&Flags, and #&Description. | CMAK .cms |
| [Profile Format] | Specifies the service profile version. | CMAK .cms .cmp |
| [Scripting&*EntryName*] | Specifies the location (path and file name) of the script (.scp) file, where *EntryName* is the name of the dial-up entry. | CMAK |
| [Server&*EntryName*] | Specifies the configuration where *EntryName* is the name of the dial-up entry. | CMAK |
| [Server&*TunnelDUN*] | Specifies the configuration to be used for all VPN connections. In this section title, *TunnelDUN* is the value of the TunnelDUN key in the [Connection Manager] section. This section uses the same format as the [Server&*EntryName*] section, including all keys valid under that section. | CMAK .cms |
| [TCP/IP&*EntryName*] | Specifies the configuration where *EntryName* is the name of the dial-up entry. | CMAK .cms |

| | | |
|---|---|---|
| [TCP/IP&*TunnelDUN*] | Specifies the TCP/IP configuration to be used for all VPN connections. In this section title, *TunnelDUN* is the value of the TunnelDUN key in the [Connection Manager] section. This section uses the same format as the [TCP/IP&*EntryName*] section, including all keys valid under that section. This section is created by the CMAK wizard, but the keys can be edited in the .cms file. | CMAK .cms |

**Keys that Control the Connection Manager User Interface**

**These keys control the appearance of the Connection Manager service profile.**

| Key | Description and values | Set in | Section |
|---|---|---|---|
| # | Controls what programs start when the custom button is clicked. The format of the value is # = *Program*, where # is an ID number from 0 to *n* and *Program* is the complete name of the program, including the extension and the path to the program. Programs start in sequence: 0 first, 1 next, and so on. Custom button programs cannot be DLLs. | CMAK .cms | [CustomButton Actions] |
| AnimatedLogo | Specifies whether an animation is used instead of a static bitmap in the Connection Manager logon dialog box.<br><br>0 (default) = Do not use animated graphic.<br><br>1 = Use animated graphic.<br><br>The [Animated Logo] section specifies the location of the animation (.spl or .swf) file. This function is provided for compatibility with Connection Manager 1.0, which supports Macromedia Flash animation files. | CMAK .cms | [Connection Manager] |
| Authenticating | Specifies the starting frame number of the animation that plays during user authentication. | .cms | [Animated Actions] |
| Connected | Specifies the starting frame number of the animation that plays when a connection is established. | .cms | [Animated Actions] |
| CustomButtonText | Specifies the text that appears on a custom button on the Connection Manager logon dialog box. Custom buttons do not appear or function when a user uses a service profile to log on to Windows. | .cms | [Connection Manager] |
| CustomButtonToolTip | Specifies the text that appears on a ToolTip for a custom button on the Connection Manager logon dialog box. Custom buttons do not appear or function when a user uses a service profile to log on to Windows. | .cms | [Connection Manager] |
| DialAutoMessage | Specifies the text that appears when the **Dial Automatically** check box has been selected. | .cms | [Connection Manager] |
| Dialing0 | Specifies the starting frame number of the animation that plays while the primary phone number is dialed. | .cms | [Animated Actions] |
| Dialing1 | Specifies the starting frame number of the animation that plays while the backup phone number is dialed. | .cms | [Animated Actions] |
| DisableICS | Specifies whether to prevent the user from sharing the connection with Internet Connection Sharing (ICS).<br><br>0 (default) = Do not suppress ICS settings.<br><br>1 = Disable ICS for this connection. The check box for enabling ICS will be unavailable.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Connection Manager] |

| | | | |
|---|---|---|---|
| Error | Specifies the starting frame number of the animation that plays at any connection error. | .cms | [Animated Actions] |
| HelpFile | Specifies the location (including name) of the custom Help file.<br><br>Default Help file = The default value for this key is blank.<br><br>When no value is specified, the standard Help file for service profiles is used. | CMAK | [Connection Manager] |
| HideAdvancedTab | Specifies whether to suppress the display of the **Advanced** tab in the properties dialog box for the service profile.<br><br>0 (default) = Do not suppress display of the **Advanced** tab.<br><br>1 = Suppress display of the **Advanced** tab.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Connection Manager] |
| HideBalloonTips | Specifies whether to suppress display of balloon-tip help for this connection.<br><br>0 (default) = Do not suppress balloon-tip help.<br><br>1 = Suppress balloon-tip help. | .cms | [Connection Manager] |
| HideDelay | Specifies the time (in seconds) to delay before displaying the **Phone Book Updates** dialog box (with progress indicators and message).<br><br>Default = Do not display the dialog box. | .cms | [Connection Manager] |
| HideDialAutomatically | Specifies whether the **Connect automatically** check box appears in the Connection Manager logon dialog box.<br><br>0 (default) = Display the check box.<br><br>1 = Do not display the check box.<br><br>If HideRememberPassword is set to 1, that key value overrides the HideDialAutomatically setting. | .cms | [Connection Manager] |
| HideDomain | Specifies whether the **Logon domain** box appears in the Connection Manager logon dialog box.<br><br>0 = Display the box.<br><br>1 = Do not display the box.<br><br>The default is 1. If the profile supports VPN connections and you have not previously set a value, CMAK sets this value to 0. | .cms | [Connection Manager] |
| HideInternetPassword | Specifies whether the **Password** box appears on the properties dialog box, on the **Internet Logon** tab.<br><br>0 (default) = Display the box.<br><br>1 = Do not display the box.<br><br>If HideInternetPassword and HideInternetUsername are both set to 1, the **Internet Logon** tab does not appear. | .cms | [Connection Manager] |

| | | | |
|---|---|---|---|
| HideInternetUsername | Specifies whether the **User name** box appears on the **Internet Logon** tab.<br><br>0 (default) = Display the box.<br><br>1 = Do not display the box.<br><br>If HideInternetPassword and HideInternetUsername are both set to 1, the **Internet Logon** tab does not appear. | .cms | [Connection Manager] |
| HidePassword | Specifies whether the **Password** box appears in the Connection Manager logon dialog box when dial-up connections are used.<br><br>0 (default) = Display the box.<br><br>1 = Do not display the box. | .cms | [Connection Manager] |
| HideRememberInternetPassword | Specifies whether the **Save Password** check box appears on the **Internet Logon** tab. If this key is not specified, the value of HideRememberPassword is used.<br><br>0 = Display check box.<br><br>1 = Do not display check box. | .cms | [Connection Manager] |
| HideRememberPassword | Specifies whether the **Save Password** check box appears on the Connection Manager logon dialog box and on the **Internet Logon** tab.<br><br>0 (default) = Display the check box.<br><br>1 = Do not display the check box.<br><br>If HideRememberPassword is set to 1, it overrides the HideDialAutomatically setting. | .cms | [Connection Manager] |
| HideTrayIcon | Specifies whether to display notification area shortcuts.<br><br>NULL = If TrayIcon is set, display the notification area shortcut specified. Otherwise, display the connection-status notification area shortcut (“blinking lights”). The NULL value is available for Microsoft Windows 2000, Windows XP, and Windows Server 2003 only.<br><br>0 = If TrayIcon is set, display the notification area shortcut specified and the connection-status notification area shortcut (“blinking lights”). Otherwise, display the Connection Manager notification area shortcut and the connection-status notification area shortcut (“blinking lights”).<br><br>1 = Do not display any notification area shortcuts.<br><br>If you choose not to display any notification area shortcuts, edit the Help file to remove references to them. | .cms | [Connection Manager] |
| HideUserName | Specifies whether the **User name** box appears in the Connection Manager logon dialog box.<br><br>0 (default) = Display the box.<br><br>1 = Do not display the box. | .cms | [Connection Manager] |

| | | | |
|---|---|---|---|
| Icon | Specifies the location and name of the icon (.ico) file for the 32 32 pixel icon, which is used, by default, for the service profile.<br><br>Default = The default value for this key is blank. When no value is specified, the standard program icon for Connection Manager is used. | CMAK | [Connection Manager] |
| Initial | Specifies the starting frame number of the animation that plays when Connection Manager starts. | .cms | [Animated Actions] |
| Logo | Specifies the location and name of the .bmp file that appears in the Connection Manager logon dialog box.<br><br>Default = The default value for this key is blank. When no value is specified, the standard logon bitmap for Connection Manager is used. | CMAK | [Connection Manager] |
| MaxDomain | Specifies the maximum length of the domain name that the user can type.<br><br>Default = 15 characters. | .cms | [Connection Manager] |
| MaxPassword | Specifies the maximum length of the password that the user can type.<br><br>Default = 127 bytes. | .cms | [Connection Manager] |
| MaxPhoneNumber | Specifies the maximum length of the phone number fields. The default for Microsoft Windows 95, Windows 98, and Windows Millennium Edition is 36; the default for Microsoft Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003 is 80. | .cms | [Connection Manager] |
| MaxUserName | Specifies the maximum length of the user name that the user can type.<br><br>Default = 256 bytes. | .cms | [Connection Manager] |
| *MenuItemName* | Specifies the command to be run from the shortcut menu. Shortcut menu commands use the format *MenuItemName = Program*, where *MenuItemName* is the descriptive name that appears on the shortcut menu and *Program* is the name of the executable file (.exe or other shell executable), including any parameters that the program requires. For example, if you wanted to add Notepad to the menu with the command name "Editor," the line would be Editor = Notepad.exe.<br><br>Default commands = Status and Disconnect.<br><br>These programs are available in a shortcut menu that appears when a user right-clicks the notification area shortcut. Applications are listed in the menu in the same sequence as they are listed in the .cms file. | CMAK | [Menu Options] |
| Movie | Specifies the location of the animation file that appears in place of the static bitmap in the Connection Manager logon dialog box. Settings in the [Animated Actions] section control how the animation plays based on the connection state. | .cms | [Animated Logo] |

| | | | |
|---|---|---|---|
| NoDialingRules | Specifies whether the user can use dialing rules on access numbers chosen from a phone book. Dialing rules check boxes are unavailable when the value of this key is set to 1.<br><br>0 (default) = Allow dialing rules.<br><br>1 = Do not allow dialing rules. | .cms | [Connection Manager] |
| NoPromptReconnect | Specifies whether the user is prompted for reconnection when the connection is inadvertently dropped.<br><br>0 (default) = Prompt the user.<br><br>1 = Do not prompt the user. | .cms | [Connection Manager] |
| *Parameter* | Specifies other animation parameters passed from software used to create the animation. An example of such a parameter is Quality = 3. | .cms | [Animated Logo] |
| Pausing | Specifies the starting frame number of the animation that plays during pauses between dials. | .cms | [Animated Actions] |
| PBLogo | Specifies the location and name of the .bmp file for the **Phone Book** dialog box.<br><br>Default = The default value for this key is blank. When nothing is specified, the standard phone book bitmap is used. | CMAK | [Connection Manager] |
| PBMessage | Specifies the text that appears in the **Phone Book** dialog box next to the **More access numbers** box when alternate access numbers are available for the selected geographic region.<br><br>Maximum = Approximately 100 characters. | CMAK | [Connection Manager] |
| PBUpdateMessage | Specifies the text that appears in the **Phone Book** dialog box during a phone book update.<br><br>Default = "Downloading new phone book." | .cms | [Connection Manager] |
| ResetPassword | Specifies whether the **New** button appears in the Connection Manager logon dialog box to enable a user to change the password. You can use this button to set up a custom password option. Type the location (including the name) of the executable (.exe) program that you use to support this process. You must provide an executable program for this option and include it as an additional file when you run the CMAK wizard. | .cms | [Connection Manager] |
| ServiceMessage | Specifies the service message that appears in the Connection Manager logon dialog box. This text is usually a customer-support number for the service.<br><br>Maximum = Approximately 50 characters. | CMAK | [Connection Manager] |
| ServiceName | Specifies the primary identifier for the profile. This text appears in the Connection Manager logon dialog box as the service name and as the name of the connection.<br><br>Maximum = Approximately 35 characters. Do not use any of the following in this name:<br><br>* \ / : ? < > \| [ ] | CMAK | [Connection Manager] |

| | | | |
|---|---|---|---|
| SmallIcon | Specifies the location and name of the icon (.ico) file for the 16 16 pixel icon, which is used, by default, for the title bar.<br><br>Default = The default value for this key is blank. When nothing is specified, the standard icon for the Connection Manager title bar is used. | CMAK | [Connection Manager] |
| TrayIcon | Specifies the location and name of the icon (.ico) file for the notification area shortcut.<br><br>Default = The default value for this key is blank. When nothing is specified, the notification area shortcut for Connection Manager is used. | CMAK | [Connection Manager] |
| Tunneling | Specifies the starting frame number of the animation that plays while a VPN connection is being established. | .cms | [Animated Actions] |
| UserNameOptional | Specifies whether the user must type a user name in the Connection Manager logon and **Internet logon** dialog boxes.<br><br>0 (default) = User name required.<br><br>1 = User name not required. | .cms | [Connection Manager] |

**Keys that Control Dial-up Settings**

**These keys control how the Connection Manager service profile makes a dial-up connection.**

| Key | Description and values | Set in | Section |
|---|---|---|---|
| CMSFile&*MergedProfile* | Specifies the location of a .cms file for a component service profile. The format is CMSFile&*MergedProfile*=*Location*, where *MergedProfile* is the name of the top-level service profile, and *Location* is the location of the .cms file for the component service profile. | CMAK | [ISP] |
| ConnectionType | Specifies the connection method to be used.<br><br>0 (default) = Connect by using a dial-up connection.<br><br>1 = Connect by using a direct connection (for example, DSL).<br><br>This key is valid only if both the Dialup key and the Direct key are set to 1. The user can change this option on the **General** tab in the properties dialog box of the service profile. | .cmp | [Connection Manager] |
| Custom_Security | Applies specific security settings on computers that are running Windows 2000, Windows XP, or Windows Server 2003.<br><br>0 = Use standard configuration.<br><br>1 (default) = Use custom configuration.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |

| | | | |
|---|---|---|---|
| CustomAuthData# | Specifies the configuration data for the Extensible Authentication Protocol (EAP) type. The format is CustomAuthData#=EAPConfigurationData, where # is the number of the specific line of data, for example:<br><br>CustomAuthData0=1AFGKLMOFGH23K<br><br>Each line can contain a maximum of 128 characters. If the configuration data has more than 128 characters, it must be specified as multiple lines, incrementing the number for each line of data. For example:<br><br>CustomAuthData0=First128Characters<br><br>CustomAuthData1=Next128Characters<br><br>CustomAuthData2=FinalCharacters<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |
| CustomAuthKey | Specifies the ID of the EAP type to be used with this connection. The format is CustAuthKey=*ID*, where *ID* is the Internet Assigned Numbers Authority (IANA) ID for the EAP type. If no ID is specified, no attempt will be made to read the CustomAuthData value.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |
| DataEncrypt | Specifies whether to force data encryption. This key is ignored for service profiles in which the EncryptionType key is set and that are installed on Windows 2000, Windows XP, or Windows Server 2003.<br><br>0 (default for dial-up entries) = Do not encrypt data.<br><br>1 (default for VPN entries) = Encrypt data. | CMAK | [Server&*EntryName*] |
| Dialup | Specifies whether dial-up connections are supported. If this key is set to 0, the **General** tab in the properties dialog box of the service profile will not appear.<br><br>0 = No.<br><br>1 (default) = Yes. | .cms | [Connection Manager] |
| Disable_LCP | Specifies whether to enable Link Control Protocol (LCP) extensions.<br><br>0 (default) = Enable.<br><br>1 = Disable. | .cms | [Server&*EntryName*] |
| DisableNbtOverIP | Specifies whether to disable the NetBIOS over TCP/IP (NetBT) routing protocol for this connection.<br><br>0 (default) = Disable NetBT.<br><br>1 = Leave NetBT enabled. | .cms | [Server&*EntryName*] |
| DNS_Address | Specifies the DNS address to use for the dial-up entry. | CMAK | [TCP/IP&*EntryName*] |
| DNS_Alt_Address | Specifies the alternate DNS address to use for the dial-up entry. | CMAK | [TCP/IP&*EntryName*] |
| DNSSuffix | Specifies the DNS suffix to use for the dial-up entry. Valid only for Windows XP and Windows Server 2003. | .cms | [TCP/IP&*EntryName*] |

| | | | |
|---|---|---|---|
| DontNegotiateMultilink | Specifies whether to allow multilinking.<br><br>0 (default) = Allow multilink.<br><br>1 = Do not allow multilink.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Server&*EntryName*] |
| DontUseRasCredentials | Specifies whether to cache the credentials of the user in Stored User Names and Passwords for access to resources after the user has logged on to the network.<br><br>0 (default) = Cache and use credentials.<br><br>1 = Do not cache and use credentials.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Server&*EntryName*] |
| DownloadDelay | Specifies the time (in seconds) to delay before downloading phone book updates. The delay time starts when the connection is established.<br><br>Default = 15 seconds.<br><br>This entry is valid only when, in the CMAK wizard, in the **Phone Book** dialog box, the **Automatically download phone book updates** check box is selected. | .cms | [Connection Manager] |
| DUN | Specifies the default network entry used when the **Dial-up Networking Entry** box for a POP is blank.<br><br>Default = *ServiceProfileServiceName*. | .cms | [Connection Manager] |
| EncryptionType | Specifies which encryption level to use.<br><br>0 = None.<br><br>1 (default for VPN) = Require.<br><br>3 (default for dial-up) = Optional.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |
| EnforceCustomSecurity | Specifies whether to require the user to choose a phone book entry that supports advanced security settings.<br><br>0 (default) = Do not enforce phone book entry selection (Valid only for Windows 95, Windows 98, Windows NT 4.0, and Windows Millennium Edition).<br><br>1 = Enforce selection of phone book entries that support advanced security settings. | CMAK | [Server&*EntryName*] |
| FilterA&*MergedProfile* | Specifies the name of the filter to apply to the phone book of the component profile. The format is FilterA&*MergedProfile* = *FilterName*. *MergedProfile* is the name of the component service profile, as specified in the References key in the .cms file, and *FilterName* is the name of the filter to apply.<br><br>The phone numbers that meet the filter criteria appear in the **Phone Book** dialog box, in the **Access numbers** list. | CMAK | [ISP] |

| | | | |
|---|---|---|---|
| FilterA&*ServiceProfile* | Specifies the name of the filter to apply to the access numbers of the top-level profile.<br><br>The selected phone numbers appear in Connection Manager, in the **Phone Book** dialog box, in the **Access numbers** list.<br><br>Default = NoSurchargeSignon | .cms | [ISP] |
| FilterB&*MergedProfile* | Specifies the name of the filter to apply to the phone book of the component profile. The format is FilterB&*MergedProfile* = *FilterName*. *MergedProfile* is the name of the component service profile, as specified in the References key in the .cms file, and *FilterName* is the name of the filter to apply.<br><br>The phone numbers that match the filter criteria appear in Connection Manager, in the **Phone Book** dialog box, in the **More access numbers** list. | CMAK | [ISP] |
| FilterB&*ServiceProfile* | Specifies the name of the filter to apply to the phone book of the top-level profile, for example:<br><br>FilterB& = SurchargeSignon<br><br>Phone numbers that match the filter criteria appear in Connection Manager, in the **Phone Book** dialog box, in the **More access numbers** list.<br><br>Default = SurchargeSignon. | .cms | [ISP] |
| Gateway_On_Remote | Specifies whether to use the default gateway on the remote private network.<br><br>0 = Do not use.<br><br>1 (default) = Use. | CMAK | [TCP/IP&*EntryName*] |
| IdleTimeout | Specifies how long (in minutes) that the connection must be idle before it is disconnected the first time that the service profile is used to connect.<br><br>Default = 10 minutes.<br><br>If this key is set to 0, idle connections are not disconnected. | .cmp | [Connection Manager] |
| IP_Address | Specifies the static IP addresses for the dial-up entry. | .cms | [TCP/IP&*EntryName*] |
| IP_Header_Compress | Specifies whether to enable IP header compression.<br><br>0 = Not enabled.<br><br>1 (default) = Enabled. | CMAK | [TCP/IP&*EntryName*] |

| | | | |
|---|---|---|---|
| IsdnDialMode | Specifies whether Connection Manager binds two channels on the same ISDN device.<br><br>0 (default) = Do not bind (use single channel only).<br><br>1 = Always bind the two channels (connect both channels automatically). If a dual channel connection cannot be made, the connection attempt will fail. This value is valid only for Windows NT 4.0, Windows 2000, Windows Millennium Edition, Windows XP, and Windows Server 2003.<br><br>2 = Attempt to bind the two channels (connect with both channels). If the two channels cannot be bound, use a single channel. This value is valid only for Windows 2000, Windows XP, and Windows Server 2003. | .cms | [Connection Manager] |
| Mask&*FilterName* | Specifies the mask value for *FilterName*. | .cms | [ISP] |
| Match&*FilterName* | Specifies the match value for *FilterName*. | .cms | [ISP] |
| Name | Specifies the location (including path and file name) of the script (.scp) file for the dial-up entry. | CMAK | [Scripting&*EntryName*] |
| Negotiate_IPX | Specifies whether to negotiate the Internetwork Packet Exchange (IPX) protocol.<br><br>0 (default) = Do not negotiate.<br><br>1 = Negotiate.<br><br>The IPX/SPX protocol is not available on Windows XP 64-Bit Edition and the 64-bit versions of Windows Server 2003. | .cms | [Server&*EntryName*] |
| Negotiate_NetBEUI | Specifies whether to negotiate the NetBIOS extended user interface (NetBEUI) protocol.<br><br>0 (default) = Do not negotiate.<br><br>1 = Negotiate | .cms | [Server&*EntryName*] |
| Negotiate_TCP/IP | Specifies whether to negotiate TCP/IP.<br><br>0 = Do not negotiate.<br><br>1 (default) = Negotiate. | .cms | [Server&*EntryName*] |
| NetworkLogon | Specifies whether to log on to the network.<br><br>0 = Do not log on. This is the default for dial-up connections.<br><br>1 = Log on. This is the default for VPN connections. | CMAK | [Server&*EntryName*] |
| PBFile | Specifies the location of the phone book (.pbk) file that the service profile uses. | CMAK | [ISP] |
| PBURL | Specifies the Uniform Resource Locator (URL) for the Phone Book Service (PBS) server. | CMAK | [ISP] |
| PW_Encrypt | Specifies whether to encrypt a password for authentication.<br><br>0 (default) = Do not encrypt.<br><br>1 = Encrypt. | CMAK | [Server&*EntryName*] |

| | | | |
|---|---|---|---|
| PW_EncryptMS | Specifies whether to require MS-CHAP version 1 for this connection.<br><br>0 = Do not require MS-CHAP v1.<br><br>1 = Require MS-CHAP v1. | CMAK | [Server&*EntryName*] |
| RedialCount | Specifies the number of redials to attempt the first time that the service profile is used to try to connect.<br><br>Default = 3. | .cmp | [Connection Manager] |
| RedialDelay | Specifies the delay (in seconds) between dial attempts. The delay is the elapsed time between the point of failure of the first call and the dialing of the first digit of the next phone number.<br><br>Default = 5 seconds. | .cms | [Connection Manager] |
| References | Specifies the service profiles merged through the CMAK wizard. Profile names are separated by commas, tabs, or spaces. Each name is used as a reference to that profile. | CMAK | [ISP] |
| RegionFile | Specifies the location of the region (.pbr) file that the service profile uses. | CMAK | [ISP] |
| Require_CHAP | Specifies whether Challenge Handshake Authentication Protocol (CHAP) is allowed.<br><br>0 = Not allowed.<br><br>1 (default) = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |
| Require_EAP | Specifies whether Extensible Authentication Protocol (EAP) is allowed.<br><br>0 (default) = Not allowed.<br><br>1 = allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |
| Require_MSCHAP | Specifies whether MSCHAP is allowed.<br><br>0 = Not allowed.<br><br>1 (default) = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |
| Require_MSCHAP2 | Specifies whether MSCHAP2 is allowed.<br><br>0 = Not allowed.<br><br>1 (default) = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |

| | | | |
|---|---|---|---|
| Require_PAP | Specifies whether Password Authentication Protocol (PAP) is allowed.<br><br>0 (default) = Not allowed.<br><br>1 = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |
| Require_SPAP | Specifies whether SPAP is allowed.<br><br>0 (default) = Not allowed.<br><br>1 = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*EntryName*] |
| Require_W95MSCHAP | Specifies whether W95MSCHAP is allowed. If the value for this key is set to 1, the value for Require_MS-CHAP must also be set to 1.<br><br>0 (default) = Not allowed.<br><br>1 = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | .cms | [Server&*EntryName*] |
| SecureClientForMSNet | Specifies whether to allow the use of NetBT. If this key is set, the value of the SecureLocalFiles key is ignored on Windows XP and Windows Server 2003.<br><br>0 (default) = Allow the use of NetBT.<br><br>1 = Do not allow the use of NetBT.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [TCP/IP&*EntryName*] |
| SecureFileAndPrint | Specifies whether to enable file and print sharing. If this key is set, the value of the SecureLocalFiles key is ignored on Windows XP and Windows Server 2003.<br><br>0 (default) = Enable.<br><br>1 = Disable. Files and print resources will not be shared.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Server&*EntryName*] |
| SecureLocalFiles | Specifies whether to share local files and allow printer sharing. If the value of the key is set to 1, it will disable NetBT on Windows 2000, Windows XP, and Windows Server 2003. This key is ignored on Windows XP and Windows Server 2003 if a value is assigned to the SecureClientForMSNet key, the SecureFileAndPrint key, or both.<br><br>0 (default) = Share local files.<br><br>1 = Do not share local files.<br><br>Valid only for Windows NT, Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [TCP/IP&*EntryName*] |

| | | | |
|---|---|---|---|
| *ServiceType* | Specifies a service type as a filter name to be applied to the phone book of the component profile. The format of the entry is *ServiceType=FilterName*, where *ServiceType* is the name as it appears in the **Phone Book** dialog box in the **Service type** list, and *FilterName* is the name of the filter to be associated with the entry. For example:<br><br>ISDN Multicast = MulticastISDN<br><br>A service type name should include no more than 25 characters. You can concatenate service type names by inserting a space between each name. At least two service types must be defined in order for phone numbers to appear in the properties dialog box for the service profile. | .cms | [Service Type] |
| Specify_IP_Address | Specifies whether to specify a static IP address for the dial-up entry.<br><br>0 (default) = Do not specify.<br><br>1 = Specify. | .cms | [TCP/IP&*EntryName*] |
| Specify_Server_Address | Specifies whether to specify static DNS or Windows Internet Name Service (WINS) addresses for the dial-up entry.<br><br>0 (default) = Do not specify.<br><br>1 = Specify. | CMAK | [TCP/IP&*EntryName*] |
| SW_Compress | Specifies whether to negotiate PPP software compression of data.<br><br>0 = Do not negotiate.<br><br>1 (default) = Negotiate. | .cms | [Server&*EntryName*] |
| TcpWindowSize | Specifies the size (in bytes) of the window offered by the connection. The window is the number of bytes of information to send before expecting a return packet. Minimum value is 4,096; maximum value is 65,500. If this key is not set, the default window size for the operating system and the connection type is used.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | .cms | [TCP/IP&*EntryName*] |
| UseTerminalWindow | Specifies whether to use a terminal window when connecting to a service.<br><br>0 (default) = Do not use a terminal window.<br><br>1 = Use a terminal window.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | .cms | [Scripting&*EntryName*] |
| Version | Specifies the version of the phone book. If you provide phone books, the PBS server uses this value to determine whether and how to update phone books. If the version listed in this key is more than five versions out of date, the PBS server downloads the complete phone book. Otherwise, the PBS server downloads a delta file. | .cms | [Connection Manager] |
| WINS_Address | Specifies the WINS address for the dial-up entry. | CMAK | [TCP/IP&*EntryName*] |
| WINS_Alt_Address | Specifies the alternate WINS address for the dial-up entry. | CMAK | [TCP/IP&*EntryName*] |

**Keys that control VPN settings**

**These keys control how the Connection Manager service profile makes a VPN connection.**

| Key | Description and values | Set in | Section |
|---|---|---|---|

| | | | |
|---|---|---|---|
| ConnectionType | Specifies the connection method to be used.<br><br>0 (default) = Connect by using a dial-up connection.<br><br>1 = Connect by using a direct connection (for example, DSL).<br><br>This key is valid only if both the Dialup key and the Direct key values are set to 1. The user can change this option in the properties dialog box for the service profile on the **General** tab. | .cmp | [Connection Manager] |
| Custom_Security | Specifies whether specific security settings should be applied on computers that are running Windows 2000, Windows XP, or Windows Server 2003.<br><br>0 = Use standard configuration.<br><br>1 (default) = Use custom configuration.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| CustomAuthData# | Specifies the configuration data for the EAP type. The format is CustomAuthData#=EAPConfigurationData, where # is the number of the specific line of data, for example:<br><br>CustomAuthData0=1AFGKLMOFGH23K<br><br>Each line can contain a maximum of 128 characters. If the ID has more than 128 characters, it must be specified as multiple lines, incrementing the # value for each line of data. For example:<br><br>CustomAuthData0=First128Characters<br><br>CustomAuthData1=Next128Characters<br><br>CustomAuthData2=FinalCharacters<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| CustomAuthKey | Specifies the ID of the EAP type to use with this connection. The format is CustAuthKey=*ID*, where *ID* is the Internet Assigned Numbers Authority (IANA) ID for the EAP type. If no ID is specified, no attempt will be made to read the CustomAuthData value.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| DataEncrypt | Specifies whether to force data encryption. This key is ignored for service profiles in which the EncryptionType key is set and that are installed on Windows 2000, Windows XP, or Windows Server 2003.<br><br>0 (default for dial-up entries) = Do not encrypt data.<br><br>1 (default for VPN entries) = Encrypt data. | CMAK | [Server&*TunnelDUN*] |
| Direct | Specifies whether direct connections are supported.<br><br>0 = No.<br><br>1 (default) = Yes.<br><br>Valid only for VPN service profiles. | .cms | [Connection Manager] |

| | | | |
|---|---|---|---|
| Disable_LCP | Specifies whether to enable Link Control Protocol (LCP) extensions.<br><br>0 (default) = Enable.<br><br>1 = Disable. | .cms | [Server&*TunnelDUN*] |
| DisableNbtOverIP | Specifies whether to disable the NetBT routing protocol for this connection.<br><br>0 (default) = Disable NetBT.<br><br>1 = Leave NetBT enabled. | .cms | [Server&*TunnelDUN*] |
| DNS_Address | Specifies the DNS address for the VPN entry. | CMAK | [TCP/IP&*TunnelDUN*] |
| DNS_Alt_Address | Specifies the alternate DNS address for the VPN entry. | CMAK | [TCP/IP&*TunnelDUN*] |
| DNSSuffix | Specifies the DNS suffix to use for the connection. Valid only for Windows XP and Windows Server 2003. | .cms | [TCP/IP&*TunnelDUN*] |
| DontNegotiateMultilink | Specifies whether to allow multilinking.<br><br>0 (default) = Allow multilinking.<br><br>1 = Do not allow multilinking.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Server&*TunnelDUN*] |
| DontUseRasCredentials | Specifies whether to cache the user credentials in Stored User Names and Passwords for access to network resources after the user has logged on to the network.<br><br>0 (default) = Cache and use credentials.<br><br>1 = Do not cache and use credentials.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Server&*TunnelDUN*] |
| EncryptionType | Specifies which encryption level to use.<br><br>0 = None.<br><br>1 (default for VPN) = Require.<br><br>3 (default for dial-up) = Optional.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| EnforceCustomSecurity | Specifies whether to require the user to select a phone book entry that supports advanced security settings.<br><br>0 (default) = Do not enforce phone book entry selection. This value is valid for Windows 95, Windows 98, Windows NT 4.0, and Windows Millennium Edition only.<br><br>1 = Enforce selection of phone book entries that support advanced security settings. | CMAK | [Server&*TunnelDUN*] |
| Gateway_On_Remote | Specifies whether to use the default gateway on the remote private network.<br><br>0 = Do not use.<br><br>1 (default) = Use. | CMAK | [TCP/IP&*TunnelDUN*] |
| IP_Address | Specifies the static IP addresses for the VPN entry. | .cms | [TCP/IP&*TunnelDUN*] |

| | | | |
|---|---|---|---|
| IP_Header_Compress | Specifies whether to enable IP header compression.<br><br>0 = Not enabled.<br><br>1 (default) = Enabled. | CMAK | [TCP/IP&*TunnelDUN*] |
| Negotiate_IPX | Specifies whether to negotiate the IPX protocol.<br><br>0 (default) = Do not negotiate.<br><br>1 = Negotiate.<br><br>The IPX/SPX protocol is not available on Windows XP 64-Bit Edition and the 64-bit versions of Windows Server 2003. | .cms | [Server&*TunnelDUN*] |
| Negotiate_NetBEUI | Specifies whether to negotiate the NetBIOS extended user interface (NetBEUI) protocol.<br><br>0 (default) = Do not negotiate.<br><br>1 = Negotiate | .cms | [Server&*TunnelDUN*] |
| Negotiate_TCP/IP | Specifies whether to negotiate TCP/IP.<br><br>0 = Do not negotiate.<br><br>1 (default) = Negotiate. | .cms | [Server&*TunnelDUN*] |
| NetworkLogon | Specifies whether to log on to the network.<br><br>0 = Do not log on. This is the default for dial-up connections.<br><br>1 = Log on. This is the default for VPN connections. | CMAK | [Server&*TunnelDUN*] |
| PW_Encrypt | Specifies whether to encrypt a password when authenticating.<br><br>0 (default) = Do not encrypt.<br><br>1 = Encrypt. | CMAK | [Server&*TunnelDUN*] |
| PW_EncryptMS | Specifies whether to require MS-CHAP version 1 for this connection.<br><br>0 = Do not require MS-CHAP v1.<br><br>1 = Require MS-CHAP v1. | CMAK | [Server&*TunnelDUN*] |
| Require_CHAP | Specifies whether Challenge Handshake Authentication Protocol (CHAP) is allowed.<br><br>0 = Not allowed.<br><br>1 (default) = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| Require_EAP | Specifies whether Extensible Authentication Protocol (EAP) is allowed.<br><br>0 (default) = Not allowed.<br><br>1 = allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |

| | | | |
|---|---|---|---|
| Require_MSCHAP | Specifies whether MSCHAP is allowed.<br><br>0 = Not allowed.<br><br>1 (default) = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| Require_MSCHAP2 | Specifies whether MSCHAP2 is allowed.<br><br>0 = Not allowed.<br><br>1 (default) = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| Require_PAP | Specifies whether Password Authentication Protocol (PAP) is allowed.<br><br>0 (default) = Not allowed.<br><br>1 = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| Require_SPAP | Specifies whether SPAP is allowed.<br><br>0 (default) = Not allowed.<br><br>1 = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| Require_W95MSCHAP | Specifies whether W95MSCHAP is allowed. If this key is set to 1, Require_MS-CHAP must also be set to 1.<br><br>0 (default) = Not allowed.<br><br>1 = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | .cms | [Server&*TunnelDUN*] |
| SecureClientForMSNet | Specifies whether to allow the use of NetBT. If this key is set, the value of the SecureLocalFiles key is ignored on Windows XP and Windows Server 2003.<br><br>0 (default) = Allow the use of NetBT.<br><br>1 = Do not allow the use of NetBT.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Server&*TunnelDUN*] |
| SecureFileAndPrint | Specifies whether to enable file and print sharing. If this key is set, the SecureLocalFiles key is ignored on Windows XP and Windows Server 2003.<br><br>0 (default) = Enable.<br><br>1 = Disable. Files and print resources will not be shared.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Server&*TunnelDUN*] |

| | | | |
|---|---|---|---|
| SecureLocalFiles | Specifies whether to share local files and allow printer sharing. If the value of the key is set to 1, it will disable NetBT on Windows 2000, Windows XP, and Windows Server 2003. This key value will be ignored on Windows XP and Windows Server 2003 if a value is assigned to the SecureClientForMSNet key, the SecureFileAndPrint key, or both.<br><br>0 (default) = Share local files.<br><br>1 = Do not share local files.<br><br>Valid only for Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Server&*TunnelDUN*] |
| Specify_IP_Address | Specifies whether to specify a static IP address.<br><br>0 (default) = Do not specify.<br><br>1 = Specify. | .cms | [TCP/IP&*TunnelDUN*] |
| Specify_Server_Address | Specifies whether to specify static DNS or WINS addresses.<br><br>0 (default) = Do not specify.<br><br>1 = Specify. | CMAK | [TCP/IP&*TunnelDUN*] |
| SW_Compress | Specifies whether to negotiate Point-to-Point Protocol (PPP) software compression of data.<br><br>0 = Do not negotiate.<br><br>1 (default) = Negotiate. | .cms | [Server&*TunnelDUN*] |
| TcpWindowSize | Specifies the size (in bytes) of the window offered by the connection. The window is the number of bytes of information to send before expecting a return packet. Minimum value is 4,096; maximum value is 65,500. If this key is not set, the default window size for the operating system and the connection type is used.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | .cms | [TCP/IP&*TunnelDUN*] |
| Tunnel | Specifies whether VPN connections are enabled for the phone numbers in the top-level phone book.<br><br>0 (default) = Not enabled.<br><br>1 = Enabled. | CMAK | [Connection Manager] |
| TunnelAddress | Specifies the IP address or fully qualified domain name for the VPN server. | CMAK | [Connection Manager] |
| TunnelDUN | Specifies the default networking entry used for VPN connections.<br><br>Default = *ServiceProfileServiceName* tunnel. | .cms | [Connection Manager] |
| TunnelFile | Specifies the relative path and file name to the VPN file. If you use macros in the relative path, make sure that the macros are recognized on all operating systems to which you want to distribute the profile. | CMAK | [Connection Manager] |

| | | | |
|---|---|---|---|
| TunnelReferences | Specifies whether VPN connections are enabled for phone numbers taken from the phone book of a component profile.<br><br>0 (default) = Do not use a VPN connection when using any of these phone numbers.<br><br>1 = Use a VPN connection when using a phone number from a component profile. | CMAK | [Connection Manager] |
| UsePreSharedKey | Specifies whether to use a preshared key for Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec) authentication.<br><br>0 (default) = Do not use a preshared key.<br><br>1 = Use a preshared key.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Networking&*TunnelDUN*] |
| VPNEntry | Specifies whether the entry is a VPN entry or a dial-up entry.<br><br>0 = Dial-up entry.<br><br>1 = VPN entry. | CMAK | [Networking&*TunnelDUN*] |
| VpnStrategy | Specifies which tunneling protocol to use when establishing a VPN connection.<br><br>1 = Use PPTP only.<br><br>2 = Try PPTP first.<br><br>3 = Use L2TP only.<br><br>4 = Try L2TP first.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | CMAK | [Networking&*TunnelDUN*] |
| WINS_Address | Specifies the WINS address for the VPN entry. | CMAK | [TCP/IP&*TunnelDUN*] |
| WINS_Alt_Address | Specifies the alternate WINS address for the VPN entry. | CMAK | [TCP/IP&*TunnelDUN*] |

**Keys that Control Required User Information**

**These keys control what information the service profile requires before the user can click . Connect**

| Key | Description and values | Set in | Section |
|---|---|---|---|
| DomainOptional | Specifies whether the user must specify a domain name when making a connection.<br><br>0 = Require a domain.<br><br>1 (default) = Do not require a domain. | .cms | [Connection Manager] |
| GlobalCredentials | Specifies whether to save the user name and password supplied for first-time use of the connection for all users of the connection.<br><br>0 (default) = Save the user name and password for a single user.<br><br>1 = Save the user name and password for all users of the connection.<br><br>Valid only for Windows XP and Windows Server 2003. | .cmp | [Connection Manager] |

| | | | |
|---|---|---|---|
| PasswordOptional | Specifies whether the user must specify a password in the Connection Manager logon and **Internet Logon** dialog boxes.<br><br>0 (default) = Password required.<br><br>1 = Password not required. | .cms | [Connection Manager] |
| UseSameUserName | Specifies whether to use a separate user name and password when connecting to the Internet service provider (for a VPN connection).<br><br>0 (default) = Use the user name and password in the properties dialog box on the **Internet Logon** tab. Service profiles built to support VPN connections contain this additional tab.<br><br>1 = Use the same user name and password when connecting to the Internet service provider as appears in the Connection Manager logon dialog box for logging on to the private service. | CMAK | [Connection Manager] |
| UseWinLogonCredentials | Specifies whether to use the user name, password, and domain that the user specified when the user logged on most recently to Windows.<br><br>0 = Do not use credentials.<br><br>1 (default) = Use credentials.<br><br>Valid only for Windows 2000, Windows XP, and Windows Server 2003. | .cms | [Connection Manager] |

**Keys that Provide Data the First Time the Service Profile is Used**

**These keys populate information in the Connection Manager user interface the first time that a service profile is used.**

| Key | Description and values | Set in | Section |
|---|---|---|---|
| Description0 | Describes the primary phone number that is provided the first time that the service profile is used. Appears in the **Phone Book** dialog box as a Point of Presence (POP) name. | .cmp | [Connection Manager] |
| Description1 | Describes the backup phone number that is provided the first time that the service profile is used. Appears in the **Phone Book** dialog box as a POP name. | .cmp | [Connection Manager] |
| DialAutomatically | Specifies whether to have the **Connect automatically** check box in the Connection Manager logon dialog box selected by default and to connect automatically the first time that the service profile is used.<br><br>0 (default) = Do not select the check box.<br><br>1 = Select the check box and connect automatically. | .cmp | [Connection Manager] |
| Domain | Specifies the logon domain name that appears automatically the first time that the service profile is used. | .cmp | [Connection Manager] |
| DUN0 | Specifies the dial-up entry used for the primary phone number the first time that the service profile is used. Value must exactly match the entry name specified in Phone Book Administrator. | .cmp | [Connection Manager] |
| DUN1 | Specifies the dial-up entry used for the backup phone number the first time that the service profile is used. Value must exactly match the entry name specified in Phone Book Administrator. | .cmp | [Connection Manager] |

| | | | |
|---|---|---|---|
| EnableLogging | Specifies whether a log file is created for the connection the first time that the service profile is used. The log file is named *ServiceProfileServiceName*.log for an all-user profile, and *ServiceProfileServiceName* (Single User).log for a single-user profile.<br><br>0 = Do not create a log file.<br><br>1 (default) = Create a log file. | .cmp | [Connection Manager] |
| InternetPassword | Specifies the password that appears automatically the first time that the user tries to connect using a VPN connection. | .cmp | [Connection Manager] |
| InternetUserName | Specifies the user name that appears automatically the first time that the user tries to connect using a VPN connection. | .cmp | [Connection Manager] |
| KeepDefaultCredentials | Specifies whether to keep first-time credentials such as user name and password in the .cmp file of the profile for persistent default display and usage for all first-time users of the profile. For single-user profiles, this information will appear only once.<br><br>0 (default) = Do not keep first-time credentials in the .cmp file.<br><br>1 = Keep first-time credentials in the .cmp file. | .cmp | [Connection Manager] |
| Password | Specifies the password that appears automatically the first time that the profile is used. | .cmp | [Connection Manager] |
| PCS | Specifies how the password will be stored for pre-population.<br><br>0 = Password stored in plaintext; pre-population is allowed.<br><br>No other value for this key is valid for pre-population. | .cmp | [Connection Manager] |
| Phone0 | Specifies the primary phone number populated for the user the first time the service profile is used and the default phone number any time another phone number is not selected. | .cmp | [Connection Manager] |
| Phone1 | Specifies the backup phone number populated for the user the first time the service profile is used. | .cmp | [Connection Manager] |
| PhoneCanonical0 | Specifies the canonical format of the primary phone number populated for the user the first time the service profile is used. | .cmp | [Connection Manager] |
| PhoneCanonical1 | Specifies the canonical format of the backup phone number populated for the user the first time the service profile is used. | .cmp | [Connection Manager] |
| PhoneCountry0 | Specifies the Telephony Application Programming Interface (TAPI) country/region code for the primary phone number set as default for the user the first time the service profile is used. | .cmp | [Connection Manager] |
| PhoneCountry1 | Specifies the Telephony Application Programming Interface (TAPI) country/region code for the backup phone number set as default for the user the first time the service profile is used. | .cmp | [Connection Manager] |
| PhoneSource0 | Specifies the service profile for the primary phone number to be used the first time the user connects. | .cmp | [Connection Manager] |
| PhoneSource1 | Specifies the service profile for the backup phone number to be used the first time the user connects. | .cmp | [Connection Manager] |
| Region0 | Specifies the default state or province from the phone book for the primary phone number the first time that the service profile is opened. | .cmp | [Connection Manager] |

| | | | |
|---|---|---|---|
| Region1 | Specifies the default state or province from the phone book for the backup phone number the first time that the service profile is opened. | .cmp | [Connection Manager] |
| RememberInternetPassword | Specifies whether the **Save password** check box on the **Internet Logon** tab is selected by default for a populated Internet password. This key is ignored if the InternetPassword key is not set in the .cmp file.<br><br>0 (default) = The check box is cleared by default.<br><br>1 = The check box is selected by default, and the password is saved. | .cmp | [Connection Manager] |
| RememberPassword | Specifies whether to have the **Save password** check box in the Connection Manager logon dialog box selected by default for a populated password. This key is ignored if the Password key is not set in the .cmp file.<br><br>0 (default) = The check box is cleared by default.<br><br>1 = The check box is selected by default, and the password is saved. | .cmp | [Connection Manager] |
| ServiceType0 | Specifies the default service type from the phone book for the primary phone number the first time that the service profile is opened. | .cmp | [Connection Manager] |
| ServiceType1 | Specifies the default service type from the phone book for the backup phone number the first time that the service profile is opened. | .cmp | [Connection Manager] |
| UseDialingRules0 | Specifies whether dialing rules are applied to the primary phone number the first time that the service profile is used.<br><br>0 = Dial the number exactly as it appears in the phone book or as it is typed.<br><br>1 (default) = Apply dialing rules to the number when displaying and dialing. | .cmp | [Connection Manager] |
| UseDialingRules1 | Specifies whether dialing rules are applied to the backup number the first time the service profile is used.<br><br>0 = Dial the number exactly as it appears in the phone book or as it is typed.<br><br>1 (default) = Apply dialing rules to the number when displaying and dialing. | .cmp | [Connection Manager] |
| UserName | Specifies the name used to authenticate the user the first time that the service profile is used. | .cmp | [Connection Manager] |
| UserNamePrefix | Specifies the prefix added to the user name. Enables access to a service that requires a realm name as a user name prefix. This key is not used in a VPN connection. | CMAK | [Connection Manager] |
| UserNameSuffix | Specifies the suffix appended to the user name. Enables access to a service that requires a realm name as a user name suffix. This key not used in a VPN connection. | CMAK | [Connection Manager] |

**Connection Manager Custom Action Keys**

**These keys control when and how a custom action is performed.**

| Key | Description and values | Set in | Section |
|---|---|---|---|

| | | | |
|---|---|---|---|
| # | Specifies a single custom action. The format of the value is # = *Program*, where # is an ID number from 0 to *n* in sequential order and *Program* is the complete name of the program, including the file name extension and the path to the program. Programs start in sequence: 0 first, 1 next, and so on. Macros, such as %UserName% and %ErrorCode%, are set in Connection Manager, and they can be used as parameters for custom action programs. | CMAK | [Pre-Init Actions]<br><br>[Pre-Connect Actions]<br><br>[Pre-Dial Actions]<br><br>[Pre-Tunnel Actions]<br><br>[Connect Actions]<br><br>[Disconnect Actions]<br><br>[On-Cancel Actions]<br><br>[On-Error Actions]<br><br>[Auto Applications] |
| #&Flags | Specifies whether a custom action runs when a specific type of connection is established. The format of the value is #&Flags = *n*, where # is the ID number of a defined custom action and n indicates how dial-up and direct connections for that custom action are handled:<br><br>If *n* = 0 (default), all connections trigger this custom action. (The default key does not appear in the .cms file.)<br><br>If *n* = 1, all direct tunnel connections trigger this custom action.<br><br>If *n* = 2 , all connections that involve dialing (direct or double) trigger this custom action.<br><br>If *n* = 4, only connections that use dial-up trigger this custom action.<br><br>If *n* = 8, only connections that involve a tunnel trigger this custom action | CMAK | [Pre-Init Actions]<br><br>[Pre-Connect Actions]<br><br>[Pre-Dial Actions]<br><br>[Pre-Tunnel Actions]<br><br>[Connect Actions]<br><br>[Disconnect Actions]<br><br>[On-Cancel Actions]<br><br>[On-Error Actions]<br><br>[Auto Applications] |

| | | | |
|---|---|---|---|
| #&Description | Specifies a short description of a custom action. The format of the value is #&Description = *Text*, where # is the number of a defined custom action and *Text* is the descriptive information for that custom action. This option can provide useful information when a synchronous custom action does not run properly.<br><br>Default = The name of the custom action program. | CMAK | [Pre-Init Actions]<br><br>[Pre-Connect Actions]<br><br>[Pre-Dial Actions]<br><br>[Pre-Tunnel Actions]<br><br>[Connect Actions]<br><br>[Disconnect Actions]<br><br>[On-Cancel Actions]<br><br>[On-Error Actions]<br><br>[Auto Applications] |

**Connection Manager Control Keys**

**Connection Manager uses these keys to locate files and determine behavior during a connection.**

| Key | Description and values | Set in | Section |
|---|---|---|---|
| AutoReconnect | Specifies whether to automatically attempt to reconnect when a connection is terminated unexpectedly.<br><br>0 (default) = Do not automatically attempt to reconnect.<br><br>1 = Attempt to reconnect. | .cms | [Connection Manager] |
| CheckOSComponents | Specifies whether Connection Manager checks the operating system components.<br><br>0 (default) = Check the operating system components only when an error occurs.<br><br>1 = Always check the operating system components on service profile initialization. | .cms | [Connection Manager] |
| CMSFile | Specifies the name of the service profile file: (*Path**ServiceProfileFileName*.cms), as specified in **File name** on the **Service and File Names** page. Do not use any of the following in this name: ! @ ^ & ( ) { } ` ´ ~ , ; <space> * = \ / : ? " < > \| . % + [ ] $ #. Do not alter this value in the .cmp file. If a different name is required, use the CMAK wizard to edit the existing profile, specify new names for both the service name and the file name, and then finish running the wizard. | CMAK | [Connection Manager] |
| DisableICS | Specifies whether to prevent the user from sharing the connection with Internet Connection Sharing (ICS).<br><br>0 (default) = Do not suppress ICS settings.<br><br>1 = Disable ICS for this connection. The check box for enabling ICS will be unavailable.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Connection Manager] |

| | | | |
|---|---|---|---|
| DoNotCheckBindings | Specifies whether to include the TCP/IP bindings when checking operating system components.<br><br>0 (default) = Include the TCP/IP bindings in the TCP/IP check and bind all that are not bound.<br><br>1 = Do not include the TCP/IP bindings when checking operating system components.<br><br>TCP/IP services that sit between the TCP/IP protocol and the adapter are protected from being checked only if this option is set to 1 in all installed service profiles. This option does not prevent other service profiles or programs from checking bindings. | .cms | [Connection Manager] |
| EnableICF | Specifies whether to enable Internet Connection Firewall (ICF) for this connection with default settings.<br><br>0 (default) = Do not adjust ICF settings.<br><br>1 = Enable ICF for the connection with default settings.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Connection Manager] |
| FileDirectory | Specifies the location of the log file. If the value is blank, the log file is created in the temp folder. | .cms | [Logging] |
| IdleThreshold | Specifies the maximum bytes per minute that can be received if an idle state is to be detected for the connection.<br><br>Default = 0.<br><br>This key is not valid on Windows NT 4.0. | .cms | [Connection Manager] |
| InternetConnection | Specifies whether to inform the operating system that this connection is a connection to the Internet.<br><br>0 (default) = Do not inform the operating system that this connection connects to the Internet for the purposes of automatically connecting when no connection is present.<br><br>1 = Inform the operating system that this connection connects to the Internet.<br><br>Valid only for Windows XP and Windows Server 2003. | .cms | [Connection Manager] |
| MaxFileSize | Specifies the maximum file size (in kilobytes) before the log is cleared. The file size is checked at initialization to prevent data loss during a connection. The default is 100. | .cms | [Logging] |
| MonitorCallingProgram | Specifies whether to handle Connection Manager as a monitored custom action when Connection Manager is called programmatically and the calling process terminates.<br><br>0 = Do not monitor.<br><br>1 (default) = Monitor. | .cms | [Connection Manager] |
| PasswordHandling | Specifies the method of handling the password during the logon process.<br><br>0 (default) = Normal (case sensitive).<br><br>1 = Change uppercase letters to lowercase.<br><br>2 = Change lowercase letters to uppercase. | .cms | [Connection Manager] |
| Version | Specifies the version of the service provider or service profile file. Do not modify this value. | .cms<br><br>.cmp | [Profile Format] |

**Connection Manager Registry Entries**
The following registry entries are associated with Connection Manager.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**\PBServerMonitor**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Services

**Version**
Windows Server 2003; Windows 2000 Server

The PBServerMonitor subkey is used to store the Performance subkey.

**\Performance**
**Registry path**
HKLM\SYSTEM\CurrentControlSet\Services\PBServerMonitor

**Version**
Windows Server 2003; Windows 2000 Server

The Performance subkey is used to map the locations of the data, library, and binary files that the PBServer performance counter uses.

**\Phone Book Service**
Registry path

**HKLM\SYSTEM\CurrentControlSet\Services\EventLog\Application**
**Version**
Windows Server 2003; Windows 2000 Server

The Phone Book Service subkey is used to map the location of the event file and to enumerate what types of databases are supported.

**\Connection Manager Administration Kit**
**Registry path**
HKLM\Software\Microsoft

**Version**
Windows Server 2003; Windows 2000 Server; Windows XP Professional

The Connection Manager Administration Kit subkey is used to map the location of the last service profile created with the wizard.

**\Connection Manager**
**Registry path**
\Software\Microsoft

**Version**
Windows Server 2003; Windows XP; Windows Millennium Edition; Windows 2000; Windows 98

This registry subkey is located under HKLM, HKCU, or both, depending on how a service profile is installed, and this subkey is used to store subkeys relating to specific service profiles.

**\Access Points**
**Registry path**
\Software\Microsoft\Connection Manager\UserInfo*ProfileName*

**Version**
Windows Server 2003; Windows XP; Windows Millennium Edition; Windows 2000; Windows 98

The Access Points subkey is used to store the *AccessPointName* subkeys in the HKCU context.

**\AccessPointName**
**Registry path**

\Software\Microsoft\Connection Manager\UserInfo*ProfileName*\AccessPoint

**Version**
Windows Server 2003; Windows XP; Windows Millennium Edition; Windows 2000; Windows 98

The *AccessPointName* subkey is used to store the Favorites connection information. The name of the registry key *AccessPointName* is the friendly name that the user gives it. The default friendly name is Home <default>.

**\ICSData**
**Registry path**
\Software\Microsoft\Connection Manager*ProfileName*

**Version**
Windows Server 2003; Windows XP

The ICSData subkey is used to enable Internet Connection Sharing (ICS) for a Connection Manager service profile. Only All-User connections can have this key. The data contained under this subkey duplicates the **UserInfo** subkey data, and this data is used when ICS is invoked and no user is logged on.

**\Mappings**
**Registry path**
\Software\Microsoft\Connection Manager

**Version**
Windows Server 2003; Windows XP; Windows Millennium Edition; Windows 2000; Windows 98

The Mappings subkey is used to store the full path of the .cmp file for each service profile. The format for the value of the registry key is:

Under HKLM: All-users profiles

Under HKCU: Single-user profiles

| Field | Description |
|---|---|
| Name | The friendly name of the service profile. |
| Type | REG_SZ |
| Data | The full path to the .cmp file of the service profile. |

**\ProfileName**
**Registry path**
\Software\Microsoft\Connection Manager

**Version**
Windows Server 2003; Windows XP; Windows Millennium Edition; Windows 2000; Windows 98

The *ProfileName* subkey is used to store system interaction subkeys, such as the WinLogon subkey for Windows logon custom action executables. The name of the registry key *ProfileName* will be the friendly name of the service profile, not the eight-character file name.

**\UserInfo**
**Registry path**
\Software\Microsoft\Connection Manager

**Version**
Windows Server 2003; Windows XP; Windows Millennium Edition; Windows 2000; Windows 98

The UserInfo subkey is used to store the *ProfileName* subkeys in the HKCU context.

**\WinLogon Actions**
**Registry path**
\Software\Microsoft\Connection Manager*ProfileName*

**Version**
Windows Server 2003; Windows XP; Windows Millennium Edition; Windows 2000; Windows 98

**The WinLogon Actions subkey is used to enable custom actions to run during the logon process. The format for the**

**value of the registry key is:**

| Field | Description |
|---|---|
| Name | The name of the executable file that will run. |
| Type | REG_DWORD |
| Data | A value that indicates the location of the executable. |

**Table Supported values for the Data field are:**

| Value | Location of executable |
|---|---|
| 0x00000000(0) | *%windir%*\system32 |
| 0x00000001(1) | Profile directory |

**Network Ports Used by Connection Manager**

The following table describes the network ports that Connection Manager commonly uses. Because administrators can customize Connection Manager service profiles, the exact network ports that service profiles use vary.

**Port Assignments for Connection Manager**

| Service Name | UDP | TCP |
|---|---|---|
| HTTP | | 80 |
| FTP | | 20/21 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Microsoft Platform SDK on MSDN for more information about remote access APIs
- VPN Technical Reference
- Resource Kit Tools
- Registry Reference
- Connection Manager Administration Kit
- Routing and Remote Access
- Microsoft Systems Architecture

# IAS Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS Technical Reference**
Internet Authentication Service (IAS) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy. IAS is included in Microsoft Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. You can configure IAS for two roles: RADIUS server and RADIUS proxy. IAS configured as a RADIUS server performs centralized connection authentication, authorization, and accounting for many types of network access, including wireless, authenticating switch, dial-up, virtual private network (VPN), and Windows Server 2003 or Microsoft Windows 2000 servers running Routing and Remote Access service. With IAS configured as a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers.

**In this subject**


- What Is IAS?
- How IAS Technology Works
- IAS Tools and Settings

# What Is IAS?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Components of an IAS infrastructure
- IAS Features
- New IAS Features Supported by Microsoft Windows Server 2003


Internet Authentication Service (IAS) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy. RADIUS is an industry standard for the client-server protocol. For more information about RADIUS, see RFC 2865, "Remote Authentication Dial-In User Service (RADIUS)" and RFC 2866, "RADIUS Accounting." For information about standards that apply to IAS in Windows Server 2003, see RFC 2868, "RADIUS Attributes for Tunnel Protocol Support" and RFC 2869, "RADIUS Extensions."

As a RADIUS server, IAS performs centralized connection authentication, authorization, accounting, and auditing (AAAA) for many types of network access including wireless, authenticating switch, and remote access dial-up and virtual private network (VPN) connections.

As a RADIUS proxy, IAS provides the routing of RADIUS messages between RADIUS clients (access servers), RADIUS proxies, and the RADIUS servers that perform AAAA for the connection attempt. When used as a RADIUS proxy, IAS is a central switching or routing point through which RADIUS access and accounting messages flow. Using IAS, you can centrally manage AAAA for organizations of all sizes, including small businesses, medium organizations, enterprise-level organizations, and Internet Service Providers (ISPs). IAS provides you with the ability to secure and manage network access across a variety of network access scenarios such as the following:

- Employees connecting to your organization network through dial-up, VPN, wireless, and wired connections, using a variety of devices, including organization computers, personal digital assistants, and non-domain member computers, such as employee-owned devices.
- Employees connecting to other networks, including the Internet and business partner networks.
- Business partners connecting to your organization network.

IAS enables the use of a heterogeneous or homogenous set of wireless, switch, remote access, or VPN equipment. IAS also supports the use of both certificate-based and password-based authentication protocols, and provides three forms of logging, including the ability to leverage the advantages of logging to a relational database with Structured Query Language (SQL) Server logging. In addition, you can outsource your remote-access infrastructure to a service provider while retaining control over user authentication, authorization, and accounting, thereby keeping sensitive organization information private and more secure on your network.

When an IAS server is a member of an Active Directory domain, IAS uses the Microsoft Active Directory directory service as its user account database and is part of a single sign-on solution. With single sign-on, users supply account credentials (user name and password or a certificate) only once during the authentication and authorization process; these credentials are then used to log on to an Active Directory domain and for network access control (authenticating and authorizing access to a network).

When IAS remote access policies are configured with Protected Extensible Authentication Protocol (PEAP) and a PEAP type as the authentication method for wireless connections, PEAP fast reconnect allows users to roam with portable devices between wireless access points without re-entering user account credentials each time the portable device associates with a new wireless access point.

You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you

cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the Domain Name System (DNS) query. With IAS in Microsoft Windows Server 2003, Enterprise Edition, and Microsoft Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range. IAS is not included on computers running the Windows Server 2003, Web Edition, operating system.

**Components of an IAS infrastructure**

There are five components to an IAS or RADIUS infrastructure: access clients, access servers (RADIUS clients), IAS servers (RADIUS servers), IAS proxies (RADIUS proxies), and user account databases. A RADIUS infrastructure is used to perform authentication, authorization and accounting of user network access attempts. Authentication is the process of verifying the credentials of the users attempting to connect to a network. The authorization process determines whether users have permission to connect to the network, and the conditions under which permission has been granted. Accounting is an option that provides record keeping of successful or failed connection attempts.

The following figure, "Components of an IAS Infrastructure," illustrates the relationships between the five components of an IAS infrastructure.

**Components of an IAS Infrastructure**




**Access Clients**

An access client is a device that requires some level of access to a larger network. Examples of access clients are dial-up or VPN clients, wireless clients, or LAN clients connected to an authenticating switch.

**Access Servers Used as RADIUS Clients**

An access server is a device that provides some level of access to a larger network. An access server using a RADIUS infrastructure is also a RADIUS client, sending connection requests and accounting messages to a RADIUS server. Examples of access servers follow:

- NASs that provide remote access connectivity to an organization network or the Internet. An example is a computer running Windows Server 2003 or Microsoft Windows 2000 that also runs the Routing and Remote Access service, providing either traditional dial-up or VPN remote access to an organization's intranet.
- Wireless access points that provide physical layer access to an organization's network, using wireless-based transmission and reception technologies.
- Switches that provide physical layer access to an organization's network, using traditional LAN technologies such as Ethernet.

**IAS Servers Used as RADIUS Servers**

An IAS or RADIUS server is a device that receives and processes connection requests or accounting messages sent by RADIUS clients or RADIUS proxies. In the case of connection requests, the RADIUS server processes the list of RADIUS attributes in the connection request. Based on a set of rules and the information in the user account database, the RADIUS server authenticates and authorizes the connection and sends back either an Access-Accept message or an Access-Reject message. The Access-Accept message can contain connection restrictions that the access server implements the duration of the connection.

**IAS Proxies and RADIUS Proxies**

An IAS or RADIUS proxy is a device that forwards or routes RADIUS connection requests and accounting messages between RADIUS clients (and RADIUS proxies) and RADIUS servers (or RADIUS proxies). The RADIUS proxy uses information within the RADIUS message, such as the User-Name or Called-Station-ID RADIUS attributes, to route the RADIUS message to the appropriate RADIUS server.

A RADIUS proxy can be used as a forwarding point for RADIUS messages when the authentication, authorization, and accounting must occur at multiple RADIUS servers in different organizations.

**User Accounts Databases**
The user account database is the list of user accounts and their properties that can be checked by a RADIUS server to verify authentication credentials and user account properties containing authorization and connection parameter information.

The user account databases that IAS can use are the local Security Accounts Manager (SAM), a Microsoft Windows NT Server 4.0 domain, or the Active Directory directory service and user accounts database included with Windows Server 2003 and Windows 2000. For Active Directory, IAS can provide authentication and authorization for user or computer accounts in the domain in which the IAS server is a member, two-way trusted domains, and trusted forests with domain controllers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

If the user accounts for authentication reside in a different type of database, IAS can be configured as a RADIUS proxy to forward the authentication request to a RADIUS server that does have access to the user account database. Different databases for Active Directory include untrusted forests, untrusted domains, or one-way trusted domains.

**IAS Features**
IAS in Windows Server 2003 has an extensive set of features. The core IAS features were introduced in Windows 2000 Server operating system; and new IAS features have been added to Windows Server 2003.

**Core IAS Features**
Core IAS features offer advantages such as multiple authentication methods and centralized administration for access servers.

**Authentication Methods**
The protection of user account credentials during the authentication of users attempting connections is an important security concern. IAS supports a variety of authentication protocols and allows you to use arbitrary authentication methods to meet your requirements for authentication.

- Password-based Point-to-Point Protocol (PPP) authentication protocols. PPP is a set of industry-standard framing and authentication protocols that enables remote access solutions to be interoperable in a multivendor network. IAS supports the authentication protocols within PPP, such as Password Authentication Protocol, CHAP, Microsoft Challenge Handshake Authentication Protocol (MS-CHAP v1), MS-CHAP version 2 (MS-CHAP v2), and Extensible Authentication Protocol (EAP).

- EAP and PEAP. An Internet standards-based infrastructure that allows the addition of arbitrary authentication methods, such as smart cards, certificates, one-time passwords, and token cards. A specific authentication method that uses the EAP infrastructure is an EAP type. IAS includes support for EAP-Message Digest 5 (MD5) and EAP-Transport Layer Security (EAP-TLS), as well as PEAP-MS-CHAP v2 and PEAP-TLS.

**Various Authorization Methods**
IAS supports a number of authorization methods and allows you to add custom methods that meet your authorization requirements. The supported authorization methods are:

- Dialed Number Identification Service (DNIS). The authorization of a connection attempt that is based on the number called. DNIS supplies the number that was called to the call receiver and is provided by most standard telephone companies.

- Automatic Number Identification/Calling Line Identification (ANI/CLI). The authorization of a connection attempt that is based on the phone number of the caller. ANI/CLI service supplies the number of the caller to the call receiver and is provided by most standard telephone companies.

- Guest authentication. The authorization of a connection when the caller does not send a user name or password during the authentication process. If unauthenticated access is enabled, the Guest account is used by default as the identity of the caller.

**Centralized user Authentication and Authorization**
To authenticate a connection request, IAS validates the connection credentials against user and computer accounts in the local SAM, a Windows NT Server 4.0 domain, or an Active Directory domain. For an Active Directory domain, IAS supports the use of Active Directory user principal names and universal groups.

To authorize a connection request, IAS uses the dial-in properties of the user account that correspond to both the connection credentials and remote access policies. One of the elements of authorization is remote access permission, which can be set both on the user or computer account and the remote access policy. Although it is relatively easy to manage remote access permission for each user account, this approach does not scale well as an organization grows. Remote access policies provide a more powerful and flexible way to manage remote access permission.

With remote access policies, you can authorize network access based on various conditions, including:

- User account membership in a group.
- The time of day or day of the week.
- The type of media by which the user is connecting (for example, wireless, Ethernet switch, modem, or VPN).
- The phone number that the user calls.
- The access server from which the request arrives.

By configuring profile settings on remote access policies, you can control many connection parameters, including:

- The use of specific authentication methods.
- The idle time-out.
- The maximum time of a single session.
- The number of links in a multilink session.
- The use of encryption and its strength.
- The use of packet filters to control what the remote access user can access when connected to the network. For example, you can use filters to control which IP addresses, hosts, and ports the user is allowed to use in sending or receiving packets.
- The creation of a compulsory tunnel that forces all packets from that connection to be securely tunneled through the Internet and then terminated in a private network.
- The virtual local area network identifier for wireless or Ethernet connections.

**Heterogeneous Access Servers**
IAS supports NASs that are compliant with RADIUS RFCs 2865 and 2866. In addition to dial-up access servers, IAS also supports the following:

- Wireless access points. By using remote access policies and the Wireless-IEEE 802.11 port type condition, IAS can be used as a RADIUS server for wireless access points that use RADIUS for authentication and authorization of wireless nodes.
- Authenticating switches. By using remote access policies and the Ethernet port type condition, IAS can be used as a RADIUS server for Ethernet network switches that use RADIUS for authentication and authorization of switch nodes.
- Computers running Windows Server 2003 or Windows 2000 and the Routing and Remote Access service. A computer running the Routing and Remote Access service can be configured as a dial-up server, a VPN server, or a VPN router for dial-up or persistent site-to-site (also known as router-to-router) connections between networks. IAS and the Routing and Remote Access service use the same remote access policies and log files. This integration provides a consistent

implementation for both IAS and the Routing and Remote Access service. It enables you to deploy Routing and Remote Access at small sites without the requirement of a separate, centralized IAS server. It also provides the ability to scale to a centralized remote access management model whenever you have multiple Routing and Remote Access servers in your organization. IAS, in conjunction with Routing and Remote Access servers, implements a single point of administration for remote access to your network for outsourced-dial, demand-dial, and VPN access. The policies within IAS at a central and large site can be exported to the independent Routing and Remote Access server in a small site.

### Note

- A site-to-site VPN connection is typically used to connect remote offices together when both routers are connected to the Internet by permanent wide area network (WAN) links, such as T1 or Frame Relay. In this configuration, the VPN connection is always available. However, when a permanent WAN link is not possible or practical, you can configure a dial-up site-to-site VPN connection.

**Centralized Administration for All Access Servers**
Support for the RADIUS standard allows IAS to control connection parameters for any NAS that implements RADIUS. The RADIUS standard also allows individual access server vendors to create proprietary extensions called vendor-specific attributes (VSAs). IAS has incorporated the extensions from a number of vendors in its dictionary of attributes. In circumstances where an attribute is not included in the IAS dictionary of attributes, additional VSAs can be created and added to the profile of individual remote access policies.

**Outsourced Dial-Up and Wireless Network Access**
Outsourced dialing (also known as wholesale dialing) provides a contract between an organization and an ISP. The ISP allows employees of the organization to connect to its network before the VPN tunnel to the private network of the organization is established. When an employee of the organization connects to the network access server of the ISP, the authentication and records of usage are forwarded to the IAS server at the organization. The IAS server enables the organization to control user authentication, track usage, and determine which employees are allowed to access the network of the ISP.

The advantages of outsourced dialing are the potential financial and administrative savings. By using an ISP's hardware and wide area network (WAN) links instead of purchasing and installing your own, you might save a great deal on infrastructure costs. If traveling or remotely located employees dial in to an ISP that has worldwide connections, making a local rather than a long distance connection, you might significantly decrease your long-distance phone bill. And by moving support requirements to the provider, you might eliminate administrative costs.

You can also outsource wireless access. A vendor can provide wireless access in a remote location and use IAS as a RADIUS proxy to forward your employee connection requests to a RADIUS server on your network for authentication and authorization.

**Centralized, Scalable Auditing and Usage Accounting**
You can use IAS as a RADIUS server to centralize the accounting data from many NASs, and you can use IAS SQL Server logging to further centralize accounting data by configuring multiple IAS servers to log to the same database on a computer running SQL Server 2000. Support for the RADIUS standard allows IAS to collect the usage (accounting) records sent by many access servers configured as RADIUS clients to the IAS server. IAS logs audit information (for example, authentication accepts and rejects) and usage information (for example, logon and logoff records) to log files. IAS supports a text-based log-file format that can be directly imported into a database, as well as SQL Server logging, where IAS passes accounting data to the SQL Server Desktop Engine (MSDE 2000) or SQL Server 2000 in XML format. After the accounting data in the database is collected, it can be queried and analyzed by non-Microsoft data-analysis software.

**IAS Integration with the Routing and Remote Access Service**
The Routing and Remote Access service can be configured to use Windows authentication and accounting, or to use RADIUS authentication and accounting. When RADIUS authentication or accounting is selected, any RFC-compliant RADIUS server can be used. However, using an IAS server is recommended to achieve the optimum level of integration in Windows Server 2003 or Windows 2000 environments and to take advantage of centralized remote access policies.

For example, in a small network environment or branch offices with a small number of remote access servers and no requirements for centralized management of remote access, the Routing and Remote Access service can be configured to use Windows authentication and accounting.

In a global enterprise with large numbers of different types of access servers deployed worldwide, centralized authentication and accounting using IAS can be beneficial. However, if a small remote office is configured with a low bandwidth connection to the global organization with the centralized IAS server, the Routing and Remote Access service servers at the branch office can be configured with Windows authentication and accounting and the remote access policy configuration can be copied from a

central location to the remote access servers of the remote office.

IAS and the Routing and Remote Access service share the same remote access policies and authentication capabilities. When the Routing and Remote Access service is configured for Windows authentication, local remote access policies are used, and logging is recorded in a local file by default. You can also configure the Routing and Remote Access service to log accounting data to a database on a computer running Microsoft SQL Server 2000.

When the Routing and Remote Access service is configured as a RADIUS client to an IAS server, the remote access policies of the IAS server are used and logging is recorded in a local file on the IAS server or, when IAS SQL Server logging is configured on the IAS server, to a SQL Server 2000 database.

Because the policies within IAS at a central large site can be exported to the independent remote access server in a small site, a consistent implementation across IAS and the Routing and Remote Access service is provided. It allows you to deploy the Routing and Remote Access service at small sites without the need for a separate, centralized IAS server; it also provides the capability to scale up to a centralized remote access management model when the need arises to do so. In this case, IAS in conjunction with remote access servers implements a single point of administration for remote access to your network for outsourced-dial, demand-dial, and VPN access.

**Graphical User Interface**
You can configure IAS by using the IAS Microsoft Management Console (MMC) snap-in, which enables you to configure local or remote IAS servers. You can add multiple IAS servers to a single snap-in, providing ease of management for many servers using one interface. If you install the Windows Server 2003 Administrative Tools Pack on a computer running Microsoft Windows XP Professional, you can use the IAS snap-in from the client computer to manage your IAS servers.

**Remote Monitoring**
You can monitor IAS by using Windows Server 2003–based tools, such as Event Viewer, Performance Logs and Alerts, and Simple Network Management Protocol (SNMP). You can also use Network Monitor to capture RADIUS messages for detailed traffic analysis and troubleshooting.

**Scalable Configurations**
You can scale IAS to network configurations of varying size, from stand-alone servers for small networks to large corporate and ISP networks. As your network grows, you can add access servers, IAS proxy servers, and IAS servers to scale up and out, exporting and importing server configurations to minimize administrative overhead. IAS logging to SQL Server databases also provides the ability to scale the logging of network session information.

**Import and Export of Configuration to Manage Multiple IAS Servers**
An IAS server configuration can be exported in part or total by using the Netsh commands for network access authentication, authorization, accounting, and auditing. For example, you can export the full configuration of an IAS server using the **netsh aaaa show config** command at the command prompt, or you can export only the remote access policies of the server using the **netsh aaaa show remote_acces_policies** command.

### Note

- The configuration of IAS servers running operating systems in the Windows 2000 Server family can be imported into products in the Windows Server 2003 family with netsh AAAA commands. The reverse, however, is not possible.

**New IAS Features Supported by Windows Server 2003**
In response to advancements in the RADIUS protocol and network access technologies, new features have been added to IAS in Windows Server 2003. For example, IAS in Windows Server 2003 provides support for: RADIUS proxy, mapping network authentication and authorization for IAS proxy. Support for IEEE 802.1X wireless and authenticating switches and PEAP for 802.11 wireless clients has also been added.

**RADIUS Proxy**
In addition to RADIUS server support, Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition provide RADIUS proxy support. With RADIUS proxy support, you can use IAS as a RADIUS message router or forwarder. Based on attributes in the incoming RADIUS message, the RADIUS proxy forwards the message to a specific RADIUS server or client.

You can configure IAS as a RADIUS proxy by using connection request processing and remote RADIUS server groups. Connection request processing is a set of rules that determine how an incoming connection request is processed: It is either authenticated and authorized by the IAS server (when the IAS server is being used as a RADIUS server), or it is forwarded to another RADIUS server for authentication and authorization (when the IAS server is being used as a RADIUS proxy). When IAS is being used as a proxy, it forwards the connection request to a member of a remote RADIUS server group.

By correctly configuring connection request processing rules and remote RADIUS server groups, an IAS server can be used as a

RADIUS server for some requests and a RADIUS proxy for others. The IAS proxy also provides the ability to balance RADIUS requests across multiple RADIUS servers, to detect both which RADIUS servers are no longer available and when they become available again.

**Mapping Network Authentication and Authorization for IAS Proxy**
When you map a user account in a remote user accounts database to a user account in a local user accounts database, the proxy component of IAS can separate the authentication and authorization of connection requests. IAS sends the authentication request to the remote IAS server and processes the authorization request locally. The IAS proxy can forward password-based user credentials to an external RADIUS server for authentication, and perform authorization against a user account in an Active Directory domain and a locally configured remote access policy.

### Note

- Members of a remote RADIUS server group can be IAS servers that authenticate connection requests by using Active Directory or they can be third-party RADIUS servers that can perform authentication by using other user account databases. Regardless of the user accounts database or RADIUS server used with account mapping, authorization is performed on the local IAS proxy.

To configure IAS to split authentication and authorization between two different IAS servers and user accounts databases, you can map the realm portion of a user name to a remote RADIUS server for authentication, even if that RADIUS server is located on another private network.

For example, IAS can authenticate a visitor from a partner organization by using the partner organization RADIUS server and user accounts database. IAS then authorizes access to your network by using connection request policy settings on your IAS server and a Windows user accounts database in an Active Directory domain that is established for visitor accounts.

You can configure the proxy component with the Remote-RADIUS-to-Windows-User-Mapping attribute in the advanced properties of a connection request policy.

**IEEE 802.1X Wireless Access Points and Authenticating Switches**
IAS provides authentication, authorization, and accounting for connections that use the link-layer standard IEEE 802.1X for wireless and switch access.

**Wireless access points**
Wireless networking technology has become more widely used since the adoption of wireless industry standards such as IEEE 802.11b, 802.11a, and 802.11g. You can use wireless networking to roam within a building or area and connect automatically to a network when you are within range of a wireless access point beacon.

While providing convenience, wireless networking technologies and wireless access points have the following security risks:

- Anyone who has a compatible wireless network adapter can gain access to the network.

- Wireless networking signals use radio waves to send and receive information. Anyone within an appropriate distance to a wireless access point can detect and receive all data sent to and from the wireless access point.

To counter the first security risk, wireless access points must require authentication and authorization of the wireless node before data can be sent to and received from the network to which the wireless access point is attached. To provide their own authentication and authorization, each wireless access point should require a user account database with each user's account credentials and a set of rules by which authorization is granted.

Because maintaining separate user accounts databases for each wireless access point is difficult to manage, it is recommended that you use wireless access points that function as RADIUS clients, and use the industry standard RADIUS protocol to send connection request and accounting messages to a central RADIUS server. The RADIUS server has access to a user account database and a set of rules for granting authorization. The RADIUS server processes the wireless access point's connection request and either accepts or rejects the connection request.

To counter the second security risk, the data sent between the wireless nodes and the wireless access points must be encrypted. Therefore, the authentication method that the wireless node uses must allow for the determination of encryption keys that are used to encrypt data.

IAS supports wireless access authentication by means of the following:

- The use of PEAP with MS-CHAP v2 for secure password-based authentication and fast reconnect when users roam with wireless computers between wireless access points that are configured as RADIUS clients to the same RADIUS server.

- The use of EAP-TLS to provide certificate-based authentication and the management of encryption keys that are used to encrypt data sent between wireless access points and wireless nodes.

- The use of the Wireless-IEEE 802.11 and Wireless-Other port types when configuring the NAS-Port-Type condition of remote access policies. By using these port types, you can create a separate remote access policy that contains connection parameters and encryption settings that are specifically designed for wireless nodes.

- By providing a means to grant guest access for a wireless client that does not have a certificate installed or a set of valid account name/password credentials.

**Authenticating switch support**

A network switch provides selective filtering of the traffic received on each port of the switch and allows better traffic management by segmenting a network. A typical switch, however, sends and receives packets to any node that is connected to the switch, and does not perform any authentication to verify that the user has permission to access the network. This can be a problem, for example, if switch ports are installed in a conference room. A person from another organization can simply connect a network adapter to the switch port and gain instant access to the network.

To prevent physical layer access to the network, newer switches can be configured to require switch node authentication and authorization before data is sent to or received from the network node that is attached to the switch. To provide its own authentication and authorization, each switch is required to have a user account database. This database contains each users authentication credentials and a set of rules by which authorization is granted. Multiple individual databases are difficult to manage; however, newer switches can serve as a RADIUS client, using the industry-standard RADIUS protocol to send connection requests and accounting messages to a central RADIUS server. The RADIUS server has access to a user account database and a set of rules for granting authorization. The RADIUS server processes the connection request from the switch and either accepts or rejects the request.

Typically, as the data sent to and from the switch node travels on the wire between the switch node and the switch, encryption of the sent data is neither required nor implemented.

IAS supports switch access authentication by the following means:

- EAP-TLS provides computer or user certificate-based authentication.

- The NAS-Port-Type condition of remote access policies is configured for the Ethernet port type. By using this port type, you can create a separate remote access policy containing connection parameters that are specifically designed for switch nodes.

Virtual local area network (VLAN) switching is used on the basis of whether or not the switch access client is authenticated.

**PEAP for 802.11 Wireless Clients**

PEAP provides protection for clients and authenticators (IAS or RADIUS servers) that are using EAP. The next generation of EAP and PEAP uses Transport Layer Security (TLS) to create end-to-end communication between client and authenticator after the identity of the authenticator is verified.

PEAP creates a session key that is used to derive keys for the encryption of authentication traffic between the wireless client and the IAS server. It also provides smooth and efficient roaming between wireless access points that are configured as RADIUS clients on the same IAS server.

You can use PEAP with two different authentication types: PEAP-MS-CHAP v2 or PEAP-TLS. PEAP-MS-CHAP v2 provides secure password authentication for 802.1X wireless and authenticating switch clients. PEAP-MS-CHAP v2 can be used to complement the security of MS-CHAP v2, and it does not require deployment of user or computer certificates on the client computer. A server certificate is required on the IAS server. With PEAP-MS-CHAP v2, you can obtain a server certificate from a third party certification authority, such as Verisign, instead of deploying your own public key infrastructure (PKI).

You can also use PEAP-TLS, which provides strong security by means of a PKI with certificates. A computer certificate installed on the IAS server is used to authenticate the IAS server and a locally-installed computer and user certificate or a smart card on the wireless client is used to authenticate the wireless client. When you use PEAP-TLS, client certificate information is encrypted, providing improved security over the use of EAP-TLS without PEAP.

**Note**

- When you deploy both PEAP and EAP unprotected by PEAP on the same network, do not use the same EAP authentication type with and without PEAP for the same network access type. For example, if you deploy PEAP-TLS on your wireless network, do not also deploy EAP-TLS. Deploying authentication methods with the same type (one with and the other without the protection of PEAP) for the same network access type creates a security issue.

**Enhanced EAP Configuration for Remote Access Policies**
In Microsoft Windows 2000 Server, you can select only a single EAP type when configuring the kind of authentication allowed by a remote access policy. All connections matching the conditions of the policy must use this single EAP type, which is selected in the policy profile settings. Additionally, the configuration of an EAP type is global for all of the remote access policies configured on the IAS server. This can cause problems when you want to either individually configure the properties for EAP types for each policy or select multiple EAP types by group or for a specific kind of network connection.

In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, these limitations are removed for IAS.

For example, you can select different computer certificates for VPN connections and EAP-TLS authentication for wireless connections, or you can select multiple EAP types for wireless connections in the circumstance where some of your wireless clients use EAP-TLS authentication, and some of them use PEAP-MS-CHAP v2.

**Network Access Quarantine Control**
Network Access Quarantine Control provides phased network access, which restricts remote clients to quarantine mode until each client is either verified as meeting or configured according to organization network access policy. Quarantine mode is applied by a VPN server running Windows Server 2003 and the Routing and Remote Access service. Quarantine mode is created using a session timer, or an IP filter, or both. The IP filter and session timer are configured in a remote access policy in the IAS snap-in. After the client computer configuration is verified as meeting organization network policy by an administrator-provided script that runs on the client, the quarantine restrictions are removed from the to the connection. To deploy this feature, you must use Connection Manager Administration Kit, the Routing and Remote Access service, and additional components deployed on clients and access servers.

**Logging to a SQL Server Database**
You can use Microsoft SQL Server 2000 to log IAS accounting information, such as user authentication requests, accounting requests, and periodic data, to a database that warehouses data from multiple IAS servers. To diminish cost and increase logging performance, you can install the Microsoft SQL Server Desktop Engine (MSDE 2000) on each IAS server, and then periodically publish records from the MSDE 2000 database on each IAS server to the central database on the SQL Server, providing comprehensive data views for troubleshooting network access authentication and report creation.

**Ignoring the Dial-In Properties of User Accounts**
To ignore the dial-in properties of user accounts, you can configure a RADIUS attribute on the profile properties of a remote access policy. The dial-in properties of the user account contain the following properties:

- Remote access permission
- Caller-ID
- Callback options
- Static IP address
- Static routes

To support multiple types of connections for which IAS provides authentication and authorization, you might need to disable the processing of user account dial-in properties. This can be done to support scenarios in which specific dial-in properties are not required.

For example, the caller-ID, callback, static IP address, and static routes properties are designed for a client that is dialing into a NAS, not for wireless access points. A wireless access point that receives these settings in a RADIUS message might not be able to process them, which could cause the wireless client to be disconnected. When IAS provides authentication and authorization for users who are both dialing in and accessing the organization network by means of wireless technology, the dial-in properties must be configured to support either dial-in connections (by setting dial-in properties) or wireless connections (by not setting dial-in properties).

You can use IAS to enable dial-in properties processing for the user account in some scenarios (such as dial-up connections)

and to disable dial-in properties processing for user account dial-in properties in other scenarios (such as wireless and authenticating switch). This is accomplished by configuring the **Ignore-User-Dialin-Properties** attribute on the **Advanced** tab of the profile settings for a remote access policy.

The **Ignore-User-Dialin-Properties** attribute is set to the following options:

- To enable user account dial-in properties processing, delete the **Ignore-User-Dialin-Properties** attribute or set it to **False**. For example, for a remote access policy that is designed for dial-in connections, no additional configuration is required.

- To disable user account dial-in properties processing, set the **Ignore-User-Dialin-Properties** attribute to the value of **True**. For example, this is set for the remote access policy that is designed for wireless or authenticating switch connections. When the dial-in properties of the user account are ignored, remote access permission is determined by the remote access permission setting in the remote access policy.

You can also use this attribute to manage network access control by means of groups and the remote access permissions on the remote access policy. By setting the **Ignore-User-Dialin-Properties** attribute to the value of **True**, the remote access permission on the user account is ignored. The disadvantage to using the **Ignore-User-Dialin-Properties** attribute in this way is that you cannot use the additional dial-in properties of caller-ID, callback, static IP address, and static routes.

**Note**

- The **Ignore-User-Dialin-Properties** attribute disables the use of all dial-in properties for the user account. Specific dial-in properties cannot be selectively disabled.

#### Configuring RADIUS Clients by IP Address Range

For IAS in Windows 2000, you must specify a RADIUS client by IP address or by DNS name. In addition, you must configure each RADIUS client separately, even if you have a number of RADIUS clients on the same subnet. While this is not an issue for typical dial-in or VPN access server configurations, wireless access points are frequently deployed in large numbers on the same subnet, and configuring each wireless access point as a RADIUS client at the RADIUS server can be time consuming.

In Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, you can use IAS to specify a group of RADIUS clients by using an IP address range. This simplifies the configuration of multiple RADIUS clients that are on the same network segment, or that are configured on a virtual local area network (VLAN). When you configure RADIUS clients by IP address range, all RADIUS clients in the range must be identically configured and use the same shared secret.

When you configure RADIUS clients using this method, the IP address range for RADIUS clients is expressed in the network prefix length notation *w.x.y.z/p*, where *w.x.y.z* is the dotted decimal notation of the address prefix and *p* is the prefix length (the number of high order bits that define the network prefix). This is also known as Classless Inter-Domain Routing notation. An example is 192.168.21.0/24. To convert from subnet mask notation to network prefix length notation, *p* is the number of high-order bits set to one in the subnet mask.

**Note**

- You cannot configure RADIUS clients by using the IP address range in IAS for Windows Server 2003, Standard Edition.

#### Computer Authentication

In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, Active Directory and IAS support the authentication of computer accounts by using standard user authentication methods, such as PPP. This allows a computer and its computer certificate to be authenticated for wireless or authenticating switch access clients.

#### Checking User Certificate Purposes

To enforce the use of specific types of user certificates for specific connection types, you can configure IAS to check the purposes, also called object identifiers (OIDs), of certificates in their Enhanced Key Usage (EKU) extensions when the client computer supplies the certificate as proof of identity.

When PEAP-TLS or EAP-TLS are configured as authentication methods for a remote access policy in IAS, the access client can send a smart card certificate or a user certificate from the client computer certificate store during the authentication process.

Because smart card certificates are more secure than certificates from the client computer certificate store, you might want to ensure that remote access VPN connections use a smart card certificate for user authentication instead of a locally-installed

user certificate. In this case, you can configure a remote access policy to require that the user certificate EKU extensions contain the **Smart Card Logon** purpose object identifier (1.3.6.1.4.1.311.20.2.2). When the VPN client offers the smart card certificate for user authentication, the IAS server checks the EKU extensions according to remote access policy. Because the **Smart Card Logon** object identifier is present in the VPN client certificate, authentication succeeds.

You can configure more than one object identifier that is required to be present in the user certificate, which the access client offers to IAS. To configure more than one object identifier, use the **Allow certificates with these OIDs** attribute on the **Advanced** tab in the properties of a remote access policy profile. No required OIDs are specified by default.

**Improved Support for Attribute Manipulation**
In Windows 2000, you can use IAS to manipulate the contents of the User-Name RADIUS attribute. Using connection request policies in IAS for Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, you can manipulate the User-Name, Called-Station-ID, and Calling-Station-ID RADIUS attributes.

**Authentication Type Remote Access Policy Condition**
You can create remote access policies by using the Authentication Type condition. You can use Authentication Type to specify connection constraints that are based on the authentication protocol or method that is used for authentication by the access client.

**Improved Support for the Class Attribute**
In Windows 2000, IAS generates a value for the Class attribute and appends it to the existing value of the Class attribute received in the RADIUS request message by NASs that send the Class attribute. The result is the Class attribute in the RADIUS response message. In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition:

- You can disable the automatic generation of a value for the Class attribute by using **Generate-Class-Attribute** on the **Advanced** tab in the properties of a remote access policy profile. Automatic generation of a value for the Class attribute is disabled by default.

- Instead of appending the generated value of the Class attribute to the existing Class attribute, IAS creates a separate Class attribute. The RADIUS response message contains both the original Class attribute and the second Class attribute that IAS generates.

### Note

- To provide the best session correlation for applications that query the accounting log, it is recommended that you use NASs that are capable of sending the Class attribute.

# How IAS Technology Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- RADIUS Protocol
- IAS Roles
- IAS as a RADIUS Server
- IAS as a RADIUS Proxy


Internet Authentication Service (IAS) is the Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy in Microsoft Windows Server 2003; Standard Edition, Windows Server 2003; Enterprise Edition, and Windows Server 2003; Datacenter Edition. As a RADIUS server, IAS performs centralized connection authentication, authorization, accounting, and auditing (AAAA) for many types of network access, including wireless, authenticating switch, dial-up and virtual private network (VPN) remote access, and router-to-router connections. As a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers. IAS supports the Internet Engineering Task Force (IETF) standards for RADIUS described in RFC 2865 and RFC 2866.

**RADIUS Protocol**

RADIUS is an industry standard protocol described in RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)," and RFC 2866, "RADIUS Accounting." RADIUS is used to provide authentication, authorization, and accounting services. A RADIUS client, — typically a dial-up server, VPN server, or wireless access point (AP) — sends user credentials and connection parameter information in the form of a RADIUS message to a RADIUS server. The RADIUS server authenticates and authorizes the RADIUS client request, and sends back a RADIUS message response. RADIUS clients also send RADIUS accounting messages to RADIUS servers. Additionally, the RADIUS standards support the use of RADIUS proxies. A RADIUS proxy is a computer that forwards RADIUS messages between RADIUS-enabled computers.

RADIUS messages are sent as UDP messages. UDP port 1812 is used for RADIUS authentication messages and UDP port 1813 is used for RADIUS accounting messages. Some network access servers (NASs) might use UDP port 1645 for RADIUS authentication messages and UDP port 1646 for RADIUS accounting messages. By default, IAS supports receiving RADIUS messages sent to both sets of UDP ports.

The RADIUS protocol is implemented solely between NASs (RADIUS clients) and IAS servers and proxies (RADIUS servers and proxies). A user's client computer does not use the RADIUS protocol when connecting to a NAS.

**RADIUS Authentication Process**

The RADIUS authentication process begins when a user attempts to access a network through a NAS that is configured as a RADIUS client to a RADIUS server.

For example, when the remote access user sends credentials using the Challenge Handshake Authentication Protocol (CHAP), the RADIUS client creates a RADIUS Access-Request message containing such attributes as the Port ID the user is accessing and the results of the CHAP authentication process (the user name, the challenge string, and the response of the remote access client).

The RADIUS Access-Request message is sent from the RADIUS client to the RADIUS server. If a response is not received within a length of time, the request is re-sent. The RADIUS client can also be configured to forward requests to an alternate server or servers in the event that the primary server is down or unreachable. An alternate server can be used either after a specified number of tries to the primary server are not responded to, or in a round-robin fashion.

In the case of the Routing and Remote Access service, multiple RADIUS servers can be added and prioritized using a scoring mechanism. If a primary RADIUS server does not respond within three seconds, the Routing and Remote Access service automatically switches to the RADIUS server with the next highest score.

After the RADIUS server receives the request, it validates the sending RADIUS client. Validation occurs by verifying that the RADIUS Access-Request message is sent from a configured RADIUS client. If the Access-Request message is sent by a valid

RADIUS client, and if the Message-Authenticator attribute is required for the RADIUS client, the value of the Message-Authenticator attribute is verified using the RADIUS shared secret. If the Message-Authenticator attribute is either missing or contains an incorrect value, the RADIUS message is silently discarded — that is, the message is discarded without logging an event in the Event Log or making an entry in the IAS accounting log.

If the RADIUS client is valid, the RADIUS server consults a database of users to find the user whose name matches the User-Name attribute in the request. The user account contains a list of requirements that must be met to allow access for the user. This list of requirements can include verification of the password, and it can also specify whether the user is allowed access.

If any condition where the authentication or authorization is not met, the RADIUS server sends a RADIUS Access-Reject message in response, indicating that this user request is not valid.

If all conditions are met, the list of configuration values for the user is placed into a RADIUS Access-Accept message that is sent back to the RADIUS client. These values include a list of RADIUS attributes and all necessary values to deliver the desired service. For Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP) service types, this can include values such as encryption types, the Class attribute, Maximum Transmission Unit (MTU), as well as the desired compression and packet filter identifiers.

**RADIUS Message Format**

The following section provides information that might be useful for the following:

- Understanding a Network Monitor capture.
- Understanding the different message formats for analyzing the accounting log.
- Entering vendor-specific attribute (VSA) numbers.

RADIUS messages sent to the RADIUS server are sent as UDP messages. RADIUS authentication messages are sent using UDP port 1812, and RADIUS accounting messages are sent on UDP port 1813. Some older NASs use UDP port 1645 for RADIUS authentication messages and UDP port 1646 for RADIUS accounting messages.

IAS supports the receiving of RADIUS messages on any configurable set of ports; however by default it is configured to receive messages on both sets of UDP ports. Exactly one RADIUS message is encapsulated in the UDP payload.

**General Packet Structure**

The figure, "General Structure of RADIUS Packet," provides a summary of the data structure of a RADIUS packet. The RADIUS client or server sends the fields from top to bottom, or from the Code field in vertical order to the Attributes.

### General Structure of RADIUS Packet




**Code field**

The Code field is 1 byte long and indicates the type of RADIUS message. A message with an invalid Code field is silently discarded. The defined values for the RADIUS Code field are listed in the following table "Values for the RADIUS Code Field."

### Values for the RADIUS Code Field

| Codes (Decimal) | Packets |
|---|---|
| 1 | Access-Request |
| 2 | Access-Accept |
| 3 | Access-Reject |
| 4 | Accounting-Request |
| 5 | Accounting-Response |
| 11 | Access-Challenge |

| | |
|---|---|
| 12 | Status-Server (experimental) |
| 13 | Status-Client (experimental) |
| 255 | Reserved |

**Identifier field**

The Identifier field is 1 byte long and is used to match a request with its corresponding response.

**Length field**

The Length field is two octets long and indicates the entire length of the RADIUS message, including the Code, Identifier, Length, and Authenticator fields, and the RADIUS attributes. The Length field can vary from 20 to 4,096 bytes.

**Authenticator field**

The Authenticator field is sixteen octets long and contains the information that the RADIUS client and server use to verify that the message came from a computer that is configured with a common shared secret

**Attributes section**

The Attributes section of the RADIUS message contains one or more RADIUS attributes, which carry the specific authentication, authorization, information, and configuration details for RADIUS messages.

**RADIUS Attributes**

The figure, "RADIUS Attribute Structure," shows the structure of each RADIUS attribute. RADIUS attributes use the common Type-Length-Value format used by other protocols.

**RADIUS Attribute Structure**




**Type field**

The Type field is 1 byte long and indicates the specific type of RADIUS attribute. The RADIUS Standard attributes are listed in the following table, "RADIUS Standard Attribute Types." For information about other RADIUS attributes and their use, see RFCs 2865 and 2866.

### RADIUS Standard Attribute Types

| Type Values | Description | Type Values | Description |
|---|---|---|---|
| 1 | User-Name | 45 | Acct-Authentic |
| 2 | User-Password | 46 | Acct-Session-Time |
| 3 | CHAP-Password | 47 | Acct-Input-Packets |
| 4 | NAS-IP-Address | 48 | Acct-Output-Messages |
| 5 | NAS-Port | 49 | Acct-Terminate-Cause |
| 6 | Service-Type | 50 | Acct-Multi-Session-Id |
| 7 | Framed-Protocol | 51 | Acct-Link-Count |
| 8 | Framed-IP-Address | 52 | Acct-Input-Gigawords |
| 9 | Framed-IP-Netmask | 53 | Acct-Output-Gigawords |
| 10 | Framed-Routing | 55 | Event-Timestamp |
| 11 | Filter-ID | 60 | CHAP-Challenge |
| 12 | Framed-MTU | 61 | NAS-Port-Type |
| 13 | Framed-Compression | 62 | Port-Limit |

| | | | |
|---|---|---|---|
| 14 | Login-IP-Host | 63 | Login-LAT-Port |
| 15 | Login-Service | 64 | Tunnel-Type |
| 16 | Login-TCP-Port | 65 | Tunnel-Medium-Type |
| 18 | Reply-Message | 66 | Tunnel-Client-Endpt |
| 19 | Callback-Number | 67 | Tunnel-Server-Endpt |
| 20 | Callback-Id | 68 | Acct-Tunnel-Connection |
| 22 | Framed-Route | 69 | Tunnel-Password |
| 23 | Framed-IPX-Network | 70 | ARAP-Password |
| 24 | State | 71 | ARAP-Features |
| 25 | Class | 72 | ARAP-Zone-Access |
| 26 | Vendor-Specific | 73 | ARAP-Security |
| 27 | Session-Time-out | 74 | ARAP-Security-Data |
| 28 | Idle-Time-out | 75 | Password-Retry |
| 29 | Termination-Action | 76 | PROMPT |
| 30 | Called-Station-Id | 77 | Connect-Info |
| 31 | Calling-Station-Id | 78 | Configuration-Token |
| 32 | NAS-Identifier | 79 | EAP-Message |
| 33 | Proxy-State | 80 | Message-Authenticator |
| 34 | Login-LAT-Service | 81 | Tunnel-Pvt-Group-ID |
| 35 | Login-LAT-Node | 82 | Tunnel-Assignment-ID |
| 36 | Login-LAT-Group | 83 | Tunnel-Preference |
| 37 | Framed-AppleTalk-Link | 84 | ARAP-Challenge-Response |
| 38 | Framed-AppleTalk-Network | 85 | Acct-Interim-Interval |
| 39 | Framed-AppleTalk-Zone | 86 | Acct-Tunnel-Packets-Lost |
| 40 | Acct-Status-Type | 87 | NAS-Port-Id |
| 41 | Acct-Delay-Time | 88 | Framed-Pool |
| 42 | Acct-Input-Octets | 90 | Tunnel-Client-Auth-ID |
| 43 | Acct-Output-Octets | 91 | Tunnel-Server-Auth-ID |
| 1 | User-Name | 45 | Acct-Authentic |
| 2 | User-Password | 46 | Acct-Session-Time |
| 44 | Acct-Session-Id | | |

**Length field**

The Length field indicates the length of the attribute, including the Type, Length, and Value fields.

**Value field**

The Value field is zero or more octets and contains information specific to the Attribute. The format and length of the Value field is based on the type of RADIUS attribute. Note that value 26 is reserved for VSAs.

**Vendor-Specific Attributes**
VSAs are available to allow vendors to support their own proprietary attributes that are not covered in RFC 2865. IAS includes VSAs from a number of vendors in its multivendor dictionary. However, this list evolves over time and new attributes, and vendors are always being added.

You can add attributes that are not in the IAS multivendor dictionary as Vendor-Specific (RADIUS standard attribute type 26) on the **Advanced** tab of a remote access policy profile. Therefore, unlike many RADIUS implementations that allow you to add attributes to their multivendor dictionary, IAS allows you to add attributes to each remote access policy profile. To use attribute type 26, an administrator must know the VSA format and the exact information to enter. The VSA formats are documented in the following section. For information about what values to enter, see your NAS documentation. The following figure, "Vendor-Specific Attribute Structure," shows the structure of the VSA structure.

### Vendor-Specific Attribute Structure




**Type value**
The Type value is set to 26 (0x1A) to indicate a VSA.

**Length value**
The Length value is set to the number of bytes in the VSA.

**Vendor-ID**
Vendor-ID is 4 octets long. The high-order octet is 0, and the low-order is 3 octets. Together, they make up the Structure and Identification of Management Information (SMI) Network Management Private Enterprise Code of the vendor.

**String field**
The String field is the VSA consisting of one or more octets. To conform to the recommendation of RFC 2865, the String field should consist of the fields as shown in the figure "Structure of the String Field."

### Structure of the String Field




**Vendor type**
The Type value is used to indicate a specific VSA for the vendor.

**Vendor length**
The Type value is set to the number of bytes in the string.

**Attribute-Specific field**
The Attribute-Specific field contains the data for the specific vendor attribute.

Vendors who do not conform to RFC 2865 use the attribute type 26 to identify a VSA but do not use the Vendor Type, Vendor Length, and Attribute-Specific fields within the String field. In this case, the VSA format appears as shown in the figure, "Structure of the String Field," earlier in this section.

When adding a VSA for a particular NAS as type 26, you need to know whether the attribute conforms to RFC 2865. For information about whether your NAS uses the VSA format documented in the figure, "Structure of the String Field," earlier in this section, see your NAS documentation.

VSAs are configured from the **Vendor-Specific Attribute Information** dialog box when adding a Vendor-Specific Attribute from the **Advanced** tab of a remote access policy profile. If the VSA format conforms to RFC 2865, use the **Yes. It conforms.** option and configure the attribute with the vendor-assigned attribute number, attribute format, and attribute value as defined in NAS documentation. If the VSA format does not conform to RFC 2865, choose **No. It does not conform,** and configure the attribute with the hexadecimal attribute value, which includes the string of the VSA format (everything after Vendor-ID) as defined in NAS documentation.

**RADIUS Message Example**
A Windows Server 2003 Point-to-Point Tunneling Protocol (PPTP) client attempts a remote access connection to a Windows

Server 2003 VPN server. The VPN server is at the IP address 10.10.210.13, and the IAS server is at the IP address 10.10.210.12.

**Access-Request Message**

The following Network Monitor capture display shows the Access-Request message sent by the VPN server to the IAS server.

```
+ IP: ID = 0x850; Proto = UDP; Len: 248
+ UDP: Src Port: Unknown, (1327); Dst Port: Unknown (1812); Length = 228 (0xE4)
  RADIUS: Message Type: Access Request(1)
  RADIUS: Message Type = Access Request
  RADIUS: Identifier = 2 (0x2)
  RADIUS: Length = 220 (0xDC)
  RADIUS: Authenticator = 8A 6F DC 03 23 5F 4B 62 CA 40 92 38 DC 75
          CB 74
  RADIUS: Attribute Type: NAS IP Address(4)
    RADIUS: Attribute type = NAS IP Address
    RADIUS: Attribute length = 6 (0x6)
    RADIUS: NAS IP address = 10.10.210.13
  RADIUS: Attribute Type: Service Type(6)
    RADIUS: Attribute type = Service Type
    RADIUS: Attribute length = 6 (0x6)
    RADIUS: Service type = Framed
  RADIUS: Attribute Type: Framed Protocol(7)
    RADIUS: Attribute type = Framed Protocol
    RADIUS: Attribute length = 6 (0x6)
    RADIUS: Framed protocol = PPP
  RADIUS: Attribute Type: NAS Port(5)
    RADIUS: Attribute type = NAS Port
    RADIUS: Attribute length = 6 (0x6)
    RADIUS: NAS port = 32 (0x20)
  RADIUS: Attribute Type: Vendor Specific(26)
    RADIUS: Attribute type = Vendor Specific
    RADIUS: Attribute length = 12 (0xC)
    RADIUS: Vendor ID = 311 (0x137)
    RADIUS: Vendor string =    _
  RADIUS: Attribute Type: Vendor Specific(26)
    RADIUS: Attribute type = Vendor Specific
    RADIUS: Attribute length = 18 (0x12)
    RADIUS: Vendor ID = 311 (0x137)
    RADIUS: Vendor string = MSRASV5.00
  RADIUS: Attribute Type: NAS Port Type(61)
    RADIUS: Attribute type = NAS Port Type
    RADIUS: Attribute length = 6 (0x6)
    RADIUS: NAS port type = Virtual
  RADIUS: Attribute Type: Tunnel Type(64)
    RADIUS: Attribute type = Tunnel Type
    RADIUS: Attribute length = 6 (0x6)
    RADIUS: Tag = 0 (0x0)
    RADIUS: Tunnel type = Point-to-Point Tunneling Protocol(PPTP)
  RADIUS: Attribute Type: Tunnel Media Type(65)
    RADIUS: Attribute type = Tunnel Media Type
    RADIUS: Attribute length = 6 (0x6)
    RADIUS: Tag = 0 (0x0)
    RADIUS: Tunnel media type = IP (IP version 4)
  RADIUS: Attribute Type: Calling Station ID(31)
    RADIUS: Attribute type = Calling Station ID
    RADIUS: Attribute length = 14 (0xE)
    RADIUS: Calling station ID = 10.10.14.226
  RADIUS: Attribute Type: Tunnel Client Endpoint(66)
    RADIUS: Attribute type = Tunnel Client Endpoint
    RADIUS: Attribute length = 14 (0xE)
    RADIUS: Tunnel client endpoint = 10.10.14.226
  RADIUS: Attribute Type: User Name(1)
    RADIUS: Attribute type = User Name
    RADIUS: Attribute length = 18 (0x12)
    RADIUS: User name = NTRESKIT\johndoe
  RADIUS: Attribute Type: Vendor Specific(26)
    RADIUS: Attribute type = Vendor Specific
    RADIUS: Attribute length = 24 (0x18)
    RADIUS: Vendor ID = 311 (0x137)
    RADIUS: Vendor string = _+-_e_$+fN<N
  RADIUS: Attribute Type: Vendor Specific(26)
    RADIUS: Attribute type = Vendor Specific
```

```
    RADIUS: Attribute length = 58 (0x3A)
    RADIUS: Vendor ID = 311 (0x137)
    RADIUS: Vendor string = _4
```

The RADIUS attributes sent by the VPN server include the framed protocol, the service type, the class, various tunnel attributes for the PPTP connection, and a series of VSAs for Microsoft CHAP (MS-CHAP v1) authentication. For more information about Microsoft vendor-specific RADIUS attributes, see RFC 2548.

**Access-Accept Message**

The following Network Monitor capture display shows the Access-Accept message sent by the IAS server to the VPN server.

```
+ IP: ID = 0xB18; Proto = UDP; Len: 248
+ UDP: Src Port: Unknown, (1812); Dst Port: Unknown (1327); Length = 228 (0xE4)
  RADIUS: Message Type: Access Accept(2)
  RADIUS: Message Type = Access Accept
  RADIUS: Identifier = 2 (0x2)
  RADIUS: Length = 220 (0xDC)
  RADIUS: Authenticator = 52 E19 98 2E F8 E2 D3 B7 3B E1 24 5B 72          55 9E
  RADIUS: Attribute Type: Framed Protocol(7)
    RADIUS: Attribute type = Framed Protocol
    RADIUS: Attribute length = 6 (0x6)
    RADIUS: Framed protocol = PPP
  RADIUS: Attribute Type: Service Type(6)
    RADIUS: Attribute type = Service Type
    RADIUS: Attribute length = 6 (0x6)
    RADIUS: Service type = Framed
  RADIUS: Attribute Type: Class(25)
    RADIUS: Attribute type = Class
    RADIUS: Attribute length = 32 (0x20)
    RADIUS: Class = <$_@
  RADIUS: Attribute Type: Vendor Specific(26)
    RADIUS: Attribute type = Vendor Specific
    RADIUS: Attribute length = 42 (0x2A)
    RADIUS: Vendor ID = 311 (0x137)
    RADIUS: Vendor string = _$_DZ,Sc7__:+RW_t-qxF                        (-+%p6
  RADIUS: Attribute Type: Vendor Specific(26)
    RADIUS: Attribute type = Vendor Specific
    RADIUS: Attribute length = 42 (0x2A)
    RADIUS: Vendor ID = 311 (0x137)
    RADIUS: Vendor string = _$_
  RADIUS: Attribute Type: Vendor Specific(26)
    RADIUS: Attribute type = Vendor Specific
    RADIUS: Attribute length = 51 (0x33)
    RADIUS: Vendor ID = 311 (0x137)
    RADIUS: Vendor string = _-
  RADIUS: Attribute Type: Vendor Specific(26)
    RADIUS: Attribute type = Vendor Specific
    RADIUS: Attribute length = 21 (0x15)
    RADIUS: Vendor ID = 311 (0x137)
    RADIUS: Vendor string =
```

The RADIUS attributes sent by the IAS server include the user name, the service type, the framed protocol, the service class, and a series of VSAs for MS-CHAP authentication.

**Validating an Incoming RADIUS Message**

For RADIUS packets for which IAS is acting as RADIUS server, IAS performs the following validation checks:

1. IAS first verifies that the RADIUS message was sent from a configured RADIUS client by checking the source IP address of the RADIUS message.

2. If the message is not from a configured RADIUS client, it is silently discarded and an event is logged in the system event log.

3. IAS checks to ensure that the packet type is valid for the port on which it was received and then checks the structure of the RADIUS message. Items checked include valid values of the Code field, the size of the RADIUS message, and the size of the RADIUS attributes in the RADIUS message. IAS also checks whether or not the attributes themselves are the expected data type. If the attribute is not the expected data type, or if the RADIUS message is malformed, the message is

discarded and an event is logged in the system event log.

4. The User-Name attribute is checked for the value of the **Ping User-Name** registry setting. If there is a match between the value of the User-Name attribute and the value of the **Ping User-Name** registry setting, IAS sends an Access-Reject for authentication requests and an accounting response for accounting requests.

5. If the Message-Authenticator attribute is present, its value is validated using the shared secret. A shared secret is a text string that serves as a password between a RADIUS server and a RADIUS client.

For RADIUS request packets for which IAS is acting as RADIUS proxy, IAS performs the following validation checks:

1. IAS first verifies that the RADIUS message was sent from a configured RADIUS client by checking the source IP address of the RADIUS message.

   If the message is not from a configured RADIUS client, it is silently discarded, and an event is logged in the system event log.

2. IAS checks to ensure that the packet type is valid for the port on which it was received and then checks the structure of the RADIUS message. Items checked include valid values of the Code field, the size of the RADIUS message, and the size of the RADIUS attributes in the RADIUS message.

   If the RADIUS message is malformed, it is silently discarded and an event is logged in the system event log.

3. IAS checks for the presence of the Proxy-State attribute and whether the value of the Proxy-State attribute corresponds to an active proxy session in the proxy session table. If the RADIUS message does not contain a valid Proxy-State attribute, it is silently discarded and an event is logged in the system event log.

4. If the Message-Authenticator attribute is present, its value is validated using the shared secret. A shared secret is a text string that serves as a password between a RADIUS server and a RADIUS client. If the Message-Authenticator attribute is required but missing, or verification of its value fails, the message is silently discarded, and an event is logged in the system event log.

For RADIUS response packets for which IAS is acting as RADIUS proxy, IAS performs the following validation checks:

1. IAS first verifies that the RADIUS message was sent from a configured remote RADIUS server by checking the source IP address of the RADIUS message.

2. If the message is not from a configured remote RADIUS server, it is silently discarded and an event is logged in the system event log.

3. IAS checks to ensure that the packet type is valid for the port on which it was received and then checks the structure of the RADIUS message. Items checked include valid values of the Code field, the size of the RADIUS message, and the size of the RADIUS attributes in the RADIUS message.

4. If the RADIUS message is malformed, it is silently discarded, and an event is logged in the system event log.

5. IAS checks for the presence of the Proxy-State attribute and whether the value of the Proxy-State attribute corresponds to an active proxy session in the proxy session table. If the RADIUS message does not contain a valid Proxy-State attribute, it is silently discarded and an event is logged in the system event log.

6. IAS checks the IP source address and UDP source port of the message against the matching entry in the proxy session table to verify that the message was sent from the remote RADIUS server to which the initial request was sent.

7. The value of the Authenticator field is checked.

   If the RADIUS message contains an invalid Authenticator field, it is silently discarded and an event is logged in the system event log. This is typically caused by mismatched shared secrets between the IAS server and the remote RADIUS server

sending the RADIUS message.

8. If the Message-Authenticator attribute is present, its value is validated using the shared secret. If the Message-Authenticator attribute is required but missing, or verification of its value fails, the message is silently discarded, and an event is logged in the system event log.

**IAS Roles**

In Microsoft Windows 2000, you can use IAS only as a RADIUS server. In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, you can use IAS as either a RADIUS server or a RADIUS proxy. In addition, you can use IAS simultaneously as both proxy and server to perform AAAA on some Access-Request messages while forwarding others to remote RADIUS server groups. For multiple domain environments, an IAS server can authenticate credentials for user accounts in the domain of which it is a member and all domains that trust this domain. To read the dial-in properties for user, you must add the computer account of the IAS server to the **RAS and IAS Servers** group for each domain.

When configuring your IAS server to function as a RADIUS server, proxy server, or both, it is important to understand how IAS processes each Access-Request message received from a RADIUS client.

**Using IAS as a RADIUS Server**

When you use IAS as a RADIUS server, Access-Request messages are authenticated with Active Directory, a Microsoft Windows NT Server 4.0 domain, or by the Security Account Manager (SAM) in Microsoft Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. Access-Request messages are authorized with the user or computer account dial-in properties and remote access policies that are configured on the IAS server. Additionally, IAS can return an Access-Accept message without authenticating and authorizing the Access-Request message.

The contents of user authentication and accounting requests are logged in a local log file or are sent to a Structured Query Language (SQL) Server database, depending on how remote access logging settings are configured in the IAS console.

You can use IAS as a RADIUS server to do the following:

- Centralize authentication, authorization, and accounting for a heterogeneous set of access servers.
- Deploy secure password authentication with Protected Extensible Authentication Protocol (PEAP) for 802.1X authenticating switches and wireless clients
- Use a Windows NT Server 4.0 domain, an Active Directory domain, or the local SAM as your user account database.
- Centralize both the configuration of remote access policies and connection logging for accounting when you are using the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 Routing and Remote Access service on multiple dial-up servers, VPN servers, or demand-dial routers.
- Outsource your dial-in, VPN, or wireless access to a service provider. The access servers use RADIUS to authenticate and authorize connections that are made by members of your organization.

**Note**

- The configuration of IAS servers running on products in the Windows 2000 Server family can be imported into products in the Windows Server 2003 family by using netsh AAAA commands. The reverse, however, is not possible. To centrally manage IAS on both Windows 2000 and Windows Server 2003, you can first make configuration changes on IAS servers running Windows 2000, and then import the new server configuration into IAS servers running Windows Server 2003.

When you configure IAS as a RADIUS server, Access-Request messages are processed locally. In this process, IAS does the following:

1. **Validates the RADIUS packet**. The incoming Access-Request message is validated for source IP address, the digital signature, valid attributes, and so on. If the RADIUS message is not valid, an event is logged in the system event log and the RADIUS Access-Request message is discarded. An Access-Reject message is not sent.

**Checks for auto-reject**. Auto Reject, also called **Ping User-Name**, is used to send an immediate Access-Reject message when the User-Name attribute in the Access-Request message matches a specific value.

Some RADIUS proxy servers and network access servers periodically send authentication and accounting requests, or ping requests, to verify that the IAS server is present on the network. These ping requests include fictional user names. When IAS processes these requests, the event and accounting logs become filled with access reject records, making it difficult to keep track of valid records. When you configure a registry entry for **Ping User-Name**, IAS matches the registry entry value against the value of the **User-Name** attribute in ping requests that other servers make. If the registry entry and the user name value match, IAS automatically rejects the request but does not create an event or accounting log entry.

2. **Performs connection request policy evaluation**. If no connection request policies are matched, an event is logged in the system event log and the RADIUS Access-Request message is discarded.

3. **Applies realm stripping rules**. IAS determines or defines the domain name and the user identity for the Access-Request message. If the User-Name attribute in the Access-Request message is not the Auto Reject name, then the user identity is determined. User identity is how IAS identifies the user for the purposes of authentication and authorization. Typically, the user identity is the string value of the User-Name RADIUS attribute. If the User-Name attribute is not present, the user identity is set to the Guest account or the account specified by the Default User Identity registry value.

   IAS can use any RADIUS attribute to identify the user. The RADIUS attribute that IAS uses to identify the user is configurable by setting the User Identity Attribute registry setting.

4. **Determines authentication server**. IAS determines whether to authenticate locally or forward to a remote RADIUS server group (In this example, the message is authenticated locally and is not forwarded.)

5. **Performs name cracking**. Name cracking is the resolution of the user identity to a user account using user principal names (UPNs), Lightweight Directory Access Protocol (LDAP), distinguished names (DNA), Canonical Names, and so on. If a user principal name is encountered by IAS, IAS performs a query to the Active Directory Global Catalog in an attempt to resolve the name. To speed up this process, a copy of the Global Catalog must be located on a domain controller within the same site as the IAS server.

   When the user identity does not contain a domain name, IAS supplies a domain name. By default, the IAS-supplied domain name is the domain for which the IAS server is a member. You can specify the IAS-supplied domain by means of the DefaultDomain registry setting.

6. **Checks for authentication plug-ins**. The existence of authentication plug-ins is checked. Authentication plug-ins are optional components created using the IAS software development kit (SDK); each plug-in can return Accept, Reject, or Continue. If an authentication plug-in returns an Accept, the user is considered to be authenticated and the account is then validated. If the authentication plug-in returns a Reject, an Access-Reject message is sent, and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings. If the authentication plug-in returns a Continue, the next plug-in is checked. If there are no more plug-ins, the user still needs to be authenticated.The authentication plug-in can also return RADIUS attributes to be included in the Access-Accept message.

7. **Checks for remote access account lockout.** After the authentication plug-ins are checked, the registry on the IAS server is read for remote access account lockout entry for the user account. If the account is locked out by way of the remote access account lockout, IAS sends an Access-Reject message back to the NAS and logs an authentication event. For more information, see "IAS Tools and Settings."

8. **Checks for PAP, CHAP, MS-CHAP** . If Password Authentication Protocol (PAP), CHAP, Microsoft Challenge Handshake Authentication Protocol (MS-CHAP v1), or MS-CHAP v2 are used to authenticate the remote access client, IAS consults an authentication sub-module based on the authentication protocol to perform the authentication. The user's credentials (user name and password) are authenticated against the user name and password of the accounts database (either a domain or the local accounts database), and the group membership of the user account is determined. The exact method of authentication varies depending on the authentication protocol.

If the authentication of the credentials is not successful, an Access-Reject message is sent and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings.

If either EAP or unauthenticated access is being used, then the user authentication process is bypassed. EAP authentication takes place later in this process. For unauthenticated access, no user authentication is performed.

9. **Validates user account**. Based on the user or computer account determined through name cracking, the user account is validated to discover whether the account is locked out (which is not the same as remote access account lockout), whether the account is disabled, and whether the user account's password has expired. If the user account is not valid, an Access-Reject message is sent and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings.

10. **Performs remote access policy evaluation**. Remote access policies configured on the IAS server are evaluated to find a policy that matches the parameters of the connection. If a matching policy is not found, an Access-Reject message is sent and an event is logged.

11. **Checks user properties and profile properties**. If the Ignore-User-Dialin-Properties attribute is set to 0, the dial-in properties of the user account and the profile properties of the matching remote access policy are evaluated against the parameters of the connection attempt to ensure that the connection attempt is allowed. If the Ignore-User-Dialin-Properties attribute is set to 1, the profile properties from the matching policy become the set of properties for the connection.

    If the connection attempt is not allowed, an Access-Reject message is sent, and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings.

12. **Checks for EAP authentication**. If Extensible Authentication Protocol (EAP) is the authentication protocol used for the connection attempt, EAP authentication takes place. The initial negotiation for EAP consists of selecting EAP as the authentication protocol and negotiating an EAP type. Based on the EAP type, the profile settings for the matching policy are checked to ensure that the EAP type is allowed. If the EAP type is not allowed with the profile settings, an Access-Reject message is sent and the authentication failure event is logged in the system event log or the IAS accounting log, depending on the configured logging settings.

    If the EAP type is allowed with the profile settings, EAP authentication for the EAP type occurs. IAS sends an EAP challenge to the NAS requesting it to start EAP negotiation. Communications between EAP modules on a RADIUS client and server are tunneled using the RADIUS protocol. After negotiation is complete, an EAP provider can return attributes that are sent back to the NAS in the Access-Accept message. If EAP authentication fails, an Access-Reject message is sent, and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings.

13. **Checks for authorization plug-ins**. The existence of authorization plug-ins is checked. Authorization plug-ins are optional components you can create using the IAS SDK. Each plug-in can return either Reject or Continue. If the authorization plug-in returns a Reject, an Access-Reject message is sent and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings. If the authorization plug-in returns Continue, the next plug-in is checked. If there are no more plug-ins, the user is considered to be authorized.

    The authorization plug-in can also return RADIUS attributes to be included in the Access-Accept message.

14. **Sends an Access-Accept**. If the dial-in properties of the user account, the profile properties of the matching remote access policy, and the conditions imposed by authorization plug-ins allow the connection attempt, an Access-Accept message is sent back to the NAS. Included with the Access-Accept message is the set of RADIUS attributes for the restrictions on the connection. In addition, an authentication success event is logged in either the system event log or the IAS authentication log, depending on the configured logging settings.

The authentication and authorization process for the Routing and Remote Access service, when configured for Windows authentication, requires steps 6 through 15 of this process. The authentication and authorization success and failure are the return values of functions called by the Routing and Remote Access service. Local event or authentication logging depends on the configured logging settings of the Routing and Remote Access service.

**Using IAS as a RADIUS Proxy**
When you use IAS as a RADIUS proxy, Access-Request messages are forwarded to a remote RADIUS server for authentication and authorization.

If you configure the IAS proxy to map user accounts between a remote RADIUS server and the local user accounts database, Access-Request messages are authenticated by the remote server and are authorized with the user or computer account dial-in properties and remote access policies that are configured on the IAS proxy.

IAS can also return an Access-Accept message without authenticating and authorizing the Access-Request message.

Accounting-Request messages are logged in a local log file (based on remote access logging settings) and forwarded to another RADIUS server for accounting.

You can use IAS as a RADIUS proxy when the following conditions exist:

- You want to provide authentication and authorization for user accounts that are not members of either the domain in which the IAS server is a member or another domain that has a two-way trust with the domain in which the IAS server is a member. This includes accounts in untrusted domains, one-way trusted domains, and other untrusted forests. Instead of configuring your access servers to send their connection requests to an IAS RADIUS server, you can configure them to send their connection requests to an IAS RADIUS proxy. The IAS RADIUS proxy uses the realm name portion of the user name and forwards the request to an IAS server in the correct domain or forest. Connection attempts for user accounts in one domain or forest can be authenticated for NASs in another domain or forest.

- IAS supports authentication across forests without a RADIUS proxy when the two forests contain only domains that consist of domain controllers running Microsoft Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. The forest functional level must be Windows Server 2003, and there must be a two-way trust relationship between forests. If you use Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) with certificates as your authentication method, however, you must use a RADIUS proxy for authentication across forests that consist of Windows Server 2003 domains.

- Your user accounts database is not a Windows user accounts database, and it is not compliant with LDAP. Additionally, you are using a non-Microsoft RADIUS server. In this case, the IAS proxy server forwards connection requests that match a specified realm name to a non-Microsoft RADIUS server in a remote RADIUS server group. The remote RADIUS server performs authentication and authorization against the non-Microsoft user accounts database. Examples of non-Microsoft user accounts databases include Novell Directory Services (NDS) and SQL databases.

- You want to process a large number of connection requests. In this case, instead of configuring your RADIUS clients to attempt to balance their connection and accounting requests across multiple RADIUS servers, you can configure them to send their connection and accounting requests to an IAS RADIUS proxy. The IAS RADIUS proxy dynamically balances the load of connection and accounting requests across multiple RADIUS servers and increases the processing of large numbers of RADIUS clients and authentications per second.

- You want to provide RADIUS authentication and authorization for outsourced service providers and minimize intranet firewall configuration. An intranet firewall is between your perimeter network (the network between your intranet and the Internet) and intranet. By placing an IAS server on your perimeter network, the firewall between your perimeter network and intranet must allow traffic to flow between the IAS server and multiple domain controllers. When replacing the IAS server with an IAS proxy, the firewall must allow only RADIUS traffic to flow between the IAS proxy and one or multiple IAS servers within your intranet.

- You are a service provider that offers outsourced dial, VPN, or wireless network access services to multiple customers. Your NASs send connection requests to the IAS RADIUS proxy. Based on the realm portion of the user name in the connection request, the IAS RADIUS proxy forwards the connection request to a RADIUS server that is maintained by the customer and can authenticate and authorize the connection attempt. .

When you configure IAS as a RADIUS proxy, Access-Request messages are forwarded to a remote RADIUS server and are processed remotely.

When you map external user accounts to a local Windows user account, authentication is performed remotely and authorization is performed locally. In this process, IAS does the following tasks:

1. **Validates the RADIUS packet**. The incoming Access-Request message is validated for source IP address, the digital

signature, valid attributes, and so on. If the RADIUS message is not valid, an event is logged in the system event log and the RADIUS Access-Request message is discarded. An Access-Reject message is not sent.

2. **Checks for auto-reject**. Auto Reject, also called **Ping User-Name,** is used to send an immediate Access-Reject message when the User-Name attribute in the Access-Request message matches a specific value.

   Some RADIUS proxy servers and NASs periodically send authentication and accounting requests, also called ping requests, to verify that the IAS server is present on the network. These ping requests include fictional user names. When IAS processes these requests, the event and accounting logs become filled with access reject records, making it more difficult to keep track of valid records. When you configure a registry entry for **Ping User-Name**, IAS matches the registry entry value against the value of the User-Name attribute in ping requests made by other servers. If the registry entry and the user name value match, IAS rejects the request but does not create an event or accounting log entry.

3. **Performs connection request policy evaluation**. Connection request policies configured on the IAS server are evaluated to find a policy that matches the parameters of the connection. If a policy is found that matches the connection request, the name of the remote RADIUS server group to which the Access-Request should be forwarded is injected into the message.

   If a matching policy is not found, an Access-Reject message is sent and an event is logged. For more information about connection request policies and policy evaluation logic, see "Connection Request Policies" later in this section.

4. **Applies realm stripping rules**. IAS determines or defines the domain name and the user identity for the Access-Request message. If the User-Name attribute in the Access-Request message is not the Auto Reject name, then the user identity is determined. User identity is the value that IAS uses to identify the user for the purposes of authentication and authorization. Typically, the user identity is the string value of the User-Name RADIUS attribute. If the User-Name attribute is not present, the user identity is set to the Guest account or the account specified by the **Default User Identity r**egistry value.

   IAS can use any RADIUS attribute to identify the user. The RADIUS attribute that IAS uses to identify the user is configurable by setting the User Identity Attribute registry setting.

5. **Determines whether to authenticate locally or forward to a remote RADIUS server group**. Based on the realm portion of the user name, the IAS proxy determines whether to forward the message or continue to process it locally.

6. **Check for authentication plug-ins**. The existence of authentication plug-ins is checked. Authentication plug-ins are optional components you can create with the IAS SDK; each plug-in can return Accept, Reject, or Continue. If an authentication plug-in returns an Accept, the user is considered to be authenticated and the account is then validated. If the authentication plug-in returns a Reject, an Access-Reject message is sent and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings. If the authentication plug-in returns a Continue, the next plug-in is checked. If there are no more plug-ins, the user still needs to be authenticated.

   The authentication plug-in can also return RADIUS attributes to be included in the Access-Accept message.

7. **Forward the Access-Request to a remote RADIUS server group**. IAS creates a RADIUS message, which includes the attributes received from the RADIUS client, with the exception of the proxy state attribute. IAS forwards the message to the remote RADIUS server group specified in the request. When a valid reply is received from the remote server, attributes from the RADIUS packet are merged with the attributes in the request. Attributes returned by the remote RADIUS server are not added if the same attribute has already been added by the connection request policy. In other words, local settings override remote settings.

8. **Maps a remote user account to a local Windows account**. If the IAS proxy server is configured to map an external user account to a local Windows account, and if the remote RADIUS server has performed authentication and returned an Access-Accept, mapping of the accounts is performed.

   If mapping is not configured, IAS proceeds to number 12 in this process, "Check for authorization plug-ins." For more information about account mapping, see "Mapping Network Authentication and Authorization" later in this section.

9. **Validates user account**. Based on the user or computer account that is determined by means of name cracking, the user account is validated to check whether the account is locked out (which is not the same as remote access account lockout), whether the account is disabled, and whether the user account's password has expired. If the user account is not valid, an Access-Reject message is sent and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings.

10. **Performs remote access policy evaluation**. Remote access policies configured on the IAS server are evaluated to find a policy that matches the parameters of the connection. If a matching policy is not found, an Access-Reject message is sent and an event is logged.

11. **Checks user properties and profile properties**. If the Ignore-User-Dialin-Properties attribute is set to 0, the dial-in properties of the user account and the profile properties of the matching remote access policy are evaluated against the parameters of the connection attempt. This is done to ensure that the connection attempt is allowed. If the Ignore-User-Dialin-Properties attribute is set to 1, the profile properties from the matching policy become the set of properties for the connection.

    If the connection attempt is not allowed, an Access-Reject message is sent, and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings.

12. **Checks for authorization plug-ins**. Authorization plug-ins are optional components created using the IAS SDK. Each plug-in can return either Reject or Continue. If an authorization plug-in is installed and it returns a Reject, an Access-Reject message is sent and the authentication failure event is logged in the system event log or the IAS authentication log, depending on the configured logging settings. If the authorization plug-in returns Continue, the next plug-in is checked. If there are no more plug-ins, the user is authorized.

    An authorization plug-in can also return RADIUS attributes to be included in the Access-Accept message.

13. **Sends an Access-Accept**. If the dial-in properties of the user account, the profile properties of the matching remote access policy, and the conditions imposed by authorization plug-ins allow the connection attempt, an Access-Accept message is sent back to the NAS. Included with the Access-Accept message is the set of RADIUS attributes for the restrictions on the connection. In addition, an authentication success event is logged in the system event log or the IAS authentication log, depending on the configured logging settings.

**IAS as a RADIUS Server**

IAS can be used as a RADIUS server to perform AAAA for connection requests at RADIUS clients. A RADIUS client can be either a NAS or a RADIUS proxy. When IAS is used as a RADIUS server, it provides the following:

- A central authentication and authorization service for all access requests that are sent by RADIUS clients.

  IAS uses a Microsoft Windows NT Server 4.0 domain, an Active Directory domain, or the local SAM to authenticate user credentials for a connection attempt. IAS uses the dial-in properties of the user account and remote access policies to authorize a connection.

- A central accounting recording service for all accounting requests that RADIUS clients send.

- Accounting requests are stored in a local log file or a SQL server database for analysis.

When IAS is used as a RADIUS server, RADIUS messages provide authentication, authorization, and accounting for network access connections in the following order:

1. Access servers, such as dial-up NASs, VPN servers, and wireless APs, receive connection requests from access clients.

2. The access server that is configured to use RADIUS as the authentication, authorization, and accounting protocol creates an Access-Request message and sends it to the IAS server.

3. The IAS server evaluates the Access-Request message.

4. If required, the IAS server sends an Access-Challenge message to the access server. The access server processes the challenge and sends an updated Access-Request to the IAS server.

5. IAS checks the user name for a domain name. If a domain name is specified and IAS is configured to access the user accounts database in the specified domain, IAS proceeds with processing the connection request. When the user name does not contain a domain name, IAS supplies one. By default, the IAS-supplied domain name is the domain of which the IAS server is a member. IAS resolves a user name without a specified domain name by using the following sequence:

    a. IAS checks the default domain registry key. If a default domain is specified, IAS authenticates the user against the domain specified in the registry key.

    b. If the IAS server is a member of a domain, IAS authenticates the user against the domain to which it is joined.

    c. If the IAS server is not a member of a domain, IAS authenticates the user against the local SAM database.

6. IAS checks the user identity in the User-Name attribute. User identity is what IAS uses to identify the user for the purposes of authentication and authorization. Typically, the user identity is the string value of the User-Name RADIUS attribute. If the User-Name attribute is not present, the user identity is set to the Guest account or the account specified by the **Default User Identity** registry value.

7. The dial-in properties of the user account are obtained by using a secure connection to a domain controller.

8. The connection attempt is authorized with both the dial-in properties of the user account and remote access policies. If the connection attempt is both authenticated and authorized, the IAS server sends an Access-Accept message to the access server.

9. If the connection attempt is either not authenticated or not authorized, or the connection request parameters do not match a connection request policy and a remote access policy, the IAS server sends an Access-Reject message to the access server.

10. The access server completes the connection process with the access client and sends an Accounting-Request message to the IAS server where the message is logged in IAS format, database-compatible format, or to a SQL Server.

11. The IAS server sends an Accounting-Response to the access server.

The access server also sends Accounting-Request messages when the connection is being established, when the access client connection is closed, and when the access server is started and stopped.

**IAS Authentication**

The authentication of access clients is an important security concern. Authentication methods typically use an authentication protocol that is negotiated during the connection establishment process. IAS also supports unauthenticated connections.

Each authentication method has advantages and disadvantages in terms of security, usability, and breadth of support. The authentication method used is determined by the configuration of the access client, the access server, and the RADIUS server. Consult your access server documentation to determine which authentication methods are supported.

You can configure IAS to accept multiple authentication methods. You can also configure your access servers to attempt to negotiate a connection by using the most secure protocol first, and then the next most secure, and so on down to the least secure. For example, the Routing and Remote Access service tries to negotiate a connection using EAP first, then MS-CHAP v2, then MS-CHAP v1 , then CHAP, and then PAP. When EAP is chosen as the authentication method, the negotiation of the EAP type occurs between the access client and the IAS server.

In addition, you can use remote access policies for various authentication methods, depending on the type of client being authenticated. For example, you can create two remote access policies, one for VPN clients and one for wireless clients—each of which uses a different authentication method. The remote access policy for VPN clients can be configured to use EAP-TLS with smart cards or certificates as the authentication method and authentication type, while the remote access policy for wireless clients can be configured to use PEAP-MS-CHAP v2, which provides secure password authentication.

**Authentication process**

A NAS functions as a RADIUS client of an IAS server. The NAS can be a dial-up server, VPN server, wireless AP, or an

authenticating switch. When a user or computer attempts to log on to a network through a NAS, the access client sends identity information, such as account credentials, to the access server. The access server sends the credentials to the RADIUS server.

When the NAS passes the account credentials to IAS, IAS authenticates the user or computer by sending the supplied credentials to a domain controller that compares the credentials with the credentials of an account in the accounts database. IAS also verifies that the account is authorized to access all or part of the network by examining properties of the account in Active Directory and by using remote access policies configured on the IAS server.

If the user is authenticated and authorized, IAS returns necessary configuration information to the RADIUS client so that it can deliver service to the user.

When a user or computer attempts to connect to a network, the authentication request is processed as follows:

1. The NAS tries to negotiate a connection with the remote access client by using the most secure protocol first, then the next least secure protocol, and so on, down to the least secure protocol (for example, the NAS tries to negotiate a connection by using EAP, MS-CHAP v2, then MS-CHAP , then CHAP, and finally PAP).

2. The NAS forwards the authentication request to an IAS server in the form of a RADIUS Access-Request message.

3. The IAS server validates the RADIUS Access-Request message.

4. If digital signatures are enabled and the verification of the digital signature fails, the IAS server silently discards the message. When the NAS does not get a response within its time-out period, it retries and then disconnects the user. If the IAS server cannot connect to the domain controller or cannot find the domain controller the user belongs to, it silently discards the message. This allows a RADIUS proxy to retransmit the request to the backup IAS server, which would then attempt to authenticate the user against the database of the domain.

5. If digital signatures are enabled and the verification of the digital signature is successful, the IAS server queries the domain controller, which validates the user credentials.

6. If the user credentials are authentic, the IAS server evaluates the connection attempt against the configured remote access policies and the dial-in properties of the user's account to decide whether to authorize the request. If the connection attempt matches the conditions of at least one policy and the user account dial-in properties, remote access policy properties, and the remote access policy profile properties authorize the connection, IAS sends a RADIUS Access-Accept message to the NAS that sent the Access-Request message. The Access-Accept message authorizes the connection but also contains connection parameters based on the remote access policy profile settings and the dial-in properties of the user account. The NAS interprets this authorization data to determine the connection parameters that the server has authorized.

If the user is not authentic or the user's attempt to connect does not match conditions in at least one policy, or matches conditions in a policy that denies access, IAS sends a RADIUS Access-Reject message to the NAS, and the NAS disconnects the user.

**Ping User-Name value**

The **Ping User-Name** registry entry specifies the fictional user name (or a user name pattern, with variables, that matches the fictional user name) sent by RADIUS proxy servers and NASs to IAS servers to test whether the IAS server is available. When IAS receives ping requests that match the **Ping User-Name** registry entry value, IAS rejects the authentication requests without processing the request. IAS does not record transactions involving the fictional user name in any log files, which makes the event log easier to interpret.

**Adding Ping User-Name to the registry**

When you add **Ping User-Name** to the registry, you must supply values for **Name**, **Type**, and **Data**, as shown in the following table, "**Ping User-Name** Values."

**Ping User-Name Values**

| Name | Type | |
|---|---|---|
| ping user-name | REG_SZ | *User name* |

To indicate more than one user name for a **Ping User-Name** value, enter a name pattern, such as a Domain Name System

(DNS) name including wildcard characters, in **Data**.

**IAS and Tunneling**
The entire process of packet encapsulation, transmission, and de-encapsulation is called tunneling. Tunneling is a method of using an internetwork infrastructure of one protocol to transfer a payload. Sometimes, the payload consists of the frames (or packets) of another protocol. Instead of being sent as the originating host produces it, the frame is encapsulated with an additional header. The additional header provides routing information so that the encapsulated payload can traverse an intermediate internetwork (also known as a transit internetwork). The encapsulated packets are then routed between tunnel endpoints over the transit internetwork. After the encapsulated payload packets reach their destination on the transit internetwork, the frame is de-encapsulated and forwarded to its final destination.

The logical path in the transit internetwork through which the encapsulated packets travel is called a tunnel.

With remote access VPN connections, there are two types of tunneling: voluntary and compulsory

**Voluntary tunneling**
When a client computer is configured with the appropriate tunneling protocol, the user or client computer can issue a VPN request to create and configure a voluntary tunnel from the client computer to the VPN server. In this case, the user's computer is a tunnel endpoint that acts as the tunnel client, and the VPN server is the tunnel server.

A simple example of voluntary tunneling occurs when a user with a dial-up connection to the Internet through an Internet service provider (ISP) connects to an organization VPN server that is also connected to the Internet. In this scenario, the user has two separate accounts, each account having its own set of credentials:

- A set of credentials for an account with an ISP
- A set of credentials for an account with the organization

To connect to the organization VPN server, the client computer uses a dial-up modem to attempt to connect to the ISP dial-up server. The user supplies his ISP account credentials, is authenticated by the ISP RADIUS server against the ISP user accounts database, and is authorized by the ISP RADIUS server to connect to the ISP and the Internet.

After the dial-up connection is established and the client computer is on the Internet, the user initiates a VPN connection to his organization VPN server. The VPN server is configured as a RADIUS client to an IAS server at the organization. The user supplies organization account credentials for the connection attempt, and the organization IAS server uses the credentials to authenticate and authorize the user against an Active Directory user accounts database. After the user is authenticated and authorized, the VPN tunnel is established and the user can access organization resources.

In another example, IAS is used in an outsourced bulk dial scenario for voluntary tunneling. In the outsourced bulk dial scenario, an ISP is providing both dial-up and non-dedicated broadband digital subscriber line (DSL) Internet access for all the employees of an organization. Rather than providing the ISP with a copy of its user accounts database, the organization and the ISP use IAS as a RADIUS proxy so that users, when connecting to the ISP, can be authenticated and authorized using the organization RADIUS servers and Active Directory user accounts database on the organization network. This provides strong security and prevents exposure of sensitive information contained in the user accounts database to the ISP and its employees.

When the organization employee attempts to connect to the organization VPN server, the following process occurs:

1. The client computer establishes the dial-up or DSL connection to the ISP NAS.
2. Based on the connection parameters, the NAS sends an Access-Request message to an IAS server that is configured as a RADIUS proxy.
3. The RADIUS proxy processes the Access-Request message and determines where to send it based on the realm portion of the User-Name attribute. The RADIUS proxy forwards the Access-Request message to the organization IAS server, which is accessible on the Internet through a firewall configured to allow traffic on the default RADIUS UDP ports of 1812 (authentication) and 1813 (accounting).
4. The organization IAS server authenticates and authorizes the connection attempt of the client computer and sends an Access-Accept message back to the RADIUS proxy.
5. The RADIUS proxy forwards the Access-Accept message to the ISP NAS and the ISP NAS connects the client computer to the Internet.

6. After it is on the Internet, the client computer initiates a VPN tunnel connection with the organization VPN tunnel server on the Internet.

7. Based on the tunnel connection parameters, the tunnel server sends an Access-Request message to the organization IAS server.

8. The organization IAS server authenticates and authorizes the connection attempt of the tunnel client and sends an Access-Accept message back to the tunnel server.

9. The tunnel server completes the tunnel creation and the tunnel client can now send packets to the organization intranet through the tunnel across the Internet.

The figure, "Voluntary Tunnel Created by a Dial-Up User," illustrates the establishment of a voluntary VPN tunnel by a dial-up user.

**Voluntary Tunnel Created by a Dial-Up User**




Voluntary tunneling is not different from other types of network access, and IAS can be used for authentication, authorization, and accounting when the tunnel server is configured as a RADIUS client to the IAS server.

**Compulsory tunneling**

Compulsory tunneling is the creation of a secure tunnel by another computer or network device on behalf of the client computer. Compulsory tunnels are configured and created automatically for the user without their knowledge or intervention. With a compulsory tunnel, the user's computer is not a tunnel endpoint. Another device between the user's computer and the tunnel server is the tunnel endpoint, acting as the tunnel client. The dial-up access server dialed by the client computer is the tunnel endpoint, acting as the tunnel client.

A number of vendors that sell dial-up access servers have implemented the ability to create a tunnel on behalf of a dial-up client. The computer or network device providing the tunnel for the client computer is known as a Front End Processor (FEP) in PPTP, a Layer Two Tunneling Protocol (L2TP), an Access Concentrator in L2TP, or an IP Security Gateway in Internet Protocol Security (IPSec).

In this section, "How IAS Works," the acronym FEP describes this functionality, regardless of the tunneling protocol. To carry out its function, the FEP must have the appropriate tunneling protocol installed and must be capable of establishing the tunnel when the client computer attempts a connection.

A corporation can contract with an ISP to deploy a nationwide set of FEPs. These FEPs can establish tunnels across the Internet to a tunnel server connected to the private network of the corporation, thereby consolidating calls from geographically diverse locations into a single Internet connection at the corporate network.

The following figure, "Compulsory Tunnel Created by a Tunneling-Enabled NAS," shows the client computer placing a dial-up call to a tunneling-enabled NAS at the ISP to authenticate against an IAS server on the other side of the tunnel.

**Compulsory Tunnel Created by a Tunneling-Enabled NAS**




The preceding figure, "Compulsory Tunnel Created by a Tunneling-Enabled NAS," illustrates how IAS is used in an outsourced bulk dial scenario for compulsory tunneling.

In the outsourced bulk-dial scenario, a dial-up client establishes a dial-up connection to an ISP that is providing tunneled access across the Internet for all the employees of an organization. Based on the dial-up connection parameters, the ISP NAS sends an Access-Request message to an IAS server. The ISP IAS server authorizes, but does not authenticate, the tunnel connection. The ISP IAS server sends an Access-Accept message to the ISP NAS with a series of VPN tunnel attributes. The ISP NAS then acts as a tunnel client and creates a VPN tunnel to the VPN tunnel server of the organization that faces the Internet.

**Note**

- Typically, IAS provides both authentication and authorization. However, it is common for the ISP IAS server to provide only authorization when the dial-in user is authenticated by the organization IAS server.

The ISP NAS then sends a PPP message to the dial-up client to restart the authentication process so that the dial-up user can be authenticated by the organization IAS server through the tunnel. The dial-up client sends the authentication information to the ISP NAS, which encapsulates the user account credentials and sends them through the tunnel to the organization tunnel server.

After the user account credentials are received by the tunnel server, the tunnel server sends an Access-Request message to the organization IAS server. The organization IAS server authenticates and authorizes the connection of the dial-up client to the tunnel server and sends an Access-Accept message to the tunnel server. The tunnel server then completes the connection to the dial-up client.

All data that is sent by the dial-up client is automatically sent through the tunnel to the tunnel server by the ISP NAS.

This configuration is known as compulsory tunneling because the client is compelled to use the tunnel created by the FEP. After the initial connection is made, all network traffic to and from the client is automatically sent through the tunnel.

In addition, you can configure IAS to instruct an FEP to tunnel different remote access clients to different tunnel servers.

Unlike the separate tunnels created for each voluntary client, a compulsory tunnel between the FEP and organization VPN server can be shared by multiple dial-up clients. When a second client dials into the same access server (the FEP) to reach a destination for which a tunnel already exists, the data traffic for the new client is carried over the existing VPN tunnel.

The following RADIUS Attributes are used to carry the tunneling information from the IAS server to the NAS.

Used in authorization only:

- Tunnel-Pvt-Group-ID
- Tunnel-Assignment-ID
- Tunnel-Preference
- Tunnel-Password
- Tunnel-Client-Auth-ID
- Tunnel-Server-Auth-ID

Used in authorization and accounting:

- Tunnel-Type (such as PPTP and L2TP)
- Tunnel-Medium-Type (X.25, ATM, Frame Relay, IP, and so on)
- Tunnel-Client-Endpt
- Tunnel-Server-Endpt

Used for accounting only:

- Acct-Tunnel-Connection

The Windows Server 2003 or Windows 2000 Routing and Remote Access service cannot be used as a FEP for compulsory tunneling.

**Authentication Methods**
The RADIUS protocol supports several PPP authentication protocols, each with its advantages and disadvantages in terms of security, usability, and breadth of support. The protocol used is determined by the configuration of the NAS device. For information about support for the various authentication protocols, see your NAS documentation, or consult your ISP or other vendor support.

The following sections focus on the advantages and disadvantages of the authentication protocols that are currently supported by IAS, and discuss how to deploy certificate-based authentication methods. The information is also useful in configuring a particular authentication method for network access authentication.

**PAP**
Password Authentication Protocol (PAP) sends a password as a plaintext string from the user's computer to the NAS device. When the NAS forwards the password as the User-Password attribute in the Access-Request message, it is encrypted using the RADIUS shared secret as an encryption key. PAP is the most flexible protocol because passing a plaintext password to the authentication server enables that server to compare the password with nearly any storage format. For example, UNIX passwords are stored as one-way encrypted strings that cannot be decrypted. PAP passwords can be compared to these strings by reproducing the one-way encryption method.

**Note**

- The use of this authentication method is not recommended for security reasons.

When PAP is enabled as an authentication protocol, user passwords are sent from a client to a NAS in plaintext form. The NAS encrypts the password, using the shared secret, and then sends it in an Access-Request message. Because a RADIUS proxy must encrypt the PAP password by using the shared secret of its forwarding RADIUS server, a RADIUS proxy must decrypt the PAP password, using the shared secret between the RADIUS proxy and the NAS. Because it uses a plaintext version of the password, PAP has a number of security vulnerabilities. Although the RADIUS protocol encrypts the password, it is transmitted as plaintext across the dial-up connection.

PAP-based authentication is enabled:

- On the appropriate remote access policy. PAP is disabled by default.
- On a remote access client.
- As an authentication protocol on the remote access server. For the Routing and Remote Access service, PAP is disabled by default.

For information about a default setting on a particular NAS, see your NAS documentation.

If you set the **User passwords to be changed at the next attempt to log on** option, users must first log on using a local area network (LAN) connection and change the password before attempting to log on with a remote access connection using PAP. PAP and CHAP do not support changing passwords during the authentication process. If the password is not first changed using a LAN connection, logon fails. An alternative method for changing passwords is to temporarily log on using MS-CHAP to change the password.

**Note**

- A malicious user or process on a RADIUS proxy can record user names and passwords for PAP connections while the password is in its plaintext form. For this reason, the use of PAP is highly discouraged, especially for virtual private network connections.

**CHAP**

Challenge Handshake Authentication Protocol (CHAP) is designed to address the concern of passing passwords in plaintext. By using CHAP, the NAS sends a random number challenge to the user's computer. The challenge and the user's password is then hashed by using Message Digest 5 (MD5). The client computer then sends the hash as a response to the NAS challenge and the NAS forwards both the challenge and response in the RADIUS Access-Request message.

When the authenticating server receives the RADIUS message, it uses the challenge and the user's password to create its own version of the response. If the version of the server matches the response supplied by the user's computer, the access request is accepted.

CHAP responses cannot be reused because NAS devices send a unique challenge each time a client computer connects to them. Because the algorithm for calculating CHAP responses is well known, it is very important that passwords be carefully chosen and sufficiently long. CHAP passwords that are common words or names are vulnerable to dictionary attacks if they can be discovered by comparing responses to the CHAP challenge with every entry in a dictionary. Passwords that are not sufficiently long can be discovered by brute force by comparing the CHAP response to sequential trials until a match to the user's response is found.

Historically, CHAP is the most common dial-up authentication protocol used. When the server does not store the same password that was used to calculate the CHAP response, it cannot calculate an equivalent response. Because standard CHAP clients use the plaintext version of the password to create the CHAP challenge response, passwords must be stored in a reversibly encrypted form on the server to validate the CHAP response of the client.

Although the IAS server supports CHAP, a Windows NT 4.0–based domain controller cannot validate CHAP requests without support for storing reversibly encrypted passwords. This support is available in Windows Server 2003 and Windows 2000; in Windows NT 4.0, this support is available after installing Service Pack 3 or later.

CHAP-based authentication is enabled when you have done the following:

- Enabled CHAP on the appropriate remote access policy. CHAP is disabled by default.

- Enabled storage of a reversibly encrypted form of the user's password. For a stand-alone server, use computer Group Policy to enable storage of reversibly encrypted passwords for all users of the computer. For Active Directory domains, Group Policy at the domain or Organizational Unit (OU) level can be used.

- Forced a reset of users' passwords so that the new password is in a reversibly encrypted form. When you enable passwords to be stored in a reversibly encrypted form, the current passwords are in a non-reversibly encrypted form and are not automatically changed. You must either reset user passwords or set user passwords to be changed the next time you log on. After the password is changed, it is stored in a reversibly encrypted form.

- Enabled CHAP on the remote access client.

- Enabled CHAP as an authentication protocol on the remote access server. For the Routing and Remote Access service, CHAP is disabled by default.

For information about a default setting on a particular NAS, see your NAS documentation.

If you set the **User passwords to be changed at the next attempt to log on** option, the user must first log on using a LAN connection and change the password before attempting to log on with a remote access connection using CHAP. CHAP and PAP do not support changing passwords during the authentication process. If the password is not first changed by using a LAN connection, logon fails. As an alternative to this method for changing passwords, you can temporarily log on by using MS-CHAP to change the password.

**MS-CHAP v1**

MS-CHAP v1, is a variant of CHAP that does not require a plaintext version of the password on the authenticating server. In MS-CHAP v1, the challenge response is calculated with a hashed version of the password and the NAS challenge. This enables authentication over the Internet to a Windows Server 2003 or Windows 2000 domain controller.

**Note**

- The use of this authentication method is not recommended for security reasons.

MS-CHAP v1 passwords are stored more securely at the server but have the same vulnerabilities to dictionary and brute force attacks as CHAP. If you deploy MS-CHAP v1, use strong passwords for all user accounts. A strong password includes the following:

- Has at least seven characters.
- Does not contain your user name, real name, or company name.
- Does not contain a complete dictionary word.
- Is significantly different from previous passwords. Incremental passwords (Password1, Password2, Password3) are not strong.
- Contains characters from each of the following four groups: uppercase letters, lowercase letters, numerals, and symbols.

With MS-CHAP v1 in Windows Server 2003, users are prompted to change their password before it expires. In addition, you can configure MS-CHAP v1 on the profile of a remote access policy by using the **Allow user to change password after it has expired** check box to allow or prevent users from changing a password *after* it has expired.

**Note**

- If the **Allow user to change password after it has expired** check box is disabled, and a user's password expires, the user is denied access and cannot change the password to regain network access.

MS-CHAP v1 is enabled as follows :

- As an authentication protocol on the remote access server. MS-CHAP v 1 is not enabled by default on the Routing and Remote Access service. For information about default settings on other NASs, see your NAS documentation.
- On the appropriate remote access policy. MS-CHAP v 1 is not enabled by default.
- On a remote access client.

By default, MS-CHAP v1 for Windows Server 2003 does not support LAN Manager authentication. This is a change from Windows 2000.

If you want to enable the use of LAN Manager authentication with MS-CHAP v1 for older operating systems, such as Microsoft Windows NT 3.5*x* and Microsoft Windows 95, you must change the value of **Allow LM Authentication** in the Windows registry.

**MS-CHAP v2**

Microsoft CHAP v2 (MS-CHAP v2) provides mutual authentication, stronger initial data encryption keys, and different encryption keys for sending and receiving. For VPN connections, Windows Server 2003 offers MS-CHAP v 2 before offering MS-CHAP v1. Windows clients that have been updated with the most recent service pack accept MS-CHAP v2 when it is offered.

MS-CHAP v2 is a one-way encrypted password, mutual authentication process that works as follows:

1. The remote access server sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string.
2. The remote access client sends a response that contains the following:
    - The user name.

- An arbitrary peer challenge string.
- A one-way encryption of the received challenge string, the peer challenge string, the session identifier, and the user's password.

3. The remote access server checks the response from the client and sends back a response containing:
   - An indication of the success or failure of the connection attempt.
   - An authenticated response based on the sent challenge string, the peer challenge string, the encrypted response of the client, and the encrypted version of the user's password.
4. The remote access client verifies the authentication response and, if correct, uses the connection. If the authentication response is not correct, the remote access client terminates the connection.
5. If a user authenticates by using MS-CHAP v2 and attempts to use an expired password, MS-CHAP prompts the user to change the password while connecting to the server. Other authentication protocols do not support this feature effectively locking out the user who used the expired password. New to Windows Server 2003 is the ability to prevent a user from changing their expired password using MS-CHAP v2 by clearing the **Allow user to change password after it has expired** check box, under the **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** check box on the **Authentication** tab on the profile of a remote access policy.

MS-CHAP v2 is enabled as follows:

- As an authentication protocol on the remote access server. MS-CHAP v2 is enabled by default on the Routing and Remote Access service. For information about default settings on other NASs, see your NAS documentation.
- On the appropriate remote access policy. MS-CHAP v2 is enabled by default.
- On the remote access client.

For information about default settings on other NASs, see your NAS documentation.

**Note**

- Note Windows 95 and Microsoft Windows 98 support MS-CHAP v2 only for VPN connections. Windows 95 and Windows 98 do not support MS-CHAP v2 for dial-up connections.

**EAP**

Extensible Authentication Protocol (EAP) is an extension to PPP that works with dial-up, PPTP, L2TP, and 802.1X authentication clients. EAP allows the addition of new authentication methods known as EAP types. Both the access client and the remote access server must support the same EAP type for successful authentication to occur.

Windows Server 2003 includes an EAP infrastructure and two EAP types, EAP-MD5 and EAP-TLS. The IAS proxy implementation in Windows Server 2003 can forward EAP messages to a RADIUS server.

When you configure EAP as an authentication method in a remote access policy on an IAS server, NASs configured with RADIUS authentication can pass EAP messages from access clients to the IAS server or proxy using the RADIUS protocol. The EAP messages sent between the access client and access server are encapsulated and formatted as RADIUS messages between the remote access server and the RADIUS server.

One advantage of EAP is that EAP types do not need to be installed at each NAS, only at the RADIUS server. In a typical use of EAP, a NAS is configured to use EAP and to use RADIUS as its authentication provider. When a connection is made, the network access client negotiates the use of EAP with the NAS. When the client sends an EAP message to the NAS, the NAS encapsulates the EAP message as a RADIUS message and sends it to its configured RADIUS server. The RADIUS server processes the message and sends a RADIUS-encapsulated EAP message back to the NAS. The NAS then forwards the EAP message to the

network access client. In this configuration, the NAS is only a pass-through device. All processing of EAP messages occurs at the network access client and the RADIUS server.

**EAP-MD5**
MD5 with CHAP (EAP-MD5) is a required EAP type that uses the same challenge-handshake protocol as PPP-based CHAP, but the challenges and responses are sent as EAP messages. A typical use for EAP-MD5 is to authenticate the credentials of remote access clients by using user name and password security systems. You can use EAP-MD5 to test EAP interoperability.

**EAP-TLS**
Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) is an EAP type that is used in certificate-based security environments. To use smart cards for remote access authentication, you can deploy a PKI with Certificate Services in Windows Server 2003 or with a non-Microsoft certification authority, and then configure the EAP-TLS authentication method in remote access policies on your IAS servers. The EAP-TLS exchange of messages provides mutual authentication, negotiation of the encryption method, and secured private key exchange between the network access client and the IAS server. EAP-TLS provides the strongest authentication and key exchange method.

Use of the TLS authentication type with EAP (EAP-TLS) reduces latency for authentication requests because, during the first authentication attempt, TLS caches certificate properties on the access client and on the IAS server. For subsequent authentication attempts, TLS uses the cached certificate properties instead of reading the certificate from the certificate store on the IAS server or client computer. In addition, the load on the IAS server that is typically caused by processing certificates is greatly reduced, improving server performance. Improved performance and response time are especially useful for wireless networks because the wireless client frequently authenticates to the network when moved from one AP to another.

EAP-TLS also supports computer authentication by using the computer account and computer certificate (also called machine certificate). For computer authentication with EAP-TLS, you must install a computer certificate on the client computer. This certificate is used to authenticate the client computer so that the computer can obtain network connectivity and Group Policy updates prior to user login. For user authentication with EAP-TLS, after a network connection is made and the user logs in, a user certificate on the client computer is necessary.

The computer certificate is installed on the IAS server computer so that during EAP-TLS authentication, the IAS server has a certificate to send to the client computer for mutual authentication, regardless of whether the client computer authenticates with a computer certificate or a user certificate. The computer and user-certificates that are submitted by the access client and the IAS server during EAP-TLS authentication, must conform to specified requirements. For more information about these requirements, see “Using Certificate-Based Authentication Methods” later in this section.

**Deploying EAP**
When you configure a remote access policy with EAP, if the EAP types displayed in the **Select EAP Providers** dialog box are selected as **(Server-Configured)**, then the configuration of the EAP type is common to the server and applies to all remote access policies that are configured on the server. Otherwise, the EAP type configuration is specific to the current remote access policy, in which case the EAP method must be configured for each policy to which you want it to apply.

EAP-TLS is not supported if the IAS server is not a member of the domain.

### Note

- When deploying EAP-TLS, if you have issued a certificate to your IAS server that has a blank **Subject**, the certificate is not available to authenticate your IAS server. To change this, you can use Certificate Templates to create a new certificate for enrollment on your IAS server. In the **Certificate Properties** dialog box, on the **Subject Name** tab, in **Subject name format**, select a value other than **None**.

EAP-based authentication is enabled when:

- EAP is enabled as an authentication protocol on the NAS.
- EAP is enabled on the IAS server and configure the EAP type on the appropriate remote access policy.
- EAP is enabled and configure EAP on a network access client.

In addition to the EAP types defined and supported in Windows Server 2003, you can include new EAP authentication methods by using the EAP Software Development Kit.

### Note

- When using EAP-TLS on computers running Microsoft Windows XP (prior to Service Pack 1), you must have user certificates stored on the computer for user authentication (instead of using smart cards). For computers running Windows Server 2003, Windows XP (Service Pack 1 or later), or Windows 2000, you can use either user certificates stored on the computer or a smart card for user authentication.

**Protected Extensible Authentication Protocol**

Protected Extensible Authentication Protocol (PEAP) is a new member of the family of EAP authentication methods. PEAP uses TLS to create an encrypted channel between an authenticating PEAP client, such as a wireless client computer, and a PEAP authenticator, such as an IAS server. PEAP does not specify an authentication type. Instead, it provides additional security for other EAP authentication protocols, such as MS-CHAP v2, which can operate within the TLS encrypted channel provided by PEAP. PEAP is used as an authentication method for 802.1X connections, such as 802.11 wireless connections and authenticating switches. PEAP is not supported for VPN or other remote access clients.

To use PEAP in your environment with Windows Server 2003, IAS requires little infrastructure change. If your wireless APs support 802.1X and EAP, then you can use PEAP without any additional changes to your wireless APs. In addition, PEAP is available for Microsoft Windows XP, Service Pack 1 and later; Windows 2000, Service Pack 4; and Pocket PC 2002.

To enhance both the EAP protocols and network security, PEAP provides the following:

- Protection for the EAP method negotiation that occurs between client and server through a TLS channel. This helps prevent an attacker from injecting packets between the client and the NAS to cause the negotiation of a less secure EAP method. The encrypted TLS channel also helps prevent denial of service attacks against the IAS server.

- Support for the fragmentation and reassembly of messages, allowing the use of EAP types that do not provide this.

- Wireless clients that can authenticate the IAS or RADIUS server. Because the server also authenticates the client, mutual authentication occurs.

- Support for computer authentication using the client computer account and computer certificate.

- Protection against the deployment of an unauthorized wireless access point (AP) or authenticating switch when the EAP client authenticates the certificate provided by the IAS server. In addition, the TLS master secret created by the PEAP authenticator and client is not shared with the access server. Because of this, the access server cannot decrypt the messages protected by PEAP.

- PEAP fast reconnect reduces the latency for authentication requests and responses between an access client and the IAS server when the access client moves from one AP to another. The IAS server caches a portion of the access client credentials and uses these cached credentials to quickly authorize the connection after the client is associated with a new AP.

**PEAP authentication process**

There are two stages in the PEAP authentication process. The first stage establishes a secure channel between the EAP client and the IAS server. The second stage provides EAP authentication between the EAP client and the IAS server. The following example describes the two stage PEAP authentication process for wireless clients:

**TLS encrypted channel**

The wireless client associates with a wireless access point. An IEEE 802.11-based association provides an Open System or Shared Key Authentication before a secure association is created between the client and AP. If Shared Key Authentication is configured for use at the AP, then 802.1X is not used. If the AP is configured for Open System Authentication, 802.1X is used. After the IEEE 802.11-based association is successfully established between the client and AP, the TLS session is negotiated with the RADIUS server. The key that is derived during this negotiation is used to encrypt all subsequent communication, and the TLS encrypted channel is established.

**EAP-authenticated communication**

An EAP negotiation in which the wireless client authenticates to the IAS server occurs through the TLS channel. The IAS server authenticates the user or client computer with the method that is determined by the EAP type and selected for use within PEAP (either TLS or MS-CHAP v2). The AP only forwards messages between wireless client and RADIUS server. The AP (or person monitoring the AP) cannot decrypt these messages because the AP is not the TLS endpoint.MS-CHAP v2

**802.11 wireless deployments using PEAP**

You can choose either of these two EAP types to use with PEAP: MS-CHAP v2 or TLS.

MS-CHAP v2 uses credentials (user name and password) for user authentication, and a certificate in the server computer certificate store for server authentication.

TLS uses either certificates installed in the client computer certificate store or a smart card to authenticate a user or client computer, and a certificate in the server computer certificate store for server authentication.

**PEAP with EAP-MS-CHAP v2**
PEAP with EAP-MS-CHAP v2 (PEAP-MS-CHAP v2) is easier to deploy than EAP-TLS because user authentication is accomplished with password-based credentials (user name and password) instead of certificates or smart cards. Only the IAS server is required to have a certificate. Additionally, the server certificate can be issued by a public CA that is trusted by the client computer, meaning that the public CA certificate already exists in the Trusted Root Certification Authority folder in the client computer certificate store. In this case, the server certificate is not downloaded and added to the client trusted root certificate store, and the user is not prompted to make a decision about whether to trust the server.

PEAP-MS-CHAP v2 provides improved security over MS-CHAP v2 by using mutual authentication, by preventing an unauthorized server from negotiating the least secure authentication method, and by providing key generation with TLS. PEAP-MS-CHAP v2 requires that the client trust the certificates provided by the server.

**PEAP with EAP-TLS**
Public Key certificates provide a much stronger authentication method than those that use password-based credentials. PEAP with EAP-TLS (PEAP-TLS) uses certificates for server authentication and either certificates or smart cards for user and client computer authentication. To use PEAP-TLS, you must deploy a PKI.

PEAP-TLS also supports computer authentication by using the computer account and computer certificate. For computer authentication with PEAP-TLS, you must install a computer certificate, on the client computer. A computer certificate installed on the client computer is used to authenticate the client computer so that the computer can obtain network connectivity and Group Policy updates prior to user login. For user authentication with PEAP-TLS after a network connection is made and the user logs in, a user certificate on the client computer is necessary.

The computer certificate is installed on the IAS server computer so that during PEAP-TLS authentication, the IAS server has a certificate to send to the client computer for mutual authentication, regardless of whether the client computer authenticates with a computer certificate or a user certificate. The computer and user certificates submitted by the access client and the IAS server during PEAP-TLS authentication must conform to the requirements specified in the section "Using Certificate-Based Authentication Methods."

**TLS caching of certificate properties**
When using the PEAP-TLS and EAP-TLS authentication methods with certificates, TLS uses cached certificate properties instead of reading the certificate from the certificate store on the IAS server.

If a certificate is deleted, its TLS credential cache entry is removed. If a certificate is renewed, its TLS credential cache entry is removed and TLS caches properties for the renewed certificate. If a certificate expires, TLS continues using outdated cached certificate information until the cache expires or is refreshed.

If you change or replace a certificate, you can refresh the TLS cache by restarting the server computer.

**Deploying PEAP**
PEAP is used as an authentication method for 802.1X-based connections and is not supported for VPN or other remote access clients. When you deploy both PEAP and EAP unprotected by PEAP on the same network, do not use the same authentication type (MS-CHAP v2 or TLS) with and without PEAP.

For example, if you deploy PEAP with EAP-TLS (PEAP-TLS), do not also deploy EAP-TLS without PEAP. Deploying authentication methods with the same type — one with and the other without external protection of PEAP or IPSec — creates a security issue.

If you need to deploy both PEAP and EAP on the same network, use different authentication types for each deployment.

An example of a deployment of both authentication methods that does not introduce a security issue is to use PEAP-MS-CHAP v2 for secure password authentication for wireless clients while using EAP-TLS as the authentication method for remote access clients using VPN connections. For this configuration, two remote access policies are configured at the IAS server, one for wireless clients and one for VPN clients, with the appropriate authentication method and type configured for each policy.

**PEAP fast reconnect**
PEAP fast reconnect enables wireless clients to move between wireless APs on the same network without being reauthenticated each time they associate with a new AP.

Wireless APs are configured as RADIUS clients to RADIUS servers. If a wireless client roams between APs that are configured as clients of the same RADIUS server, the client is not required to be authenticated with each new association.

PEAP fast reconnect reduces the response time for authentication between client and authenticator because the authentication request is forwarded from the new server to the original server. Because both the PEAP client and authenticator use previously cached TLS connection properties (the collection of which is named the TLS handle), the authenticator can quickly determine that the client connection is a reconnect.

If the original PEAP authenticator is unavailable, full authentication must occur between the client and the new authenticator. The TLS handle of the new PEAP authenticator is cached by the client. The client can cache TLS handles for multiple PEAP authenticators. For smart cards or PEAP-MS-CHAP v2 authentication, the user is asked to supply the personal identification number (PIN) or credentials, respectively.

The following tables, PEAP-MS-CHAP v2 and "PEAP-TLS" describe authentication behavior with PEAP-MS-CHAP v2 and PEAP-TLS when wireless clients roam from one wireless AP to a new AP. PEAP authentication behavior differs when wireless APs are configured either as clients to the same RADIUS server or as clients to different RADIUS servers.

**PEAP-MS-CHAP v2**

| **New AP Is RADIUS Client to Same RADIUS Server** | **New AP Is RADIUS Client to a Different RADIUS Server** |
|---|---|
| The user is not prompted for credentials each time the client computer associates with a new AP. | The user is prompted for credentials on this initial association. By default, the domain credentials of the currently logged in user are sent automatically. The next time the client computer associates with an AP that is a client to this server, user credentials are not required. |
| The RADIUS server is not required to provide a certificate. | The RADIUS server provides a certificate on this initial association so that the wireless client can authenticate to the RADIUS server. The next time the client computer associates with an AP that is a client to this server, the server is not required to be reauthenticated. |

**PEAP-TLS**

| **New AP is RADIUS Client to Same RADIUS Server** | **New AP Is RADIUS Client to a Different RADIUS Server** |
|---|---|
| The wireless client and RADIUS server are not required to exchange certificates. | The wireless client and RADIUS server exchange certificates on this initial association. The next time the client computer associates with an AP that is a RADIUS client to this RADIUS server, certificates are not exchanged. |
| The user is not prompted for a smart card PIN each time the wireless client computer associates with a new AP. | The user is prompted for a smart card PIN on this initial association. The next time the client computer associates with an AP that is a RADIUS client to this RADIUS server, the user is not prompted for the PIN. When the TLS cache expires, the user is prompted to reenter the PIN. |

PEAP fast reconnect requires the following:

- Both the PEAP client (802.11 wireless client) and PEAP authenticator (RADIUS server) must have fast reconnect enabled.
- All APs to which the PEAP client roams must be configured as RADIUS clients to a RADIUS server (the PEAP authenticator) for which PEAP is configured as the authentication method for wireless connections.
- All APs to which the PEAP client associates must be configured to prefer the same RADIUS server (PEAP authenticator) in order to avoid being prompted for credentials from every RADIUS server. If the AP cannot be configured to prefer a RADIUS server, you can configure an IAS RADIUS proxy with a preferred RADIUS server.

PEAP-based authentication is enabled when:

- You install a certificate on the IAS server that meets the minimum server certificate requirements.
- You configure PEAP as the authentication method for a remote access policy on your IAS server.
- You enable and configure PEAP on a remote access client.

**Using Certificate-Based Authentication Methods**
Certificates are used for network access authentication because they provide strong security for authenticating users and computers and eliminate the need for less secure password-based authentication methods.

Two authentication methods use certificates: EAP-TLS and PEAP. Both methods always use certificates for server authentication, and both the remote access client and the IAS server are TLS endpoints during the authentication process, providing strong security.

Depending on the authentication type configured with the authentication method, certificates might also be used for user authentication and client computer authentication. Certificates are stored in one of two places: either in the computer or user certificate store, or embedded in a computer chip on a smart card. If you deploy EAP-TLS, user authentication can be performed by using either a smart card or a certificate in the Current User certificate store on the client computer. If you deploy PEAP-TLS, the same is true. If you deploy PEAP-MS-CHAP v2, certificates are not used on the client computer; the user supplies password-based credentials when logging on to the network.

You can configure remote access policies in IAS to use EAP-TLS to perform certificate-based authentication for many types of network access, including VPN and wireless connections. In addition, you can use PEAP as the authentication method for wireless or wired 802.1X connections.

To deploy certificates to users and computers, you must first do the following:

- Design a PKI.
- Deploy a CA by using Certificate Services or by using a third party certification authority.
- Design and publish one or more certificates using Certificate Templates, which are installed with Certificate Services.
- Select one or more distribution methods (also called certificate enrollment methods) to install the certificates on computers and distribute them to users.

If you deploy PEAP-MS-CHAP v2 for wireless clients, you can use a non-Microsoft CA to obtain a server certificate.

Because user authentication is accomplished with password-based credentials with PEAP-MS-CHAP v2, you do not need to deploy a PKI or a CA on your network, however the server certificate you obtain from a third party CA must conform to server certificate requirements.

All certificates that are used for network access authentication must meet the requirements for X.509 certificates and work for connections that use Secure Sockets Layer-Transport Layer Security (SSL/TLS). After this minimum requirement is met, both server and client certificates must meet the additional requirements listed in the "Server certificate requirements" section.

**Server certificate requirements**
You can configure clients to validate server certificates by using the **Validate server certificate** option. Using PEAP-MS-CHAP v2, PEAP-TLS, or EAP-TLS as the authentication method, the client accepts the authentication attempt of the server when the certificate meets the following requirements:

- The Subject name contains a value. If you issue a certificate to your IAS server that has a blank Subject, the certificate is not available to authenticate your IAS server.
- The computer certificate on the server chains to a trusted root CA and does not fail any of the checks that are performed by CryptoAPI.
- The IAS or VPN server computer certificate is configured with the Server Authentication purpose in Enhanced Key Usage (EKU) extensions (the object identifier for Server Authentication is 1.3.6.1.5.5.7.3.1).
- The server certificate is configured with a required cryptographic service provider (CSP) value of Microsoft RSA SChannel Cryptographic Provider.
- The Subject Alternative Name (SubjectAltName) extension, if used, must contain the DNS name of the server. With PEAP and EAP-TLS, servers display a list of all installed certificates in the computer's certificate store, with the following exceptions:

- Certificates that do not contain the Server Authentication purpose in EKU extensions are not displayed.
- Certificates that do not contain a Subject name are not displayed.
- Servers do not display registry-based and smart card-logon certificates.

**Note**

- You can designate which certificate is used by the IAS or VPN server in the remote access policy profile. When you configure EAP and PEAP authentication methods that require a certificate for server authentication and you do not select a specific certificate in the **Smart Card or other Certificate Properties** dialog box, the IAS or VPN server automatically selects a certificate from the computer certificate store. If the server obtains a newer certificate, it uses the new certificate. This might cause the IAS or VPN server to use a certificate that is not correctly configured for authentication, causing authentication to fail as the result. To prevent this, always manually select a server certificate when configuring PEAP and EAP authentication methods that require one.

**Client certificate requirements**

With EAP-TLS or PEAP-TLS, the server accepts the client authentication attempt when the certificate meets the following requirements:

- The client certificate is issued by an enterprise CA or mapped to a user or computer account in Active Directory.
- The user or computer certificate on the client chains to a trusted root CA, includes the Client Authentication purpose in EKU extensions (the object identifier for Client Authentication is 1.3.6.1.5.5.7.3.2), and fails neither the checks that are performed by CryptoAPI and specified in the remote access policy nor the Certificate object identifier checks that are specified in IAS remote access policy.
- The 802.1X client does not use registry-based certificates that are either smart card-logon or password-protected certificates.
- For user certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate contains the user principal name (UPN).
- For computer certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate must contain the fully qualified domain name (FQDN) of the client, which is also called the DNS name. With PEAP-TLS and EAP-TLS, clients display a list of all installed certificates in the Certificates snap-in with the following exceptions:
    - Wireless clients do not display registry-based and smart card-logon certificates.
    - Wireless clients and VPN clients do not display password-protected certificates.
    - Certificates that do not contain the Client Authentication purpose in EKU extensions are not displayed.

**Certificate enrollment methods and domain membership**

The domain membership of computers for which you want to enroll certificates affects the certificate enrollment method that you can choose. Certificates for domain member computers can be enrolled automatically (also known as auto-enrollment), while an administrator must enroll certificates for non-domain member computers using the CA Web enrollment pages or a floppy disk. The certificate enrollment method for non-domain member computers is known as a trust bootstrap process, through which certificates are created and then manually requested or distributed securely by administrators, providing administrators the ability to build common trust between non-domain member computers and the domain.

**Domain member certificate enrollment**

If your VPN server, IAS server, or client running Windows 2000 or Windows XP Professional is a member of a domain running Windows Server 2003 and Active Directory, you can configure the auto-enrollment of computer and user certificates through Group Policy. After auto-enrollment is configured and enabled, all domain member computers receive computer certificates the next time Group Policy is refreshed whether the refresh is triggered manually by using the **gpupdate** command or automatically when logging on to the domain.

If your computer is a member of a domain where Active Directory is not installed, you can install computer certificates manually by requesting them through the Certificates snap-in.

**Note**

- Computers running Windows 2000 can auto-enroll computer certificates only.

**Non-domain member certificate enrollment**

Certificate enrollment for computers that are not domain members cannot be done with auto-enrollment. When a computer is joined to a domain, a trust is established that allows auto-enrollment to occur without administrator intervention. When a computer is not joined to a domain, trust is not established and a certificate is not issued. Trust must be established using one of the following methods:

- An administrator (who is, by definition, trusted) must request a computer or user certificate by using the CA Web enrollment tool.

- An administrator must save a computer or user certificate to a floppy disk and install it on the non-domain member computer. Or, when the computer is not accessible to the administrator (for example, a home computer connecting to an organization network by means of an L2TP/IPSec VPN connection), a domain user whom the administrator trusts can install the certificate.

- An administrator can distribute a user certificate on a smart card issued to a user (computer certificates are not distributed on smart cards).

Many network infrastructures contain VPN and IAS servers that are not domain members. For example, a VPN server in a perimeter network might not be a domain member for security purposes. In this case, a computer certificate with the Server Authentication purpose contained in the EKU extensions must be installed on the non-domain member VPN server before it can successfully negotiate L2TP/IPSec-based VPN connections with clients. Note that if the non-domain member VPN server is used as an endpoint for a VPN connection with another VPN server, EKU extensions must contain both the Server Authentication and Client Authentication purposes.

**Choosing certificate templates and certificate enrollment methods**

Windows Server 2003, Standard Edition, provides different certificate templates from Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

If you run an enterprise CA on a computer that runs Windows Server 2003, Standard Edition, you can use the following table, “Certificate Use for Windows Server 2003, Standard Edition,” to determine the following: the best certificate templates; the certificate purposes configured in the EKU extensions of the certificate; and the certificate enrollment method for your requirements.

**Certificate Use for Windows Server 2003, Standard Edition**

| Object and Domain Membership | Certificate Template | Certificate Purposes | Preferred Certificate Enrollment Method | Alternate Certificate Enrollment Method |
|---|---|---|---|---|
| VPN or IAS server, domain member | Computer | Server Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| VPN server with site-to-site connection, domain member | Computer | Server Authentication and Client Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| Windows XP Professional client, domain member | Computer | Client Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| VPN or IAS server, non-domain member | Computer | Server Authentication | CA Web enrollment tool | Install from a floppy disk |
| VPN server with site-to-site connection, non-domain member | Computer | Server Authentication and Client Authentication | CA Web enrollment tool | Install from a floppy disk |

| | | | | |
|---|---|---|---|---|
| Windows XP Professional client, non-domain member | Computer | Client Authentication | CA Web enrollment tool | Install from a floppy disk |
| User, domain user | User | Client Authentication | Auto-enrollment | Use a smart card or the CA Web enrollment tool |

If you are running an enterprise CA on a computer running Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; the 64-bit of Windows Server 2003, Enterprise Edition; or the 64-bit version of Windows Server 2003, Datacenter Edition, you can use the following table, "Certificate Use for Other Versions of Windows Server 2003," to determine the best certificate templates and certificate enrollment method for your requirements.

**Certificate Use for Other Versions of Windows Server 2003**

| Object and Domain Membership | Certificate Template | Certificate Purpose | Preferred Certificate Enrollment Method | Alternate Certificate Enrollment Method |
|---|---|---|---|---|
| VPN or IAS server, domain member | RAS and IAS Server | Server Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| Windows XP Professional client, domain member | Workstation Authentication | Client Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| VPN or IAS server, non-domain member | RAS and IAS Server | Server Authentication | CA Web enrollment tool | Install from a floppy disk |
| Windows XP Professional client, non-domain member | Workstation Authentication | Client Authentication | CA Web enrollment tool | Install from a floppy disk |

**Note**

- If your server running IAS is not a domain controller but is a member of a domain with a Windows 2000 mixed functional level, you must add the server to the access control list (ACL) of the **RAS and IAS Server** certificate template. You must also configure the correct permissions for auto-enrollment. There are different procedures for adding single servers and groups of servers to the ACL.

You can add an individual server to the ACL for the **RAS and IAS Server** certificate template by doing the following: In Certificate Templates, select the template **RAS and IAS Server**, and then add the IAS server to the template **Security** properties. For more information, see "To allow subjects to request a certificate that is based on the template" in Help and Support Center for Windows Server 2003. After you add your IAS server to the ACL, grant **Read**, **Enroll**, and **Auto-enroll** permissions.

You can manage a group of servers by adding the servers to a new global or universal group, and then adding the group to the ACL of the certificate template:

1. In Active Directory Users and Computers, create a new global or universal group for IAS servers. Next, add to the group all computers that are IAS servers, are not domain controllers, and are members of a domain with a Windows 2000 mixed functional level.

2. In Certificate Templates, select the **RAS and IAS Server** template, and then add the group you created to the template Security properties by completing the steps.

After you add your new group, grant **Read**, **Enroll**, and **Auto-enroll** permissions.

**Unauthenticated Access**

The unauthenticated access method allows remote access users to log on without having their identity verified by IAS. For example, IAS does not verify the user's name and password against an account in Active Directory when unauthenticated access is enabled. The only user validation performed with the unauthenticated access method is authorization to verify that the user has permission to access the network.

Enabling unauthenticated access presents security risks that you must carefully consider to determine whether to enable this authentication method. Three scenarios that illustrate typical security considerations for unauthenticated access include the following: Guest Access; Dialed Number Identification Service (DNIS) authorization; and Automatic Number

Identification/Calling Line Identification (ANI/CLI) authorization.

**Guest access for PPP users**

Guest access is the ability to log on to a domain without a user name or a password. Both Routing and Remote Access service and IAS must be configured to support unauthenticated access.

When a remote access server receives a connection attempt, it negotiates with the user different authentication types enabled at the server. If the client accepts one of them, it sends the appropriate credentials for the accepted authentication type. It the user refuses authentication, Routing and Remote Access service checks its properties to verify if unauthenticated access is enabled and, if enabled, forwards the Access-Request message to IAS. This Access-Request message does not contain a User-Name attribute or any other credentials.

When IAS receives the message without a User-Name attribute, it processes the message as an attempt by the user to connect as guest. In this case, IAS uses the name of the guest account in a domain as the user identity. IAS then evaluates policies to determine the right profile. If a match is found, and unauthenticated access is enabled in the profile, other authorizations are validated, and an Access-Accept message is returned. The accounting log file logs the user identity and authentication type, which can be used to determine whether the user is logged on with guest access.

You can enable Guest access by doing the following:

1. Enable unauthenticated access on the remote access server.

2. Enable unauthenticated access on the appropriate remote access policy.

3. Enable the Guest account.

4. Set the remote access permission on the Guest account to either Allow access or Control access through Remote Access Policy depending on your remote access policy administrative model.

If you do not want to enable the Guest account, do the following:

1. Create a user account and set the remote access permission to either **Allow access** or **Control access through Remote Access Policy**.

2. Set the **Default User Identity** registry value on the authenticating server (either the remote access server or the IAS server) to the name of the account.

Following is an example of guest access.

1. During PPP negotiation, the dial-in client rejects all of the PPP authentication protocols of the NAS.

2. If the NAS is configured to allowed unauthenticated access, the NAS sends an Access-Request message without the User-Name attribute and without a password. For the Windows Server 2003 or Windows 2000 Routing and Remote Access service, unauthenticated access is enabled from the **Authentication** tab on the properties of a server in the Routing and Remote Access snap-in.

3. Because the User-Name attribute is not included in the Access-Request message and by default the IAS user identity is using the User-Name attribute, the user identity is set to Guest (or the value of **Default User Identity**).

4. With the user identity of Guest and an unauthenticated connection attempt, The authentication and authorization process as discussed earlier in the chapter is performed. If the connection attempt matches a policy whose profile settings have unauthenticated access enabled and the Guest account is enabled and has the appropriate remote access permission, IAS sends an Access-Accept message to the NAS.

**DNIS authorization**

Dialed Number Identification Service (DNIS) authorization is the authorization of a connection attempt based on the number called. This attribute is referred to as Called-Station-ID. DNIS is used by standard telecommunication companies. This service returns the number called to the called party. Based on the Called-Station-ID attribute, IAS can deliver different services to dial-up/remote access users.

You can enable DNIS authorization by doing the following tasks:

1. Enable unauthenticated access on the remote access server.

2. Create a remote access policy on the authenticating server (remote access server or IAS server) for DNIS-based authorization with the Called-Station-ID condition set to the phone number.

3. Enable unauthenticated access on the remote access policy for DNIS-based authorization.

**ANI authorization**

ANI authorization is based on the number the user calls from. This attribute is referred to as Calling Station ID, or Caller ID. Based on the Calling-Station-ID attribute, IAS can deliver different services to dial-up/remote access users.

Using ANI authorization is different from using the Caller ID dial-in property of a user account. ANI authorization is performed when the user does not type any user name or password, and refuses to use any valid authentication method. In this case, IAS receives Calling-Station-ID, and no user name and password. To support ANI authorization, the Active Directory must have user accounts with Caller IDs as user names. This kind of authentication is used with the cellular phone authentication and by ISPs in Germany and Japan.

When using the Caller ID property on a user account, the user types in credentials, such as a user name and password, and uses a valid authentication method to log on. IAS authenticates the user with the credentials, and then compares the Calling-Station-ID attribute in the Access-Request to the Caller ID property of the user account to authorize the connection attempt.

You can enable ANI authorization by doing the following:

1. Enable unauthenticated access on the remote access server.

2. Enable unauthenticated access on the appropriate remote access policy for ANI/CLI-based authentication.

3. Create a user account for each number calling, for which you want to provide ANI/CLI authorization. The name of the user account must match the number that the user is dialing from. For example, If a user is dialing in from 555-0100, create a "5550100" user account.

4. Set the **User Identity Attribute** registry value to 31 on the authenticating server.

   To always use the calling number as the user identity, set the **Override User-Name** registry value to 1 on the authenticating server.

### Note

- If the Override User-Name is set to 1 and the User Identity Attribute is set to 31, then the authenticating server can perform only ANI/CLI-based authentication. The result is that normal authentication using authentication protocols such as MS CHAP v1, CHAP, and EAP is disabled.

**ANI example**

The following example explains how ANI/CLI authorization works for a dial-up client with an existing user account called 5550100. The client is dialing in from the phone number 555-0100.

1. During PPP negotiation, the dial-in client rejects all of the PPP authentication protocols of the NAS.

2. If the NAS is configured to allow unauthenticated access, the NAS sends an Access-Request message without the User-Name attribute and without a password. (For the Routing and Remote Access service, you can enable unauthenticated access on the Authentication tab in the properties of a server in the Routing and Remote Access snap-in.)

3. Because the User-Name attribute is not included in the Access-Request message, and the IAS user identity is set to use the Calling-Station-ID attribute, the user identity is set to 5550100.

4. With the user identity of 5550100 and an unauthenticated connection attempt, the authentication and authorization process is performed. IAS sends an Access-Accept message to the NAS if the following two conditions are met:

- The connection attempt matches a policy with a profile setting of unauthenticated access enabled.
- The 550100 account has the appropriate remote access permissions.

**MAC address authorization**

Media Access Control (MAC) address authorization functions in the same way as ANI authorization, but it is used for wireless clients and clients connecting to your network by using an 802.1X authenticating switch.

MAC address authorization is based on the MAC address of the network adapter installed in the user's client computer. Like ANI authorization, MAC address authorization uses the Calling-Station-ID attribute instead of user name and password or certificate-based credentials to identify the user during the connection attempt.

MAC address authorization is performed when the user does not type in any user name or password, and refuses to use any valid authentication method. In this case, IAS receives Calling-Station-ID, and no user name and password. To support MAC address authorization, the Active Directory must have user accounts with MAC addresses as user names.

MAC address authorization is enabled when you do the following:

1. Enable MAC address authorization on access servers (such as wireless APs).
2. Enable unauthenticated access on the appropriate remote access policy for MAC address-based authentication, and enable PAP.
3. Create a user account for each MAC address for which you want to provide MAC address authorization. The name of the user account must match the MAC address of the network adapter installed in the computer from which the user is connecting. The format of the password assigned to the account is determined by the network access server vendor. Review the network access server documentation to determine the appropriate password.
4. Set the **User Identity Attribute** registry value to 31 on the authenticating server. This registry value location is: **HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy**
   - The **User Identity Attribute** registry DWORD decimal value: **31**
5. To always use the MAC address as the user identity, set the **Override User-Name** registry value to 1 on the IAS server. This registry value location is: **HKLM\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy**
   - The **Override User-Name** registry DWORD decimal value: **1**

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry can severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**IAS Authorization**

You can use the authorization feature of IAS to allow or deny connection attempts that are based on the connection parameters. The following sections include in-depth information about configuring remote access policies and VSAs.

**Remote Access Policies**

In Windows NT, versions 3.5, 3.51, and 4.0, authorization is based on **Grant dial-in permission to user** option in either User Manager or the Remote Access Admin utility. An individual user configured callback options. For the Routing and Remote Access service and IAS in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000, network access authorization is granted on the basis of user account dial-in properties and remote access policies.

Remote access policies are an ordered set of rules that define how connections are either authorized or rejected. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting.

If a connection is authorized, the remote access policy profile specifies a set of connection restrictions. The dial-in properties of the user account also provide a set of restrictions. Where applicable, user account connection restrictions override the remote access policy profile connection restrictions.

Remote access policies validate several connection settings before authorizing the connection, including the following:

- Remote access permission
- Group membership
- Type of connection
- Time of day
- Authentication methods
- Advanced conditions:
    - Access server identity
    - Access client phone number or MAC address
    - Whether user account dial-in properties are ignored
    - Whether unauthenticated access is allowed

After the connection is authorized, remote access policies can also be used to specify connection restrictions, including the following:

- Idle time-out time
- Maximum session time
- Encryption strength
- IP packet filters
- Advanced restrictions:
    - IP address for PPP connections
    - Static routes

The following settings can also affect connection restrictions:

- Group membership
- Type of connection
- Time of day
- Authentication methods
- Identity of the access server
- Access client phone number or media access control (MAC) address

- Whether unauthenticated access is allowed

It is important to remember that a network connection is accepted only if the settings of the connection attempt match at least one of the configured remote access policies (subject to the conditions of the dial-in properties of the user object and the profile properties of the remote access policy). If the settings of the connection attempt do not match at least one of the remote access policies, the connection attempt is rejected regardless of the dial-in properties of the user account.

For Windows Server 2003 and Windows 2000 IAS servers, remote access policies are administered from the Routing and Remote Access snap-in (when configured for Windows authentication) or the IAS snap-in.

**Authorizing access**
You can use remote access policies to manage authorization by user or by group.

**Authorization by user**
If you are managing authorization by user, set the remote access permission on the user or computer account to either Grant access or to Deny access. Optionally, create different remote access policies based on different types of connections. For example, you might want to have one remote access policy for dial-up connections and a different remote access policy for wireless connections. Managing authorization by user is recommended only when you have a small number of user or computer accounts to manage.

If you are managing authorization by user, the basic process for authorizing a connection attempt occurs as follows:

- If the connection attempt matches all policy conditions, check the remote access permission setting of the account.
    - If the remote access permission is set to **Grant access**, apply the connection settings of the policy profile and user account.
    - If the remote access permission is set to **Deny access**, reject the connection attempt.
- If the connection attempt does not match all policy conditions, process the next remote access policy.

If the connection attempt does not match all conditions of any remote access policy, reject the connection attempt.

**Authorization by group**
If you are managing authorization by group, set the remote access permission on the user account to **Control access through Remote Access Policy** and create remote access policies that are based on different types of connections and group membership. You can use existing Active Directory groups when you configure remote access policies in IAS, or you can use the Active Directory Users and Computers snap-in to create new Windows groups upon which to base your remote access policies. For example, you might want to have one remote access policy for dial-up connections for employees (members of the Employees group) and a different remote access policy for dial-up connections for contractors (members of the Contractors group).

If you have a large number of user accounts, you can configure IAS to ignore the dial-in properties of user accounts instead of individually configuring each account to **Control access through Remote Access Policy**. You can set **Ignore-User-Dialin-Properties** on the **Advanced** tab of a remote access policy profile in the IAS console.

If you are managing authorization by group, the basic process for authorizing a connection attempt occurs as follows:

- If the connection attempt matches all policy conditions, check the remote access permission of the remote access policy.
    - If the remote access permission is set to **Grant remote access permission**, apply the connection settings of the policy profile and user account.
    - If the remote access permission is set to **Deny remote access permission**, reject the connection attempt.
- If the connection attempt does not match all policy conditions, process the next remote access policy.

If the connection attempt does not match all conditions of any remote access policy, reject the connection attempt.

**Note**

- It is not recommended that you manage network access authorization by user while also setting user account dial-in properties to **Control access through Remote Access Policy**. When you set user account dial-in properties to **Control access through Remote Access Policy**, manage network access authorization by group instead of by user.

**Local vs. Centralized Policy Management**
Because remote access policies are stored locally on either a remote access server or an IAS server, you must take the following steps in order to use centralized management of a single set of remote access policies for multiple remote access or VPN servers:

1. Install IAS as a Remote Authentication Dial-In User Service (RADIUS) server on a computer.

2. Configure IAS with RADIUS clients that correspond to each of the NASs.

3. On the IAS server, create the central set of policies that all NASs are using.

4. Configure each of the NASs as a RADIUS client to the IAS server.

After you configure a Windows Server 2003 or Windows 2000 remote access server as a RADIUS client to an IAS server, the local remote access policies stored on the remote access server are no longer used.

Centralized management of remote access policies are also used when you have remote access servers that are running Windows NT 4.0 with the Routing and Remote Access Service (RRAS). You can configure the server that is running Windows NT 4.0, using RRAS as a RADIUS client to an IAS server. You cannot configure a remote access server that is running Windows NT 4.0 without using RRAS to take advantage of centralized remote access policies.

**Dial-in properties of a user object**
In Windows Server 2003, the user object for a stand-alone server or Active Directory–based server contains a set of dial-in properties that are used when allowing or denying a connection attempt made by a user. For a stand-alone server, the dial-in properties are available on the **Dial-in** tab of the user object in the local User Manager. For an Active Directory–based server, the dial-in properties are available on the **Dial-in** tab of the user object in the Active Directory Users and Computers snap-in. The Windows NT 4.0 User Manager for Domains administrative tool cannot be used for Active Directory–based servers.

The dial-in properties for a user object are the following:

**Remote Access Permission (Dial-in or VPN)**
Use this property to set whether remote access is explicitly allowed, denied, or determined through remote access policies. If access is explicitly allowed, remote access policy conditions or user object or profile properties can override the setting. The **Control access through Remote Access Policy** option is available only on user objects for stand-alone Routing and Remote Access service servers or members of a domain that has a Windows 2000 native or Windows Server 2003 domain functional level.

### Note

- When the domain functional level in Active Directory is set to Windows 2000 mixed, the **Control access through Remote Access Policy** radio button is disabled on the **Dial-in** tab in user account **Properties**. To enable the use of **Control access through Remote Access Policy**, you must raise the domain functional level in Active Directory to either Windows 2000 native or Windows Server 2003.

By default, the Administrator and Guest accounts are set to **Control access through Remote Access Policy** on a stand-alone remote access server or in a Windows 2000 native domain. For a Windows 2000 mixed domain, the Administrator and Guest accounts are set to **Deny access** by default. New accounts created on a stand-alone remote access server or in a Windows 2000 native domain are set to **Control access through Remote Access Policy**. New accounts created in a domain with a Windows 2000 mixed domain functional level are set to **Deny access**.

**Verify Caller ID**
When Caller ID is enabled, the server verifies the caller's phone number. If the caller's phone number does not match the configured phone number, the connection attempt is denied.

Caller ID must be supported by the caller, the phone system between the caller and the remote access server, and the remote access server. On a computer running the Routing and Remote Access service, caller ID support consists of call answering equipment that provides caller ID information and the appropriate Windows driver to pass the information to the Routing and

Remote Access service.

If you configure a caller ID phone number for a user and you do not have support for the passing of caller ID information from the caller to the Routing and Remote Access service, the connection attempt is denied.

**Callback Options**
If this property is enabled, the server calls the caller back during the connection establishment at a phone number set by the caller or a specific phone number set by the administrator.

**Assign a Static IP Address**
If this property is enabled, the administrator assigns a specific IP address to the user when the connection is made.

**Apply Static Routes**
If this property is enabled, the administrator defines a series of static IP routes that are added to the routing table of the remote access server when a connection is made. This setting is designed for user accounts that Windows Server 2003 or Windows 2000 routers use for demand-dial routing.

For a user account that is a member of a Windows NT 4.0 domain or a Windows 2000 mixed domain, only the **Allow access** and **Deny access** Remote Access Permission options and **Callback Options** dial-in properties are available. In addition, the User Manager for Domains administrative tool can be used to grant or deny dial-in access and set callback options.

For a user account that is on a stand-alone server or in a Windows 2000 native domain, the callback number can be of unlimited size. For a user account that is in a Windows NT 4.0 domain or a Windows 2000 mixed domain, the callback number can only be 128 characters long. Callback numbers that are long than 128 characters are truncated during a callback connection attempt, which results in a failed callback connection.

When a Windows NT 4.0 RAS server uses a Windows 2000 native domain to obtain the dial-in properties of a user account, the **Control access through Remote Access Policy** option is retrieved as **Deny access**. Callback settings are retrieved correctly.

User accounts upgraded to Windows Server 2003 or Windows 2000 that were configured with dial-in permission enabled are set to **Allow access**. User accounts upgraded to Windows Server 2003 or Windows 2000 that were configured with dial-in permission disabled are set to **Control access through Remote Access Policy**.

A Windows NT 4.0 Remote Access Service (RAS) server does not use remote access policies. It is recommended that you upgrade Windows NT 4.0 RAS servers to take advantage of remote access policies.

**Elements of a Remote Access Policy**
A remote access policy is a named rule that consists of the following elements:

- Conditions
- Remote access permission
- Profile

**Remote access policy conditions**
Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt. If there are multiple conditions, all of the conditions must match the settings of the connection attempt in order for the connection attempt to match the policy. Not all NASs send all of the IAS server-specific attributes.

The table "Condition Attributes for a Remote Access Policy" shows the condition attributes that you can set for a remote access policy.

**Condition Attributes for a Remote Access Policy**

| Attribute Name | Description |
|---|---|
| Authentication Type | The type of authentication that is being used by the access client. Authentication types include CHAP, EAP, MS-CHAP , and MS-CHAP v2 |

| | |
|---|---|
| Called Station ID | The phone number of the NAS. This attribute is a character string. You can use pattern matching syntax to specify area codes. In order to receive called station ID information during a call, the phone line, the hardware, and the Windows Server 2003 driver for the hardware must support the passing of the called ID. Otherwise, the called station ID is manually set for each port. |
| Calling Station ID | The phone number used by the caller. This attribute is a character string. You can use pattern matching syntax to specify area codes. |
| Client Friendly Name | The name of the RADIUS client that is requesting authentication. This name is configured in **Friendly name** on the **Settings** tab in the properties of a RADIUS client in IAS. This attribute is a character string. You can use pattern matching syntax to specify client names. |
| Client IP Address | The IP address of the RADIUS client. This can be the IP address of either the originating RADIUS client or an intermediate RADIUS proxy. This attribute is a character string. You can use pattern matching syntax to specify IP networks. |
| Client Vendor | The vendor of the NAS that requests authentication. Ensure that you configure the NAS vendor when registering the RADIUS client on the IAS server. For example, the Routing and Remote Access service is the Microsoft NAS manufacturer. You can use this attribute to configure separate policies for RADIUS clients from different NAS manufacturers. This attribute is used by the IAS server. |
| Day and Time Restrictions | The day of the week and the time of day of the connection attempt. The day and time is relative to the day and time of the server providing the authorization. |
| Framed Protocol | The type of framing for incoming packets. Examples are PPP, SLIP, Frame Relay, and X.25. This attribute is used by the IAS server. |
| NAS Identifier | The name of the NAS. This attribute is a character string. You can use pattern matching syntax to specify multiple NASs. |
| NAS IP Address | The IP address of the NAS (the RADIUS client) that sent the message. This attribute is a character string. You can use pattern matching syntax to specify IP networks. |
| NAS Port Type | The type of media that is used by the access client. Examples include analog phone lines (known as async), ISDN, tunnels or virtual private networks (known as virtual), IEEE 802.11 wireless, and Ethernet switches. |
| Service Type | The type of service that is being requested. Examples include framed (such as PPP connections) and login (such as Telnet connections). For more information about RADIUS service types, see RFC 2865, "Remote Authentication Dial In User Service (RADIUS)." This attribute is used by the IAS server. |
| Tunnel Type | The type of tunnel that is being created by the requesting client. Tunnel types include the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol (L2TP), which are used by Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; Windows XP; and Windows 2000 remote access clients and demand-dial routers. You can use this condition to specify profile settings, such as authentication methods or encryption strengths for a specific type of tunneling technology. |
| Windows Groups | The names of the groups to which the user or computer account that is attempting the connection belongs. There is no condition attribute for a specific user or computer name. It is not necessary to have a separate remote access policy for each group. Instead, you can use multiple groups or nested groups to consolidate and delegate the administration of group membership. For a remote access or IAS server in a Windows 2000 native domain, you can use universal groups. You cannot use the built-in local or domain groups for the Windows Groups attribute. |

If conditions that use an IAS server attribute are evaluated against a Routing and Remote Access service server that is configured for Windows authentication, they do not match and the policy is not applied.

**Remote access permission**

IAS evaluates all remote access policies in the ordered list of remote access policies in an effort to match policy conditions with the connection attempt. If IAS can match all of the conditions of one of the remote access policies with the settings of the connection attempt, the matching remote access policy is used for the request. IAS then determines remote access permission by evaluating the setting for **If a connection request matches the specified conditions**. **If a connection request matches the specified conditions** is a setting found in the properties of a remote access policy, and can have one of the following values:

- **Deny remote access permission** . If **Deny remote access permission** is set on the policy, then remote access permission is denied.

- **Grant remote access permission**. If **Grant remote access permission** is set on the policy, then remote access permission is granted.

Remote access permission is also granted or denied for each user account. The user account remote access permission overrides the policy remote access permission. When remote access permission on a user account is set to the **Control access through Remote Access Policy** option, the policy remote access permission determines whether the user is granted access.

Granting access with either the user account permissions setting or the policy permission setting is the first step in accepting a connection. The connection attempt is then subjected to the settings of the user account properties and the policy profile properties and restrictions. If the connection attempt does not match the settings of the user account properties or the profile properties, or if specific restrictions defined in the profile prevent the connection, the connection attempt is rejected. For example, if the remote access policy profile restricts connections to use only EAP as an authentication method but the connection attempt is made using MS-CHAP v2 by a client that is not configured to use EAP, the connection attempt is rejected even when the value of **If a connection request matches the specified conditions** is **Grant remote access permission**.

By default, the **Deny remote access permission policy** permission is selected.

**Profile**
A remote access policy profile is a set of properties that are applied to a connection when the connection is granted remote access permission, either through the user account permission setting or the policy permission setting. A profile consists of the following groups of properties:

- Dial-in constraints

- IP

- Multilink

- Authentication

- Encryption

- Advanced

**Dial-in constraints**
You can set the following dial-in constraints:

- **Minutes server can remain idle before it is disconnected (Idle-Timeout)** The time after which a connection is disconnected when there is no activity. By default, this property is not set. Additionally, the Routing and Remote Access service does not disconnect an idle connection.

- **Minutes client can be connected (Session-Timeout)** The maximum time that a connection is connected. The connection is disconnected by the access server after the maximum session length. By default, this property is not set and there is no maximum session limit.

- **Allow access only on these days and at these times** The days of the week and hours of each day that a connection is allowed. If the day and time of the connection attempt do not match the configured day and time limits, the connection attempt is rejected. By default, this property is not set and there are no day or time limits. The Routing and Remote Access service does not disconnect active connections that were previously connected at a time when connection attempts were allowed.

- **Allow access to this number only (Called-Station-ID)** The specific phone number that a caller must call in order for a connection to be allowed. If the dial-in number of the connection attempt does not match the configured dial-in number, the connection attempt is rejected. By default, this property is not set, and all dial-in numbers are allowed.

- **Allow access only through these media (NAS-Port-Type)** The specific media type, such as modem (also known as

*asy*nc), ISDN (also known as virtual), or 802.11 wireless that a caller must use so that a connection is allowed. If the media of the connection attempt does not match the configured dial-in media, the connection attempt is rejected. By default, this property is not set and all media types are allowed.

**IP**

You can set IP properties that specify IP address assignment behavior. You have the following options:

- The access server must supply an IP address.
- The access client can request an IP address.
- IP address assignment is determined by the access server (this is the default setting).
- A static IP address is assigned. A static IP address assigned to the user account overrides this setting. The IP address assigned is typically used to accommodate VSAs for IP addresses.

You can also use the IP tab to define IP packet filters that apply to remote access connection traffic. This is intended for use with the Routing and Remote Access service. You can configure IP traffic that is allowed to remote access clients (output filters) or from remote access clients (input filters) on an exception basis. Either all traffic is allowed, except traffic specified by filters; or all traffic is blocked, except traffic specified by filters. Remote access policy profile filtering applies to all remote access connections that match the remote access policy.

Filters configured on the IP tab are applied to client connections on a computer running Windows Server 2003 and the Routing and Remote Access service, but are not applied to dial-on-demand connections.

**Multilink**

You can set Multilink properties that both enable Multilink and determine the maximum number of ports that a Multilink connection can use. Additionally, you can set Bandwidth Allocation Protocol (BAP) policies that both determine BAP usage and specify when extra BAP lines are dropped. The Multilink and BAP properties are specific to the Routing and Remote Access service. By default, Multilink and BAP are disabled.

The Routing and Remote Access service must have Multilink and BAP enabled in order for the Multilink properties of the profile to be enforced.

**Authentication**

You can set authentication properties to both enable the authentication types that are allowed for a connection and specify the EAP type that must be used. Additionally, you can configure the EAP type. With Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, you can specify whether users can change their expired passwords using MS-CHAP and MS-CHAP v2.

The Routing and Remote Access service must have the corresponding authentication types enabled in order for the authentication properties of the profile to be enforced.

**Encryption**

You can set encryption properties for the following encryption strengths:

- No Encryption

  When selected, this option allows an unencrypted connection. To require encryption, clear the **No Encryption** option. For L2TP over IPSec-based VPN connections, the Authentication Header (AH) is used.

- Basic

  For dial-up and PPTP-based VPN connections, Microsoft Point-to-Point Encryption (MPPE) with a 40-bit key is used. For L2TP over IPSec-based VPN connections, 40-bit DES encryption is used.

- Strong

  For dial-up and PPTP-based VPN connections, MPPE with a 56-bit key is used. For L2TP over IPSec-based VPN connections, 56-bit DES encryption is used.

- Strongest

  For dial-up and PPTP-based VPN connections, MPPE with a 128-bit key is used. For L2TP over IPSec-based VPN connections, triple DES (3DES) encryption is used.

**Advanced**

You can set **Advanced** properties to specify the series of RADIUS attributes that are sent back to the RADIUS client by the IAS server. RADIUS attributes are specific to performing RADIUS authentication and are ignored by the remote access server. By default, Framed-Protocol is set to PPP and Service-Type is set to Framed.

The only attributes that the Routing and Remote Access service uses are **Account-Interim-Interval, Framed-Protocol, Framed-MTU, Reply-Message**, and **Service-Type**.

**Configuration issues**

Following are several issues that might impact your configuration of IAS, depending on your network and needs:

- Some elements of a remote access policy correspond to RADIUS attributes that are used during RADIUS-based authentication. For remote access policies on an IAS server, verify that the NASs used are sending RADIUS attributes that correspond to the configured remote access policy conditions and profile settings. If a NAS does not send a RADIUS attribute that corresponds to a remote access policy condition or profile setting, then all RADIUS authentication requests from that NAS are denied.

- You can only use the **Generate-Session-Time-out** attribute if your user account database is a Security Accounts Manager (SAM) database or is the user account database for an Active Directory domain. When the value of **Generate-Session-Time-out** is set to **True**, the **ForceLogoff** value for a SAM database should be set to **0**. In the Local Security Settings console, **ForceLogoff** is changed to zero when **Network security: Force logoff when logon hours expire** is enabled.

- You can configure wireless connection policy so that wireless clients periodically reauthenticate. This ensures that the client Wired Equivalent Privacy (WEP) encryption keys are changed often enough to provide adequate security for the wireless connection. To configure reauthentication, set the session time-out in your remote access policy or connection request policy for wireless connections (using the Session-Time-out attribute) to the required interval (for example, 10 minutes). Additionally, configure the Termination-Action attribute with the **Attribute value** set to **RADIUS-Request**. If the Termination-Action attribute is not set to **RADIUS-Request**, wireless APs might end the connection during reauthentication. For more information, see your hardware documentation.

**Default remote access policies**

Two default remote access policies are installed with IAS in Windows Server 2003, with the following names and configurations:

1. The **Connections to Microsoft Routing and Remote Access server** policy has the following configuration:

   - Policy conditions are set to process Access-Request messages sent by RADIUS clients that are computers running the Microsoft Routing and Remote Access service. This policy condition is set by using the MS-RAS Vendor attribute with a wildcard value of **^311$**.

   - Permission is set to **Deny remote access permission**.

   - Two connection attributes are configured in profile properties: Framed-Protocol is configured with a value of **PPP**, and Service-Type is configured with a value of **Framed**.

   - Connections that use no encryption are not allowed. Only connections using basic, strong, or strongest encryption are allowed.

   - IP packet filters are configured to allow traffic only between the IAS server and network access servers configured as RADIUS clients.

   - All other profile properties are set to default values.

2. The **Connections to other access servers** policy has the following configuration:

    - The **Day-and-Time-Restrictions** condition is set to all times and all days.
    - Permission is set to **Deny remote access permission**.
    - Two connection attributes are configured in profile properties: Framed-Protocol is configured with a value of **PPP**, and Service-Type is configured with a value of **Framed**.
    - All profile properties are set to default values.

**Vendor Profiles**
Some vendors use VSAs to provide functionality that is not supported in standard attributes. IAS enables you to create or modify VSAs to take advantage of the proprietary functionality that is supported by some NAS vendors.

If you need to configure more than one VSA for a specific profile, you must arrange them in the appropriate order. If you are using filters and the order of the filters is important, use the arrow buttons to rearrange the attributes.

**Example 1**
The following example demonstrates the procedure of adding a Cisco VSA to a profile. The example illustrates only the mechanism of adding a standard-conforming VSA to a profile. Cisco VSAs are readily available through the IAS multi-vendor dictionary.

Cisco VSAs conform to the RADIUS RFC 2548 for Vendor-Specific Attributes (type 26). The following information is for a Cisco attribute to specify a primary DNS server:

- Vendor ID: 9. This is the unique ID for Cisco. When you specify that vendor, this is automatically supplied.
- Vendor Type: 1. This is the vendor-type number for VSAs that take the attribute-value pair form, referred to in Cisco documentation as "cisco-avpair."
- Data Type: String.
- Format: If the attribute is mandatory, the format is: <protocol>: attribute=value

If the attribute is optional, the attribute-value pair is separated by an asterisk (*) instead of an equal sign (=). In this example, <protocol> is a value of the Cisco "protocol" attribute for a particular type of authorization. "Attribute" and "value" represent an appropriate attribute/value (AV) pair defined in the Cisco TACACS+ specification. This allows the full set of features available for TACACS+ authorization to also be used for RADIUS.

The Cisco attribute used to specify a primary DNS server appears as follows:

```
ip:dns-servers=10.10.10.10
```

**Example 2**
The following example demonstrates how to add a 3Com/U.S. Robotics VSA to a profile.

### Note

- Example 2 is included only to illustrate the mechanism for adding a standard nonconforming VSA to a profile. 3Com/U.S. Robotics VSAs are readily available from the IAS multivendor dictionary.

U.S. Robotics VSAs do not conform to the recommended format for Vendor-Specific Attributes (type 26) in RFC 2865. Therefore, all U.S. Robotics VSAs must be entered in hexadecimal format.

The following information is for a U.S. Robotics attribute to specify a primary DNS/NBNS server:

- Vendor ID: 429. This is the unique ID for U.S. Robotics. When you specify that vendor, this ID is automatically supplied.

- Indicator: 0x900F
- Data Type: String
- Format: The VSA must be entered in hexadecimal format.

For more information about the proprietary attributes of U.S. Robotics, see your U.S. Robotics documentation.

**Accepting a Connection Attempt**

When a user attempts a connection, the connection attempt is accepted or rejected based on the following:

1. The first policy in the ordered list of remote access policies is checked. If there are no policies, the connection attempt is rejected.
2. If all the conditions of the policy do not match the connection attempt, go to next policy. If there are no more policies, reject the connection attempt.
3. If all the conditions of the policy match the connection attempt, check the value of the **Ignore-User-Dialin-Properties** attribute.
4. If the **Ignore-User-Dialin-Properties** attribute is set to **False**, check the remote access permission setting for the user attempting the connection:
   - If **Deny access** is selected, reject the connection attempt.
   - If **Allow access** is selected, apply the user account properties and profile properties:

     If the connection attempt does not match the settings of the user account properties and profile properties, reject the connection attempt.

     If the connection attempt matches the settings of the user account properties and profile properties, accept the connection attempt.
5. If the remote access permission is not set to **Allow access** or **Deny access**, the remote access permission must be set to **Control access through Remote Access Policy**. In that case, check the remote access permission setting of the policy:
   - If **Deny remote access permission** is selected, reject the connection attempt.
   - If **Grant remote access permission** is selected, apply the user account properties and profile properties:

     If the connection attempt does not match the settings of the user account properties and profile properties, reject the connection attempt.

     If the connection attempt matches the settings of the user account properties and profile properties, accept the connection attempt.
6. If the **Ignore-User-Dialin-Properties** attribute is set to **True**, check the remote access permission setting of the policy:
   - If **Deny remote access permission** is selected, reject the connection attempt.
   - If **Grant remote access permission** is selected, apply the profile properties:

     If the connection attempt does not match the settings of the profile properties, reject the connection attempt.

If the connection attempt matches the settings of the profile properties, accept the connection attempt.

**IAS Accounting**
Support for the RADIUS standard allows IAS to log the accounting data produced by NASs and IAS. User authentication and accounting requests and periodic data recorded in log files is used primarily for IAS troubleshooting, connection analysis, and billing purposes. In addition, accounting records can provide you with a useful tool when analyzing possible security problems, such as an attempt by a malicious user to access network resources. You can log accounting data in three formats, including IAS format, database-compatible format, and SQL Server Logging format.

**RADIUS Accounting**
IAS supports RADIUS accounting, with which an administrator can track network usage for auditing and billing purposes. Accounting data can also be queried by Help desk and other applications to assist with network access troubleshooting. RADIUS accounting provides the following benefits:

- Real-time data collection.
- Accounting data can be collected from a central location.
- Non-Microsoft products can be used to analyze RADIUS accounting data to provide charge-back, troubleshooting, performance, and exception reports.

When configured for accounting, IAS can log accounting data to a log file or to a SQL Server database.

When a RADIUS client is configured to use RADIUS accounting, at the start of service delivery it generates an Accounting-Start message describing the type of service being delivered and the user it is being delivered to. The message is then sent to the RADIUS Accounting server, which sends back an acknowledgment to the RADIUS client. At the end of service delivery, the client generates an Accounting-Stop message describing the type of service that was delivered and statistics (optional), such as elapsed time, input and output octets, or input and output packets. It then sends that data to the RADIUS accounting server, which sends back an acknowledgment to the RADIUS client.

The Accounting-Request message (whether for the Start or Stop message) is submitted to the RADIUS accounting server through the network. If no response is returned within a length of time, the request is re-sent a number of times. The client can also forward requests to an alternate server or servers in the event that the primary server is down or unreachable. An alternate server can be used either after a number of tries to the primary server fail, or in a round-robin fashion. If the RADIUS accounting server can not successfully record the accounting message, it does not send an Accounting-Response acknowledgment to the RADIUS client. For example, when the log file gets filled up, IAS starts discarding accounting messages. This prompts the RADIUS client to switch to the backup IAS server.

**SQL Server Logging**
You can use SQL Server logging to record user authentication and accounting requests received from one or more NASs to a data source on a computer running the Microsoft SQL Server Desktop Engine (MSDE 2000) or Microsoft SQL Server 2000. Accounting data is passed from IAS in eXtensible Markup Language (XML) format to a stored procedure in an MSDE 2000 or SQL Server 2000 database, which support both SQL and XML (SQLXML).

Logging to a SQL Server 2000 relational database instead of a standard text file (that is, either IAS format or database-compatible format) has many advantages, including the following:

- Relationships between data tables enable the flexible creation of dynamic data views by using queries and reports.
- SQL Server 2000 has high-speed optimizations that can effectively support very large, terabyte-sized databases.
- Multiple bulk copy operations can be performed concurrently against a single table in SQL Server to speed data entry.
- The Transact-SQL BACKUP and RESTORE statements are optimized to read through a database serially and write in parallel to multiple backup devices. In addition, SQL Server 2000 supports differential backups, which replicate only the data changed after the last backup.
- You can configure SQL Server logging to provide failover and redundancy, and multiple IAS servers can log to the same database, providing ease of administration.
- Data from multiple IAS servers can be stored in the same database, providing a centralized data source for many IAS

servers and the ability to achieve a more comprehensive data view for applications that query the database, such as Help desk applications.

In addition, you can improve logging performance and, in some circumstances, reduce the cost of deploying SQL Server logging by using MSDE 2000 databases installed on each IAS server.

The recommended deployment model for SQL Server logging is to configure IAS to log to a stored procedure in a MSDE 2000 database installed on the local server. You can then write a program, service, or component that periodically publishes the accounting data from the MSDE 2000 database to a central SQL Server.

**Components of IAS SQL Server logging**
IAS SQL Server logging includes several components, some of which are installed on the IAS server, and some of which are configured on the computer running SQL Server 2000.

The figure, "IAS SQL Server Logging," illustrates the components of SQL Server logging with IAS in Windows Server 2003 and a computer running SQL Server 2000.

**IAS SQL Server Logging**




**IAS server**
A computer running Windows Server 2003 and IAS configured to use SQL Server logging.

**MSDE 2000 database**
A data engine and database built and based on core SQL Server technology that is installed on the computer running Windows Server 2003 and IAS.

**Stored procedure**
A group of Transact-SQL statements compiled into a single execution plan, or program, that runs inside the MSDE 2000 database on the computer running Windows Server 2003 and IAS. The stored procedure must be named report_event. Report_event accepts and processes data from the IAS server and stores it in MSDE 2000 tables.

**Custom accounting-forwarder component**
A low-level program installed on the computer running Windows Server 2003 and IAS that automatically manages and administers the MSDE 2000 database. The custom component, service, or application should periodically publish records from the MSDE 2000 database on the IAS server to the IAS database on the computer running SQL Server. The custom component should also flag records as published, and delete previously published records from the MSDE 2000 database during off-peak hours. If there are multiple IAS servers configured for SQL Server logging, each server must have the custom component installed to process records in the local MSDE 2000 database and forward records to the SQL Server database.

**SQL server**
A computer or computer cluster running SQL Server 2000. If you have a small IAS deployment, you might need only one computer running SQL Server; if you have a large IAS deployment in an enterprise network environment, you might need more than one computer running SQL Server or a cluster of computers running SQL Server.

**IAS database**
A database on the computer running SQL Server dedicated to storing the records from all IAS servers configured for SQL Server logging.

**Configuring accounting to provide the best session correlation**
When you use an application to query your SQL Server 2000 database, it is important that sufficient data is logged to provide the application the ability to correlate related fields and information into a cohesive view of any particular user session. This is called *session correlation*. To provide the best session correlation, take the following actions:

- Use NASs that send the Class attribute in all accounting-requests. The Class attribute is sent to the RADIUS client in an Access-Accept message, and is useful for correlating Accounting-Request messages with authentication sessions.

- If the Class attribute is sent by the NAS in the accounting request messages, it can be used to match the accounting and authentication records. The combination of the attributes Unique-Serial-Number, Service-Reboot-Time, and Server-Address must be a unique identification for each authentication that the server accepts.

- Use NASs that support interim accounting.

- Use NASs that send on/off accounting messages.

- Use NASs that support the storing and forwarding of accounting data. NASs that support this feature can store accounting data in circumstances when it cannot communicate with the IAS server. When the IAS server is available, the NAS forwards the stored records to the IAS server, providing increased reliability in accounting over NASs that do not provide this feature.

- Always configure the Acct-Interim-Interval attribute in remote access policies. The Acct-Interim-Interval attribute sets the interval (in seconds) between each interim update that the NAS sends. You can add the Acct-Interim-Interval RADIUS attribute on the Advanced tab on the profile settings of the remote access policy. According to RFC 2869, the value of the Acct-Interim-Interval attribute must not be smaller than 60 seconds, or one minute, and should not be smaller than 600 seconds, or 10 minutes. For more information, see RFC 2869, "RADIUS Extensions."

- Ensure that logging of periodic status is enabled on your IAS servers.

**IAS log file**

IAS can create a log file based on the data returned by the NASs. This information is useful for keeping track of usage and correlating authentication information with accounting records (for example, to discover missing records or instances of over-billing).

IAS supports two formats of the log file: IAS format and database. Database format allows you to keep track of a predetermined set of attributes, and IAS format is more detailed and can contain information about all attributes. Use database format if you want to import the data directly into a database. The IAS format can be used if you need to record more detailed information than the database log format allows.

### Note

- The IAS log file contains all the IAS user-related events. IAS service and system-related events are recorded in the Event log files.

**IAS Authentication and Domain Functionality**

IAS authentication features and behavior vary with the functional levels of the domains in which IAS is running. Knowing the differences is useful when making decisions about a particular domain in which IAS is deployed.

Domain and forest functionality, introduced in Windows Server 2003 Active Directory, provides a way to enable domain- or forest-wide Active Directory features within your network environment. Different levels of domain functionality and forest functionality are available depending on your environment.

If all domain controllers in your domain or forest are running Windows Server 2003 and the functional level is set to Windows Server 2003, all domain- and forest-wide features are available. When Windows NT 4.0 or Windows 2000 domain controllers are included in your domain or forest with domain controllers running Windows Server 2003, Active Directory features are limited.

The concept of enabling additional functionality in Active Directory exists in Windows 2000 with mixed and native modes. Mixed-mode domains can contain Windows NT 4.0 backup domain controllers and cannot use Universal security groups, group nesting, and security ID (SID) history capabilities. When the domain is set to native mode, Universal security groups, group nesting, and SID history capabilities are available. Domain controllers running Windows 2000 Server are not aware of domain and forest functionality.

**Domain Functionality**

Domain functionality enables features that will affect the entire domain and that domain only. Four domain functional levels are available: Windows 2000 mixed (default), Windows 2000 native, Windows Server 2003 interim, and Windows Server 2003. By default, domains operate at the Windows 2000 mixed functional level.

The following table, “Domain Controller Support by Domain Functional Level,” lists the domain functional levels and their corresponding supported domain controllers.

**Domain Controller Support by Domain Functional Level**

| Domain Functional Level | Domain Controllers Supported |
|---|---|
| Windows 2000 mixed (default) | Windows NT 4.0<br>Windows 2000<br>Windows Server 2003 family |
| Windows 2000 native | Windows 2000<br>Windows Server 2003 family |
| Windows Server 2003 interim | Windows NT 4.0<br>Windows Server 2003 family |
| Windows Server 2003 | Windows Server 2003 family |

IAS is capable of authenticating access requests received through the RADIUS protocol against domains with Windows 2000 native, Windows 2000 mixed, Windows Server 2003 interim, and Windows Server 2003 domain functional levels. When you have a stand-alone IAS server, IAS can also authenticate access requests when configured with a Windows Server 2003 local accounts database.

The IAS authentication features and capabilities available to administrators are dependant upon the domain functional level of a particular domain against which the users authenticate.

**The Windows Server 2003 Domain Functional Level**
Domains configured with the Windows Server 2003 domain functional level provide the most flexibility in managing remote access through groups, including full support for all features.

**The Windows Server 2003 Interim Domain Functional Level**
Domains with a Windows Server 2003 interim functional level are mainly used for migration from Windows NT 4.0 to Windows Server 2003. For IAS, a Windows Server 2003 interim domain acts exactly like a Windows NT 4.0 domain.

For an IAS server that is a member of a Windows Server 2003 interim domain, the following authentication and remote access management features are available:

- User account dial-in properties
- Remote access permissions include only **Allow access** and **Deny access**

  Without the **Control access through Remote Access Policy** option, policy-based management using groups becomes more difficult because the user’s remote access permission overrides remote access policy permissions.
- Callback options

In order for the IAS server to access user account dial-in properties stored in Active Directory, the Internet Authentication service must run in the security context of a computer account that is a member of the **RAS and IAS Servers** security group. This assignment can be implemented through the Active Directory Users and Computers snap-in or by registering the IAS server in the IAS snap-in. You can also use the **netsh ras add registeredserver** command.

If IAS is a member of Windows NT 4.0 domain but has to authenticate users against a trusted Active Directory domain, it is not able to gain access to Active Directory because its computer account cannot become a member of the RAS and IAS Servers security group. In this case, verify that the **Everyone** group is added to the Pre-Windows 2000 Compatible Access group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, issue the **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** command on a domain controller computer and then restart the domain controller computer.

**The Windows 2000 Native Domain Functional Level**
From the network access management perspective, the following benefits are available in the Windows 2000 native domain

functional level:

- Full ability to manage remote access permissions through groups. An administrator can use the universal group feature to create a single policy for users in different domains. Nested groups can be used to organize extremely large numbers of users into smaller groups for better management.

- An ability to connect remote network to office network. You can specify routes for the remote network through Static Routes.

- Support for User Principal Names (UPNs).

Users can have the same UPN regardless of what domain the user belongs to. This indirection provides scalability that might be required in organizations that have large number of domains.

The following is a detailed list of authentication and remote management features available for an IAS server that is a member of a domain with a Windows 2000 domain functional level.

- User or computer account dial-in properties, which include:

    - All remote access permissions options are available, including **Allow access**, **Deny access**, and **Control access through Remote Access Policy**.

    - Caller-ID

    - Callback options

    - Static IP address

    - Static routes

- Support for UPNs and Universal Groups

- Support for EAP-TLS and PEAP.

In order for the IAS server to access user account dial-in properties stored in Active Directory, IAS must run in the security context of a computer account that is a member of the **RAS and IAS Servers** security group. This assignment can be implemented through the Active Directory Users and Computers snap-in or by registering the IAS server in the IAS snap-in. You can also use the **netsh ras add registeredserver** command.

**The Windows 2000 Mixed Domain Functional Level**

Domains with a Windows 2000 mixed functional level are mainly used for migration from Windows NT 4.0 to Windows Server 2003 or Windows 2000. For IAS, a Windows 2000 mixed domain acts exactly like a Windows NT 4.0 domain.

For an IAS server that is a member of a Windows 2000 mixed domain, the following authentication and remote access management features are available:

- User account dial-in properties

- Remote access permissions include only **Allow access** and **Deny access**

  **Note**

    - Without the **Control access through Remote Access Policy** option, policy-based management, using groups becomes more difficult because the user's remote access permission overrides remote access policy permissions.

- Callback options

Just as in Windows 2000 native domains, in order for the IAS server to access user account dial-in properties stored in Active

Directory, the Internet Authentication service must run in the security context of a computer account that is a member of the **RAS and IAS Servers** security group. This assignment can be implemented through the Active Directory Users and Computers snap-in or by registering the IAS server in the IAS snap-in. You can also use the **netsh ras add registeredserver** command.

If IAS is a member of Windows NT 4.0 domain but has to authenticate users against a trusted Active Directory domain, it is not able to gain access to Active Directory because its computer account cannot become a member of the RAS and IAS Servers security group. In this case, verify that the Everyone group is added to the Pre-Windows 2000 Compatible Access group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, issue the **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** command on a domain controller computer and then restart the domain controller computer.

**Stand-Alone IAS Servers**
Stand-alone IAS servers can be used in very small networks with no domains. All the users need to be defined in the local accounts database of a stand-alone server.

The following authentication features are available on a stand-alone IAS server:

- Dial-in User Account Properties
    - All remote access permission options, including **Allow access**, **Deny access**, and **Control access through Remote Access Policy**
- Caller-ID
- Callback options
- Static IP address
- Static routes

Support for UPNs, Universal Groups, and EAP-TLS is not available on a stand-alone IAS server computer.

User account dial-in properties can be administered through the properties of the Incoming Connections connection in the Network Connections folder or through the Local Users and Groups snap-in.

**Differences in IAS Implementations**
A previous version of IAS was released with the Windows NT 4.0 Option Pack. The following section describes the differences in behavior between the two versions.

**Windows NT 4.0 IAS**

- If no domain name is specified during authentication, the IAS server authenticates the user against only the local SAM database.
- IAS does not use the callback permissions for all user objects. IAS
- IAS log files are written in ASCII.

**Windows 2000 and Windows Server 2003 IAS**
IAS resolves a user name with no domain specified by using the following sequence:

- IAS determines a default domain from the registry, if one is specified there.
    - If the IAS server is a member of a domain, IAS authenticates the user against that domain.
    - If the IAS server is not a member of a domain, IAS authenticates the user against the local SAM database.
- IAS uses the callback permissions for all user objects.
- IAS log files are multi-language and are written in UTF-8.

**IAS as a RADIUS Proxy**

IAS can be used as a RADIUS proxy to provide the routing of RADIUS messages between RADIUS clients (NASs) and RADIUS servers that perform user authentication, authorization, and accounting for the connection attempt. When used as a RADIUS proxy, IAS is a central routing point through which RADIUS access and accounting messages flow. IAS records information in an accounting log about the messages that are forwarded.

The figure, "IAS as a RADIUS Proxy," illustrates how the feature can be used to relays messages between RADIUS clients (access servers) and either RADIUS servers or another RADIUS proxy.

**IAS as a RADIUS Proxy**




IAS acts as a RADIUS proxy for incoming RADIUS messages when the message matches a connection request policy that is configured to forward the request to a remote RADIUS server group.

When IAS is a RADIUS proxy between a RADIUS client and a RADIUS server, the messages that RADIUS sends for network access are forwarded as follows:

1. Access servers, such as dial-up network access servers, VPN servers, and wireless access points, receive connection requests from access clients.

2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the IAS server that is being used as the IAS RADIUS proxy.

3. The IAS RADIUS proxy receives the Access-Request message and, based on the locally configured connection request policies, determines where to forward the Access-Request message.

4. The IAS RADIUS proxy forwards the Access-Request message to the appropriate RADIUS server.

5. The RADIUS server evaluates the Access-Request message.

6. If required, the RADIUS server sends an Access-Challenge message to the IAS RADIUS proxy, where it is forwarded to the access server. The access server processes the challenge with the access client and sends an updated Access-Request to the IAS RADIUS proxy, where it is forwarded to the RADIUS server.

7. The RADIUS server authenticates and authorizes the connection attempt.

8. If the connection attempt is both authenticated and authorized, the RADIUS server sends an Access-Accept message to the IAS RADIUS proxy, where it is forwarded to the access server.

9. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends an Access-Reject message to the IAS RADIUS proxy, where it is forwarded to the access server.

10. The access server completes the connection process with the access client and sends an Accounting-Request message to the IAS RADIUS proxy. The IAS RADIUS proxy logs the accounting data and forwards the message to the RADIUS server.

11. The RADIUS server sends an Accounting-Response to the IAS RADIUS proxy, where it is forwarded to the access server.

**Access-Request Messages**
Forwarding a RADIUS Access-Request message from a RADIUS client to a RADIUS server results in the following changes to the RADIUS Access-Request message:

- Attribute manipulation rules are applied and the RADIUS attribute is modified according to the configured find and replace rules.

- All RADIUS message fields for which the RADIUS shared secret is used are recalculated. The recalculation is performed by using the shared secret of the IAS server and the remote RADIUS server to which the message is being forwarded. These RADIUS attributes include the Message Authenticator (also known as the Signature attribute), User-Name, Tunnel-Password, and MS-CHAP -MPPE-Keys. Additionally, the Authenticator field in the RADIUS header is recalculated.

- If the Access-Request message contains a Proxy-State attribute indicating that the message has been forwarded by another RADIUS proxy, the contents of the Proxy-State attribute are saved in a proxy session table and the Proxy-State attribute is rewritten with a value determined by the IAS server. The new value of the Proxy-State attribute is stored in a proxy session table.

When forwarding the Access-Request message, the set of additional RADIUS attributes as configured on the **Advanced** tab from the profile settings of the matching connection request policy is stored in the proxy session table. These attributes are not actually added to the Access-Request message but are saved for the response to the Access-Request message.

**Access-Accept Messages**
The proxy-state attribute in the Access-Accept message is used to find the corresponding entry in the proxy session table. The entry in the proxy session table indicates the IP address of the client to which the message is forwarded, the Proxy-State attribute of the original Access-Request message, and any additional RADIUS attributes to add. If an entry is not found, the Access-Accept message is discarded.

Forwarding a RADIUS Access-Accept message from a RADIUS server to a RADIUS client results in the following changes to the RADIUS Access-Accept message:

- All RADIUS message fields for which the RADIUS shared secret is used are recalculated by using the shared secret of the IAS server and the RADIUS client to which the message is being forwarded. These RADIUS attributes include the Message Authenticator (also known as the Signature attribute), Tunnel-Password, and MS-CHAP -MPPE-Keys. Additionally, the Authenticator field in the RADIUS header is recalculated.

- If the proxy session table entry contains an original Proxy-State attribute, the Proxy-State attribute in the Access-Accept message is replaced with the original Proxy-State attribute of the Access-Request message.

- If the proxy session table entry contains an additional RADIUS attributes, these RADIUS attributes are added to the Access-Accept message. If any of these RADIUS attributes are already present, their values are overwritten with new values.

**Access-Reject Messages**
The proxy-state attribute in the Access-Reject message is used to lookup the corresponding entry in the proxy session table. The entry in the proxy session table indicates the IP address of the client to which the message is to be forwarded, the Proxy-State attribute of the original Access-Request message, and any additional RADIUS attributes to add. If an entry is not found, the Access-Reject message is discarded.

Forwarding a RADIUS Access-Reject message from a RADIUS server to a RADIUS client results in the following changes to the RADIUS Access-Reject message:

- All RADIUS message fields for which the RADIUS shared secret are used are recalculated by using the shared secret of the IAS server and the RADIUS client to which the message is being forwarded. These RADIUS attributes include the Message Authenticator (also known as the Signature attribute), Tunnel-Password, and MS-CHAP -MPPE-Keys. Additionally, the Authenticator field in the RADIUS header is recalculated.

- If the proxy session table entry contains an original Proxy-State attribute, the Proxy-State attribute in the Access-Reject message is replaced with the Proxy-State attribute of the original Access-Request message.

- If the proxy session table entry contains an additional RADIUS attributes, these RADIUS attributes are added to the Access-Reject message. If any of these RADIUS attributes are already present, their values are overwritten with new values.

**Accounting-Request Messages**
When a RADIUS proxy forwards a RADIUS Accounting-Request message from a RADIUS client to a RADIUS server, it results in the following changes:

- Attribute manipulation rules are applied and the RADIUS attribute is modified according to the configured find and replace rules.

- All RADIUS message fields for which the RADIUS shared secret are used are recalculated using the shared secret of the IAS server and the remote RADIUS server to which the message is being forwarded. These RADIUS attributes include the Message Authenticator (also known as the Signature attribute) and User-Name. Additionally, the Authenticator field in the RADIUS header is recalculated.

- If the Accounting-Request message contains a Proxy-State attribute indicating that the message has been forwarded by another RADIUS proxy, the contents of the Proxy-State attribute are saved in a proxy session table and the Proxy-State attribute is rewritten with a value determined by the IAS server. The new value of the Proxy-State attribute is stored in a proxy session table.

When forwarding the Accounting-Request message, the set of additional RADIUS attributes as configured on the **Advanced** tab from the profile settings of the matching connection request policy is stored in the proxy session table. These attributes are not actually added to the Accounting-Request message but are saved for the response to the Accounting-Request message.

**Accounting-Response Messages**
The Proxy-State attribute in the Accounting-Response message is used to lookup the corresponding entry in the proxy session table. The entry in the proxy session table indicates the IP address of the client to which the message is to be forwarded, the Proxy-State attribute of the original Accounting-Request message, and any additional RADIUS attributes to add. If an entry is not found, the Accounting-Response message is discarded.

When a RADIUS proxy forwards a RADIUS Accounting-Response message from a RADIUS server to a RADIUS client, it results in the following changes to the RADIUS Accounting-Response message:

- All RADIUS message fields for which the RADIUS shared secret are used are recalculated using the shared secret of the IAS server and the RADIUS client to which the message is being forwarded, such as the Message Authenticator (also known as the Signature attribute). Additionally, the Authenticator field in the RADIUS header is recalculated.

- If the proxy session table entry contains an original Proxy-State attribute, the Proxy-State attribute in the Accounting-Response message is replaced with the original Proxy-State attribute of the Accounting-Request message.

- If the proxy session table entry contains additional RADIUS attributes, these RADIUS attributes are added to the Accounting-Response message. If any of these RADIUS attributes are already present, their values are overwritten with new values.

**Connection Request Processing**
To determine whether a specific connection attempt request or an accounting message received from a RADIUS client should be processed locally or forwarded to another RADIUS server, the IAS server uses connection request processing. Connection-request processing is a combination of the following:

- Connection request policies that determine, for any incoming RADIUS request message, whether the message is processed locally or forwarded to another RADIUS server.

- Remote RADIUS server groups that specify multiple RADIUS servers when an incoming RADIUS request message is forwarded.

**Connection Request Policies**

Connection request policies are sets of conditions and profile settings that give network administrators flexibility in configuring how incoming authentication and accounting request messages are handled by the IAS server. Using connection request policies, you can create a series of policies so that some RADIUS request messages sent from RADIUS clients are processed locally (IAS is being used as a RADIUS server) and other types of messages are forwarded to another RADIUS server (IAS is being used as a RADIUS proxy). This capability allows IAS to be implemented in many new RADIUS deployments.

With connection request policies, you can use IAS as a RADIUS server or as a RADIUS proxy, based on the time of day and day of the week, by the realm name in the request, by the type of connection being requested, by the IP address of the RADIUS client, and so on.

It is important to remember that with connection request policies, a RADIUS request message is processed only if the settings of the incoming RADIUS request message match at least one of the connection request policies. For example, if the settings of an incoming RADIUS Access-Request message do not match at least one of the connection request policies, an Access-Reject message is sent.

**Elements of a Connection Request Policy**

A connection request policy is a named rule that consists of the following elements:

- Conditions
- Profile

Connection request policies, like remote access policies, operate on the following principle:

- If the RADIUS request message matching all of the conditions of a connection request policy, the profile settings of the policy are applied to determine how the message is processed.

**Policy Conditions**

Connection request policy conditions are one or more RADIUS attributes that are compared to the attributes of the incoming RADIUS request message. If there are multiple conditions, then all conditions must match the attributes of the incoming RADIUS message so that the RADIUS request message matches the policy.

The following table, "Connection Request Policy Condition Attributes," shows the condition attributes that you can set for a connection request policy.

**Connection Request Policy Condition Attributes**

| Attribute name | Description |
|---|---|
| Called Station ID | The phone number of the NAS. This attribute is a character string. You can use pattern matching syntax to specify area codes. |
| Calling Station ID | The phone number used by the caller (the access client). This attribute is a character string. You can use pattern matching syntax to specify area codes. |
| Client Friendly Name | The name of the RADIUS client computer that is requesting authentication. This attribute is a character string. You can use pattern matching syntax to specify client names. |
| Client IP Address | The IP address of the NAS (the RADIUS client). This attribute is a character string. You can use pattern matching syntax to specify IP networks. |
| Client Vendor | The vendor of the NAS that is requesting authentication. You can use this attribute to configure separate policies for different NAS manufacturers that are configured as RADIUS clients to an IAS server. |
| Day and Time Restrictions | The day of the week and the time of day of the connection attempt. The day and time is relative to the date and time of the IAS server. |
| Framed Protocol | The type of framing for incoming packets. Examples are PPP, SLIP, Frame Relay, and X.25. |

| | |
|---|---|
| NAS Identifier | The name of the NAS. This attribute is a character string. You can use pattern matching syntax to specify NAS identifiers. |
| NAS IP Address | The IP address of the NAS (the RADIUS client). This attribute is a character string. You can use pattern matching syntax to specify IP networks. |
| NAS Port Type | The type of media used by the caller. Examples are analog phone lines (known as async), ISDN, and tunnels or virtual private networks (known as virtual). |
| Remote RADIUS to Windows User Mapping | Specifies that authorization is performed locally with remote access policies and user account properties, and authentication is performed against a remote RADIUS server. For example, visitors to your network are authenticated by their organization's RADIUS server, while a Windows user account created for the visitor at your organization is used for authorization. This configuration allows you to provide visitors with limited access to your network resources, while authenticating them against the user account database at their own organization. |
| Service Type | The type of service being requested. Examples include framed (such as PPP connections) and login (such as Telnet connections). For more information about RADIUS service types, see RFC 2865. |
| Tunnel Type | The type of tunnel being created by the requesting client. Tunnel types include the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol (L2TP) used by Windows Server 2003 and Windows 2000 remote access clients and demand-dial routers. |
| User Name | The user name used by the access client in the RADIUS message. This attribute is a character string. You can use pattern matching syntax to specify client names. |

**Policy Profile**

A connection request policy profile is a set of properties that are applied to an incoming RADIUS message. A connection request policy profile consists of the following groups of properties:

- Authentication
- Accounting
- Attribute Manipulation
- Advanced

**Authentication**

You can set the following authentication options used for RADIUS Access-Request messages:

- Authenticate requests on this server

  Use a Windows NT 4.0 domain or the Active Directory directory service, or the local SAM on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, for both authentication and the matching remote access policy and user account dial-in properties for authorization. In this case, the IAS server is being used as a RADIUS server.

- Forward requests to another RADIUS server in a remote RADIUS server group

  Forward the Access-Request message to another RADIUS server in a specified remote RADIUS server group. If the IAS server receives a valid Access-Accept message that corresponds to the Access-Request message, the connection attempt is considered authenticated and authorized. In this case, the IAS server is being used as a RADIUS proxy.

- Accept the connection attempt without performing authentication or authorization

  Do not verify the user credentials and do not check user account properties or remote access policies to verify that the user is authorized to access network resources. An Access-Accept message is immediately sent to the RADIUS client by IAS without performing authentication or authorization. This setting is used for some types of compulsory tunneling where the access client is tunneled before the users credentials are authenticated.

This authentication option cannot be used when the authentication protocol of the access client is PEAP, MS-CHAP v2, or EAP-TLS, which provide mutual authentication. In mutual authentication, the access client proves that it is a valid access client to the authenticating server (the IAS server), and the authenticating server proves that it is a valid authenticating server to the access client. When this authentication option is used, the Access-Accept message is returned. However, the authenticating server does not provide validation to the access client and mutual authentication fails.

**Accounting options**

You can set the following accounting options that are used for RADIUS Accounting-Request messages:

- Forward accounting information in a specific remote RADIUS server group

  Pass the accounting request to another RADIUS server in a specified remote RADIUS server group. In this case, the IAS server is acting as a RADIUS proxy.

  **Note**

  - IAS always records the accounting information for Accounting-Request messages in local log files on the basis of remote access logging settings.

**Attribute manipulation rules**

You can configure a set of find and replace rules that manipulate the text strings of one of the following attributes:

- User Name
- Called Station ID
- Calling Station ID

Attribute manipulation rule processing occurs for one of the preceding attributes before the RADIUS message is subject to authentication and accounting settings.

You can configure attribute manipulation rules in the profile of a connection request policy. When you configure attribute manipulation rules, you specify text that you want IAS to find in the attribute and you specify replacement text. While processing attribute manipulation rules, IAS seeks the text you specify in the rule. If the text is found in the attribute, IAS deletes the text and replaces it with the replacement text that you specify in the rule.

Configuring an attribute manipulation rule for the User Name attribute in IAS for Windows Server 2003 is equivalent to configuring realm replacement rules for IAS in Windows 2000.

If you are using the MS-CHAP v2 authentication protocol, you cannot manipulate the User Name attribute if the connection request policy is used to forward the RADIUS message. The only exception occurs when a backslash (\) character is used and the manipulation only affects the information to the left of this character. A backslash character is typically used to indicate a domain name (the information to the left of the backslash character) and a user account name within the domain (the information to the right of the backslash character). In this case, only attribute manipulation rules that modify or replace the domain name are allowed.

**Note**

- Find and replace rules apply only to a single attribute. You cannot configure find and replace rules for each attribute or add to the list of attributes available for manipulation.

**Advanced properties**

You can set advanced properties to specify the series of RADIUS attributes that are:

- Added to the RADIUS response message when the IAS server is being used as a RADIUS authentication or accounting server.

When there are attributes specified on both a remote access policy and the connection request policy, the attributes that are sent in the RADIUS response message are the combination of the two sets of attributes.

- Added to the RADIUS message when the IAS server is being used as a RADIUS authentication or accounting proxy. If the attribute already exists in the message that is forwarded, it is replaced with the value of the attribute specified in the connection request policy.

**Default Connection Request Policy**

A default connection request policy named **Use Windows authentication for all users** is created when you install IAS. This policy has the following configuration:

- The Day-and-Time-Restrictions condition is set to all times and all days.
- Authentication is set to authenticate requests on this server.
- Accounting is not set to forward accounting information to a remote RADIUS server group.
- Attribute manipulation rules are not configured.
- Advanced attributes are not configured.

The default connection request policy uses IAS as a RADIUS server. To configure an IAS server to act as a RADIUS proxy, you must also configure a remote RADIUS server group. You can create a new remote RADIUS server group while you are creating a new connection request policy with the New Connection Request Policy Wizard. You can either delete the default connection request policy or verify that the default connection request policy is the last policy processed.

**Realm Names**

When an access client sends password-based user account credentials during connection attempts with a NAS, a user name and password are included. The user name consists of two elements:

- Identification of the user account name.
- Identification of the user account location on the network.

For example, for the user name user1@microsoft.com, user1 is the user account name and microsoft.com is the location of the user account. The identification of the location of the user account is known as a realm. There are two forms of realm names:

- The realm name can be a prefix.

  For example, Microsoft\user1, where Microsoft is the name of a Windows NT domain.

- The realm name can be a suffix.

  For example, user1@microsoft.com, where microsoft.com is either a DNS domain name or the name of an Active Directory domain.

When initiating a network connection, users can be required to provide the realm name as either a prefix or a suffix when typing account credentials in the **Connect** dialog box. For example, when the owner of the account user1 connects to microsoft.com with a VPN connection, they can be required to type either MICROSOFT\user1 or user1@microsoft.com.

If you use Connection Manager Administration Kit (CMAK) to create Connection Manager profiles for distribution to users, the profile can be configured to automatically append the realm name to the user account name during connection attempts by users.

For example, you can specify a realm name and syntax when you create a custom dialing package with CMAK so that the user must specify only the user account name.

In this circumstance, if the realm name is Microsoft, the correct syntax is to specify the realm name as a prefix. When a user logs on with the account name **user1**, the user enters the string “user1,” but does not enter a realm name. The Connection Manager profile automatically appends the realm name to **user1** and sends the following user name to the NAS: **MICROSOFT\user1**.

For more information, see “Pattern Matching Syntax” later in this section.

**RADIUS User-Name Attribute**
During the authentication phase of the connection attempt, the user name is passed from the access client to the access server. The access server creates an Access-Request message, using the user name as the value for the RADIUS attribute User-Name. The access server, which is configured as a RADIUS client to an IAS server or IAS proxy, sends the Access-Request message containing the User-Name attribute to the IAS server or proxy.

The IAS server evaluates the Access-Request message against the set of configured connection request policies. Connection request policies are a set of conditions and profile settings that determine, for any incoming RADIUS request message, whether the message is processed locally or forwarded to another RADIUS server. The realm name, contained in the RADIUS User-Name attribute, can be specified in the connection request policy and assist IAS in determining how to route connection requests.

**Using Realm Names to Route Connection Requests**
You can configure a set of connection request policies that are specific to the realm name within the User-Name attribute of incoming messages. This allows to you create routing rules that forward RADIUS messages with a specific realm name to a specific set of RADIUS servers (a remote RADIUS server group) when IAS is used as a RADIUS proxy.

Before the RADIUS message is either processed locally (when IAS is being used as a RADIUS server) or forwarded to another RADIUS server (when IAS is being used as a RADIUS proxy), the User-Name attribute in the message might be modified by attribute manipulation rules that are configured in the profile of the first matching connection request policy. Attribute manipulation rules for the User-Name attribute might change the following:

- Remove the realm name from the user name (also known as realm stripping).

  For example, the user name user1@microsoft.com is changed to user1.

- Change the realm name but not its syntax.

  For example, the user name user1@microsoft.com is changed to user1@wcoast.microsoft.com.

- Change the syntax of the realm name.

  For example, the user name microsoft\user1 is changed to user1@microsoft.com.

Attribute manipulation rules are configured in the Attribute tab on the profile of a connection request policy.

IAS attribute manipulation rules use regular expression syntax.

After the User-Name attribute is modified according to attribute manipulation rules, additional settings on the profile of the first matching connection request policy are used by IAS to determine whether to process the message locally or to forward the message to a remote RADIUS server group:

- When IAS is configured as a RADIUS server, the Access-Request message is processed locally

- When IAS is configured as a RADIUS proxy, the Access-Request message is forwarded to another RADIUS server that is configured as a member of a remote RADIUS server group.

**Pattern Matching Syntax**
IAS supports regular expressions for pattern matching, which is widely used in UNIX environments. You can use this syntax to specify the conditions of remote access policy attributes and RADIUS realms, as specified in the following table, “Pattern Matching Syntax Characters.”

### Pattern Matching Syntax Characters

| Character | Description | Example |
|---|---|---|
| \ | Marks the next character as a character to match. | /”n”/ matches the character “n”. The seque”n”ce /\”n”/ matches a line feed or newline character. |

| | | |
|---|---|---|
| ^ | Matches the beginning of the input or line. | |
| $ | Matches the end of the input or line. | |
| * | Matches the preceding character zero or more times. | /zo*/ matches either “z” or “zoo.” |
| + | Matches the preceding character one or more times. | /zo+/ matches “zoo” but not “z.” |
| ? | Matches the preceding character zero or one times. | /a?ve?/ matches the “ve” in “never.” |
| . | Matches any single character except a newline character. | |
| **(** *pattern* **)** | Matches *pattern* and remembers the match. To match ( ) (parentheses), use "\(" or "\)." | |
| *x* **\|** *y* | Matches either *x* or *y*. | /z\|food?/ matches “zoo” or “food.” |
| **{** *n* **}** | Matches exactly *n* times (*n* is a nonnegative integer). | /o{2}/ does not match the “o” in “Bob,” but matches the first two instances of the letter “o” in “foooood.” |
| **{** *n* **,}** | Matches at least *n* times (*n* is a nonnegative integer). | /o{2,}/ does not match the “o” in “Bob” but matches all of the instances of the letter “o” in “foooood.” /o{1,}/ is equivalent to /o+/. |
| **{** *n* **,}** | Matches at least *n* times (*n* is a nonnegative integer). | /o{2,}/ does not match the “o” in “Bob” but matches all of the instances of the letter “o” in “foooood.” /o{1,}/ is equivalent to /o+/. |
| **{** *n* **,** *m* **}** | Matches at least *n* and at most *m* times (*m* and *n* are nonnegative integers). | /o{1,3}/ matches the first three instances of the letter “o” in “foooood.” |
| **[** *xyz* **]** | Matches any one of the enclosed characters (a character set). | /[abc]/ matches the “a” in “plain.” |
| **[^** *xyz* **]** | Matches any characters that are not enclosed (a negative character set). | /[^abc]/ matches the “p” in “p”lain. |
| **\b** | Matches a word boundary (for example, a space). | /ea*r\b/ matches the “er” in “never early.” |
| **\B** | Matches a nonword boundary. | /ea*r\B/ matches the “ear” in “never early.” |
| **\d** | Matches a digit character (equivalent to [0-9]). | |
| **\D** | Matches a nondigit character (equivalent to [^0-9]). | |
| **\f** | Matches a form feed character. | |
| **\n** | Matches a line feed character. | |
| **\r** | Matches a carriage return character. | |
| **\s** | Matches any white space character including space, tab, and form feed (equivalent to [ \f\n\r\t\v]). | |
| **\S** | Matches any non-white space character (equivalent to [^ \f\n\r\t\v]). | |
| **\t** | Matches a tab character. | |
| **\v** | Matches a vertical tab character. | |

| | | |
|---|---|---|
| **\w** | Matches any word character, including underscore (equivalent to [A-Za-z0-9_]). | |
| **\W** | Matches any nonword character, excluding underscore (equivalent to [^A-Za-z0-9_]). | |
| **** *num* | Refers to remembered matches (?*num*, where *num* is a positive integer). For example, \1 replaces what is stored in the first remembered match. This option can be used only in the **Replace** text box when configuring attribute manipulation. | |
| **/** *n* **/** | Allows the insertion of ASCII codes into regular expressions (?*n*, where *n* is an octal, hexadecimal, or decimal escape value). | |

**Examples for Remote Access Policy Attributes**
The following examples illustrate the use of the pattern matching syntax to specify remote access policy attributes:

- To specify all phone numbers within the 899 area code, the syntax is: 899.*
- To specify a range of all IP addresses that begin with 192.168.1, the syntax is: 192\.168\.1\..+

**Examples for Manipulation of the Realm Name in the User-Name Attribute**
The following examples illustrate the use of the pattern matching syntax to manipulate realm names for the User-Name attribute. In the first example, the realm name is removed before the message is forwarded by an IAS proxy to an IAS server. In the other examples, additional variations of realm name manipulation are demonstrated.

**To remove the realm portion of the User-Name attribute**
Find: @microsoft\.com

Replace:

**To replace user@example.microsoft.com with example.microsoft.com\user**
Find: (.*)@(.*)

Replace: $2\$1

**To replace domain\user with specific_domain\user**
Find: (.*)\\(.*)

Replace: specific_domain\$2

**To replace user with user@specific_domain**
Find: $

Replace: @specific_domain

**Example for RADIUS Message Forwarding by a Proxy Server**
You can create routing rules that forward RADIUS messages with a specific realm name to a specific set of RADIUS servers when IAS is used as a RADIUS proxy. A recommended syntax for routing requests based on realm name follows:

- NetBIOS name: WCOAST
- Pattern: ^wcoast\\

In the following example, wcoast.microsoft.com is a unique user principal name (UPN) suffix for the DNS or Active Directory domain wcoast.microsoft.com. Using the supplied pattern, the IAS proxy can route messages based on domain NetBIOS name or UPN suffix as follows:

- NetBIOS name: WCOAST
- UPN suffix: wcoast.microsoft.com
- Pattern: ^wcoast\\|@wcoast\.microsoft\.com$

**Remote RADIUS Server Groups**
A remote RADIUS server group is a named group that contains one or more RADIUS servers to which messages are forwarded by the RADIUS proxy. When IAS is being used as a RADIUS proxy for RADIUS request messages, a remote RADIUS server group must be specified. You can specify various settings to determine the order in which the servers are used by the proxy server. You can also specify how the proxy distributes the RADIUS messages across the servers in the group.

Each server in the group has the following settings:

- Name or address

  Each group member must have a unique name within the group. The name can be an IP address or a name that can be resolved to its IP address.

- Authentication and accounting

  When IAS is used as a RADIUS proxy, it is acting as a RADIUS client to another RADIUS server. Therefore, each group member must be configured to send RADIUS messages to the correct UDP port that is used by the RADIUS server for RADIUS traffic. Additionally, each group member must be configured for the correct shared secret. The default authentication port is 1812. The default accounting port is 1813.

- Load balancing

  A priority setting is used to indicate which member of the group is the primary server (the priority is set to 1). For group members that have the same priority, a weight setting is used to calculate how often RADIUS messages are sent to each of them. You can use additional settings to configure the way in which the IAS server detects when a group member first becomes unavailable and when it becomes available after it has been determined to be unavailable.

After a remote RADIUS server group is configured, it can be specified in the authentication and accounting settings of a connection request policy. Because of this, it is recommended that you configure a remote RADIUS server group first. Next, you can configure the connection request policy to use the newly configured remote RADIUS server group. Alternately, you can use the New Connection Request Policy Wizard to create a new remote RADIUS server group while you are creating the connection request policy.

**Authentication Across Forests**
With IAS as a RADIUS proxy configured to forward messages to remote RADIUS server groups in different forests, you can authenticate users across forests. For example, if you have NASs in a perimeter network that serves users from two forests, you can configure the NASs as RADIUS clients to a RADIUS proxy server. At the proxy server you can configure two remote RADIUS server groups, each group containing the RADIUS servers in one forest. When a user connects to a NAS, the RADIUS proxy examines the realm portion of the user name and forwards the message to the remote RADIUS server group for authentication in the Active Directory user accounts database for that forest.

**Using Remote RADIUS Server Groups to Focus Accounting**
In some circumstances, you might want to assign the task of logging RADIUS accounting data to one or more remote RADIUS server groups, while assigning the task of processing authentication requests to other remote RADIUS server groups. This configuration is especially useful in organizations that regularly query RADIUS accounting data for network access troubleshooting and report creation. If you use IAS as a RADIUS proxy to route accounting data to one or more remote RADIUS server groups, you can configure the RADIUS servers in the remote RADIUS server groups to log to a central SQL Server database. Help desk and other applications can then query this database for comprehensive data views of network access sessions. Fore more information about IAS SQL Server logging, see “IAS Accounting” earlier in this section.

**Processing a Connection Request**
When an IAS server receives a valid RADIUS message from a RADIUS client, the RADIUS message is processed based on the following:

1. The first policy in the ordered list of connection request policies is checked. If there are no policies and the RADIUS message is an Access-Request message, an Access-Reject message is immediately sent.

2. If all the conditions of the policy do not match the RADIUS message, then go to next policy. If there are no more policies and the RADIUS message is an Access-Request message, an Access-Reject message is immediately sent.

3. If all the conditions of the policy match the RADIUS message, then apply the attribute manipulation rules configured in the policy.

4. If the RADIUS message is an Access-Request message, check the authentication settings on the matching connection request policy profile:

    - If the authentication setting is to authenticate using this server, use either the local SAM in Windows Server 2003 or Active Directory for authentication and the local remote access policies and user account dial-in properties for authorization.

    - If the authentication setting is to forward the request to a specific remote RADIUS server group, forward the Access-Request message containing the attribute(s) changed by the attribute manipulation rules to the appropriate RADIUS server in the group based on the priority and weight settings.

    - If the authentication setting is to accept the connection attempt without performing authentication or authorization, send back an Access-Accept message immediately.

5. If the RADIUS message is an accounting message, check the accounting settings on the profile of the matching policy:

    - If the accounting settings specify that accounting information is to be sent to a member of a remote RADIUS server group, forward the accounting message containing the attribute(s) changed by the attribute manipulation rules to the appropriate RADIUS server in the group on the basis of priority and weight settings.

    - Based on the settings for remote access logging, log the accounting message in the IAS accounting log.

**Using IAS as a RADIUS Proxy for Load Balancing**

IAS proxies, placed between access servers and IAS servers, can be used to balance the load of a large volume of authentication traffic. The following figure, "RADIUS Proxy Load Balancing," illustrates how IAS proxies are used for load balancing authentication requests to multiple IAS servers.

**RADIUS Proxy Load Balancing**




By using IAS as a RADIUS proxy, you can evenly distribute authentication, authorization, and accounting traffic across all the IAS servers in the organization. This load balancing solution is easily scaled up and out as your organization grows, and as you add more NASs and IAS servers.

Additionally, when you use IAS as a RADIUS proxy for load balancing, there is a consistent scheme for failure detection and RADIUS server failover.

Primary characteristics of load balancing include:

1. Active Directory is configured for user accounts and groups. Ensure that all users that are making network access connections have a corresponding user account. Also organize remote access by group to take advantage of group-based remote access policies.

2. Configure IAS as a RADIUS server on each domain controller. Configure IAS to read the user account properties in Active Directory, and configure the IAS proxies as RADIUS clients to the IAS servers.

3. Configure one or more IAS proxy servers. Configure NASs as RADIUS clients to the proxy servers. Use the New Remote RADIUS Server Group Wizard to create a custom remote RADIUS server group, and create a connection request policy that forwards RADIUS request messages to the IAS servers where the realm name matches the accounts in the domain. Also delete the default connection request policy.

4. Configure each access server to use the IAS proxies for the authentication, authorization, and accounting of network connections.

**Mapping Network Authentication and Authorization**

You can use IAS proxy to split network access authorization and authentication between two user account databases. For example, you can configure IAS to authenticate visitors (such as employees of a partner organization) by using their password-based credentials stored in the external partner organization user database, while performing authorization, using an account created for the visitor, in your local Active Directory user account database.

To split authentication and authorization between two user account databases, you must create a local user account for each visitor. The visitor account user name must map to an account with the same user namein the partner organization user database. When the visitor logs on to your network, the external user database provides authentication while the local user database and remote access policies are used to process authorization.

The forwarding of authentication requests to the partner organization RADIUS server is necessary only when visitor authentication is performed with password-based credentials. Visitors can also be authenticated by mapping certificates to local user accounts. If a partner organization uses certificate-based authentication, you can authenticate visitors by trusting the partner organization certification authority (CA) and mapping visitor certificates to local user accounts. Because the visitor is authenticated locally with a certificate, it is not necessary to forward authentication requests to the remote RADIUS server at the partner organization.

**Configuring Visitor User Accounts**

If you use password-based authentication, you must create a user account for each visitor. The visitor user account must be created so that the user name in the local database and the user name in the partner organization user database match exactly. To allow visitors access to your organization network while authenticating them at their partner organization, perform the following actions to configure visitor accounts:

1. Add the user principal name suffix to the domain.

2. Create an organizational unit that contains the user accounts for visitors.

3. Create user accounts in the organizational unit for visitors. When you create the user account, the user logon name must exactly match the visitor's user logon name as configured in the user account database at the partner organization. For example, if the visitor's user logon name at Microsoft Corporation is "user"@microsoft.com, designate "user" as the User logon namein the Active Directory Users and Computers snap-in.

4. Select the correct UPN suffix from the drop-down list of suffixes. For example, if in step one you added the UPN suffix "microsoft.com," select "@microsoft.com" from the list.

5. Configure other user account properties as needed. Make sure that the user account is configured in a way that allows the visitor to have network access, but also contains necessary security restrictions. For example, in user account properties, on the Dial-in tab, you can set Remote Access Permission (Dial-in or VPN) to Allow access.

**Configuring the IAS Proxy Server**

To allow visitors access to your organization network, you need to perform the following actions on your IAS proxy server:

- Configure the partner organization RADIUS server as a member of a Remote RADIUS server group.

  Run the New Remote RADIUS Server Group Wizard, create and name a custom group, and then add the partner organization RADIUS server (or servers) to the group. While adding each server, configure the shared secret (obtained from the partner organization) on the **Authentication/Accounting** tab of the **Add RADIUS server** dialog box.

- Configure a connection request policy for visitor access.

  To determine whether a specific connection attempt request or an accounting message received from a RADIUS client should be processed locally or forwarded to another RADIUS server, the IAS server uses connection request processing. When you allow visitor access with external authentication, you must configure a minimum of two connection request policies on your IAS proxy server. One policy allows your organization employees access, while denying access to other network access attempts. Another policy allows visitor access. When you configure your policy evaluation order, place the policy or policies for your organization users last, with the visitor policy processed first.

  To configure a connection request policy, run the New Connection Request Policy Wizard, create and name a custom policy, and then add **Policy Conditions**. You can configure as many conditions as you want, but you must include the User-Name attribute. Use wildcard and other characters to configure a value that must be matched by the user principal name (UPN) suffix your visitors will use when logging on. For example, if you want to authenticate all visitors from Microsoft Corporation with this connection request policy, you can type "^.@microsoft\.com$" in **User-Name**. In **Request Processing Method**, click **Edit Profile**, and then click the **Advanced** tab. Click **Add**, and then add the attribute **Remote-RADIUS-to-Windows-User-Mapping**, with the attribute value set to **True**. This attribute tells the IAS proxy server to authenticate users (whose properties match other policy conditions) against the RADIUS servers configured in the remote RADIUS server group. On the **Attribute** tab, you can configure replacement text for the User-Name attribute, as well as configure any other attributes that you require. On the **Authentication** tab, click **Forward requests to the following remote RADIUS server group for authentication**, and then choose the Remote RADIUS server group that can authenticate your visitors.

- Configure a remote access policy for visitor access.

  Because you can create remote access policies that are based on different types of connections and group membership, you can create a remote access policy used specifically to process network access attempts by visitors. At minimum, you need two remote access policies: one for your employees and a different one for visitors.

- Ensure that NASs are configured as RADIUS clients to this IAS proxy.

  When visitors log on to your network, they do so through NASs, such as wireless APs. These APs must be configured as RADIUS clients to your IAS proxy server.

After the IAS proxy is configured to forward authentication requests and user accounts for visitors are created in Active Directory, the password-based authentication request is forwarded to the partner organization RADIUS server when the visitor logs onto your organization network. If the user is authenticated, authorization is performed against your organization Active Directory user database. Active Directory returns user attributes (values configured for the user account, such as Allow access on the **Dial-in** tab of user properties), and IAS performs policy evaluation against these local user account attributes.

**Note**

- By configuring a local user account with the UPN of an account in the remote partner organization user account database, you map one account to the other. If the partner organization changes the user name of an account you have mapped, you must also change the user name in your user account database.

**Mapping with Certificate-Based Authentication**

Certificate-based authentication is more secure than password-based authentication. In addition, when you map certificates to user accounts instead of using password-based authentication methods, an authentication request is not forwarded to the partner organization RADIUS server and user account database. For these reasons, certificate mapping enhances security and can significantly reduce logon time for users.

A certificate is mapped to a user account in one of two ways: a single certificate is mapped to a single user account (one-to-one mapping) or multiple certificates are mapped to one user account (many-to-one mapping).

**One-to-One Mapping**

For one-to-one mapping of certificates to visitor user accounts, you must perform the following tasks:

- Create a user account for each visitor.

  Because the user account itself is not mapped to the user account at the partner organization, the user account names do

not have to exactly match the user account names in the partner organization user accounts database. Configuring the local account with the same user name, however, can make it easier to configure realm manipulation rules.

- Map the certificate to a user account.

- Perform cross-certification or authorize the partner organization CA as a qualified subordinate certification authority using a computer running Windows Server 2003, Enterprise Edition and Certificate Services.

  Performing cross-certification or authorizing the partner organization CA as a qualified subordinate CA is recommended for domains with a Windows Server 2003 domain functional level and clients running Windows XP. If your domains are Windows 2000 native, it is recommended that you use a certificate trust list (CTL) instead.

- Configure your IAS server or IAS proxy server

  Because authentication requests are not forwarded to external RADIUS servers when authentication is performed with certificates, you are not required to use IAS as a proxy server.

- Configure a connection request policy for visitor access with the New Connection Request Policy Wizard.

- Create a realm manipulation rule to map the partner organization realm name, which is contained in the user's certificate, to the user's local user account.

  For example, if user@tailspintoys.com is mapped to user@microsoft.com, the realm manipulation rule must replace "@tailspintoys.com" with "@microsoft.com" using pattern-matching syntax.

- Configure a remote access policy for visitor access.

- Ensure that NASs are configured as RADIUS clients to this IAS server.

When visitors log on to your network, they do so through NASs, such as wireless APs. These APs must be configured as RADIUS clients to your IAS proxy server or IAS server.

**Many-to-One Mapping**

Many-to-one certificate mapping allows you to map many certificates to one user account. After you have trusted the enterprise root CA of a partner organization, you can map all certificates issued by the partner organization CA to one account that you create in your local domain. This solution provides ease of management for many users, and eliminates the need to create an individual user account for every visitor to your organization.

For many-to-one mapping of certificates, you must perform the following steps:

- Perform cross-certification or authorize the partner organization CA as a qualified subordinate certification authority using a computer running Windows Server 2003, Enterprise Edition and Certificate Services.

Performing cross-certification or authorizing the partner organization CA as a qualified subordinate CA is recommended for domains with a Windows Server 2003 domain functional level and clients running Windows XP. If your domains are Windows 2000 native, it is recommended that you use a certificate trust list (CTL) instead.

- Create one user account and map the partner organization certificate to the account.

- Configure your IAS server or IAS proxy server

Because authentication requests are not forwarded to external RADIUS servers when authentication is performed with certificates, you are not required to use IAS as a proxy server.

- Configure a connection request policy for visitor access with the New Connection Request Policy Wizard.

- Create a realm manipulation rule to map the partner organization realm name, which is contained in the user's certificate, to the user's local user account.

For example, if user@tailspintoys.com is mapped to user@microsoft.com, the realm manipulation rule must replace @tailspintoys.com with "@microsoft.com," using pattern-matching syntax.

- Ensure that NASs are configured as RADIUS clients to this IAS server.

- Configure a remote access policy for visitor access.

When visitors log on to your network, they do so through NASs, such as wireless APs. These APs must be configured as RADIUS clients to your IAS proxy server or IAS server.

**Note**

- For strong security between your IAS proxy server and your partner organization RADIUS servers, you can use Internet Protocol security (IPSec).

# IAS Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- IAS Tools
- IAS Settings
- IAS Registry Entries
- Related Information


You can use Internet Authentication Service (IAS) tools, such as the IAS console and the Netsh commands for IAS, to configure your IAS servers. IAS settings include the use of the Message Authenticator attribute, shared secrets, domain name, and certificate templates. Because the correct configuration of your certificate templates for use with certificate-based authentication methods is critical to the success of your IAS deployment, this document provides step-by-step information about how to configure the Computer template and the IAS and RAS Servers template in Certificate Templates for Windows Server 2003.

**IAS Tools**
You can use the following tools to configure and manage your IAS servers.

**IAS Console**
After you install IAS, you can use the IAS console to configure Internet Authentication Service on the local computer. The console is launched through **Administrative Tools**, and can be used to configure Remote Authentication Dial-In User Service (RADIUS) clients, Remote Access Logging, Remote Access Policies, Connection Request Processing, and other server properties.

The IAS console can also be loaded as a snap-in from the Microsoft Management Console (MMC), allowing you to administer IAS on the local computer and on remote computers that are running IAS.

**Netsh Command Reference for AAAA**
You can use commands in the Netsh AAAA context to show and set the configuration of the authentication, authorization, accounting, and auditing (AAAA) database used by IAS and the Routing and Remote Access service. The AAAA database is also known as the IAS database (Ias.mdb). The primary use of commands in the Netsh AAAA context is to:

- Export the configuration of one IAS server, including registry keys and the IAS database, as a Netsh script using either the **dump** command or one of the **show** commands.
- Import the configuration to another IAS server using the **netsh exec** command, and a Netsh script that contains the **set config** command.

You can run the Netsh commands for AAAA from the Windows Server 2003 command prompt or from the command prompt for the Netsh AAAA context. For these commands to work at the Windows Server 2003 command prompt, you must type **netsh aaaa** before typing commands and parameters as they appear in the following syntax. There might be functional differences between Netsh context commands on Windows 2000 and Windows Server 2003.

The following table provides the command reference formatting legend, which describes how to interpret the syntax statement for each command.

**Command Reference Formatting**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply. |
| **Bold** | Elements that the user must type exactly as shown. |
| Between brackets [()] | Optional items. |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one. |

The following is the full command reference for the Netsh commands for AAAA.

**dump**
Displays the configuration of the IAS database file (Ias.mdb) as a Netsh command script.

**Syntax**

```
dump
```

**Remarks**
The **dump** command displays the Netsh command script that you can use to duplicate the configuration of the server running IAS or the Routing and Remote Access service on which the command is executed. The Netsh command script contains the configuration of the IAS server, including the registry keys and database file, in a compressed text format as a large data block. This large data block is used by the **set config** command within the script to import the configuration of a saved data block into an existing IAS database on the same or another computer using the **netsh exec** command. To save the Netsh command script to a file, type the following at a command prompt: **netsh aaaa show config >** *Path**File.txt*.

The **dump** command is not supported on computers running Windows 2000 Server.

**set config**
Configures the IAS server and IAS database with the registry keys and database imported from the specified data block.

**Syntax**

```
set config [type= {server_settings | clients | connection_request_policies |
logging | remote_access_policies ]blob= DataBlock}
```

**Parameters**
**blob**=*DataBlock*

Required. Specifies the file that contains the configuration of the IAS server, including registry keys and IAS database, in a compressed text format as output by the **dump** or **show** commands.

**type**=

Specifies the kind of data that you want to import from the *DataBlock* to the local server configuration.

**server_settings**

Specifies that the server configuration settings are imported from the *DataBlock* to the local server. The imported settings include Server Description, system event log settings, registry keys, and ports that are used specifically for the IAS service.

**clients**

Specifies that the settings for **RADIUS Clients**, as viewed in the IAS console, are imported to the local server from *DataBlock*.

**connection_request_policies**

Specifies that the IAS settings for **Connection Request Processing**, including **Connection Request Policies** and **Remote RADIUS Server Groups**, as viewed in the IAS console, are imported to the local server from *DataBlock*.

**logging**

Specifies that the IAS settings for **Remote Access Logging**, as viewed in the IAS console, are imported to the local server from *DataBlock*.

**remote_access_policies**

Specifies that the IAS settings for **Remote Access Policies**, as viewed in the IAS console, are imported to the local server from *DataBlock*.

**Remarks**
Running the **set config** command manually is not supported. This command is used only within a Netsh command script that is created by using the **dump** or **show** commands.

To view the version of the IAS database on which the Netsh script is being run, use the **show version** command. Scripts older than those created with Windows Server 2003 are supported.

The configuration of IAS servers running on products in the Windows 2000 Server family can be imported into products in the Windows Server 2003 family that support IAS with the **set config** command. The reverse, however, is not possible.

**show clients**
Displays the RADIUS client list for the local IAS server.

**Syntax**

```
show clients
```

**Remarks**
This command dumps the RADIUS client list of the IAS server on which the command is run. In the IAS console, this client list is displayed in **RADIUS Clients**.

The Netsh command script contains the local server settings, including registry keys and the IAS database (Ias.mdb), as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh**

**exec** command. To save the Netsh command script to a file, type **netsh aaaa show clients >***Path**File.txt* at a command prompt.

This command is not available on computers running Windows 2000 Server.

**show config**
Displays the configuration of the IAS database file as a Netsh command script.

**Syntax**

```
show config
```

**Remarks**
The **show config** command is equivalent to the **dump** command.

The **show config** command displays the Netsh command script that you can use to duplicate the configuration of the server running IAS or the Routing and Remote Access service on which the command is run. The Netsh command script contains the configuration of the IAS server, including registry keys and the IAS database, as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh exec** command. To save the Netsh command script to a file, type **netsh aaaa show config >***Path**File.txt* at a command prompt.

This command is not available on computers running Windows 2000 Server.

**show connection_request_policies**
Displays the Connection Request Processing policies configuration of the IAS server in script format.

**Syntax**

```
show connection_request_policies
```

**Remarks**
This command displays Connection Request Policies for the IAS server on which the command is run. In the IAS console, these policies are displayed in **Connection Request Processing** and include **Connection Request Policies** and **Remote RADIUS Server Groups**.

The Netsh command script contains the local server settings as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh exec** command. To save the Netsh command script to a file, type **netsh aaaa show connection_request_policies>***Path**File.txt* at a command prompt.

This command is not available on computers running Windows 2000 Server.

**show logging**
Displays the logging configuration for the local IAS server.

**Syntax**

```
show logging
```

**Remarks**
This command displays the Remote Access Logging configuration for the IAS server on which the command is run. In the IAS console, this information is displayed in **Remote Access Logging**.

The Netsh command script contains the local server settings as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh exec** command. To save the Netsh command script to a file, type **netsh aaaa show logging >***Path**File.txt* at a command prompt.

This command is not available on computers running Windows 2000 Server.

**show remote_access_policies**
Displays all objects within the remote access policy.

**Syntax**

```
show remote_access_policies
```

**Remarks**
This command displays the Remote Access Policies for the IAS server on which the command is run. In the IAS console, this information is displayed in Remote Access Policies.

The Netsh command script contains the local server settings as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh exec** command. To save the Netsh command script to a file, type **netsh aaaa show remote_access_policies >** *Path**File.txt* at a command prompt.

This command is not available on computers running Windows 2000 Server.

The output of this command is an encoded data block that is either displayed in the command prompt window or saved to a text file at the folder location you specify. The following registry keys and values are contained in the data block:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn\DefaultDomain\REG_SZ**

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy\Allow LM Authentication\REG_DWORD**

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy\Default User Identity\REG_SZ**

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy\User Identity Attribute\REG_DWORD**

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy\Override User-Name\REG_DWORD**

**show server_settings**
Displays the configuration of the local IAS server as a Netsh command script.

**Syntax**

```
show server_settings
```

**Remarks**
This command displays the configuration of server settings for the IAS server on which the command is run. These settings include: server description; settings for accounting and authorization events in the system event log; ports used by the service; and registry keys and their values. To display server settings in the IAS console, right-click the server and click **Properties**.

The Netsh command script contains the local server settings as a large data block in a compressed text format. This large data block is used within the script by the **netsh exec** command to import the configuration of a saved data block into an existing IAS server on the same or another computer. To save the Netsh command script to a file, type **netsh aaaa show server_settings >** *Path\File.txt* at a command prompt.

This command is not available on computers running Windows 2000 Server.

The output of this command is an encoded data block that is either displayed in the command prompt window or saved to a text file at the folder location you specify. The following registry keys and values are contained in the data block:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IAS\Parameters\Allow SNMP Set\REG_DWORD**

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout\MaxDenials\REG_DWORD**

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout\ResetTime\REG_DWORD**

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IAS\Parameters\Ping User-Name\REG_SZ**

**show version**
Displays the version of the IAS database.

**Syntax**

```
show version
```

**Importing or Exporting Configurations**
You can export IAS server configuration in part or total by using the Netsh commands for AAAA. For example, you can export the full configuration of an IAS server using the **netsh aaaa show config** command at the command prompt, or you can export only the remote access policies of the server using the **netsh aaaa show remote_acces_policies** command. IAS configuration can then be imported by using the **netsh exec** command.

When importing or exporting configurations, keep the following in mind:

- You do not need to stop IAS on the destination computer to run the **netsh exec** command. When the command is run, IAS is automatically refreshed with the updated configuration settings.

- Do not use this procedure if the source IAS database is a higher version number than the version number of the destination IAS database. You can view the version number of the IAS database from the display of the **netsh aaaa show config** command.

### Note

- The configuration of IAS servers running on products in the Windows 2000 Server family can be imported into products in the Windows Server 2003 family with netsh AAAA commands. The reverse, however, is not possible. To centrally manage IAS on both Windows 2000 and Windows Server 2003, you can first make configuration changes on IAS servers running Windows 2000, and then import the new server configuration into IAS servers running Windows Server 2003.

- IAS configurations are encoded and compressed; however, they are not encrypted in the text file. Because IAS configurations are not encrypted in the text file, sending the text file over a network poses a security risk and is not recommended.

- You can copy a remote access configuration from a server running Windows 2000 or Windows Server 2003 and the Routing and Remote Access service to an IAS server. Use the **netsh aaaa dump** command at the remote access server to create a file that contains the configuration, copy the file to the IAS server, and then use the **netsh exec** command to import the configuration to the IAS server. Next, configure RADIUS clients on the IAS server.

**Software Development Kits**
The Microsoft Platform Software Development Kit (SDK) contains the IAS and EAP SDKs, which you can use to extend the functionality of IAS.

**IAS Software Development Kit**
The IAS SDK can be used to:

- Return custom attributes to the access server in addition to those returned by IAS. For example, you can create a customized module to assign IP addresses.
- Control the number of user network sessions.
- Import usage and audit data directly into an Open Database Connectivity (ODBC)-compliant database.
- Create customized authorization modules.
- Create customized authentication modules (non-EAP).

**EAP Software Development Kit**
The Extensible Authentication Protocol (EAP) SDK allows you to create EAP types and implement arbitrary authentication methods using EAP.

**Network Monitor**
The optional management and monitoring component Network Monitor tool can be used to capture the RADIUS messages being sent and received for detailed analysis.

For more information about setting up and using Network Monitor, see “Monitoring Network Performance” in the *Microsoft Windows Server 2003 Resource Kit Server Operations Guide*.

To view RADIUS messages, configure Network Monitor with a display filter to display only RADIUS messages by disabling all protocols except the RADIUS protocol.

**IAS Settings**
The following section provides step-by-step information about how to configure the domain name and certificate templates, as well as how to use shared secrets and other IAS features and settings.

**Configuring the Default Domain Name**
While processing connection requests, IAS examines the user name portion of the Access-Request message to determine whether a domain name has been specified. If a domain name is specified and IAS is configured to access the user accounts database in the designated domain, IAS proceeds with processing the connection request.

### Note

- Some network access servers delete or modify the domain name as specified by the user. As a result, the network access request is authenticated against the default domain, which might not be the domain for the user’s account. To resolve this problem, configure your RADIUS servers to change the user name into the correct format with the accurate domain name.

When the user name does not contain a domain name, IAS supplies one. By default, the IAS-supplied domain name is the domain of which the IAS server is a member.

IAS resolves a user name without a specified domain name in the following sequence:

1. IAS checks the default domain registry key. If one is specified, IAS authenticates the user against the domain specified in the registry key.
2. If the IAS server is a member of a domain, IAS authenticates the user against the domain to which it is joined.
3. If the IAS server is not a member of a domain, IAS authenticates the user against the local Security Accounts Manager (SAM) database.

For information about how to specify the IAS-supplied domain name, see "IAS Registry Entries."

**Configuring Certificate Templates**
The configuration of certificate templates is not a function performed within the IAS console; however, the correct configuration of your certificate templates is critical to the success of your IAS deployment. The following two sections provide server certificate requirements and client certificate requirements, as well as how to configure your certificate template in the Certificate Templates snap-in.

**Note**

- You can view the IAS console and the Certificate Templates snap-in within the same Microsoft Management Console (MMC) window. Just create a custom snap-in configuration by adding the snap-ins you want to view to the console window. To reuse your custom view, save the console configuration.

These instructions are based on the following assumptions:

- You have designed a public key infrastructure (PKI).
- Based on both Certificate Services best practices and your own network and PKI requirements, you have installed Certificate Services in Windows Server 2003.

**Note**

- If you are running an enterprise certification authority (CA) on a computer running Windows Server 2003, Enterprise Edition, Windows Server 2003, Datacenter Edition, the 64-bit version of Windows Server 2003, Enterprise Edition, or the 64-bit version of Windows Server 2003, Datacenter Edition, it is recommended that you use the **RAS and IAS Server** certificate template for server certificates and the **Workstation Authentication** template for client certificates. If you are running an enterprise CA on a computer running Windows Server 2003, Standard Edition, you can use the **Computer** certificate template for both server and client certificates. However, you must configure the certificate purposes to match the certificate function.

**Server certificate requirements**

Client computers can be configured to validate server certificates by using the **Validate server certificate** option.

With Protected EAP Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP-MS-CHAP v2), Protected EAP Transport Level Security (PEAP-TLS), or EAP-TLS as the authentication method, the client accepts the server's authentication attempt when the certificate meets the following requirements:

- The Subject name contains a value. If you issue a certificate to your IAS server that has a blank Subject, the certificate is not available to authenticate your IAS server.

  **To configure the certificate template with a Subject name**

  1. Open Certificate Templates.
  2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
  3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.
  4. In **Subject name format**, select a value other than **None**.

- The computer certificate on the server chains to a trusted root CA and does not fail any of the checks that are performed by CryptoAPI and specified in the remote access policy. When you use the **RAS and IAS Servers** certificate template, the certificate based on the template passes CryptoAPI checks.
- The IAS or VPN server computer certificate is configured with the Server Authentication purpose in Enhanced Key Usage (EKU) extensions. The object identifier for Server Authentication is 1.3.6.1.5.5.7.3.1.
- For user certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate contains the user principal name (UPN).

  **To configure the UPN name in a certificate template**

  1. Open Certificate Templates.
  2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
  3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.
  4. In **Include this information in alternate subject name**, select **User principal name (UPN)**.

- For computer certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate must contain the client's fully qualified domain name (FQDN), which is also called the DNS name.

**To configure this name in the certificate template**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.
4. In **Include this information in alternate subject name**, select **DNS name**.

With PEAP-TLS and EAP-TLS, clients display a list of all installed certificates in the Certificates snap-in, with the following exceptions:

- Wireless clients do not display registry-based and smart card-logon certificates.
- Wireless clients and virtual private network (VPN) clients do not display password-protected certificates.
- Certificates that do not contain the Client Authentication purpose in EKU extensions are not displayed.

**Note**

- You can designate which certificate is used by the IAS or VPN server in the remote access policy profile. When you configure EAP and PEAP authentication methods that require a certificate for server authentication and you do not select a specific certificate in the **Smart Card or other Certificate Properties** dialog box, the IAS or VPN server automatically selects a certificate from the computer certificate store. If the server obtains a newer certificate, it uses the new certificate. This might cause the IAS or VPN server to use a certificate that is not correctly configured for authentication, causing authentication to fail as a result. To prevent this, always manually select a server certificate when configuring PEAP and EAP authentication methods that require one.

**Client certificate requirements**

With EAP-TLS or PEAP-TLS, the server accepts the client authentication attempt when the certificate meets the following requirements:

- The client certificate is issued by an enterprise CA and maps to a user or computer account in Active Directory.
- The user or computer certificate on the client chains to a trusted root CA, includes the Client Authentication purpose in EKU extensions (the object identifier for Client Authentication is 1.3.6.1.5.5.7.3.2), and fails neither the checks that are performed by CryptoAPI and specified in the remote access policy nor the Certificate object identifier checks that are specified in IAS remote access policy. When you use the **Workstation Authentication** certificate template, the certificate based on the template passes CryptoAPI checks by default.
- The 802.1X client does not use registry-based certificates that are either smart card-logon or password-protected certificates.
- For user certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate contains the user principal name (UPN).

  **To configure the UPN in a certificate template**

  1. Open Certificate Templates.
  2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
  3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.
  4. In **Include this information in alternate subject name**, select **User principal name (UPN)**.

- For computer certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate must contain the client's fully qualified domain name (FQDN), which is also called the DNS name.

  **To configure the DNS name in the certificate template**

  1. Open Certificate Templates.
  2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.

3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.

4. In **Include this information in alternate subject name**, select **DNS name**.

With PEAP-TLS and EAP-TLS, clients display a list of all installed certificates in the Certificates snap-in, with the following exceptions:

- Wireless clients do not display registry-based and smart card-logon certificates.
- Wireless clients and VPN clients do not display password-protected certificates.
- Certificates that do not contain the Client Authentication purpose in EKU extensions are not displayed.

**Using the Message Authenticator Attribute**

When you add one or more RADIUS clients in the IAS snap-in, you configure the IP address of one RADIUS client or you configure multiple RADIUS clients using an IP address range. If an incoming RADIUS Access-Request message does not originate from an IP address of a configured RADIUS client, IAS automatically discards the message, providing protection for an IAS server. However, source IP addresses can be spoofed (substituted with other IP addresses) by malicious users.

To provide protection from spoofed Access-Request messages and RADIUS message tampering, each RADIUS message can be additionally protected with the RADIUS Message Authenticator attribute, which is described in RFC 2869, "RADIUS Extensions."

The RADIUS Message Authenticator attribute is a Message Digest 5 (MD5) hash of the entire RADIUS message. The shared secret configured on the IAS server and the RADIUS client is used as the key. If the RADIUS Message Authenticator attribute is present, it is verified. If the Access-Request message fails verification, the message is discarded by the IAS server. If the RADIUS client settings require the Message Authenticator attribute and it is not present, the RADIUS message is discarded.

With IAS in Windows Server 2003, all EAP and PEAP authentication methods use the Message Authenticator attribute by default.

**Enabling Remote Access Account Lockout**

You can use remote access account lockout to specify how many times a remote access authentication fails against a valid user account before the user is denied access. Remote access account lockout is especially important for remote access VPN connections over the Internet. An attacker on the Internet can attempt to access an organization intranet by sending credentials (valid user name, guessed password) during the VPN connection authentication process. During a dictionary attack, the attacker sends hundreds or thousands of credentials by using a list of passwords based on common words or phrases.

When remote access account lockout is enabled, a dictionary attack is thwarted after a specified number of failed attempts. As the network administrator, you must decide on two remote access account lockout variables:

- The number of failed attempts before future attempts are denied.

  After each failed attempt, a failed attempts counter for the user account is incremented. If the user account's failed attempts counter reaches the configured maximum, future attempts to connect are denied.

  A successful authentication resets the failed attempts counter when its value is less than the configured maximum. In other words, the failed attempts counter does not accumulate beyond a successful authentication.

- The frequency with which the failed attempts counter is reset.

  The failed attempts counter is periodically reset to 0. If an account is locked out after the maximum number of failed attempts, the failed attempts counter is automatically reset to 0 after the reset time.

You enable the remote access account lockout feature by changing settings in the registry on the computer that provides the authentication. If the remote access server is configured for Windows Authentication, modify the registry on the remote access server computer. If the remote access server is configured for RADIUS authentication and IAS is being used, modify the registry on the IAS server.

For more information about how to enable remote access account lockout, see "IAS Registry Entries."

**Using Shared Secrets**

A shared secret is a text string that serves as a password between:

- A RADIUS client and RADIUS server.
- A RADIUS client and a RADIUS proxy.
- A RADIUS proxy and a RADIUS server.

For a configuration that uses a RADIUS client, a RADIUS proxy, and a RADIUS server, the shared secret that is used between the RADIUS client and the RADIUS proxy can be different from the shared secret used between the RADIUS proxy and the RADIUS server.

Shared secrets are used to verify that RADIUS messages, with the exception of the Access-Request message, are sent by a RADIUS-enabled device that is configured with the same shared secret. Shared secrets also verify that the RADIUS message has not been modified in transit. The shared secret is also used to encrypt some RADIUS attributes, such as User-Password and Tunnel-Password. To provide verification for Access-Request messages, you can enable use of the RADIUS Message Authenticator attribute for both the RADIUS client configured on the IAS server and the access server.

If you specify RADIUS clients by using an IP address range, all RADIUS clients within the address range must use the same shared secret.

When creating and using a shared secret:

- Use the same case-sensitive shared secret on both RADIUS devices.
- Use a different shared secret for each RADIUS server-RADIUS client pair.
- Generate a random sequence at least 22 characters long.
- Use any standard alphanumeric and special characters.
- Use a shared secret of up to 128 characters in length. To protect your IAS server and your RADIUS clients from brute force attacks, use long shared secrets (more than 22 characters).
- Make the shared secret a random sequence of letters, numbers, and punctuation, as shown in the following table.

  **Alphanumeric and special characters**

| Group | Examples |
|---|---|
| Letters (uppercase and lowercase) | A, B, C and a, b, c |
| Numerals | 0, 1, 2, 3 |
| Symbols (all characters not defined as letters or numerals) | Exclamation point (!), asterisk (*), colon (:) |

- Change your shared secret often to protect your IAS server and your RADIUS clients from dictionary attacks.
- Be certain that the RADIUS client's shared secret and the shared secret you type when running the New RADIUS Client Wizard are identical.

The stronger your shared secret, the more secure are the attributes (for example, those used for passwords and encryption keys) that are encrypted with the shared secret. An example of a strong shared secret is 5d#>3fq4bV)J7%a3-zR13sM. Before creating your shared secret, verify that your network access servers support the use of the letters, numerals, and symbols that you have chosen.

**Note**

- When the Password Authentication Protocol (PAP) is used between an access client and an access server (a RADIUS client), the access server encrypts the PAP password by using the shared secret and sends it in an Access-Request packet. If the access server sends the Access-Request message to a RADIUS proxy, the RADIUS proxy must first decrypt the PAP password with the shared secret that was used between the RADIUS proxy and the access server. Next, it encrypts the PAP password with the shared secret that was used between the RADIUS proxy and the RADIUS server before forwarding the Access-Request message. Because a malicious user or process at a RADIUS proxy can record user names and passwords for PAP connections after they are decrypted but before they are encrypted, the use of PAP is strongly discouraged. In addition, if you are using a password-based authentication method, it is strongly recommended that you use MS-CHAP v2, MS-CHAP, or CHAP with strong passwords to provide password protection from dictionary attacks.

**IAS Registry Entries**

In Windows Server 2003, IAS supports registry values that provide the following functionality:

**Note**

- Incorrectly editing the registry can severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**HKEY_LOCAL_MACHINE\SYSTEM**
**Eventlogging**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL

Secure channel (Schannel) is a security support provider (SSP) that supports a set of Internet security protocols, such as Secure Sockets Layer (SSL) and TLS. These protocols provide identity authentication and secure, private communication through encryption. Logging of client

certificate validation failures is a secure channel event, and is not enabled on the IAS server by default.

Logging of client certificate validation failures is a secure channel event, and is not enabled on the IAS server by default.

**To enable secure channel events**
You can enable additional secure channel events by changing the registry key value from **1** (**REG_DWORD** type, data **0x00000001**) to **3** (**REG_DWORD** type, data **0x00000003**).

**Note**

- The logging of rejected or discarded authentication events is enabled by default.

**Ping User-Name**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IAS\Parameters

**Ping User-Name** can be added to the registry on your IAS servers by a member of the local Administrators group. When you add **Ping User-Name** to the registry, you must supply values for **Name**, **Type**, and **Data**.

**Note**

- **Ping User-Name** is not installed by default. You must add **Ping User-Name** to the registry. You can add an entry to the registry using Registry Editor.

**Values for Ping User-Name**

| Name | Type | Data |
|---|---|---|
| ping user-name | REG_SZ | *User name* |

To indicate more than one user name for a **Ping User-Name** value, enter a name pattern, such as a DNS name including wildcard characters, in **Data**.

**Ping User-Name** can be added to the registry key as a string value:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IAS\Parameters**

**Note**

- When you configure **Ping User-Name** in the IAS server registry, you must configure NASs and RADIUS proxies to use the same User name, or a User name that matches the name and wildcard characters, that you configured in **Data**. If the User name value configured for the NASs and RADIUS proxies do not match the value in **Data**, and IAS is configured for logging, IAS will log the RADIUS traffic, and log files will become cluttered with accounting data from ping requests.

**MaxConcurrentApi**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Netlogon\Parameters

If the IAS server is on a computer other than a domain controller and it is receiving a very large number of authentication requests per second, you can improve performance by increasing the number of concurrent authentications between the IAS server and the domain controller.

**To increase concurrent authentication**
Add a new value named **MaxConcurrentApi** to this registry key and assign to it a value from **2** through **5**.

**DefaultDomain**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn

**To specify the IAS-supplied domain**
By default, the IAS-supplied domain name is the domain for which the IAS server is a member. You can specify the IAS-supplied domain through the **DefaultDomain** registry setting.

**AccountLockout**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Parameters

**To enable remote access account lockout**
Set the **MaxDenials** entry in the registry to **1** or greater. **MaxDenials** is the maximum number of failed attempts before the account is locked out.

By default, **MaxDenials** is set to **0**, which means that remote access account lockout is disabled.

**To modify the amount of time before the failed attempts counter is reset**
Set the **ResetTime (mins)** entry in the registry to the required number of minutes.

By default, **ResetTime (mins)** is set to **0xb40**, or 2,880 minutes (48 hours).

**To manually reset a user account that has been locked out before the failed attempts counter is automatically reset**
Delete the following registry subkey that corresponds to the user's account name:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout\domain name:user name**

When the lockout count for a user account is reset to **0** due to either a successful authentication or an automatic reset, the registry subkey for the user account is deleted.

**Default User Identity**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy

User identity is the means by which IAS identifies the user for the purposes of authentication and authorization. Normally, the user identity is the string value of the User-Name RADIUS attribute. If the User-Name attribute is not present, or is present but equals null, the user identity is set to the Guest account or the account specified by the **Default User Identity** registry value:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy**

**Enable or Disable LAN Manager Authentication**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy

By default, MS-CHAP for Windows Server 2003 does not support LAN Manager authentication. This is a change from Windows 2000.

Although the use of MS-CHAP or LAN Manager authentication is not recommended for security reasons, you might need to deploy one or both of these authentication methods to support legacy clients. If you deploy MS-CHAP with change password capability enabled in IAS, you must also deploy LAN Manager authentication.

**To enable LAN Manager authentication**
If you want to enable the use of LAN Manager authentication with MS-CHAP for older Microsoft operating systems such as Windows NT 3.5 and Windows 95, you must set **Allow LM Authentication** to **1** on the authenticating server.

**To disable LAN Manager authentication**
LAN Manager authentication is disabled by default. However, if you have previously enabled it and want to disable it again, set **Allow LM Authentication** to **0**.

**Override User-Name**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy

To always use the calling number as the user identity, set the **Override User-Name** registry value to **1** on the authenticating server.

**Note**

- If you set Override User-Name to **1** and the User Identity Attribute to **31**, the authenticating server can perform only Automatic Number Identification/Calling Line Identification (ANI/CLI)-based authentication. Normal authentication by using authentication protocols (such as MS CHAP and EAP) is disabled.

**User Identity Attribute**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy

The RADIUS attribute that IAS uses to identify the user is configurable by setting the **User Identity Attribute** registry setting.

You can change this key's value to the number of the RADIUS attribute that is used for the user identity. By default, **User Identity Attribute** is set to **1**, the RADIUS type value for the User-Name RADIUS attribute.

This registry setting tells the authenticating server to use the calling number (RADIUS attribute 31, Calling-Station-ID) as the identity of the calling user. The user identity is set to the calling number only when there is no user name being supplied in the connection attempt.

**Related Information**
The following resources contain additional information that is relevant to this section.

- "Windows Server 2003 Resource Kit Tools" in the Tools and Settings Collection.

- RFC 2869 in the IETF RFC Database.

# VPN Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN Technical Reference**
Virtual private network (VPN) is a network access technology that is available as part of the Routing and Remote Access service that is included in Microsoft Windows Server 2003. VPN enables administrators to provide access to private networks at a lower cost than using leased-line connections.

VPN technology is the extension of a private network that encompasses point-to-point links across shared or public networks, such as the Internet. VPN allows remote clients to connect to an organization's network resources. VPN also allows remote offices to connect to each other to share resources and information. A VPN connection uses authentication and data encryption to help preserve the privacy and integrity of data as it traverses a public network.

**In this subject**


- What Is VPN?
- How VPN Works
- VPN Tools and Settings

# What Is VPN?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is VPN?**

### In this section


- VPN Scenarios
- VPN Connection Properties
- Tunneling Overview
- Routing for VPN
- VPN and Firewalls Overview
- Technologies Related to VPN
- Related Information


The virtual private network (VPN) technology included in Windows Server 2003 helps enable cost-effective, secure remote access to private networks. VPN allows administrators to take advantage of the Internet to help provide the functionality and security of private WAN connections at a lower cost. In Windows Server 2003, VPN is enabled using the Routing and Remote Access service. VPN is part of a comprehensive network access solution that includes support for authentication and authorization services, and advanced network security technologies.

There are two main strategies that help provide secure connectivity between private networks and enabling network access for remote users.

**Dial-up or leased line connections**
A dial-up or leased line connection creates a physical connection to a port on a remote access server on a private network. However, using dial-up or leased lines to provide network access is expensive when compared to the cost of providing network access using a VPN connection.

**VPN connections**
VPN connections use either Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol/Internet Protocol security (L2TP/IPSec) over an intermediate network, such as the Internet. By using the Internet as a connection medium, VPN saves the cost of long-distance phone service and hardware costs associated with using dial-up or leased line connections. A VPN solution includes advanced security technologies such as data encryption, authentication, authorization, and Network Access Quarantine Control.

### Note

- Network Access Quarantine Control is used to delay remote access to a private network until the configuration of the remote access computer has been examined and validated.

Using VPN, administrators can connect remote or mobile workers (VPN clients) to private networks. Remote users can work as if their computers are physically connected to the network. To accomplish this, VPN clients can use a Connection Manager profile to initiate a connection to a VPN server. The VPN server can communicate with an Internet Authentication Service (IAS) server to authenticate and authorize a user session and maintain the connection until it is terminated by the VPN client or by the VPN server. All services typically available to a LAN-connected client (including file and print sharing, Web server access, and messaging) are enabled by VPN.

VPN clients can use standard tools to access resources. For example, clients can use Windows Explorer to make drive connections and to connect to printers. Connections are persistent: Users do not need to reconnect to network resources during their VPN sessions. Because drive letters and universal naming convention (UNC) names are fully supported by VPN, most commercial and custom applications work without modification.

**VPN Scenarios**
Virtual private networks are point-to-point connections across a private or public network such as the Internet. A VPN client uses special TCP/IP-based protocols, called tunneling protocols, to make a virtual call to a virtual port on a VPN server. In a typical VPN deployment, a client initiates a virtual point-to-point connection to a remote access server over the Internet. The remote access server answers the call, authenticates the caller, and transfers data between the VPN client and the organization's private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header. The header provides routing information that enables the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data being sent is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The link in which the private data is encapsulated and encrypted is known as a VPN connection.

### A VPN Connection




There are two types of VPN connections:

- Remote access VPN
- Site-to-site VPN

**Remote Access VPN**
Remote access VPN connections enable users working at home or on the road to access a server on a private network using the infrastructure provided by a public network, such as the Internet. From the user's perspective, the VPN is a point-to-point connection between the computer (the VPN client) and an organization's server. The exact infrastructure of the shared or public network is irrelevant because it appears logically as if the data is sent over a dedicated private link.

**Site-to-Site VPN**
Site-to-site VPN connections (also known as router-to-router VPN connections) enable organizations to have routed connections between separate offices or with other organizations over a public network while helping to maintain secure communications. A routed VPN connection across the Internet logically operates as a dedicated WAN link. When networks are connected over the Internet, as shown in the following figure, a router forwards packets to another router across a VPN connection. To the routers, the VPN connection operates as a data-link layer link.

A site-to-site VPN connection connects two portions of a private network. The VPN server provides a routed connection to the network to which the VPN server is attached. The calling router (the VPN client) authenticates itself to the answering router (the VPN server), and, for mutual authentication, the answering router authenticates itself to the calling router. In a site-to site VPN connection, the packets sent from either router across the VPN connection typically do not originate at the routers.

### VPN Connecting Two Remote Sites Across the Internet




**VPN Connection Properties**
PPTP-based VPN and L2TP/IPSec-based VPN connection properties are described in the following sections.

**Encapsulation**
VPN technology provides a way of encapsulating private data with a header that allows the data to traverse the network.

**Authentication**
There are three types of authentication for VPN connections:

**User authentication**
For the VPN connection to be established, the VPN server authenticates the VPN client attempting the connection and verifies that the VPN client has the appropriate permissions. If mutual authentication is being used, the VPN client also authenticates the VPN server, providing protection against masquerading VPN servers.

The user attempting the PPTP or L2TP/IPSec connection is authenticated using Point-to-Point (PPP)-based user authentication protocols such as Extensible Authentication Protocol-Transport Layer Security (EAP-TLS), Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge-Handshake Authentication Protocol version 2 (MS-CHAP v2), Shiva Password Authentication Protocol (SPAP), and Password Authentication Protocol (PAP). For PPTP connections, you must use EAP-TLS, MS-CHAP, or MS-CHAP v2. EAP-TLS using smart cards or MS-CHAP v2 is highly recommended, as they provide mutual authentication and are the most secure methods of exchanging credentials.

**Computer authentication with L2TP/IPSec**
By performing computer-level authentication with IPSec, L2TP/IPSec connections also verify that the remote access client computer is trusted.

**Data authentication and integrity**
To verify that the data being sent on an L2TP/IPSec VPN connection originated at the other end of the connection and was not modified in transit, L2TP/IPSec packets include a cryptographic checksum based on an encryption key known only to the sender and the receiver.

**Data Encryption**
Data can be encrypted for protection between the endpoints of the VPN connection. Data encryption should always be used for VPN connections where private data is sent across a public network such as the Internet. Data that is not encrypted is vulnerable to unauthorized interception. For VPN connections, Routing and Remote Access uses Microsoft Point-to-Point Encryption (MPPE) with PPTP and IPSec encryption with L2TP.

**Address and Name Server Allocation**
When a VPN server is configured, it creates a virtual interface that represents the interface on which all VPN connections are made. When a VPN client establishes a VPN connection, a virtual interface is created on the VPN client that represents the interface connected to the VPN server. The virtual interface on the VPN client is connected to the virtual interface on the VPN server, creating the point-to-point VPN connection.

The virtual interfaces of the VPN client and the VPN server must be assigned IP addresses. The assignment of these addresses is done by the VPN server. By default, the VPN server obtains IP addresses for itself and VPN clients using the Dynamic Host Configuration Protocol (DHCP). Otherwise, a static pool of IP addresses can be configured to define one or more address ranges, with each range defined by an IP network ID and a subnet mask or start and end IP addresses.

Name server assignment, the assignment of Domain Name System (DNS) and Windows Internet Name Service (WINS) servers to the VPN connection, also occurs during the process of establishing the VPN connection.

**Tunneling Overview**
Tunneling is a method of using a network infrastructure to transfer data for one network over another network. The data (or payload) to be transferred can be the frames (or packets) of another protocol. Instead of sending a frame as it is produced by the originating node, the tunneling protocol encapsulates the frame in an additional header. The additional header provides routing information so that the encapsulated payload can traverse the intermediate network.

The encapsulated packets are then routed between tunnel endpoints over the network. The logical path through which the

encapsulated packets travel through the network is called a tunnel. After the encapsulated frames reach their destination on the network, the frame is de-encapsulated (the header is removed) and the payload is forwarded to its final destination. Tunneling includes this entire process (encapsulation, transmission, and de-encapsulation of packets).

**Tunneling**




**Tunneling Protocols**

Tunneling enables the encapsulation of a packet from one type of protocol within the datagram of a different protocol. For example, VPN uses PPTP to encapsulate IP packets over a public network such as the Internet. A VPN solution based on either PPTP or L2TP can be configured.

PPTP and L2TP depend heavily on the features originally specified for PPP. PPP was designed to send data across dial-up or dedicated point-to-point connections. For IP, PPP encapsulates IP packets within PPP frames and then transmits the encapsulated PPP-packets across a point-to-point link. PPP was originally defined as the protocol to use between a dial-up client and a network access server (NAS).

**PPTP**

PPTP allows multiprotocol traffic to be encrypted and then encapsulated in an IP header to be sent across an organization's IP network or a public IP network such as the Internet. PPTP encapsulates Point-to-Point Protocol (PPP) frames in IP datagrams for transmission over the network. PPTP can be used for remote access and site-to-site VPN connections. PPTP is documented in RFC 2637 in the IETF RFC Database.

PPTP uses a TCP connection for tunnel management and a modified version of Generic Routing Encapsulation (GRE) to encapsulate PPP frames for tunneled data. The payloads of the encapsulated PPP frames can be encrypted, compressed, or both. The following figure shows the structure of a PPTP packet containing an IP datagram.

**Structure of a PPTP Packet Containing an IP Datagram**




When using the Internet as the public network for VPN, the PPTP server is a PPTP-enabled VPN server with one interface on the Internet and a second interface on the intranet.

**L2TP**

L2TP allows multiprotocol traffic to be encrypted and then sent over any medium that supports point-to-point datagram delivery, such as IP, X.25, frame relay, or asynchronous transfer mode (ATM). L2TP is a combination of PPTP and Layer 2 Forwarding (L2F), a technology developed by Cisco Systems, Inc. L2TP represents the best features of PPTP and L2F. L2TP encapsulates PPP frames to be sent over IP, X.25, frame relay, or ATM networks. When configured to use IP as its datagram transport, L2TP can be used as a tunneling protocol over the Internet. L2TP is documented in RFC 2661 in the IETF RFC Database.

L2TP over IP networks uses User Datagram Protocol (UDP) and a series of L2TP messages for tunnel management. L2TP also uses UDP to send L2TP-encapsulated PPP frames as tunneled data. The payloads of encapsulated PPP frames can be encrypted, compressed, or both, although the Microsoft implementation of L2TP does not use MPPE to encrypt the PPP payload. The following figure shows the structure of an L2TP packet containing an IP datagram.

**Structure of an L2TP Packet Containing an IP Datagram**




**L2TP with IPSec (L2TP/IPSec)**
In the Microsoft implementation of L2TP, IPSec Encapsulating Security Payload (ESP) in transport mode is used to encrypt L2TP traffic. The combination of L2TP (the tunneling protocol) and IPSec (the method of encryption) is known as L2TP/IPSec. L2TP/IPSec is described in RFC 3193 in the IETF RFC Database.

The result after applying ESP to an IP packet containing an L2TP message is shown in the following figure.

**Encryption of L2TP Traffic with IPSec ESP**




**Routing for VPN**
Routing for remote access and site-to-site VPN connections is described in the following sections.

**Routing for Remote Access VPN Connections**
Conventional routing occurs between routers over either LAN-based shared access technologies, such as Ethernet or Token Ring, or WAN-based point-to-point technologies, such as T1 or frame relay.

**Default Routing**
The preferred method for directing packets to a remote network is to create a default route on the remote access client that directs packets to the remote network (the default configuration for VPN remote access clients). Any packet that is not intended for the neighboring LAN segment is sent to the remote network. When a connection is made, the remote access client, by default, adds a default route to its routing table and increases the metric of the existing default route to ensure that the newest default route is used. The newest default route points to the new connection, which ensures that any packets that are not addressed to the local LAN segment are sent to the remote network.

Under this configuration, when a VPN client connects and creates a new default route, Internet sites that have been accessible are no longer accessible (unless Internet access is available through the organization's intranet). This poses no problem for remote VPN clients that require access only to the organization's network. However, it is not acceptable for remote clients that need access to the Internet while they are connected to the organization's network.

**Split Tunneling**
Split tunneling enables remote access VPN clients to route corporate-based traffic over the VPN connection while sending Internet-based traffic using the user's local Internet connection. This prevents the use of corporate bandwidth for access to Internet sites.

However, a split tunneling implementation can introduce a security issue. If a remote access client has reachability to both the Internet and a private organization network simultaneously, the possibility exists that the Internet connection could be exploited to gain access to the private organization network through the remote access client. Security-sensitive companies can choose to use the default routing model to help ensure that all VPN client communications are protected by the corporate firewall.

**Routing for Site-to-Site VPN Connections**
With conventional WAN technologies, IP packets are forwarded between two routers over a physical or logical point-to-point connection. This connection is dedicated to the customer across a private data network that is provided by the WAN service provider.

With the advent of the Internet, packets can now be routed between routers that are connected to the Internet across a virtual connection that emulates the properties of a dedicated, private, point-to-point connection. This type of connection is known as a site-to-site VPN connection. Site-to-site VPN connections can be used to replace expensive long-haul WAN links with short-haul WAN links to a local Internet service provider (ISP).

A site-to-site VPN connection connects two portions of a private network. The VPN server provides a routed connection to the network to which the VPN server is attached. On a site-to-site VPN connection, the packets sent from either router across the VPN connection typically do not originate at the routers.

To facilitate routing between the sites, each VPN server and the routing infrastructure of its connected site must have a set of routes that represent the address space of the other site. These routes can be added manually, or routing protocols can be used to automatically add and maintain a set of routes.

**Site-to-Site Routing Protocols**
There are two routing protocols that can be used in a site-to-site VPN deployment:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)

**RIP**
RIP is designed for exchanging routing information within a small to medium-size network. RIP routers dynamically exchange routing table entries.

The Windows Server 2003 implementation of RIP has the following features:

- The ability to select which RIP version to run on each interface for incoming and outgoing packets.
- Split-horizon, poison-reverse, and triggered-update algorithms that are used to avoid routing loops and speed recovery of the network when topology changes occur.
- Route filters for choosing which networks to announce or accept.
- Peer filters for choosing which router's announcements are accepted.
- Configurable announcement and route-aging timers.
- Simple password authentication support.
- The ability to disable subnet summarization.

**OSPF**
OSPF is designed for exchanging routing information within a large or very large network. Instead of exchanging routing table entries like RIP routers, OSPF routers maintain a map of the network that is updated after any change to the network topology. This map, called the link state database, is synchronized between all the OSPF routers and is used to compute the routes in the routing table. Neighboring OSPF routers form an adjacency, which is a logical relationship between routers to synchronize the link state database.

**VPN and Firewalls Overview**
The routing service supports a variety of inbound and outbound packet-filtering features that block certain types of traffic. The filtering options include the following: TCP port, UDP port, IP protocol ID, Internet Control Message Protocol (ICMP) type, ICMP code, source address, and destination address. A VPN server can be placed behind a firewall or in front of a firewall. These two approaches are described in the following sections.

**VPN Server Behind a Firewall**
In the most common configuration, the firewall is connected to the Internet, and the VPN server is an intranet resource that is attached to the perimeter network. The VPN server has an interface on both the perimeter network and the intranet. In this scenario, the firewall must be configured with input and output filters on its Internet interface that allow tunnel maintenance traffic and tunneled data to pass to the VPN server. Additional filters can allow traffic to pass to Web, FTP, and other types of servers on the perimeter network. For an additional layer of security, the VPN server should also be configured with PPTP or L2TP/IPSec packet filters on its perimeter network interface.

**VPN Server in Front of a Firewall**
When the VPN server is in front of the firewall and connected to the Internet, packet filters must be added to the VPN server's Internet interface to allow only VPN traffic to and from the IP address of that interface.

For inbound traffic, when the tunneled data is decrypted by the VPN server, it is forwarded to the firewall. Through the use of its filters, the firewall allows the traffic to be forwarded to intranet resources. Because the only traffic that crosses the VPN server is generated by authenticated VPN clients, in this scenario, firewall filtering can be used to prevent VPN users from accessing specific intranet resources. Because Internet traffic allowed on the intranet must pass through the VPN server, this approach also prevents the sharing of FTP or Web intranet resources with non-VPN Internet users.

**Technologies Related to VPN**

Integrating VPN with the other network infrastructure components is an important part of VPN design and implementation. VPN has to be integrated with directory, authentication, and security services, as well as with IP address assignment and name server assignment services. Without proper design, VPN clients are unable to obtain proper IP addresses and resolve intranet names, and packets cannot be forwarded between VPN clients and intranet resources.

VPN-related technologies are described in the following sections:

- Connection Manager
- DHCP
- EAP-RADIUS
- IAS
- Name Server Assignment (DNS and WINS)
- NAT

**Connection Manager**

Connection Manager is a service profile that can be used to provide customized remote access to a network through a VPN connection. The advanced features of Connection Manager are a superset of basic dial-up networking. Connection Manager provides support for local and remote connections by using a network of points of presence (POPs), such as those available worldwide through ISPs. Windows Server 2003 includes a set of tools that enable a network manager to deliver pre-configured connections to network users. These tools are:

- The Connection Manager Administration Kit (CMAK)
- Connection Point Services (CPS)

**CMAK**

A network administrator can tailor the appearance and behavior of a connection made with Connection Manager by using CMAK. With CMAK, an administrator can develop client dialer and connection software that allows users to connect to the network by using only the connection features that the administrator defines for them. Connection Manager supports a variety of features that both simplify and enhance implementation of connection support, most of which can be incorporated using the Connection Manager Administration Kit Wizard.

CMAK enables administrators to build profiles that customize the Connection Manager installation package so that it reflects an organization's identity. CMAK allows administrators to determine which functions and features to include and how Connection Manager appears to end-users. Administrators can do this by using the CMAK wizard to build custom service profiles.

**CPS**

Connection Point Services (CPS) automatically distributes and updates custom phone books. These phone books contain one or more Point of Presence (POP) entries, with each POP supplying a telephone number that provides dial-up access to an Internet access point for VPN connections. The phone books give users complete POP information, so when they travel they can connect to different Internet POPs rather than being restricted to a single POP.

Without the ability to update phone books (a task CPS handles automatically), users would have to contact their organization's technical support staff to be informed of changes in POP information and to reconfigure their client-dialer software. CPS has two components:

- Phone Book Administrator

- Phone Book Service

**Phone Book Administrator**
Phone Book Administrator is a tool used to create and maintain the phone book database and to publish new phone book information to the Phone Book Service.

**Phone Book Service**
The Phone Book Service runs on an IIS server and responds to requests from Connection Manager clients to verify the current version of subscribers' or corporate employees' current phone books and, if necessary, downloads a phone book update to the Connection Manager client.

**DHCP**
For both PPTP and L2TP connections, the data being tunneled is a PPP frame. A PPP connection must be established before data can be sent. The VPN server must have IP addresses available in order to assign them to a VPN server's virtual interface and to VPN clients during the IP Control Protocol (IPCP) negotiation phase that is part of the process of establishing a PPP connection. The IP address assigned to a VPN client is also assigned to the virtual interface of that VPN client.

For Windows Server 2003-based VPN servers, the IP addresses assigned to VPN clients are obtained through DHCP by default. A static IP address pool can also be configured. DHCP is also used by remote access VPN clients to obtain additional configuration settings after the PPP connection is established.

**EAP-RADIUS**
EAP-RADIUS is the passing of EAP messages of any EAP type by an authenticator to a Remote Authentication Dial-In User Service (RADIUS) server for authentication. For example, for a remote access server that is configured for RADIUS authentication, the EAP messages sent between the remote access client and remote access server are encapsulated and formatted as RADIUS messages between the remote access server (the authenticator) and the RADIUS server (the authenticator).

EAP-RADIUS is used in environments where RADIUS is the authentication provider. An advantage of using EAP-RADIUS is that EAP types only need to be installed at the RADIUS server, not at each remote access server. In the case of an IAS server, only EAP types need to be installed.

In a typical use of EAP-RADIUS, a server running Routing and Remote Access is configured to use EAP and to use an IAS server for authentication. When a connection is made, the remote access client negotiates the use of EAP with the remote access server. When the client sends an EAP message to the remote access server, the remote access server encapsulates the EAP message as a RADIUS message and sends it to its configured IAS server. The IAS server processes the EAP message and sends a RADIUS-encapsulated EAP message back to the remote access server. The remote access server then forwards the EAP message to the remote access client. In this configuration, the remote access server is only a pass-through device. All processing of EAP messages occurs at the remote access client and the IAS server.

Routing and Remote Access can be configured to authenticate locally or to a RADIUS server. If Routing and Remote Access is configured to authenticate locally, all EAP methods will be authenticated locally. If Routing and Remote Access is configured to authenticate to a RADIUS server, then all EAP messages will be forwarded to the RADIUS server with EAP-RADIUS.

**IAS**
The VPN server can be configured to use either Windows or RADIUS as an authentication provider. If Windows is selected as the authentication provider, the user credentials sent by users attempting VPN connections are authenticated using typical Windows authentication mechanisms, and the connection attempt is authorized using local remote access policies.

If RADIUS is selected and configured as the authentication provider on the VPN server, user credentials and parameters of the connection request are sent as RADIUS request messages to a RADIUS server.

The RADIUS server receives a user-connection request from the VPN server and authenticates and authorizes the connection attempt. In addition to a yes or no response to an authentication request, RADIUS can inform the VPN server of other applicable connection parameters for this user such as maximum session time, static IP address assignment, and so on.

RADIUS can respond to authentication requests based on its own user account database, or it can be a front end to another database server, such as a Structured Query Language (SQL) server or a Windows domain controller (DC). The DC can be located on the same computer as the RADIUS server, or elsewhere. In addition, a RADIUS proxy can be used to forward requests to a remote RADIUS server.

IAS is the Windows implementation of a RADIUS server and proxy.

**Name Server Assignment (DNS and WINS)**
Name server assignment, the assignment of Domain Name System (DNS) and Windows Internet Name Service (WINS)

servers, occurs during the process of establishing a VPN connection. The VPN client obtains the IP addresses of the DNS and WINS servers from the VPN server for the intranet to which the VPN server is attached.

The VPN server must be configured with DNS and WINS server addresses to assign to the VPN client during IPCP negotiation. For NetBIOS name resolution, you do not have to use WINS and can enable the NetBIOS over TCP/IP (NetBT) proxy on the VPN server.

**NAT**

A network address translator (NAT) translates the IP addresses and Transmission Control Protocol/User Datagram Protocol (TCP/UDP) port numbers of packets that are forwarded between a private network and the Internet. The NAT on the private network can also provide IP address configuration information to the other computers on the private network.

PPTP-based VPN clients can be located behind a NAT if the NAT includes an editor that can translate PPTP packets. PPTP-based VPN servers can be located behind a NAT if the NAT is configured with static mappings for PPTP traffic. If the L2TP/IPSec-based VPN clients or servers are positioned behind a NAT, both client and server must support IPSec NAT traversal (NAT-T).

**Related Information**

The following resources contain additional information that is relevant to this section.

- Connection Manager Technical Reference
- DNS Technical Reference
- RFC 2637, 2661, and 3193 in the IETF RFC Database

# How VPN Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How VPN Works**
**In this section**


- VPN Architecture
- VPN Tunneling
- VPN Authentication
- VPN Encryption
- VPN Addressing and Routing
- VPN and Firewalls
- VPN and NAT
- Related Information


Microsoft Windows Server 2003 includes extensive support for virtual private network (VPN) technology, which leverages the IP connectivity of the Internet to connect remote clients and remote sites.

A VPN connection is the extension of a private network that includes links across shared or public networks, such as the Internet. VPN connections (VPNs) enable organizations to send data between two computers across the Internet in a manner that emulates the properties of a point-to-point private link.

**VPN Architecture**
Using VPNs, an organization can help secure private network traffic over an unsecured network, such as the Internet. VPN helps provide a secure mechanism for encrypting and encapsulating private network traffic and moving it through an intermediate network. Data is encrypted for confidentiality, and packets that might be intercepted on the shared or public network are indecipherable without the correct encryption keys. Data is also encapsulated, or wrapped, with an IP header containing routing information.

VPNs help enable users working at home, on the road, or at a branch office to connect in a secure fashion to a remote corporate server using the Internet. From the users perspective, the VPN is a point-to-point connection between the user's computer and a corporate server. The nature of the intermediate network, the Internet, is irrelevant to the user because it appears as if the data is being sent over a dedicated private link.

There are a number of ways to use VPN. The most common scenario is when a remote user accesses a private network across the Internet using a remote access VPN connection. In another scenario, a remote office connects to the corporate network using either a persistent or an on-demand site-to-site VPN connection (also known as a router-to-router VPN connection).

Each of these VPN scenarios can be deployed to provide connectivity over a public network, such as the Internet, or over a private intranet. VPN connections can also be deployed in an extranet scenario to communicate securely with business partners. An extranet functions as an intranet that can be securely shared with a designated business partner.

With both the remote access and site-to-site connections, VPNs enable an organization to replace long distance dial-up or leased lines with local dial-up or leased lines to an Internet service provider (ISP).

**Remote access VPN**
A remote access VPN connection is made by a remote access client. A remote access client is a single computer user who connects to a private network from a remote location. The VPN server provides access to the resources of the network to which the VPN server is connected. The packets sent across the VPN connection originate at the VPN client.

The VPN client authenticates itself to the VPN server and, for mutual authentication, the VPN server authenticates itself to the

VPN client.

**Site-to-site VPN**
A site-to-site VPN connection connects two portions of a private network or two private networks. For example, this allows an organization to have routed connections with separate offices, or with other organizations, over the Internet. A routed VPN connection across the Internet logically operates as a dedicated Wide Area Network (WAN) link.

The VPN server provides a routed connection to the network to which the VPN server is attached. On a site-to-site VPN connection, the packets sent from either router across the VPN connection typically do not originate at the routers. The calling router (the VPN client) authenticates itself to the answering router (the VPN server), and, for mutual authentication, the answering router authenticates itself to the calling router.

**Internet-based VPN Connections**
Using an Internet-based VPN connection, an organization can avoid long-distance charges while taking advantage of the global availability of the Internet.

**Remote Access VPN Connections over the Internet**
A remote access VPN connection over the Internet enables a remote access client to initiate a dial-up connection to a local ISP instead of connecting to a corporate or outsourced network access server (NAS). By using the established physical connection to the local ISP, the remote access client initiates a VPN connection across the Internet to the organization's VPN server. When the VPN connection is created, the remote access client can access the resources of the private intranet. The following figure shows remote access over the Internet.

### VPN Connecting a Remote Client to a Private Intranet




**Site-to-Site VPN Connections Over the Internet**
When networks are connected over the Internet, as shown in the following figure, a router forwards packets to another router across a VPN connection. To the routers, the VPN connection operates as a data-link layer link.

### VPN Connecting Two Remote Sites Across the Internet




**Intranet-based VPN Connections**
The intranet-based VPN connection takes advantage of IP connectivity in an organization's Local Area Network (LAN).

**Remote Access VPN Connections over an Intranet**
In some organization intranets, the data of a department, such as human resources, is so sensitive that the network segment of the department is physically disconnected from the rest of the intranet. While this protects the data of the human resources department, it creates information accessibility problems for authorized users not physically connected to the separate network segment.

VPN connections help provide the required security to enable the network segment of the human resources department to be physically connected to the intranet. In this configuration, a VPN server can be used to separate the network segments. The VPN server does not provide a direct routed connection between the corporate intranet and the separate network segment. Users on the corporate intranet with appropriate permissions can establish a remote access VPN connection with the VPN server and gain access to the protected resources. Additionally, all communication across the VPN connection is encrypted for data confidentiality. For those users who are not authorized to establish a VPN connection, the separate network segment is hidden from view.

The following figure shows remote access over an intranet.

**VPN Connection Allowing Remote Access to a Secured Network over an Intranet**




**Site-to-Site VPN Connections over an Intranet**

Two networks can be connected over an intranet using a site-to-site VPN connection. This type of VPN connection might be necessary, for example, for two departments in separate locations, whose data is highly sensitive, to communicate with each other. For instance, the finance department might need to communicate with the human resources department to exchange payroll information.

The finance department and the human resources department are connected to the common intranet with computers that can act as VPN clients or VPN servers. When the VPN connection is established, users on computers on either network can exchange sensitive data across the corporate intranet.

The following figure shows two networks connected over an intranet.

**VPN Connecting Two Networks over an Intranet**




**VPN Tunneling**

Tunneling is a network technology that enables the encapsulation of one type of protocol packet within the datagram of a different protocol. For example, Windows VPN connections can use Point-to-Point Tunneling Protocol (PPTP) packets to encapsulate and send private network traffic, such as TCP/IP traffic over a public network such as the Internet.

For PPTP and Layer Two Tunneling Protocol (L2TP), a tunnel is similar to a session. Both of the tunnel endpoints must agree to the tunnel and must negotiate configuration variables, such as address assignment, encryption, or compression parameters. In most cases, data transferred across the tunnel is sent using a datagram-based protocol. A tunnel management protocol is used as the mechanism to create, maintain, and terminate the tunnel.

After the tunnel is established, data can be sent. The tunnel client or server uses a tunnel data transfer protocol to prepare the data for transfer. For example, when the tunnel client sends a payload to the tunnel server, the tunnel client first appends a tunnel data transfer protocol header to the payload. The client then sends the resulting encapsulated payload across the network, which routes it to the tunnel server. The tunnel server accepts the packets, removes the tunnel data transfer protocol header, and forwards the payload to the target network. Information sent between the tunnel server and the tunnel client behaves similarly.

There are two types of tunneling:

- Voluntary tunneling
- Compulsory tunneling

**Voluntary Tunneling**

A user or client computer can issue a VPN request to configure and create a voluntary tunnel. In this case, the users computer is a tunnel endpoint and acts as the tunnel client.

Voluntary tunneling occurs when a client computer or routing server creates a virtual connection to the target tunnel server. To accomplish this, tunneling client software and the appropriate tunneling protocol must be installed on the client computer. For the protocols discussed in this technical reference, voluntary tunnels require an IP connection (either LAN or dial-up).

In a dial-up situation, the client must establish a dial-up connection to the network before the client can set up a tunnel. This is the most common case. The best example of this is the dial-up Internet user, who must dial an ISP and obtain an Internet

connection before a tunnel over the Internet can be created.

For a LAN-attached client computer, there is already a connection to the network that can provide routing of encapsulated payloads to the chosen LAN tunnel server. This would be the case for a client that is using an always-on broadband Internet connection.

It is a common misconception that VPN connections require a dial-up connection. They require only IP connectivity between the VPN client and VPN server. Some clients (such as home computers) use dial-up connections to the Internet to establish IP transport. This is a preliminary step in preparation for creating a tunnel and is not part of the tunnel protocol itself.

**Compulsory Tunneling**

In compulsory tunneling, a VPN-capable remote access server configures and creates a compulsory tunnel. With a compulsory tunnel, the user's computer is not a tunnel endpoint. Another device, the dial-up access server, between the user's computer and the tunnel server is the tunnel endpoint and acts as the tunnel client.

A number of vendors that sell dial-up access servers have implemented the ability to create a tunnel on behalf of a dial-up client. The computer or network device providing the tunnel for the client computer is variously known as a Front End Processor (FEP) for PPTP or an L2TP Access Concentrator (LAC) for L2TP. For the purposes of this reference, the term FEP is used to describe this functionality, regardless of the tunneling protocol. To carry out its function, the FEP must have the appropriate tunneling protocol installed and must be capable of establishing the tunnel when the client computer connects.

In compulsory tunneling, the client computer places a dial-up call to a tunneling-enabled NAS at the ISP. For example, a corporation might have contracted with an ISP to deploy a nationwide set of FEPs. These FEPs can establish tunnels across the Internet to a tunnel server connected to the organization’s private network, thus consolidating calls from geographically diverse locations into a single Internet connection at the organization network.

This configuration is known as compulsory tunneling because the client is compelled to use the tunnel created by the FEP. Once the initial connection is made, all network traffic to and from the client is automatically sent through the tunnel. With compulsory tunneling, the client computer makes a single PPP connection. When a client dials into the NAS, a tunnel is created and all traffic is automatically routed through the tunnel. An FEP can be configured to tunnel all dial-up clients to a specific tunnel server. The FEP could also tunnel individual clients, based on the user name or destination.

Unlike the separate tunnels created for each voluntary client, multiple dial-up clients can share a tunnel between the FEP and the tunnel server. When a second client dials into the access server (FEP) to reach a destination for which a tunnel already exists, there is no need to create a new instance of the tunnel between the FEP and tunnel server. Instead, the data traffic for the new client is carried over the existing tunnel. Since there can be multiple clients in a single tunnel, the tunnel is not terminated until the last user of the tunnel disconnects.

**PPTP**

Point-to-Point Tunneling Protocol (PPTP) encapsulates Point-to-Point Protocol (PPP) frames into IP datagrams for transmission over an IP-based network, such as the Internet or over a private intranet. PPTP is described in RFC 2637 in the IETF RFC Database.

PPTP uses a TCP connection, known as the PPTP control connection, to create, maintain, and terminate the tunnel. PPTP uses a modified version of Generic Routing Encapsulation (GRE) to encapsulate PPP frames as tunneled data. The payloads of the encapsulated PPP frames can be encrypted, compressed, or both.

PPTP assumes the availability of an IP network between a PPTP client (a VPN client using the PPTP tunneling protocol) and a PPTP server (a VPN server using the PPTP tunneling protocol). The PPTP client might already be attached to an IP network that can reach the PPTP server, or the PPTP client might have to use a dial-up connection to a NAS to establish IP connectivity as in the case of dial-up Internet users.

Authentication that occurs during the creation of a PPTP-based VPN connection uses the same authentication mechanisms as PPP connections, such as Extensible Authentication Protocol (EAP), Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge-Handshake Authentication Protocol version 2 (MS-CHAP v2), CHAP, Shiva Password Authentication Protocol (SPAP), and Password Authentication Protocol (PAP). PPTP inherits encryption, compression, or both of PPP payloads from PPP. For PPTP connections, EAP-Transport Layer Security (EAP-TLS), MS-CHAP, or MS-CHAP v2 must be used for the PPP payloads to be encrypted using Microsoft Point-to-Point Encryption (MPPE).

MPPE provides only link encryption between the VPN client and the VPN server. It does not provide end-to-end encryption, which is data encryption between the client application and the server hosting the resource or service that is being accessed by the client application. If end-to-end encryption is required, IPSec can be used to encrypt IP traffic from end-to-end after the PPTP tunnel is established.

**Tunnel Maintenance with the PPTP Control Connection**

There is a PPTP control connection between the IP address of the PPTP client using a dynamically allocated TCP port and the IP

address of the PPTP server using the reserved TCP port 1723. The PPTP control connection carries the PPTP call control and management messages that are used to maintain the PPTP tunnel. This includes the transmission of periodic PPTP Echo-Request and PPTP Echo-Reply messages to detect a connectivity failure between the PPTP client and PPTP server. PPTP control connection packets consist of an IP header, a TCP header, a PPTP control message, and a data-link trailer and header as shown in the following figure:

**PPTP Control Connection Packet**




The following table lists the primary PPTP control messages that are sent over the PPTP control connection. For all of the PPTP control messages, the specific PPTP tunnel is identified by the TCP connection.

**PPTP Call Control and Connection Management Messages**

| Message Type | Purpose |
|---|---|
| Start-Control-Connection-Request | Sent by the PPTP client to establish the control connection. Each PPTP tunnel requires a control connection to be established before any other PPTP messages can be issued. |
| Start-Control-Connection-Reply | Sent by the PPTP server to reply to the Start-Control-Connection-Request message. |
| Outgoing-Call-Request | Sent by the PPTP client to create a PPTP tunnel. Included in the Outgoing-Call-Request message is a Call ID that is used in the GRE header to identify the tunneled traffic of a specific tunnel. |
| Outgoing-Call-Reply | Sent by the PPTP server in response to the Outgoing-Call-Request message. |
| Echo-Request | Sent by either the PPTP client or PPTP server as a keep-alive mechanism. If the Echo-Request is not answered, the PPTP tunnel is eventually terminated. |
| Echo-Reply | The reply to an Echo-Request. The PPTP Echo-and Echo-Reply messages are not related to the ICMP Echo Request and Echo Reply messages. |
| WAN-Error-Notify | Sent by the PPTP server to all VPN clients to indicate error conditions on the PPP interface of the PPTP server. |
| Set-Link-Info | Sent by the PPTP client or PPTP server to set PPP-negotiated options. |
| Call-Clear-Request | Sent by the PPTP client, indicating that a tunnel is to be terminated. |
| Call-Disconnect-Notify | Sent by the PPTP server in response to a Call-Clear-Request or for other reasons to indicate that a tunnel is to be terminated. If the PPTP server terminates the tunnel, a Call-Disconnect-Notify is sent. |
| Stop-Control-Connection-Request | Sent by the PPTP client or the PPTP server to inform the other that the control connection is being terminated. |
| Stop-Control-Connection-Reply | The reply to the Stop-Control-Connection-Request message. |

For information about the exact structure of PPTP control messages, see RFC 2637 in the IETF RFC Database.

**PPTP Data Tunneling**

PPTP data tunneling is performed through multiple levels of encapsulation. The following figure shows the resulting structure of tunneled PPTP data.

**Tunneled PPTP Data**



**PPP and GRE encapsulation**
The initial PPP payload is encrypted and encapsulated with a PPP header to create a PPP frame. The PPP frame is then encapsulated with a modified GRE header. GRE is described in RFC 1701 and RFC 1702 in the IETF RFC Database and was designed to provide a simple, general purpose mechanism for encapsulating data sent over IP networks. GRE is a client protocol of IP using IP protocol 47.

For PPTP, the GRE header is modified in the following ways:

An acknowledgement bit is used to indicate that a 32-bit acknowledgement field is present and significant.

The Key field is replaced with a 16-bit Payload Length field and a 16-bit Call ID field. The Call ID field is set by the PPTP client during the creation of the PPTP tunnel.

A 32-bit Acknowledgement field is added.

Within the GRE header, the Protocol Type is set to 0x880B, the EtherType value for a PPP frame.

### Note

- GRE is sometimes used by ISPs to forward routing information within an ISP's network. To prevent the routing information from being forwarded to Internet backbone routers, ISPs filter out GRE traffic on the interfaces connected to the Internet backbone. As a result of this filtering, PPTP tunnels can be created using PPTP control messages, but tunneled PPTP data is not forwarded.

The resulting encapsulated GRE and PPP payload is then encapsulated with an IP header containing the appropriate source and destination IP addresses for the PPTP client and PPTP server.

**Encapsulation of the data-link layer**
To be sent on a local area network (LAN) or WAN link, the IP datagram is finally encapsulated with a header and trailer for the data-link layer technology of the outgoing physical interface. For example, when IP datagrams are sent on an Ethernet interface, the IP datagram is encapsulated with an Ethernet header and trailer. When IP datagrams are sent over a point-to-point WAN link, such as an analog phone line or ISDN, the IP datagram is encapsulated with a PPP header and trailer.

**Processing of the tunneled PPTP data**
Upon receipt of the tunneled PPTP data, the PPTP client or PPTP server:

- Processes and removes the data-link header and trailer.
- Processes and removes the IP header.
- Processes and removes the GRE and PPP headers.
- Decrypts and, if needed, decompresses the PPP payload.
- Processes the payload for receipt or forwarding.

**PPTP Packet Development**
The following figure shows the path that tunneled PPTP data takes through the Windows Server 2003 networking architecture from a VPN client over a remote access VPN connection using an analog modem. The following steps outline this process:

1. An IP datagram is submitted by its appropriate protocol to the virtual interface that represents the VPN connection using Network Driver Interface Specification (NDIS).
2. NDIS submits the packet to NDISWAN, which encrypts and optionally compresses the data and provides a PPP header consisting of only the PPP Protocol ID field. This assumes that address and control field compression were negotiated during the Link Control Protocol (LCP) phase of the PPP connection process.
3. NDISWAN submits the data to the PPTP protocol driver, which encapsulates the PPP frame with a GRE header. In the GRE header, the Call ID field is set to the appropriate value to identify the tunnel.
4. The PPTP protocol driver then submits the resulting packet to the TCP/IP protocol driver.

5. The TCP/IP protocol driver encapsulates the tunneled PPTP data with an IP header and submits the resulting packet to the interface that represents the dial-up connection to the local ISP using NDIS.

6. NDIS submits the packet to NDISWAN, which provides PPP headers and trailers.

7. NDISWAN submits the resulting PPP frame to the appropriate WAN miniport driver representing the dial-up hardware (for example, the asynchronous port for a modem connection).

**PPTP Packet Development**



**Resulting Packet Structure**

**Note**

- It is possible to negotiate an encrypted PPP connection for the dial-up connection with the ISP. This is unnecessary and not recommended because the private data being sent, the tunneled PPP frame, is already encrypted. The additional level of encryption is not needed and can impact performance.

**L2TP**

Layer Two Tunneling Protocol (L2TP) is a combination of PPTP and Layer 2 Forwarding (L2F), a technology developed by Cisco Systems, Inc. Rather than having two incompatible tunneling protocols competing in the marketplace and causing customer confusion, the Internet Engineering Task Force (IETF) mandated that the two technologies be combined into a single tunneling protocol that represents the best features of PPTP and L2F. L2TP is described in RFC 2661 in the IETF RFC Database.

L2TP encapsulates PPP frames to be sent over IP, X.25, frame relay, or ATM networks. When sent over an IP network, L2TP frames are encapsulated as User Datagram Protocol (UDP) messages. L2TP can be used as a tunneling protocol over the Internet or over private intranets.

L2TP uses UDP messages over IP networks for both tunnel maintenance and tunneled data. The payloads of encapsulated PPP frames can be encrypted or compressed (or both); however, L2TP clients do not negotiate the use of MPPE for L2TP connections. Encryption for L2TP connections is provided by IPSec Encapsulating Security Payload (ESP) in transport mode.

It is possible to create Windows-based L2TP connections that are not encrypted by IPSec. However, this does not apply to a VPN connection because the private data being encapsulated by L2TP is already not encrypted. Non-encrypted L2TP connections can be used temporarily to troubleshoot an L2TP over IPSec connection by eliminating the IPSec authentication and negotiation process.

L2TP for Windows assumes the availability of an IP network between an L2TP client (a VPN client using the L2TP tunneling protocol and IPSec) and an L2TP server (a VPN server using the L2TP tunneling protocol and IPSec). The L2TP client might already be attached to an IP network that can reach the L2TP server, or the L2TP client might have to use a dial-up connection to a NAS to establish IP connectivity as in the case of dial-up Internet users.

Authentication that occurs during the creation of L2TP tunnels must use the same authentication mechanisms as PPP connections.

An Internet-based L2TP server is an L2TP-enabled remote access server with one interface on the Internet and a second interface on a private intranet.

L2TP tunnel maintenance and tunneled data have the same packet structure.

**Tunnel Maintenance with L2TP Control Messages**

In contrast to PPTP, L2TP tunnel maintenance is not performed over a separate TCP connection. L2TP call control and management traffic is sent as UDP messages between the L2TP client and the L2TP server. In Windows, the L2TP client and the

L2TP server both use UDP port 1701.

**Note**

- The L2TP client and L2TP server in Windows always use UDP port 1701. The Windows Server 2003 L2TP server supports L2TP clients that use a UDP port other than 1701.

L2TP control messages over IP connections are sent as UDP datagrams. In the Windows Server 2003 implementation, L2TP control messages sent as UDP datagrams are sent as the encrypted payload of IPSec ESP transport mode as shown in the following figure.

**L2TP Control Message**




Because a TCP connection is not used, L2TP uses message sequencing to ensure delivery of L2TP messages. Within the L2TP control message, the Next-Received field (similar to the TCP Acknowledgment field) and the Next-Sent field (similar to the TCP Sequence Number field) are used to maintain the sequence of control messages. Out-of-sequence packets are dropped. The Next-Sent and Next-Received fields can also be used for sequenced delivery and flow control for tunneled data.

L2TP supports multiple calls for each tunnel. In the L2TP control message and the L2TP header for tunneled data is a Tunnel ID that identifies the tunnel and a Call ID that identifies a call within the tunnel.

The following table lists the primary L2TP control messages.

**L2TP Control Messages**

| Message Type | Purpose |
|---|---|
| Start-Control-Connection-Request | Sent by the L2TP client to establish the control connection. Each L2TP tunnel requires a control connection to be established before any other L2TP messages can be issued. It includes an Assigned Tunnel-ID that is used to identify the tunnel. |
| Start-Control-Connection-Reply | Sent by the L2TP server to reply to the Start-Control-Connection-Request message. |
| Start-Control-Connection-Connected | Sent in reply to a Start-Control-Connection-Reply message to indicate that tunnel establishment was successful. |
| Outgoing-Call-Request | Sent by the L2TP client to create an L2TP tunnel. Included in the Outgoing-Call-Request message is an Assigned Call ID that is used to identify a call within a specific tunnel. |
| Outgoing-Call-Reply | Sent by the L2TP server in response to the Outgoing-Call-Request message. |
| Start-Control-Connection-Connected | Sent in reply to a received Outgoing-Call-Reply message to indicate that the call was successful. |
| Hello | Sent by either the L2TP client or L2TP server as a keep-alive mechanism. If the Hello is not acknowledged, the L2TP tunnel is eventually terminated. |
| WAN-Error-Notify | Sent by the L2TP server to all VPN clients to indicate error conditions on the PPP interface of the L2TP server. |
| Set-Link-Info | Sent by the L2TP client or L2TP server to set PPP-negotiated options. |
| Call-Disconnect-Notify | Sent by either the L2TP server or L2TP client to indicate that a call within a tunnel is to be terminated. |

| | |
|---|---|
| Stop-Control-Connection-Notification | Sent by either the L2TP server or L2TP client to indicate that a tunnel is to be terminated. |

For the exact structure of L2TP control messages, see RFC 2661 in the IETF RFC Database.

**L2TP Data Tunneling**
L2TP data tunneling is performed using multiple levels of encapsulation. The following figure shows the resulting structure of tunneled L2TP over IPSec data.

### L2TP Packet Encapsulation




**L2TP encapsulation**
The initial PPP payload is encapsulated with a PPP header and an L2TP header.

**UDP encapsulation**
The encapsulated L2TP packet is then encapsulated with a UDP header with the source and destination ports set to 1701.

**IP encapsulation**
The UDP message is encrypted and encapsulated with an IPSec ESP header and trailer and an ESP Authentication (Auth) trailer.

**Encapsulation of the data-link layer**
To send on a LAN or WAN link, the IP datagram is finally encapsulated with a header and trailer for the data-link layer technology of the outgoing physical interface. For example, when an IP datagram is sent on an Ethernet interface, the IP datagram is encapsulated with an Ethernet header and trailer. When an IP datagram is sent over a point-to-point WAN link such as an analog phone line or ISDN, the IP datagram is encapsulated with a PPP header and trailer.

**Processing of the tunneled L2TP/IPSec data**
Upon receipt of the tunneled L2TP/IPSec data, the L2TP client or L2TP server:

- Processes and removes the data-link header and trailer.
- Processes and removes the IP header.
- Uses the IPSec ESP Auth trailer to authenticate the IP payload and the IPSec ESP header.
- Uses the IPSec ESP header to decrypt the encrypted portion of the packet.
- Processes the UDP header and sends the L2TP packet to the L2TP driver.
- Uses the Tunnel ID and Call ID in the L2TP header to identify the specific L2TP tunnel.
- Uses the PPP header to identify the PPP payload and forward it to the proper protocol driver for processing.

**L2TP with Internet Protocol security (L2TP/IPSec)**
Tunneling protocols such as PPTP and L2TP are implemented at the data-link layer of the Open Systems Interconnection (OSI) reference model and provide data security by helping to create secure tunnels. In contrast, the IPSec protocol is implemented at the network layer and helps secure data at the packet level. IPSec provides two security protocols: Authentication Header (AH) and ESP.

**IPSec ESP encapsulation protocol**
To provide maximum security for L2TP/IPSec packets, ESP can also be used to encapsulate IPSec packets.

**L2TP Packet Development**
The figure below shows the path that tunneled L2TP data takes through the Windows Server 2003 networking architecture from a VPN client over a remote access VPN connection using an analog modem. The following steps outline the process:

1. An IP datagram is submitted by the appropriate protocol to the virtual interface that represents the VPN connection

using NDIS.

2. NDIS submits a packet to NDISWAN, which optionally compresses and provides a PPP header consisting of only the PPP Protocol ID field. This assumes that address and control field compression were negotiated during the LCP phase of the PPP connection process.

3. NDISWAN submits the PPP frame to the L2TP protocol driver, which encapsulates the PPP frame with an L2TP header. In the L2TP header, the Tunnel ID and the Call ID are set to the appropriate value identifying the specific L2TP connection.

4. The L2TP protocol driver then submits the resulting packet to the TCP/IP protocol driver with information to send the L2TP packet as a UDP message from UDP port 1701 to UDP port 1701 with the IP addresses of the VPN client and the VPN server.

5. The TCP/IP protocol driver constructs an IP packet with the appropriate IP header and UDP header. IPSec then analyzes the IP packet and matches it with a current IPSec policy. Based on the settings in the policy, IPSec encapsulates and encrypts the UDP message portion of the IP packet using the appropriate ESP headers and trailers.

6. The original IP header with the Protocol field set to 50 is added to the front of the ESP payload.

7. Using NDIS, the TCP/IP protocol driver then submits the resulting packet to the interface that represents the dial-up connection to the local ISP using NDIS.

8. NDIS submits the packet to NDISWAN.

9. NSIDWAN provides PPP headers and trailers and submits the resulting PPP frame to the appropriate WAN miniport driver representing the dial-up hardware.

**L2TP Packet Development**




**Resulting Packet Structure**

**Note**

- It is possible to negotiate an encrypted PPP connection for the dial-up connection with an ISP. This is not necessary and not recommended because the private data being sent, the tunneled PPP frame, is already encrypted with IPSec. The additional level of encryption is not needed and can impact performance.

**VPN Authentication**

The VPN server can be configured to use either Windows or Remote Authentication Dial-In User Service (RADIUS) as an authentication provider. If Windows is selected as the authentication provider, the user credentials sent by users attempting VPN connections are authenticated using typical Windows authentication mechanisms, and the connection attempt is authorized using the VPN client's user account properties and local remote access policies.

If RADIUS is selected and configured as the authentication provider on the VPN server, user credentials and parameters of the connection request are sent as RADIUS request messages to a RADIUS server.

The RADIUS server receives a user-connection request from the VPN server and authenticates and authorizes the connection attempt. In addition to a yes or no response to an authentication request, RADIUS can inform the VPN server of other applicable connection parameters for this user such as maximum session time, static IP address assignment, and so on.

RADIUS can respond to authentication requests based on its own user account database, or it can be a front end to another database server, such as a Structured Query Language (SQL) server or a Windows domain controller (DC). The DC can be located on the same computer as the RADIUS server or elsewhere. In addition, a RADIUS server can act as a proxy client to a remote RADIUS server.

The RADIUS protocol is described in RFC 2865 and RFC 2866 in the IETF RFC Database.

The VPN server can be configured to use either Windows or RADIUS as an accounting provider. If Windows is selected as the accounting provider, the accounting information accumulates on the VPN server for later analysis. Logging options can be specified from the properties of the **Local File** or **SQL Server** objects in the **Remote Access Logging** folder in the Routing and Remote Access snap-in. If RADIUS is selected, RADIUS accounting messages are sent to the RADIUS server for accumulation and later analysis.

Most RADIUS servers can be configured to place authentication request records into an audit file. A number of third parties have written billing and audit packages that read RADIUS accounting records and produce various useful reports. For more information about RADIUS accounting, see RFC 2866 in the IETF RFC Database.

The VPN server can be managed using industry-standard network management protocols and infrastructure. The computer acting as the VPN server can participate in a Simple Network Management Protocol (SNMP) environment as an SNMP agent if the Windows Server 2003 SNMP service is installed. The VPN server records management information in various object identifiers of the Internet Management Information Base (MIB) II, which is installed with the Windows Server 2003 SNMP service. Objects in the Internet MIB II are documented in RFC 1213 in the IETF RFC Database.

**Authentication Protocols**

The following authentication protocols are used to identify VPN users and grant or deny user access to network resources based on the user's credentials.

**PAP**

Password Authentication Protocol (PAP) is a clear-text authentication scheme. The NAS requests the user name and password, and PAP returns them in clear text (unencrypted). Obviously, this authentication scheme is not secure because a malicious user could capture the user's name and password and use it to get subsequent access to the NAS and all of the resources provided by the NAS. PAP provides no protection against replay attacks or remote client impersonation once the user's password is compromised.

**SPAP**

The Shiva Password Authentication Protocol (SPAP) is a reversible encryption mechanism employed by Shiva Corporation. A computer running Windows XP Professional uses SPAP when connecting to a Shiva LAN Rover. A Shiva client that connects to a server running Routing and Remote Access also uses SPAP. Currently, this form of authentication is more secure than plaintext but less secure than CHAP or MS-CHAP.

**CHAP**

Challenge Handshake Authentication Protocol (CHAP) is an encrypted authentication mechanism that prevents transmission of the actual password on the connection. The NAS sends a challenge, which consists of a session ID and an arbitrary challenge string, to the remote client. The remote client must use the MD5 one-way hashing algorithm to return the user name and a hash of the challenge, session ID, and the client’s password. The user name is sent as plain text.

CHAP is an improvement over PAP because the clear-text password is not sent over the link. Instead, the password is used to create a hash from the original challenge. The server knows the client’s clear-text password and can, therefore, replicate the operation and compare the result to the password sent in the client’s response. CHAP protects against replay attacks by using an arbitrary challenge string for each authentication attempt. CHAP protects against remote-client impersonation by unpredictably sending repeated challenges to the remote client throughout the duration of the connection.

**MS-CHAP**

Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) is an encrypted authentication mechanism very similar to CHAP. As in CHAP, the NAS sends a challenge, which consists of a session ID and an arbitrary challenge string, to the remote client. The remote client must return the user name and an encrypted form of the challenge string, the session ID, and the MD4-hashed password. This design, which uses the MD4 hash of the password, helps provides an additional level of security because it allows the server to store hashed passwords instead of clear-text passwords or passwords that are stored using reversible encryption. MS-CHAP also provides additional error codes, including a password-expired code, and additional encrypted client-server messages that permit users to change their passwords during the authentication process. In MS-CHAP, both the client and the NAS independently generate a common initial encryption key for subsequent data encryption by MPPE.

**MS-CHAP v2**

MS-CHAP version 2 (MS-CHAP v2) is an updated encrypted authentication mechanism that provides stronger security for the exchange of user name and password credentials and determination of encryption keys. With MS-CHAP v2, the NAS sends a challenge to the client that consists of a session identifier and an arbitrary challenge string. The remote access client sends a

response that contains the user name, an arbitrary peer challenge string, and an encrypted form of the received challenge string, the peer challenge string, the session identifier, and the user's password. The NAS checks the response from the client and sends back a response containing an indication of the success or failure of the connection attempt and an authenticated response based on the sent challenge string, the peer challenge string, the encrypted response of the client, and the user's password. The remote access client verifies the authentication response and, if correct, uses the connection. If the authentication response is not correct, the remote access client terminates the connection.

Using this process, MS-CHAP v2 provides mutual authentication — the NAS verifies that the client has knowledge of the user's password, and the client verifies that the NAS has knowledge of the user's password. MS-CHAP v2 also determines two MPPE encryption keys, one for data sent and one for data received.

**EAP**

Extensible Authentication Protocol (EAP) is a PPP authentication protocol that allows for an arbitrary authentication method. EAP differs from the other authentication protocols in that, during the authentication phase, EAP does not actually perform authentication. Phase 2 for EAP only negotiates the use of a common EAP authentication method (known as an EAP type). The actual authentication for the negotiated EAP type is performed after Phase 2.

During phase 2 of PPP link configuration, the NAS collects the authentication data and then validates the data against its own user database or a central authentication database server, such as one maintained by a Windows domain controller, or the authentication data is sent to a RADIUS server.

As stated previously, most implementations of PPP provide a limited number of authentication methods. EAP is an IETF standard extension to PPP that allows for arbitrary authentication mechanisms for the validation of a PPP connection. EAP was designed to allow the dynamic addition of authentication plug-in modules at both the client and authentication server. This allows vendors to supply a new authentication scheme at any time. EAP provides the highest flexibility in authentication uniqueness and variation.

EAP is documented in RFC 2284 in the IETF RFC Database, and is supported in Windows Server 2003, Windows XP, and Windows 2000.

**EAP-MD5 Challenge**

Extensible Authentication Protocol-Message Digest 5 Challenge (EAP-MD5 Challenge) is a required EAP type that uses the same challenge handshake protocol as PPP-based CHAP, but the challenges and responses are sent as EAP messages. A typical use for EAP-MD5 Challenge is to authenticate the credentials of remote access clients by using user name and password security systems. EAP-MD5 Challenge can be used to test EAP interoperability.

**EAP-TLS**

Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) is an EAP type that is used in certificate-based security environments. If smart cards are used for remote access authentication, EAP-TLS is the required authentication method. The EAP-TLS exchange of messages provides mutual authentication, negotiation of the encryption method, and encrypted key determination between the remote access client and the authenticator. EAP-TLS provides the strongest authentication and key-determination method.

When the Routing and Remote Access service is configured to use Windows authentication, EAP-TLS is supported only when the VPN server is a member of a domain. A VPN server running as a stand-alone server or a member of a workgroup does not support EAP-TLS.

EAP-TLS is an IETF standard (RFC 2716 in the IETF RFC Database for a strong authentication method based on public-key certificates. With EAP-TLS, a client presents a user certificate to the server, and the server presents a server certificate to the client. The first provides strong user authentication to the server; the second provides assurance that the VPN client has reached a trusted VPN server. Both systems rely on a chain of trusted certification authorities (CAs) to verify the validity of the offered certificate.

The user's certificate could be stored on the VPN client computer or in an external smart card. In either case, the certificate cannot be accessed without some form of user identification (PIN number or name/password credentials) between the user and the client computer. This approach meets the something-you-know-plus-something-you-have criteria recommended by most security experts.

EAP-TLS is supported in Windows Server 2003 and Windows XP. Like MS-CHAP and MS-CHAP v2, EAP-TLS returns an encryption key to enable subsequent data encryption by MPPE.

**RADIUS**

The Remote Authentication Dial-In User Service (RADIUS) protocol is used to provide centralized administration of authentication, authorization, and accounting (AAA) and an industry-standard security infrastructure. RADIUS is defined in RFCs 2138 and 2139 in the IETF RFC Database. RADIUS enables administrators to manage a set of authorization policies, accumulate accounting information, and access an account database from a central location.

Because it is impossible to update separate user accounts on separate servers for the same user simultaneously, most administrators set up a master account database at a domain controller or on a RADIUS server. This enables the VPN server to send the authentication credentials to a central authenticating device, and the same user account can be used for both dial-up remote access and VPN-based remote access.

**VPN Encryption**

To help ensure confidentiality of the data as it traverses the shared or public transit network, it is encrypted by the sender and decrypted by the receiver. Because data encryption is performed between the VPN client and VPN server, it is not necessary to use data encryption on the communication link between a dial-up client and its Internet service provider (ISP). For example, a mobile user uses a dial-up networking connection to dial in to a local ISP. Once the Internet connection is made, the user creates a VPN connection with the corporate VPN server. If the VPN connection is encrypted, there is no need to use encryption on the dial-up networking connection between the client and the ISP.

Remote access data encryption does not provide end-to-end data encryption. End-to-end encryption is data encryption between the client application and the server that hosts the resource or service being accessed by the client application. To get end-to-end data encryption, use IPSec to help create a secure connection after the remote access connection has been made.

Data encryption for PPP or PPTP connections is available only if MS-CHAP, MS-CHAP v2, or EAP-TLS is used as the authentication protocol. Data encryption for L2TP connections relies on IPSec, which does not require a specific PPP-based authentication protocol.

The encryption and decryption processes depend on both the sender and the receiver having knowledge of a common encryption key. Intercepted packets sent along the VPN connection in the transit network are unintelligible to any computer that does not have the common encryption key. The length of the encryption key is an important security parameter. Computational techniques can be used to determine the encryption key. Such techniques require more computing power and computational time as the encryption key gets larger. Therefore, it is important to use the largest possible key size.

In addition, the more information that is encrypted with the same key, the easier it is to decipher the encrypted data. With some encryption technologies, administrators are given the option to configure how often the encryption keys are changed during a connection.

PPTP uses user-level PPP authentication methods and Microsoft Point-to-Point Encryption (MPPE) for data encryption.

**Data Encryption with MPPE**

PPTP inherits MPPE encryption, which uses the Rivest-Shamir-Adleman (RSA) RC4 stream cipher. MPPE is available only for PPTP-based VPN connections when the EAP-TLS, MS-CHAP, or MS-CHAP v2 authentication protocols are used.

For the Routing and Remote Access service, MPPE encryption strengths are configured on the Encryption tab on the properties of a remote access policy to use 40-bit (the Basic setting), 56-bit (the Strong setting), or 128-bit (the Strongest setting) encryption keys. Administrators should use 40-bit MPPE encryption keys to connect with older operating systems that do not support 56-bit or 128-bit encryption keys (this includes older Windows operating systems and operating systems from companies other than Microsoft). Otherwise, use 128-bit encryption keys. Encryption strengths for L2TP/IPSec connections use 56-bit DES (the Basic or Strong setting) or 168-bit 3DES (the Strongest setting).

Encryption keys are determined at the time of the connection. By default, the highest key strength supported by the VPN client and VPN server is negotiated during the process of establishing a connection. If the VPN server requires a higher key strength than is supported by the VPN client, the connection attempt is rejected.

MPPE was originally designed for encryption across a point-to-point link where packets arrive in the same order in which they were sent with little packet loss. For this environment, the decryption of each packet depends on the decryption of the previous packet.

For VPN connections, however, IP datagrams sent across the Internet can arrive in a different order from the one in which they were sent, and a higher proportion of packets can be lost. Therefore, for VPN connections, MPPE changes the encryption key for each packet. The decryption of each packet is independent of the previous packet. MPPE includes a sequence number in the MPPE header. If packets are lost or arrive out of order, the encryption keys are changed relative to the sequence number.

**VPN Addressing and Routing**

Based on whether or not a route is added by default, a VPN client has broad access to Internet locations or to locations on the intranet, but not to both:

- If the currently active default route is pointing to the Internet (and the gateway on the remote network is not being used), Internet locations are reachable, but only intranet locations matching the network ID corresponding to the Internet address class of the assigned IP address can be reached.

- If the currently active default route is pointing to the intranet (and the gateway on the remote network is being used), all intranet locations are reachable, but only the IP address of the VPN server and locations available through other routes can be reached on the Internet.

For most VPN clients with an Internet connection, this does not present a problem, because the client is typically engaged in either intranet communication or Internet communication, but not both.

To work around this problem, instead of having the client create a new default route when a connection is made, administrators can configure the client's routing table with specific routes that direct packets to the organization's network over the VPN connection. While connected to the intranet, the client can obtain Internet access using the default route that points to the Internet. This configuration is known as split tunneling.

**Split Tunneling**
The VPN client can obtain the routes needed for split tunneling in several ways:

- If the VPN client has a configured connection without a default route, the client adds a route that it infers from the Internet address class of the IP address assigned to it for the current connection. For a simple target network, such as a small office, this one route is sufficient to allow packets to be routed to the target network. However, for a complex network, administrators need to configure multiple routes to successfully direct packets to the remote network.

- A client running the Microsoft Windows XP or Windows Server 2003 operating systems uses a DHCPINFORM message after the connection to request the DHCP Classless Static Routes option. This DHCP option contains a set of routes that are automatically added to the routing table of the requesting client. This additional information is available only if the Windows Server 2003 DHCP server has been configured to provide the DHCP Classless Static Routes option and if the VPN server has the DHCP Relay Agent routing protocol component configured with the IP address of the DHCP server.

- If the remote access client is managed using the Connection Manager component of Windows Server 2003, the network administrator can configure routing table updates from the Routing Table Update page of the Connection Manager Administration Kit when creating the Connection Manager profile.

If none of the approaches discussed above is an option, a batch file or program can be written that updates the routing table on the client with the necessary routes to the private intranet.

**Security Considerations for Split Tunneling**
When a VPN client computer is connected to both the Internet and a private intranet and has routes that allow it to reach both networks, the possibility exists that a malicious Internet user might use the connected VPN client computer to reach the private intranet through the authenticated VPN connection. This is possible if the VPN client computer has IP routing enabled. IP routing is enabled on Windows XP-based computers by setting the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip \Parameters\IPEnableRouter registry entry to 1 (data type is REG_DWORD).

If split tunneling is required, administrators can help prevent a malicious user from gaining access over the Internet by doing the following:

- Use the Network Access Quarantine Control feature in Windows Server 2003 to check whether VPN clients have IP routing enabled and, if so, do not allow VPN access until it has been disabled.

- Use IP packet filters on the VPN remote access policy profile to discard both inbound traffic on the VPN connection that has not been sent from the VPN client and outbound traffic that is not destined to the VPN client. The default remote access policy, named "Connections to Microsoft Routing and Remote Access server in Windows Server 2003" has these packet filters configured and enabled by default.

  **Note**

  - Using the methods above does not prevent unwanted traffic if a malicious Internet user is remotely controlling the VPN client computer. To prevent this, ensure that the VPN client computer has a firewall enabled (such as Internet Connection Firewall in Windows XP) and an anti-virus program installed and running with the latest virus signature file installed. These are also settings that can be enabled and enforced when using Network Access Quarantine Control.

**DHCP Classless Static Routes Option**
Classless static routes are implemented using DHCP scope option 249. Using classless static routes, each DHCP client can be configured with the route to any destination on the network, and the subnet mask can be specified. Because each scope represents a physical subnet, the scope can be viewed as the start location for any message that is to be sent by a client to another subnet. The parameters used to configure option 249 are Destination, Mask, and Router. One or more static routes can be configured with option 249. All DHCP-enabled clients on the network can be provided with routes to all other subnets using option 249.

This option is configured as a scope option because the Router IP address, like the DHCP Router option that defines the default gateway for a DHCP client, is different for each subnet. For example, subnets A and D each use a router. The routers they use will be different, and the Router IP address will be different in each case.

**Static Routing**
Static routing requires that routing tables be configured and maintained manually. Static routers do not exchange information. Because of this limitation, when compared to dynamic routing, static routing is typically implemented in small networks or in networks that require the highest level of security.

**Auto-Static Updates**
If routing protocols are not used to update the routing tables, then the routes must be entered as static routes. The static routes that correspond to the network IDs available across the interface are entered manually or automatically. The automatic entering of static routes for demand-dial interfaces is known as making auto-static updates and is supported by the server running Routing and Remote Access. Auto-static updates are supported by Routing Information Protocol (RIP) for IP, but not by OSPF.

Auto-static refers to the automatic adding of the requested routes as static routes in the routing table. The sending of the request for routes is performed through an explicit action, either through Routing and Remote Access or the Netsh utility while the demand-dial interface is in a connected state. Auto-static updates are not automatically performed every time a demand-dial connection is made.

When instructed, a demand-dial interface that is configured for auto-static updates sends a request across an active connection to request all of the routes of the router on the other side of the connection. In response to the request, all of the routes of the requested router are automatically entered as static routes in the routing table of the requesting router. The static routes are persistent: They are kept in the routing table even if the interface becomes disconnected or the router is restarted. An auto-static update is a one-time, one-way exchange of routing information.

Administrators can automate and schedule auto-static updates by executing the update as a scheduled task. When an auto-static update is requested, the existing auto-static routes are deleted before the update is requested from other routers. If there is no response to the request, then the router cannot replace the routes it has deleted. This might lead to a loss of connectivity to remote networks.

**Dynamic Routing**
By implementing a dynamic routing protocol, such as RIP or Open Shortest Path First (OSPF), administrators can configure routers to exchange routing information with each other as needed.

**RIP**
The biggest advantage of RIP is that it is extremely simple to configure and deploy. The biggest disadvantage of RIP is its inability to scale to large or very large networks. The maximum hop count used by RIP routers is 15. Networks that are 16 hops or more away are considered unreachable. As networks grow larger in size, the periodic announcements by each RIP router can cause excessive traffic. Another disadvantage of RIP is its high recovery time. When the network topology changes, it might take several minutes before the RIP routers reconfigure themselves to the new network topology. While the network reconfigures itself, routing loops might form that result in lost or undeliverable data.

Initially, the routing table for each router includes only the networks that are physically connected. A RIP router periodically sends announcements that contain its routing table entries to inform other local RIP routers of the networks it can reach. RIP version 1 uses IP broadcast packets for its announcements. RIP version 2 can use multicast or broadcast packets for its announcements.

RIP routers can also communicate routing information through triggered updates. Triggered updates occur when the network topology changes and updated routing information is sent that reflects those changes. With triggered updates, the update is sent immediately rather than waiting for the next periodic announcement. For example, when a router detects a link or router failure, it updates its own routing table and sends updated routes. Each router that receives the triggered update modifies its own routing table and propagates the change.

Routing and Remote Access supports RIP versions 1 and 2. RIP version 2 supports multicast announcements, simple password authentication, and more flexibility in subnetted and Classless InterDomain Routing (CIDR) environments.

**OSPF**
The biggest advantage of OSPF is that it is efficient; OSPF requires very little network overhead even in very large networks. The biggest disadvantage of OSPF is its complexity; OSPF requires proper planning and is more difficult to configure and administer.

OSPF uses a Shortest Path First (SPF) algorithm to compute routes in the routing table. The SPF algorithm computes the shortest (least cost) path between the router and all the subnets of the network. SPF-calculated routes are always loop-free.

Changes to network topology are efficiently flooded across the entire network to ensure that the link state database on each router is synchronized and accurate at all times. Upon receiving changes to the link state database, the routing table is recalculated.

As the size of the link state database increases, memory requirements and route computation times increase. To address this scaling problem, OSPF divides the network into areas (collections of contiguous networks) that are connected to each other through a backbone area. Each router only keeps a link state database for those areas that are connected to the router. Area border routers (ABRs) connect the backbone area to other areas.

**Single-Adapter model**
With the single-adapter model, also known as the NBMA model, the network for the frame relay service provider (also known as the frame relay cloud) is treated as an IP network and the endpoints on the cloud are assigned IP addresses from a designated IP network ID. To ensure that OSPF traffic is received by all of the appropriate endpoints on the cloud, the frame relay interface must be configured to send unicast OSPF announcements to all of the appropriate endpoints. For the server running Routing and Remote Access, this is done by designating the interface as an NBMA network and adding OSPF neighbors.

In addition, in a spoke and hub frame relay topology, the frame relay interface for the hub router must have a router priority set to 1 or greater and the frame relay interfaces for the spoke routers must have a router priority set to 0. Otherwise, the hub router, which is the only router that can communicate with all of the spoke routers, cannot become the designated router and adjacencies cannot form across the frame relay network.

**Multiple-Adapter model**
With the multiple-adapter model, each frame relay virtual circuit appears as a point-to-point link with its own network ID, and the endpoints are assigned IP addresses from a designated IP network ID. Because each virtual circuit is its own point-to-point connection, administrators can configure the interface for the point-to-point network type.

**Virtual links**
An OSPF-routed network can be subdivided into areas, which are collections of contiguous networks. All areas are connected together through a common area called the backbone area. A router that connects an area to the backbone area is called an area border router (ABR). Normally, ABRs have a physical connection to the backbone area. When it is not possible or practical to have an ABR physically connected to the backbone area, administrators can use a virtual link to connect the ABR to the backbone.

A virtual link is a logical point-to-point connection between an ABR of an area and an ABR that is physically connected to the backbone area. For example, a virtual link is configured between the ABR of Area 2 and the ABR of Area 1. The ABR of Area 1 is physically connected to the backbone area. Area 1 is known as the transit area, the area across which the virtual link is created in order to logically connect Area 2 to the backbone.

To create a virtual link, both routers, called virtual link neighbors, are configured with the transit area, the router ID of the virtual link neighbor, matching hello and dead intervals, and a matching password.

**External routes and ASBRs**
The set of OSPF routers in an organization defines an OSPF autonomous system (AS). By default, only OSPF routes corresponding to directly-connected network segments are propagated within the AS. An external route is any route that is not within the OSPF AS. External routes can come from many sources:

- Other routing protocols such as RIP for IP (version 1 and version 2)
- Static routes
- Routes set on the router through SNMP

External routes are propagated throughout the OSPF AS through one or more autonomous system boundary routers (ASBRs). An ASBR advertises external routes within the OSPF AS. For example, if the static routes of a server running Routing and Remote Access need to be advertised, that router must be enabled as an ASBR.

**External route filters**
By default, OSPF routers acting as ASBRs import and advertise all external routes. Administrators might want to filter out external routes to keep the ASBR from advertising improper routes. External routes can be filtered on the ASBR by:

**External route source**
Administrators can configure the ASBR to accept or ignore the routes of certain external sources such as routing protocols (RIP version 2) or other sources (static routes or SNMP).

**Individual route**
Administrators can configure the ASBR to accept or discard specific routes by configuring one or multiple destination, network mask pairs.

**VPN and Firewalls**
A firewall uses packet filtering to allow or disallow the flow of specific types of network traffic. IP packet filtering provides a way for administrators to define precisely what IP traffic is allowed to cross the firewall. IP packet filtering is important when private intranets are connected to public networks, such as the Internet. There are two approaches to using a firewall with a VPN server:

- A firewall is between the VPN server and the Internet. In this configuration, the VPN server is behind the firewall.
- The VPN server is connected to the Internet and the firewall is between the VPN server and the intranet. In this configuration, the VPN server is in front of the firewall.

**VPN Server Behind a Firewall**
In the configuration shown in the following figure, the firewall is connected to the Internet and the VPN server is another intranet resource connected to the perimeter network, also known as a screened subnet or demilitarized zone (DMZ). The perimeter network is an IP network segment that typically contains resources available to Internet users such as Web servers and FTP servers. The VPN server has an interface on the perimeter network and an interface on the intranet.

In this approach, the firewall must be configured with input and output filters on its Internet and perimeter network interfaces to allow the passing of tunnel maintenance traffic and tunneled data to the VPN server. Additional filters can allow the passing of traffic to Web servers, FTP servers, and other types of servers on the perimeter network. As an added layer of security, the VPN server should also be configured with PPTP or L2TP/IPSec packet filters on its perimeter network interface as described in "VPN Server in Front of a Firewall" in this section.

Because the firewall does not have the encryption keys for each VPN connection, it can only filter on the plaintext headers of the tunneled data, meaning that all tunneled data passes through the firewall. However, this is not a security concern because the VPN connection requires an authentication process that prevents unauthorized access beyond the VPN server.

### VPN Server Behind the Firewall




**Packet Filters for a VPN Server Behind a Firewall**
If the VPN server is behind a firewall, packet filters must be configured for both an Internet interface and a perimeter network interface. In this scenario, the firewall is connected to the Internet, and the VPN server is an intranet resource that is connected to the perimeter network. The VPN server has an interface on both the perimeter network and the Internet.

**PPTP connections for the Internet interface of the firewall**
The following table shows the inbound and outbound PPTP filters on the firewall's Internet interface.

### VPN Server Behind a Firewall: PPTP Filters on the Firewall's Internet Interface

| Filter Type | Filter | Action |
|---|---|---|

| | | |
|---|---|---|
| Inbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**TCP destination port** = 1723 (0x6BB) | Allows PPTP tunnel maintenance traffic from the PPTP client to the PPTP server. |
| Inbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 47 (0x2F) | Allows tunneled PPTP data from the PPTP client to the PPTP server. |
| Inbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. If all traffic from TCP port 1723 is allowed to reach the VPN server, network attacks can emanate from sources on the Internet that use this port. Administrators should only use this filter in conjunction with the PPTP filters that are also configured on the VPN server. |
| Outbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Allows PPTP tunnel maintenance traffic from the PPTP server to the PPTP client. |
| Outbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 47 (0x2F) | Allows tunneled PPTP data from the PPTP server to the PPTP client. |
| Outbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**TCP destination port** = 1723 (0x6BB) | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. If all traffic from the VPN server is allowed to reach TCP port 1723, network attacks can emanate from sources on the Internet using this port. Administrators should only use this filter in conjunction with the PPTP filters that are also configured on the VPN server. |

**PPTP connections for the perimeter network interface of the firewall**

The following table shows the inbound and outbound PPTP filters on the firewall's perimeter network interface.

**VPN Server Behind a Firewall: PPTP Filters on the Perimeter Network Interface**

| Filter Type | Filter | Action |
|---|---|---|

| | | |
|---|---|---|
| Inbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Allows PPTP tunnel maintenance traffic from the VPN server to the VPN client. |
| Inbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 47 (0x2F) | Allows tunneled PPTP data from the VPN server to the VPN client. |
| Inbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**TCP destination port** = 1723 (0x6BB) | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. If all traffic from TCP port 1723 is allowed to reach the VPN server, network attacks can emanate from sources on the Internet using this port. |
| Outbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Allows PPTP tunnel maintenance traffic from the PPTP client to the PPTP server. |
| Outbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 47 (0x2F) | Allows tunneled PPTP data from the PPTP client to the PPTP server. |
| Outbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**TCP source port** = 1723 (0x6BB) | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. If all traffic from the VPN server is allowed to reach TCP port 1723, network attacks can emanate from sources on the Internet using this port. |

**L2TP/IPSec connections for the Internet interface of the firewall**

The following table shows the inbound and outbound L2TP/IPSec filters on the firewall's Internet interface.

**VPN Server Behind a Firewall: L2TP/IPSec Filters on the Firewall's Internet Interface**

| Filter Type | Filter | Action |
|---|---|---|
| Inbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**UDP destination port** = 500 (0x1F4) | Allows IKE traffic to the VPN server. |
| Inbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**UDP destination port** = 4500 (0x1194) | Allows IPSec NAT-T traffic to the VPN server. |
| Inbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 50 (0x32) | Allows IPSec ESP traffic to the VPN server. |

| | | |
|---|---|---|
| Outbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**UDP source port** = 500 (0x1F4) | Allows IKE traffic from the VPN server. |
| Outbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**UDP source port** = 4500 (0x1194) | Allows IPSec NAT-T traffic from the VPN server. |
| Outbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 50 (0x32) | Allows IPSec ESP traffic from the VPN server. |

No filters are required for L2TP traffic at UDP port 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted with IPSec ESP.

**L2TP/IPSec connections for the perimeter network interface of the firewall**

The following table shows the inbound and outbound L2TP/IPSec filters on the firewall's perimeter network interface.

**VPN Server Behind a Firewall: L2TP/IPSec Filters on the Firewall's Perimeter Network Interface**

| Filter Type | Filter | Action |
|---|---|---|
| Inbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**UDP source port** = 500 (0x1F4) | Allows IKE traffic from the VPN server. |
| Inbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**UDP source port** = 4500 (0x1194) | Allows IPSec NAT-T traffic from the VPN server. |
| Inbound | **Source IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 50 (0x32) | Allows IPSec ESP traffic from the VPN server. |
| Outbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**UDP destination port** = 500 (0x1F4) | Allows IKE traffic to the VPN server. |
| Outbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**UDP destination port** = 4500 (0x1194) | Allows IPSec NAT-T traffic to the VPN server. |
| Outbound | **Destination IP address** = Perimeter network interface of VPN server<br><br>**IP Protocol ID** = 50 (0x32) | Allows IPSec ESP traffic to the VPN server. |

**VPN Server in Front of a Firewall**

With the VPN server in front of the firewall and connected to the Internet, as shown in the following figure, administrators need to add packet filters to the Internet interface that allow only VPN traffic to and from the IP address of the VPN server's interface on the Internet.

For inbound traffic, when the tunneled data is decrypted by the VPN server it is forwarded to the firewall, which employs its filters to allow the traffic to be forwarded to intranet resources. Because the only traffic that is crossing the VPN server is traffic generated by authenticated VPN clients, firewall filtering in this scenario can be used to prevent VPN users from accessing specific intranet resources.

Because the only Internet traffic allowed on the intranet must go through the VPN server, this approach also prevents the sharing of intranet resources with non-VPN Internet users.

**VPN Server in Front of the Firewall**




**Packet Filters for a VPN Server in Front of a Firewall**

When a VPN server is in front of a firewall and connected to the Internet, inbound and outbound packet filters on the VPN server need to be configured to allow only VPN traffic to and from the IP address of the VPN server's Internet interface. Use this configuration if the VPN server is in a perimeter network, with one firewall positioned between the VPN server and the intranet and another between the VPN server and the Internet.

All of the following packet filters are configured, using the Routing and Remote Access snap-in, as IP packet filters on the Internet interface. Depending on the configuration decisions made when running the Routing and Remote Access Server Setup Wizard, these packet filters might already be configured.

**PPTP connections for the inbound and outbound filters**

The following table shows the VPN server's inbound and outbound filters for PPTP.

**VPN Server in Front of a Firewall: PPTP Packet Filters on the Internet Interface**

| Filter Type | Filter | Action |
|---|---|---|
| Inbound | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**TCP destination port** = 1723 | Allows PPTP tunnel maintenance to the VPN server. |
| Inbound | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**IP Protocol ID** = 47 | Allows tunneled PPTP data to the VPN server. |
| Inbound | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**TCP (established) source port** = 1723 | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. Accepts TCP traffic only when a VPN server initiates the TCP connection. |
| Outbound | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**TCP source port** = 1723 | Allows PPTP tunnel maintenance traffic from the VPN server. |

| | | |
|---|---|---|
| Outbound | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**IP Protocol ID** = 47 | Allows tunneled PPTP data from the VPN server. |
| Outbound | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**TCP (established) destination port** = 1723 | Required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. Sends TCP traffic only when a VPN server initiates the TCP connection. |

**L2TP/IPSec connections**

The following table shows the VPN server's inbound and outbound filters for L2TP/IPSec.

**VPN Server in Front of a Firewall: L2TP/IPSec Packet Filters on the Internet Interface**

| Filter Type | Filter | Action |
|---|---|---|
| Inbound | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP destination port** = 500 | Allows IKE traffic to the VPN server. |
| Inbound | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP destination port** = 1701 | Allows L2TP traffic from the VPN client to the VPN server. |
| Inbound | **Destination IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP destination port** = 4500 | Allows IPSec NAT-T traffic from the VPN client to the VPN server. |
| Outbound | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP source port** = 500 | Allows IKE traffic from the VPN server. |
| Outbound | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP source port** = 1701 | Allows L2TP traffic from the VPN server to the VPN client. |
| Outbound | **Source IP address** = Internet interface of VPN server<br><br>**Subnet mask** = 255.255.255.255<br><br>**UDP source port** = 4500 | Allows IPSec NAT-T traffic from the VPN server to the VPN client |

**VPN and NAT**

A network address translator (NAT) is a device that is typically used to provide shared access for private networks to a public network such as the Internet. Because NAT does not work with protocols that use encryption, a VPN solution that includes a NAT can add a layer of complexity to a VPN deployment.

**NAT with PPTP Connections**
If a VPN client that uses a PPTP connection is behind a NAT, the NAT must include a NAT editor that can translate PPTP traffic. The NAT editor is required because tunneled PPTP data has a GRE header rather than a TCP header or a UDP header. The NAT editor uses the Call ID field in the GRE header to identify the PPTP data stream and translate IP addresses and call IDs for PPTP data packets that are forwarded between a private network and the Internet.

**PPTP NAT editor**
The NAT/Basic Firewall routing protocol component of the Routing and Remote Access service and the Internet Connection Sharing feature of Network Connections includes a NAT editor for PPTP traffic.

**NAT with L2TP Connections**
To use L2TP-based VPN connections behind a NAT, IPSec NAT Traversal (NAT-T) must be implemented at both ends of the VPN connection.

**IPSec NAT-T**
IPSec NAT-T addresses the difficulty of using IPSec-based VPNs across a NAT. Windows Server 2003 allows an L2TP/IPSec connection to pass through a NAT. This capability is based on the latest IETF standards.

IPSec NAT-T enables IPSec peers to negotiate and communicate when they are behind a NAT. To use IPSec NAT-T, both the remote access VPN client and the remote access VPN server must support IPSec NAT-T. IPSec NAT-T is supported by the Windows Server 2003 Microsoft L2TP/IPSec VPN Client and by the L2TP/IPSec NAT-T Update for Windows XP and the L2TP/IPSec NAT-T Update for Windows 2000. During the IPSec negotiation process, IPSec NAT-T-capable peers automatically determine whether both the initiating IPSec peer (typically a client computer) and responding IPSec peer (typically a server) can perform IPSec NAT-T. In addition, IPSec NAT-T-capable peers automatically determine if there are any NATs in the path between them. If both of these conditions are true, the peers automatically use IPSec NAT-T to send IPSec-protected traffic.

**User Datagram Protocol-Encapsulating Security Payload (UDP-ESP)**
IPSec NAT-T provides UDP encapsulation of IPSec packets to enable IKE and ESP-protected traffic to pass through a NAT. IKE automatically detects that a NAT is present and uses UDP-ESP encapsulation to enable ESP-protected IPSec traffic to pass through the NAT.

**Related Information**
The following resources contain additional information that is relevant to this section.

- RFC 2637, 1701, 1702, 2661, 2865, 2866, 1213, 2284, 2716, 2138, and 2139 in the IETF RFC Database.

# VPN Tools and Settings

Updated: August 21, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN Tools and Settings**

**In this section**

Windows Server 2003 provides the following tools to troubleshoot VPN connections:

- TCP/IP Troubleshooting Tools
- Authentication and Accounting Logging
- Event Logging
- IAS Event Logging
- PPP Logging
- Tracing
- Oakley Logging
- Network Monitor

**TCP/IP Troubleshooting Tools**

The **Ping**, **Tracert**, and **Pathping** tools use ICMP Echo and Echo Reply messages to verify connectivity, display the path to a destination, and test path integrity. The **route print** command can be used to display the IP routing table. Alternately, on the VPN server, you can use the **netsh routing ip show rtmroutes** command or the Routing and Remote Access snap-in. The Nslookup tool can be used to troubleshoot DNS and name resolution issues.

**Authentication and Accounting Logging**

A VPN server running Windows Server 2003 supports the logging of authentication and accounting information for remote access VPN connections in local logging files when Windows authentication or Windows accounting is enabled. This logging is separate from the events recorded in the system event log. You can use the information that is logged to track remote access usage and authentication attempts. Authentication and accounting logging is especially useful for troubleshooting remote access policy issues. For each authentication attempt, the name of the remote access policy that either accepted or rejected the connection attempt is recorded.

Enable authentication and accounting logging from the **Settings** tab on the properties of the **Local File** object in the **Remote Access Logging** folder in the Routing and Remote Access snap-in (if the Routing and Remote Access service is configured for Windows authentication and accounting) or the Internet Authentication Service snap-in (if the Routing and Remote Access service is configured for RADIUS authentication and accounting and the RADIUS server is an IAS server)

The authentication and accounting information is stored in a configurable log file or files stored in the *SystemRoot*\System32\LogFiles folder. The log files are saved in Internet Authentication Service (IAS) or database-compatible format, meaning that any database program can read the log file directly for analysis.

If the VPN server is configured for RADIUS authentication and accounting and the RADIUS server is a computer running Windows Server 2003 and IAS, the authentication and accounting logs are stored in the *SystemRoot*\System32\LogFiles folder on the IAS server computer.

IAS for Windows Server 2003 can also send authentication and accounting information to a Structured Query Language (SQL) database.

**Event Logging**

On the **Logging** tab in the properties of a VPN server in the Routing and Remote Access snap-in, there are four levels of logging. Select **Log all events**, and then try the connection again. After the connection fails, check the system event log for events logged during the connection process. After you are done viewing remote access events, select the **Log errors and**

**warnings** option on the **Logging** tab to conserve system resources.

#### IAS Event Logging

If your VPN servers are configured for RADIUS authentication and your RADIUS servers are computers running Windows Server 2003 and IAS, check the system event log for IAS events for rejected or accepted connection attempts. IAS system event log entries contain a lot of information about the connection attempt including the name of the remote access policy that accepted or rejected the connection attempt. IAS event logging for rejected or accepted connection attempts is enabled by default and configured from the **Service** tab from the properties of an IAS server in the Internet Authentication Service snap-in.

#### PPP logging

PPP logging records the series of programming functions and PPP control messages during a PPP connection and is a valuable source of information when you are troubleshooting the failure of a PPP connection. To enable PPP logging, select the **Log additional Routing and Remote Access information** option on the **Logging** tab on the properties of a remote access server.

By default, the PPP log is stored as the Ppp.log file in the *SystemRoot*\Tracing folder.

#### Tracing

The Windows Server 2003 Routing and Remote Access service has an extensive tracing capability that you can use to troubleshoot complex network problems. You can enable the components in Windows Server 2003 to log tracing information to files using the Netsh command or through the registry.

#### Enabling Tracing with Netsh

You can use the Netsh command to enable and disable tracing for specific components or for all components.

**To enable and disable tracing for a specific component, use the following syntax:**

**netsh ras set tracing** *Component* **enabled|disabled**

where *Component* is a component in the list of Routing and Remote Access service components found in the Windows Server 2003 registry under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing. For example, to enable tracing for the RASAUTH component, the command is:

**netsh ras set tracing rasauth enabled**

To enable tracing for all components, use the following command:

**netsh ras set tracing * enabled**

#### Enabling Tracing Through the Registry

**You can configure the tracing function by changing settings in the Windows Server 2003 registry under:**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing

You can enable tracing for each Routing and Remote Access service component by setting the registry values described later. You can enable and disable tracing for components while the Routing and Remote Access service is running. Each component is capable of tracing and appears as a subkey under the preceding registry key.

To enable tracing for each component, you can configure the following registry entries for each protocol key:

**EnableFileTracing** REG_DWORD *Flag*

You can enable logging tracing information to a file by setting **EnableFileTracing** to 1. The default value is 0.

**FileDirectory** REG_EXPAND_*SZ Path*

You can change the default location of the tracing files by setting **FileDirectory** to the path you want. The file name for the log file is the name of the component for which tracing is enabled. By default, log files are placed in the *SystemRoot*\Tracing folder.

**FileTracingMask** REG_DWORD *LevelOfTracingInformationLogged*

**FileTracingMask** determines how much tracing information is logged to the file. The default value is 0xFFFF0000.

**MaxFileSize** REG_DWORD *SizeOfLogFile*

You can change the size of the log file by setting different values for **MaxFileSize**. The default value is 0x10000 (64K).

**Note**

Tracing consumes system resources and should be used sparingly to help identify network problems. After the trace is captured or the problem is identified, you should immediately disable tracing. Do not leave tracing enabled on multiprocessor computers.

**Note**

Tracing information can be complex and very detailed. Most of the time this information is useful only to Microsoft support professionals or to network administrators who are very experienced with the Routing and Remote Access service. Tracing information can be saved as files and sent to Microsoft support for analysis.

**Oakley Logging**

You can use the Oakley log to view details about the SA establishment process. The Oakley log is enabled in the registry. It is not enabled by default. To enable the Oakley log, set the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLogging registry setting to 1. The Oakley key does not exist by default and must be created.

After it is enabled, the Oakley log, which is stored in the *SystemRoot*\Debug folder, records all IPSec SA negotiations. A new Oakley.log file is created each time the IPSec Policy Agent is started and the previous version of the Oakley.log file is saved as Oakley.log.sav.

To activate the new EnableLogging registry setting after modifying its value, stop and start the IPSec Policy Agent and related IPSec services by running the following sequence of commands:

1. Stop the Routing and Remote Access service using the **net stop remoteaccess** command.
2. Stop the IPSec services using the **net stop policyagent** command.
3. Start the IPSec services using the **net start policyagent** command.
4. Start the Routing and Remote Access service using the **net start remoteaccess** command.

**Network Monitor**

Use Network Monitor, a packet capture and analysis tool supplied with Windows Server 2003, to capture and view the traffic sent between a VPN server and VPN client during the VPN connection process and during data transfer. You cannot interpret the encrypted portions of VPN traffic with Network Monitor. Network Monitor is installed as an optional networking component.

The proper interpretation of the remote access and VPN traffic with Network Monitor requires an in-depth understanding of PPP, PPTP, IPSec, and other protocols. Network Monitor captures can be saved as files and sent to Microsoft support for analysis.

# 802.11 Wireless Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**802.11 Wireless Technical Reference**
The wireless local area network (WLAN) protocol, 802.11, and associated technologies, such as the 802.1X protocol, are Institute of Electrical and Electronic Engineers (IEEE) standards supported by Microsoft to provide secure wireless networking and extend your core network infrastructure.

The IEEE 802.11 wireless standard defines the specifications for the physical layer and the media access control (MAC) sublayer for wireless communications. The 802.1X standard defines port-based, network access control used to provide authenticated network access for Ethernet networks. Although this standard was originally designed for wired Ethernet networks, it has been adapted for use on 802.11 wireless LANs.

**In this subject**


- What Is 802.11 Wireless?
- How 802.11 Wireless Works
- 802.11 Wireless Tools and Settings

# What Is 802.11 Wireless?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is 802.11 Wireless?**
The wireless local area network (WLAN) protocol, IEEE 802.11, and associated technologies, such as the 802.1X protocol and Wi-Fi Protected Access (WPA), allow secure high-speed wireless network access and mobile access to a network infrastructure. Until the recent development and wide adoption of IEEE 802.11b, also known as *Wi-Fi*, in order to obtain high-speed network access to your local area network (LAN) your network client needed to be physically connected to the LAN with some type of wiring.

**Benefits**
Some of the specific benefits of implementing a WLAN are included in the following:

**Wireless connections can extend or replace a wired infrastructure**
In situations where it is costly, inconvenient, or impossible to lay cables, wireless connections can extend or replace a wired infrastructure. This benefit includes the following:

- To connect the networks in two buildings separated by a physical, legal, or financial obstacle, you can either use a link provided by a telecommunications vendor (for a fixed installation cost and ongoing recurring costs) or you can create a point-to-point wireless link using wireless LAN technology (for a fixed installation cost, but no recurring costs). Eliminating recurring telecommunications charges can provide significant cost savings.

- Wireless LAN technologies can be used to create a temporary network, which is in place for only a specific amount of time. For example, the network needed at a convention or trade show can be a wireless network, rather than deploying the physical cabling required for a traditional Ethernet network.

- Some types of buildings, such as historical buildings, might be governed by building codes that prohibit the use of wiring, making wireless networking an important alternative.

- The wiring-free aspect of wireless LAN networking is also very attractive to homeowners who want to connect the various computers in their home together without having to drill holes and pull network cables through walls and ceilings.

**Increased productivity for the mobile employee**
This benefit includes the following:

- The mobile user whose primary computer is a portable computer can change location and always remain connected to the network. This enables the mobile user to travel to various places — meeting rooms, hallways, lobbies, cafeterias, classrooms, and so forth — yet retain access to networked data. Without wireless access, the user has to carry cabling and is restricted to working near a network jack.

- Wireless LAN networking is a perfect technology for environments where movement is required. For example, retail environments can benefit when employees use a wireless portable computer or handheld device to enter inventory information directly into the store database from the sales floor.

- Even if no wireless infrastructure is present, wireless portable computers can still form their own ad hoc networks to communicate and share data with each other.

**Easy access to the Internet in public places**
Beyond the corporate workplace, access to the Internet and even corporate sites can be made available through public wireless "hot spot" networks. For example, many airports, conference centers, and hotels provide wireless access to the Internet for their visitors. When a traveling worker reaches his or her destination, perhaps meeting a client at a corporate office, limited access can be provided to the traveling worker through the local wireless network. The network can recognize the traveling worker from another corporation and create a connection that provides Internet access to the visiting user, but is isolated from the local corporate network.

### Note

- In all these wireless LAN scenarios, it is worth noting that today's standards-based wireless LANs operate at the same speeds that were considered state of the art for wired networks just a few years ago. For example, IEEE 802.11b, a prevalent wireless LAN technology, operates at a maximum of 11 megabits per second (Mbps), or about 30 to 100 times faster than standard dial-up technologies. This data transfer rate is adequate for running a number of applications or services. In addition, ongoing advancements of wireless standards continue to increase the data transfer rate, with speeds of up to 54 Mbps for IEEE 802.11g.

#### Dependencies and Associated Technologies

In a enterprise environment, secure 802.11 wireless technology has specific dependencies and other technologies with which it is typically associated.

#### Dependencies for a secure WLAN

The dependencies of a secure WLAN include the following:

- A standard TCP/IP network environment.
- Active Directory, to store account properties and validate password-based credentials.
- Remote Authentication Dial-In User Service (RADIUS) support, to provide centralized connection authentication, authorization, and accounting.
- A certificate infrastructure, also known as a public key infrastructure (PKI), to issue and validate the certificates required for Extensible Authentication Protocol–Transport Layer Security (EAP-TLS) and Protected EAP (PEAP)–TLS authentication. TLS can use either computer certificates for authenticating the wireless client computer or smart cards or registry-based user certificates for authenticating the user of a wireless client computer.
- For PEAP-Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) authentication, computer certificates for the RADIUS servers and root CA certificates of the issuing CAs on the wireless clients (if needed).

#### Associated Technologies

The associated technologies of a secure WLAN include the following:

- WPA, as agreed upon by wireless vendors, is an interoperable interim standard until the IEEE 802.11i standard is being ratified. WPA requires secure wireless networking, addresses the issues with WEP through a software upgrade, provides a secure wireless networking solution for small office/home office (SOHO) wireless users, and provides compatibility with the upcoming IEEE 802.11i standard.
- Dynamic Host Configuration Protocol (DHCP) services, to provide automatic Internet Protocol (IP) configuration to wireless clients.
- Domain Name System (DNS) and Windows Internet Name Service (WINS) services, to provide name resolution.

### Note

- For information about IEEE 802.11b wireless networks for the home and small business, see the Microsoft Wi-Fi Web page.

#### 802.11 Wireless Network Infrastructure

The following figure shows a simple example of a WLAN using a wireless access point and some of the Windows Server 2003 technologies on which a secure WLAN is dependent.

### Example Wireless Network Infrastructure


Access
Point
Wireless Client
(Windows XP)
RADIUS
(IAS)
Active Directory

# How 802.11 Wireless Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How 802.11 Wireless Works**
The IEEE 802.11 protocol is a network access technology for providing connectivity between wireless stations and wired networking infrastructures.

By deploying the IEEE 802.11 protocol and associated technologies, you enable the mobile user to travel to various places — meeting rooms, hallways, lobbies, cafeterias, classrooms, and so forth — and still have access to networked data. Also, beyond the corporate workplace, you enable access to the Internet and even corporate sites can be made available through public wireless "hot spot" networks. Airports, restaurants, rail stations, and common areas throughout cities can be configured to provide this service.

This section provides an in-depth view of how IEEE 802.11 works, including the architecture, related protocols, and technologies.

**802.11 Architecture**
The 802.11 logical architecture contains several main components: station (STA), wireless access point (AP), independent basic service set (IBSS), basic service set (BSS), distribution system (DS), and extended service set (ESS). Some of the components of the 802.11 logical architecture map directly to hardware devices, such as STAs and wireless APs. The wireless STA contains an adapter card, PC Card, or an embedded device to provide wireless connectivity. The wireless AP functions as a bridge between the wireless STAs and the existing network backbone for network access.

An IBSS is a wireless network, consisting of at least two STAs, used where no access to a DS is available. An IBSS is also sometimes referred to as an ad hoc wireless network.

A BSS is a wireless network, consisting of a single wireless AP supporting one or multiple wireless clients. A BSS is also sometimes referred to as an infrastructure wireless network. All STAs in a BSS communicate through the AP. The AP provides connectivity to the wired LAN and provides bridging functionality when one STA initiates communication to another STA or a node on the DS.

An ESS is a set of two or more wireless APs connected to the same wired network that defines a single logical network segment bounded by a router (also known as a *subnet*).

The APs of multiple BSSs are interconnected by the DS. This allows for mobility, because STAs can move from one BSS to another BSS. APs can be interconnected with or without wires; however, most of the time they are connected with wires. The DS is the logical component used to interconnect BSSs. The DS provides distribution services to allow for the roaming of STAs between BSSs.

The following figure shows the 802.11 architecture.

**802.11 Architecture**




**802.11 Operating Modes**
IEEE 802.11 defines the following operating modes:

- Infrastructure mode

- Ad hoc mode

In both operating modes, a Service Set Identifier (SSID), also known as the *wireless network name*, identifies the wireless network. The *SSID* is a name configured on the wireless AP (for infrastructure mode) or an initial wireless client (for ad hoc mode) that identifies the wireless network. The SSID is periodically advertised by the wireless AP or the initial wireless client using a special 802.11 MAC management frame known as a *beacon frame*.

**802.11 Infrastructure Mode**

In *infrastructure mode*, there is at least one wireless AP and one wireless client. The wireless client uses the wireless AP to access the resources of a traditional wired network. The wired network can be an organization intranet or the Internet, depending on the placement of the wireless AP. An extended service set (ESS) is shown in the following figure.

**802.11 Infrastructure Mode**




**802.11 Ad Hoc Mode**

In *ad hoc mode*, wireless clients communicate directly with each other without the use of a wireless AP, as shown in the following figure.

**802.11 Wireless Clients in Ad Hoc Mode**




Ad hoc mode is also called *peer-to-peer mode*. Wireless clients in ad hoc mode form an independent basic service set (IBSS). One of the wireless clients, the first wireless client in the IBSS, takes over some of the responsibilities of the wireless AP. These responsibilities include the periodic beaconing process and the authentication of new members. This wireless client does not act as a bridge to relay information between wireless clients. Ad hoc mode is used to connect wireless clients together when there is no wireless AP present. The wireless clients must be explicitly configured to use ad hoc mode. There can be a maximum of nine members in an ad hoc 802.11 wireless network.

**802.11 Protocols and Technologies**

The 802.11-related protocols and technologies are discussed in detail in the following section:

- **802.11.** The IEEE 802.11 wireless standard defines the specifications for the physical layer and the media access control (MAC) layer.

- **802.1X.** The IEEE 802.1X standard defines port-based, network access control used to provide authenticated network access for Ethernet networks.

- **Extensible Authentication Protocol (EAP) over LAN (EAPOL).** EAP is a Point-to-Point Protocol (PPP)-based authentication mechanism that was adapted for use on point-to-point local area network (LAN) segments.

- **Wired Equivalent Privacy (WEP).** WEP provides data confidentiality services by encrypting the data sent between wireless nodes.

- **Wi-Fi Protected Access (WPA).** WPA is an interim standard until the IEEE 802.11i standard is ratified. These standards, intended to be a replacement for the WEP standard, offer more robust methods of data encryption and network authentication.

- **Wireless Auto Configuration.** The Wireless Auto Configuration feature of Windows XP and Windows Server 2003 dynamically selects the wireless network to which a connection is attempted, based either on configured preferences or default settings.

**802.11 Protocol**

The IEEE 802 standards committee defines two separate layers, the Logical Link Control (LLC) and media access control, for the Data-Link layer of the OSI model. The IEEE 802.11 wireless standard defines the specifications for the physical layer and the media access control (MAC) layer that communicates up to the LLC layer, as shown in the following figure.

**802.11 and OSI Model**




All of the components in the 802.11 architecture fall into either the media access control (MAC) sublayer of the data-link layer or the physical layer.

**802.11 MAC Frame**

The 802.11 MAC frame, as shown in the following figure, consists of a MAC header, the frame body, and a frame check sequence (FCS). The numbers in the following figure represent the number of bytes for each field.

**802.11 MAC Frame Format**




**Frame Control Field**

The Frame Control field, shown in the following figure, contains control information used for defining the type of 802.11 MAC frame and providing information necessary for the following fields to understand how to process the MAC frame. The numbers in the following figure represent the number of bits for each field.

**Frame Control Field**




A description of each Frame Control field subfield are as follows:

- **Protocol Version** provides the current version of the 802.11 protocol used. Receiving STAs use this value to determine if the version of the protocol of the received frame is supported.

- **Type and Subtype** determines the function of the frame. There are three different frame type fields: control, data, and management. There are multiple subtype fields for each frame type . Each subtype determines the specific function to perform for its associated frame type.

- **To DS and From DS** indicates whether the frame is going to or exiting from the DS (distributed system), and is only used in data type frames of STAs associated with an AP.

- **More Fragments** indicates whether more fragments of the frame, either data or management type, are to follow.

- **Retry** indicates whether or not the frame, for either data or management frame types, is being retransmitted.

- **Power Management** indicates whether the sending STA is in active mode or power-save mode.

- **More Data** indicates to a STA in power-save mode that the AP has more frames to send. It is also used for APs to indicate that additional broadcast/multicast frames are to follow.

- **WEP** indicates whether or not encryption and authentication are used in the frame. It can be set for all data frames and management frames, which have the subtype set to authentication.

- **Order** indicates that all received data frames must be processed in order.

**Duration/ID Field**
This field is used for all control type frames, except with the subtype of Power Save (PS) Poll, to indicate the remaining duration needed to receive the next frame transmission. When the sub-type is PS Poll, the field contains the association identity (AID) of the transmitting STA.

**Address Fields**
Depending upon the frame type, the four address fields will contain a combination of the following address types:

- **BSS Identifier (BSSID).** BSSID uniquely identifies each BSS. When the frame is from an STA in an infrastructure BSS, the BSSID is the MAC address of the AP. When the frame is from a STA in an IBSS, the BSSID is the randomly generated, locally administered MAC address of the STA that initiated the IBSS.

- **Destination Address (DA).** DA indicates the MAC address of the final destination to receive the frame.

- **Source Address (SA).** SA indicates the MAC address of the original source that initially created and transmitted the frame.

- **Receiver Address (RA).** RA indicates the MAC address of the next immediate STA on the wireless medium to receive the frame.

- **Transmitter Address (TA).** TA indicates the MAC address of the STA that transmitted the frame onto the wireless medium.

For more information about the address types and the contents of the address fields in the 802.11 MAC header, see the IEEE 802.11 standard at the IEEE Web site.

**Sequence Control**
The Sequence Control field contains two subfields, the Fragment Number field and the Sequence Number field, as shown in the following figure.

Sequence Control Field




A description of each Sequence Control field subfield are as follows:

- **Sequence Number** indicates the sequence number of each frame. The sequence number is the same for each frame sent for a fragmented frame; otherwise, the number is incremented by one until reaching 4095, when it then begins at zero again.

- **Fragment Number** indicates the number of each frame sent of a fragmented frame. The initial value is set to 0 and then incremented by one for each subsequent frame sent of the fragmented frame.

**Frame Body**
The frame body contains the data or information included in either management type or data type frames.

**Frame Check Sequence**
The transmitting STA uses a cyclic redundancy check (CRC) over all the fields of the MAC header and the frame body field to generate the FCS value. The receiving STA then uses the same CRC calculation to determine its own value of the FCS field to verify whether or not any errors occurred in the frame during the transmission.

**802.11 PHY Sublayer**
At the physical (PHY) sublayer, IEEE 802.11 defines a series of encoding and transmission schemes for wireless communications, the most common of which are the Frequency Hopping Spread Spectrum (FHSS), Direct Sequence Spread Spectrum (DSSS), and Orthogonal Frequency Division Multiplexing (OFDM) transmission schemes. The following figure shows the 802.11, 802.11b, 802.11a, and 802.11g standards that exist at the PHY sublayer. These standards are described in the sections that follow.

### Standards for 802.11 at the PHY Layer

| | 802.2 Logical Link Control (LLC) | | | |
|---|---|---|---|---|
| MAC | CSMA/CA | | | |
| PHY | **802.11**<br>2 Mbps<br>S-Band ISM<br>FHSS | **802.11b**<br>11 Mbps<br>S-Band ISM<br>DSSS | **802.11a**<br>54 Mbps<br>C-Band ISM<br>OFDM | **802.11g**<br>54 Mbps<br>S-Band ISM<br>OFDM |

**IEEE 802.11**
The bit rate for the original IEEE 802.11 standard is 2 Mbps using the FHSS transmission scheme and the S-Band Industrial, Scientific, and Medical (ISM) frequency band, which operates in the frequency range of 2.4 to 2.5 GHz. However, under less-than-ideal conditions, a lower bit rate speed of 1 Mbps is used.

**802.11b**
The major enhancement to IEEE 802.11 by IEEE 802.11b is the standardization of the physical layer to support higher bit rates. IEEE 802.11b supports two additional speeds, 5.5 Mbps and 11 Mbps, using the S-Band ISM. The DSSS transmission scheme is used in order to provide the higher bit rates. The bit rate of 11 Mbps is achievable in ideal conditions. In less-than-ideal conditions, the slower speeds of 5.5 Mbps, 2 Mbps, and 1 Mbps are used.

### Note

- 802.11b uses the same frequency band as that used by microwave ovens, cordless phones, baby monitors, wireless video cameras, and Bluetooth devices.

**802.11a**
IEEE 802.11a (the first standard to be ratified, but just now being widely sold and deployed) operates at a bit rate as high as 54 Mbps and uses the C-Band ISM, which operates in the frequency range of 5.725 to 5.875 GHz. Instead of DSSS, 802.11a uses OFDM, which allows data to be transmitted by subfrequencies in parallel and provides greater resistance to interference and greater throughput. This higher-speed technology enables wireless LAN networking to perform better for video and conferencing applications.

Because they are not on the same frequencies as other S-Band devices (such as cordless phones), OFDM and IEEE 802.11a provide both a higher data rate and a cleaner signal. The bit rate of 54 Mbps is achievable in ideal conditions. In less-than-ideal conditions, the slower speeds of 48 Mbps, 36 Mbps, 24 Mbps, 18 Mbps, 12 Mbps, and 6 Mbps are used.

**802.11g**
IEEE 802.11g operates at a bit rate as high as 54 Mbps, but uses the S-Band ISM and OFDM. 802.11g is also backward-compatible with 802.11b and can operate at the 802.11b bit rates and use DSSS. 802.11g wireless network adapters can connect to an 802.11b wireless AP, and 802.11b wireless network adapters can connect to an 802.11g wireless AP. Thus, 802.11g provides a migration path for 802.11b networks to a frequency-compatible standard technology with a higher bit rate. Existing 802.11b wireless network adapters cannot be upgraded to 802.11g by updating the firmware of the adapter — they must be replaced. Unlike migrating from 802.11b to 802.11a (in which all the network adapters in both the wireless clients and the wireless APs must be replaced at the same time), migrating from 802.11b to 802.11g can be done incrementally.

Like 802.11a, 802.11g uses 54 Mbps in ideal conditions and the slower speeds of 48 Mbps, 36 Mbps, 24 Mbps, 18 Mbps, 12 Mbps, and 6 Mbps in less-than-ideal conditions.

**802.1X Protocol**
The IEEE 802.1X standard defines port-based, network access control used to provide authenticated network access for Ethernet networks. This port-based network access control uses the physical characteristics of the switched LAN infrastructure to authenticate devices attached to a LAN port. Access to the port can be denied if the authentication process fails. Although this standard was designed for wired Ethernet networks, it has been adapted to 802.11 wireless LANs.

**Components of 802.1X**

IEEE 802.1X defines the following terms, as described in the following sections:

- **Port access entity.** A LAN port, also known as *port access entity (PAE)*, is the logical entity that supports the IEEE 802.1X protocol that is associated with a port. A PAE can adopt the role of the authenticator, the supplicant, or both.

- **Authenticator.** An *authenticator* is a LAN port that enforces authentication before allowing access to services accessible using that port. For wireless connections, the authenticator is the logical LAN port on a wireless AP through which wireless clients in infrastructure mode gain access to other wireless clients and the wired network.

- **Supplicant.** The *supplicant* is a LAN port that requests access to services accessible on the authenticator. For wireless connections, the supplicant is the logical LAN port on a wireless LAN network adapter that requests access to the other wireless clients and the wired network by associating with and then authenticating itself to an authenticator.

  Whether for wireless connections or wired Ethernet connections, the supplicant and authenticator are connected by a logical or physical point-to-point LAN segment.

- **Authentication server.** To verify the credentials of the supplicant, the authenticator uses an *authentication server*, which checks the credentials of the supplicant on behalf of the authenticator and then responds to the authenticator, indicating whether or not the supplicant is authorized to access the authenticator's services.

The following figure shows these components for a wireless LAN network.

**Components of IEEE 802.1X Authentication**




The authentication server can be the following:

- **A component of the access point.** In this case, the AP must be configured with the sets of user credentials corresponding to the supplicants that will be attempting to connect (it is typically not implemented for wireless APs).

- **A separate entity.** In this case, the AP forwards the credentials of the connection attempt to a separate authentication server. Typically, a wireless AP uses the Remote Authentication Dial-In User Service (RADIUS) protocol to send a connection request message to a RADIUS server.

#### Controlled and Uncontrolled Ports

The authenticator's port-based access control defines the following different types of logical ports that access the wired LAN by means of a single physical LAN port:

- **Uncontrolled Port.** The uncontrolled portallows an uncontrolled exchange between the authenticator (the wireless AP) and other networking devices on the wired network — regardless of any wireless client's authorization state. Frames sent by the wireless client are never sent using the uncontrolled port.

- **Controlled Port.** The controlled portallows data to be sent between a wireless client and the wired network only if the wireless client is authorized by 802.1X. Before authentication, the switch is open and no frames are forwarded between the wireless client and the wired network. When the wireless client is successfully authenticated using IEEE 802.1X, the switch is closed, and frames can be sent between the wireless client and nodes on the wired network.

The different types of ports are shown in the following figure.

**Controlled and Uncontrolled Ports for IEEE 802.1X**




On an authenticating Ethernet switch, the wired Ethernet client can send Ethernet frames to the wired network as soon as authentication is complete. The switch identifies the traffic of a specific wired Ethernet client using the physical port to which the Ethernet client is connected. Typically, only a single Ethernet client is connected to a physical port on the Ethernet switch.

Because multiple wireless clients contend for access to the same frequency channel and send data using the same channel, an extension to the basic IEEE 802.1X protocol is required to allow a wireless AP to identify the secured traffic of a particular wireless client. The wireless client and wireless AP do this through the mutual determination of a per-client unicast session key. Only authenticated wireless clients have knowledge of their per-client unicast session key. Without a valid unicast session key tied to a successful authentication, a wireless AP discards the traffic sent from the wireless client.

**EAP over LAN**

To provide a standard authentication mechanism for IEEE 802.1X, the Extensible Authentication Protocol (EAP) was chosen. EAP is a Point-to-Point Protocol (PPP)-based authentication mechanism that was adapted for use on point-to-point LAN segments. EAP messages are normally sent as the payload of PPP frames. To adapt EAP messages to be sent over Ethernet or wireless LAN segments, the IEEE 802.1X standard defines EAP over LAN (EAPOL), a standard encapsulation method for EAP messages.

**WEP**

WEP provides data confidentiality services by encrypting the data sent between wireless nodes. Setting a WEP flag in the MAC header of the 802.11 frame indicates that the frame is encrypted with WEP encryption. WEP provides data integrity by including an integrity check value (ICV) in the encrypted portion of the wireless frame.

WEP defines two shared keys:

- **Multicast/global key.** The *multicast/global key* is an encryption key that protects multicast and broadcast traffic from a wireless AP to all of its connected wireless clients.

- **Unicast session key.** The *unicast session key* is an encryption key that protects unicast traffic between a wireless client and a wireless AP and multicast and broadcast traffic sent by the wireless client to the wireless AP.

WEP encryption uses the RC4 symmetric stream cipher with 40-bit and 104-bit encryption keys. Although 104-bit encryption keys are not specified in the 802.11 standard, many wireless AP vendors support them.

**Note**

- Some implementations that advertise the use of 128-bit WEP encryption keys are just adding a 104-bit encryption key to the 24-bit initialization vector (IV) and calling it a 128-bit key. The IV is a field in the header of each 802.11 frame that is used during the encryption and decryption process.

**WEP Encryption**

The WEP encryption process is shown in the following figure.

**WEP Encryption Process**



To encrypt the payload of an 802.11 frame, the following process is used:

1. A 32-bit integrity check value (ICV) is calculated for the frame data.

2. The ICV is appended to the end of the frame data.

3. A 24-bit initialization vector (IV) is generated and appended to the WEP encryption key.

4. The combination of initialization vector and WEP encryption key is used as the input of a pseudo-random number generator (PRNG) to generate a bit sequence that is the same size as the combination of data and ICV.

5. The PRNG bit sequence, also known as the *key stream*, is bit-wise exclusive ORed (XORed) with the combination of data and ICV to produce the encrypted portion of the payload that is sent between the wireless access point (AP) and the wireless client.

6. To create the payload for the wireless MAC frame, the IV is added to the front of the encrypted combination of the data and ICV, along with other fields.

**WEP Decryption**

The WEP decryption process is shown in the following figure.

**WEP Decryption Process**




To decrypt the 802.11 frame data, the following process is used:

1. The initialization vector (IV) is obtained from the front of the MAC payload.

2. The IV is appended to the WEP encryption key.

3. The combination of initialization vector and WEP encryption key is used as the input of the same PRNG to generate a bit sequence of the same size as the combination of the data and the ICV. This process produces the same key stream as that of the sending wireless node.

4. The PRNG bit sequence is XORed with the encrypted combination of the data and ICV] to decrypt the combined data and ICV portion of the payload.

5. The ICV calculation for the data portion of the payload is run, and its result is compared with the value included in the incoming frame. If the values match, the data is considered to be valid (sent from the wireless client and unmodified in transit). If they do not match, the frame is silently discarded.

**Security Issues with WEP and IEEE 802.11**

The main problem with WEP is that the determination and distribution of WEP encryption keys are not defined. WEP keys must be distributed by using a secure channel outside of the 802.11 protocol. In practice, WEP keys are text strings that must be manually configured using a keyboard for both the wireless AP and wireless clients. However, this key distribution system does not scale well to an enterprise organization and is not secure.

Additionally, there is no defined mechanism for changing the WEP encryption keys either per authentication or periodically for an authenticated connection. All wireless APs and clients use the same manually configured WEP key for multiple sessions. With multiple wireless clients sending a large amount of data, an attacker can remotely capture large amounts of WEP ciphertext and use cryptanalysis methods to determine the WEP key.

The lack of a WEP key management protocol is a principal limitation to providing 802.11 security, especially in infrastructure mode with a large number of stations. Some examples of this type of network include corporate and educational institutional campuses and public places such as airports and malls. The lack of automated authentication and key determination services also affects operation in ad hoc mode, in which users might want to use in peer-to-peer collaborative communication in areas such as conference rooms.

**WPA**

Although 802.1X addresses many of the security issues of the original 802.11 standard, issues still exist with regard to weaknesses in the WEP encryption and data integrity methods. The long-term solution to these problems is the IEEE 802.11i standard, which is currently in draft form.

Until the IEEE 802.11i standard is ratified, wireless vendors have agreed on an interoperable interim standard known as Wi-Fi Protected Access (WPA). The goals of WPA are the following:

- **To require secure wireless networking.** WPA requires secure wireless networking by requiring 802.1X authentication, encryption, and unicast and multicast/global encryption key management.

- **To address WEP issues with a software upgrade.** The implementation of the RC4 stream cipher within WEP is vulnerable to known plaintext attacks. Additionally, the data integrity provided with WEP is relatively weak. WPA solves all the remaining security issues with WEP, yet only requires firmware updates in wireless equipment and an update for wireless clients. Existing wireless equipment is not expected to require replacement.

- **To provide a secure wireless networking solution for small office/home office (SOHO) wireless users.** For the SOHO, there is no RADIUS server to provide 802.1X authentication with an EAP type. SOHO wireless clients must use either shared key authentication (highly discouraged) or open system authentication (recommended) with a single static WEP key for both unicast and multicast traffic. WPA provides a pre-shared key option intended for SOHO configurations. The pre-shared key is configured on the wireless AP and each wireless client. The initial unicast encryption key is derived from the authentication process, which verifies that both the wireless client and the wireless AP have the pre-shared key.

- **To be compatible with the upcoming IEEE 802.11i standard.** WPA is a subset of the security features in the proposed IEEE 802.11i standard. All the features of WPA are described in the current draft of the 802.11i standard.

- **To be available today.** WPA upgrades to wireless equipment and for wireless clients were available beginning in February 2003.

**WPA Security Features**

WPA contains enhancements or replacements for the following security features:

- Authentication

- Encryption

- Data integrity

**Authentication**

With 802.11, 802.1X authentication is optional; with WPA, 802.1X authentication is required. Authentication with WPA is a combination of open system and 802.1X authentication, which uses the following phases:

- The first phase uses open system authentication to indicate to the wireless client that it can send frames to the wireless AP.

- The second phase uses 802.1X to perform a user-level authentication. For environments without a RADIUS infrastructure, WPA supports the use of a pre-shared key; for environments with a RADIUS infrastructure, WPA supports EAP and RADIUS.

**Encryption**

With 802.1X, rekeying of unicast encryption keys is optional. Additionally, 802.11 and 802.1X provide no mechanism to change the global encryption key that is used for multicast and broadcast traffic. With WPA, rekeying of both unicast and global encryption keys is required. The Temporal Key Integrity Protocol (TKIP) changes the unicast encryption key for every frame, and

each change is synchronized between the wireless client and the wireless AP. For the multicast/global encryption key, WPA includes a facility for the wireless AP to advertise changes to the connected wireless clients.

**TKIP**
For 802.11, WEP encryption is optional. For WPA, encryption using TKIP is required. TKIP replaces WEP with a new encryption algorithm that is stronger than the WEP algorithm, yet can be performed using the calculation facilities present on existing wireless hardware. TKIP also provides for the following:

- The verification of the security configuration after the encryption keys are determined.
- The synchronized changing of the unicast encryption key for each frame.
- The determination of a unique starting unicast encryption key for each pre-shared key authentication.

**AES**
WPA defines the use of the Advanced Encryption Standard (AES) as an optional replacement for WEP encryption. Because adding AES support by using a firmware update might not be possible for existing wireless equipment, support for AES on wireless network adapters and wireless APs is not required.

**Data Integrity**
With 802.11 and WEP, data integrity is provided by a 32-bit ICV that is appended to the 802.11 payload and encrypted with WEP. Although the ICV is encrypted, it is possible through cryptanalysis to change bits in the encrypted payload and update the encrypted ICV without being detected by the receiver.

With WPA, a method known as *Michael* specifies a new algorithm that calculates an 8-byte message integrity code (MIC) with the calculation facilities available on existing wireless hardware. The MIC is placed between the data portion of the 802.11 frame and the 4-byte ICV. The MIC field is encrypted along with the frame data and the ICV.

Michael also provides replay protection through the use of a frame counter field in the 802.11 MAC header.

### Note

- WPA requires software changes to the following:
- Wireless APs
- Wireless network adapters
- Wireless client software

**Wireless Auto Configuration**
Wireless Auto Configuration dynamically selects the wireless network to which a connection is attempted, based either on configured preferences or default settings. This process includes automatically selecting and connecting to a more preferred wireless network when it becomes available. If none of the preferred wireless networks is found nearby, Wireless Auto Configuration configures the wireless adapter so that there is no accidental connection until the wireless client roams within the range of a preferred network.

Wireless Auto Configuration corresponds to the Wireless Configuration service in Windows Server 2003 and the Wireless Zero Configuration service in Windows XP. You can use the Services snap-in (available in the Administrative Tools folder) to view the current status of (as well as stop, start, and restart) either service. You can also manage either service from the command prompt by using the **net** command. For example, to stop either service, type **net stop wzcsv** at a command prompt.

Wireless Auto Configuration minimizes the configuration that is required to access wireless networks and allows you to travel to different wireless networks without reconfiguring the network connection settings on your computer for each location. For the initial scan of available wireless networks, Wireless Auto Configuration performs the following process:

1. Wireless Auto Configuration attempts to connect to the preferred networks that appear in the list of available networks in the preferred networks preference order.
2. If there are no successful connections, Wireless Auto Configuration attempts to connect to the preferred networks that do not appear in the list of available networks in the preferred networks preference order. Thus, it can connect even when the wireless APs are configured to suppress the beaconing of the SSID of the wireless network.

3. If there are no successful connections and there is an ad hoc network in the list of preferred networks that is available, Wireless Auto Configuration tries to connect to it.

4. If there are no successful connections, and there is an ad hoc network in the list of preferred networks that is not available, Wireless Auto Configuration configures the wireless network adapter to act as the first node in the ad hoc network.

5. If there are no successful connections to preferred networks, and there are no ad hoc networks in the list of preferred networks, Wireless Auto Configuration determines the **Automatically Connect To Non-Preferred Networks** setting (located on the **Wireless Networks** tab of the wireless network connection).

6. If the **Automatically Connect To Non-Preferred Networks** setting is disabled, Wireless Auto Configuration creates a random network name and places the wireless network adapter in infrastructure mode.

   This behavior prevents the Windows XP wireless client from accidentally connecting to a wireless network that does not appear in the list of preferred networks. You then see the **One Or More Wireless Networks Are Available** message in the notification area. The wireless adapter is not connected to any wireless network, but continues to scan for preferred wireless networks every 60 seconds.

7. If the **Automatically Connect To Non-Preferred Networks** setting is enabled, Wireless Auto Configuration attempts to connect to the available networks in the order in which the wireless adapter sensed them.

   If all connection attempts fail, Wireless Auto Configuration creates a random network name and places the wireless network adapter in infrastructure mode. You then see the **One Or More Wireless Networks Are Available** message in the notification area.

For subsequent scans, Wireless Auto Configuration determines whether there are any changes in the wireless environment that require switching the connection. If the Windows wireless client is already connected to a wireless network and there is no other preferred network higher in the preference list that has not been attempted already, Wireless Auto Configuration maintains the existing connection. If the Windows wireless client is already connected to a wireless network, but a more preferred wireless network becomes available, Wireless Auto Configuration disconnects from the currently connected wireless network and attempts to connect to the more preferred wireless network.

The operation of Wireless Auto Configuration provides the following:

- The first time a wireless adapter is added to a computer running Windows XP or Windows Server 2003 and a wireless network is available, Wireless Auto Configuration prompts you with the **One Or More Wireless Networks Are Available** message in the notification area, which leads you to select a wireless network in the **Connect To Wireless Network** dialog box.

  After you select a wireless network and the connection is successful, the selected network is automatically added as a preferred network, and you are no longer prompted whenever you are within range of it. For an organization, this is the typical process for configuring the initial connection to a private wireless network. After the initial configuration, Wireless Auto Configuration connects (and then maintains the connection) to the organization's wireless network.

  When you take your portable computer to your home wireless network, to an airport, or to another location with public wireless access, Wireless Auto Configuration first attempts to connect to your preferred network. When that connection attempt fails, you are prompted again to connect to your home wireless network or to the public access wireless network.

- If there are two preferred wireless networks, and the most preferred one is not initially available, Wireless Auto Configuration configures a wireless connection to the next most preferred network. When the most preferred network eventually becomes available, Wireless Auto Configuration automatically switches the wireless client connection to it after the next scan.

- If there are no preferred networks in the list of those available, Wireless Auto Configuration attempts to configure connections to the preferred networks in their configured order, in case the wireless APs for the wireless network are configured to prohibit the beaconing of their SSID.

**Related Information**

For more information about 802.11 and related technologies, see the Microsoft Wi-Fi Web site.

# 802.11 Wireless Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**802.11 Wireless Tools and Settings**

**In this section**


- Windows Wireless Client
- Wireless AP
- 802.11 Wireless Authentication Infrastructure
- 802.11 Wireless Registry Entries


This section contains information about the tools and settings available for configuring, managing, and customizing wireless clients, and for troubleshooting wireless connectivity.

**Windows Wireless Client**

The following services and features are associated with Windows wireless clients and wireless client connections.

**Network Connections Folder**

The Network Connections folder and the messages displayed in the notification area of the task bar provide information about the state of the wireless network availability as well as client authentication and connection status. If authentication of a client requires additional information from the user, such as selecting one of multiple user certificates, a text balloon appears to instruct the user. Within the Network Connections folder, the text under the name of the connection corresponding to the wireless network adapter describes the state of the authentication.

For wireless clients running Windows XP, when you click the wireless connection, the Details dialog box displays the authentication status, the Internet Protocol (IP) address configuration, and information about the connected wireless network and the current association. When you obtain status on the connection, you can view information such as the signal speed on the **General** tab and the IP address configuration on the **Support** tab.

If the wireless adapter is assigned an Automatic Private IP Addressing (APIPA) address in the range 169.254.0.0/16 or the configured alternate static IP address, the wireless client is still associated with the wireless AP, but either authentication has failed or the Dynamic Host Configuration Protocol (DHCP) server is not available. If the authentication fails and the association is still in place, the wireless adapter is enabled and Transmission Control Protocol/Internet Protocol (TCP/IP) performs its normal configuration process. If a DHCP server is not found (either authenticated or not), Windows XP automatically configures an APIPA or alternate static address.

For a wireless client running Windows 2000, you can use the **Network and Dial-up Connections** folder to view whether the connection corresponding to your wireless network adapter is authenticated. If authentication is successful, the connection icon appears normal. If authentication has failed, the connection icon has a red X through it. To view the IP address configuration of the connection corresponding to the wireless network adapter, type **ipconfig** at a command prompt.

**Netsh**

Windows XP, Windows Server 2003, and Windows 2000 have an extensive tracing capability that you can use to troubleshoot complex problems for specific components related to wireless client authentication. You can enable the components in Windows to log tracing information to files using the **netsh** command for specific components or for all components. For more detailed information, see “Netsh,” in the section “802.11 Wireless Authentication Infrastructure” later in this document.

**Network Monitor**

You can use Microsoft Network Monitor, available in Microsoft Systems Management Server, Windows 2000 Server, and Windows Server 2003, or a commercial packet analyzer (also known as a network sniffer), to capture and view the authentication and data traffic sent and received by a wireless network adapter.

For Windows wireless client authentication, you can use Network Monitor to capture the set of frames exchanged between the wireless client computer and the wireless AP during the wireless authentication process. You can then use Network Monitor to view the individual frames and determine why the authentication failed.

Network Monitor includes 802.1X, EAPOL, and EAP parsers. A *parser* is a component included with Network Monitor that can separate the fields of a protocol header and display their structure and values. Without a parser, Network Monitor displays the hexadecimal bytes of a header, which you must parse manually.

In Windows 2000 Server and Windows Server 2003, Network Monitor is installed as an optional management and monitoring tool using Control Panel's Add or Remove Programs. After it is installed, you can run Network Monitor from the Administrative Tools folder.

**Wireless Monitor snap-in**
In Windows Server 2003, you can use the new Wireless Monitor snap-in to view wireless AP or wireless client information. There are two main views of information in the Wireless Monitor snap-in:

- Access Point Information
- Wireless Client Information

When you click Access Point Information in the console tree, the wireless network adapter scans for the available wireless APs within range and then displays them in the details pane.

You can use the list of wireless APs to determine the visibility and parameters (such as signal strength, channel, and data rates) of specific wireless APs for a specific location.

When you click Wireless Client Information in the console tree, the list of wireless events for the Wireless Zero Configuration service and the EAPOL component displays in the details pane.

You can use these events to determine how the Wireless Zero Configuration service chooses to connect to a specific wireless AP and explore the details of the authentication process.

**Wireless AP**
The following services and features are associated with wireless access points (APs). Consult your wireless AP documentation for information about the set of troubleshooting tools provided with your wireless AP.

**Panel Indicators**
Most wireless APs have one or more indicators, status lights that are visible on the housing of the wireless AP, from which you can obtain a quick assessment of the wireless AP's hardware status. For example, you might see the following:

- An indicator to show that the wireless AP has electrical power.
- An indicator to show general operation status. For example, the indicator might show whether the wireless AP is associated with any wireless clients.
- An indicator to show wireless network traffic. This indicator might blink for each frame received on the wireless network.
- An indicator to show data collisions. If the blinking of this indicator seems excessive, evaluate the performance of the link by using methods suggested by the wireless AP vendor.
- An indicator to show wired network traffic. This indicator might blink for each frame received on the wired network.

Alternatively, the wireless AP might have a liquid crystal display (LCD) panel display that shows icons indicating its current status. Consult your wireless AP documentation for information about panel indicators and their interpretation.

**Site Survey Software**
Site survey software is used to determine the coverage volume and where the data rate changes for each wireless AP. An administrator can use site survey software during the deployment of wireless APs to determine their optimal placement. It is typically installed on a wireless-capable portable computer from a CD provided by the wireless AP vendor or the wireless network adapter vendor.

If wireless clients cannot connect to a specific wireless AP, use the site survey software to perform a site survey for that wireless AP. There might have been a change in the devices that create interference and objects that interfere with signal propagation because the original site survey and wireless AP placement was done.

**SNMP support**
Many wireless APs include a Simple Network Management Protocol (SNMP) agent that supports the following SNMP

Management Information Bases (MIBs):

- IEEE 802.11 MIB
- IEEE 802.1 Port Access Entity (PAE ) MIB
- SNMP MIB (described in RFC 1157)
- SNMP MIB II (described in RFC 1213)
- Ethernet Interface MIB (described in RFC 1398)
- SNMP Bridge MIB (described in RFC 1493)
- Remote Network Monitoring MIB (described in RFC 1757)
- RADIUS Authentication Client MIB (described in RFC 2618)

The SNMP agent on the wireless AP can be used in conjunction with your existing SNMP-based network management infrastructure to monitor your wireless APs, set trap conditions, and monitor loads on your wireless APs. In SNMP, trap conditions specify which events will trigger a message to be sent by the SNMP agent to a management system, indicating that the event has occurred on the host running the agent.

**Diagnostics**
The purpose of the diagnostics is to ensure that the wireless AP is operating properly (from a hardware standpoint) and to validate its current configuration. However, the exact diagnostic facilities of a wireless AP vary from one wireless AP to another.

Diagnostics for wireless APs can be of the following forms:

- Software or a Web site provided by the wireless AP vendor.
- A command-line tool or facility, such as terminal access to the wireless AP.

**802.11 Wireless Authentication Infrastructure**
Wireless 802.11 clients and APs can use Internet Authentication Service (IAS) — the Microsoft implementation of RADIUS — to enforce authentication in an enterprise environment. The following services and features are associated with authentication infrastructure of Wireless 802.11.

**IAS Event Logging**
To troubleshoot IAS authentication attempts, view events in the Windows event logs. Ensure that event logging is enabled for all types of IAS events (such as rejected, discarded, and successful authentication events). Event logging for all these types of events are enabled by default for both Windows Server 2003 IAS and Windows 2000 IAS.

IAS events are stored in the system event log, which can be viewed in the Event Viewer snap-in. Here is an example of a successful IAS authentication event (Event ID 1).

```
User client@example.com was granted access.
Fully-Qualified-User-Name =example.com/Users/Client
NAS-IP-Address =10.7.0.4
NAS-Identifier =<not present>
Client-Friendly-Name =Building 7 Wireless AP
Client-IP-Address =10.7.0.4
NAS-Port-Type =Wireless-IEEE 802.11
NAS-Port =6
Policy-Name =Wireless Remote Access Policy
Authentication-Type =EAP
EAP-Type =Smart Card or other Certificate
```

To view s failed IAS authentication event, view all events that have Event ID 2.

Viewing the IAS events in the system event log is one of the most useful troubleshooting tools for obtaining information about failed authentications. The IAS events are also helpful when troubleshooting remote access policies. When you have multiple

remote access policies configured, the Policy-Name field in the event description records the name of the remote access policy that either accepted or rejected the connection attempt.

**Network Monitor**
Network Monitor is useful for checking whether RADIUS messages are being exchanged, and for determining the RADIUS attributes of each message.

You can use Microsoft Network Monitor—available in Microsoft Systems Management Server, Windows 2000 Server, and Windows Server 2003 — or a commercial packet analyzer (also known as a *network sniffer*) to capture and view RADIUS authentication and accounting messages that are sent to and from the IAS RADIUS server or an IAS RADIUS proxy. Network Monitor includes a RADIUS parser that you can use to view the attributes of a RADIUS message and trouble-shoot connection issues.

**SChannel Logging**
Secure channel (SChannel) is a Security Support Provider (SSP) that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) authentication protocols. SSL and TLS are the Internet standard for secure transactions.

By default, SChannel logs only error messages in the system event log. To log errors, warnings, informational, and successful events, set the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\SecurityProviders\SCHANNEL\EventLogging registry value to 4 (as a DWORD type). With SChannel logging recording all events, it is possible to obtain more information about the certificate exchange and validation process on the IAS server during EAP-TLS authentication.

**Netsh**
Windows Server 2003 and Windows 2000 have an extensive tracing capability that creates tracing files that describe the internal behavior of Windows components during the wireless client authentication and authorization process. This information is typically most useful to Microsoft support engineers, who might request that you create trace files for a connection attempt during their investigation of a support issue. You can enable the components in Windows Server 2003 to log tracing information to files by using the **netsh** command for specific components or for all components. To enable and disable tracing for a specific component, use the following command:

```
netsh ras set tracing
component
 enabled|disabled
```

In the preceding command, *component* is any of the items in the list of components found in the registry under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing. For example, to enable tracing for the IASRAD component, use the following command:

```
netsh ras set tracing iasrad enabled
```

To obtain detailed information about the EAP authentication process for Windows XP or Windows Server 2003, enable tracing for only the EAP over LAN (EAPOL) and Remote Access Service Transport Level Security (RASTLS) components. To enable tracing for EAPOL, use the following command:

```
netsh ras set tracing eapol enabled
```

To enable tracing for RASTLS, use the following command:

```
netsh ras set tracing rastls enabled
```

After these commands are issued, try the authentication process again and view the Eapol.log and Rastls.log files in the *SystemRoot*\Tracing folder. To disable tracing for EAPOL and RASTLS, use the following respective commands:

```
netsh ras set tracing eapol disabled
netsh ras set tracing rastls disabled
```

To obtain detailed information about the EAP authentication process for Windows 2000, enable tracing for the RASTLS component.

Although you can enable tracing for individual components of IAS, it is generally easier to turn tracing on for all the IAS

components at once; and Microsoft support engineers typically want to see all the trace files, rather than the trace file for an individual component. To enable tracing for all components, use the following command:

```
netsh ras set tracing * enabled
```

To disable tracing for all components, use the following command:

```
netsh ras set tracing * disabled
```

The log files that are generated are stored in the *SystemRoot*\tracing folder.

**Note**

- Tracing consumes system resources and should be used sparingly during the investigation of a support issue. After the trace is done or the problem is identified, you should disable tracing. Do not leave tracing enabled on multiprocessor computers.

**SNMP Agent**

You can use the Simple Network Management Protocol (SNMP) agent software included with Windows 2000 Server and Windows Server 2003 to monitor status information for your IAS server from an SNMP snap-in. IAS supports the RADIUS Authentication Server MIB (RFC 2619) and the RADIUS Accounting Server MIB (RFC 2621). Use **AddRemove Programs** in Control Panel to install the SNMP agent. The SNMP agent can be used in conjunction with your existing SNMP-based network management infrastructure to monitor your IAS RADIUS servers or proxies.

**Performance Logs And Alerts Snap-in**

You can use the Performance Logs And Alerts snap-in to monitor counters, create logs, and set alerts for specific IAS components and program processes. You can also use charts and reports to determine how efficiently your server uses IAS and to both identify and troubleshoot potential problems.

You can use the Performance Logs And Alerts snap-in to monitor counters within the following IAS-related performance objects:

- IAS Accounting Clients
- IAS Accounting Proxy
- IAS Accounting Server
- IAS Authentication Clients
- IAS Authentication Proxy
- IAS Authentication Server
- IAS Remote Accounting Servers
- IAS Remote Authentication Servers

For more information about how to use the Performance Logs And Alerts snap-in, see Help and Support Center for Windows Server 2003.

**802.11 Wireless Registry Entries**

Controlling the computer and user authentication behavior of Windows XP and Windows Server 2003 can be configured by using either the AuthMode registry setting or in a Windows Server 2003 Active Directory environment by using the Wireless Network (IEEE 802.11) Policies Group Policy extension. Controlling the transmission behavior of the EAPOL-Start message when authenticating can be configured by using either the SupplicantMode registry setting or the Wireless Network (IEEE 802.11) Policies Group Policy extension.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**AuthMode Registry Entry**
The setting of the AuthMode registry entry controls the computer and user authentication behavior of Windows XP and Windows Server 2003.

**Registry path**
HKEY_LOCAL_MACHINE\Software\Microsoft\EAPOL\Parameters\General\Global\AuthMode

**Version**
Windows XP and Windows Server 2003

AuthMode has the following values:

- **0 - Computer authentication mode.** If computer authentication is successful, no user authentication is attempted. If the user logon is successful before computer authentication, user authentication is performed. This is the default setting for Windows XP (prior to Service Pack 1).

- **1 - Computer authentication with re-authentication.** If computer authentication is successful, a subsequent user logon results in a re-authentication with user credentials. The user logon has to complete in 60 seconds or the existing network connectivity is terminated. The user credentials are used for subsequent authentication or re-authentication. Computer authentication is not attempted again until the user logs off the computer. This is the default setting for Windows XP Service Pack 1 (SP1) and Windows Server 2003.

- **2 - Computer authentication only.** When a user logs on, it has no effect on the connection. Only computer authentication is performed. The exception to this behavior is when a user successfully logs on, and then roams between wireless APs. In that case, user authentication is performed. For changes to this setting to take effect, restart the Wireless Zero Configuration service for Windows XP or Windows Server 2003.

**SupplicantMode Registry Entry**
The setting of the SupplicantMode registry entry specifies the transmission behavior of the EAPOL-Start message when authenticating.

**Registry path**
HKEY_LOCAL_MACHINE\Software\Microsoft\EAPOL\Parameters\General\Global\SupplicantMode

**Version**
Windows XP and Windows Server 2003

SupplicantMode has the following values:

- **1 - Do not transmit.** Specifies that EAPOL-Start messages are not sent.

- **2 - Transmit.** Determines when to send EAPOL-Start messages and, if needed, sends an EAPOL-Start message.

- **3 - Transmit per 802.1x.** Sends an EAPOL-Start message upon association to initiate the 802.1X authentication process.

**Wireless 802.11 Group Policy Settings**
To automate the configuration of network settings on wireless clients running Windows XP SP1 and Windows Server 2003, Windows Server 2003 Active Directory domains support a new Wireless Network (IEEE 802.11) Policies Group Policy extension that allows you to configure wireless network settings that are part of Computer Configuration Group Policy. The Wireless Network (IEEE 802.11) Policies Group Policy extension is found in the Group Policy Object Editor by navigating to Default Domain Policy/Computer Confiuration/Windows Settings/Security Settings/Wireless.

Wireless network settings include the list of preferred networks, WEP settings, and IEEE 802.1X settings. These settings encompass all of the items on the Association and Authentication tabs of the properties of a wireless network and additional settings. The Association and Authentication tabs on Windows XP or Windows Server 2003 computers can be found by

opening **Network Connections**, right-clicking on the wireless connection icon and clicking **Properties**, then on the Wireless Network tab, click **Properties**. The settings are downloaded to wireless client computers running Windows XP SP1 or Windows Server 2003 that are members of a Windows Server 2003 Active Directory domain, making it much easier to deploy a specific configuration for secure wireless connections.

**Wireless Network (IEEE 802.11) Policy Settings**

In the Group Policy snap-in, you can configure wireless policies by using the Computer Configuration/Windows Settings/Security Settings/Wireless Network (IEEE 802.11) Policies node.

**Note**

- These policy settings do not apply to wireless clients running Windows XP (prior to SP1) or Microsoft 802.1X Authentication Client.

By default, there are no Group Policy objects (GPOs) in the Wireless Network (IEEE 802.11) Policies node. You can create only a single wireless network policy for each Group Policy object.

The properties of a wireless network policy consist of a General tab and a Preferred Networks tab.

On the General tab, you can view the options as described in the following table.

**General Settings of a Wireless Network Policy**

| Setting | Description |
|---|---|
| **Name** | This option allows you to type a friendly name for the wireless network policy. |
| **Description** | This option allows you to type a description for the wireless network policy. |
| **Check For Policy Changes Every** | This option allows you to specify the interval in minutes after which wireless clients that are Active Directory members check for changes in the wireless network policy. |
| **Networks To Access** | This option selects the types of wireless networks with which the wireless network is allowed to create connections:<br>• Any Available Network (Access Point Preferred)<br>• Access Point (Infrastructure) Networks Only<br>• Computer-To-Computer (Ad Hoc) Networks Only |
| **Use Windows To Configure Wireless Network Settings for Clients** | This option enables Wireless Auto Configuration. |
| **Automatically Connect to Non-preferred Networks** | This option enables automatic connections to wireless networks that are not configured as preferred networks. |

On the **Preferred Networks** tab, you can view the options as described in the following table.

**Preferred Network Options of a Wireless Network Policy**

| Option | Description |
|---|---|
| **Networks** | This option displays the list of preferred wireless networks. |
| **Add/Edit/Remove** | These buttons create a new preferred wireless network, modify the settings of the selected preferred wireless network, and delete the selected preferred wireless network. |
| **Move Up/Move Down** | These buttons move the selected preferred wireless network up or down in the Networks list. |

The properties of a preferred wireless network consist of a Network Properties tab and an IEEE 802.1X tab.

On the **Network Properties** tab, you can view and configure the settings of a wireless network for a Windows wireless client that supports the Wireless Zero Configuration service. The settings are:

- **Network Name (SSID).** Displays the names of networks that are within the reception range of the wireless adapter on your computer. If selected, a network key is used to encrypt data while it is transmitted over the network.

- **Description.** This option allows you to type a description for the preferred wireless network.

- **Data Encryption (WEP Enabled).** Specifies whether Wired Equivalent Privacy (WEP) is enabled for this preferred wireless network.

- **Network Authentication (Shared Mode).** Specifies whether a network key is used to authenticate to this preferred wireless network. Under the 802.11 standard for wireless networks, when a network key is used for authentication, the network is operating in shared key authentication mode. If a network key is not used for authentication, the network is operating in open system mode.

- **The Key Is Provided Automatically.** Specifies whether a network key is automatically provided for this preferred wireless network. For example, if 802.1X is being used for dynamic key distribution.

- **This Is A Computer-to-Computer (Ad Hoc) Network; Wireless Access Points Are Not Used.** Specifies whether the preferred wireless network is a computer-to-computer (ad hoc) network, not an access point (infrastructure) network.

On the **IEEE 802.1X** tab, you can view and configure the following settings (which are equivalent to the authentication settings of a Windows wireless client):

- **Enable Network Access Control Using IEEE 802.1X.** Specifies whether 802.1X authentication is enabled. If selected, credentials such as smart cards, certificates, and passwords are used for authentication.

- **EAP Type and Settings.** Specifies the Extensible Authentication Protocol (EAP) type that is to be used.

- **Authenticate As Guest When User Or Computer Information Is Unavailable.** Specifies whether the computer attempts authentication to the network when a user is not logged on.

- **Authenticate As Computer When Computer Information Is Available.** Specifies whether the computer attempts authentication to the network as a guest when user or computer information is not available.

The following are additional settings on the IEEE 802.1X tab that do not appear on authentication settings of a Windows wireless client:

- **EAPOL-Start Message.** This option allows you to specify the transmission behavior of the EAPOL-Start message when authenticating. These settings set the SupplicantMode registry setting. You can select from the following:

    - Do Not Transmit

    - Transmit

    - Transmit Per 802.1X

- **Max Start.** This option allows you to specify the number of successive EAPOL-Start messages that are sent out when no response is received to the initial EAPOL-Start messages.

- **Start Period.** This option allows you to specify the interval, in seconds, between the retransmission of EAPOL-Start messages when no response is received to the previously sent EAPOL-Start message.

- **Held Period.** This option allows you to specify the period, in seconds, for which the authenticating client will not perform any 802.1X authentication activity after it has received an authentication failure indication from the authenticator.

- **Authentication Period.** This option allows you to specify the interval, in seconds, for which the authenticating client will wait before retransmitting any 802.1X requests after end-to-end 802.1X authentication has been initiated.

- **Computer Authentication.** This option allows you to specify how computer authentication works with user authentication. These settings set the AuthMode registry setting. There are three possible settings:

    - **With User Authentication.** When users are not logged on to the computer, authentication is performed using the computer credentials. After a user logs on to the computer, authentication is maintained using the computer credentials. If a user travels to a new wireless access point, authentication is performed using the user credentials.

    - **With User Re-Authentication.** When users are not logged on to the computer, authentication is performed using the computer credentials. After a user logs on to the computer, authentication is performed using the user credentials. When a user logs off of the computer, authentication is performed with the computer credentials. This is the recommended setting because it ensures that the connection to the wireless AP is always using the security credentials of the computer's current security context (computer credentials when no user is logged on and user credentials when a user is logged on).

    - **Computer Only.** Authentication is always performed using the computer credentials. User authentication is never performed.

For more information about Group Policy settings, see the Group Policy Settings Reference for Windows Server 2003.

**WPA upgrade**

To use the new Wi-Fi Protected Access (WPA) standard for wireless clients running Windows XP (SP1 and later) and Windows Server 2003 that are using a wireless network adapter that supports the WZC service, you must obtain and install the free download WPA Wireless Security Update. It updates the wireless network configuration dialog boxes to support new WPA options. An update to the properties of the Wireless Network (IEEE 802.11) Policies to allow configuration of WPA options is being investigated.

# Network Interoperability Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Interoperability Technologies**
Microsoft Windows Server 2003 interoperability technologies provide the core functionality that enables remote administration and directory service synchronization. The Remote Procedure Call (RPC) functionality in Windows Server 2003 enables the implementation of a comprehensive network access solution. The Directory Synchronization Services (MSDSS) and the File Migration Utility (FMU) components in Services for Netware 5.03 (SFN) enable computers running Windows Server 2003 to interoperate with computers running Novell NetWare.

**Network Interoperability Technologies Components**
The following network access technologies provide options for secure client access to an organization's network.

**Remote Procedure Call**
RPC is technology for creating distributed client/server programs. RPC is an efficient and secure interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality residing in a different process. That process can be on the same computer, on the local area network (LAN), or across the Internet.

For more information, see "RPC Technical Reference.

**Services for NetWare**
SFN provides directory object migration and synchronization between computers that are running Windows Server 2003 with Microsoft Active Directory and computers that are running Novell's NetWare operating system with NDS, eDirectory, and Bindery directory services.

For more information, see "Services for Netware Technical Reference.

**Network Interoperability Technologies Scenarios**
RPC is an IPC technique that allows client and server software to communicate. The Microsoft RPC facility is compatible with the Open Group's Distributed Computing Environment (DCE) specification for remote procedure calls and is interoperable with other DCE-based RPC systems, such as those for HP-UX and IBM AIX UNIX–based operating systems.

SFN is a synchronization, migration, and interoperability utility that enables the migration of objects from Novell's NetWare operating system with NDS, eDirectory, and Bindery directory services, to Windows Server 2003 with the Active Directory directory service.

# RPC Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RPC Technical Reference**
Microsoft Remote Procedure Call (RPC) is a powerful, robust, efficient, and secure interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality residing in a different process. That process can be on the same computer, on the local area network (LAN), or across the Internet.

In Microsoft Windows Server 2003 and Windows XP, there are three types of processes: native 32-bit processes, native 64-bit processes, and 32-bit processes running under the 32-bit process emulator on a 64-bit system (often referred to as WOW6432 processes). RPC fully supports the 64-bit version of Windows.

Microsoft RPC was initially developed as an extension to Open Software Foundation (OSF) RPC. With the exception of some of its advanced features, Microsoft RPC is interoperable with other vendors' implementations of OSF RPC.

**In this subject**


- What Is RPC?
- How RPC Works

# What Is RPC?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is RPC?**

**In this section**


- Terms and Definitions
- RPC Dependencies and Interactions


Microsoft Remote Procedure Call (RPC) is a powerful technology for creating distributed client/server programs. RPC is an interprocess communication technique that allows client and server software to communicate. The Microsoft RPC facility is compatible with the Open Group's Distributed Computing Environment (DCE) specification for remote procedure calls and is interoperable with other DCE-based RPC systems, such as those for HP-UX and IBM AIX UNIX–based operating systems.

Computer operating systems and programs have steadily gotten more complex over the years. With each release, there are more features. The growing intricacy of systems makes it more difficult for developers to avoid errors during the development process. Often, developers create a solution for their system or application when a nearly identical solution has already been devised. This duplication of effort consumes time and money and adds complexity to already complex systems.

RPC is designed to mitigate these issues by providing a common interface between applications. RPC serves as a go–between for client/server communications. RPC is designed to make client/server interaction easier and safer by factoring out common tasks, such as security, synchronization, and data flow handling, into a common library so that developers do not have to dedicate the time and effort into developing their own solutions.

**Terms and Definitions**

The following terms are associated with RPC.

**Client**

A process, such as a program or task, that requests a service provided by another program. The client process uses the requested service without having to "deal" with many working details about the other program or the service.

**Server**

A process, such as a program or task, that responds to requests from a client.

**Endpoint**

The name, port, or group of ports on a host system that is monitored by a server program for incoming client requests. The endpoint is a network-specific address of a server process for remote procedure calls. The name of the endpoint depends on the protocol sequence being used.

**Endpoint Mapper (EPM)**

Part of the RPC subsystem that resolves dynamic endpoints in response to client requests and, in some configurations, dynamically assigns endpoints to servers.

**Client Stub**

Module within a client application containing all of the functions necessary for the client to make remote procedure calls using the model of a traditional function call in a standalone application. The client stub is responsible for invoking the marshalling engine and some of the RPC application programming interfaces (APIs).

**Server Stub**

Module within a server application or service that contains all of the functions necessary for the server to handle remote requests using local procedure calls.

**RPC Dependencies and Interactions**

RPC is a client/server technology in the most generic sense. There is a sender and a receiver; data is transferred between them. This can be classic client/server (for example, Microsoft Outlook communicating with a server running Microsoft Exchange Server) or system services within the computer communicating with each other. The latter is especially common. Much of the Windows architecture is composed of services that communicate with each other to accomplish a task. Most services built into the Windows architecture use RPC to communicate with each other.

The following table briefly describes the services in Windows Server 2003 that depend on the RPC system service (RPCSS).

**Services That Depend on RPCSS**

| Service | Description |
|---|---|
| Background Intelligent Transfer Service | Transfers data between clients and servers in the background. |
| COM+ Event System | Supports System Event Notification Service (SENS), which provides automatic distribution of events to subscribing Component Object Model (COM) components. |
| COM+ System Application | Manages the configuration and tracking of COM+-based components. |
| Cryptographic Services | Provides three management services: Catalog Database Service, which confirms the signatures of Windows files; Protected Root Service, which adds and removes Trusted Root Certification Authority certificates from this computer; and Key Service, which helps enroll this computer for certificates. |
| DHCP Server | Performs TCP/IP configuration for DHCP clients, including dynamic assignments of IP addresses, specification of the WINS and DNS servers, and connection–specific Domain Name System (DNS) names. |
| Distributed Link Tracking Client | Enables client programs to track linked files that are moved within an NTFS volume to another NTFS volume on the same computer or to an NTFS volume on another computer. |
| Distributed Link Tracking Server | Enables the Distributed Link Tracking Client service within the same domain to provide more reliable and efficient maintenance of links within the domain. |
| Distributed Link Transaction Coordinator | Coordinates transactions that span multiple resource managers, such as databases, message queues, and file systems. |
| DNS Server | Enables DNS clients to resolve DNS names by answering DNS queries and dynamic update requests. |
| Error Reporting Service | Collects, stores, and reports unexpected application failures to Microsoft. |
| File Replication Service | Allows files to be automatically copied and maintained simultaneously on multiple servers. |
| Help and Support | Enables Help and Support Center to run on the computer. |
| Human Interface Device Access | Enables generic input access to Human Interface Devices (HID), which activates and maintains the use of predefined hot buttons on keyboards, remote controls, and other multimedia devices. |
| Indexing Service | Indexes contents and properties of files on local and remote computers; provides rapid access to files through flexible querying language. |
| IPSec Services | Provides end-to-end security between clients and servers on TCP/IP networks. |
| Kerberos Key Distribution Center | On domain controllers, enables users to log on to the network using the Kerberos authentication protocol. |
| Logical Disk Manager | Detects and monitors new hard disk drives and sends disk volume information to Logical Disk Manager Administrative Service for configuration. |
| Logical Disk Manager Administrative Service | Configures hard disk drives and volumes. |

| | |
|---|---|
| Messenger | Transmits net send and Alerter service messages between clients and servers. This service is not related to Windows Messenger. |
| Microsoft Software Shadow Copy Provider | Manages software-based volume shadow copies taken by the Volume Shadow Copy service. |
| Network Connections | Manages objects in the Network and Dial-Up Connections folder in which you can view local area network (LAN) and remote connections. |
| Print Spooler | Manages all local and network print queues and controls all printing jobs. |
| Protected Storage | Protects storage of sensitive information, such as private keys and prevents access by unauthorized services, processes, or users. |
| Remote Desktop Help Session Manager | Manages and controls Remote Assistance. |
| Remote Registry | Enables remote users to modify registry settings on a computer. |
| Removable Storage | Manages and catalogs removable media and operates automated removable media devices. |
| Resultant Set of Policy Provider | Enables a user to connect to a remote computer, access the Windows Management Instrumentation (WMI) database for that computer, and either verify the current Group Policy settings made for the computer or check settings before they are applied. |
| Routing and Remote Access | Enables multi-protocol LAN-to-LAN, LAN-to-wide area network (WAN), virtual private network (VPN), and network address translation (NAT) routing services for clients and servers on the network. |
| Security Accounts Manager | Upon startup, signals other services that the Security Accounts Manager (SAM) is ready to accept requests. |
| Shell Hardware Detection | Provides notifications for AutoPlay hardware events. |
| Task Scheduler | Enables a user to configure and schedule automated tasks on the computer. |
| Telephony | Provides Telephony API (TAPI) support for clients using programs that control telephony devices and IP-based voice connections. |
| Telnet | Enables a remote user to log on to a computer and run programs; supports various TCP/IP Telnet clients, including UNIX-based and Windows-based computers. |
| Terminal Services | Allows users to connect interactively to a remote computer. Remote Desktop, Fast User Switching, Remote Assistance, and Terminal Server depend on this service. |
| Terminal Services Session Directory | Enables a user connection request to be routed to the appropriate terminal server in a cluster. |
| Upload Manager | Manages the synchronous and asynchronous file transfers between clients and servers on the network. |
| Virtual Disk Service | Provides software volume and hardware volume management service. |

| | |
|---|---|
| Volume Shadow Copy | Manages and implements Volume Shadow Copies used for backup and other purposes. |
| Windows Audio | Manages audio devices for Windows-based programs. |
| Windows Image Acquisition (WIA) | Provides image acquisition services for scanners and cameras. |
| Windows Installer | Installs, repairs, and removes software according to instructions contained in .MSI files. |
| Windows Internet Name Service (WINS) | Resolves NetBIOS names for TCP/IP clients by locating network services that use NetBIOS names. |
| Windows Management Instrumentation | Provides a common interface and object model to access management information about operating system, devices, applications, and services. If this service is stopped, most Windows-based software will not function properly. |
| Wireless Configuration | Enables automatic configuration for IEEE 802.11 adapters. |
| WMI Performance Adapter | Provides performance library information from WMI providers to clients on the network. |

# How RPC Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How RPC Works**
**In this section**

- RPC Architecture
- RPC Processes and Interactions
- Network Ports Used by RPC

Microsoft Remote Procedure Call (RPC) is an interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality residing in a different process. That process can be on the same computer, on the local area network (LAN), or across the Internet. The Microsoft RPC mechanism uses other IPC mechanisms, such as named pipes, NetBIOS, or Winsock, to establish communications between the client and the server. With RPC, essential program logic and related procedure code can exist on different computers, which is important for distributed applications.

This section covers the architecture of RPC and the way in which RPC communication takes place.

**RPC Architecture**
By replacing dedicated protocols and communication methods with a robust and standardized interface, RPC is designed to facilitate communication between client and server processes. The functions contained within RPC are accessible by any program that must communicate using a client/server methodology. The following figure shows the RPC architecture.

**RPC Architecture**




The following table lists and describes the components and functions of the RPC architecture.

**RPC Components**

| Component | Description |
|---|---|
| Client or server process | Program or service that initiates or responds to an RPC request. |
| RPC stubs | Program subsystems used by a client or server to initiate an RPC request. |

| | |
|---|---|
| Marshalling engine<br><br>(NDR20 or NDR64) | Provides a common RPC interface between RPC clients and servers. NDR20 is used in a 32-bit architecture and NDR64 is optimized for a 64-bit architecture. The client and the server negotiate which marshalling engine is used for the communication. |
| Runtime application programming interface (API) | Provides a direct interface for RPC to clients or servers. RPC clients and servers typically call the runtime API to initialize RPC and prepare the data structure that is used to make RPC calls. This runtime API layer also determines if an RPC request coming from a marshalling engine or directly from a client or server is going to a local server or a remote server. The runtime API layer then routes the RPC to the Connection RPC, Datagram RPC, or Local RPC layers. |
| Connection RPC protocol engine | Used when the RPC requires a connection–oriented protocol. This layer designates the connection–oriented protocol to use if the RPC is outgoing or receives an incoming connection–oriented RPC. |
| Datagram RPC protocol engine | Used when the RPC requires a connectionless protocol. This layer designates the connectionless protocol to use if the RPC is outgoing or receives an incoming connectionless RPC. |
| Local RPC protocol engine | Used when the server and client are located on the same host. |
| Registry | Accessed when the RPC service first loads. Registry keys may specify IP port ranges and the device names of network cards that RPC should bind to. Unless APIs force its use, the registry is not used in normal RPC operations. |
| Win32 APIs<br><br>(kernel32.dll, advapi32.dll, ntdll.dll) | Kernel32.dll is a Windows NT base API client dynamic-link library (DLL) file that provides system services for managing threads, memory, and resources.<br><br>Advapi32.dll is an advanced Windows 32 base API DLL file; it is an API services library that supports security and registry calls.<br><br>Ntdll.dll is an NT layer DLL file that controls Windows NT system functions. |
| SSPI<br><br>(secur32.dll) | Provides a security interface for RPC. Negotiates the use of Kerberos, NTLM, or Secure Sockets Layer (SSL) for authentication and encryption. |
| Endpoint Mapper (EPM)<br><br>(rpcss.dll) | Rpcss.dll primarily provides the infrastructure for COM, but a portion of rpcss.dll is used for the EPM. An RPC server contacts the EPM to receive dynamic endpoints and register those endpoints in the EPM database. RPC clients contact the EPM from the protocol-engine level to resolve endpoints when establishing a connection with an unknown RPC server endpoint. |
| Active Directory | Used in the RPC client process only when the security interface specifies Kerberos or Negotiate as the security provider or when the server uses NTLM as the security provider. |
| Network stack | Used to pass RPC requests and replies between a client and a remote server. |
| Kernel | Used to pass RPC requests and replies between a client and a local server. |

**RPC Processes and Interactions**

The RPC components make it easy for clients to call a procedure located in a remote server program. The client and server each have their own address spaces; that is, each has its own memory resource allocated to data used by the procedure. The following figure shows the RPC process.

**RPC Process**




The RPC process starts on the client side. The client application calls a local stub procedure instead of code implementing the procedure. Stubs are compiled and linked with the client application during development. Instead of containing code that implements the remote procedure, the client stub code retrieves the required parameters from the client address space and delivers them to the client runtime library. The client runtime library then translates the parameters as needed into a standard Network Data Representation (NDR) format for transmission to the server.

**Note**

- There are two NDR marshalling engines within the RPC runtime library: NDR20 and NDR64. A 32-bit client initiating the communication uses the NDR20 marshalling engine; a 64-bit client can use either the NDR20 or the NDR64 marshalling engine. The same marshalling engine is used on both the client and the server side, regardless of program architecture. There is a slight decline in performance when either the client or server uses an architecture different from the other because the marshalling engine must do additional translation during the communication.

The client stub then calls functions in the RPC client runtime library (rpcrt4.dll) to send the request and its parameters to the server. If the server is located on the same host as the client, the runtime library can use the Local RPC (LRPC) function and pass the RPC request to the Windows kernel for transport to the server. If the server is located on a remote host, the runtime library specifies an appropriate transport protocol engine and passes the RPC to the network stack for transport to the server. RPC can use other IPC mechanisms, such as named pipes and Winsock, to accomplish the transport. The other IPC mechanisms allow RPC more flexibility in the way in which it completes its communications tasks. For more information about IPC mechanisms, see “Interprocess Communications” on MSDN.

The following table lists the network protocols supported by RPC and the type of RPC connection for which the protocol is used.

**RPC-Supported Network Protocols**

| Protocol | RPC Type |
|---|---|
| Transmission Control Protocol (TCP) | Connection–oriented |
| Sequenced Packet Exchange (SPX) | Connection–oriented |
| Named Pipe | Connection–oriented |
| HTTP | Connection–oriented |
| User Datagram Protocol (UDP) | Connectionless |
| Cluster Datagram Protocol (CDP) | Connectionless |

When the server receives the RPC, either locally or from a remote client, the server RPC runtime library functions accept the request and call the server stub procedure. The server stub retrieves the parameters from the network buffer and, using one of the NDR marshalling engines, converts them from the network transmission format to the format required by the server. The server stub calls the actual procedure on the server. The remote procedure then runs, possibly generating output parameters and a return value. When the remote procedure is complete, a similar sequence of steps returns the data to the client.

The remote procedure returns its data to the server stub which, using one of the NDR marshalling engines, converts output parameters to the format required for transmission back to the client and returns them to the RPC runtime library functions. The server RPC runtime library functions transmit the data to the client computer using either LRPC or the network.

The client completes the process by accepting the data over the network and returning it to the calling function. The client RPC runtime library receives the remote-procedure return values, converts the data from its NDR to the format used by the client computer, and returns them to the client stub.

For Microsoft Windows, the runtime libraries are provided in two parts: an import library, which is linked to the application, and the RPC runtime library, which is implemented as a DLL.

The server application contains calls to the server runtime library functions, which register the server's interface with the RPC runtime and, optionally, the EPM, and allow the server to accept remote procedure calls. The server application also contains the application-specific remote procedures that are called by the client applications.

**RPC Security Context Multiplexing**

Windows Server 2003 Service Pack 1 (SP1) provides RPC security context multiplexing for connection-oriented connections, such as those that use Transmission Control Protocol (TCP). This allows the RPC server to negotiate multiple security contexts over a single connection. For example, when multiple RPC clients establish a connection to an RPC server and a middle-tier RPC server resides between the clients and the destination server, the middle-tier server multiplexes the security context of the additional clients over an already established connection to the destination server. This eliminates the need for the middle-tier RPC server to use one of the exhaustible TCP ports to establish a new connection for each RPC client connecting to the RPC server

**Network Ports Used by RPC**

RPC server programs typically use dynamic port mappings to avoid conflicts with programs and protocols registered in the range of well-known TCP ports. RPC server programs associate their universally unique identifier (UUID) with a dynamic port and register the combination with the RPC EPM. The EPM provides a single point of contact for RPC clients. The RPC clients contact the EPM and use the server program's UUID to determine the port being used by the server program. The following table indicates the network ports normally used by RPC.

**Port Assignments for RPC**

| Service Name | UDP | TCP |
|---|---|---|
| HTTP | 80, 443, 593 | 80, 443, 593 |
| Named Pipes | 445 | 445 |
| RPC Endpoint Mapper | 135 | 135 |
| RPC Server Programs | <Dynamically assigned> | <Dynamically assigned> |

# Services for Netware Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services for NetWare Technical Reference**
Microsoft Services for NetWare 5.03 is a synchronization, migration, and interoperability tool made up of the following components: Directory Synchronization Services (MSDSS) and File Migration Utility (FMU).

Services for NetWare enables an organization to migrate objects from Novell's NetWare operating system with NDS, eDirectory, and Bindery directory services to the Microsoft Windows Server 2003 operating system with the Active Directory directory service. Services for NetWare also enables short-term directory service interoperability and bi-directional synchronization.

**In this subject**


- What Is Services for Netware
- How Services for Netware Works
- Services for Netware Tools and Settings

# What Is Services for Netware

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services for NetWare Scenarios**
To enhance interoperability with NetWare networks, Services for NetWare 5.03 includes Directory Synchronization Services (MSDSS) and the File Migration Utility (FMU):

- **Directory Synchronization Services** — Enables you to migrate and synchronize directory resources from NetWare servers.
- **File Migration Utility** — Enables you to migrate files from NetWare servers.

Services for NetWare 5.03 also provides Active Directory and NDS interoperability. This functionality enables you to add Active Directory to a Bindery or NDS environment. In addition, Services for NetWare provides tools that enable you to troubleshoot connectivity problems, login scripts, and password synchronization. Services for NetWare also provides utilities for monitoring network traffic.

Services for NetWare 5.03 requires Novell Client for Windows.

**Note**

- File and Print Services for NetWare (FPNW) has been removed from Services for NetWare 5.03.

**Synchronization Overview**
Directory Synchronization Services (MSDSS) performs the following functions:

- Simplifies the integration of NDS/eDirectory and Active Directory.
- Reduces directory management overhead by enabling two-way synchronization.
- Speeds the migration of NDS/eDirectory and NetWare bindery information to Windows Server 2003.

To enable these capabilities, MSDSS bi-directionally synchronizes data stored in Active Directory with NDS or eDirectory.

To provide interoperability between Windows Server 2003 directory services and NetWare 3.x bindery services, MSDSS synchronizes data stored in Active Directory with all NetWare 3.x bindery services on a one-way basis. Bindery does not support two-way synchronization. You can also establish a one-way synchronization session with NDS or eDirectory.

For more information about the following synchronization methods used by MSDSS, see "How Services for Netware Works:"

- Forward Synchronization
- Reverse Synchronization
- One-Way Synchronization
- Two-Way Synchronization
- Scheduled Synchronization
- Manual Synchronization
- Synchronization Filtering
- Password Synchronization

**Migration Overview**
The File Migration Utility (FMU) enables you to migrate files from NetWare servers to Windows Server 2003. FMU can copy files from multiple NetWare file servers to multiple file servers running Windows Server 2003.

FMU is integrated with MSDSS and enables the migration of large amounts of file data from supported versions of NetWare over to Windows Server 2003, while preserving unique file-security permissions during the migration. In addition, users can access their files during the migration process.

Before you use FMU to translate the NDS, eDirectory, or Bindery file system rights and permissions to the equivalent rights and permissions in the NTFS file system, you must run a one-time migration of NDS or Bindery directory objects to Active Directory. After you have performed the one-time migration and selected the option to migrate files, a migration log will be created that can be used later by FMU. This migration log serves as a mapping file that maps user objects in NDS to Active Directory.

For each NDS or Bindery volume, FMU calculates the effective rights, starting at the root, and tracks inheritance and filters as the migration continues through the tree. Like Novell NetWare, Microsoft Windows secures folders and files by controlling group and user access. Windows security is supported through the NTFS file system. To help preserve security on transferred NetWare files, you must transfer the NetWare files to an NTFS volume, where the effective rights of the NetWare files are converted to permissions.

**Interoperability Overview**
If you want to maintain an environment that contains both Active Directory and Bindery/NDS directory services, you can perform a phased migration and run the two systems in parallel. If you choose to run the two systems in parallel, you will be able to perform additional migration tasks (other than synchronizing the two directory services), such as replacing applications that are dependent on Novell services with applications that are compatible with Active Directory.

**Related Information**
The following resources contain additional information that is relevant to this section.

- How Services for Netware Works
- Services for Netware Tools and Settings

# How Services for Netware Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services for NetWare Architecture**
Services for NetWare 5.03 includes two components: Directory Synchronization Services and the File Migration Utility (FMU).

**Directory Synchronization Services**
Using Directory Synchronization Services (MSDSS) an organization can retain an existing Novell directory in an environment that includes computers running Windows Server 2003 and Active Directory. MSDSS synchronization enables you to choose:

- One-way synchronization to centralize directory management from Active Directory.
- Two-way synchronization to manage data from either Active Directory or a Novell directory.

**File Migration Utility**
Using FMU, an organization can migrate a Novell Bindery or NDS directory to Windows Server 2003 Active Directory.

**MSDSS**
The following sections summarize the synchronization methods that are available with Directory Synchronization Services (MSDSS):

- Forward Synchronization
- Reverse Synchronization
- One-Way Synchronization
- Two-Way Synchronization
- Scheduled Synchronization
- Manual Synchronization
- Synchronization Filtering
- Password Synchronization

**Forward Synchronization**
The process of synchronizing data from Active Directory to Novell directories is referred to as forward synchronization. The forward synchronization process efficiently queries Active Directory for any new objects and alterations to existing objects. If an existing object is modified, only its changes, not the entire object, are synchronized with NDS, eDirectory, or Bindery directory services objects. If a new object is created, then only the new object and its attributes are synchronized with NDS, eDirectory, or Bindery directory services objects. Synchronization takes place during the scheduled synchronization or when you initiate a manual synchronization.

**Reverse Synchronization**
The process of synchronizing data from Novell directories to Active Directory is referred to as reverse synchronization. Reverse synchronization is less efficient than forward synchronization. This means that network utilization during reverse synchronization is higher than it is during forward synchronization.

NDS does not currently provide a comprehensive mechanism to query changes to objects. Therefore, during reverse synchronization, an MSDSS algorithm compares all of the objects residing in the NDS directory to the objects residing in Active Directory. Only changed objects and newly created objects in NDS are copied to Active Directory.

Because network utilization during reverse synchronization can be higher, it should be scheduled to run less frequently than forward synchronization. However, if you decrease the frequency of directory synchronization in either direction, you will increase the likelihood of inconsistencies between Active Directory and NDS.

**One-Way Synchronization**
Using one-way synchronization, you can deploy Active Directory in a Novell environment without replacing existing directories or having to manage two directories. This enables you to:

- Consolidate network management where multiple directories are required.
- Manage accounts from a single directory.
- Use directory-enabled applications, devices, and services based on Active Directory.

After an initial reverse synchronization is performed (from NDS, eDirectory, or Bindery directories to Active Directory), the NDS, eDirectory, or Bindery directory objects that reside in Active Directory can be forward-synchronized on a one-way basis. When users, groups, organizational units (OUs), or any of their applicable attributes are modified in Active Directory, they will be synchronized with the NDS, eDirectory, or Bindery directories during the next scheduled synchronization.

Setting up a one-way synchronization session provides an administrator with a central directory management solution and supports additions or modifications to NDS, eDirectory, or Bindery directory services objects from Active Directory only.

**Two-Way Synchronization**
Using two-way synchronization, you can create or change objects in either Active Directory or NDS and the partner directory will be updated at the next scheduled synchronization.

You can use two-way synchronization to enable interoperability between NDS and Active Directory and to maintain an existing NDS directory and add Active Directory.

**Scheduled Synchronization**
During a scheduled synchronization, changes from one directory are replicated to another directory at scheduled intervals that are set by an administrator. The default interval for forward synchronization is every 15 minutes, 24 hours a day. The default interval for reverse synchronization is every hour, from midnight to 06:00. If a two-way synchronization session has been configured, you can define a separate schedule for each direction of the synchronization.

**Manual Synchronization**
When you manually synchronize directories, you have the option of copying recent changes from one directory to another immediately, without waiting for the next scheduled synchronization.

Usually, it is safe to let synchronization occur during the normally scheduled times. However, a security-related situation might call for an immediate manual synchronization. For example, users with access to both directory platforms might request a password reset or logon change, or user accounts might need to be immediately disabled.

When user accounts are disabled in Active Directory, access to those accounts from Active Directory is immediately restricted. However, users will still have access to their accounts from NDS until a scheduled synchronization takes place. To prevent this situation from occurring, you can perform a manual synchronization.

**Synchronization Filtering**
Using synchronization filtering, administrators can define an explicit set of object types within a session to be read during the synchronization process. Object types are collections of attributes. When a synchronization filter is applied, changes to a defined object type are synchronized on the next scheduled synchronization.

**Password Synchronization**
Password synchronization can occur only if passwords are changed from Active Directory. Passwords are synchronized when an administrator performs any of the following administrative functions:

- Initial reverse synchronization
- User creation in NDS or eDirectory
- Password changes in Active Directory

**Initial Reverse Synchronization**
Users that have been copied from NDS, eDirectory, or Bindery to Active Directory through initial reverse synchronization will be prompted to change their passwords the first time they log on to Active Directory. This new password will then be synchronized with the password attribute in NDS, eDirectory, or Bindery.

**User Creation in NDS or eDirectory**
If users are created in NDS or eDirectory (in a two-way synchronization), the new user objects are copied to Active Directory on the next scheduled reverse synchronization. When the new users successfully log on to Active Directory, they are prompted to change their passwords. The new passwords are copied to NDS or eDirectory on the next scheduled forward synchronization.

**Password Changes in Active Directory**
When users change their passwords or when an administrator resets user passwords in Active Directory, the new passwords overwrite the existing NDS, eDirectory, or Bindery passwords. This password synchronization process occurs during the next scheduled forward synchronization.

### Note

- If the **User must change password at next logon** check box is selected for any user, the password will not be synchronized from Active Directory to the Novell directory until the user logs on to Active Directory and changes the password.

Because it is impossible to access encrypted passwords stored in NDS, eDirectory, and Bindery, new user passwords are created when user objects are migrated to Active Directory. Any password changes made from NDS or eDirectory will not be synchronized with Active Directory. The new password is generated using password generation schemes that are specified by Services for NetWare.

The password scheme is applied and then copied to Active Directory before the user object is migrated. When the user accounts exist in Active Directory, users will be prompted to create new passwords when logging on to Active Directory.

Changes made to user passwords in Active Directory will be synchronized with the respective Novell user objects at the next scheduled synchronization.

User password schemes only affect user objects that exist in NDS, and use of a password scheme will not overwrite a valid user password in Active Directory. Because of this password scheme rule, it is recommended that all future password changes occur in Active Directory. Otherwise, if a user password is changed in NDS, the passwords in both directories will be inconsistent until it is changed in Active Directory.

The following table lists and describes the password schemes that are available to administrators.

**Password Schemes Available to Administrators**

| Password Schemes | Description |
|---|---|
| Set passwords to blank. | Clears all passwords. Users logging on to Active Directory for the first time will not need to specify a password. |
| Set passwords to the name of the user. | Sets passwords to user names after migration has taken place. This is the default scheme. |
| Set passwords to a random value. | Creates random passwords for each user migrated to Active Directory. The password generated by SFN does meet Windows Server 2003 default password policy. |
| Set all passwords to the following value. | Provides an option to specify and confirm the password value for all users migrating to Active Directory. |

**MSDSS Directory Schema Extensions**
Extensions to the Active Directory and NDS directory schemas are required for full MSDSS functionality.

### Note

- MSDSS directory schema extensions are implemented once and are irreversible.

A directory schema provides a description of object classes and associated attributes that can be placed in a directory. For each object class, the directory schema defines:

- The required attributes of the object class
- Any additional attributes of the object class

- The parent object class of the defined object class

**Active Directory Schema Extensions**
Extensions to the Active Directory schema are required for:

- Migration
- One-way synchronization
- Two-way synchronization

If a user who is installing MSDSS has the necessary permissions to extend the Active Directory schema, MSDSS automatically updates the Active Directory schema during the setup process. The schema update process is required only once for each Active Directory forest. If a user who is installing MSDSS does not have the necessary permissions to extend the Active Directory schema, the schema can be updated manually by a user with the appropriate schema credentials. To manually update the schema, log on to the schema master in the Active Directory forest in which MSDSS is to be installed with the appropriate schema credentials, then open a command prompt and type:

```
msiexec /I Path \msdss.msi SCHEMAONLY=1
```

**NDS Directory Schema Extensions**
Extensions to the NDS directory schema are required for two-way synchronization.

To support two-way synchronization, you must also extend the NDS schema. If you are required to extend the NDS schema when you create a new session, the system prompts you. Extending the NDS schema requires the following:

- Installation of Novell Client for Windows (also a requirement for MSDSS).
- Administrator permission to access the root object of the NDS tree.
- Access to the server that stores the master replica of the root partition (this server propagates the changes to all the servers in the NDS tree).

You can also extend the NDS or eDirectory schema manually by using the MSDSS command-line NDS schema extension tool NDSext.exe (available in the *systemroot*\System32\Directory Synchronization\Client folder). Use the following syntax:

```
ndsext extend Treename Username.Context Password
```

This command enables you to extend the NDS schema of a specified tree.

**File Migration Utility**
The File Migration Utility (FMU) enables you to migrate files from NetWare servers to Windows Server 2003. FMU can copy files from multiple NetWare file servers to multiple file servers running Windows Server 2003.

FMU is integrated with MSDSS and enables the migration of large numbers of data files from all versions of NetWare over to Windows Server 2003, while preserving unique file-security permissions during the migration. In addition, users can access their files during the migration process.

Before using FMU to translate the NDS, eDirectory, or Bindery file system rights and permissions to the equivalent rights and permissions in the NTFS file system, you must run a one-time migration of NDS or Bindery directory objects to Active Directory. This operation creates the mapping file used by FMU to map user rights to folders. After a scan confirms that all files, subdirectories, and permissions will migrate successfully, you can perform the migration process. Each source directory and file is copied to the selected target. Each migrated trustee is translated with its associated rights to the equivalent NTFS permissions and then added to the security descriptor for the migrated directory or file.

For each NDS or Bindery volume, FMU calculates the effective rights, starting at the root, and tracks inheritance and filters as the migration continues through the tree. Like Novell NetWare, Microsoft Windows secures folders and files by controlling group and user access. Windows security is supported through the NTFS file system. To help preserve security on transferred NetWare files, you must transfer the NetWare files to an NTFS volume, where the effective rights of the NetWare files are converted to permissions.

After you have performed the one-time migration and selected the option to migrate files, a migration log will be created that can be used later by FMU. This migration log serves as a mapping file that maps user objects in NDS to Active Directory.

**Accounts Required for FMU**
FMU requires the following accounts for its security context:

- The account that the user provides for NDS or Bindery access. The security context of this account is used to read the directories, files, and associated trustees that reside on the source Novell system.

- The user accounts accessed through Active Directory for users who are logged on. The security context of this account is used for creating directories and files and setting the associated security descriptor on the targeted Active Directory domain.

**Sources and Targets**
During a migration, the source is always either a volume or a directory from an NDS or Bindery server. You can have multiple mappings in any migration. One source can be mapped to multiple targets or multiple sources can be mapped to a single target. The migration target must be a shared folder on Microsoft Windows 2000 or Windows Server 2003.

You can display the source and target listings by using Active Directory Service Interfaces (ADSI) calls and then use the Lightweight Directory Access Protocol (LDAP), NDS, and NWCOMPAT providers to populate the lists. If any of these providers are missing, you might get blank listings.

**FMU and MSDSS**
MSDSS is an essential part of the file migration process because MSDSS informs FMU of the user and group relationships with the corresponding objects in NDS or Bindery. Within NDS, users, groups, organizational units, and organizations can all be trustees to directories and files.

**Note**

- Trustees are granted rights to work with directories, files, or objects. Active Directory does not support a container such as an organization to be the trustee of a directory or file. For more information about NDS or Bindery inheritance, see Novell documentation.

To maintain these rights throughout the migration of files from Novell NetWare to Microsoft Windows, MSDSS maps organizational units and organizations to Active Directory groups by creating local domain security groups for each NDS organizational unit and organization.

**MSDSS Migration Logs for FMU**
FMU requires the MSDSS-generated migration log to determine where each Novell trustee has been mapped in Active Directory. The migration log contains a list of trustee objects that were migrated to Active Directory during the MSDSS one-time migration process, and the log also lists the groups of trustees that were made members in Active Directory. Only trustees and their group associations that are listed in the migration log will be used to ensure equivalent security for files during the file migration process. The migration log contains a list of trustee objects that were migrated to Active Directory during the MSDSS one-time migration process. MSDSS can recognize the following NDS, eDirectory, or Bindery trustee objects:

- Users

- Groups

- Organizational units

- Organizations

Because only the trustee objects from NDS or eDirectory organizational units or organizations that you selected during the MSDSS one-time migration process will appear in the migration log, you will need to know which trustees have restrictions to the files you will be migrating to Windows and where these trustees reside in your NDS or eDirectory hierarchy.

You can use the migration log **Maps** option to display the maps that are loaded with the migration log. You can use **Maps** to confirm the trustee relationships to their corresponding objects in Active Directory before you perform a file migration. You can view the following maps:

- NDS Organizations and organizational units to Active Directory Group Mappings
- NDS Group to Active Directory Group Mappings
- NDS User to Active Directory User Mappings

You can use the migration log **Access Rights** option to view NetWare rights that are equivalent to NTFS permissions.

**Note**

- The NDS **Modify** right does not have an equivalent NTFS permission.

**Directory Service Interoperability**

Using the Services for NetWare directory service interoperability features, an organization can establish a synchronization strategy over a period of weeks or months, during which Active Directory interoperates with Novell Bindery, NDS, or eDirectory. This enables an organization to perform a migration to Windows Server 2003 in planned stages.

To maintain an environment that contains both Active Directory and Bindery/NDS directory services, you can perform a phased migration and run the two systems in parallel. A phased migration reduces the risk of data loss because you can migrate in managed stages and you can reverse the process if necessary. However, maintaining two separate directory services can, over time, add additional administrative costs to the migration.

In a phased migration, you use MSDSS to copy all Bindery or NDS user accounts, groups, and distribution lists, and (for NDS only) organizational units and organizations to Active Directory, while you maintain these objects in NDS or Bindery. While you gradually move resources to the Windows Server 2003-based environment, the MSDSS-provided directory synchronization enables users to continue to access those resources that remain on NetWare servers. As the changeover continues, users begin to access resources on computers running Windows Server 2003.

If you plan to perform a migration over several months or longer, two-way synchronization is recommended. Two-way synchronization should be used if the long-term goal is to interoperate using the two directory services. In this scenario, you can replace existing NetWare technologies with the latest comparable Microsoft technologies over a period of time.

After you have moved all resources to Windows Server 2003, converted all Novell services and applications to their counterparts in Active Directory, and moved object security permissions and objects that MSDSS does not migrate (such as computer accounts, printer objects, and application objects), synchronization between the two directory services might not be necessary. If not, you can delete the synchronization sessions and decommission remaining NetWare servers.

**Services for NetWare and MIIS**

Microsoft Identity Integration Server 2003 (MIIS) enables the integration and management of identity information across multiple repositories, systems, and platforms. MIIS augments Active Directory by providing broad interoperability capabilities, including: integration with a wide range of identity repositories; provisioning and synchronizing identity information, including password management, across multiple stores; and brokering changes to identity information by automatically detecting updates and sharing the changes across systems.

- MIIS also provides cross-platform identity synchronization across many different identity stores, including Novell's eDirectory. MIIS is a much more robust long-term cross-platform interoperability solution than Services for NetWare; however, MIIS does not support NDS. MIIS supports Novell's eDirectory version 8.6.2 and later.

  **Note**

  - NetWare 5.x and NetWare 6.x are supported by MIIS if the NDS directory has been upgraded to Novell's eDirectory:

MIIS currently supports the following versions of Novell's eDirectory:

- eDirectory 8.6.2
- eDirectory 8.7
- eDirectory 8.7.x

**Note**

- Support for eDirectory 8.7.x is included with Microsoft Identity Integration Server 2003, SP1.

For NetWare administrators who manage networks that include Novell software not supported by MIIS, Services for NetWare 5.03 can provide an interim Active Directory/NDS interoperability solution. For all other Novell customers, MIIS is the recommended long-term interoperability solution.

**Services for NetWare Troubleshooting**

Windows Server 2003 includes tools to help you monitor network traffic and to troubleshoot connectivity-related problems, including problems with routers, login scripts, client/server installation and configuration, and password synchronization. Among the troubleshooting tools provided with Windows Server 2003 are:

- **Netdiag** to help isolate TCP/IP-related connectivity problems.
- **Netstat** to display active TCP/IP connections, ports on which the computer is listening, Ethernet statistics, the IP routing table, IPv4 statistics (for the IP, ICMP, TCP, and UDP protocols), and IPv6 statistics (for the IPv6, ICMPv6, TCP over IPv6, and UDP over IPv6 protocols).
- **Ping** for troubleshooting IP-level connectivity between two TCP/IP computers.
- **Tracert** to trace routes and display the path between the sending host and a destination.
- **Pathping** to identify the path to a remote host and then ping the remote host for a period of time to collect and report statistics.
- **Event Viewer** to track errors and events recorded in Application, Security, and System logs.
- **Ipxroute** to display and modify information about the routing tables used by the IPX protocol.
- **Network Monitor** to monitor, record, and analyze network traffic.

For more detailed information about these and other troubleshooting tools, see "TCP/IP Tools and Settings" and "Services for Netware Tools and Settings."

**Connectivity Problems**

Most connectivity problems between computers running Windows Server 2003 and NetWare are caused by invalid configurations. If these systems are incorrectly configured, NetWare might not be able to access resources or it might access them inconsistently. Because resource access involves interaction between both Windows Server 2003-based and NetWare-based computers, you must be sure that resource access configurations are accurate on these computers.

In general, any of the following conditions could be symptomatic of protocol related issues, such as a scenario that includes IPX configured on a NetWare server while a computer running Windows Server 2003 is using IP.

- Access is denied to applications residing on NetWare servers.
- Programs fail to run and display error messages.
- Data throughput is slow.
- Users see the error message "Server not found."
- Users are denied access to network resources located on NetWare servers.
- Access to NetWare resources is limited.
- Some clients have some access to NetWare services, while others have none.

- Users can connect to servers running Windows Server 2003, but not to NetWare servers.
- Users can connect to some NetWare servers, but not to others.

To resolve some of the connectivity issues, you can verify that NetWare rights and user credentials are configured correctly by asking the following questions:

- Have you configured user accounts correctly on the NetWare file server?
- Have you set up the appropriate groups?
- Is group membership correctly configured?
- Have you set up user rights correctly for required resources?

**Note**

- For more information about NetWare configuration procedures, contact your NetWare administrator or refer to the Novell NetWare documentation.

**Router Problem**

In addition to potential NetWare and Windows Server 2003 connectivity problems, you might also experience the following router problem: NetBIOS-over-IPX packets do not propagate across routers.

If an incomplete browse list is displayed for a NetWare directory, or you cannot browse a server across a WAN link, you might have a problem with propagating packets using NetBIOS over IPX.

To facilitate the communication of a LAN-based protocol (such as NetBIOS) across an IPX network, IPX routers must be capable of propagating NetBIOS-over-IPX broadcast packets (type 20 packets), also known as IPX WAN broadcast packets. However, many router manufacturers set router parameters that prevent passing type 20 packets by default. If your router is not configured to pass type 20 packets, contact the router manufacturer for reconfiguration procedures.

**Services for NetWare Protocols**

NWLink, the IPX/SPX/NetBIOS-Compatible Transport Protocol, is Microsoft's 32-bit implementation of Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX). NWLink is included with Windows Server 2003.

NWLink provides the transport protocol to support communications with NetWare file servers. To log on to a NetWare network from Windows Server 2003, you must use a NetWare client such as Novell Client for Windows XP Professional.

**Note**

- You can also use NWLink to enable NetWare clients to connect to Windows Server 2003, Windows XP Professional, or Windows 2000 resources.

Because NWLink is Network Driver Interface Specification (NDIS)-compliant, a computer running Windows Server 2003 can also simultaneously run other protocol stacks, such as TCP/IP, for communication with computers running Windows with TCP/IP. In addition, you can bind NWLink to multiple network adapters with multiple frame types.

NWLink supports the networking application programming interfaces (APIs) NetBIOS and Windows Sockets (Winsock). These APIs enable communication among computers running Windows Server 2003 and between computers running Windows Server 2003 and NetWare bindery or NDS-based servers.

In addition to using NWLink to connect computers running Windows Server 2003 and NetWare, you can use NWLink whenever you need IPX/SPX. For example, when an IPX/SPX-based protocol is required, you can use NWLink as an alternative protocol to connect computers running Microsoft Proxy Server for Windows, Microsoft Systems Management Server (SMS), Microsoft Host Integration Server, Microsoft SQL Server, or Microsoft Exchange Server.

NWLink requires no initial client configuration on small non-routed networks.

**Note**

- The 64-bit versions of the Windows XP and Windows Server 2003 family do not support NWLink.

When used in combination with a redirector, NWLink can also function as a protocol that connects computers running the following Microsoft operating systems, if they are also running IPX/SPX:

- Windows 2000 Server
- Windows XP Professional
- Windows 2000 Professional
- Windows NT
- Windows for Workgroups 3.11
- Windows 98
- Windows 95
- Microsoft MS-DOS

**NWLink Architecture**

The following figure shows NWLink in the Windows Server 2003 architecture and the files in which each of the NWLink protocols are implemented.

**NWLink in Windows Server 2003 Architecture**




**NWLink Protocols**

NWLink provides a comprehensive set of transport and network layer protocols that enable integration with the NetWare environment. The following table shows NWLink protocols, defines their functions, and lists the associated drivers.

**NWLink Protocols**

| Protocol | Function | Driver |
|---|---|---|
| IPX | Provides connectionless datagram transfer services. | Nwlnkipx.sys |
| SPX and SPXII | Provides connection-oriented transfer services. | Nwlnkspx.sys |
| RIP | Provides route and router discovery services. | Nwlnkipx.sys |
| SAP | Collects and distributes service names and addresses. | Nwlnkipx.sys |
| NetBIOS | Provides compatible support with NetBIOS, for IPX/SPX on NetWare servers. | Nwlnknb.sys |

**IPX**

IPX is a peer-to-peer networking protocol that provides connectionless datagram transfer services and controls addressing and routing of packets of data within, and between, network segments. With connectionless transmission, it is unnecessary to set up a session each time packets are to be transmittedpackets are simply sent out on the network media. This requires less overhead than connection-oriented transmission, where it is necessary to establish a session each time packets are

transmitted. Connectionless transmission works best when data is generated intermittently, in short bursts.

Because IPX is a connectionless protocol, it does not provide for flow control or acknowledgment when the destination node receives the datagram packet. Instead, individual datagram packets travel independently to their destination. There is no guarantee that packets arrive at their destination or that they arrive in sequence. However, because transmission on local area networks (LANs) is relatively error-free, IPX is efficient in delivering short-burst data on LANs.

NWLink enables application programming for Winsock and remote procedure calls (RPCs) over Windows Sockets. IPX supports Windows Socket IDs for use by Windows Sockets applications. IPX enables NetBIOS, named pipes, mailslot, network dynamic data exchange (network DDE), RPC over NetBIOS, and RPC over named pipes. Through direct hosting, NWLink also supports other applications that use IPX. Direct hosting is a feature that enables computers to communicate over IPX while bypassing the NetBIOS layer. Direct hosting can lower overhead and increase throughput.

**IPX Frame Structure**

An IPX packet is encapsulated within the data field of an Institute of Electrical and Electronic Engineers (IEEE) frame, such as Ethernet, Token Ring, or Fiber Distributed Data Interface (FDDI), and immediately follows the media and data-link layer headers. The following figure shows the structure of an IPX header. The first 30 bytes of an IPX frame contain the header information while the remaining bytes contain the actual IPX packet data. For example, the packet data might signify a client request for a service or the data might be a server response.

**IPX Packet Header Structure for IPX**




**SPX**

SPX is a transport protocol that offers connection-oriented services over IPX. Although connection-oriented services require overhead for session setup, when a session is established, connection-oriented services require less overhead for data transmission than connectionless services. SPX works best for applications that require a continuous connection. SPX provides reliable delivery through sequencing and acknowledgments and also verifies successful packet delivery to any network destination by requesting verification from the destination node upon receipt of the data. SPX can track data transmissions consisting of a series of separate packets. If an acknowledgment request brings no response within a specified time, SPX retransmits the request eight times. If no response is received, SPX assumes the connection has failed.

**SPXII**

SPXII improves upon SPX in the following ways:

- **SPXII allows more outstanding unacknowledged packets than SPX.** SPX permits only one outstanding unacknowledged packet at a time, while SPXII allows as many outstanding packets as negotiated by the networked peers at connection setup time.

- **SPXII allows larger packets.** SPX has a maximum packet size of 576 bytes, while SPXII can use the maximum packet size of the underlying LAN. For example, SPXII can use a packet size of 1518 bytes on an Ethernet network.

- **SPXII provides a packet burst mechanism.** Packet burst, also known as burst mode, allows the transfer of multiple data packets without requiring that each packet be sequenced and acknowledged individually. By allowing multiple packets to be acknowledged once, burst mode can reduce network traffic on most IPX networks. Additionally, the packet burst mechanism monitors for dropped packets and retransmits only the missing packets.

**Note**

- In Windows Server 2003, burst mode is enabled by default.

**RIP**

NWLink uses Routing Information Protocol (RIP) over IPX (RIPX) to implement route and router discovery services used by SPX and NetBIOS/IPX (NBIPX). RIP runs on a layer equivalent to the OSI Application layer and determines the forwarding media access control (MAC) address for outbound traffic. RIP code is implemented within the file Nwlnkipx.sys.

NWLink includes the RIP protocol for Windows-based clients and for computers running Windows Server 2003 that do not have the Routing and Remote Access service installed. These computers do not forward packets as routers do, but rather use a RIP table to determine where to send packets. RIP clients, such as workstations, can locate the optimal route to an IPX network number by broadcasting a RIP GetLocalTarget route request. Each router that can reach the destination responds to the GetLocalTarget route request with a single route. Based on the RIP responses from the local routers, the sending node chooses the best router to forward the IPX packet. However, clients using the RIP protocol without Routing and Remote Access do not forward packets.

The RIPX packet header immediately follows the IPX packet header. The following figure shows the packet header for RIPX.

**IPX Packet Header Structure for RIPX**




**SAP**

Service Advertising Protocol (SAP) is the protocol that routers use to distribute the names and addresses of services running on IPX nodes. When bindery-based or NDS queries fail, SAP clients use SAP broadcasts to send the following types of messages:

- SAP GetNearestServer request for the name and address of the nearest server of a specific type.
- SAP general service request for the names and addresses of all services or of all services of a specific type.

NWLink includes a subset of the SAP protocol for Windows-based clients and for computers running Windows Server 2003 that do not have an IPX router installed.

The SAP header immediately follows the IPX header. The following figure shows the SAP header.

**IPX Packet Header Structure for SAP**




**NetBIOS over IPX**

To facilitate the operation of NetBIOS-based applications on an IPX network, NetBIOS over IPX provides the following standard NetBIOS services:

- **Datagrams.** Single packets sent without acknowledgment, for example, broadcasts.
- **Sessions.** Multiple packets sent with acknowledgments between two endpoints.
- **Name management.** Registering, querying, and releasing NetBIOS names.

The following figure shows the packet structure for NetBIOS over IPX.

**IPX Packet Header Structure for NetBIOS over IPX**




The module NWLnkNB.sys formats NetBIOS-level requests and passes them to NWLink for transmission over the network. NWLnkNB.sys includes the following performance enhancements:

- **PiggyBackAck.** A feature that facilitates acknowledgment of previous frames within the response frames.

- **Sliding window acknowledgment mechanism.** A dynamic window–sizing algorithm that enables burst mode to adjust the number of frames it can send.

As an alternative to NetBIOS over IPX, computers can communicate directly over IPX through direct hosting.

**Protocol Priority Setting**

When multiple transport protocols exist on a computer's network adapter, Windows Server 2003 negotiates network connections in the order that the network services binding list prioritizes the protocols. For example, if TCP/IP is ranked at the top of the network services binding list, Windows Server 2003 attempts to make network connections with TCP/IP before attempting to use other transport protocols. If users require network connections to servers using TCP/IP, setting the highest priority to this protocol provides the best overall performance. However, if users need connections to NetWare servers using IPX/SPX, you can improve overall performance for those users by changing the protocol priorities in their network services binding list, so that Windows Server 2003 attempts to make network connections with IPX/SPX before using the other protocols.

**NWLink Forwarder**

The forwarder component provides IPX router support. The name of the forwarder component driver is **Nwlnkfwd.sys**. This component is installed with NWLink but is used only when a computer running Windows Server 2003 is acting as an IPX router running the Routing and Remote Access service. This component operates in kernel mode.

When the IPX router software is activated, the forwarder component works with the IPX Router Manager and the filtering component to forward packets. The forwarder component obtains configuration information from the IPX Router Manager and stores a table of the best routes. When the forwarder component receives an incoming packet, it passes the packet to the filtering driver, which checks for input filters. When the component receives an outgoing packet, it passes the packet to the filtering driver, and if no outgoing filters prevent the packet from being transmitted, the filtering component passes the packet back to the forwarder component. The forwarder component then forwards the packet over the appropriate interface.

**Frame Types and Network Numbers**

NWLink uses frame types and network numbers to communicate with other computers on the same network segment and to provide correct packet routing.

**Determining Frame Type and Network Number with NWLink Auto Detect**

To connect a computer running Windows Server 2003 to a computer running IPX/SPX, a network number and frame type must be configured for each computer. The frame type and network number must be identical to the frame type and network number used on the local network segment. The NWLink Auto Detect feature automatically detects the internal network number and frame type. The Auto Detect feature, which is the recommended method for configuring network number and frame type, works in the following manner:

- When the NWLink protocol is initialized, it sends a RIPX request using a specific frame type. The RIPX request is a broadcast request for the local network. If a response is not received, NWLink sends additional requests.

- When a response to a RIPX request is received, the frame type for NWLink is set to the frame type of the response, and the IPX network number is set to the value of the Source Network number in the IPX header of the RIP response.

  **Note**

- If a computer has multiple network adapters attached to different networks, such as Token Ring, FDDI, and Ethernet, you can run Auto Detect for each adapter.

- If there are multiple RIPX responses containing multiple network numbers, Auto Detect uses a counting algorithm to determine the most likely network number.

- If there is no RIPX response to any request, NWLink sets the frame type to Ethernet 802.2 (for Ethernet network adapters) and the network number to 0.

AutoDetect uses the RIPX response to select the frame type and network number. If a network computer has an incorrect manual setting, AutoDetect will select an inappropriate frame type and network number combination for the network adapter. For example, if a computer replies with an incorrect frame type and network number, Auto Detect detects the incorrect configuration and uses it. Auto Detect might also select an incorrect network number if multiple computers respond. In that case, Auto Detect uses a heuristic algorithm to determine the most likely network number, but Auto Detect does not validate the packet information that it receives.

**Understanding Network Numbers**
You can configure NWLink with two distinctly different types of network numbers for routing purposes: the external network number and the internal network number. NWLink uses an IPX external network number to identify the computer's local network segment within the routed environment. In addition, NWLink uses a virtual network number, also known as an internal network number, to identify a logical network inside the computer. The purpose of the logical network inside the computer is to provide optimal routing to services running on the computer.

In Windows Server 2003 you can change both the external network number and internal network number from Control Panel.

**External network number**
The external network number is associated with the physical network adapter and the local area network segment (a segment in a NetWare network is analogous to a TCP/IP subnet). The external network number identifies the computer's local segment and facilitates communications within the routed environment. Accordingly, packets transmitted across the network contain the computer's external network number. Otherwise, if an internal number exists, the packets must contain the internal network number along with an appended MAC address to identify individual computers. All computers on the same network segment that use a given frame type must have the same external network number that must also be unique for each network segment.

If you replace a network adapter or connect the computer to a different segment, or if the external network number on a segment changes, you might not be able to communicate over IPX. If you lose IPX connectivity, you must reset the external network number and reboot the computer, when prompted to do so.

- To determine the external network number and frame type that have been set on a computer running Windows Server 2003, run the **ipxroute config** command at the command-line interface of that computer.

- To determine the external network number and frame type that have been set on a server running NetWare, enter **config** at the command-line interface of the NetWare server or inspect the Autoexec.ncf file.

### Note

- It is possible to run multiple IPX networks over the same physical network segment if those networks use different frame types.

**Internal Network Number**
Internal network numbers are necessary for internal routing purposes. Service-supporting hosts, which are hosts running services that SAP-based clients need to access, use internal network numbers. Service-supporting hosts use these numbers to help clients determine the optimal route for transmitting packets to services running on the service-supporting hosts. If only external network numbers are used, different routers might send a client multiple routes with the same route metrics, when only one route is actually optimal. However, if you specify an internal network number, you create a virtual network inside the service-supporting host and packets are always forwarded to the services running on the service-supporting host using an optimal path.

By default, NWLink does not identify internal network numbers. For example, when NWLink is installed, the Auto Detect feature does not detect the internal network number automatically. Instead, NWLink sets an internal network number of

00000000 on the computer's network adapter. When set to 00000000, the internal network number is not in use for that computer. This default behavior works whether NWLink is bound to a single network adapter or to multiple network adapters. However, an internal network number is required if the computer running Windows Server 2003 is acting as a server for applications that use NetWare SAP. This includes applications such as SQL Server, Exchange Server, SNA Server, or Host Integration Server (HIS) 2000.

Depending on your configuration, you might already have an internal network number. The Routing and Remote Access service and the SAP agent automatically install a unique internal network number during setup procedures. If you previously installed any of these services, you might already have a unique internal network number set for your computer's network adapter.

The default value of zero (00000000) for the internal network number is not returned when you type **ipxroute config** at the command-line interface. However, if you configure an internal network number, you can view the number using either **ipxroute config**, or by monitoring a trace of a RIP or SAP broadcast from the computer where the number is set.

**Supporting IPX Frame Types with NWLink**

Both Windows Server 2003 and NetWare support frame types for Ethernet, Token Ring, and FDDI topologies. The following table lists the frame types that Windows Server 2003 supports:

**Frame Types Supported by Windows Server 2003**

| Network Type | Supported Frame Types |
|---|---|
| Ethernet | Ethernet II, 802.2, 802.3, Ethernet Subnetwork Access Protocol (SNAP) |
| Token Ring | 802.5 and 802.5 SNAP |
| FDDI | 802.2, 802.3, and SNAP |

Frame types are based upon IEEE standards and define packet formats used by various topologies.

**Ethernet II**

The Ethernet II frame includes a 2-byte Type field that immediately follows the source address. The Type field contains a unique value, called an **EtherType**, which identifies the receiving computer for which the upper-layer protocol handles Data field contents.

**Ethernet 802.2**

Several frame types are based on the IEEE 802.3 standards, which define operations for the physical layer and the MAC sublayer. NetWare combines the IEEE 802.3 standard for the physical layer and the IEEE 802.2 standard for operations of the Logical Link Control (LLC) sublayer and refers to this combination as the Ethernet 802.2 frame type.

Ethernet II and Ethernet 802.2 frames can run on the same network. Because the value in the Ethernet II Type field always equals a number greater than 1,500 and the value in the Ethernet 802.2 Length field always equals a number of 1,500 or less, network hardware and software can distinguish between these frame types. This makes it possible for both frame types to coexist on the same network.

**Subnetwork Access Protocol**

The Subnetwork Access Protocol (SNAP) frame type is derived from the Ethernet 802.2 frame type. The SNAP portion of the frame is carried in the Data field. This frame type gives individual vendors the ability to assign unique values for protocols running on their hardware.

**Ethernet 802.3**

Like the 802.2 frame, the 802.3 frame differs from the Ethernet II frame because the 802.3 frame uses a Length field in place of the Type field in Ethernet II. It is often called 802.3 raw because it does not use the 802.2 LLC header in the Data field.

**Token Ring 802.5**

The Token Ring protocol is composed of two distinct frame subtypes, the MAC frames and non-MAC frames:

- **MAC frames.** These transmit management information for the Token Ring network. MAC frames carry signaling information and are not passed on by bridges and routers. They reside only on the LAN ring segment where the packets originate.

- **Non-MAC frames.** These transmit data for Token Ring networks. Non-MAC frames transmit data between computer stations on the ring and are passed on by bridges and routers.

**Note**

- The content of fields at the beginning of a Token Ring frame can redefine the composition of fields later in the frame.

**Related Information**

The following resources contain additional information that is relevant to this section.

- What Is Services for Netware
- Services for Netware Tools and Settings

# Services for Netware Tools and Settings

Updated: June 7, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

### In this section


- Services for NetWare Tools
- Services for NetWare Registry Entries


The "Services for NetWare Tools" section describes tools that are useful for troubleshooting Services for NetWare.

The "Services for NetWare Registry Entries" section describes the registry entries that can be configured to use the advanced features of MSDSS, such as modifying wakeup interval values and setting the maximum number of failed objects.

**Services for NetWare Tools**
Windows Server 2003 includes several tools that allow you to determine computer configuration settings and perform diagnostic tests to resolve communication problems in an IPX/SPX network environment used for Services for NetWare 5.03.

**Dsa.msc: Active Directory Users and Computers**
**Category**
An Active Directory Administrative Tools MMC snap-in that is automatically installed on all Windows Server 2003 domain controllers running Windows Server 2003.

### Version compatibility

| Can Be Run From | Can Be Run Against |
|---|---|
| Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br><br>Servers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition<br><br>Computers running:<br>• Windows XP Professional with Adminpak.msi installed | Domain controllers running:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows 2000 Server<br>• Windows 2000 Advanced Server<br>• Windows 2000 Datacenter Server |

Active Directory Users and Computers provides a graphical user interface that can be used to manage users and computers in Active Directory domains.

**Ipxroute**
**Category**
**Ipxroute** can be used to display and modify information about the routing tables used by the IPX protocol.

**Version compatibility**

**Ipxroute** is available in Windows Server 2003, Windows 2000 Server, and Windows XP Professional.

Ipxroute enables you to troubleshoot IPX connectivity problems and provides information about the current state of the stack. Ipxroute also displays information such as current IPX status, network number, MAC address, interface name, frame type, and so on.

For example, to display information at the command-line interface for all bindings on which IPX is configured, use the ipxroute command with the **config** argument, as follows:

```
C:\>ipxroute config
```

To determine if there is connectivity to a specific network, use the **ipxroute** command with the **ripout** argument as follows:

```
C:\>ipxroute ripout <network number>
```

where <*network number*> is the 8-digit hexadecimal network number.

You can also view additional information about **ipxroute** options, which might assist you in troubleshooting connectivity problems. To do this, use the help switch for **ipxroute**, as follows:

```
C:\>ipxroute /?
```

**Network Monitor**
**Category**
**Network Monitor** is a utility for monitoring, recording, and analyzing network traffic.

**Version compatibility**
Network Monitor is part of the Microsoft Systems Management Server (SMS) product. A limited version, which can capture frames that are sent to or from the network adapter of the computer on which Network Monitor is installed, is included with Windows Server 2003 and can be installed as an optional networking component. You can use NetMon to identify network traffic patterns and to perform other tasks associated with finding network problems, such as the following:

- Locating client-to-server connection problems.
- Identifying computers that make a disproportionate number of service requests.
- Capturing frames (packets) directly from the network.
- Displaying and filtering captured frames.
- Identifying unauthorized users on a network.

**Services for NetWare Registry Entries**
The following registry entries are associated with Services for NetWare.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**Note**

- For information about additional Services for NetWare registry entries for MSDSS and FMU, see the Services for Netware 5.03 online Help.

**HKEY_LOCAL_MACHINE\SYSTEM**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM CurrentControlSet\Services\MSDSS\Parameters.

**UseMemberAttribute**
**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\ CurrentControlSet\Services\MSDSS\Parameters

**Version**
The UseMemberAttribute registry setting is available in Microsoft Services for NetWare 5.03 on Windows Server 2003.

Set this value to 1 to migrate group "member"ships using the "member" attribute of the group object. If you set the value to 0, you use the "security equal to me" attribute.

When migrating groups from NDS to Active Directory, MSDSS reads the "security equal to me" attribute by default and migrates those group members to Active Directory. Some NetWare administration tools (particularly those that operate in Bindery mode) add "member"s only to the "member" attribute. This causes a discrepancy between the group memberships displayed in NDS, eDirectory, and Active Directory after some migrations. If you experience this, set the value to 1 to allow NDS and eDirectory group members that are not marked as "security equal to me" to be migrated to Active Directory.

**SyncIntervalForward or SyncIntervalReverse**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSDSS\Parameters

**Version**
The SyncIntervalForward and SyncIntervalReverse registry settings are available in Microsoft Services for NetWare 5.03 running on Windows Server 2003.

To modify the wakeup interval values, you can add the DWORD values SyncIntervalForward and SyncIntervalReverse to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSDSS\Parameters.

After these values have been added, you can indicate the wakeup interval time in seconds.

**MaxFailedObjects**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSDSS\Parameters

**Version**
The MaxFailedObjects registry setting is available in Microsoft Services for NetWare 5.03 running on Windows Server 2003.

To modify the default value of objects that cannot be retried, you can add the DWORD value MaxFailedObjects to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSDSS\Parameters.

A value of 0 in MaxFailedObjects means there is no limit for the number of failed objects.

**BackupInterval**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSDSS\Parameters

**Version**
The BackupInterval registry setting is available in Microsoft Services for NetWare 5.03 running on Windows Server 2003.

To modify the default backup interval, create a DWORD value named BackupInterval and specify a custom interval, in seconds.

**SyncEmailAddress**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSDSS\Parameters

**Version**
The SyncEmailAddress registry setting is available in Microsoft Services for NetWare 5.03 running on Windows Server 2003.

You can disable the "mail" synchronization attribute by setting the DWORD value to 0 in the SyncEmailAddress registry key.

**SamRename**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\\CurrentControlSet\Services\MSDSS\Parameters

**Version**
The SamRename registry setting is available in Microsoft Services for NetWare 5.03 running on Windows Server 2003.

Active Directory, NDS, and eDirectory handle logon names differently. Active Directory provides separate attributes for logon name (samAccountName) and object naming in the directory hierarchy (that is, the relative distinguished name). NDS and eDirectory use the same attribute (that is, CN) for both purposes. The initial release of MSDSS synchronizes the directory object name (that is, the relative distinguished name maps to the CN) but does not synchronize the logon names. MSDSS in Services for NetWare 5.03 provides an option for synchronizing logon names between Active Directory, NDS, and eDirectory.

Setting the value to 1 enables MSDSS to keep both the samAccountName and relative distinguished name attributes of a user object in Active Directory synchronized with the CN attribute of the corresponding NDS or eDirectory user object.

**Note**

- The session must be a two-way synchronization session to synchronize changes in NDS or eDirectory to Active Directory.

# Remote Access Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Access Technologies**
In the Microsoft Windows Server 2003 operating system, remote access technologies allow IT professionals to perform administrative tasks and users to run programs on a remote computer as if they were working locally.

**Remote Access Technologies Components**
Dial-up Remote Access, Telnet, and Terminal Services are included in Windows Server 2003 as remote access technologies.

**Dial-up Remote Access**
Dial-up Remote Access is a remote access technology that is available as part of the Routing and Remote Access services included in Microsoft Windows Server 2003. A dial-up remote access client uses the telecommunications infrastructure to create a temporary physical circuit or a virtual circuit to a port on a remote access server. If the remote access server accepts the dial-up connection from the client, the server forwards packets between the client and the network to which the remote access server is attached.

**Telnet**
Windows Server 2003 and Windows XP Professional include Telnet Client and Telnet Server components, which allow users to make remote connections based on the Telnet protocol. By using Telnet Client and Telnet Server, you can create a remote command console session on a host. Using a local command prompt, you can run command-line programs, shell commands, and scripts in a remote command console session just as though you were locally logged on to the host.

**Terminal Services**
The Terminal Services component of the Windows Server 2003 operating system delivers the Windows Server 2003 desktop and the latest Windows-based applications to a variety of desktops, including those that normally cannot run Windows.

Terminal Services provides the ability to host multiple, simultaneous client sessions on Windows Server 2003. Through terminal emulation, Terminal Services allows the same set of applications to run on diverse types of desktop hardware. For organizations wanting to more flexibly deploy applications and control desktop management costs, a Terminal Services architecture offers an important enhancement to the traditional two- or three-tier client-server architecture based on servers and full-scale personal computers.

Terminal Services also provides Remote Desktop for Administration, which can greatly reduce the overhead associated with remote administration. Enabled by Terminal Services technology, Remote Desktop for Administration is specifically designed for server management. Therefore, it does not install the application-sharing and multi-user capabilities or the process scheduling of the full Terminal Server component (formerly called Terminal Services in Application Server mode).

**Remote Access Technologies Scenarios**
In a dial-up remote access scenario, a remote access client uses the telecommunications infrastructure to create a temporary physical or virtual circuit to transparently connect the remote access client to a network. The physical or logical connection between the remote access server and the remote access client is created using available dial-up equipment. The dial-up equipment and network infrastructure can vary, depending on the type of connection used. For example, a dial-up connection can be made using the Public Switched Telephone Network (PSTN), Integrated Services Digital Network (ISDN), X.25, or Asynchronous Transfer Mode (ATM) over asymmetric digital subscriber line (ADSL).

Telnet Client and Telnet Server are well suited for troubleshooting and configuring remote computers, especially in mixed environments that require interoperability between different operating systems. For example, you can use Windows Server 2003 Telnet Client to connect to a Telnet server that is running another operating system, such as UNIX. Likewise, you can use a Telnet client that is running UNIX to connect to a computer running Windows Server 2003 Telnet Server. Telnet Client and Telnet Server are also ideal in situations in which memory and processor resources are minimal on a client or host or in which network bandwidth is limited. This is because Telnet Clients and Telnet Server use less memory and processor time than other remote management tools, and Telnet clients and servers transmit only plaintext (unencrypted characters) across the network.

Organizations are always searching for ways to reduce the costs of ownership and this is one of the major goals of Terminal Services. Terminal Services lets enterprises more easily and cost-effectively accomplish this goal by allowing organizations to:

- Centrally deploy and manage Windows-based applications.
- Remotely administer Windows Server 2003-based computers.

Terminal Services extends the model of distributed computing by allowing computers to operate in a server-based computing environment. With Terminal Services running on a Windows Server 2003-based computer, all client application execution, data processing, and data storage occur on the server. Applications and user desktops are transmitted over the network and displayed through terminal emulation software. Similarly, print streams, keyboard input, and mouse clicks are also transmitted over the network between the server and the terminal emulation software. Users log on and see only their individual session, which is managed transparently by the server operating system and is independent of any other client session.

Remote Desktop for Administration allows for the management of servers from any location without affecting server performance or application compatibility. In addition to the console session, up to two remote administration sessions are supported. Because it is meant as a single-user remote access solution, no Terminal Server Client Access License (CAL) is required to use Remote Desktop for Administration.

Administrators can also fully administer computers running Windows Server 2003 from computers running earlier versions of Windows by installing Remote Desktop Connection.

# Dial-up Remote Access Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Microsoft® Windows Server™ 2003 provides extensive support for technology that connects remote clients to corporate networks or the Internet. This technical reference describes the remote access connectivity type known as dial-up remote access.

Dial-up remote access is a remote access technology that is available as part of Routing and Remote Access that is included in Windows Server 2003. With dial-up remote access, a remote access client uses the telecommunications infrastructure to create a temporary physical or virtual circuit to a port on a remote access server, which is typically attached to a corporate network.

**In this subject**


- What Is Dial-up Remote Access?
- How Dial-up Remote Access Works
- Dial-up Remote Access Tools and Settings

# What Is Dial-up Remote Access?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Dial-up Remote Access?**

**In this section**


- Dial-up Remote Access vs. VPN
- Components of a Dial-up Remote Access Connection
- Components of Secure Remote Access
- Dial-up Remote Access Dependencies
- Related Information


Dial-up remote access is a remote access technology that is available as part of Routing and Remote Access that is included in Microsoft Windows Server 2003.

Organizations require a simple solution that allows employees to remotely access their corporate e-mail accounts and shared files from home or from other locations outside the corporate network. Dial-up remote access provides the solution by enabling a remote access client to use the wide area network (WAN) infrastructure to connect to a remote access server.

**Dial-up Remote Access vs. VPN**

Windows Server 2003 remote access provides two different types of remote access connectivity: dial-up remote access and virtual private network (VPN) remote access.

- With dial-up remote access, a remote access client uses the telecommunications infrastructure to create a temporary physical circuit or a virtual circuit to a port on a remote access server. After the physical or virtual circuit is created, the rest of the connection parameters can be negotiated.
- With virtual private network remote access, a VPN client uses an IP internetwork to create a virtual point-to-point connection with a remote access server acting as the VPN server. After the virtual point-to-point connection is created, the rest of the connection parameters can be negotiated.

**Demand-Dial Routing**

Demand-dial routing is not a technology, but a process by which packets are forwarded across a Point-to-Point Protocol (PPP) link. Demand-dial routing is used by remote access technologies—dial-up remote access and VPN, as well as routing technologies—unicast routing, multicast routing, and network address translation (NAT). For more information about demand-dial routing, see "Demand Dial Routing Technical Reference."

**Components of a Dial-up Remote Access Connection**

A dial-up remote access connection contains the following components:

- Remote access client
- Remote access server
- WAN infrastructure.

**Components of a Dial-up Remote Access Connection**




**Remote Access Client**

Remote access clients running Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0, Windows Millennium Edition, and Windows 98 can connect to a Windows Server 2003 remote access server. Almost any PPP remote access client, including UNIX and Macintosh, can connect to a Windows Server 2003 remote access server.

**Remote Access Server**

The Windows Server 2003 remote access server accepts dial-up connections and forwards packets between remote access clients and the network to which the remote access server is attached.

### Note

- The term remote access server as it is used in this section refers to a Windows Server 2003 computer running Routing and Remote Access and configured to provide remote access.

**Dial-up Equipment and the WAN Infrastructure**

The physical or logical connection between the remote access server and the remote access client is facilitated by dial-up equipment installed at the remote access client, the remote access server, and the WAN infrastructure. The nature of the dial-up equipment and WAN infrastructure varies, depending on the type of connection.

The following telecommunications technologies can make up the WAN infrastructure:

- Public Switched Telephone Network (PSTN)
- Integrated Services Digital Network (ISDN)
- X.25
- Asynchronous Transfer Mode (ATM) over asymmetric digital subscriber line (ADSL)

**PSTN**

PSTN, also known as Plain Old Telephone Service (POTS), is the most common network used for dial-up remote access. PSTN is the analog phone system designed to carry the minimum frequencies required to distinguish human voices.

Dial-up equipment consists of an analog modem at the remote access client and at least one analog modem at the remote access server. For large organizations, the remote access server is attached to a modem bank containing up to hundreds of modems. Because PSTN was not designed for data transmission, there are limits to the maximum bit rate of a PSTN connection. The maximum bit rate supported by PSTN connections is 33,600 bits per second or 33.6 kilobits per second (Kbps).

### Standard PSTN Connection




**Digital Links and V.90**

The maximum bit rate of PSTN depends on the range of frequencies being passed by PSTN switches and the signal-to-noise ratio of the connection. The modern-day analog phone system is only analog on the local loop, the set of wires that connect the customer to the central office (CO) PSTN switch. After the analog signal reaches the PSTN switch, it is converted to a digital signal. The analog-to-digital conversion introduces noise on the connection known as quantization noise.

When a remote access server is connected to a CO by using a digital switch based on T-Carrier or ISDN rather than an analog PSTN switch, there is no analog-to-digital conversion when the remote access server sends information to the remote access client. Because there is no quantization noise in the path back to the remote access client, there is a higher signal-to-noise ratio and, therefore, a higher maximum bit rate.

With this technology, called V.90, remote access clients can send data at 33.6 Kbps and receive data at 56 Kbps. In North America, the maximum receive bit rate is 53 Kbps due to Federal Communications Commission (FCC) power rules.

To obtain V.90 speeds, the following must be true:

- The remote access client must be using a V.90 modem.
- The remote access server must be using a V.90 digital switch and be connected to PSTN using a digital link, such as T-Carrier or ISDN.
- There cannot be any analog-to-digital conversions in the path from the remote access server to the remote access client.

**PSTN Connection with V.90**




**ISDN**

ISDN is a set of international specifications for a digital replacement of PSTN. ISDN provides a single digital network to handle voice, data, fax, and other services over existing local loop wiring. ISDN behaves like an analog phone line except that it is a digital technology with higher data rates and a much lower connection time. ISDN offers multiple channels, with each channel operating at 64 Kbps. Because the network is digital from end to end, there are no analog-to-digital conversions.

Dial-up equipment consists of an ISDN adapter for the remote access client and the remote access server. Remote access clients typically use Basic Rate ISDN (BRI) with two 64-Kbps channels; large organizations typically use Primary Rate ISDN (PRI) with 23 64-Kbps channels.

**ISDN Connection**




**X.25**

X.25 is an international standard for sending data across public packet-switching networks. Windows Server 2003 remote access supports X.25 in the following ways:

- The remote access client supports the use of X.25 smart cards, which can connect directly to the X.25 data network and use the X.25 protocol to establish connections and send and receive data. The remote access client also supports dialing into a packet assembler and disassembler (PAD) of an X.25 carrier using an analog modem.
- A remote access server only supports direct connections to X.25 networks by using an X.25 smart card.

**Note**

- X.25 smart cards are adapters that use the X.25 protocol and can directly connect to an X.25 public data network. X.25 smart cards are not related to smart cards used for authentication and secure communications. For more information about smart cards used for authentication, see "IAS Technical Reference."

**X.25 Connection**




**ATM over ADSL**

ATM is a fixed-length, cell-based, packet-switching technology that transmits data across an ATM network. Routing and Remote Access supports PPP connections over switched virtual circuit (SVC) or permanent virtual circuit (PVC) ATM connections. PPP connections over ATM require an ATM adapter with support for the physical connection to the ATM network. After it is installed, the ATM adapter appears as a dial-up device.

ADSL is a local loop technology for small business and residential customers. Although ADSL provides higher bit rates than PSTN and ISDN connections, the bit rate is not the same in the upstream and downstream directions. Typical ADSL connections offer 64 Kbps from the customer and 1.544 megabits per second (Mbps) to the customer. The asymmetric nature of the connection fits well with typical Internet use. Most Internet users receive a lot more information than they send.

ADSL equipment can appear to Windows Server 2003 as either an Ethernet interface or a dial-up interface. When an ADSL adapter appears as an Ethernet interface, the ADSL connection operates in the same way as an Ethernet connection to the Internet.

When an ADSL adapter appears as a dial-up device, ADSL provides a physical connection and the individual local area network (LAN) protocol packets are sent using ATM. An ATM adapter with an ADSL port is installed in both the remote access client and remote access server.

### ATM over ADSL Connection




**Remote Access Protocols**

Windows Server 2003 supports the following remote access protocols:

- PPP
- SLIP

**PPP**

PPP is an industry standard method of using point-to-point links to transport multi-protocol datagrams. A PPP-enabled connection can dial into remote networks through any industry-standard PPP server. PPP also permits a remote access server to receive calls from, and provide network access to, other vendors' remote access software that complies with the PPP standards. For more information about PPP, see "How Dial-up Remote Access Works." PPP is documented in RFC 1661.

**SLIP**

Serial Line Internet Protocol (SLIP) is an older remote access standard typically used by UNIX remote access servers. Routing and Remote Access supports SLIP for dial-up connections. Support for SLIP is included in the Windows remote access client to ensure interoperability with other remote access software. For more information about SLIP, see RFC 1055, A Nonstandard for Transmission of IP Datagrams Over Serial Lines: SLIP and RFC 1144, "Compressing TCP/IP Headers for Low-Speed Serial Links."

### Notes

- Windows Server 2003 and Microsoft Windows XP do not support SLIP for incoming connections.
- You must use the TCP/IP protocol and a serial COM port to connect to a SLIP server.

**LAN Protocols**
LAN protocols are the protocols used by the remote access client to access resources on the network connected to the remote access server. Windows Server 2003 remote access supports TCP/IP and AppleTalk. For more information about these protocols, see "How Dial-up Remote Access Works."

**Components of Secure Remote Access**
Because remote access is designed to transparently connect a remote access client to a network and its potentially sensitive data, the security of remote access connections is an important consideration. Windows Server 2003 remote access offers security features, such as secure user authentication, mutual authentication, data encryption, callback, caller ID, and smart card authentication. For more information about smart card authentication, see "IAS Technical Reference."

**Authentication vs. Authorization**
To understand why connection attempts are either accepted or denied, it is important to understand the distinction between authentication and authorization:

- Authentication is the verification of the credentials of the connection attempt. This process consists of sending the credentials from the remote access client to the remote access server in either plain text or an encrypted form by using an authentication protocol.

- Authorization is the determination that the connection attempt is allowed. Authorization occurs after successful authentication.

To be accepted, the connection attempt must be both authenticated and authorized. It is possible for the connection attempt to be authenticated by using valid credentials, but not authorized. In this case, the connection attempt is denied.

If a remote access server is configured for Windows Authentication, the security features of Windows Server 2003 are used to verify the credentials for authentication, and the dial-in properties of the user account and locally stored remote access policies are used to authorize the connection. If the connection attempt is both authenticated and authorized, the connection attempt is accepted. For more information, see "Remote Access Policies" later in this section.

If the remote access server is configured for Remote Authentication Dial-In User Service (RADIUS) authentication, the credentials of the connection attempt are passed to the RADIUS server for authentication and authorization. If the connection attempt is both authenticated and authorized, the RADIUS server sends an accept message back to the remote access server and the connection attempt is accepted. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends a reject message back to the remote access server and the connection attempt is rejected.

If the RADIUS server is a computer running Internet Authentication Service (IAS), the IAS server performs authentication through selected authentication features and authorization through the dial-in properties of the user account and remote access policies stored on the IAS server.

**Secure User Authentication**
Secure user authentication is obtained through the encrypted exchange of user credentials. This is possible with the PPP remote access protocol using either the Challenge Handshake Authentication Protocol (CHAP), Shiva Password Authentication Protocol (SPAP), Microsoft Challenge Handshake Authentication Protocol version 1 (MS-CHAP), Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2), or Extensible Authentication Protocol (EAP) authentication protocols. MS-CHAP, MS-CHAP v2, and EAP offer the most security. The remote access server can be configured to require a secure authentication method; if the remote access client cannot perform the required secure authentication, the connection is denied.

For more information about PPP authentication protocols, see "PPP Authentication Protocols" in How Dial-up Remote Access Works.

**Secure Authentication Scheme**
A secure authentication scheme provides protection against replay attacks, remote access client impersonation, and remote access server impersonation.

- A replay attack occurs when a person captures the packets of a successful connection attempt and then replays those packets in an attempt to obtain an authenticated connection.

- Remote access client impersonation occurs when a person takes over an existing authenticated connection. The intruder waits until the connection is authenticated and then obtains the connection parameters, disconnects the user, and takes control of the authenticated connection.

- Remote server impersonation occurs when a computer appears as the remote access server to the remote access client.

The impersonator appears to verify the remote access client credentials and then captures all of the traffic from the remote access client.

**Mutual Authentication**

Mutual authentication is obtained by authenticating both ends of the connection through the encrypted exchange of user credentials. This is possible with the PPP remote access protocol using either the EAP-Transport Level Security (EAP-TLS) or MS-CHAP v2 authentication protocols. If smart cards are used for remote access authentication, EAP-TLS is the required authentication method. During mutual authentication, the remote access client authenticates itself to the remote access server, and then the remote access server authenticates itself to the remote access client.

For more information about smart card authentication, see "IAS Technical Reference."

A remote access server does not always request authentication from the remote access client. However, in the case of a Windows Server 2003 remote access client configured for MS-CHAP v2 only or EAP-TLS only, the remote access client will force the mutual authentication of the client and server. If the remote access server does not respond to the authentication request, the connection is terminated by the client.

For more information about EAP-TLS, see "How Dial-up Remote Access Works."

**Data Encryption**

Data encryption protects the data sent between the remote access client and the remote access server. Remote access data encryption does not provide end-to-end data encryption. End-to-end encryption is data encryption between the client application and the server that hosts the resource or service being accessed by the client application. If end-to-end encryption is required, use Internet Protocol security (IPSec) to create an encrypted end-to-end connection after the remote access connection has been made.

### Note

- IPSec can also be used for encrypting a Layer Two Tunneling Protocol (L2TP) virtual private network connection. For more information, see "VPN Technical Reference."

Data encryption on a remote access connection is based on a secret encryption key known to the remote access server and remote access client. This shared secret key is generated during the user authentication process.

Data encryption is possible over dial-up remote access links when using the PPP remote access protocol and the EAP-TLS or MS-CHAP authentication protocols.

The remote access server can be configured to require data encryption. The remote access client can be configured to request the following levels of data encryption:

- No encryption (the server disconnects if it requires encryption).
- Optional encryption (connect even if there is no encryption).
- Require encryption (disconnect if the server declines).
- Maximum-strength encryption (disconnect if the server declines).

If the remote access client cannot perform the required encryption, the connection attempt is rejected.

Remote access clients and remote access servers running Windows Server 2003, Windows 2000, Windows NT 4.0, Windows 98, and Windows 95 use the Microsoft Point-to-Point Encryption (MPPE) Protocol. MPPE uses the Rivest-Shamir-Adleman (RSA) RC4 stream cipher and either 40-bit, 56-bit, or 128-bit secret keys. MPPE keys are generated from the MS-CHAP and EAP-TLS user authentication process.

**Callback**

With callback, the remote access server calls the remote access client after the user credentials have been verified. Callback is configured as part of the dial-in properties of the user account. Callback can also be configured on the server to call the remote access client back at a number specified by the user during the time of the call. A traveling user can reduce phone charges by dialing in and having the remote access server call back at the user's current location. Callback can also be configured to always call the remote access client back at a specific location, which is the secure form of callback.

**Caller ID**
Caller ID can be used to verify that the incoming call is coming from a specified phone number. Caller ID is configured as part of the dial-in properties of the user account. If the Caller ID number of the incoming connection for that user does not match the configured Caller ID, the connection is denied.

Caller ID requires that the caller's phone line, the phone system, the remote access server's phone line, and the Windows Server 2003 driver for the dial-up equipment all support Caller ID. If a Caller ID is configured for a user account and the Caller ID is not being passed from the caller to Routing and Remote Access, then the connection is denied.

Caller ID is a feature designed to provide a higher degree of security for networks that support telecommuters. The disadvantage of configuring Caller ID is that the user can dial in from a single phone line only.

**Secure Remote Access and IAS**
Internet Authentication Service (IAS) is the Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy. RADIUS is an industry standard for the client-server protocol. For more information about RADIUS, see RFC 2865, "Remote Authentication Dial-In User Service (RADIUS)" and RFC 2866, "RADIUS Accounting."

As a RADIUS server, IAS performs centralized connection authentication, authorization, accounting, and auditing (AAAA) for network access types that include dial-up remote access, VPN, wireless, and authenticating switch connections.

**Network Access Quarantine Control**
IAS Network Access Quarantine Control provides phased network access, which restricts remote clients to quarantine mode until each client is either verified as meeting, or configured according to, organization network access policy. Quarantine control can be applied when:

- The user connects directly to a dial-up remote access server on a private network.
- The user connects to an Internet Service Provider (ISP) through a modem, and then connects to a private network through a VPN connection over the Internet.

To implement Network Access Quarantine Control, you must use Connection Manager Administration Kit (CMAK), Routing and Remote Access, and additional components that include servers that perform name resolution (such as Domain Name System [DNS] servers), obtain the latest version of the Connection Manager profile, or access instructions and components required to make the remote access client comply with network policies.

For more information about Network Access Quarantine Control, see Network Access Quarantine Control in Windows Server 2003.

**Remote Access Policies**
In Windows Server 2003, remote access connections are accepted based on the dial-in properties of a user account and remote access policies. A remote access policy is a set of conditions and connection parameters that define the characteristics of the incoming connection and the set of constraints imposed on it. Remote access policies can be used to specify allowed connections conditioned by the time of day and day of the week, the Windows Server 2003 group to which the dial-in user belongs, the type of remote access client (dial-up or VPN), and so on. Remote access policies can be used to impose connection parameters such as maximum session time, idle disconnect time, required secure authentication methods, required encryption, and so on.

For information about how the processing of connection attempts involves remote access policies, see "How Dial-up Remote Access Works."

With multiple remote access policies, different sets of conditions can be applied to different remote access clients or different requirements can be applied to the same remote access client based on the parameters of the connection attempt. For example, multiple remote access policies can be used to:

- Allow or deny connections if the user account belongs to a specified group.
- Define different days and times for different user accounts based on group membership.
- Configure different authentication methods for dial-up and VPN remote access clients.
- Configure different authentication or encryption settings for Point-to-Point Tunneling Protocol (PPTP) or L2TP connections.

- Configure different maximum session times for different user accounts based on group membership.
- Send network access server-specific RADIUS attributes to a RADIUS client.
- When you have multiple Windows Server 2003 remote access or VPN servers and you want all of the servers to use a centralized set of remote access policies to authorize incoming connections, you must configure a computer to run Windows Server 2003 and IAS and then configure each remote access or VPN server as a RADIUS client to the IAS server computer.

For more information about remote access policies, see "IAS Technical Reference."

**Unauthenticated Connections**
Windows Server 2003 also supports unauthenticated PPP connections. In an unauthenticated PPP connection, the authentication phase of the PPP connection establishment is skipped. Neither the remote access client nor the remote access server exchange credentials. The use of unauthenticated PPP connections must be carefully considered because connections are allowed without verification of the identity of the remote access client.

There are two common cases in which unauthenticated connections are desired:

- When using Automatic Number Identification/Calling Line Identification (ANI/CLI) authentication, the authentication of a connection attempt is based on the phone number of the caller. ANI/CLI service returns the number of the caller to the receiver of the call and is provided by most standard telephone companies.

  ANI/CLI authentication is different from caller ID authorization. In caller ID authorization, the caller sends a valid user name and password. The caller ID that is configured for the dial-in property on the user account must match the connection attempt; otherwise, the connection attempt is rejected. In ANI/CLI authentication, a user name and password are not sent.

- When using guest authentication, the Guest account is used as the identity of the caller.

**Dial-up Remote Access Dependencies**
Dial-up remote access requires one of the following operating systems for the remote access server:

- Windows 2000 Server
- Windows Server 2003
- Windows XP (appropriately configured for incoming connections)

Dial-up remote access requires one of the following operating systems for the remote access client:

- Windows 98
- Windows Millennium Edition
- Windows 2000
- Windows 2000 Server
- Windows XP
- Windows Server 2003

Dial-up remote access requires the following protocols and interfaces:

- PPP

- Either TCP/IP or AppleTalk

**Related Information**

The following resources contain additional information that is relevant to this section.

- VPN Technical Reference
- IAS Technical Reference
- Demand Dial Routing Technical Reference
- Connection Manager Technical Reference
- RFC 1055, RFC 1144, RFC 1661, RFC 2865, and RFC 2866 in the IETF RFC Database

# How Dial-up Remote Access Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Dial-up Remote Access Works**

**In this section**


- Remote Access Server Architecture
- Dial-up Remote Access Protocols
- Dial-up Remote Access Processes and Interactions
- Related Information


Dial-up remote access is a remote access technology that is available as part of Routing and Remote Access that is included in Windows Server 2003. With dial-up remote access, a remote access client uses the telecommunications infrastructure to create a temporary physical or virtual circuit to a port on a remote access server, which is typically attached to a corporate network.

The physical or logical connection between the remote access server and the remote access client is made possible through dial-up equipment installed at the remote access client, the remote access server, and the wide area network (WAN) infrastructure. The nature of the dial-up equipment and WAN infrastructure varies, depending on the type of connection.

The following telecommunications technologies can make up the WAN infrastructure:

- Public Switched Telephone Network (PSTN)
- Integrated Services Digital Network (ISDN)
- X.25
- Asynchronous Transfer Mode (ATM) over asymmetric digital subscriber line (ADSL)

For more information about dial-up remote access configurations with each of these telecommunications technologies, see "What Is Dial-up Remote Access?."

**Remote Access Server Architecture**

To understand how dial-up remote access works in Windows Server 2003, you must first understand the architecture of the remote access server. The following figure illustrates how the remote access server components work together.

**Remote Access Server Architecture**




The following table describes the remote access server components.

**Remote Access Server Components**

| Component | Description |
|---|---|

| | |
|---|---|
| Ndis.sys | The Network Driver Interface Specification (NDIS) wrapper provides the NDIS packet-level interface to the TCP/IP and AppleTalk protocols. |
| Ndiswan.sys | The NDISWAN driver is an intermediate NDIS driver that provides an IEEE 802.3 miniport interface to protocol drivers and a protocol interface to WAN miniport drivers. NDISWAN provides framing, compression, and encryption services for remote access connections. |
| WAN miniport driver | The WAN miniport driver contains the code required to operate the dial-up equipment. To use an adapter supporting WAN media, such as Integrated Services Digital Network (ISDN) or asynchronous transfer mode (ATM) with Windows Server 2003 remote access, the adapter vendor must create a WAN miniport driver. |
| Remote access components | These components are libraries that provide remote access programming interface for applications, Point-to-Point (PPP) protocols (link control, authentication, and network control protocols), and so on. Remote access components can communicate directly with the NDISWAN driver or by accessing the Telephony API (TAPI). |
| TAPI components | These components are libraries that provide a call control programming interface for all TAPI-aware applications. TAPI components communicate directly with the NDISWAN driver to manage connections. |

Routing and Remote Access can be invoked from the command line and is available to the following application programming interfaces (APIs):

- For the remote access server, MultiProtocol Routing (MPR) API.
- For the remote access client, remote access API.

For more information about the APIs associated with remote access, see the Windows Server 2003 Software Development Kit (SDK).

**Dial-up Remote Access Protocols**

Dial-up remote access uses PPP, a data link layer protocol that uses network layer protocols, including TCP/IP and AppleTalk, to transmit data across a WAN infrastructure.

This section describes PPP and the following protocols and methods that are part of the process when using PPP.

- PPP encapsulation, including PPP over Ethernet (PPPoE)
- PPP Link Control Protocol (LCP)
- Multilink Protocol (MP), Bandwidth Allocation Protocol (BAP), and Bandwidth Allocation Control Protocol (BACP)
- PPP Network Control Protocols (NCPs)
- PPP authentication protocols

PPP performs the following functions:

- Provides multiprotocol data-link layer encapsulation.

  PPP creates frames that contain separate Internet Protocol (IP) datagrams.

- Establishes, maintains, and ends the logical link.

  The PPP protocol uses LCP to establish and configure the parameters of the data-link connection. LCP negotiation includes the authentication of the credentials of the remote access client.

- Provides protocol configuration.

  After the data-link connection has been negotiated, NCPs such as Internet Protocol Control Protocol (IPCP) and AppleTalk Control Protocol (ATCP) negotiate parameters to configure PPP peers. For example, for IPCP, an IP address is allocated to

the remote access client by the remote access server. Compression and encryption are also negotiated.

**Note**

- PPP can terminate the link at any time. Termination generally occurs due to carrier loss, authentication failure, link quality failure, time-out, or link closure by the dial-up client or system administrator. When the link is closing, PPP informs the network layer protocols so that they can take appropriate action.

**PPP Encapsulation**

PPP encapsulation uses a variant of the International Organization for Standardization (ISO) High-Level Data Link Control (HDLC) protocol to encapsulate multiprotocol datagrams as the payload of PPP frames. The PPP header and trailer are shown in the following figure.

**PPP Encapsulation**




The fields are described as follows:

- **Flag** - Set to 0x7E (bit sequence 011111110) to signify the beginning and end of a PPP frame. In successive PPP frames, only a single Flag character is used.

  **Note**

  - The Flag character must not occur in the middle of the PPP frame. Character stuffing and bit stuffing are used to ensure that the Flag character occurs at the beginning and the end of the frame. For more information, see "Dial-up Remote Access Processes and Interactions" later in this section.

- **Address** - In HDLC environments, the Address field is used to address the frame to the destination node. On a point-to-point link, the destination node does not need to be addressed. Therefore, for PPP, the Address field is set to 0xFF, the broadcast address. If both PPP peers agree to perform address and control field compression during LCP negotiation, the Address field is not included.

- **Control** - In HDLC environments, the Control field is used for data-link layer sequencing and acknowledgments. PPP does not provide link-to-link reliable data transfer. Therefore, for all PPP frames, the Control field is set to 0x03 to indicate an unnumbered information frame. If both PPP peers agree to perform address and control field compression during LCP negotiation, the Control field is not included.

- **Protocol ID** - The 2-byte Protocol ID field identifies the protocol of the PPP payload. If both PPP peers agree to perform protocol field compression during LCP negotiation, the Protocol ID field is one byte for Protocol IDs in the range of 0x00-00 to 0x00-FF.

- **PPP Payload** - An encapsulation of IP datagrams or AppleTalk datagrams.

- **Frame Check Sequence (FCS)** - A 16-bit checksum that is used to check for bit-level errors in the PPP frame. If the receiver's calculation of the FCS does not match the FCS in the PPP frame, the PPP frame is silently discarded.

The maximum size of a PPP frame, known as the maximum receive unit (MRU), is determined during the negotiation of the logical link. The default MRU is 1,500 bytes. If the MRU is negotiated lower than 1,500 bytes, a PPP host must still have the ability to receive 1,500-byte frames in the event of link synchronization failure.

The following table lists typical values for the PPP Protocol ID.

**PPP Protocol IDs**

| Protocol | Protocol ID/Compressed Value |
|---|---|
| IP | 0x00-21 / 0x21 |
| AppleTalk | 0x00-29 / 0x29 |
| Van Jacobsen Compressed TCP/IP | 0x00-2D / 0x2D |
| Multilink | 0x00-3D / 0x3D |
| Microsoft Point-to-Point Compression Protocol (MPPC) | 0x00-FD / 0xFD |
| Microsoft Point-to-Point Encryption Protocol (MPPE) | 0x00-FD / 0xFD |

If MPPE or MPPC are negotiated, then the PPP Protocol ID is set to 0x00-FD. With MPPE and MPPC both using the same PPP Protocol ID, each peer must know that the resulting PPP payload either is encrypted or compressed, or both.

- If only MPPC is negotiated, then the PPP Protocol ID is set to 0x00-FD and the PPP payload is compressed.
- If only MPPE is negotiated, then the PPP Protocol ID is set to 0x00-FD and the PPP payload is encrypted.
- If both MPPC and MPPE are negotiated, then compression always occurs before encryption. The compressed PPP frame, consisting of the PPP protocol ID field set to 0xFD and the compressed data, is then encrypted and encapsulated with another PPP header consisting of the protocol ID field set to 0xFD and a 2-byte MPPE header.

**PPPoE**

PPPoE (PPP over Ethernet) is a method of encapsulating PPP frames so that they can be sent over an Ethernet network. PPPoE was created so that Internet service providers (ISPs) that deploy a broadband Internet access technology in a bridged Ethernet topology, such as cable modems or Digital Subscriber Line (DSL), can use the per-user authentication and connection identification facilities of PPP to identify individual customer connections for accounting and billing purposes. For more information about PPPoE, see RFC 2516.

The following figure shows a PPPoE frame.

**PPPoE Frame Structure**




The following are the fields in the PPPoE frame:

- **Preamble** - An 8-byte field that consists of 7 bytes of alternating 1s and 0s (each byte is the bit sequence 10101010) to synchronize a receiving station and a 1-byte 10101011 sequence that indicates the start of a frame. The Preamble provides receiver synchronization and frame delimitation services.
- **Destination Address** - A 6-byte field that indicates the destination's address. The destination can be a unicast, a multicast, or the Ethernet broadcast address. The unicast address is also known as an individual, physical, hardware, or media access control (MAC) address. For the Ethernet broadcast address, all 48 bits are set to 1 to create the address 0xFF-FF-FF-FF-FF-FF.
- **Source Address** - A 6-byte field that indicates the sending node's unicast address.

- **EtherType** - The Ethernet II header for PPPoE frames. EtherType is set to 0x88-63 for PPPoE discovery frames and 0x88-64 for PPP session frames. For more information about Ethernet II, see RFC 894.

- **Version** - A 4-bit field that is set to the value of 1.

- **Type** - A 4-bit field that is set to the value of 1.

- **Code** - A 1-byte field that is used to identify the type of PPPoE message. There are defined values for the PPPoE frames exchanged during the discovery phase. For PPP session frames, the Code field is set to 0.

- **Session_ID** - A 2-byte field that identifies the PPPoE session ID. This field is set to 0 until a session ID is negotiated with the access concentrator (AC) during the discovery phase of the PPPoE connection.

- **Length** - A 2-byte field that is used to indicate the size, in bytes, of the PPPoE payload.

- **PPPoE Payload** - A variable-sized payload that can contain either one or more PPPoE tags for PPPoE frames sent during the discovery phase or PPP frames for the PPP session phase. PPPoE tags are information elements in type-length-value (TLV) format. Typical PPPoE tags used during the discovery phase are Service-Name (the name of the ISP or service offered by the AC) and AC-Name (the name of the AC). For a complete list of PPPoE tags and their structure, see RFC 2516.

- **Frame Check Sequence (FCS)** - A 4-byte field that provides bit-level integrity verification on the bits in the Ethernet II frame. The FCS is also called a cyclical redundancy check (CRC). The source node calculates the FCS and places the result in this field. When the destination receives the FCS, it runs the same CRC algorithm and compares its own value with the one placed in the FCS field by the source node. If the two values match, the frame is considered valid and the destination node processes it. If the two values do not match, the frame is silently discarded.

For more information about the PPPoE connection process, see "Dial-up Remote Access Processes and Interactions" later in this section.

**LCP**

LCP (Link Control Protocol) negotiates link and PPP parameters to dynamically configure the data link layer of a PPP connection. Common LCP options include the PPP MRU, the authentication protocol, compression of PPP header fields, callback, and multilink options. For more information about LCP, see RFC 1661.

**LCP Packet Structure**

LCP uses the PPP Protocol ID of 0xC0-21. The packet structure of LCP is illustrated in the following figure. Each LCP packet is a single LCP message that consists of an LCP Code field identifying the type of LCP packet; an Identifier field so that requests and replies can be matched; a Length field indicating the size of the LCP packet; and LCP packet type-specific data.

**LCP Packet Structure**




The following table lists the LCP packet types.

**LCP Packet Types**

| LCP Code | LCP Packet Type | Description |
|---|---|---|

| | | |
|---|---|---|
| 1 | Configure-Request | Sent to open or reset a PPP connection. Configure-Request contains a list of LCP options with changes to default option values. |
| 2 | Configure-Ack | Sent when all of the values of all of the LCP options in the last Configure-Request received are recognized and acceptable. When both PPP peers send and receive Configure-Acks, the LCP negotiation is complete. |
| 3 | Configure-Nack | Sent when all the LCP options are recognized, but the values of some options are not acceptable. Configure-Nack includes the offending options and their acceptable values. |
| 4 | Configure-Reject | Sent when LCP options are not recognized or not acceptable for negotiation. Configure-Reject includes the unrecognized or non-negotiable options. |
| 5 | Terminate-Request | Optionally sent to close the PPP connection. |
| 6 | Terminate-Ack | Sent in response to the Terminate-Request. |
| 7 | Code-Reject | Sent when the LCP code is unknown. The Code-Reject message includes the offending LCP packet. |
| 8 | Protocol-Reject | Sent when the PPP frame contains an unknown Protocol ID. The Protocol-Reject message includes the offending LCP packet. Protocol-Reject is typically sent by a PPP peer in response to a PPP NCP for a local area network (LAN) protocol not enabled on the PPP peer. |
| 9 | Echo-Request | Optionally sent to test the PPP connection. |
| 10 | Echo-Reply | Sent in response to an Echo-Request. The PPP Echo-Request and Echo-Reply are not related to the Internet Control Message Protocol (ICMP) Echo Request and Echo Reply messages. |
| 11 | Discard-Request | Optionally sent to exercise the link in the outbound direction. |

**LCP Options**

When using the Configure-Request, Configure-Ack, Configure-Nack, and Configure-Reject LCP packet types, the LCP data portion of the LCP packet consists of one or more LCP options, as illustrated in the following figure. Each LCP option consists of an option Type field; a Length field indicating the total length, in bytes, of the option; and the data associated with the option.

**LCP Options**




The following table lists common LCP options negotiated by Microsoft PPP peers.

**LCP Options**

| Option Name | Option Type | Option Length | Description |
|---|---|---|---|
| Maximum Receive Unit (MRU) | 1 | 4 | The maximum size (up to 65,535) of the PPP frame. The default MRU is 1,500. If neither peer is changing the default, this option is not negotiated. |

| | | | |
|---|---|---|---|
| Asynchronous Control Character Map (ACCM) | 2 | 6 | A bit map that enables (bit set to 1) or disables (bit set to 0) the use of character escapes for asynchronous links for the 32 ASCII control characters from 0x00 to 0x20. By default, character escapes are used. The ACCM bit map is set to 0x00-00-00-00 for links with XON/XOFF software flow control. |
| Authentication Protocol | 3 | 5 or 6 | Indicates the authentication protocol used during the authentication phase of the connection. Values for this field for Microsoft PPP peers are 0xC2-27 for Extensible Authentication Protocol (EAP); 0xC2-23-80 for Microsoft Challenge Handshake Authentication Protocol version 1 (MS-CHAP); 0xC2-23-81 for MS-CHAP version 2; 0xC2-23-05 for Message Digest 5 Challenge Handshake Authentication Protocol (MD5-CHAP); 0xC0-27 for Shiva Password Authentication Protocol (SPAP); and 0xC0-23 for PAP. |
| Magic Number | 5 | 6 | A random number chosen to distinguish a peer and detect looped back lines. |
| Protocol Compression | 7 | 2 | A flag indicating that the PPP Protocol ID be compressed to a single octet when the 2-byte Protocol ID is in the range of 0x00-00 to 0x00-FF. |
| Address and Control Field Compression | 8 | 2 | A flag indicating that the PPP Address field (always set to 0xFF) and the PPP Control field (always set to 0x03) must be removed from the PPP header. |
| Callback | 13 or 0x0D | 3 | A 1-octet indicator of how callback is to be determined. For remote access clients and servers running Microsoft Windows 32-bit operating systems, the callback option octet is set to 0x06, indicating that the callback is determined during Callback Control Protocol (CBCP) negotiation. |

**NCP**

After the link and PPP parameters have been negotiated with LCP, the PPP peers then use a series of NCPs (Network Control Protocols) to negotiate the parameters of individual LAN protocols. Microsoft PPP supports the NCPs listed in the following table.

### Network Control Protocols

| Protocol | Description |
|---|---|
| IPCP | IPCP negotiates the use of IP. |
| ATCP | ATCP negotiates the use of AppleTalk. |
| Compression Control Protocol (CCP) | CCP packets are exchanged to configure MPPC (Microsoft Point-to-Point Compression Protocol) and MPPE (Microsoft Point-to-Point Encryption Protocol). |
| BACP (Bandwidth Allocation Control Protocol) | Working with MP (Multilink Protocol) and BAP (Bandwidth Allocation Protocol), BACP elects a favored peer in case both PPP peers request to add or remove a connection at the same time. |

**IPCP**

IPCP (Internet Protocol Control Protocol) as used by Microsoft PPP peers is documented in RFC 1332 and RFC 1877. IPCP negotiates IP-based parameters to dynamically configure a TCP/IP-based PPP peer across a point-to-point link. Common IPCP options include an IP address and the IP addresses of Domain Name System (DNS) servers.

**Packet Structure**

IPCP uses the PPP Protocol ID of 0x80-21. The packet structure of IPCP is exactly the same as LCP, except only packet types 1 through 7 are defined. For Configure-Request, Configure-Ack, Configure-Nack, and Configure-Reject IPCP packet types, the IPCP data portion of the IPCP packet consists of one or more IPCP options. Each IPCP option consists of an Option Type field; an Option Length field indicating the total length, in bytes, of the option; and the data associated with the option.

**Negotiated Options**

The following table lists the IPCP options negotiated by Microsoft PPP peers.

### IPCP Options

| Option Name | Option Type | Option Length | Description |
|---|---|---|---|

| | | | |
|---|---|---|---|
| IP compression protocol | 2 | 4 | The Van Jacobsen TCP compression protocol. |
| IP address | 3 | 6 | The IP address to be allocated to the remote access client. |
| Primary DNS server address | 129 or 0x81 | 6 | The primary DNS server for the remote access client. |
| Primary NBNS server address | 130 or 0x82 | 6 | The primary NBNS (WINS) server for the remote access client. |
| Secondary DNS server address | 131 or 0x83 | 6 | The secondary DNS server for the remote access client. |
| Secondary NBNS server address | 132 or 0x84 | 6 | The secondary NBNS (WINS) server for the remote access client. |

There are no IPCP options for these common TCP/IP configuration items:

- Subnet mask

  The subnet mask is assumed by the remote access client to be the class-based subnet mask of the IP address that is allocated to the remote access client.

- Default gateway

  The default gateway IP address is not allocated by the remote access server. However, a default route is created on the remote access client, which points to the remote access connection. If a default route already exists in the routing table, then the metric of the existing default route is increased and a new default route is added with a lower metric. This is the default behavior for remote access clients running Windows 32-bit operating systems. This behavior can be modified by disabling the **UseDefault Gateway on Remote Network** setting on the TCP/IP properties of a remote access client's phone book entry or dial-up connection object.

- DNS domain name

  The DNS domain name configured from the TCP/IP protocol properties on the remote access server is not negotiated during IPCP. For Windows Server 2003 remote access clients, the DNS domain name can be obtained through a DHCPInform message. For more information, see "TCP/IP Configuration for Remote Access Clients" later in this section.

**ATCP**

ATCP (AppleTalk Control Protocol) as used by Microsoft PPP peers is documented in RFC 1378. ATCP negotiates AppleTalk-based parameters to dynamically configure an AppleTalk-based PPP peer across a point-to-point link. Common ATCP options include AppleTalk address and server information.

**Packet Structure**

ATCP uses the PPP Protocol ID of 0x80-29. The packet structure of ATCP is exactly the same as the packet structure for LCP, except that only packet types 1 through 7 are defined. For Configure-Request, Configure-Ack, Configure-Nack, and Configure-Reject ATCP packet types, the ATCP data portion of the ATCP packet consists of one or more ATCP options. Each ATCP option consists of an option Type field; an option Length field indicating the total length, in bytes, of the option; and the data associated with the option.

**Negotiated Options**

The following table lists the ATCP options negotiated by Microsoft PPP peers.

### ATCP Options

| Option Name | Option Type | Option Length | Description |
|---|---|---|---|
| AppleTalk Address | 1 | 6 | Negotiates the AppleTalk network and node numbers. |
| Server Information | 3 | 16 | Used to convey information about the remote access server. |

**CCP**

CCP (Compression Control Protocol) is documented in RFC 1962. CCP negotiates parameters to dynamically configure, enable, and disable data compression algorithms between PPP peers across a point-to-point link. Common CCP options include an organization identifier and the use of MPPC.

**Packet Structure**
CCP uses the PPP Protocol ID of 0x80-FD. The packet structure of CCP is exactly the same as the packet structure for LCP, except only packet types 1 through 7 are defined. For Configure-Request, Configure-Ack, Configure-Nack, and Configure-Reject CCP packet types, the CCP data portion of the CCP packet consists of one or more CCP options. Each CCP option consists of an option Type field; an option Length field indicating the total length, in bytes, of the option; and the data associated with the option.

**Negotiated Options**
The following table lists the CCP options negotiated by Microsoft PPP peers.

### CCP Options

| Option Name | Option Type | Option Length | Description |
|---|---|---|---|
| Organization Unique Identifier | 0 | 6 or larger | Used to negotiate an organization's proprietary compression protocol. |
| MPPC | 18 or 0x12 | 6 | Used to indicate the use of MPPC, MPPE, and the encryption strength. |

**MPPE and MPPC**
With CCP option 18, Microsoft PPP peers negotiate both MPPC and MPPE at the same time. The option data field for CCP option 18 is 4 bytes (32 bits) long. Bits within this data field are used as flags to indicate:

- Whether compression is enabled (0x00-00-00-01).
- Whether 40-bit session keys are derived from the LAN Manager version of the user's password (0x00-00-00-10).
- Whether 40-bit session keys are derived from the Windows NT version of the user's password (0x00-00-02-00).
- Whether 56-bit session keys are derived from the Windows NT version of the user's password (0x00-00-00-80).
- Whether 128-bit session keys are derived from the Windows NT version of the user's password (0x00-00-00-40).
- Whether the encryption keys are refreshed with each PPP frame (0x01-00-00-00).

For multiple choices, the flag values are added together. For example, for compression (0x00-00-00-01) and 128-bit encryption keys (0x00-00-00-40), the resulting 32-bit option data field is set to 0x00-00-00-41.

### Note

- The Encryption Control Protocol (ECP) is used to negotiate a specific encryption method and is documented in RFC 1968. However, for Microsoft PPP peers, the only supported encryption is MPPE, which is negotiated during CCP with the negotiation of MPPC. Therefore, Microsoft PPP peers do not use ECP.

**MP, BAP, and BACP**
Windows Server 2003 remote access supports the Multilink protocols listed in the following table.

### Multilink Protocols

| Protocol | Description |
|---|---|
| MP | Allows multiple physical links to appear as a single logical link over which data can be sent and received. |
| BAP | Used to dynamically add or remove additional links to an MP connection. |
| BACP | Elects a favored peer in case both PPP peers request to add or remove a connection at the same time. |

**MP**
MP is defined in RFC 1990 and used to aggregate multiple physical links into a single logical link. A good example is the aggregation of both B-channels of an ISDN Basic Rate Interface (BRI) connection. MP fragments, sequences, and re-orders

alternating packets sent across multiple physical connections so that the end result is a single logical link with the combined bandwidth of all of the aggregated physical links. MP is the recommended method of combining multiple B-channels of a BRI connection because the support for bonding, the combining of ISDN B-channels through hardware support, can be specific to the ISDN adapter. MP can be used for any ISDN adapter. MP must be supported on both sides of the connection.

The following figure illustrates the structure of an MP frame. The payload of a Multilink PPP packet is either a fragment of a PPP frame or the entire PPP frame. Multilink PPP fragmentation need not occur if the Multilink PPP packet fits within the MRU of the link. To prevent incorrect ordering of the datagrams or fragments across multiple links, additional fields are used between the PPP Protocol field and the IP datagram. Multilink PPP uses the PPP Protocol ID of 0x00-3D.

**Multilink PPP**




RFC 1717 defines two different packet formats for short sequence numbers and long sequence numbers. The sequence number is used to prevent the misordering of frames that are sent across multiple links, not to sequence fragments.

For short sequence numbers, a two-byte Delimitation/Sequence # field consists of four bits used for delimitation and 12 bits used for the sequence number. Within the delimitation field are two flags. The first bit (the Beginning bit) is an indicator that this fragment begins a sequence of fragments corresponding to a packet. The second bit (the Ending bit) is an indicator that this fragment ends a sequence of fragments corresponding to a packet. The other bits in the first four bits of the short sequence number header are set to 0.

For a PPP frame that is sent without fragmentation, both the Beginning and Ending bits are set. For a PPP frame that is larger than the MRU of the physical link, the PPP frame is fragmented, and each fragment is sent as a separate PPP packet. MP performs a data-link layer fragmentation that is not related to IP fragmentation.

For long sequence numbers, a four-byte Delimitation/Sequence # field consists of eight bits (one byte) used for delimitation, and 24 bits (3 bytes) used for the sequence number. Within the delimitation field, the same bits as the short sequence number header define the Beginning bit and the Ending bit. The other bits in the first byte of the long sequence number header are set to 0. The long sequence number header is used by default, unless the short sequence number is chosen during LCP negotiation.

The following table lists Multilink LCP options negotiated by Microsoft PPP peers. For information about other Multilink options, see RFC 1990.

**Multilink LCP Options**

| Option Name | Option Type | Option Length | Description |
|---|---|---|---|
| Multilink Maximum Receive Reconstructed Unit | 17 or 0x11 | 4 | Specifies the number of octets that a peer can reconstruct when performing reassembly of fragmented MP frames. |
| Short Sequence Number Header Format | 18 or 0x12 | 2 | Specifies the use of the short sequence number in the MP header. |
| Multilink Endpoint Discriminator | 19 or 0x13 | 9 | A unique system identifier to differentiate links from two PPP peers with the same authenticated name. |

**BAP**

Although MP allows for multiple physical links to be aggregated, MP does not provide a mechanism to adapt to changing conditions by adding extra links when needed or terminating extra links when unneeded. This additional capability is provided by BAP and BACP, as defined in RFC 2125. BAP is a PPP control protocol that is used on an MP connection to dynamically manage links. BAP uses the PPP Protocol of ID 0xC0-2D.

For example, an MP and BAP-enabled remote access client and remote access server create an MP connection consisting of a single physical link. As the use of the single link rises to a configured level, the remote access client uses a BAP Call-Request

message to request an additional link. The BAP Call-Request message specifies the type of link desired, such as analog phone, ISDN, or X.25. The remote access server then sends a BAP Call-Response message that contains the phone number of an available port on the remote access server of the same type specified by the remote access client in the BAP Call-Request.

When use on the second link drops to a specified level, either the remote access client or the remote access server can send a BAP Link-Drop-Query-Request message to drop the link.

BAP also supports a Callback-Request message in which the requesting peer specifies the link type and the number to call back. For more information about BAP messages, see RFC 2125.

The following table lists the BAP LCP option negotiated by Microsoft PPP peers.

**BAP LCP Option**

| Option Name | Option Type | Option Length | Description |
|---|---|---|---|
| BAP Link Discriminator | 23 or 0x17 | 4 | A unique number used to identify a link in a Multilink PPP connection. |

**BACP**

BACP is a PPP NCP that negotiates a single option: the election of a favored peer. If both peers of an MP and BAP-enabled connection send BAP Call-Request or BAP Link-Drop-Query-Request messages at the same time, the favored peer is the peer whose requests are implemented.

BACP uses the PPP Protocol ID of 0xC0-2B. The packet structure of BACP is exactly the same as the packet structure for LCP, except that only packet types 1 through 7 are defined. For Configure-Request, Configure-Ack, Configure-Nack, and Configure-Reject BACP packet types, the BACP data portion of the BACP packet consists of the single BACP Favored-Peer option listed in the following table.

**BACP Favored Peer Option**

| Option Name | Option Type | Option Length | Description |
|---|---|---|---|
| Favored-Peer | 1 | 6 | A randomly allocated 4-byte magic number used to elect a favored BAP peer. The favored peer is the peer with the lowest magic number. |

**PPP Authentication Protocols**

After both PPP peers agree on a single, specific PPP authentication protocol, authentication for PPP is accomplished through that protocol.

Windows Server 2003 remote access supports the PPP authentication protocols listed in the following table.

**PPP Authentication Protocols**

| Protocol | Description | Security Level |
|---|---|---|
| PAP ( Password Authentication Protocol) | Uses plaintext passwords. It is typically negotiated if the remote access client and remote access server cannot negotiate a more secure form of validation. | The least secure authentication protocol.<br><br>Does not protect against replay attacks, remote client impersonation, or remote server impersonation. |
| SPAP (Shiva Password Authentication Protocol) | A reversible encryption mechanism employed by Shiva remote access servers. | More secure than PAP, but less secure than CHAP or MS-CHAP.<br><br>Does not protect against remote server impersonation. |

| | | |
|---|---|---|
| CHAP (Challenge Handshake Authentication Protocol) | A challenge-response authentication protocol that uses the industry-standard Message Digest 5 (MD5) hashing scheme to encrypt the response. | An improvement over PAP and SPAP in that the password is not sent over the PPP link.<br><br>Requires a plaintext version of the password to validate the challenge response.<br><br>Does not protect against remote server impersonation. |
| MS-CHAP (Microsoft Challenge Handshake Authentication Protocol version 1) | A nonreversible, encrypted password authentication protocol that is similar to CHAP. | More secure than CHAP.<br><br>Requires an MD4 hash of the password to validate the challenge response.<br><br>Does not protect against remote server impersonation |
| MS-CHAP v2 (Microsoft Challenge Handshake Authentication Protocol version 2) | An upgrade of MS-CHAP. | Provides stronger security than MS-CHAP.<br><br>A comparison of MS-CHAP and MS-CHAP v2 features appears later in this section. |
| EAP (Extensible Authentication Protocol) | Allows for arbitrary authentication of a remote access connection through the use of authentication schemes known as EAP types. | Offers the strongest security by providing the most flexibility in authentication variations. |

For information about authentication processes for PPP authentication protocols, see "Dial-up and Remote Access Processes and Interactions" later in this section.

**PAP**

PAP is included in Windows Server 2003 so that:

- Remote access clients running Windows 32-bit operating systems can connect to older remote access servers that do not support a secure authentication protocol
- Remote access clients running Microsoft operating systems that do not support a secure remote access protocol can connect to remote access servers running Windows 32-bit operating systems.

**Note**

- To make your remote access server more secure, ensure that PAP is disabled. As a consequence, however, older remote access clients not running Microsoft operating systems that do not support secure authentication protocols will be unable to connect.

**SPAP**

A Windows Server 2003 remote access client can use SPAP to authenticate itself to a Shiva remote access server. A remote access client running Windows 32-bit operating systems can use SPAP to authenticate itself to a Windows Server 2003 remote access server.

**CHAP**

CHAP is used by vendors of dial-in servers and clients. CHAP is supported by both the Windows Server 2003 remote access server and remote access client.

CHAP protects against replay attacks by using an arbitrary challenge string per authentication attempt. Instead of sending the password over the link, the password is used to create a one-way hash from a challenge string. The server, knowing the client's password, can duplicate the operation and compare the result with that sent in the client's response.

CHAP requires that local or domain passwords are stored in a reversibly encrypted form.

**MS-CHAP**
As in CHAP, the remote access server sends a challenge that consists of a session ID and an arbitrary challenge string to the remote client. The remote client must return the user name and an MD4 hash of the challenge string, the session ID, and the MD4-hashed password.

One difference between CHAP and MS-CHAP is that, in CHAP, the plaintext version of the password must be available to validate the challenge response. With MS-CHAP, the remote access server requires only the MD4 hash of the password to validate the challenge response. In Windows Server 2003, the user's password is stored as an MD4 hash and in a reversibly encrypted form. When CHAP is used, the remote access server decrypts the reversibly encrypted password to validate the remote access client's response.

MS-CHAP also allows for error codes including a "password expired" code and password changes. MS-CHAP protects against replay attacks by using an arbitrary challenge string per authentication attempt.

If MS-CHAP is used as the authentication protocol and MPPE is negotiated, then shared secret encryption keys are generated by each PPP peer. MS-CHAP also provides a set of messages that allows a user to change the password during the user authentication process.

**MS-CHAP v2**
Windows Server 2003 includes support for MS-CHAP v2, which provides stronger security for remote access connections. The following table compares the security features of MS-CHAP to those of MS-CHAP v2.

**MS-CHAP and MS-CHAP v2 Security Feature Comparison**

| MS-CHAP | MS-CHAP v2 |
|---|---|
| LAN Manager encoding of the response used for backward compatibility with older Microsoft remote access clients is cryptographically weak. | LAN Manager-encoded responses are not allowed. |
| LAN Manager encoding of password changes is cryptographically weak. | LAN Manager-encoded password changes are not allowed. |
| Only one-way authentication is possible. The remote access client cannot verify that it is dialing in to its organization's remote access server or a masquerading remote access server. | Two-way authentication, also known as mutual authentication, is provided. The remote access client receives verification that the remote access server that it is dialing in to has access to the user's password. |
| With 40-bit encryption, the cryptographic key is based on the user's password. Each time the user connects with the same password, the same cryptographic key is generated. | The cryptographic key is always based on the user's password and an arbitrary challenge string. Each time the user connects with the same password, a different cryptographic key is used. |
| A single cryptographic key is used for data sent in both directions on the connection. | Separate cryptographic keys are generated for transmitted and received data. |

**EAP**
EAP is an extension to PPP that allows for arbitrary authentication mechanisms to be employed for the validation of a PPP connection. With PPP authentication protocols such as MS-CHAP and SPAP, a specific authentication mechanism is chosen during the link establishment phase. Then, during the connection authentication phase, the negotiated authentication protocol is used to validate the connection. The authentication protocol itself is a fixed series of messages sent in a specific order.

With EAP, the authentication mechanism is not chosen during the link establishment phase. Instead, each PPP peer negotiates to perform EAP during the connection authentication phase. After the connection authentication phase is reached, the PPP peers must first negotiate the use of a specific EAP authentication scheme known as an EAP type. After the EAP type is agreed on, EAP allows for an open-ended conversation between the remote access client and the remote access server that can vary based on the parameters of the connection. The conversation consists of requests for authentication information and the responses. The length and detail of the authentication conversation depends on the EAP type.

For example, when EAP is used with security token cards, the remote access server can separately query the remote access client for a name, PIN, and card token value. As each query is asked and answered, the user passes through another level of authentication. When all questions have been answered satisfactorily, the user is authenticated and permitted access to the network.

The use of EAP is negotiated during LCP negotiation by specifying the authentication protocol LCP option (type 3) and the authentication protocol 0xC2-27. After LCP negotiation is complete, EAP messages use the PPP Protocol ID of 0xC2-27.

Windows Server 2003 includes support for the MD5-Challenge and Extensible Authentication Protocol-Transport Level Security (EAP-TLS) EAP types.

EAP is designed to allow authentication plug-in modules at both the client and server ends of a connection. By installing an EAP type library file on both the remote access client and the remote access server, a new EAP type can be supported. This presents vendors with the opportunity to supply a new authentication scheme at any time. EAP provides the most flexibility in authentication uniqueness and variations.

**Note**

- Protected Extensible Authentication Protocol (PEAP) is used as an authentication method for 802.11 wireless client computers, but is not supported for other remote access clients. For more information, see the "IAS Technical Reference" and "802.11 Wireless Technical Reference."

**EAP Infrastructure**

EAP is a set of internal components that provide architectural support for EAP types. For successful authentication, both the remote access client and authenticator must have the same EAP authentication module installed. Windows Server 2003 provides two EAP types: MD5-Challenge and EAP-TLS. EAP-TLS is available only for members of a domain. In addition, EAP contains EAP-RADIUS, which is not an EAP type, but is used to send EAP messages to a RADIUS server for authentication.

You can also install additional EAP types. The components for an EAP type must be installed on every remote access client and every authenticator.

**MD5-Challenge**

MD5-Challenge is the CHAP authentication mechanism used within the EAP framework. Rather than negotiating to perform MD5-Challenge authentication during the link establishment phase, the authenticator and peer negotiate to use EAP during the connection authentication phase.

MD5-Challenge is a required EAP type and can be used to test EAP interoperability. Like CHAP, MD5-Challenge requires that local or domain passwords are stored in a reversibly encrypted form.

**EAP-TLS**

The Transport Layer Security (TLS) protocol, based on Secure Sockets Layer (SSL), allows applications to communicate securely. TLS provides authentication (user and data), data integrity, and data confidentiality services. To achieve these services, TLS specifies a framework that allows the following:

- Client and two-way authentication using symmetric or asymmetric encryption.
- Negotiation of the specific encryption algorithm (the cipher-suite).
- Secured exchange of encryption keys to be used for encrypting messages.
- Message integrity and user authentication using a message authentication code.

For more information about TLS, see RFC 2246. For more information about EAP-TLS, see RFC 2716.

EAP-TLS uses the TLS protocol for authentication during the establishment of a PPP connection. With EAP-TLS, mutual authentication between the PPP client and the authenticator is performed through the exchange and verification of certificates. The client attempting the connection sends a user certificate, and the authenticator sends a machine certificate. If smart cards are used for remote access authentication, EAP-TLS is the required authentication method. For more information about smart card authentication, see "IAS Technical Reference."

EAP-TLS is supported only on Windows Server 2003 remote access server computers that are members of a Windows Server 2003 mixed or native domain. Stand-alone Windows Server 2003 remote access servers do not support EAP-TLS.

**EAP-RADIUS**

EAP-RADIUS passes EAP messages of any EAP type by the remote access server to a RADIUS server for authentication. The EAP messages sent between the remote access client and remote access server are encapsulated and formatted as RADIUS messages between the remote access server and the RADIUS server. The remote access server becomes a pass-through device passing EAP messages between the remote access client and the RADIUS server. All processing of EAP messages occurs at the remote access client and the RADIUS server.

EAP-RADIUS is used in environments in which Remote Authentication Dial-In User Service (RADIUS) is used as the

authentication provider. An advantage of using EAP-RADIUS is that EAP types do not need to be installed at each remote access server, only at the RADIUS server.

In a typical use of EAP-RADIUS, the remote access server is configured to use EAP and to use RADIUS as its authentication provider. When a connection attempt is made, the remote access client negotiates the use of EAP with the remote access server. When the client sends an EAP message to the remote access server, the remote access server encapsulates the EAP message as a RADIUS message and sends it to its configured RADIUS server. The RADIUS server processes the EAP message and sends a RADIUS-encapsulated EAP message back to the remote access server. The remote access server then forwards the EAP message to the remote access client.

For more information about RADIUS support for EAP, see RFC 2869.

**CBCP**
The Microsoft implementation of PPP includes the optional use of CBCP (Callback Control Protocol) immediately after authentication.

CBCP negotiates the use of callback in those cases in which the remote access server, after authenticating the remote access client, terminates the physical connection, waits a specified amount of time, and then calls the remote access client back at either a static or dynamically configured phone number. Common CBCP options include the phone number that will be used by the remote access server to call back the remote access client.

**Packet Structure**
CBCP uses the PPP Protocol ID of 0xC0-29. The packet structure of CBCP is exactly the same as the packet structure for LCP; however, only the Callback-Request (type 1), Callback-Response (type 2), and Callback-Ack (type 3) types are used. For all CBCP packet types, the CBCP data portion of the CBCP packet consists of one or more CBCP options. Each CBCP option consists of an option Type field; an option Length field indicating the total length, in bytes, of the option; and the data associated with the option.

**Negotiated Options**
The following table lists the CBCP options negotiated by Microsoft PPP peers.

**CBCP Options**

| Option Name | Option Type | Option Length | Description |
|---|---|---|---|
| No callback | 1 | 2 | No callback is used for the connection. |
| Callback to a user-specified number | 2 | variable | The user of the remote access client computer determines the callback number. |
| Callback to an administrator-defined number | 3 | variable | Settings on the remote access server determine the callback number. |
| Callback to any of a list of numbers | 4 | variable | The remote access server calls back to one of a list of phone numbers. |

**Dial-up Remote Access Processes and Interactions**
This section describes the following processes that take place when establishing a dial-up remote access connection.

- Remote access connection
- TCP/IP configuration for Remote Access clients
- PPP connection
- Connection attempt processing with Remote Access policies
- Broadcast Name Resolution

**Remote Access Connection**
Connections are established by remote access clients that call the remote access programming interface, which, in turn, uses TAPI to pass call connection information to the dial-up equipment. After the physical connection is made, TAPI is no longer

used; other remote access components negotiate the connection with link, authentication, and network control protocols by communicating directly with NDISWAN.

After a remote access connection is established, protocol drivers can communicate over that connection using standard NDIS calls such as NdisSend(). NdisSend() calls for dial-up connections are forwarded to NDISWAN, which then determines the appropriate device and port, performs compression and encryption, provides PPP framing, and then forwards the completed frame to the WAN miniport driver. The WAN miniport driver then forwards the frame to the dial-up adapter.

All inbound remote access client connections, initiated by remote access clients to the remote access server, are represented as a single adapter called the remote access server interface. For each outbound remote access client connection, initiated by the remote access server, a separate interface is created.

To accept calls, the remote access server instructs each WAN miniport driver to indicate when it goes into a line-up state. When the call is placed, the WAN miniport driver passes the line-up state indicator through NDISWAN to the TAPI components. TAPI returns a call handle to NDISWAN to be used to refer to the physical connection, and then NDISWAN and the remote access components negotiate the rest of the remote access connection.

**IP and AppleTalk Routing with Remote Access Server**

After the remote access connection is established, the remote access client can begin sending LAN protocol traffic to the remote access server or to locations beyond the remote access server. When the remote access client sends LAN protocol traffic that is not destined for the remote access server, the remote access server must forward the LAN traffic to its appropriate destination. To accomplish this, the remote access server must have forwarding capabilities enabled on its routable protocols and act as an IP and AppleTalk router.

When Routing and Remote Access is installed and enabled to provide point-to-LAN remote access connectivity, it enables forwarding between the installed LAN adapters and the WAN miniport interface.

The following figure illustrates the remote access server architecture as it appears when routing IP packets. For simplicity, only IP routing is shown. However, AppleTalk routing operates in a similar way.

**IP Routing on the Remote Access Server**




**Packet Processing from Remote Access Clients**

The following process describes how IP packets sent by the remote access client are forwarded by the remote access server.

1. Depending on the dial-up technology, either the entire PPP frame is received by the WAN hardware and passed as a single frame to the appropriate WAN miniport driver, or individual bits of the PPP frame are passed to the appropriate WAN miniport driver.

2. The WAN miniport driver passes the PPP frame to Ndiswan.sys.

3. Ndiswan.sys verifies the PPP checksum and uses the PPP protocol ID to determine that it is an IP datagram.

4. The IP datagram is passed to the TCP/IP protocol driver.

5. The TCP/IP protocol driver, which is enabled for IP forwarding, determines a forwarding interface and an IP address based on the destination IP address in the IP datagram and the contents of its routing table.

6. To forward the IP datagram using the LAN adapter, the TCP/IP protocol calls NDIS with an NdisSend(), along with instructions to send it using the LAN adapter.

7. NDIS forwards the IP datagram to the appropriate LAN miniport driver.

8. The LAN miniport forwards the IP datagram to the LAN adapter through NDIS.

The end result is that packets from the remote access client are forwarded using the same IP routing process used for all IP routing. The success of the IP forwarding process depends on whether the remote access server can find a suitable entry in the IP routing table. Therefore, either the remote access server is configured with a default gateway, or the remote access server has specific routes to all the locations on the intranet to which the remote access server is attached. Specific routes can be added through static routes, or by enabling a routing protocol on the remote access server.

**Packet Forwarding to Remote Access Clients**
The following process describes how IP packets sent by intranet hosts to the remote access client are forwarded by the remote access server.

1. The LAN adapter passes a frame to its appropriate LAN miniport driver through NDIS. For information about how an IP datagram is forwarded to the MAC address of the remote access server, see "TCP/IP On-Subnet and Off-Subnet Addressing" in the next section.

2. The LAN miniport driver passes the IP datagram to the TCP/IP protocol driver through NDIS.

3. The TCP/IP protocol driver, which is enabled for IP forwarding, determines a forwarding interface and IP address based on the destination IP address in the IP datagram and the contents of its routing table. When the remote access client connects, a host route is created in the IP routing table for the IP address allocated to the remote access client that points to the remote access server interface.

4. To forward the IP datagram using the WAN adapter, the TCP/IP protocol calls NDIS with an NdisSend() with instructions to send it using NDISWAN and a specific connection handle.

5. NDISWAN resolves the connection handle to a specific device and port, adds a PPP header and trailer, and forwards the IP datagram to the appropriate WAN miniport driver through NDIS.

6. The WAN miniport driver forwards the IP datagram to the WAN adapter through NDIS.

The end result is that packets from intranet hosts are forwarded using the same IP routing process used for all IP routing. The success of the IP forwarding process depends on whether the IP addresses of remote access clients can be reached from the hosts on the intranet.

**TCP/IP On-Subnet and Off-Subnet Addressing**
The remote access server can be configured for on-subnet or off-subnet addressing. This addressing configuration determines how an IP node on a subnet to which the remote access server is attached resolves the MAC address of the remote access server.

- On-subnet addressing is the allocation of IP addresses to remote access clients that are in a range defined by a subnet to which the remote access server is attached. On-subnet addressing uses a subset of addresses of an attached subnet.

- Off-subnet addressing is the allocation of IP addresses to remote access clients that are not in a range defined by a subnet to which the remote access server is attached. Off-subnet addressing uses a separate subnet address space that is unique to the intranet.

**On-Subnet Addressing and Proxy ARP**
With on-subnet addressing, remote access clients are logically on the same subnet as a subnet attached to the remote access server. Proxy Address Resolution Protocol (ARP) is used by the remote access server to receive IP datagrams that are forwarded to remote access clients.

There are two cases in which Proxy ARP is used:

- When the remote access server is configured to use Dynamic Host Configuration Protocol (DHCP) to obtain addresses for IP-based remote access clients.

- When the remote access server is configured with a static IP address pool that consists of address ranges that are a subset of the addresses for a subnet to which the remote access server is attached.

In either case, the remote access clients are logically on the same subnet as the remote access server. Therefore, IP nodes on that subnet forwarding IP datagrams to a remote access client perform a direct delivery by sending a broadcast ARP Request frame for the remote access client's IP address.

The remote access client cannot respond to the ARP Request because the remote access server does not forward the ARP Request frame to the remote access client, and the remote access client does not have a MAC address corresponding to the remote access connection.

Therefore, the remote access server responds with an ARP Reply frame with its own MAC address. The node forwarding the packet then sends the IP datagram to the remote access server's MAC address. The remote access server then uses the IP routing process to forward the IP datagram across the dial-up connection to the remote access client.

**Off-Subnet Addressing and IP Routing**
With off-subnet addressing, remote access clients are logically on a separate subnet that can be reached across the remote access server. In this case, Proxy ARP is not used. The remote access server is acting as a router between the subnet of the remote access clients and the subnets to which the remote access server is attached. IP nodes on the LAN-based subnets attached to the remote access server that forward IP datagrams to a remote access client perform an indirect delivery by sending a broadcast ARP Request frame for the remote access server's IP address.

To reach remote access clients from IP nodes on the intranet, routes that represent the address ranges of the IP address pool and point to the LAN interface of the remote access server must be present in intranet routers.

When the first TCP/IP-based remote access client connects, routes corresponding to the off-subnet address ranges that point to the remote access server interface are added to the IP routing table of the remote access server. If the remote access server is configured with an IP routing protocol, the new routes are advertised to neighboring routers by using the advertising process of the configured routing protocol. If the remote access server is not configured with an IP routing protocol, routes corresponding to the off-subnet address ranges that point to the remote access server's LAN interface must be added to the routers of the intranet.

**TCP/IP Configuration for Remote Access Clients**
This section describes how the remote access server for Windows Server 2003 allocates network configuration parameters for TCP/IP-based remote access clients.

**IPCP Configuration**
To configure a TCP/IP-based remote access client with IPCP, the remote access server allocates an IP address and assigns the IP addresses of DNS and WINS servers.

**IP Address Allocation**
To allocate an IP address to a remote access client, the remote access server is either configured to use DHCP to obtain IP addresses, or with a static IP address pool.

**DHCP and Automatic Private IP Addressing**
When the remote access server is configured to use DHCP to obtain IP addresses, Routing and Remote Access instructs the DHCP client component to obtain 10 IP addresses from a DHCP server. The remote access server uses the first IP address obtained from DHCP for the remote access server interface, and subsequent addresses are allocated to TCP/IP-based remote access clients as they connect. IP addresses freed due to disconnected remote access clients are reused.

When all 10 IP addresses are used, the remote access server uses the DHCP client component to obtain 10 more. You can modify the number of IP addresses obtained at a time by changing a registry entry. For more information about registry entries for dial-up remote access, see "Dial-up Remote Access Tools and Settings."

With the Windows NT 4.0 remote access server, the DHCP-allocated addresses are recorded and reused when remote access is restarted. The Windows Server 2003 and Windows 2000 Server remote access servers now release all DHCP-allocated IP addresses using DHCPRELEASE messages each time the service is stopped.

If the remote access server initially starts using DHCP-allocated addresses and the DHCP server becomes unavailable, then an IP address cannot be allocated to additional TCP/IP-based remote access clients.

If a DHCP server is not available when Routing and Remote Access is started, then the DHCP client returns 10 addresses in the range of 169.254.0.1 to 169.254.255.254 to allocate to remote access clients. The address range 169.254.0.0/16 is used for Automatic Private IP Addressing (APIPA). APIPA addresses for point-to-LAN remote access connectivity work only if the network to which the remote access server computer is attached is also using APIPA addresses. If the local network is not using APIPA addresses, remote access clients are only able to obtain point-to-point remote access connectivity. For more information about APIPA, see "DHCP Technical Reference."

If a DHCP server becomes available, then the next time Routing and Remote Access needs IP addresses, DHCP addresses are

obtained. These DHCP addresses are allocated to remote access clients that connected after the addresses were obtained.

The remote access server uses a specific LAN interface to obtain DHCP-allocated IP addresses for remote access clients. You can select which LAN interface to use from the **IP** tab on the properties of a remote access router in the Routing and Remote Access snap-in. By default, **Allow RAS to select adapter** is selected, which means that Routing and Remote Access randomly picks a LAN interface to use.

**Static IP Address Pool**
When a static IP address pool is configured, the remote access server uses the first IP address in the first address range for the remote access server interface, and subsequent addresses are allocated to TCP/IP-based remote access clients as they connect. IP addresses freed due to disconnected remote access clients are reused.

If an address range in the static IP address pool is for off-subnet addresses, either enable an appropriate routing protocol on the remote access server or add the routes corresponding to the IP address ranges to the routers of your intranet. For more information, see "TCP/IP On-Subnet and Off-Subnet Addressing" earlier in this section.

**DNS and WINS Address Assignment**
As part of the IPCP negotiation, the remote access server assigns the IP addresses of DNS and WINS servers. The set of DNS and WINS server IP addresses that are assigned to the remote access client depends on the following factors:

- Whether DNS and WINS server IP address assignment is prohibited.
- Whether DNS and WINS server IP addresses for remote access clients are globally configured by using data stored in the registry.
- Whether the remote access server has more than one LAN interface.
- Whether the DNS and WINS server IP addresses for the remote access server are configured statically or are obtained through DHCP.

**Prohibiting DNS and WINS IP Address Assignment**
You can prevent the remote access server from assigning DNS and WINS IP addresses by changing a registry entry.

**Configuring Global DNS and WINS IP Address Assignment**
You can globally configure DNS and WINS server IP addresses for remote access clients by entering the IP addresses in the registry.

For more information about registry entries for dial-up remote access, see "Dial-up Remote Access Tools and Settings."

**Multiple LAN Interfaces**
If the DNS and WINS server IP address assignment is not prohibited or globally configured, then the remote access server allocates the DNS and WINS server IP addresses of a LAN interface on the remote access server to remote access clients. If there is only one LAN interface, which is the typical configuration of a dial-up remote access server, the remote access server allocates the DNS and WINS server IP addresses of the single LAN interface to remote access clients. If there is more than one LAN interface, the DNS and WINS server IP addresses of a specific LAN interface must be determined.

With multiple LAN interfaces, which is the typical configuration of a VPN remote access server, the remote access server, by default, picks one LAN interface randomly during startup and uses the DNS and WINS server IP addresses of the chosen LAN interface to allocate to remote access clients. To override this behavior, you can select the desired LAN interface through the **IP** tab on the properties of a remote access router in the Routing and Remote Access snap-in. By default, **Allow RAS to select adapter** is selected.

**Static Configuration or DHCP**
After the LAN adapter for DNS and WINS server IP address assignment has been determined:

- If the LAN adapter has a static IP configuration, then the IP addresses of the statically-configured DNS and WINS servers are allocated to remote access clients.
- If the LAN adapter obtained its IP configuration using DHCP, then the IP addresses of the DHCP-obtained DNS and WINS servers are allocated to remote access clients.

The following figure illustrates the way in which the remote access server determines the set of DNS and WINS server IP addresses to assign to remote access clients during IPCP negotiation.

**Determining DNS and WINS Server IP Address**




**Overriding IPCP-Allocated DNS and WINS Server IP Addresses with DHCPInform**

After IPCP is completed, Windows Server 2003 remote access clients send their remote access servers a DHCPInform message. DHCPInform is a DHCP message used by DHCP clients to obtain DHCP options. While PPP remote access clients do not use DHCP to obtain IP addresses for the remote access connection, Windows Server 2003 remote access client's use the DHCPInform message to obtain DNS server IP addresses, WINS server IP addresses, and a DNS domain name. The DHCPInform message is sent after the IPCP negotiation is concluded.

The DHCPInform message received by the remote access server is then forwarded to a DHCP server. The remote access server forwards DHCPInform messages only if it has been configured with the DHCP Relay Agent, as discussed in the following section. The response to the DHCPInform message is forwarded back to the requesting remote access client.

If the DHCPInform response contains DNS and WINS server IP address options, then these new values override what was allocated during IPCP. When the remote access client is a Windows Server 2003 remote access client and the DHCPInform response contains a DNS domain name, the DNS domain name is used as the per-adapter DNS suffix for the remote access connection of the remote access client. For more information on per-adapter DNS suffixes, see the "DNS Technical Reference."

**Remote Access Server and the DHCP Relay Agent**

To facilitate the forwarding of DHCPInform messages between remote access clients and DHCP servers, the remote access server uses the DHCP Relay Agent, a component of the Windows Server 2003 remote access router. To configure the remote access server to use the DHCP Relay Agent, add the **Internal** interface to the DHCP Relay Agent IP routing protocol with the Routing and Remote Access snap-in.

If the remote access server is using DHCP to obtain IP addresses for remote access clients, then the remote access server uses the DHCP Relay Agent to forward DHCPInform messages to the DHCP server of the selected LAN interface, which is found on the **IP** tab on the properties of a remote access router in the Routing and Remote Access snap-in.

If the remote access server is using a static IP address pool to obtain IP addresses for remote access clients, then the DHCP Relay Agent must be configured with the IP address of at least one DHCP server. Otherwise, DHCPInform messages sent by remote access clients are silently discarded by the remote access server.

**PPP Connection**

The following four phases must be completed successfully before the PPP connection is ready to transfer user data:

1. PPP configuration
2. Authentication
3. Callback
4. Protocol configuration

**Phase 1: PPP Configuration**

PPP configures PPP protocol parameters using LCP. During the initial LCP phase, each peer negotiates communication options that are used to send data and include:

- PPP parameters, such as MRU, Address and Control field compression, and Protocol ID compression.

- Which authentication protocols are used to authenticate the remote access client. An authentication protocol is selected but not implemented until the authentication phase.

- Multilink options.

**Preventing the Flag Character from Occurring in the Middle of a PPP Frame**
PPP employs two different methods to prevent the occurrence of the Flag character from appearing in the middle of a frame.

**PPP on Asynchronous Links**
On asynchronous links, such as analog phone lines, PPP uses a technique called character stuffing to prevent the occurrence of the Flag character within the PPP frame. If the Flag character (0x7E) occurs elsewhere in the PPP frame, the sender replaces it with the sequence 0x7D-5E. The 0x7D character is known as the PPP Escape character. If the PPP Escape character occurs, the sender replaces it with the sequence 0x7D-5D. The receiver translates 0x7D-5E sequences back to 0x7E and 0x7D-5D sequences back to 0x7D.

In addition, character values less than 0x20 can be modified to prevent the serial drivers from interpreting them as control characters. If negotiated by LCP, characters below 0x20 are modified by sending the sequence: 0x7D-[Original character with the 6th bit complemented]. For example, the byte 0x01 would be transmitted as 0x7D-0x21.

**PPP on Synchronous Links**
With synchronous links, such as T-Carrier, ISDN, or other digital links, a technique called bit stuffing is used to prevent the occurrence of the Flag character within the PPP frame. The Flag character is the bit sequence 01111110. Bit stuffing ensures that six 1 bits in a row occur only when the Flag character is sent. To accomplish this, bit stuffing sends the Flag character unmodified and elsewhere inserts a 0 bit whenever a sequence of five 1 bits occurs. Therefore, the bit sequence 111111 is encoded on the medium as 1111101 and the bit sequence 111110 is encoded as 1111100. Bit stuffing means that a byte can be encoded on the medium as more than eight bits, but the stuffed bits are added and removed by the synchronous link hardware.

**PPPoE Connection**
The PPPoE connection process has the following two phases:

1. A discovery phase, in which a client computer uses PPPoE frames to discover the presence of an AC, a device that terminates the cable modem or DSL connection and provides access to the Internet, and to determine a PPPoE session ID.

2. A PPP session phase, in which a PPP connection is established and used for data transfer in the same way as a dial-up or VPN-based PPP connection.

**Phase 1: PPPoE Discovery**
The PPPoE discovery phase consists of the following four PPPoE frames:

1. The PPPoE Active Discovery Initiation (PADI) frame is sent by the PPPoE client to the Ethernet broadcast address (0xFF-FF-FF-FF-FF-FF). Within the Ethernet payload, the Code field is set to 9; the Session ID is set to 0; and there is a single Service-Name PPPoE tag, as well as other tags as needed. If the network connection in the Network Connections folder corresponding to the broadband Internet adapter has been configured with a service name, that service name is sent. Otherwise, the PADI frame is sent with a null service name.

2. The PPPoE Active Discovery Offer (PADO) frame is sent by the AC to the unicast MAC address of the PPPoE client. Within the Ethernet payload, the Code field is set to 7; the Session ID is set to 0; and there are the AC-Name and Service-Name tags, as well as other tags as needed. If the network connection in the Network Connections folder corresponding to the broadband Internet adapter has not been configured with a service name, it is automatically set to the value of the Service-Name tag in the PADO frame.

3. The PPPoE Active Discovery Request (PADR) frame is sent by the PPPoE client to the unicast MAC address of the AC. Within the Ethernet payload, the Code field is set to 25; the Session ID is set to 0; and there is a Service-Name tag and other tags as needed.

4. The PPPoE Active Discovery Session (PADS) frame is sent by the AC to the unicast MAC address of the PPPoE client. Within the Ethernet payload, the Code field is set to 101; the Session ID field is set to the session ID for the PPP session of the PPPoE client; and there is a Service-Name tag, as well as other tags as needed.

To terminate the PPPoE session, either the PPPoE client or the AC can send a PPPoE Active Discovery Terminate (PADT) frame, which contains the Code field set to 167 and the session ID set to the session being terminated.

**Phase 2: PPP Session**

After the PPPoE discovery process is complete, a PPP connection is negotiated and network protocol data, such as IP datagrams, are sent over the PPPoE connection. The following figure shows a PPPoE frame that contains a PPP frame.

### Structure of a PPPoE Frame Containing a PPP Frame




Because of the additional PPPoE overhead, the maximum size of PPP frames that can be sent over a PPPoE connection is 1494 bytes.

**LCP Negotiation**

During the LCP negotiation, LCP packets are exchanged between PPP peers to negotiate a set of options and option values when sending data. The LCP negotiation is actually two separate dialogues between two PPP peers (Peer 1 and Peer 2):

1. Peer 1 asks, negotiates, and then receives confirmation of the LCP options that are used when sending data to Peer 2. This dialogue starts with Peer 1 sending a Configure-Request message and ends when Peer 2 sends a Configure-Ack message.

2. Peer 2 asks, negotiates, and then receives confirmation of the LCP options that are used when sending data to Peer 1. This dialogue starts with Peer 2 sending a Configure-Request message and ends when Peer 1 sends a Configure-Ack message.

Peer 1 and Peer 2 do not have to use the same set of LCP options.

When a PPP peer sends its initial Configure-Request, the response can be any of the following:

- A Configure-Nack message because one or more options have unacceptable values.

- A Configure-Reject message because one or more of the options are unknown or not negotiable.

- A Configure-Ack message because all of the options have acceptable values.

When a PPP peer receives a Configure-Nack message or Configure-Reject message, it sends a new Configure-Request message with modified options or option values. When a Configure-Ack message is received, the PPP peer is ready to send data.

The following figure shows a LCP negotiation process for Peer 1 using the options W, X, Y, and Z.

### Sample LCP Negotiation




The following sequence describes the LCP negotiation show in the figure:

1. Peer 1 sends a Configure-Request message requesting option W; option X set to 100; option Y set to 0; and option Z. Options W and Z are flag options.

2. Peer 2 does not understand option Z, so it sends a Configure-Reject message containing option Z.

3. Peer 1 sends a new Configure-Request message requesting option W; option X set to 100; and option Y set to 0.

4. Peer 2 prefers that option X be set to 200, so it sends a Configure-Nack message containing option X and its preferred value.

5. Peer 1 sends a new Configure-Request message requesting option W; option X set to 200; and option Y set to 0.

6. Peer 2 sends a Configure-Ack message.

Each time Peer 1 sends a new Configure-Request message, it changes the Identifier value in the LCP header so that Configure-Request messages can be matched with their responses.

This process describes only how Peer 1 sends data to Peer 2. A separate LCP negotiation must be performed so that Peer 2 can be configured to send data to Peer 1. Very often, the LCP packets for the two dialogues are intermixed during the connection process. Peer 1 is configuring the way it sends data at the same time as Peer 2.

**Phase 2: Authentication**
After LCP is complete, the authentication protocol agreed upon by the remote access server and the remote access client is implemented. The nature of this traffic is specific to the PPP authentication protocol.

**PAP Negotiation**
The use of PAP is negotiated during LCP negotiation by specifying the authentication protocol LCP option (type 3) and the authentication protocol 0xC0-23. After LCP negotiation is complete, PAP messages use the PPP protocol ID of 0xC0-23.

PAP is a simple exchange of messages:

1. The remote access client sends a PAP Authenticate-Request message to the remote access server containing the remote access clients user name and clear text.

2. The remote access server checks the user name and password and sends back either a PAP Authenticate-Ack message when the user's credentials are correct, or a PAP Authenticate-Nak message when the user's credentials are not correct.

**SPAP Negotiation**
The use of SPAP is negotiated during LCP negotiation by specifying the authentication protocol LCP option (type 3) and the authentication protocol 0xC0-27. After LCP negotiation is complete, SPAP messages use the PPP protocol ID of 0xC0-27.

Like PAP, SPAP is a simple exchange of messages:

1. The remote access client sends an SPAP Authenticate-Request message to the remote access server containing the remote access client's user name and encrypted password.

2. The remote access server decrypts the password, checks the user name and password, and sends back either an SPAP Authenticate-Ack message when the user's credentials are correct, or an SPAP Authenticate-Nak message with a reason why the user's credentials are not correct.

**CHAP Negotiation**
The use of CHAP is negotiated during Phase 1 by specifying the authentication protocol LCP option (type 3), the authentication protocol 0xC2-23, and the algorithm 0x05. After LCP negotiation is complete, CHAP messages use the PPP Protocol ID of 0xC2-23.

The challenge handshake authentication process for CHAP is an exchange of three messages:

1. The remote access server sends a CHAP Challenge message containing a session ID and an arbitrary challenge string.

2. The remote access client returns a CHAP Response message containing the user name in clear text and a hash of the challenge string, session ID, and the client's password using the MD5 one-way hashing algorithm.

3. The remote access server duplicates the hash and compares it to the hash in the CHAP Response. If the hashes are the same, the remote access server sends back a CHAP Success message. If the hashes are different, a CHAP Failure message is sent.

**MS-CHAP Negotiation**
The use of MS-CHAP is negotiated during LCP negotiation by specifying the authentication protocol LCP option (type 3), the authentication protocol 0xC2-23, and the algorithm 0x80. After LCP negotiation is complete, MS-CHAP messages use the PPP protocol ID of 0xC2-23.

The challenge handshake authentication process for MS-CHAP is an exchange of three messages:

1. The remote access server sends an MS-CHAP Challenge message containing a session ID and an arbitrary challenge string.

2. The remote access client returns an MS-CHAP Response message containing the user name in clear text and a hash of the challenge string, session ID, and the MD4 hash of the client’s password using the MD4 one-way hashing algorithm.

3. The remote access server duplicates the hash and compares it to the hash in the MS-CHAP Response. If the hashes are the same, the remote access server sends back an MS-CHAP Success message. If the hashes are different, an MS-CHAP Failure message is sent.

**MS-CHAP v2 Negotiation**
The use of MS-CHAP v2 is negotiated during LCP negotiation by specifying the authentication protocol LCP option (type 3), the authentication protocol 0xC2-23, and the algorithm 0x81. After LCP negotiation is complete, MS-CHAP messages use the PPP protocol ID of 0xC2-23.

The challenge handshake authentication process for MS-CHAP v2 is an exchange of three messages:

1. The remote access server sends an MS-CHAP v2 Challenge message to the remote access client that consists of a session identifier and an arbitrary challenge string.

2. The remote access client sends an MS-CHAP v2 Response message that contains:

    - The user name.

    - An arbitrary peer challenge string.

    - A one-way encryption of the received challenge string, the peer challenge string, the session identifier, and the user’s password.

3. The remote access server checks the MS-CHAP v2 Response message from the client and sends back an MS-CHAP v2 Response message containing:

    - An indication of the success or failure of the connection attempt.

- An authenticated response based on the sent challenge string, the peer challenge string, the client's encrypted response, and the user's password.

4. The remote access client verifies the authentication response and if it is correct, uses the connection. If the authentication response is not correct, the remote access client terminates the connection.

**Client Verification with EAP-MD5**
After the connection authentication phase is reached, the following process verifies the client:

1. The authenticator sends an EAP-Request message requesting the identity of the client.

2. The client sends its user ID to the authenticator as an EAP-Response message.

3. The authenticator sends an EAP-Request message containing the MD5 challenge string.

4. The client sends the MD5 hash of its user ID and password to the authenticator as an EAP-Response message.

5. On validation of the user ID and password, the authenticator sends a Success message to the client.

**Phase 3: Callback**
The Microsoft implementation of PPP includes an optional callback phase using CBCP immediately after authentication. The dial-in properties of the user account must be enabled for callback; otherwise, the remote access client cannot receive a call. Either the remote access client or the remote access server must specify the callback number.

If a connection is implementing callback, both PPP peers hang up and the remote access server calls the remote access client at the negotiated number.

**Phase 4: Protocol Configuration**
After the PPP is configured and callback is complete (optional), network layer protocols can be configured. With remote access on Windows 32-bit operating systems, the remote access server sends the remote access client Configuration-Request messages for all of the LAN protocols enabled for remote access on the remote access server. The remote access client either continues the negotiation of the LAN protocols enabled at the remote access client or sends an LCP Protocol-Reject message containing the Configuration-Request message.

IPCP and ATCP each go through a negotiation process very similar to LCP negotiation to configure their corresponding network layer protocol. CCP packets are exchanged to configure MPPC and MPPE.

**Connection Attempt Processing with Remote Access Policies**
To process a connection attempt, the parameters of the connection attempt are compared to the user name, password, and dial-in properties of the user account and the configured remote access policies.

Some general characteristics of remote access connection attempt processing are:

- If a connection attempt does not use a valid user name and password, then the connection attempt is denied.

- If there are no configured policies, then all connection attempts are denied.

- If the connection attempt does not match any of the remote access policies, then the connection attempt is denied.

- If the remote access permission of the user account for the remote access user is set to **Deny Access**, the connection attempt is always denied for that remote access user.

- The only time that a connection attempt is allowed is when it matches the conditions of a remote access policy; remote access permission is enabled either through the dial-in properties of the user account or through the remote access permission of the remote access policy (assuming the user's remote access permission is set to control access through remote access policies); and the parameters of the connection attempt match or conform to the parameters and conditions of the dial-in properties of the user account and the remote access policy profile properties.

The following figure illustrates the processing of remote access connection attempts using the dial-in properties of the user account and remote access policies. The figure assumes that the user name and password sent during the authentication process match a valid user account.

**Connection Attempt Processing**




For more information about remote access policies and IAS, see "IAS Technical Reference."

**Broadcast Name Resolution**

Enabling broadcast name resolution on a remote access server allows a remote access client to resolve full computer names and NetBIOS names of computers and other resources on a remote network that does not have either a DNS or a WINS server configured. This feature requires no configuration of the remote access clients; it is enabled by default on the remote access server.

When broadcast name resolution is enabled, the remote access server uses a NetBIOS over TCP/IP (NetBT) proxy to resolve names as follows:

1. A remote access client that must resolve a name to an IP address broadcasts a NetBIOS Name Query packet across all interfaces.

2. The remote access server receives the NetBIOS Name Query packet and checks its cache for the appropriate mapping of name to IP address.

   - If the remote access server has the appropriate mapping, it sends a NetBIOS Name Query Response packet to the remote access client.

   - If the remote access server does not have the appropriate mapping, it broadcasts the NetBIOS Name Query packet on all of its LAN interfaces. The resource whose name is being resolved sends a NetBIOS Name Query Response packet to the remote access client.

   After the remote access server routes the NetBIOS Name Query Response packet from the resource to the remote access client, the mappings remain in the cache of the server for a limited period of time. The default period of time is 10 minutes.

3. After this point, the remote access client uses the IP address of the resource to send packets to that resource.

As a result, a remote access client can resolve names of resources on a remote network as if the remote access client were on the same LAN as the resources.

**Note**

- By default, broadcast name resolution allows remote access clients to resolve names on a remote network, but it does not allow resources on a LAN to resolve the names of remote access clients. Remote access servers with broadcast name resolution enabled do not, by default, broadcast NetBIOS Name Query packets to interfaces that connect to the Internet. Unless you change this behavior, a remote access client cannot resolve the name of another remote access client. If you have a network that averages 10 connections or fewer, you might want to change this behavior.

**Related Information**

The following resources contain additional information that is relevant to this section.

- VPN Technical Reference
- IAS Technical Reference
- IPv4 Multicasting Technical Reference
- Demand Dial Routing Technical Reference
- DHCP Technical Reference
- DNS Technical Reference
- Connection Manager Technical Reference
- IETF RFC Database

# Dial-up Remote Access Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up Remote Access Tools and Settings**

**In this section**


- Dial-up Remote Access Tools
- Dial-up Remote Access Registry Entries


The "Dial-up Remote Access Tools" section describes tools that are useful for auditing, customizing, and troubleshooting dial-up remote access connections.

The "Dial-up Remote Access Registry Entries" section describes the registry entries that enable the configuration of Internet Protocol (IP) addresses and the storage of Point-to-Point Protocol (PPP) settings.

**Dial-up Remote Access Tools**

The following tools are associated with dial-up remote access.

**Connection Manager**

**Category**

Connection Manager is client connection software that is available as part of the operating system.

**Version compatibility**

Connection Manager is available in Windows Server 2003 and Windows Server 2000 and is compatible with Windows XP, Windows Millennium Edition, Windows 98 and Windows NT 4.0.

Connection Manager is client connection software that administrators can customize and distribute to users. The customized and distributed Connection Manager software includes a service profile, which contains all of the files required to install Connection Manager.

Connection Manager also contains the Connection Manager Administration Kit (CMAK) component and Connection Point Services (CPS).

**Connection Manager Administration Kit**

CMAK is a wizard that creates the Connection Manager service profile and then builds the service profile as a compressed, self-installing executable file that administrators can distribute to users.

**Connection Point Services**

CPS enables you to automatically distribute and update custom phone books. These phone books contain one or more point of presence (POP) entries, with each POP supplying a telephone number that provides dial-up access to an Internet access point. The phone books give users complete POP information, so when they travel they can connect to different Internet access points rather than being restricted to a single POP.

Without the ability to update phone books (a task CPS handles automatically), users would have to contact technical support to be informed of changes in POP information and to reconfigure their client dialer software.

CPS contains two major components:

- Phone Book Administrator (PBA)
- Phone Book Services (PBS)

**Phone Book Administrator**

PBA creates and maintains phone book files for use with Connection Managerservice profiles. Each phone book is a collection of POPs provided by either the administrator or one or more Internet service providers (ISPs) that the organization uses to provide Internet access. Each POP provides a local access number and connection settings for a specific region within a country/region or dependency.

**Phone Book Services**

PBS is an Internet Information Services (IIS) extension. When a service profile that is configured to check for phone book updates establishes a connection, Connection Manager queries the PBS server for a phone book update. PBS compares the version of the phone book that the service profile is using with the most recent files in the phone book database and then sends any updates to the service profile.

For more information about Connection Manager, see the "Connection Manager Technical Reference."

**Netdiag**
**Category**
Netdiag is a diagnostic tool available as part of the operating system as Netdiag.exe.

**Version compatibility**
Netdiag is available in Windows Server 2003, Windows XP, and Windows Server 2000.

Netdiag helps isolate networking and connectivity problems by performing a series of tests to determine the state of your network client. These tests and the network status information they expose give network administrators and support personnel a more direct means of identifying and isolating network problems. In addition, because this tool does not require parameters or switches to be specified, support personnel and network administrators can focus on analyzing the output rather than on training users about how to use the tool.

For more information about Netdiag, see "Windows Support Tools" in the Tools and Settings Collection.

**Netsh**
**Category**
Netsh is a command-line scripting tool available as part of the operating system.

**Version compatibility**
Netsh is available in Windows Server 2003, Windows 2000 Server, and all versions of Windows that include TCP/IP.

Netsh enables you to display or modify the network configuration of a local or remote computer. You can also use Netsh to enable the components in Windows Server 2003 to log tracing information to files. You can enable and disable tracing for specific components or for all components.

For more information about Netsh, see "Command-Line References" in the Tools and Settings Collection.

**Network Monitor**
**Category**
Network Monitor is a tool for monitoring, recording, and analyzing network traffic. Network Monitor is part of Microsoft Systems Management Server (SMS). A limited version, which can capture frames that are sent to or from the network adapter of the computer on which Network Monitor is installed, is included with Windows Server 2003 and can be installed as an optional networking component.

**Version compatibility**
Network Monitoris available in Windows Server 2003 and Windows Server 2000.

Network Monitor, a packet capture and analysis tool, can capture and view the traffic sent between a remote access server and a remote access client during the remote access connection process and during data transfer. Network Monitor captures all PPP packets sent over a serial link, including connection establishment and PPP-encapsulated user data.

### Note

- If compression or encryption are used, then the PPP payload is not interpreted by Network Monitor. Compressed or encrypted payloads are indicated by the PPP protocol ID of 3D (assuming protocol ID compression). To see the structure of user data within PPP payloads, disable compression and encryption.

To use Network Monitor to interpret unencrypted remote access traffic, you must be familiar with PPP. For more information about PPP and Network Monitor, see "Dial-up Remote Access Protocols" and "Dial-up Remote Access Processes and Interactions" in How Dial-up Remote Access Works. Network Monitor captures can be saved as files and sent to Microsoft Product Support Services for analysis.

To capture PPP packets with Network Monitor, you set the capture network to the network corresponding to the dial-up connection. You can use Network Monitor to:

- Troubleshoot the PPP connection establishment process.

- Ensure that PPP payloads are being encrypted.

- Ensure that PPP payloads are being compressed.

When using Network Monitor, keep the following in mind:

- Captured PPP frames do not contain a Flag character but do contain an Ethernet-like source address and destination address. This behavior is due to Network Monitor's receipt of packets from Ndiswan.sys, an intermediate NDIS driver that provides an IEEE 802.3 miniport interface to protocol drivers. Ndiswan.sys looks like an Ethernet adapter to protocols.

  For each PPP frame, the Ethernet-like source and destination addresses are both set to either **SEND** or **RECV** to indicate that the PPP frame was either sent or received by the computer on which the capture was taken. The **SEND** and **RECV** addresses do not necessarily identify the traffic of a remote access server or remote access client. If the capture was taken on the remote access server, then **SEND** frames were sent by the remote access server and **RECV** frames were sent by the remote access client. If the capture was taken on the remote access client, then **SEND** frames were sent by the remote access client and **RECV** frames were sent by the remote access server.

- Captured PPP frames contain an Address or Control field regardless of whether address and control field compression are negotiated.

- Protocol ID compression is usually negotiated with Microsoft PPP peers, making the PPP Protocol ID a single byte when possible.

- Use Network Monitor display to view only the traffic of desired protocols. For example, to view the Internet Protocol Control Protocol (IPCP) negotiation only, set the display filters to disable the display of all protocols except IPCP.

**Sample PPP Connection Captured with Network Monitor**
The following printout is an example of a PPP connection establishment process captured with Network Monitor showing only the frame summaries. The entries are indented to improve readability.

```
1 8.726 SEND SEND LCP
        Config Req Packet, Ident = 0x00, Length = 36
2 8.796 RECV RECV LCP
        Config Req Packet, Ident = 0x00, Length = 25
3 8.796 SEND SEND LCP
        Config Ack Packet, Ident = 0x00, Length = 25
4 8.816 RECV RECV LCP
        Config Reject Packet, Ident = 0x00, Length = 17
5 8.816 SEND SEND LCP
        Config Req Packet, Ident = 0x01, Length = 23
6 8.886 RECV RECV LCP
        Config Ack Packet, Ident = 0x01, Length = 23
7 8.886 SEND SEND LCP
        Ident Packet, Ident = 0x02, Length = 18
8 8.886 SEND SEND LCP
        Ident Packet, Ident = 0x03, Length = 23
9 8.886 RECV RECV PPPCHAP
        Challenge, ID = 0x 1: Challenge
10 8.886 SEND SEND PPPCHAP
        Challenge, ID = 0x 1: Response, administrator
11 8.976 RECV RECV PPPCHAP
        Challenge, ID = 0x 1: Success
12 8.976 RECV RECV CBCP
        Callback Request, Ident = 0x01
13 8.976 SEND SEND CBCP
        Callback Response, Ident = 0x01
14 8.996 RECV RECV CBCP
        Callback Acknowledgement, Ident = 0x01
15 8.996 SEND SEND CCP
        Configuration Request, Ident = 0x04
16 8.997 SEND SEND IPCP
        Configuration Request, Ident = 0x05
17 8.997 RECV RECV CCP
        Configuration Request, Ident = 0x01
```

```
18 9.017 RECV RECV IPCP
        Configuration Request, Ident = 0x02
19 9.037 RECV RECV NBFCP
        Configuration Request, Ident = 0x04
20 9.147 SEND SEND CCP
        Configuration Acknowledgement, Ident = 0x01
21 9.147 SEND SEND IPCP
        Configuration Acknowledgement, Ident = 0x02
22 9.167 SEND SEND LCP
        Protocol Reject Packet, Ident = 0x07, Length = 32
23 9.237 RECV RECV CCP
        Configuration Reject, Ident = 0x04
24 9.237 RECV RECV IPCP
        Configuration Reject, Ident = 0x05
25 9.237 SEND SEND IPCP
        Configuration Request, Ident = 0x08
26 9.287 RECV RECV IPCP
        Configuration No Acknowledgement, Ident = 0x08
27 9.287 SEND SEND IPCP
        Configuration Request, Ident = 0x0A
28 9.327 RECV RECV IPCP
        Configuration Acknowledgement, Ident = 0x0A
29 10.729 SEND SEND CCP
        Configuration Request, Ident = 0x04
30 10.960 RECV RECV CCP
        Configuration Reject, Ident = 0x04
31 10.960 SEND SEND CCP
        Configuration Request, Ident = 0x0B
32 10.960 RECV RECV CCP
        Configuration Acknowledgement, Ident = 0x0B
```

The trace was captured on the remote access client. Therefore, the **SEND** frames were sent from the remote access client and the **RECV** frames were sent from the remote access server. In this trace, you can see the four phases of the establishment of the PPP connection:

- Phase 1: PPP configuration is performed in frames 1 through 8 by using the exchange of Link Control Protocol (LCP) configuration packets.
- Phase 2: Authentication, in which the user's credentials are verified, is performed in frames 9 through 11.
- Phase 3: Callback is performed in frames 12 through 14.
- Phase 4: Protocol configuration, in which compression, encryption, and IP are configured, is performed in frames 15 through 38.

In addition to the summary view, Network Monitor can also expand frames for detailed analysis. For example, frame 1 from this trace is displayed as:

```
FRAME: Base frame properties
FRAME: Time of capture = Nov 18, 1998 15:23:6.967
FRAME: Time delta from previous physical frame: 0 milliseconds
FRAME: Frame number: 1
FRAME: Total frame length: 50 bytes
FRAME: Capture frame length: 50 bytes
FRAME: Frame data: Number of data bytes remaining = 50 (0x0032)
PPP: Link Control Protocol Frame (0xC021)
PPP: Destination Address = SEND_
PPP: Source Address = SEND_
PPP: Protocol = Link Control Protocol
LCP: Config Req Packet, Ident = 0x00, Length = 36
LCP: Code = Configuration Request
LCP: Identifier = 0 (0x0)
LCP: Length = 36 (0x24)
LCP: Options: ASYNC.MAP:00 00 00 00-MAGIC#:0x0C05-PROT.COMP- ADR/CF.COMP-
CALL.BACK:Unkn---
LCP: ASYNC.MAP:00 00 00 00
LCP: Option Type = Async Control Character Map
LCP: Option Length = 6 (0x6)
```

```
LCP: Async Control Character Map = 00 00 00 00
LCP: MAGIC#:0x0C05
LCP: Option Type = Majic Number
LCP: Option Length = 6 (0x6)
LCP: Magic Number = 3077 (0xC05)
LCP: PROT.COMP
LCP: Option Type = Protocol Field Compression
LCP: Option Length = 2 (0x2)
LCP: ADR/CF.COMP
LCP: Option Type = Address and Control Field Compression
LCP: Option Length = 2 (0x2)
LCP: CALL.BACK:Unkn
LCP: Option Type = Callback
LCP: Option Length = 3 (0x3)
LCP: CallBack = 0x06
LCP: Multilink Maximum Receive Reconstructed Unit
LCP: Option Type = 0x11
LCP: Option Length = 4 (0x4)
LCP: Multilink Endpoint Discriminator
LCP: Option Type = 0x13
LCP: Option Length = 9 (0x9)
```

**PPP Tracing**

**Category**

PPP tracing is an operating system tool available in Routing and Remote Access.

**Version compatibility**

PPP tracing is supported in Windows Server 2003 and Windows 2000 Server. PPP tracing is the same as the PPP logging feature in Windows NT 4.0 and earlier.

To use PPP tracing to troubleshoot PPP connection failures, you must be familiar with PPP. For more information about PPP and PPP tracing, see "Dial-up Remote Access Protocols" and "Dial-up Remote Access Processes and Interactions" in How Dial-up Remote Access Works.

Tracing is a facility of Windows Server 2003 remote access and routing components that allows you to optionally enable and disable the recording of programming code and network events to a file. PPP tracing logs the PPP packets exchanged during the PPP connection establishment process and is valuable for troubleshooting the failure of a PPP connection.

By default, the PPP log is stored in the Ppp.log file in the *systemroot*\Tracing folder. The PPP log generated by PPP tracing contains the programming calls and contents of PPP packets for PPP control protocols. PPP tracing cannot be used to view PPP user data sent across the connection.

**Sample PPP Trace**

The following printout is an excerpt from a PPP trace of a PPP connection establishment process. The entries are indented to improve readability.

```
[1472] 15:57:50:094: Line up event occurred on port 5
[1472] 15:57:50:104: Starting PPP on link with IfType=0x0,IPIf=0x0
[1472] 15:57:50:104: RasGetBuffer returned ae70054 for SendBuf
[1472] 15:57:50:104: FsmInit called for protocol = c021, port = 5
[1472] 15:57:50:104: ConfigInfo = 273e
[1472] 15:57:50:104: APs available = 1
[1472] 15:57:50:104: FsmReset called for protocol = c021, port = 5
[1472] 15:57:50:104: Inserting port in bucket # 5
[1472] 15:57:50:104: Inserting bundle in bucket # 6
[1472] 15:57:50:104: FsmOpen event received for protocol c021 on port 5
[1472] 15:57:50:104: FsmThisLayerStarted called for protocol = c021,
            port = 5
[1472] 15:57:50:104: FsmUp event received for protocol c021 on port 5
[1472] 15:57:50:104: <PPP packet sent at 11/04/1998 23:57:50:104
[1472] 15:57:50:104: <Protocol = LCP, Type = Configure-Req, Length =
            0x2f, Id = 0x0, Port = 5
[1472] 15:57:50:104: <C0 21 01 00 00 2D 02 06 00 00 00 00 03 05 C2 23
            |.!...-........#|
[1472] 15:57:50:104: <80 05 06 72 5F 50 9A 07 02 08 02 0D 03 06 11 04
            |...r_P.........|
[1472] 15:57:50:104: <06 4E 13 09 03 00 60 08 3E 46 07 17 04 00 03 00
            |.N....`.>F......|
[1472] 15:57:50:104: InsertInTimerQ called portid=6,Id=0,Protocol=c021,
            EventType=0,fAuth=0
[1472] 15:57:50:104: InsertInTimerQ called portid=6,Id=0,Protocol=0,
```

```
            EventType=3,fAuth=0
[1472] 15:57:50:104: >PPP packet received at 11/04/1998 23:57:50:104
[1472] 15:57:50:104: >Protocol = LCP, Type = Configure-Req, Length =
            0x26, Id = 0x0, Port = 5
[1472] 15:57:50:104: >C0 21 01 00 00 24 02 06 00 00 00 00 05 06 00 00
            |.!...$..........|
[1472] 15:57:50:104: >C0 05 07 02 08 02 0D 03 06 11 04 06 4E 13 09 03
            |._..........N...|
[1472] 15:57:50:104: >00 60 08 52 F9 D8 00 00 00 00 00 00 00 00 00 00
            |.`.R............|
```

The last three lines are a hexadecimal display of the same LCP packet in frame 1 of the preceding Network Monitor trace. To understand this frame, you must manually parse this frame according to the PPP and LCP packet structure. An example of the parsing of this PPP frame is listed in the following table.

**Parsing of the LCP Configure-Request**

| Bytes | Meaning |
|---|---|
| C0 21 | PPP Protocol ID for LCP. |
| 01 | LCP code for a Configure-Request. |
| 00 | LCP identifier for this Configure-Request. |
| 00 24 | Length, in bytes, of the LCP packet (36 bytes long). |
| 02 | LCP option for Asynchronous Control Character Map (ACCM). |
| 06 | Length, in bytes, of the ACCM option. |
| 00 00 00 00 | Data for the ACCM option. |
| 05 | LCP option for the magic number. |
| 06 | Length, in bytes, of the magic number option. |
| 00 00 C0 05 | Data for the magic number option. |
| 07 | LCP option for protocol compression. |
| 02 | Length, in bytes, of the protocol compression option. |
| 08 | LCP option for address and control field compression. |
| 02 | Length, in bytes, of the address and control field compression option. |
| 0D | LCP option for callback. |
| 03 | Length, in bytes, of the callback option. |
| 06 | Callback option data. |
| 11 | LCP option for the Multilink Maximum Receive Reconstructed Unit. |
| 04 | Length, in bytes, of the Multilink Maximum Receive Reconstructed Unit option. |
| 06 4E | Multilink Maximum Receive Reconstructed Unit option data. |
| 13 | LCP option for the Multilink Endpoint Discriminator option. |
| 09 | Length, in bytes, of the Multilink Endpoint Discriminator option. |
| 03 00 60 08 52 F9 D8 | Multilink Endpoint Discriminator option data. |

**Rasdiag Category**

Rasdiag is a diagnostic tool for remote access clients that is available as part of the Windows XP operating system as Rasdiag.exe.

**Version Compatibility**
Rasdiag is available in Windows XP.

Rasdiag collects diagnostic information about remote services and places that information in a file. You can use this tool to work with Microsoft Product Support Services to troubleshoot remote connection issues by taking a snapshot of the configuration data and capturing an attempted remote connection.

Rasdiag is not described in the Support Tools section of Windows XP Help. For more information about Rasdiag, see "Windows Support Tools" in the Tools and Settings Collection.

**Rasdial**
**Category**
Rasdial is a command-line tool that automates the process of dialing up a network connection and is available as part of the operating system.

**Version Compatibility**
Rasdial is available in Windows Server 2003, Windows XP, and Windows Server 2000.

Rasdial creates a simple batch file that connects to a dial-up network connection by specifying an entry in the current rasphone.pbk file, located in the ../Documents and Settings/*username*/ folder for an individual user connection, and the ../Documents and Settings/All Users/ folder for all user connections. You can also use Rasdial to connect to or disconnect from a network connection.

Rasdial is different from Autodial:

- Rasdial does not automatically detect your network connection status before dialing a connection and does not depend on the Remote Access Auto Connection Manager service. Rasdial requires that there is an existing phone book entry that can be used or modified.

Rasdial is different from Rasphone:

- Rasdial provides unattended dialing. It does not support the following features available in Rasphone: prefix/suffix, auto-redial, change password, retry authentication, statistics, operator-assisted dialing, and connections requiring Terminal mode input.

For more information about Rasdial, see "Windows Support Tools" in the Tools and Settings Collection.

**Dial-up Remote Access Registry Entries**
The following registry entries are associated with dial-up remote access.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_LOCAL_MACHINE\SYSTEM**
The following registry key is located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services \RemoteAccess\Parameters\Ip

**Ip**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services \RemoteAccess\Parameters\

**Version**
The Ip registry setting is available in Windows Server 2003 and Windows 2000 Server.

When the remote access server is configured to use Dynamic Host Configuration Protocol (DHCP) to obtain IP addresses, Routing and Remote Access instructs the DHCP client component to obtain 10 IP addresses from a DHCP server. When all 10 IP addresses are used, the remote access server uses the DHCP client component to obtain 10 more. You can modify the number of IP addresses obtained at a time by changing the value of **InitialAddressPoolSize**.

If you do not want the remote access server to assign Domain Name System (DNS) and Windows Internet Name Service (WINS) IP addresses, set the values of **SuppressDNSNameServers** and **SuppressWINSNameServers** to **1**.

To globally configure DNS and WINS server IP addresses for remote access clients, enter the IP addresses in the values of **DNSNameServers** and **WINSNameServer**.

**HKEY_LOCAL_MACHINE\SYSTEM**
The following registry key is located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP

**PPP**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\

**Version**
The PPP registry setting is available in Windows Server 2003 and Windows 2000 Server.

Routing and Remote Access stores PPP settings at this location.

# Telnet Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

All versions of Microsoft Windows Server 2003 operating systems include Telnet Client and Telnet Server components. Using Telnet Client and Server, you can create a remote command console session on a host. You can run command line programs, shell commands, and scripts in a remote command console session just as though you were locally logged on to the host and using a local command prompt window.

This subject describes how you can use Telnet for remote administration, the functions that Telnet performs, and how Telnet servers and clients interact.

**In this subject**


- What Is Telnet?
- How Telnet Works
- Telnet Tools and Settings

# What Is Telnet?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 and Windows XP Professional include Telnet Client and Telnet Server, which allow users to make remote connections based on the Telnet protocol. Using Telnet Client and Server, you can create a remote command console session on a host. Using a local command prompt window, you can run command line programs, shell commands, and scripts in a remote command console session just as though you were locally logged on to the host. Thus, having Telnet client and server software solves two networking problems. It allows for interoperability between disparate operating systems, and it facilitates administration of remote systems, saving administrative time and network resources.

**Interoperability**
Telnet Client and Server are well suited for troubleshooting and configuring remote computers, especially in mixed environments that require interoperability between different operating systems. For example, you can use Windows Server 2003 Telnet Client to connect to a Telnet server that is running on another operating system such as UNIX. Likewise, you can use a Telnet client that is running on UNIX to connect to a computer running Windows Server 2003 Telnet Server. Telnet Client and Server are also ideal in situations where memory and processor resources are minimal on a client or host or where network bandwidth is limited. This is because computers running Telnet clients and servers use less memory and processor time than other remote management tools, and Telnet clients and servers transmit only plaintext (unencrypted characters) across the network.

**Remote Administration**
Remote administration is a method of managing one or more remote computers from a single location. In a large organization, you can use remote administration to centrally manage hundreds or even thousands of computers located in other buildings or even in other cities. In a small organization, you can use remote administration to manage a single server located in an adjacent office.

By its nature, remote administration lowers the total cost of ownership (TCO) by making system management easier and more efficient. Using remote administration, server operators and technicians can manage and troubleshoot servers without having to locally log on to the server, thereby lowering the cost of on-site support. Remote administration also assists help desk technicians in solving problems more quickly by letting them take control of a user's computer.

**Telnet and Other Technologies**
Telnet is closely related to two other technologies that provide an alternative for remote command console sessions, depending on your needs for distributed computing, security, configuration requirements, and so on.

**Windows Server 2003 Terminal Services**
Telnet and Terminal Services are similar in that they are both used for remote sessions. However, Terminal Services extends the model of distributed computing by allowing client computers to operate in a server-based computing environment. Whereas Telnet only allows terminal emulation between a Telnet server and client, Terminal Services running on Windows Server 2003-based computers allows clients to run applications, while data processing, and data storage occur on the server. Applications and user sessions are transmitted over the network and displayed via terminal emulation software. Similarly, print streams, keyboard input, and mouse clicks are also transmitted over the network between the server and the terminal emulation software. Each user logs on and sees only their individual session, which is managed transparently by the server operating system and is independent of any other client session. You might want to consider using Windows Server 2003 Terminal Services if you require more extensive distributed computing.

**Windows Services for UNIX 3.5**
Windows Services for UNIX 3.5 includes a Telnet character-mode client that provides functionality not included in Windows Server 2003. The Windows Services for UNIX character-mode client supports both stream mode and console mode. It also provides for logging and additional configuration settings.

Windows Services for UNIX 3.5 includes two Telnet servers:

- The default, Windows-based Telnet server, which is functionally similar to the one included with all versions of Microsoft Windows since Windows 2000

- The Interix telnetd

Only one of these Telnet servers can be enabled at a time. By default, neither Telnet server is enabled for security reasons.

The Windows Services for UNIX 3.5 Telnet server accepts logons from a variety of clients, including the Telnet clients shipped with Windows 2000, Windows NT, Windows 95, and Windows 98, as well as a variety of character-mode terminal clients from virtually any operating system. Additionally, it can be configured to meet specific site requirements to improve security, simplify logons, support stream or console mode, and so forth.

The Windows Services for UNIX 3.5 Telnet server should be familiar to users of Windows 2000, Windows XP, and Windows Server 2003. It is essentially the same as the server included in Windows XP Professional and Windows Server 2003, and it is very similar to the one included in Windows 2000 editions. It uses the Windows command shell (Cmd.exe) as the default shell. You can start and stop this server from either the Services MMC (Services.msc) or from the Windows Services for UNIX Administration MMC (Sfumgmt.msc). If you are using Windows and UNIX in a mixed networking environment, you might want to consider installing Windows Services for UNIX 3.5 to extend the capabilities of Telnet.

# How Telnet Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Understanding Telnet
- Telnet Architecture
- Telnet Limitations
- Related Information


All Windows versions include Telnet Client and Telnet Server components. Using Telnet Client and Server, you can create a remote command console session on a host. You can run command line programs, shell commands, and scripts in a remote command console session just as though you were locally logged on to the host and using a local command prompt window.

Windows Server 2003 Telnet Client and Server are well suited for troubleshooting and configuring remote computers, especially in mixed environments that require interoperability between different operating systems. For example, you can use Telnet Client to connect to a Telnet server that is running on another operating system such as UNIX. Likewise, you can use a Telnet client that is running on UNIX to connect to a computer running Telnet Server. Windows Server 2003 Telnet Client and Server are also ideal in situations where memory and processor resources are minimal on a client or host or where network bandwidth is limited. This is because computers running Telnet clients and servers use less memory and processor time than other remote management tools, and Telnet clients and servers transmit only plaintext (unencrypted characters) across the network.

**Understanding Telnet**

Before using the Windows Server 2003 Telnet tools, you should consider the following:

- Windows Server 2003 Telnet Client and Server are based on the Telnet protocol, which specifies a method for transmitting and receiving unencrypted ASCII characters (plaintext) across a network. Understanding how the protocol works, and how Telnet clients and servers use the Telnet protocol, helps you manage Telnet connections.

- The Windows Server 2003 Telnet tools have several inherent limitations that affect the types of remote management tasks you can perform and the level of security that is in effect when you perform those tasks. Understanding these limitations helps you determine when and when not to use the Telnet tools.

- You can configure Telnet Server settings by using the Windows Server 2003 Telnet administration tool (Tlntadmn.exe) and the registry editor (Regedit.exe). Although the default Telnet Server settings are sufficient for most Telnet client connections, you might need to change the default settings to better suit your organization. Examples of Telnet Server settings include: authentication type, default port assignment for Telnet connections, maximum number of client connections, and maximum number of failed logon attempts.

- By default, members of the local administrators group can log on to a Telnet server. However, you might not want all Telnet users to have full administrative control of the host they log on to. In this case, you can use a Telnet clients group to grant users Telnet logon rights without granting them any administrative rights on the host. To configure these user rights from the graphical user interface, you must use the Active Directory Users and Groups snap-in or the Local Users and Groups snap-in. You can also use the **Net User** and **Net Group** commands to configure user rights from the command line.

- You can configure several optional settings when you use Telnet Client to establish a Telnet session on a host. Depending on the type of Telnet server you are logging on to, and how the Telnet server is configured, you might need to enable or change some of these optional settings. Examples of Windows Server 2003 Telnet Client settings include: client-side logging, terminal type, port assignment, and alternate user name for logon.

- You can manage active Telnet sessions on a host by using the Windows Server 2003 Telnet administration tool. Some of the administrative tasks you can perform include: terminating Telnet sessions, sending console messages to users with active Telnet sessions, and listing Telnet session information (for example, user name, logon time, idle time, and client IP address).

### Note

- The information in this document refers to the Telnet Client and Telnet Server components that are installed by default with Windows Server 2003 and Windows XP Professional operating systems.

**Telnet Architecture**

Most network operating systems provide a Telnet client and a Telnet server. Telnet clients and servers are small executable programs that allow a local computer (a client) to access services and programs on a remote computer (a host). Telnet clients and servers, including Windows Server 2003 Telnet Client and Telnet Server, are based on the Telnet protocol, which is a subset of the TCP/IP suite and is described in RFC 854. The Telnet protocol specifies two general mechanisms: how Telnet clients and servers establish a connection across a network and how they transmit and receive information across a network.

You usually run a Telnet client program on a local computer: for example, a workstation that you are logged on to. You usually run a Telnet server program on a remote computer: for example, a host you want to administer. Telnet client programs initiate connections with Telnet servers. Telnet servers run in the background on a host, listening for Telnet clients to request a connection.

**Common Telnet Features**

Because Telnet clients and servers are based on the same standard protocol, all Telnet clients and servers have several features in common. These common features are what make Telnet clients and servers well suited for performing remote administration tasks in environments that require interoperability among disparate operating systems. In short, the Telnet protocol makes it possible for you to connect a computer running Windows Server 2003 Telnet Client to a UNIX Telnet server. The key features that make this interoperability possible include the following:

**Common communication protocols**

All Telnet clients and servers use TCP/IP as the underlying communication protocol. This makes Telnet clients and servers particularly useful for remotely administering computers across the Internet or within wide area networks (WANs) that are connected to the Internet. If your network does not support TCP/IP, you will not be able to use a Telnet client or server.

**Common communication ports**

TCP port 23 is reserved for Telnet client and server communication. By default, most Telnet clients initiate communication on port 23, and most Telnet servers listen on port 23 for connection requests. You can change the default port assignments with some Telnet client and server software, such as Windows Server 2003 Telnet Client and Server, but port 23 is the universally accepted port for Telnet communication.

**Common character set for communication**

All Telnet clients and servers are capable of transmitting and receiving a predefined character set consisting of standard ASCII character codes and ASCII control codes. All Telnet clients and servers transmit the ASCII codes across a network in unencrypted form (plaintext).

**Common implementation of the Network Virtual Terminal**

All Telnet clients and servers implement a network virtual terminal (NVT). The NVT is responsible for translating operating system-specific instructions (keyboard codes or display codes) into a consistent set of codes that all Telnet clients and servers can transmit and receive. The NVT is what makes Telnet clients and servers capable of communicating with each other regardless of which operating system they are using.

**Creating a Telnet Connection**

To create a Telnet connection between a Telnet client and a Telnet server, you must do the following:

- **Start the Telnet Server program on the host.** On Windows Server 2003, Telnet Server (Tlntsvr.exe) runs as a service. You can start the service manually every time you want to connect to a host, or you can configure the service so that it starts every time your computer starts. Telnet clients cannot connect to a host unless a Telnet server program (or service) is running and listening for connection requests.

- **Run the Telnet Client program on the local computer.** When you run Windows Server 2003 Telnet Client (Telnet.exe), you must specify the host to which you want to connect. You can also configure several optional connection

settings and features.

When you run a Telnet client, it makes a connection request to the host. If a Telnet server responds to the request, the Telnet client and server negotiate the details of the connection, such as flow control settings, window size, and terminal type. After the connection details are successfully negotiated, and logon credentials are validated, the Telnet server program creates a Telnet command console session.

On Windows Server 2003, each Telnet command console session consists of two processes: Tlntsess.exe and Cmd.exe. Tlntsess.exe is responsible for managing the Telnet session. Cmd.exe is the command interpreter, or shell program, that runs commands, programs, or scripts on the host.

**Note**

- Cmd.exe is the default command interpreter for a Windows Server 2003 Telnet command console session. However, you can configure the Windows Server 2003 Telnet Server program to use as a default any command interpreter or shell program that is installed on the host.

**Running Programs Remotely Using a Telnet Connection**

After you establish a Telnet connection with Telnet Server, the following message appears in the command prompt window on the client:

```
*= = = =
```

# Telnet Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Telnet Tools
- Telnet Registry Entries


Telnet Server tools and settings determine how Telnet Server handles auditing, authentication, idle session time-out, and other remote command console session options. Usually, you do not need to configure Telnet Server options to connect a Telnet client to Windows Server 2003-based Telnet Server: the default Telnet Server options are compatible with most Telnet clients. However, you must configure Telnet Server options if you want to do any of the following:

- Audit logon and logoff information.
- Disable NTLM or password authentication, or change the default domain for authenticating unqualified user names (by default, the domain in which the machine account resides is used to authenticate unqualified user names).
- Prohibit authentication of user accounts in trusted domains, which restricts Telnet access to users whose user accounts are stored only in the local Security Accounts Manager (SAM) database (by default, Telnet Server authenticates user accounts in trusted domains and the local SAM database).
- Change the default shell, or command interpreter, that is used for Telnet sessions (Cmd.exe is the default shell).
- Specify an IP address on which you want the Telnet Server program to listen for connection requests.
- Change the mode of operation from console mode to stream mode.
- Ensure that all programs started in a Telnet session terminate when you disconnect a Telnet session.
- Change the TCP port on which Telnet Server listens for a connection (by default, Telnet servers listen on TCP port 23).
- Change the maximum number of Telnet sessions that Telnet Server will accept (the default is 2).
- Change the maximum number of logon attempts before a user is disconnected (the default is 3).
- Disable idle session time-out, or change the idle session time-out value (the default is 1 hour).
- Disable Alt key mapping (by default, pressing Ctrl-A simulates the Alt key).

**Telnet Tools**
The following tools are associated with Telnet Server.

**Telnet.exe: Telnet Command Prompt**
**Category**
The Telnet command prompt tool is included with the Windows Server 2003 and Windows XP operating systems.

**Version compatibility**
Use this command on computers running Windows Server 2003 or Windows XP.

Once all of the settings and options are configured, you can use Telnet.exe to initiate and conduct a Telnet session. You can create a Telnet connection, configure Telnet.exe options, and use all Telnet.exe features by using the Telnet command prompt.

The Telnet command prompt is useful if you are performing quick maintenance tasks on several different hosts or you need to use advanced Telnet options and features.

You can access the Telnet command prompt by running the **Telnet** command without any command-line parameters. You can also access the Telnet command prompt by typing the Telnet escape character during an active Telnet session. The default escape character is Ctrl+].

After you start the Telnet command prompt, the following message appears:

```
Welcome to Microsoft Telnet Client
Escape Character is 'Ctrl+]'
Microsoft Telnet >
```

You can close the Telnet command prompt by using the **Quit** command.

**Telnet.exe with command-line parameters**

You can create a Telnet connection and configure some Telnet.exe options by using the **Telnet** command in conjunction with various command-line parameters. Using the **Telnet** command with command-line parameters is helpful if you are creating Telnet connections within a script or batch file or you do not need to use advanced Telnet client options and features. When you use **Telnet** with command-line parameters, you can use a single command to create a connection with a host. The command-line syntax for Telnet.exe is:

```
telnet [-a][-eescape_char][-flog_file][-luser_name][-tterm]host [port]
```

The command-line parameters are described in the following table.

**Telnet.exe Command-Line Parameters**

| Parameter | Description |
|---|---|
| **-a** | Instructs Telnet.exe to log on to the host using the credentials of the user who is currently logged on to the client. |
| **-e** *escape_char* | Specifies an escape character, which displays the Telnet command prompt. The default escape character is Ctrl+]. |
| **-f** *log_file* | Creates a client-side log file and turns on client-side logging for the current session. The *log_file* parameter must consist of a path and file name. |
| **-l** *user_name* | Instructs Telnet.exe to log on to the host using the user account that is specified in *user_name*. The user account specified in *user_name* must have Telnet logon rights on the host. |
| **-t** *term* | Specifies the terminal type. The default terminal type is ANSI. Other valid terminal types include VT52, VT100, and VTNT. |
| *host* | Specifies the host with which you want to create a Telnet connection. The *host* parameter can be a NetBIOS name, a fully qualified domain name, or an IP address. |
| *port* | Specifies the TCP port on which you want to create a Telnet connection. The default Telnet port is 23. |

For example, the following command uses the credentials of the user who is currently logged on to the client to create a Telnet connection on port 23 with a host named server01:

```
Telnet Server01
```

Likewise, the following example creates the same Telnet connection and enables client-side logging to a log file named c:\telnet_logfile:

```
telnet -f c:\telnet_logfile server01
```

The connection with the host remains active until you exit the Telnet session (by using the **Exit** command), or you use the Telnet Server administration tool to terminate the Telnet session on the host.

**TIntadmn.exe: Telnet Administration**

**Category**
The Telnet administration command-line tool is included with the Windows Server 2003 and Windows XP operating systems.

**Version compatibility**
Run this command on computers running Windows Server 2003 or Windows XP.

Tlntadmn.exe is a command-line tool, and is installed by default when you install Windows Server 2003. Unlike the Telnet Server administration tool in Windows 2000, Tlntadmn.exe is a noninteractive tool that must be run with various command-line parameters. Because Tlntadmn.exe is noninteractive, and uses command-line parameters to configure Telnet Server options, you can use Tlntadmn.exe in scripts or batch files to automate Telnet Server configuration tasks. You must be a member of the Administrators local group to use the Telnet Server administration tool.

For more information about **Tlntadmn.exe**, see "Telnet commands" in the Command Line References in the Tools and Settings Collection.

**Telnet Registry Entries**
The following registry entries are associated with Telnet.

You can configure most Windows Server 2003 Telnet Server options by using the Telnet Server administration tool (Tlntadmn.exe). However, you can only configure some options by using the registry editor (Regedit.exe). You should use the registry editor to configure only the Telnet Server options that cannot be configured with the Telnet Server administration tool.

You must be a member of one of the following groups to use Regedit.exe: Administrators, Server Operators, and Power Users. You can use Regedit.exe to change registry settings on a local or a remote computer. However, only members of the Administrators group can use Regedit.exe to configure registry settings on a remote computer.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft**
The following registry entries are located under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\TelnetServer\1.0\.

**AllowTrustedDomain**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\TelnetServer\1.0.\

**Version**
The AllowTrustedDomain entry is included in Windows Server 2003 and Windows XP.

You can prevent the Telnet Server program from authenticating users on trusted domains by configuring this registry entry.

By default, the Telnet Server program authenticates user accounts in trusted domains and in the local SAM database. Preventing the Telnet Server program from authenticating user accounts in trusted domains restricts Telnet access to only those users whose user accounts are in the local SAM database. By default, the AllowTrustedDomain registry entry has a value of 1. To prevent Telnet Server from authenticating user accounts in trusted domains, you must set this registry entry to 0.

**DefaultShell**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\TelnetServer\1.0.\

**Version**
The DefaultShell entry is included in Windows Server 2003 and Windows XP.

You can change the default shell, or command interpreter, that the Telnet Server program uses for a Telnet session by configuring this registry entry.

By default, Telnet Server runs all commands in the Windows Server 2003-based command interpreter (Cmd.exe). You can change this to any command interpreter that is installed on the host. You must provide a path and file name for the command interpreter.

**ListenToSpecificIpAddr**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\TelnetServer\1.0.\

**Version**
The ListenToSpecificIpAddr entry is included in Windows Server 2003 and Windows XP.

You can configure the Telnet Server program so it listens for connection requests that are sent to a specific IP address. This is useful if a host has several network adapters, and you want to limit Telnet connections to only one of the network adapters. It is also useful if you have a firewall, and you want to filter Telnet traffic through the firewall to only a few IP addresses.

By default, this registry entry has the value INADDR_ANY, which instructs Telnet Server to listen for Telnet connection requests that are sent to all IP addresses assigned to the host. You can change the value of this registry entry to any IP address that is assigned to the host.

# Terminal Services in Windows Server 2003 Technical Reference

Updated: November 14, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Terminal Services component of the Microsoft Windows Server 2003 operating system delivers the Windows Server 2003 desktop and the latest Windows-based applications to a wide variety of desktops, including those that normally would not be able to run Windows. Through terminal emulation, Terminal Services allows the same set of applications to run on diverse types of desktop hardware. For organizations wanting to more flexibly deploy applications and control desktop management costs, a Terminal Services architecture offers an important enhancement to the traditional two or three tier client-server architecture based on servers and full-scale personal computers.

Terminal Services extends the model of distributed computing by allowing PCs to operate in a server-based computing environment. With Terminal Services running on a Windows Server 2003-based computer, all client application execution, data processing, and data storage occur on the server. Applications and user desktops are transmitted over the network and displayed via terminal emulation software. Similarly, print streams, keyboard input, and mouse clicks are also transmitted over the network between the server and the terminal emulation software. Each user logs on and sees only their individual session, which is managed transparently by the server operating system and is independent of any other client session.

| **Note** |
| --- |
| In Windows Server 2008 R2, Terminal Services was renamed Remote Desktop Services. To find out what's new in this version and to find the most up-to-date resources, visit the Remote Desktop Services page on the Windows Server TechCenter. |

### In this section

- What Is Terminal Services?
- How Terminal Services Works
- Terminal Services Tools and Settings

#### Remote administration

- Remote Administration Using Terminal Services: Remote Desktop for Administration, Remote Desktops snap-in, Remote Desktop Connection
- Enable or disable Remote Desktop
- Troubleshooting General Remote Desktop Error Messages
- Remote Desktop Connection (Terminal Services Client 6.0)

#### Related downloads

- Terminal Services Scaling and Performance on x64-Based Versions of Windows Server 2003
- What's New in Terminal Services for Windows Server 2003 Service Pack 1
- Terminal Services Management Pack for Operations Manager 2007

#### Related resources

- Troubleshooting Remote Desktop Licensing Issues
- Terminal Services in Windows Server 2003 Online Help

- Hosting Applications with Terminal Server (Windows Server 2003 Deployment Guide)
- Getting Started with Terminal Server
- Configuring Authentication and Encryption in Terminal Services
- Remote Desktop Services in Windows Server 2008 R2
- Terminal Services in Windows Server 2008

# What Is Terminal Services?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Terminal Services?**

**In this section**


- Reducing the Total Cost of Ownership
- Common Terminal Services Scenarios
- Interactions with Other Technologies
- Related Information


Terminal Services provides the ability to host multiple, simultaneous client sessions on Windows Server 2003. Terminal Services is capable of directly hosting compatible multi-user client desktops running on a variety of Windows-based and non Windows-based computers. Standard Windows-based applications do not need modification to run on the terminal server, and all standard Windows Server 2003-?based management infrastructure and technologies can be used to manage the client desktops. In this way, corporations can take advantage of the rich choice of applications and tools offered by today's Windows operating system environment.

**Reducing the Total Cost of Ownership**

Organizations are always searching for ways to reduce the costs of ownership and this is one of the major goals of Terminal Services. Terminal Services lets enterprises more easily and cost-efficiently accomplish this goal by allowing organizations to:

- Centrally deploy and manage Windows-based applications.
- Remotely administer Windows Server 2003-based computers.

**Centrally deploy and manage Windows-based applications.**

Terminal Services provides centralized deployment and management of 32-bit Windows-based applications to Windows-based terminals, remote users, or local PC desktops. Using Terminal Services, companies can ensure that all clients are using the current version of an application because the software is installed once on a server computer, rather than on every desktop throughout the company. This model reduces the costs and challenge of updating desktop machines, especially for remotely located desktops or branch office environments. In addition, Terminal Services features such as Remote Control can simplify application support.

Organizations can use Terminal Server mode to deliver Windows-based applications to heterogeneous desktop environments, over local area network (LAN), wide area network (WAN) and dial-up connections. This is a cost-effective way to deploy line-of-business applications that are frequently updated, hard to install, or need to be accessed over low-bandwidth connections.

Terminal Services acts as a convenient bridging tool for earlier desktops migrating to Windows XP Professional. It allows the Windows XP desktop experience to be delivered "virtually" to non-PC desktops and PCs that need hardware upgrades before they can run a full Windows XP operating system locally. Terminal Services clients are available for many different desktop platforms including Microsoft MS-DOS, Windows-based PCs, Macintosh, UNIX and others. By letting users access current applications on hardware that might otherwise be of little use, Terminal Services can help companies that are gradually replacing older machines. (Connectivity to MS-DOS, Macintosh, and UNIX-based machines requires a third party add-on such as Citrix MetaFrame).

**Remotely administer Windows Server 2003-based computers**

Remote Desktop for Administration can greatly reduce the overhead associated with remote administration. Enabled by Terminal Services technology, Remote Desktop for Administration is specifically designed for server management. Therefore, it does not install the application-sharing and multi-user capabilities or the process scheduling of the full Terminal Server component (formerly called Terminal Services in Application Server mode). As a result, Remote Desktop for Administration can be used on an already busy server without noticeably affecting performance and CPU utilization, which makes it a convenient and efficient service for remote management.

Remote Desktop for Administration is used to remotely manage Windows Server 2003 servers. This mode is designed to provide operators and administrators with remote access to typical back-end servers and domain controllers. The administrator has access to the graphical user interface-based tools that are available in the Windows environment, even if he or she is not using a Windows-based computer to administer the server. Administrators can securely manage their Windows Server 2003-based computers over any network connection from any device using the Terminal Services Client software. This lets an administrator perform tasks such as directory maintenance, virus scans, backups, reboots and even promote a server to be a domain controller—all from a remote location.

Remote Desktop for Administration allows for the management of servers from any location without affecting server performance or application compatibility. In addition to the console session, up to two remote administration sessions are supported; since this is meant as a single-user remote access solution, no Terminal Server Client Access License (CAL) is required to use Remote Desktop for Administration.

Administrators can also fully administer computers running Windows Server 2003 family operating systems from computers running earlier versions of Windows by installing Remote Desktop Connection.

**Note**

- Remote Desktop for Administration is disabled by default in Windows Server 2003 family operating systems.

#### Common Terminal Services Scenarios

The way in which you plan to use Terminal Server has an effect on how you deploy it. The following scenarios outline the various ways in which you might use Terminal Server in your organization. Use the information in these scenarios to clarify how you are going to use Terminal Server in your organization.

#### Hosting Line-of-Business Applications

If your organization or certain groups within your organization use specialized applications to do their work, it might be beneficial to host the applications with Terminal Server. For example, you might decide to use Terminal Server in the following situations:

- **Custom applications.** If your line-of-business application is developed in-house or especially for your organization, and it tends to require frequent updating or repair, deploying the application once on a terminal server can simplify administration of the application. This is especially useful if your environment is geographically dispersed or if you are deploying Terminal Server to centrally serve your organization's branch offices.

- **Large central data pool.** Applications that rely on access to a single data source often run better on a terminal server because large amounts of data do not have to travel across the network to users. Instead, the data processing takes place on the server. Only the keystrokes and display information travel across the network, so you can use lower bandwidth connections. This is especially useful if the users of the data pool are located remotely, for example in a branch office with a slow connection to the database server.

- **Task workers.** In environments where you want workers to access only the application that they need to perform their jobs, you can centralize the administration of these users by using Terminal Server.

- **IT Admin tools.** System administrators can perform tasks that require Enterprise Administrator or Domain Administrator permissions on a terminal server that has the necessary tools installed on it. They can also run their desktop applications on their local computer without these permissions.

- **Upgrading operating systems.** If your organization uses a line-of-business application that is not optimized for your desktop operating system, you can host the application with Terminal Server rather than change operating systems.

#### Hosting the Desktop

You can use Terminal Server to host user's entire desktop environments, so that when users log on, they see their usual desktop environments or desktop environments especially designed for their remote use. In this situation, users can open and close the applications they choose in the same way that they access applications from the Windows desktop on the local computer. You can host the desktop in the situations discussed in the following sections.

#### Remote users

Hosting the desktop with Terminal Server can provide consistency and better performance for users in remote locations because large amounts of application data are not being transmitted over the connection. For example, you might use Terminal Server in the following situations:

- **Bandwidth-constrained locations.** In areas where high bandwidth is not available or cost-effective, deploying applications on a terminal server can improve performance for users who are connecting to the network from a remote location.

- **Mobile users.** For users who are traveling and tend to access the corporate network over connections of varying bandwidth, Terminal Server can provide a more consistent experience.

#### Client heterogeneity

If your organization is in the process of converting all users to the same platform or upgrading desktop hardware, you can use Terminal Server to quickly deliver the most up-to-date version of the operating system and applications to the user while enabling you to spread the desktop platform or hardware conversion over a longer period of time. You can also deliver a highly controlled standard desktop to users by using Terminal Server, as illustrated in more detail in the following list:

- **Mixed-platform environment.** If you have users who require applications based on operating systems other than Windows to perform their jobs, but your organization is transitioning to or requires a Windows-based desktop, you can host the desktop with Terminal Server. This requires the use of third-party software in conjunction with Terminal Server.

- **Upgrading hardware.** If your organization is planning to upgrade to Windows XP on the desktop, but not all of the desktop hardware is compatible, you can use Terminal Server to host the desktops of users who have older hardware while you are in the process of upgrading the hardware. All users can have the same desktop environment and run the latest versions of the applications designed for Windows XP regardless of their desktop hardware.

- **Highly controlled environments.** In situations where you want to deliver a standardized and controlled environment to users, you can host the desktop with Terminal Server to centralize management.

#### Hardware Considerations

You can reduce hardware costs by hosting applications with Terminal Server and using thin client devices or older hardware on your user's desktops.

- **Use of thin clients.** Windows Powered thin clients (sometimes called Windows-based terminals) offer an alternative to personal computers and traditional "green screen" terminals by enabling easy remote access to productivity and line-of-business applications that are hosted on Windows-based terminal servers.

- **Extending the life of older hardware.** Rather than replacing older hardware that is no longer capable of running new Windows-based applications, you can use that older hardware much like a thin client to access the desktop and applications on the server rather than on the local computer.

#### Using the Remote Desktop Web Connection

The Remote Desktop Web Connection is an ActiveX control that provides virtually the same functionality as the executable version of Remote Desktop Connection, but it delivers this functionality over the Web even if the executable version is not installed on the client computer. When hosted in a Web page, the ActiveX Client Control allows a user to log on to a terminal server through a TCP/IP Internet or intranet connection and view a Windows desktop inside Internet Explorer.

The Remote Desktop Web Connection provides an easy way to offer Terminal Server through a URL. Consider using the Remote Desktop Web Connection in the following situations:

- **Roaming users.** Users who are away from their computers can use Remote Desktop Web Connection to gain secure access to their primary workstations from any computer running Windows and Internet Explorer, provided they can reach the target computers.

- **Delivery of extranet applications.** You can use Remote Desktop Web Connection to allow business partners or customers access to internal applications over the Internet. Users who gain access in this manner do not need to reconfigure their computers, and they do not gain access to your internal network.

- **Deployment transition.** You can deploy the Remote Desktop Web Connection quickly and use it while you are deploying your full Remote Desktop Connection infrastructure.

#### Interactions with Other Technologies

Terminal Services depends on, or interacts with, the following technologies:

- RDP Protocol - On the server, RDP uses its own video driver to render display output by constructing the rendering information into network packets using RDP protocol and sending them over the network to the client. On the client, RDP receives rendering data and interprets the packets into corresponding Microsoft Win32 graphics device interface API calls. For the input path, client mouse and keyboard events are redirected from the client to the server. On the server, RDP uses its own virtual keyboard and mouse driver to receive these keyboard and mouse events.

- Terminal Server Interactions with RDP




**Related Information**

The following resources contain additional information that is relevant to this section.

- How Terminal Services Works

- Terminal Services Tools and Settings

- For more information about Terminal Services, see MSDN.

# How Terminal Services Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Terminal Services Works**

**In this section**


- Terminal Services Architecture
- Terminal Services Physical Structure
- Terminal Services Processes and Interactions
- Network Ports Used by Terminal Services
- Related Information


Terminal Services provides the ability to host multiple, simultaneous client sessions on Microsoft Windows Server 2003. Terminal Server is capable of directly hosting compatible multi-user client desktops running on a variety of Windows-based and non Windows-based hardware. Standard Windows-based applications do not need modification to run on the Terminal Server, and all standard Windows Server 2003-based management infrastructure and technologies can be used to manage the client desktops.

**Terminal Services Architecture**

Terminal Services consists of four components: the Windows Server 2003 multi-user kernel, the Remote Desktop client, the Terminal Services Licensing service, and Session Directory Services. Specifically:

Multi-user kernel. The multi-user kernel extensions, originally developed for Windows NT 4.0 Server, Terminal Server Edition, have been enhanced and fully integrated as a standard part of the Windows Server 2003 family kernel. These are resident on the server at all times, regardless of whether Terminal Services is enabled or not.

Remote Desktop client: The client software is an application that establishes and maintains the connection between a client and a server computer running Terminal Services.

Terminal Services Licensing service: This system allows terminal servers to obtain and manage terminal server client access license (TS CAL) tokens for devices and users connecting to a terminal server.

Session Directory Services: The session directory (SD) keeps a list of sessions indexed by user name, and allows a user to reconnect to the terminal server where the user's disconnected session resides and resume that session.

**Terminal Services Architecture**




The following table describes the Terminal Services architecture components.

**Terminal Services Components**

| Component | Description |
|---|---|
| CSRSS.exe | The Client-Server Runtime Subsystem is the process and thread manager for all logon sessions. |
| RdpDD.sys | Captures the Windows user interface and translates it into a form that is readily converted by RDPWD into the RDP protocol |
| RdpWD.sys | Unwraps the multi-channel data and then transfers it to the appropriate session. |
| SMSS.exe | Session Manager creates and manages all sessions. |
| Termsrv.exe | Manages client connections and initiates creation and shutdown of connection contexts. |
| Termdd.sys | The RDP protocol, which listens for RDP client connections on a TCP port. |
| Tdtcp.sys | Packages the RDP protocol onto the underlying network protocol, TCP/IP. |
| Wlnotify.dll | Runs in the session's WinLogon process to create processes in the user session. |
| Win32k.sys | Manages the Windows GUI environment by taking the mouse and keyboard inputs and sending them to the appropriate application. |
| WinLogon.exe | This system service handles user logons and logoffs and processes the special Windows key combination Ctrl-Alt-Delete. WinLogon is responsible for starting the Windows shell (which is usually Windows Explorer). |

**Terminal Services Architecture**

As the Windows Server 2003 Terminal Server boots and loads the core operating system, the Terminal Server service (termsrv.exe) is started and begins waiting for session connections. Each connection is given a unique session identifier or "SessionID" to represent an individual session to the Terminal Server, and each process created within a session is "tagged" with the associated SessionID to differentiate its namespace from any other session namespaces.

The console session (Terminal Server keyboard, mouse, and video) is always the first to load, is treated as a special-case client connection, and is assigned SessionID0. The console session starts as a normal Windows Server 2003 session, with the configured Windows display, mouse, and keyboard drivers loaded.

After creating the console session, the Terminal Server service then calls the Windows Session Manager (SMSS.EXE) to create two idle client sessions, which then await client connections. To create the idle sessions, the Session Manager starts the Client-Server Run-time Subsystem (CSRSS.EXE), and a new SessionID is assigned to that process. The CSRSS process also invokes the WinLogon process (WINLOGON.EXE) and the Windows Manager and GDI kernel module (Win32k.sys) under the newly associated SessionID.

The Windows image loader recognizes this Win32k.sys as a SessionSpace loadable image by a predefined bit set in the image header. It then relocates the code portion of the image into physical memory with pointers from the virtual kernel address space for that session if Win32k.sys has not already been loaded. By design, it always attaches to a previously loaded image's code (Win32k.sys) if one already exists in memory (that is, from any active application or session). The data (or non-shared) section of this image is then allocated to the new session from a newly created SessionSpace pageable kernel memory section.

Unlike the console session, Terminal Server client sessions are configured to load separate drivers for the display, keyboard, and mouse. The display driver is the Remote Desktop Protocol (RDP) display device driver (rdpdd.dll), and the mouse and keyboard drivers are replaced with the RDP driver Rdpwd.sys. These drivers allow the RDP client session to be both available and interactive, remotely. Finally, Terminal Server also invokes a connection listener thread for the RDP protocol (Termdd.sys), which listens for RDP client connections on a TCP port.

At this point, the CSRSS process exists under its own SessionID namespace, with its data instantiated per process as necessary. Any processes created from within this SessionID will execute within the SessionSpace of the CSRSS process automatically. This prevents processes with different SessionIDs from accessing another sessions data.

**Terminal Services Physical Structure**

Terminal Services provides remote access to a Windows desktop through "thin client" software, allowing the client computer to serve as a terminal emulator. It provides an effective and reliable way to distribute Windows-based programs, providing a single point of installation with multiple users having access to the Windows Server 2003 operating system desktop, where they can run programs, save files, and use network resources as if they were sitting at that computer.

For computers running Windows Server 2003 operating systems, the Terminal Services client program (Remote Desktop Connection) is already installed. Windows Server 2003 operating systems also include Terminal Services Client software for computers running 16- and 32-bit operating systems.

A Terminal Services client can exist in a variety of forms. Thin-client hardware devices that run an embedded Windows-based operating system can run the Terminal Services client software to connect to a server computer running Terminal Services. Windows, Macintosh, or UNIX computers can run Terminal Services client software to connect to a Terminal Services server to display Windows-based applications. This combination of Terminal Services clients provides access to Windows-based applications from virtually any operating system.

**Terminal Server Licensing**

The Windows Server 2003 operating system family provides a client license management system known as Terminal Server Licensing. This system allows terminal servers to obtain and manage terminal server client access license (TS CAL) tokens for devices and users connecting to a terminal server. Terminal Server Licensing is a component service of Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. It can manage unlicensed, temporarily licensed, and client-access licensed clients, and supports terminal servers that run Windows Server 2003 as well as the Windows 2000 Server operating system. This greatly simplifies the task of license management for the system administrator, while minimizing under- or over-purchasing of licenses for an organization. Terminal Server Licensing is used only with Terminal Server and not with Remote Desktop for Administration.

**Session Directory**

Terminal Services is a technology that lets users run Microsoft Windows-based applications on a remote Windows Server 2003-based computer. In a Terminal Server-based computing environment, all application execution and data processing occur on the server. In a load balanced environment, a farm of terminal servers have incoming session connections distributed in a balanced manner across the servers in the farm. The session directory (SD) keeps a list of sessions indexed by user name, and allows a user to reconnect to the terminal server where the user's disconnected session resides and resume that session.

### Terminal Services Physical Structure




### Terminal Services Components

| Component | Description |
|---|---|
| Terminal Server | Hosts applications for client computers. |
| Terminal Server Licensing Service | Issues TS Device CAL token to a requesting terminal server. |
| Session Directory | Maintains a list of the user names associated with the session IDs connected to the servers in a load balanced Terminal Server cluster. |

| | |
|---|---|
| Remote Desktop Connection | Client software that allows a connection to one remote computer. |
| Remote Desktops MMC Snap-in | Client software used by administrators to connect to multiple remote computers simultaneously. |
| Remote Desktop Web Connection | ActiveX RDP Web browser client. |

**Terminal Services Processes and Interactions**
Terminal Services transmits only the user interface of the program to the client, with the client computer connecting through the network, sending keystroke and mouse-movement information over the Remote Desktop Protocol to the Terminal Server. It then sends the client screen information in the form of simple (and bandwidth-friendly) GDI events, backed with bitmap information if required to properly display the desktop state.

Each user logs on and sees only their individual session, which is managed transparently by the server operating system and is independent of any other client session. The Terminal Server provides virtual Windows session management, so users can essentially treat that session as their own personal computer.

**Remote Desktop Clients**
There are three clients for Terminal Services: Remote Desktop Connection, Remote Desktops Snap-In, and Remote Desktop Web Connection. The client software is a very small software application that establishes and maintains the connection between a client and a server computer running Terminal Services. Each client transmits all input from the user to the server, such as keystrokes and mouse movements, and all output from the server such as application display information and print streams. Remote Desktop Web Connection provides most of the same functionality as the Remote Desktop Connection software; but it does not require a private network, or virtual private network connection. Remote Desktop Web Connection is covered in the next section.

### Connection using Remote Desktop Connection




**Remote Desktop Connection and Remote Desktops Snap-In**
The Terminal Services client has a new name, Remote Desktop Connection. Remote Desktop Connection can be installed and run on any Windows 95, Windows 98, Windows Millennium Edition, or Win32 platform (non-Windows-based clients are supported by the Citrix Metaframe add-on).

The client initiates a connection to the Terminal Server through TCP port 3389. The Terminal Server RDP listener thread detects the session request and creates a new RDP stack instance to handle the new session request. The listener thread hands over the incoming session to the new RDP stack instance and continues listening on the TCP port for further connection attempts. Each RDP stack is created as the client sessions are connected to handle negotiation of session configuration details.

First, an encryption level is established for the session. The Terminal Server initially supports three encryption levels: low, medium, and high.

Low encryption encrypts only packets being sent from the client to the Terminal Server. This “input only” encryption is to protect the input of sensitive data like a user’s password. Medium encryption encrypts outgoing packets from the client the same as low-level encryption, but also encrypts all display packets being returned to the client from the Terminal Server. This method of encryption secures sensitive data as it travels over the network to be displayed on a remote screen. Both low and medium encryption use the RC4 algorithm with a 40-bit key. High encryption encrypts packets in both directions, to and from the client, but uses the industry standard, non-exportable 128-bit high-level encryption.

At this point, prior to any logon being presented to the end user, the licensing details are negotiated. First, the client secures a Windows Server 2003 Client access license. Second, a Windows Server 2003 desktop license is verified on the machine that is connecting, and if no license can be validated, a connectivity license is provided to a non-Windows client to allow connection to the Terminal Server.

After session details have been negotiated, the server RDP stack instance, for this connection, is mapped to an existing idle Win32k user session, and the user is prompted with the Windows Server 2003 logon screen. If autologon is configured, the encrypted user name and password is passed to the Terminal Server and logon proceeds. If no idle Win32k sessions currently

exist, the Terminal Server service calls the Session Manager (SMSS) to create a new user space for the new session. Much of the Win32k user session is using shared code and does load noticeably faster after one instance has previously loaded.

After the user types her or his user name and password, encrypted packets are sent to the Terminal Server. The Winlogon process performs the necessary account authentication to ensure that the user has sufficient credentials to log on and then passes the user's domain and user name to the Terminal Server service, which maintains a domain/user name, SessionID list. If a SessionID is already associated with this user, they are remapped into the existing namespace, and the previous session stack is reloaded. Otherwise, the connection proceeds as normal and the Terminal Server service creates a new domain/user name, SessionID mapping. If for some reason more than one session is active for this user, the list of sessions is displayed and the user decides which one to make the reconnection.

The Remote Desktops snap-in is ideal for administrators who are remotely administering multiple servers or terminal servers. You can create Remote Desktop connections to multiple terminal servers or to computers running Windows 2000 Server or Windows Server 2003 family operating systems with the Remote Desktops MMC snap-in. A navigable tree display provides easy switching between connections.

**Remote Desktop Web Connection**

Remote Desktop Web Connection provides virtually the same functionality as Remote Desktop Connection, but delivers this functionality over the Web. When embedded in a Web page, Remote Desktop Web Connection can establish a Remote Desktop session with a remote computer, even if Remote Desktop Connection is not installed on the client computer. Remote Desktop Web Connection must be installed on a Web server with Internet Information Services (IIS) and Active Server Pages (ASP) enabled.

The Remote Desktop Web Connection is an ActiveX control that offers full feature parity with the standard Remote Desktop Connection client software offered in Windows XP and the Windows Server 2003 family of operating systems. The Remote Desktop Web Connection, like the Remote Desktop Connection which is installed by default, allows for connectivity to a Terminal Server. When installing the Remote Desktop Web connection on a specified computer the ActiveX control and sample ASP pages are installed. Internet Information Services is the delivery mechanism for the Remote Desktop Web Connection ActiveX control. This ActiveX control launches a Remote Desktop session to a Terminal Server computer from within Internet Explorer.




**How Remote Desktop Web Connection works**

Remote Desktop Web Connection connects to the Terminal Server as follows:

1. The user opens a Web browser and requests the initial Remote Desktop Web Connection (DHTML) login page.

2. The IIS server sends the page, and if this if the first time the client has connected, the user is also prompted to download the Remote Desktop ActiveX Control.

3. The user populates the connection information which includes the Terminal Server name.

4. The client computer creates a connection directly to the Terminal Server computer by using port 3389.

**Connection Process**

Remote Desktop Web Connection is a Web application that consists of an ActiveX control and DHTML pages. The DHTML pages provided only serve as a way to pass parameters to the Remote Desktop ActiveX control. Once the information is passed from the DHTML page to the ActiveX control, the Web page and the Web server no longer has a role in the process. Internet Information Services only serves as a mechanism for delivering the ActiveX control. When the user clicks **Connect**, the parameters from the DHTML page are passed to the control and a Remote Desktop connection to the Terminal Server is initiated.

When you connect to the IIS computer that is serving up the Remote Desktop Web Connection page, you are connecting over port 80. Upon connection to the Web page, the ActiveX control is downloaded to your client computer and stored in the default location for downloaded controls in Internet Explorer - %systemroot\Downloaded Program Files. From the supplied sample

Web page, the name of the Terminal Server and the display resolution are passed as parameters to the ActiveX control. After these parameters are passed, the connect method on the control is called, and then a session is launched to the Terminal Server computer. The Active X control on the client computer then creates a connection directly to the Terminal Server computer over TCP port 3389.

### Note

- The Web client is the same as the full Remote Desktop Connection client without the entire configuration interface. It obtains these properties from the Remote Desktop Web Connection page, and not by any communication with the IIS computer itself.

**Session Disconnect**
If a user decides to disconnect the session, the processes and all virtual memory space remain and are paged off to the physical disk if memory is required for other processes. Because the Terminal Server keeps a mapping of domain/user names and SessionIDs, when the same user reconnects, the existing session is loaded and made available again. An additional benefit of RDP is that of being able to change session screen resolutions, depending on what the user requests for the session. For example, let's say a user had previously connected to a Terminal Server session at 800 x 600 resolution and disconnected. If the user then moves to a different computer that only supports 640 x 480 resolution and reconnects to the existing session, the desktop is redrawn to support the new resolution.

**Automatic Reconnection**
Automatic Reconnection adds resilience to the Remote Desktop Connection client is Windows Server 2003. It is designed to recover from temporary connection losses due to network problems. Automatic Reconnection enables disconnected Terminal Services sessions to automatically re-authenticate to a Terminal Server without prompting the user for credentials.

Mobile users can greatly benefit from this feature. For example, with the previous versions of Terminal Server client, a user working with a wireless laptop which momentarily loses connectivity would be disconnected from the session and greeted with the message "The connection was ended because of a network error. Please try connecting to the remote computer again." Now with Automatic Reconnection, the user resumes the original session without needing to re-enter a password.

**Reconnection Process:**
The Automatic Reconnection process is managed using a cookie that is sent to the client. When reconnection is initiated on the client, it sends this cookie to the server as a token for validating the connection. The auto reconnection cookie is generated at the server, and is flushed and regenerated any time a user logs in to a session or when a session is reset. This ensures that after the user connects to the session from a different computer, the original computer cannot reconnect. The server also invalidates and updates the cookie at hourly intervals, sending an updated cookie to the client as long as the session is active.

**User Logoff**
Once a user logs off from the session, all processes associated with the SessionID are terminated and any memory allocated to the session is released. Of course, if the user was running a 32-bit application like Microsoft Word and logged off from the session, the application would remain in memory until the very last user exited from the application.

**Terminal Services Licensing service**
The Windows Server 2003 operating system family provides a client license management system known as Terminal Server Licensing. This system allows terminal servers to obtain and manage terminal server client access license (TS CAL) tokens for devices and users connecting to a terminal server. Terminal Server Licensing is a component service of Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. It can manage unlicensed, temporarily licensed, and client-access licensed clients, and supports terminal servers that run Windows Server 2003 as well as the Windows 2000 Server operating system. This greatly simplifies the task of license management for the system administrator, while minimizing under- or over-purchasing of licenses for an organization. Terminal Server Licensing is used only with Terminal Server and not with Remote Desktop for Administration.




**Terminal Server Licensing**

Terminal Server for Windows Server 2003 (known as Application Server mode in Windows 2000 Server) provides application deployment and management for users on a variety of devices through its application server mode. Each device or user who initiates a session on a terminal server running Windows Server 2003 must be licensed with one of the following:

1. Windows Server 2003 Terminal Server Device Client Access License.

2. Windows Server 2003 Terminal Server User Client Access License.

3. Windows Server 2003 Terminal Server External Connector.

Note that additional licenses might be needed, such as Microsoft or other application, operating system, and Client Access licenses. The licenses in the preceding list are required even if other add-on products are used on top of Windows Server 2003.

The Terminal Services Licensing service is only associated with licensing for a terminal server client. It is not used to license any other application or service, and does not replace or interoperate with the licensing service for any other component, or alter your rights and obligations under any End User License Agreement (EULA). The Terminal Server Licensing service is not a replacement for purchasing a TS CAL.

TS CAL tokens are electronic representations of real licenses, but they are not actual licenses themselves. Therefore if a license token is lost, it does not mean that you have lost an actual license. If you have the documentation to prove that you have bought an actual license, the license token can be re-issued. Conversely, just because you have a license token does not mean that it necessarily maps to an actual legal license.

Terminal Services Licensing is designed to manage these license tokens to allow an administrator to more accurately assess an organization's licensing requirements. However, there are a few situations in which a license token will not map to an actual license. The administrator should determine if this is the case, and if necessary, purchase extra licenses (but not install the corresponding license tokens) to account for this discrepancy.

**How Terminal Services Licensing Works**

All communication during the licensing process occurs between the client and the terminal server, and between the terminal server and the license server. The terminal server client never communicates directly with the license server.

When a client device attempts to connect to a terminal server in Per Device mode, the terminal server determines if the client has a license token. Terminal server clients store license tokens in the following registry key:

HKEY_LOCAL_MACHINE\Software\Microsoft\MSLicensing

If a client has no license token, the terminal server attempts to contact a license server from its list of discovered license servers. If no contact is made, the terminal server restarts the discovery process. If no license server responds, the device can not connect to the terminal server unless it is operating within the terminal server grace period.

When a license server responds, the terminal server requests a temporary token for the device because this is the first time the device has connected to a terminal server. The terminal server then pushes this temporary token to the device. After a user has provided valid credentials resulting in a successful logon, the terminal server instructs the license server to mark the issued temporary token as validated.

The next time a user attempts to connect to a terminal server in Per Device mode from this device, the terminal server requests a Windows Server 2003 TS Device CAL token for this device. If the license server has available TS Device CAL tokens, the

license server removes one token from the available pool, marks it as issued to the device, logs the device name, the user name of the device, and the date issued, and then pushes this TS Device CAL token to the device.

If the license server has no TS Device CAL tokens, it will first look to any other license server in its domain, workgroup, or site. License servers maintain information about where other accessible license servers exist, and if they have license tokens. If another license server is accessible that does have inventory, the first license server will request a license token from the second license server and deliver it to the terminal server, which then passes the token to the client device. If there are no available TS Device CAL tokens, the device will continue to connect with the temporary token.

Temporary tokens allow devices to connect for 90 days, and will then expire. TS Device CALs, while representing perpetual licenses, are set to expire 52-89 days from the date they are issued. The terminal server always attempts to renew these tokens 7 days prior to their expiration. The purpose of this system is to recover TS Device CAL tokens that are lost due to events such as hardware failure or operating system reinstallation.

**Client License Distribution Per User**
When a terminal server is configured in Per User mode, the terminal server must be able to locate a license server after the grace period has expired. While it is possible to install TS Per User CAL tokens on a license server, there is currently no method of assigning a TS Per User CAL token to a particular user account.

**Client License Distribution for External Connector**
There is currently no support in Terminal Server Licensing or the Microsoft Clearinghouse for the External Connector. In order to use an External Connector license, you will need to configure your terminal server in Per User mode.

**Terminal Server Licensing Model**
Terminal Server Licensing operates between several components as shown in the previous figure, including the Terminal Server Licensing-enabled license server, the Microsoft Certificate Authority and License Clearinghouse, one or more terminal servers, and terminal server clients. A single license server can support multiple terminal servers. There can be one or more license servers in a domain, or throughout a site.

**Microsoft Certificate Authority and License Clearinghouse**
The Microsoft Clearinghouse is the facility Microsoft maintains to activate license servers and to issue client license key packs to license servers. A client license key pack is a digital representation of a group of client access license tokens. The Microsoft Clearinghouse is accessed through the Terminal Services Licensing administrative tool. It can be reached directly over the Internet, through a Web page, or by phone.

**License Server**
A license server is a computer on which Terminal Server Licensing is installed. A license server stores all TS CALs license tokens that have been installed for a group of terminal servers and tracks the license tokens that have been issued. One license server can serve many terminal servers simultaneously. A terminal server must be able to connect to an activated license server in order for permanent license tokens to be issued to client devices. A license server that has been installed but not activated will only issue temporary license tokens.

**Terminal Server**
A terminal server is a computer on which the Terminal Server service is installed. It provides clients access to Windows-based applications running entirely on the server and supports multiple client sessions on the server. As clients connect to a terminal server, the terminal server determines if the client needs a license token, requests a license token from a license server, and then delivers that license token to the client.

**Terminal Services Session Directory Service**
In a Terminal Server-based computing environment, all application execution and data processing occur on the server computer. In a load balanced environment, terminal servers are grouped into farms, with each farm being represented to client machines as a single computer name with one IP address. The device performing the load balancing redirects incoming session connections to each machine in the farm according to its load balancing algorithm. Using a load-balancing solution with Terminal Server distributes sessions across the servers in the farm for improved performance. Terminal Services Session Directory, which is available with Windows Server 2003, Enterprise Edition, works with your load-balancing solution.

Session Directory is a load balancing feature that enables users to easily reconnect to a disconnected session on a server farm running Terminal Services. The Session Directory is a database that tracks user session's that are running on load-balanced terminal servers. It provides information when a user reconnects (after disconnecting intentionally or because of a network failure) to ensure that the user reconnects to the same session rather than starting a new session. Session Directory, which can support several thousand sessions, is also cluster-aware. Session Directory is compatible with the Windows Server 2003 load balancing service and is supported by third-party external load balancer products.




Terminal Services Session Directory management works with a load balancing service to ensure that users are transparently reconnected to the original server hosting their disconnected Terminal Server session.

The components of Terminal Services Session Directory are:

- A network load-balancing solution
- Two or more Terminal Servers logically grouped into a Terminal Server cluster
- A Session Directory server

Network Load Balancing, a clustering technology included in the Windows Server 2003, Enterprise Edition, and Windows Server 2003, Data Center Edition, enables servers to deliver high performance and failover protection. Network Load Balancing distributes IP traffic to multiple copies (or instances) of a TCP/IP service, such as Terminal Server or IIS, each running on an individual host within the cluster. The clients access the cluster using one or more virtual IP addresses; from there the NLB cluster evenly partitions the client requests among the hosts. From the clients point of view, the cluster appears to be a single server. As enterprise traffic needs increase, network administrators can simply add an additional server and incorporate it into the cluster.

In a load balanced environment, each Terminal Services cluster is represented to client machines as a single computer name with one IP address. The device or service performing its load balancing redirects connections to individual servers (nodes) in the cluster according to its load balancing algorithm.

Load balancing pools the processing resources of several servers using the TCP/IP networking protocol whereas Session Directory keeps track of the disconnected sessions on the cluster and ensures that users are reconnected to those same sessions.

Session Directory members can only be comprised of Enterprise and Datacenter Terminal Servers. Computers running Windows Server 2003, Standard Edition, are unable to join a Session Directory because they lack the required clustering technology; however a computer running Windows Server 2003, Standard Edition, can operate as the Session Directory server.

**How Session Directory Works**

1. An incoming connection to the cluster is load balanced to one node, which provides the logon prompt.
2. When the user logs on to the Terminal Server cluster, the Terminal Server receiving the initial client logon request sends the user name to the Session Directory server.
3. The Session Directory server checks the user name against its database and sends the result to the requesting server. The Session Directory database is a jet database containing a list of sessions indexed by user name.
4. If the user has no disconnected sessions, the log on process continues at the server hosting the initial connection.
5. If the user has a disconnected session on another serer, the initial hosting server sends the client information necessary for the client to continue authentication against the server hosting the disconnected session. The transition from one server to the other is transparent to the user.
6. When the user logs on to the disconnected session, the Session Directory is updated.

The Terminal Server Session Directory database is updated and queried by the Terminal Servers whenever users log on, log off, or disconnect from a session.

**Note**

- The Session Directory redirection feature is dependent on the version of the Remote Desktop Connection client used. The Remote Desktop client released with Windows XP is the minimum level necessary to function properly with Session Directory. Windows 2000 and Windows NT 4.0 Terminal Server clients are not Session Directory-aware, so client sessions will not be redirected to a disconnected session when that session is on an alternate node.

**Network Load Balancing**

Session Directory exists to solve a problem for Terminal Server clusters. These clusters can be load balanced using various solutions. This section discusses some general practices for any solution in use for load balancing. Here are some general considerations for Terminal Server clusters:

Consider splitting network traffic between two network adapters: one for Terminal Services connections and the other for access to other network resources and infrastructure. This allows for network access to the server in case the adapter bound to the cluster becomes unavailable.

For easier administration, place all load-balanced Terminal Servers into an Organizational Unit (OU) and apply Group Policy settings to that OU.

Home directories and other user data storage will need to be configured in such a way that the users can easily access their data no matter which server they are logged into.

Network Load Balancing (NLB) was previously installed by adding it as a service to a network connection and then configuring that network component on each node individually. NLB Manager in Windows Server 2003 is a new centralized snap-in added to provide easier configuration and maintenance for NLB configurations.

Although it is still possible to configure an NLB cluster by modifying network connection properties directly, the best practice is to use NLB Manager. In addition, the use of both NLB Manager and modifying network properties to configure an NLB cluster is not recommended.

**Remote Desktop for Administration**

You can use Remote Desktop for Administration to manage a network remotely using a configuration similar to the one shown in the following illustration.




Remote Desktop for Administration provides remote access to the server desktop by using the Terminal Services Remote Desktop Protocol (RDP) on port 3389. RDP transmits the user interface to the client session, and also transmits keyboard and mouse clicks from the client to the server. You can create up to two simultaneous remote connections. Each session you log on to is independent of other client sessions as well as the server console session. In essence, you can use Remote Desktop for Administration to log on to the server remotely as though you were logged on locally.

If you need to connect to the server console session remotely (for example, to access applications that direct only their user interface to the console), either use the Remote Desktops snap-in or use Remote Desktop Connection from the command line. When you attempt to connect to the console session, whether remotely or locally, you will be notified if there is already another user connected to the console session. The notification message will be shown after your logon credentials are validated, and will include information about the user who is logged on to the console session, including user name, location of logon (local or remote), and the state of the session (in use, locked, or idle).

**Note**

- Be aware of the security implications of remote logons. Users who log on remotely can perform tasks as though they were sitting at the console. For this reason, you should ensure that the server is behind a firewall. You should require all users who make remote connections to use a strong password.

The connection to Remote Desktop for Administration uses TCP/IP, either over an existing network connection or by remote access. A remote access server running one of the Windows Server 2003 family of operating systems provides two different types of remote access connectivity:

Network and Dial-up Connections

Virtual private networking

The following illustration shows how you can connect to a computer running one of the Windows Server 2003 family of operating systems from a remote location using remote access.




**Administering Windows Server 2003 family operating systems remotely**
After you are connected to a computer running a Windows Server 2003 family operating system, you can use Remote Desktop for Administration to remotely administer the server and its local computers. Remote Desktop for Administration gives you access to a variety of administrative tools used to configure and manage computers. Through a Terminal Services session, you can access Microsoft Management Console (MMC), Active Directory, Systems Management Server, network configuration tools, and most other administrative tools.

Remote Desktop for Administration is extremely useful because it provides remote access to most configuration settings, including Control Panel, which usually cannot be configured remotely. Also, using Remote Desktop for Administration can be particularly convenient for diagnosing a problem and testing multiple solutions quickly.

You can access the servers from anywhere in the world by using a wide-area network (WAN), a virtual private network (VPN), or a dial-up connection. You can start time-consuming batch administrative jobs (for example, tape backups), disconnect, and later reconnect to the corporate network to check progress.

Server application and operating system upgrades can be completed remotely as well as tasks that are not usually possible unless you are sitting at the console, such as domain controller promotion/demotion and disk defragmentation. Server file system tasks such as copying large files and virus scanning are much more efficient when performed within a Remote Desktop for Administration session, rather than using utilities that are executed from a client computer.

Administrative tasks are quicker and more intuitive than using command line utilities, although it is still possible to open a command shell.

### Note

- For some third-party applications, pop-up messages cannot be seen in a Terminal Services session. This is because there is a different security context or desktop for the connected session that does not display the application's pop-up messages. The pop-up messages in these instances will go directly to the console. If you need to see these messages, connect to the console session using Remote Desktop Connection from the command line or the Remote Desktops snap-in.

**Network Ports Used by Terminal Services**
Terminal Server uses RDP to communicate between client and server. RDP works only across a TCP/IP connection, such as a local area network (LAN), wide area network (WAN), dial-up, Integrated Services Digital Network (ISDN), digital subscriber line (DSL), or virtual private network (VPN) connection. You can still use other protocols, such as Internetwork Packet Exchange (IPX) or NetBIOS Extended User Interface (NetBEUI), as the transport protocol for non–Terminal Server traffic, such as network file or printer sharing, or between the client portion of a client-server application and its server.

### Port Assignments for Terminal Services

| Service Name | UDP | TCP |
|---|---|---|
| Remote Desktop Protocol (RDP) | | 3389 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- What Is Terminal Services?
- Terminal Services Tools and Settings
- For more information about Terminal Services, see MSDN.

# Terminal Services Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services Tools and Settings**

**In this section:**


- Terminal Services Tools
- Terminal Services Group Policy Settings
- Terminal Services WMI Classes
- Network Ports Used by Terminal Services
- Related Information


This section summarizes the tools and settings associated with Terminal Services.

**Terminal Services Tools**

The following tools are associated with Terminal Services.

**Change.exe: Change logon**

**Category**

This tool is included in all Microsoft Windows Server 2003 operating systems except the Web edition.

**Version compatibility**

This tool will run on all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

Enables or disables session logons to a server that has Terminal Server enabled. The change logon command disables logons from client sessions other than the system console. Users that are currently logged on are not affected. Client sessions are always re-enabled when you restart the system. If you are connected to the terminal server from a remote location and disable client sessions, and if you log off before re-enabling client sessions, you will not be able to reconnect. You need to logon at the system console in order to re-enable sessions.

**Change.exe: Change port**

**Category**

This tool is included in all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

**Version compatibility**

This tool will run on all Windows Server 2003 operating systems except Windows Server 2003, Web edition..

Changes the COM port mappings for MS-DOS application compatibility. Most MS-DOS applications support only COM1 though COM4 serial ports. Change port maps a serial port to a different port number, allowing applications that cannot access high-numbered COM ports to access the serial port. For example, to map COM12 to COM1 for use by a MS-DOS application, type change port com12=com1. Remapping works only for the current session.

You can run Change port without any parameters to display the available COM ports and the current COM port mappings.

**Change.exe: Change user**

**Category**

This tool is included in all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

**Version compatibility**

This tool will run on all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

Changes the .ini file mapping for the current user. Use change user /install before installing an application to create .ini files for the application in the Terminal Server system directory. These files are used as master copies for the user-specific .ini files.

After installing the application, use change user /execute to revert to normal .ini file mapping.

**Cprofile.exe: Clean profile**

**Category**
This tool is included in all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

Cleans the specified profiles of wasted space and removes user-specific file associations from the registry when disabled.

A terminal server uses file associations to determine which application to use to access files of various types. File types are registered using Windows Explorer.

Per-user file associations allow each user to have a different application associated with a specific file type. For example, one user could have .doc files associated with Microsoft Word and another user could have .doc files associated with Windows WordPad.

If user-specific file associations are enabled, Clean profile only removes the unused space from the user profile. If user-specific file associations are disabled, Clean profile also removes the corresponding registry entries.

**Flattemp.exe: Flat Temp**

**Category**
This tool is included in all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

Enables or disables flat temporary folders. After each user has a unique temporary directory, use flattemp /enable to enable flat temporary directories.

The default method for creating temporary folders for multiple users (usually pointed to by the TEMP and TMP environment variables) is to create subfolders in the \Temp folder, using the logonID as the subfolder name. For example, if the TEMP environment variable points to C:\Temp, the temporary folder assigned to the user logonID DonHall is C:\Temp\DonHall. Using flattemp, you can point directly to the \Temp folder and prevent subfolders from forming. This is useful when you want the user temporary folders to be contained in home directories, whether on a terminal server local drive or on a shared network drive. You should use this command only when each user has a separate temporary folder.

**Logoff.exe: Logoff**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Logs off a user from a session and deletes the session from the server. You can always log off from the session to which you are currently logged on. You must, however, have Full Control permission to log off users from other sessions.

Logging off a user from a session without warning can result in loss of data at the user's session. You should send a message to the user using the msg command to warn the user before taking this action.

If no ID or name for the session is specified, logoff logs off the user from the current session. If you specify a session name, it must be an active one.

When you log off a user, all processes end and the session is deleted from the server.

You cannot log off a user from the console session.

**Msg.exe: Msg**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Sends a message to a user. If you do not specify a name for the user or session, msg displays an error message. When specifying the name of a session, it must be an active one.

The user must have send message access permission to send a message.

**Mstsc.exe: Remote Desktop Connection**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Mstsc.exe is the command-line tool that launches Remote Desktop Connection. Remote Desktop Connection (formerly known as the Terminal Services client) is installed by default on all Windows Server 2003 family operating systems. You can use Remote Desktop Connection to connect to terminal servers, or to the desktop of a computer running one of the Windows Server 2003 family operating systems or Windows XP for remote administration.

Remote Desktop Connection allows you to create and configure your connection, save your connection settings to a file, and open and edit your saved connection files, all in the same program.

**Query.exe: Query process**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Query process displays information about processes running on a terminal server. You can use this command to find out which programs a specific user is running, and also which users are running a specific program.

If you do not specify the user name, session name, or program name, query process displays only the processes belonging to the current user.

If a session is specified, it must identify an active session. You can use wildcards to identify the process.

Query process returns the following information:

- The user who owns the process
- The session that owns the process
- The ID of the session
- The name of the process
- The state of the process
- The ID of the process

**Query.exe: Query session**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Query session displays information about sessions on a terminal server. The list includes information not only about active sessions but about other sessions that the server runs..

A user can always query the session to which the user is currently logged on. To query other sessions, the user must have Query Information access permission.

If you do not specify a session using session name, user name, or session id, query session displays information about all active sessions in the system.

**Query.exe: Query termserver**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Query termserver displays a list of all terminal servers on the network.

Query termserver searches the network for all attached terminal servers and returns the following information:

- The name of the server.
- The network (and node address if the /address option is used).

**Query.exe: Query user**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Query user displays information about user sessions on a terminal server. You can use this command to find out if a specific user is logged on to a specific terminal server. Query user returns the following information:

- The name of the user
- The name of the session on the terminal server
- The session ID
- The state of the session (active or disconnected)
- The idle time (the number of minutes since the last keystroke or mouse movement at the session)
- The date and time the user logged on

If you use query user without specifying a user name, session name, or session ID, a list of all users who are logged on to the server is returned. Alternatively, you can also use query session to display a list of all sessions on a server.

**Reset.exe: Reset session**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Enables you to reset (delete) a session from the terminal server. You can always reset your own sessions, but you must have Full Control access permission to reset another user's session.

Be aware that resetting a user's session without warning can result in loss of data at the session.

You should reset a session only when it malfunctions or appears to have stopped responding.

**Shadow.exe: Shadow**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Shadow enables you to remotely control an active session of another user. You can either view or actively control the session. If you choose to actively control a user's session, you will be able to input keyboard and mouse actions to the session.

You can always remotely control your own sessions (except the current session), but you must have Full Control access permission to remotely control another session. You can also initiate remote control using Terminal Services Manager.

Before monitoring begins, the server warns the user that the session is about to be remotely controlled, unless this warning is disabled. Your session might appear to be frozen for a few seconds while it waits for a response from the user.

Your session must be capable of supporting the video resolution used at the session you are remotely controlling or the operation fails.

The console session can neither remotely control another session nor can it be remotely controlled by another session.

**Tscc.msc: Terminal Services Configuration snap-in**

**Category**
This tool is included in all Windows Server 2003 operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

A Terminal Services connection provides the link clients use to log on to a session on the server. A TCP/IP connection is configured when Terminal Services is installed. Using Terminal Services Configuration, you can change the default properties of the connection or add new connections.

When you open Terminal Services Configuration you will see that a connection has already been configured. This is called the RDP-TCP connection. Typically, this is the only connection that needs to be configured for clients to connect to the server for Remote Desktop for Administration or application sharing with Terminal Server. Only one RDP (Remote Desktop Protocol) connection can be configured for each network adapter. If you want to configure additional RDP connections, you must install additional network adapters.

With Terminal Services Configuration, you can reconfigure the properties of the RDP-TCP connection, which includes limiting the amount of time client sessions can remain active on the server, setting protection levels for encryption, and selecting which permissions you want users and groups to have. Some connection properties can also be configured on a per-user basis using Terminal Services Group Policies or the Terminal Services extension to Local Users and Groups. For example, you can set different session time limits for each user when you use the Terminal Services extension to Local Users and Groups. Using Terminal Services Configuration, you can only set session time limits on a per-connection basis, which means the same time limit applies to all users who log on to the server using the connection.

**Tscon.exe: Tscon**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 family operating systems and on all Windows XP operating systems.

Tscon allows you to connect to another session. You must have Full Control access permission or Connect special access permission to connect to another session. You cannot connect to the console session

If you do not specify a password in the password parameter, and the target session belongs to a user other than the current one, tscon fails.

**Tsdiscon.exe: Tscon**

**Category**
This tool is included in all Windows Server 2003 family operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Tsdiscon disconnects a session from a terminal server. You must have Full Control permission to disconnect another user from a session.

If no session ID or session name is specified, tsdiscon disconnects the current session. The console session cannot be disconnected.

Any applications that were running when you disconnected the session are automatically running when you reconnect to that session with no loss of data.

**Tskill.exe: Tskill**

**Category**
This tool is included in all Windows Server 2003 operating systems and in all Windows XP operating systems.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems and on all Windows XP operating systems.

Tskill ends a process. You can use tskill to end only those processes belonging to you, unless you are an administrator. Administrators have full access to all tskill functions and can end processes running in other user sessions.

When all processes running in a session end, the session also ends.

**Tsmmc.msc: Remote Desktops snap-in**

**Category**
This tool is included in all Windows Server 2003 operating systems. For Windows XP, this tool is available in the Windows Server 2003 Administration Tools Pack.

**Version compatibility**
This tool will run on all Windows Server 2003 and Windows XP operating systems.

Windows Server 2003 provides the ability to connect to the console (session 0) of a computer by using a Remote Desktops connection. Because of this, a Connect to Console check box is available in the user interface, but this feature only works when you connect to a Windows Server 2003-based computer.

It is still possible to use the Remote Desktops tool to connect to a computer that is running Microsoft Windows NT 4.0, Terminal Server Edition, or Windows 2000 Server with Terminal Services enabled, but this creates a regular Remote Desktop Protocol (RDP) session. The option to connect to the console session is ignored.

The Remote Desktops snap-in is ideal for administrators who are remotely administering multiple servers or terminal servers. You can create Remote Desktop connections to multiple terminal servers or to computers running Windows 2000 Server or Windows Server 2003 family operating systems with the Remote Desktops snap-in. A navigable tree display provides easy switching between connections.

By default, the Remote Desktops snap-in connects you to the console session of the computer you specify in the connection. To run a specific program on connection, create a new connection that specifies the program name, and ensure that the default behavior is not selected.

**Tsprof.exe: Terminal Services profile**

**Category**
This tool is included in all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems except Windows Server 2003, Web edition.

Copies the user configuration information, which is displayed in the Terminal Services extensions to Local Users and Groups and Active Directory Users and Computers, from one user to another. Terminal Services profile can also set the profile path for a user.

**Tsadmin.exe: Terminal Server administration**

**Category**
This tool is included in all Windows Server 2003 operating systems except Windows Server 2003, Web Edition.

**Version compatibility**
This tool will run on all Windows Server 2003 operating systems except Windows Server 2003, Web Edition.

Use Terminal Services Manager to view information about terminal servers that reside in trusted domains. Use this tool to monitor users, sessions, and applications on each terminal server, and to carry out assorted actions to manage the server.

When a user creates a session by connecting to a terminal server from a client computer, the session appears in the Session list in Terminal Services Manager. In addition, the name of the user who logs on by using the session appears in the Users list. Any applications run in the user's session can be monitored on the Processes list. Therefore, you can oversee all users, sessions, and processes on a terminal server from one location.

**Terminal Services Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with Terminal Services.

**Group Policy Settings Associated with Terminal Services**

| Group Policy Setting | Description |
|---|---|
| Keep-Alive Connections | Specifies whether persistent connections are allowed. |
| Automatic reconnection | Specifies whether to allow Remote Desktop Connection clients to automatically reconnect to Terminal Services sessions if their network link is temporarily lost. |
| Restrict Terminal Services users to a single remote session | Specifies whether to restrict users to a single remote Terminal Services session. |
| Enforce Removal of Remote Desktop Wallpaper | Specifies whether desktop wallpaper is displayed to remote clients connecting via Terminal Services. |
| Deny log off of an administrator logged in to the console session | Specifies whether to allow an administrator attempting to connect to the console of a server to log off an administrator currently logged on to the console. |
| Allow Time Zone Redirection | Specifies whether to allow the client computer to redirect its time zone settings to the Terminal Services session. |
| Do not allow clipboard redirection | Specifies whether to prevent the sharing of clipboard contents (clipboard redirection) between a remote computer and a client computer during a Terminal Services session. |
| Do not allow smart card device redirection | Specifies whether to prevent the mapping of smart card devices in a Terminal Services session. |
| Allow audio redirection | Specifies whether users can choose where to play the remote computer's audio output during a Terminal Services session. |
| Do not allow COM port redirection | Specifies whether to prevent the redirection of data to client COM ports from the remote computer in a Terminal Services session. |
| Do not allow client printer redirection | Specifies whether to prevent the mapping of client printers in Terminal Services sessions. |
| Do not allow LPT port redirection | Specifies whether to prevent the redirection of data to client LPT ports during a Terminal Services session. |
| Do not allow drive redirection | Specifies whether to prevent the mapping of client drives in a Terminal Services session. |
| Do not set default client printer to be default printer in a session | Specifies whether the client default printer is automatically set as the default printer in a Terminal Services session. |
| Always prompt client for password upon connection | Specifies whether Terminal Services always prompts the client for a password upon connection. |
| Set client connection encryption level | Specifies whether to enforce an encryption level for all data sent between the client and the remote computer during a Terminal Services session. |
| Secure Server (Require Security) | Specifies whether a Terminal Server requires secure RPC communication with all clients or allows unsecured communication. |
| Limit number of connections | Specifies whether Terminal Services limits the number of simultaneous connections to the server. |

| | |
|---|---|
| Limit maximum color depth | Specifies the maximum color resolution (color depth) for Terminal Services connections. |
| Allow users to connect remotely using Terminal Services | Specifies whether to allow users to connect remotely using Terminal Services. |
| Do not allow local administrators to customize permissions | Specifies whether to disable the administrator rights to customize security permissions in the Terminal Services Configuration tool (tscc.msc). |
| Remove Windows Security item from Start menu | Specifies whether to remove the Windows Security item from the Settings menu on Terminal Services clients. |
| Remove Disconnect option from Shut Down dialog | Specifies whether to remove the Disconnect option from the Shut Down Windows dialog box on Terminal Services clients. |
| License Server Security Group | Specifies the terminal servers and license servers to which a Terminal Server License Server offers licenses. |
| Prevent License Upgrade | Specifies how a License Server distributes license upgrades to terminal servers running Windows 2000. |
| Do not delete temp folder upon exit | Specifies whether Terminal Services retains a user's per-session temporary folders at logoff. |
| Do not use temp folders per session | Specifies whether to prevent Terminal Services from creating session-specific temporary folders. |
| Terminal Server IP Address Redirection | Specifies how client devices are directed when reconnecting to an existing terminal server session. |
| Join Session Directory | Specifies whether Terminal Services uses a Session Directory for tracking user sessions, allowing a group of terminal servers to locate and connect a user back to an existing session. |
| Session Directory Server | Specifies whether to configure a server as a Session Directory Server for Terminal Services sessions on your network. |
| Session Directory Cluster Name | Specifies the Cluster Name for the terminal server, associating it with other servers in the same logical group. |
| Set time limit for disconnected sessions | Specifies a time limit for disconnected Terminal Services sessions. |
| Sets a time limit for active Terminal Services sessions | Specifies a time limit for active Terminal Services sessions. |
| Sets a time limit for active but idle Terminal Services sessions | Specifies a time limit for active but idle Terminal Services sessions. |
| Allow reconnection from original client only | Specifies whether to allow users to reconnect to a disconnected Terminal Services session using a computer other than the original client computer. |
| Terminate session when time limits are reached | Specifies whether to terminate a timed-out Terminal Services session instead of disconnecting it. |
| Set path for TS Roaming Profiles | Specifies whether Terminal Services uses the specified network path for roaming user profiles. |
| TS User Home Directory | Specifies whether Terminal Services uses the specified network share or local directory path as the root of the user's home directory for a Terminal Services session. |
| Sets rules for remote control of Terminal Services user sessions | Specifies the level of remote control permitted in a Terminal Services session. |
| Start a program on connection | Configures Terminal Services to run a specified program automatically upon connection. |

For more information about Group Policy settings, see the Group Policy Settings Reference for Windows Server 2003.

**Terminal Services WMI Classes**

The following table lists and describes the WMI classes that are associated with Terminal Services.

### WMI Classes Associated with Terminal Services

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_TerminalService | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSSessionDirectory | \\root\Cimv2 | Windows Server 2003 |
| Win32_TerminalServiceSetting | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSGeneralSetting | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSLogonSetting | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSSessionSetting | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSEnvironmentSetting | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSRemoteControlSetting | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSClientSetting | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSNetworkAdapterSetting | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSPermissionsSetting | \\root\Cimv2 | Windows Server 2003 |
| Win32_TSAccount | \\root\Cimv2 | Windows Server 2003 |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by Terminal Services**

Terminal Server uses RDP to communicate between client and server computers. RDP works only across a TCP/IP connection, such as a local area network (LAN), wide area network (WAN), dial-up, Integrated Services Digital Network (ISDN), digital subscriber line (DSL), or virtual private network (VPN) connection. You can still use other protocols, such as Internetwork Packet Exchange (IPX) or NetBIOS Extended User Interface (NetBEUI), as the transport protocol for non-Terminal Server traffic, such as network file or printer sharing, or between the client portion of a client-server application and its server.

### Port Assignments for Terminal Services

| Service Name | UDP | TCP |
|---|---|---|
| Remote Desktop Protocol | | 3389 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- How Terminal Services Works
- What Is Terminal Services?
- For more information about Terminal Services, see MSDN.

# Data Streaming and E-mail Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data Streaming and E-mail Technologies**

The use of streaming video and audio media is becoming an increasingly desirable feature in network technologies. As networks become more saturated with traffic, the timely delivery of latency-sensitive data, such as the data generated by video conferencing and Internet Protocol (IP) telephony programs, becomes more problematic. ATM and QoS are two technologies that you can use to facilitate the timely delivery of latency-sensitive data.

POP3 is an e-mail protocol that can be used facilitate mail delivery to computers that are not part of a network or computers that do not have an e-mail server on the local network.

**ATM**

Most network technologies use variable-length packet mechanisms to transfer data, and the available network bandwidth is divided into fixed channels or slots synchronized by a timing mechanism. This method is inefficient, and as the quantity of latency-sensitive data places increasing demands on the finite network bandwidth, the inefficiency becomes more problematic.

Asynchronous transfer mode (ATM) is a recent innovation in broadband networking technology that can be used to build high-speed networks. ATM differs from other existing LAN and WAN technologies because it is specifically designed to use bandwidth resources at maximum efficiency and support high-speed communications. To do this, ATM uses small, fixed-length cells to efficiently structure and parcel data for transfer. Because devices that communicate on an asynchronous network negotiate the rate at which they will communicate, based on physical hardware limitations and the ability to maintain a reliable flow of information within the network, they do not have the same types of restrictions as most existing network technologies.

ATM has been demonstrated at data rates up to 9.6 gigabits per second (Gbps) over Synchronous Optical NETwork (SONET), an international specification for fiber-optic communication. ATM is a scalable solution that can simultaneously accommodate data, voice, audio, fax, and video information types.

ATM provides versatility by means a set of features, including ATM-LAN emulation, address resolution protocol (ARP) address resolution service (MARS), and Encapsulated Ethernet over a permanent virtual connection (PVC):

- The ATM LAN emulation client enables existing programs and protocols designed for use on Ethernet and Token Ring networks to run without modification over an ATM network.
- In ATM, the LAN emulation (LANE) client module conforms to the ATM Forum LANE 1.0 specification. The ARP/MARS service enables the Microsoft TCP/IP stack to resolve ATM addresses to hardware addresses for more direct and efficient use of ATM media.
- Encapsulated Ethernet over PVC enables the transmission of Ethernet packets across PVC on an ATM network adapter.

**POP3**

Post Office Protocol version 3 (POP3) service enables a server to host e-mail accounts for subscribers of Internet service providers (ISP) or for small businesses. The most common use is for home ISP clients or small businesses that cannot locally implement and support a mail server, Simple Mail Transfer Protocol (SMTP) service, and related mail handling infrastructure. The POP3 protocol is used to retrieve mail that is being held by a mail server that offers SMTP service. The POP3 service performs the tasks of downloading messages, moving the messages from a folder in the file system of a mail server to clients on small networks, or to stand-alone home computers. Typically, after mail is downloaded from the mail server, the mail server deletes the mail from its folder.

**QoS**

As the finite bandwidth of a network becomes more congested with traffic, it becomes increasingly problematic to deliver latency-sensitive data in a timely manner. Quality of Service (QoS) is a set of service requirements that a network must meet to ensure an adequate service level for data transmission. The goal of QoS is a guaranteed delivery system for network traffic, such as Internet Protocol (IP) packets. A QoS guarantee indicates a service level that a program requires to transmit data at a specified rate and to deliver it within a specified time. Because QoS provides some guarantee of sufficient network resources, it gives a shared network a level of service similar to that of a dedicated network.

QoS facilitates the deployment of media-rich applications, such as video conferencing and Internet Protocol (IP) telephony, without adversely affecting network throughput. QoS achieves more efficient use of network resources by differentiating

between subsets of data such as fax data and latency-sensitive streaming media. QoS can also improve the performance of mission-critical software such as Enterprise Resource Planning (ERP) applications.

# ATM Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ATM Technical Reference**
Asynchronous transfer mode (ATM) refers to a number of related technologies, including software, hardware, and connection media. ATM is specifically designed to support high-speed communication. ATM differs from other existing LAN and WAN technologies by allowing networks to use bandwidth resources at maximum efficiency while maintaining Quality of Service (QoS) for users and programs.

The basic components of ATM are the computers that are connected to the ATM network and the devices that both connect these computers and ensure that data is transferred successfully. The computers that are connected to the ATM network are called endstations. Routers, Digital Subscriber Line Access Multiplexers (DSLAM) and ATM switches are examples of devices that connect end stations and ensure that data is transferred successfully.

**In this subject**


- What Is ATM?
- How ATM Works

# What Is ATM?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is ATM?**

**In this section**


- ATM Overview
- Connection Types: LAN vs. ATM


The cost of maintaining separate, specialized networks for computer, voice, and video is high. To reduce networking costs, ATM enables integration of all of these services on a single network and the combination of existing networks into a single infrastructure. In particular, Windows operating systems provide rich connectivity using Asynchronous Transfer Mode (ATM) while maintaining support for legacy systems.

To support ATM, Network Driver Interface Specification (NDIS) has been updated with ATM commands. Because many applications do not yet use ATM services, Windows Server 2003 includes support for LAN Emulation (LANE) for LAN applications, such as Ethernet. Similarly, IP over ATM support has been added, eliminating the additional header cost of LAN packets. Winsock 2.0 native ATM has also been added to support the many applications that use Windows Sockets (Winsock).

**ATM Overview**

ATM is a connection-oriented, unreliable (does not acknowledge the receipt of cells sent), virtual circuit packet switching technology. Unlike most connectionless networking protocols, ATM is a deterministic networking system — it provides predictable, guaranteed quality of service. From end to end, every component in an ATM network provides a high level of control. ATM technology includes:

- **Scalable performance.** ATM can send data across a network quickly and accurately, regardless of the size of the network. ATM works well on both very low and very high-speed media.
- **Flexible, guaranteed Quality of Service (QoS).** ATM allows the accuracy and speed of data transfer to be specified by the client. This feature distinguishes ATM from other high-speed LAN technologies such as gigabit Ethernet. The QoS feature of ATM also supports time dependent (or isochronous) traffic. Traffic management at the hardware level ensures that quality service exists end-to-end. Each virtual circuit in an ATM network is unaffected by traffic on other virtual circuits. Small packet size and a simple header structure ensure that switching is done quickly and that delays due to high traffic are minimized.
- **Unobstructed speed.** ATM imposes no architectural speed limitations. Its pre-negotiated virtual circuits, fixed-length cells, message segmentation and re-assembly in hardware, and hardware-level switching all help support extremely fast forwarding of data.
- **Integration of different traffic types.** ATM supports integration of voice, video, and data services on a single network. ATM over Asymmetric Digital Subscriber Line (ADSL) enables residential access to these services.

**Connection Types: LAN vs. ATM**

Traditional LANs, such as Ethernet and Token Ring, use a connectionless, unreliable approach, that cannot guarantee successful transmission when sending information across the network. Likewise, TCP/IP data transfers between networks are connectionless and unreliable. ATM, which is a connection-oriented, circuit-based technology, differs from the traditional approaches to networking.

**Traditional LAN**

In a traditional LAN, each client uses a network adapter card, which has a software driver. Above that driver is a protocol driver, such as TCP/IP. The protocol driver bundles information into frames of varying size and gives each bundle an appropriate header. As a result, when the adapter gains access to the media, the data packets are sent on the media to a destination hardware address. Traditional LAN technologies do not guarantee that data arrives on time or in the proper order. While Ethernet and Token Ring can detect errors, they provide no service guarantees and are not responsible for the recovery of missing or corrupted data packets.

Because the end stations are joined by a common medium, each end station on the traditional LAN recognizes the frames, or packets, of data put on the wire by each of the others, regardless of whether the frame is passed sequentially from one station to the next (as in a ring topology) or broadcast to all stations simultaneously (as with Ethernet). Each station has an adapter card, which processes the frame and examines the destination address. If the address applies to that computer, the frame is checked for errors. If there are no errors, the adapter initiates a hardware interrupt and passes the frame to the network adapter driver. The following figure, Traditional LAN: Connectionless Data Transmittal of a Packet, shows an example of a traditional LAN.

**Traditional LAN: Connectionless Data Transmittal of a Packet**




Because a traditional LAN is connectionless, it cannot provide mechanisms that can guarantee successful transmission. For example, it cannot determine the status of the destination adapter to ensure that it can receive a frame. It cannot ensure that bandwidth is available throughout the transmission. Unanticipated blockage, due to the media access control scheme of shared access technologies, can hinder a traditional LAN technology from supporting time-sensitive applications such as video or voice traffic. Traditional LANs can use upper-level protocol drivers to verify packet transmission (retransmitting, if necessary), partition big messages into smaller ones, and use time stamps for synchronization. However, these services add time to the transmission, and none of them provides end-to-end QoS guarantees.

**Traditional Internetworking**

If the destination address is remote rather than local, the chances of transmission failure increase. If a router on an Ethernet network detects a broadcast meant for another network, the router accepts the packet and passes it on using TCP/IP. A TCP/IP datagram is packet-switched to its destination individually. The header of each contains a globally significant switching address. This address allows a routing decision to be made each time the packet is forwarded, and packets to the same destination might follow completely different paths to get there, jumping over networks that use different underlying technology. No connection is required, but no delivery is guaranteed. The following figure, Two Packets Taking Different Routes Through a Traditional LAN, shows an example of two packets taking different routes through a traditional LAN.

**Two Packets Taking Different Routes Through a Traditional LAN**




Like an Ethernet data transfer, a routed data transfer cannot offer guarantees because bandwidth is never reserved ahead of time. The packets being sent over TCP/IP are simply transmitted on the wire and routed. While this allows flexibility in routing around obstructions, network performance can vary a great deal depending on conditions at the routers and on the amount of network traffic.

**ATM Networks**

In contrast to connectionless transmission protocols, ATM is connection-oriented. An ATM endpoint establishes a defined path known as a virtual channel (VC), also called virtual circuit, to the destination endpoint prior to sending any data on the network. It then sends a series of same-size frames, called cells, along the virtual channel towards the destination.

While establishing the connection, the ATM endpoint also negotiates a QoS contract for the virtual channel. The QoS contract spells out the bandwidth, maximum transit delay, acceptable variance in the transit delay, and so forth, that the VC provides, and this contract extends from one endpoint to the other through all of the intermediate ATM switches.

The path of ATM traffic is established at the outset, and the switching hardware merely needs to examine a simple header to identify the proper path. Beyond specifying a path, ATM allows a location to establish a full duplex connection (traffic travels in both directions) with multiple locations at the same time. Note, however, that ATM is an unreliable transmission protocol because it does not acknowledge the receipt of cells sent. As with LANs, missing or corrupted information must be detected and corrected by upper-layer protocols.

The following figure, "ATM Virtual Channel and Packet Transmission," illustrates ATM virtual channel and packet transmission.

**ATM Virtual Channel and Packet Transmission**




**Network Speed**

Unlike Ethernet networks, ATM has no inherent speed limit, and its efficiency is not affected by the distance that the data has to travel. In addition, ATM establishes the pathway for a particular series of packets at the outset and ATM switches make minimal switching decisions thereafter. To travel across the ATM network, data is segmented into same-size cells, and encapsulated with a header that contains information about switching, congestion, and error-checking.

Cells are transmitted in order, and the ATM network uses Virtual Path Identifier and Virtual Channel Identifier (VPI/VCI) numbers in the ATM header to forward them efficiently. A switch reads the header, compares the VPI/VCI to its switching table to determine the correct output port and new VPI/VCI, and then forwards the cell. All the addressing information that the ATM switch needs is contained in the header and is always found in the same place. This makes the forwarding task simple to implement in hardware by, reducing latency. Moreover, with ATM from end to end, there is no data translation required if a packet must travel from a LAN through a WAN to reach a destination LAN. The following figure, "ATM Fixed-Length Cells," shows two ATM end stations sending fixed-length cells from A to B (although ATM traffic is bi-directional).

**ATM Fixed-Length Cells**




Because ATM uses small (53-byte), fixed-length cells that require less logic to process, the network spends no time determining where a particular cell begins and ends. The small cell size ensures that delays in forwarding cells are minimized. Because the cell size is so predictable, buffer usage and analysis algorithms can be simplified and optimized.

Traditional LAN technologies, such as Ethernet, have inherent speed limitations Either the underlying infrastructure (the cable) or the segment length must be changed to support fast traffic. However, unlike Ethernet and Token Ring, ATM has no such imposed limitations. If you can invent a faster physical layer — if you can design a quicker method of transmitting data from one place to another over one wire or many wires — ATM can work over that physical layer and at those new speeds. In addition, ATM allows information with different requirements and from different nodes to be transmitted nearly simultaneously without conflict.

ATM places fixed-length cells on the media when the data is produced according to the parameters of a negotiated connection. ATM can simultaneously handle the needs of isochronous (time-dependent) traffic, such as voice and video, and non-isochronous traffic, such as LAN data.

# How ATM Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How ATM Works**

**In this section**


- ATM Architecture
- ATM Components
- ATM Cell Structure
- Virtual Paths and Virtual Channels
- Quality of Service
- ATM Addresses
- ATM Connection Types
- LAN Emulation
- TCP/IP Over ATM
- Additional Resources


Asynchronous transfer mode (ATM) is one of many network transmission protocols included in Windows Server 2003. The most commonly used transmission protocol included in Windows Server 20003 is TCP/IP, which is a connectionless protocol. As such, TCP/IP cannot offer some of the advantages that a connection-oriented, virtual circuit, packet-switching technology, such as ATM, can. Unlike most connectionless networking protocols, ATM is a deterministic networking system — it provides predictable, guaranteed quality of service.

The ideal environment in which to use ATM is one that combines computer, voice, and video networking into a single network, and the combination of existing networks into a single infrastructure.

**ATM Architecture**

ATM is a combination of hardware and software that can provide either an end-to-end network or form a high-speed backbone. The structure of ATM and its software components comprise the ATM architecture, as the following illustration shows. The primary layers of ATM are the physical layer, the ATM layer, and the ATM Adaptation layer.

**ATM Architectural Diagram**




Each layer and sublayer is described briefly in the following table, “ATM Layers.”

**ATM Layers**

| Layer | Function |
|---|---|

| | |
|---|---|
| ATM Adaptation | The ATM Adaptation layer facilitates the use of packets larger than a cell. Packets are segmented by the ATM interface, transmitted individually, and then reassembled on the receiving end. The ATM Adaptation Layer includes the Segmentation and Reassembly and Convergence sublayers. |
| ATM | The ATM layer regulates cells and cell transport and establishes and releases Virtual Circuits. The ATM layer has no sublayers |
| Physical | The Physical layer represents the physical medium and regulates Physical layer functions such as voltages and bit timing. The Physical layer consists of the Transmission Convergence and the Physical Medium Dependent sublayers |

**Physical Layer**

The physical layer provides for the transmission and reception of ATM cells across a physical medium between two ATM devices. This can be a transmission between an ATM endpoint and an ATM switch, or it can be between two ATM switches. The physical layer is subdivided into a Physical Medium Dependent sublayer and Transmission Convergence sublayer.

**PMD Sublayer**

The Physical Medium Dependent (PMD) sublayer is responsible for the transmission and reception of individual bits on a physical medium. These responsibilities encompass bit timing, signal encoding, interacting with the physical medium, and the cable or wire itself.

ATM does not rely on any specific bit rate, encoding scheme or medium and various specifications for ATM exist for coaxial cable, shielded and unshielded twisted pair wire, and optical fiber at speeds ranging from 64 kilobits per second to 9.6 gigabits per second. In addition, the ATM physical medium can extend up to 60 kilometers or more by using single-mode fiber and long-reach lasers. Thus it can readily support wide-range connectivity, including a private metropolitan area network. The independence of ATM from a particular set of hardware constraints has allowed it to be implemented over radio and satellite links.

**Transmission Convergence Sublayer**

The Transmission Convergence (TC) sublayer functions as a converter between the bit stream of ATM cells and the PMD sublayer. When transmitting, the TC sublayer maps ATM cells onto the format of the PDM sublayer, such as the DS-3 interface or Synchronous Optical Network (SONET) frames. Because a continuous stream of bytes is required, unused portions of the ATM cell stream are filled by idle cells. These idle cells are identified in the ATM header and are silently discarded by the receiver. They are never passed to the ATM layer for processing.

The TC sublayer also generates and verifies the Header Error Control (HEC) field for each cell. On the transmitting side, it calculates the HEC and places it in the header. On the receiving side, the TC sublayer checks the HEC for verification. If a single bit error can be corrected, the bit is corrected, and the results are passed to the ATM layer. If the error cannot be corrected (as in the case of a multibit error) the cell is silently discarded.

Finally, the TC sublayer delineates the ATM cells, marking where ATM cells begin and where they end. The boundaries of the ATM cells can be determined from the PMD layer formatting or from the incoming byte stream using the HEC field. The PMD performs the HEC validation per byte on the preceding 4 bytes. If it finds a match, the next ATM cell boundary is 48 bytes away (corresponding to the ATM payload). The PMD performs this verification several times to ensure that the cell boundaries have been determined correctly.

**ATM Layer**

The ATM layer provides cell multiplexing, demultiplexing, and VPI/VCI routing functions. The ATM layer also supervises the cell flow to ensure that all connections remain within their negotiated cell throughput limits. If connections operate outside their negotiated parameters, the ATM layer can take corrective action so the misbehaving connections do not affect connections that are obeying their negotiated connection contract. The ATM layer also maintains the cell sequence from any source.

The ATM layer multiplexes and demultiplexes and routes ATM cells, and ensures their sequence from end to end. However, if a cell is dropped by a switch due to congestion or corruption, it is not the responsibility of the ATM layer to correct the dropped cell by means of retransmission or to notify other layers of the dropped cell. Layers above the ATM layer must detect the lost cell and decide whether to correct it or disregard it.

In the case of interactive voice or video, a lost cell is typically disregarded because it takes too long to resend the cell and place it in the proper sequence to reconstruct the audio or video signal. A significant number of dropped cells in time-dependent services, such as voice or video, results in a choppy audio or video playback, but the ATM layer cannot correct the problem unless a higher Quality of Service is specified for the connection.

In the case of data (such as a file transfer), the upper layer application must sense the absence of the cell and retransmit it. A file with randomly missing 48-bytes chunks is a corrupted file that is unacceptable to the receiver. Because operations, such as file transfers, are not time dependent, the contents of the cell can be recovered by incurring a delay in the transmission of the

file corresponding to the recovery of the lost cell.

**ATM Layer Multiplexing and Demultiplexing**
ATM layer multiplexing blends all the different input types so that the connection parameters of each input are preserved. This process is known as traffic shaping.

ATM layer demultiplexing takes each cell from the ATM cell stream and, based on the VPI/VCI, either routes it (for an ATM switch) or passes the cell to the ATM Adaptation Layer (AAL) process that corresponds to the cell (for an ATM endpoint).

**ATM Adaptation Layer**
The ATM Adaptation Layer (AAL) creates and receives 48-byte payloads through the lower layers of ATM on behalf of different types of applications. Although there are five different types of AALs, each providing a distinct class of service, Windows Server 2003 supports only AAL5. ATM Adaptation is necessary to link the cell-based technology at the ATM Layer to the bit-stream technology of digital devices (such as telephones and video cameras) and the packet-stream technology of modern data networks (such as frame relay, X.25 or LAN protocols such as TCP/IP or Ethernet).

**AAL5**
AAL5 provides a way for non-isochronous (time-dependent), variable bit rate, connectionless applications to send and receive data. AAL5 was developed as a way to provide a more efficient transfer of network traffic than AAL3/4. AAL5 merely adds a trailer to the payload to indicate size and provide error detection. AAL5 is the preferred AAL when sending connection-oriented or connectionless LAN protocol traffic over an ATM network. Windows Server 2003 supports AAL5.

AAL5 provides a straightforward framing at the Common Part Convergence Sublayer (CPCS) that behaves more like LAN technologies, such as Ethernet. The following figure, "Breakdown of an AAL5 Cell Header and Payload," shows a detailed breakdown of an AAL5 Cell Header and Payload, followed by a detailed description of each of the components.

### Breakdown of an AAL5 Cell Header and Payload




With AAL5, there is no longer a dual encapsulation. The service class frames cells at the CPCS, but not at the Segmentation and Reassembly sublayer to minimize overhead. It also uses a bit in the Payload Type (PT) field of the ATM header rather than a separate SAR framing.

AAL5 is the AAL of choice when sending connection-oriented (X.25 or Frame Relay) or connectionless (IP or IPX) LAN protocol traffic over an ATM network.

**AAL5 CPCS Sublayer**
The preceding figure, "Breakdown of an AAL5 Cell Header and Payload," shows the framing that occurs at the AAL5 CPCS sublayer. (Note that only a trailer is added.).

**CPCS PDU Payload**
The block of data that an application sends. The size can vary from 1 byte to 65,535 bytes. The packet assembler/disassembler (PAD) consists of padding bytes of variable length (0-47 bytes), which create a whole number of cells by making the CPCS PDU payload length a multiple of 48 bytes.

**User-to-User Indication**
Transfers information between AAL users.

**Common Part Indicator**
Currently used only for alignment processes so that the AAL5 trailer falls on a 64-bit boundary.

**Length of CPCS PDU Payload Field**
Indicates the length of the CPCS PDU payload in bytes. The length does not include the PAD.

**Cyclic Redundancy Check (CRC)**
A 32-bit portion of the trailer that performs error checking on the bits in the CPCS PDU. The AAL5 CRC uses the same CRC-32 algorithm used in 802.*x*-based networks such as Ethernet and Token Ring.

**AAL5 SAR sublayer**
Byte-by-byte, the preceding figure, "Breakdown of an AAL5 Cell Header and Payload," shows the framing that occurs at the AAL5 SAR sublayer. There is no SAR header or trailer added. On the transmitting side, the AAL5 SAR sublayer merely segments the CPCS PDU into 48-byte units and passes them to the ATM layer for the final ATM header.

On the receiving side, the sublayer reassembles a series of 48-byte units and passes the result to the CPCS. The AAL5 SAR uses the third bit in the PT field to indicate when the last 48-byte unit in a CPCS PDU is being sent. When the ATM cell is received with the third bit of the PT field set, the ATM layer indicates this fact to the AAL; the AAL then begins a CRC and length-checking analysis of the full CPCS PDU.

**ATM Components**
To support native Asynchronous Transfer Mode (ATM), NDIS has been updated with native ATM commands. Because many applications do not yet use native ATM services, LANE support was added for LAN applications, such as Ethernet. Similarly, Microsoft has added IP over ATM support, thereby eliminating the additional header cost of LAN packets. Microsoft also added Winsock 2.0 native ATM to support the many applications that use Windows Sockets (Winsock).

Furthermore, circuit connectivity has been added to theTelephony Application Programming Interface TAPI (connection management protocol) to provide complete ATM support. TAPI can now make and receive calls and can redirect them to ATM circuits or from circuits into devices or other network types. Examples include Microsoft DirectShow, as well as Point to Point Protocol (PPP) over ATM as the dial-up remote access protocol on ATM. Using the raw channel access (RCA) kernel streaming filter, TAPI can be used to connect a data stream to the RCA filter containing video, and send it over an ATM circuit, as shown in the following figure, "Windows ATM Services."

**Windows ATM Services**




These enhancements allow applications to exploit ATM services, such as Quality of Service (QoS), and with the use of TAPI, achieve a high level of integration between established multimedia features and network protocols.

**ATM Call Manager**
The ATM signaling component, also known as the UNI call manager, handles virtual channel creation and management. This section describes how the ATM call manager does its job, specifically the handling of both permanent and switched virtual connections.

**How the Call Manager Differentiates PVCs and SVCs**
Permanent virtual connections (PVCs) are almost identical to switched virtual connections (SVCs), but each PVC must be manually configured, device by device, by an administrator. In contrast, SVCs are dynamically configured when they are established. Each device — from a starting end station through switches to another end station — independently determines its role in supporting a virtual connection It also determines what device to forward the request to, and whether or not it can guarantee the requested Quality of Service at that time. PVC resource allocations are set aside the moment they are first configured, whether or not they are used immediately. SVC resource allocations are allocated dynamically.

The SVC and PVC values are both stored in the internal tables of the ATM call manager, the ATM adapter, and the ATM switch, and the kind of values stored in those tables are identical. The difference between the two kinds of connections lies in how the connection values are handled at initialization. At that time, the ATM call manager checks the registry for any PVCs. If it finds one, it stores its VC number, along with other VC information such as Quality of Service, the process ID (or more generically, the service access point), and the source and destination addresses. It uses a single bit to designate that it is a PVC and not an SVC.

During initialization, the ATM adapter does not know about PVCs. Until someone (typically an administrator) configures an application to use a PVC, applications are not aware of PVCs either. When an application wants to use a PVC, it issues an ATM command through its provided interfaces. The request specifies the destination address, the Quality of Service, and the virtual

connection number (among other information). Up to this point, the PVC is handled exactly as if it were a request for an SVC. The call manager receives the request and checks the information received against the entries in its internal table of VCs. If it finds a match, and the match is designated as a PVC in the PVC field in its table, the call manager then handles the rest of the process differently from how it handles an SVC request.

A typical SVC request initiates two commands; the first determines whether the adapter can handle another VC, and the second activates the VC along the path of network components. A PVC request, however, works a little differently. When the call manager receives a request specifying a PVC, it behaves as if the PVC has already been established from end to end. It sends the two initiating calls in rapid succession to the ATM adapter. The ATM adapter never detects that it is working with a PVC. It obtains the Quality of Service and other information from the setup commands and determines how to shape the traffic. From that point, the PVC functions identically to an SVC.

**Data Encapsulation over PVC**

Support for Ethernet and IP data encapsulation over ATM using VC Multiplexing or LLC/SNAP AAL5 PVCs is new to Windows Server 2003. This support extends to the encapsulation and transport of IP or Ethernet packets over ATM between clients that are connected to a supporting infrastructure using a permanent virtual connection. ATM supports Ethernet and IP data encapsulation by acting as a bridging Ethernet or routing adapter for the TCP/IP protocol.

ATM AAL5 PVC support is similar in concept to ATM LANE, except where LANE uses SVCs instead of PVCs.
Windows Server 2003 supports the two encapsulation methods, LLC Encapsulation and VC Multiplexing as described in RFC 2684. Both Ethernet and IP protocols are supported by using either encapsulation method for bridged and routed protocol data units (PDUs). For example PPP over Ethernet, (PPPoE), Layer 2 Tunneling Protocol (L2TP ), Ethernet, or Ethernet encapsulated in IP are some of the protocols supported by using ATM AAL5 PVC support in Windows Server 2003.

ATM AAL5 is implemented as an NDIS intermediate driver, as shown in the following figure, "(NGFD_ATM90) ATM AAL5 NDIS Intermediate Driver."

**(NGFD_ATM90) ATM AAL5 NDIS Intermediate Driver**




There are several common scenarios for using ATM AAL5 PVCs. One example is using ATM AAL5 PVC support for remote connectivity from a home or small office that uses an internal ADSL mode. In Windows Server 2003 you configure the ADSL mode as an ATM AAL5 PVC connection, listed as Microsoft Ethernet PVC - RFC 2684 in Windows Server 2003. As shown in the following figure, "(NGFD_ATM91) ADSL Connectivity with ATM AAL5 PVC ,"the ADSL modem connects by using the Public Switched Telephone Network (PSTN) to a Digital Subscriber Line Access Multiplexer (DSLAM) located at your service provider, most likely the central office of your local telephony carrier. The DSLAM either bridges the encapsulated data directly to a network or connects to an external bridge, router, or ATM switch located at your service provider. A connection can then be made to the targeted network, such as your corporate office or the Internet.

**(NGFD_ATM91) ADSL Connectivity with ATM AAL5 PVC**




**ATM LAN Emulation Module**

LAN emulation client services are included in the Windows Server 2003 operating system. When Plug and Play detects an ATM adapter and installs the appropriate driver, the LANE client is also installed by default. This permits full LANE connectivity without the need for configuration, provided that the following conditions exist:

- The switch has LANE services available and turned on.
- The LANE services configuration has an enabled default, Emulated Local Area Network (ELAN).

For centralized administration and ease of configuration, this LANE client implementation allows configuration of only the ELAN name. All other ELAN configuration information — such as the Maximum Transmissible Unit (MTU), ELAN type, and LAN Emulation Server (LES) — is obtained from the LAN Emulation Configuration Server (LECS). If a default ELAN is enabled in the LECS, no configuration is required.

**ATMARP and ARP MARS**
IP over ATM support is included with Windows Server 2003. IP over ATM exposes many features of ATM so that TCP/IP can make use of them directly. With this support, applications written to use TCP/IP can also make use of ATM.

In addition, the client that runs ATMARP (IP over ATM) in Windows Server 2003 supports multicast address resolution by using Multicast Address Resolution Service (MARS). This client contains an ATM ARP/MARS service that enables Windows to act as both an ATMARP server and a MARS with integrated Multicast Server (MCS). The MARS setup allows configuration of a ranges of addresses; the service acts as an MCS for all those addresses.

**API Support: Winsock 2.0, TAPI, and NDIS 5.0**
All Windows Server 2003 enhancements to SNMP are possible due to extensions to the operating system. The chief extension is a connection-oriented service added to NDIS version 5.0. NDIS 5.0 includes Connection-Oriented NDIS (CoNDIS), a new NDIS API extension for the support of connection-oriented media. These new APIs enable applications and protocols to create VCs and specify Quality of Service for those VCs. CoNDIS supports multiple call managers to enable various media-specific signaling needs, including an ATM-specific call manager. In addition, CoNDIS supports point-to-multipoint connections for efficient multicast services, as shown in the following figure "Supported CoNDIS Multicast Services."

**Supported CoNDIS Multicast Services**




Two components operate on top of NDIS, integrating ATM services with the rest of the operating system and exposing ATM services through well-known APIs. Winsock 2.0 now has direct ATM support by using the Winsock ATM Service Provider. Winsock support provides direct access to ATM services from user-mode applications. With the addition of IP over ATM support, Windows Sockets applications that use TCP/IP as a transport protocol can be run over ATM networks and inter-operate with standard LAN-based IP clients.

**NDIS 5.0 ATM Miniport Drivers**
Although NDIS 5.0 supports both connectionless and connection-oriented network adapter drivers, only the connection-oriented drivers are of use in an ATM network.

Connection-oriented miniport drivers are always deserialized; that is, they serialize the operation of their own miniport functions and queue all incoming packets internally rather than relying on NDIS to perform the same functions. This improves full-duplex performance if the critical sections of the driver are kept small.

While NDIS library continues to support legacy NDIS 3.0 network adapter drivers, only NDIS 4.0 and 5.0 miniport drivers can take advantage of the enhanced functionality and performance characteristics of the NDIS library support for network adapter drivers.

NDIS has several connection-oriented features. It also contains additional features such as binary compatibility, improved

power management by means of Wake On LAN support (which enables a network adapter to turn on a client from a low power state based on packet receipt), and checksum performance in hardware rather than in software. As a result, driver performance is improved over all network types.

**TAPI**

Telephony Application Programming Interface (TAPI) is responsible for connection setup and other operating system functions related to telephony. In Windows Server 2003, TAPI has been expanded to support telephony over connection-oriented media such as ATM. While TAPI does not handle data directly, it can create a circuit and connect that circuit to another device.

By redirecting calls, TAPI provides more than just high bandwidth and good connectivity. The TAPI component of CoNDIS maps (or proxies) the TAPI call management functions to NDIS 5.0 call management functions, allowing a connection from another medium to be redirected to or from ATM. For example, TAPI can redirect calls to a data handler such as the raw channel access filter, or DirectShow components.

### Note

- DirectShow allows hardware and software vendors to create individual multimedia modules called filters. Multiple filters can be connected by the use of pins and a filter graph. TAPI connects different components, and DirectShow uses the same approach to enable filters and devices to connect to each other.

**PPP over ATM**

With the advent of Digital Subscriber Line (DSL) technologies, high-speed network access from the home and small office environment is becoming more of a reality. Several standards are being developed in these areas, including Asymmetric DSL (ADSL) and Universal ADSL (UADSL or DSL Lite). These technologies operate over a local loop, typically copper wire running between the public telephone network and the home. In most areas, this local loop connects directly to an ATM core network run by a telephone company.

Without changing protocols, ATM over the DSL service preserves the high-speed characteristics and QoS guarantees, which are available in the core. These guarantees create the potential for an end-to-end ATM network to the residence or small office. This network model provides several advantages, including:

- Protocol transparency
- Support for multiple classes of QoS with service guarantees
- Bandwidth scalability
- An evolution path to newer DSL technologies

Adding the Point-to-Point Protocol (PPP) over this end-to-end architecture adds functionality and usefulness. PPP provides the following additional advantages:

- Authentication
- Open System Interconnection (OSI) Layer 3 address assignment
- Multiple concurrent sessions to different destinations
- OSI Layer 3 protocol transparency
- Encryption and compression

These enhancements provide high bandwidth, even over a telephone line with an ATM adapter and this new level of integrated ATM support. In addition, with the adoption of PPP over ATM, little change is required at the ISP level, as telephone companies and ISPs typically use PPP.

Finally, if each VC carrying a PPP session carries only one session, each destination has its own authenticated PPP session, providing per-VC authentication. This provides an extra measure of security. Using Null Encapsulation over AAL5 can further reduce overhead because the protocol multiplexing is provided in PPP.

**PPP over ATM and NDISWAN**
In earlier versions of Windows operating systems, the Network Driver Interface Specification Wide Area Network (NDISWAN) component both supported operation of standard protocol stacks over WAN media and acted as the PPP engine. As explained earlier, in Windows Server 2003 the NDISWAN component has been extended and the TAPI proxy component added to provide this same support over NDIS 5.0 connection-oriented media, such as ATM.

At initialization, the NDISWAN component, acting as a client to the TAPI proxy, registers itself as the stream handler for PPP data. When the user starts dial-up networking to connect to a network, the dial-up networking module communicates with TAPI to make the phone call. When the request is made on an ATM device, TAPI does two things:

- Using the TAPI proxy and NDIS 5.0, it uses a call manager to make the telephone call by way of the ATM adapter.
- When the call connects, it redirects the connection from the adapter by using NDIS 5.0 to NDISWAN.

NDISWAN then handles further network (PPP) negotiation, and ultimately, by means of the LAN and TCP/IP stacks, it connects the user's computer to the remote network. The important thing to note here is that TAPI makes the call and then gives the resulting connection to another process, in this case NDISWAN.

This connectability enables several new types of applications, such as DVD-quality streaming video, real-time process control, a common standard for both LAN and WAN, and integrated software that pulls together aspects of TV, telephony, and data streams. Many of these applications make use of the raw channel access filter and DirectShow technology.

**Support for Raw Channel Access Filtering: DirectShow**
DirectShow technology was developed to better integrate multimedia services and to help multimedia developers more easily customize the operating system to their needs. DirectShow allows hardware and software vendors to create individual multimedia modules called filters. Multiple filters can be connected by the use of pins and a filter graph. TAPI connects different components, and DirectShow uses the same approach to enable filters and devices to connect to each other.

The following figure, "Windows COM-Based DirectStreaming," shows Windows COM-based DirectStreaming by which a Windows Server 2003-based application can handle many categories of real-time inputs.

**Windows COM-Based DirectStreaming**




DirectShow has an RCA filter — a simple module that exposes the raw data, whether it is voice, video, or other, to any device that wants to handle it. With NDIS 5.0, the RCA filter can be connected to TAPI. NDIS 5.0 can export ATM VCs as DirectShow pins.

**Raw Channel Access Filtering**
The Windows Server 2003 support for raw channel access filtering comes by means of the CoNDIS 5.0 driver. The NDIS proxy sets up a call at the ATM layer, running from the proxy to the call manager and on to the client. Unlike many voice or video feeds, this connection is made by using AAL5 rather than AAL1. The analog data transferred through the filter becomes digital data, which is packaged and handled as any other data over an AAL5 connection.

**DirectShow Application**
An example of using DirectShow and raw channel access support in Windows ATM services is a video-streaming application that delivers current weather information over the telephone. Customers can simply dial a number for the recorded information.

The following steps outline this process:

1. At initialization, the raw channel access filter registers as the stream handler for voice data.

2. A user calls a number to get the current weather information.

3. TAPI receives the call. TAPI redirects the incoming call to the raw channel access filter because it is a voice call.

4. NDIS 5.0 maps the DirectShow pin to the VC number.

5. DirectShow searches the filter graph, and the stream starts.

The following figure, "Weather Report Application Using DirectShow RCA Filter," shows how a weather report application uses the DirectShow raw channel access filter. The path shows how the data is routed through the various protocols from telephony data to the application layer.

**Weather Report Application Using DirectShow RCA Filter**




**IP Phone Access**

Similarly, a user can make a telephone call that is routed across a traditional LAN. This can enable such things as IP-based telephones. Again, TAPI handles the incoming call and uses NDIS 5.0 to connect it to a pin. DirectShow then reformats the data using a real-time protocol filter that goes through UDP/IP to reach an Ethernet card. The resulting connection allows a telephone user to talk to a computer user. This ATM-based network integration crosses previous boundaries between telephone and computer networks.

**ATM Cell Structure**

At either a private or a public User-Network Interface (UNI), an ATM cell always consists of a 5-byte header followed by a 48-byte payload. The header is composed of six elements, each detailed in the following figure, "Cell Header Structure."

**Cell Header Structure**




**Generic Flow Control**

The Generic Flow Control (GFC) field is a 4-bit field that was originally added to support the connection of ATM networks to shared access networks such as a Distributed Queue Dual Bus (DQDB) ring. The GFC field was designed to give the User-Network Interface (UNI) 4 bits in which to negotiate multiplexing and flow control among the cells of various ATM connections. However, the use and exact values of the GFC field have not been standardized, and the field is always set to 0000.

**Virtual Path Identifier**

The Virtual Path Identifier (VPI) defines the virtual path for this particular cell. VPIs for a particular virtual channel connection are discovered during the connection setup process for switched virtual connection (SVC) connections and manually configured for permanent virtual connection (PVC) connections. At the UNI, the VPI length of 8 bits allows up to 256 virtual

paths. VPI 0 exists by default on all ATM equipment and is used for administrative purposes such as signaling to create and delete dynamic ATM connections.

**Virtual Channel Identifier**
The Virtual Channel Identifier (VCI) defines the virtual channel within the specified virtual path for this particular cell. Just as with VPIs, VCIs are also discovered during the connection setup process for switched virtual connection (SVC) connections and manually configured for permanent virtual connection (PVC) connections. The VCI length of 16 bits allows up to 65,536 virtual channels for each virtual path. VCIs from 0 to 15 are reserved by the International Telecommunication Union ITU and VCIs from16 to 32 are reserved by the ATM Forum for each virtual path. These reserved VCIs are used for signaling, operation and maintenance, and resource management.

The combination of VPI and VCI values identifies the virtual channel for a specified ATM cell. The VPI/VCI combination provides the ATM forwarding information that the ATM switch uses to forward the cell to its destination. The VPI/VCI combination is not a network layer address such as an IP or IPX network address.

The VPI/VCI combination acts as a local identifier of a virtual channel and is similar to the Logical Channel Number in X.25 and the Data Link Connection Identifier (DLCI) in Frame Relay. At any particular ATM endpoint or switch, the VPI/VCI uniquely identifies a virtual connection to the next ATM endpoint or switch. The VPI/VCI pair need not match the VCI/VPI used by the final destination ATM endpoint.

The VPI/VCI combination is unique for each transmission path, that is for each cable or connection to the ATM switch. However, two different virtual channels on two different ports on an ATM switch can have the same VPI/VCI without conflict.

**Payload Type Indicator**
The Payload Type Indicator is a 3-bit field. Its bits are used as follows:

- The first bit indicates the type of ATM cell that follows. A first bit set to 0 indicates user data; a bit set to 1 indicates operations, administration and management (OA&M) data.

- The second bit indicates whether the cell experienced congestion in its journey from source to destination. This bit is also called the Explicit Forward Congestion Indication (EFCI) bit. The second bit is set to 0 by the source. If an interim switch experiences congestion while routing the cell, it sets the bit to 1. After it is set to 1, all other switches in the path leave this bit value at 1.

- Destination ATM endpoints can use the EFCI bit to implement flow control mechanisms to throttle back on the transmission rate until cells with an EFCI bit set to 0 are received.

- The third bit indicates the last cell in a block for AAL5 in user ATM cells. For non-user ATM cells, the third bit is used for OA&M functions.

**Cell Loss Priority**
The Cell Loss Priority (CLP) field is a 1-bit field used as a priority indicator. When it is set to 0, the cell is high priority and interim switches must make every effort to forward the cell successfully. When the CLP bit is set to 1, the interim switches sometimes discard the cell in congestion situations. The CLP bit is very similar to the Discard Eligibility bit in Frame Relay.

An ATM endpoint sets the CLP bit to 1 when a cell is created to indicate a lower priority cell. The ATM switch can set the CLP to 1 if the cell exceeds the negotiated parameters of the virtual channel connection. This is similar to bursting above the Committed Information Rate in Frame Relay.

**Header Error Control**
The Header Error Control (HEC) field is an 8-bit field that allows an ATM switch or ATM endpoint to correct a single-bit error or to detect multi-bit errors in the first 4 bytes of the ATM header. Multi-bit error cells are silently discarded. The HEC only checks the ATM header and not the ATM payload. Checking the payload for errors is the responsibility of upper layer protocols.

**Virtual Paths and Virtual Channels**
ATM uses virtual paths and channels to logically divide the bandwidth of the transmission path.

The following figure, “Channels within a Path inside the Transmission Medium,” shows how virtual paths and channels are divided within the transmission medium.

### Channels Within a Path Inside the Transmission Medium




**Transmission Path**
The transmission path consists of the physical cable connected to a particular port of an ATM switch. The cable has a defined bandwidth, such as 155 megabits per second for an Optical Carrier-3 (OC-3) optical fiber link.

**Virtual Path**
The bandwidth of the transmission path is logically divided into separate virtual paths and identified using the VPI in the ATM header. Each virtual path is allocated a fixed amount of bandwidth. Virtual paths do not dynamically vary their bandwidths beyond what has been allocated.

**Virtual Channel**
The bandwidth of a virtual path is logically divided into separate virtual channels using a virtual channel identifier in the ATM header. Unlike virtual paths, virtual channels share the bandwidth dynamically within a virtual path.

**Switching Hierarchy**
The transmission path to virtual path to virtual channel hierarchy is the basis for ATM switching. ATM can switch cells at the transmission path, virtual path and virtual channel level.

**Switching at the transmission path level**
Switching at the transmission path level allows an ATM switch to determine which output port to use to forward the cell.

**Switching at the virtual path level**
Switching at the virtual path level allows entire groups of virtual channels to be switched at the same time. Virtual path switching is similar to the telephone system cross-connect switching of entire groups of telephone calls based on the area code of the phone number. The switching occurs based on the area code, not the 7-digit individual phone number.

When performing virtual path switching, an ATM only takes action on the virtual path identifier in the ATM cell header. This ability to bypass the rest of the header makes virtual path switching faster than virtual channel switching.

ATM virtual path switching most often occurs within the public networks of ATM service providers because this virtual path switching allows ATM service providers to aggregate bundles of virtual channels along high-speed backbone links. These aggregate channels create trunk line structures very similar to those used in telephone networks.

**Switching at the virtual channel level**
Switching at the virtual channel level allows for precise switching and bandwidth allocation. Virtual channel switching resembles switching a phone call to its final 7-digit location; that is, ATM switching is based on the entire VPI/VCI, just as the final phone switching is based on the entire 10-digit phone number (3-digit area code plus 7-digit individual phone number).

ATM virtual channel switching occurs within both private and public networks. A switch must analyze both the virtual path identifier and the virtual channel identifier to make a switching decision.

**Quality of Service**
As part of the negotiated connection, ATM endpoints establish a service contract that guarantees a specific quality of service. These Quality of Service (QoS) guarantees are not offered by traditional LAN technologies.

With a traditional LAN, any notion of service guarantee is based on priority, where one transmission receives delivery preference over others. Because the sending end station does not know the condition of the network or the recipient of the data prior to transmission (traditional LANs are connectionless), traffic is subject to delay at routers and elsewhere. These unforeseen delays make bandwidth availability and delivery times difficult to predict. While higher-priority traffic typically reaches its destination prior to lower-priority traffic, it is possible that the higher priority traffic arrives too late for isochronous traffic.

### Note

- The term "QoS" applies to several forms of quality guarantees, including ATM QoS and generic quality of service. Of the three, only ATM QoS is implemented at the hardware level.

ATM offers precise and explicit service guarantees that are not based on a relative structure (such as priority). With ATM, a data supplier can request a specific bandwidth, maximum delay, and delay variation tolerance. Each ATM switch then determines whether or not it can meet the request after taking current allocations into consideration. If it can accommodate the transmission, it guarantees the service level and allocates the necessary resources. With ATM, the service contract is enforced, and the bandwidth is allocated at the hardware level. All switches between the sender and receiver agree to the service level

before the contract is granted. The source station hardware, also having agreed to the contract, is responsible for shaping the traffic to fit the connection contract before it enters the network.

ATM offers five service categories.

**Constant Bit Rate (CBR)**
Specifies a fixed bit rate. Data is sent in a steady stream with low cell loss. This is an expensive service because the granted bandwidth must be allocated, whether or not it is actually used. CBR is typically used for circuit emulation. This category is supported in Windows Server 2003.

**Variable Bit Rate (VBR)**
Specifies a throughput capacity over time, but data is not sent at a constant rate. This also specifies low cell loss. It is available in two varieties, real-time VBR for isochronous applications and non-real-time VBR for all others.

**Available Bit Rate (ABR)**
Ensures a guaranteed minimum capacity but allows data to be sent at higher capacities when the network is free. ABR adjusts the rate of transmission based on feedback, which specifies low cell loss. ABR provides better throughput than VBR, but is less expensive than CBR. It is important to note that ABR has only recently been fully defined, and not all hardware and software support this service category. It is part of the UNI 4.0 specification.

**Unspecified Bit Rate (UBR)**
Does not guarantee bandwidth or throughput, and cells can be dropped. A UBR connection does not have a contract with the ATM network. This category is supported in Windows Server 2003.

**Weighted Unspecified Bit Rate (WUBR)**
The newest service category put forward by the ATM Forum, and it functions by assigning different processing priorities to different types of traffic, similar to a traditional connectionless LAN. Each such type of traffic is carried across a different connection; cells in connections with lower priority are dropped before those with higher priority.

Guaranteed QoS allows ATM to support time-sensitive (isochronous) applications, such as video and voice, as well as more conventional network traffic. While 100-megabit Ethernet and other high-speed networks can provide comparable bandwidth, only ATM can provide the QoS guarantees required for real-time telephony, VCR-quality video streaming, CD-quality sound, smooth videoconferencing, and other delay-sensitive voice and video applications.

QoS is so vital to the industry that several initiatives are underway to provide QoS support for connectionless TCP/IP–based networks. While these solutions are useful, they require that all nodes on the network participate, which can be difficult to guarantee on heterogeneous networks. Because these solutions shape the traffic in software, latency and variations in delay are sometimes introduced. This is not the case with ATM.

Most importantly, the acceptance of ATM as a common standard for both LANs and WANs enables an enterprise to deploy QoS applications and integrated services. The deployment of ATM/Asymmetric Data Subscriber Line (ADSL) to the home enables residential access to these services. ADSL uses existing copper twisted-pair telephone lines to transmit broadband data to the home, without requiring recabling or a new telephone infrastructure. This extends the reach of ATM networks from the home desktop to the business desktop and everywhere in between.

**ATM Addresses**
ATM addresses are needed to support the use of virtual connections through an ATM network. At the simplest level, ATM addresses are 20 bytes in length and composed of three distinct parts. The following figure, "Simplified View of ATM Addressing," shows the three parts of the 20-byte ATM address.

### Simplified View of ATM Addressing




This ATM address breaks down into the following three basic parts:

**ATM switch identifier**
The first 13 bytes identify a particular switch in the ATM network. The use of this portion of the address can vary considerably depending on which address format is in use. Each of the three major ATM addressing schemes in use provides information about ATM switch location differently. The three formats are the data country/region code (DCC) format, international code designator (ICD) format, and the E.164 format proposed by the ITU Telecommunication Standardization Sector (ITU-T) for international telephone numbering use in broadband Integrated Services Digital Network (ISDN) networks.

**Adapter MAC address**
The next 6 bytes identify a physical endpoint, such as a specific ATM adapter, using a media access control (MAC) layer address that is physically assigned to the ATM hardware by the manufacturer. The use and assignment of MAC addresses for ATM hardware is identical to MAC addressing for other Institute of Electrical and Electronic Engineers (IEEE) 802.*x* technologies, such as Ethernet and Token Ring.

**Selector**
The last byte is used to select a logical connection endpoint on the physical ATM adapter.

Although all ATM addresses fit this basic three-part structure, there are significant differences in the exact format of the first 13 bytes of any given address, depending on the addressing format that is being used or whether the ATM network is for public or private use.

In summary, the 20-byte ATM address is in the hierarchical format starting with the switch at the highest level, down to the adapter, and then down to the logical endpoint.

All of the three ATM address formats that are currently in widespread use (DCC, ICD, and E.164) include the following characteristics:

- Compliance with the Network Service Access Point addressing plan as proposed by the Open Standards Interconnection protocol suite of the International Standards Organization.
- Each can be used to establish and interconnect privately-built ATM networks that support switched virtual connections (SVCs).

**Addressing in Detail**
The type of ATM address used depends on whether the addresses are for a public or private ATM network. ATM addresses are used to establish virtual channels between ATM endpoints. The following figure, "Primary ATM Address Formats," shows the three primary address formats.

**Primary ATM Address Formats**




The most important fields from these address formats are listed in the following table, "Primary ATM Address Format Fields."

**Primary ATM Address Format Fields**

| Address Fields | Function |
|---|---|
| AFI | The single-byte authority and format identifier (AFI) identifies the type of address. |
| DCC | The defined values are 45 for E.164, 47 for ICD and 39 for DCC addresses. |
| AA | This single byte identifies the-specific part (DSP) of the address. |
| Reserve | Reserved for future use. |
| RD | 2 bytes of routing domain information. |
| Area | 2 bytes of area identifier. |
| ESI | 6 bytes of end-system identifier, which is an IEEE 802.x MAC address. |
| SEL | 1 byte of NSAP selector. |

| | |
|---|---|
| ICD | 2 bytes of ICD. |
| E.164 | 8 bytes (16 digits) of the ISDN telephone number. |

**ATM Connection Types**
ATM connections between endpoints are not distinguished only by their various Quality of Service parameters and the formats of their addressing schemes. They also fall into one of two larger categories: point-to-point connections and point-to-multipoint connections. The connection types that any particular ATM connection uses depend on how ATM signaling builds its connection.

**Signaling**
Signaling components exist at the end station and at the ATM switch. The signaling layer of ATM software is responsible for creating, managing, and terminating switched virtual connections (SVCs). There are two signaling standards:

**User Network Interface**
The ATM standard wire protocol implemented by the signaling software

**Network Network Interface**
The way one ATM switch signals another ATM switch

### ATM Signaling




**Point-to-Point Connection**
When process that uses ATM seeks to connect to another process elsewhere on the network, it asks the signaling software to establish an SVC. To do this, the signaling software sends an SVC creation request to the ATM switch using the ATM adapter and the reserved signaling VC. Each ATM switch forwards the request to another switch until the request reaches its destination. An ATM switch determines which switch to send the request to next based on the ATM address for the connection and the internal network database (routing tables) of the switch. Each switch also determines if the need of the request for service category and Quality of Service can be met. At any point in this process, a switch can refuse the request.

If all the switches along the path can support the virtual channel as requested, the destination end station receives a packet that contains the VC number. From that point on, the process that uses ATM can interact with the destination process directly by sending packets to the VPI/VCI that identify the specified VC.

The ATM adapter shapes data traffic for each VC to match the contract made with the ATM network. If too much data is sent for any reason, the ATM switch can ignore — and lose — the data in favor of providing bandwidth to another contract or set of contracts. This is true for the entire breadth of the network; if bandwidth or speed exceeds the limits established by the contract, any device, including the ATM adapter, can simply drop the data. If this happens, the end stations concerned are not notified of the cell loss.

**Point-to-Multipoint Connection**
Unlike a standard LAN environment, ATM is a connection-oriented medium that has no inherent capabilities for broadcasting or multicasting packets. To provide this ability, the sending node can create a virtual channel to all destinations and send a copy of the data on each virtual channel. However, this is highly inefficient. A more efficient way to do this is through point-to-multipoint connections. Point-to-multipoint connects a single source endpoint, known as the root node, to multiple destination endpoints, known as leaves. Wherever the connection splits into two or more branches, the ATM switches copy cells to the multiple destinations.

Point-to-multipoint connections are unidirectional; the root can transmit to the leaves, but the leaves cannot transmit to the root or to each other on the same connection. Leaf-to-node and leaf-to-leaf transmission requires a separate connection. One reason for this limitation is the simplicity of AAL5 and the inability to interleave cells from multiple payloads on a single

connection.

**LAN Emulation**
LAN Emulation (LANE) is a group of software components that allows ATM to work with legacy networks and applications. With LAN emulation, you can run your traditional applications and protocols that use LAN on an ATM network without modification.

LANE makes the ATM protocol layers appear to be an Ethernet or Token Ring LAN to overlying protocols and applications. LAN emulation provides an intermediate step between fully exploiting ATM and not using ATM at all. LANE can increase the speed of data transmission for current applications and protocols when ATM is used over high speed media; unfortunately, LANE does not take advantage of native ATM features such as QoS. However, LANE does allow your current system and software to run on ATM, and it facilitates communication with nodes attached to legacy networks.

**LANE Architecture**
LANE consists of two primary components: the LANE client and the LANE services. The LANE client allows LAN protocols and applications that use LAN to function as if they were communicating with a traditional LAN. It exposes LAN functionality at its top edge (to users) and native ATM functionality at its bottom (to the ATM protocol layers).

The LANE services are a group of native ATM applications that hide the connection-oriented nature of ATM from connectionless legacy protocols. These services maintain the databases necessary to map LAN addresses to ATM addresses, thus allowing the LANE clients to create connections and send data.

The LANE services components can reside anywhere on an ATM network, but most ATM switches are included with LANE services components installed. Therefore, for practical purposes, LANE services reside on an ATM switch or group of switches.

The three primary LANE services are the LAN emulation configuration server (LECS), the LANE server (LES), and the Broadcast and Unknown server (BUS). The LECS distributes configuration information to clients, allowing them to register on the network. The LES manages one or more Emulated LANs (ELANs). It is responsible for adding members to the ELAN, maintaining a list of all the members of ELAN, and handling address resolution requests for the LANE clients. The BUS handles broadcast and multicast services, as shown in the following figure. "LANE Client, LECS, LES, and BUS."

**LANE Client, LECS, LES, and BUS**




**Lane Client Interaction with ATM Network**
When the LANE client seeks to join the network, the first thing it must do is find the LECS because the LECS gives the client the address of the LES managing the ELAN that it seeks to join. Without the LES address, the client cannot communicate with other members of the ELAN. At initialization the client has neither established a connection to any ATM switch nor to the switch or other entity containing the LECS. The client must establish an ATM connection, preferably a connection directly to the configuration server.

If the ATM network only has a single ATM switch, and the switch contains all the LANE services, then finding the LECS is easy. However, if the network has multiple switches, the local switch to which the LANE client has immediate access might not have LANE services running on it. Fortunately, LANE includes several established mechanisms for a LANE client to discovery the LECS.

**LECS Discovery**
The LANE client can use any of the following techniques when attempting to connect to the LECS:

- It can use a well-known ATM address, defined in the ATM protocol.
- It can use a well-known VC.
- It can query using the Integrated Local Management Interface (ILMI).

Both the well-known ATM address and the well-known VC are standardized. Most switches and clients are preconfigured with this information. In most cases, the LANE client can find the LECS using one of these methods. However, if the well-known

values have been changed at the end station or at the switch, either type of discovery becomes unsuccessful.

If this happens, the LANE client can fall back on ILMI, a protocol standard (similar to Simple Network Management Protocol) designed for ATM administrative and configuration purposes. ILMI provides a query function that the LANE client can use to find the LECS address, and then set up a VC to it.

After the client has discovered the LECS and connected to it, the client asks the LECS to provide configuration information to allow it to connect to a particular ELAN. It does this by sending one or more pieces of information about the desired ELAN, such as the LAN type (Ethernet or Token Ring), the maximum packet size, and the name of the LAN.

The LECS takes the information from the LANE client and reads its table of ELANs to find a match. When it locates the correct ELAN, it returns that address to the LANE client.

**LES Address Matching**
With the information provided by the LECS, the LANE client can join the ELAN. To do this, it sends an emulated LAN address and its true ATM address to the LES. The LES registers this information. From then on, the LANE client can send and receive data over the ATM network as if it were using a normal LAN.

When the LANE client receives a request from a protocol (such as TCP/IP, IPX, or NetBEUI) to send information to another point in the ELAN, it sends the destination LAN address to the LES. The LES looks for a match in its database, and then returns the true ATM address to the LANE client. The client then sets up a normal VC between itself and the destination, and subsequent data traffic is sent directly on this VC without any further intervention by the LES or the other LANE services. While this address resolution request is being processed, interim traffic is sent to the BUS and copied from there to all stations in the ELAN.

If the LES does not find a match for the destination address, the data is sent to the Broadcast and Unknown Server (BUS). The BUS attempts to deliver the data to the unknown client.

**BUS Distribution**
The BUS does two things: it handles distribution of data to unknown clients and it emulates LAN broadcast services. If the LES cannot find a particular ELAN client, the data is sent to the BUS for distribution, and the BUS forwards it to all the clients of the ELAN.

The BUS also handles broadcasts. It registers its address with the LES identical to any other client. It registers under the address of F (x16), which is the normal LAN address for a broadcast message. When a LANE client protocol wants to broadcast a message to the entire LAN, it addresses the message to F (x16) and passes it on. The LEC sends this address to the LES for resolution, and the LES returns the ATM address of the BUS. The LEC can then send the message to the BUS. The BUS maintains a list of all clients on the ATM network and sends the message to all clients. The BUS service is typically located in the same piece of equipment with the LES.

**Integrated Local Management Interface**
The ILMI resides on an ATM switch and provides diagnostic, monitoring, and configuration services to the User-Network Interface. The ILMI is defined by the ATM Forum and uses the Simple Network Management Protocol (SNMP) and a Management Information Base (MIB). It runs over AAL3/4 or AAL 5 with a default VPI/VCI of 0/16.

The ILMI MIB contains data describing the physical layer, the local VPCs and VCCs, network prefixes, administrative and configuration addresses, ATM layer statistics, and the ATM layer itself. The most common client-oriented function of the ILMI is to assist a client during LECS discovery.

**TCP/IP over ATM**
The protocol for classical IP over ATM (sometimes abbreviated as CLIP/ATM) is a well-established standard spelled out in RFC 1577 and subsequent documents. Windows Server 2003 provides a full implementation of this standard.

The IP over ATM approach provides several advantages over ELAN solutions. The most obvious advantages are its ability to support QoS interfaces, its lower overhead (as it requires no MAC header), and its lack of a frame size limit.

**IP over ATM Architecture**
IP over ATM is a group of components that do not necessarily reside in one place. When they don't reside in one place, the services are not typically on an ATM switch. For the purposes of this discussion, it is assumed the IP over ATM server services reside on a Windows Server 2003 server.

The core components required for IP over ATM are roughly the same as those required for LANE, as both approaches require the mapping of a connectionless medium to a connection-oriented medium, and vice versa. In IP over ATM, these services are provided by an IP ATMARP server for each IP subnet. This server maintains a database of IP and ATM, and provides configuration and broadcast services.

**IP over ATM Components**
IP over ATM is a very small layer between the ATM protocol and the TCP/IP protocol. As with LANE, the client emulates

standard IP to the TCP/IP protocol at its top edge while simultaneously issuing native ATM commands to the ATM protocol layers underneath.

IP over ATM is often preferred to LANE because it is faster than LANE. A primary reason for this performance advantage is that IP over ATM adds almost no additional header information to packets as they are handed down the stack. After it has established a connection, the IP over ATM client typically can transfer data without modification.

As with LANE, IP over ATM is handled by two main components: the IP over ATM server and the IP over ATM client. The IP over ATM server is composed of an ATMARP server and Multicast Address Resolution Service (MARS). The ATMARP server provides services to map network layer IP unicast addresses to ATM addresses, while MARS provides similar services for broadcast and multicast addresses. Both services maintain IP address databases just as LANE services do.

The IP over ATM server can reside on more than one computer, but the ATMARP and MARS databases cannot be distributed. You can have one IP over ATM server handle ATMARP traffic, and one handle MARS. If, however, you divided the ATMARP Server between servers, it effectively creates two different IP networks. All IP over ATM clients in the same logical IP subnet (LIS) need to be configured to use the same ATMARP server. Traditional routing methods are used to route between logical IP subnets, even if they are on the same physical network.

Windows Server 2003 includes fully integrated ATMARP and MARS servers.

**ATMARP Server**
The IP over ATM client and ATMARP server go through a process similar to the LANE client and the LECS when a client joins the network and discovers other network members. As with LANE, after an address is found, native ATM takes over and TCP/IP packets are sent across a VC from end station to end station. There is, however, a major difference in how the IP over ATM client discovers the ATMARP server.

**ATMARP Server Discovery**
Because the ATMARP server typically resides on a server rather than on an ATM switch, it is not possible to use ILMI or a well-known VC to discover its address. In fact, there is no default IP over ATM mechanism for server discovery. To start using IP over ATM, an administrator must find the ATM address of the appropriate ATMARP server and manually configure each IP over ATM client with this address. In a single ATM switch network, this is not much of a problem, but in larger networks it can become a demanding job. To ease configuration in smaller networks, ATM ARP/MARS servers and ATM ARP/MARS clients running Windows Server 2003 use a default address.

After the ATMARP server has been discovered, the IP over ATM client can use this server to resolve IP to ATM address mappings and communicate with other computers. The ATMARP server supports only unicast traffic. To send packets to a broadcast address or multicast list, the IP over ATM client uses MARS.

**MARS**
Mimicking the role of the BUS in LAN emulation, MARS handles distribution of broadcast and multicast messages to all the members of the network or to all members of a multicast group. Because of the potential for blockage, MARS provides two modes of operation. When an ATMARP client receives a request to send a packet to a multicast or broadcast IP address, it sends a request to the MARS to resolve this address to a list of clients that are members of that group. In one mode of operation, the MARS return a list of all ATM addresses to which the group address resolves. The client then creates a Point-to-Multipoint (PMP) ATM connection to all of these addresses and forwards the packet on that connection.

The other mode of operation involves a multicast server (MCS). The MCS registers interest in one or more multicast groups with the MARS. The MCS receives information describing the membership of that group, as well as updates when clients join or leave that group. When a client requests a group address resolution from the MARS, the MARS simply returns the single address of the MCS. The packet is then sent to the MCS, which creates the PMP connection and distributes the packet to all members of the group.

The following figures are examples of the VCs that are created in each of these modes.

**IP Multicast over ATM Connections Without MCS**




**IP Multicast over ATM Connections with MCS**




The disadvantage of the first method, in which each client sending packets to the group creates its own PMP connection to all other members of the group, is the large number of virtual channels required. The disadvantage of the second method, which uses the MCS, is that the MCS becomes both a central point of failure and a potential blockage because it distributes all multicast packets for all of the groups it serves.

**IP over ATM Operation**

IP over ATM faces the same problems, and relies on the same basic tools and fixes as LANE. In particular, it faces the issues of address resolution and broadcasting.

In normal ATM, SVC connections are established by sending a connection request containing the ATM address of the destination endpoint to the ATM switch. Before an IP endpoint can create an SVC in this manner, the endpoint must resolve the IP address of the destination to an ATM address.

Typically, when an Ethernet host needs to resolve an IP address to an Ethernet MAC address, it uses an ARP broadcast query frame. Hardware broadcasting is not done in ATM. The Address Resolution Protocol (ARP) of the ATMARP server resolves IP addresses to ATM addresses.

An ATM endpoint that needs to resolve an IP address sends an ATMARP request to the ATMARP server for its LIS. The ATMARP request contains the ATM address, IP address, and the requested IP address of the sender. If the ATMARP server accepts the requested IP address, it sends back an ATMARP response containing the requested ATM address. If the requested IP address is not found, the ATMARP servers send back a negative ATMARP reply, unlike the procedure in an ELAN, which sends an unresolved address to the LANE BUS. This behavior allows an ARP requestor to distinguish between an unknown address and non-functioning ATMARP server.

The end result is a three-way mapping from the IP address to an ATM address to a VPI/VCI pair. The IP address and ATM address are required to create a VC. The IP address and VPI/VCI then are required to send the subsequent cells containing data across the VC.

An ATM endpoint creates SVCs to other ATM endpoints within its LIS. For an ATM endpoint to resolve an arbitrary IP address, it must be configured with the ATM address of the ATMARP server in its LIS.

When the server starts, an ATM endpoint establishes a VC with the ATMARP server using ATM signaling. As soon as the VC is opened with the server, the server sends the ATM endpoint an ATMARP request. When the ATM endpoint sends the response, the ATMARP server has the ATM and IP address of the new ATM endpoint. In this way, the ATMARP server builds its table of ATM to IP address mappings.

**IP over ATM Client Initialization**

In Windows Server 2003, IP over ATM does not require the use of an Inverse ARP. Instead, the client goes directly to the server to register itself on the network. Because the process is automatic, no human intervention is required to initialize the client. Depending on whether the client address is mapped to a static address or a dynamic address, the procedure varies between the following two approaches.

**With a static IP address**
The following example details each step in establishing an IP over ATM connection for a single IP over ATM client with a static IP address. First, the client initializes and gets an ATM address from the ATM switch. The client then connects to the ATM ARP/MARS server and joins the broadcast group. The IP-to-ATM address mapping of the client is also added to the ATMARP server database. The client is now ready to contact other hosts and begin data transfer.

**With DHCP**
Establishing an IP over ATM connection for a single IP over ATM client using Dynamic Host Configuration Protocol (DHCP) is similar but not identical. First the client initializes and gets an ATM address from the ATM switch. Then the client connects to the ATM ARP/MARS server and joins the broadcast group. The client connects to the multicast server (MCS) and sends a DHCP request. The MCS broadcasts the DHCP request to all members of the broadcast group.

When the DHCP server receives the request, it sends a DHCP reply to the MCS. The MCS then broadcasts the reply to the broadcast group. The client receives the DHCP reply and then registers its IP and ATM addresses with the ATM ARP/MARS server. The client is now ready to contact other hosts and begin data transfer.

**Logical IP Subnets**
A LAN-based IP internetwork consists of a series of cabling plants separated by IP routers. The cabling plants connect the hosts of an IP network or subnet together and the routers connect the networks and subnets to each other. An IP host on a particular network can send IP packets directly to a host on the same network by addressing the packet to the media access control (MAC) address of the destination host. An IP host can send IP packets to hosts on other networks by addressing the packet to the MAC address of the router. This paradigm can connect hundreds (or thousands) of hosts on the same network to hundreds (or thousands) more on other networks. While this configuration works well for connectionless, broadcast-based technologies such as Ethernet and Token Ring, it is important to use care when attempting to create the same situation with ATM. The following figure, "Two LISs Running on a Single Switch," shows two LISs running on a single switch.

**Two LISs Running on a Single Switch**




Before a single IP packet can be sent, a connection must be created between the source and destination at the ATM layer. On an ATM network using SVC, the path must be negotiated between switches so that the sender has a valid VPI/VCI address to which to send the ATM cells. While possible, it is not very practical to have hundreds (or thousands) of ATM endpoints on the same IP network. A host, such as a network server, cannot send an arbitrarily large number of VCs to other hosts in the network. More VCs mean more overhead and resources for the both the operating systems and the hardware (ATM adapters and switches) of the IP network. If connecting across an ATM service provider, more VCs also mean more cost.

The logical IP subnet (LIS) is a way of constraining the number of ATM endpoints in an IP network or subnet. The LIS is a group of IP hosts that share a common IP network number; these hosts communicate with each other directly using ATM virtual channels. Different logical IP subnets can be created on the same ATM switch to create a virtual IP internetwork.

Note the example in the previous figure, "Two LISs Running on a Single Switch." When hosts in LIS 131.107.56.0/24 are ready to communicate among each other, they establish a VC with each other (direct delivery). When hosts in LIS 131.107.56.0/24 are ready to communicate with hosts in LIS 131.107.68.0/24, they establish a VC with the router and send an IP packet to the router (indirect delivery). The router then establishes its own VC with the destination host and forwards the IP packet.

An IP router belongs to multiple LISs and is configured with multiple IP addresses and subnet masks. If the router has a single ATM interface (and therefore a single ATM address), it can either use the single ATM address (with the unique End System Identifier) or use multiple ATM addresses by varying the last byte in the 20-byte ATM address (the SEL field). To give clients a secondary point of access, a router uses multiple addresses primarily in the case of a server failure, when a route is no longer functional.

**Services at an ATM Switch**
The overall scheme of ATM protocols, hardware, and interconnections comes together as shown in the following figure, "Overview of ATM Architecture from Desktop to WAN." The hardware includes ATM clients, ATM switches, and the blades that function as edge devices between a device or application that uses ATM and another protocol or environment. The connections include SVCs and permanent virtual connections (PVCs), all handled by the local switch shown at the center of the diagram.

In addition, the following figure, "Overview of ATM Architecture from Desktop to WAN," shows LANE clients joining the local switch, as well as a client being redirected from Local Switch 2 by the LECS to the BUS and LES services on Local Switch 1. The bottom layer is the client section, showing all the ARP, ATM, and other clients of the local switch. (With the exception of a remote access client, which is shown dialing into the switch at the top of the diagram). The various forms of edge devices are shown in the central section, and remote connections are shown at the top.

**Overview of ATM Architecture from Desktop to WAN**




**Additional Resources**

The following resources contain additional information that is relevant to this section.

- ITU-T Q.2610 Broadband Integrated Services Digital Network Signaling on the ITU Web site.ITU-T Q.2931 Broadband Integrated Services Digital Subscriber Signaling System on the ITU Web site.
- RFC 1577 in the RFC Database for information about Classical IP and ARP over ATM.
- RFC 2225 in the RFC Database for information about ATM Signaling support for IP over ATM.
- RFC 1661in the RFC Database for information about ATMF RFC1661 PPP
- RFC 2022 in the RFC Database for information about Multicast in UNI 3.0 and UNI 3.1 Networks

# POP3 Service Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**POP3 Service Technical Reference**
The POP3 service enables a server to host e-mail accounts and includes tools to administer the servers, domains, and mailboxes.

The POP3 service performs the tasks of message download and request handling on a Microsoft Windows-based server, where message download consists of transmitting the messages from a folder in the file system to clients and request handling is performed according to the POP3 protocol, which defines how the server responds to requests sent from an e-mail client.

**In this subject**


- What Is POP3 Service?
- How POP3 Service Works
- POP3 Service Tools and Settings

# What Is POP3 Service?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is POP3 Service?**

**In this section**


- Common POP3 Scenarios
- Other E-Mail Retrieval Technologies
- POP3 Dependencies
- Related Information


The POP3 service supplements existing functionality in Windows Server 2003 provided by the SMTP service, which receives e-mail messages. The POP3 service makes e-mail messages available for download from a server, enabling a server to host e-mail accounts and provide basic e-mail access.

The POP3 service performs the tasks of message download and request handling on a Windows-based server, where message download consists of transmitting the messages from a folder in the file system to clients and request handling is performed according to the POP3 protocol, which defines how the server responds to requests sent from an e-mail client.

Basic e-mail access depends on the POP3 protocol implemented by the POP3 service. The POP3 protocol, described in RFC 1939, "Post Office Protocol - Version 3," a standard for retrieving e-mail messages from servers, is stable, widely used, and supports basic e-mail access. The POP3 protocol supports message download to a client after the Simple Mail Transfer Protocol (SMTP) delivers messages received by means of the Internet. SMTP transmits e-mail from the originating server to the destination server using a protocol defined in the proposed standard, RFC 2821, "Simple Mail Transfer Protocol." For more information about the POP3 and SMTP protocols, see RFC 1393 and RFC 2821 in the IETF RFC Database.

The POP3 service is a scaleable e-mail server that meets the basic e-mail access needs of businesses and organizations. Any client that supports RFC 1939 can be configured to use the POP3 service to download e-mail. The POP3 service supports the following client authentication methods: Active Directory integrated authentication, Local Windows accounts authentication, and Encrypted Password file authentication. Because the POP3 service mail store exists in the file system, message storage and retrieval do not require specific database technologies on the server.

**Common POP3 Scenarios**

The POP3 service is designed to meet the needs of businesses and organizations with a small number of users needing basic e-mail access. However, the POP3 service is capable of scaling up to meet the needs of Internet service providers (ISPs) offering basic e-mail access to a large number of users.

**Small Businesses and Organizations**

Many small businesses have an investment in a Windows Server infrastructure that is suitable for an e-mail platform like POP3. POP3 enables organizations to host e-mail services for their Internet domain and support basic e-mail service. E-mail clients receive e-mail from the POP3 server and send e-mail by using the SMTP service. Groupware and scheduling functionality are not available with the POP3 service.

**Internet Service Providers**

ISPs offer e-mail service to customers as an added incentive; therefore the cost of e-mail service must be low enough that it does not impact the bottom line. Because POP3 use is so widespread, system administrators are familiar with existing POP3-compatible technology. The POP3 service in Windows Server 2003 supports a variety of authentication methods and scales up well to meet the demands of the ISP. Additionally, ISPs with large user bases can take advantage of the Active Directory directory service to store account information.

**Other E-Mail Retrieval Technologies**

POP3 is not the only protocol to support message retrieval. Internet Message Application Protocol version 4 (IMAP4) is similar in purpose to POP3. The IMAP4 protocol defines an extensive command set, including message searching, downloading of individual message body parts, and message tagging. IMAP4 implementations are not as common as POP3 because of the complexity of the IMAP4 protocol. Consequently, the use of IMAP4 is not as widespread as POP3 although the use of IMAP4 is growing as users demand more features from ISPs. For more information about IMAP4, see RFC 3501, "Internet Message

Access Protocol -Version 4, Rev 1" in the IETF RFC Database.

Messaging API (MAPI) is a protocol used with Microsoft Exchange Server and various e-mail clients for feature-rich messaging, including groupware and scheduling. One of the capabilities of MAPI is message retrieval for clients, such as Microsoft Outlook 2003. MAPI provides access to messages and folders and the associated properties of messaging objects. Message retrieval occurs in the context of a remote procedure call (RPC) from the client to the Exchange Server store, with the message contents transmitted in the response to the RPC call.

**POP3 Dependencies**

The POP3 service has the following dependencies:

- **Authentication Methods.** Active Directory integrated authentication, Local Windows accounts authentication, or Encrypted Password file authentication are the authentication methods that may be used with the POP3 service. All methods allow plaintext authentication, permitting credentials to be transmitted without encryption. If Active Directory-integrated authentication is chosen, Active Directory must be in place and functioning properly.

- **NTFS or FAT file system.** The POP3 service uses the file system as the mail store and takes advantage of NTFS access control lists (ACLs) to secure the mailboxes. POP3 service supports other file systems but NTFS provides the access control features that secure the mailboxes, and is strongly recommended.

- **Mail delivery from the Internet.** The SMTP service is a component that is installed by the POP3 service and receives incoming mail for the POP3 service.

**Related Information**

The following resources contain additional information that is relevant to this section.

- NTFS Technical Reference

- RFC 1939, "Post Office Protocol - Version 3" in the IETF RFC Database

- RFC 2821, "Simple Mail Transfer Protocol" in the IETF RFC Database

- RFC 3501, "Internet Message Access Protocol -Version 4, Rev 1" in the IETF RFC Database

# How POP3 Service Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How POP3 Service Works**

**In this section**


- POP3 Architecture
- POP3 Protocol
- POP3 Physical Structure
- POP3 Processes and Interactions
- Network Ports Used by POP3
- Related Information


The POP3 service enables a server to host e-mail accounts and includes tools to administer the servers, domains, and mailboxes.

The POP3 service performs the tasks of message download and request handling on a Windows-based server, where message download consists of transmitting the messages from a folder in the file system to clients and request handling is performed according to the POP3 protocol, which defines how the server responds to requests sent from an e-mail client.

The SMTP service receives e-mail from the Internet, saves the e-mail to the Queue folder, and notifies the SMTP delivery service for POP3 of the arrival of e-mail. The delivery service then moves the e-mail to the POP3 mail store, where it is available for download to POP3 e-mail clients.

The ideal environment for the POP3 service consists of the following:

- Windows Server 2003 is installed and functioning properly.
- Server hardware is sized appropriately for the number of users and disk space requirements.
- A suitable authentication configuration for the installation is implemented, using, Active Directory integrated, local Windows accounts or encrypted password file authentication.
- The volume that holds the mail store is formatted with the NTFS file system.
- Quotas are enabled on the volume that will host the mailboxes.
- The Internet domain name is registered with a naming authority.
- The SMTP service is configured as a mail exchange host with the Internet Service Provider.
- Mail Exchanger (MX) records have propagated through the Domain Name System.

**Note**

- Mail Exchanger (MX) records and mail exchange host configurations are not required for all installations. Consult your ISP to determine how your server should be configured.

This section covers how the service works with existing services and components of Windows Server 2003. The POP3 protocol

is presented to explain the command set implemented by the POP3 service. The physical structure of the POP3 e-mail system and the mail store is shown to illustrate the processes and interactions of the POP3 service and the SMTP service.

**POP3 Architecture**

The POP3 service uses the SMTP service, an existing component of Internet Information Services (IIS), to provide inbound message delivery as well as outbound message sending for POP3 clients. The figure POP3 Architecture shows the modularity of the POP3 service and the relationship with IIS and the file system.

**POP3 Architecture**




The POP3 service is the interface between e-mail clients and the mail store. POP3 listens on TCP port 110 for connections from e-mail clients, authenticates the client, and manages the connection with the client.

The authentication store is the repository of user information needed to authenticate the user. The store can be the Active Directory database, the local SAM database, or the encrypted password file for the user. The authentication module accesses the authentication store to verify the credentials submitted by the client to the POP3 service.

The Mail Storage Access API is the common interface to the mail store for all processes. The POP3 service, the SMTP delivery service for POP3, and the POP3 Server Administrator use the API to access the mail store.

The mail store uses the file system for storage. The mail store is typically located on the same server as the POP3 service, but it should be located on a different local or network volume than the operating system to avoid potential disk space problems. For large mail stores, the mail store can be placed on a Network Attached Storage (NAS) device and accessed by one or more servers running the POP3 service. Even though the mail store is contained in the file system, it is accessed by using the Mail Storage Access API.

The SMTP delivery service for POP3 is the component that transfers e-mail from the SMTP service to the user mailboxes. The delivery service is notified by the SMTP service when new e-mail arrives. New e-mail is delivered to the mail store by means of the Mail Storage Access API.

The POP3 Server Administrator is the interface used by the tools to configure the server, domains, and mailboxes. The server administrator accesses the mail store by means of the Mail Storage Access API to create or modify domains and mailboxes. Modifications to the service itself are made through calls to the administration interface within the POP3 service.

With the exception of the authentication store, the installation of the POP3 service installs all of the components shown in the figure POP3 Architecture. The POP3 service installation copies the files to the server and registers the dynamic-link libraries. The file names are listed in the table POP3 Components along with a description of the file. The files are copied to the *systemroot*\system32\pop3server directory, unless noted otherwise, and the installation path is added to the system path variable.

The following table lists the components of the POP3 service.

**POP3 Components**

| Component | Description |
|---|---|

| | |
|---|---|
| P3admin.dll | The interface to the mail store that supports the creation, deletion, verification, and renaming of mailboxes and domains. The POP3 Server Administrator uses this functionality. |
| P3store.dll | The interface to the mail store that supports mailbox locking, message creation and deletion, and creation, modification, and retrieval of the hashed password. The POP3 service, SMTP delivery service for POP3, and POP3 Server Administrator use this functionality. |
| Pop3auth.dll | The interface to the authentication methods. |
| Pop3evt.dll | The interface to the System log. |
| Pop3svc.exe | The Microsoft POP3 Service, which is started by the Service Control Manager. Pop3svc initializes the socket and thread pools, starts the listener, and manages the client connections to the service. |
| Winpop.exe | Command line tool to administer the POP3 service. Installation location is *systemroot*\system32. |
| P3server.msc | POP3 service console file used when the Microsoft Management Console (MMC) console is used. Installation location is *systemroot*\system32. |
| Pop3perf.dll | Contains methods that collect performance counter data from the POP3 service and make it available for display in the System Monitor. |
| Pop3pi.h | Header file that contains the list of performance counter indices used by the POP3 service. |
| Pop3snap.dll | The interface to the MMC. |
| Pop3perf.ini | Initialization files for the performance counters, which are not dynamically initialized. |

**POP3 Protocol**

The POP3 protocol is an application layer protocol, similar to the SMTP in the figure POP3 Protocol Architecture. Since POP3 uses TCP as the transport layer, it belongs on the same side of the figure as SMTP.

**POP3 Protocol Architecture**




RFC 1939, "Post Office Protocol -Version 3" defines a set of commands and responses for accessing a mailbox on a server and retrieving the messages. The POP3 service implements RFC 1939.

The POP3 protocol defines not only the commands and responses between a POP3 client and server but also the state of the session. The state of the session defines whether a POP3 command is valid in that context. The protocol defines commands for authentication, message retrieval, and message management in a series of three states. The three states of a POP3 session are:

- Authorization
- Transaction
- Update

POP3 commands submitted by the POP3 client are valid only in the context of the current state and the state changes as the session progresses. Typically, session state progresses from the authorization state, to the transaction state, and finally to the

update state. However, a session can go from the authorization state to being disconnected if a client sends the QUIT command while in the authorization state.

The authorization state is the state of the session once a connection has been established. The connection between the client and server uses the TCP three-way handshake to establish a network connection. The table Authorization State lists the commands sent by the client to the POP3 server.

**Authorization State**

| Command | Description |
|---|---|
| USER *username* | Requires an argument that specifies the mailbox name. The USER command is used when plaintext authentication is allowed. |
| PASS *password* | Must be preceded by the USER command. Requires the argument that specifies the password for the mailbox account. The response will indicate logon failure if an incorrect password was supplied, another user is already logged on to the mailbox, or the account is otherwise locked. |
| AUTH or AUTH [NTLM\|APOP] | If no argument is passed by the client, the response from the server contains the authentication methods in effect. The POP3 service accepts NTLM or APOP as valid arguments. |

The transaction state is the second state of the POP3 session. The commands valid in the transaction state provide information about the e-mail messages, return the contents of the messages, and mark the messages for deletion. The table Transaction State lists the commands sent by the client to the POP3 server. The commands in the table are valid when the POP3 session is in the transaction state.

**Transaction State**

| Command | Description |
|---|---|
| STAT | Returns a response with the number of messages in the mailbox and the byte count for each message. |
| LIST or LIST *n* | If no argument is supplied, the server returns a listing for each message containing the message length and message marker. If an argument is supplied, the server returns a listing for the specified message, it if exists. |
| RETR *n* | Returns the contents of the specified message, if it exists and has not been previously marked for deletion. Returns an error if no message id is specified. |
| DELE *n* | Returns a success response if the argument is an existing message not previously marked for deletion. This command does not delete the message. |
| NOOP | Returns a positive response but the server performs no action. |
| RSET | Returns a positive response after it removes the deletion mark from messages previously marked for deletion. |
| TOP *n* | Returns the number of specified lines of a message not previously marked for deletion. |
| UIDL or UIDL *n* | Returns a globally-unique identifier (GUID) for each message in the mailbox or a specified message, if the message was not previously marked for deletion. |

The update state is the last state of a POP3 session. The receipt of the QUIT command signals the server to delete any messages marked for deletion. The QUIT command can also be used in the authorization state to end a session before it enters the transaction state. In this situation, the connection to the server is terminated. For more information about the POP3 protocol, see RFC 1939.

**Note**

- Messages were marked for deletion in the transaction state by the DELE command. The actual deletion did not occur when the messages were marked for deletion.

The table Update State lists the single command sent by the client to end the session. The server displays the greeting message in response to this command.

**Update State**

| Command | Description |
|---|---|
| QUIT | Causes the greeting message to be displayed. Any messages previously marked for deletion by the DELE command are deleted when the QUIT is sent. |

**POP3 Physical Structure**

The physical structure of the POP3 mail system is shown in the figure Physical Structure of E-mail Systems. It shows both end systems running the POP3 and SMTP service. A message is shown as it travels from the sender through the Internet to the destination by using the protocols shown.

**Physical Structure of E-mail Systems**




The major components of the physical structure are client systems, mail servers, and the Internet. A client system can be a personal computer, laptop, or personal digital assistant. The client system must have a POP3 client capable of communicating with an SMTP service for sending mail and a POP3 service for receiving mail. The typical client configuration includes a setting for the SMTP server and the POP3 server, either as IP addresses or as resolvable host names.

The servers are shown with an SMTP service to send mail to and receive mail from the Internet. The Internet provides the backbone of networks used to carry the message from the originating system to the destination system. The delivery service, which is not shown, is used to deliver mail to the POP3 mailboxes.

The SMTP service and SMTP delivery service for POP3 work together to deliver a message sent by means of the Internet. As shown in the figure, all message transmission uses the SMTP protocol until the message reaches the destination host. The POP3 protocol is used by the POP3 service to download e-mail messages to the client.

**Various Installations of the POP3 Service**

The most common POP3 server installation is shown in the figure Typical POP3 Server Installation. The POP3 service and the message store are on one server and the mail store is on a volume on the same server.

**Typical POP3 Server Installation**




The clients shown in the figure connect to the single server. The POP3 service and the message store are on the server and the mail store is on a volume on the same server. Ideally, the mail store is located on a separate volume from the operating system. The processes and interactions of the server are contained within the server, with the possible exception of the authentication routines, which might make calls to other servers depending on the authentication method employed by the POP3 service.

The use of a remote store with the POP3 service is shown in the figure POP3 Server with Remote Store. The remote store is a

NAS device that provides large message stores with room to grow.

**POP3 Server with Remote Mail Store**




In this installation, the mail store is accessing a disk that is on a NAS device. This installation also supports one or more domains and all clients access a single server to retrieve e-mail. The advantage of this installation is the ability to grow the store as needed, as long as the hardware is capable of adequately supporting the user's. Any incoming e-mail to the POP3 users from the Internet is delivered to the SMTP service on the POP3 server and written to the mail store on the NAS device.

In the figure Multiple POP3 Servers, three servers access a remote store on a NAS device. The POP3 clients connect to different servers to balance the load across the servers. Adding more servers allows the POP3 service to support more users and provides a level of redundancy.

**Multiple POP3 Servers**




The only notable difference in this installation is that any incoming e-mail to the POP3 servers from the Internet is delivered to the SMTP server designated to receive e-mail for the domain, which then delivers the mail to the POP3 mail store. Any outbound messages are delivered to the Internet by the SMTP server designated for this purpose, although the same server typically receives and sends messages for a domain. Because there is more than one POP3 server, a server should be assigned the task of receiving inbound e-mail.

**Note**

- If you are using encrypted password file authentication, you must use the same globally unique identifier (GUID) on each mail server. To do so, select one mail server, identify its GUID, and then configure all other mail servers to use the same GUID.

**POP3 Processes and Interactions**

The mail delivery process is how e-mail travels from the originating server to the destination server. The process described in the figure POP3 E-mail Delivery shows the steps and the protocol used for each leg of the journey.

**POP3 E-mail Delivery**




In the figure POP3 E-mail Delivery, e-mail is sent to someone@example.com.

1. The e-mail is submitted for the user someone@example.com.

2. The SMTP service resolves the e-mail domain, example.com, to a mail server on the Internet, mailserver1.example.com.

3. E-mail is routed to the example.com domain where it is received by the SMTP service on mailserver1.example.com.

4. The SMTP service places the e-mail in the Queue folder on the destination mail server. The delivery service is notified of new e-mail for the domain example.com.

5. The delivery service moves e-mail to the recipients mailbox, P3_someone.mbx, in the mail store.

6. The user checks for e-mail by connecting to the mailbox for someone@example.com. The POP3 service verifies the credentials for the user and then either accepts or denies the connection.

7. If the user authenticates successfully, the e-mail is downloaded to the user's computer.

**Preparing to Retrieve E-Mail**

The POP3 service initializes its resources when the service is started by the Service Control Manager. Before the service can accept requests, it must do the following:

- Initialize the endpoint for connections and listen for requests

- Create a thread pool

- Create a socket pool

**Listening for Connection Requests**

In order for clients to connect to the service, the server must listen on some communication endpoint for a connection request. The POP3 service listens, by default, on TCP port 110 for incoming connections, unless it is configured to listen on another

port. The service expects clients connecting to the POP3 service to use the POP3 protocol during the session. When a TCP connection is established, the server sends a greeting message that identifies itself as a POP3 server. In the past, attempts have been made to exploit vulnerabilities in software using the version information included in the greeting message. The greeting message can be configured to hide system information.

**Creating a POP3 Service Thread Pool**
The POP3 service creates a pool of threads at initialization that wait on the I/O Completion Port for I/O Completion packets. I/O completion ports are used when thread pools process asynchronous I/O requests for an application. When an I/O Completion packet is posted to the port, the packet is passed to a thread to perform work for the packet. The I/O Completion packet has a handle to a data structure called the I/O Context. The I/O Context contains information about the type of I/O, callback functions, and other information specific to the type of object, such as an inactivity timer.

**Creating a POP3 Service Socket Pool**
The POP3 service creates a pool of sockets to handle incoming connection requests. While the service is running, the service maintains a pool of sockets sufficient to handle the connection requests while not consuming resources needlessly. There are four settings that are specific to the socket pool. The first setting is the minimum number of sockets that the POP3 service creates at startup. The *socketmin* setting is chosen to support a reasonable number of connections and allow the server to start up quickly. The default value of *socketmin* is 10.

The *socketthreshold* value is the number of unused sockets in the pool; the *socketthreshold* value is also the size of the block of sockets to create when new sockets are created. When the number of sockets remaining in the pool reaches the *socketthreshold* value, the server adds more sockets to the pool in anticipation of connection requests. If the threshold value is reached, the POP3 service creates new sockets in increments equal to the threshold setting size. If the *socketthreshold* value is greater than the remaining number of sockets until the maximum value is reached, the service will create only as many sockets as possible. The default value is 5.

The maximum number of sockets, *socketmax*, is the total number of sockets the POP3 service will create. As the POP3 service increase the size of the pool, the number of sockets will approach the maximum value. Once the maximum setting is reached, the POP3 service will not create any additional sockets. At this point, any incoming requests are placed in a queue. The size of the queue is determined by the *socketbacklog* setting. If the number of connections exceeds the number represented by *socketmax +socketbacklog*, any incoming connection requests are refused.

When a connection is closed, the socket is closed and the total number of sockets in use is decremented by 1. If the total number of available sockets is less than *socketthreshold* and *socketmax*, a new block of sockets is created. The server always maintains a sufficient number of sockets for incoming connections.

### Note

- The POP3 service manages connections to the server and automatically disconnects idle connections.

**Storing Messages**
The POP3 mail store maintains an organizational structure of message domains and mailboxes for the POP3 service, using the existing file system on the server. This organizational structure permits the POP3 service to host a single e-mail domain or many e-mail domains on a server. Mail domains are folders in the file system that exist at the same level in the directory and contain any mailboxes created in the domain. Mailboxes contain the e-mail messages for the recipient.

### Mail Store




The mail store is a directory on the local file system or a Universal Naming Convention (UNC) path, or the mail store can exist on a NAS device. In the latter configuration, multiple POP3 servers can access the same mail store. The POP3 service can use the FAT file system but the recommended file system is NTFS. By using NTFS, the mail store offers logging, recovery, and disk quotas. In the event of a server crash, NTFS can recover data by replaying its log files.

Disk quotas are a feature of the NTFS file system that can be helpful for managing the POP3 mail store. The POP3 mail store is

not a single-instance storage system, in which the server keeps only a single copy of a message and a separate table of recipients indexed by a message identifier. When a single-instance storage system stores a message sent to more than one recipient on a server, the recipient's access the same instance of the message in the mail store. When one recipient deletes the message, the message remains in the mail store until all other recipients have deleted the message.

When a message is received by the SMTP service, each recipient of a message receives a separate copy of the message from the delivery service. The use of disk quotas ensures that individual mailboxes and the mail store in general do not use excessive or unanticipated amounts of disk space and adversely affect the performance of the POP3 service or the server. A disk use policy that encourages users to download messages to a local disk and not use the POP3 mail store as a permanent storage location can be implemented through the use of disk quotas.

**Mail Root**
The mail root is the first organizational segment, shown in the figure Mail Store. The POP3 mail root is a directory path that identifies the first part of the path to the mail store. By default, the mail root is defined as *system_drive*\InetPub\Mailroot\Mailbox. All domains will be created under the mail root.

**Mail Domain**
The mail domain is the second segment of the organizational structure. When a domain is created by one of the configuration tools, the directory is created under the mail root. The domain name must be unique within the mail root and as a property of the SMTP service, where the domain is stored as a property in the IIS metabase. The domain name must:

- Have fewer than 255 characters
- Consist of alphabetic and numeric characters in the set {a-z, 0-9, -}
- Be registered with an Internet Naming Authority

**Mailbox**
The mailbox is the third segment of the organizational structure. The mailboxes are shown in the diagram as the folders on the right hand side. The mailboxes for a domain are logically grouped together by the domain because the path to the mailbox includes the domain name. Mailboxes for user accounts are created in a directory in a POP3 domain to hold the e-mail message delivered by the SMTP delivery service for POP3. The organizational structure permits mailbox names to be used more than once in different domains so a new domain can have mailboxes with the same names as those shown in the diagram. The mailbox folder name is P3_*mailbox name*.mbx, where the mailbox name must:

- Have fewer than 64 characters
- Cannot consist of characters in the set {@ ( ) / \ [ ] : ; , " . * = | ? + }
- Be unique in the context of the domain

When a mailbox is created, it has the following contents:

- Lock file
- Quota file

The mailbox lock file prevents simultaneous access to the mailbox by an e-mail client. The lock file is usually an empty file with one notable exception. The lock file contains the hashed password if Encrypted Password file authentication is enabled on the server. The hashed password is compared to the hash value sent by client and the user is granted or denied access based on the comparison.

The mailbox will also contain a quota file if the authentication method is Active Directory integrated or local Windows account authentication. For these two methods, the quota file is created when the mailbox is created. If Encrypted Password file authentication is the authentication method, the quota file must be created by using the command-line tool Winpop. The security identifier (SID) for the account is contained in the quota file and file ownership of all messages is assigned to the SID.

**Message Files**
The message files are the last item of the organizational structure. All e-mail message files are created using a naming convention with an .eml extension appended to identify the file as a message. Because the volume of messages delivered to the POP3 service can be large, a naming convention was implemented that allows for 10,000 messages per millisecond. The

naming convention names each e-mail message by using a timestamp consisting of the year, month, day, hour, minute, second, and millisecond plus a 4-digit counter. Using the message file format, message files appear as P3##..##.eml, where ##..## is a string of 26 digits derived by using the naming convention.

**Authenticating Clients**

The authentication module in the POP3 service is the interface to the authentication methods. The authentication module interacts with an abstraction referred to as an authentication database. The authentication process might refer to a file in the mailbox directory for authentication rather than query a database. The available choices for authentication method are:

- Active Directory integrated authentication
- Local Windows accounts authentication
- Encrypted Password file authentication

**Paths to Authentication**

The three possible paths to authentication are plaintext, NTLM, and authenticated POP (APOP). The plaintext authentication consists of the following sequence:

1. Client sends USER *mailbox*.
2. Server send +OK.
3. Client sends PASS *password*.
4. Server send +OK.

The NTLM sequence consists of the client sending an AUTH command, such as AUTH NTLM or AUTH APOP.

If the client sends the AUTH command with one of the authentication methods as an argument, the response and subsequent data are base-64 encoded to obscure the conversation. The client-server sequence still happens as described in RFC 1939, when AUTH is used but the commands and the responses must be converted to and from base64 encoding.

The ensuing discussion of authentication necessitates some mention of passwords being transmitted in the clear. NTLM requires the use of encrypted conversations during authentication. If NTLM is enabled, the plaintext transmission of passwords is not allowed. If the client sends the AUTH command, the server will respond with a listing of the accepted authentication mechanisms, APOP and NTLM. NTLM authentication is a challenge/response mechanism supported by the POP3 service. When the POP3 client connects to the server with the AUTH NTLM command, the server sends a random string of bytes (challenge) to the client, which the client encrypts with a hash of the user password. The client then sends the encrypted challenge back to the server where the response is decrypted with the password hash stored on the server. If the server is able to derive the original challenge from the client response, then the client is successfully authenticated.

APOP authentication is a challenge/response process used with POP3 and Internet Message Access Protocol 4 (IMAP4) servers. When a client connects to the mail server, the server appends an ASCII character string to the server greeting message. The client uses the string and the user password to create a hash value that is transmitted to the server, where it is compared to the value in the encrypted password file. When APOP is enabled, Encrypted Password file authentication uses the MD5 hashing algorithm in conjunction with CRAM, referred to as MD5-CRAM.

**Authentication Methods**

Active Directory integrated authentication is an option if the POP3 server is installed in an Active Directory domain on a domain controller or member server. Active Directory integrated authentication supports both plaintext authentication and NTLM authentication. The authentication store can be an Active Directory database or a global catalog derived from an Active Directory database.

Local Windows accounts authentication is an option when the POP3 server is not a member of an Active Directory domain and user accounts are created and maintained on the local server. Local Windows authentication supports both plaintext and NTLM authentication. When local accounts are created for POP3 user accounts, the authentication store is the local SAM database.

Encrypted Password file authentication is an option when the POP3 server is not a member of an Active Directory domain and user accounts are not created on the local server. A hash of the password is created using the Message Digest 5 (MD5) hashing algorithm and saved in the lock file in the mailbox when Encrypted Password file authentication is chosen as an authentication method. During authentication, the client sends a hash of the user password which is compared to the hashed password stored

in the lock file. This type of authentication is called a Challenge Response Authentication Mechanism (CRAM). If the hashed passwords match, the authentication of the user is successful. Encrypted Password file authentication supports both plaintext and APOP authentication.

**POP3 Service and Authentication Module**




The authentication process shown in the figure Architecture of POP3 Service and Authentication Module is a conceptual representation of the authentication.

**Note**

- MD5 is a hashing algorithm, based on its predecessors, MD2 and MD4. CRAM is a challenge-response authentication mechanism, based on the hashing of a password known only to the server and the user. When using the MD5-CRAM mechanism, a client transmits a digest of a password in response to a challenge from the server. The hashed password is sent across the network to be used in the authentication of a user.

**Requesting E-mail Delivery**

In the figure POP3 Service and Mail Store, the Mail Storage Access API provides an abstraction to the mail store and is shown as an interface to the mail store.

**POP3 Service and Mail Store**




If the client has successfully authenticated with the server, then the session progresses to the Transaction state. Other than the NOOP command, all of the commands sent by the client in the Transaction state require the POP3 service to access the mail store. The POP3 service reads from and writes to the mail store by means of the Mail Storage Access API.

The Mail Storage Access API creates an abstraction to the storage environment. This abstraction offers a consistent interface to the mail store and allows other storage possibilities in future releases. The Mail Storage Access API is called by other modules in the POP3 service, including the SMTP delivery service for POP3 and the POP3 Server Administrator.

When the client successfully authenticates with the service, the POP3 service modifies the lock file in the users mailbox directory to prevent another session from opening the same mailbox. The mailbox "lock" file is named "lock" and this file is created when the mailbox is created. When the mailbox is locked, the file name is reversed to "kcol." The mailbox lock does not interfere with the mail delivery service as it delivers messages.

**Delivering Messages**

In the figure POP3 Service and Mail Delivery, the delivery of a message from the Internet to the POP3 mail store is illustrated. The message is shown passing through the SMTP delivery service for POP3 to the mail store, where it can be downloaded

using the POP3 service.

**POP3 Service and Mail Delivery Module**




The SMTP delivery service for POP3 delivers messages to the individual mailboxes in the mail store. The delivery service replaces the default SMTP mail delivery module. When new mail arrives, the delivery service is notified by the SMTP service that the SMTP service has put messages in the SMTP Queue folder. The delivery service is an event sink that is triggered by the arrival of new mail. The SMTP delivery service for POP3 uses the same Mail Storage Access API to write the message files as the POP3 service uses to retrieve messages from the message store. If a message is delivered to a mailbox that is currently in use, the message is not available during this session and must be downloaded during a subsequent session.

The mail is delivered to the mailbox after it is retrieved from the Queue folder. A message sent to multiple recipients in a POP3 domain will be retrieved from the Queue folder as multiple messages, one for each recipient. The POP3 mail store does not implement single-instance message store, therefore each mailbox receives a copy of the message. Each message file is named using the naming convention described in the section on the physical structure of the POP3 service.

If quotas are in place and the user has exceeded the disk space quota, the message is returned to the sender if the **Deny disk space once the quota has been met** option is enabled. If the nondelivery report (NDR) cannot be sent to the original sender, the message is left in the Badmail folder.

**Administering the Server, Domain, and Mailbox**

The POP3 Server Administrator provides control over the POP3 service, domains, and mailboxes. Through the POP3 Server Administrator, all settings for the server, domain, and mailbox are available. The POP3 Server Administrator stores settings for the authentication method, socket and thread pool, and port number in the registry of the server. The architecture diagram does not depict the registry but these settings are registry entries.

In the figure POP3 Service and Server Administrator, the interaction with other components of the POP3 service is shown. The Server Administrator uses the Mail Storage Access API to make changes to mail store, such as creating new domains and users. The Server Administrator also creates and modifies settings in the IIS metabase.

**POP3 Service and Server Administrator**




**Creating a Domain**

The POP3 Server Administrator uses the Mail Storage Access API to create, modify, or delete domains in the mail store. When a new domain is created, the server administrator creates a directory entry for the domain under the mail root. There are three steps performed for the domain creation:

- Create a directory for the domain.
- Set the permissions on the directory.
- Add the domain name as an SMTP domain listing in the IIS metabase.

The new domain name is added as an SMTP property for SMTP in the IIS metabase. When a new domain is created to be used by multiple POP3 servers, the server administrator creates a directory entry for the domain and adds the domain name to the IIS metabase on the server. Subsequent installation of other servers accessing the same store will only add the domain name to the IIS metabase on that server.

**Deleting a Domain**

When a domain is deleted, the directory entry is deleted from the directory and the domain name is removed from the metabase. The deletion of a POP3 domain in an environment with multiple servers has some caveats that are worth noting. The first installation of a domain creates the mail store and adds the domain name to the metabase; the first deletion of a domain removes the mail store and deletes the domain name from the metabase. The domain name must be deleted from the metabase on other servers by deleting the domain on those servers.

**Creating a Mailbox**

The POP3 Server Administrator uses the Mail Storage Access API to create, modify, or delete mailboxes in the mail store. When a new mailbox is created, a directory is created under the domain for the mailbox. The following actions occur when the mailbox is created:

- Create the mailbox in a directory in the POP3 domain.
- Set permissions on the mailbox. When the mailbox is created, only local and domain administrators and the local network service have permissions to read from and write to the mailbox.
- Create an associated account. If the service is using Active Directory integrated or local Windows account authentication, an associated account should be created. No associated account is created for encrypted password file authentication.
- Add account to groups, if necessary. If the local Windows account authentication method is chosen, the user created for the associated account will be added to the “POP3 Users” group, which has no rights to log on locally to the server. If Active Directory integrated authentication is chosen, the user created for the associated account will have typical Active Directory domain user rights.

**Deleting a Mailbox**
When an existing mailbox is deleted, the mailbox and its contents are deleted from the server. The POP3 Server Administrator deletes the associated account in Active Directory or the SAM database only if the option to delete the associated account was specified. Otherwise, the administrator will be prompted about the deletion of an associated account. If Encrypted Password authentication is being used, no associated account was created.

**Locking or Unlocking Domain Mailboxes**
The POP3 Server Administrator can lock or unlock individual mailboxes or all mailboxes in a domain. When a mailbox is created, a lock file is created in the mailbox directory. The POP3 Server Administrator locks the mailbox by changing the name of the "lock" file from "lock" to "kcol." The POP3 service uses the same functionality to lock a mailbox when a session is underway to prevent simultaneous access to a mailbox.

The POP3 Server Administrator uses the Mail Storage API to access the mail store. The POP3 Server Administrator configures the IIS and SMTP service components through the IIS metabase.

**Network Ports Used by POP3**
By default, the POP3 service listens on the well-known TCP port 110 for incoming connections from e-mail clients. There is no UDP port for the POP3 service. The administrator of the POP3 service can specify a different TCP port using the administration tools, provided that the new port number is not already reserved for another protocol.

### Port Assignments for POP3

| Service Name | UDP | TCP |
|---|---|---|
| POP3 Service | None | 110 or an unassigned port chosen by the administrator. |

**Related Information**
The following resources contain additional information that is relevant to this section.

- RFC 1939, "Post Office Protocol -Version 3" in the IETF RFC Database
- RFC 2821, "Simple Mail Transfer Protocol" in the IETF RFC Database

# POP3 Service Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- POP3 Service Tools
- POP3 Service Registry Entries
- POP3 Service WMI Classes
- Network Ports Used by POP3 Service


The following tools, registry settings, and Windows Management Instrumentation (WMI) classes are associated with the POP3 service.

**POP3 Service Tools**
The following tools are used to administer the POP3 service.

**Winpop.exe: Winpop**
**Category**
Operating system tool

**Version compatibility**
All versions of Winpop in Microsoft Windows Server 2003 are identical.

Winpop is the POP3 service command-line tool that sets or retrieves the service settings for the POP3 service. Winpop can also perform the domain or mailbox activities essential to the operation of the POP3 installation.

For more information about Winpop.exe, see "Command Line References" in the Tools and Settings Collection.

**P3server.msc: POP3 snap-in**
**Category**
Operating system tool

**Version compatibility**
All versions of Microsoft Management Console (MMC) and the POP3 snap-in in Windows Server 2003 are identical.

The Microsoft Management Console POP3 snap-in provides a user interface to perform the same administrative activities as Winpop.exe. The POP3 snap-in lets the administrator see all of the domains and mailboxes graphically and administer POP3 remotely.

**POP3 Service Web Administration**
**Category**
Web interface component

**Version compatibility**
All versions of POP3 Service Web Administration in Windows Server 2003 are identical.

POP3 Service Web Administration provides remote management capabilities by using a Web browser interface as a client. From this interface, the administrator has access to settings for the Server Properties and the Domains and Mailboxes. POP3 Service Web Administration is automatically installed when you install the E-Mail Service from Add/Remove Windows Components, or you can install it manually from Add/Remove Windows Components. However, it is not installed by using Configure Your Server Wizard to configure the Mail server role. In order to use POP3 Service Web Administration, you need to make sure that Web Interface for Remote Administration is also installed on the same server. You can access the interface by entering `https://servername:8098` in a Web browser or in the Run dialog box at a client computer, where `servername` is the name of the server on which you installed POP3 Service Web Administration.

**POP3 Service Registry Entries**
The following registry entries are associated with the POP3 service.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values might be written to the registry. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**POP3 Server**
The following registry entries are located under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Pop3 Service\

**ConsoleFile**
**Registry path**
HKLM\SOFTWARE\Microsoft\Pop3 Service\

**Version**
Windows Server 2003

Stores the path to the MMC file (.msc), installed in the *windir*\system32 directory.

**CreateUser**
**Registry path**
HKLM\SOFTWARE\Microsoft\Pop3 Service\

**Version**
Windows Server 2003

Stores a value that indicates whether an associated user account will be created in Active Directory or the local Security and Accounts Manager (SAM) database when a mailbox is created. The value that is set depends on whether the **Automatically create a user** check box is selected.

**Greeting**
**Registry path**
HKLM\SOFTWARE\Microsoft\Pop3 Service\

**Version**
Windows Server 2003

Stores the greeting message for the POP3 service. The default message, **Microsoft Windows POP3 Server Version 1.0 <FQDN of server> ready,** can be changed to hide the system information. This key does not exist by default and must be created. For more information about creating this registry key, see "Advanced Mail Server Configurations" on the Microsoft Windows Server System Web site.

**InstallDir**
**Registry path**
HKLM\SOFTWARE\Microsoft\Pop3 Service\

**Version**
Windows Server 2003

Stores the path to the executables and DLLs installed by the POP3 server installation, with the exception of winpop.exe and p3server.msc, which are installed in the *windir*\system32 directory. The default installation path is *windir*\system32\pop3server.

**Logging Level**
**Registry path**
HKLM\SOFTWARE\Microsoft\Pop3 Service\

**Version**
Windows Server 2003

Stores the logging level for the POP3 service. The default value is 1 and the range of values is from 0 to 3, as shown in the following table. The events are written to the event log and can be viewed by using the Event Viewer.

| Value | Explanation |
|---|---|
| 0 | No events are logged by the POP3 service. |

| | |
|---|---|
| 1 | Only critical events are logged by the POP3 service. |
| 2 | Critical and warning events are logged by the POP3 service. |
| 3 | Critical, warning, and informational events are logged by the POP3 service. |

**MailRoot**
**Registry path**
HKLM\SOFTWARE\Microsoft\Pop3 Service\

**Version**
Windows Server 2003

Stores the path to the mail root directory, typically *systemdrive*\inetpub\mailroot. The POP3 e-mail domains are created as subdirectories underneath the mail root.

**RequireSPA**
**Registry path**
HKLM\SOFTWARE\Microsoft\Pop3 Service\

**Version**
Windows Server 2003

Stores the setting for Secure Password Authentication (SPA) on the server. The POP3 server checks the value at startup to determine if SPA is required for local Windows account or Active Directory authentication. If SPA is enabled, plain text transmission of passwords is not allowed. This key is not created until SPA is activated; it is disabled by default.

**Pop3Svc**
The following registry entries are located under HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Pop3Svc\

**Description**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the textual description as it appears in the Service Control Manager (SCM). The default value of the string is "Provide POP3 downloading service."

**DisplayName**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the display name of the POP3 service as it appears in the SCM. The default value of the string is "Microsoft POP3 Service."

**ErrorControl**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the default configuration setting for the POP3 service, used by the Service Control Manager (SCM). The default value is 1.

**ImagePath**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the image path value, the path to the POP3 service dynamic link library (DLL). The default value is

*windir*\System32\POP3Server\pop3svc.dll.

**MaxNumberofSockets**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the value of the maximum number of sockets POP3 service can create. The default value is 40.

**MinNumberofSockets**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the value of the minimum number of sockets the POP3 service creates at startup. The default value is 10.

**ObjectName**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the account name in which the POP3 service runs on the server. The default account name is "NT AUTHORITY\NetworkService."

**Pop3Port**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the port number the POP3 service listens on for incoming connection requests. By default, the POP3 service listens on TCP port 110 for incoming connections. If this value is changed, the port number must also be changed on the clients to match the server setting.

**SocketBacklog**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the value of the socket backlog that the POP3 service permits when the maximum number of sockets allowed is in use. The socket backlog is the number of pending connections for the POP3 service. When the number of sockets reaches *maxnumberof sockets+socketbacklog*, new client connections for the POP3 service are rejected. The default value is 5.

**SocketsThreshold**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows 2003 Server

Stores the sockets threshold value the POP3 service uses as a guideline. The sockets threshold setting determines how many sockets the POP3 service creates at a time. If the number of sockets in existence does not exceed the maximum number of sockets (*MaxNumberofSockets*), the server creates new sockets in a block size equal to the sockets threshold value. The threshold value also indicates when the POP3 service needs to create more sockets. If the number of available sockets reaches the threshold value, the POP3 service creates new sockets. The default value is 5.

**Start**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**

Windows Server 2003

Stores the number of attempts the SCM makes to start the POP3 Service. The default value is 3.

**ThreadCountPerCPU**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the thread count value for the POP3 service. By default, the POP3 service creates two threads per CPU to handle client requests. The number of threads per processor multiplied by the number of processors equals the number of simultaneous requests the POP3 service can process.

**Type**
**Registry path**
HKLM\System\CurrentControlSet\Services\Pop3Svc\

**Version**
Windows Server 2003

Stores the default configuration setting for the POP3 service used by the SCM. The default value is 16.

**Application**
The following registry entries are located under
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Eventlog\Application\Pop3Svc

### Note

- There are identical keys located under HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Eventlog\Application\Pop3 Service. These keys, however, are used by the Simple Mail Transfer Protocol (SMTP) service and are not used by the POP3 Service.

**CategoryCount**
**Registry path**
HKLM\System\CurrentControlSet\Services\Eventlog\Application\Pop3Svc

**Version**
Windows Server 2003

Stores the number of event categories supported by the POP3 service. The default value is 2.

**CategoryMessageFile**
**Registry path**
HKLM\System\CurrentControlSet\Services\Eventlog\Application\Pop3Svc

**Version**
Windows Server 2003

Stores language-dependent strings that describe the event categories. The default value for this string is *windir*\System32\POP3Server\pop3evt.dll.

**EventMessageFile**
**Registry path**
HKLM\System\CurrentControlSet\Services\Eventlog\Application\Pop3Svc

**Version**
Windows Server 2003

Stores language-dependent strings that describe the events. The default value for this string is *windir*\System32\POP3Server\pop3evt.dll.

**TypesSupported**
**Registry path**
HKLM\System\CurrentControlSet\Services\Eventlog\Application\Pop3Svc

**Version**

Windows Server 2003

Stores a value used to mask certain bits of an integer to indicate the type of events supported. The default value is 7, indicating that the first 3 bits are set to 1.

**POP3 Service WMI Classes**

The following table lists and describes a WMI class that is associated with the POP3 service. The class is in the *root\cimv2* namespace.

**WMI Classes Associated with POP3**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_PerfFormattedData_Pop3Svc_POP3Service | \root\cimv2 | Windows Server 2003 |

For more information that describes many WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by POP3 Service**

The POP3 service listens on the well-known TCP port 110 for incoming connections from clients. The administrator of the POP3 server might assign other values by using the administration tools, provided that the new port number is not already assigned to another protocol.

**Port Assignments for POP3**

| Service Name | UDP | TCP |
|---|---|---|
| POP3 service | None | 110 or an unassigned port chosen by the administrator. |

# QoS Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**QoS Technical Reference**
The increase in network traffic in recent years has intensified the competition for network resources. New applications, multimedia applications in particular, place high demands on networks and can have a negative impact on mission-critical applications, such as Enterprise Resource Planning (ERP) systems. Traffic bottlenecks and delays can frustrate end users and put network administrators in reaction mode.

Quality of Service (QoS) is an industry-wide group of standards and mechanisms that support preferential treatment for certain traffic flows. Microsoft Windows Server 2003 provides QoS mechanisms for prioritizing traffic as it is generated by an application. The application can be written to apply QoS parameters to the traffic, or traffic management tools can be used to apply the QoS parameters.

**In this subject**


- What Is QoS?
- How QoS Works
- QoS Tools and Settings

# What Is QoS?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is QoS?**

**In this section**


- QoS Concepts
- Uses and Benefits of QoS
- Network Technologies and Support for QoS
- Windows Server 2003 QoS
- Related Information


Quality of Service (QoS) for networks is an industry-wide set of standards and mechanisms for ensuring high-quality performance for critical applications. By using QoS mechanisms, network administrators can use existing resources efficiently and ensure the required level of service without reactively expanding or over-provisioning their networks.

Traditionally, the concept of quality in networks meant that all network traffic was treated equally. The result was that all network traffic received the network's best effort, with no guarantees for reliability, delay, variation in delay, or other performance characteristics. With best-effort delivery service, however, a single bandwidth-intensive application can result in poor or unacceptable performance for all applications. The QoS concept of quality is one in which the requirements of some applications and users are more critical than others, which means that some traffic needs preferential treatment.

**QoS Concepts**

The goal of QoS is to provide preferential delivery service for the applications that need it by ensuring sufficient bandwidth, controlling latency and jitter, and reducing data loss. The following table describes these network characteristics.

**Network Characteristics Managed by QoS**

| Network Characteristic | Description |
|---|---|
| Bandwidth | The rate at which traffic is carried by the network. |
| Latency | The delay in data transmission from source to destination. |
| Jitter | The variation in latency. |
| Reliability | The percentage of packets discarded by a router. |

The Internet Engineering Task Force (IETF) defines two major models for QoS on IP-based networks: Integrated Services (Intserv) and Differentiated Services (Diffserv). These models encompass several categories of mechanisms that provide preferential treatment to specified traffic. The following table describes the three general categories of QoS mechanisms.

**Note**

- Different versions of Windows support different combinations of these mechanisms. Windows XP and Windows Server 2003 do not support resource reservation.

| Category of QoS Mechanisms | Description |
|---|---|
| Admission control | Determine which applications and users are entitled to network resources. These mechanisms specify how, when, and by whom network resources on a network segment (subnet) can be used. |

| | |
|---|---|
| Traffic control | Regulate data flows by classifying, scheduling, and marking packets based on priority and by shaping traffic (smoothing bursts of traffic by limiting the rate of flow). Traffic control mechanisms segregate traffic into service classes and control delivery to the network. The service class assigned to a traffic flow determines the QoS treatment the traffic receives. |

The Intserv model integrates resource reservation and traffic control mechanisms to support special handling of individual traffic flows. The Diffserv model uses traffic control to support special handling of aggregated traffic flows.

**Uses and Benefits of QoS**

Network administrators can use QoS to guarantee throughput for mission-critical applications so that their transactions can be processed in an acceptable amount of time. Network administrators can also use QoS to manage User Data Protocol (UDP) traffic. Unlike Transmission Control Protocol (TCP), UDP is an inherently unreliable protocol that does not receive feedback from the network and, therefore, cannot detect network congestion. Network administrators can use QoS to manage the priority of applications that rely on UDP, such as multimedia applications, so that they have the required bandwidth even in times of network congestion, but do not overwhelm the network.

QoS provides the following benefits:

- Gives administrators control over network resources and allows them to manage the network from a business, rather than a technical, perspective.
- Ensures that time-sensitive and mission-critical applications have the resources they require, while allowing other applications access to the network.
- Improves user experience.
- Reduces costs by using existing resources efficiently, thereby delaying or reducing the need for expansion or upgrades.

**Network Technologies and Support for QoS**

QoS depends on support throughout the network. To achieve QoS from sender to receiver, all of the network elements through which a traffic flow passes — such as network interface cards, switches, routers, and bridges — must support QoS. If a network device along this path does not support QoS, the traffic flow receives the standard first-come, first-served treatment on that network segment.

Network technologies such as Frame Relay, asynchronous transfer mode (ATM), and the more traditional local area network (LAN) technologies (including Ethernet, Token Ring, and 802.11 wireless LAN) support QoS mechanisms. ATM, in particular, offers a high degree of support for QoS. Because ATM is a connection-oriented networking technology, it can establish a service contract that guarantees a specific quality of service and can allocate network resources. ATM enforces the service contract and allocates bandwidth at the hardware level. For more information about ATM QoS, see "How ATM Works."

Service level guarantees and resource allocations such as those provided by ATM are not supported by the more traditional IEEE 802 LAN technologies. These LAN technologies do not support QoS mechanisms at the hardware level, but at higher levels of the Open Systems Interconnection (OSI) reference model only. For these types of technologies, QoS is primarily based on the concept of priority, in which one data transmission receives delivery preference over other transmissions.

Traditional IEEE 802 LAN technologies are connectionless, so a host that sends a transmission cannot detect the state of the network or the state of the destination host prior to the transmission. Traffic on these network technologies, therefore, is subject to delay at points along the way, making bandwidth availability and delivery time difficult to predict. Although high-priority traffic typically arrives at its destination before lower-priority traffic, high-priority traffic cannot be guaranteed to arrive within a specified time.

Although QoS is more difficult to implement on connectionless networks, there is growing interest in developing QoS for IP-based networks.

**Windows Server 2003 QoS**

QoS in Windows XP and Windows Server 2003 is different from QoS in Windows 2000. Windows 2000 supports both the Intserv and Diffserv models and includes resource reservation, admission control, and traffic control mechanisms. Windows XP and Windows Server 2003 support only the Diffserv model and include traffic control mechanisms.

The current industry trend is toward the use of the Diffserv model because those solutions that use the Intserv model have proven to be difficult to scale in large networks as well as difficult to design and develop. The Diffserv model is designed to be scalable; it offers a simple method of providing differentiated services for IP traffic.

QoS support in Windows Server 2003 has changed for the following reasons:

- The Generic QoS (GQoS) application programming interface (API) is complex for developers to use.
- Many applications cannot determine their traffic requirements in advance.
- Resource allocation does not scale well in large networks.

**Windows 2000 QoS versus Windows Server 2003 QoS**
This section briefly describes the QoS mechanisms available in Windows 2000, Windows XP, and Windows Server 2003.

**Resource Reservation**
In Windows 2000, RSVP works with other components of the Intserv architecture to request network resources for traffic that has known characteristics. Windows XP and Windows Server 2003 do not support RSVP.

**Admission Control**
The Windows 2000 implementation of admission control is QoS Admission Control Service (ACS). ACS works with Subnet Bandwidth Manager (SBM) to manage congested subnets on IEEE 802-type networks. SBM is a signaling protocol used by ACS to indicate when a QoS request can be received. ACS then determines whether sufficient bandwidth is available and grants or rejects the request. Windows XP and Windows Server 2003 do not support ACS and SBM.

**Traffic Control**
In Windows 2000, as well as Windows XP and Windows Server 2003, the Generic Packet Classifier (GPC) classifies packets and the QoS Packet Scheduler schedules and marks packets and shapes traffic flows. Windows 2000, Windows XP, and Windows Server 2003 support prioritization at Layer 2 (Data Link layer) and at Layer 3 (Network layer) of the OSI reference model.

**QoS APIs**
The two QoS application programming interfaces (APIs) included with Windows 2000, General Quality of Service (GQoS) and Traffic Control (TC), are also supported in Windows XP and Windows Server 2003. GQoS is an API used by developers to create QoS-aware applications. In Windows XP and Windows 2003, however, function calls to the GQoS API do not generate RSVP messages. The TC API can be used by traffic management applications running with administrative privileges to apply QoS to applications that are not QoS-aware.

The following table summarizes the QoS mechanisms and APIs available in Windows 2000 compared to Windows XP and Windows Server 2003.

### QoS Mechanisms and APIs

| QoS Mechanism | Windows 2000 | Windows XP and Windows Server 2003 |
|---|---|---|
| RSVP | Yes | No |
| ACS | Yes | No |
| SBM | Yes | No |
| Packet Classifier | Yes | Yes |
| Packet Scheduler | Yes | Yes |
| GQoS | Yes | Yes, with limited functionality |
| TC | Yes | Yes |

**Related Information**
The following resources contain additional information that is relevant to this section.

- ATM Technical Reference

# How QoS Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How QoS Works**
**In this section**


- Traffic Control Overview
- Windows Server 2003 QoS Architecture
- QoS Protocols
- QoS Interfaces
- QoS Example
- QoS for Slow Links
- Related Information


The recent rapid growth in the number of bandwidth-intensive and time-sensitive applications, such as IP telephony, video conferencing, and other multimedia applications, often puts a strain on networks and keeps network administrators busy balancing capacity and performance. Windows Server 2003 Quality of Service (QoS) is a group of components that can differentiate traffic flows so that higher-priority traffic receives preferential treatment. Network administrators can use QoS to efficiently manage their resources and to ensure that high-priority applications receive preferential treatment without adversely affecting the performance of more typical network traffic. Various network technologies support QoS, but the focus of this document is on QoS for IP-based networks.

Implementing QoS involves combining a set of technologies defined by the Internet Engineering Task Force (IETF) and the Institute of Electrical and Electronic Engineers (IEEE). These technologies are designed to alleviate the problems caused by shared network resources and finite bandwidth. Although the concept of QoS encompasses a variety of standards and mechanisms, QoS for Windows Server 2003 IP-based networks is centered on traffic control, which includes mechanisms for prioritization and traffic shaping (the smoothing of traffic bursts).

QoS can be used in any network environment in which bandwidth, latency, jitter, and data loss must be controlled for mission-critical software, such as Enterprise Resource Planning (ERP) applications, or for latency-sensitive software, such as video conferencing, IP telephony, or other multimedia applications. QoS also can be used to improve the throughput of traffic that crosses a slow link, such as a dial-up connection.

All the network elements along the path that prioritized traffic takes must support QoS. Such network elements include the sending and receiving hosts, Layer 2 (Data Link layer) network devices (bridges and switches), and Layer 3 (Network layer) network devices (routers), including routers used for wide area network (WAN) links. If a network device along this path does not support QoS, the traffic flow receives the standard first-come, first-served treatment on that network segment.

This section describes how traffic control in Windows Server 2003 works. It covers the following:

- Overview of traffic prioritization and shaping.
- Description of Windows Server 2003 traffic control components.
- Description of the protocols used for Layer 2 and Layer 3 packet marking.
- Overview of the application programming interfaces (APIs) that invoke QoS mechanisms.
- Example of how prioritized traffic traverses a network.

- Description of QoS support for slow links.

**Traffic Control Overview**
Applications generate traffic at varying rates and typically require the network to carry traffic at the rate it is generated. Some applications are more tolerant than others of delays, variation in transmission rate, and data loss. Based on characteristics such as these, network traffic can fall into one of the following two types:

**Elastic traffic**
Elastic traffic adapts easily to changing network conditions. Delivery of elastic traffic is slow when bandwidth is limited and faster when bandwidth is abundant. Elastic traffic is usually generated by transaction-oriented applications, such as bulk data transfers.

**Real-time traffic**
Real-time traffic is limited in its ability to adapt to changing network conditions and delays can significantly reduce intelligibility and usefulness. Real-time traffic is generated primarily by real-time applications, such as video conferencing, that require dedicated bandwidth.

These two types of traffic have different delivery requirements. Traffic control provides the means for differentiating these types of traffic so that their delivery requirements can be met.

**Traffic Control Mechanisms**
Traffic control is a collection of mechanisms that segregate traffic into the appropriate service types and regulate its delivery to the network. Traffic control involves classifying, shaping, scheduling, and marking traffic.

**Classifying Traffic**
Duringclassification, packets are separated into distinct data flows and then directed to the appropriate queue on the forwarding interface. Queues are based on service type. The algorithm that services a queue determines the rate at which traffic is forwarded from the queue.

Traffic control in Windows Server 2003 supports the following service types:

**Best effort**
Best-effort is the standard service level in many IP-based networks. It is a connectionless model of delivery that provides no guarantees for reliability, delay, or other performance characteristics.

**Controlled load**
Controlled load data flows are treated similarly to best-effort data flows in unloaded (uncongested) conditions. This means that a very high percentage of transmitted packets will be delivered successfully to the receiving end node and a very high percentage of packets will not exceed the minimum delay in delivery. Controlled load service provides high quality delivery without guaranteeing minimum latency. For more information about controlled load service types, see the IETF-defined Request for Comment (RFC) 2211 in the IETF RFC Database.

**Guaranteed service**
Guaranteed service provides high-quality delivery with guaranteed minimum latency. The impact of guaranteed traffic on the network is heavy, so guaranteed service is typically used only for traffic that does not adapt easily to change. For more information about guaranteed service types, see RFC 2212 in the IETF RFC Database.

**Network control**
Network control, the highest service level, is designed for network management traffic.

**Qualitative service**
Qualitative service is designed for applications that require prioritized traffic handling but cannot quantify their QoS traffic requirements. These applications typically send traffic that is intermittent or burst-like in nature. For this service type, the network determines how the data flow is treated.

**Shaping and Scheduling Traffic**
During traffic shaping, the bursts in traffic are smoothed over time to create an even data flow. Traffic can be shaped only when it is transmitted, not when it is received.

During scheduling, a QoS component, the QoS Packet Scheduler, negotiates simultaneous requests for network access and determines which queue receives priority.

**Marking Traffic**
Markingmeans specifying a priority level in a packet header. Marking can be performed at Layer 2 of the Open Systems Interconnection (OSI) model (in a tag appended to the media access control [MAC] header) by using 802.1p priorities or at Layer 3 (in the IP header) by using Differentiated Services (Diffserv) priorities, as described later in this document. The priority

level maps to the service type a packet is to receive from network devices.

**Invoking Traffic Control**
In Windows 2000, applications use the Generic QoS (GQoS) API to apply QoS parameters to traffic generated by the application. Applications that use the GQoS API are said to be QoS-aware. In Windows Server 2003, applications can still call the GQoS API to prioritize traffic, but resource reservation is no longer supported.

A lower-level API, the Traffic Control (TC) API, provides prioritization and traffic shaping support at the host level rather than at the application level. Network administrators can use traffic management tools that call the TC API to apply QoS parameters on behalf of an application that is not QoS-aware.

**Windows Server 2003 QoS Architecture**
Traffic control in Windows Server 2003 consists of the TC API, the Generic Packet Classifier (GPC), and the QoS Packet Scheduler. The following diagram illustrates how the traffic control components fit together in Windows Server 2003.

### Traffic Control Architecture

In this diagram, the host computer is running a non-QoS-aware application that is generating traffic. The network administrator uses a traffic management program that calls the TC API to apply QoS parameters to the traffic. Traffic control on the sending host classifies, schedules, and marks the packets to accommodate the QoS service type specified by the traffic management program.

The Generic Packet Classifier classifies packets by service type; the QoS Packet Scheduler schedules the packets, marks them with the appropriate priority, and then sends the prioritized traffic to the network media.

**Traffic Control API**
Traffic control uses standardized QoS parameters to segment packets into sequences and to regulate the sequences. The Traffic Control (TC) API (Traffic.dll) applies the QoS parameters to the appropriate packets. Developers and network administrators use the TC API to specify the traffic that is to receive preferential treatment, the traffic that is to be treated alike, and the QoS parameters that define the preferential treatment. Traffic control uses the following specifications to define traffic:

- A flowspec, which is the list of QoS parameters such as service type and token rate (the permitted transmission rate of packets) that applies to a sequence of packets. A flowspec defines the preferential treatment.

- A flow, which is the sequence of packets that is subject to a flowspec. All of the packets in a flow receive the same treatment.

- A filterspec, which is the list of attributes, including source IP address, destination IP address, and port number, that classifies packets into a single flow.

Using the filterspec and flowspec, traffic control filters packets into the appropriate flow and applies the appropriate QoS parameters to the flow. QoS-aware devices, such as switches and routers, use the filterspec to determine how to handle packets.

**Generic Packet Classifier**
The Generic Packet Classifier (Msgpc.sys) classifies packets generated by an application so that the packets can subsequently be prioritized for sending across the network.

When an application or traffic management program invokes QoS, the Generic Packet Classifier on the sending host determines the service type (defined by the flowspec) to which each individual packet belongs and maps each packet to a queue (flow) based on the service type. After packets are assigned to a queue, the QoS Packet Scheduler can manage the queue in accordance with the QoS parameters specified in the flowspec.

### Note

- In Windows Server 2003, unlike Windows 2000, the GQoS API directly calls the Generic Packet Classifier. In Windows 2000, the GQoS API calls the QoS RSVP Service Provider, which is not supported in Windows Server 2003.

**QoS Packet Scheduler**
The QoS Packet Scheduler (Psched.sys) is a kernel-level component that marks and schedules packets that the Generic Packet Classifier has assigned into queues. The QoS Packet Scheduler enforces QoS parameters for a given flow to achieve traffic shaping.

The QoS Packet Scheduler retrieves packets from the queues and marks them with a priority and rate of flow. The rate of flow is used to pace the transmission of packets to the network, thus mitigating the inherent "send it all right now" nature of IP transmissions. The priority is used to determine the order in which packets are submitted to the network during periods of congestion, thus spreading bursts of traffic over time and creating a more uniform traffic flow.

Using the priority and rate of flow, the QoS Packet Scheduler determines the delivery schedule for each queue and negotiates competition between any queued packets that simultaneously request access to the network.

To ensure preferential treatment throughout the network, packets must be marked so that network devices along the path can detect the priority and handle the packets appropriately. Packets are marked with an 802.1p priority for prioritization by Layer 2 devices and with a Diffserv priority for prioritization by Layer 3 devices. The QoS Packet Scheduler provides 802.1p marking. TCP/IP provides Diffserv marking.

The QoS Packet Scheduler is not automatically installed with Windows Server 2003. You install this component on any host on which you want to perform traffic marking or shaping. Use Network Connections to display the properties for the connection where you want to install the QoS Packet Scheduler; install the QoS Packet Scheduler as a new service.

**Note**

- To use the QoS Packet Scheduler for 802.1p (Layer 2) marking, the network adapter must support 802.1p. On the **Advanced** tab for the properties of the network adapter (available through Network Connections or Device Manager), you must select the option to enable 802.1p support.

**QoS Protocols**

Windows Server 2003 supports both 802.1p (Layer 2) and Diffserv (Layer 3) priority marking. Network devices that support these protocols use the markings to provide end-to-end preferential treatment.

**802.1p**

The IEEE 802.1p signaling standard defines traffic prioritization at Layer 2 of the OSI model. Layer 2 network devices, such as switches, that adhere to this standard can group incoming packets into separate traffic classes.

The 802.1p standard is used to prioritize packets as they traverse a network segment (subnet). When a subnet becomes congested, causing a Layer 2 network device to drop packets, the packets marked for higher priority receive preferential treatment and are serviced before packets with lower priorities.

The 802.1p priority markings for a packet are appended to the MAC header. In Windows Server 2003, the QoS Packet Scheduler performs 802.1p marking.

On Ethernet networks, 802.1p priority markings are carried in Virtual Local Area Network (VLAN) tags. The IEEE 802.1q standard defines VLANs and VLAN tags. This standard specifies a 3-bit field for priority in the VLAN tag, but it does not define the values for the field. The 802.1p standard defines the values for the priority field. This standard defines eight priority classes (0 - 7). Network administrators can determine the actual mappings, but the standard makes general recommendations.

The VLAN tag is placed inside the Ethernet header, between the source address and either the Length field (for an IEEE 802.3 frame) or the EtherType field (for an Ethernet II frame) in the MAC header. The 802.1p marking determines the service level that a packet receives when it crosses an 802.1p-enabled network segment.

The following diagram illustrates the location of an 802.1p priority within a VLAN tag.

**Structure of a VLAN Tag**




Layer 2 priority values are defined by Group Policy. Windows Server 2003 provides a default priority value for each service type on a host. Sophisticated switches, however, might direct hosts or routers to use different mappings. The QoS Packet Scheduler in Windows Server 2003 uses the default priority values listed in the following table.

**Default Priority Values for Service Types**

| Service Type | Priority Marking |
|---|---|

| | |
|---|---|
| Non-conforming packets[1] | 0 |
| Best-effort | 0 |
| Controlled load | 4 |
| Guaranteed service | 5 |
| Network control | 7 |
| Qualitative | 0 |

[1] This service type is applied to packets that do not conform to the flowspec.

These default values can be modified by applying a Computer Configuration Group Policy setting for the host at Administrative Templates\Network\QoS Packet Scheduler\Layer-2 priority value. These Group Policy settings create corresponding registry subkeys and values at HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\Psched\UserPriorityMapping.

Implementing 802.1p marking on a host requires the following:

- The network adapter and device driver support 802.1p.
- The option to enable 802.1p support is selected.
- The service types are enabled with Group Policy.
- For a customized priority value, the default priority value has been changed with a Group Policy setting.

**Differentiated Services**

Diffserv (Differentiated Services) is a protocol that defines traffic prioritization at Layer 3 of the OSI model. Layer 3 network devices, such as routers, that support this protocol use Diffserv markings to identify the forwarding treatment, or per-hop behavior (PHB), that marked traffic is to receive. Diffserv markings for a packet are placed in the IP header. In Windows Server 2003, TCP/IP performs Diffserv marking.

RFC 2475 defines the architecture for Diffserv. RFC 2474 defines the bits in the Diffserv field. The Type of Service (TOS) field in Internet Protocol version 4 (IPv4) headers and the Traffic Class field in Internet Protocol version 6 (IPv6) headers are redefined for Diffserv values.

RFC 2474 redefines the IPv4 TOS octet as 6 bits for the Diffserv value, also known as Diffserv code point or DSCP, followed by 2 unused bits. The following figure illustrates how Diffserv redefines the TOS field.

**Diffserv Redefinition of TOS Field**




The first 3 bits of the TOS field were formerly used for IP precedence, which indicated the importance of a packet. DSCP values are backward-compatible with IP precedence values, which means that legacy routers that support only IP precedence can interpret DSCP values.

As with Layer 2 prioritization, DSCP values are defined by Group Policy. Windows Server 2003 provides a default DSCP value for each service type on a host, but some Layer 3 network devices can direct hosts to use different values. Windows Server 2003 uses the default DSCP values listed in the following table.

**Default DSCP Values for Service Types for Conforming Packets**

| Service Type | Priority Marking |
|---|---|
| Best-effort | 0 |
| Controlled load | 24 |

| | |
|---|---|
| Guaranteed | 40 |
| Network control | 48 |
| Qualitative | 0 |

Packets that do not conform to the flowspec have a default priority of 0.

These defaults can be modified by applying a Computer Configuration Group Policy setting for the host at Administrative Templates\Network\QoS Packet Scheduler\. The **DSCP value of conforming packets** subcategory supports settings for packets that conform to the flowspec. These Group Policy settings create corresponding registry subkeys and values at HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingConforming. The **DSCP value of non-conforming packets** subcategory supports settings for packets that do not conform to the flowspec. These Group Policy settings create corresponding registry subkeys and values at HKLM\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingNonConforming.

Diffserv differs from 802.1p prioritization in that Layer 3 devices can aggregate flows. Routers, therefore, need to distinguish a comparatively small number of aggregated flows instead of the thousands or millions of individual flows that compose them, and can provide better service to the packets that have higher priority.

**QoS Interfaces**

Windows Server 2003 provides two application programming interfaces (APIs) for assigning QoS parameters to traffic. Application developers use the GQoS API to apply QoS parameters at the application level. Network administrators use traffic management tools that call the TC API to apply QoS parameters at the host level.

**Generic QoS API**

In Windows 2000, the GQoS API was the foundation for QoS. It provided access to several QoS mechanisms, including resource reservation mechanisms and traffic control mechanisms. In Windows XP and Windows Server 2003, the GQoS API still provides access to traffic control mechanisms, but it no longer supports resource reservation. Existing calls to invoke resource reservations are still supported, but they are treated differently: a traffic flow is immediately created locally without a reservation request and the traffic from the application is marked and scheduled as though a reservation had been requested successfully.

GQoS is part of Windows Sockets 2.0 (Winsock2). Most applications use the GQoS API to invoke the QoS capabilities in Windows. For more information about the GQoS API, see “QoS API” on MSDN.

**Traffic Control API**

The TC API provides access to traffic control mechanisms that regulate network traffic on the local host. These components regulate traffic within the kernel of the local host and, through prioritization and scheduling, across the network.

QoS-aware applications (those that call the GQoS API) invoke the TC API indirectly. Network administrators can use traffic management programs to invoke this API directly on behalf of applications that are not QoS-aware. A lower-level API than the GQoS API, the TC API requires administrative privileges.

Unlike the GQoS API, the TC API allows traffic from a number of sources (on the same sending host) to be aggregated into a single traffic control flow. For example, all of the traffic to a specific destination IP address can be included in a single flow, regardless of the source port and destination port. The GQoS API, on the other hand, allows only traffic from a single conversation, which is defined by source and destination IP addresses, source and destination ports, and protocol (Transmission Control Protocol [TCP] or User Datagram Protocol [UDP]) to be included in a single flow.

For more information about the TC API, see “Application-Driven QoS Components” on MSDN.

**QoS Example**

The following diagram illustrates what happens when QoS is invoked for traffic generated by an application. In this diagram, an application running on Client A is generating traffic and sending it to Client B. In this example, the application is not QoS-aware, which means it does not use the GQoS API to specify QoS parameters. Instead, the network administrator uses a traffic management tool to define the QoS parameters for the host.

### Invoking QoS on a Host




In this example, the following occurs:

1. Client A on Network A is running an application that is generating traffic. The application is not QoS-aware, so the traffic is destined for best-effort service.

2. The network administrator has used a traffic management tool (Traffic Control Monitor, for example) to define flows (the filterspec) for each interface on Client A and the QoS parameters that apply to each flow (the flowspec) for preferential treatment.

3. As the application generates traffic, the GPC (Generic Packet Classifier) on Client A classifies the traffic into queues by service type.

4. The QoS Packet Scheduler on Client A manages the queues determined by the GPC by providing preferential treatment for higher-priority traffic. The QoS Packet Scheduler takes the following actions:

   - Reduces traffic bursts by smoothing transmission peaks over a period of time (that is, shapes traffic).

   - Schedules packets for transmission.

   - Marks packets for Layer 3 and, if 802.1p is enabled for the network adapter, Layer 2 prioritization, as specified by QoS Packet Scheduler Group Policy.

   - Sends the packets to the network media.

5. When the traffic reaches the Layer 3 router, the router uses the DSCP marking in the IP header to determine how to forward it.

6. When the traffic reaches the Layer 2 switch, the switch uses the 802.1p marking in the MAC header to give the traffic preferential treatment.

**QoS for Slow Links**
When slow links, such as dial-up modems, are used to connect to the Internet or another remote network, the network can suffer significant performance problems. Windows Server 2003 provides QoS support for these slow links.

**QoS for Internet Connection Sharing Clients**
Some home or small office networks use Internet Connection Sharing (ICS) to connect to the Internet or to another network. With ICS, multiple computers on a small network can access Internet resources by using the Internet connection of the ICS host computer. In these ICS networks, typically the computers are connected to a relatively fast internal network (for example, 100 megabits per second [Mbps] Ethernet). The connection to the Internet, however, can be a slow link, such as a dial-up connection.

When an ICS client computer (one that uses the shared Internet connection provided by the ICS host computer) is on a fast network and connects to a computer on a different fast network through a slow link connection, a speed mismatch occurs. This speed mismatch can result in a significant transmission delay on ICS networks. In this situation, the TCP receive window on the ICS client computer is set to a large value based on the high speed of the local small network. The TCP receive window

specifies the number of bytes remaining in a memory buffer that stores incoming data on a connection. The larger the window size, the more packets a sender can send at one time before waiting for acknowledgements and a newly advertised receive window size.

For a slow link, the ICS client sends the large value of the TCP receive window to the Internet server. When the Internet server responds to the ICS client's request for data, it starts sending data at a slow rate. If packets are not lost, the Internet server gradually increases the sending rate until it sends almost a full window of packets at a time. This is normal TCP connection behavior.

However, the disparity between the large amount of data that the ICS client reports it can receive and the smaller amount of data that can be accommodated by the slow link can cause data to back up in a potentially large queue at the slow link. Large queues can have a negative effect on the performance of other traffic across the slow link.

To avoid such problems caused by slow links, the ICS host computer automatically sets the TCP receive window for communications across the slow link to a size that accommodates the slow link, thus overriding the ICS client computer's default TCP receive window size. The smaller window size does not adversely affect traffic performance because it is the same size as the window that the receiving computer would use if it were connected directly to the slow link. The QoS Packet Scheduler on the ICS host computer makes this window adjustment on behalf of the ICS client computers.

**QoS for Modems**

Many home computers still connect to the Internet through slow, dial-up modem links with connection speeds less than 56 kilobits per second (Kbps). In spite of the speed limitation, users sometimes run several programs that access the Internet concurrently. For example, they might be browsing the Internet, downloading files, receiving e-mail, joining a chat session, or running streaming audio or video. Most programs that support these activities use TCP and open their own connections.

If the QoS Packet Scheduler is not installed on the computer, the first program that uses the slow link initially has exclusive use of the link. Exclusive use allows the program's connection to reach a steady state, which can result in the transfer of a full TCP receive window of data. When the next program starts to transfer data, its connection is subject to the TCP slow-start algorithm that limits how much unacknowledged data can be in transit. Because data is already in transit on the first connection, the second connection takes much longer to reach a steady state and the transfer rate is slower than for the first connection.

To resolve these issues, the QoS Packet Scheduler in Windows XP and Windows Server 2003 uses a Deficit Round Robin (DRR) queue-servicing scheme for slow links. The DRR scheme allocates queues and assigns data flows to each queue. The queues are serviced in a round-robin fashion, which improves responsiveness and performance for all network communications.

The DRR scheme was available in Windows 2000 and is used automatically in Windows XP and Windows Server 2003 when the operating system detects a slow link.

**Integrated Services over Slow Links**

Special mechanisms are provided to perform traffic shaping on slow links, such as 28.8-kilobytes per second (KBps) modem links. On such links, large packets can occupy the link long enough to delay small audio packets that must be sent on the same link. This can cause problems with audio quality. To avoid this problem, traffic control fragments large packets at the link layer, sending only one fragment at a time. Latency-sensitive audio packets can then be inserted in between the larger packet's fragments, thus reducing audio latency and improving audio quality.

Integrated Services over Slow Links (ISSLOW) is a queuing mechanism that is used to optimize slow (low-capacity) network interfaces by reducing latency. In particular, it is designed for interfaces that forward traffic to modem links, ISDN bearer (B) channels, and sub-T1 links.

A typical packet occupies a modem link for up to half a second. Other packets queued behind it can experience significant delays. Packets that are long enough in length to exceed the maximum tolerable delay for their QoS flow are fragmented before transmission through the link, so that high-priority packets can be inserted between the fragments of the larger packet, and meet the required QoS parameters for speedy transmission. The following figure illustrates a Point-to-Point Protocol (PPP) link using ISSLOW.

**PPP Link with ISSLOW**




For example, a PPP link is carrying best-effort traffic in addition to guaranteed service level flow "A." The PPP link capacity is 100 KBps and the average latency is 100 milliseconds. The best-effort traffic is consuming the majority of the bandwidth on the link, which is starving flow "A" of the resources it requires to provide guaranteed service levels. In this example, flow "A" cannot tolerate a delay of more than 145 milliseconds. As the best-effort traffic fills the queue, packets from flow "A" arrive. The first best-effort packet (10 kilobits) is fragmented into 2-kilobit packets. The 8-kilobit packet from flow "A" is fragmented into 2-

kilobit packets as well, and the flow "A" fragments are inserted between the best-effort fragments in order to meet the required latency guarantees of flow "A".

**Related Information**

The following resources contain additional information that is relevant to this section.

- ATM Technical Reference
- *Quality of Service* on MSDN

For more information about QoS and Traffic Control, see the following RFCs in the IETF RFC Database.

- RFC 2211, "Specification of the Controlled-Load Network Element Service."
- RFC 2212, "Specification of Guaranteed Quality of Service."
- RFC 2474, "Definition of the Differentiated Services Field."
- RFC 2475, "An Architecture for Differentiated Service."

# QoS Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**QoS Tools and Settings**

**In this section**

- QoS Tools
- QoS Group Policy Settings and Registry Entries

**QoS Tools**

The following tool is associated with Quality of Service (QoS).

**Tcmon.exe: Traffic Control Monitor**

**Category**

This tool is included in the Resource Kit Tools for Windows 2000, Windows XP, and Windows Server 2003.

**Version compatibility**

Runs on Windows 2000, Windows XP, and Windows Server 2003.

Tcmon.exe is a tool for viewing and setting traffic control information for local or remote network adapters. The Traffic Control Monitor is a graphical user interface (GUI) tool that interacts with the Traffic Control application programming interface (API).

Several components must be installed for the Traffic Control Monitor to run successfully. The Traffic Control Monitor is installed automatically when Windows Resource Kit tools are installed. The following additional components must be installed separately:

- Comctl32.ocx
- Msvbm60.dll
- Tccom.exe

**Note**

- Tccom.exe can be called by applications or scripts in Visual Basic Scripting Edition (VBScript) or other scripting languages to remotely configure traffic control on a host.

For information about how to install these components, see Windows Resource Kit Tools Help for Traffic Control Monitor.

**QoS Group Policy Settings and Registry Entries**

The following Group Policy settings and registry entries are associated with QoS. Each registry setting has a corresponding Group Policy setting. The registry settings appear at four locations in the registry; most of the settings are repeated in three of those locations. The registry structure matches the structure of the Computer Configuration Group Policy settings under Administrative Templates in the Group Policy tool. Each of the tables in this section describes the Group Policy settings and the registry entries in one of those four locations.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The registry entries are associated with Windows 2000, Windows XP, and Windows Server 2003. For more information about these registry entries, see the Group Policy Settings Reference for Windows Server 2003.

The following table describes three of the Group Policy settings that are associated with the QoS Packet Scheduler that do not

set alternative Differentiated Services Code Point (DSCP) values. These Group Policy settings are located in Computer Configuration Group Policy at Administrative Templates\Network\QoS Packet Scheduler; they control registry entries located under HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\Psched.

**Group Policy Settings Associated with QoS Packet Scheduler**

| Group Policy Setting | Description | Registry Entry |
|---|---|---|
| Limit outstanding packets | Specifies the maximum number of outstanding packets permitted on the system. When the number of outstanding packets reaches this limit, the QoS Packet Scheduler postpones all submissions to network adapters until the number falls below this limit. | MaxOutstandingSends |
| Limit reservable bandwidth | Specifies the maximum, total percentage of link bandwidth that all QoS-aware applications on this computer can reserve. If no QoS-aware application reserves any bandwidth, then this setting has no effect and all bandwidth is available for use by other applications. If a QoS-aware application reserves more bandwidth than it uses, then the unused, reserved bandwidth is available for use by other applications. The reservation does not ensure that the bandwidth will be available to the QoS-aware application because applications that are not QoS-aware might consume too much bandwidth. For example, the default value of the registry entry reserves 80 percent of the bandwidth for QoS-aware applications, but this does not prevent other applications from using more than the remaining 20 percent of the bandwidth. | NonBestEffortLimit |
| Set timer resolution | Determines the smallest unit of time, in microseconds, used by the QoS Packet Scheduler when scheduling packets for transmission. The QoS Packet Scheduler cannot schedule packets for transmission more frequently than permitted by the value of this entry. | TimerResolution |

The following table describes Group Policy settings that are associated with QoS and that set alternative DSCP values for packets that conform to the QoS traffic flow specification. You can use these Group Policy settings to specify alternative Layer-3 DSCP values, which are used by Layer 3 network devices to identify each of five service types. These Group Policy settings are located in Computer Configuration Group Policy at Administrative Templates\Network\QoS Packet Scheduler\DSCP value of conforming packets; they control registry entries located under HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingConforming.

**Group Policy Settings Associated with DSCP Values of Conforming Packets**

| Group Policy Setting | Description | Registry Entry |
|---|---|---|
| Best effort service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the best effort service type. | ServiceTypeBest Effort |
| Controlled load service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the controlled load service type. | ServiceTypeControlledLoad |
| Guaranteed service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the guaranteed service type. | ServiceTypeGuaranteed |
| Network control service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the network control service type. | ServiceTypeNetworkControl |
| Qualitative service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the qualitative service type. | ServiceTypeQualitative |

The following table describes Group Policy settings that are associated with QoS and that set alternative DSCP values for packets that do not conform to the QoS traffic flow specification. You can use these Group Policy settings to specify alternative Layer-3 DSCP values, which are used by Layer 3 network devices to identify each of five service types. These Group Policy settings are located in Computer Configuration Group Policy at Administrative Templates\Network\QoS Packet Scheduler\DSCP value of non-conforming packets; they control registry entries located under HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\Psched\DiffservByteMappingNonConforming.

**Group Policy Settings Associated with DSCP Values of Non-conforming Packets**

| Group Policy Setting | Description | Registry Entry |
|---|---|---|
| Best effort service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the best effort service type. | ServiceTypeBestEffort |
| Controlled load service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the controlled load service type. | ServiceTypeControlledLoad |
| Guaranteed service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the guaranteed service type. | ServiceTypeGuaranteed |
| Network control service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the network control service type. | ServiceTypeNetworkControl |
| Qualitative service type | Defines the DSCP value that the QoS Packet Scheduler inserts into packets to request the qualitative service type. | ServiceTypeQualitative |

The following table describes Group Policy settings that are associated with QoS and that set alternative Layer-2 priority values. You can use these Group Policy settings to specify alternative Layer-2 priority values, which Layer 2 network devices use to identify each of five service types. These Group Policy settings are located in Group Policy at Administrative Templates\Network\QoS Packet Scheduler\Layer-2 priority value; they control registry entries located under HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\Psched\UserPriorityMapping.

**Group Policy Settings Associated with Layer-2 Priority Values**

| Group Policy Setting | Description | Registry Entry |
|---|---|---|
| Non-conforming packets | Defines the priority value that the QoS Packet Scheduler inserts into the Layer-2 header of packets that do not conform to the flow specification. | ServiceTypeNonConforming |
| Best effort service type | Defines the priority value that the QoS Packet Scheduler inserts into the Layer-2 header of packets to request the best effort service type. | ServiceTypeBestEffort |
| Controlled load service type | Defines the priority value that the QoS Packet Scheduler inserts into the Layer-2 header of packets to request the controlled load service type. | ServiceTypeControlledLoad |
| Guaranteed service type | Defines the priority value that the QoS Packet Scheduler inserts into the Layer-2 header of packets to request the guaranteed service type. | ServiceTypeGuaranteed |
| Network control service type | Defines the priority value that the QoS Packet Scheduler inserts into the Layer-2 header of packets to request the network control service type. | ServiceTypeNetworkControl |
| Qualitative service type | Defines the priority value that the QoS Packet Scheduler inserts into the Layer-2 header of packets to request the qualitative service type. | ServiceTypeQualitative |

# Network Security Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Learn about network security technologies, including Windows Firewall.

**Network Security Technologies**
Windows Server 2003 provides two network security technologies: Windows Firewall, which provides host-based inspection and filtering of Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic, and Ipsec, which provides inspection, filtering, encryption, and authentication of IP traffic.

**Windows Firewall**
Windows Firewall provides host-firewall protection on computers running Windows Server 2003 with Service Pack 1 (SP1) and Windows XP with Service Pack 2 (SP2). As a host firewall, Windows Firewall runs on each of your servers and clients, providing protection from network attacks that pass through your perimeter network or originate inside your organization, such as Trojan horse attacks, worms, or any other type of malicious program spread through unsolicited incoming traffic.

# Windows Firewall Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To protect your network from the growing number of malicious attacks, you need to adopt a defense-in-depth strategy. A defense-in-depth strategy incorporates a variety of network security technologies into your security architecture and implements those technologies in a layered structure that extends from your perimeter network (outer layer) to every computer in your organization (inner layer). By providing host firewall protection on the innermost layer of your network, Windows Firewall, a new security component in Microsoft® Windows Server™ 2003 with Service Pack 1 (SP1), can be an effective part of your defense-in-depth security strategy.

**In this subject**


- What Is Windows Firewall?
- How Windows Firewall Works
- Windows Firewall Tools and Settings

# What Is Windows Firewall?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Common Windows Firewall Scenarios
- Technologies Related to Windows Firewall


Windows Firewall provides host-firewall protection on computers running Windows Server 2003 with Service Pack 1 (SP1) and Windows XP with Service Pack 2 (SP2). As a host firewall, Windows Firewall runs on each of your servers and clients, providing protection from network attacks that pass through your perimeter network or originate inside your organization, such as Trojan horse attacks, worms, or any other type of malicious program spread through unsolicited incoming traffic.

The following figure shows how Windows Firewall works in conjunction with perimeter network firewalls.




Windows Firewall inspects and filters all IP version 4 (IPv4) and IP version 6 (IPv6) network traffic. It is a stateful firewall, which means it tracks the state of each network connection and determines whether incoming traffic is allowed or blocked. Windows Firewall blocks incoming traffic unless it is in response to a request by the host (in which case, it is solicited traffic) or has been specifically allowed (in which case, it has been added to the Windows Firewall exceptions list). Aside from a few Internet Control Message Protocol (ICMP) messages, Windows Firewall allows all outgoing traffic.

Windows Firewall is designed to be a supplemental security solution. You cannot use Windows Firewall as a perimeter firewall. Windows Firewall should be part of a comprehensive security architecture that implements a variety of security technologies, such as border routers, perimeter firewalls, intrusion detection systems, virtual private networking (VPN), IEEE 802.1X authentication for wireless and wired connections, and Internet Protocol security (IPsec).

**Common Windows Firewall Scenarios**

With a few exceptions, Windows Firewall can be enabled on all configurations of Windows Server 2003 with SP1. Therefore, it is recommended that you enable Windows Firewall on every server in your organization, including bastion hosts and other servers in your perimeter network, mobile and remote clients that connect to your network and run Windows Server 2003, and all servers in your internal network. In addition, Windows Firewall can be enabled for any network architecture. Therefore, it is also recommended that you enable Windows Firewall regardless of the way in which you have designed and implemented your perimeter network or your internal network.

**Common Configuration Tasks**

When it is enabled in its default configuration, Windows Firewall blocks all unsolicited incoming network traffic on all network connections. Although blocking unsolicited incoming traffic reduces your attack surface and increases your level of security, it can cause programs and system services that are acting as servers, listeners, or peers to stop working properly. On clients, the number of programs and system services that do this is small because clients are typically configured as service consumers, not service providers. On servers, the number of programs and system services that do this can be large because, by their nature, servers are service providers; they must receive unsolicited incoming traffic.

Therefore, you usually need to configure some Windows Firewall settings on servers so that the appropriate unsolicited traffic is not blocked by Windows Firewall. The configuration changes that you need to make vary according to the server role and the programs and system services that are running on the server.

In addition, you need to configure Windows Firewall settings on every server that you intend to manage remotely. If you are managing the server by using the Administration Tools Pack on a client running Windows XP, then you need to configure Windows Firewall settings on the server to ensure that Windows Firewall allows the unsolicited traffic from the client computer to reach the programs and system services that are running on the server. You might also need to configure Windows Firewall settings on the client if the server initiates traffic to the client (for example, when an e-mail server notifies an e-mail client that new messages have arrived).

**Restrictions and Limitations**

You should disable Windows Firewall in the following scenarios.

**Servers Running Routing and Remote Access**

Do not run Windows Firewall if you are running Routing and Remote Access. Routing and Remote Access uses its own firewall and it cannot be used with Windows Firewall.

**Computers Running a Perimeter Firewall**

Do not run Windows Firewall on a server that is running a perimeter firewall, such as Microsoft Internet Security and Acceleration (ISA) Server 2004. The protection that Windows Firewall provides in this situation is redundant and unnecessary. In addition, Windows Firewall can cause perimeter firewalls such as ISA Server to function improperly.

**Computers Running Non-Microsoft Host Firewalls**

Do not use Windows Firewall on a computer that is running a non-Microsoft host firewall. Although it is possible to run two host firewalls on a single computer, it is not recommended that you do so. Host firewall implementations vary widely and there is no guarantee that a non-Microsoft host firewall and Windows Firewall will work well together.

**Computers Acting as Domain Controllers**

Do not permanently run Windows Firewall on a domain controller. You can use Windows Firewall on a temporary basis (for example, to protect a domain controller from attack while you install security updates, virus signatures, or security software).

In addition, you might want to disable Windows Firewall if a server requires you to open numerous ports or allow a large number of applications and system services to receive unsolicited traffic. In this case, it might make sense to disable Windows Firewall because a significant volume of network traffic will be allowed to pass through Windows Firewall anyway. By disabling Windows Firewall, you eliminate the operational overhead associated with Windows Firewall configuration and maintenance. However, you should closely evaluate the design of any client or server that requires you to open numerous ports. Servers that are configured for numerous roles or to provide numerous services can be a critical point of failure in your organization and might indicate poor infrastructure design.

**Technologies Related to Windows Firewall**

Windows Server 2003 provides several network security technologies that are related to and complement Windows Firewall. It is recommended that you use these technologies in conjunction with Windows Firewall to ensure that your security architecture is built on a solid defense-in-depth strategy.

**IPsec**

IPsec helps protect networks from active and passive attacks by securing IP packets through the use of packet filtering, cryptographic security services, and the enforcement of trusted communications. IPsec supports network-level data integrity, data confidentiality, data origin authentication, and replay protection. Because IPsec is integrated at the Internet layer (layer 3), it provides security for all IPv4-based traffic, and because IPsec is applied transparently to applications, there is no need to configure separate security for each application that uses TCP/IP. In addition, IPsec helps provide defense-in-depth against:

- Network-based attacks from untrusted computers, attacks that can result in the denial-of-service of applications, services, or the network.
- Data corruption.
- Data theft.
- User-credential theft.
- Administrative control of servers, other computers, and the network.

IPsec and Windows Firewall are not mutually exclusive security technologies. The use of IPsec in conjunction with Windows Firewall increases the security of the communication on your internal network and reduces the attack surface on each of your clients and servers. For example, on a secured Web server, you can use Windows Firewall to block all unsolicited incoming traffic except Web traffic and you can use IPsec to secure all Web traffic as it travels across the network.

**VPN**

VPN is a network access technology that is part of the Routing and Remote Access service in Windows Server 2003. You can use VPN to securely connect remote clients and remote offices to an organization's private network. VPN technology provides access to private networks at a lower cost than leased-line connections, which create a physical connection to a port on a remote access server on a private network.

A VPN connection uses authentication and data encryption to help preserve the privacy and, in some cases, integrity of data as it traverses a public network. VPN connections use either Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol/Internet Protocol security (L2TP/IPsec) over an intermediate network, such as the Internet. By using the Internet as a connection medium, VPN saves the cost of long-distance phone service and hardware associated with the use of dial-up or leased line connections. A VPN solution includes advanced security technologies, such as data encryption, user and computer authentication, authorization, and Network Access Quarantine Control.

VPN clients can use standard tools to access resources. For example, clients can use Windows Explorer to access shared files and folders and connect to printers. Connections are persistent; users do not need to reconnect to network resources during their VPN sessions. Because drive letters and universal naming convention (UNC) names are fully supported by VPN, most commercial and custom applications work without modification.

# How Windows Firewall Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


Windows Firewall Terminology

Windows Firewall Architecture

Windows Firewall Interfaces

Windows Firewall Protocols

Windows Firewall Connection Types

Windows Firewall Service Dependencies

Windows Firewall Physical Structures

Windows Firewall Processes and Interactions


Windows Firewall is a stateful host-firewall technology that inspects and filters all IP version 4 (IPv4) and IP version 6 (IPv6) network traffic. As a stateful firewall, Windows Firewall tracks the state of each network connection and determines whether the unsolicited incoming traffic should be allowed or dropped. Windows Firewall blocks incoming traffic unless the traffic is a response to a request by the host (in which case, it is solicited traffic) or it is specifically allowed (in which case, it has been added to the Windows Firewall exceptions list). Exceptions can be specified by program name, UDP or TCP port number, or system service name. Aside from a few Internet Control Message Protocol (ICMP) messages, Windows Firewall allows all outgoing traffic.

| **Note** |
| --- |
| By default, running dcpromo.exe to configure a computer as a Domain Controller turns off Windows Firewall. Use the Security Configuration Wizard to re-enable Windows Firewall with the port exceptions required to support the Domain Controller role. Alternatively, you can use Microsoft Knowledge Base article 555381, "How to Configure Windows Server 2003 Firewall for a Domain Controller" (http://go.microsoft.com/fwlink/?linkid=91224) to manually enable Windows Firewall with the port exceptions required to support the Domain Controller role. |

**Windows Firewall Terminology**

Before you review the components and processes that are associated with Windows Firewall, it is helpful to understand Windows Firewall terminology.

### Domain profile

A set of Windows Firewall settings that are applied when a computer is connected to a network that contains the domain controllers for the domain in which its computer account resides. The domain profile is one of two profiles used by Windows Firewall to apply settings to a computer.

### Exception

A port, program, or system service that is allowed to receive unsolicited traffic.

### Exceptions list

A rules store that specifies which ports, programs, or system services are allowed to receive unsolicited traffic. To allow unsolicited incoming traffic through Windows Firewall, you add ports, programs, and system services to the exceptions list.

### Notification

A **Windows Security Alert** dialog box that appears when a program that is not listed in the exceptions list attempts to listen for unsolicited traffic on a port. If you have administrative rights on a computer, the notification displays the option to add the program to the exceptions list, which will allow the program to receive unsolicited traffic. If you do not have administrative rights on a computer, the notification displays a warning that a program is attempting to listen on a port.

A **Windows Security Alert** dialog box does not appear when a system service attempts to listen for unsolicited incoming

traffic; it appears only when a program attempts to listen for unsolicited incoming traffic.

**Port**

A logical communication endpoint representing a program or a system service that listens for and receives IP packets. Ports are specified by a positive 16-bit decimal number and by the type of traffic (either UDP or TCP) that is expected to pass through the port. Most ports are predefined and well-known for specific services, such as DHCP and DNS. However, some ports are created dynamically by a server and assigned to the programs and system services that need to listen for incoming traffic from clients.

**Program**

A software program that is usually started by a user and runs under the user's account. Programs run only while the user is logged on to a computer and usually consist of one or more executable (.exe) files and one or more dynamic-link library (.dll) files.

| **Note** |
| --- |
| Some system services run within their own .exe file and are started by a user (for example, the Telnet service, which runs in Tlntsvr.exe). Typically, these system services run under a privileged account, such as Local Service. They do not run under the user's account. In Windows Firewall, any system service that runs within its own .exe file is considered a program. |

**Scope**

A Windows Firewall setting that you configure for an exception. The scope setting controls from which addresses unsolicited traffic is allowed to originate. By default, the scope of an exception is any address, which includes any computer on the Internet. You can change the scope of an exception to locally reachable addresses or a list of individual IPv4 addresses or IPv4 address ranges.

**Standard profile**

A set of Windows Firewall settings that are applied when a computer is not connected to a network that contains the domain controllers in which its computer account resides (for example, a public network, such as the Internet). The standard profile is one of two profiles used by Windows Firewall to apply settings to a computer.

**System service**

A software program that is usually not started by a user and runs under a privileged account, such as LocalSystem. System services run even when a user is not logged on to a computer and usually run as a separate process within Svchost.exe.

**Windows Firewall Architecture**

Windows Firewall relies on several networking components, including the Windows Firewall/Internet Connection Sharing service (also known as the SharedAccess service), the network address translation (NAT) driver (Ipnat.sys), the IPv4-based TCP/IP driver (Tcpip.sys) and the IPv6-based TCP/IP driver (Tcpipv6.sys), and the Windows Sockets (Winsock) driver (Winsock.dll). Windows Firewall also relies on several administrative tools that allow you to configure Windows Firewall settings. The following figure shows the relationship among these components.




The following describes each of the architectural components.

**Administrative tools**

There are several administrative tools that interact with the Windows Firewall/Internet Connection Sharing service and allow you to configure Windows Firewall settings. These tools include Windows Firewall in Control Panel (Firewall.cpl), the **netsh firewall** command, and Group Policy with the Group Policy Object Editor.

### NAT driver

The NAT driver (Ipnat.sys) provides a data store for Windows Firewall. Windows Firewall uses the data store, known as the NAT Mapping Table, to store connection state information for dynamically-created exceptions. The connection state information typically consists of a 5-tuple entry, which includes the protocol, the source and destination port numbers, and the source and destination IP addresses. Without a data store, Windows Firewall would not be a stateful firewall.

| **Note** |
| --- |
| Although Windows Firewall uses the NAT driver, Windows Firewall does not provide any network address translation. |

### TCP/IP driver

The TCP/IP driver (Tcpip.sys) controls the flow of information between a network adapter and a program or system service. As incoming traffic flows through the TCP/IP driver, the traffic is inspected by the NAT driver. The NAT driver processes the traffic based on the entries in the Windows Firewall exceptions list. If the traffic matches an exception, the NAT driver determines that the traffic is allowed; the packets continue through the TCP/IP driver. If the traffic does not match an exception, the NAT driver determines that the traffic is unsolicited; the packets are dropped and do not continue through the TCP/IP stack. Neither the NAT driver nor the TCP/IP driver sends a notification to the sender when packets are dropped (this is sometimes referred to as a silent discard).

### Winsock driver

The Winsock driver is responsible for assigning and binding ports to a program or system service. Programs and system services use this driver when they need to listen for incoming traffic.

### Windows Firewall/Internet Connection Sharing service

The Windows Firewall/Internet Connection Sharing service runs in Svchost.exe. The Windows Firewall/Internet Connection Sharing service is responsible for processing the Windows Firewall exceptions that you configure for network traffic. The Windows Firewall/Internet Connection Sharing service then relays the exceptions to the NAT driver, which inspects the traffic that is flowing through the TCP/IP driver.

### Winsock driver

The Windows Firewall/Internet Connection Sharing service provides all of the processing (filtering) that determines whether unsolicited incoming traffic is allowed or dropped.

#### Windows Firewall Interfaces

There are two interfaces related to Windows Firewall: the firewall-hook interface and the Windows Firewall application programming interfaces (APIs). The following figure shows these interfaces and how they relate to Windows Firewall components.




The firewall-hook interface is used by the TCP/IP driver to communicate with Windows Firewall, which instructs the TCP/IP driver about how to process incoming and outgoing packets. There should be only one firewall-hook interface on Windows XP and Windows Server 2003 operating systems because it is not recommended that software developers implement a firewall-hook driver on these operating systems. To provide additional firewall functionality on Windows XP and Windows Server 2003, software developers should create a Network Driver Interface Specification (NDIS) intermediate miniport driver to manage packets sent and received across a firewall. For information about firewall-hook drivers, see MSDN (http://go.microsoft.com/fwlink/?LinkId=45023).

The Windows Firewall APIs provide a public interface for programmatically configuring Windows Firewall settings. Windows

Firewall APIs and Internet Connection Sharing (ICS) APIs are implemented in Hnetcfg.dll. Using the Windows Firewall API, you can:

- Enable and disable Windows Firewall.
- Add and remove programs and system services from the exceptions list.
- Add and remove TCP and UDP ports from the exceptions list.
- Enable and disable preconfigured system service exceptions.
- Configure ICMP settings.
- Configure the log file.

The following table describes the Windows Firewall APIs.

| Interface | Description |
|---|---|
| **INetFwOpenPort** | Provides access to the properties of a port opened in Windows Firewall. |
| **INetFwOpenPorts** | Provides access to a collection of ports opened in Windows Firewall. |
| **INetFwService** | Provides access to the properties of a service that may be authorized to listen through Windows Firewall. |
| **INetFwRemoteAdminSettings** | Provides access to the settings that control remote administration. |
| **INetFwServices** | Provides access to a collection of services that may be authorized to listen through Windows Firewall. |
| **INetFwAuthorizedApplication** | Provides access to the properties of a program that has been authorized to open ports in Windows Firewall. |
| **INetFwAuthorizedApplications** | Provides access to a collection of programs authorized to open ports in Windows Firewall. |
| **INetFwIcmpSettings** | Provides access to the settings that control ICMP packets. |
| **INetFwProfile** | Provides access to a Windows Firewall settings profile. |
| **INetFwPolicy** | Provides access to a Windows Firewall policy. |
| **INetFwMgr** | Provides access to the Windows Firewall settings for a computer. |

For more information about Windows Firewall APIs, see MSDN.

**Windows Firewall Protocols**

Windows Firewall provides the stateful filtering of TCP/IP traffic (IPv4 and IPv6) that uses the TCP transport protocol. It also provides the "pseudo-stateful" filtering of TCP/IP traffic that uses the UDP transport protocol. ICMP traffic is not statefully filtered; rather, ICMP traffic is allowed or blocked based on Windows Firewall settings (for example, you can explicitly allow or deny incoming echo requests or outgoing destination unreachable messages by configuring Windows Firewall settings). Because Windows Firewall is tied directly to the TCP/IP architecture of Windows, it does not provide any filtering of non-TCP/IP protocols, such as IPX/SPX or AppleTalk.

With the exception of some File Transfer Protocol (FTP) traffic, Windows Firewall does not use Application layer information to statefully filter traffic. FTP is a special case because of the way in which an FTP server establishes the data channel for an FTP file transfer. During a typical FTP user session, an FTP client initiates a control channel with an FTP server. When the FTP client transfers a file from the FTP server, the FTP server tries to establish a data channel with the FTP client by initiating communication on a TCP port different from the one used for the control channel. This can cause most firewalls running on the FTP client computer to drop the data channel packets coming from the server because they appear to be unsolicited. To overcome this problem, Windows Firewall uses the Application Layer Gateway Service to provide dynamic port mapping for the FTP data channel, thereby facilitating the stateful filtering of FTP traffic.

| Note |
| --- |
| Passive (PASV) FTP is not treated as a special case and is not handled by the Application Layer Gateway Service. PASV FTP traffic is statefully filtered by Windows Firewall in the same way other TCP/IP traffic is statefully filtered. |

Although Windows Firewall does not use Application layer information to statefully filter traffic, Windows Firewall does use Application layer information to identify and manage some traffic.

**IPv6 Tunneled Traffic**

IPv6 tunneled traffic is an IPv6 packet that has been encapsulated with an IPv4 header so that hosts that support both IPv4 and IPv6 can communicate using IPv6 even when separated by an IPv4 infrastructure, such as the Internet or a private intranet. The tunneled IPv6 packet bypasses the IPv4 processing of Windows Firewall. However, after the IPv4 header of the packet is removed, the decapsulated IPv6 packet is processed by Windows Firewall just like all other IPv6 traffic.

**AH/ESP**

IPsec uses two transforms, the Authentication Header (AH) and the Encapsulating Security Payload (ESP) header and trailer, to encapsulate and secure IP packets or payloads. These transforms wrap the IP payload with header or header and trailer information and provide security properties, such as data origin authentication, data integrity, data confidentiality, and replay protection.

Traffic protected with AH or ESP is processed by the IPsec components of Windows before it is processed by Windows Firewall. The IPsec components validate the security information in the AH or ESP header and trailer, decrypt the payload if needed, and remove the AH or ESP header and trailer. The resulting IP packet is handed back to the TCP/IP driver, which is then handed to Windows Firewall for processing. However, you can change this behavior so that IPsec-secured traffic bypasses Windows Firewall by configuring the **Windows Firewall: Allow authenticated bypass** Group Policy setting. When you enable this setting, you specify groups of IPsec-secured computers whose traffic bypasses all Windows Firewall processing and overrides any existing policies or settings that would otherwise block the traffic.

**GRE**

The Generic Routing Encapsulation (GRE) protocol is used to encapsulate the data channel of Point-to-Point Tunneling Protocol (PPTP) traffic. Any PPTP traffic that uses a GRE header to encapsulate Point-to-Point Protocol (PPP) frames will pass directly through Windows Firewall. Non-PPTP traffic that uses GRE is filtered by Windows Firewall.

**IGMP**

Internet Group Management Protocol (IGMP) is an IP multicast maintenance protocol. IP multicast provides the one-to-many delivery of traffic without the overhead and performance impact that is associated with broadcast delivery and multiple unicast delivery. IP multicast is used primarily for audio and video teleconferencing and gaming. Hosts and routers send IGMP messages so that IP multicast-enabled routers can determine the set of multicast host groups that have members on a subnet. All IGMP traffic passes directly through Windows Firewall.

**PGM**

The Pragmatic General Multicast (PGM) protocol is a scalable multicast protocol that enables receivers to detect loss, request retransmission of lost data, or notify a program of unrecoverable loss. PGM is often referred to as reliable multicast; it is a receiver-reliable protocol, which means the receiver is responsible for ensuring all data is received, absolving the sender of reception responsibility. Any traffic that uses the PGM protocol passes directly through Windows Firewall.

**Windows Firewall Connection Types**

Windows Firewall can be bound to any network connection that is listed in Network Connections in Control Panel. This includes the following connection types.

**Local area connections**

This includes wired and wireless network connections that use Ethernet, Token Ring, cable modem, 1394, Digital Subscriber Line (DSL) modem, Fiber Distributed Data Interface (FDDI), IP over Asynchronous Transfer Mode (ATM), and WAN technologies such as T1 and Frame Relay.

**Dial-up connections**

This includes network connections that use X.25, modem, and Integrated Services Digital Network (ISDN) technologies.

**VPN connections**

This includes virtual private networking (VPN) connections that use PPTP or Layer Two Tunneling Protocol (L2TP).

### Incoming connections

This includes incoming connections that use dial-up (modem, ISDN, X.25) and VPN (PPTP, L2TP). However, Windows Firewall settings for incoming connections can only be configured globally. You cannot configure Windows Firewall settings on a per-connection basis for incoming connections.

| **Note** |
| --- |
| Windows Firewall cannot be enabled on direct connections, such as serial cable connections, infrared link, and DirectParallel cable connections. |

#### Windows Firewall Service Dependencies

Windows Firewall runs in the Windows Firewall/Internet Connection Sharing service and is dependent on several subsidiary services. If any of the subsidiary services are not running or cannot be started, Windows Firewall might not start or run properly. Windows Firewall depends on the following services.

### Application Layer Gateway Service

The Application Layer Gateway Service (sometimes known as the ALG service) is required if you enable Windows Firewall on a computer that is an FTP client or FTP server that does not use PASV FTP. The Application Layer Gateway Service listens for outgoing FTP traffic from an FTP client. It then extracts the port from which the FTP client is expecting to receive data and creates an appropriate dynamic port mapping for the FTP data channel.

The Application Layer Gateway Service and Windows Firewall interact as follows:

- If the Application Layer Gateway Service is disabled and you try to enable Windows Firewall, Windows Firewall will start, but FTP traffic that does not use PASV FTP might fail.
- If you stop the Application Layer Gateway Service while Windows Firewall is running, Windows Firewall will continue to run, but FTP traffic that does not use PASV FTP might fail.
- If the Application Layer Gateway Service is stopped and its startup type is set to Manual, then the Application Layer Gateway Service will attempt to start when you enable Windows Firewall.

### DCOM Server Process Launcher

The DCOM Server Process Launcher service (sometimes known as the Dcomlaunch service) is new in Windows Server 2003 with Service Pack 1 (SP1); it is responsible for starting all DCOM-related services. Because Windows Firewall relies on DCOM, the DCOM Server Process Launcher service must be running when you start Windows Firewall.

The DCOM Server Process Launcher service and Windows Firewall interact as follows:

- If the DCOM Server Process Launcher is disabled and you try to enable Windows Firewall, Windows Firewall will not start because the Windows Firewall/Internet Connection Sharing service cannot start.
- You cannot stop the DCOM Server Process Launcher service, so it is not possible to configure the DCOM Server Process Launcher service if it is started and Windows Firewall is enabled.
- If the startup type for the DCOM Server Process Launcher service is set to Manual, and the DCOM Server Process Launcher is not started, you will not be able to enable Windows Firewall. If you attempt to enable Windows Firewall, a Windows Firewall dialog box will be displayed to indicate that another process that requires the network address translation driver (Ipnat.sys) is running.

### Network Connections

The Network Connections service (sometimes known as the Netman service) manages all network connections that are created and configured in Network Connections in Control Panel. This service is also responsible for displaying network status in the notification area on the desktop (the area on the taskbar to the right of the taskbar buttons). Windows Firewall uses this service to identify and manage network connections.

The Network Connections service and Windows Firewall interact as follows:

- You cannot enable Windows Firewall if the Network Connections service is disabled.
- If you stop the Network Connections service while Windows Firewall is enabled, Windows Firewall will stop running and become disabled.
- If the Network Connections service is stopped and its startup type is set to Manual, the Network Connections service will attempt to start when you enable Windows Firewall. If the Network Connections service cannot start, Windows Firewall cannot be enabled.

**Network Location Awareness**

The Network Location Awareness service (sometimes known as the Nla service) collects and stores network configuration information, such as changes to the names and locations of IP addresses and domain names. The Network Location Awareness service component supports the Network Location Service, which provides the infrastructure that informs programs and the operating system of the network connections on a computer. Windows Firewall uses the Network Location Awareness service to determine which Windows Firewall profile to use.

The Network Location Awareness service and Windows Firewall interact as follows:

- If the Network Location Awareness service is disabled and you try to enable Windows Firewall, Windows Firewall will start but cannot determine which profile to use. In this case, Windows Firewall uses the settings for the standard profile.
- If you stop the Network Location Awareness service while Windows Firewall is running, Windows Firewall continues to run and use the settings for the profile that it is currently using.
- If the Network Location Awareness service is stopped and its startup type is set to Manual, then the service will attempt to start when you enable Windows Firewall.

**Remote Procedure Call**

The Remote Procedure Call service (sometimes known as the RpcSs service) is a secure interprocess communication (IPC) mechanism that enables data exchange and invocation of functionality that resides in a different process. That different process can be on the same computer, on the local area network (LAN), or across the Internet. The Remote Procedure Call service serves as the RPC Endpoint Mapper (EPM) and Service Control Manager (SCM).

The Remote Procedure Call service and Windows Firewall interact as follows:

- You cannot enable Windows Firewall if the Remote Procedure Call service is disabled.
- If you stop the Remote Procedure Call service while Windows Firewall is enabled, Windows Firewall will stop running and become disabled.
- If the Remote Procedure Call service is stopped and its startup type is set to Manual, then the service will attempt to start when you enable Windows Firewall. If the Remote Procedure Call service cannot start, Windows Firewall cannot be enabled.

**Windows Management Instrumentation**

Windows Management Instrumentation (WMI) is an infrastructure for building management programs and instrumentation released as part of the current generation of Microsoft operating systems. The Windows Management Instrumentation service (sometimes known as the Winmgmt service) provides a common interface and object model to access management information about operating system, devices, programs, and system services. Windows Firewall uses the Windows Management Instrumentation service to store per-connection settings and legacy (ICF) settings.

- You cannot enable Windows Firewall if the Windows Management Instrumentation service is disabled.
- If you stop the service while Windows Firewall is enabled, Windows Firewall will stop running and become disabled.

- If the service is stopped and its startup type is set to Manual, then the service will attempt to start when you enable Windows Firewall. If the Windows Management Instrumentation service cannot start, Windows Firewall cannot be enabled.

**Windows Firewall Physical Structures**
Windows Firewall relies on two physical structures: the NAT Mapping Table, which contains information about the traffic passing through a network connection, and a rules store, which contains the configuration settings (rules) that are used to filter traffic. The state table allows Windows Firewall to track the activity of a given network connection; it is the mechanism that makes Windows Firewall a stateful firewall. The rules store provides a persistent data store for Windows Firewall configuration settings; it is the mechanism that Windows Firewall uses to apply consistent filtering rules to a network connection.

**NAT Mapping Table**
Windows Firewall uses the NAT Mapping Table to store the state of network connections. The NAT Mapping Table is managed by the NAT driver (Ipnat.sys) and is used by both Windows Firewall and Routing and Remote Access to store connection information.

**TCP Connections**
TCP connections are stored in a standard 5-tuple entry format in the NAT Mapping Table. The 5-tuple entry consists of the following data:

- Protocol. This is either TCP or UDP.
- Source address. This is the IP address of the host that initiated the communication session.
- Source port. This is the TCP or UDP port number that the host used to initiate the communication session.
- Destination address. This is the IP address of the computer with which the host initiated traffic.
- Destination port. This is the TCP or UDP port number to which the host sent traffic.

A 5-tuple entry is created in the NAT Mapping Table for every incoming communication session that uses TCP. A separate mapping table is maintained for each network connection (for example, if you have two network adapters in a computer, there is a mapping table for each network adapter). The timeout for established but inactive TCP connections is 24 hours. In other words, the state table entry for an inactive TCP connection remains in the NAT Mapping Table for 24 hours unless the connection is terminated and closed using the standard TCP close sequence.

**UDP Connections**
UDP connections that use ports 1 through 1024 are also stored in the standard 5-tuple entry format in the NAT Mapping Table. The 5-tuple entry is the same as the entry used for TCP connections. The 5-tuple entry is used to track UDP connections on ports 1 through 1024 because these ports are usually designated for use by specific operating system services. Using a 5-tuple entry to store the state of all connections to these ports decreases the likelihood of an attack on these critical system services.

UDP connections that use ports higher than 1024 are stored in a 3-tuple entry format in the NAT Mapping Table. The 3-tuple format, sometimes known as loose source mapping, consists of the following data:

- Protocol. This is always UDP.
- Source address. This is the IP address of the host that initiated the communication session.
- Source port. This is the UDP port number that the host used to initiate the communication session.

A 3-tuple entry is created in the NAT Mapping Table for every incoming communication session that uses a UDP port greater than 1024. A separate mapping table is maintained for each network connection (for example, if you have two network adapters in a computer, there is a mapping table for each network adapter).

Because UDP is a connectionless protocol and has no sequence numbers or flags, it has no mechanism for terminating and closing a connection. Therefore, the timeout for UDP connections is much shorter than that for TCP connections. The timeout for established but inactive UDP connections is 60 seconds. In other words, if an established UDP connection is inactive for 60 seconds, the state table entry for that connection is removed from the NAT Mapping Table. This behavior applies to UDP connections on all ports.

There are some exceptions to the 60-second timeout for UDP connections. When a computer sends a multicast or broadcast message, Windows Firewall waits for up to 3 seconds for a unicast response. In essence, the state table entry for multicast or broadcast connections is maintained for up to 3 seconds; if the multicast or broadcast connection is inactive for 3 seconds (that is, there is no response), then the state table entry is deleted from the NAT Mapping Table. Dynamic Host Configuration Protocol (DHCP) multicast messages are a special case and are exempt from the 3-second timeout.

**ICMP Traffic**

The Windows Firewall state table does not store any information about ICMP traffic. Incoming ICMP traffic is accepted or rejected according to the configured Windows Firewall ICMP settings.

**Rules Store**

A rules store contains the list of rules used by a firewall to determine whether unsolicited incoming traffic is allowed or blocked. A typical rules store is created around the notion of "implicit deny," which means that all unsolicited incoming traffic is blocked (denied) unless there is an explicit rule allowing the traffic through the firewall. Furthermore, a typical rules store is processed sequentially from top to bottom: that is, the firewall compares the characteristics of unsolicited incoming traffic against each rule, one at a time, until a rule is found that allows the traffic (in which case, the traffic passes through the firewall) or the end of the rules list is reached (in which case, the traffic is blocked). Creating and maintaining this type of rules store can be difficult because the order of the rules is important and it is relatively easy to create a rule that inadvertently allows all traffic through the firewall.

Windows Firewall uses the notion of implicit deny, but it does not use a strictly sequential or ordered rules store. When you turn on Windows Firewall in its default configuration, all unsolicited incoming TCP and UDP traffic is blocked. In other words, you must create explicit rules to allow unsolicited incoming traffic to pass through Windows Firewall. However, you do not need to create the rules in any particular order because the rules are not processed sequentially.

| **Note** |
| --- |
| With the exception of some ICMP messages, outgoing traffic is not blocked by Windows Firewall. |

In most cases, the rules used by Windows Firewall to filter unsolicited incoming traffic are a union of the Windows Firewall settings you configure using Windows Firewall in Control Panel, the **netsh firewall** command, local Group Policy settings, and domain-based Group Policy settings. You cannot configure ordered rules or rules that specify a precedence for specific protocols, ports, programs, or IP addresses. The only time the resultant rules are not determined by a union of all settings is when Group Policy settings conflict with settings that you configured locally through Windows Firewall in Control Panel or the **netsh firewall** command. In this case, the resultant rules are still determined by a union, but the domain-based Group Policy settings take precedence over any locally-configured settings (including local Group Policy settings) and the local Group Policy settings take precedence over settings configured through Windows Firewall in Control Panel and the **netsh firewall** command.

| **Note** |
| --- |
| Exceptions that are created programmatically using the Windows Firewall API behave in the same manner as exceptions that are created through Windows Firewall in Control Panel or the **netsh firewall** command. If there is a conflict between an exception that is created programmatically and an exception that is created with Group Policy, the Group Policy settings will take precedence. |

Windows Firewall stores its rules in two locations in the registry:

```
HKLM\Software\Policies\Microsoft\WindowsFirewall
```

This subkey contains the settings you configure through Group Policy (both local Group Policy and domain-based Group Policy).

```
HKLM\System\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy
```

This subkey contains the settings you configure through Windows Firewall in Control Panel, the **netsh firewall** command, and scripts.

**Windows Firewall Processes and Interactions**

This section describes how Windows Firewall processes network traffic and the way in which Windows Firewall interacts with other operating system components.

**Windows Firewall and Post-Setup Security Updates**

By default, Windows Firewall and the Windows Firewall/Internet Connection Sharing service are turned off in Windows Server 2003 with SP1. However, there are instances when Windows Firewall and the Windows Firewall/Internet Connection

Sharing server are turned on automatically on a temporary basis. This occurs when the Post-Setup Security Updates feature runs. The Post-Setup Security Updates feature is designed to reduce the risk of a new server being attacked from the time it is first connected to the network until the most recent security updates are installed from Windows Update. The new feature appears in the graphical user interface as a single dialog box the first time a member of the Administrators group logs on to a new server. The dialog box provides links for installing current security updates and configuring automatic updates. In addition, the Post-Setup Security Updates feature temporarily turns on Windows Firewall, which blocks all inbound connections other than those that are specifically opened during setup or by policy. When you close the Post-Setup Security Updates dialog box, Windows Firewall is turned off.

Post-Setup Security Updates is available only when you perform a clean installation of Windows Server 2003 with SP1, including slip-stream installations. You cannot run Post-Setup Security Updates from the **Start** menu or from the command line (it starts automatically), and it is not available when:

- Windows Firewall is enabled or disabled during setup by using an unattended installation answer file or a Netfw.inf file.
- Windows Firewall is enabled or disabled through Group Policy before the Post-Setup Security Updates dialog box appears.

In addition, Post-Setup Security Updates is not available when you upgrade an existing Windows Server 2003 operating system by adding a service pack or if you upgrade a Windows 2000 Server operating system to Windows Server 2003 with SP1.

**What Happens When Post-Setup Security Updates Is Closed?**
If Windows Update or any other configuration change causes your computer to restart before you click the **Finish** button on the **Post-Setup Security Updates** dialog box, it reopens the next time an administrator logs on to the server.

If you close the **Post-Setup Security Updates** dialog box using ALT+F4 or Task Manager, no change is made to the configuration of Windows Firewall. The tests the server uses to determine whether the **Post-Setup Security Updates** dialog box should be displayed run again the next time a user logs on.

If you close Post-Setup Security Updates by clicking the **Finish** button on the **Post-Setup Security Updates** dialog box, the following occurs:

- If you did not change any Windows Firewall settings while the **Post-Setup Security Updates** dialog box was displayed, a confirmation dialog box appears explaining that inbound connections will now be opened and giving you the opportunity to confirm that you are done with any post-setup security updates. When the action is confirmed, Post-Setup Security Updates attempts to disable Windows Firewall and stop and disable the Windows Firewall/Internet Connection Sharing service.
- If Windows Firewall is disabled successfully, a registry value is set to suppress Post-Setup Security Updates in the future. It is possible that Windows Firewall is disabled successfully, but the attempt to stop the Windows Firewall/Internet Connection Sharing service fails.
- If Windows Firewall settings cannot be changed, a dialog box appears explaining that no changes will be made to inbound connection settings. Post-Setup Security Updates is not suppressed and the tests to determine whether Post-Setup Security Updates should be displayed will be run again the next time a user logs on.
- If Windows Firewall was explicitly enabled or disabled while the Post-Setup Security Updates was displayed, a dialog box appears indicating that no changes will be made to inbound connection settings. These changes could have been made by Group Policy settings or by opening Windows Firewall in Control Panel and clicking **OK** to confirm the firewall settings. A registry value is set to suppress Post-Setup Security Updates in the future.
- If the Windows Firewall/Internet Connection Sharing service was stopped or disabled while the **Post-Setup Security Updates** dialog box was displayed, a dialog box appears indicating that no changes will be made to inbound connection settings. A registry value is set to suppress Post-Setup Security Updates in the future.
- If Internet Connection Sharing was enabled while the **Post-Setup Security Updates** dialog box was displayed, a confirmation dialog box appears indicating that inbound connections will now be opened and giving you the opportunity to confirm that you are done with any post-setup security updates. When the action is confirmed, Post-Setup Security Updates attempts to disable Windows Firewall and stop the Windows Firewall/Internet Connection Sharing service.

- If the state of the firewall cannot be determined, a dialog box appears indicating that no changes will be made to inbound connection settings. Post-Setup Security Updates is not suppressed and the tests to determine whether the **Post-Setup Security Updates** dialog box should be displayed will be run again the next time a user logs on.

| **Note** |
| --- |
| The text that appears in the **Post-Setup Security Updates** dialog box is not refreshed if Windows Firewall status changes after the dialog box is first displayed. If the status of Windows Firewall changes after it appears and before the **Finish** button is clicked, the text may state that all inbound connections are blocked when, in fact, they are not. When you click **Finish**, Post-Setup Security Updates checks the status of Windows Firewall again before displaying a dialog box indicating any changes to be made on closure. |

**When Is the Post-Setup Security Updates Dialog Box Displayed?**

Because the Post-Setup Security Updates feature runs automatically and cannot be started manually, you can use the following information to determine whether Post-Setup Security Updates will run on your server:

| Test Question | Answer=Yes | Answer=No |
| --- | --- | --- |
| Is the logged-on user an administrator? | Continue to the next test. | Skip the remaining tests and do not display the **Post-Setup Security Updates** dialog box. These tests run again the next time a user logs on. |
| Is this is a new installation of a version of Windows Server 2003 that includes Service Pack 1 (not an upgrade)? | Continue to the next test. | Skip the remaining tests and do not display the **Post-Setup Security Updates** dialog box. The registry value is set to suppress Post-Setup Security Updates in the future. |
| Has Post-Setup Security Updates been suppressed in the registry? | Skip the remaining tests and do not display the **Post-Setup Security Updates** dialog box. | Continue to the next test. |
| Is the Windows Firewall/Internet Connection Sharing service running? | Continue to the next test. | Repeat this test for two minutes. If the service has still not started, do not display the **Post-Setup Security Updates** dialog box. These tests are run again the next time a user logs on. |
| Has Windows Firewall been explicitly enabled or disabled for the current Windows Firewall profile?<br><br>(Windows Firewall may have been enabled or disabled using an unattended-setup script at the time of installation or through the application of Group Policy settings or by opening Windows Firewall in Control Panel and clicking **OK** to confirm the settings.) | Skip the remaining tests and do not display the **Post-Setup Security Updates** dialog box. A registry value is set to suppress Post-Setup Security Updates in the future. | If Windows Firewall is enabled and the user did not enable it, display the **Post-Setup Security Updates** dialog box.<br><br>If the status of Windows Firewall cannot be determined, do not display the **Post-Setup Security Updates** dialog box. These tests are run again the next time a user logs on. |

**Startup Protection**

Because there is a brief period of time during startup when a computer is connected to a network but Windows Firewall is not yet running, Windows Firewall uses a special startup policy that allows it to perform stateful packet filtering while your computer is starting. The startup policy, which you cannot configure, allows your computer to perform basic networking startup tasks using DHCP and the DNS service to communicate with a domain controller to obtain Group Policy updates. The startup policy is in effect only until the Windows Firewall/Internet Connection Sharing service starts, at which time Windows Firewall uses its user-configured settings and discards the special startup policy settings.

**Windows Firewall Filtering**

When an incoming packet reaches your computer, Windows Firewall inspects the packet and determines whether it meets

criteria specified in the Windows Firewall exceptions list (the combination of the NAT Mapping Table and the Windows Firewall rules store). If the packet matches an entry in the list, Windows Firewall passes the packet to the TCP/IP protocol for further processing. If the packet does not match an entry in the list, Windows Firewall silently discards the packet and creates an entry in the Windows Firewall log file (if the log file is enabled on the **Log Settings** dialog box in Windows Firewall). The entries in the exceptions list can consist of program names, system service names, TCP ports, and UDP ports. There is no way to create an entry in the exceptions list based on the IP Protocol field in the IP header.

There are only two conditions under which traffic can pass through Windows Firewall:

- When the incoming traffic corresponds to a recent outgoing request, in which case the response traffic is then considered incoming solicited traffic.

An example of this occurs when a DNS Name Query Request message is sent to a DNS server. Windows Firewall adds an entry to the NAT Mapping Table so that the corresponding DNS Name Query Response message sent by the DNS server can be passed to the TCP/IP protocol for further processing.

- When you configure Windows Firewall for exceptions, in which case you allow a computer using Windows Firewall to accept unsolicited incoming traffic when acting as a server, a listener, or a peer.

An example of this occurs when your computer is acting as a Web server. In this case, you must configure Windows Firewall to allow Web traffic so that the computer can respond to requests from Web clients. To do this, you configure an exception, which becomes a rule in the Windows Firewall rules store. You can configure exceptions based on specific programs, in which case the ports required by the program are opened and closed dynamically according to the program's needs, or on TCP or UDP ports, in which case the ports are opened even if the programs or system services are not using them.

**Windows Firewall Profile Determination**

Windows Firewall filters traffic according to the settings that are stored in one of two profiles: the domain profile or the standard profile. Windows Firewall uses a network determination algorithm to determine which profile to use.

The network determination algorithm uses connection-specific DNS suffix information to determine whether a computer is connected to a managed network containing the domain in which the computer is a member. To do this, the algorithm compares the following DNS suffix information:

- The last-received Group Policy update DNS name. This is stored in the HKEY_Local_Machine\Software\Microsoft\Windows\CurrentVersion\Group Policy\History\NetworkName registry entry and represents the connection-specific DNS suffix of the connection over which the last Group Policy update was received.

- The connection-specific DNS suffixes of the currently connected connections (those that are assigned an IP address) that are not Point-to-Point Protocol (PPP) or Serial Line Internet Protocol (SLIP)-based, such as dial-up or VPN network connections.

Using this DNS suffix information, the algorithm performs the following analysis:

- If the computer is not a member of a domain, then the computer is always attached to another network.

- If the last-received Group Policy update DNS name matches any of the connection-specific DNS suffixes of the currently connected connections on the computer that are not PPP or SLIP-based, then the computer is attached to a managed network.

- If the last-received Group Policy update DNS name does not match any of the connection-specific DNS suffixes of the currently connected connections on the computer that are not PPP or SLIP-based, then the computer is attached to another network.

Windows Server 2003 uses this network determination process during startup and when it is informed by the Network Location Awareness service that network settings on the computer have changed. To determine which profile to use, Windows Firewall applies this network determination process as follows:

- If the connection-specific DNS suffix of a currently connected connection on the computer that is not PPP or SLIP-based

(such as an Ethernet or 802.11 wireless network adapter) matches the value of the HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Group Policy\History\NetworkName registry entry, Windows Firewall uses the domain profile.

- If the connection-specific DNS suffix of a currently connected connection on the computer that is not PPP or SLIP-based does not match the value of the HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Group Policy\History\NetworkName registry entry, Windows Firewall uses the standard profile.

Although this network determination algorithm works well in most situations, configuration variations can cause Windows Server 2003 to determine that it is always attached to the managed network or create temporary configuration problems when roaming between different portions of the same managed network that assign different DNS suffixes with DHCP.

**Always Attached to the Managed Network**
The connection-specific DNS suffix of the connection over which the last set of Group Policy updates were received is determined from its TCP/IP configuration, which is typically configured using DHCP and the DNS Domain Name DHCP option (DHCP option number 15). You can also manually configure connection-specific DNS suffixes.

If the connection-specific DNS suffixes for the computer's connections are manually configured, then the manually-specified value overrides the value from the DNS Domain Name DHCP option. If the manually-configured DNS suffix matches the last-received Group Policy update DNS name, then Windows will always determine that the computer is attached to the managed network.

This can create a problem for Windows Firewall Group Policy settings if an administrator configures Group Policy settings so that Windows Firewall is disabled when the computer is attached to the managed network (this practice is highly discouraged unless another host-based firewall is being used) and Windows Firewall is enabled when the computer is attached to another network. If Windows Server 2003 always determines that the computer is attached to the managed network, Windows Firewall will be disabled when the computer is attached to the Internet, thereby making the computer accessible to network attacks.

**Roaming Between Different Portions of the Same Managed Network**
Temporary network determination problems can also occur when computers roam between portions of a managed network that assign different DNS domain names using DHCP.

For example, a laptop computer with only a wireless LAN connection is a member of the noam.corp.example.com Active Directory domain. DHCP servers in all of the offices of the organization in North America assign the DNS name noam.corp.example.com. However, in Europe, the Active Directory domain is europe.corp.example.com and the DNS servers assign the DNS name europe.corp.example.com. When the user connects a laptop computer to the wireless network in an office in Europe, the computer's last-received Group Policy update DNS name no longer matches the DNS suffix of the computer's wireless connection and Windows determines that the laptop computer is on another network, even though the laptop computer is connected to a portion of the managed network.

However, this condition lasts only as long as it takes for Group Policy settings to be updated. When the computer updates Group Policy settings, the last-received Group Policy update DNS name is reset to the new connection-specific DNS suffix and the computer is then determined to be on the managed network.

Continuing our example, the laptop computer updates Group Policy settings from a domain controller in the noam.corp.example.com domain and the last-received Group Policy update DNS name is set to europe.corp.example.com, which now matches the connection-specific DNS suffix of the wireless connection. When the network determination algorithm is run again, it determines that the computer is connected to a managed network.

The workaround for this issue is to have all of the DHCP servers in an organization assign a common DNS domain name, rather than a region-specific DNS domain name. For our example, the workaround is to have all the DNS servers assign the name "example.com." Therefore, the last-received Group Policy update DNS name will always match the connection-specific DNS suffix of a LAN connection that is attached to the organization network.

The problem of temporarily determining that the computer is on another network when it is actually on the managed network can also occur with the following configurations:

- The DNS Domain Name DHCP option is changed, assigning a domain suffix that does not match the last-received Group Policy update DNS name.

- The DNS Domain Name DHCP option is removed, resulting in a blank connection-specific DNS suffix for the computer's LAN connection.

- The connection-specific DNS suffix for the computer's LAN connection is manually configured and does not match the last-received Group Policy update DNS name.

**Scope Resolution**

Scope settings define from which addresses incoming traffic is allowed to originate for a specific exception. By default, exceptions have unrestricted scope, which means that exceptions apply to all unsolicited incoming IPv4 and IPv6 traffic no matter from where the traffic originates. You can change this default behavior and restrict the scope of an exception by using the **My network (subnet) only** scope option or by using the **Custom list** scope option. The **My network (subnet) only** scope option changes the default behavior so that an exception applies only to unsolicited incoming traffic that originates from an IPv4 or IPv6 address that is directly reachable from your computer. The **Custom list** scope option changes the default behavior so that an exception applies only to specific IPv4 addresses or a range of IPv4 addresses. You cannot use the **Custom list** scope option to specify an IPv6 address or a range of IPv6 addresses.

**Resolving Directly Reachable Computers**

When you use the **My network (subnet) only** scope option to restrict the scope of an exception, Windows Firewall does not perform any IP configuration testing or subnet mask interpretation to determine whether the traffic is coming from a computer that is a neighbor on a locally attached subnet. Instead, Windows Firewall queries the IPv4 and IPv6 routing tables to determine whether the source IPv4 or IPv6 address of the incoming packet can be reached directly. You can see all destinations that are considered directly reachable by typing the **route print** command at a command prompt. For the IPv4 routing table, all IPv4 addresses that match the routes in which the IPv4 address of the **Gateway** column equals the IPv4 address of the **Interface** column are considered directly reachable. For the IPv6 routing table, all IPv6 addresses that match routes in which the **Gateway** column is set to **On-link** are considered directly reachable. Therefore, the set of directly reachable addresses depends on your networking configuration, as specified by the IPv4 and IPv6 configuration of LAN-based connections (such as Ethernet and 802.11 wireless), dial-up connections, and broadband Internet connections. In some Internet configurations, all destinations are considered directly reachable.

For example, for a computer that is directly connected to a private home network only, the set of directly reachable unicast addresses is confined to those that match the IPv4 network ID of the private subnet. If the network connection is configured with an IPv4 address of 192.168.0.99 with a subnet mask of 255.255.0.0, only traffic from IPv4 addresses in the range 192.168.0.0 to 192.168.255.255 is allowed.

In another example, for a computer that is directly connected to both a private home network and the Internet through a cable modem, the set of directly reachable unicast addresses are those that match either the network ID of the private subnet or the subnet of the cable modem provider. For example, if the private network connection is configured with an IPv4 address of 192.168.0.1 and a subnet mask of 255.255.0.0 and the cable modem connection is configured with an IPv4 address of 131.107.17.211 and a subnet mask of 255.255.255.0, traffic received by either network connection is allowed from IPv4 addresses in the ranges from 192.168.0.0 to 192.168.255.255 and from 131.107.17.0 to 131.107.17.255.

**Resolving Specific IP Addresses and Address Ranges**

When you use the **Custom list** scope option to restrict the scope of a exception, you must provide a list of specific IPv4 addresses or a list of IPv4 address ranges (which usually correspond to subnets).

The IPv4 addresses are specified in dotted decimal notation; each address is separated by a comma. IPv4 address ranges are specified using a dotted decimal subnet mask or a prefix length; each address range is separated by a comma. A dotted decimal subnet mask is typically specified as an IPv4 network ID (such as 10.47.81.0/255.255.255.0) or by an IPv4 address within a range (such as 10.47.81.231/255.255.255.0). A network prefix length is typically specified as an IPv4 network ID (such as 10.47.81.0/24) or an IPv4 address within the range (such as 10.47.81.231/24).

Windows Firewall will ignore any IPv6 addresses or IPv6 address ranges that you add to the **Custom list** scope option.

# Windows Firewall Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Windows Firewall Tools
- Windows Firewall Group Policy Settings
- Windows Firewall Automated Installation Tools and Settings
- Windows Firewall Scripting Reference
- Netsh Commands for Windows Firewall
- Windows Firewall Log File and Security Log Settings


**Windows Firewall Tools**
The following tools are associated with Windows Firewall.

**Firewall.cpl: Windows Firewall Control Panel Tool**
**Category**
Control Panel tool included in the operating system.

**Version compatibility**
Available on Windows XP with Service Pack 2 (SP2) and Windows Server 2003 with Service Pack 1 (SP1).

You can use Windows Firewall in Control Panel to configure Windows Firewall settings on a local computer. This includes enabling and disabling Windows Firewall, specifying the type of information that is written to the Windows Firewall log file, and configuring the way Windows Firewall handles network traffic to and from specific programs, services, and ports.

**Gpedit.msc: Group Policy Object Editor**
**Category**
Microsoft Management Console (MMC) snap-in included in the operating system.

**Version compatibility**
Available on Windows 2000, Windows XP, and Windows Server 2003.

You can use the Group Policy Object Editor to configure Windows Firewall Group Policy settings. Windows Firewall Group Policy settings are configured on a per-computer basis and are stored in Administrative Templates (.adm files). You can configure Windows Firewall Group Policy settings on a local Group Policy object (GPO) or an Active Directory GPO. Group Policy settings that are configured on a local GPO are applied to the local computer only. Group Policy settings that are configured on an Active Directory GPO can be linked to an Active Directory container, such as a domain, organizational unit, or site.

**Gpmc.msc: Group Policy Management Console**
**Category**
MMC snap-in included in the operating system.

**Version compatibility**
Available on Windows XP with SP2 and Windows Server 2003 with SP1. Also available for download for Windows XP with SP1 and Windows Server 2003.

You can use the Group Policy Management Console (GPMC) to manage GPOs that contain Group Policy settings for Windows Firewall. This includes managing GPOs for multiple domains and sites within one or more forests. This snap-in uses a simplified user interface (UI) with drag-and-drop support. It can be used to perform many other administrative tasks, such as backup, restore, import, copy, and reporting of GPOs.

**Netsh.exe: Netsh**
**Category**

Command-line tool included in the operating system.

**Version compatibility**
Runs on Windows XP and the Windows Server 2003 family.

You can use the Netsh Firewall context to monitor and configure Windows Firewall settings. Several new commands have been added to the Netsh Firewall context for Windows Server 2003 with SP1. The new commands correspond to new Windows Firewall configuration settings.

You can also use Netsh to apply Windows Firewall settings that are stored in a Netfw.inf file. For more information, see "*Windows Firewall Netfw.inf Settings*" later in this section.

**Windows Firewall Group Policy Settings**
If your organization uses Group Policy, use the Windows Firewall Group Policy settings to configure Windows Firewall. The Group Policy Object Editor (Gpedit.msc) provides access to the Windows Firewall settings. The settings are stored within the Group Policy Object Editor at Computer Configuration/Administrative Templates/Network/Network Connections/Windows Firewall.

The following table lists the Windows Firewall Group Policy settings:

| Setting | Description |
|---|---|
| **Windows Firewall: Allow authenticated IPSec bypass** | Allows unsolicited incoming messages from specified systems that authenticate using the IPsec transport. |
| **Windows Firewall: Protect all network connections** | Turns on Windows Firewall, which replaces Internet Connection Firewall on all computers that are running Windows Server 2003 with SP1. |
| **Windows Firewall: Do not allow exceptions** | Specifies that Windows Firewall blocks all unsolicited incoming messages. This Group Policy setting overrides all other Windows Firewall settings that allow such messages. |
| **Windows Firewall: Define program exceptions** | Allows you to view and change the program exceptions list defined by Group Policy. Windows Firewall uses two program exception lists: one is defined by Group Policy settings and the other is defined by the Windows Firewall component in Control Panel. |
| **Windows Firewall: Allow local program exceptions** | Allows administrators to use the Windows Firewall component in Control Panel to define a local program exceptions list. |
| **Windows Firewall: Allow remote administration exception** | Allows remote administration of this computer using administrative tools such as MMC and Windows Management Instrumentation (WMI). To do this, Windows Firewall opens TCP ports 135 and 445. Services typically use these ports to communicate using remote procedure calls (RPC) and Distributed Component Object Model (DCOM). This policy setting also allows SVCHOST.EXE and LSASS.EXE to receive unsolicited incoming messages and allows hosted services to open additional dynamically-assigned ports, typically in the range of 1024 to 1034. |
| **Windows Firewall: Allow file and printer sharing exception** | Allows file and printer sharing. To do this, Windows Firewall opens UDP ports 137 and 138, and TCP ports 139 and 445. |

| | |
|---|---|
| **Windows Firewall: Allow ICMP exceptions** | Defines the set of Internet Control Message Protocol (ICMP) message types allowed by Windows Firewall. Utilities can use ICMP messages to determine the status of other computers. For example, Ping uses the echo request message. If you do not enable the "Allow inbound echo request" message type, Windows Firewall blocks echo request messages sent by Ping running on other computers, but it does not block outbound echo request messages sent by Ping running on this computer. |
| **Windows Firewall: Allow Remote Desktop exception** | Allows this computer to receive Remote Desktop requests. To do this, Windows Firewall opens TCP port 3389. |
| **Windows Firewall: Allow UPnP framework exception** | Allows this computer to receive unsolicited UPnP messages sent by network devices, such as routers with built-in firewalls. To do this, Windows Firewall opens TCP port 2869 and UDP port 1900. |
| **Windows Firewall: Prohibit notifications** | Prevents Windows Firewall from displaying notifications to the user when a program requests that Windows Firewall add the program to the program exceptions list. |
| **Windows Firewall: Allow logging** | Allows Windows Firewall to record information about the unsolicited incoming messages that it receives. |
| **Windows Firewall: Prohibit unicast response to multicast or broadcast requests** | Prevents this computer from receiving unicast responses to its outgoing multicast or broadcast messages. |
| **Windows Firewall: Define port exceptions** | Allows you to view and change the port exceptions list defined by Group Policy. Windows Firewall uses two port exception lists: one is defined by Group Policy settings and the other is defined by the Windows Firewall component in Control Panel. |
| **Windows Firewall: Allow local port exceptions** | Allows administrators to use the Windows Firewall component in Control Panel to define a local port exceptions list. |

**Windows Firewall Automated Installation Tools and Settings**

There are two automated installation technologies that you can use to deploy or configure Windows Firewall. One technology uses an information file known as Netfw.inf to automate the configuration of Windows Firewall settings. The other technology uses an unattended installation answer file to automate the configuration of Windows Firewall settings.

**Windows Firewall Netfw.inf Settings**

You can use the Windows Firewall information file (Netfw.inf) to deploy Windows Firewall in a corporate environment; however, this automated installation technology is typically used by OEMs to configure new computers in a manufacturing environment.

The Netfw.inf file configures Windows Firewall by modifying registry settings. To configure a Netfw.inf file, you must use standard .inf file structure and syntax.

The Netfw.inf file is found in the following location: %windir%\Inf\Netfw.inf.

**Note**

When you restore Windows Firewall default settings, the settings that are specified in Netfw.inf are reapplied to your computer and all existing settings are deleted.

The following is an example of a Netfw.inf file:

```
[version]
Signature = "$Windows NT$"
DriverVer = MM/DD/YYYY,1.2.3456.7890

[DefaultInstall]
AddReg=ICF.AddReg.DomainProfile
AddReg=ICF.AddReg.StandardProfile

[ICF.AddReg.DomainProfile]

[ICF.AddReg.StandardProfile]
```

**Note**

The first two sections contain version and configuration information and do not need to be modified. The sections that require modification are [ICF.AddReg.DomainProfile] and [ICF.AddReg.StandardProfile].

**[ICF.AddReg.DomainProfile]**
This section is used for defining changes to Windows Firewall's default configuration when a computer is connected to a network that contains its domain.

**[ICF.AddReg.StandardProfile]**
This section is used for defining changes to Windows Firewall's default configuration when a computer is not connected to a network that contains its domain. If a computer is not a member of a domain, then Windows Firewall will always enforce the configuration stored in the standard profile.

**Configuration Options**
The following settings can be defined in the Windows Firewall Netfw.inf file:

- Operational mode
- Disable notifications
- Block unicast responses to multicast and broadcast packets
- Enable remote administration
- Allow ICMP messages
- Open ports
- Allow programs

**Note**

All of the settings defined in the Windows Firewall Netfw.inf file will be applied to all of the computer's network interfaces; you cannot use the Netfw.inf file to open ports or allow ICMP messages for individual interfaces. In addition, you cannot use the Netfw.inf file to define logging settings.

**Operational Mode**
Windows Firewall can be placed in one of the three operational modes by adding the following entries to the Netfw.inf file:

| Mode | Section | Entry |
|---|---|---|

| | | |
|---|---|---|
| **On** | [ICF.AddReg.DomainProfile] | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile", "DoNotAllowExceptions",0x00010001,0 |
| | | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile", "EnableFirewall",0x00010001,1 |
| | [ICF.AddReg.StandardProfile] | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile", "DoNotAllowExceptions",0x00010001,0 |
| | | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile", "EnableFirewall", 0x00010001,0x00000001 |
| **Off** | [ICF.AddReg.DomainProfile] | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile", "DoNotAllowExceptions",0x00010001,1 |
| | | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile", "EnableFirewall",0x00010001,1 |
| | [ICF.AddReg.StandardProfile] | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile", "DoNotAllowExceptions",0x00010001,1 |
| | | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile", "EnableFirewall",0x00010001,1 |
| **Off with no excep-tions** | [ICF.AddReg.DomainProfile] | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile", "DoNotAllowExceptions", 0x00010001,0 |
| | | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile", "EnableFirewall",0x00010001,0 |
| | [ICF.AddReg.StandardProfile] | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile", "DoNotAllowExceptions",0x00010001,0 |
| | | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile", "EnableFirewall",0x00010001,0 |

**Disable Notifications**

Add the following entries to the Windows Firewall Netfw.inf file to disable notifications:

| Section | Entry |
|---|---|
| [ICF.AddReg.DomainProfile] | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile","DisableNotifications", 0x00010001,1 |
| [ICF.AddReg.StandardProfile] | HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile","DisableNotifications", 0x00010001,1 |

**Block Unicast Responses to Multicast and Broadcast Packets**

Add the following entries to the Windows Firewall Netfw.inf file to disable unicast responses to multicast and broadcast packets:

| Section | Entry |
|---|---|

| | |
|---|---|
| [ICF.AddReg.DomainProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile", "DisableUnicastResponsesToMulticastBroadcast", 0x00010001,1` |
| [ICF.AddReg.StandardProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile", "DisableUnicastResponsesToMulticastBroadcast", 0x00010001,1` |

**Enable Remote Administration**

Add the following entries to the Windows Firewall Netfw.inf file to enable remote administration:

| Section | Entry |
|---|---|
| [ICF.AddReg.DomainProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile\RemoteAdminSettings", "Enabled",0x00010001,1` |
| [ICF.AddReg.StandardProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile\RemoteAdminSettings", "Enabled",0x00010001,1` |

Add the following entries to the Windows Firewall Netfw.inf file to define the default scope for remote administration:

| Section | Entry |
|---|---|
| [ICF.AddReg.DomainProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile\RemoteAdminSettings", "RemoteAddresses",0x00000000,scope` |
| [ICF.AddReg.StandardProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile\RemoteAdminSettings", "RemoteAddresses",0x00000000,scope` |

**Scope Definition**

You can define the set of IP addresses from which unsolicited incoming traffic is allowed when you enable remote administration, open a port, or allow a program. This set of IP addresses from which unsolicited incoming traffic is allowed is referred to as the scope of the exception. There are three options when defining the scope for a Windows Firewall exception:

- **All IP Addresses** – This is the default scope for a Windows Firewall exception and it allows unsolicited incoming traffic that matches the exception from any computer. In the Windows Firewall Netfw.inf file, making an entry's scope element an asterisk ("*") will result in a scope of all IP addresses for the entry.

- **Local Subnet Only** – This scope allows unsolicited incoming traffic that matches the exception from any computer on the same subnet as the network connection on which the traffic was received through Windows Firewall, while dropping unsolicited incoming traffic from all other computers. When a computer's subnet changes, the set of allowed IP addresses dynamically changes to match the new subnet. In the Windows Firewall Netfw.inf file, making an entry's scope element **LocalSubnet** will result in a local subnet only scope for the entry.

- **Custom** – This scope is a list of IPv4 addresses and address ranges that typically correspond to subnets. Unsolicited incoming traffic that matches the exception and originates from a computer with an IPv4 address in the defined list is allowed through Windows Firewall. Unsolicited incoming traffic from computers with IPv4 addresses that are not in the list is dropped. A custom scope can include the local subnet (using the **LocalSubnet** string), IPv4 addresses, and IPv4 address ranges, but cannot include IPv6 addresses or IPv6 address ranges. For IPv4 address ranges, you can specify the range using a dotted decimal subnet mask or a prefix length (w.x.y.z/n). When you use a dotted decimal subnet mask, you can specify the range as an IPv4 network ID (such as 10.47.81.0/255.255.255.0) or by using an IPv4 address within the range (such as 10.47.81.231/255.255.255.0). When you use a network prefix length, you can specify the range as an IPv4 network ID (such as 10.47.81.0/24) or by using an IPv4 address within the range (such as 10.47.81.231/24).

| **Note** |
|---|

| |
|---|
| Do not use spaces between the entries in the list of sources or the entire list will be ignored and Windows Firewall will use the default scope of any source IPv4 address. |

**Allow ICMP Messages**

Add the following entries to the Windows Firewall Netfw.inf file to allow certain ICMP messages:

| Section | Entry |
|---|---|
| [ICF.AddReg.DomainProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile\IcmpSettings", "ICMP Message Type",0x00010001,1` |
| [ICF.AddReg.StandardProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile\IcmpSettings", "ICMP Message Type",0x00010001,1` |

The following table lists the permitted values for certain ICMP message types:

| ICMP Message Type | Number | Description |
|---|---|---|
| **AllowOutboundPacketTooBig** | 2 | When an IPv6 packet is too large to be forwarded, data will be dropped and a computer will reply to the sender with a “Packet Too Big” message. |
| **AllowOutboundDestinationUnreachable** | 3 | Data that fails to reach this computer due to an error will be discarded and reported with a “Destination Unreachable” message that explains the failure. |
| **AllowOutboundSourceQuench** | 4 | When the rate of a transmission exceeds a computer’s ability to process incoming data, data will be dropped and the sender will be asked to transmit more slowly. |
| **AllowRedirect** | 5 | Data sent from a computer will be rerouted. |
| **AllowInboundEchoRequest** | 8 | Messages sent to a computer will be repeated back to the sender. This is commonly used for troubleshooting (for example, to ping a computer). |
| **AllowInboundRouterRequest** | 10 | A computer will respond to router discovery messages. |
| **AllowOutboundTimeExceeded** | 11 | When a computer discards a packet because its hop count was exceeded or it ran out of time to assemble fragments of a packet, it will reply to the sender with a “Time Exceeded” message. |
| **AllowOutboundParameterProblem** | 12 | When a computer discards data it has received due to a problematic header, it will reply to the sender with a “Parameter Problem” error message. |
| **AllowInboundTimestampRequest** | 13 | Data sent to a computer can be responded to with a confirmation message indicating the time that the data was received. |
| **AllowInboundMaskRequest** | 17 | A computer will listen for and respond to requests for a network subnet mask. |

**Open Ports**

Add the following entries to the Windows Firewall Netfw.inf file to add static ports to the exceptions list:

| Section | Entry |
|---|---|

| | |
|---|---|
| [ICF.AddReg.DomainProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile\GloballyOpenPorts\List", "port number:protocol",0x00000000, "port number:protocol:scope:mode:port's friendly name"` |
| [ICF.AddReg.StandardProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile\GloballyOpenPorts\List", "port number:protocol",0x00000000,"port number:protocol:scope:mode:port's friendly name"` |

In the two preceding entries, the following elements must be defined:

- **Port Number** – A port is specified by the combination of a protocol and a port number. The port number must be between 1 and 65535, inclusively.
- **Protocol** – A protocol is specified by the combination of a protocol and a port number. The protocol must be either TCP or UDP.
- **Scope** – See "*Scope Definition*" earlier in this section.
- **Mode** – The two permitted values for this element are enabled and disabled. If a port's entry is enabled, the port is statically opened in Windows Firewall. If a port's entry is disabled, the port is not statically opened in Windows Firewall.
- **Port's Friendly Name** – This is the description that will be used to represent the entry in Windows Firewall in Control Panel. It should provide an indication of why the port is statically opened, such as "Web Server (TCP 80)" or "Telnet Server (TCP 23)."

**Allow Programs**

Add the following entries to the Windows Firewall Netfw.inf file to add programs to the exceptions list:

| Section | Entry |
|---|---|
| [ICF.AddReg.DomainProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile\AuthorizedApplications\List", "program's image path",0x00000000, "program's image path:scope:mode:program's friendly name"` |
| [ICF.AddReg.StandardProfile] | `HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile\AuthorizedApplications\List", "program's image path",0x00000000, "program's image path:scope:mode:program's friendly name"` |

In the two preceding entries, the following elements must be defined:

**Program's Image Path** – This is the fully qualified path for the file to be added to Windows Firewall's default exception list. It may include environmental variables, such as %ProgramFiles%.

**Scope** – See "*Scope Definition*" earlier in this section.

**Mode** –The two permitted values for this element are enabled and disabled. If a program's entry is enabled, ports for the program are dynamically opened in Windows Firewall. If a program's entry is disabled, ports for the program are not dynamically opened in Windows Firewall.

**Program's Friendly Name** – This is the name that will be used to represent the entry in the Windows Firewall user interface. It should include the product name and publisher, such as MSN Messenger version 6.1.

**Windows Firewall Unattended Installation**

You can use an unattended installation answer file, such as Unattend.txt, to automate the configuration of Windows Firewall settings during an unattended installation. Unattended installation is typically used in a corporate environment when you are configuring new computers or deploying a new operating system.

To configure Windows Firewall settings during an unattended installation, you must include the [WindowsFirewall] section along with one or more of the following user-defined sections in your answer file:

- [WindowsFirewall.profile_name]
- [WindowsFirewall.program_name]
- [WindowsFirewall.service_name]
- [WindowsFirewall.portopening_name]
- [WindowsFirewall.icmpsetting_name]

| **Important** |
| --- |
| Applications that are not already installed cannot be added using unattended setup. You must use Netfw.inf to add applications. |

| **Note** |
| --- |
| When you restore Windows Firewall default settings, any settings that you specified in an answer file are deleted and are not restored. To restore the Windows Firewall settings that you configured during installation with an answer file, you must configure the settings manually. |

**[WindowsFirewall]**

The [WindowsFirewall] section contains entries for specifying the log file settings and which user-defined profiles to use.

The [WindowsFirewall] section contains the following entries:

| Entry | Description |
| --- | --- |
| **Profiles** | Specifies the names of the user-defined profiles to use for configuring Windows Firewall (domain, standard, or both domain and standard). |
| **LogFile** | Specifies the location and file name of the Windows Firewall log file. By default, the log file is named Pfirewall.log. |
| **LogSize** | Specifies the maximum size of the Windows Firewall log file. |
| **LogDroppedPackets** | Specifies whether to log dropped packets to the Windows Firewall log file. |
| **LogConnections** | Specifies whether to log connections to the Windows Firewall log file. |

**[WindowsFirewall.profile_name]**

The [WindowsFirewall.profile_name] section is a user-defined section that is referenced by the [WindowsFirewall] section to make changes to Windows Firewall's default configuration.

The [WindowsFirewall.profile_name] section contains the following entries:

| Entry | Description |
| --- | --- |
| **Type** | Specifies the type of profile to use for changing the default configuration of Windows Firewall. |
| **Mode** | Specifies whether to enable or disable Windows Firewall. |
| **Exceptions** | Specifies whether to enable or disable the Windows Firewall exceptions list. |
| **Notifications** | Specifies whether to enable or disable notifications. |
| **MulticastBroadcastResponse** | Specifies whether to enable or disable multicast and broadcast packets. |

| | |
|---|---|
| **AllowedPrograms** | Specifies a list of programs that will not be blocked by Windows Firewall. |
| **Services** | Specifies a list of services that will not be blocked by Windows Firewall. |
| **PortOpenings** | Specifies a list of open ports that will not be blocked by Windows Firewall. |
| **IcmpSettings** | Specifies a list of ICMP message types that will not be blocked by Windows Firewall. |

**[WindowsFirewall.program_name]**

The [WindowsFirewall.program_name] section is a user-defined section that can be used to add programs to the Windows Firewall exceptions list.

The [WindowsFirewall.program_name] section contains the following entries:

| Entry | Description |
|---|---|
| **Program** | Specifies the path of a program to be added to the exceptions list. This is a required entry. |
| **Name** | Specifies the name of a program to be added to the exceptions list. This is a required entry. |
| **Mode** | Specifies whether to enable or disable an entry in the exceptions list. |
| **Scope** | Defines the set of limits on which computers (IP addresses) are allowed to send traffic through the specified exception (program, service). Mode must be set to **1** (on). |
| **Addresses** | Specifies the addresses for an entry in the exceptions list. |

**[WindowsFirewall.service_name]**

Windows Firewall opens static ports used by services in the exceptions list of the current profile. Only services that actually require unsolicited, incoming traffic should be added to the exceptions list. You must add the [WindowsFirewall.service_name] section in the [WindowsFirewall.profile_name] section.

The [WindowsFirewall.service_name] section contains the following entries:

| Entry | Description |
|---|---|
| **Type** | Specifies the type of service to use for changing the default configuration of Windows Firewall. |
| **Mode** | Specifies whether to enable or disable an entry in the exceptions lists. |
| **Scope** | Defines the set of limits on which computers (IP addresses) are allowed to send traffic through the specified exception. The value of the Mode entry must equal **1** (on). |
| **Addresses** | Specifies the addresses for an entry in the exceptions list. |

**[WindowsFirewall.portopening_name]**

A static port may need to be opened for a Windows service to receive unsolicited, incoming traffic. To support these scenarios, you can add static ports to the Windows Firewall exceptions list by using the [WindowsFirewall.portopening_name] section. You must add this section to the [WindowsFirewall.profile_name] section.

The [WindowsFirewall.portopening_name] section contains the following entries:

| Entry | Description |
|---|---|
| **Protocol** | Specifies the protocol of a port. The protocol must be either TCP or UDP. |
| **Port** | Specifies the port number. The port number must be between 1 and 65535, inclusive. |
| **Name** | Specifies the friendly name of a port to be added to the exceptions list. This descriptive name is used to represent the entry for Windows Firewall in Control Panel. |
| **Mode** | Specifies whether to enable or disable an entry in the exceptions list. |
| **Scope** | Defines the set of limits on which computers (IP addresses) are allowed to send traffic through the specified exception. The value of the Mode entry must equal **1** (on). |

| | |
|---|---|
| **Addresses** | Specifies the addresses for an entry in the exceptions list. |

**[WindowsFirewall.icmpsetting_name]**
The default configuration for Windows Firewall blocks all ICMP message types; however, you can modify this behavior by adding entries to the Windows Firewall exceptions list that enable certain ICMP message types. You must include the [WindowsFirewall.icmpsetting_name] section in the [WindowsFirewall.profile_name] section.

The [WindowsFirewall.icmpsetting_name] section contains the following entries:

| Entry | Description |
|---|---|
| **Type** | Specifies the type of ICMP message to enable. |
| **Mode** | Specifies whether to enable or disable the ICMP message type. |

**Windows Firewall Scripting Reference**
This section describes the methods and properties associated with Windows Firewall scripting. You can use them in a script, such as Microsoft Visual Basic Scripting Edition (VBScript) or JScript, to configure Windows Firewall settings. All of these methods and properties are implemented in the Hnetcfg.dll. They are grouped in this section according to the following categories:

- Policy
- Profile
- Remote administration
- ICMP
- Port
- Application
- Service

**Policy**
The following table lists the scripting properties used to access a Windows Firewall policy:

| Property | Description |
|---|---|
| **LocalPolicy** | A read-only element that accesses the local firewall policy. This property is retrieved through the HNetCfg.FwMgr COM Object. |

**Profile**
The following table lists the scripting methods and properties used to access and configure a Windows Firewall profile:

| Method or Property | Description |
|---|---|
| **AuthorizedApplications** | A read-only element that accesses the collection of authorized applications in a profile. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **CurrentProfile** | A read-only element that accesses the current Windows Firewall profile. This property is retrieved through the HNetCfg.FwMgr COM Object (**LocalPolicy** property). |
| **CurrentProfileType** | A read-only element that accesses the type of Windows Firewall profile currently in effect. This property is retrieved through the HNetCfg.FwMgr COM Object. |

| | |
|---|---|
| **ExceptionsNotAllowed** | A read-write element that accesses a Boolean value which is TRUE if Windows Firewall should not allow exceptions. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **FirewallEnabled** | A read-write element that accesses a Boolean value which is TRUE if Windows Firewall is enabled. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **GetProfileByType** | Gets the Windows Firewall profile of the requested type. This method is retrieved through the HNetCfg.FwMgr COM Object (**LocalPolicy** property). |
| **GloballyOpenPorts** | A read-only element that accesses the collection of globally-opened ports in a profile. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **IcmpSettings** | A read-only element that accesses the ICMP settings in a profile. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **NotificationsDisabled** | A read-write element that accesses a Boolean value which is TRUE if interactive notifications are disabled. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **RemoteAdminSettings** | Accesses the object that contains the remote administration settings. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **Services** | A read-only element that accesses the collection of services in a profile. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **Type** | A read-only element that accesses the type of a profile. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **UnicastResponsesToMulticastBroadcastDisabled** | A read-write element that accesses a Boolean value which is TRUE if Windows Firewall should not allow unicast responses to multicast and broadcast traffic. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |

**Remote Administration**

The following table lists the scripting properties used to access the settings that control remote administration:

| Property | Description |
|---|---|
| **Enabled** | A read-write element that accesses a Boolean value which is TRUE if the settings controlling remote administration are currently enabled. This property is retrieved through the HNetCfg.FwMgr COM Object (**RemoteAdminSettings** property). |
| **IpVersion** | A read-write element that accesses the IP version for which remote administration is authorized. This property is retrieved through the HNetCfg.FwMgr COM Object (**RemoteAdminSettings** property). |
| **RemoteAddresses** | Accesses the set of remote addresses from which remote administration is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**RemoteAdminSettings** property). |
| **Scope** | A read-write element that controls the network scope from which remote administration is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**RemoteAdminSettings** property). |

**ICMP**

The following table lists the scripting methods and properties used to access and configure the settings controlling ICMP packets:

| Method or Property | Description |
|---|---|
| **AllowInboundEchoRequest** | A read-write element that accesses a Boolean value which is TRUE if InboundEchoRequest is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **AllowInboundMaskRequest** | A read-write element that accesses a Boolean value which is TRUE if InboundMaskRequest is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **AllowInboundRouterRequest** | A read-write element that accesses a Boolean value which is TRUE if InboundRouterRequest is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **AllowInboundTimestampRequest** | A read-write element that accesses a Boolean value which is TRUE if InboundTimestampRequest is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **AllowOutboundDestinationUnreachable** | A read-write element that accesses a Boolean value which is TRUE if OutboundDestinationUnreachable is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **AllowOutboundPacketTooBig** | A read-write element that accesses a Boolean value which is TRUE if OutboundPacketTooBig is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **AllowOutboundParameterProblem** | A read-write element that accesses a Boolean value which is TRUE if OutboundParameterProblem is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **AllowOutboundSourceQuench** | A read-write element that accesses a Boolean value which is TRUE if OutboundSourceQuench is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **AllowOutboundTimeExceeded** | A read-write element that accesses a Boolean value which is TRUE if OutboundTimeExceeded is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **AllowRedirect** | A read-write element that accesses a Boolean value which is TRUE if Redirect is allowed. This property is retrieved through the HNetCfg.FwMgr COM Object (**IcmpSettings** property). |
| **IsIcmpTypeAllowed** | Determines whether the specified ICMP type is allowed. This method is retrieved through the HNetCfg.FwMgr COM Object. |

| **Note** |
|---|
| All of the methods and properties associated with ICMP are common to both IPv4 and IPv6. |

**Port**

The following table lists the scripting methods and properties used to access and configure a port that has been opened in Windows Firewall:

| Method or Property | Description |
|---|---|
| **_NewEnum** | Returns an object supporting IEnumVARIANT that can be used to iterate through all the ports in a collection. This property is retrieved through the HNetCfg.FwMgr COM Object [**GloballyOpenPorts (Profile)** property]. |

| | |
|---|---|
| **Add** | Adds a new port to a collection. This method is retrieved through the HNetCfg.FwMgr COM Object [**GloballyOpenPorts (Profile)** property]. |
| **BuiltIn** | A read-only element that accesses a Boolean value which is TRUE if a port is defined by the system. This property is retrieved through the HNetCfg.FWOpenPort COM Object. |
| **Count** | A read-only element yielding the number of items in a collection of ports. This property is retrieved through the HNetCfg.FwMgr COM Object [**GloballyOpenPorts (Profile)** property]. |
| **Enabled** | A read-write element that accesses a Boolean value which is TRUE if the settings for a port are currently enabled. This property is retrieved through the HNetCfg.FWOpenPort COM Object. |
| **IpVersion** | A read-write element that accesses the IP version of a port. This property is retrieved through the HNetCfg.FWOpenPort COM Object. |
| **IsPortAllowed** | Determines whether an application can listen for inbound traffic on a specified port. This method is retrieved through the HNetCfg.FwMgr COM Object. |
| **Item** | Returns a specified port if it is within a collection. This method is retrieved through the HNetCfg.FwMgr COM Object [**GloballyOpenPorts (Profile)** property]. |
| **Name** | A read-write element that accesses the friendly name of a port. This property is retrieved through the HNetCfg.FWOpenPort COM Object. |
| **Port** | A read-write element that accesses the port number of a port. This property is retrieved through the HNetCfg.FWOpenPort COM Object. |
| **Protocol** | A read-write element that accesses the protocol type setting of a port. This property is retrieved through the HNetCfg.FWOpenPort COM Object. |
| **RemoteAddresses** | Accesses a set of remote addresses from which a port can listen for traffic. This property is retrieved through the HNetCfg.FWOpenPort COM Object. |
| **Remove** | Removes a port from a collection. This method is retrieved through the HNetCfg.FwMgr COM Object [**GloballyOpenPorts (Profile)** property]. |
| **Scope** | A read-write element that controls the network scope from which a port can listen for traffic. This property is retrieved through the HNetCfg.FWOpenPort COM Object. |

**Application**

The following table lists and gives a short description of the scripting methods and properties used to access and configure an application that has been added to the Windows Firewall exceptions list.

| Method or Property | Description |
|---|---|
| **_NewEnum** | Returns an object supporting IEnumVARIANT that can be used to iterate through all the applications in a collection. This property is retrieved through the HNetCfg.FwMgr COM Object (**AuthorizedApplications** property). |
| **Add** | Adds a new application to a collection. This method is retrieved through the HNetCfg.FwMgr COM Object (**AuthorizedApplications** property). |
| **Count** | A read-only element yielding the number of items in a collection of applications. This property is retrieved through the HNetCfg.FwMgr COM Object (**AuthorizedApplications** property). |
| **Enabled** | A read-write element that accesses a Boolean value which is TRUE if the settings for an application are currently enabled. This property is retrieved through the HNetCfg.FwAuthorizedApplication COM Object. |
| **IpVersion** | A read-write element that accesses the IP version for an application. This property is retrieved through the HNetCfg.FwAuthorizedApplication COM Object. |
| **IsPortAllowed** | Determines whether an application can listen for inbound traffic on a specified port. This method is retrieved through the HNetCfg.FwMgr COM Object. |

| | |
|---|---|
| **Item** | Returns a specified application if it is within a collection. This method is retrieved through the HNetCfg.FwMgr COM Object (**AuthorizedApplications** property). |
| **Name** | A read-write element that accesses the friendly name of an application. This property is retrieved through the HNetCfg.FwAuthorizedApplication COM Object. |
| **ProcessImageFileName** | A read-write element that accesses the process image file name of an application. This property is retrieved through the HNetCfg.FwAuthorizedApplication COM Object. |
| **RemoteAddresses** | Accesses the set of remote addresses from which an application can listen for traffic. This property is retrieved through the HNetCfg.FwAuthorizedApplication COM Object. |
| **Remove** | Removes an application from a collection. This method is retrieved through the HNetCfg.FwMgr COM Object (**AuthorizedApplications** property). |
| **Scope** | A read-write element that controls the network scope from which an application can listen for traffic. This property is retrieved through the HNetCfg.FwAuthorizedApplication COM Object. |

**Service**

The following table lists the scripting methods and properties used to access and configure a service that has been authorized to listen through Windows Firewall:

| Method or Property | Description |
|---|---|
| **_NewEnum** | Returns an object supporting IEnumVARIANT that can be used to iterate through all the services in a collection. This property is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |
| **Count** | A read-only element yielding the number of items in a collection of services. This property is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |
| **Customized** | A read-only element that indicates whether at least one of the ports associated with a service has been customized. This property is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |
| **Enabled** | A read-write element that accesses a Boolean value which is TRUE if all ports associated with the service are enabled. This property is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |
| **GloballyOpenPorts** | A read-only element that accesses the collection of globally-opened ports associated with a service. This property is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |
| **Item** | Returns a specified service if it is within a collection. This method is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |
| **Name** | A read-only element that accesses the friendly name of a service. This property is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |
| **RemoteAddresses** | Accesses the set of remote addresses from which a service can listen for traffic. This property is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |
| **Scope** | A read-write element that controls the network scope from which a service can listen. This property is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |
| **Services** | A read-only element that accesses the collection of services in a profile. This property is retrieved through the HNetCfg.FwMgr COM Object (**CurrentProfile** property or the **GetProfileByType** method). |
| **Type** | A read-only element that accesses the type of a service. This property is retrieved through the HNetCfg.FwMgr COM Object (**Services** property). |

For more information about Windows Firewall interfaces, see MSDN.

**Netsh Commands for Windows Firewall**

You can run these commands from the Windows Server 2003 family command prompt or from the command prompt for the **netsh firewall** context. For these commands to work at the Windows Server 2003 family command prompt, you must type **netsh firewall** before typing commands and parameters as they appear in the following syntax. There might be functional differences between **netsh** context commands on Windows 2000 and the Windows Server 2003 family.

You cannot use the **netsh firewall** commands to remotely configure Windows Firewall settings, and you cannot use the **dump** command to create a script based on the current Windows Firewall configuration.

To view help for a command at the command prompt, type the following:

*CommandName* **/?**

where *CommandName* is the name of the command.

- add allowedprogram
- add portopening
- delete allowedprogram
- delete portopening
- set allowedprogram
- set icmpsetting
- set logging
- set multicastbroadcastresponse
- set notifications
- set opmode
- set portopening
- set service
- show (all commands)

**add allowedprogram**

The **add allowedprogram** command is used to add a program-based exception.

**Syntax**

**add allowedprogram** [[**program** =] *path*] [[**name** =] *name*] [[**mode** =] {**ENABLE**|**DISABLE**}] [[**scope** =] {**ALL**|**SUBNET**|**CUSTOM**}] [[**addresses** =] *addresses*] [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}]

**Parameters**

**[[program =] *path*]**

Specifies the path and file name of a program to be added to the exceptions list. This is a required entry.

**[[name =] *name*]**

Specifies the name of a program to be added to the exceptions list. This is a required entry.

**[[mode =] {ENABLE|DISABLE}]**

Specifies whether to enable or disable a program. This is an optional entry.

- **ENABLE** - Allow through Windows Firewall (default).
- **DISABLE** - Do not allow through Windows Firewall.

**[[scope =] {ALL|SUBNET|CUSTOM}]**

Defines the set of limits on which computers (IP addresses) are allowed to send traffic through the specified program. This is an optional entry.

- **ALL** - Allow all traffic through Windows Firewall (default).
- **SUBNET** - Allow only local network (subnet) traffic through Windows Firewall.
- **CUSTOM** - Allow only specified traffic through Windows Firewall.

**[[addresses =] *addresses*]**
Specifies the custom scope addresses in the exceptions list. This is an optional entry.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**
Specifies the names of the profiles used to configure Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

**Remarks**
**scope** must be **CUSTOM** to specify **addresses**.

**Examples**
The following examples show how the **add allowedprogram** command and preceding parameters can be used to add a program-based exception:

**add allowedprogram** *C:\MyApp\MyApp.exe MyApp* **ENABLE**

**add allowedprogram** *C:\MyApp\MyApp.exe MyApp* **DISABLE**

**add allowedprogram** *C:\MyApp\MyApp.exe MyApp* **ENABLE CUSTOM** *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**add allowedprogram program** = *C:\MyApp\MyApp.exe* **name** = *MyApp* **mode** = **ENABLED**

**add allowedprogram program** = *C:\MyApp\MyApp.exe* **name** = *MyApp* **mode** = **DISABLE**

**add allowedprogram program** = *C:\MyApp\MyApp.exe* **name** = *MyApp* **mode** = **ENABLE scope** = **CUSTOM** *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**add portopening**
The **add portopening** command is used to create a port-based exception.

**Syntax**
**add portopening** [[**protocol** =] {**TCP**|**UDP**|**ALL**}] [[**port** =] **1**-**65535**] [[**name** =] *name*] [[**mode** =] {**ENABLE**|**DISABLE**}] [[**scope** =] {**ALL**|**SUBNET**|**CUSTOM**}] [[**addresses** =] *addresses*] [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}] [[**interface** =] *name*]

**Parameters**
**[[protocol =] {TCP|UDP|ALL}]**
Specifies the protocol of a port. The protocol must be TCP, UDP, or all. This is a required entry.

- **TCP** - Transmission Control Protocol (TCP).
- **UDP** - User Datagram Protocol (UDP).
- **ALL** - All protocols.

**[[port =] 1-65535]**
Specifies the port number. The port number must be between 1 and 65535, inclusive. This is a required entry.

**[[name =] *name*]**
Specifies the friendly name of a port to be added to the exceptions list. This descriptive name is used to represent the entry for Windows Firewall in Control Panel. This is a required entry.

**[[mode =] {ENABLE|DISABLE}]**
Specifies whether to enable or disable a port in the exceptions lists. This is an optional entry.

- **ENABLE** - Allow through Windows Firewall (default).
- **DISABLE** - Do not allow through Windows Firewall.

**[[scope =] {ALL|SUBNET|CUSTOM}]**
Defines the set of limits on which computers (IP addresses) are allowed to send traffic through the specified port. This is an optional entry.

- **ALL** - Allow all traffic through Windows Firewall (default).
- **SUBNET** - Allow only local network (subnet) traffic through Windows Firewall.
- **CUSTOM** - Allow only specified traffic through Windows Firewall.

**[[addresses =] *addresses*]**
Specifies the custom scope addresses in the exceptions list. This is an optional entry.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**
Specifies the names of the profiles used to configure Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

**[[interface =] *name*]**
Specifies the interface name. This is an optional entry.

**Remarks**
**profile** and **interface** cannot be specified together; **scope** and **interface** cannot be specified together. **scope** must be **CUSTOM** to specify **addresses**.

**Examples**
The following examples show how the **add portopening** command and preceding parameters can be used to create a port-based exception:

**add portopening TCP 80** *MyWebPort*

**add portopening UDP 500** *IKE* **ENABLE ALL**

**add portopening ALL 53** *DNS* **ENABLE CUSTOM** *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**add portopening protocol** = **TCP port** = **80name** = *MyWebPort*

**add portopening protocol** = **UDP port** = **500name** = *IKE* **mode** = **ENABLE scope** = **ALL**

**add portopening protocol** = **ALL port** = **53name** = *DNS* **mode** = **ENABLE scope** = **CUSTOM addresses** = *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**delete allowedprogram**
The **delete allowedprogram** command is used to delete an existing program-based exception.

#### Syntax

**delete allowedprogram** [[**program** =] *path* [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}]

#### Parameters

**[[program =] *path***

Specifies the path and file name of the program to be deleted from the exceptions list. This is a required entry.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**

Specifies the names of the profiles used for configuring Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

#### Examples

The following examples show how the **delete allowedprogram** command and preceding parameters can be used to delete an existing program-based exception:

**delete allowedprogram** *C:\MyApp\MyApp.exe*

**delete allowedprogram program** = *C:\MyApp\MyApp.exe*

#### delete portopening

The **delete portopening** command is used to delete a port-based exception.

#### Syntax

**delete portopening** [[**protocol** =] {**TCP**|**UDP**|**ALL**}] [[**port** =] **1**-**65535**] [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}] [[**interface** =] *name*]

#### Parameters

**[[protocol =] {TCP|UDP|ALL}]**

Specifies the protocol of the port to be deleted from the exceptions list. The protocol must be TCP, UDP, or all. This is a required entry.

- **TCP** - Transmission Control Protocol (TCP).
- **UDP** - User Datagram Protocol (UDP).
- **ALL** - All protocols.

**[[port =] 1-65535]**

Specifies the number of the port to be deleted from the exceptions list. The port number must be between 1 and 65535, inclusive. This is a required entry.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**

Specifies the names of the profiles used for configuring Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

**[[interface =] *name*]**

Specifies the interface name. This is an optional entry.

**Remarks**
**profile** and **interface** cannot be specified together.

**Examples**
The following examples show how the **delete portopening** command and preceding parameters can be used to delete a port-based exception:

**delete portopening TCP 80**

**delete portopening UDP 500**

**delete portopening protocol** = **TCP port** = **80**

**delete portopening protocol** = **UDP port** = **500**

**set allowedprogram**
The **set allowedprogram** command is used to modify the settings of a program-based exception.

**Syntax**
**set allowedprogram** [[**program** =] *path*] [[**name** =] *name*] [[**mode** =] {**ENABLE**|**DISABLE**}] [[**scope** =] {**ALL**|**SUBNET**|**CUSTOM**}] [[**addresses** =] *addresses* [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}]

**Parameters**
**[[program =] *path*]**
Specifies the path and file name of a program to be added to the exceptions list. This is a required entry.

**[[name =] *name*]**
Specifies the name of a program to be added to the exceptions list. This is a required entry.

**[[mode =] {ENABLE|DISABLE}]**
Specifies whether to enable or disable a program. This is an optional entry.

- **ENABLE** - Allow through Windows Firewall (default).
- **DISABLE** - Do not allow through Windows Firewall.

**[[scope =] {ALL|SUBNET|CUSTOM}]**
Defines the set of limits on which computers (IP addresses) are allowed to send traffic through the specified program. This is an optional entry.

- **ALL** - Allow all traffic through Windows Firewall (default).
- **SUBNET** - Allow only local network (subnet) traffic through Windows Firewall.
- **CUSTOM** - Allow only specified traffic through Windows Firewall.

**[[addresses =] *addresses*]**
Specifies the custom scope addresses in the exceptions list. This is an optional entry.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**
Specifies the names of the profiles used to configure Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

**Remarks**
**scope** must be **CUSTOM** to specify **addresses**.

**Examples**

The following examples show how the **set allowedprogram** command and preceding parameters can be used to modify the settings of a program-based exception:

**set allowedprogram** *C:\MyApp\MyApp.exe MyApp* **ENABLE**

**set allowedprogram** *C:\MyApp\MyApp.exe MyApp* **DISABLE**

**set allowedprogram** *C:\MyApp\MyApp.exe MyApp* **ENABLE CUSTOM** *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**set allowedprogram program** = *C:\MyApp\MyApp.exe* **name** = *MyApp* **mode** = **ENABLED**

**set allowedprogram program** = *C:\MyApp\MyApp.exe* **name** = *MyApp* **mode** = **DISABLE**

**set allowedprogram program** = *C:\MyApp\MyApp.exe* **name** = *MyApp* **mode** = **ENABLE scope** = **CUSTOM** *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**set icmpsetting**

The **set icmpsetting** command is used to specify Internet Control Message Protocol (ICMP) traffic that has been added to the exceptions list.

**Syntax**

**set icmpsetting** [[**type** =] {**2**|**3**|**4**|**5**|**8**|**9**|**11**|**12**|**13**|**17**|**ALL**}] [[**mode** =] {**ENABLE**|**DISABLE**}] [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}] [[**interface** =] *name*]

**Parameters**

**[[type =] {2|3|4|5|8|9|11|12|13|17|ALL}]**

Specifies the type of ICMP message to enable. This is a required entry.

- **2** - Allow outbound packet too big.
- **3** - Allow outbound destination unreachable.
- **4** - Allow outbound source quench.
- **5** - Allow redirect.
- **8** - Allow inbound echo request.
- **9** - Allow inbound router request.
- **11** - Allow outbound time exceeded.
- **12** - Allow outbound parameter problem.
- **13** - Allow inbound timestamp request.
- **17** - Allow inbound mask request.
- **ALL** - All types.

**[[mode =] {ENABLE|DISABLE}]**

Specifies whether to enable or disable the ICMP message type. This is an optional entry.

- **ENABLE** - Allow through Windows Firewall (default).
- **DISABLE** - Do not allow through Windows Firewall.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**

Specifies the names of the profiles used to configure Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

**[[interface =] *name*]**
Specifies the interface name. This is an optional entry.

**Remarks**
**profile** and **interface** cannot be specified together; **type 2** and **interface** cannot be specified together.

**Examples**
The following examples show how the **set icmpsetting** command and preceding parameters can be used to specify ICMP traffic that has been added to the exceptions list:

**set icmpsetting 8**

**set icmpsetting 8 ENABLE**

**set icmpsetting ALL DISABLE**

**set icmpsetting type** = **8**

**set icmpsetting type** = **8mode** = **ENABLE**

**set icmpsetting type** = **ALL mode** = **DISABLE**

**set logging**
The **set logging** command is used to specify Windows Firewall logging options.

**Syntax**
**set logging** [[**filelocation** =] *path*] [[**maxfilesize** =] **1**-**32767**] [[**droppedpackets** =] {**ENABLE**|**DISABLE**}] [[**connections** =] {**ENABLE**|**DISABLE**}]

**Parameters**
**[[filelocation =] *path*]**
Specifies the location and file name of the Windows Firewall log file. By default, the log file is named Pfirewall.log. This is a required entry.

**[[maxfilesize =] 1-32767]**
Specifies the maximum size (in kilobytes) of the Pfirewall.log file. This is an optional entry.

**[[droppedpackets =] {ENABLE|DISABLE}]**
Specifies whether to log dropped packets to the Pfirewall.log file. This is an optional entry.

- **ENABLE** - Log.
- **DISABLE** - Do not log.

**[[connections =] {ENABLE|DISABLE}]**
Specifies whether to log connections to the Pfirewall.log file. This is an optional entry.

- **ENABLE** - Log.
- **DISABLE** - Do not log.

**Remarks**
At least one parameter must be specified.

**Examples**

The following examples show how the **set logging** command and preceding parameters can be used to specify Windows Firewall logging options:

**set logging** *%windir%\pfirewall.log* **4096**

**set logging** *%windir%\pfirewall.log* **4096 ENABLE**

**set logging filelocation** = *%windir%\pfirewall.log* **maxfilesize** = **4096**

**set logging filelocation** = *%windir%\pfirewall.log* **maxfilesize** = **4096droppedpackets** = **ENABLE**

#### set multicastbroadcastresponse

The **set multicastbroadcastresponse** command is used to specify the unicast response to a multicast or broadcast request.

#### Syntax

**set multicastbroadcastresponse** [[**mode** =] {**ENABLE**|**DISABLE**}] [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}]

#### Parameters

**[[mode =] {ENABLE|DISABLE}]**

Specifies whether to enable or disable multicast and broadcast packets. This is a required entry.

- **ENABLE** - Allow responses to multicast/broadcast traffic through Windows Firewall.
- **DISABLE** - Do not allow responses to multicast/broadcast traffic through Windows Firewall.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**

Specifies the names of the profiles used to configure Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

#### Examples

The following examples show how the **set multicastbroadcastresponse** command and preceding parameters can be used to specify the unicast response to a multicast or broadcast request:

**set multicastbroadcastresponse ENABLE**

**set multicastbroadcastresponse DISABLE**

**set multicastbroadcastresponse mode** = **ENABLE**

**set multicastbroadcastresponse mode** = **DISABLE**

#### set notifications

The **set notifications** command is used to specify the behavior of Windows Firewall notifications.

#### Syntax

**set notifications** [[**mode** =] {**ENABLE**|**DISABLE**}] [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}]

#### Parameters

**[[mode =] {ENABLE|DISABLE}]**

Specifies whether to enable or disable notifications. This is a required entry.

**ENABLE** - Allow notifications from Windows Firewall.

**DISABLE** - Do not allow notifications from Windows Firewall.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**

Specifies the names of the profiles used to configure Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

**Examples**

The following examples show how the **set notifications** command and preceding parameters can be used to specify the behavior of Windows Firewall notifications:

**set notifications ENABLE**

**set notifications DISABLE**

**set notifications mode** = **ENABLE**

**set notifications mode** = **DISABLE**

**set opmode**

The **set opmode** command is used to specify the operating mode of Windows Firewall, either globally or for a specific connection (interface).

**Syntax**

**set opmode** [[**mode** =] {**ENABLE**|**DISABLE**}] [[**exceptions** =] {**ENABLE**|**DISABLE**}] [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}] [[**interface** =] *name*]

**Parameters**

**[[mode =] {ENABLE|DISABLE}]**

Specifies whether to enable or disable Windows Firewall. This is a required entry.

**ENABLE** - Enable Windows Firewall.

**DISABLE** - Disable Windows Firewall.

**[[exceptions =] {ENABLE|DISABLE}]**

Specifies whether to enable or disable the Windows Firewall exceptions list. This is a required entry.

**ENABLE** - Enable the Windows Firewall exceptions list.

**DISABLE** - Disable the Windows Firewall exceptions list.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**

Specifies the names of the profiles used to configure Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

**[[interface =] *name*]**

Interface name (optional).

**Remarks**

**profile** and **interface** cannot be specified together; **exceptions** and **interface** cannot be specified together.

**Examples**

The following examples show how the **set opmode** command and preceding parameters can be used to specify the operating mode of Windows Firewall:

**set opmode ENABLE**

**set opmode ENABLE DISABLE**

**set opmode mode** = **ENABLE**

**set opmode mode** = **ENABLE exceptions** = **DISABLE**

#### set portopening

The **set portopening** command is used to modify the settings of a port-based exception.

#### Syntax

**set portopening** [[**protocol** =] {**TCP**|**UDP**|**ALL**}] [[**port** =] **1**-**65535**] [[**name** =] *name*] [[**mode** =] {**ENABLE**|**DISABLE**}] [[**scope** =] {**ALL**|**SUBNET**|**CUSTOM**}] [[**addresses** =] *addresses*] [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}] [[**interface** =] *name*]

#### Parameters

**[[protocol =] {TCP|UDP|ALL}]**

Specifies the protocol of a port. The protocol must be TCP, UDP, or all. This is a required entry.

- **TCP** - Transmission Control Protocol (TCP).
- **UDP** - User Datagram Protocol (UDP).
- **ALL** - All protocols.

**[[port =] 1-65535]**

Specifies the port number. The port number must be between 1 and 65535, inclusive. This is a required entry.

**[[name =] *name*]**

Specifies the friendly name of a port to be added to the exceptions list. This descriptive name is used to represent the entry for Windows Firewall in Control Panel. This is a required entry.

**[[mode =] {ENABLE|DISABLE}]**

Specifies whether to enable or disable a port in the exceptions lists. This is an optional entry.

- **ENABLE** - Allow through Windows Firewall (default).
- **DISABLE** - Do not allow through Windows Firewall.

**[[scope =] {ALL|SUBNET|CUSTOM}]**

Defines the set of limits on which computers (IP addresses) are allowed to send traffic through the specified port. This is an optional entry.

- **ALL** - Allow all traffic through Windows Firewall (default).
- **SUBNET** - Allow only local network (subnet) traffic through Windows Firewall.
- **CUSTOM** - Allow only specified traffic through Windows Firewall.

**[[addresses =] *addresses*]**

Specifies the custom scope addresses in the exceptions list. This is an optional entry.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**

Specifies the names of the profiles used to configure Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.

- **ALL** - All profiles.

**[[interface =] *name*]**
Specifies the interface name. This is an optional entry.

**Remarks**
**profile** and **interface** cannot be specified together; **scope** and **interface** cannot be specified together. **scope** must be **CUSTOM** to specify **addresses**.

**Examples**
The following examples show how the **set portopening** command and preceding parameters can be used to modify the settings of a port-based exception:

**set portopening TCP 80** *MyWebPort*

**set portopening UDP 500** *IKE* **ENABLE ALL**

**set portopening ALL 53** *DNS* **ENABLE CUSTOM** *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**set portopening protocol** = **TCP port** = **80name** = *MyWebPort*

**set portopening protocol** = **UDP port** = **500name** = *IKE* **mode** = **ENABLE scope** = **ALL**

**set portopening protocol** = **ALL port** = **53name** = *DNS* **mode** = **ENABLE scope** = **CUSTOM addresses** = *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**set service**
The **set service** command is used to enable or disable the predefined file and printer sharing, remote administration, Remote Desktop, and UPnP exceptions.

**Syntax**
**set service** [[**type** =] {**FILEANDPRINT**|**REMOTEADMIN**|**REMOTEDESKTOP**|**UPNP**|**ALL**}] [[**mode** =] {**ENABLE**|**DISABLE**}] [[**scope** =] {**ALL**|**SUBNET**|**CUSTOM**}] [[**addresses** =] *addresses* [[**profile** =] {**CURRENT**|**DOMAIN**|**STANDARD**|**ALL**}]

**Parameters**
**[[type =] {FILEANDPRINT|REMOTEADMIN|REMOTEDESKTOP|UPNP|ALL}]**
Specifies the type of service to enable. This is a required entry.

**FILEANDPRINT** - File and printer sharing.

**REMOTEADMIN** - Remote administration.

**REMOTEDESKTOP** - Remote assistance and remote desktop.

**UPNP** - UPnP framework.

**ALL** - All types.

**[[mode =] {ENABLE|DISABLE}]**
Specifies whether to enable or disable a service. This is a required entry.

**ENABLE** - Allow through Windows Firewall (default).

**DISABLE** - Do not allow through Windows Firewall.

**[[scope =] {ALL|SUBNET|CUSTOM}]**
Defines the set of limits on which computers (IP addresses) are allowed to send traffic through the specified service. This is an optional entry.

- **ALL** - Allow all traffic through Windows Firewall (default).
- **SUBNET** - Allow only local network (subnet) traffic through Windows Firewall.
- **CUSTOM** - Allow only specified traffic through Windows Firewall.

**[[addresses =] *addresses***
Specifies the custom scope addresses in the exceptions list. This is an optional entry.

**[[profile =] {CURRENT|DOMAIN|STANDARD|ALL}]**
Specifies the names of the profiles used to configure Windows Firewall. This is an optional entry.

- **CURRENT** - Current profile (default).
- **DOMAIN** - Domain profile.
- **STANDARD** - Standard profile.
- **ALL** - All profiles.

**Remarks**
**scope** is ignored if **mode** is **DISABLE**. **scope** must be **CUSTOM** to specify **addresses**.

**Examples**
The following examples show how the **set service** command and preceding parameters can be used to enable or disable the predefined service exceptions:

**set service FILEANDPRINT**

**set service REMOTEADMIN ENABLE SUBNET**

**set service REMOTEDESKTOP ENABLE CUSTOM** *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**set service type** = **FILEANDPRINT**

**set service type** = **REMOTEADMIN mode** = **ENABLE scope** = **SUBNET**

**set service type** = **REMOTEDESKTOP mode** = **ENABLE scope** = **CUSTOM addresses** = *157.60.0.1,172.16.0.0/16,10.0.0.0/255.0.0.0*,LocalSubnet

**show (all commands)**
The following **show** commands are used to display the current configuration of Windows Firewall:

| Command | Description |
|---|---|
| **show allowedprogram** | Displays the programs that have been added to the exceptions list. |
| **show config** | Displays the local configuration information. |
| **show currentprofile** | Displays the current profile. |
| **show icmpsetting** | Displays the ICMP settings. |
| **show logging** | Displays the logging settings. |
| **show multicastbroadcastresponse** | Displays multicast or broadcast response settings. |
| **show notifications** | Displays the current settings for notifications. |
| **show opmode** | Displays the operational mode. |
| **show portopening** | Displays the ports that have been added to the exceptions list. |
| **show service** | Displays the services. |
| **show state** | Displays the current state information. |

**Remarks**
To use the **show** command parameter, [[**verbose** =] {**ENABLE**|**DISABLE**}], to display more detailed information about the configuration, set the parameter to **ENABLE**. Verbose mode is disabled by default.

**Examples**

**show state**

**show state verbose** = **ENABLE**

**Windows Firewall Log File and Security Log Settings**
Windows Firewall records event data in two locations: the Event Viewer security log and the Windows Firewall log file. Event Viewer security log entries provide a record of Windows Firewall configuration changes, startup status, and behavior. The Windows Firewall log file entries provide a record of network traffic events, such as dropped packets and successful connections. Neither of the logging mechanisms provides intrusion detection or security breach alerting.

**Security Log Entries**
Windows Firewall writes entries to the security log when a computer is started and when a program or system service attempts to listen for unsolicited incoming traffic but is blocked. These entries provide information about the status and configuration of Windows Firewall, including information about the applications and ports that permit traffic through Windows Firewall. These entries also provide information about which ports and protocols a program or system services is trying to use so you can configure the necessary exceptions in Windows Firewall. These security log entries are viewed with Event Viewer, which can filter the entries by Event IDs. The Event IDs associated with Windows Firewall are in the range of 848 through 861.

| **Note** |
|---|
| Windows Firewall events are written to the event log any time the Windows Firewall/Internet Connection Sharing service is running, even if Windows Firewall is turned off (disabled). |

The following table lists the Event IDs associated with Windows Firewall:

| Event ID | Description |
|---|---|
| 848 | Displays the startup configuration of Windows Firewall. |
| 849 | Displays an application exception configuration. |
| 850 | Displays a port exception configuration. |
| 851 | Displays a change made to the application exceptions list. |
| 852 | Displays a change made to the application exceptions list. |
| 853 | Displays a change made to the Windows Firewall operation mode. |
| 854 | Displays a change made to Windows Firewall logging settings. |
| 855 | Displays a change made to ICMP settings. |
| 856 | Displays a change made to the **Windows Firewall: Prohibit unicast response to multicast or broadcast requests** Group Policy setting. |
| 857 | Displays a change made to the remote administration setting. |
| 858 | Displays Windows Firewall Group Policy settings have been applied. |
| 859 | Displays Windows Firewall Group Policy settings have been removed. |
| 860 | Displays a change made to a profile. |
| 861 | Displays an application or service attempting to listen for incoming traffic. |

**Event ID 848**
The following table lists the entries associated with Event ID 848:

| Entry | Possible Values | Notes |
|---|---|---|
| **Group Policy applied** | Yes/No | Specifies whether Group Policy is applied. |

| | | |
|---|---|---|
| **Profile used** | Domain/Local | Specifies if the profile is from the domain or the local computer. |
| **Interface** | All interfaces/<interface name> | Specifies the network adapter (interface) to which the settings apply. |
| **Operational mode** | On/Off/On with no exceptions | Specifies which mode Windows Firewall is in. |
| **File and Printer Sharing** | Enabled/Disabled | Specifies whether File and Printer Sharing is enabled or disabled in the exceptions list. |
| **Remote Desktop** | Enabled/Disabled | Specifies whether Remote Desktop is enabled or disabled in the exceptions list. |
| **UPnP Framework** | Enabled/Disabled | Specifies whether UPnP is enabled or disabled in the exceptions list. |
| **Allow remote administration** | Enabled/Disabled | Specifies whether Remote Assistance is enabled or disabled in the exceptions list. |
| **Allow unicast responses to multicast/broadcast traffic** | Enabled/Disabled | Specifies whether Windows Firewall will allow unicast traffic that is in response to multicast or broadcast traffic through Windows Firewall. |
| **Log dropped packets** | Enabled/Disabled/"-" | Specifies the log file setting. |
| **Log successful connections** | Enabled/Disabled/"-" | Specifies the log file setting. |
| **Allow incoming echo request** | Enabled/Disabled/"-" | Specifies the ICMP setting. |
| **Allow incoming timestamp request** | Enabled/Disabled/"-" | Specifies the ICMP setting. |
| **Allow incoming mask request** | Enabled/Disabled/"-" | Specifies the ICMP setting. |
| **Allow incoming router request** | Enabled/Disabled/"-" | Specifies the ICMP setting. |
| **Allow outgoing destination unreachable** | Enabled/Disabled/"-" | Specifies the ICMP setting. |
| **Allow outgoing source quench** | Enabled/Disabled/"-" | Specifies the ICMP setting. |
| **Allow outgoing parameter problem** | Enabled/Disabled/"-" | Specifies the ICMP setting. |
| **Allow outgoing time exceeded** | Enabled/Disabled/"-" | Specifies the ICMP setting. |
| **Allow redirect** | Enabled/Disabled/"-" | Specifies the ICMP setting. |

| **Note** |
|---|
| The hyphen ("-") is used to indicate that this setting has not been configured. |

**Event ID 849**

The following table lists the entries associated with Event ID 849:

| Entry | Possible Values | Notes |
|---|---|---|
| **Policy origin** | Local Policy/Group Policy | Specifies whether the application was added to the exceptions list through local policies or domain-wide Group Policy. |

| | | |
|---|---|---|
| **Profile used** | Standard/Domain | Specifies the profile in which the application is listed as an exception. |
| **Name** | <file name> | Specifies the display name, if any, of the executable file. |
| **Path** | <path> | Specifies the path to the application. |
| **State** | Enabled/Disabled | Specifies whether the application is enabled or disabled in the exceptions list. |
| **Scope** | All interfaces/Local subnet/<custom scope> | Specifies the conditions under which the application is processed as an exception. |

**Event ID 850**

The following table lists the entries associated with Event ID 850:

| Entry | Possible Values | Notes |
|---|---|---|
| **Policy origin** | Local Policy/Group Policy | Specifies whether the port was added to the exceptions list through local policies or domain-wide Group Policy. |
| **Profile used** | Standard/Domain | Specifies the profile in which the port is listed as an exception. |
| **Interface** | All interfaces/<interface name> | Specifies the network adapter (interface) to which the settings apply. |
| **Name** | <name> | Specifies the name of the port. |
| **Port number** | <port number> | Specifies the number of the port. |
| **Protocol** | TCP/UDP | Specifies the protocol of the port. |
| **State** | Enabled/Disabled | Specifies whether the port is enabled or disabled in the exceptions list. |
| **Scope** | All interfaces/Local subnet/<custom scope> | Specifies the conditions under which the port is processed as an exception. |

**Event ID 851**

The following table lists the entries associated with Event ID 851:

| Entry | Possible Values | Notes |
|---|---|---|
| **Policy origin** | Local policy/Group Policy | Specifies whether the change was made through local policies or domain-wide Group Policy. |
| **Profile changed** | Standard/Domain | Specifies the profile in which the change occurred. |
| **Change type** | Add/Remove/Modify | Specifies whether the application was added or removed from the exceptions list, or whether exception list settings were modified for the application. |
| **New Settings: Name** | <name> | Specifies the new display name, if any, of the executable file. |
| **New Settings: Path** | <path> | Specifies the new path to the application. |
| **New Settings: State** | Enabled/Disabled | Specifies whether the application is currently enabled or disabled in the exceptions list. |

| | | |
|---|---|---|
| **New Settings: Scope** | All interfaces/Local subnet/<custom scope> | Specifies the new conditions under which the application is processed as an exception. |
| **Old Settings: Name** | <name> | Specifies the old display name, if any, of the executable file. |
| **Old Settings: Path** | <path> | Specifies the old path to the application. |
| **Old Settings: State** | Enabled/Disabled | Specifies whether the application was previously enabled or disabled in the exceptions list. |
| **Old Settings: Scope** | All interfaces/Local subnet/<custom scope> | Specifies the old conditions under which the application was processed as an exception. |

**Event ID 852**

The following table lists the entries associated with Event ID 852:

| Entry | Possible Values | Notes |
|---|---|---|
| **Policy origin** | Local policy/Group Policy | Specifies whether the change was made through local policies or domain-wide Group Policy. |
| **Profile changed** | Standard/Domain | Specifies the profile in which the change occurred. |
| **Change type** | Add/Remove/Modify | Specifies whether the port was added or removed from the exceptions list, or whether exception list settings were modified for the port. |
| **Interface** | All interfaces/<interface name> | Specifies the network adapter (interface) to which the settings apply. |
| **New Settings: Name** | <name> | Specifies the new name of the port. |
| **New Settings: Port number** | <port number> | Specifies the new number of the port. |
| **New Settings: Protocol** | TCP/UDP | Specifies the new protocol of the port. |
| **New Settings: State** | Enabled/Disabled | Specifies whether the port is enabled or disabled in the exceptions list. |
| **New Settings: Scope** | All interfaces/Local subnet/<custom scope> | Specifies the new conditions under which the port is processed as an exception. |
| **Old Settings: Name** | <name> | Specifies the old name of the port. |
| **Old Settings: Port number** | <port number> | Specifies the old number of the port. |
| **Old Settings: Protocol** | TCP/UDP | Specifies the old protocol of the port. |
| **Old Settings: State** | Enabled/Disabled | Specifies whether the port was previously enabled or disabled in the exceptions list. |

**Event ID 853**

The following table lists the entries associated with Event ID 853:

| Entry | Possible Values | Notes |
|---|---|---|
| **Policy origin** | Local policy/Group Policy | Specifies whether the change was made through local policies or domain-wide Group Policy. |
| **Profile changed** | Standard/Domain | Specifies the profile in which the change occurred. |
| **Interface** | All interfaces/interface name | Specifies the network adapter (interface) to which the new setting applies. |
| **New Setting: Operation mode** | On/On with no exceptions/Off | Specifies which mode Windows Firewall is in currently. |
| **Old Setting: Operation mode** | On/On with no exceptions/Off | Specifies which mode Windows Firewall was in previously. |

**Event ID 854**

The following table lists the entries associated with Event ID 854:

| Entry | Possible Values | Notes |
|---|---|---|
| **Policy origin** | Local policy/Group Policy | Specifies whether the change was made through local policies or domain-wide Group Policy. |
| **Profile changed** | Standard/Domain | Specifies the profile in which the change occurred. |
| **New Settings: Log dropped packets** | Enabled/Disabled/"-" | Specifies the new log file setting. |
| **New Settings: Log successful connections** | Enabled/Disabled/"-" | Specifies the new log file setting. |
| **Old Setting: Log dropped packets** | Enabled/Disabled/"-" | Specifies the old log file setting. |
| **Old Setting: Log successful connections** | Enabled/Disabled/"-" | Specifies the old log file setting. |

| **Note** |
|---|
| The hyphen ("-") is used to indicate that this setting has not been configured. |

**Event ID 855**

The following table lists the entries associated with Event ID 855:

| Entry | Possible Values | Notes |
|---|---|---|
| **Policy origin** | Local policy/Group Policy | Specifies whether the change was made through local policies or domain-wide Group Policy. |
| **Profile changed** | Standard/Domain | Specifies the profile in which the change occurred. |
| **Interface** | All interfaces/interface name | Specifies the network adapter (interface) to which the new setting applies. |
| **New Setting: Allow incoming echo request** | Enabled/Disabled/"-" | Specifies the new ICMP setting. |
| **New Setting: Allow incoming timestamp request** | Enabled/Disabled/"-" | Specifies the new ICMP setting. |

| | | |
|---|---|---|
| **New Setting: Allow incoming mask request** | Enabled/Disabled/"-" | Specifies the new ICMP setting. |
| **New Setting: Allow incoming router request** | Enabled/Disabled/"-" | Specifies the new ICMP setting. |
| **New Setting: Allow outgoing destination unreachable** | Enabled/Disabled/"-" | Specifies the new ICMP setting. |
| **New Setting: Allow outgoing source quench** | Enabled/Disabled/"-" | Specifies the new ICMP setting. |
| **New Setting: Allow outgoing parameter problem** | Enabled/Disabled/"-" | Specifies the new ICMP setting. |
| **New Setting: Allow outgoing time exceeded** | Enabled/Disabled/"-" | Specifies the new ICMP setting. |
| **New Setting: Allow redirect** | Enabled/Disabled/"-" | Specifies the new ICMP setting. |
| **Old Setting: Allow incoming echo request** | Enabled/Disabled/"-" | Specifies the old ICMP setting. |
| **Old Setting: Allow incoming timestamp request** | Enabled/Disabled/"-" | Specifies the old ICMP setting. |
| **Old Setting: Allow incoming mask request** | Enabled/Disabled/"-" | Specifies the old ICMP setting. |
| **Old Setting: Allow incoming router request** | Enabled/Disabled/"-" | Specifies the old ICMP setting. |
| **Old Setting: Allow outgoing destination unreachable** | Enabled/Disabled/"-" | Specifies the old ICMP setting. |
| **Old Setting: Allow outgoing source quench** | Enabled/Disabled/"-" | Specifies the old ICMP setting. |
| **Old Setting: Allow outgoing parameter problem** | Enabled/Disabled/"-" | Specifies the old ICMP setting. |
| **Old Setting: Allow outgoing time exceeded** | Enabled/Disabled/"-" | Specifies the old ICMP setting. |
| **Old Setting: Allow redirect** | Enabled/Disabled/"-" | Specifies the old ICMP setting. |

| **Note** |
|---|
| The hyphen ("-") is used to indicate that this setting has not been configured. |

**Event ID 856**

The following table lists the entries associated with Event ID 856:

| Entry | Possible Values | Notes |
|---|---|---|
| **Unicast response to multicast broadcasts is** | On/Off | Specifies whether the option is enabled or disabled. |

**Event ID 857**

The following table lists the entries associated with Event ID 857:

| Entry | Possible Values | Notes |
|---|---|---|

| | | |
|---|---|---|
| **Policy origin** | Local policy/Group policy | Specifies whether the change was made through local policies or domain-wide Group Policy. |
| **Profile changed** | Standard/Domain | Specifies the profile in which the change occurred. |
| **New setting: Allow remote administration** | Enabled/Disabled | Specifies whether the option is enabled or disabled. |
| **Old setting: Allow remote administration** | Enabled/Disabled | Specifies whether the option is enabled or disabled. |

**Event ID 858**

| **Note** |
|---|
| There are no entries associated with Event ID 858. |

**Event ID 859**

| **Note** |
|---|
| There are no entries associated with Event ID 859. |

**Event ID 860**

The following table lists the entries associated with Event ID 860:

| Entry | Possible Values | Notes |
|---|---|---|
| **Active profile** | Standard/Domain | Specifies if the profile has switched from standard to domain, or vice versa. |

**Event ID 861**

The following table lists the entries associated with Event ID 861:

| Entry | Possible Values | Notes |
|---|---|---|
| **Name** | <name> | Specifies the name of the application or service. |
| **Path** | <path> | Specifies the path to the application or service. |
| **Process identifier** | <process identifier> | Specifies the identifier (label) used to determine the process involved in the listening activity. |
| **User account** | <user account> | Specifies the user's account. |
| **User domain** | <user domain> | Specifies the user's domain. |
| **Service** | Yes/No | Specifies whether the listener is a service. |
| **RPC server** | Yes/No | Specifies whether the listener is on an RPC server. |
| **IP version** | IPv4/IPv6 | Specifies the IP version of the application or service. |
| **IP protocol** | TCP/UDP | Specifies the IP protocol being used by the application or service. |
| **Port number** | <port number> | Specifies the port being used by the application or service. |
| **Allowed** | Yes/No | Specifies if the application or service is allowed to listen for unsolicited incoming network traffic. |
| **User notified** | Yes/No | Specifies whether the user was notified of the listening activity. |

**Windows Firewall Log File Entries**

When the Windows Firewall log file is enabled, Windows Firewall generates a plaintext security log file (Pfirewall.log), which is found in %Windir%\pfirewall.log. The security log has two sections: the header and the body.

The following table lists the entries contained in the header:

| Item | Description | Example |
|---|---|---|
| #Version: | Displays which version of the Windows Firewall security log is installed. | 1.5 |
| #Software: | Displays the name of the security log. | Microsoft Windows Firewall |
| #Time: | Indicates that all of the timestamps in the log are in local time. | Local |
| #Fields: | Displays a static list of fields that are available for security log entries, if data is available. These fields are listed in the following table.<br><br>**Note** The hyphen (-) is used for fields for which no information is available. | src-ip |

The body is the report of information gathered about traffic or attempts to cross Windows Firewall. The body of the security log is a dynamic list; new entries appear at the bottom of the log.

The following table lists the information contained in the body of the log file:

| Item | Description | Example |
|---|---|---|
| Date | Displays the year, month, and day that the recorded transaction occurred. Dates are recorded in the format:<br><br>YYYY-MM-DD | 2001-01-27 |
| Time | Displays the hour, minute, and seconds at which the recorded transaction occurred. Times are recorded in the format:<br><br>HH:MM:SS | 21:36:59 |
| Action | Displays which operation was observed by Windows Firewall. The options available are OPEN, OPEN-INBOUND, CLOSE, DROP, and INFO-EVENTS-LOST. An INFO-EVENTS-LOST action indicates the number of events that occurred but were not recorded in the log. | OPEN |
| Protocol | Displays the protocol that was used for the communication. The options available are TCP, UDP, ICMP, and a protocol number for packets that are not TCP, UDP, or ICMP. | TCP |
| src-ip | Displays the source IP address (the IP address of the computer attempting to establish communications). | 192.168.0.1 |
| dst-ip | Displays the destination IP address of a communication attempt. | 192.168.0.1 |
| src-port | Displays the source port number of the sending computer. Only TCP and UDP display a valid src-port entry. All other protocols display a src-port entry of -. | 4039 |
| dst-port | Displays the port number of the destination computer. Only TCP and UDP display a valid dst-port entry. All other protocols display a dst-port entry of -. | 53 |
| size | Displays the packet size, in bytes. | 60 |

| | | |
|---|---|---|
| tcpflags | Displays the TCP control flags found in the TCP header of an IP packet:<br>• **A**ck Acknowledgment field significant<br>• **F**in No more data from sender<br>• **P**sh Push function<br>• **R**st Reset the connection<br>• **S**yn Synchronize sequence numbers<br><br>**U**rg Urgent Pointer field significant | FAP |
| tcpsyn | Displays the TCP sequence number in the packet. | 1315819770 |
| tcpack | Displays the TCP acknowledgement number in the packet. | 2515999782 |
| tcpwin | Displays the TCP window size, in bytes, in the packet. | 64240 |
| icmptype | Displays a number that represents the Type field of the ICMP message. | 8 |
| icmpcode | Displays a number that represents the Code field of the ICMP message. | 0 |
| info | Displays an information entry that depends on the type of action that occurred. For example, an INFO-EVENTS-LOST action will result in an entry of the number of events that occurred but were not recorded in the log from the time of the last occurrence of this event type. | 23 |
| Path | Displays the direction of the communication. The options available are SEND, RECEIVE, FORWARD, and UNKNOWN. | RECEIVE |

# Storage Technologies Collection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Storage Technologies Collection**
Much of the data stored on Microsoft Windows-based servers is critically important for the day-to-day activities in business organizations. Because it is important that employees are able to access and store data reliably, data storage and recovery are a concern for many IT managers. In response to the expanding needs for storage in distributed computing environments, Windows Server 2003 includes many storage technologies that are designed to help you store, access, and manage data on servers.

Storage technologies include the following components:

- File Services Technologies
- File Systems Technologies
- Disks and Volumes
- Storage Services Technologies

**File Services Technologies**
File services are the underlying technologies that enable file servers to share data within an organization. These essential services allow users to store and share data on servers across the network.

Windows Server 2003 offers a number of file server solutions, such as disk quotas, Shadow Copies for Shared Folders, Distributed File System (DFS), and File Replication service (FRS), for enhancing the manageability, scalability, and availability of file servers.

**File Systems Technologies**
A file system is the structure in which files are named, stored, and organized. Windows Server 2003 supports the following file systems: FAT16, FAT32, and NTFS. You can use any combination of these file systems on a hard disk, but each volume on a hard disk can use only one file system. By default, Windows Server 2003 uses the NTFS file system.

NTFS provides performance, reliability, and advanced features not found in any version of FAT. For example, NTFS includes built-in security features such as file and folder permissions and the Encrypting File System (EFS), which is the technology used to store encrypted files on NTFS volumes. NTFS also provides support for volumes up to 256 terabytes in size, support for disk quotas and compression, and support for mounted drives.

**Disks and Volumes**
Windows Server 2003 offers two types of disks for storing information about your server: basic and dynamic. Basic disks use the same disk structures as those used by the Microsoft MS-DOS operating system and all previous versions of Microsoft Windows. Dynamic disks were introduced in Windows 2000 and they provide features that basic disks do not, such as the ability to create volumes that span multiple disks, including fault-tolerant mirrored and RAID-5 volumes.

**Storage Services Technologies**
In response to the expanding needs for storage in distributed computing environments, Windows Server 2003 includes several storage technologies that are designed to help you store, access, and manage data on servers.

**Storage Technologies Scenarios**
This section describes some of the most common scenarios for using storage technologies.

**File Services Scenarios**
The following sections describe the scenarios in which organizations deploy file services in Windows Server 2003.

**DFS**
DFS is used in organizations that want to achieve the following goals:

- Provide an intuitive way for users to access multiple file servers throughout the organization.

- Make data on multiple file servers appear as though it were available on a single file server.
- Make data available in multiple sites so that users in each site use fast, inexpensive bandwidth to access the data.
- Reduce delays that occur when users access heavily used shared folders.
- Provide fault-tolerant access to shared folders.
- Consolidate file servers or migrate data without affecting how users locate data.

**Disk Quotas**
Disk quotas are used in organizations that want to achieve the following goals:

- Limit the amount of data that users can store on a particular volume.
- Monitor how much disk space each user is using.

**FRS**
FRS is used in organizations that want to achieve the following goals:

- Ensure data availability if a file server fails or is taken offline for maintenance.
- Make data available in multiple sites to provide inexpensive access to users within each site.

**Shadow Copies for Shared Folders**
Shadow Copies for Shared Folders is used in organizations that want to achieve the following goals:

- Provide access to previous versions of files on the file server.
- Reduce the cost associated with restoring files for users.

**File Systems Scenarios**
**FAT or FAT32 file systems**
The FAT and FAT32 file systems are used in organizations that want to achieve the following goals:

- Provide compatibility with older operating systems, such as Windows 98.
- Provide a way to switch between different versions of the Windows operating system on the same computer by using a dual-boot configuration.

**NTFS file system**
The NTFS file system is used in organizations that want to achieve the following goals:

- Keep files secure. The NTFS file system is more secure than either FAT or FAT32.
- Use the advanced features provided by NTFS, such as disk quotas, file and folder permissions, encryption, large volume support, and sparse file management.

**Disk and Volume Scenarios**
The following sections describe the scenarios in which organizations use basic and dynamic disks and volumes in Windows Server 2003.

**Basic disks and volumes**
Basic disks and volumes are used in organizations that want to achieve the following goals:

- Provide simple data storage.

- Provide separate volumes for operating system and business data so that when a new version of the operating system is released, the boot or system volume can be reformatted and a new operating system installed, leaving the business data, located on the second volume, untouched.

- Provide separate volumes that are individually secured to limit access to specific, authorized users.

**Dynamic disks and volumes**
Dynamic disks and volumes are used in organizations that want to achieve the following goals:

- Increase the size of an existing disk by extending it onto a second disk.

- Increase the speed at which data is read and written.

- Provide fault-tolerance for business critical data that will be read more often than it is written by creating a RAID-5 volume.

- Provide good write performance with fault-tolerance for business critical data by creating a mirrored volume.

- Provide a way to easily move data from a smaller disk to a larger disk in the same computer by using a mirrored volume.

**Storage Services Scenarios**
The following sections describe the scenarios in which organizations use storage services technologies in Windows Server 2003.

**Virtual Disk Service Scenarios**
The Virtual Disk Service is used in organizations that want to achieve the following goals:

- Manage LUNs on diverse hardware storage devices.

- Manage disks and volumes.

- Manage end-to-end storage operations.

**Removable Storage Scenarios**
Removable Storage is used in organizations that want to achieve the following goals:

- Manage stand-alone drive libraries.

- Manage automated libraries.

**Remote Storage Scenarios**
Remote Storage is used in organizations that want to achieve the following goals:

- Conserve disk space on managed volumes.

- Extend disk space on managed volumes.

- Generate multiple media copies of removable media.

- Replace damaged removable media while Remote Storage is running.

- Recover from loss of Remote Storage metadata.

# File Services Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

File services are the underlying technologies that enable file servers to share data within an organization. File servers running basic services, such as folder sharing and permissions, are easy to set up, but many organizations require a higher level of manageability, scalability, and availability for their file servers. The Microsoft Windows Server 2003 operating system provides files services such as Distributed File System (DFS), disk quotas, File Replication Service (FRS), and Shadow Copies for Shared Folders to help organizations meet these needs.

For information about File and Storage Services in Windows Server 2012, see File and Storage Services Overview. For information about File Services in Windows Server 2008 R2, see File Services for Windows Server 2008 R2.

**File Services Components**
File Services technologies include the following components.

**DFS**
A DFS namespace links together shared folders on different servers to create a hierarchical structure that behaves like a single high-capacity hard disk. Users can navigate the namespace without having to know the physical server names or shared folders that host the data. DFS also provides increased availability, storage scalability, load sharing, and simplified maintenance.

**Disk Quotas**
Windows Server 2003 provides disk quota functionality that tracks quotas on a per-user, per-volume basis. After you enable the warning level and limit, they apply to all users who own files stored on the volume. Any user who creates a new file is automatically assigned the current warning level and limit.

**FRS**
FRS is a multimaster replication service used to replicate files and folders in the SYSVOL shared folder on domain controllers and in DFS shared folders. FRS works by detecting changes to file and folders in a replica set, and then replicating those changes to other replica members, which are connected in a replication topology. Because FRS is a multimaster replication service, any member of a replica set can generate changes. In addition, FRS can resolve file and folder conflicts to make data consistent among the replica members.

**Shadow Copies for Shared Folders**
Shadow Copies for Shared Folders provides point-in-time copies of files that are located on file servers. With Shadow Copies for Shared Folders, users can quickly recover deleted or changed files stored on the network without administrator assistance, increasing productivity and reducing administrative costs.

**File Services Scenarios**
The following sections describe the scenarios in which organizations deploy file services in Windows Server 2003.

**DFS**
DFS is used in organizations that want to achieve the following goals:

- Provide an intuitive way for users to access multiple file servers throughout the organization.
- Make data on multiple file servers appear as though it were available on a single file server.
- Make data available in multiple sites so that users in each site use fast, inexpensive bandwidth to access the data.
- Reduce delays that occur when users access heavily used shared folders.
- Provide fault-tolerant access to shared folders.
- Consolidate file servers or migrate data without affecting how users locate data.

**Disk Quotas**
Disk quotas are used in organizations that want to achieve the following goals:

- Limit the amount of data that users can store on a particular volume.
- Monitor how much disk space each user is using.

**FRS**

FRS is used in organizations that want to achieve the following goals:

- Ensure data availability if a file server fails or is taken offline for maintenance.
- Make data available in multiple sites to provide inexpensive access to users within each site.

**Shadow Copies for Shared Folders**

Shadow Copies for Shared Folders is used in organizations that want to achieve the following goals:

- Provide access to previous versions of files on the file server.
- Reduce the cost associated with restoring files for users.

# Disk Quotas Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk Quotas Technical Reference**
Disk quotas track and control disk space usage for NTFS volumes, allowing administrators to control the amount of data that each user can store on a specific NTFS volume. By using disk quotas, you can configure Microsoft Windows Server 2003 to do the following:

- Log an event when a user exceeds a specified disk space warning level. The warning level specifies the point at which a user is nearing the quota limit.
- Prevent further use of disk space or log an event when a user exceeds a specified disk space limit.

Disk quotas are not available for FAT volumes.

**In this subject**


- What Are Disk Quotas?
- How Disk Quotas Work
- Disk Quotas Tools and Settings

# What Are Disk Quotas?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Common Core Scenarios for Using Disk Quotas
- Dependencies on Other Technologies
- Related Information


Disk quotas track and control disk space usage for NTFS volumes, allowing administrators to control the amount of data that each user can store on a specific NTFS volume. By using disk quotas, you can configure Microsoft Windows Server 2003 to do the following:

- Log an event when a user exceeds a specified disk space warning level. The warning level specifies the point at which a user is nearing the quota limit.
- Prevent further use of disk space or log an event when a user exceeds a specified disk space limit.

If users are sharing file disk space on a server, they can consume considerable disk space when saving documents, graphics, and other files. This can be problematic if some users need to use more disk space than other users, resulting in some users not having enough disk space to store their files on the volume.

An administrator can solve this problem by enacting a Group Policy or a computer policy that limits the amount of disk space each user has on the volume. The administrator can also adjust quotas to the user's needs because user's are tracked on a per-user basis, per-volume.

**Common Core Scenarios for Using Disk Quotas**
Disk quotas can be used in several ways, depending on the needs and resources of your organization.

**Limited number of servers**
When you have a limited number of file servers and applications installed on those servers, you can use disk quotas to track disk usage. If you have a large number of users accessing a small group of file servers, you can set disk quotas to allocate the amount of space each user is allowed on a particular server. The disk quotas apply to all users whether they are local or remote users of the server.

**Limited space on a server volume**
When you have limited space on a server volume with multiple local and remote users accessing applications on that volume, you can set a specific amount of hard disk space for each user by enabling disk quotas.

**Monitoring but not limiting disk space usage**
If you have plenty of disk space on a volume, as well as enough servers in your organization, you might not want to enforce disk quotas. However, by enabling disk quotas, you can monitor the disk space usage of each user without restricting any user from accessing additional space when a limit is reached. You can then adjust the quota according to the demands of the job. Later, you can increase or decrease allocated space as necessary.

**Dependencies on Other Technologies**
Disk quotas rely on the following technologies:

**NTFS**
The server volume must be formatted as NTFS. Disk quotas do not work on FAT volumes.

**Active Directory**
The Active Directory directory Service is required if you want to set a Group Policy for a computer or for a group of users. You must be the system administrator in the domain in which the server resides. Disk Quotas obtains a list of the user account names from the domain controller or from Local Users and Groups in the Computer Management console.

**Group Policy**

Using Group Policy you can create and apply a disk quota policy setting to a set of targeted users and computers in an Active Directory environment.

The following figure, "Relationship of Disk Quotas to Other Windows Server 2003 Technologies," shows the relationship of disk quotas to other technologies in Windows Server 2003.

**Relationship of Disk Quotas to Other Windows Server 2003 Technologies**




In the preceding diagram, "Relationship of Disk Quotas to Other Windows Server 2003 Technologies," an administrator accesses an NTFS volume on a server running Windows Server 2003 to set disk quotas. When the administrator uses Disk Management to set disk quotas, Active Directory returns a list of users in the domain. Based on the names in the list, the administrator can create a disk quota Group Policy or computer policy for each user on that particular NTFS volume, regardless of where the user is located.

**Related Information**

The following resources contain additional information that is relevant to this section.

- NTFS Technical Reference
- Active Directory Collection
- Core Group Policy Technical Reference

# How Disk Quotas Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk Quota Architecture**

The following figure, "Disk Quota Architecture," shows the architecture of disk quotas on an NTFS volume.

**Disk Quota Architecture**




The following table, "Components of Disk Quota Architecture," describes the components of the preceding figure, "Disk Quota Architecture."

**Components of Disk Quota Architecture**

| Architecture Component | Description |
|---|---|
| Active Directory Domain | Domain where the NTFS volume must reside. Disk quotas are based on all users within the domain. |
| Administrator Interface | Volume Quota Property Page user interface where administrators set Group Policy for disk quotas on an NTFS volume. |
| User Policy | Applies disk quota Group Policy at the computer or user level for the NTFS volume. Group Policy is determined by the users and groups in the Active Directory domain account. |
| Dskquota.dll | Implements the disk quota Group Policy on the client side. The dll is loaded when the client computer processes the policy. |
| Dskquoui.dll | Implements the user interface associated with managing disk quotas. |
| System Log | Records an audit entry whenever disk quota events occur, such as a user approaching a limit. |
| Quota Settings | Disk quota settings enabled by the administrator. |
| Volume Quota Property Page | User interface page allowing administrators to set and manage disk quotas. |
| Ntioapi_h | Windows Server 2003 header file identifying NT systems. The header includes functions for setting disk quotas and querying volumes for disk quotas. |
| NtQueryQuotaInformationFile | Function declared in ntioapi_h. Sets information for users on the volume. |

| | |
|---|---|
| NtSetQuotaInformationFile | Function declared in ntioapi_h. Queries quota information for users on the volume. |
| NtQueryVolumeInformationFile | Function declared in ntioapi_h. Queries the volume quota information to determine if quotas are enabled. |
| NtSetVolumeInformationFile | Function declared in ntioapi_h. Identifies the volume ID. Sets volume quota information. |
| Registry | System-defined database in which information and configuration data for disk quotas is stored. |
| Security Identifier | A variable-length number that uniquely identifies a user in a domain. |
| User File | File owned by a user saving the file to the local NTFS volume. |

When a user saves a file to a share on an NTFS volume, dskquota.dll accesses Group Policy to determine what, if any, disk quotas are applicable. Group Policy checks the user's name from a list of user names in the domain to determine if that user is a member of the domain.

A user who is an administrator and wants to set disk quotas can do so by using Disk Management in the Computer Management console. When this is done, dskquota.dll recognizes that the administrator is logged on. Dskquota.dll calls dskquoui.dll and opens the Volume Quota Property Page, allowing the administrator to create and administer disk quotas. A list of user names and associated security identifiers (SIDs) are provided to the administrator from the domain controller. After a disk quota policy is created, modified, and saved, it is written to the registry. Policy data is applied to the user of the NTFS volume and to a user's SID. Each file a user saves to a volume has an associated SID that identifies that file with the user. The SID in the file is then compared to the SID of the user from the Group Policy.

If the user is a member of the domain, Group Policy determines if disk quotas are enabled. If the user is approaching the disk quota limit, an event is written to the System log. If enforcement is enabled, the administrator can access the System log to determine what follow-up actions to take, such as preventing the user from accessing more disk space or allocating additional disk space to the user. A user who has not reached the warning level or disk quota limit can save the file to the local NTFS volume.

#### Disk Quota Interfaces

Disk quota relies on several disk management interfaces to correctly process disk quota control and logging of events. The following table, "Disk Quota Management Interfaces," lists the interfaces used in setting and managing disk quotas on an NTFS volume.

**Disk Quota Management Interfaces**

| Disk Management Interface | Description |
|---|---|
| IDiskQuotaControl | Controls disk quota facilities of an NTFS volume. |
| IDiskQuotaEvents | Calls by the system whenever significant quota events occur. |
| IDiskQuotaUser | Represents single user quota entry in the volume quota information file. |
| IDiskQuotaBatch | Enables the administrator to add multiple quota user objects to a container. |
| IEnumDiskQuotaUsers | Enumerates user quota entries on the volume. |
| IUnknown | Provides pointers to other interfaces in any given object. |
| WMI | Windows Management Instrumentation. Enables identification and manipulation of objects. |

**Disk Quota Interface**




When the Component Object Model (COM) accesses disk quotas, the IUnknown pointer is invoked. IUnknown has no functions for disk quota operations, so it calls QueryInterface to ask an object for an interface by which the disk quota operations can be invoked. WMI is then used to access the five-disk quota interfaces.

**Disk Quota Physical Structure**

The following figure, "Disk Quota Physical Structure," shows the physical structure of the environment in which disk quotas are used.

**Disk Quota Physical Structure**




**Disk Quota Physical Structure Components**

| Component | Description |
|---|---|
| Domain Controller | A server in an Active Directory domain running Windows Server 2003 that manages security for a domain. The domain controller must be in the same domain as the domain users. |
| Domain Users | All users in the domain who have access to the NTFS volume. Users must be in the same domain as the domain controller and the NTFS volume, and can be local or remote. |
| Disk Quota Policy | Applies disk quota policy at the computer or user level for the NTFS volume. The policy is determined for the users and groups in the Active Directory domain account. |
| NTFS Volume | NTFS-formatted volume. The volume can be a separate server from the domain controller, or be a partition on the domain controller. |
| User Accounts | A list of users in the domain, provided by the domain controller |

In an Active Directory domain, when an NTFS volume is accessed by a domain user, Group Policy applies any disk quotas to

those users if the administrator has enabled disk quotas. If a user is saving a file to the NTFS volume, Group Policy determines if disk quotas are enabled for the user. Disk quotas check the list of names against the names provided by the domain controller of the domain to determine if disk quotas are applicable to the user. If they are, disk quotas can do any of the following:

- Allow the user to save the file to the NTFS volume
- Prevent a user who has reached the quota from writing to the volume
- Allow a user who has reached the quota to write to the volume.
- Warn the user that the disk quota limit has been reached.

**Disk Quota Processes and Interactions**

Disk quotas are depend on NTFS. The information in the following sections discuss how Active Directory domains are used to provide information to set and store disk quotas, as well as how NTFS tracks disk quotas, and how disk quotas can be transferred from one volume to another.

Disk quotas are tracked on a per-user, per-volume basis; users are charged only for the files they own. Quotas are tracked per volume, even if the hard disk contains multiple volumes. If multiple shared folders are on the same volume, the quotas apply to all shared folders equally, and a user's total use of all shared folders cannot exceed the assigned quota on that volume.

**Setting Disk Quotas**

The Disk Management snap-in in the Computer Management console allows an administrator to set a disk quota Group Policy. Dskquota.dll recognizes an administrator who logs on and opens the Computer Management console. The administrator is provided with a list of names from the domain controller. These names are then matched to the security identifier (SID) for current volume users, and the administrator is allowed to enable disk quotas. The administrator can set a policy that controls the use of a computer, domain, or site to enable quota management on individual volumes. The policy is applied to local or remote users as long as the users are located within the same domain as the NTFS volume. At the same time, the administrator can register an event for when a disk quota is reached by a user, such as simply warning the user that a limit has been reached, or preventing the user from saving additional files to the volume.

Enforcement of the user or computer policy is optional. When users reach an enforced disk quota limit, the system responds as though the physical space on the volume were exhausted. When users reach an unenforced limit they can continue to write to the volume as long as physical space is available, but an event is logged to the System log.

**Storing Disk Quotas on a Volume**

NTFS volumes contain a master file table (MFT), which stores information about all of the files on the volume. Within the MFT are metadata files of attributes of files on the volume.

One of the file attributes of the MFT is $Extend metadata record, one function of which is storing disk quota information for a file. The $Extend attribute is further defined by $Quota. When disk quotas are enabled, NTFS uses a Quota bit in the Standard Information file record attribute, and copies the file size to the Standard Information attribute. Because the attribute is considered a resident attribute, an attribute which can fit within the MFT file record for that file, the MFTcan quickly scan the record to determine the quota assigned to a user.

Disk quotas are stored in the $Quota metadata record. $Quota uses two indexes, $O and $Q. The first index, $O, indexes the owner of a file by the owner's SID. The second index, $Q, keeps track of the default quotas, thresholds, and quota attributes. When a user is assigned a quota or adds files to a volume with disk quotas enabled, the user's SID is added to the $O index then assigned an entry in the $Q index specifying the quota limit.

**Tracking File Sizes**

Disk quotas track the file size against the allocated disk file using the SID to look up the user's disk quota entry allowance.

SIDs are stored in a user's profile within the NTFS structure on the volume. The user's name is obtained from a list of names in the domain in which the volume resides. NTFS tracks the user's volume usage by tagging files and directories with the SID. When a user saves a file to the NTFS volume with disk quotas enabled, the SID in the file is resolved against the name in a cache file of stored SID names and quota information on the volume. When the file or directory is saved to the volume, the size of files and directories a user owns count against the set quota limit.

**Transferring Disk Quota Settings Between Volumes**

Transferring disk quotas between NTFS volumes saves an administrator from having to create identical disk quotas on multiple volumes.

When disk quotas are exported from one volume and imported on into another, disk quotas creates a transfer file, which is imported into another volume. The following figure, "Transferring Settings Between Volumes,"shows the information included in a Quota Settings Transfer file.

**Transferring Settings Between Volumes**




In the preceding figure,, the GUID contains a record for each user for whom the disk quotas policy applies. The SID is a unique structure that identifies the users or groups. cbSID identifies the size of the SID buffer. These settings are applied to an export file, and applied to a different volume when imported into that volume.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Core Group Policy Technical Reference
- NTFS Technical Reference

# Disk Quotas Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Disk Quota Tools
- Disk Quota Registry Entries
- Disk Quota Group Policy Settings
- Disk Quota WMI Classes
- Related Information


The following tools, registry settings, Group Policy settings, and Windows Management Instrumentation (WMI) classes are associated disk quotas on NTFS volumes.

**Disk Quota Tools**
Disk Quota Tools describes the tools installed with Microsoft Windows Server 2003 as well as downloadable support tools. The tools in this list are those used to create and maintain disk quotas on NTFS volumes.

**Dskmgmt.msc: Disk Management**
**Category**
Operating system tool.

**Version compatibility**
Runs on Windows Server 2003.

Use the Disk Management snap-in for MMC to set disk quota limits on NTFS-formatted volumes for Windows Server 2003.

**Dskuse.exe: User Disk Tool**
**Category**
Downloadable support tool.

**Version compatibility**
Runs on Windows Server 2003. Works on NTFS volumes only.

Scans a single directory, a directory tree, or an entire drive and reports the amount of space used by each user or all users.

To find more information about Dskuse, see "Resource Kit Tools" in Tools and Settings Collection

**Fsutil.exe: File Systems Utilities**
**Category**
Operating System tool.

**Version compatibility**
Windows Server 2003. Requires NTFS as the file system on the volume.

Perform file system tasks such as disabling long file names, checking whether a volume is flagged as dirty, viewing NTFS-related information about a volume, creating hard links, and managing quotas and sparse files.

To find more information about Fsutil, see "Command Line References" in Tools and Settings Collection.

**Disk Quotas Registry Entries**
The registry entries listed in this section are associated with file systems.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft

Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**DiskQuota**
The file registry entries for DiskQuota can be found under HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota.

**ApplyToRemovableMedia**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota

**Version**
Microsoft Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Extends the disk quota policies to NTFS file system volumes on removable media.

To find more information about this registry entry, see "Registry Reference" in Tools and Settings Collection.

**Enable**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota

**Version**
Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Enables and disables disk quota management on all NTFS file system volumes of the computer, and prevents users from changing disk quota settings.

To find more information about this registry entry, see "Registry Reference" in Tools and Settings Collection.

**Enforce**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota

**Version**
Windows NT 4.0 Server and Windows NTNT 4.0 Workstation and later.

Determines whether disk quota limits are enforced, and prevents users from changing the setting.

To find more information about this registry entry, see "Registry Reference" in Tools and Settings Collection.

**Limit**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota

**Version**
Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Specifies the number of units in the default disk quota limit for new users of an NTFS file system volume.

To find more information about this registry entry, see "Registry Reference" in Tools and Settings Collection.

**LimitUnits**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota

**Version**
Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Specifies the units used in the default disk quota limit for new users of an NTFS volume.

To find more information about this registry entry, see "Registry Reference" in Tools and Settings Collection.

**LogEventOverLimit**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota

**Version**
Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Determines whether the system records an event in the local Application Log when users reach the disk quota limit on a volume, and prevents users from changing the logging setting.

To find more information about *registry setting*, see "Registry Reference" in Tools and Settings Collection.

**LogEventOverThreshold**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota

**Version**
Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Determines whether the system records an event in the local Application Log when users reach the disk quota warning level, and prevents users from changing the logging setting.

To find more information about this registry entry, see "Registry Reference" in Tools and Settings Collection.

**Threshold**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota

**Version**
Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Specifies the number of units in the default disk quota warning level for new users of an NTFS volume.

To find more information about this registry entry, see "Registry Reference" in Tools and Settings Collection.

**ThresholdUnits**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows NT\DiskQuota

**Version**
Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Specifies the units for the default disk quota warning level for new users of an NTFS volume.

To find more information about this registry entry, see "Registry Reference" in Tools and Settings Collection.

**Disk Quotas Group Policy Settings**
The following table, "Group Policy Settings Associated with Disk Quotas," lists and describes the Group Policy settings that are associated with disk quotas.

**Group Policy Settings Associated with Disk Quotas**

| Group Policy Setting | Description |
|---|---|
| DiskQuotaEnable | Enables disk quotas on an NTFS volume. This policy is not used on FAT volumes. |
| DiskQuotaLimit | Specifies the disk quota limit and warning level. This policy is not used on FAT volumes. |
| DiskQuotaLogEventOverThreshold | Specifies whether an application log event is created when a user exceeds disk quota limit. This policy is not used on FAT volumes. |
| DiskQuotaApplyToRemovableMedia | Specifies whether disk quotas policies are applied to removable media. This policy is not used on FAT volumes. |

To find more information about Group Policy settings, see "Group Policy Settings Reference" in Tools and Settings Collection.

**Disk Quota WMI Classes**
The following table, "WMI Classes Associated with Disk Quotas," lists and describes the WMI classes that are associated with disk quotas.

**WMI Classes Associated with Disk Quotas**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_DiskQuota | //root/cimv2 | Windows XP and Windows Server 2003 |

| | | |
|---|---|---|
| Win32_QuotaSetting | //root/cimv2 | Windows XP and Windows Server 2003 |
| Win32_VolumeQuota | //root/cimv2 | Windows Server 2003 |
| Win32_VolumeQuotaSetting | //root/cimv2 | Windows XP and Windows Server 2003 |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Related Information**

The following resources contain additional information that is relevant to this section.

- For more information about NTFS, see the NTFS Technical Reference.

# DFS Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DFS Technical Reference**
Distributed File System (DFS) allows administrators to group shared folders located on different servers by transparently connecting them to one or more DFS namespaces. A DFS namespace is a virtual view of shared folders in an organization. Using the DFS tools, an administrator selects which shared folders to present in the namespace, designs the hierarchy in which those folders appear, and determines the names that the shared folders show in the namespace. When a user views the namespace, the folders appear to reside on a single, high-capacity hard disk. Users can navigate the folders in the namespace without needing to know the server names or shared folders hosting the data. DFS also provides many other benefits, including fault tolerance and load-sharing capabilities, making it ideal for all types of organizations.

**In this subject**


- What Is DFS?
- How DFS Works
- DFS Tools and Settings

# What Is DFS?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Common DFS Scenarios
- Technologies Related to DFS
- DFS Dependencies
- Related Information


One of the goals of most information technology (IT) groups is to manage file server resources efficiently while keeping them available and secure for users. As networks expand to include more users and servers—whether they are located in one site or in geographically distributed sites—administrators find it increasingly difficult to keep users connected to the files they need. On one hand, distributing resources across a network makes them more available to more people and promotes cross-organizational efforts. On the other hand, storing files on different file servers located throughout an organization makes it difficult for users to know where to look for information. Administrators also find it difficult to keep track of all the servers and all of the people who use those servers. The task of swapping out an old server becomes a major communication chore when users across an organization must be notified to update links and file paths.

To help administrators address these problems, Windows Server 2003 includes Distributed File System (DFS). DFS allows administrators to group shared folders located on different servers by transparently connecting them to one or more DFS namespaces. A DFS namespace is a virtual view of shared folders in an organization. Using the DFS tools, an administrator selects which shared folders to present in the namespace, designs the hierarchy in which those folders appear, and determines the names that the shared folders show in the namespace. When a user views the namespace, the folders appear to reside on a single, high-capacity hard disk. Users can navigate the namespace without needing to know the server names or shared folders hosting the data. DFS also provides other benefits, including the following:

**Simplified data migration**
DFS simplifies the process of moving data from one file server to another. Because users do not need to know the name of each physical server or shared folder that contains the data, administrators can physically move data to another server without needing to reconfigure applications and shortcuts and without needing to reeducate users about where they can find their data. This minimizes the impact of server consolidation on users. It also allows administrators to deploy additional file servers and present the folders on those new servers as new folders within an existing namespace.

**Increased availability of file server data**
In the event of a server failure, DFS refers client computers to the next available server, so users can always access shared folders without interruption.

**Load sharing**
DFS provides a degree of load sharing by mapping a given logical name to shared folders on multiple file servers. For example, suppose that \\Company\StockInfo is a heavily used shared folder. Administrators can use DFS to associate this location with multiple shared folders on different servers, even if the servers are located in different sites.

**Security integration**
Administrators do not need to configure additional security for DFS namespaces because file and folder security is enforced by existing the NTFS file system and shared folder permissions on each target. For example, a user navigating a DFS namespace is permitted to access only the files or folders for which he or she has appropriate NTFS or shared folder permissions.

**Common DFS Scenarios**
DFS is commonly used in the following scenarios:

**Server Consolidation**
Many organizations today are consolidating older file servers throughout the organization into fewer, larger, more powerful file servers. Consolidation reduces the cost of managing multiple file servers and increases the efficiency of storage allocation and backup tasks. Organizations that have implemented DFS can perform server consolidations without impacting the way users' access data. The following figure illustrates this benefit.

### How DFS Eliminates the Impact of Server Consolidation on Users




**Publishing Applications**

DFS is commonly used to publish applications to users throughout the organization. Using DFS in this scenario provides a number of benefits, such as the ability to use multiple servers to host application data and distribute the load across servers. A feature in DFS known as "least expensive target selection" ensures that users are connected to the closest server. The following figure illustrates a DFS namespace used to publish applications in an organization based in San Francisco with offices in Miami and Dallas.

### Using DFS to Publish Applications




This organization has three types of software:

- Business-critical software and operating systems that must be available at all times.
- Previous versions of software that are still in use in the Dallas branch office.
- Multimedia software used primarily in San Francisco.

The organization uses four servers to host the business-critical software and operating systems, including two servers in the San Francisco site. Using two servers to host the applications ensures that a failure on one server does not cause the data to become unavailable. All users can access this software at \\Software\Public\Applications, and users are automatically directed

to the server in their site (San Francisco, Dallas, or Miami).

Because the archived software is used only in the Dallas office and the data is not business-critical, only a single server hosts that data. The multimedia software is not business-critical, but the organization uses two servers for this software to improve server response times because the client portion of the multimedia software accesses files from the server.

**Increasing Data Availability**
As described in the scenario for publishing applications, administrators can use DFS to increase the availability of data by storing the data on multiple servers. DFS makes this process transparent by presenting to users what appears to be a single folder in the namespace. Administrators can use File Replication service (FRS) or some other replication method, such as the Windows Resource Kit Tool Robocopy, to keep the data synchronized on the servers. If one of the servers hosting data is unavailable, clients are referred to another server that hosts the data.

**Technologies Related to DFS**
File Replication service (FRS) can be used to keep data in DFS shared folders synchronized among servers. However, DFS and FRS are two separate technologies, and DFS does not require FRS. You can use other replication methods, such as manual copying, the Resource Kit Tool Robocopy, or other replication tools to keep DFS shared folders synchronized. Conversely, if you want to use FRS to keep data in shared folders synchronized, you must use DFS.

**DFS Dependencies**
DFS has the following dependencies:

- **Active Directory replication.** Domain-based DFS requires that Active Directory replication is working properly so that the DFS object resides on all domain controllers in the domain.

- **Server Message Block (SMB).** Clients must access DFS root servers by using the SMB protocol.

- **Remote Procedure Call (RPC) service and Remote Procedure Call Locater service.** The DFS tools use RPC to communicate with the DFS service running on DFS root servers.

- **Distributed File System service dependencies.** The Distributed File System service must be running on all DFS root servers and domain controllers so that DFS can work properly. This service depends on the following services:

    - The Server service, Workstation service, and Security Accounts Manager (SAM) service on DFS root servers. The Distributed File System service also requires an NTFS volume to store the physical components of DFS on root servers.

    - The Server service and Workstation service on domain controllers.

**Related Information**
The following resources contain additional information that is relevant to this section.

- FRS Technical Reference

- Resource Kit Tools in the Tools and Settings Collection for information about Robocopy.

# How DFS Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- DFS Terminology
- DFS Client and Server Compatibility
- Characteristics of Namespace Types
- DFS Architecture
- DFS Physical Structures and Caches
- DFS Processes and Interactions
- DFS Protocols
- DFS Interfaces
- Network Ports Used by DFS
- Related Information


Distributed File System (DFS) allows administrators to group shared folders located on different servers by transparently connecting them to one or more DFS namespaces. A DFS namespace is a virtual view of shared folders in an organization. Using the DFS tools, an administrator selects which shared folders to present in the namespace, designs the hierarchy in which those folders appear, and determines the names that the shared folders show in the namespace. When a user views the namespace, the folders appear to reside on a single, high-capacity hard disk. Users can navigate the namespace without needing to know the server names or shared folders hosting the data. DFS also provides many other benefits, including fault tolerance and load-sharing capabilities, making it ideal for all types of organizations. For more information about the scenarios in which DFS is commonly used, see "What Is DFS?".

For information about DFS on later versions of Windows Server, see DFS Namespaces and DFS Replication Overview.

These sections provide an in-depth view of how DFS works in an optimal environment. An optimal environment for DFS is defined as follows:

- Domain Name System (DNS) and Active Directory replication are working properly.
- Sites and site costs in Active Directory are configured properly.
- The domain controller acting as the primary domain controller (PDC) emulator is working properly.
- The Distributed File System (DFS) service is running on all domain controllers and root servers.
- Client computers are running one of the following operating systems: Windows NT 4.0 with Service Pack 6a, Windows 2000 with Service Pack 4 or later, Windows XP with Service Pack 1 or later, or Windows Server 2003.
- Client computers are properly joined to the domain.
- No firewalls block remote procedure call (RPC) ports used by DFS and the DFS root server that hosts DFS. These ports are described in "Network Ports Used by DFS" later in this section.
- The **Bridge all site links** option in Active Directory must be enabled. (This option is available in the Active Directory Sites and Services snap-in.) Turning off **Bridge all site links** can affect the ability of DFS to refer client computers to target computers that have the least expensive connection cost. An Intersite Topology Generator that is running Windows Server 2003 relies on the **Bridge all site links** option being enabled to generate the intersite cost matrix that DFS requires for its site-costing functionality.

If you turn off this option, you must create site links between the Active Directory sites for which you want DFS to calculate accurate site costs. Any sites that are not connected by site links will have the maximum possible cost. For more information about site link bridging, see "Active Directory Replication Topology Technical Reference."

- Windows 2000 domain controllers and Windows 2000 DFS root servers do not have the **restrictanonymous** registry entry set to 2. (This registry entry is located at HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Control/Lsa.) For more information about this registry entry, see "DFS Client and Server Compatibility" and "How Site Discovery Works" later in this section.

#### DFS Terminology

Before you review the DFS components and processes, it is helpful to understand DFS terminology. The following sections provide terms and definitions, an illustration of DFS namespaces, and a brief description of the client and server roles in DFS.

#### DFS Terms and Definitions

The following terms are used to describe the basic components of DFS.

**DFS namespace**

A virtual view of shared folders on different servers as provided by DFS. A DFS namespace consists of a root and many links and targets. The namespace starts with a root that maps to one or more root targets. Below the root are links that map to their own targets.

**DFS link**

A component in a DFS path that lies below the root and maps to one or more link targets.

**DFS path**

Any Universal Naming Convention (UNC) path that starts with a DFS root.

**DFS root**

The starting point of the DFS namespace. The root is often used to refer to the namespace as a whole. A root maps to one or more root targets, each of which corresponds to a shared folder on a separate server. The DFS root must reside on an NTFS volume. A DFS root has one of the following formats: *ServerName**RootName* or *DomainName**RootName*.

**domain-based DFS namespace**

A DFS namespace that has configuration information stored in Active Directory. The path to access the root or a link starts with the host domain name. A domain-based DFS root can have multiple root targets, which offers fault tolerance and load sharing.

**link referral**

A type of referral that contains a list of link targets for a particular link.

**link target**

The mapping destination of a link. A link target can be any UNC path. For example, a link target could be a shared folder or another DFS path.

**referral**

A list of targets, transparent to the user, which a DFS client receives from DFS when the user is accessing a root or a link in the DFS namespace. The referral information is cached on the DFS client for a time period specified in the DFS configuration.

**root referral**

A type of referral that contains a list of root targets for a particular root.

**root target**

A physical server that hosts a DFS namespace. A domain-based DFS root can have multiple root targets, whereas a stand-alone DFS root can only have one root target. Root targets are also called root servers.

**stand-alone DFS namespace**

A DFS namespace whose configuration information is stored locally in the registry of the root server. The path to access the root or a link starts with the root server name. A stand-alone DFS root has only one root target. Stand-alone roots are not fault tolerant; when the root target is unavailable, the entire DFS namespace is inaccessible. You can make stand-alone DFS roots fault tolerant by creating them on server clusters.

#### DFS Namespaces Illustrated

The following figure illustrates a physical view of file servers and shared folders in the Contoso.com domain. Without a DFS

namespace in place, users need to know the names of six different file servers, and they need to know which shared folders reside on each file server.

**Example of Physical Servers and Shared Folders in Contoso.com**




When the IT group in Contoso.com implements DFS, they must first decide the type of namespace to implement. Windows Server 2003 offers two types of namespaces: stand-alone and domain-based. The IT group also chooses a root name, which is similar to the shared folder name in a Universal Naming Convention (UNC) path *ServerName**SharedFolderName*.

The following figure illustrates two namespaces as users would see them. Notice how the address format differs — one begins with a server name, Software, and the other begins with a domain name, Contoso.com. These differences illustrate the two types of roots: stand-alone roots, which begin with a server name, and domain-based roots, which begin with a domain name. Valid formats for domain names include *NetBIOSDomainName**RootName* and *DNSDomainName**RootName*.

**Examples of DFS Namespaces**




A root can have one or more root targets. A root target, also known as a root server, is a physical server that hosts a namespace. A domain-based root can have multiple root servers to provide redundancy and high availability, whereas a stand-alone root can only have one root server. (Stand-alone roots created on server clusters provide high availability.) Root servers are transparent to users, who see only a single folder, whereas administrators can view the physical root servers by using the Distributed File System snap-in, the Dfscmd.exe command-line tool, or the Dfsutil.exe command-line tool (available as part of the Windows Support Tools). The following figure illustrates roots and root servers.

**Roots and Root Servers**




Below the root, administrators can create multiple links. For example, if an administrator creates a link named Templates, users browsing the domain-based namespace shown above would see the path \\Contoso.com\Public\Templates. If an administrator wants to create a deeper hierarchy in the namespace, the administrator can use subfolder names in the link name, such as Templates\Presentations or Templates\Specs. Users who browse the namespace would see \\Contoso.com\Public\Templates. If the user opens the Templates folder, the user would see two subfolders: Presentations and Specs.

Each link can correspond to one or more link targets. A link target can be any UNC path. For example, a link target could be a shared folder (such as \\NOAMFS3\Users), a folder underneath a shared folder, or a path to another DFS namespace (such as \\Contoso.com\Public\Users). Link targets are transparent to users but are visible to administrators by using the DFS tools.

**Overview of Client and Server Roles in DFS**

In addition to understanding the namespace components (roots, root targets, links, and link targets), it is also helpful to understand the client and server roles involved in DFS. The following sections summarize client and server roles in DFS. These roles are described in more detail later in this document.

**DFS clients**

DFS clients are the computers used to access namespaces. DFS clients are typically end-user computers; computers running Windows-

based server operating systems can also access namespaces. DFS clients in this section are assumed to run one of the operating systems described in "DFS Client and Server Compatibility" later in this section.

**Domain controllers**
Domain controllers play numerous roles in DFS:

- DFS clients access domain controllers to get a list of trusted domains and domain controllers for those domains.
- When a DFS client first attempts to access a domain-based namespace, a domain controller provides a list of root servers to the client. This list of root servers is known as a root referral.
- Domain controllers store DFS metadata in Active Directory about domain-based namespaces. DFS metadata consists of information about entire namespace, including the root, root targets, links, link targets, and settings. By default, root servers that host domain-based namespaces periodically poll the domain controller acting as the primary domain controller (PDC) emulator master to obtain an updated version of the DFS metadata and store this metadata in memory.
- Whenever an administrator makes a change to a domain-based namespace, the change is made on the domain controller acting as the PDC emulator master and is then replicated (via Active Directory replication) to other domain controllers in the domain.
- Domain controllers acting as the Intersite Topology Generator generate site cost information. This information describes the relative cost of the connection between two sites. DFS uses this information to sort targets in order of lowest cost in referrals when least-expensive target selection is enabled.

**Root servers**
Root servers (also called root targets) are the servers on which administrators create stand-alone and domain-based namespaces. A root server can be a member server or a domain controller. Root servers play the following roles in DFS:

- For stand-alone namespaces, the DFS metadata is stored in the registry of the root server. For domain-based namespaces, DFS metadata is stored in Active Directory. Root servers also keep the DFS metadata for each namespace in a memory cache.
- Root servers host shared folders that act as the root (known as the root folder) and subfolders (under the root folders) that correspond to every link (known as link folders) in those namespaces.
- When a client attempts to access a root or link in a stand-alone namespace, or a link in a domain-based namespace, the root server provides a referral to the client.

**Link targets**
Link targets are typically shared folders or folders within shared folders. Link targets can be served by any network file system that is accessible by a UNC path, such as Server Message Block (SMB), NetWare Core Protocol (NCP) for NetWare, or Network File System (NFS) for UNIX. (The client computers must have the appropriate redirector installed to access link targets.) The UNC path can lead to shared folders in any workgroup, shared folders within the same domain as the namespace, shared folders in trusted domains, and shared folders in trusted forests.

Shared folders that are specified as link targets have no special settings that indicate that they are part of a DFS namespace. All existing shared folder permissions and NTFS permissions on the shared folder still apply when users access the shared folder through the namespace.

A link target can also be a DFS path in another namespace. For example, the Software link in \\Contoso.com\Public\Software might have a link target of \\Software\Public, which is a root within a stand-alone namespace. When using DFS paths as link targets, it is important to ensure that client failover works correctly. For more information, see "Linking to Different DFS Namespaces" later in this section.

**DFS Client and Server Compatibility**
The following table describes the types of clients and servers that can act as DFS clients, root servers, and link targets.

### DFS Compatibility

| Platform | Act As DFS Clients? | Act As a Root Server? | Act As a Link Target? |
|---|---|---|---|

| | | | |
|---|---|---|---|
| Microsoft Windows Server 2003, Web Edition | Yes | Yes<br><br>Can host one stand-alone namespace or one domain-based namespace per server. | Yes |
| Windows Server 2003, Standard Edition | Yes | Yes<br><br>Can host one stand-alone namespace or one domain-based namespace per server. | Yes |
| Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition | Yes | Yes<br><br>Can host multiple stand-alone namespaces and multiple domain-based namespaces per server. | Yes |
| Windows Storage Server 2003 | Yes | Yes<br><br>Can host multiple stand-alone namespaces and multiple domain-based namespaces per server. | Yes |
| Windows XP | Yes | No | Yes |
| Windows Preinstallation Environment (WinPE) | Yes<br><br>Can access stand-alone namespaces but cannot access domain-based namespaces. | No | No |
| Windows 2000 Server family | Yes | Yes<br><br>Can host one stand-alone namespace or domain-based namespace per server. | Yes |
| Windows 2000 Professional | Yes | No | Yes |
| Windows NT Server 4.0 with Service Pack 6a | Yes | Yes<br><br>Can host a single stand-alone namespace per server. | Yes |
| Windows NT Workstation 4.0 with Service Pack 6a | Yes | No | Yes |
| Windows 98 | Yes<br><br>Client for stand-alone DFS included; install the Active Directory client extension for Windows 95 or Windows 98 to access domain-based DFS namespaces. | No | Yes |

When evaluating client compatibility, review the following important considerations:

- Client computers must be members of a local or trusted domain before they can access a domain-based namespace by using the format *NetbiosDomainName**RootName*. If clients are members of a workgroup or an untrusted domain and can resolve DNS names, they can access domain-based namespaces by using the format *DNSDomainName**RootName*. For information about how clients determine the list of trusted domains, see "DFS Physical Structures and Caches on DFS Clients."

- Link targets can be shared folders served by other protocols, such as NetWare Core Protocol (NCP) for NetWare and Network File System (NFS) for UNIX, but client computers must have the appropriate redirector installed to access those link targets.

- In organizations that have a large number of domains and forest trusts, clients might have difficulty accessing link targets in other domains or forests. For more information, see "DFS-Related Processes on Clients" later in this section.

- If a Windows 2000 Server root server has the **restrictanonymous** registry entry set to 2, and the root server is not a domain controller, then DFS clients in a workgroup cannot access DFS namespaces that are hosted on the root server. (This registry entry

is located at HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Control/Lsa.)

**Caution**

- Incorrectly editing the registry might severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

#### Characteristics of Namespace Types

The following table provides a general overview of the differences between domain-based and stand-alone namespaces.

**Characteristics of Namespace Types**

| Characteristic | Domain-based | Stand-Alone |
|---|---|---|
| Path to namespace | *NetbiosDomainName**RootName*<br><br>*DNSDomainName**RootName* | *ServerName**RootName* |
| Where DFS metadata is stored | In Active Directory. Also stored in a memory cache on the root server. | In the registry of the root server. Also stored in a memory cache on the root server. |
| Group memberships required to create or delete namespaces | DFS administrators must be members of the Domain Admins group to create or delete domain-based namespaces or have delegated permissions to the DFS-Container object in Active Directory. | DFS administrators must be members of the local Administrators group on the local server to create or delete stand-alone namespaces. |
| Group memberships required to administer namespaces | DFS administrators must be members of the local Administrators group on each of the root servers to add or delete links, link targets, roots, and root targets.* | DFS administrators must be members of the local Administrators group on the local server to add or delete links, link targets, roots, and root targets. |
| DFS namespace size restrictions | We recommend that you keep the size of the DFS object in Active Directory below 5 megabytes (MB) by using fewer than 5,000 links in domain-based namespaces. | The largest recommended namespace size for a stand-alone namespace is 50,000 links. |
| Supported methods to ensure DFS root availability | Create multiple root targets in the same domain. | Create a stand-alone root on a server cluster. |
| Supported methods to ensure link target availability | Create multiple link targets and replicate files by using one of the following methods:<br>• Enabling FRS<br>• Copying files manually or by using scripts<br>• Using another replication tool | Create multiple link targets and replicate files by using one of the following methods:<br>• Copying files manually or by using scripts<br>• Using another replication tool |

*To ensure that the administrator can reliably administer the namespace, we recommend that the administrator be a member of the local Administrators group on each root server. If a DFS administration tool sends a request to a root server on which the administrator is not a member of the local Administrators group, the request will fail. If the administrator is a member of the local Administrators group on some but not all root servers, the administrator's ability to administer the namespace will be inconsistent.

#### DFS Architecture

The DFS service (Dfssvc.exe) is the core component of the DFS architecture and runs on root servers and domain controllers. The primary functions of the DFS service include handling referrals, managing namespaces, and communicating with the DFS driver (Dfs.sys).

The components of the DFS architecture on DFS clients and root servers are illustrated in the following figure. In this figure, the DFS architecture of the domain controller is simplified to show only the DFS object.

**DFS Client and Root Server Architecture**




The following figure illustrates the DFS architecture of a domain controller and a simplified view of the DFS client and root server architecture. Note that domain controllers use DFS architecture similar to root servers; this is because domain controllers play a role in referring client computers to domain-based roots. It is also possible for domain controllers to host namespaces and play the role of root server. In this case, the domain controller also hosts the DFS metadata cache (regardless of namespace type) and the stand-alone DFS metadata in its registry (for stand-alone namespaces).

**DFS-Related Architecture on Domain Controllers**




The following tables describe the components of the DFS architecture.

**Components of DFS Tools Described in the Figure Titled "DFS Client and Root Server Architecture"**

| Component | Description |
|---|---|
| Dfsshlex.dll | Displays the DFS tab when you use a DFS client to view the properties of a folder in a DFS namespace using Windows Explorer. |

| | |
|---|---|
| Distributed File System snap-in (Dfsgui.msc and Dfsui.dll) | The graphical user interface tool used for administering DFS namespaces. By default, the Distributed File System snap-in (Dfsgui.msc) is installed on servers running Windows Server 2003. You can administer namespaces remotely from a computer running Windows XP with Service Pack 1 by installing the Windows Server 2003 Administration Tools Pack, which the includes Distributed File System snap-in. |
| Distributed File System command-line utility (Dfscmd.exe) | The DFS command-line tool used to perform basic DFS tasks, such as configuring DFS roots, links, and targets. Also supports scripting. Dfscmd.exe is available on all servers running a member of the Windows Server 2003 family. |
| Distributed File System utility (Dfsutil.exe) | The DFS command-line tool used to perform basic and advanced DFS tasks, such as enabling root scalability mode, least expensive target selection, and same-site target selection. When installed on DFS clients, Dfsutil.exe can be used to view and clear the referral cache (PKT cache), domain cache (SPC cache), and MUP cache. Also supports scripting. Dfsutil.exe is included with the Windows Support Tools and is available on the Microsoft Download Center Web site. |
| Netapi32.dll | Contains the NetDFS*xxx* APIs that are used to administer remote root servers. Also contains the APIs that are used to view and clear the referral cache (PKT cache), domain cache (SPC cache), and MUP cache on clients. |

**Components of DFS Client Architecture Described in the Figure Titled "DFS Client and Root Server Architecture"**

| Component | Description |
|---|---|
| Domain cache | Contains the domain name referrals and domain controller referrals that are stored in memory on each client computer. The domain cache stores NetBIOS and DNS domain names for the local domain, all trusted domains in the forest, domains in trusted forests, and the mappings of the domain names to the domain controllers. The domain cache (also called SPC cache) is maintained in Mup.sys. |
| I/O Manager | Plays a role in handling the file I/O in redirecting the UNC paths to Mup.sys. I/O Manager is part of the core kernel (Ntoskrnl.exe) of the operating system. |
| Mrxsmb.sys | Handles communications to the root servers, domain controllers, and Windows-based file servers that use the CIFS protocol. |
| Mup.sys | Implements DFS client support and redirector selection. "Mup" stands for "multiple UNC provider."<br><br>Mup.sys is a networking component that handles input/output (I/O) requests for a file or device that has a UNC name (that is, a name that begins with \\). If the UNC name is a DFS path, Mup.sys resolves it to the physical UNC path. After the path is resolved, or if the path was not a DFS path, Mup.sys determines the local redirector that handles the UNC path. |
| Ntlanman.dll | Handles the establishment of drive letter and unnamed connections for DFS and SMB paths. |
| Nwrdr.sys and Mrxdav.sys | The redirectors for Netware and Web Distributed Authoring and Versioning (WebDAV), respectively. |
| Referral cache | Stores the root referrals and link referrals received by the client when it accesses a namespace. The referral cache (also known as the PKT cache) is maintained in Mup.sys. |
| Workstation service (Wkssvc.dll) | Sends domain controller information and domain names to Mup.sys. Mup.sys uses this information to generate the referrals in the client's domain cache. This service is also a component of the Server Message Block (SMB) client. Also manages the loading and unloading of the Mrxsmb.sys. |

**Components of the DFS Root Server and Domain Controller Architecture Described in the Figure "DFS-Related Architecture on Domain Controllers"**

| Component | Description |
|---|---|
| Dfs.sys | The DFS driver. Dfs.sys runs only on Windows-based server operating systems and reflects referral requests from DFS clients to the DFS service running in user mode. Dfs.sys also handles the processing of links when they are encountered during file system access. |
| DFS metadata cache | The in-memory copy of DFS metadata maintained in Dfssvc.exe on root servers. DFS metadata consists of information about entire namespace, including the root, root targets, links, link targets, and settings. |

| | |
|---|---|
| Dfssvc.exe | The DFS service. Provides server-side support for NetDFS*xxx* APIs that configures and maintains DFS namespaces. The DFS service is also responsible for maintaining an up-to-date version of the DFS metadata and for giving referrals to clients who attempt to access the namespace. |
| Domain-based root referral cache and site caches | The in-memory copies of domain-based root referrals and site information that enable Dfssvc.exe to perform fast lookups. These caches are described in the next section. |
| Registry | Where the DFS metadata is stored for stand-alone DFS namespaces. |
| Srv.sys | The SMB Service server driver. Plays a role in DFS by passing on referral requests from DFS clients to Dfs.sys. |
| Active Directory | Where the DFS object is stored. The DFS object stores the DFS metadata for a domain-based namespace. |

**DFS Physical Structures and Caches**

The following figure illustrates the physical structures and caches on domain controllers, root servers, and DFS clients. These physical structures and caches are described in the sections that follow.

**DFS Physical Structures and Caches**




**DFS Physical Structures and Caches on Domain Controllers**

Domain controllers play an important role for domain-based namespaces. They store the DFS metadata for domain-based namespaces, and they provide referrals to DFS clients to help them access domain-based namespaces. There are a number of physical structures and caches that support these processes; these structures and caches are illustrated in the following figure.

**Physical Structures and Caches on Domain Controllers**




The following sections describe each of the physical structures and caches on domain controllers. Domain controllers can also host roots; in this case, the domain controller also contains the physical structures and caches described in "DFS Physical Structures and Caches on Root Servers" later in this section.

**DFS Object in Active Directory**

The DFS object stores the DFS metadata for a domain-based namespace. The DFS object is created in Active Directory when you create a domain-based root, and Active Directory replicates the entire DFS object to all domain controllers in a domain. When a client requests a referral for a domain-based namespace, the domain controller first checks its domain-based root referral cache (described later in this section) for an existing referral. If none exists, the domain controller locates the DFS object for that namespace and uses the metadata in the object to create the referral.

The following figure illustrates the location of the DFS objects in Active Directory. Each DFS object (class **fTDfs**) corresponds to a domain-based namespace.

**Hierarchy of DFS Objects in Active Directory**




Each DFS object has the following four important attributes:

- **pKT.** A binary representation of the DFS metadata for this root.
- **pKTGuid.** The globally unique identifier (GUID) of the DFS metadata.
- **remoteServerName.** Lists the root targets for the root.
- **Security descriptor.** Controls who can modify the DFS object. The security descriptor on the DFS-Configuration object controls who can create or delete a DFS object.

It is important to be aware of the size of the DFS object. An entry in the metadata uses approximately 400 bytes per DFS link. (This figure does not include any comments that you add to describe the namespace.) A mid-sized DFS namespace with 100 links requires approximately 40 kilobytes (KB). Any change to the namespace causes the entire DFS object to be replicated to all domain controllers in the domain. We recommend that the DFS object for a domain-based namespace be less than 5 megabytes (MB) (which is equivalent to approximately 5,000 links) to reduce the impact of network traffic originating from changes to the namespace.

For more information about how domain-based root servers discover changes to the DFS object in Active Directory, see "How DFS Discovers Changes to the Namespace" later in this section.

**DFS Registry Entries**

DFS supports several registry entries for configuring DFS behavior on domain controllers. DFS registry entries are located in the registry under the following subkeys:

- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DFS
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters

For a list of DFS registry entries, see "DFS Tools and Settings."

**Domain Name Referral Cache**
A domain name referral contains the NetBIOS and DNS names of the local domain, all trusted domains in the forest, and domains in trusted forests. A DFS client requests a domain name referral from a domain controller to determine the domains in which the clients can access domain-based namespaces.

By default, domain controllers store domain name referrals in a memory cache for 12 hours; this period is defined by the **DomainNameIntervalinSeconds** registry entry on the domain controller. Restarting the DFS service on the domain controller will cause this cache to be cleared. If a domain name changes, or if domains are added or removed from the forest or trusted forests, these changes are reflected in domain name referrals after 12 hours or after this cache is cleared.

**Domain Controller Referral Cache**
A domain controller referral contains the NetBIOS and DNS names of the domain controllers for the list of domains it has cached. A DFS client requests a domain controller referral from a domain controller (in the client's domain) to determine which domain controllers can provide a referral for a domain-based namespace.

Domain controllers store these domain controller referrals in a memory cache for 10 minutes; you cannot change this setting. Restarting the DFS service on the domain controller will cause this cache to be cleared. If a domain controller name changes, or if domain controllers are promoted or demoted from the domain, these changes are reflected in domain controller referrals after 10 minutes or after this cache is cleared.

**Domain-based Root Referral Cache**
A domain-based root referral contains a list of root targets that host a given domain-based namespace. Domain controllers provide domain-based root referrals to clients that attempt to access a domain-based namespace.

By default, domain controllers store domain-based root referrals in a memory cache for 15 minutes; this period is defined by the **RootReferralTimeoutInSeconds** registry entry on the domain controller. Restarting the DFS service on the domain controller will cause this cache to be cleared. If root targets are added or removed, or if settings are changed on the root, these changes are reflected in domain-based root referrals after 15 minutes or after this cache is cleared.

### Note

- The domain-based root referrals in this memory cache do not store targets in any particular order. The targets are sorted according to the target selection method only when requested from the client. Also, these referrals are based on DFS metadata stored on the local domain controller, not the PDC emulator master.

**Client Site Cache**
When a client requests a referral from a domain controller, the DFS service on the domain controller uses the site information defined in Active Directory (through the DSAddressToSiteNames API) to determine the site of the client, based on the client's IP address. DFS stores this information in the client site cache, which contains mappings of client IP addresses to site names. DFS uses this cache to quickly determine the site that a client belongs to.

By default, domain controllers keep entries in the client site cache for up to 24 hours. This period is defined by the **SiteSupportIntervalInSeconds** registry entry (12 hours) on the domain controller, multiplied by two. If a site name changes, or if a client moves from one site to another and uses the same IP address, the site information in this cache will be out-of-date for a maximum of twice the value of the **SiteSupportIntervalInSeconds** registry entry (24 hours) or until the DFS service is restarted on the domain controller.

When a cleanup process begins after the configured time interval, entries are handled as follows:

- If a site name that corresponds to a client was accessed since the last cleanup process, the entry is refreshed.
- If a site name that corresponds to a client has not been accessed since the last cleanup process, the entry is purged.

By default, DFS can store 200,000 entries in the client site cache. This limit is determined by the **MaxClientsToCache** registry entry.

**Target Site Cache**
The DFS service on domain controllers uses the site information defined in Active Directory (through the DSAddressToSiteNames API) to determine the site for domain-based root targets and domain controllers based on their current IP addresses. DFS stores this information in its target site cache, which contains mappings of root server names and domain controller names to site names. DFS uses this information to quickly determine the site that a root server or domain controller belongs to.

By default, domain controllers keep entries in the target site cache for up to 24 hours. This interval is defined by the **SiteSupportIntervalInSeconds** registry entry value (12 hours) on the domain controller, multiplied by two. After the configured time interval, DFS refreshes the site information in this cache. If a site name changes or a root target moves from one site to another, the site information in this cache will be out-of-date for a maximum of twice the value of the **SiteSupportIntervalInSeconds** registry entry (24 hours) or until the DFS service is restarted on the domain controller.

#### Site Cost Cache

The site cost cache on domain controllers contains a mapping of sites to their associated cost information as defined in Active Directory. (Costs are configured by administrators and are based on the expense of transferring data over the connection.) DFS populates this cache only when least expensive target selection is enabled. DFS uses the site cost cache to quickly determine the connection cost associated with domain controllers and domain-based root targets so that DFS can sort targets in order of lowest cost in referrals.

The DFS service obtains site cost information for domain controllers and domain-based root targets as needed to fulfill client referral requests. By default, entries in the site cost cache are maintained for up to 12 hours. This period is defined by the **SiteSupportIntervalInSeconds** registry entry value (12 hours) on the local domain controller. If the connection cost between two sites is changed in Active Directory, the site cost information in this cache will be out-of-date for a maximum of the **SiteSupportIntervalInSeconds** registry entry (12 hours) or until the DFS service is restarted.

The **QuerySiteCostTimeoutinSeconds** registry entry on domain controllers is also used for this cache to specify a time-out period for site cost queries. After the time-out period, the DsQuerySitesByCost API calls for site cost information in Active Directory will fail. By default, the time-out period is 30 seconds. If DFS cannot determine a site cost within 30 seconds, DFS assumes the maximum possible cost for that site.

#### DFS Physical Structures and Caches on Root Servers

The physical structures and caches on a root server vary according to the type of namespace the root server hosts (domain-based or stand-alone). Servers running Windows 2003 Server, Enterprise Edition, and Windows Server 2003, Datacenter Edition, can host multiple stand-alone and domain-based roots.

The following figure illustrates the physical structures and caches on root servers.

**Physical Structures and Caches on Root Servers**




The following sections describe each of the physical structures and caches on root servers.

#### Stand-Alone DFS Metadata in the Registry

The stand-alone DFS metadata contains information about the root, root target, links, link targets, and settings defined for each stand-alone namespace. This metadata is maintained in the registry of the root server at HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfs\Roots\Standalone.

### Note

- Domain-based root servers have a registry entry for each root under KEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfs\Roots\Domain, but these entries do not contain the domain-based DFS metadata. When the DFS service starts on a root server, the service checks this path for registry entries that correspond to domain-based roots. If these entries exist, the root server polls the PDC emulator master to obtain the DFS metadata for each domain-based namespace and stores the metadata in memory.

#### Root and Link Folders

Each root and link in a namespace has a physical representation on an NTFS volume on each root server that hosts the namespace. The DFS root corresponds to a shared folder, known as the root folder, on the root server. When you use the DFS tools to add links to the root, DFS creates special folders under each root folder. These folders, called link folders, are actually reparse points, and they display an error message similar to the following message if you try to access them on the local server:

"E:\Public\GroupData is not accessible. The network location cannot be reached."

DFS clients that access the link folders from across the network are redirected to the appropriate link target by using the process described in "The Referral Process" later in this section.

The following figure illustrates volume E:\ on the local storage of a root server. The volume contains root folder and link folders for the \\Contoso.com\Public namespace.

**Link Folders on a Root Server**




For more information about how these folders are created, see "Root and Link Creation Process" later in this section.

**DFS Registry Entries**

DFS supports several registry entries for configuring DFS behavior on root servers. The DFS service registry entries are located in the registry under the following subkeys:

- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DFS
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters

For a list of these DFS registry entries, see "DFS Tools and Settings."

**DFS Metadata Cache**

The DFS metadata cache on root servers contains information about the root, root targets, links, link targets, and settings defined in a namespace. For domain-based namespaces, this metadata is also stored in the DFS object in Active Directory. For stand-alone namespaces, this metadata is also stored in the root server's registry.

This cache is updated each time the root server polls the DFS object in Active Directory (for domain-based namespaces) and the registry (for stand-alone namespaces). For more information, see "How DFS Discovers Changes to the Namespace" later in this section.

**Domain-based Root Referral Cache**

When a client computer contacts a domain-based DFS root server directly by using the syntax *ServerName**RootName*, the root server provides a domain-based root referral to the client computer. One scenario in which this occurs is when a client attempts to access a link target that points to a path within another domain-based namespace. For more information, see "Linking to Different DFS Namespaces" later in this section.

By default, root servers store these domain-based root referrals in a memory cache for 15 minutes; this period is defined by the **RootReferralTimeoutInSeconds** registry entry on the root server. Restarting the DFS service on the root server will cause this cache to be cleared. If root targets are added or removed, or if settings are changed on the root, these changes are reflected in domain-based root referrals (provided by root servers) after 15 minutes or after the cache is cleared.

**Note**

- The domain-based root referral cache does not store targets in any particular order. The targets are sorted according to the target selection method only when requested from the client. Also, these referrals are based on DFS metadata stored on the local domain controller, not the PDC emulator master.

**Client Site Cache**

When a client requests a referral from a root server, the DFS service on the root server uses the site information defined in Active Directory (via the DSAddressToSiteNames API) to determine the site of the client, based on the client's IP address. DFS stores this information in the root server's client site cache, which contains mappings of client IP addresses to site names. DFS uses this cache to quickly determine the site that a client belongs to.

By default, root servers keep entries in the client site cache for up to 24 hours. This period is defined by the **SiteSupportIntervalInSeconds** registry entry (12 hours) on the root server, multiplied by two. If a site name changes, or if a client moves from one site to another and uses the same IP address, the site information in this cache will be out-of-date for a maximum of twice the value of the **SiteSupportIntervalInSeconds** registry entry (24 hours) or until the DFS service is restarted on the root server.

When a cleanup process begins after the configured time interval, entries are handled as follows:

- If a site name that corresponds to a client was accessed since the last cleanup process, the entry is refreshed.
- If a site name that corresponds to a client has not been accessed since the last cleanup process, the entry is purged.

By default, DFS can store 200,000 entries in the client site cache. This limit is determined by the **MaxClientsToCache** registry entry.

**Target Site Cache**
When the DFS service starts on a root server, it uses the site information defined in Active Directory (through the DSAddressToSiteNames API) to determine the site for all link targets, based on the current IP address of each link target. The DFS service also obtains site information when new link targets are added. DFS caches this information in the root server's target site cache. The target site cache contains mappings of link target server names to site names. DFS uses this information to quickly determine the site that a link target belongs to.

By default, root servers keep entries in the target site cache for up to 24 hours. This interval is defined by the **SiteSupportIntervalInSeconds** registry entry value (12 hours) on the root server, multiplied by two. After the configured time interval, DFS refreshes the site information in this cache. If a site name changes or a link target server moves from one site to another, the site information in this cache will be out-of-date for a maximum of twice the value of the **SiteSupportIntervalInSeconds** registry entry (24 hours) or until the DFS service is restarted.

**Site Cost Cache**
The site cost cache on root servers contains a mapping of sites to their associated cost information, as defined in Active Directory. (Costs are based on the expense of transferring data over the connection.) DFS populates this cache only when least expensive target selection is enabled. DFS uses the site cost cache to quickly determine the connection cost associated with link targets, so that DFS can sort link targets in referrals in order of lowest cost.

The DFS service obtains site cost information for link targets (through the DsQuerySitesByCost API) as needed to fulfill client referral requests. By default, root servers keep entries in the site cost cache for up to 12 hours. This period is defined by the **SiteSupportIntervalInSeconds** registry entry value (12 hours) on the root server. If the connection cost between two sites is changed in Active Directory, the site cost information in this cache will be out-of-date for a maximum of the **SiteSupportIntervalInSeconds** registry entry (12 hours) or until the DFS service is restarted.

The **QuerySiteCostTimeoutinSeconds** registry entry on root servers is also used for this cache to specify a time-out period for site cost queries. After the time-out period, the DsQuerySitesByCost API calls for site cost information in Active Directory will fail. By default, the time-out period is 30 seconds. If DFS cannot determine a target's site cost within 30 seconds, DFS assumes the maximum possible cost for the target server.

**DFS Physical Structures and Caches on DFS Clients**
DFS clients running Windows XP and Windows 2000 store referrals from domain controllers and root servers to improve performance. A client can use the referrals in its cache when reaccessing a target in the namespace, rather than recontacting domain controllers and root servers for referrals. DFS clients also have registry entries used for customizing DFS behavior on the client.

The following figure illustrates the physical structures and caches on DFS clients.

**Caches on DFS Clients**




These physical structures and caches are described in the sections that follow. For information about how clients update their caches, see "DFS-Related Processes on Clients" later in this section.

**DFS Registry Entries**
DFS registry entries are located in the DFS client's registry under the following subkeys:

- HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\LanmanWorkstation\
- HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\System\DFSClient\.

For a list of DFS registry entries, see "DFS Tools and Settings.

**Domain Cache on Clients**
The domain cache on clients contains two types of referrals:

- **Domain name referrals,** which are the NetBIOS and DNS domain names for the client's local domain, trusted domains in the forest, and domains in trusted forests.

- **Domain controller referrals,** which are the mappings of the domain names to the domain controllers that host those domains.

You can view the domain cache on a client computer by using the Dfsutil.exe command-line tool with the **/spcinfo** parameter. The following text illustrates the output displayed when you use this command.

```
[*][dc-01.contoso.com]
[*][CONTOSO]
[*][contoso.com]
[+][contoso.com]
        [-dc-02.contoso.com]
        [+dc-01.contoso.com]
        [-dc-04.contoso.com]
        [-dc-03.contoso.com]
[-][AFRICA]
[-][EUROPE]
[+][CONTOSO]
        [-DC-02]
        [+DC-01]
        [-DC-03]
        [-DC-04]
[-][europe.contoso.com]
[-][africa.contoso.com]
```

The preceding sample output can be interpreted as follows:

- Entries preceded by [*] are provided by the Workstation service.

- The [+] next to the domain controller named DC-01 under [contoso.com] and [CONTOSO] indicates that it is the active domain controller for that domain. The client will contact this domain controller as needed to obtain domain name referrals, domain controller referrals, and domain-based root referrals.

- The client has already contacted DC-01 to receive domain name referrals for Contoso.com, Europe.Contoso.com, and Africa.Contoso.com.

- The client has already contacted DC-01 to receive domain controller referrals for the Contoso.com domain.

For information about how clients update their domain caches, see “DFS-Related Processes on Clients” later in this section.

**MUP Cache**

The multiple UNC provider (MUP) cache stores information about which redirector, such as DFS, SMB, or WebDAV, is required for each UNC path that a client computer attempts to access. Entries in the MUP cache are held for 15 minutes. You can use Dfsutil.exe’s **/PurgeMupCache** parameter to clear the MUP cache. This might be necessary when a folder is changed from an SMB shared folder to a WebDAV or DFS root folder or vice versa.

For information about how clients update their MUP caches, see “DFS-Related Processes on Clients” later in this section.

**Referral Cache on Clients**

DFS clients store root referrals and link referrals in the referral cache (also called the PKT cache). These referrals allow clients to access the root and links within a namespace. You can view the contents of the referral cache by using Dfsutil.exe with the **/pktinfo** parameter on a DFS client. The following figure illustrates the output displayed when you use this command on a Windows XP client computer.

**Sample Referral Cache Output**

```
--mup.sys--
5 entries...
1 Entry: \contoso.com\netlogon
  ShortEntry: \contoso.com\netlogon
  Expires in 0 seconds
  UseCount: 0 Type:0x1 ( DFS )
     0:[\dc-01.contoso.com\netlogon] State:0x31 (ACTIVE)
     1:[\dc-02.contoso.com\netlogon] State:0x21 ( )
     2:[\dc-64.contoso.com\netlogon] State:0x21 ( )
     3:[\dc-04.contoso.com\netlogon] State:0x21 ( )

  Entry: \contoso.com\SysVol
1 ShortEntry: \contoso.com\SysVol
  Expires in 180 seconds
  UseCount: 0 Type:0x1 ( DFS )
     0:[\dc-02.contoso.com\SysVol] State:0x31( ACTIVE )
     1:[\dc-01.contoso.com\SysVol] State:0x21( )
     2:[\dc-64.contoso.com\SysVol] State:0x21( )
     3:[\dc-04.contoso.com\SysVol] State:0x21( )

  Entry: \contoso.com\public
2 ShortEntry: \contoso.com\public
  Expires in 0 seconds
  UseCount: 0 Type:0x81 ( REFERRAL_SVC DFS )
     0:[\root-dfs-01\public] State:0x19 ( ACTIVE )
     1:[\root-dfs-02\public] State:0x09 ( )

  Entry: \software\public
3 ShortEntry: \software\public
  Expires in 300 seconds
  UseCount: 0 Type:0x81 ( REFERRAL_SVC DFS )
     0:[\SOFTWARE\Public] State:0x19 ( ACTIVE )

  Entry: \software\public\Applications
4 ShortEntry: \software\public\Applications
  Expires in 1800 seconds
  UseCount: 0 Type:0x1 ( DFS )
     0:[\software2\apps$] State:0x31 ( ACTIVE )
     1:[\software4\apps$] State:0x21 ( )
     2:[\software3\apps$] State:0x21 ( )
```

The previous figure shows four types of referrals: (1) NETLOGON and SYSVOL referrals, (2) domain-based root referrals, (3) stand-alone root referrals, and (4) link referrals. Each of these referrals contains the following information:

**Entry and ShortEntry**
Entry specifies the full name of the path. ShortEntry specifies the short name (eight characters followed by a period and a three more characters) of the path. Root servers and domain controllers running Windows Server 2003 use the full name when populating Entry and ShortEntry.

**Expires in seconds**
Specifies the Time to Live of the referral. Zero (0) seconds indicates that the referral has expired and that the client will obtain a new referral the next time the client tries to access the target server.

**UseCount and Type**
UseCount is the number of open connections and files for this path. If a client has a mapped drive to a DFS path, its use count goes up by 1. On clients running Windows 2000, Windows XP, and Windows Server 2003, you cannot clear a referral until its use count is 0.

Type indicates the type of referral. Some common types include:

- **Type 0x81 ( REFERRAL_SVC DFS ).** Indicates a root referral.
- **Type 0x1 ( DFS ).** Indicates a link referral that has physical shared folders as link targets.
- **Type 0x10 ( OUTSIDE_MY_DOM).** Indicates a link referral whose link target is a path in a domain-based namespace.

**Target list**
The target list contains the root targets or link targets that correspond to the path specified in Entry. Targets are listed in order, according to the target selection method configured for the root or link. A target marked Active indicates that the client either has used this target or will use this target the next time the user tries to access this portion of the namespace.

**DFS Processes and Interactions**
DFS is a complex service that involves many processes and interactions between clients, root servers, and domain controllers. The major categories of the processes and interactions, along with the names of the relevant sections, are described as follows.

**Processes Related to Referrals**
The referral process is the process by which domain controllers and root servers refer clients to different targets in the namespace. Numerous processes take place so that DFS can provide an optimal list of targets in the referral. These processes are described in the following sections:

- See “How Site Discovery Works” for information about how DFS determines the sites to which clients and targets belong.
- See “How Target Selection Works” for information about the logic that DFS uses to sort targets in a referral. There are three target selection methods: default target selection, same-site target selection, and least expensive target selection.

- See “How DFS Works in Environments Without WINS” for information about how DFS can use either NetBIOS names or DNS names of targets in referrals.

- See “The Referral Process” for the actual steps involved in the referral process.

- See “How DFS Is Used During the Logon Process” for information about the two special types of referrals, SYSVOL and NETLOGON referrals, which are used when a client first logs on to the domain.

- See “Linking to Different DFS Namespaces” for information about how link targets can point to other namespaces (instead of physical shared folders) and the referral process involved when a client navigates from one DFS namespace to another.

- See “What Happens During Failover” for information about how clients react when a target in a referral is not available and how clients react when they have a session enabled with a target and the target server fails.

**Creating and Changing Namespaces**
These following sections describe what happens when you create or change a namespace.

- See “Root Folder and Link Folder Creation Process” for information about how DFS-related folders are created on root servers.

- See “How DFS Discovers Changes to the Namespace” for information about what happens when a namespace changes. This section also provides information about where DFS data is stored so that you can determine how long it takes before clients can see an updated view of the namespace.

- See “How Root Scalability Mode Works” for information about a special mode in which domain-based root servers discover namespace changes by contacting the closest domain controller.

**What Happens During DFS Operations**
See the section “What Happens When You Check Target Status” for information about how DFS works when you check target status by using the Distributed File System snap-in.

**DFS Interoperability**
See the sections “How DFS Works on Server Clusters” and “How DFS Works with Offline Files” for information about how DFS operates on server clusters and the Offline Files feature.

**Client Processes**
See the section “DFS-Related Processes on Clients” for information about the processes that DFS clients perform to update their domain caches, MUP caches, and referral caches. Entries in these caches are used during the referral process.

**How Site Discovery Works**
Site discovery, also called site awareness, is the process that DFS uses to determine which site a root target, link target, or client belongs to. A site is one or more TCP/IP subnets as defined in Active Directory. DFS uses this site information to determine the order in which to sort the targets in a referral.

DFS determines the site of targets and clients as follows:

- When the DFS service starts, DFS determines the site of each target by converting the target’s name to an IP address (using the GetHostbyName API) and then querying Active Directory (using the DSAddressToSiteNames API) to convert the IP address to a site name based on the subnet-to-site mappings as defined in Active Directory.

- When a client attempts to access the DFS namespace, DFS determines the site of a client by querying Active Directory (using the DSAddressToSiteNames API) to convert the client’s IP address to a site name based on the subnet-to-site mappings as defined in Active Directory.

Site discovery in Windows Server 2003 works regardless of the operating system that the target server is running. For example, if a link target is running Windows NT 4.0 or a non-Windows operating system, DFS can determine the target’s site using its IP address. For server clusters and multihomed clients and targets, site discovery works as follows:

- If a target is part of a virtual server in a server cluster, DFS uses the virtual server’s IP address to determine the target’s site.

- When a target server is multihomed (that is, the server has multiple IP addresses), the root server or domain controller determines the target’s site by using the GetHostbyName API, and the first IP address returned is what is used to determine the site.

- When a multihomed client requests a referral, the root server or domain controller uses the source IP address from the client's network request, which could be arbitrary based on the client's network configuration. Or, if a client request goes through a NAT gateway, the root server or domain controller uses the source IP address of the request as rewritten by the NAT gateway.

After root servers and domain controllers obtain site information about client computers and target servers, they store the information in the following memory caches:

- **Target site cache.** The target site cache contains mappings of target server names to site names.

- **Client site cache.** The client site cache contains mappings of client IP addresses to site names.

Site information remains in the client site cache and the target site cache for up to two times the duration specified in the **SiteSupportIntervalInSeconds** registry entry (set to 12 hours, by default). Because site information is cached and not immediately updated after it changes, it is important to understand how out-of-date cached site information affects site discovery for target servers and client computers.

- Target site information can become out-of-date when a target server is moved from one site to another and the target's IP address changes so that the target belongs to a subnet that is mapped to a different site in Active Directory. Changing site names can also cause target site information to become out-of-date. Eventually DFS discovers the new site information when the old site information is updated in the target site cache, which occurs within 24 hours, by default. Until site information is refreshed in the cache, DFS might sort affected targets in referrals based on their old site information, possibly causing clients to connect to target servers outside of their own site or to targets that are not closest to the client computer.

- Client site information can become out-of-date when a client computer moves to another site but keeps the same IP address, when the client's subnet is moved to a different site, or when the site names change. In these scenarios, DFS discovers new site information for the client within 24 hours, by default. Until site information is refreshed in the cache, DFS might sort targets in referrals based on the client's old site information, possibly causing the client to contact target servers outside of its own site or to targets that are not closest to the client.

When reviewing how site discovery works in Windows Server 2003, it is also important to understand how site discovery is affected when your organization uses a combination of root servers and domain controllers running Windows Server 2003 and Windows 2000 Server. These issues are described in the sections that follow.

**Differences in where site information is stored**

In Windows 2000 Server, DFS stores a copy of site information for root and link targets in the DFS metadata stored in the DFS object in Active Directory (for domain-based namespaces) and in the registry (for stand-alone namespaces). In environments that have domain-based DFS namespaces hosted on multiple root servers running Windows 2000 Server and Windows Server 2003, the following conditions can arise:

- Root servers and domain controllers running Windows Server 2003 ignore any site information in the DFS metadata. They obtain site information dynamically from Active Directory.

- Root servers and domain controllers running Windows 2000 Server retrieve the target's site information from the DFS metadata if the target was created by using Windows 2000 Server. However, if the target was created by using Windows Server 2003, no site information is stored for that target in the DFS metadata. As a result, Windows 2000 Server cannot determine the site of such a target. If this occurs, a referral from a Windows 2000 Server could lead the client computer to a random target, possibly outside of the client's site.

To enable root servers and domain controllers running Windows 2000 to sort targets based on site, you can use the **/UpdateWin2kStaticSiteTable** parameter in Dfsutil.exe to update the static site information for all root and link targets in the DFS object in Active Directory. This tool must be run each time a target server's site information is changed, such as after the target is moved to another site or after the target's site name changes.

When all root servers and domain controllers are running Windows Server 2003, you can use the **/PurgeWin2kStaticSiteTable** parameter in Dfsutil.exe to remove site information from the DFS object in Active Directory. We recommend using this method to decrease the size of the DFS object in Active Directory.

**How the restrictanonymous registry entry affects site discovery**

Root servers running Windows 2000 Server or Windows Server 2003 and that are not domain controllers cannot determine a DFS client computer's site when the **restrictanonymous** registry entry is set to 2 on domain controllers running Windows 2000 Server. (This registry entry is located at HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Control/Lsa.) As a result, these DFS root servers

sort targets in link referrals and root referrals randomly, regardless of the namespace type (stand-alone or domain-based), target selection method, or client operating system.

#### How Target Selection Works

When a client computer attempts to access a root or link in the namespace, a domain controller or root server provides a referral to the client. The referral contains a list of target servers, and this list is sorted according to the currently configured target selection method. These methods are described using the following figure as an example.

**Example of Client and Target Sites**




**Note**

- Out-of-date or misconfigured site information can cause DFS to sort referrals incorrectly. For more information about how DFS caches site information and when this information is refreshed, see "How Site Discovery Works" earlier in this section.

#### Default Target Selection

By default, DFS places target servers in the referral in the following order:

1. Targets in the same site as the client are listed in random order at the top of the referral.
2. Targets outside of the client's site are listed in random order.

Using the previous figure as an example, when default target selection is enabled, DFS places the target servers in the referral in the following order:

**Default Referral Order**




If no same-site target servers are available, the client computer is referred to a random target server no matter how expensive the connection is or how distant the target is. This is the default DFS behavior in Windows 2000 Server and Windows Server 2003.

#### Same-Site Target Selection

In this method, also called in-site target selection, the referral contains the targets that are in the same site as the client. These same-site targets are listed in random order. If no same-site targets exist, clients in that site do not receive a referral and cannot access that portion of the namespace.

Using the figure "Example of Client and Target Sites" as an example, when same-site target selection is enabled, DFS places the target servers in the referral in the following order:

**Same-Site Target Selection Referral Order**

You enable this setting on a DFS root or on individual links in the namespace by using Dfsutil.exe with the **/InSite** parameter. If you enable this setting on the root, root referrals and link referrals return only targets that are in the same site as the client. If no root or link targets exist in the client's site, then no referral is returned and the client cannot access that portion of the namespace. If this

setting is enabled on a link, and no link targets exist in the client's site, then no referral is returned and the client cannot access the link.

Same-site target selection works differently in Windows 2000 Server and Windows Server 2003. In Windows 2000 Server, if this setting is enabled on a root, the setting applies to all links but not the root itself. If a client attempts to access a namespace but no root targets exist in the client's site, the out-of-site root targets are returned in the root referral. Like Windows Server 2003, DFS in Windows 2000 Server does not return link referrals if no link targets exist in the client computer's site.

**Least Expensive Target Selection**
In this method, also called closest site selection or site costing, DFS uses the DsQuerySitesByCost API to retrieve minimum communication cost between two sites from Active Directory. The following figure illustrates sites and their associated site link costs.

**Example of Sites and Site Link Costs**




When least expensive target selection is enabled, DFS places targets in the referral in the following order:

1. Targets in the same site as the client are listed in random order at the top of the referral.

2. Targets outside of the client's site are listed in order of lowest cost to highest cost. Referrals with the same cost are grouped together and within each group the targets are listed in random order.

Using the previous figure as an example, when least expensive target selection is enabled, DFS places the targets in the referral in the following order:

**Least Expensive Target Selection Referral Order**




If the DsQuerySitesByCost API call fails, no cost information is returned to the root server. In this case, DFS assumes the maximum possible cost for that target.

Least expensive target selection works in both stand-alone and domain-based namespaces, as long as all root servers and all domain controllers are running Windows Server 2003. These requirements ensure that:

- All domain controllers can provide root referrals based on site cost.

- All root servers can provide root and link referrals based on site cost.

- The domain controllers acting as Intersite Topology Generators can calculate site cost information. (One domain controller in each site is selected as the Intersite Topology Generator, and the role is automatically given to a domain controller running Windows Server 2003, if one is available for the site.)

Least expensive target selection is not available in the following situations:

- When a stand-alone namespace is hosted on a root server that is not part of any domain.

- When a root server or domain controller is running Windows 2000 Server.

- The **Bridge all site links** option in Active Directory is disabled. (This option is available in the Active Directory Sites and Services snap-in.) Turning off **Bridge all site links** can affect the ability of DFS to sort targets in referrals in order of lowest cost. An Intersite Topology Generator that is running Windows Server 2003 relies on the **Bridge all site links** option being enabled to generate the intersite cost matrix that DFS requires for its site-costing functionality. If you turn off this option, you must create site links between any Active Directory sites for which you want DFS to calculate accurate site costs. Any sites that are not connected by site links will have the maximum possible cost. For more information about site link bridging, see "Active Directory Replication Topology Technical Reference.

In these situations, DFS uses the default target selection method. For more information about site costs, see "DNS Support for Active Directory Technical Reference."

**Situations in Which Clients Access Unexpected Targets**

Even when least expensive or same-site target selection is enabled, clients might still access high-cost or out-of-site targets. Also, when default target selection is enabled, a client computer might access a target server outside of its own site, even though a same-site target exists. These problems are typically caused by one of the following situations:

- The same-site target is temporarily unavailable (due to server or network issues), and the client fails over to the next target, which could be outside of the client's site.

- The out-of-site or high-cost target has incorrect site information that makes DFS interpret the target as being in the client's site. This problem can occur when subnets are not configured correctly. If no same-site targets exist and a client unexpectedly chooses a high-cost target, it might be caused by an incorrect site cost setting.

- The client's IP address is not in a subnet that is defined in Active Directory and so DFS cannot obtain site information about the client.

- The target's IP address is not in a subnet that is defined in Active Directory and so DFS cannot obtain site information about the target.

- The client is using a cached referral that contains outdated information about the site of the target server.

- Site information has changed, but the old site information is still cached on the root server or domain controller in the target site cache, client site cache, or site cost cache.

- The DFS object is not up-to-date when the root server polls a domain controller. This can be caused by Active Directory replication latency or failure.

- The **Bridge all site links** option is disabled, as described earlier.

- Site awareness is not working correctly because the **restrictanonymous** registry entry located at HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Control/Lsa is set to 2 on Windows 2000 domain controllers. If this registry entry is set to 2, DFS root servers that are not domain controllers (and are running either Windows 2000 Server or Windows Server 2003) randomly sort the targets in a referral, regardless of the namespace type (stand-alone or domain-based), target selection method, or client operating system.

**How DFS Works in Environments Without WINS**

The default behavior of DFS is to use NetBIOS names for all target servers in the namespace. This allows clients that support NetBIOS-

only name resolution to locate and connect to targets in a DFS namespace. Administrators can use NetBIOS names when specifying target names and those exact paths are added to the DFS metadata. For example, an administrator can specify a target \\FS1\Users, where FS1 is the NetBIOS name of a server whose DNS or FQDN name is FS1.contoso.com.

Organizations that do not use NetBIOS and WINS can still use DFS, but before setting up namespaces the administrators must create the **DFSDnsConfig** registry entry on all root servers and then restart the DFS service on all root servers. Administrators must then use the DNS names for when adding all targets to the namespace. When these steps are complete, the referrals will contain the DNS names of targets accessed by clients. If a namespace already exists, the administrator must perform the following steps:

1. Export the namespace to a data file by using Dfsutil.exe.
2. Delete the namespace.
3. In the exported data file, change the NetBIOS names to DNS names for all targets.
4. Recreate all root targets by using DNS names.
5. Import the updated data file.

For more information about the **DFSDnsConfig** registry entry and Dfsutil.exe, see "DFS Tools and Settings.

**The Referral Process**
The referral process describes the steps that client computers, root servers, and domain controllers perform when a client attempts to access a target server in a namespace. The following sections provide general and in-depth descriptions of the referral process for domain-based and stand-alone namespaces. These processes assume that all targets are available. For more information about how clients respond when a target is unavailable, see "What Happens During Failover" later in this section.

**Overview of the Referral Process**
The general steps that occur when a client accesses a domain-based or stand-alone namespace are described below. These processes assume the following:

- The client's domain cache contains the necessary domain name referrals and domain controller referrals.
- The client's referral cache does not contain existing referrals for the targets that the client is attempting to access.
- The first root target and link target in each referral are available.
- The link target is a shared folder on a physical server, not a DFS path in another namespace. The referral process that occurs when a link points to a DFS path is described in "Linking to Different DFS Namespaces" later in this section.

**Simplified Referral Process for Domain-based Namespaces**

1. A user attempts to access a link in a domain-based namespace, such as by typing \\Contoso.com\Public\Software in the **Run** dialog box.
2. The client computer sends a query to the active domain controller to discover a list of root targets for the domain-based namespace.
3. The domain controller returns a list of root targets defined for the requested namespace.
4. The client selects the first root target in the referral and sends a query to the root server for the requested link.
5. The root server constructs a list of link targets in the referral. The order in which the link targets are sorted depends on the target selection method. The root server sends the referral to the client.
6. The client establishes a connection to the first link target in the list.

**Simplified Referral Process for Domain-based Namespaces**




**Simplified Referral Process for Stand-Alone Namespaces**

1. A user attempts to access a DFS path, such as by typing \\Software\Public\Antivirus in the **Run** dialog box.

2. The DFS client computer sends a query to the root server that hosts the stand-alone namespace.

3. The root server returns a root referral to the client. The root referral contains the single root target.

4. The client sends a query to the root server for the requested link.

5. The root server constructs a list of link targets in the referral. The order in which the link targets are sorted depends on the target selection method. The root server sends the referral to the client.

6. The client establishes a connection to the first link target in the list.

**Simplified Referral Process for Stand-Alone Namespaces**




**Referral Process for Domain-based Namespaces**

The referral process for domain-based namespaces is as follows:

1. A user attempts to access a link in a domain-based namespace, such as by typing \\Contoso.com\Public\Software in the **Run** dialog box.

2. The client checks its domain cache for an existing domain name referral for the Contoso.com domain. If this referral is in the domain cache, the client proceeds to step 3. If no domain name referral is in the domain cache, the client connects to the IPC$ shared folder of the active domain controller in the context of the LocalSystem account and sends a domain name referral request with a null string. The domain controller returns the list of local and trusted domains to the client.

3. The client checks its referral cache for the longest prefix match from a previous referral. This might be a root referral (\\Contoso.com\Public) or a link referral (\\Contoso.com\Public\Software). If the client finds a valid entry, it goes to step 5 or 6 depending on the type of the referral. If not, the client continues to the next step.

4. The client checks its domain cache for an existing domain controller referral for the Contoso.com domain. If this referral is in the cache, the client proceeds to step 5. If no domain controller referral is in the domain cache, the client connects to the IPC$ shared folder of the active domain controller in the context of the LocalSystem account and sends a domain controller referral request containing the appropriate domain name (Contoso.com). The domain controller returns the list of domain controllers in the Contoso.com domain. The domain controllers in the clients site are at the top of the list. If least-expensive target selection is enabled, domain controllers outside of the targets site are sorted by lowest cost. If same-site target selection is enabled, DFS ignores this setting and lists the remaining domain controllers in random order.

5. The client checks its referral cache for an existing domain-based root referral. If this referral is in the cache, the client proceeds to

step 6. If no domain-based root referral exists in the referral cache, the client connects to the IPC$ shared folder of the active domain controller in the context of the LocalSystem account and requests a domain-based root referral. The domain controller determines the clients site and returns a list of root targets. By default, the root targets in the clients site are at the top of the list, followed by the remaining root targets in random order. If least-expensive target selection is enabled, the remaining root targets are ordered by lowest cost. If same-site target selection is enabled, only root servers in the clients site are listed in the referral.

6. The client chooses the first root target in the domain-based root referral. The client connects to the root server and navigates the subfolders under the root folder. When a client encounters a link folder, the root server sends a STATUS_PATH_NOT_COVERED status message to the client, indicating that this is a link folder that requires redirection.

7. The client checks its referral cache for an existing link referral. If this link referral is in the cache, the client computer connects to the first link target in the list. If the link referral is not in the cache, the client connects to the IPC$ shared folder of the root server in the context of the LocalSystem account and requests a link referral from the root server. The root server returns a list of link targets. At the top of the target list are the link targets that are in the same site as the client. The root server puts the remaining link targets in the target list using one of the following methods:

   - By default, the link targets outside of the clients site are put in random order.
   - If same-site target selection is enabled, the root server does not add out-of-site link targets to the referral.
   - If least-expensive target selection is enabled, the root server checks its site cost cache to determine whether cost information exists for the connection between the clients site and the link targets site. If the site cost data is not in the cache, the root server contacts the domain controller acting as the Intersite Topology Generator and uses the DsQuerySitesByCost API to retrieve site cost data. The root server then sorts the remaining link targets by lowest cost.

**Referral Process for Stand-Alone Namespaces**

The referral process for stand-alone namespaces is as follows:

1. A user attempts to access a DFS path, such as by typing \\Software\Public\Antivirus in the **Run** dialog box.

2. The client checks its domain cache and determines that the DFS path is not a domain-based DFS path (that is, Software is not the name of a domain).

3. The client checks its referral cache to determine if an existing stand-alone root referral is in the cache. If the root referral is in the cache, the client proceeds to Step 5. If the referral is not in the cache, the client connects to the IPC$ shared folder of the root server in the context of the LocalSystem account and sends a root referral request. The root server responds with a root referral so that all future requests are funneled through DFS.

4. The client connects to the root server and navigates the subfolders under the root folder. When a client encounters a link folder, the root server sends a STATUS_PATH_NOT_COVERED status message to the client, indicating that this is a link folder that requires redirection.

5. The client checks its referral cache for an existing link referral. If the link referral is in the referral cache, the client connects to the first link target in the list. If the link referral is not in the referral cache, the client connects to the IPC$ shared folder of the root server in the context of the LocalSystem account and requests a link referral from the root server. The root server returns a list of link targets. At the top of the target list are the link targets that are in the same site as the client. The root server puts the remaining link targets in the target list using one of the following methods:

   - By default, the link targets outside of the clients site are put in random order.
   - If same-site target selection is enabled, the root server does not add out-of-site link targets to the referral.
   - If least-expensive target selection is enabled, the server checks its site cost cache to determine whether cost information exists for the connection between the clients site and the link targets site. If the site cost data is not in the cache, the root server contacts the domain controller acting as the Intersite Topology Generator and uses the DsQuerySitesByCost API to retrieve site cost data. The root server then sorts the remaining link targets by lowest cost.

**How DFS Is Used During the Logon Process**

When a client logs on to a domain, the client contacts a domain controller and requests a special type of referral, known as a SYSVOL

or NETLOGON referral, to gain access to the scripts and policies stored in the SYSVOL and NETLOGON shared folders on domain controllers. (Specifically, the client requests SYSVOL and NETLOGON referrals from the active domain controller, which is preceded by a + under the appropriate domain in the client's domain cache.) These referrals map the paths *DomainName*\Sysvol and *DomainName*\Netlogon to a list of SYSVOL and NETLOGON shared folders for all domain controllers. Clients store these referrals in their referral cache. To see examples of these referrals, see the figure "Sample Referral Cache Output" earlier in this section.

DFS objects do not exist in Active Directory for these shared folders; instead, the domain controller simply recognizes that it is responsible for these paths and responds to queries of the form *DomainName*\Sysvol and *DomainName*\Netlogon. These referrals are distinct in that although the targets are hosted on a domain controller, the targets are not roots themselves and do not appear in any of the DFS tools, nor can the SYSVOL and NETLOGON shared folders host links.

Domain controllers generate SYSVOL and NETLOGON referrals each time a client requests a referral. By default, the list of domain controllers listed in a SYSVOL or NETLOGON referral are sorted as follows:

1. All domain controllers in the client's site are grouped in random order at the top of the list.

2. Domain controllers outside of the client's site are listed in random order.

It is possible to configure DFS to sort the domain controllers outside of the client's site in order of lowest cost. You can enable this feature by adding the **SiteCostedReferrals** registry entry on each domain controller and then restarting the DFS service on each domain controller. The DFS service then obtains site cost information for all domain controllers and stores this information in its site cost cache.

**Linking to Different DFS Namespaces**

When you create a link and specify the link target, you can use any UNC path, including a path to another namespace. For example, the Apps link in a stand-alone namespace \\Dev\Team\Apps might have a link target of \\Contoso.com\Public\Software, which is a link within a domain-based namespace. These types of links are often called interlinks, and DFS namespaces that point to other namespaces are often called cascaded namespaces.

Linking to different namespaces requires that certain rules be followed so that clients can fail over correctly if one of the targets is unavailable. The rules are as follows:

- DFS clients can perform a maximum of eight hops through other DFS namespaces. This means that a given link can point to a DFS path, and that path can point to another DFS path, and so on, up to a maximum of eight times.

- A link that points to a different namespace can correspond to only one of the following options:

    - A link target can point to one or more stand-alone DFS paths anywhere in the stand-alone DFS namespace. For example, a link target can be \\Software\Public (a DFS path with just the root) or \\Software\Public\Applications (a DFS path with a root and link).

    - A link target can point to a single domain-based DFS path anywhere in the domain-based DFS namespace. For example, a link target can be \\Contoso.com\Public or \\Contoso.com\Public\Users.

The referral process that occurs when a client accesses a link whose target is a DFS path is as follows:

1. A user attempts to access a link that points to another namespace.

2. If necessary, the DFS client requests domain and domain controller referrals (for domain-based namespaces) and root referrals (for both types of namespaces) by following the steps described in "The Referral Process" earlier in this section.

3. The root server provides a link referral. Because the link points to another namespace, the link target in the referral shows the DFS path, not the physical path.

4. The client attempts to access the DFS path provided in the link referral and requests the necessary referrals.

5. The client receives the referrals and accesses the intended target.

The following figure illustrates a client's referral cache after the client attempts to access the Apps link within a stand-alone namespace (\\Dev\Team\Apps) that points to a link in a domain-based namespace (\\Contoso.com\Public\Software). Note that the link target in the link referral for Apps is listed as Contoso.com\Public\Software and that the type is shown as type 0x10, OUTSIDE_MY_DOM. This type is used when a link points to a path in a domain-based namespace.

**DFS Client Referral Cache**

Root referrel for stand-alone namespace \\Dev\Team

Link referral for link named "Apps"

Root referral for \\Contoso.com\Public

Link referral for \\Contoso.com\Public\Software

```
Entry: \dev\team
ShortEntry: \dev\team
Expires in 300 seconds
UseCount: 0 Type:0x81 ( REFERRAL_SVC DFS )
   0:[\dev\team] State:0x19 ( ACTIVE )

Entry: \dev\team\apps
ShortEntry: \dev\team\apps
Expires in 1800 seconds
UseCount: 0 Type:0x10 ( OUTSIDE_MY_DOM )
   0:[\contoso.com\public\software] State:0x21 ( )

Entry: \contoso.com\public
ShortEntry: \contoso.com\public
Expires in 300 seconds
UseCount: 0 Type:0x81 ( REFERRAL_SVC DFS )
   0:[\fs-03\Public] State:0x19 ( ACTIVE )
   1:[\fs-02\Public] State:0x09 ( )
   2:[\fs-01\Public] State:0x09 ( )

Entry: \contoso.com\public\Software
ShortEntry: \contoso.com\public\Software
Expires in 1800 seconds
UseCount: 0 Type:0x1 ( DFS )
   0:[\dc-01\SW] State:0x31 ( ACTIVE )
   1:[\fs-03\SW] State:0x21 ( )
   2:[\fs-01\SW] State:0x21 ( )
   3:[\fs-02\SW] State:0x21 ( )
```

**What Happens During Failover**

DFS failover is the process in which clients attempt to access another target in a referral after one of the targets fails or is no longer part of the namespace. Failover can occur in five scenarios:

- A client attempts to access the currently active target in a referral, but the target is unavailable.
- A client attempts to access a domain controller to request a domain referral or a domain-based root referral, but the currently active domain controller in the referral is unavailable.
- The DFS client has a session enabled with a target, and the target server fails.
- The DFS client has a session enabled with a target, and the hard disk on the server fails or the shared folder is disabled.
- A referral expires for a DFS client running Microsoft Windows XP with Service Pack 2 (SP2) or Windows Server 2003 with Service Pack 1 (SP1), and the previous active target is not in the new referral.

When reviewing the different types of failovers, note the following:

- Clients must access a domain-based namespace by using the format *DomainName**RootName*. If a client accesses a domain-based namespace directly on the root server (*RootServer**RootName*), root target failover does not occur.
- DFS failover is only performed when a client opens a file or folder. If a client has files or folders open and attempts to read or write to them when the target server is unavailable, the application will receive a failure on that operation.

The following scenarios describe different types of failures and what is involved in the failover process in each case.

**The active target in a referral is unavailable**

When a DFS client accesses a namespace, the client always chooses the first target in the target list in the referral. If the first target is available, a session setup is performed and client credentials are passed to the server, unless a prior connection already exists with the selected target. If the selected target cannot be accessed, the client attempts to access the next target in the list, and so on, until all targets are exhausted. If all targets are unavailable, then the client cannot access that portion of the DFS namespace.

**Note**

- When DFS attempts to access a target in the target list, and the first target or subsequent targets are not available, DFS does not attempt to access those targets again during later failovers until it has reached the bottom of the list.

**A domain controller fails**

When a client attempts to request a domain-based root referral or a domain name referral from a domain controller, and the first domain controller (for a given domain) in the client's domain cache is unavailable, the client will fail over to the next domain controller in the cache. If all domain controllers for the given domain are unavailable, the client cannot access the domain-based DFS namespace for that domain.

**A client is accessing the target server when it becomes unavailable**
In the following scenario, assume that a client application is accessing a target in the namespace when the server hosting the physical shared folder loses power or drops off the network. If the referral contains more than one target, the client will attempt a failover the next time the client attempts to open any file or folder in that portion of the DFS namespace. The amount of time it takes to detect that the hosting computer is offline depends on the protocol that the client is using. Many protocols, such as TCP/IP, account for slow and loosely connected WAN links, and, as such, might have retry counts of up to two minutes before the protocol times out. After the protocol returns an error, DFS returns a failure to the application. The next time the application attempts to open a file or folder in that portion of the namespace, DFS will fail over to the next target. If the DFS client reaches the bottom of the list, it will start again at the top until it has attempted to connect to all targets.

**A client is accessing the target server when a hard drive failure occurs or folder sharing is stopped**
In the following scenario, assume that a client is accessing a target in the namespace when the server hosting the physical shared folder loses the hard disk containing the folder, or an administrator stops sharing the folder. If the referral contains more than one target, the client will attempt a failover as in the previous scenario. However, in this scenario, the server hosting the shared folder is still responding to the client request; therefore, the failover to a new target is much faster than the previous scenario because DFS does not rely on the protocol to detect a failure.

**A referral expires for a DFS client running Windows XP with SP2 or Windows Server 2003 with SP1**
The behavior of the Time to Live value has changed for DFS clients running Microsoft Windows XP with SP2 or Windows Server 2003 with SP1. These clients do not renew the Time to Live value each time they access a target using a cached referral. As a result, the referral expires after the Time to Live value lapses, regardless of whether the client has accessed or is currently accessing the target. In terms of DFS failover, this new behavior can cause the DFS client to fail over to a new target if the previously active target is not in the new referral, such as when the namespace is changed to remove the target that the client had accessed.

For more information about the Time to Live value, see “How DFS Discovers Changes to the Namespace” later in this section.

**Root Folder and Link Folder Creation Process**
When you first create a DFS namespace, you specify a shared folder to use as the root folder. When you add links to the namespace, DFS creates a special folder, called a link folder, under the root folder to represent each link. The hierarchy of the root folder and link folders on the root server’s local storage matches the actual namespace. If you create a link by using the format *FolderName\FolderName\LinkName*, DFS creates the hierarchy of empty folders and then creates the link folder. Root and link folders are illustrated in the following figure.

### Roots, Links, Root Folders, and Link Folders




When you use a DFS tool to create a link, DFS checks whether a folder that has the same name as the link appears under the root. DFS then takes one of the following actions:

- If no folder exists, DFS creates the link folder and sets a reparse point on the folder.
- If an empty folder with the same name as the link exists, DFS sets a reparse point on that folder.
- If a non-empty folder with the same name as the link exists, DFS renames the non-empty folder to DFS.*GUIDLinkName*, creates a new link folder, and sets a reparse point on the folder. An example folder name is DFS.cf13c05f-5c10-4879-9acb-04ced8f46c7aTemplates, where cf13c05f-5c10-4879-9acb-04ced8f46c7a is the *GUID* and Templates is the *LinkName*.

The reparse points set on the link folders prevent data from being stored in the link folders and prevent administrators from deleting link folders by using traditional methods, such as by using Windows Explorer or the command prompt. If an administrator needs to remove a link folder, the administrator can use Dfsutil.exe to remove all DFS reparse points under a specified root folder and then delete the link by using the DFS tools. When the DFS service is started, it checks that the link folders are configured correctly and recreates link folders and reparse points as necessary. It also cleans up stale reparse points that are no longer part of any namespace.

When you add a new root target to an existing domain-based namespace, you specify an existing shared folder or create a new shared folder to use as the root. DFS creates the link folders on the new root target after the root server obtains the DFS metadata by polling the PDC emulator master (or the closest domain controller if root scalability mode is enabled). When a root server obtains the DFS metadata after polling and detects that a link has been removed, the root server removes the reparse point and the link folder for the deleted link.

**How DFS Discovers Changes to the Namespace**
To ensure that clients have an up-to-date view of the namespace, DFS must maintain a consistent version of the namespace on all root

servers. A change to the namespace can include the following types of changes:

- Adding, changing, or removing root targets, links, and link targets.
- Changing the settings on a namespace, including the Time to Live for roots and links, target selection method (for example, changing to least expensive or same-site target selection), enabling root scalability mode, and so forth.
- Enabling or disabling a referral to a link target.

The following sections describe how changes are discovered in stand-alone and domain-based namespaces.

**Note**

- These sections assume that site information is up-to-date in the client site cache, target site cache, and site cost cache on root servers and domain controllers. If these caches contain outdated site information, DFS might not sort targets in referrals correctly. Client and target site information can be out of date for up to 24 hours, and site cost information can be out-of-date up to 12 hours. For more information about how site information is updated, see "How Site Discovery Works" earlier in this section.

#### How Changes Are Discovered in Stand-Alone Namespaces

Changes made to a stand-alone namespace are updated immediately in the DFS metadata in the root server's registry and in the root server's in-memory DFS metadata cache. Therefore, stand-alone root servers typically provide up-to-date root referrals and link referrals.

Clients do not request new root and link referrals until the referrals expire or are cleared from their cache. You can use the following table to determine the maximum length of time that out-of-date referrals are used in stand-alone namespaces.

**When Referrals to Stand-Alone Namespaces Expire**

| Where Referrals Are Cached | Referral Type | Default Expiration Time | Where Expiration Time Is Stored |
|---|---|---|---|
| DFS clients | Stand-alone root referral | 5 minutes | Time to Live of root (configured using the DFS tools) |
| DFS clients | Link referral | 30 minutes | Time to Live of link (configured using the DFS tools) |

For DFS clients that are not running Windows XP with Service Pack 2 (SP2) or Windows Server 2003 with SP1, the Time to Live for a referral determines the earliest time that a client will request a new referral, but only if the existing referral expires before it is accessed again. Clients that use a cached referral will renew the Time to Live of the referral and thus use the referral indefinitely until the client's referral cache is cleared or the client is restarted.

This behavior has changed for clients running Windows XP with SP2 or Windows Server 2003 with SP1. Specifically, the Time to Live value is not reset each time a client accesses a target using a cached referral. Instead, the referral expires after the Time to Live value lapses. This change has several effects:

- Clients running Windows XP with SP2 or Windows Server 2003 with SP1 will request referrals more frequently than other DFS clients, which can cause moderately increased load on the stand-alone DFS root server.
- Because they request new referrals more frequently, clients running Windows XP with SP2 or Windows Server 2003 with SP1 will discover namespace updates more quickly than other DFS clients.

#### How Changes Are Discovered in Domain-based Namespaces

The update process for domain-based namespaces is more complex because there are multiple locations of the DFS metadata — the DFS metadata is stored in the DFS object in Active Directory on each domain controller and is stored in a memory cache on multiple root servers. To ensure that all root servers discover changes to the namespace, DFS uses two methods of keeping root servers up to date:

- Root server notification
- Periodic polling of the PDC emulator master

These methods are described in the sections that follow.

**Root server notification**

To maintain a consistent namespace across root servers, DFS depends on the domain controller acting as the primary domain controller (PDC) emulator master to be the gatekeeper for all updates to the namespace. Having all root servers running Windows Server 2003 poll the PDC emulator master after the namespace changes ensures that root servers can obtain the updated DFS metadata relatively quickly without needing to wait for Active Directory replication to replicate the DFS object to all domain controllers.

Root server notification in Windows Server 2003 works as follows:

1. The DFS tool that an administrator uses to make the change contacts the closest root server for that namespace.

2. The root server updates the DFS metadata in the DFS object in Active Directory on the PDC emulator master.

3. The root server sends a change notification message to all other root servers hosting the namespace.

4. When root servers running Windows Server 2003 receive the change notification message, they poll the PDC emulator master to obtain an updated version of the DFS metadata. (Root servers running Windows 2000 Server ignore change notification messages and poll the PDC emulator master every hour.) If a Windows Server 2003 root server does not receive the change notification message, for example, if the root server is temporarily offline, it will poll the PDC emulator master at the next polling interval.

This update process can cause an increased load on the PDC emulator master in namespaces that have more than 16 root targets and in namespaces that change frequently. In these environments, we recommend enabling root scalability mode. For more information, see "How Root Scalability Mode Works" later in this section.

**Periodic polling of the PDC emulator master**
By default, root servers poll the PDC emulator master during the following events:

- When the root server starts up or the DFS service is restarted.

- When any DFS API is called against the root server. For example, using the Dfscmd.exe command-line tool with the **/**view *****dfsname******dfsshare* **/**fullparameter causes the root server to poll the PDC emulator master. (This command calls the NetDfsEnum API.)

- At the interval specified in the **SyncIntervalInSeconds** registry entry (on the root server), which is one hour by default. Organizations that have stable namespaces can increase this value so that root servers contact the PDC emulator master less frequently.

If the PDC emulator master is unavailable when a root server attempts polling, the root server will poll its closest domain controller. If the DFS object has replicated from the PDC emulator master to the domain controller closest to the root server, the root server receives the up-to-date version of the DFS object at the next polling.

**When changes to domain-based namespaces are reflected in referrals**
Even if the DFS metadata (stored in a memory cache on root servers and present in the DFS object in Active Directory on domain controllers) is identical and up-to-date, DFS root servers and domain controllers can provide outdated domain-based root referrals until those referrals expire or are cleared from the domain-based root referral cache. On the other hand, link referrals are up-to-date as soon as the root server receives an updated copy of the DFS object. After the root servers and domain controllers begin providing up-to-date referrals, clients might continue to use outdated referrals until the referrals expire or are flushed from the client's referral cache.

You can use the following table to determine the maximum length of time that out-of-date referrals are used in domain-based namespaces. This table assumes that the root servers and domain controllers have received the most recent version of the DFS metadata.

**When Referrals to Domain-based Namespaces Expire**

| Where Referrals Are Cached | Referral Type | Default Expiration Time | Where Expiration Time Is Stored |
|---|---|---|---|
| Domain controller | Domain-based root referral | 15 minutes | **RootReferralTimeoutInSeconds** registry entry |
| Root server | Domain-based root referral | 15 minutes | **RootReferralTimeoutInSeconds** registry entry |

| | | | |
|---|---|---|---|
| DFS client | Domain-based root referral | 5 minutes | Time to Live of root (configured using the DFS tools) |
| DFS client | Link referral | 30 minutes | Time to Live of link (configured using the DFS tools) |

Based on the default values in the previous table, domain controllers and root servers can provide updated domain-based root referrals after 15 minutes, assuming that the domain controller or root server has received the most recent version of the DFS object in Active Directory. Any latency in Active Directory replication could cause additional delay.

If settings on a link change, root servers can provide up-to-date link referrals as soon as the root server has obtained an updated version of the DFS metadata (in the DFS object) from the PDC emulator master. As described earlier, root servers contact the PDC emulator master soon after they receive the change notification message. If root servers are connected by high-speed network connections to each other and to the PDC emulator master, any changes to links are reflected in referrals almost immediately. A root server will provide out-of-date information in a link referral if any of the following conditions apply:

- The root server does not receive a change notification message due to network problems or other issues. In this case, the root server will contact the PDC emulator to obtain the updated DFS metadata at the polling interval described earlier. (Note that root servers running Windows 2000 Server ignore change notification messages and poll the PDC emulator master every hour.)

- Root scalability mode is enabled. In this case, root servers obtain the DFS metadata from the closest domain controller every hour, so you must factor in an additional one-hour delay plus the time required for Active Directory replication to replicate the updated DFS object to each domain controller.

For DFS clients that are not running Windows XP with SP2 or Windows Server 2003 with SP1, the Time to Live for a referral determines the earliest time that a client will request a new referral, but only if the existing referral expires before it is accessed again. Clients that use a cached referral will renew the Time to Live of the referral and thus use the referral indefinitely until the client's referral cache is cleared or the client is restarted.

This behavior has changed for clients running Windows XP with SP2 or Windows Server 2003 with SP1. Specifically, the Time to Live value is not reset each time a client accesses a target using a cached referral. Instead, the referral expires after the Time to Live value lapses. This change has several effects:

- Clients running Windows XP with SP2 or Windows Server 2003 with SP1 will request referrals more frequently than other DFS clients, which can cause moderately increased load on the domain-based DFS root servers and domain controllers.

- Because they request new referrals more frequently, clients running Windows XP with SP2 or Windows Server 2003 with SP1 will discover namespace updates more quickly than other DFS clients.

**How Root Scalability Mode Works**

Root scalability mode allows organizations to use more than the recommended 16 root targets per domain-based root. When root scalability mode is enabled, DFS root servers do not send change notification messages to other root servers when the namespace changes or poll the PDC emulator master every hour. Instead, they poll their closest domain controller every hour to discover updates to the namespace.

Root scalability mode reduces network traffic to the PDC emulator master at the expense of faster updates to all root servers. Updates are still made to the DFS object in Active Directory on the PDC emulator master, but root servers do not discover those changes until the updated DFS object replicates (using Active Directory replication) to the closest domain controller for each root server. After root scalability mode is enabled, clients might have an inconsistent view of the namespace until the most up-to-date version of the DFS object is replicated to all domain controllers and the DFS root servers poll a domain controller to discover the changes.

Root scalability mode does not entirely eliminate traffic to the PDC emulator master. The PDC emulator master is still used as follows:

- As described earlier, changes to the namespace are still made on the PDC emulator master.

- If the closest domain controller is the PDC emulator master, then the root server will poll the PDC emulator master every hour.

- When the DFS service starts on the root server (when the server reboots, for example), the DFS service polls the PDC emulator master. On the next polling interval, the root server polls the closest domain controller to get an initial copy of the DFS metadata. This behavior occurs because the root scalability setting is stored in the DFS object in Active Directory, so when a root server first starts up, it polls the PDC emulator master, determines that root scalability mode is enabled, and begins polling the domain controller at regular intervals.

- If a domain-based root has root targets running Windows 2000 Server and Windows Server 2003, root servers running

Windows 2000 Server poll the PDC emulator master every hour.

**What Happens When You Check Target Status**

You can use the Distributed File System snap-in to check the status of roots, links, and root and link targets. The snap-in verifies the namespace elements as follows:

- When you check the status of a root or link, the snap-in calls the GetFileAttributes API on each target to verify that they can be accessed. If this succeeds, the snap-in calls the GetFileAttributes API on the root folder or link folder to verify that it can be accessed.

- When you check the status of a root target or link target, the snap-in calls the GetFileAttributes API on the target to verify that the target can be accessed.

The following tables list the possible status indicators for DFS roots, links, and targets. These status indicators remain on the namespace element until you restart the snap-in or choose the **Refresh** command on the **Action** menu.

**DFS Status Flags for Roots and Links**

| Root or Link | Description |
|---|---|
| Folder icon (yellow) | No known status. |
| Folder icon with green check mark | All targets are reachable on the network and their referrals are enabled. |
| Folder icon with yellow exclamation point | Indicates one of the following conditions:<br>• All corresponding targets are reachable on the network but one or more targets has its referral disabled.<br>• Some corresponding targets are not reachable on the network. |
| Folder icon with red X | Indicates one of the following conditions<br>• All corresponding targets are not reachable on the network.<br>• All corresponding targets have their referrals disabled.<br>• The root path is unreachable. |

**DFS Status Flags for Targets**

| Target | Description |
|---|---|
| Folder icon (yellow) | No known status. |
| Folder icon with green check mark | Target is reachable on the network. |
| Folder icon with red X | Target is not reachable on the network. |

**How DFS Works on Server Clusters**

A server cluster is a group of individual computer systems, called nodes, working together cooperatively to provide increased computing power and continuous availability of business-critical applications or resources. A cluster can be configured so that the workload is distributed among the group, and if one of the nodes fails, another node automatically assumes its duties.

Using DFS in conjunction with server clusters can result in the increased availability of stand-alone roots in an organization. DFS is supported on server clusters as follows:

- **Link targets.** Link targets can consist of any File Share resource on shared cluster storage and any shared folder on a cluster node's local (non-shared) storage. The fact that a link target exists on a server cluster is transparent to DFS, so you can select these link targets as you would link targets on non-clustered file servers. If a link target is hosted on a cluster node's local storage, and the node fails or is taken offline, the link target is not available.

- **Stand-alone roots.** Stand-alone roots are supported on shared cluster storage and on a cluster node's local storage. When a

stand-alone root is created on shared cluster storage, the root is available as long as one of the nodes in the server cluster is available. If a stand-alone root is hosted on a cluster node's local storage, and the node fails or is taken offline, the namespace is not available.

- **Domain-based roots.** Domain-based roots are not supported on shared cluster storage. Domain-based roots are supported only on a cluster node's local storage. If a domain-based root is hosted on a cluster node's local storage, and the node fails or is taken offline, the namespace is not available unless there is another root target (either on another node's local storage or on another server in the domain).

If you plan to create stand-alone roots on shared cluster storage, it is important to understand the resource dependencies for DFS roots and the processes related to the creation and failover of DFS roots. The following sections describe these issues. For more information about how server clusters work in general, see "Server Clusters Technical Reference."

**Resource Dependencies of Stand-Alone Roots**

Using Cluster Administrator, you choose the File Share resource type to create a stand-alone DFS root on a server cluster's shared storage. The resource group that hosts the DFS root must contain the following resources:

- File Share resource (select the **DFS Root** option)
- Network Name resource (this is the virtual server name that is used by the clients to access this DFS root)
- IP Address resource (for the network name)
- Physical Disk resource (hosts the DFS root)

The dependencies for these resources must be configured correctly so that DFS can fail over if the node hosting the DFS root fails. The dependency tree is illustrated in the following figure.

**File Share Resource Dependency Tree**




**How Stand-Alone Roots on Server Clusters Work**

When the DFS service starts on a cluster node, the DFS service uses the Cluster APIs to determine whether the current node is part of a cluster. This process is important for two reasons

- Stand-alone roots on cluster nodes can be put in standby mode by using the NetDfsSetInfo API. When a root is put on this state, the DFS service does not provide referrals for it. Note that only stand-alone roots on cluster nodes can be put on this state. This state is not supported for normal stand-alone or domain-based roots. When the root is in standby mode, the operations of removing the root and bringing the cluster to online mode are supported.
- The referrals for stand-alone roots on cluster nodes must contain the respective virtual server name rather than the name of the node that is currently hosting the root.

When a stand-alone root is first created on shared cluster storage, the DFS service (on the node that hosts the resource group that contains the DFS root resource) creates the DFS metadata for the root in a registry entry at KEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfs\Roots\Standalone. The Cluster service propagates this registry information to the cluster database on the quorum resource, which is a resource that maintains the configuration data that is necessary for the recovery of the resource group. When a link is added or removed, or when the settings on the namespace are changed, the DFS service on the hosting node makes these changes in the DFS metadata in the registry. These changes to the registry are propagated to the cluster database. The physical folder acting as the root folder and any link folders are maintained on the shared cluster storage.

DFS registers a callback function with the Cluster service when a root is first created on a node and when the root is created on another node during cluster failover. The callback function is important because it allows DFS to determine whether the root is created on shared cluster storage. When DFS registers a callback function, the Cluster service enumerates all the File Share resources in the

cluster, and DFS determines if the root it is currently creating is a cluster resource.

If the node that hosts the DFS root fails, the following actions occur on another node (for example, Node 2):

1. The Cluster service reads the root's registry settings from the quorum resource and places these settings in the registry (of Node 2).

2. The Cluster service brings the resources online in the order of dependency.

3. The Cluster service creates an SMB shared folder on the root folder in the shared cluster storage.

4. The Cluster service informs DFS about the new root information in the registry and starts the DFS service if it is not already running.

5. DFS creates the internal data structures necessary for the DFS service to provide referrals for this root.

If a client attempts to access the root between steps 3 and 4, the client cannot access links under the root. This problem occurs because the DFS root is not yet created, so the SMB redirector claims the UNC path, causing the client's MUP cache to contain entries indicating that the resource is an SMB shared folder, and not a DFS root. Clearing the client's MUP cache by using Dfsutil.exe's **/PurgeMupCache** parameter will correct this problem.

**How DFS Works with Offline Files**
By using the Offline Files feature, you can make shared folders that correspond to DFS root targets and link targets available offline to client computers running Windows XP Professional or Windows Server 2003. The following sections describe two important interactions between Offline Files and DFS. These sections assume that the client computers are running either Windows XP Professional or Windows Server 2003.

**Offline Settings on Root Targets and Link Targets**
The Offline Files feature provides three settings for shared folders:

- Only the files and programs that users specify will be available offline.

- All files and programs on the share will be available offline.

- No files or programs on the share will be available offline.

Offline settings for root targets and link targets are set on the individual shared folders that are specified as targets. Although link targets appear subordinate to root targets in the namespace, link targets do not inherit or otherwise use the offline settings set on root targets. In addition, if a root or link has multiple targets, and those targets have different offline settings, the client will use whatever settings are applied to the target to which the client connects. Therefore, it is important for administrators to apply offline settings consistently to all targets for a DFS root or link.

**How the Offline Files Feature Interprets DFS Paths**
The Offline Files feature does not distinguish DFS paths from UNC paths. This can cause the client to interpret the entire namespace as unavailable if a target is down when a client attempts to access it. For example, if \\Contoso.com\Public is a domain-based root with several root targets and numerous links, the Offline Files feature interprets this namespace as a single server named \\Contoso.com. If a client is accessing or attempts to access a target in the \\Contoso.com\Public namespace, and the target is unavailable, the client interprets the entire namespace as unavailable and will attempt to open a user's locally cached files (if they exist). The client cannot access any target in the namespace until the target comes back online. The client will check every 10 minutes to detect whether the target has come back online. Alternatively, the user can use the Synchronization Manager to attempt to synchronize the offline files with those stored on the network. The user can also use Synchronization Manager to go back online without synchronizing changes.

**DFS-Related Processes on Clients**
The following sections describe how DFS clients running Windows XP and Windows 2000 update the memory caches described in "DFS Physical Structures and Caches on DFS Clients" earlier in this section.

**How Clients Update the Domain Cache**
DFS clients periodically discover new domains in the local forest and in trusted forests. This discovery process, which occurs every 15 minutes by default, runs against a domain controller from the domain that hosts the client's computer account. To avoid real-time queries to domain controllers in the domain, the referrals received during the discovery process are stored in a special table, called the domain cache or SPC cache. As a result of this process, clients can more quickly distinguish queries for fully qualified domain names from fully qualified computer names.

The process by which DFS clients update the domain cache is as follows:

1. When a DFS client starts up, the Workstation service calls the DFS client driver (Mup.sys) with information about the NetBIOS

and DNS names of the domain to which the client is joined. The Workstation service also provides the name of one domain controller in that domain. For this example, assume that the domain controller is called ClientDC.

2. The client saves the domain information and opens the IPC$ connection (\\ClientDC\IPC$) to the domain controller and sends a request with an empty name to request a list of trusted domains (in the client's forest and in trusted forests).

   If the client cannot connect to the domain controller or the domain controller fails in its response, the failure is sent to the Workstation service. This causes the Workstation service to find another domain controller for the client's domain and to pass that information to the DFS client. If this also fails, the Workstation service stops attempting to locate a domain controller. Regardless of whether previous attempts have succeeded, the Workstation service performs this process every 15 minutes (to ensure that DFS has a current working domain controller). The **DfsDcNameDelay** registry entry on the DFS client specifies the 15-minute interval.

3. If the referral request for the domain name information is successful, the client fills in its domain cache. This cache contains the domain information returned to the client as a response to the empty referral request. This information is updated every time the Workstation service calls the DFS client driver.

4. Any UNC request on the client comes to DFS first before any of the network redirectors. DFS checks its domain cache to determine whether the name is a domain name. If it is (meaning that the UNC path request is for a domain-based namespace), DFS checks its referral cache to detect whether it already has a referral to this namespace. If no referral exists, and the name is a domain name, the client proceeds to the next step.

5. The client determines whether the domain name is expanded. An expanded domain name is one in which the domain controllers are cached for the domain. If the domain name is not expanded, the client requests a domain controller referral from ClientDC.

DFS clients request new domain name referrals and domain controller referrals every 15 minutes, by default. You can configure this interval by using the **DfsDcNameDelay** registry entry on the client. Domain name referrals and domain controller referrals remain in the client's domain cache until the client receives updated referrals at the next interval or the client computer is restarted.

**Note**

- Without the proper domain name referrals and domain controller referrals, clients cannot access domain-based namespaces. This problem can occur when clients are logged on using cached credentials. DNS problems can also cause a client to be missing referrals.

If an organization has a large number of trusted domains and forests, it is possible that clients will not be able to access all domain-based namespaces within the trusted domains and forests. If a client can access a link target in another trusted domain or trusted forest by using the target's UNC path, the client can also access the link target by using its DFS path, but only if the list of domains fits into the client's buffer. By default, DFS clients send a 4-KB (2,048 Unicode character) buffer to a domain controller when requesting domain name referrals. If the list of domains is too large to fit into the 4-KB buffer, DFS clients automatically increase their buffer size to accept the list of domains, up to a maximum of 56 KB.

If the list of domains exceeds 56 KB, DFS puts as many domains in the buffer as it can until the buffer reaches 56 KB. DFS then writes an entry with the ID 14536 in the system event log of the domain controller that provided the domain referral. When populating the buffer, DFS gives preference to local and explicitly trusted domains by filling the buffer with their names first. Consequently, by creating explicit trust relationships with domains that host important DFS namespaces, you can minimize the possibility that those domain names might be dropped from the list that is returned to the client.

**Note**

- To make sure that clients can access link targets in other trusted domains or trusted forests, use DNS names for all link targets and configure DFS to use fully qualified domain names in referrals. For more information about using DNS names, see "How DFS Works in Environments Without WINS" earlier in this section.

**How Clients Update the MUP Cache**

When a Windows-based client attempts to access a UNC path, Windows sends the request to the multiple UNC provider (Mup.sys) to identify which redirector should handle the request. When Mup.sys receives the request, Mup.sys determines whether the path is in a DFS namespace. If it is, Mup.sys passes the request to DFS. If the path is a shared folder or WebDAV folder, for example, Mup.sys checks its internal cache, known as the MUP cache, to see whether the connection had been made previously. Mup.sys then sends the request to each redirector that handles each request synchronously and attempts to identify a resource on the network that matches the request. After all redirectors return, Mup.sys chooses the redirector that the application will use (based on response and priority).

**How Clients Update the Referral Cache**

When users access DFS roots and links, the client sends a referral request to the DFS server (either a domain controller or root server, depending on the type of referral requested). The DFS server puts the referral response in a 4 KB buffer (2,048 Unicode characters) and returns the buffer to the client. Root servers and domain controllers send referrals up to 4 KB if there is room for at least one target in the referral. If the paths of the targets are so long that the 4-KB buffer cannot hold even one target, DFS makes the client request a larger response limit (up to 56 KB).

After a client receives a referral, the client stores the referral in the referral cache and accesses the first available target (either a particular root target or link target) in the list of targets provided in the referral. The client will continue to access that target until one of the following occurs:

- The client computer is restarted, which clears the referral cache.
- The user manually clears the referral cache by using Dfsutil.exe with the **/pktflush** parameter.
- The target server becomes unavailable.
- The Time to Live value for the root or link referral expires. Windows Server 2003 uses a default Time to Live value of 300 seconds (5 minutes) for root referrals and 1800 seconds (30 minutes) for link referrals.

After any of these events, the client will obtain a new referral the next time it attempts to access the target. If the client continues to access the target within the Time to Live value, the Time to Live value is renewed and the client never requests a new referral. If the Time to Live value is not renewed, it is decremented by 4 minutes (240 seconds) until it reaches zero. If the Time to Live value is set less than 240 seconds, the referral will not expire for 240 seconds.

**DFS Protocols**

DFS uses the Common Internet File System (CIFS) for communication between DFS clients, root servers, and domain controllers. CIFS is an extension of the Server Message Block (SMB) file sharing protocol. Examining network captures of CIFS communications between a DFS client and server is helpful for understanding and troubleshooting DFS processes. The following sections illustrate two network captures created when a client receives different types of referrals.

**Network Capture of a Client Receiving Domain Name Referrals and Domain Controller Referrals**

The following network capture (taken using Microsoft Network Monitor) shows what happens when a client first receives domain name referrals and domain controller referrals. This network capture was taken on a domain controller (named DC-01) after the client (named DFSCLIENT) was restarted. Only SMB frames are shown. Note that this network capture is an example of client and domain controller communication and might differ from network captures taken in your environment.

### Network Capture of a Client Receiving Domain Name Referrals and Domain Controller Referrals

| Frame | Protocol | Description | Src Other Addr | Dst Other Addr |
|---|---|---|---|---|
| 13 | SMB | C negotiate, Dialect = NT LM 0.12 | DFSCLIENT | DC-01.CONTOSO.CO |
| 14 | SMB | R negotiate, Dialect # = 5 | DC-01.CONTOSO.CO | DFSCLIENT |
| 15 | SMB | C session setup & X | DFSCLIENT | DC-01.CONTOSO.CO |
| 18 | SMB | R session setup & X | DC-01.CONTOSO.CO | DFSCLIENT |
| 20 | SMB | C tree connect & X, Share = \\DC-01.CONTOSO.COM\IPC$ | DFSCLIENT | DC-01.CONTOSO.CO |
| 22 | SMB | R tree connect & X, Type = IPC | DC-01.CONTOSO.CO | DFSCLIENT |
| 23 | SMB | C transact2 NT Get DFS Referral | DFSCLIENT | DC-01.CONTOSO.CO |
| 24 | SMB | R transact2 NT Get DFS Referral (response to frame 23) | DC-01.CONTOSO.CO | DFSCLIENT |
| 25 | SMB | C transact2 NT Get DFS Referral | DFSCLIENT | DC-01.CONTOSO.CO |
| 26 | SMB | R transact2 NT Get DFS Referral (response to frame 25) | DC-01.CONTOSO.CO | DFSCLIENT |
| 27 | SMB | C transact2 NT Get DFS Referral | DFSCLIENT | DC-01.CONTOSO.CO |
| 28 | SMB | R transact2 NT Get DFS Referral (response to frame 27) | DC-01.CONTOSO.CO | DFSCLIENT |

The following sections provide a partial printout of the SMB data for each frame and explain the processes that occur in each frame. These processes are explained in more detail in “DFS Physical Structures and Caches on DFS Clients” and “The Referral Process” earlier in this section.

**Frames 20 and 22: Client connects to IPC$**

The client connects to IPC$ on DC-01.

**Frame 23: Client requests domain name referral**

The client sends a request with an empty name to the domain controller to request a domain name referral.

```
SMB: Transaction parameters
SMB: DFS Max Referral Level = 3 (0x3)
SMB: DFS Request Filename =
```

**Frame 24: Domain controller provides a domain name referral**

The domain controller provides a domain name referral that contains the names of all trusted domains in the client’s forest and in trusted forests. In this particular example, there is only one domain.

```
SMB: Transaction data
SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = Unknown Server Type
SMB: DFS TimeToLive = 600 (0x258)
SMB: DFS Special Name =\contoso.com
SMB: DFS Number of Expanded Names = 0 (0x0)
SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = Unknown Server Type
SMB: DFS TimeToLive = 600 (0x258)
SMB: DFS Special Name =\CONTOSO
SMB: DFS Number of Expanded Names = 0 (0x0)
```

**Frame 25: Client requests a domain controller referral**

The client requests a domain controller referral for Contoso.com. Note that the client is sending the DNS name (Contoso.com).

```
SMB: Transaction parameters
SMB: DFS Max Referral Level = 3 (0x3)
SMB: DFS Request Filename = \contoso.com
```

**Frame 26: Domain controller provides a domain controller referral**

The domain controller provides domain controller referrals for the three domain controllers in the Contoso.com domain. The domain controller names are returned as DNS names.

```
SMB Transaction data
SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = Unknown Server Type
SMB: DFS TimeToLive = 600 (0x258)
SMB: DFS Special Name =\contoso.com
SMB: DFS Number of Expanded Names = 3 (0x3)
SMB: DFS Expanded Name =\Dc-01.contoso.com
SMB: DFS Expanded Name =\Dc-03.contoso.com
SMB: DFS Expanded Name =\Dc-02.contoso.com
```

**Frame 27: Client requests a domain controller referral**

The client requests a domain controller referral for Contoso, the NetBIOS name of the domain.

```
SMB: Transaction parameters
SMB: DFS Max Referral Level = 3 (0x3)
SMB: DFS Request Filename =\CONTOSO
```

**Frame 28: Domain controller provides a domain controller referral**

The domain controller provides domain controller referrals for the three domain controllers in the Contoso.com domain. The domain controller names are returned as NetBIOS names.

```
SMB Transaction data
SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = Unknown Server Type
SMB: DFS TimeToLive = 600 (0x258)
SMB: DFS Special Name =\
SMB: DFS Number of Expanded Names = 3 (0x3)
SMB: DFS Expanded Name =\Dc-01
SMB: DFS Expanded Name =\Dc-03
SMB: DFS Expanded Name =\Dc-02
```

**Network Capture of a Client Accessing a Domain-based Namespace**

The following network capture shows what happens when a client accesses a link (Software) in the \\Contoso.com\Public namespace. This network capture was taken on a DFS client (named DFSCLIENT) and includes referrals from a domain controller (named DC-01), a root server (ROOT-DFS-03), and a link target server (NOAM-FS-1). Only SMB frames are shown, and some frames were removed for brevity. Note that this network capture is an example of client, root server, and link target communication and might differ from network captures taken in your environment.

### Network Capture of a Client Accessing a Domain-based Namespace

| Frame | Protocol | Description | Src Other Addr | Dst Other Addr |
|---|---|---|---|---|
| 352 | SMB | C tree connect & X, Share = \\DC-01.CONTOSO.COM\IPC$ | DFSCLIENT | DC-01.CONTOSO.CO |
| 353 | SMB | R tree connect & X, Type = IPC | DC-01.CONTOSO.CO | DFSCLIENT |
| 354 | SMB | C transact2 NT Get DFS Referral | DFSCLIENT | DC-01.CONTOSO.CO |
| 355 | SMB | R transact2 NT Get DFS Referral (response to frame 354) | DC-01.CONTOSO.CO | DFSCLIENT |
| 365 | SMB | C negotiate, Dialect = NT LM 0.12 | DFSCLIENT | ROOT-DFS-03 |
| 366 | SMB | R negotiate, Dialect # = 5 | ROOT-DFS-03 | DFSCLIENT |
| 369 | SMB | C session setup & X | DFSCLIENT | ROOT-DFS-03 |
| 370 | SMB | R session setup & X | ROOT-DFS-03 | DFSCLIENT |
| 374 | SMB | C tree connect & X, Share = \\ROOT-DFS-03\PUBLIC | DFSCLIENT | ROOT-DFS-03 |
| 375 | SMB | R tree connect & X, Type = A: | ROOT-DFS-03 | DFSCLIENT |
| 377 | SMB | C transact2 Query path info, File = \contoso.com\public | DFSCLIENT | ROOT-DFS-03 |

| Frame | Protocol | Description | Src Other Addr | Dst Other Addr |
|---|---|---|---|---|
| 561 | SMB | C transact2 Query path info, File = \contoso.com\public\Software | DFSCLIENT | ROOT-DFS-03 |
| 562 | SMB | R transact2 - NT error, System, Error, Code = (599) STATUS_PATH_NOT_COVERED | ROOT-DFS-03 | DFSCLIENT |
| 563 | SMB | C tree connect & X, Share = \\ROOT-DFS-03\IPC$ | DFSCLIENT | ROOT-DFS-03 |
| 564 | SMB | R tree connect & X, Type = IPC | ROOT-DFS-03 | DFSCLIENT |
| 565 | SMB | C transact2 NT Get DFS Referral | DFSCLIENT | ROOT-DFS-03 |
| 567 | SMB | R transact2 NT Get DFS Referral (response to frame 565) | ROOT-DFS-03 | DFSCLIENT |
| 577 | SMB | C negotiate, Dialect = NT LM 0.12 | DFSCLIENT | NOAM-FS-1 |
| 578 | SMB | R negotiate, Dialect # = 5 | NOAM-FS-1 | DFSCLIENT |
| 587 | SMB | C session setup & X | DFSCLIENT | NOAM-FS-1 |
| 588 | SMB | R session setup & X | NOAM-FS-1 | DFSCLIENT |
| 589 | SMB | C tree connect & X, Share = \\NOAM-FS-1\APPS | DFSCLIENT | NOAM-FS-1 |
| 590 | SMB | R tree connect & X, Type = A: | NOAM-FS-1 | DFSCLIENT |
| 591 | SMB | C transact2 Query path info, File = | DFSCLIENT | NOAM-FS-1 |
| 592 | SMB | R transact2 Query path info (response to frame 591) | NOAM-FS-1 | DFSCLIENT |
| 593 | SMB | C transact2 Findfirst, File = * | DFSCLIENT | NOAM-FS-1 |
| 594 | SMB | R transact2 Findfirst (response to frame 593) | NOAM-FS-1 | DFSCLIENT |

The following sections provide a partial printout of the SMB data for each frame and explain the processes that occur in each frame. These processes are explained in more detail in "The Referral Process" earlier in this section.

**Frames 352-353: Client connects to IPC$**
The client connects to IPC$ on DC-01.

**Frame 354: Client requests a root referral**
The client requests a root referral for the \\Contoso.com\Public namespace.

```
SMB: Transaction parameters
SMB: DFS Max Referral Level = 3 (0x3)
SMB: DFS Request Filename = \contoso.com\public
```

**Frame 355: Domain controller provides a root referral**
The domain controller provides a root referral that contains three root servers, \\Root-DFS-03\Public, \\Root-DFS-02\Public, and \Root-DFS-01\Public.

```
SMB Transaction data
SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = SMB Server
SMB: DFS TimeToLive = 300 (0x12C)
SMB: DFS Filename =\contoso.com\public
SMB: DFS 8.3 Filename =\contoso.com\public
SMB: DFS Sharename =\Root-DFS-03\public
SMB: Servicesite GUID = 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = SMB Server
SMB: DFS TimeToLive = 300 (0x12C)
SMB: DFS Filename =\contoso.com\public
SMB: DFS 8.3 Filename =\contoso.com\public
SMB: DFS Sharename =\Root-DFS-02\public
SMB: Servicesite GUID = 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = SMB Server
SMB: DFS TimeToLive = 300 (0x12C)
SMB: DFS Filename =\contoso.com\public
SMB: DFS 8.3 Filename =\contoso.com\public
SMB: DFS Sharename =\Root-DFS-01\public
SMB: Servicesite GUID = 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

**Frames 365-377: Client connects to Root-DFS-03**
The client establishes a connection with the first root server in the referral, Root-DFS-03.

**Frame 561: Client navigates to Software link folder**

The client navigates to the Software link folder on the root server.

**Frame 562: Root server provides STATUS_PATH_NOT_COVERED message**
The root server responds with the STATUS_PATH_NOT_COVERED message, which indicates that this is a link folder and that the client must request a link referral.

**Frames 563-564: Client connects to IPC$ on root server**
The client connects to IPC$ on the root server.

**Frame 565: Client requests a link referral**
The client requests a link referral to the Software link.

```
SMB: Transaction parameters
SMB: DFS Max Referral Level = 3 (0x3)
SMB: DFS Request Filename =\contoso.com\public\Software\
```

**Frame 567: Root server provides a link referral**
The root server sends a link referral that contains three link targets for the Software link: \\Noam-FS-1\Apps, \\Noam-FS-3\Apps, and \\Noam-FS-2\Apps.

```
SMB Transaction data
SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = Unknown Server Type
SMB: DFS TimeToLive = 1800 (0x708)
SMB: DFS Filename =\contoso.com\public\software
SMB: DFS 8.3 Filename =\contoso.com\public\software
SMB: DFS Sharename =\noam-fs-1\apps
SMB: Servicesite GUID = 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = Unknown Server Type
SMB: DFS TimeToLive = 1800 (0x708)
SMB: DFS Filename =\contoso.com\public\software
SMB: DFS 8.3 Filename =\contoso.com\public\software
SMB: DFS Sharename =\noam-fs-3\apps
SMB: Servicesite GUID = 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

SMB: DFS Version 3 Referral
SMB: DFS Version Number = 3 (0x3)
SMB: DFS Server Type = Unknown Server Type
SMB: DFS TimeToLive = 1800 (0x708)
SMB: DFS Filename =\contoso.com\public\software
SMB: DFS 8.3 Filename =\contoso.com\public\software
SMB: DFS Sharename =\noam-fs-2\apps
SMB: Servicesite GUID = 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

**Frames 577-589: Client connects to \\Noam-FS-1\Apps**
The DFS client sets up a session with the first link target in the referral, \\Noam-FS-1\Apps.

**DFS Interfaces**
The following APIs are associated with DFS. For more information about these APIs, see DFS in the Microsoft Platform SDK on MSDN.

### DFS APIs

| API Name | Description |
|---|---|
| NetDfsAdd | Creates a new Dfs link or adds targets to an existing link in a DFS root. |
| NetDfsAddFtRoot | Creates a new domain-based root. If the domain-based root already exists, the function adds the specified root target to the root. |
| NetDfsAddStdRoot | Creates a new stand-alone root. |
| NetDfsEnum | Enumerates all links in the named root. It also lists all roots in a domain and all roots on a server. |
| NetDfsGetClientInfo | Returns the information about cached referrals on a client for a given DFS root or link path. |
| NetDfsGetInfo | Retrieves information about a root or a link in the named root. |

| | |
|---|---|
| NetDfsRemove | Removes a link target from a link or removes a link. If the specified target is the last target of the link, then the NetDfsRemove function also removes the link. |
| NetDfsRemoveFtRoot | Removes the specified root target from a domain-based root. If the target is the last target of the root, the function also deletes the root. |
| NetDfsRemoveFtRootForced | Removes the specified root target from a domain-based root, even if the server is offline. If the target is the last root target, the function also deletes the root. |
| NetDfsRemoveStdRoot | Deletes the stand-alone root. |
| NetDfsSetClientInfo | Modifies information about cached referrals on a client for a given DFS root or link path. |
| NetDfsSetInfo | Sets or modifies the information associated with a link in the named root, or with the root itself. |

**Network Ports Used by DFS**

DFS uses the following network ports.

**Network Ports Used by DFS**

| Service Name | Relevant Computers | UDP | TCP |
|---|---|---|---|
| NetBIOS Name Service | Domain controllers; root servers that are not domain controllers; servers acting as link targets; client computers acting as link targets | 137 | 137 |
| NetBIOS Datagram Service | Domain controllers; root servers that are not domain controllers; servers acting as link targets; client computers acting as link targets | 138 | |
| NetBIOS Session Service | Domain controllers; root servers that are not domain controllers; servers acting as link targets; client computers acting as link targets | | 139 |
| LDAP Server | Domain controllers | 389 | 389 |
| Remote Procedure Call (RPC) endpoint mapper | Domain controllers | | 135 |
| Server Message Block (SMB) | Domain controllers; root servers that are not domain controllers; servers acting as link targets; client computers acting as link targets | 445 | 445 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- Active Directory Replication Topology Technical Reference
- DNS Support for Active Directory Technical Reference
- FRS Technical Reference
- Server Clusters Technical Reference

# DFS Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- DFS Tools
- DFS Registry Entries
- DFS Group Policy Settings
- Network Ports Used by DFS
- Related Information


**DFS Tools**
The following tools are associated with Distributed File System (DFS).

**Dfscmd.exe: Distributed File System command-line tool**

**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
Windows Server 2003.

Dfscmd allows administrators to perform and script basic Distributed File System (DFS) tasks, such as configuring DFS roots, links, and targets.

**Dfsgui.msc: Distributed File System snap-in**

**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
Windows Server 2003.

Dfsgui.msc is also part of the Windows Server 2003 Administration Tools Pack; you can install this pack on computers running Windows XP with Service Pack 1 (SP1) or later.

Dfsgui.msc allows administrators to create DFS namespaces, add and remove targets, set the Time to Live for targets, enable and disable referrals, enable replication for DFS root or link targets, create custom replication schedules and topologies, and so forth.

**Dfsutil.exe: Distributed File System utility**

**Category**
Windows Support Tools for Windows Server 2003.

**Version compatibility**
Windows XP, Windows 2000, and Windows Server 2003.

Dfsutil allows administrators to perform advanced DFS tasks, such as enabling root scalability mode, least expensive target selection, and same-site target selection. Dfsutil is also useful for determining the size of a namespace, exporting or importing namespaces, checking the site name of a computer or IP address, adding and removing root targets, and updating site information for root servers running Windows 2000 Server. When installed on DFS clients, Dfsutil can be used to view and clear the referral cache (PKT cache), domain cache (SPC cache), and MUP cache.

**DFS Registry Entries**
The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It

is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following registry entries are associated with DFS.

For more information about registry entries, see the Registry Reference in the Tools and Settings Collection.

**DFS Domain Controller Registry Entries**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs on domain controllers. These registry entries are not used on DFS root servers unless the root server is also a domain controller. All entries are REG_DWORD.

**DomainNameIntervalInSeconds**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters\

**Version**
Domain controllers running Windows Server 2003.

Specifies how long (in seconds) domain controllers cache domain name referrals. The default value is 43,200 seconds (12 hours). The minimum is 600 seconds (10 minutes). There is no maximum value.

**SiteCostedReferrals**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters\

**Version**
Domain controllers running Windows Server 2003.

When set to 0 (the default), SYSVOL and NETLOGON referrals contain domain controllers in the client's site listed first in random order, followed by a random list of domain controllers. When set to 1, SYSVOL and NETLOGON referrals sort domain controllers in order of lowest cost. Domain controllers in the clients site are at the top of the referral list, followed by domain controllers sorted by lowest cost.

**DFS Domain Controller and Root Server Registry Entries**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs on root servers and domain controllers. All entries are REG_DWORD.

**DfsDnsConfig**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\

**Version**
Domain controllers and root servers running Windows Server 2003.

When set to 1, specifies that this server will use fully qualified domain names (FQDN) in referrals. When set to 0 (the default), specifies that this server will use NetBIOS names in referrals.

**LdapTimeoutValueInSeconds**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters\

**Version**
Domain controllers and root servers running Windows Server 2003.

Determines the time-out value (in seconds) for DFS LDAP calls. The default value is 30 seconds, the minimum value is 3 seconds, and the maximum value is 300 seconds (5 minutes).

**MaxClientsToCache**

**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters\

**Version**
Domain controllers and root servers running Windows Server 2003.

Specifies the maximum number of entries that can be maintained in the client site cache. When this limit is hit, unused cached entries are dropped. Client site entries are regenerated as needed. The default value is 200,000 entries, and the minimum value is 20,000 entries. There is no maximum value.

**QuerySiteCostTimeoutinSeconds**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters\

**Version**
Domain controllers and root servers running Windows Server 2003.

Specifies the time-out period (in seconds) after which DsQuerySitesByCost API calls for site cost information in Active Directory will fail. By default, the time-out period is 30 seconds. The minimum value is 3 seconds, and the maximum value is 300 seconds (5 minutes). If DFS cannot determine a target's site cost within this time-out period, DFS assumes the maximum possible cost for the target.

**RootReferralTimeoutInSeconds**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters\

**Version**
Domain controllers and root servers running Windows Server 2003.

Specifies how long (in seconds) domain controllers and root servers cache domain-based root referrals. The default and minimum value is 900 seconds (15 minutes). There is no maximum value.

**SiteSupportIntervalInSeconds**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\Parameters\

**Version**
Domain controllers and root servers running Windows 2000 Server and Windows Server 2003.

Specifies how long (in seconds) entries are maintained in the client site cache, target site cache, and site cost cache. The default value is 43,200 seconds (12 hours), and the minimum value is 300 seconds (5 minutes). Due to internal DFS processes, the entries in the client site cache and target site can remain for up to two times the value of this registry entry.

**SyncIntervalinSeconds**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dfs\

**Version**
Domain controllers and root servers running Windows Server 2003.

Specifies how often domain-based root servers and domain controllers poll the primary domain controller (PDC) emulator master to obtain updated DFS metadata. The minimum value is 900 seconds (15 minutes), and the default value is 3600 seconds (60 minutes).

**DFS Client Registry Entries**
The following registry entries are available for DFS clients running Windows XP and Windows 2000 Server. All entries are REG_DWORD.

**Client Referral Time to Live**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DfsClientDriver\Paramaters

**Version**
Client computers running Windows XP and Windows 2000.

Specifies referral Time To Live (TTL) values, overriding those provided by the server.

There are six keys that you can specify to control referral TTL values:

- **MinReferralTTL**The minimum TTL for non-domain referrals, in seconds. The default is **300** (five minutes).
- **MaxReferralTTL**The maximum TTL for non-domain referrals, in seconds. The default is **10080** (one week).
- **MinDomainTTL**The minimum TTL for domain referrals, in seconds. The default is **900** (15 minutes).
- **MaxDomainlTTL**The maximum TTL for domain referrals, in seconds. The default is **10080** (one week).
- **ReferralTTLMultiplier**The number of failed attempts to validate a non-domain referral before the client will stop testing the referral. The minimum value is **1**; the maximum is **100**; the default is **4**.
- **DomainTTLMultiplier**The number of failed attempts to validate a domain referral before the client will stop testing the referral. The minimum value is **1**; the maximum is **100**; the default is **4**.

**DfsDcNameDelay**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanworkstation\

**Version**
Client computers running Windows XP and Windows 2000.

Specifies the interval at which DFS clients discover new domains in the local forest and in trusted forests. This discovery process runs against a domain controller from the domain that is hosting the client's computer account. The default value for this entry is 900 seconds (15 minutes).

**DisableDFS**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Mup\

**Version**
Client computers running Windows Server 2003, Windows XP and Windows 2000.

When set to 0 (the default), clients can access DFS namespaces, including the SYSVOL and NETLOGON shared folders on domain controllers. When set to 1, disables DFS processing on the client.

**ProviderCacheTimeoutInMinutes**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Mup\

**Version**
Client computers running Windows Server 2003 or Windows XP.

Specifies the length of time that a multiple UNC provider (MUP) cache entry is held until it is reevaluated. The default and minimum value is 15 minutes, and the maximum value is 720 minutes (12 hours).

**DFS Group Policy Settings**
The following table lists and describes the Group Policy settings are associated with DFS.

### Group Policy Settings Associated with DFS

| Group Policy Setting and Path | Description |
|---|---|

| | |
|---|---|
| Computer Configuration\Administrative Templates\Network\Sets how often a DFS Client discovers DC's | Configures how often a DFS client attempts to discover domain controllers on the network. By default, a DFS client attempts to discover domain controllers every 15 minutes. This value is specified in minutes. If you disable this setting or do not configure it, the default value (15 minutes) is used.<br><br>The minimum value you can select is 15 minutes. If you specify a value of less than 15 minutes, the default value of 15 minutes is used. |
| User Configuration\Administrative Templates\Shared Folders\Allow DFS roots to be published | Determines whether the user can publish DFS roots in Active Directory. If you enable this setting or do not configure it, users can use the "Publish in Active Directory" option to publish DFS roots as shared folders in Active Directory. If you disable this setting, the user cannot publish DFS roots in Active Directory, and the "Publish in Active Directory" option is disabled.<br><br>The default is to allow shared folders to be published when this setting is not configured. |

For more information about Group Policy settings, see the Group Policy Settings Reference in the Tools and Settings Collection.

**Network Ports Used by DFS**

The following table lists the network ports used by DFS.

**Network Ports Used by DFS**

| Service Name | Relevant Computers | UDP | TCP |
|---|---|---|---|
| NetBIOS Name Service | Domain controllers; root servers that are not domain controllers; servers acting as link targets; client computers acting as link targets | 137 | 137 |
| NetBIOS Datagram Service | Domain controllers; root servers that are not domain controllers; servers acting as link targets; client computers acting as link targets | 138 | |
| NetBIOS Session Service | Domain controllers; root servers that are not domain controllers; servers acting as link targets; client computers acting as link targets | | 139 |
| LDAP Server | Domain controllers | 389 | 389 |
| Remote Procedure Call (RPC) endpoint mapper | Domain controllers | | 135 |
| Server Message Block (SMB) | Domain controllers; root servers that are not domain controllers; servers acting as link targets; client computers acting as link targets | 445 | 445 |

**Related Information**

The following resources contain additional information that is relevant to this section.

- DNS Support for Active Directory Technical Reference
- Group Policy Settings Reference in the Tools and Settings Collection

# FRS Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FRS Technical Reference**
File Replication service (FRS) is a technology that replicates files and folders stored in the SYSVOL shared folder on domain controllers and Distributed File System (DFS) shared folders. When FRS detects that a change has been made to a file or folder within a replicated shared folder, FRS replicates the updated file or folder to other servers. Because FRS is a multimaster replication service, any server that participates in the replication of a shared folder can generate changes. In addition, FRS can resolve file and folder conflicts to make data consistent among servers.

**In this subject**


- What Is FRS?
- How FRS Works
- FRS Tools and Settings

# What Is FRS?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Common FRS Scenarios
- Technologies Related to FRS
- FRS Dependencies
- Related Information


File Replication service (FRS) is a technology that replicates files and folders stored in the SYSVOL shared folder on domain controllers and Distributed File System (DFS) shared folders. When FRS detects that a change has been made to a file or folder within a replicated shared folder, FRS replicates the updated file or folder to other servers. Because FRS is a multimaster replication service, any server that participates in replication can generate changes. In addition, FRS can resolve file and folder conflicts to make data consistent among servers.

By keeping files and folders synchronized across servers, FRS enables organizations to increase the availability of data. If one server becomes unavailable, the files are still available, because they exist on another server. Using multiple servers to host data also helps organizations that have offices in multiple geographic locations, because clients can access servers in or closest to their current site and do not need to use expensive WAN links to access data.

There are numerous methods for keeping files synchronized on servers. Although SYSVOL requires FRS, DFS shared folders can be kept synchronized by using methods other than FRS, such as manual copying, Robocopy, or other replication tools. FRS provides numerous benefits that other replication methods do not. These benefits include the following:

**Authenticated RPC with encryption**
To provide secure communications, FRS uses Kerberos authentication protocol for authenticated remote procedure call (RPC) to encrypt the data sent between members of a replica set.

**Compression**
To save network bandwidth, FRS compresses files in the staging folder by using NTFS compression. Files sent between servers participating in replication, known as replica members, remain compressed when transmitted over the network.

**Conflict resolution**
FRS resolves file and folder name conflicts to make data consistent among the replica members. If identically named files are created or modified on two or more replica members, FRS uses a “last writer wins” rule; this means that the most recently created or modified version of a file becomes the version that is replicated to the other replica members. If identically named folders are created on two or more replica members, FRS identifies the conflict during replication and renames the folder that was most recently created. Both folders are replicated to all servers in the replica set, and administrators can merge the contents of two folders or take some other measure to reestablish the single folder.

**Continuous replication**
FRS provides continuous replication, subject to replication schedule, server load, and network load. When a file or folder is changed and closed, FRS begins replicating the changed file or folder to other replica members within three seconds.

**Fault-tolerant replication path**
FRS does not rely on broadcast technology, and it can provide fault-tolerant distribution by way of multiple connection paths between members. If a replica member is unavailable, the data will take a different route. FRS prevents an identical file from being sent more than once to any replica member.

**Replication scheduling**
You can schedule replication to occur at specified times and durations as needed by your organization. For example, scheduling replication to occur during evening hours can reduce the cost of transmitting data over expensive WAN links. Replicating data during off-hours also frees up network bandwidth for other uses.

**Replication integrity**
Files are replicated only after they have been changed and closed. FRS relies on the update sequence number (USN) journal to log records of files that have changed on a replica member. As a result, FRS does not lose track of a changed file even if a

replica member shuts down abruptly. After the replica member comes back online, FRS replicates new or updated files that originated from other replica members, as well as replicating locally created or updated files that occurred before the shutdown. This replication takes place according to the replication schedule.

**Common FRS Scenarios**
FRS is commonly used in the following scenarios:

**SYSVOL Replication**
Every domain controller has a shared folder in its local file system that is the file system component of Active Directory. This shared folder, named SYSVOL, contains files and folders that must be available and synchronized between domain controllers in a domain, including:

- The NETLOGON shared folder, which includes system policies and user-based logon and logoff scripts for non-Windows Server 2003 and non-Windows 2000 network clients, such as clients running Windows 95, Windows 98, and Windows NT 4.0.
- Windows Server 2003 and Windows 2000 system policies.
- Group Policy settings (templates), including Group Policy settings for domain controllers running Windows Server 2003 or Windows 2000.

When you add, remove, or modify the contents in the SYSVOL shared folder, FRS replicates the changed contents to the SYSVOL shared folders on all other domain controllers in the domain. The following figure illustrates how SYSVOL shared folders are kept in sync on domain controllers in the corp.contoso.com domain.

**SYSVOL Replication**




**Publishing Applications**
FRS and DFS can be combined to distribute and publish applications. DFS allows administrators to group shared folders located on different servers by transparently connecting them to one or more hierarchical namespaces that behave like a single high-capacity hard disk. Users can navigate the namespace without having to know the physical server names or shared folders hosting the data. Administrators can also host data on multiple servers in multiple sites; users see a single copy of the data and DFS can refer users to the closest server to access the data.

When using DFS and FRS to publish applications, an administrator defines the topology used to distribute the data. The topology defines the path along which files are replicated; common topologies for distributing applications are hub-and-spoke, a multilevel tree, and redundant hub-and-spoke. New applications are typically deployed or updated on a single computer (usually the hub or root of the topology or another centrally located computer), and FRS then replicates the files to the spoke computers. The following figure illustrates a hub-and-spoke topology where applications are deployed on the hub and then replicated to the spokes.

**Publishing Applications Using a Hub-and-Spoke FRS Topology**




**Publishing Data**

DFS and FRS can also be used to publish data across an organization. Common examples of data include documents, diagrams, and operational procedures. This scenario often uses similar topologies (hub-and-spoke, a multilevel tree, or a redundant hub-and-spoke) as those used for publishing applications.

Another method of publishing data is known as reverse publication. In reverse publication, data flows from a number of different servers to a central server. Reverse publication is often used to gather log files and reports from individual computers and collect them in one central location. Reverse publication is also used in backup scenarios where data is collected from multiple servers and then backed up from the central server.

The following figure illustrates the reverse publication method.

**Reverse Publication**




**Ensuring That Personal Folders Are Available**

Organizations that store users' personal folders on file servers can increase the availability of those folders by using DFS and FRS to store the folders on multiple servers. The following figure illustrates a DFS namespace, \\Corp.Contoso.Com\Users, hosted by two servers, ServerA and ServerB. DFS equally distributes the client load between the two servers, and FRS keeps the data synchronized on the two servers. If one of the servers fails, DFS refers clients to the remaining server. Even when one of the servers is unavailable, users can continue to access the \\Corp.Contoso.com\Users namespace.

**How DFS and FRS Keep Data Available**




**Technologies Related to FRS**

FRS is closely related to the following two technologies.

**Active Directory replication**

Active Directory replication is the process by which the changes that are made to Active Directory objects on one domain controller are automatically synchronized with other domain controllers. This replication method does not use FRS and does not include SYSVOL replication. However, FRS replicates SYSVOL by using the connection object topology and schedule that the Knowledge Consistency Checker (KCC) creates for Active Directory replication. In addition, FRS has its own Active Directory objects that are replicated using Active Directory replication.

**DFS**

FRS can be used to keep data in DFS shared folders synchronized among replica members. However, DFS and FRS are two separate technologies, and DFS does not require FRS. You can use other replication methods, such as manual copying, the Windows Resource Kit tool Robocopy, or other replication tools to keep DFS shared folders synchronized. Conversely, if you want to use FRS to keep data in shared folders synchronized, you must use DFS.

**FRS Dependencies**

FRS has the following dependencies:

- **Active Directory replication**. FRS requires that Active Directory replication is working properly so that FRS Active Directory objects reside on all domain controllers in the domain.

- **DFS**. If you want to use FRS to keep data in folders synchronized on multiple servers, you must first set up a DFS namespace. (This dependency is not applicable to SYSVOL.)

- **DNS**. FRS requires that DNS is properly designed and deployed so that FRS can correctly resolve DNS names of replica members. (DNS is also required for Active Directory replication.)

- **Kerberos authentication**. FRS requires that Kerberos authentication is properly designed and deployed.

- **NTFS**. To detect changes to files and folders, FRS relies on the USN journal on NTFS volumes. FRS is not supported on FAT file system volumes.

- **Remote procedure call (RPC)**. FRS requires Internet Protocol (IP) connectivity and the Remote Procedure Call (RPC) for communication between replication partners and communication with domain controllers.

**Related Information**

The following resources contain additional information that is relevant to this section:

- For information about Active Directory, see "Active Directory Replication Topology Technical Reference."

# How FRS Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- FRS Terminology
- FRS Architecture
- FRS Protocols
- FRS Interfaces
- FRS Physical Structures
- FRS Processes and Interactions
- Network Ports Used by FRS
- Related Information


File Replication service (FRS) is a technology that replicates files and folders stored in the SYSVOL shared folder on domain controllers and Distributed File System (DFS) shared folders. When FRS detects that a change has been made to a file or folder within a replicated shared folder, FRS replicates the updated file or folder to other servers. Because FRS is a multimaster replication service, any server that participates in replication can generate changes. In addition, FRS can resolve file and folder conflicts to make data consistent among servers.

These sections provide an in-depth view of how FRS works in an optimal environment. An optimal environment for FRS is defined as follows:

- DNS and Active Directory replication are working properly.
- FRS Active Directory objects are properly configured and present on all domain controllers.
- Kerberos authentication is working properly.
- No firewalls prevent replication from working.
- All servers that participate in replication (known as replica members) have adequate free space on NTFS file system volumes to support the FRS physical structures.
- All replica members have hardware that is sized appropriately for the roles that the servers perform and the hardware is functioning properly.
- No replica members have network, disk, RAM, or CPU bottlenecks.
- All members are running Windows Server 2003 with Service Pack 1 (SP1) or the pre-SP1 release of Ntfrs.exe (as described in article 823230, "Issues That Are Fixed in the Pre-Service Pack 1 Release of Ntfrs.exe," in the Microsoft Knowledge Base).
- For SYSVOL replica sets, the number of domain controllers in a domain falls well below the recommended limit of 1,200 domain controllers.
- No replica members are running disk defragmentation, backup, or antivirus software that is known to be incompatible

with FRS.

- The update sequence number (USN) journal and staging folder on all replica members are sized according to Microsoft's recommendations.
- The system clocks on all replica members are set to within 30 minutes of each other.
- No files are blocking replication by being held open for an extended period of time on any replica member.
- Changes to a particular file are only made on one replica member so that no conflict resolution is required.
- Administrators do not attempt to copy files from one replica member to another in an attempt to help replication.

**FRS Terminology**

Before you review the FRS components and processes, it is helpful to understand FRS terminology. The following sections provide a brief introduction and illustration of the basic components of FRS as well as a glossary of terms.

**Basic FRS Components**

When you set up DFS replication, you choose a shared folder whose contents you want to replicate. For the SYSVOL shared folder on domain controllers, replication is established automatically during the domain controller promotion process. The set of data to be replicated to all servers is known as the replica tree, and the servers to receive this replica tree are known as replica members. The group of replica members that participates in the replication of a given replica tree is known as a replica set. Replica members are connected in a topology, which is a framework of connections between servers and defines the replication path. There are a number of common topologies, such as full mesh, ring, and hub and spoke.

The following figure illustrates these terms.

**Basic FRS Components (Ring Topology)**




It is also helpful to understand the terminology used to describe the relationship between two members. This terminology is illustrated in the following figure.

**Replication Partner Terminology**




As illustrated in the figure, any two replica members with a direct connection between them are known as direct replication partners. In the figure, ServerA and ServerB are direct replication partners, as are ServerB and ServerC. Although changes that originate on ServerA are eventually replicated to ServerC, these two servers do not have a direct connection to each other, so they are known as transitive replication partners.

When a changed file or folder is received on a server, either because the change originated there or because the server received the file or folder from another partner, that server becomes the upstream partner, also known as inbound partner, for any direct replication partner that has not yet received the change. For example, in the figure titled "Replication Partner Terminology," a change originates on ServerA, and ServerA is considered the upstream or inbound partner for ServerB. ServerB is considered the downstream (also known as outbound) partner for ServerA.

**Glossary of FRS Terms**

The following terms are used to describe the components and processes of FRS.

**Active Directory object**

A distinct set of named attributes that represents a network resource. FRS uses Active Directory objects to represent servers that participate in replica sets and the topology that FRS uses to replicate data.

**Authoritative restore (also called D4)**

The process whereby one replica member is made the authoritative member, and all the data in its replica tree is replicated to all other replica members. This procedure should be used only in extreme circumstances under the supervision of your support provider or Microsoft Product Support Services. This process is informally referred to as a "D4" in reference to the registry setting used to invoke this process.

**Burflags**

Short for "backup restore flags," an FRS-related registry entry used to change the behavior of the FRS service at startup. The values for this registry entry are D2 (for nonauthoritative restores) and D4 (for authoritative restores).

**Change order**

A message that contains information about a file or folder that has changed on a replica member. The change order is sent to the member's outbound partners. If the outbound partners accept the change, the partners request the associated staging file. After installing the changed file in their individual replica trees, the partners propagate the change order to their outbound partners.

**Directed change order**

A change order that is directed to a single outbound partner and primarily produced when the partner is doing a vvjoin, such as during initial synchronization or a nonauthoritative restore (D2).

**File event time**

The time at which a change to a file or folder is made. This might not be the same as the file create or last-write time. For example, restoring a file from a backup tape preserves the file create and last-write times, but the file event time is the time when the actual file restoration was performed.

**File GUID**

An identifying property of a file or folder in a replica tree. FRS creates and manages file globally unique identifiers (GUIDs), which, along with the replication version number and event time, are stored in the file ID table in the FRS database. Each file

and folder stores its file GUID as part of its NTFS attributes; therefore, corresponding files and folders across all replica set members have the same file GUID.

**File ID table**

A table in the FRS database that contains an entry with version and identity information for each file and folder in the replica tree.

**File version number**

A property of a file and folder in a replica tree that is incremented each time the file or folder is updated. The file version number is used to resolve concurrent updates originating from more than one member of the replica set. The version number is only incremented by the member that originated the file update. Other members that propagate the update do not change the version number.

**Filter**

A setting that excludes subfolders (and their contents) or files from replication.

**FRS debug logs**

Text files in the *systemroot*\Debug folder that store FRS transaction and event detail.

**Inbound connection**

For a given replica member, a component of the NTFRS Member object in Active Directory that identifies inbound partners. An inbound connection exists for each inbound partner.

**Inbound log**

A table in the FRS database that stores pending change orders to be processed. As entries are processed, acknowledgments are sent to the inbound partners.

**Knowledge Consistency Checker (KCC)**

A built-in process that runs on all domain controllers and generates replication topology for the Active Directory forest. The KCC creates separate replication topologies depending on whether replication is occurring within a site (intrasite) or between sites (intersite). The KCC also dynamically adjusts the topology to accommodate new domain controllers, domain controllers moved to and from sites, changing costs and schedules, and domain controllers that are temporarily unavailable.

**Link target**

In a DFS namespace, the mapping destination of a link, typically a shared folder on a server. FRS can be used to keep link targets (shared folders) synchronized on different servers.

**Local change order**

A change order that is created because of a change to a file or folder on the local server. The local server becomes the originator of the change order and constructs a staging file.

**MD5 checksum**

A cryptographically secure one-way hashing algorithm that is used by FRS to verify that a file on each replica member is identical.

**Morphed folder**

A folder that FRS creates when resolving a name conflict between two or more identically named directories that originate on different replica members. FRS identifies the conflict during replication, and the receiving member protects the original copy of the folder and renames (morphs) the later inbound copy of the folder. The morphed folder names have a suffix of "_NTFRS_xxxxxxxx," where "xxxxxxxx" represents eight random hexadecimal digits.

**Nonauthoritative restore (also called D2)**

The process whereby a given replica member is reinitialized by obtaining a replica tree from another replica member. This process is often used to resynchronize a replica member's replica tree with one of its inbound partners, such as after a server failure. This process is informally referred to as a "D2" in reference to the registry setting used to invoke this process.

**Originator GUID**

A globally unique identifier (GUID) that is associated with each replica member. All change orders produced by a given replica member carry the replica member's originator GUID, which is saved in the file ID table.

**Outbound connection**

For a given replica member, a component of the NTFRS Member object in Active Directory that identifies outbound partners. An outbound connection exists for each outbound partner.

**Outbound log**

A table in the FRS database that stores pending change orders to be sent to outbound partners. The changes can originate locally or come from an inbound partner. These change orders are eventually sent to all outbound replica partners.

**Parent GUID**

The globally unique identifier (GUID) of the parent folder that contains a particular file or folder in the replica tree.

**Preinstall folder**

A hidden subfolder under the root of the replica tree. When a newly created file or folder is replicated to a downstream partner, the file or folder is first created in the preinstall folder on the downstream partner. After the file or folder is completely replicated, it is renamed to its target location in the replica tree. This process is used so that partially constructed files are not visible in the replica tree.

**Reanimation change order**

A change order that FRS uses to reconcile a recently received change order against a previously deleted file or folder.

**Remote change order**

A change order received from an inbound (or upstream) partner that originated elsewhere in the replica set.

**Retry change order**

A change order that is in some state of completion but was blocked for some reason and must be retried later.

**Schedule**

The frequency at which data replicates.

**Staging folder**

A folder that acts as a queue for changed files and folders to be replicated to downstream partners.

**Staging file**

A file that FRS constructs in the staging folder as a result of a file or folder change. FRS constructs staging files by using backup application programming interfaces (APIs) and then replicates the staging files to downstream partners. Downstream partners use restore APIs to reconstruct the staging files in the preinstall folder before renaming the files into the replica tree. Staging files are compressed to reduce the network bandwidth used during replication.

**SYSVOL**

On a domain controller, a shared folder that stores a copy of the domain's public files, including system policies and Group Policy settings, that are replicated to all other domain controllers in the domain by FRS.

**Update sequence number (USN)**

A monotonically increasing sequence number for each volume. The USN is incremented every time a modification is made to a file on the volume.

**USN journal**

On NTFS volumes, a persistent log that tracks all changes on the volume, including file creations, deletions, and changes. The USN journal has a configurable maximum size and is persistent across reboots and system crashes. FRS uses the USN journal to monitor changes made in the replica tree.

**USN journal wrap**

An error that occurs when large numbers of files change so quickly that the USN journal must discard the oldest changes

(before FRS has a chance to detect the changes) to stay within the specified size limit. To recover from a journal wrap, you must perform a nonauthoritative restore on the server to synchronize its files with the files on the other replica members.

### Version vector

A vector of volume sequence numbers (VSNs), with one entry per replica set member. All change orders carry the originator GUID of the originating member and the associated VSN. As each replica member receives the update, it tracks the VSN in a vector slot that is assigned to the originating member. This vector describes whether the replica tree is current with each member. The version vector is then used to filter updates from inbound partners that might have already been applied. The version vector is also delivered to the inbound partner when the two members join. When a new connection is created, the version vector is used to scan the file ID table for more recent updates that are not seen by the new outbound partner.

### Version vector join (vvjoin)

The process in which a downstream partner joins with an upstream partner for the first time.

#### FRS Architecture

The File Replication service (Ntfrs.exe) is the primary component of the FRS architecture. This service and its related components are illustrated in the following figure.

### FRS Architecture




The following table describes the components related to the FRS architecture.

### Components of the FRS Architecture

| Component | Description |
|---|---|
| Active Directory | Where FRS topology and schedule information is stored. FRS periodically polls Active Directory to retrieve updated information about the FRS topology and schedule. FRS uses Lightweight Directory Access Protocol (LDAP) to connect to Active Directory using port 389. |
| Esent.dll | The dynamic link library of the Jet database. |
| File Replication service event log (NtFrs.Evt) | The log where events and error messages related to FRS are logged in Event Viewer. The FRS event log file is located in *Systemroot*\System32\Config\. |
| FRS debug logs | Text files where FRS stores transaction and event details. The FRS debug logs are typically used for troubleshooting FRS. |

| | |
|---|---|
| FRS Jet database (Ntfrs.jdb) | The database where FRS transactions are stored. |
| Ntfrs.exe | The executable file for the File Replication service. Runs as LocalSystem on all domain controllers and members of DFS replica sets. |
| Ntfrsapi.dll | The dynamic link library file that allows other applications to get information about FRS via remote procedure call (RPC). These applications include Active Directory Installation Wizard (Dcpromo), Ntfrsutl.exe (part of the Windows Support Tools), Windows Backup, and Volume Shadow Copy service. |
| Ntfrsprf.dll | The dynamic link library file that allows System Monitor to connect to FRS and collect performance-related information via local remote procedure call (LRPC). |
| Ntfrsres.dll | The dynamic link library file where event log messages are stored. When events occur, the event log messages are displayed in the File Replication service event log. |
| Registry | Where FRS configuration data is stored. Some FRS-related registry keys are recognized when changes to those keys occur; other registry keys are read only when the File Replication service is first started or restarted. |
| Replica tree | The contents of a folder that is replicated among replica members in a replica set. |
| Staging folder | A folder that acts as a queue for changed files to be replicated to downstream partners. |
| USN journal (also called NTFS change journal) | A persistent log of changes made to files on an NTFS volume. NTFS uses the change journal to track information about added, deleted, and modified files for each volume.The USN journal describes the nature of any changes to files on the volume. When any file or folder is created, modified, or deleted, NTFS adds a record to the USN journal for that volume. |

**FRS Protocols**
FRS stays at application layer of Open Systems Interconnection (OSI) model, and uses RPC as transport medium for communication between replica members.

**FRS Interfaces**
Windows Server 2003 includes an FRS writer for the Volume Shadow Copy service. The FRS writer allows Volume Shadow Copy service-compatible backup programs, such as Windows Backup, to make point-in-time, consistent backups of the replica tree. Using a Volume Shadow Copy service-compatible backup program ensures that replica tree backups do not contain partial files (which can occur when a replica member receives an updated file and begins overwriting the previous version of the file in the replica tree with the updated version of the file). This process, known as the install process, creates a window during which the file is half-written. The window exists because FRS must move the updated file from the preinstall folder and overwrite the existing file in the replica tree.

Before a Volume Shadow Copy service-compatible backup program takes a shadow copy of a replica tree, the program instructs FRS to stop requesting new work items from the install command server queue. After all currently active installations are complete, FRS enters a frozen state during which no installations occur. Shadow copies made while FRS is in a frozen state will not contain partial files. The install service thaws either when the backup program signals a thaw or at a predetermined time-out period.

For more information about the Volume Shadow Copy service, see "Volume Shadow Copy Service Technical Reference."

**FRS Physical Structures**
FRS requires a number of physical structures to be stored in Active Directory and on each replica member. These structures are illustrated in the following figure and are described in the sections that follow.

### Overview of FRS Physical Structures




**FRS Files and Folders on Replica Members**

Each replica member contains a number of files and folders used by FRS for SYSVOL and DFS replica sets. The following figure illustrates the files and folders that are stored on a DFS replica member.

**FRS Files and Folders on DFS Replica Members**




The following figure illustrates the files and folders stored on domain controllers. Note how the location of these folders differs from DFS replica members.

**FRS Files and Folders on Domain Controllers**




The following sections briefly describe each of these folders.

**Replica tree on DFS replica members**

The replica tree is the set of files and folders that are kept in sync on all replica members. The replica tree begins at the replica root, which is the shared folder where all the replicated data is stored. The replica tree can exist on any NTFS volume for DFS replica members, and FRS does not impose limits on the size of the replica tree.

**SYSVOL folder domain controllers**

SYSVOL is a folder that stores numerous FRS-related folders and a copy of the domain's public files, including system policies and Group Policy settings that are replicated to all other domain controllers in the domain by FRS. The actual replica root begins at the *systemroot*\SYSVOL\domain folder, but the folder that is actually shared is the *systemroot*\SYSVOL\sysvol folder. These folders appear to contain the same content because SYSVOL uses junction points (also called reparse points). A junction point is a physical location on a hard disk that points to data that is located elsewhere on the hard disk or on another storage device. The following table lists the folders in SYSVOL that contain junction points and the locations to which these junction points resolve:

**SYSVOL Junction Points**

| Location of Junction Point | Where the Junction Point Resolves |
|---|---|
| SYSVOL\staging areas*domainname* | SYSVOL\staging\domain |
| sysvol*domainname* | SYSVOL\domain |

**Staging folder**

The staging folder is a hidden folder that acts as a queue for changed files to be replicated to downstream partners. After the changes are made to a file and the file is closed, FRS creates the file in the staging folder by using backup application programming interfaces (APIs) and replicates the file according to schedule. Files in the staging folder are called staging files. Any further use of corresponding files in the replica tree does not prevent FRS from replicating the staging files to other members. In addition, if the file is replicated to multiple downstream partners or to members with slow data links, using staging files ensures that the underlying files in the replica tree can still be accessed.

The default size of the staging folder is approximately 660 megabytes (MB), the minimum size is 10 MB, and the maximum size is 2 terabytes. For performance and disk space purposes, the staging folder can be located on a volume other than where the replica tree is stored.

For more information about the staging folder, see "How the Staging Folder Works" later in this section.

**Preinstall folder**

The preinstall folder is a hidden system folder named DO_NOT_REMOVE_NtFrs_PreInstall_Directory. When a downstream partner begins receiving a new file (or folder), the file is installed in the preinstall folder on the downstream partner so that a partially constructed file is not added to the replica tree. After the downstream partner receives the entire file in the preinstall folder, the file is renamed to its target location in the replica tree.

**FRS Jet database**

FRS creates a Microsoft Jet database at *systemroot*\ntfrs\jet\ to store FRS transactions. FRS does not impose size limits on the Jet database, but replication stops working if the volume where the database is stored runs out of disk space. For performance and disk space purposes, the Jet database can be moved to a volume other than where the replica tree is stored.

For more information about the FRS Jet database, see "FRS Database Tables" later in this section.

**FRS debug logs**
FRS creates debug logs in the *systemroot*\Debug folder. By default, the debug logs store FRS transaction and event detail in sequentially numbered files from Ntfrs_0001.log through Ntfrs_0005.log. The characteristics of the debug logs are determined by the values of several registry entries. These values allow you to set the number of log files used, the number of entries per log, the level of detail logged, and so forth.

FRS creates another log file, Ntfrsapi.log, to track events that occur during the promotion and demotion of domain controllers. Information in the Ntfrsapi.log file includes which server was chosen as the parent for the initial replication (seeding) process and the creation of registry keys. Because errors during the promotion and demotion processes are also tracked in this file, Ntfrsapi.log is useful for troubleshooting problems where the SYSVOL and Netlogon shared folders are not shared correctly.

For performance and disk space purposes, the debug logs can be moved to a volume other than the one on which the replica tree is stored.

For more information about the FRS debug logs, see "How FRS Debug Logs Work" later in this section.

**Pre-existing data folder**
The pre-existing folder, named NtFrs_PreExisting___See EventLog, is an optional folder that is located under the replica root. If the pre-existing data folder is present on a replica member, FRS created it after one of the following events:

- The server was added to a replica set but the server already had one or more files in the shared folder that became the replica tree. In this case, FRS moved that data into the pre-existing data folder and then replicated the replica tree from one of the upstream partners to the new member.

- The replica member had a nonauthoritative restore (also called D2) performed on it by an administrator. This process is used to bring a replica member back up to date with its partners after problems such as assertions in the FRS service, corruption of the local FRS Jet database, journal wrap errors, and other replication failures. When you perform a nonauthoritative restore on a server, FRS moves the existing data in the replica tree to the pre-existing data folder and then receives the replica tree from one of the upstream partners.

- The server was prestaged before it was added to the replica set. During the prestaging process, files in the replica tree are temporarily moved from the replica tree to the pre-existing data folder. For more information about the prestaging process, see "How Prestaging Works" later in this section.

Only one pre-existing data folder can exist at a time. If one of the previously listed events occurs, causing the pre-existing data folder to be created, and then another one of the events occurs, the previous pre-existing data folder is deleted and replaced with another pre-existing data folder.

**FRS Database Tables**
The Jet database (Ntfrs.jdb) used by FRS has a set of five tables:

- Connection record table

- Version vector table

- File ID table

- Inbound change order table

- Outbound change order table

You can view the contents of three of these tables (the file ID table, inbound change order table, and outbound change order table) by using Ntfrsutl.exe (a Windows Support Tool), though doing so is typically only necessary for troubleshooting purposes. Windows Server 2003 provides several Windows Support Tools scripts, Topchk.cmd, Connstat.cmd, and Iologsum.cmd, that can present the output of Ntfrsutl in a more readable format.

The following sections briefly describe the tables in Ntfrs.jdb.

**Connection record table**
The connection record table tracks the state of each inbound and outbound connection for the replica set and tracks the delivery state of each outbound change order.

**Version vector table**
The version vector table contains a vector of volume sequence numbers (VSNs), one entry per replica member. All change orders carry the originator globally unique identifier (GUID) of the originating member and the associated volume sequence number. As each replica set member receives the update, it tracks the VSN in a vector slot assigned to the originating member. This vector now describes how up-to-date this member's replica tree is with respect to each member. The version vector is then used to filter out updates from inbound partners that might have already been applied. The version vector is also delivered to the inbound partner when two members join. When a new connection is created, the version vector is used to scan the file ID table for more recent updates not seen by the new outbound partner.

**File ID table**
The file ID table lists all files in the replica sets on the local server of which FRS is aware. Data stored in the file ID table includes the file name, file GUID, NTFS file ID, parent GUID, parent file ID, version number, event time, and originator GUID. You can view entries in the file ID table by using Ntfrsutl.exe with the **idtable** parameter.

**Inbound change order table (inbound log)**
The inbound change order table, referred to as the inbound log, stores pending change orders to be processed. As entries are processed, acknowledgments are sent to the inbound partners. State variables in the record track the progress of a change order and determine where to continue after a system crash or a retry of a failed operation. Data stored in the inbound log includes the change order's GUID, file name, originator GUID, file GUID, version number, and event time. You can view change orders in the inbound log by using Ntfrsutl.exe with the **inlog** parameter.

**Outbound change order table (outbound log)**
The outbound change order table, referred to as the outbound log, stores pending change orders to be sent to outbound partners. By default, change orders remain in the outbound log for seven days, even if all replication partners have received the change. Also in the outbound log is the leading (next change) and trailing (last acknowledged) index for each partner. You can view change orders in the inbound log by using Ntfrsutl.exe with the **outlog** parameter.

**FRS Change Orders**
A change order is a message that contains information about a file or folder that has changed on a replica member. Change orders are stored in the inbound and outbound logs on each replica member and contain information such as the change order's GUID, the originator GUID, the file GUID, the file name, event time, and file version number.

You can view change orders in the inbound and outbound logs by using Ntfrsutl.exe with the **inlog** and **outlog** parameters. You can also use the Windows Support Tools script lologsum.cmd to parse the Ntfrsutl output into a more readable format.

FRS uses the following five types of change orders:

**Local change order**
A change order that is created because of a change to a file or folder on the local server. The local server becomes the originator of the change order and constructs a staging file.

**Remote change order**
A change order received from an upstream partner that originated elsewhere in the replica set.

**Retry change order**
A change order that is in some state of completion but was blocked for some reason and must be retried later. This is a change order property in that both local and remote change orders can become retry change orders.

**Directed change order**
A change order that is directed to a single outbound partner and primarily produced when the partner is performing an initial sync (vvjoin) or a nonauthoritative restore (D2).

**Reanimation change order**
A change order that FRS uses to reconcile a recently received change order against a previously deleted file or folder.

**FRS Registry Entries**
FRS registry entries are located in the registry under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs. For a list of FRS registry entries, see "FRS Tools and Settings."

**File Replication Service Event Log**
The File Replication service event log contains events logged by FRS. It is created when FRS is first enabled on a server, and the

creation process involves setting registry keys under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Eventlog\File Replication Service. The event log file name is NtFrs.Evt, which is located under *systemroot*\System32\Config.

**FRS Objects in Active Directory**

For FRS to function properly, certain critical Active Directory objects (as well as their attributes and parent containers) must exist in Active Directory. These objects, which define a replica set's topology, schedule, and filters, are created in Active Directory when you use the Distributed File System snap-in to configure replication for DFS or when you promote a server to a domain controller using the Active Directory Installation Wizard (Dcpromo.exe). Active Directory replication is used to replicate these objects to all domain controllers in a domain, and missing or corrupted objects can cause FRS replication to fail.

**Hierarchy of FRS Objects in Active Directory for DFS Replica Sets**

The following figure illustrates the hierarchy of FRS-related Active Directory containers and objects for DFS replica sets. The figure also illustrates how several of these objects are linked together using reference attributes. For example, in the NTFRS Member object, there is an attribute called **frsComputerReference** that refers to the computer object for the server.

**FRS Containers and Objects in Active Directory for DFS Replica Sets**




The following sections briefly describe the Active Directory objects, containers, and their attributes for DFS replica sets. The class name of each object is shown in parentheses.

**NTFRS Settings object**

The NTFRS Settings object (class **nTFRSSettings**) is used as a container for the NTFRS Replica Set objects. The NTFRS Settings object can contain other NTFRS Settings objects; therefore, it provides a way to form a hierarchy to better organize the NTFRS Replica Set objects.

**NTFRS Replica Set object**

Every NTFRS Replica Set object (class **nTFRSReplicaSet**) represents a set of computers that replicate a specified folder tree and a common set of data among them. There is one NTFRS Replica Set object for every replica set. The NTFRS Replica Set object must be directly under an NTFRS Settings object.

The most commonly used attributes on this object are **fRSReplicaSetType** (set to 3 for DFS replica sets), **fRSFileFilter**, **fRSDirectoryFilter**, and **Schedule**. If you set the **Schedule** attribute, it applies to all the NTDS Connection objects in the replica set that do not have a **Schedule** attribute.

**NTFRS Member object**

Every NTFRS Member object (class **nTFRSMember**) corresponds to a computer that is part of the replica set. The relationship between the NTFRS Member object and the computer object is indicated by the **frsComputerReference** attribute. The NTFRS Member object can contain one or more NTDS Connection objects that define the inbound partners that a member replicates from. NTDS Connection objects refer to other member objects in the same replica set object using the **fromServer** attribute.

**NTDS Connection object**

NTDS Connection objects (class **nTDSConnection**) define the inbound and the outbound partners of a replica member. NTDS Connection objects are located under the NTFRS Member object in the domain naming context for DFS replica sets.

The NTDS Connection object is inbound to the NTFRS Member object that it is located under, and it is outbound from the NTFRS Member object that its **fromServer** attribute points to.

**Computer object**
Each computer object represents a computer in the domain. The relationship between the NTFRS Member object and the computer object is indicated by the **frsComputerReference** attribute.

**NTFRS Subscriptions container**
The NTFRS Subscriptions object (class **nTFRSSubscriptions**) is similar to the NTFRS Settings object in that it is primarily used as a container to group NTFRS Subscriber objects. The **fRSWorkingPath** attribute defines the location of the Ntfrs.jdb file, which is typically located in the %SystemRoot%\Ntfrs folder.

**NTFRS Subscriber object**
Every NTFRS Subscriber object (class **nTFRSSubscriber**) under a computer's computer object corresponds to a replica set that the computer is a member of. The **fRSMemberReference** attribute of the NTFRS Subscriber object points to the NTFRS Member object of the replica set that it corresponds to. Every NTFRS Subscriber object also has both an **fRSRootPath** attribute that specifies the folder tree to replicate and an **fRSStagingPath** attribute that specifies the folder to store the staging files under.

**Hierarchy of FRS Objects in Active Directory for SYSVOL Replica Sets**
SYSVOL replica sets have a different hierarchy of objects in Active Directory than do DFS replica sets. This hierarchy is illustrated in the following figure.

**FRS Containers and Objects in Active Directory for SYSVOL Replica Sets**




The following sections describe any attributes and relationships that are specific to SYSVOL replica sets. In some cases, the objects for DFS and SYSVOL replica sets are identical.

**NTFRS Settings object**
Same as for DFS replica sets.

**NTFRS Replica Set object**
Same as for DFS replica sets except that only one NTFRS Replica Set object can be of the SYSVOL type in a domain. In this case, the **fRSReplicaSetType** attribute is set to 2.

**NTFRS Member object**
In SYSVOL replica sets, the **serverReference** attribute of the NTFRS Member object points to the NTDS Settings objects (class **nTDSDSA**) that contain the NTDS Connection objects that this member replicates from.

**NTDS Connection object**

For SYSVOL replica sets, FRS uses both manually generated connection objects and connection objects that are generated by Knowledge Consistency Checker (KCC) that are located in the NTDS Settings object in the configuration naming context. (You can use the Active Directory Sites and Services snap-in to view these connection objects.) These connection objects are also used during Active Directory replication.

For SYSVOL replica sets, the NTDS Connection object is inbound to the NTFRS Member object that corresponds to the NTDS Settings object that the NTDS Connection object is located under. It is outbound from the NTFRS Member object that corresponds to the NTDS Settings object that its **fromServer** attribute points to.

**Computer object**
Same as for DFS replica sets.

**NTFRS Subscriptions container**
Same as for DFS replica sets.

**NTFRS Subscriber object**
Same as for DFS replica sets.

**FRS Processes and Interactions**
The following figure illustrates four categories of FRS processes and interactions described in this section.

**FRS Processes and Interactions**




The following sections describe the four categories of processes and interactions.

**How changes to topology, schedule, filters, and settings are detected**
FRS periodically polls Active Directory to obtain new or updated information about replica set topology, schedule, and file and folder filters. FRS also polls the registry to detect updated settings. For information about the Active Directory and registry polling process and the intervals at which polling occurs, see “How FRS Polling Works” later in this section.

**How new and existing partners communicate**
Replica members communicate with each other to determine whether new or changed files and folders need to be replicated from one member to another. When a downstream partner joins an upstream partner for the first time, the process is known as a version vector join (vvjoin). For more information, see “How the Vvjoin Process Works” later in this section. Routine communication between members is described in “How the Partner Join Process Works” later in this section.

**How replica sets are created and how members are added and removed**
The creation of new replica sets and the addition and removal of replica members involves many processes. For example, when a new replica member is added to a replica set, the FRS physical structures are created on the member and in Active Directory, the FRS service is started, registry keys are created, the member polls the registry and Active Directory, the new member vvjoins with an upstream partner, and so on.

The following sections describe these processes and interactions:

- What Happens When You Create a DFS Replica Set
- "What Happens When SYSVOL Is Created During Domain Controller Promotion"
- How Prestaging Works
- What Happens When You Remove a Member from a Replica Set

**How changed files and folders are tracked and replicated**
A number of processes occur when a file or folder is changed, closed, and then replicated from one member to another. Change orders are generated and flow from one member to another, staging files are built in the staging folder, and file and folder filters prevent files and folders with specified names or extensions from replicating. Replication takes place according to a schedule, change orders are periodically removed from the outbound log, and details about these processes are written to the FRS debug logs. You can find information about these processes in the following sections:

- Types of Changes That Trigger Replication
- How Files and Folders Are Replicated
- How the Staging Folder Works
- How Replication Schedules Work
- How File and Folder Filters Work
- How the Outbound Log Cleanup Process Works
- How FRS Debug Logs Work

**How FRS Polling Works**
FRS periodically polls Active Directory (using LDAP) to obtain new or updated information about replica set topology, schedules, and file and folder filters. (This information is stored in the Active Directory objects described in "FRS Objects in Active Directory" earlier in this section.) FRS also polls the registry to detect updated settings, such as staging directory size, USN journal size, debug log settings, and so forth. Polling takes place after FRS starts up, such as when the host computer or the service starts (when **net start ntfrs** is typed at the command line, for example) and during regular intervals after that.

The following sections describe the polling process and the intervals at which polling occurs.

**How Active Directory Polling Works**
The Active Directory polling process is as follows. Note that if any of these steps fails, replication does not occur. Failure is typically caused by network connectivity issues or missing, damaged, or duplicate Active Directory objects.

1. FRS locates the computer object for the server it is running on and enumerates the NTFRS Subscriber objects.
2. For a given subscriber, FRS follows the **fRSMemberReference** attribute to the NTFRS Member object
3. Next, FRS reads the NTFRS Replica Set object (always directly above the NTFRS Member object) to determine whether it is SYSVOL or DFS replica set.

After FRS determines the replica type (SYSVOL or DFS), FRS enumerates the connection objects. This process is different for SYSVOL and DFS.

**How FRS enumerates connection objects for DFS replica sets**

1. FRS finds all the NTDS Connection objects that represent an inbound (or upstream) connection to this member. (All NTDS Connection objects that represent an inbound connection are directly under the NTFRS Member object.)

2. FRS finds all the NTDS Connection objects that represent an outbound (or downstream) connection to this member. (All NTDS Connection objects that represent an outbound connection have this member's name in the **fromServer** attribute.)

3. FRS forms a list of NTFRS Member objects that have an inbound or outbound connection with this member. This list is a combined list for all the replica sets that this server is a member of.

4. FRS forms a list of the computers corresponding to the members in the member list. All members have an **frsComputerReference** attribute pointing to the computer object.

5. FRS finds the DNS name of each of the computers in the computer list. This is the name that FRS uses to bind to its inbound and outbound partners.

**How FRS enumerates connection objects for SYSVOL replica sets**

1. FRS finds all the NTDS Connection objects that represent an inbound connection to this member. The **serverReference** attribute on the NTFRS Member object points to the NTDS Settings object. (All the NTDS Connection objects that represent an inbound connection are directly under the NTDS Settings object.)

2. FRS finds all NTDS Connection objects that represent an outbound connection to this member. (All NTDS Connection objects that represent an outbound connection have this member's name in the **fromServer** attribute.)

3. FRS forms a list of member objects that have an inbound connection or outbound connection with this member. Forming the list of members is a two-step process for SYSVOL replica sets. First, FRS forms the list of NTDS Settings objects, and then FRS finds the NTFRS Member objects that point to each of these NTDS Settings objects. This list is a combined list for all the replica sets that this server is a member of.

4. FRS forms a list of the computers corresponding to the members in the member list. All members have an **frsComputerReference** attribute pointing to the computer object.

5. FRS finds the DNS name of each of the computers in the computer list. This is the name that FRS uses to bind to its inbound and outbound partners.

**How Registry Polling Works**

FRS polls Active Directory and the registry at the same interval. Some settings stored in Active Directory are also stored in the registry, but the Active Directory settings override the registry settings. Though some registry changes are detected and applied during polling, other changes are not applied until the FRS service is stopped and restarted. For more information about registry entries that require the service to be restarted, see Registry Reference in Tools and Settings Collection.

**Polling Intervals**

Each time a replica member is started or the FRS service is started, FRS polls the registry and the computer object in Active Directory in eight short intervals. If no changes in Active Directory are detected, FRS polls in long intervals until it detects configuration or subscriber list changes in Active Directory. The lengths of the short and long intervals are as follows:

- For domain controllers, the default short and long polling intervals are five minutes each. (Domain controllers always use the short interval, regardless of the long interval setting.)

- For member servers, the default short polling interval is five minutes and the long polling interval is 60 minutes.

### Note

- These polling intervals are specified in the **DS Polling Long Interval in Minutes** and **DS Polling Short Interval in Minutes** registry values. To view the current interval and the length of the long and short polling interval, use the Ntfrsutl.execommand with the **poll** parameter.

If there have been no FRS configuration changes in Active Directory after the completion of eight short polling intervals, FRS automatically begins polling according to the long polling interval. If one of the following events occurs, the polling cycle is reset:

- A replica member is added or removed.
- A connection is added or removed.
- A schedule is changed.
- A file or folder filter is changed.

Unlike changes in Active Directory, registry changes do not affect the polling cycle.

**How the Vvjoin Process Works**

A version vector is an array that contains pairings of the member originator GUID of the originating member and the VSN of the change. The VSN is a monotonically increasing sequence number assigned to each change that originates on a given replica member. Although the VSN is similar to the USN, FRS does not store the actual USN of the change in the version vector because a replica set can be moved to a different volume, or the volume could be formatted and the replica tree restored from backup. In these cases, the USN could be reset.

The following graphic illustrates a version vector on three servers, ServerA, ServerB, and ServerC.

**Example of Three Version Vectors**




Each entry in the version vector is considered the high water mark as to the most recent change order received from the corresponding originator. For example, the version vector on ServerA indicates that this member has received and processed all changes up to VSN 9 originating on ServerC, VSN 12 originating on ServerD, and VSN 32 originating on ServerE. ServerB's version vector shows that its replica tree is not as up-to-date as the replica tree on ServerA. The new member, ServerF, has an empty version vector.

A version vector join (vvjoin) is the process by which a downstream partner joins with an upstream partner for the first time. During this process, the two members compare version vectors to determine which files need to be replicated from the upstream partner to the downstream partner. Vvjoins can occur in a number of scenarios, which are explained using the following figure as an example.

**Sample Replica Members for Vvjoin Scenarios**




In this figure, ServerA and ServerB are members of a replica set. ServerA is upstream and ServerB is downstream for connection C. Using this figure as an example, a vvjoin can occur in the following scenarios:

**Scenario A**

Connection C is new, or connection C is deleted and a new connection is created in its place.

**Scenario B**

Connection C is disabled and later enabled.

**Scenario C**

ServerA polled a domain controller that has not seen the NTDS Connection object for connection C yet. At this point ServerA acts as though the connection is deleted. Later the NTDS Connection object for connection C replicates to this domain controller and ServerA reads it. The effect is similar to Scenario B.

**Scenario D**

ServerB is a newly added member of the replica set.

**Scenario E**

ServerA or ServerB performs a nonauthoritative restore (D2).

**Scenario F**

ServerA performs an authoritative restore (D4), and ServerB performs a nonauthoritative restore (D2).

**Scenario G**

ServerB is offline for longer than the **Outlog Change History In Minutes** registry entry, which is 10080 minutes (seven days) by default.

Vvjoins take place when the schedule first opens between the two members in the previous scenarios. The type of vvjoins performed in these scenarios can be either full or optimized vvjoins. These vvjoins are explained in the following sections.

#### When Optimized Vvjoins Occur

In scenarios A, B, C, and D, FRS attempts to perform an optimized vvjoin instead of a full vvjoin. An optimized vvjoin occurs when the upstream partner's outbound log contains all the changes that the downstream partner is missing.

Optimized vvjoins are desirable in scenarios where a downstream partner (ServerB in the previous figure) needs to vvjoin and has almost all files in the replica tree already. For example, assume the connection between ServerA and ServerB is new, but ServerB was offline for two days and only missed changes from a prior upstream partner, say ServerC, for those two days. In this case, the new upstream partner (ServerA) retains the change orders in the outbound log for seven days and sends its outbound log to the downstream partner (ServerB) instead of enumerating its entire database. ServerA also sends its existing staging files, so ServerA needs to regenerate staging files only for staging files that were purged if the staging folder reached 90 percent of its limit or that were purged during the staging folder cleanup process. The staging folder cleanup process is described in "How the Outbound Log Cleanup Process Works" later in this section.

#### When Full Vvjoins Occur

If the upstream partner's outbound log does not contain all the changes that the downstream partner is missing, a full vvjoin occurs. A full vvjoin requires the upstream partner to enumerate its entire database before determining whether to send changed files and folders to the downstream partner. As a result, the full vvjoin process is not as efficient as an optimized vvjoin, but a full vvjoin does not necessarily mean that the upstream partner sends the entire replica across the network to the downstream partner. The following sections describe the scenarios in which full vvjoins are performed and whether all or part of the data is replicated.

#### When a new replica member is prestaged and added to the replica set

In scenario D, the new member (ServerB) has been prestaged with a restored version of the replica tree. When you prestage files, you use a backup program to back up the replica tree and restore it on the new member. Using a backup program is required to preserve the file GUIDs and security descriptors.

Although it might seem that an optimized vvjoin would work in this scenario, optimized vvjoins do not look for prestaged content. Therefore, you must either clear the outbound log of the direct upstream partners or wait until the replica set is at least seven days old (or the duration specified in the **Outlog Change History In Minutes** registry value). Either of these steps will trigger a full vvjoin, which looks for prestaged content and replicates to the new member only the files and folders that have changed since the new member was prestaged. If these steps are not followed, an optimized vvjoin is performed instead, causing all the data in the replica to be replicated across the network to the new member.

**Note**

- The **Outlog Change History In Minutes** registry entry does not take effect until the next polling interval or until you use the **Ntfrsutl poll /now** command.

For more information about prestaging, see "How Prestaging Works" later in this section.

#### When a new replica member is not prestaged and the replica set is more than seven days old

In scenario D, if the new member (ServerB) has not been prestaged with a restored copy of the replica tree, and the replica set is more than seven days old (or older than the duration specified in the **Outlog Change History In Minutes** registry value), ServerA enumerates its entire database, sends directed change orders for every file and folder in the replica tree, generates staging files for every file in the replica tree, and then sends all files in the replica tree across the network. In this scenario, the full vvjoin requires a lot of CPU time, disk I/O, and network bandwidth.

**An existing replica member was offline for more than seven days**
In scenario G, if a server rejoins a replica set after being offline for more than seven days (or longer than the duration specified in the **Outlog Change History In Minutes** registry value), the upstream partner enumerates its database but only sends across the network files and folders that have changed since the server was offline.

For more information about what happens when a server is offline for more than seven days, see "How the Outbound Log Cleanup Process Works" later in this section.

**How the Partner Join Process Works**
The partner join process is a handshake between two existing partners before they begin replicating. The process begins after any of the following events:

- The connection schedule opens.
- The FRS service is started while the schedule is open.
- A connection is reestablished after the loss of network connectivity or server shutdown while the schedule is open.

During the partner join process, which is similar to a server message block (SMB) session setup, the partners negotiate capability level, exchange version vectors, and test for compression support. The schedule must be open for a partner join to occur. On a given connection, the upstream partner has the leading (next change) and trailing (last acknowledged) indexes and calculates what files need to be replicated to the downstream partner.

**What Happens When You Create a DFS Replica Set**
After you use the Configure Replication Wizard in the Distributed File System snap-in to configure replication for DFS link targets, the following events take place:

1. FRS-related Active Directory objects and their attributes are created on the server acting as the primary domain controller (PDC) emulator master. (Note that these objects will eventually replicate to all domain controllers.)
2. The FRS service is started on all members and set to automatic, if the members are not currently part of another replica set. The File Replication service event log is also created the first time the service starts.
3. Active Directory is immediately polled after the FRS service starts. If the FRS service was already running, the poll takes place according to the current polling interval. (After FRS detects a change, it begins the eight short polling intervals as described in "How FRS Polling Works" earlier in this section.)
4. If the Active Directory objects have replicated to the domain controller that FRS polls, FRS verifies that the replica root is a valid path, verifies that the staging directory is a valid path (and creates it if necessary), and checks the database path and creates the database. If the Active Directory objects are not yet present on the domain controller that FRS polls, there is a delay in this process until the Active Directory objects are replicated to this domain controller.

The final steps of the process differ for the initial master and the remaining replica members.

**For initial master:**

1. On the initial master, FRS walks the replica tree, adds the files and folders to the file ID table, and stamps a file GUID on every file and folder.
2. FRS declares itself as online and sends a join command to the other members of the replica set. This confirms that FRS is ready to replicate.

**For the remaining replica members:**

1. The remaining replica members come online in the seeding state, which indicates that they are ready to seed (replicate in) the data from an online server.
2. Eventually the remaining replica members vvjoin with other online members.

**What Happens When SYSVOL is Created During Domain Controller Promotion**
The SYSVOL shared folder is built by the Active Directory Installation Wizard (Dcpromo) during the installation of Active Directory. The process is as follows:

1. Dcpromo calls FRS to prepare for promotion. If FRS is already running on the server to be promoted, Dcpromo stops the FRS service.

2. Dcpromo deletes information from previous demotion or promotions (primarily FRS-related registry keys).

3. The Net Logon service stops sharing the SYSVOL shared folder (if it exists), and the **SysvolReady** registry entry is set to 0 (false).

4. Dcpromo creates the SYSVOL folder and the necessary subfolders and junction points.

5. The FRS service is started.

6. Dcpromo makes a call to FRS to start a promotion thread that sets the necessary registry keys.

7. Dcpromo reboots the server.

8. When the server is restarted, FRS detects that the server is a domain controller and then checks the registry for the **SysVol Information is Committed** registry entry. Because this entry is set to 0 (false), FRS creates the necessary Active Directory objects and then populates information from the registry to the Active Directory objects and creates the reference attributes as necessary.

9. FRS begins to source the SYSVOL content from the computer that is identified in the **Replica Set Parent** registry entry. (This key is temporary and is deleted after SYSVOL has been successfully replicated.) The connection to the server specified in Replica Set Parent is also temporary. This connection, called a volatile connection, is used to perform the initial vvjoin so that the new domain controller does not need to rely on Active Directory replication for the new connection objects to be replicated.

10. When SYSVOL is finished replicating, FRS sets the **SysvolReady** registry entry to 1 (true), and then the Net Logon service shares the SYSVOL folder and publishes the computer as a domain controller.

**How Prestaging Works**
Prestaging is the process of restoring a version of the replica tree on a new server before adding it to a replica set, thus avoiding the need to send the replica tree across the network. Prestaging works for both DFS and SYSVOL replica sets, though the process is slightly different:

- Prestaging a DFS replica member requires you to restore a backup of the replica tree on a server before you add the server to an existing replica set.

- Prestaging a domain controller requires you to restore a system state backup (which includes the SYSVOL shared folder) to be used as the data source when promoting a server running Windows Server 2003 to a domain controller.

**How Prestaging DFS Replica Sets Works**
The following procedure describes how FRS responds when you prestage a DFS replica set.

1. Replication is enabled between two members, Server1 and Server2. The initial master, Server1, does not have any data in its replica tree when replication is enabled.

2. After you copy files into the replica tree (using \\Server1\Apps as an example replica root), FRS generates staging files and sends change orders to the downstream partner (Server2). FRS also generates an MD5 (a hash algorithm) checksum during the staging file generation and saves the result in the file ID table on Server1 and in the change order sent to Server2.

3. When Server2 processes this change order, it saves the MD5 checksum in the file ID table on Server2. This process is the only way an MD5 checksum is saved in the file ID table and the presence of the MD5 is necessary to avoid overhead

when new members are added later.

When step 3 is finished, the replicated files exist on both Server1 and Server2, and both file ID tables have MD5 checksums for each file and folder in the replica tree. (You can verify this by using the Ntfrsutil.exe command with the **idtable** parameter. Search for "MD5CheckSum" in the output.)

4. Windows Backup or another backup program is used to back up the contents of the replica tree from either Server1 or Server2. The backup is then restored to another server, for example, Server3.

5. If fewer than seven days have passed since the replica set containing Server1 and Server2 was created, the outbound log must be cleared so that a full vvjoin is triggered when the next member joins. (Setting the **Outlog Change History In Minutes** registry entry to 0 clears the outbound log.)

6. When Server3 is added to the replica set using the Distributed File System snap-in, FRS on Server3 moves all files from the restored target folder (\\Server3\Apps) to the pre-existing data folder, and then initiates a full vvjoin from each computer that Server3 has inbound NTDS Connection objects. Note that the vvjoin process is serialized, meaning that the downstream partner performs a vvjoin with one upstream member at a time. Subsequent vvjoins are much faster because the downstream partner will typically have a complete (or nearly complete) replica tree after the first vvjoin completes.

   The key requirement in this step is that Server3 has inbound connections from an upstream partner, Server1 or Server2 in this case, whose file ID table contains MD5 checksums for files contained in the replica set of interest.

7. FRS on Server1 enumerates all the files and folders in its file ID table and sends directed (that is, single target) change orders to Server3. Because the file ID table has an MD5 checksum, it is included in the change order. As Server3 processes these change orders, this server takes the file GUID for the file or folder from the change order and attempts to locate the corresponding file in the pre-existing data folder. If the server locates the file, it re-computes the MD5 checksum on the content of that file, compares the result to the MD5 checksum it received in the change order and, if they match, uses the pre-existing file instead of attempting to obtain the file from Server1. If Server3 does not find the file, or if the MD5 checksum or attributes do not match, the server obtains the file from Server1. Any change to the file content, such as to the access control lists or data streams, can cause an MD5 mismatch and the file is obtained from Server1 or other upstream partner.

   When all replication activity has settled out, the file ID tables on all three servers have an identical MD5 checksum and identical file content in the replicated folder.

**How Prestaging SYSVOL Works**

The process for prestaging SYSVOL is part of the overall process of using backup media to create additional domain controllers. FRS follows the same method of determining whether the restored files are identical to those on the upstream partner as it does for DFS replica members. Specifically, the following requirements must be met for SYSVOL prestaging to be successful:

- FRS must have constructed MD5 checksum data for the files in the SYSVOL tree. MD5 checksum data is constructed after one of the following events occurs:

    - Every file in SYSVOL is modified after there are two or more domain controllers in the domain.

    - All data is moved to a non-replicated folder outside of SYSVOL and then moved back into SYSVOL after there are two or more domain controllers in the domain.

- The system state backup must contain all files in SYSVOL. If any files in the system state backup are out of date, those files are replicated across the network to the new member.

- The outbound log on the upstream partner must be cleared if SYSVOL is less than seven days old. As described in "How the Vvjoin Process Works" earlier in this section, if the replica set is less than seven days old, FRS attempts an optimized vvjoin, which does not look for prestaged content and will cause the entire replica tree to be replicated across the network to the new member. To work around this, you can clear the outbound log on all upstream partners, or you can clear the outbound log on one upstream partner and then do one of the following steps to make sure that the new

member vvjoins with this partner:

- Create a single inbound connection to the member with the cleared outbound log.
- Set up connection schedules such that the parent with the cleared outbound log is the only schedule that is open.
- Stop the FRS service on all other upstream partners.

**Note**

- When prestaging SYSVOL, do not clear the outbound log on bridgehead servers. Bridgehead servers typically have many outbound connections, and clearing the outbound log can cause full vvjoins to occur for every connection that has pending outbound change orders. If this occurs, the bridgehead server can experience a significant decrease in FRS performance and significantly increased CPU, memory, and bandwidth usage. To avoid this situation, we recommend that you prestage from an upstream replica member that has few outbound connections.

**What Happens When You Remove a Member from a Replica Set**

When a member is removed from a replica set, certain events occur both on the removed member and in Active Directory. The following sections describe what happens when a domain controller is demoted (and thus removed from the SYSVOL replica set) and what happens when you remove a member from a DFS replica set.

**Demoting a Domain Controller**

Dcpromo (via Ntfrsapi.dll) performs the following steps during a domain controller demotion:

1. Dcpromo prepares for demotion by:
   a. Stopping FRS.
   b. Cleaning out old demotion state in the registry (if any).
   c. Restarting FRS.
   d. Binding to Active Directory using a different domain controller in the domain. The binding occurs because the demotion will invalidate the local domain controller and thus Active Directory replication would not take place after the system restarts. (If this is the last domain controller in the domain then this step is not performed.)
2. Dcpromo starts the demotion by:
   a. Setting the SYSVOL **Replica Set Command** registry entry to Delete.
   b. Instructing FRS to tombstone the SYSVOL replica set.
3. Dcpromo commits the demotion by:
   a. Stopping the FRS service and setting the service to manual.
   b. Setting the **Sysvol Ready** registry entry to 0.
   c. Deleting the NTFRS Subscription object and NTFRS Replica Set object for SYSVOL.

**Removing a Member from a DFS Replica Set**

The following process occurs when you use the Distributed File System snap-in to remove a replica member from a replica set.

1. If the topology is hub and spoke and the hub member is removed, the topology type is changed to custom.
2. The NTFRS Subscriber object and its container are deleted.

3. The connections are adjusted based on current topology preference. For example, if the topology is a ring topology, a ring topology is reformed using the remaining members.

4. The connections from the removed member to other members are deleted.

5. The NTFRS Member object is deleted.

**Types of Changes That Trigger Replication**

Because a file is a basic unit of replication, any change to a file will cause the entire file to be replicated after the handle to the file is closed. (Replication takes place either immediately or when the replication schedule opens.) The following table provides a comprehensive description of all changes that can cause a file to be replicated. This table includes the code logged in the FRS outbound logs to indicate the type of change that caused replication to occur. These changes must be followed by a close record in the USN journal before replication begins.

**Types of Changes That Trigger Replication**

| Change Code in Outbound Log | Explanation |
|---|---|
| Create | A file or folder was created. |
| Delete | A file or folder was deleted. |
| RenNew | A file or folder was renamed. (This type of change also applies to file or folders deleted locally using Windows Explorer, in which case the file or folder is renamed and moved to the Recycle Bin.) |
| DatOvrWrt | Main file data stream was overwritten. |
| DatExt | Main file data stream was extended. |
| DatTrunc | Main file data stream was truncated. |
| Info | Basic information was changed (hidden and read-only attributes, last write time, and so on) Note that some types of basic info changes, such as archive flag or the last access time, do not cause files to be replicated. |
| Oid | Object ID change. |
| StreamNam | Alternate data stream name change. |
| StrmOvrWrt | Alternate data stream was overwritten. |
| StrmExt | Alternate data stream was extended. |
| StrmTrunc | Alternate data stream was truncated. |
| EAChg | Extended file attribute was changed. |
| Security | File permissions, auditing, or ownership were changed. |
| CompressChg | File compression attribute changed. |
| EncryptChg | File encryption changed. Note this type of change triggers replication only when the file is changed from encrypted to not encrypted. |
| Reparse | Reparse point changed. |

The following changes do not trigger replication:

- Changes to a file or folder's last access time

- Changes to a file or folder's archive bit

- Changes to encrypted files
- Changes related to reparse points

The process that FRS uses to determine whether a file has changed is as follows:

1. FRS monitors the USN journal for changes. When FRS detects a close record for a file (a handle being closed will trigger this), FRS gathers relevant information about the recently closed file, including the file's attributes and MD5 checksum, from the file ID table.

2. FRS computes an MD5 checksum for the recently closed file. The MD5 checksum is calculated based on the file's data, including its security descriptors. File attributes are not included in this calculation. If the MD5 checksum and attributes of the recently closed file are identical to the information about the file stored in the file ID table, the file is not replicated. If either the MD5 checksum or the attributes are different, FRS begins the change order process described in "How Files and Folders Are Replicated" later in this section.

Certain types of programs can make identical updates a file without actually changing the file's content (in other words, the file's MD5 checksum for the last change is identical to the previous change to that same file). These programs typically include backup, defragmentation, and antivirus programs that were not written to be compatible with FRS. Setting file system policy on files by using Group Policy can also cause frequent identical updates that trigger replication. These types of updates include:

- Overwriting a file with a copy of the same file.
- Setting the same ACLs on a file multiple times.
- Restoring an identical copy of the file over an existing one.

FRS suppresses excessive replication of files caused by these types of updates and logs FRS event ID 13567 in the File Replication service event log. More specifically, this event is logged when FRS detects that 15 identical updates were made to FRS replicated files within a one-hour period and this condition has occurred over three consecutive hours. The duplicate changes might have occurred on a single file 15 times, or 15 unique files one time each, or any combination between those two.

Note that these thresholds apply only to event logging; FRS suppression works continuously. In addition, FRS suppression does not apply to folders. Therefore, frequent updates to folders affect replication and staging even when suppression is turned on.

**How Files and Folders Are Replicated**

The following figure describes the change order flow process, which begins when a file or folder is added to a replica tree or when a file or folder is changed and then closed in a replica tree. The sections that follow describe each step in the process. The differences in the flow process for file deletions are also described.

### How Change Orders Are Processed




**Step 1: NTFS creates a change journal entry**
When a file in a replica set is changed and closed, NTFS makes an entry in the USN journal. The USN journal records changes to all files on the NTFS volume (such as file creations, deletions, and modifications). In Windows Server 2003, the default journal size is 128 MB (or 512 MB in Windows Server 2003 with SP1 or the pre-SP1 release of Ntfrs.exe) and is persistent across restarts and crashes. If the FRS service is stopped or fails, this will have no effect on the replication of FRS content. Even with the service stopped, file changes are still recorded at the file system level in the USN journal. When the FRS service is restarted, stored USN journal entries will trigger replication.

**Step 2: FRS monitors the USN journal**
FRS monitors the USN journal for changes that apply to the replica trees on ServerA. Only closed files are checked. File and folder filters are applied against changes in the files and folders in the replica tree.

**Step 3: Aging cache is used**
The aging cache is a three-second delay designed to catch additional changes to a file. This prevents replication when a file is undergoing rapid updates.

**Step 4: FRS creates entries in the inbound log and file ID table**
ServerA records the change as a change order in its inbound log. It also creates an entry in the file ID table so that recovery can take place if a crash occurs. The inbound log contains change orders arriving from all inbound partners. Change orders are logged in the order that they arrive. Each change order contains information about a change to a file or folder on a replica member, such as the name of the file or the time it was changed. This information is used to construct a message about the change.

**Step 5: FRS creates the staging file in the staging folder**
ServerA uses backup APIs to create a backup of the changed file or folder in ServerA's staging folder. These backup files, known as staging files, encapsulate the data and attributes associated with a replicated file or folder. By creating the staging file in the staging folder, FRS ensures that file data can be supplied to partners regardless of any activity that might prevent access to the original file. All staging files in the staging folder are compressed to save disk space and network bandwidth during replication.

At step 1 if a file is deleted, FRS does not create a staging file.

**Step 6: FRS creates the entry in the outbound log**
ServerA updates the outbound log. The outbound log contains change orders generated for a specified replica set. These change orders can originate locally or come from an inbound partner. Change orders recorded in the outbound log are eventually sent to all outbound partners.

If a file is deleted, the change order contains the event time of the deletion and the tombstone lifetime (by default 60 days). When the tombstone for the deleted file expires, the file's entry in the file ID table is purged.

**Step 7: FRS sends a change notification**
ServerA sends a change notification to ServerB.

**Step 8: FRS records and acknowledges the change notification**
ServerB stores the change order in its inbound log and file ID table. If ServerB decides to accept the change order, it sends a change order acknowledgement (CO ACK) to ServerA to replicate the modified file's staging file.

**Step 9: FRS replicates the staging file**
ServerB copies the staging file from ServerA to the staging folder on ServerB. ServerB then writes to its outbound log so other outbound partners can pick up the change.

If the change order is for a deleted file, no staging file is fetched. The change notification for the deleted file is added to ServerB's outbound log,

**Step 10: FRS constructs the staging file and moves it into the replica tree**
ServerB uses restore APIs to reconstruct (that is, restore) the file or folder in the preinstall folder, and then FRS renames the file or folder into the replica tree.

If the change order was for a file deletion, ServerB deletes the file from the replica tree.

**How the Staging Folder Works**
The staging folder is a queue for changes to be replicated to downstream partners. After the changes are made to a file and the file is closed, the file content is compressed, written to the staging folder, and replicated according to schedule. Any further use of that file does not prevent FRS from replicating the staging file to other members.

The following sections describe various aspects of the staging folder.

**Staging folder size**
The size of the staging folder governs the maximum amount of disk space that FRS can use to hold those staging files and the maximum file size that FRS can replicate. The default size of the staging folder is approximately 660 MB, the minimum size is 10 MB, and the maximum size is 2 terabytes. The largest file that FRS can replicate is determined by the staging folder size on both the upstream partner and downstream partners and whether the compressed replicated file can be accommodated by the current staging folder size. Therefore, the largest file that FRS can replicate is 2 terabytes, assuming that the staging folder size is set to the maximum on upstream and downstream partners.

**Staging folder compression**
FRS replica members running Windows 2000 Server with Service Pack 3 (SP3) or later or Windows Server 2003 compress the files replicated among them. Compression reduces the size of files in the staging folder on the upstream partners, over the network between compression-enabled partners, and in the staging folder of downstream partners prior to files being moved into their final location.

All files and subfolders within the staging folder are compressed. However, compression is not enabled on the staging folder itself.

**How the staging folder stays below its maximum size limit**
When FRS tries to allocate space for a staging file and is not successful because the size of the files in the staging folder has reached 90 percent of the staging folder size, FRS starts to remove files from the staging folder. Staged files are removed (in the order of the longest time since the last access) until the size of the staging folder has dropped below 60 percent of the staging folder limit. Additionally, staging files for downstream partners that have been inaccessible for more than seven days are deleted. As a result, FRS does not stop replicating if the staging folder runs out of free space. This means that if a downstream partner goes offline for an extended period of time, the offline member does not cause the upstream partner's staging folder to fill with accumulated staging files.

### Note

- During the process of deleting staging files to reach the 60 percent limit, FRS does not delete staging files that have change orders in the inbound log pending install, because they are still needed by this member.

The fact that FRS removes files from the staging folder does not mean that the underlying file is deleted or will not be replicated. The change order in the outbound log still exists, and the file will eventually be sent to downstream partners when they process the change order. However, before replication can take place, the upstream partner must recreate the staging file in the staging folder, which can affect performance. Recreating the staging file can also cause a replication delay if the file on the upstream partner is in use, preventing FRS from creating the staging file.

**Staging folder cleanup**
Staging files are cleaned up during the outbound log cleanup process. Specifically, staging files are deleted from the staging folder after all replica members have acknowledged the change orders that generated the staging files. For more information about the outbound log cleanup process and its effect on the staging folder, see "How the Outbound Log Cleanup Process

Works" later in this section.

**How Replication Schedules Work**
The replication schedule defines the periods during which replication takes place. The following sections describe various aspects of replication schedules.

**SYSVOL Schedules**
The SYSVOL schedule is an attribute associated with each NTFRS Replica Set object and with each NTDS Connection object. FRS replicates SYSVOL using the same intrasite connection objects and schedule built by the KCC for Active Directory replication. FRS uses two replication protocols for SYSVOL:

- **SYSVOL connection within a site**. The connection is always considered to be on; any schedule is ignored and changed files are replicated immediately.

- **SYSVOL connection between sites**. SYSVOL replication is initiated between two intersite members at the start of the 15-minute interval, assuming the schedule is open. The connection is treated as a trigger schedule. The upstream partner ignores its schedule and responds to any request by the downstream partner. When the schedule closes, the upstream partner unjoins the connection only after the current contents of the outbound log, at the time of join, have been sent and acknowledged.

The SYSVOL replication schedule can be configured so that replication takes place four, two, one, or zero times per hour.

- When set to four times per hour, replication starts at 0:00, 0:15, 0:30, and 0:45. This schedule is essentially continuous replication.

- When set to two times per hour, replication occurs at 15-minute intervals starting at 0:15 and 0:45.

- When set to one time per hour, replication starts at 0:00 and ends at 0:15 assuming that all changes have been replicated.

It is possible to change the SYSVOL schedule, but this should be done only after careful consideration of the implications and alternatives. Also, note that schedules exist on both site links and NTDS Connection objects. If you change the schedule on an NTDS Connection objects, the connection then becomes a manual connection that cannot be managed by KCC.

**DFS Replica Set Schedules**
For DFS replica sets, FRS uses the NTDS Connection objects, topology, and schedule built by the Distributed File System snap-in. The snap-in sets the schedule as follows:

- When you view the properties of a link and set the schedule in the ***Linkname* Properties** dialog box, the snap-in applies the schedule to all NTDS Connection objects for this replica set. The snap-in does not apply the schedule to the NTFRS Replica Set object.

- If you view individual connections in the **Customize Topology** dialog box and modify the schedule on a connection, the snap-in applies that schedule to the NTDS Connection object for the connection.

- If you add a new member to an existing replica set, the new member's NTDS Connection objects get the default "always on" schedule.

- If you use the snap-in to set schedules on individual connections, and then you view the schedule from the ***Linkname* Properties** dialog box, the snap-in displays the schedule for the last enumerated NTDS Connection object. If you click **OK** in this dialog box, the schedule shown is applied to all NTDS Connection objects. To avoid this, click **Cancel** to close the dialog box.

When replication is set to available, replication occurs continuously. Unlike SYSVOL schedules, DFS replica set schedules are on/off schedules. As long as the schedule is open, the connection stays in the joined state and change orders are sent out as they are processed by the upstream partner. Any queued change orders in the upstream partner's outbound log are sent first. When the schedule closes, FRS unjoins the connection.

**How Local Time Affects the Schedule**
Replication schedules are stored in Coordinated Universal Time (UTC). However, when you view the schedule, the schedule is shown in the local time of the server. Daylight savings time causes the schedule to shift by an hour for any server that is

located in an affected time zone.

**How File and Folder Filters Work**
Filters exclude subfolders (and their contents) or files from replication. You exclude subfolders by specifying their name, and you exclude files by using wildcard characters to specify file names and extensions. By default, no subfolders are excluded. The default file filters exclude the following files from replication:

- File names starting with a tilde (~) character
- Files with .bak or .tmp extensions

Filtering takes place when FRS monitors the USN journal for changed files and folders. If a new or updated file meets a filter that excludes it from replication, FRS ignores the change in the USN journal and does not generate a change order.

Because filtering occurs at the USN journal, filters act as exclusion filters only for new files and folders added to a replica set. They have no effect on existing files in the replica set. For example, if you change the existing file filter from "*.tmp, *.bak" to "*.old, *.bak," FRS does not go through the replica set and exclude all files that match *.old, nor does it go through the replica set and begin to replicate all files that match *.tmp. After the filter change, new files matching *.old that are added to the replica set are not replicated. New files matching *.tmp that are added to the replica set are replicated.

Any file in the replica set that was excluded from replication under the old file filter (such as Test.tmp, created when the old filter was in force) is automatically replicated under the new file filter only after the file is modified. Likewise, changes to any file that was not excluded from replication under the old filter (such as Test.old, created when the old filter was in force) continue to replicate under the new filter, until you explicitly delete the file.

These rules apply in the same manner to the folder exclusion filter. If a folder is excluded, all subfolders and files under that folder are also excluded.

Regardless of the filters you set, FRS always excludes the following from replication:

- NTFS mounted drives.
- Files encrypted by using EFS.
- Any reparse points except those associated with DFS namespaces. If a file has a reparse point used for Hierarchical Storage Management (HSM) or Single Instance Store (SIS), FRS replicates the underlying file but not the reparse point.
- Files on which the temporary attribute has been set.

**How the Outbound Log Cleanup Process Works**
The outbound log cleanup process runs periodically to remove change orders from the outbound log. The process is scheduled by the outlog thread, and the first cleanup is scheduled to run 60 seconds after the outlog thread starts. Service initialization code starts the outlog thread.

At every scheduled cleanup, the time it took to perform the last cleanup is calculated and the next scheduled time is set. The time before the next scheduled cleanup is set to 50 multiplied by the time it took to run the last cleanup. This interval is then adjusted so that it is not less than 60 seconds, it is not more than the **Outlog Change History In Minutes** registry entry (which is seven days, by default), and it is not more than eight hours. As a result, the cleanup process tries not to take more than 2 percent of the processing time. For large outbound logs, the cleanup can take several minutes to run.

The cleanup process does not need to run every time it is scheduled. The cleanup process only runs for the replica sets that have received at least one change order acknowledgement (ACK) since the last time cleanup ran on that replica set. In situations where a replica member does not receive ACKs from any partners for a long time, there is an interval after which the cleanup process runs on that replica member, even if it has not received an ACK from its partners. The interval is either the amount specified in the **Outlog Change History In Minutes** registry entry (seven days, by default) or eight hours, whichever is smaller.

The cleanup function goes through all the change orders in the outbound log and does the following.

1. Deletes all vvjoin change orders and their corresponding staging files if they have been acknowledged by the partner they were meant for. (Vvjoin change orders are deleted because they are directed change orders, which means they are generated for a specific partner and are of no use after they are acknowledged by that partner.)

2. Deletes staging files for all the change orders that have been acknowledged by all the partners. Staging files are not kept longer than required; only change orders are kept. Staging files are regenerated as needed.

3. Deletes change orders that have been in the outbound log for more than the **Outlog Change History In Minutes** registry entry. If there is a connection that has not yet sent the change order that is being deleted, that connection is marked for deletion at the next poll. The deletion of these connections will take them out of the list of active connections and trigger cleanup on the change orders that were backed up for them. Because these connections are not actually deleted from Active Directory, they will reappear at the second poll and will rejoin and go through a complete vvjoin. This functionality provides a way for FRS to clean up change orders in the outbound log and staging files for members that have not accepted changes for a long time.

**How FRS Debug Logs Work**

FRS creates text-based logs in the *systemroot*\Debug folder to help you debug problems. The Ntfrs log files store transaction and event detail in sequentially numbered files (Ntfrs_0001.log through Ntfrs_0005.log). Transactions and events are written to the log with the highest sequence number in existence at that time. The characteristics of the log files are determined by the values of several registry entries in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters subkey. After the number of logs specified by the value of the **Debug Log Files** registry entry has been filled, the lowest log version is deleted and the remaining log file names are decreased by *n*–1 to make room for a new log file.

Log detail is controlled by the value of the **Debug Log Severity** registry entry, ranging from 0 to 5, with 5 providing the most detail. Log size is determined by the value of the **Debug Maximum Log Messages** registry entry. The default value of 20,000 lines for **Debug Maximum Log Messages** results in a log file of approximately 2 MB, for a total of 10 MB of logs (**Debug Log Files** * [**Debug Maximum Log Message** * 120]). Setting **Debug Log Messages** to 50,000 results in log files of approximately 5 MB and 25 MB of total log space with default settings.

Solving problems by using the Ntfrs logs requires ensuring that the value of the **Debug Log Severity** registry entry is set high enough to capture the events needed to identify the problem. Severity settings range from 0 to 5 and are cumulative, meaning that a setting of 4 includes log events with a severity of 0 to 3. Error logs hold a severity setting of 0 and are always displayed. Setting the severity level at 5 will cause every action to be logged and can make finding more important information difficult. The default log setting is 2.

All log records have a date and time stamp and an identifying string annotated by **::** with a letter or pair of letters between each colon. These identifying strings make it easy for you to identify errors, warning messages, and milestone events in the log files and filter out less important information. For example, a log entry that has a **:U:** as an identifying string includes information related to the USN journal. A log entry containing information about Active Directory polling will have an identifying string of **:DS:**. A tracking log entry has an identifying string of **:T:** and summarizes a change order that has finished or is in the process of updating a given file or folder. The following table lists the log record identifiers.

**Log Record Identifiers**

| Identifier | Description |
|---|---|
| :: | Change order trace records |
| ++ | Continuation records |
| :DS: | DS access entries |
| :FK: | FRS registry Key entries |
| :H: | Log header entries |
| :S: | Service startup and shutdown entries |
| :SC: | Service controller entries |
| :SR: | Send / Receive command server entries |
| :T: | Tracking record |
| :U: | USN journal entries |

| | |
|---|---|
| :V: | Version Vector join entries |
| :X: | Communication entries |

Tracking record entries (**:T:**) can be a helpful way to identify and understand problems that can occur during the change order process. A tracking record entry will tell you what files have been changed and where the change originated.

The following tracking record entry describes a remote change order that is creating a new file called test_file in Replica-A. The version number is zero.

```
7/31-08:40:08 :T: CoG: d42cda60 CxtG: 000001b7 [RemCo   ] Name: test_file
7/31-08:40:08 :T: EventTime: Mon Jul 31, 2003 08:40:04 Ver: 0
7/31-08:40:08 :T: FileG: ceff96a6-5c9f-433a-989c841454a1593b FID: 61a70000 0000036c
7/31-08:40:08 :T: ParentG: 1a89f4e1-a0c0-43e4-aedbe869f767f372 Size: 00000000 00000008
7/31-08:40:08 :T: OrigG: 2eea81b4-f92d-4941-9f269d4bbdd7ea05   Attr: 00000020
7/31-08:40:08 :T: LocnCmd: Create State: IBCO_COMMIT_STARTED         ReplicaName:
Replica-A (1)
7/31-08:40:08 :T: CoFlags: 0000040c   [Content Locn NewFile ]
7/31-08:40:08 :T: UsnReason: 00000002 [DatExt ]
```

The individual fields that comprise the tracking log entry are described in the following table. In some cases only the first DWORD of a GUID is actually displayed in the log.

**Tracking Log Entries**

| **Identifier** | **Description** |
|---|---|
| :T: | Identifying string. |
| CoG: | Change order GUID - Uniquely identifies a create/delete/rename/modify action for a file. |
| CxtG: | Connection GUID - Identifies the connection object in the topology connecting an upstream member to the member that delivered this change order. |
| [ ] - RemCo<br>[ ] - RemCo, Abort<br>[ ] - LclCo<br>[ ] LclCo, Abort | Identifies a remote change order.<br>Identifies a remote changer order that was aborted.<br>Identifies a local change order.<br>Identifies a local change order that was aborted. |
| Name: | File name. |
| EventTime: | Time on the originating member at which the change was performed. |
| Ver: | Version number of the file. Increases by one each time a local change order is created. |
| FileG: | File GUID - Uniquely identifies the file or folder and is used as the NTFS object ID on the file or folder. The corresponding file/folder on each replica member has the same file GUID. |
| FID: | File ID - The NTFS volume-specific file ID (file reference number). |
| ParentG: | Parent GUID - The GUID of the parent folder that contains this file or folder. |
| Size: | The approximate size of the file or folder noted in hexadecimal. |
| OrigG: | Originator GUID - The GUID associated with the member of the replica set that originated this update. |
| Attr: | File attributes - The attribute flags for the file/folder. |
| LocnCmd: | Location command - One of Create, Delete, NoCmd, MoveDir; indicating that the file is being created, deleted, updated, or is changing parent folders. |

| | |
|---|---|
| State: | The change order state - One of IBCO_STAGING_RETRY, IBCO_FETCH_RETRY, IBCO_INSTALL_RETRY, IBCO_COMMIT_STARTED, IBCO_INSTALL_REN_RETRY, IBCO_INSTALL_DEL_RETRY, and IBCO_FETCH_REQUESTED indicating that the change order is being retried later because of inability to complete the fetch of the staging file, or inability to install the change to the file. Finished change orders have a state of IBCO_COMMIT_STARTED. |
| ReplicaName: | The name of the replica set containing this file or folder. |
| CoFlags: | Change 0rder flags<br>• Abort - Set when CO is being aborted.<br>• VVAct - Set when VV activate request is made.<br>• Content - Valid content command.<br>• Locn - Valid location command.<br>• LclCo- CO is locally generated.<br>• Retry - CO needs to retry.<br>• InstallInc - Local install not completed.<br>• Refresh - CO is an upstream-originated file refresh request.<br>• OofOrd - Don’t check/update version vector.<br>• NewFile - If CO fails, delete IDTable entry.<br>• DirectedCo - This CO is directed to a single connection.<br>• DemandRef - CO is a downstream demand for refresh.<br>• VVjoinToOri - CO is from vvjoin to originator.<br>• MorphGen - CO generated as part of name morph resolution.<br>• MoveinGen - This CO was generated as part of a sub-dir MOVEIN.<br>• OidReset - All CO did was reset the object identifier back to FRS defined value.<br>• CmpresStage - The stage file for this CO is compressed. |

| | |
|---|---|
| UsnReason: | Flags set in the NTFS change log describing modifications to the file.<br>• Close - Change log close record.<br>• Create - File or folder was created.<br>• Delete - File or folder was deleted.<br>• RenNew - File or folder was renamed.<br>• DatOvrWrt - Main file data stream was overwritten.<br>• DatExt - Main file data stream was extended.<br>• DatTrunc - Main file data stream was truncated.<br>• Info - Basic info change (attrib, last write time, and so forth).<br>• Oid - Object ID change.<br>• StreamNam - Alternate data stream name change.<br>• StrmOvrWrt - Alternate data stream was overwritten.<br>• StrmExt - Alternate data stream was extended.<br>• StrmTrunc - Alternate data stream was truncated.<br>• EAChg - Extended file attribute was changed.<br>• Security - File access permissions changed.<br>• IndexableChg - File change requires reindexing.<br>• Hlink - Hard link change.<br>• CompressChg - File compression attribute changed.<br>• EncryptChg - File encryption changed.<br>• Reparse - Reparse point changed. |

The following tracking log entry describes a local change order (a change order originating on the computer where the log was produced) that is updating the same file, test_file. The version number is now 1. Notice that the originator GUID is different from that of the tracking log entry above. The file GUID and parent GUID of both log entries are the same for both change orders because the same file is involved and it has not changed parent directories.

```
7/31-08:56:55 :T: CoG: cd55ad6f CxtG: 37b12c93  [LclCo ]  Name: test_file
7/31-08:56:55 :T: EventTime: Mon Jul 31, 2003 08:56:52   Ver:  1
7/31-08:56:55 :T: FileG: ceff96a6-5c9f-433a-989c841454a1593b FID:  61a70000 0000036c
7/31-08:56:55 :T: ParentG: 1a89f4e1-a0c0-43e4-aedbe869f767f372 Size: 00000000 00000200
7/31-08:56:55 :T: OrigG: 8f759ded-e611-43c4-be05c10138dfdea4 Attr: 00000020
7/31-08:56:55 :T: LocnCmd:   NoCmd    State: IBCO_COMMIT_STARTED      ReplicaName:
Replica-A (1)
7/31-08:56:55 :T: CoFlags:   00000024   [Content LclCo ]
7/31-08:56:55 :T: UsnReason: 00000002   [DatExt ]
```

**Network Ports Used by FRS**

FRS uses the following two network ports:

**Network Ports Used by FRS**

| Service Name | UDP | TCP |
|---|---|---|
| LDAP | | 389 |
| RPC | | Dynamic |

By default, FRS replication over remote procedure calls (RPCs) occurs dynamically over an available port by using RPC Endpoint Mapper (also known as Remote Procedure Call Server Service or RPCSS) on port 135; the process is the same for Active Directory replication.

**Related Information**

The following resources contain additional information that is relevant to this section:

- Active Directory Replication Topology Technical Reference
- Volume Shadow Copy Service Technical Reference

# FRS Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- FRS Tools
- FRS Registry Entries
- Network Ports Used by FRS


**FRS Tools**
The following tools are associated with File Replication service (FRS).

**Connstat.cmd: FRS Connection Status Report**
**Category**
Windows Support Tools for Windows Server 2003 and Windows 2000

**Version compatibility**
Runs on and collects information from servers running Windows 2000 Server or Windows Server 2003. Requires PERL 5.005 521e or 522e on the server where ConnStat is running.

ConnStat processes the output generated when the NtfrsUtl tool is run with the **sets** parameter. Using the NtfrsUtl sets output as its data source, ConnStat creates a report that summarizes the status of the FRS replication connections to and from a specified computer.

**Dfsgui.msc: Distributed File System snap-in**
**Category**
Windows Server 2003 operating system tool

**Version compatibility**
Runs on Windows Server 2003. This snap-in is also part of the Windows Server 2003 Administration Tools Pack; you can install this pack on computers running Windows XP with Service Pack 1 (SP1) or later.

The Distributed File System snap-in is used to enable replication for DFS root or link targets, create custom replication schedules and topologies, set replication filters and connection priorities, and so forth.

**FRSDiag.exe: FRS Diagnostics**
**Category**
Diagnostic tool for Windows Server 2003 and Windows 2000. This tool is available from the Microsoft Download Center on the Microsoft Web site.

**Version compatibility**
Runs on any computer running Windows 2000 Server, Windows 2000 Professional, Windows Server 2003, or Windows XP. Requires that NET Framework version 1.1 be installed on the computer where FRSDiag is running. Collects information from servers running Windows 2000 and Windows Server 2003.

FRSDiag is an easy-to-use tool that replaces the Health_Chk tool used to gather FRS information. FRSDiag provides a graphical user interface for choosing which information to gather and analyzes the results to identify common FRS and Active Directory problems.

**FrsFlags.vbs: File Replication Service Flags**
**Category**
Resource Kit Tools for Windows Server 2003

**Version compatibility**
Runs on Windows Server 2003.

FrsFlags is a command-line tool that enables and disables the Install Override feature of FRS. It sets and clears the FrsFlags attribute of replica objects in Active Directory.

**Health_chk.cmd: FRS Health Check**

**Category**
Windows Support Tools for Windows Server 2003 and Windows 2000

**Version compatibility**
Runs on and collects information from servers running Windows 2000 Server or Windows Server 2003.

Health_Chk retrieves FRS-related state information from the specified domain controller or server. It collects and provides administrators with in-depth data that can be used to monitor or troubleshoot an FRS system. Health_Chk calls a series of tools for the output they produce. It can be run against a remote server, though running the script directly on the server to be analyzed enables the tools to gather additional useful data that is not available when the script is run remotely.

**Iologsum.cmd: FRS Inbound and Outbound Logs Summary Report Tool**
**Category**
Windows Support Tools for Windows Server 2003 and Windows 2000

**Version compatibility**
Runs on and collects information from servers running Windows 2000 Server or Windows Server 2003. Requires PERL 5.005 521e or 522e on the server where IOLogSum is running.

IOLogSum processes the output generated when the NtfrsUtl tool is run with the **inlog**, **outlog**, or **idtable** parameters. Running NtfrsUtl **inlog** or NtfrsUtl **outlog** creates a summary inbound or outbound NtfrsUtl log report showing the pending change orders for the replica member. Running NtfrsUtl **idtable** creates a report with version information for each file in the replica tree; this version information tracks the last change order accepted for that particular file or directory. Running IOLogSum then sorts and formats this NtfrsUtl output.

**Ntfrsutl.exe:File Replication Utility**
**Category**
Windows Support Tools for Windows Server 2003 and Windows 2000

**Version compatibility**
Runs on and collects information from servers running Windows 2000 Server or Windows Server 2003.

NtfrsUtl dumps the internal tables, thread, and memory information for FRS.

**Sonar.exe: Sonar**
**Category**
Resource Kit Tools for Windows Server 2003 and Windows 2000. This tool is available from the Microsoft Download Center on the Microsoft Web site.

**Version compatibility**
Runs on any computer running Windows 2000 Server, Windows 2000 Professional, Windows Server 2003, and Windows XP. Requires .NET Framework version 1.1 on the computer where Sonar is running. Can be used to monitor servers running Windows 2000 and Windows Server 2003.

Sonar is a tool designed for monitoring key statistics about FRS members in a replica set. Sonar can collect basic status information from FRS running on either Windows 2000 (all service packs) or Windows Server 2003. You can use Sonar to easily watch coarse statistics on a replica set, and to monitor traffic levels, backlogs, free space, and other issues. Sonar allows definition of filters that define rules for which rows to display, and also allows definition of column sets that can be viewed.

**Topchk.cmd: DFS and SYSVOL Replication Topology Analysis Tool**
**Category**
Windows Support Tools for Windows Server 2003 and Windows 2000

**Version compatibility**
Runs on and collects information from servers running Windows 2000 Server or Windows Server 2003. Requires PERL 5.005 521e or 522e on the server where TopChk is running.

TopChk processes the output generated when the NtfrsUtl tool is run with the **ds** parameter. Using the NtfrsUtl **ds** output as its data source, TopChk creates a report that summarizes the replication topology of one or more replica sets.

**Ultrasound.exe: Ultrasound**
**Category**
Monitoring and troubleshooting tool for Windows Server 2003 and Windows 2000. This tool is available from the Microsoft Download Center on the Microsoft Web site.

**Version compatibility**
The user interface component of Ultrasound can be run on computers running Windows 2000 Server, Windows 2000 Professional, Windows Server 2003, or Windows XP. Ultrasound requires that .NET Framework version 1.1 and Microsoft Data

Access Components (MDAC) 2.6 be installed on the computer where the user interface component is running. Ultrasound also requires a database server running either Microsoft SQL Server 2000 Desktop Engine (MSDE) or SQL Server 2000 with SP3a. Ultrasound is used to monitor FRS on servers running Windows 2000 and Windows Server 2003 on which you install Ultrasound WMI providers.

Ultrasound is a powerful monitoring and troubleshooting tool used to measure the health of FRS replica sets. Ultrasound shows health ratings and historical information about replica sets and allows administrators to monitor the progress of replication and detect problems that can cause replication to become backlogged or stopped. Ultrasound also provides detailed views for troubleshooting and a framework for customizing alerts and views to suit an organization's needs.

**FRS Registry Entries**
The following registry entries are associated with FRS. The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs. The headings reflect the path beyond Ntfrs, such as Parameters or Parameters\Backup/Restore\Process at Startup, and the entries under each path are alphabetized.

**Parameters**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters.

**Compress Staging Files**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Disables the compression of staging files in the staging folder. Staging files are compressed by default.

**Debug Disable**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Disables the FRS debug log files. By default, FRS records its actions in debug log files. These log files are named NtFrs_000*n*.log and are located in the *Systemroot*\Debug directory. These files are typically used to investigate FRS replication problems.

To find more information aboutthis registry entry, see Registry Reference in Tools and Settings Collection.

**Debug Log File**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Specifies where FRS stores its debug log files. If you change the value of this entry, FRS creates new debug log files in the directory specified, but it does not move existing debug log files to the new location. Also, FRS does not delete or maintain log files in the original location.

To find more information aboutthis registry entry, see Registry Reference in Tools and Settings Collection.

**Debug Log Files**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Specifies how many FRS debug log files can be active simultaneously. FRS creates a set of sequential log files named

NtFrs_000*n*.log. The files are numbered, beginning with 0001, so that the log with the highest number contains the most recently logged entries. The log files are retained until the number of files reaches the value of this entry, and then FRS deletes NtFrs_0001.log and renumbers all remaining logs accordingly.

To find more information aboutthis registry entry, see Registry Reference in Tools and Settings Collection.

**Debug Log Severity**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Specifies the level of detail that FRS records in its debug log files (NtFrs_000*n*.log). FRS assigns severity levels to different log print statements in the FRS code. The value of this entry determines which severity levels are written to the log. The higher the value of this entry, the more detail the log records.

If the value of this entry is 0, only the most severe events are recorded in the log. If the value of this entry is 5, all events are recorded in the log. The default value is set to 2 so that tracking records are displayed, as well as warning and error messages.

To find more information aboutthis registry entry, see Registry Reference in Tools and Settings Collection.

**Debug Maximum Log Messages**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Specifies how many entries can be stored in each FRS debug log file (NtFrs_000*n*.log). When the number of event entries reaches the value of **Debug Maximum Log Messages**, the system closes the current log file and starts a new one. The log files are numbered sequentially, so when FRS closes NtFrs_0001.log, it begins NtFrs_0002.log.

To find more information aboutthis registry entry, see Registry Reference in Tools and Settings Collection.

**DS Polling Short Interval in Minutes**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

The interval at which FRS polls Active Directory at service startup or after configuration changes. FRS performs eight short polling intervals and then switches to long polling intervals if it is not interrupted by configuration changes that cause the short polling sequence to start over.

**DS Polling Long Interval in Minutes**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

The interval at which FRS polls Active Directory for configuration changes after eight short polling intervals have finished without interruption. FRS performs eight short polling intervals and then switches to long polling intervals if it is not interrupted by configuration changes. Configuration changes cause the short polling sequence to begin again.

**Enable Sharing Violation Logging**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003

Enables FRS to log sharing violation events (event ID 13573) in the File Replication service event log. When this is enabled, FRS keeps track of change orders that encounter sharing violations during install and other scenarios, and then logs event ID 13573 if the same change order encounters sharing violations equal to the **Sharing Violation Events Counter** registry entry per hour.

**Max Sharing Violation Event**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003

Specifies the maximum number of sharing violation events that are posted to the File Replication service event log in one hour. The default value is 10, which means that FRS logs ten sharing violation events per hour at most. Note that if multiple change orders encounter sharing violations for the same file, FRS posts a single event to the log.

**Ntfs Journal Size in MB**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Controls the size of the update sequence number (USN) journal. In Windows Server 2003, the default journal size is 128 megabytes (MB); in the pre-SP1 release of Ntfrs.exe, the default USN journal size is 512 MB. This setting applies to all volumes hosting an FRS replica tree. When the USN journal is set to 128 MB on a server that contains more than 100,000 files (or 400,000 files if the server is running the pre-SP1 release of Ntfrs.exe), we recommend that you increase the USN journal size by 128 MB per 100,000 additional files managed by FRS on the server.

To increase the size of the USN journal, you must stop and restart the NTFRS service. To decrease the size of the USN journal, you must reformat all volumes containing FRS replicated content.

**Outlog Change History In Minutes**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Controls outbound log change-order retention. FRS retains change orders even after they are sent to all current downstream partners. This feature permits FRS to vvjoin with a new downstream partner from the outbound log and avoids a full file ID table scan. As a result, this feature is helpful in environments where the topology changes frequently and during a rollout, when many new members come online within a short time.

**RPC TCP/IP Port Assignment**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Specifies the port that FRS uses for replication traffic. By default, FRS replication over remote procedure calls (RPCs) occurs dynamically over an available port by using RPC Endpoint Mapper (also known as RPCSS) on port 135. You can override this default functionality and specify the port that all FRS replication traffic passes through. When you do so, you can limit replication to a static port. The preferred port range to select from is 49152 to 65535, which is the range for private and dynamic ports that is specified by Internet Assigned Numbers Authority (IANA).

**Sharing Violation Events Counter**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003

Specifies how often FRS logs sharing violation events in the File Replication service event log. The default is 10, which means that FRS logs a sharing violation event for every ten sharing violations that a change order encountered.

**Staging Space Limit in KB**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Specifies the maximum amount of disk space used for the staging folder. The size of the staging folder governs the maximum amount of disk space that FRS can use to hold those staging files and the maximum file size that FRS can replicate. When FRS tries to allocate space for a staging file and is not successful because the amount of space in use has reached 90 percent of the staging space limit parameter, FRS starts to delete staging files. Staged files are deleted (in the order of the longest time since the last access) until the amount of space in use has dropped below 60 percent of the staging space limit parameter.

To find more information aboutthis registry entry, see Registry Reference in Tools and Settings Collection.

**Suppress Identical Updates to Files**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Specifies whether FRS tries to identify and suppress updates that do not change the content of a file (everything that is used to calculate the MD5 and attributes).

**Working Directory**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters

**Version**
Windows Server 2003 and Windows 2000

Specifies the location of the Ntfrs.jdb file and associated log files. The Jet database should reside on an NTFS volume.

**Parameters\Access Checks\Get Ds Polling Interval**
The following registry entries are located under
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Get Ds Polling Interval.

**Access checks are [Enabled or Disabled]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Get Ds Polling Interval

**Version**
Windows Server 2003 and Windows 2000

Specifies whether a user or process can get information about the Active Directory polling interval (by using Ntfrsutl.exe with the **poll** parameter, for example). Set to Enabled by default.

**Access checks require [Full Control or Read]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Get Ds Polling Interval

**Version**
Windows Server 2003 and Windows 2000

Specifies the permissions with which a user or process can get information about the Active Directory polling interval. Set to Read by default.

**Parameters\Access Checks\Get Internal Information**
The following registry entries are located under
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Get Internal Information.

**Access checks are [Enabled or Disabled]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Get Internal Information

**Version**
Windows Server 2003 and Windows 2000

Specifies whether a user or process can get internal information about FRS (by using Ntfrsutl.exe, for example). Set to Enabled by default.

**Access checks require [Full Control or Read]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Get Internal Information

**Version**
Windows Server 2003 and Windows 2000

Specifies the permissions with which a user or process can get internal information about FRS (by using Ntfrsutl.exe, for example). Set to Full Control by default.

**Parameters\Access Checks\Get Perfmon Data**
The following registry entries are located under
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Get Perfmon Data.

**Access checks are [Enabled or Disabled]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Get Perfmon Data

**Version**
Windows Server 2003 and Windows 2000

Specifies whether a user or process can collect performance-related information about FRS. Set to Enabled by default.

**Access checks require [Full Control or Read]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Get Perfmon Data

**Version**
Windows Server 2003 and Windows 2000

Specifies the permissions with which a user or process can collect performance-related information about FRS. Set to Read by default.

**Parameters\Access Checks\Set Ds Polling Interval**
The following registry entries are located under
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Set Ds Polling Interval.

**Access checks are [Enabled or Disabled]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Set Ds Polling Interval

**Version**
Windows Server 2003 and Windows 2000

Specifies whether a user or process can change the Active Directory polling interval. Set to Enabled by default.

**Access checks require [Full Control or Read]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Set Ds Polling Interval

**Version**
Windows Server 2003 and Windows 2000

Specifies the permissions with which a user or process can change the Active Directory polling interval. Set to Full Control by default.

**Parameters\Access Checks\Start Ds Polling**
The following registry entries are located under
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Start Ds Polling.

**Access checks are [Enabled or Disabled]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Start Ds Polling

**Version**
Windows Server 2003 and Windows 2000

Specifies whether a user or process can trigger Active Directory polling (by using Ntfrsutl.exe with the **poll /now** parameter, for example). Set to Enabled by default.

**Access checks require [Full Control or Read]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Start Ds Polling

**Version**
Windows Server 2003 and Windows 2000

Specifies the permissions with which a user or process can trigger Active Directory polling (by using Ntfrsutl.exe with the **poll /now** parameter, for example). Set to Read by default.

**Parameters\Access Checks\Writer Commands**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Writer Commands.

**Access checks are [Enabled or Disabled]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Writer Commands

**Version**
Windows Server 2003 and Windows 2000

Specifies whether the FRS VSS writer can issue commands. Set to Enabled by default.

**Access checks require [Full Control or Read]**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Access Checks\Writer Commands

**Version**
Windows Server 2003 and Windows 2000

Specifies the permissions with which the FRS VSS writer can issue commands. Set to Full Control by default.

**Parameters\Backup/Restore\Process at Startup**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\ Backup/Restore\Process at Startup.

**Burflags**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\Backup/Restore\Process at Startup

**Version**
Windows Server 2003 and Windows 2000

Short for “backup restore flags,” Burflags is used to perform nonauthoritative or authoritative restores on FRS members of DFS or SYSVOL replica sets.

**Parameters\SysVol Seeding**
**Replica Set Parent**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\Parameters\SysVol Seeding

**Version**
Windows Server 2003 and Windows 2000

Specifies the RPC binding string for the parent computer from which to seed during the SYSVOL creation process. This registry entry applies only to SYSVOL replica sets.

**Network Ports Used by FRS**
FRS uses the following two network ports:

### Network Ports Used by FRS

| Service Name | TCP |
|---|---|
| LDAP | 389 |
| RPC | Dynamic |

By default, FRS replication over remote procedure calls (RPCs) occurs dynamically over an available port by using RPC Endpoint Mapper (also known as Remote Procedure Call Server Service or RPCSS) on port 135; the process is the same for Active Directory replication.

# File Systems Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File Systems Technologies**

**In this overview**


- File Systems Components
- File Systems Deployment Scenarios


A file system is a part of the operating system on a volume and it determines how files are named, stored, and organized on basic or dynamic disks. A file system manages files and folders as well as the information required to locate and access these items by local and remote users. The Microsoft Windows Server 2003 operating system supports the NTFS and FAT file systems.

When you format a volume on a hard disk, a file system is always required to access and manage data. You can put only one file system on a volume. If you partition your hard disk into multiple volumes, each volume can have a different file system.

All file systems used by Windows Server 2003 organize hard disks based on cluster size, which is determined by the number of sectors contained by the cluster. When a file is written to disk, the file system tries to find enough contiguous clusters to hold the file. If the file system cannot find contiguous clusters, it saves the file to noncontiguous clusters. The result is file fragmentation. Disk Defragmenter is a file system tool that organizes the clusters on the volume to be more contiguous, which, in turn, improves file system performance.

**File Systems Components**

NTFS is the preferred file system for Windows Server 2003. NTFS addresses many of the limitations of FAT file system-based file systems. It provides improved file security, as well as performance, reliability, and advanced features not found in any version of FAT. Use NTFS wherever possible to gain the maximum benefits from Windows Server 2003.

Although NTFS is the preferred file system for hard disks, Windows Server 2003 uses FAT12 for formatting floppy disks and FAT32 for formatting flash media and DVD-RAM discs.

**FAT**

The two versions of FAT supported by Windows Server 2003 are FAT16 and FAT32. FAT16 is an older file system. It was widely used on Windows NT 4.0, and supports a 16-bit file allocation table. FAT16 provides optimal performance on volumes that are less than 2 gigabytes (GB). FAT32 is the version of FAT introduced with Windows 95 OEM Service Release 2 (OSR2) and Windows 98. FAT32 supports a 32-bit allocation table, and supports smaller cluster sizes than FAT16, in order to more efficiently allocate space. On Windows Server 2003, FAT32 can support volumes up to 32 GB. FAT is also the only file system supported for floppy disks and other removable media.

Some earlier Microsoft Windows operating systems as well as removable disks and floppy disks support only the FAT file system. FAT16 and FAT32 volumes are limited to 4 gigabytes (GB) and 32 GB respectively, so you must use NTFS to create volumes larger than 32 GB.

**Disk Defragmenter**

Disk Defragmenter is a tool that is used to improve the overall performance of disk input or output on either NTFS or FAT-formatted volumes. Disk Defragmenter analyzes local volumes and consolidates fragmented files and folders so that each occupies a single, contiguous space on the volume. As a result, a system can access files and folders and save new ones more efficiently. By consolidating files and folders, Disk Defragmenter also consolidates a volume's free space, making it less likely that new files will be fragmented.

**File Systems Scenarios**

This section describes a few scenarios in which NTFS, FAT, and Disk Defragmenter are used. Before formatting a volume for a file system, you should determine what the volume will be used for, and which operating system will be mounted on that volume. The following scenarios might help you when you are deciding how to format volumes.

**NTFS Scenarios**

This section describes scenarios in which NTFS should be the file system on a volume.

**Increasing reliability**

When a computer is restarted after a system failure, NTFS uses its log file and checkpoint information to restore the file system. In the event of a bad-sector error, NTFS dynamically remaps the cluster containing the bad sector and allocates a new cluster for the data. It also marks the cluster as bad so that it is no longer used. By formatting a POP3 mail server with NTFS, the mail store can offer logging and recovery. In the event of a server failure, NTFS can recover data by using its log files.

**Increasing security**
NTFS allows you to set permissions on a file or folder. You can also specify the groups and users who can access the file or folder, and set specific levels of access (for example, some users can only open the file, while other users can open it and modify it). NTFS also supports Encrypting File System (EFS) technology. Unauthorized users are prevented from accessing encrypted files, even if the unauthorized user has physical access to the computer. If a POP3 mail server is formatted with NTFS, NTFS provides increased security for the mail store (security that is not available with the FAT file system).

**Supporting large volumes**
Clusters are the smallest amount of disk space used to store a file. When you format a volume as NTFS, you can choose the cluster size of the volume. The smaller the cluster size, the more efficiently a disk stores information. You can create an NTFS volume which is up to 16 terabytes in size using the default cluster size (4 kilobytes (KB)) for large volumes, or you can create NTFS volumes of up to 256 terabytes using the maximum cluster size of 64 KB. NTFS also supports larger files and more files per volume than FAT. In addition, NTFS manages disk space more efficiently than FAT by using smaller cluster sizes. Using smaller clusters reduces wasted space on hard disks. NTFS supports the many capabilities of dynamic disks for managing large storage requirements.

**Limited space on a volume**
If your organization has limited space on a volume, NTFS provides support for increasing storage on a server with limited disk space.

- Disk quotas allow you to track and control user disk space usage for NTFS volumes.
- NTFS supports compression as well as adding unallocated space from the same disk or from another disk to increase the size of an NTFS volume.
- Mounted volumes are volumes attached to an empty folder on an NTFS volume. Mounted drives function the same way as any other volume, but are assigned a label or name instead of a drive letter. Mounted drives more easily withstand system changes that occur when devices are added or removed from a computer. They are not subject to the 26-volume limit imposed by drive letters, so you can use them for access to more than 26 volumes on your computer.

**Using features available only in NTFS**
NTFS has a number of features that are not available if you are using a FAT file system. These include:

- **Distributed link tracking.** Distributed link tracking ensures that the integrity of shortcuts and OLE links. You will not break shortcut or OLE links, even if you rename source files, move them to NTFS volumes on different computers within a Windows Server 2003 or Windows 2000 domain, or change the computer name or folder name that stores the target.
- **Sparse files.** Sparse files provide a method for reserving disk space for files that contain meaningful data as well as large sections of data composed of zeros. If an NTFS file is marked as sparse, then NTFS allocates disk clusters only for the data explicitly specified by the application. The result is efficient file system storage and access.
- **NTFS change journal.** A log of changes that are made to files on a volume. NTFS maintains the log on each volume by tracking information about files that are added, deleted, or modified.
- **Hard links.** Hard links are NTFS-based pointers that allow multiple paths to refer to the same file or directory on NTFS volumes. By creating hard links, a single file can appear to exist in multiple folders (but you do not have to take up the space necessary to duplicate the file). By creating hard links, you can:
    - Make a single file appear to reside in different folders.
    - Use multiple names to refer to a single file that resides in a single folder.
    - Use multiple names to refer to files that appear to reside different folders (in reality the links all point to a single file that resides in a single folder).

**Using Windows Server 2003 features that require NTFS**
Windows Server 2003 includes a number of features that require NTFS. A few of these features include:

- **Volume Shadow Copy Service.** Service that provides an infrastructure for creating highly accurate, point-in-time shadow copies. These copies of a single volume or multiple volumes can be made without affecting the performance of a production server. The Volume Shadow Copy service can produce accurate shadow copies by coordinating with business applications, backup applications, and storage hardware.

- **Distributed File System (DFS).** Strategic storage management solution in Windows Server 2003 that allows you to logically group shared files and folders on different servers. With DFS, users no longer need to know the actual physical location of files in order to access them. For example, if you have marketing files on multiple servers in a domain, you can use DFS to make it appear as though all of the marketing files are on a single server.

- **File System Replication (FRS).** Technology that replicates files and folders stored in the SYSVOL shared folder on domain controllers and Distributed File System (DFS) shared folders. When FRS detects that a change has been made to a file or folder within a replicated shared folder, FRS replicates the updated file or folder to other servers.

In addition, NTFS is required before you can promote a server running Windows Server 2003 to a domain controller that hosts the Active Directory directory service.

**FAT Scenarios**
This section describes scenarios in which FAT should be the file system on a volume.

**Mounting a server with a small volume**
Small partitions (less than 512 megabytes) on a server work better with FAT. It uses a small amount of space on the partition, but still performs well.

**Running multiple operating systems**
If you are running multiple operating systems (for example, you are running Windows 98 in addition to Windows Server 2003), you must use FAT16 or FAT32 as Windows 98 is only supported by the FAT file system.

**Using FAT with older operating systems**
Both Windows 95 and Windows 98 use the FAT file system. These operating systems do not support NTFS.

**Using FAT with removable disks**
You must use FAT16 or FAT32 when formatting removable disks (for example, floppy disks). It is not possible to format a floppy disk with the NTFS file system because the file system itself requires too much disk space.

**Disk Defragmenter Scenarios**
This section describes scenarios in which Disk Defragmenter should be run on a volume.

**Installing new software**
Running Disk Defragmenter on a volume after installing software organizes the clusters on the volume.

**Rewriting, updating and deleting files**
Running Disk Defragmenter on a regular basis consolidates space to create more contiguous free clusters on the volume, which allows disk heads to read files and write updates more quickly.

**Operating System Compatibility with NTFS and FAT**
Not all Microsoft operating systems support NTFS or all versions of FAT. The following table lists common operating systems and indicates which file system they support.

### Operating System Compatibility with NTFS and FAT

| Operating System | FAT16 | FAT32 | NTFS |
|---|---|---|---|
| Windows XP | ● | ● | ● |
| Windows Server 2003 | ● | ● | ● |
| Windows 2000 | ● | ● | ● |
| Windows NT 4.0 | ● | | ● |

| | | | |
|---|---|---|---|
| Windows 95 OSR2, Windows 98, and Windows Millennium Edition | • | • | |
| Windows 95 (prior to OSR2) | • | | |
| MS-DOS | • | | |

# NTFS Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NTFS Technical Reference**

A file system is a part of the operating system that determines how files are named, stored, and organized on a volume. A file system manages files and folders, and the information needed to locate and access these items by local and remote users. The Microsoft Windows Server 2003 operating system supports both the FAT and NTFS file systems.

NTFS provides performance, reliability, and advanced features not found in FAT. For example, NTFS includes built-in security features such as file and folder permissions and the Encrypting File System (EFS), which is the technology used to store encrypted files on NTFS volumes. NTFS also provides support for volumes up to 256 terabytes in size, support for disk quotas and compression, and support for mounted drives.

**In this subject**


- What Is NTFS?
- How NTFS Works
- NTFS Tools and Settings

# What Is NTFS?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is NTFS?**

**In this section**


- Common NTFS Scenarios
- Dependencies with Other Technologies
- Related Technology
- Related Information


A file system is a part of the operating system that determines how files are named, stored, and organized on a volume. A file system manages files and folders, and the information needed to locate and access these items by local and remote users. Microsoft Windows Server 2003 supports both the FAT and NTFS file systems.

NTFS allows you to gain the maximum benefits for the needs of today's enterprise business environments from Windows Server 2003, such as increased security, more robust and reliable performance, as well as a design for greater storage growth, features not found in FAT.

**Common NTFS Scenarios**

This section describes a few scenarios in which NTFS should be used as the file system on a server running Windows Server 2003.

**Increasing reliability**

NTFS uses its log file and checkpoint information to restore the consistency of the file system when the computer is restarted in the event of a system failure. In the event of a bad-sector error, NTFS dynamically remaps the cluster containing the bad sector and allocates a new cluster for the data, as well as marking the cluster as bad and no longer using it. For example, by formatting a POP3 mail server with NTFS, the mail store can offer logging and recovery. In the event of a server crash, NTFS can recover data by replaying its log files.

**Increasing security**

NTFS allows you to set permissions on a file or folder, and specify the groups and users whose access you want to restrict or allow, and then select the type of access. NTFS also supports the Encrypting File System (EFS) technology used to store encrypted files on NTFS volumes. Any intruder who tries to access your encrypted files is prevented from doing so, even if that intruder has physical access to the computer. For example, a POP3 mail server, when formatted with an NTFS file system, provides increased security for the mail store, security that would not be available should the server be formatted with the FAT file system.

**Supporting large volumes**

NTFS allows you to create an NTFS volume up to 16 terabytes using the default cluster size (4 KB) for large volumes. You can create NTFS volumes up to 256 terabytes using the maximum cluster size of 64 KB. NTFS also supports larger files and more files per volume than FAT. In addition, NTFS manages disk space more efficiently than FAT by using smaller cluster sizes. For example, a 30-GB NTFS volume uses 4-KB clusters. The same volume formatted by using FAT32 uses 16-KB clusters. Using smaller clusters reduces wasted space on hard disks. NTFS supports the many capabilities of dynamic disks for managing large storage requirements.

**Limited space on a volume**

If your organization has limited space on a volume, NTFS provides support for increasing storage on a server with limited disk space.

- Disk quotas allow you to track and control user disk space usage for NTFS volumes.
- NTFS supports compression as well as adding unallocated space from the same disk or from another disk to increase the size of an NTFS volume.

- Mounted volumes allow you to mount a volume at any empty folder on a local NTFS volume if you run out of drive letters or need to create additional space that is accessible from an existing folder.

#### Using features available only in NTFS

NTFS has a number of features that are not available if you are using a FAT file system. These include:

- **Distributed link tracking.** Maintains the integrity of shortcuts and OLE links. You can rename source files, move them to NTFS volumes on different computers within a Windows Server 2003 or Windows 2000 domain, or change the computer name or folder name that stores the target without breaking the shortcut or OLE links.

- **Sparse files.** Large, consecutive areas of zeros. NTFS manages sparse files by tracking the starting and ending point of the sparse file, as well as its useful (non-zero) data. The unused space in a sparse file is made available as free space.

- **NTFS change journal**. Provides a persistent log of changes made to files on a volume. NTFS maintains the change journal by tracking information about added, deleted, and modified files for each volume.

- **Hard links.** NTFS-based links to a file on an NTFS volume. By creating hard links, you can have a single file in multiple folders without duplicating the file. You can also create multiple hard links for a file in a folder if you use different file names for the hard links. Because all of the hard links reference the same file, applications can open any of the hard links and modify the file.

#### Using Windows Server 2003 features that require NTFS

Windows Server 2003 includes a number of features that require NTFS as the file system. A few of these features include:

- **Volume Shadow Copy service.** Service that provides an infrastructure for creating highly accurate, point-in-time shadow copies. These copies of a single volume or multiple volumes can be made without affecting the performance of a production server. The Volume Shadow Copy Service can produce accurate shadow copies by coordinating with business applications, backup applications, and storage hardware.

- **Distributed File System (DFS)**. Strategic storage management solution in Windows Server 2003 that enables you to group shared folders located on different servers logically by transparently connecting them to one or more hierarchical namespaces.

- **File System Replication (FRS)** Technology that replicates files and folders stored in the SYSVOL shared folder on domain controllers and Distributed File System (DFS) shared folders. When FRS detects that a change has been made to a file or folder within a replicated shared folder, FRS replicates the updated file or folder to other servers.

In addition, NTFS is required before you can promote a server running Windows Server 2003 to a domain controller that hosts the Active Directory directory service.

If the volume is not formatted with the NTFS file system, these Windows Server 2003 features will not be available.

### Note

- Although NTFS is the preferred file system for hard disks, NTFS cannot be used on removable media. Instead, Windows Server 2003 uses FAT12 for formatting floppy disks, and FAT32 for formatting flash media and DVD-RAM discs.

#### Operating System and NTFS Compatibility

NTFS is not supported on versions of Microsoft Windows earlier than Windows NT 4.0 and Windows 2000 Professional or MS-DOS. The table Operating System and NTFS Compatibility shows which operating systems support NTFS.

### Operating System and NTFS Compatibility

| Operating System | NTFS |
|---|---|
| Windows XP | • |
| Windows Server 2003 | • |

| | |
|---|---|
| Windows 2000 | • |
| Windows NT 4.0 | • |
| Windows 95 OSR2, Windows 98, and Windows Millennium Edition | |
| Windows 95 (prior to OSR2) | |
| MS-DOS | |

**Note**

- Computers running Windows NT 4.0 require Service Pack 4 or later to access NTFS volumes previously mounted by Windows 2000, Windows XP, or Windows Server 2003.

**Dependencies on Other Technologies**
NTFS depends on the following technologies:

**Basic Disks and Volumes**
Basic disks and basic volumes are the storage types most often used with Windows operating systems. The term basic disk refers to a disk that contains basic volumes, such as primary partitions and logical drives. The term basic volume refers to a partition on a basic disk. Basic disks, which are found in both x86-based and Itanium-based computers, provide a simple storage solution that can accommodate changing storage requirements.

**Dynamic Disks and Volumes**
Dynamic disks can use the master boot record (MBR) or GUID partition table (GPT) partitioning scheme. All volumes on dynamic disks are known as dynamic volumes. Dynamic disks were first introduced with Windows 2000 and provide features that basic disks do not.

**Related Technology**
NTFS is related to the following technology:

**FAT file system**
The File Allocation Table (FAT) file system is an older file system that relies on an allocation table to keep track of files and folders on a volume. Windows Server 2003 supports both FAT16 and FAT32 file systems.

**Related Information**

- Basic Disks and Volumes Technical Reference
- Disk Quotas Technical Reference
- Dynamic Disks and Volumes Technical Reference
- FRS Technical Reference

# How NTFS Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How NTFS Works**

**In this section**


- NTFS Architecture
- NTFS Physical Structure
- NTFS Processes and Interactions
- Related Information


A file system is a required part of the operating system that determines how files are named, stored, and organized on a volume. A file system manages files and folders, and the information needed to locate and access these items by local and remote users.

Microsoft Windows Server 2003 supports the NTFS file system on basic and dynamic disks. Basic disks and volumes are the storage types most often used with Windows operating systems. Dynamic disks offer greater flexibility for volume management because they use a database to track information about dynamic volumes on the disk and about other dynamic disks in the computer.

During the format of a volume you can choose the type of file system for the volume. When you choose the NTFS file system, the formatting process places the key NTFS file data structures on the volume, regardless of whether it is a basic or dynamic volume.

**NTFS Architecture**

During format and setup of a volume file system on a hard disk, a master boot record (MBR) is created. The MBR contains a small amount of executable code called the master boot code as well as a partition table for the disk. When a volume is mounted, the MBR executes the master boot code and transfers control to the boot sector on the disk, allowing the server to boot the operating system on the file system of that specific volume.

**Note**

- The partition table contains a number of fields used to describe the partition. One of these fields is the System ID field, which defines the file system, such as NTFS, on the partition. For NTFS volumes, the system ID is 0x07.

The figure NTFS Architecture shows the architecture of this process.

**NTFS Architecture**




The following table describes the components of an NTFS file system.

**NTFS Architecture Components on an x86-based System**

| Component | Component Description |
|---|---|
| Hard disk | Contains one or more partitions. |
| Boot sector | Bootable partition that stores information about the layout of the volume and the file system structures, as well as the boot code that loads Ntdlr. |
| Master Boot Record | Contains executable code that the system BIOS loads into memory. The code scans the MBR to find the partition table to determine which partition is the active, or bootable, partition. |
| Ntldlr.dll | Switches the CPU to protected mode, starts the file system, and then reads the contents of the Boot.ini file. This information determines the startup options and initial boot menu selections. |
| Ntfs.sys | System file driver for NTFS. |
| Ntoskrnl.exe | Extracts information about which system device drivers to load and the load order. |
| Kernel mode | The processing mode that allows code to have direct access to all hardware and memory in the system. |
| User mode | The processing mode in which applications run. |

**NTFS Physical Structure**

The following information describes how clusters and sectors are organized on an NTFS volume, how the boot sector on the volume determines the file system, and how the Master File Table (MFT) organizes structures on the volume.

**Clusters and Sectors on an NTFS Volume**

A cluster (or allocation unit) is the smallest amount of disk space that can be allocated to hold a file. All file systems used by Windows Server 2003 organize hard disks based on cluster size, which is determined by the number of sectors (units of storage on a hard disk) that the cluster contains. For example, on a disk that uses 512-byte sectors, a 512-byte cluster contains one sector, whereas a 4-kilobyte (KB) cluster contains eight sectors.

Computers access certain sectors on a hard disk during startup to determine which operating system to start and where the partitions are located. The data stored on these sectors varies depending on the computer platform.

**Sequence of Clusters on an NTFS Volume**

Clusters on an NTFS volume are numbered sequentially from the beginning of the partition into logical cluster numbers. NTFS stores all objects in the file system using a record called the Master File Table (MFT), similar in structure to a database.

On NTFS volumes, clusters start at sector zero; therefore, every cluster is aligned on the cluster boundary. Contiguous clusters for file storage allow for faster processing of a file.

**Note**

- Floppy disks do not use NTFS and are always formatted as FAT.

**Limitations of Cluster Sizes on an NTFS Volume**

Because NTFS uses different cluster sizes depending on the size of the volume, each file system has a maximum number of clusters it can support. The smaller the cluster size, the more efficiently a disk potentially stores information because unused space within a cluster cannot be used by other files. And the more clusters a file system supports, the larger the volumes you can create and format by using a particular file system. NTFS uses smaller cluster sizes, which makes it a more efficient file organization structure.

The table Default NTFS Cluster Sizes lists NTFS volume and default cluster sizes.

**Default NTFS Cluster Sizes**

| Volume Size | NTFS Cluster Size |
|---|---|
| 7 megabytes (MB)–512 MB | 512 bytes |
| 513 MB–1,024 MB | 1 KB |
| 1,025 MB–2 GB | 2 KB |
| 2 GB–2 terabytes | 4 KB |

**Maximum Sizes on an NTFS Volume**

Before you format an NTFS volume, evaluate the types of files to be stored on the volume so that you can determine whether to use the default cluster size.

When formatting NTFS volumes, you can specify a cluster size of up to 64 KB using the Disk Management snap-in. If you format a volume, but do not specify a cluster size, default values are used. If you want to change the cluster size after the volume is formatted, you must reformat the volume.

Before you choose a cluster size other than the default, note the following important limitations:

- For Microsoft Windows NT, Windows 2000, Windows XP, and Windows Server 2003, the cluster size of FAT16 volumes ranging from 2 gigabytes (GB) through 4 GB is 64 KB, which can create compatibility issues with some applications. For example, setup programs do not compute free space properly on a volume with 64-KB clusters and cannot run because of a perceived lack of free space. For this reason, you can use either NTFS or FAT32 to format volumes larger than 2 GB.

- Because file compression is not supported on cluster sizes greater than 4 KB, the default NTFS cluster size for Windows Server 2003 never exceeds 4 KB.

In theory, the maximum NTFS volume size is 264 clusters minus 1 cluster. However, the maximum NTFS volume size as implemented in Windows Server 2003 is $2^{32}$ clusters minus 1 cluster. For example, using 64-KB clusters, the maximum NTFS volume size is 256 terabytes minus 64 KB. Using the default cluster size of 4 KB, the maximum NTFS volume size is 16 terabytes minus 4 KB.

**Note**

- If you use large numbers of files in an NTFS folder (300,000 or more), disable short-file name generation for better performance, and especially if the first six characters of the long file names are similar.

The table NTFS Size Limits lists NTFS size limits.

**NTFS Size Limits**

| Description | Limit |
|---|---|
| Maximum file size | Architecturally: 16 exabytes minus 1 KB ($2^{64}$ bytes minus 1 KB)<br><br>Implementation: 16 terabytes minus 64 KB ($2^{44}$ bytes minus 64 KB) |

| | |
|---|---|
| Maximum volume size | Architecturally: $2^{64}$ clusters minus 1 cluster<br><br>Implementation: 256 terabytes minus 64 KB ( $2^{32}$ clusters minus 1 cluster) |
| Files per volume | 4,294,967,295 ($2^{32}$ minus 1 file) |

**Partition Tables on MBR and GUID disks**

Master boot record (MBR) disks use both basic and dynamic volumes. Because partition tables on MBR disks support partition sizes only up to 2 terabytes, you must use dynamic volumes to create NTFS volumes over 2 terabytes. Windows Server 2003 manages dynamic volumes in a special database instead of in the partition table; therefore dynamic volumes are not subject to the 2-terabyte physical limit imposed by the partition table. Dynamic NTFS volumes can be as large as the maximum volume size supported by NTFS. Itanium-based computers that use GUID partition table (GPT) disks also support NTFS volumes larger than 2 terabytes.

**Organization of an NTFS Volume**

The figure Organization of an NTFS Volume illustrates how NTFS organizes structures on a volume.

### Organization of an NTFS Volume




The following table describes each of the organizational structures on the NTFS volume.

### NTFS Volume Components

| Component | Description |
|---|---|
| NTFS Boot Sector | Contains the BIOS parameter block that stores information about the layout of the volume and the file system structures, as well as the boot code that loads Windows Server 2003. |
| Master File Table | Contains the information necessary to retrieve files from the NTFS partition, such as the attributes of a file. |
| File System Data | Stores data that is not contained within the Master File Table. |
| Master File Table Copy | Includes copies of the records essential for the recovery of the file system if there is a problem with the original copy. |

**Boot Sectors**

On MBR disks, the boot sector, which is located at the first logical sector of each partition, is a critical disk structure for starting your computer. It contains executable code and the data required by the code, including information that the file system uses to access the volume. The boot sector is created when you format a volume. At the end of the boot sector is a 2-byte structure called a signature word or end of sector marker, which is always set to 0x55AA. On computers running Windows Server 2003, the boot sector on the active partition loads into memory and starts Ntldr, which loads the boot menu if multiple versions of Windows are installed, or loads the operating system if only one operating system is installed.

GUID partition table (GPT) disks are similar to MBR disks, except they use primary and backup partition structures to provide redundancy. These structures are located at the beginning and the end of the disk. GPT identifies these structures by their logical block address (LBA) rather than by their relative sectors.

A boot sector consists of the following elements:

- An x86-based CPU jump instruction.
- The original equipment manufacturer identification (OEM ID).
- The BIOS parameter block (BPB), a data structure.
- The extended BPB.
- The executable boot code (or bootstrap code) that starts the operating system.

All Windows Server 2003 boot sectors contain the preceding elements regardless of the type of disk (basic disk or dynamic disk).

**Components of a Boot Sector**

The MBR transfers CPU execution to the boot sector, so the first three bytes of the boot sector must be valid, executable x86-based CPU instructions. This includes a jump instruction that skips the next several nonexecutable bytes.

Following the jump instruction is the 8-byte OEM ID, a string of characters that identifies the name and version number of the operating system that formatted the volume. To preserve compatibility with MS-DOS, Windows Server 2003 records "NTFS" in this field.

### Note

- You might also see the OEM ID "MSWIN4.0" on disks formatted by Windows 95 and "MSWIN4.1" on disks formatted by Windows 95 OEM Service Release 2 (OSR2), Windows 98, and Windows Millennium Edition. Windows Server 2003 does not use the OEM ID field in the boot sector except for verifying NTFS volumes.

Following the OEM ID is the BPB, which provides information that enables the executable boot code to locate Ntldr. The BPB always starts at the same offset, so standard parameters are in a known location. Disk size and geometry variables are encapsulated in the BPB. Because the first part of the boot sector is an x86 jump instruction, the BPB can be extended in the future by appending new information at the end. The jump instruction needs only a minor adjustment to accommodate this change. The BPB is stored in a packed (unaligned) format.

**NTFS Boot Sector**

The table Boot Sector Sections on an NTFS Volume describes the boot sector of a volume that is formatted with NTFS. When you format an NTFS volume, the format program allocates the first 16 sectors for the boot sector and the bootstrap code.

### Boot Sector Sections on an NTFS Volume

| Byte Offset | Field Length | Field Name |
|---|---|---|
| 0x00 | 3 bytes | Jump instruction |
| 0x03 | 8 bytes | OEM ID |
| 0x0B | 25 bytes | BPB |
| 0x24 | 48 bytes | Extended BPB |
| 0x54 | 426 bytes | Bootstrap code |
| 0x01FE | 2 bytes | End of sector marker |

On NTFS volumes, the data fields that follow the BPB form an extended BPB. The data in these fields enables Ntldr to find the MFT during startup. On NTFS volumes, the MFT is not located in a predefined sector. For this reason, NTFS can move the MFT if there is a bad sector in the current location of the MFT. However, if the data is corrupted, the MFT cannot be located, and Windows Server 2003 assumes that the volume has not been formatted.

The following example illustrates the boot sector of an NTFS volume that is formatted by using Windows Server 2003. The printout is formatted in three sections:

- Bytes 0x00– 0x0A are the jump instruction and the OEM ID (shown in bold print).
- Bytes 0x0B–0x53 are the BPB and the extended BPB.
- The remaining code is the bootstrap code and the end of sector marker (shown in bold print).

```
Physical Sector: Cyl 0, Side 1, Sector 1
00000000: EB 52 90 4E 54 46 53 20 - 20 20 20 00 02 08 00 00 .R.NTFS ..... ..
00000010: 00 00 00 00 00 F8 00 00 - 3F 00 FF 00 3F 00 00 00 ........?...?...
00000020: 00 00 00 00 80 00 80 00 - 1C 91 11 01 00 00 00 00 ................
00000030: 00 00 04 00 00 00 00 00 - 11 19 11 00 00 00 00 00 ................
00000040: F6 00 00 00 01 00 00 00 - 3A B2 7B 82 CD 7B 82 14 ........:.{..{..
00000050: 00 00 00 00 FA 33 C0 8E - D0 BC 00 7C FB B8 C0 07 .....3.....|....
```

The table BPB and Extended BPB Fields on NTFS Volumes describes the fields in the BPB and the extended BPB on NTFS volumes. The fields starting at 0x0B, 0x0D, 0x15, 0x18, 0x1A, and 0x1C match those on FAT16 and FAT32 volumes. The sample values correspond to the data in this example.

**BPB and Extended BPB Fields on NTFS Volumes**

| Byte Offset | Field Length | Sample Value | Field Name and Definition |
|---|---|---|---|
| 0x0B | 2 bytes | 00 02 | **Bytes Per Sector**. The size of a hardware sector. For most disks used in the United States, the value of this field is 512. |
| 0x0D | 1 byte | 08 | **Sectors Per Cluster**.The number of sectors in a cluster. |
| 0x0E | 2 bytes | 00 00 | **Reserved Sectors**. Always 0 because NTFS places the boot sector at the beginning of the partition. If the value is not 0, NTFS fails to mount the volume. |
| 0x10 | 3 bytes | 00 00 00 | Value must be 0 or NTFS fails to mount the volume. |
| 0x13 | 2 bytes | 00 00 | Value must be 0 or NTFS fails to mount the volume. |
| 0x15 | 1 byte | F8 | **Media Descriptor**. Provides information about the media being used. A value of F8 indicates a hard disk and F0 indicates a high-density 3.5-inch floppy disk. Media descriptor entries are a legacy of MS-DOS FAT16 disks and are not used in Windows Server 2003. |
| 0x16 | 2 bytes | 00 00 | Value must be 0 or NTFS fails to mount the volume. |
| 0x18 | 2 bytes | 3F 00 | Not used or checked by NTFS. |
| 0x1A | 2 bytes | FF 00 | Not used or checked by NTFS. |
| 0x1C | 4 bytes | 3F 00 00 00 | Not used or checked by NTFS. |
| 0x20 | 4 bytes | 00 00 00 00 | The value must be 0 or NTFS fails to mount the volume. |
| 0x24 | 4 bytes | 80 00 80 00 | Not used or checked by NTFS. |
| 0x28 | 8 bytes | 1C 91 11 01 00 00 00 00 | **Total Sectors**. The total number of sectors on the hard disk. |
| 0x30 | 8 bytes | 00 00 04 00 00 00 00 00 | **Logical Cluster Number for the File $MFT**. Identifies the location of the MFT by using its logical cluster number. |
| 0x38 | 8 bytes | 11 19 11 00 00 00 00 00 | **Logical Cluster Number for the File $MFTMirr**. Identifies the location of the mirrored copy of the MFT by using its logical cluster number. |
| 0x40 | 1 byte | F6 | **Clusters Per MFT Record**. The size of each record. NTFS creates a file record for each file and a folder record for each folder that is created on an NTFS volume. Files and folders smaller than this size are contained within the MFT. If this number is positive (up to 7F), then it represents clusters per MFT record. If the number is negative (80 to FF), then the size of the file record is 2 raised to the absolute value of this number. |

| | | | |
|---|---|---|---|
| 0x41 | 3 bytes | 00 00 00 | Not used by NTFS. |
| 0x44 | 1 byte | 01 | **Clusters Per Index Buffer**. The size of each index buffer, which is used to allocate space for directories. If this number is positive (up to 7F), then it represents clusters per MFT record. If the number is negative (80 to FF), then the size of the file record is 2 raised to the absolute value of this number. |
| 0x45 | 3 bytes | 00 00 00 | Not used by NTFS. |
| 0x48 | 8 bytes | 3A B2 7B 82 CD 7B 82 14 | **Volume Serial Number**. The volume's serial number. |
| 0x50 | 4 bytes | 00 00 00 00 | Not used by NTFS. |

**Master File Table**

When you format a volume with NTFS, Windows Server 2003 creates an MFT and metadata files on the partition. The MFT is a relational database that consists of rows of file records and columns of file attributes. It contains at least one entry for every file on an NTFS volume, including the MFT itself.

The MFT stores the information required to retrieve files from the NTFS partition.

**MFT and Metadata Files**

Because the MFT stores information about itself, NTFS reserves the first 16 records of the MFT for metadata files (approximately 16 KB), which are used to describe the MFT. Metadata files that begin with a dollar sign ($) are described in the table Metadata Files Stored in the MFT. The remaining records of the MFT contain the file and folder records for each file and folder on the volume.

**Metadata Files Stored in the MFT**

| System File | File Name | MFT Record | Purpose of the File |
|---|---|---|---|
| Master file table | $Mft | 0 | Contains one base file record for each file and folder on an NTFS volume. If the allocation information for a file or folder is too large to fit within a single record, other file records are allocated as well. |
| Master file table mirror | $MftMirr | 1 | Guarantees access to the MFT in case of a single-sector failure. It is a duplicate image of the first four records of the MFT. |
| Log file | $LogFile | 2 | Contains information used by NTFS for faster recoverability. The log file is used by Windows Server 2003 to restore metadata consistency to NTFS after a system failure. The size of the log file depends on the size of the volume, but you can increase the size of the log file by using the Chkdsk command. |
| Volume | $Volume | 3 | Contains information about the volume, such as the volume label and the volume version. |
| Attribute definitions | $AttrDef | 4 | Lists attribute names, numbers, and descriptions. |
| Root file name index | . | 5 | The root folder. |
| Cluster bitmap | $Bitmap | 6 | Represents the volume by showing free and unused clusters. |
| Boot sector | $Boot | 7 | Includes the BPB used to mount the volume and additional bootstrap loader code used if the volume is bootable. |

| | | | |
|---|---|---|---|
| Bad cluster file | $BadClus | 8 | Contains bad clusters for a volume. |
| Security file | $Secure | 9 | Contains unique security descriptors for all files within a volume. |
| Upcase table | $Upcase | 10 | Converts lowercase characters to matching Unicode uppercase characters. |
| NTFS extension file | $Extend | 11 | Used for various optional extensions such as quotas, reparse point data, and object identifiers. |
| | | 12–15 | Reserved for future use. |

The data segment locations for both the MFT and the backup MFT, $Mft and $MftMirr, respectively, are recorded in the boot sector. The $MftMirr is a duplicate image of either the first four records of the $Mft or the first cluster of the $Mft, whichever is larger. If any MFT records in the mirrored range are corrupted or unreadable, NTFS reads the boot sector to find the location of the $MftMirr. NTFS then reads the $MftMirr and uses the information in $MftMirr instead of the information in the MFT. If possible, the correct data from the $MftMirr is written back to the corresponding location in the $Mft.

**MFT Zone**

To prevent the MFT from becoming fragmented, NTFS reserves 12.5 percent of volume by default for exclusive use of the MFT. This space, known as the MFT zone, is not used to store data unless the remainder of the volume becomes full.

Depending on the average file size and other variables, as the volume fills to capacity, either the MFT zone or the unreserved space on the volume becomes full first.

- Volumes that have a small number of large files exhaust the unreserved space first.
- Volumes with a large number of small files exhaust the MFT zone space first.

In either case, fragmentation of the MFT occurs when one region or the other becomes full. You can change the size of the MFT zone for newly created volumes by to correspond to a percentage of the volume to be used as the MFT zone. The MFT zone sizes follow:

- Setting 1, the default, reserves approximately 12.5 percent of the volume.
- Setting 2 reserves approximately 25 percent.
- Setting 3 reserves approximately 37.5 percent.
- Setting 4 reserves approximately 50 percent.

In most computers, the default setting of 1 is adequate. The default setting accommodates volumes with an average file size of 8 KB. Storing a large number of smaller files might necessitate that you increase the size of the MFT zone for new volumes.

After you increase the size of the MFT zone, NTFS does not immediately allocate space to accommodate the size of the new MFT zone. Instead, NTFS exhausts the original reserved space before increasing the size of the MFT zone. When the original space is exhausted, NTFS looks for the next contiguous space large enough to hold the additional MFT zone, which can cause the MFT to become fragmented. You can adjust the zone size for the MFT if the defaults do not fit your needs.

**NTFS File Record Attributes**

Every allocated sector on an NTFS volume belongs to a file. Even the file system metadata is part of a file. NTFS views each file (or folder) as a set of file attributes. File elements such as its name, its security information, and even its data are file attributes. Each attribute is identified by an attribute type code and an optional attribute name.

File and folder records are 1 KB each and are stored in the MFT, the attributes of which are written to the allocated space in the MFT. Besides file attributes, each file record contains information about the position of the file record in the MFT.

When a file's attributes can fit within the MFT file record for that file, they are called resident attributes. Attributes such as file name and time stamp are always resident. When the amount of information for a file does not fit in its MFT file record, some

file attributes become nonresident. Nonresident attributes are allocated one or more clusters of disk space. A portion of the nonresident attribute remains in the MFT and points to the external clusters. NTFS creates the Attribute List attribute to describe the location of all attribute records. The table NTFS File Attribute Types lists the file attributes currently defined by NTFS.

**NTFS File Attribute Types**

| Attribute Type | Description |
|---|---|
| Standard Information | Information such as access mode (read-only, read/write, and so forth) timestamp, and link count. |
| Attribute List | Locations of all attribute records that do not fit in the MFT record. |
| File Name | A repeatable attribute for both long and short file names. The long name of the file can be up to 255 Unicode characters. The short name is the 8.3, case-insensitive name for the file. Additional names, or hard links, required by POSIX can be included as additional file name attributes. |
| Data | File data. NTFS supports multiple data attributes per file. Each file typically has one unnamed data attribute. A file can also have one or more named data attributes. |
| Object ID | A volume-unique file identifier. Used by the distributed link tracking service. Not all files have object identifiers. |
| Logged Tool Stream | Similar to a data stream, but operations are logged to the NTFS log file just like NTFS metadata changes. This attribute is used by EFS. |
| Reparse Point | Used for mounted drives. This is also used by Installable File System (IFS) filter drivers to mark certain files as special to that driver. |
| Index Root | Used to implement folders and other indexes. |
| Index Allocation | Used to implement the B-tree structure for large folders and other large indexes. |
| Bitmap | Used to implement the B-tree structure for large folders and other large indexes. |
| Volume Information | Used only in the $Volume system file. Contains the volume version. |

NTFS creates a file record for each file and a folder record for each folder created on an NTFS volume. The MFT includes a separate file record for the MFT itself. These file and folder records are 1 KB each and are stored in the MFT. The attributes of the file are written to the allocated space in the MFT. Besides file attributes, each file record contains information about the position of the file record in the MFT. The figure MFT Entry with Resident Record shows the contents of an MFT record for a small file or folder. Small files and folders (typically, 900 bytes or smaller) are entirely contained within the file's MFT record.

**MFT Entry with Resident Record**




Typically, each file uses one file record. However, if a file has a large number of attributes or becomes highly fragmented, it might need more than one file record. If this is the case, the first record for the file, the base file record, stores the location of the other file records required by the file.

Folder records contain index information. Small folder records reside entirely within the MFT structure, while large folders are organized B-tree structures and have records with pointers to external clusters that contain folder entries that cannot be contained within the MFT structure.

The benefit of using B-tree structures is evident when NTFS enumerates files in a large folder. The B-tree structure allows NTFS to group, or index, similar file names and then search only the group that contains the file, minimizing the number of disk accesses needed to find a particular file, especially for large folders. Because of the B-tree structure, NTFS outperforms FAT for large folders because FAT must scan all file names in a large folder before listing all of the files.

**Last Access Time**
Each file and folder on an NTFS volume contains an attribute called Last Access Time. This attribute shows when the file or folder was last accessed, such as when a user performs a folder listing, adds files to a folder, reads a file, or makes changes to a file. The most up-to-date Last Access Time is always stored in memory and is eventually written to disk within two places:

- The file's attribute, which is part of its MFT record.
- A directory entry for the file. The directory entry is stored in the folder that contains the file. Files with multiple hard links have multiple directory entries.

The Last Access Time on disk is not always current because NTFS looks for a one-hour interval before forcing the Last Access Time updates to disk. NTFS also delays writing the Last Access Time to disk when users or programs perform read-only operations on a file or folder, such as listing the folder's contents or reading (but not changing) a file in the folder. If the Last Access Time is kept current on disk for read operations, all read operations become write operations, which impacts NTFS performance.

### Note

- File-based queries of Last Access Time are accurate even if all on-disk values are not current. NTFS returns the correct value on queries because the accurate value is stored in memory.

NTFS eventually writes the in-memory Last Access Time to disk as follows.

**Within the file's attribute**
NTFS typically updates a file's attribute on disk if the current Last Access Time in memory differs by more than an hour from the Last Access Time stored on disk, or when all in-memory references to that file are gone, whichever is more recent. For example, if a file's current Last Access Time is 1:00 P.M., and you read the file at 1:30 P.M., NTFS does not update the Last Access Time. If you read the file again at 2:00 P.M., NTFS updates the Last Access Time in the file's attribute to reflect 2:00 P.M. because the file's attribute shows 1:00 P.M. and the in-memory Last Access Time shows 2:00 P.M.

**Within a directory entry for a file**
NTFS updates the directory entry for a file during the following events:

- When NTFS updates the file's Last Access Time and detects that the Last Access Time for the file differs by more than an hour from the Last Access Time stored in the file's directory entry. This update typically occurs after a program closes the handle used to access a file within the directory. If the program holds the handle open for an extended time, a lag occurs before the change appears in the directory entry.
- When NTFS updates other file attributes such as Last Modify Time, and a Last Access Time update is pending. In this case, NTFS updates the Last Access Time along with the other updates without additional performance impact.

### Note

- NTFS does not update a file's directory entry when all in-memory references to that file are gone.

If you have an NTFS volume with a high number of folders or files, and a program is running that briefly accesses each of these in turn, the I/O bandwidth used to generate the Last Access Time updates can be a significant percentage of the overall I/O bandwidth.

**Multiple Data Streams**
A data stream is a sequence of bytes. An application populates the stream by writing data at specific offsets within the stream. The application can then read the data by reading the same offsets in the read path. Every file has a main, unnamed stream associated with it, regardless of the file system used.

However, NTFS supports additional named data streams in which each data stream is an alternate sequence of bytes as illustrated in the figure Unnamed and Named Streams. Applications can create additional named streams and access the streams by referring to their names. This feature permits related data to be managed as a single unit. For example, a graphics program can store a thumbnail image of bitmap in a named data stream within the NTFS file containing the image.

### Unnamed and Named Streams




FAT volumes support only the main, unnamed stream, so if you try to copy or move Streamexample.doc to a FAT volume or floppy disk, you receive an error message.

**NTFS Processes and Interactions**

The following sections describe NTFS processes and interactions.

**Mounting an NTFS Volume**

When mounting an NTFS volume, the MBR executes code to start up the boot sector. The boot sector then executes additional code to mount the volume.

**Master Boot Code Startup Process**

The MBR contains a small amount of executable code called the master boot code, the disk signature, and the partition table for the disk. During startup, the master boot code performs the following activities:

1. Scans the partition table for the active partition.
2. Finds the starting sector of the active partition.
3. Loads a copy of the boot sector from the active partition into memory.
4. Transfers control to the executable code in the boot sector.

**Boot Sector Startup Process**

Computers use the boot sector to run instructions during startup. The initial startup process is summarized in the following steps:

1. The system BIOS and the CPU initiate the power-on self test (POST).
2. The BIOS finds the boot device, which is typically the first disk the BIOS finds, unless the controller is configured to boot from a different disk.
3. The BIOS loads the first physical sector of the boot device into memory and transfers CPU execution to that memory address.

If the boot device is on a hard disk, the BIOS loads the MBR. The master boot code in the MBR loads the boot sector of the active partition, and transfers CPU execution to that memory address. On computers that are running Windows Server 2003, the executable boot code in the boot sector finds Ntldr, loads it into memory, and transfers execution to that file.

**Note**

- Windows Server 2003 cannot start up from a spanned, striped, or RAID-5 volume on dynamic disks. These disk structures cannot be registered into the MBR partition table; therefore, a system volume that uses these structures cannot start.

If drive A contains a floppy disk, the system BIOS loads the first sector (the boot sector) of the disk into memory. If the disk is a startup disk (formatted by MS-DOS with core operating system files applied), the boot sector loads into memory and uses the executable boot code to transfer CPU execution to Io.sys, a core MS-DOS operating system file. If the floppy disk is not a startup disk, the executable boot code displays an error message.

**Note**

- These messages do not appear on normally functioning systems that are configured to look for the startup files on drive C first. On many computers, an option in the CMOS setup program allows the user to set the sequence of installed disks that the system searches to find the startup files.

If you get similar errors when trying to start the computer from the hard disk, the boot sector might be corrupted.

Initially, the startup process is independent of disk format and operating system. The unique characteristics of operating and file systems become important when the boot sector's executable boot code starts.

**Formatting Volumes**
During volume format, Windows Server 2003 places key NTFS file system structures on the volume, including the boot sector and the MFT as well as replacing Ntldr. Formatting also aligns clusters at the cluster size boundary.

Formatting a volume will check the integrity of all sectors on the volume during the process, as well as allow you to change the cluster size used on the volume. If a volume is formatted using Quick format, the file system structure on the volume is created, but the integrity of every sector in the volume is not checked.

**Converting Volumes**
Windows Server 2003 can convert previous versions of NTFS to the new version of NTFS used in Windows Server 2003.

**Converting NTFS Volumes Formatted By Using Windows 2000**
When Windows Server 2003 first mounts an NTFS volume that was formatted by using Windows 2000, Windows Server 2003 converts the NTFS volume to NTFS 3.1. The conversion consists of changing the NTFS version from 3.0 to 3.1. No other changes are made to existing metadata or files on the volume. However, Windows Server 2003 uses a different header style for new files created on NTFS 3.1 volumes. As a result of this change, some non-Microsoft imaging programs cannot create images of NTFS 3.1 volumes. Contact the manufacturer of your imaging program to find out if a version is available that supports NTFS 3.1 volumes in Windows Server 2003.

Computers running Windows NT 4.0 with Service Pack 4 or later or Windows 2000 can access NTFS 3.1 volumes without any conversion or additional service packs. Also, note that NTFS 3.1 is identical in Windows XP and Windows Server 2003.

**Converting NTFS Volumes Formatted By Using Windows NT 4.0 and Earlier**
When you upgrade the operating system from Windows NT 4.0 to Windows Server 2003, all local volumes formatted by using the version of NTFS used in Windows NT 4.0 and earlier are upgraded to NTFS 3.1. The upgrade occurs when Windows Server 2003 mounts the volume for the first time after Windows Server 2003 Setup is completed. (The upgrade does not take place during Setup.) Any NTFS volumes that are removed or turned off during Setup, or added after Setup, are converted when Windows Server 2003 mounts the volumes.

The Ntfs.sys driver performs the conversion by determining which version of NTFS is used on the volume and converting the volume if necessary. The conversion takes only a few seconds on any size volume and consists of the following new records in the master file table:

- $Secure, which contains unique security descriptors for all files within a volume.
- $Extend, which is used for extensions such as quotas, reparse point data, and object identifiers. The conversion process also adds three new files the to $Extend directory:
    - $Quota, used for disk quotas.
    - $Reparse, used for reparse points.
    - $ObjID, used for distributed link tracking.

Both $Secure and $Extend take the place of previously unused master file table (MFT) records, so sufficient space always exists in the volume for these two records. However, $Quota, $Reparse, and $ObjID are new additions to the MFT, and you must have enough free space in the volume to contain these files, or the conversion fails.

If the conversion fails, the volume is still available, but you can only perform NTFS-related tasks that were available in Windows NT 4.0 or earlier. To convert the volume to NTFS 3.1, you must free disk space by deleting or moving files and then dismount the volume.

**Note**

- Removable media that is formatted by using the previous version of NTFS is upgraded after the installation or upgrade process, or when you insert the media and Windows Server 2003 mounts it.

**Limitations in Converting Volumes**

The conversion is a one-way process. After you convert a volume to NTFS, you cannot reconvert the volume to FAT without backing up your data, reformatting the volume as FAT, and then restoring your data. There should also be a certain amount of free space on the volume and sufficient memory to update the cache.

The following limitations also apply to conversion of a volume from FAT to NTFS:

- In multiple-boot configurations, NTFS volumes are accessible only by using Windows NT 4.0 with Service Pack 4 or later, Windows 2000, Windows XP, or Windows Server 2003.

- When you install Recovery Console onto a volume that is formatted for either the FAT16 or FAT32 file systems, and then convert the volume to NTFS, the Recovery Console no longer runs. This problem occurs because the file-system-specific boot files (in the cmdcons folder of the system volume) that are used to run Recovery Console are not valid for a volume that has been converted to NTFS. You can re-install Recovery Console from the Windows Server 2003 operating system disk after the conversion. You can also use the Windows Server 2003 operating system disk to start Recovery Console.

- Because formatting in Windows Server 2003 aligns FAT data clusters at the cluster size boundary, conversion can preserve the cluster size for the size of the volume (up to 4 KB) instead of using the 512-byte cluster size used in Windows 2000 for converted volumes. The table Cluster Sizes for Volumes Converted to NTFS lists cluster sizes for volumes converted to NTFS.

  **Cluster Sizes for Volumes Converted to NTFS**

| Original FAT Cluster Size | Converted NTFS Cluster Size |
|---|---|
| 512 bytes | 512 bytes |
| 1 KB | 1 KB |
| 2 KB | 2 KB |
| 4 KB and larger | 4 KB |

**File Naming**

Windows Server 2003 supports both long and short file names on NTFS volumes.

**File Names in Windows Server 2003**

Every time you create a file with a long file name, NTFS creates a second file entry that has a similar 8.3 short file name. A file with an 8.3 short file name has a file name containing 1 to 8 characters and a file name extension containing 1 to 3 characters. The file name and file name extension are separated by a period.

File names in Windows Server 2003 can be up to 255 characters and can contain spaces, multiple periods, and special characters that are not allowed in MS-DOS file names. Windows Server 2003 makes it possible for other operating systems to access files that have long names by generating an MS-DOS-readable (8.3) name for each file. These MS-DOS-readable names also enable MS-DOS-based and Windows 3.*x*–based applications to recognize and load files that have long file names. When a program saves a file on a computer running Windows Server 2003, both the 8.3 file name and long file name are retained.

**Note**

- The 8.3 format means that files can have between 1 and 8 characters in the file name. The name must start with a letter or a number and can contain any characters except the following:

- . " / \ [ ] : ; | = , * ? (space)

- An 8.3 file name typically has a file name extension that is from one to three characters long and has the same character restrictions. A period separates the file name from the file name extension.

- Several special file names are reserved by the system and cannot be used for files or folders: CON, AUX, COM1, COM2, COM3, COM4, LPT1, LPT2, LPT3, PRN, NUL

**How NTFS Generates Short File Names**

In Windows Server 2003, both FAT and NTFS use the Unicode character set, which contains several prohibited characters that MS-DOS cannot read, for their names. To generate a short MS-DOS-readable file name, Windows Server 2003 deletes all of these characters from the long file name and removes any spaces. Because an MS-DOS-readable file name can have only one period, Windows Server 2003 also removes extra periods from the file name. If necessary, Windows Server 2003 truncates the file name to six characters and appends a tilde (**~**) and a number. For example, each non-duplicate file name is appended with **~1**. Duplicate file names end with **~2**, then **~3**, and so on. After the file names are truncated, the file name extensions are truncated to three or fewer characters. Finally, when displaying file names at the command line, Windows Server 2003 translates all characters in the file name and extension to uppercase.

**Note**

- You can permit extended characters by using the **fsutil behavior set** command. You must restart the computer before this setting takes effect. For more information about using the **fsutil behavior set** command, see the topic Fsutil: behavior in Help and Support Center in Windows Server 2003.

When five or more files exist that can result in duplicate short file names, Windows Server 2003 uses a slightly different method for creating short file names. For the fifth and subsequent files, Windows Server 2003:

- Uses only the first two letters of the long file name.
- Generates the next four letters of the short file name by mathematically manipulating the remaining letters of the long file name.
- Appends **~1** (or another number, if necessary, to avoid a duplicate file name) to the result.

This method substantially improves performance when Windows Server 2003 must create short file names for a large number of files with similar long file names. Windows Server 2003 uses this method to create short names for files on both FAT and NTFS volumes.

The following table shows the short file names for files created by six tests.

**Short File Names Created by Windows Server 2003 — Example One**

| Long File Name | Short File Name |
|---|---|
| This is test 1.txt | THISIS~1.TXT |
| This is test 2.txt | THISIS~2.TXT |
| This is test 3.txt | THISIS~3.TXT |
| This is test 4.txt | THISIS~4.TXT |
| This is test 5.txt | THA1CA~1.TXT |
| This is test 6.txt | THA1CE~1.TXT |

If the long file names in the preceding table are created in a different order, their short file names are different, as shown in the following table.

**Short File Names Created by Windows Server 2003 — Example Two**

| Long File Name | Short File Name |
|---|---|
| This is test 2.txt | THISIS~1.TXT |

| | |
|---|---|
| This is test 3.txt | THISIS~2.TXT |
| This is test 1.txt | THISIS~3.TXT |
| This is test 4.txt | THISIS~4.TXT |
| This is test 5.txt | THA1CA~1.TXT |
| This is test 6.txt | THA1CE~1.TXT |

When you delete a file, its short file name is also deleted. When you create new files in the same folder, Windows Server 2003 might re-use short file names that have been deleted. For instance, in Example 1, if you delete the file "This is test 1.txt," and then create a new file called "This is test 7.txt," its short file name becomes THISIS~1.TXT.

If you have a large number of files (300,000 or more) in a folder, and the files have long file names with the same initial characters, the time required to create the files increases. The increase occurs because NTFS bases the short file name on the first six characters of the long file name. In folders with more than 300,000 files, the short file names start to conflict after NTFS uses all of the 8.3 names that are similar to the long file names. Repeated conflicts between a generated short file name and existing short file names cause NTFS to regenerate the short file name from 6 to 8 times.

**Compression of Files and Folders**

NTFS supports compression on individual files, all files within a folder, and all files within NTFS volumes. Because compression is implemented within NTFS, any Windows-based program can read and write compressed files without determining the compression state of the file. Compression is set in a bit within the file header, while information about compression is stored in the Data file attribute.

Files and directories are compressed and decompressed by passing FSCTL_SET_COMPRESSION code to DeviceIoControl. A compressed file or directory then has the flag FILE_ATTRIBUTE_COMPRESSED associated with it. Applications can determine a file or directory's compression state using GetFileAttributes.

When a program opens a compressed file, NTFS decompresses only the portion of the file being read and then copies that data to memory. By leaving data in memory uncompressed, NTFS performance is not impacted when it reads or modifies data in memory. NTFS compresses the modified or new data in the file when the data is later written to disk.

The compression algorithms in NTFS support cluster sizes of up to 4 KB. When the cluster size is greater than 4 KB on an NTFS volume, none of the NTFS compression features are available.

**Moving and Copying Files or Folders**

Moving and copying files and folders can change their compression state. The resulting compression state depends on whether you move or copy the files and whether you move the files between NTFS volumes or to FAT volumes.

NTFS supports compression on one file, all files in a directory, or all files on a volume. Compression is set in a bit within the file header, while information about compression is stored in the Data file attribute. When the bit is set, the system compresses the file when it is saved and decompresses it as needed.

Compression adds overhead to the system because a compressed NTFS file is decompressed, copied, and then recompressed as a new file even when the file is copied in the same computer. Any change to the compression attribute is applied to the files you specify for moving or copying. If you compress all files in the volume, the process might take a few minutes to finish, depending on the size of the volume, the number of files to compress, and the speed of the computer. The delay occurs because Windows Server 2003 must change the compression state of every folder on the volume and compress or uncompress every file on the volume.

Changing the compression state of folders is relatively fast because for each folder Windows Server 2003 changes only the compression attribute. However, compressing or uncompressing every file on the volume takes longer because NTFS must read data in its current form (compressed or uncompressed) from the disk, convert the data to its new form in memory, and then write the data back to disk.

**Moving Files or Folders within an NTFS Volume**

When you move an uncompressed file or folder to another folder on the NTFS volume, the file remains uncompressed. The figure Moving an Uncompressed File to a Compressed Folder illustrates the result of moving an uncompressed file to a compressed folder.

### Moving an Uncompressed File to a Compressed Folder




When you move a compressed file to an uncompressed folder, the file remains uncompressed after the move. The figure Moving a Compressed File to an Uncompressed Folder illustrates the result of moving a compressed file or folder to a compressed folder.

**Moving a Compressed File to an Uncompressed Folder**




**Copying Files or Folders within an NTFS Volume**

Copying a file to a folder takes on the compression attribute of the target folder.

If you copy a compressed file to an uncompressed folder, the file is uncompressed when it is copied to the folder, as shown in the figure Copying a Compressed File to an Uncompressed Folder.

**Copying a Compressed File to an Uncompressed Folder**




When you copy a file to a folder that already contains a file of the same name, the copied file takes on the compression attribute of the target file, as shown in the figure Copying a File to a Folder that Contains a File of the Same Name.

**Copying a File to a Folder that Contains a File of the Same Name**




**Copying Files between FAT and NTFS Volumes**

Files copied from a FAT folder to an NTFS folder take on the compression attribute of the target folder. Compressed files copied from an NTFS volume to a FAT volume or floppy disk are uncompressed.

**Mounted Drives on NTFS Volumes**

Mounted drives, also known as volume mount points or drive paths, are volumes attached to an empty folder on an NTFS volume. Mounted drives function the same way as any other volume, but are assigned a label or name instead of a drive letter. Mounted drives are robust against system changes that occur when devices are added or removed from a computer. They are not subject to the 26-volume limit imposed by drive letters, so you can use them for access to more than 26 volumes on your computer.

The version of NTFS included with Windows Server 2003 must be used on the host volume. However, the volume to be mounted can be formatted in any file system supported by Windows Server 2003.

One volume can host multiple mounted drives, providing a way for you to easily extend the storage capacity of any particular volume on a Windows Server 2003 system. Users on the local computer or users who connect to it over a network can continue to use the same drive letter for access to the volume, but multiple volumes can be in use simultaneously from that drive letter.

Only NTFS volumes can hold a mounted drive, although any local drive can be mounted on one.

**Implementing Mounted Drives**
NTFS mounted drives are implemented by using reparse points and are subject to their restrictions. Reparse points are a collection of user-defined data. The format of this data is understood by the application which stores the data, and a file system filter, which you install to interpret the data and process the file. When an application sets a reparse point, it stores this data, plus a reparse tag, which uniquely identifies the data it is storing. When the file system opens a file with a reparse point, it attempts to find the file system filter associated with the data format identified by the reparse tag. If a file system filter is found, the filter processes the file as directed by the reparse data. If a file system filter is not found, the file open operation fails.

The following restrictions apply to reparse points:

- Reparse points can be established for a directory, but the directory must be empty. Otherwise, NTFS fails to establish the reparse point. In addition, you cannot create directories or files in a directory that contains a reparse point.

- Reparse points and extended attributes are mutually exclusive. NTFS cannot create a reparse point when the file contains extended attributes, and it cannot create extended attributes on a file that contains a reparse point.

- Reparse point data cannot exceed 16 kilobytes. Setting a reparse point fails if the amount of data to be placed in the reparse point exceeds this limit.

**Hard Links**
A hard link is an NTFS-only based link to a given file. When you create a hard link to a file on an NTFS volume, NTFS adds a directory entry for the hard link without duplicating the original file. By creating hard links you can:

- Use the same file name as the original file but appear in different folders.

- Use different file names from the original file but appear in the same folder.

- Use different file names from the original file and appear in different folders.

Because a hard link is a directory entry for a file, an application can modify a file by using any of its hard links. Applications that use any other hard link can detect the changes. However, directory entries for hard links are updated only when a user accesses a file by using the hard link. For example, if a user opens and modifies a file by using its hard link, and the size of the original file changes, the hard link that is used to access the file also shows the new size.

Hard links do not have security descriptors; instead, the security descriptor belongs to the original file to which the hard link points. Thus, if you change the security descriptor of any hard link, you actually change the underlying file's security descriptor. All hard links that point to the file allow the newly specified access. You cannot give a file different security descriptors on a per-hard-link basis.

Hard links use the Win32 function CreateHardLink to create hard links between files.

**Distributed Link Tracking**
Distributed link tracking ensures that shell shortcuts and OLE links continue to work after the target file is renamed or moved. When you create a shortcut to a file on an NTFS volume, distributed link tracking stamps a unique object identifier (ID) into the target file, known as the link source. Information about the object ID is also stored within the referring file, known as the link client. Distributed link tracking uses this object ID to locate the link source in any combination of the following events that occur on NTFS volumes within a Windows Server 2003-based domain:

- The link source is renamed.

- The link source is moved to another folder on the same volume or to a different volume on the same computer.

- The link source is moved from one shared network folder to another shared network folder on different computers within the same domain.

- The computer containing the link source is renamed.
- The name of the shared network folder containing the link source has changed.
- The volume containing the link source is moved to another computer within the same domain.

**Note**

- Distributed link tracking works only on NTFS volumes in computers running Windows 2000, Windows XP, or Windows Server 2003. The NTFS volumes cannot be on removable media.

Distributed link tracking attempts to maintain even those links that do not occur within a domain: cross-domain, within a workgroup, or on a single computer that is not connected to a network. Links can always be maintained in these events when a link source is moved within a computer, or when the network shared folder on the link source computer is changed. Typically, links can be maintained when the link source is moved to another computer; however, this form of tracking is less reliable over time.

Distributed link tracking uses different services for client and server:

- The Distributed Link Tracking Client service runs on all Windows 2000-based and Windows Server 2003-based computers. In computers that are not part of a network, the Client service performs all activities related to link tracking.
- The Distributed Link Tracking Server service runs on Windows 2000 and Windows Server 2003 domain controllers. The Server service maintains information relating to the movement of link sources. Because of this service and the information it maintains, links within a domain are more reliable than those outside a domain. For computers that run in a domain, the Distributed Link Tracking Client service takes advantage of this information by communicating with the Distributed Link Tracking Server service.

The Distributed Link Tracking Client service monitors activity on NTFS volumes and stores maintenance information in a file called Tracking.log, which is located at the root of each volume in a hidden folder called System Volume Information. This folder is protected by permissions that allow only the system to have access to it. The System Volume Information folder is also used by other Windows Server 2003 services such as Indexing Service.

**Sparse Files**

Sparse files provide a method of saving disk space for files that contain meaningful data, as well as large sections of data composed of zeros. If an NTFS file is marked as sparse, then NTFS allocates disk clusters only for the data explicitly specified by the application. Non-specified ranges of the file are represented by non-allocated space on the disk. When a sparse file is read from allocated ranges, the data is returned as it was stored. Data read from non-allocated ranges is returned as zeros.

File system application programming interfaces (APIs) allow for the file to be copied or backed as actual bits and sparse stream ranges.File system APIs also allow for querying allocated ranges. Programs that implement these APIs then need only to read allocated ranges to recover all data in the file. The result is efficient file system storage and access. The figure Sparse Data Storage shows how data is stored with and without the sparse file attribute set.

**Sparse Data Storage**




For example, the properties of a file might show that the file is a 1-GB sparse file. Although the file is 1 GB, it occupies only 64 KB of disk space.

### Note

- Only NTFS volumes mounted by Windows 2000, Windows XP, or the Windows Server 2003 family support sparse files. If you copy or move a sparse file to a FAT volume or an NTFS volume mounted by an operating system other than those listed previously, the file is built to its originally specified size. If the required space is not available, the operation fails.

**Disk Quotas**

You can enable disk quotas to restrict the amount of volume space users take up on remote or local computers with NTFS file systems. Disk quotas uses names from the domain in which the server resides. An administrator can then set disk quotas against those users in the domain.

For additional information about Disk Quotas, see the Disk Quotas Technical Reference.

**NTFS Change Journal**

As files, folders, and other NTFS objects are added, deleted, and modified, NTFS enters change journal records in streams, one for each volume on the computer.

The total size of all the records currently in the journal varies, but there is a configurable maximum size. The change journal can exceed the maximum size until the size reaches an outer threshold, at which point a portion of the oldest records are deleted until the change journal is restored to its maximum size. The maximum size of the change journal is configurable but cannot be reduced, only increased.

The change journal conveys significant scalability benefits to applications that might otherwise need to scan an entire volume for changes. File system indexing, replication managers, virus scanners, and incremental backup applications can benefit from using the change journal.

The change journal is much more efficient than time stamps or file notifications for determining changes in a particular namespace. Applications that must rescan an entire volume to determine changes can now scan once and subsequently refer to the change journal. The I/O cost depends on how many files have changed, not on how many files exist on the volume.

The APIs are fully documented and can be leveraged by independent software vendors (ISVs). Microsoft uses the change journal in Windows Server 2003 components such as the Indexing Service and File Replication Service. ISVs can use this feature to enhance the scalability and robustness of a range of products including backup, antivirus, and auditing tools.

**NTFS File System Recoverability**

NTFS is a recoverable file system that guarantees the consistency of the volume by using standard transaction logging and recovery techniques. In the event of a system failure, NTFS runs a recovery procedure that accesses information stored in a transaction log file. The NTFS recovery procedure guarantees that the volume is restored to a consistent state. Transaction logging requires very little overhead.

**Recovering NTFS File Structures**

NTFS views each operation that modifies a file on a volume as a transaction and manages each one as an integral unit. NTFS

might also break a single complex operation into multiple transactions. After a transaction is started, it is either completed, or if an event occurs that causes the operation to fail, it is rolled back, and the NTFS volume returns to its state before the transaction began. Events that can cause an operation to fail include bad sectors, transient low-memory conditions, and disconnected devices.

To ensure that a transaction can either be completed or rolled back, NTFS performs the following steps for each transaction:

1. Records the metadata operations of a transaction in a log file cached in memory.
2. Records the actual metadata operations in memory.
3. Marks the transaction in the cached log file as committed.
4. Flushes the log file to disk.
5. Flushes the actual metadata operations to disk.

Steps 4 and 5 occur in a lazy fashion after the transaction is completed, meaning that the flush operations are not tied to the transaction itself. Instead, NTFS modifies the log and metadata quickly in memory, and then flushes later at a convenient time to boost performance.

NTFS guarantees that the log records containing the metadata operations of the transaction are written to disk before the metadata that is modified in the transaction is written to disk. After NTFS updates the cache, NTFS commits the transaction by recording in the cached log file that the transaction is complete. After the cached log file is flushed to disk, all committed transactions are guaranteed to be completed, even if the system fails before the changes are written to disk.

**Note**

- Applications can specify the FILE_FLAG_WRITE_THROUGH Win32 flag to instruct the system to write through any intermediate cache and go directly to disk. The system can still cache write operations, but cannot lazily flush them.

If a system failure occurs, NTFS has enough information in the log to complete or abort any partial NTFS transaction. During recovery operations, NTFS redoes each committed transaction found in the log file. Then NTFS locates in the log file the transactions that were not committed at the time of the system failure and undoes each metadata operation recorded in the log file. Because NTFS flushes the log to disk before any metadata changes are written to disk, NTFS has complete information available about any metadata changes that need to be rolled back during recovery.

**Note**

- NTFS uses transaction logging and recovery to guarantee that the volume structure is not corrupted. For this reason, all file system data is accessible after a system failure. NTFS guarantees user data only if the program used to create the data uses the FILE_FLAG_WRITE_THROUGH Win32 flag. If the program does not use this flag, user data can be lost due to a system failure. If a system failure does occur, NTFS shows either the previous data, the new data, or zeros. Users do not see random data on the volume as the result of a crash.

**Caching and Data Recovery**

The cache is the area of RAM that contains the most recently used data. When you write data to disk, the lazy-writetechnique in Windows Server 2003 indicates that the data is written when it is still in the cache. Cache memory can also be on the disk controller, such as cache memory available on SCSI controllers, or on the disk unit, such as cache memory available on Advanced Technology Attachment (ATA) disks. The following information can help you decide whether to enable the disk or the controller cache:

- Write caching improves disk performance, particularly if large amounts of data are being written to the disk.
- Control of the write-back cache is a firmware function provided by the disk manufacturer. See the documentation supplied with the disk or disk controller. You cannot configure the write-back cache from Windows Server 2003.
- Write caching does not impact the reliability of the file system's own metadata as long as the firmware provided by the disk manufacturer honors write-through requests issued by the NTFS driver. NTFS instructs the disk device driver to ensure that metadata is written whether or not write caching is enabled. Non-metadata is typically written to disk and

can be cached.

- Read caching in the disk does not affect the reliability of a file system.

**Cluster Remapping**
When NTFS detects a bad sector, NTFS dynamically remaps the cluster containing the bad sector — a recovery technique called cluster remapping — and allocates a new cluster for the data. If the error occurred during a read, NTFS returns a read error to the calling program, and the data is lost. If the error occurs during a write, NTFS writes the data to the new cluster, and no data is lost.

NTFS puts the address of the cluster containing the bad sector in the bad cluster file, $BadClus, in the MFT so that the bad sector is not reused.

**Disk Recovery Operations**
NTFS ensures the integrity of all NTFS volumes by performing disk recovery operations whenever a volume is mounted after the computer is restarted or after the volume is dismounted. NTFS also uses a technique called cluster remapping to minimize the effects of a bad sector on an NTFS volume.

NTFS views each operation that modifies a file on a volume as a transaction and manages each one as an integral unit. NTFS might also break a single complex operation into multiple transactions. After a transaction is started, it is either completed, or if an event occurs that causes the operation to fail, it is rolled back, and the NTFS volume returns to its state before the transaction began. Events that can cause an operation to fail include bad sectors, transient low-memory conditions, and disconnected devices.

NTFS guarantees that the log records containing the metadata operations of the transaction are written to disk before the metadata that is modified in the transaction is written to disk. After NTFS updates the cache, NTFS commits the transaction by recording in the cached log file that the transaction is complete. After the cached log file is flushed to disk, all committed transactions are guaranteed to be completed, even if the system crashes before the changes are written to disk.

If a system failure occurs, NTFS has enough information in the log to complete or abort any partial NTFS transaction. During recovery operations, NTFS redoes each committed transaction found in the log file. Then NTFS locates in the log file the transactions that were not committed at the time of the system failure and undoes each metadata operation recorded in the log file. Because NTFS flushes the log to disk before any metadata changes are written to disk, NTFS has complete information available about any metadata changes that need to be rolled back during recovery.

**Cleanup Operations on Windows NT–Based Volumes**
Because files on volumes formatted by using the version of NTFS included with Windows Server 2003 can be read and written to by Windows NT 4.0 Service Pack 4 or later, Windows Server 2003 might need to perform cleanup operations to ensure the consistency of the data structures of a volume after it is mounted on a computer running Windows NT.

Windows Server 2003 does not perform cleanup operations on volumes previously mounted by using Windows 2000 or Windows XP.

Cleanup operations affect the following features:

**Reparse points**
Computers running Windows NT 4.0 or earlier cannot access files that have reparse points, so no cleanup operations are necessary. Reparse points are files or directories that have blocks of data called reparse data associated with them.

**Disk quotas**
If disk quotas are turned off, Windows Server 2003 performs no cleanup operations. If disk quotas are turned on, Windows Server 2003 cleans up the quota information by rebuilding the index. If a user exceeds the disk quota while the NTFS volume is mounted by a Windows NT 4.0 SP4 or later system, and disk quotas are strictly enforced, all further disk allocations of data by that user using Windows Server 2003 fail. The user can still read and write data to any existing file but cannot increase the size of a file. However, the user can delete and shrink files. When usage falls below the assigned disk quota, disk allocations of data can resume.

**Encryption**
Encrypted files cannot be accessed by computers that are running Windows NT 4.0 or earlier, so no cleanup operations are necessary.

**Sparse files**
Computers running Windows NT 4.0 or earlier cannot access sparse files, so no cleanup operations are necessary.

**Change journal**

Computers that are running Windows NT 4.0 or earlier do not log file changes in the change journal. When Windows Server 2003 starts, the change journals on volumes accessed by Windows NT are reset to indicate that the journal history is incomplete. Applications that use the change journal must be able to accept incomplete journals.

**Object identifiers**
Windows Server 2003 maintains two references to the object identifier: one on the file and one in the volume-wide object identifier index. If you delete a file that has an object identifier, Windows Server 2003 must scan and clean up the entry in the index.

**POSIX Compliance**
NTFS provides a several features to support the Portable Operating System Interface (POSIX) standard, which is defined by the Institute of Electrical and Electronic Engineers (IEEE) standard 1003.1-1990 (also known as ISO/IEC 9945-1:1990).

NTFS includes the following POSIX-compliant features.

**Case-sensitive naming**
For example, POSIX interprets README.TXT, Readme.txt, and readme.txt as separate files.

**Hard links**
A file can have more than one name. This allows two different file names, which can be in different folders on the same volume, to point to the same data.

**Additional time stamps**
These show when the file was last accessed or modified.

The POSIX subsystem included with Windows NT and Windows 2000 is not included with Windows Server 2003. A new subsystem supporting the broad functionality found on most UNIX systems beyond the POSIX.1 standard is shipped as part of Interix 2.2. The Interix subsystem can be certified to the NIST FIPS 151-2 POSIX Conformance Test Suite.

### Note

- You must use Interix-based programs to manage file names that differ only in case. You cannot use standard Windows Server 2003 command-line tools (such as **copy**, **del**, and **move**, or their equivalents in Windows Explorer or My Computer) to manage file names that differ only in case.

**Related Information**

- Basic Disks and Volumes Technical Reference
- Disk Quotas Technical Reference
- Dynamic Disks and Volumes Technical Reference

# NTFS Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NTFS Tools and Settings**
**In this section**


- NTFS Tools
- NTFS Registry Entries
- NTFS Group Policy Settings
- NTFS WMI Classes


The following tools, registry settings, Group Policy settings, and Windows Management Instrumentation (WMI) classes are associated with NTFS volumes running Microsoft Windows Server 2003.

**NTFS Tools**
NTFS Tools describes the tools installed with Windows Server 2003 as well as downloadable support tools. The tools in this list are those used to create and maintain NTFS volumes.

**Cacls.exe: Set Access Control List**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Cacls.exe runs only on an NTFS volume.

Cacls.exe displays and modifies ACLs of files or folders.

For more information about Cacls, see Command Line References in Tools and Settings Collection.

**Chkdsk.exe: Check Disk**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Chkdsk.exe works on both NTFS and FAT volumes.

Chkdsk is a command-line tool that verifies the logical integrity of a file system on a Windows Server 2003 volume. Setting parameters on the command line will list and fix any errors.

For more information about Chkdsk, see Command Line References in Tools and Settings Collection.

**Chklnks.exe: Link Check Wizard**
**Category**
Resource Kit tool.

**Version compatibility**
Windows Server 2003. Chklnks.exe works only on an NTFS volume.

Chklnks.exe is a command-line tool that scans all the link (shortcut) files on a computer to determine whether or not the shortcuts point to existing applications or documents.

For more information about Chklnks, see Resource Kit Tools in Tools and Settings Collection.

**Chkntfs.exe: Chkntfs**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Chkntfs.exe works on both NTFS and FAT volumes.

Displays or specifies whether automatic system checking is schedule to be run on an NTFS volume when the computer is started.

For more information about Chkntfs, see Command Line References in Tools and Settings Collection.

**Cipher.exe: Cipher**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Cipher.exe works only on an NTFS volume.

Displays or alters the encryption of folders and files on NTFS volumes.

For more information about Cipher, see Command Line References in Tools and Settings Collection.

**Compact.exe: Compact**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Compact.exe works only on an NTFS volume.

Displays and alters the compression of files or directories on NTFS volumes.

For more information about Compact, see Command Line References in Tools and Settings Collection.

**Compress.exe: Compress Files**
**Category**
Resource Kit tool.

**Version compatibility**
Windows Server 2003. Compress.exe works only on an NTFS volume.

Compress.exe is a command-line tool to create compressed copies of one or more files.

For more information about Compress, see Resource Kit Tools in Tools and Settings Collection.

**Convert.exe: Convert**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Convert.exe works on both NTFS and FAT volumes.

Converts a FAT16 or FAT32 volume to NTFS. Conversion is a one-way process. You can convert a FAT file system to NTFS, but you cannot convert an NTFS volume to a FAT file system.

For more information about Convert , see Command Line References in Tools and Settings Collection.

**Defrag.exe: Defragmentation**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Defrag.exe runs on both NTFS and FAT volumes.

Arranges clusters on a disk in an optimized order.

For more information about Defrag, see Command Line References in Tools and Settings Collection.

**DiskPart.exe: Disk Partition**
Disk Partition is a command-line tool that allows you to manage disks and partitions by using scripts or direct input from a command prompt.

**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. DiskPart includes a few 64-bit parameters that are only available on Itanium-based computers.

Locates and consolidates fragmented boot files, data files, and folders on local or remote volumes.

For more information about DiskPart, see Command Line References in Tools and Settings Collection.

**Expand.exe: Expand Compressed Files**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Expand.exe works only on NTFS volumes.

Extracts a file from a cabinet (.CAB) or compressed file.

For more information about Expand, see Command Line References in Tools and Settings Collection.

**Format.exe: Format**
**Category**
Operating system tool

**Version compatibility**
Windows Server 2003. Works on both NTFS and FAT volumes.

Format prepares a volume on the specified disk to accept Windows files.

For more information about Format, see Command Line References in Tools and Settings Collection.

**Fsutil.exe: File Systems Utilities**
**Category**
Operating System tool.

**Version compatibility**
Windows Server 2003. Runs on both NTFS and FAT file systems.

Perform file system tasks such as disabling long file names, checking whether a volume is flagged as dirty, viewing NTFS-related information about a volume, creating hard links, and managing quotas and sparse files.

For more information about Fsutil, see Command Line References in Tools and Settings Collection.

**Hlscan.exe: Hard Link Display Tool**
**Category**
Resource Kit tool.

**Version compatibility**
Windows Server 2003. Hlscan.exe works only on an NTFS volume.

Hlscan.exe is a command-line tool that displays hard links on an NTFS volume or in specified files or directories of the volume.

For more information about Compress, see Resource Kit Tools in Tools and Settings Collection.

**Mountvol.exe: Mount Volume**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Mountvol.exe works only on an NTFS volume.

Creates, deletes, or lists a mounted drive.

For more information about Mountvol, see Command Line References in Tools and Settings Collection.

**Recovery Console**
**Category**
Operating system tool.

**Version compatibility**
Windows Server 2003. Recovery Console works on both NTFS and FAT volumes.

The Recovery Console is a text-mode command-line interpreter that you can use for basic troubleshooting and maintenance. It is separate from the Windows Server 2003 command prompt and grants limited local hard disk access for both FAT and NTFS volumes.

**Vrfydsk.exe: Verify Disk**
**Category**
Resource Kit tool.

**Version compatibility**
Windows Server 2003. Vrfydsk.exe works on both NTFS and FAT volumes.

Vrfydsk is a command-line tool that is a command-line tool that checks the consistency and integrity of the file system (FAT or NTFS) on a specified volume.

For more information about DmDiag, see Resource Kit Tools in Tools and Settings Collection.

**NTFS Registry Entries**
The following registry entries are associated with NTFS.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**FileSystem**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\FileSystem.

**NtfsAllowExtendedCharacterIn8dot3Name**
**Registry Path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem

**Version**
Microsoft Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Allows generation of short (8-character name, 3-character extension) file names on NTFS file systems.

For more information about this registry entry, see the Registry Reference in Tools and Settings Collection.

**NtfsDisable8dot3NameCreation**
**Registry Path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem

**Version**
Microsoft Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Disables generation of short (8-character name, 3-character extension) file names on NTFS file systems.

For more information about this registry entry, see the Registry Reference in Tools and Settings Collection.

**NTFSDisableLastAccessUpdate**
**Registry Path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem.

**Version**
Microsoft Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Disables last accessed timestamp of a file.

For more information about this registry entry, see the Registry Reference in Tools and Settings Collection.

**NTFSEncryptionService**
**Registry Path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem

**Version**
Used in Windows 2000 only.

Specifies which encryption service the NTFS file system used. Encrypting File System (EFS) is the encryption service built into Windows Server 2003.

For more information about this registry entry, see the Registry Reference in Tools and Settings Collection.

**NtfsMftZoneReservation**
**Registry Path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem

**Version**
Microsoft Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Reserves the amount of space allocated to the MFT. This registry entry is not automatically created. You must manually add this key to the registry entry.

Because this is an optional entry that you must add, this registry entry is not documented in the Registry Reference for Windows Server 2003.

**System**
The following registry entry is located under HKEY_LOCAL_MACHINE\SYSTEM\

**MountedDevices**
**Registry Path**
HKEY_LOCAL_MACHINE\SYSTEM\

**Version**
Windows Server 2003.

Stores database information for all NTFS partitions on a volume. This allows NTFS to maintain the persistent names of volumes.

For more information about this registry entry, see the Registry Reference in Tools and Settings Collection.

**NTFS Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with NTFS.

### Group Policy Settings Associated with NTFS

| Group Policy Setting | Description |
|---|---|
| DLT_AllowDomainMode | Enables programs to track linked files that are moved within an NTFS volume, to another NTFS volume on the same computer, or to an NTFS volume on another computer. |
| DiskQuotaEnable | Enables disk quotas on an NTFS volume. This policy is not used on FAT volumes. |
| DiskQuotaLimit | Specifies the disk quota limit and warning level. This policy is not used on FAT volumes. |
| DiskQuotaLogEventOverThreshold | Specifies whether an application log event will be created when a user exceeds his or her disk quota limit. This policy is not used on FAT volumes. |
| DiskQuotaApplyToRemovableMedia | Specifies whether disk quotas policies are to be applied to removable media. This policy is not used on FAT volumes. |
| ExplorerNoResolveSearch | Prevents the system from conducting a comprehensive search of the target drive to resolve a shortcut. FAT volumes do not use tracking features. |
| ExplorerNoResolveTrack | Prevents the system from using NTFS tracking features to resolve a shortcut. FAT volumes do not use tracking features. |

For more information about Group Policy settings, see Group Policy Settings Reference in Tools and Settings Collection.

**NTFS WMI Classes**
The following table lists and describes the WMI classes that are associated with NTFS.

### WMI Classes Associated with NTFS

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_DiskDriveToDiskPartition | //root/cimv2 | Windows 2000 and later |
| Win32_DiskPartition | //root/cimv2 | Windows 2000 and later |

| | | |
|---|---|---|
| Win32_LogicalDisk | //root/cimv2 | Windows 2000 and later |
| Win32_LogicalDiskRootDirectory | //root/cimv2 | Windows 2000 and later |
| Win32_LogicalDiskToPartition | //root/cimv2 | Windows 2000 and later |
| Win32_DiskDrivePhyiscalMedia | //root/cimv2 | Windows 2000 and later |
| Win32_Volume | //root/cimv2 | Windows Server 2003 |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

# FAT Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FAT Technical Reference**

A file system is a part of the operating system that determines how files are named, stored, and organized on a volume. A file system manages files and folders, and the information needed to locate and access these items by local and remote users. The Microsoft Windows Server 2003 operating system supports both the file allocation table (FAT) and NTFS file systems.

Some earlier Microsoft Windows operating systems, as well as removable disks and floppy disks, support only the FAT file system. For magnetic disks you should use NTFS due to its reliability to ensure metadata integrity for fast recoverability, security, compression, and encryption.

**In this subject**


- What Is FAT?
- How FAT Works
- FAT Tools and Settings

# What Is FAT?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Common FAT Scenarios**
This section describes a few scenarios in which the FAT file systems are used.

**Mounting a server with a small volume**
If you have a small partition (less than 512 megabytes) on a server, the FAT file system imposes the least overhead on a small capacity partition, but still performs well.

**Running multiple operating systems**
In a multiple-boot configuration with NTFS volumes, you also need a FAT volume if you want to access any of the NTFS volumes from one of the other NTFS volumes. For example, to start a Windows Server 2003-based computer in Microsoft MS-DOS, Microsoft Windows 3.*x*, or Microsoft Windows 95, you must use FAT16. For a multiple-boot configuration that has Microsoft Windows 95 OEM Service Release 2 (OSR2), Microsoft Windows 98, or Microsoft Windows Millennium Edition (Me), you must use FAT32.

**Using FAT with older operating systems**
Microsoft Windows 95 and Microsoft Windows 98 use the FAT file system. These operating systems do not support NTFS.

**Using FAT with removable disks**
For removable disks that can be disconnected or ejected unexpectedly, you must use FAT16 or FAT32.

**Operating System and FAT Compatibility**
File systems are designed to allow you to format the volume of a hard disk in the way that is most appropriate for your use. Before formatting a volume for a file system, you can determine what the volume will be used for, and which operating system will be mounted on that volume. Sometimes you need to run more than one operating system on your hard disk. The following table shows which FAT file systems are supported by which operating systems.

### Operating System and File System Compatibility

| Operating System | FAT16 | FAT32 |
|---|---|---|
| Windows XP | ● | ● |
| Windows Server 2003 | ● | ● |
| Windows 2000 | ● | ● |
| Windows NT 4.0 | ● | |
| Windows 95 OSR2, Windows 98, and Windows Me | ● | ● |
| Windows 95 (prior to OSR2) | ● | |
| MS-DOS | ● | |

**Dependencies on Other Technologies**
FAT is dependent upon the following technology.

**Basic disks and volumes**
Basic disks and basic volumes are the storage types most often used with Windows operating systems. The term basic disk refers to a disk that contains basic volumes, such as primary partitions and logical drives. The term basic volume refers to a partition on a basic disk. Basic disks, which are found in both x86-based and Itanium-based computers, provide a simple storage solution that can accommodate changing storage requirements.

**Related Technology**
FAT is related to the following technology:

**NTFS file system**
NTFS allows you to gain the maximum benefits for the needs of today's enterprise business environments from Windows Server 2003, such as increased security, more robust and reliable performance, as well as a design for greater storage growth,

features not found in FAT.

**Related Information**

- Basic Disks and Volumes Technical Reference
- NTFS Technical Reference

# How FAT Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- FAT Architecture
- FAT Physical Structure
- FAT Processes and Interactions
- Related Information


A file system is a required part of the operating system on a volume that determines how files are named, stored, and organized. A file system manages files and folders, and the information needed to locate and access these items by local and remote users.

Microsoft Windows Server 2003 supports the file allocation table (FAT) file system on basic disks and on readable/writable disks. Basic disks and volumes are the storage types most often used with Windows operating systems.

During the format of a volume, you can choose the type of file system for the volume. When you choose the FAT file system, the formatting process places key file data structures on the volume to tell the hard disk which file system the operating system uses.

**FAT Architecture**

During format and setup of a volume file system on a hard disk, a master boot record (MBR) is created. The MBR contains a small amount of executable code called the master boot code as well as a partition table for the disk. When a volume is mounted, the MBR executes the master boot code and transfers control to the boot sector on the disk, allowing the server to start the operating system on the file system of that specific volume.

**Note**

- The partition table contains a number of fields used to describe the partition. One of these fields is the System ID field, which defines the file system, such as FAT, on the partition. For FAT16 volumes, the system ID is 0x04; for FAT32 volumes, the system ID is 0x0B.

The following figure shows the architecture of this process.

**FAT Architecture**



The following table describes the components of the FAT file system.

**FAT Components**

| Component | Component Description |
|---|---|
| Hard disk | Contains one or more partitions. |
| Boot sector | Bootable partition that stores information about the layout of the volume and the file system structures, as well as the boot code that loads Ntdlr. |
| Master Boot Record | Contains executable code that the system BIOS loads into memory. The code scans the MBR to find the partition table to determine which partition is the active, or bootable, partition. |
| Ntldlr.dll | Switches the CPU to protected mode, starts the file system, and then reads the contents of the Boot.ini file. This information determines the startup options and initial boot menu selections. |
| Fastfat.sys | System file driver for FAT16 and FAT32. |
| Ntoskrnl.exe | Extracts information about which system device drivers to load and the load order. |
| Kernel Mode | The processing mode that allows code to have direct access to all hardware and memory in the system. |
| User Mode | The processing mode in which applications run. |

**FAT Physical Structure**

The following information describes the physical structure of the FAT file system, and includes information about how clusters affect the FAT file system, as well as descriptions of the partition table.

**Clusters and Sectors on a FAT Volume**

A cluster (or allocation unit) is the smallest amount of disk space that can be allocated to hold a file. All file systems used by Windows Server 2003 organize hard disks based on cluster size, which is determined by the number of sectors (units of storage on a hard disk) that the cluster contains. For example, on a disk that uses 512-byte sectors, a 512-byte cluster contains one sector, whereas a 4-kilobyte (KB) cluster contains eight sectors.

Computers access certain sectors on a hard disk during startup to determine which operating system to start and where the partitions are located. The data stored on these sectors varies depending on the computer platform.

**Sequence of Clusters on a FAT Volume**

Clusters are numbered sequentially from the beginning of the partition into logical cluster numbers. Because FAT file system data clusters are located after the BIOS Parameter Blocks (BPB), reserved sectors, and two FAT structures, FAT formatting cannot guarantee that data clusters are aligned on a cluster boundary.

**Limitations of Cluster Sizes on a FAT Volume**

Because FAT16 and FAT32 each use different cluster sizes depending on the size of the volume, each file system has a maximum number of clusters it can support. The smaller the cluster size, the more efficiently a disk stores information because unused space within a cluster cannot be used by other files. And the more clusters a file system supports, the larger the volumes you can create and format by using a particular file system.

The size of each cluster is a multiple of the sector size. The numerals in the names, FAT16 and FAT32 refer to the number of bits required for a file allocation table entry as follows:

- FAT16 uses a 16-bit file allocation table entry (216 clusters).
- FAT32 uses a 32-bit file allocation table entry. However, FAT32 reserves the first 4 bits of a FAT32 file allocation table entry, which means FAT32 has a theoretical maximum of $2^{28}$ clusters.
- Floppy disks are always formatted as FAT.

The following table provides a comparison of FAT16 and FAT32 volume and default cluster sizes.

**Default FAT Cluster Sizes**

| Volume Size | FAT16 Cluster Size | FAT32 Cluster Size |
|---|---|---|
| 7 megabytes (MB)–16 MB | 2 KB | Not supported |
| 17 MB–32 MB | 512 bytes | Not supported |
| 33 MB–64 MB | 1 KB | 512 bytes |
| 65 MB–128 MB | 2 KB | 1 KB |
| 129 MB–256 MB | 4 KB | 2 KB |
| 257 MB–512 MB | 8 KB | 4 KB |
| 513 MB–1,024 MB | 16 KB | 4 KB |
| 1,025 MB–2 gigabytes (GB) | 32 KB | 4 KB |
| 2 GB–4 GB | 64 KB | 4 KB |
| 4 GB–8 GB | Not supported | 4 KB |
| 8 GB–16 GB | Not supported | 8 KB |
| 16 GB–32 GB | Not supported | 16 KB |
| 32 GB–2 terabytes | Not supported | Not supported |

Windows Server 2003 formats FAT32 volumes up to 32 GB regardless of cluster size. To format volumes larger than 32 GB, you must use NTFS. However, Windows Server 2003 can mount FAT32 volumes larger than 32 GB that were created by other operating systems.

**Maximum Sizes on FAT Volumes**

Before you format a volume, you can evaluate the types of files to be stored on the volume so that you can determine whether to use the default cluster size. If you format a volume but do not specify a cluster size, default values are used. If you want to change the cluster size after the volume is formatted, you must reformat the volume.

FAT16 and FAT32 have the following size limitations:

- FAT volumes smaller than 16 MB are formatted as FAT12.
- FAT16 volumes larger than 2 GB are not accessible from computers running MS-DOS, Windows 95, Windows 98, Windows Millennium Edition (Me), and many other operating systems. This limitation occurs because these operating systems do not support cluster sizes larger than 32 KB, which results in the 2 GB limit.
- In theory, FAT32 volumes can be about 8 terabytes; however, the maximum FAT32 volume size that Windows Server 2003 can format is 32 GB. Therefore, you must use NTFS to format volumes larger than 32 GB. However, Windows Server 2003 can read and write to larger FAT32 volumes formatted by other operating systems.

The largest possible file for a FAT32 volume is 4 GB minus 1 byte. FAT32 contains 4 bytes per cluster in the file allocation table; FAT16 contains 2 bytes per cluster; and FAT12 contains 1.5 bytes per cluster. A FAT32 volume must have at least 65,527 clusters

FAT16 supports a maximum of 65,524 clusters per volume. The following table lists FAT16 size limits.

**FAT16 Size Limits**

| Description | Limit |
|---|---|
| Maximum file size | Tested: 4 GB minus 1 byte ($2^{32}$ bytes minus 1 byte) |
| Maximum volume size | Tested: 4 GB |

| | |
|---|---|
| Files per volume | Approximately 65,536 ($2^{16}$ files) |
| Maximum number of files and folders within the root folder | 512 (Long file names can reduce the number of available files and folders in the root folder.) |

A FAT32 volume must have a minimum of 65,527 clusters. Windows Server 2003 can format FAT32 volumes up to 32 GB, but it can mount larger FAT32 volumes created by other operating systems. The following table lists FAT32 size limits.

**FAT32 Size Limits**

| Description | Limit |
|---|---|
| Maximum file size | Tested: 4 GB minus 1 byte ($2^{32}$ bytes minus 1 byte) |
| Maximum volume size | Tested: 32 GB (implementation) |
| Files per volume | 4,177,920 |
| Maximum number of files and subfolders within a single folder | 65,534 (The use of long file names can significantly reduce the number of available files and subfolders within a folder.) |

**Organization of a FAT Volume**

A volume formatted with FAT is organized as illustrated in the following figure.

**Organization of a FAT Volume**




The following table describes each of the organizational structures on a FAT volume.

**FAT Volume Components**

| Component | Description |
|---|---|
| Boot Sector | Contains the BIOS parameter block that stores information about the layout of the volume and the file system structures, as well as the boot code that loads Windows Server 2003. |
| Reserved Sectors | The number of sectors that precede the start of the first FAT, including the boot sector. |
| FAT 1 | Original FAT. |
| FAT 2 (Duplicate) | Backup copy of the FAT. |
| Root folder | Describes the files and folders in the root of the partition. |
| Other folders and all files | Contains the data for the files and folders within the file system. |

**Boot Sectors on MBR and GPT Disks**

On MBR disks (disks that contain a master boot record), the boot sector, which is located at the first logical sector of each partition, is a critical disk structure for starting your computer. It contains executable code and the data required by the code, including information that the file system uses to access the volume. The boot sector is created when you format a volume. At the end of the boot sector is a 2-byte structure called a signature word or end of sector marker, which is always set to 0x55AA. On computers running Windows Server 2003, the boot sector on the active partition loads into memory and starts Ntldr, which loads the boot menu if multiple versions of Windows are installed or loads the operating system if only one operating system is installed.

GUID partition table (GPT) disks are similar to MBR disks, except they use primary and backup partition structures to provide redundancy. These structures are located at the beginning and the end of the disk. GPT identifies these structures by their logical block address (LBA) rather than by their relative sectors.

A boot sector consists of the following elements:

- An x86-based CPU jump instruction.
- The original equipment manufacturer identification (OEM ID).
- The BIOS parameter block (BPB), a data structure.
- The extended BPB.
- The executable boot code (or bootstrap code) that starts the operating system.

All Windows Server 2003 boot sectors contain the preceding elements regardless of the type of disk (basic disk or dynamic disk).

**Components of a Boot Sector**
The MBR transfers CPU execution to the boot sector, so the first three bytes of the boot sector must be valid, executable x86-based CPU instructions. This includes a jump instruction that skips the next several nonexecutable bytes.

Following the jump instruction is the 8-byte OEM ID, a string of characters that identifies the name and version number of the operating system that formatted the volume. To preserve compatibility with MS-DOS, Windows Server 2003 records “MSDOS5.0” in this field on FAT16 and FAT32 disks.

### Note

- You might also see the OEM ID "MSWIN4.0" on disks formatted by Windows 95 and "MSWIN4.1" on disks formatted by Windows 95 OEM Service Release 2 (OSR2), Windows 98, and Windows Me. Windows Server 2003 does not use the OEM ID field in the boot sector except for verifying NTFS volumes.

Following the OEM ID is the BPB, which provides information that enables the executable boot code to locate Ntldr. The BPB always starts at the same offset, so standard parameters are in a known location. Disk size and geometry variables are encapsulated in the BPB. Because the first part of the boot sector is an x86 jump instruction, the BPB can be extended in the future by appending new information at the end. The jump instruction needs only a minor adjustment to accommodate this change. The BPB is stored in a packed (unaligned) format.

**FAT16 Boot Sector**
The following table describes the sections of a boot sector of a volume formatted with the FAT16 file system.

### Boot Sector Sections on a FAT16 Volume

| Byte Offset | Field Length | Field Name |
|---|---|---|
| 0x00 | 3 bytes | Jump instruction |
| 0x03 | 8 bytes | OEM ID |
| 0x0B | 25 bytes | BPB |
| 0x24 | 26 bytes | Extended BPB |
| 0x3E | 448 bytes | Bootstrap code |
| 0x01FE | 2 bytes | End of sector marker |

The following example illustrates a hexadecimal printout of the boot sector on a FAT16 volume. The printout is formatted in three sections:

- Bytes 0x00– 0x0A are the jump instruction and the OEM ID (shown in bold print).
- Bytes 0x0B– 0x3D are the BPB and the extended BPB.
- The remaining section is the bootstrap code and the end of sector marker (shown in bold print).

```
Physical Sector: Cyl 0, Side 1, Sector 1
00000000: EB 3C 90 4D 53 44 4F 53  - 35 2E 30 00 02 40 01 00   .<.MSDOS5.0..@..
00000010: 02 00 02 00 00 F8 FC 00  - 3F 00 40 00 3F 00 00 00   ........?.@.?...
00000020: 01 F0 3E 00 80 00 29 A8  - 8B 36 52 4E 4F 20 4E 41   ..>...)..6RNO NA
00000030: 4D 45 20 20 20 20 46 41  - 54 31 36 20 20 20 33 C0   ME    FAT16   3.
00000040: 8E D0 BC 00 7C 68 C0 07  - 1F A0 10 00 F7 26 16 00   ....|h......&..
00000050: 03 06 0E 00 50 91 B8 20  - 00 F7 26 11 00 8B 1E 0B   ....P.. ..&.....
00000060: 00 03 C3 48 F7 F3 03 C8  - 89 0E 08 02 68 00 10 07   ...H........h...
00000070: 33 DB 8F 06 13 02 89 1E  - 15 02 0E E8 90 00 72 57   3.............rW
00000080: 33 DB 8B 0E 11 00 8B FB  - 51 B9 0B 00 BE DC 01 F3   3.......Q.......
00000090: A6 59 74 05 83 C3 20 E2  - ED E3 37 26 8B 57 1A 52   .Yt... ...7&.W.R
000000A0: B8 01 00 68 00 20 07 33  - DB 0E E8 48 00 72 28 5B   ...h. .3...H.r([
000000B0: 8D 36 0B 00 8D 3E 0B 02  - 1E 8F 45 02 C7 05 F5 00   .6...>....E.....
000000C0: 1E 8F 45 06 C7 45 04 0E  - 01 8A 16 24 00 EA 03 00   ..E..E.....$....
000000D0: 00 20 BE 86 01 EB 03 BE  - A2 01 E8 09 00 BE C1 01   . ..............
000000E0: E8 03 00 FB EB FE AC 0A  - C0 74 09 B4 0E BB 07 00   .........t......
000000F0: CD 10 EB F2 C3 50 4A 4A  - A0 0D 00 32 E4 F7 E2 03   .....PJJ...2....
00000100: 06 08 02 83 D2 00 A3 13  - 02 89 16 15 02 58 A2 07   .............X..
00000110: 02 A1 13 02 8B 16 15 02  - 03 06 1C 00 13 16 1E 00   ................
00000120: F7 36 18 00 FE C2 88 16  - 06 02 33 D2 F7 36 1A 00   .6........3..6..
00000130: 88 16 25 00 A3 04 02 A1  - 18 00 2A 06 06 02 40 3A   ..%.......*...@:
00000140: 06 07 02 76 05 A0 07 02  - 32 E4 50 B4 02 8B 0E 04   ...v....2.P.....
00000150: 02 C0 E5 06 0A 2E 06 02  - 86 E9 8B 16 24 00 CD 13   ............$...
00000160: 0F 83 05 00 83 C4 02 F9  - CB 58 28 06 07 02 76 11   .........X(...v.
00000170: 01 06 13 02 83 16 15 02  - 00 F7 26 0B 00 03 D8 EB   ..........&.....
00000180: 90 A2 07 02 F8 CB 42 4F  - 4F 54 3A 20 43 6F 75 6C   ......BOOT: Coul
00000190: 64 6E 27 74 20 66 69 6E  - 64 20 4E 54 4C 44 52 0D   dn't find NTLDR.
000001A0: 0A 00 42 4F 4F 54 3A 20  - 49 2F 4F 20 65 72 72 6F   ..BOOT: I/O erro
000001B0: 72 20 72 65 61 64 69 6E  - 67 20 64 69 73 6B 0D 0A   r reading disk..
000001C0: 00 50 6C 65 61 73 65 20  - 69 6E 73 65 72 74 20 61   .Please insert a
000001D0: 6E 6F 74 68 65 72 20 64  - 69 73 6B 00 4E 54 4C 44   nother disk.NTLD
000001E0: 52 20 20 20 20 20 20 00  - 00 00 00 00 00 00 00 00   R       ........
000001F0: 00 00 00 00 00 00 00 00  - 00 00 00 00 00 00 55 AA   ..............U.
```

The following two tables illustrate the layout of the BPB and the extended BPB for FAT16 volumes. The sample values correspond to the data in this example of a hexadecimal printout of the boot sector on a FAT16 volume.

**BPB Fields for FAT16 Volumes**

| Byte Offset | Field Length | Sample Value | Field Name and Definition |
|---|---|---|---|
| 0x0B | 2 bytes | 00 02 | **Bytes Per Sector**.The size of a hardware sector. Valid decimal values for this field are 512, 1024, 2048, and 4096. |
| 0x0D | 1 byte | 40 | **Sectors Per Cluster**.The number of sectors in a cluster. Because FAT16 can track only a limited number of clusters (up to 65,524), FAT16 supports large volumes by increasing the number of sectors per cluster. The default cluster size for a volume depends on the volume size. Valid decimal values for this field are 1, 2, 4, 8, 16, 32, 64, and 128. |
| 0x0E | 2 bytes | 01 00 | **Reserved Sectors**. The number of sectors that precede the start of the first FAT, including the boot sector. The value of this field is typically 1. |
| 0x10 | 1 byte | 02 | **Number of FATs**. The number of copies of the FAT on the volume. The value of this field is typically 2. |
| 0x11 | 2 bytes | 00 02 | **Root Entries**. The total number of 32-byte file and folder name entries that can be stored in the root folder of the volume. On a typical hard disk, the value of this field is 512. One entry can be used as a volume label, and files and folders with long names use multiple entries per file. The largest number of file and folder entries is typically 511; however, if long file names are used, entries usually run out before you reach that number. |
| 0x13 | 2 bytes | 00 00 | **Small Sectors**. The number of sectors on the volume represented in 16 bits (< 65,536). For volumes larger than 65,536 sectors, this field has a value of zero and the **Large Sectors** field is used instead. |

| | | | |
|---|---|---|---|
| 0x15 | 1 byte | F8 | **Media Descriptor**. Provides information about the media being used. A value of 0xF8 indicates a hard disk and 0xF0 indicates a high-density 3.5-inch floppy disk. Media descriptor entries are a legacy of MS-DOS FAT16 and FAT12 disks and are not used in Windows Server 2003. |
| 0x16 | 2 bytes | FC 00 | **Sectors Per FAT**. The number of sectors occupied by each FAT on the volume. The computer uses this number and the number of FATs and reserved sectors to determine where the root directory begins. The computer can also determine where the user data area of the volume begins based on the number of entries in the root directory (512). |
| 0x18 | 2 bytes | 3F 00 | **Sectors Per Track**. Part of the apparent disk geometry used on a low-level formatted disk. |
| 0x1A | 2 bytes | 40 00 | **Number of Heads**. Part of the apparent disk geometry used on a low-level formatted disk. |
| 0x1C | 4 bytes | 3F 00 00 00 | **Hidden Sectors**. The number of sectors on the volume before the boot sector. This value is used during the boot sequence to calculate the absolute offset to the root directory and data areas. |
| 0x20 | 4 bytes | 01 F0 3E 00 | **Large Sectors**. If the value of the **Small Sectors** field is zero, this field contains the total number of sectors in the FAT16 volume. If the value of the **Small Sectors** field is not zero, the value of this field is zero. |

**Extended BPB Fields for FAT16 Volumes**

| Byte Offset | Field Length | Sample Value | Field Name and Definition |
|---|---|---|---|
| 0x24 | 1 byte | 80 | **Physical Drive Number**. Related to the BIOS physical drive number. Floppy drives are identified as 0x00 and physical hard disks are identified as 0x80, regardless of the number of physical disk drives. Typically, this value is set prior to issuing an INT 13h BIOS call to specify the device to access. This value is only relevant if the device is a boot device. |
| 0x25 | 1 byte | 00 | **Reserved**. FAT16 volumes are always set to zero. |
| 0x26 | 1 byte | 29 | **Extended Boot Signature**. A field that must have the value 0x28 or 0x29 to be recognized by Windows Server 2003. |
| 0x27 | 4 bytes | A8 8B 36 52 | **Volume Serial Number**. A random serial number that is created when a volume is formatted and that helps to distinguish between disks. |
| 0x2B | 11 bytes | NO NAME | **Volume Label**. A field that was once used to store the volume label. The volume label is now stored as a special file in the root directory. |
| 0x36 | 8 bytes | FAT16 | Not used by Windows Server 2003. |

**FAT32 Boot Sector**

The FAT32 boot sector is structurally very similar to the FAT16 boot sector, but the FAT32 BPBcontains additional fields. The FAT32 extended BPB uses the same fields as FAT16, but the offset addresses of these fields within the boot sector are different than those found in FAT16 boot sectors. Volumes formatted in FAT32 are not readable by operating systems that are incompatible with FAT32.

The following table describes the sections of a boot sector of a volume formatted with the FAT32 file system.

**Boot Sector Sections on a FAT32 Volume**

| Byte Offset | Field Length | Field Name |
|---|---|---|
| 0x00 | 3 bytes | Jump instruction |
| 0x03 | 8 bytes | OEM ID |
| 0x0B | 53 bytes | BPB |
| 0x40 | 26 bytes | Extended BPB |
| 0x5A | 420 bytes | Bootstrap code |

| | | |
|---|---|---|
| 0x01FE | 2 bytes | End of sector marker |

The following example illustrates a hexadecimal printout of the boot sector on a FAT32 volume. The printout is formatted in three sections:

- Bytes 0x00– 0x0A are the jump instruction and the OEM ID (shown in bold print).
- Bytes 0x0B– 0x59 are the BPB and the extended BPB.
- The remaining section is the bootstrap code and the end of sector marker (shown in bold print).

```
Physical Sector: Cyl 878, Side 0, Sector 1
00000000: EB 58 90 4D 53 44 4F 53  - 35 2E 30 00 02 10 24 00   .X.MSDOS5.0...$.
00000010: 02 00 00 00 00 F8 00 00  - 3F 00 FF 00 3F 00 00 00   ........?...?...
00000020: 1D 91 11 01 2A 22 00 00  - 00 00 00 00 02 00 00 00   ....*"..........
00000030: 01 00 06 00 00 00 00 00  - 00 00 00 00 00 00 00 00   ................
00000040: 80 00 29 F1 9E 5E 5E 4E  - 4F 20 4E 41 4D 45 20 20   ..)..^^NO NAME
00000050: 20 20 46 41 54 33 32 20  - 20 20 33 C9 8E D1 BC F4     FAT32   3.....
00000060: 7B 8E C1 8E D9 BD 00 7C  - 88 4E 02 8A 56 40 B4 08   {......|.N..V@..
00000070: CD 13 73 05 B9 FF FF 8A  - F1 66 0F B6 C6 40 66 0F   ..s......f...@f.
00000080: B6 D1 80 E2 3F F7 E2 86  - CD C0 ED 06 41 66 0F B7   ....?.......Af..
00000090: C9 66 F7 E1 66 89 46 F8  - 83 7E 16 00 75 38 83 7E   .f..f.F..~..u8.~
000000A0: 2A 00 77 32 66 8B 46 1C  - 66 83 C0 0C BB 00 80 B9   *.w2f.F.f.......
000000B0: 01 00 E8 2B 00 E9 48 03  - A0 FA 7D B4 7D 8B F0 AC   ...+..H...}.}...
000000C0: 84 C0 74 17 3C FF 74 09  - B4 0E BB 07 00 CD 10 EB   ..t.<.t.........
000000D0: EE A0 FB 7D EB E5 A0 F9  - 7D EB E0 98 CD 16 CD 19   ...}....}.......
000000E0: 66 60 66 3B 46 F8 0F 82  - 4A 00 66 6A 00 66 50 06   f`f;F...J.fj.fP.
000000F0: 53 66 68 10 00 01 00 80  - 7E 02 00 0F 85 20 00 B4   Sfh.....~.... ..
00000100: 41 BB AA 55 8A 56 40 CD  - 13 0F 82 1C 00 81 FB 55   A..U.V@........U
00000110: AA 0F 85 14 00 F6 C1 01  - 0F 84 0D 00 FE 46 02 B4   .............F..
00000120: 42 8A 56 40 8B F4 CD 13  - B0 F9 66 58 66 58 66 58   B.V@......fXfXfX
00000130: 66 58 EB 2A 66 33 D2 66  - 0F B7 4E 18 66 F7 F1 FE   fX.*f3.f..N.f...
00000140: C2 8A CA 66 8B D0 66 C1  - EA 10 F7 76 1A 86 D6 8A   ...f..f....v....
00000150: 56 40 8A E8 C0 E4 06 0A  - CC B8 01 02 CD 13 66 61   V@............fa
00000160: 0F 82 54 FF 81 C3 00 02  - 66 40 49 0F 85 71 FF C3   ..T.....f@I..q..
00000170: 4E 54 4C 44 52 20 20 20  - 20 20 20 00 00 00 00 00   NTLDR      .....
00000180: 00 00 00 00 00 00 00 00  - 00 00 00 00 00 00 00 00   ................
00000190: 00 00 00 00 00 00 00 00  - 00 00 00 00 00 00 00 00   ................
000001A0: 00 00 00 00 00 00 00 00  - 00 00 00 00 0D 0A 43 61   ..............Ca
000001B0: 6E 6E 6F 74 20 73 74 61  - 72 74 2E 20 20 52 65 6D   nnot start.  Rem
000001C0: 6F 76 65 20 6D 65 64 69  - 61 2E FF 0D 0A 44 69 73   ove media....Dis
000001D0: 6B 20 65 72 72 6F 72 FF  - 0D 0A 50 72 65 73 73 20   k error...Press 
000001E0: 61 6E 79 20 6B 65 79 20  - 74 6F 20 72 65 73 74 61   any key to resta
000001F0: 72 74 0D 0A 00 00 00 00  - 00 AC CB D8 00 00 55 AA   rt............U.
```

The following two tables illustrate the layout of the BPB and the extended BPB for FAT32 volumes. The sample values correspond to the data in this example of a hexadecimal printout of the boot sector on a FAT32 volume.

**BPB Fields for FAT32 Volumes**

| Byte Offset | Field Length | Sample Value | Field Name and Definition |
|---|---|---|---|
| 0x0B | 2 bytes | 00 02 | **Bytes Per Sector**. The size of a hardware sector. Valid decimal values for this field are 512, 1024, 2048, and 4096. |
| 0x0D | 1 byte | 10 | **Sectors Per Cluster**.The number of sectors in a cluster. The default cluster size for a volume depends on the volume size. Valid decimal values for this field are 1, 2, 4, 8, 16, 32, 64, and 128. The Windows Server 2003 implementation of FAT32 allows for creation of volumes up to a maximum of 32 GB. However, larger volumes created by other operating systems (Windows 95 OSR2 and later) are accessible in Windows Server 2003. |
| 0x0E | 2 bytes | 24 00 | **Reserved Sectors**. The number of sectors that precede the start of the first FAT, including the boot sector. |

| | | | |
|---|---|---|---|
| 0x10 | 1 byte | 02 | **Number of FATs**. The number of copies of the FAT on the volume. The value of this field is always 2. |
| 0x11 | 2 bytes | 00 00 | **Root Entries (FAT12/FAT16 only)**. For FAT32 volumes, this field must be set to zero. |
| 0x13 | 2 bytes | 00 00 | **Small Sectors (FAT12/FAT16 only)**. For FAT32 volumes, this field must be set to zero. |
| 0x15 | 1 byte | F8 | **Media Descriptor**. Provides information about the media being used. A value of 0xF8 indicates a hard disk and 0xF0 indicates a high-density 3.5-inch floppy disk. Media descriptor entries are a legacy of MS-DOS FAT16 disks and are not used in Windows Server 2003. |
| 0x16 | 2 bytes | 00 00 | **Sectors Per FAT (FAT12/FAT16 only)**. For FAT32 volumes, this field must be set to zero. |
| 0x18 | 2 bytes | 3F 00 | **Sectors Per Track**. Contains the "sectors per track" geometry value for disks that use INT 13h. The volume is broken down into tracks by multiple heads and cylinders. |
| 0x1A | 2 bytes | FF 00 | **Number of Heads**. Contains the "count of heads" geometry value for disks that use INT 13h. For example, on a 1.44-MB, 3.5-inch floppy disk this value is 2. |
| 0x1C | 4 bytes | 3F 00 00 00 | **Hidden Sectors**. The number of sectors on the volume before the boot sector. This value is used during the boot sequence to calculate the absolute offset to the root directory and data areas. This field is generally only relevant for media that are visible on interrupt 13h. It must always be zero on media that are not partitioned. |
| 0x20 | 4 bytes | 1D 91 11 01 | **Large Sectors**. Contains the total number of sectors in the FAT32 volume. |
| 0x24 | 4 bytes | 2A 22 00 00 | **Sectors Per FAT (FAT32 only)**. The number of sectors occupied by each FAT on the volume. The computer uses this number and the number of FATs and reserved sectors (described in this table) to determine where the root directory begins. The computer can also determine where the user data area of the volume begins based on the number of entries in the root directory. |
| 0x28 | 2 bytes | 00 00 | Not used by Windows Server 2003. |
| 0x2A | 2 bytes | 00 00 | **File System Version (FAT32 only)**. The high byte is the major revision number; the low byte is the minor revision number. This field supports the ability to extend the FAT32 media type in the future with concern for old FAT32 drivers mounting the volume. Both bytes are zero in Windows Server 2003, Windows 2000, and Windows Me and earlier. |
| 0x2C | 4 bytes | 02 00 00 00 | **Root Cluster Number (FAT32 only)**. The cluster number of the first cluster of the root directory. This value is typically, but not always, 2. |
| 0x30 | 2 bytes | 01 00 | **File System Information Sector Number (FAT32 only)**. The sector number of the File System Information (FSINFO) structure in the reserved area of the FAT32 volume. The value is typically 1. A copy of the FSINFO structure is kept in the Backup Boot Sector, but it is not kept up-to-date. |
| 0x32 | 2 bytes | 06 00 | **Backup Boot Sector (FAT32 only)**. A value other than zero specifies the sector number in the reserved area of the volume where a copy of the boot sector is stored. The value of this field is typically 6. No other value is recommended. |
| 0x34 | 12 bytes | 00 00 00 00 00 00 00 00 00 00 00 00 | **Reserved (FAT32 only)**. Reserved space for future expansion. The value of this field must always be zero. |

**Extended BPB Fields for FAT32 Volumes**

| Byte Offset | Field Length | Sample Value | Field Name and Definition |
|---|---|---|---|

| | | | |
|---|---|---|---|
| 0x40 | 1 byte | 80 | **Physical Drive Number**. Related to the BIOS physical drive number. Floppy drives are identified as 0x00 and physical hard disks are identified as 0x80, regardless of the number of physical disk drives. Typically, this value is set prior to issuing an INT 13h BIOS call to specify the device to access. This value is only relevant if the device is a boot device. |
| 0x41 | 1 byte | 00 | **Reserved**. FAT32 volumes are always set to zero. |
| 0x42 | 1 byte | 29 | **Extended Boot Signature**. A field that must have the value 0x28 or 0x29 to be recognized by Windows Server 2003. |
| 0x43 | 4 bytes | F1 9E 5E 5E | **Volume Serial Number**. A random serial number that is created when a volume is formatted and that helps to distinguish between disks. |
| 0x47 | 11 bytes | NO NAME | **Volume Label**. A field that was once used to store the volume label. The volume label is now stored as a special file in the root directory. |
| 0x52 | 8 bytes | FAT32 | **System ID**. A text field with a value of FAT32. |

**FAT 1 and FAT 2**

The file allocation table, FAT 1, identifies each cluster in the volume as one of the following:

- Unused
- Cluster in use by a file
- Bad cluster
- Last cluster in a file

Because the file allocation table is so important, a copy of the table, FAT 2, is maintained on the volume. This allows the file system driver or consistency-checking program to access the duplicate copy if it cannot access FAT 1.

**FAT Root Folder**

The root folder describes the files and folders in the root of the partition. It has a maximum number of available entries fixed at 512. The maximum number of entries on a floppy disk depends on the size of the disk.

**FAT Root Folder Structure**

| Root Folder Entry | Size | Description |
|---|---|---|
| Name | 11 bytes | Name in 8.3 format. |
| Attribute byte | 1 byte | Information about the entry. |
| Creation time and date | 5 bytes | Time and date file was created. |
| Last access date | 2 bytes | Date file was last accessed. |
| Last modification time and date | 4 bytes | Time and date file was last modified. |
| First cluster | 2 bytes | Starting cluster number in the file allocation table. |
| File size | 4 bytes | Size of the file. |

**Note**

- Three bytes in each 32-byte folder entry are held in reserve.

The information stored in the root folder is used by all operating systems that support FAT. Windows Server 2003 can store additional timestamps in a FAT folder entry. These timestamps show when the file was created or last accessed.

The attribute byte for each entry in a folder describes what kind of entry it is. For example, one bit indicates that the entry is for a subfolder, and another bit marks the entry as a volume. Typically, the operating system controls the settings of these bits.

The attribute byte includes four bits that can be turned on or off by the user — archive, system, hidden, and read-only.

**FAT Processes and Interactions**
FAT is affected by many factors, such as cluster size, fragmentation level, and the use of programs such as antivirus software.

**Mounting a FAT Volume**
When mounting a FAT volume, the MBR executes code to start up the boot sector. The boot sector then executes additional code to mount the volume.

**Master Boot Code Startup Process**
The MBR contains a small amount of executable code called the master boot code, the disk signature, and the partition table for the disk. During startup, the master boot code performs the following activities:

1. Scans the partition table for the active partition.
2. Finds the starting sector of the active partition.
3. Loads a copy of the boot sector from the active partition into memory.
4. Transfers control to the executable code in the boot sector.

**Boot Sector Startup Process**
Computers use the boot sector to run instructions during startup. The initial startup process is summarized in the following steps:

1. The system BIOS and the CPU initiate the power-on self test (POST).
2. The BIOS finds the boot device, which is typically the first disk the BIOS finds, unless the controller is configured to boot from a different disk.
3. The BIOS loads the first physical sector of the boot device into memory and transfers CPU execution to that memory address.

If the boot device is on a hard disk, the BIOS loads the MBR. The master boot code in the MBR loads the boot sector of the active partition, and transfers CPU execution to that memory address. On computers that are running Windows Server 2003, the executable boot code in the boot sector finds Ntldr, loads it into memory, and transfers execution to that file.

### Note

- Windows Server 2003 cannot start up from a spanned, striped, or RAID-5 volume on dynamic disks. These disk structures cannot be registered into the MBR partition table; therefore, a system volume that uses these structures cannot start.

If drive A contains a floppy disk, the system BIOS loads the first sector (the boot sector) of the disk into memory. If the disk is a startup disk (formatted by MS-DOS with core operating system files applied), the boot sector loads into memory and uses the executable boot code to transfer CPU execution to Io.sys, a core MS-DOS operating system file. If the floppy disk is not a startup disk, the executable boot code displays an error message.

### Note

- These messages do not appear on normally functioning systems that are configured to look for the startup files on drive C first. On many computers, an option in the CMOS setup program allows the user to set the sequence of installed disks that the system searches to find the startup files.

If you get similar errors when trying to start the computer from the hard disk, the boot sector might be corrupted.

Initially, the startup process is independent of disk format and operating system. The unique characteristics of operating and file systems become important when the boot sector's executable boot code starts.

**File Processing on FAT Clusters**

When a file is saved to a FAT-formatted volume, it stores file information in clusters on the hard disk. If a file requires space greater than the cluster's size, FAT continues to store file information in the next available cluster until all information about the file is stored. The following figure shows an example of how FAT stores and retrieves file information from clusters.

**File Processing on a FAT Volume**




In this figure, 10 clusters (0 through 9) contain 3 files. One file occupies clusters 2, 3, 6, and 8; a second file occupies clusters 4 and 5; and a third file occupies cluster 7. The starting cluster number is the address of the first cluster used by the file. Each cluster contains a pointer to the next cluster in the file, or an end-of-file indicator (0xFFFF), which indicates that this cluster is the end of the file.

The attribute byte for each entry in a folder describes what kind of entry it is. For example, one bit indicates that the entry is for a subfolder, and another bit marks the entry as a volume. Typically, the operating system controls the settings of these bits.

The attribute byte includes four bits that can be turned on or off by the user — archive, system, hidden, and read-only.

**FAT16 File System**

FAT16 works best on small disks with simple folder structures. It is included in Windows Server 2003 for the following reasons:

- It provides backward compatibility in the form of an upgrade path for earlier versions of Windows-compatible products.
- It is compatible with most other operating systems.

Because FAT16 uses a larger cluster size, a FAT16 volume does not efficiently manage disk space when relatively small files are placed on it. A small file still uses a full cluster, even if the file does not require all the space allocated by the cluster. The cluster, once used by a file, cannot be used by any other file. This results in wasted space. Therefore, FAT16 is not recommended for volumes larger than 511 MB, and you cannot use FAT16 on volumes larger than 4 gigabytes(GB).

In the file allocation table of a FAT16 volume, files are given the first available location on the volume.

**FAT32 File System**

Because FAT32 requires 4 bytes to store cluster values, many internal and on-disk data structures have been revised or expanded. Most programs are unaffected by these changes; however, disk tools that read the on-disk format must be updated to support FAT32.

The most significant difference between FAT16 and FAT32 is the maximum number of clusters supported, which in turn affects a volume's maximum size and storage efficiency. FAT32 breaks the 4-GB volume limitation of FAT16 by extending the maximum number of clusters, and therefore works well on large disks with complex folder structures.

Due to the greater number of available clusters within FAT32, each cluster can be made smaller for a particular volume, increasing the efficiency of data storage. For example, FAT16 volumes between 2 and 4 GB use a 64-KB cluster, whereas FAT32 volumes between 16 GB and 32 GB use a 16-KB cluster.

The largest possible file for a FAT32 volume is 4 GB minus 1 byte. FAT32 contains 4 bytes per cluster in the file allocation table; FAT16 contains 2 bytes per cluster; and FAT12 contains 1.5 bytes per cluster. A FAT32 volume must have at least 65,527 clusters.

**Formatting Volumes**

During volume format, Windows Server 2003 places key file system structures on the volume, including the boot sector as well as replacing Ntldr. A check of the integrity of all sectors on the volume is performed, and you have the opportunity to change the cluster size used on the volume. If a volume is formatted using Quick format, the file system structure on the volume is created, but the integrity of every sector in the volume is not checked.

The version of FAT that Setup uses depends on the size of the volume. For volumes smaller than 2 GB (2048 MB), Setup uses FAT16. For volumes 2 GB or larger, Setup uses FAT32. For volumes 32 GB or larger, Setup uses NTFS and does not offer FAT.

Disk Management requires volumes on dynamic disks and GPT disks to be formatted as NTFS.

**Converting Volumes**

Windows Server 2003 can convert FAT16 or FAT32 to NTFS. This type of conversion is a one-way process. After you convert a volume to NTFS, you cannot reconvert the volume to FAT without backing up your data, reformatting the volume as FAT, and then restoring your data.

During the formatting process, Windows Server 2003 aligns FAT data clusters at the cluster size boundary. Convert.exe can preserve the cluster size for the size of the volume (up to 4 KB) instead of using the 512-byte cluster size used in Windows 2000 for converted volumes. The following table lists cluster sizes for volumes converted to NTFS.

**Cluster Sizes for Volumes Converted to NTFS**

| Original FAT Cluster Size | Converted NTFS Cluster Size |
|---|---|
| 512 bytes | 512 bytes |
| 1 KB | 1 KB |
| 2 KB | 2 KB |
| 4 KB and larger | 4 KB |

**Lack of Support for Compression**

Unlike NTFS, FAT does not support compression of files. If you move or copy a file from a folder on a FAT volume to a folder on an NTFS volume, the file takes on the attributes of the folder on the NTFS volume. Compressed files copied or moved from an NTFS volume to a FAT volume automatically uncompress, because FAT does not support compression.

**How Cluster Fragmentation Affects Performance**

Lack of disk space can become an issue on a volume, especially when multiple users access the volume to create, open, modify, or delete files.

Fragmentation occurs when the virtual clusters of a file are on non-contiguous logical clusters. Because the file is spread out across the disk, access is slower, so fragmentation affects the performance of the file system. As users create new files or change the size of existing ones, the file system uses additional free clusters. If the file system cannot use contiguous clusters, it will find any available cluster on the disk. The closer that clusters for each file are located on the hard disk, the more efficiently the disk can retrieve the information. Disk defragmentation is a process that groups all the virtual clusters of a file into contiguous logical clusters, and is usually done in order to improved disk drive performance.

**File Naming**

Windows Server 2003 supports both long and short file names on FAT volumes.

**File Names on FAT Volumes**

Files created or renamed on FAT volumes use attribute bits to support long file names in a way that does not interfere with how MS-DOS gains access to the volume.

When you create a file that has a long file name, Windows Server 2003 creates a conventional 8.3 name for the file and one or more secondary folder entries for the file, one for each set of 13 characters in the long file name. Each secondary folder entry stores a corresponding part of the long file name in Unicode. MS-DOS accesses the file by using the conventional 8.3 file name contained in the folder entry for the file.

Windows Server 2003 marks the secondary folder entries as part of a long file name by setting the volume ID, read-only, system, and hidden attribute bits. MS-DOS typically ignores folder entries with all these attribute bits set.

The following figure shows all of the folder entries for the file Thequi~1.fox, which has a long name, "The quick brown.fox." The long name is in Unicode and each character in the name uses 2 bytes in the folder entry. The attribute field for the long-name entries has the value 0x0F. The attribute field for the short name has the value 0x20.

**Long File Name on a FAT Volume**




**Note**

- Windows NT and Windows Server 2003 do not use the same algorithm to create long and short file names as Windows 95, Windows 98, and Windows Me. However, on computers that use a multiple-boot configuration to start these operating systems, files that you create while running one operating system can be accessed while running another.

By default, Windows Server 2003 supports long file names on FAT volumes. You can disable long file names on FAT volumes if you use MS-DOS–based disk tools regularly on the computer. These tools might either eliminate the long file names created by Windows Server 2003 or delete the files that have long file names.

**Note**

- It is advised that you do not disable long file names on FAT volumes.

**How Windows Server 2003 Creates File Names**

File names in Windows Server 2003 can be up to 255 characters and can contain spaces, multiple periods, and special characters that are not allowed in MS-DOS file names. Windows Server 2003 makes it possible for other operating systems to access files that have long names by generating an MS-DOS-readable (8.3) name for each file. These MS-DOS-readable names also enable MS-DOS-based and Windows 3.*x*–based applications to recognize and load files that have long file names. When a program saves a file on a computer running Windows Server 2003, both the 8.3 file name and long file name are retained.

**Note**

- The 8.3 format means that files can have between 1 and 8 characters in the file name. The name must start with a letter or a number and can contain any characters except the following:

- . " / \ [ ] : ; | = , * ? (space)

- An 8.3 file name typically has a file name extension that is from one to three characters long and has the same character restrictions. A period separates the file name from the file name extension.

- Several special file names are reserved by the system and cannot be used for files or folders: CON, AUX, COM1, COM2, COM3, COM4, LPT1, LPT2, LPT3, PRN, and NUL.

When you create a file that has a long file name, Windows Server 2003 creates a conventional 8.3 name for the file and one or more secondary folder entries for the file, one for each set of 13 characters in the long file name. Each secondary folder entry stores a corresponding part of the long file name in Unicode. MS-DOS accesses the file by using the conventional 8.3 file name contained in the folder entry for the file.

Windows Server 2003 marks the secondary folder entries as part of a long file name by setting the volume ID, read-only, system, and hidden attribute bits. MS-DOS typically ignores folder entries with all these attribute bits set.

**Generating Short File Names**

In Windows Server 2003, FAT uses the Unicode character set, which contains several prohibited characters that MS-DOS cannot read, for their names. To generate a short MS-DOS-readable file name, Windows Server 2003 deletes all of these

characters from the long file name and removes any spaces, but ensures that the generated short file name is unique. Because an MS-DOS-readable file name can have only one period, Windows Server 2003 also removes extra periods from the file name.

**Related Information**

- Basic Disks and Volumes Technical Reference
- NTFS Technical Reference

# FAT Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- FAT Tools
- FAT Registry Entries
- FAT WMI Classes


The following tools, registry settings, and Windows Management Instrumentation (WMI) classes are associated with FAT volumes.

**FAT Tools**
The following tools are associated with FAT. The tools in this list are used to create and maintain FAT volumes.

**Chkdsk.exe: Check Disk**
**Category**
Operating system tool

**Version compatibility**
Windows Server 2003. Chkdsk.exe works on FAT and NTFS volumes.

Chkdsk is a command-line tool that verifies the logical integrity of a file system on a Windows Server 2003 volume. Setting parameters on the command line will list and fix any errors.

For more information about Chkdsk, see Command Line References in the Tools and Settings Collection.

**Chkntfs.exe: Chkntfs**
**Category**
Operating system tool

**Version compatibility**
Windows Server 2003. Chkntfs.exe works on FAT and NTFS volumes.

Displays or specifies whether automatic system checking is schedule to be run on a FAT16 or FAT32 volume when the computer is started.

For more information about Chkntfs, see Command Line References in the Tools and Settings Collection.

**Defrag.exe: Defragmentation**
**Category**
Operating system tool

**Version compatibility**
Windows Server 2003. Defrag.exe works on FAT and NTFS volumes.

Arranges clusters on a disk in an optimized order.

For more information about Defrag, see Command Line References in the Tools and Settings Collection.

**Diskpart.exe: Disk Partition**
**Category**
Operating system tool

**Version compatibility**
Windows Server 2003. Diskpart includes a few 64-bit parameters that are only available on Itanium-based computers. Diskpart.exe works on both FAT and NTFS volumes.

Command-line tool that allows you to manage disks and partitions by using scripts or direct input from a command prompt. Diskpart locates and consolidates fragmented boot files, data files, and folders on local or remote volumes.

For more information about Diskpart, see Command Line References in the Tools and Settings Collection.

**Format.exe: Format**
**Category**
Operating system tool

**Version compatibility**
Windows Server 2003. Format.exe works on FAT and NTFS volumes.

Prepares a volume on the specified disk to accept Windows files.

For more information about Format, see Command Line References in the Tools and Settings Collection.

**Fsutil.exe: File Systems Utilities**
**Category**
Operating System tool

**Version compatibility**
Windows Server 2003. Fsutil.exe works on FAT and NTFS volumes.

Performs file system tasks such as disabling long file names, checking whether a volume is flagged as dirty, viewing NTFS-related information about a volume, creating hard links, and managing quotas and sparse files.

For more information about Fsutil, see Command Line References in the Tools and Settings Collection.

**Recovery Console**
**Category**
Operating system tool

**Version compatibility**
Windows Server 2003. Recovery Console works on FAT and NTFS volumes.

The Recovery Console is a text-mode command-line interpreter that you can use for basic troubleshooting and maintenance. It is separate from the Windows Server 2003 command prompt and grants limited local hard disk access for FAT volumes.

**Vrfydsk.exe: Verify Disk**
**Category**
Resource Kit tool

**Version compatibility**
Windows Server 2003. Vrfydsk.exe works on FAT and NTFS volumes.

Command-line tool that checks the consistency and integrity of the file system on a specified volume.

For more information about Vrfydsk, see Resource Kit Tools in the Tools and Settings Collection.

**FAT Registry Entries**
The following registry entries are associated with the FAT file system.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**FileSystem**
The following registry entry is located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\FileSystem.

**Win95TruncatedExtensions**
**Registry Path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem

**Version**
Microsoft Windows NT 4.0 Server and Windows NT 4.0 Workstation and later.

Truncates extensions on FAT partitions, allowing you to work with 3-character extension files as well as 4-character extension files (for example, *.htm is treated identically to *.html).

For more information about this registry entry, see the Registry Reference in the Tools and Settings Collection.

**FAT WMI Classes**

The following table lists and describes the WMI classes that are associated with file systems.

**WMI Classes Associated with FAT**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_AutochkSetting | \root\cimv2 | Windows XP or later |
| Win32_DefragAnalysis | \root\cimv2 | Windows Server 2003 |
| Win32_DiskDriveToDiskPartition | \root\cimv2 | Windows 2000 and later |
| Win32_DiskPartition | \root\cimv2 | Windows 2000 and later |
| Win32_LogicalDisk | \root\cimv2 | Windows 2000 and later |
| Win32_LogicalDiskRootDirectory | \root\cimv2 | Windows 2000 and later |
| Win32_LogicalDiskToPartition | \root\cimv2 | Windows 2000 and later |
| Win32_DiskDrivePhyiscalMedia | \root\cimv2 | Windows 2000 and later |
| Win32_Volume | \root\cimv2 | Windows Server 2003 |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

# Disk Defragmenter Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk Defragmenter Technical Reference**
Defragmentation is the process of consolidating fragmented files and folders. Disk Defragmenter is a Microsoft Windows Server 2003 tool that analyzes storage areas on a disk, referred to as volumes, and consolidates fragmented files and folders so that each occupies a single, contiguous space on the volume. As a result, a system can access and save files and folders more efficiently. By consolidating files and folders, Disk Defragmenter also consolidates the free space on a volume, making it less likely that new files will be fragmented.

Disk Defragmenter can defragment volumes that are formatted with the FAT16, FAT32, and NTFS file systems. Windows Server 2003 offers two choices for defragmenting disks: the Disk Defragmenter snap-in and a new command-line version of the tool, Defrag.exe.

**In this subject**


- What Is Disk Defragmenter?
- How Disk Defragmenter Works
- Disk Defragmenter Tools and Settings

# What Is Disk Defragmenter?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Disk Defragmenter?**

**In this section**


- Common Scenarios for Disk Defragmenter
- Disk Defragmenter Interactions
- Related Information


When you create, copy, or move a file on your hard disk, the operating system looks for blocks of contiguous free disk space where it can write the file. When the operating system cannot find enough contiguous space, it looks for free space scattered throughout the hard disk to which the file can be written in pieces. As a result, pieces of the file may end up scattered over the hard disk rather than in one contiguous chunk. When a file is written to non-contiguous free space, it is known as a fragmented file.

All Windows file systems are susceptible to fragmentation, regardless of the type of file system or the type of disk. When a volume contains many fragmented files and folders, it requires several additional disk drive reads and more movement of the disk drive heads to collect the various pieces, resulting in Windows taking much longer to read a file. Similarly, creating new files and folders also takes longer when the free space available on the volume is scattered. If the free space is not consolidated, then subsequent files written to the volume will more likely be fragmented.

Disk Defragmenter is both a Microsoft Management Console (MMC) snap-in and a command-line utility that can be used to optimize NTFS- or FAT-formatted volumes. Disk Defragmenter rearranges files so that files and free space are consolidated. A reduction in fragmentation results in a reduction in the amount of mechanical movement required to locate all clusters of a file.

**Common Scenarios for Disk Defragmenter**

Disk Defragmenter is commonly used in the following scenarios.

**Installing new software**

Running Disk Defragmenter on a volume after installing software organizes the clusters on the volume.

**Rewriting, updating, and deleting files**

Running Disk Defragmenter on a regular basis consolidates space to create more contiguous free clusters on the volume, which allows disk heads to read files and write updates more quickly.

**Disk Defragmenter Interactions**

Disk Defragmenter interacts with the following technologies.

**NTFS and FAT file systems**

You can use Disk Defragmenter to reduce fragmentation and improve the overall performance of disk input and output (disk I/O) on either NTFS- or FAT-formatted volumes by performing online defragmentation of the master file table (MFT) on NTFS volumes or by working on volumes with any cluster size.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Basic Disks and Volumes Technical Reference
- Dynamic Disks and Volumes Technical Reference
- NTFS Technical Reference
- FAT Technical Reference

# How Disk Defragmenter Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Disk Defragmenter Works**

**In this section**


- Disk Defragmenter Architecture
- Physical Structure of a Volume Using Disk Defragmenter
- Disk Defragmenter Processes and Interactions
- Related Information


Fragmentation causes a disk subsystem to perform more data-seeking operations, slowing the rate at which data is transferred and resulting in sluggish disk performance. Fragmentation can occur under the following conditions.

- You create a file, but the volume does not have a group of contiguous clusters that is large enough to contain the entire file. Therefore, the file is broken into fragments rather than residing in contiguous clusters on the volume.
- You edit a file so that it outgrows its existing space on the volume. When a file uses all the clusters in a group of contiguous, free clusters, the file is then broken into fragments that are stored in free clusters elsewhere on the disk.

To achieve the best results from Disk Defragmenter, run Disk Defragmenter frequently. Do the following before running Disk Defragmenter.

- Verify that the volume has 15 percent of its space free for the utility to use. Windows Server 2003 uses this free space as a sorting area for file fragments.
- Delete any unwanted or temporary files.
- Terminate any file system backup processes before running the utility. Backup processes cause the disk defragmentation process to pause.

This section covers the architecture of Disk Defragmenter, the physical layout of a volume, the process and interactions of defragmentation, and files excluded by the utility.

**Disk Defragmenter Architecture**

Disk Defragmenter interacts with the local file system to move files and consolidate free space, resulting in improved I/O performance. When a fragmented file is requested through an I/O call, the disk heads must move to different tracks on the disk to read all of the fragments. A file that is spread across thousands of fragments may require thousands of I/O calls to the disk controller and extensive movement of the disk head to complete a read operation, resulting in slower performance of the server. To reduce both the I/O calls and the head movement required to read the file, Disk Defragmenter moves the fragments of the file to clusters that are contiguous.

The following diagram, Disk Defragmenter Architecture, shows how the Disk Defragmenter snap-in and command-line tool use the DeviceIOControl file system application programming interface (API) to locate fragments of a file and to move files. The file system API provides access to the layout of the files on the disk and the free space. After the fragments of the file are located, the file system control code is used to move the fragment to a new location.

**Disk Defragmenter Architecture**




In the previous illustration, a disk has three volumes: C, V, and F. Two volumes are formatted with the NTFS file system and one volume is formatted with the FAT file system. Windows Server 2003, the Disk Defragmenter snap-in, and the Disk Defragmenter command-line tool are installed on the C volume. The following occurs when you run the Disk Defragmenter snap-in on the V volume.

1. Disk Defragmenter analyzes the V volume and determines the extent to which the volume is fragmented.

2. Disk Defragmenter uses file system API calls to move and organize the files.

3. Disk Defragmenter writes the reorganized files back to the V volume.

**Physical Structure of a Volume Using Disk Defragmenter**

The physical structure of the volume determines how efficiently the space is used by the file system. The smaller the cluster size, the more efficiently a disk stores information; this is because unused space within a cluster cannot be used by other files. Because the file system has limits on the number of clusters it can support, it chooses the default cluster size of the volume.

### Note

- NTFS uses default cluster sizes unless another cluster size is specified by the user.

Cluster size does not determine where files are written to the volume. Disk Defragmenter organizes clusters, regardless of their size, to more efficiently use resources.

**Default Cluster Sizes**

A cluster (or allocation unit) is the smallest amount of disk space that can be allocated to hold a file. All file systems used by Windows Server 2003 organize hard disks based on cluster size, which is determined by the number of sectors (units of storage on a hard disk) that the cluster contains. For example, on a disk that uses 512-byte sectors, a 512-byte cluster contains one sector, whereas a 4-kilobyte (KB) cluster contains eight sectors.

The following tables give the default cluster sizes for NTFS and FAT file systems.

### Default NTFS Cluster Sizes

| Volume Size | NTFS Cluster Size |
|---|---|
| 7 MB – 512 MB | 512 bytes |
| 513 MB – 1,024 MB | 1 KB |
| 1,025 MB – 2 GB | 2 KB |
| 2 GB – 2 terabytes | 4 KB |

The FAT16 file system supports a maximum of 65,524 clusters per volume, in spite of the 4-GB maximum volume size. FAT32 volumes must have a minimum of 65,527 clusters on a 32-GB volume, which is the largest FAT32 volume that Windows Server 2003 can format.

**Default FAT Cluster Sizes**

| Volume Size | FAT16 Cluster Size | FAT32 Cluster Size |
|---|---|---|
| 7 MB – 16 MB | 2 KB | Not supported |
| 17 MB – 32 MB | 512 bytes | Not supported |
| 33 MB – 64 MB | 1 KB | 512 bytes |
| 65 MB – 128 MB | 2 KB | 1 KB |
| 129 MB – 256 MB | 4 KB | 2 KB |
| 257 MB – 512 MB | 8 KB | 4 KB |
| 513 MB – 1,024 MB | 16 KB | 4 KB |
| 1,025 MB – 2 GB | 32 KB | 4 KB |
| 2 GB – 4 GB | 64 KB | 4 KB |
| 4 GB – 8 GB | Not supported | 4 KB |
| 8 GB – 16 GB | Not supported | 8 KB |
| 16 GB – 32 GB | Not supported | 16 KB |
| 32 GB – 2 terabytes | Not supported | Not supported |

**Volume Structure Before and After Running Disk Defragmenter**

The following figure shows a general example of how Disk Defragmenter organizes file fragments to create contiguous clusters.

**Volume Structure Before and After Running Disk Defragmenter**




Two files, File01 and File02, are saved in nine clusters on a volume. Each cluster-sized part of a file contains an indicator to tell the disk head on which cluster the next part of the file is located. For example, File01 uses four clusters. Contained in cluster 1 with the first chunk of the file's data is a pointer that indicates that the next part of the file is located in cluster 3; cluster 3 indicates that the next part of the file is located in cluster 6; cluster 6 indicates that the next part of the file is located in cluster 8. Cluster 8 has an indicator, 0xFFFF, that marks that cluster as containing the last part of the file. Because the parts are spread out across four chunks on the volume, the disk head must move more to retrieve the parts of the file than if the file were consolidated in one chunk. This extra movement results in decreased performance.

After running Disk Defragmenter, file parts are organized such that more parts of each file are contiguous. Thus, File01 now uses clusters 0 – 3. Because the disk head has less distance to travel, it has less work to do, and performance is increased. In addition, the free space available in clusters 2, 4, 7, and 9 before disk defragmentation has been consolidated; there is now a larger free space in clusters 6 – 9.

**Disk Defragmenter Processes and Interactions**

The amount of time that defragmentation takes depends on several factors, including the size of the volume, the number and

size of files on the volume, the amount of fragmentation, and the available local system resources. You can find the fragmented files and folders before defragmenting them by analyzing the volume first. You can then see how fragmented the volume is and decide whether to defragment the volume.

The defragmentation process can run even when files are in use by another process. The movement of a file by Disk Defragmenter is invisible to the other process and is managed by the file system APIs. There will be some performance impact on the system if Disk Defragmenter is run while backups are in progress, but there is no risk of data loss due to the use of Disk Defragmenter.

This section covers the following topics.

- Process limitations
- Analyzing volumes
- Updating the master file table (MFT)
- Optimizing startup

**Process Limitations**

Disk Defragmenter cannot be used under the following circumstances.

- You cannot defragment volumes that the file system has marked as dirty, which indicates possible corruption. Running Chkdsk or Fsutil on a volume in advance of running Disk Defragmenter can help to determine whether a volume is dirty.
- You cannot run Disk Defragmenter from the command line and through the Disk Defragmenter snap-in simultaneously. If you run Disk Defragmenter from the command line and then open the Disk Defragmenter snap-in, the defragmentation options in the Disk Defragmenter snap-in are unavailable. Conversely, if you are using the snap-in to defragment a volume and then run Disk Defragmenter from the command line, the command-line tool fails.

**Files Excluded by Disk Defragmenter**

All Windows file systems have on-disk structures that are excluded by Disk Defragmenter. Files are excluded because it may not be possible to move the file after the operating system has started up or because the benefit of consolidating the file is negligible. The following table lists files excluded by Disk Defragmenter.

**Files Excluded by Disk Defragmenter**

| Component | Description |
|---|---|
| Bootsect.dos | Holds the contents of the boot sector, disk sector 0 from the bootable hard drive volume. |
| Safeboot.fs, Safeboot.csv, Safeboot.rsv | Contain an internal verification of data structures. |
| Hiberfil.sys | File to which the system's physical memory is written during hibernation. On resuming from hibernation, the BIOS reads Hiberfil.sys to restore the state of the computer to its pre-hibernation state. Because the location of the Hibernate file is determined very early in the startup process, it cannot be moved. It can, however, be defragmented safely at startup using an enterprise-level defragmenter. |
| Memory.dmp | Holds the contents of the kernel address space when a computer restarts due to an error. |
| Paging file | Used by the operating system as storage space when swapping data in and out of random access memory (RAM). When memory requirements exceed the amount of physical RAM, Virtual Memory Manager transfers the oldest data stored in RAM to the paging file. The paging file can be defragmented during the boot process only. |
| FAT directories | Lists the contents of a directory on FAT16 and FAT32 file systems. |

| | |
|---|---|
| Change journal | Provides a persistent log of changes made to files on an NTFS volume. NTFS maintains the change journal by tracking information about added, deleted, and modified files for each volume. |

**Note**

- Although Disk Defragmenter does not defragment FAT-formatted directories, it will defragment directories on NTFS-formatted volumes.

**Analyzing Volumes**

By using the Disk Defragmenter snap-in, you can analyze the volume before you defragment and view the fragmentation report to see how many files and folders are fragmented. If 10 percent or more of the files and folders are fragmented, it is recommended that you defragment the volume. If the volume is less than 10 percent fragmented, you can still defragment it.

The following figure shows the Disk Defragmenter snap-in.

**Disk Defragmenter snap-in**




The upper part of the screen lists the volumes on the local computer and allows you to select a volume to analyze and defragment. The lower part of the screen displays a graphical representation of the fragmentation on the volume. The colors indicate the condition of the volume:

- Red areas show fragmented files.
- Blue areas show contiguous (unfragmented) files.
- White areas show free space on the volume. White areas on an NTFS volume might also represent the MFT zone (the space set aside to keep as much of the MFT as possible contiguous).
- Green areas show files that cannot be moved. The green areas usually represent the paging file, but on NTFS volumes, green areas might also represent space used by the NTFS change journal and the NTFS log file.

By comparing the estimated disk usage before defragmentation to the estimated disk usage after defragmentation, you can see the improvement in your volume. The defragmentation report provides further details on the fragmentation state of the volume. Analyze volumes regularly and defragment them when Disk Defragmenter recommends it.

**Updating the Master File Table**

On NTFS volumes, the MFT represents an index for every file on the volume and contains an entry for every file, including the MFT itself. With the introduction of Windows Server 2003, the MFT can be defragmented while the server is online; it is no longer excluded by Disk Defragmenter. Although Disk Defragmenter attempts to defragment the MFT after it becomes

fragmented, it cannot move the first 16 KB of the file, which contains the metadata for the file system. If space is available, the remainder of the MFT can be moved as one fragment to a location on the volume. The NTFS file system reserves one-eighth of the disk space for the MFT to use. This area of the disk is known as the MFT zone and it is set aside to keep as much of the MFT as possible contiguous.

The following steps are the processes by which Disk Defragmenter and file system defragment an NTFS volume and update the MFT.

1. Locate the fragments of each file throughout the disk.
2. Copy the fragments to contiguous locations on the disk.
3. Verify that the copy is an exact duplicate of the original file.
4. Update the MFT in NTFS with the new location of the beginning of the file.
5. Free up the space previously used by the file.

### Note

- When file fragments are moved to contiguous clusters, the files within a folder are not necessarily grouped in contiguous clusters by the defragmentation process. It should not be assumed that a logical grouping of files in a folder translates into a physical grouping of files on the volume.

#### Optimizing Startup

The Boot Optimize functionality uses Disk Defragmenter to move startup files into a contiguous location at the start of the volume. Disk Defragmenter moves the files in a specified order to a location that affords the earliest full fit on the volume for all of the files. This placement reduces the number and distance of disk head movements required to read startup files into memory, allowing the computer to boot more quickly.

#### Related Information

The following resources contain additional information that is relevant to this section.

- NTFS Technical Reference
- FAT Technical Reference

# Disk Defragmenter Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk Defragmenter Tools and Settings**

**In this section**


- Disk Defragmenter Tools
- Disk Defragmenter Registry Entries
- Disk Defragmenter Group Policy Settings
- Disk Defragmenter WMI Classes
- Related Information


**Disk Defragmenter Tools**

The following tools are associated with Disk Defragmenter.

**Defrag.exe: Disk Defragmenter**

**Category**

Operating system tool

**Version compatibility**

All versions of the Disk Defragmenter command-line tool in the Windows Server 2003 family are identical.

Disk Defragmenter displays a blinking cursor while it is analyzing and defragmenting a volume. When Disk Defragmenter finishes analyzing a volume, it displays the analysis report. When Disk Defragmenter finishes defragmenting a volume, it displays the defragmentation report. After the process completes, Disk Defragmenter exits to the command prompt. The command-line version of Disk Defragmenter will analyze a volume only if the **-a** option is used.

For more information about Defrag.exe, see "Command Line References" in the Tools and Settings Collection.

**Dfrg.msc: Disk Defragmenter snap-in**

**Category**

Operating system tool

**Version compatibility**

All versions of the Disk Defragmenter snap-in in the Windows Server 2003 family are identical.

The Disk Defragmenter snap-in is divided into two main areas. The upper part lists the volumes on the local computer and allows you to select a volume to analyze and defragment. The lower part displays a graphical representation of the fragmentation on the volume. The colors indicate the condition of the volume:

- Red areas show fragmented files.
- Blue areas show contiguous (unfragmented) files.
- White areas show free space on the volume. White areas on an NTFS volume might also represent the MFT zone. For more information about the MFT zone, see "NTFS Technical Reference".
- Green areas show files that cannot be moved. The green areas usually represent the paging file, but on NTFS volumes, green areas might also represent space used by the NTFS change journal and the NTFS log file.

**Disk Defragmenter Registry Entries**

The following registry entries are associated with Disk Defragmenter.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It

is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

For more information about registry entries, see "Registry Reference" in the Tools and Settings Collection.

**BootOptimizeFunction**
The following registry entries are located under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg\BootOptimizeFunction.

**Enable**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg\BootOptimizeFunction

**Version**
Windows Server 2003

Accepted values for this entry are Y or N. If the entry is set to Y, Windows automatically optimizes the file location for boot optimization. This optimization occurs automatically if the system is idle for 10 minutes. Boot optimization improves startup time by locating startup files in contiguous clusters on the volume, reducing the movement of the disk head when reading the volume.

**LcnEndLocation**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg\BootOptimizeFunction

**Version**
Windows Server 2003

Stores the logical cluster number (LCN) that specifies the end location of the files moved by Disk Defragmenter for boot optimization.

**LcnStartLocation**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg\BootOptimizeFunction

**Version**
Windows Server 2003

Stores the LCN that specifies the starting location of the files moved by Disk Defragmenter for boot optimization.

**OptimizeComplete**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg\BootOptimizeFunction

**Version**
Windows Server 2003

Stores the value that indicates whether the optimize operation is complete.

**OptimizeError**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg\BootOptimizeFunction

**Version**
Windows Server 2003

Stores an error message if an error was encountered during the optimization process.

**Dfrg**
The following registry entries are located under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg.

**FreeSpaceErrorLevel**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg

**Version**
Windows Server 2003

Sets the percentage of volume capacity that must be free space for Disk Defragmenter to use as a sorting area. The default value is 15.

**PathName**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg

**Version**
Windows Server 2003

Stores the path to the resource DLL for defragmentation, usually *systemroot*\system32.

**ResourceDllName**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dfrg

**Version**
Windows Server 2003

Stores the name of the resource DLL for defragmentation, usually *systemroot*\system32\DfrgRes.dll.

**Disk Defragmenter Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with Disk Defragmenter.

### Group Policy Settings Associated with Disk Defragmenter

| Group Policy Setting | Description |
|---|---|
| Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\Back up files and directories | Allows the Administrator to bypass access control list (ACL) check when obtaining the file handle during the defragmentation procedure. |

For more information about Group Policy settings, see "Group Policy Settings Reference" in the Tools and Settings Collection.

**Disk Defragmenter WMI Classes**
The following table lists and describes the WMI classes that are associated with Disk Defragmenter. The Win32_Volume class contains the Defrag method, used to defragment a FAT or NTFS volume.

### WMI Classes Associated with Disk Defragmenter

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_Volume | \root\cimv2 | Windows Server 2003 |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Related Information**
The following resources contain additional information that is relevant to this section.

- FAT Technical Reference
- NTFS Technical Reference

# Disks and Volumes

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disks and Volumes**

The Microsoft Windows Server 2003 operating system offers two types of disks for storing information on your server: basic and dynamic. Basic and dynamic disks offer flexibility in the types of volumes you use to store data. In x86-based computers, basic disks are the default type of storage and use the same disk structures as those used by the Microsoft MS-DOS operating system and all previous versions of Microsoft Windows. Dynamic disks provide features that basic disks do not, such as the ability to create volumes that span multiple disks, including fault-tolerant volumes. Dynamic disks were introduced in Microsoft Windows 2000, and use a database to track information about dynamic volumes on the disk and about other dynamic disks in the computer. Because each dynamic disk in a computer stores a replica of the dynamic disk database, Windows Server 2003 can repair a corrupted database on one dynamic disk by using the database on another dynamic disk.

When using basic disks in x86-based computers, you are limited to creating four primary partitions per disk, or three primary partitions and one extended partition with unlimited logical drives. Primary partitions and logical drives on basic disks are known as basic volumes. All volumes on dynamic disks are known as dynamic volumes. All dynamic disks in a computer automatically belong to the same disk group.

The term partition style refers to the method that Windows Server 2003 uses to organize partitions on the disk. All x86-based computers use the partition style known as the master boot record (MBR). The MBR contains a partition table that describes where the partitions are located on the disk. Because MBR is the only partition style available on x86-based computers, it is used automatically and you do not need to choose this style.

Itanium-based computers running the 64-bit versions of Windows Server 2003 use a new partition style called the globally unique identifier (GUID) partition table (GPT). The GPT partition style supports partitions up to 2 terabytes and 128 partitions per disk.

**Note**

- You can use the 64-bit versions of Windows Server 2003 to *manage* MBR disks and GPT disks. However, you cannot *start* the 64-bit versions of Windows Server 2003 from an MBR disk.

After you install a new disk, you must choose a partition style and storage type to use on the disk. Your choices vary according to which operating system you are running and whether the computer is an x86-based computer or an Itanium-based computer. The following table lists the storage types and partition styles that are available for each version of Windows Server 2003.

**Storage Types and Partition Styles Available in Windows Server 2003**

| Operating System | Basic Volumes | Dynamic Simple, Spanned, and Striped Volumes | Dynamic Mirrored and RAID-5 Volumes | MBR Partition Style | GPT Partition Style |
|---|---|---|---|---|---|
| Windows Server 2003, Standard Edition | ● | ● | ● | ● | |
| Windows Server 2003, Enterprise Edition | ● | ● | ● | ● | |
| Windows Server 2003, Datacenter Edition | ● | ● | ● | ● | |
| Windows Server 2003, Web Edition | ● | ● | ● | ● | |
| The 64-bit version of Windows Server 2003, Enterprise Edition | ● | ● | ● | ● | ● |
| The 64-bit version of Windows Server 2003, Datacenter Edition | ● | ● | ● | ● | ● |

The introduction of GPT makes understanding the partition styles a bit more challenging, but most disk-related tasks are unchanged. You can still use basic disks and dynamic disks as you did in Windows 2000, and these storage types are available on disks that use either partition style.

The Disk Management snap-in differentiates between partition styles by referring to disks that use the master boot record as MBR disks and disks that use the GUID partition table as GPT disks. The following figure shows how Disk Management displays GPT and MBR disks in an Itanium-based computer.

**How Disk Management Displays GPT and MBR Disks on an Itanium-based Computer**



# Basic Disks and Volumes Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Basic disks and basic volumes are the storage types most often used with Microsoft Windows operating systems. The term basic disk refers to a disk that contains basic volumes, such as primary partitions and logical drives. The term basic volume refers to a partition on a basic disk. Basic disks, which are found in both x86-based and Itanium-based computers, provide a simple storage solution that can accommodate changing storage requirements. Basic disks support clustered disks, Institute of Electrical and Electronics Engineers (IEEE) 1394 disks, and universal serial bus (USB) removable drives.

In x86-based computers running Windows Server 2003, basic disks use the same Master Boot Record (MBR) partition style as the disks used by the Microsoft MS-DOS operating system and all previous versions of Microsoft Windows.

Itanium-based computers also support basic disks, but you can choose from two partition styles (MBR or GPT) for each basic disk. You can create up to 128 volumes on an MBR or GPT disk. The partition style determines the operating systems that can access the disk.

The following topics within the Basic Disks and Volumes Technical Reference describe how basic disks and volumes work, their architecture, and the tools you can use to manipulate them.

**In this subject**


- What Are Basic Disks and Volumes?
- How Basic Disks and Volumes Work
- Basic Disks and Volumes Tools and Settings

# What Are Basic Disks and Volumes?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Basic Disk and Volume Scenarios


Basic disks and basic volumes are the storage types most often used with Windows operating systems. The term basic disk refers to a disk that contains basic volumes, such as primary partitions and logical drives. The term basic volume refers to a partition on a basic disk. Basic disks, which are found in both x86-based and Itanium-based computers, provide a simple, elegant storage solution that can accommodate changing storage requirements. Basic disks support clustered disks, Institute of Electrical and Electronics Engineers (IEEE) 1394 disks, and universal serial bus (USB) removable drives.

In x86-based computers running Windows Server 2003, basic disks use the same Master Boot Record (MBR) partition style as the disks used by Microsoft MS-DOS, and all previous versions of Microsoft Windows.

Itanium-based computers also support basic disks, but you can choose from two partition styles (MBR or GPT) for each basic disk. You can create up to 128 volumes on an MBR or GPT disk. The partition style determines the operating systems that can access the disk.

Before you can create simple volumes, spanned volumes, or volumes that use redundant array of independent disks (RAID) technology (striped volumes, mirrored volumes, and RAID-5 volumes) you must convert a basic disk to a dynamic disk.

Windows Server 2003 supports the following types of basic volumes:

- Primary partitions (master boot record (MBR) and GUID partition table (GPT) disks)
- Logical drives within extended partitions (MBR disks only)

The number of basic volumes you can create on a basic disk depends on the partition style of the disk:

- On MBR disks, you can create up to four primary partitions, or you can create up to three primary partitions and one extended partition. Within the extended partition, you can create up to 128 logical drives.
- On GPT disks, you can create up to 128 partitions. Because GPT disks do not limit you to four partitions, extended partitions and logical drives are not available on GPT disks.
- If you want to add more space to existing primary partitions and logical drives, you can extend the volume using the **extend** command in DiskPart.

**Basic Disk and Volume Scenarios**

Basic disks and volumes can be scaled to match your storage needs. Basic disks and volumes are commonly used in the following scenarios.

**Home or business desktop computer with one disk**

Most home and business users require a basic disk and one basic volume for storage, and do not require a computer with volumes that span multiple disks or that provide fault-tolerance. This is the best choice for those who require simplicity and ease of use.

**Home or business desktop computer with one disk and more than one volume**

If a home or small business user wants to upgrade the operating system without losing their personal data, they should store the operating system in a separate location from their personal data. In this scenario, a basic disk with two or more basic volumes is required. The user can install the operating system on the first volume, creating a boot volume or system volume, and use the second volume to store data. When a new version of the operating system is released, the user can reformat the boot or system volume and install the new operating system. Their personal data, located on the second volume, remains untouched.

**Business server with one disk and multiple volumes and logical drives**

If a small business operates a file server and requires multiple volumes for file sharing and file security, the system administrator can create up to three primary partitions and one extended partition with up to 128 logical drives. In this scenario, each of the partitions and logical drives receives its own drive letter so that each of these volumes can be individually secured to limit access to specific, authorized users.

For example, perhaps each department within this business requires its own volume. The business could create individual volumes and grant permissions to members of those departments. Data shared by members of human resources, for example, could be kept separate from the data used by members of the accounting, sales, or marketing departments.

As storage needs and the importance of the data stored on disk increase, the logical next step for this business would be to add additional disks to the server, convert the disks to dynamic, and then create fault-tolerant mirrored- or RAID-5 volumes.

# How Basic Disks and Volumes Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Basic Disks and Volumes Architecture
- Basic Disk and Volume Interfaces
- Basic Disks and Volumes Physical Structure
- Disk Sectors on GPT Disks
- Basic Disks and Volumes Processes and Interactions


Basic disks and basic volumes are the storage types most often used with Windows operating systems. The term basic disk refers to a disk that contains basic volumes, such as primary partitions and logical drives. The term basic volume refers to a partition on a basic disk. Basic disks, which are found in both x86-based and Itanium-based computers, provide a simple, elegant storage solution that can accommodate changing storage requirements. Basic disks support clustered disks, Institute of Electrical and Electronics Engineers (IEEE) 1394 disks, and universal serial bus (USB) removable drives. Before you can use a basic disk, it must have a disk signature and be formatted with either the File Allocation Table (FAT) or NT file system (NTFS) file systems.

An optimal environment for basic disks and volumes is defined as follows:

- Windows Server 2003 operating system is installed and functioning properly.
- The basic disk is functioning properly and displays the Healthy status in the Disk Management snap-in.

The following sections provide an in-depth view of how basic disks and volumes work in an optimal environment.

**Basic Disks and Volumes Architecture**

Basic disks and volumes rely on the Logical Disk Manager (LDM) and Virtual Disk Service (VDS) and their associated components. These components enable you to perform tasks such as converting basic disks into dynamic disks, and creating fault-tolerant volumes. The following diagram shows the LDM and VDS components.

**Logical Disk Manager and Virtual Disk Service Components**




The following table lists the LDM and VDS components and provides a brief description of each.

**Logical Disk Manager and Virtual Disk Service Components**

| Component | Description |
|---|---|
| Disk Management snap-in:<br><br>Dmdlgs.dll<br><br>Dmdskmgr.dll<br><br>Dmview.ocx<br><br>Diskmgmt.msc | Binaries that comprise the Disk Management snap-in user interface. |
| DiskPart command line utility:<br><br>Diskpart.exe | A scriptable alternative to the Disk Management snap-in. |
| Mount Manager command line:<br><br>Mountvol.exe | A command line utility that can be used to create, delete, or list volume mount points |

| | |
|---|---|
| Virtual Disk Service:<br><br>Vds.exe<br><br>Vdsutil.dll | A program used to configure and maintain volume and disk storage. |
| Virtual Disk Service provider for basic disks and volumes:<br><br>Vdsbas.dll | The Virtual Disk Service calls into the basic provider when configuring basic disks and volumes. |
| Virtual Disk Service provider for dynamic disks and volumes:<br><br>Vdsdyndr.dll | The Virtual Disk Service calls into the dynamic provider when configuring dynamic disks and volumes. |
| Dmboot.sys<br><br>Dmconfig.dll<br><br>Dmintf.dll<br><br>Dmio.sys<br><br>Dmload.sys<br><br>Dmremote.exe<br><br>Dmutil.dll | Drivers and user mode components used to configure dynamic disks and volumes and perform I/O. |
| Logical Disk Administrator service:<br><br>Dmadmin.exe | The VDS provider for dynamic disks and volumes uses the interfaces exposed by this service to configure dynamic disks. |
| Logical Disk Service:<br><br>Dmserver.dll | A service that detects and monitors new hard disk drives and sends disk volume information to Logical Disk Manager Administrative Service for configuration. If this service is stopped, dynamic disk status and configuration information might become outdated. If this service is disabled, any services that explicitly depend on it will fail to start. |
| Basic disk I/O driver:<br><br>Ftdisk.sys | A driver that manages all I/O for basic disks. Other system components, such as mount point manager, call into this driver to get information about basic disk volumes. |
| Mount point manager driver:<br><br>Mountmgr.sys | A binary that tracks drive letters, folder mount paths and other mount points for volumes. Assigns a unique volume mount point of the form \??\Volume<*GUID*> to each volume, in addition to any drive letters or folder paths that have been assigned by the user. Ensures that a volume will get the same drive letter each time the computer boots, and also tries to retain a volume's drive letter when the volume's disk is moved to a new computer. |
| Partition manager:<br><br>Partmgr.sys | A filter driver that sits on top of the disk driver. All disk driver requests pass through the partition manager driver. This driver creates partition devices and notifies the volume managers of partition arrivals and removals. Exposes IOCTLs that return information about partitions to other components, and allow partition configuration. |

**Basic Disk and Volume Interfaces**

Basic disks and volumes can be managed using the Disk Management snap-in. The following table lists the control codes used

by the Disk Management snap-in when managing disks and volumes.

For more information about the interfaces used by Disk Management when managing disks and volumes, see *Disk Management Reference* on MSDN.

**Disk Management Control Codes**

| Control Code | Operation |
|---|---|
| IOCTL_DISK_CREATE_DISK | Initializes the specified disk and disk partition table using the specified information. |
| IOCTL_DISK_DELETE_DRIVE_LAYOUT | Removes the boot signature from the master boot record. |
| IOCTL_DISK_GET_DRIVE_GEOMETRY_EX | Retrieves information about the physical disk's geometry. |
| IOCTL_DISK_GET_DRIVE_LAYOUT_EX | Retrieves information about the number of partitions on a disk and the features of each partition. |
| IOCTL_DISK_GET_LENGTH_INFO | Retrieves the length of the specified disk, volume, or partition. |
| IOCTL_DISK_GET_PARTITION_INFO_EX | Retrieves partition information for AT and EFI (Extensible Firmware Interface) partitions. |
| IOCTL_DISK_GROW_PARTITION | Enlarges the specified partition. |
| IOCTL_DISK_SET_DRIVE_LAYOUT_EX | Partitions a disk. |
| IOCTL_DISK_SET_PARTITION_INFO_EX | Sets the disk partition type. |

For more information about the Disk Management control codes, see *Disk Management Control Codes* on MSDN.

**Basic Disks and Volumes Physical Structure**

Basic disks can use either the master boot record (MBR) or GUID partition table (GPT) partitioning style. x86-based computers use disks with the MBR partitioning style and Itanium-based computers use disks with the GPT partitioning style.

The following figure compares a basic MBR disk to a basic GPT disk.

**Comparison of MBR and GPT Disks**




A comparison of MBR and GPT disks is listed in the following table.

**Comparison of MBR and GPT Disks**

| Characteristic | MBR Disk (x86-based Computer) | GPT Disk (Itanium-based Computer) |
|---|---|---|
| Number of partitions on basic disks | Supports up to:<br>• Four primary partitions per disk, or<br>• Three primary partitions and an extended partition with up to 128 logical drives | Supports up to 128 partitions |
| Compatible operating systems | Can be read by:<br>• Microsoft MS-DOS<br>• Microsoft Windows 95<br>• Microsoft Windows 98<br>• Microsoft Windows Millennium Edition<br>• Windows NT, all versions<br>• Windows 2000, all versions<br>• Windows XP<br>• Windows Server 2003, all versions for x86-based computers and Itanium-based computers | Can be read by:<br>• Windows XP 64-Bit Edition<br>• The 64-bit version of Windows Server 2003, Enterprise Edition<br>• The 64-bit version of Windows Server 2003, Datacenter Edition |
| Maximum size of basic volumes | 2 terabytes | 2 terabytes |
| Partition tables (copies) | Contains one copy of the partition table. | Contains primary and backup partition tables for redundancy and checksum fields for improved partition structure integrity. |
| Locations for data storage | Stores data in partitions and in unpartitioned space. Although most data is stored within partitions, some data might be stored in hidden or unpartitioned sectors created by OEMs or other operating systems. | Stores user and program data in partitions that are visible to the user. Stores data that is critical to platform operation in partitions that the 64-bit versions of Windows Server 2003 recognize but do not make visible to the user. Does not store data in unpartitioned space. |
| Troubleshooting methods | Uses the same methods and tools used in Windows 2000. | Uses tools designed for GPT disks. (Do not use MBR troubleshooting tools on GPT disks.) |

**Master Boot Record on Basic Disks**

The master boot record (MBR), the most important data structure on the disk, is created when the disk is partitioned. The MBR contains a small amount of executable code called the master boot code, the disk signature, and the partition table for the disk. At the end of the MBR is a 2-byte structure called a signature word or end of sector marker, which is always set to 0x55AA. A signature word also marks the end of an extended boot record (EBR) and the boot sector.

The disk signature, a unique number at offset 0x01B8, identifies the disk to the operating system. Windows Server 2003 uses the disk signature as an index to store and retrieve disk information, such as drive letters, in the registry.

**Master boot code**

The master boot code performs the following activities:

1. Scans the partition table for the active partition.
2. Finds the starting sector of the active partition.
3. Loads a copy of the boot sector from the active partition into memory.
4. Transfers control to the executable code in the boot sector.

If the master boot code cannot complete these functions, the system displays a message similar to one of the following:

```
Invalid partition table.
Error loading operating system.
Missing operating system.
```

**Note**

- Floppy disks and removable disks, such as Iomega Zip disks, do not contain an MBR. The first sector on these disks is the boot sector. Although every hard disk contains an MBR, the master boot code is used only if the disk contains the active primary partition.

**Partition Table on Basic Disks**

The partition table, which is a 64-byte data structure that is used to identify the type and location of partitions on a hard disk, conforms to a standard layout, independent of the operating system. Each partition table entry is 16 bytes long, with a maximum of four entries. Each entry starts at a predetermined offset from the beginning of the sector, as follows:

- Partition 1 0x01BE (446)
- Partition 2 0x01CE (462)
- Partition 3 0x01DE (478)
- Partition 4 0x01EE (494)

The following example shows a partial printout of an MBR revealing the partition table from a computer with three partitions. When there are fewer than four partitions on a disk, the remaining partition table fields are set to the value 0.

```
000001B0:                                                 80 01                  ..
000001C0: 01 00 07 FE BF 09 3F 00 - 00 00 4B F5 7F 00 00 00   ......?...K....
000001D0: 81 0A 07 FE FF FF 8A F5 - 7F 00 3D 26 9C 00 00 00   .........=&....
000001E0: C1 FF 05 FE FF FF C7 1B - 1C 01 D6 96 92 00 00 00   ................
000001F0: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00         ..............
```

The following figure provides an example of how to interpret the sector printout of the partition table. The Boot Indicator, System ID, Relative Sectors, and Total Sectors values correspond to the values described in the table titled "Partition Table Fields" later in this section.

**Interpreting Data in the Partition Table**

```
                    System ID
                                Relative Sectors
                                                Total Sectors
                                                              Boot Indicator
000001B0:                                                     80 01                  ..
000001C0: 01 00 07 FE BF 09 3F 00  - 00 00 4B F5 7F 00 00 00   ......?...K.▯...
000001D0: 81 0A 07 FE FF FF 8A F5  - 7F 00 3D 26 9C 00 00 00   ........▯.=&....
000001E0: C1 FF 05 FE FF FF C7 1B  - 1C 01 D6 96 92 00 00 00   ................
000001F0: 00 00 00 00 00 00 00 00  - 00 00 00 00 00 00         ..............
```

The following table describes the fields in each entry in the partition table. The Sample Values correspond to the first partition table entry shown in the previous example. The Byte Offset values correspond to the addresses of the first partition table entry. There are three additional entries whose values can be calculated by adding 10h to the byte offset value specific for each additional partition table entry (for example, add 20h for partition table entry 3 and 30h for partition table entry 4).

The following table and sections following the table provide additional detail about these fields.

**Partition Table Fields**

| Byte Offset | Field Length | Sample Value1 | Field Name and Definition |
|---|---|---|---|
| 0x01BE | 1 byte | 80 | **Boot Indicator**. Indicates whether the volume is the active partition. Legal values include:<br>• 00. Do not use for booting.<br>• 80. Active partition. |
| 0x01BF | 1 byte | 01 | **Starting Head**. |
| 0x01C0 | 6 bits | 01 *2 | **Starting Sector**. Only bits 0–5 are used. The upper two bits, 6 and 7, are used by the Starting Cylinder field. |
| 0x01C1 | 10 bits | 00 * | **Starting Cylinder**. Uses 1 byte in addition to the upper 2 bits from the Starting Sector field to make up the cylinder value. The Starting Cylinder is a 10-bit number that has a maximum value of 1023. |
| 0x01C2 | 1 byte | 07 | **System ID**. Defines the volume type. See the table titled "System ID Values" later in this section for sample values. |
| 0x01C3 | 1 byte | FE | **Ending Head**. |
| 0x01C4 | 6 bits | BF * | **Ending Sector**. Only bits 0–5 are used. The upper two bits, 6 and 7, are used by the Ending Cylinder field. |
| 0x01C5 | 10 bits | 09 * | **Ending Cylinder**. Uses 1 byte in addition to the upper 2 bits from the Ending Sector field to make up the cylinder value. The Ending Cylinder is a 10-bit number, with a maximum value of 1023. |
| 0x01C6 | 4 bytes | 3F 00 00 00 | **Relative Sectors**. The offset from the beginning of the disk to the beginning of the volume, counting by sectors. |
| 0x01CA | 4 bytes | 4B F5 7F 00 | **Total Sectors**. The total number of sectors in the volume. |

1. Numbers larger than one byte are stored in little endian format, or reverse-byte ordering. Little endian format is a method of storing a number so that the least significant byte appears first in the hexadecimal number notation. For example, the sample value for the Relative Sectors field in the previous table, 3F 00 00 00, is a little endian representation of 0x0000003F. The decimal equivalent of this little endian number is 63.

2. Sample values marked with an asterisk (*) do not accurately represent the value of the fields, because the fields are either 6 bits or 10 bits and the data is recorded in bytes.

**Boot Indicator Field**

The first element of the partition table, the Boot Indicator field, indicates whether the volume is the active partition. Only one primary partition on the disk can have this field set. See the previous table for the acceptable values.

It is possible to have different operating systems and different file systems on different volumes. By using disk configuration tools, such as the Windows Server 2003-based Disk Management and DiskPart, or the MS-DOS-based Fdisk, to designate a primary partition as active, the Boot Indicator field for that partition is set in the partition table.

**System ID Field**

Another element of the partition table is the System ID field. It defines which file system, such as FAT16, FAT32, or NTFS, was used to format the volume. The System ID field also identifies an extended partition, if one is defined. Windows Server 2003 uses the System ID field to determine which file system device drivers to load during startup. The following table identifies the values for the System ID field.

**System ID Values**

| Partition Type | ID Value |
|---|---|
| 0x01 | FAT12 primary partition or logical drive (fewer than 32,680 sectors in the volume) |

| | |
|---|---|
| 0x04 | FAT16 partition or logical drive (32,680–65,535 sectors or 16 MB–33 MB) |
| 0x05 | Extended partition |
| 0x06 | BIGDOS FAT16 partition or logical drive (33 MB–4 GB) |
| 0x07 | Installable File System (NTFS partition or logical drive) |
| 0x0B | FAT32 partition or logical drive |
| 0x0C | FAT32 partition or logical drive using BIOS INT 13h extensions |
| 0x0E | BIGDOS FAT16 partition or logical drive using BIOS INT 13h extensions |
| 0x0F | Extended partition using BIOS INT 13h extensions |
| 0x12 | EISA partition or OEM partition |
| 0x42 | Dynamic volume |
| 0x84 | Power management hibernation partition |
| 0x86 | Multidisk FAT16 volume created by using Windows NT 4.0 |
| 0x87 | Multidisk NTFS volume created by using Windows NT 4.0 |
| 0xA0 | Laptop hibernation partition |
| 0xDE | Dell OEM partition |
| 0xFE | IBM OEM partition |
| 0xEE | GPT partition |
| 0xEF | EFI System partition on an MBR disk |

Windows Server 2003 does not support multidisk volumes that are created by using Windows NT 4.0 or earlier, and that use System ID values 0x86, 0x87, 0x8B, or 0x8C.

Before you upgrade from Windows NT Server 4.0 to Windows Server 2003, you must first back up and then delete all multidisk volumes. After you complete the upgrade, create dynamic volumes and restore the data. If you do not delete the multidisk volumes before beginning Setup, you must use the Ftonline tool, which is part of Windows Support Tools, to access the volume after Setup completes.

If you are upgrading from Windows 2000 to Windows Server 2003, you must convert the multidisk volumes to dynamic before you begin Setup, or Setup does not continue.

MS-DOS can only access volumes that have a System ID value of 0x01, 0x04, 0x05, or 0x06. However, you can delete volumes that have the other values listed in the table titled "System ID Values" earlier in this section by using Disk Management, DiskPart, or the MS-DOS tool Fdisk.

**Starting and Ending Cylinder, Head, and Sector Fields**

The Starting and Ending Cylinder, Head, and Sector fields (collectively known as the CHS fields) are additional elements of the partition table. These fields are essential for starting the computer. The master boot code uses these fields to find and load the boot sector of the active partition. The Starting CHS fields for non-active partitions point to the boot sectors of the remaining primary partitions and the extended boot record (EBR) of the first logical drive in the extended partition, as shown in the figure titled "Interpreting Data in the Partition Table" earlier in this section.

Knowing the starting sector of an extended partition is very important for low-level disk troubleshooting. If your disk fails, you need to work with the partition starting point (among other factors) to retrieve stored data.

The Ending Cylinder field in the partition table is 10 bits long, which limits the number of cylinders that can be described in the partition table to a range of 0 through 1,023. The Starting Head and Ending Head fields are each one byte long, which limits the field range from 0 through 255. The Starting Sector and Ending Sector fields are each six bits long, which limits the range of these fields from 0 through 63. However, the enumeration of sectors starts at 1 (not 0, as for other fields), so the maximum number of sectors per track is 63.

Because Windows Server 2003 supports hard disks that are low-level formatted with a standard 512-byte sector, the maximum disk capacity described by the partition table is calculated as follows:

Maximum capacity = sector size cylinders (10 bits) heads (8 bits) sectors per track (6 bits).

Using the maximum possible values yields:

512 1024 256 63 (or 512 x $2^{24}$) = 8,455,716,864 bytes or 7.8 gigabytes (GB).

Windows Server 2003 and other Windows-based operating systems that support BIOS INT 13h extensions can access partitions that exceed the first 7.8 GB of the disk by ignoring the Starting and Ending CHS fields in favor of the Relative Sectors and Total Sectors fields.

Windows 2000 and Windows Server 2003 ignore the Starting and Ending CHS fields regardless of whether the partition exceeds the first 7.8 GB of the disk. However, Windows Server 2003 must place the appropriate values in the Starting and Ending CHS fields because Windows 95, Windows 98, and Windows Millennium Edition (which all support BIOS INT 13h extensions) use the Starting and Ending CHS fields if the partition does not exceed the first 7.8 GB of the disk. These fields are also required to maintain compatibility with the BIOS INT 13h for startup.

MS-DOS and other Windows operating systems that do not support BIOS INT 13h extensions ignore partitions that exceed the 7.8 GB boundary because these partitions use a System ID that is recognized only by operating systems that support BIOS INT 13h extensions.

Both the operating system and the computer must support BIOS INT 13h extensions if you want to create partitions that exceed the first 7.8 GB of the disk.

**Relative Sectors and Total Sectors Fields**

The Relative Sectors field represents the offset from the beginning of the disk to the beginning of the volume, counting by sectors, for the volume described by the partition table entry. The Total Sectors field represents the total number of sectors in the volume.

Using the Relative Sectors and Total Sectors fields (resulting in a 32-bit number) provides eight more bits than the CHS scheme to represent the total number of sectors. This allows you to create partitions that contain up to 232 sectors. With a standard sector size of 512 bytes, the 32 bits used to represent the Relative Sectors and Total Sectors fields translates into a maximum partition size of 2 terabytes (or 2,199,023,255,552 bytes).

**Note**

- For more information about the maximum partition size that each file system supports, see "NTFS Technical Reference."

The following figure shows the MBR, partition table, and boot sectors on a basic disk with four partitions. The definitions of the fields in the partition table and the extended partition tables are the same.

**Detail of a Basic Disk with Four Partitions**




**Extended Boot Record on Basic Disks**

An extended boot record (EBR), which consists of an extended partition table and the signature word for the sector, exists for each logical drive in the extended partition. It contains the only information on the first side of the first cylinder of each logical drive in the extended partition. The boot sector in a logical drive is usually located at either Relative Sector 32 or 63. However, if there is no extended partition on a disk, there are no EBRs and no logical drives.

The first entry in an extended partition table for the first logical drive points to its own boot sector. The second entry points to the EBR of the next logical drive. If no further logical drives exist, the second entry is not used and is recorded as a series of zeros. If there are additional logical drives, the first entry of the extended partition table for the second logical drive points to its own boot sector. The second entry of the extended partition table for the second logical drive points to the EBR of the next logical drive. The third and fourth entries of an extended partition table are never used.

As shown in the following figure, the EBRs of the logical drives in the extended partition are a linked list. The figure shows three logical drives on an extended partition, illustrating the difference in extended partition tables between preceding logical drives and the last logical drive.

**Detail of an Extended Partition on a Basic Disk**



With the exception of the last logical drive on the extended partition, the format of the extended partition table, which is described in the following table, is repeated for each logical drive: the first entry identifies the logical drive's own boot sector and the second entry identifies the next logical drive's EBR. The extended partition table for the last logical drive has only its own partition entry listed. The second through fourth entries of the last extended partition table are not used.

**Contents of Extended Partition Table Entries**

| Entry | Contents |
|---|---|
| First | Information about the current logical drive in the extended partition, including the starting address for the boot sector preceding the data. |
| Second | Information about the next logical drive in the extended partition, including the address of the sector that contains the EBR for the next logical drive. If no additional logical drives exist, this field is not used. |
| Third | Not used. |
| Fourth | Not used. |

The fields in each entry of the extended partition table are identical to the MBR partition table entries. For more information about partition table fields, see the table titled "Partition Table Fields" earlier in this section.

The Relative Sectors field in an extended partition table entry shows the number of bytes that are offset from the beginning of the extended partition to the first sector in the logical drive. The number in the Total Sectors field refers to the number of sectors that make up the logical drive. The value of the Total Sectors field equals the number of sectors from the boot sector defined by theextended partition table entry to the end of the logical drive.

Because of the importance of the MBR and EBR sectors, it is recommended that you run disk-scanning tools regularly and that you regularly back up all your data files to protect against losing access to a volume or an entire disk.

**Disk Sectors on GPT Disks**

GUID partition table (GPT) disks use primary and backup partition structures to provide redundancy. These structures are located at the beginning and the end of the disk. GPT identifies these structures by their logical block address (LBA) rather than by their relative sectors. Using this scheme, sectors on a disk are numbered from 0 to *n*-1, where *n* is the number of sectors on the disk.

As shown in the following figure, the first structure on a GPT disk is the Protective MBR in LBA 0, followed by the primary GPT header in LBA 1. The GPT header is followed by the primary GUID partition entry array, which includes a partition entry for each partition on the disk.

Partitions on the disk are located between the primary and backup GUID partition entry arrays. The partitions must be placed within the first usable and last usable LBAs, as specified in the GPT partition header.

**Partition Structures on a GPT Disk**




**Protective MBR**

The Extensible Firmware Interface (EFI) specification requires that LBA 0 be reserved for compatibility code and a Protective MBR. The Protective MBR has the same format as an existing MBR, and it contains one partition entry with a System ID value of 0xEE. This entry reserves the entire space of the disk, including the space used by the GPT header, as a single partition. The Protective MBR is included to prevent disk utilities that were designed for MBR disks from interpreting the disk as having available space and overwriting GPT partitions. The Protective MBR is ignored by EFI; no MBR code is run.

The following example shows a partial printout of a Protective MBR.

```
000001B0: 00 00 00 00 00 00 00 00 - 04 06 04 06 00 00 00 00   ................
000001C0: 02 00 EE FF FF FF 01 00 - 00 00 FF FF FF FF 00 00   ................
000001D0: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
000001E0: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
000001F0: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 55 AA   ..............U.
```

The following table describes the fields in each entry in the Protective MBR.

**Protective MBR in GPT Disks**

| Byte Offset | Field Length | Sample Value1 | Field Name and Definition |
|---|---|---|---|
| 0x01BE | 1 byte | 00 | **Boot Indicator**. Must be set to 00 to indicate that this partition cannot be booted. |
| 0x01BF | 1 byte | 00 | **Starting Head**. Matches the Starting LBA of the single partition. |
| 0x01C0 | 1 byte | 02 | **Starting Sector**. Matches the Starting LBA of the single partition. |
| 0x01C1 | 1 byte | 00 | **Starting Cylinder**. Matches the Starting LBA of the GPT partition. |
| 0x01C2 | 1 byte | EE | **System ID**. Must be EE to specify that the single partition is a GPT partition. If you move a GPT disk to a computer running Windows 2000 with Service Pack 1 or greater or Windows Server 2003, the partition is displayed as a **GPT Protective Partition** and cannot be deleted. |

| | | | |
|---|---|---|---|
| 0x01C3 | 1 byte | FF | **Ending Head**. Matches the Ending LBA of the single partition. If the Ending LBA is too large to be represented here, this field is set to FF. |
| 0x01C4 | 1 byte | FF | **Ending Sector**. Matches the Ending LBA of the single partition. If the Ending LBA is too large to be represented here, this field is set to FF. |
| 0x01C5 | 1 byte | FF | **Ending Cylinder**. Matches the Ending LBA of the single partition. If the Ending LBA is too large to be represented here, this field is set to FF. |
| 0x01C6 | 4 bytes | 01 00 00 00 | **Starting LBA**. Always set to 1. The Starting LBA begins at the GPT partition table header, which is located at LBA 1. |
| 0x01CA | 4 bytes | FF FF FF FF | **Size in LBA**. The size of the single partition. Must be set to FF FF FF FF if this value is too large to be represented here. |

1. Numbers larger than one byte are stored in little endian format, or reverse-byte ordering. Little endian format is a method of storing a number so that the least significant byte appears first in the hexadecimal number notation.

**GPT Partition Table Header**

The GPT header defines the range of logical block addresses that are usable by partition entries. The GPT header also defines its location on the disk, its GUID, and a 32-bit cyclic redundancy check (CRC32) checksum that is used to verify the integrity of the GPT header.

GPT disks use a primary and a backup GUID partition table (GPT) header:

- The primary GPT header is located at LBA 1, directly after the Protective MBR.
- The backup GPT header is located in the last sector of the disk. No data follows the backup GPT header.

EFI verifies the integrity of the GPT headers by using a CRC32 checksum, which is a calculated value that is used to test data for the presence of errors. If the primary GPT header is corrupted, the system checks the backup GPT header checksum. If the backup checksum is valid, then the backup GPT header is used to restore the primary GPT header. This restoration process works in reverse if the primary GPT header is valid but the backup GPT header is corrupted. If both the primary and backup GPT headers are corrupted, then the 64-bit versions of Windows Server 2003 cannot access the disk.

**Note**

- Do not use disk editing tools such as DiskProbe to make changes to GPT disks because any change that you make renders the checksums invalid, which might cause the disk to become inaccessible. To make changes to GPT disks, do either of the following:
- Use Diskpart.efi in the firmware environment.
- Use Diskpart.exe or Disk Management in the 64-bit versions of Windows Server 2003.

The following example shows a partial printout of a GPT header.

```
00000000: 45 46 49 20 50 41 52 54 - 00 00 01 00 5C 00 00 00   EFI PART....\...
00000010: 27 6D 9F C9 00 00 00 00 - 01 00 00 00 00 00 00 00   'm..............
00000020: 37 C8 11 01 00 00 00 00 - 22 00 00 00 00 00 00 00   7.......".......
00000030: 17 C8 11 01 00 00 00 00 - 00 A2 DA 98 9F 79 C0 01   .............y..
00000040: A1 F4 04 62 2F D5 EC 6D - 02 00 00 00 00 00 00 00   ...b/..m........
00000050: 80 00 00 00 80 00 00 00 - 27 C3 F3 85 00 00 00 00   ........'.......
00000060: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
```

The following table describes the fields in the GPT header.

**GUID Partition Table Header**

| Byte Offset | Field Length | Sample Value1 | Field Name and Definition |
|---|---|---|---|

| | | | |
|---|---|---|---|
| 0x00 | 8 bytes | 45 46 49 20 50 41 52 54 | **Signature**. Used to identify all EFI-compatible GPT headers. The value must always be 45 46 49 20 50 41 52 54. |
| 0x08 | 4 bytes | 00 00 01 00 | **Revision**. The revision number of the EFI specification to which the GPT header complies. For version 1.0, the value is 00 00 01 00. |
| 0x0C | 4 bytes | 5C 00 00 00 | **Header Size**. The size, in bytes, of the GPT header. The size is always 5C 00 00 00 or 92 bytes. The remaining bytes in LBA 1 are reserved. |
| 0x10 | 4 bytes | 27 6D 9F C9 | **CRC32 Checksum**. Used to verify the integrity of the GPT header. The 32-bit cyclic redundancy check (CRC) algorithm is used to perform this calculation. |
| 0x14 | 4 bytes | 00 00 00 00 | **Reserved**. Must be 0. |
| 0x18 | 8 bytes | 01 00 00 00 00 00 00 00 | **Primary LBA**. The LBA that contains the primary GPT header. The value is always equal to LBA 1. |
| 0x20 | 8 bytes | 37 C8 11 01 00 00 00 00 | **Backup LBA**. The LBA address of the backup GPT header. This value is always equal to the last LBA on the disk. |
| 0x28 | 8 bytes | 22 00 00 00 00 00 00 00 | **First Usable LBA**. The first usable LBA that can be contained in a GUID partition entry. In other words, the first partition begins at this LBA. In the 64-bit versions of Windows Server 2003, this number is always LBA 34. |
| 0x30 | 8 bytes | 17 C8 11 01 00 00 00 00 | **Last Usable LBA**. The last usable LBA that can be contained in a GUID partition entry. |
| 0x38 | 16 bytes | 00 A2 DA 98 9F 79 C0 01 A1 F4 04 62 2F D5 EC 6D | **Disk GUID**. A unique number that identifies the partition table header and the disk itself. |
| 0x48 | 8 bytes | 02 00 00 00 00 00 00 00 | **Partition Entry LBA**. The starting LBA of the GUID partition entry array. This number is always LBA 2. |
| 0x50 | 4 bytes | 80 00 00 00 | **Number of Partition Entries**. The maximum number of partition entries that can be contained in the GUID partition entry array. In the 64-bit versions of Windows Server 2003, this number is equal to 128. |
| 0x54 | 4 bytes | 80 00 00 00 | **Size of Partition Entry**. The size, in bytes, of each partition entry in the GUID partition entry array. Each partition entry is 128 bytes. |
| 0x58 | 4 bytes | 27 C3 F3 85 | **Partition Entry Array CRC32**. Used to verify the integrity of the GUID partition entry array. The 32-bit CRC algorithm is used to perform this calculation. |
| 0x5C | 420 bytes | | **Reserved**. Must be 0. |

1. Numbers larger than one byte are stored in little endian format, or reverse-byte ordering. Little endian format is a method of storing a number so that the least significant byte appears first in the hexadecimal number notation.

**GUID Partition Entry Array**

Similar to the partition table on MBR disks, the GUID partition entry array contains partition entries that represent each partition on the disk. The 64-bit versions of Windows Server 2003 create an array that is 16,384 bytes, so the first usable block must start at an LBA greater than or equal to 34. (LBA 0 contains the protective MBR; LBA 1 contains the GPT header; and LBAs 2 through 33 are used by the GUID partition entry array.)

Each GPT disk contains two GUID partition entry arrays:

- The primary GUID partition entry array is located after the GUID partition table header and ends before the first usable LBA.

- The backup GUID partition entry array is located after the last usable LBA and ends before the backup GUID partition table header.

A CRC32 checksum of the GUID partition entry array is stored in the GPT header. When a new partition is added, this checksum is updated in the primary and backup GUID partition entries, and then the GPT header size checksum is updated.

**GUID Partition Entry**

A GUID partition entry defines a single partition and is 128 bytes long. Because the 64-bit versions of Windows Server 2003 create a GUID partition entry array that has 16,384 bytes, you can have a maximum of 128 partitions on a basic GPT disk.

Each GUID partition entry begins with a partition type GUID. The 16-byte partition type GUID, which is similar to a System ID in the partition table of an MBR disk, identifies the type of data that the partition contains and identifies how the partition is used. The 64-bit versions of Windows Server 2003 recognize only the partition type GUIDs described in the following table, and do not mount any other type of partition. However, original equipment manufacturers (OEMs) and independent software vendors (ISVs), as well as other operating systems might define additional partition type GUIDs.

**Partition Type GUIDs**

| Partition Type | GUID Value |
|---|---|
| Unused entry | {00000000–0000–0000–0000–000000000000} |
| EFI System partition | {28732AC1–1FF8–D211–BA4B–00A0C93EC93B} |
| Microsoft Reserved partition | {16E3C9E3–5C0B–B84D–817D–F92DF00215AE} |
| Primary partition on a basic disk | {A2A0D0EB–E5B9–3344–87C0–68B6B72699C7} |
| LDM Metadata partition on a dynamic disk | {AAC80858–8F7E–E042–85D2–E1E90434CFB3} |
| LDM Data partition on a dynamic disk | {A0609BAF–3114–624F–BC68–3311714A69AD} |

The following example illustrates a partial hexadecimal printout of the GUID partition entry array on a basic GPT disk. This printout shows three partition entries: an EFI System partition, a Microsoft Reserved partition, and a primary partition. The partition type GUIDs are bold and match the entries in the previous table.

```
00000000: 28 73 2A C1 1F F8 D2 11 - BA 4B 00 A0 C9 3E C9 3B   (s*......K...>.;
00000010: C0 94 77 FC 43 86 C0 01 - 92 E0 3C 77 2E 43 AC 40   ..w.C.....<w.C.@
00000020: 3F 00 00 00 00 00 00 00 - CC 2F 03 00 00 00 00 00   ?........./......
00000030: 00 00 00 00 00 00 00 00 - 45 00 46 00 49 00 20 00   ........E.F.I. .
00000040: 73 00 79 00 73 00 74 00 - 65 00 6D 00 20 00 70 00   s.y.s.t.e.m. .p.
00000050: 61 00 72 00 74 00 69 00 - 74 00 69 00 6F 00 6E 00   a.r.t.i.t.i.o.n.
00000060: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
00000070: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
00000080: 16 E3 C9 E3 5C 0B B8 4D - 81 7D F9 2D F0 02 15 AE   ....\..M.}.-....
00000090: 80 BC 80 FC 43 86 C0 01 - 50 7B 9E 5F 80 78 F5 31   ....C...P{._.x.1
000000A0: CD 2F 03 00 00 00 00 00 - D0 2A 04 00 00 00 00 00   ./.......*......
000000B0: 00 00 00 00 00 00 00 00 - 4D 00 69 00 63 00 72 00   ........M.i.c.r.
000000C0: 6F 00 73 00 6F 00 66 00 - 74 00 20 00 72 00 65 00   o.s.o.f.t. .r.e.
000000D0: 73 00 65 00 72 00 76 00 - 65 00 64 00 20 00 70 00   s.e.r.v.e.d. .p.
000000E0: 61 00 72 00 74 00 69 00 - 74 00 69 00 6F 00 6E 00   a.r.t.i.t.i.o.n.
000000F0: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
00000100: A2 A0 D0 EB E5 B9 33 44 - 87 C0 68 B6 B7 26 99 C7   ......3D..h..&..
00000110: C0 1B 0B 00 44 86 C0 01 - F1 B3 12 71 4F 75 88 21   ....D......qOu.!
00000120: D1 2A 04 00 00 00 00 00 - 4E 2F 81 00 00 00 00 00   .*......N/......
00000130: 00 00 00 00 00 00 00 00 - 42 00 61 00 73 00 69 00   ........B.a.s.i.
00000140: 63 00 20 00 64 00 61 00 - 74 00 61 00 20 00 70 00   c. .d.a.t.a. .p.
00000150: 61 00 72 00 74 00 69 00 - 74 00 69 00 6F 00 6E 00   a.r.t.i.t.i.o.n.
00000160: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
00000170: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
```

The following table illustrates the layout of a GUID partition entry. The sample values correspond to the EFI System partition entry in the preceding example.

**GUID Partition Entry**

| Byte Offset | Field Length | Sample Value1 | Field Name and Definition |
|---|---|---|---|

| | | | |
|---|---|---|---|
| 0x00 | 16 bytes | 28 73 2A C1 1F F8 D2 11 BA 4B 00 A0 C9 3E C9 3B | **Partition Type GUID**. Identifies the type of partition. The partition type GUID in this example identifies Microsoft Reserved partitions. For a description of partition type GUIDs, see the table titled Partition Type GUIDs earlier in this section. |
| 0x 10 | 16 bytes | C0 94 77 FC 43 86 C0 01 92 E0 3C 77 2E 43 AC 40 | **Unique Partition GUID**. A unique ID created for each partition entry. |
| 0x 20 | 8 bytes | 3F 00 00 00 00 00 00 00 | **Starting LBA**. The starting LBA of the partition that is defined by this partition entry. |
| 0x 28 | 8 bytes | CC 2F 03 00 00 00 00 00 | **Ending LBA**. The ending LBA of the partition that is defined by this partition entry. |
| 0x 30 | 8 bytes | 00 00 00 00 00 00 00 00 | **Attribute Bits**. Describe how the partition is used. For a description of the attribute used by the 64-bit versions of Windows Server 2003, see the table titled "GUID Partition Entry Attributes Used by the 64-Bit Editions of Windows on Itanium-based Computers" later in this section. |
| 0x 38 | 72 bytes | EFI system partition | **Partition Name**. A 36-character Unicode string that can be used to name the partition. |

1. Numbers larger than one byte are stored in little endian format, or reverse-byte ordering. Little endian format is a method of storing a number so that the least significant byte appears first in the hexadecimal number notation.

#### GUID Partition Entry Attributes

GUID partition entry attributes are descriptors for how a partition is used. The attributes are specified within a 64-bit value, so EFI supports up to 64 different attributes. The 64-bit versions of Windows Server 2003 use two attributes as described in the following table.

**GUID Partition Entry Attributes Used by the 64-Bit Editions of Windows on Itanium-based Computers**

| Bits | Description |
|---|---|
| Bit 0 | Specifies that this partition is required for the platform to function. All original equipment manufacturer (OEM) partitions must have this bit set to protect the OEM partition from being overwritten by the disk tools supplied with Windows Server 2003. |
| Bit 60 | Marks the partition as read-only. Used only for primary basic partitions of type {EBD0A0A2-B9E5-4433-87C0-68B6B72699C7}. |
| Bit 62 | Marks the partition as hidden. Used only for primary basic partitions of type {EBD0A0A2-B9E5-4433-87C0-68B6B72699C7}. |
| Bit 63 | Prevents the system from assigning a default drive letter to the partition. Used only for primary basic partitions of type {EBD0A0A2-B9E5-4433-87C0-68B6B72699C7}. |

#### Boot Sectors on GPT Disks

Boot sectors on GPT disks are similar to boot sectors on MBR disks, except that EFI ignores all x86 code in the boot sector. Instead, EFI uses its own file system driver to read the BIOS parameter block (BPB) and then mount the volume.

#### Basic Disks and Volumes Processes and Interactions

The following basic disk and volume processes and interactions assume that your computer has at least one basic disk and that the basic disk is functioning properly.

#### Creating a Basic Volume

When you create a basic volume, the Virtual Disk Service (VDS) uses the IOCTL, IOCTL_DISK_SET_DRIVE_LAYOUT_EX, to set the drive layout, and add a new partition.

VDS calls the disk driver. Partition manager (Partmgr.sys), sits on top of the disk driver as a filter). Partition manager creates the partition device and notifies the volume manager that there is a new partition. The volume manager announces the volume device to Plug and Play and to the system. Plug and Play notifies the mount point manager (Mountmgr.sys) that a new volume

has arrived and the mount point manager sets up a drive letter as long as AutoMount is enabled.

**Note**

- AutoMount is disabled by default on Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

### What Happens When a Basic Volume Is Created




**Reading or Writing to a Basic Volume**

When you read or write data to basic disk or volume, the input/output (I/O) is sent through the file system to the volume manager. The volume manager creates a new I/O request packet (IRP) and sends it to the partition manager (Partmgr.sys) in the disk stack, which relays it to the disk driver. The volume manager waits for the completed IRP to return from the disk stack. Then, the volume manager completes and returns the file system IRP.

The following figure illustrates the processes involved when reading or writing to a basic volume.

**Basic Disk and Volume Input/Output (I/O) Processes**




**Converting a Basic Disk into a Dynamic Disk**

You can convert a basic disk to dynamic by using Disk Management or by using DiskPart, a command-line tool that provides the same functions as Disk Management.

When you convert a disk to dynamic, the following events occur:

- All existing primary partitions and logical drives become simple volumes.
- The disk joins the local disk group and receives a copy of the dynamic disk database.

**Note**

- For certain disks, the menu command to convert the disk to dynamic is unavailable in Disk Management.

You can convert basic disks to dynamic at any time. In most cases, you do not need to restart your computer to complete the conversion. However, you must restart the computer if the disks you are converting contain any of the following volumes:

- **System volume (x86-based computers only).** The system volume contains hardware-specific files such as Ntldr and Boot.ini. These files are needed to load Windows Server 2003 in x86-based computers.
- **Boot volume.** The boot volume contains the Windows Server 2003 operating system and its support files. In x86-based computers, the boot volume can be, but does not have to be, the same volume as the system volume. In Itanium-based computers, the boot volume is never the same volume as the EFI System partition.
- **Volumes that contain the paging file.** The paging file is a hidden file on the hard disk that Windows Server 2003 uses to hold parts of programs and data files that do not fit in memory.

**Note**

- When you convert MBR disks that contain the system, boot, or paging file volumes to dynamic, you are prompted to

restart the computer two times. You must restart the computer both times to complete the conversion.

As shown in the following figure the Disk Management snap-in identifies the system and boot volumes, as well as those that contain the paging file in the graphical and disk list views. If you have a combined system and boot volume that also contains the paging file (the most common scenario), then only **System** is shown.

**How Disk Management Identifies Separate System, Boot, and Paging File Volumes for an x86-based Computer**




The list volume command in DiskPart shows the system, boot, and paging file volumes as follows:

```
Volume ###   Ltr   Label        Fs     Type        Size        Status     Info
----------   ---   -----------  -----  ----------  ---------   ---------  --------
Volume 0      G                       DVD-ROM           0 B   Healthy
Volume 1      C                 NTFS  Partition      2048 MB   Healthy    System
Volume 2      E                 NTFS  Partition     17418 MB   Healthy
Volume 3      F                 NTFS  Partition      4003 MB   Healthy    Boot
Volume 4      H                 NTFS  Partition       751 MB   Healthy    Pagefile
```

Even after you convert a disk to dynamic, some types of primary partitions do not become dynamic volumes. These partitions retain their partition entries in the partition table and are shown as primary partitions in Disk Management. These partitions are:

- Known OEM partitions (usually displayed in Disk Management as EISA Configuration partitions).
- EFI System partitions on GPT disks.

**Before Converting Disks to Dynamic**

Converting a disk to dynamic changes the partition layout on the disk and creates the dynamic disk database. The result of these changes is increased flexibility for volume management in Windows Server 2003. However, these changes are not easily reversed, and the structure of dynamic disks is not compatible with some operating systems. Therefore, you must consider the following issues before you convert disks to dynamic.

**If the disk contains partitions displayed as Healthy (Unknown) in Disk Management**

Do not convert a disk to dynamic if it contains unknown partitions created by other operating systems. Windows Server 2003

converts unrecognized partitions to dynamic, making them unreadable to other operating systems.

**If a disk contains shadow copies**
If you use separate disks to store the source volume and the shadow copies, and you convert the disk that contains the source volume to dynamic, the shadow copies are lost. Shadow copies are retained only if the source files and the shadow copies are stored on the same volume.

**If the disk contains an OEM partition that is not at the beginning of the disk**
Do not convert a disk to dynamic if it contains an OEM partition that is not at the beginning of the disk. (In Disk Management, an OEM partition usually appears as an EISA Configuration partition.) When you convert a disk to dynamic, Windows Server 2003 preserves the OEM partition only if it is the first partition on the disk. Otherwise the partition is deleted during the conversion to dynamic.

**If you want to extend a dynamic volume**
You can extend dynamic volumes that do not retain their partition entries in the partition table. The following volumes retain their entries in the partition table and cannot be extended:

- The system volume and boot volume of the operating system that you used to convert the disk to dynamic.
- Any basic volume that was present on the disk when you converted the disk from basic to dynamic by using the version of Disk Management included with Windows 2000.
- Simple volumes on which you run the DiskPart command **retain**. This command adds a partition entry to the partition table. However, after you use this command, you can no longer extend the volume.

**Note**

- The **retain** command adds an entry to the partition table of an MBR disk only for simple volumes that are contiguous, start at cylinder-aligned offsets, and are an integral number of cylinders in size. If a volume does not meet these requirements, the **retain** command fails.
- The following examples describe volumes on which the **retain** command will succeed:
- The simple volume is contiguous and starts at the beginning of the disk.
- The simple volume was present on the disk when the disk was converted to dynamic.

The only way to add more space to the system or boot volume on a dynamic disk is to back up all data on the disk, repartition and reformat the disk, reinstall Windows Server 2003, convert the disks to dynamic, and then restore the data from backup.

The following volumes do not have partition entries and can be extended:

- Simple volumes and spanned volumes created from unallocated space on a dynamic disk.
- A basic volume that meets the following criteria:
    - The basic volume is not the system or boot volume.
    - The basic volume is on a disk that was converted from basic disk to dynamic disk by using Windows Server 2003.

Although striped, mirrored, or RAID-5 volumes do not have entries in the partition table, Windows Server 2003 does not support extending them. You can add more space to a striped, mirrored, or RAID-5 volume by backing up the data, deleting the volume, recreating the volume by using Windows Server 2003, and then restoring the data.

**If you want to install Windows Server 2003 on a dynamic volume**
You can install Windows Server 2003 only on dynamic volumes that retain their partition entries in the partition table. The only dynamic volumes listed in the partition table are the following:

- The system volume and boot volume of the operating system (Windows Server 2003 or Windows 2000) that you used to convert the disk to dynamic. The system volume and boot volume can be simple or mirrored volumes.

- Any basic volume that was present on the disk when you used Windows 2000 to convert the disk from basic to dynamic.

- Simple volumes on which you run the DiskPart command **retain**. This command adds a partition entry to the partition table so that you can install Windows Server 2003 on the simple volume.

- A basic mirror set that was converted to a dynamic mirrored volume by using Windows 2000. If you break this mirrored volume into two simple volumes, you can also install Windows Server 2003 on either simple volume because they both retain their partition entries.

Because these dynamic volumes retain their partition entries, you can install Windows Server 2003 on them. However, you cannot extend any of these volumes because you can only extend volumes that do not have entries in the partition table.

**If you want to access the disk by using Windows Millennium Edition or earlier, or Windows NT 4.0**
If you plan to move the disk after you convert it to dynamic, note that you can access dynamic disks only from computers that are running Windows 2000, Windows XP Professional, Windows XP 64-Bit Edition, or Windows Server 2003. You cannot access dynamic disks from computers running Windows NT 4.0 or earlier.

When moving disks, note that access to dynamic disks is further restricted by the partition style used on the dynamic disk:

- **Dynamic MBR disks**. Only computers running Windows 2000, Windows XP Professional, Windows XP 64-Bit Edition, or Windows Server 2003 can access dynamic MBR disks.

- **Dynamic GPT disks**. Only Itanium-based computers running Windows XP 64-Bit Edition or the 64-bit versions of Windows Server 2003 can access dynamic GPT disks.

### Note

- Volumes on dynamic MBR and GPT disks are available across a network to computers running MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0 or earlier, Windows XP, and Windows Server 2003.

**If a disk or computer contains multiple copies of Windows XP Professional, Windows Server 2003 or Windows 2000**
Do not convert a disk to dynamic if it contains multiple copies of Windows XP Professional, Windows Server 2003, or Windows 2000. Even though these operating systems support dynamic disks, they require certain registry entries that allow them to start from dynamic disks. If the operating systems are installed on the same disk and you use one of the operating systems to convert the disk to dynamic, the registry of the other operating system becomes out-of-date because the drivers required to start the operating system from a dynamic disk are not loaded. Therefore, you can no longer start the other operating system.

One way that you can use dynamic disks with Windows XP Professional, Windows Server 2003, and Windows 2000 in a multiple-boot configuration is to install each operating system to a different disk. However, startup problems can also occur if you boot from one of the operating systems and then convert the disks that contain the other operating systems to dynamic.

To ensure that each operating system can start, start each operating system and then convert only the disk that contains the current boot volume to dynamic. For example, install Windows XP Professional on disk 1 and Windows Server 2003 on disk 2. Use Windows XP Professional to convert disk 1 to dynamic, and then use Windows Server 2003 to convert disk 2 to dynamic. By using this method, you ensure that the registries are updated for each boot volume.

**Disks That Cannot Be Converted to Dynamic**
Windows Server 2003 Setup and Disk Management ensure that disks initialized by Windows Server 2003 can be converted to dynamic. However, on some disks the conversion fails or the **Convert to Dynamic Disk** command is not available when you right-click a basic disk. The following conditions prevent you from converting a basic disk to dynamic.

**Cluster disks**
You cannot convert cluster disks to dynamic if they are connected to shared SCSI or Fibre Channel buses. Windows Cluster service cannot read disks that are dynamic and makes dynamic disks unavailable to programs or services that are dependent on these disk resources in the server cluster. For this reason, the option to convert these disks to dynamic is unavailable.

You must use Veritas Volume Manager to use dynamic disks with Cluster service.

**Removable disks**

You cannot use dynamic disks on the following:

- Removable media, such as Iomega Zip or Jaz disks, CDs, DVDs, or floppy disks.
- Disks that use universal serial bus (USB) or IEEE 1394 (also called FireWire) interfaces.

**Disks with sectors larger than 512 bytes**
A sector is a unit of storage on a hard disk. The majority of hard disks use 512-byte sectors. Windows Server 2003 supports converting basic disks to dynamic only if the sector size of the basic disk is 512 bytes.

**GPT disks with non-contiguous partitions**
If an unknown partition lies between two known partitions on a GPT disk, you cannot convert the disk to dynamic. Unknown partitions are created by operating systems or utilities that use partition type GUIDs that the 64-bit versions of Windows Server 2003 do not recognize.

**MBR disks that do not have space for the dynamic disk database**
An MBR disk requires 1 MB of free space at the end of the disk to be used for the dynamic disk database. Windows Server 2003 and Windows 2000 automatically reserve 1 MB or one cylinder, whichever is greater, when creating partitions on a disk, but in rare cases, disks with partitions created by other operating systems might not have this space available. If this space is not available, you cannot convert the disk to dynamic.

To convert the disk to dynamic, you must back up or move the data, delete the partitions, recreate the partitions, restore the data, and then convert the disk to dynamic. By using Windows Server 2003 to create the partitions, you ensure that the necessary space is available for the dynamic disk database.

This limitation does not affect GPT disks because the database is created in its own partition with space borrowed from the Microsoft Reserved partition.

# Basic Disks and Volumes Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Basic Disk and Volume Tools
- Basic Disk and Volume Registry Entries
- Basic Disk and Volume WMI Classes


The following tools, registry settings, and Windows Management Instrumentation (WMI) classes are associated with basic disks and volumes.

**Basic Disk and Volume Tools**
The following tools are associated with basic disks and volumes.

**Bootcfg.exe: Boot Configuration Tool**
**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
All versions of the Boot Configuration Tool in the Windows Server 2003 family are identical.

You can use the Boot Configuration Tool to query, configure, or change settings in the Boot.ini file on your computer.

**Dmdiag.exe: Disk Manager Diagnostics**
**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
All versions of Dmdiag in the Windows Server 2003 family are identical.

Dmdiag displays the following information for the computer on which it is run:

- Computer name and operating system version
- Physical disk to disk type
- Mount points
- Logical Disk Manager (LDM) file versions
- Drive letter usage
- List of devices
- Symbolic links
- Disk partition information

**Dskmgmt.msc: Disk Management**
**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
All versions of the Disk Management snap-in in the Windows Server 2003 family are identical.

You can use the Disk Management snap-in to remotely manage disks and volumes on other computers running Windows 2000, Windows XP or Windows Server 2003

**DiskPart.exe: DiskPart**
**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
All versions of DiskPart in the Windows Server 2003 family are identical. DiskPart includes a few 64-bit parameters that are only available on Itanium-based computers.

DiskPart is a text-mode command interpreter that enables you to manage objects (disks, partitions, or volumes) by using scripts or direct input from a command prompt.

**Format.exe: Format**
**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
All versions of Format in the Windows Server 2003 family are identical.

Format prepares a volume on the specified disk to accept Windows files.

**FTonline.exe: Fault-tolerant Disk Mounter**
**Category**
Windows Server 2003 support tool.

**Version compatibility**
FTonline enables an administrator to mount and recover files from fault-tolerant disks created in previous versions of Windows. This tool is useful if you did not upgrade your disks to dynamic disks, or failed to back up your data before installing Windows Server 2003. You can install FTonline using the Support Tools setup program located in the \Support\Tools folder on the Windows XP Professional and the Windows Server 2003 family of operating systems CDs.

**Mountvol.exe: Mountvol**
**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
All versions of Mountvol in the Windows Server 2003 family are identical. Mountvol includes a few 64-bit parameters that are only available on Itanium-based computers.

Mountvol creates, deletes, or lists volume mount points. Mountvol enables you to link volumes without using drive letters.

**SecInspect.exe: Sector Inspector**
**Category**
Windows Server 2003 command-line tool.

**Version compatibility**
SecInspect is a command-line diagnostics tool that enables administrators to view the contents of master boot records, boot sectors, and IA64 GUID partition tables. Additional features include creating hexadecimal dumps of binary files and backup/restore of sector ranges. For more information about this tool, see the Help that comes with the tool. To find this tool, see Tool Updates in Tools and Settings Collection.

**Basic Disk and Volume Registry Entries**
The following registry entries are associated with basic disks and volumes. For more information about the registry, see the Registry Reference in Tools and Settings Collection.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following sections describe the basic disk and volume registry entries that are listed below \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\. The headings represent the next level in the path after \Services. For example, Vds\Debuglog is equal to \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Vds\Debuglog.

**Vds\Debuglog**
**Level**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Vds\Debuglog\Level

**Version**
Windows Server 2003 family.

Sets the level of logging for VDS.

Use Registry Editor to add a decimal value named Level. Stop and restart the service after changing.

### VDS Log Options

| Item Logged | Bitmask Number |
|---|---|
| Errors | 1 |
| Warnings | 2 |
| Trace | 4 |
| Information | 8 |

Set the value by adding the bitmask numbers of the values you want to log. For example, decimal 3 logs errors (1) and warnings (2). Decimal 9 logs errors (1) and information (8). The range for this value is 0-15. Set the value to 0xF (decimal 15) to log all items.

**Dmadmin\Parameters**
**EnableDynamicConversionFor1394**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\dmadmin\Parameters\EnableDynamicConversionFor1394

**Version**
Windows Server 2003 family.

Allows a user to convert an Institute of Electrical and Electronics Engineers, Inc. (IEEE) 1394 (FireWire) disk to a dynamic disk on Windows Server 2003 or earlier operating systems. Converting 1394 disks to dynamic is not a tested or supported scenario. This registry entry is provided for compatibility with a Windows 2000 beta that supported this functionality.

EnableDynamicConversionFor1394 is a DWORD Value, with a range of 0-1. Default is zero. To modify, use Registry Editor. Stop and restart the service after modifying.

**Basic Disk and Volume WMI Classes**
The following table lists and describes the WMI classes that are associated with basic disks and volumes.

### WMI Classes Associated with Basic Disks and Volumes

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_DiskDrive | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_DiskDrivePhyiscalMedia | \root\cimv2 | Windows Server 2003 family |
| Win32_DiskDriveToDiskPartition | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_DiskPartition | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_LogicalDisk | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_LogicalDiskToPartition | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_MappedLogicalDisk | \root\cimv2 | Windows Server 2003 family |
| Win32_PhysicalMedia | \root\cimv2 | Windows Server 2003 family |
| Win32_Volume | \root\cimv2 | Windows Server 2003 family |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

# Dynamic Disks and Volumes Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Like basic disks, which are the most commonly used storage type found on computers running Microsoft Windows, dynamic disks can use the master boot record (MBR) or GUID partition table (GPT) partitioning scheme. All volumes on dynamic disks are known as dynamic volumes. Dynamic disks were first introduced with Windows 2000 and provide features that basic disks do not, such as the ability to create volumes that span multiple disks (spanned and striped volumes), and the ability to create fault tolerant volumes (mirrored and RAID-5 volumes).

Dynamic disks offer greater flexibility for volume management because they use a database to track information about dynamic volumes on the disk and about other dynamic disks in the computer. Because each dynamic disk in a computer stores a replica of the dynamic disk database, Windows Server 2003can repair a corrupted database on one dynamic disk by using the database on another dynamic disk.

The following topics describe how dynamic disks and volumes work, their architecture, and the tools you can use when working with them:

**In this subject**


- What Are Dynamic Disks and Volumes?
- How Dynamic Disks and Volumes Work
- Dynamic Disks and Volumes Tools and Settings

# What Are Dynamic Disks and Volumes?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Are Dynamic Disks and Volumes?**

**In this section**


- Types of Dynamic Volumes
- Dynamic Disk and Volume Scenarios


Like basic disks, which are the most commonly used storage type found on computers running Microsoft Windows, dynamic disks can use the master boot record (MBR) or GUID partition table (GPT) partitioning scheme. All volumes on dynamic disks are known as dynamic volumes. Dynamic disks were first introduced with Windows 2000 and provide features that basic disks do not, such as the ability to create volumes that span multiple disks (spanned and striped volumes), and the ability to create fault tolerant volumes (mirrored and RAID-5 volumes).

Dynamic disks offer greater flexibility for volume management because they use a database to track information about dynamic volumes on the disk and about other dynamic disks in the computer. Because each dynamic disk in a computer stores a replica of the dynamic disk database, Windows Server 2003 can repair a corrupted database on one dynamic disk by using the database on another dynamic disk.

The location of the database is determined by the partition style of the disk.

- On MBR disks, the database is contained in the last 1 megabyte (MB) of the disk.
- On GPT disks, the database is contained in a 1-MB reserved (hidden) partition known as the Logical Disk Manager (LDM) Metadata partition.

All online dynamic disks in a computer must be members of the same disk group, which is a collection of dynamic disks. A computer can have only one dynamic disk group, also called the primary disk group. Each disk in a disk group stores a replica of the same dynamic disk database. A disk group uses a name consisting of the computer name plus a suffix of Dg0. The disk group name is stored in the registry.

When you move dynamic disks to a computer that has existing dynamic disks, you must import the dynamic disks to merge the databases on the moved disks with the databases on the existing dynamic disks.

The disk group name on a computer never changes, as long as the disk group contains dynamic disks. If you remove the last disk in the disk group or convert all dynamic disks to basic, the registry entry remains. However, if you create a dynamic disk again on that computer, a new disk group name is generated. The computer name in the disk group remains the same, but the suffix is Dg1 instead of Dg0.

When you move a dynamic disk to a computer that has no dynamic disks, the dynamic disk retains its disk group name and ID from the original computer and uses them on the local computer.

For more information about converting basic disks to dynamic disks, including the limitations of dynamic disks, see "Converting a Basic Disk into a Dynamic Disk" in How Basic Disks and Volumes Work.

**Types of Dynamic Volumes**

A dynamic volume is a volume that is created on a dynamic disk. Dynamic volume types include simple, spanned, and striped volumes. Windows Server 2003 also supports mirrored and RAID-5 volumes, which are fault tolerant. Fault tolerance is the ability of computer hardware or software to make sure that your data is still available, even if there is a hardware failure.

Regardless of the partition style used (MBR or GPT), you can create up to 1000 dynamic volumes per disk group, although boot time increases as the number of volumes increases. The recommended number of dynamic volumes is 32 or fewer per disk group.

The following sections describe the different types of dynamic volumes.

**Simple Volumes**

Simple volumes are the dynamic-disk equivalent of the primary partitions and logical drives found on basic disks. When creating simple volumes, keep these points in mind:

- If you have only one dynamic disk, you can create only simple volumes.
- You can increase the size of a simple volume to include unallocated space on the same disk or on a different disk. The volume must be unformatted or formatted by using NTFS. You can increase the size of a simple volume in two ways:
    - By extending the simple volume on the same disk. The volume remains a simple volume, and you can still mirror it.
    - By extending a simple volume to include unallocated space on other disks on the same computer. This creates a spanned volume.

    **Note**

If the simple volume is the system volume or the boot volume, you cannot extend it.

**Spanned Volumes**
Spanned volumes combine areas of unallocated space from multiple disks into one logical volume. The areas of unallocated space can be different sizes. Spanned volumes require two disks, and you can use up to 32 disks. When creating spanned volumes, keep these points in mind:

- You can extend only NTFS volumes or unformatted volumes.
- After you create or extend a spanned volume, you cannot delete any portion of it without deleting the entire spanned volume.
- You cannot stripe or mirror spanned volumes. For more information about striped or mirrored volumes, see “Striped Volumes” or “Mirrored Volumes” later in this section.
- Spanned volumes do not provide fault tolerance. If one of the disks containing a spanned volume fails, the entire volume fails, and all data on the spanned volume becomes inaccessible. The reliability for a spanned volume is less than the least reliable disk in the set.

**Striped Volumes**
Striped volumes improve disk input/output (I/O) performance by distributing I/O requests across disks. Striped volumes are composed of stripes of data of equal size written across each disk in the volume. They are created from equally sized, unallocated areas on two or more disks. In Windows Server 2003, the size of each stripe is 64 kilobytes (KB) and cannot be changed.

Striped volumes cannot be extended or mirrored and do not offer fault tolerance. If one of the disks containing a striped volume fails, the entire volume fails, and all data on the striped volume becomes inaccessible. The reliability for the striped volume is less than the least reliable disk in the set.

**Mirrored Volumes**
A mirrored volume is a fault-tolerant volume that provides a copy of a volume on another disk. Mirrored volumes provide data redundancy by duplicating the information contained on the volume. The two disks that make up a mirrored volume are known as mirrors. Each mirror is always located on a different disk. If one of the disks fails, the data on the failed disk becomes unavailable, but the system continues to operate by using the unaffected disk.

Mirrored volumes are available only on computers running the Windows 2000 Server family or Windows Server 2003.

**RAID-5 Volumes**
A RAID-5 volume is a fault-tolerant volume that stripes data and parity across three or more disks. Parity is a calculated value that is used to reconstruct data if one disk fails. When a disk fails, Windows Server 2003 continues to operate by recreating the data that was on the failed disk from the remaining data and parity. RAID-5 volumes are available only on computers running the Windows 2000 Server family or Windows Server 2003.

**Dynamic Disk and Volume Scenarios**
To gain the maximum benefit from dynamic disks and volumes, it is best to use them in computers with more than one disk,

which allows you to scale the storage to match your needs. Dynamic disks and volumes are commonly used in the following scenarios.

**Create a spanned volume to increase volume size**
Spanned volumes are typically created by a user who has at least two disks in their computer. If the user doesn't require the high read throughput of a striped volume or the fault tolerance offered by mirrored or RAID-5 volumes, the user can create a spanned volume.

If the data volume on a user's primary disk gets too full, the user can extend that volume onto a second disk in their computer. This enables the user to create a big volume that uses space on two disks. If necessary, they can extend the volume to cover up to 32 disks. If the user has some space left on their primary disk after performing configuration changes, they can use that space and combine it with space on the second disk, creating a spanned volume.

**Create a spanned volume to combine two or three small disks into a large volume**
If a user has two or more small disks in their computer, they can combine those disks into a single large volume by creating a spanned volume. To the user, the space spread across the disks would look and function like a single volume.

**Create a striped volume to accommodate high read/write throughput**
Striped volumes are typically created by the user who has at least two disks in their computer. If the user requires high read/write throughput but does not require the fault-tolerance offered by mirrored or RAID-5 volumes, the user can create a striped volume.

**Create a RAID-5 volume to protect critical data**
RAID-5 volumes are typically created by the user who requires fault-tolerance and who has at least three disks in their computer. If one of the disks in the RAID-5 volume fails, the data on the remaining disks, along with the parity information, can be used to recover the lost data. RAID-5 volumes are well-suited to storing data that will need to be read frequently but written to less frequently. Database applications that read randomly work well with the built-in load balancing of a RAID-5 volume.

**Create a mirrored volume to protect critical data**
Mirrored volumes are typically created by the user who requires fault-tolerance and who has two disks in their computer. If one disk fails, the user always has a copy of their data on the second disk. Mirrored volumes provide better write performance than RAID-5 volumes.

**Create a mirror to migrate data to a larger disk**
If a user has run out of room on the simple volume where they store data and there is no room left on the disk to extend the volume, they can move this data to a larger disk instead of creating a spanned volume. The user can create a mirrored volume using the simple volume containing their data and a larger disk. After creating the mirrored volume, the user can break the mirrored volume and extend the new volume to fill the larger disk, leaving a complete copy of the original volume with available space for new data. The space on the original volume can be reclaimed for other uses.

# How Dynamic Disks and Volumes Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Dynamic Disks and Volumes Work**

**In this section**


- Dynamic Disks and Volumes Architecture
- Dynamic Disk and Volume Physical Structure
- Dynamic Disk and Volume Processes and Interactions
- Types of RAID Volumes


Like basic disks which are the most commonly used storage type found on computers running Microsoft Windows, dynamic disks can use the master boot record (MBR) or GUID partition table (GPT) partitioning scheme. All volumes on dynamic disks are known as dynamic volumes. Dynamic disks were first introduced with Windows 2000 and provide features that basic disks do not, such as the ability to create volumes that span multiple disks (spanned and striped volumes), and the ability to create fault tolerant volumes (mirrored and RAID-5 volumes).

Dynamic disks offer greater flexibility for volume management because they use a database to track information about dynamic volumes on the disk and about other dynamic disks in the computer. Because each dynamic disk in a computer stores a replica of the dynamic disk database, Windows Server 2003 can repair a corrupted database on one dynamic disk by using the database on another dynamic disk.

An optimal environment for basic disks and volumes is defined as follows:

- Windows Server 2003 operating system is installed and functioning properly.
- The dynamic disks are functioning properly and they display the Online status in the Disk Management snap-in.
- The dynamic volumes display the Healthy status in the Disk Management snap-in.

The following sections provide an in-depth view of how basic disks and volumes work in an optimal environment.

**Dynamic Disk and Volume Architecture**

Dynamic disks and volumes rely on the Logical Disk Manager (LDM) and Virtual Disk Service (VDS) and their associated components. These components enable you to perform tasks such as converting basic disks into dynamic disks, and creating fault-tolerant volumes. The following diagram shows the LDM and VDS components.

**Logical Disk Manager and Virtual Disk Service Components**




The following table lists the LDM and VDS components and provides a brief description of each.

**Logical Disk Manager and Virtual Disk Service Components**

| Component | Description |
|---|---|
| Disk Management snap-in:<br><br>Dmdlgs.dll<br><br>Dmdskmgr.dll<br><br>Dmview.ocx<br><br>Diskmgmt.msc | Binaries that comprise the Disk Management snap-in user interface. |
| DiskPart command line utility:<br><br>Diskpart.exe | A scriptable alternative to the Disk Management snap-in. |
| Mount Manager command line:<br><br>Mountvol.exe | A command line utility that can be used to create, delete, or list volume mount points. |

| | |
|---|---|
| Virtual Disk Service:<br><br>Vds.exe<br><br>Vdsutil.dll | A program used to configure and maintain volume and disk storage. |
| Virtual Disk Service provider for basic disks and volumes:<br><br>Vdsbas.dll | The Virtual Disk Service calls into the basic provider when configuring basic disks and volumes. |
| Virtual Disk Service provider for dynamic disks and volumes:<br><br>Vdsdyndr.dll | The Virtual Disk Service calls into the dynamic provider when configuring dynamic disks and volumes. |
| Dmboot.sys<br><br>Dmconfig.dll<br><br>Dmintf.dll<br><br>Dmio.sys<br><br>Dmload.sys<br><br>Dmremote.exe<br><br>Dmutil.dll | Drivers and user mode components used to configure dynamic disks and volumes and perform I/O. |
| Logical Disk Administrator service:<br><br>Dmadmin.exe | The VDS provider for dynamic disks and volumes uses the interfaces exposed by this service to configure dynamic disks. |
| Logical Disk Service:<br><br>Dmserver.dll | A service that detects and monitors new hard disk drives and sends disk volume information to Logical Disk Manager Administrative Service for configuration. If this service is stopped, dynamic disk status and configuration information might become outdated. If this service is disabled, any services that explicitly depend on it will fail to start. |
| Basic disk I/O driver:<br><br>Ftdisk.sys | A driver that manages all I/O for basic disks. Other system components, such as mount point manager, call into this driver to get information about basic disk volumes. |
| Mount point manager driver:<br><br>Mountmgr.sys | A binary that tracks drive letters, folder mount paths and other mount points for volumes. Assigns a unique volume mount point of the form \??\Volume<*GUID*> to each volume, in addition to any drive letters or folder paths that have been assigned by the user. Ensures that a volume will get the same drive letter each time the computer boots, and also tries to retain a volume's drive letter when the volume's disk is moved to a new computer. |
| Partition manager:<br><br>Partmgr.sys | A filter driver that sits on top of the disk driver. All disk driver requests pass through the partition manager driver. This driver creates partition devices and notifies the volume managers of partition arrivals and removals. Exposes IOCTLs that return information about partitions to other components, and allow partition configuration. |

**Dynamic Disk and Volume Physical Structure**

Dynamic disks can use either the master boot record (MBR) or GUID partition table (GPT) partitioning style. x86-based

computers use disks with the MBR partitioning style and Itanium-based computers use disks with the GPT partitioning style.

The following diagram compares a dynamic MBR disk to a dynamic GPT disk.

**Comparison of Dynamic MBR and GPT Disks**




The following table compares dynamic MBR and GPT disks.

**Comparison of Dynamic MBR and GPT Disks**

| Characteristic | MBR Disk(x86-based Computer) | GPT Disk(Itanium-based Computer) |
|---|---|---|
| Number of volumes on dynamic disks | Supports up to 1000 volumes per disk group. | Supports up to 1000 volumes per disk group. |
| Compatible operating systems | Can be read by:<br>• Windows 2000, all versions<br>• Windows XP<br>• Windows Server 2003, all versions for x86-based computers and Itanium-based computers | Can be read by:<br>• Windows XP 64-Bit Edition<br>• The 64-bit version of Windows Server 2003, Enterprise Edition<br>• The 64-bit version of Windows Server 2003, Datacenter Edition |
| Maximum size of dynamic volumes | Supports the maximum volume size of the file system used to format the volume.<br>Up to 64 terabytes for a striped or spanned volume using 32 disks. | Supports the maximum volume size of the file system used to format the volume.<br>Up to 64 terabytes for a striped or spanned volume using 32 disks. |

| | | |
|---|---|---|
| Partition tables (copies) | Contains one copy of the partition table. | Contains primary and backup partition tables for redundancy and checksum fields for improved partition structure integrity. |
| Locations for data storage | Stores data in partitions and in unpartitioned space. Although most user and program data is stored within partitions, some system metadata might be stored in hidden or unpartitioned sectors created by OEMs or other operating systems. | Stores user and program data in partitions that are visible to the user. Stores system metadata that is critical to platform operation in partitions that the 64-bit versions of Windows Server 2003 recognize but do not make visible to the user. Does not store any data in unpartitioned space. |
| Troubleshooting methods | Uses the same methods and tools used in Windows 2000. | Uses tools designed for GPT disks. (Do not use MBR troubleshooting tools on GPT disks.) |

**Dynamic Disk and Volume Processes and Interactions**

The following section discusses the different processes that are used by dynamic disks and volumes and discusses the ways in which those processes interact. This section assumes that your computer has at least three dynamic disks and that the dynamic disk is functioning properly.

**Creating a simple volume**

Creating a simple volume involves the Virtual Disk Service (VDS) and Logical Disk Manager (LDM). A simple volume is a dynamic volume made up of disk space from a single dynamic disk. When you use Disk Management or DiskPart to create a simple volume, they call the VDS API, which sends an IOCTL to the Volume Manager to create the volume. Volume Manager creates the simple volume and documents it in the dynamic disk's database. Next, Volume Manager sends information about the new volume to Plug and Play. Plug and Play sends information about the new volume to the Mount Manager and to VDS. Mount Manager assigns a drive letter to the volume and after VDS sends the information to Disk Management or DiskPart, the volume is available and ready for use.

### What Happens When a Simple Volume Is Created




**Partition Entries on MBR Dynamic Disks**

Like basic disks, dynamic disks contain an MBR that includes the master boot code, the disk signature, and the partition table for the disk. However, the partition table on a dynamic disk does not contain an entry for each volume on the disk because volume information is stored in the dynamic disk database. Instead, the partition table contains entries for the system volume, boot volume (if it is not the same as the system volume), and one or more additional partitions that cover all the remaining unallocated space on the disk. All these partitions use System ID 0x42, which indicates that these partitions are on a dynamic disk. Placing these partitions in the partition table prevents MBR-based disk utilities from interpreting the space as available for new partitions.

### Note

- In Windows 2000, the partition entries for existing basic volumes were preserved in the partition table when the disk was

converted to dynamic. These entries prevented the converted dynamic volumes from being extended. This limitation has been removed from Windows Server 2003 for all converted volumes except the boot and system volumes. Partition entries for all other converted volumes are removed from the partition table, and therefore these volumes can be extended.

The following example shows a partial printout of an MBR on a dynamic disk that contains four simple volumes: the system volume, the boot volume, and two data volumes. Note, however, that the partition table contains entries for only three partitions. The first entry is the system volume, which is marked as active. The second entry is the boot volume, and the third entry is the container partition for the two data volumes on the disk. All entries are type 0x42, which specifies dynamic volumes.

```
000001B0:                                                 80 01   .....,Dc!.!.....
000001C0: 01 00 42 FE 7F 04 3F 00 - 00 00 86 FA 3F 00 00 00   ..B..?.....?...
000001D0: 41 05 42 FE FF 02 C5 FA - 3F 00 7E 04 7D 00 00 00   A.B.....?.~.}...
000001E0: C1 03 42 FE FF FF 43 FF - BC 00 58 53 54 00 00 00   ..B...C...XST...
000001F0: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 55 AA   ..............U.
```

**Partition Entries on Dynamic GPT Disks**

The following example illustrates a partial hexadecimal printout of a GUID partition entry array on a dynamic GPT disk. The GUID partition entry array shows the Microsoft Reserved partition plus additional entries that appear only on dynamic GPT disks:

- The LDM Metadata partition is a 1-megabyte hidden partition that stores the dynamic disk database, which contains information about all dynamic disks and volumes installed on the computer.
- The LDM Data partition acts as a container for dynamic volumes. Individual dynamic volumes do not contain entries in the GUID partition entry array.

The partition type GUIDs are bold and match the entries in the table titled “Partition Type GUIDs” later in this section.

```
00000000: 16 E3 C9 E3 5C 0B B8 4D - 81 7D F9 2D F0 02 15 AE   ....\..M.}.-....
00000010: 31 C3 97 A6 A4 9F 1D 44 - 85 61 15 49 4A E9 7C 24   1......D.a.IJ.|$
00000020: 22 08 00 00 00 00 00 00 - 21 00 01 00 00 00 00 00   ".......!.......
00000030: 00 00 00 00 00 00 00 00 - 4D 00 69 00 63 00 72 00   ........M.i.c.r.
00000040: 6F 00 73 00 6F 00 66 00 - 74 00 20 00 72 00 65 00   o.s.o.f.t. .r.e.
00000050: 73 00 65 00 72 00 76 00 - 65 00 64 00 20 00 70 00   s.e.r.v.e.d. .p.
00000060: 61 00 72 00 74 00 69 00 - 74 00 69 00 6F 00 6E 00   a.r.t.i.t.i.o.n.
00000070: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
00000080: AA C8 08 58 8F 7E E0 42 - 85 D2 E1 E9 04 34 CF B3   ...X.~.B.....4..
00000090: 66 F2 3F 3A 09 D9 EA 49 - B1 32 75 D5 98 04 3C 34   f.?:...I.2u...<4
000000A0: 22 00 00 00 00 00 00 00 - 21 08 00 00 00 00 00 00   ".......!.......
000000B0: 00 00 00 00 00 00 00 00 - 4C 00 44 00 4D 00 20 00   ........L.D.M. .
000000C0: 6D 00 65 00 74 00 61 00 - 64 00 61 00 74 00 61 00   m.e.t.a.d.a.t.a.
000000D0: 20 00 70 00 61 00 72 00 - 74 00 69 00 74 00 69 00    .p.a.r.t.i.t.i.
000000E0: 6F 00 6E 00 00 00 00 00 - 00 00 00 00 00 00 00 00   o.n.............
000000F0: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
00000100: A0 60 9B AF 31 14 62 4F - BC 68 33 11 71 4A 69 AD   .`..1.bO.h3.qJi.
00000110: E2 33 A2 82 3A 5E D5 4C - AE 8E 4B EC 6B 76 4D ED   .3..:^.L..K.kvM.
00000120: 22 00 01 00 00 00 00 00 - 09 77 11 01 00 00 00 00   ".......w......
00000130: 00 00 00 00 00 00 00 00 - 4C 00 44 00 4D 00 20 00   ........L.D.M. .
00000140: 64 00 61 00 74 00 61 00 - 20 00 70 00 61 00 72 00   d.a.t.a. .p.a.r.
00000150: 74 00 69 00 74 00 69 00 - 6F 00 6E 00 00 00 00 00   t.i.t.i.o.n.....
00000160: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
00000170: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00   ................
```

**Partition Type GUIDs**

| Partition Type | GUID Value |
|---|---|
| Unused entry | {00000000–0000–0000–0000–000000000000} |
| EFI System partition | {28732AC1–1FF8–D211–BA4B–00A0C93EC93B} |

| | |
|---|---|
| Microsoft Reserved partition | {16E3C9E3–5C0B–B84D–817D–F92DF00215AE} |
| Primary partition on a basic disk | {A2A0D0EB–E5B9–3344–87C0–68B6B72699C7} |
| LDM Metadata partition on a dynamic disk | {AAC80858–8F7E–E042–85D2–E1E90434CFB3} |
| LDM Data partition on a dynamic disk | {A0609BAF–3114–624F–BC68–3311714A69AD} |

**GUID Partition Entry Attributes**

GUID partition entry attributes are descriptors for how a partition is used. The attributes are specified within a 64-bit value, so EFI supports up to 64 different attributes. The 64-bit versions of Windows Server 2003 use two attributes as described in the following table.

**GUID Partition Entry Attributes Used by the 64-Bit Editions of Windows on Itanium-based Computers**

| Bits | Description |
|---|---|
| Bit 0 | Specifies that this partition is required for the platform to function. All original equipment manufacturer (OEM) partitions must have this bit set to protect the OEM partition from being overwritten by the disk tools supplied with Windows Server 2003. |
| Bit 60 | Marks the partition as read-only. Used only for primary basic partitions of type {EBD0A0A2-B9E5-4433-87C0-68B6B72699C7}. |
| Bit 62 | Marks the partition as hidden. Used only for primary basic partitions of type {EBD0A0A2-B9E5-4433-87C0-68B6B72699C7}. |
| Bit 63 | Prevents the system from assigning a default drive letter to the partition. Used only for primary basic partitions of type {EBD0A0A2-B9E5-4433-87C0-68B6B72699C7}. |

**Types of RAID Volumes**

A redundant array of independent disks (RAID) is a fault-tolerant disk configuration in which part of the physical storage capacity contains redundant information about data stored on the disks. The redundant information is either parity information (in the case of a RAID-5 volume), or a complete, separate copy of the data (in the case of a mirrored volume). The redundant information enables regeneration of the data if one of the disks or the access path to it fails, or a sector on the disk cannot be read.

Windows Server 2003 supports three types of software RAID configurations:

- **Striped volumes** use RAID-0, which stripes data across multiple disks. RAID-0 does not offer fault tolerance, but it does offer increased performance.
- **Mirrored volumes** use RAID-1, which provides redundancy by creating two identical copies of a volume.
- **RAID-5 volumes** use RAID-5, which stripes parity information across multiple disks. This parity information can be used to recreate data stored on a failed disk.

**Note**

- Fault tolerance is never an alternative to performing regular backups.

Use DiskPart or Disk Management to configure and repair mirrored volumes and RAID-5 volumes. The following figure shows a mirrored and a RAID-5 volume with failed redundancy status. The mirrored and RAID-5 volumes are in failed redundancy because one of the disks that makes up the volumes is offline.

**Mirrored and RAID-5 Volumes That Have Failed Redundancy Status**




**Striped Volumes**

Striped volumes improve input/output (I/O) performance by distributing I/O requests across two or more disks. Striped volumes are composed of stripes of data of equal size written across each disk in the volume. They are created from equally sized, unallocated areas on two or more disks. For Windows Server 2003, the size of each stripe is 64 kilobytes (KB).

The disks in a striped volume do not need to be identical, but there must be unused space available on each disk that you want to include in the volume.

You cannot increase the size of a striped volume after it is created. To change the size of a striped volume, you must first complete the following steps:

1. Back up the data.
2. Delete the striped volume by using Disk Management or DiskPart.
3. Create a new, larger, striped volume by using Disk Management or DiskPart.
4. Restore the data to the new striped volume.

Striped volumes do not contain redundant information. Therefore, the cost per gigabyte on a striped volume is identical to that for the same amount of storage configured from a contiguous area on a single disk. If one disk fails, the entire striped volume fails and no data can be recovered. The reliability for the striped volume is less than the least reliable disk in the set.

Striped volumes are used for performance reasons. In general, striped volumes work well when you need to distribute disk I/O operations. Access to the data on a striped volume is usually faster than access to the same data would be on a single disk, because the I/O is spread across more than one disk. Therefore, Windows Server 2003 can be seeking on more than one disk at the same time and can have simultaneous read or write operations.

A striped volume works well in the following situations:

- When users need rapid read or write access to large databases or other data structures.
- When collecting data from external sources at very high transfer rates. This is especially useful when collection is done

asynchronously.

- When multiple independent applications require access to data stored on the striped volume. When the operating system supports asynchronous multithreading, which helps with load balancing of disk read and write operations.

**Mirrored Volumes**

A mirrored volume provides an identical twin of the selected volume. All data written to the mirrored volume is written to both volumes, which results in disk capacity of only 50 percent. Because dual-write operations can degrade system performance, many mirrored volume configurations use duplexing, which means that each disk in the mirrored volume resides on its own disk controller. The benefit of duplexing is that you reduce the risk of a single point of failure: if one disk controller fails, the other controller (and the disk on that controller) continues to operate normally. If you do not use two controllers, a failed controller makes both volumes in a mirrored volume inaccessible until the controller is replaced.

**Note**

- If one disk in a mirrored volume fails, the computer continues to run and the mirrored volume is still accessible. However, the mirrored volume is no longer fault-tolerant, so you need to replace the failed disk or controller as soon as possible. If your computer supports hot-swappable hard disks, you do not need to restart the computer to install a new disk and resynchronize the mirror.

Almost any volume can be mirrored, including the system and boot volumes. However, you cannot mirror the EFI System partition on GPT disks. In addition, you cannot add disk space to a mirrored volume to increase the size of the volume later.

**Advantages of Mirrored Volumes**

Random disk-read operations on a mirrored volume are more efficient than random disk-read operations on a single volume. Windows Server 2003 has the capacity to load balance read operations across the disks. With current SCSI and Fibre Channel technology, two disk read operations can be done simultaneously.

When one of the volumes that makes up a mirrored volume fails, the mirrored volume is said to have lost redundancy. A mirrored volume that loses redundancy impacts system performance the least, because the remaining disk contains all of the data. No data recomputation is needed to run the system. When you configure your boot volume on a mirrored volume, you do not have to reinstall Windows Server 2003 to restart the computer after a disk failure.

When compared to a RAID-5 volume, a mirrored volume:

- Has a lower entry cost because it requires only two disks, whereas a RAID-5 volume requires three or more disks.
- Requires less system memory.
- Provides good overall performance.
- Does not degrade performance during a failure except when high-volume read operations are performed. However, if a single write error occurs, redundancy is lost.

A mirrored volume works well in the following situations:

- When extremely high data reliability is required. A duplexed mirrored volume has the best data reliability because the entire I/O subsystem is duplicated.
- When you have heavy write loads that need fault tolerance. In this case, mirrored volumes perform better than RAID-5 volumes.
- When simplicity is important. Mirrored volumes are simple to understand and easy to set up.

**Disadvantages of Mirrored Volumes**

Disk-write operations on mirrored volumes are less efficient because data must be written to both disks. This performance penalty is minor, however, because writing to both disks usually takes place concurrently. In many situations, an end-user application is not affected by data being written do both disks.

Another performance penalty occurs when you resynchronize a mirrored volume. Resynchronization is the process by which a mirrored volume's mirrors are made to contain identical data. During resynchronization, performance is affected because the computer is performing many I/O operations to copy the data.

Mirrored volumes are the least efficient at maximizing storage space. Because the data is duplicated, the space requirements for a mirrored volume are higher than for a RAID-5 volume.

**Best Practices for Configuring Mirrored Volumes**
To a large extent, how you configure your mirrored volumes depends on the number of disks and controllers that you want to have on the computer running Windows Server 2003. The following are general guidelines for configuring mirrored volumes:

- Keep data volumes separate from boot volumes for better performance. Configuring your boot volume on a disk (and controller) that does not contain data sets gives you better performance.

- Do not put the paging file on a mirrored volume. The paging file does not need to be redundant and can decrease the mirrored volume's performance due to frequent disk-writing operations. Instead, put the paging file on a striped or simple volume.

- For additional protection, put each disk in a mirrored volume on its own disk controller. When you use a mirrored volume for your system or boot volumes, you can make the configuration more fault-tolerant by putting each disk member of the mirrored volume on a separate controller. This approach allows you to survive controller or disk failures. Putting each disk member of the mirrored volume on a separate channel of a multichannel controller does not make the controller fault tolerant. However, this approach might improve performance.

- Use identical disks when putting the system or boot volume on a mirrored volume. Although it is not necessary to use identical disks or to have the same volumes on each disk, it is strongly recommended that you use identical disks and controllers if you put your system and boot volume on a mirrored volume.

**Note**

- After you mirror your system volume, you must test your configuration by starting Windows Server 2003 from each volume to ensure that you can start Windows Server 2003 if one of the disks fails. Startup problems can occur if the disks use different geometries or if the system volumes are at different offsets on the disk.

**Creating Mirrored Volumes**
To create a mirrored volume, use the Disk Management snap-in or the DiskPart command-line tool. You can create a mirrored volume in two ways:

- Add a mirror to an existing simple volume on a dynamic disk. You must have an area of unused space on a different dynamic disk at least as large as the original simple volume. If you do not have a dynamic disk with enough unallocated space, the **Add Mirror** command is unavailable.

- Create a new mirrored volume from unallocated space on two dynamic disks. The amount of disk space used for each half of the mirrored volume must be equal. If you have less unallocated space on one disk than the other, the mirrored volume can be no larger than the smaller of the two unallocated spaces.

In either case, if you have unallocated space left over, you can use the space to create other volumes.

**Mirroring the System and Boot Volumes in x86-based Computers**
To ensure that your x86-based computer can load Windows Server 2003 if one of the disks or controllers fails, you can mirror the system and boot volumes.

- Mirroring the system volume makes an exact copy of the volume that contains the hardware-specific files needed to load Windows Server 2003.

- Mirroring the boot volume makes an exact copy of the volume that contains the Windows Server 2003 operating system.

The boot and system volumes can be separate volumes on the same disk, separate volumes on different disks, or they can be the same volume on the same disk. In addition, the system and boot volumes can be mirrored to a different disk on the same

controller or to a different disk on a different controller. The following figure illustrates some common configurations for mirroring system and boot volumes.

**Common Configurations of Mirrored System Volumes and Boot Volumes (x86-based Computers)**




**Guidelines for Mirroring System or Boot Volumes in x86-based Computers**

Before you mirror the system or boot volume in an x86-based computer, note the following guidelines:

- Use care when selecting Advanced Technology Attachment (ATA) disks for a mirrored system volume. Although using ATA disks is supported, the recovery procedure is more complicated when the master disk on the primary integrated device electronics (IDE) channel fails. In this case, you must move the disk with the remaining mirror to the primary IDE channel and set its jumper to master position.

- Do not mirror the system volume by using an ATA disk with a SCSI disk because startup problems can occur if one of the disks fails.

- If you use duplexed SCSI controllers, make sure to use identical controllers from the same manufacturer.

- You must test the mirrored system volume before a failure to ensure that the computer can start from the remaining mirror.

**Mirroring the Boot Volume and Replicating the EFI System Partition in Itanium-based Computers**

To ensure that your Itanium-based computer can load Windows Server 2003 if one of the disks or controllers fails, you can mirror the boot volume. Mirroring the boot volume makes an exact copy of the volume that contains the Windows Server 2003 operating system. In addition to mirroring the boot volume, you should also replicate the EFI system partition. If you do not replicate the EFI system partition, and the disk holding it fails, you will not be able to boot the computer, even if there is a good, remaining boot volume on another disk.

The process to replicate the EFI system partition involves creating a new EFI system partition on a second GUID partition table (GPT) disk. Because the second EFI system partition is empty, you must copy the contents from original EFI system partition into the second EFI system partition. After replicating the EFI system partition, you must use Bootcfg.exe to add the appropriate boot entries to the NVRAM to point to the copy of the EFI system partition you placed on the second disk. Later, if you make any changes to the original EFI system partition, you must manually replicate those changes in the second EFI system partition.

The boot volume and the EFI system partition can be on the same disk, or they can be the on different disks. The following figure illustrates the most common configuration.

**Common Configurations of Mirrored Boot Volumes and Replicated EFI System Partitions (Itanium-based Computers)**




**Prior to mirroring** **After mirroring the boot volume and replicating the EFI system partition**

**Guidelines for Mirroring the Boot Volume and Replicating the EFI System Partition in Itanium-based Computers**

Before you mirror the boot volume or replicate the EFI system partition in an Itanium-based computer, you must note the following guidelines:

- Use care when selecting Advanced Technology Attachment (ATA) disks for a replicated EFI sys partition. Although using ATA disks is supported, the recovery procedure is more complicated when the master disk on the primary integrated device electronics (IDE) channel fails. In this case, you must move the disk with the remaining EFI system partition to the primary IDE channel and set its jumper to master position.

- Do not replicate the EFI system partition by using an ATA disk with a SCSI disk, because startup problems can occur if one of the disks fails.

- If you use duplexed SCSI controllers, make sure to use identical controllers from the same manufacturer.

- You must test the mirrored boot volume and replicated EFI system partition before a failure, to ensure that the computer can start from the remaining boot volume and EFI system partition.

**RAID-5 Volumes**

Using three or more disks, a RAID-5 volume dedicates the equivalent of the space of one disk in the RAID-5 volume for storing parity stripes, but distributes the parity stripes across all the disks in the group. The data and parity information are arranged on the volume so that they are always on different disks.

Implementing a RAID-5 volume requires a minimum of three disks. The disks do not need to be identical, but there must be equally sized blocks of unallocated space available on each disk in the set. The disks can be on the same or different controllers. However, neither the system volume nor boot volume can be on a RAID-5 volume.

### Note

- As with striped volumes, you cannot add disks to a RAID-5 volume if you need to increase the size of the volume later.

If one of the disks in a RAID-5 volume fails, none of the data is lost. When a read operation requires data from the failed disk, the system reads all of the remaining good data stripes in the stripe and the parity stripe. Each data stripe is subtracted (by using **XOR**) from the parity stripe; the order is not important. The result is the missing data stripe.

When the system needs to write a data stripe to a disk that has failed, it reads the data stripes on the other disks. The system uses the data stripes on the remaining disks to calculate the parity. Because the data stripe on the failed disk is unavailable, it is not written; the system only updates the parity stripe.

**Advantages of RAID-5 Volumes**

RAID-5 volumes work well for storing data that will need to be read frequently but written to less frequently and also work well in the following situations:

- In large query or database mining applications where reads occur much more frequently than writes. Performance degrades as the percentage of write operations increases. Database applications that read randomly work well with the built-in load balancing of a RAID-5 volume.

- Where a high degree of fault tolerance is required without the expense (incurred by the additional disk space required) of a mirrored volume. A RAID-5 volume is significantly more efficient than a mirrored volume when larger numbers of disks are used. The space required for storing the parity information is equivalent to 1/Number of disks, so a 10-disk array uses 1/10 of its capacity for parity information. The disk space that is used for parity decreases as the number of disks in the array increases.

**Disadvantages of RAID-5 Volumes**
RAID-5 volumes are not well suited for most write-intensive workloads because a single write is likely to generate a two disk reads (one to read the old data and one to read the old parity information) and two writes (one to update the data and a second to update the parity information).

For example, a RAID-5 volume is not well suited for the following situations:

- For hosting applications that require high-speed data collection. This type of application requires continuous high-speed disk writes, which do not work well with the asymmetrical I/O balance inherent in RAID-5 volumes and the extra I/Os required to write the parity stripe.

- In transaction-processing database applications in which records are continually updated, such as in financial applications where balances are frequently updated.

If a disk that is part of a RAID-5 volume fails, read operations for data stripes on that disk are substantially slower than for a single disk. The software has to read all of the other disks in the set to calculate the data.

A RAID-5 volume requires more system memory than a mirrored volume. In addition, regenerating a RAID-5 volume negatively impacts performance more than regenerating a mirrored volume does.

**Guidelines for Configuring RAID-5 Volumes**
When configuring a RAID-5 volume, buy disks based on:

- **Performance**. RAID-5 performance improves for each disk that you use.

- **Percentage of usable storage**. The space lost to parity information decreases with each additional disk.

- **Cost per gigabyte**. A RAID-5 volume requires a minimum of three disks. Buying three large disks might cost less per gigabyte, but buying six smaller disks results in better performance and more available disk space because less space is used for parity.

The following table compares two RAID-5 configurations that provide the same disk capacities. The configuration that uses six disks is the more efficient storage solution in terms of capacity and performance.

**Comparison of Two RAID-5 Configurations**

| Comparative Feature | 3 Disks at 36.4 GB Each | 6 Disks at 18.2 GB Each |
|---|---|---|
| Total capacity | 109.2 GB | 109.2 GB |
| Spaced used for parity | 36.0 GB | 18.6 GB |
| Available disk space | 73.2 GB | 90.6 GB |

Follow these guidelines for configuring RAID-5 volumes:

- Do not configure your system volume or your boot volume on a RAID-5 volume. In addition, keep the RAID-5 volume on a different controller and disk than your system and boot volume. Using separate controllers improves performance and can accelerate recovery from hardware failures.

- Do not put the paging file on a RAID-5 volume. The paging file does not need to be redundant and can decrease the RAID-5 volume's performance due to frequent disk-writing operations. Instead, put the paging file on a striped or simple volume.

**Fault-Tolerant Hardware and Software**
You can create a RAID-5 volume using hardware- or software-based solutions. With hardware-based RAID, an intelligent disk controller handles the creation and regeneration of redundant information on the disks that make up the RAID-5 volume. The Windows Server 2003 family of operating systems provides software-based RAID, where the creation and regeneration of redundant information on the disks in the RAID-5 volume is handled by the Logical Disk Manager (LDM). In either case, data is stored across all members in the disk array.

In general, hardware-based RAID offers performance advantages over software-based RAID because hardware-based RAID incurs no overhead on the system processor. For example, you can improve data throughput significantly by implementing RAID-5 through hardware that does not use system software resources. Read and write performance and total storage size can be further improved by using multiple disk controllers.

Some hardware-based RAID arrays support hot swapping, which enables you to replace a failed disk or controller while the computer is still running Windows Server 2003. Consider the following points when you evaluate a fault-tolerant hardware or software solution:

- Hardware fault tolerance provides better performance.
- Hardware fault tolerance offers features such as hot sparing, in which additional disks are attached to the controller and left in standby mode. If a failure occurs, the controller uses one of the spare disks to replace the bad disk.
- Software fault tolerance is less expensive.

Regardless of whether you implement fault tolerance by using hardware, software, or both, implementing fault tolerance does not reduce the need for backups.

# Dynamic Disks and Volumes Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dynamic Disks and Volumes Tools and Settings**

**In this section**


- Dynamic Disk and Volume Tools
- Dynamic Disk and Volume Registry Entries
- Dynamic Disk and Volume WMI Classes


The following tools, registry settings, and Windows Management Instrumentation (WMI) classes are associated with dynamic disks and volumes.

**Dynamic Disk and Volume Tools**

The following tools are associated with dynamic disks and volumes.

**Bootcfg.exe: Boot Configuration Tool**

**Category**

Windows Server 2003 operating system tool.

**Version compatibility**

All versions of the Boot Configuration Tool in the Windows Server 2003 family are identical.

You can use the Boot Configuration Tool to query, configure, or change settings in the Boot.ini file on your computer.

**Dmdiag.exe: Disk Manager Diagnostics**

**Category**

Windows Server 2003 operating system tool.

**Version compatibility**

All versions of Dmdiag in the Windows Server 2003 family are identical.

Dmdiag displays the following information for the computer on which it is run:

- Computer name and operating system version
- Physical disk to disk type
- Mount points
- Logical Disk Manager (LDM) file versions
- Drive letter usage
- List of devices
- Symbolic links
- Disk partition information

**Dskmgmt.msc: Disk Management**

**Category**

Windows Server 2003 operating system tool.

**Version compatibility**

All versions of the Disk Management snap-in in the Windows Server 2003 family are identical.

You can use the Disk Management snap-in to remotely manage disks and volumes on other computers running Windows 2000, Windows XP or Windows Server 2003

**DiskPart.exe: DiskPart**
**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
All versions of DiskPart in the Windows Server 2003 family are identical. DiskPart includes a few 64-bit parameters that are only available on Itanium-based computers.

DiskPart is a text-mode command interpreter that enables you to manage objects (disks, partitions, or volumes) by using scripts or direct input from a command prompt.

**Format.exe: Format**
**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
All versions of Format in the Windows Server 2003 family are identical.

Format prepares a volume on the specified disk to accept Windows files.

**FTonline.exe: Fault-tolerant Disk Mounter**
**Category**
Windows Server 2003 support tool.

**Version compatibility**
FTonline enables an administrator to mount and recover files from fault-tolerant disks created in previous versions of Windows. This tool is useful if you did not upgrade your disks to dynamic disks, or failed to back up your data before installing Windows Server 2003. You can install FTonline using the Support Tools setup program located in the \Support\Tools folder on the Windows XP Professional and the Windows Server 2003 family of operating systems CDs.

**Mountvol.exe: Mountvol**
**Category**
Windows Server 2003 operating system tool.

**Version compatibility**
All versions of Mountvol in the Windows Server 2003 family are identical. Mountvol includes a few 64-bit parameters that are only available on Itanium-based computers.

Mountvol creates, deletes, or lists volume mount points. Mountvol enables you to link volumes without using drive letters.

**SecInspect.exe: Sector Inspector**
**Category**
Windows Server 2003 Resource Kit tool.

**Version compatibility**
SecInspect is a command-line diagnostics tool that enables administrators to view the contents of master boot records, boot sectors, and IA64 GUID partition tables. Additional features include creating hexadecimal dumps of binary files and backup/restore of sector ranges. For more information about this tool, see the Help that comes with the tool. To find this tool, see Resource Kit Tool Updates in Tools and Settings Collection.

**Dynamic Disk and Volume Registry Entries**
The following registry entries are associated with dynamic disks and volumes. For more information about the registry, see the Registry Reference in Tools and Settings Collection.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following sections describe the dynamic disk and volume registry entries that are listed below \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\. The headings represent the next level in the path after \Services. For example, Vds\Debuglog is equal to \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Vds\Debuglog.

**Vds\Debuglog**
**Level**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Vds\Debuglog\Level

**Version**
Windows Server 2003 family.

Sets the level of logging for VDS.

Use Registry Editor to add a decimal value named Level. Stop and restart the service after changing.

### VDS Log Options

| Item Logged | Bitmask Number |
|---|---|
| Errors | 1 |
| Warnings | 2 |
| Trace | 4 |
| Information | 8 |

Set the value by adding the bitmask numbers of the values you want to log. For example, decimal 3 logs errors (1) and warnings (2). Decimal 9 logs errors (1) and information (8). The range for this value is 0-15. Set the value to 0xF (decimal 15) to log all items.

**Dmadmin\Parameters**
**EnableDynamicConversionFor1394**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\dmadmin\Parameters\EnableDynamicConversionFor1394

**Version**
Windows Server 2003 family.

Allows a user to convert a Institute of Electrical and Electronics Engineers, Inc. (IEEE) 1394 (FireWire) disk to a dynamic disk on Windows Server 2003 or earlier operating systems. Converting 1394 disks to dynamic is not a tested or supported scenario. This registry entry is provided for compatibility with a Windows 2000 beta that did support this functionality.

EnableDynamicConversionFor1394 is a DWORD Value, with a range of 0-1. Default is zero. To modify, use Registry Editor. Stop and restart the service after modifying.

**Dynamic Disk and Volume WMI Classes**
The following table lists and describes the WMI classes that are associated with dynamic disks and volumes.

### WMI Classes Associated with Dynamic Disks and Volumes

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_DiskDrive | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_DiskDrivePhyiscalMedia | \root\cimv2 | Windows Server 2003 family |
| Win32_DiskDriveToDiskPartition | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_DiskPartition | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_LogicalDisk | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_LogicalDiskToPartition | \root\cimv2 | Windows NT Server 4.0 SP4 and later |
| Win32_MappedLogicalDisk | \root\cimv2 | Windows Server 2003 family |
| Win32_PhysicalMedia | \root\cimv2 | Windows Server 2003 family |
| Win32_Volume | \root\cimv2 | Windows Server 2003 family |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

# Storage Services Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Storage Services Technologies**
Much of the data stored on Microsoft Windows-based servers is critically important for the day-to-day activities in business organizations. Because it is important that employees are able to access and store data reliably, data storage and recovery are a concern for many IT managers. In response to the expanding needs for storage in distributed computing environments, Windows Server 2003 includes several storage technologies that are designed to help you store, access, and manage data on servers.

**In this overview**

- Storage Services Components
- Storage Services Scenarios

**Storage Services Components**
This section discusses the components that make up Storage Services:

- Virtual Disk Service
- Remote Storage
- Removable Storage

**Virtual Disk Service**
For Windows 2000, each company that made storage devices (for example, hard disks, Peripheral Component Interconnect (PCI) redundant array of independent disks (RAID) adapters, and storage arrays) provided its own application for installing and managing that storage device. As a result, users had to run separate applications for each type of storage device that they used on their computer system. This made it difficult to manage different types of storage devices. To address this issue, Windows Server 2003 introduced Virtual Disk Service (VDS). VDS is a set of application programming interfaces (APIs) that provides a single interface for managing disks. VDS provides an end-to-end solution for managing storage hardware and disks, and for creating volumes on those disks.

**Remote Storage**
Remote Storage is a component that you can use to extend disk space without adding more hard disks. Remote Storage automatically copies eligible files on your local volumes to a library of magnetic tapes or magneto-optical disks. Remote Storage then monitors the amount of space that is available on your local volumes. When the amount of available space on a local volume falls below the level that you designate, Remote Storage automatically removes the content from a sufficient number of eligible files and migrates this content to an attached storage device, which frees up disk space on the volume.

**Removable Storage**
Removable Storage is a component that you can use to track your removable storage media (tapes and optical disks) and manage the hardware libraries that contain them (changers and jukeboxes).

With Removable Storage, you can:

- Label, catalog, and track media.
- Control library drives, slots, and doors.
- Perform drive-cleaning operations.

Removable Storage works together with your data-management applications such as Backup. You use data-management applications to manage the actual data stored on the media. Removable Storage makes it possible for multiple applications to share the same storage media resources, which can reduce your costs. It also provides a common interface for managing those

resources, which makes resource management much more efficient.

**Storage Services Scenarios**
This section describes a few scenarios in which you can use Virtual Disk Service, Removable Storage, and Remote Storage to solve unique storage problems in your organization.

**Virtual Disk Service Scenarios**
VDS is commonly used in the following scenarios.

**Managing LUNs on Hardware Storage Devices**
VDS defines a set of features which provide a virtual view of any storage hardware. You can use the VDS APIs to write a storage management application that can manage any other storage hardware RAID arrays and logical unit numbers (LUNs) that have VDS providers. For example, you can use the VDS APIs to create or delete simple, spanned, striped, mirrored, or striped with parity LUNs.

**Managing Disks and Volumes**
The VDS APIs enable applications to manage disks and volumes. You can create your own application that uses the VDS APIs or use applications such as the Disk Management snap-in or the DiskPart command-line tool (both use the VDS APIs) to perform tasks such as creating or deleting simple, striped, spanned, mirrored or RAID-5 volumes, formatting volumes, assigning drive letters to volumes, and converting basic disks into dynamic disks.

**Managing End-To-End Storage Operations**
The VDS APIs enable applications to manage end-to-end storage operations. Using the VDS APIs, your application can create a LUN in the storage array, unmask the LUN to a host, create and format a volume on the LUN, and assign a drive letter to the volume.

**Remote Storage Scenarios**
Remote Storage is commonly used in the following scenarios.

**Conserving disk space on managed volumes**
You can use Remote Storage to ensure that disk space on the managed volume where your system files are located does not fall below a certain level. This is important if you want to make sure that operating system efficiency is not compromised because of decreased disk space.

**Extending disk space on managed volumes**
You can use Remote Storage to make additional space available for new, more frequently accessed files, while retaining older, less frequently accessed files that you might need to access some time in the future.

**Generating multiple media copies of removable media**
You can use Remote Storage to ensure that one or more copies of your migrated files exist on different media in case the original storage medium (disk or tape) becomes unusable.

**Replacing damaged removable media while Remote Storage is running**
You can use Remote Storage to ensure that when you replace a copy of your migrated files that has become damaged or unusable with another copy of the same files, Remote Storage will, at the same time, automatically generate additional copies on different media.

**Recovering from loss of Remote Storage metadata**
You can use Remote Storage to ensure that you can retrieve any metadata that is associated with your Remote Storage files.

**Removable Storage Scenarios**
Removable Storage is commonly used in the following scenarios.

**Managing stand-alone drive libraries**
You can use Removable Storage to manage multiple single-drive libraries, such as CD-ROM or DVD-ROM drives. A library consists of data-storage media and the device that is used to read from and write to the media – for example, tape and a stand-alone tape drive. You can also use Removable Storage to organize all the media in your libraries into different media pools, and move media between media pools in order to provide the appropriate amount of data storage that your applications require.

**Managing automated libraries**
You can use Removable Storage to manage automated libraries. Automated libraries are automated units that hold multiple tapes or disks, and some have multiple drives. These libraries are sometimes called *changers* or *jukeboxes*, and commonly use robotic subsystems to move media stored in the library's storage slots. You can also use Removable Storage to manage a combination of single-drive and automated libraries.

# Virtual Disk Service Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual Disk Service Technical Reference**
For Microsoft Windows 2000, each company that made storage devices (for example, hard disks, PCI RAID adapters, storage arrays) provided its own application for installing and managing that storage device. As a result, users had to run separate applications for each type of storage device they used on their computer system. This made managing the different types of storage devices difficult. To address this issue, Microsoft Windows Server 2003 introduced Virtual Disk Service (VDS). VDS is a set of application programming interfaces (APIs) that provides a single interface for managing disks. VDS provides an end-to-end solution for managing storage hardware and disks, and for creating volumes on those disks.

VDS uses two sets of providers to manage your storage devices. The built-in VDS software providers enable you to manage disks and volumes at the operating system level. Hardware providers supplied by the hardware vendor enable you to manage hardware RAID arrays.

Each hardware vendor writes a VDS hardware provider that translates the general-purpose VDS APIs into specific instructions for their hardware. VDS simplifies storage management because storage applications no longer need to take into account the specific hardware being targeted. Because of VDS, you no longer need to use a different storage application to manage each different storage device. Instead, you can use one compatible storage application to manage any hardware that has a VDS hardware provider.

Windows Server 2003 components that work with VDS include the Disk Management snap-in, the DiskPart command-line tool, and the DiskRAID command-line tool (Diskraid.exe), which is available in the Windows Server 2003 Deployment Kit. DiskRAID requires at least one VDS provider supplied by a storage vendor.

**In this subject**


- What Is Virtual Disk Service?
- How Virtual Disk Service Works
- Virtual Disk Service Tools and Settings

# What Is Virtual Disk Service?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Common Virtual Disk Service Scenarios
- Technologies Related to Virtual Disk Service
- Virtual Disk Service Dependencies
- Related Information


For Windows 2000, each company that made storage devices (for example, hard disks, PCI RAID adapters, storage arrays) provided its own application for installing and managing that storage device. As a result, users had to run separate applications for each type of storage device they used on their computer system. This made managing the different types of storage devices difficult. To address this issue, Windows Server 2003 introduced Virtual Disk Service (VDS). VDS is a set of application programming interfaces (APIs) that provides a single interface for managing disks. VDS provides an end-to-end solution for managing storage hardware and disks, and for creating volumes on those disks.

VDS uses two sets of providers to manage your storage devices. The built-in VDS software providers enable you to manage disks and volumes at the operating system level. Hardware providers supplied by the hardware vendor enable you to manage hardware RAID arrays.

Each hardware vendor writes a VDS hardware provider that translates the general-purpose VDS APIs into specific instructions for their hardware. VDS simplifies storage management because storage applications no longer need to take into account the specific hardware being targeted. Because of VDS, you no longer need to use a different storage application to manage each different storage device. Instead, you can use one compatible storage application to manage any hardware that has a VDS hardware provider.

Windows Server 2003 components that work with VDS include the Disk Management snap-in, the DiskPart command-line tool, and the DiskRAID command-line tool (Diskraid.exe), which is available in the Windows Server 2003 Deployment Kit. DiskRAID requires at least one VDS provider supplied by a storage vendor.

Depending on the underlying disk subsystem, a disk can be a single physical hard disk (sometimes called a spindle), or it can be a group of hard disks bound together at the controller level and presented to the operating system as a single disk. This type of disk is often referred to as a virtual disk or logical unit number (LUN), and typically uses some form of hardware-based RAID to ensure data availability if one of the disks fails.

If your hardware vendor supplies a VDS hardware provider, you can use DiskRAID to manage individual physical disks, also called drives, in a LUN. The following example shows two DiskRAID printouts. The first printout shows information about the LUNs in a given disk subsystem. The second printout shows information about the physical disks that correspond to LUN 0.

**How DiskRAID Displays LUNs in a Disk Subsystem**

```
LUN ###  Status      Health        Type        Size    Sub  Dev
-------  ----------  ------------  --------  --------  ---  ---
LUN 0    Online      Healthy       Stripe     1064 GB    0    3
LUN 1    Online      Healthy       Parity      200 GB    0    4
```

**How DiskRAID Displays the Disks Associated With LUN 0**

```
Drive ###  Status      Health          Size      Free    Bus  Slot  Flags
---------  ----------  ------------  --------  --------  ---  ----  -----
Drive 0    Online      Healthy          66 GB     66 GB    0     5
Drive 1    Online      Healthy          66 GB     35 GB    0     0
Drive 2    Online      Healthy          66 GB     35 GB    1     0
Drive 3    Online      Healthy          66 GB     66 GB    1    11
Drive 4    Online      Healthy          66 GB     66 GB    0     7
Drive 5    Online      Healthy          66 GB     66 GB    0     8
Drive 6    Online      Healthy          66 GB     35 GB    1     3
```

```
Drive 7     Online       Healthy          66 GB      66 GB     1      7
Drive 8     Online       Healthy          66 GB      35 GB     0      2
Drive 9     Online       Healthy          66 GB      66 GB     1      9
```

Windows Server 2003 supplies VDS software providers for basic and dynamic disks and volumes so that you can use tools such as the Disk Management snap-in and DiskPart command-line tool to manage LUNs. However, the Disk Management snap-in and the DiskPart command-line tool display LUNs as disks, as shown in the following figure. Unlike DiskRAID, you cannot use these tools to view or manage the underlying physical disks in a LUN. However, using the Disk Management snap-in or DiskPart command-line tool, you can create partitions and volumes, convert disks to dynamic, and so forth. In the following figure, disk 2 (volume G:) and disk 3 (volume H:) are both on LUNs.

**How Disk Management Displays LUNs As Disks**




**Common Virtual Disk Service Scenarios**
VDS is commonly used in the following scenarios.

**Managing LUNs on Hardware Storage Devices**
VDS defines a set of features which provide a virtual view of any storage hardware. You can use the VDS APIs to write a storage management application that can manage any other storage hardware RAID arrays and LUNs that have VDS providers. For example, you can use the VDS APIs to create or delete simple, spanned, striped, mirrored, or striped with parity LUNs.

**Managing Disks and Volumes**
The VDS APIs enable applications to manage disks and volumes. You can create your own application that uses the VDS APIs or use applications such as the Disk Management snap-in or the DiskPart command-line tool (both use the VDS APIs) to perform tasks such as creating or deleting simple, striped, spanned, mirrored or RAID-5 volumes, formatting volumes, assigning drive letters to volumes, and converting basic disks into dynamic disks.

**Managing End-To-End Storage Operations**
The VDS APIs enable applications to manage end-to-end storage operations. Using the VDS APIs, your application can create a LUN in the storage array, unmask the LUN to a host, create and format a volume on the LUN, and assign a drive letter to the volume.

**Technologies Related to Virtual Disk Service**
VDS is closely related to the following technologies:

**Volume Shadow Copy Service (VSS)**

The Volume Shadow Copy service provides an infrastructure for creating consistent point-in-time copies known as shadow copies. The Volume Shadow Copy service can produce consistent shadow copies by coordinating with business applications, file-system services, backup applications, and storage hardware.

**Logical Disk Manager (LDM)**
The Logical Disk Manager (LDM) provides the drivers and user-mode components that are used to configure dynamic disks and volumes on computers running the Windows XP, Windows 2000, and Windows Server 2003 operating systems.

**Virtual Disk Service Dependencies**
VDS has the following dependencies:

- **OEM Hardware Providers**. VDS requires a hardware provider that is supplied by the hardware manufacturer in order to manage storage hardware.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Disks and Volumes
- Basic Disks and Volumes Technical Reference
- Dynamic Disks and Volumes Technical Reference

# How Virtual Disk Service Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual Disk Service Architecture**

VDS provides three interfaces: a single interface between the application layer and the service, and two interfaces between the service and provider programs in the data layer. The following illustration shows the application-to-service interface and the service-to-provider interface.

**Virtual Disk Service Architecture**




VDS coordinates with the file-system functions, synchronizes provider activities, and coordinates between applications. Being between the application and provider, VDS presents uniform functionality to applications even though some of the underlying providers might lack such uniformity.

VDS implements the following common functionality: formatting volumes, adding and removing mount points or drive letters, and managing unallocated disks (disks that have no partition information). VDS also returns event notifications to registered applications.

**Software and Hardware Providers**

VDS supplies two provider interfaces, one for a software provider and one for a hardware provider. Each provider implements a different portion of the VDS API.

- A *software provider* is a host-based program that is supported by a kernel-mode driver in the storage input/output (I/O) stack. The provider-kernel runtime interacts with the Plug and Play (PnP) Manager at discovery time to claim each disk. Software providers operate on volumes, disks, and disk partitions.

  VDS includes two software provider types: basic and dynamic. The basic software provider manages basic disks and offers no fault-tolerant binding. The dynamic software provider manages dynamic disks and offers fault management, where applicable. Software provider behavior is consistent with the behavior of basic and dynamic disks on the host. For example, if the operating system of a given host supports fault-tolerant dynamic disks, VDS also supports fault-tolerant dynamic disks on the host.

- A *hardware provider* is software that translates the VDS APIs into instructions that are specific to the storage hardware. Hardware providers implement the methods used to manage a storage subsystem—a hardware disk array or adapter card that enables the creation of logical disks configured to enhance performance, data availability, or data recovery. Many RAID cabinet manufacturers have produced hardware providers designed for use with VDS.

  The capabilities of a hardware provider depend on the capabilities of the underlying storage hardware. Therefore, the degree to which each manufacturer implements the VDS API can vary. For instance, manufacturers can implement as many of the features that are defined in the VDS API as are necessary to optimize hardware configurations, monitor and dynamically tune performance, or automate fault management.

Hardware providers offer several configuration options that are not available to the software providers. The most notable is the automagic configuration model, which presents an attribute-based view of storage to each application. Binding hints, such as “mostly reads” or “fast crash recovery required” replace the complexity of binding physical storage into virtual storage. Each hardware provider performs extent mapping, space allocation, and binding-type selection based on the hints submitted by an application. For the complete hardware provider description, including the available configuration options, see the documentation supplied by your hardware subsystem manufacturer.

**Virtual Disk Service Interfaces**

Storage management applications can use VDS methods to query and configure host-based disks and RAID storage. Typical storage management applications use VDS to solve configuration management and monitoring problems. These applications range from dedicated storage-management systems to applications that require better control over configuration or fault management. The following Windows Server 2003 applications use VDS:

- **The Disk Management snap-in**. This snap-in configures and manages disks controlled by a host computer. System administrators and end users can

use the snap-in to query and configure local (or remote) disks and volumes.

- **The DiskPart utility**. DiskPart is a scriptable command-line utility that configures and manages disks, volumes, and partitions. Although DiskPart is available for Windows XP and Windows 2000, these versions do not support VDS.

- **The DiskRAID utility**. DiskRAID is a scriptable command-line utility that configures and manages hardware RAID subsystems. This utility can interact with any storage hardware accompanied by a VDS hardware provider. DiskRAID ships with the Windows Server 2003 Technical Reference and the Windows Server 2003 Deployment Kit.

VDS objects provide interfaces that expose methods for querying, configuring, and maintaining storage devices, as shown in the following diagram.

**VDS Interfaces Architecture**




The following tables list the interfaces implemented by the VDS service, the VDS software providers, and the VDS hardware providers.

**VDS Interfaces Implemented by Applications**

| Interface | Description |
|---|---|
| IVdsAdviseSink | Receives VDS notifications. |

**Interfaces Implemented by VDS**

| Interface | Description |
|---|---|
| IVdsAdmin | Registers a provider with VDS. |
| IVdsService | Provides service-level methods for working with VDS. |
| IVdsServiceLoader | Launches VDS. |
| IVdsVolumeMF | Performs access path and file system operations on volume objects. |

**VDS Interfaces Implemented by Software Providers**

| Interface | Description |
|---|---|
| IEnumVdsObject | Enumerates through a set of VDS objects of a given type. |
| IVdsAdvancedDisk | Creates and deletes partitions, and modifies partition attributes, such as read-only, hidden, and so forth.<br><br>Used most often with OEM partitions.<br><br>For ordinary partitions, use IVdsVolume. |
| IVdsAsync | Manages asynchronous operations. |

| | |
|---|---|
| IVdsCreatePartitionEx | Creates a new partition on a basic disk.<br><br>Used most often with OEM partitions or for aligning partitions.<br><br>For ordinary partitions, use IVdsVolume. |
| IVdsDisk | Queries and configures basic and dynamic disks. |
| IVdsPack | Queries and configures packs containing disks, and creates volumes. |
| IVdsProvider | Returns a hardware or software provider's properties. |
| IVdsSwProvider | Performs software-provider operations. |
| IVdsRemovable | Queries and ejects removable disks, such as CD-ROMs. |
| IVdsVolume | Creates and deletes volume plexes, and modifies volume attributes. |
| IVdsVolumePlex | Queries the containing volume and repairs bad extents. |

**VDS Interfaces Implemented by Hardware Providers**

| Interface | Description |
|---|---|
| IEnumVdsObject | Enumerates through a set of VDS objects of a given type. |
| IVdsAsync | Manages asynchronous operations. |
| IVdsController | Exposes methods for performing query and configuration operations on a controller. |
| IVdsDrive | Provides methods for performing query and configuration operations on a drive. |
| IVdsHwProvider | Provides methods for performing query, re-enumeration, and refresh operations on a hardware provider. |
| IVdsLun | Provides methods for performing query and configuration operations on a LUN. |
| IVdsLunPlex | Provides methods for performing query and configuration operations on a LUN plex. A *LUN plex* is the collection of disks that make up the LUN. Except for mirrored LUNs, each LUN can have only one plex. Each plex contains a copy of the data on the LUN. New plexes can be added to a LUN, and, with the exception of the original plex, existing plexes can be removed. VDS supports four LUN plex types: simple, spanned, striped, and striped with parity. |
| IVdsMaintenance | Provides methods for performing maintenance operations on a subsystem, controller, drive, or LUN. |
| IVdsProvider | Returns a hardware or software provider's properties. |
| IVdsSubSystem | Provides methods for performing query and configuration operations on a subsystem. |

**Virtual Disk Service Processes and Interactions**

Some of the common VDS processes and interactions are discussed in the following sections:

- Managing Servers and Desktops
- Retrieving a List of Providers
- Determining the Number of Controllers on the Subsystem

**Managing Servers and Desktops**

VDS makes it easy for you to manage the storage on remote servers. Applications that are written for VDS, such as the Disk Management snap-in, use the appropriate services to connect to and manage the storage on those systems.

To remotely manage a computer running Windows Server 2003 from another computer running Windows Server 2003, the following process occurs. The VDS Loader on the first Windows Server 2003 computer loads VDS on the managed computer. VDS calls CoCreateInstance, which passes an interface pointer back to the Disk Management snap-in, which then allows remote management of the computer.

The following diagram illustrates the components that are used when you remotely manage the storage devices on computers that are running Windows Server 2003.

**Components Used When Managing Remote Servers**




**Retrieving a List of Providers**

When you manage storage on a server, it is often useful to retrieve the list of hardware and software providers. To get the list of providers, you can use an application, such as DiskRAID, to call CoCreateInstance with the class identifier of the VDS Loader and request the IVdsServiceLoader interface. Using the IVdsServiceLoader interface, the application makes a call to the LoadService() method which loads VDS.

The application then calls the WaitForServiceReady() method to wait for the service to fully load before making other calls. The LoadService() method loads the providers. VDS queries the registry and retrieves the list of registered software and hardware providers. VDS loads all providers listed in the registry. The LoadService() method returns the IVdsService interface, which contains the QueryProviders method. This method returns a COM enumerator interface.

The application then uses the COM enumerator interface to get an IUnknown interface on the provider object (call Next*() on the enumerator interface), and calls the QueryInterface method on the IUnknown interface to get specific provider interface, such as IVdsSwProvider or IVdsHwProvider. This provider interface can be used to query the software and hardware providers and retrieve the providers properties such as the name, type, and version. The following diagram summarizes the process of retrieving a list of providers.

**How an Application Gets the List of Providers**




**Determining the Number of Controllers on the Subsystem**

When managing storage on a server, it is often useful to determine how many controllers are on the storage subsystems attached to the server. To determine the number of controllers that are on the subsystem, an application, such as DiskRAID, uses the IVdsService interface to call the QueryProviders method and get a COM enumerator interface. Calling the Next() method on the COM enumerator interface provides an IUnknown pointer to a provider object. The application can call the QueryInterface() method on the IUnknown interface and request the appropriate provider interface, in this case, IVdsHwProvider.

With the hardware provider interface, the application can query the hardware subsystems by calling the QuerySubSystems() method to get a subsystem enumerator interface. Calling the Next() method on this interface returns an IUnknown interface on the subsystem object. The application can then call the QueryInterface() method to request a specific interface, such as the IVdsSubSystem interface. The application obtains the list of controllers for a subsystem by calling the QueryControllers method on the IVdsSubSystem interface. The following diagram summarizes the process of determining the number of

controllers on the subsystem.

**How an Application Determines the Number of Controllers on the Subsystem**




**Related Information**

The following resources contain additional information that is relevant to this section.

- Disks and Volumes
- Basic Disks and Volumes Technical Reference
- Dynamic Disks and Volumes Technical Reference

# Virtual Disk Service Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Virtual Disk Service Tools
- Virtual Disk Service Registry Entries


**Virtual Disk Service Tools**
The following tools are associated with Virtual Disk Service (VDS).

**DiskPart.exe: DiskPart**
**Category**
Windows Server 2003 operating system tool

**Version compatibility**
All versions of DiskPart in the Windows Server 2003 family are identical. DiskPart includes a few 64-bit parameters that are only available on Itanium-based computers.

DiskPart is a text-mode command interpreter that you can use to manage objects (disks, partitions, or volumes) by using automated scripts or commands from a command prompt.

**DiskRAID.exe: RAID configuration tool**
**Category**
Windows Server 2003 Resource Kit tool

**Version compatibility**
DiskRAID is not included in the Windows Server 2003 operating system. DiskRAID is included with the Windows Server 2003 Resource Kit tools.

DiskRAID is a text-mode command interpreter. You can use DiskRAID to manage hardware devices that have VDS hardware providers by using automated scripts or commands from a command prompt. You can use DiskRAID to configure and manage redundant array of independent disks (RAID) storage subsystems.

To find more information about DiskRAID, see Resource Kit Tools in Tools and Settings Collection.

**Dskmgmt.msc: Disk Management**
**Category**
Windows Server 2003 operating system tool

**Version compatibility**
All versions of the Disk Management snap-in in the Windows Server 2003 family are identical.

You can use the Disk Management snap-in to remotely manage disks and volumes on other computers running Windows 2000, Windows XP or Windows Server 2003.

**Mountvol.exe: Mountvol**
**Category**
Windows Server 2003 operating system tool

**Version compatibility**
All versions of the Mountvol command-line tool in the Windows Server 2003 family are identical. Mountvol includes a few 64-bit parameters that are only available on Itanium-based computers.

Mountvol creates, deletes, or lists volume mount points. You can use Mountvol to link volumes without using drive letters.

**Virtual Disk Service Registry Entries**
The following registry entries are associated with VDS. For more information about the registry, see the Registry Reference in Tools and Settings Collection.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are

not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

The following sections describe the VDS registry entries that are listed below \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Vds. The headings represent the next level in the path after \Vds. For example, Debuglog is equal to \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Vds\Debuglog.

**Debuglog**
**Level**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Vds\Debuglog\Level

**Version**
Windows Server 2003 family

Sets the level of logging for VDS.

Use Registry Editor to add a decimal value named Level. Stop and restart the service after changing.

### VDS Log Options

| Item Logged | Bitmask Number |
|---|---|
| Errors | 1 |
| Warnings | 2 |
| Trace | 4 |
| Information | 8 |

Set the value by adding the bitmask numbers of the values you want to log. For example, decimal 3 logs errors (1) and warnings (2). Decimal 9 logs errors (1) and information (8). The range for this value is 0-15. Set the value to 0xF (decimal 15) to log all items.

**SoftwareProviders**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Vds\SoftwareProviders

**Version**
Windows Server 2003 family

Stores the globally unique identifiers (GUIDs) and attributes of the software providers that are registered with VDS.

VDS supplies a class identifier (Clsid) for the service loader and service objects, but only the service loader Clsid is public. The service is initialized when the providers, a calling application, and the service perform the following tasks:

- Each new provider invokes the IVdsAdmin::RegisterProvider method during installation to register with VDS. The call creates a registry key under either the SoftwareProviders or the HardwareProviders key (depending on the provider type), which is identified by the object GUID of the provider. The Clsid of the provider object is contained under this key. It includes the provider's name, version, and version GUID.

### Note

- Object GUIDs of the provider are persistent. Software and hardware object GUIDs are not.

- An application calls the CoCreateInstance function and passes the service loader Clsid as an argument. With a pointer to the service loader object, the application can start VDS locally or remotely by passing the desired computer name as a parameter to the IVdsServiceLoader::LoadService method.

- When the initial application attaches to the service, VDS first calls CoCreateInstance on each Clsid found under the SoftwareProviders and HardwareProviders registry keys, and then calls the IVdsProviderPrivate::OnLoad method on each provider to start the programs.

**HardwareProviders**
**Registry path**
\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Vds\HardwareProviders

**Version**
Windows Server 2003 family

Stores the globally unique identifiers (GUIDs) and attributes of the non-Microsoft hardware providers that are registered with VDS. VDS requires these GUIDs to be able to connect to and manage the storage hardware. When VDS starts, it loads all registered hardware providers.

VDS supplies a class identifier (Clsid) for the service loader and service objects, but only the service loader Clsid is public. The service is initialized when the providers, a calling application, and the service perform the following tasks:

- Each new provider invokes the IVdsAdmin::RegisterProvider method during installation to register with VDS. The call creates a registry key under either the SoftwareProviders or the HardwareProviders key (depending on the provider type), which is identified by the object GUID of the provider. The Clsid of the provider object is contained under this key. It includes the provider's name, version, and version GUID.

### Note

- Object GUIDs of the provider are persistent; software and hardware object GUIDs are not.

- An application calls the CoCreateInstance function, and passes the service loader Clsid as an argument. With a pointer to the service loader object, the application can start VDS locally or remotely by passing the desired computer name as a parameter to the IVdsServiceLoader::LoadService method.

- When the initial application attaches to the service, VDS first calls CoCreateInstance on each Clsid found under the SoftwareProviders and HardwareProviders registry keys, and then calls the IVdsProviderPrivate::OnLoad method on each provider to start the programs.

# Remote Storage Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Storage Technical Reference**
You can use Remote Storage to extend disk space on your server without adding more hard disks. Remote Storage automatically copies eligible files on your local volumes to a library of magnetic tapes or magneto-optical disks. Remote Storage then monitors the amount of space that is available on your local volumes.

When the amount of available space on a local volume falls below the level that you designate, Remote Storage automatically removes the content (data) from a sufficient number of eligible files and migrates this content to an attached storage device, thus freeing up disk space on the volume. When data is removed from a file, the logical size of the file remains the same, and to the user, the file appears unchanged, but the physical size of the file is reduced to 1 KB. Content in other eligible files is not removed until more disk space is needed. When you need to open a file whose data has been removed, the data is automatically recalled from remote storage.

Remote Storage runs on Microsoft Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and the 64-bit versions of these operating systems. Remote Storage is also available on Microsoft Windows 2000 Server. However, Remote Storage is not installed by default. To install Remote Storage, you can select Remote Storage during Setup, or you can later manually install Remote Storage through **Add or Remove Programs** in Control Panel. You can only use Remote Storage to manage NTFS volumes.

**In this subject**


- What Is Remote Storage?
- How Remote Storage Works
- Remote Storage Tools and Settings

# What Is Remote Storage?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Remote Storage?**

**In this section**


- Common Scenarios for Remote Storage
- Remote Storage Dependencies on or Interactions with Other Technologies
- Remote Storage Logical Diagrams
- Related Information


You can use Remote Storage to extend disk space on your server without adding more hard disks. Remote Storage automatically copies eligible files on your local volumes to a library of magnetic tapes or magneto-optical disks. Remote Storage then monitors the amount of space that is available on your local volumes.

When the amount of available space on a local volume falls below the level that you designate, Remote Storage automatically removes the content (data) from a sufficient number of eligible (or premigrated) files and migrates this content to an attached storage device, thus freeing up disk space on the volume. When data is removed from a file, the logical size of the file remains the same, and to the user, the file appears unchanged, but the physical size of the file is reduced to 1 KB. Content in other eligible files is not removed until more disk space is needed. When you need to open a file whose data has been removed, the data is automatically recalled from remote storage. To find more information about these processes, see "Remote Storage Logical Diagrams" in this section. See also "Remote Storage Terms and Definitions" and "Remote Storage Processes and Interactions" in "How Remote Storage Works."

Remote Storage runs on Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and the 64-bit versions of these operating systems. Remote Storage is also available on Windows 2000 Server. However, Remote Storage is not installed by default. To install Remote Storage, you can select Remote Storage during Setup, or you can later manually install Remote Storage through **Add or Remove Programs** in Control Panel. You can only use Remote Storage to manage NTFS volumes.

**Levels of data storage**

Remote Storage data storage is hierarchical, with two defined levels. The upper level, called local storage, includes the NTFS disk volumes of the computer running Remote Storage. The lower level, called remote storage, includes the automated media library or stand-alone tape or disk drive that is connected to the server.

Remote Storage supports all SCSI-class 4–mm tape, 8–mm tape, digital linear tape (DLT), and magneto-optical devices that are supported by Removable Storage. (To find more information about Removable Storage, see "What Is Removable Storage?." To find more information about the relationship between Remote Storage and Removable Storage, see "Remote Storage Processes and Interactions" in "How Remote Storage Works.") Remote Storage does not support quarter-inch cartridge (QIC) tape libraries, rewritable compact disks (CD-RWs), or rewritable DVDs (DVD-RWs).

**Retrieval of stored files**

When you need to access a file on a volume managed by Remote Storage, you simply open the file as usual. If the data for the file is no longer on your local volume, Remote Storage recalls the data from a media library. Because this can take more time than usual, Remote Storage removes the data only from those files on your local volumes that you are least likely to need, based on criteria that you set.

**Coordination with other tools**

Remote Storage uses Removable Storage to access the applicable tapes that are contained in libraries. Remote Storage also works with Backup for data recovery, and with Task Scheduler to schedule file copy operations. For information about Removable Storage, see "Removable Storage Technical Reference." For information about Backup, see "Backup Technical Reference."

Remote Storage also provides certain data-recovery features, including the ability to generate multiple copies of data in remote storage.

**Business Benefits**
Remote Storage provides several significant business benefits, including the following:

- A low-cost solution for archiving files that are seldom accessed, but should still be available.
- A low-cost solution that enables virtual expansion of local storage space.
- Transparent automatic access to data in remote storage.
- Automation of the labor-intensive overhead associated with daily manual data-management operations.
- Centralized sharing of remote storage devices for multiple volumes.

### Note

- Using Remote Storage to copy files to remote storage is not the same as backing up your files. You should follow a regular schedule of data backups, including backing up the contents of the local volumes that Remote Storage manages. You should also back up the Remote Storage database and other program files that are located in the System32\RemoteStorage folder.

### Note

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. In addition, do not add a volume to Remote Storage that contains directories that are part of an FRS replica set. Otherwise, you might severely impact system performance and possibly cause data loss within your media library.
- FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

**Common Scenarios for Remote Storage**
Remote Storage is especially useful in the following scenarios.

(To find more information about the terms used in these scenarios, see "Remote Storage Terms and Definitions" in "How Remote Storage Works." See also "Remote Storage Processes and Interactions" in the same topic.)

**Conserving disk space on managed volumes**
In this scenario, you want to ensure that disk space on the managed volume where your system files are located does not fall below a certain level because you are concerned that operating system efficiency will be severely compromised if this is allowed to happen.

Remote Storage initiates Automatic File Truncation whenever the level of free space on a managed volume becomes less than the specified Desired Free Space. Eligible premigrated files are truncated in order to create additional free space on this volume. The premigrated files are quickly identified and truncated, based on the last access date, and are moved to an attached storage device from which they can later be retrieved if necessary. Placeholders for these files are created and reside in local storage, and the files themselves appear unchanged to the user. Automatic File Truncation continues until the free space on the volume is greater than the Desired Free Space.

To find more information about this process, see "Remote Storage Logical Diagrams" later in this section.

**Extending disk space on managed volumes**
In this scenario, you want to make additional space available for new, more frequently accessed files, while retaining older, less frequently accessed files that you might need to access some time in the future.

By using Scheduled File Truncation, you can force the truncation of premigrated files. The premigrated files that meet the selection criteria are truncated, regardless of the available space on the managed volume. Premigrated files are checked to make sure that they have not been modified since they were premigrated. If a premigrated file has been modified, then the file is not considered premigrated and the file is returned to normal file status. Additionally, the data for the premigrated file in remote storage (called the "unnamed data attribute") is marked as "deleted" to facilitate the reclaiming of space on the device.

With Scheduled File Truncation, you can manage disk space proactively. The action can be scheduled to reduce volume usage before a known volume-intensive event, such as the installation of a very large application.

To find more information about this process, see "Remote Storage Logical Diagrams" later in this section.

**Generating multiple media copies of removable media**
In this scenario, you want ensure that there exists one or more copies of your migrated files on different media in case the original storage medium (disk or tape) becomes unusable.

Media in remote storage are copied to protect data. Copies are updated at the same time as the original media are updated to ensure that a current copy of data is available.

Media copies are replicas that can be substituted for the original media. Data is migrated only to the original media, but can be recalled from either an original or a copy.

Copies are made before the original media are completely filled and are updated as the original is updated. When a media copy is complete, it can be removed from the library located with other media in the copy set, and a new media copy is started.

Remote Storage creates media copies only if a library has two or more available drives. Other media, for which copies cannot be made, must be protected by different means, such as strict physical or environmental safeguards. To replace an original with a copy, you must manually perform the action. You should uniquely identify media copies in order to avoid confusion among different versions.

When a media copy is used to replace damaged or lost remote storage media, Remote Storage automatically copies the media again to ensure a complete set of copies.

**Replacing damaged removable media while Remote Storage is running**
In this scenario, you want to ensure that when you replace a copy of your migrated files with another copy of the same files, Remote Storage will, at the same time, automatically generate additional copies on different media.

When a media copy is used to replace damaged or lost remote storage media, Remote Storage automatically copies the media again to ensure a complete set of copies.

Remote Storage creates media copies only if a library has two or more available drives. Other media, for which copies cannot be made, must be protected by different means, such as strict physical or environmental security. To replace an original with a copy, you must manually perform the action. You should uniquely identify media copies in order to avoid confusion among different versions.

**Recovering from loss of Remote Storage metadata**
In this scenario, you want to ensure that you can retrieve any metadata associated with your Remote Storage files.

Remote Storage metadata exists in files that can be protected by the primary backup system (Backup) for Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition. Windows Server 2003 file security ensures that only the account used by Remote Storage has access to the metadata. Therefore, the administrator must assign appropriate user rights to this account so that only certain individuals (or groups) have access to this metadata. Otherwise, the metadata might be at risk, and if deleted, would no longer be retrievable.

**Selecting the Appropriate Storage Medium: Two Different Scenarios**
As part of deciding whether or not to use Remote Storage, it is important to select the appropriate storage medium. Remote Storage supports magnetic tape and optical discs as media.

The following section compares the advantages and disadvantages of using magnetic tape or optical discs for your remote storage media and then presents two different scenarios as examples of the decision-making process.

*Advantages of Using Magnetic Tape*

The advantages of using magnetic tape include:

- Magnetic tape drives and media are less expensive than optical disc changers and media.
- Magnetic tape drives and media take up less room per megabyte than optical changers and associated media.
- Magnetic tapes are available in several different sizes.

*Disadvantages of Using Magnetic Tape*

The disadvantages of using magnetic tape include:

- A magnetic tape changer takes longer to recall files that have been migrated to Remote Storage. (Magnetic tapes are sequential access media. Depending on the location of a file on the tape, it can take longer to recall a file.)
- The shelf life of magnetic tapes is substantially shorter than that of optical discs.
- Files in Remote Storage cannot be recalled while a media copy to magnetic tape is in progress.
- Magnetic tapes require the following special environmental considerations:
    - Backup media lasts longer in cool, humidity-controlled locations.
    - Backup media must be stored in an area free of magnetic fields, such as those near the backs of computer terminals and analog telephone equipment.

*Advantages of Using Optical Media*

The advantages of using optical media include:

- Optical disc changers recall files from Remote Storage faster than magnetic tape changers. (Optical discs are random access media. Regardless of where the files exist on a disc, discs enable faster seek times while recalling files.)
- Optical media has a longer shelf life than magnetic tape.
- Files in Remote Storage can be recalled while a media copy is in progress.
- Optical discs do not require special environmental considerations.

*Disadvantages of Using Optical Media*

The disadvantages of optical media include:

- Optical disc changers and media are more expensive than magnetic tape drives and media.
- Optical disc changers and media take up more room per megabyte than tape changers and associated media.
- Optical media are available in fewer different sizes.

The following scenarios provide examples of when tape or optical media might be appropriate for an organization.

**Magnetic Tape Scenario**
You are a network administrator at an urgent care medical facility. One of your responsibilities is to maintain the health records of patients who are treated in your facility. A typical patient is treated and released without the need for follow-up care. The patient files are not likely to be recalled from Remote Storage, and recalling these files quickly is not necessary. You decide to use magnetic tape as the most cost-effective storage solution for your organization.

**Optical Disc Scenario**
Your role in the IT department of a large insurance company is to archive vast amounts of customer information. It is extremely important that this information be readily accessible and that the media on which the information is stored have as long a shelf life as possible. You decide to use optical media as the most practical storage solution for your organization.

**Remote Storage Dependencies on or Interactions with Other Technologies**
Remote Storage depends on, or interacts with, the following technologies:

- Microsoft Management Console
- Removable Storage

- Backup (or similar, non-Microsoft data-management programs)
- Win32 tape and disk management APIs
- Media libraries
- The registry
- Event Viewer
- Windows Explorer
- Windows security
- NTFS
- Task Scheduler
- Microsoft Systems Management Server (SMS)
- Active Directory

To find more information about this topic, see "Remote Storage Architecture" in "How Remote Storage Works."

**Remote Storage Logical Diagrams**
The following figures show how files are migrated to, and later recalled from, remote storage.

In the first figure, files on a user's hard drive are premigrated when their unnamed data attributes are copied to an attached storage device. These files are now eligible for truncation.

**Remote Storage Logical Diagram - Files Are Premigrated**

**Step 1: Files Are Premigrated**




In the second figure, the files that have been premigrated are automatically truncated and then migrated to an attached storage device when the level of available space on the hard drive falls below the Desired Free Space.

**Remote Storage Logical Diagram - Files Are Truncated and Migrated**

**Step 2: Premigrated Files Are Truncated and Migrated**




In the third figure, the files are recalled from the attached storage device when the user accesses the files on the hard drive. The files on the hard drive appear unchanged to the user, even though the data was removed during Step 2.

**Step 3: User Recalls Migrated Files**




**Related Information**

The following resources contain additional information that is relevant to this section.

- Resource Kit Tools Help in the Tools and Settings Collection.
- Command-Line References in the Tools and Settings Collection.

# How Remote Storage Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Remote Storage Works**

**In this section**


- Remote Storage Terms and Definitions
- Remote Storage Architecture
- Remote Storage Protocols
- Remote Storage Interfaces
- Remote Storage Processes and Interactions
- Network Ports Used by Remote Storage
- Related Information


You can use Remote Storage to extend disk space on your server without adding more hard disks. Remote Storage automatically copies eligible files on your local volumes to a library of magnetic tapes or magneto-optical disks. Remote Storage then monitors the amount of space that is available on your local volumes.

When the amount of available space on a local volume falls below the level that you designate, Remote Storage automatically removes the content (data) from a sufficient number of eligible (or premigrated) files and migrates this content to an attached storage device, thus freeing up disk space on the volume. When data is removed from a file, the logical size of the file remains the same, and to the user, the file appears unchanged, but the physical size of the file is reduced to 1 KB. Content in other eligible files is not removed until more disk space is needed. When you need to open a file whose data has been removed, the data is automatically recalled from remote storage. To find more information about these processes, see "What Is Remote Storage?" in "What is Remote Storage?" See also "Remote Storage Terms and Definitions" and "Remote Storage Processes and Interactions" in this section.

Remote Storage runs on Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and the 64-bit versions of these operating systems. Remote Storage is also available on Windows?2000 Server., However, Remote Storage is not installed by default. To install Remote Storage, you can select Remote Storage during Setup, or you can later manually install Remote Storage through **Add or Remove Programs** in Control Panel. You can only manage Remote Server on NTFS volumes.

**Levels of data storage**

Remote Storage data storage is hierarchical, with two defined levels. The upper level, called local storage, includes the NTFS disk volumes of the computer running Remote Storage. The lower level, called remote storage, includes the automated media library or stand-alone tape or disk drive that is connected to the server.

Remote Storage supports all SCSI-class 4–mm tape, 8–mm tape, digital linear tape (DLT), and magneto-optical devices that are supported by Removable Storage. (To find more information about Removable Storage, see "What Is Removable Storage? To find more information about the relationship between Remote Storage and Removable Storage, see "Remote Storage Processes and Interactions" in this section.) Remote Storage does not support quarter-inch cartridge (QIC) tape libraries, rewritable compact disks (CD-RWs), or rewritable DVDs (DVD-RWs).

**Retrieval of stored files**

When you need to access a file on a volume managed by Remote Storage, you simply open the file as usual. If the data for the file is no longer on your local volume, Remote Storage recalls the data from a media library. Because this can take more time than usual, Remote Storage removes the data only from those files on your local volumes that you are least likely to need, based on criteria that you set.

**Coordination with other tools**

Remote Storage uses Removable Storage to access the applicable tapes contained in libraries. Remote Storage also works with

Backup for data recovery, and with Task Scheduler to schedule file copy operations. For information about Removable Storage, see "Removable Storage Technical Reference". For information about Backup, see "Backup Technical Reference."

Remote Storage also provides certain data-recovery features, including the ability to generate multiple copies of data in remote storage.

**Business Benefits**
Remote Storage provides several significant business benefits, including the following:

- A low-cost solution for archiving files that are seldom accessed, but should still be available.
- A low-cost solution that enables virtual expansion of local storage space.
- Transparent automatic access to data in remote storage.
- Automation of the labor-intensive overhead associated with daily manual data-management operations.
- Centralized sharing of remote storage for multiple volumes.

### Note

- Using Remote Storage to copy files to remote storage is not the same as backing up your files. You should follow a regular schedule of data backups, including backing up the contents of the local volumes that Remote Storage manages. You should also back up the Remote Storage database and other program files located in the System32\RemoteStorage folder.

### Notes

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. In addition, do not add a volume to Remote Storage that contains directories that are part of an FRS replica set. Otherwise, you might severely impact system performance and possibly cause data loss within your media library.
- FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

The following sections provide an in-depth view of how Remote Storage works in an optimal environment. An optimal environment for Remote Storage is defined as follows:

- Windows Server 2003 is correctly installed and Remote Storage is enabled.
- All other dependencies are in place, properly designed and deployed, and functioning normally.
- Server hardware is sized appropriately and there are no disk, CPU, memory, or network bottlenecks that affect the performance of the technology.

**Remote Storage Terms and Definitions**
The following terms are associated with Remote Storage.

**Managed volume**
A local volume formatted with NTFS file system 5.0 whose disk space is managed by Remote Storage. Remote Storage frees up disk space by automatically moving infrequently accessed files to a remote storage device.

**Desired Free Space**
The amount of free space that you specify a volume should have at all times during normal use.

**Eligible file**
A file that has been premigrated, and is therefore ready to be truncated.

**File selection criteria**
User-defined criteria that determine which files Remote Storage will premigrate. One such criterion might be that all files that are .doc files should be premigrated.

**Unnamed data attribute**
The default data stream of an NTFS file, sometimes referred to as $DATA. The unnamed data attribute is copied to remote storage during premigration.

**Premigrated file**
A file whose unnamed data attribute has been copied to a remote storage device in preparation for truncation. Until truncation and migration take place, this file remains unchanged on the managed volume.

**Truncated file**
A file on a managed volume from which Remote Storage has removed data so that the physical file size is reduced to 1 KB. This file appears unchanged to the user.

**Migrated file**
A file that Remote Storage has copied from local storage to remote storage. The resulting file in remote storage is not visible to users. The file remaining in local storage is truncated.

**Automatic File Truncation**
The process in which data is removed from a sufficient number of premigrated files when available disk space on a managed volume falls below the required level. This process occurs without user intervention.

**Scheduled File Truncation**
The process in which a user forces data to be removed from premigrated files even if available disk space has not fallen below the required level.

**Placeholder**
An identifier for a file that has been migrated from local storage to remote storage. Placeholders are not visible to users.

**Remote storage**
The term "remote storage" (lower case) refers to infrequently accessed data and the media on which it is stored until the data is needed again.

**Removable Storage**
A service that you can use to manage removable media (such as tapes and discs) and storage devices (libraries). Removable Storage enables applications to access and share the same media resources. For a complete listing of Removable Storage terms and definitions, see MSDN and type Removable Storage Glossary in the "Search for" text box.

**Remote Storage Architecture**
From an architectural standpoint, Remote Storage is divided into three areas:

- **Remote Storage components:** The Remote Storage Microsoft Management Console (MMC) snap-in and the Remote Storage service.
- **Windows components and services:** Win32 tape and disk APIs, tape and disk drivers, changer drivers, Event Viewer, and the registry.
- **Libraries:** A robotic or automated library, stand-alone tape drive, and stand-alone disk driver, for example.

The following figure shows the components within these three areas and how the areas and components interact.

**Remote Storage Architectural Diagram**




The following table lists the components that are included in Remote Storage.

**Remote Storage Components**

| Component | Description |
|---|---|
| Remote Storage MMC snap-in | Administrative interface to Remote Storage. The Remote Storage snap-in enables you to:<br>• Set the device type while running the Remote Storage Setup Wizard the first time you start the snap-in, and subsequently view media state and utilization.<br>• Set Remote Storage system-wide feature options.<br>• View information about Remote Storage activity.<br>• Recover from media disasters. |
| Remote Storage service | Two system services that implement Remote Storage functionality.<br>The primary service starts automatically. The second service is used for recall notifications. |

The following table lists the items that are included in Windows components and services.

**Windows Components and Services**

| Service | Description |
|---|---|
| Win32 tape and disk management APIs | Comprises four different APIs: CreateFile, DeleteFile, GetDiskFreeSpace, and GetDiskFreeSpaceEx.<br>To find more information about these APIs, see "How Removable Storage Works" in "How Removable Storage Works." |

| | |
|---|---|
| Tape and disk driver | Enables a specific device, such as a stand-alone tape or disk drive, to communicate with the operating system. Although a device might be installed on your system, Windows cannot use the device until you have installed and configured the appropriate driver. If a device is listed in the Windows Catalog, a driver is usually included with Windows. Device drivers load automatically (for all enabled devices) when a computer is started, and thereafter run in a manner transparent to users.<br><br>(The Windows Catalog website has hardware and software compatibility information for the Windows 2000 Server family, Windows XP, and the Windows Server 2003 family.) |
| Changer driver | Enables a specific device, such as an automated or robotic library, to communicate with the operating system. Although a device might be installed on your system, Windows cannot use the device until you have installed and configured the appropriate driver. If a device is listed in the Windows Catalog, a driver is usually included with Windows. Device drivers load automatically (for all enabled devices) when a computer is started, and thereafter run in a manner transparent to users. |
| Event Viewer | A component that you can use to view and manage event logs, gather information about hardware and software problems, and monitor security events. Event Viewer maintains logs about program, security, and system events.<br><br>Recalls of files from Remote Storage are logged in the Windows event log. |
| The registry | A database repository for information about the configuration of a computer. The registry contains information that Windows continually references during operation, such as:<br><br>• Profiles for each user.<br>• Programs that are installed on the computer and the types of documents each can create.<br>• Property settings for folders and program icons.<br>• Hardware that exists on the system.<br>• Ports that are being used.<br><br>The registry is organized hierarchically as a tree and is made up of keys and their subkeys, hives, and entries. |

The following table lists the components that are included in libraries.

**Libraries**

| Library | Description |
|---|---|
| Robotic or automated library | Automated units that hold multiple tapes or disks. Some have multiple drives. These libraries are sometimes called *changers* or *jukeboxes*, and commonly use robotic subsystems to move media stored in the library's storage slots. Automated libraries can also consist of other hardware components that are managed by Removable Storage, such as doors, inject/eject ports, cleaner cartridges, and bar-code readers. |
| Stand-alone tape drive | A single-drive, nonautomated unit that holds a single tape. With this device, you manually insert a tape into the unit. |
| Stand-alone magneto-optical disk drive | A single-drive, nonautomated unit that holds a single disk. With this device, you manually insert a disk into the unit. |

**Remote Storage Protocols**

Remote Storage does not utilize any networking protocols.

**Remote Storage Interfaces**

There are no Remote Storage interfaces available.

**Remote Storage Processes and Interactions**

The following is a general description of how Remote Storage functions at a basic level. (To find more information about the terms used in this section, see "Remote Storage Terms and Definitions" earlier in this section.)

Remote Storage is a hierarchical storage management application that migrates data from *upper-level storage* to *lower-level storage*. Upper-level storage is commonly known as local storage: frequently accessed data that is stored locally on high-performance disks. Lower-level storage is also known as remote storage: infrequently accessed data that is stored on less expensive media until it is needed again. Hierarchical storage management reduces the costs of storing large quantities of data, while ensuring that the data is still accessible. The following table outlines when to use the two levels of Remote Storage.

**Levels of Remote Storage**

| Considerations | Upper-level Storage | Lower-level Storage |
|---|---|---|
| **Storage type** | Local or primary | Remote storage |
| **How often is data needed?** | Frequently | Less frequently |
| **Media type** | High-performance disks | CD-ROMs<br>DVD-ROMs<br>Magnetic tape, which includes:<br>• Digital audio tape (DAT)<br>• Digital linear tape (DLT)<br>Magneto-optical disk drive |
| **Storage cost** | More expensive | Less expensive |

Regardless of where you store the data, the file system namespace continues to provide users with access to the file. For example, when a user double-clicks a file in My Documents, Remote Storage retrieves the file from its storage location and restores it to the file system.

For users to access files that are stored on offline media, the media must be placed back online before the files are accessible. Therefore, these files require more time to access than files that are stored on online media.

The following is a more in-depth examination of how Remote Storage works at a basic level.

To prepare for the need to free storage space on a local server, Remote Storage regularly premigrates (copies) the unnamed data attribute of all eligible NTFS files on local storage to remote storage. After this is done, the files on local storage are considered to be premigrated. The unnamed data attribute of an NTFS file contains information that is tracked by NTFS. This information is needed to manage the data. To find more information about this process, see "What Is Remote Storage?" in "What Is Remote Storage?"

When the amount of available space on a local volume falls below the Desired Free Space, Remote Storage automatically removes the content (data) from a sufficient number of eligible files and migrates this content to an attached storage device, thus freeing up disk space on the volume. This process is referred to as truncation. When data is removed from a file, the disk space used by the file is reduced to 1 KB, although the file itself appears unchanged. At the same time, a "placeholder" is generated for each such file and this placeholder serves as an identifier for its newly-created counterpart on the attached storage device. Remote Storage marks the placeholder with the flag, FILE_ATTRIBUTE_OFFLINE. Content in other eligible files is not removed until more disk space is needed. When you need to open a file whose data has been removed, the data is automatically recalled from remote storage.

Even though Remote Storage changes the physical size of the file on local storage, the logical size and the date and time (created, last modified, last accessed) of the file remain unchanged. Users or applications (for example, disk quotas) that use file size as a determinant of a particular action are not affected by Remote Storage data migration.

For these processes and interactions to work as described, Removable Storage should reside in an optimal environment. An optimal environment for Removable Storage is defined as follows:

- Windows Server 2003 is correctly installed and the Removable Storage service is running.

- All other dependencies are in place, properly designed and deployed, and functioning normally.
- Server hardware is sized appropriately and there are no disk, CPU, memory, or network bottlenecks that affect the performance of the technology.

**Remote Storage Technologies**
The following is a more detailed description of specific technologies, features or attributes that either form a part of, or are closely identified with, Remote Storage.

**Removable Storage**
Remote Storage uses Removable Storage for copying data to online libraries that contain removable media. For example, if a volume that is being managed by Remote Storage meets the requirements for migrating data to secondary storage, Removable Storage moves media into a slot so that Remote Storage can migrate data to that media. To find more information about this topic, see "How Removable Storage Works" in "How Removable Storage Works."

**Disk Management Attribute**
In Remote Storage, you can define a single, system-wide, disk management attribute that can be applied to each managed volume. This attribute is referred to as the Desired Free Space setting and establishes a threshold that triggers Automatic File Truncation. The purpose of Automatic File Truncation is to free up disk space on a managed volume when the amount of space falls below a certain percentage. The default setting is five percent, but this can be adjusted to any percentage that you specify. In addition, you can set the level of Desired Free Space for several different managed volumes all at the same time.

### Notes

- It is important to realize that although the Desired Free Space setting is used to determine when files that have already been premigrated should be truncated, it is not used to determine how many files should be premigrated before the need for truncation occurs.
- In addition, by default, Remote Storage does not premigrate the Windows system files (or any files marked with the system attribute).

**Premigration**
Files are selected for premigration based on the following criteria:

- Path name
- File names, including wildcard characters in names
- Inclusion or exclusion of premigration specifications

The following criteria can be applied to an entire volume, but not to an individual file:

- Last access date or inactivity period
- Owner
- File size
- System and hidden files

You use these criteria to specify which files will be premigrated and, consequently, become eligible files. Eligible files are files that are truncated as the need for additional disk space arises. Remote Storage premigrates all files that meet the criteria and tracks these files to allow rapid file truncation. Even though Remote Storage changes the physical size of the file in local storage, the logical size and the dates and times (created, last modified, last accessed) of the file remain unchanged.

Remote Storage applies the inclusion and exclusion rules in such a manner that the more specific a rule, the higher its priority. For example, the rules of "Exclude *.*" and "Include \PROG*.EXE" causes the file "\Prog\Test.exe" to be included and "\Test.exe" to be excluded.

**Automatic File Truncation**
Remote Storage initiates Automatic File Truncation whenever the level of free space in a managed volume becomes less than the Desired Free Space. Eligible premigrated files are truncated to free space on local storage. The premigrated files are quickly identified and truncated, based on the last access date. The file truncation process continues until the free space on the volume is greater than the Desired Free Space.

If the premigrated file has not been modified, then the file is truncated and a placeholder for the file is created. If the premigrated file has been modified, the file is no longer considered premigrated and is returned to "normal file" status. In addition, the capacity flow meter is decremented and the data for the premigrated file in remote storage is marked as deleted to facilitate space reclamation. Reclamation is the process of deleting from the Remote Storage media pool any files that are no longer necessary, thereby freeing up additional storage space.

### Note

- Storage maintains a flow meter capacity metric that monitors the amount of data that is contained in remote storage, including overflowed and shelved data. Copying data from local storage to remote storage increases the value of the capacity metric. Removing local storage references to data in remote storage decreases the value of the capacity metric.

**Scheduled File Truncation**
Scheduled File Truncation is the process in which a user forces data to be removed from premigrated files even if available disk space has not fallen below the required level. Premigrated files are checked to make sure that they have not been modified since they were premigrated. If the premigrated file has been modified, then the file is not considered premigrated and the file is returned to normal file status. Additionally, the data for the premigrated file in remote storage is marked as deleted to facilitate space reclamation.

You can use Scheduled File Truncation to manage disk space proactively. The action can be scheduled to reduce volume usage in advance of a known volume-intensive event, such as the installation of a very large application.

**Validation of Managed Files**
You can use Validate Managed Files to validate premigrated files and placeholders, thus making sure that all the files on your managed volumes point to valid and correct data in remote storage. Validation includes detecting placeholders moved between volumes, disconnected placeholders, and modified premigrated files. Corrective actions are taken immediately to synchronize the placeholders and premigrated files with remote storage. Validation also recalculates volume statistics. Validation is automatically performed two hours after a backup program is used to restore a remote storage file, but can also be scheduled manually.

You should perform validation a regular basis to validate all files on a given volume and to correct any inconsistencies. You should also perform validation after restoring files on a local volume or after disk errors have occurred on a volume.

**File Recall**
The file is recalled on the first read, write, or memory map request. If the data is on tape or disk in remote storage, but the tape or disk is not in the library (storage device), the recall does not succeed and returns the message STATUS_FILE_IS_OFFLINE. In order to successfully recall the file, you need to manually insert the tape or disk into the library and then try the recall again.

### Notes

- Remote Storage does not affect a request that is made against a premigrated file. The data for the file exists locally on the NTFS volume and can be accessed normally. However, the client making the request on the server might time out. Client computers have their own configured time-out that is independent of Remote Storage, and if accessing the data takes longer than the time-out, the input/output request is unsuccessful for the client.

- Changing attributes, access control lists (ACLs), or named streams in a file does not recall the file.

**Runaway Recall Detection and Prevention**
*Runaway recalls* can be prevented by specifying that no more than a certain number of files can be recalled by a user from a volume in an hour. An administrator defines this setting, both in terms of enabling or disabling the feature, and in specifying the number of recall requests that are handled in a particular period of time.

**Renaming, Moving, or Copying Placeholders and Premigrated Files**
Migrated files (and their associated placeholders) and premigrated files can only be renamed on the same volume. Files cannot be renamed across volume boundaries. It is important to note that renaming a file does not cause the data in remote storage to be recalled.

However, it is possible to copy or move placeholders across volume boundaries. When this occurs, the data in remote storage is recalled. At the end of a *move* operation between volumes, the original placeholder is deleted and the file information that has been recalled and is now present on the other volume once again assumes a premigrated status. At the end of a *copy* operation between volumes, the original placeholder is also deleted, but there now exist two files, the first on the original volume and the second on the other. The same process applies when copying a placeholder to a different location on the same volume: the data in remote storage is recalled, the original placeholder is deleted, and there now exist two identical files, both on the same volume.

### Note

- In order to prevent file information from being recalled when moving or copying a file to a different volume on the same computer, you should use Backup to back up and restore the file to its new location, rather than moving or copying the file directly. Despite being moved or copied, the file will continue to point to the correct location in remote storage.

**Discontinuing Volume Management**
If necessary or desirable, you can specify that Remote Management no longer manage a volume. In doing this, you are given two choices:

- Remote Storage discontinues managing the local volume, but leaves data in remote storage. Data that is already in remote storage can be recalled as before. No automatic freeing of disk space takes place.

- Remote Storage recalls data from remote storage, and then discontinues managing the local volume. You should only select this option if your local volume has adequate disk space for the recalled data.

If you have adequate disk space to recall data from remote storage, but the option of recalling all data is disabled, click the **Refresh** button in order to update volume statistics. After you have updated the volume statistics, the option should be enabled.

**Restoring Placeholders from Backup and Disconnected Placeholders**
The data in remote storage is accessible only through a placeholder. If the placeholder is deleted, that data is inaccessible. An upper-level storage backup application, such as Backup, protects placeholders. If the placeholder is deleted, lost, or destroyed, the only recourse is to retrieve a copy from the backup media. If a placeholder is restored from backup, the placeholder is reconnected to that data. This is the same mechanism that you must use if any other file on the volume is deleted, lost, or destroyed.

Disconnected placeholders are placeholders whose file contents have been removed from remote storage. These placeholders might have been restored from backup after the space in remote storage was reclaimed, or the data within remote storage might be physically unavailable because of media failures. Disconnected placeholders do not point to valid locations in remote storage. By validating placeholders, you remove disconnected placeholders from the system and synchronize the remaining placeholders with remote storage.

**Interaction with Other Reparse Point Types**
Volume mount points and symbolic links cannot be changed into placeholders. The system explicitly ignores these types of files. Any file system scan that encounters a mount point or symbolic link does not follow the link.

**Placeholders Binding to a Specific Remote Storage Engine**
Placeholders identify the specific remote storage system that contains the managed data, binding a placeholder to a specific Remote Storage engine. Placeholders cannot be moved between multiple remote storage systems.

**Windows Server 2003 Family Security Integration**
Remote Storage supports NTFS security features. Remote Storage recalls files only if the user has valid access to a placeholder. Because a placeholder is identified as a special remote storage reparse point, the security of the reparse point mechanism determines the security of the placeholder.

Administrators and users undergo the normal security checks for Windows users. The Remote Storage Administrative user interface uses Windows permissions to grant or deny access to the interface. Only user accounts in the Administrators group are allowed access to the UI functions.

**Grouping of Files**
Files are grouped together by their volume and by the time at which they were originally premigrated. The files remain grouped during all remote storage operations. No other grouping of files is supported.

**Replication**

By using the file selection criteria, the administrator can exclude from the Manage Files action any files that are being replicated.

**Remote Storage Engine**
Remote storage is the part of the system that holds the copies of the data. The Remote Storage engine consists of libraries, drives, and the media used to store the data.

The Remote Storage engine service runs only on computers that are running Windows Server 2003, Datacenter Edition; Windows Server 2003, Enterprise Edition; or Windows 2000 Server. Only one Remote Storage engine can exist on a server.

**Flow meter capacity metric**
Remote Storage maintains a flow meter capacity metric that monitors the amount of data that is contained in remote storage, including data that has overflowed onto other media or is located on media that are no longer in the device and have been shelved. Copying data from local storage to remote storage increases the value of the capacity metric. Removing local storage references to data in remote storage decreases the value of the capacity metric.

Remote Storage increases the value of the flow meter capacity metric whenever one of the following occurs:

- A new file is managed.
- A previously deleted placeholder is restored to a volume and reconnected to remote storage.

Remote Storage decreases the value of the flow meter capacity metric whenever one of the following occurs:

- A premigrated file is modified.
- A placeholder or premigrated file is deleted.
- A file is unmanaged.

There might be a delay in changing the value of the flow meter capacity metric because detection of these conditions can happen at a later time.

**Media pool for remote storage**
All remote storage in Remote Storage exists in a single storage or media pool. A media pool is defined as a unit of storage that is administered by Removable Storage and composed of homogeneous storage media. A media pool is a self-contained storage area with homogeneous characteristics — for example, random access, sequential access, read/write, and write-once. To find more information about media pools, see "How Removable Storage Works" in "How Removable Storage Works."

Remote Storage supports one storage library within the single storage pool. The library must have at least one drive. To support media copies, it must have two or more drives. When multiple drives exist within a library, all drives must be identical in type.

### Note

- Relocation is the movement of data from one remote storage pool to another. Because Remote Storage only supports a single remote storage pool, relocation is not supported in Remote Storage.

**Handling remote storage that is full**
Remote Storage provides the following options for handling situations in which remote storage is physically full:

- **Stop managing new files.** Remote Storage becomes a recall-only system. This mode is invoked only after any possible space reclamation is performed.
- **Shelf storage.** Remote Storage cycles removable media to shelf (offline storage). Recalls of files stored on shelved media require physical intervention.

Remote storage is considered full either when there is not enough space on the media or when the media are not online. In either case, you must physically intervene to allow Remote Storage to continue migrating files. A notice is issued when physical remote storage use reaches the warning level that is defined by the administrator.

**Shelf media for Remote Storage**
Shelf media for Remote Storage are media that are not currently accessible by the system. The medium with the data on it can be removed from the device and the media management subsystem tracks its location. Shelf media allow a system to have more data in remote storage than the device can physically contain.

For removable media, you can remove media and optionally replace them with blank media. Remote Storage continues to premigrate data by using the new media. The removed media are considered shelf media. If a user needs access to data that is on shelf media, you must place the shelf media in the library, and the new media become shelf media. Media do not need to be full to be shelf media.

**Upgrading Remote Storage libraries**
Removable Storage provides Remote Storage with the ability to replace a library with another library that contains identical media and data read/write mechanisms. You must physically move the media from the old library to the new library, that is, physically replace the library, and then use the Remote Storage snap-in to register configuration changes within Remote Storage.

**Remote Storage Device Support**
Remote Storage supports all of the same 4-mm, 8-mm, and digital linear tape (DLT) and magneto-optical disc devices that Removable Storage supports. Because Remote Storage supports only a single media type, a device cannot be a mixed media device.

Removable Storage uses Windows kernel-mode drivers to support the management of tape and optical disc devices. These device drivers provide the backbone of device support for both Removable Storage and Remote Storage. The Windows Logo program identifies device drivers that have been approved for use with the Windows Server 2003 family.

Remote Storage cooperatively shares storage libraries with external applications by using a common media management facility, such as Removable Storage.

### Note

- Remote Storage does not support using small-capacity removable disk devices such as the Iomega Zip drive.

**Remote Storage Integration in Windows Server 2003**
Remote Storage is integrated closely with the NTFS file system in Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition. The following table provides details of this integration and describes the interfaces between Remote Storage and these operating systems.

### Remote Storage and Windows Server 2003 Integration

| Interface | Description |
|---|---|
| Placeholder format and reparse points | A placeholder has the system-defined $REPARSE_POINT attribute set with information that can identify and retrieve the unnamed data attribute from remote storage. The unnamed data attribute is in remote storage and its length is zero bytes. The reparse point type is “Remote Storage”.<br><br>The recall mechanism is based on the reparse point being identified by NTFS and is an efficient mechanism that uses this information when recalling a file. |
| File size of a placeholder | The logical file size of the unnamed data attribute of a placeholder is the same size as it was before it was truncated. The physical size of the unnamed data attribute of a placeholder is zero bytes. |
| Disk quotas and remote storage | Disk quotas do not change. Disk quotas are based on the logical file size that is not changed by Remote Storage. Therefore, the amount of space in use on a volume is not reduced or changed by having data in remote storage. |
| Remote Storage stage change | Remote Storage supports a Distributed Component Object Model (DCOM) interface–based mechanism that can change the system status. |

| | |
|---|---|
| Event log usage | Major Remote Storage system events are logged in the Windows event log. You can use the Windows Event Viewer to see the logged events.<br><br>The following events are optionally logged in the Windows event log:<br><br>• Recalled files<br>• Managed files<br>• Truncated files<br>• Mounted media<br>• Run jobs<br>• Scanned files |
| Windows Server 2003 Job Scheduler | Remote Storage uses the Windows Server 2003 Job Scheduler to schedule Remote Storage jobs. Job status can be monitored from within the Remote Storage snap-in.<br><br>Depending on the Job Scheduler, a job window can be specified that limits the amount of time spent in a single scan. This is useful because the system might have a large amount of data to manage and the scan might initially take too much time. A bookmark is kept where the scan stopped so the scan can continue from that point the next time it is run. |
| Windows Server 2003 registry usage | Remote Storage uses the registry to keep persistent information about startup, including the programs needed to initialize Setup. |
| Windows Server 2003 client time-out on recall | The registry entry HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\Services\LanmanWorkstation\parameters \OfflineFileTimeOutIntervalInSeconds controls the time-out period that the client uses when sending requests to a server. The default is value is 900 (or 15 minutes). You can increase this value on all Windows clients that open files on the volumes managed by Remote Storage. |

**Multiple Remote Storage Interactions**
A Remote Storage installation is limited to a single system running Windows Server 2003, Enterprise Edition or Windows Sever 2003, Datacenter Edition. There is no limit, however, to the number of stand-alone Remote Storage instances that can run on a homogeneous Windows network.

**Moving data from one Remote Storage installation to another**
To move managed data from one Remote Storage installation to another, you recall managed data and move the data to another volume managed by Remote Storage.

**Choosing Products That Work with Remote Storage**
Because it manages data transparently for applications and users, Remote Storage does not require that the other applications running on NTFS be completely compatible with it, nor does it require that the applications contain specific code to recognize that Remote Storage is running. However, products that regularly open many files might cause a great deal of data recall, thus reducing the benefits provided by Remote Storage. Thus, for best results, use applications that are compatible with Remote Storage.

### Note

- Although the features included with Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition are compatible with Remote Storage, many other applications are not and can result in time-outs or runaway recalls. In all cases, check with your independent software vendor (ISV) in order to determine if they provide a version that is compatible with Remote Storage.

**Backup Programs and Remote Storage**
A backup product that is compatible with Remote Storage must perform the following actions:

- Guarantee that a full backup of the file is kept. The backup product must recognize the difference between backing up the

placeholder and backing up the full data of the file, and must keep the full backup copy.

- Back up a placeholder (reparse point) correctly.
- Use the FILE_FLAG_OPEN_BACKUP_SEMANTICS parameter when opening a file to back it up.
- Properly use FILE_FLAG_OPEN_NO_RECALL.

Backup applications that are compatible with Remote Storage protect placeholder data and do not cause recalls.

Remote Storage relies on the primary backup to protect placeholder data. You must be aware of when placeholder backups occur. Specifically, if backup rotations are used, you must know when the primary backup databases contain a complete copy of the data and when they contain placeholder data.

Backup applications that open placeholders to read and copy without using the FILE_FLAG_OPEN_BACKUP_SEMANTICS parameter cause the data for the file to be recalled. If these products are performing a full system backup, all placeholders are replaced with recalled data.

**Antivirus Programs and Remote Storage**
Antivirus programs typically open files to check for infection by viruses. These products sometimes cause migrated data to be recalled. Antivirus programs that are compatible with Remote Storage can operate in one of the following ways:

- Detect placeholders and skip the check on such files. This assumes that the scanner checked the data before premigrating it. This might be somewhat hazardous because Remote Storage manages the unnamed data stream and does not manage named data streams. Because data in remote storage remains unchanged, an additional check for viruses is unnecessary.
- Open the file with the FILE_OPEN_NO_RECALL parameter. This allows the scanner to check the unnamed data stream by using local storage without recalling the data. This assumes a linear search of the file data, which might not always be the case.
- Do not update the last access date and time.

**Document Management Programs**
Many document management products open the data that they are managing, scan for key data, and then store references in additional files or alternate data streams, leaving the original data unchanged. A document management product that is compatible with Remote Storage can open placeholders with the FILE_OPEN_NO_RECALL parameter, which allow the migrated unnamed data to be read but not recalled.

In addition, after the index files are generated, some document management products provide a file viewer much like Windows Explorer. If the product is compatible with Remote Storage, when it opens each file, it experiences synchronous recalls of each file and any associated delay. The product also indicates which files are migrated and allows the user to issue asynchronous recalls of the data, while exhibiting improved performance.

Policies for managing files of the document management system exclude index files and other special files so that they are quickly accessible to the document management system.

### Note

- The Indexing service that is included with the Windows Server 2003 family is compatible with Remote Storage.

**Remote Storage Protection and Recovery**
Media in remote storage are copied to protect data. Copies are updated as the original media are updated to ensure that a current copy of data is available.

**Media copies**
Media copies are replicas that can be substituted for the original media. Data is migrated only to the original media, but can be recalled from either an original or a copy.

Copies are made before the original media are completely filled and are updated as the original is updated. When a media copy is complete, it can be removed from the library that is located with other media in the copy set, and a new media copy is

started.

Remote Storage creates media copies only if a library has two or more available drives. Other media, for which copies cannot be made, must be protected by different means, such as strict physical or environmental safeguards. To replace an original with a copy, you must manually perform the action. You should uniquely identify media copies in order to avoid confusion among different versions.

When a media copy is used to replace damaged or lost remote storage media, Remote Storage automatically copies the media again to ensure a complete set of copies.

Partial media copies are made by the system to protect the data. These partial copies can be removed from the library and are updated when they are reinserted in the library. If the source of the media copy is on the shelf, the source media must be reinserted in the library for the completion of the media copy.

**Copy sets and set rotation**
Remote Storage supports up to three copy sets of remote storage media. Media for the additional sets can be moved to shelf storage when they are synchronized. You can view the status of media copies.

**Media for media copies**
If a tape library has two or more available drives, Remote Storage uses scratch media for making a media copy. You can remove the updated media copies after several new media copies have been made. This is especially useful on weekends, during vacations, or following initial startup. After a copy is made, it must be returned to the library to be updated.

**Metadata Protection**
Remote Storage metadata exists in files that can be protected by the primary backup system (Backup) for Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition. Windows Server 2003 file security ensures that only the account used by Remote Storage has access to the metadata. It is therefore necessary for the administrator to assign appropriate user rights to this account so that only certain individuals (or groups) have access to this metadata. Otherwise, the metadata might be at risk, and if deleted, would no longer be retrievable.

**Network Ports Used by Remote Storage**
Remote Storage does not utilize any network ports.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Resource Kit Tools Help in the Tools and Settings Collection.
- Command-Line References in the Tools and Settings Collection.
- Group Policy Settings Reference in the Tools and Settings Collection.
- To find more information about Removable Storage interfaces, see MSDN and type “Disk Management Functions” in the “Search For” text box.

# Remote Storage Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Storage Tools and Settings**

**In this section**


- Remote Storage Tools
- Remote Storage Registry Entries
- Remote Storage Group Policy Settings
- Remote Storage WMI Classes
- Network Ports Used by Remote Storage
- Related Information


Remote Storage enables you to extend disk space on your server without adding more hard disks. Remote Storage automatically copies eligible files on your local volumes to a library of magnetic tapes or magneto-optical disks. Remote Storage then monitors the amount of space available on your local volumes.

When the amount of available space on a local volume falls below the level that you designate, Remote Storage automatically removes the content (data) from a sufficient number of eligible (or premigrated) files and migrates this content to an attached storage device, thus freeing up disk space on the volume. When data is removed from a file, the logical size of the file remains the same, and to the user, the file appears unchanged, but the physical size of the file is reduced to 1 KB. Content in other premigrated files is not removed until more disk space is needed. When you need to open a file whose data has been removed, the data is automatically recalled from remote storage. To find more information about these processes, see What Is Remote Storage? in "What is Remote Storage?" See also How Remote Storage Works and How Remote Storage Works in "How Remote Storage Works."

Remote Storage runs on Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and the 64-bit versions of these operating systems. Remote Storage is also available on Windows 2000 Server. However, Remote Storage is not installed by default. To install Remote Storage, you can select Remote Storage during Setup, or you can later manually install Remote Storage through **Add or Remove Programs** in Control Panel. You can only use Remote Storage to manage NTFS volumes.

**Levels of data storage**

Remote Storage data storage is hierarchical, with two defined levels. The upper level, called local storage, includes the NTFS disk volumes of the computer running Remote Storage. The lower level, called remote storage, includes the automated media library or stand-alone tape or disk drive that is connected to the server.

Remote Storage supports all SCSI-class 4–mm tape, 8–mm tape, digital linear tape (DLT), and magneto-optical devices that are supported by Removable Storage. (To find more information about Removable Storage, see What Is Removable Storage? To find more information about the relationship between Remote Storage and Removable Storage, see How Remote Storage Works in "How Remote Storage Works.") Remote Storage does not support quarter-inch cartridge (QIC) tape libraries, rewritable compact disks (CD-RWs), or rewritable DVDs (DVD-RWs).

**Retrieval of stored files**

When you need to access a file on a volume managed by Remote Storage, you simply open the file as usual. If the data for the file is no longer on your local volume, Remote Storage recalls the data from a media library. Because this can take more time than usual, Remote Storage removes the data only from those files on your local volumes that you are least likely to need, based on criteria that you set.

**Coordination with other tools**

Remote Storage uses Removable Storage to access the applicable tapes that are contained in libraries. Remote Storage also works with Backup for data recovery, and with Task Scheduler to schedule file copy operations. For information about Removable Storage, see Removable Storage Technical Reference For information about Backup, see Backup Technical

## Reference

Remote Storage also provides certain data-recovery features, including the ability to generate multiple copies of data in remote storage.

**Business benefits**
Remote Storage provides several significant business benefits, including the following:

- A low-cost solution for archiving files that are seldom accessed, but should still be available.
- A low-cost solution that enables virtual expansion of local storage space.
- Transparent automatic access to data in remote storage.
- Automation of the labor-intensive overhead that is associated with daily manual data-management operations.
- Centralized sharing of remote storage devices for multiple volumes.

### Note

- Using Remote Storage to copy files to remote storage is not the same as backing up your files. You should follow a regular schedule of data backups, which includes backing up the contents of the local volumes that Remote Storage manages. You should also back up the Remote Storage database and other program files located in the System32\RemoteStorage folder.

### Note

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. In addition, do not add a volume to Remote Storage that contains directories that are part of an FRS replica set. Otherwise, you might severely impact system performance and possibly cause data loss within your media library.
- FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

**Remote Storage Tools**
The following tools are associated with Remote Storage.

**Remote Storage Microsoft Management Console (MMC) Snap-in**
**Category**
Part of Windows Server 2003.

**Version compatibility**
Runs on Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and the 64-bit versions of these operating systems. It is also available on Windows 2000 Server.

The Remote Storage MMC snap-in is an administrative interface to Remote Storage. With the Remote Storage snap-in, you can:

- Set the device type while running the Remote Storage Setup Wizard the first time you start the snap-in, and subsequently view media state and utilization.
- Set Remote Storage system-wide feature options.
- View information about Remote Storage activity.
- Recover from media disasters.
- Create and submit jobs.

**Rss.exe: Rss**
**Category**
Part of Windows Server 2003.

**Version compatibility**
Runs on Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and the 64-bit versions of these operating systems. It is also available on Windows 2000 Server.

You can use Rss.exe to manage Remote Storage from the command line. Using the **rss** command, you can run batch scripts for applications that will allow them to access Remote Storage directly.

To find more information about Remote Storage command-line parameters, see Command-Line References in the Tools and Settings Collection.

**Rsdiag.exe: Remote Storage Diagnostic Tool**
**Category**
Part of Windows Server 2003. However, it is not installed by default. You must install Rsdiag.exe separately from the Support\Reskit\Stormgmt\Rstool folder on the Windows Server 2003 operating system disc.

**Version compatibility**
Runs on Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and the 64-bit versions of these operating systems. It is also available on Windows 2000 Server.

You can use Rsdiag.exe to examine remote storage databases and display diagnostic information. Rsdiag.exe can output the contents of remote storage databases and property files to text files. You can use Rsdiag.exe and Rsdir.exe together to diagnose problems in the databases and to recover lost data.

To find more information about Rsdiag.exe, see Resource Kit Tools Help in the Tools and Settings Collection.

**Rsdir.exe: File Information Tool**
**Category**
Part of Windows Server 2003. However, it is not installed by default. You must install Rsdir.exe separately from the Support\Reskit\Stormgmt\Rstool folder on the Windows Server 2003 operating system disc.

**Version Compatibility**
Runs on Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and the 64-bit versions of these operating systems. It is also available on Windows 2000 Server.

You can use Rsdir.exe to examine reparse points and display information for managed files.

To find more information about Rsdir.exe, see Resource Kit Tools Help in the Tools and Settings Collection.

**Remote Storage Registry Entries**
The following registry entry is associated with Remote Storage.

The information provided here is a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**TapeTypesToSupport**
**Registry path**
HKEY_LOCAL_MACHINE\CurrentControlSet\Services\Remote_Storage_Server\Parameters\

**Version**
Windows Server 2003

**Data Type**
REG_MULTI_SZ

Each string in this multi-string value corresponds to an additional tape device that you want to support (in addition to those that are supported by default: 4-mm, 8-mm and DLT). Each string contains a number that corresponds to the tape device serial number as it appears in STORAGE_MEDIA_TYPE enumeration. For more information about STORAGE_MEDIA_TYPE enumeration, see http://go.microsoft.com/fwlink/?linkid=99314.

The only currently supported value is "86", which represents the LTO Ultrium tape device. Add this value if you want to use an LTO Ultrium tape device with Remote Storage. You need to add the value prior to completing the Remote Storage Setup Wizard, which starts the first time that you run the Remote Storage MMC snap-in.

**Remote Storage Group Policy Settings**
There are no Group Policy settings associated with Remote Storage.

**Remote Storage WMI Classes**
There are no WMI classes associated with Remote Storage.

**Network Ports Used by Remote Storage**
Remote Storage does not utilize any network ports.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Resource Kit Tools Help in the Tools and Settings Collection.
- Command-Line References in the Tools and Settings Collection.

# Removable Storage Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Removable Storage Technical Reference**
The data storage and management features in the Microsoft Windows Server 2003 operating system provide you with various ways to manage and store data. With Removable Storage, a primary component of this feature set, you can track your removable storage media (tapes and optical disks) and manage the hardware libraries, such as changers and jukeboxes, which contain them.

With Removable Storage, you can:

- Label, catalog, and track media.
- Control library drives, slots, and doors.
- Perform drive-cleaning operations.

Removable Storage works together with your data-management applications such as Backup. You use data-management applications to manage the actual data stored on the media. Removable Storage makes it possible for multiple applications to share the same storage media resources, which can reduce your costs. It also provides a common interface for managing those resources, so that you can manage your storage media more efficiently.

**In this subject**


- What Is Removable Storage?
- How Removable Storage Works
- Removable Storage Tools and Settings

# What Is Removable Storage?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Removable Storage?**

**In this section**


- Common Scenarios for Removable Storage
- Removable Storage Dependencies on or Interactions with Other Technologies
- Removable Storage Logical Diagram
- Related Information


The data storage and management features in the Windows Server 2003 operating system provide you with various ways to manage and store data. With Removable Storage, a primary component of this feature set, you can track your removable storage media (tapes and optical disks) and manage the hardware libraries, such as changers and jukeboxes, which contain them.

With Removable Storage, you can:

- Label, catalog, and track media.
- Control library drives, slots, and doors.
- Perform drive-cleaning operations.

Removable Storage works together with your data-management applications such as Backup. You use data-management applications to manage the actual data stored on the media. Removable Storage makes it possible for multiple applications to share the same storage media resources, which can reduce your costs. It also provides a common interface for managing those resources, so that you can manage your storage media more efficiently.

Removable Storage organizes all the media in your libraries into different media pools. A media pool is a logical collection of removable media that have the same management policies. Media pools are used by applications to control access to specific tapes or discs within libraries that are managed by Removable Storage. Removable Storage also moves media between media pools in order to provide the amount of data storage that your applications require. To find more information about media pools, see "How Removable Storage Works" in "How Removable Storage Works."

You cannot use Removable Storage to manage volumes, such as for media siding or striping. Also, you cannot use Removable Storage to manage files, such as for data backup or disk-extender operations. These services are performed by data-management applications such as Backup or Remote Storage. Remote Storage is not available on computers running Windows XP Professional; Windows Server 2003, Web Edition; or Windows Server 2003, Standard Edition.

You must run all your data-management applications on the same computer that connects to your library. Removable Storage does not support multiple data-management applications running on different computers that are connected to the same library.

Removable Storage is configured to start automatically when you start your computer. It is possible to change the service so that you can start it manually, but this is highly discouraged. Disabling the service causes several applications that are included in Windows Server 2003 (such as Backup and Remote Storage) to become inoperative.

**Common Scenarios for Removable Storage**

Removable Storage is commonly used in the following scenarios:

- Managing stand-alone drive libraries
- Managing automated libraries

**Managing stand-alone drive libraries**
In this scenario, you use Removable Storage to manage multiple single-drive libraries, such as CD-ROM or DVD-ROM drives.

In its simplest form, a library consists of data-storage media and the device that is used to read from and write to the media – for example, tape and a stand-alone tape drive. The group of libraries and associated media that you manage with a Removable Storage installation is called a *Removable Storage system*.

The main benefit to using Removable Storage is its efficiency: it is much easier to manage multiple libraries with a single tool, the Removable Management MMC snap-in, on behalf of different data-management applications, than it is to individually manage the same libraries with different sets of tools from those applications.

A further benefit of Removable Storage is that it organizes all the media in your libraries into different media pools, and also moves media between media pools in order to provide the appropriate amount of data storage your applications require.

To find more information about media pools, see "How Removable Storage Works" in "How Removable Storage Works."

**Managing automated libraries**
In this scenario, you use Removable Storage to manage one or more automated libraries.

Automated libraries are automated units that hold multiple tapes or disks, and some have multiple drives. These libraries are sometimes called *changers* or *jukeboxes*, and commonly use robotic subsystems to move media stored in the library's storage slots.

You can also use Removable Storage to manage a combination of single-drive and automated libraries.

The benefits are the same as the ones described in the previous scenario.

**Administrative Interface**
The Removable Storage Microsoft Management Console (MMC) snap-in is an administrative interface that you can use to manage both stand-alone drive libraries and automated libraries.

Using the Removable Storage snap-in, you can:

- Create media pools and set media pool properties.
- Insert and eject media in an automated library.
- Mount and dismount media.
- Clean tape drives.
- View the state of media and libraries.
- Enable and disable drives and libraries.
- Perform library inventories
- Set security permissions for users.
- Complete or refuse operator requests
- Cancel work queue items.

To find more information about these tasks, see "How Removable Storage Works" in "How Removable Storage Works."

**Removable Storage Dependencies on or Interactions with Other Technologies**
Removable Storage depends on, or interacts with, the following technologies:

- Microsoft Management Console (MMC)
- Backup (or similar, non-Microsoft data-management programs)

- Win32 tape and disk management application programming interfaces (APIs)
- The registry
- Event Viewer
- Group Policy
- Media libraries

To find more information about this topic, see "How Removable Storage Works" in "How Removable Storage Works."

**Removable Storage Logical Diagram**
The first part of the following figure shows the inherent complexity in using multiple applications to manage multiple devices (each containing a different media-type) without the aid of Removable Storage.

The second part shows how you can reduce this complexity by using Removable Storage as the common interface for managing multiple devices.

### Removable Storage Logical Diagram




**Related Information**
The following resources contain additional information that is relevant to this section.

- Resource Kit Tools Help in the Tools and Settings Collection.
- Command-Line References in the Tools and Settings Collection.
- Group Policy Settings Reference in the Tools and Settings Collection.
- To find more information about Removable Storage interfaces, see MSDN and type "Disk Management Functions" in the "Search For" text box.

# How Removable Storage Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Removable Storage Works**

**In this section**


- Removable Storage Terms and Definitions
- Removable Storage Architecture
- Removable Storage Protocols
- Removable Storage Interfaces
- Removable Storage Processes and Interactions
- Network Ports Used by Removable Storage
- Related Information


The data storage and management features in the Windows Server 2003 operating system provide you with various ways to manage and store data. With Removable Storage, a primary component of this feature set, you can track your removable storage media (tapes and optical disks) and manage the hardware libraries, such as changers and jukeboxes, which contain them.

With Removable Storage, you can:

- Label, catalog, and track media.
- Control library drives, slots, and doors.
- Perform drive-cleaning operations.

Removable Storage works together with your data-management applications such as Backup. You use data-management applications to manage the actual data stored on the media. Removable Storage makes it possible for multiple applications to share the same storage media resources, which can reduce your costs. It also provides a common interface for managing those resources, so that you can manage your storage media more efficiently.

For these processes and interactions to work as described, Removable Storage should reside in an optimal environment. An optimal environment for Removable Storage is defined as follows:

- Windows Server 2003 is correctly installed and the Removable Storage service is running.
- All other dependencies are in place, properly designed and deployed, and functioning normally.
- Server hardware is sized appropriately and there are no disk, CPU, memory, or network bottlenecks that affect the performance of the technology.

**Removable Storage Terms and Definitions**

For a complete listing of Removable Storage Terms and Definitions, see MSDN and type **Removable Storage Glossary** in the **Search for** text box.

**Removable Storage Architecture**

From an architectural standpoint, Removable Storage is divided into three areas:

- **Removable Storage components:** The Removable Storage Microsoft Management Console (MMC) snap-in, the Removable Storage Application Programming Interface (API) and service, the Removable Storage database, and a data-management application (such as Backup).

- **Windows components and services:** Win32 tape and disk APIs, tape and disk drivers, changer drivers, Event Viewer, the registry, and Group Policy.

- **Libraries:** A robotic or automated library, a stand-alone tape drive, or a stand-alone DVD-RW drive, for example.

The following figure shows the components within these three areas and how the areas and components interact.

**Removable Storage Architectural Diagram**




The following table lists the components that are included in Removable Storage.

**Removable Storage Components**

| Component | Description |
|---|---|

| | |
|---|---|
| Removable Storage MMC snap-in | Provides an administrative interface to Removable Storage. Administrators can use the snap-in to insert and eject media, perform inventories, mount and dismount media, clean drives and check status information. Removable Storage objects can be physical objects (such as libraries, drives, and media) or groups of media (called media pools).<br><br>If only one removable storage device is being utilized, it is unlikely that you will ever be required to use the snap-in because each removable storage client application can take care of routine tasks by itself. For example, a simple backup application should not require a user to go to the snap-in to place media in a drive. In general, you will only need to use the snap-in if a system is running more than one removable-storage client application (for example, Backup and a document management system), or if the system has a sophisticated robotic library that requires administrative management.<br><br>The MMC snap-in can help administrators perform the following tasks:<br><br>● Complete or refuse operator requests.<br><br>● Point and click operations such as inserting media, ejecting media, drive cleaning, and dismounting media, moving media from one pool to another.<br><br>● Cancel work queue items.<br><br>● Enable and disable drives and libraries. |
| Removable Storage API | Provides a programmatic interface for applications to use to access the services and media in a Removable Storage system. It is a dynamic-link library (DLL) that contains a group of functions that you can use to build your data management application. These functions perform the following operations:<br><br>● Mount and dismount media.<br><br>● Clean drives.<br><br>● Insert and eject media.<br><br>● Inventory libraries.<br><br>● Enable and disable libraries, drives, and media.<br><br>● Access information in the Removable Storage database.<br><br>When you develop your media management application using the Removable Storage API, you gain the following benefits:<br><br>● Multiple applications can share the same library, drive, and media resources.<br><br>● Code written to the Removable Storage API for a specific library can execute on a wide range of libraries, both those that exist now and those that are introduced in the future.<br><br>● A single computer can track multiple types of media.<br><br>● A single computer can track media that are inside or outside a media library unit. |
| Removable Storage service | Manages and catalogs removable media and operates automated removable media services. If this service is stopped, programs that are dependent on Removable Storage, such as Backup and Remote Storage, will operate more slowly. If this service is disabled, any services that explicitly depend on it will fail to start.<br><br>The Removable Storage service depends on the Remote Procedure Call (RPC) service to function correctly. |

| | |
|---|---|
| Removable Storage database | Contains information about the objects in a Removable Storage system. Removable Storage maintains the data in the database and updates it whenever the administrator, an application, or a device makes a change in a Removable Storage system. For example, if the administrator wants to mount a piece of media in a library, Removable Storage mounts the media and updates the database to reflect the change. The database includes the following information:<br>• Library, drive, and media configuration and state.<br>• Media pool configuration and contents.<br>• Library requests.<br>• Library worklists.<br>• Operator requests. |
| Backup (storage management application) | Helps you protect data from accidental loss if your system's hardware or storage media fails. For example, you can use Backup to create a duplicate copy of the data on your hard disk and then archive the data on another storage device. The backup storage medium can be a logical drive such as your hard disk, a separate storage device such as a removable disk, or an entire library of disks or tapes organized and controlled by a robotic changer. If the original data on your hard disk is accidentally erased or overwritten, or becomes inaccessible because of a hard disk malfunction, you can easily restore the data from the archived copy. |

The following table lists the items that are included in Windows components and services.

**Windows Components and Services**

| Service | Description |
|---|---|
| Win32 tape and disk management APIs | Comprises four separate APIs: CreateFile, DeleteFile, GetDiskFreeSpace, and GetDiskFreeSpaceEx.<br>To find more information about these APIs, see "Removable Storage Interfaces" later in this section. |
| Tape and disk drivers | Enables a specific device, such as a stand-alone tape or disk drive, to communicate with the operating system. Although a device might be installed on your system, Windows cannot use the device until you have installed and configured the appropriate driver. If a device is listed in the Windows Catalog, a driver is usually included with Windows. Device drivers load automatically (for all enabled devices) when a computer is started, and then run in a manner that is transparent to users.<br>(The Windows Catalog website has hardware and software compatibility information for the Windows 2000 Server family, Windows XP, and the Windows Server 2003 family.) |
| Changer driver | Enables a specific device, such as an automated or robotic library, to communicate with the operating system. Although a device might be installed on your system, Windows cannot use the device until you have installed and configured the appropriate driver. If a device is listed in the Windows Catalog, a driver is usually included with Windows. Device drivers load automatically (for all enabled devices) when a computer is started, and thereafter run in a manner transparent to users. |
| Event Viewer | Provides a user interface for you to view and manage event logs, gather information about hardware and software problems, and monitor security events. Event Viewer maintains logs about program, security, and system events. |

| | |
|---|---|
| Registry | Contains information about the configuration of a computer. The registry database contains information that Windows continually references during operation, such as:<br>• Profiles for each user.<br>• Programs installed on the computer and the types of documents each can create.<br>• Property settings for folders and program icons.<br>• Hardware that exists on the system.<br>• Ports that are being used.<br><br>The registry is organized hierarchically as a tree and is made up of keys and their subkeys, hives, and entries. |
| Group Policy | Specifies settings for groups of users and computers. These groups include software policies, scripts, user documents and configurations, application deployment, and security configurations.<br><br>To find more information about how Removable Storage uses Group Policy, see Removable Storage Tools and Settings in "Removable Storage Tools and Settings." See also, Group Policy Settings Reference in the Tools and Settings Collection. |

The following table lists the components that are included in Libraries.

**Libraries**

| Library | Description |
|---|---|
| Robotic or automated library | Automated units that hold multiple tapes or disks. Some have multiple drives. These libraries are sometimes called *changers* or *jukeboxes*, and commonly use robotic subsystems to move media stored in the library's storage slots. Automated libraries can also consist of other hardware components that are managed by Removable Storage, such as doors, inject/eject ports, cleaner cartridges, and bar-code readers. |
| Stand-alone tape drive | A single-drive, nonautomated unit that holds a single tape. With this device, you manually insert a tape into the unit. |
| Stand-alone disk drive | A single-drive, nonautomated unit that holds a single disk. With this device, you manually insert a disk into the unit. |

**Removable Storage Protocols**
Removable Storage does not utilize any networking protocols.

**Removable Storage Interfaces**
The following APIs are associated with Removable Storage.

(To find more information about Removable Storage interfaces, see MSDN and type **Removable Storage Manager Functions** in the **Search for** text box.)

**Note**

- The Removable Storage functions were defined when the technology was previously named NTMS. As a result, many function names contain the prefix "Ntms". Similarly, the "scratch" pool was renamed "free", and the "foreign" pool was renamed "unrecognized". The term "partition" was replaced with the term "side".

Removable Storage can automatically clean tape drives that become dirty through use. You can use the following functions to control drive cleaning.

**Cleaner Management Functions**

| API Name | Description |
| --- | --- |
| CleanNtmsDrive | Queues a cleaning request for the specified drive for cleaning. |
| EjectNtmsCleaner | Ejects the cleaning cartridge from the currently reserved cleaner slot. |
| InjectNtmsCleaner | Allows a cleaner cartridge to be inserted into the specified library unit. |
| ReleaseNtmsCleanerSlot | Removes an existing slot reservation for a cleaning cartridge. The slot can then be used for data cartridges. |
| ReserveNtmsCleanerSlot | Reserves a single slot in a library unit for a drive cleaner cartridge. |

Removable Storage uses a database to store information about the removable media devices and cartridges in a system. Like any system resource, this database should be backed up in case it becomes lost or corrupted. You can use the following functions to backup and recover the Removable Storage database.

**Database Backup and Recovery Functions**

| API Name | Description |
| --- | --- |
| ExportNtmsDatabase | Creates a consistent set of database files in the Removable Storage database directory. |
| ImportNtmsDatabase | Causes Removable Storage to import the database files from the database Export directory at the next restart of the Removable Storage. |

You can use the following functions to get notifications when Removable Storage database objects change.

**Database Notification Functions**

| API Name | Description |
| --- | --- |
| CloseNtmsNotification | Closes the specified open notification channel. |
| OpenNtmsNotification | Opens a channel to receive Removable Storage object change notifications for objects of the specified type. |
| WaitForNtmsNotification | Waits for the next object change notification. |

Library control functions control injecting, ejecting, and moving media. They also enable and disable drive and changer resources. You can use the following library control functions.

**Library Control Functions**

| API Name | Description |
| --- | --- |
| AccessNtmsLibraryDoor | Unlocks the door of the specified library. If the library is busy, Removable Storage queues the request and returns successfully. |
| CancelNtmsLibraryRequest | Cancels outstanding Removable Storage requests, such as calls to the **CleanNtmsDrive** function. If the library is busy, Removable Storage queues the cancellation and returns success. |
| DeleteNtmsDrive | Deletes a drive from the Removable Storage database. The drive must have a **dwOperationalState** of NTMS_NOT_PRESENT. |
| DeleteNtmsLibrary | Deletes a library, and all the devices contained in the library, from the Removable Storage database. All media in the library is moved to the offline library. |
| DeleteNtmsRequests | Deletes a request or a list of requests from the Removable Storage database. Library or operator requests that are in a completed, failed, refused, or canceled state are removed. Submitted requests, queued requests, waiting requests, and in-progress requests cannot be deleted. |
| DismountNtmsDrive | Queues a command to move the media in the specified drive to its storage slot. This function should be paired with the **MountNtmsMedia** function. |

| | |
|---|---|
| EjectDiskFromSADrive | Ejects the media that is in a stand-alone removable drive. |
| EjectNtmsMedia | Ejects the specified medium from the port of the current library. If the library is busy, Removable Storage queues **EjectNtmsMedia** and returns a state of success. |
| GetNtmsRequestOrder | Gets the order that the specified request will be processed in the library queue. |
| GetNtmsUIOptions | Obtains the list of computer names to which the specified type of user interface is being directed for the given object. A call to **GetNtmsUIOptions** returns the list of destinations for the instance determined by the *lpObjectId* and *dwType* parameters.<br><br>If there are no destinations in the list for the specified instance, the function returns ERROR_SUCCESS along with a list length of zero. |
| GetVolumesFromDrive | Retrieves the volume and drive letter for a given Removable Storage media drive. |
| IdentifyNtmsSlot | Identifies the media in the specified slot in a library. The command returns a state of success when the identification is complete. |
| InjectNtmsMedia | Allows media to be inserted into the port of the specified library. If the library is busy, Removable Storage queues **InjectNtmsMedia** and returns a state of success. |
| InventoryNtmsLibrary | Queues an inventory of the specified library. If the library is busy, Removable Storage queues **InventoryNtmsLibrary** and returns success. |
| SetNtmsUIOptions | Modifies the list of computer names to which the specified type of UI is being directed for the given object. |
| SetNtmsRequestOrder | Sets the order that the specified request will be processed in the library queue. |

Your media label library must support the following functions. Removable Storage uses these functions to communicate with media label libraries.

**Media Label Library Functions**

| API Name | Description |
|---|---|
| ClaimMediaLabel | Determines whether a specified media label was created by the media's associated application. |
| MaxMediaLabel | Determines the maximum size of the media label for the applications that are supported by the media label library. |

Media services functions are used for mounting and dismounting media and managing media pools. You can use the following media services functions.

**Media Services Functions**

| API Name | Description |
|---|---|
| AddNtmsMediaType | Adds the specified media type to the specified library if there is not currently a relation in the library object. The function then creates the system media pools if they do not exist. |
| AllocateNtmsMedia | Allocates a piece of available media. |
| ChangeNtmsMediaType | Moves the specified Physical Media Identifier (PMID) to the specified target media pool and sets the PMID's media type identifier to the media type of the target media pool.<br><br>(A PMID is defined as a representation of media.) |
| CreateNtmsMedia | Creates a PMID and side (or sides) for a new piece of offline media. The media is placed in the media pool that is specified for the *lpPhysicalMedia* parameter.<br><br>A tape or disk side is where actual information is stored. |

| | |
|---|---|
| CreateNtmsMediaPool | Creates a new application media pool. |
| DeallocateNtmsMedia | Deallocates the side that is associated with the specified logical media. |
| DecommissionNtmsMedia | Moves a side from the Available state to the Decommissioned state. |
| DeleteNtmsMedia | Deletes a physical piece of offline media from Removable Storage by removing all references to the specified media from the database. |
| DeleteNtmsMediaPool | Deletes the specified application media pool. |
| DeleteNtmsMediaType | Deletes the specified media type relation from the specified library, provided that the library does not contain any physical media objects of the specified media type. |
| DismountNtmsMedia | Queues a command to move the specified media in a drive to its storage. This function should be paired with the **MountNtmsMedia** function. |
| GetNtmsMediaPoolName | Retrieves the specified media pool's full name hierarchy. |
| MountNtmsMedia | Synchronously mounts one or more pieces of media. |
| MoveToNtmsMediaPool | Moves the specified medium from its current media pool to the specified media pool. |
| SetNtmsMediaComplete | Marks a piece of logical media as complete. |
| SwapNtmsMedia | Swaps the sides associated with the two specified Logical Media Identifiers (LMIDs). The specified LMIDs must be in the same media pool.<br><br>(An LMID is defined as a logical representation of a physical storage location. The Media Service generally defines sides of the logical media for use by the application. The intent of logical media is to allow the underlying physical medium of the data to change without affecting the application.) |

You can use the following functions to enumerate, get, and set information about Removable Storage objects.

**Object Management Functions**

| API Name | Description |
|---|---|
| DisableNtmsObject | Disables the specified Removable Storage object. |
| EnableNtmsObject | Enables the specified object. |
| EnumerateNtmsObject | Enumerates the Removable Storage objects contained in the *lpContainerId* parameter. |
| GetNtmsObjectAttribute | Retrieves the extended attribute (named Private Data) from the specified Removable Storage object. |
| GetNtmsObjectInformation | Returns the information structure for the specified object. |
| GetNtmsObjectSecurity | Reads the security descriptor for the specified Removable Storage object. |
| SetNtmsObjectAttribute | Creates an extended attribute (named Private Data) in the specified Removable Storage object. |
| SetNtmsObjectInformation | Changes the information structure of the specified object. |
| SetNtmsObjectSecurity | Writes the security descriptor for the specified Removable Storage object. |

Removable Storage identifies media in a system by reading an on-media identifier (OMID). You can use the following function to update the RSM database when a new OMID is written.

**On-Media-Identifier Management Functions**

| API Name | Description |
|---|---|

| | |
|---|---|
| UpdateNtmsOmidInfo | Updates the Removable Storage database with label information immediately after writing to the newly allocated medium. |

You can use the following functions to manage and control requests to system operators.

**Operator Request Functions**

| API Name | Description |
|---|---|
| CancelNtmsOperatorRequest | Cancels the specified Removable Storage operator request. |
| SatisfyNtmsOperatorRequest | Completes the specified Removable Storage operator request. |
| SubmitNtmsOperatorRequest | Submits a Removable Storage operator request. |
| WaitForNtmsOperatorRequest | Waits for the specified Removable Storage operator request. |

Before calling any other functions, an application must open a Removable Storage session, and the session must be closed before the application exits. You can use the following session management functions to open and close Removable Storage sessions.

**Session Management Functions**

| API Name | Description |
|---|---|
| CloseNtmsSession | Closes the specified Removable Storage session. |
| OpenNtmsSession | Sets up a session with a Removable Storage server. |

**Removable Storage Processes and Interactions**

The following is a general description of how Removable Storage functions at a basic level.

As noted earlier, Removable Storage works together with your data-management applications, such as Backup. You use data-management applications to manage the actual data stored on the media. Removable Storage makes it possible for multiple programs to share the same storage media resources, which can reduce your costs.

The Removable Storage MMC snap-in is an administrative interface to Removable Storage and is used to perform a number of important tasks. With the Removable Storage MMC snap-in, you can:

- Create media pools and set media pool properties.
- Insert and eject media in an automated library.
- Mount and dismount media.
- Clean tape drives.
- View the state of media and libraries.
- Enable and disable drives and libraries.
- Perform library inventories
- Set security permissions for users.
- Complete or refuse operator requests
- Cancel work queue items.

Removable Storage organizes all the media in your libraries into different media pools. Removable Storage also moves media between media pools in order to provide the amount of data storage that your applications require.

You cannot use Removable Storage to manage volumes, such as for media siding or striping. Also, you cannot use Removable Storage to manage files, such as for data backup or disk-extender operations. These services are performed by data-management applications such as Backup or Remote Storage. Remote Storage is not available on computers running Windows XP Professional, Windows Server 2003, Web Edition; or Windows Server 2003, Standard Edition.

You must run all your data-management programs on the same computer that connects to your library. Removable Storage does not support multiple data-management programs running on different computers that are connected to the same library.

For these processes and interactions to work as described, Removable Storage should reside in an optimal environment. An optimal environment for Removable Storage is defined as follows:

- Windows Server 2003 is correctly installed and the Removable Storage service is running.
- All other dependencies are in place, properly designed and deployed, and functioning normally.
- Server hardware is sized appropriately and there are no disk, CPU, memory, or network bottlenecks that affect the performance of the technology.

**Removable Storage Technologies**

The following is a more detailed description of specific technologies, features or attributes that either form a part of, or are closely identified with, Removable Storage.

**Libraries**

In its simplest form, a library consists of data storage media and the device used to read from and write to the media. The group of libraries and associated media that are managed by a Removable Storage installation is called a *Removable Storage system*. There are two major types of libraries:

- **Automated libraries** are automated units that hold multiple tapes or disks, and some have multiple drives. These libraries are sometimes called *changers* or *jukeboxes*, and commonly use robotic subsystems to move media stored in the library's storage slots. Automated libraries can also consist of other hardware components that are managed by Removable Storage, such as doors, inject/eject ports, cleaner cartridges, and bar-code readers.
- **Stand-alone drive libraries**, or *stand-alone drives*, are single-drive, nonautomated units, such as tape or CD-ROM drives, that hold a single tape or disk. With these, you manually insert a tape or disk into the unit.

Not only can Removable Storage manage multiple libraries, but it can also track offline media that are not currently contained in a library. These media might be on a shelf or in a drawer, for example. Every tape or disk that is offline can belong to a library.

**Managing libraries**

Removable Storage provides comprehensive control and management of libraries. You can enable or disable entire libraries or only specific drives within a library. You can change the media type for a library. You can also open the door on an automated library and set various time-out values.

Removable Storage attempts to automatically configure all libraries, including automated libraries, when you first start Removable Storage and when you add or remove a library from a computer that is running Removable Storage.

Removable Storage can always automatically configure stand-alone drives, but you can use the automatic configuration feature of Removable Storage for automated libraries only if you connect all drives on the same SCSI bus as the associated media changer, or if all drives have unique serial numbers. Also, you must use automated libraries that support drive-element address reporting. For information on whether your library supports this feature, consult the library manufacturer.

### Note

- Because not all library configurations support the automatic configuration feature, Removable Storage provides a method for manually configuring libraries. However, do not use this method unless it is absolutely necessary because manually configured libraries do not recognize system configuration changes. Manually configured libraries also involve registry changes.
- Removable Storage numbers slots and drives starting from either 0 or 1. However, the device (firmware) has its own addresses that it uses. The Mcd.sys driver maps these addresses to the more user-friendly numbers (starting from 0 or 1) used by Removable Storage.

**Inventorying libraries**
You can create an inventory of all media in a library. Removable Storage supports two inventory methods:

- **Fast inventory.** Removable Storage checks the storage slots and reads the on-media identifier on media in slots that were previously empty.

- **Full inventory.** Removable Storage mounts each tape or disk in your library and reads the on-media identifier.

**Media Pools**
A media pool is a collection of tapes or disks that have the same management properties. All media in a Removable Storage system belong to a media pool, and each media pool holds only one type of media. Data management programs use media pools to gain access to specific tapes or disks within a library.

Using media pools, you can define properties that apply to a set of media. This is useful because Removable Storage allows multiple programs to share the same media within a single library. A library can include media from different media pools, each with different properties. A single media pool can span multiple libraries. You can also create hierarchies of media pools, or media pools that contain other media pools. For example, you can create a media pool for each specific media type required by a program, and then create another media pool that contains this collection of media pools. Media pools can contain either media or other media pools, but not both.

A Removable Storage system provides two classes of media pools: system and application.

System media pools include free media pools, unrecognized media pools, and import media pools. Removable Storage creates one free, one unrecognized, and one import media pool for each media type in your Removable Storage system. The system media pools are used to hold media that are not currently being used by an application.

Application media pools are created by data-management applications such as Backup and Remote Storage, the latter for computers running Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; the 64-bit version of Windows Server 2003, Enterprise Edition; or the 64-bit version of Windows Server 2003, Datacenter Edition.

The different media pools are defined as follows:

- **Unrecognized media pools** contain blank (new) media and media that Removable Storage does not recognize. You should immediately move a new tape or disk from an unrecognized media pool to a free media pool so that the tape or disk can be used by applications, or remove it from the library.

Unrecognized media are automatically deleted from the Removable Storage database when they are ejected from a library.

- **Import media pools** contain media that Removable Storage recognizes in the database, but that have not been used before in a particular Removable Storage system. For example, media in an import media pool could be media from one office location that are introduced into a Removable Storage system at another office location.

You can move media from import media pools to free media pools or application media pools so applications can use them.

- **Free media pools** contain media that are not currently in use by applications and do not contain useful data. Media in free media pools are available for use by applications.

You can configure application media pools to automatically draw media from free media pools when there are not sufficient media available in a particular application media pool. If you do not implement this configuration, you must manually move media from a free media pool when needed.

- **Application media pools** determine which media can be accessed by which applications. You or data management applications create them.

Media in an application media pool are controlled by that application or by an administrator. An application can use more than one media pool, and more than one application can share a single media pool. For example, Backup might use one media pool for a full backup and another media pool for an incremental backup, each containing a different media type.

There can be any number of application media pools in a Removable Storage system. Media that are currently reserved for use

by an application, called *allocated media*, cannot be moved between media pools. Allocation controls how the media are used by applications.

The following figure shows how media can move from one media pool to another.

**How Media Can Move From One Media Pool To Another**




**Managing media pools**

After configuring your libraries, you can configure your media pools. Before creating a new media pool, you should answer the following questions:

- What media types are available in your Removable Storage system?
- Should media be automatically drawn from a free media pool when there are not sufficient media in an application media pool? (It is recommended that you implement this configuration.)
- Should media be automatically returned to a free media pool when they are no longer needed by an application?

**Note**

- You cannot delete system media pools.

**Managing tapes and disks in a media pool**

After you create and configure your media pools, you can move media to a specific media pool. It is important that you have a sufficient amount of media in a free media pool so that applications can use media in this pool when needed.

Next, you can use a data management application such as Backup to allocate media to the application media pool created for (or by) that application. Allocating media to an application media pool reserves all media in that media pool for exclusive use by that application. See the documentation for your particular data management application for information on how to reserve media for that application.

Removable Storage provides comprehensive control of tapes or disks in your library. You can enable or disable specific media in your library, insert and eject media, and mount and dismount media.

A tape or disk can be inserted in or ejected from an automated library using either a library door or an inject/eject port. A library door allows unrestricted access to the media inside. You can add or remove media directly from a storage slot using a door access. An inject/eject port allows controlled access to the media inside. You insert a tape or disk in the port and then the library uses the drive bay to move the media to a storage slot.

A tape or disk can be mounted on or dismounted from a particular drive in a library. You can use the same media many times before dismounting them. You can also mount and dismount the same media many times before deallocating them.

**Note**

- Removable Storage requires that CDs are formatted with the CD-ROM file system and that rewritable optical media (such as magneto-optical disks) are formatted with the FAT, FAT32, or NTFS file system.
- Windows does not directly support the use of media that can be written to only once, such as write once, read-many (WORM), CD-R, and DVD-R disks. If you plan to use these, you must acquire other software that supports these

technologies.

**Media Identification**
Removable Storage uses two identification methods to track media and perform inventories: on-media identifiers and bar codes.

**On-media identifiers**
On-media identifiers are electronically recorded on media the first time a tape or disk is inserted into a library. The next time the medium is inserted into the library, Removable Storage uses the on-media identifier to recognize and track the tape or disk.

An on-media identifier has two parts: a label type and a label ID. The label type identifies the specific format used to record information on the tape or disk, such as Microsoft Tape Format (MTF). The label ID is a unique identifier for the tape or disk in the Removable Storage system. If Removable Storage does not recognize the label type, it adds the medium to an unrecognized media pool. If Removable Storage recognizes the label type, but does not recognize the label ID, it puts the media in an import media pool. If Removable Storage recognizes the label type and the label ID, it updates the Removable Storage database to reflect that the medium is online and available for use.

### Note

- Removable Storage does not record on-media identifiers for read-only or write-once optical media, such as CD-ROM, DVD-ROM, CD-R, DVD-R, or WORM disks. Instead, Removable Storage uses the volume and serial number information that is already associated with the disk.

**Bar codes**
If your library supports bar codes, Removable Storage can use them to identify media. Media with bar codes also have on-media identifiers. Removable Storage can use either the bar codes or the on-media identifiers. Using bar codes to track media is generally much faster because each tape or disk does not have to be mounted to read the on-media identifier.

**Media States**
Each tape or disk in a Removable Storage system has a state associated with it that determines what operations can be performed with the media. When media are inserted into a library, Removable Storage uses the on-media identifier or bar code to determine its state.

There are two sets of media states that govern media movement and usage by Removable Storage: physical states and side states.

**Physical states**
The physical states identify operations performed that involve media location or movement. The following table lists and describes the five possible physical states.

### Physical States

| Physical State | Description |
|---|---|
| Idle | Tape or disk is currently in a storage slot in an automated library or is shelved offline. |
| In-use | Tape or disk is in the process of being moved. |
| Loaded | Tape or disk is mounted in a drive and data is available to be read or written. |
| Mounted | Tape or disk is mounted in a drive, but data is not yet available to be read or written. |
| Unloaded | Tape or disk has been dismounted and is ready to be removed from a drive. |

**Side states**
A tape or disk side is where actual information is stored. Side states identify operations that are performed that involve media usage rather than location or movement. The following table lists and describes the nine possible side states.

### Side States

| Side State | Description |
|---|---|

| | |
|---|---|
| Allocated | Side is reserved for use by an application. The side is not available to any other application. |
| Available | Side is currently available for use by an application. |
| Completed | Side is in use, but cannot be used for write operations. The side is full. |
| Decommissioned | Side is no longer available for use because it has reached its allocation maximum. |
| Imported | Side's label type is recognized, but its label ID is not recognized by Removable Storage. |
| Incompatible | The media type of the side is not compatible with the library. The tape or disk should be removed from the library. |
| Reserved | Side is unavailable for allocation except by a single application. This applies only to two-sided media where one side has already been allocated. |
| Unprepared | Side has been placed in a free media pool, but does not yet have a free media label. This is a temporary state. |
| Unrecognized | Side's label type and label ID are not recognized by Removable Storage. |

The relationship between the various media pools and side states is as follows: Import media pools can only hold media sides in the Imported state. Unrecognized media pools can only hold media sides that are in the Unrecognized state. Free media pools can only hold media sides that are in the Available (and transitional unprepared) state.

**Managing Operator Requests and Queued Work**

The work queue lists all library requests, or work items that are initiated by an application or by Removable Storage. A request to mount a tape in a library, for example, results in a mount work item listed in the work queue.

The following table describes the states that a work item can have.

**Work Item States**

| Work State | Description |
|---|---|
| Completed | The work item has been completed successfully. |
| Failed | Removable Storage was unable to complete the work item. |
| In Process | Removable Storage is actively operating on the work item. |
| Queued | The work item is ready to be carried out, but Removable Storage has not yet begun to operate on it. |
| Waiting | One or more resources needed to satisfy the request are currently in use. |
| Cancelled | A user has cancelled the work item. |

An operator request is a message that asks you to perform a specific task. Operator requests can be issued by Removable Storage or by a program that is aware of Removable Storage, such as Backup. Removable Storage can generate an operator request when any of the following occur:

- An application initiates a mount request for a tape or disk that is offline.
- There are no available media online. That is, a program requires media and none exist in the appropriate application media pool or in a free media pool.
- A library has failed and requires servicing.
- A drive needs cleaning and no usable cleaner cartridges are available in the library.

You can respond to operator requests by either completing or refusing the operator request. When you refuse an operator request, Removable Storage notifies the application that generated the request. By default, Removable Storage stores an operator request in the work queue for 72 hours after you respond to the request.

The following table describes the states that an operator request can have.

**Operator Request States**

| Operator request state | Description |
|---|---|
| Completed | You have indicated, or Removable Storage has detected, that the operator request was completed. |
| Refused | You have indicated that the operator request will not be performed. |
| Submitted | Removable Storage or an application has submitted an operator request for action. |

You can change how operator requests are displayed on your computer, delete operator requests, cancel any pending operation listed in the work queue, and change the order of mount requests in the work queue.

**Network Ports Used by Removable Storage**

Removable Storage does not utilize any network ports.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Resource Kit Tools Help" in the Tools and Settings Collection.
- "Command-Line References" in the Tools and Settings Collection.
- "Group Policy Settings Reference" in in the Tools and Settings Collection.
- To find more information about Removable Storage interfaces, see MSDN and type "Disk Management Functions" in the "Search For" text box..

# Removable Storage Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Removable Storage Tools and Settings**

**In this section**


- Removable Storage Tools
- Removable Storage Registry Entries
- Removable Storage Group Policy Settings
- Removable Storage WMI Classes
- Network Ports Used by Removable Storage
- Related Information


The data storage and management features in the Windows Server 2003 operating system provide you with various ways to manage and store data. With Removable Storage, a primary component of this feature set, you can track your removable storage media (tapes and optical disks) and manage the hardware libraries, such as changers and jukeboxes, which contain them.

With Removable Storage, you can:

- Label, catalog, and track media.
- Control library drives, slots, and doors.
- Perform drive-cleaning operations.

Removable Storage works together with your data-management applications such as Backup. You use data-management applications to manage the actual data stored on the media. Removable Storage makes it possible for multiple applications to share the same storage media resources, which can reduce your costs. It also provides a common interface for managing those resources, so that you can manage your storage media more efficiently.

Removable Storage organizes all the media in your libraries into different media pools. A media pool is a logical collection of removable media that have the same management policies. Media pools are used by applications to control access to specific tapes or discs within libraries that are managed by Removable Storage. Removable Storage also moves media between media pools in order to provide the amount of data storage that your applications require. To find more information about media pools, see "How Removable Storage Works" in "How Removable Storage Works."

You cannot use Removable Storage to manage volumes, such as for media siding or striping. Also, you cannot use Removable Storage to manage files, such as for data backup or disk-extender operations. These services are performed by data-management applications such as Backup or Remote Storage. Remote Storage is not available on computers running Windows XP Professional; Windows Server 2003, Web Edition; or Windows Server 2003, Standard Edition.

You must run all your data-management applications on the same computer that connects to your library. Removable Storage does not support multiple data-management applications running on different computers that are connected to the same library.

Removable Storage is configured to start automatically when you start your computer. It is possible to change the service so that you can start it manually, but this is highly discouraged. Disabling the service causes several applications that are included in Windows Server 2003 (such as Backup and Remote Storage) to become inoperative.

**Removable Storage Tools**

The following tool is associated with Removable Storage.

**Removable Storage Microsoft Management Console (MMC) Snap-in**
**Category**
Part of Windows Server 2003.

**Version compatibility**
Also part of Windows XP and can be used to perform administrative tasks on locally attached storage devices.

The Removable Storage MMC snap-in is an administrative interface to Removable Storage. Administrators can use the snap-in to insert and eject media, perform inventories, mount and dismount media, clean drives and check status information. Removable Storage objects can be physical objects (such as libraries, drives, and media) or logical groups of media (called media pools).

### Note

- If only one removable storage device is being utilized, it is unlikely that you will ever be required to use the snap-in because each removable storage client application can take care of routine tasks by itself. For example, a simple backup application should not require a user to go to the snap-in to place media in a drive. In general, you will only need to use the snap-in if a system is running more than one removable storage client application (for example, Backup and a document management system), or if the system has a sophisticated robotic library that requires administrative management.

The MMC snap-in can help administrators perform the following tasks:

- Create media pools and set media pool properties.
- Insert and eject media in an automated library.
- Mount and dismount media.
- Clean tape drives.
- View the state of media and libraries.
- Enable and disable drives and libraries.
- Perform library inventories.
- Set security permissions for users.
- Complete or refuse operator requests.
- Cancel work queue items.

To find more information about Microsoft Management Console, see “Resource Kit Tools Help” in the Tools and Settings Collection.

**Removable Storage Registry Entries**
The following registry entries are associated with Removable Storage.

The information provided here is a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

To find more information about the Removable Storage registry entries, see “Registry Reference” in the Tools and Settings Collection.

**AutoCfg**
**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Config\Changer%d\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_DWORD

This entry specifies whether the changer has been manually configured by an administrator.

**Changer**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Config\

**Version**
Windows 2000, Windows XP, Windows Server 2003

The Config key contains a subkey for each changer (for example, Changer0) and a subkey for each stand-alone drive. Each changer subkey contains an entry for each drive bay in the changer (for example, DriveBay0). For each drive bay entry that has the value "???", replace that value with the device name of the drive in that bay, such as "Tape3". Do not use an initial "\" or "." (period) character"."

**CurrentConfigPassed**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Config\Changer%d\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_DWORD

This entry is set when configuration is complete.

**Description**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\NTMS\OMID\Tape*\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_SZ

This entry contains the description for the media label library (MLL). A media label library is defined as a dynamic-link library (DLL) that can interpret the format of a media label written by a Removable Storage application.

**DriveBay%d**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Config\Changer%d\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_SZ

This entry contains information about a drive in a specific drive bay (for example, CdRom0, Tape1, and so on).

**EventMessageFile**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\EventLog\System\Removable Storage Service\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_EXPAND_SZ

This entry stores the location of the file that contains event identifiers for the events generated by the source program, and text

describing each event. Typically, event identifiers are included in the dynamic-link library (DLL) or executable file for the source program.

**NumDrives**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Config\Changer%d\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_DWORD

This entry shows the number of drives in the changer.

**Path**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\NTMS\OMID\Tape*\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_EXPAND_SZ

This entry contains the path and DLL for the media label library (MLL).

**ServiceDll**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Parameters\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_EXPAND_SZ

This entry informs Service Control Manager of the service DLL (along with its path).

**ShutdownTime**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Parameters\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_DWORD

This entry determines the amount of inactive time that can pass before the Removable Storage service shuts down.

**TypesSupported**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\EventLog\System\Removable Storage Service\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_DWORD

This entry indicates newly introduced media types. These are in addition to the standard media types that the operating system recognizes by default.

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NtmsSvc\Config**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Config\.

**DriveList**
**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Config\Changer%d\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_SZ

This entry lists the drives in the changer.

**DriveList**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtmsSvc\Config\Standalone\

**Version**
Windows 2000, Windows XP, Windows Server 2003

**Data Type**
REG_SZ

This entry lists the stand-alone drives that are connected to the computer.

**Removable Storage Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with Removable Storage. These settings permit or prohibit the use of the Removable Storage MMC snap-in.

**Group Policy Settings Associated with Removable Storage Management**

| Group Policy Setting | Description |
|---|---|
| Not Configured | If this policy is not configured, the setting of the **Restrict users to the explicitly permitted list of snap-ins** policy determines whether this snap-in is permitted or prohibited. |
| Enabled | The snap-in is permitted. |
| Disabled | The snap-in is prohibited. |

**Note**

- The path to Removable Storage Management is User Configuration\Administrative Templates\Windows Components\Microsoft Management Console\Restricted/Permitted snap-ins. (Restricted/Permitted snap-ins is the full name for the folder.)

To find more information about Removable Storage Group Policy settings, see "Group Policy Settings Reference" in the Tools and Settings Collection.

**Removable Storage WMI Classes**
There are no WMI Classes associated with Removable Storage.

**Network Ports Used by Removable Storage**
Removable Storage does not utilize any network ports.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Resource Kit Tools Help in the Tools and Settings Collection.
- Registry Reference in the Tools and Settings Collection.
- "Command-Line References" in the Tools and Settings Collection.
- Group Policy Settings Reference in Tools and Settings Collection.
- To find more information about Removable Storage interfaces, see MSDN and type "Disk Management Functions" in the

“Search For” text box.

# Windows Security Collection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Security Collection**

As organizations expand the availability of network data, applications, and systems, it becomes more challenging to ensure the security of the network infrastructure. Security technologies in the Microsoft Windows Server 2003 operating system enable organizations to better protect their network resources and organizational assets in increasingly complex environments and business scenarios.

**Fundamental Security Principles**

Windows Server 2003 security technologies address fundamental security requirements that help meet the complex security needs of organizations of all types and sizes. Windows Server 2003 security is based on the following fundamental principles:

- **Identification.** To help ensure that only the appropriate users and computers have access to resources, it is first necessary to identify users and computers on the network. This involves setting up and maintaining account information for users and computers, preferably in a single, easy-to-access location so that it is easy to set up, modify, and maintain. The user name generally is a unique identifier.

- **Authentication.** The authentication process validates the authentication data of a user or computer against the information in a database. This authentication data can include the user name, logon domain, password, and other credentials. After a user or computer has been authenticated, the operating system examines the privileges that are assigned to the user account. The information relating to the user in the account database is used to create an access token, which is then used to determine the access rights of a user or computer during the current session.

- **Authorization and access control.** Access rights to a given resource are validated based on access control lists (ACLs) associated with the resource. The contents of the access token are compared to the contents of the ACL in order to determine the rights of the user in regard to the resource.

- **Confidentiality.** Confidentiality helps prevent the intentional or unintentional disclosure of data or of the actions that a user is performing on the data — for example, a withdrawal from a bank account. Confidentiality is typically accomplished by means of encryption as data crosses exposed portions of a network.

- **Integrity**. Integrity services help to ensure that the content of a message or data file has not been modified when it travels over a network.

- **Nonrepudiation.** Nonrepudiation, an extension of authentication and integrity, prevents a user from denying, after the fact, that they sent a message or signed a document. It can also be used to prove that the message was sent, that it was delivered, and that it was received.

- **Trusts.** Logical relationships are established between domains, by means of trusts, to allow pass-through authentication, in which one domain accepts the logon authentications of the other domain. A trust either allows or disallows authentication traffic to flow between two or more domains.

- **Audit entries.** Audit entries represent data that is recorded in the security event log of a server or workstation when specified system, application, and security-related events take place. Audit entries provide valuable data about system operations, which can be used to identify system use and misuse, and to diagnose system behavior.

**Security Architecture**

The Windows Server 2003 security infrastructure consists of the following components:

- **Logon and authentication technologies.** Logon and authentication technologies include a variety of protocols, including Kerberos version 5 authentication, NTLM, Secure Sockets Layer/Transport Layer Security (SSL/TLS), and Digest; as well as features such as Stored User Names and Passwords that enable single sign-on (SSO) and reduced sign-on (RSO).

- **Authorization and access control technologies.** The ACL-based impersonation model and a new roles-based protected subsystem model enable extremely flexible and manageable authorization and access control strategies.

- **Data security technologies.** Encrypting File System (EFS), Internet Protocol security (IPSec), system key utility (Syskey), and Routing and Remote Access Services (RRAS) provide additional security for data under a variety of special circumstances.

- **Group Policy technologies.** Group Policy options that can enhance security management include security policy and software restriction policies.

- **Trust technologies.** Trusts can be established between domains and across forests to improve security and business processes for complex organizations.

- **Public key infrastructure (PKI) technologies.** Certificates, Certificate Services, and certificate policy-enabled qualified subordination can be used to support a variety of application-specific security solutions.

Each of these sets of technologies can be used in conjunction with the other sets of technologies — such as networking and storage — to enable secure network-enabled business processes.

**Logon and Authentication Technologies**

Applications integrated with Windows Server 2003 and Active Directory directory service use the built-in features of the operating system to perform authentication. Authentication mechanisms are fundamental to system security.

Authentication mechanisms verify the identity of users who attempt to log on to a domain or to access network resources. Authentication, the act of proving the identity of a user to an application or resource, is typically done by cryptographically signing data with a key only the user knows.

The server side of the authentication exchange compares the signed data with a known cryptographic key to validate the authentication attempt. Because the cryptographic keys are stored in a secure central location, the authentication process is scalable and maintainable. In the Windows operating system, the recommended locations for storing identity information, including the cryptographic keys that are the credentials of the users, is the Local Security Authority (LSA) of a client computer or Active Directory on a network.

Windows Server 2003 implements a number of authentication protocols, including Kerberos v 5 authentication, NTLM, secure channel, Passport, and Digest. These protocols enable authentication of users, computers, and services in a variety of network scenarios, including private networks based on Active Directory, networks based on Microsoft Windows NT 4.0 domains, and clients and servers based on non-Windows operating systems, and also across the public Internet.

Windows makes it possible to use a single user identity across a local network by locally caching user credentials in the Local Security Authority (LSA) of a client computer. When a user logs on to a domain, Windows authentication packages transparently use the credentials to provide single sign-on when authenticating to network resources on servers that trust the user's account domain.

However, as more companies deploy Internet applications that are used by partners and customers, the authentication mechanisms they currently use to provide single sign-on to their intranets are not always adequate. In fact, the technologies required for Internet-facing Web application single sign-on are significantly different. For this reason, the authentication process they provide is referred to by a different term, Web single sign-on.

Sometimes it is not possible to achieve single sign-on. In these cases, reduced sign-on (RSO) is preferred. This reduces the number of times a user must enter their credentials to access resources and applications.

Stored User Names and Passwords improves the SSO and RSO experience by simplifying the management and use of multiple sets of logon credentials, including X.509 certificates used with smart cards and Passport credentials. The credentials — part of the user's profile — are stored until needed as an encrypted part of a user's local profile and can even roam with the user if the user's network policy supports Roaming Profiles.

Authentication in Windows Server 2003 also includes two-factor authentication, such as smart cards, which are a tamper-resistant and portable way to provide security solutions for tasks such as client authentication, logging on to a Windows Server 2003 domain, code signing, and helping to secure e-mail.

**Authorization and Access Control Technologies**

How you choose to give and manage appropriate access to individual users and groups of users is the key to your resource authorization and access control strategy. Users must be authorized to access specific resources, but only at the level of access

they need.

Authorization and access control in Windows are based on user rights and permissions. User rights grant specific privileges and logon rights to users and groups in your computing environment. Permissions define the type of access granted to a user or group for an object or object property.

Windows Server 2003 manages user rights and permissions based on the following two primary authorization models:

- The ACL-based impersonation model
- A new roles-based protected subsystem model

Logging on does not automatically give users access to the resources that they require. With the impersonation model, an authentication package authenticates the user and then builds an access token that represents the security context for that user. The security context represents the identity and groups that a user is a member of, as well as the privileges or rights that are granted to that user to perform specific tasks that affect an entire computer system.

Many users and computers have identical needs for access to a network resource. For this reason, the ACL-based impersonation model allows you to place users and computers in security groups in order to manage the rights and permissions of many users and computers at the same time.

If your organization has Active Directory, it can serve as the focal point for managing and administering user accounts, security group memberships, security policies, and organizational resources such as computers, printers, and servers. Organizations that use Active Directory as the authoritative store for identity, authentication, and authorization information can extend this access control framework to other applications, systems, and, by means of Lightweight Directory Access Protocol (LDAP) compliance, to other platforms.

After the authentication package generates the security context, an application or service can be given the security context to use for accessing resources. In other words, the service can use the user's security context to obtain access to resources instead of using the service's own context.

Another component of the impersonation model is the object-based ACL. The ACL for an object indicates what level of access any number of security principals have to the object. For example, the ACL on a file might define that one user has the ability to read the file while another user has the ability to both read and delete the file.

The operating system, specifically the Windows Object Manager, is the final arbitrator of access to objects that are using the impersonation model. The Object Manager does this by comparing the security context of the user against the ACL on the object. Because the operating system is the final arbitrator, the ACL-based impersonation model has security characteristics that are good for scenarios in which multiple applications can access a system resource.

Windows Server 2003, Enterprise Edition introduces a new roles-based protected subsystem authorization model, a feature called Authorization Manager. In the roles-based protected subsystem model, users are assigned to roles and the security settings that authorize users for specific sets of tasks are based on authorization rules that allow you to apply fine-grained control over the mapping between access control and the structure and tasks performed in your organization. For example, under the roles-based protected subsystem model, a bank teller is allowed to enter account deposits and withdrawals for a customer. If this employee suddenly becomes a loan officer, the application rules can dynamically allow him or her to view a customer's bank balances and record new loans, but prevent him or her from making deposits or withdrawals on behalf of a customer.

Windows Server 2003 provides complete authorization auditing support. For the ACL-based impersonation model, the Object Manager reports audits for resource access according to the audit configuration. This configuration can be enforced by means of Group Policy or manually on each server. The roles-based protected subsystem model supports both run-time auditing and policy change auditing. Run-time auditing uses either fail or success audits to report application initialization, context creation and deletion, and object access.

**Data Security Technologies**

Applications consume authentication and authorization data to allow or deny access to a user. However, the business requirements of an organization — for example, the need to support mobile users who travel with sensitive corporate data or the need to communicate sensitive data securely over the public Internet — make additional data protection strategies necessary.

IT professionals can use the data security technologies in Windows Server 2003 to complement the data security features that are coded into an application. Windows Server 2003 includes the following disk-based encryption features:

- **Encrypting File System (EFS).** Stored data (stored online or offline) can be protected by means of EFS and digital signatures. By using EFS, you can encrypt data as it is stored on disk. EFS uses public key encryption to encrypt local NTFS file system data.

- **The system key utility (Syskey).** Syskey provides additional protection against password-cracking software. It uses strong encryption techniques to secure account password information that is stored in the Security Accounts Manager (SAM) database or in directory services.

Additional Windows Server 2003 features can protect network data as it passes in and out of your site (across intranets, extranets, or an Internet gateway). These features include:

- **Internet Protocol Security (IPSec).** IPSec is a suite of industry-standard cryptography-based protection services and security protocols that are used to protect local area network (LAN) traffic. IPSec provides computer-level authentication, as well as data encryption, for virtual private network (VPN) connections that use the Layer 2 Tunneling Protocol (L2TP).

- **Routing and Remote Access.** The Routing and Remote Access service for the Windows Server family is a full-featured software router and an open platform for routing and internetworking. It offers routing services to businesses in LAN and wide area network (WAN) environments, or over the Internet, by using secure VPN connections.

- **Internet Authentication Service (IAS).** Internet Authentication Service (IAS) in Windows Server 2003 provides security and authentication for dial-in users and for local networks by means of 802.1x.

**Group Policy Technologies**

Windows Server 2003 Group Policy technologies enable you to efficiently manage groups of users and computers by configuring Group Policy objects (GPOs) and linking them to an organizational unit, domain, or site.

Group Policy allows an organization to more effectively manage its applications and operating environments on a per-computer or per-user basis.

Group Policy in Windows Server 2003 includes extensions for the following:

- Administrative templates

- Software installation

- Security settings

- Software restrictions

- IPSec policy

- Scripts

- Wireless network policies

- Folder redirection

- Internet Explorer maintenance

- Disk quotas

- Quality of service

- Remote installation services

You can use security settings policies as part of your overall security implementation to help secure domain controllers, servers, clients, and other resources in your organization. Security settings policies are used to manage the following aspects of

security:

- User authentication to a network or computer
- Resources that users are permitted to access
- Whether to record a user's or group's actions in the Event log
- Membership in a group

This assists administrators in the following tasks:

- Ensuring that corporate security policies are enforced across desktops and servers.
- Keeping desktops and servers up-to-date with the latest security patches.
- Monitoring systems for potential security compromises.

Software restriction policies further enhance an organization's ability to enforce security policy on its computers and network by limiting the opportunities for malicious code to enter an operating system and impact an organization.

Users can receive hostile code in many forms, ranging from native Windows executable files (.exe), to macros in documents (such as .doc files), to scripts (such as .vbs files). Attackers often use social engineering methods, such as viruses and worms, to cause users to activate their code. If such code is activated, it can generate denial-of-service attacks on the network, send sensitive or private data to the Internet, put the security of the computer at risk, or damage the contents of the hard disk drive.

To help protect network computers from both hostile code and unknown or unsupported software, administrators can use software restriction policies to define which applications are allowed or not allowed to run on a target computer. To help organizations address the problem of unknown code, administrators can use software restriction policies to perform the following tasks:

- Counteract computer viruses.
- Regulate which ActiveX controls can be downloaded.
- Run only digitally signed scripts.
- Enforce the policy that only approved software is installed on system computers.
- Lock down a computer.

For more information about Group Policy, see the Group Policy Collection.

**Trust Technologies**

The goals of SSO and RSO are largely based on trusts establish between domains and forests. However, SSO and RSO can be difficult to establish and maintain when organizations are continuously evolving with mergers, acquisitions, and divestitures, as well as a wide variety of partnerships with business partners.

The trust technologies in Windows Server 2003 help organizations address these business requirements and enhance their ability to offer and maintain SSO and RSO. When a trust exists between two domains, the authentication mechanisms for each domain trust the authentication data coming from the other domain. Trusts help provide for controlled access to shared resources in a resource domain (the trusting domain) by verifying that incoming authentication requests come from a trusted authority (the trusted domain).

Applications integrated with Windows Server 2003 and Active Directory use the built-in features of the operating system to establish and maintain trust for a wide variety of business requirements and scenarios, including forest trusts and cross forest trusts.

Windows Server 2003 fully audits trust configuration at a detailed level. Auditable events include the creation, deletion, and modification of trusts.

**PKI Technologies**

A different form of trust can be implemented by using a public key infrastructure (PKI). A PKI is a system of digital certificates, certification authorities (CAs) and registration authorities (RAs) that verify and authenticate the validity of each party involved in an electronic transaction by means of the use of public key cryptography. For this to occur, however, each party must trust the issuer of the certificate, most often a CA.

For Windows Server 2003 users, computers, and services, trust in a CA is established when a copy of the root certificate, as well as a valid certification path, exists in the trusted root certification authorities store. A valid certification path means that none of the certificates in the certification path has been revoked or has expired. The process of importing a root certificate into the trusted root certification authorities store is the most straightforward way to establish a PKI trust for a computer and the applications hosted by that computer.

While the root store certification authority mechanism is fairly easy to deploy, it might not meet the security requirements of complex scenarios for establishing a federated trust between two organizations. One solution is to use qualified subordination.

Qualified subordination allows you to place certificate issuance constraints on subordinate CAs and to place usage constraints on the certificates they issue. By using qualified subordination, you can focus subordinate CAs according to specific certification needs and administer your PKI more efficiently. Qualified subordination also allows you to establish trust between CAs in separate trust hierarchies. This type of trust relationship is also called cross-certification. With this trust relationship, qualified subordination is not limited to subordinate CAs. Trust between hierarchies can be established using a subordinate CA in one hierarchy and the root CA in another hierarchy.

**Security Scenarios**

As organizations grow and evolve, the need for security technologies that help to protect system resources becomes more critical. For example:

- A growing number of users have laptop computers which they are just as likely to use away from the office as at a network-connected desktop.
- Many organizations are making business processes more efficient by eliminating paper documents and replacing them with electronic documents that may be more difficult to track or control.
- Many organizations are increasing their use of contractors and consultants, and entering partnerships that require a high level of information sharing.

These types of scenarios require organizations to conduct in-depth security audits and evaluations on a regular basis to ensure that existing and potential security threats are addressed to the extent possible.

It is important to evaluate the potential security implications for every network application and technology that you deploy. It is also important to consider the security technologies in Windows Server 2003 any time that you plan to introduce new capabilities in your organization. Security technologies in Windows Server 2003 provide many enhanced security advantages, including the following:

- Enhanced single sign on by means of Kerberos authentication and domain trusts.
- Greater control over access control by means of applications that apply roles-based authorization.
- Increased flexibility and security for employees who are traveling and accessing the network remotely, by means of the use of EFS and Routing and Remote Access connectivity.
- Improved password protection by means of PKI-enabled smart card authentication or security policy settings that require users to set complex passwords and change them regularly.
- The ability to create security groups and GPOs that allow contractors and business partners limited access to set ACLs on carefully defined sets of resources.

# Logon and Authentication Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logon and Authentication Technologies**
The Microsoft Windows Server 2003 operating system implements a default set of authentication protocols, including Negotiate, Kerberos, NTLM, Schannel (secure channel), and Digest, as part of an extensible architecture. These protocols enable authentication of users, computers, and services; the authentication process, in turn, enables authorized users and services to access resources in a secure manner.

**Authentication Concepts**
Authentication is a process for verifying the identity of an object or person. When you authenticate an object, the goal is to verify that the object is genuine. When you authenticate a person, the goal is to verify that you are not dealing with an imposter.

Similarly, in a networking context, authentication is the act of proving identity to a network application or resource. Typically, identity is proven by a cryptographic operation that uses either a key only the user knows — as with public key cryptography — or a shared key. The server side of the authentication exchange compares the signed data with a known cryptographic key to validate the authentication attempt.

Storing the cryptographic keys in a secure central location makes the authentication process scalable and maintainable. Active Directory directory service is the recommended and default technology for storing identity information (including the cryptographic keys that are the user's' credentials). Active Directory is required for default NTLM and Kerberos implementations.

Authentication techniques range from a simple logon, which identifies users based on something that only the user knows — like a password, to more powerful security mechanisms that use something that the user has — like tokens, public key certificates, and biometrics. In a business environment, users might access multiple applications on many types of servers within a single location or across multiple locations. For these reasons, authentication must support environments for other platforms, for Windows Server 2003, and for other Windows operating systems.

**Authentication and Authorization: A Travel Analogy**
A travel analogy can help explain how authentication works. For example, when you cross an international border, a border guard asks for credentials and you present a passport. The process is two-fold:

- The guard authenticates the passport by verifying that it was issued by a security authority that the local government trusts — trusts, at least, to issue passports — and by verifying that the passport has not been modified.
- The guard authenticates you by verifying that your face matches the face of the person pictured on the passport.

If the passport proves to be valid and you prove to be its owner, you are who you say you are.

Transitive trust between security authorities is the foundation of authentication; the type of authentication that takes place at an international border is based on trust. The local government does not know you, but it trusts that your government does. When your government issued your passport, it did not know you either. It trusted the agency that issued your birth certificate. The agency that issued your birth certificate, in turn, trusted the physician who signed the certificate. The physician witnessed your birth and stamped the certificate with your footprint, direct proof of your identity. Trust that is transferred in this way, through trusted intermediaries, is transitive.

Similarly, transitive trust is the foundation for network security in Windows Server 2003. A trust relationship flows throughout a set of domains, such as a domain tree, and forms a relationship between a domain and all domains that trust that domain. For example, if domain A has a transitive trust with domain B, and if domain B trusts domain C, then domain A trusts domain C.

There is a difference between authentication and authorization. With authentication, the system proves that you are who you say you are; however with authorization, the system verifies that you have rights to do what you want to do. To take the border analogy to the next step, merely authenticating that you are the proper owner of a valid passport does not necessarily authorize you to enter a country. Residents of a particular country are allowed to enter another country by simply presenting a passport only in situations where the country being entered grants blanket permission for all citizens of that particular country to enter.

Similarly, you can grant all users from a certain domain permissions to access a resource. Any user who belongs to that

domain has access to the resource, just as Canada allows U.S. citizens to enter Canada. However, U.S. citizens attempting to enter Brazil or India would find that they cannot enter those countries merely by presenting a passport because both of those countries require visiting U.S. citizens to have a valid visa. Thus, authentication does not guarantee access (or authorization).

**Single Sign-On**

Single sign-on (SSO) makes it possible for users to access resources over the network without having to repeatedly supply their credentials. Forcing users to enter multiple passwords — or to enter a single password multiple times — to access system resources is inefficient and provides an unpleasant user experience.

Users of the Windows Server 2003 operating systems must supply their credentials only once to access network resources; subsequent authentication is transparent to users. SSO allows users to authenticate once with the system in order to access all applications and data sources that they are authorized to use without entering another account ID or password. The Microsoft Windows platform capitalizes on the ability to use a single user identity (maintained by Active Directory) across the network by locally caching user credentials in the operating system's Local Security Authority (LSA). When a user logs on to the domain, Windows authentication packages transparently use the credentials to provide SSO when authenticating to network resources.

**Authentication Architecture**

**Windows Server 2003 Authentication Services**

Windows Server 2003 provides a method for applications to authenticate users by using the Security Support Provider Interface (SSPI) to abstract calls to authentication. Thus, developers do not need to understand the complexities of specific authentication protocols or build authentication protocols into their applications.

For best results, however, before starting the application design process, developers should understand the different types of authentication, the effects of using different protocols, and the interfaces they support. Then they can specify their preferred authentication protocol and use the authentication protocols provided by the Windows platform.

If the preferred protocol is not in Windows Server 2003, developers can use a custom Security Support Provider (SSP). Implementing authentication protocols on an application-by-application basis is inefficient and can introduce security flaws across the organization.

**Security Subsystem Architecture**

Windows Server 2003 operating systems include a set of security components that make up the Windows security model. These components ensure that applications cannot gain access to resources without authentication and authorization.

The Local Security Authority (LSA) is a protected subsystem that authenticates and logs users on to the local computer. In addition, LSA maintains information about all aspects of local security on a computer (these aspects are collectively known as the local security policy), and it provides various services for translation between names and security identifiers (SIDs).

The security subsystem keeps track of the security policies and the accounts that are in effect on a computer system. In the case of a domain controller, these policies and accounts are the ones that are in effect for the domain in which the domain controller is located. These policies and accounts are stored in Active Directory.

The local security policy identifies the following:

- Which domains are trusted to authenticate logon attempts.
- Who can have access to the system and in what way (for example, interactively, over the network, or as a service).
- Who is assigned what rights.
- What security auditing is performed.
- What the default memory quotas are for paged and non-paged memory pool usage.

**LSA Architecture**




The LSA security subsystem provides services for validating access to objects, checking user rights, and generating audit messages. A local procedure call (LPC) occurs between components on the same system. A remote procedure call (RPC) occurs between components on different systems, as do LDAP communications.

In general, the LSA performs the following functions:

- Manages local security policy.
- Provides interactive user authentication services.
- Generates access tokens.
- Manages the audit policy and settings.

### LSA Components for All Systems

| Component | Description |
|---|---|
| Winlogon.exe | The executable file responsible for managing secure user interactions. Winlogon initiates the logon process for Windows Server 2003, Microsoft Windows 2000, Microsoft Windows XP, and Microsoft Windows NT 4.0 operating systems. |
| Msgina.dll | The default dynamic-link library (DLL) module that operates in the security context of Winlogon. The Graphical Identification and Authentication (GINA) is responsible for processing secure attention sequence (SAS) events and activating the user's shell. The GINA can generate SAS events in some instances. |

| | |
|---|---|
| Netlogon.dll | Some of the services that Net Logon performs include:<br>• Net Logon maintains the computer's secure channel (not to be confused with Schannel) to a domain controller.<br>• Net Logon passes the user's credentials through a secure channel to the domain controller and returns the domain SIDs and user rights for the user.<br>• In the Windows 2000 Server and Windows Server 2003 operating systems, the Net Logon service publishes service resource records in the Domain Name System (DNS) and uses DNS to resolve names to the Internet Protocol (IP) addresses of domain controllers.<br>• Net Logon also implements the replication protocol based on RPC for synchronizing Windows NT 4.0 primary domain controllers (PDCs) and backup domain controllers (BDCs). |
| Msv1_0.dll | The NTLM authentication protocol. This protocol authenticates clients that do not use Kerberos authentication. |
| Schannel.dll | The Secure Sockets Layer (SSL) and Transport Layer Security (TLS) authentication protocol. This protocol provides authentication over an encrypted channel instead of a less-secure clear channel. |
| Kerberos.dll | The Kerberos V5 authentication protocol. This protocol provides authentication using Kerberos protocol instead of plaintext, NTLM, or digest method. |
| Wdigest.dll | Simple challenge-and-response protocol that provides increased security over basic authentication. |
| Lsasrv.dll | The LSA Server service, which both enforces security policies and acts as the security package manager for the LSA. |
| Samsrv.dll | The Security Accounts Manager (SAM), which stores local security accounts, enforces locally stored policies, and supports APIs. |
| Secur32.dll | The multiple authentication provider that holds all of the components together. |

**LSA Components Unique to Domain Controllers**

| Component | Description |
|---|---|
| Kdcsvc.dll | The Kerberos Key Distribution Center (KDC) service, which is responsible for providing ticket-granting tickets to clients. |
| Ntdsa.dll | The directory service module, which supports the Windows Server 2003 and Windows 2000 replication protocol and LDAP, and manages partitions of data. |

**Security Support Provider Interface**

The Security Support Provider Interface (SSPI) is the foundation for authentication in Windows Server 2003. SSPI is the implementation of the Generic Security Service API (GSSAPI) in Windows Server 2003. For more information about GSSAPI, see RFCs 2743 and 2744 in the IETF RFC Database.

The default Security Support Providers in Windows Server 2003 — Negotiate (SPNEGO), Kerberos, NTLM, Schannel, and Digest authentication protocols — are plugged into the SSPI in the form of DLLs. Additional SSPs can be plugged in if they are interoperable with the SSPI.

**SSPI in Authentication**

The SSPI in Windows Server 2003 provides a mechanism that carries authentication tokens over the existing protocol, thus eliminating the need for communicating parties to specify a network protocol for use during authentication. When two parties need to be authenticated so that they can communicate, the requests for authentication are routed to the SSPI, which completes the authentication process, regardless of the network protocol currently in use.

**SSPI Architecture**




Any application can make a request of the SSPI. Each of these requests goes through the SSPI, for example:

- Winlogon sends requests to the Local Security Authority (LSA), which obtains tickets to access the local computer.
- Internet Explorer obtains tickets to access information about a Web site.
- An LDAP client obtains tickets to enable access to information in an x500 directory, such as Active Directory.

Some applications integrate use of the SSPI into their authentication designs. When such applications need to authenticate a user, they ask the SSPI to complete the authentication by using specific parameters, such as the user's name and secret key (if they were not cached from a previous logon), and the service that the user wants to access. The SSPI receives the request, examines its contents, and passes the request to the SSP specified in the request. The SSP handles the details of the authentication and sends a success or failure message back to the SSPI, which relays the message back to the application that initially requested authentication. SSPI enables an application to use various security models available on a computer or network without changing the interface to the security system.

For example, if a client has a service ticket for a server for which the preferred authentication protocol is Kerberos, and the client wants to issue that server a remote procedure call (RPC), the client's Kerberos SSP generates a request message based on standard GSS-API (RFC 1964) calls, which the application sends to the server. When the server receives the request, the server passes the ticket to the Kerberos SSP. The client routes the reply token to its SSP, which then evaluates and manages the contents of the reply. The server and the client know nothing about the contents of the token that is used to enable secure communication, because this information is handled by the SSP. The details of authentication are handled behind the scenes, and the results of the authentication process are passed back to the applications.

**Security Support Provider Selection**

The SSPI in Windows Server 2003 can use any of the five protocols — Negotiate, Kerberos, NTLM, Schannel, and Digest — that it supports. However, because not all operating systems support the same SSP packages that Windows Server 2003 supports; clients and servers must negotiate to use a protocol that they both support. Windows Server 2003 prefers clients to use Kerberos, a strong standards-based protocol, when possible, but continues to allow clients that do not support Kerberos, such as Windows NT 4.0 clients, to authenticate.

Before authentication can take place, however, the two communicating computers must agree on a protocol that they both can support. For any protocol to be usable through the SSPI, each party must have the appropriate SSP. For example, in order for a client and server to use the Kerberos authentication protocol, they must both support Kerberos V5. Windows Server 2003 uses the function **EnumerateSecurityPackages** to identify which SSPs are supported on a computer and what the capabilities of those SSPs are.

The selection of an authentication protocol can be handled in one of two ways, each of which is described in more detail below:

- A single authentication protocol can be specified.
- The negotiate option can be used to allow the parties to attempt to find an acceptable protocol. This is based on the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO, RFC 2478).

**Single Authentication Protocol Specified**

When a single acceptable protocol is specified, the server allows only one acceptable authentication method for the exchange. If the client does not support the protocol specified, the communication fails. When a single acceptable protocol is specified, the authentication exchange takes place as follows:

1. The client requests access to a service.

2. The server replies with an authentication challenge in which it specifies the protocol that be used.

3. The client examines the contents of the reply and checks to determine whether it supports the specified protocol. If the client does support the specified protocol, the authentication continues. If the client does not support the protocol, the authentication fails, regardless of whether the client is authorized to access the resource.

For example, if a Windows Server 2003 client received a reply that specified the Kerberos protocol, the authentication exchange could proceed. On the other hand, if a Windows NT 4.0 client received the same reply, the authentication exchange would fail, because computers running versions of Windows earlier than Windows 2000 do not support Kerberos authentication. If the specified protocol fails, then negotiation is needed.

**Negotiate Option Is Used**
When the authentication begins with the option to negotiate for an authentication protocol, the SPNEGO exchange takes place as follows:

1. The client requests access to a service.

2. The server replies with a list of authentication protocols that it can support and an authentication challenge based on the protocol that is its first choice. For example, the server might list Kerberos and NTLM, and send a Kerberos challenge.

3. The client examines the contents of the reply and checks to determine whether it supports any of the specified protocols.

    - If the client supports the preferred protocol, authentication proceeds.

    - If the client does not support the preferred protocol, but does support one of the other protocols listed by the server, the client lets the server know which authentication protocol it supports, and the authentication proceeds.

    - If the client does not support any of the listed protocols, the authentication exchange fails.

For example, if the server specified Kerberos and NTLM, and returned a Kerberos challenge, a Windows Server 2003 client could immediately authenticate by using Kerberos. A Windows NT 4.0 client could negotiate to complete the authentication exchange by using NTLM.

**Kerberos SSP**
The Kerberos SSP (Kerberos.dll) is the first choice for authentication in Windows Server 2003. The Kerberos SSP uses the Kerberos V5 authentication protocol (RFC 1510).

Because Kerberos is the default authentication protocol for Windows Server 2003, all domain services support the Kerberos SSP, which includes:

- Active Directory queries using the Lightweight Directory Access Protocol (LDAP).

- Remote server or workstation management using RPC calls.

- Print services.

- Client-server authentication.

- Remote file access using Common Internet File System/server message block (CIFS/SMB).

- Distributed file system management and referrals.

- Intranet authentication to Internet Information Services (IIS).

- Security authority authentication for Internet Protocol security (IPSec).

- Certificate requests to Certificate Services for domain users and computers.

**NTLM SSP**
Windows Server 2003 supports the NTLM SSP, Msv1_0.dll, to enable clients running versions of Windows earlier than Windows 2000 to authenticate. NTLM is the default authentication protocol for Windows NT 4.0. The NTLM SSP includes the NTLM and NTLMv2 authentication protocols.

Windows Server 2003 can use the NTLM SSP for the following:

- Client/server authentication
- Print services
- File access using CIFS/SMB
- Secure RPC/DCOM-based services

**Negotiate SSP**
Microsoft Negotiate is a security support provider (SSP) that acts as an application layer between the Security Support Provider Interface (SSPI) and the other SSPs. When an application calls in to the SSPI to log on to a network, it can specify an SSP to process the request. If the application specifies Negotiate, Negotiate analyzes the request and picks the best SSP to handle the request based on customer-configured security policy.

Negotiate implements RFC 2478, "The Simple and Protected GSS-API Negotiation Mechanism."

Currently, the Negotiate security package selects between Kerberos and NTLM. Negotiate selects Kerberos by default unless Kerberos cannot be used by one of the systems involved in the authentication or the calling application did not provide sufficient information to use Kerberos.

**Secure Channel (Schannel) SSP**
The Schannel SSP is used for access to Web-enabled services, such as mail or personal information served on Web pages.

The Schannel SSP, which includes a suite of four authentication protocols, uses public key certificates to authenticate parties. When authenticating parties, Schannel SSP selects one of the four protocols in the following order of preference:

- Transport Layer Security (TLS).
- Secure Socket Layer (SSL) version 3.0.
- Private Communications Technology (PCT). PCT is disabled by default in Windows Server 2003.
- SSL version 2.0.

The protocol that is selected is the preferred authentication protocol that both parties can support. For example, if a server supports all four Schannel protocols and the client supports only SSL 3.0 and SSL 2.0, the authentication process uses SSL 3.0.

**Digest SSP**
Digest SSP (Wdigest.dll) is the preferred authentication protocol for certain applications, such as LDAP and HTTP. Digest SSP is used for the following:

- Internet Explorer (IE) and Internet Information Services (IIS) access.
- LDAP queries.

**Passport SSP**
Windows Server 2003 contains no specific Passport SSP. Instead, passport authentication is enabled by an Internet Server Application Programming Interface (ISAPI) DLL called Passport.dll that captures and processes any cookie information that contains passport credentials.

**Authentication Protocols**
Each authentication protocol is composed of a series of steps, used by two or more parties, that verifies the identity of one or more of the parties.

The authentication process is vulnerable to numerous types of security attacks. For example, an attacker who has access to authentication information can steal keys to the network, modify messages, or replay old messages.

The authentication protocols in Windows Server 2003, which include Kerberos, NTLM, Schannel, Digest, and Passport, are designed to address potential security threats by verifying the identities of one or more of the communicating parties.

**Windows Server 2003 Authentication Protocols**

| Authentication Protocol | Description |
|---|---|
| Kerberos V5 authentication | An industry standard protocol that is used with either a password or a smart card for interactive logon. It is also the default method of network authentication for services in Windows 2000, Windows XP, and Windows Server 2003. |
| Digest authentication | An industry standard that is used in Windows Server 2003 for Lightweight Directory Access Protocol (LDAP) and Web authentication that transmits credentials across the network as an MD5 hash or message digest. |
| Negotiate | Negotiate implements RFC 2478, "The Simple and Protected GSS-API Negotiation Mechanism."<br><br>Currently, the Negotiate security package selects between Kerberos and NTLM. By default, Negotiate selects Kerberos unless Kerberos cannot be used by one of the systems involved in the authentication, or the calling application did not provide sufficient information to use Kerberos. |
| NTLM authentication | A challenge response protocol that is used to provide compatibility with versions of Windows earlier than Windows 2000. |
| Passport authentication | A user-authentication service that provides a Web-based authentication mechanism enabling Web-services to offer a single sign-in service. |
| Secure Sockets Layer/Transport Layer Security (SSL/TLS) authentication | A protocol that is used for Web-based server authentication, such as when a user attempts to access a secure Web server. |

Each of these protocols is used in different ways in Windows Server 2003 to promote secure communication in a network environment that is not secure.

**Logon Authentication Scenarios**

It is not possible to use a computer that is running Windows Server 2003 without using at least some system services. To use a system service, the account must be a client of the service, and all clients must first be authenticated by the service. Authentication requires the service to know the user's credentials.

On computers running Windows Server 2003, system services can run under various accounts depending on the way that the applications are written. For example, a computer can access domain resources by using the network service, local system, or domain user accounts. Domain users who log on interactively must also be assigned the appropriate permissions before they can access the computer's services.

**Interactive Logon**

The logon process begins when a user enters credentials into the Log On to Windows dialog box or when the user inserts a smart card into the smart card reader. Credentials consist of a user's user name and password for local logons and the user's user name, password, and domain for domain logons. The process confirms the user's identification to the security database on the user's local computer or to an Active Directory domain. This mandatory logon process cannot be turned off for users in a domain.

Users can perform an interactive logon to a computer in either of two ways:

- Locally, when the user has direct physical access to the computer.
- Remotely, through Terminal Services, in which case the logon is further qualified as remote interactive.

After an interactive logon, Windows runs applications on the user's behalf and the user can interact with those applications.

Users can perform an interactive logon by using a local user account or a domain account to log on to a computer.

**Local Logon**
A local logon grants a user access to Windows resources on the local computer. A local logon requires that the user have a user account in the Security Accounts Manager (SAM) on the local computer. The SAM protects and manages user and group information in the form of security accounts stored in the local computer registry. The computer can have network access, but it is not required. Local user account and group membership information is used to manage access to local resources.

**Domain Logon**
A domain logon grants a user access to local and domain resources. A domain logon requires that the user have a user account in Active Directory. The computer must have an account in the Active Directory domain and be physically connected to the network. Users must also have the user rights to log on to a local computer or a domain. Domain user account information and group membership information is used to manage access to domain and local resources.

**Smart Card Logon**
Logging on to a domain with a smart card provides a strong form of authentication because smart cards use cryptography-based identification and proof of possession when authenticating a user to a domain.

To log on to a domain with a smart card, a user does not need to press CTRL+ALT+DEL and enter credentials into the Log On to Windows dialog box. The user merely inserts the smart card into the smart card reader and the computer prompts for the user's personal identification number (PIN) instead of the user name and password.

Smart cards are a tamper-resistant, portable way to provide security solutions for tasks such as client authentication, code signing, securing e-mail, and logging on to a Windows Server operating systems domain.

Smart cards provide:

- Tamper-resistant storage for protecting private keys and other forms of personal information.
- Isolation of security-critical computations involving authentication, digital signatures, and key exchange from other parts of the computer that do not have a "need to know." The security critical operations are all performed on the smart card.
- Portability of credentials and other private information between computers at work, home, or on the road.

**Network Logon**
A network logon can only be used after a user, service, or computer authentication has taken place. During network logon, the process does not use the logon dialog boxes, such as the Log On to Windows dialog box, to collect data. Instead, previously established credentials or another method to collect credentials is used. This process confirms the users identification to any network service that the user is attempting to access. This process is typically invisible to the user unless alternate credentials need to be provided.

To provide this type of authentication, the security system includes these authentication mechanisms:

- Kerberos V5
- Public key certificates
- Secure Sockets Layer/Transport Layer Security (SSL/TLS)
- Digest
- NTLM (for compatibility with Windows NT 4.0-based systems)

# Digest Authentication Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Digest Authentication Technical Reference**
Digest Authentication is a standards-based authentication protocol, implemented through wdigest.dll, that provides for authentication, between Microsoft Windows operating environments and operating environments other than Windows, over the Internet. It is described in RFC 2617: HTTP Authentication: Basic and Digest Access Authentication. Wdigest.dll was first included as an authentication package in the Microsoft Windows XP operating system. The Digest security support provider (SSP) implements both RFC 2617 and RFC 2831: Using Digest Authentication as a SASL Mechanism. Digest Authentication, as a simple authentication and security layer (SASL) mechanism, is used primarily for Lightweight Directory Access Protocol (LDAP) authentication.

Digest Authentication works in situations where the Kerberos protocol does not. Digest Authentication offers single sign-on only to a single Web URL protection space. If users navigate to a different Web site, or even to a different server in the same site, they will usually be prompted to enter credentials again.

This section explains what Digest Authentication is and how it works, and it explains which tools are available for configuring Digest Authentication. This technology is also referred to in Internet Information Services (IIS) 6.0 documentation as Advanced Digest Authentication.

**In This Section**


- What is Digest Authentication?
- How Digest Authentication Works
- Digest Authentication Tools and Settings

# What is Digest Authentication?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Digest Authentication?**

**In This Section**


- Common Digest Authentication Scenarios
- Technologies Related to Digest
- Digest Authentication Dependencies


WDigest.dll was introduced in the Windows XP operating system The Digest Authentication protocol is designed for use with Hypertext Transfer Protocol (HTTP) and Simple Authentication Security Layer (SASL) exchanges, as documented in RFCs 2617 and 2831. These exchanges require that parties that seek to authenticate must demonstrate their knowledge of secret keys. This process improves upon earlier versions of HTTP authentication, in which users provide passwords that are not encrypted when they are sent to a server, leaving them vulnerable to capture by attackers, or that are encrypted but sent in an expensive, ongoing, Secure Sockets Layer (SSL) session.

Digest Authentication has similar security characteristics to the proprietary NTLM protocol. Both Digest Authentication and NTLM are challenge/response protocols. Challenge/response protocols require an authenticating server to generate a challenge containing some amount of unpredictable data. A client then uses a key derived from the user's password to encrypt the challenge and forms a response. The server, or a trusted service such as Active Directory, can verify that the user possesses the correct password by comparing the client's encrypted response to a stored response based on the credential associated with the user in Active Directory or in the server account database for local users. If the responses match, the user is authenticated.

**Limitations of Digest Authentication**

SSL and Transport Layer Security (TLS) are often used to protect Digest Authentication from an offline attack against the Digest Authentication challenge/response.

Digest Authentication offers single sign-on only to a single Web URL protection space. If users navigate to a different Web site, or even to a different server in the same site, they will usually be prompted to enter credentials again.

**Digest Authentication Protocol Standards**

Digest Authentication is a standards-based authentication protocol that provides for authentication, between Windows operating environments and operating environments other than Windows, over the Internet. It is described in RFC 2617. Windows Server 2003 implements Digest Authentication as a simple authentication and security layer (SASL) mechanism that is used primarily for LDAP authentication, as described in RFC 2831. SASL is method for adding authentication support to connection-based protocols.

**Common Digest Authentication Scenarios**

Many people think of Digest Authentication as a protocol that is used with Web browsers for authenticating users browsing the Internet. However, Digest Authentication is also a general purpose protocol that can be used for authentication, and by using SASL, it can provide integrity protection. For example, you can use Digest Authentication for:

- Authenticated client access to a Web site
- Authenticated client access using SASL
- Authenticated client access with integrity protection to a directory service using LDAP

**Authenticated client access to a Web site**

Digest Authentication can be used to provide user authentication when users access pages on a Web server.

**Authenticated client access to a directory service using LDAP**

Digest Authentication can be used as a SASL mechanism for any protocol that has a SASL profile. Using it as a SASL mechanism is a convenient way to support a single authentication mechanism for Web, e-mail, LDAP, and other protocols.

**Authenticated client access with integrity protection to a directory service using LDAP**
When using SASL, integrity protection can be added.

This is not an exhaustive list. The ability to access Digest through the SSPI interface means that developers can take advantage of Digest Authentication for just about any application. Because Digest Authentication uses a *message digest* function to enable authentication, it works well with devices such as personal digital assistants (PDAs) that have little processor power and that need to have authenticated access to resources to read e-mail, access appointment information on an Exchange server, or view Web pages on a server running services such as Internet Information Services (IIS).

The Windows Server 2003 implementation of Digest Authentication is integrated with Active Directory. This integration enables you to centrally manage accounts, but you must consider how accounts for employees, partners, and customers are provisioned and maintained. Because Digest Authentication is user name- and password-based, you must think about how you manage user accounts, password changes, password resets, and entitlements for your employees, partners, and customers.

**Technologies Related to Digest Authentication**
The following diagram shows how Digest Authentication fits with other technologies in Windows Server 2003. Depending on whether the client application or server application are user-mode or kernel-mode applications, they will use either Secur32.dll or Ksecdd.sys respectively via SSPI calls to communicate with the Local Security Authority Subsystem (LSASS).

### Digest Authentication Architecture




Below is a description of the components that participate in Digest Authentication.

### Security Subsystem Components used in Digest Authentication

| Component | Description |
|---|---|
| Wdigest.dll | The SSP that implements an industry standard and that is used in Windows Server 2003 for LDAP and Web authentication. Digest Authentication transmits credentials across the network as an MD5 hash or message digest and thus provides increased security over basic authentication. |
| Ksecdd.sys | The Kernel Security Device Driver is used to communicate with LSASS in user mode. |
| Lsasrv.dll | The LSA Server service, which both enforces security policies and acts as the security package manager for the LSA. |

| | |
|---|---|
| Netlogon.dll | The Net Logon service performs Digest Authentication- relevant services:<br>• Net Logon passes the user's credentials through a secure channel to the domain controller and returns the domain SIDs and user rights for the user. |
| Secur32.dll | The Secur32.dll component is the multiple authentication provider that implements SSPI for user mode applications. |
| Samsrv.dll | The Security Accounts Manager (SAM) stores local security accounts, enforces locally stored policies, and supports APIs. |

SSPI is an application interface that provides the security services for Windows Server 2003. SSPI supports the Digest Authentication SSP, wdigest.dll, which is the preferred authentication protocol for some applications and which uses protocols such as LDAP and HTTP.

The Digest SSP is used for the following:

- Internet Explorer (IE) and IIS access
- LDAP queries

**Digest Authentication Dependencies**
Digest Authentication depends on several related technologies and resources to function properly. The following section describes these technologies and resources and summarizes how they relate to Digest Authentication.

**Active Directory Domains**
Digest Authentication is not supported in earlier operating systems, such as Windows NT.

- Users and services must have a valid Active Directory domain account.
- The Web server must be a member of the same forest as the user accounts.

For more information about Active Directory, see the Active Directory Technical Reference.

**Operating Systems**
Wdigest.dll is built into Windows Server 2003 and Windows XP. If a client or server is running an earlier operating system, it will not use the Digest Authentication described in How Digest Authentication Works, but the older implementation of Digest.

**Windows Server 2003 Domain Controllers**
All domain controllers for the domains of the users and services using Digest Authentication must be running on Windows Server 2003 to use the latest implementation of Digest Authentication. This requirement is because of password hash storage in Active Directory. The domains do not need to be configured for the Windows Server 2003 domain functional level.

If any of the domain controllers in the account domains are running Windows 2000 Server, then subauthentication, which requires reversible encryption, is required for Digest Authentication to work.

# How Digest Authentication Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Digest Authentication Works**

**In This Section**


- Digest SSP Architecture
- Digest Authentication Processes and Interactions
- Network Ports Used by Digest Authentication
- Related Information


**Digest SSP Architecture**

The Windows Server 2003 operating system implements the Digest Authentication protocol as a security support provider (SSP), a dynamic-link library (DLL) that is supplied with the operating system. Which SSP is used depends on the capabilities of the computer on the other side of the connection and the preferences of the individual application that is being used.

The Microsoft Security Support Provider Interface (SSPI) is the foundation for authentication in Windows Server 2003. Applications and infrastructure services that require authentication use SSPI to provide it.

SSPI is the implementation of the Generic Security Service API (GSSAPI) in Windows Server 2003. For more information about GSSAPI, see RFC 2743 and RFC 2744 in the IETF RFC Database.

The default SSPs in Windows Server 2003 — Negotiate (SPNEGO), Kerberos, NTLM, Schannel, and Digest Authentication protocols — are incorporated into the SSPI as DLLs. Additional SSPs can be incorporated if they can operate with the SSPI.

**SSPI Architecture**




The SSPI in Windows Server 2003 provides a mechanism that carries authentication tokens over the existing communication channel between the client and server. When two parties need to be authenticated so that they can communicate securely, the requests for authentication are routed to the SSPI, which completes the authentication process, regardless of the network protocol currently in use. The SSPI returns transparent binary large objects, these are passed between the applications, at which point they can be passed to the SSPI layer on that side. Thus, the SSPI enables an application to use various security models available on a computer or network without changing the interface to the security system.

The following table describes the SSP components that are plugged into the SSPI. Each of the protocols in the table is used in different ways in Windows Server 2003 to promote secure communication in an unsecure network environment.

**SSP Layer Components**

| Component | Description |
|---|---|
| Kerberos V5 Authentication | An industry-standard protocol that is used with either a password or a smart card for interactive logon. It is also the preferred authentication method for services in Windows 2000 Server, Windows 2000 Professional, and Windows Server 2003. |

| | |
|---|---|
| NTLM Authentication | A challenge-response protocol that is used to provide compatibility with versions of Windows earlier than the Windows 2000 operating systems. |
| Digest Authentication | An industry standard that is used in Windows Server 2003 for Lightweight Directory Access Protocol (LDAP) and Web authentication. Digest Authentication transmits credentials across the network as an MD5 hash or message digest. |
| Schannel | An SSP that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) Internet standard authentication protocols. Schannel is used for Web-based server authentication such as when a user attempts to access a secure Web server. |
| Negotiate | An SSP that can be used to negotiate a specific authentication protocol. When an application calls into SSPI to log on to a network, it can specify an SSP to process the request. If the application specifies Negotiate, Negotiate analyzes the request and picks the best SSP to handle the request based on customer-configured security policies. |

**Digest SSP**

The Digest SSP, which is implemented through wdigest.dll, is one of the authentication protocols for some applications such as LDAP and HTTP. Depending on your deployment scenario, it might be the preferred protocol.

The Digest SSP is used for the following:

- Internet Explorer (IE) and Internet Information Services (IIS) access
- LDAP queries

**Digest Authentication Processes and Interactions**

**Digest Authentication Scheme**

### Overview of a Digest Authentication Scheme That Is Using HTTP




The Digest Authentication process involves the following steps:

1. The client sends the server a request for an access-protected resource.
2. The server receives the request from the client and recognizes that the request is for information or a service that requires authentication. The server checks whether the client has provided authentication information. Because this is the client's first request, it does not find any. The server responds by returning an HTTP WWW-Authenticate header to the client. This header includes the following:
    - The domain for which the server must supply credentials.
    - Information about whether authentication is required. If the exchange is Secure Authentication Socket Layer (SASL)-based, integrity and confidentiality might also be specified.
    - Server-side data, which often includes items such as a timestamp, nonce and an HTML tag that references the Uniform Resource Indicator (URI) requested, and is protected in the server's private key.
    - Hints about the security context.
    - A list of ciphers that the server supports.

- Specification of the encryption algorithm that is used in this authentication protocol. If the authentication is SASL-based, rather than HTTP-based, RC4 is the algorithm that is used in Windows Server 2003, although this is extensible and could be designed to use other algorithms.

The header is returned with a 401 message code, indicating that the client's request has generated an unauthorized access error.

3. The client receives the server's challenge and gathers the required credentials, either by accessing information in existing credentials or by prompting the user. The client's challenge response is then hashed with the user's secret key. This response includes the following parts:

   - The original nonce protected with the user's secret key, proving the client's knowledge of the user's password.
   - The user name and domain.
   - The original nonce and timestamp sent by the server, the number of times this number has been used, and a new random number, called a nonce number, which is unique and is used only once. To prevent replay attacks nonce counts are used.
   - The URI, which is the location of the resource that the client initially requested.
   - Information about whether authentication, integrity, and confidentiality is used.
   - Specification of the algorithms that are used.

4. The server receives the message and, as part of the authentication of the challenge response, the following checks are made:

   - That the challenge response is to a challenge issued by the server, and not another source.
   - That the challenge response message is properly formatted and the server-enforced options are valid.
   - That the nonce number has not been used before in challenge responses, to ensure that this is not a replay attack.

   If any of these checks fail, the server fails the authentication of the request.

   If the response meets the server's criteria, it relays the client's response to a domain controller, which completes the authentication. The Digest SSP checks whether the Digest Authentication response is valid and ensure that access restrictions, such as an expired password, do not apply.

   If the check is successful, the domain controller notifies the server that the client is authentic. A Digest Authentication session key is also returned to the originating server. If the check is not successful, the originating server sends the client an error message. After a successful authentication, the originating server allows the client to access the resource indicated by the URI.

   The Digest Authentication session key returned by the authenticating domain controller is cached by the originating server for use in authenticating future requests.

**Authentication of Subsequent Requests**

**Example Subsequent Request Using HTTP**




For subsequent requests, the client simply needs to send an authorization header with the security context reference to the server, because the security context has already been established with the server.

**Digest Authentication Headers**

Because Digest Authentication is accomplished through other protocols such as HTTP and LDAP, what is defined next is the challenge and challenge response that applications will transport to the Digest SSP.

**WWW-Authenticate Response Header**

When the server receives a request for an access-protected resource, the server replies with a "401 Unauthorized" status code and a WWW-Authenticate header with the digest scheme below. This section describes the directives that make up a Digest Authentication challenge.

RFC 2831 requires the size of a Digest Authentication challenge must be less than 2048 bytes. The Digest SSP enforces this for all Digest Authentication challenges. Following digest scheme is used:

Challenge = "Digest" digest-challenge

HTTP Authentication Digest-challenge = 1# (realm=*realm* | [domain="*domainname*"] | nonce="*nonce-value*" | [opaque="*opaque-value*"] | [stale=("true" | "false")] | [algorithm=("MD5" | "MD5-sess" | token] | [qop=1#("auth" | "auth-int" | token)] | [auth-parm]

SASL Digest-challenge = 1# ("*realm*" | nonce="*nonce-value*" | [qop=1#("auth" | "auth-int" | "auth-conf" | token")] | [stale=(true | false)] | [maxbuf=*maxbuf*] | [charset=*charset*] | [algorithm=("MD5" | "MD5-sess" | token] | [cipher=("3des" | "des" | "rc4-40" | "rc4" | "rc4-56" | token) | [auth-parm]

**Note**

- RFC 2617-compliant Digest Authentication challenges and responses must also comply with RFC 2616: Hypertext Transfer Protocol -- HTTP/1.1 quoted string requirements. This requirement particularly affects the use of backslash (\) and embedded double quotes. Both must be preceded (escaped) with a backslash.

- For example, *domain\username* according to RFC 2616 is read as *domainusername*. This reading is important because if an application sends information in this format rather than as *domain\\username,* authentication fails.

- However, because this is a known issue with *domain\username* , if authenticating with backslash encoding fails, Digest SSP attempts to authenticate the response and assumes that the backslash is part of the string. This behavior can be turned off by setting the ServerCompat registry key.

A Digest Authentication challenge can contain the following directives, which are defined in the relevant RFCs.

**Realm**

An implementation-defined string that is displayed to users if Digest Authentication is prompting users for credentials. This string helps users know which user name and password is required.

**Domainname (optional, not used in-SASL)**

A space-separated list of URIs that are contained in quotation marks and that define the protection space.

**Nonce-value**

A unique, encoded value that is generated by the server for each challenge. This value must not be altered by the client.

**Opaque-value (optional, not used in-SASL)**

A string of data specified by the server that contains a reference for the security context that is being established.

**Stale (optional)**

A flag that when set to TRUE indicates that the previous request from the client was rejected because the nonce value was stale, but the digest was valid. When not set to TRUE or not present then the user name and/or password are invalid and new values must be obtained.

**Algorithm (optional)**
A string indicating that the pair of algorithms used to produce the digest and checksum. If MD5 is not present, it is assumed. Digest SSP supports MD5 and MD5-Sess. For optimal performance, use MD5-Sess.

**QOP-(optional)**
Should be used by all RFC 2617-compliant implementations. The quality of protection (QOP) determines the level of security that is used when sending and receiving messages. The possible values for the QOP directive are described in the following table.

**QOP Options**

| Value | Description |
|---|---|
| "auth" | Authentication only. |
| "auth-int" | Authentication and integrity checking by using signatures. |
| "auth-conf" | (SASL only) Authentication, integrity and confidentiality checking by using signatures and encryption. See Ciphers. |

**Auth-parm (optional)**
Allows for future extensions of the RFC. Unrecognized directives must be ignored.

- **Maxbuf (optional, SASL only)**

  Number indicating the size of the largest buffer the server is able to receive when using "auth-int" or "auth-conf." The default is 65536.

- **Charset (optional)**

  Specified in RFC 2831, the charset directive can be set to UTF-8 if the server can process UTF-8–encoded user names and realms.

  **Note**

  - Digest SSP uses the charset directive for both HTTP and SASL mode. This directive allows the use of UTF-8 in HTTP mode. To maintain backwards compatibility with clients that do not understand the charset directive, the realm in the challenge is encoded in the standard ISO 8859-1 (Latin-1) character set. If the client understands the charset directive, then the username and realm (and password) can be encoded in UTF-8 in the challenge response sent back to the server. If the client does not understand the charset directive, then the standard ISO 8859-1 (Latin-1) must be used.

- **Cipher (SASL only)**

  The list of ciphers that the server supports. This element can be present in a Digest SASL challenge only if the QOP directive specifies auth-conf.

**Authorization Request Header**
The client must send the server a Digest Authentication challenge response when it receives a Digest Authentication challenge in response to a request for a resource. The client must send the request again and supply an Authorization Request header with the challenge response. This section describes the directives that make up a Digest Authentication challenge response.

RFC 2831 requires that the size of a Digest Authentication challenge response must be less than 4096 bytes. The Digest SSP enforces this for all Digest Authentication challenge responses. The following digest scheme is used:

Credentials = "Digest" digest-response

HTTP Authentication Digest-response = 1# (username="*username*"| realm | nonce | uri=digest-uri | response="*response* "| [algorithm] | [cnonce=*cnonce*] | [opaque] | [qop=*message-qop*] | [nc=*nonce-count*] | [auth-param])

SASL Digest-response = 1# (username="*username*"| realm | nonce | [cnonce=*cnonce*] | [nc=*nonce-count*] | [qop=*message-qop*] | uri=digest-uri | response="*response*" | [maxbuf=*maxbuf*] | [charset=*charset*] | cipher=*cipher* | [authzid=*authzid*] | [auth-param])

**Username**
The account name of the security principal requesting the resource.

**Realm**
The name of the domain that contains the account indicated by the user name.

**Nonce**
The nonce received in the Digest Authentication challenge.

**Digest-URI**
The URI of the requested resource.

**Response**
A value computed according to RFC 2617. An accurate response is conclusive proof that the user's password is known on the client side.

**Algorithm (Optional, not used in-SASL)**
The algorithm received in the Digest Authentication challenge.

**Cnonce (Optional)**
A unique, encoded value that is generated by the client for each challenge response, and that is used to avoid chosen plaintext attacks, provides some message integrity protection, and provides mutual authentication. This authentication is provided mutually in that the server proves it knows the user's secret information and not in that the server proves its identity. The nonce must be specified if a QOP directive is sent.

**Opaque (Optional, not used in-SASL)**
The opaque value received in the Digest Authentication challenge.

**Message-QOP (Optional)**
The quality of protection.

**Nonce-count (optional)**
The nonce count. The number of times the client has sent a challenge response by using the nonce. The nonce must be specified if a QOP directive is sent.

**Auth-parm (optional)**
Allows for future extensions of RFC. Unrecognized directives must be ignored.

- **Maxbuf (optional, SASL only)**

  Number indicating the size of the largest buffer that the client is able to receive. The default is 65536.

- **Charset (optional)**

  If the client understands the charset directive, it can respond by using UTF-8–encoded values.

  **Note**

  - Digest SSP uses the charset directive for both HTTP and SASL mode. This allows the use of UTF-8 in HTTP mode. To maintain backwards compatibility with clients that do not understand charset directive, the realm in the challenge is encoded in the standard ISO 8859-1 (Latin-1) character set. If the client understands the charset directive, then the username and realm (and password) can be encoded in UTF-8 in the challenge response sent back to the server. If the client does not understand the charset directive, then the standard ISO 8859-1 (Latin-1) must be used.

- **Cipher (SASL only)**

  The cipher selected by the client. This element must be present in a Digest SASL challenge if the QOP directive specifies auth-conf.

- **Authzid (optional, SASL only)**

  The authorization ID is specified in RFC 2222 and encoded in UTF-8. If the authenticating user has sufficient privilege and the server supports it, the server uses this ID after authentication for making all accesses and access checks. If the server

does not support authorization ID, then authentication fails if the client specifies it.

**Authentication-Info Header**
The server sends an HTTP response with an Authentication-Info Header with Security Context Reference back to the client. This section describes the directives that make up an Authentication-Info Header.

AuthenticationInfo = "Authentication-Info" : auth-info

Auth-info = 1# (nextnonce=*noncevalue* | [qop=*message-qop*] | [rspauth=*response-digest*] | [cnonce=*cnonce*] | [nc=*nonce-count*])

**Noncevalue**
The nonce that the server requires the client to use for future authentication responses. This nonce is not supported by Digest SSP (wdigest.dll).

**Message-QOP (Optional)**
The quality of protection applied by the server.

**Response-digest (Optional)**
Supports mutual authentication. This authentication is provided mutually in that the server proves it knows the user's secret information and not in that the server proves its identity.

**Cnonce (Optional)**
A unique, encoded value generated by the client for each challenge response and that is used to avoid chosen plaintext attacks, provide mutual authentication, and provide some message integrity protection. The nonce must be specified if a QOP directive is sent.

**Nonce-count (optional)**
The nonce count. The number of times the client has sent a challenge response by using the nonce. The nonce must be specified if a QOP directive is sent.

**SASL**
The Digest SASL mechanism defined in RFC 2831 describes the use of Digest Authentication as a SASL mechanism. Integrity protection on the application protocol messages is added as a new feature after the authentication exchange.

**Quality of Protection and Ciphers**
The quality of protection determines the level of security that is used when sending and receiving messages. The Digest SSP selects a level of integrity or confidentiality based on the requirements of the client and server.

When using Digest Authentication as a SASL mechanism, applications can specify that messages be kept confidential by using encryption. The server specifies the ciphers that it supports for encrypted messages, and the client selects a cipher from this list.

**Quality of Protection**
The quality of protection, identified by the QOP directive, is first specified by the server in the Digest Authentication challenge, and then confirmed by the client in the challenge response. If the client requires a quality of protection that the server does not support, the client must terminate the authentication.

The possible values for the QOP directive are described in the following table.

### QOP Options

| Value | Description |
|---|---|
| "auth" | Authentication only. |
| "auth-int" | Authentication and integrity checking by using signatures. |
| "auth-conf" | (SASL only) Authentication, integrity and confidentiality checking by using signatures and encryption. |

**Ciphers**
The following material is valid only when Digest is used as a SASL mechanism. Ciphers and encryption are not available for HTTP authentication by using Digest Authentication.

A Digest SASL challenge's cipher directive contains a list of the ciphers supported by a server that requires confidential communications. The cipher directive can only be included if the QOP directive is set to "auth-conf." The ciphers recognized by Digest Authentication are listed in the following table, from strongest to weakest.

**Ciphers That Are Supported by Digest Authentication**

| Cipher Value | Description |
|---|---|
| “rc4” | The RC4 cipher with a 128-bit key. |
| “3des” | The triple DES cipher in CBC mode with EDE with the same key for each E stage (two keys mode). Total key length is 112 bits. |
| “rc4-56” | The RC4 cipher with a 56-bit key. |
| ““rc4-40““ | The RC4 cipher with a 40-bit key. |
| “des” | The Data Encryption Standard (DES) cipher in cipher block chaining (CBC) mode with a 56-bit key. |

When a server application requests confidentiality (qop=auth-conf) the Digest SSP generates a challenge that offers the client all of the ciphers that are installed on the server and supported by Digest Authentication. The Digest client selects the strongest available cipher.

**Password Hashes**

The use of precalculated hashes is a new feature that allows system operators to skip the use of reversible encrypted passwords on the domain controller. A precalculated hash is formed for a user when the user’s password is changed.

Active Directory stores the following password-related materials:

- NTLM hash created by the NTLM SSP
- Kerberos key created by the Kerberos SSP
- Digest hashes created by the Digest SSP

When account passwords are modified (during either account creation or password reset/change), Security Account Manager (SAM) requests the authentication packages so that it can provide updated credential information based on the new password.

When passwords are changed, credential update notifications are sent to the authentication packages that have registered for credential updates.

**Digest Hashes**

Digest Authentication verifies passwords by computing a hash of the account name, realm property, and password. The specific format is H(username:realm:password) where H() is the MD5 hash operation. As this form shows, the plaintext password is required to compute the MD5 hash. To use Digest Authentication in previous versions of Windows Server, reversible encryption has to be set on an account. This setting allows the security provider to compute the digest hash when it is authenticating a server request by using Digest Authentication. In Windows Server 2003, a new digest hash is created and stored in Active Directory when a password is modified on an account. This digest hash contains several pre-computed versions of the H(username:realm:password) hash and permits these hashes to be authenticated without requiring reversible encryption to be enabled on an account.

The specific, precomputed forms are:

**Digest Hashes**

| Username | Realm | Password |
|---|---|---|
| SamAccountName | FlatDomainName for Account’s Domain | password |
| SamAccountName | DNSDomainName for Account’s Domain | password |
| User Principal Name (UPN) | NULL string | password |
| NetBiosName | NULL string | password |
| SamAccountName | “Digest” | password |

| | | |
|---|---|---|
| User Principal Name (UPN) | "Digest" | password |
| NetBiosName | "Digest" | password |

Where NetBiosName is FlatDomainName\SamAccountName.

Because the SamAccountName, UPN, and NetBiosName can have mixed cases, additional hashes are computed with these in lower case, upper case, and the original case used when account created.

**Note**

- Because only three versions of the user name are hashed, other variants of the user name do not work. For example, user, USER, and User might be created, but if presented with UsEr authentication would fail.

During the authentication of a request from a server, the challenge response sent from the client is compared to the pre-computed hashes stored on the user account. If a match is found, the server is notified by the domain controller that the challenge response is valid. If no match is found, a check is made by the domain controller to see if Reversible Encryption is enabled for the user account; if it is enabled, the domain controller computes the Digest hash with the plaintext version of the account's password, and a final validity test is performed by the domain controller on the challenge response.

**Network Ports Used by Digest Authentication**

Digest Authentication is not a network protocol. It uses the application protocol, so the port that Digest Authentication challenges and responses are sent by varies based on which application requests Digest Authentication and what port that application uses to communicate.

**Related Information**

The following resources contain additional information that is relevant to this section.

- RFC 2616 in the IETF RFC Database
- RFC 2617 in the IETF RFC Database
- RFC 2743 in the IETF RFC Database
- RFC 2744 in the IETF RFC Database
- RFC 2831 in the IETF RFC Database

# Digest Authentication Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Digest Authentication Tools and Settings**

**In this section**


- Digest Authentication Tools
- Digest Authentication Registry Entries
- Digest Group Policy Settings
- Network Ports Used by Digest Authentication


**Digest Authentication Tools**

You can use the following tools to administer and troubleshoot Digest Authentication.

**Dsa.msc: Active Directory Users and Computers**

**Category**

Active Directory Users and Computers is a Microsoft Management Console (MMC) snap-in that automatically installs when you install Active Directory. This tool also ships with the Administration Tools Pack (Adminpak.msi).

You can access the tool from the **Start** menu: Click **Start**, then click **Programs**,then click **Administrative Tools,** and then click **Active Directory Users and Computers**.

**Version compatibility**

Active Directory Users and Computers runs on domain controllers that are running the Windows Server 2003 or Windows 2000 operating systems. In both of these operating systems, MMC provides a window in the user interface (UI) where you can add, configure, and control items. Active Directory Users and Computers is the snap-in that you can use to administer and publish information in Active Directory.

The Windows Server 2003 version of Active Directory Users and Computers can target domain controllers that are running Windows Server 2003 or Windows 2000 Server.

On administrative workstations that are running Windows XP Professional or Windows 2000, you can install the Windows Server 2003 Administration Tools Pack (Adminpak.msi) from the i386 directory on the Windows Server 2003 CD. This version of the Administration Tools Pack encrypts and signs Lightweight Directory Access Protocol (LDAP) traffic between the administrative tool clients and domain controllers.

**Note**

- You cannot run the Windows Server 2003 Administration Tools Pack (Adminpak.msi) on a computer that is running Windows XP Professional, Windows XP Home Edition, or Windows XP 64-Bit Edition Version 2003 without Windows XP Service Pack 1 (SP1).

You can use Active Directory Users and Computers to manage the properties that are listed in the following table, which are associated with objects in Active Directory. Your management of the properties affects the digest hashes for these objects.

**Active Directory Users and Computers Object Management**

| Property | Changes That Affect Digest Authentication |
|---|---|
| **User Objects Account Tab:**<br><br>Store passwords by using reversible encryption | An optional account property that can be set to keep an encrypted version of the user account's password on the domain controller. It is encrypted so that only the LSA can decrypt to access the Plaintext password for the account. |

You can find more information about Active Directory Users and Computers on Microsoft TechNet.

**Eventvwr.msc: Event Viewer**
**Category**
Event Viewer is included in Windows Server?2003, Windows XP, and Windows 2000.

**Version compatibility**
Event Viewer is supported by Windows Server 2003, Windows XP, and Windows 2000.

All account logons that occur on computers that are using Digest Authentication might be present in the security log. To find these events you need to enable auditing of account logon and logon events for user authentication. The following table lists event IDs and information potentially associated with Digest Authentication. Only relevant event information is present in the event log.

### Digest Authentication Security Log Events

| Event ID | Event Type | Description |
|---|---|---|
| 537 | Logon/Logoff<br>• Failure Audit | Logon Failure:<br>• Reason: An error occurred during logon<br>• User Name:<br>• Domain:<br>• Logon Type:<br>• Logon Process: WDIGEST<br>• Authentication Package: WDigest<br>• Workstation Name: |

| | | |
|---|---|---|
| 540 | Logon/Logoff<br>• Success Audit | Successful Network Logon:<br>• User Name:<br>• Domain:<br>• Logon ID:<br>• Logon Type:<br>• Logon Process: WDIGEST<br>• Authentication Package: WDigest<br>• Workstation Name:<br>• Logon GUID:<br>• Caller User Name:<br>• Caller Domain:<br>• Caller Logon ID:<br>• Caller Process ID:<br>• Transited Services:<br>• Source Network Address:<br>• Source Port: |
| 576 | Logon/Logoff<br>• Success Audit | Special privileges assigned to new logon:<br>• User Name:<br>• Domain:<br>• Logon ID:<br>• Privileges: |
| 680 | Account Logon<br>• Success Audit<br>• Failure Audit | Logon attempt by: WDigest<br>• Logon account:<br>• Source Workstation:<br>• Error Code: |

For more information about "Event Viewer", see "Event Viewer" on Microsoft TechNet.

**Netmon.exe: Network Monitor**
**Category**
A limited version of Network Monitor is included in Windows Server 2003, Windows XP, and Windows 2000. The full version

of Network Monitor is included with Microsoft Systems Management Server.

**Version compatibility**
Network Monitor is supported for Windows Server 2003, Windows XP, and Windows 2000.

Network Monitor enables you to capture network traces, which can be used in troubleshooting most network issues.

**Digest Authentication Registry Entries**
The following registry entries are associated with Digest Authentication:

- ClientCompat
- Negotiate
- ServerCompat
- UTF8HTTP
- UTF8SASL

The information here is provided as a reference for you to use when troubleshooting or verifying that the required settings are applied. We recommend that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. Modifications to the registry can result in irrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**ClientCompat**
**Registry path**
HKLM System\CurrentControlSet\Control\SecurityProviders\WDigest\

**Version**
Windows Server 2003, Windows XP

This entry controls the client option to quote the QOP directive because some servers require the client to quote the QOP directive. This entry does not exist in the registry by default. The default value is 1. The client will quote the QOP directive value in the challenge response that is sent over to the server. This behavior can be turned off by setting this value to 0.

**Negotiate**
**Registry path**
HKLM System\CurrentControlSet\Control\SecurityProviders\WDigest\

**Version**
Windows Server 2003, Windows XP

This entry controls adding Digest Authentication to the list of negotiable SSPs. The default value is 0.

**ServerCompat**
**Registry path**
HKLM System\CurrentControlSet\Control\SecurityProviders\WDigest\

**Version**
Windows Server 2003, Windows XP

This entry controls server to accept client challenge responses that do not encode backslash. This entry does not exist in the registry by default. The default value is 1. If authentication with backslash encoding fails, Digest SSP attempts to authenticate the response and assumes that the backslash is part of the string. This behavior can be turned off by setting this value to 0.

**UTF8HTTP**
**Registry path**
HKLM System\CurrentControlSet\Control\SecurityProviders\WDigest\

**Version**
Windows Server 2003, Windows XP

This entry controls UTF-8 encoding for HTTP mode. ISO 8859-1 Latin is used but does not support non US languages. This is a

port of SASL option to allow use of UTF-8 encoding to the HTTP mode. The default value is 1.

**UTF8SASL**
**Registry path**
HKLM System\CurrentControlSet\Control\SecurityProviders\WDigest\

**Version**
Windows Server 2003, Windows XP

This entry controls UTF-8 encoding for SASL mode. The default value is 1.

**Digest Authentication Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with Digest Authentication.

### Group Policy Settings Associated with Digest Authentication

| Group Policy Setting | Description |
|---|---|
| Account Policies:<br><br>Password Policy: Store passwords by using reversible encryption for all users in the domain | An optional account property that can be set to keep an encrypted version of the user account's password on the domain controller. It is encrypted so that only the LSA can decrypt to access the Plaintext password for the account. |

For more information about Group Policy settings, see the "Group Policy Settings Reference for Windows Server 2003" in Tools and Settings Collection.

**Network Ports Used by Digest Authentication**
Digest Authentication is not a network protocol. It uses the application protocol, so the port that Digest Authentication challenges and responses are sent by varies based on which application requests Digest Authentication and what port that application uses to communicate.

# Interactive Logon Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive Logon Technical Reference**
The Microsoft Windows Server 2003 operating system requires all users to log on to access local and network resources. Windows secures resources by implementing the interactive logon process, which provides authentication of users.

**In this subject**


- What is Interactive Logon?
- How Interactive Logon Works
- Interactive Logon Tools and Settings

# What is Interactive Logon?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Interactive Logon Scenarios
- Related Information


Windows Server 2003 requires that all users must validate their identities to successfully log on to the computer. The process of validating a user's identity is called authentication.

Security in the Windows Server 2003 operating system controls the use of local and network resources through the interrelated mechanisms of authentication and authorization. After a user is authenticated, Windows Server 2003 uses the authorization and access control technologies to implement the second phase of protecting resources: determining if an authenticated user is authorized to access a resource.

Windows Server 2003 can authenticate users and computers that have accounts either in the Security Accounts Manager (SAM) or the Active Directory directory service.

**Interactive Logon Scenarios**

An interactive logon to a computer can be performed either locally, when the user has direct physical access, or remotely, through Terminal Services, in which case the logon is further qualified as remote interactive. After an interactive logon, Windows runs applications on the user's behalf and the user can interact with those applications.

Users can perform an interactive logon by logging on to a computer using a local user account or by using their domain account.

The following figure shows local and domain logons.

**Interactive Logon Types**




**Local Logon**

A local logon requires that the user have a user account in the SAM on the local computer. The SAM protects and manages user and group information in the form of security accounts stored in the local computer registry. The computer can have network access, but it is not required. Local user account and group membership information is used to manage access to local resources.

**Domain Logon**

A domain logon requires that the user have a user account in the domain's Active Directory. The computer must be joined to the domain and have a network connection to the domain. Users must also have rights to log on to a local computer or a domain. Domain user account and group membership information is used to manage access to domain and local resources.

**Smart Card Logon**

Windows Server 2003 can also be configured to support special domain logon processes that require the use of a hardware token, such as a smart card. A smart card logon is a form of enhanced-security logon that requires both a physical identifier (the smart card) and a personal identification number (PIN).

**Related Information**

The following resources contain additional information that is relevant to this section.

- Logon and Authentication Technologies
- Authorization and Access Control Technologies
- Kerberos Authentication Technical Reference

# How Interactive Logon Works

Updated: March 31, 2015

Applies To: Windows 7, Windows 8, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Vista

**In this section**


- Interactive Logon Architecture
- Interactive Logon Components
- Interactive Logon Protocols
- Interactive Logon Processes and Interactions
- Network Ports Used by Interactive Logon
- Related Information


The interactive logon process is the first step in user authentication and authorization. Interactive logon is mandatory in the Windows Server 2003, Windows XP, Microsoft Windows 2000, and Windows NT 4.0 operating systems. Interactive logon provides a way to identify authorized users and determine whether they are allowed to log on and access the system.

This section describes the interactive logon architecture, the process of an interactive logon, and the way in which the Windows Server 2003 interactive logon process locks the workstation and helps to keep it secure.

**Interactive Logon Architecture**

The Windows Server 2003 interactive logon architecture includes the following components:

- Winlogon
- Graphical Identification and Authentication (GINA) dynamic-link library (DLL)
- Local Security Authority (LSA)
- Authentication packages (NTLM and Kerberos)

For more information about these components, see Interactive Logon Components later in this section.

Windows Server 2003 interactive logons begin with the user pressing CTRL+ALT+DEL to initiate the logon process. The CTRL+ALT+DEL keystroke is called a secure attention sequence (SAS); Winlogon registers this sequence during the boot process to keep other programs and processes from using it. The GINA DLL generates the logon dialog box. The following figure shows the Windows Server 2003 logon dialog box.

**Windows Server 2003 Logon Dialog Box**

A user who logs on to a computer using either a local or domain account must enter a user name and password, which form the user's credentials and are used to verify the user's identity. For smart card logons, a user's credentials are contained on the card's security chip, which is read by an external device, a smart card reader. During a smart card logon, a user enters a personal identification number (PIN) instead of a user name, domain, and password.

**Local Logon**

Local logons give users access to local computer applications and resources but not to domain applications and resources. When users log on locally, their identities are validated by authentication packages to local account information stored in the Security Accounts Manager (SAM) database. The SAM operates in the security context of the LSA; it protects and manages user and group information in the form of security accounts stored in the local computer registry. Because user accounts are stored on the local computer, network access is not required for local logons. However, if a computer has a network connection and a user logs on to a local account, there is no interaction with the network.

Local logons can be performed on Windows client operating systems, such as Windows XP and Windows 2000 Professional. Windows server operating systems, such as Windows NT Server, Windows 2000 Server, and Windows Server 2003, also permit local logons.

The following figure shows the local logon architecture.

**Interactive Local Logon**




A successful local logon begins when a user presses CTRL+ALT+DEL. Winlogon and the GINA DLL collect the user's credentials and then send the credentials to the LSA. The LSA verifies the user's identity and then returns a logon success and the user's access token to Winlogon and the GINA DLL. Winlogon and the GINA DLL then activate the user's shell by creating a new process, such as Explorer.exe.

**Domain Logon**

Domain logons give users access to resources throughout the domain. Domain user accounts are stored in an Active Directory domain. Active Directory is deployed on each domain controller, and domain user accounts are replicated throughout a domain.

Before a user can log on to a computer by using a domain account, the computer must be joined to a domain. If the computer has access to a network connection, the user can log on to a domain if the user has an account in the domain's Active Directory.

The computer must transparently authenticate to the domain's Active Directory. This form of logon is called a computer logon. Both users and computers are considered equal security principals in Active Directory; to be granted access to network

resources, both must be able to verify their identities.

Users can use a domain account to log on to Windows client operating systems, such as Windows 8. Windows server operating systems, such as Windows Server 2012 R2, also permit domain logons. Only server operating systems can function as domain controllers and deploy Active Directory.

On a domain-joined computer, Windows is hard-coded to show only the last logged on user or **Other user**. Additional tiles for other users to log on are available only for computers joined to a workgroup.

The following figure shows the domain logon architecture.

**Interactive Domain Logon**




Unlike a local logon in which the local LSA validates the user, during a domain logon, the LSA on a domain controller validates the user. The LSA evaluates the user's credentials to determine if the logon should be processed as a logon to a local account or a logon to a domain account. After determining the logon type, either the NTLM or Kerberos authentication package validates the user. If the authenticating domain controller is a computer running Windows 2000 or Windows Server 2003, the LSA will use Kerberos, the default authentication package for domain and network logons. The LSA uses NTLM to process domain logons in Windows NT 4.0 mixed environments.

**Single Sign On**

Single sign-on (SSO) uses credentials collected during an interactive domain logon to allow the user to authenticate to a network one time and, thereafter, to have access to all authorized network resources without additional authentication. The network resources can range from hardware devices, such as printers, to applications, files, and other types of data, all of which may be located throughout an enterprise on servers of various types, possibly in different domains and running different operating systems.

**Cached Credentials**

After a successful domain logon, information is cached; this means that later a user can log on to the computer with the domain account even if the domain controller that authenticated the user is not available. Because the user has already been authenticated, Windows uses the cached credentials to log the user on locally. For example, if a mobile user logs on to a portable computer that is a domain member with a domain account and then takes the portable computer to a location where the domain is unavailable, Windows will attempt to use the cached credentials from the last successful logon with a domain account to locally log on the user and allocate access to local computer resources.

**Interactive Logon Components**

The components of the interactive logon process are responsible for helping to establish secure user authentication. This section describes the roles and functions of the user logon components.

**Winlogon**

Winlogon (%windir%\System32\Winlogon.exe) is the executable file responsible for managing secure user interactions. Winlogon initiates the logon process for Windows Server 2003, Windows 2000, Windows XP, and Windows NT 4.0.

Winlogon is responsible for the following tasks.

- Desktop lockdown
- Standard SAS recognition
- SAS routine dispatching
- User profile loading
- Screen saver control
- Multiple network provider support

**Desktop Lockdown**

Winlogon helps prevent unauthorized users from gaining access to system resources by locking the computer desktop. At any time, Winlogon is in one of three possible states: logged on, logged off, or workstation locked, as shown in the following figure.

**Winlogon States**




Winlogon switches between three desktops—Winlogon, screen saver, and default—depending on its state and user activity. The following table lists and describes each of these desktops.

**Winlogon Desktops**

| Desktop | Description |
|---|---|
| Winlogon | Also called the secure desktop. There are five dialog boxes in the Winlogon desktop, which are described in the following table. |
| Screen saver | Used when a screen saver is running. |
| Default | Displayed when a user has successfully logged on to Windows Server 2003 and the user's shell has been activated. The user's applications run in this desktop. |

The following table lists and describes the five dialog boxes in the Winlogon desktop.

**Winlogon Desktop Dialog Boxes**

| Desktop | Description |
|---|---|
| Welcome to Windows | Displayed when a computer has initially booted and no user has entered a SAS or when a user has logged off. |

| | |
|---|---|
| Log On to Windows | Displayed after a user presses CTRL+ALT+DEL when no user is currently logged on. |
| Windows Security and Change Password | Displayed when:<br>• A logged-on user presses CTRL+ALT+DEL when the workstation is not locked.<br>OR<br>• When the workstation is not locked, after a user presses CTRL+ALT+DEL and clicks **Change Password**. |
| Computer Locked | Displayed when the workstation is locked either explicitly by the user or when a screen saver is set as secure and is dismissed. |
| Unlock Computer | Displayed when a user presses CTRL+ALT+DEL when the workstation is locked. |

By switching between the desktops for different logon tasks, Winlogon ensures that its processes will be invisible to applications not involved in the logon process, preventing logon data from being compromised.

**Standard SAS Recognition**
Winlogon registers the default SAS during the boot process, which ensures that no other program can access the same key sequence and compromise the logon process by impersonating Winlogon and its logon dialog boxes.

**SAS Routine Dispatching**
When Winlogon recognizes a SAS event or the GINA delivers a SAS, Winlogon calls one of the SAS processing functions of the GINA.

**User Profile Loading**
After a successful logon, Winlogon loads user profiles into the HKEY_CURRENT_USER registry key.

**Screen Saver Control**
Winlogon monitors keyboard and mouse activity to determine when to activate screen savers. When a screen saver that is set as secure is dismissed, Winlogon treats the workstation as locked and the GINA displays the **Computer Locked** dialog box. When a screen saver not set as secure is dismissed, the user is able to access the application desktop without being prompted to reenter credentials.

**Multiple Network Provider Support**
If there are multiple network providers installed on a Windows-based system, they can be included in the authentication process and in password-updating operations. This inclusion lets additional network providers gather identification and authentication information during a logon, using the secure desktop.

**GINA**
The GINA is a DLL module that operates in the security context of Winlogon. Winlogon loads the GINA early in the boot process. The GINA is responsible for processing SAS events and activating the user's shell. The GINA can generate SAS events in some instances.

In Windows Server 2003, Windows XP, and Windows 2000, Msgina.dll (%windir%\System32\Msgina.dll) is the default GINA; it can be replaced to support specific and unique authentication methods. GINA customization is enabled to accommodate the use of authentication hardware tokens, such as retinal scanners and proprietary smart card solutions.

**LSA**
The LSA (%windir%\System32\lsass.exe) is a protected security subsystem that helps create secure user interactions in Windows Server 2003. Winlogon and the GINA call the LSA to process logon credentials.

The following components involved in user logon run in the security context of the LSA:

- Authentication packages
- SAM
- Active Directory

## LSA Responsibilities

| Role | Description |
| --- | --- |
| Application programming interfaces (APIs) | The LSA provides APIs for:<br>• Packages to create a user's logon token after authentication.<br>• Applications to call directly for user logon.<br>• Microsoft Security Support Provider Interface (SSPI) requests and the routing of requests to the appropriate authentication package.<br>• Applications to enumerate available authentication packages and query package capabilities.<br>• Applications to resolve Security Identifier-to-Name and Name-to-Security Identifier.<br>• Applications to store secret information in the system. |
| Logon session state | Manages the state of the logon session for any user authenticated by the authentication packages. |
| Notification to installed authentication packages | Notifies installed authentication packages of user logon, logoff, and credential updates. |
| Authentication packages | Loads and unloads authentication packages. |
| Encryption assistance | Includes subcomponents that provide encryption services to users of the local computer. |

The following figure shows the LSA architecture. Winlogon, which runs as a separate process from lsass.exe, is shown here to illustrate the way in which the logon process interacts with the LSA.

**LSA Architecture**




| Note |
| --- |
| Some components shown in this figure run only on domain controllers or non-domain controllers. The SAM runs only on non-domain controllers and Windows NT computers; Directory Services runs only on domain controllers. |

As the figure shows, all calls from Winlogon are sent by using the Secur32.dll LsaLogonUser call. LsaLogonUser returns the call through local procedure call (LPC) to the LSA Server service. The LSA Server service then determines which authentication package to use to process the logon.

You can set user rights for local computers by using the Local Security Policy snap-in of the Microsoft Management Console (MMC). On Windows Server 2003 non-domain controllers, the Group Policy User Rights Assignment contains information about which users are authorized to perform different tasks, including logging on to the system locally. Domain controllers use the Domain Security Policy snap-in of the MMC to set user rights for domain computers. The User Rights Assignment on domain controllers contains information that applies to the entire domain.

**Authentication Packages**

Authentication packages on the user's local computer communicate with server authentication packages to authenticate users. Windows Server 2003 applies default authentication packages for user authentication and also supports custom authentication packages.

The following table lists the Windows Server 2003 authentication packages used for interactive logons.

**Interactive Logon Authentication Packages**

| Name | Associated Protocol | Environment |
|---|---|---|
| Kerberos version 5 (V5) | Kerberos.dll | Windows 2000, Windows XP, and Windows Server 2003 |
| NTLM | MSV1_0.dll | Windows NT 4.0 and mixed environments |

**SAM**

The SAM stores information about local user accounts in the Windows registry. Passwords are encrypted by the NTLM authentication package. The outcome of the encryption is a hashed password transformed into ciphertext, a string of numbers and letters that appears meaningless. The hashing process occurs by means of a hashing algorithm. NTLM uses the same algorithm to encrypt and decrypt a user's password.

**Active Directory**

Domain account information is stored in Active Directory on the domain controllers. Within a particular domain, each domain controller has an identical copy of the domain's Active Directory. Users can therefore authenticate to any domain controller and their logon information will be recognized throughout the domain.

**Interactive Logon Protocols**

In network environments, authentication protocols are used to authenticate users, computers, and groups. For local logons, Windows Server 2003 supports the NTLM authentication protocol. For domain logons, Windows Server 2003 supports the Kerberos V5 and NTLM authentication protocols. Windows Server 2003 also supports custom authentication packages.

**Logon Protocols**

| Component | Description |
|---|---|
| Kerberos.dll | • Kerberos V5 authentication package authentication protocol.<br>• Standard Internet protocol and default network authentication protocol in Windows Server 2003.<br>• Performs authentication of users and computers in Windows 2000, Windows XP, and Windows Server 2003. |
| MSV1_0.dll | • NTLM authentication package authentication protocol.<br>• Default network authentication protocol in Windows NT 4.0.<br>• Performs authentication of users and computers in Windows NT 4.0 mixed environments.<br>• Included in Windows Server 2003 for compatibility with Windows NT 4.0 mixed environments. |

For more information about how the Kerberos V5 protocol implements user authentication, see Kerberos Authentication Technical Reference.

**Interactive Logon Processes and Interactions**
This section describes the events that occur during an interactive logon.

**Local Logon**
The following figure shows the local logon process.

### Local Logon Process




The GINA specifies the Negotiate authentication package when it calls into the LSA. Negotiate must then choose an authentication package to process the logon. Negotiate sends the credentials to Kerberos, the default authentication package beginning with Windows Server 2003. However, the Kerberos authentication package cannot process local logons, so it returns an error to Negotiate. Negotiate then calls NTLM to authenticate the user by comparing the received credentials with those hashed in the SAM.

If the credentials are valid, the LSA generates an access token for the user based upon user rights assigned to the user's account and LsaLogonUser returns the logon success and the user's access token to Winlogon and the GINA. The GINA then activates the user's shell and Winlogon switches to the default desktop. If the credentials are invalid, the LSA returns a logon failure, the GINA displays an error message and prompts the user to present valid credentials, and Winlogon remains in the Winlogon desktop.

**Domain Logon**
Domain logons can only be performed from computers that are joined to a domain. Domain credentials consist of a user's domain account user name, password, and the name of the domain. The local computer's LSA chooses the appropriate authentication package to use based on the domain's environment.

The following figure shows the process that occurs when the local computer can reach a domain controller to authenticate the user. If a domain controller is not available, a cached logon occurs.

### Domain Logon Process




As with the local logon process, the Negotiate authentication package routes the authentication request to the default authentication package, Kerberos.

The domain client chooses Kerberos, and Kerberos validates the user's credentials by contacting the domain controller. The LSA on the domain controller returns the logon success or failure to the local computer's LSA. If the domain logon succeeds, the local LSA generates an access token for the user based upon user rights assigned to the user's account, and LsaLogonUser returns the logon success and the user's access token to Winlogon and the GINA. The GINA then activates the user's shell, and Winlogon switches to the default desktop. If the credentials are invalid, the LSA returns a logon failure, the GINA displays an error message and prompts the user to present valid credentials, and Winlogon remains in the Winlogon desktop.

| **Note** |
| --- |
| In Windows NT 4.0 mixed environments, the NTLM authentication package validates the user's credentials by contacting the domain controller. Windows NT 4.0 stores user accounts in the SAM. |

**Smart Card Logons**

Windows Server 2003 can be configured to support smart card logons. All smart card logons are processed by the Kerberos authentication package. During a smart card logon, the local computer checks to ensure the user's PIN is valid. The domain controller then authenticates the user based on the user's credentials in the smart card.

**Cached Logons**

Windows Server 2003 supports cached logons. The cached credentials of the last 10 users who have successfully logged on to a domain account can be used to log a user on locally if the authenticating domain controller becomes unavailable.

**Network Ports Used by Interactive Logon**

Because the logon process can be deployed across various network boundaries, it can span one or more firewalls. The following table lists the three main configurable ports used by interactive logon.

**Port Assignments for Interactive Logon**

| Service Name | UDP | TCP |
| --- | --- | --- |
| LSA RPC port | Dynamic RPC | Dynamic RPC |

| | | |
|---|---|---|
| Kerberos V5 port | 88 | 88 |
| NTLM port | Dynamic | Dynamic |

You can use the Registry Editor to modify the registry to apply fixed ports for NTLM and the LSA.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Kerberos Authentication Technical Reference
- Authorization and Access Control Technologies

# Interactive Logon Tools and Settings

Updated: November 10, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive Logon Tools**
The following tools are associated with interactive logon.

**EventViewer.msc: Event Viewer**
**Category**

Event Viewer is included as part of the Windows Server 2003, Windows 2000, and Windows XP operating systems.

**Version compatibility**

Event Viewer is supported for Windows Server 2003, Windows XP, and Windows 2000.

The Event Log service starts automatically when you start Windows. The Event Log is a compilation of system events. There are three Event Logs: an application log, a system log, and a security log. All users can view application and system logs. Only administrators can gain access to security logs.

By default, security logging is turned off. You can use Group Policy to enable security logging. The administrator can also set auditing policies in the registry that cause the system to halt when the security log is full.

The system and security logs will contain logon error codes and other events related to interactive logon and network authentication.

The following table lists the event IDs that relate to interactive logon.

**Interactive Logon Events**

| Event ID | Description |
|---|---|
| 528 | Successful logon. |
| 529 | Logon failure. Unknown user name or bad password. |
| 530 | Logon failure. Account logon time restriction violation. |
| 531 | Logon failure. The account is currently disabled. |
| 532 | Logon failure. The specified user account has expired. |
| 533 | Logon failure. The user is not allowed to log on at this computer. |
| 534 | Logon failure. The user has not been granted the requested logon type at this computer. |
| 535 | Logon failure. The specified account's password has expired. |
| 536 | Logon failure. The NetLogon component is not active. |
| 537 | Logon failure. An unexpected error occurred during logon. |
| 538 | User logoff. This event is generated when the logoff process is complete. A logoff is considered complete when the associated logon session object is deleted, which occurs after all tokens associated with the logon session are closed. This can take an arbitrarily long time; this event should not be used to calculate the total logon duration. Instead, use event 551. |
| 539 | Logon failure. Account locked out. |
| 540 | Successful network logon. |
| 541 | IPSec security association established. |

| | |
|---|---|
| 542 | IPSec security association ended. Mode: Data Protection (Quick mode). |
| 543 | IPSec security association ended. Mode: Key Exchange (Main mode). |
| 544 | IPSec security association establishment failed because peer could not authenticate. The certificate trust could not be established. |
| 545 | IPSec peer authentication failed. |
| 546 | IPSec security association establishment failed because peer sent invalid proposal. |
| 547 | IPSec security association negotiation failed. |
| 548 | Logon failure. Domain security identifier (SID) is inconsistent. This event is generated when a user account from a trusted domain attempts to authenticate, but the domain SID does not match the SID stored in the Trusted Domain Object (TDO). |
| 549 | Logon failure. All SIDs were filtered out. During authentication across forests, SIDs corresponding to untrusted namespaces are filtered out. This event is generated when all SIDs are filtered. This event is generated on the Kerberos Key Distribution Center (KDC).<br><br>This event is not generated on Windows Server 2003. |
| 550 | Notification message that can indicate a possible denial-of-service attack. |
| 551 | User-initiated logoff. This event is generated when the user initiates the logoff process. When the logoff process is complete, event 538 is logged. |
| 552 | Successful logon. This event is generated when a user logs on with explicit credentials while already logged on as another user. This event is logged when using the RunAs tool. |
| 553 | Logon failure. This event is generated when an authentication package detects a replay attack. |

**Lusrmgr.msc: Local Users and Groups**

### Category

Local Users and Groups is included as part of the Windows Server 2003, Windows XP Professional, and Windows 2000 operating systems and is available in the Windows Server 2003 Administration Tools Pack (Adminpak.msi).

### Version compatibility

Local Users and Groups is supported for administering local groups on Windows Server 2003, Windows XP Professional, and Windows 2000.

Local Users and Groups manages users and groups of users for your computer. You can create new users and groups, add users to groups, remove users from groups, disable user and group accounts, and reset passwords.

You might need to be logged on as an administrator or a member of the Administrators group to perform some tasks.

**Secpol.msc: Local Security Policy**

### Category

Local Security Policy is included as part of the Windows Server 2003, Windows XP, and Windows 2000 operating system tools and is available in the Windows Server 2003 Administration Tools Pack.

### Version compatibility

Local Security Policy is supported for Windows Server 2003 and Windows 2000.

You can set security settings for the local computer in Local Security Policy.

For more information about security policy settings specific to interactive logons, see the “Interactive Logon Group Policy Settings” section later in this topic.

**Whoami**

### Category

Whoami is a command-line tool in Windows Server 2003. This tool is also included in the Windows 2000 Resource Kit.

### Version compatibility

Whoami is supported for Windows Server 2003 and Windows 2000.

You can use whoami to display the complete contents of a user's access token in the command window.

Whoami displays:

- User name and SID.
- Groups and their SIDs.
- Privileges and their status (for example, enabled or disabled).
- Logon ID.

For more information about whoami, see "Command Line References" in Tools and Settings Collection.

**Interactive Logon Registry Entries**
The following registry entries are associated with interactive logon.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_LOCAL_MACHINE\SYSTEM**
The following registry entries are located under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\.

**Bounds**
Registry Path

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

Version

The Bounds registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The Bounds entry specifies thresholds for managing the length of the kernel-mode Local Security Authority (LSA) audit queue. The audit queue stores kernel-mode events destined for the security log in Event Viewer.

The value of this entry is an 8-byte binary field. The value of the first four bytes specifies the maximum number of items that can be held in the audit queue (the upper bound). When the number of audits exceeds this value, the LSA discards all new audits until the number of audits remaining in the queue reaches the lower bound, as specified by the value of the last four bytes.

**CrashOnAuditFall**
### Registry Path

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

### Version

The CrashOnAuditFall registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The CrashOnAuditFall entry directs the system to halt when it cannot record new events in the security log in Event Viewer. This feature prevents unauthorized activities from occurring when they cannot be recorded in the security log.

The system also uses this entry to indicate that this feature has been triggered (a value of 2). When the value of this entry is 2, only members of the Administrators group can log on to the computer. This restricted state lets an Administrator log on to resolve the problem and to reset the value of this entry to 1.

Typically, the system cannot record security events because the security log in Event Viewer is full or because the internal queue to the log has reached the maximum established by the Bounds value.

This entry does not exist in the registry by default. You can add it by using the registry editor, Regedit.exe.

**Everyoneincludesanonymous**

### Registry Path

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

### Version

The Everyoneincludesanonymous registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The Everyoneincludesanonymous entry specifies whether the local Everyonegroup of users includes anonymous users. (The Everyone group is a generic group of users that includes all users.) Specifically, this entry determines whether the LSA includes the SID of the Everyone group in the security token of an anonymous user.

**Kerberos**

### Registry Path

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

### Version

The Kerberos registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The Kerberos subkey stores configuration data for the Kerberos version 5 (V5) authentication protocol in Windows Server 2003. The Kerberos V5 protocol is the primary security protocol for authentication within a domain. It verifies the identity of the user and the system.

**MSV1_0**

### Registry Path

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

### Version

The MSV1_0 registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The MSV1_0 subkey stores information about the computer’s security settings.

**NoLMHash**

### Registry Path

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

### Version

The NoLMHash registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The NoLMHash value specifies whether Security Accounts Manager (SAM) stores the LAN Manager (LM) hash of the user’s password. The LM hash of the user's password is required to authenticate down-level clients that cannot use NTLM or NTLM version 2 (NTLMv2) authentication.

The following table lists the possible settings for the NoLMHash value.

### NoLMHash Settings

| Value | Meaning |
|---|---|
| 0 | SAM stores the message digest of the user's password. |
| 1 | SAM does not store the message digest of the user’s password. |

**NotificationPackages**

### Registry Path

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

### Version

The NotificationPackages registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The NotificationPackages value specifies the dynamic-link libraries (DLLs) that are loaded or called when passwords are set or changed. This entry is often used to load libraries for services such as File and Print Services for NetWare.

You can use this entry to implement custom DLLs that specify alternate criteria for valid passwords, such as complex password rules. Notification packages that implement password verification are recommended for all primary and backup domain controllers (relevant for Windows NT 4.0 environments). You can also write a custom password filter to fulfill special security requirements.

**RestrictAnonymous**

**Registry Path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

**Version**

The RestrictAnonymous registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The RestrictAnonymous value restricts anonymous users from displaying lists of users and from viewing security permissions.

The following table lists the possible settings for the RestrictAnonymous value.

**RestrictAnonymous Settings**

| Value | Meaning |
|---|---|
| 0 | Disabled. Anonymous users are not restricted. |
| 1 | Enabled. Users who log on anonymously (a logon referred to as a null session connection) cannot display lists of domain user names or share names. Also, these users cannot view security permissions; they cannot use all of the features of Windows Explorer, Local Users and Groups, and other programs that enumerate users or shares. |
| 2 | Anonymous users have no access without explicit anonymous permissions. |

This entry does not exist in the registry by default. You can add it by using the registry editor, Regedit.exe.

**HKEY_LOCAL_MACHINE\SAM**

The following registry entry is located under HKEY_LOCAL_MACHINE\SAM.

**SAM**

**Registry Path**

HKEY_LOCAL_MACHINE\

**Version**

The SAM registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The SAM key contains information used by the SAM (Security Accounts Manager). Servers that are running Windows Server 2003 and are acting as domain controllers use this key to store user and group information pertaining to the directory service restore mode.

Domain controllers maintain user and group information in Active Directory to manage user security. However, Windows NT 4.0 and earlier, as well as Windows Server 2003-based computers that are not part of a Windows Server 2003 domain, use SAM. All of the data in this key is in binary form.

| **Note** |
|---|
| The HKEY_LOCAL_MACHINE\Security\SAM subkey stores a duplicate of the information in this key. The information in the SAM subkey cannot be viewed by users of Windows XP Professional and should not be edited directly in Windows Server 2003. Do not alter or add entries in the SAM subkey. Doing so could prevent users from logging on to their computers or the domain, and it could require you to restore the entire system. |

**HKEY_CURRENT_USER\Software**

The following registry entries are located under HKEY_CURRENT_USER\Software\Windows NT\CurrentVersion\Winlogon\.

**ReportDC**

**Registry path**

HKEY_USERS\.DEFAULT\Software\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The ReportDC registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The ReportDC value stores the user setting for the domain controller message. (The computer setting is stored in the value of HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\WindowsNT\CurrentVersion\Winlogon\ReportControllerMissing.)

This entry determines, in part, whether the system displays the following message when it cannot contact the domain controller that stores a user's roaming user profile:

"A domain controller for your domain could not be contacted. You have been logged on using cached account information. Changes made to your profile since you last logged on might not be available."

This entry stores the setting of the **Do not display this message again** check box on the domain controller message dialog box. When the check box is selected, the system sets the value of this entry to 0.

The following table lists the possible settings for the ReportDC value.

**ReportDC Settings**

| Value | Meaning |
|---|---|
| 0 | Do not display the domain controller message. (The check box is selected.) |
| 1 (or not in the registry) | Display the domain controller message. (The check box is cleared.) |

Two registry entries manage the display of the domain controller message: ReportDC and ReportControllerMissing. The message is displayed only when both entries are set to their display setting, that is, when ReportDC is either set to 1 or does not appear in the registry and when ReportDomainControllerMissing appears in the registry with a value of TRUE. By default, the message is not displayed.

**RunLogonScriptSync**

**Registry Path**

HKEY_CURRENT_USER\SOFTWARE\Microsoft\WindowsNT\CurrentVersion\Winlogon\

**Version**

The RunLogonScriptSync registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The RunLogonScriptSync value specifies whether the system waits for the logon script to finish running before it starts Windows Explorer and creates the desktop.

The following table lists the possible settings for the RunLogonScriptSync value.

**RunLogonScriptSync Settings**

| Value | Meaning |
|---|---|
| 0 | The logon script and Windows Explorer can run simultaneously. |
| 1 | Windows Explorer does not start until the logon script has finished running. |

This entry can be superceded by Group Policy settings included in Windows Server 2003. When the **Run logon scripts synchronously** Group Policy setting is enabled (in the Computer Configuration or User Configuration policy folders), the system ignores this entry. The configuration of the **Run logon scripts synchronously** policy setting is stored in the values of RunLogonScriptSync in HKEY_LOCAL_MACHINE and of RunLogonScriptSync in HKEY_CURRENT_USER.

**HKEY_LOCAL_MACHINE\SOFTWARE**

The following registry entries are located under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\.

**AutoRestartShell**

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The AutoRestartShell registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The AutoRestartShell entry specifies whether the Windows user interface (typically, Explorer.exe) restarts automatically if it stops unexpectedly.

#### CachePrimary

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The CachePrimary registry setting appears in Windows Server 2003, Windows 2000, and Windows XP.

Windows Server 2003 does not use this entry. The entry remains in the registry to support programs designed for Windows NT 4.0 and earlier.

| **Note** |
| --- |
| Do not delete this entry from the registry or change its value. These changes can cause serious, unexpected results. |

#### DCacheUpdate

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The DCacheUpdate registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The DCacheUpdate entry stores domain names in binary form for internal Winlogon programming code to use.

| **Note** |
| --- |
| Winlogon stores the value of this entry for its own use. Do not delete this entry or change its value. If you do, Windows Server 2003 might not operate properly. |

#### DefaultDomainName

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The DefaultDomainName registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The DefaultDomainName entry stores the name of the domain to which the user most recently logged on successfully. The value of this entry appears in the **Log On to Windows** dialog box generated by the Graphical Identification and Authentication (GINA) DLL the next time the dialog box is displayed.

#### DefaultUserName

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The DefaultUserName registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The DefaultUserName entry stores the last user name entered in the **Log On to Windows** dialog box generated by the GINA. The value of this entry appears in the dialog box the next time it is displayed. It is also used in automated logons and unattended setup.

#### DontDisplayLastUserName

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The DontDisplayLastUserName registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The DontDisplayLastUserName entry specifies whether a user name appears in the **Log On to Windows**dialog box generated by the GINA.

By default, Windows Server 2003 displays the user name of the last user who logged on successfully (as stored in the value of DefaultUserName) in the dialog box. If the value of this entry is 1, then the **User name** box in the **Log On to Windows** dialog box is blank.

**RunLogonScriptSync**

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The RunLogonScriptSync registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The RunLogonScriptSync entry specifies whether the system waits for the logon script to finish running before it starts Windows Explorer and creates the desktop.

This entry does not exist in the registry by default. You can add it by using the registry editor, Regedit.exe.

**ScreenSaverGracePeriod**

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The ScreenSaverGracePeriod registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The ScreenSaverGracePeriod entry specifies when password protection of a screen saver becomes effective. This entry specifies the delay between the appearance of a password-protected screen saver and the enforcement of the password requirement.

Password protection of a screen saver is not effective immediately. By default, a brief period elapses within which the user can use the mouse or the keyboard to stop the screen saver without entering the password. This delay is designed to minimize the disruption that results when the screen saver starts while the user is working.

You can add this entry to the registry to adjust the length of the delay. To make password protection effective immediately, set the value of this entry to 0x0.

This entry does not exist in the registry by default. You can add it by using the registry editor, Regedit.exe.

**Shell**

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The Shell registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The Shell entry specifies the programs that provide the user interface to the operating system.

By default, Winlogon starts the programs specified in the value of Userinit, including Userinit.exe. Userinit.exe starts the user interface program. However, if Winlogon cannot start the programs specified in the value of Userinit, then Winlogon directly runs the programs specified in the value of this entry.

**ShowLogonOptions**

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The ShowLogonOptions registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The ShowLogonOptions specifies whether logon options are displayed in the **Log on to Windows** dialog box when the dialog box opens.

The **Log on to Windows** dialog box has an **Options** button that alternately hides and displays the **Domain** box and the **Log on using dial-up connection** options.

#### Welcome

**Registry Path**

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\

**Version**

The Welcome registry setting is available in Windows Server 2003, Windows 2000, and Windows XP.

The Welcome entry specifies the text that appears in the caption bar beside the title of the **Log On to Windows**, **Windows Security**, **Computer Locked**, and **Unlock Computer** dialog boxes.

This entry does not exist in the registry by default. You can add it by using the registry editor, Regedit.exe.

#### Interactive Logon Group Policy Settings

The following table lists and describes the Group Policy settings that are associated with Logon.

**Group Policy Settings Associated with Interactive Logon**

| Group Policy Setting | Description |
|---|---|
| Password Policy:<br>• Enforce password history<br>• Maximum password age<br>• Minimum password age<br>• Minimum password length<br>• Password must meet complexity requirements<br>• Store password using reversible encryption for all users in the domain | Changes to the Password Policy settings control:<br>The strength and complexity required of every user's password |
| Audit Policy:<br>• Audit account logon events<br>• Audit account management<br>• Audit logon events | Changes to the Audit Policy settings control:<br>• Auditing of logons and logoffs<br>• Auditing of password and permissions changes |

| | |
|---|---|
| User Rights Assignment:<br>• Access the computer from the network<br>• Deny logon as a batch job<br>• Deny logon as a service<br>• Deny logon locally<br>• Deny logon through terminal services<br>• Logon as a batch job<br>• Logon as a service<br>• Logon locally | Changes to the User Rights Assignment settings control:<br>• Which users are allowed or disallowed to log on to perform different tasks, including logging on as a batch job and a service<br>• Which users are allowed or disallowed to log on locally or through Terminal Services, as well as who can access the computer from the network |

| Security Options: | Changes to the Security Options settings control: |
|---|---|
| • Accounts: Limit local accounts use of blank passwords to console logon only<br>• Domain member: Maximum machine account password age<br>• Domain member: Require strong (in Windows 2000 or later) session key<br>• Interactive logon: Do not display last user name<br>• Interactive logon: Do not require CTRL+ALT+DEL<br>• Interactive logon: Message Text for users attempting to log on<br>• Interactive logon: Message title for users attempting to log on<br>• Interactive logon: Number of previous logons to cache (in case the domain controller is not available)<br>• Interactive logon: Require domain controller authentication to unlock workstation<br>• Interactive logon: Smart card removal behavior<br>• Recovery console: Allow automatic administrative logon<br>• Shutdown: Allow system to be shut down without having to log on | • Message text and title displayed by the GINA during an interactive logon<br>• Domain member settings<br>• Authentication settings, including allowing or disallowing blank passwords and password age |

For more information about Group Policy settings, see the "Group Policy Settings Reference for Windows Server 2003" in Tools and Settings Collection.

**Interactive Logon WMI Classes**

The following table lists the WMI classes that are associated with interactive logon.

**WMI Classes Associated with Interactive Logon**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_LogonSession | \root\cimv2 | Windows Server 2003<br>Windows XP |
| Win32_LogonSessionMappedDisk | \root\cimv2 | Windows Server 2003<br>Windows XP |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports Used by Interactive Logon**

Because the logon process can be deployed across various network boundaries, it can span one or more firewalls. The following table lists the three main configurable ports used by interactive logon.

**Port Assignments for Interactive Logon**

| Service Name | UDP | TCP |
|---|---|---|
| LSA RPC port | Dynamic RPC | Dynamic RPC |
| Kerberos V5 port | 88 | 88 |
| NTLM port | Dynamic | Dynamic |

You can use the Registry Editor tool to modify the registry to apply fixed ports for NTLM and the LSA.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Kerberos Authentication Tools and Settings

# Kerberos Authentication Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Kerberos Authentication Technical Reference**
The Microsoft Windows Server 2003 operating system implements the Kerberos version 5 authentication protocol. Windows Server 2003 also implements extensions for public key authentication. The Kerberos authentication client is implemented as a security support provider (SSP) and can be accessed through the Security Support Provider Interface (SSPI). Initial user authentication is integrated with the Winlogon single sign-on architecture. The Kerberos Key Distribution Center (KDC) is integrated with other Windows Server 2003 security services running on the domain controller. The KDC uses the domain's Active Directory directory service database as its security account database. Active Directory is required for default Kerberos implementations.

This subject will explain what Kerberos authentication is and how the Kerberos V5 protocol and extensions supported by Windows Server 2003 work.

**In this subject**


- What Is Kerberos Authentication?
- How the Kerberos Version 5 Authentication Protocol Works
- Kerberos Authentication Tools and Settings

# What Is Kerberos Authentication?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Kerberos Authentication?**

**In this section**


- Kerberos Authentication Benefits
- Kerberos V5 Protocol Standards
- Supported Extensions to the Kerberos V5 Protocol
- Technologies Related to Kerberos Authentication
- Kerberos Authentication Dependencies


The Kerberos version 5 authentication protocol provides a mechanism for authentication — and mutual authentication — between a client and a server, or between one server and another server.

Windows Server 2003 implements the Kerberos V5 protocol as a security support provider (SSP), which can be accessed through the Security Support Provider Interface (SSPI). In addition, Windows Server 2003 implements extensions to the protocol that permit initial authentication by using public key certificates on smart cards.

The Kerberos Key Distribution Center (KDC) uses the domain's Active Directory directory service database as its security account database. Active Directory is required for default NTLM and Kerberos implementations.

The Kerberos V5 protocol assumes that initial transactions between clients and servers take place on an open network in which packets transmitted along the network can be monitored and modified at will. The assumed environment, in other words, is very much like today's Internet, where an attacker can easily pose as either a client or a server, and can readily eavesdrop on or tamper with communications between legitimate clients and servers.

Microsoft's implementation of the Kerberos V5 protocol is:

**The default authentication package for Windows Server 2003.**
The Kerberos V5 protocol became the default authentication package with Windows 2000. Windows Server 2003 still supports NTLM for non-Kerberos clients such as the Windows NT Server 4.0 operating system.

**Based on RFC 1510 and draft revisions.**
The Kerberos protocol is a mature, widely used, open standard. Microsoft's implementation of the Kerberos V5 protocol adheres to the defined RFC standards, and thus provides interoperability with other implementations.

**Extensible.**
Kerberos architecture allows you to specify additional or alternate security methods. Also, the default shared secret key process can be supplemented with private/public key pairs by using smart cards.

**Kerberos Authentication Benefits**
The Kerberos V5 protocol is more secure, more flexible, and more efficient than NTLM. The benefits gained by using Kerberos authentication are:

**Delegated authentication.**
Windows services impersonate a client when accessing resources on the client's behalf. In many cases, a service can complete its work for the client by accessing resources on the local computer. Both NTLM and the Kerberos V5 protocol provide the information that a service needs to impersonate its client locally. However, some distributed applications are designed so that a front-end service must impersonate clients when connecting to back-end services on other computers. The Kerberos V5 protocol includes a proxy mechanism that enables a service to impersonate its client when connecting to other services. No equivalent is available with NTLM.

**Interoperability.**
Microsoft's implementation of the Kerberos V5 protocol is based on standards-track specifications that are recommended to the Internet Engineering Task Force (IETF). As a result, the implementation of the Kerberos V5 protocol in Windows

Server 2003 lays a foundation for interoperability with other networks in which the Kerberos V5 protocol is used for authentication.

**More efficient authentication to servers.**
With NTLM authentication, an application server must connect to a domain controller in order to authenticate each client. With the Kerberos V5 authentication protocol, on the other hand, the server is not required to go to a domain controller. Instead, the server can authenticate the client by examining credentials presented by the client. Clients can obtain credentials for a particular server once and then reuse those credentials throughout a network logon session. Renewable session tickets replace pass-through authentication. For more information about what renewable session tickets are and how they work, please see "How the Kerberos Version 5 Authentication Protocol Works."

**Mutual authentication.**
By using the Kerberos protocol, a party at either end of a network connection can verify that the party on the other end is the entity it claims to be. Although NTLM enables servers to verify the identities of their clients, NTLM does not enable clients to verify a server's identity, nor does NTLM enable one server to verify the identity of another. NTLM authentication was designed for a network environment in which servers were assumed to be genuine. The Kerberos V5 protocol makes no such assumption.

**Kerberos V5 Protocol Standards**
The Kerberos authentication protocol originated at MIT more than a decade ago, where it was developed by engineers working on Project Athena. The first public release was the Kerberos version 4 authentication protocol. After extensive industry review of that protocol, the protocol's authors developed and released the Kerberos version 5 authentication protocol.

The Kerberos V5 protocol is now on a standards track with the IETF. The implementation of the protocol in Windows Server 2003 closely follows the specification defined in Internet RFC 1510. In addition, the mechanism and format for passing security tokens in Kerberos messages follows the specification defined in Internet RFC 1964.

The Kerberos V5 protocol specifies mechanisms to:

- **Authenticate user identity.** When a user wants to gain access to a server, the server needs to verify the user's identity. Consider a situation in which the user claims to be, for example, Alice@tailspintoys.com. Because access to resources are based on identity and associated permissions, the server must be sure the user really has the identity it claims.

- **Securely package the user's name.** The user's name — that is, the User Principal Name (UPN): Alice@tailspintoys.com, for example — and the user's credentials are packaged in a data structure called a ticket.

- **Securely deliver user credentials.** After the ticket is encrypted, messages are used to transport user credentials along the network.

### Note

- Although the Kerberos protocol authenticates the user's identity, it does not authorize access. This is an important distinction. Tickets in other contexts, such as drivers' licenses, often both prove identity and authorize actions or access. A Kerberos ticket only proves that the user is who the user claims to be. After the user's identity is verified, the Local Security Authority will authorize or deny resource access.

**Keys: Privacy Through Encryption**
Kerberos messages are encrypted with various encryption keys to ensure that no one can tamper with the client's ticket or with other data in a Kerberos message.

Possible Kerberos keys include:

**Long-term key.**
A key — known only to the target server and the KDC — with which the client's ticket is encrypted.

**Client/server session key.**
A short-term, single-session key that is created by the KDC and used to encrypt the client-to-server and server-to-client messages after identity and authorization have been confirmed.

**KDC/user session key.**
The KDC and the user share a secret encryption key as well, which is used, for example, to encrypt the message to the client containing a session key.

**The Kerberos V5 protocol can use both symmetric and asymmetric encryption**
Because most Kerberos encryption methods are based on keys that can be created only by the KDC and the client, or by the KDC and a network service, the Kerberos V5 protocol is said to use symmetric encryption. That is, the same key is used to encrypt and decrypt messages.

Microsoft's implementation of the Kerberos protocol can also make limited use of asymmetric encryption. A private/public key pair can be used to encrypt or decrypt initial authentication messages from a network client or a network service.

**Kerberos Authenticator Prevents Packet Replay**
The Kerberos authentication mechanism creates and securely delivers an authenticator — usually based on unique timestamps — along with the client's ticket. The authenticator is unique and valid for only one use. This limited use minimizes the possibility of someone obtaining and re-using (replaying) the client ticket, perhaps in an attempt to steal and use the client's identity.

**Supported Extensions to the Kerberos V5 Protocol**
Windows Server 2003 implements extensions to the Kerberos V5 protocol. The extensions enable initial authentication by using public key certificates instead of conventional shared secret keys. This enhancement enables the protocol to support interactive logon with a smart card. The extensions for public key authentication are based on a draft specification being worked on by the IETF working group.

**Technologies Related to Kerberos Authentication**
The following diagram shows how Kerberos authentication fits with other technologies in Windows Server 2003. Depending on whether the client or server applications are user-mode or kernel-mode applications, they use either Secur32.dll or Ksecdd.sys, respectively, by means of SSPI calls to communicate with the Local Security Authority Subsystem (LSASS).

### Kerberos Architecture




The following table is a description of the components that participate in Kerberos Authentication.

### Security Subsystem Components Used in Digest Authentication

| Component | Description |
|---|---|
| Kerberos.dll | The SSP that implements an industry-standard protocol that is used with either a password or a smart card for interactive logon. It is also the preferred authentication method for services in Windows 2000 and Windows Server 2003. |
| Kdcsvc.dll | The Kerberos Key Distribution Center (KDC) service, which is responsible for providing ticket-granting tickets to clients. |

| | |
|---|---|
| Ksecdd.sys | The Kernel Security Device Driver is used to communicate with LSASS in user mode. |
| Lsasrv.dll | The LSA Server service, which both enforces security policies and acts as the security package manager for the LSA. |
| Secur32.dll | The Secur32.dll component is the multiple authentication provider that implements SSPI for user mode applications. |

Windows Server 2003 implements the Kerberos V5 authentication protocol as an SSP, which is a dynamic-link library (DLL) supplied with the operating system. The system uses the Kerberos SSP, Kerberos.dll, as its first choice for authentication. After the LSA establishes a security context for an interactive user, another instance of the Kerberos SSP can be loaded by a process running in the user's security context in order to support the signing and sealing of messages.

Because the Kerberos protocol is the preferred authentication protocol for Windows Server 2003, all domain services support the Kerberos SSP, including:

- Active Directory queries that use the Lightweight Directory Access Protocol (LDAP).
- Remote server or workstation management using RPC calls.
- Print services.
- Client-server authentication.
- Remote file access that uses Common Internet File System/server message block (CIFS/SMB).
- Distributed file system management and referrals.
- Intranet authentication to Internet Information Services (IIS).
- Security authority authentication for Internet Protocol security (IPSec).
- Certificate requests to Certificate Services for domain users and computers.

**Kerberos Authentication Dependencies**
This section reviews dependencies and summarizes how each dependency relates to Kerberos authentication.

**Operating System**
Kerberos authentication relies on client functionality that is built in to the Windows Server 2003 operating system, the Microsoft Windows XP operating system, and the Windows 2000 operating system. If a client, domain controller, or target server is running an earlier operating system, it cannot natively use Kerberos authentication.

**TCP/IP Network Connectivity**
For Kerberos authentication to occur, TCP/IP network connectivity must exist between the client, the domain controller, and the target server. For more information about TCP/IP, see "TCP/IP Technical Reference."

**Domain Name System**
The client uses the fully qualified domain name (FQDN) to access the domain controller. DNS must be functioning for the client to obtain the FQDN.

For best results, do not use Hosts files with DNS. For more information about DNS, see "DNS Technical Reference."

**Active Directory Domain**
Kerberos authentication is not supported in earlier operating systems, such as the Microsoft Windows NT 4.0 operating system. You must be using user and computer accounts in the Active Directory directory service to use Kerberos authentication. Local accounts and Windows NT domain accounts cannot be used for Kerberos authentication.

**Time Service**
For Kerberos authentication to function correctly, all domains and forests in a network should use the same time source so that the time on all network computers is synchronized. An Active Directory domain controller acts as an authoritative source of time for its domain, which guarantees that an entire domain has the same time. For more information, see "Windows Time Service Technical Reference."

**Service Principal Names**

Service principal names (SPNs) are unique identifiers for services running on servers. Every service that uses Kerberos authentication needs to have an SPN set for it so that clients can identify the service on the network. If an SPN is not set for a service, clients have no way of locating that service. Without correctly set SPNs, Kerberos authentication is not possible. For more information about user-to-user authentication, see "How the Kerberos Version 5 Authentication Protocol Works," and search for "The User-to-User Authentication Process."

# How the Kerberos Version 5 Authentication Protocol Works

Updated: October 7, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How the Kerberos Version 5 Authentication Protocol Works**

**In this section**


- Kerberos SSP Architecture
- Kerberos Physical Structure
- Kerberos Processes and Interactions
- Network Ports Used by the Kerberos V5 Protocol
- Related Information


The Kerberos Network Authentication Service version 5, defined in RFC 1510, provides a means of verifying the identities of principals on an open, potentially insecure network. This section discusses how the RFC standard Kerberos version 5 authentication protocol is used in Windows Server 2003.

This section is divided into four subsections:

- **Kerberos SSP Architecture** illustrates how the Microsoft Security Support Provider Interface (SSPI) in Windows Server 2003 provides a mechanism of generic routines that can access the Kerberos Security Support Provider (SSP).
- **Kerberos Physical Structure** discusses components of the Windows Server 2003 implementation of Kerberos authentication. These components include keys, tickets, and the Key Distribution Center (KDC).
- **Kerberos V5 Authentication Protocol Processes and Interactions** illustrates how the Kerberos protocol is used in different situations, details what is included in Kerberos messages, and discusses other related technologies. The examples will show the complete process, including related components and processes that are not defined by the Kerberos Network Authentication Service. Some processes—such as how to find an authentication service, what credentials information is passed, and where credentials information is stored—are specific to Windows systems and might or might not differ in other implementations of the Kerberos protocol.
- **Network Ports Used by the Kerberos V5 Protocol** tabulates networks ports used during Kerberos authentication.

**Kerberos SSP Architecture**

Windows Server 2003 implements the Kerberos V5 authentication protocol as a Security Support Provider (SSP), a dynamic-link library (DLL) supplied with the operating system. Windows Server 2003 includes an SSP for NTLM authentication as well. By default, both SSPs are loaded by the Local Security Authority (LSA) on a Windows Server 2003 computer when the system boots. The system can use either SSP to authenticate network logons and client/server connections. Which SSP is used depends on the capabilities of the computer on the other side of the connection and the preferences of the individual application that is being used.

The Microsoft SSPI is the foundation for authentication in Windows Server 2003. That is, applications and infrastructure services that require authentication use SSPI to provide it.

The SSPI is the implementation of the Generic Security Service API (GSSAPI) in Windows Server 2003. For more information about GSSAPI, see RFC 2743 and RFC 2744 in the IETF RFC Database.

The default SSPs in Windows Server 2003—Negotiate (SPNEGO), Kerberos, NTLM, Schannel, and Digest authentication protocols—are plugged into the SSPI in the form of DLLs. Additional SSPs can be plugged in if they can interoperate with the SSPI.

**SSPI Architecture graphic**




The SSPI in Windows Server 2003 provides a mechanism that carries authentication tokens over the existing communication channel between the client and server. When two parties need to be authenticated so that they can communicate securely, the requests for authentication are routed to the SSPI, which completes the authentication process, regardless of the network protocol currently in use. The SSPI returns transparent binary large objects, which are then passed to the other side of the connection by the application, at which point they can be passed to the SSPI layer on that side. Thus, the SSPI enables an application to use various security models available on a computer or network without changing the interface to the security system.

The following table describes the SSP components that are plugged into the SSPI. Each of the protocols in the table is used in different ways in Windows Server 2003 to promote secure communication in an insecure network environment.

**SSP Layer Components**

| Component | Description |
|---|---|
| Kerberos V5 authentication | An industry-standard protocol that is used with either a password or a smart card for interactive logon. It is also the preferred authentication method for services in Windows 2000 and Windows Server 2003. |
| NTLM authentication | A challenge-response protocol that is used to provide compatibility with versions of Windows earlier than Windows 2000. |
| Digest authentication | An industry standard that is used in Windows Server 2003 for Lightweight Directory Access Protocol (LDAP) and Web authentication. Digest transmits credentials across the network as an MD5 hash or message digest. |
| Schannel | An SSP that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) Internet standard authentication protocols. Schannel is used for Web-based server authentication such as when a user attempts to access a secure Web server. |
| Negotiate | An SSP that can be used to negotiate a specific authentication protocol. When an application calls into SSPI to log on to a network, it can specify an SSP to process the request. If the application specifies Negotiate, Negotiate analyzes the request and picks the best SSP to handle the request based on customer-configured security policy. |

**Kerberos SSP**

Windows Server 2003 implements the Kerberos V5 authentication protocol as an SSP, a DLL supplied with the operating system. The system uses the Kerberos SSP, Kerberos.dll, as its first choice for authentication. After the LSA establishes a security context for an interactive user, another instance of the Kerberos SSP can be loaded by a process running in the user's security context to support the signing and sealing of messages.

Because the Kerberos protocol is the preferred authentication protocol for Windows Server 2003, all domain services support the Kerberos SSP, including:

- Active Directory queries using the Lightweight Directory Access Protocol (LDAP)
- Remote server or workstation management using RPC calls
- Print services
- Client-server authentication
- Remote file access using the Common Internet File System/Server Message Block (CIFS/SMB)
- Distributed file system management and referrals

- Intranet authentication to Internet Information Services (IIS)
- Security authority authentication for Internet Protocol Security (IPSec)
- Certificate requests to Certificate Services for domain users and computers

**Kerberos Physical Structure**
Kerberos authentication provides a mechanism for mutual authentication between a client and a server on an open network—a network in which packets transmitted along the network could be monitored and modified at will. In order to provide secure authentication, Kerberos authentication uses symmetric keys, encrypted objects, and Kerberos services.

This subsection will cover the following topics:

- Keys Used in Kerberos Authentication
- Kerberos Tickets
- The Authenticator
- Credentials Cache
- Key Distribution Center
- The Kerberos Realm

**Kerberos Components**




**Keys Used in Authentication**
**Why are keys needed?**
The Kerberos protocol relies heavily on an authentication technique involving shared-secret keys. The basic shared-secret concept is quite simple: If a secret is known by only two people, then either person can verify the identity of the other by confirming that the other person knows the secret.

For example, suppose that Alice often sends messages to Bob, and that Bob needs to ensure that a message from Alice really has come from Alice before he acts on its information. They decide to solve their problem by selecting a password and agreeing to share the secret password between the two of them, but not with anyone else. If a message purported to be from

Alice can somehow demonstrate that the sender knows the password, Bob can verify that the sender is indeed Alice.

The only question left for Alice and Bob to resolve is how Alice will show that she knows the password. She could include it somewhere in her messages, perhaps in a signature block at the end—Alice, Our$ecret. This would be simple and efficient and would be effective if Alice and Bob could be sure that no one else is reading their mail. Unfortunately, their messages pass over a network used by people like Carol, who uses a network analyzer to scan traffic in hope that one day she might spot a password. Thus, Alice must not prove that she knows the secret by including it in her message. To keep the password secret, she must show that she knows it without revealing it.

The Kerberos protocol solves this problem with secret key cryptography. Instead of sharing a password, communication partners share a cryptographic key, and they use knowledge of this key to verify one another's identity. In order for the technique to work, the shared key must be symmetric—that is, a single key must be capable of both encryption and decryption. One party proves knowledge of the key by encrypting a piece of information; the other proves knowledge of the key by decrypting the information.

Kerberos authentication relies on several keys and key types for encryption. Key types can include long-term symmetric keys, long-term asymmetric keys, and short-term symmetric keys. The authentication protocol was designed to use symmetric encryption, meaning that the same shared key is used by the sender and the recipient for encryption and decryption.

**Long-Term Symmetric Keys: User, System, Service, and Inter-realm Keys**
The long-term symmetric keys are derived from a password. The plaintext password is transformed into a cryptographic key by passing the text of the password through a cryptographic function. (All implementations of the Kerberos version 5 authentication protocol must support DES-CBC-MD5. Other algorithms are permissible.) The result of the cryptographic function is the key.

**User keys**
When a user is created, the password is used to create the user key. In Active Directory domains, the user key is stored with the user's object in the Active Directory. At the workstation, the user key is created when the user logs on.

**System keys**
When a workstation or a server joins a Windows domain, it receives a password. In the same manner as a user account, the system account's password is used to create the system key.

**Service keys**
Services use a key based on the account password they use to log on.

All KDCs in the same realm use the same service key. This key is based on the password assigned to the krbtgt account. Every Active Directory domain will have this built-in account.

**Inter-realm keys**
In order for cross-realm authentication to occur, the KDCs must share an inter-realm key. The realms can then trust each other because they share the key.

Active Directory domains that have a parent-child relationship share an inter-realm key. This inter-realm key is the basis of transitive trusts in Windows 2000 and Windows Server 2003. If a shortcut trust is created, the two domains will exchange a key specifically for their trust.

**Kerberos SSP encryption key lengths**
Kerberos SSP supports different encryption types and key lengths depending on the task to be completed and the options specified. Although the size of the key determines the degree of protection the key provides, key size does not significantly impact the size of the ticket. The following table lists the key lengths that Kerberos SSP supports for the various types of encryption.

### Key Lengths for Encryption Types

| Encryption Algorithm | Key Length |
|---|---|
| RC4-HMAC | 128 |
| DES-CBC-CRC | 56 |
| DES-CBC-MD5 | 56 |

**Long-Term Asymmetric Keys: Public Key**
Currently, public-key certificates stored on smart cards are the only long-term asymmetric keys in the Microsoft implementation of Kerberos authentication.

**Short-Term Symmetric Keys: Session Keys**
The session keys used for ticket-granting tickets (TGTs) and service tickets are short-lived and used only as long as that session or service ticket is valid.

**Kerberos Tickets**
The main component of Kerberos authentication is the ticket. The Kerberos messages are used to request and deliver tickets. There are two types of tickets used in Kerberos authentication, TGTs and service tickets.

**Kerberos Ticket Requests**
The Kerberos client sends ticket requests to the KDC:

| Requested Ticket Type | KDC Service That Receives Request |
|---|---|
| TGT | Authentication service |
| Service ticket | Ticket-granting service |

The ticket requests include:

- Requested properties (flags), such as whether the ticket is renewable.
- Requested encryption method.

**Ticket-Granting Tickets**
The KDC responds to a client's authentication service request by returning a service ticket for itself. This special service ticket is called a ticket-granting ticket (TGT). A TGT enables the authentication service to safely transport the requester's credentials to the ticket-granting service.

A TGT is:

- A user's initial ticket from the authentication service.
- Used to request service tickets.
- Meant only for use by the ticket-granting service.

TGTs are encrypted with a key shared by the KDCs. The client cannot read tickets. Only KDC servers can read TGTs to secure access to user credentials, session keys, and other information. Like an ordinary service ticket, a TGT contains a copy of the session key that the service (in this case the KDC) will use in communicating with the client. The TGT is encrypted with the KDC's long-term key.

From the client's point of view, a TGT is just another ticket. Before it attempts to connect to any service, the client first checks its credentials cache for a service ticket to that service. If it does not have one, it checks the cache again for a TGT. If it finds a TGT, the client fetches the corresponding TGS session key from the cache, uses this key to prepare an authenticator (described later in this document), and sends both the authenticator and the TGT to the KDC, along with a request for a service ticket for the service. In other words, gaining admission to the KDC is no different from gaining admission to any other service in the domain—it requires a session key, an authenticator, and a ticket.

From the KDC's point of view, TGTs enable it to avoid the performance penalties of looking up a user's long term key every time the user requests a service. The KDC looks up the user's long-term key only once, when it grants an initial TGT. For all other exchanges with this client, the KDC can decrypt the TGT with its own long-term key, extract the session key, and use that to validate the client's authenticator.

**Service tickets**
A service ticket enables the ticket-granting service (TGS) to safely transport the requester's credentials to the target server or service.

The KDC responds to the client's request to connect to a service by sending both copies of the session key to the client. The client's copy of the session key is encrypted with the key that the KDC shares with the client. The service's copy of the session key is embedded, along with information about the client, in a data structure called a service ticket. The entire structure is then encrypted with the key that the KDC shares with the service. The ticket—with the service's copy of the session key safely inside—becomes the client's responsibility to manage until it contacts the service.

A service ticket is used to authenticate with services other than the TGS and is meant only for the target service.

A service ticket is encrypted with a service key, which is a long-term key shared by the KDC and the target service. Thus, although the client manages the service ticket, the client cannot read it. Only the KDC and the target service can read tickets, enabling secure access to user credentials, the session key, and other information.

**Note**

- The KDC is simply providing a ticket-granting service. It does not keep track of its messages to make sure they reach the intended address. No harm will be done if the KDC's messages fall into the wrong hands. Only someone who knows the client's secret key can decrypt the client's copy of the session key. Only someone who knows the server's secret key can read what is inside the ticket.

When the client receives the KDC's reply, it extracts the ticket and the client's copy of the session key, putting both aside in a secure cache (located in volatile memory, not on disk). When the client wants admission to the server, it sends the server a message that consists of the ticket, which is still encrypted with the server's secret key, and an authenticator, which is encrypted with the session key. The ticket and authenticator together are the client's credentials to the server.

**Benefits of service tickets**

The server does not have to store the session key that it uses in communicating with clients. It is the client's responsibility to hold a ticket for the server in its credentials cache and present the ticket each time it wants access to the server. Whenever the server receives a service ticket from a client, the server can use its secret key to decrypt the ticket and extract the session key. When the server no longer needs the session key, it can discard it.

The client does not need to go back to the KDC each time it wants access to this particular server. Service tickets can be reused. To guard against the possibility that someone might steal a copy of a ticket, service tickets have an expiration time that is specified by the KDC in the ticket's data structure.

How long a ticket is valid depends on the policy for the realm. Although the RFC recommends a maximum ticket lifetime of one day, in both MIT Kerberos 5 Release 1.3.1 and the Windows 2000 or Windows Server 2003 implementations of the Kerberos protocol, tickets are good for no longer than 10 hours by default, about the length of a normal logon session. When the user logs off, the credentials cache is flushed and all service tickets—as well as all session keys—are destroyed.

**What Is in a Ticket?**

For our purpose here, it is enough to list the fields in a ticket and to describe the information they contain. The exact data structures for tickets as well as messages can be found in RFC 1510.

The first three fields in a ticket are not encrypted. The information is in plaintext so that the client can use the information to manage tickets in its cache.

**Ticket Contents**

| Field Name | Description |
|---|---|
| Ticket Version Number | 5 |
| Realm | Name of the realm (domain) that issued the ticket. A KDC can issue tickets only for servers in its own realm, so this is also the name of the server's realm. |
| Server Name | Name of the server. |
| | The following fields are encrypted with the server's secret key. |
| Flags | Options that specify how and when the ticket can be used. See table below for details about ticket flags. |
| Key | The session key used to encrypt and decrypt client and target server messages. |
| Client Realm | Name of the requester's realm. |
| Client Name | Requester's name. |
| Transited | Lists the Kerberos realms that took part in the authentication of the client during cross-realm authentication. |

| | |
|---|---|
| Authentication Time | Time of initial authentication by the client.<br><br>The KDC places a timestamp in this field when it issues a TGT. When it issues tickets based on a TGT, the KDC copies the authentication time of the TGT to the authentication time of the ticket.<br><br>An application could have a policy to reject tickets that were based on "old" TGTs, requiring the client to obtain a new TGT and then a new service ticket. |
| Start Time | Time after which the ticket is valid. |
| End Time | Ticket's expiration time. |
| Renew Till | (Optional) Maximum end time that can be set in a ticket with a RENEWABLE ticket flag. |
| Client Address | (Optional) One or more addresses from which the ticket can be used. If omitted, the ticket can be used from any address. |
| Authorization-Data | (Optional) This field contains authorization data for the client.<br><br>In MIT's implementation of the Kerberos protocol, this field normally contains access restrictions. For example, if this were a ticket for a print server sent by a client wishing to print a file, the file name could be included here, restricting the print server's access on behalf of the client to only this file.<br><br>In the Microsoft implementation, this field is significant because it contains the user's SID and the SIDs of groups the client belongs to—data which is used to build the client's access token (the privilege attributes for the client).<br><br>The Kerberos service does not interpret the contents of this field. Interpretation is left up to the service requesting credentials. |

The following table lists and describes the Kerberos ticket Flag field. The Flag field is a bit-field in which options are set by turning a particular bit on (1) or off (0). Although the field is 32 bits long, only 11 ticket flags are of interest to Kerberos administrators.

**Ticket Flags**

| Flag | Description |
|---|---|
| FORWARDABLE | (TGT only). Tells the ticket-granting service that it can issue a new TGT—based on the presented TGT—with a different network address based on the presented TGT. |
| FORWARDED | Indicates either that a TGT has been forwarded or that a ticket was issued from a forwarded TGT. |
| PROXIABLE | (TGT only). Tells the ticket-granting service that it can issue tickets with a network address that differs from the one in the TGT. |
| PROXY | Indicates that the network address in the ticket is different from the one in the TGT used to obtain the ticket. |
| MAY-POSTDATE | (TGT only). Tells the ticket-granting service (TGS) that a postdated ticket can be issued. |
| POSTDATED | Indicates that the ticket has been postdated. |
| INVALID | Indicates that the ticket is invalid and must be validated by the KDC before use. A postdated ticket is flagged as invalid until its valid start time. |
| RENEWABLE | Used in combination with the End Time and Renew Till fields to cause tickets with long life spans to be renewed at the KDC periodically. |
| INITIAL | Indicates that a ticket was issued using the authentication service (AS) exchange and not issued based on a TGT. |
| PRE-AUTHENT | Indicates that the client was authenticated by the KDC before a ticket was issued. This flag usually indicates the presence of an authenticator in the ticket. It can also flag the presence of credentials taken from a smart card logon. |

| | |
|---|---|
| HW-AUTHENT | Indicates that the initial authentication required the use of hardware. |

**What information clients have about tickets**

A client needs to have some information about what is inside tickets and TGTs in order to manage its credentials cache. When the KDC returns a ticket and session key as the result of an authentication service (AS) or ticket-granting service (TGS) exchange, it packages the client's copy of the session key in a data structure that includes the information in the following ticket fields: Authentication Time, Start Time, End Time, and Renew Till. The entire structure is encrypted in the client's key and returned with the KRB_AS_REP or KRB_TGS_REP messages.

**Kerberos authorization data for Windows accounts**

The Kerberos protocol describes a number of fields to be used in the authentication exchange. One of these, the optional authorization data field, is intended to be specific to the end service, in this case, Microsoft authorization. The Microsoft implementation follows the Kerberos specification (RFC 1510) by using the authorization field to store information about user identity and group membership. This information is stored in Microsoft authorization data that is generated by a domain controller during an AS or TGS request. The Microsoft authorization data includes the following groups of information:

- **Credential information.** Information about the ticket and the client, such as Group IDs, which list relative IDs (RIDs) that belong to the client for his or her domain and groups.
- **Client information.** Ensures that the ticket belongs to the client who is submitting it.
- **Supplemental credentials.** Returned by the KDC service to enable PKINIT.
- **Signatures.** The Microsoft authorization data contains two digital signatures, one for the domain controller and one for the server offering the service. These signatures are used to prevent clients or untrusted services from generating their own forged Microsoft authorization data.
- **PAC-request pre-authentication data.** Normally, the Microsoft authorization data is included in every pre-authenticated ticket received from an AS request. However, a client can also explicitly request — with the privilege attribute certificate (PAC)-request pre-authentication data field — to include or not include SIDs. If the field is present, the account SIDs will be included. If the field is not present, SIDs will not be included.

For more information about Microsoft authorization data present in Kerberos tickets issued by Windows domain controllers, see "Utilizing the Windows 2000 Authorization Data in Kerberos Tickets for Access Control to Resources," a Windows Development (General) topic on MSDN.

**How the KDC limits a ticket's lifetime**

Each ticket has a start time and an expiration time. At any time after the start time but before the expiration time, a client holding a ticket for a service can present the ticket and gain access to the service, no matter how many times the client has used the ticket before. In order to reduce the risk that a ticket or the corresponding session key might be compromised, administrators can set the maximum lifetime for tickets. The maximum lifetime for tickets setting is an element of Kerberos policy.

When a client asks the KDC for a ticket to a service, the client can request a specific start time. If this time is missing from the request or is a time in the past, the KDC sets the ticket's start time field to the current time.

Whether or not a client specifies a start time, its request must include a desired expiration time. The KDC determines the value in a ticket's End Time field by adding the maximum ticket life fixed by Kerberos policy to the value in the ticket's Start Time field. It then compares the result with the requested expiration time. Whichever time is sooner becomes the ticket's end time.

**What happens when tickets expire**

The KDC does not notify clients when service tickets or TGTs are about to expire. Furthermore, other than keeping short-term records needed to prevent replay attacks, it does not keep track of transactions with clients.

If a client presents an expired service ticket when requesting a connection to a server, the server returns an error message. The client must request a new service ticket from the KDC. After a connection is authenticated, however, it no longer matters whether the service ticket remains valid. Service tickets are used only to authenticate new connections with servers. Ongoing operations are not interrupted if the service ticket used to authenticate the connection expires during the connection.

If a client presents an expired TGT when requesting a service ticket from the KDC, the KDC responds with an error message. The client must request a new TGT, and to do that it needs the user's long-term key. If the client did not cache the user's long-term key during the initial logon process, the client might have to ask the user for a password and derive the long-term key.

**Renewable TGTs**
When tickets are renewable, session keys are refreshed periodically without issuing a completely new ticket. If Kerberos policy permits renewable tickets, the KDC sets a RENEWABLE flag in every ticket it issues and sets two expiration times in the ticket. One expiration time limits the life of the current instance of the ticket; the second expiration time sets a limit on the cumulative lifetime of all instances of the ticket.

The expiration time for the current instance of the ticket is held in the End Time field. As with non-renewable tickets, the value in the End Time field equals the value in the Start Time field plus the value of the maximum ticket life specified by Kerberos policy. A client holding a renewable ticket must send it—presenting a fresh authenticator as well—to the KDC for renewal before the end time is reached. When the KDC receives a ticket for renewal, it checks the value of a second expiration time held in the Renew Till field. This value is set when the ticket is first issued. It equals the value in the tickets Start Time field plus the value of the maximum cumulative ticket life specified by Kerberos policy. When the KDC renews the ticket, it checks to determine if the renew-till time has not yet arrived. If it has not, the KDC issues a new instance of the ticket with a later end time and a new session key.

This means that administrators can set Kerberos policy so that tickets must be renewed at relatively short intervals—every day, for example. When tickets are renewed, a new session key is issued, minimizing the value of a compromised key. Administrators can also set cumulative ticket life for a relatively long period—one week or one month, for example. At the end of that time, the ticket expires and is no longer valid for renewal.

**The Authenticator**
Whenever a client sends a ticket to a target server—whether the server is the TGS or some other network server or resource—the client also includes an authenticator in the message. The authenticator verifies that the client listed as the sender in the ticket really is the ticket's source.

**Why is an authenticator necessary?**
The target server can trust the contents of the ticket because the ticket is encrypted with the target server's secret key. However, the target server cannot trust that the ticket was really sent by the client specified in the ticket. The ticket could have been stolen and is now being included in an imposter's message.

**How does the authenticator work?**

- The authenticator is encrypted with the session key created by the KDC to be used between the client and the target server. Only the client and the target server can access the session key.
- The target server uses its secret key to decrypt the ticket, finds the session key inside the ticket, and uses it to decrypt the authenticator.
- If the target server can successfully decrypt the authenticator and if the authenticator's data is accurate, then the target server will trust the source of the ticket.

**The authenticator's timestamp**
The timestamp is perhaps the most important piece of data in the authenticator. A domain's Kerberos policy and policies established by applications typically require that the timestamp be within minutes of the time on the target server. Otherwise, the authenticator and the ticket will be rejected. This helps prevent message replay.

Although no authenticator is included in the initial authentication service request message (explained later in this document), the message does include the client's timestamp in the pre-authentication field.

**What is in an authenticator?**
The following table lists and describes authenticator fields. The exact data structures for authenticators as well as messages can be found in RFC 1510.

### Authenticator Fields

| Field | Description |
|---|---|
| Checksum | A checksum of the data from the ticket-granting service request (KRB_TGS_REQ) or application server request (KRB_AP_REQ) message that the authenticator was a part of. This helps validate that the message sender was really the client. |
| Subkey | In a KRB_AP_REQ, this field specifies a subkey that the target server and the client will use to encrypt further communications, instead of using the TGS-supplied session key. |

| | |
|---|---|
| Sequence Number | An optional field that could be used to prevent replays. This sequence number might be based on the client nonce, for example. |
| Authorization Data | This optional field can include authorization data for specific applications. This is not the same as the Authorization Data field that carries the user's privilege attribute certificate (PAC). |

**Credentials Cache**

On computers running Windows 2000, Windows XP, or Windows Server 2003, tickets and keys obtained from the KDC are stored in a credentials cache, an area of volatile memory protected by the LSA. The credentials cache is never paged to disk. All objects stored there are destroyed when a security principal logs off or when the system is shut down.

The credentials cache is managed by the Kerberos SSP, which runs in the LSA's security context. Whenever tickets and keys need to be obtained or renewed, the LSA calls the Kerberos SSP to accomplish the task.

### Note

- In Windows XP and Windows 2000 with Service Pack 2 or later, TGT renewal is triggered when the TGT is used within 5 minutes of its expiration.
- In Windows Server 2003, periodically the system will automatically renew expiring TGTs.

The LSA also keeps a copy of an interactive user's hashed password. If the user's TGT expires during a logon session, the Kerberos SSP uses the LSA's copy of the hashed password to obtain a new TGT without interrupting the user's logon session. The password is not stored permanently on the computer, and the local copy of the hashed password is destroyed when the user's logon session is destroyed.

Hashed passwords for services and computers are handled differently than above. As in Windows NT, hashed passwords are stored in a secure area of the computer's registry. The registry is also used to store hashed passwords for user accounts on the local system, but local accounts are used only for access to computers in standalone mode, never for network access.

**Key Distribution Center**

The Key Distribution Center (KDC) is a service that runs on a physically secure server. The KDC maintains a database with account information for all security principals in its realm (the Kerberos equivalent of a Windows Server 2003 domain). Along with other information about each security principal, the KDC stores a cryptographic key known only to the security principal and the KDC. This key—also called the long-term key—is used in exchanges between the security principal and the KDC. In most implementations of the protocol, the long-term key is derived from a user's logon password.

As in other implementations of the Kerberos protocol, Microsoft implements the KDC as a single process that provides two services:

**Authentication service (AS)**

The AS issues TGTs good for admission to the ticket-granting service in its domain. Before network clients can get tickets for services, each client must get an initial TGT from the authentication service in the user's account domain.

**Ticket-granting service (TGS)**

The TGS issues tickets good for admission to other services in the TGS's domain or to the ticket-granting service of a trusted domain. When a client wants access to a service, it must contact the ticket-granting service in the service's account domain, present a TGT, and ask for a ticket. If the client does not have a TGT valid for admission to that ticket-granting service, it must get one through a referral process that begins at the ticket-granting service in the user account's domain and ends at the ticket-granting service in the service account's domain.

Windows Server 2003 implements the Key Distribution Center (KDC) as a domain service. It uses the domain's Active Directory as its account database and gets some information about users from the global catalog.

The KDC is located on every domain controller, as is the Active Directory directory service. Both services are started automatically by the domain controller's Local Security Authority (LSA) and run in the process space of the LSA. Neither service can be stopped. Windows Server 2003 ensures availability of these services by allowing each domain to have several domain controllers, all peers. Any domain controller can accept authentication requests and ticket-granting requests addressed to the domain's KDC.

The security principal name used by the KDC for an Active Directory domain is krbtgt, as specified by RFC 1510. An account for this security principal is created automatically when a new domain is created. The account cannot be deleted, nor can the name be changed.

A password is automatically assigned to the krbtgt account. The password for the krbtgt account is used to derive a secret key for encrypting and decrypting the TGTs that the KDC issues. Passwords for domain trusts are either automatically assigned or provided when the trust is established. The password for a domain trust is used to derive an inter-realm key for encrypting referral tickets. (Referral tickets are discussed later in this document.) Although passwords for domain trusts are changed automatically, passwords for trusts with non-Windows realms and the krbtgt accounts must be changed either manually or through a script.

All instances of the KDC within a domain use the domain account for the security principal krbtgt. A client addresses messages to a domain's KDC by including both the service's security principal name, krbtgt, and the name of the domain. Both items of information are also used in tickets to identify the issuing authority. For information about name forms and addressing conventions, see RFC 1510 in the IETF RFC Database.

**Session Key Distribution: the Session Key**

If each client were to need a unique key for every service, and if each service were to need a unique key for every client, key distribution could quickly become a tough problem to solve. Furthermore, the need to store and protect so many keys on so many computers would present an enormous security risk. Hypothetically, a network of *n* accounts would require $n(n-1)/2$ keys, where *n* is an integer greater than 1. Thus, a network of only 10 users, services, and systems would need 45 keys; a network of 100 accounts would need 4,950 keys. Every new user, service, or system would generate a new key pair for each interaction.

The name of the Kerberos protocol suggests how it solves the problem of key distribution. Kerberos (Cerberus), in Greek mythology, was the fierce, three-headed dog that guarded the gates of the Underworld. The Kerberos protocol—like Kerberos the guard—has three heads: a client, a server, and a trusted third party to mediate between the other two. The trusted intermediary in the protocol is the Key Distribution Center (KDC).

When a client wants to connect to a server, the client sends a request to the KDC, and the KDC distributes a unique, short-term session key for the two parties to use when they authenticate each other. The server's copy of the session key is encrypted in the server's long-term key. The client's copy of the session key is encrypted in the client's long-term key.

**KDC Configuration**

RFC 1510 recommends the following values for KDC configuration:

### KDC configuration

| Configuration Element | RFC 1510 Recommendation | Active Directory Domain Default Setting |
|---|---|---|
| **Maximum ticket lifetime** | One day | 600 minutes (10 hours) |
| **Maximum renewable lifetime** | One week | Seven days |
| **Empty addresses** | Only when suitable restrictions appear in authorization data | Support empty addresses |

For more information about Kerberos policy settings, see Account Policy Settings.

For more information about KDC configuration, see RFC 1510 in the IETF RFC Database.

**Pre-authentication**

By default, the KDC requires all accounts to use pre-authentication. However, pre-authentication can be disabled for individual accounts when necessary for compatibility with other implementations of the protocol.

**Active Directory as a KDC account database**

The account database that the KDC needs in order to obtain information about security principals is provided in Windows Server 2003 by the domain's Active Directory. Each principal is represented by an account object. The encryption key used in communicating with a user, computer, or service is stored as an attribute of that security principal's account object.

Only domain controllers are Active Directory servers. Because each domain controller keeps a writeable copy of the directory, you can create accounts, reset passwords, and modify group membership at any domain controller. Changes made to one replica of the directory are automatically propagated to all other replicas. Windows 2000 and Windows Server 2003 do not, however, implement the Kerberos replication protocol. Instead, the systems replicate the information store for Active Directory using a proprietary multi-master replication protocol over a secure channel between replication partners.

Physical storage of account data is managed by the Directory System Agent (DSA), a protected process integrated with the LSA on the domain controller. Clients of the directory service are never given direct access to the data store. Any client wanting

access to directory information must use one of the supported Active Directory Service Interfaces (ADSI) to connect to the DSA and then search for, read, and write directory objects and their attributes.

Requests to access an object or attribute in the directory are subject to validation by Windows 2000 or Windows Server 2003 access control mechanisms. Like file and folder objects in the NTFS file system, objects in Active Directory are protected by Access Control Lists (ACLs) that specify who can access the object and in what way. Unlike files and folders, however, Active Directory objects have an ACL for each of their attributes. Thus, attributes for sensitive account information can be protected by more restrictive permissions than those granted for other attributes of the account.

The most sensitive information about an account is, of course, its password. Although an account object's password attribute stores an encryption key derived from a password, not the password itself, this key could be just as useful to an intruder. Therefore, access to an account object's password attribute is granted only to the account holder, never to anyone else, not even administrators. Only processes with Trusted Computer Base privilege—processes running in the security context of the LSA—are allowed to read or change password information.

In order to hinder an offline attack by someone with access to a domain controller's backup tape, an account object's password attribute is further protected by a second encryption using a system key. Administrators are given the option to choose where the system key is stored and which of several algorithms is used to encrypt password attributes. You can store the system key on removable media so that it can be safeguarded separately, or store it—protected by a dispersal mechanism—on the domain controller.

**Service Principal Names**

SPNs can only be registered by a domain administrator, with one exception—a system account can register an SPN for its computer account.

In Windows 2000, SPNs were canonicalized to the SAM account name—for example, Server1, Server1$. This caused issues when clients requested tickets for a service with a non-canonical name. The system was unable to detect the existing cached service ticket and would request a new service ticket for the service. In Windows Server 2003, on the other hand, the SPN requested is used without canonicalization.

SPNs are unique identifiers for services running on servers. Every service that will use Kerberos authentication needs to have an SPN set for it so that clients can identify the service on the network. If an SPN is not set for a service, then clients will have no way of locating that service. Without properly set SPNs, Kerberos authentication is not possible.

An SPN is registered in Active Directory under a user account as an attribute called **Service-Principal-Name**. The SPN is assigned to the account under which the service the SPN identifies is running. Any service can look up the SPN for another service. When a service wants to authenticate to another service, it uses that service's SPN to differentiate it from all of the other services running on that computer.

This is why SPNs are so crucial to constrained delegation (discussed later). When you set up a computer for delegation, one step of the process is to list the services on other computers that the computer is allowed to delegate to. This list forms a type of ACL. The services running on the other computers are identified by the SPNs that are issued to those services.

Additionally, in Windows Server 2003, KDCs will not issue a service ticket for an account that does not have an SPN. If a service account were simply a user account with a human-generated password, then that account would be more vulnerable to an offline dictionary attack. For an account without an SPN, the KDC will return KDC_ERR_S_PRINCIPAL_UNKNOWN. However, the context of the error will be KRB_ERR_MUST_USE_USER2USER, which has the description of "Server principal valid for user-to-user only."

In general, SPNs should be set at the time of account creation. Setting an SPN on an account is as important—from a security standpoint—as creating that account. Special accounts that were created for services are useless until an SPN is set on those accounts for whatever service will be running under them.

The following table lists, in alphabetical order, the built-in SPNs that are recognized for computer accounts.

**Built-in SPNs Recognized for Computer Accounts**

| SPN | SPN | SPN | SPN |
|---|---|---|---|
| alerter | http | policyagent | scm |
| appmgmt | ias | protectedstorage | seclogon |
| browser | iisad | rasman | snmp |

| | | | |
|---|---|---|---|
| cifs | min | remoteaccess | spooler |
| cisvc | messenger | replicator | tapisrv |
| clipsrv | msiserver | rpc | time |
| dcom | mcsvc | rpclocator | trksvr |
| dhcp | netdde | rpcss | trkwks |
| dmserver | netddedsm | rsvp | ups |
| dns | netlogon | samss | w3svc |
| dnscache | netman | scardsvr | wins |
| eventlog | nmagent | scesrv | www |
| eventsystem | oakley | schedule | |
| fax | plugplay | | |

These SPNs are recognized for computer accounts if the computer has a host SPN. Unless they are explicitly placed on objects, a host SPN can substitute for any of the above SPNs.

**Requirements for setting an SPN**

Because multiple services can run simultaneously under the same account, setting an SPN requires four unique pieces of information. These four pieces of information uniquely identify any service running on a network and can be used to mutually authenticate to any service.

For each SPN that is set, the following information is required:

- The type of service, formally called a service class. This enables you to differentiate between multiple services running under the same account.
- The account under which the service is running.
- The computer on which the service is running, including any aliases that point to that computer.
- The port on which the service is running (optional if the default port for the service of that type is used such as port 80 for HTTP).

The syntax of an SPN itself is service/hostname:port, where:

- Service is the service class of the SPN.
- Hostname is the computer to which the SPN belongs.
- Port is the port on which the service that the SPN is registered to runs.

**KDC key version numbers**

Users can change their passwords, and therefore their keys, at any time. This poses a problem for all existing credentials that are encrypted with the older key, because the credentials would no longer be usable if the user expected to use the new key. Although this is merely inconvenient to users—who can simply get new tickets—long-running services such as batch jobs would be interrupted if there was no one available to get a new ticket with the new key for the job. To address this problem, clients keep a copy of their older key for a period of time after switching to a new password.

When parties exchange encrypted communications, they need a way to determine which keys to use to decrypt the data they have received. If there is only one key, this is easy. However, when there are many keys, some of which correspond to the same account—such as when passwords are changed—then this task becomes difficult. One way to find the right key is to try all available keys, but this wastes time and computing resources, particularly if the client has many keys to try.

To streamline this process, encrypted data includes a header that has information including a sort of serial number that is associated with the key and the encryption type used to encrypt the data. When a client receives an encrypted communication, it includes information about whose key set to use, as well as which particular key version to use to decrypt the data. This way, even if a ticket is encrypted with an older key, the recipient can still decrypt it by looking at the key version number, finding the correct key version, and decrypting the ticket. Eventually, tickets with outdated keys expire, and the system gradually shifts to using the new key in all cases, creating no problems for the user.

New to Windows Server 2003 is the read-only attribute **ms-DS-KeyVersionNumber**. When a password is changed, the key version number is incremented and the current value can be queried using **ms-DS-KeyVersionNumber**. Windows 2000 KDCs will always issue tickets with a key version number equal to 1. Windows Server 2003 KDCs only keep track of the latest key version number for a principal account because there is no reason for the KDC to use outdated keys.

Key version numbers also enable Windows Server 2003 to provide more informative errors when a message cannot be decrypted. In general cases, when information can not be decrypted, the system provides the error KRB_AP_ERR_MODIFIED, meaning that the encrypted information is not accessible, but with no further information. For example, the data could have been modified in transit, the data could have been malformed initially, or there could be no available key. Each of these possible causes requires troubleshooting at different locations and using different techniques.

On the other hand, using key version numbers can enable the system to communicate the source of the failure. For example, if the client has no key to match the key version number included in the header, the system responds with the error KRB_AP_ERR_BADKEYVER. This error message specifically informs the client that the appropriate key version is not available. The client can then resolve the failure by resetting the password, which resets the keys.

The key version number is generated each time a password is changed on long-lived keys—such as the key that results from a user password, the keys for computer accounts, or the keys for domain trusts. Key version numbers also are provided because they are required for compliance with the Kerberos RFC (1510).

Key version numbers are constructed from the update logon timestamp attribute which is updated in the SAM when a password is changed. The update logon timestamp is a 32-bit value that identifies the version of the keys generated from the password. Key version number functionality based on update logon timestamps is only available in domains that have Windows Server 2003 domain functionality.

**Kerberos Realms and Active Directory Domains**
With the release of Windows 2000, Active Directory domains became the implementation of Kerberos realms.

The RFC 1510 describes Kerberos realms and their organization as follows:

"The Kerberos protocol is designed to operate across organizational boundaries. A client in one organization can be authenticated to a server in another. Each organization wishing to run a Kerberos server establishes its own 'realm.'"

"Realms are typically organized hierarchically. Each realm shares a key with its parent and a different key with each child. If an inter-realm key is not directly shared by two realms, the hierarchical organization enables an authentication path to be easily constructed. If a hierarchical organization is not used, it might be necessary to consult some database in order to construct an authentication path between realms."

"Although realms are typically hierarchical, intermediate realms might be bypassed to achieve cross-realm authentication through alternative authentication paths (these might be established to make communication between two realms more efficient). It is important for the end-service to have information about which realms were transited when deciding how much faith to place in the authentication process. To facilitate this decision, a field in each ticket contains the names of the realms that were involved in authenticating the client."

It is clear from the RFC's discussion that the concept of Kerberos realms has similarities—such as organizational boundaries and trusts—to attributes of the Windows NT domains. Active Directory domains use inter-realm-key-based trusts and a hierarchical design. These inter-realm (also called inter-domain) keys are the basis of transitive trust in Active Directory domains.

### Note

- In literature documenting Windows operating systems, the term "domains" might be used in place of place of the term "realms." In such cases, the term "realms" will be used to denote non-Windows Kerberos realms in the environment.

In Kerberos authentication, Active Directory domains fulfill all the requirements of Kerberos realms.

**Kerberos Processes and Interactions**
This subsection will cover the following topics:

- Kerberos V5 exchange and message summary
- Windows Server 2003–based examples:
    - Local logon
    - Domain logon
    - Single domain authentication
    - Simple cross-domain authentication
    - Cross-realm authentication and shortcut trusts
    - Untrusted domain
    - Delegation
- Kerberos V5 message details:
    - Authentication service exchange
    - Ticket-granting service exchange
    - Client/server authentication exchange
    - Credentials message
- Required technologies for Active Directory domains:
    - Kerberos and IP transport: UDP/TCP
    - DNS name resolution

**Kerberos V5 Exchange and Message Summary**
The RFC standard Kerberos version 5 authentication protocol communication sequences consist of six (five required and one optional) messages. These messages compose three types of exchanges (also known as sub-protocols), which are examined more closely later in this section.

### Kerberos Exchange and Message Summary




**The Authentication Service Exchange**
**1. Kerberos authentication service request (KRB_AS_REQ)**
The client contacts the Key Distribution Center's authentication service for a short-lived ticket (a message containing the client's identity and—for Windows clients—SIDs) called a ticket-granting ticket (TGT). This happens at logon.

**2. Kerberos authentication service response (KRB_AS_REP)**

The authentication service (AS) constructs the TGT and creates a session key the client can use to encrypt communication with the ticket-granting service (TGS). The TGT has a limited lifetime. At the point that the client has received the TGT, the client has not been granted access to any resources, even to resources on the local computer.

Why use a TGT? Could the AS simply issue a ticket for the target server? Yes, but if the AS issued tickets directly, the user would have to enter a password for every new server/service connection. Issuing a TGT with a short lifespan (typically 10 hours) gives the user a valid ticket for the ticket-granting service, which in turn issues target-server tickets. The TGT's main benefit is that the user only has to enter a password once, at logon.

**The Ticket-Granting Service Exchange**
**3. Kerberos ticket-granting service request (KRB_TGS_REQ)**
The client wants access to local and network resources. To gain access, the client sends a request to the TGS for a ticket for the local computer or some network server or service. This ticket is referred to as the service ticket or session ticket. To get the ticket, the client presents the TGT, an authenticator, and the name of the target server (the Server Principal Name or SPN).

**4. Kerberos ticket-granting service response (KRB_TGS_REP)**
The TGS examines the TGT and the authenticator. If these are acceptable, the TGS creates a service ticket. The client's identity is taken from the TGT and copied to the service ticket. Then the ticket is sent to the client.

### Note

- The TGS cannot determine if the user will be able to get access to the target server. It simply returns a valid ticket. Authentication does not imply authorization.

**The Client/Server Exchange**
**5. Kerberos application server request (KRB_AP_REQ)**
After the client has the service ticket, the client sends the ticket and a new authenticator to the target server, requesting access. The server will decrypt the ticket, validate the authenticator, and for Windows services, create an access token for the user based on the SIDs in the ticket.

**6. Kerberos application server response (optional) (KRB_AP_REP)**
Optionally, the client might request that the target server verify its own identity. This is called mutual authentication. If mutual authentication is requested, the target server will take the client computer's timestamp from the authenticator, encrypt it with the session key the TGS provided for client-target server messages, and send it to the client.

**Local Logon Example**
Users can log on to computers with either a domain account or an account local to the computer. When the user logs on with an account local to the computer, the user's credentials are authenticated using the local account database. Because the local computer does not act as a KDC, and because local logon does not require network access (for example, to contact a KDC), local authentication uses NTLM to authenticate the account.

### Local Logon Sequence




1. After the Graphical Identification and Authentication (GINA) dynamic-link library has collected the user's logon information, it passes that information to the LSA for authentication.

2. The LSA simply passes the information on to the SSPI.

3. The SSPI does not determine whether the user is logging on locally or through a domain account. Because the Kerberos SSP is the default authentication provider, the authentication request is first passed to the Kerberos SSP.

4. The Kerberos SSP verifies that the logon target name is the same as the local computer name. If so, the Kerberos SSP returns an error message to the SSPI that no logon servers are available, just as if it had checked the network for a KDC and found none.

5. The SSPI now sends the request to the next security provider, NTLM. The authentication process used by the NTLM SSP is identical to versions used in Windows NT.

**Domain Logon Example**

If the user is a member of a domain, then the account database is on a domain controller. If you are logging on to one of the domain controllers for the domain you belong to, then the process it the same as for local logon as described above. If instead you are logging on to a member server or a domain controller of another domain, then Kerberos authentication will used if possible.

Users never directly access the system. The system impersonates the user and accesses resources based on the resource permissions granted to the user. This is a standard security feature of all versions of Windows NT, Windows 2000, and Windows Server 2003. The user logging on to a workstation never directly accesses even local resources. The local system impersonates the user just as a remote system would.

Because the local system relies on impersonation, the user who logs on to the domain must be authenticated as a valid local user. The only new component added to the domain logon process is Kerberos authentication prior to building an access token.

The authentication process for a domain user to access their computer is very similar to the process used to authenticate access to network resources—that is, the user must obtain a valid ticket for the local workstation.

- The Kerberos client requests and then receives a TGT from the KDC.

- The Kerberos client uses the TGT to request and then receive a service ticket for the local workstation from the KDC.

- The service ticket for a network resource would be encrypted with the system or service key depending on whether the resource is a system or service. The workstation has a system key created when the computer joined the domain. The service ticket for the workstation is encrypted with this key.

- The local LSA builds an access token from the credentials contained in the service ticket and then grants or denies the user access.

**How Logging on the Workstation Works**

Suppose the user has a network account in the domain tailspintoys.com. The user's workstation also belongs to tailspintoys.com. The user logs on to the network by pressing CTRL+ALT+DEL, which is the Secure Attention Sequence (SAS) on computers with a standard Windows configuration.

In response to the SAS, the workstation's Winlogon service switches to the logon desktop and calls the GINA DLL, a component loaded in Winlogon's process. The GINA is responsible for collecting logon data from the user, packaging it in a data structure, and sending everything to the LSA for verification. Other parties can develop replacement GINAs, but in this case Winlogon has loaded the standard component supplied with the operating system, MSGINA.DLL. Winlogon calls it, and the GINA displays the standard logon information dialog box.

The user types a user name and password. From the domains in a drop-down list, the user selects tailspintoys.com. When the user clicks OK to dismiss the dialog box, MSGINA returns the logon information to Winlogon. Winlogon sends the information to the LSA for validation by calling LsaLogonUser. This procedure is standard for any user logon regardless of the authentication method used.

At this point, the LSA begins to use the Kerberos V5 authentication protocol.

**Kerberos Keys**




A user key is derived from a password. After it receives a data structure with the user's logon data, the LSA converts the plaintext password to a cryptographic key by passing the text of the password through a cryptographic function. (All implementations of the Kerberos version 5 authentication protocol must support DES-CBC-MD5. Other algorithms are permissible.) The result of the cryptographic function is the user key.

The LSA saves the user key in the user credentials cache, where it can be retrieved when needed for TGT renewal or for NTLM authentication to servers that are not capable of Kerberos authentication.

In order to validate the user's logon information and set up the user's logon session on the computer, the LSA must obtain:

- A TGT good for admission to the ticket-granting service.
- A service ticket good for admission to the computer.

The LSA obtains these tickets by working through the Kerberos SSP, which exchanges messages directly with the KDC in tailspintoys.com. (The workstation has located a KDC through a DNS query. Every domain controller is a KDC and registers its role with DNS.)

The TGT is obtained in an authentication service (AS) exchange; the service ticket is obtained in a ticket-granting service (TGS) exchange.

**Obtaining a TGT in the authentication service exchange**

### Message 1: Authentication Service Request




The Kerberos client on the workstation sends the message KRB_AS_REQ to the KDC in tailspintoys.com.

The message includes:

- The user principal name.
- The name of the account domain.
- Pre-authentication data encrypted with the user's key derived from the user's password.

**Retrieving the user key**

The KDC gets its copy of the user key from the user's record in its account database. When it receives a request from the Kerberos client on the user's workstation, the KDC searches its database for the user, pulls up the account record, and takes the user key from a field in the record.

This process—computing one copy of the key from a password, fetching another copy of the key from a database—actually

takes place only once, when a user initially logs on to the network. Immediately after accepting the user's password and deriving the user's long-term key, the Kerberos client on the workstation requests a service ticket and TGS session key that it can use in subsequent transactions with the KDC during this logon session.

**Verifying the user's identity**

The KDC decrypts the pre-authentication data and evaluates the timestamp inside. If the timestamp passes the test, the KDC can be assured that the pre-authentication data was encrypted with the user key and thus verify that the user is genuine.

After it has verified the user's identity, the KDC creates credentials that the Kerberos client on the workstation can present to the ticket-granting service. For more information about how domain controllers create credentials in a Windows environment, see "How Access Tokens Work" in the Access Tokens Technical Reference.

### Message 2: Authentication Service Reply




The KDC replies with KRB_AS_REP containing a service ticket for itself. This special service ticket is called a ticket-granting ticket (TGT). Like an ordinary service ticket, a TGT contains a copy of the session key that the service (in this case the KDC) will use in communicating with the user. The message that returns the TGT to the user also includes a copy of the session key that the user can use in communicating with the KDC. The TGT is encrypted in the KDC's long-term key. The user's copy of the session key is encrypted in the user's long-term key.

### KRB_AS_REP Message Contents




The message includes:

- A TGS session key for the user to use with the TGS, encrypted with the user key derived from the user's password.
- A TGT for the KDC in tailspintoys.com, encrypted with the TGS key.

  The TGT includes:

  - A TGS session key for the KDC to use with the user.
  - Authorization data for the user.

  The authorization data includes:

  - The SID for the user's account.
  - SIDs for security groups which the user is a member of in the domain tailspintoys.com.

- SIDs for universal groups in the enterprise that include either the user or one of the domain groups the user is a member of.

When the client receives the KDC's reply to its initial request, the client uses its cached copy of the user key to decrypt its copy of the session key. It can then discard the user key derived from the user's password, for it is no longer needed. In all subsequent exchanges with the KDC, the client uses the TGS session key. Like any other session key, this key is temporary, valid only until the TGT expires or the user logs off. For that reason, the TGS session key is often called a logon session key.

From the client's point of view, a TGT is just another ticket. Before the client attempts to connect to any service, the client first checks the user credentials cache for a service ticket to that service. If it does not have one, it checks the cache again for a TGT. If it finds a TGT, the LSA fetches the corresponding logon session key from the cache, uses this key to prepare an authenticator, and sends both the authenticator and the TGT to the KDC, along with a request for a service ticket for the service. In other words, gaining admission to the KDC is no different from gaining admission to any other service in the domain—it requires a session key, an authenticator, and a ticket (in this case, a TGT).

From the KDC's point of view, TGTs enable the KDC to avoid the performance penalties of looking up a user's long term key every time the user requests a service. The KDC looks up the user key only once, when it grants an initial TGT. For all other exchanges with this user, the KDC can decrypt the TGT with its own long-term key, extract the logon session key, and use that to validate the user's authenticator.

**Obtaining a service ticket in the ticket-granting service exchange**

### Message 3: Ticket-Granting Service Request




The Kerberos client sends the message KRB_TGS_REQ to the KDC in tailspintoys.com.

The message includes:

- The name of the target computer.
- The name of the target computer's domain.
- The user's TGT.
- An authenticator encrypted with the session key the user shares with the KDC.

### Message 4: Ticket-Granting Service Reply




The KDC responds to the user's request to connect to a server by sending both copies of the session key to the user. The user's copy of the session key is encrypted with the key that the KDC shares with the user. The workstation's copy of the session key is embedded, along with information about the user, in a data structure called a service ticket. The entire structure is then encrypted with the key that the KDC shares with the user. The ticket—with the user's copy of the session key safely inside—becomes the workstation's responsibility to manage.

**KRB_TGS_REP Message Contents**




The KRB_TGS_REP message includes:

- A session key for the user to share with the computer encrypted with the session key the user shares with the KDC.
- The user's service ticket to the computer, encrypted with the computer's secret key.

  The service ticket includes:

  - A session key for the computer to share with the user.
  - Authorization data copied from the user's TGT.

When the Kerberos client receives the KDC's reply, it extracts the ticket and the user's copy of the session key, putting both in the user credentials cache (located in volatile memory, not on disk).

**Note**

- The KDC is simply providing a ticket-granting service. It does not keep track of its messages to make sure they reach the intended address. No harm will be done if the KDC's messages fall into the wrong hands. Only someone who knows the user's secret key can decrypt the user's copy of the session key. Only someone who knows the client's secret key can read what is inside the ticket.

**Getting the User's Credentials to Winlogon**




After credentials reach the workstation, the Windows Server 2003 access token creation process is the same as that of Windows NT versions. The LSA on the workstation receives the user's service ticket, decrypts the service ticket with the system key stored in its credentials cache, and then extracts the authorization data. The privilege attribute certificate (PAC) is taken from the service ticket and used to create the user's access token. The LSA then queries the local SAM database to discover whether the user is a member of any security groups local to the computer, and whether memberships in those groups grant the user any special rights on the local computer. It adds any SIDs returned by this query to the list taken from the ticket's authorization data. The entire list is then used to build an access token, and a handle to the access token is returned to Winlogon, along with an identifier for the user's logon session and confirmation that the logon information was valid.

Winlogon creates a window station and several desktop objects for the user, attaches the user's access token, and starts the shell process the user will use to interact with the computer. The user's access token is subsequently inherited by any application process that the user starts during the logon session.

When the user logs off, the credentials cache is flushed and all service tickets—as well as all session keys—are destroyed.

**Single Domain Authentication Example**

The process used to access resources and services on another system works basically the same as the logon example above except that the service ticket will need to be sent to the remote system. Except for the Kerberos authentication, security access checks are nearly identical with those of Windows NT 4.0 and earlier versions.

In the following example, both the user and service are in tailspintoys.com, the user has already logged on, and the user has requested and received a ticket for the workstation. The client wants to access \\server\shared folder to read a file. The process follows this sequence:

- Client and server negotiate a security package to use for authentication. They choose Kerberos.
- The client sends the service ticket (containing user credentials in the PAC) to the server.
- If the server accepts the ticket (that is, is able to decrypt the ticket with its secret key), then the server creates an access token for the user based on the PAC.
- The client redirector sends a server message block (SMB) message requesting file access.
- Server security compares file permissions with the user's credentials and grants or denies access.

## After Workstation Logon




When the user attempts to use the service on the remote server, the user has already logged on to the local workstation. Thus, the user's credentials cache will have a TGT and a session key for the domain the user's account resides in. This TGT is used to request a ticket for the service.

**Message 3: KRB_TGS_REQ**




Because the TGT will expire within a few hours, the Kerberos client must check the expiration time before requesting additional tickets. If the TGT is expired, then the Kerberos client must send a KRB_AS_REQ to get another TGT.

The Kerberos client on the user's workstation requests credentials for the service by sending the KDC a Kerberos ticket-granting service request (KRB_TGS_REQ). This message includes the user's name, an authenticator encrypted with the user's logon session key, the TGT obtained in the AS exchange from workstation logon, and the name of the service for which the user wants a ticket.

**Note**

- The KRB_TGS_REQ message will typically be sent to the same domain controller that issued the TGT. However, TGS requests can be made to any domain controller. Thus, if the original KDC becomes unavailable, the client can discover a new KDC through a DNS query and then send a KRB_TGS_REQ there.

When the KDC receives KRB_TGS_REQ, it decrypts the TGT with its own secret key, extracting the user's TGS session key (logon session key). It uses the session key to decrypt the authenticator and evaluates that. If the authenticator passes the test, the KDC extracts the user's authorization data from the TGT and creates another session key for the client to use with the service. The KDC encrypts one copy of this new session key with the user's TGS session key. It embeds another copy of the session key in a ticket, along with the user's authorization data, and encrypts the ticket with the service's key. The KDC then sends these credentials back to the client in a Kerberos ticket-granting service reply (KRB_TGS_REP).

**Message 4: KRB_TGS_REP**




When the Kerberos client receives the reply, it uses the user's TGS session key to decrypt the session key to use with the service, and stores the key in its credentials cache. Then it extracts the ticket to the service and stores that in its cache.

**Message 5: Kerberos Application Request (KRB_AP_REQ)**




The Kerberos client on the user's workstation requests service from the service by sending the service a Kerberos application request (KRB_AP_REQ).

**KRB_AP_REQ Message Contents**




This message contains:

- An application option flag indicating whether to use session key. (The setting of this flag is one of the options in configuring Kerberos applications. The user is never asked.)
- An application option flag indicating whether the client wants mutual authentication.
- The service ticket obtained in the TGS exchange.
- An authenticator encrypted with the session key for the service.

**Message 6: KRB_AP_REP (Optional)**




The service receives KRB_AP_REQ, decrypts the ticket, and extracts the user's authorization data and the session key. The service uses the session key to decrypt the user's authenticator and then evaluates the timestamp inside. If the authenticator passes the test, the service looks for a mutual authentication flag in the client's request. If the flag is set, the service uses the session key to encrypt the time from the user's authenticator and returns the result in a Kerberos application reply (KRB_AP_REP). If the flag is not set, then no response is needed.

When the client on the user's workstation receives KRB_AP_REP, it decrypts the service's authenticator with the session key it shares with the service and compares the time returned by the service with the time in the client's original authenticator. If the

times match, the client knows that the service is genuine.

**Building an access token from a Kerberos service ticket**
After credentials reach the server, the Windows Server 2003 token creation process is the same as that of earlier Windows versions. After the server authenticates the client using Kerberos authentication, the PAC is taken from the service ticket and used to create the user's access token.

**Remote resource authorization**
After the service has an access token for the user, it still needs to verify that the user is authorized to do what the user wants to do. Except for the Kerberos authentication, security access checks are nearly identical with those of Windows NT 4.0 and earlier versions:

1. The client redirector sends an SMB message requesting file access.

2. Server security compares file permissions with the user's credentials and grants or denies access.

3. If access is granted, the connection proceeds.

During the connection, the session key can be used to encrypt application data, or the client and server can share another key for this purpose.

**Simple Cross-Realm Authentication and Examples**
The functions of the KDC are divided into two distinct services: an authentication service whose job is to issue TGTs, and a ticket-granting service whose job is to issue service tickets. This division of labor enables the Kerberos protocol to operate across domain boundaries. A client can get a TGT from the authentication service of one domain and use it to get service tickets from the ticket-granting service of another domain.

**The Basic TGS Referral Process**
Clients typically need to request access to services outside of a single Kerberos realm. Whether in a pure Active Directory forest or in a mixed Kerberos environment, the referral process is the same. When a client requests a service ticket for a server in a remote Kerberos realm, the request is sent to the KDC in the client account's realm. The KDC determines that the server is in another realm, so it cannot issue a service ticket. This can only be done by a KDC in the target server's realm. So, instead of issuing the service ticket, the KDC in the client account's realm issues a TGS referral.

The TGS referral message is a normal KRB_TGS_REP message containing:

- A TGT for another trusted realm.

- Information indicating the KRB_TGS_REP is a referral to another TGS.

The client recognizes that the KRB_TGS_REP is a referral message and sends the TGT in a new KRB_TGS_REQ message to the remote realm to get the service ticket it needs.

### Cross-Realm Authentication Between Two Realms




To understand how cross-realm authentication works, first consider the simplest case: a network with only two realms, East and West. If administrators for these realms are members of the same organization, or if for some other reason they are willing to treat the other realm's users as their own, they can enable authentication across realm boundaries simply by sharing an inter-realm key. (In Active Directory domains this happens automatically when two domains establish a trust relationship.) After the inter-realm key is shared, the ticket-granting service of each realm is registered as a security principal with the other realm's KDC. As a result, the ticket-granting service in each realm can treat the ticket-granting service in the other realm as just another service, something for which properly authenticated clients can request and receive service tickets.

When a user with an account in West wants access to a server with an account in East, the process is:

1. The Kerberos client on the user's workstation sends a request for a service ticket to the ticket-granting service in the user

account's realm, West.

2. The ticket-granting service in West determines that the desired server is not a security principal in its realm, so it replies by sending the client a referral ticket—a TGT encrypted with the inter-realm key that the KDC in West shares with the KDC in East.

3. The client uses the referral ticket to prepare a second request for a service ticket, and this time sends the request to the ticket-granting service in the server account's realm, East.

4. The ticket-granting service in East uses its copy of the inter-realm key to decrypt the referral ticket. If decryption is successful, it sends the client a service ticket to the desired server in its domain.

**Three-Realm Referral Example**

The referral process is more complicated on networks with more than two realms. In theory, the KDC in each realm could establish a direct link to the KDC in every other realm on the network, in each case sharing a different inter-realm key. In practice, however, the number and complexity of these relationships could become unmanageable, especially on a large network. The Kerberos protocol solves the problem by making direct links unnecessary. A client in one realm can get a ticket to a server in another realm by traveling a referral path through one or more intermediate realms.

**Three-Realm Referral**




For example, consider a network with three realms, east.tailspintoys.com, west.tailspintoys.com, and tailspintoys.com. The KDC in east.tailspintoys.com does not share an inter-realm key with the KDC in west.tailspintoys.com, but both KDCs do share inter-realm keys with tailspintoys.com. In this case, when a user with an account in west.tailspintoys.com wants access to a server with an account in east.tailspintoys.com, the referral path begins at the KDC for the user account's realm, west.tailspintoys.com, passes through an intermediate realm, tailspintoys.com, and ends at the KDC for the server account's realm, east.tailspintoys.com. The client must send its request for a service ticket three times, to three different KDCs. The three TGS exchanges are:

1. The client requests the KDC for west.tailspintoys.com to give it a ticket to the server in east.tailspintoys.com.

   The KDC for west.tailspintoys.com replies by sending the client a referral ticket to the KDC for tailspintoys.com. This ticket is encrypted in the inter-realm key that west.tailspintoys.com shares with tailspintoys.com.

2. The client requests the KDC for tailspintoys.com to give it a ticket to the server in east.tailspintoys.com.

   The KDC for tailspintoys.com replies by sending the client a referral ticket to the KDC for east.tailspintoys.com. This ticket is encrypted in the inter-realm key that tailspintoys.com shares with east.tailspintoys.com.

3. The client requests the KDC for east.tailspintoys.com to give it a ticket to the server in east.tailspintoys.com.

   The KDC for east.tailspintoys.com can decrypt the TGT from tailspintoys.com. The TGT has the user's credentials and

verifies that they have not been tampered with because they were encrypted with a key only the two KDC have. The KDC for east.tailspintoys.com sends back a ticket for the server.

For a detailed look at the complete authentication process, see the next example, "Cross-Realm Authentication and Shortcut Trusts."

**Cross-Realm Authentication and Shortcut Trusts Example**
What happens if there are multiple trusted Kerberos realms between the client's logon realm and the target server's realm? In this case, the client receives a TGS referral for the nearest trusted realm along the shortest path to the target realm. The client must send the KRB_TGS_REQ message requesting a service ticket for the target server to each intervening realm. A KDC in each intervening realm issues another TGS referral to the client until the client has a TGT for a KDC in the target realm. The following example describes an Active Directory domain environment.

**Three Domain Cross-Realm Authentication Example**
Suppose the user logs on at a workstation in west.tailspintoys.com. The user wants to connect to a shared file on a server in east.tailspintoys.com. Two steps occur before the cross-realm authentication steps:

- The user uses Explorer to request a connection to \\server\shared folder.
- DNS resolves that the target server is located in east.tailspintoys.com.

**Cross-Realm Authentication: Three Domains**




To authenticate the user's identity to the server, the user needs a service ticket. Only the KDC in east.tailspintoys.com can issue this ticket. Although the user needs a TGT for that KDC, its initial ticket request needs to go to the local KDC for west.tailspintoys.com. The Kerberos SSP on the user's workstation must obtain a TGT for the server's domain and then use the TGT to get a ticket for the server.

This involves a referral process that proceeds as follows:

1. The user's Kerberos client sends a KRB_TGS_REQ message to a local KDC in west.tailspintoys.com, specifying server@east.tailspintoys.com as the target.

   The message includes:

   - The name of the target computer.
   - The name of the target computer's domain.
   - A TGT for admission to the KDC in west.tailspintoys.com.
   - An authenticator encrypted with the TGS session key the user shares with that KDC.

The Kerberos client compares the location with its own workstation's domain. Because they are different, the client sets a flag in the KDC Options field of the KRB_TGS_REQ message for NAME_CANONICALIZE, which indicates to the KDC that server might be in another realm.

2. The user's local KDC (for west.tailspintoys.com) receives the KRB_TGS_REQ and sends back a referral ticket for tailspintoys.com. The local KDC:

   a. Notes the flag setting.

   b. Verifies the user's TGT.

   c. Looks up the path to the server—in this case, a computer object in east.tailspintoys.com.

   d. Builds a tree of known domains until it has a structure including east.tailspintoys.com.

   e. Queries LSA for a trust list.

   f. Constructs the shortest path between its domain and east.tailspintoys.com based on the domain tree and the trust list.

   g. Constructs a new TGT for the user.

      Unless there is a direct trust between west.tailspintoys.com and east.tailspintoys.com (which there is not in this example), the TGT is issued for the next intervening domain along the shortest path to east.tailspintoys.com. In this example, the next intervening domain is tailspintoys.com.

      The TGT also has a referral flag set, so the KDC will be informed that the KRB_TGS_REQ is coming from another realm. This flag also tells the KDC to fill in the Transited Realms field.

   h. Constructs a KRB_TGS_REP message containing these four items:




   - A TGT for tailspintoys.com, encrypted with the inter-realm key shared between west.tailspintoys.com and tailspintoys.com.

   - A session key for the user to share with the KDC in tailspintoys.com, encrypted with the user's TGS session key.

   - A pre-authentication field data type of KERB-PA-SERV-REFERRAL, which tells the user's Kerberos client that the enclosed ticket is a referral TGT.

- Pre-authentication data (the realm name for the referral — in this case, tailspintoys.com).

3. The user's Kerberos client receives the KRB_TGS_REP and sends KRB_TGS_REQ to the KDC in tailspintoys.com. The Kerberos client:

   a. Notes that this is a referral and gets the referral realm (tailspintoys.com).

   b. Sends a KRB_TGS_REQ message to a KDC in tailspintoys.com with the TGT, the NAME_CANONICALIZE flag set, and a request for a service ticket to server.

   The message includes:

   - The name of the target computer.
   - The name of the target computer's domain.
   - A TGT for admission to the KDC in tailspintoys.com.
   - An authenticator encrypted with the session key the user shares with that KDC.

4. A KDC in tailspintoys.com goes through the same process as the KDC in west.tailspintoys.com, returning a TGT for east.tailspintoys.com in a TGS referral to the user.




5. The user's Kerberos client sends a KRB_TGS_REQ along with the TGT it received from the tailspintoys.com KDC to a KDC in east.tailspintoys.com.

6. The KDC in east.tailspintoys.com replies with KRB_TGS_REP.

   The message includes:

   - A session key for the user to share with the server, encrypted with the session key the user shares with the KDC in east.tailspintoys.com.
   - A service ticket for admission to the server encrypted with the server key shared with the KDC.

   The service ticket includes:

- A session key for the server to share with the user.
- Authorization data copied from the user's TGT, plus data for the local domain, east.tailspintoys.com.

The authorization data includes:

- The SID for the user's account.
- SIDs for groups in west.tailspintoys.com that include the user.
- SIDs for universal groups that include either user or one of the user's groups in west.tailspintoys.com.
- SIDs for groups in east.tailspintoys.com that include the user, one of the user's groups in west.tailspintoys.com, or one of their universal groups.

7. The user's Kerberos client sends the server a KRB_AP_REQ message.

   The message includes:

   - The user principal name.
   - A service ticket to the server.
   - An authenticator encrypted with the session key the user shares with the server.

**Note**

- If the user had asked for mutual authentication, then the server would reply with KRB_AP_REP. The reply message would include an authenticator encrypted with the session key the server shares with the user.

The same process takes place when a user logs on to a remote domain and needs a service ticket for the local workstation. The KDC in the logon domain issues a referral to the next realm between it and the workstation's realm, and so on, until the client can request a service ticket from the KDC in the workstation's realm.

**Effects of Shortcut Trusts**

You can optimize the inter-realm authentication process in your Active Directory forest by shortening the trust path with shortcut trusts. In this example, the user has an account in it.west.tailspintoys.com. The user has logged on to the domain account and has collected a TGT for the KDC in it.west.tailspintoys.com. Suppose the user needs access to a document stored in a public folder on a server in the domain it.east.tailspintoys.com.

**Forest Configuration**




Without shortcut trusts, the referral path to retrieve a service ticket would be:

1. it.west.tailspintoys.com

2. west.tailspintoys.com
3. tailspintoys.com
4. east.tailspintoys.com
5. it.east.tailspintoys.com

The following illustration shows how the referral path would change with the addition of a shortcut trust.

**Forest Configuration with Shortcut Trusts**




If Shortcut Trust 1 is created, then the referral path to retrieve a service ticket would become:

1. it.west.tailspintoys.com
2. west.tailspintoys.com
3. east.tailspintoys.com
4. it.east.tailspintoys.com

**Do shortcut trusts always affect the shortest referral path?**
If the shortest path to a realm would be to use an explicit trust that is positioned in the tree below the current domain, that path will not be used. This is because the shortest-path algorithm can only make use of explicit trusts that are at or above the current position in the tree.

If only Shortcut Trust 2 is created then the referral path to retrieve a service ticket would remain:

1. it.west.tailspintoys.com
2. west.tailspintoys.com
3. tailspintoys.com
4. east.tailspintoys.com
5. it.east.tailspintoys.com

This shortcut trust would only shorten the path from test.it.west.tailspintoys.com to east.tailspintoys.com and it.east.tailspintoys.com. But from test.it.west.tailspintoys.com and east.tailspintoys.com it would reduce the number of exchanges from five to two.

**Untrusted Domain Examples**

To use Kerberos authentication in a remote untrusted domain, the user must supply credentials as part of the connection request. Otherwise, the KRB_TGS_REQ message fails, and the system tries the next security provider, typically NTLM.

**Client Does Not Use Alternative Credentials**
Suppose user 1 is logged on to west.tailspintoys.com and wants to access a shared file on user 2's Windows Server 2003 server in sales.contoso.com, an untrusted domain.

1. User 1 opens Windows Explorer and simply tries to map a network drive to \\User2\shared folder without providing the credentials with which to connect.

2. The KRB_TGS_REQ goes to user 1's local KDC just as it would if sales.contoso.com were a trusted domain.

3. When the KDC attempts to build a domain tree and trust list, the attempt fails because there is no trust.

4. User 1 might or might not get an error message because the system automatically reverts to the next security provider. If NTLM is available, it is used. If NTLM authentication is successful, user 1 might connect, but Kerberos authentication will not have been used.

**Client Uses Alternative Credentials to Send KRB_AS_REQ Directly to Untrusted KDC**
User 1 can enter credentials. The Kerberos client responds differently when the user includes new credentials.

- A different ticket cache is created for each different set of credentials.

  **Note**

  - If user 1 connects to an untrusted resource using different credentials (for example, Administrator), a new ticket cache is created for these credentials. The tickets in this cache are not currently viewable with either Kerberos Tray or Kerberos List.

- If user 1 has used the credentials recently, a ticket cache might already exist. If so, the client checks for an existing TGT for the untrusted domain. It cannot use an existing TGT from user 1's normal ticket cache.

- The client sends a KRB_AS_REQ message to get a TGT from the KDC in the target server's realm. If the supplied credentials are correct, this request is processed normally, without the need for any TGS referrals.

**Delegation Examples**
**The Need for Delegation**
A client might need to let an application or a service connect to other servers or services on its behalf. A client might use a front-end server, for example, that then needs to authenticate with a back-end server. The front-end server needs to authenticate to the back-end server with the client's credentials, because if it authenticated under its own service account, it would have different authorization than the user.

A multi-tier client/server application presents a special situation for the Kerberos protocol. In this kind of application, a client might connect to a server that must itself connect to a second server on the back end. For this to happen, the first server must have a ticket to the second. Ideally, the ticket should limit the first server's access on the second server to whatever the client—instead of the server—is authorized to do.

The Kerberos protocol deals with this situation through a mechanism called delegation of authentication. Essentially, the client delegates authentication to a server by telling the KDC that the server is authorized to represent the client. The concept is similar to the concept of impersonation in Windows NT.

Traditionally, the Kerberos protocol enabled delegation to be accomplished using one of two mechanisms:

- **Proxy tickets.** The client can get a ticket for the back-end server and then give it to the front-end server. Tickets obtained in this way—by a client for a proxy—are called proxy tickets. The difficulty with proxy tickets is that the client must have information about the name of the back-end server.

- **Forwarded tickets.** The client can give the front-end server a TGT that the server can use to request tickets as needed. Tickets obtained in this way—with credentials forwarded by a client—are called forwarded tickets. This second method overcomes the difficulty (that the client must have information about the name of the back-end server) presented by the

first method.

Windows Server 2003 supports two new extensions for delegation:

- **Protocol transition extension.** The protocol transition extension allows a kerberized service (that is, a service designed to use Kerberos authentication features) to obtain a Kerberos service ticket on behalf of a Kerberos principal to the service itself without requiring the principal to initially authenticate to the KDC with a credential.

- **Constrained delegation extension.** The constrained delegation extension allows a kerberized service to obtain service tickets (under the delegated user's identity) to a subset of other services after it has been presented a service ticket obtained either through KRB_TGS_REQ or the protocol transition extension.

Whether the KDC will allow clients to obtain proxy tickets or forwardable TGTs is a matter for Kerberos policy.

### Note

- Technically, both proxying and forwarding are proxies, so different Kerberos documents might discuss both types as proxy. In other documents, forwarding the user's identity is also known as delegation. In many Microsoft documents, including the *Microsoft Windows 2000 Server Resource Kit Distributed Systems Guide*, both proxying and forwarding are referred to as types of delegation. The terminology is not critical, but you should be aware that terms are used differently by different organizations.

#### What Delegation Does for the Client

Delegation of authentication allows the client to send its identity in the form of a Kerberos ticket to the front-end server. The front-end server can then impersonate the client and authenticate with the back-end server as if the front-end server were the client.

Delegation is not limited to a single pair of a front-end and a back-end server. A client can delegate its identity to a service that can then authenticate with any number of back-end services. It is also possible for the client to delegate its identity to Service A, and for Service A to in turn delegate the client's identity to Service B, for Service B to delegate the client's identity to Service C, and so on.

The simpler case—in which the client delegates its identity to a front-end server so that the front-end server can authenticate in the client's security context to a back-end server—is more common.

#### Delegation of Authentication

In Windows NT, a service can impersonate clients when accessing resources on the computer where the service process is running. In Windows 2000 and Windows Server 2003, a service can use the Kerberos protocol's delegation feature to impersonate clients on other computers, as well. For delegation of authentication to work, three conditions must be met:

- The computers that host the client process, the service process, and the processes for back-end services must all be running Windows 2000, Windows XP, or Windows Server 2003 in an Active Directory domain.

- The client's account must be enabled for delegation.

- The service's account must be enabled for delegation.

#### To configure a client's account for delegation

1. Right-click the object that represents the user in Active Directory Users and Computers.

2. Click **Properties**, and then click the **Account** tab.

3. In the **Account options** list, be sure that the **Account is sensitive and cannot be delegated** option is not selected.

#### For accounts in domains that are Windows 2000 mixed, Windows 2000 native, or Windows Server 2003 interim functional level

To configure a service's account for delegation, determine whether the service runs under the host computer's Local System account or under its own domain user account.

**To configure a service's account for delegation if the service runs under the Local System account**

1. Right-click the computer object in Active Directory Users and Computers.
2. Click **Properties**.
3. Click the **General** tab.
4. Select **Trust computer for delegation**.

**To configure a service's account for delegation if the service runs under its own domain user account**

1. Right-click the user object in Active Directory Users and Computers.
2. Click **Properties**.
3. Click the **Account** tab.
4. In the **Account options** list, select **Account is trusted for delegation**.

**For accounts in domains that are Windows Server 2003 functional level**
To configure a service's account for delegation, determine whether the service runs under the host computer's Local System account or under its own domain user account.

**To configure a service's account for delegation if the service runs under the Local System account**

1. Right-click the computer object in Active Directory Users and Computers.
2. Click **Properties**.
3. Click the **Delegation** tab.

   **Note**

   - If the service's account is missing the **Delegation** tab and you have confirmed the domain is **Windows Server 2003** functional level, then the SPN for the service account has not been registered. Use Setspn to register the service account's SPN.
   - If the **Trust this user for delegation to any service (Kerberos only)** option is selected, then the account is using unconstrained delegation, which is not recommended.
   - If the **Trust this user for delegation to specified services only** option is selected, then the account is using constrained delegation.
   - If the **Use Kerberos Only** option is selected, then the account is using constrained delegation without protocol transition.
   - If the **Use any authentication protocol** option is selected, then the account is using constrained delegation with protocol transition.

**To configure a service's account for delegation if the service runs under its own domain user account**

1. Right-click the user object in Active Directory Users and Computers.
2. Click **Properties**.
3. Click the **Delegation** tab.

4. If **Do not trust this computer for delegation** is selected, select the appropriate option:

    - If the **Trust this computer for delegation to any service (Kerberos only)** option is selected, then the account is using unconstrained delegation, which is not recommended.
    - If the **Trust this computer for delegation to specified services only** option is selected, then the account is using constrained delegation.
    - If the **Use Kerberos Only** option is selected, then the account is using constrained delegation without protocol transition.
    - If the **Use any authentication protocol** option is selected, then the account is using constrained delegation with protocol transition.

**Differences Between Proxy and Forwarding**
Both proxy and forwarding are types of delegation. The credentials message (KRB_CRED) is used to deliver both proxy tickets and forwarded tickets. The KRB_CRED message is detailed later in this topic.

**Proxy is limited use.**
The client gives a front-end server a ticket and a key for a specific back-end server. The front-end server can authenticate as the client with this specific service only.

**Forwarding is general use.**
The client sends a TGT and a key for a TGS to the front-end server. The front-end server requests service tickets to wherever it needs to on the client's behalf. The service that receives the client's TGT is the client as far as the KDC is concerned.

Proxy requires that the client have information about the name of the back-end or target server or service. Rather than require that the client keep track of all possible back-end server names, most services use forwarded tickets, as described below.

When the user's identity is forwarded to another service, a special type of authentication called user-to-user authentication might be required. This special case is discussed later in this section.

**Proxy**
For a service to handle a request for a client, the service might need to complete that request in the context of the client. Proxy allows the client to send its identity in the form of a Kerberos ticket to the front-end server. The front-end server can then impersonate the client and authenticate with the back-end server as if the front-end server were the client.

**What are proxy tickets?**
Service tickets obtained by a client for a service to impersonate it are called proxy tickets. When the KDC issues a TGT to a client, it checks Kerberos policy to see if proxy tickets are allowed. If they are, the KDC sets the PROXIABLE flag in the TGT that it issues to the client.

The client obtains a proxy ticket by presenting a TGT to the ticket-granting service and asking for a ticket to the back-end server. The client's request includes a flag signaling that it wants a proxy ticket and also includes the name of the server that will represent the client. When the KDC receives the client's request, it creates a ticket for the back-end server, sets the PROXY flag in the ticket, and sends it back to the client. The client then sends the ticket to the front-end server, which uses the ticket to access the back-end server.

**Client sends a front-end service a service ticket and session key for a back-end service**
In a proxy situation, the client wants a front-end service to authenticate as the client with a specific back-end service. Delegation is limited to a specific target. For example, the client knows that it must authenticate with SQL. It also knows that SQL will need to use the client's credentials to authenticate with a specific database on another server.

To authenticate with these services, the client requests a service ticket for the database. In the request, the client specifies that this will be a proxy ticket and gives the KDC the name of the SQL server. The client sends the proxy ticket to the SQL server along with a session key for the database. The SQL server then authenticates with the database as if it were the client.

### Note

- Proxy requires that the client know the name of the back-end or target server or service. Rather than require that the client keep track of all possible back-end server names, most services use forwarded tickets, as is described in the following proxy process.

The illustrated message sequence in the following figure assumes that Client has already authenticated with Proxy. Now, Client wants Proxy to access Server on its behalf.

**Messages in the Proxy Process**




The message sequence for the proxy ticket is:

1. Client sends a request for a proxy ticket for Server to the TGS in the KDC. The request includes the client's TGT, a request for the ticket to be flagged as a proxy, and the addresses for which the ticket is valid. The addresses are those of the client computer and Proxy.

2. The TGS sends the proxy ticket and a session key for Server to Client.

3. Proxy needs the Server ticket and a valid Server session key. Client sends both of these to Proxy in a KRB_CRED (Kerberos credentials) message.

4. Proxy can now attempt to access Server as the workstation. The KRB_AP_REQ message is identical to the one the client would have sent if it were trying to access Server directly. The message contains the Server service ticket and an authenticator.

   Server decrypts the ticket with its secret key, finds the session key, and uses the session key to decrypt the authenticator. If this succeeds, Proxy is authenticated as Client.

5. If the authenticator passes the test, the service looks for a mutual authentication flag in the client's request. If the flag is set, the service uses the session key to encrypt the time from the user's authenticator and returns the result in a Kerberos application reply (KRB_AP_REP). If a mutual authentication flag is not set, no response is needed.

**Forwarding**

The difficulty with proxy tickets is that the client must have information about the name of the back-end server. This difficulty is overcome by a second method of delegation—forwarding—that enables the client to give the front-end server a TGT that it can use to request tickets as needed.

**What are forwarded tickets?**

Forwardable TGTs obtained from a client that enable the service to impersonate the client are called forwarded tickets.

If a client wants to delegate the task of obtaining tickets for back-end servers to a front-end server, it must ask the KDC for a forwardable TGT. It does this through an AS exchange, indicating to the KDC the name of the server that will act on the client's behalf. If Kerberos policy permits forwarding, the KDC creates a TGT for the front-end server to use in the client's name, sets the FORWARDABLE flag, and sends the TGT back to the client. The client then forwards the TGT to the front-end server.

When the front-end server requests a ticket to the back-end server, it presents the client's TGT to the KDC. When the KDC issues a ticket, it sees the FORWARDABLE flag in the TGT, sets the FORWARDED flag in the ticket, and returns the ticket to the front-end server.

**How the client's TGT and TGS session key are sent to front-end services**

Proxy is limited in that the client must have information about the target server's name. Forwarding relieves the client of this burden. Rather than ask the KDC for a service ticket specifically for the server, the client sends its TGT—with the FORWARDED flag set—and TGS session key to the delegate server.

When the delegate server has the client's forwarded TGT and TGS session key, it can request from the KDC directly any service tickets it needs for any additional servers or services. The client does not need to be involved. With the client's forwarded TGT, the delegate server can fully impersonate the client, and can request service tickets on the client's behalf and in the client's security context for any services.

**The Forwarded Ticket and Delegation**
In Windows 2000, Windows XP, and Windows Server 2003, Kerberos delegation allows a service running under the local system security context to take on the client's identity and authenticate as the client to another service.

A server that has a client's forwarded ticket *is* the client. Thus, the forwarded ticket enables the server to request new tickets from the TGS. Because a TGT is required to request new tickets, the forwarded ticket is the client's TGT. Then, when a client forwards its TGT to another server, the client is said to have delegated its identity.

There are two forwarding situations:

- The client delegates its identity to a service that logged on with its own credentials, has its own user account, and knows its own password.

- The client delegates its identity to a service that was started in a user context, is using that user's credentials, but did not log on and does not have information about the user's password (for example, DCOM's "activate as activator"). This situation requires a special authentication mode called user-to-user authentication, which is discussed in the next section.

### Messages in the Forwarding Process




The normal message sequence for a forwarded ticket is illustrated above.

1. The client requests a TGT with valid addresses for the client and the delegate. The TGT is requested as forwardable.

2. The KDC sends a TGT with valid addresses for the client and the delegate. The TGT is flagged as forwardable.

3. The client uses the TGT to request service ticket for the delegate server.

4. The KDC sends a service ticket for the delegate server.

5. The client uses the service ticket to connect to the delegate server and requests mutual authentication.

6. The client authenticates the delegate server.

   If the authenticator passes the test, the service looks for a mutual authentication flag in the client's request. If the flag is set, the service uses the session key to encrypt the time from the user's authenticator and returns the result in a Kerberos application reply (KRB_AP_REP). If no mutual authentication flag is set, then no response is needed.

   Up to this point, the messages are the same as a normal client-server authentication.

7. After the user is authenticated and authorized, the client sends the delegate server its TGT and a session key to be used in messages to the TGS. The KRB_CRED message is used to deliver the TGT and session key.

The delegate server now has the user's TGT and the correct session key for the user/TGS session; so Delegate can now request tickets to other resources in the user's name.

The delegate server requests and receives a service ticket for Server from the KDC. In the request, the delegate server sends the user's TGT and an authenticator encrypted with the user/TGS session key.

8. The delegate server sends the service ticket to Server. The ticket contains the user's credentials. The accompanying authenticator is encrypted with the TGS-supplied Delegate/Server session key.

   Server receives the authentication request and authenticates Delegate—that is, authenticates the user the delegate server is impersonating. Server will then build an access token based on the user's identity.

9. If the authenticator passes the test, the service looks for a mutual authentication flag in the delegate's request. If the flag is set, the service uses the session key to encrypt the time from the user's authenticator and returns the result in a Kerberos application reply (KRB_AP_REP). If no mutual authentication flag is set, then no response is needed.

**User-to-User Authentication**

The Kerberos authentication protocol supports authentication between clients, as specified in RFC 1510, when a client is running as a service that another client needs to access. In versions of Windows earlier than Windows 2000, NTLM is the only protocol that supports this type of user-to-user authentication.

It is often the case that a computer or service running in one of the new security contexts provided by Windows Server 2003 must run in a different security context to authenticate to a remote computer. For example, a local service running in the current user context would authenticate to a remote system as an Anonymous user.

User-to-user is the most complex Kerberos authentication. The following figure illustrates some typical user-to-user situations.

**Typical User-to-User Situation**




A typical situation requiring user-to-user authentication is illustrated above.

1. A user logs on and starts an application called Service A and forwards this application its identity. Service A can now access network resources on the user's behalf.

2. Then Service A, acting as the user, now starts another application called Service B. Service B shares Service A's identity of the user.

Three network entities now share the user's identity: the client, Service A, and Service B. Service B is unique, however, in that it is the only one of the three that did not have to log on.

When Service A started Service B, it did so with the user's existing identity. Service B will be associated with the user in Active Directory services, but Service B does not have information about the user's password or user key. Service B already has a copy of the user's TGT and the client/TGS session key, so it never needs to have information about the user's password.

**Why is this a problem?**
How will the user be able to authenticate with Service B? In normal authentication, the user would send Service B a ticket supplied by the TGS and an authenticator encrypted with the User/Service B session key. Service B should decrypt the ticket, get the session key, and decrypt the authenticator.

The problem is that Service B cannot decrypt the ticket because it does not have a copy of the user key. The solution to the problem is user-to-user authentication.

**The User-to-User Authentication Process**




As illustrated above, the user-to-user authentication process follows this sequence:

1. The user requests a ticket for Service B. The TGS checks with Active Directory to find the user account associated with Service B. Because Service B is impersonating, that account would be the user. Thus, the TGS checks the user object, gets the user key, and uses it to encrypt the service ticket.

2. The user sends the service ticket to Service B.

3. Even though Service B has the user's credentials, it does not have information about the user's password, cannot calculate the secret key, and thus cannot decrypt the ticket. The result is that the user cannot authenticate with Service B and so cannot use the application.

   Service B responds to the KRB_AP_REQ message with an error. With the error, Service B sends the user its TGT (which is the user's TGT in this case), along with a request for a new ticket encrypted with the session key from Service B's TGT rather than Service B's secret key.

4. The user requests a new ticket for Service B from the TGS. In the request message, the user sends its TGT, Service B's TGT in the Additional Tickets field, and the ENC-TKT-IN-SKEY flag, which is an optional request to encrypt the ticket with the additional ticket's session key.

   The TGS associates the additional ticket with Service B, because that is the target server in the ticket request. The KDC looks up Service B, finds that it is associated with the user object, gets the user key, and decrypts Service B's TGT. The TGS

gets the session key from Service B's TGT, creates a new ticket for the user to use, and encrypts it with the session key.

5. The user sends the new KRB_AP_REQ to Service B and authentication proceeds normally.

**What happens when another user tries to connect to Service B?**
The situation is not very different. Service B is using user 1's identity but does not have information about user 1's password, so it cannot decrypt a presented service ticket.

When user 2 asks the KDC for a service ticket to Service B, the KDC will send user 2 a service ticket that will be encrypted with user 1's secret key, because Service B is impersonating user 1.

When user 2 presents the ticket to Service B, Service B will respond with the same error message as when user 1 originally tried to connect. The error message will contain Service B's TGT (that is, user 1's TGT) and request a new service ticket encrypted with its session key.

Thus, it does not matter who tries to connect to Service B. The service ticket is encrypted with user 1's key regardless of the requester's identity. Because Service B never has information about user 1's password, it never has the user key and cannot decrypt the ticket.

### Note

- User-to-user tickets are not cached.

**Kerberos Protocol Transition**
Kerberos protocol transition allows a kerberized service (that is, a service designed to use Kerberos authentication features) to obtain a service ticket to itself on behalf of a Kerberos principal because no user credential is required for the transition to happen.

Kerberos protocol transition is enabled programmatically through the Win32 API LsaLogonUser(). A new logon message type (KerbS4ULogon) and a new logon information structure (KERB_S4U_LOGON) have been defined to be passed into the LsaLogonUser() call. A successful call into LsaLogonUser() will return a Windows access token. For more information about LsaLogonUser(), see "LSA Authentication" in the Microsoft Platform SDK on MSDN.

Two variants of an access token can be returned depending on whether the caller process possesses the user right, Act as part of the operating system:

- Without Act as part of the operating system, a security impersonation level of SecurityIdentification is returned.

- With Act as part of the operating system, a security impersonation level of SecurityImpersonation is returned.

An identification token allows the process to determine the identity and privileges associated with a user whereas an impersonation token also allows the process to take on the token user's identity and operate under that user's identity. With an impersonation level token, the caller process can use the service ticket obtained through protocol transition to obtain service tickets to other kerberized services if the service account has been configured correctly to use the constrained delegation extension.

The programmatic mechanism described above is the way that applications can be enabled to take advantage of protocol transition. In addition, the applications should also be configured to run as kerberized services.

**Kerberos Constrained Delegation**
In the Windows 2000 delegation model, the KDC does not limit the scope of services that a Kerberos principal's identity can be delegated to. That is, after a service account is trusted for delegation, it can request service tickets on behalf of a given user to any other service accounts. This unrestricted way of performing delegation does not provide for an application to specify a subset of service accounts that it deems trustworthy to delegate to. Thus, applications can be exposed to broader impersonation risks.

Impersonation might span across resource domains that have differing levels of security policy requirements, and some of those security policies might be weaker than the application's security requirements. From the domain administrator's point of view, it is risky to enable unconstrained Kerberos delegation in the enterprise, because there is no way to exclude untrusted servers from participating in delegation.

With constrained delegation, on the other hand, domain administrators can configure service accounts to only delegate to specific sets of service accounts. In Windows Server 2003, the **ms-DS-Allowed-To-Delegate-To** attribute is added to service

accounts to help enforce constrained delegation. The **ms-DS-Allowed-To-Delegate-To** attribute lists the service principal names (SPNs) of other service accounts that a given service account is allowed to delegate to. When a Windows Server 2003 KDC processes a service ticket request via the constrained delegation extension, it will verify that the target service account is one that is listed in the **ms-DS-Allowed-To-Delegate-To** attribute.

Thus, with the addition of the constrained delegation extension, Windows Server 2003 supports two modes of delegation:

**Unconstrained delegation**
Legacy mode of delegation is still supported only if the user has initially entered a user credential to obtain a TGT. This is not any different from the Windows 2000 implementation.

**Constrained delegation**
Can be enabled for a Kerberos principal when the principal has a TGT that is acquired through the normal process of authenticating to the KDC when:

- A user has initially entered a user credential to obtain a TGT.
- Another service account has obtained a service ticket on behalf of the principal through the protocol transition extension.

Constrained delegation is the only delegation mode supported with protocol transition and only works in the boundary of a domain.

You can configure a service account to use protocol transition in conjunction with constrained delegation by choosing **Trust this user for delegation to specified services only** and **Use any authentication protocol** on the **Delegation** tab of the account. If these options are selected for a service account, the account can use the protocol transition extension to obtain a service ticket enables it to obtain service tickets to a pre-configured subset of kerberized services. An additional requirement is that the impersonated user account must not be marked as **Account is sensitive and cannot be delegated** on the **Account** tab.

The following sections detail the contents of individual messages in each of the three Kerberos exchanges.

**Authentication Service Exchange Message Details**

### Authentication Service Exchange




**Message 1: The Authentication Service Request**
The initial message sent to the AS requesting a TGT is the KRB_AS_REQ message. This message and its corresponding response contain the most detail. Subsequent messages have less because they contain similar fields and data.

**KRB_AS_REQ message contents**
Although the same message structure is used for the initial KRB_AS_REQ message and for the KRB_TGS_REQ message (in which the client asks for a ticket for a target server), some of the fields are not be used for the KRB_AS_REQ message.

### Authentication Service Request Message Fields

| Field | Description |
|---|---|
| Protocol Version | 5 |
| Message Type | KRB_AS_REQ |

| | |
|---|---|
| **Pre-authentication Data** | The intent of these fields is to provide pre-authentication before the KDC sends the client a ticket. This is to prevent brute force or dictionary attacks on the user's password. Note that these fields are used in the KRB_TGS_REQ message for the TGT and the authenticator. |
| PAData Type | PA-AS-REQ or PA-PK-AS-REQ. |
| PAData Value | The AS_REQ includes a client timestamp encrypted with the user key. In this case, the data type will be PA-AS-REQ.<br><br>If a Smart Card is used, the user's public certificate signed with the user's private key will be placed in this field. The data type will be PA-PK-AS-REQ in this case. |
| **Request Body** | |
| KDC Options | The client can request that the TGT have certain optional features. The KDC options are detailed in the table below. |
| Client Name | Name of the requester. |
| Realm | Realm (Active Directory domain) of the requester. |
| Server Name | In the KRB_AS_REQ, this will be the KDC name. The client is specifically requesting a ticket for the TGS. |
| From | (Optional) If the ticket were postdated (not supported by Windows 2000, but is supported in Windows Server 2003), this field would specify the time from which the ticket would be valid. |
| Till | This is the requested expiration time. The KDC does not have to honor this request if the requested expiration time violates the domain's Kerberos policy. |
| Renew Time | (Optional) The requested renewal time. |
| Nonce | A random number generated by the client. The nonce can be used as the basis for sequence numbering or as an additional authenticator. The nonce supplied here will be returned in the encrypted portion of the KRB_AS_REP message. The client will compare the nonces to make sure they match. |
| Encryption Type | The desired encryption algorithm to be used. |
| Addresses | IP addresses from which ticket will be valid. |
| Encrypt Authorization Data | Not used for KRB_AS_REQ. |
| Additional Tickets | Not used for KRB_AS_REQ. |

The Encrypt Authorization Data field is not likely to be used in Microsoft implementations because the user key—or a private key if sending a certificate after logging on with a smart card—is used to encrypt a timestamp. The Additional Tickets field is only used in Kerberos authentication in some special cases, such as user-to-user authentication for distributed applications.

**KDC option flags for the KRB_AS_REQ and KRB_TGS_REQ messages**

These flags are very similar to the ticket flags discussed earlier (see the Ticket Flags table in the Kerberos Tickets topic), because you use KDC options to define the requested ticket flags. Many of these options are used in both the KRB_AS_REQ and KRB_TGS_REQ messages. Some are used only in the KRB_TGS_REQ message.

**Key Distribution Center Option Flags**

| Flag Bit | Flag Value | Flag Meaning |
|---|---|---|
| 1 | FORWARDABLE | Ticket can be forwarded. Allowed addresses are specified in the message's Addresses field. |
| 2 | FORWARDED | Ticket is a forwarded ticket. |

| | | |
|---|---|---|
| 3 | PROXIABLE | Ticket can be proxied. |
| 4 | PROXY | Ticket is a proxy ticket. |
| 5 | ALLOW POSTDATE | Ticket can be post-dated. When the service wanted the ticket to be activated, it would send a TGS request with the VALIDATE flag set (below). Not supported in Windows 2000, but supported in Windows Server 2003. |
| 6 | POST-DATED | Ticket is post-dated. Not supported in Windows 2000, but supported in Windows Server 2003. |
| 8 | RENEWABLE | Ticket can be renewed. |
| 15 | NAME_CANONICALIZE | Indicates to the KDC that the client is prepared to receive a reply with a different client or server principal name than was sent in the request. This flag is used in referrals. It also alerts the KDC that the target server might be in another realm.<br><br>The KDC option flag NAME_CANONICALIZE is important for cross-realm authentication. |
| 27 | RENEWABLE OK | This flag means that it is okay to issue a renewable ticket if a ticket with the requested expiration time cannot be provided. However, it does not mean that the ticket requested should be renewable. That flag is set by flag bit 8. |
| 28 | ENC-TKT-IN-SKEY | Encrypt ticket in session key. Normally, tickets are encrypted with the target server key. However, in user-to-user authentication, the ticket is encrypted with the session key taken from a provided TGT. This flag is used in that situation. |
| 30 | RENEW | This is not used for the KRB_AS_REQ message. If a ticket is flagged RENEWABLE and needs to be renewed, this flag would be set and the ticket needing renewal would be included with the KRB_TGS_REQ message. |
| 31 | VALIDATE | Only used with postdated tickets. Not supported in Windows 2000, but supported in Windows Server 2003. |

**Message 2: The Authentication Service Response**

The KRB_AS_REP message returns the TGT and a session key to the client.

The TGT is encrypted with the TGS's secret key, so only the TGS can decrypt it. This is important because the TGT contains the SIDs associated with the user and will be used to create the user's access token. This sensitive information needs to be protected from tampering.

The message also contains the short-lived session key that the client will use to encrypt data to the TGS or decrypt data from the TGS.

The session key is actually included in the message twice, as is true of much of the message's data. The client needs the session key, so it is included and encrypted in the user's secret key. The user will then be able to decrypt and use the session key. Another copy of the session key is included within the service ticket for the TGS (the TGT).

When the client has a valid TGT and a session key for the TGS, the client can request additional tickets for the local workstation or for network servers or services.

**KRB_AS_REP message contents**

This data structure is used for both the KRB_AS_REP message and the KRB_TGS_REP message. Thus, the field contents might change depending on the message type.

### Authentication Server Reply Message Fields

| Field | Description |
|---|---|
| Protocol Version | 5 |
| Message Type | KRB_AS_REP |

| | |
|---|---|
| Pre-authentication data | For the KRB_AS_REP message, this field will normally be empty.<br><br>One exception is if the client logged on with a smart card. In this case, because the user did not enter a name and password, the LSA does not have the user's password. So, the user's NTLM password and credentials are passed back to the LSA in this field (hashed and encrypted with the user's key). |
| Client Realm | Realm or domain of the requester. |
| Client Name | Name of the requester. |
| Ticket | The encrypted ticket is placed here. For the KRB_AS_REP message, this will be the TGT encrypted with the TGS's secret key.<br><br>Every domain controller will be a KDC and therefore will run the AS and TGS services. The client can connect to any TGS in the domain to request service tickets. Because the TGT is encrypted with the TGS's secret key, and because every TGS in the domain will need to be able to decrypt the TGT, every TGS's secret key must be identical. Thus, every domain controller will have a user account called krbtgt. This is the service account for the KDC, and its password (randomly selected and periodically changed) is used to create the TGS's secret key. |
| | The rest of the message fields are encrypted with the user's key.<br><br>The client cannot read the encrypted ticket and might need to verify the ticket information. Furthermore, the client needs to have information about start times and end times so it can determine when to request renewal or replacement tickets. Therefore, the seven fields from Authentication Time to Client Addresses are all included in the ticket and copied here into these fields. |
| Key | This is the session key the user will use to encrypt and decrypt TGS messages. |
| Last Requested | The last time a ticket was requested. This is similar to a "last logged on" time. This can be used to track how frequently a client is requesting tickets. A Kerberos policy can be set that limits how frequently ticket requests can be made. Such limits can help prevent brute force attacks. |
| Nonce | The nonce from the KRB_AS_REQ nonce field will be copied here. |
| Key Expiration | When the user's key will expire. This is used for password aging. |
| Flags | These are the flags set in the ticket, based on the flags requested in the KRB_AS_REQ message and the domain's Kerberos policy. |
| Authentication Time | The time the ticket was issued. |
| Start Time | (Optional) At what time the ticket is valid. |
| End Time | At what time the ticket expires (although it can be renewed if the ticket is renewable). |
| Renew Till | (Optional) At what time the ticket absolutely expires (cannot be renewed past this time). |
| Server Realm | Requested server's realm (domain). |
| Server Name | Requested server's name |

| | |
|---|---|
| Client Addresses | (Optional) Addresses from which the ticket will be valid. This is important in several situations. Normally, a ticket is only valid if sent from a specific address.<br><br>In Windows 2000, the KDC did not propagate the addresses to the ticket. Because of this, two new registry keys that affect Client Addresses were created for Windows Server 2003:<br><br>• **KdcUseClientAddresses**. Not created by default, this setting will configure the KDC to propagate client addresses to tickets.<br><br>• **KdcDontCheckAddresses.** Not created by default, this setting will configure the KDC to skip checking client addresses. This feature might be used in network address translation (NAT) scenarios.<br><br>If the client computer is multi-homed, the ticket needs to be valid from every IP-enabled adapter, so the addresses are added here (and in the address field in the ticket).<br><br>If a ticket is proxied or forwarded, it will be used on the client's behalf from another address (or perhaps several addresses). These addresses will also be included in this field. |

**Ticket-Granting Service Exchange Message Details**

### Ticket-Granting Service Exchange




**Message 3: The Ticket-Granting Service Request**

After the client has received a TGT, the client can begin requesting service tickets for local or network resources.

The primary contents of the KRB_TGS_REQ message are the TGT, the name of the target server for which a service ticket is requested, and the authenticator.

**KRB_TGS_REQ message contents**

As noted earlier in this section, the KRB_TGS_REQ message uses the same structure as the KRB_AS_REQ message, so many of the fields are used identically. The most important differences are the pre-authentication fields (PAData Type and PAData Value). The TGT and the authenticator are placed in the PAData Value field.

- The authenticator is encrypted with the session key.
- The TGT is encrypted with the TGS's secret key (the key based on the krbtgt user account).

The same message structure is used for the initial KRB_AS_REQ message and for the KRB_TGS_REQ message. However, some fields used in the KRB_TGS_REQ message were not used in the KRB_AS_REQ message. The following table lists and describes the fields used in the ticket-granting service request message.

### Ticket-Granting Service Request Message Fields

| Field | Description |
|---|---|
| Protocol Version | 5 |
| Message Type | KRB_TGS_REQ |

| | |
|---|---|
| **Pre-authentication Data** | The intent of these fields is to provide some authentication before the KDC sends the client a ticket. This is to prevent brute force or dictionary attacks on the user password. Note that this field is used in the KRB_TGS_REQ message for the TGT and the authenticator. |
| PAData Type | PA-TGS-REQ |
| PAData Value | TGT and authenticator. |
| **Request Body** | |
| KDC Options | The client can request that the ticket have certain optional features. The requested options take the form of a bit string. (See the Key Distribution Center Option Flags table above.) |
| Client Name | (Optional) Name of the requester. |
| Realm | Realm or domain of the requester. |
| Server Name | (Optional) Requested server name. |
| From | (Optional) If the ticket were postdated (not supported by Windows 2000, but supported in Windows Server 2003), this field would specify the time from which the ticket would be valid. |
| Till | This is the requested expiration time. The KDC does not have to honor this request, if the requested expiration time violates the domain's Kerberos policy. |
| Renew Time | (Optional) The requested renewal time. |
| Nonce | A random number generated by the client. The nonce can be used as the basis for sequence numbering or as an additional authenticator. The nonce supplied here will be returned in the encrypted portion of the KRB_TGS_REP message. The client will compare the nonces to make sure they match. |
| Encryption Type | The desired encryption algorithm to be used. |
| Addresses | (Optional) IP addresses from which ticket will be valid. |
| Encrypt Authorization Data | (Optional) In the KRB_TGS_REQ message, this will specify a key to use to encrypt any pre-authentication data. |
| Additional Tickets | (Optional) In a KRB_TGS_REQ for user-to-user authentication, a TGT will be included in this field. The TGS will use the session key from this TGT — instead of the service's secret key — to encrypt the service ticket. |

The TGS decrypts the ticket with the server key, then retrieves the session key from the ticket and uses it to decrypt the authenticator.

Only the TGT and the authenticator are encrypted in this message. A checksum is computed using the contents of fields in the request body, and that checksum is encrypted within the authenticator.

**Message 4: The Ticket-Granting Service Response**

The TGS response includes the ticket from the requested target server and a session key for the client and the target server with which to encrypt and decrypt their session messages.

The ticket is encrypted with the target server's secret key, so the client cannot decrypt and modify the ticket.

The message containing the session key for use with the target server is encrypted with the session key being used by the client and the TGS. The client can decrypt the new session key and use it to encrypt the authenticator that it will send to the target server along with the new ticket.

**KRB_TGS_REP message contents**

As noted earlier in this section, the KRB_TGS_REP and KRB_AS_REP messages are identically structured and the fields are similarly used. The major difference is that the ticket in the Ticket field is not a TGT. Instead, it is the service ticket for the target server.

### Ticket-Granting Server Reply Message Fields

| Field | Description |
|---|---|
| Protocol Version | 5 |
| Message Type | KRB_TGS_REP |
| Pre-authentication Data | In a KRB_TGS_REP message, application-specific data might be placed here. |
| Client Realm | Realm or domain of the requester. |
| Client Name | Name of the requester |
| Ticket | The encrypted ticket is placed here. For the KRB_TGS_REP message, this will be the service ticket encrypted with the target server's secret key. |
| | The message fields below are encrypted with the TGS session key. |
| Key | This is the session key the user will use with the application server. |
| Last Requested | The last time a ticket was requested. This is similar to a "last logged on" time. This can be used to track how frequently a client is requesting tickets. A Kerberos policy can be set that limits how frequently ticket requests can be made. Such limits can help prevent brute force attacks. |
| Nonce | The nonce from the KRB_TGS_REQ nonce field will be copied here. |
| Key Expiration | (Optional) When the user's key will expire. This is used for password aging. |
| Flags | These are the flags set in the ticket, based on the flags requested in the KRB_TGS_REQ message and on the domain's Kerberos policy. |
| Authentication Time | The time the ticket was issued. |
| Start Time | (Optional) At what time the ticket is valid. |
| End Time | At what time the ticket expires (although it can be renewed if the ticket is renewable). |
| Renew Till | (Optional) At what time the ticket absolutely expires (cannot be renewed past this time). |
| Server Realm | Requested server's realm (domain). |
| Server Name | Requested server's name. |
| Client Addresses | (Optional) Addresses from which the ticket will be valid. If a ticket is proxied or forwarded, it will be used on the client's behalf from another address (or perhaps several addresses). These addresses will also be included in this field. |

**Client/Server Authentication Exchange Message Details**

**Client/Server Authentication Exchange**



**Message 5: The Application Server Request**
After the client has requested and received a TGT from the AS, and has requested and received a service ticket from the TGS, the client is ready to send the ticket to the target server.

The KRB_AP_REQ message contains an authenticator encrypted with the session key that the client and target server share, the service ticket encrypted with the target server's secret key, and the optional mutual authentication request. The mutual authentication response, which takes the form of a KRB_AP_REP message, is only necessary if the client requires verification that the target server is truly the recipient of the KRB_AP_REQ and subsequent messages.

### Note

- If the target service requires user-to-user authentication, the initial service ticket will be rejected along with a copy of the target service's TGT. The client will then make a new KRB_TGS_REQ message containing the target service's TGT and requesting a new service ticket that is encrypted with the session key contained in the target service's TGT, instead of with the target service's secret key.

**KRB_AP_REQ Message Contents**
### Application Server Request Message Fields

| Field | Description |
|---|---|
| Protocol Version | 5 |
| Message Type | KRB_AP_REQ |
| **Application Options Fields** | |
| Use Session Key | During user-to-user authentication, after the client has obtained a new ticket for the target service, the client will send a new KRB_AP_REQ with this flag set. This tells the target service to use its session key to decrypt the ticket. (See the User-To-User Authentication section above.) |
| Mutual Authentication Required | If this flag is set, the target server will respond with the KRB_AP_REP message, which will authenticate the target server to the client. |
| **Other Message Fields** | |
| Ticket | The service ticket for the target server, encrypted in either the target server's secret key or in its session key, depending on whether user-to-user authentication is required. |
| Authenticator | The client's timestamp and other data, encrypted with the session key that the client and target server share. |

**Message 6: The Optional Application Server Response**
This message is optional unless mutual authentication is required. The message is used if the client needs to verify the target server's identity. This is requested in an application options field (Mutual Authentication Required) of the KRB_AP_REQ message. If this is requested, the target server will take the client computer's timestamp from the authenticator, encrypt it with the session key that the TGS provided for client/target server messages, and send it to the client.

### Note

- Most Windows services request mutual authentication.

How do the contents of the KRB_AP_REQ message prove to the client that the target server was the actual recipient of the client's authentication request? The authenticator is encrypted with the client/server session key. The session key is inside the service ticket. The service ticket is encrypted with the secret key shared by the server and the TGS. Thus, the only way the server can get access to the client's timestamp in the authenticator is to have information about the shared secret key, and the server would only have information about the secret key if it truly was the server and not an imposter.

To summarize, the steps in the mutual authentication process are:

1. The server receives the service ticket and the authenticator in KRB_AS_REQ.
2. The server uses its secret key to decrypt the ticket.
3. The server retrieves the session key from the ticket.
4. The server decrypts the authenticator with the session key.
5. The server retrieves the client's timestamp from the authenticator.
6. The server creates the mutual authentication message, includes the client's timestamp, encrypts the message with the session key, and sends KRB_AS_REP to the client.
7. The client decrypts this message with the session key, and thus the server is authenticated.

**KRB_AP_REP message contents**

The mutual authentication message is sent to authenticate the target server to the client.

The main information is simply a copy of the client's timestamp from the authenticator—that was previously included in the KRB_AP_REQ message—encrypted with the session key that the client and the target server share.

The application might require a new key for encrypting subsequent messages, replacing the session key, in which case the new key would need to be given to the server.

### Application Server Reply Message Fields

| Field | Description |
|---|---|
| Protocol Version | 5 |
| Message Type | KRB_AP_REP |
| | The following message fields are encrypted with the session key |
| Client Time | Current time on Client (from the authenticator). |
| CUSEC | Millisecond part of client time (from the authenticator). |
| Subkey | (Optional) Specifies a key to encrypt client sessions with application server. |
| Sequence Number | (Optional) Application-specific, so used if sequence number was specified in the authenticator. |

**Credentials Message Details**

### Credentials Message




The Kerberos credentials message is used to transport a user's ticket and session key securely to the service impersonating it. The ticket—because it is already encrypted—can simply be sent, but the session key will need to be encrypted.

**KRB_CRED encryption**

- The user's service ticket for Server is encrypted with Server's key (the key that Server shares with the KDC).
- The user's session key for Server is encrypted with the session key that the client and Delegate share. Server needs to use the session key, so it needs to be able to decrypt it.

**KRB_CRED message contents**

### Credentials Message Fields

| Field | Description |
|---|---|
| Protocol Version | 5 |
| Message Type | KRB_CRED |
| Ticket | The encrypted ticket is placed here. For the KRB_TGS_REP message, this will be the service ticket encrypted with the target server's secret key. |
| | The following message fields are encrypted with the session key that the client and Delegate share. |
| Key | This is the session key the client will use with the application server. |
| Principal Realm | (Optional) Name of the delegated principal's realm (domain). |
| Principal Name | (Optional) Delegated principal's name. |
| Flags | (Optional) These are the flags set in the ticket. |
| Authentication Time | (Optional) The time the ticket was issued. |
| Start Time | (Optional) At what time the ticket is valid. |
| End Time | (Optional) At what time the ticket expires (although it can be renewed if the ticket is renewable). |
| Renew Till | (Optional) At what time the ticket absolutely expires (cannot be renewed past this time). |
| Server Realm | (Optional) Requested server's realm (domain). |
| Server Name | (Optional) Requested server's name. |
| Client Addresses | (Optional) Addresses from which the ticket will be valid. If a ticket is proxied or forwarded, it will be used on the client's behalf from another address (or perhaps several addresses). These addresses will also be included in this field. |
| Nonce | The nonce from the KRB_TGS_REQ Nonce field will be copied here. |
| Timestamp | Time when KRB_CRED was generated. |
| USEC | Millisecond part of timestamp. |
| Sender Address | (Optional) Addresses, application-specific. |
| Recipient Address | (Optional) Addresses, application-specific. |

To summarize, when proxying or forwarding, the credentials message contains three necessary items to impersonate the user. When proxying, the KRB_CRD message contains the:

- Proxy service ticket.

- User's session key with target server.
- Information that a normal user would have as the result of AS exchange.

When forwarding, the KRB_CRED message contains the:

- User's forwarded TGT.
- User's TGS session key.
- Information that a normal user would have as the result of TGS exchange.

With a valid ticket, session key, and other client information, the service is ready to impersonate the client, whether proxying or forwarding.

**Kerberos and IP Transport: UDP/TCP**
According to RFC 1510, when a client contacts the KDC, it should send a User Datagram Protocol (UDP) datagram to port 88 at the KDC's IP address. The KDC should respond with a reply datagram to the sending port at the sender's IP address.

UDP is a connectionless transport protocol, making it a logical choice where an exchange of messages must precede a connection. UDP is also well suited to applications with one message and one response, such as the exchanges between clients and the KDC, so long as each message fits into a single datagram. However, UDP works best when datagrams are transmitted as single units—when each datagram fits into one frame. The capacity of a frame varies with the medium. The Maximum Transmission Unit (MTU) for an Ethernet frame is 1500 octets. Thus, if the physical network is Ethernet, it follows that Kerberos messages sent as UDP datagrams can carry up to 1500 octets of data.

On the other hand, Windows authorization data can easily total more than 1500 octets. Because Windows authorization data is needed only by computers with the Windows operating system, Windows-specific data is omitted from tickets to computers with other operating systems. Messages containing tickets to computers with other operating systems are well within the limits of UDP transport, so that is how they are transmitted. Messages containing tickets for computers running Windows are likely to exceed the limit, so those messages are transmitted using the Transmission Control Protocol (TCP), which has much greater capacity.

This physical limit constitutes two different transport possibilities when computers with the Windows operating system send Kerberos messages on an Ethernet network.

- Messages below the 1465 maximum UDP payload are sent as normal UDP datagrams without incident.
- Messages that exceed the 1465 byte limit establish a TCP connection using port 88.

**Finding the KDC: DNS Name Resolution**
RFC 1510 specifies that IP transport should be used for messages between clients and the KDC. When the Kerberos SSP on a client computer wants to send an initial authentication service (AS) request, it needs to find an IP address for the KDC in the user account's domain.

To do that, the Windows client needs the Domain Name System (DNS) name of a server where the KDC service is running. If the DNS name can be resolved to an IP address, that is where the Kerberos SSP sends its message. Otherwise, the Kerberos SSP generates an error indicating that it can find no such domain.

**Why Do Windows Systems Use DNS?**

- The client computer needs the IP address of the KDC in the user account's domain.
- Windows systems use DNS to find IP addresses.
- Every Active Directory domain controller runs the KDC service. Domain controllers are advertised in DNS service locator records (SRV RR) as defined in RFC 2052.

Therefore, all that is needed to find a domain's KDC is to locate a domain controller's IP address by querying DNS for the SRV RR for domain controllers.

**How the Kerberos SSP locates a KDC**

1. If the Netlogon service is running, the Kerberos SSP asks Netlogon for the name of a domain controller (using the DsrGetDcName internal API).

2. If the domain controller's IP address is not found in the local cache, DNS is queried. The response will contain the DNS names of one or more domain controllers, along with their IP addresses.

   Clients that do not support the SRV record type in their DNS resolver can do a standard host record (A RR) lookup on *ActiveDirectorydomainname*. The response will contain the IP addresses of all domain controllers for the domain.

3. If the computer is a domain controller, the Kerberos client assumes that the KDC is running locally. If the KDC does not start or is not running, the Kerberos SSP finds another KDC in the domain.

4. After the client locates a domain controller, the KRB_AS_REQ message is sent to the KDC service on that domain controller.

**Retrying**
The Kerberos client retries up to three times to contact a KDC. It initially waits 10 seconds for a reply, and will wait an additional 10 seconds on each retry. After a failure, the Kerberos client forces discovery of another KDC.

### Note

- In Active Directory domains, the KDC service runs on every domain controller. In addition to being KDC servers, domain controllers are also Lightweight Directory Access Protocol (LDAP) servers. Both services are registered in DNS service locator records (SRV resource records). Clients can find a domain controller by querying DNS for SRV resource records with the name _ldap._tcp.dc._msdcs.*DnsDomainName*. The KDC service can be located by querying DNS for SRV resource records with the name _kerberos._udp.*DnsDomainName*. Clients that do not support the SRV record type in their DNS resolver can query for a host record (an A resource record) with the name of the domain.

**Locating MIT Kerberos KDCs**
Computers running Windows 2000, Windows XP, and Windows Server 2003 can participate in Kerberos realms that are based on the MIT Kerberos version 5 reference implementation.

If the specified logon realm is determined to be an MIT Kerberos realm, then the Kerberos SSP retrieves the list of MIT KDCs from the registry.

The servers in the list are cycled through on each new request for a KDC without any attempt to verify if the KDC is running on the server.

Computers running Windows Server 2003 can participate in Kerberos realms that are not Active Directory domains. In this case, the KDC will not be on an Active Directory domain controller, so the DNS names for KDC servers must be stored in the client computer's registry. The Kerberos SSP looks in the registry for the DNS domain name of the user's realm, and then resolves the name to an IP address by querying a DNS server.

Windows Server 2003 includes Kerberos Setup (ksetup.exe), which you can use to configure clients to participate in Kerberos realms that are not Active Directory domains. To find more information about Kerberos Setup, see "Security Management Tools" in the Windows Support Tools section of the Tools and Settings Collection.

**Network Ports Used by the Kerberos V5 Protocol**
### Network Ports Used During Kerberos Authentication

| Service Name | UDP | TCP |
|---|---|---|
| DNS | 53 | 53 |
| Kerberos | 88 | 88 |

**Related Information**
The following resources contain additional information that is relevant to this section.

- Active Directory Collection

- DNS Technical Reference
- Account Policy Settings
- Domain and Forest Trusts Technical Reference
- Windows Time Service Technical Reference

# Kerberos Authentication Tools and Settings

Updated: June 26, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Kerberos Authentication Tools and Settings**

### In this section


- Kerberos Protocol Tools
- Kerberos Protocol Registry Settings
- Kerberos Authentication Group Policy Settings
- Network Ports Used by the Kerberos V5 Protocol


**Kerberos Protocol Tools**

The following tools are associated with administering and troubleshooting the Kerberos version 5 authentication protocol.

**Domain.msc: Active Directory Domains and Trusts**

**Category**

Active Directory Domains and Trusts is an MMC snap-in that is installed automatically on computers running Windows Server 2003 when you install Active Directory. You can also use this tool on non-domain controllers by installing the Windows Server 2003 Administration Tools pack. You can start Active Directory Domains and Trusts in the following way: Click **Start**, then click **Programs**,then click **Administrative Tools,** and then click **Active Directory Domains and Trusts.**

**Version compatibility**

This tool is compatible with domain controllers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. You can also use Active Directory Domains and Trusts with the Windows Server 2003 Administration Tools pack to remotely administer Active Directory from a computer that is not a domain controller, such as a computer running Windows XP Professional. You can use this tool to target any domain controller in either Windows 2000 Server or Windows Server 2003 domains.

### Note

- When using Windows Server 2003 Active Directory administrative tools to connect to a domain controller running Windows 2000, you must first be sure that the Windows 2000–based domain controller to which you are connecting has Service Pack 3 or later installed. This is because Windows Server 2003 administrative tools sign and encrypt all Lightweight Directory Access Protocol (LDAP) traffic by default.

Active Directory Domains and Trusts provides a graphical interface in which you can view all domains in the forest. Using this tool, an administrator can manage each of the domains in the forest, trust relationships between domains, configure the functional level for each domain or forest, and configure the alternative user principal name (UPN) suffixes for a forest.

Active Directory Domains and Trusts can be used to create, view, and modify most trust-related tasks. It can be used to target all Active Directory domain controllers and can verify all Active Directory trust types. Trust verification takes place between two domains by enumerating all of the domain controllers in each domain. If you choose to have Active Directory Domains and Trusts create both sides of the trust at once, the trust password is automatically generated.

You can find more information about Active Directory Domains and Trusts on Microsoft TechNet.

**Dsa.msc: Active Directory Users and Computers**

**Category**

Active Directory Users and Computers is a Microsoft Management Console (MMC) snap-in that is automatically installed when you install Active Directory. Dsa.msc is also included with the Administration Tools Pack (Adminpak.msi).

You can access the tool from the **Start** menu: Click **Start**, click **Programs**,click **Administrative Tools,** and then click **Active Directory Users and Computers**.

**Version compatibility**

Active Directory Users and Computers runs on domain controllers that are running Windows Server 2003 or Windows 2000.

In both of these server systems, MMC provides a window in the user interface where you can add, configure, and control items. Active Directory Users and Computers is the MMC tool that you can use to administer and publish information in the directory.

The Windows Server 2003 version of Active Directory Users and Computers can target domain controllers that are running Windows Server 2003 or Windows 2000.

On administrative workstations that are running Windows XP Professional or Windows 2000, you can install the Windows Server 2003 Administration Tools Pack (Adminpak.msi) from the i386 folder on the Windows Server 2003 CD. This version of the Administration Tools Pack encrypts and signs LDAP traffic between the administrative tool clients and domain controllers.

**Note**

- You cannot run the Windows Server 2003 Administration Tools Pack (Adminpak.msi) on a computer that is running Windows XP Professional, Windows XP Home Edition, or Windows XP 64-Bit Edition Version 2003 without Windows XP Service Pack 1 (SP1).

You can use Active Directory Users and Computers to manage the following properties associated with objects. Your managing the properties will affect Kerberos V5 authentication for these objects.

**Active Directory Users and Computers Object Management**

| **Property** | **Changes That Affect Kerberos V5 Authentication** |
|---|---|
| **Computer objects** | |
| Trust computer for delegation | The service can impersonate a user to use other network services. Allows services running as local system to request services from other services. |
| **Computer objects delegation tab options** | **(This tab will only appear in domains with Windows Server 2003 Functional Level.)** |
| Do not trust this computer for delegation | |
| Trust this computer for delegation to any service (Kerberos only) | Not recommended. This enables unconstrained delegation. |
| Trust this computer for delegation to specified services only | • Use Kerberos only: Enables constrained delegation without protocol transition.<br>• Use any authentication protocol: Enables constrained delegation with protocol transition |
| **User or service objects account tab options** | |
| Account is trusted for delegation | The account is enabled for delegation. This is a security-sensitive setting. Accounts with this option enabled should be tightly controlled. This setting enables a service that runs under the account to assume a client's identity and authenticate as that user to remote servers on the network. |
| Account is sensitive and cannot be delegated | When this flag is set, the security context of the user is not delegated to a service even if the service account is set as trusted for Kerberos delegation. |
| Use DES encryption types for this account | Restrict this principal to use only Data Encryption Standard (DES) encryption types for keys. If a user account is configured to use DES encryption, a Windows 2000, Windows XP, or Windows Server 2003–based client requests a ticket-granting ticket (TGT) by using the DES-CBC-MD5 encryption type. |

| | |
|---|---|
| Do not require Kerberos pre-authentication | Overrides the default setting that the KDC requires all accounts to use pre-authentication. The default setting makes offline password-guessing attacks very difficult. You can choose to override the default setting for individual accounts when necessary for compatibility with other implementations of the protocol. |
| **User or service objects delegation tab options** | **(This tab will only appear on accounts with registered SPNs in domains with Windows Server 2003 Functional Level.)** |
| Do not trust this user for delegation | |
| Trust this user for delegation to any service (Kerberos only) | Not recommended. This enables unconstrained delegation. |
| Trust this user for delegation to specified services only | • Use Kerberos only: Enables constrained delegation without protocol transition.<br>• Use any authentication protocol: Enables constrained delegation with protocol transition. |

You can find more information about Active Directory Users and Computers on Microsoft TechNet.

**Eventvwr.msc: Event Viewer**
**Category**
Event Viewer is included in Windows Server 2003, Windows XP, and Windows 2000.

**Version compatibility**
Event Viewer is supported for Windows Server 2003, Windows XP, and Windows 2000.

The system and security logs contain Kerberos error codes and other events related to authentication.

The following table lists event IDs and information that can be associated with Kerberos authentication. An event log contains only relevant event information. For example, only failure audits have Kerberos error codes; smartcard logons have certificate information.

**Security Log Events That Might Contain Kerberos Error Codes**

| Event ID | Account Logon Event Type | Event Information Potentially Associated with Kerberos Authentication |
|---|---|---|

| | | |
|---|---|---|
| 672 | • Success audit (Windows 2000 and Windows Server 2003)<br>• Failure audit (Windows Server 2003) | Authentication Ticket Request:<br>• User Name<br>• Supplied Realm Name<br>• User ID<br>• Service Name<br>• Service ID<br>• Ticket Options<br>• Result Code: *Kerberos error code*<br>• Ticket Encryption Type<br>• Pre-Authentication Type<br>• Client Address<br>• Certificate Issuer Name<br>• Certificate Serial Number<br>• Certificate Thumbprint |
| 673 | • Success audit (Windows 2000 and Windows Server 2003)<br>• Failure audit (Windows Server 2003) | Service Ticket Request:<br>• User Name<br>• User Domain<br>• Service Name<br>• Service ID<br>• Ticket Options<br>• Ticket Encryption Type<br>• Client Address<br>• Failure Code: *Kerberos Error Code*<br>• Logon GUID<br>• Transited Services |

| | | |
|---|---|---|
| 675 | • Failure audit | Pre-authentication Failed:<br>• User Name<br>• User ID<br>• Service Name<br>• Pre-authentication Type<br>• Failure Code: *Kerberos error code*<br>• Client Address |
| 676 | • Failure audit (Obsolete in Windows Server 2003; both success and failure audits use event ID 672.) | Authentication Ticket Request Failed:<br>• User Name<br>• Supplied Realm Name<br>• Service Name<br>• Ticket Options<br>• Failure Code: *Kerberos error code*<br>• Client Address |
| 677 | • Failure audit (Obsolete in Windows Server 2003; both success and failure audits use event ID 673.) | Service Ticket Request Failed:<br>• User Name<br>• User Domain<br>• Service Name<br>• Ticket Options<br>• Failure Code: *Kerberos error code*<br>• Client Address |

**Kerberos V5 Authentication Protocol Error Messages Generated by Windows Server 2003**

| Kerberos Error Number | Kerberos Error Code | Description |
|---|---|---|
| 0x3 | KDC_ERR_BAD_PVNO | Requested protocol version number not supported. |
| 0x6 | KDC_ERR_C_PRINCIPAL_UNKNOWN | Client not found in Kerberos database. |
| 0x7 | KDC_ERR_S_PRINCIPAL_UNKNOWN | Server not found in Kerberos database. |
| 0x8 | KDC_ERR_PRINCIPAL_NOT_UNIQUE | Multiple principal entries in database. |
| 0xA | KDC_ERR_CANNOT_POSTDATE | Ticket not eligible for postdating. |

| | | |
|---|---|---|
| 0xB | KDC_ERR_NEVER_VALID | Requested start time is later than end time. |
| 0xC | KDC_ERR_POLICY | KDC policy rejects request. |
| 0xD | KDC_ERR_BADOPTION | KDC cannot accommodate requested option. |
| 0xE | KDC_ERR_ETYPE_NOSUPP | KDC has no support for encryption type. |
| 0xF | KDC_ERR_SUMTYPE_NOSUPP | KDC has no support for checksum type. |
| 0x10 | KDC_ERR_PADATA_TYPE_NOSUPP | KDC has no support for pre-authentication data type. |
| 0x12 | KDC_ERR_CLIENT_REVOKED | Client's credentials have been revoked. |
| 0x17 | KDC_ERR_KEY_EXPIRED | Password has expired - change password to reset. |
| 0x18 | KDC_ERR_PREAUTH_FAILED | Pre-authentication information was invalid. |
| 0x19 | KDC_ERR_PREAUTH_REQUIRED | Additional pre-authentication required. |
| 0x1B | KDC_ERR_MUST_USE_USER2USER | Server principal valid for user-to-user only. |
| 0x1C | KDC_ERR_PATH_NOT_ACCPETED | KDC Policy rejects transited path. |
| 0x1D | KDC_ERR_SVC_UNAVAILABLE | A service is not available. |
| 0x1F | KRB_AP_ERR_BAD_INTEGRITY | Integrity check on decrypted field failed. |
| 0x20 | KRB_AP_ERR_TKT_EXPIRED | Ticket expired. |
| 0x21 | KRB_AP_ERR_TKT_NYV | Ticket not yet valid. |
| 0x22 | KRB_AP_ERR_REPEAT | Request is a replay. |
| 0x23 | KRB_AP_ERR_NOT_US | The ticket isn't for us. |
| 0x24 | KRB_AP_ERR_BADMATCH | Ticket and authenticator do not match. |
| 0x25 | KRB_AP_ERR_SKEW | Clock skew too great. |
| 0x28 | KRB_AP_ERR_MSG_TYPE | Invalid message type. |
| 0x29 | KRB_AP_ERR_MODIFIED | Message stream modified. |
| 0x34 | KRB_ERR_RESPONSE_TOO_BIG | Response too big for UDP, retry with TCP. |
| 0x3C | KRB_ERR_GENERIC | Generic error (description in e-text). |
| 0x44 | KDC_ERR_WRONG_REALM | User-to-user TGT issued different KDC. |

To find more information about "Event Viewer", see "Event Viewer" on Microsoft TechNet.

**Kerbtray.exe: Kerberos Tray**
**Category**
Kerberos Tray is included in the Windows Server 2003 Resource Kit and the Windows 2000 Resource Kit.

**Version compatibility**
Kerberos Tray is supported for Windows Server 2003, Windows XP, and Windows 2000.

Kerberos Tray is a graphical user interface tool that displays ticket information for a computer running Microsoft's implementation of the Kerberos version 5 authentication protocol.

You can view and purge the ticket cache by using the Kerberos Tray tool icon located in the notification area of the desktop. By positioning the cursor over the icon, you can view the time left until the initial ticket-granting ticket (TGT) expires. The icon also changes in the hour before the Local Security Authority (LSA) renews the ticket.

To find more information about Kerberos Tray, see "Windows Server 2003 Resource Kit Tools Help in the Tools and Settings Collection."

**Klist.exe: Kerberos List**
**Category**
Kerberos List is included in the Windows Server 2003 Resource Kit and the Windows 2000 Resource Kit.

**Version compatibility**
Kerberos List is supported for Windows Server 2003, Windows XP, and Windows 2000.

Kerberos List is a command-line tool that is used to view and delete Kerberos tickets granted to the current logon session. To use Kerberos List to view tickets, you must run the tool on a computer that is a member of a Kerberos realm.

When Kerberos List is run from a client, it shows the:

- Ticket-granting ticket (TGT) to a Kerberos Key Distribution Center (KDC) in Windows.
- Ticket-granting ticket (TGT) to Ksserver on UNIX.

**Parameters**
Kerberos List uses the following syntax:

```
klist [tickets | tgt | purge] [-?]
```

**Tickets**
Lists the current cached tickets of services to which you have authenticated since logging on. Displays the following attributes of all cached tickets:

| Option | Description |
|---|---|
| End Time | Time at which that the ticket becomes invalid. After a ticket is past this time, it cannot be used to authenticate to a service. |
| KerbTicket Encryption Type | Encryption type used to encrypt the Kerberos ticket. |
| Renew Time | Maximum lifetime of a renewable ticket (see TicketFlags in the table below). To continue using this ticket, you must renew it before reaching the established End Time and before the expiration date established in RenewUntil. |
| Server | Server and domain for the ticket. |

**tgt**
Lists the initial Kerberos ticket-granting ticket (TGT). Displays the following attributes of the currently cached ticket:

| Option | Description |
|---|---|
| AltTargetDomainName | Name supplied to InitializeSecurityContext that generated this ticket, typically a service principal name (SPN). |
| DomainName | Domain name of the service. |
| End Time | Time when the ticket becomes invalid. When a ticket is past the end time, it cannot be used to authenticate to a service. |
| FullServiceName | Canonical name of the account principal for the service. |
| KeyExpirationTime | Expiration time from the KDC reply. |
| RenewUntil Maximum lifetime of a renewable ticket (see TicketFlags) | To continue using a ticket, you must renew it. Tickets must be renewed before the expiration time set in End Time and in RenewUntil. |

| | |
|---|---|
| ServiceName | A TGT is a ticket for the Key Distribution Center (KDC) service. The service name for a TGT is "krbtgt." |
| Start time | Time when the ticket becomes valid. |
| TargetDomainName | For a cross-realm ticket, this is the realm, rather than the issuing realm, in which the ticket is good. |
| TargetName | Service name for which the ticket was requested. This is the name of a servicePrincipalName property on an account in the directory. |
| TicketFlags | Kerberos ticket flags set on the current ticket in hexadecimal. Kerberos Tray displays these flags on the Flags tab. |
| TimeSkew | The reported time difference between the client computer and the server computer for a ticket. |

**purge**
Will list each ticket and enables you to delete specific tickets. When you choose **Yes** to a ticket, then purge tickets destroys the ticket that you have cached, so use this with caution. It might stop you from being able to authenticate to resources. If this happens, you must log off and then log on again.

**-?**
Displays command-line help.

To find more information about Kerberos List, see "Windows Server 2003 Resource Kit Tools Help" in the Tools and Settings Collection.

**Ksetup.exe: Kerberos Setup**
**Category**
Kerberos Setup is included in the Windows Server 2003 Support Tools.

**Version compatibility**
Kerberos Setup is supported for Windows Server 2003.

Kerberos Setup is a command-line tool you can use to configure Windows for Kerberos V5 interoperability. Kerberos Setup configures a client connected to a server running Windows Server 2003 to use a server running the Kerberos V5 authentication protocol. After Kerberos Setup configures the client, the client then uses a Kerberos V5 realm instead of a Windows Server 2003 domain. This provides a single sign-on to the Key Distribution Center (KDC) and a local client account connected to a computer running Windows Server 2003.

Administrators can use Kerberos Setup to:

- Set up a realm entry for a Kerberos V5 realm.
- Set up a list of KDCs for the Kerberos V5 realm for which you set up a realm entry.
- Set up a kpasswd server for the Kerberos V5 realm for which you set up a realm entry.
- Set up local account to Kerberos V5 account mappings.
- Set the computer's password in the Kerberos realm.
- Change a user's password in a Kerberos V5 realm.

To find more information about Kerberos Setup, see "Support Tools Help" in the Tools and Settings Collection.

**Ktpass.exe: Kerberos Keytab Setup**
**Category**
Kerberos Keytab Setup is included in the Windows Server 2003 Support Tools.

**Version compatibility**
Kerberos Keytab Setup is supported for Windows Server 2003.

Kerberos Keytab Setup is a command-line tool that enables an administrator to configure a non-Windows Kerberos service as a security principal in the Windows Server 2003 Active Directory. Kerberos Keytab Setup configures the server principal name for the host or service in Active Directory and generates an MIT-style Kerberos keytab file containing the shared secret key of the service. The tool enables non-Windows brand operating system services that support Kerberos authentication to use the interoperability features provided by the Kerberos KDC service that runs on Windows Server 2003.

To find more information about Kerberos Keytab Setup, see "Support Tools Help" in the Tools and Settings Collection.

**Netdom.exe: Windows Domain Manager**
**Category**
Windows Domain Manager is included in Windows Server 2003, Windows 2000, and the Windows Server 2003 Administration Tools Pack (Adminpak.msi).

**Version compatibility**
Windows Domain Manager is supported for Windows Server 2003, Windows XP, and Windows 2000.

On administrative workstations that are running Windows XP Professional, you can install the Windows Server 2003 Administration Tools Pack (Adminpak.msi) from the i386 directory on the Windows Server 2003 CD. Please see the note in the Active Directory Users and Computers subsection above regarding the Administration Tools Pack.

The Windows Server 2003 version of Windows Domain Manager can target domain controllers that are running Windows Server 2003 or Windows 2000.

This tool enables you to manage Active Directory domains and trust relationships from the command line.

Functions of Windows Domain Manager that relate to Kerberos authentication include:

- Establishing one-way or two-way trust relationships between domains, such as:
    - From an Active Directory domain to an Active Directory domain in another enterprise (also know as an uplevel external trust).
    - Between two Active Directory domains in an enterprise (a shortcut trust).
    - The Windows Server 2003 or Windows 2000 Server half of an interoperable Kerberos realm.
- Managing trust relationships between domains, including:
    - Enumerating direct and indirect trust relationships.
    - Viewing and changing some attributes on a trust — for example, transitivity to non-Windows Kerberos realms.

To find more information about Windows Domain Manager, see "Support Tools Help" in the Tools and Settings Collection.

**Netmon.exe: Network Monitor**
**Category**
A limited version of Network Monitor is included in Windows Server 2003, Windows XP, and Windows 2000. The full version of Network Monitor is included with Microsoft Systems Management Server.

**Version compatibility**
Network Monitor is supported for Windows Server 2003, Windows XP, and Windows 2000.

Network Monitor enables you to capture network traces which can be used in troubleshooting most network issues.

**Setspn.exe: Manipulate Service Principal Names for Accounts**
**Category**
Setspn is included in the Windows Server 2003 Support Tools.

**Version compatibility**
Setspn is supported for Windows Server 2003 and Windows XP Professional.

This command-line tool sets service principal names. SPNs are used to locate a target principal name for running a service. Because SPNs are security-sensitive, SPNs can only be set for user objects if you have domain administrator credentials.

It is not usually necessary to modify SPNs. They are set up by a computer when it joins a domain and when services are

installed on the computer. In some cases, however, this information can become stale. For example, if the computer name is changed, the SPNs for installed services would need to be changed to match the new computer name. Also, some services and applications might require manual modification of a service account's SPN information to correctly authenticate.

Domain administrators can use Setspn to:

- View the current SPNs.
- Reset the account's default SPNs.
- Add SPNs.
- Delete supplemental SPNs.

To find more information about Setspn, see "Support Tools Help" in the Tools and Settings Collection.

**Kerberos Protocol Registry Settings**
Registry settings in the following hives are associated with the Kerberos protocol:

- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Domains
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\HostToRealm
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\UserList
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Eventlog\System
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, instead of editing the registry directly, use Group Policy or other Windows tools such as MMC to accomplish tasks. If you must edit the registry, use extreme caution.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Domains**
The Domains subkey stores information about non-Windows Kerberos realms.

To change the value of the entries in this subkey, use Kerberos Setup (Ksetup.exe), a tool included in Windows Server 2003 Support Tools. Do not use the registry editor.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\HostToRealm**
This subkey stores Host to Realm mapping information. The subkey does not exist in the registry by default.

**SpnMappings**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\HostToRealm

**Version**
Windows Server 2003 and Windows XP

This entry is used to create a HostToRealm mapping table.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters**
The Parameters subkey stores configuration options for the Kerberos V5 authentication protocol in Windows Server 2003.

**AllowTgtSessionKey**
**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry controls whether session keys are exported with initial or cross-realm TGTs. This entry does not exist in the registry by default. The default value is false due to security concerns.

**CacheS4UTickets**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003

This entry enables and disables Service-for-User (S4U) caching. This entry does not exist in the registry by default. The default value is true.

**ClientIpAddresses**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry controls adding IP addresses in KRB_AS_REQ, thus forcing the Caddr field to contain IP addresses in all tickets. This entry does not exist in the registry by default. The default value is false, due to potential DHCP client and network address translation (NAT) issues.

**DefaultEncryptionType**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry controls the default encryption type for PreAuth. This entry does not exist in the registry by default. The default value is KERB_ETYPE_RC4_HMAC_NT.

**FarKdcTimeout**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry controls the time-out interval, in minutes, that is used to invalidate a domain controller from a different site in the domain controller cache. This entry does not exist in the registry by default. The default value is 10 minutes.

**KdcBackoffTime**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry specifies a time value, in seconds, between successive calls to the Key Distribution Center (KDC) if the previous call failed. This entry does not exist in the registry by default. The default value is 5 seconds.

**KdcSendRetries**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry controls the number of retry attempts that a client makes in order to contact a KDC. This entry does not exist in the registry by default. The default value is 3.

**KdcWaitTime**

**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry specifies a time value, in seconds, that is used to time out the Winsock calls. This entry does not exist in the registry by default. The default value is 5 seconds.

**KerbDebugLevel**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry controls the verboseness level of debug log macros. This entry does not exist in the registry by default. The default value is 0, which is no logging; however, if you are using a checked build, the default value is 1, which is error logging.

### Common Debug Values

| Verboseness Level | Value |
|---|---|
| Error s | 0x00000001 |
| Warnings | 0x00000002 |
| Tracing | 0x00000004 |
| API tracing | 0x00000008 |
| Credential related tracing | 0x00000010 |
| Security Context tracing | 0x00000020 |
| Logon Session tracing | 0x00000040 |
| Logon tracing | 0x00000100 |
| KDC tracing | 0x00000200 |
| Detailed Security Context tracing | 0x00000400 |
| Time related tracing | 0x00000800 |
| User related tracing | 0x00001000 |
| Leak related tracing | 0x00002000 |
| WinSock related tracing | 0x00004000 |
| SPN cache tracing | 0x00008000 |
| S4U Errors | 0x00010000 |
| S4U tracing | 0x00020000 |
| Loopback tracing | 0x00080000 |
| Ticket renewal tracing | 0x00100000 |
| User to User tracing | 0x00200000 |
| Locks tracing | 0x01000000 |

**LogLevel**
**Registry path**

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

If this entry is set to anything non-zero, all Kerberos errors are logged in the system event log. This entry does not exist in the registry by default. The default value is 0.

**LogToFile**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry enables debug tracing and logging to a file. This entry does not exist in the registry by default. The default value is false.

**MaxPacketSize**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry controls the maximum size, in bytes, for using User Datagram Protocol (UDP). If the packet size is bigger than this value, TCP is used. This entry does not exist in the registry by default. The default value is 1465 bytes.

**MaxReferralCount**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry specifies the number of KDC referrals that a client pursues before the client gives up. This entry does not exist in the registry by default. The default value is 3.

**MaxTokenSize**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry specifies the maximum value, in bytes, of the Kerberos token size. Use this entry to allow query context attributes to be modified to return a value large enough for tickets containing large numbers of groups. It is recommended that this value remain less than 65,000. This entry does not exist in the registry by default. The default value is 12,000 bytes.

**NearKdcTimeout**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry specifies the time-out interval, in minutes, that is used to invalidate a domain controller in the same site in the domain controller cache. This entry does not exist in the registry by default. The default value is 30 minutes.

**RequestOptions**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003 and Windows XP

This entry indicates whether there are additional options that must be sent as KDC options in ticket-granting service (TGS) requests (KRB_TGS_REQ). This entry is meant to accommodate future modifications to the KDC options and can be any RFC 1510 value. This entry does not exist in the registry by default. The default value is 0x10000.

**RetryPDC**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003

This entry controls whether the primary domain controller (PDC) will be contacted for password-expiry errors for KRB_AS_REQ. This entry does not exist in the registry by default. The default value is false.

**S4UCacheTimeout**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003

This entry controls the lifetime, in minutes, of the S4U negative cache entries, which are used to restrict how many S4U proxy requests are made from a specific computer. This entry does not exist in the registry by default. The default value is 15 minutes.

**S4UTicketLifetime**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003

This entry controls the lifetime, in minutes, of tickets obtained by S4U proxy requests. This entry does not exist in the registry by default. The default value is 15 minutes.

**SkewTime**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003

This entry controls the maximum time difference, in minutes, that is permitted between server, client, and KDC. This entry does not exist in the registry by default. The default value is 5 minutes.

**SpnCacheTimeout**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003 and Windows XP

This entry specifies the lifetime, in minutes, of the service principal name (SPN) cache entries. This entry does not exist in the registry by default. The default value is 15 minutes.

**StartupTime**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003, Windows XP, and Windows 2000

This entry specifies the time, in seconds, that Windows waits for the KDC to start before Windows gives up. This entry does not exist in the registry by default. The default value is 120 seconds.

**TgtRenewalTime**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters

**Version**
Windows Server 2003

This entry controls the amount of time, in seconds, that a Kerberos client waits before it tries to renew a TGT before the ticket expires. Only applies to initial TGTs. This entry does not exist in the registry by default. The default value is 600 seconds.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\UserList**
The UserList subkey stores entries that associate a Kerberos security principal to a local Windows Server 2003 user account.

Computers that are running Windows Server 2003 can use another KDC — instead of a KDC in an Active Directory domain — to administer authentication. For ease of use, you can map a Kerberos security principal, such as the name of a principal or a realm, to a local Windows user account.

This subkey stores mappings that you enter when you use the **/MapUser** parameter with Kerberos Setup (Ksetup.exe), a tool that is included in Windows Server 2003 Support Tools. Ksetup.exe adds the entries to the registry.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Eventlog\System**
This subkey stores the event log configuration information.

**Kerberos**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Eventlog\System

**Version**
Windows Server 2003, Windows XP, and Windows 2000

The Kerberos subkey stores configuration information for Event Viewer about the log sources specified by the Kerberos protocol in the registry path. The configuration information tells Event Viewer where to find and how to display Kerberos authentication-related events. Event Viewer displays those events as part of the log specified by the system-log name in the registry path.

**Sources**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Eventlog\System

**Version**
Windows Server 2003, Windows XP, and Windows 2000

The Sources entry specifies that Kerberos writes events to this log.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc**
This subkey stores KDC configuration information.

**ExistingConnectionTimeout**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc

**Version**
Windows Server 2003 and Windows 2000

This entry controls the time-out, in seconds, for input/outputs (I/Os) on this context until the first byte of data is received. This entry does not exist in the registry by default. The default value is 60 seconds.

**KdcDebugLevel**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc

**Version**
Windows Server 2003

This entry controls the verboseness level of debug log macros. This entry does not exist in the registry by default. The default value is 0, which is no logging; however, if you are using a checked build, the default value is 1, which is error logging.

### Common Debug Values

| Verboseness Level | Value |
|---|---|
| Error s | 0x00000001 |
| Warnings | 0x00000002 |
| Tracing | 0x00000004 |
| API tracing | 0x00000008 |

| | |
|---|---|
| Credential related tracing | 0x00000010 |
| Security Context tracing | 0x00000020 |
| Logon Session tracing | 0x00000040 |
| Logon tracing | 0x00000100 |
| KDC tracing | 0x00000200 |
| Detailed Security Context tracing | 0x00000400 |
| Time related tracing | 0x00000800 |
| User related tracing | 0x00001000 |
| Leak related tracing | 0x00002000 |
| WinSock related tracing | 0x00004000 |
| SPN cache tracing | 0x00008000 |
| S4U Errors | 0x00010000 |
| S4U tracing | 0x00020000 |
| Loopback tracing | 0x00080000 |
| Ticket renewal tracing | 0x00100000 |
| User to User tracing | 0x00200000 |
| Locks tracing | 0x01000000 |
| Use Extended Errors | 0x10000000 |

**KdcDontCheckAddresses**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc

**Version**
Windows Server 2003

This entry controls whether the KDC checks if the client IP addresses match one of the IP addresses in the TGT Caddr field.

- If the entry does not exist or if DWORD = 0, client addresses are checked.
- If DWORD = 1, client addresses are not checked.

The client must send the IP address in the KRB_TGS_REQ, which is not done by default in Windows operating systems. To enable sending IP addresses as a Kerberos client, see “ClientIpAddresses” earlier in this guide. This entry does not exist in the registry by default. The default value is true due to potential DHCP and NAT issues.

**KdcExtraLogLevel**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc

**Version**
Windows Server 2003

This entry controls extra KDC logging in event logs and audits. This entry does not exist in the registry by default. The default value is 2.

### Extra KDC Log Levels

| Log Level | Value |
| --- | --- |
| Audit SPN unknown errors | 0x1 |
| Log detailed PKINIT1 errors | 0x2 |
| Log all KDC errors with KLIN information | 0x4 |

1 PKINIT is an Internet Engineering Task Force (IETF) Internet draft for "Public key Cryptography for Initial Authentication in Kerberos."

**KdcIssueForwardedTickets**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc

**Version**
Windows Server 2003

This entry controls whether the KDC issues forwarded tickets. This entry does not exist in the registry by default. The default value is true.

**KdcUseClientAddresses**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc

**Version**
Windows Server 2003

This entry controls whether the KDC accepts client IP addresses in the KRB_AS_REQ:

- If the entry does not exist or if DWORD = 0, client IP addresses are not accepted.
- If DWORD = 1, client IP addresses are accepted.

In order for the KDC to propagate the client IP addresses, the client must send them in the KRB_TGS_REQ, which does not occur by default in Windows operating systems. To enable sending IP addresses as a Kerberos client, see "ClientIpAddresses" earlier in this guide. This entry does not exist in the registry by default. The default value is false due to potential DHCP and NAT issues.

**KdcUseRequestedEtypesForTickets**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc

**Version**
Windows Server 2003

This entry controls whether the encryption type for the service ticket is:

- The first encryption type in the requested encryption types, if true.
- The strongest encryption type that the server supports, if false.

This entry does not exist in the registry by default. The default value is false.

**MaxDatagramReplySize**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc

**Version**
Windows Server 2003

This entry controls the maximum size, in bytes, of a UDP packet for the KRB_TGS_REP and KRB_AS_REP, before the KDC returns a KRB_ERR_RESPONSE_TOO_BIG error, which requires the client to change to TCP. This entry does not exist in the registry by default. The default value is 1465 bytes.

**Kerberos Authentication Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with Kerberos V5 Authentication.

**Group Policy Settings Associated with Kerberos V5 Authentication**

| Group Policy Setting | Description |
|---|---|
| **User Rights Assignment:** | |
| Impersonate a client after authentication | Windows 2000 security setting that was first introduced in SP4 for Windows 2000. When you assign this user right to a user, you permit programs that run on behalf of that user to impersonate a client. This security setting helps to prevent unauthorized servers from impersonating clients that connect to it through methods such as remote procedure calls (RPC) or named pipes.<br><br>By default, members of the Administrators group and the System account are assigned this user right. The following components also are assigned this user right:<br><br>• Services that are started by the Service Control Manager<br><br>• Component Object Model (COM) servers that are started by the COM infrastructure and that are configured to run under a specific account |
| **Kerberos Policy:** | |
| Enforce user logon restrictions | Determines whether the KDC validates every request for a session ticket against the user rights policy on the target computer. When this policy is enabled, the user requesting the session ticket must have the right either to Log on locally or to Access this computer from network. Validation of each request is optional because the extra step takes time and might slow network access to services. By default, this policy is enabled. |
| Maximum lifetime for service ticket | Determines the maximum amount of time (in minutes) that a ticket granted for a service (that is, a session ticket) can be used to access the service. If the setting is zero minutes, the ticket never expires. Otherwise, the setting must be greater than ten minutes and less than the setting for Maximum lifetime for user ticket. By default, the setting is 600 minutes (10 hours). |
| Maximum lifetime for user ticket | Determines the maximum amount of time (in hours) that a user's TGT can be used. When a user's TGT expires, a new one must be requested or the existing one must be renewed. By default, the setting is ten hours. |
| Maximum lifetime for user ticket renewal | Determines the longest period of time (in days) that a TGT can be used if it is repeatedly renewed. By default, the setting is seven days. |
| Maximum tolerance for computer clock synchronization | Determines the maximum difference (in minutes) that Kerberos will tolerate between the time on a client's clock and the time on a server's clock while still considering the two clocks synchronous. By default, the setting is five minutes. |

For more information about these Group Policy settings, see "Account Policy Settings."

#### Network Ports Used by the Kerberos V5 Protocol

### Network Ports Used During Kerberos Authentication

| Service Name | UDP | TCP |
|---|---|---|
| DNS | 53 | 53 |
| Kerberos | 88 | 88 |

# TLS/SSL Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TLS/SSL Technical Reference**

How can you secure data that is being sent between applications across an untrusted network? The Transport Layer Security (TLS) protocol, Version 1, provides one answer to this common security question. TLS, along with the Secure Sockets Layer (SSL) protocol, Versions 2.0 and 3.0 authenticates and secures data transfers by using certificate-based authentication and symmetric encryption keys.

This section explains what TLS and SSL are, how TLS and SSL use the Security Support Provider Interface (SSPI) and Schannel SSP to secure the transfer of data, and which tools and settings are available for configuring TLS and SSL.

**In this section**


- What is TLS/SSL?
- How TLS/SSL Works
- TLS/SSL Tools and Settings

# What is TLS/SSL?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is TLS/SSL?**

### In this section

- Common TLS and SSL Scenarios
- Technologies Related to Schannel
- TLS and SSL Dependencies

One problem when you administer a network is securing data that is being sent between applications across an untrusted network. You can use TLS/SSL to authenticate servers and clients and then use it to encrypt messages between the authenticated parties.

The Transport Layer Security (TLS) protocol, Secure Sockets Layer (SSL) protocol, versions 2.0 and 3.0, and the Private Communications Transport (PCT) protocol are based on public key cryptography. The Security Channel (Schannel) authentication protocol suite provides these protocols. All Schannel protocols use a client/server model.

In the authentication process, a TLS/SSL client sends a message to a TLS/SSL server, and the server responds with the information that the server needs to authenticate itself. The client and server perform an additional exchange of session keys, and the authentication dialog ends. When authentication is completed, SSL-secured communication can begin between the server and the client using the symmetric encryption keys that are established during the authentication process.

For servers to authenticate to clients, TLS/SSL does not require server keys to be stored on domain controllers or in a database, such as the Microsoft Active Directory directory service. Clients confirm the validity of a server's credentials with a trusted root certification authority's (CA's) certificates, which are loaded when you install Microsoft Windows Server 2003. Therefore, unless user authentication is required by the server, users do not need to establish accounts before they create a secure connection with a server.

**History and Standards for TLS and SSL**

SSL was developed by Netscape Communications Corporation in 1994 to secure transactions over the World Wide Web. Soon after, the Internet Engineering Task Force (IETF) began work to develop a standard protocol that provided the same functionality. They used SSL 3.0 as the basis for that work, which became the TLS protocol. The implementation of the TLS protocol in Windows Server 2003 closely follows the specification defined in Request for Comments (RFC) 2246, The TLS Protocol Version 1.0. For more information about TLS, see RFC 2246 in the IETF RFC database.

TLS and SSL are most widely recognized as the protocols that provide secure HTTP (HTTPS) for Internet transactions between Web browsers and Web servers. TLS/SSL can also be used for other application level protocols, such as File Transfer Protocol (FTP), Lightweight Directory Access Protocol (LDAP), and Simple Mail Transfer Protocol (SMTP). TLS/SSL enables server authentication, client authentication, data encryption, and data integrity over networks such as the World Wide Web.

**Differences between TLS and SSL**

Although there are some slight differences between SSL 3.0 and TLS 1.0, this reference refers to the protocol as TLS/SSL.

### Note

- Although their differences are minor, TLS 1.0 and SSL 3.0 are not interchangeable. If the same protocol is not supported by both parties, the parties must negotiate a common protocol to communicate successfully.

**TLS Enhancements to SSL**

- The keyed-Hashing for Message Authentication Code (HMAC) algorithm replaces the SSL Message Authentication Code (MAC) algorithm.

  HMAC produces more secure hashes than the MAC algorithm. The HMAC produces an integrity check value as the MAC does, but with a hash function construction that makes the hash much harder to break. For more information about the HMAC, see "Hash Algorithms in The Handshake Layer in TLS/SSL Architecture" in How TLS/SSL Works.

- TLS is standardized in RFC 2246.
- Many new alert messages are added.
- In TLS, it is not always necessary to include certificates all the way back to the root CA. You can use an intermediary authority.
- TLS specifies padding block values that are used with block cipher algorithms. RC4, which is used by Microsoft, is a streaming cipher, so this modification is not relevant.
- Fortezza algorithms are not included in the TLS RFC, because they are not open for public review. (This is Internet Engineering Task Force (IETF) policy.)
- Minor differences exist in some message fields.

**Benefits of TLS/SSL**

TLS/SSL provides numerous benefits to clients and servers over other methods of authentication, including:

- Strong authentication, message privacy, and integrity
- Interoperability
- Algorithm flexibility
- Ease of deployment
- Ease of use

**Strong authentication, message privacy, and integrity**

TLS/SSL can help to secure transmitted data using encryption. TLS/SSL also authenticates servers and, optionally, authenticates clients to prove the identities of parties engaged in secure communication. It also provides data integrity through an integrity check value. In addition to protecting against data disclosure, the TLS/SSL security protocol can be used to help protect against masquerade attacks, man-in-the-middle or bucket brigade attacks, rollback attacks, and replay attacks.

**Interoperability**

TLS/SSL works with most Web browsers, including Microsoft Internet Explorer and Netscape Navigator, and on most operating systems and Web servers, including the Microsoft Windows operating system, UNIX, Novell, Apache (version 1.3 and later), Netscape Enterprise Server, and Sun Solaris. It is often integrated in news readers, LDAP servers, and a variety of other applications.

**Algorithm flexibility**

TLS/SSL provides options for the authentication mechanisms, encryption algorithms, and hashing algorithms that are used during the secure session.

### Note

- Data can be encrypted and decrypted, but you cannot reverse engineer a hash. Hashing is a one-way process. Running the process backward will not create the original data. This is why a new hash is computed and then compared to the sent hash.

**Ease of deployment**

Many applications use TLS/SSL transparently on a Windows Server 2003 operating system. You can use TLS for more secure browsing when you are using Internet Explorer and Internet Information Services (IIS) and, if the server already has a server certificate installed, you only have to select the check box.

**Ease of use**

Because you implement TLS/SSL beneath the application layer, most of its operations are completely invisible to the client. This allows the client to have little or no knowledge of the security of communications and still be protected from attackers.

**Limitations of TLS/SSL**
There are a few limitations to using TLS/SSL, including:

**Increased processor load**
This is the most significant limitation to implementing TLS/SSL. Cryptography, specifically public key operations, is CPU-intensive. As a result, performance varies when you are using SSL. Unfortunately, there is no way to know how much performance you will lose. The performance varies, depending on how often connections are established and how long they last. TLS uses the greatest resources while it is setting up connections.

**Administrative overhead**
A TLS/SSL environment is complex and requires maintenance; the system administrator must configure the system and manage certificates.

**Common TLS/SSL Scenarios**
Many people think of TLS and SSL as protocols that are used with Web browsers to browse the Internet more securely. However, they are also general purpose protocols that can be used whenever authentication and data protection are necessary. For example, you can use TLS/SSL for:

- SSL-secured transactions with an e-commerce Web site
- Authenticated client access to an SSL-secured Web site
- Remote access
- SQL access
- E-mail

This is not an exhaustive list. In fact, the ability to access these protocols through Security Service Provider Interface (SSPI) means that you can use them for just about any application. Many applications are being modified to take advantage of the features of TLS/SSL.

**SSL-secured transactions with an e-commerce Web site**
This situation is a typical use of SSL between a browser and a Web server. An example is an e-commerce shopping site where clients need to provide their credit card numbers. The protocol first confirms that the certificate of the Web site is valid, and then sends the client's credit card information as cipher text. For this type of transaction, where the server's certificate comes from a trusted source, authentication only occurs for the server. TLS/SSL must be enabled for the Web page, such as an order form, where the data transactions occur.

**Authenticated client access to an SSL-secured Web site**
Both the client and server need certificates from a mutually-trusted certification authority (CA). With Schannel, client certificates can be mapped on a one-to-one or many-to-one basis to their Windows Server 2003 user or computer accounts, and they can be managed by Active Directory Users and Computers. Users can then be authenticated to a Web site without needing to supply a password.

Many-to-one mapping has several uses. For example, if you want to give several users access to confidential material, you can create a group, map the users' certificates to the group, and give the group the necessary permissions to the material.

In one-to-one mapping, the server has a copy of the client's certificate; whenever the client logs in; the server verifies that they are identical. This one-to-one mapping is typically used for private material, such as a banking Web site where only one individual has the right to view a personal account.

**Remote access**
In this situation, telecommuting is a common use for Schannel. You can use this technology to provide authentication and data protection when users remotely log in to Windows-based systems or networks. Users can more securely access their e-mail or enterprise applications from home or while traveling, reducing the risk of exposure of the information to anyone on the Internet.

**SQL access**
With Microsoft SQL Server, you can require authentication of the client when connecting to the server running SQL Server. Either the client or server can be configured to require encryption of the data that is transferred between them. Very sensitive information, such as financial or medical databases, can be protected to prevent unauthorized access and disclosure of information about the network.

**E-mail**
When using Exchange servers, you can use Schannel to help protect data as it moves from server to server on the intranet or Internet. Full end-to-end security might require the use of Secure/Multipurpose Internet Mail Extensions (S/MIME); however, helping to protect data in a server-to-server exchange allows companies to use the Internet to securely transfer e-mail among divisions within the same company, subsidiaries, and partners. This can be done regardless of whether S/MIME is used.

**Technologies Related to Schannel**
In the Windows Server 2003 operating system, the Schannel Security Support Provider SSP authentication protocol suite contains TLS. The following diagram shows where Schannel fits among these and other technologies in Windows Server 2003. Client applications or server applications use Secur32.dll by way of SSPI calls to communicate with the Local Security Authority Subsystem (LSASS).

### Schannel Architecture




Below is a description of the components that participate in TLS/SSL.

### Security Subsystem Components used in TLS/SSL Authentication

| Component | Description |
|---|---|
| Schannel.dll | The Transport Layer Security (TLS) and Secure Sockets Layer (SSL) authentication protocol. This protocol provides authentication over an encrypted channel instead of a less-secure clear channel. |
| Lsasrv.dll | The LSA Server service, which both enforces security policies and acts as the security package manager for the LSA. |
| Netlogon.dll | TLS-relevant services that the Net Logon service performs:<br>• Net Logon passes the user's certificate information through an SSL-secured channel to the domain controller to map the user certificate to a user account. |
| Secur32.dll | The multiple authentication provider that implements SSPI. |

| | |
|---|---|
| Ntdsa.dll | The directory service module, which supports the Windows Server 2003 and Windows 2000 replication protocol and LDAP, and manages partitions of data. |

The authentication protocol suite is enabled by the Schannel .SSP, which is supported by the SSPI application programming interface (API) that provides the security services for Windows Server 2003 operating systems.

**TLS and SSL Dependencies**

TLS and SSL depend on several related technologies and resources to function properly. The following section describes these technologies and resources and summarizes how they relate to TLS/SSL authentication.

**Operating system**

TLS and SSL authentication relies on client functionality that is built into the Microsoft Windows Server 2003 operating system, the Microsoft Windows XP operating system, and the Microsoft Windows 2000 operating system. If a client or server is running an operating system that does not support TLS/SSL, it cannot use TLS/SSL authentication.

**TCP/IP network connectivity**

For TLS or SSL authentication to occur, there must be TCP/IP network connectivity between the client, and the target server. For more information about TCP/IP, see "TCP/IP Technical Reference."

**Active Directory domain**

When a server application requires client authentication, Schannel automatically attempts to map the certificate that is supplied by the client to a user account. You can authenticate users who log on with a client certificate by manually creating mappings, which relate the certificate information to a Windows user account. After you create and enable a certificate mapping, each time a client presents a client certificate; your server application automatically associates that user with the appropriate Windows user account. You must use user accounts in the Active Directory directory service if you wish to use Active Directory certificate mapping. For more information about Active Directory, see "Active Directory Collection."

**Trusted certificate authorities**

Because authentication relies on digital certificates, certification authorities (CAs) such as Verisign or Microsoft Certificate Server are an important part of TLS/SSL. A CA is a mutually-trusted third party that confirms the identity of a certificate requestor (usually a user or computer), and then issues the requestor a certificate. The certificate binds the requestor's identity to a public key. CAs also renew and revoke certificates as necessary. For example, if a client is presented with a server's certificate, the client computer might try to match the server's CA against the client's list of trusted CAs. If the issuing CA is trusted, the client will verify that the certificate is authentic and has not been tampered with. For more information about PKI Technologies, see "PKI Technologies."

# How TLS/SSL Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How TLS/SSL Works**

**In this section**


- Schannel SSP Architecture
- TLS/SSL Architecture
- TLS/SSL Processes and Interactions
- Network Ports Used by TLS/SSL
- Related Information


TLS/SSL authenticates and secures data transfers by using certificate-based authentication and symmetric encryption keys. This section discusses how the RFC-standard TLS protocol is used in the Windows Server 2003 operating system.

This section is divided into five subsections:

- **Schannel SSP Architecture** illustrates how the Microsoft Security Support Provider Interface (SSPI) in Windows Server 2003 provides a mechanism of generic routines that can access the Secure Channel (Schannel) Security Support Provider (SSP).
- **TLS/SSL Architecture** discusses the Handshake and Record Layer and related sub-protocols of TLS/SSL, as well as the Schannel session cache.
- **TLS/SSL Protocol Processes and Interactions** illustrates full handshake protocol, application data flow, resuming a secure session, and renegotiation with what is included in TLS messages. Some processes — such as certificate mapping, and how internally renegotiations are handled — are specific to Windows systems and might or might not differ in other implementations of the TLS protocol.
- **Network Ports Used by the TLS/SSL** tabulates networks ports that are used for TLS/SSL
- **Related Information** lists links of related information.

The Windows Server 2003 operating system can use three related security protocols to provide authentication and secure communications over the Internet:

- Transport Layer Security (TLS) 1.0
- Secure Sockets Layer 3.0
- Secure Sockets Layer (SSL) 2.0

All three protocols provide authentication through the use of certificates and secure communication through a variety of possible cipher suites. The generic term *cipher suite* refers to a combination of protocols such as key exchange, bulk encryption, and message integrity. Because authentication relies on digital certificates, certification authorities (CAs) like Verisign are an important part of Secure Channel (Schannel). A CA is a mutually trusted third party that confirms the identity of a certificate requestor (usually a user or computer), and then issues the requestor a certificate. The certificate binds the requestor's identity to a public key. CAs also renew and revoke certificates as necessary. For example, if a client is presented with a server's certificate, the client computer might try to match the server's CA against the client's list of trusted CAs. If the issuing CA is trusted, the client will verify that the certificate is authentic and has not been tampered with. Microsoft Internet

Explorer and Internet Information Services (IIS) make use of these protocols, and preferably TLS, for Secure Hypertext Transfer Protocol (HTTPS).

This document focuses on TLS, because TLS is replacing SSL and PCT. TLS is standardized in RFC 2246 IETF RFC database.

**Note**

- Throughout this document *TLS/SSL* refers to the common protocol features of Transport Layer Security and Secure Sockets Layer. *TLS* refers to only Transport Layer Security, and *SSL* refers to only Secure Sockets Layer.

**Schannel SSP Architecture**

The Windows Server 2003 operating system implements the TLS/SSL protocols as a Security Support Provider SSP, a dynamic-link library (DLL) called Schannel that is supplied with the operating system. Which SSP is used depends on the capabilities of the computer on the other side of the connection and the configuration of the individual application that is being used.

The Microsoft Security Support Provider Interface (SSPI) is the foundation for authentication in Windows Server 2003. That is, applications and infrastructure services that require authentication use SSPI to provide it.

The SSPI is the implementation of the Generic Security Service API (GSSAPI) in Windows Server 2003. For more information about GSSAPI, see RFC 2743 and RFC 2744 in the IETF RFC database.

The default SSPs in Windows Server 2003 — Kerberos, NTLM, Digest, Schannel, and Negotiate authentication protocols — are incorporated into the SSPI in the form of DLLs. Additional SSPs can be incorporated if they can interoperate with the SSPI.

**SSPI Architecture**




In the Windows Server 2003 operating system, SSPI provides a mechanism that carries authentication tokens over the existing communication channel between the client and server. When two parties need to be authenticated so that they can communicate more securely, the requests for authentication are routed to the SSPI, which completes the authentication process, regardless of the network protocol currently in use. The SSPI returns transparent binary large objects, and then these are passed between the applications, at which point they can be passed to the SSPI layer on that side. Thus, the SSPI enables an application to use various security models available on a computer or network without changing the interface to the security system.

The following table describes the SSP components that are plugged into the SSPI. Each of the protocols in the table is used in different ways in Windows Server 2003 to promote more secure communication in an insecure network environment.

**SSP Layer Components**

| Component | Description |
|---|---|
| Kerberos V5 authentication | An industry-standard protocol that is used with either a password or a smart card for interactive logon. It is also the preferred authentication method for services in Windows 2000 and Windows Server 2003. |
| NTLM authentication | A challenge-response protocol that is used to provide compatibility with versions of Windows earlier than Windows 2000. |
| Digest authentication | An industry standard that is used in Windows Server 2003 for Lightweight Directory Access Protocol (LDAP) and Web authentication. Digest transmits credentials across the network as a Message Digest 5 (MD5) hash or message digest. |

| | |
|---|---|
| Schannel | An SSP that implements SSL and TLS. Schannel is used for applications used in cross-organization environments, such as Web-based server authentication, in which a user attempts to access a secure Web server or corporate access using VPN. |
| Negotiate | An SSP that can be used to negotiate a specific authentication protocol. When an application calls into SSPI to log on to a network, it can specify an SSP to process the request. If the application specifies Negotiate, Negotiate analyzes the request and picks the best SSP to handle the request, based on customer-configured security policy. |

**Secure Channel SSP**

You can use the Secure Channel (Schannel) SSP for access to Web-enabled services, such as e-mail or personal information served on Web pages.

The Schannel SSP uses public key certificates to authenticate parties. It includes four authentication protocols in its suite. When authenticating parties, it will select one of the four protocols in the following order of preference:

- TLS version 1.0
- SSL version 3.0
- PCT. PCT is turned off by default in Windows Server 2003. PCT has been superseded by Secure Sockets Layer 3.0 and the TLS protocol. The Schannel SSP supports PCT 1.0 for backward compatibility only, but this protocol might not be available in future releases.
- SSL version 2.0

Schannel then selects the most preferred authentication protocol that both parties can support. For example, if a server supports all four Schannel protocols and the client supports only SSL 3.0 and PCT, Schannel uses SSL 3.0 for authentication.

**TLS/SSL Architecture**

The Schannel authentication protocol suite is based on public key cryptography. The Schannel suite includes Transport Layer Security (TLS), Secure Sockets Layer (SSL) version 3.0, SSL version 2.0, and Private Communications Transport (PCT). All Schannel protocols are based on a client/server model. An Schannel client sends a message to a server, and the server responds with the information needed to authenticate itself. The client and server perform an additional exchange of session keys, and the authentication dialogue ends. When authentication is completed, secure communication can begin between the server and the client using the secret keys established during the authentication process.

Schannel does not require server keys to be stored on domain controllers or in a database, such as Active Directory. Clients, however, must be able to confirm the validity of credentials with a trusted authority. Schannel validates the credentials with the root CA's certificates, which are loaded when you install Windows Server 2003. Therefore, users do not need to establish accounts before authenticating and creating a secure connection with a server.

The TLS/SSL security protocol is layered between the application protocol layer and the TCP/IP layer, where it can secure and send application data to the transport layer. Because it works between the application layer and the transport layer, TLS/SSL can support multiple application layer protocols.

TLS/SSL assumes that a connection-oriented transport, typically TCP, is in use. The protocol allows client/server applications to detect the following security risks:

- Message tampering
- Message interception
- Message forgery

The TLS/SSL protocol can be divided into two layers. The first layer consists of the application protocol and the three Handshake sub-protocols: the Handshake Protocol, the Change Cipher Spec Protocol, and the Alert Protocol. The second layer is the Record Protocol. The following figure illustrates the various layers and their components.

**TLS/SSL Protocol Layers**




**The Handshake Protocols**

The Handshake protocols of the TLS/SSL protocol are responsible for establishing or resuming secure sessions. The main goals of this layer are to:

- Negotiate cipher suites and compression algorithms.
- Authenticate the server to the client and, optionally, authenticate the client to the server through certificates and public or private keys.
- Exchange random numbers and a pre-master secret. Together with some further data, these values will be used to create the shared secret key that the Record Layer will use to hash and encrypt application data. The shared secret key is called the Master Secret.

The three Handshake sub-protocols are:

- **Handshake**. This sub-protocol is used to negotiate session information between the client and the server. The session information consists of a session ID, peer certificates, the cipher specification to be used, the compression algorithm to be used, and a shared secret that is used to generate keys.
- **Change Cipher Spec**. This sub-protocol changes the keying material that is used for encryption between the client and server. The keying material is raw data that is used to create keys for cryptographic use. The Change Cipher Spec sub-protocol consists of a single message to tell other party in the TLS/SSL session, who is known as the peer, that the sender wants to change to a new set of keys. The key is computed from the information that is exchanged by the Handshake sub-protocol.
- **Alert**. This sub-protocol uses messages to indicate a change in status or an error condition to the peer. There are a wide variety of alerts to notify the peer of both normal and error conditions. A full list can be found in RFC 2246, "The TLS Protocol Version 1.0." Alerts are commonly sent when the connection is closed, an invalid message is received, a message cannot be decrypted, or the user cancels the operation.

**Handshake Protocol Functions**

The Handshake protocol provides a number of very important security functions. It performs a set of exchanges that starts authentication and negotiates the encryption, hash, and compression algorithms.

**Authentication**

A certificate is a digital form of identification that is usually issued by a certification authority (CA) and contains identification information, a validity period, a public key, a serial number, and the digital signature of the issuer. For authentication purposes, the Handshake Protocol uses an X.509 certificate to provide strong evidence to a second party that helps prove the identity of the party that holds the certificate and the corresponding private key.

A CA is a mutually trusted third party that confirms the identity of a certificate requestor (usually a user or computer), and then issues the requestor a certificate. The certificate binds the requestor's identity to a public key. CAs also renew and revoke certificates as necessary. For example, if a client is presented with a server's certificate, the client computer might try to match the server's CA against the client's list of trusted CAs. If the issuing CA is trusted, the client will verify that the certificate is authentic and has not been tampered with. Finally, the client will accept the certificate as proof of identity of the server.

**Encryption**

There are two main types of encryption: *symmetrickey* (also known as *shared secret key*) and *asymmetrickey* (also known as *public key* or *public-private key*). TLS/SSL uses symmetric key for bulk encryption and public key for authentication and key exchange.

- **Symmetric Key (also known as Private Key)**. In symmetric key encryption, the same key is used to encrypt and decrypt the message. If two parties want to exchange encrypted messages securely, they must both possess a copy of the same symmetric key. Symmetric key cryptography is often used for encrypting large amounts of data because it is computationally faster than asymmetric cryptography. Typical algorithms include Data Encryption Standard (DES), Triple DES (3-DES), RC2, RC4, and Advanced Encryption Standard (AES).

- **Asymmetric Key, also known as Public Key**. Asymmetric key encryption uses a pair of keys that have been derived simultaneously through a complex mathematical process. One of the keys is made public, typically by asking a CA to publish the public key in a certificate for the certificate-holder (also called the subject). The private key is kept secret by the subject and never revealed to anyone. The keys work together, with one being used to perform the inverse operation of the other: If the public key is used to encrypt data, only the private key of the pair can decrypt it; if the private key is used to encrypt, the public key must be used to decrypt. This relationship allows a public key encryption scheme to do two important things. First, anyone can obtain the public key for a subject and use it to encrypt data that only the user with the private key can decrypt. Second, if a subject encrypts data using its private key, anyone can decrypt the data by using the corresponding public key. This is the foundation for digital signatures. The most common algorithm is Rivest, Shamir, and Adleman (RSA).

TLS/SSL uses public key encryption to authenticate the server to the client and, optionally, the client to the server. Public key cryptography is also used to establish a session key. The session key is used in symmetric algorithms to encrypt the bulk of the data with the faster, less processor-intensive symmetric key encryption.

#### Hash Algorithms

During the handshake process, the client and server agree on the hash algorithm. A hash is a one-way mapping of values to a smaller set of representative values, so that the size of the resulting hash is smaller than the original message and the hash is unique to the original data. A hash is similar to a fingerprint: a fingerprint is unique to the individual and is much smaller than the original person. Hashing is used to establish data integrity during transport. Two common hash algorithms are Message Digest 5 (MD5) and Standard Hash Algorithm 1 (SHA-1). MD5 produces a 128-bit hash value and SHA-1 produces a 160-bit value.

### Note

- Data can be encrypted and decrypted, but you cannot reverse engineer a hash. Hashing is a one-way process. Running the process backward does not recreate the original data. This is why a new hash is computed and then compared to the sent hash.

The hash algorithm includes a value that is used to check the integrity of the transmitted data. This value is established by using either a message authentication code (MAC) or a hashed message authentication code (HMAC). The MAC uses a mapping function to represent the message data as a fixed-length, preferably smaller, value and then hashes the message. The MAC ensures that the data has not been modified during transmission. The difference between a MAC and a digital signature is that a digital signature is also an authentication method. SSL uses a MAC.

The HMAC is similar to the MAC but uses a hash algorithm in combination with a shared secret key. The shared secret key is appended to the data to be hashed. This makes the hash more secure because both parties must have the same shared secret key to prove the data is authentic. TLS uses an HMAC. For more information about HMAC, see RFC 1024 in the IETF RFC Database.

#### The Change Cipher Spec Protocol

The Change Cipher Spec Protocol signals a transition of the cipher suite to be used on the connection between the client and server. This protocol is composed of a single message which is encrypted and compressed with the current cipher suite. This message consists of a single byte with the value 1. Message after this will be encrypted and compressed using the new cipher suite.

#### The Alert Protocol

The Alert Protocol includes event-driven alert messages that can be sent from either party. Following an alert message, the session is either ended or the recipient is given the choice of whether or not to end the session. The alerts are defined in the TLS specification in RFC 2246. The following table lists the possible alert messages and their descriptions. Schannel SSP will only generate these alert messages at the request of the application.

### Event-Driven Alert Messages from the Alert Sub-protocol

| Alert Message | Description |
|---|---|

| | |
|---|---|
| close_notify | Notifies the recipient that the sender will not send any more messages on this connection. |
| unexpected_message | Received an inappropriate message This alert should never be observed in communication between proper implementations. This message is always fatal. |
| bad_record_mac | Received a record with an incorrect MAC. This message is always fatal. |
| decryption_failed | Decryption of a TLSCiphertext record is decrypted in an invalid way: either it was not an even multiple of the block length or its padding values, when checked, were not correct. This message is always fatal. |
| record_overflow | Received a TLSCiphertext record which had a length more than 2^14+2048 bytes, or a record decrypted to a TLSCompressed record with more than 2^14+1024 bytes. This message is always fatal. |
| decompression_failure | Received improper input, such as data that would expand to excessive length, from the decompression function. This message is always fatal. |
| handshake_failure | Indicates that the sender was unable to negotiate an acceptable set of security parameters given the options available. This is a fatal error. |
| bad_certificate | There is a problem with the certificate, for example, a certificate is corrupt, or a certificate contains signatures that cannot be verified. |
| unsupported_certificate | Received an unsupported certificate type. |
| certificate_revoked | Received a certificate that was revoked by its signer. |
| certificate_expired | Received a certificate has expired or is not currently valid. |
| certificate_unknown | An unspecified issue took place while processing the certificate that made it unacceptable. |
| illegal_parameter | Violated security parameters, such as a field in the handshake was out of range or inconsistent with other fields. This is always fatal. |
| unknown_ca | Received a valid certificate chain or partial chain, but the certificate was not accepted because the CA certificate could not be located or could not be matched with a known, trusted CA. This message is always fatal. |
| access_denied | Received a valid certificate, but when access control was applied, the sender did not proceed with negotiation. This message is always fatal. |
| decode_error | A message could not be decoded because some field was out of the specified range or the length of the message was incorrect. This message is always fatal. |
| decrypt_error | Failed handshake cryptographic operation, including being unable to correctly verify a signature, decrypt a key exchange, or validate a finished message. |
| export_restriction | Detected a negotiation that was not in compliance with export restrictions; for example, attempting to transfer a 1024 bit ephemeral RSA key for the RSA_EXPORT handshake method. This message is always fatal. |
| protocol_version | The protocol version the client attempted to negotiate is recognized, but not supported. For example, old protocol versions might be avoided for security reasons. This message is always fatal. |
| insufficient_security | Failed negotiation specifically because the server requires ciphers more secure than those supported by the client. Returned instead of handshake_failure. This message is always fatal. |
| internal_error | An internal error unrelated to the peer or the correctness of the protocol makes it impossible to continue, such as a memory allocation failure. The error is not related to protocol. This message is always fatal. |
| user_canceled | Cancelled handshake for a reason that is unrelated to a protocol failure. If the user cancels an operation after the handshake is complete, just closing the connection by sending a close_notify is more appropriate. This alert should be followed by a close_notify. This message is generally a warning. |

| | |
|---|---|
| no_renegotiation | Sent by the client in response to a hello request or sent by the server in response to a client hello after initial handshaking. Either of these would normally lead to renegotiation; when that is not appropriate, the recipient should respond with this alert; at that point, the original requester can decide whether to proceed with the connection. One case where this would be appropriate would be where a server has spawned a process to satisfy a request; the process might receive security parameters (key length, authentication, and so on) at startup and it might be difficult to communicate changes to these parameters after that point. This message is always a warning. |

**The Record Layer**

As specified by RFC 2246, the Record Layer might have four functions:

- It fragments the data coming from the application into manageable blocks (and reassemble incoming data to pass up to the application). Schannel SSP does not support fragmentation at the Record Layer.
- It compresses the data and decompresses incoming data. Schannel SSP does not support compression at the Record Layer.
- It applies a Message Authentication Code (MAC), or hash/digest, to the data and uses the MAC to verify incoming data.
- It encrypts the hashed data and decrypts incoming data.

**Record Protocol Functions**

The Record Protocol receives and encrypts data from the application layer and delivers it to the Transport Layer. Then it takes the data, fragments it to a size appropriate to the cryptographic algorithm, optionally compresses it (or, for data received, decompresses it), applies a MAC or HMAC and then encrypts (or decrypts) the data using the information negotiated during the Handshake Protocol. HMAC is supported only by TLS.

The resulting encrypted data is then passed to the transport layer.

**Record Layer Encryption States**

The record layer changes encryption states during the session setup. That is, different encryption methods are used at different stages of the process.

- **No Encryption Used**. Prior to the negotiating a cipher suite and deciding which security methods to use, the state is Null. No message authentication or encryption is performed.
- **Public Key Encryption Used** (Uses public-private key pair). Once a cipher suite is negotiated and certificates are exchanged, The record layer hashes outgoing data with the appropriate MAC (typically both MD5 and secure hash algorithm (SHA-1), as required for RSA signing) and encrypts it with the sender's private key. Incoming data is decrypted with the sender's public key.
- **Symmetric Key Encryption Used** (Uses shared session key rather than public-private key pair). Near the end of the Handshake sequence, the handshake protocol exchanges data that allows both the client and the server to compute a Master Secret and then to derive the Client Write MAC Secret, Server Write MAC Secret, Client Write Key, and Server Write Key from it.

  When these keys are available, the client and server send a message to Change Cipher Spec. At that point, the client and server stop using the public-private key pair and start using the shared session keys derived from the Master Secret.

**TLS/SSL Processes and Interactions**

This section provides a detailed explanation of the TLS/SSL protocol, specifically the handshake protocol, its associated messages and alerts, and the record protocol.

The first three operations listed below take place using the Handshake Protocol. The parties negotiate which cipher suite they will use, authenticate, exchange keys, and exchange application data.

**Handshake and Cipher Suite Negotiation**

First, the client and server contact each other and choose a common cipher suite. The suite will include:

- A key exchange method, which determines how the shared master key will be exchanged.

- A bulk encryption method which determines how application data will be encrypted.
- A MAC, which determines how application data will be hashed and signed to prove integrity.

**Authentication**
The server always authenticates its identity to the client. However, the client might not need to authenticate with the server, depending on the application. The authentication method (primarily which digital certificate format will be used) depends on the negotiated cipher suite.

**Key Exchange**
After choosing a cipher suite, the client and server exchange a key (or the information needed to create a key) that they will use for data encryption, which again depends on the negotiated cipher suite's requirements. The key exchange operation requires the following things:

- The random values are created called the Client Random and the Server Random.
- The client generates a random Pre-Master Secret.

The client must transmit the Pre-Master Secret securely to the server.

- With RSA key exchange, the Pre-Master Secret is encrypted with the server's public key. The client gets this key from the server's certificate.
- With EDH key exchange, the pre-master secret is the result of the EDH operation. In this case the ephemeral DH key is signed by the server's private key.

Both client and server use the Pre-Master Secret to create a shared Master Secret. This is done by hashing the pre-master secret together with the ClientRandom and ServerRandom values. The Master Secret is used to create four keys that are shared by the client and server:

- **Client Write MAC Secret.** This key is added to client message hashes. The client uses the key to create the initial hash. The server uses it to authenticate client messages.
- **Server Write MAC Secret.** This key is added to server message hashes. The server uses the key to create the initial hash. The client uses it to authenticate server messages.
- **Client Write Key**. The client uses this key to encrypt messages. The server uses the key to decrypt client messages.
- **Server Write Key.** The server uses this key to encrypt messages. The client uses the key to decrypt server messages**.**

The last operation takes place at the Record Layer using the Record Protocol.

**Application Data Exchange**
The client application and the server application communicate with each other. All data is encrypted using the negotiated bulk encryption method.

**TLS/SSL Message Sequence Overview**
The process of authenticating and establishing an encrypted channel with a server using the Schannel authentication protocol involves the following steps:

### Full TLS Handshake




The Handshake protocols manage this process of exchanging key generation material, and it uses the Handshake Protocol.

Once hashing and encryption keys are ready for use, the Record Layer takes over. It uses the Record Protocol to secure application data, using the keys created in the Handshake process.

**The Full Handshake Protocol**

The handshake protocol is a series of sequenced messages that negotiate the security parameters of a data transfer session.

**Client Hello Messages**

The Client Hello is typically the first message in the TLS/SSL session setup sequence. The protocol allows the server to request a hello, but this is application-specific and not part of the normal sequence.

### Client Hello Messages




**Client Hello Message**

The client initiates a session by sending a Client Hello message to the server. The Client Hello message contains:

- **Version Number**. The version number of the highest version that the client supports. This is sent by the client to the server. Version 2 is used for SSL 2.0, version 3 for SSL 3.0, and version 3.1 for TLS. Although the IETF RFC for TLS is TLS version 1.0, the protocol uses 3.1 in the version field to indicate that it is a later version, with more functionality than SSL 3.0.

- **Client Random**. A 4-byte number that consists of the client's date and time, plus a 28-byte cryptographically-generated pseudorandom number. This is used in the calculation of the Master Secret from which the encryption keys are derived.

- **(Optional) Session Identification**. A byte string used to identify an active or resumable session state. Enables the client to resume a previous session. Resuming a previous session can be useful, because creating a new session requires processor-intensive public key operations that can be avoided by resuming an existing session with its established session keys. Previous session information, identified by the SessionID, is stored in the respective client and server session caches. This is empty if the client is not resuming a session.

- **Cipher Suite**. The list of cipher suites available on the client. An example of a cipher suite is TLS_RSA_WITH_DES_CBC_SHA, where TLS is the protocol version, RSA is the algorithm that will be used for the key exchange, DES_CBC is the encryption algorithm (using a 56-bit key in CBC mode), and SHA-1 is the hash function.

  **Note**

  - Microsoft applications typically specify RSA as the key exchange algorithm, RC4 as the encryption method, and MD5 as the Message Authentication Codes.

- **Compression Algorithm**. The client's supported compression algorithms. Compression is optional and is one of the

four operations that the Record Protocol might perform on all data after the Hello sequence is finished.

**Server Session Cache Allows Session Resume**

The Client Hello is used to either start a new session or to resume an existing session. If the Client Hello includes a Session ID, the client is attempting to resume a session.

The server maintains a session cache to allow fast resumption of recent sessions, similar to a ticket cache in Kerberos. Resuming a session does not require certificate exchanges, so it is much quicker than the normal Hello sequence.

The application determines whether client's are allowed to resume sessions and how long idle session IDs are valid.

The server uses CryptoAPI to manage both the session ID and the certificate cache.

### Note

- The Client Hello can be initiated at any time during an existing session and is not limited to just session initialization. The server application might request a new Hello periodically to request client authentication based on the resource that is requested. Either the client application or the server might request a new handshake to refresh encryption keys.

**TLS 1.0 Cipher Suites**

Schannel supports the cipher suites in the following table for TLS 1.0. In this table, the following acronyms are used:

CBC = cipher block chaining

DES = Data Encryption Standard

DHE = Ephemeral Diffie-Hellman

DSS = Digital Signature Standard

The suites are listed in order of preference.

**TLS 1.0 Cipher Suites**

| Key Exchange | Cipher | Hash |
|---|---|---|
| RSA | RC4 128 | MD5 |
| RSA | RC4 128 | SHA-1 |
| RSA | 3DES EDE CBC | SHA-1 |
| DHE DSS | 3DES EDE CBC | SHA-1 |
| RSA | DES CBC | SHA-1 |
| DHE DSS | DES CBC | SHA-1 |
| RSA Export 1024 | RC4 56 | SHA-1 |
| RSA Export 1024 | DES CBC | SHA-1 |
| DHE DSS Export 1024 | DES CBC | SHA-1 |
| RSA Export | RC4 40 | MD5 |
| RSA Export | RC2 CBC 40 | MD5 |
| RSA | None | MD5 |
| RSA | None | SHA-1 |

**SSL 3.0 Cipher Suites**

Schannel supports the cipher suites listed under TLS 1.0 Cipher Suites.

#### SSL 2.0 Cipher Suites

Schannel supports the following cipher suites for SSL 2.0. The suites are listed in order from most secure to least secure.

### SSL 2.0 Cipher Suites

| Cipher | Hash |
|---|---|
| RC4 128 | MD5 |
| DES 192 EDE3 CBC | MD5 |
| RC2 CBC 128 CBC | MD5 |
| DES 64 CBC | MD5 |
| RC4 128 Export 40 | MD5 |

#### Server Responses to Client Hello

### Server Responses to Client Hello




#### Server Hello Message

The server responds with a Server Hello message. The Server Hello message includes:

- **Version Number**. The server sends the highest version number that is supported by both sides. This is the protocol version that will be used during the connection.
- **Server Random[32]**. ServerRandom[32] is a 4-byte representation of the server's date and time plus a 28-byte, cryptographically-generated, pseudorandom number. This number, along with the Client Random, is used by both the client and the server to generate the Master Secret from which the encryption keys will be derived.
- **Session Identification** (if any). This can be one of three choices.
    - **New session ID.** The client did not indicate a session to resume, so a new ID is generated. A new session ID is also generated when the client indicates a session to resume, but the server can't or won't resume that session.
    - **Resumed Session ID**. The ID is the same as indicated in the Client Hello. The client indicated a specific session ID to resume and the server is willing to resume that session.
    - **Null**. This is a new session, but the server is not willing to resume it at a later time, so no ID is returned.
- **Cipher Suite**. The server chooses the strongest cipher that both the client and server support. If there are no cipher suites that both parties support, the session is ended with a *handshake failure* alert.
- **Compression Algorithm**. If used, specifies the compression algorithm to use.

#### Server Certificate Message

The server sends its certificate to the client. The server certificate contains the server's public key. The client uses this key to authenticate the server and to encrypt the Premaster Secret. The Server Certificate message includes:

- **The server's certificate list**. The first certificate in the list is the server's X.509v3 certificate that contains the server's public key.
- **Other validating certificates**. All other validating certificates, up to but not including the root certificate from the CA, signed by the CA.

The client also checks the name of the server in the certificate to verify that it matches the name the client used to connect. For example, if the user types http://www.contoso.com as the URL in the browser, the certificate contains a subject name of www.contoso.com or *.contoso.com. Internet Explorer warns the user if these names do not match; other client applications typically fail the connection directly.

**(Optional) Server Key Exchange Message**
The server creates and sends a temporary key to the client. This key can be used by the client to encrypt the Client Key Exchange message later in the process. The step is only required when the server's certificate does not contain a public key that is suitable for key exchange or when the cipher suite mandates the use of an ephemeral key for the key exchange operation. This message is also used when the server uses a DSS certificate. In this case, the server public key is not suitable for key exchange, and so an ephemeral DH key is created by the server and sent in the Server Key Exchange message.

In these situations, the server must create and send a temporary key that is used instead of the server's public key. The key that is created depends on the cipher suite. With export versions of RSA that do not allow a public key greater than 512 bits, the temporary shorter key is signed with the unusable public key for authenticity.

Why is this important? The client will need this key to encrypt the Premaster Secret, whichis discussed below, in the Client Key Exchange message.

### Note

- This message is not used in non-export versions of Microsoft applications, since non-export RSA certificates will always include the server's public key in its certificate.

**(Optional) Client Certificate Request Message**
The server must always present its certificates to the client, but the client is not always required to authenticate itself. Therefore, the client is not always required to send its certificates to the server. If the server does not require client authentication, then this message is not sent.

This step might be used for Web sites such as a banking Web site, where the server must confirm the identity of the client before providing sensitive information. If the application requires mutual authentication, the server sends a Client Certificate Request. The Client Certificate Request message includes:

- The type of certificate required (typically RSA or DSS)
- A list of acceptable CAs

**Server Hello Done Message**
This message indicates that the server is finished and awaiting a response from the client. This message has no content. It signals that the Server Hello sequence is finished.

**Client Responses to Server Hello**
The client receives the server's message and checks the certification hierarchy of the server's certificate.

### Client Responses to Server Hello




**Client Certificate Message (Required if Requested)**
If the server sent a Client Certificate Request, the client sends its certificate to the server for client authentication. The client's certificate contains the client's public key. The Client Certificate message includes theclient's **certificate list**. The first certificate in the list is the client's X.509v3 certificate that contains the client's public key. After that are other validating certificates, up to but not including the root certificate from the CA, signed by the CA.

#### Client Key Exchange Message

The client sends a Client Key Exchange message after computing the premaster secret using the two random values that are generated during the Client Hello message and the Server Hello message. Before it is transmitted to the server, the premaster secret is encrypted by the public key from the server's certificate. Both parties compute the master secret locally and derive the session key from it.

If the server can decrypt this data and complete the protocol, the client is assured that the server has the correct private key. This step is crucial to prove the authenticity of the server. Only the server with the private key that matches the public key in the certificate can decrypt this data and continue the protocol negotiation.

The Client Key Exchange message includes:

- **Client's protocol version**. The server will verify that it matches the original value sent in the Client Hello message. This measure guards against rollback attacks. Rollback attacks work by manipulating the message in order to cause the server and the client to use a less secure, earlier version of the protocol.

- **Pre-master secret**. This is the client-generated number (48-byte for RSA), encrypted with the server's public key, that is used with the Client Random and the Server Random to create the Master Secret.

#### Certificate Verify Message (Required if Client Certificate Is Sent)

This message is sent only if the client previously sent a Client Certificate message. The client is authenticated by using its private key to sign a hash of all the messages up to this point. The server verifies the signature with the public key of the signer, which ensures that it was signed with the client's private key. This message contains a long signature to verify the client's certificate, if one was requested and sent.

For RSA, the signature consists of:

- An MD5 hash of all previous handshake messages.

- An SHA-1 hash of all previous handshake messages.

- Both hashes, which are concatenated and encrypted with the client's private key.

For DSS, the signature consists of:

- An SHA-1 hash of all previous handshake messages.

- Encrypted with the client's private key.

#### Change Cipher Spec Message

The Change Cipher Spec message notifies the server that all future messages including the Client Finished message are encrypted using the keys and algorithms just negotiated. At this point, client and server are authenticated and the client has sent the pre-master secret. Both the client and the server have calculated the Master Secret. Up until now, however, any encryption has used the client's or server's private/public keys. The Change Cipher Spec message tells the server that the client is ready to use the Write Key for all further encryption.

### Note

- The Change Cipher Spec message is not part of the Handshake Protocol. The TLS RFC defines a separate protocol for Change Cipher Spec. This allows developers to define a specific Change Cipher Spec message.

#### Computing the Master Secret Key and Subsequent Keys

The Handshake sequence securely exchanges data that is used to create the Master Secret. This secret, however, is never used by itself in encryption; rather, several keys are derived from the Master Secret. These keys are then used for message authentication and encryption.

### Note

- The details in this section and the next (hashing in the Record Layer) provide a glimpse inside the computational "black box" and are meant to illustrate the complexity that is necessary to secure communications. The key material exchange

specifics change over time as new algorithms are developed.

**Computing the Master Secret**
The Master Secret results from the pseudorandom material that is exchanged in the Client Hello message and Server Hello message that is discussed earlier in the sections on the Client and Server Final messages. Both the server and the client create the Master Secret key. It is never exchanged.

To create the Master Secret key, the system passes the following as variables to a pseudorandom function (PRF):

- The 48-byte Pre-Master Secret.
- The literal phrase "master secret"
- The concatenation of Client Random number and Server Random number.

The result is the Master Secret.

**MAC Secret and Write Encryption Keys**
To enable message encryption and hashing, four keys need to be created on both the server and the client:

- Client Write MAC Secret
- Server Write MAC Secret
- Client Write Key
- Server Write Key

These keys are never exchanged.

To produce these keys, the Master Secret is passed to the PRF until enough output has been generated to yield a value that is divided into six parts. (The last two parts are only generated for non-export block ciphers.)

**Client Write MAC Secret**
This key is added to client message hashes. The client uses the key to create the initial hash. The server uses it to authenticate client messages.

**Server Write MAC Secret**
This key is added to server message hashes. The server uses the key to create the initial hash. The client uses it to authenticate server messages.

**Client Write Key**
The client uses this key to encrypt messages. The server uses the key to decrypt client messages.

**Server Write Key**
The server uses this key to encrypt messages. The client uses the key to decrypt server messages.

**Finished Message**
The Finished message is a hash of the entire conversation which provides further authentication of the client. This message is a complicated verification that the client who is sending the Finished message is truly the client who started the Hello conversation. It is the first message that the Record Layer encrypts and hashes with the Write Key and the Write MAC Secret.

**Construction of the message at the Handshake Protocol**
To generate the message contents, the system passes the following as variables to a pseudorandom function (PRF):

- The Master Secret.
- The literal phrase "client finished".
- The concatenation of an MD5 hash of all previous handshake messages and an SHA-1 hash of all previous handshake messages.

The Handshake protocol passes the result to the Record Layer.

**Completing the message at the Record Layer**
The Record Layer hashes the data using HMAC with the Client Write MAC Secret, which is derived from the Master Secret. Then, the Record Layer encrypts the data with the Client Write Key, which is also derived from the Master Secret.

**Server Finished Messages**
If the server has a copy of the private half of its public key pair, it can decrypt the ClientKeyExchange message and generate a matching set of encryption and MAC keys. The server responds to the client with a request to communicate with secret key cryptography, so that both parties have acknowledged that they will use the cryptography and keys that they have agreed upon.

### Server Finished Messages




**Change Cipher Spec Message**
This message notifies the client that the server will begin encrypting messages with the keys that were just negotiated. The server switches its record layer security state to symmetric encryption using the session keys.

**Finished Message**
This message is a hash of the entire exchange to this point using the session key and the MAC Secret. If the client can successfully decrypt this message and validate the contained hashes, the client is assured that the TLS/SSL handshake was successful, and the keys computed on the client machine match those computed on the server.

This message is the same as the Client Finished message with these differences:

- The literal phrase “server finished” is used instead of “client finished”.
- The hashed handshake messages include the client’s Client Finished message.
- The server uses the Server Write MAC Secret and Server Write Key for hashing and encrypting at the Record Layer.

**Application Data Flow**
So far in this process, client and server applications have been authenticated through the use of certificates and public/private keys, a Premaster Secret and other data have been exchanged to create the Master Secret, and both the server and the client have successfully proven that they have not changed identities throughout the Hello sequence. Now the applications can begin to communicate, using the established keys and parameters. The Record Layer fragments, compresses, hashes, and encrypts all further communications. When necessary, it decrypts, verifies, decompresses, and reassembles the communication

### Record Protocol




**The Record Protocol**
When the record protocol receives the data from the application layer, it might perform the following tasks:

- Fragments the data into blocks or reassembles fragmented data into its original structure. Schannel does not support fragmentation at the Record Layer.
- Numbers the sequence of data blocks in the message to protect against attacks that attempt to reorder data.
- Compresses or decompresses the data using the compression algorithm negotiated in the handshake protocol. Schannel does not support compression at the Record Layer.

- Encrypts or decrypts the data using the encryption keys and cryptographic algorithm negotiated during the handshake protocol.

- Applies an HMAC (or, for SSL 3.0, a MAC) to outgoing data. It then computes the HMAC and verifies that it is identical to the value that was transmitted in order to check data integrity when a message is received.

Once the record protocol has completed its operations on the data, it sends the data to the application for transmission. If the data is incoming, it is sent to the appropriate process for reception.

**HMAC Hashing in the Record Layer**

Fragmented, compressed application data is hashed as it passes through the Record Layer before being encrypted.

**Note**

- The details in this section provide a glimpse inside the computational "black box" and are meant to illustrate the complexity necessary to secure communications. The key material exchange specifics change over time as new algorithms are developed.

The MAC consists of two values that are hashed:

- The secret key value: Client or Server Write MAC Secret (derived from the Master Secret).

- A concatenation of:

    - The record's sequence number.

    - Packet type.

    - Protocol version.

    - Packet length.

    - The non-encrypted application data.

The result of this hashing and the original compressed data fragment are then encrypted and passed down to TCP.

When the message is decrypted at its destination, a new hash is computed, based on the compressed fragment and the MAC Secret. The new hash is compared to the hash that was sent in the message. If they match, the message's data integrity is verified.

**Note**

- Data can be encrypted and decrypted, but you cannot reverse engineer a hash. Hashing is a one-way process. Running the process backward does not recreate the original data. This is why a new hash is computed and then compared to the sent hash.

**Resuming a Secure Session**

Once a secure session has been established between a client and a server, the client can attempt to resume the session in the future, using the same keys used in the previous session. If the server is able to resume the session, then the abridged version of the Handshake Protocol below will occur. If not, then a full handshake will occur.

**Resume Session Messages**




- The client sends a Client Hello message using the Session ID of the session to be resumed.
- The server checks its session cache for a matching Session ID. If a match is found and the server is able to resume the session, it sends a Server Hello message with the Session ID.

  **Note**

  - If a session ID match is not found, the server generates a new session ID and the TLS client and server perform a full handshake.

- The server must send a Change Cipher Spec message to notify the client that the server will begin encrypting messages with a new set of encryption and MAC keys, based on the cached Master Secret and the Client Random Server Random values from the current connection. This ensures that every connection has its own set of encryption keys. The server switches its record layer security state to symmetric encryption using the session keys.
- The server sends a Server Finished message.
- The client must send a Change Cipher Spec message to notify the server that the client will start using the session keys for hashing and encrypting messages.
- The client sends a Client Finished message.
- The client and server can now resume application data exchange over the secure channel.

**Renegotiation**
**Renegotiation Methods**
Renegotiation can originate from either the client or the server with the following exceptions:

- No client-side renegotiation in Windows 2000 and Windows XP.
- No renegotiation in SSL 2.0.

**Server Renegotiation Scenarios**
The server can request renegotiation in two situations:

- After the initial handshake, the client requests an access-protected resource. The server can initiate a renegotiation to ask for client authentication. This is the case where the server sends the Hello Request. When the client receives a Hello Request, it begins a new handshake by sending a Client Hello message.
- To refresh keys. This is performed in the same way as the second case, above.

**Note**

- Schannel does not allow the client to ignore the Hello request. The client must initiate a new handshake by sending a

Client Hello message, or the server will close the connection with a fatal error.

**Typical Server Renegotiation**

**Server Initiated Renegotiation**




Server-initiated renegotiation uses the following procedure:

- The client and server successfully complete a full TLS/SSL handshake.
- The client requests for an access-protected resource.
- The server sends a Hello Request message to the client.
- The client receives the Hello Request message.

  **Note**

  - Schannel does not allow the client to ignore the Hello request, but other client's might. The client must initiate a new handshake by sending a Client Hello message or the Windows server closes the connection.

- The client sends a Client Hello message to the server to initiate a new full handshake.

**Client Renegotiation Scenarios**

The client can also initiate a renegotiation anytime by sending a Client Hello message. This is usually performed to refresh encryption keys.

**Typical Client Renegotiation**

**Client Initiated Renegotiation**




Client-initiated renegotiation uses the following procedure:

- The client and the server successfully complete a full SSL handshake.
- The client needs to refresh encryption keys. The client sends a Client Hello message to the server.
- The server receives the Client Hello message. The server has no way of knowing that this message is not application data, and so it passes the message on. It is returned as an "extra" buffer containing the unprocessed Client Hello message with a flag signaling that it is a renegotiation.

  **Note**

  - The server might not respond to the renegotiation request right away. In the most general case, the server might

have already sent some application data at the same time that the client is requesting the renegotiation, or it might decide to send some data before responding to the renegotiation request. However, the client will never send more application data after it has initiated a renegotiation.

- The server sends the Server Hello group of messages to the client.
- The client receives the Server Hello messages. The client has no way of knowing that this message is not application data, and so it passes the message on. It is returned as an "extra" buffer containing the unprocessed server hello messages with a flag signaling that it is a renegotiation
- The client sends the next set of handshake messages, which are sent to the server, and the handshake proceeds as usual.

**The Schannel Session Cache**
The first time that a client connects to a server, a full handshake is performed. Once this is complete, the Master Secret, Cipher Suite, and certificates are stored in the session cache on the respective client and server machines. TLS/SSL supports a reconnect operation that can be used to enable a client and server to resume a previously-negotiated session. This allows subsequent connections between the same client and server to use a reconnect handshake rather than a full handshake. This is desirable in many cases because the TLS/SSL reconnect requires far less CPU and network bandwidth because there is no need for a key exchange or certificate exchange. But there are cases in which reconnects can be troublesome; such as during performance testing or when the client never reconnects to the same Web server.

Schannel supports reconnects by use of the Schannel session cache because it keeps a list of previous TLS/SSL sessions that are established with the current authenticated logon data that is used by a security principal to establish its identity. The sessions are referenced by the TLS/SSL session ID. Schannel currently maintains up to ten thousand cached sessions for a maximum of ten hours.

**How Schannel Uses Certificate Mapping**
When a server application requires client authentication, Schannel automatically attempts to map the certificate that is supplied by the client to a user account. You can authenticate users who log on with a client certificate by creating mappings, which relate the certificate information to a Windows user account. After you create and enable a certificate mapping, each time a client presents a client certificate; your server application automatically associates that user with the appropriate Windows user account.

**Types of Mapping**
In most cases, a certificate is mapped to a user account in one of two ways: a single certificate is mapped to a single user account (one-to-one mapping), or multiple certificates are mapped to one user account (many-to-one mapping). By default schannel will use the following four certificate mapping methods, in order of preference.

**S4U Certificate mapping**
This is new for Windows Server 2003. Schannel uses the UPN from the client certificate and attempts an S4U logon using the Kerberos security package. The CA that issued the client certificate must have NTAuth trust. The domain of the server and the user must be running in Windows Server 2003 function level. This is the only mapping method that works even when the server is in a different forest than the client user.

**User Principal Name mapping**
Enterprise CAs place an entry, called a user principal name (UPN), in each certificate. The UPN looks very much like an e-mail name. The UPN is unique within an Active Directory forest. The UPN finds the user's account in Active Directory, and that account is logged on. UPN mappings are implicit and therefore do not require administrators to create and maintain.

Schannel uses the UPN from the client certificate, and queries Active Directory for a matching user account. The CA that issued the client certificate must have NTAuth trust.

**One-to-One Mapping also known as subject/issuer mapping**
A method that maps a single user certificate to a single Active Directory user account. Unlike UPN mapping, the administrator must explicitly map a certificate to a user's account. This certificate can come from any source, as long as the root CA for that certificate is trusted for client authentication. The client certificate can be renewed, etc without affecting this process as long as the subject and issuer name fields do not change.

For example, suppose that users obtain certificates from a CA that is approved by your company. You then take these user certificates and map them to the Active Directory user accounts. This allows users to connect to the server, using TLS from anywhere by providing their client certificate.

Schannel builds an "altSecurityId" string from the client's certificate's subject and issuer name fields, and queries Active

Directory for a user account with a matching **altSecurityId** property.

**Many-to-One Mapping**
Many-to-one mapping involves mapping many certificates to a single user account. Many-to-one mapping works the same way as one-to-one mapping, but it maps all client certificates issued by a given CA to the same user account. This certificate can come from any source, as long as the root CA for that certificate is trusted for client authentication.

Schannel builds an "altSecurityId" string from the client's certificate's issuer name field, and queries Active Directory for a user account with a matching **altSecurityId** property.

**Network Ports Used by TLS/SSL**
**Port Assignments for Common Applications over TLS/SSL**

| Service Name | TCP |
|---|---|
| smtp | 25 |
| https | 443 |
| nntps | 563 |
| ldaps | 636 |
| ftps-data | 989 |
| ftps | 990 |
| telnets | 992 |
| imaps | 993 |
| pop3s | 995 |
| ms-sql-s | 1433 |
| mfst-gc-ssl | 3269 |
| tftps | 3713 |

**Related Information**

- RFC 2246 in the IETF RFC Database
- RFC 3546 in the IETF RFC Database
- RFC 2817 in the IETF RFC Database

# TLS/SSL Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TLS/SSL Tools and Settings**

### In this section


- TLS/SSL Tools
- TLS/SSL Registry Entries
- TLS/SSL Group Policy Settings
- Network Ports Used by Schannel


**TLS/SSL Tools**

The following tools are associated with TLS/SSL.

**Dsa.msc: Active Directory Users and Computers Category**

Active Directory Users and Computers is a Microsoft Management Console (MMC) that is automatically installed when you install Active Directory. This tool also ships with the Administration Tools Pack (Adminpak.msi).

You can access the tool from the **Start** menu: To do this, click **Start**, point to **Programs**,point to **Administrative Tools,** and then click **Active Directory Users and Computers**.

**Version compatibility**

Active Directory Users and Computers runs on domain controllers that run Windows Server 2003 or Windows 2000 operating systems. You can use MMC to administer and publish information in the directory.

The Windows Server 2003 version of Active Directory Users and Computers can target domain controllers that are running Windows Server 2003 or Windows 2000 Server.

On administrative workstations that are running Windows XP Professional or Windows 2000, you can install the Windows Server 2003 Administration Tools Pack (Adminpak.msi) from the i386 folder on the Windows Server 2003 CD. This version of the Administration Tools Pack encrypts and signs Lightweight Directory Access Protocol (LDAP) traffic between the administrative tool client's and domain controllers.

### Note

- You cannot run the Windows Server 2003 Administration Tools Pack (Adminpak.msi) on a computer that is running Windows XP Professional, Windows XP Home Edition, or Windows XP 64-Bit Edition Version 2003 without Windows XP Service Pack 1 (SP1).

You can use Active Directory Users and Computers to manage the following properties that are listed in the following table, which are associated with objects in Active Directory. Any changes you make affect certificate mapping for these objects.

**Active Directory Users and Computers Object Management**

| Property | Changes That Affect Schannel |
|---|---|
| **Computer objects** | |
| Name Mapping Task | Can add, edit or remove certificates. |

| User objects | |
|---|---|
| Name Mapping Task | Can add, edit or remove certificates. |
| Published Certificates Tab | Lists the X.509 certificates published for the user account. Can view, remove, and copy to file listed certificates. Can add new certificates from the local certificate store or from a DER Encoded Binary X509 (*.cer) or PKCS #7 (*.p7b) file. |

You can find more information about Active Directory Users and Computers on the TechNet Web site.

**Eventvwr.msc: Event Viewer**
**Category**
Event Viewer is included in the Windows Server 2003, Windows XP, and Windows 2000 operating systems.

**Version compatibility**
Event Viewer is supported for the Windows Server 2003, Windows XP, and Windows 2000 operating systems.

The system log contains Secure Channel (Schannel) events that are related to authentication.

### Schannel Events

| Event ID | Severity | Description |
|---|---|---|
| 36864 | Informational | The Schannel security package has loaded successfully. |
| 36865 | Error | A fatal error occurred while opening the system *cryptographic subsystem* cryptographic module. Operations that require the SSL or TLS cryptographic protocols will not work correctly. The error code is *error code*. |
| 36866 | Error | The Schannel security package has failed to load. Operations that require the SSL or TLS cryptographic protocols will not work correctly. |
| 36867 | Informational | Creating an SSL [client\| server] credential. |
| 36868 | Informational | The SSL [client\| server] credential's private key has the following properties:<br>• CSP name<br>• CSP type<br>• Key name<br>• Key Type<br>• Key Flags<br><br>The attached data contains the certificate. |
| 36869 | Error | The SSL [client\| server] credential's certificate does not have a private key information property attached to it. This most often occurs when a certificate is backed up incorrectly and then later restored. This message can also indicate a certificate enrollment failure. |
| 36870 | Error | A fatal error occurred when attempting to access the SSL [client\| server] credential private key.<br><br>The error code returned from the cryptographic module is *error code*. |
| 36871 | Error | A fatal error occurred while creating an SSL [client\| server] credential. |

| | | |
|---|---|---|
| 36872 | Warning | No suitable default server credential exists on this system. This will prevent server applications that expect to make use of the system default credentials from accepting SSL connections. An example of such an application is the directory server. Applications that manage their own credentials, such as Microsoft Internet Information Services (IIS), are not affected by this. |
| 36873 | Error | No supported cipher suites were found when initiating an SSL connection. This indicates a configuration problem with the client application or the installed cryptographic modules. The SSL connection request has failed. |
| 36874 | Error | An SSL connection request was received from a remote client application, but none of the cipher suites supported by the client application are supported by the server. The SSL connection request has failed. |
| 36875 | Warning | The remote server has requested SSL client authentication, but no suitable client certificate could be found. An anonymous connection will be attempted. This SSL connection request might succeed or fail, depending on the server's policy settings. |
| 36876 | Error | The certificate received from the remote server has not validated correctly. The error code is *error code*. The SSL connection request has failed. The attached data contains the server certificate. |
| 36877 | Warning | The certificate received from the remote client application has not validated correctly. The error code is *error code*. The attached data contains the client certificate. |
| 36878 | Warning | The certificate received from the remote client application is not suitable for direct mapping to a client system account, possibly because the authority that issuing the certificate is not sufficiently trusted. The error code is *error code*. The attached data contains the client certificate. |
| 36879 | Warning | The certificate received from the remote client application was not successfully mapped to a client system account. The error code is *error code*. This is not necessarily a fatal error, as the server application might still find the certificate acceptable. |
| 36880 | Informational | An SSL [client\| server] handshake completed successfully. The negotiated cryptographic parameters are as follows.<br>• Protocol<br>• Cipher<br>• Cipher strength<br>• MAC<br>• Exchange<br>• Exchange strength |
| 36881 | Error | The certificate received from the remote server has expired. The SSL connection request has failed. The attached data contains the server certificate. |
| 36882 | Error | The certificate received from the remote server was issued by an untrusted certificate authority. Because of this, none of the data contained in the certificate can be validated. The SSL connection request has failed. The attached data contains the server certificate. |
| 36883 | Error | The certificate received from the remote server has been revoked. This means that the certificate authority that issued the certificate has invalidated it. The SSL connection request has failed. The attached data contains the server certificate. |
| 36884 | Error | The certificate received from the remote server does not contain the expected name. It is therefore not possible to determine whether we are connecting to the correct server. The server name we were expecting is *servername*. The SSL connection request has failed. The attached data contains the server certificate. |

| | | |
|---|---|---|
| 36885 | Warning | When asking for client authentication, this server sends a list of trusted certificate authorities to the client. The client uses this list to choose a client certificate that is trusted by the server. Currently, this server trusts so many certificate authorities that the list has grown too long. This list has thus been truncated. The administrator of this machine should review the certificate authorities trusted for client authentication and remove those that do not really need to be trusted. |

To find more information about "Event Viewer", see "Event Viewer" on TechNet.

**Netmon.exe: Network Monitor**
**Category**
A limited version of Network Monitor is included in Windows Server 2003, Windows XP, and Windows 2000 operating systems. The full version of Network Monitor is included with Microsoft Systems Management Server.

**Version compatibility**
Network Monitor is supported for Windows Server 2003, Windows XP, and Windows 2000.

Network Monitor enables you to capture network traces which can be used in troubleshooting most network issues.

**TLS/SSL Registry Entries**
The following registry subkeys and entries can help you administer and troubleshoot TLS/SSL, but they apply more to Schannel SSP than TLS/SSL. They can help you verify that the required settings are applied.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, instead of editing the registry directly, use Group Policy or other Windows tools such as MMC to accomplish tasks. If you must edit the registry, use extreme caution.

**CertificateMappingMethods**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**
Windows Server 2003

This entry does not exist in the registry by default. The default value is that all four certificate mapping methods are supported.

**Ciphers**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This subkey controls use of symmetric algorithms.

**DES 56/56**
This subkey controls use of DES 56 bit algorithm.

**NULL**
This subkey controls use of no encryption.

**RC2 128/128**
This subkey controls use of RC2 128 bit algorithm.

**RC2 40/128**
This subkey controls use of RC2 40 bit algorithm.

**RC2 56/128**
This subkey controls use of RC2 56 bit algorithm.

**RC4 128/128**
This subkey controls use of RC4 128 bit algorithm.

**RC4 40/128**
This subkey controls use of RC4 40 bit algorithm.

**RC4 56/128**
This subkey controls use of RC4 56 bit algorithm.

**Triple DES 168/168**
This subkey controls use of 3DES 168 bit algorithm.

The default for these ciphers is enabled.

To disable, create the **Enabled** entry in the appropriate subkey. (This entry does not exist in the registry by default.) After you have created the entry, change the DWORD value to 0. When you disable any algorithm, you disallow all cipher suites that use that algorithm. To re-enable the cipher, change the DWORD value to 0xffffffff.

**ClientCacheTime**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This entry controls the time to expire client side cache entries in milliseconds. A value of 0 turns off secure connection caching. This entry does not exist in the registry by default. The default values are:

### Default Client Cache Time

| Windows Version | Time |
|---|---|
| Windows NT 4.0 with Service Pack 6a | 2 minutes |
| Windows NT 4.0 with Service Pack 6a and Q265369 | 60 minutes |
| Windows 2000 | 2 minutes |
| Windows 2000 with Service Pack 2 or greater | 10 hours |
| Windows XP | 10 hours |
| Windows Server 2003 | 10 hours |

**Fipsalgorithmpolicy**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\LSA**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This entry controls FIPS compliance. The default is 0.

### FIPS Cipher Suites

| Protocol Version | Key Exchange | Cipher | Hash |
|---|---|---|---|
| SSL 3.0 | RSA | DES CBC | SHA-1 |
| SSL 3.0 | RSA | 3DES EDE CBC | SHA-1 |
| SSL 3.0 | RSA | Export 1024 DES CBC | SHA-1 |
| TLS 1.0 | RSA | DES CBC | SHA-1 |
| TLS 1.0 | RSA | 3DES EDE CBC | SHA-1 |
| TLS 1.0 | RSA | Export 1024 DES CBC | SHA-1 |

**Hashes**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This subkey controls use of hash algorithms.

**MD5**
This subkey controls use of MD5 as hashing algorithm.

**SHA**
This subkey controls use of SHA-1 as hashing algorithm.

The default for these ciphers is enabled.

To disable, create the **Enabled** entry in the appropriate subkey. (This entry does not exist in the registry by default.) After you have created the entry, change the DWORD value to 0. When you disable any algorithm, you disallow all cipher suites that use that algorithm. To re-enable the cipher, change the DWORD value to 0xffffffff.

**IssuerCacheSize**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**
Windows Server 2003

This entry controls the size of issuer cache and is used with issuer mapping. Starting with Windows Server 2003 operating systems, Schannel attempts to map all of the issuers in the client's certificate chain—not just the one that directly issues the client certificate. When the issuers do not map to an account which is the typical case, the server might attempt to map the same issuer name over and over, hundreds of times a second. To prevent this, Windows Server 2003 has a negative cache, so if an issuer name does not map to an account, then it is added to the cache and Schannel will not attempt to map the issuer name again until the cache entry expires. This registry entry specifies the cache size. This entry does not exist in the registry by default. The default value is 100.

**IssuerCacheTime**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**
Windows Server 2003

This entry controls the length of cache timeout interval in milliseconds. Starting with Windows Server 2003operating systems, Schannel attempts to map all of the issuers in the client's certificate chain—not just the one that directly issues the client certificate. In the case where the issuers do not map to an account which is the typical case, the server might attempt to map the same issuer name over and over, hundreds of times a second. To prevent this, Windows Server 2003 has a negative cache, so if an issuer name does not map to an account, then it is added to the cache and Schannel will not attempt to map the issuer name again until the cache entry expires. This cache is kept for performance reasons, so that the system does not keep trying to map the same issuers over and over. This entry does not exist in the registry by default. The default value is 10 minutes.

**KeyExchangeAlgorithm**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This subkey controls use of key exchange algorithms.

**Diffie-Hellman\Enabled**
The subkey controls use of DH for key exchange.

**PKCS**
This subkey controls use of RSA for key exchange.

The default for these ciphers is enabled.

To disable, create the **Enabled** entry in the appropriate subkey. (This entry does not exist in the registry by default.) After you have created the entry, change the DWORD value to 0. When you disable any algorithm, you disallow all cipher suites that use that algorithm. To re-enable the cipher, change the DWORD value to 0xffffffff.

**MaximumCacheSize**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**

Windows Server 2003, Windows XP, Windows 2000

This entry controls the maximum number of cache elements. Setting MaximumCacheSize to 0 disables the server-side session cache and prevents reconnects. Increasing MaximumCacheSize above the default values causes Lsass.exe to consume additional memory. Each session cache element typically requires 2-4k bytes of memory. This entry does not exist in the registry by default. The default value is 10,000 elements.

**PCT 1.0**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This subkey controls use of Private Communications Transport PCT.

**Client**
This subkey controls use of PCT on client

**Server**
This subkey controls use of PCT on server

To disable, create the **Enabled** entry in the appropriate subkey. (This entry does not exist in the registry by default.) After you have created the entry, change the DWORD value to 0. To re-enable the protocol, change the DWORD value to 0xffffffff.

**DisabledByDefault**
This entry controls disabling PCT by default. This entry does not exist in the registry by default.

**SendTrustedIssuerList**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**
Windows Server 2003

This entry controls the flag controlling sending of list of trusted issuers. In the case of servers that trust hundreds of certificate authorities for client authentication, there are too many issuers for the server to be able to send them all to the client when requesting client authentication. In this situation, this registry key can be set, and instead of sending a partial list, Schannel will not send any to the client.

Not sending a list of trusted issuers might impact what the client sends when asked for a client certificate. For example, when Internet Explorer receives a request for client authentication, it only displays the client certificates that chain up to one of the certificate authorities that is sent by the server. If the server did not send a list, then Internet Explorer displays all of the client certificates that are installed on the client machine. This behavior might be desirable, when PKI environments include cross certificates, the client and server certificates will not have the same Root CA and therefore, Internet Explorer cannot chose a certificate that chains up to on of the server's CAs. By configuring the server to not send a trusted issuer list then Internet Explorer will send all its certificates.

This entry does not exist in the registry by default. This value is true by default.

**ServerCacheTime**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This entry controls the time to expire server side cache entries in milliseconds. A value of 0 disables the server-side session cache and prevents reconnects. Increasing ServerCacheTime above the default values causes Lsass.exe to consume additional memory. Each session cache element typically requires 2-4k bytes of memory. This entry does not exist in the registry by default. The default values are:

### Default Server Cache Time

| Windows Version | Time |
|---|---|
| Windows NT 4.0 with Service Pack 6a | 2 minutes |

| | |
|---|---|
| Windows NT 4.0 with Service Pack 6a and Q265369 | 5 minutes |
| Windows 2000 | 2 minutes |
| Windows 2000 with Service Pack 2 or greater | 10 hours |
| Windows XP | 10 hours |
| Windows Server 2003 | 10 hours |

**SSL 2.0**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This subkey controls use of SSL 2.0.

**Client**
This subkey controls use of SSL 2.0 on the client.

**Server**
This subkey controls use of SSL 2.0 on the server.

To disable, create the **Enabled** entry in the appropriate subkey. (This entry does not exist in the registry by default.) After you have created the entry, change the DWORD value to 0. To re-enable the protocol, change the DWORD value to 0xffffffff.

**DisabledByDefault**
Flag to disable SSL 2.0 by default.

**SSL 3.0**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This subkey controls use of SSL 3.0.

**Client**
This subkey controls use of SSL 3.0 on client

**Server**
This subkey controls use of SSL 3.0 on server

To disable, create the **Enabled** entry in the appropriate subkey. (This entry does not exist in the registry by default.) After you have created the entry, change the DWORD value to 0. To re-enable the protocol, change the DWORD value to 0xffffffff.

**TLS 1.0**
**Registry path**
**HKLM SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols**

**Version**
Windows Server 2003, Windows XP, Windows 2000

This subkey controls use of TLS 1.0.

**Client**
This subkey controls use of TLS 1.0 on client

**Server**
This subkey controls use of TLS 1.0 on server

To disable, create the **Enabled** entry in the appropriate subkey. (This entry does not exist in the registry by default.) After you have created the entry, change the DWORD value to 0. To re-enable the protocol, change the DWORD value to 0xffffffff.

**TLS/SSL Group Policy Settings**
The following table lists and describes the Group Policy settings that are associated with Schannel.

**Group Policy Settings Associated with Schannel**

| Group Policy Setting | Description |
|---|---|
| Security options:<br><br>System cryptography: Use FIPS compliant algorithms for encryption, hashing and signing | Changes to this setting determine whether Schannel will only support the TLS protocol as a client (and as a server, if applicable) and only use:<br><br>• Triple Data Encryption Standard (3DES) encryption algorithm for the TLS traffic encryption<br><br>• Rivest-Shamir-Adleman (RSA) public key algorithm for the TLS key exchange and authentication<br><br>• Secure Hash Algorithm version 1 (SHA-1) hashing algorithm for the TLS hashing<br><br>Both the client and the server must support these algorithms and TLS to communicate using a secure channel application. For example, if you enable this policy setting, you will also need to configure Internet Explorer to use TLS (which is Off by default) to connect using Secure Hypertext Transfer Protocol (HTTPS) to a server with this setting. |

For more information about Group Policy settings, see the "Group Policy Settings Reference for Windows Server 2003" in Tools and Settings Collection.

**Network Ports Used by Schannel**

**Port Assignments for Common Applications over TLS/SSL**

| Service Name | TCP |
|---|---|
| smtp | 25 |
| https | 443 |
| nntps | 563 |
| ldaps | 636 |
| ftps-data | 989 |
| ftps | 990 |
| telnets | 992 |
| imaps | 993 |
| pop3s | 995 |
| ms-sql-s | 1433 |
| mfst-gc-ssl | 3269 |
| tftps | 3713 |

# Authorization and Access Control Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authorization and Access Control Technologies**
Security in the Microsoft Windows Server 2003 operating system controls the use of system and network resources through the interrelated mechanisms of authentication and authorization. After a user is authenticated, Windows Server 2003 uses the authorization and access control technologies to implement the second phase of protecting resources: determining if an authenticated user has the correct authorization to access a resource. Shared resources — resources available to users and groups other than the resource's owner — need to be protected from unauthorized use. In the Windows Server 2003 access control model, users and groups are assigned rights that inform the operating system what each user and group can and cannot do. Shared resources are assigned permissions that enable resource managers to enforce access control in two ways: by denying access to unauthorized users, and by limiting the extent of access provided to authorized users.

**Authorization and Access Control Technologies Architecture**
The authorization and access control model used in Windows Server 2003 is based on the following concepts:

**User-based authorization**
Every application that a user starts runs in that user's security context, not in the application's security context. Applications can also run in a restricted security context, with fewer privileges and more limited access than their user's security context.

**Discretionary access to securable objects**
The user who owns a securable object can control who has permission to use it and in what way. An object's owner can give permission for different kinds of access to particular users or groups of users. Owners can also allow or deny other users access to individual properties of certain types of objects, as well as to the entire object.

**Inheritance of permissions**
You can control permissions for new objects created in a container object by setting inheritable permissions on the container. The permissions that you set on a container are also inherited by existing objects in the container as well as newly created objects.

**Administrative privileges**
You can control which users or groups have the right to perform various administrative functions or to take any action that affects systemwide resources. Domain administrators can use Group Policy to manage privileges on several computers at once or even on all computers joined to a domain.

**Auditing of system events**
The auditing feature detects attempts to circumvent protections on resources and creates an audit trail of administrative actions on the system. If another administrator changes the auditing policy so that failed logon attempts are no longer audited, the log shows this event too. You can also use Group Policy to centrally control who is allowed to manage security logs on computers joined to a domain, as well as to control such configuration options as log size and retention method.

**How Access Control Works**
The access control model used by the Windows Server 2003 operating system ensures authorized use of its objects by security principals. Security principals include users and groups. Security principals perform actions on objects, which include files, folders, printers, registry keys, Active Directory directory service objects, and other types of objects.

Each object has an owner that grants permissions to security principals. During the access control check, these permissions are examined to determine which security principals can access the object and how they can access it.

Object owners generally grant permissions to security groups rather than to individual users. Users added to existing groups adopt the permissions of that group. Object owners also often define permissions for container objects, rather than individual objects, to ease access control management. If an object (such as a folder) can hold other objects (such as subfolders and files), it is called a container. In a hierarchy of objects, the relationship between a container and its content is expressed by referring to the container as the parent and to an object in the container as the child.

Although users are the ones who attempt to access a shared resource, programs actually perform the operations. More specifically, the executable program sequence within a process, called a thread, runs program instructions. A process can have several threads, all executing at the same time. In describing the access control model, the subject taking action is always a thread, never a user, program, or process.

Although only threads can act on system objects, threads do not carry their own security identifiers (SIDs). A thread assumes the rights and privileges of the security principal that initiated the process. When a user logs on, the system creates an access

token for that user. The access token contains the user's SID, the SIDs for any groups the user belongs to, and the user's privileges. This token provides the security context for whatever actions the user executes on that computer.

Threads do not access files in the same way that users do. Threads interact with objects through one of several application programming interfaces (APIs) that are provided by the operating system. For example, the thread that opens a file is probably executing code with the following instruction:

```
hfile=CreateFile(pszFile,GENERIC_WRITE,0,NULL,OPEN_EXISTING,0,NULL);
```

The second argument in the call to CreateFile() specifies a desired set of access rights, GENERIC_WRITE, which indicates to the operating system that the thread wants to open the file and modify it. Other APIs work in a similar fashion. The calling process must signal its intentions for an object by specifying a desired level of access.

Before allowing a subject to proceed with the action it intends to carry out on an object, the operating system's security subsystem performs an access check to determine whether the subject is authorized to perform the action. The access check systematically compares information in two key data structures:

- The subject's access token, which contains a SID for the user and additional SIDs for groups that the user belongs to.
- The object's security descriptor, which contains a discretionary access control list (DACL) with a list of access control entries (ACEs) that specify the access rights that are allowed or denied to particular users or groups, each of whom is identified by a SID.

The security subsystem checks the object's DACL, looking for ACEs that apply to the user and group SIDs in the subject's access token. The system examines each ACE in order until it finds one that either allows or denies access to the user or one of the user's groups, or until there are no more ACEs to check. It is possible for a DACL to have several ACEs that apply to the token's SIDs. If this occurs, the access rights granted by each ACE accumulate until the total matches or exceeds the level of access that the subject has requested. For example, if one ACE grants read access to a group and another ACE grants write access to a user who is a member of the group, the user can have both read/write access to the object. If the access check reaches the end of the DACL and the desired access is still not explicitly allowed or denied, the security subsystem denies access.

The following figure shows the authorization and access control process.

**Authorization and Access Control Process**




**Authorization and Access Control Technologies Components**

The following table provides a summary of the components that comprise the authorization and access control model.

**Components of Authorization and Access Control**

| Components | Description |
|---|---|
| Security Principal | A user, group, or computer. Security principals have accounts. Local accounts are managed by the Security Accounts Manager (SAM) on the computer. Domain accounts are managed by Active Directory. |

| | |
|---|---|
| Security Identifier (SID) | The Windows security model identifies account objects — such as users, groups, computers, and domain trusts — by SIDs. SIDs are domain-unique values, built when the user or group is created or when the computer or trust is registered with the domain. The components of a SID follow a hierarchical convention: A SID contains parts that identify the revision number, the authority — such as the domain — that issued the SID, and a variable number of sub authority or relative identifier (RID) values that uniquely identify the security principal relative to the issuing authority. |
| Access Tokens | After a user logs on and is authenticated, the system creates an access token for the user containing the SID of the user (the primary SID), and the SIDs of all the domain groups the user is a member of. The system uses this access token to determine whether to grant the user access to system resources. |
| Permissions | Authority to perform an operation or a set of operations on an object or object property. Permissions are granted or denied by an object's owner. Because access to an object is at the owner's discretion, the type of access control used in Windows Server 2003 is called discretionary access control.<br><br>Permissions are applied to any secured objects such as files, Active Directory objects, or registry objects. Permissions can be granted to any account — a user, group, or computer. You can assign permissions for objects to:<br><br>● Groups, users, and special identities in the domain.<br><br>● Groups and users in that domain and any trusted domains.<br><br>● Local groups and users on the computer where the object resides.<br><br>The permissions that can be attached to an object depend on the type of object. For example, the permissions that can be attached to a file are different from those that can be attached to a registry key. |
| User Rights | Authority to perform an operation that affects an entire computer rather than a particular object. User rights (also known as privileges) are assigned by administrators to individual users or groups as part of the security settings for the computer. Although user rights can be managed centrally through Group Policy, they are applied locally. Users can (and usually do) have different user rights on different computers. User rights grant specific privileges and logon rights to users and groups in your computing environment. |
| Security Descriptors | A data structure containing the security information associated with a securable object, such as a file or a printer. |
| Access Control Lists (ACLs) | If permissions are configured for an object, its security descriptor contains a discretionary access control list (DACL) with SIDs for the users and groups who are allowed or denied access. If auditing is configured for the object, its security descriptor also contains a system access control list (SACL) that controls how the security subsystem audits attempts to access the object. |

How authorization and access control functions also depends upon related technologies in the Windows Server 2003 security architecture. The following table provides a summary of these related technologies.

**Technologies Related to Authorization and Access Control**

| **Related Components** | **Description** |
|---|---|
| Authentication | During the authentication process, a user is identified to the local or trusted domain by presentation of credentials, usually in the form of a user name and password. Assuming these credentials are acceptable, the system creates an access token for the user. |

| | |
|---|---|
| Auditing | Establishing an audit policy allows you to track potential security problems, help to ensure user accountability, and provide evidence in the event of a security breach.<br><br>The most common types of events to be audited are:<br><br>● Access to objects, such as files and folders.<br><br>● Management of user accounts and group accounts.<br><br>● Users logging on to and logging off from the system.<br><br>You can then view these security-related events in the security log with the Event Viewer. |
| Trusts | By default, all users in a specific Microsoft Windows NT 4.0 or Active Directory domain can connect to resources contained within that domain. This means that a single domain can, by default, provide network connectivity, for users that are members of that domain, to all resources that reside within that domain. To expand resource sharing beyond the boundaries of a single domain, you will need to use trust relationships. Trust relationships provide a security channel for data to flow across domains or forests. |
| SIDHistory | In Active Directory, domain migration or restructuring across trusts is made considerably easier as a result of an attribute of Active Directory security principals called SIDHistory. SIDHistory is used to store the former SIDs of moved objects such as users and security groups. When a user is moved, the SIDHistory attribute of the user object is updated with the former SID. When the user then logs onto the system, the system retrieves the entries in the SIDHistory and adds them to the user access token. In this way SIDHistory ensures that migrated users can continue to access resources even when the resources are located in a trusting domain. |
| SID Filtering | SID filtering prevents domains from accepting SIDs with domain SIDs from outside the sender's domain. Applying SID filtering to trusts can prevent malicious users who have domain administrator level access in the trusted domain from granting, to themselves or other user accounts in their domain, elevated user rights to the trusting domain. |

**Authorization and Access Control Technologies Deployment Scenarios**
Common deployment scenarios for controlling access to resources include the User/ACL method, the Account group/ACL method, the Account group/Resource group method, and the role-based authorization method.

These methods are all based on a planning and deployment process that involves first identifying your users by workgroup, job function, or a combination of workgroup and job function. You can then identify the different types of resources that users might access, such as departmental or job-specific data. This allows you to plan for different levels of access. These access control methods rely on the following concepts:

- **Account group** An account group is a security group whose members are user or computer accounts, all of which require the same permissions for a resource

- **Resource group** A resource group is a security group that has been added to the access control list (ACL) of a resource and granted a specific set of access permissions

- **Principle of Least Access** The principle of least access states that users must have access to the software, data, and devices required to perform their daily duties, but must not have access to local or network resources that are not required for their job tasks. This minimizes potential security risks

In Windows Server 2003, there are four basic strategies for controlling access to resources:

**User/ACL Method**
Security principal accounts are added directly to the ACL for a resource. The ACL is then configured to allow the appropriate permissions for each security principal. This option is appropriate only for managing basic access to a limited number of resources for a small number of users. Otherwise, the complexity and administrator time needed to set and maintain the appropriate permissions becomes burdensome.

**Account Group/ACL Method**
Security groups, rather than individual security principal accounts, are added to the resource ACL, and the group is given a set

of access permissions. The Account group/ACL method is scalable because security groups can be nested if your domain is configured for Microsoft Windows 2000 or Windows Server 2003 native mode functional levels. If groups from multiple domains or from multiple forests require identical access permissions, they can be grouped together into one all-encompassing security group, and you can add this security group to the resource ACL. However, if different groups require different access permissions, this method requires more administrative effort from the resource owner. In this case, each group must be added to the ACL separately, and access permissions must be granted to each group.

**Account Group/Resource Group Method**

Users with similar access requirements are grouped into account groups, and the account groups are added to a resource group that is granted specific resource access permissions. The Account group/resource group method is often most appropriate for large organizations with many shared resources because it is scalable and can be maintained in all environments and at all domain functional levels. This method can be used in any environment, but it is most useful when you cannot nest groups because, for example, your domains are at Windows 2000 mixed functional levels or you need to share resources with Windows NT version 4.0 operating system domains. The Account group/resource group method is also preferable if your resources require a stable set of common permissions.

**Role-based Authorization Method**

Users with similar roles are authorized to perform predefined sets of tasks based on scripts called authorization rules. This allows you to apply fine-grained control over the mapping between access control and the structure and tasks performed in your organization. To use role-based authorization, however, programs need to be written to take advantage of this capability. Additionally, your environment must be operating at the Windows Server 2003 domain functional level.

# Security Descriptors and Access Control Lists Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Descriptors and Access Control Lists Technical Reference**

**In this subject**


- What Are Security Descriptors and Access Control Lists?
- How Security Descriptors and Access Control Lists Work
- Security Descriptors and Access Control Lists Tools and Settings


All objects in the Active Directory directory service, and all securable objects on the network, have security descriptors to help control access to the objects. Security descriptors contain access control lists (ACLs), and they include information about who owns an object, who can access it and in what way, and what types of access are audited.

You can use this access control model to individually secure objects such as files and folders, Active Directory objects, registry keys, and printers, as well as devices, ports, services, processes, and threads. Because of this individual control, you can adjust the security of objects to meet the needs of your organization, delegate authority over objects or attributes, and create custom objects or attributes that require unique security protections to be defined.

# What Are Security Descriptors and Access Control Lists?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Technologies Related to Security Descriptors and Access Control Lists
- Related Information


All objects in Active Directory, and all securable objects on a local computer or on the network, have security descriptors to help control access to the objects. Security descriptors include information about who owns an object, who can access it and in what way, and what types of access are audited. Security descriptors, in turn, contain the access control list (ACL) of an object, which includes all of the security permissions that apply to that object. An object's security descriptor can contain two types of ACLs:

- A discretionary access control list (DACL) that identifies the users and groups who are allowed or denied access
- A system access control list (SACL) that controls how access is audited

You can use this access control model to individually secure objects such as files and folders, Active Directory objects, registry keys, and printers, as well as devices, ports, services, processes, and threads. Because of this individual control, you can adjust the security of objects to meet the needs of your organization, delegate authority over objects or attributes, and create custom objects or attributes that require unique security protections to be defined.

**Technologies Related to Security Descriptors and ACLs**
Security descriptors and ACLs are closely related to the technologies and components of the Windows Server 2003 security architecture that are described in the following sections.

**Authentication and Access Tokens**
Access tokens are created by the security system, and they contain security information about users who have logged on and been authenticated. When a user requests access to an object, the access token of the account requesting access is compared to the object's DACL.

**Permissions**
Each permission that an object's owner grants to a particular user or group is stored as an access control entry (ACE) in a DACL that is part of the object's security descriptor. In the user interface, ACEs are displayed as Permission Entries.

**Auditing**
If auditing is configured for an object, the object's security descriptor also contains a SACL that controls how the security subsystem audits attempts to access the object.

The following figure shows the relationship of security descriptors and ACLs to other key components of the authorization and access control model.

**Relationship of Security Descriptors and ACLs to Other Authorization and Access Control Components**




**Related Information**

The following resources contain additional information that is relevant to this section:

- Security Identifiers Technical Reference
- Permissions Technical Reference
- Logon and Authentication Technologies

# How Security Descriptors and Access Control Lists Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Security Descriptor Architecture
- ACL Architecture
- Inheritance
- Related Information


All objects in Active Directory, and all securable objects on a local computer or on the network, have security descriptors to help control access to the objects. Security descriptors include information about who owns an object, who can access it and in what way, and what types of access are audited. Security descriptors, in turn, contain the access control list (ACL) of an object, which includes all of the security permissions that apply to that object. An object's security descriptor can contain two types of ACLs:

- A discretionary access control list (DACL) that identifies the users and groups who are allowed or denied access
- A system access control list (SACL) that controls how access is audited

The access control model that is used by the Windows Server 2003 operating system is administered at the object level by setting different levels of access, or permissions, to objects. If permissions are configured for an object, its security descriptor contains a DACL with security identifiers (SIDs) for the users and groups that are allowed or denied access. If auditing is configured for the object, its security descriptor also contains a SACL that controls how the security subsystem audits attempts to access the object.

You can use this access control model to individually secure objects such as files and folders, Active Directory objects, registry keys, and printers, as well as devices, ports, services, processes, and threads. Because of this individual control, you can adjust the security of objects to meet the needs of your organization, delegate authority over objects or attributes, and create custom objects or attributes that require unique security protections to be defined.

The following sections define the basic elements of security descriptors and ACLs, look at where access control information comes from, discuss how access control permissions are inherited, and describe how ACLs work in an optimal environment. An optimal environment for ACLs is defined as follows:

- Security on objects is designed using the theory of least privilege. Users are assigned the least level of access that they need to perform their job functions.
- Groups, rather than individual user accounts, are used to grant or deny access to resources.
- A consistent model is used for assigning permissions.
- Permissions on Active Directory objects are not assigned to domain local groups.

**Securable Objects**

A securable object is an object that can have a security descriptor. All named Windows objects are securable. Some unnamed objects, such as process and thread objects, can also have security descriptors. Each type of securable object defines its own set of specific access rights and its own mapping of generic access rights.

The following are the common securable objects:

- Files or folders on an NTFS file system

- Active Directory objects
- Registry keys
- Network shares
- Local or remote printers
- Windows services
- Named pipes
- Anonymous pipes
- Processes
- Threads
- File-mapping objects
- Access tokens
- Window-management objects (window stations and desktops)
- Interprocess synchronization objects (events, mutexes, semaphores, and waitable timers)
- Job objects
- Distributed Component Object Model (DCOM) objects

**Security Descriptor Architecture**
Access control information that is associated with an object is contained in the object's security descriptor. When a user tries to access the object, the operating system examines the object's security descriptor to determine whether the access will be granted. Before allowing a subject to proceed with the action it intends to carry out on an object, the operating system's security subsystem performs an access check to determine whether the subject is authorized to perform the action on that object.

The information that is included in a security descriptor depends on the type of object and how it was created. In general, security descriptors can include information about the following:

- Which user owns the object
- Which users and groups are allowed or denied access
- Which users and groups' access should be audited
- How objects in a container inherit access control information from the container

**Parts of a Security Descriptor**
A security descriptor is a binary data structure of variable length. The data structure has the following parts:

**Header**
The header contains a revision number and a set of control flags that describe characteristics of the security descriptor, such as the memory layout, which elements are present, and how particular elements were added or modified.

**Owner**
The Owner field contains the SID for the object's owner. The owner of an object can modify permissions and give other users

the right to take ownership.

**Primary Group**
The Primary Group field contains the SID for the owner's primary group. This information is used only by the POSIX subsystem, and it is ignored by the rest of Windows Server 2003.

**DACL**
The DACL is a list of zero or more access control entries (ACEs). Each ACE in the list has a header that specifies whether the ACE allows or denies access, a SID that specifies a particular user or group, and an access mask that lists the operations that are allowed or denied. The content of the DACL is controlled by the object's owner. The owner can share control with other users by granting them Change Permissions (WRITE_DAC) permission. Note that it is possible for a security descriptor to have no DACL (also known as a NULL DACL); this gives unconditional access to everyone. A security descriptor with an empty DACL gives no access to anyone. For more information, see "No DACL vs. Empty DACL" later in this section.

**SACL**
The SACL is similar to the DACL except that the SACL is used to audit rather than control access to an object. When an audited action occurs, the operating system records the event in the security log. Each ACE in a SACL has a header that indicates whether auditing is triggered by success or failure or both, a SID that specifies a particular user or security group to monitor, and an access mask that lists the operations to audit. The content of the SACL is controlled by security administrators for the local system. Security administrators are users who have been assigned the Manage Auditing and Security Log (SeSecurityPrivilege) privilege. By default, this privilege is assigned to the built-in Administrators group.

**Layout in Memory**
A security descriptor's layout in memory can be either self-relative or absolute. Which of the two formats a particular security descriptor uses is indicated by a control flag in the security descriptor's header.

A self-relative security descriptor is stored in a contiguous block of memory, and the address for each part of the descriptor is expressed as an offset from the beginning of the memory block. To locate any part of the security descriptor, a process must be given only the address of the area in memory where the security descriptor begins. The process determines where the part in question begins by adding the memory offset for the part to the memory address for the security descriptor. (Hence, the name "self-relative"; the address for a part is relative to the address for the whole.)

The following figure illustrates the layout of a security descriptor in self-relative format.

**Security Descriptor in Self-Relative Format**




Self-relative layout is used for security descriptors on objects that must be stored on disk, transmitted by a communications protocol, or copied in memory.

A security descriptor in absolute format contains pointers to its information, not the information itself. This format enables each component of the security descriptor to be allocated separately and assembled on the fly. It is used for objects that do not require persistent storage, such as objects that exist only during the lifetime of a single process. The following figure illustrates the layout of a security descriptor in absolute format.

**Security Descriptor in Absolute Format**




**Security Descriptor Control Flags**
A security descriptor's header contains a set of control flags that qualify the meaning of the security descriptor or its components. In Windows 2000 and later operating systems, control flags play an important role in the automatic propagation of inheritable security information from parent (container) objects to child (contained) objects.

Security descriptor control flags are stored in a bit field, and they are turned on or off by setting individual bits. The following table lists the security descriptor control flags and their meaning.

**Security Descriptor Control Flags**

| Flag | Meaning |
| --- | --- |
| SE_DACL_AUTO_INHERITED | **Windows 2000 and later:** Inheritable ACEs in this object's DACL have been propagated to existing child objects.<br><br>This flag is not set in security descriptors for Windows NT, which does not support automatic propagation of inheritable ACEs. |
| SE_DACL_DEFAULTED | The DACL was provided by a default mechanism.<br><br>This flag can affect how the operating system treats the DACL with respect to inheritance. The operating system ignores this flag if SE_DACL_PRESENT is not set. |
| SE_DACL_PRESENT | The security descriptor has a DACL.<br><br>**Windows 2000 and later:** If this flag is not set (that is, if the security descriptor has no DACL), SE_DACL_PROTECTED must be set. Otherwise, the operating system considers the security descriptor to be invalid. |
| SE_DACL_PROTECTED | **Windows 2000 and later:** The security descriptor's DACL cannot be modified by inheritable ACEs.<br><br>If this flag is not set, the security descriptor inherits information from the security descriptor of the parent object. |
| SE_GROUP_DEFAULTED | The primary group SID was provided by a default mechanism. |
| SE_OWNER_DEFAULTED | The owner SID was provided by a default mechanism. |
| SE_SACL_AUTO_INHERITED | **Windows 2000 and later:** Inheritable ACEs in this object's SACL have been propagated to existing child objects.<br><br>This flag is not set in security descriptors for Windows NT, which does not support automatic propagation of inheritable ACEs. |
| SE_SACL_DEFAULTED | The SACL was provided by a default mechanism.<br><br>This flag can affect how the operating system treats the SACL with respect to inheritance. The operating system ignores this flag if SE_SACL_PRESENT is not set. |
| SE_SACL_PRESENT | The security descriptor has a SACL. |
| SE_SACL_PROTECTED | **Windows 2000 and later:** The security descriptor's SACL cannot be modified by inheritable ACEs. |
| SE_SELF_RELATIVE | The security descriptor is in self-relative format with all information in a contiguous block of memory. If this flag is not set, the security descriptor is in absolute format. |

**Where Access Control Information Comes From**

Access control information is first written to an object's security descriptor when the object is created. This information can later be modified. In both cases, the information that goes into a security descriptor can come from one of the following sources:

- The subject
- The Object Manager
- The parent object

When a subject creates a new object, it can assign a security descriptor to the object. If the subject does not assign a security descriptor, the operating system uses access control information inherited from the parent object to create one. If no information is available to inherit, the operating system uses default access control information provided by the object manager for the particular type of object that the subject wants to create.

After an object is created, information in its security descriptor can be changed by the object's owner or by another user who

has the owner's permission. Changes can also come from the parent object when that object's owner modifies its security descriptor. Every time the security descriptor on a container object is changed, the object manager propagates any changes that are marked as inheritable to all objects in the container.

**Subjects**
A subject is a thread executing in the security context provided by an access token. A subject's access token contains information that can be used in security descriptors for new objects that the subject creates and existing objects that the subject modifies. This information includes the following fields.

**Default Owner**
This field of an access token contains the SID that identifies the security principal that, by default, becomes the owner of objects that the subject either creates or takes ownership of. Typically, the default owner is the individual user that is represented by the access token. When the SID in the Default Owner field of the subject's access token is copied to the Owner field of an object's security descriptor, SE_OWNER_DEFAULTED is set in the security descriptor's control flags.

### Note

- When users who are members of the local Administrators group access objects on Windows Server 2003, the Default Owner field in the user's access token contains the SID for the Administrators group, not the SID for the user. A similar rule applies to users who access objects in Active Directory. If the user is a member of the Domain Admins group, the Default Owner field in the user's access token contains the SID for the Domain Admins group. In both cases, objects that the user creates or takes ownership of are owned by the group, not by the individual user.

- On Windows XP Professional, the Default Owner field always contains the user's SID. This is a change from Windows 2000 and earlier operating systems, in which the rules were the same for both Windows 2000 Server and Windows XP Professional.

**Primary Group**
This field of an access token contains the SID for the user's primary group. When the SID in the Primary Group field of the subject's access token is copied to the Primary Group field of an object's security descriptor, SE_GROUP_DEFAULTED is set in the security descriptor's control flags.

**Default DACL**
This field of an access token might contain a DACL, or it might be empty. If the subject's access token contains a DACL, it can be used as the DACL in a new object's security descriptor. If the Default DACL field is used in the security descriptor for a new object, SE_DACL_DEFAULTED is set in the security descriptor's control flags.

**Object Managers**
The object manager can provide default security information for new objects. In Windows Server 2003, the only object manager that supplies default permissions for new objects is Active Directory.

For more information, see "DACLs for Newly Created Active Directory Objects" later in this section.

**Parent Objects**
Some objects can contain other objects. For example, an NTFS folder object can contain file objects and other folder objects. A registry key object can contain subkey objects. An Active Directory organizational unit (OU) object can contain other OU objects as well as user objects, group objects, and computer objects. Terminal objects contain Window Station objects that contain Desktop objects that contain Window objects. Any object that is contained by another object is called a "child" object. A "child" object's container is its "parent" object.

Child objects can inherit access control information from their parent object. For example, suppose that the administrator for a server creates a file share with one folder, Public$. The administrator creates this folder so that users can have a place to store information that they want to share. With this purpose in mind, the administrator sets the permissions, which are implemented as ACEs, as shown in the following figure.

### Permission Entries on Public Folder (Owner: Administrators)




None of the permissions that are listed in this figure were acquired through inheritance. This is because the administrator cleared the check box for **Inherit from parent the permission entries that apply to child objects**. Clearing the check box sets the security descriptor control flag SE_DACL_PROTECTED, which protects a child object's DACL by blocking inheritance from the parent object's DACL.

Permissions that are acquired through inheritance are called "inherited" permissions. Permissions that are not inherited, but are instead defined directly on an object, are called "explicit" permissions. One way to tell an explicit permission from an inherited permission is to select an entry in the **Permission Entries** list and read the text that is displayed after the list. In this figure, the second entry is selected, and the text after the list says **This permission is defined directly on this object**. In other words, the permission is explicit, not inherited.

The text in this figure also says **This permission is inherited by child objects**. Permissions on a parent object that apply to child objects are called "inheritable" permissions. To see which of the permissions that are set on a parent object are inheritable, examine the **Apply to** column of **Permission Entries**. If **Apply to** says **This object only** (or, for folder objects, **This folder only**), the permission is not inherited by child objects. Of the four permissions that are shown in this figure, three are inheritable and one is not.

To see how inheritable permissions become inherited permissions, suppose that Alice creates a subfolder in Public$. Alice is an engineer, so she names her folder Engineering Data. Because this new object is a child of Public$, its DACL inherits permissions from the DACL on Public$. The new object's permissions are shown in the following figure.

**Permission Entries on Engineering Data Folder (Owner: Alice)**




Note that Alice has not cleared the check box for **Inherit from parent the permission entries that apply to child objects**, so inheritable permissions in the parent object's DACL are inherited by the child object's DACL. Inherited permissions are indicated in **Permission Entries** by a disabled (unavailable) symbol at the beginning of each entry. The permission is still effective; all that is disabled is the ability to modify the entry. Because inherited permissions are defined on a parent object, they can be changed only by modifying the parent object's DACL.

Even though inherited permissions cannot be changed, the owner of a child object can add explicit permissions to the object's DACL. For example, suppose Alice decides that inherited permissions that are given to Creator Owner are too restrictive because they allow only the user who creates a file to make changes to the file. She wants all members of the Engineering group to be able to edit and add information to the Engineering Data folder, so she explicitly gives this group Modify permission for all objects in the folder. Alice also feels that people in her company's marketing department will misuse information in Engineering Data, so she decides to explicitly deny the Marketing group full control of (and therefore all access to) the folder, subfolders, and files. The results of Alice's changes to the access control settings are shown in the following figure.

**Modified Permission Entries on Engineering Data Folder**




The list of permission entries in this figure now includes two explicit permissions, both with enabled symbols indicating that

the entries can be edited. Note that explicit permissions appear at the top of the list. Permissions are listed in the order in which they will be processed during an access check. Because explicit permissions are listed before inherited permissions, they are processed first. The assumption is that the owner of a child object adds explicit permissions in order to qualify inherited permissions. For example, in this figure, an inherited permission allows Everyone to read the folder, subfolders, and files. Alice has added an explicit permission that denies all access to a subset of the group Everyone — the Marketing group. The explicit deny entry is placed before any inherited entries; therefore, it is processed before any inherited entries. For more information about ACE order, see "Order of ACEs in a DACL" later in this section.

Alice applied both of her explicit permissions to **This folder, subfolders, and files** because she wants child objects in her folder to inherit the new permissions. If the Engineering Data folder were on a server running Windows NT, the new explicit permissions would affect only new objects created in her folder, because the access control model used in Windows NT supports inheritance only at the time that an object is created. To change permissions on existing objects in a Windows NT system, Alice would need to take the additional step of setting explicit permissions on each existing subfolder and file. What is new in Windows 2000 and later is inheritance after the time of creation. New or changed inheritable permissions in the DACL on a parent object are automatically propagated to existing child objects every time the DACL on the parent object changes. If Alice's folder were on a system running Windows 2000 or later, the entry denying Marketing permission to access the Engineering Data folder would be propagated to subfolders as soon as Alice clicked **Apply** in the **Advanced Security Settings** dialog box.

Automatic propagation of inheritable permissions is a powerful capability because you can use it to change permissions on an entire tree of objects by changing permissions on the top-level object in the tree. For example, suppose a member of the Administrators group realizes that information in the Public$ folder and its subfolders is proprietary, not public. In that case, it is inappropriate to allow Everyone access because that group includes the Guest account. The administrator decides that it is more prudent to allow access only to Authenticated Users, a group that includes only users whose identities can be authenticated by the server or by a trusted security authority. To tighten security on the folder, the administrator simply needs to change Everyone to Authenticated Users for two entries in the permissions list. The change is shown in the following figure.

**Modified Permission Entries on Public Folder**




When the administrator applies the changes to the access control settings on the parent object, all inheritable permissions in the parent object's DACL are propagated to DACLs on subfolders and files. The impact on the Engineering Data folder is shown in the following figure.

**Permission Entries Propagated to Engineering Data Folder**




Note that propagation of inheritable permissions from the Public$ folder to the Engineering Data folder does not change explicit permissions in the child's DACL. Only inherited permissions are replaced when inheritable permissions are propagated to existing child objects. However, if explicit permissions are also inheritable, the propagation process reapplies them as it moves down the tree. For example, both explicit permissions that Alice added to the DACL on her folder are inheritable by child objects in the folder. As the propagation process moves downward from Alice's folder, it picks up these additional inheritable permissions and applies them to the DACL of any child object that it finds.

The owner of a parent object can choose to overwrite explicit permissions that are defined on child objects. This is done by selecting the **Replace permission entries on all child objects with entries shown here that apply to child objects** check box in the **Advanced Security Settings** dialog box. When the owner of a parent object chooses this option, the propagation process removes explicit permissions from the DACLs on all child objects. It also sets the option **Inherit from parent the permission entries that apply to child objects** on all child objects, removing any protection from inheritance that might have been set by the objects' owners.

**ACL Architecture**

An ACL is an ordered list of ACEs that define the protections that apply to an object and its properties. Each ACE identifies a security principal and specifies a set of access rights that are allowed, denied, or audited for that security principal.

An object's security descriptor can contain two ACLs:

- A DACL that identifies the users and groups that are allowed or denied access
- A SACL that controls how access is audited

The following figure shows the data structure for an ACL.

**Data Structure of an ACL**

| | |
|---|---|
| ACL Size | ACL Revision |
| ACE Count | |
| ACE [1] | |
| ACE [...] | |
| ACE [*n*] | |

**Parts of an ACL**

The individual parts of an ACL are as follows.

**ACL size**

The number of bytes of memory allocated for the ACL. The size of an ACL varies with the number and size of its ACEs. The maximum size of an ACL is 64K, or approximately 1,820 ACEs, depending on the size of the ACEs. However, for performance reasons, the maximum size is not practical.

**ACL revision**

The revision number for the ACL's data structure. The structure of an ACL is the same for all revisions, but the structure of ACEs in the ACL can vary.

**ACE count**
The number of ACEs in the ACL. A value of zero means the ACL has no ACEs — it is empty; therefore, access-checking can stop. Note that an empty DACL differs from a NULL DACL. An empty DACL gives no access to anyone, but a NULL DACL gives unconditional access to everyone, and should therefore be avoided. For more information, see "No DACL vs. Empty DACL" later in this section.

**ACEs**
An ordered list containing zero or more ACEs. During an access check, ACEs are processed in the order in which they are listed.

**How the System Uses ACLs**
When a thread tries to access a securable object, the system either grants or denies access. If the object does not have a DACL, the system grants access; otherwise, the system looks for ACEs in the object's DACL that apply to the thread. Each ACE in the object's DACL specifies the access rights that are allowed or denied for a security principal or logon session.

The system compares the trustee in each ACE to the trustees that are identified in the thread's access token. An access token contains SIDs that identify the user and the group accounts to which the user belongs. A token also contains a logon SID that identifies the current logon session.

The system examines each ACE in sequence until one of the following events occurs:

- An access-denied ACE explicitly denies any of the requested access rights to one of the trustees listed in the thread's access token.
- One or more access-allowed ACEs for trustees listed in the thread's access token explicitly grant all the requested access rights.
- All ACEs have been checked and there is still at least one requested access right that has not been explicitly allowed, in which case, access is implicitly denied.

Because the system stops checking ACEs when the requested access is explicitly granted or denied, the order of ACEs in a DACL is important.

**Access Control Entries**
All ACEs include the following access control information:

- A SID that identifies a user or group
- An access mask that specifies access rights
- A set of bit flags that determine whether child objects can inherit the ACE
- A flag that indicates the type of ACE

**ACE Types**
Windows 2000 and Windows Server 2003 support six types of ACEs. Three are generic ACE types that can be present in ACLs that are attached to all securable objects. The three remaining ACE types are object specific and can occur only in ACLs for Active Directory objects. The following tables list the generic and object-specific ACE types.

**Generic ACE Types**

| Type | Description |
|---|---|
| Access-denied | Used in a DACL to deny access |
| Access-allowed | Used in a DACL to allow access |
| System-audit | Used in a SACL to log attempts to access |

**Object-Specific ACE Types**

| Type | Description |
|---|---|
| Access-denied, object-specific | Used in a DACL to deny access to a property or property set or to limit inheritance to a specified type of child object |
| Access-allowed, object-specific | Used in a DACL to allow access to a property or property set or to limit inheritance to a specified type of child object |
| System-audit, object-specific | Used in a SACL to log attempts to access a property or property set or to limit inheritance to a specified type of child object |

Generic and object-specific ACEs are fundamentally alike. What sets them apart is the degree of control they offer over inheritance and object access.

Generic ACEs offer limited control over the kinds of child objects that can inherit them. Essentially, they can distinguish only between containers and noncontainers. For example, the DACL on a Folder object in NTFS can include a generic ACE that allows a group of users to list the folder's contents. This is an operation that can be performed only on container objects; therefore, the ACE that allows the operation can be flagged as a CONTAINER_INHERIT_ACE. Only container objects in the folder (that is, only other Folder objects) inherit the ACE. Noncontainer objects (that is, File objects) do not inherit the ACE of the parent object.

Object-specific ACEs offer a greater degree of control over the types of child objects that can inherit them. For example, an OU object's ACL can have an object-specific ACE that is marked for inheritance only by User objects. Other types of objects, such as Computer objects, will not inherit the ACE. This capability is why object-specific ACEs are called object specific. Their inheritance can be limited to specific types of child objects.

There are similar differences in how the two categories of ACE types control access to objects. Generic ACEs apply to an entire object. If a generic ACE gives a particular user Read access, the user can read all the information that is associated with the object — both data and properties. This is not a serious limitation for most object types. File objects, for example, have few properties, which are all used for describing characteristics of the object rather than for storing information. Most of the information in a File object is stored as object data; therefore, there is little need for separate controls on a file's properties.

Object-specific ACEs can apply to any individual property of an object or to a set of properties. These ACE types are used only in ACLs for Active Directory objects, which, unlike other object types, store most of their information in properties. It is often desirable to place independent controls on each property of an Active Directory object, and object-specific ACEs make that possible. For example, when you define permissions for a User object, you can use one object-specific ACE to allow Principal Self (that is, the user) Write access to the Phone-Home-Primary (homePhone) property, and you can use other object-specific ACEs to deny Principal Self access to the Logon-Hours (logonHours) property and other properties that set restrictions on the user account.

**Structure of a Generic ACE**

All three generic ACE types — access denied, access allowed, and system audit — have the same data structure, which is illustrated in the following figure.

### Structure of a Generic ACE

| | |
|---|---|
| ACL Size | ACL Type |
| Inheritance and Audit Flags | |
| Access Mask | |
| SID | |

The individual parts of an ACE are as follows.

**ACE Size**

The number of bytes of memory that are allocated for the ACE.

**ACEType**

Specifies whether the ACE allows, denies, or monitors access.

**Inheritance and Audit Flags**

A set of bit flags that control inheritance and auditing. For information about inheritance flags, see "Inheritance" later in this section. The following table describes audit flags.

### Audit Flags

| Flag | Meaning |
|---|---|

| | |
|---|---|
| FAILED_ACCESS_ACE_FLAG | Meaningful only in system-audit and system-audit object ACEs. The access mask specifies operations that are logged when they fail. |
| SUCCESSFUL_ACCESS_ACE_FLAG | Meaningful only in system-audit and system-audit object ACEs. The access mask specifies operations that are logged when they succeed. |

**Access Mask**
A 32-bit value whose bits correspond to access rights for the object. Bits can be set either on or off, but the setting's meaning depends on the ACE type. For example, if the bit that corresponds to the right to read permissions is turned on, and the ACE type is Deny, the ACE denies the right to read the object's permissions. If the same bit is set on but the ACE type is Allow, the ACE grants the right to read the object's permissions.

**SID**
Identifies a user or group whose access is controlled or monitored by this ACE.

**Structure of an Object-Specific ACE**
The following figure shows the structure of an object-specific ACE.

**Structure of an Object-Specific ACE**

| | |
|---|---|
| ACE Size | ACE Type |
| Inheritance and Audit Flags | |
| Access Mask | |
| Object Type | Inherited Object Type |
| Inheritance and Audit Flags | |

The fields for ACE Size, ACEType, Inheritance and Audit Flags, Access Mask, and SID are identical to like elements in the data structure for a generic ACE. The key differences between a generic and an object-specific ACE are as described in the following sections.

**Object flags**
Object flags indicate whether Object Type or Inherited Object Type is present. The following table shows the three possible object flags.

**Object Flags**

| Flag | Meaning |
|---|---|
| 0 (no flags) | Neither Object Type nor Inherited Object Type is present. In this case, the ACE applies to the entire object. It is, in effect, a generic ACE. |
| ACE_OBJECT_TYPE_PRESENT | ACE applies to a property, property set, or extended right, or it controls the ability to create a particular type of child object. |
| ACE_INHERITED_OBJECT_TYPE_PRESENT | ACE can be inherited only by a particular type of child object. |

**Object Type**
Object Type contains a globally unique identifier (GUID) that identifies one of the following:

- schemaIDGUID. When the schema ID GUID is specified in the ACE, the ACE controls who can create a child object from the identified class schema within a container. The SID portion of the ACE identifies a user or group who can create this type of child object. The ACE's access mask contains the object-specific access right ADS_RIGHT_DS_CREATE_CHILD.

- propertyGUID. When a property GUID is specified in the ACE, the ACE controls the ability to read or write a particular property or property set. The SID portion of the ACE identifies a user or group who can read or write the property or property set. The ACE's access mask contains either ADS_RIGHT_DS_READ_PROP or ADS_RIGHT_DS_WRITE_PROP.

- rightsGUID. When an extended rights GUID is specified in the ACE, the ACE controls the right to perform the operation that is associated with the extended right. The SID portion of the ACE identifies a user or group who has the extended right. The ACE's access mask contains ADS_RIGHT_DS_CONTROL_ACCESS. For a list of Extended rights GUIDs, see the Extended Rights Reference.

**Inherited Object Type**
Inherited Object Type contains a GUID that identifies the type of child object that can inherit the ACE. Inheritance is also

controlled by the ACE's Inheritance Flags and by any protection against inheritance placed on the child object in its Security Descriptor Control Flags.

**DACLs for Newly Created Objects**
The operating system uses the following rules to set the DACL in the security descriptors for most types of new securable objects:

1. The object's DACL is the DACL from the security descriptor, which is specified by the creating process. The operating system merges any inheritable ACEs into the DACL unless SE_DACL_PROTECTED is set in the security descriptor control flags. The operating system then sets the SE_DACL_PRESENT security descriptor control flag.

2. If the creating process does not specify a security descriptor, the operating system builds the object's DACL from inheritable ACEs in the parent object's DACL. The operating system then sets the SE_DACL_PRESENT security descriptor control flag.

3. If the parent object has no inheritable ACEs, the operating system asks the object manager to provide a default DACL. The operating system then sets the SE_DACL_PRESENT and SE_DACL_DEFAULTED security descriptor control flags.

4. If the object manager does not provide a default DACL, the operating system checks the subject's access token for a default DACL. The operating system then sets the SE_DACL_PRESENT and SE_DACL_DEFAULTED security descriptor control flags.

5. If the subject's access token does not have a default DACL, the new object is assigned no DACL, which allows unconditional access by everyone. The SE_DACL_PRESENT security descriptor control flag is not set.

**DACLs for Newly Created Active Directory Objects**
The method that is used to build a DACL for a new Active Directory object is slightly different from the method that is used to build DACLs for other object types. There are two key differences:

- The rules for creating a DACL distinguish between generic inheritable ACEs and object-specific inheritable ACEs in the parent object's security descriptor. Generic inheritable ACEs can be inherited by all types of child objects. Object-specific inheritable ACEs can be inherited only by the type of child object to which they apply.

- The Active Directory schema can provide a security descriptor. Each object class that is defined in the schema has a **defaultSecurityDescriptor** attribute. If neither the creating process nor inheritance from the parent object provides a DACL for a new Active Directory object, the operating system uses the DACL in the default security descriptor that is specified by the schema.

The operating system uses the following rules to set the DACL in the security descriptor for new Active Directory objects:

1. The object's initial DACL is the DACL from the security descriptor that is specified by the creating process.

2. The operating system merges inheritable ACEs on the parent object into the new object's DACL. It then sets the SE_DACL_PRESENT security descriptor control flag.

3. If the parent object has no inheritable object-specific ACEs for the type of object being created, the operating system uses the default DACL from the Active Directory schema for that object type. It then sets the SE_DACL_PRESENT and SE_DACL_DEFAULTED security descriptor control flags.

4. If the Active Directory schema does not specify a default DACL for the object type, the operating system checks the subject's access token for a default DACL. It then sets the SE_DACL_PRESENT and SE_DACL_DEFAULTED security descriptor control flags.

5. If the subject's access token does not have a default DACL, the new object is assigned no DACL, which gives unconditional access to everyone. The SE_DACL_PRESENT security descriptor control flag is not set.

**SACLs for Newly Created Objects**
The operating system uses the following rules to set the SACL in the security descriptors for new securable objects:

1. If the creating process provides an explicit SACL, the operating system places it in the object's security descriptor. The operating system merges any inheritable ACEs into the SACL unless SE_SACL_PROTECTED is set in the security descriptor control flags. It then sets the SE_SACL_PRESENT security descriptor control flag.

2. If the creating process does not provide an explicit SACL, the operating system builds the object's SACL from inheritable ACEs in the parent object's SACL. It then sets the SE_SACL_PRESENT security descriptor control flag.

3. If the parent object has no inheritable ACEs, the operating system asks the object manager to provide a default SACL. It then sets the SE_SACL_PRESENT and SE_SACL_DEFAULTED security descriptor control flags.

4. If the object manager does not provide a default SACL, the new object is assigned no SACL.

**Inheritance**

Inheritance is the process that propagates ACEs from the ACL of a parent object to the ACL of a child object. In Windows Server 2003, inheritable ACEs can be propagated from parent to child when one of the following events takes place:

- A new child object is created.

- The DACL on the parent object is modified.

- The SACL on the parent object is modified.

In this model, any object can be the child of another object. Only container objects can be parents. The ACL for a container object can carry ACEs that are not effective on the container but are present only for the purpose of inheritance — only so that they can be passed down to subsequent levels of objects in the container hierarchy until they reach a child object that is not a container, where they become effective ACEs.

The mechanism for inheritance depends on two things: a set of inheritance flags in each ACE and a set of inheritance rules that are built into the operating system.

**Inheritance Flags**

The header for an ACE contains a set of inheritance flags that control how the ACE is inherited and how the ACE affects a child object that inherits it. The following table lists the inheritance flags.

### Inheritance Flags

| Flag | Meaning |
|---|---|
| INHERITED_ACE | **Windows 2000 and later**: The ACE was inherited from a parent object's DACL or SACL.<br><br>This flag is not set in an explicit ACE — an ACE defined directly on the object. |
| INHERIT_ONLY_ACE | Indicates that this is an inherit-only ACE. The ACE is ignored during an access check but can be propagated to child objects.<br><br>If this flag is not set, the ACE is an effective ACE — an ACE that is processed during an access check.<br><br>Both effective and inherit-only ACEs can be inherited. Whether an ACE is inherited depends on the state of the flags for OBJECT_INHERIT_ACE and CONTAINER_INHERIT_ACE. |
| CONTAINER_INHERIT_ACE | Container objects inherit this ACE as an effective ACE. When the ACE is inherited by a container object, the operating system clears the flag for INHERIT_ONLY_ACE. |
| OBJECT_INHERIT_ACE | Noncontainer objects inherit this ACE as an effective ACE. When the ACE is inherited by a noncontainer object, the operating system clears the flag for INHERIT_ONLY_ACE.<br><br>Container objects also inherit this ACE but only for the purpose of inheritance. When the ACE is inherited by a container object, the operating system sets the flag for INHERIT_ONLY_ACE. |

| | |
|---|---|
| NO_PROPAGATE_INHERIT_ACE | If a child object inherits an ACE where this flag is set, the operating system clears the flags for OBJECT_INHERIT_ACE and CONTAINER_INHERIT_ACE. This prevents the ACE from being inherited by subsequent generations of objects. |

**Rules of Inheritance**

The operating system interprets the inheritance flags and other inheritance information according to the rules of ACE inheritance that are outlined in the following table. These rules function the same way for both DACLs and SACLs. When the operating system propagates inheritable ACEs to child objects, it observes the preferred (or canonical) order. After ACEs have been propagated, the system sets the INHERITED_ACE flag in all inherited ACEs.

The following table lists the rules of inheritance from parent ACEs to child ACLs.

**Rules of Inheritance**

| Parent ACE Inheritance Flags | Effect on Child ACL |
|---|---|
| No flags | None. |
| OBJECT_INHERIT_ACE only | **Noncontainer child objects:** Inherited as an effective ACE.<br><br>**Container child objects:** Containers inherit an inherit-only ACE unless the NO_PROPAGATE_INHERIT_ACE flag is also set. |
| CONTAINER_INHERIT_ACE only | **Noncontainer child objects:** No effect on the child object.<br><br>**Container child objects:** The child object inherits an effective ACE. The inherited ACE is inheritable unless the NO_PROPAGATE_INHERIT_ACE flag is also set. |
| CONTAINER_INHERIT_ACE and OBJECT_INHERIT_ACE | **Noncontainer child objects:** Inherited as an effective ACE.<br><br>**Container child objects:** The child object inherits an effective ACE. The inherited ACE is inheritable unless the NO_PROPAGATE_INHERIT_ACE flag is also set. |

If an inherited ACE is an effective ACE for the child object, the system maps any generic rights to the specific rights for the child object. Similarly, the system maps generic SIDs, such as CREATOR_OWNER, to the appropriate SID. If an inherited ACE is an inherit-only ACE, any generic rights or generic SIDs are left unchanged so that they can be mapped appropriately when the ACE is inherited by the next generation of child objects.

When a container object inherits an ACE that is both effective on the container and inheritable by its descendants, the container might inherit two ACEs. This occurs if the inheritable ACE contains generic information. The container inherits an inherit-only ACE containing the generic information, and an effective ACE in which the generic information is mapped.

An object-specific ACE has an Inherited Object Type field that can contain a GUID to identify the type of object that can inherit the ACE. If the field does not contain a GUID, the inheritance rules are the same as for a standard ACE. If the field contains a GUID, the ACE is inheritable by objects that match the GUID if OBJECT_INHERIT_ACE is set and by containers that match the GUID if CONTAINER_INHERIT_ACE is set.

**Order of ACEs in a DACL**

The preferred order of ACEs in a DACL is called the “canonical” order. For Windows 2000 and Windows Server 2003, the canonical order is the following:

- All explicit ACEs are placed in a group before any inherited ACEs.
- Within the group of explicit ACEs, access-denied ACEs are placed before access-allowed ACEs.
- Inherited ACEs are placed in the order in which they are inherited. ACEs that are inherited from the child object’s parent come first, then ACEs inherited from the grandparent, and so on up the tree of objects.

The following figure shows the canonical order of ACEs.

**Canonical Order of ACEs**




The canonical order ensures that an explicit access-denied ACE is enforced regardless of any explicit access-allowed ACE. This means that the object's owner can define permissions that allow access to a group of users, such as Everyone, but deny access to a subset of the group, such as Marketing. If the object's ACEs are in canonical order, the ACE that denies Marketing comes before the ACE that allows Everyone. During an access check, the operating system steps through the ACEs in the order in which they appear in the object's DACL, so that the deny ACE is processed before the allow ACE. As a result, users who are members of the Marketing group are denied access. Everyone else is allowed access to the object.

The canonical order also ensures that all explicit ACEs are processed before any inherited ACE. This is consistent with the concept of discretionaryaccess control: access to a child object is at the discretion of the child's owner, not the parent's owner. The owner of a child object can define permissions directly on the child that modify the effects of inherited permissions. For example, suppose a parent object has an inheritable ACE that denies access to Marketing. And suppose the owner of a child object defines an explicit ACE that allows access to a subset of Marketing — let's say a user named Bob. If the child object's ACEs are in canonical order, the explicit ACE that allows Bob access comes before any inherited ACE, including the inherited ACE that denies access to Marketing. During an access check, the operating system reaches the ACE that allows Bob access before it gets to the ACE that denies access to Marketing. As a result, Bob is allowed access to the object even though he is a member of the Marketing group. Other members of the Marketing group are denied access.

**What Happens During an Upgrade from Windows NT 4.0**

When you upgrade a computer from Windows NT 4.0 to Windows Server 2003, the security descriptors on NTFS and registry objects must also be upgraded to reflect the new canonical order. This change does not take place during setup. Instead, the DACLs on existing objects are upgraded the first time that their security descriptors are modified after the operating system upgrade. Due to the automatic propagation of inheritable ACEs, the DACLs on all objects in the hierarchy below a modified object are also converted to the new canonical order.

The security descriptors for objects that have been converted contain the security descriptor control flags SE_DACL_AUTO_INHERITED and SE_SACL_AUTO_INHERITED.

The operating system uses the following rules for propagating inherited ACEs to child objects:

- If a child object with no DACL inherits an ACE, the result is a child object with a DACL containing only the inherited ACE.
- If a child object with an empty DACL inherits an ACE, the result is a child object with a DACL containing only the inherited ACE.
- If you remove an inheritable ACE from a parent object, automatic inheritance removes any copies of the ACE that are inherited by child objects.
- If automatic inheritance results in the removal of all ACEs from a child object's DACL, the child object has an empty DACL rather than no DACL.

**No DACL vs. Empty DACL**

A security descriptor with no DACL (also known as a NULL DACL) gives unconditional access to everyone. A security descriptor with an empty DACL gives no access to anyone.

In Windows NT, if a developer wanted to allow everyone unconditional access to an object, his code could create the object with a NULL DACL — that is, without a DACL. In Windows 2000 and later, it is still possible to create objects with a NULL DACL, but code that does this must set the SE_DACL_PROTECTED security descriptor control flag to prevent the object from acquiring a DACL through inheritance. In the worst case, forgetting to protect a NULL DACL can result in an empty DACL if the parent object has no inheritable permissions. The result would be to allow no access to anyone, exactly the opposite of what the object's creator intended.

Best practice is not to use NULL DACLs at all. If you are developing for Windows 2000 or later and want to give everyone unconditional access to an object, create a DACL with one ACE that grants Everyone full control. In Windows Server 2003, the Everyone group no longer includes Anonymous Logon; therefore you might also want to include a second ACE that grants Anonymous Logon full control.

**Related Information**

The following resources contain additional information that is relevant to this section:

- Security Descriptors and Access Control Lists Tools and Settings

# Security Descriptors and Access Control Lists Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Security Descriptors and Access Control Lists Tools
- Related Information


This section contains information about the tools that are associated with security descriptors and access control lists (ACLs).

**Security Descriptors and Access Control Lists Tools**
The following tools are associated with security descriptors and ACLs.

**Cacls.exe: Cacls**
**Category**
This tool ships with the Windows Server 2003.

**Version compatibility**
This tool is compatible with computers running Windows Server 2003.

Cacls is a command-line tool that displays or modifies discretionary access control lists (DACLs) on specified files.For more information about Cacls, see "Cacls" in Help and Support Center in Windows Server 2003.

**RoboCopy.exe: Robust File Copy Utility**
**Category**
This tool is included in the Windows Server 2003 Resource Kit Tools.

**Version compatibility**
This tool is compatible with computers running Windows XP Professional or Windows Server 2003 operating systems.

An administrator can use the Robust File Copy Utility (RoboCopy) to copy files with a fine level of control over which file attributes are copied. For example, you can specify whether to copy files with or without their NTFS file system ACLs, file ownership information, or file auditing information.

This command-line tool simplifies the task of maintaining an identical copy of a folder tree in multiple locations, either on the same computer or in separate network locations. RoboCopy can provide time-efficient maintenance of mirror images of large folder trees on network servers that are separated by slow or unreliable wide area network (WAN) links.By default, RoboCopy ignores source file attributes when selecting files to copy. It copies any file matching specified conditions, regardless of the attribute settings of the file.

For more information about RoboCopy, see "Robocopy.exe" in Resource Kit Tools Help in Tools and Settings Collection.

**Showacls.exe: Show ACLs**
**Category**
This tool is included in the Windows Server 2003 Resource Kit Tools.

**Version compatibility**
This tool is compatible with computers running Windows XP Professional or Windows Server 2003 operating systems.

Show ACLs (ShowACLs) is a command-line tool that enumerates access rights for files, folders, and trees. It allows masking so that you can enumerate only specific ACLs. ShowACLs works on NTFS partitions only.

You can also use ShowACLs to view permissions for a particular user. ShowACLs does this by enumerating the local and global groups to which the particular user belongs and matching the user's security identifier (SID) — and the SIDs of the groups to which the user belongs — to the SIDs in each access control entry (ACE).

For more information about Showacls, see "Showacls.exe" in Resource Kit Tools Help in Tools and Settings Collection.

**Subinacl.exe: SubInACL**
**Category**
This tool is included in the Windows Server 2003 Resource Kit Tools.

**Version compatibility**
This tool is compatible with computers running Windows XP Professional or Windows Server 2003 operating systems.

SubInACL is a command-line tool that enables administrators to obtain security information about files, registry keys, and services and transfer this information from user to user, from local or global group to group, and from domain to domain.

For example, if a user moves from one domain (DomainA) to another domain (DomainB), the administrator can replace DomainA\User with DomainB\User in the security information for the user's files. This gives the user access to the same files from the new domain.

SubInACL enables administrators to do the following:

- Display security information that is associated with files, registry keys, or services. This information includes owner, group, permission ACL, discretionary access control list (DACL), and system access control list (SACL).
- Change the owner of an object.
- Replace the security information for one identifier (account, group, well-known SID) with that of another identifier.

Migrate security information about objects. This is useful if you reorganize a network's domains and need to migrate the security information for files from one domain to another. For more information about Subinacl, see "Subinacl.exe" in Resource Kit Tools Help in Tools and Settings Collection.

**Related Information**
The following resources contain additional information that is relevant to this section:

- "Resource Kit Tools Help" in Tools and Settings Collection for information about RoboCopy.exe
- "Resource Kit Tools Help" in Tools and Settings Collection for information about Showacls.exe
- "Resource Kit Tools Help" in Tools and Settings Collection for information about Subinacl.exe

# Access Tokens Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Access Tokens Technical Reference**
When a user is authenticated, the Local Security Authority (LSA) creates an access token for that user. An access token contains a security identifier (SID) for the user, all of the SIDs for the groups to which the user belongs, and the user's privileges. If you add a user to a group after the user access token has been issued, or you modify privileges assigned to the user account, the user must log off and then log on again before the access token will be updated.

This subject will explain what an access token is in the Microsoft Windows Server 2003 operating system, and how the system uses access tokens to grant access to resources.

**In this subject**


- What Are Access Tokens?
- How Access Tokens Work
- Access Tokens Tools and Settings

# What Are Access Tokens?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Are Access Tokens?**

**In this section**


- Technologies Related to Access Tokens
- Related Information


An access token is a protected object that contains information about the identity and privileges associated with a user account.

When a user logs on interactively or tries to make a network connection to a computer running Windows, the logon process authenticates the user's logon credentials. If authentication is successful, the logon process returns a security identifier (SID) for the user and a list of SIDs for the user's security groups. The Local Security Authority (LSA) on the computer uses this information to create an access token — in this case, the primary access token — that includes the SIDs returned by the logon process as well as a list of privileges assigned by local security policy to the user and to the users security groups.

After LSA creates the primary access token, a copy of the access token is attached to every process and thread that executes on the user's behalf. Whenever a thread or process interacts with a securable object or tries to perform a system task that requires privileges, the operating system checks the access token associated with the thread to determine the level of authorization for the thread.

There are two kinds of access tokens, primary and impersonation. Every process has a primary token that describes the security context of the user account associated with the process. A primary access token is typically assigned to a process to represent the default security information for that process. Impersonation tokens, on the other hand, are usually used for client/server scenarios. Impersonation tokens enable a thread to execute in a security context that differs from the security context of the process that owns the thread.

**Technologies Related to Access Tokens**

Access tokens are closely related to the following technologies.

**Other authorization and access control components**

Access tokens contain account and group SIDs as well as privileges for the account. User rights and permissions to access objects such as Active Directory objects, files, and registry settings are granted to these accounts and groups:

- **Security principals.** Security principals include users or groups with either local or domain accounts, and computer accounts created when a computer running Windows NT, Windows 2000, Windows XP, or a member of the Windows Server 2003 family joins a domain. User rights and permissions to access objects such as Active Directory objects, files, and registry settings are assigned to security principals.

- **Security identifiers (SIDs).** Every account and group is automatically assigned a security identifier (SID) when the account or group is created.

- **Security descriptors and access control lists (ACLs).** A security descriptor is associated with each securable object. If permissions are configured for an object, the object's security descriptor contains a discretionary access control list (DACL), with SIDs for the user's and groups that are allowed or denied access.

- **User rights and permissions.** User rights and permissions to access objects such as Active Directory objects, files, and registry settings are assigned to security principals. When a system task that requires privileges is attempted, the operating system checks the access token associated with the thread or process to determine the level of authorization for that thread or process.

**Authentication**

A user account enables a user to log on to computers and domains with an identity that can be authenticated by the computer or domain.

**Active Directory**
Accounts and groups that are created in an Active Directory domain are stored in the Active Directory database and managed using Active Directory tools.

**Group Policy**
User rights can be assigned to Active Directory groups through Group Policy. Password and Account Lockout policy can be assigned to users through Group Policy.

The diagram below shows the relationship of access tokens to other authorization and access control components.

**Relationship of Access Tokens to Other Authorization and Access Control Components**




**Related Information**
The following resource contains additional information that is relevant to this section.

- Logon and Authentication Technologies

# How Access Tokens Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Access Tokens Work**

**In this section**


- Access Tokens Interfaces
- Access Tokens Physical Structure
- Access Tokens Processes and Interactions
- Related Information


When a user is authenticated, the Local Security Authority (LSA) creates an access token — in this case, a primary access token — for that user. An access token contains a security identifier (SID) for the user, all of the SIDs for the groups to which the user belongs, and the user's privileges. If you add a user to a group after the user's access token has been issued, or modify privileges assigned to the user account, the user must log off and then log on again before the access token will be updated.

Whenever a thread or process interacts with a securable object or tries to perform a system task that requires privileges, the operating system checks the effective access token to determine its level of authorization. If a thread is impersonating, the effective token is usually taken to be the token on the thread. If a thread interacting with the securable object is not impersonating, then the token on the process is examined for the access decision.

Thus, there are two kinds of access tokens, primary and impersonation. Every process has a primary token that describes the security context of the user account associated with the process. A primary access token is typically assigned to a process to represent the default security information for that process. Impersonation access tokens, on the other hand, are usually used for client/server scenarios. Impersonation tokens enable a thread to execute in a security context that differs from the security context of the process that owns the thread.

**Access Tokens Interfaces**

The tables below list APIs used to create, modify, or query access tokens.

**Token Creation APIs**

| API Name | Description |
|---|---|
| CreateRestrictedToken | Creates a new access token that is a restricted version of an existing access token. |
| DuplicateToken | Creates a new impersonation token that duplicates one already in existence. |
| DuplicateTokenEx | Creates a new access token that duplicates an existing token. This function can create either a primary token or an impersonation token. |
| OpenProcessToken | Opens the access token associated with a process. |
| OpenThreadToken | Opens the access token associated with a thread. Tokens with the anonymous impersonation level cannot be opened. |

**Token Modification APIs**

| API Name | Description |
|---|---|
| AdjustTokenGroups | Enables or disables groups already present in the specified access token. Access to TOKEN_ADJUST_GROUPS is required to enable or disable groups in an access token. You cannot enable a group that has the SE_GROUP_USE_FOR_DENY_ONLY attribute. |

| | |
|---|---|
| AdjustTokenPrivileges | Enables or disables existing privileges in the specified access token. Enabling or disabling privileges in an access token requires TOKEN_ADJUST_PRIVILEGES access. |
| ImpersonateAnonymousToken | Enables the specified thread to impersonate the system's anonymous logon token. To ensure that a token matches the operating system's concept of anonymous access, this function should be called before attempting network access to generate an anonymous token on the remote server. Anonymous tokens do not include the Everyone Group SID unless the system default has been overridden. |
| ImpersonateSelf | Obtains an access token that impersonates the security context of the calling process. The token is assigned to the calling thread. For this function to succeed, the DACL protecting the process token must grant the TOKEN_DUPLICATE right to itself. |
| RevertToSelf | Terminates the impersonation of a client application. |
| SetThreadToken | Assigns an impersonation token to a thread. The function can also cause a thread to stop using an impersonation token. |
| SetTokenInformation | Sets various types of information for a specified access token. The information that this function sets replaces existing information. The calling process must have appropriate access rights to set the information. Token-type information can be set only when an access token is created. |

**Token Query APIs**

| API Name | Description |
|---|---|
| CheckTokenMembership | Determines whether a specified SID is enabled in an access token. |
| GetTokenInformation | Retrieves a specified type of information about an access token. The calling process must have appropriate access rights to obtain the information. |
| IsTokenRestricted | Indicates whether a token contains a list of restricted SIDs. |
| PrivilegeCheck | Determines whether a specified set of privileges is enabled in an access token. The PrivilegeCheck function is typically called by a server application to check the privileges of a client's access token. |

For more information about token APIs, search for Authorization Structures in the Authorization reference in the Microsoft SDK documentation on MSDN.

#### Access Tokens Physical Structure

### Access Token




An access token contains a complete description of the security context for a process or thread, including the information in the following table.

**Access Token Components**

| Component | Description |
|---|---|
| User | The SID for the user's account. If the user logs on to an account on the local computer, the user's SID is taken from the account database maintained by the local Security Accounts Manager (SAM). If the user logs on to a domain account, the SID is taken from the Object-SID property of the User object in Active Directory. |
| Groups | A list of SIDs for security groups that include the user. The list also includes SIDs from the SID-History property of the User object representing the user's account in Active Directory. |
| Privileges | A list of privileges held on the local computer by the user and by the user's security groups. |
| Default Owner | The SID for the user or security group who, by default, becomes the owner of any object that the user either creates or takes ownership of. |
| Primary Group | The SID for the user's primary security group. This information is used only by the POSIX subsystem and is ignored by the rest of Windows Server 2003. |
| Default Discretionary Access Control List (DACL) | A built-in set of permissions that the operating system applies to objects created by the user if no other access control information is available. The default DACL grants Full Control to Creator Owner and System. |
| Source | The process that caused the access token to be created, such as Session Manager, LAN Manager, or Remote Procedure Call (RPC) Server. |
| Type | A value indicating whether the access token is a primary or impersonation token. |
| Impersonation Level | A value that indicates to what extent a service can adopt the security context of a client represented by this access token. |
| Statistics | Information about the access token itself. The operating system uses this information internally. |
| Restricting SIDs | An optional list of SIDs added to an access token by a process with authority to create a restricted token. Restricting SIDs can limit a thread's access to a level lower than what the user is allowed. |
| TS Session ID | A value that indicates whether the access token is associated with the Terminal Services client session. |

| | |
|---|---|
| Session Reference | Reserved for internal use. |
| SandBox Inert | Nonzero if the token includes the SANDBOX_INERT flag. |
| Audit Policy | Since Windows Server 2003, used for per user auditing. |
| Origin | Introduced with Windows Server 2003. If the token resulted from a logon using explicit credentials, then the token will contain the ID of the logon session that created it. If the token resulted from network authentication, then this value will be zero. |

**SID Attributes in an Access Token**

Each user and group SID in an access token can have one of two attributes that control how the system uses the SID in an access check. The following table lists the attributes that control access checking.

### SID Attributes

| Attribute | Description |
|---|---|
| SE_GROUP_ENABLED | A SID with this attribute is enabled for access checks. When the system performs an access check, it checks for access control entries (ACEs) that apply to the SID. |
| SE_GROUP_USE_FOR_DENY_ONLY | Since Windows 2000, a SID with this attribute is a deny-only SID. When the system performs an access check, it checks for ACEs that deny access to the SID. It ignores ACEs that allow access for the SID. |

Both attributes are mutually exclusive. If one attribute is set, the other cannot be set. If neither attribute is set, the SID is ignored. Moreover, no process is ever allowed to remove a deny-only attribute from a SID.

**SIDHistory**

SIDHistory is used to store the former SIDs of moved objects such as user's and security groups.

Prior to Windows 2000, restructuring a domain generally meant the creation of new accounts that needed to be placed in the same groups the old accounts were placed in. Since Windows 2000, domain restructuring is made considerably easier as a result of SIDHistory, which is a new attribute of Active Directory security principals.

APIs (or tools and support utilities built from these APIs) update the SIDHistory attribute of the object, representing it in Active Directory with its former SID as part of the move operation for the user or group. When the user then logs on to the system, not only the new SID, but also the old SID retrieved from the SIDHistory attribute, are added to the user's access token and used to determine the user's group memberships. The SIDs of the groups of which the user is a member through either the new SID or the old SID are then also added to the access token, together with any SIDHistory those groups might have.

Because groups can be moved, they can also have SIDHistory. The system also retrieves the SIDHistory attributes of all the groups of which the user is a member and then adds these attributes to the user's access token.

Because these SIDHistory entries in the token look to the system like normal group memberships during authorization checks, the token can grant appropriate access even on earlier versions of the system.

### Note

- The SIDHistory attribute can only be updated in native mode Windows 2000 or Windows 2003 domains, which has the effect of requiring all migration operations relying on SIDHistory to have a native mode target for restructure.

**Access Tokens Processes and Interactions**

Clients never get direct access to network resources; instead, network services build a client's access token and then impersonate the client to execute the requested operations. The Windows operating system kernel uses security identifiers in the access token to verify whether the user is authorized the desired access rights to the target object.

There are two kinds of access tokens, primary and impersonation.

**Primary Tokens**

Every process has a primary token that describes the security context of the user account associated with the process. The primary token may be assigned to a process to represent the default security information for that process.

**Creating a Primary Token**

The access token is built from the SID list acquired through authentication. Group membership can be evaluated non-transitively (as was done with Windows NT) or transitively. Since Windows 2000, chaining of groups has been possible in native domains. This allows a member of group A to be a member of group B if group A is a member of group B.

**Creating the SID List for User Account to Access System**




**Creating User SID Lists**

The following example is for a user logging onto a system. This occurs during authentication either during pass thru authentication for NTLM or when creating the ticket granting ticket and service ticket for Kerberos authentication. The user's domain controller will do two or more of the following:

1. Authenticate the user.

2. Only if the user's domain is in native mode, retrieve the user's SID history and add it to the SID list.

3. Retrieve the user's global group membership.

    a. If the user's domain is in mixed mode, retrieve global group membership of the user's SID and add these to the SID list.

    b. If the user's domain is in native mode, **transitively** retrieve the global group membership of both the user's SID and the user's SID history and add these to the SID list.

4. If the user's domain is in native mode, pass the list of SIDs to the global catalog.

    a. The global catalog will **transitively** retrieve the universal group membership for each SID in the list and add these to the SID list.

    b. Return the list back to the domain controller.

5. If the system's domain is in native mode, **transitively** retrieve the domain local group membership for each SID in the list and add these to the SID list.

6. Return the list to the system.

**Building the access token**

The system places the SIDs received from the domain controller and all the built-in groups for each SID into the access token. These built-in groups will include the following:

**Logon type**

- For specific logon types, a specific SID is added for that logon type, such as INTERACTIVE, or BATCH, for example.
- If using NTLM logon, the system adds a specific SID.

**Implicit Membership**

- In the case of the token for a "NT AUTHORITY\SYSTEM" process, the SID for "Administrators" is added to the token. Therefore SYSTEM implicitly belongs to "Administrators."
- The SID "Everyone" is added to every token, except for Anonymous tokens when the security policy option "Let Everyone permissions apply to Anonymous user's" is disabled.
- "Authenticated Users" is added to non-anonymous tokens.
- There is also a "NONE" group, which is added to local tokens on non-domain controllers; "NONE" is the SID for the primary group for local accounts.

**Summary of Conditions Required for SIDs to Be in SID List**

| **SID Type** | **User's Domain** | | **System's Domain** | |
|---|---|---|---|---|
| | **Mixed** | **Native** | **Mixed** | **Native** |
| SID History | No | Yes | No | Yes |
| Built-in | No | No | Yes, system's built-in | Yes, system's built-in |
| Domain Local | No | No | Yes, if system is DC in User's domain | Yes, transitive |
| Global | Yes, not transitive | Yes, transitive | | |
| Universal | No | Yes, transitive | No | No |

**Impersonation Tokens**

Impersonation tokens are access tokens that have been created to capture the security information of a client process, thus enabling a server to impersonate the client process in security operations.

Impersonation is the ability of a thread to execute in a security context that is different from the context of the process that owns the thread. Impersonation is designed to meet the security requirements of client/server applications. When running in a client's security context, a service "is" the client, to some degree. One of the service's threads uses an access token representing the client's credentials to obtain access to the objects to which the client has access.

The primary reason for impersonation is to cause access checks to be performed against the client's identity. Using the client's identity for access checks can cause access to be either restricted or expanded, depending on what the client has permission to do. For example, suppose a file server has files containing confidential information, and each of these files is protected by a DACL. To prevent a client from obtaining unauthorized access to information in these files, the service can impersonate the client before accessing the files.

Every process has a primary token that describes the security context of the user account associated with the process. In an ordinary application process, the user associated with the process is the human user who started the application. In a service process, however, services run under their own accounts and have identities in their own right.

System services installed with the operating system run under the Local System, local service, or Network Service account. Other services can be configured to run three different account types: Local System account, service accounts on the local system, or service accounts in Active Directory.

The primary token for a service is associated with the control thread in the service process and with all other threads acting on behalf of the service. It identifies the service's account, group memberships, and privileges. The service uses this information during access checks when the service requests access to objects needed for its own housekeeping.

Then, when a service accepts a client, it creates a thread to do the work and associates the client's access token with the worker

thread. Once associated in this way, the client's access token is an impersonation token that identifies the client, the client's groups, and the client's privileges. The impersonation token's information is used during access checks when the thread requests access to resources on the client's behalf. When impersonation is over, the thread reverts to using the primary token and returns to operating in the service's security context instead of the client's.

**The Mechanism for Impersonation**




**Impersonation Levels**

A client process can control to what extent a service is able to act as the client by selecting an impersonation level when it connects to the service. By choosing an impersonation level, the client tells the service to what extent it can impersonate the client. In other words, the impersonation levels described in the table below define to what extent the client agrees to let the service "be" the client.

Users cannot select an impersonation level. One is specified as Security Quality of Service (SQoS) information in the code for client/server applications. There are four impersonation levels: anonymous, identify, impersonate, and delegate. The anonymous level has never been supported. Prior to Windows 2000, only the identify and impersonate levels were supported. Windows 2000 added support for the delegate level.

**Impersonation Levels**

| Impersonation Level | Description |
|---|---|
| Anonymous | The client is anonymous to the service. The service can impersonate the client but the impersonation token does not contain any information about the client. |
| Identify | The service can get the identity of the client and can use this information in its own security mechanism, but it cannot impersonate the client. |
| Impersonate | The service can impersonate the client. If the service is on the same computer as the client process, it can access network resources as the client. If the service is on a remote computer, it can impersonate the client only when accessing resources on the service's computer. |
| Delegate | The service can impersonate the client not only when it accesses resources on the service's computer but also when it accesses resources on other computers. This level is supported only in Windows 2000 and later versions of the Windows operating system. |

**Configuring Clients and Services for Delegation**

Impersonation works at the delegate level only under the following conditions:

- The computers hosting the client process, the service process, and processes for any backend services must all be running Windows 2000 or later in a Windows 2000 or later domain. This is because delegation requires the Kerberos authentication protocol.

- The client's user account must be enabled for delegation.

- The service's account must be enabled for delegation.

The way you configure the service account depends on whether the service runs under the Local System account on a computer or under its own domain user account.

- If the service is configured to run under the Local System account, the computer where the service runs must be trusted for delegation. (Note that when you trust a computer for delegation, you enable delegation for all services that run under the Local System account on the computer. If an unwary administrator installs an untrusted service on the computer and configures it to run as Local System, it too can access network resources while impersonating other user's. A better practice is to configure services that use delegation to run under their own domain user accounts managed by domain administrators.)

- If the service is configured to run under its own domain user account, the service's user account must be enabled to act as a delegate.

**Restricted Tokens**
A restricted token is a primary or impersonation access token that has been modified by the CreateRestrictedToken function. A process or impersonating thread running in the security context of a restricted token is restricted in its ability to access securable objects or perform privileged operations. In Windows 2000 and later, an application can start a child process in a restricted security context so that code executing in the child process has more limited access to resources than the application's user has. For example, when you use Internet-browsing software to display a Web page in an untrusted security zone, code associated with the Web page can execute with fewer privileges on your computer than you have. (This feature is not used by Microsoft Internet Explorer version 5.0 and earlier.)

Applications can create restricted security contexts for child processes and impersonation threads by creating a restricted token and assigning it to the child process or the impersonating thread. Restricted tokens are created by removing privileges, by applying the deny-only attribute to SIDs, or by adding a list of restricting SIDs to the original access token.

When a restricted process or thread tries to access a securable object, the system performs two access checks by:

- Reading the token's enabled and deny-only SIDs.

- Reading the list of restricting SIDs.

Access is granted only if both access checks allow the requested access rights.

For more information about restricted tokens, see Access Control in the Microsoft Platform SDK on MSDN.

**Related Information**
The following resource contains additional information that is relevant to this subject.

- Logon and Authentication Technologies

# Access Tokens Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Access Tokens Tools and Settings**

**In this section**


- Access Tokens Tools
- Access Tokens Registry Entries
- Access Tokens Group Policy Settings
- Access Tokens WMI Classes


**Access Tokens Tools**

The following tools are associated with access tokens.

**Dsa.msc: Active Directory Users and Computers**

**Category**

Active Directory Users and Computers is a Microsoft Management Console (MMC) snap-in that is installed automatically when you install Active Directory. This tool also ships with the Administration Tools Pack (Adminpak.msi).

You can access the tool from the **Start** menu: Click **Start**, then click **Programs**,then click **Administrative Tools**,and then click **Active Directory Users and Computers**.

**Version Compatibility**

Active Directory Users and Computers runs on domain controllers that are running Windows Server 2003 and Windows 2000. In both of these server systems, MMC provides a window in the user interface where you can add, configure, and control items. Active Directory Users and Computers is the MMC snap-in that you can use to administer and publish information in the directory.

The Windows Server 2003 version of Active Directory Users and Computers can target domain controllers that are running Windows Server 2003 or Windows 2000.

On administrative workstations that are running Windows XP Professional or Windows 2000, you can install the Windows Server 2003 Administration Tools Pack (Adminpak.msi) from the i386 directory on the Windows Server 2003 CD. This version of the Administration Tools Pack encrypts and signs LDAP traffic between the administrative tool clients and domain controllers.

**Note**

- You cannot run the Microsoft Windows Server 2003 Administration Tools Pack (Adminpak.msi) on a computer that is running Windows XP Professional, Windows XP Home Edition, or Windows XP 64-Bit Edition Version 2003 without Windows XP Service Pack 1 (SP1).

You can manage the following objects and their associated properties with this tool, which in turn will affect access tokens created for these objects.

**Active Directory Users and Computers Object Management**

| Object Type | Changes That Affect Access Tokens |
|---|---|
| User objects | Configure user or service account for delegation. |
| Computer objects | Configure computer account for delegation. This will affect services running under the local System account. |

To find more information about Active Directory Users and Computers, see Windows Server 2003 Tools Help in the Tools and Settings Collection.

**Ntrights.exe: Ntrights**
**Category**
Ntrights is included in the Windows Server 2003 Resource Kit and the Windows 2000 Resource Kit.

**Version compatibility**
Ntrights is supported for Windows Server 2003, Windows XP Professional, and Windows 2000.

Ntrights is a command-line tool that enables you to assign or revoke a right for a user or group of users on a local or remote computer. You can also place an entry that notes the change in the event log of the computer.

Ntrights is useful in unattended or automated installations during which you might want to change the default rights. You can also use the tool in situations where you need to change a right in an existing installation, but you cannot access and log on to all computers.

To find more information about Ntrights, see Windows Server 2003 Resource Kit Tools Help in the Tools and Settings Collection.

**Services.msc: Services**
**Category**
Services is a Microsoft Management Console (MMC) snap-in that is automatically installed when you install Windows Server 2003. There are two ways to access the tool:

- Click **Start**, then click **Programs**,then click **Administrative Tools**,and then click **Services**.

Or,

- Right-click **My Computer**, and then click **Manage**.

**Version compatibility**
Services runs on systems that are running Windows Server 2003, Windows XP, and Windows 2000.

Services can configure the security context of services that impact access tokens. Services affects access tokens when you use it to:

- Manage the services on your computer.
- Set up recovery actions to take place if a service fails.
- Create custom names and descriptions for services so you can easily identify them.

To find more information about Services, see Windows Server 2003 Tools Help in the Tools and Settings Collection.

**Showpriv.exe: Show Privilege**
**Category**
Show Privilege is included in the Windows Server 2003 Resource Kit and the Windows 2000 Resource Kit.

**Version Compatibility**
Show Privilege is supported for Windows Server 2003, Windows XP Professional, and Windows 2000.

Show Privilege is a command-line tool that displays the rights assigned to users and groups. The tool must be run locally on the target computer. To display users and groups that have domain privileges, Show Privilege must be run on a domain controller. The following table shows the privileges specific to access tokens.

### Access Token Privileges

| Privilege Name | Equivalent Security Policy User Right Setting | Description |
|---|---|---|
| SeCreateTokenPrivilege | Create a token object | Allows a process to create an access token. |
| SeAssignPrimaryTokenPrivilege | Replace a process-level token | Allows a process that has this privilege to replace the access token associated with a process. |

| | | |
|---|---|---|
| SeImpersonatePrivilege | Impersonate a client after authentication | Allows a process to impersonate. |

To find more information about Show Privilege, see Windows Server 2003 Resource Kit Tools Help in the Tools and Settings Collection.

**Whoami.exe: Whoami**
**Category**
Whoami is a command-line tool included in the Windows Server 2003 family. The tool is also included in the Windows 2000 Resource Kit.

**Version Compatibility**
Whoami is supported for Windows Server 2003 and Windows 2000.

You can use this command-line tool to display the complete contents of the access token in the command window. For the current user's security context, Whoami can display, for example:

- User name and security identifier (SID)
- Groups and their SIDs
- Privileges and their status (for example, enabled or disabled)
- Logon ID

To find more information about Services, see Windows Server 2003 Tools Help in the Tools and Settings Collection.

**Access Tokens Registry Entries**
The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, instead of editing the registry directly, use Group Policy or other Windows tools such as MMC to accomplish tasks. If you must edit the registry, use extreme caution.

The following registry settings that affect access tokens cannot be modified by using Group Policy or other Windows tools.

**EveryoneIncludesAnonymous**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

**Version compatibility**
**EveryoneIncludesAnonymous** is supported for Windows Server?2003 and Windows XP.

This registry setting controls whether the Everyone SID is included in the access token generated for an anonymous user.

### EveryoneIncludesAnonymous Settings

| Setting | Effect |
|---|---|
| 0 | (default) Do not include the Everyone SID in the access token generated for an anonymous user. |
| 1 | Include the Everyone SID in the access token generated for an anonymous user. |

**RestrictAnonymous**
**Registry path**
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\

**Version compatibility**
**RestrictAnonymous** is supported for Windows Server 2003, Windows XP, and Windows 2000.

This registry setting restricts anonymous users from displaying lists of users, and from viewing security permissions.

### RestrictAnonymous Settings

| Setting | Effect |
|---|---|

| | |
|---|---|
| 0 | (default) Anonymous users are not restricted. Rely on default permissions. |
| 1 | Do not allow enumeration of Security Accounts Manager (SAM) accounts and shares. |
| 2 | In Windows 2000, do not include the Everyone SID in the access token generated for an Anonymous user.<br><br>Not supported in Windows Server 2003. |

**Results of Anonymous User Settings**

**Anonymous User: Windows 2000**

| Restrict Anonymous Setting | Can Enumerate Local SAM Accounts and Shares? | Can Access Other Securable Objects If: |
|---|---|---|
| 0 | Yes | Anonymous or Everyone is granted access by the object's access control list (ACL). |
| 1 | No | Anonymous or Everyone is granted access by the object's ACL. |
| 2 | No | Anonymous is explicitly granted access by the object's ACL. |

**Anonymous User: Windows Server 2003 and Windows XP**

| Restrict Anonymous Setting | EeveryoneIncludesAnonymous Setting | Can Enumerate Local SAM Accounts and Shares? | Can Access Other Securable Objects If: |
|---|---|---|---|
| 0 | 0 | Yes | Anonymous is explicitly granted access by the object's ACL. |
| 0 | 1 | Yes | Anonymous or Everyone is granted access by the object's ACL. |
| 1 | 0 | No | Anonymous is explicitly granted access by the object's ACL. |
| 1 | 1 | No | Anonymous or Everyone is granted access by the object's ACL. |

**Effects of Anonymous User Settings Entered in a Domain Controller's Registry**

**Ability of anonymous users to enumerate account information**

There is no local SAM on a domain controller. Thus, **RestrictAnonymous** does not control the ability of anonymous users to enumerate account information. Instead, access to account information is controlled by ACLs on account objects in Active Directory.

**Ability of anonymous users to enumerate shared resources**

Anonymous users will not be able to enumerate shared resources or pipes if **RestrictAnonymous** is set to equal 1.

**Ability of Anonymous Users to Access Active Directory Data on Windows 2000 Domain Controllers**

| Restrict Anonymous Setting | Pre-Windows 2000 Compatible Access Security Group Membership | Access to Any Active Directory Data |
|---|---|---|
| 0 or 1 | No | No |
| 0 or 1 | Yes | Yes, if Everyone is a member of this group. |
| 2 | No | No |
| 2 | Yes | No |
| 2 | Yes, Anonymous must be explicitly a member. | Yes |

**Ability of Anonymous Users to Access Active Directory Data on Windows Server 2003 Domain Controllers**

| EeveryoneIncludesAnonymous Setting | Pre-Windows 2000 Compatible Access Security Group Membership | Access to Any Active Directory Data |
|---|---|---|
| 0 | No | No |
| 0 | Yes | Yes, if Anonymous is also a member of this group. |
| 1 | Yes | Yes, even if Anonymous is not a member of this group as long as Everyone is a member of this group. |

**Note**

- Both Everyone and Anonymous are members of Pre-Windows 2000 Compatible Access group by default in Windows Server 2003.

#### Access Tokens Group Policy Settings

The following table lists and describes the Group Policy settings that are associated with access tokens.

**Group Policy Settings Associated with Access Tokens**

| Group Policy Setting | Description |
|---|---|
| User Rights Assignment:<br>• Create a token object<br>• Replace a process level token | Changes to these settings control:<br>• Calling APIs to create tokens.<br>• Whether a process can replace a token. |
| Audit Policy:<br>• Audit policy change<br>• Audit privilege use<br>• Audit process tracking | Changes to this setting will:<br>• Generate audits when rights are assigned with one of the tools discussed earlier.<br>• Enable audit privilege use. Will log when SeAssignPrimaryTokenPrivilege was used.<br>• Create an audit for assigning a primary token that contains the two processes involved and the identity of the token assigned. |
| Security Options:<br>• Network access: Let Everyone permissions apply to anonymous users | Changes to this setting will affect whether Everyone is in the token for anonymous users. |

For more information about these Group Policy settings, see Account Policy Settings.

#### Access Tokens WMI Classes

The following table lists and describes the Windows Management Information (WMI) classes that are associated with access tokens. These WMI classes are shipped with Windows Server 2003.

**WMI Classes Associated with Access Tokens**

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Win32_TokenGroups | \root\cimv2 | Windows Server 2003<br>Windows XP |

| | | |
|---|---|---|
| Win32_TokenPrivileges | \root\cimv2 | Windows Server 2003<br><br>Windows XP |

For more information about these WMI classes, see the WMI reference in the SDK documentation on MSDN.

# Permissions Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Permissions Technical Reference**

**In this subject**


- What Are Permissions?
- How Permissions Work
- Permissions Tools and Settings


Permissions are a key component of the Microsoft Windows Server 2003 security architecture that you can use to manage the process of authorizing users, groups, and computers to access objects on a network.

Permissions enable the owner of each securable object, such as a file, Active Directory directory service object, or registry key, to control who can perform an operation or a set of operations on the object or object property. Because access to an object is at the discretion of the object's owner, the type of access control that is used in Windows Server 2003 is called discretionary access control.

# What Are Permissions?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Are Permissions?**

**In this section**


- Technologies Related to Permissions
- Related Information


Permissions are a key component of the Windows Server 2003 security architecture that you can use to manage the process of authorizing users, groups, and computers to access objects on a network.

Permissions enable the owner of each securable object, such as a file, Active Directory object, or registry key, to control who can perform an operation or a set of operations on the object or object property. Because access to an object is at the discretion of the object's owner, the type of access control that is used in Windows Server 2003 is called discretionary access control.

Permissions are different from user rights in that permissions are attached to objects and user rights apply to user accounts. Administrators use user rights (also known as privileges) to assign specific privileges and logon rights to groups or users. These rights authorize users to perform specific actions, such as logging on to a system interactively or backing up files and directories.

On computers, user rights enable administrators to control who has the authority to perform operations that affect an entire computer, rather than a particular object. Administrators assign user rights to individual users or groups as part of the security settings for the computer. Although user rights can be managed centrally through Group Policy, they are applied locally. Users can (and usually do) have different user rights on different computers.

Permissions are expressed in the security architecture as access control entries (ACEs). The following figure shows the relationship of permissions to other key components of the authorization and access control model.

**The Role of Permissions (as ACEs) in the Authorization and Access Control Model**




**Technologies Related to Permissions**

Permissions are closely related to the following technologies and security components.

**File systems (FAT and NTFS)**

The file system determines how files on a volume are named, stored, and organized. A file system manages files and folders and the information that is needed by local and remote users to locate and access these items.

**Security principals**

A security principal can be a user or computer account or a group of these accounts — that is, any entity that the security system recognizes. User accounts can be used by human users as well as by autonomous processes. Permissions to access

objects such as Active Directory objects, files, and registry settings are granted to security principals.

**Access tokens**
After a user logs on and is authenticated, the system creates an access token for the user that contains the security identifier (SID) of the user and the SIDs of all the domain groups that the user is a member of. Every process that the user creates contains the user's access token, which is then used to determine whether to grant that user access to a system resource in a given logon session.

**Security descriptors and access control lists**
If permissions are configured for an object, the object's security descriptor contains a discretionary access control list (DACL) with SIDs for the users and groups who are allowed or denied access.

**Related Information**
The following resources contain additional information that is relevant to this section:

- Security Principals Technical Reference
- Access Tokens Technical Reference
- Security Descriptors and Access Control Lists Technical Reference

# How Permissions Work

Updated: June 21, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

**How Permissions Work**

**In this section**


- Permissions
- Conflicts Between User Rights and Permissions
- Related Information


Permissions are a key component of the Windows Server 2003 security architecture that you can use to manage the process of authorizing users, groups, and computers to access objects on a network.

Permissions enable the owner of each secured object, such as a file, Active Directory object, or registry key, to control who can perform an operation or a set of operations on the object or object property. Because access to an object is at the owner's discretion, the type of access control that is used in Windows Server 2003 is called discretionary access control. An owner of an object always has the ability to read and change permissions on the object.

Permissions are applied to secured objects, such as files and folders, Active Directory objects, services, or registry objects. Permissions can be granted to a user, group, or computer. You can assign permissions on objects to the following:

- Groups, users, and special identities in the domain
- Groups and users in the domain and any trusted domains
- Local groups and users on the computer where the object resides

The permissions that are attached to an object depend on the type of object. For example, the permissions that can be attached to a file are different from the permissions that can be attached to a registry key.

**Optimal Environment**

The following sections provide general guidelines for how permissions should be applied to files and folders and Active Directory objects in an optimal environment.

**Permissions on Files and Folders**

- Shared folders are stored on NTFS file system volumes. You can set permissions at the file level only if the files are stored on an NTFS volume. On a FAT or FAT32 volume, you can set permissions for shared folders but not for files and folders within a shared folder. Moreover, share permissions on a FAT or FAT32 volume restrict network access only, not access by users working directly on the computer.
- Security templates are used wherever possible, rather than individual permissions.
- Default permission entries are used on file system objects, particularly on system folders and root folders. Changing default permissions can cause unexpected access problems or reduce security.

**Permissions on Active Directory Objects**

- The same set of permissions is assigned to multiple objects wherever possible. Access control lists (ACLs) in Windows Server 2003 feature single instancing: If multiple objects have identical ACLs, Active Directory stores only one instance of the ACL.
- Default permissions on Active Directory objects are left unchanged. Changing default permissions can cause unexpected

access problems or reduce security.

- The number of access control entries (ACEs) that apply to child objects is kept to a minimum. When you use the **Apply Onto** option to control inheritance, be aware that not only will the specified objects inherit that ACE, but all child objects will also receive a copy of that ACE. If there are enough objects that will get copies of this ACE, the increased amount of data can cause performance problems on your network.

**Permissions**

A permission is authorization to perform an operation on a specific object, such as a file. Permissions can be granted by owners and by anyone with the permission to grant permissions, which normally includes administrators on the system.

If you own an object, you can grant any user or security group any permission on that object, including the permission to take ownership.

When permission to perform an operation is not explicitly granted, it is implicitly denied. For example, if Alice allows the Marketing group, and only the Marketing group, permission to read her file, users who are not members of the Marketing group are implicitly denied access.

Permissions can also be explicitly denied. For example, Alice might not want Bob to be able to read her file, even though he is a member of the Marketing group. She can exclude Bob by explicitly denying him permission to read the file. This is normally how explicit denies are used — to exclude a subset (such as Bob) from a larger group (such as Marketing) that has been given permission to perform an operation. Note that use of explicit denials, while possible, increase the complexity of the authorization policy and can create unexpected errors. For example, you might want to allow domain administrators to perform an action but deny domain users. If you attempt to implement this by explicitly denying domain users, you also deny any domain administrators who are also domain users. Though it is sometimes necessary, you can and should avoid the use of explicit denies in most cases.

Each permission that an object's owner grants to a particular user or group is stored as an ACE in a discretionary access control list (DACL) that is part of the object's security descriptor.

**Permissions and Security Descriptors**

Every container and object on the network has a set of access control information attached to it that is called a security descriptor. This information controls the type of access that is allowed to users and groups. The security descriptor is automatically created along with the container or object that is created. For example, each file has a security descriptor.

An object's security descriptor contains its permissions. Permissions are associated with, or assigned to, specific users and groups. For example, for a file Temp.dat, the Administrator group might be assigned Read, Write, and Delete permissions, while the Operator group might be assigned Read and Write permissions only.

Each assignment of permissions to a user or group is called a permission entry, which is a type of ACE. The entire set of permission entries in a security descriptor is known as a permission set or ACL. Therefore, for the file named Temp.dat, the permission set includes two permission entries, one for the Administrator group and one for the Operator group.

**Ownership**

Every object has an owner, whether in an NTFS volume or in Active Directory. The owner controls how permissions are set on the object and to whom permissions are granted.

By default, in Windows Server 2003, the owner is the creator of the object. If the creator of an object is a member of the Administrators group, the Administrators group is the owner. The owner can always change permissions on an object, even when denied all access to the object.

Ownership can be taken by the following:

- Any user with the **Take ownership of files or other objects** user right, which can be granted to any user. By default, the Administrators group is given the **Take ownership of files or other objects** user right.

  An administrator who wants to repair or change permissions on a file must begin by taking ownership of the file.

- Anyone or any group who has the **Take ownership** permission on the object in question

- A user who has the **Restore files and directories** user right

Ownership can be transferred in the following ways:

- The current owner can grant the **Take ownership** permission to another user, allowing that user to take ownership at any time. The user must actually take ownership to complete the transfer.
- An administrator can take ownership.
- A user who has the **Restore files and directories** user right can assign ownership to any user or group.

#### Explicit vs. Inherited Permissions

There are two types of permissions, explicit permissions and inherited permissions:

- Explicit permissions are those that are set by default when the object is created or by user action.
- Inherited permissions are those that are propagated to a child object from a parent object. Inherited permissions ease the task of managing permissions and ensure consistency of permissions among all objects in a given container.

By default, objects that are created in a container inherit the permissions from that container when the objects are created. For example, when you create a folder called MyFolder, all subfolders and files that are created in MyFolder automatically inherit the permissions from MyFolder. Therefore, MyFolder has explicit permissions, while all subfolders and files in it have inherited permissions.

### Notes

- Inherited Deny permissions do not prevent access to an object if the object has an explicit Allow permission entry.
- Explicit permissions take precedence over inherited permissions, even inherited Deny permissions.

#### Access Masks

In an ACE, permissions are represented by one or more bits in a 32-bit value called an access mask. When a thread requests access to an object, it specifies the type of access that it desires by using an access mask as well. During an access check, the operating system compares the desired access mask that is supplied by the thread with the access mask in each ACE of the object's DACL, looking for bits that match. The following figure illustrates the layout of an access mask.

### Layout of an Access Mask

| 31 | 30 | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| GR | GW | GE | GA | Reserved | | | | AS | Standard Access Rights | | | | | | | Object-Specific Access Rights | | | | | | | | | | | | | | | |

Key
GR Generic Read
GW Generic Write
GE Generic Execute
GA Generic All
AS Right to access SACL

Each bit corresponds to an access right — a particular operation or set of operations that can be performed on the object. Turning a bit on in a desired access mask signals that the thread requests the right to perform the corresponding operation. Turning a bit on in the access mask for an ACE signals that the corresponding operation is either allowed or denied, depending on the type of ACE.

Although many of the access rights in an access mask correspond to permissions that can be set through the user interface, not all of them do. Some can be set only by a process, such as the process that creates an object, and they cannot be manipulated through the user interface. Windows Server 2003 uses an access mask format that includes the following types of access rights.

#### Generic access rights

These rights apply to all objects, but they mean different things for different object types. Each object type maps each generic access right to a different set of standard and object-specific access rights. The following table lists generic access rights.

### Generic Access Rights

| Constant in Win32 API | Meaning |
|---|---|

| | |
|---|---|
| GENERIC_ALL | Read, write, and execute access |
| GENERIC_EXECUTE | Execute access |
| GENERIC_READ | Read access |
| GENERIC_WRITE | Write access |

**Standard access rights**
These rights are more specific than generic access rights, but they still apply to operations that are available for most object types. The following table lists standard access rights.

**Standard Access Rights**

| Constant in Win32 API | Meaning |
|---|---|
| DELETE | The right to delete the object |
| READ_CONTROL | The right to read the information in the object's security descriptor, not including the information in the system access control list (SACL) |
| SYNCHRONIZE | The right to use the object for synchronization. Some object types do not support this access right. |
| WRITE_DAC | The right to modify the DACL in the object's security descriptor |
| WRITE_OWNER | The right to change the owner in the object's security descriptor |

**SACL access right**
This right controls the ability to read or change an object's SACL, which controls auditing for the object. The operating system allows access to the SACL only if the subject's access token includes the **Manage auditing and security log** (SeSecurityPrivilege) user right.

**Object-specific access rights**
Each type of securable object defines its own set of access rights. For a list of the specific access rights that apply to each type of object, see the Microsoft Platform SDK link on MSDN.

**Permissions for Files and Folders**
Folder permissions include Full Control, Modify, Read & Execute, List Folder Contents, Read, and Write. Each of these permissions consists of a logical group of special permissions that are listed and defined in the following table.

**Permissions for Files and Folders**

| Permission | Description |
|---|---|
| Traverse Folder/ Execute File | For folders: Traverse Folder allows or denies moving through folders to reach other files or folders, even if the user has no permissions on the traversed folders. (Applies to folders only.) Traverse folder takes effect only when the group or user is not granted the **Bypass traverse checking** user right in the Group Policy snap-in. (By default, the Everyone group is given the **Bypass traverse checking** user right.)<br><br>For files: Execute File allows or denies running program files. (Applies to files only.)<br><br>Setting the Traverse Folder permission on a folder does not automatically set the Execute File permission on all files in that folder. |
| List Folder/ Read Data | List Folder allows or denies viewing file names and subfolder names in the folder. List Folder only affects the contents of that folder; it does not affect whether the folder you are setting the permission on will be listed. (Applies to folders only.)<br><br>Read Data allows or denies viewing data in files. (Applies to files only.) |
| Read Attributes | Allows or denies viewing the attributes of a file or folder, such as read-only and hidden. Attributes are defined by NTFS. |

| | |
|---|---|
| Read Extended Attributes | Allows or denies viewing the extended attributes of a file or folder. Extended attributes are defined by programs and may vary by program. |
| Create Files/ Write Data | Create Files allows or denies creating files in the folder. (Applies to folders only.)<br><br>Write Data allows or denies making changes to the file and overwriting existing content. (Applies to files only.) |
| Create Folders/ Append Data | Create Folders allows or denies creating folders in the folder. (Applies to folders only.)<br><br>Append Data allows or denies making changes to the end of the file but not changing, deleting, or overwriting existing data. (Applies to files only.) |
| Write Attributes | Allows or denies changing the attributes of a file or folder, such as read-only or hidden. Attributes are defined by NTFS.<br><br>The Write Attributes permission does not imply creating or deleting files or folders, it only includes the permission to make changes to the attributes of a file or folder. To allow (or deny) create or delete operations, see Create Files/Write Data, Create Folders/Append Data, Delete Subfolders and Files, and Delete. |
| Write Extended Attributes | Allows or denies changing the extended attributes of a file or folder. Extended attributes are defined by programs and may vary by program.<br><br>The Write Extended Attributes permission does not imply creating or deleting files or folders, it only includes the permission to make changes to the attributes of a file or folder. To allow (or deny) create or delete operations, see Create Files/Write Data, Create Folders/Append Data, Delete Subfolders and Files, and Delete. |
| Delete Subfolders and Files | Allows or denies deleting subfolders and files, even if the Delete permission has not been granted on the subfolder or file. (Applies to folders.) |
| Delete | Allows or denies deleting the file or folder. If you do not have Delete permission on a file or folder, you can still delete it if you have been granted Delete Subfolders and Files on the parent folder. |
| Read Permissions | Allows or denies reading permissions of the file or folder, such as Full Control, Read, and Write. |
| Change Permissions | Allows or denies changing permissions of the file or folder, such as Full Control, Read, and Write. |
| Take Ownership | Allows or denies taking ownership of the file or folder. The owner of a file or folder can always change permissions on it, regardless of any existing permissions that protect the file or folder. |
| Synchronize | Allows or denies different threads to wait on the handle for the file or folder and synchronize with another thread that may signal it. This permission applies only to multithreaded, multiprocess programs. |

You cannot access an encrypted file without the Encrypting File System (EFS) key, even if you have the necessary permissions.

The following table lists the special permissions and their corresponding file and folder permissions.

**Special Permissions**

| Special Permissions | Full Control | Modify | Read & Execute | List Folder Contents (folders only) | Read | Write |
|---|---|---|---|---|---|---|
| Traverse Folder/Execute File | ● | ● | ● | ● | | |
| List Folder/Read Data | ● | ● | ● | ● | ● | |
| Read Attributes | ● | ● | ● | ● | ● | |
| Read Extended Attributes | ● | ● | ● | ● | ● | |
| Create Files/Write Data | ● | ● | | | | ● |

| | | | | | | |
|---|---|---|---|---|---|---|
| Create Folders/Append Data | ● | ● | | | | ● |
| Write Attributes | ● | ● | | | | ● |
| Write Extended Attributes | ● | ● | | | | ● |
| Delete Subfolders and Files | ● | | | | | |
| Delete | ● | ● | | | | |
| Read Permissions | ● | ● | ● | ● | ● | ● |
| Change Permissions | ● | | | | | |
| Take Ownership | ● | | | | | |
| Synchronize | ● | ● | ● | ● | ● | ● |

You should also be aware of the following:

- Groups or users that are granted Full Control on a folder can delete any files in that folder, regardless of the permissions protecting the file.

- Although List Folder Contents and Read & Execute appear to have the same special permissions, these permissions are inherited differently. List Folder Contents is inherited by folders but not files, and it should only appear when you view folder permissions. Read & Execute is inherited by both files and folders, and it is always present when you view file or folder permissions.

- In Windows Server 2003, by default, Everyone does not include Anonymous Users; therefore, permissions that are applied to Everyone do not affect Anonymous Users.

**Permissions on a Shared Folder**

Access on a shared folder is determined through two sets of permission entries; the permissions that are set on the share (called share permissions) and the permissions that are set on the folder (called NTFS file and folder permissions). Share permissions are often used for managing computers with FAT32 file systems or for managing other computers that do not use the NTFS file system.

Share permissions and NTFS permissions are independent in the sense that neither changes the other. The final access permissions on a shared folder are determined by taking into consideration both the share permission and the NTFS permission entries. The more restrictive permissions are then applied.

The following table suggests permissions that a security-conscious administrator can grant to the Users group for certain shared folder types. The recommended permissions have been tested and they work correctly, but there are alternative approaches. For example, some experienced administrators prefer to always set share permissions to Full Control for Everyone and to rely entirely on NTFS permissions to restrict access.

**Permissions for File Servers**

| **Folder Type** | **Share Permissions** | **NTFS Permissions** |
|---|---|---|
| Public folder. A folder that can be accessed by everyone. | Grant Change permission to the Users group. | Grant Modify permission to the Users group. |

| | | |
|---|---|---|
| Drop folder. A folder where users can drop confidential reports or homework assignments that only the group manager or instructor can read. | Grant Change permission to the Users group.<br><br>Grant the Full Control permission to the group manager. | Grant the Write permission for the users' group that is applied to This Folder only. (This option is available on the Advanced page.)<br><br>If each user needs to have certain permissions to the files that he or she dropped, you can create a permission entry for the Creator Owner well-known security identifier (SID) and apply it to Subfolder and files only. For example, you can grant the Read and Write permission to the Creator Owner SID on the drop folder and apply it to all subfolders and files. This grants the user who dropped or created the file (the Creator Owner) the ability to read and write to the file. The Creator Owner can then access the file through the Run command using \\ServerName\DropFolder\FileName.<br><br>Note that as creator of the file, the user has the ability to Read and Change the permission on the file independent of any permissions that are granted explicitly.<br><br>Grant the Full Control permission for the group manager. |
| Application folder. A folder containing applications that can be run over the network. | Grant Read permission to the Users group. | Grant Read, Read & Execute, and List Folder Content permissions to the Users group. |
| Home folders. Individual folders for each user. Only the user has access to the folder. | Grant the Full Control permission to each user for their respective folders. | Grant the Full Control permission to each user for their respective folders. |

In addition, you should know the following information:

- NTFS permissions affect access both locally and remotely. NTFS permissions apply regardless of protocol. Share permissions, by contrast, apply only to network shares. Share permissions do not restrict access to any local user, or to any terminal server user, of the computer on which you have set share permissions. Therefore, share permissions do not provide privacy between users on a computer that is used by several users, nor on a terminal server that is accessed by several users.

- By default, Everyone does not include Anonymous Users; therefore, permissions that are applied to Everyone do not affect Anonymous Users. This default behavior is new for Windows Server 2003.

- A user who places a file on a share is the creator of that copy of the file and therefore has the ability to read and change the permissions on the file independent of the permissions that are inherited.

- When possible, you should attempt to grant permissions to groups instead of to specific user accounts.

**Permissions on Active Directory Objects**

You can manage permissions on objects in Active Directory at two levels:

- **At the object level**. Permissions that are allowed or denied at the object level apply to the entire object. For example, you can set an object-level permission on an organizational unit (OU) object that allows a particular group, such as Account Operators, to create child objects in the OU.

- **At the property level**. Permissions that are allowed or denied at the property level apply only to specific properties. For example, you can set a property-level permission on a User object that allows a particular user, such as Principal Self (that is, the user that is represented by the object), to change the **Home Address** property of the User object.

There is an important difference between how permissions are propagated from container objects in NTFS and how they are propagated from container objects in Active Directory. If you own a folder object in an NTFS volume, you can control who has access to the folder and all that it contains by setting a permission on the folder that applies not only to the folder but to all

objects in the folder. For example, you can add a permission to a Folder object that allows the Marketing group read access. If you apply the permission to **This folder, subfolders and files**, it is propagated to all objects that are contained in the folder. The Marketing group then has read access to the folder and to all its subfolders and files.

Propagation of permissions works differently for Active Directory objects. If you own a container object in Active Directory, you can allow access to certain types of objects in the container without allowing access to other types of objects in the container. For example, if you own an OU object, you can add a permission that allows the Marketing group read access and then apply the permission to a particular type of child object that is contained by the OU. If you apply it to Contact objects, for example, the permission is propagated only to Contact objects but not to User objects or to any other type of child object. In other words, permissions that are propagated from container objects in Active Directory can be object specific, whereas permissions that are propagated from container objects in an NTFS volume cannot be object specific.

**Permissions on Registry Keys**

Registry permissions are set on container objects only, and the individual keys inherit their permissions from the container. A registry key has two basic permissions: Full Control and Read. The Read permission includes the special permissions Read Control, Query Value, Notify, and Enumerate Subkeys. The Full Control permission includes all the special registry permissions, as shown in the following table. You can use Regedit.exe or the Subinacl.exe command-line tool to modify registry permissions.

**Special Registry Permissions**

| Permission | Description |
|---|---|
| Query Value | Allow the registry key value to be read. |
| Set Value | Allow the value of an existing key to be written. |
| Create Subkey | Allow the creation of subkeys. |
| Enumerate Subkeys | Allow the enumeration of subkeys. |
| Notify | Notify permission is required to request change notifications for registry keys or subkeys. |
| Create Link | Reserved by the operating system. |
| Delete | Allow the key to be deleted. |
| Write DACL | Allow modification of the container's DACL. |
| Write Owner | Allow modification of the container's owner. |
| Read Control | Allow the SACL to be read. |

**Permissions on Services**

Each service has a DACL containing the permissions on the service. You can view or modify the DACL by using Subinacl.exe. You can also export the security configuration of a computer into a security template and view the DACL by using the Security Templates Microsoft Management Console (MMC) snap-in. The following table shows the permissions for services.

**Permissions on Services**

| Permission | Description |
|---|---|
| Full Control | Grant full control of the service. |
| Query Template | Allow the service configuration to be viewed. |
| Change Template | Allow the service configuration to be modified. |
| Query Status | Allow the Service Control Manager (SCM) to be queried for the service status. |
| Enumerate Dependents | Allow the dependent services to be displayed. |
| Start | Allow the service to be started. |
| Stop | Allow the service to be stopped. |

| | |
|---|---|
| Pause and Continue | Allow the service to be paused and resumed. |
| Interrogate | Allow the service to respond to status queries. |
| User Defined Control | Allow special instructions to be given to the service. |
| Delete | Allow the service to be deleted. |
| Read Permissions | Allow the DACL of the service to be viewed. |
| Change Permissions | Allow the DACL of the service to be changed. |
| Take Ownership | Allow the owner of the service to be changed. |

**Conflicts Between User Rights and Permissions**

User rights enable administrators to assign specific privileges and logon rights to groups or users. These rights authorize users to perform specific actions, such as logging on to a system interactively or backing up files and directories. User rights are different from permissions because user rights apply to user accounts and permissions are attached to objects.

Administrators can use user rights to manage who has the authority to perform operations that span an entire computer, rather than a particular object. Administrators assign user rights (also known as privileges) to individual users or groups as part of the security settings for the computer. Although user rights can be managed centrally through Group Policy, they are applied locally. Users can (and usually do) have different user rights on different computers.

Unlike permissions, which are granted by an object's owner, user rights are assigned as part of the local security policy for the computer.

There are two types of user rights: privileges, such as the right to back up files and directories, and logon rights, such as the right to log on to a system locally. For more information about user rights and their descriptions and default settings, see "Security Policy Settings."

Conflicts between privileges and permissions normally occur only in situations where the rights that are required to administer a system overlap the rights of resource ownership. When rights conflict, a privilege overrides a permission.

For example, to create a backup of files and folders, backup software must be able to traverse all folders in an NTFS volume, list the contents of each folder, read the attributes of every file, and read data in any file that has its archive attribute set. It is impractical to arrange this access by coordinating with the owner of every file and folder; therefore, the required rights are included in the **Back up files and directories** (SeBackupPrivilege) privilege, which is assigned by default to two built-in groups: Administrators and Backup Operators. Any user who has this privilege can access all files and folders on the computer for the purpose of backing up the system. The same default permissions granted to Backup Operators that allow them to back up and restore files also make it possible for Backup Operators to use the group's permissions for other purposes, such as reading another user's files or installing Trojan horse programs. Therefore, the backup operators group should be limited to only highly trusted user accounts that require the ability to back up and restore computers.

The ability to take ownership of files and other objects is another case where an administrator's need to maintain the system takes priority over an owner's right to control access. Normally, you can take ownership of an object only if its current owner gives you permission to do so. Owners of NTFS objects can allow another user to take ownership by granting the other user Take Ownership permission; owners of Active Directory objects can grant another user Modify Owner permission. A user who has this privilege can take ownership of an object without the current owner's permission. By default, the privilege is assigned only to the built-in Administrators group. It is normally used by administrators to take and reassign ownership of resources when their current owner is no longer available.

**Related Information**

The following resources contain additional information that is relevant to this section:

- Microsoft Platform SDK link on the Web Resources page
- Security Policy Settings
- Permissions Tools and Settings

# Permissions Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Permissions Tools and Settings**

**In this section**

- Permissions Tools

This section contains information about the tools that are associated with permissions.

**Permissions Tools**

The following tools are associated with permissions.

**Permcopy.exe: Share Permissions Copy**

**Category**

This tool is included in Windows Server 2003 Resource Kit Tools.

**Version compatibility**

This tool is compatible with computers running Windows XP Professional or Windows Server 2003.

Share Permissions Copy is a command-line tool that copies share-level permissions (Full Control, Read, and Change) from one share to another.

To find more information about Share Permissions Copy, see "Permcopy.exe" in "Resource Kit Tools" in the Tools and Settings Collection.

**Perms.exe: User File Permissions Tool**

**Category**

This tool is included in Windows Server 2003 Resource Kit Tools.

**Version compatibility**

This tool is compatible with computers running Windows Server 2003.

The User File Permissions Tool is a command-line tool that displays user access permissions for a file or directory on an NTFS file system volume. It queries the permissions that are associated with a specific access control entry (ACE), displaying only the permissions that are granted by that particular ACE.

To find more information about the User File Permissions Tool, see "Perms.exe" in "Resource Kit Tools" in the Tools and Settings Collection.

**Takeown.exe: Take Ownership**

**Category**

This tool ships with Windows Server 2003.

**Version compatibility**

This tool is compatible with computers running Windows Server 2003.

Take Ownership is a command-line tool that displays or modifies discretionary access control lists (DACLs) on specified files.

To find more information about "Take Ownership," see "Takeown" in "Command-line reference A-Z" in Help and Support Center in Windows Server 2003.

# Security Principals Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Principals Technical Reference**

**In this subject**


- What Are Security Principals?
- How Security Principals Work
- Security Principals Tools and Settings


Security principals include the following:

- Any entity that can be authenticated by the system, such as a user account, a computer account, or a thread or process that runs in the security context of a user or computer account.
- Security groups of these accounts.

Every security principal is automatically assigned a security identifier (SID) when it is created.

Accounts and security groups that are created in an Active Directory directory service domain are directory objects, and they can be used to manage access to domain resources.

Local user accounts and security groups are created on a local computer, and they can be used to manage access to resources on that computer. Local user accounts and security groups are stored in and managed by the Security Accounts Manager (SAM) on the local computer.

In Microsoft Windows Server 2003, security principals are a foundation for controlling access to securable resources.

# What Are Security Principals?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Are Security Principals?**

**In this section**


- Accounts and Security Groups
- Technologies Related to Security Principals
- Related Information


Security principals include the following:

- Any entity that can be authenticated by the system, such as a user account, a computer account, or a thread or process that runs in the security context of a user or computer account.
- Security groups of these accounts.

Every security principal is automatically assigned a security identifier (SID) when it is created.

In Windows Server 2003, security principals are a foundation for controlling access to securable resources.

**Accounts and Security Groups**

Accounts and security groups that are created in an Active Directory domain are directory objects, and they can be used to manage access to domain resources.

Local user accounts and security groups are created on a local computer, and they can be used to manage access to resources on that computer. Local user accounts and security groups are stored in and managed by the Security Accounts Manager (SAM) on the local computer.

**User Accounts**

A user account uniquely identifies a person using a computer system, thereby allowing the system to enforce the appropriate authorization to allow or deny access to resources. User accounts can be created in Active Directory and also on local computers. User accounts are used to:

- Represent, identify, and authenticate the identity of a user. A user account enables a user to log on to computers and domains with a unique identity that can be authenticated by the computer or domain.
- Authorize (grant or deny) access to resources. After a user has been authenticated, the user is authorized (granted or denied) access to resources based on the permissions that are assigned to that user on the resource.

In addition, the actions that are carried out by a user account can be audited.

**Security Groups**

A security group is a collection of user accounts, computer accounts, and other groups that can be managed as a single unit from a security perspective. Groups can be either Active Directory based or local to a particular computer. In Windows Server 2003, there are several built-in security groups that are preconfigured with the appropriate rights and permissions to perform specific tasks. Additionally, you can (and, typically, will) also create a security group for each unique combination of security requirements that applies to multiple users in your organization.

You use groups in Active Directory to manage rights and permissions to domain resources.

Local groups exist in the SAM database on local computers (on all Windows-based computers) except domain controllers. You use local groups to manage rights and permissions only to resources on the local computer.

Using security groups to manage access control solves the following problems. Using security groups:

- Simplifies administration. You can assign a common set of rights, a common set of permissions, or both to many accounts at one time, rather than assigning them to each account individually. Also, when users transfer jobs or leave the organization, permissions are not tied to their user accounts, making permission reassignment or removal easier.

- Allows you to implement a role-based access-control model for granting permissions by using groups with different scopes (local, global, domain local, and universal) for the appropriate purposes.

- Minimizes the size of access control lists (ACLs) and speeds security checking. A security group has its own SID; therefore, the group SID can be used to specify permissions for a resource. In an environment with more than a few thousand users, if the SIDs of individual user accounts are used to specify access to a resource, the ACL of that resource can become unmanageably large, and the time needed for the system to check permissions to the resource can become unacceptable.

**Technologies Related to Security Principals**

Security principals are closely related to the following components and technologies in Windows Server 2003.

**Authorization and Access Control Components**

- **SIDs**. Each security principal is automatically assigned a SID when it is created.

- **Access Tokens**. After a user logs on and is authenticated, the system creates an access token for the user that contains the SID of the user, the SIDs of all the domain and local security groups that the user is a member of, and some well-known SIDs. Every process the user creates inherits the user's access token, which is then used to determine whether to grant that user access to a system resource in a given logon session.

- **Security Descriptors and ACLs**. A security descriptor is a data structure that is associated with each securable object. The security descriptor contains a discretionary access control list (DACL), with SIDs for the user's and groups that are allowed or denied access to that object.

- **Permissions**. Permissions to access securable objects such as Active Directory objects, files, and registry settings are granted to security principals.

**Authentication**

A user account enables a user to log on to computers and domains with an identity that can be authenticated by the computer or domain.

**Active Directory**

Accounts and security groups that are created in an Active Directory domain are stored in the Active Directory database and managed using Active Directory tools.

The following diagram shows the Windows Server 2003 authorization and access control process. In this diagram, the subject (a thread in a process that is initiated by a user) attempts to access an object, such as a shared folder. The information in the user's access token is compared to the access control entries (ACEs) in the object's security descriptor, and the access decision is made. The SIDs of security principals are used in the users access token and in the ACEs in the object's security descriptor.

### Authorization and Access Control Process




**Related Information**

The following resources contain additional information that is relevant to this section:

- Security Identifiers Technical Reference
- Access Tokens Technical Reference
- Security Descriptors and Access Control Lists Technical Reference
- Permissions Technical Reference
- Logon and Authentication Technologies

# How Security Principals Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Security Principals Work**

**In this section**


- Active Directory Accounts and Security Groups
- Local Accounts and Security Groups
- Related Information


Security principals include the following:

- Any entity that can be authenticated by the system, such as a user account, a computer account, or a thread or process that runs in the security context of a user or computer account.
- Security groups of these accounts.

Every security principal is automatically assigned a security identifier (SID) when it is created.

Security principals that are created in an Active Directory domain are Active Directory objects, which can be used to manage access to domain resources. Local users and security groups are created on a local computer and can be used to manage access to resources on that computer. Local user accounts and groups are managed by the Security Accounts Manager (SAM) on the local computer.

**Active Directory Accounts and Security Groups**

Active Directory user accounts and computer accounts can represent a physical entity, such as a computer or person. User accounts can also be used as dedicated service accounts for some applications. Groups are used to collect user accounts, computer accounts, and other groups into manageable units.

In Windows Server 2003 there are several built-in accounts and security groups that are preconfigured with the appropriate rights and permissions to perform specific tasks. For Active Directory, there are two types of administrative responsibility. Service administrators are responsible for maintaining and delivering the directory service, including domain controller management and directory service configuration. Data administrators are responsible for maintaining the data that is stored in the directory service and on domain member servers and workstations.

It is important to understand which default accounts and groups are service administrators. Service administration accounts and groups have the most widespread power in the network environment and require the most protection.

**User Accounts**

User accounts can be created, disabled, reset, and deleted by using the Active Directory Users and Computers snap-in in Microsoft Management Console (MMC) and by using command-line tools. In addition, built-in accounts are created when you create the Active Directory domain.

**Built-in User Accounts**

The Users container in Active Directory Users and Computers displays the built-in user accounts: Administrator, Guest, KRBTGT, and Support_388945a0. These built-in user accounts are created automatically when you create the domain. The HelpAssistant account is installed when a Remote Assistance session is established.

Each built-in account has a different combination of rights and permissions. The Administrator account has the most extensive rights and permissions over the domain, while the Guest account has limited rights and permissions.

Some of the administrative accounts that are listed in these tables are protected by a background process that periodically checks and applies a specific security descriptor, which is a data structure that contains security information that is associated with a protected object. This process ensures that any successful unauthorized attempt to modify the security descriptor on one of the administrative accounts or groups will be overwritten with the protected settings.

This security descriptor is present on the AdminSDHolder object. This means that if you want to modify the permissions on one

of the service administrator groups or on any of its member accounts, you must modify the security descriptor on the AdminSDHolder object so that it will be applied consistently. Be careful when making these modifications, because you are also changing the default settings that will be applied to all of your protected administrative accounts.

The following table describes each default user account on domain controllers running Windows Server 2003.

**Built-in User Accounts on Domain Controllers Running Windows Server 2003**

**Administrator**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-500 |
| **Type** | User |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | N/A |
| **Default member of** | Administrators, Domain Admins, Enterprise Administrators, Domain Users<br><br>(Membership in Domain Users is due to the fact that the Primary Group ID of all user accounts is Domain Users.) |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Yes |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Notes** | This account is the first account created during installation of the operating system. The account cannot be deleted or locked out, but it can be renamed or disabled. It is a member of the Administrators group and cannot be removed from that group.<br><br>The Administrator account is a default member of the Administrators, Domain Admins, Enterprise Admins, Group Policy Creator Owners, and Schema Admins groups in Active Directory. The Administrator account is known to exist on many versions of Windows; therefore, renaming or disabling this account makes it more difficult for malicious users to try and gain access to it.<br><br>When the Administrator account is disabled, it can still be used to gain access to a domain controller using Safe Mode. |

**Guest**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-501 |
| **Type** | User |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | Domain Guests, Guests |

| | |
|---|---|
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Can be moved out, but it is not recommended. |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Notes** | A user account for people who do not have individual accounts. This user account does not require a password. By default, the Guest account is disabled, and it is recommended that it stay disabled.<br><br>You can set rights and permissions for the Guest account just like any user account. By default, the Guest account is a member of the built-in Guests group and the Domain Guests global group, which allows a user to log on to a domain. |

**HelpAssistant**

| Attribute | Value |
|---|---|
| **Type** | User |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Notes** | The primary account that is used to establish a Remote Assistance session. This account has limited access to the computer. The HelpAssistant account is managed by the Remote Desktop Help Session Manager service and is automatically deleted if no Remote Assistance requests are pending. |

**KRBTGT**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-502 |
| **Type** | User |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | Domain Users (This membership is due to the fact that the Primary Group ID of all user accounts is Domain Users.) |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Can be moved out, but it is not recommended. |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Notes** | A service account that is used by the Key Distribution Center (KDC) service. |

**Support_388945a0**

| Attribute | Value |
|---|---|
| **Type** | User |
| **Default container** | CN=Users, DC=<domain>, DC= |

| | |
|---|---|
| **Notes** | The Support_388945a0 account enables Help and Support Service interoperability with signed scripts. This account is primarily used to control access to signed scripts that are accessible from within Help and Support Services. Administrators can use this account to delegate the ability for an ordinary user, who does not have administrative access over a computer, to run signed scripts from links embedded within Help and Support Services. These scripts can be programmed to use the Support_388945a0 account credentials instead of the user's credentials to perform specific administrative operations on the local computer that otherwise would not be supported by the ordinary user's account. When the delegated user clicks on a link in Help and Support Services, the script executes under the security context of the Support_388945a0 account. This account has limited access to the computer and is disabled by default. |

**Account Options**

Each Active Directory user account has a number of account options. The following table describes options that you can use to configure password settings and security-specific information for user accounts.

**Account Options for User Accounts**

| Account Option | Description |
|---|---|
| User must change password at next logon | Forces user to change their password the next time that the user logs on to the network. Use this option when you want to ensure that the user will be the only person to know their password. |
| User cannot change password | Prevents users from changing their password. Use this option when you want to maintain control over a user account, such as for a guest or temporary account. |
| Password never expires | Prevents a user password from expiring. It is recommended that service accounts have this option enabled and use strong passwords. |
| Store passwords using reversible encryption | Provides support for applications that use protocols that require knowledge of the plaintext form of the user's password for authentication purposes.<br><br>This option is required when using Challenge-Handshake Authentication Protocol (CHAP) in Internet Authentication Services (IAS), and when using Digest Authentication in Internet Information Services (IIS). |
| Account is disabled | Prevents user from logging on with the selected account. Many administrators use disabled accounts as templates for common user accounts. |
| Smart card is required for interactive logon | Requires that a user possess a smart card to log on to the network interactively. The user must also have a smart card reader attached to their computer and a valid personal identification number (PIN) for the smart card. |
| Account is trusted for delegation | Allows a service running under this account to perform operations on behalf of other user accounts on the network. A service running under a user account (otherwise known as a service account) that is trusted for delegation can impersonate a client to gain access to resources, either on the computer where the service is running or on other computers. In a forest that is set to the Windows Server 2003 functional level, this setting is found on the Delegation tab, and it is available only for accounts that have been assigned service principal names (SPNs), as set using the **setspn** command from Windows Support Tools. This is a security-sensitive capability, and it should be assigned cautiously.<br><br>This option is available only on domain controllers running Windows Server 2003 where the domain functionality is set to Windows 2000 mixed or Windows 2000 native. On domain controllers running Windows Server 2003 where the domain functional level is set to Windows Server 2003, the **Delegation** tab is used to configure delegation settings. The **Delegation** tab appears only for accounts that have an assigned SPN. |
| Account is sensitive and cannot be delegated | Allows control over a user account, such as for a guest or temporary account. This option can be used if this account cannot be assigned for delegation by another account. |
| Use DES encryption types for this account | Provides support for the Data Encryption Standard (DES). DES supports multiple levels of encryption, including Microsoft Point-to-Point Encryption (MPPE) Standard (40-bit and 56-bit), MPPE Strong (128-bit), Internet Protocol security (IPSec) DES (40-bit), IPSec 56-bit DES, and IPSec Triple DES (3DES). |

| | |
|---|---|
| Do not require Kerberos preauthentication | Provides support for alternate implementations of the Kerberos protocol. Domain controllers running Windows 2000 or Windows Server 2003 can use other mechanisms to synchronize time. Because preauthentication provides additional security, use caution when enabling this option. |

**Active Directory Groups**

Groups are used to collect user accounts, computer accounts, and other groups into manageable units. Working with groups instead of with individual users helps simplify network maintenance and administration.

There are two types of groups in Active Directory: distribution groups and security groups. You can use distribution groups to create e-mail distribution lists and security groups to assign permissions to shared resources.

**Distribution Groups**

Distribution groups can be used only with e-mail applications (such as Exchange) to send e-mail to collections of users. Distribution groups are not security enabled, which means that they cannot be listed in discretionary access control lists (DACLs). Distribution groups are not covered in this Technical Reference subject.

**Security Groups**

When they are used with care, security groups provide an efficient way to assign access to resources on your network. Using security groups, you can:

- Assign user rights to security groups in Active Directory.

  User rights are assigned to a security group to determine what members of that group can do within the scope of a domain (or forest). User rights are automatically assigned to some security groups when Active Directory is installed to help administrators define a person's administrative role in the domain. For example, a user who is added to the Backup Operators group in Active Directory has the ability to back up and restore files and directories that are located on each domain controller in the domain.

  This is possible because, by default, the user rights **Backup files and directories** and **Restore files and directories** are automatically assigned to the Backup Operators group. Therefore, members of this group inherit the user rights that are assigned to that group.

  You can assign user rights to security groups, by using Group Policy, to help delegate specific tasks. You should always use discretion when assigning delegated tasks because an untrained user who is assigned too many rights can potentially cause significant harm to your network.

- Assign permissions to security groups on resources.

  Permissions should not be confused with user rights. Permissions are assigned to the security group on the shared resource. Permissions determine who can access the resource and the level of access, such as Full Control. Some permissions that are set on domain objects are automatically assigned to allow various levels of access to default security groups such as the Account Operators group or the Domain Admins group.

  Security groups are listed in DACLs that define permissions on resources and objects. When assigning permissions for resources (file shares, printers, and so on), administrators should assign those permissions to a security group rather than to individual users. The permissions are assigned once to the group, instead of several times to each individual user. Each account that is added to a group receives the rights that are assigned to that group in Active Directory and receives the permissions that are defined for that group at the resource.

Like distribution groups, security groups can also be used as an e-mail entity. Sending an e-mail message to the group sends the message to all the members of the group.

**Converting Between Security and Distribution Groups**

A group can be converted from a security group to a distribution group, and vice versa, at any time, but only if the domain functional level is set to Windows 2000 native or higher. No groups can be converted while the domain functional level is set to Windows 2000 mixed.

### Note

- Although a contact can be added to a security group as well as to a distribution group, contacts cannot be assigned rights and permissions. Contacts in a group can be sent e-mail.

**Group Scope**

Groups are characterized by a scope that identifies the extent to which the group is applied in the domain tree or forest. There are three group scopes that are defined by Active Directory: Universal, Global, and Domain Local.

### Note

- In addition to these three scopes, the default groups in the Builtin container have a Group Scope of Builtin Local. Their Group Scope and Group Type cannot be changed.

The scope of the group defines where the group can be granted permissions. Universal groups are only available when the domain functional level is set to Windows 2000 native or Windows Server 2003.

The following table lists the three group scopes and more information about each scope for a security group.

**Group Scopes**

| Scope | Possible Members | Scope Conversion | Can Grant Permissions | Possible Member of |
|---|---|---|---|---|
| Universal | Accounts from any domain in the same forest<br><br>Global groups from any domain in the same forest<br><br>Other universal groups from any domain in the same forest | Can be converted to domain local scope<br><br>Can be converted to global scope as long as the group does not contain any other universal groups | On any domain in the same forest or trusting forests | Other universal groups in the same forest<br><br>Domain local groups in the same forest or trusting forests<br><br>Local groups on computers in the same forest or trusting forests |
| Global | Accounts from the same domain<br><br>Other global groups from the same domain | Can be converted to universal scope as long as the group is not a member of any other global group | On any domain in the same forest or trusting domains or forests | Universal groups from any domain in the same forest<br><br>Other global groups from the same domain<br><br>Domain local groups from any domain in the same forest or from any trusting domain |
| Domain Local | Accounts from any domain or any trusted domain<br><br>Global groups from any domain or any trusted domain<br><br>Universal groups from any domain in the same forest<br><br>Other domain local groups from the same domain<br><br>Accounts, global groups, and universal groups from other forests and from external and Windows NT 4.0 domains | Can be converted to universal scope as long as the group does not contain any other domain local groups | Within the same domain | Other domain local groups from the same domain<br><br>Local groups on computers in the same domain, excluding built-in groups that have well-known SIDs |

**Default Groups**

Default groups, such as the Domain Admins group, are security groups that are created automatically when you create an Active Directory domain. You can use these predefined groups to help control access to shared resources and to delegate specific domain-wide administrative roles.

Many default groups are automatically assigned a set of user rights that authorize members of the group to perform specific actions in a domain, such as logging on to a local system or backing up files and folders. For example, a member of the Backup Operators group has the right to perform backup operations for all domain controllers in the domain.

When you add a user to a group, the user receives all the user rights that are assigned to the group and all the permissions that are assigned to the group on any shared resources.

Default groups are located in the Builtin container and the Users container in Active Directory Users and Computers. The Builtin container contains groups that are defined with domain local scope. The Users container contains groups that are defined with global scope and groups that are defined with domain local scope. You can move groups that are located in these containers to other groups or organizational units (OU) within the domain, but you cannot move them to other domains.

Some of the administrative groups that are listed in these tables and all members of these groups are protected by a background process that periodically checks and applies a specific security descriptor, which is a data structure that contains security information that is associated with a protected object. This process ensures that any successful unauthorized attempt to modify the security descriptor on one of the administrative accounts or groups will be overwritten with the protected settings.

This security descriptor is present on the AdminSDHolder object. This means that if you want to modify the permissions on one of the service administrator groups or on any of its member accounts, you must modify the security descriptor on the AdminSDHolder object so that it will be applied consistently. Be careful when making these modifications because you are also changing the default settings that will be applied to all of your protected administrative accounts.

The following tables provide descriptions of the default groups that are located in the Builtin and Users containers.

**Account Operators**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-548 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | Allow log on locally<br><br>Shut down the system |
| **Notes** | By default, this built-in group has no members, and it can create and manage users and groups in the domain, including its own membership and that of the Server Operators. This group is a service administrator because it can modify Server Operators, which in turn can modify domain controller settings. As a best practice, leave the membership of this group empty and do not use it at all for any delegated administration. This group cannot be renamed, deleted, or moved. |

**Administrators**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-544 |

| | |
|---|---|
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | Administrator, Domain Admins, Enterprise Admins |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | Access this computer from the network<br>Adjust memory quotas for a process<br>Back up files and directories<br>Bypass traverse checking<br>Change the system time<br>Create a pagefile<br>Debug programs<br>Enable computer and user accounts to be trusted for delegation<br>Force a shutdown from a remote system<br>Increase scheduling priority<br>Load and unload device drivers<br>Allow log on locally<br>Manage auditing and security log<br>Modify firmware environment values<br>Profile single process<br>Profile system performance<br>Remove computer from docking station<br>Restore files and directories<br>Shut down the system<br>Take ownership of files or other objects |

| | |
|---|---|
| **Notes** | The Administrators group has built-in capabilities that give its members full control over the system. This group cannot be renamed, deleted, or moved. This built-in group controls access to all the domain controllers in its domain, and it can change the membership of all administrative groups. Its own membership can be modified by the default service administrator groups Administrators and Domain Admins in the domain, as well as by the Enterprise Admins group. This group has the special privilege to take ownership of any object in the directory or any resource on a domain controller. This account is considered a service administrator because its members have full access to the domain controllers in the domain. |

**Backup Operators**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-551 |
| **Type** | Builtin Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | Back up files and directories<br><br>Allow log on locally<br><br>Restore files and directories<br><br>Shut down the system |
| **Notes** | Backup Operators can back up and restore all files on a computer, regardless of the permissions that protect those files. Backup Operators also can log on to the computer and shut it down. This group cannot be renamed, deleted, or moved. By default, this built-in group has no members, and it can perform backup and restore operations on domain controllers. Its membership can be modified by the default service administrator groups Administrators and Domain Admins in the domain, as well as by the Enterprise Admins group. It cannot modify the membership of any administrative groups. While members of this group cannot change server settings or modify the configuration of the directory, they do have the permissions needed to replace files (including operating system files) on domain controllers. Because of this, members of this group are considered service administrators. |

**Cert Publishers**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-517 |
| **Type** | Domain Local |
| **Default container** | CN=Users, DC=<domain>, DC= |

| | |
|---|---|
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | No |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | No default user rights |
| **Notes** | A group that includes all computers that are running an enterprise certification authority. Cert Publishers are authorized to publish certificates for User objects in Active Directory. |

**Domain Admins**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-512 |
| **Type** | Domain Global |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | Administrator |
| **Default member of** | Administrators |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Yes |
| **Safe to delegate management of this group to non-Service admins?** | No |

| | |
|---|---|
| **Default User Rights** | Access this computer from the network<br>Adjust memory quotas for a process<br>Back up files and directories<br>Bypass traverse checking<br>Change the system time<br>Create a pagefile<br>Debug programs<br>Enable computer and user accounts to be trusted for delegation<br>Force a shutdown from a remote system<br>Increase scheduling priority<br>Load and unload device drivers<br>Allow log on locally<br>Manage auditing and security log<br>Modify firmware environment values<br>Profile single process<br>Profile system performance<br>Remove computer from docking station<br>Restore files and directories<br>Shut down the system<br>Take ownership of files or other objects |
| **Notes** | A group whose members are authorized to administer the domain. By default, the Domain Admins group is a member of the Administrators group on all computers that have joined a domain, including the domain controllers. Domain Admins is the default owner of any object that is created in Active Directory for the domain by any member of the group. If members of the group create other objects, such as files, the default owner is the Administrators group. This group controls access to all domain controllers in a domain, and it can modify the membership of all administrative accounts in the domain. Its own membership can be modified by the service administrator groups Administrators and Domain Admins in its domain, as well as the Enterprise Admins group. This is a service administrator account because its members have full access to a domain's domain controllers. |

**Domain Computers**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-515 |
| **Type** | Domain Global |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | All computers joined to the domain, excluding domain controllers |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Yes (but not required) |

| | |
|---|---|
| **Safe to delegate management of this group to non-Service admins?** | Yes |
| **Default User Rights** | No default user rights |
| **Notes** | The group includes all computers that have joined the domain, excluding domain controllers. |

**Domain Controllers**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-516 |
| **Type** | Domain Global |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | Computer accounts for all Domain Controllers of the domain |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | No |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | No default user rights |
| **Notes** | The group includes all domain controllers in the domain. New domain controllers are added to this group automatically. |

**Domain Guests**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-514 |
| **Type** | Domain Global |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | Guest |
| **Default member of** | Guests |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Can be moved out but it is not recommended |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | No default user rights |
| **Notes** | This group includes the domain's built-in Guest account. |

**Domain Users**

| Attribute | Value |
|---|---|

| | |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-513 |
| **Type** | Domain Global |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | All user accounts in the domain |
| **Default member of** | Users |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Yes |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | No default user rights |
| **Notes** | This group includes all user accounts in a domain. When you create a user account in a domain, it is added to this group automatically. |

**Enterprise Admins**

| **Attribute** | **Value** |
|---|---|
| **Well-Known SID/RID** | S-1-5-<root domain>-519 |
| **Type** | Universal (if Domain is in Native-Mode) else Domain Global |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | Administrator |
| **Default member of** | Administrators |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Yes |
| **Safe to delegate management of this group to non-Service admins?** | No |

| | |
|---|---|
| **Default User Rights** | Access this computer from the network<br>Adjust memory quotas for a process<br>Back up files and directories<br>Bypass traverse checking<br>Change the system time<br>Create a pagefile<br>Debug programs<br>Enable computer and user accounts to be trusted for delegation<br>Force shutdown from a remote system<br>Increase scheduling priority<br>Load and unload device drivers<br>Allow log on locally<br>Manage auditing and security log<br>Modify firmware environment values<br>Profile single process<br>Profile system performance<br>Remove computer from docking station<br>Restore files and directories<br>Shut down the system<br>Take ownership of files or other objects |
| **Notes** | This group exists only in the root domain of an Active Directory forest of domains. It is a universal group if the domain is in native mode; it is a global group if the domain is in mixed mode. The group is authorized to make forest-wide changes in Active Directory, such as adding child domains. By default, the only member of the group is the Administrator account for the forest root domain. This group is automatically added to the Administrators group in every domain in the forest, providing complete access to the configuration of all domain controllers. This group can modify the membership of all administrative groups. Its own membership can be modified only by the default service administrator groups in the root domain. This account is considered a service administrator. |

**Group Policy Creators Owners**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-520 |
| **Type** | Domain Global |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | Administrator |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | No |

| | |
|---|---|
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | No default user rights |
| **Notes** | This group that is authorized to create new Group Policy objects in Active Directory. By default, the only member of the group is Administrator. |

**Guests**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-546 |
| **Type** | Builtin Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | Guest |
| **Default member of** | Guest, Domain Guests |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | No default user rights |
| **Notes** | A built-in group. By default, the only member is the Guest account. The Guests group allows occasional or one-time users to log on with limited privileges to a computer's built-in Guest account. |

**Incoming Forest Trust Builders**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-557 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | No default user rights |
| **Notes** | This group cannot be renamed, deleted, or moved. |

**Network Configuration Operators**

| Attribute | Value |
| --- | --- |
| **Well-Known SID/RID** | S-1-5-32-556 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | Yes |
| **Default User Rights** | No default user rights |
| **Notes** | This group cannot be renamed, deleted, or moved. |

**Performance Log Users**

| Attribute | Value |
| --- | --- |
| **Well-Known SID/RID** | S-1-5-32-559 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | Yes |
| **Default User Rights** | No default user rights |
| **Notes** | This account cannot be renamed, deleted, or moved. |

**Performance Monitor Users**

| Attribute | Value |
| --- | --- |
| **Well-Known SID/RID** | S-1-5-32-558 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |

| | |
|---|---|
| **Safe to delegate management of this group to non-Service admins?** | Yes |
| **Default User Rights** | No default user rights |
| **Notes** | This group cannot be renamed, deleted, or moved. |

**Pre–Windows 2000 Compatible Access**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-554 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | Refer to the Notes column |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | Access this computer from the network<br><br>Bypass traverse checking |
| **Notes** | In Windows 2000, if you choose the Pre–Windows 2000 Compatible Permissions mode, Everyone (including Anonymous) is a member, and if you choose the Windows 2000–only permissions mode, the membership is empty. In Windows Server 2003, if you choose the Pre–Windows 2000 Compatible Permissions mode, Everyone and Anonymous are members, and if you choose the Windows 2000-only permissions mode, Authenticated Users are members. |

**Print Operators**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-550 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Cannot be moved |

| | |
|---|---|
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | Allow log on locally<br><br>Shut down the system |
| **Notes** | A built-in group that exists only on domain controllers. By default, the only member is the Domain Users group. Print Operators can manage printers and document queues. This group cannot be renamed, deleted, or moved. |

**RAS and IAS Servers**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<domain>-553 |
| **Type** | Domain Local |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Yes |
| **Safe to delegate management of this group to non-Service admins?** | Yes |
| **Default User Rights** | No default user rights |
| **Notes** | By default, this group has no members. Computers that are running the Routing and Remote Access service are added to the group automatically. Members of this group have access to certain properties of User objects, such as Read Account Restrictions, Read Logon Information, and Read Remote Access Information. |

**Remote Desktop Users**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-555 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | Yes |

| | |
|---|---|
| **Default User Rights** | No default user rights |
| **Notes** | This group cannot be renamed, deleted or moved. |

**Schema Admins**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-<root domain>-518 |
| **Type** | Universal (if Domain is in Native-Mode) else Domain Global |
| **Default container** | CN=Users, DC=<domain>, DC= |
| **Default members** | Administrator |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Yes |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | No default user rights |
| **Notes** | This group exists only in the root domain of an Active Directory forest of domains. It is a universal group if the domain is in native mode; it is a global group if the domain is in mixed mode. The group is authorized to make schema changes in Active Directory. By default, the only member of the group is the Administrator account for the forest root domain. This group has full administrative access to the schema. The membership of this group can be modified by any of the service administrator groups in the root domain. This account is considered a service administrator because its members can modify the schema, which governs the structure and content of the entire directory. |

**Server Operators**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-549 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | Yes |
| **Safe to move out of default container?** | Cannot be moved |

| | |
|---|---|
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | Back up files and directories<br><br>Change the system time<br><br>Force shutdown from a remote system<br><br>Allow log on locally<br><br>Restore files and directories<br><br>Shut down the system |
| **Notes** | This group exists only on domain controllers. By default, the group has no members. Server Operators can log on to a server interactively, create and delete network shares, start and stop services, back up and restore files, format the hard disk of the computer, and shut down the computer. This group cannot be renamed, deleted, or moved. By default, this built-in group has no members, and it has access to server configuration options on domain controllers. Its membership is controlled by the service administrator groups Administrators and Domain Admins in the domain, as well as the Enterprise Admins group. It cannot change any administrative group memberships. This is a service administrator account because its members have physical access to domain controllers and they can perform maintenance tasks (such as backup and restore), and they have the ability to change binaries that are installed on the domain controllers. |

**Terminal Server License Servers**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-561 |
| **Type** | BuiltIn Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | None |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | Yes |
| **Notes** | This group cannot be renamed, deleted, or moved. |

**Users**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-545 |
| **Type** | Builtin Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | Authenticated Users, Domain Users |
| **Default member of** | Domain Users (This membership is due to the fact that the Primary Group ID of all user accounts is Domain Users.) |

| | |
|---|---|
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | No |
| **Default User Rights** | No default user rights |
| **Notes** | After the initial installation of the operating system, the only member is the Authenticated Users group. When a computer joins a domain, the Domain Users group is added to the Users group on the computer. Users can perform tasks such as running applications, using local and network printers, shutting down the computer, and locking the computer. Users can install applications that only they are allowed to use if the installation program of the application supports per-user installation. This group cannot be renamed, deleted, or moved. |

**Windows Authorization Access Group**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-32-560 |
| **Type** | Builtin Local |
| **Default container** | CN=BuiltIn, DC=<domain>, DC= |
| **Default members** | Enterprise Domain Controllers |
| **Default member of** | None |
| **Protected by ADMINSDHOLDER?** | No |
| **Safe to move out of default container?** | Cannot be moved |
| **Safe to delegate management of this group to non-Service admins?** | Yes |
| **Notes** | This group cannot be renamed, deleted, or moved. |

**Special Identities**

In addition to the groups in the Users and Builtin containers, servers running Windows Server 2003 include several special identities. For convenience, these identities are generally referred to as groups. These special groups do not have specific memberships that can be modified, but they can represent different users at different times, depending on the circumstances.

Although the special identities can be assigned rights and permissions to resources, the memberships cannot be modified or viewed. Group scopes do not apply to special identities. Users are automatically assigned to these special identities whenever they log on or access a particular resource.

The special groups are described in the following tables.

**Anonymous Logon**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-7 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A user who has logged on anonymously. |

**Authenticated User**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-11 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A group that includes all users whose identities were authenticated when they logged on. Membership is controlled by the operating system. |

**Batch**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-3 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A group that implicitly includes all users who have logged on through a batch queue facility such as task scheduler jobs. Membership is controlled by the operating system. |

**Creator Group**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-3-1 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A placeholder in an inheritable access control entry (ACE). When the ACE is inherited, the system replaces this SID with the SID for the primary group of the object's current owner. The primary group is used only by the Portable Operating System Interface for UNIX (POSIX subsystem. |

**Creator Owner**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-3-0 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A placeholder in an inheritable ACE. When the ACE is inherited, the system replaces this SID with the SID for the object's current owner. |

**Dialup**

| Attribute | Value |
|---|---|

| | |
|---|---|
| **Well-Known SID/RID** | S-1-5-1 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | Used for accounts that are connected to the computer over a dial-up connection. |

**Digest Authentication**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-64-21 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |

**Enterprise Domain Controllers**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-9 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A group that includes all domain controllers in an Active Directory forest of domains. Membership is controlled by the operating system. |

**Everyone**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-1-0 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A group that includes Authenticated Users and Guest. Whenever a user logs on to the network, the user is automatically added to the Everyone group. On computers running Windows 2000 and earlier, the Everyone group included the Anonymous Logon group as a default member, but in Windows Server 2003, the Everyone group contains only Authenticated Users and Guest; it no longer includes Anonymous Logon by default (although this can be changed).<br><br>Membership is controlled by the operating system. |

**Interactive**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-4 |
| **Object Class** | Foreign Security Principal |

| | |
|---|---|
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A group that includes all users who have logged on interactively. Membership is controlled by the operating system. |

**Local Service**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-19 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | The Local Service account is similar to an authenticated user account. The Local Service account has the same level of access to resources and objects as members of the Users group. This limited access helps safeguard your system if individual services or processes are compromised. Services that run as the Local Service account access network resources as a null session with anonymous credentials. The name of the account is NT AUTHORITY\LocalService. This account does not have a password. |

**LocalSystem**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-18 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A service account that is used by the operating system. The Local System account is a powerful account that has full access to the system and acts as the computer on the network. If a service logs on to the Local System account on a domain controller, that service has access to the entire domain. Some services are configured by default to log on to the Local System account. Do not change the default service setting. The name of the account is LocalSystem. This account does not have a password. |

**Network**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-2 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A group that implicitly includes all users who are logged on through a network connection. Membership is controlled by the operating system. |

**Network Service**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-20 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | The Network Service account is similar to an authenticated user account. The Network Service account has the same level of access to resources and objects as members of the Users group. This limited access helps safeguard your system if individual services or processes are compromised. Services that run as the Network Service account access network resources using the credentials of the computer account. The name of the account is NT AUTHORITY\NetworkService. This account does not have a password. |

**NTLM Authentication**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-64-10 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |

**Other Organization**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-1000 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A group that implicitly includes all users who are logged on to the system through a dial-up connection. Membership is controlled by the operating system. |

**Principal Self**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-10 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |

| | |
|---|---|
| **Description** | A placeholder in an ACE on a user, group, or computer object in Active Directory. When you grant permissions to Principal Self, you grant them to the security principal that is represented by the object. During an access check, the operating system replaces the SID for Principal Self with the SID for the security principal that is represented by the object. |

**Proxy**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-8 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | Does not currently apply: this SID is not used in Windows 2000 Server or Windows Server 2003. |

**Remote Interactive Logon**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-14 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | Represents all users who are currently logged on to the computer using a Remote Desktop connection. This group is a subset of the Interactive group. Access tokens that contain the Remote Interactive Logon SID also contain the Interactive SID. |

**Restricted Code**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-12 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | Used by a process that is executing in a restricted security context, such as running an application with the RunAs service. When code executes at the restricted security level, the Restricted SID is added to the user's access token. |

**SChannel Authentication**

| Attribute | Value |
|---|---|
| **Well-Known SID/RID** | S-1-5-64-14 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |

**Service**

| Attribute | Value |
| --- | --- |
| **Well-Known SID/RID** | S-1-5-6 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A group that includes all security principals that have logged on as a service. Membership is controlled by the operating system. |

**Terminal Server User**

| Attribute | Value |
| --- | --- |
| **Well-Known SID/RID** | S-1-5-13 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |
| **Description** | A group that includes all users who have logged on to a Terminal Services server. Membership is controlled by the operating system. |

**This Organization**

| Attribute | Value |
| --- | --- |
| **Well-Known SID/RID** | S-1-5-15 |
| **Object Class** | Foreign Security Principal |
| **Default Location in Active Directory** | cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain> |

**Local Accounts and Security Groups**

You manage local user accounts and security groups by using Local Users and Groups, which is located in Computer Management, a collection of administrative tools that you can use to manage a single local or remote computer. You can use Local Users and Groups to secure and manage user accounts and groups that are stored locally on your computer. A local user or group can be assigned permissions and rights on a particular computer and that computer only. Local Users and Groups is available on the following client and server operating systems:

- Client computers running Windows 2000 Professional or Windows XP Professional.
- Member servers running Windows 2000 Server or Windows Server 2003.
- Stand-alone servers running Windows 2000 Server or Windows Server 2003.

You can use Local Users and Groups to limit the ability of users and groups to perform certain actions by assigning them rights and permissions. A right authorizes a user to perform certain actions on a computer, such as backing up files and folders or shutting down a computer. A permission is a rule that is associated with an object (usually a file, folder, or printer), and it regulates which users can have access to the object and in what manner.

You cannot use Local Users and Groups to view local users and groups after a member server has been promoted to a domain controller. However, you can use Local Users and Groups on a domain controller to target remote computers (that are not domain controllers) on the network.

Use Active Directory Users and Computers to manage users and groups in Active Directory.

**Local User Accounts**

The Users folder that is located in the Local Users and Groups snap-in in MMC displays the default user accounts as well as the user accounts that you create. The default user accounts are created automatically when you install a stand-alone server or

member server running Windows Server 2003. The following table describes the default user accounts on servers running Windows Server 2003.

**Default Local User Accounts on Servers Running Windows Server 2003**

| Default User Account | Description |
|---|---|
| Administrator account | The Administrator account has full control of the server and can assign user rights and access control permissions to users as necessary. This account must be used only for tasks that require administrative credentials. It is highly recommended that you set up this account to use a strong password.<br><br>The Administrator account is a member of the Administrators group on the server. The Administrator account can never be deleted or removed from the Administrators group, but it can be renamed or disabled. Because the Administrator account is known to exist on many versions of Windows, renaming or disabling this account will make it more difficult for malicious users to try and gain access to it.<br><br>The Administrator account is the account you use when you first set up the server. You use this account before you create an account for yourself.<br><br>Even when the Administrator account has been disabled, it can still be used to gain access to a computer by using Safe Mode. |
| Guest account | The Guest account is used by people who do not have an actual account on the computer. A user whose account is disabled, but not deleted, can also use the Guest account. The Guest account does not require a password. The Guest account is disabled by default, but you can enable it.<br><br>You can set rights and permissions for the Guest account just like any user account. By default, the Guest account is a member of the default Guests group, which allows a user to log on to a server. Additional rights, as well as any permissions, must be granted to the Guests group by a member of the Administrators group. The Guest account is disabled by default, and it is recommended that it stay disabled. |
| Support_388945a0 account | The Support_388945a0 account enables Help and Support Service interoperability with signed scripts. This account is primarily used to control access to signed scripts that are accessible from within Help and Support Services. Administrators can use this account to delegate the ability for an ordinary user, who does not have administrative access over a computer, to run signed scripts from links that are embedded in Help and Support Services. These scripts can be programmed to use the Support_388945a0 account credentials instead of the user's credentials to perform specific administrative operations on the local computer that otherwise would not be supported by the ordinary user's account. When the delegated user clicks on a link in Help and Support Services, the script will execute under the security context of the Support_388945a0 account. This account has limited access to the computer and is disabled by default. |
| HelpAssistant account (installed with a Remote Assistance session) | The primary account that is used to establish a Remote Assistance session. This account is created automatically when you request a Remote Assistance session, and it has limited access to the computer. The HelpAssistant account is managed by the Remote Desktop Help Session Manager service, and it will be automatically deleted if no Remote Assistance requests are pending. |

**Local Groups**

The Groups folder that is located in Local Users and Groups displays the default local groups as well as the local groups that you create. The default local groups are automatically created when you install a stand-alone server or a member server running Windows Server 2003. Belonging to a local group gives a user the rights and abilities to perform various tasks on the local computer.

You can add local user accounts, domain user accounts, computer accounts, and security groups to local groups. However, you cannot add local user accounts and local groups to domain group accounts.

The following table provides descriptions of the default groups that are located in the Groups folder, and it lists the assigned user rights for each group. These rights are assigned within the local security policy.

**Default Local Groups on Servers Running Windows Server 2003**

| Group | Description | Default User Rights |
|---|---|---|

| | | |
|---|---|---|
| Administrators | Members of this group have full control of the server and can assign user rights and access control permissions to users as necessary. The Administrator account is also a default member. When this server is joined to a domain, the Domain Admins group is automatically added to this group. Because this group has full control of the server, add users with caution. | Access this computer from the network<br><br>Adjust memory quotas for a process<br><br>Allow log on locally<br><br>Allow log on through Terminal Services<br><br>Back up files and directories<br><br>Bypass traverse checking<br><br>Change the system time<br><br>Create a pagefile; Debug programs<br><br>Force shutdown from a remote system<br><br>Increase scheduling priority<br><br>Load and unload device drivers<br><br>Manage auditing and security log<br><br>Modify firmware environment variables<br><br>Perform volume maintenance tasks<br><br>Profile single process |

| | | |
|---|---|---|
| | | Profile system performance<br><br>Remove computer from docking station<br><br>Restore files and directories<br><br>Shut down the system<br><br>Take ownership of files or other objects |
| Backup Operators | Members of this group can back up and restore files on the server, regardless of any permissions that protect these files. This is because the right to perform a backup takes precedence over all file permissions. File permissions cannot change security settings. | Access this computer from the network<br><br>Allow log on locally<br><br>Back up files and directories<br><br>Bypass traverse checking<br><br>Restore files and directories<br><br>Shut down the system |
| DHCP Administrators (installed with the DHCP Server service) | Members of this group have administrative access to the Dynamic Host Configuration Protocol (DHCP) Server service. This group provides a way to assign limited administrative access to the DHCP server only, while not providing full access to the server. Members of this group can administer DHCP on a server using the DHCP console or the Netsh command, but they are not able to perform other administrative actions on the server. | No default user rights |
| DHCP Users (installed with the DHCP Server service) | Members of this group have read-only access to the DHCP Server service. This allows members to view information and properties that are stored at a specified DHCP server. This information is useful to support staff when they need to obtain DHCP status reports. | No default user rights |
| Guests | Members of this group have a temporary profile created at log on, and when the member logs off, the profile is deleted. The Guest account (which is disabled by default) is also a default member of this group. | No default user rights |
| HelpServicesGroup | This group allows administrators to set rights that are common to all support applications. By default, the only group member is the account that is associated with Microsoft support applications, such as Remote Assistance. Do not add users to this group. | No default user rights |

| | | |
|---|---|---|
| Network Configuration Operators | Members of this group can make changes to TCP/IP settings and renew and release TCP/IP addresses. This group has no default members. | No default user rights |
| Performance Monitor Users | Members of this group can monitor performance counters on the server locally and from remote clients without being a member of the Administrators or Performance Log Users groups. | No default user rights |
| Performance Log Users | Members of this group can manage performance counters, logs, and alerts on the server locally and from remote clients without being a member of the Administrators group. | No default user rights |
| Power Users | Members of this group can create user accounts and then modify and delete the accounts that they have created. They can create local groups and then add or remove users from the local groups that they have created. They can also add or remove users from the Power Users, Users, and Guests groups. Members can create shared resources and administer the shared resources that they have created. They cannot take ownership of files, back up or restore directories, load or unload device drivers, or manage security and auditing logs. | Access this computer from the network<br><br>Allow log on locally<br><br>Bypass traverse checking<br><br>Change the system time<br><br>Profile single process<br><br>Remove computer from docking station<br><br>Shut down the system |
| Print Operators | Members of this group can manage printers and print queues. | No default user rights |
| Remote Desktop Users | Members of this group can remotely log on to a server. | Allow log on through Terminal Services |
| Replicator | The Replicator group supports replication functions. The only member of the Replicator group should be a domain user account that is used to log on the Replicator services of a domain controller. Do not add user accounts of actual users to this group. | No default user rights |
| Users | Members of this group can perform common tasks, such as running applications, using local and network printers, and locking the server. Users cannot share directories or create local printers. By default, the Domain Users, Authenticated Users, and Interactive groups are members of this group. Therefore, any user account that is created in the domain becomes a member of this group. | Access this computer from the network<br><br>Allow log on locally<br><br>Bypass traverse checking |

| | | |
|---|---|---|
| WINS Users (installed with WINS service) | Members of this group are permitted read-only access to Windows Internet Name Service (WINS). This allows members to view information and properties that are stored at a specified WINS server. This information is useful to support staff when they need to obtain WINS status reports. | No default user rights |

**Related Information**

The following resources contain additional information that is relevant to this section:

- Security Principals Tools and Settings

# Security Principals Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Principals Tools and Settings**

**In this section**


- Security Principals Tools
- Related Information


This section contains information about the tools that are associated with security principals.

**Security Principals Tools**

The following tools are associated with security principals.

**Dsa.msc: Active Directory Users and Computers**

**Category**

Active Directory Users and Computers is an Active Directory Administrative Tools Microsoft Management Console (MMC) snap-in. This tool is installed automatically when you install Active Directory, and it is available on the **Start** menu under **Programs** and then **Administrative Tools**. This tool also ships with the Windows Server 2003 Administration Tools Pack (Adminpak.msi).

**Version compatibility**

Active Directory Users and Computers runs on domain controllers that are running Windows Server 2003 and Windows 2000 Server.

The Windows Server 2003 version of Active Directory Users and Computers (which you can install on a domain controller or on an administrative workstation by using the Administration Tools Pack) can target domain controllers that are running Windows Server 2003 and Windows 2000 Server.

You can manage the following objects and their properties with this tool:

- User objects
- Group objects
- Computer objects

**Dsadd.exe: Dsadd**

**Category**

This tool ships with Windows Server 2003.

**Version compatibility**

This tool is compatible with domain controllers running Windows Server 2003.

You can use this tool to add specific types of objects to the directory.

For more information about Dsadd, see “Dsadd” in “Command-line reference A-Z” in Help and Support Center in Windows Server 2003.

**Dsget.exe: Dsget**

**Category**

This tool ships with Windows Server 2003.

**Version compatibility**

This tool is compatible with domain controllers running Windows Server 2003.

You can use this tool to display the selected properties of a specific object in the directory.

For more information about Dsget, see “Dsget” in “Command-line reference A-Z” in Help and Support Center in

Windows Server 2003.

**Dsmod.exe: Dsmod**
**Category**
This tool ships with Windows Server 2003.

**Version compatibility**
This tool is compatible with domain controllers running Windows Server 2003.

You can use this tool to modify an existing object of a specific type in the directory.

For more information about Dsmod, see “Dsmod” in “Command-line reference A-Z” in Help and Support Center in Windows Server 2003.

**Dsmove.exe: Dsmove**
**Category**
This tool ships with Windows Server 2003.

**Version compatibility**
This tool is compatible with domain controllers running Windows Server 2003.

You can use this tool to move a single object, within a domain, from its current location in the directory to a new location or to rename a single object without moving it in the directory tree.

For more information about Dsmove, see “Dsmove” in “Command-line reference A-Z” in Help and Support Center in Windows Server 2003.

**Dsquery.exe: Dsquery**
**Category**
This tool ships with Windows Server 2003.

**Version compatibility**
This tool is compatible with domain controllers running Windows Server 2003.

You can use this tool to query Active Directory according to specified criteria.

For more information about Dsquery, see “Dsquery” in “Command-line reference A-Z” in Help and Support Center in Windows Server 2003.

**Dsrm.exe: Dsrm**
**Category**
This tool ships with Windows Server 2003.

**Version compatibility**
This tool is compatible with domain controllers running Windows Server 2003.

You can use this tool to delete an object of a specific type or any general object from the directory.

For more information about Dsrm, see “Dsrm” in “Command-line reference A-Z” in Help and Support Center in Windows Server 2003.

**Ldifde.exe: Ldifde**
**Category**
This tool ships with Windows Server 2003.

**Version compatibility**
This tool is compatible with domain controllers running Windows Server 2003.

You can use Ldifde to create, modify, and delete directory objects on domain controllers running Windows Server 2003. You can also use Ldifde to extend the schema, export Active Directory user and group information to other applications or services, and populate Active Directory with data from other directory services.

For more information about Ldifde, see “Ldifde” in “Command-line reference A-Z” in Help and Support Center in Windows Server 2003.

**Ldp.exe: Active Directory Administration Tool**
**Category**
This tool ships with the Windows Server 2003 Support Tools.

**Version compatibility**

This tool is compatible with domain controllers running Windows Server 2003.

Ldp.exe is a Lightweight Directory Access Protocol (LDAP) graphical user interface (GUI) tool that you can use to perform operations, such as connect, bind, search, modify, add, and delete, against any LDAP-compatible directory, such as Active Directory. You can also use Ldp.exe to view objects that are stored in Active Directory, along with their metadata, such as security descriptors and replication metadata.

To find more information about Ldp.exe, see "Support Tools Help" in Tools and Settings Collection.

**Lusrmgr.msc: Local Users and Groups**
**Category**
Local Users and Groups is an MMC snap-in. This tool is included as part of the Windows Server 2003, Windows XP Professional, and Windows 2000 operating systems, and it is available in the Administration Tools Pack (Adminpak.msi).

**Version compatibility**
Local Users and Groups is supported for administering local groups on Windows Server 2003, Windows XP Professional, and Windows 2000.

Local Users and Groups manages users and groups of users for your computer. You can create new users and groups, add users to groups, remove users from groups, disable user and group accounts, and reset passwords.

You might need to be logged on as an administrator or a member of the Administrators group to perform some tasks.

For more information about Local Users and Groups, search for "Local users and groups overview" in Help and Support Center in Windows Server 2003.

**Net localgroup**
**Category**
This tool ships with Windows Server 2003.

**Version compatibility**
This tool is compatible with computers running Windows Server 2003.

You can use this tool to add, display, or modify local groups. When it is used without parameters, Net localgroup displays the name of the server and the names of local groups on the computer.

For more information about Net localgroup, search for "Net localgroup" in Help and Support Center in Windows Server 2003.

**Whoami**
**Category**
Whoami is a command-line tool in Windows Server 2003. This tool is also included in the Windows 2000 Resource Kit.

**Version compatibility**
Whoami is supported for Windows Server 2003 and Windows 2000.

You can use Whoami to display the complete contents of a user's access token in the command window.

Whoami displays the following:

- User name and security identifier (SID)
- Groups and their SIDs
- Privileges and their status (for example, enabled or disabled)
- Logon ID

To find more information about Whoami, see "Command-line reference A-Z" in Help and Support Center in Windows Server 2003.

**Related Information**
The following resources contain additional information that is relevant to this section:

- How Security Principals Work

- Support Tools Help in Tools and Settings Collection for information about Ldp.exe

# Security Identifiers Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Identifiers Technical Reference**
A security identifier (SID) is a unique name (alphanumeric value) that identifies an object — such as a user, group, or computer account — on a network. Each account or group is issued a unique SID when it is first created.

During the authentication process, a user is identified to the local or trusted domain by presentation of credentials, usually in the form of a user name and password. Assuming these credentials are acceptable, the system creates an access token for the user that contains the SID of the user (the primary SID) and the SIDs of all the domain groups of which the user is a member. The system uses this access token to determine whether to grant the user access to system resources.

SIDs are a fundamental building block of the Microsoft Windows Server 2003 security model, which controls the use of network resources through the interrelated mechanisms of authentication and authorization. The authorization and access control technologies in Windows Server 2003 use SIDS, along with other security features and components (such as access tokens, security descriptors, and access control lists), to help determine if an authenticated user has the correct authorization to access a resource.

**In this subject**


- What Are Security Identifiers?
- How Security Identifiers Work
- Security Identifiers Tools

# What Are Security Identifiers?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Are Security Identifiers?**

**In this section**


- Security Identifiers Dependencies and Interactions
- Related Information


A security identifier (SID) is a value of variable length that is used to uniquely identify a security principal or security group.

Each account or group has a unique SID that is issued by an authority, such as a Windows domain controller, and stored in a security database. The system generates the SID that identifies a particular account or group at the time the account or group is created. When a SID has been used as the unique identifier for a user or group, it can never be used again to identify another user or group.

Each time a user logs on, the system creates an access token for that user. The access token contains the user's SID, the SIDs for any groups the user belongs to, and the user's privileges. This token provides the security context for whatever actions the user executes on that computer.

In addition to the uniquely created, domain-specific SIDs that are assigned to specific users and groups, there are well-known SIDs that identify generic groups and generic users. For example, the Everyone and World SIDs identify a group that includes all users. Well-known SIDs have values that remain constant across all operating systems.

SIDs are a fundamental building block of the Windows security model. They work together with specific components of the authorization and access control technologies in the Windows Server 2003 security infrastructure to help protect access to network resources and provide a more secure computing environment.

**Security Identifiers Dependencies and Interactions**

The Windows Server 2003 operating system uses SIDs in the following authorization and access control components:

**Security principals**

A security principal can be a user or computer account, or a group of these accounts — any entity that the security system recognizes. User accounts can be used by human users as well as autonomous processes. A security principal is automatically assigned a security identifier when it is created.

Although all security principals incorporate a unique SID, it is important to note that SIDs are generated and stored differently depending on the type of account they are associated with:

- The SID for a local account or group is generated by the Local Security Authority (LSA) on the computer and stored with other account information in a secure area of the registry.
- The SID for a domain account or group is generated by the domain security authority and stored as an attribute of the User or Group object in Active Directory directory service.

**Access tokens**

One SID in an access token identifies the user. Other SIDs identify the security groups to which the user belongs.

**Security descriptors and ACLs**

One SID in an object's security descriptor identifies the object's owner. Another SID identifies the owner's primary group. Within the security descriptor are access control lists (ACLs) that contain access control entries (ACEs). Each ACE contains a SID that identifies the user or group for which access is allowed, denied, or audited.

These related security components are discussed in depth in separate sections of the Microsoft Windows Server 2003 Technical Reference.

The following diagram shows the relationship of SIDs to other key components of the authorization and access control model.

**Security Identifiers in the Authorization and Access Control Process**




**Related Information**

The following resources contain additional information that is relevant to this section:

- Security Principals Technical Reference
- Access Tokens Technical Reference
- Security Descriptors and Access Control Lists Technical Reference
- Permissions Technical Reference
- Logon and Authentication Technologies

# How Security Identifiers Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Security Identifiers Work**

**In This Section**


- Security Identifier Architecture
- Relative Identifier Allocation
- Security Identifiers and Globally Unique Identifiers
- Well-Known SIDs


Users refer to accounts by using the account name, but the operating system internally refers to accounts by their security identifiers (SIDs). For domain accounts, the SID of a security principal is created by concatenating the SID of the domain with a relative identifier (RID) for the account. SIDs are unique within their scope (domain or local) and are never reused.

The system generates the SID that identifies a particular account or group at the time the account or group is created. The SID for a local account or group is generated by the Local Security Authority (LSA) on the computer and it is stored with other account information in a secure area of the registry. The SID for a domain account or group is generated by the domain security authority and stored as an attribute of the User or Group object in Active Directory directory service.

For every local account and group, the SID is unique for the computer where it was created; no two accounts or groups on the computer ever share the same SID. Likewise, for every domain account and group, the SID is unique within an enterprise. This means that the SID for an account or group created in one domain will never match the SID for an account or group created in any domain in the enterprise.

SIDs always remain unique. Security authorities never issue the same SID twice, and they never reuse SIDs for deleted accounts. For example, if a user with a user account in a Windows Server 2003 domain leaves her job, an administrator deletes her account, including the SID that identifies the account. If she later returns to a different job at the same company, an administrator creates a new account, and the Windows Server 2003 operating system generates a new SID. The new SID does not match the old one; so none of the users access from her old account is transferred to the new account. Her two accounts represent two completely different security principals.

**Security Identifier Architecture**

A security identifier is a data structure in binary format that contains a variable number of values. The first values in the structure contain information about the SID structure. The remaining values are arranged in a hierarchy, like a telephone number, and identify the SID-issuing authority (for example, the Windows Server 2003 operating system), the SID-issuing domain, and a particular security principal or group. The following figure illustrates the structure of a SID.

**SID Structure**




The individual values of a SID are described in the table below.

**SID Components**

| Component | Description |
|---|---|
| Revision | Indicates the version of the SID structure that is used in a particular SID. The structure used in all SIDs that are created by a Windows Server 2003 operating system and earlier versions is revision level 1. |

| | |
|---|---|
| Identifier Authority | Identifies the highest level of authority that can issue SIDs for a particular type of security principal. For example, the identifier authority value in the SID for the Everyone group is 1 (World Authority). The identifier authority value in the SID for a specific Windows NT, Windows 2000, or Windows Server 2003 account or group is 5 (NT Authority). |
| Subauthorities | Holds the most important information in a SID. Contained in a series of one or more subauthority values. All values up to, but not including, the last value in the series collectively identify a domain in an enterprise. This part of the series is called the domain identifier. The last value in the series, which is called the relative identifier (RID), identifies a particular account or group relative to a domain. |

The components of a SID are easier to visualize when SIDs are converted from binary to string format by using standard notation:

```
S-R-X-Y1-Y2-Yn-1-Yn
```

In this notation, the components of a SID are as follows:

**SID Components in String Format**

| Component | Definition |
|---|---|
| S | Indicates that the string is a SID |
| R | Revision level |
| X | Identifier authority value |
| Y | A series of subauthority values, where *n* is the number of values |

The SID's most important information is contained in the series of subauthority values. The first part of the series (-Y1-Y2-Y*n*-1) is the domain identifier. This element of the SID becomes significant in an enterprise with several domains, because the domain identifier differentiates SIDs that are issued by one domain from SIDs that are issued by all other domains in the enterprise. No two domains in an enterprise share the same domain identifier.

The last item in the series of subauthority values (-Y*n*) is the relative identifier. It distinguishes one account or group from all other accounts and groups in the domain. No two accounts or groups in any domain share the same relative identifier.

For example, the SID for the built-in Administrators group is represented in standardized SID notation as the following string:

```
S-1-5-32-544
```

This SID has four components:

- A revision level (1)
- An identifier authority value (5, NT Authority)
- A domain identifier (32, Builtin)
- A relative identifier (544, Administrators)

SIDs for built-in accounts and groups always have the same domain identifier value: 32. This value identifies the domain Builtin, which exists on every computer that is running a version of the Windows NT, Windows 2000, or Windows Server 2003 operating systems. It is never necessary to distinguish one computer's built-in accounts and groups from another computer's built-in accounts and groups because they are local in scope: local either to a single computer; or in the case of domain controllers for a network domain, local to several computers that are acting as one. However, built-in accounts and groups do need to be distinguished from one another within the scope of the Builtin domain; therefore, the SID for each account and group has a unique relative identifier. A relative identifier value of 544 is unique to the built-in Administrators group. No other account or group in the Builtin domain has a SID with a final value of 544.

In another example, consider the SID for the global group Domain Admins. Every domain in an enterprise has a Domain Admins group, and the SID for each group is different. The following is the SID for the Domain Admins group in the Contoso domain (Contoso\Domain Admins):

```
S-1-5-21-1004336348-1177238915-682003330-512
```

The SID for Contoso\Domain Admins has:

- A revision level (1)
- An identifier authority (5, NT Authority)
- A domain identifier (21-1004336348-1177238915-682003330, Contoso)
- A relative identifier (512, Domain Admins)

The SID for Contoso\Domain Admins is distinguished from the SIDs for other Domain Admins groups in the same enterprise by its domain identifier: 21-1004336348-1177238915-682003330. No other domain in the enterprise uses this value as its domain identifier. The SID for Contoso\Domain Admins is distinguished from the SIDs for other accounts and groups created in the Contoso domain by its relative identifier, 512. No other account or group in the domain has a SID with a final value of 512.

**Relative Identifier Allocation**

Where accounts and groups are stored in an account database that is managed by a local Security Accounts Manager (SAM), it is fairly easy for the system to generate a unique relative identifier for each account and group that it creates on a stand-alone computer. The SAM on a stand-alone computer can simply keep track of the relative identifier values that it has used before and make sure that it never uses them again.

In a network domain, however, generating unique relative identifiers is a more complex process. Windows Server 2003 network domains can have several domain controllers. Each domain controller stores Active Directory account information. This means that, in a network domain, there are as many copies of the account database as there are domain controllers. In addition to this, every copy of the account database is a master copy. New accounts and groups can be created on any domain controller. Changes that are made to Active Directory on one domain controller are replicated to all other domain controllers in the domain. The process of replicating changes in one master copy of the account database to all other master copies is called a multimaster operation.

The process of generating unique relative identifiers is a single-master operation. One domain controller is assigned the role of relative identifier (RID) master, and it allocates a sequence of relative identifiers to each domain controller in the domain. When a new domain account or group is created in one domain controller's replica of Active Directory, it is assigned a SID. The relative identifier for the new SID is taken from the domain controller's allocation of relative identifiers. When its supply of relative identifiers begins to run low, the domain controller requests another block from the RID master.

Each domain controller uses each value in a block of relative identifiers only once. The RID master allocates each block of relative identifier values only once. This process assures that every account and group created in the domain has a unique relative identifier.

**Security Identifiers and Globally Unique Identifiers**

When a new domain user or group account is created, Active Directory stores the account's SID in the Object-SID (objectSID) property of a User or Group object. It also assigns the new object a globally unique identifier (GUID), which is a 128-bit value that is unique not only in the enterprise, but also across the world. GUIDs are assigned to every object that is created by Active Directory, not just User and Group objects. Each object's GUID is stored in its Object-GUID (objectGUID) property.

Active Directory uses GUIDs internally to identify objects. For example, the GUID is one of an object's properties that is published in the global catalog. Searching the global catalog for a User object's GUID produces results if the user has an account somewhere in the enterprise. In fact, searching for any object by Object-GUID might be the most reliable way of finding the object you want to locate. The values of other object properties can change, but the Object-GUID never changes. When an object is assigned a GUID, it keeps that value for life.

If a user moves from one domain to another, the user gets a new SID. The SID for a group object does not change because groups stay in the domain where they were created. However, if people move, their accounts can move with them. If an employee moves from North America to Europe, but stays in the same company, an administrator for the enterprise can move the employee’s User object from, for example, Contoso\Noam to Contoso\Europe. If the administrator does this, the User object for the account needs a new SID. The domain identifier portion of a SID that is issued in Noam is unique to Noam; so the SID for the user’s account in Europe has a different domain identifier. The relative identifier portion of a SID is unique relative to the domain; so if the domain changes, the relative identifier also changes.

When a User object moves from one domain to another, a new SID must be generated for the user account and stored in the

Object-SID property. Before the new value is written to the property, the previous value is copied to another property of a User object, SID-History (sIDHistory). This property can hold multiple values. Each time a User object moves to another domain, a new SID is generated and stored in the Object-SID property and another value is added to the list of old SIDs in SID-History. When a user logs on and is successfully authenticated, the domain authentication service queries Active Directory for all the SIDs that are associated with the user — the user's current SID, the user's old SIDs, and the SIDs for the user's groups. All these SIDs are returned to the authentication client and are included in the user's access token. When the user tries to gain access to a resource, any one of the SIDs in the access token, including one of the SIDs in SID-History, can allow or deny the user access.

If you allow or deny user's access to a resource based on their jobs, you should allow or deny access to a group, not an individual. That way, when users change jobs or move to other departments, you can easily adjust their access by removing them from certain groups and adding them to others. However, if you allow or deny an individual user access to resources, you probably want that user's access to remain the same no matter how many times the user's account domain changes. The SID-History property makes this possible. When a user changes domains, there is no need to change the access control list (ACL) on any resource. If an ACL has the user's old SID but not the new one, the old SID is still in the user's access token, listed among the SIDs for the user's groups, and the user is granted or denied access based on the old SID.

The reason for using SIDs, and not GUIDs, is for backward compatibility. Windows NT uses SIDs to identify users and groups in ACLs on resources; so you do not have to change ACLs on your resources when you upgrade. In Windows 2000 and Windows Server 2003, ACLs continue to identify users and groups by SID, not GUID — even ACLs on resources in Active Directory. For example, a user gains access to a Group Policy object based on one of the user's SIDs, not on the GUID for the User object.

**Well-Known SIDs**
The values of certain SIDs are constant across all systems. They are created during installation of the operating system or domain and are called well-known SIDs because they identify generic users or generic groups. For example, well-known SIDs identify the following users and groups:

**Everyone (S-1-1-0)**
The generic group Everyone automatically includes everyone who uses the computer, even anonymous guests. The identifier authority value for this SID is 1 (World Authority). It has only one subauthority value, 0 (Null RID).

**Creator Owner (S-1-3-0)**
The generic user Creator Owner is a placeholder in an inheritable access control entry (ACE). When the ACE is inherited, the system replaces the SID for Creator Owner with the SID for the object's current owner. The identifier authority value for this SID is 3 (Creator Authority). It has only one subauthority value, 0 (Null RID).

**Principal Self (S-1-5-10)**
The generic user Principal Self is a placeholder in an ACE on a User, Group, or Computer object in Active Directory. When you grant permission to Principal Self, you grant it to the security principal represented by the object. During an access check, the operating system replaces the SID for Principal Self with the SID for the security principal represented by the object. The identifier authority for this SID is 5 (NT Authority). It has only one subauthority value, 10 (Self RID).

There are universal well-known SIDs that are meaningful on all secure systems that use this security model, including operating systems other than Windows. In addition, there are well-known SIDs that are meaningful only on Windows operating systems.

The following table lists the universal well-known SIDs.

### Universal Well-known SIDs

| Universal Well-Known SID | Value | Identifies |
|---|---|---|
| Null SID | (S-1–0–0) | A group with no members. This is often used when a SID value is not known. |
| World | (S-1–1–0) | A group that includes all users. |
| Local | (S-1–2–0) | Users who log on to terminals that are locally (physically) connected to the system. |
| Creator Owner ID | (S-1–3–0) | A security identifier to be replaced by the security identifier of the user who created a new object. This SID is used in inheritable ACEs. |

| | | |
|---|---|---|
| Creator Group ID | (S-1–3–1) | A security identifier to be replaced by the primary-group SID of the user who created a new object. Use this SID in inheritable ACEs. |

The following table lists the predefined identifier authority constants. The first four values are used with universal well-known SIDs; the last value is used with well-known SIDs in Windows XP, Windows NT, Windows 2000, and Windows Server 2003 operating systems.

**Predefined Identifier Authority Constants**

| Identifier Authority | Value | SID String Prefix |
|---|---|---|
| SECURITY_NULL_SID_AUTHORITY | 0 | S-1–0 |
| SECURITY_WORLD_SID_AUTHORITY | 1 | S-1–1 |
| SECURITY_LOCAL_SID_AUTHORITY | 2 | S-1–2 |
| SECURITY_CREATOR_SID_AUTHORITY | 3 | S-1–3 |

The following RID values are used with universal well-known SIDs. The Identifier authority column shows the prefix of the identifier authority with which you can combine the RID to create a universal well-known SID.

**RID Values Used with Universal Well-known SIDs**

| Relative Identifier Authority | Value | Identifier Authority |
|---|---|---|
| SECURITY_NULL_RID | 0 | S-1–0 |
| SECURITY_WORLD_RID | 0 | S-1–1 |
| SECURITY_LOCAL_RID | 0 | S-1–2 |
| SECURITY_CREATOR_OWNER_RID | 0 | S-1–3 |
| SECURITY_CREATOR_GROUP_RID | 1 | S-1–3 |

The SECURITY_NT_AUTHORITY (S-1–5) predefined identifier authority produces SIDs that are not universal and are meaningful only in Windows XP, Windows NT, Windows 2000, and Windows Server 2003 installations. The following table lists the well-known SIDs.

**Well-Known Security Identifiers**

| SID | Display Name | Description |
|---|---|---|
| S-1-5-1 | Dialup | A group that includes all users who are logged on to the system by means of a dial-up connection. |
| S-1-5-2 | Network | A group that includes all users who are logged on by means of a network connection. Access tokens for interactive users do not contain the Network SID. |
| S-1-5-3 | Batch | A group that includes all users who have logged on by means of a batch queue facility, such as task scheduler jobs. |
| S-1-5-4 | Interactive | A group that includes all users who log on interactively. A user can start an interactive logon session by logging on directly at the keyboard, by opening a Remote Desktop connection from a remote computer, or by using a remote shell such as telnet. In each case, the user's access token contains the Interactive SID. If the user logs on using a Remote Desktop connection, the user's access token also contains the Remote Interactive Logon SID. |
| S-1-5-5-*X*-*Y* | Logon Session | A logon session. The *X* and *Y* values for these SIDs uniquely identify a particular logon session. |
| S-1-5-6 | Service | A group that includes all security principals that have logged on as a service. |

| | | |
|---|---|---|
| S-1-5-7 | Anonymous Logon | A user who has connected to the computer without supplying a user name and password.<br><br>The Anonymous Logon identity is different from the identity used by Internet Information Services (IIS) for anonymous Web access. IIS uses an actual account—by default, IUSR_*ComputerName*—for anonymous access to resources on a Web site. Strictly speaking, such access is not anonymous; the security principal is known even though unidentified people are using the account. IUSR_*ComputerName* (or whatever you name the account) has a password, and IIS logs the account on when the service starts. As a result, the IIS "anonymous" user is a member of Authenticated Users while Anonymous Logon is not. |
| S-1-5-8 | Proxy | Does not currently apply: this SID is not used in the Windows 2000 Server or Windows Server 2003 operating systems. |
| S-1-5-9 | Enterprise Domain Controllers | A group that includes all domain controllers in a forest of domains. |
| S-1-5-10 | Self | A placeholder in an ACE on a user, group, or computer object in Active Directory. When you grant permissions to Self, you grant them to the security principal represented by the object. During an access check, the operating system replaces the SID for Self with the SID for the security principal represented by the object. |
| S-1-5-11 | Authenticated Users | A group that includes all users and computers whose identities have been authenticated. Authenticated Users does not include Guest even if the Guest account has a password.<br><br>**Note:** This group includes authenticated security principals from any trusted domain, not just the current domain. |
| S-1-5-12 | Restricted | An identity used by a process that is executing in a restricted security context. In Windows XP and Windows Server 2003 operating systems, a software restriction policy can assign one of three security levels to executable code: unrestricted, restricted, or disallowed. When code executes at the restricted security level, the Restricted SID is added to the user's access token. |
| S-1-5-13 | Terminal Server User | A group that includes all users who log on to a server with Terminal Services enabled: specifically, a server that is in Terminal Services version 4.0 application compatibility mode. |
| S-1-5-14 | Remote Interactive Logon | A group that includes all users who log on to the computer by using a Remote Desktop Connection. This group is a subset of the Interactive group. Access tokens that contain the Remote Interactive Logon SID also contain the Interactive SID. |
| S-1-5-18 | System (or LocalSystem) | An identity that is used locally by the operating system and by services configured to log on as LocalSystem.<br><br>System is a hidden member of Administrators. That is, any process running as System has the SID for the built-in Administrators group in its access token.<br><br>When a process that is running locally as System accesses network resources, it does so by using the computer's domain identity. Its access token on the remote computer includes the SID for the local computer's domain account plus SIDs for security groups that the computer is a member of, such as Domain Computers and Authenticated Users. |
| S-1-5-19 | LocalService | An identity used by services that are local to the computer, have no need for extensive local access and do not need authenticated network access. Services that run as LocalService access local resources as ordinary users, and access network resources as anonymous users. As a result, a service that runs as LocalService has significantly less authority than a service that runs as LocalSystem — both locally and on the network. |
| S-1-5-20 | NetworkService | An identity used by services that have no need for extensive local access but do need authenticated network access. Services running as NetworkService access local resources as ordinary users and access network resources using the computer's identity. As a result, a service that runs as NetworkService has the same network access as a service that runs as LocalSystem but has significantly reduced local access. |

| | | |
|---|---|---|
| S-1-5-*domain*-500 | Administrator | A user account for the system administrator. Every computer has a local Administrator account and every domain has a domain Administrator account.<br><br>The Administrator account is the first account created during operating system installation. The account cannot be deleted, disabled, or locked out, but it can be renamed.<br><br>By default, the Administrator account is a member of the Administrators group and cannot be removed from that group. |
| S-1-5-*domain*-501 | Guest | A user account for people who do not have individual accounts. Every computer has a local Guest account and every domain has a domain Guest account.<br><br>By default, Guest is a member of Everyone and Guests. The domain Guest account is also a member of the Domain Guests and Domain Users groups.<br><br>Unlike Anonymous Logon, Guest is a real account and can be used to log on interactively. By default, the account is enabled on Windows XP and disabled on all editions of the Windows Server 2003 operating system. The Guest account does not require a password, but can have one. |
| S-1-5-*domain*-502 | krbtgt | A user account that is used by the Key Distribution Center (KDC) service. The account exists only on domain controllers. |
| S-1-5-*domain*-512 | Domain Admins | A global group whose members are authorized to administer the domain. By default, the Domain Admins group is a member of the Administrators group on all computers that have joined the domain, including domain controllers.<br><br>Domain Admins is the default owner of any object that is created in the domain's Active Directory by any member of the group. If members of the group create other objects, such as files, the default owner is the Administrators group. |
| S-1-5-*domain*-513 | Domain Users | A global group that includes all users in a domain. When you create a new User object in Active Directory, the user is automatically added to this group. |
| S-1-5-*domain*-514 | Domain Guests | A global group that, by default, has only one member, the domain's built-in Guest account. |
| S-1-5-*domain*-515 | Domain Computers | A global group that includes all computers that have joined the domain, excluding domain controllers. |
| S-1-5-*domain*-516 | Domain Controllers | A global group that includes all domain controllers in the domain. New domain controllers are added to this group automatically. |
| S-1-5-*domain*-517 | Cert Publishers | A global group that includes all computers that host an enterprise certification authority.<br><br>Cert Publishers are authorized to publish certificates for User objects in Active Directory. |
| S-1-5-*root domain*-518 | Schema Admins | A group that exists only in the forest root domain. It is a universal group if the domain is in native mode and a global group if the domain is in mixed mode. The Schema Admins group is authorized to make schema changes in Active Directory. By default, the only member of the group is the Administrator account for the forest root domain. |

| | | |
|---|---|---|
| S-1-5-*root domain*-519 | Enterprise Admins | A group that exists only in the forest root domain. It is a universal group if the domain is in native mode, a global group if the domain is in mixed mode.<br><br>The Enterprise Admins group is authorized to make changes to the forest infrastructure, such as adding child domains, configuring sites, authorizing DHCP servers, and installing enterprise Certificate Authorities.<br><br>By default, the only member of Enterprise Admins is the Administrator account for the forest root domain. The group is a default member of every Domain Admins group in the forest. |
| S-1-5-*domain*-520 | Group Policy Creator Owners | A global group that is authorized to create new Group Policy objects in Active Directory. By default, the only member of the group is Administrator.<br><br>Objects that are created by members of Group Policy Creator Owners are owned by the individual user who creates them. In this way, the Group Policy Creator Owners group is unlike other administrative groups, such as Administrators and Domain Admins. Objects that are created by members of these groups are owned by the group rather than by the individual. |
| S-1-5-*domain*-553 | RAS and IAS Servers | A domain local group. By default, this group has no members. Computers that are running the Routing and Remote Access service are added to the group automatically.<br><br>Members of this group have access to certain properties of User objects, such as Read Account Restrictions, Read Logon Information, and Read Remote Access Information. |

The following RIDs are relative to each domain.

**RIDs Relative to Each Domain**

| RID | Identifies |
|---|---|
| DOMAIN_USER_RID_ADMIN | The administrative user account in a domain. |
| DOMAIN_USER_RID_GUEST | The guest-user account in a domain. Users who do not have an account can automatically log on to this account. |
| DOMAIN_GROUP_RID_ADMINS | The domain administrators' group. This account exists only on systems running Windows 2000 Server or Windows NT Server. |
| DOMAIN_GROUP_RID_USERS | A group containing all user accounts in a domain. All users are automatically added to this group. |
| DOMAIN_GROUP_RID_GUESTS | The guest-group account in a domain. |
| DOMAIN_GROUP_RID_COMPUTERS | The domain computers' group. All computers in the domain are members of this group. |
| DOMAIN_GROUP_RID_CONTROLLERS | The domain controllers' group. All Domain controllers in the domain are members of this group. |
| DOMAIN_GROUP_RID_CERT_ADMINS | The certificate publishers' group. Computers running Certificate Services are members of this group. |
| DOMAIN_GROUP_RID_SCHEMA_ADMINS | The schema administrators' group. Members of this group can modify the Active Directory schema. |
| DOMAIN_GROUP_RID_ENTERPRISE_ADMINS | The enterprise administrators' group. Members of this group have full access to all domains in the Active Directory forest. Enterprise administrators are responsible for forest-level operations such as adding or removing new domains. |
| DOMAIN_GROUP_RID_POLICY_ADMINS | The policy administrators' group. |

The following table provides examples of domain-relative RIDs that are used to form well-known SIDs for local groups.

**Examples of RIDs for Local Groups**

| RID | Identifies |
| --- | --- |
| DOMAIN_ALIAS_RID_ADMINS | A local group that is used for administration of the domain. |
| DOMAIN_ALIAS_RID_USERS | A local group that represents all users in the domain. |
| DOMAIN_ALIAS_RID_GUESTS | A local group that represents guests of the domain. |
| DOMAIN_ALIAS_RID_POWER_USERS | A local group that is used to represent a user or set of users who expect to treat a system as if it were their personal computer rather than as a workstation for multiple users. |
| DOMAIN_ALIAS_RID_ACCOUNT_OPS | A local group that exists only on systems running Windows 2000 Server or Windows NT Server. This local group permits control over non-administrator accounts. |
| DOMAIN_ALIAS_RID_SYSTEM_OPS | A local group that exists only on systems running Windows 2000 Server or Windows NT Server. This local group performs system administrative functions, not including security functions. It establishes network shares, controls printers, unlocks workstations, and performs other operations. |
| DOMAIN_ALIAS_RID_PRINT_OPS | A local group that exists only on systems running Windows 2000 Server or Windows NT Server. This local group controls printers and print queues. |
| DOMAIN_ALIAS_RID_BACKUP_OPS | A local group that is used for controlling assignment of file backup-and-restoring user rights. |
| DOMAIN_ALIAS_RID_REPLICATOR | A local group responsible for copying security databases from the primary domain controller to the backup domain controllers. These accounts are used only by the system. |
| DOMAIN_ALIAS_RID_RAS_SERVERS | A local group that represents remote access and servers running Internet Authentication Service (IAS). This group permits access to various attributes of user objects. |
| DOMAIN_ALIAS_RID_PREW2KCOMPACCESS | A local group that exists only on systems running Windows 2000 Server and Windows Server 2003. This group provides access rights equal to anonymous access under Windows NT, which is Everyone access. |

# Security Identifiers Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Identifiers Tools**
The following tools are associated with security identifiers (SIDs).

**Getsid.exe: Get Security ID**
**Category**
Included with Windows Server 2003 Support Tools.

**Version compatibility**
This tool is compatible with domain controllers running Windows Server 2003.

Use this command-line tool to compare the user SIDs of two accounts. For example, you can use Getsid.exe to compare account SIDs between domain controllers when you suspect user database corruption.

For more information about Getsid.exe, see "Support Tools Help" in the Tools and Settings Collection.

**Sidwalk.exe, Sidwalk.msc, and Showaccs.exe: SIDWalker Security Administration tools**
**Category**
Included with Windows Server 2003 Support Tools.

**Version compatibility**
These tools are compatible with domain controllers running Windows Server 2003.

Use these tools to help you migrate or merge computers from a Windows NT 4.0 resource domain to a Windows Server 2003 domain. You can also use these tools to help manage server resources that you move between domains. The SIDWalker toolset allows you to change the access control lists (ACLs) on objects that were previously owned by accounts that have been moved or deleted. You can use these tools to either delete or replace every occurrence of an old security identifier (SID) with the corresponding new SID.

For more information about the SIDWalker Security Administration tools, see "Support Tools Help" in the Tools and Settings Collection.

**Netdom.exe: Windows Domain Manager**
**Category**
A command-line tool that is included in Windows Server 2003 Support Tools.

**Version compatibility**
This tool is compatible with Windows Server 2003.

You can use Netdom.exe to apply or remove SID filtering. SID filtering is set on outgoing external or forest trusts to prevent malicious users who have domain-level or enterprise-level administrator access in a trusted forest from granting (to themselves or to other user accounts in their forest) elevated user rights to a trusting forest. SID filtering does this by preventing misuse of the SID-History attribute on security principals in the trusted forest.

For more information about SID filtering, see "Security Considerations for Trusts."

For more information about Netdom.exe, see "Support Tools Help" in the Tools and Settings Collection.

**Ntdsutil.exe: Ntdsutil**
**Category**
A command-line tool that is included in Windows Server 2003.

**Version compatibility**
This tool is compatible with Windows Server 2003.

You can use Ntdsutil with the **security account management** command to check for and resolve duplicate SIDs in Active Directory — in the unlikely situation that duplicates are issued. Only experienced administrators should use this tool.

For more information about Ndtsutl, see "Command-Line References" in the Tools and Settings Collection.

# Data Security Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data Security Technologies**
In a distributed computing environment, data exists in many locations. In some situations, data is stored on a centralized server. Other situations may call for a decentralized data storage scenario, where each user maintains his own data on his own computer. The most common scenario is a hybrid of these two, where some data is stored on a centralized server and retrieved by client computers, while other data is stored only on local computers and never shared.

All of these scenarios present some security issues to consider. In most environments, at the very least some of this data must be kept secure. To help provide data security in a variety of scenarios, there are several security technologies that can help protect sensitive data during storage and network transmission.

**Data Security Architecture**
There is no single data security component that protects all data against all forms of attack. There are, however, several technologies that can act in concert to provide a comprehensive security infrastructure that protects against many types of attacks. This type of attack defense is known as "defense in depth" and is the most appropriate method to ensure that sensitive data is not compromised. The data security components that provide the various data safeguards are described later in this overview.

**Data Security Components**
There are three main data security components in the Microsoft Windows operating system: Internet Protocol security (IPSec), Encrypting File System (EFS), and Syskey.

**Internet Protocol Security**
Internet Protocol security (IPSec) is a framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services. The implementation of IPSec in the Microsoft Windows Server 2003 operating system is based on standards developed by the Internet Engineering Task Force (IETF) IPSec working group.

IPSec provides aggressive protection against private network and Internet attacks through end-to-end security. The only computers that must know about IPSec protection are the sender and receiver in the communication. In Windows Server 2003, IPSec provides the ability to protect communication between workgroups, local area network computers, domain clients and servers, branch offices (which might be physically remote), extranets, and roving clients.

Although IPSec provides strong data security when data is transmitted on an IP network, it does not protect the data when it is stored by a computer on a disk drive. To help protect stored data, EFS must be used.

**Encrypting File System**
With Encrypting File System (EFS), you can store data securely. EFS does this by encrypting data in selected NTFS file system files and folders.

Because EFS is integrated with the file system, it is easy to manage, difficult to attack, and transparent to the user. This is particularly useful for securing data on computers that may be vulnerable to theft, such as mobile computers. However, EFS is useful in helping to protect sensitive data stored on any computer.

EFS is designed to store data securely on local computers. As such, it does not support the secure transmission of files over a network. Other technologies, such as Internet Protocol security (IPSec), can be used in conjunction with EFS to provide a larger solution.

**Syskey**
Syskey is used to manage a computer's startup key, which protects all of a system's master keys. Users have their own master key, which is a cryptographic key that is used to protect other keys used by applications and services. At system startup, a startup key is used to decrypt the user's master key, which is then used to encrypt all of the private keys on the computer, including private keys that are used for EFS. Startup keys are automatically generated and used for computers in a domain, but must be manually configured on stand-alone computers. Startup key security can be increased by storing the key on removable media or by requiring a system startup password. You can make those changes with Syskey.

Syskey is useful in protecting EFS keys that are stored on the local hard drive. Although it does not protect the EFS data itself, it protects the keys that can be used to access that data. By using Syskey in combination with EFS, you employ more than one layer of security between an attacker and your encrypted data.

**Data Security Deployment Scenarios**
There are several scenarios where data security technologies will help provide important safeguards for data. The following are two of the most common.

**Portable laptops during travel**
Laptops are, by their very nature, designed to be taken on trips. Modern laptops often contain large hard drives and powerful processors, enabling owners to work while they travel. These laptops can store any information the owner wants, including sensitive information such as trade secrets, undisclosed financial reports, personnel data, and so forth.

Unfortunately, laptops get stolen frequently. Without physical control of the computer, the user has no way to protect the data it contains. An attacker has numerous ways to extract the data from the laptop. If the data is unprotected, it's a very simple matter.

To protect against these types of attacks, you should use both EFS and Syskey. EFS encrypts sensitive files so that they are cryptographically protected. This provides very strong protection, but the EFS key that is used for decryption must also be protected. Syskey is used to provide that strong protection of the EFS private keys. In combination, these two technologies make it computationally infeasible for an attacker to decrypt your data.

**Insecure network infrastructure**
Most modern building tenants share hollow walls with their neighbors. This is often a security concern because network wiring frequently runs through these common walls. An attacker can easily connect their own computer to a network without being detected by simply making a hole on the other side of the wall and splicing into the network wiring.

The insecure network infrastructure may not always be a building's hollow walls. For example, remote users might use virtual private networks (VPNs) to communicate with the office. This VPN communication frequently uses the Internet to relay the data. Because the Internet is an insecure network itself, the data could easily be compromised during transmission.

This type of threat can be easily mitigated by using IPSec. You should use IPSec in this scenario to encrypt all network traffic that might contain sensitive data, such as file transfers and database interaction. Less sensitive network transmissions, such as identification broadcasts and Web surfing, might not require IPSec protection. IPSec can be configured to include or exclude any desired data to provide the appropriate amount of security.

# Encrypting File System Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Encrypting File System (EFS) is a Microsoft technology that makes it possible for users to encrypt files on their computers and control who can decrypt or recover these files. When data is encrypted, it cannot be read or modified by intruders, even if they gain physical access to the hard disk on a computer.

**In this subject**


- What is Encrypting File System?
- How Encrypting File System Works
- Encrypting File System Tools and Settings

# What is Encrypting File System?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is Encrypting File System?**

**In this section**


- Problems Solved by EFS
- EFS Scenarios
- EFS Interactions and Dependencies


Encrypting File System (EFS) provides the core file encryption technology for storing files on NTFS file system volumes. Security features such as logon authentication and file permissions protect network resources from unauthorized access. However, anyone with physical access to a computer such as a stolen laptop can bypass security in the operating system by installing a new operating system on that computer. In this way, sensitive data can be exposed. Encrypting sensitive files by means of EFS adds another layer of security. When files are encrypted, data in them is protected even if an attacker has full access to the data storage system on the computer.

**Problems Solved by EFS**

EFS addresses security concerns raised by tools available on other operating systems that allow users to access files stored on an NTFS volume without an access check. With EFS, data in NTFS files is encrypted on disk. The encryption technology used is public key–based and runs as an integrated system service, making it easy to manage, difficult to attack, and transparent to the user. If a user attempting to access an encrypted NTFS file has the private key associated with that file, the user is able to open the fileandwork with it transparently as a normal document. A user without the private key to the file is denied access.

Because EFS is tightly integrated with NTFS, file encryption and decryption are transparent. When users open a file, it is decrypted by EFS as data is read from disk. When users save the file, EFS encrypts the data as it is written to disk. Authorized users might not even realize that the files are encrypted because they can work with the files as they normally do.

**EFS Scenarios**

EFS allows users to store confidential information on a computer when people who have physical access to a computer might otherwise compromise that information, intentionally or unintentionally. EFS is especially useful for securing sensitive data on portable computers, such as those commonly used by consultants or sales people who frequently conduct business away from the offices of an organization. It can also be useful on computers shared by several users, such as in banks or medical facilities. Both kinds of systems are susceptible to attack by techniques that circumvent the restrictions of access control lists (ACLs). In a shared system, an attacker can gain access by starting a different operating system. An attacker can also steal a computer, remove a hard disk, place the hard disk in another system, and gain access to the files stored on the disk. Files encrypted by EFS, however, appear as unintelligible characters when the attacker does not have the decryption key.

In its default configuration, EFS enables users to start encrypting files with no administrative effort. From the point of view of the user, encrypting a file is simply a matter of setting a file attribute. The encryption attribute can also be set for a file folder. This means that any file created in or added to the folder is automatically encrypted.

Only authorized users and designated data recovery agents (DRAs) can decrypt encrypted files. Other system accounts that have permissions for a file — even the Take Ownership permission — cannot open the file without authorization. Even the local administrator account cannot open the file if that account is not designated as a DRA. If an unauthorized user tries to open an encrypted file, access is denied.

**What Can Be Encrypted**

Individual files and file folders (or sub-folders) on NTFS volumes can be encrypted. Although it is common to refer to file folders as having the encryption attribute set as "encrypted," the folder itself is not encrypted. When encryption is set for a folder, EFS automatically encrypts all new files created in the folder. All files copied or moved into the folder Offline Files can also be encrypted.

**Note**

- When offline files are encrypted, the entire offline files database is encrypted rather than individual files. Individual files do not display the encryption attribute. The database is encrypted using the startup key for the system.

System files and any files in the *systemroot* folder or its subfolders cannot be encrypted. No files or directories in a roaming user profile can be encrypted. A file cannot be both compressed and encrypted. Being compressed does not prevent encryption, but when the file is encrypted, it is uncompressed.

**Note**

- Encrypting the temp directory can cause some applications to not function properly, and therefore is not recommended.

**Recovering Data Protected with EFS**

EFS supports data recovery in the sense that it makes it possible for designated data recovery agents (DRAs) to decrypt files that user has encrypted. A DRA is established by default on Windows 2000 systems. The DRA is optional on Windows XP Professional and Windows Server 2003 in order to provide organizations with greater flexibility in implementing data recovery strategies. With Windows XP Professional and Windows Server 2003, one or more DRAs can be established for individual computers, for a domain, or for a combination of individual computers and the domain. However, in no case is the private key of any user revealed to a recovery agent.

**EFS Interactions and Dependencies**

EFS is available only on computers running Windows Server 2003, Windows XP Professional, Windows 2000 Server, and Windows 2000 Professional, and using the NTFS file system.

However, there are differences between how EFS is implemented on Windows Server 2003 and Windows XP Professional compared to Windows 2000. In particular, systems running Windows Server 2003 and Windows XP Professional do not require the existence of a DRA, whereas EFS is disabled on systems running Windows 2000 unless there is a DRA.

On computers running Windows 2000, an EDRP is automatically configured on stand-alone computers when the local administrator initially logs on. This policy makes the local Administrator account the DRA for the computer. Recovery keys and a self-signed file recovery certificate are automatically generated. The local Administrator can change the default EDRP if needed.

In Windows 2000 Active Directory environments, a default recovery policy is configured for a domain when the first domain controller is set up. The default recovery policy uses a self-signed certificate to make the domain administrator account the DRA. The domain Administrator can change the default EDRP if needed.

The following interactions and dependencies also affect how EFS can be used:

- You can use EFS to encrypt or decrypt data on a remote computer, but you cannot use it to encrypt data sent over the network unless Web Distributed Authoring and Versioning (WebDAV) encryption is configured. Otherwise, IPSec should be used to protect data across the network.

- You cannot encrypt system files or folders.

- You cannot encrypt compressed files and folders until you decompress them.

- If you encrypt an entire folder, temporary copies of encrypted files that it contains are also encrypted.

- If you copy a file into an encrypted folder encrypts the file, but if you move it into the folder the file remains encrypted or unencrypted, just as it was before you copied the file.

- If you move or copy EFS files to another file system, encryption is removed, but if you back them up, the encryption is preserved.

- When you encrypt a file, other file permissions are unaffected. An administrator, for instance, can still delete your EFS file even though the administrator cannot open it.

- You must have EFS certificates available to use EFS. If you do not currently have a public key infrastructure (PKI), you can use self-signed certificates. If you have a PKI, you can configure them to provide EFS certificates.

EFS in Windows Server 2003 and Windows XP Professional includes the following features that are not available on systems running Windows 2000:

- Additional users can be authorized to access encrypted files.
- Certificates can be checked for revocation status when encrypted files are shared. (revocation is checked only when a user is added to an encrypted file)
- Offline files can be encrypted.
- The Advanced Encryption Standard (AES) and DES (3DES) encryption algorithms are supported.
- Encrypted files can be stored in Web folders using WebDav.
- EFS can be used with Windows Server 2003 clusters.
- File recovery policy can be configured with greater flexibility.

# How Encrypting File System Works

Updated: November 11, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Encrypting File System Works**

**In this section**


- Encrypting File System Architecture
- Physical Structure of an Encrypted File
- Encrypting File System Processes and Interactions
- Related Information


Encrypting File System (EFS) provides file encryption for NTFS files stored on disk. EFS supplements other security features such as logon authentication or file permissions that protect network resources from unauthorized access. Encrypting sensitive files by using EFS adds a layer of security that protects data even if an attacker gains full access to the computer where the data is stored.

EFS functions equally well on standalone systems running Windows XP Professional and Windows Server 2003, and in Active Directory domain environments. On standalone systems, EFS is automatically enabled but can be disabled by editing local security policy. In Active Directory environments, EFS can be managed using the domain security policy.

In a domain controller, a default recovery policy is automatically put in place when the administrator logs on to the computer for the first time, making the enterprise administrator the data recovery agent for computers that are members of the domain. For standalone computers, a default data recovery agent is not created. As a security best practice, if you decide to use the default data recovery agent, you should export the private key and save it separately and then remove the private key from the domain controller. However, it is strongly recommended that you create at least one data recovery agent in case the original account or credentials used to encrypt data with EFS become unavailable. For example, you can create a key pair on a smart card, and then set that certificate as the data recovery agent instead of the default data recovery agent. This enables you to recover data by using the smart card without exposing the private key to other software.

This section is divided into three subsections:

- **Encrypting File System Architecture** describes and illustrates the relationships between the operating system components that make up EFS in both user mode and kernel mode.
- **Physical Structure of an Encrypted File** describes the fields that are added to a data file when it is encrypted and the role of each field.
- **Encrypting File System Processes and Interactions** describes the processes used by EFS to encrypt, decrypt, recover and share data.

**Encrypting File System Architecture**

In order to provide transparent encryption, EFS relies on several components that operate at the user mode and kernel mode levels of the operating system:

**EFS Components**

| User Mode Components | Kernel Mode Components |
|---|---|

| | |
|---|---|
| • Efsadu .dll<br>• Feclient.dll<br>• Local Security Authority (LSA)<br>• CryptoAPI<br>• Cryptographic service providers (CSPs) | • NTFS<br>• Ksecdd<br>• FIPS.sys |

The following diagram illustrates the relationships among these components.

**Relationship of EFS Components**




One additional kernel mode component, the WebDAV Redirector, is required to manage encrypted files on Web Distributed Authoring and Versioning (WebDAV) file shares.

For more information about the encryption and decryption processes, see "Encrypting File System Processes and Interactions." The following sections describe the components of EFS in detail.

**Efsadu**

The Efsadu dynamic link library (DLL) provides the EFS user interface. Efsadu is called when a user access the Advanced Properties page for a file or folder to modify file or folder encryption options. The options processed through Efsadu include whether to apply the encryption to file or folder only, or to a folder and all subfolders and files that it contains. EFSadu also contains the user interface for adding users to an encrypted file.

**Feclient**

When a user modifies EFS options for a file or folder, or when an application attempts to access an encrypted file on an NTFS volume, the Win32 application programming interface (API) passes the resulting EFS-related calls to the Feclient DLL. Feclient then calls the EFS remote procedure call (RPC) interfaces in the LSA.

This two-stage process for calling the EFS routines in LSA is required because EFS supports remote storage of encrypted files on servers running Windows Server 2003 or servers running Windows 2000. The role of Feclient is to determine which client computer or server has the requested file and to then contact the LSA on that computer or server.

**Note**

- Users can also modify EFS options from the command line by using Cipher.exe. For more information about Cipher.exe, see "Encrypting File System Tools and Settings."

**Local Security Authority**

The Local Security Authority (LSA) performs the majority of user mode EFS operations. The LSA responds to requests from the kernel to create a file encryption key (FEK) and encapsulate it with the proper user and recovery agent public keys. The FEK is returned to the kernel and used to encrypt a file. The LSA also performs the reverse operation of unwrapping a FEK and returning it to the kernel to decrypt a file.

During the encapsulation process, the recovery policy for the local computer is checked to determine whether the FEK should also be encapsulated by one or more recovery agent public keys. During decryption, the recovery policy is checked again to see if additional or new recovery certificates have been added; if so, the LSA reencapsulates the FEK using the new recovery agent public keys. The LSA also checks to determine whether the EFS certificate belonging to the current user has changed, which would also require the FEK to be reencapsulated.

To improve performance, the LSA maintains a cache of handles (or pointers) to private keys used for encapsulating and extracting FEKs. These handles point to private keys that remain in the cryptographic service provider (CSP) and are, therefore, protected from unauthorized access.

**CryptoAPI**
CryptoAPI consists of a set of functions that allow applications to encrypt or digitally sign data while providing protection for private key data. Applications can use the functions in CryptoAPI independently of the underlying implementation.

EFS uses CryptoAPI for user mode cryptographic operations; CryptoAPI also provides underlying security services for secure channels and code signing. CryptoAPI supports public key and symmetric-key operations such as key generation, key management and secure storage, key exchange, encryption, decryption, hashing, digital signatures, and verification of signatures. In addition, application developers can use CryptoAPIservices to add cryptography to their 32-bit Windows applications.

**Cryptographic Service Providers**
The public-private key pairs for EFS users and recovery agent accounts are obtained from the Microsoft Base cryptographic service provider (CSP), also called the RSA base provider, or the Microsoft Enhanced CSP. These CSPs are included with Windows Server 2003 and Windows XP Professional.

### Note

- Only Microsoft CSPs for RSA are supported for EFS. EFS does not support the Microsoft smart card CSP.

Depending on the CSP you select, you can select one of three encryption algorithms to encrypt and decrypt the contents of files:

- **Advanced Encryption Standard (AES) algorithm.** This is the default encryption standard for Windows Server 2003 and Windows XP Professional with Service Pack 1 installed. AES uses a 256-bit key for encryption and decryption.

- **DESX.** An enhanced variant of the Data Encryption Standard (DES).

- **Triple-DES (3DES).** 3DES, which is compliant with Federal Information Processing Standards (FIPS 140-1 Level 1), offers significantly stronger encryption using a 128-bit or 168-bit key. 3DES is enabled through a Group Policy setting. All EFS FIPS mode encryption processes are completed using 3DES.

The encryption algorithms available to EFS depend on which version of Windows is running on the computer:

- Windows XP with SP1 or later can encrypt or decrypt files using DESX, 3DES, or AES.

- Versions of Windows XP Professional earlier than SP1 can use the DESX or the Triple-DES (3DES) algorithm for EFS encryption and decryption.

- Windows 2000 Professional and Windows 2000 Server use the DESX algorithm for EFS encryption and decryption.

### Note

- When 3DES is enabled in FIPS mode, it is used as the encryption algorithm for IP Security as well as for EFS. When 3DES is enabled, all new encryption processes are completed by using 3DES, but DESX and 3DES are always available for decryption, regardless of the encryption policy.

If a user attempts to view files that have been encrypted using an encryption algorithm not supported by the operating system, the file appears to be corrupted or filled with random characters.

**NTFS**
The NTFS driver contains the majority of the EFS code in the kernel. This driver responds to file system controls (FSCTLs) from the LSA to encrypt or decrypt a file. When a file is opened, the driver passes the wrapped FEK to Ksecdd to be unwrapped by the LSA. The driver then decrypts the data when it is read from the file and encrypts the data when it is written to the file.

The EFS API in the NTFS driver also provides a set of functions for reading and writing encrypted files in raw mode. Raw mode allows direct reading and writing of the ciphertext (text that has been encrypted using an encryption key). This allows the Windows backup utilities to read and write the encrypted data for backup purposes, while preserving the confidentiality of the data.

To improve performance for files opened multiple times in quick succession, the NTFS driver maintains a cache of unwrapped FEKs that it receives from the LSA. By default, FEKs are cached for five seconds; the cache time is configurable from two to thirty seconds.

**Ksecdd**
Ksecdd is a very thin component that NTFS calls to communicate with the LSA. Ksecdd is used to set up local procedure call (LPC) communications to the LSA.

**FIPS.sys**
Microsoft Kernel Mode Cryptographic Module (FIPS.sys) is a FIPS 140-1 Level 1 compliant, general-purpose, software-based cryptographic module located at the kernel mode level of the Windows operating system. It runs as a kernel mode export driver (kernel-mode DLL) and includes several different cryptographic algorithms in an easy-to-use cryptographic module accessible by other kernel mode drivers. It can be linked into other kernel mode services to permit the use of FIPS 140-1 Level 1 compliant cryptography.

### Note

- Federal Information Processing Standard (FIPS) 140-1 is a security implementation designed for certifying cryptographic software. FIPS 140-1–validated software is required by the U.S. Government and requested by other prominent organizations.

In Windows 2000, Windows XP Professional, and Windows Server 2003, both IPSEC and EFS use FIPS.sys to encrypt traffic packet data and file contents, respectively, provided that you use FIPS-compliant algorithms.

**WebDAV Redirector**
The Web Distributed Authoring and Versioning (WebDAV) client redirector is responsible for managing encrypted files on WebDAV file shares. The redirector works by retrieving the file from the server and then caching it locally while edits are performed. When the file is closed, the redirector stores the file on the server. While the file is cached locally, the redirector calls the NTFS driver to read and write the file.

When the file is stored on the server, the client redirector performs a backup read on the locally cached file and takes the data and stores it on the server. When a file is retrieved from the server, the redirector performs a backup write to the cached file which takes the encrypted data and restores it to a normal, but encrypted, NTFS file.

**Physical Structure of an Encrypted File**
EFS uses a combination of public key and symmetric key encryption to ensure that files are protected from all but the most computationally infeasible methods of attack.

EFS follows the standard cryptographic procedure of key encipherment. Data is encrypted using a symmetric file encryption key (FEK) for speed; the FEK is then secured asymmetrically for maximum security. When a user requests that a file be encrypted, EFS uses a uniquely generated FEK to encrypt a file and then encrypts the FEK itself by using the public key from the public key certificate belonging to the user. The encrypted FEK is stored in a file header called EFS metadata. When a user requests decryption, EFS decrypts the FEK using the user's private key, and then uses the FEK to decrypt the file.

An encrypted file contains encrypted data and a header with fields to store copies of the encrypted FEK for authorized users and designated data recovery agents. The structure of an encrypted file is shown in the following figure.

### Encrypted File Structure




**Data Decryption Field**

An encrypted file contains a minimum of one stored FEK, the FEK encrypted by using the public key belonging to the initial encryptor. The storage field for this encrypted FEK is called the data decryption field(DDF). Additionally, if an EFS-encrypted file is shared, a copy of the FEK is encrypted by using the public key belonging to the newly authorized user, and the encrypted FEK is stored in another DDF. In this way, there is a separate data decryption file for every user of an encrypted file.

**Data Recovery Field**

If the effective security policy for a computer designates one or more data recovery agents, copies of the FEK are encrypted for each data recovery agent by using the public key of each data recovery agent, and the copies are stored in another file header field called the data recovery field (DRF).

**Key Storage**

When an encrypted file is saved, Windows XP Professional and Windows Server 2003 automatically provide four levels of encryption to protect the encryption keys that protect the file, and a fifth level, the startup key, can be configured. These levels operate as follows:

- **Level One:** EFS generates a unique FEK for every file that is encrypted.
- **Level Two:** EFS uses the public key in the EFS certificate of the user to encrypt the FEK. The certificate containing the public key is stored by default in the certificate store of the computer. The corresponding private key, used to decrypt the FEK, is stored in an encrypted form in the profile of the corresponding user or data recovery agent account in the RSA folder.
- **Level Three:** The Data Protection API generates the master key for the user that is used to encrypt the private keys belonging to the user.
- **Level Four:** The Data Protection API generates a symmetric password encryption key, derived from the credentials of the user, which encrypts the master key belonging to the user.
- **Level Five:** A startup key (also called the syskey) can be used to protect all master keys as well as a variety of other secrets that are stored on computers. At system startup, the startup key is used to encrypt all of the private keys on the computer, including private keys that are used for EFS. Startup keys are automatically generated and used for computers in a domain, but must be manually configured on stand-alone computers. Startup key security can be increased by storing the key on removable media or by requiring a system startup password.

Windows XP Professional and Windows Server 2003 provide several additional configuration options that increase security. To prevent loss of access to your master keys on stand-alone computers, you can create a password reset disk (PRD). You can also implement options to prevent your computer from entering hibernation mode and to delete paging files at system shutdown so that data is not needlessly exposed.

**Certificate and Public Key Storage**

Windows Server 2003 and Windows XP Professional store public key certificates for a user in the personal certificate store of that user. Certificates can be stored in plaintext because they are public information, and they are digitally signed by

certification authorities to protect against tampering.

User certificates are located in the Documents and Settings*username*\Application Data\Microsoft\SystemCertificates\My\Certificates folder for each user profile. These certificates are written to the personal store of the user in the system registry each time the user logs on to the computer. For roaming profiles, the certificates of users are located in a shared folder on a server configured by the domain administrator and follow users when they log on to different computers in the domain.

**Encrypting File System Processes and Interactions**
The following processes are associated with EFS:

- Encryption
- Decryption
- File recovery
- Multi-user sharing

**Encrypting a File with EFS**
When a user encrypts an existing file, the following process occurs:

1. EFS opens the file for exclusive access.
2. EFS creates a temporary file in the current directory and carries out the following series of operations:
    a. Each data stream in the source file is copied to the temporary file for backup purposes.
    b. Each data stream in the original file is next truncated to a length of zero, and its length is then set back to its original size. This essentially deletes all the data in the stream.
    c. EFS writes the encryption metadata to the original file. At this point EFS has the plaintext data in a temporary file and an empty source file that is marked encrypted (because of the presence of the EFS metadata).
3. A file encryption key (FEK) is randomly generated and used to encrypt the file with the AES, 3DES, or DESX algorithms, depending on the version of the operating system and the effective security policy.
4. A data decryption field (DDF) is created to contain the FEK encrypted by using the user's public key. In addition:
    - If a data recovery agent has been designated through Group Policy, a data recovery field (DRF) is created to contain the FEK encrypted by using RSA and the public key of the data recovery agent.
    - If there are multiple data recovery agents, a copy of the FEK is encrypted by using the public key of each data recovery agent, and a DRF is created to store each encrypted FEK.

    **Note**

    - The file recovery property in the certificate is an example of an enhanced key usage (EKU) field. An EKU extension and extended property specify and limit the valid uses of a certificate. File Recovery is one of the EKU fields defined by Microsoft as part of the Microsoft public key infrastructure (PKI).
5. EFS writes the encrypted data, along with the DDF and any DRFs, back to the file. Because symmetric encryption does not add data to a file, file sizes change very little after encryption. The encryption metadata, consisting primarily of encrypted FEKs, is quite small (depending on how many users are added to the file).

    **Note**

- The size of EFS metadata is limited to 256 KB.

6. The plaintext temporary file is deleted.

**Note**

- Data from deleted files might not be erased when the file is deleted. The **cipher /w** command can be used to remove data from available unused disk space on the entire volume.

When a user saves a file to a folder that has been configured for encryption, the encryption process is similar, except that no temporary file is created.

The following figure shows the EFS encryption process when a data recovery agent is being used.

**EFS Encryption with a Data Recovery Agent**




**Decrypting a File with EFS**

Conversion to plaintext follows a similar process as for encryption, in that a temporary file is created and the source file is written to it. The source file is then truncated and the data is read from the temporary file and written back to the source file in the clear. When an application accesses an encrypted file, decryption proceeds as follows:

1. NTFS recognizes that the file is encrypted, retrieves the DDF, and passes it to EFS.

2. EFS retrieves the private key belonging to the user from the profile of the user and uses it to decrypt the DDF and obtain the FEK.

3. EFS then uses the FEK to decrypt sections of the file as needed by the application.

**Note**

- Because EFS uses cipher block chaining, when an application opens a file, only those sections of the file that the application is using are decrypted. The behavior is different if the user removes the encryption attribute from the file. In this case, the entire file is decrypted and rewritten as plaintext.

4. EFS returns the decrypted data to NTFS, which then sends the data to the requesting application.

The following figure illustrates the decryption process:

**EFS Decryption**




**Recovering a File**

Data recovery is similar to the decryption process except:

- The DRF is used instead of the DDF.
- EFS decrypts the DRF with the private key of a data recovery agent to recover the FEK.

The recovery process is as follows:

1. NTFS recognizes the file as encrypted and sends a request to EFS.
2. EFS retrieves the DRF (instead of the DDF) and decrypts the DRF (instead of the DDF) with the private key of the data recovery agent (instead of the private key of the user) to obtain the FEK.
3. EFS uses the FEK to decrypt the file
4. NTFS then completes the file request and sends the data to the requesting application.

The following figure illustrates the file recovery process.

**EFS File Recovery**




**Authorizing Users to Access Encrypted Files**

Users can share encrypted files with other local, domain, and trusted domain users. Authorizing user access to encrypted files is a separate process from sharing files for network access by using share-level security and access control lists (ACLs).

Therefore, sharing access to encrypted files involves the configuration of both file system ACL permissions, and EFS permissions.

Because there is no method for issuing a certificate and assigning a key pair for a group, groups cannot be authorized for access: Only individual user accounts can be authorized for access to an encrypted file.

After a file has been encrypted, additional users can be authorized to access encrypted files by opening the Advanced Attributes dialog box from the Properties dialog box for the file.

EFS certificates can also be imported to the profile of a user. If the EFS certificate is self-signed, it is added to the Trusted People container. If the EFS certificate was issued by a certification authority CA, it is added to the Other People container. Certificates added to either container appear in the Select User dialog box.

**Note**

- If you import a certificate belonging to a user, you have the option to manually select its location. If you import a self-signed certificate, be sure to add it to Trusted People. Certificates in Trusted People are the only certificates that are not chain validated, because these certificates are already trusted. Self-signed certificates always fail chain validation because no CA was involved in certifying them, so placing them anywhere other than in Trusted People makes them unusable.

Before a user can be authorized to access an encrypted file and be added to the file, EFS needs to determine whether the certificate can be trusted. When a certificate for a user is selected, EFS attempts to validate the certificate chain. If the chain validation fails, CryptoAPI also checks to see if the certificate is in Trusted People store. If not, the certificate cannot be used.

Sharing an encrypted file can be illustrated with a simple example. Assume that Alice already has access to the encrypted file. She is either the original encryptor, or she has been authorized for access. Alice wants to authorize Bob to access the file. To accomplish this, EFS performs the following steps:

1. EFS obtains the private key belonging to Alice from her user profile to decrypt the FEK contained in the data decryption field of the file.

2. EFS obtains the public key belonging to Bob from his certificate and uses it to encrypt the FEK.

3. The encrypted FEK belonging to Bob is stored in a new data decryption field with the file.

**Note**

- At no time in this process is the file itself decrypted, so it is not directly at risk. However, the FEK is briefly decrypted. But because EFS performs this operation in nonpaged memory, the decrypted FEK is never paged and so is never exposed.

The following Figure shows what happens when Alice shares an encrypted file with Bob.

**EFS Access Authorization**




Users electing to share encrypted files should remember the following:

- Shared EFS files are not file shares. If authorized users need to access shared EFS files over the network, a file share or a Web folder is required. Alternatively, users can establish remote sessions with computers that store encrypted files by using Terminal Services.

- Any user who is authorized to decrypt a file can authorize other users to access the file. Granting access is not limited to the file owner. Caution users to share files only with trusted accounts, because those accounts can authorize other accounts. Removing the Write permission from a user or group of users can prevent this problem, but it also prevents the user or group from modifying the file.

- EFS sharing requires that the users who will be authorized to access the encrypted file have EFS certificates. These certificates can be located in roaming profiles or in the user profiles on the computer on which the file to be shared is stored , or they can be stored in and retrieved from Active Directory.

- EFS sharing of an encrypted file often means that the file will be accessed across the network. To enhance the security of file data, it is best if Web folders are used for encrypted file storage whenever possible.

- If a user chooses to remotely access an encrypted file stored on a file share and to authorize other users to access the file, the authorization process and requirements are the same as on the local computer. Additionally, EFS must impersonate the user in order to perform this operation, and all of the requirements for remote EFS operations on files stored on file shares apply.

- If a user chooses to remotely access an encrypted file stored on a Web folder and to authorize other users to access the file, the file is automatically transmitted to the local computer in ciphertext. The authorization process takes place on the local computer with the same requirements as for encrypted files stored locally.

**Validating Certificate Chains**

Imported self-signed certificates for EFS are automatically placed in Trusted People, and there is no certificate chain to validate, so the user can be added to the encrypted file.

**Note**

- Certificates are not verified for valid chains or revocation status during encryption or decryption.

However, if the certificate was signed by a CA, EFS attempts to build a certification chain and validate the certificate only when

adding a user to the file. If the chain ends with an untrusted root CA, EFS will not use the certificate, and the user is not added to the file.

For more information about how chain validation works, see the "Certificate Services Technical Reference."

**Related Information**

The following resources contain additional information that is relevant to this section.

- How NTFS Works
- How Certificates Work

# Encrypting File System Tools and Settings

Updated: September 10, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encrypting File System Tools and Settings**

**In this section**


- Encrypting File System Tools
- Encrypting File System Registry Entries
- Encrypting File System Group Policy Settings
- Related Information


**Encrypting File System Tools**

The following tools are associated with Encrypting File System.

**Cipher.exe: Cipher**

**Category**

Cipher is an operating system command-line tool.

**Version compatibility**

This tool is compatible with Windows 2000, Windows XP, Windows Server 2003.

Allows a user or administrator to display or alter the encryption of files. In addition to encrypting or decrypting a file or folder, Cipher can be used to update the file encryption keys or the keys of the data recovery agent (DRA) should there be a change in the data recovery policy.

When used with the **/w** switch, Cipher can also remove data from portions of the volume it can access that have not been allocated to files or directories. Cipher does not lock the drive, so other programs can obtain space on the drive which cipher cannot erase. Because the /w option writes to a large portion of the volume, it might take a long time to complete and should only be used when necessary.

For more information about Cipher, see “Command-Line References” in the Tools and Settings Collection.

**Efsinfo.exe: Encrypting File System Information**

**Category**

Encrypting File System Information is a Windows Server 2003 command-line tool.

**Version compatibility**

This tool is compatible with Windows 2000, Windows XP, Windows Server 2003.

Encrypting File System Information displays information about files and folders encrypted with Encrypting File System (EFS) on partitions that use the NTFS file system. Options include displaying encryption information about the files and folders in the current folder, recovery agent information, and certificate thumbnail information.

For more information about EFSinfo, see “Command-Line References” in the Tools and Settings Collection.

**Xcopy.exe: Xcopy**

**Category**

Xcopy is a command line tool that ships with Windows Server 2003 and Windows XP Professional.

**Version compatibility**

This tool is compatible with Windows 2000, Windows XP Professional, Windows Server 2003.

Encrypted files are copied from Web folders in the same way that plaintext files are copied from file shares. The Xcopy command does not require any special parameters. The file is transmitted in ciphertext and remains encrypted on the local computer if possible. The encryption status for files copied from Web folders is the same as for files copied locally.

Encrypted files are copied from Web folders in the same way that plaintext files are copied from file shares. The Copy and Xcopy commands do not require any special parameters. The file is transmitted in ciphertext and remains encrypted on the

local computer if possible. The encryption status for files copied from Web folders is the same as that for files copied locally.

For more information about Xcopy, see "Command-Line References" in the Tools and Settings Collection.

**SecPol.msc: Local Security Settings Snap-in**
**Category**
Local Security Settings is a Microsoft Management Console (MMC) snap-in that ships with Windows Server 2003, Windows 2000 Server, and Windows XP Professional.

**Version compatibility**
This tool is compatible with Windows 2000, Windows XP Professional, Windows Server 2003.

Local Security Settings is compatible with Windows Server 2003 and Windows 2000 Server, and can be used to EFS data recovery agents on computers running Windows Server 2003, Windows XP Professional, and Windows 2000.

To find more information about the Local Security Settings snap-in, see "Security Policy Settings."

**Encrypting File System Registry Entries**
The following registry entries are associated with Encrypting File System.

- HKEY_CURRENT_USER\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys
- HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys
- HKEY_CURRENT_USER\Software\Microsoft\Cryptography\CertificateTemplateCache
- HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer

The following information is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_CURRENT_USER\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys**
The following registry entries are created only when EFS is used for the first time.

**CertificateHash**
**Registry path**
HKEY_CURRENT_USER\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys\CertificateHash

**Version**
Windows Server 2003, Windows 2000, and Windows XP.

This registry entry contains the certificate hash used by the current user to encrypt and decrypt data.

**Flag**
**Registry path**
HKEY_CURRENT_USER\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys\Flag

**Version**
Windows Server 2003, Windows 2000, and Windows XP.

This registry entry identifies the encryption algorithm used to encrypt and decrypt new EFS files.

| **Note** |
| --- |
| For more information about the encryption algorithms that can be used with EFS, see "How Encrypting File System Works" in the Technical Reference. |

**KeyCacheValidationPeriod**
**Registry path**
HKEY_CURRENT_USER\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys\KeyCacheValidationPeriod

**Version**

Windows Server 2003, Windows 2000, and Windows XP.

This registry entry identifies the length of time in seconds a cached EFS certificate is assumed to be valid before revalidation is required.

**HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys**
The following registry entries control how EFS is used on the local machine.

**UserCacheSize**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys\UserCacheSize

**Version**
Windows Server 2003, Windows 2000, and Windows XP.

This registry entry identifies the number of user EFS certificates to cache.

**EFSDomainGPOCreated**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys\EFSDomainGPOCreated

**Version**
Windows Server 2003, Windows 2000, and Windows XP.

This registry entry records whether domain-based EFS Group Policy settings are in effect.

**EFSBlob**
**Registry path**
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\EFS\CurrentKeys\EFSBlob

**Version**
Windows Server 2003, Windows 2000, and Windows XP.

This registry entry contains data about the file recovery certificates that have been registered for the system.

**HKEY_CURRENT_USER\Software\Microsoft\Cryptography\CertificateTemplateCache**
The template cache is used to cache copies of the domain's certificate templates which are distributed by Active Directory Domain Services. When your domain administrator configures certificate templates for your organization this registry key will contain sub-keys for each template in the domain.

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\EFS**
The following registry entries are added when domain Group Policy is used to enable or disable EFS.

**EFSConfiguration**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\EFS\EFSConfiguration

**Version**
Windows Server 2003, Windows 2000, and Windows XP

This registry setting makes it possible to disable and re-enable EFS on a local computer using Group Policy.

**LastGoodEFSConfiguration**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\EFS\LastGoodEFSConfiguration

**Version**
Windows Server 2003, Windows 2000, and Windows XP

This registry setting is added when EFS is disabled through domain Group Policy. This registry setting is deleted if EFS is re-enabled using domain Group Policy.

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer**
**Registry path**
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer

**Version**
Windows Server 2003, Windows 2000, and Windows XP

This registry setting can be used to display **Encrypt** and **Decrypt** options on the Windows Explorer shortcut menu when a user

right-clicks a file or folder.

**Encrypting File System Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with Encrypting File System.

**Group Policy Settings Associated with Encrypting File System**

| Group Policy Setting | Description |
|---|---|
| Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Autoenrollment | Can be used to enroll certificates automatically, renew expired certificates, update pending certificates, and remove certificates that have been revoked. In addition, this setting can be used to block certificate auto-enrollment. |
| Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Encrypting File System | Can be used to disable or enable encryption of files using EFS. The default setting is enabled. In addition, this setting can be used to register one or more data recovery agents for use with EFS or to remove data recovery agents if they are no longer desired. |
| Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Automatic Certificate Request Settings | Can be used to configure automatic certificate request settings for a specific certificate template for a domain by using the Automatic Certificate Request Setup Wizard. The request will be processed automatically at the first occurrence of any of the following: a user logs on, Group Policy is refreshed, or a computer joins the domain and is subject to a Group Policy setting. |

To find more information about these Group Policy settings, see "Group Policy Settings Reference" in the Tools and Settings Collection.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Certificates Tools and Settings" in Certificates Technical Reference
- "Certificate Services Tools and Settings" in Certificate Services Technical Reference

# IPSec Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Internet Protocol security (IPSec) in the Microsoft Windows Server 2003 operating system helps protect networks from active and passive attacks by securing IP packets through the use of packet filtering, cryptographic security services, and the enforcement of trusted communications. IPSec is integrated within the security framework of Windows Server 2003, which allows you to use Active Directory domains to provide identity and manage trust relationships between users and computers in different departments within an organization.

This subject provides a high-level description of IPSec that includes the scenarios for which it is intended and not intended.

**In This Subject**


- What Is IPSec?
- How IPSec Works
- IPSec Tools and Settings

# What Is IPSec?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Is IPSec?**

**In This Subject**


- IPSec Scenarios
- IPSec Dependencies
- IPSec and ICF
- Related Information


Internet Protocol security (IPSec) is a framework of open standards for helping to ensure private, secure communications over Internet Protocol (IP) networks through the use of cryptographic security services. IPSec supports network-level data integrity, data confidentiality, data origin authentication, and replay protection. Because IPSec is integrated at the Internet layer (layer 3), it provides security for almost all protocols in the TCP/IP suite, and because IPSec is applied transparently to applications, there is no need to configure separate security for each application that uses TCP/IP.

IPSec helps provide defense-in-depth against:

- Network-based attacks from untrusted computers, attacks that can result in the denial-of-service of applications, services, or the network
- Data corruption
- Data theft
- User-credential theft
- Administrative control of servers, other computers, and the network.

You can use IPSec to defend against network-based attacks through a combination of host-based IPSec packet filtering and the enforcement of trusted communications.

IPSec is integrated with the Windows Server 2003 operating system and it can use the Active Directory directory service as a trust model. You can use Group Policy to configure Active Directory domains, sites, and organizational units (OUs), and then assign IPSec policies as required to Group Policy objects (GPOs). In this way, IPSec policies can be implemented to meet the security requirements of many different types of organizations.

This section describes the solution that IPSec is intended to provide by providing information about core IPSec scenarios, IPSec dependencies, and related technologies.

The following figure shows an Active Directory-based IPSec policy being distributed to two IPSec peers and IPSec-protected communications being established between those two peers.

**Two IPSec Peers Using Active Directory-based IPSec Policy**




The MicrosoftWindows implementation of IPSec is based on standards developed by the Internet Engineering Task Force (IETF) IPSec working group. For a list of relevant IPSec RFCs, see the "Related Information" section later in this subject.

**IPSec Scenarios**

IPSec is a general-purpose security technology that can be used to help secure network traffic in many scenarios. However, you must balance the need for security with the complexity of configuring IPSec policies. Additionally, due to a lack of suitable standards, IPSec is not appropriate for some types of connectivity. This section describes IPSec scenarios that are recommended, IPSec scenarios that are not recommended, and IPSec scenarios that require special consideration.

**Recommended Scenarios for IPSec**

IPSec is recommended for the following scenarios:

- Packet filtering
- End-to-end security between specific hosts
- End-to-end traffic through a Microsoft Internet Security and Acceleration (ISA) Server-secured network address translator
- Secure server
- Layer Two Tunneling Protocol (L2TP) over IPSec (L2TP/IPSec) for remote access and site-to-site virtual private network (VPN) connections
- Site-to-site IPSec tunneling with non-Microsoft IPSec gateways

**Packet Filtering**

IPSec can perform host-based packet filtering to provide limited firewall capabilities for end systems. You can configure IPSec to permit or block specific types of unicast IP traffic based on source and destination address combinations and specific protocols and specific ports. For example, nearly all the systems illustrated in the following figure can benefit from packet filtering to restrict communication to only specific addresses and ports. You can strengthen security by using IPSec packet filtering to control exactly the type of communication that is allowed between systems.

**Filtering Packets by Using IPSec**




As illustrated in this figure:

- The internal network domain administrator can assign an Active Directory-based IPSec policy (a collection of security settings that determines IPSec behavior) to block all traffic from the perimeter network (also known as a demilitarized zone [DMZ], demilitarized zone, or screened subnet).
- The perimeter network domain administrator can assign an Active Directory-based IPSec policy to block all traffic to the

internal network.

- The administrator of the computer running Microsoft SQL Server on the internal network can create an exception in the Active Directory-based IPSec policy to permit structured query language (SQL) protocol traffic to the Web application server on the perimeter network.

- The administrator of the Web application server on the perimeter network can create an exception in the Active Directory-based policy to permit SQL traffic to the computer running SQL Server on the internal network.

- The administrator of the Web application server on the perimeter network can also block all traffic from the Internet, except requests to TCP port 80 for the HyperText Transfer Protocol (HTTP) and TCP port 443 for HTTPS (HTTP over Secure Sockets Layer/Transport Layer Protocol [SSL/TLS]), which are used by Web services. This provides additional security for traffic allowed from the Internet in case the firewall was misconfigured or compromised by an attacker.

- The domain administrator can block all traffic to the management computer, but allow traffic to the perimeter network.

You can also use IPSec with the IP packet-filtering capability or NAT/Basic Firewall component of the Routing and Remote Access service to permit or block inbound or outbound traffic, or you can use IPSec with the Internet Connection Firewall (ICF) component of Network Connections, which provides stateful packet filtering. However, to ensure proper Internet Key Exchange (IKE) management of IPSec security associations (SAs), you must configure ICF to permit UDP port 500 and port 4500 traffic needed for IKE messages.

**End-to-End Security Between Specific Hosts**

IPSec establishes trust and security from a unicast source IP address to a unicast destination IP address (end-to-end). For example, IPSec can help secure traffic between Web servers and database servers or domain controllers in different sites. As shown in the following figure, only the sending and receiving computers need to be aware of IPSec. Each computer handles security at its respective end and assumes that the medium over which the communication takes place is not secure. The two computers can be located near each other, as on a single network segment, or across the Internet. Computers or network elements that route data from source to destination are not required to support IPSec.

**Securing Communications Between a Client and a Server by Using IPSec**




The following figure shows domain controllers in two forests that are deployed on opposite sides of a firewall. In addition to using IPSec to help secure all traffic between domain controllers in separate forests, as shown in the figure, you can use IPSec to help secure all traffic between two domain controllers in the same domain and between domain controllers in parent and child domains.

**Securing Communications Between Two Domain Controllers in Different Forests by Using IPSec**




**End-to-End Traffic Through an ISA-Secured Network Address Translator**

Windows Server 2003 supports IPSec NAT Traversal (NAT-T). IPSec NAT-T allows traffic to be secured by IPSec and also to be translated by a network address translator. For example, you can use IPSec transport mode to help secure host-to-host traffic through a computer that is running ISA Server and that is functioning as a network address translator if ISA (or any other NAT device) does not need to inspect the traffic between the two hosts. IPSec transport mode is used to protect traffic between hosts and it can provide security between computers that are on the same local area network (LAN) or connected by private wide area network (WAN) links. In the following figure, a computer running Windows Server 2003 and Microsoft Internet Security and Acceleration (ISA) Server is functioning as a network address translator. The IPSec policy on Server A is configured to secure traffic to the IP address of Server B, while the IPSec policy on Server B is configured to secure traffic to the external IP address of the computer running ISA Server.

**Securing Communications Through an ISA-Secured NAT by Using IPSec NAT-T**




**Secure Server**

You can require IPSec protection for all client computers that access a server. In addition, you can set restrictions on which computers are allowed to connect to a server running Windows Server 2003. The following figure shows IPSec in transport mode securing a line of business (LOB) application server.

**Securing an Application Server by Using IPSec**




In this scenario, an application server in an internal corporate network must communicate with clients running Windows 2000 or Windows XP Professional; a Windows Internet Name Service (WINS) server, Domain Name System (DNS) server, and Dynamic Host Configuration Protocol (DHCP) server; Active Directory domain controllers; and a non-Microsoft data backup

server. The users on the client computers are company employees who access the application server to view their personal payroll information and performance review scores. Because the traffic between the clients and the application server involves highly sensitive data, and because the server should only communicate with other domain members, the network administrator uses an IPSec policy that requires ESP encryption and communication only with trusted computers in the Active Directory domain.

Other traffic is permitted as follows:

- Traffic between the WINS server, DNS server, DHCP server, and the application server is permitted because WINS servers, DNS servers, and DHCP servers must typically communicate with computers that run on a wide range of operating systems, some of which might not support IPSec.

- Traffic between Active Directory domain controllers and the application server is permitted, because using IPSec to secure communication between domain members and their domain controllers is not a recommended usage.

- Traffic between the non-Microsoft data backup server and the application server is permitted because the non-Microsoft backup server does not support IPSec.

**L2TP/IPSec for Remote Access and Site-to-Site VPN Connections**

You can use L2TP/IPSec for all VPN scenarios. This does not require the configuration and deployment of IPSec policies. Two common scenarios for L2TP/IPSec are securing communications between remote access clients and the corporate network across the Internet and securing communications between branch offices.

### Note

- Windows IPSec supports both IPSec transport mode and tunnel mode. Although VPN connections are commonly referred to as "tunnels," IPSec transport mode is used for L2TP/IPSec VPN connections. IPSec tunnel mode is most commonly used to help protect site-to-site traffic between networks, such as site-to-site networking through the Internet.

**L2TP/IPSec for remote access connections**

A common requirement for organizations is to secure communications between remote access clients and the corporate network across the Internet. Such a client might be a sales consultant who spends most of the time traveling, or an employee working from a home office. In the following figure, the remote gateway is a server that provides edge security for the corporate intranet. The remote client represents a roaming user who requires regular access to network resources and information. An ISP is used as an example to demonstrate the path of communication when the client uses an ISP to access the Internet. L2TP/IPSec provides a simple, efficient way to build a VPN tunnel and help protect the data across the Internet.

**Securing Remote Access Clients by Using L2TP/IPSec**

**L2TP/IPSec for site-to-site VPN connections**

A large corporation often has multiple sites that require communication — for example, a corporate office in New York and a sales office in Washington. In this case, L2TP/IPSec provides the VPN connection and helps protect the data between the sites. In the following figure, the router running Windows Server 2003 provides edge security. The routers might have a leased line, dial-up, or other type of Internet connection. The L2TP/IPSec VPN tunnel runs between the routers only and provides protected communication across the Internet.

**Establishing an L2TP/IPSec VPN Tunnel Between Sites**

**Site-to-Site IPSec Tunneling with Non-Microsoft Gateways**

For interoperability with gateways or end systems that do not support L2TP/IPSec or Point-to-Point Tunneling Protocol (PPTP) VPN site-to-site connections, you can use IPSec in tunnel mode. When IPSec tunnel mode is used, the sending gateway encapsulates the entire IP datagram by creating a new IP packet that is then protected by one of the IPSec protocols. The following figure illustrates site-to-site IPSec tunneling.

**Establishing an IPSec Gateway-to-Gateway Tunnel Between Sites**




In this figure, traffic is being sent between a client computer in a vendor site (Site A) and a File Transfer Protocol (FTP) server at the corporate headquarters site (Site B). Although an FTP server is used for this scenario, the traffic can be any unicast IP traffic. The vendor uses a non-Microsoft IPSec-enabled gateway, while corporate headquarters uses a gateway running Windows Server 2003. An IPSec tunnel is used to secure traffic between the non-Microsoft gateway and the gateway running Windows Server 2003.

**Scenarios for Which IPSec Is Not Recommended**

IPSec policies can be quite complex to configure and manage. Additionally, IPSec can incur performance overhead to establish and maintain secure connections, and it can incur network latency. In some deployment scenarios, the lack of standard methods for user authentication and address assignment make IPSec an unsuitable choice. Because IPSec depends on IP addresses for establishing secure connections, you cannot specify dynamic IP addresses. It is often necessary for a server to have a static IP address in IPSec policy filters. In large network deployments and in some mobile user cases, using dynamic IP addresses at both ends of the connection can increase the complexity of IPSec policy design. For these reasons, IPSec is not recommended for the following scenarios:

- Securing communication between domain members and their domain controllers
- Securing all traffic in a network
- Securing traffic for remote access VPN connections using IPSec tunnel mode

**Securing Communication Between Domain Members and their Domain Controllers**

Using IPSec to help secure traffic between domain members (either clients or servers) and their domain controllers is not recommended because:

- If domain members were to use IPSec-secured communication with domain controllers, increased latency might occur, causing authentication and the process of locating a domain controller to fail.
- Complex IPSec policy configuration and management is required.
- Increased load is placed on the domain controller CPU to maintain SAs with all domain members. Depending on the number of domain members in the domain controller's domain, such a load might overburden the domain controller.

**Securing All Traffic in a Network**

In addition to reduced network performance, using IPSec to help secure all traffic in a network is not recommended because:

- IPSec cannot secure multicast and broadcast traffic.
- Traffic from real-time communications, applications that require Internet Control Message Protocol (ICMP), and peer-to-peer applications might be incompatible with IPSec.
- Network management functions that must inspect the TCP, UDP, and other protocol headers are less effective, or cannot function at all, due to IPSec encapsulation or encryption of IP payloads.

**Securing Traffic for Remote Access VPN Connections by Using IPSec Tunnel Mode**

IPSec tunnel mode is not a recommended technology for remote access VPN connections, because there are no standard methods for user authentication, IP address assignment, and name server address assignment. Using IPSec tunnel mode for gateway-to-gateway VPN connections is possible using computers running Windows Server 2003. But because the IPSec tunnel is not represented as a logical interface over which packets can be forwarded and received, routes cannot be assigned to use the IPSec tunnel and routing protocols do not operate over IPSec tunnels. Therefore, the use of IPSec tunnel mode is only recommended as a VPN solution for site-to-site VPN connections in which one end of the tunnel is a non-Microsoft VPN server or security gateway that does not support L2TP/IPSec. Instead, use L2TP/IPSec or PPTP for remote access VPN connections.

**IPSec Uses That Require Special Considerations**
The following scenarios merit special consideration, because they introduce an additional level of complexity for IPSec policy configuration and management:

- Securing traffic over IEEE 802.11wireless networks
- Securing traffic in home networking scenarios
- Securing traffic in environments that use dynamic IP addresses

**Securing Traffic Sent over 802.11 Networks**
You can use IPSec transport mode to protect traffic sent over 802.11 wireless networks. However, IPSec is not the recommended solution for providing security for corporate 802.11 wireless LAN networks. Instead, it is recommended that you use either 802.11 Wired Equivalent Privacy (WEP) encryption or Wi-Fi Protected Access (WPA) and IEEE 802.1X authentication.

To use IPSec to help secure traffic sent over 802.11 networks, you must ensure that client computers and servers support IPSec. Configuration management and trust are also required on client computers and servers when IPSec is used. Because many computers on a network do not support IPSec or are not managed, it is not appropriate to use IPSec alone to protect all 802.11 corporate wireless LAN traffic.

**Securing Traffic in Home Networking Scenarios**
Although IPSec is not optimized for use in general home networking scenarios, when network security administrators deploy IPSec with appropriate scripts and support tools, it can be used effectively on home computers for specific scenarios.

IPSec can be used to connect home computers to a corporate intranet for remote access. Network security administrators can use scripts and support tools to deploy IPSec on the home computers of employees who require secure connectivity to the corporate network. For example, an administrator can use a Connection Manager profile to deploy an L2TP/IPSec-based VPN connection on home computers. Employees can then establish IPSec-secured connections across the Internet to the corporate network by using the VPN client built-in to Network Connections.

### Note

- In some cases, non-Microsoft VPN or firewall clients might disable the IPSec service, which is required for IPSec to function. If you encounter this problem, it is recommended that you contact the VPN or firewall vendor.

IPSec is not recommended for end users in general home networking scenarios for the following reasons:

- The IPSec policy configuration user interface (IP Security Policy Management) is intended for professional network security administrators, rather than for end users. Improper policy configuration can result in blocked communications, and if problems occur, built-in support tools are not yet available to aid end users in troubleshooting.
- Some home networking applications use broadcast and multicast traffic, for which IPSec cannot negotiate security.
- Many home networking scenarios use a wide range of dynamic IP addresses.
- Many home networking scenarios involve the use of a network address translator. To use IPSec across a NAT, both IPSec peers must support IPSec NAT-T.

**Securing Traffic in Environments That Use Dynamic IP Addresses**
IPSec depends on IP addresses for establishing secure connections, and it is often necessary for a server to have a static IP address in IPSec policy filters. In large network deployments and in some mobile user cases, using dynamic IP addresses at

both ends of the connection can increase the complexity of IPSec policy design.

**IPSec Dependencies**
There is no single optimal environment for IPSec. However, there are dependencies that are critical to the successful deployment of IPSec. This section describes how the following two IPSec dependencies affect the deployment of IPSec:

- Active Directory (if your deployment requires the use of Active Directory-based IPSec policies, rather than local IPSec policies)
- Successful mutual authentication

**Active Directory**
For organizations with large numbers of computers that must be managed in a consistent way, it is best to distribute IPSec policies by using Group Policy to configure Active Directory domains, sites, and organizational units (OUs), and then assigning IPSec policies as required to Group Policy objects (GPOs). Although you can assign local IPSec policies to computers that are not members of a trusted domain, distributing IPSec policies and managing IPSec policy configuration and trust relationships is much more time-consuming for computers that are not members of a trusted domain.

If you do use Active Directory-based IPSec policies, IPSec policy design and management must take into account the delays that result from the replication of Group Policy data from domain controllers to domain members. Often, the first step in troubleshooting a problem with IPSec connectivity is to determine whether the computer in question has the most current Group Policy assignment. To do this, you must be a member of the local Administrators group on the computer for which troubleshooting is being performed.

**Successful Mutual Authentication**
For IPSec-secured communications to be established, there must be successful mutual authentication between IPSec peers. IPSec requires the use of one of the following authentication methods: Kerberos version 5, an X.509 version 3 computer certificate issued by a public key infrastructure (PKI), or a preshared key. The two IPSec peers must use at least one common authentication method or communication will fail. Make sure that you choose an authentication method that is appropriate for your environment.

When you deploy IPSec to negotiate security for upper-layer protocols such as TCP connections, failures to communicate are often caused by the failure of IPSec to mutually authenticate the two communication endpoints. Authentication might succeed for some computers and fail for others due to issues within the authentication system itself (typically, Kerberos or public key certificates, rather than preshared keys, because preshared keys are not recommended). For these reasons, it is important to evaluate how the dependency of IPSec connectivity on authentication affects your environment and support practices. Additional training is recommended so that administrators can quickly determine whether a connectivity problem is caused by an IPSec authentication failure and understand how to investigate the authentication system.

**IPSec and ICF**
IPSec is similar to the ICF feature of Network Connections. However, there are important differences between these two technologies as well. It is important to understand the similarities and differences, so that you can deploy IPSec where it is truly needed and obtain the maximum benefits from the security that IPSec provides. This section describes similarities and differences between IPSec and ICF.

Microsoft ICF is a locally managed, stateful host firewall that, by default, discards all incoming packets except those sent in response to packets sent by the host. One primary difference between IPSec and ICF is that IPSec provides complex static filtering based on IP addresses, while ICF provides stateful filtering for all addresses on a network interface. For example, when you use IPSec, you can configure a filter to block all inbound traffic that is sent to a specific IP address over a specific protocol and port (for example, "Block all inbound traffic from the Internet to TCP port 135"). You can also configure exemptions to permit specific types of traffic from specific source IP addresses. When you use ICF, you can configure exemptions to permit traffic based solely on the port, regardless of the source IP address.

It is recommended that you use ICF when you want a firewall for a network interface that can be accessed through the Internet. It is recommended that you use IPSec when you want to secure traffic over upper-layer protocols or when you need to allow access only to a group of trusted computers. Note that it is easier to configure ICF to permit traffic over a certain port than it is to configure an IPSec policy.

IPSec is not a full-featured host firewall. However, it does provide the ability to centrally manage policies that can permit, block, or secure unicast IP traffic based on specific addresses, protocols, and ports. Some of the functions found in standard firewalls that IPSec does not provide include stateful inspection, application protocol awareness, intrusion detection, and packet logging. Although IPSec lacks some features of firewalls, the packet blocking and filtering it provides can be effective in helping to limit the spread of viruses and thwart specific attacks known to use specific ports. You can also use IPSec to prevent specific

applications and services from being used on the network.

**Related Information**

The following resources contain additional information that is relevant to this section.

- Active Directory Collection
- How IPSec Policy Extension Works
- How TCP/IP Works
- How VPN Works
- How 802.11 Wireless Works

The following RFCs and Internet Drafts are relevant to IPSec. To find the RFCs, type the appropriate RFC number in the IETF RFC Database. To find the Internet Drafts, type the appropriate keyword in the IETF Internet Drafts database.

- RFC 2401: Security Architecture for the Internet Protocol
- RFC 2402: IP Authentication Header
- RFC 2406: IP Encapsulating Security Payload (ESP)
- RFC 2408: Internet Security Association and Key Management Protocol (ISAKMP)
- RFC 2409: The Internet Key Exchange (IKE)
- UDP Encapsulation of IPSec Packets (draft-ietf-ipsec-udp-encaps-02.txt)
- Negotiation of NAT-Traversal in the IKE (draft-ietf-ipsec-nat-t-ike-02.txt)

# How IPSec Works

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How IPSec Works**

**In this section**


- IPSec Architecture
- IPSec Protocols
- IPSec Processes and Interactions
- Network Ports and Protocols Used by IPSec
- Related Information


In the Microsoft Windows Server 2003 operating system, Internet Protocol security (IPSec) helps provide defense-in-depth against network-based attacks from untrusted computers. IPSec provides protection from attack in host-to-host, virtual private network (VPN), site-to-site (also known as gateway-to-gateway or router-to-router), and secure server environments. You can configure IPSec policies to meet the security requirements of a computer, an organizational unit, a domain, a site, or a global organization.

IPSec uses packet filtering and cryptography. Cryptography provides user authentication, ensures data confidentiality and integrity, and enforces trusted communication. The strong cryptographic-based authentication and encryption support that IPSec provides is especially effective for securing traffic that must traverse untrusted network paths, such as those on a large corporate intranet or the Internet. IPSec also is especially effective for securing traffic that uses protocols and applications that do not provide sufficient security for communications.

To successfully deploy IPSec for Windows Server 2003, you must ensure the following:

- If your scenario requires Active Directory-based IPSec policy (a collection of IPSec rules that determine IPSec behavior), the Active Directory directory service and Group Policy must be configured correctly on the corporate network, appropriate trusts must be defined, and appropriate permissions must be applied. Although Group Policy applies to both users and computers, IPSec policy is a computer configuration Group Policy setting that applies only to computers.
- Each computer that will establish IPSec-secured communications must have an IPSec policy assigned. This policy must be compatible with the IPSec policy that is assigned to other computers with which that computer must communicate.
- Authentication must be configured correctly and an appropriate authentication method must be specified in the IPSec policy so that mutual authentication can occur between IPSec peers.
- Routers, firewalls, or other filtering devices must be configured correctly to permit IPSec protocol traffic on all parts of the corporate network, if IPSec negotiation messages and IPSec-secured traffic must pass through these devices.
- Computers must run operating systems that automatically support IPSec or must have appropriate client software installed.
- If computers are running different versions of the Microsoft Windows operating system (for example, Windows Server 2003, the Microsoft Windows XP operating system, and the Microsoft Windows 2000 operating system), you must address the compatibility of the IPSec policies.
- If clients must establish IPSec-secured connections with servers, those servers must be adequately sized to support those connections. If necessary, you can use IPSec hardware offload network adapters.
- The number of IPSec policies are kept to a minimum, and the IPSec policies are made as simple as possible.

- Systems administrators who will configure and support IPSec must be properly trained and must be members of the appropriate administrative groups.

**IPSec Architecture**
Several Requests for Comments (RFCs) define the architecture and components of IPSec. These components and their interrelationship comprise the logical architecture of IPSec. This section briefly describes the fundamental components of the IPSec logical architecture and then explains how these components are implemented in Windows Server 2003.

For comprehensive descriptions of IPSec architecture and components, see the RFCs that are listed in "Related Information" later in this section.

**Logical Architecture**
The IPSec architecture can be categorized into four main areas:

- Security Associations
- SA and key management support
- IPSec protocols
- Algorithms and methods

**Security Associations**
Security Associations (SAs) are a combination of a mutually agreeable policy and keys that defines the security services, mechanisms, and keys used to protect communications between IPSec peers. Each SA is a one-way or simplex connection that provides security services to the traffic that it carries.

Because SAs are defined only for one-way communication, each IPSec session requires two SAs. For example, if both IPSec protocols, Authentication Header (AH) and Encapsulating (ESP), are used for an IPSec session between two peers, then four SAs would be required.

SAs for IPSec-secured communications require two databases: a security policy database (SPD) and security association database (SAD). The SPD stores the security requirements or policy requisites for an SA to be established. It is used during both inbound and outbound packet processing. IPSec checks inbound packets to ensure that they have been secured according to policy. Outbound packets are secured according to policy.

The SAD contains the parameters of each active SA. The Internet Key Exchange (IKE) protocol automatically populates the SAD. After an SA is established, the information for each SA is stored in the SAD. The following figure shows the relationship between SAs, the SPD, and the SAD.

**SA, SPD, and SAD Architecture**




**SA and Key Management**
IPSec requires SA and key management support. The Internet Security Association and Key Management Protocol (ISAKMP) defines the framework for authentication and key exchange by providing procedures for negotiating, establishing, changing, and deleting SAs. It does not define the actual key exchange: it merely provides the framework.

IPSec requires support for both manual and automatic management of SAs and keys. IKE is the default automated key management protocol for IPSec. IKE is a hybrid protocol that incorporates parts of the Oakley key exchange protocol and the SKEME keying techniques protocol. The following figure shows the relationship between the ISAKMP, IKE, Oakley, and SKEME protocols.

**ISAKMP, IKE, Oakley, and SKEME Protocol Architecture**




The Oakley protocol uses the Diffie-Hellman key exchange or key agreement algorithm to create a unique, shared, secret key, which is then used to generate keying material for authentication or encryption. For example, such a shared secret key could be used by the DES encryption algorithm for the required keying material. A Diffie-Hellman exchange can use one of a number of groups that define the length of the base prime numbers (key size) which are created for use during the key exchange process. The longer the number, the greater the key strength. Well-known groups include Groups 1, 2, and 14.

The following figure shows the relationship between the Oakley protocol, the Diffie-Hellman algorithm, and well-known Diffie-Hellman key exchange groups.

**Oakley Protocol, Diffie-Hellman Key Exchange Algorithm, and Well-Known Diffie-Hellman Groups Architecture**




The Oakley protocol defines several modes for the key exchange process. These modes correspond to the two negotiation phases defined in the ISAKMP protocol. For phase 1, the Oakley protocol defines two principle modes: main and aggressive. IPSec for Windows does not implement aggressive mode. For phase 2, the Oakley protocol defines a single mode, quick mode.

**IPSec Protocols**

To provide security for the IP layer, IPSec defines two protocols: Authentication Header (AH) and Encapsulating Security Payload (ESP). These protocols provide security services for the SA. Each SA is identified by the Security Parameters Index (SPI), IP destination address, and security protocol (AH or ESP) identifier.

The SPI is a unique, identifying value in an SA that is used to distinguish among multiple SAs on the receiving computer. For example, IPSec communication between two computers requires two SAs on each computer. One SA services inbound traffic and the other services outbound traffic. Because the addresses of the IPSec peers for the two SAs are the same, the SPI is used to distinguish between the inbound and outbound SA. Because the encryption keys differ for each SA, each SA must be uniquely identified.

The following figure shows the relationship between the SA, SPI, IP destination address, and security protocol.

**IPSec Protocols and SA Architecture**




**Algorithms and Methods**

The IPSec protocols use authentication, encryption, and key exchange algorithms. Two authentication or keyed hash algorithms, HMAC-MD5 (Hash Message Authentication Code - MD5) and HMAC-SHA-1, are used with both the AH and ESP protocols, The DES and 3DES encryption algorithms are used with ESP. The following figure shows the relationship between the authentication and encryption algorithms and the AH and ESP protocols.

**IPSec Protocols and Algorithms for Authentication and Encryption**




The authentication methods for IPSec, as defined by the IKE protocol, are grouped into three categories: digital signature, public-key, and pre-shared key. The following figure shows the relationship between the IKE protocol and the authentication methods.

**IKE Protocol and Authentication Methods Architecture**




**Windows Server 2003 IPSec Architecture and Components**

The components and architecture of Windows Server 2003 IPSec are based on the IPSec RFCs. The basic IPSec architecture for Windows Server 2003 has the following components: Active Directory, a Policy Agent, the IKE protocol, an IPSec driver, and a TCP/IP driver.

The following figure illustrates how these components interact.

**Windows Server 2003 IPSec Architecture**




The following table describes each of these components.

**IPSec Components**

| Component | Description |
|---|---|
| Active Directory | Windows Server 2003 Active Directory stores domain-wide IPSec policies for computers that are members of the domain. Active Directory-based IPSec policies are polled and retrieved by the Policy Agent. |
| Policy Agent | The Policy Agent retrieves IPSec policy from an Active Directory domain, a configured set of local policies, or a local cache. The Policy Agent then distributes authentication and security settings to the IKE component and the IP filters to the IPSec driver. |
| IKE | IKE receives authentication and security settings from the Policy Agent and waits for requests to negotiate IPSec SAs. When requested by the IPSec driver, IKE negotiates both kinds of SAs (main mode and quick mode) with the appropriate endpoint requested by the IPSec driver based on the policy settings obtained from the Policy Agent. After negotiating an IPSec SA, IKE sends the SA settings to the IPSec driver. |

| | |
|---|---|
| IPSec driver | The IPSec driver monitors and secures outbound unicast IP traffic and monitors, decrypts, and validates inbound unicast IP traffic. After the IPSec driver receives the filters from the Policy Agent, it determines which packets are permitted, which are blocked, or which are secured. For secure traffic, the IPSec driver either uses active SA settings to secure the traffic or requests that new SAs be created. The IPSec driver is bound to the TCP/IP driver to provide IPSec processing for IP packets that pass through the TCP/IP driver. |
| TCP/IP driver | The TCP/IP driver is the Windows Server 2003 implementation of the TCP/IP protocol. It is a kernel-mode component that is loaded from the tcpip.sys file during startup. |

The architecture of the Policy Agent, IKE protocol, and IPSec driver are described in more detail in the following sections.

**Policy Agent Architecture**

The Policy Agent retrieves IPSec policy information, handles the internal interpretation and processing of the policy, and sends it to the other IPSec components that require the information to perform security services. The Policy Agent has the following components: policy store, Policy Agent service, local registry, local cache, and Interface Manager.

The following figure shows the architecture of the Policy Agent.

**Policy Agent Architecture**




The following table briefly describes the Policy Agent components.

**Policy Agent Components**

| Component | Description |
|---|---|
| Policy store | The IPSec policy store maintains both IPSec policy descriptions and interfaces to applications and other tools that provide policy data management. The policy store accesses IPSec policy data that is stored in either the local registry or in Active Directory. |
| Policy Agent | The IPSec Policy Agent controls the retrieval and distribution of IPSec policy and maintains the data about the configured policy for the IPSec driver and IKE. |
| Local registry | The local registry stores the locally configured IPSec policies, the local cache, and other IPSec settings. |
| Local cache | The local cache stores IPSec policies after they are downloaded from an Active Directory domain controller by the Policy Agent. |
| Interface Manager | Interface Manager manages a list that contains items that correspond to each physical and logical network adapter on the system. |

The following sections provide additional detail about each of these components.

**Policy store**

The policy store organizes IPSec policy data and stores it in a format that the Policy Agent can use. In Windows Server 2003, policy data can be stored in the following:

- Active Directory

- Local and remote registry
- A file (for exporting and importing only)

In addition to providing an interface that UI services can use to store policy in each of these media, the policy store does the following:

- Provides policy data for default IPSec policies
- Checks policy information for consistency
- Retrieves policy version information

The policy store reads and writes policy information both to and from persistent storage and is aware of shared policy-setting dependencies. This ensures that all policies using shared settings are marked as changed when they are modified and that Windows Server 2003 IPSec components download the modified policies.

**Policy Agent**

The Policy Agent retrieves IPSec policy information and delivers it to other IPSec components that require this information to perform security services, as shown in the following illustration.

**Policy Agent Service Retrieving and Delivering IPSec Policy Information**




The Policy Agent performs the following tasks:

- Retrieves the appropriate IPSec policy (if one has been assigned) from Active Directory if the computer is a domain member or from the local registry if the computer is not a member of a domain
- Determines filter list order
- Delivers the assigned IPSec policy information (IP filters) to the IPSec driver
- Delivers both main mode and quick mode settings to IKE
- Polls for changes in policy configuration. If the computer is a member of a domain, policy retrieval occurs when the computer starts, at the interval specified in the IPSec policy, and at the default Winlogon polling interval. You can also manually poll Active Directory for policy by using the **gpupdate /target:computer** command.

If there are no IPSec policies in Active Directory or the registry, or if the IPSec Policy Agent cannot connect to Active Directory, the IPSec Policy Agent waits for policy to be assigned or activated.

The Policy Agent appears in the list of computer services in the Services snap-in under the name IPSEC Services and starts automatically as part of the initialization of the Local Security Authority (LSA) service.

**Local registry**

Each computer running Windows XP or a Windows Server 2003 has only one local Group Policy object, often called the local computer policy. Using this local Group Policy object allows Group Policy settings to be stored on individual computers regardless of whether they are members of an Active Directory domain.

The local registry maintains the IPSec policy configuration in the following registry key and its subkeys: **HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\IPSec**.

If you assign local IPSec policies and you do not assign Active Directory-based IPSec policies, the local policies are stored in the following registry key: **HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\IPSec\Policy\Local**.

If you assign Active Directory-based IPSec policies, the policies are read from Active Directory and stored in the local cache.

**Local cache**
The local cache is part of the registry. If you assign an IPSec policy in Active Directory, the policy is stored in and read from Active Directory. A copy of the current policy in Active Directory is maintained in a cache in the local registry at: **HKEY_LOCAL_MACHINE\Software\Policies \Microsoft\Windows\IPSec\Policy\Cache**.

If a computer to which an IPSec policy in Active Directory is assigned cannot connect to the domain, the cached copy of the policy in Active Directory is applied. When the computer reconnects to the domain, new policy information for that computer replaces old, cached information. You cannot configure or manage the cached copy of an IPSec policy in Active Directory.

**Interface Manager**
Interface Manager maintains a list of physical and logical network adapters on the computer and notifies the Policy Agent when interface and address changes occur. Interface Manager also maintains a complete list of generic filters. Generic filters are filters that are configured to use My IP Address either as a source address or as a destination address. Generic filters are saved in the appropriate IPSec policy storage location with either a source address or a destination address of 0.0.0.0 and a corresponding subnet mask of 255.255.255.255.

**IKE Module Architecture**
The IKE module receives authentication and security settings from the Policy Agent and waits for requests to negotiate SAs. When the IKE module receives a request to negotiate an SA from the IPSec driver, the IKE module negotiates both kinds of SAs (the main mode SA and the quick mode SA) with the appropriate endpoint based on the request of the IPSec driver that the policy settings obtained from the Policy Agent. After it has negotiated an SA, the IKE module sends the SA settings to the IPSec driver.

The IKE module has the following components:

- CryptoAPI
- Diffie-Hellman Cryptographic Service Provider (CSP)
- RSA CSP
- Certificate store
- Security Support Provider Interface (SSPI)
- Kerberos Security Support Provider (SSP)

The following figure shows the architecture of the IKE module.

**IKE Module Architecture**




- A cryptographic service provider (CSP) is an independent software module that provides implementations of cryptographic standards and algorithms. The CSP carries out the cryptographic functions of CrytoAPI, creating keys, destroying them, and using them to perform a variety of cryptographic operations.

- The following table briefly describes the IKE module components.

**IKE Module Components**

| Component | Description |
|---|---|
| CryptoAPI | CryptoAPI provides a set of functions that allows applications based on Windows to encrypt or digitally sign data in a flexible manner while providing protection for the user's sensitive private key data. Actual cryptographic operations are performed by independent modules known as CSPs.<br><br>The IKE negotiation must be encrypted. This encryption is limited by what can be configured in IPSec policy. The standard CryptoAPI functions for keyed hashing (using HMAC-MD5 and HMAC-SHA1) and data encryption (using DES and 3DES) are used |
| Diffie-Hellman CSP | The Diffie-Hellman CSP contains the implementation of the Diffie-Hellman key exchange and determination algorithm. IKE uses only the Microsoft Base or Enhanced CSP for Diffie-Hellman. However, the Diffie-Hellman calculation can be accelerated using the CryptoAPI exponentiation offload interface (OffloadModExpo), as documented in the CryptoAPI SDK. |
| RSA CSP | The RSA CSP contains the implementation of the Rivest-Shamir-Adleman (RSA) cryptographic algorithms. When certificate authentication is selected, IKE checks the CryptoAPI default provider to see if it is capable of performing RSA 512-bit digital signatures. If so, then IKE uses this default CSP. If not, IKE enumerates the RSA providers, selects hardware-based providers first and ensures that they can provide 512-bit signatures. IKE performs these actions to open the certificate store.<br><br>The CSP that is used for signature operations during IKE negotiation is specified by the certificate selected during the IKE negotiation; the certificate's associated CSP for the private key signature is used. As long as the RSA CSP supports the NOHASHID flag for the **CryptSignHash( )** API call, IKE can use that CSP for private key signing of IKE payloads. Because the enumeration process does not permit IKE to know if the CSP supports the NOHASHID option, it is possible to choose a certificate that appears valid, but whose CSP does not allow IKE to construct the proper signature. |
| Certificate store | The certificate store is a permanent storage location where certificates, certificate revocation lists, and certificate trust lists are stored. The certificate store is a physical store on the Windows Server 2003 computer, but it is viewed logically as belonging to the user account of the currently logged-on user, a service account, or the computer account. IKE can use only the computer account, usually referred to as the computer store. You can view the computer store by using the Certificates snap-in. |
| SSPI | The SSPI enables network applications to access one of several security providers to establish authenticated connections and exchange data securely over those connections. |
| Kerberos SSP | The Kerberos SSP contains an implementation of the Kerberos security protocol. The Kerberos SSP is an SSPI provider |

**IPSec Driver Architecture**

The IPSec driver is a kernel-mode component that monitors and secures IP packets. In addition to the Policy Agent and IKE, the IPSec driver uses the following components: the Security Association Database (SAD), the Security Policy Database (SPD), the

TCP/IP driver, TCP/IP applications, and the network interface.

The IPSec driver matches IP packet information with the IP filters that are configured in the active SPD. If traffic must be secured, the IPSec driver either uses the appropriate SA to determine how to provide packet security or requests that the IKE module negotiate SAs to be used to provide packet security. After the IPSec driver determines which SA to use, it creates and validates encrypting, decrypting, and hashing to create or interpret the AH and ESP headers on an IPSec-protected packet.

The following figure shows the IPSec driver architecture and how the driver interacts with other components in Windows Server 2003.

**IPSec Driver Architecture**




The following table briefly describes the IPSec driver components.

**IPSec Driver Components**

| Component | Description |
|---|---|
| SAD | The SAD is a database in the IPSec driver that contains the parameters associated with each active SA. This database is populated automatically from the IKE module. |
| SPD | The SPD is a database in the IPSec driver that specifies the filter lists and associated settings that determine the status of all inbound or outbound IP traffic. Inbound packets are checked to ensure that they have been secured according to policy. Outbound packets are permitted, blocked, or secured, according to policy. For secured traffic, the security policy that is used is the negotiated SA, which is stored in the SAD. |
| TCP/IP driver | The TCP/IP driver is the Windows Server 2003 implementation of the TCP/IP protocol. It is a kernel-mode component that is loaded from the Tcpip.sys file during startup. |
| TCP/IP applications | TCP/IP applications use TCP/IP and access TCP/IP network services through an appropriate network API, such as Windows Sockets, NetBIOS, or Remote Procedure Call (RPC). |
| Network interface | The network interface is the logical or physical interface over which IP packets are sent and received. The details of the Network Driver Interface Specification (NDIS) interface, the network adapter driver, and the physical media over which the IP packets are sent and received are beyond the scope of this subject. |

**Policy Data Structure**

The data in a policy indicates the desired protection for the traffic between computers on a network. The data is made up of various computer-related attributes (for example, IP address and port number), the communication methods allowed (for example, algorithms and key lengths), and the IKE key negotiation and management. The policy store updates and stores the policy data. The Policy Agent retrieves the stored policy data and makes it available to all IPSec components.

As shown in the following figure, an IPSec policy contains several subsets of information.

**IPSec Policy Structure**




The following table describes the components of an IPSec policy.

**IPSec Policy Components**

| Component | Description |
|---|---|
| Policy-wide parameters | Policy-wide parameters specify the polling interval used to detect changes in policy.<br><br>The policy-wide parameters are configured on the **General** tab in the properties of an IPSec policy. |
| ISAKMP policy | The ISAKMP policy contains IKE parameters, such as encryption key lifetimes, and other settings. The ISAKMP policy settings are configured in the **Key Exchange Settings** dialog box, which is available from the **Advanced** button on the **General** tab in the properties of an IPSec policy.<br><br>The ISAKMP policy also contains a list of security methods for protecting the identity of IPSec peers during authentication. These methods are, listed in order of preference, and are configured in the **Key Exchange Security Methods** dialog box, which is available from the **Methods** button on the **Key Exchange Settings** dialog box. |
| IPSec rules | IPSec rules contain a statement that associates a filter list with a filter action, an authentication method, an IPSec mode, and other settings. Typically, an IPSec rule is configured for a specific purpose (for example, "Block all inbound traffic from the Internet to TCP port 135"). You can define many IPSec rules in a single IPSec policy.<br><br>IPSec rules are configured on the **Rules** tab in the properties of an IPSec policy. |

IPSec rules associate IKE negotiation parameters with one or more IP filters. The following table describes the components of an IPSec rule.

**IPSec Rule Components**

| Component | Description |
|---|---|
| Filter list | The filter list contains one or more predefined filters that describe the types of traffic to which an action (permit, block, or secure) is applied.<br><br>The filter list is configured on the **IP Filter List** tab in the properties of an IPSec rule within an IPSec policy. |

| | |
|---|---|
| Filter action | The filter action defines the security requirements for the data transmission. A filter action can be configured to permit traffic, block traffic, or negotiate secure communications with IPSec for packets matching the filter list. If security negotiation is selected, you must also configure security methods and their order: whether initial incoming unsecured traffic should be accepted, whether unsecured communication with computers that do not support IPSec should be allowed, and whether to use perfect forward secrecy (PFS). PFS is a mechanism that determines whether the existing keying material for a master key can be used to derive a new session key. Session key PFS performs a new Diffie-Hellman key exchange to generate new master key keying material instead of using master key keying material to derive more than one session key.<br><br>The negotiation settings are configured on the **Filter Action** tab in the properties of an IPSec rule within an IPSec policy. |
| Authentication method(s) | An IPSec rule contains one or more authentication methods, listed in order of preference, that are used for protection during IKE negotiations. The available authentication methods are the Kerberos v5 protocol, the use of a certificate issued from a specified certification authority (CA), or a preshared key.<br><br>The negotiation data is configured on the **Authentication Methods** tab in the properties of an IPSec rule within an IPSec policy. |
| Tunnel endpoint | A setting that specifies whether traffic is tunneled and, if it is, specifies the tunnel endpoint, which is the tunneling computer that is closest to the IP traffic destination, as specified by the associated IP filter list. Two rules are required to describe an IPSec tunnel. For the outbound traffic rule, the tunnel endpoint is the IP address or subnet of the IPSec peer on the other end of the tunnel. For the inbound traffic rule, the tunnel endpoint is an IP address or subnet configured on the local computer.<br><br>The tunnel endpoint is configured on the **Tunnel Setting** tab in the properties of an IPSec rule within an IPSec policy. |
| Connection type | The connection type setting specifies whether the rule applies to only local area network (LAN) connections, to only dial-up connections, or to both types of connections.<br><br>The interface applicability is configured on the **Connection Type** tab in the properties of an IPSec rule within an IPSec policy. |

**Note**

- The Kerberos v5 authentication method is not supported on computers running the Microsoft Windows XP Home Edition operating system.

**Default response rule**

Each policy has a default response rule. The default response rule has the IP filter list of **<Dynamic>** and the filter action of **Default Response** when the list of rules is viewed with the IP Security Policies snap-in. The default response rule cannot be deleted, but it can be deactivated. It is activated for all of the default policies and in the IP Security Policy wizard.

IPSec uses the default response rule to ensure that the computer responds to requests for secure communication. If an active policy does not include a rule for a computer requesting secure communication, IPSec applies the default response rule and negotiates security. For example, if Host A intends to communicate securely with Host B, but Host B does not have an inbound filter defined for Host A, then IPSec uses the default response.

For the defense response rule, you can configure only the security methods for secure traffic and the authentication methods. The following table lists the default security methods for the default response rule.

**Default Security Methods for the Default Response Rule**

| Type | AH Integrity | ESP Confidentiality | ESP Integrity |
|---|---|---|---|
| Encryption and Integrity | <None> | 3DES | SHA1 |
| Custom | <None> | 3DES | MD5 |
| Custom | <None> | DES | SHA1 |
| Custom | <None> | DES | MD5 |

| | | | |
|---|---|---|---|
| Custom | SHA1 | <None> | <None> |
| Custom | MD5 | <None> | <None> |

You can configure the security methods and their preference order on the **Security Methods** tab in the properties of the default response rule in the IP Security Policies snap-in.

The default response rule works in the following way:

1. If the IKE module receives a request to negotiate security, it queries the Policy Agent for a matching filter for traffic to and from the source and destination address of the ISAKMP message. If a matching filter is explicitly configured, the IKE negotiation is based on the settings of the associated rule.

2. If no matching filter is found and the default response rule is not activated, IKE negotiation fails.

3. If no matching filter is found and the default response rule is activated, then IKE dynamically creates an IP filter within the Policy Agent that corresponds to the traffic specification of the incoming ISAKMP message. IKE authenticates and negotiates security based on the settings on the **Authentication Methods** and **Security Methods** tabs for the default response rule.

You configure the default authentication method for the default response rule by using the IP Security Policy wizard.

**Example: Default response rule**
Typically, the default response rule is used when a group of servers are configured with policy to secure communications between themselves and any IP address and to accept unsecured communication, but respond using secured communications. The client computers are configured with the default response rule. When the clients communicate with each other, the traffic is not secured. When the clients communicate with the server, the traffic is secured (with the exception of the initial packet sent by the client to the server).

For example, a client computer can reliably exchange data with a server by using Transmission Control Protocol (TCP). A TCP connection is established through the exchange of three TCP segments: SYN (synchronize), SYN-ACK (synchronize-acknowledgement), and ACK (acknowledgement).

Because the client computer does not have an explicit rule to secure traffic to the server, the client computer sends an unsecured SYN segment to the server. The server receives the SYN segment and checks its filters. It finds the filter that requires secure traffic to and from any IP address. However, the filter action settings also allow the server to accept unsecured communication, responding with secured communication. Consequently, the SYN segment sent by the client and received by the server is sent to the TCP/IP protocol on the server.

The TCP/IP protocol on the server responds by sending a SYN-ACK segment back to the client. This IP packet is sent to the IPSec driver on the server. Checking the filter settings to require secured communications, the IPSec driver notes that there are no active SAs between itself and the client. The IPSec driver requests that the IKE module negotiate SAs for secure communication.

The IKE module on the server sends an ISAKMP message to the client to begin the process of negotiating SAs for secure communications. When the ISAKMP message is received on the client computer, the IKE module on the client computer queries the filter lists in the Policy Agent. Because it finds no explicit filter match and the default response rule is activated, the IKE module creates a dynamic filter for traffic to and from the client and server computer.

IKE negotiation continues based on the policy settings of the server and the client. After IKE negotiation is finished, the server sends the secured TCP SYN-ACK segment to the client. The client responds with a secured TCP ACK segment and secured TCP data can be exchanged between the client and the server.

**IPSec Protocols**
IPSec is integrated at the IP layer (layer 3) of the TCP/IP stack, so it provides security for almost all protocols in the TCP/IP suite. Because IPSec is applied to all applications, you do not need to configure separate security for each application that uses TCP/IP.

### IPSec Protocol Architecture




**IPSec AH and ESP Protocols**

The IPSec protocols — Authentication Header (AH) and Encapsulating Security Payload (ESP) — provide data and identity protection for each IP packet. The AH protocol is an IPSec protocol that provides data origin authentication, data integrity, and anti-replay protection for the entire packet (the IP header and the data payload carried in the packet, except fields in the IP header that are allowed to change in transit). AH can be used alone, in combination with the ESP protocol, or in IPSec tunnel mode. IPSec tunnel mode is used to protect site-to-site (gateway-to-gateway) traffic between networks, such as site-to-site networking through the Internet. The sending gateway encapsulates the entire IP datagram by adding a new IP header and then protects the new packet using one of the IPSec protocols. Windows Server 2003 supports IPSec tunnel mode for configurations where Layer Two Tunneling Protocol (L2TP) cannot be used.

The ESP protocol is an IPSec protocol that provides data confidentiality, data origin authentication, data integrity, and anti-replay protection for the ESP payload. The ESP protocol can be used alone, in combination with the AH protocol, or in IPSec tunnel mode.

**IPSec AH and ESP Protocols in IPSec Transport Mode**

IPSec protocols provide data and identity protection for each IP packet by adding their own security protocol header to each packet. This section describes how AH and ESP protect IP packets when IPSec is used in transport mode.

You use IPSec in transport mode to protect traffic in end-to-end communications scenarios (for example, for communications between a client and a server). You can also use IPSec in transport mode for basic packet filtering (that is, to statically permit or block traffic based on source and destination address combinations and on the IP protocol and TCP and UDP ports).

IPSec transport mode encapsulates the original IP payload with an IPSec header (AH or ESP).

**AH transport mode**

The AH protocol provides data origin authentication, data integrity, and anti-replay protection *for the entire packet* (both the IP header and the data payload carried in the packet), except for the fields in the IP header that are allowed to change in transit. AH does not provide data confidentiality, which means that it does not encrypt the data. The data is readable, but protected from modification and spoofing.

### AH Transport Mode Packet Structure

As shown in the figure, data integrity and authentication are provided by the placement of the AH header between the IP header and the IP packet payload. The AH protocol uses keyed hash algorithms to sign the packet for integrity. The AH protocol is identified in the IP header with an IP protocol ID of 51. This protocol can be used alone or with the ESP protocol.

The following table describes the AH header fields.

### AH Header Fields

| AH Header Field | Function |
|---|---|
| Next header | Identifies the IP payload by using the IP protocol ID. For example, a value of 6 represents TCP. |
| Length | Indicates the length of the AH header. |
| SPI | Used in combination with the destination address and the security protocol (AH or ESP) to identify the correct SA for the communication. The receiver uses this value to determine with which SA the packet is identified. |

| | |
|---|---|
| Sequence number | Provides anti-replay protection for the packet. The sequence number is a 32-bit, incrementally increasing number (starting from 1) that indicates the packet number sent over the SA for the communication. The sequence number cannot repeat for the life of the quick mode security association. The receiver checks this field to verify that a packet for a security association with this number has not already been received. If one has been received, the packet is rejected. |
| Authentication data | Contains the integrity check value (ICV), also known as the message authentication code, which is used to verify both data origin authentication and data integrity. The receiver calculates the ICV value and checks it against this value (which is calculated by the sender) to verify integrity. The ICV is calculated over the IP header, the AH header, and the IP payload. |

**ESP transport mode**

The ESP protocol provides data origin authentication, data integrity, anti-replay protection, and the option of confidentiality *for the IP payload only*. ESP in transport mode does not protect the entire packet with a cryptographic checksum nor does it protect the IP header.

**ESP Transport Mode Packet Structure**




As shown in the figure, the ESP header is placed before the IP payload and an ESP trailer and ESP trailer and authentication data field are placed after the IP payload. The ESP protocol is identified in the IP header with the IP protocol ID of 50.

The following table describes the ESP header fields.

**ESP Header Fields**

| ESP Header Field | Function |
|---|---|
| SPI | When used in combination with the destination address and the security protocol (AH or ESP), the SPI identifies the SA for the communication. The receiver uses this value to determine the SA with which this packet should be identified. |
| Sequence number | Provides anti-replay protection for the packet. The sequence number is a 32-bit, incrementally increasing number (starting from 1) that indicates the packet number sent over the quick mode SA for the communication. The sequence number cannot repeat for the life of the quick mode SA. The receiver checks this field to verify that a packet for an SA with this number has not already been received. If one has been received, the packet is rejected. |

The following table describes the ESP trailer fields.

**ESP Trailer Fields**

| ESP Trailer Field | Function |
|---|---|
| Padding | Padding of 0 to 255 bytes is used to ensure that the encrypted payload is on byte boundaries required by encryption algorithms. |
| Padding length | Indicates the length of the Padding field in bytes. The receiver uses this field to remove padding bytes after the encrypted payload with the padding bytes has been decrypted. |
| Next header | Identifies the type of data in the payload, such as TCP or UDP. |

The following table describes the ESP authentication trailer field.

**ESP Authentication Trailer Field**

| ESP Authentication Trailer Field | Function |
|---|---|
| Authentication data | Contains the ICV, also known as the message authentication code, which is used to verify both message authentication and integrity. The receiver calculates the ICV value and checks it against this value (which is calculated by the sender) to verify integrity. The ICV is calculated over the ESP header, the payload data, and the ESP trailer. |

**IPSec AH and ESP Protocols in IPSec Tunnel Mode**

You use IPSec tunnel mode primarily to protect traffic between sites that must traverse an untrusted path. For example, you can use tunnel mode to do the following:

- Establish gateway-to-gateway tunnels between sites, when the gateways or end systems do not support L2TP/IPSec Virtual Private Network (VPN) connections.
- Protect traffic end-to-end when one endpoint of the communication does not support IPSec. You can send protected traffic to a computer that supports tunnel mode and that is placed immediately in front of the computer that does not support IPSec.

With tunnel mode, an entire IP packet is encapsulated with an AH or ESP header and an additional IP header. The IP addresses of the outer IP header are the tunnel endpoints, and the IP addresses of the encapsulated IP header are the ultimate source and destination addresses.

You can use both the ESP and AH protocols in combination when tunneling to provide both confidentiality for the tunneled IP packet and integrity and authentication for the entire packet.

**AH tunnel mode**

AH tunnel mode encapsulates an IP packet with an AH and IP header and signs the entire packet for integrity and authentication.

### AH Tunnel Mode Packet Structure




**ESP tunnel mode**

ESP tunnel mode encapsulates an IP packet with both an ESP and IP header and an ESP authentication trailer.

### ESP Tunnel Mode Packet Structure




The signed portion of the packet indicates where the packet has been signed for integrity and authentication. The encrypted portion of the packet indicates what information is protected with confidentiality.

Because a new header for tunneling is added to the packet, everything that comes after the ESP header is signed (except for the ESP authentication trailer) because it is now encapsulated in the tunneled packet. The original header is placed after the ESP header. The entire packet is appended with an ESP trailer before encryption occurs. All data that follows the ESP header, except for the ESP authentication trailer, is encrypted. This includes the original header, which is now considered to be part of the data portion of the packet.

The entire ESP payload is then encapsulated within the new tunnel header. The tunnel header is not encrypted because it is used only to route the packet from origin to tunnel endpoint.

If the packet is being sent across a public network, it is routed to the IP address of the gateway for the receiving intranet. The gateway decrypts the packet, discards the ESP header, and uses the original IP header to route the packet to the intranet computer.

**IPSec Processes and Interactions**

This section describes the following IPSec processes and interactions:

- Policy Agent initialization
- Policy data retrieval and distribution
- IKE main mode and quick mode negotiation
- Key protection
- IPSec Driver Processes

**Policy Agent Initialization**
When the Policy Agent is started by the Local Security Authority (LSA) service, it performs a number of initialization steps and then begins a cycle during which it waits for a number of events to be signaled.

At startup, the Policy Agent checks the registry for the value that indicates the location of the computer's IPSec policy in Active Directory. To find the value, it checks the following registry key **HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\IPSec\GPTIPSECPolicy\DSIPSECPolicyPath**.

If the computer is a member of a domain and an IPSec policy has been configured in Active Directory, this value was set during computer initialization. If the computer is not a member of a domain or there is no IPSec policy in Active Directory, the Policy Agent finds no value.

After initializing interval values for event processing, the Policy Agent starts the Interface Manager (IM). The IM establishes the event processing that needs to be signaled whenever a network interface configuration or status changes. The Policy Agent then starts the IPSec driver.

Next, IKE starts. At this point, additional initialization occurs. This includes determining through the IPSec driver whether strong cryptography is available.

In preparation for loading the IPSec policy from either Active Directory or the local registry, the Policy Agent initializes the policy state and polling interval. It then loads the IPSec policy from the appropriate store, as described in the following section. After the policy is loaded, the Policy Agent begins a service loop.

If, during the service wait cycle, the Policy Agent encounters either an unexpected event or the Policy Agent service stop event, the Policy Agent shuts down IKE, stops the IPSec driver, cleans up local data, and then stops itself.

**Policy Data Retrieval and Distribution**
The Policy Agent loads the static IPSec policy from one of the stores if one of the following situations occurs:

- The Policy Agent has initialized and started IKE and the IPSec driver.
- An event has signaled the Policy Agent to reload the policy.
- A polling interval timeout occurs when the Policy Agent unsuccessfully accesses the current policy location to check for updates.

The Policy Agent first attempts to retrieve the policy from Active Directory. If it successfully does so, the Policy Agent supplies the retrieved policy data to IKE and the IPSec driver. The Policy Agent then copies the IPSec policy to the registry cache.

If the Policy Agent fails to retrieve the policy from Active Directory, it checks the local cache for the IPSec policy data. If the local cache contains policy data, the Policy Agent supplies the data to IKE and the IPSec driver.

If the Policy Agent fails to retrieve policy from Active Directory and the local cache, it checks the registry for local policy data. If the policy is successfully retrieved and the data is successfully provided to IKE and the IPSec driver, the Policy Agent is placed in the Local Downloaded policy state.

If the Policy Agent successfully loads a policy, the polling interval for policy checks is set from the value in the policy data. If the Policy Agent fails to load a policy, it goes into the Initial policy state and sets the polling interval to the default value.

While loading the static IPSec policy, the Policy Agent notes the state of the static policy data. If a service loop timeout occurs, the Policy Agent uses the state information to determine activity.

**IKE Main Mode and Quick Mode Negotiation**
The IKE component is a Policy Agent service. IKE is started by the Policy Agent, restarted by the Policy Agent as needed, and shut down when the Policy Agent shuts down. All policy data that IKE requires for operation is provided by the Policy Agent.

IKE establishes a combination of mutually agreeable policy and keys that defines the SA; the security services, protection mechanisms, and cryptographic keys between communicating peers.

To ensure successful and secure communication, IKE performs a two-phase negotiation operation. Phase 1 negotiation is known as main mode negotiation and Phase 2 is known as quick mode negotiation. The IKE main mode SA (also known as the ISAKMP SA) protects the IKE negotiation itself. The SAs created during the second IKE negotiation phase are known as the quick mode SAs (also known as IPSec SAs). The quick mode SAs protect application traffic. Two quick mode SAs are negotiated, one for inbound and one for outbound traffic.

**Main Mode Negotiation**
IKE performs main mode negotiation with an IPSec peer to establish protection mechanisms and keys for subsequent use in protecting quick mode IKE communications. IKE main mode negotiation occurs in three parts:

- Part one: Negotiation of protection mechanisms
- Part two: Diffie-Hellman exchange
- Part three: Authentication

Main mode negotiation consists of the exchange of a series of six ISAKMP messages. An ISAKMP message is the payload of a (User Datagram Protocol) UDP message with the source and destination UDP ports set to 500 (or 4500). An ISAKMP message has an ISAKMP header and one or more ISAKMP payloads as defined in RFC 2408.

Both the initiator and the responder in the exchange send three messages. The initiator is the IPSec peer that initiates secure communications by sending the first message. The responder, which sends the second message, is the IPSec peer with which the initiator is requesting secure communications.

The following table shows the first four main mode messages, which are not encrypted.

**Main Mode Messages 1 Through 4**

| Main Mode Message | Sender | Payload |
|---|---|---|
| 1 | Initiator | ISAKMP header, Security Association (contains proposals) |
| 2 | Responder | ISAKMP header, Security Association (contains a selected proposal) |
| 3 | Initiator | ISAKMP header, Key Exchange (contains Diffie-Hellman key), Nonce, additional payloads (depending on authentication method) |
| 4 | Responder | ISAKMP header, Key Exchange (contains Diffie-Hellman key), Nonce, additional payloads (depending on authentication method) |

The first four main mode messages contain the following ISAKMP payloads:

- **Security Association**. The Security Association payload sent in message 1 is a list of proposed protection mechanism for the main mode SA. The Security Association payload sent in message 2 is a specific protection suite for the main mode SA that is common to both IPSec peers. It is selected by the responder.
- **Key Exchange**. The Key Exchange payload is sent in message 3 by the initiator and in message 4 by the responder and contains Diffie-Hellman key determination information for the Diffie-Hellman key exchange process.
- **Nonce**. The Nonce payload is sent in messages 3 and 4 and contains a nonce, which is a pseudorandom number that is used only once. The initiator and responder each send their own unique nonces. Nonces are used to provide replay protection.

Depending on the authentication method that is selected in the IPSec policy, messages 3 and 4 might contain additional

payloads. The payloads of all messages beyond the first four messages are encrypted and vary based on the authentication method selected.

**Note**

- It is important to understand the differences in negotiation behavior between initiating an IKE main mode negotiation and quick mode negotiation and responding to one, and rekeying an existing one. The IKE RFC 2409 requires that rekeys can be performed by either peer (in either direction) at any time, regardless of the security association lifetimes negotiated. Therefore, the computer that initiates a negotiation might become the responder and these roles might alternate many times. Some of these differences are due to behavior required for interoperability and some are caused by enforcement of policy settings.

**Part One: Negotiation of protection mechanisms**

When initiating an IKE exchange, IKE proposes protection mechanisms based on the applied security policy. Each proposed protection mechanism includes attributes for encryption algorithms, hash algorithms, authentication methods, and Diffie-Hellman groups. The first part of the main mode is contained in main mode messages 1 and 2.

The following table lists the protection mechanism attribute values that are supported by Windows Server 2003 IKE. These values are described in more detail in later sections.

**Main Mode Attribute Values Supported by IKE**

| Main Mode Attribute | Value |
|---|---|
| Encryption algorithm | DES, 3DES |
| Integrity algorithm | HMAC-MD5, HMAC-SHA1 |
| Authentication method | Kerberos v5, public key certificate, preshared key |
| Diffie-Hellman group | Group 1, Group 2, Group 14 (2048) |

The encryption algorithm, integrity algorithm, and Diffie-Hellman group are configured as one of multiple key exchange security methods.

The initiating IKE peer proposes one or more protection suites in the same order as they appear in the applied security policy. If one of the protection suites is acceptable to the responding IKE peer, the responder selects it for use and responds to the initiator with its choice. Because the responding IKE peer might not be running Windows Server 2003 or Windows 2000 and is selecting the first proposed protection suite that is acceptable, the protection suites in the applied security policy should be configured in the order of most secure to least secure.

**Part Two: Diffie-Hellman exchange**

After a protection suite is negotiated, IKE queries a Diffie-Hellman CSP through CryptoAPI to generate a Diffie-Hellman public and private key pair based on the negotiated Diffie-Hellman group. The Diffie-Hellman public key is sent to the IKE peer in an ISAKMP Key Exchange payload. Main mode negotiation part 2 is contained in main mode messages 3 and 4.

The cryptographic strength of a Diffie-Hellman key pair is related to its prime number length (key size). Windows Server 2003 IKE supports the following Diffie-Hellman groups:

- Group 1 (768 bits)
- Group 2 (1024 bits)
- Group 14 (2048 bits)

  **Note**

  - For enhanced security, Windows Server 2003 IPSec includes Diffie-Hellman Group 14, which provides 2048 bits of keying strength. However, Diffie-Hellman Group 14 is not currently supported in Windows 2000 or Windows XP for general IPSec policy use. For updated information about the availability of Diffie-Hellman Group 14 for L2TP/IPSec connections for Windows XP and Windows 2000, search for article 818043, L2TP/IPSec NAT-T Update for Windows XP and Windows 2000, in the Microsoft Knowledge Base.

After the Diffie-Hellman public keys are exchanged, IKE accesses CryptoAPI to compute the shared key based on the mutually agreeable authentication method.

The following figure shows the Diffie-Hellman exchange between the IPSec peers and the relationship between IKE, CryptoAPI, and the Diffie-Hellman CSP.

**IKE Diffie-Hellman Key Exchange**




The Diffie-Hellman exchange occurs in the following steps:

1. On each IPSec peer, IKE requests that the first Diffie-Hellman CSP generate a Diffie-Hellman public and private key pair based on the Diffie-Hellman group selected during main mode messages 1 and 2.

2. The public portion of the Diffie-Hellman public and private key pair is returned to IKE.

3. The initiator sends the Diffie-Hellman public key to the responder in a Key Exchange payload (main mode message 3).

4. The responder sends its Diffie-Hellman public key to the initiator in a Key Exchange payload (main mode message 4).

5. On each IPSec peer, the IKE component sends the other IPSec peer's Diffie-Hellman public key to the Diffie-Hellman CSP.

6. The Diffie-Hellman CSP computes the shared secret and returns the value to the IKE component.

**Part Three: Authentication**

IKE supports three methods of authentication:

- Kerberos v5

- Public key certificate

- Preshared key

The authentication that occurs is a computer-based authentication, also known as machine-based authentication. The authentication process verifies only the identity of the computers, not the individual using the computer when the authentication process occurs.

**Kerberos v5 authentication**

Kerberos v5 authentication is the default authentication standard in Windows Server 2003 and Windows 2000 domains. Any computer in the domain or a trusted domain can use this method of authentication.

**Note**

- In Windows Server 2003, the Kerberos protocol is no longer a default exemption. Therefore, if you want to enable Kerberos authentication, you must create filters in the IPSec policy that explicitly allow such traffic. For more information, see article 810207, "IPSec Default Exemptions Are Removed in Windows Server 2003" in the Microsoft Knowledge Base.

Windows IKE Kerberos authentication is based on the Generic Security Service (GSS) API IKE authentication method, which is described in the Internet draft entitled "A GSS-API Authentication Method for IKE."

The following table lists the ISAKMP messages exchanged during a Kerberos authentication main mode negotiation.

**Kerberos Authentication Method Main Mode Messages**

| Main Mode Message | Sender | Payload |
|---|---|---|
| 1 | Initiator | ISAKMP header, Security Association (contains proposals) |
| 2 | Responder | ISAKMP header, Security Association (contains a selected proposal) |
| 3 | Initiator | ISAKMP header, Key Exchange (contains Diffie-Hellman key), Nonce, Initiator Kerberos Token |
| 4 | Responder | ISAKMP header, Key Exchange, Nonce, Responder Kerberos Token |
| 5* | Initiator | ISAKMP header, Identification, Initiator Hash |
| 6* | Responder | ISAKMP header, Identification, Responder Hash |

* ISAKMP payloads of message are encrypted.

Authentication occurs when:

- Each peer authenticates the other peer's Kerberos token: The responder verifies the initiator's Kerberos token and the initiator verifies the responder's Kerberos token.
- Each peer's hash is calculated and verified: The responder verifies the initiator's hash and the initiator verifies the responder's hash.

The hash calculation is performed over the following:

- The Diffie-Hellman public values of the initiator and responder
- The initiator and responder cookies
- The ISAKMP payloads of message 2
- The identity name string for the IPSec peer
- The IPSec peer's Kerberos token

To perform Kerberos authentication, Windows Server 2003 uses a Kerberos SSP that is accessible through the SSPI. The following figure shows the exchange of Kerberos tokens between the IKE peers and the relationship between the IKE, SSPI, and the Kerberos SSP.

**IKE Authentication with Kerberos**




Kerberos authentication is performed in the following steps:

1. IKE on the initiator accesses the SSPI to initialize a security context.
2. The initiator's Kerberos SSP creates a Kerberos token and returns it to the IKE component.

3. The initiator sends the initiator's Kerberos token and computer identity to the responder in main mode message 3.

4. The responder IKE component accesses the SSPI to acquire a security context.

5. The responder's Kerberos SSP creates a Kerberos token and returns it to IKE.

6. The responder sends the responder Kerberos token and computer identity to the sender in main mode message 4.

7. On each IPSec peer, IKE creates the hash for the next main mode message to be sent and then requests that SSPI sign the hash with the Kerberos session key.

8. The Kerberos SSP returns the signature to the IKE component.

9. The initiator sends the signed initiator hash to the responder in main mode message 5.

10. The responder sends the signed responder hash to the initiator in main mode message 6.

11. On each IPSec peer, IKE accesses the SSPI to compute the hash for the other peer. The initiator accesses the SSPI to compute the responder hash. The responder accesses the SSPI to compute the initiator hash.

12. The Kerberos SSP returns the computed hash to the IKE component where the hash value is verified to complete IKE authentication. The initiator compares its calculated responder hash with the responder hash received in main mode message 6. The responder compares its calculated initiator hash with the initiator hash received in main mode message 5.

The following sections describe the IKE certificate selection and acceptance process. If you decide to use certificates for IKE authentication, understanding this process and its requirements is integral to ensuring proper deployment.

**Public key certificate authentication**

Windows IKE performs public key certificate authentication during main mode in compliance with RFC 2409. IKE uses CryptoAPI to retrieve the computer certificate, verify peer certificates and certificate chains, check certificate revocation, and create and verify digital signatures. All certificate, certificate chain, and signature information is exchanged in main mode messages, as shown in the following table.

**Certificate-based IKE Authentication Main Mode Messages**

| Main Mode Message | Sender | Payload |
|---|---|---|
| 1 | Initiator | ISAKMP header, Security Association (contains proposals) |
| 2 | Responder | ISAKMP header, Security Association (contains a selected proposal) |
| 3 | Initiator | ISAKMP header, Key Exchange (contains Diffie-Hellman key), Nonce |
| 4 | Responder | ISAKMP header, Key Exchange (contains Diffie-Hellman key), Nonce, Certificate Request |
| 5* | Initiator | ISAKMP header, Identification, Certificate, Certificate Request, Signature |
| 6* | Responder | ISAKMP header, Identification, Certificate, Signature |

* ISAKMP payloads of message are encrypted.

**IKE certificate selection process**

When IKE negotiates to use certificates for authentication, the following process is used to select a computer certificate:

1. The list of trusted roots is prepared. This is the list of the CA root names provided by the peer in the Certificate Request Payloads (CRPs), and it matches the CA root names configured in the list of trusted roots in the appropriate authentication method of the IPSec policy. If there are no matching CA root names, all trusted CA root names from the appropriate authentication method are used.

2. IKE searches the computer store for an IPSec certificate that chains to any of the trusted CA roots identified in step 1. An IPSec certificate contains an Enhanced Key Usage (EKU) attribute with a value equal to the IP security IKE intermediate object identifier (OID) 1.3.6.1.5.5.8.2.2.

3. For each certificate chain found, checks are performed to verify the following:

   - The certificate chain does not have any trust errors.
   - The certificate chain is not a root-only chain.
   - The computer certificate has a private key.
   - The computer certificate has an RSA type public/private key pair.
   - The computer certificate has a public key length that is greater than 512 bits.
   - The computer certificate has a Digital Signature key usage.
   - The computer certificate is not a CA signing certificate that is used to issue certificates.
   - The certificate chain passes certificate revocation list (CRL) checking, which is performed by default or if the value of the **StrongCRLCheck** registry subkey is set to **1** or **2**. For more information about CRL checking, see "IPSec CRL checking" later in this section.

   If all of these checks succeed, IKE selects the certificate chain to be sent to the IPSec peer. If any of these checks fails, IKE continues to search for another IPSec type certificate, using the same list of root CA names.

4. If a valid computer certificate chain is not located, IKE retries the process, from step 2. Although it uses the same list of root CA names, IKE does not search for an IPSec type certificate.

5. If a valid computer certificate chain is still not found and if the list of root CA names in step 1 is a subset of the names allowed by the local IPSec policy, IKE retries, from step 2. This time, IKE uses the entire list of root CA names allowed by the local authentication method.

   This step is required for successful authentication when cross-certificates are used to establish trust relationships.

6. After IKE selects a computer certificate, it includes all intermediate certificates in the chain up to the root, except for the root CA certificate. A certificate chain in PKCS#7 format is then sent to the IPSec peer. If there are no intermediate CAs, only the computer certificate is sent.

   If a computer certificate cannot be selected, the authentication fails.

   **Note**

   - If IKE negotiates with another computer running Windows Server 2003, or with other Microsoft IKE implementations that use IPSec NAT-T (such as Microsoft L2TP/IPSec VPN Client), a special method to avoid fragmentation of ISAKMP UDP packets might be implemented. Otherwise, the ISAKMP message that contains the certificate chain will likely be fragmented as the packet is transmitted.

**IKE certificate acceptance process**

1. IKE receives the peer's certificates or certificate chains and verifies that the peer's certificates chain up to any of the root CAs in the appropriate authentication method of the local IPSec policy.

2. For each peer certificate chain, checks are performed to verify that:

- The computer certificate Subject Name or Subject AltName is consistent with the peer's ID field passed in the IKE negotiation.
- The computer certificate chain does not have any trust errors. If there is a trust error, the peer authentication fails.

3. If the two checks in step 2 succeed, checks are performed for the peer certificate chain to verify that:

   - The certificate chain passes CRL checking, which is performed by default or if the value of the **StrongCRLCheck** registry subkey is set to **1** or **2**.
   - The computer certificate has an RSA type public/private key pair.
   - The computer certificate has a public key length that is greater than 512 bits.
   - The computer certificate has a Digital Signature key usage.

   If any of these checks fails, the peer authentication fails.

4. If certificate-to-account mapping is enabled in the IPSec policy for the certificate root CA of the peer, IKE calls the Windows secure channel (Schannel) APIs to perform the mapping. Schannel completes the mapping and builds an access token for the computer account. This access token is automatically evaluated against the **Access this computer from the network** or the **Deny this computer access from the network** logon right defined in Group Policy Security settings.

   If the logon right evaluation fails, the peer authentication fails.

### IPSec CRL checking

If you use certificate-based authentication, you can also enable IPSec certificate revocation list (CRL) checking. By default, in Windows XP and Windows Server 2003, IPSec CRLs are automatically checked during IKE certificate authentication, but a fully successful CRL check is not required for the certificate to be accepted. However, if enhanced security is required, a fully successful CRL check is also required. CRL checking can cause delays in authentication or unnecessary failures, and some non-Microsoft PKI systems might not support it. You can disable IPSec CRL checking or specify a stronger level of IPSec CRL checking by using the Netsh IPSec context or by modifying the registry.

To disable IPSec CRL checking or specify a different level of IPSec CRL checking, use the following command:

```
netsh ipsec dynamic set config strongcrlcheck value={0 | 1 | 2}
```

#### To enable IPSec CRL checking through the registry

**Note**

- Incorrectly editing the registry might severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

1. Under the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent**\ key, add a new Oakley subkey, with a DWORD entry named **StrongCRLCheck**.

2. Assign this entry any value from 0 through 2, where:

   - A value of **0** disables CRL checking (this is the default for Windows 2000).
   - A value of **1** causes CRL checking to be attempted and certificate validation to fail only if the certificate is revoked

(this is the default for Windows XP and Windows Server 2003). Other failures that are encountered during CRL checking (such as the revocation URL not being reachable) do not cause certificate validation to fail.

- A value of **2** enables strong CRL checking, which means that CRL checking is required and that certificate validation fails if any error is encountered during CRL processing. Set this registry value for enhanced security.

3. Do one of the following:

   - Restart the computer.
   - Stop and then restart the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

Note that IPSec CRL checking does not guarantee that certificate validation fails immediately when a certificate is revoked. There is a delay between the time that the revoked certificate is placed on an updated and published CRL and the time when the computer that performs the IPSec CRL checking retrieves this CRL. The computer does not retrieve a new CRL until the current CRL has expired or until the next time the CRL is published. By default, IKE requests that CryptoAPI wait 15 seconds to complete the CRL retrieval. If the CRL cannot be retrieved at that time, IKE either ignores the error (if the value of the **StrongCRLCheck** registry subkey is set to **1**, or it causes authentication to fail (if the value of **StrongCRLCheck** is set to **2**). CRLs are cached in memory and in **\Documents and Settings*UserName*\Local Settings\Temporary Internet Files** by CryptoAPI. Because CRLs persist across computer restarts, if a CRL cache problem occurs, restarting the computer does not resolve the problem.

**Excluding the CA name from certificate requests**

If you use certificate authentication to establish trust between IPSec peers, you can also use Windows Server 2003 to exclude CA names from certificate requests. Excluding the CA name prevents a malicious user from learning sensitive information about the trust relationships of a computer, such as the name of the company that owns the computer and the domain membership of the computer (if an internal PKI is being used). Although excluding the CA name from certificate requests enhances security, computers with multiple certificates from different roots might require the CA root names to select the correct certificate. Also, some non-Microsoft IKE implementations might not respond to a certificate request that does not include a CA name. For these reasons, excluding the CA name from certificate requests might cause IKE certificate authentication to fail in certain cases.

**Certificate-to-account mapping**

In Windows Server 2003, a specific group of computers can be authorized to use IPSec when either Kerberos v5 or certificates are used for IKE authentication. This capability enables much stronger peer authentication and allows IPSec to be used to restrict network access to a server. When you enable IPSec certificate-to-account mapping, the IKE protocol associates (that is, maps) a computer certificate to a computer account in an Active Directory domain or forest, and then retrieves an access token, which includes the list of computer security groups. This process ensures that the certificate offered by the IPSec peer corresponds to an active computer account in the domain, and that the certificate is one that should be used by that computer.

Certificate-to-account mapping can be used only for computer accounts that are in the same forest as the computer performing the mapping. This provides much stronger authentication than simply accepting any valid certificate chain. For example, you can use this capability to restrict access to computers that are within the same forest. Certificate-to-account mapping, however, does not ensure that a specific trusted computer is allowed IPSec access.

If the certificate-to-account mapping process is not completed properly, authentication will fail and IPSec-protected connections will be blocked.

**Preshared key authentication**

Preshared key authentication requires that each IKE peer use a predefined and shared key to authenticate the IKE exchange.

The following table describes the main mode messages for preshared key authentication.

**Preshared Key Authentication Main Mode Messages**

| Main Mode Message | Sender | Payload |
|---|---|---|
| 1 | Initiator | ISAKMP header, Security Association (contains proposals) |

| | | |
|---|---|---|
| 2 | Responder | ISAKMP header, Security Association (contains a selected proposal) |
| 3 | Initiator | ISAKMP header, Key Exchange (contains Diffie-Hellman key), Nonce |
| 4 | Responder | ISAKMP header, Key Exchange (contains Diffie-Hellman key), Nonce |
| 5* | Initiator | ISAKMP header, Identification, Initiator Hash |
| 6* | Responder | ISAKMP header, Identification, Responder Hash |

* ISAKMP payloads of message are encrypted.

Messages 5 and 6 contain an initiator and responder hash calculated with the preshared key. Each IPSec peer authenticates the other peer's packet by decrypting and verifying the hash inside the packet (the hash inside the packet is a hash of the preshared key).

**Note**

- Preshared keys are easily implemented but can be compromised if they are not used correctly. Microsoft does not recommend the use of preshared key authentication because the key value is not securely stored, and it is, therefore, difficult to keep secret. Preshared key authentication is provided for interoperability purposes and compliance with RFC standards.

#### Quick Mode Negotiation

When main mode negotiation completes or an existing quick mode SA expires, IKE begins quick mode negotiation. During quick mode negotiation, IKE queries the Policy Agent for information required to perform the appropriate filter actions, including whether the IPSec mode is tunnel or transport, whether the protocol is ESP or AH or both, and which encryption and hashing algorithms are proposed or accepted.

#### Quick Mode SA negotiation

The quick mode negotiation process is implemented as defined in RFC 2409. All quick mode negotiation messages are protected with the main mode SA that was established during the main mode negotiation. Each successful quick mode negotiation establishes two quick mode SAs. One SA is inbound and the other SA is outbound.

The following table lists the quick mode messages exchanged by two IPSec peers running Windows IPSec.

**Quick Mode Messages**

| Quick Mode Message | Sender | Payload |
|---|---|---|
| 1* | Initiator | ISAKMP header, Security Association (contains proposals and secure traffic description) |
| 2* | Responder | ISAKMP header, Security Association (contains a selected proposal) |
| 3* | Initiator | ISAKMP header, Hash |
| 4* | Responder | ISAKMP header, Notification |

* ISAKMP payloads of message are encrypted.

The four quick mode messages contain the following payloads:

- **Security Association.** This payload contains a list of proposals and encryption and hashing algorithms (AH or ESP, DES or 3DES, and HMAC-MD5 or HMAC-SHA1) for securing the traffic and a description of the traffic that is protected. This description might include IP addresses, IP protocols, TCP ports, or UDP ports, and is based on the matching filter of the initiator. The Security Association in the second quick-mode message includes a Security Association payload that contains the chosen method of securing the traffic.

- **Hash.** This payload provides verification and replay protection.

- **Notification.** This payload has a connected notify message. This message is requested and sent between two IPSec peers running Windows Server 2003. Quick mode message 4 with the Notification payload is not required by the IKE standard and is used to prevent the initiator from sending IPSec-protected packets to the responder before the responder is ready to receive them.

- For more information about the contents of ISAKMP payloads for quick mode exchanges, see RFC 2407 and RFC 2408 in the IETF RFC Database.

Windows Server 2003 IPSec supports the filter action choices listed in the following table.

**Filter Action Choices**

| Filter Action Choices | ESP Encryption/Integrity Algorithm | AH |
|---|---|---|
| Encryption and integrity | 3DES/HMAC-SHA1 | None |
| Integrity only | None/HAMC-SHA1 | None |
| Custom | DES, 3DES, or none/HMAC-MD5, HMAC-SHA1, or none | HMAC-MD5 or HMAC-SHA1 |

**Note**

- Computers running Windows Server 2003 and Windows XP support the 3DES and DES algorithms and do not require installation of additional components. However, computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use 3DES. If a computer running Windows 2000 is assigned a policy that uses 3DES encryption, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the security method defaults to the weaker DES algorithm.

**Generating and regenerating session key material**

IKE generates session keys for both the inbound and outbound quick mode SAs based on the main mode shared master key and nonce material exchanged during the quick mode negotiation. Additionally, Diffie-Hellman key exchange material can also be exchanged and used to enhance the cryptographic strength of the IPSec session key.

**Key Protection**

The following features enhance the base prime numbers (keying material) and the strength of the keys for master and session keys.

**Key Lifetimes**

Key lifetimes control when a new key is generated. A key lifetime allows you to force automatic key regeneration after a specific interval of either kilobytes (KB) or seconds, whichever occurs first. This ensures that even if an attacker is able to decipher part of a communication protected by one key, new keys protect the remainder of the communication. Whenever a key lifetime is reached, the SA is also renegotiated and the key is refreshed or regenerated. For this reason, key lifetimes are also referred to as SA lifetimes. You can specify key lifetimes for the master key and for session keys.

The master key lifetime corresponds to the main mode SA lifetime created by the IKE main mode negotiation. It is configured in terms of time and number of quick mode negotiations and it applies to all security rules in the IPSec policy. The main mode SA and master key typically have a long lifetime (the default value is eight hours). Master keys are much more resource-intensive to generate than session keys because they require reauthentication and additional Diffie-Hellman exchanges.

The session keys correspond to the quick mode SAs that are used to protect program traffic. The session keys and quick mode SAs are quickly derived from the master key by IKE quick mode negotiation. Session keys are used to protect data and they have lifetimes based on the amount of data sent and the amount of time elapsed since the key started being used. Typically, session keys have shorter lifetimes than master keys. The default values are 100,000 KB (approximately 100 megabytes) or one hour.

When session keys are refreshed, new quick mode SAs replace the old ones. Quick mode SA session keys must be refreshed before either the data or time lifetime expires; otherwise, traffic is discarded. A session key will be deleted if the quick mode SAs become idle (by default, SAs become idle after five minutes). Because the master key lives longer than the session key, it allows new quick mode SAs and session keys to be established quickly. To prevent data loss, the quick mode SA session key is generated shortly before the main mode SA expires. The main mode SA and corresponding master key are not deleted when session keys and quick mode SAs are deleted, unless the specified number of quick mode SA negotiations has elapsed.

For example, if a communication takes one hour (3,600 seconds) and if you specify the minimum session key lifetime of five minutes (300 seconds), more than 12 keys are generated to complete the communication. If a 100 MB file is transferred over a fast corporate LAN using an IPSec security association with a 100 MB and one hour lifetime, at least one, if not two, session rekeys occur.

**Session Key Refresh Limit**
Repeated rekeying from a session key can compromise the Diffie-Hellman shared secret. Consequently, you might want to limit the number of quick mode session keys that can be derived from a main mode negotiation.

If you have enabled master key PFS, the session key limit is disregarded. Setting a session key limit to 1 is identical to enabling master key PFS. If both a master key lifetime and a session limit are specified, whichever limit is reached first causes a new main mode negotiation. By default, IPSec policy does not specify a session limit.

**Diffie-Hellman Groups**
Diffie-Hellman groups are used to determine the length of the base prime numbers (key material) for the Diffie-Hellman exchange. The cryptographic strength of any key derived from a Diffie-Hellman exchange depends, in part, on the strength of the Diffie-Hellman group on which the prime numbers are based. When a stronger group is used, the key that is derived from a Diffie-Hellman exchange is stronger and more difficult for an attacker to break.

IKE negotiates which group to use, ensuring that there are not any negotiation failures that result from a mismatched Diffie-Hellman group between the two peers.

If session key PFS is enabled, a new Diffie-Hellman key is negotiated during the first quick mode SA negotiation. This new key removes the dependency of the session key on the Diffie-Hellman exchange that is performed for the master key.

Both the initiator and responder must have session key PFS enabled, or negotiation fails.

The Diffie-Hellman group is the same for both the main mode and quick mode SA negotiations. When session key PFS is enabled, even though the Diffie-Hellman group is set as part of the main mode SA negotiation, it affects any rekeys during session key establishment.

**Perfect Forward Secrecy**
Unlike key lifetimes, PFS determines how a new key is generated, rather than when it is generated. Specifically, PFS ensures that the compromise of a single key permits access only to data that is protected by it, not necessarily to the entire communication. To achieve this, PFS ensures that a key used to protect a transmission cannot be used to generate additional keys. In addition, if the key that was used was derived from specific keying material, that material cannot be used to generate other keys.

**Master key PFS**
In Windows Server 2003 IPSec, you can configure the number of times quick mode SAs can be created based on a single main mode SA. If you enable master key PFS, the IKE allows only a single quick mode SA for each main mode SA. By default, master key PFS is disabled, so there is no limit to the number of quick mode SAs that can be created from one main mode SA. To derive a new quick mode SA, a new main mode negotiation is performed, which includes a new Diffie-Hellman exchange and a new authentication process.

**Session key PFS**
Whenever a quick mode SA requires renegotiation, IKE determines whether a session key PFS is specified in the corresponding filter rule. If it is, IKE additionally generates a new Diffie-Hellman key and exchanges it with the IKE peer during quick mode negotiation. By performing another Diffie-Hellman key exchange, IKE provides additional cryptographic strength to quick mode key generation beyond that already contributed by the main mode SA. Performing additional Diffie-Hellman exchanges requires additional computational resources and might affect IPSec performance.

**IPSec Driver Processes**
The IPSec driver does not participate in IP packet processing until the first time the Policy Agent informs the driver that there is an active IPSec policy. If no IPSec policy is active, the IPSec driver does not participate in inbound and outbound IP traffic processing.

**IPSec Driver Responsibilities**
The IPSec driver is responsible for the following:

- Maintaining the SAD and SPD.
- Checking each IP packet to determine whether it matches a policy filter. When a match is found and an SA must be created, the IPSec driver invokes the IKE module. After the IKE module completes its negotiations, an SA is returned to the IPSec driver and stored in the SAD.
- Implementing the IPSec policy as specified in the SA (for example, using a specific hashing method for outbound packets, verifying the integrity of inbound packets, and using a specific method for encrypting and decrypting).
- Tracking the length of time a specific key is in use and requesting a new key from IKE as necessary.

- Tracking the number of bytes that have been transformed (that is, hashed or encrypted) for each SA and requesting a new key from the IKE module if the byte count allowed by the SA is exceeded.

- For each secured inbound packet that contains an AH or ESP header, parsing the packet for the SPI to determine the SA.

- For each non-secured inbound packet, checking the filter list in the SPD to determine whether the packet is permitted or discarded:

    - The filter list can contain an inbound permit filter if the corresponding filter action is set to **Permit** or **Negotiate security** and either the **Accept unsecured communication, but always respond with IPSec** or **Allow unsecured communication with non IPSec-aware computer** options are enabled. The IPSec driver sends unmodified permitted packets to the TCP/IP driver for additional processing.

    - The packet is discarded either because the filter action is set to Block or the filter action is set to Negotiate security and unsecured communications are not allowed.

- Handling hardware offload of cryptographic functions by skipping cryptographic processing on packets processed by offload network adapters and managing offloaded SAs.

- Handling network layer issues such as path maximum transmission unit (PMTU) discovery.

- Creating the SPI that the responder uses to identify the appropriate SA for an inbound packet.

- Deleting expired SAs.

- Providing the implementation of the AH and ESP protocols.

  For outbound traffic that must be secured; the IPSec driver, based on the parameters of the SA, calculates and places the AH or ESP or both headers and trailer on the IP packet before sending it to the TCP/IP driver. For inbound traffic that contains an AH or ESP header, the IPSec driver processes the header and, if it is valid, sends the authenticated and decrypted packet without the AH or ESP headers and trailer back to the TCP/IP driver.

**IPSec Driver Communication**

The following types of communications occur between the Policy Agent, IPSec driver, and IKE module:

- The Policy Agent adds a set of filters to the IPSec driver. Each filter is accompanied by a policy identifier (a GUID) and an index. The index indicates to IPSec what weight to assign to the filter. A lower index indicates higher precedence and higher weight.

- When the Policy Agent deletes a set of filters, it deletes all associated outbound SAs first and then allows the inbound SAs to expire.

- The Policy Agent can receive a set of usage statistics from the IPSec driver. These statistics include the number of packets sent and received for both AH and ESP protocols, the number of SAs, the number of rekeys, and the number of bad packets.

- The IKE module adds SAs as the result of successful negotiation of keys.

- The IKE module can expire a specified SA.

**Packet Processing**

The IPSec driver receives the active IP filter list from the Policy Agent, as shown in the following illustration, and then attempts to match every inbound and outbound packet against the filters in the list.

### IPSec Driver Matching an IP Filter List




When a packet matches a filter, the IPSec driver applies the filter action. When a packet does not match any filters, the IPSec driver passes the packet back without modification to the TCP/IP driver to be received or transmitted.

If the filter action permits transmission, the packet is received or sent with no modifications. If the action blocks transmission, the packet is discarded. If the action requires the negotiation of security, main mode and quick mode SAs are negotiated.

The negotiated quick mode SA and keys are used with both outbound and inbound processing. The IPSec driver stores all current quick mode SAs in a database. The IPSec driver uses the SPI field to match the correct SA with the correct packet.

When an outbound IP packet matches the IP filter list with an action to negotiate security, the IPSec driver queues the packet and then notifies IKE, which begins security negotiations with the destination IP address of that packet. If several outbound packets are going to the same destination and match the same filter before IKE has finished the negotiation, then only the last packet sent is saved.

The following sections describe the basic inbound packet and outbound packet processing that the IPSec driver performs in transport mode.

**Inbound packet processing**

The following figure illustrates this process.

**Basic Inbound Packet Process**




**Note**

- The inbound packet process applies only to local host unicast traffic (traffic with the unicast destination address of the host) when there is an active IPSec policy.

Basic inbound packet processing for transport mode occurs in the following sequence:

1. IP packets are sent from the network interface to the TCP/IP driver.

2. The TCP/IP driver sends the IP packet to the IPSec driver.

3. If the inbound packet is IPSec-protected, the IPSec driver looks up the SA in the SAD.

4. If the inbound packet is not IPSec-protected, the IPSec driver checks the packet for a filter match by looking up the filters in the SPD.

5. After the IPSec-protected inbound packet is authenticated and decrypted, the AH or ESP or both headers are removed and the packet is sent to the TCP/IP driver. If a packet that is not IPSec-protected is permitted by policy, that packet is sent

to the TCP/IP driver.

6. The TCP/IP driver performs IP packet processing as needed and sends the application data to the TCP/IP application.

Detailed inbound packet processing for transport mode occurs in the following sequence:

1. The TCP/IP driver sends the unicast packet to the IPSec driver.

2. If the packet is ISAKMP, the unmodified packet is sent back to the TCP/IP driver.

   **Note**

   - To modify the default filtering behavior for Windows Server 2003 IPSec, you can use the Netsh IPSec context or modify the registry. For more information, see "Default exemptions to IPSec filtering" later in this section.

3. If hardware offload processing was performed, the IPSec driver checks to determine whether the hardware processing was successful.

4. If the hardware processing was not successful, an event is logged and the packet is discarded.

5. The packet is parsed to determine whether an AH or ESP header or both are present.

6. If the packet does not contain an AH or ESP header, the packet is compared to the filter list for a match.

7. If a filter match is not found, the unmodified packet is sent to the TCP/IP driver.

8. If a filter match is found, the IPSec driver attempts to find an SA based on the packet contents.

9. If an SA is not found, the matching filter is checked to determine if it is an inbound permit filter.

10. If the matching filter is an inbound permit filter, the unmodified packet is sent to the TCP/IP driver.

11. If the matching filter is not an inbound permit filter, the packet is discarded.

12. If an SA is found, it is checked to determine whether it is a soft SA. A soft SA is one in which the **Negotiate security** filter action is enabled, but there is no authentication or encryption being performed because the computer with which communication occurs is not running IPSec. This process is also known as fallback to clear. Even though the packet is not being protected, an SA without an AH or ESP header is still maintained in the SAD. Soft SAs and fallback to clear are possible only when **Allow unsecured communication with non IPSec-aware computer** is selected on the **Security methods** tab in the properties of a filter action.

13. If the SA is a soft SA, the unmodified packet is sent to the TCP/IP driver.

14. If the SA is not a soft SA, the packet is discarded.

15. If the packet contains an AH or ESP header (or both), the header is parsed for the SPI.

16. The SPI is used to look up the SA in the SAD.

17. If the SA corresponding to the SPI is not found in the SAD, a Bad SPI event is logged and the packet is discarded.

18. If the SA corresponding to the SPI is found in the SAD, the current time is used to update the SA's last used time. The time is used for aging the SA.

19. The SA is checked to determine whether cryptographic processing for the SA was offloaded to hardware. For packets that

have been processed by hardware offload, steps 20 and 21 are skipped.

20. The packet is authenticated or decrypted or both. This process involves verifying the HMAC in the AH or ESP header, processing the other fields in the AH and ESP headers and trailer, and decrypting the ESP payload.

21. If cryptographic processing is unsuccessful, an event is logged and the packet is discarded.

22. The AH or ESP headers and ESP trailer are removed.

23. If the SA for this packet is a tunnel SA (using either AH or ESP tunnel mode), the decapsulated packet is reinjected into the TCP/IP driver and the original packet is discarded. By reinjecting the decapsulated packet, the TCP/IP driver treats it as if it were received from the network adapter.

24. If the SA for this packet is not a tunnel SA, the IP packet, with the AH and ESP headers removed, is sent back to TCP/IP driver for additional processing.

**Outbound packet processing**

Basic outbound packet processing is shown in the following figure.

**Basic Outbound Packet Process**




Basic outbound packet processing for transport mode occurs in the following sequence:

1. Application data is sent to the TCP/IP driver from the TCP/IP application.

2. The TCP/IP driver sends an IP packet to the IPSec driver.

3. The IPSec driver checks the packet for a filter match by looking up the filters in the SPD.

4. The IPSec driver checks the packet for an active SA by looking up the SAs in the SAD. Based on the SA, the traffic is authenticated or encrypted or both.

5. If the traffic must be protected and there is not an SA, the IPSec driver requests that IKE create the appropriate SAs. The IP packet is then held until the SA is established and can be IPSec framed.

6. The IP packet is sent back to the TCP/IP driver.

7. The TCP/IP driver sends the IP packet to the network interface.

Detailed outbound packet processing for transport mode occurs in the following sequence:

1. The TCP/IP driver sends the unicast outbound packet to the IPSec driver.

2. If the packet is ISAKMP, the unmodified packet is sent back to the TCP/IP driver.

3. The IPSec driver attempts to find a filter that matches the packet. If a filter is not found, the unmodified packet is sent

back to the TCP/IP driver.

4. If a filter match is found, the IPSec driver attempts to find an SA that matches the packet.

5. If an SA is not found, the filter action is checked. If the filter action is set to **Negotiate security**, the IPSec driver requests that the IKE module negotiate the appropriate SAs.

6. If the IKE negotiation is successful, the IKE module informs the IPSec driver of the new SA and the IPSec driver looks up the SA again.

7. If the IKE negotiation is not successful, the packet is discarded.

8. If the filter action is set to **Permit**, the unmodified packet is sent back to the TCP/IP driver. Otherwise, the packet is discarded.

9. If an SA is found in the SAD, the current time is used to update the SA's last used time. The time is used for aging the SA.

10. The SA is checked to determine whether it is about to expire. If the SA is about to expire, the IPSec driver informs the IKE module to initiate a quick mode or Phase 2 rekey of the quick mode SA.

11. The SA is checked to determine whether it has expired. If the SA has expired, the packet is discarded.

12. The Don't Fragment (DF) flag in the IP header of the packet is checked. If the DF flag is set to 1, the size of the IP packet with the proposed AH or ESP or both headers and trailer is calculated.

13. If the size of the IP packet with the proposed IPSec overhead is larger than the path maximum transmission unit (PMTU) for the destination IP address, the IPSec driver indicates a packet-too-large condition for the packet and the unmodified packet is sent back to the TCP/IP driver. The packet-too-large condition allows the TCP/IP driver to either adjust the PMTU for the destination or, in the case of transit traffic, inform the sending host with an Internet Control Message Protocol (ICMP) Destination Unreachable-Fragmentation Needed and DF Set message that includes the new PMTU. The packet is eventually discarded by the TCP/IP driver.

14. If the DF flag is not set to 1, or if it is set to 1 and the additional IPSec overhead is not greater than the current PMTU for the destination, blank AH or ESP both headers and trailer are constructed (based on the settings for the SA).

15. The IPSec driver checks to determine whether the hardware offload is capable of offloading the SA for this packet. If so, the IPSec driver checks to determine whether the SA for the packet was offloaded to the hardware.

16. If the SA was offloaded to the hardware, an offload status is set on the packet and the modified packet with blank AH or ESP or both headers and trailer is sent to the TCP/IP driver.

17. If the SA has not been offloaded to the hardware, the IPSec driver accesses NDIS with instructions to add the SA to the hardware offload network interface.

18. If hardware offload is not enabled or the SA has not been offloaded to the hardware, the IPSec driver performs the cryptographic processing and adds the appropriate values in the fields of the AH or ESP or both headers and trailer.

19. The IPSec driver sends the modified packet to the TCP/IP driver.

**Default exemptions to IPSec filtering**

In Windows Server 2003, the default filtering exemptions have been removed for Kerberos, Resource Reservation Setup Protocol (RSVP), and multicast and broadcast traffic, but remain for ISAKMP traffic, and inbound multicast and broadcast traffic.

To modify the default filtering behavior for Windows Server 2003 IPSec, you can use the Netsh IPSec context or modify the registry.

To modify the default filtering behavior using Netsh, use the following command:

```
netsh ipsec dynamic set config ipsecexempt value={ 0 | 1 | 2 | 3}
```

Depending on which exemptions you want, specify the appropriate values as follows:

- A value of **0** specifies that multicast, broadcast, RSVP, Kerberos, and ISAKMP traffic are exempt from IPSec filtering. This is the default filtering behavior for Windows 2000 (with Service Pack 3 and earlier service packs) and Windows XP.

  **Note**

  - Use this setting only if it is required for compatibility with Windows 2000 and Windows XP. If Kerberos traffic is exempted from filtering, an attacker can bypass other IPSec filters by using either UDP or TCP source port 88 to access any open port. Many port scan tools will not detect this because these tools do not allow setting the source port to 88 when checking for open ports.

- A value of **1** specifies that Kerberos and RSVP traffic are not exempt from IPSec filtering (multicast, broadcast, and ISAKMP traffic are exempt).

- A value of **2** specifies that multicast and broadcast traffic are not exempt from IPSec filtering (RSVP, Kerberos, and ISAKMP traffic are exempt).

- A value of **3** specifies that only ISAKMP traffic is exempt from IPSec filtering. This is the default filtering behavior for Windows Server 2003, Windows 2000 (with Service Pack 4 and later service packs) and Windows XP (with Service Pack 1 and later service packs).

If you change the value for this setting, you must restart the computer for the new value to take effect.

**To modify the default filtering behavior by using the registry**

1. In Regedit, under the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IPSEC** key, add a new DWORD entry named **NoDefaultExempt**.

2. Assign this entry any value from **0** through **3**.

3. Restart the computer.

The filtering behaviors for each value are equivalent to those noted above for the **netsh ipsec dynamic set config ipsecexempt value=***x*command.

The following table summarizes the equivalent filters that are implemented if all default exemptions to IPSec filtering are enabled (that is, if **NoDefaultExempt** is **0**). When the IP address is specified, the subnet mask is 255.255.255.255. When the IP address is **Any**, the subnet mask is 0.0.0.0.

**Equivalent Filters When NoDefaultExempt=0**

| Source Address | Destination Address | Protocol | Source Port | Destination Port | Filter Action |
|---|---|---|---|---|---|
| My IP Address | Any IP Address | UDP | Any | 88 | Permit |
| Any IP Address | My IP Address | UDP | 88 | Any | Permit |
| Any IP Address | My IP Address | UDP | Any | 88 | Permit |
| My IP Address | Any IP Address | UDP | 88 | Any | Permit |
| My IP Address | Any IP Address | TCP | Any | 88 | Permit |
| Any IP Address | My IP Address | TCP | 88 | Any | Permit |

| | | | | | |
|---|---|---|---|---|---|
| Any IP Address | My IP Address | TCP | Any | 88 | Permit |
| My IP Address | Any IP Address | TCP | 88 | Any | Permit |
| My IP Address | Any IP Address | UDP | 500 | 500 [1] | Permit |
| Any IP Address | My IP Address | UDP | 500 | 500 | Permit |
| My IP Address | Peer IP Address | UDP | 4500 | 4500 [2] | Permit |
| Peer IP Address | My IP Address | UDP | 4500 | 4500 | Permit |
| My IP Address | Any | 46 (RSVP) | | | Permit |
| Any IP Address | My IP Address | 46 (RSVP) | | | Permit |
| Any IP Address | <multicast> [3] | | | | Permit |
| My IP Address | <multicast> | | | | Permit |
| Any IP Address | <broadcast>[4] | | | | Permit |
| My IP Address | <broadcast> | | | | Permit |
| <All IPv6 protocol traffic> [5] | | | | | Permit |

[1]In order for IPSec transport mode to be negotiated through an IPSec tunnel mode SA, ISAKMP traffic cannot be exempted if it needs to pass through the IPSec tunnel first.

[2] When IPSec NAT-T is performed, the filter exemption for UDP port 4500 is automatically generated based on the source and destination IP addresses used during the initial part of the IKE negotiation on UDP port 500. This dynamic permit filter for port 4500 is displayed in the IP Security Monitor snap-in, under **Quick Mode\Specific Filters**, and in the output for the **netsh ipsec dynamic show qmfilter** command.

[3]Multicast traffic is defined as the class D range, with a destination address range of 224.0.0.0 with a 240.0.0.0 subnet mask, which corresponds to the range of addresses from 224.0.0.0 to 239.255.255.255.

[4]Broadcast traffic is defined as a destination address of 255.255.255.255 (the limited broadcast address) or as having the host ID portion of the IP address set to all 1's (the subnet broadcast address).

[5]IPSec does not support filtering for IP version 6 (IPv6) packets, except when IPv6 packets are encapsulated with an IPv4 header.

Windows Server 2003 IPSec does not support specific filters for broadcast protocols or ports, nor does it support multicast groups, protocols, or ports. Because IPSec does not negotiate security for multicast and broadcast traffic, these types of traffic are dropped if they match a filter with a corresponding filter action to negotiate security. A filter with a source address of **Any IP Address** and a destination address of **Any IP Address** can block or permit all multicast and broadcast traffic. By default (and if **the NoDefaultExempt** registry key is set to a value of **2** or **3**), outbound multicast or broadcast traffic will be matched against a filter with a source address of **My IP Address** and a destination address of **Any IP Address**. More specific unicast IP address filters that block, permit, or negotiate security for unicast IP traffic should be configured in the same IPSec policy to achieve appropriate security.

**Hardware acceleration (offloading)**

Hardware acceleration is accomplished by offloading specific processing tasks that are normally completed by an operating system component to the network adapter. Some network adapters can perform IPSec cryptographic functions, such as encryption and decryption of data and the calculation and verification of message authentication codes.

When the NDIS interface binds, the offload capability of the network interface is queried. During outbound packet processing, after an SA is created a check is made to ensure that the network interface can offload cryptographic functions, support transport-over-tunnel functionality, and support IP header options. If not, the packet cannot be offloaded. A check is also made to determine whether the SA for the packet being offloaded is a soft SA. A soft SA is an SA in which no authentication or encryption is being performed because the computer with which communication occurs is not running IPSec. Because no AH

or ESP headers need to be processed, hardware offloading is unnecessary.

If hardware offloading is enabled, a check is made per packet to determine whether the SA for an outbound packet has already been offloaded to the offload adapter. If so, the existing offloaded SA is used. If the SA is not yet offloaded and the offload has not previously failed for this SA, an attempt is made to offload the SA. However, the attempt to offload the SA is made asynchronously. The IPSec driver does not wait for the SA offload to be successful before continuing to process the packet. This causes the first packet to always be cryptographically processed by the IPSec driver, with the cryptographic processing of following packets occurring on the hardware offload network adapter.

Typically, IPSec network offload adapters do not accelerate the IKE negotiation. However, some SSL offload adapters might be capable of processing the IKE Diffie-Hellman calculation in hardware. To determine whether your SSL offload adapter can do so, see the manufacturer's documentation.

Windows 2000, Windows XP, and Windows Server 2003 provide hardware acceleration APIs in the Windows Driver Development Kit (DDK) as part of TCP/IP Task Offload. For more information about these APIs, see *Task Offload* on MSDN.

**Network Ports and Protocols Used by IPSec**
The following table lists the network ports and protocols used by IPSec.

### IPSec Port and Protocol Assignments

| Protocol | Protocol ID | UDP | TCP |
|---|---|---|---|
| ESP | 50 | N/A | N/A |
| AH | 51 | N/A | N/A |
| ISAKMP | N/A | 500 (4500) | N/ A |

The following sections describe how to configure routers, firewalls, or other filtering devices to ensure that traffic that is sent over IPSec protocols can pass through these devices. Additional considerations for IPSec NAT traversal are also described.

**Firewall Filters**
In order for IPSec-secured communications to take place through a firewall or other filtering device, you must configure the firewall to permit IPSec traffic on UDP source and destination port 500 (ISAKMP) and IP Protocol 50 (ESP). You might also need to configure the firewall to permit IPSec traffic on IP protocol 51 (AH) to permit troubleshooting by IPSec administrators and to allow the traffic to be inspected while it is still IPSec-encapsulated.

Ensure that the firewall filter can permit or track fragments for ISAKMP. IKE with certificate or Kerberos authentication requires ISAKMP packets to be fragmented because the ISAKMP protocol uses UDP. ISAKMP messages that are larger than the local interface MTU are automatically fragmented by IP. If only certificate authentication is used, Windows Server 2003 implements a method to avoid IKE message fragmentation. When certificate authentication is used for communication between computers running Windows Server 2003 IPSec and Windows XP IPSec or Windows 2000 IPSec, fragmentation is required.

You must also allow ISAKMP to be initiated from either a source or destination IP address. RFC 2408 specifies that the ISAKMP protocol must be able to negotiate security in either direction. Stateful filtering that allows only one computer to initiate IKE to a responder typically times out and deletes the stateful inbound filter in the firewall. As a result, IKE cannot rekey IPSec security associations, and IPSec connectivity is lost.

**IPSec NAT-T**
In Windows 2000 and Windows XP, if traffic between the client and a server must pass through a network address translator (NAT), then IPSec cannot secure the traffic (the IKE negotiation will fail when translated by a NAT). Windows Server 2003 provides support for version 2 of a new IETF Internet draft called IPSec NAT-T. IPSec NAT-T allows IPSec ESP packets in either transport mode or tunnel mode to pass through NATs that allow UDP traffic. In this design, IKE automatically detects NATs and uses UDP-ESP encapsulation on UDP port 4500 to enable traffic to pass through a network address translator. The Windows Server 2003 implementation of IPSec NAT-T also supports PMTU discovery for UDP-ESP encapsulation. This new functionality allows you to secure servers running Windows Server 2003, when clients are behind a network address translator. IPSec NAT-T does not support the use of AH across network address translators.

If you are using IPSec NAT-T to secure a server, it is recommended that you do not create UDP port 4500 filters in the IPSec policy that is assigned to the server. The IPSec driver recognizes UDP port 4500 traffic and detects the associated UDP-ESP quick mode SA. However, if you are using firewalls or filtering routers to filter traffic for the IPSec-secured server, then you must configure the firewalls or filtering routers to permit the UDP-ESP traffic.

To configure firewalls or filtering routers to permit traffic on UDP source and destination port 4500, use the following settings to create a filter called "Permit IPSec NAT-T ISAKMP traffic on UDP port 4500":

- Source address = *SpecificIPAddress*
- Destination address = *SpecificIPAddress*
- Protocol = UDP
- Source port = **Any** or **4500** (The network address translator might translate source port 4500 to a different source port)
- Destination port = 4500

On computers that are running Windows 2000 or Windows XP, you can install an update package that allows L2TP/IPSec clients that are behind network address translators to use IPSec NAT-T. The NAT-T functionality provided in this update package meets the specifications of IIETF RFC 3193, "Securing L2TP using IPSec," and version 2 of the "UDP Encapsulation of IPSec Packets" and "Negotiation of NAT-Traversal in the IKE" Internet drafts and is compatible with Windows Server 2003 IPSec NAT-T. However, using IPSec in general transport mode for NAT-T is not supported on computers running Windows XP or Windows 2000, even when this update is installed, because PMTU discovery is not provided by UDP-ESP traffic. For more information about the IPSec update package for Windows 2000 and Windows XP, see article 818043, "L2TP/IPSec NAT-T Update for Windows XP and Windows 2000," in the Microsoft Knowledge Base.

**Related Information**
The following RFCs and Internet Drafts are relevant to IPSec. To find the RFCs, type the appropriate RFC number in the IETF RFC Database. To find the Internet Drafts, type the appropriate keyword in the IETF Internet Drafts Database.

- RFC 2401: Security Architecture for the Internet Protocol
- RFC 2402: IP Authentication Header
- RFC 2406: IP Encapsulating Security Payload (ESP)
- RFC 2408: Internet Security Association and Key Management Protocol (ISAKMP)
- RFC 2409: The Internet Key Exchange (IKE)
- UDP Encapsulation of IPSec Packets (draft-ietf-ipsec-udp-encaps-02.txt)

Negotiation of NAT-Traversal in the IKE (draft-ietf-ipsec-nat-t-ike-02.txt)

# IPSec Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec Tools and Settings**

**In this section**


- IPSec Tools
- IPSec Registry Entries
- Related Information


Use the following Internet Protocol security (IPSec) tools and registry setting to enable, configure, and manage IPSec on a computer running the Microsoft Windows Server 2003 operating system.

**IPSec Tools**

To manage IPSec policy in Windows Server 2003, you use three tools: the IP Security Policy Management snap-in, the Netsh IPSec context, and the Resultant Set of Policy (RSoP) snap-in. To monitor IPSec performance, you use the IP Security Monitor snap-in. If you need additional troubleshooting functionality, you can use audit logging and detailed IKE logging in Event Viewer and Network Monitor.

As the following table shows, several of the tools that you used to create and manage IPSec policies in previous versions of the Microsoft Windows operating system have been changed, replaced, or are no longer available in Windows Server 2003.

**IPSec Tool Changes in Windows Server 2003**

| Tool | Where It Was Previously Available | Changes in Windows Server 2003 |
|---|---|---|
| IPSeccmd | Support Tools folder of the Microsoft Windows XP operating system CD | Not included in Windows Server 2003. Use the Netsh IPSec context instead. |
| IPSecmon.exe | Microsoft Windows 2000 | Replaced by the new IP Security Monitor snap-in. |
| IPSecpol.exe | *Microsoft Windows 2000 Server Resource Kit* | Not included in Windows Server 2003. |
| Netdiag.exe | Support Tools folder of the Windows Server 2003 operating system CD | No longer includes IPSec functionality. Use the Netsh IPSec context instead. |

Despite these changes, you can share IPSec policies created on computers running the Windows XP, Windows 2000, and Windows Server 2003 operating systems among any of those operating systems.

**IP Security Policy Management Snap-In**

**Category**

The IP Security Policy Management snap-in is included in Windows Server 2003.

**Version compatibility**

The IP Security Policy Management snap-in is available with the Microsoft Windows XP, Windows 2000 Server, and Windows Server 2003 operating systems.

To create and manage IPSec policy, you primarily use the IP Security Policy Management snap-in that is available in the Microsoft Management Console (MMC). You can use IP Security Policy Management to create, modify, and store local IPSec policies or IPSec policies based on the Active Directory directory service. Additionally, you can use IP Security Policy Management to modify IPSec policy on remote computers.

To manage local IPSec policy, you can use either IP Security Policy Management or the command line, depending on your deployment needs. It is recommended that you use IP Security Policy Management to manage IPSec policy for Active Directory-based IPSec policy,

You use different methods to access IP Security Policy Management, depending on whether the IPSec policy is Active Directory-based or local.

To access Active Directory-based IPSec policy, do either of the following on the computer from which you want to manage policy:

- Start IP Security Policy Management from the appropriate organizational unit (OU) in Active Directory (Group Policy).
- Add IP Security Policy Management for Active Directory-based IPSec policy to MMC.

To access local IPSec policy for a computer, do any of the following on the computer for which you want to manage policy:

- Start IP Security Policy Management from Local Security Policy.
- Add IP Security Policy Management for local IPSec policy to MMC.
- Add Group Policy Object Editor for local IPSec policy to MMC.

To create an IPSec policy, a user or process must be logged on to the computer as a member of the Domain Admins group or the local Administrators group or must be running with local system privileges.

For more information about Group Policy settings, see "Group Policy Settings Reference for Windows Server 2003" in the Tools and Settings Collection.

**Resultant Set of Policy (RSoP) Snap-In**
**Category**
RSoP is included as part of the operating system.

**Version compatibility**
Logging mode is available on Windows XP and later operating systems. Planning mode requires that you have a Windows Server 2003 computer as a domain controller.

To view IPSec policy assignments for a computer or for members of a Group Policy container, use the RSoP addition to Group Policy. You can use information about IPSec policy assignments to troubleshoot policy precedence issues and to plan your deployment.

To view IPSec policy assignments in RSoP, you must first open the RSoP MMC console, and then run a query. RSoP provides two types of queries: logging mode queries (for viewing IPSec policy assignments for a computer) and planning mode queries (for viewing IPSec policy assignments for members of a Group Policy container).

For more information about RSoP, see “How Core Group Policy Works” in Core Group Policy Technical Reference.

**Logging mode queries**
To view all of the IPSec policies that are assigned to an IPSec client, you run an RSoP logging mode query. When you run a logging mode query, RSoP retrieves policy information from the Windows Management Instrumentation (WMI) repository on the target computer, and then displays this information in the RSoP console. In this way, RSoP provides a view of the policy settings that are being applied to a computer at a given time.

The query results show the precedence of each IPSec policy assignment, so that you can quickly determine which IPSec policies are assigned but are not being applied and which IPSec policy is being applied. The RSoP console also displays detailed settings (the filter rules, filter actions, authentication methods, tunnel endpoints, and connection type) for the IPSec policy that is being applied.

**Planning mode queries**
To view all of the IPSec policies that are assigned to members of a Group Policy container, you run an RSoP planning mode query. A planning mode query can be useful if, for example, you are planning a company reorganization and you want to move computers from one OU to a new OU. By supplying the appropriate information and then running a planning mode query, you can determine which IPSec policies are assigned but are not being applied to the new OU and which IPSec policy is being applied. In this way, you can identify which policy would be applied if you were to move the computers to the new OU. As with logging mode queries, when you run a planning mode query, the RSoP console displays detailed policy settings for the IPSec policy that is being applied.

When you run a planning mode query, RSoP retrieves the names of the target user, computer, and domain controller from the WMI repository on the domain controller. WMI then uses the Group Policy Data Access Service (GPDAS) to create the policy settings that would be applied to the target computer, based on the RSoP query settings that you entered. RSoP reads the policy settings from the WMI repository on the domain controller, and then displays this information in the RSoP console user

interface.

**IP Security Monitor Snap-In**
**Category**
IP Security Monitor is part of the operating system.

**Version compatibility**
In Windows 2000, IP Security Monitor was implemented as an executable program (IPSecmon.exe). In Windows XP and Windows Server 2003, IP Security Monitor is implemented as an MMC console. In Windows Server 2003, IP Security Monitor includes enhancements that allow you to:

- Monitor IPSec information for your local computer and for remote computers.
- View details —, including the policy name, description, date last modified, store, path, OU, and Group Policy object name — about active IPSec policies.
- View main mode and quick mode generic filters and specific filters.
- View main mode and quick mode statistics.
- Customize refresh rates and use Domain Name System (DNS) name resolution for filter and security association output.
- Search for specific main mode or quick mode filters that match any source or destination IP address, a source or destination IP address on your local computer, or a specific source or destination IP address.

**Netsh.exe: Netsh**
**Category**
Netsh is a command-line tool.

**Version compatibility**
Although you can use netsh commands with both Windows 2000 Server and Windows Server 2003, specific IPSec-related capabilities were added in Windows Server 2003.

The Netsh commands for IPSec provide a fully equivalent alternative to the console-based management and diagnostic capabilities provided by the IP Security Policy Management and IP Security Monitor consoles. You can use Netsh commands for IPSec to script IPSec policy creation, display details about IPSec policies, and change the IPSec configuration for troubleshooting. In addition, administering IPSec from the command line is useful when you want to extend the security and manageability of IPSec. For example, you can use Netsh commands for IPSec to enable IPSec driver event logging, set default traffic exemptions, and configure computer startup security.

For more information about using Netsh IPSec, see the "Command-Line Reference" in the Tools and Settings Collection.

**Eventvwr.msc: Audit Logging in Event Viewer**
**Category**
Event Viewer is part of the operating system.

**Version compatibility**
Event Viewer is available with Windows XP, Windows 2000 Server, and Windows Server 2003.

You can view the success or failure of Internet Key Exchange (IKE) negotiations in the Event Viewer security log. To view these events, enable success or failure auditing for the **Audit logon events** audit policy for your domain or local computer.

To disable audit logging, see "DisableIKEAudits" in IPSec Registry Entries.

**Eventvwr.msc: Detailed IKE Logging in Event Viewer**
**Category**
Event Viewer is part of the operating system.

**Version compatibility**
Event Viewer is available with Windows XP, Windows 2000 Server, and Windows Server 2003.

Enabling audit logging for IKE events and viewing the events in Event Viewer is the fastest and simplest way to troubleshoot failed main mode or quick mode negotiations. However, some scenarios might require a more detailed analysis of the IKE main mode negotiation and quick mode negotiations for troubleshooting. If the audit failure events do not provide enough information, you can enable tracing for IKE negotiations. The IKE tracing log is a very detailed log intended for troubleshooting

IKE interoperability under controlled circumstances. Expert knowledge of the Internet Security Association and Key Management Protocol (ISAKMP) RFC 2408 and IKE RFC 2409 is required to interpret this log.

The IKE tracing log appears as the *systemroot*\Debug\Oakley.log file. The log has a fixed size of 50,000 lines and will overwrite as necessary. Each time the IPSec service is started, a new Oakley.log file is created and the previous version of the Oakley.log file is saved as Oakley.log.sav. When the Oakley.log file becomes full, it is saved as Oakley.log.bak, and a new Oakley.log file is created.

Many IKE negotiations can occur simultaneously. Therefore, to capture a more easily interpreted log, minimize the number of negotiations and log for as short a period of time as possible.

In Windows Server 2003, you can enable or disable the IKE tracing log dynamically while the IPSec service is running. You do this by using Netsh IPSec. For more information about how to enable or disable the IKE tracing log, see "Command-Line Reference" in the Tools and Settings Collection.

**Netmon.exe: Network Monitor**
**Category**
Network Monitor is available in Microsoft Systems Management Server or with Windows 2000 Server and Windows Server 2003.

**Version compatibility**
You can use Network Monitor to capture and view packets in Windows XP, Windows 2000, or Windows Server 2003.

To view IPSec and other network communication, you can install and use Network Monitor.

### Note

- You can use the version of Network Monitor that is provided with Windows Server 2003 to view only the network traffic that is sent to or from the computer on which Network Monitor is installed. To view network traffic that is sent to or from another computer and is routed through your computer (using the Routing and Remote Access service), you must use the version of Network Monitor that is provided with Systems Management Server.

The version of Network Monitor that is provided with Windows Server 2003 includes parsers for the ISAKMP (IKE), Authentication Header (AH), and Encapsulating Security Payload (ESP) protocols. The Network Monitor parsers for ESP can parse inside the ESP packet only if null-encryption is being used and the full ESP packet is captured. Network Monitor cannot parse the encrypted portions of IPSec-secured ESP traffic when encryption is performed in software. However, if encryption is performed by an IPSec hardware offload network adapter, the ESP packets are decrypted when Network Monitor captures them and, as a result, they can be parsed and interpreted into the upper-layer protocols. If you need to diagnose ESP software-encrypted communication, you must disable ESP encryption and use ESP-null encryption by changing the IPSec policy on both computers.

**IPSec Registry Entries**
In Windows Server 2003, you can use the Netsh IPSec command-line tool to perform many of the tasks that you might have performed previously by modifying the registry.

When you enable success or failure auditing for the **Audit logon events** audit policy, IPSec records the success or failure of each main mode and quick mode negotiation and the establishment and termination of each negotiation as separate events. Keep in mind, however, that enabling this type of auditing can cause the security log to fill with IKE events. In Windows 2000, you cannot disable auditing of IKE events. In Windows Server 2003, however, you can disable auditing of IKE events by modifying the registry.

The following information is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows Server 2003 tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**DisableIKEAudits**
**Registry path**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Audit\

**Version compatibility**
Windows Server 2003

When you enable success or failure auditing for the **Audit logon events** audit policy, IPSec records the success or failure of each main mode and quick mode negotiation and the establishment and termination of each negotiation as separate events. However, enabling this type of auditing can cause the security log to fill with IKE events. For example, for servers that are connected to the Internet, attacks on the IKE protocol can fill the security log with IKE events. IKE events can also fill the security log for servers that use IPSec to secure traffic to many clients. To avoid this, you can disable auditing for IKE events in the security log by modifying the registry.

To disable auditing of IKE events in the security log, you must first create the **DisableIKEAudits** key and set the registry setting to a value of **1**.After making this change to the registry, you must either restart the computer or stop and then restart the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

**Note**

- Stopping and restarting the IPSec service can disconnect all of the computers that are using IPSec from the computer on which the service is stopped and it can prevent further communication with that computer. If you restart the IPSec service immediately, the retransmit behavior of TCP might cause the TCP-based communication to resume after new IKE and IPSec Security Associations (SAs) are established.

For more information about this registry entry, see the "Registry Reference for Windows Server 2003" in the Tools and Settings Collection.

**Related Information**

The following resources contain additional information that is relevant to this section.

- IPSec Policy Extension Technical Reference
- "Command-Line Reference for Windows Server 2003" in the Tools and Settings Collection
- "Registry Reference for Windows Server 2003" in the Tools and Settings Collection

# PKI Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PKI Technologies**

Organizations need enhanced security for data and strong credentials for identity management. You can use certificates to secure data and manage identification credentials from users and computers both within and outside your organization.

A public key infrastructure (PKI) is the combination of software, encryption technologies, processes, and services that enable an organization to secure its communications and business transactions. The ability of a PKI to secure communications and business transactions is based on the exchange of digital certificates between authenticated users and trusted resources.

You can design a PKI solution to meet the following security and technical requirements of your organization:

- **Confidentiality.** You use a PKI to encrypt data that is stored or transmitted.
- **Integrity.** You use a PKI to digitally sign data. A digital signature helps you identify whether another user or process modified the data.
- **Authenticity.** A PKI provides several authenticity mechanisms. Authentication data passes through hash algorithms, such as Shivest Hash Algorithm 1 (SHA1), to produce a message digest. The message digest is then digitally signed by using the sender's private key to prove that the message digest was produced by the sender.
- **Nonrepudiation.** When data is digitally signed, the digital signature provides proof of the integrity of the signed data and proof of the origin of the data. A third party can verify the integrity and origin of the data at any time. This verification cannot be refuted by the owner of the certificate that digitally signed the data.

**PKI Technologies Architecture**

The architecture of a PKI involves implementing various interdependent technologies and processes to make it possible to issue, validate, renew, and revoke certificates. These technologies include:

- One of more servers running Certificate Services and that provide certificate enrollment, revocation and other certificate management services.
- Active Directory directory service or another directory service that provides account management, policy distribution, and certificate publication services.
- Domain controllers that can authenticate end users and computers when they request certificates.
- Domain client computers and users, who request, receive, and use certificates for specific purposes. Although certificates can also be used by services and by non-domain clients, in most Windows PKI environments, domain users and computers are the primary recipients and users of certificates. In some cases, the domain client can be a subordinate CA that requests and receives a certificate authorizing it to issue certificates of its own.

The following figure illustrates the architecture of PKI technologies.

**PKI Technologies Architecture**




**PKI Technologies Components**

The key technology components of a PKI and their relation to the PKI architecture diagram are described in the following table.

**Components of a PKI**

| Components | Description |
|---|---|
| Certificates | Provide the foundation of a PKI. Digital certificates are electronic credentials that are associated with a public key and a private key that an organization uses to authenticate users. Certificates are created on servers running Certificate Services and stored on clients and in a directory such as Active Directory. |
| Certificate templates | Can be used to define the content and purpose of a digital certificate, including issuance requirements; implemented extensions, such as application policy or extended key usage; and enrollment permissions for certificates that a CA issues. Certificate templates are stored in the Active Directory and used by enterprise CAs to provide the default attributes for a certificate. |
| Certificate Services | The part of the core operating system that allows a business to act as its own CA, and issue and manage digital certificates. Certificate Services includes tools to manage issued certificates, publish CA certificates and CRLs, configure CAs, import and export certificates and keys, and recover archived private keys. |
| CAs | Servers on which Certificate Services has been configured to issue, validate, and manage certificates. Windows Server 2003 supports multiple levels of a CA hierarchy and a cross-certified trust network. This includes offline and online CAs. |
| Certificate Revocation Lists | List of certificates that a CA considers no longer usable. Certificates have a specified lifetime, but CAs can reduce this lifetime by a process known as certificate revocation. Publishers can use any kind of directory service, including X.500, Lightweight Directory Access Protocol (LDAP), or directories in a specific operating system, including Active Directory, to store CRLs. Publishers can also publish CRLs on Web servers. |
| Certificate policy and practice statements | The two documents that outline how a CA and its certificates are to be used, the degree of trust that can be placed in these certificates, legal liabilities if the trust is broken, and so on. These documents can also define or impact PKI designs, operations, and usage, including how a CA is configured, how client requests are processed, and guidelines and procedures for revoking certificates. |
| Certificate policies | Configurable limitations on the scope of a certificate. Certificate policies can be implemented as required and allowed certification path length, the range of namespaces that are permitted or excluded by a qualified subordinate CA, the extent to which an organization trusts the identity presented in a certificate, and the applications that can be used in conjunction with certain certificates. |
| Certificate and CRL repositories | A directory service or other location where certificates are stored and published. In a Windows Server 2003 domain environment, the Active Directory is the most likely publication point for certificates issued by Windows Server 2003–based CAs. |
| PKI-enabled applications | Examples of PKI-enabled applications include: Encrypting File System (EFS), Microsoft Internet Explorer, Microsoft Money, Internet Information Services (IIS), Routing and Remote Access, Microsoft Outlook, and Microsoft Outlook Express. Also included are a variety of third-party applications that work with Windows 2000 and Windows Server 2003 Certificate Services. |

**PKI Technologies Scenarios**

A PKI consists of one or more CAs that play specific roles. One or more CAs of each type can be combined in a CA hierarchy.

**Roles in a CA Hierarchy**

Each CA in a CA hierarchy is assigned a role, which is determined by the CA's location in the CA hierarchy. Common roles in a

CA hierarchy include a root CA, a policy CA, and an issuing CA.

**Root CAs**
A root CA is the highest CA in a CA hierarchy and is the trust point for all certificates that are issued by the CAs in the CA hierarchy. If a user, computer, or service trusts a root CA, they implicitly trust all certificates that are issued by all other CAs in the CA hierarchy.

A root CA is different from all other CAs in that it issues its own certificate. This means that the Issuer and Subject fields of the certificate contain the same distinguished name. A root CA only issues certificates to other CAs that are directly subordinate to it.

**Policy CAs**
A policy CA is typically located on the second-tier of a CA hierarchy, directly beneath the root CA. In this scenario, the root CA is often referred to as a parent CA, because the root CA issued a Subordinate Certification Authority certificate to the policy CA. In fact, any CA that issues a certificate to another CA is referred to as a parent CA. The CA that receives the certificate from a parent CA is known as a subordinate CA.

The role of a policy CA is to describe the policies and procedures that an organization implements to secure its PKI, the processes that validate the identity of certificate holders, and the processes that enforce the procedures that manage certificates. A policy CA issues certificates only to other CAs. The CAs that receive these certificates must uphold and enforce the policies that the policy CA defined.

It is not mandatory to use policy CAs unless different divisions, sectors, or locations of your organization require different issuance policies and procedures. However, if your organization requires different issuance policies and procedures, you must add policy CAs to the hierarchy to define each unique policy. For example, an organization can implement one policy CA for all certificates that it issues internally to employees and another policy CA for all certificates that it issues to non-employees.

**Issuing CAs**
An issuing CA is typically located on the lowest tier of a CA hierarchy. An issuing CA issues certificates to other computers, users, network devices, services, or other issuing CAs. An issuing CA is always online.

The parent CA for an issuing CA can be a root CA, a policy CA or another issuing CA. The issuing CA must enforce the policies and procedures that are described in the policy CA above the issuing CA in the CA hierarchy.

**Types of CA Hierarchies**
You can use one of two CA models: a root hierarchy or a cross-certification trust. Windows Server 2003 networks recognize and support both models.

**Root Hierarchies**
In a root CA hierarchy, all of the CAs in the organization's CA hierarchy are chained to a common root CA. The following figure illustrates a root CA hierarchy.

**Root CA Hierarchy**




A root CA hierarchy:

- Enhances security and scalability. It protects the upper layers of the CA hierarchy from network attacks by removing the upper layers of the CA hierarchy from the network.

- Provides flexible administration to the CA hierarchy. You can use role separation to delegate CA management to separate administration groups in an organization.

- Supports commercial CAs.All commercial CAs — such as VeriSign, GTE, Thawte, and RSA — implement trusted root CA hierarchies.

- Supports most applications.Applications such as Microsoft Internet Explorer and Netscape Communicator support certificates that root CA hierarchies issue, as do IIS and Apache Web servers.

**Cross-Certification Trusts**

In a cross-certification trust, a CA in one organization's root CA hierarchy issues a subordinate CA certificate to a CA in another organization's CA hierarchy. The following figure illustrates a cross-certification hierarchy.

**Cross-Certification Trust**




A cross-certification trust:

- Provides interoperability between businesses and between products. When cross-certification is implemented, the certificates are logically chained to the trusted root CA of the organization that is evaluating the presented certificate.

- Joins disparate PKI domains. You can issue a Cross Certification Authority certificate from any CA in your organization's hierarchy to any CA in a partner organization's CA hierarchy.

- Assumes complete trust of a foreign CA hierarchy. Cross-certification does not enforce any constraints on the certificates that a partner organization issues. You must implement qualified subordination to implement constraints on those certificates.

# CA Certificates Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certification authority (CA) certificates are certificates issued by a CA to itself or to a second CA for the purpose of creating a defined relationship between the two CAs. A certificate that is issued by a CA to itself is referred to as a trusted root certificate because it is intended to establish a point of ultimate trust for a CA hierarchy. Once the trusted root has been established, it can be used to authorize subordinate CAs to issue certificates on its behalf. Although the relationship between CAs is most commonly hierarchical, CA certificates can also be used to establish trust relationships between CAs in two different public key infrastructure (PKI) hierarchies. In all of these cases, the CA certificate is critical to defining the certificate path and usage restrictions for all end entity certificates issued for use in the PKI.

**In this subject**


- What Are CA Certificates?
- How CA Certificates Work
- CA Certificates Tools and Settings

# What Are CA Certificates?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certification authority (CA) certificates are certificates that are issued by a CA to itself or to a second CA for the purpose of creating a defined relationship between the two CAs.

A certificate that is issued by a CA to itself is referred to as a trusted root certificate, because it is intended to establish a point of ultimate trust for a CA hierarchy.

Once the trusted root has been established, it can be used to authorize subordinate CAs to issue certificates on its behalf.

Although the relationship between CAs is most commonly hierarchical, CA certificates can also be used to establish trust relationships between CAs in two different public key infrastructure (PKI) hierarchies.

In all of these cases, the CA certificate is critical to defining the certificate path and usage restrictions for all end entity certificates issued for use in the PKI.

**Problems Solved by CA Certificates**

When an end entity uses a certificate, a trust relationship must be verified between the end entity certificate and the root CA. This trust relationship between a root CA and an end entity certificate is verified by validating the contents of all of the certificates in the certificate chain up to the root CA.

The guidelines and procedures that have been established for the PKI define the trust. This includes the ability of end entity certificates to be used for certain purposes and prevented from being used for other purposes. These guidelines and procedures are implemented in a number of ways. For example, the trust and the security of the PKI can be established in two ways:

- The steps taken to ensure the physical security of the server that hosts the CA.
- The manner in which administrative roles for the CA are delegated.

However, the appropriate configuration of CA certificates for the organization's needs is one of the most powerful tools that an organization has to implement appropriate PKI security.. CA certificates contain special configuration data that regulate the CAs to which they are issued. These configuration options can:

- Define the organizational namespace in which certificates issued by the subordinate CA can be issued and trusted.
- Specify the acceptable uses of certificates issued by the subordinate CA.
- Define the issuance guidelines that must be followed in order for a certificate issued by the subordinate CA to be considered valid.
- Create a managed trust between separate certification hierarchies.

**Core CA Certificate Use Scenarios**

Trust and acceptable use are integrally connected: How much you trust a certificate should help you decide whether it is acceptable for a predetermined use at a specific time in a specific location.

This is the most basic security and business principle in PKI deployment: you must create defined trust relationships within an enterprise.. Once this is done, you can use defined trusts to enable certificate types of connections while limiting others. For example, they can:

- Prevent certificates from being used with unintended applications.
- Apply consistent certificate issuance policies.
- Apply consistent name formatting in issued certificates.

- Prevent the implementation of unauthorized subordinate CAs.

The following are some of the trust scenarios that can be enabled by properly configured CA certificates:

- Simple trust within an enterprise PKI
- Cross trust between two enterprise CAs (restricted)
- Bridge trust for multiple enterprise PKIs (unrestricted)
- Certificate trust continuity when the certificate policy changes

**Simple Trust Within an Enterprise**

Simple trusts between CAs in a single PKI are the most common type of certificate trust relationships. They impose clear hierarchical roles and relationships between the root CA, intermediate policy CAs, the CAs that issue end entity certificates (if any), and the end-entity certificates themselves.

The following diagram illustrates a simple PKI trust hierarchy in a single organization.




**Cross Trust Between Two CAs**

Cross trust between two CAs allows an organization to extend a simple PKI trust hierarchy to secondary PKI trust hierarchies within the same organization or within other organizations. You can implement cross trust by using special CA certificates called *cross certificates* and a strategy called *qualified subordination*. Qualified subordination allows you to place certificate issuance constraints on subordinate CAs and to place usage constraints on the certificates they issue. With qualified subordination, you can focus subordinate CAs according to specific certification needs and administer your public key infrastructure (PKI) more efficiently. With qualified subordination, each CA in a cross trust relationship can define and implement rules that do the following:

- Define the namespaces for which your PKI hierarchy will issue and accept trusted certificates
- Specify the acceptable uses of certificates issued by a cross-certified subordinate CA
- Define what issuance practices must be followed for a certificate issued by the qualified subordinate CA to be considered valid.

For example, suppose that, due to a merger, Organization 1 wants to share certain certificate-protected resources with users in the engineering department of Organization 2. They can accomplish this by creating a cross trust with a specific CA in Organization 2, but limit the locations and applications in organization 1 for which Organization 2's certificates can be used.

These cross CA trusts are one-directional. If Organization 2 wants to grant Organization 1 certificate users privileges within Organization 2, they must establish a separate cross trust in order to enable this certificate use.

The following figure illustrates one typical cross CA trust relationship. There are two-way trust relationships between the CAs in both Organization 1 and Organization 2, but trust between the two organizations is managed using two separate one-way trusts, so that each organization can determine the degree and type of trust it wants to share with the other organization.




Cross certificates restrict the certificates that are acceptable to your organization. They can be issued by either the qualified subordinate CA or by CAs that chain through the qualified subordinate CA,. The configuration options that you apply in the cross certificates between the two organizations allow you to:

- Limit trust to all CAs in your organization
- Limit trust to specific CAs in your organization.
- Restrict the name space available to certificate holders from the other organization
- Add new CAs to an existing CA hierarchy
- Apply application policy that defines the purposes for which certificates issued by a subordinate CA can be used
- Apply policy constraints
- Map CA policy in one organization to CA policy in another organization
- Create trust between two organizations
- Create trust between multiple organizations

**CA Certificate Dependencies and Interactions**

A number of factors can impact interoperability between a Windows Server 2003 CA and the CAs of other organizations. The interoperability of Windows Server 2003 PKIs with non-Microsoft PKIs is dependent on how the following features and standards are implemented in both PKIs:

- Standards support
- Algorithm and key length support
- CRL distribution points
- Authority information access (AIA)
- Authority key identifier (AKI)
- Certificate extensions
- Extended key usage (EKU)
- Directory integration

You must determine whether the other PKIs with which you need to establish trust support these features and standards, and how you can accommodate differences.

**Standards Support**

A number of technical standards provide a basis for interoperability between Windows Server 2003 and other PKI applications. To promote non-Microsoft interoperability with the Windows Server 2003 application programming interface (API), Microsoft

supports the following standards:

- **Public Key Infrastructure for X.509 Certificates (PKIX)**. Defines interoperable PKI standards for the Internet.
- **X.509**. Describes the standard format of a certificate.
- **Public-Key Cryptography System (PKCS)**. Provides a standard for public key message exchanges.
- **Transport Layer Security (TLS)**. Provides a secure and authenticated channel between hosts on the Internet above the transport layer.
- **Secure/Multipurpose Mail Extensions (S/MIME)**. Serves as a standard for secure e-mail across the Internet.
- **Kerberos authentication protocol**. Provides a symmetric key framework for authentication in large networks.
- **Personal Computer/Smart Card (PC/SC)**. Serves as a standard for integrating smart cards and smart card readers.

Most PKI vendors have adopted many or all of these PKI standards. Different vendors, however, can implement the standards in different ways. While it might be possible to link external PKI implementations to yours, this might involve making some changes to your existing design. For this reason, it is strongly recommended that you evaluate the external PKI to determine whether it meets all your critical requirements.

**Hardware Support**

Windows Server 2003 uses CryptoAPI to abstract hardware-based key management from applications, and it uses the PC/SC standard instead of PKCS #11 to communicate with smart cards and readers. Many non-Microsoft CAs have their own cryptographic APIs and use PKCS #11 to interface to hardware tokens such as smart cards. Because the Windows 2000 and Windows Server 2003 operating systems require hardware devices to support Plug and Play and power management features, and PC/SC includes support for these ease-of-use features, Windows Server 2003 does not support PKCS #11.

### Note

- The Windows Server 2003 PKI can use non-Microsoft cryptographic service providers (CSPs), and can enroll users for certificates that have keys that were generated by non-Microsoft CSPs.

**Algorithm and Key Length Support**

Windows 2000 and Windows Server 2003 support several well-known cryptographic algorithms that are also supported by other PKI products. Key length, in turn, is determined in part by the cryptographic algorithm that you select. When you install a Windows Server 2003 CA, you can select CA key lengths from 384 bits to 16,384 bits, depending on the CSP that you select. In a typical deployment, user certificates have 1,024-bit keys and root CAs have 4,096-bit keys.

In general, the lengths of public and private keys do not impact interoperability, as long as both environments support a common range of key lengths. If one PKI supports large public keys and another does not, however, the two cannot exchange symmetric keys or sign and verify data.

### Note

- While key exchange and digital signature operations performed between PKIs do not require the same public and private key lengths, symmetric key algorithms do. In addition, if different key lengths are used, both key lengths must be supported in both environments.

When PKIs do not support the same key lengths, some applications cannot decrypt data that other applications have encrypted. In addition, the PKIs might not be able to establish secure communications channels between applications if the applications cannot agree on symmetric key lengths, as required by protocols such as Secure Sockets Layer (SSL) and TLS.

If a PKI uses public key cryptography based on an algorithm such as RSA, all PKI operations can be accomplished with only one key pair. However, single key pairs might not meet the security requirements of your organization or its choice of algorithm. For this reason, Windows Server 2003 supports both single key pairs and dual key pairs. A good PKI is flexible enough to allow as many or as few key pairs as are required by applications.

If one PKI operates according to the number of keys that applications use, it can impact interoperability with other PKIs. For

example, an e-mail application could sign a message with a signature-only key and include the associated certificate in a message sent to a recipient without also sending an encryption certificate as part of the message. The recipient might then be unable to discover the encryption certificate of the sender to reply with an encrypted message back to the sender.

**CRL Distribution Points**
With certificate revocation lists (CRLs), the optional CRL distribution point (CDP) extension in a certificate identifies how revocation information for the certificate can be obtained. If a CRL distribution point is not always available, certificate chain building can be delayed, causing inconvenience for the user. If a CRL is not available at the distribution point that has been specified in the certificate, CRL retrieval might even fail and the certificate will be considered invalid.

You need to compare any non-Microsoft CRL support with the Windows Server 2003 CRL support. For example, the non-Microsoft PKIs might not support the Windows Server 2003 CRL process, which includes the use of delta CRLs. Conversely, the Windows Server 2003 PKI might not support the methods of the non-Microsoft PKI for processing CRL information. Also, a few non-Microsoft PKIs support partitioned CRLs, which most PKIs, including Windows Server 2003 PKIs, do not support.

**Authority Information Access**
The Authority Information Access (AIA) extension is a pointer to the most currently published CA parent certificate of a CA. The Windows Server 2003 PKI uses this extension to assist in building trust chains to validate client certificates.

Some non-Microsoft PKIs do not provide the AIA extension. In this case, parent certificates must be distributed to domain clients so that the certificates are available before the chain building process begins. Cross certificates must also be available locally on domain clients, because there is no information in a certificate that specifies where it can be found.

**Authority Key Identifier**
The Authority Key Identifier (AKI) extension provides a means to identify the public key of the CA that validates the signature on a CRL. This identification is based on either the subject key identifier (SKI) or the issuer name and serial number from the certificate that is issued by the CRL issuer. The AKI extension is useful in cases when a CRL issuer has more than one signing key.

An organization that expects its PKI certificates to be used by other Windows Server 2003 PKIs must populate the Authority Key Identifier extension with a unique key identifier and an issuer name and serial number. The Windows Server 2003 PKI attempts to construct certificate chains by using the issuer name and serial number in the AKI first, and then the subject key identifier.

**Certificate Extensions**
Not all certificate extensions are universally recognized. If a CA does not recognize a certificate extension in a request and it has been marked critical, it rejects the certificate. Unless you intend to limit the use of the certificate to a specific application that understands the critical extension, avoid putting critical extensions in certificates because it limits interoperability.

**Extended Key Usage**
The Extended Key Usage (EKU) extension indicates the purposes for which the public key that is contained in the certificate can be used. The Windows Server 2003 PKI uses the EKU extension to indicate certificates that support special functions, such as Internet protocol security (IPSec) and encrypting file system (EFS) file encryption backup. The EKU extensions of other organizations may be used for different purposes.

**Directory Integration**
Windows Server 2003 PKI certificates can be published to any directory or repository, although the default CA exit module only supports Active Directory. However, by default, a Windows Server 2003 PKI relies on Active Directory and LDAP for authentication, including smart card logons and certificate autoenrollment, as well as for certificate management.

With Microsoft Certificate Services, certificates issued by a non-Microsoft CA can be associated with a Windows Server 2003 user account stored in Active Directory. This is possible because applications such as Internet Explorer and Internet Information Services (IIS) can be used to authenticate a user to an account stored in Active Directory, based on the user principal name (UPN) information in a certificate.

The account to which the certificate maps provides information about user access rights on the server. This is an extremely powerful feature for Web-based applications and non-Microsoft CAs because it combines strong authentication by means of public key technology with the native authorization model of Windows Server 2003. For example, to enable extranet and remote access scenarios without requiring the application and certificate to manage access rights, administrators can use certificates from partner companies and map them to accounts in Active Directory by means of one-to-one or many-to-one mappings.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Certificates Technical Reference

- Certificate Services Technical Reference

# How CA Certificates Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certification authority (CA) certificates are certificates that are issued by one CA to another CA. These CA certificates become a part of the certificate trust hierarchy, the certificate path from end-entity certificates to the trusted root CA certificate.

The first CA certificate issued in a public key infrastructure (PKI) is a root certificate, issued by a CA to itself. Once a root CA has been created, it can be used to issue, sign, and validate CA certificates that are issued to other CAs.

Most commonly, root CAs are used to issue CA certificates to subordinate CAs in a PKI hierarchy. These subordinate CAs, in turn, can issue their own CA or end-entity certificates.

However, CA certificates can also be used to establish trust between two or more PKI hierarchies. CA certificate-based trust relationships can connect PKIs in one organization, in two organizations, or spanning multiple organizations.

Because of their critical role in establishing trust between CAs and in the certificate validation process, CA certificates are extremely powerful and critical elements of an organization's security strategy. For this reason, CA certificates are typically configured with a variety of policy constraints to strictly define their acceptable use and to prevent their unacceptable use.

The information in the CA Certificates Technical Reference is interrelated with the information in the Certificate Services Technical Reference. The information in these two Technical Reference documents should be taken together to gain a full understanding of how Certificate Services can be implemented in Microsoft Windows 2000 and Windows Server 2003 environments.

**Ideal Environment for CA Certificates**
In a PKI, trust is established when a valid certification path can be traced from an end-entity certificate to an issuing CA and ending with a trusted root certificate. Trust in CA certificates is established by carefully defining and controlling:

- Where and how certificates can be used.
- The applications that can be configured to use certificates.
- Requirements that must be met before a certificate is issued.
- Acceptable formats for subject names in certificates.
- Equivalence between certificate policies in multiple organizations.

For more information about the mechanisms that are available to define and configure CA certificate restrictions, see the following section, "CA Certificate Architecture."

**CA Certificate Architecture**
In many ways, the architecture of CA certificates is identical to the architecture of end-entity certificates. However, CA certificates also have some unique features.

**CA Certificate Shared Features**
All certificates that are issued in Windows Server 2003 and earlier PKIs are structured to comply with standards established by the Public-Key Infrastructure (X.509) Working Group (PKIX) of the Internet Engineering Tasks Force (IETF). The following figure shows the contents of X.509 version 3 certificates, as implemented in Windows Server 2003.

**Structure of an X.509 Certificate**




X.509 version 3 certificates are the current certificate format in a Windows Server 2003 PKI. Version 3 certificates support the following fields that have been supported since X.509 version 1:

- **Subject.** Provides the name of the computer, user, network device, or service that the CA issues the certificate to. The subject name is commonly represented by using an X.500 or Lightweight Directory Access Protocol (LDAP) format.

- **Serial Number.** Provides a unique identifier for each certificate that a CA issues.

- **Issuer.** Provides a distinguished name for the CA that issued the certificate. The issuer name is commonly represented by using an X.500 or LDAP format. For a root CA, the Issuer and Subject are identical. For all other CA certificates and for end entity certificates, the Subject and Issuer will be different.

- **Valid From.** Provides the date and time when the certificate becomes valid.

- **Valid To.** Provides the date and time when the certificate is no longer considered valid.

  **Note**

  - The date when an application or service evaluates the certificate must fall between the Valid From and Valid To fields of the certificate for the certificate to be considered valid.

- **Public Key.** Contains the public key of the key pair that is associated with the certificate.

  In addition to the fields defined in X.509 version 1, X.509 version 3 certificates include optional fields or extensions that provide additional functionality and features to the certificate. These extensions are not necessarily included in each certificate that a CA issues:

- **Subject alternative name.** A subject can be presented in many different formats. For example, if the certificate must include a user's account name in the format of an LDAP distinguished name, e-mail name, and a user principal name (UPN), you can include the e-mail name or UPN in a certificate by adding a subject alternative name extension that includes these additional name formats. Subject alternative name is only used in end entity certificates, not in CA certificates.

- **Basic constraints.** This X509 version 3 extension is used to distinguish between end-entity certificates and CA certificates. There are still some PKI clients today that do not recognize basic constraints, which can make it possible for an end-entity to act as a CA. Windows Server 2003 and Windows 2000 operating systems honor basic constraints in accordance with Internet Engineering Task Force (IETF) Request for Comments (RFC) 2459 and will reject CA certificates that do not contain this extension.

- **Name constraints.** This extension restricts the namespaces that are permitted or excluded by a qualified subordinate CA and its subordinates when issuing certificates.

- **Policies.** Defines the list of acceptable issuance and application policies for certificate usage. These policies are identified in the certificate by object identifiers (also known as OIDs).

- **Policy mapping.** Allows a policy from one domain to be mapped onto a policy of another domain.

- **Policy constraints.** Restricts the subordination levels in a certificate hierarchy to which a policy is applied. These extensions are used in conjunction with issuance and application policies only.

- **Application policy.** Defines which applications can be used in conjunction with certain certificates.

- **Application policy mapping**. Identifies equivalence between the application policies of two organizations that cross certify by using certificate application policies.

- **Cross certificate distribution points. I**dentifies where cross certificates related to a particular certificate can be obtained and how often the cross certificates at that location are updated.

- **CRL distribution points (CDP).** Provides one or more URLs where the application or service can retrieve a certificate revocation list (CRL) from. Used when an application or service must determine whether a certificate has been revoked before its validity period has expired.

- **Authority Information Access (AIA).** Provides one or more URLs from where an application or service can retrieve the issuing CA certificate. Used to validate the certificate of the CA that issued the certificate — also referred to as the parent CA — for revocation and validity.

- **Enhanced Key Usage (EKU).** Defines which applications can be used in conjunction with certain certificates. Because some implementations of public key infrastructure (PKI) applications might not understand application policies, both application policies and enhanced key usage sections appear in certificates issued by a Microsoft CA

Many of these extensions are configured using the Certification Authority snap-in, Certreq.exe, and Certutil.exe, which are described in the CA Certificate Tools section. They are also discussed in greater detail in the Certificate Services Technical Reference. Other extensions can be configured by creating CAPolicy.inf and Poliy.inf files that must be installed on a CA as part of the CA setup and renewal process. These extensions and their configuration options are discussed in greater detail in the following section.

#### CA Certificate Extensions

Many people assume that simply accepting the default configuration options presented by the Setup user interface when configuring a new root CA or a subordinate CA automatically provides a degree of security appropriate for their requirements. In fact, few organizations have identical security requirements.

The CA setup routines only present a portion of the security options that can be used to configure new root and subordinate CAs. Many end-entity certificate configuration options can be included in certificate requests. However, many of the most important configuration constraints must exist on the server hosting the CA before the CA is installed or renewed. These CA-based certificate constraints make it possible to address a variety of security needs and concerns. These CA certificate-based constraint extensions are configured using two files, CAPolicy.inf for root CAs and Policy.inf for subordinate CAs.

**Note**

- Some of these extensions can also be configured using the command line tools Certutil.exe and Certreq.exe. For a complete understanding of the capabilities of Certutil and Certreq, see the Windows Server 2003 Help.

The following section will describe all the options available in these two files so that you can create CAPolicy.inf and Policy.inf files tailored to your specific needs.

#### Overview of CAPolicy.inf and Policy.inf

A CAPolicy.inf file is a configuration file that defines the extensions, constraints, and other configuration settings that are applied to a root CA certificate and all certificates issued by the root CA. The CAPolicy.inf file must be installed on a host server before the setup routine for the root CA begins. When the security restrictions on a root CA are to be modified, the root certificate must be renewed and an updated CAPolicy.inf file must be installed on the server before the renewal process begins.

A Policy.inf file is a configuration file that defines the constraints that are applied to a subordinate CA certificate and all certificates issued by the subordinate CA. The constraints can include basic constraints, name constraints, application policies, and certificate policies. If a subordinate CA is being used to issue cross certificates connecting two different PKIs, these restrictions serve an essential role in ensuring that the cross certification process does not detract from security.

CAPolicy.inf and Policy.inf files:

- Are created and defined manually by an administrator.

- Are utilized during the creation of root and subordinate CA certificates.

- Are defined on the signing CA where you sign and issue the certificate, not the CA where the request is generated.

There are, however, important differences between CAPolicy.inf and Policy.inf files. For example, CAPolicy.inf files must be placed in the *Systemroot* folder before installation begins. Policy.inf, on the other hand, can exist in any folder on the requesting computer. In addition, the Policy.inf file can use any file name, as long as the syntax is correct.

CAPolicy.inf and Policy.inf files make it possible to specify and configure a wide variety of CA attributes and options. This makes it very difficult to define a default template that administrators should use for all purposes. Therefore, the following section will describe all the options in order to enable you to create an .inf file tailored to your specific needs.

The following terms are used to describe the .inf file structure:

- A *section* is an area in the .inf file that covers a logical group of keys. Section names in .inf files are identified by appearing in brackets. Many, but not all, sections in an .inf file are used to configure certificate extensions.
- A *key* is the name of an entry. It appears to the left of the equal sign.
- A *value* is the parameter. It appears to the right of the equal sign.

In the following example, [RequestAttributes] is the section, AttributeName1 is the key, and AttributeValue1 is the value:

```
[RequestAttributes]
AttributeName1 = AttributeValue1
```

The following table describes the possible sections of CAPolicy.inf and Policy.inf files. It is followed by a more detailed discussion of the configuration options that can be configured using each section of these files.

| **Section** | **Explanation** | **Applies to** |
|---|---|---|
| [Version] | Identifies the file as an .inf file. | Setup and renewal of root and subordinate CAs |
| [RequestAttributes] | Defines the certificate template to use when creating new certificates. | Setup and renewal of root and subordinate CAs |
| [NameConstraintsExtension] | Defines the permitted and excluded namespaces for certificates issued by a subordinate CA. | Setup and renewal of root and subordinate CAs |
| [PolicyExtension] | Lists user-defined policies | Setup and renewal of root and subordinate CAs |
| [Policy MappingsExtension] | List of user-defined policy mappings. | Setup and renewal of root and subordinate CAs |
| [PolicyConstraintsExtension] | Specifies the depth of the CA hierarchy that requires and inhibits policy mapping. | Setup and renewal of root and subordinate CAs |
| [ApplicationPolicyExtension] | Defines the applications or services a certificate can be used for. An alternate to the Enhanced Key Usage extension | Setup and renewal of root and subordinate CAs |
| [ApplicationPolicyMappingsExtension] | Lists user-defined application policy mappings. | Setup and renewal of root and subordinate CAs |
| [ApplicationPolicyConstraintsExtension] | Defines the depth of the CA hierarchy that requires and inhibits policy mapping. | Root and subordinate CAs |
| [BasicConstraintsExtension] | Sets the maximum subordinate CA path length | Root and subordinate CAs |
| [CrossCertificateDistributionPointsExtension] | Identifies where cross certificates related to a particular certificate can be obtained and how often that location is updated. | Primarily on subordinate CAs that are used to issue cross certificates. |
| [EnhancedKeyUsageExtension] | Defines the applications or services a certificate can be used for. | An alternate to the Application Policy extension |
| [AuthorityInformationAccess] | Specifies the authority information access points . | Initial setup and renewal of a root CA |
| [CRLDistributionPoints} | Used to specify CRL distribution points (CDPs) | Initial setup and renewal of a root CA |
| [certsrv_server] | Used to specify renewal key length, the renewal validity period, and the certificate revocation list (CRL) validity period | New CA setup and CA renewal. |

| | | |
|---|---|---|
| [NewRequest] | Used .inf files that serve as templates for new certificate requests. | New CA setup and CA renewal. |
| [RequestAttribute] | Used in conjunction with [NewRequest] to define additional attributes associated with template-based certificate requests. | New CA setup and CA renewal. |

The policy constraint options that you use, and where you choose to implement them, depends on your unique requirements.

**Note**

- The use of each policy option can be configured as mandatory (Critical = True, Critical=Yes, or Critical=0) or optional (Critical = False, Critical=No, Critical=1). If an extension is marked as critical and the validating computer cannot parse the extension, it must reject the certificate chain. These values are not case sensitive.

**Version**

The version section identifies the file as an .inf file.

If present, Version must contain the Signature key. The Signature key must be equal to the fixed string "$Windows NT$"

```
Signature= "$Windows NT$"
```

**Request Attributes**

Request attributes are used to pass significant name and value pairs to the policy module of a CA. Any pairs of significant can be passed to a policy module, including pairs that are of use to a customer policy module developed specifically for an organization. For example, request attributes can be used to specify the certificate template to use when creating a new certificate, such as:

```
CertificateTemplate = CrossCA
```

In addition, it can also be used to predefine attributes and values associated with any new certificate, in the format:

```
AttributeName1 = AttributeValue1
AttributeName2 = AttributeValue2
```

**Name Constraints**

Name Constraints is an optional extension that allows organizations to designate which namespaces are either permitted or excluded for certificates issued by a subordinate CA. When you create a new CA, you can define name constraints for the CA by configuring CAPolicy.inf. These name constraints apply to all certificates that are issued in the PKI. Similarly, when you create a Cross Certification Authority certificate, you define name constraints in the Policy.inf file to apply to all certificates issued in the second PKI that will be considered valid in your PKI hierarchy.

**Note**

- Name constraint validation is performed on the CA, not on the client. However, your clients must run Windows XP or Windows Server 2003 in order to use name constraints.

When a CA that is configured with name constraints receives a request, it compares the names present in the subject and the subject alternate name fields of the request to the configured name constraints, to determine whether the namespace is permitted or excluded.

The following table includes examples of permitted and excluded name constraint settings.

| Element | Description | Sample Values |
|---|---|---|
| NameConstraints<br><br>Excluded Name Constraints Extension | Identifies whether permitted name types and excluded name types have been configured. | Include = NameConstraintsPermitted<br><br>Exclude = NameConstraintsExcluded |

| | | |
|---|---|---|
| [NameConstraintsPermitted] | List of user-defined permitted name formats. | DNS = .corp.contoso.com<br><br>email = @contoso.com<br><br>email = @exchange.contoso.com<br><br>UPN = .corp.contoso.com<br><br>UPN = @corp.contoso.com<br><br>DIRECTORYNAME = "DC=com, DC=Contoso, DC=Corp"<br><br>DIRECTORYNAME = "C=US, S=WA, L=Redmond, O=Contoso"<br><br>DIRECTORYNAME = "O=Contoso, OU=NorthAmerica" |
| [NameConstraintsExcluded] | List of user defined excluded name types | DNS = .development.contoso.com<br><br>email = @development.contoso.com<br><br>UPN = .development.contoso.com<br><br>UPN = @development.contoso.com<br><br>DIRECTORYNAME = "DC=com, DC=Contoso, DC=Development"<br><br>DIRECTORYNAME = "C=US, O=Contoso, OU=Development" |

Before deploying a PKI, organizations need to decide which individual clients and business units are able to enroll for and use certain certificates.

For many organizations, the selected users, computers, and services are members of specific Active Directory domains and subdomains.

Name constraints can be based on any of the following types of name formats:

- X.500 directory name
- DNS domain name
- Universal Resource Identifier (URI)
- E-mail and user principal name
- Internet protocol (IP) address

Additional name constraint forms may be enforced using the Other Name forms, which are identified by name and object identifier that are encoded with UTF8 or ASN.1. Other Names forms can be used to represent names that are outside of the standard name types, such as Rfc822Name, DNS, or X.500 Directory Name. Other Names consist of an object identifier and a binary large object. The most common Other Name is the Universal Principal Name (UPN), which has the constant OID_NT_PRINCIPAL_NAME as the object identifier. The version 2 "Directory E-mail Replication" template certificates and the Version 1 "Domain controller" certificates have a subjectaltname extension that contains the constant OID_NTDS_REPLICATION as other name type and the globally-unique identifier (GUID) representation of the domain controller as the binary large object. In Windows 2000, only the UPN was present in the other name constraint. However, in Windows Server 2003 operating systems, the other name constraint has been expanded to support even the domain controller GUIDs. The following is the syntax where the first part is the object identifier, followed by a qualifier identifying the encoding of the following data as UTF8, octet, or ASN.1:

```
OtherName=1.2.3.4.99.100,{utf8}ssss
OtherName=1.2.3.4.99.101,{octet}ABCD
OtherName=1.2.3.4.99.102,"{asn}BAgAAQIDBAUGBw=="
OtherName=1.3.6.1.4.1.311.25.1
```

One scenario where this might be useful is when you have cross-forest smartcard logon, wherein the other domain has Windows 2000 domain controllers. In this case, you can provide an additional name constraint as "OtherName=1.3.6.1.4.1.311.25.1" to allow all Domain E-Mail Replication certificates and restrict the namespace using the DNS name constraint. It is technically possible to restrict the GUID of a domain controller as another name but that would not be a generic constraint but a specific constraint to only one domain controller.

**Relative Distinguished Name Constraints**

Relative distinguished names are used to identify the names of objects stored in directories, such as Active Directory or X.500 directories.

Using relative distinguished name constraints allows you to restrict a qualified subordinate CA to issue certificates only to specific users or computers in Active Directory by using the object's relative distinguished name. The distinguished name can be very specific, where the permitted or excluded relative distinguished name refers to a specific object, or the distinguished name can be a wildcard value that references all the objects within a specific container, domain, or organizational unit.

The name constraint must be included in either the Policy.inf file used to create a qualified subordinate CA or in the CAPolicy.inf file used when installing a new subordinate CA. The following table shows some examples of LDAP name constraints and defines what objects the constraints refer to.

| Relative Distinguished Name Constraint | Includes |
|---|---|
| DC=example,DC=contoso,DC=com | All objects in the example.contoso.com domain |
| CN=Users,DC=example,DC=Contoso,DC=com | All objects in the Users container in the example.contoso.com domain |
| CN=Brian Komar,CN=Users,DC=contoso,DC=com | The user object Brian Komar in the Users container in the contoso.com domain |
| OU=Marketing,DC=Contoso,DC=Com | All objects in the Marketing organizational unit (OU) of the Contoso.com domain |
| CN=Servers,CN=Sites,CN=Configuration,DC=Contoso,DC=Com | All domain controllers in the contoso.com forest |
| OU=Domain Controllers, DC=Tacoma,DC=Contoso,DC=Com | All domain controllers in the tacoma.contoso.com domain. |

The range of characters in the allowed name constraint is defined by the relative distinguished name type. Most relative distinguished names (such as O=, OU=,CN=, and so on) support international (UTF8) characters, where relative distinguished names that include (DC=,C=,email, and so on) support on IA5 character types. A relative name constraint will appear in either the permitted or excluded name constraints using the following format:

```
DIRECTORYNAME = "DC=NWtraders, DC=com"
```

**DNS Name Constraints**

You can apply the DNS namespaces that your network uses for name resolution as name constraints for a qualified subordinate CA. When the qualified subordinate CA receives a certificate request, it compares the DNS name that is associated with the computer requesting the certificate to its DNS name constraints and decides whether or not to issue a certificate.

DNS name constraints compare any DNS names that are contained in certificate requests to the list of permitted and excluded DNS name constraints. As with LDAP distinguished name constraints, the name in the certificate request must match one of the permitted names for the certificate to be issued.

The DNS domain names that are accepted by the subordinate CA must follow the standard DNS naming conventions as specified in RFCs 1034 and 1035, and are listed in either the Policy.inf or CAPolicy.inf file used. Therefore, the DNS name constraint cannot contain international character sets. For more information about RFCs 1034 and 1035 see the entries in the IETF RFC Database.

When configuring a DNS name constraint, you can either designate a specific DNS host name, for example, host.example.contoso.com, or you can designate a DNS namespace. For example, the DNS name constraint .example.contoso.com indicates all hosts that have a DNS name that ends with example.contoso.com. The preceding period (.) indicates that any DNS name that ends with example.contoso.com matches the name constraint.

**Note**

- This does not include a host that does not have a period preceding example.contoso.com. For example, the host named ourexample.contoso.com would not result in a match when compared to the name constraint .example.contoso.com.

The following table demonstrates the results for submitted DNS names against configured DNS name constraints.

| Submitted DNS Name | DNS Name Constraint | Evaluation |
|---|---|---|
| www.example.contoso.com | .example.contoso.com | Match |
| www.south.example.contoso.com | .example.contoso.com | Match |
| example.contoso.com | .example.contoso.com | No Match |
| Ourexample.contoso.com | example.contoso.com | No Match |
| www.example.contoso.com | example.contoso.com | Match |

**Uniform Resource Identifier Constraints**

Uniform Resource Identifiers (URIs) are used to identify resources on the Internet by means of identifiers such as URL, FTP, HTTP, telnet,

mailto, news, and gopher. URLs in CAPolicy.inf should use the replaceable parameter syntax that also appears in the **Extensions** tab of the Certification Authority snap-in. If you don't use the replaceable parameter syntax, then CDP and AIA extensions in renewed CA root certificates might point to the same location that was specified in the previous CA root certificate.

When validating the URI names in a certificate request, the qualified subordinate CA ignores the protocol element in the URI, such as http:// or ftp://, and uses only the domain or host names.

The URI naming conventions supported for name constraints must follow the syntax specified in RFC 2396. For more information about name constraints, see RFC 2396 in the IETF RFC Database.

When a qualified subordinate CA validates a URI name, it ignores the protocol element in the URI. For example, if the submitted URI is http://www.example.contoso.com and the URI name constraint is URL=ftp://.example.contoso.com, the actual comparison performed is between the host name www.example.contoso.com and the name constraint .example.contoso.com. In this case, the result would be a match, even though the protocols differ.

As with DNS constraints, a URI constraint can reference either a host or a domain. If a preceding period is included in the URI constraint, the domain name must include the entire suffix defined in the name constraint. For example, if the URI constraint is .example.contoso.com, then matches are achieved for west.example.contoso.com and for east.example.contoso.com.

The most frequent application of URI name constraints is the validation of certificate requests for Web Server certificates, where the Web servers URL is submitted in the certificate request. For example, if the Web Server certificate request has a name of http://www.contoso.com and the qualified subordinate CAs permitted URI name constraint is URL=http://www.contoso.com, then this URI constraint would permit the certificate request for the Web Server certificate.

A URI name constraint will appear in either the permitted or excluded name constraints using the following format:

```
URL = "http://.contoso.com"
```

**RFC 822, E-mail, and UPN Constraints**

You can specify RFC 822, e-mail and UPN name constraints for an individual subject, such as person@example.contoso.com, or you can specify constraints for all subjects whose e-mail names or UPNs end in a specific name, such as @example.contoso.com. Typically, you need to specify e-mail or UPN name constraints for all subjects whose e-mail addresses and UPNs end in a specific name.

RFC 822 and e-mail constraints must follow RFC 822 naming conventions and use IA5 encoding of character strings. User Principal Names (UPNs) use the same syntax but are encoded using UTF8. Although RFC 822 names, UPNs, and e-mail addresses share the same syntax, you must provide separate examples in the name constraints listing of Policy.inf to differentiate between requests using e-mail addresses and requests using UPN names.

The following table shows some examples of RFC 822, e-mail, and UPN constraints and how they are evaluated.

| Submitted Address | UPN/E-Mail Constraint | Evaluation |
|---|---|---|
| jsmith@nwtraders.com | jsmith@nwtraders.com | Match |
| bjsmith@nwtraders.com | jsmith@nwtraders.com | Match |
| jsmith@nwtraders.com | @nwtraders.com | Match |
| jbsmith@nwtraders.com | bjsmith@nwtraders.com | No Match |

**Note**

- UPN name constraints should always be entered as two separate entries in the list of name constraints. One entry should include the at sign (@, such as in @nwtraders.com); the second entry should replace the at sign with a period (.), so this entry would be .nwtraders.com. This format allows for the possibility of having a UPN of subdomain.nwtraders.contoso.

**IP Address Constraints**

IP address constraints allow you to specify either specific IP addresses or ranges of IP addresses that a CA can successfully receive certificates from a qualified subordinate CA. The format of the IP addresses in the constraints must follow either RFC 791 (for IPv4 addresses) or RFC 2460.

When a certificate request is received by a qualified subordinate CA, the source IP address in the certificate request is compared to the IP address constraints to determine if the IP address is permitted or excluded by the name constraints.

To specify a range of IP addresses, you must specify both the IP address and associated subnet mask in each constraint. For example, if you wanted to create an IP address constraint for the 192.168.3.0 network, the constraint would be 192.168.3.0/255.255.255.0.

An IP address constraint will appear in either the permitted or excluded name constraints using the following format:

```
IPADDRESS = 192.168.3.0/255.255.255.0
IPADDRESS = 192.168.2.254/255.255.255.255
IPADDRESS = 172.16.8.0/255.255.248.0
```

The first constraint includes all computers in the 192.168.3.0 network. The second constraint is a host-specific constraint that only includes the host with IP address 192.168.2.254. The final constraint is a Classless Internet Domain Routing (CIDR) example that includes all hosts with addresses between 172.16.8.1 and 172.16.15.254.

Name constraints can be configured to result in the following outcomes:

- **Permitted.** The certificate request contains all names that are listed as permitted in the CA name constraints extension of the issuer.
- **Not permitted.** The certificate request contains a name that is not listed as permitted in the name constraints extension of the issuer.
- **Excluded.** The certificate request contains a name that is listed as excluded in the name constraints extension of the issuer.

When name constraints are added to a CA certificate, the following rules are applied to the subject name and alternate subject name fields:

- **Excluded namespaces take precedence over permitted namespaces.** A qualified subordinate CA will not issue a certificate to a user within an excluded namespace even if the user is also within a permitted namespace. For example, a user might be within the permitted Active Directory namespace **.xyz.com** but also within the excluded DNS namespace **.uvw.xyz.com**. The excluded DNS namespace overrides the permitted Active Directory namespace and the certificate request of the user fails.
- **If the name constraints extension exists in a CA certificate, all name constraints must be present in the appropriate format.** Any name formats that are not included are considered to be wildcards that match all possibilities. For example, if the DNS name constraint is absent, the entry is treated as DNS="".

  **Note**

  - You can disable this requirement by modifying the following registry setting: **HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\SystemCertificates\Root\ProtectedRoots** . For more information, see "CA Certificate Tools and Settings."
- **All name constraints are considered, even if they are not specified.** No precedence is applied to the listed name constraints. For this reason, name constraints that are not present are treated as wildcards. For example if you only restrict the DNS name space, the Name Constraints extension sets the remaining name constraints to allow all name spaces.
- **Name constraints are applied to the Subject Name extension and any existing Subject Alternate Name extensions.** For example, if a user can be identified by a DNS domain name and an alternate e-mail name, name constraints apply to both.
- **Name constraints apply to all names contained in a certificate request.** Each name in the subject or subject alternate name extensions must match at least one of the name constraints listed for that name type. A certificate request that includes a subject name or subject alternate name that does not match a listed name type is rejected.
- **Name constraints are not case sensitive.** The comparison is not case-sensitive if the name or part of the name is encoded with an IA5 string. If it is encoded as a Unicode or UTF8 string, the comparison will be a binary match of the content.
- **A qualified subordinate CA's name constraints cannot permit names excluded by its parent CA.** For example, if the parent CA excludes the DNS name .contoso.com, then its qualified subordinate CA cannot permit the DNS sub-domain .example.contoso.com.
- **Constraints apply only when the namespace types specified as name constraints are present in the certificate request.** If no namespace of the specified types is in the certificate request, the certificate is acceptable. For example, a CA is configured with the following name constraints:

```
[NameConstraintsExtension]
Include = NameConstraintsPermitted
Exclude = NameConstraintsExcluded

[NameConstraintsPermitted]
DirectoryName = "DC=Contoso, DC=Com"
email = @contoso.com
UPN = .contoso.com
UPN = @contoso.com

[NameConstraintsExcluded]
```

In this configuration, if the CA receives a request with the e-mail name of user1@northwindtraders.com, the request would be rejected.

Likewise, the request would be rejected if the subject alternate name contained user1@northwindtraders.com while the subject name contained CN=user1,OU=Contractors,DC=Contoso,DC=Com. Remember that all subject names must match the name constraints for a certificate to be issued. The default policy as mentioned previously requires all names to be defined unless otherwise configured on the certification authority.

**Note**

- If you do not specify permitted or excluded names, you are creating a wildcard (*) name constraint for those name constraint formats.

When the name constraint processing is completed, the validation process will result in one of the following outcomes:

- **Permitted.** The end certificate contains a name that is listed as permitted in an issuer's name constraints extension.
- **Not permitted**. The end certificate contains a name that is not listed as permitted in an issuers name constraints extension.
- **Excluded**. The end certificate contains a name that is listed as excluded in an issuer's name constraints extension.
- **Not Defined**. The issuer certificate does not list a constraint for a specific name type, such as Directory Name or IP Address.

**Note**

- Windows Server 2003 CAs will have constraints set to permit all namespaces for a particular name type, unless specified otherwise in the name constraints extension in the Policy.inf

**Policy**

You can use issuance policies to define the extent to which your organization trusts the identity presented in a certificate. For example, you can set an issuance policy stipulating that you only trust certificates that were issued during a face-to-face meeting with a network administrator, such as when a smart card certificate is issued.

**Note**

- *Issuance policy* is the Microsoft term for Certificate Policies defined in part 4.2.1.5 in RFC 2459 and is contained in the certificatePolicies extension of a certificate.

An object identifier must describe every issuance policy that you define. The inclusion of an issuance policy object identifier in an issued certificate indicates that the certificate was issued in a manner that meets the issuance requirements associated with the issuance policy object identifier.

**Note**

- Issuance policy is only available on Windows Server 2003 CAs. Windows 2000 does not provide issuance policy.

The following table includes examples of issuance policy settings.

| Element | Description | Sample Values |
|---|---|---|
| [PolicyStatementExtension] | Lists the policies that have been defined by the organization, and whether they are optional or mandatory | Policies = MediumAssurancePolicy, LowAssurancePolicy, RequireCertificatePolicy |

| | | |
|---|---|---|
| [MediumAssurancePolicy] | Requires that each policy that is included must be identified by an object identifier. In addition, a Notice or URL key can be included for each policy. | OID = 1.3.6.1.4.1.311.21.8.1308480216.1095248452.3429159812.3095816201.1472143667.3857372322<br><br>Notice = "Legal policy statement text" |
| [LowAssurancePolicy] | Requires that each policy that is included must be identified by an object identifier. In addition, a Notice or URL key can be included for each policy. | OID = 1.3.6.1.4.1.311.21.8.1308480216.1095248452.3429159812.3095816201.1585704526.3613270325<br><br>URL = "http://http.site.com/some where/default.asp"<br><br>URL = "ftp://ftp.site.com/some where else/default.asp"<br><br>Notice = "Limited use policy statement text."<br><br>URL = "ldap://ldap.site.com/some where else again/default.asp" |
| [RequireCertificatePolicy] | Requires that each policy that is included must be identified by an object identifier. In addition, a Notice or URL key can be included for each policy. | OID = 1.3.6.1.4.1.311.21.15<br><br>URL = http://extra.site.com/Extra Policy/default.asp<br><br>[oidpolicy]<br><br>OID = 1.3.6.1.4.1.311.21.55 |
| [LegalPolicy] | Requires that an organization's legal policy documents must be identified with at least one object identifier, and zero or more Notice and URL keys. Legal policy statements usually need to be identified in all certificates an organization issues or accepts. | OID = 1.3.6.1.4.1.311.21.43<br><br>Notice = "Legal policy statement text" |

| | | |
|---|---|---|
| [LimitedUsePolicy] | Used by many organizations to limit the allowed uses of certificates they issue or accept. These usage limitations can be identified with at least one object identifier and 0 or more Notice and URL keys. | OID = 1.3.6.1.4.1.311.21.47<br><br>URL = "http://http.site.com/some where/default.asp"<br><br>URL = "ftp://ftp.site.com/some where else/default.asp"<br><br>Notice = "Limited use policy statement text."<br><br>URL = "ldap://ldap.site.com/some where else again/default.asp" |
| [ExtraPolicy] | Used by organizations to define additional policies beyond legal and use restrictions. These can also be defined with at least one object identifier and 0 or more Notice and URL keys. | OID = 1.3.6.1.4.1.311.21.53<br><br>URL = http://extra.site.com/Extra Policy/default.asp |

You can use a specific certificate template to define one or more issuance policy object identifiers that need to be included in any certificates issued. Windows Server 2003 includes four predefined issuance policies:

- **All Issuance (2.5.29.32.0).** The all issuance policy indicates that the issuance policy contains all other issuance policies. Typically, this object identifier is only assigned to CA certificates. A root CA automatically implies that all issuance policies are defined.

  **Note**

  - The x.y.z portion of the object identifier in these examples is a randomly-generated numeric sequence that is unique for each Windows Server 2003 forest.

- **Low Assurance (1.3.6.1.4.1.311.21.8.** *x.y.z.* **1.400).** The low assurance object identifier represents certificates that are issued with no additional security requirements.

- **Medium Assurance (1.3.6.1.4.1.311.21.8** *.x.y.z.* **1.401).** The medium assurance object identifier represents certificates that have additional security requirements for issuance. For example, a smart card certificate that is issued in a face to face meeting with a smart card issuer might be considered a medium assurance certificate and contain the medium assurance object identifier.

- **High Assurance (1.3.6.1.4.1.311.21.8** *.x.y.z.* **1.402).** The high assurance object identifier represents certificates that are issued with the highest security. For example, the issuance of a key recovery agent certificate might require additional background checks and a digital signature from a designated approver, because a person holding this certificate can recover private key material from a Windows Server 2003, Enterprise Edition CA.

In addition, you can create your own object identifiers to represent custom issuance policies. For example, two organizations that are involved in a purchaser and seller relationship can define custom object identifiers to represent digital signature certificates for specific purchase amounts. In such a case, an object identifier can be defined for purchase between $100,000 and $500,000 and another object identifier can be defined for purchases greater than $500,000. Applications can then use these object identifiers to recognize whether a person had the appropriate signing authority for a specific volume purchase.

You can provide additional information about the issuance policies implemented at a CA by configuring policy qualifiers. Policy qualifiers

provide information directly, or provide links to information, that describe the purpose of the issuance policy. The following code shows both formats of the policy qualifier:

```
[LegalPolicy]
OID = 1.3.6.1.4.1.311.21.43
Notice = "Legal policy statement text"
URL = "http://www.example.contoso.com/policy/isspolicy.asp"
```

When viewing the certificate in an application, a person can follow the URL and read the description of the issuance policy in a Web browser. If a text message is used as a policy qualifier, the application will display the message.

It is recommended that URLs rather than notice text be implemented for certificate practice statements. This allows practice statements to be modified without having to re-issue the CA certificate. In addition, URLs do not dramatically increase the size of certificates compared to notice text.

**Policy Mappings**

You can use policy mappings to explicitly state the comparability of certificate policies in two PKI hierarchies. Policy mapping is commonly used to cross-certify two PKIs, such as between two organizations.

In many cases, the administrators of two PKIs define their own policies and object identifiers. In some cases, these policies are identical, but in most cases there are small differences between them. For example, one organization might stipulate that one physical form of identification is sufficient to grant a certificate request, while a second organization requires three forms of physical identification to grant a similar request. In these cases, you need to negotiate with the administrators of the other PKI to define terms of equivalence before the cross-certified relationship can be established. This is called policy mapping.

Policy mapping enables interoperability between two organizations that apply similar issuance and application policies, but have deployed different object identifiers.

The following table includes examples of policy mappings.

| Element | Description | Sample Values |
|---|---|---|
| [PolicyMappingsExtension] | Each entry maps one foreign policy object identifier to a local policy object identifier. | 1.3.6.1.4.1.311.21.53 = 1.2.3.4.87<br><br>1.3.6.1.4.1.311.21.54 = 1.2.3.4.89 |

The first object identifier is foreign, the second is local. If the policy object identifier (for example, 1.2.3.4) of one company represents a specific function, and the policy object identifier of another company (for example, 11.22.33.44) represents the same function, they can be mapped, so that 11.22.33.44 and 1.2.3.4 are interchangeable.

The qualified subordinate CA that contains this mapping is called the *issuer CA* and the subordinate CA whose policies have been mapped is called the *subject CA*. In mapping some or all of the policies of the subject CA to the policies of the issuer, the issuer CA effectively subordinates the subject CA. The result of this mapping is that users and computers in the issuer CA trust hierarchy can use their own certification paths to validate certificates from users and computers in the subject CA trust hierarchy. The separate trust hierarchy can be within the same intranet or in separate PKI environments over an extranet.

Policy mappings can be applied by:

- Identifying the trust hierarchies that are to be linked with a trust relationship.
- Establishing equivalence in the assurance levels that are used by the two trust hierarchies involved in the trust relationship.
- Obtaining the issuance and application policy object identifiers used in both trust hierarchies involved in the trust relationship.
- Mapping the issuance and application policy object identifiers in the separate trust hierarchies, and defining policy constraints in the CA certificate request for the qualified subordinate CA that you are installing in your trust hierarchy.
- Installing the qualified subordinate CA with the policies, policy mappings, and policy constraints in your trust hierarchy.

**Policy Constraints**

The policy constraints extension specifies the depth of the CA hierarchy that requires and inhibits policy mapping. Require policy constraints define the levels of a PKI hierarchy where foreign certificates with mapped policy can be used. Inhibit policy constraints specify the boundaries beyond which these certificates cannot be used.

The following table includes examples of policy constraints settings.

| Element | Description | Sample Values |
|---|---|---|
| [PolicyConstraintsExtension] | Specifies the depth of the CA hierarchy that requires policy mapping | RequirePolicyMapping = 3 |

| | | |
|---|---|---|
| | Specifies the depth of the CA hierarchy that inhibits policy mapping. | InhibitPolicyConstraints = 5 |

You can refine policy mapping by setting parameters for how the issuance policy defined in qualified subordination affects other CAs below the qualified subordinate CA. These parameters can help to limit unplanned trust relationships. The following two settings define this relationship:

- **Require explicit policy.** Specifies the number of certificates that can exist in the hierarchy below a certificate before an explicit policy must exist. For example, if the explicit policy is configured with a setting of 3, the defined issuance policy must exist for three layers of the hierarchy. The CA on which the qualified subordination is defined is the first level.

- **Inhibit policy mapping.** Specifies the number of additional certificates that can appear in the path before policy mapping is no longer permitted. For example, an inhibit policy mapping value of 3 restricts the policy mapping to only three levels of CAs below the qualified subordinate CA.

**Application Policy**

Certificates provide important information that is not specific to an application. However, you might need to define which applications can be used in conjunction with certain certificates. Application policy allows you to ensure that certificates are only used with the applications that you specify.

Application policies included in certificates allow applications to determine if a certificate can be used when authenticating a user, encrypting data, or signing a device driver. An application can be coded to only accept certificates that contain specific application policies. When the application receives signed information from a user, the application will review the certificate associated with the private key used to sign the information and verify that the certificate chain has the required object identifier as a valid application policy.

Each application policy has one object identifier, and can include Notice and URL keys, as in the examples in the following table.

| Element | Description | Sample Values |
|---|---|---|
| [ApplicationPolicyStatementExtension] | Specifies the application policies that apply to the specified certificates. | Policies = AppLegalPolicy, AppLimitedUsePolicy, AppExtraPolicy, AppOIDPolicy |
| [AppLegalPolicy] | Specifies specific legal requirements and implications of the certificate. The notice spells out these legal requirements and implications in terms understandable by the user. | OID = 1.3.6.1.4.1.311.21.54<br><br>Notice = "Application Legal policy statement text" |
| [AppLimitedUsePolicy] | Many certificates will be issued for restricted uses, such as for signing purchase orders up to a certain value. The sample URLs would provide explanatory text for these usage restrictions. | OID = 1.3.6.1.4.1.311.21.58<br><br>URL = "http://http.contoso.com/application some where/default.asp"<br><br>URL = "ftp://ftp.contoso.com/application some where else/default.asp"<br><br>Notice = "Application Limited use policy statement text."<br><br>URL = "ldap://ldap.contoso.com/application some where else again/default.asp" |
| [AppExtraPolicy] | Many certificates will be issued for restricted uses, such as for signing purchase orders up to a certain value. The sample URLs would provide explanatory text for these usage restrictions. | OID = 1.3.6.1.4.1.311.21.64<br><br>URL = http://extra.contoso.com/Application Extra Policy/default.asp |
| [Appoidpolicy] | An optional format to use for an application policy. | OID = 1.3.6.1.4.1.311.21.66 |

Application policy is Microsoft-specific and is treated much like Extended Key Usage. Because some implementations of PKI applications might not understand application policies, both application policies and enhanced key usage sections appear in certificates issued by a Microsoft CA.

If a certificate has an extension containing an application policy and also has an EKU extension, the EKU extension is ignored. If, however, a certificate has only an EKU extension, the EKU extension is treated like an application policy extension. If a certificate has an application policy extension and an EKU property, the effective policy for the certificate is the common policy between the EKU property object identifiers and the application policy object identifiers.

**Note**

- You can view all defined application policy object identifiers in the Certificate Templates console. To do this, in the console tree, right-click **Certificate Templates** and click **View Object Identifiers**.

After you define the object identifiers for your organization's certificate policies, obtain the complementary object identifiers from the partner organization. Obtain the partner's object identifiers because the object identifiers differ between the two organizations.

The following Application Policy object identifiers are recognized by Windows Server 2003.

- Server Authentication (1.3.6.1.5.5.7.3.1)
- Client Authentication (1.3.6.1.5.5.7.3.2)
- Code Signing (1.3.6.1.5.5.7.3.3)
- Secure E-mail (1.3.6.1.5.5.7.3.4)
- IP Security End System (1.3.6.1.5.5.7.3.5)
- IP Security Tunnel Endpoint (1.3.6.1.5.5.7.3.6)
- IP Security User (1.3.6.1.5.5.7.3.7)
- Time Stamping (1.3.6.1.5.5.7.3.8)
- IP Security IKE Intermediate (1.3.6.1.5.5.8.2.2)
- All Application Policies (1.3.6.1.4.1.311.10.12.1)
- Microsoft Trust List Signing (1.3.6.1.4.1.311.10.3.1)
- Qualified Subordination (1.3.6.1.4.1.311.10.3.10)
- Key Recovery (1.3.6.1.4.1.311.10.3.11)
- Document Signing (1.3.6.1.4.1.311.10.3.12)

**Note**

- A list of available object identifiers can be viewed in the **Certificate Templates** console. To do this, in the console tree, right-click **Certificate Templates** and click **View Object Identifiers**. Alternatively, custom application policy object identifiers can be defined and used in custom-developed applications.

**Application Policy Mappings**

When qualified subordination is configured between two CAs that use certificate application policies, you must map the object identifiers between the two organizations in the policy file that you create. Application policy mapping ensures that only authorized application policy object identifiers from a partner organization are allowed in certificates that the partner organization issues. Application policy mapping associates the partner organization's application policy object identifiers with application policy object identifiers that are defined in your organization's PKI.

The following example shows how application policy mappings are configured in CAPolicy.inf or a Policy.inf file. Each entry maps one foreign application policy object identifier to one local application policy object identifier. The first object identifier is foreign, the second is local.

| Element | Description | Sample Values |
|---|---|---|
| [ApplicationPolicyMappingsExtension] | List of user defined application policy mappings. | 1.3.6.1.4.1.311.21.64 = 1.2.3.4.98<br><br>1.3.6.1.4.1.311.21.65 = 1.2.3.4.100 |

**Application Policy Constraints**

The application policy constraints extension specifies the depth of the CA hierarchy that requires and inhibits policy mapping. Require application policy constraints define the levels of a PKI hierarchy where foreign certificates with mapped policy can be used. Inhibit application policy constraints specify the boundaries beyond which these certificates cannot be used.

Policy qualifiers are typically URLs that provide information directly or provide links to information that describe the purpose of the certificate policy. The following table shows examples of application policy constraints:

| Element | Description | Sample Values |
|---|---|---|
| [ApplicationPolicyConstraintsExtension] | Defines the depth of the CA hierarchy that requires application policy mapping | RequirePolicyMapping = 6 |
| | Defines the depth of the CA hierarchy that inhibits policy constraints | InhibitPolicyConstraints = 10 |

**Basic Constraints**

A basic constraint defines which CAs your organization trusts in a partner's CA hierarchy by limiting the path length for a certificate chain.

The following table shows examples of basic constraints:

| Element | Description | Sample Values |
|---|---|---|
| [BasicConstraintsExtension] | Defines the path limits for validating certificate chains. | PathLength = 3 |
| [EnhancedKeyUsageExtension] | | OID = 1.3.6.1.5.5.7.3.4 Secure E-mail<br>OID = 1.3.6.1.5.5.7.3.1 Server Authentication |

Basic constraints allow an application to determine whether a certificate is a CA certificate, which can then be used by the certificate chain engine to build certification paths, or an end-certificate, which cannot.

You can also use basic constraints to limit the maximum number of CA certificates that can be included in a CA path. For example, setting a path length of 0 in the basic constraints section of the CAPolicy.inf file allows only certificates issued by that specific CA to be included in the CA path. A path length of 2 allows only a total of three CA certificates in a certification path. In the latter case, any certification paths that include more than three CAs are discarded.

Use basic constraints if you do not want to trust additional CAs that are created lower in the CA hierarchy of your organization. You can also use basic constraints in cross-certified relationships if you trust your business partner and the certificates from all their existing CAs, but you do not want to trust certificates from any additional CAs that they authorize.

**Cross Certificate Distribution Points**

The cross certificate distribution point (CCDP) extension identifies where cross certificates related to a particular certificate can be obtained and how often that location is updated. Windows XP and later operating systems use this extension for the discovery of cross-certificates that might be used during the path discovery and chain building process.

**Note**

- The CCDP extension is unique to Windows CAs.

The following table shows examples of configuration options for cross certificate distribution point extensions.

| Element | Description | Sample Values |
|---|---|---|
| CrossCertDistPoints | One or more locations where the cross certificates related to a particular certificate can be found. | URL = http://%1/Public/My CA.crt<br>URL = ftp://contoso.com/Public/MyCA.crt<br>URL = file://\\%1\Public\My CA.crt |
| syncDeltaTime | An optional integer field that specifies the delta between when this location will be refreshed. If syncDeltaTime is present, CrossCertDistPointNames must also be present. | SyncDeltaTime = 24 |

This extension should be non-critical, but support for this extension by CAs and applications is recommended.

**Enhanced Key Usage**

Application policies are sometimes called *extended key usage* or *enhanced key usage*. Because some implementations of PKI applications might not understand application policies, both application policies and enhanced key usage sections appear in certificates issued by a Microsoft CA.

If a certificate has an extension containing an application policy and also has an EKU extension, the EKU extension is ignored. If, however, a certificate has only an EKU extension, the EKU extension is treated like an application policy extension. If a certificate has an application policy extension and an EKU property, the effective policy for the certificate is the common policy between the EKU property object identifiers and the application policy object identifiers.

This section can be used during CA setup or CA certificate renewal, but is only used if you are setting up a root CA or renewing a root CA certificate. The Enhanced Key Usage extension for a subordinate CA is determined by the CA which issued the subordinate CA's certificate. An example of this section is:

```
[EnhancedKeyUsage]
OID=1.2.3.4.5
OID=1.2.3.4.6
```

**Authority Information Access**
You can specify the authority information access points in CAPolicy.inf. This section is only used if you are setting up a root CA or renewing a root CA certificate. Subordinate CA authority information access extensions are determined by the CA which issued the subordinate CA's certificate. This section can be used during CA setup or CA certificate renewal.

Multiple URLs are supported. HTTP, FILE, FTP, and LDAP URLs are supported. If no URLs are specified (in other words, if the Authority Information Access section is empty), the default Authority Information Access extension will be suppressed.

The syntax is: `[AuthorityInformationAccess]URL="http://CompanyWebSite/Public/myCA.crl"`

**CRL Distribution Points**
You can specify CRL distribution points (CDPs) in CAPolicy.inf. Any CDP in CAPolicy.inf will take precedence for certificate verifiers over the CDP's specified in the CA policy module.

Some additional notes about the CRL section:

- Multiple URLs are supported.
- HTTP, file, FTP, and LDAP URLs are supported.
- This section can be used during CA setup or CA certificate renewal.
- This section is only used if you are setting up a root CA or renewing a root CA certificate. Subordinate CA CRL distribution point extensions are determined by the CA that issued the subordinate CA's certificate.
- URLs with spaces must be surrounded by quotes.
- If no URLs are specified (in other words, if the CRL Distribution Point section is empty), the default CRL Distribution Point extension will be suppressed.

If you want to specify the CDP using CAPolicy.inf, a sample of the syntax is:

```
[CRLDistributionPoint]
URL="http://CompanyWebSite/Public/myCA.crl"
```

Replacement tokens are used to retain the configuration of distribution points flexible. You can use replacement tokens in the CAPolicy.inf file and in the Certification Authority MMC in CA Extensions.

A replacement token consists of the percent sign and a number. This behavior occurs if you use replacement tokens in the Certificate Services MMC or if you use them in a Certutil.exe command. If replacement tokens are used in a batch file, and you use the percent sign (%), you must use another escape sign when needed, because the Windows shell typically interprets a percent sign as a command-line parameter.

The mapping of replacement tokens is different in versions of Windows later than Windows 2000 Server. You can use the following tokens for CRLDistributionPoint, AuthorityInformationAccess, and CrossCertificateDistributionPointsExtension URLs. The following table describes the replacement token options for Windows 2000 and Windows Server 2003.

| **Token name** | **Description** | **Windows 2000 map value** | **Windows Server 2003 map value** |
|---|---|---|---|
| ServerDNSName | The DNS name of the CA server | %1 | %1 |
| ServerShortName | The NetBIOS name of the CA server | %2 | %2 |
| CaName | The name of the CA | %3 | %3 |
| Cert_Suffix | The renewal extension of the CA | %4 | Not applicable. |
| CertificateName | | Not applicable. | %4 |
| Domain_Name | The location of the domain root in Active Directory | %5 | N/A |
| (Not used) | | N/A | %5 |
| ConfigurationContainer | The location of the configuration container in Active Directory | %6 | %6 |

| | | | |
|---|---|---|---|
| CATruncatedName | The "sanitized" name of the CA, 32 characters with a hash on the end | %7 | %7 |
| CRLNameSuffix | The renewal extension for the CRL | %8 | %8 |
| DeltaCRLAllowed | | | %9 |
| CDPObjectClass | | | %10 |
| CAObjectClass | | | %11 |

The Certification Server setup process replaces all %number% sequences with the appropriate value.

**certsrv_server**

Another section of CApolicy.inf that is not used frequently is [certsrv_server], which is used to specify renewal key length, the renewal validity period, and the certificate revocation list (CRL) validity period for a CA that is being renewed or installed.

Options that can be configured in this section include:

- **RenewalKeyLength**, which sets key size for renewal only. This is only used when CA renewal generates new keys.
- **RenewalValidityPeriod** and **RenewalValidityPeriodUnits**, which establish the lifetime of new root CA certificate when renewing the old CA certificate.
- **CRLPeriod** and **CRLPeriodUnits**, which establish the validity period for the full CRL, while **CRLDeltaPeriod** and **CRLDeltaPeriodUnits** establish the validity period for the delta CRL.
- **RenewalKeyLength**, **RenewalValidityPeriod** and **RenewalValidityPeriodUnits**, which are only processed when the CA is being renewed.
- **LoadDefaultTemplates**, which is not processed when the CA is a subordinate CA.

An example would be:

```
[certsrv_server]
RenewalKeyLength=2048
RenewalValidityPeriod=Years
RenewalValidityPeriodUnits=5
CRLPeriod=Days
CRLPeriodUnits=2
CRLDeltaPeriod=Hours
CRLDeltaPeriodUnits=4
```

**NewRequest**

The [NewRequest] section is mandatory for an .inf file that acts as a template for a new certificate request. If this section is missing, the following error message is displayed:

**INF file line not found 0xe0000102 (INF: -536870654)**

This section requires at least one key with a value. If this section is empty and has no keys, the following error message is displayed:

**Incorrect function. 0x1 (WIN32: 1)**

The following options can be associated with a new certificate request.

**EncipherOnly**

This parameter has an impact on the functional capabilities of the public-private key pair. If the value is set to **True**, the key can exclusively be used for encryption. If this is set to **False**, the key can be used for encryption and other purposes. This parameter refers only to the key type in CryptoAPI and has only an indirect relationship to the Key Usage extension that might be included in an issued x.509 certificate. The default value of the certificate enrollment control is set to False. This setting only affects AT_KEYEXCHANGE key types and then correspondingly limits the key usage to key encryption or data encryption. The digital signature and non-repudiation key usages will be disallowed.

| **Summary Information** | **Details** |
|---|---|
| Syntax | EncipherOnly={Boolean} |
| Values | **True** or **False** |
| Default value | **False** |
| Sample | EncipherOnly = True |

| | |
|---|---|
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | Yes (through the key usage extension) |

**Exportable**

This parameter is ignored when the **UseExistingKeySet** key is set to **True** because you can set the exportable flag only when a new key is created. You cannot change this flag for an existing key. If this attribute is set to **True**, the private key can be exported together with the certificate. To ensure a high level of security, private keys should not be exportable, however in some cases it might be required to make the private key exportable if several computers or users must share the same private key. For example, a Web-server that runs Secure Sockets Layer (SSL) and is published with Microsoft Internet Security and Acceleration Server to the Internet requires a certificate with an exportable key, because the certificate and key must be installed on the Web-server and the computer running ISA Server as well. Also, if you want to manage and back up the keys, they must be exportable to provide for this functionality.

This setting correlates with the template configuration if a Windows Server 2003 enterprise CA is used. A certificate template administrator can explicitly define if private keys should be exportable or not. The template setting for exportable has no effect on the exportability properties of a key generated through the **certreq.exe -new** command.

| **Summary Information** | **Details** |
|---|---|
| Syntax | Exportable = {Boolean} |
| Values | **True** or **False** |
| Default value | **False** |
| Sample | Exportable = **True** |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**KeyContainer**

It is not recommended to set this parameter for new requests where new key material is generated. The key container is automatically generated and maintained by the system.

For requests where the existing key-material should be used, this value can be set to the key-container name of the existing key.

If the **UseExistingKeySet** key is set to **True**, the key container must be set explicitly in the .inf file or in the renewal certificate.

| **Summary Information** | **Details** |
|---|---|
| Syntax | KeyContainer = {Key_containerName} |
| Values | Random string value (GUID) |
| Default value | None |
| Sample | KeyContainer = {C347BD28-7F69-4090-AA16-BC58CF4D749C} |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**KeyLength**

The key length defines the length of the public and private key. The key length has an impact on the security-level of the certificate. Greater key length provides usually a higher security level, however some applications might have limitations regarding the key length.

| **Summary Information** | **Details** |
|---|---|
| Syntax | KeyLength = {integer} |
| Values | Any valid key length that is supported by the cryptographic service provider |
| Default value | 1024 (depends on CSP) |

| | |
|---|---|
| Sample | KeyLength=2048 |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**KeySpec**
KeySpec determines if the key can be used for signatures, for exchange (encryption) or for both. If the **KeySpec** is set to a value of 2, the **EncipherOnly** key is ignored.

| Summary Information | Details |
|---|---|
| Syntax | KeySpec = {integer} |
| Values | AT_EXCHANGE: 1<br>AT_SIGNATURE: 2 |
| Default value | 2 |
| Sample | KeySpec=1 |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**KeyUsage**
Key usage defines what the certificate key should be used for. The value is a bit field that is composed of the key usage flags as they are defined in the Windows Platform Software Development Kit.

### Note

- The key usage extension has an indirect dependency with the extended key usage extension, so these two extensions need to align. For example, if you specify "Encrypting File System" as the extended key usage, you should also specify "Key Encipherment" as the key usage.

To combine multiple key usages, summarize the values of the individual key usages by performing hexadecimal arithmetic. The default key usage is set to **Digital Signature (0x80)** and **Non-Repudiation (0x40)**. The hexadecimal sum of 0x80 and 0x40 results in 0xC0.

The key usage parameter is primarily useful if certificate requests are submitted to stand-alone CAs. Only enterprise CAs have certificate key usages predefined through certificate templates. The key usage extension is set to Critical by default.

| Summary Information | Details |
|---|---|
| Syntax | KeyUsage = {hexadecimal_value} |
| Values | CERT_DIGITAL_SIGNATURE_KEY_USAGE: 0x80<br>CERT_NON_REPUDIATION_KEY_USAGE: 0x40<br>CERT_KEY_ENCIPHERMENT_KEY_USAGE: 0x20<br>CERT_DATA_ENCIPHERMENT_KEY_USAGE: 0x10<br>CERT_KEY_AGREEMENT_KEY_USAGE: 0x08<br>CERT_KEY_CERT_SIGN_KEY_USAGE: 0x04<br>CERT_OFFLINE_CRL_SIGN_KEY_USAGE:0x02<br>CERT_CRL_SIGN_KEY_USAGE: 0x02<br>CERT_ENCIPHER_ONLY_KEY_USAGE: 0x01 |

| | |
|---|---|
| Default value | 0xC0 for AT_SIGNATURE (default)<br>For AT_KEYEXCHANGE, EncipherOnly = TRUE: 0x30<br>For AT_KEYEXCHANGE, EncipherOnly = FALSE: 0xf0 |
| Sample | KeyUsage=0xa0 |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | Yes |

**MachineKeySet**

This key is important when you need to create certificates that are owned by the computer and not a user. The key-material that is generated is maintained in the security context of the security principal (user or computer account) that has created the request. When a CA administrator creates a certificate request on behalf of a computer, the key-material must be created in the machines security context and not the administrators' security context. Otherwise, the computer cannot access its private key, since it would be in the administrator's security context.

You must set this key to **True**, if you are creating requests for domain controllers, Web servers or any other service that runs in the computer's security context.

| **Summary Information** | **Details** |
|---|---|
| Syntax | MachineKeySet = {Boolean} |
| Values | **True** or **False** |
| Default value | FALSE |
| Sample | MachineKeySet = True |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**PrivateKeyArchive**

The PrivateKeyArchive setting works only if the corresponding RequestType is set to CMC**,** because only the CMC request format allows for securely transferring the requester's private key to the CA for key archival. Only the Windows Server 2003 enterprise CA supports private key archival. The Windows Server 2003 enterprise CA must be online and accessible to retrieve the CA encryption certificate directly during this process.

| **Summary Information** | **Details** |
|---|---|
| Syntax | PrivateKeyArchive = {Boolean} |
| Values | **True** or **False** |
| Default value | **False** |
| Sample | PrivateKeyArchive = True |
| Supported with CA type | Windows Server 2003 |
| Can be manipulated in a pending request | No |

**ProviderName**

The provider name is the display name of the certificate service provider (CSP). If you do not know the provider name of the CSP you are using, run certutil –csplist at a command line prompt. The command displays the names of all CSPs that are available on the local system.

| **Summary Information** | **Details** |
|---|---|
| Syntax | ProviderName = {CSP_stringname} |
| Values | The descriptive name of the certificate service provider |

| | |
|---|---|
| Default value | Microsoft Strong Cryptographic Provider |
| Sample | ProviderName=Microsoft RSA SChannel Cryptographic Provider |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**ProviderType**

The provider type is used to select specific providers based on specific algorithm capability such as "RSA Full". If you do not know the provider type of the CSP you are using, type certutil –csplist at a command-line prompt. The command displays the provider type of all CSPs that are available on the local system.

| Summary Information | Details |
|---|---|
| Syntax | ProviderType = {integer} |
| Values | The number that describes the provider type |
| Default value | 12 |
| Sample | ProviderType=13 |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**RenewalCert**

If a CA administrator needs to renew a certificate that exists on the system where the certificate request is generated, the CA administrator must specify its certificate hash as a value for this key. You can only renew certificates that have not expired. Certificates that have expired cannot be renewed and must be replaced with a new certificate.

**Note**

- If you do not know the certificate hash, use the Certificates MMC Snap-In and look at the certificate that should be renewed. Open the certificate properties and see the **Thumbprint** attribute of the certificate. Certificate renewal requires either a PKCS#7 or a CMC request format.

| Summary Information | Details |
|---|---|
| Syntax | RenewalCert={CertificateHash} |
| Values | The certificate hash of any certificate that is available at the computer where the certificate request is created |
| Default value | Not applicable. |
| Sample | RenewalCert=4EDF274BD2919C6E9EC6A522F0F3B153E9B1582D (You must use double quotes if white space is used between pairs of hexadecimal values) |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**Requestername**

The requester name can be specified for certificate requests if the **RequestType** is set to PKCS7 or CMC. If the **RequestType** is set to PKCS10, this key will be ignored. The **Requestername** can only be set as part of the request. CA administrators cannot manipulate the **Requestername** in a pending request.

| Summary Information | Details |
|---|---|
| Syntax | Requestername=samAccountName in Active Directory |

| | |
|---|---|
| Values | The domain name of the user who requested the certificate. |
| Default value | Not applicable. |
| Sample | Requestername = "DOMAINNAME\username" |
| Supported with CA type | Windows 2000 (with a PKCS#7 request type only)<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**RequestType**

The **RequestType** determines the standard that is used to generate and send the certificate request.

| **Summary Information** | **Details** |
|---|---|
| Syntax | RequestType={string value} |
| Values | CMC<br>PKCS10<br>PKCS10-<br>PKCS7 |
| Default value | PKCS10 |
| Sample | RequestType = CMC |
| Supported with CA type | Windows 2000 (except CMC request type)<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**Silent**

By default, the **Silent** option gives the CSP access to the interactive user desktop of the CA administrator and request information such as a smartcard personal identification number (PIN) from the CA administrator. If this key is set to **True**, the CSP must not interact with the desktop and will be blocked from displaying any user interface to the CA administrator.

| **Summary Information** | **Details** |
|---|---|
| Syntax | Silent={Boolean} |
| Values | **True** or **False** |
| Default value | **False** |
| Sample | Silent = **True** |
| Supported with CA type | Windows 2000<br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**Subject**

Several applications rely on the subject information in a certificate. Thus, it is recommended to specify a value for this key. If the subject is not set here, you should include a subject name as part of the Subject Alternate Name certificate extension.

| **Summary Information** | **Details** |
|---|---|
| Syntax | Subject={String} |
| Values | Relative distinguished name string values |
| Default value | Not applicable. |

| | |
|---|---|
| Sample | Subject=CN=computer1.contoso.com<br><br>Subject=CN=John Smith,CN=Users,DC=Contoso,DC=com x |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | Yes |

**UseExistingKeySet**

The *UseExistingKeySet* parameter is used to specify that an existing key pair should be used in building a certificate request. If this key is set to **True**, you must also specify a value for the *RenewalCert* key or the **KeyContainer** name. You must not set the key to *Exportable* because you cannot change the properties of an existing key. In this case, no key-material is generated when the certificate request is built.

| Summary Information | Details |
|---|---|
| Syntax | UseExistingKeySet={Boolean} |
| Values | **True** or **False** |
| Default value | **False** |
| Sample | UseExistingKeySet=**True** |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**UserProtected**

If this key is set to **True**, a password must be set and typed in whenever a private key is used. The CA administrator can choose to protect the key with a password when the key is generated or choose to display a password dialog box only when the key is used within an application.

| Summary Information | Details |
|---|---|
| Syntax | UserProtected={Boolean} |
| Values | **True** or **False** |
| Default value | **False** |
| Sample | UserProtected=**True** |
| Supported with CA type | Windows 2000<br><br>Windows Server 2003 |
| Can be manipulated in a pending request | No |

**[RequestAttributes]**

This section is used in conjunction with the **New Request** attribute to define additional attributes associated with the certificate request. All keys from this section can be specified as part of the `certreq -submit -attrib` command. For example, instead of specifying the CertificateTemplate in the .inf instruction file, a CA administrator could enter the following at a command-line prompt:

```
certreq -submit -attrib "CertificateTemplate:Domaincontroller" dcname.req
```

**CertificateTemplate**

The value of this key instructs the CA which certificate template to use when processing the certificate request. The name must be specified as the common name of the certificate template or the object identifier of the template, not the template display name.

| Summary Information | Details |
|---|---|
| Syntax | CertificateTemplate = {Templatename} |
| Values | Any valid certificate template name |
| Default value | Not applicable. |
| Sample value | Critical = “DomainController” |

| | |
|---|---|
| Can be manipulated in a pending request | Yes |

**SAN**

The subject alternate name (SAN) can be specified in various ways. An SPN is specified as a UPN value.

| Summary Information | Details |
|---|---|
| Syntax | SAN = "{tag}={value}&{tag}={value}]..."<br><br>Always include double quotes around values |
| Values | A list of valid subject alternate name extension tags.<br><br>*email, upn, dns, guid, url, ipaddress, oid*<br><br>Alternatively, any valid object identifier can be specified. |
| Default value | Not applicable. |
| Sample value | See the following table for sample values. |
| Can be manipulated in a pending request | Yes |

The value is a string type by default but can be prefixed with a tag in section brackets ({}) that determines the data type of the value explicitly as exemplified in the following table:

| SAN string: | Explanation |
|---|---|
| 2.5.29.17={octet}MBmCF1cySzMtQk8tREMuY29udG9zbzluY29t | An ASN.1 string is put into the SAN as with one or several encoded extensions.<br><br>The octet is a wrapper around a base64 string that is encoded in binary format |
| 1.2.3.4={utf8}SomeReadableText | An ASN.1 string is put into the SAN as with one or several encoded extensions.<br><br>The value is a string encoded with a UTF8 wrapper. |
| 2.5.29.17={asn}=3019821757 | An ASN.1 string is put into the SAN as with one or several encoded extensions. |
| email=John.Smith@contoso.com | The user's e-mail address is put into the SAN. |
| dns=w2k3bodc.contoso.com | The DNS name is set to computers FQDN. |
| dn="CN=W2K3BODC,OU=Domain Controllers,DC=contoso,DC=com" | The distinguished name is set to the domain controller's distinguished name. |
| url="http://www.contoso.com/default.html" | The URL is set to an HTTP URL. |
| ipaddress=172.134.10.134 | The IP is put into the SAN. |
| oid=2.5.29.17 | A given object identifier is put into the SAN. |
| upn=John.Smith@contoso.com | The UPN name represents the user. |
| guid=f7c3ac41-b8ce-4fb4-aa58-3d1dc0e36b39 | A GUID is put into the SAN. |

**Combining Extensions**

Many of the configuration options in the CA certificate extensions that are described above overlap. For example, you might need to decide whether to configure certain policy restrictions in the Policy Extension or in the Application Policy Extension. In addition, with many of these extensions and restrictions, organizations might need to decide whether to implement them at the root CA level or in a subordinate CA.

As a general rule, policy should be configured as high in the PKI hierarchy as possible. However, it is rare to put constraints in the root CA certificate itself. Instead, if a policy applies to the entire PKI, such policy constraints should be applied at the root level, when the CA certificate for a subordinate CA is created.

Policies that apply only to a portion of the PKI can be applied by establishing multiple subordinate policy CAs below the root CA.

Policy constraints are the most useful of the constraint options, both internally and externally. Name and path constraints might not provide you with sufficient protection in certain cross-certified relationships if the security standards of your business partners are not as strong as yours. For example, Organization A might have a policy that all user certificates must be approved manually by an administrator, while

Organization B approves certificate requests automatically as long as the user has an e-mail account. The security administrators for Organization A must ensure that the certificates that users in Organization B use to access resources meet the higher security standards that are necessary in Organization A. They can accomplish this by using policy constraints.

You can use name constraints to limit unplanned domain-based trusts. If your organization has more than one domain — such as contoso.com and subdomain.contoso.com — you might only want cross-certificates to map to CAs in the contoso.com domain and not to CAs in the subdomain.contoso.com domain. On the other hand, if you have five different domains and want your cross-certificates to apply to three of them, path constraints provide a more flexible solution.

Use path constraints if transitive trusts create problems for your organization — for example, if you have an environment in which users can freely install subordinate CAs and you do not have strong security guidelines governing CA creation and management.

Use policy mapping if the organization that you are cross-certifying with has policies that are similar to those of your organization. Policy mapping is less useful when the policies of your organization are stricter than the policies of the other organization, or vice versa. If this is the case, use policy constraints to restrict your trust relationship instead of using policy mapping.

**CA Certificate Processes and Interactions**

Two types of processes and interactions are associated with CA certificates:

- Processes and interactions that are common to end entity certificates as well as CA certificates.
- Processes and interactions that are unique to CA certificates.

All certificates are validated essentially the same way -- by a series of checks involving

- **Certificate discovery.** The process of collecting CA certificates from Group Policy, Enterprise Policy, and AIA pointers in issued certificates, and the process of enrolling certificates.
- **Path validation.** The process by which public key certificates and their issuer certificates are processed in a hierarchical fashion until the certificate chain terminates at a trusted, self-signed certificate. Typically, this is the root CA certificate.
- **Revocation checking.** Each certificate in the certificate chain is verified to ensure that none of the certificates are revoked. The revocation checking can take place either in conjunction with the chain building process or after the chain is built.

For more information about these processes, see "How Certificates Work" in the Technical Reference.

Where CA certificates differ significantly from end entity certificates is how they are created, and the relationships between them established, from the trusted root CA certificate to subordinate CA certificates and cross certificates. The following sections describe these processes:

- How root CA certificates are created.
- How root CA certificates are renewed.
- How cross certificates are created.
- How subordinate CA certificates are created.
- How policies are applied for qualified subordination.

**How Root CA Certificates are Created**

The creation of a root CA certificate is inseparable from the process of setting up Certificate Services on the first CA in a PKI hierarchy. Certificate Services provides components that process subsequent certificate requests and enable you to manage certificate use. The root certificate, on the other hand, provides the foundation and establishes the basic rules that government certificate issuance and use. Where the root certificate defines standards for what is acceptable and unacceptable in the PKI hierarchy, Certificate Services applies those standards throughout its processes.

An issued certificate typically inherits properties from a certificate template that is provided by the issuing CA. (Some examples are certificate lifetime, the distribution point of the CRL, and other parameters.) Since the root CA requires a certificate for itself, the root CA must self-sign the root CA certificate because there is no parent CA that could issue the CA certificate.

Before the CA certificate is generated, you must have completed any custom configuration of the CAPolicy.inf file that is relevant to the CA certificate. The CAPolicy.inf file should contain all configuration information that is required to generate the self signed CA certificate according to the organization's needs.

**Process for Creating Root CA Certificates**




To ensure that the CAPolicy.inf file is correctly processed:

- The ASCII text file must be available on the local computer before the CA setup procedure starts or before CA certificate renewal is attempted.
- The file must be placed in the *Systemroot* folder on the local computer on which the CA is installed.

For the specific syntax, see "CA Certificate Physical Structure," in this document.

**Note**

- There is no syntax or error-checking message mechanism if the CAPolicy.inf file is not in the correct format.

The public and private keys are then generated by the CSP. If you use the default CSP, the keys are written to the local computer's key store. If you did not use a hardware security module (HSM), the key is generated by CryptoAPI and is stored in the profile of the system account on the local computer. The length of time that is required to generate the key depends on both the size of the key that is generated and the CPU performance of the local computer. If an HSM is installed and selected, the key is generated in the HSM and stored according to the HSM specific architecture. Since no certificate templates are available on a stand-alone CA server that is a member of a workgroup, the CA certificate needs to be built from configuration information in the registry.

The following default key usage extension values are added to the CA certificate:

- Digital signature
- Certificate signing
- Certificate offline CRL signing
- Certificate CRL signing

However, if a root CA is installed on a computer that is running either a Windows 2000 or a Windows Server 2003 operating system, and that computer is a domain member, the CA inherits the Enhanced Key Usage extension settings from the CA template in Active Directory, even if the CA is installed as a stand-alone CA. If no Active Directory is available, the Enhanced Key Usage settings are also taken out of the configuration that is available in the registry.

**How Root CA Certificates are Renewed**

Root CA certificates need to have a longer lifetime than any other certificates in a PKI hierarchy. However, even root CA certificates will expire eventually. If they are not renewed, then every other certificate in the hierarchy is invalidated as well.

In addition, root CA certificates can be renewed before this is necessary in order to:

- Change the key used by the CA. (This is used in case of actual or suspected compromise.)
- Add certificate policies to the CA for qualified subordination.
- Change the CDP or Authority Information Access (AIA) paths.
- Partition a CRL.

Note that it is good practice to change the CA key at each renewal

Before the renewal process begins, the CAPolicy.inf file must be edited to reflect any changes, such as using a larger encryption key.

**Note**

- With root CA certificates, you cannot change the cryptographic service provider, reduce the key size, or shorten the validity period.

This updated information is then combined with the unchanged configuration information in the expiring root CA certificate to create the renewed root CA certificate, as illustrated in the following figure.

**Process for Renewing Root CA Certificates**




Both the expiring root CA certificate and the renewed root CA certificate continue to play critical roles in the PKI hierarchy:

- The original root CA certificate remains the ultimate foundation of trust for the hierarchy, and helps to validate the certificate chains for all certificates that have been issued under the original hierarchy.
- The renewed root CA certificate provides the foundation of trust for all certificates that are issued in the hierarchy from the renewal date forward.

To support both of these scenarios, by default, the root CA renewal process creates a pair of cross CA certificates that establish the trust relationship between the original and renewed root certificate:

- One verifies that the original root CA certificate trusts the renewed CA certificate
- The second verifies that the renewed CA certificate trusts the original root certificate.

The following figure illustrates this relationship.

**How Root CA Certificates Use Cross Certificates**




**Note**

- Although the certificates linking the original and renewed root certificate are referred to as cross certificates, they are not the same as the cross certificates used to establish trust between two separate PKI hierarchies.

Stand-alone CAs generate self -signed cross certificates when CA keys are changed. A cross certificate is generated for each key transition, for the period where the lifetimes of each root certificate overlap.

**How Subordinate CA Certificates Are Created**

The process for creating subordinate CA certificates in many ways resembles the process for creating a root CA certificate. Policy.inf must be in place on the server before you start Certificate Services setup, and the configuration data in Policy.inf is combined with configuration options that you select during Certificate Services setup.

At this point the two processes diverge, and the setup code for a subordinate CA generates a subordinate CA certificate request, which must be processed on the parent CA. After the parent CA has processed the request and generated the subordinate CA certificate, it can be transferred using the Certification Authority snap-in or Certutil.exe onto the subordinate CA. The following figure illustrates this process.

**Process for Creating Subordinate CA Certificates**




The process for obtaining a subordinate CA certificate from the root CA differs, depending on whether the parent CA is available online.

- If a parent CA is available online, the request can be sent directly to the root CA over the network.
- If a parent CA is not available online, the request must be saved to a file, which must be physically transferred to the parent CA to be processed.

At a minimum, the parent CA should provide a file that contains the subordinate CA's newly-issued certificate and, preferably, its full certification path. If a subordinate CA certificate does not include the full certification path, you must install the parent CA's certificate in the Intermediate Certification Authorities certificate store of the computer (if the parent CA is not a root CA), as well as the certificates of any other intermediate CA in the chain. You must also install the certificate of the root CA in the chain into the Trusted Root Certification Authorities store. These certificates should be installed in the certificate store before the CA certificate is installed on the new subordinate CA so that the new subordinate CA can build a valid CA chain when it starts.

**How Policies Are Applied for Qualified Subordination**

You can use CA certificates to implement qualified subordination, the process used to clearly define which certificates issued in one CA hierarchy are trusted in another CA hierarchy. Qualified subordination, which is also commonly called cross certification, can be used to link PKIs when two organizations merge or following an acquisition, or to facilitate closer business relationships between two organizations such as joint development agreements or consulting arrangements.

Qualified subordination is applied by issuing CA certificates from a CA in one PKI to a CA in a second PKI in order to establish a one-way trust relationship. If a CA in the second PKI issues a similar CA certificate to a CA in the first PKI, there will be a two-way trust relationship between the two hierarchies. Most likely, administrators in each hierarchy would configure these cross certificates using basic, policy, naming, and application constraints to limit which certificates are accepted from the other PKI, and the purposes for which they can be used.

**Note**

- In Windows 2000 operating systems, true cross-certification of CA hierarchies is not possible. Instead, you can define Certificate Trust Lists (CTLs) that trust specific CAs and use restricted certificates.

For example, consider the two PKIs illustrated in the following figure.

**How Cross Certificates Establish Trust Between PKI Hierarchies**




In this example, cross certificates are issued by Subordinate CA 2 to Subordinate CA A, which makes it possible for user certificates issued by Issuing CA A to be used in the PKI anchored by Root CA 1, and by Subordinate CA A to Subordinate CA 2, which makes it possible for certificates issued by Issuing CA 1 to be used in the PKI anchored by Root CA 2.

Cross CA certificates such as this can also be issued from one root CA to a second root CA, by a root to a subordinate CA, by a subordinate CA to a root, or by a third "bridge" root CA to a root or subordinate CA in each of the organizations that wish to have a cross relationship. However, in any of these cases, it is likely that various name and policy constraints will be used to define and limit the usage of certificates issued in the partner's CA hierarchy.

The key to understanding the use of cross certificates is that when a certificate issued by Issuing CA 1 is presented for use in the root CA 2 hierarchy, it is not validated against root CA 2. Instead, its certificate validation chain now includes the Sub CA A cross certificate, Subordinate CA 2, and Root CA 1. However, the cross certificate itself also must be validated, so the user certificate still ends up being validated for use --indirectly, in the case of the partner organization) for use.

To prevent undesired trust from being established, you can use:

- Basic constraints. These define which CAs your organization trusts in a partner's CA hierarchy by limiting the path length for a certificate chain.. For example, in the cross certificate issued by Subordinate CA 2, basic constraints could prevent the acceptance and use of certificates issued by any CAs beyond Subordinate CAB — or certificates issued by Subordinate CA B itself.

- Name constraints. These prevent acceptance of certificates from the other organization in undesired namespaces. For example, if Subordinate CA 1 serves certificate requests for users and computers in a domain for a separate division of the organization, it would be undesirable for certificates issued in Organization A for use in a relationship with the organization supported by Subordinate CA 2 to be used elsewhere in the organization.

- Application constraints. These prevent the use of the partner organization's certificates from being used for unapproved purposes. The cross certificate issued by Subordinate CA A might specify that certificates issued in the partner organization can only be used for e-mail signing, or for computer authentication, and specifically exclude other uses, such as key recovery.

- Policy constraints. These define the level of trust or assurance associated with the partner organization's certificates. In many organizations, certificates can have significant legal implications, such as for signing purchase orders up to, but not exceeding, certain amounts. For example, Sub CA Cross Certificate can include a policy constraint specifying that certificates issued in the partner organization can only used to sign purchase orders up to $1,000. In a case such as this, it would likely also be valuable to include a policy mapping to an Organization 2 policy defining policy limits, such as legal liability, associated with certificates issued in the Root CA 2 hierarchy.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "How Certificates Work" in Certificates Technical Reference

- "How Certificate Services Works" in Certificate Services Technical Reference

# CA Certificates Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- CA Certificate Tools
- CA Certificate Registry Entries
- CA Certificate Group Policy Settings


**CA Certificate Tools**
The following tools are associated with CA certificates.

**Certreq.exe: Certreq**

**Category**
Certreq ships with the Windows Server 2003 operating system tools and with the Windows Server 2003 Adminpak.

**Version Compatibility**
Certreq is compatible with Windows Server 2003 and Windows 2000 Server, and can be used to manage the certificate containers for users, computers, and services on computers running Windows 2000, Windows XP, and Windows Server 2003.Certreq enables you to submit, retrieve, create, and accept certificate requests that are sent to a Windows Server 2003 CA. You can also use Certreq to create and sign cross-certificate requests. You can also place the Certreq command syntax in a batch file to script certificate requests.

To find more information about Certreq, see "Command Line References" in Tools and Settings Collection.

**Certutil.exe: Certutil**

**Category**
Certutil is a command-line program that is installed as part of Certificate Services and with the Windows Server 2003 Adminpak.

**Version Compatibility**
Certutil can be used on Windows 2000 Server CAs and Windows Server 2003 CAs.

You can use Certutil to extract and display CA configuration information, configure Certificate Services, back up and restore CA components, and verify certificates, key pairs, and certificate chains.

To find more information about Certutil, see "Command Line References" in Tools and Settings Collection.

**Certsrv.msc: Certification Authority Snap-in**

**Category**
The Certification Authority snap-in ships with Windows 2000 Server and Windows Server 2003.

**Version Compatibility**
The Certification Authority snap-in is compatible with Windows Server 2003 and Windows 2000 Server.

The Certification Authority snap-in can be used to manage multiple CAs. It allows you to perform a variety of administrative tasks including:

- Starting and stopping the CA.
- Backing up and restoring the CA.
- Changing exit and policy modules.
- Viewing the CA certificate.

- Installing or reinstalling a CA certificate for the CA.
- Setting security permissions and delegating administrative control for the CA.
- Revoking certificates.
- Viewing or modifying CRL distribution points.
- Publishing CRLs and scheduling CRL publication.
- Configuring the types of certificates that are to be issued by the CA.
- Viewing information about certificates that have been issued.
- Viewing information about certificates that have been revoked.
- Viewing pending certificate requests.
- Approving or denying pending certificate requests.
- Viewing failed certificate requests.
- Renewing the CA's certificate.

**CA Certificate Registry Entries**

CA certificate-related registry entries correlate to the physical view of the certificate-related data that can be viewed by using the Certificates snap-in.

The registry settings in this section are a subset of the registry settings in "Certificate Services Tools and Settings." This subset makes it possible to monitor and manage key configuration options associated with CA certificates.

The following registry keys are associated with CA certificates that were distributed via Group Policy:

- **HKEY_CURRENT_USER\Software\Microsoft\SystemCertificates**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\SystemCertificates**

The following registry keys are associated with CA certificates that were not distributed via Group Policy:

- **HKEY_CURRENT_USER\Software\Policies\Microsoft\SystemCertificates**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\EnterpriseCertificates**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\SystemCertificates**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc**

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as MMC, to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_Current_User\Software\Microsoft\SystemCertificates**

**HKEY_Current_User\Software\Microsoft\SystemCertificates** contains registry settings for user certificates that have been distributed by means other than Group Policy. These keys also contain registry settings relating to the CA certificates that

support the user certificates.

The majority of the registry subkeys under SystemCertificates contain binary large objects that pertain to:

- **Certificates.** These entries identify the certificates associated with the registry entry.
- **CRLs.** These entries identify the certificate revocation lists (CRLs) associated with the registry entry.
- **CTLs.** These entries identify the certificate trust lists (CTLs) associated with the registry entry.

The following registry subkeys are located under SystemCertificates. Additional subkeys — which might appear under some registry subkeys — appear below under the registry subkey that they correspond to.

These subkeys appear under SystemCertificates:

**AuthRoot**

**Registry path**
**HKEY_Current_User\Software\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**AuthRoot** contains data about certificates, CRLs, and CTLs from non-Microsoft root CAs.

**CA**

**Registry path**
**HKEY_Current_User\Software\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**CA** contains data about certificates, CRLs, and CTLs from intermediate CAs.

**Root**

**Registry path**
**HKEY_Current_User\Software\ Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Root** contains data about trusted root CA certificates, CRLs, and CTLs.

**trust**

**Registry path**
**HKEY_Current_User\Software\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Trust** contains data about enterprise trust certificates, CRLs, and CTLs.

**HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates**
**HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates** contains registry settings for user certificates that have been distributed by using Group Policy. These keys also contain registry settings relating to the CA certificates that support the user certificates.

The majority of the registry entries contain binary large objects that pertain to:

- **Certificates.** These entries identify the certificates associated with the registry entry.
- **CRLs.** These entries identify the CRLs associated with the registry entry.

- **CTLs.** These entries identify the CTLs associated with the registry entry.

The following registry entries are located under SystemCertificates Additional subkeys — which might appear under some registry subkeys — will be detailed below under the registry subkeys that they correspond to.

**AuthRoot**

**Registry path**
**HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**AuthRoot** contains data about certificates, CRLs, and CTLs from non-Microsoft root CAs.

**CA**

**Registry path**
**HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**CA** contains data about certificates, CRLs, and CTLs from intermediate CAs.

**root**

**Registry path**
**HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Root** contains data about trusted root CA certificates, CRLs, and CTLs.

**trust**

**Registry path**
**HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Trust** contains data about enterprise trust certificates, CRLs, and CTLs

**HKEY_Local_Machine\Software\Microsoft\EnterpriseCertificates**
**HKEY_Local_Machine\Software\Microsoft\EnterpriseCertificates** contains registry settings for computer certificates that are downloaded from the Active Directory by the auto enrollment feature.

**AuthRoot**

**Registry path**
**HKEY_Current_User\Software\Microsoft\EnterpriseCertificates\AuthRoot**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**AuthRoot** contains data about certificates, CRLs, and CTLs from non-Microsoft root CAs.

**CA**

**Registry path**
**HKEY_Current_User\Software\Microsoft\EnterpriseCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**CA** contains data about certificates, CRLs, and CTLs from intermediate CAs.

**NTAuth**

**Registry Path**
**HKEY_Local_Machine\Software\Microsoft\EnterpriseCertificates\NTAuth**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Any CA that issues smart card logon or domain controller certificates must publish its CA certificate into the **NTAuth** store in Active Directory. Windows 2000 and Windows 2003 CAs automatically publish their CA certificates in Active Directory. Certificates published by third parties must be imported into the **NTAuth** store. After you put the non-Microsoft CA in the **NTAuth** store, domain-based Group Policy adds this registry key to all computers in the domain.

For more information about the **NTAuth** store, see "How Certificate Services Works" in the Windows Server 2003 Technical Reference.

**Root**

**Registry path**
**HKEY_Current_User\Software\ Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Root** contains data about trusted root CA certificates, CRLs, and CTLs.

**trust**

**Registry path**
**HKEY_Current_User\Software\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Trust** contains data about enterprise trust certificates, CRLs, and CTLs.

**HKEY_Local_Machine\Software\Microsoft\SystemCertificates**
**HKEY_Local_Machine\Software\Microsoft\SystemCertificates** contains registry settings for computer certificates that have been distributed by means other than Group Policy. These keys also contain registry settings relating to the CA certificates that support the computer certificates.

**AuthRoot**

**Registry path**
**HKEY_Local_Machine\Software\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**AuthRoot** contains data about certificates, CRLs, and CTLs from non-Microsoft root CAs.

**CA**

**Registry path**
**HKEY_Local_Machine\Software\Microsoft\SystemCertificates\CA**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**CA** contains data about certificates, CRLs, and CTLs from intermediate CAs.

**ROOT**

**Registry path**
**HKEY_Local_Machine\Software\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**ROOT** contains data about trusted root CA certificates, CRLs, and CTLs.

**trust**

**Registry path**
**HKEY_Local_Machine\Software\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Trust** contains data about enterprise trust certificates, CRLs, and CTLs.

**HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates**
**HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates** contains registry settings for computer certificates that have been distributed using Group Policy. These keys also contain registry settings relating to the CA certificates that support the computer certificates.

The majority of registry entries contain binary large objects that pertain to:

- **Certificates.** These entries identify the certificates associated with the registry entry.
- **CRLs.** These entries identify the CRLs associated with the registry entry.
- **CTLs.** These entries identify the CTLs associated with the registry entry.

The following registry entries are located under **SystemCertificates**. Additional subkeys — which might appear under some registry subkeys — will be detailed below under the registry subkeys that they correspond to.

**AuthRoot**

**Registry path**
**HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**AuthRoot** contains data about certificates, CRLs, and CTLs from non-Microsoft root CAs.

**ca**

**Registry path**
**HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Ca** contains data about certificates, CRLs, and CTLs from intermediate CAs.

**Certificates**

**Registry path**
**HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\SystemCertificates\Root\Certificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Certificates** contains a unique key for each root CA certificate that is based on the hash of the certificate.

**ProtectedRoots**

**Registry path**
**HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\SystemCertificates\Root\ProtectedRoots**

**Version**

Windows Server 2003 and Windows XP Professional Service Pack 2

**ProtectedRoots** can be used to:

1. Disable user root store trust. This prevents CryptoAPI applications from using the user root store in building trusted certificate chains.

2. Disable user root store changes. This prevents users from adding root CAs to the user trusted root store. This value can also be set through Group Policy in Windows Server 2003.

3. Disable user root store purge. This prevents the user from removing root CAs from the user trusted store that are also in the local computer trusted root store.

4. Disable the requirement for NTAuth policy processing. This disables the requirement for an issuing CA to be present in the NTAuth store of the local computer. This value can be set using Group Policy.

5. Disable name constraint enforcement for undefined name types. By default, Windows XP Professional with Service Pack 2 and Windows Server 2003 reject undefined name types in a name constraint validation. Setting this value forces the acceptance of all name forms that are not explicitly defined.

### Note

- The option to disable name constraint enforcement is not configurable using Group Policy.

**Root**

**Registry path**
**HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Root** contains data about trusted root CA certificates, CRLs, and CTLs.

**Trust**

**Registry path**
**HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates**

**Version**
Windows Server 2003, Windows XP, and Windows 2000

**Trust** contains data about enterprise trust certificates, CRLs, and CTLs.

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc**
The registry entries in this hive correlate to CA certificate configuration information that can be configured and viewed using the Certificates and Certification Authority snap-ins, as well as the Certutil command-line tool.

**CertSvc\Configuration\CAName**
The registry settings in this hive contain information relating to the configuration of the certification authority (CA) installed on the server.

**CACertFileName**

**Registry Path**
**\CertSvc\Configuration** *CAName*\ **CACertFileName**

**Version**
Windows Server 2003 and Windows 2000 Server

On subordinate CAs, this setting identifies the name and location of the root certificate.

**CACertHash**

**Registry Path**
**\CertSvc\Configuration** *CAName*\ **CACertHash**

**Version**
Windows Server 2003 and Windows 2000 Server

On subordinate CAs, this setting identifies the hash algorithms — SSH-1, for example — that are used for all CA certificates.

**CACertPublicationURLs**

**Registry Path**
**\CertSvc\Configuration** *CAName*\ **CACertPublicationURLs**

**Version**
Windows Server 2003 only

This setting identifies the URL of the authority information access point where a client can find a CA certificate. Because the authority information access point is the location of the certificate that was used to sign the certificate, and because a root CA issues the CA certificate to itself, you should leave this setting blank for the root CA.

**CAKeyUsage**

**Registry Path**
**\CertSvc\Configuration** *CAName*\ **CAKeyUsage**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting makes it possible to disable the digital signature key usage on stand-alone CAs. By default, a stand-alone CA certificate will contain digital signature, certificate signing, and CRL signing key usage values.

### Note

- Enterprise CAs enforce the key usage based on the subordinate CA template settings.

**CAServerName**

**Registry Path**
**\CertSvc\Configuration** *CAName*\ **CAServerName**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting records the DNS name of the server on which the CA is running.

**CAType**

**Registry Path**
**\CertSvc\Configuration** *CAName*\ **CAType**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting records whether the CA is a root, enterprise, or stand-alone CA.

**CommonName**

**Registry Path**
\ **CertSvc\Configuration** \ *CAName*\ **CommonName**

**Version**
Windows Server 2003 and Windows 2000 Server

The common name becomes part of the certificate issuer name and is also part of the CRL and AIA if replacement tokens are used. The common name must not exceed 64 characters in length.

### Note

- Each space in the name will actually use three characters in the total length because of how escape characters are written (%20).

#### CRLFlags

**Registry Path**
\ **CertSvc\Configuration** \ *CAName*\ **CRLFlags**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting allows you to manage cross CA certificate creation. By default, cross CA certificates are generated using predefined extensions in the registry. If desired, this setting can be used to generate cross certificates using the Cross CA certificate template instead. This setting can also be used to disable automatic cross CA certificate generation. .

#### CRLPublicationURLs

**Registry Path**
\ **CertSvc\Configuration** \ *CAName*\ **CRLPublicationURLs**

**Version**
Windows Server 2003 and Windows 2000 Server

Configuring this setting ensures that the correct CRL and AIA information is embedded in each issued certificate so that the certificate's signature and revocation status can be verified.

For all CA types (online or offline, root or subordinate, enterprise or stand-alone), the configuration of the AIA extension and the CRL distribution point extension is critical. If they are not configured correctly or if they contain invalid extensions, there may be unexpected problems. For example, smart card login attempts may not work, there may be invalid e-mail digital signatures, or there may be Web sites that are not trusted.

### Note

- Changes to CRL distribution points and AIA extensions take effect only after the CA is restarted, and the extensions appear in certificates issued only after the changes are applied. The CRL and AIA configuration processes are different on computers that are running Windows Server 2003 and Windows 2000 Server.

#### DSConfigDN

**Registry Path**
\ **CertSvc\Configuration** \ *CAName*\ **DSConfigDN**

**Version**
Windows Server 2003 and Windows 2000 Server

The distinguished name in DSConfigDNidentifies the namespace of the domain that the CA belongs to. When a root  CA is configured as a stand-alone CA, the distinguished name should be the same as the namespace that will be used for the enterprise CA. The distinguished name becomes part of the certificate issuer name and is also part of the CRL and AIA if replacement tokens are used. It is also used by several variables that are used to set the CRL and AIA.

### Note

- After you use this command to change the registry key, a new CRL and any new CA certificates that are issued must be republished. Only new certificates that are issued after you use the previous command will have this information available. It is important to note that you must reissue and republish any subordinate CA certificate if it was issued before you changed the registry key.

#### EnforceX500NameLengths

**Registry Path**
**\CertSvc\Configuration** *CAName*\ **EnforceX500NameLengths**

**Version**

Windows Server 2003 and Windows 2000 Server

This setting makes it possible for certificates issued by the CA to bypass X.500 naming restrictions. Windows Server 2003 CAs enforce X.500 naming standards by default. It is possible that deep organizational unit (OU) paths might exceed normal length restrictions. If your organization's PKI does not need to be compatible with a non-Microsoft PKI, you can configure this registry setting to bypass the X.500 name length restriction.

**HighSerial**

**Registry Path**
**\CertSvc\Configuration** *CAName*\ **HighSerial**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting can be used to enable alternate forms of fixed length serial number generation. In Windows 2000 CAs, two types of fixed-length serial numbers are generated. In a Windows Server 2003 CA, three types of fixed-length serial numbers are generated. The default and alternate forms are the same as in Windows 2000. The alternate forms can be used to work around serial number encoding issues in certain non-Microsoft PKI applications.

### Note

- Any length hexadecimal string can be set in the registry (but there must be an even number of digits). The number of bytes used from the registry will be reduced if it would overflow a total of 19 bytes for the serial number. The high byte is manipulated as described previously to avoid problems with certain non-Microsoft applications. The IETF standards specify a maximum of 20 byte serial numbers.

**SubjectTemplate**

**Registry Path**
**\CertSvc\Configuration** *CAName*\ **SubjectTemplate**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting contains an ordered list of the subject relative distinguished name elements that are allowed in the **Subject** field of certificates issued by the CA.

This setting can only be set to a small, fixed list of relative distinguished name elements supported by the CA. If, during request processing, a listed relative distinguished name field is empty, or if the field is not populated by the request **Subject** field or by the policy module, the element will not be included. If the registry value is completely empty, the binary subject encoding from the request is passed through to the issued certificate unmodified.

**CertSvc\Configuration\CAName\CSP**

**Provider**

**Registry Path**
**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration** \ *CAName*\ **CSP\Provider**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting lists the CSP that is used to generate the CA certificate key material.

**Hash Algorithm**

**Registry Path**
**HKLM\System\CurrentControlSet\Services\CertSvc\Configuration\\CAName** *CSP*\ **HashAlgorithm**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting identifies the hash algorithm that is used for hashing and signing CA certificate contents.

**CertSvc\Configuration\CAName\PolicyModules**

**Active**

**Registry Path**
**CertSvc\Configuration** *CAName*\ **PolicyModules\Active**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting lists the active policy modules.

**CertSvc\Configuration\CAName\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy**

**DisableExtensionList**

**Registry Path**
**CertSvc\Configuration** *CAName*\
**PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\DisableExtensionList**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting disables the adding of a certificate extension to a certificate that is included by default in certificates issued by an enterprise CA. An example of such a certificate extension could be the S/MIME capabilities extension.

**EditFlags**

**Registry Path**
**CertSvc\Configuration** *CAName*\ **PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\EditFlags**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting enables you to identify the critical and non-critical extensions associated with policy modules, such as enabling basic constraints. The flags that can be configured include:

- **EDITF_ATTRIBUTEENDDATE.** Controls certificate validity time by certificate request. This feature is disabled by default for enterprise CAs, and enabled by default for stand-alone CAs.

  **Note**

  - Although the validity period of a certificate is normally set by the CA, the CA can be configured in a way that allows the request to specify the validity period. This CA configuration option can only reduce the validity period of certificates to be issued by a request.

- **EDITF_ENABLEAKIISSUERNAME.** Used to enable or disable the issuer name and issuer serial number.

  **Note**

  - You should disable the issuer name and issuer serial number if you renew CA keys when you renew CA certificates or plan to use cross-certificates.

**EnableEnrolleeRequestExtensionList**

**Registry Path**
**CertSvc\Configuration** *CAName*\
**PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\EnableEnrolleeRequestExtensionList**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting lists the object identifiers (also known as OIDs) of the rules extensions that are defined by the organization that must appear in the issued certificate. These settings apply to all requests submitted to a stand-alone CA, and to requests submitted to an enterprise CA that use a template which specifies that the request must supply the subject information.

**EnableRequestExtensionList**

**Registry Path**
**CertSvc\Configuration** *CAName*\ **PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\EnableRequestExtensionList**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting lists the object identifiers of the rules extensions defined by the organization that must appear in the issued certificate. By default, extensions that are in the submitted request are added to the CA's database, but are marked as disabled, so they will not appear in the issued certificate.

**RequestDisposition**

**Registry Path**
**CertSvc\Configuration** *CAName*\ **PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\RequestDisposition**

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to specify what the policy module should do with incoming and resubmitted requests. Options include:

- Process all certificate requests that meet applicable policy and template processing requirements.
- Place all initial certificate requests in a pending queue, and process all certificate requests that are resubmitted by an administrator, assuming the requests meet applicable policy and template processing requirements.

**RevocationType**

**Registry Path**
**CertSvc\Configuration** *CAName*\ **PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\RevocationType**

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to specify an alternative certificate revocation service to function with a Windows Server 2003 CA. This setting enables CRL URLs from the following: Lightweight Directory Access Protocol (LDAP), File Transfer Protocol (FTP), and file locations; CDP extensions; and Active Server Page (ASP)-based revocation.

**RevocationURL**

**Registry Path**
**CertSvc\Configuration** *CAName*\ **PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\RevocationURL**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the URL for an alternative certificate revocation service that has been configured to function with a Windows Server 2003 CA.

The replacement variables that can be used in the revocation URL are listed in the following table.

**Registry Variables for Alternative Revocation Services**

| Alternative Revocation Service | Registry Variable |
|---|---|
| SERVERDNSNAME | 1 |
| SERVERSHORTNAME | 2 |
| SANITIZEDCANAME | 3 |
| CERTFILENAMESUFFIX | 4 |

| | |
|---|---|
| DOMAINDN | 5 |
| CONFIGDN | 6 |
| SANITIZEDCANAMEHASH | 7 |
| CRLFILENAMESUFFIX | 8 |
| CRLDELTAFILENAMESUFFIX | 9 |
| DSCRLATTRIBUTE | 10 |
| DSCACERTATTRIBUTE | 11 |
| DSUSERCERTATTRIBUTE | 12 |
| DSKRACERTATTRIBUTE | 13 |
| DSCROSSCERTPAIRATTRIBUTE | 14 |

In order for the revocation service to function, the application, service, or account connecting to this URL must have Read permissions in the Certification Authority snap-in. If IIS is using a local account, follow the steps for enabling anonymous access in IIS and allowing Anonymous Read access to the CA.

### Note

- Enabling Anonymous Read access to the CA might expose privacy or security risks.

#### SubjectAltName

**Registry Path**
**CertSvc\Configuration** *CAName*\ **PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\SubjectAltName**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting uses an organizational unit for the SubjAltName extension of an issued certificate. This setting is almost never used.

#### SubjectAltName2

**Registry Path**
**CertSvc\Configuration** *CAName*\ **PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\SubjectAltName2**

**Version**
Windows Server 2003 and Windows 2000 Server

This setting makes it possible for a stand-alone CA to place the e-mail address of the authenticated user making the certificate request in the SubjAltName extension of an issued certificate. This setting is rarely used.

#### CA Certificate Group Policy Settings

The following public key Group Policy settings are associated with CA certificates:

- **Trusted Root Certification Authorities** for adding trusted root CA certificates to the Trusted Root Certification Authorities store.

- **Enterprise Trust** for configuring certificate trust lists.

#### Trusted Root Certification Authorities

When you install an enterprise root CA or a stand-alone root CA, the certificate of the CA is added automatically to the Trusted Root Certification Authorities Group Policy object for the domain. You also can add certificates for other root CAs to the Trusted Root Certification Authorities Group Policy object.

Trusted Root Certification Authorities Group Policy can be configured in the following location in the Group Policy MMC snap-

in:

Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Trusted Root Certification Authorities

The root CA certificates that you add become trusted root CAs for computers within the scope of the Group Policy object. For example, if you want to use a non-Microsoft CA as a root CA in a certification hierarchy, you must add the certificate for the non-Microsoft CA to the Trusted Root Certification Authorities Group Policy object.

To add a certificate for the root CA to the Trusted Root Certification Authorities Group Policy, in the **Public Key Policies** node, right-click **Trusted Root Certification Authorities**, click **All Tasks**, and then click **Import**. When the Certificate Import Wizard appears, use the wizard to import a certificate file for the certificate of the root CA and add it to Group Policy. The certificate is added to the **Trusted Root Certification Authorities** store of all computers within the scope of Group Policy the next time it is refreshed on each computer.

You can also use Trusted Root Certification Authorities Group Policy to control the changes users can make to trusted root CA options.

To modify these settings, right-click **Trusted Root Certification Authorities**, and then click **Properties**.

To prevent users making any changes to the trusted root CA store, you can clear the **Allow users to select new root Certification Authorities (CAs) to trust** check box.

To limit client trust of alternative certificate stores, under **Client computers can trust the following certificate stores**, select one of the following options:

- **Third-Party Root Certification Authorities and Enterprise Root Certification Authorities**
- **Enterprise Root Certification Authorities**

In addition, the following settings under **To perform certificate-based authentication of users and computers, CAs must meet the following criteria:** can be used to manage certificate-based authentication by certificate holders:

- **Registered in Active Directory only**
- **Registered in Active Directory and compliant with name constraints requirements for user principal names (UPNs)**

**Note**

- For more information about importing non-Microsoft root CA certificates, see the information for **HKEY_Local_Machine\Software\Microsoft\EnterpriseCertificates\NTAuth**, earlier in this document.

#### Certificate Trust Lists

You can create certificate trust lists (CTLs) to trust specific CAs and to restrict the uses of certificates issued by the CAs. For example, you might use a CTL to trust certificates that are issued by a commercial CA and restrict the permitted uses for those certificates. You might also use CTLs to control trust on an extranet for certificates that are issued by CAs that are managed by your business partners. You can configure CTLs for computers and for users.

Before administrators can create CTLs, they must have a valid trust list signing certificate — such as the Administrator certificate or the Trust List Signing certificate that is issued by enterprise CAs. The trust list signing private key for the administrator is used to sign the CTL for integrity. If the trust list signing certificate for an administrator is invalid, all CTLs that have been created and signed by that administrator also are invalid.

Certificate Trust List Group Policy is configured as part of Public Key Policies\Enterprise Trust in the Security Settings of the Group Policy object for a user or computer in a domain, site, or OU. Certificate Trust List Group Policy can be configured in the following location in the Group Policy MMC snap-in:

Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Enterprise Trust

You use the Certificate Trust Wizard to configure Certificate Trust List Group policy.

To activate the wizard, right-click **Enterprise Trust**, click **New** and then click **Certificate Trust List**. The wizard enables you to configure the following options:

- **Valid duration.** An optional lifetime for the CTL. If you do not specify a lifetime, the CTL expires when the trust list signing certificate expires.

- **Designate Purposes.** Enables you to restrict the purposes for which certificates are trusted. The CTL establishes trust only for certificates that are valid for the selected purposes. A certificate might support all of the listed purposes.

- **Add Purpose.** Enables you to add purposes to the **Designate Purposes** box. This also requires you to enter an object identifer for the new purpose.

- **Current CTL Certificates.** Displays the certificates of the root CAs that are to be trusted by this CTL. Certificates with certification paths to the root CA are trusted for all designated purposes specified by the CTL.

- **Add from Store.** Adds a root certificate from the Trusted Root Certification Authorities store.

- **Add from File.** Adds a root CA's certificate from a file.

- **Remove.** Deletes the certificate that is selected in the **Current CTL Certificates** box.

- **View Certificate.** Enables you to view the certificates that are selected in the **Current CTL Certificates** box.

- **Use this certificate.** Displays the trust list signing certificate for the private key that is to be used to sign the CTL.

- **Select from Store.** Adds a trust list signing certificate from the Personal store for the administrator.

- **Select from File.** Adds the trust list signing certificate from a file.

- **View Certificate.** Enables you to view the certificate listed in the **Use this certificate** box.

- **Add a timestamp to the data.** Adds a timestamp to the CTL. The timestamp is used to determine the valid lifetime of the CTL. If a timestamp is not used, the computer clock is used instead.

- **Timestamp service URL.** Identifies the location of the timestamp service that is to be used for the timestamp.

- **Friendly Name.** The optional name that appears in the Microsoft Management Console when the CTL is displayed.

- **Description.** An optional description to describe the CTL.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Certificates Tools and Settings" in Certificates Technical Reference

- "Certificate Services Tools and Settings" in Certificate Services Technical Reference

# Certificates Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Digital certificates are electronic credentials that are used to certify the online identities of individuals, computers, and other entities on a network. Digital certificates function similarly to identification cards such as passports and drivers' licenses. They are issued by certification authorities (CAs) that must validate the identity of the certificate holder, both before the certificate is issued and when the certificate is used. Common uses include business scenarios requiring authentication, encryption, and digital signing.

**In this subject**


- What Are Certificates?
- How Certificates Work
- Certificates Tools and Settings

# What Are Certificates?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**What Are Certificates?**

**In this section**


- Uses of Certificates
- Certificate Dependencies
- Related Information


Digital certificates are electronic credentials that are used to certify the identities of individuals, computers, and other entities on a network.

Digital certificates function similarly to identification cards such as passports and drivers' licenses. For example, passports and drivers' licenses are issued by recognized government authorities, whereas digital certificates are issued by recognized certification authorities (CAs).

When someone requests a passport or driver's license, the government authority verifies the identity of the requester, certifies that the requester meets all requirements to receive the card, and then issues the card. Before a certificate can be issued, a CA or CA administrator must verify the requester's identity, determine that they meet all requirements to receive the certificate, and then issue the certificate.

Like an identification card such as a driver's license or passport, a digital certificate can be used to verify the identity of its owner. When the certificate is presented to others, they help verify the identity of its owner based on the quality of the contents of the certificate:

- Personal information that helps identify the owner.
- The signature of the issuing authority. For digital certificates, the issuing authority is the CA.
- Information needed to identify and contact the issuing authority.

In addition, the quality of a certificate is enhanced if it:

- Is designed to be tamper-resistant and difficult to counterfeit.
- Is issued by an authority that can revoke the certificate at any time (for example, if the employee to whom the certificate was issued is no longer employed by the organization).
- Can be checked for revocation by contacting the issuing authority.

**Uses of Certificates**

Private and public networks are being used with increasing frequency to communicate sensitive data and complete critical transactions. This has created a need for greater confidence in the identity of the person, computer, or service on the other end of the communication. In addition, these valuable communications must be protected while they are on the network. Although accounts and passwords provide a certain level of assurance in the identity of the entity on the other end of the network, they offer little or no protection while data is in transit. In comparison, digital certificates and public key encryption provide an enhanced level of authentication and privacy to digital communications.

Certificates and certificate services from numerous vendors are being used to strengthen the security of a variety of applications and user scenarios. Windows Server 2003 Certificate Services makes possible strong security based on public key encryption that can enhance a variety of internal and external applications for your organization, including:

- Authentication.

- Encryption.
- Data integrity.

**Authentication**

Authentication is crucial to secure and reliable communication. Each party to a communication must be able to prove their own identity to those with whom they communicate, and in turn must be able to verify the identity of the parties at the other end of the communication.

This process can be challenging when both parties are in the same location. Authentication of identity on a network can be even more difficult because the communicating parties do not physically meet as they communicate. This makes it potentially easier for an unethical person to intercept messages that are meant to be private or to pretend that they are another person or entity.

Digital certificates and public key encryption provide an enhanced means of verifying identity, which makes it difficult for an entity to impersonate another entity. Digital certificates help verify identity because the data in a certificate includes the public cryptographic key from the certificate subject's public and private key pair. A message signed with its sender's private key can be verified by the message's recipient as authentic by using the sender's public key, which can be found on a copy of the sender's certificate. Verifying a signature by using a public key from a certificate proves that the signature was produced using the certificate subject's private key. If the sender has been vigilant and has kept the private key secret, the receiver can be confident in the identity of the message sender.

A few of the ways certificates are used to provide authentication are:

- Authentication of a user to a secure Web site via the Transport Layer Security (TLS) or the Secure Sockets Layer (SSL) protocol.
- Authentication of a server to a user via TLS.
- Logging on to a Windows Server 2003 domain.
- Authentication of a client on a wireless network.
- Authentication of a client across the Internet to create a virtual private network (VPN).
- Internet Protocol security (IPSec).

**Encryption**

Communications over a network, such as the Internet, are subject to possible monitoring by unknown and, perhaps, malicious users. Public networks are treacherous for unencrypted sensitive information because anyone can access the network and analyze the data being transmitted between two points. Even private local area networks (LANs) are vulnerable to determined efforts by intruders to acquire physical access to the network. Consequently, if sensitive information is transmitted between computing devices on any type of network, users will almost certainly want to use some sort of encryption to keep their data private.

Encryption is the process of disguising a message or data in such a way as to hide its substance. It can be thought of as locking something valuable into a strongbox with a key. Conversely, decryption can be compared to opening the box and retrieving the valuable item. On computers, sensitive data in the form of e-mail messages, files on a disk, and files being transmitted across the network can be encrypted by using a key. Encrypted data and the key used to encrypt data are both unintelligible.

Public key encryption is not used to encrypt large amounts of data; instead, data is typically protected with a private key and that private key in turn is encrypted with the public key of the recipient of the data. The encrypted secret key will then be transmitted to the recipient along with the encrypted data. The recipient will use the private key to decrypt the secret key. The secret key will then be used to decrypt the message itself.

Certificates enable privacy for data that is transmitted using a number of different methods. Some of the commonly used privacy-enabling protocols that use certificates are:

- Secure Multipurpose Internet Mail Extensions (S/MIME).

- TLS.
- Encrypting File System (EFS)

**Data Integrity**
An increasing number of digital documents require strong evidence that the data has not been altered since it was signed and confirmation of the identity of the person or entity who signed the data. A digital signature helps ensure the integrity and origin of data, which are essential for secure e-commerce transactions.

Digital signatures are typically used when data is distributed in plaintext, or unencrypted, form. In these cases, although the sensitivity of the message itself might not warrant encryption, there could be a compelling reason to ensure that the data is in its original form and has not been sent by an impostor. In a distributed computing environment, plaintext can conceivably be read or altered by anyone on the network who has access to it, whether authorized or not.

**Certificate Dependencies**
To use certificates in Windows Server 2003, you need to set up a public key infrastructure (PKI), which consists of one or more CAs that are linked in a hierarchy. These CAs and the PKI are required to manage certificate issuance, validation, renewal, and revocation in one or more organizations.

You can use a third-party PKI with Windows Server 2003, or you can establish your own PKI based on Windows Server 2003 Certificate Services.

If you are using Active Directory to manage user and computer accounts on the network, your enterprise CAs can use information contained in Active Directory — such as user account names, security group memberships, and certificate templates — to simplify certificate enrollment, renewal, revocation, and validation.

In addition, Active Directory stores public key Group Policy for distribution to all computers that are within the scope of a Group Policy object. Public key Group Policy enables you to control which CAs are to be trusted in the enterprise, specify alternative EFS recovery agents, and configure automatic enrollment and renewal of certificates for computers running Windows 2000, Windows XP, and Windows Server 2003 — all from a central administration point.

Active Directory also supports mapping certificates to network user accounts for authenticating clients and controlling access to network resources. Using smart cards for the user logon process is a special case of certificate mapping that extends the Kerberos v5 authentication protocol to include authentication of users on the basis of certificates and private keys that are stored on smart cards. Using smart cards for the user logon process provides enhanced security for user authentication and a single set of user credentials for logging on locally or remotely over a network.

The following figure illustrates the role of domain client certificates in an Active Directory environment.

**Domain Client Certificates in Active Directory**




**Related Information**
The following resources contain additional information that is relevant to this section:

- IPSec Technical Reference
- Logon and Authentication Technologies

# How Certificates Work

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Certificates Work**

**In this section**


- Certificate Architecture
- Certificate Processes and Interactions
- Related Information


Digital certificates are electronic credentials that are used to assert the online identities of individuals, computers, and other entities on a network. Digital certificates function similarly to identification cards such as passports and drivers licenses. They are issued by certification authorities (CAs) that must validate the identity of the certificate-holder both before the certificate is issued and when the certificate is used. Common uses include business scenarios requiring authentication, encryption, and digital signing.

The following sections provide an in-depth view of how certificates work in an optimal environment. An optimal environment for certificates includes the following:

- A public key infrastructure (PKI) that has been designed, planned, and deployed with appropriate security measures to minimize the risk that CAs can be compromised.
- Certificate configurations have been planned so that only intended users can initiate certificate use, and certificates can only be used for the purposes intended.

**Certificate Architecture**

A digital certificate binds a user, computer, or service's identity to a public key by providing information about the subject of the certificate, the validity of the certificate, and applications and services that can use the certificate. Certificates issued in Windows Server 2003 and earlier PKIs are structured to meet these objectives based on standards established by the Public-Key Infrastructure (X.509) Working Group (PKIX) of the Internet Engineering Tasks Force (IETF). The following figure shows the contents of X.509 version 3 certificates as implemented in Windows Server 2003.

**X.509 Version 3 Certificate**




X.509 version 3 certificates are the current certificate format in a Windows Server 2003 PKI. Version 3 certificates support the following fields that have been supported since X.509 version 1:

- **Subject.** Provides the name of the computer, user, network device, or service that the CA issues the certificate to. The

subject name is commonly represented by using an X.500 or Lightweight Directory Access Protocol (LDAP) format.

- **Serial Number.** Provides a unique identifier for each certificate that a CA issues.
- **Issuer.** Provides a distinguished name for the CA that issued the certificate. The issuer name is commonly represented by using an X.500 or LDAP format.
- **Valid From.** Provides the date and time when the certificate becomes valid.
- **Valid To.** Provides the date and time when the certificate is no longer considered valid.

  **Note**

  - The date when an application or service evaluates the certificate must fall between the Valid From and Valid To fields of the certificate for the certificate to be considered valid.

- **Public Key.** Contains the public key of the key pair that is associated with the certificate.

In addition to the version 1 fields, X.509 version 3 certificates include extensions that provide additional functionality and features to the certificate. These extensions are optional and are not necessarily included in each certificate that the CA issues:

- **Subject alternative name.** A subject can be presented in many different formats. For example, if the certificate must include a user's account name in the format of an LDAP distinguished name, e-mail name, and a user principal name (UPN), you can include the e-mail name or UPN in a certificate by adding a subject alternative name extension that includes these additional name formats.
- **CRL distribution points (CDP).** When a user, service, or computer presents a certificate, an application or service must determine whether the certificate has been revoked before its validity period has expired. The CDP extension provides one or more URLs where the application or service can retrieve the certificate revocation list (CRL) from.
- **Authority Information Access (AIA).** After an application or service validates a certificate, the certificate of the CA that issued the certificate — also referred to as the parent CA — must also be evaluated for revocation and validity. The AIA extension provides one or more URLs from where an application or service can retrieve the issuing CA certificate.
- **Enhanced Key Usage (EKU).** This attribute includes an object identifier (OID) for each application or service a certificate can be used for. Each OID is a unique sequence of numbers from a worldwide registry.
- **Certificate policies.** Describes what measures an organization takes to validate the identity of a certificate requestor before it issues a certificate. An OID is used to represent the validation process and can include a policy-qualified URL that fully describes the measures taken to validate the identity.

**Certificate Templates**

Windows 2000 and Windows Server 2003 Enterprise CAs use certificate templates, stored in the Active Directory directory service, to provide the default attributes for a certificate. These attributes include authorized uses for the certificate, the cryptographic algorithms used with the certificate, the format of the subject, the public key length, issuance requirements, and the certificate lifetime.

Certificate templates are configured on an enterprise CA and are applied to accept or reject incoming certificate requests. Only enterprise CAs can issue certificates based on certificate templates. Certificate template information is stored in Active Directory in the container

`CN=Certificate Templates,CN=Public Key Services,CN=Services,CN=Configuration, DC=` *ForestRootNameDN*

where *ForestRootNameDN* is the LDAP distinguished name of the forest root domain.

Computers running Windows Server 2003 support two types of certificate templates: version 1 and version 2. Servers running Windows 2000 only support the issuance of certificates that are based on version 1 certificate templates.

When the first enterprise CA is installed in the forest, version 1 templates are created by default. Unlike version 2 templates,

these cannot be modified or removed. In Windows 2000, certificate management was very limited because only the templates' security permissions could be set.

Version 1 templates are provided for backward compatibility, and support many general needs for subject certification. For example, there are certificates that allow Encrypting File System (EFS) encryption, client authentication, smart card logon, or server authentication.

Version 2 certificate templates allow you to define specific attributes for certificates that meet your organization's business needs. Version 2 template definitions are stored in Active Directory, although you can create and modify version 2 templates at any computer running Windows Server 2003 or Windows XP Professional with the Windows Server 2003 Administration pack installed. Certificates based on version 2 templates can only be issued by a CA running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.

**Types of Certificate Templates**

Only a few certificate templates are assigned to a fresh CA installation by default. Several predefined certificate templates and custom certificate templates must be assigned explicitly if needed.

There are two categories of certificate templates: certificate templates issued to users and certificate templates issued to computers. Only computers can use certificates that are issued to computers; likewise, only users can use certificates that are issued to users. Another way to distinguish between certificate templates is based on how they are used:

- **Single function.** A certificate template can be highly restricted and only be used for a single function. For example, you can use a Basic EFS certificate template only to encrypt and decrypt files that are protected by using EFS.

- **Multiple functions.** You can use a certificate template for multiple functions. For example, you can use a user certificate template to encrypt and decrypt files, authenticate with a server, and send and receive secure e-mail by using the same certificate.

After the upgrade to Windows Server 2003 certificate templates is completed, the following preconfigured certificate templates are listed in the Certificate Templates Microsoft Management Console (MMC).

**Windows Server 2003 Certificate Templates**

| Name | Description | Key Usage | Subject Type | Published to Active Directory |
|---|---|---|---|---|
| Administrator | Allows trust list signing and user authentication | Signature and encryption | User | Yes |
| Authenticated Session | Subject can authenticate to a Web server | Signature | User | No |
| Basic EFS | Used by EFS to encrypt data | Encryption | User | Yes |
| CA Exchange | Used to store keys that are configured for private key archival | Encryption | Computer | No |
| CEP Encryption | Allows the holder to act as a registration authority (RA) for simple certificate enrollment protocol (SCEP) requests | Encryption | Computer | No |
| Code Signing | Used to digitally sign software | Signature | User | No |
| Computer | Allows a computer to authenticate itself on the network | Signature and encryption | Computer | No |
| Cross-Certification Authority | Used in cross-certification and qualified subordination | Signature | CrossCA | Yes |
| Directory E-mail Replication | Used to replicate e-mail within Active Directory | Signature and encryption | DirEmailRep | Yes |

| | | | | |
|---|---|---|---|---|
| Domain Controller | All-purpose certificates held by domain controllers | Signature and encryption | DirEmailRep | Yes |
| Domain Controller Authentication | Used to authenticate Active Directory computers and users | Signature and encryption | Computer | No |
| EFS Recovery Agent | Allows the subject to decrypt files previously encrypted with EFS | Encryption | User | No |
| Enrollment Agent | Used to request certificates on behalf of another subject | Signature | User | No |
| Enrollment Agent (Computer) | Used to request certificates on behalf of another computer subject | Signature | Computer | No |
| Exchange Enrollment Agent (Offline request) | Used to request certificates on behalf of another subject and supply the subject name in the request | Signature | User | No |
| Exchange Signature Only | Used by Microsoft Exchange Key Management Service to issue certificates to Exchange users for digitally signing e-mail | Signature | User | No |
| Exchange User | Used by Microsoft Exchange Key Management Service to issue certificates to Exchange users for encrypting e-mail | Encryption | User | Yes |
| IPSEC | Used by IP Security (IPSec) to digitally sign, encrypt, and decrypt network communication | Signature and encryption | Computer | No |
| IPSEC (Offline request) | Used by IP Security (IPSec) to digitally sign, encrypt, and decrypt network communication when the subject name is supplied in the request | Signature and encryption | Computer | No |
| Key Recovery Agent | This certificate can recover private keys archived on the certification authority. | Encryption | KRA | Yes |
| Root Certification Authority | Used to prove the identity of the root certification authority | Signature | CA | Yes |
| Router (Offline request) | Used by a router when requested through SCEP from a CA that holds a CEP Encryption certificate | Signature and encryption | Computer | No |
| Smartcard Logon | Allows the holder to authenticate using a smart card | Signature and encryption | User | No |
| Smartcard User | Allows the holder to authenticate and protect e-mail using a smart card | Signature and encryption | User | Yes |
| Subordinate Certification Authority | Used to prove the identity of the root certification authority, issued by the parent or root certification authority | Signature | CA | Yes |
| Trust List Signing | The holder can digitally sign a trust list. | Signature | User | No |
| User | Certificate to be used by users for e-mail, EFS, and client authentication | Signature and encryption | User | Yes |
| User Signature Only | Allows users to digitally sign data | Signature | User | No |

| | | | | |
|---|---|---|---|---|
| Router (Offline request) | Used by a router when requested through SCEP from a CA that holds a CEP Encryption certificate | Signature and encryption | Computer | No |
| Smartcard Logon | Allows the holder to authenticate using a smart card | Signature and encryption | User | No |
| Smartcard User | Allows the holder to authenticate and protect e-mail using a smart card | Signature and encryption | User | Yes |
| Subordinate Certification Authority | Used to prove the identity of the root certification authority, issued by the parent or root certification authority | Signature | CA | Yes |
| Trust List Signing | The holder can digitally sign a trust list. | Signature | User | No |
| User | Certificate to be used by users for e-mail, EFS, and client authentication | Signature and encryption | User | Yes |

**Certificate Template Options**

The following sections detail the specific information that can be configured for a version 2 certificate template by using the Certificate Templates Microsoft Management Console (MMC) snap-in.

**Note**

- Because certificate templates exist only once within a forest, all changes made to the templates apply on a forest-wide level. If several CAs exist within one forest, templates can be assigned to CAs. Nevertheless, all CAs within a forest use one set of templates.

The **General** tab is the first tab that appears when a certificate template is duplicated. The **General** tab allows administrators to configure the following attributes of the certificate template:

- **Template display name,** The display name shown in the Certificate Templates, Certificates, and Certification Authority snap-ins.

- **Template name.** The name of the certificate template object created in the CN=Certificate Templates,CN=Public Key Services,CN=Services,DC=*ForestRootDomain* container.

- **Validity period.** Defines the validity period for an issued certificate. The validity period must not be greater than the validity period of the CA's certificate.

- **Renewal period.** The time before the validity period expires when the certificate will be renewed, if reenrollment is supported for the certificate template. By default, the minimum renewal period is 80 percent of the certificate lifetime or 6 weeks, whichever is greater.

- **Publish options.** Determines whether the certificate can be published to Active Directory. Certificates are published as an attribute of the security principal contained in the subject of the issued certificate.

- **Reenrollment option.** If the certificate template is published to Active Directory, this prevents reenrollment if a valid certificate of the same certificate template exists for the security principal indicated in the subject attribute of the certificate.

  **Note**

  - Re-enrollment settings mainly affect auto-enrollment design in a Windows Server 2003 network.

**Request Handling**

The options on the **Request Handling** tab define the purpose of the certificate template, the supported cryptographic service providers (CSPs), minimum key length, exportability, auto-enrollment settings, and whether strong private key protection should be required.

#### Certificate Purposes

The certificate purpose defines the intended primary use of the certificate. The certificate purpose can be one of four settings:

- **Encryption.** A certificate with this purpose will contain cryptographic keys for encryption and decryption.
- **Signature.** A certificate with this purpose will contain cryptographic keys for signing data only.
- **Signature and encryption.** A certificate with this purpose covers all primary uses of a certificate's cryptographic key, including encryption of data, decryption of data, initial logon, or digitally signing data.
- **Signature and smartcard logon.** A certificate with this purpose allows for initial logon with a smart card, and digitally signing data; it cannot be used for data encryption.

#### Archive Settings

When subject's lose their private keys due to the user profile becoming corrupted or being lost, or when a computer is stolen, any information that was persistently encrypted by using the corresponding public key is inaccessible. To help avoid this, Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition make it possible for a subject's encryption keys to be archived in the CA database when certificates are issued. If subject's lose their encryption keys, the information can be retrieved from the database and securely provided to the subject's. This allows the encrypted information to be recovered instead of lost.

### Note

- Signing keys cannot be archived. If signing keys are lost, the user must request a new set of signing keys.

The following key archival settings are defined on the **Request Handling** tab:

- **Archive subject's encryption private key.** If the issuing Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition CA is configured for key archival, the subject's encryption key will be archived.
- **Allow private key to be exported.** When this option is specified, the subject's private key can be exported for backup or transportation.
- **Deleting revoked or expired certificates (do not archive).** If a certificate is renewed with autoenrollment due to expiration or revocation, the previously issued certificate is removed from the subject's certificate store. By default, the certificate is archived.
- **Include symmetric algorithms allowed by the subject.** When the subject requests the certificate, they can supply a list of supported symmetric algorithms. This option allows the issuing CA to include those algorithms in the certificate, even if they are not recognized or supported by that server. The algorithms are used by applications like secure e-mail.

To enable key archival and recovery, the following configuration options must be set in the certificate template and at the Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition CA:

- The specific certificate template must be configured to allow key archival.
- One or more key recovery agents must be identified on the CA, and key recovery agent certificates must be issued to those agents.
- Key archival must be configured on the CA.

#### User Input Settings

The **Request Handling** tab also allows several user input settings to be defined for a certificate template. These settings can be used to implement strong private key protection. These settings include:

- **Enroll subject without requiring any user input.** This option allows auto-enrollment without any user interaction and is the default setting for both computer and user certificates. This option must not be enabled for computer certificates.

- **Prompt the user during enrollment.** Although useful when testing initial auto-enrollment deployments, typically, this option is disabled. By disabling this option, users do not have to provide any input for the installation of a certificate based on the certificate template.

- **Prompt the user during enrollment and require user input when the private key is used.** This option enables the user to set a strong private key protection password on the user's private key when the key is generated, and requires the user to use it whenever the certificate and private key are used. This option must not be enabled for computer certificates or smart card user certificates.

  **Note**

  - To enable smart card auto-enrollment, the **Prompt the user during enrollment** check box must be selected so that the user is prompted to insert the smart card in the smart card reader when required.

**Other Request Handling Settings**

In addition to key archival settings, some general options that affect all certificates, including those that do not enable key archival, can be defined. These include:

- **Minimum key size.** This specifies the minimum size, in bits, of the key that will be generated for this certificate.

- **Cryptographic service providers.** This is a list of CSPs that will be used to enroll certificates for the given template. The list is compiled from all CSPs that are available on the computer that is used to maintain the certificate template. Selecting one or more CSPs configures the certificate to only work with those CSPs. If you do not select a specific CSP, the certificate enrollment will work with any installed CSP, but will use the first CSP from the enumeration list. The CSP must be installed on the client workstation for the CSP to be used during enrollment. If a specific CSP is chosen and not available on a client computer, enrollment will fail. If you need to restrict the CSP choice to a specific third party CSP, you must have the CSP installed on the computer that is used to configure the certificate template.

  **Note**

  - For more information, see "CSPs" later in this section.

**Subject Name**

When configuring a certificate template, subject name guidelines must be defined. The subject name of a certificate identifies the user, computer, or service that the certificate represents. This is included in the issued certificate template and must uniquely identify the subject. The subject name can either be built automatically during the certificate request by using the authentication name of the subject or explicitly defined by the subject and included in the certificate request. Windows obtains the attributes as defined in the CA's registry from Active Directory for automatic building. To provide the name manually, the subject supplies that information in the certificate request, for example by using the Web-based enrollment pages.

A number of options can be included with the subject name, in addition to specific configuration settings for the subject name when the subject name is built from Active Directory information during the certificate request process. The format of the subject name can be defined as:

- **None.** Does not enforce any name format for this field.

- **Common name.** The CA creates the subject name from the common name (CN) of the requestor obtained from Active Directory. These should be unique within a domain, but might not be unique within an enterprise.

- **Fully distinguished name.** The CA creates the subject name from the fully distinguished name obtained from Active Directory. This guarantees that the name is unique within an enterprise.

In addition, you can choose to include the e-mail name in the subject name. This information is populated from the e-mail attribute of an account, and is included with either the CN or fully distinguished name as part of the subject name.

In addition to the subject name, you can also choose which name formats are included in the alternate subject name attributes of issued certificates. The alternate subject name formats that are available include:

- **E-mail name.** If the e-mail name field is populated in the Active Directory user object, that e-mail name will be used for user accounts.

  **Note**

  - The e-mail name is required for user certificates. If the e-mail name is not populated for a user in Active Directory, the certificate request by that user will fail.

- **Domain Name System (DNS) name.** The FQDN of the subject that requested the certificate is used for computer accounts.

- **User principal name (UPN).** The UPN is part of the Active Directory user object and will be used for user accounts.

- **Service principal name (SPN).** The SPN is part of the Active Directory computer object and will be used for computer accounts.

Usually, a subject cannot request a certificate that uses a nonmatching subject name. For example, user1@example.com would not be allowed to request a certificate with a subject name of user2@example.com.

The only subject that can request a certificate for another user is one who holds a certificate based on the Enrollment Agent template. That subject can request certificates on behalf of any other subject. For example, an enrollment agent can request Smart Card User or Smart Card Logon certificates on behalf of other users.

The only case in which a subject can request a certificate with a different subject name is when the request holds a certificate based on the Enrollment Agent template. The Enrollment Agent template allows a subject to request a certificate on behalf of another subject.

#### Issuance Requirements

The **Issuance Requirements** tab allows higher assurance-level certificates to be placed in a pending state until the certificate is reviewed before it is issued. The following settings configure the authentication and signature requirements for issuance certificates that are based on a template:

- **CA certificate manager approval.** All certificate requests are placed into the pending container of the CA for a certificate manager to issue or deny. Any defined certificate manager can issue or deny a certificate request of this type.

- **This number of authorized signatures.** This setting requires the Certificate Management Messages over CMS (CMC) certificate request to be digitally signed by one or more authorized parties before it can be issued.

  **Note**

  - Defining the number of authorized signatures enables several other configuration parameters.

- **Policy type required in signature.** This option is only enabled when the **This number of authorized signatures** value is set. This option defines which specific application policy, issuance policy, or both are required in the signing certificate. This is how the CA determines whether the signature is appropriate for authorizing the issuance of the subject's certificate.

- **Application policy.** This option specifies the application policy that must be included in the signing certificate used to sign the certificate request. It is enabled when **Policy type required in signature** is set to either **Application policy** or **Both application and issuance policy**.

- **Issuance policy.** This option specifies the issuance policy that must be included in the signing certificate used to sign the certificate request. This option is enabled when **Policy type required in signature** is set to either **Issuance policy** or **Both application and issuance policy**.

#### Superseded Templates

The **Superseded Templates** tab allows existing certificate templates to be superseded with a modified version 2 certificate template. The following scenarios are supported:

- A version 2 certificate template can supersede a version 1 certificate template.
- A version 2 certificate template can supersede a version 2 certificate template.
- A common version 2 certificate template can supersede multiple existing certificate templates.

#### Extensions

The **Extensions** tab allows an administrator to define specific application policies, issuance policies, certificate subject types, and key usage attributes for a certificate template. The following sections detail the specifics on each form of extension defined in a certificate template.

#### Application Policies

Application policies are settings that inform a target that the subject holds a certificate that can be used to perform a specific task. They are represented in a certificate by an OID that is defined for a given application. This OID is included in the issued certificate. When a subject presents its certificate, the certificate can be examined by the target to verify the application policy and determine whether the subject can perform the requested action.

The following table shows some of the more commonly used application policies and their related OIDs.

**Commonly Used Application Policies and OIDs**

| Application Policy | OID |
|---|---|
| Client Authentication | 1.3.6.1.5.5.7.3.2 |
| CA Encryption Certificate | 1.3.6.1.4.1.311.21.5 |
| Smart Card Logon | 1.3.6.1.4.1.311.20.2.2 |
| Document Signing | 1.3.6.1.4.1.311.10.3.12 |
| File Recovery | 1.3.6.1.4.1.311.10.3.4.1 |
| Key Recovery | 1.3.6.1.4.1.311.10.3.11 |
| Microsoft Trust List Signing | 1.3.6.1.4.1.311.10.3.1 |
| Qualified Subordination | 1.3.6.1.4.1.311.10.3.10 |
| Root List Signer | 1.3.6.1.4.1.311.10.3.9 |
| Encrypting file system | 1.3.6.1.4.1.311.10.3.4 |

By restricting which application policies are defined in a certificate template, a certificate cannot be used for undesired transactions. For example, a certificate with the Secure E-mail OID cannot be used for client authentication.

Because some implementations of PKI applications might not understand application policies, both application policies and EKU fields appear in certificates that a Microsoft CA issues. EKU is similar to application policy in that EKU also defines the functions of certificate.

#### Certificate Policies

Certificate policies, also referred to as issuance policies, define the measures that are used to identify the subject of the certificate and thereby define the level of assurance for an issued certificate. For example, your organization might require a face-to-face meeting before the certificate is issued to provide for a higher level of assurance for the issued certificate.

**Note**

- Certificate policies refer to the certificate policies extension identifier described in RFC 2459.

Issuance policies are represented in a certificate by an OID that is defined at the CA. This OID is included in the issued certificate. When a subject presents its certificate, the certificate can be examined by the target to verify the issuance policy and determine whether that level of issuance policy is sufficient for the subject to perform the requested action.

The following table describes the three default certificate policy OIDs included in Windows Server 2003.

**Default certificate policy OIDs in Windows Server 2003.**

| OID type | Description |
|---|---|
| Low assurance | Provides no additional mechanism to identify the subject of the certificate. For example, a certificate that is issued based only on the credentials provided can be a low assurance certificate. |
| Medium assurance | Requires additional validation of the certificate's subject. For example, a smart card certificate might require an administrator to have a face-to-face meeting with an employee before issuing the smart card to an employee. |
| High assurance | Requires research into the subject's identity. For example, a high assurance certificate might require that an organization perform a background check on an employee before issuing the certificate. |

**Note**

- The low, medium, and high assurance OIDs are unique for each Active Directory forest.

One or more issuance policies can be set for a certificate template. Because these issuance policies are simply text labels with an associated OID, no options are associated with them. The only special issuance policy is **All issuance policies**, which indicates that this policy includes all others. This is normally reserved for certificates held by CAs.

Issuance policies are often used when configuring qualified subordination (cross-certification) between PKI hierarchies to ensure that certificates that are recognized by another organization's PKI meet your organization's issuance requirements.

You can also configure a certificate template to increase the issuance security of a certificate by requiring the user or computer to provide additional forms of identification for the certificate request. The additional forms of identification can include providing photo identification, meeting face to face with a local registration authority, or signing the certificate request with a previously issued signing certificate. These can be implemented by using the following issuance options:

- **Certificate Manager Approval.** This setting sets all certificate requests to a pending state until a certificate manager issues or denies the certificate request. The certificate manager must first validate the identity of the certificate requestor before issuing or denying the certificate request. In some cases, the certificate manager will record any forms of identification that the user presents into a custom certificate issuance database application.

- **Signing Requests.** An existing certificate can sign a certificate request to increase the issuance security. You can configure a certificate template to require a signature by using a certificate request with a specific application policy OID, certificate policy OID, or combination of application and certificate policy OIDs. The assumption here is that the possession of the private key associated with the signing certificate increases the issuance security of the certificate request.

**Certificate Template Information**

The certificate template information, also referred to as the certificate subject type, defines the purpose of a certificate template. Six subject types are recognized by Windows Server 2003 CAs:

- Key recovery agent
- Directory e-mail replication
- Cross-certified CA
- CA
- Computer
- User

The certificate template information extension cannot be edited. If an administrator requires a specific subject type to be applied to a certificate, the administrator should clone from a certificate template that includes the required subject type.

**Key Usage**
A certificate enables the subject to perform a specific task. To help control the usage of a certificate outside its intended purpose, restrictions are automatically placed on certificates. Key usage is a restriction method that determines what a certificate can be used for. It allows the administrator to issue certificates that can only be used for specific tasks or certificates that can be used for a broad range of functions. If no key usage is specified, the certificate can be used for any purpose.

For signatures, key usage can be limited to one or more of the following purposes:

- Digital signature
- Signature is a proof of origin (nonrepudiation)
- Certificate signing
- CRL signing

For encryption key usage, the following options are available:

- Key exchange without key encryption
- Key exchange only with key encryption

**Security**
The **Security** tab allows you to define the discretionary access control list (DACL) for a specific certificate template. The permissions that can be assigned to a certificate template include:

- **Full Control.** This permission allows a security principal to modify all attributes of a certificate template, including the permissions for the certificate template.
- **Read.** This permission allows a security principal to see the certificate template when enrolling for certificates. It is required for a security principal to enroll or auto-enroll a certificate; it is required by the certificate server to find the certificate templates in Active Directory.
- **Write.** This permission allows a security principal to modify the attributes of a certificate template, including the permissions assigned to the certificate template.
- **Enroll.** This permission allows a security principal to enroll for a certificate based on the certificate template. To enroll for a certificate, the security principal must also have Read permissions for the certificate template.
- **Autoenroll.** This permission allows a security principal to receive a certificate through the auto-enrollment process. Auto-enrollment permissions require that the user has both Read and Enroll permissions in addition to the Auto-enroll permission.

  **Note**

  - It is recommended that certificate template permissions be assigned only to global groups or universal groups. Because the certificate template objects are stored in the Configuration naming context, you cannot assign permissions using domain local groups. To simplify administration, it is never recommended to assign certificate template permissions to individual user or computer accounts.

**Certificate Template Objects Defined in the Windows Schema**
The Certificate Templates container contains the certificate templates that are defined within an Active Directory forest. Each certificate template is of the class pKICertificate.

Each template is stored in the following location in the Configuration naming context:

```
CN=<name of template>,CN=Certificate Templates,CN=Public Key
Services,CN=Services,CN=Configuration,DC= ForestRootDomain
```

The following version 1 attributes are defined in the Active Directory schema.

**Version 1 Attributes**

| Attribute Name | Description |
|---|---|
| cn | Common name of the certificate type |
| distinguishedName | The common name of the certificate type |
| displayName | The display name of a certificate type |
| pKIExtendedKeyUsage | An array of extended key usage object identifiers (OIDs) for a certificate type |
| pKIDefaultCSPs | The list of default CSPs for this certificate type |
| pKICriticalExtensions | The list of critical extensions |
| revision | The major version of the template |
| templateDescription | The description of the template |
| flags | General enrollment flags (see below) |
| pKIDefaultKeySpec | |
| NTSecurityDescriptor | Security Descriptor name |
| pKIKeyUsage | Key Usage extension |
| pKIMaxIssuingDepth | Basic Constraints |
| pKIExpirationPeriod | Certificate validity period |
| pKIOverlapPeriod | Certificate renewal period |

The following version 2 attributes are defined in the Active Directory schema.

**Version 2 Attributes**

| Attribute name | Description |
|---|---|
| msPKI-Template-Schema-Version | The schema version of the template |
| msPKI-Template-Minor-Revision | The minor version of the template |
| msPKI-RA-Signature | The number of RA signatures required on a request referencing this template |
| msPKI-Minimal-Key-Size | The minimal key size required |
| msPKI-Template-Cert-Template-OID | The OID of this template |
| msPKI-Supersede-Templates | The OID of the template that this template supersedes |
| msPKI-RA-Policies | The RA issuer policy OIDs required in certificates used to sign a request |

| | |
|---|---|
| msPKI-RA-Application-Policies | The RA application policy OIDs required in certificates used to sign a request |
| msPKI-Certificate-Policy | The certificate issuer policy OIDs are placed in the OID_CERT_POLICIES extension by the policy module |
| msPKI-Certificate-Application-Policy | The certificate application policy OIDs are placed in the OID_APPLICATION_CERT_POLICIES extension by the policy module |
| msPKI-Enrollment-Flag | Enrollment flags (see below) |
| msPKI-Private-Key-Flag | Private key flags (see the following section) |
| msPKI-Certificate-Name-Flag | Subject name flags (see the following section) |

**Flags**

The following enrollment flags are defined in the Active Directory schema.

**Enrollment Flags**

| Flag | Description |
|---|---|
| CT_FLAG_INCLUDE_SYMMETRIC_ALGORITHMS<br>0x00000001 | Include the S/MIME symmetric algorithms in the requests |
| CT_FLAG_PEND_ALL_REQUESTS<br>0x00000002 | All certificate requests are pended |
| CT_FLAG_PUBLISH_TO_KRA_CONTAINER<br>0x00000004 | Publish the certificate to the key recovery agent (KRA) container in the Active Directory |
| CT_FLAG_PUBLISH_TO_DS<br>0x00000008 | Publish the resultant certificate to the userCertificate property on the user object in the Active Directory |
| CT_FLAG_AUTO_ENROLLMENT<br>0x00000020 | This certificate can be issued through auto-enrollment |
| CT_FLAG_PREVIOUS_APPROVAL_VALIDATE_REENROLLMENT<br>0x00000040 | A previously issued certificate will validate subsequent enrollment requests |
| CT_FLAG_DOMAIN_AUTHENTICATION_NOT_REQUIRED<br>0x00000080 | Domain authentication is not required |
| CT_FLAG_USER_INTERACTION_REQUIRED<br>0x00000100 | User interaction is required to enroll using auto-enrollment |
| CT_FLAG_ADD_TEMPLATE_NAME<br>0x00000200 | Add the szOID_CERTTYPE_EXTENSION (template name) extension Note: This flag will only be set on version 1 certificate templates and only for Windows 2000 CAs |
| CT_FLAG_REMOVE_INVALID_CERTIFICATE_FROM_PERSONAL_STORE<br>0x00000400 | Remove invalid (expired or revoked) certificates from the personal store on local client computers during auto-enrollment |

The following subject name flags are defined in the Active Directory schema.

**Subject Name Flags**

| Flag | Description |
| --- | --- |
| CT_FLAG_ENROLLEE_SUPPLIES_SUBJECT<br>0x00000001 | The enrolling application must supply the subject name |
| CT_FLAG_ENROLLEE_SUPPLIES_SUBJECT_ALT_NAME<br>0x00010000 | The enrolling application must supply the subjectAltName in the certificate request |
| CT_FLAG_SUBJECT_REQUIRE_DIRECTORY_PATH<br>0x80000000 | The subject name should be the full distinguished name (DN) based on the Active Directory path |
| CT_FLAG_SUBJECT_REQUIRE_COMMON_NAME<br>0x40000000 | The subject name should be the common name |
| CT_FLAG_SUBJECT_REQUIRE_EMAIL<br>0x20000000 | The subject name must include the e-mail name |
| CT_FLAG_SUBJECT_REQUIRE_DNS_AS_CN<br>0x10000000 | The subject name must include the DNS name as the common name |
| CT_FLAG_SUBJECT_ALT_REQUIRE_DNS<br>0x08000000 | The subject alternative name must include the DNS name |
| CT_FLAG_SUBJECT_ALT_REQUIRE_EMAIL<br>0x04000000 | The subject alternative name must include the e-mail name |
| CT_FLAG_SUBJECT_ALT_REQUIRE_UPN<br>0x02000000 | The subject alternative name must include the UPN |
| CT_FLAG_SUBJECT_ALT_REQUIRE_DIRECTORY_GUID<br>0x01000000 | The subject alternative name must include the directory GUID (used by domain controllers) |
| CT_FLAG_SUBJECT_ALT_REQUIRE_SPN<br>0x00800000 | The subject alternative name must include the service principal name (SPN) |
| CT_FLAG_SUBJECT_ALT_REQUIRE_DIRECTORY_GUID<br>0x01000000 | The subject alternative name must include the directory GUID |
| CT_FLAG_SUBJECT_ALT_REQUIRE_SPN<br>0x00800000 | The subject alternative name requires SPN |

The following private key flags are defined in the Active Directory schema.

**Private Key Flags**

| Flag | Description |
| --- | --- |
| CT_FLAG_ALLOW_PRIVATE_KEY_ARCHIVAL<br>0x00000001 | Archival of the private key is allowed or required |
| CT_FLAG_EXPORTABLE_KEY<br>0x00000010 | Mark the private key as exportable |

The following general flags are defined in the Active Directory schema.

**General Flags**

| Flag | Description |
|---|---|
| CT_FLAG_MACHINE_TYPE<br>0x00000040 | This is a machine certificate |
| CT_FLAG_IS_CA<br>0x00000080 | This is a CA certificate |
| CT_FLAG_IS_CROSS_CA<br>0x00000800 | This is a cross-CA certificate |
| CT_FLAG_IS_DEFAULT<br>0x00010000 | Default certificate type that is set on all version 1 templates that cannot be modified |
| CT_FLAG_IS_MODIFIED<br>0x00020000 | The certificate type has been modified (read only) |
| CT_MASK_SETTABLE_FLAGS<br>0x0000ffff | Flag that can be set for general flags |

**CSPs**

Microsoft CryptoAPI provides a secure interface to the installable CSP modules that perform all cryptographic operations and manage private keys for all Windows Server 2003 cryptographic services, including Certificate Services.

Vendors can develop hardware or software CSPs that support a wide range of cryptographic operations and technologies. However, Microsoft must certify and digitally sign all CSPs. CSPs do not work in Windows Server 2003 unless they have been digitally signed by Microsoft.

**Note**

- For more information about obtaining a Microsoft digital signature for a CSP, see the Platform Software Development Kit (SDK) on MSDN.

Windows Server 2003 includes the following Microsoft CSPs:

- **Microsoft Base Cryptographic Service Provider:** The Base CSP provides a broad set of basic cryptographic functions. It is not subject to United States government cryptography export restrictions and can be exported to other countries (subject to general United States export restrictions in addition to the import restrictions of other countries). The Base CSP uses Rivest-Shamir-Adelman (RSA) technology, which is licensed from RSA Data Security, Inc.

- **Microsoft Strong Cryptographic Service Provider:** The Strong CSP is an extension of the Base CSP available with Windows 2000 and Windows Server 2003. It supports all the algorithms in the Enhanced CSP, but the key lengths are the same as in the Base CSP. It supports RSA, Data Encryption Standard (DES), 3DES, RC2, and RC4 in addition to (Message-Digest Algorithm) MD5 and Secure Hash Algorithm 1 (SHA1).

- **Microsoft Enhanced Cryptographic Service Provider:** The Enhanced CSP provides the same capabilities as the Base CSP, but enforces stronger security by supporting longer key lengths and additional cryptographic algorithms. The Enhanced CSP also uses RSA technology.

The following table highlights differences between the Base CSP and the Enhanced CSP. The public key lengths shown in the table are the default key lengths.

**Comparison of Microsoft Base CSP and Microsoft Enhanced CSP**

| Algorithm | Base CSP | Strong CSP | Enhanced CSP |
|---|---|---|---|

| | | | |
|---|---|---|---|
| RSA public key signature algorithm | Key length: 512 bits. | Key length: 1,024 bits. | Key length: 1,024 bits. |
| RSA public key exchange algorithm | Key length: 512 bits. | Key length: 1,024 bits. | Key length: 1,024 bits. |
| RC2 block encryption algorithm | Key length: 40 bits. | Key length: 128 bits. | Key length: 128 bits. Salt length: Settable. |
| RC4 stream encryption algorithm | Key length: 40 bits. | Key length: 128 bits. | Key length: 128 bits. Salt length: Settable. |
| DES | Key length: 56 bits. | Key length: 56 bits. | Key length: 56 bits. |
| Triple DES (2-key) | Not supported. | Key length: 112 bits. | Key length: 112 bits. |
| Triple DES (3-key) | Not supported. | Key length: 168 bits. | Key length: 168 bits. |
| SHA1 | Key length: 160 bits. | Key length: 160 bits. | Key length: 160 bits. |
| Message Digest 2 (MD2) | Key length: 128 bits. | Key length: 128 bits. | Key length: 128 bits. |
| Message Digest 4 (MD4) | Key length: 128 bits. | Key length: 128 bits. | Key length: 128 bits. |
| Message Digest 5 (MD5) | Key length: 128 bits. | Key length: 128 bits. | Key length: 128 bits. |
| SSL3 (SHAMD5) | Key length: 288 bits. | Key length: 288 bits. | Key length: 288 bits. |
| Message Authentication Code (MAC) | Key length: 512 bits. | Key length: 1,024 bits. | Key length: 1,024 bits. |

Public keys used with the Base, Strong, and Enhanced CSPs for digital signatures can be up to 16,384 bits long. However, the public keys used for key encryption and key exchange (to protect secret keys) are limited to a maximum of 1,024 bits for the Base and Strong CSPs, and 16,384 bits for the Enhanced CSP. In addition, the symmetric keys for the encryption algorithms in the Base CSP are limited to shorter key lengths, resulting in significantly weaker cryptographic security. Overall, the key lengths and the encryption algorithms in the Enhanced CSP provide far stronger cryptographic security.

For the Base, Strong, and Enhanced CSP, public keys used for signing or key exchange can be a minimum of 384 bits long. However, public keys of at least 1,024 bits are recommended, unless they adversely affect computer performance.

The default public key length of the Strong and Base CSPs is 512 bits, and the default public key length of the Enhanced CSPs is 1,024 bits. Windows Server 2003 Certificate Services uses the default public key lengths of the CSP, although any key length the CSP supports can be configured.

The Strong and Enhanced CSPs are compatible with the Base CSP except that the CSPs can generate only RC2 or RC4 keys of the default key length. The default symmetric key length for RC2 and RC4 in the Base CSP is 40 bits. The default symmetric length for RC2 and RC4 in the Strong and Enhanced CSP is 128 bits. Therefore, the Enhanced CSP cannot create keys with Base CSP–compatible key lengths. However, the Strong and Enhanced CSPs can import RC2 and RC4 keys of up to 128 bits. Therefore, the Strong and Enhanced CSPs can import and use 40-bit keys that were generated by using the Base CSP.

Other Microsoft CSPs include:

- **Microsoft Base DSS Cryptographic Service Provider.** This CSP provides data signing and signature verification capability by using the SHA and Digital Signature Algorithm (DSA).

- **Microsoft Base DSS and Diffie-Hellman Cryptographic Service Provider.** This CSP provides a superset of the DSS Cryptographic Provider and also supports Diffie-Hellman key exchange, hashing (message digests), data signing, and signature verification by using the SHA and DSA algorithms.

- **Microsoft Base DSS and Diffie-Hellman/Schannel Cryptographic Service Provider.** This CSP offers various cryptographic services that are required for data integrity, session key exchange, and authentication during secure Web communications by using the SSL and Transport Layer Protocol (TLS) protocols.

- **Microsoft RSA/Schannel Cryptographic Provider:** This CSP is similar to the Microsoft Base DSS and Diffie-Hellman/Schannel Cryptographic Provider in that it works with SSL3 and TLS1, but uses the RSA suite of algorithms rather than Diffie-Hellman.

- **Smart Card Cryptographic Service Providers:** Windows Server 2003 includes smart card CSPs from a number of vendors, including Schlumberger, Gemplus, and Infineon.

**Local Certificate Storage**
In Windows Server 2003, public-key objects such as certificates, CRLs, and certificate trust lists (CTLs) are stored in certificate stores for use by users, services, and computers. The Windows Server 2003 certificate stores include physical stores and logical stores.

Physical certificate stores can be local or remote:

- Local objects are stored in the registry of the computer.

- The user object in Active Directory has a virtual store on the userCertificate attribute,

Many public key objects in the physical stores are shared among users, services, and computers through the use of logical certificate stores. These logical certificate stores group certificates together in logical, functional categories for users, computers, and services and provide pointers to the physical certificate stores. The use of logical certificate stores eliminates the need to store duplicates of common public key objects, such as trusted root certificates, CTLs, and CRLs for users, computers, and services.

However, some certificates, CTLs, or CRLs are issued for use only by an individual service, user, or local computer. Therefore, users, computers, and services also have individual stores that provide a place to store certificates, CTLs, or CRLs that are not shared in common. For example, a user can request and obtain a certificate or a CRL, which appears in the individual's logical store and is physically stored in the user's unique certificate store in the registry. Such individual user certificates and CRLs are not shared with local computers or with services.

In addition, some public-key objects, such as trusted root certificates and CTLs, can be distributed through Public Key Group Policy. Public key objects that are distributed through Group Policy are stored in special areas of the registry and appear in the logical stores for users, computers, and services. When you use Group Policy, separate CTLs can be created for users and computers. The CTLs for users are not shared with services or the computer. However, the CTLs for computers are shared with users and services.

The logical certificate stores include the following categories for users, computers, and services:

- **Personal** Contains individual certificates for the user, service, or computer. For example, when an enterprise CA issues you a User certificate, the certificate is installed in the Personal store for your user account.

- **Trusted Root Certification Authorities** Contains certificates for root CAs. Certificates with a certification path to a root CA certificate are trusted by the computer for all valid purposes of the certificate.

- **Disallowed Certificates** These are certificates that a user has explicitly decided not to trust using either Software Restriction policy or by clicking "Do not trust this certificate" when the decision is presented in an e-mail application or a Web browser.

- **Third-Party Root Certification Authorities** Trusted root certificates from certification authorities other than Microsoft and your organization.

- Enterprise Trust Contains CTLs. Certificates with a certification path to a CTL are trusted by the computer for purposes specified in the CTL.

- **Intermediate Certification Authorities** Contains certificates for CAs that are not trusted root certificates (for example, certificates of subordinate CAs), but that are required to validate certification paths. This store also contains CRLs for use by the user, service, or computer, and cross certificates..

- **Active Directory User Object** Contains certificates that are published in Active Directory for the user. This store appears in the Certificates snap-in for users only, not for computers or services.

- **Request** Contains pending or rejected certificate requests. This store appears only in the Certificates snap-in after a certificate request has been made for the user, computer, or service.

- **SPC** Contains certificates for software publishers that are trusted by the computer. Software that has been digitally

signed by publishers with certificates in this store is downloaded without prompting the user. By default, this store is empty. When Microsoft Internet Explorer downloads software that has been signed by a software publisher for the first time, users are prompted to choose whether they want to trust all software that is signed by this publisher. If a user chooses to trust all software signed by the publisher, the publisher's software publisher certificate (SPC) is added to the SPC store. This store appears in the Certificates snap-in for the local computer only, not for users or services.

- **Trusted People** Certificates issued to people or end entities that are explicitly trusted. Most often these are self-signed certificates or certificates explicitly trusted in an application such as Microsoft Outlook.

- **Other People** Certificates issued to people or end entities that are implicitly trusted. These certificates must be part of a trusted certification hierarchy. Most often these are cached certificates for services like Encrypting File System, where certificates are used for creating authorization for decrypting an encrypted file.

- **Certificate Enrollment Requests** Pending or rejected certificate requests.

- **Server Authentication** Certificates that server programs use to authenticate themselves to clients.

- **Client Authentication** Certificates that client programs use to authenticate themselves to servers.

- **Code Signing** Certificates associated with key pairs used to sign active content.

- Secure E-mail Certificates associated with key pairs used to sign e-mail messages.

- **Encrypting File System** Certificates associated with key pairs that encrypt and decrypt the symmetric key used for encrypting and decrypting data by the encrypting file system.

- **File Recovery** Certificates associated with key pairs that encrypt and decrypt the symmetric key used for recovering encrypted data by the encrypting file system.

  **Note**

  - Other applications can also create their own certificate stores.

Changes to the logical certificate stores are made to the appropriate physical stores that are located in either the registry or Active Directory. Because you use only the logical certificate store for a user, service, or computer, you do not have to track where the certificates are actually stored or edit the registry to manage the certificate stores. However, it is a good idea to track changes made to the physical and logical stores; this information can be useful when you try to determine why certificate-enabled actions succeed or fail.

**Viewing Physical and Logical Stores**

Certificate management tools distinguish between the physical structure of certificate stores and their logical abstraction. The logical abstraction, which simplifies the physical structure, was implemented to group certificates by function or purpose and to enable administrators to understand certificate stores in a simpler way.

A number of tools are available to view the physical and logical stores:

- Internet Explorer shows only the logical certificate stores.

- The Certificates snap-in shows the logical structure by default; however, it is possible to view the physical structure. On the **View** menu, click **Options**, and then click **Show the following: Physical certificate stores**.

- The Registry Editor (Regedit.exe) displays only the physical certificate store structure.

- The Certutil command-line tool (Certutil.exe) refers to certificate stores by labels that are equal to the store names in the registry or LDAP directory.

  **Note**

- For more information about the certificate management tools and registry stores for certificates, CRLs, and CTLs, see "Resource Kit Tools" in Tools and Settings Collection.

**Private Keys Stored in User Profiles**
Users' private keys are stored on the local file system in the following locations:

- $:\Documents and Settings\%UserName%\Application Data\Microsoft\Crypto\RSA
- $:\Documents and Settings\%UserName%\Application Data\Microsoft\Crypto\DSS

Computer private keys are stored on the local file system in the following locations:

- $:\Documents and Settings\All users\Application Data\Microsoft\Crypto\RSA\Machine Keys
- $:\Documents and Settings\All users\Application Data\Microsoft\Crypto\DSS\Machine Keys

The **Security** properties page indicates who can access this private key.

Windows associates certificates with their owner. By default, every subject that has permissions to log on to the computer or run as a service owns a set of certificate stores.

The number of certificates per store is limited only by memory and hard disk space. For certificates that are stored in the registry, the limiting factor is the total size of the registry. For information about the physical store location, see the individual store description later in this section.

**Personal Certificates Stored in User Profiles**
Certificates that are stored in the personal certificate store belong to the current user. These certificates in appear in the file system in:

```
%USERPROFILE%\Application Data\Microsoft\Systemcertificates\My
```

Access to the personal certificate store is restricted to the user who owns the certificate.

The personal store receives certificates from auto-enrollment or when a user requests a certificate manually. If a user imports a personal certificate, the certificate is stored in the Personal container.

This store maintains certificates of any type. A user who owns an S/MIME certificate, an EFS, certificate and a client-authentication certificate, for example, will see these certificates in his or her personal certificate store.

For roaming profiles, the users certificates are located on the share that hosts the user's profile. Because user certificates and private keys are part of the user's profile, this information is transmitted over the network when a roaming profile is read from or written to the server share.

**Certificate Processes and Interactions**
Certificates are credentials, and like other forms of credentials they need to be created before they can be used. After they have been created, there needs to be a means of validating their authenticity and verifying that they can be used for the desired purpose. After certificates have been created and validation processes are in place, certificates can be used for the following basic cryptographic operations:

- Public key encryption
- Message digest functions
- Digital signing

**How Certificates Are Created**
Certificates are issued by a CA, which can be any trusted service or entity willing to verify and validate the identities of those to whom it issues certificates and their association with specific keys. Companies might issue certificates to employees, schools might issue certificates to students, and so on. Of course, a CA's public key must be trustworthy or the certificates it issues will not be trusted. Because anyone can become a CA, certificates are only as trustworthy as the authority that issues the underlying keys.

The following six steps describe the process of requesting and issuing a certificate.

1. Generate a key pair The applicant generates a public and private key pair or is assigned a key pair by some authority in his or her organization.

2. Collect required information The applicant collects whatever information the CA requires to issue a certificate. The information could include the applicant's e-mail address, birth certificate, fingerprints, notarized documents — whatever the CA needs to be certain that the applicant is who he or she claims to be. CAs with stringent identification requirements produce certificates with high assurance — that is, their certificates generate a high level of confidence. CAs themselves are said to be of high, medium, or low assurance.

3. Request the certificate The applicant sends a certificate request, consisting of his or her public key and the additional required information, to the CA. The certificate request, which is signed with the client's public key, can also be encrypted by using the CA's public key. Many requests are made by using e-mail, but requests can also be sent by postal or courier service — for example, when the certificate request itself must be notarized.

4. Verify the information The CA applies whatever policy rules it requires to verify that the applicant should receive a certificate. As with identification requirements, a CA's verification policy and procedures influence the amount of confidence generated by the certificates it issues.

5. Create the certificate The CA creates and signs a digital document containing the applicant's public key and other appropriate information. The signature of the CA authenticates the binding of the subject's name to the subject's public key. The signed document is the certificate.

6. Send or post the certificate The CA sends the certificate to the applicant or posts the certificate in a directory, as appropriate.

**How Certificates Are Validated**

Before it trusts certificates, Windows performs a validation check to ensure that certificates are valid and that they have a valid certification path.

The status of a public key certificate is determined through three distinct, but interrelated, processes implemented in CryptoAPI:

- Certificate discoveryThe process of collecting CA certificates from Group Policy, Enterprise Policy, and AIA pointers in issued certificates, and the process of enrolling certificates.

- Path validationThe process by which public key certificates and their issuer certificates are processed in a hierarchical fashion until the certificate chain terminates at a trusted, self-signed certificate. Typically, this is a root CA certificate.

- Revocation checkingEach certificate in the certificate chain is verified to ensure that none of the certificates are revoked. The revocation checking can take place either in conjunction with the chain building process or after the chain is built.

**Chain Building**

Chain building is the process of building a trust chain, or certification path, from the end certificate to a root CA that is trusted by the security principal.

The chain-building process will validate the certification path by checking each certificate in the certification path from the end certificate to the root CA's certificate. The certificates are retrieved from the Intermediate Certification Authorities store, the Trusted Root Certification Authorities store, or from a URL specified in the AIA attribute of the certificate. If CryptoAPI discovers a problem with one of the certificates in the path, or if it cannot find a certificate, the certification path is discarded as a nontrusted certification path.

To improve performance, CryptoAPI will store subordinate CA certificates in the Intermediate Certification Authorities store so that future requests for the certificate can be satisfied from the store, rather than accessing the certificate through a URL.

**Certificate Storage**

Windows Server 2003 and Windows XP store certificates locally on the computer or hardware device that requested the certificate, or — in the case of a user — on the computer or device that the user used to request the certificate. There are

separate stores known as the machine store, used by the computer, and the user store or My store used by the currently logged-on user. A certificate store will often contain numerous certificates, possibly issued from a number of different CAs.

Certificate chains are formed by looking at certificates available in multiple certificate stores. The certificate chain engine must determine what scope of certificate stores to search when building certificate chains. By default, the chain engine searches the ROOT, TRUST, CA, and MY certificate stores.

In addition to the default stores, the certificate chain engine can be configured to use different stores, such as restricted root, restricted trust, restricted other, and additional stores. For example, the Trusted People store is used by both Outlook and EFS when searching for certificates. Alternatively, an application can create its own store for certificate storage or even call additional revocation providers registered with CryptoAPI.

**Purpose**
The certificate chain engine builds all possible certificate chains. The entire graph of certificate chains is constructed and then ordered by the "quality" of the chain. The best-quality chain for a given end certificate is returned to the calling application as the default chain.

Each chain is built by using a combination of the certificates available in the certificate stores and certificates available from published URL locations. Each certificate in the chain is assigned a status code. The status code indicates whether the individual certificate is:

- Signature-valid Is the signature valid?
- Time-valid Are the certificate start and expiration dates properly configured, has the start date not occurred yet, or has the certificate expired?
- Expired Has the certificate expired?
- Revoked Has the certificate been revoked?
- Time-nested Have any of the certificates that are higher in the PKI hierarchy expired?
- Any other restrictions on the certificate For example, is the certificate being used for a purpose other than has been intended?

Each status code has a precedence assigned to it. For example, an expired certificate has a higher precedence than a revoked certificate. This is because an expired certificate should not be checked for revocation status.

If different status codes are assigned to the certificates in a certificate chain, the status code with the highest precedence is applied to the certificate chain and propagated into the certificate chain status.

**Path validation**
For each certificate in the chain, the certificate chain engine must select a certificate of the issuing CA. This process, known as path validation, is repeated until a self-signed certificate is reached (typically, this is a root CA certificate). CryptoAPI treats root certificates as the absolute determinant of trust in trust decisions.

There are different processes that can be used to select the certificate for an issuing CA. The actual process that is used is based on whether the certificate currently being investigated has the Authority Key Identifier (AKI) extension defined. Inspection of the AKI extension will lead to one of three matching processes being implemented:

- Exact match If the AKI extension contains the issuer's user name and issuer serial number, only certificates that match on user name and serial number will be chosen in the chain-building process. As a further test, the issuer name on the issued certificate must match the subject name on the issuer certificate.
- Key match If the AKI extension only contains public key information, only certificates that contain the indicated public key in the Subject Key Identifier (SKI) extension will be chosen as valid issuers.

  **Note**

  - The public key information in the AKI extension and in the SKI extension is the hash of the public key. Therefore, both hashes must have been calculated by using the same algorithm. Otherwise, no key match will be

determined even though the public key used for the hashes matches. The default hash algorithm used by the Microsoft CA and CryptoAPI is SHA-1 when no SKI exists in the certificate. However, when an SKI exists, both MD5 and SHA-1 algorithms are supported, but the AKI and SKI must use the same algorithm.

- Name match. If no information exists in the AKI, or if the AKI does not exist in the certificate, the certificate will be marked as "name match." In name matching, the subject name of a certificate must match the issuer name in the current certificate in order for the certificate to be chosen as a valid issuer. Because the data is stored in a binary format, the name-matching process is case-sensitive.

In all cases, even if a matching certificate is not found in the store, the current certificate will still be marked as "exact match", "key match" or "name match," because this describes the attempted match rather than the attained match.

**Caching**

To increase performance, the certificate chain engine uses a least recently used (LRU) caching scheme. This scheme creates a cache entry for each certificate it encounters as it builds the certificate chain. Each cache entry includes the status of the certificate, so the best certificate chain can be built from cached items on subsequent calls to the chaining API without having to determine the status of each certificate again. After a certificate has been added to the cache, it will not be removed until it expires or is revoked.

During the path validation process, valid cached certificates will always be selected. If valid cached certificates are not found, a store search will be performed. For issuer certificates and CRLs, URL retrieval can be required to download the certificates and CRLs from the distribution point indicated in the URL.

All certificates are stored in the cache when the certificates are selected from a store or from a URL. The only difference is the location where the cached certificates are stored. Certificates can be stored in the following locations:

- CA store All certificates retrieved from any WinHTTP-supported URLs via the AIA extension are cached in the CA store.

- Local file system If the retrieval URL begins with LDAP, FTP, or HTTP, the certificate (or CRL) is also cached in the local file system.

In some cases, the certificate might be cached in all three locations. For example, when a certificate is retrieved from an HTTP URL, the certificate will be cached in memory, the CA store, and in the local file system.

**Revocation**

The certificate chain engine performs basic revocation checking during chain building, but the process differs between Windows 2000 and Windows 2003 and Windows XP.

For Windows 2000, post-processing revocation is used. This means that the CRL checking is performed after the chain is built. The certificate chain engine will check each certificate in the chain to determine whether the certificate has been revoked and the reason for the revocation.

Windows XP and Windows Server 2003 use revocation checking while building the certificate chain. Rather than waiting until the chain is built, each certificate is examined as the chain is built. If a revoked certificate is discovered in the chain, the chain is assigned a lower quality value.

### Note

- The existence of a revoked certificate in a certificate chain does not preclude the chain from being presented to the calling application as the best-quality certificate chain. The best-quality chain might not necessarily be a trustworthy chain.

**Certificate path validation**

The path validation process ensures that a valid certification path can be established for a given end certificate. A valid certification path is defined as an end-user certificate that can be traced along a certificate chain to a trusted root CA. Each CA certificate in the path must be discovered and subsequently validated until a final authority such as a root CA is obtained. During the validation process, a certificate can be deemed invalid (or not trusted) for many reasons. The reasons include:

- The time period configured in the certificate is not valid. This occurs when the start and expiration dates are improper, have not occurred yet, or are expired.

  **Note**

    - An expired CA certificate in the certification path reduces the quality of the path, but it does not invalidate the path. In a Windows Server 2003 PKI, a certification path can be valid as long as the CA certificate was valid at the time the certificate was issued. For example, a third-party CA might issue a certificate with a lifetime that extends past the CA's certificate's expiration date. After the CA's certificate expires, the certification path for the certificate is still valid and the certificate is trusted as long as all other validation criteria are met. In Windows 2000, however, CryptoAPI enforces time nesting rules by default where a child certificate must have a validity period shorter than the parent. Signatures, however, might be considered valid past the date of the certificate lifetime.

- The certificate is unrecognized. If the certificate format is improper, does not conform to the X.509 version 1 through version 3 standard for digital certificates, the certificate is discarded.

- The information in certificate fields is improper or incomplete.

- The certificate's digital thumbprint and signature fail the integrity check, indicating that the certificate has been tampered with or corrupted.

- The certificate's status was determined to be revoked.

- The issuing CA is not in either a trusted certification hierarchy or a CTL.

- The root CA for the certification path is not in the Trusted Root Certification Authorities store.

- The certificate is not permitted for the intended use as specified in a CTL.

- The certificate contains a critical extension that is not understood by the application.

**Critical extensions**
The CryptoAPI engine does not enforce critical extensions in certificates, only CRLs. The CryptoAPI engine and programming model place the burden of parsing critical extensions on the calling application; therefore, an application must understand and enforce the critical extension when evaluating a certificate. The CryptoAPI revocation provider, however, does evaluate critical extensions in CRLs and therefore would reject a CRL with a critical extension that is not known or understood.

**Certificate status checking**
All certificates in a certificate chain can be processed to verify that none of the certificates has been revoked. Certificate chain validation is, of course, optional from an application standpoint and might not be enforced by CryptoAPI. Windows by default checks certificate revocation status via CRLs, because the CRL processing engine is the native revocation provider included with CryptoAPI. When this functionality has been invoked, each certificate in the certificate chain is checked against the CRL published in the CRL Distribution Point (CDP) extension in the certificate. If the certificate is found to be included in the CRL, the certificate is then considered revoked. Additionally, third-party revocation providers can be registered with CryptoAPI to add support for additional protocols that check revocation status including Online Certificate Status Protocol (OCSP), Simple Certificate Validation Protocol (SCVP), and XML Key Management Specification (XKMS).

When a certificate's status is verified by using a CRL, several steps must be performed by an application to check the status of the certificates in the certificate chain. These steps are performed against each certificate in the chain. These steps include:

1. Verify the signature of the CRL. The CRL must be signed so that the application can determine whether it trusts the CRL issuer to issue CRLs.

   **Note**

   - Windows Server 2003, Windows 2000, and Windows XP can only verify a CRL that was signed by the same private key used to sign the issued certificate. These versions of Windows do not support CRLs signed by an entity other than the CA that signed the issued certificate.

2. Verify that the CRL has not expired. A CRL is considered expired if the current date is later than the date contained in the next update field of the CRL. If the certificate that is being checked has expired, the application must verify that the CRL's

issue date follows the effective date of the certificate, but precedes the certificate's expiration date.

3. Search the list of revoked certificates to verify that either the target certificate is not included or its revocation date is later than the effective date.

When a third-party revocation provider supporting OCSP has been registered, an OCSP responder will be used for certificate status checking; in this case, the process is slightly modified.

1. Verify that the response indicates the certificate is valid. A valid response indicates that the certificate has not been revoked.

2. Verify the signature on the OCSP response. This activity includes developing and processing a path that establishes that the certificate issuer or a trust point trusted the responder for the express purpose of issuing responses.

Neither Windows XP nor Windows 2000 contains an OCSP client component by default. However, a third-party OCSP client can be installed as a revocation provider to CryptoAPI. OCSP responders can be located by using the AIA extension in the certificate as defined by RFC 2459. The Windows Server 2003 CA supports the OCSP responder location to be included in the AIA extension of certificates. Multiple revocation providers can be added to CryptoAPI depending on revocation requirements. For additional information about revocation providers, see the Platform SDK on MSDN.

No matter what process is used to verify certificate validity, if the status check fails any of the above checks for any certificate in the certificate chain, the certificate chain will be rejected.

**Basic constraint validation**
Basic constraints give an administrator the ability to designate whether an issued certificate is a CA certificate or an end certificate. This validation is used to verify that a CA certificate can be used by the certificate chain engine to build certification paths.

An additional use of the basic constraint extension is to limit the maximum number of CA certificates that can be included under a given CA. For example, a path length of zero only allows end certificates to be issued by that specific CA. Likewise; a path length of two in the basic constraints extension will only allow three CA certificates in a certification path. In this example, any certification paths discovered with more than three CAs in the path will be discarded.

**Name constraint validation**
A CA certificate can contain name constraints that are applied to all certificate requests made to the CA. Each request is compared to the list of permitted and excluded constraints to determine whether the certificate should be considered permitted, not permitted, excluded, or not defined.

### Note

- Name constraint validation can only be performed by Windows XP and Windows Server 2003 clients. Name constraints are not evaluated by Windows 2000 clients. If you require that name constraints be applied, you can indicate that the extensions are critical, which should result in the chain being discarded by an application conforming to RFC 2459.

For example, a permitted constraint could allow all DNS names that end in contoso.com. This would include DNS names such as contoso.com and xcontoso.com. If you only wanted DNS names from the contoso.com DNS name space, you could use the permitted constraint .contoso.com. This constraint would permit x.contoso.com but exclude xcontoso.com.

When name constraints are present in a CA certificate, the following rules are applied to the subject name and alternate subject name entries.

- If the name constraints extension exists in a CA certificate, all name constraints should be present in the extension. Any name constraints that are not included are considered wildcards that will match all possibilities. For example, if the DNS name constraint were absent, the entry would be treated as DNS=.

- All name constraints will be considered. There is no precedence applied to the listed name constraints. It is for this reason that name constraints that are not present are treated as wildcards.

- An excluded name constraint will take precedence over a permitted name constraint

- Name constraints are applied to the subject name extension and any existing subject alternate name extensions.

- Name constraints apply to all names contained in an end certificate. Each name in the subject or subject alternate name extensions should match at least one of the name constraints listed for that name type. A subject name or subject alternate name that does not match a listed name type will be rejected. Note that most client name spaces are not included in a CA certificate and generally do not apply.

- Name constraints are case-sensitive if the names are stored in ASCII or Unicode format.

Name restrictions must be enforced across the following alternative name information entries in the subject name: Other Name (NT Principal Name only); RFC 822 Name; DNS Name; URL; Directory Name, and IP address.

When the certificate chain engine validates an end certificate for name constraints, it will arrive at one of the following results:

- **Permitted** The end certificate contains a name that is listed as permitted in an issuer's name constraints extension.

- **Not permitted** The end certificate contains a name that is not listed as permitted in an issuer's name constraints extension.

- **Excluded** The end certificate contains a name that is listed as excluded in an issuer's name constraints extension

- **Not Defined** The issuer certificate does not list a constraint for a specific name type (such as Directory Name or IP Address)

**Policy constraint validation**

A policy constraint allows a CA administrator to ensure that specific constraints are met when a certificate is issued or used by an application. The policy constraint ensures that all certificates issued by the CA implement the required policy constraints.

By definition, a root CA implements all policies. This applies to both enterprise and standalone CAs. At some point down the hierarchy, a CA can have one or more policies defined. After a CA certificate is encountered with any policy OIDs, all certificates below that CA in the hierarchy must also have a subset of those policy OIDs. A certificate chain with no valid policy set will be considered invalid, whereas one with no policy OIDs at all will be considered valid and matching the "any policy" OID. The above is valid only for application policies and not issuance polices. The absence of the certificatePolicies extension in a CA other than the root CA implies that there is no issuance policy.

There are two policies currently used with a Windows Server 2003 CA: issuance policy and application policy. The policy applied to a certificate affects how the validity of a certificate chain is determined.

**Issuance policy**

Issuance policies, referred to as the certificate policies extension in RFC 2459, allow a CA administrator to define the circumstances and requirements for certificate issuance. The assurances required for certificate issuance can vary based on certificate templates. An issuance policy defines a set of administrative rules that must be followed when issuing a certificate. The rules are defined in the certificate by an OID defined by the CA. Each certificate issued by the CA will include the OID. The OID is not defined in the CA, but instead in the certificate template.

### Note

- Issuance policy is only available to Windows Server 2003 and Windows XP clients. Issuance policy is not recognized by Windows 2000 clients.

When an entity presents a certificate to an application, the application can examine the certificate to verify the issuance policy and determine if the issuance policy is sufficient to perform the action requested by the subject.

Currently, two types of constraints are defined: **Require explicit policy** and **Inhibit policy mapping**.

- **Require explicit policy** specifies the number of certificates that can exist in the hierarchy below the current certificate before an explicit policy must exist.

- **Inhibit policy mapping** specifies the number of additional certificates that can appear in the path before policy mapping is no longer permitted.

Policies can also be mapped to other policies on a one-to-many basis. Policy mapping allows interoperability between two organizations that implement similar policies, but have deployed different OIDs. If one company's policy OID (for example, 1.2.3.4) stands for a specific function and another company's policy OID (for example, 11.22.33.44) stands for the same function, these two OIDs can be mapped. In this example, 11.22.33.44 can be mapped to 1.2.3.4.

**Application policy**

Certificates provide key information that is not specific to an application; however, the ability to decide which certificates can be used for certain functions is important. Application policy allows you to issue certificates widely and restrict their usage to only their intended purposes.

Application policies are settings that inform a target that the subject holds a certificate that can or cannot be used to perform a specific task. They are represented in a certificate by an OID that is defined by the CA. This OID is included in all issued certificates. When a subject presents its certificate, it can be examined by the target to verify the application policy and determine if it can perform the requested action.

Application policy is Microsoft-specific and is treated much like EKU

- If a certificate has an extension containing an application policy and has an EKU extension, the EKU extension is ignored.
- If there is only an EKU extension, it is treated like an application policy extension.

If there is an application policy extension (and no EKU extension) and an EKU property on the certificate, the intersection of the EKU property OIDs and the application policy OIDs is taken and the result is the effective policy for the certificate.

If the application policy extension is absent, CryptoAPI will function like any other RFC 2459–compliant client. If it is present, CryptoAPI will implement the application policy rules. It is recommended that implementers make the extension non-critical extension so that its presence should be benign to other clients.

**Note**

- Policy mapping can only be used with the application policy extension. If the EKU extension is used, policy mapping is not possible because EKU does not support policy mapping. Non-Windows clients and applications might not understand this extension or use it as designed.

**Basic Cryptographic Operations**

A certificate contains information that identifies the certificate's owner (called the subject) as an entity on the network. A certificate also contains the owner's public key. Furthermore, a certificate identifies the CA (called the issuer) that issued the certificate.

A CA uses its private key to digitally sign each certificate it issues. To create the digital signature, the CA generates a message digest from the certificate, encrypts the digest with its private key, and includes the digital signature as part of the certificate. Anyone can use the message digest function and the CA's public key to verify the certificate's integrity. If a certificate becomes corrupted or someone tampers with it, the message digest for the altered certificate does not match the digest in the CA's digital signature. The following figure shows how a certificate is signed by the issuing CA.

**Digital Signature for a Certificate**




Information contained in the Subject Public-Key Value field of X.509 version 3 certificates specifies the cryptography operations for which the public key and private key set can be used. Public key security systems commonly support the following basic cryptography operations:

- Digital signing of electronic data to verify data origin and the integrity of data.
- Secret key encryption to protect symmetric secret encryption transmitted and shared over networks.

The following table briefly describes the certificate purposes.

**Basic Cryptographic Operations**

| Cryptographic Operations | Intended use |
|---|---|
| Signature | Data signing, authentication, nonrepudiation |
| Encryption | Data encryption and decryption |
| Signature and encryption | Data encryption and decryption, digital data signing, authentication |
| Signature and smart card logon | Smart card logon, digital data signing |

The public key and private key set can be used to provide a variety of specific security functions for information security technologies. These specific functions of certificates are listed in the keyUsage and EKU extensions. Common specific security functions for public key technology include the following:

- Secure mail to provide authentication, confidentiality, integrity, and nonrepudiation for e-mail communications.
- Secure Web communications to provide authentication, integrity, and confidentiality between Web clients and servers.
- Code signing to provide integrity and nonrepudiation for executable code to be distributed on the Internet or intranets.
- Local network logon or remote access logon to authenticate users of network resources.
- IPSec authentication to authenticate clients that do not use Kerberos authentication or shared secret passwords for IPSec communications.

**Public Key Encryption**

In public key encryption, different keys are used to encrypt and decrypt information. The first key is a private key (a key that is known only to its owner), while the second key — called the public key — can be made known and available to other entities on the network.

The two keys are different but complementary in function. For example, a user's public key can be published in a certificate in a directory so that it is accessible to other people in the organization. The sender retrieves the recipient's certificate from Active Directory, retrieves the public key from the certificate, and then encrypts their communication by using this public key. Information that is encrypted with the public key can be decrypted only by using the corresponding private key of the set, which remains with its owner. The following figure shows basic encryption and decryption with asymmetric keys.

**Encryption and Decryption with Asymmetric Keys**




Public key encryption plays an increasingly important role in providing strong, scalable security on intranets and the Internet. Public key encryption is commonly used to perform the following functions:

- Encrypt symmetric secret keys to protect the symmetric keys during exchange over the network or while being used, stored, or cached by operating systems.
- Create digital signatures to provide authentication and nonrepudiation for online entities.
- Create digital signatures to provide data integrity for electronic files and documents.

The RSA digital signature process also uses private keys to encrypt information to form digital signatures. For RSA digital signatures, only the public key can decrypt information encrypted by the corresponding private key of the set.

**Message Digest Functions**

Message digest functions, also called hash functions, are frequently used in conjunction with asymmetric keys to further strengthen public key encryption. Message digests are commonly 128 bits to 160 bits in length and provide a unique digital identifier for each digital file or document. Two copies of a document will have the same message digest, but if even one of the bits for the document changes, the message digest changes. The following figure shows the basic message digest process.

**Example of the Message Digest Process**




Message digests are commonly used in conjunction with public key technology to create digital signatures or "digital thumbprints" that are used for authentication, integrity, and nonrepudiation. Message digests also are commonly used with digital signing technology to provide data integrity for electronic files and documents.

For example, to provide data integrity for e-mail messages, message digests can be generated from the completed mail message, digitally signed with the originator's private key, and then transmitted with the messages. The recipient of the message can then do the following to check the integrity of the message:

- Use the same message digest function to compute a digest for the message.
- Use the originator's public key to verify the signed message digest.
- Compare the new message digest to the original digest.

If the two message digests do not match, the recipient knows the message was altered or corrupted. The following figure shows a basic integrity check process with a digitally signed message digest.

**Example of an Integrity Check with a Digitally Signed Message Digest**




Because the message digest is digitally signed with the sender's private key, it is not feasible for an intruder to intercept the message, modify it, and create a new valid encrypted message digest to send to the recipient.

**Digital Signatures**

A common use of public key encryption is to provide digital signatures. Just as handwritten signatures or physical thumbprints

are commonly used to uniquely identify people for legal proceedings or transactions, so digital signatures are commonly used to identify electronic entities for online transactions. A digital signature uniquely identifies the originator of digitally signed data and also ensures the integrity of the signed data against tampering or corruption.

One possible method for creating a digital signature is for the originator of data to create the signature by encrypting all of the data with the originator's private key and enclosing the signature with the original data. Anyone with the originator's public key can decrypt the signature and compare the decrypted message to the original message. Because only someone with the private key can create the signature, the integrity of the message is verified when the decrypted message matches the original. Even if an intruder intercepts and alters the original message while it is in transit, the intruder cannot create a new valid signature. If an intruder alters the signature during transit, the signature will not be properly verified and therefore will be invalid.

However, encrypting all data to provide a digital signature is impractical for three reasons:

- The cipher text signature is the same size as or greater than the corresponding plaintext, so message sizes are doubled, which consumes large amounts of bandwidth and storage space.

- Public key encryption places heavy computational loads on computer processors, so network and computer performance can be significantly degraded.

- Encrypting the entire contents of information produces large amounts of ciphertext, which can be used for cryptanalysis attacks, especially known plaintext attacks (where certain parts of the encrypted data, such as e-mail headers, are known beforehand to the attacker).

Digital signature algorithms use more efficient methods to create digital signatures. The most common types of digital signatures today are created by signing message digests with the originator's private key to create a digital thumbprint of the data. Because only the message digest is signed, the signature is usually much shorter than the data that was signed. Therefore, digital signatures place a relatively low load on computer processors during the signing process, consume insignificant amounts of bandwidth, and produce small amounts of ciphertext for cryptanalysis.

**Related Information**
The following resources contain additional information that is relevant to this section:

- Logon and Authentication Technologies

# Certificates Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate Tools and Settings**

**In this section**

- Certificate Tools
- Certificate Registry Entries
- Certificate Group Policy Settings

**Certificate Tools**

The following tools are associated with certificates.

**Certmgr.msc: Certificates Snap-in**

**Category**

Certificates is a Microsoft Management Console (MMC) snap-in that ships with Windows Server 2003 and Windows 2000 Server.

**Version compatibility**

Certificates is compatible with Windows Server 2003 and Windows 2000 Server, and can be used to manage the certificate stores for users, computers, and services on computers running Windows Server 2003, Windows XP, and Windows 2000.

You can use Certificates to:

- View information about certificates, such as certificate contents and the certification path.
- Import certificates into a certificate store.
- Move certificates between certificate stores.
- Export certificates and, optionally, export private keys (if key export is enabled).
- Delete certificates from certificate stores.
- Request certificates from an enterprise certification authority (CA) for the Personal certificate store.

  **Note**

  - The ability to perform some tasks will depend on the capabilities of the PKI configuration and environment.

To find more information about the “Certificates” snap-in, see “Certificates” on Microsoft TechNet.

**Certtempl.msc: Certificate Templates Snap-in**

**Category**

Certificate Templates is an MMC snap-in that ships with Windows Server 2003.

**Version compatibility**

Certificate Templates only runs on Windows XP and Windows Server 2003, but can be used to manage certificate templates in a Windows 2000 Active Directory environment.

Certificate Templates enables administrators to duplicate, rename, and manage certificate templates. In Windows Server 2003, Certificate Templates also enables you to modify existing certificate template properties — such as certificate validity period, renewal period, cryptographic service providers (CSPs), key size, and key archival.

In addition, administrators or users with the appropriate permissions can use Certificate Templates to establish and apply enrollment policies (including autoenrollment), issuance policies, and application policies.

For more information about the Certificate Templates snap-in, see "Certificate Templates How to" on Microsoft TechNet.

**Certreq.exe: Certreq**

**Category**
Certreq is a command-line tool that ships with the Windows Server 2003 operating system tools and with the Windows Server 2003 Adminpak

**Version compatibility**
Certreq is compatible with Windows Server 2003 and Windows 2000 Server, and can be used to manage the certificate stores for users, computers, and services on computers running Windows Server 2003, Windows XP, and Windows 2000.

Certreq enables you to submit, retrieve, create, and accept certificate requests that are sent to a Windows Server 2003 CA. You can also use Certreq to create and sign requests for cross-certificates. You can also place the Certreq command syntax in a batch file to script certificate requests.

To find more information about Certreq, see "Command-Line References" in Tools and Settings Collection.

**Certutil.exe: Certutil**

**Category**
Certutil is a command-line tool that is installed as part of Certificate Services.

**Version compatibility**
Certutil can be used on Windows Server 2003 and Windows 2000 Server CAs.

You can use Certutil to extract and display CA configuration information, configure Certificate Services, back up and restore CA components, and verify certificates, key pairs, and certificate chains.

To find more information about Certutil, see "Command-Line References" in Tools and Settings Collection.

**Certificate Registry Entries**
Certificate-related registry entries correlate to the physical view of the certificate-related data that can be viewed by using the Certificates snap-in.

The following registry keys are associated with certificates:

- HKEY_Current_User\Software\Microsoft contains data about user certificates that have not been distributed by using Group Policy.
- HKEY_Current_User\Software\Policies\Microsoft contains data and settings for user certificates that have been distributed by using Group Policy.
- HKEY_Local_Machine\Software\Microsoft contains data about computer certificates that have not been distributed using Group Policy.
- HKEY_Local_Machine\Software\Policies\Microsoft contains data and settings for computer certificates that have been distributed by using Group Policy.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as MMC, to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**HKEY_Current_User\Software\Microsoft**
HKEY_Current_User\Software\Microsoft contains registry settings for user certificates that have been distributed by means other than Group Policy. These settings are stored in the following subkeys:

- HKEY_Current_User\Software\Microsoft\Cryptography

- HKEY_Current_User\Software\Microsoft\SystemCertificates

**HKEY_Current_User\Software\Microsoft\Cryptography**
The following registry entries are located under HKEY_Current_User\Software\Microsoft\Cryptography.

**Autoenrollment**

**Registry path**
HKEY_Current_User\Software\Microsoft\Cryptography\

**Version**
Windows Server 2003, Windows 2000, and Windows XP

This setting is used to manage event logging and cached directory service data when user certificate autoenrollment has been enabled.

**AEExpress**

**Registry path**
HKEY_Current_User\Software\Microsoft\Cryptography\Autoenrollment

**Version**
Windows Server 2003 and Windows XP

You can add AEExpress on a per-user basis if the default 60-second delay is not desired. With this registry setting, the autoenrollment balloon UI appears at each logon or Group Policy refresh interval.

### Note

Using this subkey in a normal production environment is not recommended. If it is used, it must be created on a per-user basis. Computer certificates do not support user interaction and should not be configured to require this setting.

**HKEY_Current_User\Software\Microsoft\SystemCertificates**
The following registry subkeys are located under SystemCertificates. The majority contain binary large objects that pertain to:

- **Certificates.** These entries identify the certificates associated with the registry entry.
- **CRLs.** These entries identify the certificate revocation lists (CRLs) associated with the registry entry.
- **CTLs.** These entries identify the certificate trust lists (CTLs) associated with the registry entry.

Additional subkeys — which might appear under some registry subkeys — will be detailed below under the registry subkeys that they correspond to.

**ACRS**

**Registry path**
HKEY_Current_User\Software\Microsoft\SystemCertificates

**Version**
Windows 2000 and Windows XP

Automatic Certificate Request Settings (ACRS) in Group Policy is used by Windows 2000 to specify certificates to autoenroll. This registry setting has been replaced in Windows XP and Windows Server 2003 by certificate template autoenrollment. ACRS is still available for backward compatibility, but is not supported for user certificates.

**AddressBook**

**Registry path**
HKEY_Current_User\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

AddressBook contains data about certificates, CRLs, and CTLs relating to people. Unlike certificates registered in the

TrustedPeople container, these certificates are not explicitly trusted.

**AuthRoot**

**Registry path**
HKEY_Current_User\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

AuthRoot contains data about certificates, CRLs, and CTLs from non-Microsoft root CAs.

**CA**

**Registry path**
HKEY_Current_User\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

CA contains data about certificates, CRLs, and CTLs from intermediate CAs.

**Disallowed**

**Registry path**
HKEY_Current_User\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Disallowed contains data about disallowed certificates. For example, this store will contain data about certificates that have been rejected as untrustworthy during the Authenticode software installation process.

**My**

**Registry path**
HKEY_Current_User\Software\ Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

My contains data about a user's personal certificates. However, physical storage of this data for the user, which is viewable by using the Certificates snap-in, has been moved to the following Documents and Settings folder:

Application Data\Microsoft\SystemCertificates

**REQUEST**

**Registry path**
HKEY_Current_User\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows 2000, and Windows XP

REQUEST contains data about pending certificate requests. When a pending certificate request has been approved or rejected, this registry data is removed. Data about pending certificate requests is also stored for the user in the following Documents and Settings folder:

Application Data\Microsoft\SystemCertificates\My

**Root**

**Registry path**
HKEY_Current_User\Software\ Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Root contains data about trusted root CA certificates, CRLs, and CTLs.

**trust**

**Registry path**
HKEY_Current_User\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Trust contains data about enterprise trust certificates, CRLs, and CTLs.

**TrustedPeople**

**Registry path**
HKEY_Current_User\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

TrustedPeople contains data about certificates from other users where there is explicit trust.

**TrustedPublisher**

**Registry path**
HKEY_Current_User\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

TrustedPublisher contains data about certificates that have been explicitly accepted as trustworthy. For example, it will contain data about certificates that are accepted as trustworthy during the Authenticode software installation process.

**HKEY_Current_User\Software\Policies\Microsoft**
HKEY_Current_User\Software\Policies\Microsoft contains registry settings for user certificates that have been distributed by using Group Policy. These settings are stored in the following subkeys:

- HKEY_Current_User\Software\Policies\Microsoft\Cryptography
- HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**HKEY_Current_User\Software\Policies\Microsoft\Cryptography**
This set of registry entries is used when Group Policy-based autoenrollment is used to manage the availability of user certificates.

**Autoenrollment**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\Cryptography

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Autoenrollment implements Group Policy for autoenrollment of user certificates.

**AEEventLogLevel**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\Cryptography\Autoenrollment\

**Version**
Windows Server 2003 and Windows XP

AEEventLogLevel enables enhanced logging of autoenrollment processes for users.

**HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates**
The following registry entries are located under SystemCertificates. The majority contain binary large objects that pertain to:

- **Certificates.** These entries identify the certificates associated with the registry entry.
- **CRLs.** These entries identify the CRLs associated with the registry entry.
- **CTLs.** These entries identify the CTLs associated with the registry entry.

Additional subkeys — which might appear under some registry subkeys — will be detailed below under the registry subkeys that they correspond to.

**AddressBook**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

AddressBook contains data about certificates, CRLs, and CTLs relating to other people. Unlike certificates registered in the TrustedPeople container, these certificates are not explicitly trusted.

**AuthRoot**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

AuthRoot contains data about certificates, CRLs, and CTLs from third-party root CAs.

**CA**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

CA contains data about certificates, CRLs, and CTLs from intermediate CAs.

**Disallowed**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Disallowed contains data about disallowed certificates. For example, it will contain data about certificates that are rejected as untrustworthy during the Authenticode software installation process

**EFS**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

EFS contains data about certificates enabling the encrypting file system for the user.

**root**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Root contains data about trusted root CA certificates, CRLs, and CTLs.

**trust**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Trust contains data about enterprise trust certificates, CRLs, and CTLs

**TrustedPeople**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

TrustedPeople contains data about certificates from other users where there is explicit trust.

**TrustedPublisher**

**Registry path**
HKEY_Current_User\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

TrustedPublisher contains data about certificates that have been explicitly accepted as trustworthy. For example, this store will contain data about certificates that are accepted as trustworthy during the Authenticode software installation process.

**HKEY_Local_Machine\Software\Microsoft**
HKEY_Local_Machine\Software\Microsoft contains registry settings for computer certificates that have been distributed by means other than Group Policy. These settings are stored in the following subkeys:

- HKEY_Local_Machine\Software\Microsoft\Cryptography
- HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**HKEY_Local_Machine\Software\Microsoft\Cryptography**

**AutoEnrollment**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\Cryptography

**Version**
Windows Server 2003 and Windows XP

This setting is used to manage event logging level and cached directory service data when computer certificate autoenrollment has been enabled.

**Calais**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\Cryptography

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Calais contains settings that enable various vendors' smart cards and smart card readers.

**Defaults**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\Cryptography

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Contains information about different specific and generic CSPs listed in the Provider and Provider Types categories.

**IEDirtyFlags**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\Cryptography

**Version**
Windows versions prior to Windows 2000

This legacy setting was used to record certificates and keys that are incompletely removed from Internet Explorer.

**OID**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\Cryptography

**Version**
Windows Server 2003, Windows XP, and Windows 2000

OID is the hive used by the CryptRegisterOIDFunction, CryptRegisterDefaultOIDFunction, and CryptRegisterOIDInfo application programming interfaces (APIs) to register data.

**Protect**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\Cryptography

**Version**
Windows Server 2003, Windows XP, and Windows 2000

You can use Protect to define providers for use with the ProtectedStorage service.

**RNG**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\Cryptography

**Version**
Windows Server 2003, Windows XP, and Windows 2000

RNG is used by the Windows Random Number Generator to enhance random number generation from one computer startup to another.

**Services**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\Cryptography

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Services identifies the predefined physical stores for certificates associated with Windows services.

**HKEY_Local_Machine\Software\Microsoft\SystemCertificates**

**AddressBook**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

AddressBook contains data about certificates, CRLs, and CTLs relating to other people. Unlike certificates registered in the TrustedPeople container, these certificates are not explicitly trusted

**AuthRoot**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

AuthRoot contains data about certificates, CRLs, and CTLs from non-Microsoft root CAs.

**CA**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

CA contains data about certificates, CRLs, and CTLs from intermediate CAs.

**Disallowed**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Disallowed contains data about disallowed certificates. For example, it will contain data about certificates that are rejected as untrustworthy during the Authenticode software installation process.

**My**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

My contains data about the certificates, CRLs, CTLs, and private keys associated with a computer account.

**Recovery**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003 and Windows XP

Recovery stores the certificate for a user's password reset disk.

**REQUEST**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

REQUEST contains data about pending certificate requests. When a pending certificate request has been approved or rejected, this registry information is removed.

**ROOT**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

ROOT contains data about trusted root CA certificates, CRLs, and CTLs.

**SPC**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Not supported

SPC is a legacy registry store that was used by Internet Explorer 3.0 for migrating Software Publisher Certificates. It is no longer used.

**trust**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Trust contains data about enterprise trust certificates, CRLs, and CTLs.

**TrustedPeople**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

TrustedPeople contains data about certificates from other users where there is explicit trust.

**TrustedPublisher**

**Registry path**
HKEY_Local_Machine\Software\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

TrustedPublisher contains data about certificates that have been explicitly accepted as trustworthy. For example, this store will contain data about certificates that are accepted as trustworthy during the Authenticode software installation process.

**HKEY_Local_Machine\Software\Policies\Microsoft**
HKEY_Local_Machine\Software\Policies\Microsoft contains registry settings for computer certificates that have been distributed by using Group Policy. These settings are stored in the following subkeys:

- HKEY_Local_Machine\Software\Policies\Microsoft\Cryptography
- HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**HKEY_Local_Machine\Software\Policies\Microsoft\Cryptography**

**Autoenrollment**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\Cryptography

**Version**
Windows Server 2003 and Windows XP

Autoenrollment implements Group Policy for autoenrollment of computer certificates. The default autoenrollment behavior is as follows:

Approximately 60 seconds after logon, the balloon UI is displayed. If no user interaction is explicitly defined on the certificate template, no UI will be displayed to the user. This delay is incorporated to enable speedy application and shell response times during the logon and starting of the client computer.

**AEEventLogLevel**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\Cryptography\Autoenrollment

**Version**
Windows Server 2003 and Windows XP

AEEventLogLevel enables enhanced logging of autoenrollment processes for computers.

**HKEY_Local_Machine\Software\Policies\Microsoft\System Certificates**
The following registry entries are located under SystemCertificates. The majority contain binary large objects that pertain to:

- **Certificates.** These entries identify the certificates associated with the registry entry.
- **CRLs.** These entries identify the CRLs associated with the registry entry.
- **CTLs.** These entries identify the CTLs associated with the registry entry.

Additional subkeys — which might appear under some registry subkeys — will be detailed below under the registry subkeys that they correspond to.

**ACRS**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows 2000 and Windows XP

Automatic Certificate Request Settings (ACRS) in Group Policy is used by Windows 2000 to specify certificates to autoenroll. This registry setting has been replaced in Windows XP and Windows Server 2003 by certificate template autoenrollment; however, ACRS is still supported for backward compatibility.

**AddressBook**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

AddressBook contains data about certificates, CRLs, and CTLs relating to other people. Unlike certificates registered in the TrustedPeople container, these certificates are not explicitly trusted.

**AuthRoot**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

AuthRoot contains data about certificates, CRLs, and CTLs from non-Microsoft root CAs.

**ca**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Ca contains data about certificates, CRLs, and CTLs from intermediate CAs.

**Disallowed**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Disallowed contains data about disallowed certificates. For example, it will contain data about certificates that are rejected as untrustworthy during the Authenticode software installation process.

**EFS**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

EFS contains data about certificates, CRLs, and CTLs that enable the use of Encrypting File System (EFS) on the computer.

**Root**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Root contains data about trusted root CA certificates, CRLs, and CTLs.

**Trust**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

Trust contains data about enterprise trust certificates, CRLs, and CTLs.

**TrustedPeople**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

TrustedPeople contains data about certificates from other users where there is explicit trust.

**TrustedPublisher**

**Registry path**
HKEY_Local_Machine\Software\Policies\Microsoft\SystemCertificates

**Version**
Windows Server 2003, Windows XP, and Windows 2000

TrustedPublisher contains data about certificates that have been explicitly accepted as trustworthy. For example, it will contain

data about certificates that are accepted as trustworthy during the Authenticode software installation process.

**Certificate Group Policy Settings**

The following table lists and describes the Group Policy settings that are associated with certificates.

**Group Policy Settings Associated with Certificates**

| Group Policy Setting | Description |
|---|---|
| Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Autoenrollment | Can be used to enroll certificates automatically, renew expired certificates, update pending certificates, and remove certificates that have been revoked. In addition, this setting can be used to block certificate autoenrollment. |
| Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Encrypting File System | Can be used to add or create a data recovery agent for use with EFS. |
| Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Automatic Certificate Request Settings | Can be used to configure automatic certificate request settings for a specific certificate template for a domain by using the Automatic Certificate Request Setup Wizard. The request will be processed automatically at the first occurrence of any of the following: a user logs on, Group Policy is refreshed, or a computer joins the domain and is subject to a Group Policy setting. |
| Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Trusted Root Certification Authorities | Can be used to add a new trusted root CA certificate to a Group Policy object (GPO) for a domain. For a root CA certificate to be imported, the root certificate must be in a PKCS #12 file, in a PKCS #7 file, or in binary-encoded X.509 v3 certificate files. |
| Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Enterprise Trust | Can be used to add a new enterprise trust policy to a GPO for a domain. You do this by using the Certificate Trust List Wizard to create a new CTL for the GPO or assigning an existing CTL to the GPO. Acceptable file formats from which you can import a certificate are:<br>• X.509 v3 certificate files (.cer, .crt)<br>• PKCS #7 files (.spc, .p7b)<br>• Microsoft serialized certificate stores (.sst) |

To find more information about these Group Policy settings, see “Group Policy Settings Reference” in the Tools and Settings Collection.

# Certificate Services Technical Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In order to realize the enhanced security made possible by certificates, organizations need a cost-effective, efficient, secure way to manage the distribution and use of certificates. In the Microsoft Windows Server 2003 operating system, you use Certificate Services to distribute certificates to users, computers, and services. Certificate Services enables you to configure certificates, establish and verify the identity of certificate holders, set policies that control how a certificate is to be used, revoke certificates if they should no longer be considered valid, and publish certificate revocation lists (CRLs) that help users verify a certificate's status.

**In this subject**


- What Is Certificate Services?
- How Certificate Services Works
- Certificate Services Tools and Settings

# What Is Certificate Services?

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Common Certificate Services Scenarios
- Technologies Related To Certificate Services
- Certificate Services Dependencies
- Related Information


Windows Server 2003 Certificate Services provides customizable services for creating and managing public key certificates used in software security systems employing public key technologies.

Organizations use certificates to enhance security by binding the identity of a person, device, or service to a corresponding private key. However, in order to realize the enhanced security made possible by certificates, organizations need a cost effective, efficient, secure way to manage the distribution and use of certificates.

A server configured as a certification authority (CA) provides the management features needed to regulate certificate distribution and use. A CA:

- Configures the format and content of certificates and issues certificates to users, computers, and services.
- Establishes and verifies the identities of certificate holders.
- Sets policies that control how certificates are to be used.
- Revokes certificates if they should no longer be considered valid and publishes certificate revocation lists (CRLs) to be used by certificate verifiers.
- Logs all request, issuance, renewal, and revocation transactions.

Certificate Services is the Windows Server 2003 service that provides the core functionality for Windows Server 2003 CAs. Certificate Services provides customizable services for managing certificates for a particular CA and for the enterprise.

**Common Certificate Services Scenarios**

Managing certificates and CAs involves the following processes:

- **Issuing certificates to users and computers.** The issuance process includes obtaining and validating information about the intended recipient of the certificate, placing policy restrictions designated by the organization in certificates that are issued, and publishing the certificates to a directory.
- **Managing certificate lifetimes.** Because all certificates have a limited life, certificates need to be renewed or allowed to expire. The renewal process is similar to the issuance process, but typically involves fewer security checks. Thus, when an organization develops its renewal strategy, it should balance security concerns against potential disruptions to users.
- **Revoking certificates and verifying revocation status.** Some certificates need to be invalidated before their expiration date. Effective certificate revocation and revocation verification processes are critical to the security of an organization's public key infrastructure (PKI).

For a more detailed technical description of these processes, see "How Certificate Services Works."

**Issuing Certificates to Users and Computers**

A certificate can be used to authenticate the identity of the computer or user that owns the certificate, as well as to encrypt or sign data. Because of this, the certificate issuer must ensure that certificate holders are known entities and meet certain criteria.

In Windows Server 2003, identity verification can be based on:

- Manual verification of a client's credentials by a Certificate Manager.
- Automatic approval based on credentials in Active Directory directory service.

When identity verification is complete, a certificate is constructed based on the certificate policy set by the CA administrator.

Certificate policies are sets of rules that are used when processing certificate requests, issuing certificates, revoking certificates, and publishing certificate revocation lists (CRLs). Certificate policies constitute a combination of administrative policies and configuration settings on the CA:

- **Administrative policies.** For example, an administrative policy might grant commercial certificates only if applicants present identification in person. Another administrative policy might grant credentials based on e-mail requests.
- **Configuration settings on the CA.** When you install Certificate Services, the CA is configured with a default set of rules and settings. These define CA-specific settings such as the CA's certificate, the CA's default issuance behavior, and the CA's key recovery agents. Along with the CA, you can also install a number of preconfigured certificate templates, which define what information a certificate request must have and how to process incoming requests for a certificate based on a particular template.

The combination of applying CA settings and certificate template settings, plus the defined administrative guidelines, are documented in the certificate policy and practice statements that govern the operation of a CA.

When a certificate request has been processed and the certificate has been created, it needs to be published and made available for use. At this point, Certificate Services returns the certificate to the client to be placed in the client's certificate store on the local computer, and data about the certificate is recorded in the certificates database on the CA. In addition, the certificate can be published to Active Directory to make it more accessible to others.

**Managing Certificate Lifetimes**

All certificates have a limited lifetime that needs to be defined by an administrator in a certificate template before the first certificate is issued. The certificate lifetimes that an administrator defines can have an impact on the security of a PKI for the following reasons:

- Over time, encryption keys become more vulnerable to attack. In general, the longer amount of time that key pairs are in use, the greater the risk that keys can be compromised.
- When a CA certificate expires, all certificates that are issued by that CA for validation also expire. This is known as time nesting.

End-entity certificates expire when the issuing CA certificate reaches the end of its lifetime, unless:

- The end-entity certificate is renewed and linked to a CA certificate with a longer lifetime.
- The end-entity certificate was revoked before the CA certificate expiration date is reached.

**Revoking Certificates and Verifying Revocation Status**

Unlike a certificate that reaches the end of its lifetime and is not renewed, a certificate might need to be designated unusable before its lifetime is complete. This is done by a process called revocation.

A certificate manager should revoke a certificate if the recipient of the certificate breaks the rules that are defined in the certificate practice statement or if the private key that is associated with the certificate is compromised. For example, a rule might specify that certificates can only be issued to a current employee of the organization or to an employee who performs a specific task. If that employee leaves the organization or is reassigned, the certificate should be revoked. Similarly, a certificate should be revoked if the private key associated with the certificate has been compromised. In either case, the goal is to prevent unapproved use of the certificate and its associated key.

**Note**

- Even when a certificate has been revoked, it is up to the party relying on the certificate to decide whether to accept or reject the certificate for an intended purpose. For example, some parties would accept an expired passport as proof of a bearer's identity, although they would not allow an expired passport to be used to cross a national border.

Verifying revocation status is an important corollary to certificate revocation. When a certificate is issued and used, certificate policies determine whether a certificate requester or certificate user meets the criteria set for the certificate. However, when revocation verification takes place and a certificate is determined to have been revoked, all other certificate policy verifications are overridden. Because a certificate that was valid yesterday might not be acceptable today, the revocation verification process is critical to the integrity of the PKI.

#### Technologies Related to Certificate Services

Certificate Services is an enabling technology that can be used in a wide variety of applications. For example, it can be used to enhance the security of technologies such as:

- **Authentication.** Certificates and Certificate Services extend and enhance, rather than replace, the security provided by password-based authentication.

- **Access control.** Certificates and Certificate Services make it possible to limit user access to highly sensitive data.

- **Digital signing.** Certificates and Certificate Services make it possible to positively identify a user who modifies data by applying digital signatures to files and documents.

- **Wireless networking.** Certificates and Certificate Services can be used to enhance the authentication process for wireless clients and deter unauthorized users from accessing a wireless network or from viewing data as it crosses between wireless access points and client computers.

- **Virtual private networks (VPNs).** Certificates and Certificate Services can be used to authenticate clients that are allowed to initiate VPN connections to an organization's network. In addition, certificates and Certificate Services can be used to deter unauthorized users from viewing or intercepting data as it crosses the public Internet through a VPN connection.

Because Certificate Services is an enabling technology, technology strategists should consider the business problems faced by their organizations to determine how Certificate Services can help solve those problems.

#### Certificate Services Dependencies

Using Certificate Services in Windows involves three types of dependencies:

- What Windows version you are using.

- Whether the computer is joined to a domain, and the version of Active Directory that the organization is using.

- Whether the organization's PKI will be linked to any other organization's PKI.

#### Windows Version Dependencies

In Windows Server 2003, Certificate Services includes a number of new and enhanced features. Windows Server 2003 Enterprise Edition includes some enhanced features that are not available with Windows Server 2003 Standard Edition. In addition, using Windows XP Professional clients enhances an organization's ability to take advantage of many of these new features.

Many features and capabilities require that you run Windows Server 2003 on the server and run Windows XP Professional on the client. The following table illustrates key differences in capabilities that are available in environments using different client and server operating systems.

**Feature Dependencies in Certificate Services**

| | Version 2 Templates | Key Archival and Recovery | Autoenrollment | Delta CRL | Qualified Subordination |
|---|---|---|---|---|---|
| Description | New features and options including customizable enrollment, validity, issuance, application, and key usage policies | Archival and recovery of private keys associated with a certificate request | Enables the issuance of certificates to computers and users | Lists of certificates whose status has changed since the last full (base) CRL was compiled can be published more frequently with less impact on network capacity. | Administrator can add restrictions to a CA certificate. These restrictions are translated into certificate extensions within issued certificates. |
| Windows Server 2003 Enterprise Edition | Enterprise CA only | Yes | User and computer | Yes | Yes |
| Windows Server 2003 Standard Edition | | | Computer only | Yes | Yes |
| Windows 2000 Server | | | Computer only | | |
| Windows XP Professional client with Windows 2000 Server CA | | | Computer only | | |
| Windows XP Professional client with Windows Server 2003 CA | Yes | Yes | User and computer | Yes | Yes |
| Windows 2000 Professional client with a Windows 2000 Server CA | | | Computer only | | |
| Windows 2000 Professional client with a Windows Server 2003 CA | By Web enrollment support only | By Web enrollment support or programmatic script | Computer only | | |

**Active Directory Dependencies**

If you are using Active Directory to manage user and computer accounts on the network, your enterprise CAs can use information contained in Active Directory — such as user account names, security group memberships, and certificate templates — to simplify certificate enrollment, renewal, revocation, and validation.

The Active Directory schema in Windows Server 2003 has been enhanced to enable a number of new Certificate Services capabilities — version 2 templates, key recovery and archival, and delta CRLs, for example — that were described in the previous section.

In addition, Active Directory stores public key Group Policy for distribution to all computers that are within the scope of the policy. Public key Group Policy enables you to control which CAs are to be trusted in the enterprise, to specify alternative Encrypting File System (EFS) recovery agents, and to configure automatic enrollment and renewal of certificates for computers running Windows 2000, Windows XP, and Windows 2003 — all from a central administration point.

Active Directory also supports mapping certificates to network user accounts for authenticating clients and controlling access to network resources. Using smart cards for the user logon process is a special case of certificate mapping that extends the Kerberos V5 authentication protocol to include authentication of users on the basis of certificates and private keys that are stored on smart cards. Using smart cards for the user logon process provides enhanced security for user authentication and a single set of user credentials for logging on locally or remotely over a network.

The Active Directory schema in Windows Server 2003 has been enhanced to enable a number of new Certificate Services capabilities — for example, version 2 templates, key recovery and archival, and delta CRLs — that were described in the previous section.

**Non-Microsoft Dependencies**

Windows Server 2003 Certificate Services is based on common industry standards defined by the Internet Engineering Task Force (IETF) Public Key Infrastructure (PKIX) working group and others. Windows Server 2003 Certificate Services supports the following industry standards:

- **RFC 2313 (PKCS #1).** A method for encrypting data using Rivest-Shamir-Adleman (RSA) public key cryptography.
- **RFC 2527.** A framework for creating certificate policies and certification practice statements for CAs and PKIs.
- **RFC 2528.** Representation of Key Exchange Algorithm (KEA) Keys in Internet X.509 Certificates.
- **RFC 2559.** Addresses operational protocols that make it possible to retrieve data regarding X.509 certificates and CRLs from an LDAP directory.
- **RFC 2560.** Guidelines for support of Online Certificate Status Protocol (OCSP).

  **Note**

  - This requires the use of non-Microsoft OCSP software.
- **RFC 2585.** Specifies the conventions for using the File Transfer Protocol (FTP) and the Hypertext Transfer Protocol (HTTP) to obtain certificates and CRLs from PKI repositories.
- **RFC 2587.** Defines auxiliary object classes that need to be supported in an LDAP schema to support PKIX requirements.
- **RFC 2255.** Guidelines for formatting LDAP URLs. Clarifies how LDAP URLs are resolved.
- **RFC 2797.** Defines a certificate management protocol using Certificate Management Messages over CMS.
- **RFC 3280.** Formatting of X.509 v3 certificates and X.509 v2 CRLs for use on the Internet.

In addition, Windows Server 2003 Certificate Services also supports industry standards that enable the following:

- **Transport Layer Security (TLS).** Provides a secure and authenticated channel between hosts on the Internet above the transport layer.
- **Secure Multi-purpose Internet Mail Extensions (S/MIME).** Serves as a standard for secure e-mail across the Internet.
- **Kerberos V5 authentication protocol.** Provides a symmetric key framework for authentication in networks.

A PKI based on Certificate Services can generally be linked to an external organization's PKI, even if the external PKI is not based on Certificate Services, so long as the external PKI supports the above standards. A detailed comparison of how each PKI implements these standards is generally required in order to determine compatibility.

**Related Information**

The following resources contain additional information that is relevant to this section.

- IPSec Technical Reference
- Logon and Authentication Technologies

# How Certificate Services Works

Updated: September 22, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Certificate Services Architecture
- Certificate Services Protocols
- Certificate Services Interfaces
- Certificate Services Physical Structure
- Certificate Services Processes and Interactions
- Related Information


Certificate Services provides customizable services for issuing and managing certificates used in software security systems employing public key technologies. You can use Certificate Services to:

- Create certification authorities (CAs) which will receive certificate requests.
- Verify the information in a certificate request.
- Verify the identity of the certificate requester.
- Issue and renew certificates.
- Revoke certificates.
- Publish certificate revocation lists (CRLs).

Additionally, in Windows Server 2003, you can use Certificate Services to:

- Enroll users for certificates from the CA using the Web, or the Microsoft Management Console (MMC) Certificates snap-in, or transparently through autoenrollment.
- Take advantage of programmatic and command-line scripting to perform common tasks such as enrollment processing.
- Help simplify — by using certificate templates and CA policy — the choices a certificate requester has to make when requesting a certificate.
- Take advantage of the Active Directory directory service to publish trusted root certificates, publish issued certificates, and publish CRLs.
- Implement the ability to log on to an Active Directory domain using a smart card.

An optimal environment for Certificate Services includes one or more servers running Certificate Services. The servers running Certificate Services have been planned and deployed to enable:

- Efficient verification of certificates through a process called chain building. With the chain-building process, all certificates in a CA hierarchy are verifiable and lead to a root CA that provides an ultimate anchor of trust.

- Security of all CAs in the hierarchy by means of proper configuration and delegation of administrative responsibilities. In addition, it might be desirable to take the root CA offline in order to more stringently protect the hierarchy from compromise.

- Well defined processes and procedures for validating certificate requests and renewals, as well as for revoking certificates.

The following sections provide an in-depth view of how Certificate Services works in an optimal environment.

**Certificate Services Architecture**
In order to properly configure and manage Certificate Services, you need to understand its major components, the parameters of those components, and dependencies. The following figure shows a functional block diagram of how Certificate Services is implemented in Windows Server 2003.

**Windows Server 2003 Certificate Services Functional Diagram**




The Certificate Services engine, exit modules, and policy modules of Windows Server 2003 Certificate Services work in conjunction with CryptoAPI to perform a variety of tasks, which include:

- Submitting certificate requests (client module).

- Processing certificate requests (Certificate Services engine).

- Verifying whether requesters are qualified to receive certificates (policy module).

- Creating and issuing certificates for qualified requesters (Certificate Services engine).

- Generating a private key and distributing it to the requester's protected certificate store (CryptoAPI and cryptographic service providers).

- Managing the private key for all cryptography operations (CryptoAPI).

- Distributing the certificates that are issued to qualified clients and, optionally, publishing certificates to Web pages, public folders, or Active Directory (exit module).

- Publishing periodic CRLs to Active Directory and, optionally, to Web pages or public folders (exit module).

- Storing all certificate transactions for the audit trail (certificates database).

  **Note**

  - In Windows Server 2003, all cryptographic functions and private key management are performed by CryptoAPI in conjunction with cryptographic service providers (CSPs). CryptoAPI also manages any interactions with hardware security modules (HSMs), if present. Any system service or application can request cryptographic services by using CryptoAPI.

The following sections describe these Certificate Services components:

- Client module
- Certificate Services engine
- Policy module
- Exit module
- Certificates database
- CryptoAPI

**Client Module**
The client module is the software that is used by the end user to generate a certificate request, send the request, and receive the finished certificate. Windows Server 2003 Certificate Services provides CA services for X.509 v3 certificate clients. It supports all Windows clients, as well as any other clients that can generate X.509 v3–compliant certificate requests.

**Certificate Services Engine**
The Certificate Services engine, whose file name is Certsrv.exe, is the core component of Certificate Services. The engine acts as a broker for all requests it receives from clients and regulates the flow of information between the other components during the processing of a request and the generation of a certificate. At each stage of processing, the engine interacts with the various modules to ensure appropriate action is taken based on the state of the request.

The Certificate Services engine supports the following CA operations:

- Issuance
- Publication
- Revocation
- CRL maintenance
- Other operations useful for certificates, such as key archival

Certificate Services operates in two modes: enterprise and stand-alone. In both modes, Certificate Services performs two main operations: receiving requests and publishing requests. Requests are received by a Distributed Component Object Model (DCOM) interface that clients call to submit requests. Requests are published by a number of different methods, including:

- Stand-alone Web enrollment with IIS and Certificate Services
- Enterprise Web enrollment with IIS and Certificate Services
- Enterprise enrollment with an enrollment agent
- Enterprise direct enrollment

These enrollment methods are described later in this section.

**Request Handling**
The default entry module processes these types of requests:

- Public Key Cryptography Standard (PKCS) #10 certificate requests
- Certificate Management protocol using CMS made through remote procedure calls (RPCs) or HTTP

The entry module is a dynamic-link library (DLL) that cannot be customized. Windows Server 2003 services usually use RPCs to submit certificate requests to enterprise CAs. However, the Web Enrollment Support pages use HTTP to submit certificate requests to CAs.

Intermediary applications can be written to handle certificate requests from different types of clients, across multiple transports, or according to policy-specific criteria. IIS is one intermediary application that accepts HTTP-based requests and makes DCOM calls on behalf of users to the CA interfaces. Intermediaries can also check on the status of a previously submitted request and obtain the Certificate Services configuration information.

Certificate requests to Certificate Services are placed in a pending queue until they are approved or denied by the policy module.

**Note**

- You can develop custom certificate enrollment applications that submit RPC or HTTP requests to Certificate Services. For more information about developing custom applications for Windows Server 2003 Certificate Services and about the required certificate request format, see "Using Certificate Enrollment Control" in the Cryptography collection of the Microsoft Platform SDK on MSDN.

**Policy Module**

After Certificate Services has parsed a certificate request, it loads the policy module and tells it where to find the request. The default policy module then examines the request, validating the different fields against its policy.

The policy module contains the set of rules governing issuance, renewal, and revocation of certificates. This policy is created from hard-coded values, registry settings, and — in the enterprise mode — certificate templates. Certificate Services provides a callback interface for the policy module. The callback interface enables the module to request information about a certificate request from the certificates database.

**Note**

- The policy module does not directly access the request; all access is through the Certificate Services engine.

The policy module determines whether a certificate request is approved, denied, or queued (left pending) for an administrator to approve or deny issuance. It does this by verifying the information in the request against various sources, such as a directory service, an external legacy database, or information from an outside authority. In addition, the policy module can send alerts to the appropriate administrator if manual (offline) approval of the request is required.

You can create custom policy modules to insert additional certificate attributes and extensions that might be required by a client application. Extension handlers work in tandem with the policy module to set rules extensions on a certificate. Each extension handler acts as a template for the rules extensions that should appear in a certificate. The policy module must load the appropriate extension handler when it is needed. For example, information such as a job title and a signing limit can be inserted and used by an online purchasing form to determine whether the user can sign for the amount requested.

The policy module can use additional information included in the certificate request to incorporate requested attributes in the issued certificate. For example, certificate requests to stand-alone CAs must include all information about the requested certificate. Therefore, the policy module incorporates this information into each issued certificate. However, enterprise CAs use certificate templates to specify certificate attributes, so certificate requests to enterprise CAs require less information.

After it has validated all the fields, the policy module will create some of the necessary certificate fields, such as the validity period, and determine whether the request should be approved or denied. The policy module then notifies the Certificate Services engine of the result.

Windows Server 2003 Certificate Services includes default policy modules for enterprise and stand-alone CAs. In addition, you can build custom policy modules to address special needs.

**Enterprise CA policy**

By default, an enterprise CA policy always issues a certificate or denies a request immediately. Enterprise CA policy uses Active Directory to determine the identity of the requester, and then automatically determines whether the requester has security permissions to receive a certificate of the type being requested. However, on a per-template basis, you can specify that certificate requests are directed to a pending queue so that an administrator can approve or deny them.

**Stand-alone CA policy**

By default, stand-alone CA policy sends certificate requests to a pending queue so that an administrator can approve or deny them. If you prefer, you can configure stand-alone CA policy to automatically approve all certificate requests. However, because

a stand-alone CA does not verify the identity of requesters in Active Directory, there is no automatic means to verify the identity and validity of the certificate requester. Therefore, if a stand-alone CA is configured to approve certificate requests automatically, it can pose a significant security risk.

**Custom policy modules**
Policy module settings are fully customizable. You can customize policy enforcement by:

- Modifying the installed policy module using the Certification Authority snap-in to create a custom policy module that then makes an explicit call to the default policy module.
- Developing your own policy modules.
- Acquiring a non-Microsoft policy module when one is available.

After you customize the policy module, you can configure Certificate Services to only call the new policy module.

For more information about how to customize policy modules, see "Certificate Enrollment Control" in the Cryptography collection of the Microsoft Platform SDK on MSDN.

**Exit Module**
After the Certificate Services engine has signed and published either a newly created certificate or new CRL, it notifies the exit module of this event. The engine provides a callback interface for the exit module, which allows the exit module to query the certificates database about issued certificates and other data. Exit modules publish completed certificates and CRLs through a number of transports or protocols. For more information, see "Certificate Services Protocols" later in this document.

### Note

- CA administrators can add in additional responses to notifications by creating a custom exit module. One example might be: Call this pager when a certificate with the user name of the CEO is issued.

The exit module packages the issued certificate in the appropriate transport mechanism or protocol and distributes it to the location specified in the request. Certificate requests can specify that the certificate be distributed to Lightweight Directory Access Protocol (LDAP) directory services, file systems, or URLs. The exit module also delivers CRLs to CRL distribution points.

By default, the server notifies each exit module installed on the server whenever a certificate or CRL is published. The default enterprise exit module publishes certificates and CRLs to Active Directory. The default stand-alone exit module publishes certificates and CRLs to the local file system.

Because Windows Server 2003 Certificate Services supports multiple exit modules, you can use the certification authority snap-in to install multiple exit modules for a CA. For example, you can install exit modules that post certificates to legacy open database connectivity (ODBC) databases or to non-Microsoft LDAP directory services.

As with policy module settings, exit module settings are fully customizable. For more information about customizing exit modules, see "Writing Custom Exit Modules" in the Cryptography collection of the Microsoft Platform SDK on MSDN.

**Certificates Database**
The certificates database records all certificate transactions. The certificates database:

- Tracks each certificate request, the requester, and whether the request was granted or denied.
- Records information for the issued certificate, such as the private key, serial number and expiration date.
- Provides a complete audit trail for each certificate from request to expiration.
- Flags and tracks certificates that are revoked by CA administrators.

You can use the Certification Authority snap-in to manage the audit trail.

### Note

- To audit CA activity, use the Certification Authority snap-in and Event Log.

For more information about the certificates database, see “Certificate Services Physical Structure” later in this document.

**CryptoAPI**
Certificate Services uses CryptoAPI for all cryptographic operations, including key generation and exchange, data encryption and decryption, hashing, digital signatures, and signature verification.

CryptoAPI is a generic application programming interface to cryptographic modules, called cryptographic service providers (CSPs). These CSPs are modular in nature and are replaceable by system administrators. During installation, the CA administrator can choose which CSP Certificate Services will use. For most default installations of Windows Server 2003 and Certificate Services, this will be the Microsoft RSA-enhanced provider.

CryptoAPI can also be configured with a hardware-based CSP. If a hardware-based CSP is installed, the CA Administrator can configure Certificate Services to use this CSP.

**Note**

- CSPs cannot access any CryptoAPI interfaces and operate on Windows unless they have received a digital signature from Microsoft. The CSP digital signature is verified each time the CSP is loaded, and each time it attempts to access the CryptoAPI interface. If the signature verification fails, no cryptographic functions are accessible to applications.

Applications can use the functions in CryptoAPI without needing information about the underlying security implementation, just as an application can use a graphics library without needing information about the particular graphics hardware configuration.

**Certificate Services Protocols**
Two protocols, DCOM and LDAP, provide key Certificate Services functionality.

Certificate Services uses the DCOM protocol for enrollment. Certificate Services performs two main operations: receiving requests and publishing requests. Requests are received by a DCOM interface that clients call to submit requests. Requests are published by a number of different methods, which are discussed in How Certificate Enrollment and Renewal Requests are Processed later in this document.

In the enterprise mode, both Certificate Services and clients use a directory service, such as Active Directory, to provide and access services for a domain. Both Certificate Services and clients retrieve templates from Active Directory. Certificate Services also accesses Active Directory by LDAP to verify the client and to store and retrieve certificates.

In addition, Certificate Services can distribute certificates through any transport mechanism. That is, it can accept certificate requests from an applicant and post certificates to the applicant through HTTP, RPC, disk file, or custom transport.

**Certificate Services Interfaces**
Certificate Services publishes the following external DCOM interfaces:

- **ICertRequest and ICertRequest2.** These interfaces accept requests for new and existing certificates.

- **ICertAdmin and ICertAdmin2.** These interfaces provide administrative functions, including request management (modification, denial, resubmission, and revocation), importing certificates, key archival, CRL management, server configuration, user role retrieval, and certificate data retrieval.

- **ICertView and ICertView2.** These interfaces allow properly authorized clients to create a customized or complete view of the Certificate Services certificates database.

- **ICertPassage.** The ICertPassage RPC protocol provides an interface to a CA for submitting a PKCS #10 request and receiving a certificate in a PKCS #7 response.

  **Note**

  - ICertPassage has been superseded in Windows Server 2003 by ICertRequest, but is still used by Windows 2000 clients.

Access to these interfaces is available to developers through licensable APIs. Each of these published interfaces has one of four terminating interfaces inside Certificate Services:

- **ICertRequestD and ICertRequestD2**
- **ICertAdminD and ICertAdminD2**

Certificate Services also uses five internal interfaces, which are not accessible to developers:

- **ICertPolicy.** Certificate Services determines whether to deny, send to a pending queue, or issue a certificate request through this interface, by calling the policy module.
- **ICertServerPolicy.** The policy module queries and modifies certificate request data through this interface.
- **ICertExit.** Certificate Services notifies any exit modules through this interface upon certificate or CRL issuance.
- **ICertServerExit.** An exit module can query the Certificate Services engine about issued certificates through this interface.
- **ICertDB.** Certificate Services uses this interface to access the certificates database.

The following figure shows the relationship between the published external interfaces, the terminating interfaces, and the internal interfaces for Windows Server 2003 Certificate Services.

**Certificate Services Interfaces**




**Certificate Services Physical Structure**

When people think of Certificate Services physical structure, they typically think in terms of the number of servers running Certificate Services, how these servers are linked, whether these servers are dedicated exclusively to Certificate Services or perform multiple functions, and so on. Many of these issues are addressed in PKI design documents. For more information about these issues, see "PKI Technologies."

The physical structure of Certificate Services itself can best be described in terms of where certificates and certificate-related data are stored. Certificate Services uses a number of places to store certificates and configuration data:

- **The certificates database and CA log file**
- **The registry**
- **Active Directory**
- **The file system**

Identical or closely related certificate data can appear in several — even in all — of these locations. For example, parts of the configuration information from the configuration container in Active Directory are replicated to the clients certificate store. Certificates that are stored in the Active Directory Configuration container are deployed to all clients across the forest. And, if a certificate is registered in the Configuration container and the domain security Group Policy object (GPO), a certificate can appear twice on the client.

**Certificates Database**

Certificate Services uses an Extensible Storage Engine (ESE) database to store all certificate requests in any status, all issued certificates, certificate extensions, certificate attributes, and CRLs. The database is composed of four tables: the Request and Certificates table, the Extensions table, the Attributes table, and the CRL table.

These tables maintain status information (receipt, parsing, authorization, signing, and dispatch) as the service processes a certificate request. When Certificate Services receives a request, it creates a row in the certificates database and begins to parse the request and verify the request signatures.

The database also stores all certificates and CRLs issued by the service so that administrators can track, audit, and archive service activity. As Certificate Services parses a certificate request, it fills in the fields of the certificates database from the fields in the request. After the request has been fully parsed, the policy module is notified and the context handle corresponding to the request is passed as a parameter. In addition, the database is used by the Certificate Services engine to store pending revocations before publishing them in a CRL.

Because the certificates database is a transaction database, it includes certificate log files, which record all certificate transactions. By default, the certificates database and the CA log files are installed at the following location:

*systemroot*\system32\certlog

You can choose another location to install either the certificates database or the CA log files — including installing the database and log files separately on different drives — when you are installing a CA.

**Note**

- There is a one-to-one mapping between the tables in the certificates database.

The following table describes the contents of the certificates database.

**Schema of the Certificates Database**

| Column Name | Localized Name | Type | Maximum Length |
|---|---|---|---|
| Request.RequestID | Request ID | Long | 4 -- Indexed |
| Request.RawRequest | Binary Request | Binary | 65536 |
| Request.RawArchivedKey | Archived Key | Binary | 65536 |
| Request.KeyRecoveryHashes | Key Recovery Agent Hashes | String | 8192 |
| Request.RawOldCertificate | Old Certificate | Binary | 16384 |
| Request.RequestAtributes | Request Attributes | String | 32768 |
| Request.RequestType | Request Type | Long | 4 |
| Request.RequestFlags | Request Flags | Long | 4 |
| Request.StatusCode | Request Status Code | Long | 4 |
| Request.Disposition | Request Disposition | Long | 4 -- Indexed |
| Request.DispositionMessage | Request Disposition Message | String | 8192 |
| Request.SubmittedWhen | Request Submission Date | Date | 8 |
| Request.ResolvedWhen | Request Resolution Date | Date | 8 -- Indexed |
| Request.RevokedWhen | Revocation Date | Date | 8 |
| Request.RevokedEffectiveWhen | Effective Revocation Date | Date | 8 -- Indexed |
| Request.RevokedReason | Revocation Reason | Long | 4 |

| | | | |
|---|---|---|---|
| Request.RequesterName | Requester Name | String | 2048 -- Indexed |
| Request.CallerName | Caller Name | String | 2048 |
| Request.SignerPolicies | Signer Policies | String | 8192 |
| Request.SignerApplicationPolicies | Signer Application Policies | String | 8192 |
| Request.Officer | Officer | Long | 4 |
| Request.DistinguishedName | Request Distinguished Name | String | 8192 |
| Request.RawName | Request Binary Name | Binary | 4096 |
| Request.Country | Request Country/Region | String | 8192 |
| Request.Organization | Request Organization | String | 8192 |
| Request.OrgUnit | Request Organization Unit | String | 8192 |
| Request.CommonName | Request Common Name | String | 8192 |
| Request.Locality | Request City | String | 8192 |
| Request.State | Request State | String | 8192 |
| Request.Title | Request Title | String | 8192 |
| Request.GivenName | Request First Name | String | 8192 |
| Request.Initials | Request Initials | String | 8192 |
| Request.SurName | Request Last Name | String | 8192 |
| Request.DomainComponent | Request Domain Component | String | 8192 |
| Request.EMail | Request E-mail Address | String | 8192 |
| Request.StreetAddress | Request Street Address | String | 8192 |
| Request.UnstructuredName | Request Unstructured Name | String | 8192 |
| Request.UnstructuredAddress | Request Unstructured Address | String | 8192 |
| Request.DeviceSerialNumber | Request Device Serial Number | String | 8192 |
| RequestID | Issued Request ID | Long | 4 -- Indexed |
| RawCertificate | Binary Certificate | Binary | 16384 |
| CertificateHash | Certificate Hash | String | 128 -- Indexed |
| CertificateTemplate | Certificate Template | String | 254 -- Indexed |
| EnrollmentFlags | Template Enrollment Flags | Long | 4 |
| GeneralFlags | Template General Flags | Long | 4 |
| SerialNumber | Serial Number | String | 128 -- Indexed |
| IssuesrNadeID | Issuer Name ID | Long | 4 |
| NotBefore | Certificate Effective Date | Date | 8 |
| NotAfter | Certificate Expiration Date | Date | 8 -- Indexed |

| | | | |
|---|---|---|---|
| SubjectKeyIdentifier | Issued Subject Key Identifier | String | 128 |
| RawPublicKey | Binary Public Key | Binary | 4096 |
| PublicKeyLength | Public Key Length | Long | 4 |
| PublicKeyAlgorithm | Public Key Algorithm | String | 254 |
| PublicKeyAlgorithmParameters | Public Key Algorithm Parameters | Binary | 4096 |
| UPN | User Principal Name | String | 2048 -- Indexed |
| DistinguishedName | Issued Distinguished Named | String | 8192 |
| RawName | Issued Binary Name | Binary | 4096 |
| Country | Issued Country/Region | String | 8192 |
| Organization | Issued Organization | String | 8192 |
| OrgUnit | Issued Organization Unit | String | 8192 |
| CommonName | Issued Common Name | String | 8192 -- Indexed |
| Locality | Issued City | String | 8192 |
| State | Issued State | String | 8192 |
| Title | Issued Title | String | 8192 |
| GivenName | Issued First Name | String | 8192 |
| Initials | Issued Initials | String | 8192 |
| SurName | Issued Last Name | String | 8192 |
| DomainComponent | Issued Domain Component | String | 8192 |
| EMail | Issued E-mail Address | String | 8192 |
| StreetAddress | Issued Street Address | String | 8192 |
| UnstructuredName | Issued Unstructured Name | String | 8192 |
| UnstructuredAddress | Issued Unstructured Address | String | 8192 |
| DeviceSerialNumber | Issued Device Serial Number | String | 8192 |

**Registry**

Certificate Services uses the registry to store all of its configuration data, as well as other data. This registry key is:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc

For more information about relevant registry settings, see "Certificate Services Tools and Settings."

**Active Directory**

Various PKI-related containers such as certification authorities, enrollment services, templates, object identifiers (OIDs), and the authority information access and CRL distribution point extensions are created when the first enterprise CA is created in a forest. The permissions on the containers are also set at that time. These containers are replicated to every domain controller within a forest.

Installing an enterprise CA creates the following containers.

**"Trusted root CA and CRL distribution point containers"**

Before you can install an issuing CA, you must publish the root CA's certificate to the Active Directory trusted root CA container, and publish the root CA's CRL to the Active Directory CRL distribution point container. This will cause all domain

members (including the issuing CA) to import the root CA's certificate into their own root stores and will enable them to verify the revocation status of any certificates issued by the root CA. The common name of the trusted root CA container is:

CN=Certification Authorities,CN=Public Key Services,CN=Services,CN=Configuration,DC=

The common name of the CRL distribution point container is:

CN=CDP,CN=Public Key Services,CN=Services,CN=Configuration,DC=

**Authority information access container**
The authority information access container is the publication point for the most currently published CA certificates for root and intermediate CAs. The authority information access extension in a certificate identifies the location of the authority information access container, which helps clients find CA certificates dynamically during chain building.

The common name of the authority information access container is:

CN=AIA,CN=Public Key Services,CN=Services,CN=Configuration,DC=

**NTAuthCertificates container**
Publishing CA certificates to the NTAuthCertificates leaf object indicates that these CAs are trusted to both:

- Issue authentication (logon) certificates for any user in the forest.
- Enable logon for smart cards, IIS mapping, and Extensible Authentication Protocol-Transport Layer Security (EAP-TLS).

The common name of the NTAuthCertificates leaf object is:

CN=NTAuthCertificates,CN=Public Key Services,CN=Services,CN=Configuration,DC=?

**Certificate Templates container**
Installing the first enterprise CA in the forest also installs 29 certificate template objects in Windows Server 2003 or 24 certificate template objects in Windows 2000. The common name for the Certificate Templates container is:

CN=Certificate Templates,CN=Public Key Services,CN=Services,CN=Configuration,DC=

For more information about certificate templates in Windows Server 2003 and Windows 2000, see the "Certificates Technical Reference."

**OID container**
Windows Server 2003 adds an object identifier (OID) container to the configuration container. Because OIDs are not hard-coded in version 2 templates, the OID container is required to work with version 2 templates. The common name for the OID container is:

CN=OID,CN=Public Key Services,CN=Services,CN=Configuration,DC=

Windows Server 2003 includes four predefined issuance policy OIDs:

- **All issuance (2.5.29.32.0).** The all issuance policy OID indicates that the issuance policy contains all other issuance policies. Typically, this object identifier is only assigned to CA certificates.
- **Low assurance (1.3.6.1.4.1.311.21.8. *x.y.z.* 1.400).** The low assurance OID is used to represent certificates that are issued with no additional security requirements.

  **Note**

  - The *x.y.z* portion of the OID is a randomly generated numeric sequence that is unique for each Windows Server 2003 forest.
- **Medium assurance (1.3.6.1.4.1.311.21.8. *x.y.z* .1.401).** The medium assurance OID is used to represent certificates that have additional security requirements for issuance. For example, a smart card certificate that is issued in a face-to-face meeting with a smart card issuer might be considered a medium assurance certificate and contain the medium assurance object identifier.
- **High assurance (1.3.6.1.4.1.311.21.8. *x.y.z* .1.402).** The high assurance OID is used to represent certificates that are

issued with the highest security. For example, the issuance of a key recovery agent certificate might require additional background checks and a digital signature from a designated approver because a person holding this certificate can recover private key material from a Windows Server 2003, Enterprise Edition CA.

**Key recovery agent container**

Key recovery agent certificates are published to the key recovery agent container in Active Directory when enrollment occurs. Because a CA can issue multiple key recovery agent certificates, each key recovery agent certificate will be added to the multi-valued user attribute of the CA object.

The common name for the key recovery agent container is:

CN=KRA,CN=Public Key Services,CN=Services,CN=Configuration,DC=

### Note

- The key recovery agent container only appears under Windows Server 2003.

**File System**

By default, the only directory in the file system that is used by Certificate Services is *systemdir*\CertLog. This is where the certificates database and request logs are stored. Some organizations prefer to create two separate directories, one for the database and one for the log files.

Certificate Services and its associated tools produce several logs which can be used for troubleshooting:

- Certificate Services logs to *systemroot*\certsrv.log (when debug logging is enabled).
- Certutil.exe logs to *systemroot*\certutil.log.
- Certification Authority snap-in logs to windir\certmmc.log.

Certificates and CRLs also are stored on the local file system. During the installation of the root CA, the root certificate is saved to the following shared folders:

- \\Localhost\Certenroll
- \\Locahost\Certconfig
- The SystemCertificates\My folder of the local computer, which is located at:

  \\Documents and Settings*username*\Application Data\Microsoft\SystemCertificates\My
- The trusted root certification authorities container in the local computer registry, which is located at:

  HKEY_Local_Machine\Software\Microsoft\SystemCertificates\Root

The initial CRL is published in the following shared folders:

- \\Localhost\Certenroll
- The intermediate certification authorities container in the registry of the local computer.

The CA certificate of a stand-alone CA is stored in the following shared folders:

- \\Localhost\Certenroll
- \\Locahost\Certconfig

#### Contents of the CertConfig and CertEnroll Shared Folders

More than one certificate file exists in the \CertConfig and \CertEnroll shared folders after a period of time. If the CA name is used as part of a file name, the sanitized CA name adds additional escape characters in order to accommodate any extended ASCII characters in the file name. The escape characters appear in the file name as %20.

The following table explains the certificate file name extensions and their purposes.

**Certificate Paths and File Name Extensions**

| Sample File Name | Description |
|---|---|
| \\Localhost\CertConfig\Certsrv.txt | This is the configuration file for the CA. |
| \\Localhost\CertConfig\Certsrv.bak | If the CA is reinstalled, the original CA configuration file is renamed with the .bak extension. |
| \\Localhost\CertConfig\SubCAname.req<br>\\Localhost\CertConfig\SubCAname(1).req | Request file that is used to generate the CA certificate. Request files are used only for subordinate CAs. Request files are generated with the same base file name suffix as certificates. |
| *systemroot*\\SubCAname.req<br>*systemroot*\\SubCAname(1).req | If no shared folder was created during the CA setup procedure and Active Directory is used to publish the CA's configuration information, request files are written to the systemroot drive instead of to the \\Localhost\CertConfig file. |
| \\Localhost\CertConfig\CAname.crt<br>\\Localhost\CertEnroll\CAname.crt | This is the concatenated name of the original root CA certificate. |
| \\Localhost\CertConfig\CAname(1).crt<br>\\Localhost\CertEnroll\CAname(1).crt | When a root CA certificate is renewed, the version number (in parentheses) is appended to the certificate name. |
| \\Localhost\CertConfig\CAname(0-1).crt<br>\\Localhost\CertEnroll\CAname(0-1).crt<br>\\Localhost\CertConfig\CAname(1-0).crt<br>\\Localhost\CertEnroll\CAname(1-0).crt | When a root certificate is renewed, a cross-certificate is automatically created from the old root certificate, which is trusted by all clients, to the new root certificate, which is not yet trusted. If the old root certificate is version 0 and the new root is version 1, the CA name is appended with the version number of the old root certificate and the version number of the new root certificate (in parenthesis and separated by a dash). If the cross-certificate provides cross trust from the new to the old CA certificates, then the order of the numbers in parenthesis is reversed. |
| \\Localhost\CertEnroll\CAname.crl | This is the name format for the original base revocation list for a CA. |
| \\Localhost\CertEnroll\CAname(1).crl | When an updated base CRL is created, the updated version is indicated in parentheses. Thus, (1) indicates the first updated base CRL. |
| \\Localhost\CertEnroll\Caname+.crl | A plus sign is added to the base CRL name to indicate that this is a delta CRL. |

As illustrated in this table, an important function of the CertConfig and CertEnroll shared folders is to provide links between original and renewed root certificates, and original and updated CRLs. For root certificates, these links are provided by cross-certificates. The cross-certificates are automatically generated when Certificate Services starts after renewing a root CA certificate with a new key. Cross-certificates are not created for subordinate CAs, nor are they created when a root certificate is renewed with the same key. If you upgrade from Windows 2000 Server after renewing a root CA certificate with a new key, the cross-certificate is generated the first time that Certificate Services starts after you upgrade to Windows Server 2003.

#### Certificate Services Processes and Interactions

Certificates are a powerful security technology. Therefore, it is important to understand and carefully manage the processes performed by Windows Server 2003 Certificate Services. These processes are described in the following subsections:

- **How certificates are created.** This describes the core functioning of Windows Server 2003 Certificates Services.

- **Certificate enrollment and renewal.** Beyond the basic process associated with the creation of a new certificate, there can be significant differences in how enrollment and renewal are handled, depending on whether a stand-alone or enterprise CA is being used, whether manual or autoenrollment is used, and which versions of the operating system are involved.

- **How certificates are revoked.** In Windows Server 2003, certificate revocation and management have been enhanced with the addition of delta CRLs.

- **How key pairs are stored.**

  **Note**

  - Another important closely related process, certificate validation, is described in "How Certificates Work."

**How Certificates Are Created**

Certificates are created only after a certificate request has been successfully processed. The process of validating or denying a certificate request involves all of the key components of Windows Server 2003 Certificate Services:

- The Certificates Services engine
- Policy module
- Exit module
- Certificates database

In addition, significant roles are played by the client requesting the certificate, CryptoAPI, and, typically, a directory such as Active Directory. The following figure illustrates these components and their relation to each other.

**Components Involved in Certificate Creation**




The following steps are involved in creating a certificate:

**Note**

- This process applies to most common PKCS #10 requests or Certificate Management protocol using CMS requests.

1. When a user initiates a certificate request, Enrollment Control (Xenroll.dll) uses a cryptographic service provider (CSP) on the client's computer to generate a public key and private key pair for the user.

2. After a key pair has been generated, Xenroll builds a certificate request based on a certificate template.

3. The user's public key is sent with the user's identifying information to the Certificate Services engine.

4. Meanwhile, a copy of the request is placed in the request folder for the requester. In the case of a user, this folder is:

   C:\Documents and Settings*username*\Application Data\Microsoft\SystemCertificates\Request\

5. The CA authenticates the user.

6. The Certificate Services engine passes the request to the policy module, which might perform an additional access check

by querying Active Directory or another reliable source such as IIS.

7. The policy module also checks the registry for a list of certificate templates for the CA. The policy module checks the list for version numbers to verify that it has the most recent versions of all certificate templates. If the request references a more recent version of the certificate template than is currently available, the request fails.

8. If the access check is successful, the Certificate Services engine creates a new row in the Request and Certificates table of the certificates database. A request ID is entered into the database to help track the request.

9. Additional attributes of the certificate request, such as the extensions associated with the request, are entered into the appropriate fields in the database.

10. The Certificate Services engine validates the digital signature used to sign the certificate request. This validation process involves comparing all the certificates in the chain leading up to the root CA to a reliable source, such as Active Directory or IIS. This validation process also involves confirming that none of the certificates in the chain has been revoked.

11. The policy module applies a certificate template to fill out the data associated with the request. This process adds information such as usage restrictions and validity periods, as well as subject name format, to the data in the original request. Some of this data is fixed and therefore copied directly from the template, some comes from the certificate request, and some is constructed on a per user basis (concatenating the friendly name of the user into common name format, for example).

    **Note**

    - When a certificate is being renewed, data that is not being modified is simply reused.

12. If a certificate template has been set to pending, an administrator must verify related information in the request — such as the identity of the requester — which is not controlled by the policy module.

13. The policy module returns the request to the Certificate Services engine, noting whether the request should be issued, should be denied, or is pending.

14. If the request has been approved, the Certificate Services engine constructs the certificate with the appropriate subject, validity period, and extensions.

15. The Certificate Services engine encodes, signs, and validates the certificate from information in the database to verify that the certificate is valid for the entire certificate chain and all defined purposes, and performs another revocation check before it issues the certificate.

16. The issued certificate is sent to the exit module, which saves the certificate to the certificates database and, depending on the template configuration, publishes it to the directory service.

    **Note**

    - If this is a local forest, the certificate is published to the user object. If this is a forest trust, the certificate is published to the contact object.

17. The exit module constructs and signs a response, which is returned to Xenroll.dll on the client. This response can include:

    - A CMS response.

    - A certificate BLOB.

    - A message indicating whether the certificate request has been approved, has been denied, or is pending.

- An error message.
- Request ID.

Xenroll places the certificate in the following folder:

C:\Documents and Settings*username*\Application Data\Microsoft\SystemCertificates\My\Certificates

**Certificate Enrollment and Renewal**
Certificate enrollment and renewal in PKIs based on Windows Server 2003 Certificate Services varies depending on:

- The types of enrollment.
- The types of clients.
- Optional components involved in the enrollment and renewal process.

**Types of Enrollment**
Certificate Services supports three certificate enrollment and renewal methods:

- Autoenrollment
- Manual
- Custom

The methods used to process certificate enrollment and renewal will vary depending on the level of verification the CA must perform to authenticate the certificate requester's identity.

Some enrollment and renewal requests can be approved or denied quickly by configuring the CA to simply verify whether the requester has a valid e-mail address or is a member of a domain.

If the certificate request involves a high-security application, on the other hand, an authorized enrollment agent might have to verify the requesters identity through other agencies or by conducting background checks, In such situations, manual certificate enrollment and renewal processes would be appropriate.

The enrollment and renewal options available to you will depend on the types of clients and CAs that you have.

**Autoenrollment**
Autoenrollment provides a quick and simple way to issue user and computer certificates and to benefit from applications that can use PKI.

Automatic certificate enrollment and renewal can be configured using the Automatic Certificate Request Setup Wizard (only for Windows Server 2003 and Windows XP user and computer certificates). Certificate autoenrollment also works in a Windows Server 2003 Terminal Services session if you use a Remote Desktop Protocol (RDP) 5.1 client.

On computers that are running Windows XP, you can automatically enroll users and computers for certificates, including smart card–based certificates. In contrast, the Windows 2000 Server family only supports certificate autoenrollment for computer certificates.

Certificate autoenrollment is based on the combination of Group Policy settings and version 2 certificate templates. This combination allows certificate enrollment and renewal in the background for computers and users at any time when you apply Group Policy.

To perform autoenrollment, the certificate requester must be registered and authenticated as either a user or computer in Active Directory.

You can configure computer autoenrollment for computer, domain controller, and IPSec certificates.

Certificate autoenrollment occurs at the end of the logon process. It also occurs every eight hours while a user is logged on. Users can also manually activate or cancel autoenrollment in the Certificates snap-in. When activated, autoenrollment

determines through Active Directory queries whether any certificates need to be autoenrolled for the user or computer. If the computer is not networked, autoenrollment is not activated.

When autoenrollment is activated, Certificate Services first checks for any pending autoenrollment requests. If there are any pending requests, an enterprise CA can return a pending status with the request identifications. Autoenrollment attempts to complete pending request identifications by querying the CA that designated the requests as pending. If the certificate is denied, autoenrollment fails the request and an error is generated in the CA event log. If the certificate still has not been issued after a specified period of time — configurable in the registry — then the request is considered to have failed. If a certificate is issued, then it is added to the following folder:

Documents and Settings*username*\Application Data\Microsoft\SystemCertificates

Certificates that are automatically issued are also automatically renewed from the enterprise CA that issued the certificates. Certificates can be manually renewed with the Certificate Renewal Wizard or through the Certificate Services Web Enrollment Support pages.

You can strengthen the security of autoenrollment by:

- Specifying in the certificate template that an additional signature with a specific OID must be included in the request.
- Requiring that a smart card be used to complete the autoenrollment process.
- Configuring the CA to place incoming requests in a pending queue and defining CA officers to verify and approve pending requests.

**Manual enrollment**
Manual enrollment can be configured using the Certificate Request and Certificate Renewal wizards, or using Windows Server 2003 Certificate Services Web Enrollment Support pages.

**Certificate Request and Certificate Renewal wizards**
You can request or renew certificates for Windows Server 2003 and Windows XP users and computers by using the Certificate Request Wizard and the Certificate Renewal Wizard that are available in the Certificates snap-in. These wizards do not function unless an enterprise CA is online to process and issue certificate requests and renewal requests. The ACLs for the certificate templates determine which user accounts or computer accounts can enroll for the various types of certificates.

You can request or renew certificates with an existing private and public key set or request that certificates be issued or renewed with a new private and public key set. However, do not renew certificates with the same private and public key sets if the maximum safe key lifetime would be exceeded.

**Manual enrollment using Web Enrollment Support**
You can have users submit their certificate requests by using a Web page. When a request is received by a CA, the CA verifies whether the requester is qualified to receive a certificate and then either approves or denies the request.

The Windows Server 2003 Certificate Services Web Enrollment Support pages are composed of Active Server Pages and ActiveX controls that provide a Web-based user interface to a CA. By default, the Web Enrollment Support pages are automatically installed on the computer where the CA is installed, but you can also install the Web Enrollment Support pages on another computer running Windows Server 2003.

You can use the Web Enrollment Support pages to perform the following tasks:

- Request and obtain a basic user certificate.
- Request and obtain other types of certificates by using advanced options.
- Request a certificate by using a certificate request file.
- Renew certificates by using a certificate renewal request file.
- Save a certificate request to a file.
- Save the issued certificate to a file.

- Check on pending certificate requests.
- Retrieve the CA's certificate.
- Retrieve the latest CRL from the CA.
- Enroll for smart card certificates on behalf of other users (for use by trusted administrators).

The Web Enrollment Support pages that are installed for stand-alone CAs are similar to the pages that are installed for enterprise CAs, but they differ in that stand-alone CAs do not use certificate templates. For stand-alone CAs, all information about the certificate, including information about the requester, must be specified in the certificate request. The Web Enrollment Support pages for stand-alone CAs support a number of types of certificates that have much of the same functionality as certificate types that are based on templates.

You can deploy stand-alone CAs and Web Enrollment Support pages to issue most of the types of certificates that enterprise CAs can issue. However, certificates for logging on by using smart cards and for autoenrollment require an enterprise CA to issue and renew the certificates.

**Custom enrollment and renewal applications**
The standard enrollment and renewal methods that are available in Windows Server 2003 can meet a wide range of needs. However, if you have special needs, you can develop custom certificate enrollment and renewal applications. Windows Server 2003 Certificate Services supports industry-standard certificate requests by using remote procedure call (RPC) requests or HTTP requests. You can develop custom applications that make certificate requests to Certificate Services. For more information about developing custom applications with Windows Server 2003 Certificate Services, see "Programming Certificate Services" in the Cryptography collection of the Microsoft Platform SDK on MSDN.

**Client Requirements for Enrollment**
You can enroll version 2 templates with any computer running Windows XP or later through the default enrollment methods, including the Certificates snap-in, autoenrollment, Web Enrollment Support pages, or command-line tools.

A computer running Windows 2000 cannot use the Certificates snap-in to enroll a version 2 template. However, any client that is running Microsoft Internet Explorer 5.01 or later can use a version 2 template to enroll certificates through the Web enrollment methods and a downloaded ActiveX control. To download the ActiveX control on a client computer, it is necessary to log on as an administrator or power user on the local computer. In addition, clients running Windows 2000 can enroll version 2 templates through a connection to a server running Windows Server 2003 that has Terminal Services enabled.

### Note

- An enrollment agent that enrolls certificates that are based on version 2 templates requires either a Windows XP or Windows Server 2003 enrollment station. There is no support to enroll version 2 templates with an enrollment agent on a Windows 2000 enrollment station. Nevertheless, certificates that are based on version 2 templates that have been enrolled through an enrollment agent on either a Windows XP or Windows Server 2003 enrollment station can be used on a Windows 2000 client computer.

If certificates have been enrolled with a Windows 2000 PKI where only version 1 templates were available, there is no immediate need to re-enroll or renew these certificates with version 2 templates.

The following three tables list different enrollment methods that are supported on computers that are running Windows 2000, Windows XP, or the Windows Server 2003 family. You can use scripted enrollment with support of CAPICOM and Xenroll. For more information, see "Programming Certificate Services" in the Cryptography collection of the Microsoft Platform SDK on MSDN.

**Certificate Request Wizard Enrollment**

| | **Version 1 template** | **Version 2 template** |
|---|---|---|
| Self enrollment on a Windows 2000 workstation | Yes | No |
| Self enrollment through a Windows Server 2003 Terminal Server session | Yes | Yes |
| Self enrollment on a Windows XP workstation | Yes | Yes |

| | | |
|---|---|---|
| Enrollment agent on a Windows 2000 workstation | No | No |
| Enrollment agent through a Windows Server 2003 Terminal Server session | No | No |
| Enrollment agent on a Windows XP workstation | No | No |

**Web Enrollment**

| | **Version 1 template** | **Version 2 template** |
|---|---|---|
| Self enrollment on a Windows 2000 workstation | Yes | Yes |
| Self enrollment through a Windows Server 2003 Terminal Server session | Yes | Yes |
| Self enrollment on a Windows XP workstation | Yes | Yes |
| Enrollment agent on a Windows 2000 workstation | Yes | No |
| Enrollment agent through a Windows Server 2003 Terminal Server session | Yes | Yes |
| Enrollment agent on a Windows XP workstation | Yes | Yes |

**Scripted Enrollment**

| | **Version 1 template** | **Version 2 template** |
|---|---|---|
| Self enrollment on a Windows 2000 workstation | Yes | Yes |
| Self enrollment through a Windows Server 2003 Terminal Server session | Yes | Yes |
| Self enrollment on a Windows XP workstation | Yes | Yes |
| Enrollment agent on a Windows 2000 workstation | Yes | Yes |
| Enrollment agent through a Windows Server 2003 Terminal Server session | Yes | Yes |
| Enrollment agent on a Windows XP workstation | Yes | Yes |

**Note**

- Because a PKI is a forest resource, the Active Directory site structure is not taken into account when any kind of certificate is requested and issued. An Active Directory–integrated certificate requester enumerates all registered enrollment services in Active Directory and sends its request to a CA that can enroll the certificate type that the user wants. The client does not necessarily choose the closest CA, from a network perspective. Because of this, you should verify that the CA deployment ensures that any client has reliable network connectivity with a CA.

#### Client Requirements for Renewal

When a certificate comes to the end of its lifetime, it must be renewed or replaced to ensure that the certificate owner is able to continue with the certificate's purpose.

In certificate renewal, the renewal requester already owns a certificate. The renewal takes the information of the existing certificate into account when the renewal request is submitted. A certificate can either be renewed with a new key or the existing key can be used for the renewed certificate.

If a certificate was enrolled with a version 2 template, it cannot be renewed if it was based on a version 1 template. However, a certificate that was enrolled with a version 1 template can be renewed with a certificate that was made from a version 2 template.

The following three tables list the different certificate renewal methods that are supported on computers that are running Windows 2000, Windows XP, or the Windows Server 2003 family.

**Certificate Renewal Wizard Enrollment**

| | Version 1 templates | Version 2 templates |
|---|---|---|
| Self renewal on a Windows 2000 workstation | Yes | No |
| Self renewal on a Windows 2000 workstation with Windows Server 2003 Terminal Server session | Yes | Yes |
| Self renewal on a Windows XP workstation | Yes | Yes |
| Renewal with an enrollment agent on a Windows 2000 workstation | No | No |
| Renewal with an enrollment agent on a Windows 2000 workstation with Windows Server 2003 Terminal Server session | No | No |
| Renewal with an enrollment agent on a Windows XP workstation | No | No |

**Web Enrollment**

| | Version 1 templates | Version 2 templates |
|---|---|---|
| Self renewal on a Windows 2000 workstation | No | No |
| Self renewal on a Windows 2000 workstation with Windows Server 2003 Terminal Server session | No | No |
| Self renewal on a Windows XP workstation | No | No |
| Renewal with an enrollment agent on a Windows 2000 workstation | No | No |
| Renewal with an enrollment agent on a Windows 2000 workstation with Windows Server 2003 Terminal Server session | No | No |
| Renewal with an enrollment agent on a Windows XP workstation | No | No |

**Scripted Enrollment**

| | Version 1 templates | Version 2 templates |
|---|---|---|
| Self renewal on a Windows 2000 workstation | Yes | Yes |
| Self renewal on a Windows 2000 workstation with Windows Server 2003 Terminal Server session | Yes | Yes |
| Self renewal on a Windows XP workstation | Yes | Yes |
| Renewal with an enrollment agent on a Windows 2000 workstation | No | No |
| Renewal with an enrollment agent on a Windows 2000 workstation with Windows Server 2003 Terminal Server session | No | No |
| Renewal with an enrollment agent on a Windows XP workstation | Yes | Yes |

**Enrollment and Renewal Using Optional Components**

Certificate enrollment and renewal rarely take place without the use of optional components, including Active Directory or another directory service, IIS, or another client acting as an enrollment agent. The following sections discuss the enrollment and renewal process for enterprise and stand-alone CAs where these optional components are being used.

**Enterprise direct enrollment**

The most basic enrollment method in the enterprise mode is direct enrollment. This is where the client makes a DCOM call directly to Certificate Services and passes the request over this call.

The client will have already retrieved any templates from Active Directory, and will have created the request, using an appropriate template. The following figure shows this configuration and the relationships between Certificate Services, Active

Directory, and the client.

**Enterprise Direct Enrollment**




**Enterprise indirect enrollment with enrollment agent**

Another likely enrollment method in enterprise mode is for clients to use an enrollment agent. An enrollment agent is a person who is authorized with a special Enrollment Agent certificate to request certificates on behalf of a client. Most uses of enrollment agents will be for smart card clients, which can store the private key off the computer. The following figure shows this configuration and the relationship between the enrollment agent, Certificate Services, and Active Directory.

**Enterprise Enrollment with Enrollment Agent**




Before any enrollment occurs, the enrollment agent will have retrieved any certificate templates from the registry of the server that hosts the CA. When a user needs to enroll, the enrollment agent uses the certificate template to create a certificate request and signs the request with the Enrollment Agent certificate. The enrollment agent computer then makes an authenticated DCOM call to Certificate Services. The enrollment agent then sends the certificate request over this DCOM call.

Unlike in the Web enrollment configurations, the enrollment agent computer does not impersonate the client. Instead, the DCOM call made to Certificate Services will have the enrollment agent computer as the caller name. This is important to note, because the certificate request will have the client's identifying information, which will not match the DCOM caller name. Certificate Services would normally reject this mismatch. However, the Enrollment Agent certificate used to sign the request contains the Enrollment Agent OID, which is explicitly trusted to issue certificates on behalf of others. As a result, Certificate Services will accept the request.

**Stand-alone Web enrollment**

The most common method of enrollment in the stand-alone mode is by an HTML interface, provided by a set of Web pages running on a server running IIS. This is called Web enrollment.

IIS can be installed on the same computer as Certificate Services or on a separate computer. In either case, when the client needs a certificate and must enroll in a hierarchy, it connects to IIS by HTTP and authenticates to the server in a manner required by IIS. The client then creates a certificate request and sends it over the HTTP connection to the server by way of a remote DCOM interface. This HTTP connection is most often secured by Transport Layer Security (TLS). IIS then receives the

request and passes it by way of a local DCOM interface to Certificate Services.

The following figure shows this configuration and the relationship between IIS, Certificate Services, and the client.

**Stand-alone Web Enrollment**




**Enterprise Web enrollment**
Enterprise mode Web enrollment also can be configured with IIS and Certificate Services running on the same computer or on separate computers.

Before enrollment occurs, the client will have retrieved any certificate templates from Active Directory. When the client needs to enroll, it will make an HTTP connection to IIS and authenticate in some manner. The computer running IIS must be trusted for delegation in Active Directory. This allows IIS to impersonate the client and make an authenticated DCOM call to Certificate Services. IIS then passes the request over this DCOM call.

The following figure shows the relationship between IIS, Certificate Services, and the client.

**Enterprise Web Enrollment**




**How Certificates Are Revoked**
To help maintain the integrity of an organization's public key infrastructure (PKI), the administrator of a CA must revoke the subject's certificate if the subject leaves the organization, if the subject's private key has been compromised, or if some other security-related event dictates that it is no longer desirable to have the subject's certificate considered valid. When a certificate is revoked by a CA, it is added to that CA's certificate revocation list (CRL).

CRLs are complete, digitally-signed lists of unexpired certificates that have been revoked. The CRL is retrieved by clients who can then cache it (based on the configured lifetime of the CRL) and use it to verify certificates presented for use.

Windows Server 2003 Certificate Services supports industry standard X.509 version 2 CRLs. Each CA maintains a CRL for the certificates it issues and publishes the CRL to CRL distribution points, which can include Web pages, network shares, or Active Directory. An X.509 version 3 certificate usually contains the CRL distribution point for its issuing CA.

**Note**

- Non-Microsoft products can use CRLS published by CAs running on Windows operating systems to provide revocation information in other formats, such as Online Certificate Status Protocol (OCSP), Simple Certificate Validation Protocol

(SCVP), and XML Key Management Specification (XKMS).

By default, enterprise CAs publish CRLs once per week to Active Directory. By default, stand-alone CAs publish CRLs once per week to the following folder on the CA server:

*systemroot*\Certsrv\Certenroll

**Delta CRLs**
Microsoft Windows Server 2003 can publish large CRLs, generating large amounts of directory replication traffic. For a CA with a large population of users, this list can become as large as 1 MB. Many applications require up-to-date revocation status information, which requires the CA to publish the CRLs frequently. Even when there are no new revocation changes, the CA has to republish the entire list so that applications have the latest information.

For this reason, a new type of CRL called a delta CRL has been introduced with Windows Server 2003 Certificate Services. Delta CRLs are CRLs that contain the list of changes in revocation status that have occurred since the last base CRL was published.

**Note**

- Delta CRLs are not necessarily relative to the last published CRL, but always relative to a designated base CRL. Delta CRLs are not aggregated, and the client always uses a pair of the latest delta CRL and a valid base CRL.

It is important that a relying party (client) has access to both a base and delta CRL. If the base CRL is not available for the client's, an error will occur with the delta CRL's processing. When an error occurs, the client falls back to the base CRL.

**Reasons for Revocation**
When a certificate is revoked, a certificate issuer can specify why the certificate was revoked. This is done by specifying a revocation reason code. The reason codes are defined by RFC 3280 and include:

- **KeyCompromise.** The token or disk location where the private key associated with the certificate has been compromised and might be in the possession of an unauthorized individual. This can include cases where a laptop is stolen or a smart card is lost.

- **CACompromise.** The token or disk location where the CA's private key is stored has been compromised and might be in the possession of an unauthorized individual. When a CA's private key is revoked, all certificates issued by the CA that have been signed using that private key are considered revoked, as well.

- **AffiliationChanged.** The user has terminated his or her relationship with the organization indicated in the **DistinguishedName** property of the certificate. This reason code is typically used when an individual has been terminated or has resigned from an organization. You do not have to revoke a certificate when a user changes departments, unless your security policy requires that a different certificate be issued by a departmental CA.

- **Superseded.** A replacement certificate has been issued to a user, and the reason does not fall under the previous reasons. This reason code is typically used when a smart card fails, the user forgets a password, or the user has changed their legal name.

- **CessationOfOperation.** If a CA is decommissioned, the CA's certificate should be revoked with this reason code. Do not revoke the CA's certificate if the CA no longer issues new certificates, yet still publishes CRLs for the currently issued certificates.

- **CertificateHold.** A temporary revocation that indicates that a CA will not vouch for a certificate at a specific point in time. After a certificate is revoked with a CertificateHold reason code, the certificate can then be revoked with another reason code, or unrevoked and returned to use.

  **Note**

  - Although CertificateHold allows a certificate to be unrevoked, it is not recommended to place a hold on a certificate, as it becomes difficult to determine if a certificate was valid for a specific time.

- **RemoveFromCRL.** If a certificate is revoked with the CertificateHold reason code, it is possible to unrevoke a certificate. The unrevoking process still lists the certificate in the CRL, but with the reason code set to RemoveFromCRL.

**Note**

- This is specific to the CertificateHold reason and is only used in delta CRLs.

- **Unspecified.** You can revoke a certificate without providing a specific reason code. Although it is possible to revoke a certificate with the Unspecified reason code, this is not recommended, because it does not provide an audit trail as to why a certificate is revoked.

**The Revocation Checking Process**

Prior to checking the status of a certificate, the client checks the certificate to ensure that the certificate is trusted and valid. For more information about how certificate trust is established, see How Certificates Work.

Every issued certificate has a defined period during which the issuing CA will vouch for the validity of the certificate. The validity period is defined in two fields in the issued certificate:

- **NotBefore.** This field defines the date and time on which the certificate's validity period begins.

- **NotAfter.** This field defines the last date on which the certificate is considered valid by the issuing CA. The NotAfter date cannot be set to a date later than the NotAfter date defined in the issuing CA's certificate.

If CRL-based status checking is used and the certificate is found to have a valid time, the presented certificate is examined to see if it contains a CRL distribution point extension. The CRL distribution point extension indicates both the protocol that must be used to retrieve the CRL (HTTP, FTP, LDAP, or FILE) and the location where the CRL is stored (represented as a URL).

Using this information, CryptoAPI searches the local certificate stores and the local cache for any CRL signed by the issuer (that is the CA) of the certificate that is to be validated. A cached version of a current CRL will always be used (so long as it is valid), rather than downloading the same CRL again. The following logic is used to evaluate the CRL:

- If a CRL is found and the certificate's serial number is listed in the CRL, then the certificate will be considered revoked.

- If the CRL is expired and the certificate is listed in the CRL with any reason other than CertificateHold, the certificate will be considered to be revoked and no attempt will be made to retrieve a new CRL.

- If the certificate is not listed in the CRL, or if the reason code is CertificateHold**,** then a new CRL will be retrieved from the URLs listed in the certificate's CRL distribution point extension. The new CRL is fetched only if it is past the NextUpdate field in the currently held CRL. The new CRL is checked to determine if the certificate is revoked. If the original reason was CertificateHold, the CRL is checked to determine if the certificate is unrevoked by looking for the RemoveFromCRL reason code.

- If the CRL cannot be obtained, the client will generate a Server offline error.

**CRL Format**

A version 2 CRL includes the following fields:

- **Version.** The version of formatting used for the CRL.

- **Issuer.** The issuing CA's distinguished name (DN) represented using an X.500 DN such as CN=TestCA,OU=PKI,DC=nwtraders,DC=com.

- **This Update.** The date that this CRL was issued. Universal Time Coordinate (UTC) format is used for dates up to the year 2049. Generalized Time format is used for dates of 2050 and beyond.

- **Next Update.** The date at which the next CRL will be issued. The same date formatting rules are used as for This Update.

- **Revoked Certificates.** A sequence of zero or more revoked certificates with the following fields represented for each revoked certificate.

- **Certificate Serial Number.** The serial number assigned by the issuing CA for each revoked certificate.
- **Revocation Date.** The date at which the revocation took place.
- **CRL Entry Extension.** An optional field that contains version 2 extensions for CRLs that provides additional information about a single CRL entry.
- **ReasonCode.** Identifies the reason for the certificate revocation. Reasons can include: Unspecified, KeyCompromise, CACompromise, AffiliationChanged, Superseded, CessationOfOperation, CertificateHold, and RemoveFromCRL.
- **HoldInstructionCode.** This option, only defined in the PKI x.509 format, indicates the action to be taken if the CertificateHold is entered in the ReasonCode field. This extension is not supported by a Windows Server 2003 CA.
- **InvalidityDate.** The date on which it is known, or suspected, that the private key associated with the certificate was compromised.
- **CertificateIssuer.** An optional field used to indicate the CA that issued the certificate if indirect CRLs are implemented. A Windows Server 2003 CA does not use this extension, but it is supported by Windows XP clients.
- **AuthorityKeyIdentifier.** A numeric representation of the issuer name and serial number from the CRL issuer's certificate that is used in the certification chain building process.
- **IssuerAltName.** An optional field used to associate alternative identities with the issuer. This can contain a DNS name or a Uniform Resource Indicator (URI). A Windows Server 2003 CA does not use this extension, but it is supported by Windows XP clients.
- **CRLNumber.** A monotonically increasing sequence number that indicates the version of the CRL since the first CRL was published.
- **IssuingDistributionPoint.** Identifies the CRL distribution point if CRL partitioning is implemented. For example, the CRL can be partitioned so that end-entity, CA certificate, or specific reason code revocations are published at a specific distribution point. A Windows Server 2003 CA does not use this extension, but it is supported by Windows XP clients.

  **Note**

  - Windows Server 2003 does not support partitioning CRLs by reason code.
- **deltaCRLIndicator.** Indicates that the CRL is a delta CRL. This also specifies the minimum CRL number of base CRL with which the delta CRL can be combined to obtain the complete state of the CA.
- **Signature (Issuer Signature Algorithm).** Identifies the algorithm used to certify the CRL.
- **SignatureValue.** The actual digital signature computed using the indicated Issuer Signature Algorithm on the ASN.1 DER encoded CRL. The value is encoded as a bit string and omits the signature algorithm and SignatureValue fields.

#### CRL Distribution Point Extensions

When a CA issues a certificate, it can place a CRL distribution point extension in the new certificate that points to the relevant CRL for the certificate containing the extension only.

For example, the root CA issues certificates with a CRL distribution point that points only to the root CA's CRL. A first level subordinate CA issues certificates with a CRL distribution point that points only to the first level subordinate CA's CRL. Finally, a second level CA issues certificates with a CRL distribution point that points only to the second level CA's CRL.

**Note**

- For more information, see Trusted root CA and CRL distribution point containers earlier in this subject.

When the default revocation provider is used, the full certificate chain must be built first, then each certificate in the chain is searched for a CRL distribution point extension to fetch a CRL that will prove the revocation status of that one certificate only. The Windows client does not support a certificate status model where one CA can supply information about revoked certificates that were issued by a different CA. This is because the key that was used to sign the certificate in question must be the same key used to sign the CRL that is used to verify revocation status of the certificate.

The CRL distribution point might point to both primary and alternative locations for CRL retrieval. When a client must acquire the latest CRL, the CRL distribution point extension in a certificate is checked and the CRL retrieval is attempted by contacting the URLs in the CRL distribution point extension. If the name resolution of a fully qualified domain name fails, or the CRL is unavailable, the client tries the next CRL distribution point from the CRL distribution point extension field.

**Note**

- If the client is able to resolve the name of the CRL distribution point but the CRL is not physically available, it might take a significant period of time before the timeout threshold is reached. Windows XP provides enhancements in reducing the timeout threshold as well as CRL unavailability detection.

If the CRL distribution point extension is not available in a certificate, then CryptoAPI will only check the local certificate store and cache for a CRL. If a local CRL is available, it will be checked despite the absence of a CRL distribution point extension. However, a specific application must specify whether or not to demand a revocation check of a certificate.

**CRL Publication**

A CRL distribution point URL can specify several protocols for CRL access. These include:

- **HTTP.** HTTP websites can store the updated CRLs.
- **LDAP.** A Lightweight Directory Access Protocol database can contain the updated CRLs.
- **File.** CRLs can be accessed using the Server Message Block (SMB) protocol.
- **FTP.** File Transfer Protocol servers can be used to distribute updated CRLs. The CRL, by default, is published in two locations by a Microsoft enterprise CA:
    - http://*CAName/*certenroll/*CRLName*
    - LDAP:///CN=*CAName*,CN=*CAComputerName*,CN=CDP,CN=Public Key Services,CN=Services,CN=Configuration, DC=*ForestRootDomain*,DC=*TLD*

When a CRL is published, the CRL will take on different names, depending on whether the CA's certificate has been renewed, and whether the existing private keys are re-used when the certificate is renewed. The following table outlines some examples of CRL names that are automatically generated for a Windows Server 2003 CA.

**Auto-generated CRL File Names**

| Scenario | Name of the CRL File | Name of Delta CRL File | CA Version |
|---|---|---|---|
| A certification authority named IssuingCA that has never had its CA certificate renewed | issuingca.crl | issuingca+.crl | 0.0 |
| A certification authority named IssuingCA that has been renewed once with the same key | issuingca.crl | issuingca+.crl | 0.0 |
| A certification authority named IssuingCA that has been renewed once with a new key | issuingca.crl<br>issuingca(1).crl | issuingca+crl<br>issuingca(1)+.crl | 2.2 |
| A certification authority named IssuingCA that has been renewed three times, twice with a new key | issuingca(3).crl | issuingca(3)+.crl | 3.3 |

**How Key Pairs Are Stored**

Windows Server 2003 and Windows XP store a user's public key certificates in the user's personal certificate store. Certificates

are stored in plaintext because they are public information, and they are digitally signed by certification authorities to protect against tampering.

User certificates for each user profile are located in the following folder:

*systemroot*\Documents and Settings*username*\Application Data\Microsoft\SystemCertificates\My\Certificates

These certificates are written to the users personal store in the system registry each time the user logs on to the computer. If roaming profiles are being used, users' certificates are stored on the file server where the roaming user profile is stored

Private keys for the Microsoft RSA-based cryptographic service providers (CSPs), including the Base CSP and the Enhanced CSP, are located in the user profile under:

*systemroot*\Documents and Settings*username*\Application Data\Microsoft\Crypto\RSA

In the case of a roaming user profile, private keys reside in the RSA folder on the file server where the roaming user profile is stored and are downloaded to the user's computer until the user logs off or the computer is restarted.

Because private keys must be protected, all files in the RSA folder are automatically encrypted by using a random symmetric key called the user's master key. The user's master key is 64 bytes in length and is generated by a strong random number generator. 3DES keys are derived from the master key and are used to protect private keys. The master key is generated automatically and is periodically renewed. It encrypts each file in the RSA folder automatically as the file is created.

The Data Protection API automatically encrypts the user's master key or keys. Master keys are stored in the user profile under

*systemroot*\Documents and Settings*username*\Application Data\Microsoft\Protect

For a domain user who has a roaming profile, the master key is located in the user's profile and is downloaded to the user's profile on the local computer until the computer is restarted.

**Certificate Services Key Storage**
Certificate Services must maintain a key pair for signing certificate requests. The public part of that key pair resides in a certificate located in the My folder of the local computer. The private part of that key pair resides in the cryptographic service provider (CSP) containers.

As shown in the following figure, Certificate Services can have multiple certificates that contain the same public key and, thus, use the same private key. The reasons for this might be because of certificate expiration or varying certificate attributes. There can be up to a maximum of 65,536 Certificate Services certificates.

Each private key has only one active base CRL and one active delta CRL that contains revoked certificates that were signed by that key. Expired CRLs can be automatically deleted upon CRL publication if there is an existing, non-expired CRL for that private key.

### Relationship of Certificates, Private Keys, and CRLs




**Key Archival and Key Recovery**
Subjects can lose their private keys in a variety of ways, such as accidental deletion or deliberate misuse. An administrator might also want to recover the key of a particular subject to access data protected by that key. Key recovery can be used whenever the key archival process has stored the subject's private key.

Windows Server 2003, Enterprise Edition can be configured to archive the private key of a specific certificate when the certificate is issued. This private key archive allows the key to be recovered at a later time if the private key is lost. This process is implemented in two separate phases: key archival and key recovery.

The following sections detail how the key archival and key recovery processes work in Windows Server 2003.

**Key Archival Process**

1. The client finds a CA in the Active Directory (enrollment services container in the configuration partition) and makes an authenticated DCOM connection to the CA.

2. The client requests the CA's encryption certificate.

3. The CA sends the encryption certificate to the client.

4. The client validates that the CA's encryption certificate has been signed by the same key as the CA signing certificate and performs a revocation status check. This ensures that only the intended CA can decrypt the certificate request containing the private key.

5. The client encrypts the private key corresponding to the request with the CA encryption certificate public key, builds a Certificate Management protocol using CMS (CMC) request, and sends the full CMC certificate request to the CA.

6. The CA validates that the encrypted private key cryptographically pairs with the public key in the certificate request.

7. The CA validates the signature on the request with the public key in the request.

8. The CA encrypts the private key from the user request with a random 3DES symmetric key and then encrypts the symmetric key with one or more key recovery agent public keys.

9. The CA saves the encrypted key BLOB containing the encrypted private key and the symmetric key encrypted with one or more key recovery agent public keys to the CA database.

10. The CA processes the certificate request normally.

11. The CA responds to the client with a full PKI response in CMC.

    **Note**

    - Only RSA encryption keys can be archived in the CA database. Signature only keys as well as non-RSA key pairs will not be archived. Denied and resubmitted requests will also not archive private keys.

**Key Recovery Process**

Key recovery is initiated by an end-entity that has lost access to a private key. The end-entity is typically a user or an administrator of a computer who notifies a key recovery agent about lost private key material. The key recovery agent must have a key recovery agent certificate that was used to protect the symmetric key on the private key of the user. The key recovery agent can also be a Certificate Manager (an officer on the CA that issued the certificate in question to the user).

**Note**

- Each organization should determine whether key recovery agent and Certificate Manager should be combined into one role or should be separate roles.

The actual key recovery process first retrieves the PKCS #7–encrypted recovery BLOB from the database. The encrypted BLOB is protected by ACLs to ensure that only a valid Certificate Manager is able to copy the BLOB from the database. Encrypted BLOBs in the database contain the issuer name and serial number of each key recovery agent certificate for key recovery agent identification purposes during recovery.

The recovery BLOB consists of wrapping the encrypted PKCS #7 certificate in the database in another signed PKCS #7 certificate. This enables a number of certificates to be included in the recovery BLOB. The returned certificates include the full chain of the user certificate that is to be recovered, the chain of the signing CA certificate (which might differ from the CA certificate under which the user certificate was issued), and the key recovery agent certificates to which the key was encrypted.

**Related Information**

The following resource contains additional information that is relevant to this section:

- Certificates Technical Reference

# Certificate Services Tools and Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In this section**


- Certificate Services Management Tools
- Certificate Services Registry Settings
- Certificate Services Group Policy Settings


**Certificate Services Management Tools**
The following tools are associated with Certificate Services.

**Certreq.exe: Certreq**

**Category**
Certreq ships with the Windows Server 2003 operating system tools and with the Windows Server 2003 Adminpak.

**Version Compatibility**
Certreq is compatible with Windows Server 2003 and Windows 2000 Server, and can be used to manage the certificate containers for users, computers, and services on computers running Windows 2000, Windows XP, and Windows Server 2003.Certreq enables you to submit, retrieve, create, and accept certificate requests that are sent to a Windows Server 2003 CA. You can also use Certreq to create and sign cross-certificate requests. You can also place the Certreq command syntax in a batch file to script certificate requests.

To find more information about Certreq, see “Command Line References” in Tools and Settings Collection.

**Certutil.exe: Certutil**

**Category**
Certutil is a command-line program that is installed as part of Certificate Services.

**Version Compatibility**
Certutil can be used on Windows 2000 Server CAs and Windows Server 2003 CAs.

You can use Certutil to extract and display CA configuration information, configure Certificate Services, back up and restore CA components, and verify certificates, key pairs, and certificate chains.

To find more information about Certutil, see “Command Line References” in Tools and Settings Collection.

**Certsrv.msc: Certification Authority Snap-in**

**Category**
The Certification Authority snap-in ships with Windows 2000 Server and Windows Server 2003.

**Version Compatibility**
The Certification Authority snap-in is compatible with Windows Server 2003 and Windows 2000 Server.

The Certification Authority snap-in can be used to manage multiple CAs. It allows you to perform a variety of administrative tasks including:

- Starting and stopping the CA.
- Backing up and restoring the CA.
- Changing exit and policy modules.
- Viewing the CA certificate.

- Installing or reinstalling a CA certificate for the CA.
- Setting security permissions and delegating administrative control for the CA.
- Revoking certificates.
- Viewing or modifying certificate revocation list (CRL) distribution points.
- Publishing CRLs and scheduling CRL publication.
- Configuring the types of certificates that are to be issued by the CA.
- Viewing information about certificates that have been issued.
- Viewing information about certificates that have been revoked.
- Viewing pending certificate requests.
- Approving or denying pending certificate requests.
- Viewing failed certificate requests.
- Renewing the CA's certificate.

To find more information about the Certification Authority snap-in, see "Use the Certification Authority snap-in" on Microsoft TechNet.

**Certificate Services Registry Entries**
Registry settings in the following hive are associated with Certificate Services:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc

The Certificate Services-related registry entries correlate to CA and certificate configuration information that can be configured and viewed using the Certificates and Certification Authority snap-ins, as well as the Certutil command-line tool. In addition, other Certificate Services-related registry settings correlate to configuration data that can be viewed directly in certificates issued by the CA.

In addition, user root store trust options can be set on client computers in the following location:

HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\SystemCertificates\Root\ProtectedRoots

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as MMC, to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

**CertSvc**
The registry settings in this hive are all created automatically as part of Certificate Services.

**Description**

**Registry Path**
\CertSvc\Description

**Version**
Windows Server 2003 and Windows 2000 Server

This setting provides a reader-friendly description of Certificate Services. The default description reads:

"Creates, manages, and removes X.509 certificates for applications such as S/MIME and SSL. If this service is stopped, certificates will not be created. If this service is disabled, any services that explicitly depend on it will fail to start."

**DisplayName**

**Registry Path**
\CertSvc\DisplayName

**Version**
Windows Server 2003 and Windows 2000 Server

This setting provides the display name for Certificate Services.

**ImagePath**

**Registry Path**
\CertSvc\ImagePath

**Version**
Windows Server 2003 and Windows 2000 Server

This setting identifies where Certserv.exe has been installed on the computer running Windows Server 2003.

**ObjectName**

**Registry Path**
\CertSvc\ObjectName

**Version**
Windows Server 2003 and Windows 2000 Server

This setting specifies the security context used by Certificate Services. The default is the security context of the Local System account.

**Start**

**Registry Path**
\CertSvc\Start

**Version**
Windows Server 2003 and Windows 2000 Server

This setting specifies whether Certificate Services is started automatically when Windows Server is started.

**CertSvc\Configuration**
The registry settings in this hive contain information relating to the configuration of the Certificate Services service on the server.

**Active**

**Registry Path**
\CertSvc\Configuration\Active

**Version**
Windows Server 2003 and Windows 2000 Server

This setting records the sanitized name of the active CA on the server. The process of sanitizing the CA name is necessary to remove characters that are not valid for file names, registry key names, or distinguished name values, or that are not valid for technology-specific reasons.

**ConfigurationDirectory**

**Registry Path**
\CertSvc\Configuration\ConfigurationDirectory

**Version**
Windows Server 2003 and Windows 2000 Server

This setting specifies the location of the CA configuration file.

**DBDirectory**

**Registry Path**
\CertSvc\Configuration\DBDirectory

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines where the CA database is located in the root CA's file system. The CA must be able to get the appropriate path from the registry when the CA starts up.

**DBFlags**

**Registry Path**
\CertSvc\Configuration\DBFlags

**Version**
Windows Server 2003 and Windows 2000 Server

This setting contains options that control the logging behavior of the certificates database and the size of the database log.

**DBLastIncrementalBackup**

**Registry Path**
\CertSvc\Configuration\DBLastIncrementalBackup

**Version**
Windows Server 2003 and Windows 2000 Server

This setting records the time of the last incremental backup of the database.

**DBLastFullBackup**

**Registry Path**
\CertSvc\Configuration\DBLastFullBackup

**Version**
Windows Server 2003 and Windows 2000 Server

This setting records the time of the last full backup of the database.

**DBLogDirectory**

**Registry Path**
\CertSvc\Configuration\DBLogDiretory

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines where the CA's transaction log files are located in the root CA's file system. The CA must be able to get the path from the registry when the CA starts up.

**DBSessionCount**

**Registry Path**
\CertSvc\Configuration\DBSessionCount

**Version**
Windows Server 2003 and Windows 2000 Server

This setting controls the number of concurrent sessions to the certificates database that can be made. By default, the Windows Server 2003 CA allows only 20 concurrent sessions to the certificates database, which is sufficient for most operations. The CA itself can use several connections, and client enrollment requests or management tools used to view the database can account for additional sessions.

**DBSystemDirectory**

**Registry Path**
\CertSvc\Configuration\DBSystemDirectory

**Version**

Windows Server 2003 and Windows 2000 Server

This setting defines where the CA database is located in the root CA's file system. The CA must be able to get the appropriate path from the registry when the CA starts up.

**DBTempDirectory**

**Registry Path**
\CertSvc\Configuration\DBTempDirectory

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines where the temp directory for the CA database is located in the root CA's file system. The CA must be able to get the appropriate path from the registry when the CA starts up.

**LDAPFlags**

**Registry Path**
\CertSvc\Configuration\LDAPFlags

**Version**
Windows Server 2003 and Windows 2000 Server

You use this setting to control the security of Lightweight Directory Access Protocol (LDAP) connections. The default value is to sign all LDAP traffic and to not use SSL to secure LDAP connections. The alternative setting is to disable LDAP signing and to use SSL to secure LDAP connections.

**SetupStatus**

**Registry Path**
\CertSvc\Configuration\SetupStatus

**Version**
Windows Server 2003 and Windows 2000 Server

This setting documents the Certificate Services configuration options that have been installed.

**Version**

**Registry Path**
\CertSvc\Configuration\Version

**Version**
Windows Server 2003 and Windows 2000 Server

This setting identifies the version of Certificate Services installed on the system.

**CertSvc\Configuration\CAName**
The registry settings in this hive contain information relating to the configuration of the certification authority (CA) installed on the server.

**CACertFileName**

**Registry Path**
\CertSvc\Configuration*CAName*\CACertFileName

**Version**
Windows Server 2003 and Windows 2000 Server

On subordinate CAs, this setting identifies the name and location of the root certificate.

**CACertHash**

**Registry Path**
\CertSvc\Configuration*CAName*\CACertHash

**Version**
Windows Server 2003 and Windows 2000 Server

On subordinate CAs, this setting identifies the SHA-1 hashes used for all CA certificates.

**CACertPublicationURLs**

**Registry Path**
\CertSvc\Configuration*CAName*\CACertPublicationURLs

**Version**
Windows Server 2003 only

This setting identifies the URL of the authority information access point where a client can find a CA certificate. Because the authority information access point is the location of the certificate that was used to sign the certificate, and because a root CA issues the CA certificate to itself, you should leave this setting blank for the root CA.

**CAServerName**

**Registry Path**
\CertSvc\Configuration*CAName*\CAServerName

**Version**
Windows Server 2003 and Windows 2000 Server

This setting records the DNS name of the server on which the CA is running.

**CAType**

**Registry Path**
\CertSvc\Configuration*CAName*\CAType

**Version**
Windows Server 2003 and Windows 2000 Server

This setting records whether the CA is a root, enterprise, or stand-alone CA.

**CAXchgCertHash**

**Registry Path**
\CertSvc\Configuration*CAName*\CAXchgCertHash

**Version**
Windows Server 2003 and Windows 2000 Server

This setting identifies the hash used for the encryption certificates.

**CAXchgOverlapPeriod**

**Registry Path**
\CertSvc\Configuration*CAName*\CAXchgOverlapPeriod

**Version**
Windows Server 2003 and Windows 2000 Server

You use this setting to specify whether the overlap period for encryption certificates will be defined in days, weeks, months, or years.

**CAXchgOverlapPeriodUnits**

**Registry Path**
\CertSvc\Configuration*CAName*\CAXchgOverlapPeriodUnits

**Version**
Windows Server 2003 and Windows 2000 Server

You use this setting to specify the number of days, weeks, months, or years that encryption certificates will overlap.

**CAXchgValidityPeriod**

**Registry Path**
\CertSvc\Configuration*CAName*\CAXchgValidityPeriod

**Version**
Windows Server 2003 and Windows 2000 Server

You use this setting to specify whether the validity period for encryption certificates will be defined in days, weeks, months, or years.

**CAXchgValidityPeriodUnits**

**Registry Path**
\CertSvc\Configuration*CAName*\CAXchgValidityPeriodUnits

**Version**
Windows Server 2003 and Windows 2000 Server

You use this setting to specify the number of days, weeks, months, or years that encryption certificates will be valid.

**CertEnrollCompatible**

**Registry Path**
\CertSvc\Configuration*CAName*\CertEnrollCompatible

**Version**
Windows Server 2003 and Windows 2000 Server

This setting enables certificate compatibility with a legacy ActiveX control.

**ClockSkewMinutes**

**Registry Path**
\CertSvc\Configuration*CAName*\ClockSkewMinutes

**Version**
Windows Server 2003 and Windows 2000 Server

This setting controls the time variance allowed for certificate issuance, certificates, and CRLs. If the times on the clocks on two computers are further apart than this registry setting, a newly enrolled certificate would not be immediately valid and the client would have to wait several minutes before being able to use the new certificate. The default setting is 10 minutes. If you expect the clocks of two computers to be more than 10 minutes apart, you can increase this setting.

**CommonName**

**Registry Path**
\CertSvc\Configuration*CAName*\CommonName

**Version**
Windows Server 2003 and Windows 2000 Server

This setting lists the sanitized name of the server on which the CA is running. The process of sanitizing the CA name is necessary to remove characters that are not valid for file names, registry key names, or distinguished name values, or that are not valid for technology-specific reasons. The CommonName setting is used by several variables that you use to set the CRL and authority information access.

**CRLDeltaOverlapPeriod**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLDeltaOverlapPeriod

**Version**
Windows Server 2003 only

You use this setting to specify whether the overlap period for delta CRLs will be defined in days, weeks, months, or years.

**CRLDeltaOverlapPeriodUnits**

**Registry Path**
\CertSvc\Configuration*CAName*\CLRDeltaOverlapPeriodUnits

**Version**

Windows Server 2003 only

You use this setting to specify the number of days, weeks, months, or years that delta CRLs can overlap.

**CRLDeltaPeriod**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLDeltaPeriod

**Version**
Windows Server 2003 only

You use this setting to specify whether delta CRL lifetimes will be defined in days, weeks, months, or years. The delta CRL publication interval setting is similar to the CRL publication interval setting. If a CA is to be offline, you should disable delta CRL publication.

**CRLDeltaPeriodUnits**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLDeltaPeriodUnits

**Version**
Windows Server 2003 only

You use this setting to specify the number of days, weeks, months, or years that delta CRLs can be valid for. If a CA is to be offline, you should disable delta CRL publication.

**CRLEditFlags**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLEditFlags

**Version**
Windows Server 2003 and Windows 2000 Server

This setting enables or disables each field of the Authority Key Identifier (AKI) extension that is automatically generated by the CA for every new certificate.

**CRLFlags**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLFlags

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the oldest unexpired base CRL and includes details about critical and non-critical extensions associated with CRLs.

**CRLNextPublish**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLNextPublish

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the next time the CA expects to publish a CRL.

**CRLOverlapPeriod**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLOverlapPeriod

**Version**
Windows Server 2003 and Windows 2000 Server

This setting specifies whether the overlap period for CRLs will be defined in days, weeks, months, or years.

**CRLOverlapUnits**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLOverlapUnits

**Version**
Windows Server 2003 and Windows 2000 Server

You use this setting to specify the number of days, weeks, months, or years that CRLs can overlap. When a large number of certificates are revoked, such as during an employee layoff, the delta CRL size might increase significantly because of the large number of entries, and almost all clients will refer to the older base CRL. You can reduce the size of the overlap period to speed the propagation process for the new base CRL and help minimize the size of delta CRLs.

**CRLPeriod**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLPeriod

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to specify whether the validity period of a CRL will be defined in days, weeks, months, or years.

**CRLPeriodUnits**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLPeriodUnits

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to specify the number of days, weeks, months, or years, that a CRL will be valid.

**CRLPublicationURLs**

**Registry Path**
\CertSvc\Configuration*CAName*\CRLPublicationURLs

**Version**
Windows Server 2003 only

This setting identifies CRL distribution points where a client can find the CRL that is related to a certificate. The CRL distribution point of a root CA should be empty. For intermediate and issuing CAs, you use the Certification Authority snap-in to define CRL distribution points. These values are then stored in the registry.

A Windows Server 2003 CA maintains all distribution points in this registry key, regardless how the distribution point is accessed. In comparison, Windows 2000 CAs store CRL distribution point data in three different registry keys depending on the publication protocol that is being used: RevocationCRLURL, LDAPRevocationCRLURL, or FileRevocationCRLURL.

**DSConfigDN**

**Registry Path**
\CertSvc\Configuration*CAName*\DSConfigDN

**Version**
Windows Server 2003 and Windows 2000 Server

This setting records the fully qualified domain name of the domain namespace that the CA belongs to. The fully qualified distinguished name becomes part of the certificate issuer name and is also used by several variables that are used to set the CRL and authority information access. On an enterprise CA, this is set automatically. On an offline CA that is not a domain member, you must manually set the value of DSConfigDN so that the fully qualified distinguished name appears in certificates and CRLs.

**DSDomainDN**

**Registry Path**
\CertSvc\Configuration*CAName*\DSDomainDN

**Version**
Windows Server 2003 and Windows 2000 Server

This setting identifies the domain that is a parent to the configuration container used by the CA.

**Enabled**

**Registry Path**
\CertSvc\Configuration*CAName*\Enabled

**Version**
Windows Server 2003 and Windows 2000 Server

This setting identifies whether the CA is enabled.

**EnforceX500NameLengths**

**Registry Path**
\CertSvc\Configuration*CAName*\EnforceX500NameLengths

**Version**
Windows Server 2003 and Windows 2000 Server

This setting makes it possible for certificates issued by the CA to bypass X.500 naming restrictions. Windows Server 2003 CAs enforce X.500 naming standards by default. It is possible that deep OU paths might exceed normal length restrictions. If your organization's PKI does not need to be compatible with a non-Microsoft PKI, you can configure this registry setting to bypass the X.500 name length restriction.

**ForceTeletex**

**Registry Path**
\CertSvc\Configuration*CAName*\ForceTeletex

**Version**
Windows Server 2003 and Windows 2000 Server

This setting enables the ability to automatically change subject relative distinguished names between printable string, teletex, UTF-8, and Unicode representations.

**HighSerial**

**Registry Path**
\CertSvc\Configuration*CAName*\HighSerial

**Version**
Windows Server 2003 only

This setting enables you to choose between three different algorithms for generating fixed-length certificate serial numbers. These serial numbers can be from 10 bytes long up to 19 bytes long. This setting also allows you to generate new certificate serial numbers based on a new set of random numbers after each restart or to use the same set of random data for each certificate, even after the computer hosting the CA is restarted.

**InterfaceFlags**

**Registry Path**
\CertSvc\Configuration*CAName*\InterfaceFlags

**Version**
Windows Server 2003 and Windows 2000 Server

This setting controls whether network interfaces are enabled for performing CA database backups, processing certificate requests, and other administrative operations.

**KRACertCount**

**Registry Path**
\CertSvc\Configuration*CAName*\KRACertCount

**Version**
Windows Server 2003 only

This setting records the number of times a key recovery certificate has been issued from the CA.

**KRACertHash**

**Registry Path**
\CertSvc\Configuration*CAName*\KRACertHash

**Version**
Windows Server 2003 only

This setting defines the hash used to create key recovery certificates.

**KRAFlags**

**Registry Path**
\CertSvc\Configuration*CAName*\KRAFlags

**Version**
Windows Server 2003 only

This setting specifies the extensions that control the key recovery behavior of the CA.

**LogLevel**

**Registry Path**
\CertSvc\Configuration*CAName*\LogLevel

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines how much data to store in the CA log.

**MaxIncomingAllocSize**

**Registry Path**
\CertSvc\Configuration*CAName*\MaxIncomingAllocSize

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to limit the amount of allocated memory associated with any parameter of a certificate request.

**MaxIncomingMessageSize**

**Registry Path**
\CertSvc\Configuration*CAName*\MaxIncomingMessageSize

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to set a maximum limit on the size of a certificate request.

**PolicyFlags**

**Registry Path**
\CertSvc\Configuration*CAName*\PolicyFlags

**Version**
Windows Server 2003 and Windows 2000 Server

This setting specifies whether an optional policy module is being used.

### Note

- This setting is seldom used. However, there is a separate setting in CertSvc/Configuration/*CAName*/PolicyModules that should be used instead to enable the use of a custom policy module.

**RoleSeparationEnabled**

**Registry Path**
\CertSvc\Configuration*CAName*\RoleSeparationEnabled

**Version**
Windows Server 2003 and Windows 2000 Server

This setting determines whether role separation is enabled on the CA.

The separation of CA administrative roles enables you to assign only a single role to a user. This feature is valuable for large enterprises, because the separation of roles ensures that the compromise of a user's account does not compromise the entire CA administered by the user. Role separation is an important operating principle in Common Criteria environments. If this setting is enabled, the CA will not permit an administrator to perform multiple CA roles, such as CA administrator or certificate operator, which are defined in CA security configurations.

### Note

- Roles defined as operating system privileges — including Backup Operator and Auditor — are not affected by the RoleSeparationEnabled setting.

**Security**

**Registry Path**
\CertSvc\Configuration*CAName*\Security

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to specify the security descriptor that controls access to the CA.

**SetupStatus**

**Registry Path**
\CertSvc\Configuration*CAName*\SetupStatus

**Version**
Windows Server 2003 and Windows 2000 Server

This setting documents the configuration options that have been installed for this CA.

**SignedAttributes**

**Registry Path**
\CertSvc\Configuration*CAName*\SignedAttributes

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to specify the certificate request attributes that must be signed in order for a certificate request to be approved.

**SubjectTemplate**

**Registry Path**
\CertSvc\Configuration*CAName*\SubjectTemplate

**Version**
Windows Server 2003 and Windows 2000 Server

This setting contains an ordered list of the subject relative distinguished name elements that are allowed in the Subject field of certificates issued by the CA.

This setting can only be set to a small, fixed list of relative distinguished name elements supported by the CA. If during request processing a listed relative distinguished name field is empty, or if the field is not populated by the request Subject field or by

the policy module, the element will not be included. If the registry value is completely empty, the binary subject encoding from the request is passed through to the issued certificate unmodified.

**UseDS**

**Registry Path**
\CertSvc\Configuration*CAName*\UseDS

**Version**
Windows Server 2003 and Windows 2000 Server

This setting specifies whether the directory service is to be used for CRL publication.

**ValidityPeriod**

**Registry Path**
\CertSvc\Configuration*CAName*\ValidityPeriod

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to specify whether the validity period of certificates issued by the CA will be defined in days, weeks, months, or years. The default value depends on the type of certificate.

**ValidityPeriodUnits**

**Registry Path**
\CertSvc\Configuration*CAName*\ValidityPeriodUnits

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to define the number of days, weeks, months, or years that a certificate issued by the CA will be valid. The validity period for a certificate cannot be greater than the validity period of the CA that issued the certificate. The default value depends on the type of certificate.

**ViewAgeMinutes**

**Registry Path**
\CertSvc\Configuration*CAName*\ViewAgeMinutes

**Version**
Windows Server 2003 and Windows 2000 Server

This setting controls the refresh rate for the transaction view of the CA database. The default is 10 minutes.

**ViewIdleMinutes**

**Registry Path**
\CertSvc\Configuration*CAName*\ViewIdleMinutes

**Version**
Windows Server 2003 and Windows 2000 Server

This setting controls the refresh rate for the transaction view of the CA database when the transaction view is idle. The default is eight minutes.

**CertSvc\Configuration\CAName\CSP**
The registry settings in this hive contain information relating to the configuration of the cryptographic service provider (CSP) for the CA.

**HashAlgorithm**

**Registry Path**
\CertSvc\Configuration*CAName*\CSP\HashAlgorithm

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the hash algorithm that is used for hashing and signing certificate contents. This value, which is set when the CA is installed, is also recorded in the CA certificate.

**MachineKeyset**

**Registry Path**
\CertSvc\Configuration*CAName*\CSP\MachineKeyset

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the key sets that are stored on the computer.

**Provider**

**Registry Path**
\CertSvc\Configuration*CAName*\CSP\Provider

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the CSP used to generate the CA certificate and the key associated with the certificate. This value is set when the CA is installed.

**ProviderType**

**Registry Path**
\CertSvc\Configuration*CAName*\CSP\ProviderType

**Version**
Windows Server 2003 and Windows 2000 Server

This setting indicates the CSP family associated with the provider. When an application connects to a CSP of a particular type, each of the CryptoAPI functions operates by default in a way prescribed by the family that corresponds to that CSP type.

**CertSvc\Configuration\CAName\EncryptionCSP**
The registry settings in this hive contain information relating to the configuration of the encryption CSP for the CA.

**EncryptionAlgorithm**

**Registry Path**
\CertSvc\Configuration*CAName*\EncryptionCSP\EncryptionAlgorithm

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the hash algorithm that is used by the CA for encryption.

**KeySize**

**Registry Path**
\CertSvc\Configuration*CAName*\EncryptionCSP\KeySize

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the size of the secret key used by the encryption hash algorithm.

**MachineKeyset**

**Registry Path**
\CertSvc\Configuration*CAName*\EncryptionCSP\MachineKeyset

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the encryption key sets that are stored on the computer.

**Provider**

**Registry Path**
\CertSvc\Configuration*CAName*\EncryptionCSP\Provider

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the CSP used by the CA for encryption.

**ProviderType**

**Registry Path**
\CertSvc\Configuration*CAName*\EncryptionCSP\ProviderType

**Version**
Windows Server 2003 and Windows 2000 Server

This setting indicates the CSP family associated with the provider. When an application connects to a CSP of a particular type, each of the CryptoAPI functions operates by default in a way prescribed by the family that corresponds to that CSP type.

**CertSvc\Configuration\CAName\ExitModules**
The registry settings in this hive contain information relating to the configuration of the exit modules for the CA.

**Active**

**Registry Path**
\CertSvc\Configuration*CAName*\ExitModules\Active

**Version**
Windows Server 2003 and Windows 2000 Server

This setting lists the active exit modules.

**CertSvc\Configuration\CAName\ExitModules\CertificateAuthority_MicrosoftDefault.Exit**

**PublishCertFlags**

**Registry Path**
CertSvc\Configuration*CAName*\ExitModules\CertificateAuthority_MicrosoftDefault.Exit\PublishCertFlags

**Version**
Windows Server 2003 and Windows 2000 Server

This setting identifies whether certificates are published to a virtual root directory for use by Internet Information Services (IIS).

**CertSvc\Configuration\CAName\ExitModules\CertificateAuthority_MicrosoftDefault.Exit\SMTP**

**EventFilter**

**Registry Path**
CertSvc\Configuration*CAName*\ExitModules\CertificateAuthority_MicrosoftDefault.Exit\SMTP\EventFilter

**Version**
Windows Server 2003 and Windows 2000 Server

This setting identifies whether Simple Mail Transfer Protocol (SMTP) e-mail notifications are being sent out. For example, notification messages can be sent out when certificates are issued, when certificate requests are pending, when certificates are revoked, and so on.

**SMTPAuthenticate**

**Registry Path**
CertSvc\Configuration*CAName*\ExitModules\CertificateAuthority_MicrosoftDefault.Exit\SMTP\SMTPAuthenticate

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to enable a Windows Server 2003 CA to communicate with an Exchange Server by way of SSL.

**SMTPServer**

**Registry Path**
CertSvc\Configuration*CAName*\ExitModules\CertificateAuthority_MicrosoftDefault.Exit\SMTP\SMTPServer

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the SMTP server being used to send out CA-related messages. The SMTP registry hive also includes the following sub-hives, which can be used to issue programmatically-defined SMTP messages for each type of event:

- **CRLIssued**
- **Denied**
- **Issued**
- **Pending**
- **Revoked**
- **Shutdown**
- **Startup**

Each of the above sub-hives includes the following:

- **BodyArg.** Programmatically defined contents of the message body.
- **BodyFormat.** Programmatically defined format of the contents of the message body.
- **TitleArg.** Programmatically defined title of the SMTP message.
- **TitleFormat.** Programmatically defined format of the title of the SMTP message.

**CertSvc\Configuration\CAName\PolicyModules**
The registry settings in this hive contain information relating to the configuration policy modules for the CA.

**Active**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\Active

**Version**
Windows Server 2003 and Windows 2000 Server

This setting lists the active policy modules.

**CertSvc\Configuration\CAName\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy**

**CAPathLength**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\CAPathLength

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the maximum number of levels of subordinate CAs that can be added underneath this CA. Path length constraints can be applied to CAs at any level in a hierarchy. If a path length constraint is being applied to a root CA, it needs to be configured using the CAPolicy.inf file as part of the CA's installation or renewal. If the path length constraint is being applied to a subordinate CA, the constraint must be configured in the registry of the parent CA.

**DefaultSMIME**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\DefaultSMIME

**Version**
Windows Server 2003 and Windows 2000 Server

This setting lists — by object identifier (OID) — the supported hash algorithms, along with the bit length supported for each algorithm.

**DisableExtensionList**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\DisableExtensionList

**Version**
Windows Server 2003 and Windows 2000 Server

This setting disables the adding of a certificate extension to a certificate that is included by default in certificates issued by an enterprise CA. An example of such a certificate extension could be the S/MIME capabilities extension.

**EditFlags**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\EditFlags

**Version**
Windows Server 2003 and Windows 2000 Server

This setting enables you to identify the critical and non-critical extensions associated with policy modules, such as enabling basic constraints. The flags that can be configured include:

- **EDITF_ATTRIBUTEENDDATE.** Controls certificate validity time by certificate request. This feature is disabled by default for enterprise CAs, and enabled by default for stand-alone CAs.

  **Note**

  - Although the validity period of a certificate is normally set by the CA, the CA can be configured in a way that allows the request to specify the validity period. This CA configuration option can only reduce the validity period of certificates to be issued by a request.

- **EDITF_ENABLEAKIISSUERNAME.** Used to enable or disable the Issuer name and Issuer serial number.

  **Note**

  - You should disable the issuer name and issuer serial number if you renew CA keys when you renew CA certificates or plan to use cross-certificates.

**EnableEnrolleeRequestExtensionList**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\EnableEnrolleeRequestExtensionList

**Version**
Windows Server 2003 and Windows 2000 Server

This setting lists the OIDs of the rules extensions defined by the organization that must appear in the issued certificate. These settings apply to all requests submitted to a stand-alone CA, and to requests submitted to an enterprise CA that use a template which specifies that the request must supply the subject information.

**EnableRequestExtensionList**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\EnableRequestExtensionList

**Version**

Windows Server 2003 and Windows 2000 Server

This setting lists the OIDs of the rules extensions defined by the organization that must appear in the issued certificate. By default, extensions that are in the submitted request are added to the CA's database, but are marked as disabled, so they will not appear in the issued certificate.

**RequestDisposition**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\RequestDisposition

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to specify what the policy module should do with incoming and resubmitted requests. Options include:

- Process all certificate requests that meet applicable policy and template processing requirements.
- Place all initial certificate requests in a pending queue, and process all certificate requests that are resubmitted by an administrator, assuming the requests meet applicable policy and template processing requirements.

**RevocationType**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\RevocationType

**Version**
Windows Server 2003 and Windows 2000 Server

You can use this setting to specify an alternative certificate revocation service to function with a Windows Server 2003 CA. This setting enables CRL URLs from the following: LDAP, FTP, and file locations; CDP extensions; and ASP-based revocation.

**RevocationURL**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\RevocationURL

**Version**
Windows Server 2003 and Windows 2000 Server

This setting defines the URL for an alternative certificate revocation service that has been configured to function with a Windows Server 2003 CA.

### Note

- In order for a revocation service to work with a Windows Server 2003 CA, the application, service, or account connecting to this URL must have Read permissions in the Certification Authority snap-in.

The replacement variables that can be used in the revocation URL are listed in the following table.

**Registry Variables for Alternative Revocation Services**

| Alternative Revocation Service | Registry Variable |
|---|---|
| SERVERDNSNAME | 1 |
| SERVERSHORTNAME | 2 |
| SANITIZEDCANAME | 3 |
| CERTFILENAMESUFFIX | 4 |
| DOMAINDN | 5 |

| | |
|---|---|
| CONFIGDN | 6 |
| SANITIZEDCANAMEHASH | 7 |
| CRLFILENAMESUFFIX | 8 |
| CRLDELTAFILENAMESUFFIX | 9 |
| DSCRLATTRIBUTE | 10 |
| DSCACERTATTRIBUTE | 11 |
| DSUSERCERTATTRIBUTE | 12 |
| DSKRACERTATTRIBUTE | 13 |
| DSCROSSCERTPAIRATTRIBUTE | 14 |

In order for the revocation service to function, the application, service, or account connecting to this URL must have Read permissions in the Certification Authority snap-in. If IIS is using a local account, follow the steps for enabling anonymous access in IIS and allowing Anonymous Read access to the CA.

**Note**

- Enabling Anonymous Read access to the CA might expose privacy or security risks.

**SubjectAltName**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\SubjectAltName

**Version**
Windows Server 2003 and Windows 2000 Server

This setting uses an OID for the SubjAltName extension of an issued certificate. This setting is almost never used.

**SubjectAltName2**

**Registry Path**
CertSvc\Configuration*CAName*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy\SubjectAltName2

**Version**
Windows Server 2003 and Windows 2000 Server

This setting makes it possible for a stand-alone CA to place in the SubjAltName extension of an issued certificate the e-mail address of the authenticated user making the certificate request. This setting is rarely used.

**CryptoAPI Policy**

**ProtectedRoots**

**Registry Path**
HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\SystemCertificates\Root\ProtectedRoots

**Version**
Windows Server 2003

This setting is used to disable or enable:

- CryptoAPI applications to use the user root store in building trusted certificate chains.
- Users to add root CAs to the user trusted root store.

  **Note**

  - In Windows Server 2003, this value can also be set through Group Policy.

- Users to remove root CAs from the user trusted store that are also the local computer trusted root store.
- The requirement for NTAuth policy processing.

  **Note**

  - In Windows Server 2003, this value can also be set through Group Policy.

- Name constraint enforcement for undefined name types. By default, Windows Server 2003 rejects undefined name types in a name constraint validation. Setting this value will accept all name forms that are not explicitly defined.

For information about managing these registry settings through Group Policy, see "Trusted Root Certification Authorities" later in this document.

#### Certificate Services Group Policy Settings

Most features of a public key infrastructure and Certificate Services work without your having to configure Public Key Group Policy settings. However, you must configure Public Key Group Policy if you want to configure:

- **Automatic Certificate Request Settings** for configuring autoenrollment for computer certificates.
- **Trusted Root Certification Authorities** for adding trusted root CA certificates to the Trusted Root Certification Authorities store.
- **Enterprise Trust** for configuring CTLs.
- **Encrypted Data Recovery Agents** for configuring Encrypting File System (EFS) recovery agents. (This is the only container that appears for Local Computer policy.)

#### Automatic Certificate Enrollment

Certificate autoenrollment is based on the combination of Group Policy settings and version 2 certificate templates. This combination enables the Windows XP Professional or Windows Server 2003 client to enroll users when they log on to their domain or to a computer when it starts, and to periodically update enrollments between these events.

Autoenrollment Group Policy is an element of Public Key Policies in the Security Settings of the Group Policy object for a user or computer in a domain, site, or OU. You can configure autoenrollment policy on the **Properties** page of the **Autoenrollment Settings** item. Options that you can configure are:

- **Do not enroll certificate automatically.** Blocks attempts to configure autoenrollment for the user or computer.
- **Enroll certificates automatically.** Enables certificate autoenrollment.
- **Renew expired certificates, update pending certificates, and remove revoked certificates.** Enables automatic certificate renewal, and cleans up expired certificates.
- **Updates certificates that use certificate templates.** Updates certificates as needed to conform to the associated certificate templates.

#### Trusted Root Certification Authorities

When you install an enterprise root CA or a stand-alone root CA, the certificate of the CA is added automatically to the Trusted Root Certification Authorities Group Policy for the domain. You also can add certificates for other root CAs to Trusted Root Certification Authorities Group Policy. The root CA certificates that you add become trusted root CAs for computers within the scope of the Group Policy object. For example, if you want to use a non-Microsoft CA as a root CA in a certification hierarchy, you must add the certificate for the non-Microsoft CA to the Trusted Root Certification Authorities Group Policy.

To add a certificate for the root CA to the Trusted Root Certification Authorities Group Policy, in the **Public Key Policies** node, right-click **Trusted Root Certification Authorities**, and then click **All Tasks** and **Import**. When the Certificate Import Wizard appears, use the wizard to import a certificate file for the certificate of the root CA and add it to Group Policy. The certificate is added to the **Trusted Root Certification Authorities** store of all computers within the scope of Group Policy the next time it is

refreshed on each computer.

You can also use Trusted Root Certification Authorities Group Policy to control the changes users can make to trusted root CA options.

To modify these settings, right-click **Trusted Root Certification Authorities**, and then select **Properties**.

To prevent users making any changes to the trusted root CA store, you can clear **Allow users to select new root Certification Authorities (CAs) to trust**.

To limit client trust of alternative certificate stores, under **Client computers can trust the following certificate stores** select one of the following options:

- **Third-Party Root Certification Authorities and Enterprise Root Certification Authorities**
- **Enterprise Root Certification Authorities**

In addition, the following settings under **To perform certificate-based authentication of users and computers, CAs must meet the following criteria:** can be used to manage certificate-based authentication by certificate holders:

- **Registered in Active Directory only**
- **Registered in Active Directory and compliant with name constraints requirements for user principal names (UPNs)**

The Trusted Root Certification Authorities Group Policy settings are the counterparts to the CryptoAPI Policy\ProtectedRoots registry settings described earlier in this document.

**Certificate Trust Lists**

You can create certificate trust lists (CTLs) to trust specific CAs and to restrict the uses of certificates issued by the CAs. For example, you might use a CTL to trust certificates that are issued by a commercial CA and restrict the permitted uses for those certificates. You might also use CTLs to control trust on an extranet for certificates that are issued by CAs that are managed by your business partners. You can configure CTLs for computers and for users.

Before administrators can create CTLs, they must have a valid trust list signing certificate — such as the Administrator certificate or the Trust List Signing certificate that is issued by enterprise CAs. The trust list signing private key for the administrator is used to sign the CTL for integrity. If the trust list signing certificate for an administrator is invalid, all CTLs that have been created and signed by that administrator also are invalid.

Certificate Trust List Group Policy is configured as part of Public Key Policies\Enterprise Trust in the Security Settings of the Group Policy object for a user or computer in a domain, site, or OU.

You use the Certificate Trust Wizard to configure Certificate Trust List Group policy.

To activate the wizard, you right-click the **Enterprise Trust** node, and then click **New** and **Certificate Trust List**. The wizard enables you to configure the following options:

- **Valid duration.** An optional lifetime for the CTL. If you do not specify a lifetime, the CTL expires when the trust list signing certificate expires.
- **Designate Purposes.** The CTL establishes trust only for certificates that are valid for the selected purposes. A certificate might support all of the listed purposes, but you can restrict the purposes for which certificates are trusted.
- **Add Purpose.** Enables you to add purposes to the **Designate Purposes** box. This also requires you to enter an OID for the new purpose.
- **Current CTL Certificates.** Displays the certificates of the root CAs that are to be trusted by this CTL. Certificates with certification paths to the root CA are trusted for all designated purposes specified by the CTL.
- **Add from Store.** Adds a root certificate from the Trusted Root Certification Authorities store.
- **Add from File.** Adds a root CA's certificate from a file.

- **Remove.** Deletes the certificate that is selected in the **Current CTL Certificates** box.
- **View Certificate.** Enables you to view the certificates that are selected in the **Current CTL Certificates** box.
- **Use this certificate.** Displays the trust list signing certificate for the private key that is to be used to sign the CTL.
- **Select from Store.** Adds a trust list signing certificate from the Personal store for the administrator.
- **Select from File.** Adds the trust list signing certificate from a file.
- **View Certificate.** Enables you to view the certificate listed in the **Use this certificate** box.
- **Add a timestamp to the data.** Adds a timestamp to the CTL. The timestamp is used to determine the valid lifetime of the CTL. If a timestamp is not used, the computer clock is used instead.
- **Timestamp service URL.** Identifies the location of the timestamp service that is to be used for the timestamp.
- **Friendly Name.** The optional name that appears in MMC when the CTL is displayed.
- **Description.** An optional description to describe this CTL.

**EFS Recovery Agents**

A recovery agent is an individual authorized to decrypt data that was encrypted by another user. Recovery agents do not need any other permissions to function in this role. Recovery agents are useful, for example, when employees leave the company and their remaining data needs to be decrypted.

By default, the local Administrator users account for the first domain controller that is installed in the domain is the EFS recovery account for that domain. You can specify alternative recovery agents for EFS. Before you can add a recovery agent for a domain, you must ensure that each recovery agent has been issued a recovery agent certificate.

To add recovery agents for a domain, their certificates need to be added to the existing recovery policy. There are two ways to accomplish this. You can:

- Add recovery agent certificates that are published in Active Directory to the recovery policy.
- Add recovery agent certificates from a file that is located on a disk or in a shared folder that is available from the computer where you are configuring Public Key settings.

Encrypted Data Recovery Agent Group Policy is configured as part of Public Key Policies\Encrypting File System in the Security Settings of the Group Policy object for a user or computer in a domain, site, or OU.

You configure Encrypted Data Recovery Agent Group Policy by using the Add Recovery Agent Wizard. To activate the wizard, right-click the **Encrypted Data Recovery Agents** node and then click **Add**.

You can use the wizard to configure the following:

- **Recovery agents.** Displays the certificates you choose for recovery agents.
- **Browse Directory.** Locate and add a recovery agent certificate for a user account by browsing in Active Directory. Use this option when the certificate is published in Active Directory.
- **Browse Folders.** Locate and add a recovery agent certificate for a user account by browsing folders and files.

When you select **Encrypted Data Recovery Agents**, the EFS recovery agent certificates that are applied by Group Policy appear in the details pane of the console. These are the recovery agent certificates that are used by EFS within the scope of Group Policy.

**Note**

- Encrypted Data Recovery Agents is the only Certificate Services–related container that appears in local computer Group Policy.

# Trust Technologies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Trust Technologies**

By default, all users in a specific Windows domain can be authenticated to resources contained within that domain. In this way, a domain can provide its users with secured access to all resources in that domain. To expand that access to include resources beyond the boundaries of a single domain, you need to use trust relationships. Trust relationships provide a mechanism for one domain to allow access to resources based on the logon authentications of another domain.

**Trust Concepts**

Trusts in Microsoft Windows NT version 4.0 differ from trusts in the Microsoft Windows 2000 and Windows Server 2003 operating systems. In domains that have domain controllers running Windows NT 4.0 Server and earlier, trusts are limited to two domains, and the trust relationship is one-way (only one of the two domains trusts the other domain) and nontransitive (does not extend to any other domains trusted by the two domains). With domains that have domain controllers running Windows Server 2003 or Windows 2000 Server, all trusts created within a forest are two-way (both domains trust each other) and transitive (extends to all domains in the forest). Domains that have domain controllers running Windows 2000 Server or Windows Server 2003 (domains that use the Active Directory directory service) are also referred to as Active Directory domains.

To understand trust technology, you should be familiar with basic trust concepts including the nature and purpose of trust relationships, the role of trusted authorities, the variety of trust paths, and trust transitivity.

**Trust Relationships**

A trust relationship (also called a trust) is a logical relationship established between domains to allow authentication and authorization to shared resources. The authentication process verifies the identity of the user, and the authorization process determines what the user is permitted to do on a computer system or network. Once a user requesting access to a resource computer in another domain has been authenticated by the resource domain, the resource computer compares the user's credentials to the permissions assigned within its security descriptor to help determine the user's level of authorization to that resource. A security descriptor contains access control lists (ACLs) that identify the users and groups that are assigned or denied access permissions on a resource.

In its simplest form, a trust acts as a technological drawbridge, by either allowing or disallowing authentication traffic to flow between two or more domains. When a trust relationship is created between two domains, traffic is allowed over the bridge, permitting the sharing of resources between them, as shown in the following illustration.

**Simple Trust Relationship**




The domain where the user account requesting access is located is referred to as the trusted domain. The domain that contains a shared resource that a user account is trying to access is referred to as the trusting domain.

Severing a trust removes the bridge through which authentication traffic flows, and removes all relationships to any trusted authorities located in the other domain. When this occurs, no authentication traffic originating from a user in the formerly trusted domain can cross over to the formerly trusting domain; it is no longer possible to share information across domain boundaries, as is shown in the following figure.

**Severed Trust Relationship**




**Trusted Authorities**

Active Directory domains do not unconditionally accept credentials coming from other domains; they accept credentials only from trusted authorities. In Active Directory, domain controllers act as trusted authorities for all security principals located in their domain and all security principals located within trusted domains, as long as a valid trust is present between those domains.

A trusted authority, in Active Directory, is analogous to a certificate granting agency that distributes certificates by which a sub agency can verify its authenticity. For example, when a passport is issued by a passport authority in one country (a domain controller in a trusted domain), customs officials physically located in another country (domain controllers in a trusting domain), trust the authority that issued the passport, therefore they trust the validity of the passport.

Conversely, a company might create a specific means by which to identify employees, such as employee badges issued by its various departments (domains). The manufacturing department (domain controllers in a trusting domain) might trust identification badges issued by the sales department (domain controllers in a trusted domain) for important transactions, but other companies (domains or forests that do not have a trust relationship with that company or its departments) would not trust them because they do not inherently trust the authority that issued them.

In Windows NT 4.0 or Active Directory domains, domain administrators act as the chief executive officer of the certificate granting agency by defining the policy regarding what external authorities (domains) are trusted in their domain. For example, a domain administrator in a trusting domain first identifies which other domains should be trusted to access shared resources in the domain. The administrator then establishes the trust relationships that provide a path by which authentication requests from trusted domains can travel and be verified. In this way, trusts define for each domain which authentication requests are valid. For more information about trust paths, see “Trust Paths” later in this overview.

**Trust Paths**

The direction that a trust is assigned determines the trust path used for authentication. A trust path is defined by the series of trust relationships that authentication requests must follow between domains. Before a user can access a resource in another domain, the security system on domain controllers running Windows NT 4.0, Windows 2000 Server, and Windows Server 2003 must determine whether the trusting domain (the domain containing the resource the user is trying to access) has a trust relationship with the trusted domain (the user’s logon domain). To determine this, the security system computes the trust path between a domain controller in the trusting domain and a domain controller in the trusted domain. All domain trust relationships have only two domains: the trusting domain and the trusted domain. In the following figure, the trust path is indicated by arrows showing a one-way direction of trust and the direction of access:

**Trust Path in a One-Way Trust**




**One-Way Trust**

A one-way trust is a unidirectional authentication path created between two domains (trust flows in one direction, and access flows in the other). This means that in a one-way trust between a trusted domain and a trusting domain, users or computers in the trusted domain can access resources in the trusting domain. However, users in the trusting domain cannot access resources in the trusted domain. Some one-way trusts can be either nontransitive or transitive, depending on the type of trust being created.

**Two-Way Trust**

A two-way trust can be thought of as a combination of two, opposite-facing one-way trusts, so that, the trusting and trusted domains both trust each other (trust and access flow in both directions). This means that authentication requests can be passed between the two domains in both directions. Some two-way relationships can be either nontransitive or transitive depending on the type of trust being created. All domain trusts in an Active Directory forest are two-way, transitive trusts. When a new

child domain is created, a two-way, transitive trust is automatically created between the new child domain and the parent domain.

**Trust Transitivity**

Transitivity determines whether a trust can be extended beyond the two domains between which it was formed. A transitive trust extends trust relationships to other domains; a nontransitive trust does not extend trust relationships to other domains. Each time you create a new domain in a forest, a two-way, transitive trust relationship is automatically created between the new domain and its parent domain. If child domains are added to the new domain, the trust path flows upward through the domain hierarchy, extending the initial trust path created between the new domain and its parent.

Transitive trust relationships thus flow upward through a domain tree as it is formed, creating transitive trusts between all domains in the domain tree. A domain tree can therefore be defined as a hierarchical structure of one or more domains, connected by transitive, bidirectional trusts, that forms a contiguous namespace. Multiple domain trees can belong to a single forest.

Authentication requests follow these extended trust paths, so accounts from any domain in the forest can be authenticated by any other domain in the forest. Consequently, with a single logon process, accounts with the proper permissions can access resources in any domain in the forest.

With a nontransitive trust, the flow is restricted to the two domains in the trust relationship and does not extend to any other domains in the forest. A nontransitive trust can be either a two-way trust or a one-way trust.

**Trust Architecture**

The architecture of trusts includes various interdependent technologies and processes that, when implemented together, provide cross domain resource sharing. These technologies include authentication, authorization, the Net Logon service and Active Directory, which all rely on DNS or Windows Internet Name Service (WINS) to locate domain service records. The following figure shows the various technologies that make up the trust architecture.

### Trust Architecture




Active Directory provides the foundation for trust relationships by supplying a central management system through which domain administrators can grant access to trusted authorities in other domains. Trusts then provide the mechanism through which authentication traffic between domains or forests must travel to obtain authorization for a given shared resource. Authentication and authorization are the security enforcement mechanisms within this resource access management environment.

The authentication process and the Net Logon service provide pass-through validation of credentials for users or distributed applications located in other domains, and rely on either the NTLM or Kerberos version 5 authentication protocol to transport all trust related traffic. Applications can use these protocols to securely and seamlessly integrate with the Active Directory user store for authentication, and can also use these protocols to secure application data.

**Trust Components**

The various components of trust technologies and the related components that are an important part of the Windows Server 2003 security architecture determine how trusts function. The following table provides a summary view of these components.

### Trust Technologies and Related Components

| Related components | Description |
|---|---|
| Security Identifiers (SIDs) | The Windows security model identifies account objects such as users, groups, computers, and domains by SIDs. SIDs are domain-unique values, built when the user or group is created, or when the computer or trust is registered with the domain. The components of a SID follow a hierarchical convention: A SID contains parts that identify the revision number, the authority — such as the domain — that issued the SID, and a variable number of sub authority or relative identifier (RID) values that uniquely identify the security principal relative to the issuing authority. |

| | |
|---|---|
| Access Tokens and Authentication | An essential component of the Windows-based trust model is authentication, which involves identifying the user to the local or trusted domain by presenting credentials, usually in the form of a user name and password. Assuming these credentials are acceptable, the system creates an access token for the user that contains the SID of the user (the primary SID), and the SIDs of all the local and domain groups of which the user is a member. Every process the user creates, for example by running an application, carries the user's access token. |
| Security Descriptors and Authorization | Once authenticated, the user can attempt to gain access to resources from any domain in the forest by using the Active Directory authorization process. This process determines what a user is permitted to do on a resource by using the access token to determine whether and at what level to grant the user access to system resources. The counterpart of the user's access token is the security descriptor attached to resources such as files or printers. A security descriptor contains a discretionary access control list (DACL), which consists of a list of access control entries (ACEs). Each ACE consists of a SID that identifies a security principal, together with an indicator of the specific access on that resource that is granted or denied to that security principal. |
| SIDHistory | In Active Directory, domain migration or restructuring across trusts is made considerably easier by an attribute on Active Directory security principals called SIDHistory. SIDHistory stores the former SIDs of moved objects such as users and security groups. When a user is moved, the SIDHistory attribute of the user object is updated with the former SID. When the user then logs on to the system, the system retrieves the entries in the user object and the user's SIDHistory and adds them to the user access token. In this way SIDHistory ensures that migrated users can continue to access resources located in a trusting (resource) domain, even though the user's new domain does not have a trust relationship with the resource domain. |
| SID Filtering | SID filtering prevents domains from accepting SIDs with domain SIDs from outside the sender's domain. Applying SID filtering to trusts can prevent malicious users who have domain administrator level access in the trusted domain from granting, to themselves or other user accounts in their domain, elevated user rights in the trusting domain. |

**Trust Deployment Scenarios**
There are three scenarios in Active Directory in which specific types of trusts can be used to help accommodate different resource sharing needs of an organization. These include using trusts within a forest (intra-forest trusts), using trusts across forests (inter-forest trusts), and using trusts to collaborate with Kerberos realms.

**Intra-Forest Trusts**
Intra-forest trusts are transitive trusts that can be used only within a single forest, and include tree-root, parent-child, and shortcut trusts.

**Tree-root trusts**
By default, two-way, transitive trusts are automatically created when a new domain is added to a domain tree or forest root domain by using the Active Directory Installation Wizard. When a new domain tree is created in an existing forest, a new tree-root trust is established.

**Parent-child trusts**
When a new child domain is added to an existing domain tree, a new parent and child trust is established. Authentication requests made from subordinate domains flow upward through their parent to the trusting domain.

**Shortcut trusts**
You can use shortcut trusts to improve user logon times between two domains that are located in two separate domain trees within the same forest. Authentication requests must first travel a trust path between domain trees, and in a complex forest this can take time. Using shortcut trusts helps to speed up authentication in this situation.

**Inter-Forest Trusts**
Inter-forest trusts can be nontransitive or transitive, depending on the type of inter-forest trust used, and can only be created between domains located in different forests or realms. Inter-forest trust types include external trusts and forest trusts; both of these trust types must be manually created.

**External trusts**
External trusts are nontransitive and can be created between Active Directory domains in different forests or between an Active Directory domain and a Windows NT 4.0 domain.

**Forest trusts**
With Windows Server 2003 forests, you can link two different forests to create a one-way or two-way transitive trust relationship. A two-way forest trust is used to form a transitive trust relationship between every domain in both forests. Forest trusts can be created only between two Windows Server 2003 forests and cannot be implicitly extended to a third forest.

**Kerberos Realm Trusts**

A realm trust can be established between any non-Windows-based operating system Kerberos version 5 realm and a Windows 2000 Server or Windows Server 2003 domain. This trust relationship allows cross-platform interoperability with security services based on other Kerberos version 5 implementations, such as that from the Massachusetts Institute of Technology. Realm trusts can be manually switched back and forth between nontransitive and transitive by using the Netdom.exe tool. Realm trusts can be either one-way or two-way.

# Domain and Forest Trusts Technical Reference

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2012, Windows Server 2012 R2

Trust technology is the foundation for the security architecture in Microsoft Windows Server 2003 networks using the Active Directory directory service. Trusts enable service administrators to implement an authentication and authorization strategy for sharing resources across domains or forests and provide a mechanism for centralizing management of multiple domains and forests.

**In this subject**


- What Are Domain and Forest Trusts?
- How Domain and Forest Trusts Work
- Domain and Forest Trust Tools and Settings
- Security Considerations for Trusts

# What Are Domain and Forest Trusts?

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Trust Scenarios
- Technologies Related to Trusts


Most organizations that have more than one domain have a legitimate need for users to access shared resources located in a different domain. Controlling this access requires that users in one domain can also be authenticated and authorized to use resources in another domain. To provide authentication and authorization capabilities between clients and servers in different domains, there must be a trust between the two domains. Trusts are the underlying technology by which secured Active Directory communications occur, and are an integral security component of the Windows Server 2003 network architecture.

When a trust exists between two domains, the authentication mechanisms for each domain trust the authentications coming from the other domain. Trusts help provide for controlled access to shared resources in a resource domain (the trusting domain) by verifying that incoming authentication requests come from a trusted authority (the trusted domain). In this way, trusts act as bridges that allow only validated authentication requests to travel between domains.

How a specific trust passes authentication requests depends on how it is configured; trust relationships can be one-way, providing access from the trusted domain to resources in the trusting domain, or two way, providing access from each domain to resources in the other domain. Trusts are also either nontransitive, in which case trust exists only between the two trust partner domains, or transitive, in which case trust automatically extends to any other domains that either of the partners trusts.

In some cases, trust relationships are automatically established when domains are created; in other cases, administrators must choose a type of trust and explicitly establish the appropriate relationships. The specific types of trusts used and the structure of the resulting trust relationships in a given trust implementation depend on such factors as how the Active Directory directory service is organized, and whether different versions of Windows coexist on the network.

**Trust Scenarios**

It is possible to create a number of different domain and forest trust configurations, depending on the Active Directory structure of the organization. Windows Server 2003 domains and forests can trust other Windows Server 2003 domains and forests, as well as Windows 2000 and Windows NT 4.0 domains. For example, trust configurations vary in nature and complexity in each of the following scenarios:

**Trusts within a single Windows 2000 Server or Windows Server 2003 forest**

By default, all domain trusts within a single Active Directory forest are two-way, transitive trusts. There are three types of transitive trusts that are used within a single Windows 2000 Server or Windows Server 2003 forest. The first is the tree-root trust, which is created by default when you create a new domain tree by using the Active Directory Installation Wizard. The two-way transitive nature of intra-forest trusts such as the tree-root trust allows all domains in one tree to trust all domains in any other tree within the same forest.

The second type of trust is a parent-child trust. It is created automatically when you create a new domain in an existing domain tree by using the Active Directory Installation Wizard. When a new child domain is created, a parent-child trust is established between the new domain and the domain that immediately precedes it in the namespace hierarchy.

The last type of trust that can be used between trees is a shortcut trust, and is used to speed up access times to resources in a domain that is deep within the tree hierarchy of another domain.

**Trusts between two Windows Server 2003 forests**

It is possible to extend the transitivity of domain trusts within a single Windows Server 2003 forest to another Windows Server 2003 forest by manually creating a one-way or two-way forest trust. A forest trust is a transitive trust between a forest root domain and a second forest root domain. A one-way forest trust allows all users in one forest to trust all domains in the other forest; a two-way forest trust forms a transitive trust relationship between every domain in both forests. The transitivity of forest trusts is limited to the two forest partners; the forest trust does not extend to additional forests trusted by either of the partners.

**Trusts across Windows Server 2003 and Windows 2000 forests**

Windows Server 2003 forest trusts cannot be created between a Windows Server 2003 forest and a Windows 2000 forest. You can, however, manually create a trust relationship between any domain in a Windows Server 2003 forest and any domain in a Windows 2000 forest by using one-way or two-way external trusts. External trusts are nontransitive and provide for access to resources in another domain outside the forest that is not already joined by a forest trust.

**Trusts between Windows Server 2003 or Windows 2000 domains and Windows NT 4.0 domains**
You can manually create a one-way or two-way external trust between Windows Server 2003 or Windows 2000 domains and Windows NT 4.0 domains so that users from either domain can be authenticated to access resources in the other domain.

**Trusts between Windows 2000 or Windows Server 2003 domains and non-Windows Kerberos realms**
Windows 2000 or Windows Server 2003 domains can be configured to trust non-Windows-brand operating system Kerberos realms, and non-Windows Kerberos realms can be configured to trust Windows Server 2003 domains by manually creating one-way or two-way realm trusts. Realm trusts can also be configured to be either nontransitive or transitive, depending on the level of interoperability you require with UNIX or Massachusetts Institute of Technology implementations of the Kerberos version 5 protocol.

When the direction of a one-way trust is from a non-Windows Kerberos realm to a Windows Server 2003 domain, the user in the Windows Server 2003 domain can access resources in the non-Windows Kerberos realm. When the direction of trust is from a Windows Server 2003 domain to a non-Windows Kerberos realm, users in the non-Windows Kerberos realm can access the resources in the Windows Server 2003 domain.

**Technologies Related to Trusts**
Trusts depend on the NTLM and Kerberos authentication protocols and on Windows-based authorization and access control mechanisms to help provide a secured communications infrastructure across Active Directory domains and forests. The following diagram illustrates how authentication and authorization technologies relate to trusts and other components of the Windows distributed security model.

### Trusts and the Windows Distributed Security Model




**Applications and Net Logon**
Both applications and the Net Logon service are components of the Windows distributed security channel model. Applications integrated with Windows Server 2003 and Active Directory use authentication protocols to communicate with the Net Logon service so that a secured path can be established over which authentication can occur.

**Authentication Protocols**
Active Directory domain controllers authenticate users and applications by using one of two protocols: either the Kerberos version 5 authentication protocol or the NTLM authentication protocol. When two Active Directory domains or forests are connected by a trust, authentication requests made using these protocols can be routed to provide access to resources in both forests.

**NTLM**
The NTLM protocol is the default protocol used for network authentication in the Windows NT 4.0 operating system. For compatibility reasons, it is used by Active Directory domains to process network authentication requests that come from earlier Windows-based clients and servers. Computers running Windows 2000, Windows XP or Windows Server 2003 use NTLM only when authenticating to servers running Windows NT 4.0 and when accessing resources in Windows NT 4.0 domains.

When the NTLM protocol is used between a client and a server, the server must contact a domain authentication service on a domain controller to verify the client credentials. The server authenticates the client by forwarding the client credentials to a domain controller in the client account domain. The authentication protocol of choice for Active Directory authentication requests, when there is a choice, is Kerberos version 5. When the Kerberos protocol is used, the server does not have to contact the domain controller. Instead, the client gets a ticket for a server by requesting one from a domain controller in the server account domain; the server validates the ticket without consulting any other authority.

**Kerberos Version 5 Protocol**
The Kerberos version 5 protocol is the default authentication protocol used by computers running Windows 2000, Windows XP Professional, or Windows Server 2003. This protocol is specified in RFC 1510 and is fully integrated with Active Directory, server message block (SMB), HTTP, and remote procedure call (RPC), as well as the client and server applications that use these protocols. In Active Directory domains, the Kerberos protocol is used to authenticate logons when any of the

following conditions is true:

- The user who is logging on uses a security account in an Active Directory domain.
- The computer that is being logged on to is a Windows 2000, Windows XP or Windows Server 2003–based computer.
- The computer that is being logged on to is joined to an Active Directory domain.
- The computer account and the user account are in the same forest.
- The computer from which the user is trying to access resources is located in a non-Windows Kerberos realm.

If any computer involved in a transaction does not support the Kerberos version 5 protocol, the NTLM protocol is used.

**Authorization and Access Control**
Authorization and trust technologies work together to help provide a secured communications infrastructure across Active Directory domains or forests. Authorization determines what level of access a user has to resources in a domain. Trusts facilitate cross-domain authorization of users by providing a path for authenticating users in other domains so their requests to shared resources in those domains can be authorized.

Once an authentication request made to a resource in a trusting domain is validated by the trusted domain, it is passed to the targeted resource computer, which determines, based on its access control configuration, whether to authorize the specific request made by the user, service, or computer in the trusted domain. In this way, trusts provide the mechanism by which validated authentication requests are passed to a trusting domain, while access control mechanisms on the resource computer determine the final level of access granted to the requestor in the trusted domain.

**Note**

- "Access to resources" in any discussion of trust relationships always assumes the limitations of access control.

# How Domain and Forest Trusts Work

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Trust Architecture
- Trust Processes and Interactions
- Network Ports used by Trusts


Active Directory provides security across multiple domains or forests through domain and forest trust relationships. Before authentication can occur across trusts, Windows must first determine whether the domain being requested by a user, computer or service has a trust relationship with the logon domain of the requesting account. To make this determination, the Windows security system computes a trust path between the domain controller for the server that receives the request and a domain controller in the domain of the requesting account.

The access control mechanisms provided by Active Directory and the Windows distributed security model provide an optimal environment for the operation of domain and forest trusts. For these trusts to function properly, every workstation and server must have a direct trust path to a domain controller in the domain in which it is located. The trust path is implemented by the Net Logon service through an authenticated remote procedure call (RPC) connection to the trusted domain authority, which is the domain controller. In addition, a secured channel extends to other Active Directory domains through interdomain trust relationships. This secured channel is used to obtain and verify security information, including security identifiers (SIDs) for users and groups.

This section describes the inner workings of trusts, discusses the various types of trusts that can be used, provides a detailed list of referral and logon processes that involve trusts, and lists the network ports that trusts can use.

**Trust Architecture**

Applications integrated with Windows Server 2003 and Active Directory use components of the operating system to establish and maintain trusts. A number of these components help form the trust architecture that provides an effective communication infrastructure for Active Directory. These include authentication protocols, the Net Logon service, the Local Security Authority (LSA), and the Trusted Domain Objects (TDOs) stored in Active Directory. These trust components are shown in the following illustration.

**Trust Components**




**NTLM Protocol (Msv1_0.dll)**

The NTLM authentication protocol is dependent on the Net Logon service on domain controllers for client authentication and authorization information. This protocol authenticates clients that do not use Kerberos authentication. NTLM uses trusts to pass authentication requests between domains.

**Kerberos Protocol (Kerberos.dll)**
The Kerberos V5 authentication protocol is dependent on the Net Logon service on domain controllers for client authentication and authorization information. The Kerberos protocol connects to an online Key Distribution Center (KDC) and the Active Directory account store for session tickets. The Kerberos protocol also uses trusts for cross-realm ticket-granting services (TGS) and to validate Privilege Attribute Certificates (PACs) across a secured channel. The Kerberos protocol performs cross-realm authentication only with non-Windows-brand operating system Kerberos realms such as an MIT Kerberos realm and does not need to interact with the Net Logon service.

**Net Logon (Netlogon.dll)**
The Net Logon service maintains a secured channel from a Windows-based computer to a domain controller. It is also used in the following trust related processes:

- Trust setup and management – Net Logon helps maintain trust passwords, gathers trust information and verifies trusts by interacting with the LSA process and the TDO. For Forest trusts, the trust information includes the Forest Trust Information (FTInfo) record, which includes the set of namespaces that a trusted forest claims to manage, annotated with a field that indicates whether each claim is actually trusted by the trusting forest.

- Authentication – Supplies user credentials over a secured channel to a domain controller and returns the domain SIDs and user rights for the user

- Domain controller location – Helps with finding or locating domain controllers in a domain or across domains

- Pass-through validation – Credentials of users in other domains are processed by Net Logon. When a trusting domain needs to verify the identity of a user, it passes the user's credentials through Net Logon to the trusted domain for verification

- Privilege Attribute Certificate (PAC) verification – When a server using the Kerberos protocol for authentication needs to verify the PAC in a service ticket, it sends the PAC across the secure channel to its domain controller for verification.

**LSA (Lsasrv.dll)**
The Local Security Authority (LSA) is a protected subsystem that maintains information about all aspects of local security on a system (collectively known as local security policy) and provides various services for translation between names and identifiers. The LSA security subsystem provides services in both kernel mode and user mode for validating access to objects, checking user privileges, and generating audit messages. LSA is responsible for checking the validity of all session tickets presented by services in trusted or untrusted domains.

**Tools**
Administrators can use Active Directory Domains and Trusts, Netdom and Nltest to expose, create, remove or modify trusts. Active Directory Domains and Trusts is the Microsoft Management Console (MMC) that is used to administer domain trusts, domain and forest functional levels, and user principal name suffixes. The Netdom and Nltest command line tools can be used to find, display, create and manage trusts. These tools communicate directly with the LSA authority on a domain controller.

**Active Directory**
The Active Directory directory service stores information about objects on a network and makes this information available to users and network administrators. Trusts can be used to extend the reach of Active Directory domains and forests to other domains, forests or realms within an organization.

**Trusted Domain Object**
Each domain or forest trust within an organization is represented by a Trusted Domain Object (TDO) stored in the System container within its domain.

**TDO Contents**
The information contained in a TDO can vary depending on whether a TDO was created by a domain trust or by a forest trust. When a domain trust is created, attributes such as the DNS domain name, domain SID, trust type, trust transitivity, and the reciprocal domain name are represented in the TDO. Forest trust TDOs store additional attributes to identify all of the trusted namespaces from the partner forest. These attributes include domain tree names, user principal name (UPN) suffixes, service principal name (SPN) suffixes, and security ID (SID) namespaces.

Because trusts are stored in Active Directory as TDOs, all domains in a Windows Server 2003 forest have knowledge of the trust relationships that are in place throughout the forest. Similarly, when two or more forests are joined together through forest trusts, the forest root domains in each forest have knowledge of the trust relationships that are in place throughout all of the domains in trusted Windows Server 2003 forests. External trusts to a Windows NT 4.0 domain do not create TDOs in Active

Directory.

**TDO Passwords**

Both domains in a trust relationship share a password, which is stored in the TDO object in Active Directory. As part of the account maintenance process, every thirty days, the trusting domain controller changes the password stored in the TDO. Because all two-way trusts are actually two one-way trusts going in opposite directions, the process occurs twice for two-way trusts. A trust has a trusting and a trusted side. On the trusted side, any writable domain controller can be used for the process. On the trusting side, the PDC emulator performs the password change. To change a password, the domain controllers complete the following process:

1. The primary domain controller (PDC) emulator in the trusting domain creates a new password.

   A domain controller in the trusted domain never initiates the password change; it is always initiated by the trusting domain PDC emulator.

2. The PDC emulator in the trusting domain sets the OldPassword field of the TDO object to the current NewPassword field.

3. The PDC emulator in the trusting domain sets the NewPassword field of the TDO object to the new password.

   Keeping a copy of the previous password makes it possible to revert to the old password if the domain controller in the trusted domain fails to receive the change, or if the change is not replicated before a request is made that uses the new trust password.

4. The PDC emulator in the trusting domain makes a remote call to a domain controller in the trusted domain asking it to set the password on the trust account to the new password.

5. The domain controller in the trusted domain changes the trust password to the new password.

6. On each side of the trust, the updates are replicated to the other domain controllers in the domain. In the trusting domain, the change triggers an urgent replication of the trusted domain object.

The password is now changed on both domain controllers. Normal replication distributes the TDO objects to the other domain controllers in the domain. However, is possible for the domain controller in the trusting domain to change the password without successfully updating a domain controller in the trusted domain. This might occur because a secured channel, which is required to process the password change, could not be established. It is also possible that the domain controller in the trusted domain might be unavailable at some point during the process and might not receive the updated password.

To deal with situations in which the password change is not successfully communicated, the domain controller in the trusting domain never changes the new password unless it has successfully authenticated (set up a secured channel) using the new password. This is why both the old and new passwords are kept in the TDO object of the trusting domain. Because a password change is not finalized until authentication using the password succeeds, the old, stored password can be used over the secured channel until the domain controller in the trusted domain receives the new password, thus enabling uninterrupted service.

If authentication using the new password fails because the password is invalid, the trusting domain controller tries to authenticate using the old password. If it authenticates successfully with the old password, it resumes the password change process within 15 minutes.

Trust password updates need to replicate to the domain controllers of both sides of the trust within 30 days. If the trust password is changed after 30 days and a domain controller then only has the N-2 password, it cannot use the trust from the trusting side and cannot create a secure channel on the trusted side.

**Practical Limitations of Trusts**

While Windows Server 2003 can support a large number of trusts between many security authorities, and there are some limitations to trust functionality. These limitations are associated with the number of TDOs, the length of trust paths, and the ability of clients to discover available trusts.

**Number of TDOs**

TDOs are associated with domains, and as the number of TDOs increases, the processing performance of these links declines. Testing indicates that the time to complete operations related to TDOs, such as authentication across domains, deteriorates noticeably if the Active Directory implementation in an organization contains more than 2400 TDOs. Few organizations,

however, require such a large number of trusts, and Windows Server 2003 places no absolute limit on the number of trust links that a domain can maintain with other domains or forests.

**Trust Path Limits**
There is a limit to the number of trust links that can be used effectively in a Windows Server 2003 network. In Windows Server 2003, Kerberos clients can traverse a maximum of 10 trust links to locate a requested resource in another domain. If the trust path between the domains exceeds this limit, the attempt to access the domain fails. This problem can be addressed by creating external trusts to target domains in order to bypass excessively long trust paths.

**Trust Search Limits**
When a client searches out a trust path, the search is limited to trusts that are established directly with a domain and those that are transitive within a forest. A child domain, for example, cannot use an external trust between its parent domain and a domain in another forest. This is because a complete trust path must be constructed before a client can begin to work its way along the path to a requested resource domain. Windows Server 2003 does not support blind referrals; if a client cannot identify a trust path, it will not attempt to seek access to a requested resource in another domain. Child domains cannot identify external trusts or realm trusts to security authorities outside of their forest because the TDOs that represent those trusts are published only within the domain that shares the trust, not to Global Catalogs. Clients are therefore unaware of trusts that their domains share with security authorities other than their parent or child domains, and are thus unable to use them to access resources.

**Trust Flow**
The flow of secured communications over trusts determines the elasticity of a trust: how you create or configure a trust determines how far the communication extends within a forest or across forests. The flow of communication over trusts is determined by the direction of the trust (one-way or two-way) and the transitivity of the trust (transitive or nontransitive).

**One-Way and Two-Way Trusts**
Trust relationships that are established to enable access to resources can be either one-way or two-way. A one-way trust is a unidirectional authentication path created between two domains. In a one-way trust between Domain A and Domain B, users in Domain A can access resources in Domain B. However, users in Domain B cannot access resources in Domain A. Some one-way trusts can be either nontransitive or transitive depending on the type of trust being created.

All domain trusts in an Active Directory forest are two-way, transitive trusts. When a new child domain is created, a two-way, transitive trust is automatically created between the new child domain and the parent domain. In a two-way trust, Domain A trusts Domain B and Domain B trusts Domain A. This means that authentication requests can be passed between the two domains in both directions. Some two-way relationships can be nontransitive or transitive depending on the type of trust being created. An Active Directory domain can establish a one-way or two-way trust with:

- Windows Server 2003 domains in the same forest.
- Windows Server 2003 domains in a different forest.
- Windows NT 4.0 domains.
- Kerberos V5 realms.

**Transitive and Nontransitive Trusts**
Transitivity determines whether a trust can be extended outside of the two domains with which it was formed. A transitive trust can be used to extend trust relationships with other domains; a nontransitive trust can be used to deny trust relationships with other domains.

Each time you create a new domain in a forest, a two-way, transitive trust relationship is automatically created between the new domain and its parent domain. If child domains are added to the new domain, the trust path flows upward through the domain hierarchy extending the initial trust path created between the new domain and its parent domain. Transitive trust relationships flow upward through a domain tree as it is formed, creating transitive trusts between all domains in the domain tree.

Authentication requests follow these trust paths, so accounts from any domain in the forest can be authenticated by any other domain in the forest. With a single logon process, accounts with the proper permissions can access resources in any domain in the forest. The following figure shows that all domains in Tree 1 and Tree 2 have transitive trust relationships by default. As a result, users in Tree 1 can access resources in domains in Tree 2 and users in Tree 1 can access resources in Tree 2, when the proper permissions are assigned at the resource.

### Default Transitive Trust Relationships




In addition to the default transitive trusts established in a Windows Server 2003 forest, by using the New Trust Wizard you can manually create the following transitive trusts.

- **Shortcut trust.** A transitive trust between domains in the same domain tree or forest that is used to shorten the trust path in a large and complex domain tree or forest.
- **Forest trust.** A transitive trust between one forest root domain and another forest root domain.
- **Realm trust.** A transitive trust between an Active Directory domain and a Kerberos V5 realm.

A nontransitive trust is restricted to the two domains in the trust relationship and does not flow to any other domains in the forest. A nontransitive trust can be a two-way trust or a one-way trust.

Nontransitive trusts are one-way by default, although you can also create a two-way relationship by creating two one-way trusts. Nontransitive domain trusts are the only form of trust relationship possible between:

- A Windows Server 2003 domain and a Windows NT domain
- A Windows Server 2003 domain in one forest and a domain in another forest (when not joined by a forest trust)

By using the New Trust Wizard, you can manually create the following nontransitive trusts:

- **External trust.** A nontransitive trust created between a Windows Server 2003 domain and a Windows NT, Windows 2000, or Windows Server 2003 domain in another forest. When you upgrade a Windows NT domain to a Windows Server 2003 domain, all existing Windows NT trusts are preserved intact. All trust relationships between Windows Server 2003 domains and Windows NT domains are nontransitive.
- **Realm trust.** A nontransitive trust between an Active Directory domain and a Kerberos V5 realm.

**Trust Types**

Although all trusts enable authenticated access to resources, trusts can have different characteristics. The types of domains included in the trust relationship affect the type of trust that is created. For example, a trust between two child domains in different forests is always an external trust, but trusts between two Windows Server 2003 forest root domains can be either external trusts or forest trusts.

Two types of trusts are created automatically when you use the Active Directory Installation Wizard. Four other types of trusts can be manually created by using either the New Trust Wizard or the Netdom command-line tool.

**Automatic Trusts**
By default, two-way transitive trusts are automatically created when a new domain is added to a domain tree or forest root domain by using the Active Directory Installation Wizard. The two default trust types are parent-child trusts and tree-root trusts.

**Parent-child trust**
A parent-child trust relationship is established whenever a new domain is created in a tree. The Active Directory installation process automatically creates a trust relationship between the new domain and the domain that immediately precedes it in the namespace hierarchy (for example, corp.tailspintoys.com is created as the child of tailspintoys.com). The parent-child trust relationship has the following characteristics:

- It can exist only between two domains in the same tree and namespace.
- The parent domain is always trusted by the child domain.
- It must be transitive and two-way. The bidirectional nature of transitive trust relationships allows the global directory information in Active Directory to replicate throughout the hierarchy.

**Tree-root trust**
A tree-root trust is established when you add a new domain tree to a forest. The Active Directory installation process automatically creates a trust relationship between the domain you are creating (the new tree root) and the forest root domain. A tree-root trust relationship has the following restrictions:

- It can be established only between the roots of two trees in the same forest.
- It must be transitive and two-way.

**Manual Trusts**
In Windows Server 2003 there are four trust types that must be created manually: shortcut trusts are used for optimization between domain trees in the same forest; external, realm, and forest trusts help provide interoperability with domains outside of the forest, with other forests, or with realms. These trust types must be created by using the New Trust Wizard or the Netdom command-line tool.

**Shortcut Trusts**
Shortcut trusts are one-way or two-way transitive trusts that can be used when administrators need to optimize the authentication process. Authentication requests must initially travel a trust path between domain trees. A trust path is the series of domain trust relationships that must be traversed in order to pass authentication requests between any two domains. In a complex forest, the time required to traverse the trust path can impact performance. You can significantly reduce this time by using shortcut trusts.

Shortcut trusts speed up logon and access times to resources in a domain that is deep within the hierarchy of another domain tree. The following figure illustrates trust relationships between two trees in a Windows Server 2003 forest.

**Trusts Within a Single Windows 2000 Server or Windows Server 2003 Forest**




In this example, the following conditions are true:

- A two-way transitive tree-root trust exists between corp.tailspintoys.com and corp.wingtiptoys.com.
- Parent-child trust relationships exist between parent and child domains in both trees of the forest. Because the two trees trust each other, all domains in Tree 1 and Tree 2 have access to resources in all domains in the forest.
- A one-way shortcut trust exists between the rome.europe.corp.tailspintoys.com and usa.corp.wingtiptoys.com domains. This creates an alternate path by which users in the usa.corp.wingtiptoys.com domain can access resources in the rome.europe.corp.tailspintoys.com domain. The shortcut trust shortens the trust path needed to reach the intended domain by cutting the number of hops required for authentication from three (rome.europe.corp.tailspintoys.com to europe.corp.tailspintoys.com to corp.tailspintoys.com to corp.wingtiptoys.com) to one (rome.europe.corp.tailspintoys.com to usa.corp.wingtiptoys.com), which increases the speed of authentication.
- Because the shortcut trust is one-way, authentication requests made from the rome.europe.corp.tailspintoys.com domain to the usa.corp.wingtiptoys.com domain still need to travel the longer trust path.

When a two-way shortcut trust is established between two domains located in separate domain trees, the time needed to fulfill authentication requests originating in either domain is reduced. For example, when a two-way trust is established between the usa.corp.wingtiptoys.com domain and the rome.europe.corp.tailspintoys.com domain, authentication requests made from either domain to the other can utilize the new, two-way shortcut trust path.

**External Trusts**

An external trust is a trust relationship that can be created between Active Directory domains that are in different forests or between an Active Directory domain and a Windows NT 4.0 or earlier domain. An external trust relationship has the following characteristics:

- It is nontransitive.
- It must be established manually in each direction to create a two-way external trust relationship. In Windows Server 2003 you can create both sides of the external two-way trust at once by using the New Trust Wizard.
- It enforces SID filter quarantining by default in Windows Server 2003. External trusts created from the trusting domain use SID filter quarantining to verify that incoming authentication requests made from security principals in the trusted domain contain only SIDs of security principals in the trusted domain. SID filter quarantining ensures that any misuse of the SID history attribute on security principals (including inetOrgPerson) in the trusted forest cannot pose a threat to the integrity of the trusting forest.

External trusts provide access to resources in a domain outside of the forest that is not already joined by a forest trust. The following illustration shows how external trusts can be used between a Windows Server 2003 forest and a Windows 2000

forest.

**External trusts Between a Windows Server 2003 Forest and a Windows 2000 Forest**




In this example, two external trust relationships exist between domains in the Windows Server 2003 forest and the Windows 2000 forest. The direction of the one-way external trust arrow indicates that the sales.corp.worldwideimporters.com domain trusts the rome.europe.corp.tailspintoys.com domain, which means that users in the rome.europe.corp.tailspintoys.com domain can access resources in the sales.corp.worldwideimporters.com domain.

The direction of the two-way external trust arrow indicates that both the europe.corp.tailspintoys.com domain and the sales.corp.worldwideimporters.com domain trust each other. This relationship allows for authentication requests to be passed between the two domains, coming from either direction, to any shared resources in those two domains.

This configuration allows:

- Users in the europe.corp.tailspintoys.com domain to access resources in the sales.corp.worldwideimporters.com domain
- Users in the sales.corp.worldwideimporters.com domain to access resources in the europe.corp.tailspintoys.com domain
- Users in the rome.europe.corp.tailspintoys.com domain to access resources in the sales.corp.worldwideimporters.com domain

It does not allow:

- Users in the rome.europe.corp.tailspintoys.com domain or europe.corp.tailspintoys.com domain to access resources in the corp.worldwideimporters.com domain
- Users in the sales.corp.worldwideimporters.com domain to access resources in either the corp.tailspintoys.com or rome.europe.corp.tailspintoys.com domain

When a trust is established between a domain in a forest and a domain outside of that forest, security principals from the external domain can access resources in the internal domain. Active Directory creates a foreign security principal object in the internal domain to represent each security principal from the trusted external domain. These foreign security principals can become members of domain local groups in the internal domain. Directory objects for foreign security principals are created by Active Directory and should not be manually modified. You can view foreign security principal objects from Active Directory Users and Computers by enabling advanced features.

The following figure illustrates a trust configuration that includes an external trust relationship between a domain in a Windows Server 2003 forest and a Windows NT 4.0 domain.

**External Trust Between a Windows NT 4.0 Domain and a Child Domain in a Windows Server 2003 Forest**




This configuration allows users in the europe.corp.tailspintoys.com domain to access resources in the Windows NT 4.0 domain. It does not allow the Windows NT 4.0 domain to access resources in the europe.corp.tailspintoys.com domain or for any trusted domains that the europe.corp.tailspintoys.com domain trusts.

All forest trusts in Windows Server 2003 Active Directory enforce SID filtering. Applying SID filtering to trusts helps to prevent malicious users who have domain administrator level access in the trusted domain from granting to themselves, or other user accounts in their domain, elevated user rights to the trusting domain.

**Realm Trusts**

A trust relationship can be established between a non-Windows Kerberos realm and a Windows 2000 or Windows Server 2003 domain. This trust relationship allows cross-platform interoperability with security services based on other Kerberos V5 implementations.

A non-Windows Kerberos realm trust relationship has the following characteristics:

- It is used only by the Kerberos V5 authentication protocol. It is not used by NTLM or other authentication protocols.
- It is one-way by default. One-way trust relationships must be established in each direction to create a two-way trust relationship. In Windows Server 2003 you can create both sides of the two-way realm trust at once by using the New Trust Wizard.
- It is nontransitive by default, but can be made transitive.
- When the direction of trust is from a non-Windows Kerberos realm to a Windows Server 2003 domain, the non-Windows Kerberos realm trusts all security principals in the Windows Server 2003 domain.
- When the direction of trust is from a Windows Server 2003 domain to a non-Windows Kerberos realm, account mappings in Active Directory are used to map a foreign Kerberos identity in a trusted non-Windows Kerberos realm to a local account identity in a Windows Server 2003 domain.
- The Windows Server 2003 domain uses only the account to which the non-Windows principal is mapped (the proxy account) to evaluate access to domain objects that have security descriptors. This is required because non-Windows Kerberos tickets do not contain all of the authorization data that is needed for Windows Server 2003. All such Windows Server 2003 proxy accounts can be used in groups and on access control lists (ACLs) to control access on behalf of the non-Windows security principal. MIT account mappings are managed through Active Directory Users and Computers.
- When the direction of trust is from a non-Windows Kerberos realm to a Windows Server 2003 domain, users in the Windows Server 2003 domain can access resources in the non-Windows Kerberos realm, as is the case for Realm 1 in the following illustration. When the direction of trust is from a Windows Server 2003 domain to a non-Windows Kerberos realm, users in the non-Windows Kerberos realm can access the resources in the Windows Server 2003 domain, as is the case for Realm 1 and Realm 2.

**Transitive and Non-Transitive Realm Trusts from a Domain in a Windows Server 2003 Forest**




This configuration allows:

- Users in Realm 1 and Realm 2 to access resources in the europe.corp.tailspintoys.com domain
- Users in the europe.corp.tailspintoys.com domain to access resources in Realm 1

It does not allow users in the europe.corp.tailspintoys.com domain to access resources in Realm 2.

**Note**

- In Windows 2000, if you create a non-Windows Kerberos realm trust relationship by using Active Directory Domains and Trusts, the trust is one-way and nontransitive. To work around this, you can use the Netdom tool (Netdom.exe) to establish two-way, transitive, non-Windows Kerberos realm trust relationships. In Windows Server 2003, you can use the New Trust Wizard to create one-way or two-way and transitive or nontransitive realm trusts.

Realm trusts are not as comprehensive as domain trusts. For instance, user principals from the Kerberos realm do not have the authorization data that Windows Server 2003-based services need for access control. In order for this authorization data to be added to the user's ticket, an account mapping mechanism is used. Selected domain user accounts are used to provide identity for Kerberos principals in the trusted realm. These mappings are kept on the SecurityID property on the user account object. They can be managed through Active Directory Users and Computers.

**Forest Trusts**

Forest trusts help you to manage a segmented Windows Server 2003 Active Directory infrastructure within your organization by providing support for accessing resources and other objects across multiple Windows Server 2003 forests. Forest trusts are useful for Application Service Providers, companies undergoing mergers or acquisitions, collaborative business extranets, and companies seeking a solution for administrative autonomy.

Forest trusts can provide the following benefits:

- Simplified management of resources across two Windows Server 2003 forests by reducing the number of external trusts necessary to share resources.
- Complete two-way trust relationships with every domain in each forest.
- Use of user principal name (UPN) authentication across two forests.
- Use of both the Kerberos V5 and NTLM authentication protocols to improve the trustworthiness of authorization data transferred between forests.
- Flexibility of administration. Administrative tasks can be unique to each forest.
- Secured data within each forest. Sensitive data can be protected so that only users within that forest can access it.
- Isolated directory replication within each forest. Schema changes, configuration changes, and the addition of new domains to a forest have forest-wide impact only within that forest, not in a trusting forest.
- Enforcement of SID filtering in Windows Server 2003. SID filtering ensures that any misuse of the SID history attribute on

security principals (including inetOrgPerson) in the trusted forest cannot pose a threat to the integrity of the trusting forest.

Enforcing SID filtering on forest trusts does not prevent use of SID history in migrations to domains within the same forest and does not affect your universal group access control strategy.

By using forest trusts you can link two different Windows Server 2003 forests to form a one-way or two-way transitive trust relationship. A forest trust allows administrators to federate two Active Directory forests with a single trust relationship to provide a seamless authentication and authorization experience across the forests.

A forest trust can be created only between a forest root domain in one Windows Server 2003 forest and a forest root domain in another Windows Server 2003 forest. Forest trusts can be created between two forests only and cannot be implicitly extended to a third forest. This means that if a forest trust is created between Forest 1 and Forest 2, and another forest trust is created between Forest 2 and Forest 3, Forest 1 does not have an implicit trust with Forest 3. The following figure shows two separate forest trust relationships between three Windows Server 2003 forests in a single organization.

**Two Forest Trusts Between Three Windows Server 2003 Forests**




In this example, a two-way transitive forest trust exists between the forest root domains in Forest 1 and Forest 2, and another two-way transitive forest trust exists between the forest root domains in Forest 3 and Forest 2. This configuration allows:

- Users in Forest 2 to access resources in any domain in either Forest 1 or Forest 3
- Users in Forest 3 to access resources in any domain in Forest 2
- Users in Forest 1 to access resources in any domain in Forest 2

This configuration does not allow users in Forest 1 to access resources in Forest 3 or vice versa. To allow users in both Forest 1 and Forest 3 to share resources, a two-way transitive trust must be created between the two forests.

If a one-way forest trust is created between two forests, members of the trusted forest can utilize resources located in the trusting forest. However, the trust operates in only one direction. For example, when a one-way, forest trust is created between Forest 1 (the trusted forest) and Forest 2 (the trusting forest), members of Forest 1 can access resources located in Forest 2, but members of Forest 2 cannot access resources located in Forest 1 using the same trust.

Consequently, as shown in the example, a two-way forest trust between two forests allows members from either forest to utilize resources located in the other forest; domains in each respective forest trust domains in the other forest implicitly.

There are specific requirements that must be met for using forest trusts. Before you can create a forest trust, you need to verify that all domain controllers in both forests are running Windows Server 2003. You also need to verify that you have the correct Domain Name System (DNS) infrastructure in place and that you have established the appropriate functionality level for Active Directory. This means that the forest must be raised to the Windows Server 2003 functional level and that you cannot install additional Windows 2000 or Windows NT Server 4.0 domain controllers in the forest.

Forest trusts can only be created when one of the following DNS configurations is in place in your infrastructure:

- A single root DNS server is the root DNS server for both forest DNS namespaces: the root zone contains delegations for each of the DNS namespaces and the root hints of all DNS servers include the root DNS server.
- Where there is no shared root DNS server, and the root DNS servers for each forest DNS namespace are running a member of the Windows Server 2003 family, DNS conditional forwarders are configured in each DNS namespace to route queries for names in the other namespace.
- Where there is no shared root DNS server, and the root DNS servers for each forest DNS namespace are not running a member of the Windows Server 2003 family, DNS secondary zones are configured in each DNS namespace to route

queries for names in the other namespace.

To create a forest trust, the administrator creating the trust must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory. Each trust is assigned a password that the administrators in both forests must know. Members of Enterprise Admins in both forests can create the trusts in both forests at once and, in this scenario, a password that is cryptographically random is automatically generated and written for both forests.

Members of the Incoming Forest Trust Builders group can create one-way, incoming forest trusts. For example, members of this group residing in Forest A can create a one-way, incoming forest trust from Forest B. This one-way, incoming forest trust allows users in Forest A to access resources located in Forest B. Members of this group are granted the permission Create Inbound Forest Trust on the forest root domain. This group has no default members.

All Windows NT Server 4.0, Windows 2000, Windows XP clients can use the forest trust for authentication. However, Windows NT Server 4.0 clients support only NTLM and they do not support user principal names (UPNs) for logging on to the network. Windows Server 2003 forest trusts cannot be created between a Windows Server 2003 forest and a Windows 2000 forest.

**Trust Processes and Interactions**
Many inter-domain and inter-forest transactions depend on domain or forest trusts in order to complete various tasks. This section describes the processes and interactions that occur as resources are accessed across trusts and authentication referrals are evaluated.

**Overview of Authentication Referral Processing**
When a request for authentication is referred to a domain, the domain controller in that domain must determine whether a trust relationship exists with the domain from which the request comes, as well as the direction of the trust and whether the trust is transitive or nontransitive, before it authenticates the user to access resources in the domain. The authentication process that occurs between trusted domains varies according to the authentication protocol in use. The Kerberos V5 and NTLM protocols in Windows 2000 Server and Windows Server 2003 process referrals for authentication to a domain differently, as do the other authentication protocols, such as Digest, and SChannel, that Windows 2000 Server and Windows Server 2003 support.

**Kerberos V5 Referral Processing**
If the client uses Kerberos V5 for authentication, it requests a ticket to the server in the target domain from a domain controller in its account domain. The Kerberos Key Distribution Center (KDC) acts as a trusted intermediary between the client and server; it provides a session key that enables the two parties to authenticate each other. If the target domain is different from the current domain, the KDC follows a logical process to determine whether an authentication request can be referred:

1. Is the current domain trusted directly by the domain of the server that is being requested?

    - If yes, send the client a referral to the requested domain.
    - If no, go to the next step.

2. Does a transitive trust relationship exist between the current domain and the next domain on the trust path?

    - If yes, send the client a referral to the next domain on the trust path.
    - If no, send the client a logon-denied message.

**NTLM Referral Processing**
If the client uses NTLM for authentication, the initial request for authentication goes directly from the client to the resource server in the target domain. This server creates a challenge to which the client responds. The server then sends the user's response to a domain controller in its computer account domain. This domain controller checks the user account against its security accounts database. If the account does not exist in the database, the domain controller determines whether to perform pass-through authentication, forward the request, or deny the request by using the following logic:

1. Does the current domain have a direct trust relationship with the user's domain?

    - If yes, the domain controller sends the credentials of the client to a domain controller in the user's domain for pass-

through authentication.

- If no, go to the next step.

2. Does the current domain have a transitive trust relationship with the user's domain?

    - If yes, pass the authentication request on to the next domain in the trust path. This domain controller repeats the process by checking the user's credentials against its own security accounts database.

    - If no, send the client a logon-denied message.

**Other Authentication Protocol Referral Processing**
Windows 2000 Server and Windows Server 2003 also support both Digest and Schannel authentication across forest trusts. This enables technologies such as Web servers to complete their tasks in a multiple domain environment. The Secure Sockets Layer (SSL) protocol and certificate mapping to Active Directory user accounts use Schannel, and Internet Information Services (IIS) uses Digest. Digest and SChannel authentication between trusted domains might be required in cases where a firewall between trusting forests or domains allows only HTTP traffic through. Digest can be carried in HTTP headers.

The Windows 2000 Server and Windows Server 2003 authentication negotiation determines the authentication protocol that is used over trusts, as long as the generic pass-through mechanism can find the appropriate domain controller in the target domain. Administrators can also choose to implement a security support provider (SSP) for another authentication protocol. If the SSP is compatible with the Windows 2000 Server and Windows Server 2003 distributed systems, it will work over trusts across domain boundaries in the same way that standard Windows 2000 Server and Windows Server 2003 SSPs work across domains.

**Kerberos-Based Processing of Authentication Requests Over Domain Trusts**
Active Directory domains within the same forest are implicitly connected by two-way, transitive trusts, so that authentication requests from one domain to another are routed to provide a seamless access to resources across domains. Domain controllers running Windows 2000 or Windows Server 2003 authenticate users and applications using one of two authentication protocols: Kerberos V5 or NTLM. Once authenticated in their own domain, users can attempt to gain access to resources from any domain in the forest using the Active Directory authorization process.

To access a shared resource in another domain by using Kerberos authentication, a user's computer first requests a ticket from a domain controller in its account domain to the server in the trusting domain that hosts the requested resource. This ticket is then issued by an intermediary trusted by both the user's computer and the server. The user's computer presents this trusted ticket to the server in the trusting domain for authentication.

The following illustration and corresponding steps provide a detailed description of the Kerberos authentication process that is used when a computer running Windows 2000 Professional, Windows 2000 Server, Windows XP Professional, or a member of the Windows Server 2003 family attempts to access resources from a server located in another domain.

**Kerberos Authentication Process Over Domain Trusts**




1. User1 logs on to Workstation1 using credentials from the europe.corp.tailspintoys.com domain. As part of this process, the authenticating domain controller issues User1 a ticket-granting ticket (TGT). This ticket is required for User1 to be authenticated to resources. The user then attempts to access a shared resource (\\fileserver1\share) on a file server located in the asia.corp.tailspintoys.com domain.

2. Workstation1 contacts the Kerberos Key Distribution Center (KDC) on a domain controller in its domain (ChildDC1) and requests a service ticket for the FileServer1 service principal name (SPN).

3. ChildDC1 does not find the SPN in its domain database and queries the global catalog to see if any domains in the forest contain this SPN. The global catalog sends the requested information back to ChildDC1.

4. ChildDC1 sends the referral to Workstation1.

5. Workstation1 contacts a domain controller in ForestRootDC1 (its parent domain) for a referral to a domain controller (ChildDC2) in the Child2 domain. ForestRootDC1 sends the referral to Workstation 1.

6. Workstation1 contacts the KDC on ChildDC2 and negotiates a ticket for the user to gain access to FileServer1.

7. Once Workstation1 has a service ticket, it sends the ticket to FileServer1, which reads the user's security credentials and constructs an access token accordingly.

Each domain has its own set of security policies governing access to resources. These policies do not cross from one domain to another.

**Kerberos-Based Processing of Authentication Requests Over Forest Trusts**
When two Windows Server 2003 forests are connected by a forest trust, authentication requests made using the Kerberos V5 or NTLM protocols can be routed between forests to provide access to resources in both forests.

When a forest trust is first established, each forest collects all of the trusted namespaces in its partner forest and stores the information in a TDO. Trusted namespaces include domain tree names, user principal name (UPN) suffixes, service principal name (SPN) suffixes, and security ID (SID) namespaces used in the other forest. TDO objects are replicated to the global catalog.

Before authentication protocols can follow the forest trust path, the service principal name (SPN) of the resource computer must be resolved to a location in the other forest. An SPN can be one of the following: the Domain Name System (DNS) name of a host, the DNS name of a domain, or the distinguished name of a service connection point object.

When a workstation in one forest attempts to access data on a resource computer in another forest, the Kerberos authentication process contacts the domain controller for a service ticket to the SPN of the resource computer. Once the domain controller queries the global catalog and determines that the SPN is not in the same forest as the domain controller, the domain controller sends a referral for its parent domain back to the workstation. At that point, the workstation queries the

parent domain for the service ticket and continues to follow the referral chain until it reaches the domain where the resource is located.

The following illustration and corresponding steps provide a detailed description of the Kerberos authentication process that is used when computers running Windows 2000 Professional, Windows XP Professional, Windows 2000 Server, or a member of the Windows Server 2003 family attempt to access resources from a computer located in another forest.

**Kerberos Authentication Process Over Forest Trusts**




1. User1 logs on to Workstation1 using credentials from the europe.corp.tailspintoys.com domain. The user then attempts to access a shared resource on FileServer1 located in the usa.corp.wingtiptoys.com forest.

2. Workstation1 contacts the Kerberos Key Distribution Center (KDC) on a domain controller in its domain (ChildDC1) and requests a service ticket for the FileServer1 SPN.

3. ChildDC1 does not find the SPN in its domain database and queries the global catalog to see if any domains in the tailspintoys.com forest contain this SPN. Because a global catalog is limited to its own forest, the SPN is not found. The global catalog then checks its database for information about any forest trusts that are established with its forest, and, if found, it compares the name suffixes listed in the forest trust trusted domain object (TDO) to the suffix of the target SPN to find a match. Once a match is found, the global catalog provides a routing hint back to ChildDC1. Routing hints help direct authentication requests toward the destination forest, and are only used when all traditional authentication channels (local domain controller and then global catalog) fail to locate a SPN.

4. ChildDC1 sends a referral for its parent domain back to Workstation1.

5. Workstation1 contacts a domain controller in ForestRootDC1 (its parent domain) for a referral to a domain controller (ForestRootDC2) in the forest root domain of the wingtiptoys.com forest.

6. Workstation1 contacts ForestRootDC2 in the wingtiptoys.com forest for a service ticket to the requested service.

7. ForestRootDC2 contacts its global catalog to find the SPN, and the global catalog finds a match for the SPN and sends it back to ForestRootDC2.

8. ForestRootDC2 then sends the referral to usa.corp.wingtiptoys.com back to Workstation1.

9. Workstation1 contacts the KDC on ChildDC2 and negotiates the ticket for User1 to gain access to FileServer1.

10. Once Workstation1 has a service ticket, it sends the service ticket to FileServer1, which reads User1's security credentials and constructs an access token accordingly.

**Network Ports used by Trusts**
Because trusts must be deployed across various network boundaries, they might have to span one or more firewalls. When this is the case, you can either tunnel trust traffic across a firewall or open specific ports in the firewall to allow the traffic to pass through. In Windows Server 2003, this procedure is simplified through configurable remote procedure call (RPC) ports for the main trust services. The two main configurable RPC ports are the Local Security Authority RPC port and the Net Logon RPC port.

**The Local Security Authority (LSA) RPC port.**
This port is used for trust creation and other access to the LSA policy database.

**The Net Logon RPC port.**
This port is used for NTLM authentication and secured channel communications.

The following table shows the list of ports that might need to be opened before you establish trusts.

### Ports Required for Trusts

| Task | Outbound Ports | Inbound Ports | From–To |
|---|---|---|---|
| Set up trusts on both sides from the internal forest | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution — portmapper (135 TCP)<br>Net Logon fixed port | N/A | Internal domain domain controllers–External domain domain controllers (all ports) |
| Trust validation from the internal forest domain controller to the external forest domain controller (outgoing trust only) | LDAP (389 UDP)<br><br>Microsoft SMB (445 TCP)<br><br>Endpoint resolution — portmapper (135 TCP)<br>Net Logon fixed port | N/A | Internal domain domain controllers–External domain domain controllers (all ports) |
| Use Object picker on the external forest to add objects that are in an internal forest to groups and DACLs | N/A | LDAP (389 UDP and TCP)<br><br>Windows NT Server 4.0 directory service fixed port<br><br>Net Logon fixed port<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution portmapper (135 TCP) | External server–Internal domain PDCs (Kerberos)<br><br>External domain domain controllers–Internal domain domain controllers (Net Logon) |
| Set up trust on the external forest from the external forest | N/A | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP) | External domain domain controllers–Internal domain domain controllers (all ports) |
| Use Kerberos authentication (internal forest client to external forest) | Kerberos (88 UDP) | N/A | Internal client–External domain domain controllers (all ports) |

| | | | |
|---|---|---|---|
| Use NTLM authentication (internal forest client to external forest) | N/A | Endpoint resolution – portmapper (135 TCP)<br>Net Logon fixed port | External domain domain controllers–Internal domain domain controllers (all ports) |
| Join a domain from a computer in the internal network to an external domain | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution — portmapper (135 TCP)<br>Net Logon fixed port<br><br>Windows NT Server 4.0 directory service fixed port | N/A | Internal client–External domain domain controllers (all ports) |

| **Note** |
|---|
| Port 445 and potentially TCP port 139 are required for backward compatibility. |

To specify the services that you want to run on a fixed port, you must appropriately configure the registry for that port.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

Settings for the Local Security Authority (LSA) RPC port are stored in the **TCP/IP Port** entry in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** registry key.

Settings for the Net Logon RPC port are stored in the **DCTcpipPort** entry in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters** registry key.

# Domain and Forest Trust Tools and Settings

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Trust Tools
- Trust WMI Classes
- Network Ports used by Trusts
- Related Information


Administrators can use a number of methods to configure and manage trust relationships in Active Directory environments. This section discusses the tools that can be used to create, view and modify trusts, as well as the WMI classes and network ports that are associated with trusts.

**Trust Tools**

The two principal Microsoft tools that can be used to create and manage trusts are Active Directory Domains and Trusts (Domain.msc) and Windows Domain Manager, also called Netdom (Netdom.exe). In addition, there are three other Microsoft tools that administrators can use to help troubleshoot trust related issues. These are Nltest (Nltest.exe), Network Connectivity Tester, also called Netdiag (Netdiag.exe) and the Domain Controller Diagnostic tool, also called Dcdiag (Dcdiag.exe). Each trust management tool has its own primary purpose, but can also be used for other purposes. The following table briefly describes these tools and their different purposes.

**Trust Tools and Purposes**

| Microsoft Tool | Primary Purpose | Other Purpose |
|---|---|---|
| Active Directory Domains and Trusts | Create and manage trusts | Manage functional levels (Windows Server 2003 only) and user principal name suffixes |
| Netdom | Manage domains and trust relationships from the command-line | Join computers to a domain and manage computer accounts |
| Nltest | Test secured channels | View or record trust information and verify trusts |
| Netdiag | Test network health | Help troubleshoot client connectivity by testing the secured channel |
| Dcdiag | Test domain controller health | Verify some trust information and monitor trust related replication |

Netdom, Nltest, Dcdiag, and Netdiag are command-line support tools that can be installed from the Support\Tools folder on the Windows Server 2003 product CD. (To do this, in the Support\Tools folder, right-click the Suptools.msi file and then click Install.) The following table lists common trust related tasks and indicates which tools can be used to perform them. Detailed descriptions of each tool follow the table.

**Trust Tools Comparison by Task**

| Trust Task | Active Directory Domains and Trusts | Netdom | NLTest | Netdiag | Dcdiag |
|---|---|---|---|---|---|
| Create an external trust | • | • | | | |
| Create a realm trust | • | • | | | |

| | | | | | |
|---|---|---|---|---|---|
| Create a shortcut trust | ● | ● | | | |
| Create a forest trust | ● | | | | |
| Create both sides of a trust at once | ● | ● | | | |
| Troubleshoot client network connectivity over a secured channel | | | | ● | |
| Batch manage trusts | | ● | | | |
| View and record all trust relationships | ● | | ● | | |
| Change transitivity of a non-Windows Kerberos realm trust | ● | ● | | | |
| Change the direction of a trust (one-way to two-way and back) | ● | ● | | | |
| Enable or disable selective authentication for external and forest trusts | ● | ● | | | |
| Change the routing status of a name suffix | ● | ● | | | |
| Enable or disable an existing name suffix from routing | ● | ● | | | |
| Verify all trust types | ● | ● | ● | | |
| Verify only external trusts | | | | | ● |
| Reset a secured channel | | ● | ● | | |
| Remove a trust | ● | ● | | | |
| Enable or disable SID Filtering | ● | ● | | | |
| Display the contents of a Forest Trust Information (FTInfo) record | | ● | | | |

**Domain.msc: Active Directory Domains and Trusts**

**Category**
Active Directory Domains and Trusts is a Microsoft Management Console (MMC) snap-in that is installed automatically on computers running Windows Server 2003 when you install Active Directory. You can open Active Directory Domains and Trusts by clicking **Start**, **Programs**, **Administrative Tools,** and clicking **Active Directory Domains and Trusts.**

**Version compatibility**
This tool is compatible with domain controllers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. You can also use Active Directory Domains and Trusts with the Windows Server 2003 Administration Tools pack to remotely administer Active Directory from a computer that is not a domain controller, such as a computer running Windows XP Professional. You can use this tool to target any domain controller in either Windows 2000 Server or Windows Server 2003 domains.

### Note

- When using Windows Server 2003 Active Directory administrative tools to connect to a domain controller running Windows 2000 you must first make sure that the Windows 2000–based domain controller to which your are connecting has Service Pack 3 or later installed. This is because Windows Server 2003 administrative tools sign and encrypt all LDAP traffic by default. If business reasons do not permit the installation of Service Pack 3 or later on domain controllers running Windows 2000 it is possible to disable this default behavior.

Active Directory Domains and Trusts provides a graphical interface in which you can view all domains in the forest. Using this tool, an administrator can manage each of the domains in the forest, trust relationships between domains, configure the functional level for each domain or forest, and configure the alternative user principal name (UPN) suffixes for a forest.

Active Directory Domains and Trusts can be used to accomplish most trust related tasks. It can be used to target all Active Directory domain controllers and can verify all Active Directory trust types. Trust verification takes place between two domains by enumerating all of the domain controllers in each domain. If you choose to have Active Directory Domains and Trusts create both sides of the trust at once, the trust password is automatically generated.

For the various trust related tasks that can be performed using this tool, see the table Trust Tools Comparison by Task earlier in this section. For more information, see Help in Active Directory Domains and Trusts.

**Netdom.exe: Windows Domain Manager**

**Category**
The Windows Domain Manager command-line tool (Netdom) is included when you install Windows Server 2003 Support Tools.

**Version compatibility**
This tool is compatible with computers running Windows XP Professional, Windows Server 2003, Standard Edition; Windows Server 2003, Web Edition, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

Netdom is a command-line tool that allows you to create and manage Active Directory trust relationships (except forest trusts) and can help reduce the number of steps needed to create a trust by using Active Directory Domains and Trusts. You can also use the Netdom command line tool to complete batch management of trusts, join computers to domains, verify trusts (including forest trusts) and secured channels, and obtain information about the status of trusts.

Netdom can be targeted at all Active Directory domain controllers and can verify all Active Directory trust types. Verification is accomplished between two domains by enumerating the domain controllers in each domain. If you choose to have Netdom create both sides of the trust at once the trust password is automatically generated.

For the various trust related tasks that you can perform using this tool, see the table Trust Tools Comparison by Task earlier in this section. To find more information about Netdom, see “Support Tools Help” in the Tools and Settings Collection.

**Nltest.exe: NLTest**

**Category**
The NLTest command-line tool is included when you install Windows Server 2003 Support Tools.

**Version compatibility**
This tool is compatible with computers running Windows XP Professional, Windows Server 2003, Standard Edition; Windows Server 2003, Web Edition, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

You can use the NLTest command-line tool to perform trust-related network administrative tasks such as testing the trust relationship between a Windows–based computer that is a member of a domain and the domain controller on which its computer account is located. In domains where an external trust is defined, NLTest can be used to test the trust relationship between all domain controllers in the trusting domain and a domain controller in the trusted domain. Nltest can also be used to verify any secured channel.

To view the various trust related tasks that can be performed using this tool, see the table Trust Tools Comparison by Tasks earlier in this section.. To find more information about NLTest, see “Support Tools Help” in the Tools and Settings Collection.

**Netdiag.exe: Network Connectivity Tester**

**Category**
The Network Connectivity Tester command-line tool (Netdiag) is included when you install Windows Server 2003 Support Tools.

**Version compatibility**
This tool is compatible with computers running Windows XP Professional, Windows Server 2003, Standard Edition; Windows Server 2003, Web Edition, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

The Netdiag command-line tool examines .dll files, output from other tools, and the system registry to find potential problems. You can use Netdiag to troubleshoot connectivity over the secured channel that exists between a workstation and a domain controller.

For the various trust related tasks that can be performed using this tool, see the table Trust Tools Comparison by Task earlier in this section. To find more information about Netdiag, see “Support Tools Help” in the Tools and Settings Collection.

**Dcdiag.exe: Domain Controller Diagnostic Tool**

**Category**
The Domain Controller Diagnostic Tool command-line tool (Dcdiag) is included when you install Windows Server 2003 Support Tools.

**Version compatibility**
This tool is compatible with computers running Windows XP Professional, Windows Server 2003, Standard Edition; Windows Server 2003, Web Edition, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

You can use the Dcdiag command line tool to verify external trusts. The Dcdiag tool cannot be used to verify Kerberos-based trust relationships; to verify Kerberos-based trust relationships, the recommended method is to use the Netdom tool. Using Dcdiag you can; scope your external trust verification by site or by domain controller, check for trust establishment, check secured channel setup, and check for ticket validity between each pair of domain controllers. By default, errors are flagged. In verbose mode, the successes are printed as well.

For the various trust related tasks that can be performed using this tool, see the table Trust Tools Comparison by Task earlier in this section. To find more information about Dcdiag, see "Support Tools Help" in the Tools and Settings Collection.

**Trust WMI Classes**
Windows Management Instrumentation (WMI) provides access to information about certain objects in a Windows 2000 Server or Windows Server 2003 operating system. WMI providers and classes represent the managed resources on a computer and are used by administrators and developers for scripting and monitoring purposes.

TrustMon is the WMI provider that can access health information about the trusts between domains and the secured channels between domain controllers. Enterprise Administrators can use the TrustMon provider and its associated WMI classes to monitor and troubleshoot trusts remotely and to create scripts that will help automate trust related administrative tasks. Consequently, developers can use the TrustMon provider and classes to create trust related event monitoring applications that will alert an administrator when an important incident occurs and, when appropriate, supply the necessary troubleshooting details.

TrustMon registers three WMI classes that can be used to retrieve information about trust properties for a given domain. These are Microsoft_TrustProvider, Microsoft_DomainTrustStatus and Microsoft_LocalDomainInfo. The following table lists and describes these WMI classes.

### WMI Classes Associated with Trusts

| Class Name | Namespace | Version Compatibility |
|---|---|---|
| Microsoft_TrustProvider | root\microsoftactivedirectory | Windows 2000 Server and Windows Server 2003 |
| Microsoft_DomainTrustStatus | root\microsoftactivedirectory | Windows 2000 Server and Windows Server 2003 |
| Microsoft_LocalDomainInfo | root\microsoftactivedirectory | Windows 2000 Server and Windows Server 2003 |

For more information about these WMI classes, see the WMI SDK documentation on MSDN.

**Network Ports used by Trusts**
Because trusts must be deployed across various network boundaries, they might have to span one or more firewalls. When this is the case, you can either tunnel trust traffic across a firewall or open specific ports in the firewall to allow the traffic to pass through. In Windows Server 2003, this procedure is simplified through configurable remote procedure call (RPC) ports for the main trust services. The two main configurable RPC ports are the Local Security Authority RPC port and the Net Logon RPC port.

**The Local Security Authority (LSA) RPC port.**
This port is used for trust creation and other access to the LSA policy database.

**The Net Logon RPC port.**
This port is used for NTLM authentication and secured channel communications.

The following table shows the list of ports that might need to be opened before you establish trusts.

### Ports Required for Trusts

| Task | Outbound Ports | Inbound Ports | From–To |
|---|---|---|---|

| | | | |
|---|---|---|---|
| Set up trusts on both sides from the internal forest | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution — portmapper (135 TCP)<br>Net Logon fixed port | N/A | Internal domain domain controllers–External domain domain controllers (all ports) |
| Trust validation from the internal forest domain controller to the external forest domain controller (outgoing trust only) | LDAP (389 UDP)<br><br>Microsoft SMB (445 TCP)<br><br>Endpoint resolution — portmapper (135 TCP)<br>Net Logon fixed port | N/A | Internal domain domain controllers–External domain domain controllers (all ports) |
| Use Object picker on the external forest to add objects that are in an internal forest to groups and DACLs | N/A | LDAP (389 UDP and TCP)<br><br>Windows NT Server 4.0 directory service fixed port<br><br>Net Logon fixed port<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution portmapper (135 TCP) | External server–Internal domain PDCs (Kerberos)<br><br>External domain domain controllers–Internal domain domain controllers (Net Logon) |
| Set up trust on the external forest from the external forest | N/A | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP) | External domain domain controllers–Internal domain domain controllers (all ports) |
| Use Kerberos authentication (internal forest client to external forest) | Kerberos (88 UDP) | N/A | Internal client–External domain domain controllers (all ports) |
| Use NTLM authentication (internal forest client to external forest) | N/A | Endpoint resolution – portmapper (135 TCP)<br>Net Logon fixed port | External domain domain controllers–Internal domain domain controllers (all ports) |
| Join a domain from a computer in the internal network to an external domain | LDAP (389 UDP and TCP)<br><br>Microsoft SMB (445 TCP)<br><br>Kerberos (88 UDP)<br><br>Endpoint resolution — portmapper (135 TCP)<br>Net Logon fixed port<br><br>Windows NT Server 4.0 directory service fixed port | N/A | Internal client–External domain domain controllers (all ports) |

To specify the services that you want to run on a fixed port, you must appropriately configure the registry for that port.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

Settings for the Local Security Authority (LSA) RPC port are stored in the **TCP/IP Port** entry in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** registry key.

Settings for the Net Logon RPC port are stored in the **DCTcpipPort** entry in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters** registry key.

**Related Information**

The following resources contain additional information that is relevant to this section.

- "Support Tools Help" in the Tools and Settings Collection
- Kerberos Authentication Technical Reference
- Networking Collection

# Security Considerations for Trusts

Updated: November 19, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**In this section**


- Potential Threats to Interforest Trusts
- Security Settings for Interforest Trusts
- Minimum Administrative Credentials for Securing Trusts
- Trust Security and Other Windows Technologies
- Related Information


The threat scenarios outlined in this section apply only to trusts made between two forests (also known as interforest trusts), including external and forest trusts. All other trusts made within a forest (also known as intraforest trusts), including parent-child, tree-root, and shortcut trusts, are optimally secured by default and do not require further planning to mitigate any known threat. As with intraforest trusts, there are no known threats to realm trusts that require mitigation.

You should be familiar with these threats before you deploy or configure a Windows Server 2003 network environment.

**Potential Threats to Interforest Trusts**

There are two potential threats to interforest trust relationships in Windows Server 2003. These threats can disrupt or undermine the integrity of interforest trusts.

- **Attack on trusting forest by malicious user in a trusted forest.** A malicious user with administrative credentials who is located in a trusted forest could monitor network authentication requests from the trusting forest to obtain the security ID (SID) information of a user who has full access to resources in the trusting forest, such as a Domain or Enterprise Administrator. SID filtering is set on all trusts by default to help prevent malicious users from succeeding with this form of attack. For more information about how SID filtering works, see "Security Settings for Interforest Trusts."

- **Attack on shared resources in a trusting forest by malicious users in another organization's forest.** Creating an external or forest trust between two forests essentially provides a pathway for authentications to travel from the trusted forest to the trusting forest. While this action by itself does not necessarily create a threat to either forest, because it allows all secured communications to occur over the pathway, it creates a larger surface of attack for any malicious user located in a trusted forest. Selective authentication can be set on interforest trusts to help minimize this attack surface area. For more information about how to mitigate this threat, see "Security Settings for Interforest Trusts."

**Security Settings for Interforest Trusts**

There are two security settings in Windows Server 2003 that can be used to enhance the integrity of communications made over interforest trusts. SID filtering helps prevent malicious users with administrative credentials in a trusted forest from taking control of a trusting forest. Selective authentication lessens the attack surface by restricting the quantity of authentication requests that can pass through an interforest trust.

**SID Filtering**

SID filtering is set on all trusts to prevent malicious users who have domain or enterprise administrator level access in a trusted forest from granting (to themselves or other user accounts in their forest) elevated user rights to a trusting forest. It does this by preventing misuse of the attributes containing SIDs on security principals (including inetOrgPerson) in the trusted forest. One common example of an attribute that contains a SID is the SID history attribute (sIDHistory) on a user account object. The SID history attribute is typically used by domain administrators to seamlessly migrate the user and group accounts that are held by a security principal from one domain to another.

When security principals are created in a domain, the domain SID is included in the SID of the principal to identify the domain in which it was created. The domain SID is important because the Windows security subsystem uses it to verify the identity of the security principal, which in turn determines what resources in the domain the principal can access.

**How SID History is used to migrate accounts**
Domain administrators can simplify account migration by using the SID history attribute to migrate permissions, either automatically by using the Active Directory Migration Tool (ADMT) or manually by adding SIDs from an old user or group account to the SID history attribute of the new, migrated account. With either method, the new account retains the same level of permissions or access to resources as the old account. If domain administrators could not use the SID history attribute in this way, they would have to determine and reapply permissions on each network resource to which the old account had access. For more information about the SID history attribute, see "Trust Security and Other Windows Technologies."

**How SID History can be used to elevate privileges**
Although SID history has legitimate and important uses, it can also pose a security threat when used to exploit an unprotected trust. A malicious user with administrative credentials who is located in a trusted forest could monitor network authentication requests from the trusting forest to obtain the SID information of a user, such as a domain or enterprise administrator, who has full access to resources in the trusting forest. After obtaining the SID of an administrator from the trusting forest, a malicious user with administrative credentials can add that SID to the SID history attribute of a security principal in the trusted forest and attempt to gain full access to the trusting forest and the resources within it.

This method of gaining access by granting unauthorized user rights to a user is known as an elevation of privilege attack. In an elevation of privilege attack, an attacker might apply the SID of a domain administrator located in a trusting forest to the SID history attribute of the attacker's own account located in a trusted forest, get a ticket that would automatically include the new SID, and then use the ticket to access resources in the trusting forest. When the attacker requests the use of a resource, the access control mechanism considers all SIDs in the authorization data to determine if the principal has the rights to complete the requested action.

In an external trust scenario, a malicious user who has domain administrator credentials in the trusted domain is a threat to the entire trusting forest. In a forest trust scenario, a malicious user who has domain or enterprise administrator credentials in the forest root domain of the trusted forest is a threat to the entire trusting forest. Although the concept of elevating privileges by modifying SIDs is relatively easy to understand, it is quite difficult to implement. Attackers can use various technologies together with SID history to accomplish an elevation of privilege attack.

**Application Programming Interfaces (APIs)**
Windows includes APIs that facilitate account migration. These APIs are not exposed and can only be accessed on a system that has been patched to allow access to them. In this case, the APIs could be misused to add SIDs for a user from one domain to the SID history of a user in another domain. This is unlikely because these APIs require domain administrative credentials for both domains, including the domain being attacked. In order to overcome that security measure, malicious users would need to get the password of an account with domain administrative credentials before adding the SID. Attackers with access to such an account could more easily use it to accomplish their ultimate goal, rather than having to carry out an elevation of privilege attack to achieve the goal.

**Disk Editors**
A disk editor, such as the DiskProbe utility included in the Windows Server 2003 support tools, could also be used to mount an attack. A malicious user could boot into another operating system and use a disk editor to modify an offline Active Directory database. With a disk editor, the user could modify the SID history attribute, modify replication attributes so the change would be replicated, and calculate a new directory checksum so as to prevent Active Directory from detecting that the directory was improperly modified. This attack is difficult not only because it requires physical access to a domain controller, but all of the tasks are technically complex.

**Debuggers**
Debuggers can also be used maliciously. A user could attach a debugger to a domain controller and use it to modify the copy of the directory loaded in memory. This method is also technically sophisticated. It requires the attacker to have unrestricted physical access, be a domain administrator, and be able to use system APIs to modify system-level code that is not publicly available.

**How interforest trusts use SID Filtering**
Even though a sophisticated attacker might be able to exploit the SID history attribute to gain unauthorized access to resources, SID filtering can provide an effective countermeasure. An outgoing external or forest trust created from a trusting domain can use SID filtering to verify that incoming authentication requests made by security principals in the trusted domain contain only SIDs of security principals in the trusted domain.

The trust compares the SIDs of the requesting security principal to the domain SID of the trusted domain. Any SIDs from domains other than the trusted domain are removed, or filtered. When this happens the request for authentication will fail and the resource will not be accessed.

A stricter form of SID filtering is SID filter quarantining. When a SID filter quarantine is applied to a trusted domain (using the trust relationship between the two domains), only SIDs from the trusted domain are allowed to traverse the trust relationship.

This prevents inbound communications (across the trust relationship) from the trusted domain to claim an identity that belongs to any other domain.

SID filter quarantining was designed to be applied to external trust relationships. It should not be applied to forest trust relationships, trusts within a domain, or trusts within a forest that has a forest functional level of Windows 2000.

External trusts created from domain controllers running Windows 2000 Service Pack 3 (or earlier) do not enforce SID filter quarantining by default. To improve the security of Active Directory forests, domain controllers running Windows Server 2003 and Windows 2000 Service Pack 4 or later enable SID filter quarantining on all new outgoing interforest trusts by default.

**Note**

- The following restrictions and recommendations apply to using SID filtering to mitigate security threats to external and forest trusts:

- You cannot turn off the default behavior that enables SID filter quarantining for newly created external trusts.

- To further secure your forest, you should consider enabling SID filter quarantining on all existing external trusts that were created by domain controllers running Windows 2000 Service Pack 3 (or earlier). You can do this by using Netdom.exe to enable SID filter quarantining on existing external trusts, or by recreating these external trusts from a domain controller running Windows Server 2003 or Windows 2000 Service Pack 4 (or later). Domain controllers running Windows NT Server 4.0 do not take part in the trust creation process when existing domain controllers in the same domain are running Windows 2000 or Windows Server 2003.

**How SID filtering impacts operations**

SID filtering can affect your existing Active Directory infrastructure in the following two ways:

- Users who use SID history data for authorization to resources in the trusting domain no longer have access to those resources.

  Because SID history that contains SIDs from any domain other than the trusted domain is removed from authentication requests made from the trusted domain, the user is denied access to resources that reference the old SID.

- Universal group access control strategy between forests will require changes.

  If you typically assign universal groups from a trusted forest to access control lists (ACLs) on shared resources in the trusting domain, SID filtering will have a significant impact on your access control strategy.

  Because universal groups must adhere to the same SID filtering guidelines as other security principal objects (that is, the universal group object SID must also contain the originating domain SID), you should verify that any universal groups that are assigned to shared resources in the trusting domain were created in the trusted domain.

  If the universal group in the trusted forest was not created in the trusted domain, even though it might contain users from the trusted domain as members, authentication requests made by members of that universal group will be filtered and discarded. Therefore, before assigning access to resources in the trusting domain for users in the trusted domain, you should confirm that the universal group containing the trusted domain users was created in the trusted domain.

  **Note**

  - The default SID filtering on forest trusts does not prevent migrations to domains within the same forest (intraforest migration) from using SID history, and it will not affect your intraforest universal group access control strategy. You should not apply SID filter quarantining on trusts between forests (that is, forest trusts or interforest trusts).

**Disabling SID Filter Quarantining on External Trusts**

Although it reduces the security of your forest (and is therefore not recommended), you can disable SID filter quarantining for an external trust by using the Netdom.exe tool. You should consider disabling SID filter quarantining only in the following situations:

- You have an equally high level of confidence in the administrators who have physical access to domain controllers in the

trusted domain and the administrators with such access in the trusting domain.

- You have a strict requirement to assign universal groups to resources in the trusting domain, even when those groups were not created in the trusted domain.

- Users have been migrated to the trusted domain with their SID histories preserved, and you want to grant them access to resources in the trusting domain based on the SID history attribute.

Only domain administrators or enterprise administrators can modify SID filtering settings. To disable SID filter quarantining for the trusting domain, type a command using the following syntax at a command-prompt:

**Netdom trust** *TrustingDomainName* **/domain:** *TrustedDomainName* **/quarantine:No /usero:** *domainadministratorAcct* **/passwordo:** *domainadminpwd*

To re-enable SID filtering, set the **/quarantine:** command-line option to **Yes**. For more information about Netdom, see "Domain and Forest Trust Tools and Settings."

**Allowing SID History to Traverse Forest Trusts**

If users are migrated from one domain to another in different forests, you may want to allow the migrated users to access resources in their original forest using their migrated (SID history) credentials. The default SID filtering applied to forest trusts prevents user resource access requests from traversing the trusts with the credentials of the original domain. If you want to enable users to use the credentials that were migrated from their original domain, you can allow SID history to traverse forest trusts by using the Netdom command.

Only domain administrators or enterprise administrators can modify SID filtering settings. To allow SID history credentials to traverse a trust relationship between two forests, type a command using the following syntax at a command-prompt:

**Netdom trust** *TrustingDomainName* **/domain:** *TrustedDomainName* **/enablesidhistory:Yes /usero:** *domainadministratorAcct* **/passwordo:** *domainadminpwd*

To re-enable the default SID filtering setting across forest trusts, set the **/enablesidhistory:** command-line option to **No**. For more information about Netdom, see "Domain and Forest Trust Tools and Settings."

**Note**

- The same security considerations for removing SID filter quarantining from external trusts apply to allowing SID history to traverse forest trusts.

**Selective Authentication**

Selective authentication is a security setting that can be set on interforest trusts. It provides Active Directory administrators who manage a trusting forest more control over which groups of users in a trusted forest can access shared resources in a trusting forest. This increased control is especially important when administrators need to grant access to shared resources in their organization's forest to a limited set of users located in another organization's forest, because creating an external or forest trust provides a pathway for all authentication requests to travel between forests.

While this action by itself does not necessarily cause a threat to either forest, because all secured communications occur over the pathway, an external or forest trust exposes a larger surface to attack by any malicious user located in a trusted forest. Selective authentication helps to minimize this exposed area by enabling Active Directory administrators to grant a new authentication permission — to computer objects in the resource domain — for specific user accounts located in another organization's forest.

This new authentication permission is set on the security descriptor of the computer object located in Active Directory, not on the security descriptor physically located on the resource computer in the trusting forest. Controlling authentication in this way provides an extra layer of protection to shared resources by preventing them from being randomly accessed by any authenticated user working in a different organization, unless the user has been granted this permission explicitly by someone with write access to the computer object in Active Directory.

**Note**

- To enable selective authentication on forest trusts, the trusting forest in which shared resources are located must have the forest functional level set to Windows Server 2003. To enable selective authentication on external trusts, the trusting domain in which shared resources are located must have the domain functional level set to Windows 2000 native.

**Authentication Settings for Interforest Trusts**

When the forest functional level is set to Windows Server 2003, Active Directory Domains and Trusts recognizes three authentication settings that can be set on interforest trusts: Domain-wide Authentication, Forest-wide Authentication, and Selective Authentication. The following table describes these authentication settings in more detail.

Authentication Settings Used on Interforest Trusts

| Authentication Setting | Interforest Trust Type | Description |
|---|---|---|
| Domain-wide Authentication | External | Permits unrestricted access by any users in the trusted domain to all available shared resources located in the trusting domain. This is the default authentication setting for external trusts. |
| Forest-wide Authentication | Forest | Permits unrestricted access by any users in the trusted forest to all available shared resources located in any of the domains in the trusting forest. This is the default authentication setting for forest trusts. |
| Selective Authentication | External and Forest | Restricts access over an external or forest trust to only those users in a trusted domain or forest who have been explicitly given authentication permissions to computer objects (resource computers) residing in the trusting domain or forest. This authentication setting must be manually enabled. |

**Domain- and forest-wide authentication settings**

Domain and forest-wide authentication provide an unrestricted pathway in which all authentication requests made by users in a trusted forest will be successfully authenticated by a domain controller in the trusting forest where the shared resource is located. The following illustration shows how the domain controller in the trusting forest authenticates and routes authentication requests made by users in the trusted forest.

**Authentication Requests Are Authenticated and Routed**




When domain or forest-wide authentication is enabled, users who are authenticated over an interforest trust are automatically provided the Authenticated Users SID of the trusting forest in their authorization data. The Authenticated Users SID is used to grant many of the default rights for users in a forest.

Because the Authenticated Users group is a computed group, and its SID is added on the server to which the user authenticates, you cannot change the membership of the group. Because of this, after you set up a interforest trust, users from the other forest receive some default rights to all of the resources in the trusting forest that are accessible by the Authenticated Users group. Consequently, you might not want to use the domain-wide or forest-wide authentication setting if the trust is set up only to allow access to a small subset of users who are in the trusted forest.

Once an authentication request made to a resource in a trusting forest is validated by the trusted forest, it is routed to the targeted resource computer, which determines, based on its access control configuration, whether to authorize the specific request made by the user, service, or computer in the trusted forest. In this way, interforest trusts provide the mechanism by which validated authentication requests are passed to a trusting forest, while access control mechanisms on the resource computer determine the final level of access granted to the requestor in the trusted forest.

Once the authentication request reaches the resource computer, that computer must determine whether the user who initiated the authentication request is authorized to access the shared resource it is hosting. It does this by comparing the SIDs provided in the access token of the authentication request with the SIDs in the security descriptor that has been granted access to the

shared resource.

**Selective authentication setting**

Unlike domain and forest-wide authentication, selective authentication provides a more restrictive pathway in which only authentication requests made from users in a trusted forest, who have been granted access to the Active Directory objects hosting the resources in a trusting forest, can be authenticated by the domain controller in the resource domain. The following illustration shows how the domain controller in the trusting forest can restrict the flow of authentication requests made by users in the trusted forest to shared resources in the trusting forest when those users have not been granted the appropriate permissions.

**Authentication Requests Are Not Authenticated or Routed**




**How selective authentication affects domain controller behavior**

When selective authentication is enabled, all authentication requests made over a trust to the trusting forest are tagged with an identifier that subjects the request to closer scrutiny by the domain controllers in the domain where the target resource is located. This is important because the mere presence of this identifier triggers the domain controller in the resource domain to first check whether the user requesting the access has been given explicit permissions to the Active Directory object where the resource is hosted. The appropriate permission to the resource object in Active Directory is verified before the domain controller sends a service ticket back to the requesting workstation in the trusted forest. The following illustration summarizes how domain controllers in the trusting forest are affected when selective authentication is enabled.

**Effect of Selective Authentication on Domain Controllers in the Trusting Forest**




The steps by which domain controllers process authentication requests when selective authentication is enabled are as follows:

1. **All domain controllers in the domain where the trust is established detect authentication requests made from the other organization.**

When selective authentication is enabled all domain controllers in the forest root domain of a trusting forest (with a forest trust) or all domain controllers in a resource domain in the trusting forest (with an external trust) are aware that all authentication requests coming in over that trust are coming from a different organization. This means that when a user from a different organization authenticates across a trust with the selective authentication option enabled, the domain controller that receives the request in the trusting domain recognizes this user as other than an authenticated user located in the trusting forest.

2. **Receiving DC in the domain where trust is established tags an identifier to the authorization data of the trusted user.**

   Once a domain controller receives the request it adds an identifier to the authorization data of the trusted user. This identifier is known as the Other Organization SID. The Other Organization SID is used to identify users making requests across a trust with selective authentication enabled.

3. **All domain controllers in the resource domain detect requests originating from the other organization, and the receiving domain controller in the resource domain performs an access check.**

   When a domain controller in a resource domain detects an Other Organization SID in the authorization for the trusted user, the domain controller is alerted to first check whether the requesting user has been granted explicit permissions to the Active Directory object where the resource is hosted. The appropriate permission to the resource object in Active Directory must be verified before the domain controller can authenticate the access request and send a service ticket back to the requesting workstation in the trusted forest.

   This verification process works by checking the access permissions set on the Allowed to Authenticate permission on the discretionary access control list (DACL) on the computer object in Active Directory. When you set selective authentication on an outgoing interforest trust, you need to manually assign access to the Allowed to Authenticate permission on each computer object, which physically represents a member server to which you want users in the trusted forest to authenticate.

   You must grant access to this permission in order for users located in a trusted forest to authenticate to the shared resources hosted by the server computer. If this permission has not been granted to a user whose authorization data contains the Other Organization SID, the service ticket to the server hosting the shared resource is not granted and the authentication process fails.

   **Note**

   - The following restrictions apply to the Allowed to Authenticate authorization permission:
   - By default, only members of the Account Operators, Administrators, Domain Administrators, Enterprise Administrators, and SYSTEM security groups located in the trusting domain can modify the Allowed to Authenticate authorization permission.
   - When Kerberos is used for authentication, the Allowed to Authenticate permission should be enabled on computer accounts when the service that hosts the resource is running as Local System or Network Service. If the service that hosts the resource is running as a Domain Service account, the permission should be configured on that account. For example, if the SQL service is running under the account SQLServer, the SQL Server account has to have the Allowed to Authenticate permission enabled for users to be able to authenticate to that account across trusts that have selective authentication enabled.
   - When NTLM is used for authentication, the Allowed to Authenticate permission should be granted to the computer account, even if the service that you want to connect to is using a domain user account.

When a Windows NT 4.0, Windows 2000 Server, or Windows Server 2003 member server receives an authentication request from a user, regardless of whether that user is identified with the Other Organization SID, the member server automatically adds the Authenticated Users SID for the trusting forest to the authorization data for that user.

**How selective authentication affects Windows Server 2003 member server authentication**

When users who are located in a domain within the trusting forest (or in a trusted domain that does not have selective authentication enabled) authenticate to local Windows Server 2003 member servers, the Authenticated Users SID and a new identifier indicating that they belong to the same organization are placed in their authorization data. This new identifier is known as the This Organization SID. Only the Other Organization SID or the This Organization SID can be present in the authorization data of an authenticated user along with the Authenticated Users SID. The following table describes the SIDs that can be added to the authorization data of a user throughout the entire selective authentication process and shows when they are added.

SIDs Added to Authorization Data During the Selective Authentication Process

| SID | Purpose | When Added to User Authorization Data |
|---|---|---|
| Authenticated Users | Adds default access rights to resources in the trusting forest. Because this SID is mandatory for all incoming authentication requests, it prevents all anonymous access to resources in the trusting forest. | Added when any incoming authentication request to a member server is made from a user who has been identified with the Other Organization SID. This SID is always applied by the member server to the authorization data of an incoming user, in addition to one of the other SIDs in this table. |
| This Organization | Identifies users who access resources locally within a trusting forest or across a trust that is not marked for selective authentication. | Added when users who are located in any of the domains within a trusting forest authenticate to a local Windows Server 2003 member server. This SID and the Authenticated Users SID are added by the member server. The Other Organization SID cannot be added to this authorization data of this user. |
| Other Organization | Identifies users coming across a trust that is marked for selective authentication. | Added when any incoming authentication request is made from a user located in a trusted forest to a domain controller in the domain in the trusting forest where the trust is established. This SID is added by the domain controller that initially receives the authentication request. The Authenticated Users SID is not added to the authorization data of the user until the member server has authenticated the user. |

**Note**

- The Allowed to Authenticate permission can be set on computer objects that represent member servers running Windows NT Server 4.0, Windows 2000 Server, and Windows Server 2003. However, only member servers running Windows Server 2003 can provide the This Organization SID to the authorization data of a user.

**Processing authentication requests made over forest trusts with selective authentication enabled**

Before an authentication request made with the Kerberos version 5 authentication protocol can follow the forest trust path, the service principal name (SPN) of the resource computer must be resolved to a location in the other forest. A SPN is a multicomponent name that is used to identify a service that is associated with a computer account. When a workstation in one forest attempts to access data on a resource computer in another forest, the Kerberos authentication process contacts the domain controller for a service ticket to the SPN of the resource computer. Once the domain controller queries the global catalog and determines that the SPN is not in the same forest as the domain controller, the domain controller sends a referral back to the workstation for its parent domain. The workstation queries its parent domain for the service ticket and continues to follow the referral chain until it reaches the domain where the resource is located.

Once the workstation reaches a domain controller in the root domain of the resource forest where selective authentication is enabled, the domain controller adds the Other Organization SID to the authorization data for the user. After the Other Organization SID is added, the domain controller in the resource domain checks the Allowed to Authenticate permission on the computer object (located in Active Directory) that is hosting the resource. If the user has been granted access the domain controller issues a service ticket back to the workstation. This is an important part of the selective authentication process, because if the user does not have permission to authenticate to the computer object, a service ticket will not be granted and access to the target resource computer fails. In this way, the selective authentication setting restricts what authentication requests can pass through a trust and obtain a ticket for a target resource computer.

The following figure and corresponding steps provide a detailed description of the Kerberos authentication process that is used when computers running Windows 2000 Professional, Windows XP Professional, Windows 2000 Server, or Windows Server 2003 attempt to access resources from a computer located in another forest when selective authentication is enabled.

**Kerberos authentication process over a forest trust with selective authentication enabled**




This illustration incorporates images from Active Directory Users and Computers in the Microsoft Management Console to help you better understand at what point during the authentication process a domain controller queries the Allowed to Authenticate permission for the resource object.

1. Acctuser1 logs on to Workstation1 using credentials from the Sales.tailspintoys.com domain. The user then attempts to access a shared resource on FileServer1 located in the WingtipToys forest. Acctuser1 is also a member of the Accounting global group in the Sales.tailspintoys.com domain.

2. Workstation1 contacts the Key Distribution Center (KDC) on a domain controller in its domain (TailspinDC1) and requests a service ticket for the FileServer1 SPN. FileServer1 is located in the Mktg.wingtiptoys.com domain and is a member server.

3. TailspinDC1 does not find the SPN in its domain database and queries the global catalog to see if any domains in the TailspinToys forest contain this SPN. Because a global catalog is limited to its own forest, the SPN is not found. The global catalog then checks its database for information about any forest trusts that are established with its forest, and, if any are found, it compares the name suffixes listed in the forest trust trusted domain object (TDO) to the suffix of the target SPN to find a match. Once a match is found, the global catalog provides a routing hint back to TailspinDC1.

4. Using the information in the routing hint, TailspinDC1 sends a referral for a domain controller located in the forest root domain of the WingtipToys forest back to Workstation1.

5. Workstation1 contacts a domain controller in the forest root domain (WingtipDC1) of the Corp.wingtiptoys.com forest and requests a service ticket to the Fileserver1 resource computer.

6. WingtipDC1 contacts its global catalog to find the SPN, and the global catalog finds a match for the SPN and sends it back to WingtipDC1.

7. Because the authentication request originated from the TailspinToys forest and selective authentication is enabled WingtipDC1 applies the Other Organization SID to the user's authorization data and then sends a referral for the mktg.wingtiptoys.com domain back to Workstation1.

8. Workstation1 contacts the KDC on WingtipDC2 and attempts to negotiate a ticket for Acctuser1 to gain access to FileServer1.

9. WingtipDC2 detects the Other Organization SID in the authorization data of Acctuser1, which requires the domain controller to first locate the computer object of the resource computer (Fileserver1) before providing a ticket back to the requesting computer.

10. Once the computer object is located, the domain controller must check whether the user requesting access to the resource computer has been explicitly granted the **Allowed to Authenticate** permission on the Fileserver1 computer object located in the Mktg.wingtiptoys.com domain. This process must complete successfully before WingtipDC1 can provide a ticket back to the requesting computer.

    **Note**

    - In this example, Acctuser1 is a member of the Accounting group in the TailspinToys forest and that group has been explicitly granted the **Allowed to Authenticate** permission to this computer object residing in Active Directory.

11. After WingtipDC1 determines that Acctuser1 has the required permission to Filerserver1, it sends a service ticket back to Workstation1 permitting it to gain access to FileServer1.

12. Now that Workstation1 has a service ticket, it sends the server service ticket to FileServer1, which reads the security credentials for Acctuser1 and constructs an access token accordingly.

**Impact of Selective Authentication**

Because all verification of incoming interforest authentication requests is done locally on the receiving domain controller in the trusting forest, access to resources in the trusting forest is likely to be extremely limited for a broad set of users on the network (which is the purpose of this security setting). Consequently, implementing selective authentication might require user education, particularly due to the following reasons:

- Users browsing network resources through My Network Places to resources located in a trusting forest might get access denied messages when attempting to access those resources.

- Resources in the trusting forest that were once available to users in a trusted forest might no longer be available.

**Note**

- As a security best practice it is recommended that resource administrators in a trusting forest remove the default access rights granted to the Authenticated Users group in all of the shared resources in the trusting forest. This practice will help to further minimize the possibility of authenticated users inside and outside of a forest from accessing protected resources.

**Enabling and Disabling Selective Authentication**

Selective authentication must be manually enabled or disabled by using Active Directory Domains and Trusts or the Netdom.exe tool. To enable selective authentication for the trusting domain, type a command using the following syntax at a command-prompt:

**Netdom trust** *TrustingDomainName* **/domain:** *TrustedDomainName* **/SelectiveAuth:Yes /usero:** *domainadministratorAcct* **/passwordo:** *domainadminpwd*

To disable selective authentication, set the **/SelectiveAuth:** command-line option to **No**. For more information about the Netdom tool, see "Domain and Forest Trust Tools and Settings."

You can enable or disable selective authentication only from the trusting side of a trust. If the trust is a two-way trust, you can also enable or disable selective authentication in the trusted domain by using the credentials of the domain administrator for the trusted domain and reversing the values of TrustingDomainName and TrustedDomainName in the command.

### Note

- You can further increase security in a forest by enabling selective authentication on all external and forest trusts used to connect forests that are not managed by the IT department in your organization.

#### Minimum Administrative Credentials for Securing Trusts

The following table lists the minimum administrative credentials necessary to enable or disable the SID filtering and selective authentication security settings in Active Directory.

Minimum Administrative Credentials Necessary To Secure Trusts

| SID Filter Quarantining - External Trusts | Allowing SID History - Forest Trusts | Selective Authentication - External Trusts | Selective Authentication - Forest Trusts |
|---|---|---|---|
| You must be a member of the Domain Admins group in the trusting domain or the Enterprise Admins group in the trusting forest. | You must be a member of the Domain Admins group in the forest root domain or the Enterprise Admins group in the trusting forest | You must be a member of the Domain Admins group in the trusting domain or the Enterprise Admins group in the trusting forest. | You must be a member of the Domain Admins group in the forest root domain or the Enterprise Admins group in the trusting forest |

#### Trust Security and Other Windows Technologies

Trust security can be affected or enhanced by other technologies built into the Windows operating system. These include SID history and the Lightweight Directory Access Protocol (LDAP).

#### SID History

SID History is used to store the former SIDs of moved objects such as users and security groups. Prior to Windows 2000, restructuring a domain generally meant the creation of new accounts that needed to be placed in the same groups the old accounts were placed in. Since Windows 2000, domain restructuring is considerably easier because of SIDHistory, which is an attribute of Active Directory security principals.

As part of the move operation for the user or group, Application Programming Interfaces (APIs), or tools and support utilities built from these APIs, update the SIDHistory attribute of the object, representing it in Active Directory with its former SID. When the user next logs on to the system, not only the new SID, but also the old SID retrieved from the SIDHistory attribute is added to the user's access token and serves to determine the user's group memberships. The SIDs of the groups of which the user is a member (through either the new SID or the old SID) are also added to the access token, together with any SIDHistory those groups might have.

Groups have a SIDHistory attribute because they also can be moved. The system retrieves the SIDHistory attributes of all the groups of which the user is a member and adds these attributes to the user's access token. Because these SIDHistory entries in the token look to the operating system like normal group memberships during authorization checks, the token can grant appropriate access even on earlier versions of the operating system.

Because old SIDs can be stored in SIDHistory, the user can continue to access resources after a move without requiring administrators to modify ACLs. Once a migration is complete, if all needed resources are in the new domain, use of SIDHistory can be phased out if desired.

### Note

- The SIDHistory attribute can be updated only in native mode Windows 2000 or Windows Server 2003 domains, which requires all migration operations relying on SIDHistory to have a native mode target.

For more information about the security threat that exploits SID history, see "Security Settings for Interforest Trusts."

#### LDAP Sign and Encrypt

When using Windows Server 2003, secure LDAP traffic is enabled so that Active Directory administrative tools sign and encrypt all LDAP traffic by default. In this way, secured LDAP traffic enhances the security of communications that can occur

between trusts.

Active Directory Domains and Trusts signs and encrypts all LDAP traffic by default. Signing LDAP traffic guarantees that the packaged data comes from a known source and that it has not been tampered with. Active Directory administrative tools in Windows 2000 do not sign and encrypt all LDAP traffic by default. In order to maintain a secure network, it is strongly recommended that you upgrade all domain controllers running Windows 2000 to Service Pack 3 or later.

You can use the corresponding Active Directory administrative tools in Windows 2000 to manage Windows 2000 domain controllers that do not have the Windows 2000 Server Service Pack 3 or later installed. However, traffic is not signed and encrypted by LDAP on domain controllers running Windows 2000.

**Related Information**
The following resources contain additional information that is relevant to this section.

- Domains and Forests Technical Reference
- Kerberos Authentication Technical Reference
- Security Identifiers Technical Reference
- Security Principals Technical Reference

# Tools and Settings Collection

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tools and Settings Collection**
The following resources provide technical information about tools and settings related to Windows Server 2003.

**Command-Line References**
Click one of the following links to view online versions of the searchable Help files of technical information about the command-line tools included with the Windows Server 2003 operating system: Command line reference

- Command-Line Reference for Windows Server 2003, Standard Edition
- Command-Line Reference for Windows Server 2003, Enterprise Edition
- Command-Line Reference for Windows Server 2003, Datacenter Edition

**Event and Error Messages**
Click the following link to go to an online database of events and errors for Windows Server 2003:

- Windows Server 2003 Events and Errors

**Group Policy Settings Reference**
Click the following link to go to the download site for a spreadsheet that lists the full set of Group Policy settings described in Administrative Template (.adm) files shipped with Windows Server 2003:

- Group Policy Settings Reference for Windows Server 2003

**Registry Reference**
Click the following link to go to a searchable online reference providing detailed descriptions of the Windows Server 2003 registry, including many entries that cannot be edited by using Windows Server 2003 tools or programming interfaces:

- Registry Reference for Windows Server 2003

**Resource Kit Tools**
You can install the Resource Kit tools on a local computer. Click the following link to go to the download site for the Windows Server 2003 Resource Kit tools, including a searchable Help file of technical information about this collection of tools (included with the Windows Server 2003 Deployment Kit) that can help you deploy, configure, maintain, and troubleshoot Windows Server 2003:

- Windows Server 2003 Resource Kit Tools Download

To view this technical reference information locally, download the Windows Resource Kit Tools for Windows Server 2003 and install them on a computer running either Windows XP Professional or Windows Server 2003.

To view Resource Kit Tools Help in Windows XP:

- In Help and Support Center, under **Pick a Task**, click **Use tools to view your computer information and diagnose problems**, and then click **Windows Resource Kit Tools**.

To view Resource Kit Tools Help in Windows Server 2003:

- Help and Support Center, under **Support Tasks**, click **Tools**, and then click **Windows Resource Kit Tools**.

**Resource Kit Tool Updates**
Click the following link to go to the download site for updated Windows Server 2003 Resource Kit tools. System administrators, developers, and power users can use these tools to streamline a wide range of administrative tasks.

- Windows Server 2003 Resource Kit Tool Updates

**Support Tools Help**
A searchable Help file of technical information about the Support Tools included with Windows Server 2003. Support Tools are intended to assist Microsoft support personnel and network administrators in diagnosing and resolving computer problems.

Click the following link to go to the online version of Support Tools Help:

- Windows Support Tools

To view Support Tools Help from a local computer, install the Support Tools from the \Support\Tools folder on the Windows Server 2003 installation disc on a computer running either Windows XP Professional or Windows Server 2003.

To view Support Tools Help in Windows XP:

- In Help and Support Center, under **Pick a Task**, click **Use tools to view your computer information and diagnose problems**, and then click **Support Tools**.

To view Resource Kit Tools Help in Windows Server 2003:

- In Help and Support Center, under **Support Tasks**, click **Tools**, and then click **Windows Support Tools**.

**Windows Firewall Settings**
Click the following link to go to a searchable online reference that tells you how to configure Windows Firewall so it works with Windows Server 2003 remote administration tools, optional components, server roles, and services. Most remote administration tools, optional components, server roles, and services will not function properly when Windows Firewall is turned on unless you configure Windows Firewall settings.

- Windows Firewall Settings

# Windows Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introduction to Windows Support Tools**
The Windows Support Tools assist support personnel and network administrators in managing their networks and troubleshooting problems. They are not installed with the Windows operating system; you must install them separately from the \Support\Tools folder of the Windows operating system CD. This Help file provides information on the tools and shortcuts for opening or running these tools.

To install Windows Support Tools, see Install Windows Support Tools on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=62270).

**Getting Help on Support Tools**

### To find Help for a tool

- Click the Alphabetical List of Tools topic on the **Contents** tab to display a list of tools by the file names of the tools.
- Click a category on the **Contents** tab and then click the file name of the tool.
- Use the **Index** tab to locate a tool by either tool name or file name.

Each tool is covered in a main Help topic. Links to associated topics covering syntax, examples, or other features of the tool are available at the top of each topic for that tool. Tools with Windows interfaces may include a separate Help file available from the Help menu in the tool window. For command-line tools, Help is also available by typing *FileName* **/?** at the command prompt.

An extensive Support Tools Glossary is available from the **Contents** tab. Links to glossary definitions that appear as pop-up windows are formatted in underlined dark green text.

**Troubleshooting with Support Tools**
Many Support Tools provide diagnostic features that are useful for troubleshooting Windows operating system configurations. However, because these tools have been designed primarily for system administrators and support personnel, their output may be difficult to analyze. In these cases, you can consult the Microsoft Windows Deployment and Resource Kits Web site for troubleshooting information.

**Understanding Notation and Terminology**
The following topics cover the conventions for usage and notation that are observed in this document:

Procedural Conventions

Notational Conventions

This documentation assumes you are already familiar with the Windows operating system. For more general information about Windows, including keyboard equivalents to menu and mouse actions, see Windows Help.

**Printing Topics**
When you print from HTML Help, a dialog box opens asking whether you want to print the selected topic or to print the selected heading and all subtopics. Printing the selected topic is recommended. If you print a heading and subtopics, you may encounter error messages and special formatting will be lost, but the topics still print.

Printing from Help and Support Center (HSC) results in similar behavior, except the dialog box only opens when a heading or topic has subtopics.

# Notational Conventions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Notational Conventions**

| Convention | Meaning |
|---|---|
| **bold** | In syntax, characters that you type exactly as shown, including commands and parameters. In text, menu names and menu commands are also bold. |
| **`bold monospace`** | Commands that you must type exactly as shown to get the results being discussed. |
| *italic* | Variables for which you supply a specific value. For example, *Filename.ext* represents any valid file name. |
| Initial Capitals (Filename.ext) | Names of files should begin with an initial capital letter, for example, Filename.ext. Paths and folders can be uppercase, lowercase, or mixed, according to how they actually appear in a standard installation of the application or the operating system. |
| ALL CAPITALS | Used for acronyms. |
| `monospace` | Examples of code. |
| [ ] (square brackets) | In syntax descriptions, square brackets enclose optional items. If you include the item, type only the information between them, not the square brackets themselves. |
| {*choice1* \| *choice2*} (braces) | In syntax descriptions, braces enclose items which require a choice, such as {**yes** \| **no**}. Type only one of the choices, not the braces or the dividing line. |

# Procedural Conventions

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Procedural Conventions**

| Convention | Meaning |
|---|---|
| type | An instruction to **type** information means to press the key or keys and then press the ENTER key. |
| select | An instruction to **select** information means to highlight folders, file names, text boxes, menu bars, and options, or to select options in a dialog box. |
| + | A plus sign ( + ) between two or more key names indicates that you must press the keys at the same time; for example, ALT + TAB. |
| , | A comma ( , ) between two or more key names indicates that you must press each key consecutively; for example, ALT, F, X. |
| Note | Alerts you to supplementary information. |
| Important | Alerts you to essential information. |
| Caution | Alerts you to possible data loss, breaches of security, or other more serious problems. |
| Operating system reference conventions | The following conventions in this documentation are used to mean the indicated groups of Windows products.<br>1. References to "Windows Server 2003" include the following products:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>2. References to "Windows Server 2003 family" include the following products:<br>• Windows Server 2003, Standard Edition<br>• Windows Server 2003, Enterprise Edition<br>• Windows Server 2003, Datacenter Edition<br>• Windows Server 2003, Web Edition |

# Related Information on the Internet

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Related Information on the Internet**
There are many Microsoft Web sites that provide information and updates regarding Microsoft products and operating systems, including Windows Server 2003, Windows XP, Windows 2000, Windows NT, and Windows 98.

If you have an Internet connection and a Web browser, you can click the following links to visit these sites.

**Windows Deployment and Resource Kits Web Site**
Microsoft Windows Deployment and Resource Kits Web site

**Windows Web Sites**
Microsoft Windows Server 2003 Web site

Microsoft Windows XP Web site

Microsoft Windows Web site

Microsoft Windows 2000 Server Web site

Microsoft Windows NT Server Web site

Microsoft Windows NT Workstation Web site

Microsoft Product Support Services Web site

Microsoft Windows Hardware Compatibility List Web site

**Other Microsoft Web Sites of Interest**
Microsoft Knowledge Base Search Web site

Microsoft Internet Explorer Web site

Microsoft TechNet Web site

MSDN Web site

**Microsoft Home Page**
For any other information about Microsoft products, search the Microsoft Web site.

# New and Removed Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New and Removed Support Tools**
Windows Server 2003 Support Tools include new tools that are not available in Windows 2000 Server Support Tools, and tools have been removed from Windows Server 2003 Support Tools that were available in Windows 2000 Server Support Tools.

**Installation of Windows Server 2003 SP1 Support Tools**
You can install the version of Windows Support Tools that ships with Windows Server 2003 Service Pack 1 (SP1) by using Suptools.msi only on computers running Windows Server 2003 with SP1. You can install individual tool executable files on computers running Windows Server 2003 with no service pack installed and on computers running Windows XP Professional by copying the executable files to the computers. However, you cannot use Suptools.msi to install the entire tool set on any computer that is not running Windows Server 2003 with SP1.

**Tools Added to Support Tools**
The tools in the following table have been added to the Windows Server 2003 Support Tools.

| Tool name | Descriptive name | Primary category |
|---|---|---|
| Addiag.exe | Application Deployment Diagnostics | Software and System Deployment |
| Apmstat Overview | Advanced Power Management Status | Configuration and Management |
| Bindiff Overview | Binary File Difference Finder | File and Folder Management |
| Bitsadmin Overview | BITS Administration Utility | File and Folder Management |
| Cabarc Overview | Cabinet Tool | File and Folder Management |
| Connstat Overview | FRS Connection Status Report | Disk and Data Management |
| Dhcploc Overview | DHCP Server Locator Utility | Network Services Management |
| Diruse Overview | Directory Disk Usage | File and Folder Management |
| Dmdiag Overview | Disk Manager Diagnostics | File and Folder Management |
| Efsinfo Overview | Encrypting File System Information | File and Folder Management |
| Exctrlst Overview | Extensible Performance Counter List | Performance Monitoring |
| Filever Overview | File Version | File and Folder Management |
| Ftonline Overview | Fault Tolerant Disk Mounter | File and Folder Management |
| Getsid Overview | Get Security ID | Network Services Management |
| Health_chk Overview | FRS Connection Status Report | Disk and Data Management |
| Httpcfg Overview | HTTP Configuration Utility | Internet Services Management Tools |
| Iasparse Overview | IAS Parse Tool | Network Services Management |
| Iologsum Overview | FRS Health Check | Disk and Data Management |
| Netcap Overview | Network Monitor Capture Utility | Network Services Management |
| Portqry Overview | Port Query | Network Services Management |
| Setspn Overview | Manipulate Service Principal Names for Accounts | Active Directory Management |

| | | |
|---|---|---|
| Spcheck Overview | Service Pack Checker | Configuration and Management |
| Topchk.cmd | DFS and SYSVOL Replication Topology Analysis | Disk and Data Management |

**Debugging Tools**
The following debugging tools have been moved from Windows Support Tools to Microsoft Debugging Tools:

- Apcompat.exe
- Dumpchk.exe
- Kill.exe

To obtain debugging tools and documentation, and for information about debuggers, see Microsoft Debugging Tools.

**Tools Removed from Support Tools**
The tools in the following table have been removed from Windows Server 2003 Support Tools:

| Tool name | Descriptive name | Disposition |
|---|---|---|
| Apcompat.exe | Application Compatibility Toolkit | Microsoft Debugging Tools. |
| Dommig.exe | Domain Migration Tool | See Active Directory Migration Tool (ADMT) on the Windows Server 2003 operating system CD. |
| Dumpchk.exe | Dump Check Utility | Obsolete. See Microsoft Debugging Tools. |
| Kill.exe | Process Kill Utility | Replaced by TaskKill.exe, which is in %windir%\system32. |
| Msinfo32.exe | Msinfo32 | Moved to Program Files\Common Files\Microsoft Shared\MSInfo. |
| Pmon.exe | Process Resource Monitor | Use Pviewer Overview instead. Pmon.exe is included in the Windows Server 2003Resource Kit Tools. |
| Pptpclnt.exe and Pptpsrv.exe | Point-to-Point Tunneling Protocol Ping Utilities (PPTP Ping) | Removed. |
| Reg.exe | Registry Console Tool | Moved to %windir%\system32. |
| Tlist.exe | Task List Viewer | Replaced by Tasklist.exe, which is in %windir%\system32. |
| Wsremote.exe | Wsremote | Use Remote Overview instead. |

# Alphabetical List of Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

Tools are listed by their file name, followed by their full name. If the full name begins with a different letter than the file name, the tool is also listed by its full name, followed by its file name.

**A**

Acldiag Overview

Replmon Overview

Search Overview

Addiag.exe

ADSI Edit (adsiedit.msc) Overview

Apmstat Overview

**B**

Bindiff Overview

Bitsadmin Overview

Browstat Overview

**C**

Cabarc Overview

Clonepr Overview

Connstat Overview

**D**

Dcdiag Overview

Depends Overview

Devcon Overview

Topchk.cmd

Dfsutil Overview

Dhcploc Overview

Diruse Overview

Dskprobe Overview

Dmdiag Overview

Dnscmd Overview

Dnslint Overview

Dsacls Overview

Dsastat Overview

Dskprobe Overview

**E**

Efsinfo Overview

Exctrlst Overview

**F**

Filever Overview

Ftonline Overview

Connstat Overview

Health_chk Overview

Iologsum Overview

**G**

Getsid Overview

Gflags.exe

**H**

Health_chk Overview

Httpcfg Overview

**I**

Iologsum Overview

Iasparse Overview

**J**

No entries

**K**

Ksetup Overview

Ktpass Overview

**L**

Ldp Overview

**M**

Setspn Overview

Poolmon Overview

Memsnap Overview

Movetree.exe

Msicuu Overview

Msizap Overview

**N**

Netcap Overview

Netdiag Overview

Netdom Overview

Nltest Overview

Ntfrsutl Overview

**O**

No entries

**P**

Poolmon Overview

Portqry Overview

Pviewer Overview

**Q**
No entries

**R**
Remote Overview

Repadmin Overview

Replmon Overview

Rsdiag Overview

Rsdir Overview

**S**
Search Overview

Sdcheck Overview

Setspn Overview

Sidwalker Security Administration Tools

Sidwalker Security Administration Tools

Spcheck Overview

**T**
Topchk.cmd

**U**
No entries

**V**
No entries

**W**
Windiff Overview

Netdom Overview

Msicuu Overview

Msizap Overview

**X**
Xcacls Overview

**Y**
No entries

**Z**
No entries

# Active Directory Management Support Tools

Updated: July 13, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Active Directory Management Tools for Windows Server 2003**
This topic provides links to Active Directory Management Tools for Windows Server 2003. These tools are available in Windows Support Tools that is included on the Windows Server 2003 operating system disk on Microsoft Download Center.

For more information about tools for managing Active Directory, see How to Administer Microsoft Windows Client and Server Computers Locally and Remotely (http://go.microsoft.com/fwlink/?LinkID=177813).

For more information about tools for managing Active Directory Domain Services in Windows Server 2008 and Windows Server 2008 R2, see Active Directory Domain Services Command Reference (http://go.microsoft.com/fwlink/?LinkId=197167).

The following topics provide more information about the tools available in Windows Support Tools for managing Active directory in Windows Server 2003:

- Acldiag Overview
- ADSI Edit (adsiedit.msc)
- Clonepr Overview
- Dsacls Overview
- Dsastat Overview
- Ldp Overview
- Movetree.exe
- Ntfrsutl Overview
- Repadmin Overview
- Replmon Overview
- Sdcheck Overview
- Search Overview
- Setspn Overview

# Acldiag Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The AclDiag command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Acldiag.exe: ACL Diagnostics**
This command-line tool detects and reports discrepancies in the access control lists (ACLs) of objects in Active Directory. It can also reapply a security delegation template to an ACL, eliminating special permissions and restoring incomplete delegations.

With AclDiag, you can:

- Display the access control entries (ACEs) in the ACL, and inheritance and audit settings.
- Display the effective permissions of users and groups to an Active Directory object.
- Compare the ACL for an object in Active Directory to the default permissions defined in the schema.
- Compare the ACL of an Active Directory object to a delegation template.
- Reapply the delegation template to the ACL of an Active Directory object.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
For more information about Active Directory, see the Active Directory Overview.

**System Requirements**
The following are the system requirements for Acldiag:

- Windows 2000, Windows XP Professional, or Windows Server 2003.
- The user must have read permissions on Active Directory objects. To reapply a delegation template, the user must have modify permissions to the Active Directory object.

**File Required**

- Acldiag.exe

**See Also**
**Concepts**
Acldiag Remarks
Acldiag Syntax
Acldiag Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview

ADSI Edit (adsiedit.msc)
Acldiag Overview

# Acldiag Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Acldiag Remarks**



AclDiag displays only the permissions of objects the user has the right to view. Because Group Policy objects are virtual objects that have no distinguished name, this tool cannot be used on them.

For general purpose ACL reporting and setting from the command prompt, you can also use DsAcls, another Windows Support Tool.

**See Also**

**Concepts**

Acldiag Overview
Acldiag Syntax
Acldiag Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)

# Acldiag Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Acldiag Syntax**

Acldiag Overview | Acldiag Remarks | Acldiag Syntax | Acldiag Examples

**acldiag** *ObjectDN* [**/schema**] [**/chkdeleg**] [**/geteffective:**{*User* | *Group* | *}] [**/fixdeleg**] [**/skip**] [**/tdo**]

**Parameters**

**Note**

- If you specify an object without additional parameters, AclDiag lists the access control entries (ACEs) in the ACL, and inheritance and audit settings.

***LDAP-URL***

Identifies the Active Directory object to investigate. Enter the LDAP URL for an object in Active Directory. The LDAP URL format consists of the name of the LDAP server followed by the distinguished name of the object. The string must be enclosed in quotation marks. For example:

**"LDAP://domain.test.microsoft.com/CN=Test Admin,CN=Users,DC=domain,DC=test,DC=microsoft,DC=com"**

**/geteffective:{*User* | *Group* | *}**

Adds an effective rights diagnosis to the display. The effective rights diagnosis displays the effective permissions to the object held by specified users or groups. Effective permissions are the permissions that are enforced after precedence is applied and conflicts in rights are resolved.

| Value | Description |
|---|---|
| *User* \| *Group* | Displays the effective permissions held by the specified user or group. |
| * | Displays the effective permissions of all users and groups in the access control list (ACL) for the object. |

**/schema**

Adds a schema diagnosis to the display. The schema diagnosis reports whether the object ACL includes the ACEs that are in the schema defaults.

**/chkdeleg**

Adds a delegation diagnosis to the display. The delegation diagnosis reports whether the object ACL includes the ACEs that are in the delegation template. A status of **misconfigured** indicates that at least one, but not all, ACEs in a delegation template (and in the schema default) are included in the ACL.

**/fixdeleg**

Directs AclDiag to reapply the delegation template to the object ACL, eliminating special permissions and restoring incomplete delegations. When the specified object inherits delegated permissions, this parameter reapplies the delegation template to the object for which the delegated permissions are explicitly defined. The **/fixdeleg** parameter is interactive, it gives the user an opportunity to fix each misconfigured delegation.

**Note**

- This parameter is effective only when used with the **/chkdeleg** parameter. Without **/chkdeleg**, **/fixdeleg** is ignored, but AclDiag does not report an error.

**/skip**

Omits the security description from the display. The security description is a list of the ACEs in the object ACL.

**/tdo**

Displays output in tab-delimited format. Fixed-width format is the default. Tab-delimited format is useful when the output is destined for a database or spreadsheet.

**See Also**

**Concepts**

Acldiag Overview
Acldiag Remarks
Acldiag Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)

# Acldiag Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Acldiag Examples**

Acldiag Overview | Acldiag Remarks | Acldiag Syntax | Acldiag Examples

**Example 1: Display the ACL of a user object in Active Directory**

To display the ACL of a user object in Active Directory, type:

```
acldiag CN="Test Admin",CN=Users,DC=domain1,DC=test,DC=fourthcoffee,DC=com
```

**Example 2: Display a Schema Analysis of a Computer Object in Active Directory**

To display a schema analysis of a computer object in Active Directory, type:

```
acldiag CN=MACHINE-TEST,CN=Computers,DC=domain1,DC=test,DC=fourthcoffee,DC=com /schema
```

**Example 3: Display ACL, Permissions, and Delegation Information for a Computer Object**

To display the ACL, the effective permissions for all users and groups in the security descriptor, and the delegation analysis of a computer object in tab-delimited format, type:

```
acldiag CN=MACHINE-TEST,CN=Computers,DC=domain1,DC=test,DC=fourthcoffeeDC=com /chkdeleg /geteffective:* /tdo
```

**Example 4: Reapply a Delegation Template to a Group Object**

To reapply a delegation template to a group object, type:

```
acldiag CN="Domain Computers",CN=Users,DC=domain1,DC=test,DC=fourthcoffee,DC=com /chkdeleg /fixdeleg
```

**See Also**

**Concepts**

Acldiag Overview
Acldiag Remarks
Acldiag Syntax
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)

# ADSI Edit (adsiedit.msc)

Updated: April 25, 2012

Applies To: Windows SBS 2008, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

Active Directory® Service Interfaces Editor (ADSI Edit) is a Lightweight Directory Access Protocol (LDAP) editor that you can use to manage objects and attributes in Active Directory. ADSI Edit (adsiedit.msc) provides a view of every object and attribute in an Active Directory forest. You can use ADSI Edit to query, view, and edit attributes that are not exposed through other Active Directory Microsoft Management Console (MMC) snap-ins: Active Directory Users and Computers, Active Directory Sites and Services, Active Directory Domains and Trusts, and Active Directory Schema.

This topic includes the following sections:


- Installing ADSI Edit
- Using ADSI Edit
- Adding ADSI Edit to MMC
- Missing Commands
- Other Topics with ADSI Edit Usage Scenarios


**Note**

Another LDAP editor that Microsoft provides is Ldp. To learn more about Ldp, see Ldp Overview (http://go.microsoft.com/fwlink/?LinkId=143517). For an example of Ldp being used, see article 224543 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=23064).

**Installing ADSI Edit**

To install ADSI Edit on computers running Windows Server® 2003 or Windows® XP operating systems, install Windows Server 2003 Support Tools from the Windows Server 2003 product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

On servers running Windows Server 2008 or Windows Server 2008 R2, ADSI Edit is installed when you install the Active Directory Domain Services (AD DS) role to make a server a domain controller. You can also install Windows Server 2008 Remote Server Administration Tools (RSAT) on domain member servers or stand-alone servers. For specific instructions, see Installing or Removing the Remote Server Administration Tools Pack (http://go.microsoft.com/fwlink/?LinkId=143345).

To install ADSI Edit on computers running Windows Vista® with Service Pack 1 (SP1) or Windows 7, you must install RSAT. For more information and to download RSAT, see article 941314 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=116179).

**Note**

- Adsiedit.msc will not run unless the Adsiedit.dll file is registered. This happens automatically if the support tools are installed. However, if the support tool files are copied instead of installed, you must run the **regsvr32** command to register Adsiedit.dll before you run the Adsiedit.msc snap-in. To register adsiedit.dll, type the following command (you must navigate to the directory containing the adsiedit.dll file):

  ```
  regsvr32 adsiedit.dll
  ```

- You can run ADSI Edit from a client computer or server. The computer does not have to be a member of a domain. However, to see domain objects using Adsiedit.msc, you must have the rights to view the Active Directory domain that you connect to. By default, members of the Domain Users group have these rights. To modify objects using ADSIEdit, you must have at least the Edit permission on the Active Directory objects that you want to change. By default, members of the Domain Admins group have this permission.

**Using ADSI Edit**

ADSI Edit (Adsiedit.msc) is an MMC snap-in. You can add the snap-in to any .msc file through the **Add/Remove Snap-in** menu option in MMC, or just open the Adsiedit.msc file from Windows Explorer. The following figure illustrates the ADSI Edit interface. In the console tree on the left, you can see the major partitions Domain, Configuration, and Schema. The figure shows the Builtin container of the Contoso.com domain selected. In the details pane on the right, you can see the Builtin groups of Active Directory.

**Note**

Adsiedit.msc automatically attempts to load the current domain to which the user is logged on. If the computer is installed in a workgroup or otherwise not logged on to a domain, the message "The specified domain does not exist" displays repeatedly. To resolve this issue, you may want to open an MMC, add the ADSI Edit snap-in, make connections as appropriate, and then save the console file.




**ADSI Edit Node**

To view the following commands, in the console tree click the **ADSI Edit** node, click the **Action** menu, and then click one of the following:

- Connect To
- Refresh

**Connect To**

The **Connection Settings** dialog box appears. You can use the **Connection Settings** dialog box to create a connection point to an object in Active Directory. The following text boxes are located in the **Connection Settings** dialog box:

**Name**. You should not change the text in this box because it might cause an error when you attempt to make a connection. The text in this box is updated automatically, if necessary, when you configure or select a Connection Point.

**Path**. Displays the URL for the selected object. It cannot be edited. If the path is not correct, click **Cancel**, and then select the correct object.

**Connection Point Section**

Click either **Select or type a Distinguished Name or Naming Context** or **Select a well known Naming Context**.

- If you click the **Select or type a Distinguished Name or Naming Context** radio button, type the distinguished name of the object that will be the connection point in Active Directory. For example, if your domain name is contoso.com and you want to connect to the **Users** container, type **cn=users,dc=contoso,dc=com**.
- If you click the **Select a well known Naming Context** radio button, select the directory partition that will be the connection point in Active Directory in the list of partitions in the selection menu.

**Tip**

Previous LDAP connections are remembered by the ADSI Edit tool. In versions earlier than Windows Server® 2008, the tool automatically attempts to load the current domain to which the user is logged on. If the computer is installed in a workgroup or otherwise not logged on to a domain, the message "The specified domain does not exist" appears repeatedly. To avoid these issues, open Mmc.exe, add the ADSI Edit snap-in manually, make any connections that are appropriate for you with whatever credentials are necessary, and then save the console file. This gives you your own default console that works with ADSI Edit.

**Computer Section**
Specifies whether you connect to the local computer or a remote computer. Click either **Default (domain or server that you are logged in to**) or **Select or type a domain or server**. You can enter the domain name or computer name in Domain Name System (DNS) format or NetBIOS format, or you can enter an IP address.

**Advanced Button**
Click the **Advanced** button to specify alternate credentials or alternate port numbers or to change the protocol that is used to connect to Active Directory. The **Advanced** dialog box contains the following text boxes:

**Specify Credentials**. Use this box to specify alternate credentials. Unless otherwise specified, the currently logged-on user's credentials are used.

**Port Number**. Type a port number if you do not want to use the default port for the LDAP or the LDAP Global Catalog protocol. The default LDAP port is 389. The default port for the Global Catalog is 3268.

**Protocol Section**
Click either **LDAP** or **Global Catalog**. The URL in **Path** might change, based on your selection. LDAP is used by default. To view the **Path** box, on the **Advanced** tab, click **OK** or **Cancel**.

**Refresh**
To update the object from Active Directory, right-click an object, and then click **Refresh**. The **Refresh** command removes the current objects in the container and repopulates the container with updated information from Active Directory.

**Directory Partition Node**
To view the following commands, select the directory partition node that you want to manage (that is, **Domain**, **Configuration**, and **Schema**), click the **Action** menu, and then click one of the following:

- Settings
- Remove
- Update Schema Now
- New
- Refresh

**Settings**
The **Connection Settings** dialog box appears, which provides the same configuration options as previously discussed in the Connect To section.

**Remove**
Removes the connection point that connects ADSI Edit to a directory partition or container within Active Directory. This command affects only what is shown in the ADSI Edit console. To remove objects from Active Directory, use the **Delete** command.

**Update Schema Now**
Reloads the schema information from Active Directory into the local computer's cache.

**New**
Click **Query** to create a new query. For more information about creating LDAP queries, see LDAP Query Basics (http://go.microsoft.com/fwlink/?LinkId=143553).

**Refresh**
To update the object from Active Directory, right-click an object, and then click **Refresh**. The **Refresh** command removes the current objects in the container and repopulates the container with updated information from Active Directory.

**Object Node**
To view the following commands, click an object in the details pane (for example, **Account Operators** is an object in the previous figure), click the **Action** menu, and then click one of the following:

- Move
- New Connection From Here
- New
- Delete
- Rename
- Refresh

**Move**

Moves the object to another container in Active Directory. Opens a dialog box that you can use to select the destination container.

**New Connection From Here**
Creates a new connection point node and adds it to the console.

**New**
The **New** menu command reveals another menu that contains the **Object** command, which creates a new child object in the selected container. This command opens a set of chained dialog boxes that begins with the class of the object. If you do not have the appropriate permissions to create an object in the selected container, no classes will be listed. After you select a class, a dialog box opens for each required attribute. In the final dialog box, click **More** to view and edit any optional attributes.

**Delete**
Deletes the selected object from Active Directory. A dialog box appears asking you to confirm the deletion. This command does not appear in the menu if you do not have permissions to delete an object from Active Directory.

**Rename**
Changes the name of the object in Active Directory.

**Refresh**
To update the object from Active Directory, right-click an object, and then click **Refresh**. The **Refresh** command removes the current objects in the container and repopulates the container with updated information from Active Directory.

**Adding ADSI Edit to MMC**

If you are running ADSI Edit on a computer that is not logged on to a domain or if you want to create a customized MMC, you may want to add the ADSI Edit snap-in to the console.

**To add the ADSI Edit Snap-in to MMC**

1. Open your existing console or create a new console. To create a new console, click **Start**, click **Run**, type **mmc**, and click **OK**, or at a command line, type **mmc**, and then press ENTER.
2. Click **Add/Remove Snap-in**, and then click **Add**.
3. In the **Add Standalone Snap-in** dialog box, click **ADSI Edit** in the list. If ADSI Edit does not appear here, see Installing ADSI Edit at the beginning of this topic.
4. Click **Add**, click **Close**, and then click **OK**.

**Missing Commands**
The **Action** menus in MMC are context sensitive. If you do not have permission to perform an action, the action might not appear in the menu.

**Other Topics with ADSI Edit Usage Scenarios**
Although ADSI Edit is not intended for regular management of your Active Directory environment, there are instances in which you may need to use it. The following topics include procedures that use ADSI Edit.

- Weaken security using ADSI Edit (http://go.microsoft.com/fwlink/?LinkId=143425)
- Manage an AD LDS Instance Using ADSI Edit (http://go.microsoft.com/fwlink/?LinkId=143426)
- Using ADSI Edit to Edit Active Directory Attributes (http://go.microsoft.com/fwlink/?LinkId=143427)
- Article 822444 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=143428)
- Step 1: Create a PSO (http://go.microsoft.com/fwlink/?LinkId=143555)

**See Also**
**Concepts**
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
Acldiag Overview
**Other Resources**
ADSIEdit videos on Microsoft Showcase

# ADSI Edit UI: Graphical Representation

Published: June 11, 2010

Updated: April 26, 2012

Applies To: Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

## ADSI Edit

Active Directory® Service Interfaces Editor (ADSI Edit) is a Lightweight Directory Access Protocol (LDAP) editor that you can use to manage objects and attributes in Active Directory. ADSI Edit (adsiedit.msc), a Microsoft Management Console (MMC) snap-in, provides a view of every object and attribute in an Active Directory forest. You can use ADSI Edit to query, view, and edit attributes that are not exposed through the other Active Directory MMC snap-ins: Active Directory Users and Computers, Active Directory Sites and Services, Active Directory Domains and Trusts, and Active Directory Schema.

### Installing ADSI EDIT

**Windows Server 2003/Windows XP** - Install Windows Server 2003 Support Tools from the Windows Server 2003 product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114).
**Windows Server 2008/Window Server 2008 R2** - ADSI Edit is installed when you install the Active Directory Domain Services (AD DS) role to make a server a domain controller. You can also install Windows Server 2008 Remote Server Administration Tools (RSAT) on domain member servers or stand-alone servers.
**Windows Vista/Windows 7** - Install RSAT. For more information and to download RSAT, see article 941314 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=116179).

**Using ADSI Edit**
Add the ADSI Edit (adsiedit.msc) snap-in to any .msc file through the Add/Remove Snap-in menu option in MMC, or open the Adsiedit.msc file in Windows Explorer.

**Missing Commands**
The Action menus in MMC are context sensitive. If you do not have permission to perform an action, the action might not appear in the menu.



Adsiedit.msc automatically attempts to load the current domain to which the user is logged on.

Adding ADSI Edit to MMC
If you are running ADSI Edit on a computer that is not logged on to a domain or if you want to create a customized MMC, you may want to add the ADSI Edit snap-in to the console.

### ADSI Edit UI

#### Three Main Nodes

**ADSI Edit**
To view the commands in the ADSI Edit node, click the Action menu, and then click Connect To.




**Name** – The display name for the connection point.

**Path** – Displays the URL for the selected object.

**Connection Point** – Select either Distinguished Name or Naming Context.

**Computer** – Specifies whether you connect to the local computer or a remote computer.

**Advanced** – Click to specify alternate credentials or alternate port numbers, or to change the protocol used to connect to Active Directory.

**Directory Partition**
To view the following commands, select the directory partition node that you want to manage (that is, Domain, Configuration, and Schema), and then click the Action menu.




**Settings** – Displays the same information as the UI for the ADSI Edit node.

**Remove** – Displays the URL for the selected object.

**Update Schema Now** – Select either Distinguished Name or Naming Context.

**New** – Specifies whether you connect to the local computer or a remote computer.

**Refresh** – To update the object from Active Directory, right-click an object, and then click Refresh.

**Object**
To view the following commands, click an object in the details pane.




**Move** – Moves the object to another container in Active Directory.

**New Connection from Here** – Creates a new connection point node and adds it to the console

**New** – Reveals another menu that contains the Object command, which creates a new child object in the selected container.

**Delete** – Deletes the selected object from Active Directory.

**Rename** – Changes the name of the object in Active Directory.

**Refresh** – To update the object from Active Directory, right-click an object, and then click Refresh.

**See Also**
**Concepts**

ADSI Edit (adsiedit.msc)

**Other Resources**

ADSI Edit videos on Microsoft Showcase

# Clonepr Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The ClonePrincipal tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Clonepr: ClonePrincipal**
This tool assists administrators in migrating users from Windows NT to Windows 2000 or Windows Server 2003 by creating clones of the Windows NT 4.0 users and groups in the new Windows 2000 or Windows Server 2003 environment.

**Concepts**
Customers deploying Windows 2000 or Windows Server 2003 might want to migrate users incrementally to a new Windows 2000 or Windows Server 2003 environment without impacting their existing Windows NT 4.0 production environment. This allows users to be introduced gradually to Windows 2000 or Windows Server 2003, and allows them an emergency fallback to their old account during the trial period. It enables administrators to consolidate multiple groups from different source domains into a single destination group and old resource domains into a single organizational unit. It eliminates the need to update group memberships and access control lists (ACLs), preserving uninterrupted network access for users in the new domain structure.

ClonePrincipal must be run on a destination domain controller running Windows 2000 or Windows Server 2003.

ClonePrincipal performs tasks related to those performed by another Windows 2000 and Windows Server 2003 Support Tool, Movetree.exe, with the following differences:

- ClonePrincipal is exclusively for interforest use. It copies security Principals between forests only. MoveTree is exclusively for intraforest use. It moves security Principals within the same forest only. This means that for ClonePrincipal, the source domain can be Windows NT 4.0 or Windows 2000 (mixed or native mode), while MoveTree relies on a Windows 2000 source domain.

- ClonePrincipal duplicates objects from the source domain in the target domain rather than moving them to it, so the source domain is unaffected. MoveTree moves objects, which means that the operations it performs destroy objects in the source domain. Both, however, add the old security identifier (SID) to the SIDHistory.

- ClonePrincipal does not maintain users' existing passwords, and MoveTree does.

To move users or groups within a Windows 2000 or Windows Server 2003 domain (for example, from one organizational unit to another), use Active Directory Users and Computers, a Microsoft Management Console snap-in that is part of Windows 2000 and Windows Server 2003.

**Files Required**

- Clonepr.dll. Component Object Model (COM) object with methods to support ClonePrincipal operations.

- Clonegg.vbs. Sample script that clones all the Global Groups in a domain.

- Cloneggu.vbs. Sample script that clones all the Global Groups and Users in a domain.

- Clonelg.vbs. Sample script that clones all the Local Groups in a domain.

- Clonepr.vbs. Sample script that clones a single security Principal.

- Sidhist.vbs. Sample script that adds the SID of a source account to the SIDHistory of a destination account.

**For More Information**
For more information about ClonePrincipal, read Clonepr.doc in the Program Files\Support Tools directory.

### Note

- You need Microsoft Word or a compatible viewer to view Clonepr.doc. You can download a free version of Microsoft Word Viewer from the Microsoft Office Download Center.

For more information about Active Directory, visit the Active Directory page on the Microsoft Windows 2000 Server Web site.

**See Also**
**Concepts**
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Dsacls Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Dsacls command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Dsacls.exe: DsAcls**
This command-line tool displays and changes permissions (access control entries) in the access control list (ACL) of objects in Active Directory.

Note

- The ACEs that you add by using DsAcls must be object-specific permissions that override the default permissions defined in the Active Directory schema for that object type. Do not add ACEs unless you are well-informed about security for Active Directory objects.

**Corresponding UI**
DsAcls is the command-line equivalent of the **Security** tab in the **Properties** dialog box for an Active Directory object in Active Directory tools, such as Active Directory Users and Computers. You can use either tool to view and change permissions to an Active Directory object.

**Concepts**
For more information about Active Directory, see Deploying Active Directory.

**System Requirements**
The following are the system requirements for DsAcls:

- DsAcls runs on Windows 2000, Windows XP Professional, and Windows Server 2003.
- To view an ACL, the user must have read permissions on Active Directory objects. To change an ACL, the user must have write permissions to the Active Directory object.

**File Required**

- Dsacls.exe

**See Also**
**Concepts**
Dsacls Syntax
Dsacls Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Dsacls Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DsAcls Examples**

**Examples of valid permissions**

```
SDRCWDWO;;user
```

Delete, read security information, change security information, and change ownership permissions on objects of type "user."

```
CCDC;group;
```

Create child and delete child permissions to create or delete objects of type "group."

```
RPWP;telephonenumber;
```

Read property and write property permissions on telephone number property.

**See Also**

**Concepts**

Dsacls Overview
Dsacls Syntax
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Dsacls Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DsAcls Syntax**

DsAcls uses the following general syntax:

**dsacls** "[*****Computer*****]*ObjectDN*" [**/A**] [**/D** *PermissionStatement* [*PermissionStatement*]...] [**/G** *PermissionStatement* [*PermissionStatement*]...] [**/I:**{**T** | **S** | **P**}] [**/N**] [**/P:**{**Y** | **N**}] [**/R** {*User* | *Group*} [{*User* | *Group*}]...] [**/S** [**/T**]] [**/?**]

**Parameters**

**Note**

- If you specify an object without additional parameters, DsAcls displays the access control entries (ACEs) in the ACL.

**"[*Computer*\]*ObjectDN*"**

Identifies the Active Directory object to investigate. Type the distinguished name of the object. To specify an object on a remote computer, type the computer name followed by the distinguished name. This parameter must be enclosed in quotation marks. For example:

```
"CN=Jeff Akers,CN=Users,DC=domain,DC=test,DC=microsoft,DC=com"
```

or

```
"\\Server01\CN=Jeff Akers,CN=Users,DC=domain,DC=test,DC=microsoft,DC=com"
```

**/A**

Adds ownership and auditing information to the display.

**/D PermissionStatement[*PermissionStatement*]**

Denies the specified permissions to the user or group.

You can deny permissions to multiple users in each **/D** command, for example:

```
/D Domain1\User1:CCDC Domain1\User2:DC;computer
```

**/G PermissionStatement[*PermissionStatement*]**

Grants specified permissions to the user or group.

You can grant permissions to multiple users in each **/G** command, for example:

```
/G Domain1\User1:CCDC Domain1\User2:DC;computer
```

**/I:{T | S | P}**

Specifies the objects to which the permissions are applied. This parameter determines whether the permissions are inheritable. **T** is the default.

| Value | Description |
|---|---|
| **T** | This object and subobjects. |
| **S** | Subobjects only. |
| **P** | Propagate inheritable permissions one level only. |

**/N**

Provides that the specified ACE replaces the ACEs in the ACL. By default, the ACE is added to the ACL.

**/P:{Y | N}**

Determines whether the object can inherit permissions from its parent objects. If you omit this parameter, the inheritance properties of the object are not changed.

| Value | Description |
|---|---|

| | |
|---|---|
| **Y** | The object is protected and cannot inherit permissions. |
| **N** | The object is not protected and can inherit permissions. |

**>Note**

- This parameter changes a property of the object, not of an ACE. To determine whether an ACE is inheritable, use the **/I** parameter.

**/R {*User* | *Group*} [{*User* | *Group*}]...**
Deletes all ACEs for the specified users or groups. *User* can be specified as *User***@***Domain* or as *Domain******User*. *Group* can be specified as *Group***@***Domain* or as *Domain******Group*.

You can delete ACEs for multiple users and groups in a single **/R** parameter, for example:

```
/R Domain1\User1 Domain1\User2
```

**/S**
Restores the security on the object to the default for that object class as defined in Active Directory schema.

**/T**
Restores the security on the tree of objects to the default for each object class. This parameter is valid only with the **/S** parameter.

**/?**
Displays help for DsAcls.

**Syntax for PermissionStatement**
*PermissionStatement* values use the following format:

{*User* | *Group*}**:***Permissions*[**;**{*ObjectType* | *Property*}][**;***InheritedObjectType*]

**Parameters**
**{*User* | *Group*}**
Specifies the user or group to whom the rights apply. *User* can be specified as *User***@***Domain* or *Domain******User*. *Group* can be specified as *Group***@***Domain* or *Domain******Group*.

***Permissions***
Type one or more of the following values (without spaces).

**Generic Permissions**

| Value | Description |
|---|---|
| **GR** | Generic Read |
| **GE** | Generic Execute |
| **GW** | Generic Write |
| **GA** | Generic All |

**Specific Permissions**

| Value | Description |
|---|---|
| **SD** | Delete. |
| **DT** | Delete an object and all of its children. |
| **RC** | Read security information. |
| **WD** | Change security information. |

| | |
|---|---|
| **WO** | Change owner information. |
| **LC** | List the children of an object. |
| **CC** | Create child object.<br><br>If {*ObjectType* \| *Property*} is not specified to define a specific child-object type, this applies to all types of child objects; otherwise, it applies to the specified child-object type. |
| **DC** | Delete a child object.<br><br>If {*ObjectType* \| *Property*} is not specified to define a specific child-object type, this applies to all types of child objects; otherwise, it applies to the specified child-object type. |
| **WS** | Write to self object.<br><br>This is meaningful only on Group objects and when {*ObjectType* \| *Property*} is a "member." |
| **RP** | Read property.<br><br>If {*ObjectType* \| *Property*} is not specified to define a specific property, this applies to all properties of the object; otherwise, it applies to the specified property of the object. |
| **WP** | Write property.<br><br>If {*ObjectType* \| *Property*} is not specified to define a specific property, this applies to all properties of the object; otherwise, it applies to the specified property of the object. |
| **CA** | Control access right.<br><br>If {*ObjectType* \| *Property*} is not specified to define the specific extended right for control access, this applies to all meaningful control accesses on the object; otherwise, it applies to the specific extended right for that object. |
| **LO** | List the object access.<br><br>Can be used to grant list access to a specific object if List Children (**LC**) is not granted to the parent as well. Can also be denied on specific objects to hide those objects if the user/group has **LC** permission on the parent.<br><br>Note<br><br>• Active Directory does not enforce this permission by default. The Active Directory must be configured to check for this permission. |

***{ObjectType | Property}***

Limits the permission to the specified object type or property. Enter the display name of the object type or of the property. If an object type or property is not specified, the permission applies to all object types and properties.

For example, the following command permits the user to create all types of child objects:

```
/G Domain\User:CC
```

However, the following command permits the user to create only child computer objects:

```
/G Domain\User:CC;computer
```

***InheritedObjectType***

Limits inheritance of the permission to the specified type of object. Enter the display name of the object type. If an object type is not specified, the permission can be inherited by all object types. This parameter is used only when permissions are inheritable.

For example, the following command permits all types of objects to inherit the permission:

```
/G Domain\User:CC
```

However, the following command permits only user objects to inherit the permission:

```
/G Domain\User:CC;;user
```

**See Also**

**Concepts**

Dsacls Overview
Dsacls Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Dsastat Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Dsastat command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Dsastat.exe: Directory Services Utility**
This command-line diagnostic tool compares and detects differences between naming contexts on domain controllers.

DsaStat can be used to compare two directory trees across replicas within the same domain or, in the case of a global catalog, across different domains. The tool retrieves capacity statistics such as megabytes per server, objects per server, and megabytes per object class, and performs comparisons of attributes of replicated objects.

### Note

- You must run DsaStat from the command window.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
The user specifies the targeted domain controllers and additional operational parameters from the command line. DsaStat determines whether domain controllers in a domain have a consistent and accurate image of their own domain. In the case of Global Catalogs, DsaStat checks to see if the Global Catalog has a consistent image with domain controllers in other domains. As a complement to the replication-monitoring tools, Repadmin Overview and Replmon Overview, DsaStat can be used to ensure that domain controllers are up to date with one another.

**System Requirements**
The following are the system requirements for DsaStat:

- Windows Server 2003 or
- Windows XP Professional or
- Windows 2000

**File Required**

- Dsastat.exe

**See Also**
**Concepts**
Dsastat Syntax
Dsastat Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Dsastat Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DsaStat Syntax**

DsaStat uses the following syntax:

**dsastat** [**/loglevel:**{**INFO**|**TRACE**|**DEBUG**}] [**/output:**{**SCREEN**|**FILE**|**BOTH**}] [**/s:***ServerName*[*PortNumber*] [**;***ServerName*[*PortNumber*]...]] [**/t:**{**TRUE**|**FALSE**}] [**/sort:**{**TRUE**|**FALSE**}] [**/p:***PageSize*] [**/scope:**{**BASE**|**ONELEVEL**|**SUBTREE**}] [**/b:***SearchPath*] [**/filter:***LdapFilter*] [**/gcattrs:**{[*LDAPAttributes*[*Attribute;Attribute* [**Auto**]|[**All**]} [**/u:***UserName*] [**/pwd:***Password*] [**/d:***Domain*]

**Parameters**

**/loglevel:{INFO|TRACE|DEBUG}**
Specifies the extent of logging that is performed during execution. The default option is INFO.

**/output:{SCREEN|FILE|BOTH}**
Specifies where the results of DsaStat are displayed. If **FILE** is selected, the output is saved to a file named Dsastat.log.*nnnnn*. The default option is **BOTH**.

**/s: *ServerName*[*PortNumber*][;*ServerName*[*PortNumber*]...]**
Specifies the names of the servers on which the comparison is performed. Separate each server name with a semicolon (**;**). The server name can be appended with the TCP/IP port number. If *PortNumber* is 3268, the comparison is performed against the global catalog of that server. The default port number is the default LDAP port (389).

**/t:{TRUE|FALSE}**
Specifies whether to perform a statistics comparison or a full-content comparison. Setting **/t:** to **TRUE** performs a statistical comparison. Setting **/t:** to **FALSE** performs a full-content comparison. The default is **TRUE**. The statistical comparison merely counts the objects. It does not compare the attributes of the objects that have been retrieved from the domain controllers and it ignores the **/gcattrs** parameter. A full-content comparison retrieves every object with attribute values and performs a comparison of the attributes of the same object from different domain controllers.

**/sort:{TRUE|FALSE}**
Determines whether the search operations are performed with sorting based on object GUID. Setting **/sort:** to **TRUE** performs sorted queries. Setting **/sort:** to **FALSE** does not perform sorted queries. The default value is **FALSE**. Performing a sorted query by setting this option to **TRUE** has a negative impact on performance; however, it complements the full-content comparison because it ensures that objects are returned in nearly the same order from different servers and improves the performance of full-content comparison.

**/p: *PageSize***
Sets the number of entries to be returned per page for ldap_search operation. The valid range for *PageSize* is 1-999. The default value is 64. If the directory objects are large (such as user objects with many of their attributes populated with values), it is advisable to use a smaller page size.

**/scope:{BASE|ONELEVEL|SUBTREE}**
Specifies the extent of scope for the search operation. The default option is **SUBTREE**.

**/b: *SearchPath***
Sets the distinguished name of the base search path. This allows DsaStat to perform the comparison against any subtree of the directory.

**/filter: *LdapFilter***
Sets LDAP filter used in the LDAP search operation. The filter must conform to LDAP search syntax as specified in RFC 2254. The default filter is (objectclass=*).

**/gcattrs:{[*LDAPAttributes*[*Attribute;Attribute...*]]| [ObjectClass]|[Auto]|[All]}**
Specifies attributes to be returned from search. This option is used only if the comparison parameter **/t** is set to **FALSE**: *LDAPAttributes* specifies any LDAP attribute available. Multiple attributes may be specified separated by semicolons. **ObjectClass** specifies that no attributes be displayed. **Auto** specifies that only the attributes replicated to the global catalog be displayed. **All** specifies that all attributes contained in an object be displayed

**/u: *UserName***
Specifies the user name to use for the query.

**/pwd: *Password***
Specifies the password to use for authenticating the user name. Must be used with the **/u** parameter.

**/d: *Domain***
Specifies the domain to use for authenticating the user name. Must be used with the **/u** parameter.

**See Also**

**Concepts**

Dsastat Overview
Dsastat Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Dsastat Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DsaStat Examples**

**Example 1: Object Count Comparison**

Your network has the following domains and servers:

- **Domain**, with two domain controllers:
  - **DomainS1** (which is also a global catalog)
  - **DomainS2**

  The distinguished name of the domain is **DC=Domain,DC=com**.

- **Child.Domain**, a child domain of Domain, with one domain controller:
  - **ChildS1** (which is also a global catalog)

  The distinguished name of the domain is **DC=Child,DC=Domain,DC=com**.

To compare the number of objects in the directory for the domain **Domain**, type the following at the command line:

```
dsastat -s:DomainS1;DomainS2 -b:DC=Domain,DC=com -gcattrs:objectclass -
p:999
```

This search merely keeps a count of the different types of objects in each of the replicas and then compares the count for each object class.

**Example 2: Full Comparison**

Your network has the following domains and servers:

- **Domain**, with two domain controllers:
  - **DomainS1** (which is also a global catalog)
  - **DomainS2**

  The distinguished name of the domain is **DC=Domain,DC=com**.

- **Child.Domain**, a child domain of Domain, with one domain controller:
  - **ChildS1** (which is also a global catalog)

  The distinguished name of the domain is **DC=Child,DC=Domain,DC=com**.

To perform a full comparison of each object in the domain **Domain**, type the following at the command line:

```
dsastat -s:DomainS1;DomainS2 -b:DC=Domain,DC=com -gcattrs:all -sort:true -
t:false -p:16
```

This search does a full content search and comparison of all the objects in the domain **Domain**. Because the default filter is **"**

**(objectclass=*)**," all objects are retrieved by this search. The objects are also sorted based on objectGUID to improve the performance of the comparison operation.

**Example 3: User Comparison**
Your network has the following domains and servers:

- **Domain**, with two domain controllers:
    - **DomainS1** (which is also a global catalog)
    - **DomainS2**

  The distinguished name of the domain is **DC=Domain,DC=com**.

- **Child.Domain**, a child domain of Domain, with one domain controller:
    - **ChildS1** (which is also a global catalog)

  The distinguished name of the domain is **DC=Child,DC=Domain,DC=com**.

To perform a comparison of all users in the Sales organizational unit in the **Domain** domain, type the following at the command line:

```
dsastat -s:DomainS1;DomainS2 -b:OU=Sales,DC=Domain,DC=com -gcattrs:all -
sort:true -t:false -p:16 -filter:"(&(objectclass=user)(!objectClass=computer))"
```

This example illustrates the specification of the base search path at a subtree of the domain, in this case the organizational unit named **Sales**. In addition, the filter specifies that the comparison is only concerned with user objects and not computer objects. Note that since computer objects are derived from user objects in the class hierarchy, a search filter specifying **objectclass=user** will return both user and computer objects.

**Example 4: DC Comparison**
Your network has the following domains and servers:

- **Domain**, with two domain controllers:
    - **DomainS1** (which is also a global catalog)
    - **DomainS2**

  The distinguished name of the domain is **DC=Domain,DC=com**.

- **Child.Domain**, a child domain of Domain, with one domain controller:
    - **ChildS1** (which is also a global catalog)

  The distinguished name of the domain is **DC=Child,DC=Domain,DC=com**.

To perform a comparison between a domain controller in the **Child** domain and a global catalog server in the domain **Domain**, type the following at the command line. (Note that the global catalog server **DomainS1** in the domain **Domain** holds a partial replica of all objects from the domain **Child**.)

```
dsastat -s:DomainS1:3268;ChildS1 -b:DC=Child,DC=Domain,DC=com -
gcattrs:objectclass -p:999
```

This search compares the global catalog on **DomainS1** and checks to see if it has a consistent partial image of the domain

**Child**. This search does not perform a full-content comparison, but simply checks that the object count is the same between the global catalog and the domain controller **ChildS1** in the domain **Child**.

**Example 5: Full Content Comparison**
To perform a full content comparison between a domain controller in the **Child** domain and a global catalog server in the domain **Domain**, type the following at the command line:

```
            dsastat -s:DomainS1:3268;ChildS1 -b:DC=Child,DC=Domain,DC=com -gcattrs:auto
-sort:true -t:false -p:16
```

In this example, DsaStat performs a full content comparison between all the objects that have been replicated to the global catalog. Specifying the **auto** option for **/gcattrs** ensures that only the attributes replicated to the global catalog are retrieved.

**See Also**
**Concepts**
Dsastat Overview
Dsastat Syntax
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Ldp Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Ldp GUI tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Ldp.exe: LDP Tool**
This GUI tool is a Lightweight Directory Access Protocol (LDAP) client that allows users to perform operations (such as connect, bind, search, modify, add, delete) against any LDAP-compatible directory, such as Active Directory. LDP is used to view objects stored in Active Directory along with their metadata, such as security descriptors and replication metadata.

### Note

- LDP is a GUI-based, Windows Explorer–like utility with a scope pane on the left that is used for navigating through the Active Directory namespace, and a details pane on the right that is used for displaying the results of the LDAP operations. Any text displayed in the details pane can be selected with the mouse and "copied" to the Clipboard.

**Corresponding UI**
LDP comes with its own user interface. For more information about the LDP user interface, read LDP.doc in the Program Files\Support Tools directory.

### Note

- You need Microsoft Word or a compatible viewer to view Ldp.doc. You can download a free version of Microsoft Word Viewer from the Microsoft Office Download Center.

**Concepts**
It is highly recommended that you read the Understanding LDAP white paper on the Microsoft Web site before continuing with this document.

**System Requirements**
The following are the system requirements for LDP:

- A connection to an LDAP server
- Windows Server 2003

**File Required**

- Ldp.exe

**See Also**
**Concepts**
Ldp Remarks
Ldp UI
Ldp Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview

ADSI Edit (adsiedit.msc)
Acldiag Overview

# Ldp Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**LDP Remarks**

**Finding Required and Optional Values for an Attribute**

The schema defines objects as well as the attributes and permissible values for each. Schema classes that contain attribute information about objects can be viewed. To view this information, search cn=Schema,cn=Configuration,dc=ForestRootDomain for classSchema objects.

**Understanding Bind Options for LDAP Authentication**

There are several authentication methods available in LDP that allow a client to bind to an LDAP server. The best method depends on several factors.

| Method | Description |
|---|---|
| Simple | Cleartext password. Try not to use this, as it is not secure. |
| MSN | MSN (Microsoft Network) authentication. This package brings up a dialog box to prompt the user for MSN credentials. |
| DPA | Normandy authentication, new MSN authentication. Same usage as MSN. |
| NTLM | Windows NT domain authentication. Use NULL credentials and attempt to use default logged-in user credentials. |
| Sicily | Negotiate with the server for any of: MSN, DPA, NTLM. Should be used for LDAPv2 servers only. |
| Negotiate | Use GSSAPI Negotiate package to negotiate security package of either Kerberos V5 or NTLM (or any other package the client and server negotiate). Pass in NULL credentials to specify default logged-in user. If Negotiate package is not installed on server or client, this will fall back to Sicily negotiation. |

**See Also**

**Concepts**

Ldp Overview
Ldp UI
Ldp Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Ldp UI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**LDP UI**

**LDP Dialog Box**
The LDP dialog box consists of two panes. The left navigation pane lists the base object and any child objects. The right details pane lists the results of the LDAP operations.

**Connection Menu**
**Connection** menu options:

| Option | Description |
|---|---|
| **Connect** | Opens a dialog box that you can use to open a session with a specified LDAP server. A connection must be established with an LDAP server before any other LDAP commands can be run. Enter the appropriate port number for the service you are connecting to. By default, LDAP uses TCP for a connection-oriented session. To use UDP for a connectionless session, select the **Connectionless** check box. By default, a successful connection results in the display of the RootDSE information in the details pane. |
| **Bind** | Opens a dialog box you can use to authenticate a specified LDAP server. Enter a user name and password of an account that has permissions to the LDAP server. If you enter an invalid user name, then you will be connected with anonymous credentials. As a shortcut, use the **Bind** option without using the **Connect** option to connect, and authenticate with the server last connected to. |
| **Advance** | Opens a dialog box that you can use to set options for the authentication methods. See **Bind Options** under the **Options** menu. |
| **Disconnect** | Terminates an open session with a specified LDAP server. Closing the LDP application automatically disconnects any open sessions. |
| **New** | Keeps the currently connected session, but clears the details pane. The keyboard shortcut for this action is CTRL+N. By default, the details pane is accumulative. |
| **Open** | Opens text files and places the information in the details pane. |
| **Save** | Saves changes to a previously saved file. |
| **Save as** | Saves the contents of the details pane to a text file. Use the **Open** command to view the contents of this file in the details pane later. |
| **Print** | Prints the details pane. |

Note

- LDP can connect to only one LDAP server at a time. If you use the **Connect** command while connected to a server, the current session is disconnected.

**Browse Menu**
**Browse** menu options:

**Add**
Opens a dialog box that you can use to add objects to Active Directory. You must enter the full distinguished name of the object, as well as all the mandatory attributes for the class of object being added.

| Option | Description |
|---|---|
| **DN** | Enter the full distinguished name of the new object. |

| | |
|---|---|
| **Attribute** | Enter the required or optional attribute. |
| **Values** | Enter the value(s) associated with the attribute. Separate multiple values for a single attribute with a semicolon. No spaces are required. |
| **Enter** | Adds the entered attribute and value(s) to the **Entry List** section of the dialog box, and clears the **Attribute** and **Value** fields. Continue entering attributes and values until all required and desired optional attributes are on the Entry List. |
| **Insert File** | Opens a dialog box that allows you to use a text file with the appropriate attributes and values. |
| **Entry List** | Displays entered attributes and values. |
| **Edit** | Opens a dialog box that you can use to enter changes to the selected entry from the Entry List. |
| **Remove** | Deletes the selected entry from the Entry List. |
| **Extended** | Select this check box if the object being added is part of an extended control. |
| **Synchronous** | If selected (the default), requires LDP to wait for a response from the destination server before continuing. If you clear this check box, LDP continues before a response is received. Clear this check box when slow WAN connections are causing LDP commands to time out. |
| **Run** | Adds the current attributes and values in the Entry List to Active Directory. If LDP encounters any errors, then the object is not added and the error message displays in the details pane. |
| **Delete** | Opens a dialog box that you can use to delete an object from Active Directory. Attributes can be deleted only if they are defined as optional and contain no values. Use the **Edit** command to delete an attribute's values. |

**DN**

| Option | Description |
|---|---|
| **Dn** | Enter the full distinguished name of the new object. |
| **Extended** | Select this check box if the object being deleted is part of an extended control. |
| **Synchronous** | If selected (the default), requires LDP to wait for a response from the destination server before continuing. If you clear this check box, LDP continues before a response is received. Clear this check box when slow WAN connections are causing LDP commands to time out. |
| **Recursive (client)** | Deletes all objects in a container, but does not delete the container. |

**Modify**

Opens a dialog box you can use to change attributes of an object stored in Active Directory.

| Option | Description |
|---|---|
| **Dn** | Enter the full distinguished name of the new object. |
| **Attribute** | Enter the required or optional attribute. |
| **Values** | Enter the value(s) associated with the attribute. Separate multiple values for a single attribute with a semicolon. No spaces are required. |
| **Insert Files** | Opens a dialog box that allows you to use a text file with the appropriate attributes and values. |
| **Enter** | Adds the entered attribute and value(s) to the **Entry List** section of the dialog box, and clears the **Attribute** and **Value** fields. Continue entering attributes and values until all required and desired optional attributes are on the Entry List. |

| | |
|---|---|
| **Operation Section** | **Add**, **Delete**, or **Replace**. Choose **Add** to add a new value to an existing attribute. Choose **Delete** to permanently remove an attribute from the listed object. Attributes containing data cannot be deleted. Also, attempting to delete required attributes results in an error. Choose **Replace** to replace an existing value with another, or to change listed values for an existing attribute. |
| **Entry List Section** | Displays existing attributes and values for an object. |
| **Edit** | Opens a dialog box you can use to changes the selected entry from the Entry List. |
| **Synchronous** | If selected (the default), requires LDP to wait for a response from the destination server before continuing. If you clear this check box, LDP continues before a response is received. Clear this check box when slow WAN connections are causing LDP commands to time out. |
| **Extended** | Select this check box if the object being modified is part of an extended control. |
| **Run** | Sends the edited values in the Entry List to Active Directory. |

**Modify RDN**

Opens a dialog box you can use to change the relative distinguished name of an object. This option is designed to modify leaf objects only. If you rename the container portion of the distinguished name, then the object moves to the container that is named.

| Option | Description |
|---|---|
| **Old DN** | Enter the current distinguished name of the object. |
| **New DN** | Enter the new distinguished name for the object. |
| **Delete Old** | If selected (the default), removes the old distinguished name from the LDAP directory. |
| **Synchronous** | If selected (the default), requires LDP to wait for a response from the destination server before continuing. If you clear this check box, LDP continues before a response is received. Clear this check box when slow WAN connections are causing LDP commands to time out. |
| **Extended rename** | Select this check box if the object being renamed is part of an extended control. |
| **Run** | Sends the change to Active Directory. |

**Search**

Opens a dialog box that you can use to create a customized search filter and to perform the search on the directory information tree. The search base must be specified as a distinguished name, and the filter must be a valid LDAP filter. Items returned from a search are separated by **>>** characters.

| Option | Description |
|---|---|
| **Base DN** | Enter a distinguished name to specify where the search starts. |
| **Filter** | Enter the search criteria separated by LDAP search filters. Enter attributes and values to find an object or set of objects.<br><br>Note: LDAP search filters are defined in RFC 2254, and in the Knowledge Base Article Q255602 "Browsing and Querying Using the LDP Utility" at Microsoft Product Support Services. |
| **Scope Section** | Specifies how many levels the search encompasses. |
| **Base** | Searches the base object only. |
| **One Level** | Searches objects immediately subordinate to the base object, but does not search the base object. |

| | |
|---|---|
| **Subtree** | Searches the entire subtree, from the base object down to all child objects. |
| **Options Button** | Opens the **Search Options** dialog box. You can use it to apply filters that allow some entries and exclude others from the search, and that allow you to control the way the search is processed. See the **Options** menu for more details. |
| **Run** | Click this button to send the search request to Active Directory. |

**Compare**

Opens a dialog box you can use to compare the value of an attribute of an object with a specified value. The result returned is either true or false.

| Option | Description |
|---|---|
| **DN** | Enter the full distinguished name of the object whose value(s) will be compared. |
| **Attribute** | Enter attribute to be compared. |
| **Values** | Enter the value(s) that will be compared with the existing value in Active Directory. Separate multiple values for a single attribute with a semicolon. No spaces are required. |
| **Synchronous** | If selected (the default), requires LDP to wait for a response from the destination server before continuing. If you clear this check box, LDP continues before a response is received. Clear this check box when slow WAN connections are causing LDP commands to time out. |
| **Run** | Starts the comparison. |

**Extended Op**

Opens a dialog box you can use to submit an extended operation to an LDAP Directory by specifying an LDAP object identifier (OID) and an applicable value.

| Option | Description |
|---|---|
| **Oid** | Enter the object ID number. |
| **Data** | Enter the value of the object identifier (OID) attribute. |
| **Controls** | See **Controls Option** under the **Options** menu. |
| **Send** | Submits the extended operation to Active Directory. |

**GetLastError**

Calls the LDAP **Getlasterror** function.

**Security**

Opens a dialog box you can use to submit an extended operation to an LDAP directory by specifying an LDAP object identifier (OID) and an applicable value.

| Option | Description |
|---|---|
| **Security Descriptor** | Opens a dialog box you can use to view access permissions on an object. |
| **User Rights** | Not implemented yet. |

**Replication**

Opens a dialog box you can use to submit an extended operation to an LDAP directory by specifying an LDAP object identifier (OID) and an applicable value.

| Option | Description |
|---|---|
| **View Metadata** | Opens a dialog box you can use to view replication metadata of an object. This command is useful to identify whether the objects are up to date and replicated between domain controllers. |

**Process Pending**

Opens a dialog box that shows the list of requests that are not finished processing.

**View Menu**

The **View** menu options:

| Option | Description |
|---|---|
| **Status Bar** | Shows or hides the status bar located along the bottom of the LDAP window. |
| **Tree** | Opens a dialog box you can use to specify the base object in the navigation pane. You can enter a default base object in the **General Options** dialog box in the **Auto Base DN Query** field. By default, this field is blank and no object is listed in the navigation pane. |
| **Enterprise Configuration** | Opens the **Live Enterprise Tree** dialog box, which shows a graphic display of all domains and domain controllers in the enterprise. The dialog box indicates whether the domain controllers are online or offline by marking offline domain controllers with a red "X". |
| **Auto Refresh (min):** | Enter the number of minutes for LDP to wait before refreshing the display. |
| **Refresh** | Manually refreshes the display. |

**Options Menu**

**Options** menu options:

**Bind Options**

| Option | Description |
|---|---|
| **Function Type** | Specifies a category of authentication for LDP to use when choosing authentication methods. |
| **Generic** | Specifies to use a standard authentication protocol. |
| **Simple** | Specifies to use no authentication protocol and to send the password in clear text. |
| **Extended** | Not available. |
| **Method** | Selects the type of authentication that LDP uses when passing credentials. |
| **Synchronous** | Specifies that the authenticating server must respond immediately to requests. This option works only with simple authentication. |
| **Use auth. identit** | Allows the use of alternate authentication credentials. All authentication methods except simple require synchronous calls. |

**Search Options**

| Option | Description |
|---|---|
| **Time Limit** | Enter the number of milliseconds that the search can take on the server. By default the maximum is 120 seconds. |
| **Size Limit** | Enter the maximum number of bytes that the search can return. Entering a null value does not place a maximum size on the data returned. |
| **Timeout (s)** | Enter the number of seconds that LDP waits for the LDAP server to respond to a search request. |
| **Timeout (ms)** | Enter the number of milliseconds that LDP waits for the LDAP server to respond to a search request. |
| **Page Size** | Enter the maximum size, in bytes, of each page of returned data. |
| **Attributes** | Specifies which attributes to return in the search. Separate multiple attributes with a semicolon. Use the wildcard character (*****) to indicate all attributes. |

| | |
|---|---|
| **Search Call Type** | Specifies a call type to use in the search. If the search will take some time, you can select **async** to allow you to perform other tasks while waiting for the search to complete. |
| **Attributes Only** | Select this check box to return only attributes of objects. The distinguished name is not returned. |
| **Chase Referrals** | Performs a search for objects found in external LDAP directories. By default, the objects' trusts of external LDAP directories return only a referral instead of the actual object. |
| **Display Results** | Displays a detailed list of objects returned by the search. By default, only a success or failure and the number of objects found display. |
| **Sort Keys** | Opens the **Sort Keys Option** dialog box. See **Sort Keys** under the **Options** menu. |
| **Controls** | Opens the **Controls Option** dialog box. |

**Pending Options**

Opens a dialog box you can use to place filters on the list of processes that have not yet completed.

| Option | Description |
|---|---|
| **All search results** | Specifies that all search results display. |
| **Blocking** | Clear this check box to set a time limit. |
| **Time Limit (sec):** | Enter a time limit in seconds. |
| **Time Limit (millisec):** | Enter a time limit milliseconds. |

**General Options**

**Value Parsing Section**

| Option | Description |
|---|---|
| **Binary** | Displays the LDAP information in its native numerical format. |
| **String** | Converts the LDAP information from its native format to ASCII characters so that it is more readable when displayed. This is the default setting. Values that are too long to be converted are still displayed in binary form. |
| **LDAP Version Section** | Specifies which version of LDAP the server is using. The default is version 3. |
| **DN Processing Section** | Converts the distinguished names, displayed in component parts, by extending the data types that LDP returns when performing a command. |

**Buffer Size Section**

| Option | Description |
|---|---|
| **Page** | Enter the number of returned lines to display per command. |
| **Line** | Enter the number of returned characters to display per command. |
| **Auto default NC query** | Specifies that LDP queries the default naming context when a connection to the LDAP server is made. The default naming context is the RootDSE. This setting is used when the distinguished name value in the **View\|Tree** dialog box is left blank. |

**Virtual List View (VLV) Section**

| Option | Description |
|---|---|

| | |
|---|---|
| **Auto VLV browse when container size is greater than:** | Select this check box to display a virtual list view whenever the object count is greater than the value displayed in the field. The default value is 100. |

**Connection Options**

Opens a dialog box you can use to change the value of any option.

| Option | Description |
|---|---|
| **Option Name** | Enter the name of the option whose value will be reset. |
| **Value** | Enter the new value for the specified option. |
| **Set** | Sends the information to the LDAP directory. |

**TLS Options**

| Option | Description |
|---|---|
| **StartTLS** or **StopTLS** | Starts or stops a secure session with the LDAP server that uses Transport Level Security (TLS). |

**Controls Option**

In the **Controls** dialog box, enter values to extend the functionality of LDAP.

The object identifier must be specified when implementing a control. To obtain a list of object identifiers, view the supportedControls property in the RootDSE of a domain controller. Individual controls are described in the Understanding LDAP white paper on the Microsoft Web site

Note

- Only server controls can be sent to a server. Client controls work only with LDAP APIs.
- To view a list of extended LDAP controls, see the Knowledge Base article Q222560 "Windows 2000 Extended LDAP Controls" at Microsoft Product Support Services.

**Sort Keys Option**

Sort Keys is a type of control that formats the display of search results. To format the display of your search results, in the **Sort Keys** dialog box, enter an attribute type.

For more information find sortKeyRequestControl in the Understanding LDAP white paper on the Microsoft Web site.

Utilities Menu

**Utilities** menu options:

| Option | Description |
|---|---|
| **Large Integer Converter** | To convert long integers into high and low parts, enter a value in the **String** field. |

Note

- LDP can connect to only one LDAP server at a time. Using the **Connect** command while connected to a server disconnects the current session.

**See Also**
**Concepts**
Ldp Overview
Ldp Remarks
Ldp Examples
Alphabetical List of Tools
Search Overview
Replmon Overview

Repadmin Overview
Movetree.exe
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Ldp Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**LDP Examples**
**Example 1: Add a new object to an LDAP directory**
To add a user to Active Directory by using LDP:

1. On the **Browse** menu, click **Add**.

2. In the **Add** dialog box, in the **Dn** input box, enter the distinguished name of the new object.

3. In the **Edit Entry** section, add the new attributes and values. Click **Enter** after typing in each attribute and associated value:

| Attribute | Value |
|---|---|
| userAccountControl | 512 |
| ObjectClass | User |
| SamAccountName | Testuser2 |

4. To add the attribute or value combination to the Entry List box, click **Edit**.

5. To add the information to Active Directory using LDAP APIs when all the attributes are entered, click **Run**.

**Example 2: Search an LDAP directory**
To search Active Directory by using LDP:

- On the **Browse** menu, click **Search**.

  The **Search** dialog box opens.

In the following searches, the LDAP directory is Active Directory. For user names, it contains a **givenName** attribute for first names, and an attribute of **sn** for last names.

- To search for all users that have a first name of John and a last name of either Smith or Jones, type the following in the **Filter** field:

  **(&(objectClass=user)(givenName=John)(|(sn=Smith)(sn=Jones))))**

- To search for users that have a last name of Jones, but filter out those users that have a first name of John or Jane, and also filter out users that have not logged on at least 100 times, type the following in the **Filter** field. (The exclamation point (!) is the NOT operator.)

  **(&(objectClass=user)(sn=Jones)(!givenName=John)(!givenName=Jane)(!logonCount<=100))**

  **Note**

  - To search for reserved characters as part of an attribute value, you must precede the reserved character with an escape character. Use the following escape characters to represent the associated characters:

| Character | Escape characters |
|---|---|

| | |
|---|---|
| * | 2a |
| ( | 28 |
| ) | 29 |
| \ | 5c |
| NUL | 00 |

- To search for all of the users whose display names end in a close parenthesis character, type the following in the **Filter** field:

  **(&(objectClass=user)(displayName=*\29))**

- To search for all users who have a surname that starts with the letter J, type the following in the **Filter** field. (Queries support the wildcard character (*****).)

  **(&(objectClass=user)(sn=j*))**

- To search for users whose home directories are G:\Accounting, type the following in the **Filter** field. (The attribute name is **home-directory**.)

  **(&(objectClass=user)(home-directory=G:\5cACCOUNTING*))**

**Example 3: Control the returns on a search of an LDAP directory**
In this example, you use the **Search Options** dialog box to control which attributes are displayed in a search. Just a few attributes are displayed; the rest are filtered and do not display in the details window of LDP.

To control which attributes display:

- In the **Attributes** input box, type the attributes to display:

  **"memberof;range:1-20"objectClass;objectGUID**

  **Notes**

  - Separate attributes with a semicolon. No spaces are necessary. All other attributes will be filtered out of the display.

  - A range is specified for the **memberof** attribute. Because a semicolon is used, the entire section must be separated from the rest of the attributes by quotation marks.

All searches display only the **memberof**, **objectClass**, and **objectGUID** attributes in the details pane. To return all attributes, replace any existing list of attributes with the wildcard character (*****).

**Example 4: View replication metadata for an object**
To list the replication metadata for an object in Active Directory:

1. On the **Browse** menu, click **Replication|View Metadata**.

   The **View Metadata** dialog box opens.

2. In the **Object Dn** field, enter the distinguished name of the object, then click **OK**.

   The output will look similar to the output in the following table.

| AttID | Ver | Loc.USN | Originating DSA | Org.USN | Org.Time/Date |
|---|---|---|---|---|---|

| | | | | | |
|---|---|---|---|---|---|
| 0 | 1 | 3693 | 9fad4c38-2d76-44b2-84f6-f2fe384f8450 | 3693 | 2000-12-29 09:15.02 |
| 3 | 1 | 3693 | 9fad4c38-2d76-44b2-84f6-f2fe384f8450 | 3693 | 2000-12-29 09:15.02 |
| d | 1 | 3693 | 9fad4c38-2d76-44b2-84f6-f2fe384f8450 | 3693 | 2000-12-29 09:15.02 |
| 20001 | 1 | 3693 | 9fad4c38-2d76-44b2-84f6-f2fe384f8450 | 3693 | 2000-12-29 09:15.02 |
| 20002 | 1 | 3693 | 9fad4c38-2d76-44b2-84f6-f2fe384f8450 | 3693 | 2000-12-29 09:15.02 |

**See Also**

**Concepts**

Ldp Overview
Ldp Remarks
Ldp UI
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Movetree.exe

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Movetree command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Movetree.exe: Active Directory Object Manager**
This command-line tool enables administrators to move Active Directory objects such as contacts between domains in a single forest. Operations such as these are performed to support domain consolidation or organizational restructuring operations.

### Important

- Use the Active Directory Migration Tool (ADMT) included on your Windows Server 2003 CD-ROM to move Active Directory objects such as user, computer, group, and OU. MoveTree should **only** be used in special cases to handle objects (such as contacts) which cannot currently be migrated by ADMT.

**Corresponding UI**
To move users or groups within a Windows Server 2003 domain (for example, from one organizational unit to another), use Active Directory Users and Computers, a Microsoft Management Consolesnap-in that is part of Windows Server 2003.

Note

- You must install Adminpak.msi before you can see the Active Directory Users and Computers snap-in in Administrative Tools.

**Concepts**
For more information about Active Directory, see the Active Directory Overview.

System Requirements

The following are the system requirements for MoveTree:

- Windows XP Professional or Windows Server 2003
- Administrator rights

**Files Required**

- Movetree.dll
- Movetree.exe

**See Also**
**Concepts**
Movetree Remarks
Movetree Syntax
Movetree Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview

Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Movetree Remarks

Updated: December 28, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MoveTree Remarks**
**Before Using MoveTree**
Before using MoveTree take the following steps to maintain peak performance:

1. Review all Group Policy objects that apply to a particular organizational unit, and make a note of the Group Policy settings they contain.

2. Recreate the Group Policy objects, linked to the moved organizational unit in the new domain, with the desired settings.

3. Make sure to remove the Group Policy objects linked from the old domain.

**MoveTree Limitations**
While MoveTree can move some Active Directory objects between domains, certain objects cannot be moved. MoveTree is also unable to move certain associated data that may exist externally to Active Directory. For details about limitations, see the following sections.

**Local and Domain Global Group Limitations**
Local and Domain Global groups are moved during a MoveTree operation, but their members are not. This is because of group membership rules. Therefore, it is important to save and recreate the memberships of Domain Local and Global groups to maintain the existing resource access permissions.

**Universal Group Limitations**
Universal groups including their members, are moved during a MoveTree operation.

**Computer Objects Limitations**
Computer objects are not moved during a MoveTree operation. Use Netdom Overview, another Windows Support Tool, to move computer accounts between domains and to join computers to domains.

**Associated Data**
Associated data that is not moved during a MoveTree operation includes policies, profiles, logon scripts, and users' personal data. Use additional scripts or management tools, such as the Remote Administration Scripts (included in the *Windows 2000 Resource Kit*), in conjunction with MoveTree to perform these additional steps.

### Important

- Remote Administration Scripts in the *Windows 2000 Resource Kit* are examples only. They might require code modification to be used with Windows Server 2003 operating systems. Use this feature only with the assistance of Microsoft Product Support professionals.

MoveTree cannot move the following objects:

- System objects (identified by the objectClass being marked as systemOnly)

- Objects in the configuration or schema naming contexts

- Objects in the special containers in the domain: Builtin, ForeignSecurityPrincipal, System, LostAndFound

- Domain controllers or any object whose parent is a domain controller

- Any object with the same name as an object that already exists in the target domain

MoveTree can fail because of some of the following error conditions:

- The source domain controller cannot transfer the RID role owner.

- The source object is locked due to another operation in progress. (For example, if another user is currently creating child objects under the source object that is selected for the move operation.)
- Either the source or destination domain have invalid credentials.
- The destination knows the source object is deleted but the source does not. (For example, the source object had been deleted on a different domain controller, but because of replication latency the source domain controller has not yet received the deletion event.)
- There is a failure at the destination domain controller (for example, Disk Full).
- A Security Accounts Manager (SAM) constraint is met. (For example, Duplicate SAM Account Name or source object password length does not meet the password restrictions in the target domain.)
- The source and destination have a schema mismatch.

When a MoveTree Operation Is Paused or Halted

During a MoveTree operation, if the process is paused or halted, then any objects that have yet to be moved remain in an *orphan* container in the *Lost And Found* container in the source domain. The *Lost And Found* container can be viewed in the Active Directory Users and Computers snap-in (a Windows administrative tool) when the Advanced View menu option is selected. The *orphan* container is named using the globally unique identifier (GUID) of the parent container being moved and can be readily identified; it will contain the objects that were selected for the MoveTree operation.

For example, if you are moving the organizational unit Sales, which has an object GUID of {123-abc}, and the MoveTree operation is halted, then the tree structure looks like this:

```
Lost + Found
         {123-abc}
               Sales
```

**MoveTree ErrorLevels**
MoveTree returns ErrorLevel 0 for success and ErrorLevels 1 through 5 for different kinds of failure. These values can be used as criteria for branching when the tool is used in a batch file; see Movetree Examples in the *MoveTree Examples* topic.

The following table describes MoveTree error levels.

| Error Level | Meaning |
|---|---|
| 0 | Success |
| 1 | Error – command-line syntax |
| 2 | Error – directory conflict (duplicate names, insufficient privilege, name conflict, immovable object) |
| 3 | Error - network error (DC unavailable) |
| 4 | Error – system resource (Low VM, disk space) |
| 5 | Error – internal processing error |

**See Also**
**Concepts**
Movetree.exe
Movetree Syntax
Movetree Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview

Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Movetree Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MoveTree Syntax**

MoveTree uses the following syntax:

**movetree** {**/start** | **/continue** | **/check**} **/s** *SrcDSA* **/d** *DstDSA* **/sdn** *SrcDN* **/ddn** *DstDN* [**/u** [*Domain*\]*UserName* **/p** *Password*] [**/verbose**]

**Parameters**

**/start**

Starts a MoveTree operation. This command includes a **/check** operation by default. MoveTree tries to continue the operation until it is complete. If there is a network fault or if the destination domain controller becomes unavailable, then MoveTree pauses the operation. If an operation has been paused, it can be continued by using the **/continue** command.

**/continue**

Continues the execution of a previously paused or failed MoveTree operation. This allows the MoveTree operation to continue even if a network fault or a domain controller error has interrupted the initial operation. Specifying **/sdn** *SrcDN* is optional for this command.

**/check**

Performs a test run of the MoveTree operation checking the whole tree without moving any objects. This enables the administrator to determine if there is sufficient disk space on the destination server, if there are any conflicts with object names or if there are any objects that could not be moved (such as Domain Local or Global groups). The administrator can then take remedial action before performing the actual move.

The **/check** command returns an error if any of the following conditions are met:

- The user does not have the necessary permissions to create objects in the destination container.
- The destination server does not have sufficient disk space to continue the operation.
- A relative distinguished name conflict exists on the destination server.
- There is a samAccountName conflict for any object that would be moved.
- Any objects cannot be moved because they are built-in accounts, or they are either a Domain Local or a Global group.
- Any computer objects would be moved. To move computer accounts and join the computers to the domain, use Netdom Overview, a Windows Support Tool.

**/s *SrcDSA***

Specifies the fully qualified primary DNS name of the source server in the domain from which the objects are being moved (for example, server1.marketing.microsoft.com). Required for all MoveTree commands.

**/d *DstDSA***

Specifies the fully qualified primary DNS name of the destination server in the domain to which the objects are being moved (for example, server2.sales.microsoft.com). Required for all MoveTree commands.

**/sdn *SrcDN***

Specifies the distinguished name of the source subtree (the object being moved) (for example, OU=Promotions,DC=Marketing,DC=Microsoft,DC=Com). Required for the **/start**, **/startnocheck**, and **/check** commands; optional for **/continue**.

**/ddn *DstDN***

Specifies the distinguished name of the destination subtree (to which the object is being moved) (for example, OU=Promotions,DC=Sales,DC=Microsoft,DC=Com). Required for all MoveTree commands.

**/u [*Domain*\]*UserName* /p *Password***

Runs MoveTree under the credentials of an alternate user name and password. Optionally, *Domain* can be specified as well. If these optional arguments are not provided, MoveTree uses the credentials of the currently logged-on user.

**/verbose**

Runs MoveTree in verbose mode, which displays more details about the operation as it runs. Optional.

**/? or /help**

Displays this information on a command-line syntax screen.

**See Also**

**Concepts**

Movetree.exe
Movetree Remarks
Movetree Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Movetree Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MoveTree Examples**
These examples assume the following scenario:

In the Marketing domain, there is a Server1 server and a Promotions organizational unit. In the Sales domain, there is a Server2 server. You want to move the Promotions organizational unit from the Marketing domain to the Sales domain, and to rename the new organizational unit Sales Promotions.

**Example 1: Perform a Test MoveTree Operation**
You want to move the Promotions organizational unit from the Marketing domain to the Sales domain and to rename the Promotions organizational unit Sales Promotions. To do a test run, and only perform the move if the test executes without errors, type the following at the command line:

```
movetree /start /s server1.marketing.microsoft.com /d server2.sales.microsoft.com/sdn
OU=Promotions,DC=Marketing,DC=Microsoft,DC=Com /ddn OU=Sales Promotions,DC=Sales,DC=Microsoft,DC=Com
```

**Example 2: Move Tree Without Test**
To move the Promotions organizational unit from the Marketing domain to the Sales domain and then rename the Promotions organizational unit to Sales Promotions without doing a test run first, type the following at the command line:

```
movetree /startnocheck /s server1.marketing.microsoft.com /d server2.sales.microsoft.com /sdn
OU=Promotions,DC=Sales,DC=Microsoft,DC=Com /ddn OU=Sales Promotions,DC=Sales,DC=Microsoft,DC=Com
```

**Example 3: Resume a Failed MoveTree Operation**
A previous MoveTree operation between Server1 in the Marketing domain and Server2 in the Sales domain failed while the objects were being moved into the Sales Promotion organizational unit. To resume the failed operation, type the following at the command line:

```
movetree /continue /s server1.marketing.microsoft.com /d server2.sales.microsoft.com /ddn OU=Sales
Promotions,DC=Sales,DC=Microsoft,DC=Com
```

**Example 4: Test a MoveTree Operation**
Eventually you want to move the Promotions organizational unit from the Marketing domain to the Sales domain, renaming it Sales Promotions. To do a test run and get verbose output to study before you perform the actual move, using the credentials Microsoft\administrator with the password "********", type the following at the command line:

```
movetree /check /s server1.marketing.microsoft.com /d server2.sales.microsoft.com /sdn
OU=Promotions,DC=Sales,DC=Microsoft,DC=Com /ddn OU=Sales Promotions,DC=Sales,DC=Microsoft,DC=Com /verbose /u
Microsoft\administrator /p ********
```

**Example 5: Use MoveTree in a Batch File**
You want to move the Promotions organizational unit from the Marketing domain to the Sales domain and rename the Promotions organizational unit Sales Promotions. To use a batch file to do a test run, and only perform the move if the test executes without errors, create a batch file with the following content:

```
movetree /check
 /s server1.marketing.microsoft.com
 /d server2.sales.microsoft.com
 /sdn OU=Promotions,DC=Marketing,DC=Microsoft,DC=Com
 /ddn OU=Sales Promotions,DC=Sales,DC=Microsoft,DC=Com
if errorlevel 0 goto start
goto exit
:start
movetree /start
 /s server1.marketing.microsoft.com
 /d server2.sales.microsoft.com
 /sdn OU=Promotions,DC=Marketing,DC=Microsoft,DC=Com
 /ddn OU=Sales Promotions ,DC=Sales,DC=Microsoft,DC=Com
:exit
```

For more information about errors in MoveTree, see Movetree Remarks in *MoveTree Remarks*.

**See Also**
**Concepts**

Movetree.exe
Movetree Remarks
Movetree Syntax
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Ntfrsutl Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Ntfrsutl command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Ntfrsutl.exe: File Replication Utility**
This command-line tool dumps the internal tables, thread and memory information for the NT File Replication Service (NTFRS). It runs against local and remote servers.

The recovery setting for NTFRS in Service Control Manager (SCM) can be critical to locating and keeping important log events on the computer. This tool provides a convenient method of reviewing those settings.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
The File Replication Service (FRS) is a key element of the Microsoft Distributed File System (DFS) technologies, and was first introduced in Windows 2000. It is used by Active directory for SYSVOL replication, and by DFS for automatic replication of content between 'alternate targets'.

NTFRS creates text-based log files in the *Systemroot*\debug folder for troubleshooting NTFRS replication problems. NTFRS builds two types of text-based log files.

- The Ntfrsapi.log file contains events that take place during promotion and demotion of a server, namely, the creation of the NTFRS registry keys.

- The Ntfrs_nnnn.log debug log files store transaction and event details in the Ntfrs_nnnn.log through Ntfrs_0005.log files. The most recent NTFRS transactions and events are written to the log with the highest version number in existence at the time.

The Registry entries that control NTFRS logging reside in the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Ntfrs\Parameters registry key.

The number of log files residing in the Debug folder is determined by the Debug Log File entry. Once the log entries are written (determined by the Debug Maximum Log Messages registry entry) , a new log file, Ntfrs_000n-1.log is created. Once the correct number of logs is created (determined by the Debug Log File registry entry), the lowest-version log file is deleted and the remaining log file names are decremented by 1 to make room for a new log file.

**System Requirements**
To access the internal NTFRS information, the user needs specified levels of access to the following registry keys on the target server.

```
       HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Ntfrs\Parameters\
Access Checks\
```

| | |
|---|---|
| Get Internal Information | Full control |
| Get Ds Polling Interval | Read |
| Set Ds Polling Interval | Full Control |

**File Required**

- Ntfrsutl.exe

**See Also**

**Concepts**

Ntfrsutl Syntax
Ntfrsutl Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Ntfrsutl Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NTFRSUtl Examples**

Example 1: Determine NTFRS Polling Interval

In this example, you want to determine the polling interval for file replication. At the command prompt, type:

```
C:\Program Files\SupportTools>ntfrsutl poll wrkstn-1
```

Your output will look similar to the one below:

```
Current Interval:      5 minutes
Short Interval  :      5 minutes
Long Interval   :     60 minutes
```

You can see that the tool ran successfully providing the local interval information that was requested.

Example 2: Determine NTFRS API Version Information

In this example, you want to find information about the current NTFRS application program interface (API). At the command prompt, type:

```
C:\Program Files\SupportTools>ntfrsutl version
```

Your output will look similar to the one below:

```
NtFrsApi Version Information
   NtFrsApi Major      : 0
   NtFrsApi Minor      : 0
   NtFrsApi Compiled on: Mar 27 2002 18:20:58
NtFrs Version Information
   NtFrs Major        : 0
   NtFrs Minor        : 0
   NtFrs Compiled on  : Mar 27 2002 18:21:38
   Latest changes:
   Windows Server 2003 Server
OS Version 5.1 (3620) -
SP (0.0) SM: 0x0112  PT: 0x03
Processor:  INTEL Level: 0x0006  Revision: 0x0803  Processor num/mask: 1/0000000
1
```

The information provided NTFRS versions as well as operating system versions and processor information.

**See Also**

**Concepts**

Ntfrsutl Overview
Ntfrsutl Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview

Cabarc Overview
Bitsadmin Overview
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Ntfrsutl Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NTFRSUtl Syntax**

NTFRSUtl Syntax 1

NTFRSUtl Syntax 2

DS

Sets

Version

Poll

**NTFRSUtl Syntax 1**

Enumerates the idtable/configtable/inlog/outlog of the service.

**Syntax**

- **ntfrsutl**[**idtable**|**configtable**|**inlog**|**outlog**] [*Computer*]

**Parameters**

***Computer***

Connects to the NTFRS service on the specified computer.

**NTFRSUtl Syntax 2**

Lists the NTFRS memory usage of the service.

**Syntax**

- **ntfrsutl**[**memory**|**threads**|**stage**] [*Computer*]

**Parameters**

***Computer***

Connects to the NTFRS service on the specified computer.

**DS**

Lists the NTFRS service's view of the DS.

**Syntax**

- **ntfrsutl ds**[*Computer*]

**Parameters**

***Computer***

Connects to the NTFRS service on the specified computer.

**Sets**

Lists the active replication sets.

**Syntax**

- ntfrsutl sets [*Computer*]

**Parameters**

***Computer***

Connects to the NTFRS service on the specified computer.

**Version**

List the API and NTFRS service versions.

**Syntax**

- **ntfrsutl version**[*Computer*]

**Parameters**

***Computer***

Connects to the NTFRS service on the specified computer.

**Poll**
Lists the current polling intervals.

**Syntax**
**ntfrsutl poll**[**/quickly**[=[*N*]]][**/slowly**[=[*N*]]][**/now**][*Computer*]

**Parameters**
**/quickly[ =[ *N*]]**
Polls quickly.

| | |
|---|---|
| **quickly** | Polls quickly until stable configuration is retrieved. |
| **quickly=** | Polls quickly every default minutes. |
| **quickly=***N* | Polls quickly every *N* minutes. |

**/slowly[ =[ *N*]]**
Polls slowly.

| | |
|---|---|
| **slowly** | Polls slowly until stable configuration is retrieved. |
| **slowly=** | Polls slowly every default minutes. |
| **slowly=***N* | Polls slowly every *N* minutes. |

**/now**
Polls now.

***Computer***
Connects to the NTFRS service on the specified computer.

**See Also**
**Concepts**
Ntfrsutl Overview
Ntfrsutl Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Repadmin Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Repadmin command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Repadmin.exe: Replication Diagnostics Tool**
This command-line tool assists administrators in diagnosing replication problems between Windows domain controllers.

Administrators can use Repadmin to view the replication topology (sometimes referred to as RepsFrom and RepsTo) as seen from the perspective of each domain controller. In addition, Repadmin can be used to manually create the replication topology (although in normal practice this should not be necessary), to force replication events between domain controllers, and to view both the replication metadata and up-to-dateness vectors.

Repadmin.exe can also be used for monitoring the relative health of an Active Directory forest. The operations **replsummary**, **showrepl**, **showrepl /csv**, and **showvector /latency** can be used to check for replication problems.

Usually, the Knowledge Consistency Checker (KCC) manages the replication topology for each naming context held on domain controllers.

| **Important** |
| --- |
| During the normal course of operations, there is no need to create the replication topology manually. Incorrect use of this tool can adversely impact the replication topology. The primary use of this tool is to monitor replication so that problems such as offline servers or unavailable LAN/WAN connections can be identified. |

**Corresponding UI**
Repadmin is used to troubleshoot replication issues in Active Directory. Many of the actions performed at the command line with Repadmin can be accomplished by using Replmon Overview.

**Concepts**
For more information about replication, see Active Directory Replication Technologies.

**System Requirements**
The following are the system requirements for Repadmin:

- Windows XP Professional, Windows Server 2003, or Windows Server 2003 with Service Pack 1 (SP1)
- Administrator rights on the domain controller

**File Required**

- Repadmin.exe

**See Also**
**Concepts**
Repadmin Remarks
Repadmin Syntax
Repadmin Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Movetree.exe
Ldp Overview

Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Repadmin Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Repadmin Remarks**

**Repadmin Terminology**

The following terminology is used in discussing Repadmin syntax:

- *NamingContext* refers to the directory partitions that Active Directory comprises. Naming contexts include the three Read/Write naming contexts — domain, schema, and configuration — and the optional Read-only naming contexts that are present on domain controllers that are global catalog servers. A naming context can also be an application directory partition. A naming context is specified as a distinguished name, which indicates its hierarchical relationship to the forest root domain (for example, DC=MyDomain,DC=Microsoft,DC=Com).

- *GUID* (globally unique identifier) refers to the 128-bit number used to uniquely identify objects stored in the directory (for example, fa1a9e6e-2e14-11d2-aa9b-bbfc0a30094c). The GUID is sometimes referred to in the syntax line as a universally unique identifier (UUID). For the purposes of Repadmin, these two terms are synonymous.

- *DN* is an X.500 distinguished name (for example, CN=Server1,CN=Servers,CN=Default-First-Site-Name,CN=Sites,CN=Configuration,DC=Microsoft,DC=Com).

Difference Between the objectGUID and the InvocationID

In the Repadmin Examples, DC object GUID and DC InvocationID returned by **showrepl** and other commands show identical hexadecimal values. However, these two values identify different objects. DC object GUID is a unique identifier for the NTDS Settings object on the domain controller. The value of DC object GUID does not change unless Active Directory is removed from the domain controller and subsequently reinstalled. DC Invocation ID identifies the directory database on the domain controller. This value changes during the process to restore Active Directory from backup. When a domain controller is first installed, the values for these two identifiers are the same; however, each time Active Directory is restored from backup, the Invocation ID value changes.

**See Also**

**Concepts**

Repadmin Overview
Repadmin Syntax
Repadmin Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Repadmin Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Repadmin Examples**

**Example 1: Display the replication partners of a server**

The following example uses the **showrepl** operation of Repadmin to display the replication partners of Server1. This command is also used to find the objectGUID and InvocationID for a server for use with other operations.

No parameters are required for the **showrepl** operation. A remote connection is assumed; therefore, the server name (*DC* in the syntax) is included.

Type the following at the command prompt:

```
repadmin /showrepl server1.microsoft.com
```

Press Enter and the following output is displayed:

```
repadmin /showrepl server1.microsoft.com
Building7a\server1
DC Options : IS_GC
Site OPtions: (none)
DC object GUID : 405db077-1e28-4825-b225-c5bb9af6f50b
DC invocationID: 405db077-1e28-4825-b225-c5bb9af6f50b

==== INBOUND NEIGHBORS ======================================

CN=Schema,CN=Configuration,DC=microsoft,Dc=com
    Building7b\server2 via RPC
         objectGuid: e55c6c75-75bb-485a-a0d3-020a44c3afe7
         last attempt @ 2002-09-09 12:25.35 was successful.


CN=Configuration,DC=microsoft,Dc=com
    Building7b\server2 via RPC
          objectGuid: e55c6c75-75bb-485a-a0d3-020a44c3afe7
         last attempt @ 2002-09-09 12:25.10 was successful.


DC=microsoft,Dc=com
     Building7b\server2 via RPC
          objectGuid: e55c6c75-75bb-485a-a0d3-020a44c3afe7
          last attempt @ 2001-09-09 12:25.11 was successful.
```

**Example 2: Initiate a replication event between two replication partners**

The following example uses the **replicate** operation of Repadmin to make Server2 initiate replication of the domain directory partition for microsoft.com from Server1. In this example, Server1 is the source server and Server2 is the destination server.

The required parameters for the **replicate** operation are the name of the server that will receive changes (*DestDC* in the syntax), the name of the directory partition (*NamingContext* in the syntax), and the name of the server that will send the changes (*SrcDC* in the syntax).

Type the following at the command prompt:

```
repadmin /replicate server2.microsoft.com server1.microsoft.com dc=microsoft,dc=com
```

Press Enter and the following output is displayed:

```
server2.microsoft.com
Sync from server1.microsoft.com to server2.microsoft.com completed successfully.
```

**Example 3: Initiate a replication event for a specified directory partition with all of its replication partners**

The following example uses the **syncall** operation of Repadmin to make Server1 initiate replication of the domain directory partition for microsoft.com from all of its source replication partners in the same site and to make all of the source replication partners initiate replication for microsoft.com from all of their source replication partners in the same site, and so on.

The required parameter for the **syncall** operation is the server name (*DC* in the syntax). The name of the directory partition (*NamingContext* in the syntax) that will be synchronized is also included in this example. If this name is not included, only the configuration partition is synchronized.

Type the following at the command prompt:

```
repadmin /syncall server1.microsoft.com dc=microsoft,dc=com
```

Press Enter and the following output is displayed:

```
repadmin /syncall server1 dc=microsoft,dc=com
Syncing partition: dc=microsoft,dc=com
CALLBACK MESSAGE: The following replication is in progress:
    From: fea22f1d-a456-4f70-aa06-bedbda29e7eb._msdcs.microsoft.com
    To  : 5c02bcaf-86d9-4bed-811e-d17a5cebf8bb._msdcs.microsoft.com
CALLBACK MESSAGE: The following replication completed successfully:
    From: fea22f1d-a456-4f70-aa06-bedbda29e7eb._msdcs.microsoft.com
    To  : 5c02bcaf-86d9-4bed-811e-d17a5cebf8bb._msdcs.microsoft.com
CALLBACK MESSAGE: SyncAll Finished.
SyncAll terminated with no errors.
```

**Example 4: Display the highest Update Sequence Number on a server**

The following example uses the **showutdvec** operation of Repadmin to show the highest USNs for a specified directory partition on each replication partner. In this example, there are only two replication partners and the directory partition is the domain directory partition for the microsoft.com domain.

The only required parameter for the **showutdvec** operation is the distinguished name of the directory partition (*NamingContext* in the syntax). A remote connection is assumed so a server name (*DC_LIST* in the syntax) is also included.

Type the following at the command prompt:

```
repadmin /showutdvec . dc=microsoft,dc=com server2.microsoft.com
```

Press Enter and the following output is displayed:

```
repadmin running command /showutdvec against server localhost

Caching GUIDs.
..
Building7b\Server1    @ USN    295458 @ Time 2002-09-09 19:33:42
Building7b\Server2    @ USN    338194 @ Time 2002-09-09 19:38:16
```

**Example 5: View unreplicated changes between two servers**

The following example uses the **showchanges** operation of Repadmin to view changes that have not yet replicated between Server1 and Server2. In this example Server1 is the source server and is sending the changes while Server2 is the destination server and is receiving the changes.

This is one implementation of the **showchanges** operation. For another implementation of this operation see Example 6: Create a file to determine what changes have occurred over a period of time.

The required parameters for this version of the **showchanges** operation are the objectGuid of the directory partition on the source server (referred to as *SourceDCObjectGUID* in the syntax line) and the name of the directory partition (referred to in the syntax line as *NamingContext*). A remote connection is assumed so the destination server name (referred to in the syntax line as *DestDC*) is also included.

Type the following at the command prompt:

```
repadmin /showchanges server2.microsoft.com 5c02bcaf-86d9-4bed-811e-d17a5cebf8bb dc=microsoft,dc=com
```

Press Enter and the following output is displayed:

```
server2.microsoft.com
```

```
Building starting position from destination server rktlabdc2.rktlabdom.com

Source Neighbor:
dc=microsoft,dc=com
==== INBOUND NEIGHBORS ======================================

dc=microsoft,dc=com
    Building7b\Server1 via RPC
        DC object GUID: 5c02bcaf-86d9-4bed-811e-d17a5cebf8bb
        Address: 5c02bcaf-86d9-4bed-811e-d17a5cebf8bb._msdcs.microsoft.com
        DC invocationID: 064152bc-f8e8-404f-bd64-cdd9bb3958cb
        SYNC_ON_STARTUP DO_SCHEDULED_SYNCS WRITEABLE
        USNs: 296048/OU, 296048/PU
        Last attempt @ 2002-09-09 20:03:53 was successful.

Destination's up-to-date vector:
064152bc-f8e8-404f-bd64-cdd9bb3958cb @ USN 296163
fea22f1d-a456-4f70-aa06-bedbda29e7eb @ USN 338287

==== SOURCE DC: 5c02bcaf-86d9-4bed-811e-d17a5cebf8bb._msdcs.microsoft.com ====

No Changes
```

**Example 6: Create a file to determine what changes have occurred over a period of time**

The following example uses the **showchanges** operation of Repadmin to create a file that records replication changes. By running the **showchanges** operation later you can compare the file created earlier to the current replication state.

This is one implementation of the **showchanges** operation. For another implementation of this operation see Example 5: View unreplicated changes between two servers.

The only required parameter for this version of the **showchanges** operation is the name of the directory partition (*NamingContext* in the syntax) on which the check should be performed. In this example, the check is performed remotely so the server name (*SourceDC* in the syntax) is included as well as the **/cookie:** parameter, along with the name of the file to be created.

Type the following at the command prompt:

```
repadmin /showchanges dc=microsoft,dc=com server2.microsoft.com /cookie:microsoft.txt
```

Press Enter and the following output is displayed:

```
pDcList->szSp.. server2
Using cookie from file test.txt (132 bytes)

==== SOURCE DC: server2 ====

Objects returned: 2
(0) modify CN=00000000000000000000000000000000,CN=VolumeTable,CN=FileLinks,CN=Sy
stem,DC=rktlabdom,DC=com
    1> objectGUID: c4955e2f-ab7c-4f96-bdb6-bf29b97ce3df
    1> instanceType: 0x4 = ( IT_WRITE )
    1> seqNotification: 130
(1) modify DC=..SerialNo-server1.microsoft.com,DC=microsoft.com,CN=MicrosoftDN
S,CN=System,DC=microsoft,DC=com
    1> objectGUID: 0422b130-bf39-4549-aeea-64ed264d10c2
    1> instanceType: 0x4 = ( IT_WRITE )
    1> dnsRecord: <32 byte blob>
New cookie written to file microsoft.txt (132 bytes)
```

**Example 7: Display the connection objects for a server**

The following example uses the **showconn** operation of Repadmin to show connection objects for a server.

No parameters are required for **showconn** operation. In this example, a remote connection is assumed so the server name (*DC_LIST* in the syntax) is specified. All connection objects for Server2 are shown.

Type the following at the command prompt:

```
repadmin /showconn server2.microsoft.com
```

Press Enter and the following output is displayed:

```
repadmin running command /showconn against server server2.microsoft.com

Show Connection Objects
Base DN:   CN=Building7b,CN=Sites,CN=Configuration,DC=microsoftDc=com
==== KCC CONNECTION OBJECTS ============================================
Connection --
    Connection name : b415ba00-8d8d-429b-8a3d-21fd06a99a6c
    Server DNS name : server2.microsoft.com
    Server DN  name : CN=NTDS
Settings,CN=SERVER2,CN=Servers,CN=Building7b,CN=Sites,CN=Configuration,DC=microsoft,DC=
com
 Source: Building7b\Server2
         No Failures.
 TransportType: intrasite RPC
 options:  isGenerated
  Reason:  RingTopology
          Replica link has been added.
  ReplicatesNC: DC=DomainDnsZones,DC=microsoft,DC=com
  Reason:  RingTopology
          Replica link has been added.
  ReplicatesNC: CN=Configuration,DC=microsoft,DC=com
  Reason:  RingTopology
          Replica link has been added.
  ReplicatesNC: DC=microsoft,DC=com
  Reason:  RingTopology
          Replica link has been added.
  ReplicatesNC: CN=Schema,CN=Configuration,DC=microsoft,DC=com
  Reason:  RingTopology
          Replica link has been added.
Connection --
    Connection name : 2357ff7a-4e54-46e2-a387-2e35b0560ab7
    Server DNS name : server2.microsoft.com
    Server DN  name : CN=NTDS
Settings,CN=RKTLABDC2,CN=Servers,CN=Building7b,CN=Sites,CN=Configuration,DC=microsoft,D
C=com
        Source: Building7b\Server2
                No Failures.
        TransportType: intrasite RPC
        options:  isGenerated
        ReplicatesNC: DC=DomainDnsZones,DC=microsoft,DC=com
        Reason:  RingTopology
                Replica link has been added.
        ReplicatesNC: DC=ForestDnsZones,DC=microsoft,DC=com
        Reason:  RingTopology
                Replica link has been added.
        ReplicatesNC: CN=Configuration,DC=microsoft,DC=com
        Reason:  RingTopology
                Replica link has been added.
        ReplicatesNC: DC=microsoft,DC=com
        Reason:  RingTopology
                Replica link has been added.
        ReplicatesNC: CN=Schema,CN=Configuration,DC=microsoft,DC=com
        Reason:  RingTopology
                Replica link has been added.
2 connections found.
```

**Example 8: Display the replication signature for a server**

The following example uses the **showsig** operation of Repadmin to show the replication signature for a server.

No parameters are required for the **showsig** operation. In this example, a remote connection is assumed so the server name (*DC_LIST* in the syntax) is specified.

Type the following at the command prompt:

```
repadmin /showsig server1.microsoft.com
```

Press Enter and the following output is displayed:

```
server1.microsoft.com
Building7a\server1

Current DC invocationID: 415db077-1e28-4588-b255-c5bb9af6f50b (current)

No retired signatures.
```

**Example 9: Display the context handles for the replication process**
The following example uses the **showctx** operation of Repadmin to show the open connections to the server that are established by remote servers.

No parameters are required for the **showctx** operation. This example specifies the server name (*DC_LIST* in the syntax) on which the check should be performed.

Type the following at the command prompt:

```
repadmin /showctx server2.microsoft.com
```

Press Enter and the following output is displayed:

```
server1.microsoft.com
Caching GUIDs.
..
3 open context handles.

NTDSAPI client @ 10.193.36.209 (PID 1528) (Handle 0x20b8a38)
    bound, refs=1, last used 2002-09-09 14:21:41

Building7a\Server1 @ 10.193.36.210 (PID 460) (Handle 0x20b8c40)
    bound, refs=1, last used 2002-09-09 14:23:22

NTDSAPI client @ 10.193.36.210 (PID 2976) (Handle 0x20bb960)
    bound, refs=2, last used 2002-09-09 14:23:28
```

**Example 10: Showing the replication status of a forest using replsummary and wildcard characters.**
The following example uses the **replsummary** operation and a wildcard character to show a summary of the replication status for all the domain controllers in the forest with a name beginning with 'FOURTH'.

Type the following at the command prompt:

```
repadmin /replsummary FOURTH*
```

Press Enter and the following output is displayed:

```
repadmin /replsummary FOURTH*
Replication Summary Start Time: 2002-09-18 14:54:49

Beginning data collection for replication summary, this may take awhile:

Source DC           largest delta  fails/total  %%  error
 FOURTH-CLT-DC-01             54m:57s    0 /   9    0
 FOURTH-DC-05                 41m:23s    0 / 175    0
 FOURTH-DC-06                 55m:08s    0 /  66    0
 FOURTH-DC-07                 09m:29s    0 /  97    0
 FOURTH-DC-08             18h:05m:02s   56 /  56  100  (1722) The RPC server is
unavailable.
 FOURTH-DC-09                 56m:47s    0 /  12    0
 FOURTH-DC-10                 55m:10s    0 /  13    0
 FOURTH-DC-11                 56m:48s    0 /  46    0
 FOURTH-DC-12                 57m:09s    0 /  34    0
 FOURTH-DC-13                 55m:52s    0 /  64    0
 FOURTH-DC-14                 55m:52s    0 /  85    0
 FOURTH-DC-15                 09m:21s    0 /  50    0
 FOURTH-DC-16                 58m:02s    0 /  37    0
 FOURTH-DC-17                 57m:00s    0 /  48    0
 FOURTH-DC-18                 57m:00s    0 /  22    0
 FOURTH-DC-19                 58m:02s    0 /  22    0
 FOURTH-DC-20                 08m:05s    0 /  22    0
```

```
 FOURTH-DC-21                  05m:43s    0 /  22    0
 FOURTH-DC-22             09h:06m:29s    0 / 199    0
 FOURTH-DC-23                  09m:29s    0 /  34    0
 FOURTH-DC-24                  56m:48s    0 / 184    0
 FOURTH-DC-25                  41m:22s    0 /  71    0
 FOURTH-DC-26                  08m:16s    0 / 108    0
 FOURTH-DC-30                  55m:13s    0 /  47    0
 FOURTH-DC-31                  41m:23s    0 /  56    0
 FOURTH-SVC-DC-40              57m:02s    0 /   9    0
 FOURTH-TK-DC-27               07m:02s    0 /  54    0
 FOURTH-TK-DC-28               08m:01s    1 / 161    0  (8461) The replication operation
was preempted.
 FOURTH-TK-DC-29               55m:10s    0 / 115    0

Experienced the following operational errors trying to retrieve replication
information:
          58 - fourth-dc-08.fourthcoffee.com
```

**Example 11: Showing the attributes of a specific object.**
The following example uses the **showattr** operation to show the attributes of a specific object in the Active Directory.

Type the following at the command prompt:

```
repadmin /showattr fsmo_dnm: ncobj:config: /subtree /filter:(objectClass=crossRef)
/atts:nCName,dnsRoot,net,dnsRoot,net,biosname,systemFlags /homeserver:FOURTH-DC-26
```

Press Enter and the following output is displayed:

```
repadmin running command /showattr against server aseanl-test2.fourthcoffee.com

DN: CN=Enterprise Configuration,CN=Partitions,CN=Configuration,DC=fourthcoffee,DC=com
    1> nCName: CN=Configuration,DC=fourthcoffee,DC=com
    1> dnsRoot: fourthcoffee.com
    1> systemFlags: 0x1 = ( FLAG_CR_NTDS_NC )
DN: CN=fourthcoffee,CN=Partitions,CN=Configuration,DC=fourthcoffee,DC=com
    1> nCName: DC=fourthcoffee,DC=com
    1> dnsRoot: fourthcoffee.com
    1> systemFlags: 0x3 = ( FLAG_CR_NTDS_NC | FLAG_CR_NTDS_DOMAIN )
DN: CN=Enterprise Schema,CN=Partitions,CN=Configuration,DC=fourthcoffee,DC=com
    1> nCName: CN=Schema,CN=Configuration,DC=fourthcoffee,DC=com
    1> dnsRoot: fourthcoffee.com
    1> systemFlags: 0x1 = ( FLAG_CR_NTDS_NC )
DN: CN=866c366e-7877-49e5-8121-
67eefd488551,CN=Partitions,CN=Configuration,DC=fourthcoffee,DC=com
    1> nCName: DC=DomainDnsZones,DC=fourthcoffee,DC=com
    1> dnsRoot: DomainDnsZones.fourthcoffee.com
    1> systemFlags: 0x5 = ( FLAG_CR_NTDS_NC | FLAG_CR_NTDS_NOT_GC_REPLICATED )
DN: CN=a07b7cea-1754-478c-84a5-
c7b127ba5024,CN=Partitions,CN=Configuration,DC=fourthcoffee,DC=com
    1> nCName: DC=ForestDnsZones,DC=fourthcoffee,DC=com
    1> dnsRoot: ForestDnsZones.fourthcoffee.com
    1> systemFlags: 0x5 = ( FLAG_CR_NTDS_NC | FLAG_CR_NTDS_NOT_GC_REPLICATED )
```

**See Also**
**Concepts**
Repadmin Overview
Repadmin Remarks
Repadmin Syntax
Alphabetical List of Tools
Search Overview
Replmon Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Repadmin Syntax

Updated: October 17, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### Repadmin.exe Syntax

Repadmin uses the following general syntax:

**repadmin** *Operation Parameters* [**/rpc**] [**/ldap**][**/u:***Domain**User*] [**/pw:**{*Password* | *****}]

**/rpc**
This parameter forces Repadmin to use an RPC session for network communications.

**/ldap**
This parameter forces Repadmin to use an LDAP session for network communications if posable. If the LDAP connection fails then Repadmin attempts to connect using an RPC session. This is the default behavior for Repadmin.

**/u**
Specifies the user name that has permissions to perform operations in Active Directory.

**/pw**
Specifies the password for the user name entered with the **/u** parameter.

#### The DC_LIST Parameter

This section explains the syntax of the semi-global *dc_list* parameter

{*dc_name***|** * |*partial_server_name**| **site:***site_name* **|gc: |fsmo**_*type*:[*name*]}

***dc_name***
Specifies the host name of a domain controller.

*****
Standard wildcard character. When this is used it returns all domain controllers in the enterprise. Improper use can cause a significant increase in network traffic.

***partial_server_name****
Partial wildcard character matching. For example, if given the partial domain controller name "woodgrovebank*", the wildcard character would return woodgrovebank-dc1 and woodgrovebank-dc2, but not woodgrove-diff-name.

**site: *site_name***
The **site:** parameter takes the name of an Active Directory site (*site_name*), and returns all domain controllers in that site.

**gc:**
Queries all Global Catalog Servers in the enterprise.

**fsmo_ *type*:[*name*]**
Specifies a group of domain controllers to query by operations master (also know as flexible single master operations or FSMO) type. Valid operations master types follow.

| Operations Master Type | Description |
|---|---|
| **fsmo_pdc**:[*name*] | Runs the repadmin.exe command against the PDC Emulator. The *name* parameter takes a naming context. |
| **fsmo_rid**:[*name*] | Runs the repadmin.exe command against the RID operations master. The *name* parameter takes a naming context. |
| **fsmo_im**:[*name*] | Runs the repadmin.exe command against the IM operations master. The *name* parameter takes a naming context. |
| **fsmo_istg**:[*name*] | Runs the repadmin.exe command against the ISTG operations master. The *name* parameter takes a site DN. |

| | |
|---|---|
| **fsmo_dnm**: | Runs the repadmin.exe command against the DNM operations master. |
| **fsmo_schema**: | Runs the repadmin.exe command against the Schema operations master. |

**Operations**

Repadmin bind

Repadmin bridgeheads

Repadmin checkprop

Repadmin dsaguid

Repadmin failcache

Repadmin istg

Repadmin kcc

Repadmin latency

Repadmin notifyopt

Repadmin options

Repadmin queue

Repadmin querysites

Repadmin regkey

Repadmin removelingeringobjects

Repadmin replicate

Repadmin replsingleobj

Repadmin replsummary

Repadmin showattr

Repadmin showcert

Repadmin showchanges

Repadmin showconn

Repadmin showctx

Repadmin showism

Repadmin showmsg

Repadmin showncsig

Repadmin showobjmeta

Repadmin showoutcalls

Repadmin showproxy

Repadmin showrepl

Repadmin showsig

Repadmin showtime

Repadmin showtrust

Repadmin showutdvec

Repadmin showvalue

Repadmin siteoptions

Repadmin syncall

Repadmin viewlist

Repadmin oldhelp

**Repadmin bind**
Connects to, and displays the replication features for, a directory partition on a domain controller.

**repadmin /bind** [*DC_LIST*]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller, or a list of domain controllers separated by a space. See above for detailed syntax.

**Repadmin bridgeheads**
Lists the bridgehead servers for a specified site.

**repadmin /bridgeheads** [*DC_LIST*] [**/verbose**]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller, or a list of domain controllers separated by a space. See above for detailed syntax.

**/verbose**
Lists detailed information.

**Repadmin checkprop**
Compares properties of specified domain controllers to determine if they are up-to-date with each other. The source domain controller contains the original information that needs to be checked. The destination domain controller data is compared to the source domain controller data.

**repadmin /checkprop** [*DC_LIST*] *NamingContext OriginatingDCInvocationID OriginatingUSN*

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller, or a list of domain controllers separated by a space. See above for detailed syntax.

***NamingContext***
Specifies the distinguished name of the directory partition on the source domain controller.

***OriginatingDCInvocationID***
Specifies the unique hexadecimal number that identifies an object on a source domain controller. The InvocationID can be retrieved by using the **showrepl** operation.

***OriginatingUSN***
Specifies the USN for the object on the source domain controller. The USN is for the object whose InvocationID is already listed.

**Repadmin dsaguid**
Returns a server name when given a GUID.

**repadmin /dsaguid** [*DC_LIST*] [*GUID*]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller, or a list of domain controllers separated by a space. See above for detailed syntax.

***GUID***
Specifies the unique hexadecimal number that identifies the domain controller. The GUID can be retrieved by using the **showrepl** operation.

**Repadmin failcache**
Displays a list of failed replication events detected by the Knowledge Consistency Checker (KCC).

**repadmin /failcache** [*DC_LIST*]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space. See above for detailed syntax.

**Repadmin istg**
Returns the server name of the ISTG server for a specified site.

**repadmin /istg** [*DC_LIST*] [**/verbose**]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller, or a list of domain controllers separated by a space. See above for detailed syntax.

**/verbose**
Lists detailed information.

**Repadmin kcc**
Forces the KCC to recalculate replication topology for a specified domain controller. By default this recalculation occurs every 15 minutes.

**repadmin /kcc** [*DC_LIST*] [**/async**]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space. See above for detailed syntax.

**/async**
Specifies that replication will be asynchronous. This means that Repadmin starts the replication event, but it does not expect an immediate response from the destination domain controller. Use this parameter to start **kcc** and not wait for it to finish.

**Repadmin latency**
Displays the amount of time between replications using the ISTG Keep Alive timestamp.

**Notes**

- These time stamps are not used in Windows Server 2003 Native Mode, and results are inaccurate.
- In Windows Server 2003 Native Mode, the **repadmin /showutdvec /latency** is used in place of the operation.

**repadmin /latency** [*DC_LIST*] [**/verbose**]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller, or a list of domain controllers separated by a space. See above for detailed syntax.

**/verbose**
Lists detailed information.

**Repadmin notifyopt**
Displays and sets the notification timing settings for replication of a specified directory partition.

**repadmin /notifyopt** [*DC_LIST*] *NamingContext* [**/first:***Value*] [**/subs:***Value*]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller, or a list of domain controllers separated by a space. See above for detailed syntax.

***NamingContext***
Specifies the distinguished name of the directory partition on the source domain controller.

***first***
The number of seconds after a change is made before the domain controller notifies its first replication partner that there is a change.

***subs***
Once the first replication partner is notified of a change, the **subs** parameter specifies the number of seconds to wait before notifying the next replication partner.

**Repadmin options**
Changes default Knowledge Consistency Checker (KCC) behavior on the NTDS Settings object of a domain controller.

Syntax for **/options** is available by typing:

**repadmin /experthelp**

| **Important** |
| --- |
| Setting these options might impair or disable the ability of the KCC to calculate the topology automatically. These options are intended for use by advanced administrators for special situations. If you set these options, document the settings carefully to ensure that the KCC functions as expected and can be restored to default functionality if necessary. |

**repadmin /options** [*DC*] [{+|-}IS_GC] [{+|-}DISABLE_INBOUND_REPL]

[{+|-}DISABLE_OUTBOUND_REPL] [{+|-}DISABLE_NTDSCONN_XLATE]

**Parameters**
***DC***
Specifies the host name of the domain controller on which to set the NTDS Settings options.

**{+|-}IS_GC**
Adds or removes the global catalog.

**{+|-}DISABLE_INBOUND_REPL**
Stops (+) or restarts (-) inbound replication.

**{+|-}DISABLE_OUTBOUND_REPL**
Stops (+) or restarts (-) outbound replication.

**{+|-}DISABLE_NTDSCONN_XLATE**
Prevents connections from forming replication partnerships.

**Repadmin queue**
Displays tasks waiting in the replication queue.

**repadmin /queue** [*DC_LIST*]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller, or a list of domain controllers separated by a space. See above for detailed syntax.

**Repadmin querysites**
Uses routing information to determine cost of a route from a specified site to another specified site or sites.

**repadmin /querysites** *FromSiteRDN ToSite1RDN* [*ToSite2RDN*...

**Parameters**
**Note**

- The Relative Distinguished Names (RDNs) used in this command are case-sensitive.

***FromSiteRDN***
Specifies the RDN of the site from which the cost is calculated.

***ToSiteRDN***

Specifies the RDN of the site to which the cost is calculated.

**Note**

- The **querysites** parameter does not allow the use of alternate credentials.

**Repadmin regkey**

Changes the value of an entry in a registry key. This command is available only in the version of Repadmin that is included with Windows Support Tools in Windows Server 2003 Service Pack 1 (SP1). This command can be applied only on domain controllers running Windows Server 2003 with SP1.

**repadmin /regkey** [*DC_LIST*]{**+**|**-**}*value*]

**Parameters**

***DC_LIST***

Specifies the host name of a domain controller, a list of domain controllers separated by a space, or * for all domain controllers in the forest. See above for detailed syntax.

**{+|-}*value*]**

Specifies enabling (+) or disabling (-) the registry entry specified in *value*. The following table describes the value that can be applied on domain controllers running Windows Server 2003 with SP1.

| Value | Description |
|---|---|
| **strict** | The **Strict Replication Consistency** entry in **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** |

**Repadmin removelingeringobjects**

Uses an authoritative domain controller to compare the directory of a domain controller (destination) that is suspected of having lingering objects against the directory of a domain controller (source) that is designated as a reference source for up-to-date values for the domain of the destination. When the advisory mode parameter is used, this command provides a list of found lingering objects. When the advisory mode parameter is not used, this command removes lingering objects from the destination domain controller.

Syntax for **/removelingeringobjects** is available by typing:

**repadmin /experthelp**

| **Important** |
|---|
| To use the **repadmin /removelingeringobjects** command, both source and destination domain controllers must be running Windows Server 2003. |

**repadmin /removelingeringobjects** *Destination_DC_List SourceDCGUID NamingContextDN* **/advisory_mode**

**Parameters**

| **Note** |
|---|
| "Destination" in the context of this syntax does not indicate a destination replication partner, but rather a source replication partner that is suspected of having lingering objects. The term "destination" identifies the server as the focus of the lingering object removal process. |

***Destination_DC_List***

Specifies the host name of a destination domain controller, a list of domain controllers separated by a space, or * for all domain controllers in the enterprise. See above for detailed syntax.

***SourceDCGUID***

Specifies the GUID of the source (reference) domain controller that will be used to identify objects that are outdated (lingering) on the destination. Obtain the GUID by running **repadmin /showrepl** against the source domain controller that you are using as the reference server.

***NamingContextDN***

Specifies the distinguished name of the directory partition from which the lingering objects are to be removed.

**advisory_mode**
Prints a list of objects that are found in the directory of the destination domain controller that are not found in the directory of the source domain controller. Objects are not removed if this parameter is used.

#### Repadmin replicate

Starts a replication event for the specified directory partition between the source and destination domain controllers. The source UUID can be determined when viewing the replication partners by using the Repadmin showrepl operation.

**Syntax 1**

- **repadmin /replicate** *Destination_DC_LIST Source_DC_NAME* [**/force**] [**/async**] [**/full**] [**/addref**]

**Syntax 2**

- **repadmin /replicate** *Destination_DC_LIST* [**/force**] [**/async**] [**/full**] [**/addref**] [**/allsources**]

**Parameters**

***Destination_DC_LIST***
Specifies the host name of the destination domain controller (Directory Server Agent) with which you want to replicate. See above for detailed syntax on the *DC_LIST* parameter. This parameter cannot be used with **/allsource**.

***Source_DC_NAME***
Specifies the host name of the source domain controller with which you want to replicate. This parameter accepts a GUID, GUID-based DNS name, or the name of a server object.

**/force**
This parameter is used to override the Disable Replication option on a server.

**/async**
Specifies that the replication will be asynchronous. This means that Repadmin starts the replication event, but it does not expect an immediate response from the destination domain controller. Use this parameter when there are slow links between domain controllers.

**/full**
Forces a full replication of all objects from the destination domain controller.

**/addref**
Directs the source to check for a notification entry on the source. If the source does not have a notification entry for this destination, one is added.

**/allsources**
A given destination can have multiple sources for the same naming context. Directs the destination to sync with all sources instead of just one. This parameter cannot be used with *Destination_DC_LIST*

**/readonly**
This parameter is ignored by the **/replicate** operation.

#### Repadmin replsingleobj

Replicates a single object between any two domain controllers that have partitions in common. The two domain controllers do not have a replication agreement. Replication agreements can be shown by using the Repadmin **/showrepl** command.

- **repadmin /replsingleobject** [*DC_LIST*] *DsaSourceGUID ObjectDN*

**Parameters**

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

***DsaSourceGUID***
Specifies the unique hexadecimal number that identifies the object that will be replicated. The objectGUID can be retrieved by using the **showrepl** operation.

***ObjectDN***
Specifies the distinguished name of the object.

#### Repadmin replsummary

The **replsummary** operation quickly and concisely summarizes the replication state and relative health of a forest.

**repadmin /replsummary** [*DC_LIST*] [**/bysrc**] [**/bydest**] [**/errorsonly**] [**/sort**:{*delta* | *partners* | *failures* | *error* | *percent*}]

**Parameters**

***DC_LIST***

Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

**/bysrc**

This parameter shows the output of the **repadmin.exe/replsummary** command, from the perspective of the replication source, in the form of a table. This means that a given source server is "pulled on" by multiple client computers. The table is sorted in order of the source computers that are having the most problems, across all the clients in the forest. This parameter does not display the destination domain controller.

**/bydest**

This parameter shows the output of the **repadmin.exe/replsummary** command, from the perspective of the replication destination, in the form of a table. This means that the "client" is pulling the changes from a replication source or "server", and the table shows the inbound systems, and what problems they are having with their partners. The table is sorted in the order of the inbound systems that are having the most problem with inbound replication, across all the possible partners in the forest. This parameter does not display the source domain controller.

**/errorsonly**

Shows only the domain controllers where the partner error is non-zero.

**/sort:{*delta*| *partners*| *failures*| *error*| *percent*}**

Sorts the **replsummary** table by the specified column heading.

**Notes**

- The **/bysrc** and **/bydest** parameters may be specified at the same time. If this is done, then Repadmin displays the **/bysrc** table first and the **/bydest** table next.
- If the parameters **/bysrc** and **/bydest** are both absent, Repadmin picks the best one, and displays the one with the least number of partner errors.

**Repadmin showattr**

The **/showattr** operation displays the attributes of an object.

**repadmin /showattr** *DC_LIST* [*OBJ_LIST*] [*OBJ_LIST_OPTIONS*] [**/atts:***attribute,attribute,...*] [**/allvalues**] [**/long**] [**/nolongblob**] [**/nolongblob**] [**/nolongfriendly**] [**/dumpallblob**]

**Parameters**

***DC_LIST***

Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax. This is a required parameter.

***OBJ_LIST***

This parameter takes a DN or a special keyword that expands into a DN. The keywords are:

- **Ncobj:config:**
- **Ncobj:schema:**
- **Ncobj:domain:** This key word is the domain partition DN of the home server.
- **Dsaobj:** This keyword is the NTDS settings object of the home server.

***OBJ_LIST_OPTIONS***

The OBJ_LIST_OPTION parameter is required to perform a generic LDAP search from the command line. The parameter requires a Base DN, with the ability to use a search modifier option. The valid search modifier options are as follow:

- **/filter:**<ldap_filter>

- **/base**
- **/subtree**
- **/onelevel**

See the LDAP documentation for the LDAP filter syntax.

**/atts:<*attributes,attributes,...***
Returns only the attributes specified. Separate each listed attribute with a comma.

**/allvalues**
For an attribute, the tool only displays 20 values unless this flag is specified, in which case it shows all values.

**/long**
Displays one value per line.

**/nolongblob**
This parameter is ignored by the **/showattr** operation.

**/nolongfriendly**
This parameter is ignored by the **/showattr** operation.

**/friendlyblob[:*attributeattribute...*]**
If any attributes are specified, only those attributes named are displayed in friendly format. Otherwise all known blobs are formatted.

**/nolongfriendly**
Does not format blobs.

**/dumpallblob**
Dumps the blob in a default byte-by-byte format if there is not a friendly formatted interpretation available for it.

**/nodumpallblob**
Does not show the blob if there is not a formatted interpretation available.

**Note**

- A blob in this context means an attribute that is not a simple type, like a string or an integer. A blob is a complex structured type that is stored as binary bytes. In order to make sense of the blob, a program must interpret it and format it. A friendly blob is a blob that the program knows about and can format in an understandable way. The program has a list of blobs that it understands.

#### Repadmin showcert

Displays the server certificates loaded on a specified domain controller.

**repadmin /showcert** *DC_LIST*

#### Parameters

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax. This is a required parameter.

#### Repadmin showchanges

Displays changes from a specified directory partition or changes to a specified object. Syntax 1 saves changes to a directory partition. If this information is saved to a file the **getchanges** operation can be run again for comparison. Syntax 2 lists changes to a specified object.

#### Syntax

**repadmin/showchanges** *SourceDC NamingContext* [**/cookie:***File*] [**/atts:***attribute,attribute,...*]

#### Syntax

**repadmin /showchanges** *DestDC SourceDCObjectGUID NamingContext* [**/verbose**] [**/statistics**] [**/noincremental**]

[**/objectsecurity**] [**/ancestors**] [**/atts:** *attribute1,attribute2,...*] [**/filter:** *ldap filter*]

**Parameters**

***DestDC***
Specifies the host name of the destination domain controller from which to enumerate the host domain controllers.

***SourceDC***
Specifies the host name of the domain controller that hosts the directory partition whose changes you want to view.

***NamingContext***
Specifies the distinguished name of the directory partition.

**/cookie:** ***File***
Specifies a name for the file to which list changes are saved.

**/atts:** ***attribute1attributes...***
Returns only the attributes specified. Separate each listed attribute with a comma.

***SourceDCObjectGUID***
Specifies the unique hexadecimal number that identifies the object whose changes will be listed. The objectGUID can be retrieved by using the **showrepl** operation.

**/verbose**
Lists detailed information.

**/statistics**
Displays a summary of information about changes instead of a list of individual changes.

**/noincremental**
Returns changes in value change format, which lists current values for attributes as well what attributes have been added or deleted. If not specified, changes are returned in attribute change format, which shows only the current value of the attribute.

**/objectsecurity**
Overrides the need for the Get Changes right to the directory partition. By default this right is needed to run the *GetChanges* parameter. However, only changes that the currently logged on user has the rights to view are displayed.

**/ancestors**
Returns changes in USN order.

**/filter:** ***ldap filter***
Returns only those changes that meet the filter requirements.

**Note**

- The information from Syntax 1 can be saved to a file for later comparison.

**Repadmin showconn**

Displays the connection objects for a specified domain controller. Default is local site.

**repadmin /showconn** [*DC_LIST*] [*ServerRDN* | *ContainerDN* | *DC_GUID*] [**/From:***ServerRDN*] [**/intersite**]

**Parameters**

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

***ServerRDN***
Specifies the RDN of a server.

***ContainerDN***
Specifies the distinguished name of a container.

***DC_GUID***
Specifies the unique hexadecimal number that identifies the domain controller. The GUID can be retrieved by using the **showrepl** operation.

**/From:** ***ServerRDN***

Displays only the connection objects that exist to the specified server. By default all connection objects are displayed.

**/intersite**
Displays only those connection objects that are between sites.

#### Repadmin showctx

Displays a list of computers that have opened sessions with a specified domain controller.

**repadmin /showctx** [*DC_LIST*] [**/nocache**]

**Parameters**

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

**/nocache**
Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings.

#### Repadmin showism

Displays intersite messaging routes calculated by the Intersite Messaging Service. This operation cannot be executed remotely.

**repadmin /showism** [*TransportDN*] [**/verbose**]

**Parameters**

***TransportDN***
Specifies whether the mail server is using SMTP or RPC to send messages.

**/verbose**
Lists detailed information.

**Note**

- The **repadmin.exe/showism** cannot be executed against a remote server.

#### Repadmin showmsg

Displays the error message for a given error number.

**repadmin /showmsg** {*Win32Error* | *DSEventID* **/NTDSMSG**}

**Parameters**

***Win32Error***
Returns a short description of the given Win32 error code.

***DSEventID* /NTDSMSG**
Returns the actual event log text for the specified event ID.

#### Repadmin showncsig

Each domain controller has a naming context signature list. **repadmin.exe/showncsig** displays a list of the removed application partition GUIDs. An application partition can be configured to be held or not held on a particular domain controller using **ntdsutil**.

**repadmin /showncsig** [*DC_LIST*]

**Parameters**

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

#### Repadmin showobjmeta

Displays the replication metadata for a specified object stored in Active Directory, such as attribute ID, version number, originating and local Update Sequence Number (USN), and originating server's GUID and Date and Time stamp. By comparing the replication metadata for the same object on different domain controllers, an administrator can determine whether replication has taken place.

**repadmin /showobjmeta** [*DC_LIST*] *ObjectDN* [**/nocache**] [**/linked**]

**Parameters**

***DC_LIST***

Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

***ObjectDN***
Specifies the distinguished name of the object.

**/nocache**
Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings.

**/linked**
Displays metadata associated with, but not stored with, the specified object.

#### Repadmin showoutcalls

Displays calls that have not yet been answered, made by the specified server to other servers.

▸ **repadmin /showoutcalls** [*DC_LIST*]

#### Parameters

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

#### Repadmin showproxy

Lists cross domain move proxy objects. When an object is moved to another domain, a marker is left in the old domain indicating that the object used to be there. This is called the proxy.

#### Syntax

▸ **repadmin /showproxy** [*DC_LIST*] [*NamingContext*] [**matchstring**]

#### Syntax

▸ **repadmin /showproxy** [*DC_LIST*] [*ObjectDN*] [**matchstring**] **/movedobject**

#### Parameters

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

***NamingContext***
Specifies the distinguished name of the directory partition on the source domain controller.

***ObjectDN***
Specifies the distinguished name of the object.

**matchstring**
Specifies a filter for the output. Type a string of characters that must be present in the distinguished name in order to display the object.

***movedobject***
Displays a history of information from the original domain on a moved object after it has reached the new domain.

#### Repadmin showrepl

Displays the replication partners for each directory partition on the specified domain controller. Helps the administrator build a visual representation of the replication topology and see the role of each domain controller in the replication process.

▸ **repadmin /showrepl** [*DC_LIST*] [*SourceDCObjectGUID*] [*NamingContext*] [**/verbose**] [**/nocache**] [**/repsto**] [**/conn**] [**/csv**] [**/all**] [**/errorsonly**] [**/intersite**]

#### Parameters

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

***SourceDCObjectGUID***
Specifies the unique hexadecimal number that identifies the object whose replication events will be listed.

***NamingContext***
Specifies the distinguished name of the directory partition.

**/verbose**
Lists detailed information.

**/nocache**
Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings.

**/repsto**
Lists the domain controllers that pull replication information from the specified directory partition. To see the outbound neighbors, specify **/repsto** or **/all**.

**/conn**
Displays the connection objects associated with each link.

**/csv**
Displays the output of the **repadmin showrepl** operation in a Comma Separated Variable (CSV) format for viewing and analysis in Microsoft Excel. Repadmin supports redirection of screen output to a file.

**/all**
Displays all replication partners.

**/errorsonly**
Only shows the partnership if it has an error associated with it.

**/intersite**
Only shows this partnership if the source server is in a different site than our own.

#### Repadmin showsig

Displays the retired invocation IDs on a domain controller. A domain controller changes its invocation ID on being restored or when re-hosting an application partition.

**repadmin /showsig** [*DC_LIST*]

#### Parameters

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

#### Repadmin showtime

Converts a directory service time value to string format for both the local and the UTC time zones.

**repadmin /showtime** [*DSTimeValue*]

#### Parameters

***DSTimeValue***
Specifies the time value that needs to be converted.

### Note

- With parameters omitted, **repadmin /showtime** displays the current system time in both the directory service format and string format.

#### Repadmin showtrust

Lists all domains trusted by a specified domain.

**repadmin /showtrust** [*DC_LIST*]

#### Parameters

***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

#### Repadmin showutdvec

Displays the highest Update Sequence Number (USN) for the specified domain controller. This information shows how up-to-date a replica is with its replication partners.

**repadmin /showutdvec** *DC_LIST NamingContext* [**/nocache**] [**/latency**]

**Parameters**

***DC_LIST***

Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

***NamingContext***

Specifies the distinguished name of the directory partition.

**/nocache**

Specifies that GUIDs are left in hexadecimal form. By default GUIDs are translated into strings.

**/latency**

Sorts the information by the time required to complete the replication. By default the information is sorted by USN.

**Repadmin showvalue**

Displays the values of the type, last modified time, originating DC, and distinguished name of a specified object.

**repadmin /showvalue** [*DC_LIST*] *ObjectDN* [*AttributeName*] [*ValueDN*] [**/nocache**]

**Parameters**

***DC_LIST***

Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

***ObjectDN***

Specifies the distinguished name of the object.

***AttributeName***

Specifies a single attribute whose value you want to display.

***ValueDN***

Specifies the distinguished name of the attribute that is displayed.

**/nocache**

Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings.

**Repadmin siteoptions**

Changes default Knowledge Consistency Checker (KCC) behavior in a site.

Syntax for **/siteoptions** is available by typing:

**repadmin /experthelp**

| **Important** |
| --- |
| Setting these options might impair or disable the ability of the KCC to calculate the topology automatically. These options are intended for use by advanced administrators for special situations. If you set these options, document the settings carefully to ensure that the KCC functions as expected and can be restored to default functionality if necessary. |

**repadmin /siteoptions** [*DC*] [**/site:***SiteName*]

[{+|-}IS_AUTO_TOPOLOGY_DISABLED]

[{+|-}IS_TOPL_CLEANUP_DISABLED] [{+|-}IS_TOPL_MIN_HOPS_DISABLED]

[{+|-}IS_TOPL_DETECT_STALE_DISABLED]

[{+|-}IS_INTER_SITE_AUTO_TOPOLOGY_DISABLED]

[{+|-}IS_GROUP_CACHING_ENABLED] [{+|-}FORCE_KCC_WHISTLER_BEHAVIOR] [{+|-}FORCE_KCC_W2K_ELECTION] [{+|-}IS_RAND_BH_SELECTION_DISABLED]

[{+|-}IS_SCHEDULE_HASHING_ENABLED]

[{+|-}IS_REDUNDANT_SERVER_TOPOLOGY_ENABLED]

{+|-}W2K3_IGNORE_SCHEDULES {+|-}W2K3_BRIDGES_REQUIRED

**Parameters**

***DC***

Specifies the host name of the domain controller on which to set the site options.

***SiteName***

The site that the option setting affects.

**{+|-}IS_AUTO_TOPOLOGY_DISABLED**

When enabled, prevents the KCC from generating intrasite connections on all servers in the site. Disable this option if you use manual connections and do not want the KCC to build connections automatically.

**{+|-}IS_TOPL_CLEANUP_DISABLED**

When enabled, prevents the KCC from removing connection objects that it does not need. Disable this option if you want to take responsibility for removing old redundant connections. Alternatively, to control or augment the topology, you can use manual connections, which the KCC does not delete.

**{+|-}IS_TOPL_MIN_HOPS_DISABLED**

When enabled, prevents the KCC from generating optimizing connections in the ring topology of intrasite replication. Optimizing connections reduce the replication latency in the site and disabling them is not recommended.

**{+|-}IS_TOPL_DETECT_STALE_DISABLED**

Prevents the KCC from excluding servers that are unreachable from the topology; that is, the KCC does use an alternate server to reroute replication. Use this option only if network communication is very unstable and brief outages are expected.

**{+|-}IS_INTER_SITE_AUTO_TOPOLOGY_DISABLED**

Prevents the KCC that functions as the intersite topology generator (ISTG) from generating connections for intersite replication. Use this option when you want to create manual intersite connections (disable the ISTG) but retain the KCC to generate intrasite connections.

**{+|-}IS_GROUP_CACHING_ENABLED**

Disables Universal Group Membership Caching in the site.

**{+|-}FORCE_KCC_WHISTLER_BEHAVIOR**

Implements KCC operation that is consistent with Windows Server 2003 forest functional level. This option can be set if all domain controllers in the site are running Windows Server 2003.

**{+|-}FORCE_KCC_W2K_ELECTION**

Off by default. When enabled, implements the Windows 2000 Server method of ISTG selection.

**{+|-}IS_RAND_BH_SELECTION_DISABLED**

Implements the Windows 2000 Server method of selecting a single bridgehead server per directory partition and transport.

**{+|-}IS_SCHEDULE_HASHING_ENABLED**

Spreads replication start times randomly across the entire schedule interval rather than just the first quarter of the interval.

**{+|-}IS_REDUNDANT_SERVER_TOPOLOGY_ENABLED**

Creates redundant connections between sites before a failure takes place. When enabled, disables KCC failover. Requires that automatic detection of failed connections also be disabled (+IS_TOPL_DETECT_STALE_DISABLED).

**{+|-}W2K3_IGNORE_SCHEDULES**

When the forest functional level Windows Server 2003 or Windows Server 2003 interim is in effect, provides KCC control of the ability to ignore schedules (replication occurs at the designated intervals and is always available).

**{+|-}W2K3_BRIDGES_REQUIRED**

When the forest functional level Windows Server 2003 or Windows Server 2003 interim is in effect, provides KCC control of the ability to enable or disable site link bridging.

For more information about using site options, see Planning Chapter 3, "Planning the Physical Structure for a Branch Office Deployment" in the Windows Server 2003 Active Directory Branch Office Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=28523).

**Repadmin syncall**

Synchronizes a specified domain controller with all replication partners.

▸ **repadmin /syncall** *DC* [*NamingContext*] [*Flags*]

**Parameters**

***DC***
Specifies the host name of the domain controller to synchronize with all replication partners.

***NamingContext***
Specifies the distinguished name of the directory partition.

***Flags***
Performs specific actions during the replication.

| Flag | Description |
|---|---|
| **/a** | Abort if any server is unavailable. |
| **/A** | Sync all naming contexts that are held on the home server. |
| **/d** | Identify servers by distinguished name in messages. |
| **/e** | Enterprise, cross sites. |
| **/h** | Print this help screen. |
| **/i** | Iterate indefinitely. |
| **/l** | Perform showrepl on each server pair in path instead of syncing. |
| **/j** | Synchronize adjacent servers only. |
| **/p** | Pause for possible user abort after every message. |
| **/P** | Push changes outward from home server. |
| **/q** | Run in quiet mode, suppress call back messages. |
| **/Q** | Run in very quiet mode, report fatal errors only. |
| **/s** | Do not synchronize. |
| **/S** | Skip initial server response check. |

**Repadmin viewlist**
Displays a list of domain controllers.

**repadmin /viewlist** [*DC_LIST*] [*OBJ_LIST*]

**Parameters**
***DC_LIST***
Specifies the host name of a domain controller or a list of domain controllers separated by a space that the object will be replicated to. See above for detailed syntax.

***OBJ_LIST***
this parameter takes a DN or a special keyword that expands into a DN. The keywords are:

- **Ncobj:config:**
- **Ncobj:schema:**
- **Ncobj:domain:**, This key word is the domain partition DN of the home server.
- **Dsaobj:**, This keyword is the NTDS settings object of the home server

**Repadmin oldhelp**
This displays a list of the operations that have been deprecated in this version of Repadmin.

**See Also**

**Concepts**

Repadmin Overview
Repadmin Remarks
Repadmin Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Replmon Overview

Updated: April 18, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Replmon graphical user interface (GUI) tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

The Replmon graphical user interface tool was removed from Windows Server 2008 and later. Repadmin is still available for troubleshooting replication. For more information on using Repadmin, see Monitoring and Troubleshooting Active Directory Replication Using Repadmin and Getting Over Replmon.

**Replmon.exe: Active Directory Replication Monitor**
This GUI tool enables administrators to view the low-level status of Active Directory replication, force synchronization between domain controllers, view the topology in a graphical format, and monitor the status and performance of domain controller replication.

You can use ReplMon to do the following:

- See when a replication partner fails.
- View the history of successful and failed replication changes for troubleshooting purposes.
- View the properties of directory replication partners.
- Create your own applications or scripts written in Microsoft Visual Basic Scripting Edition (VBScript) to extract specific data from Active Directory.
- View a snapshot of the performance counters on the computer, and the registry configuration of the server.
- Generate status reports that include direct and transitive replication partners, and detail a record of changes.
- Find all direct and transitive replication partners on the network.
- Display replication topology.
- Poll replication partners and generate individual histories of successful and failed replication events.
- Force replication.
- Trigger the Knowledge Consistency Checker (KCC) to recalculate the replication topology.
- Display changes that have not yet replicated from a given replication partner.
- Display a list of the trust relationships maintained by the domain controller being monitored.
- Display the metadata of an Active Directory object's attributes.
- Monitor replication status of domain controllers from multiple forests.

**Corresponding UI**
ReplMon provides its own user interface. See Replmon UI for more information.

**Concepts**

For more information about deploying and using Active Directory, see the Active Directory Overview Web site.

**System Requirements**
The following are the system requirements for ReplMon:

- ReplMon must be installed on a computer running Windows Server 2003. The computer can be a domain controller, member server, member workstation, or standalone computer.

**Files Required**

- Comctl32.ocx
- Comdlg32.ocx
- Ctl3d32.dll
- Iadstools.dll — component DLL that hosts the functions used by ReplMon
- Replmon.exe
- Tabctl32.ocx

### For More Information

Installing ReplMon in a directory other than the default might result in errors. For more information, see "Intalling Replmon" in Replmon Remarks. For information about the IADsTools, see IADsTools Documentation (iadstools.doc) in the Program Files\Support Tools directory.

### Note

- You need Microsoft Word or a compatible viewer to view Iadstools.doc. You can download a free version of Microsoft Word Viewer from the Microsoft Office Download Center.

**See Also**
**Concepts**
Replmon Remarks
Replmon UI
Replmon Examples
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Replmon Remarks

Updated: April 18, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ReplMon Remarks**
The Replmon graphical user interface tool was removed from Windows Server 2008 and later. Repadmin is still available for troubleshooting replication. For more information on using Repadmin, see Monitoring and Troubleshooting Active Directory Replication Using Repadmin and Getting Over Replmon.

Installing ReplMon

If you choose to install Support Tools in a directory other than the default (\Program Files\Support Tools), and if you install Support Tools in a directory that contains extended characters, the required Visual Basic controls might fail to register. If this occurs, the work-around is to remove Support Tools, and then to install them in the \Program Files\Support Tools folder.

Context-Sensitive Menus

ReplMon is designed for use with context-sensitive menus. Replmon UI and Replmon Examples specify the menu options as the method for accessing commands. For a shortcut, right-click objects in the ReplMon window to display the context-sensitive menus.

Creating a List of Monitored Servers

The list of monitored servers in ReplMon can be created by adding each server individually. If a number of servers must be monitored on a regular basis, however, the task of building the list can be automated by creating a file that lists the servers to be monitored. The file is created either by using a text editor, or in ReplMon by saving a list of servers already being monitored. The file is then opened in ReplMon.

**To create a list of monitored servers and save it as a text file**

1. Open a text editor.

2. Type a list of host names for the servers. Each server must be on its own line.

3. If alternative credentials are needed, add a comma after *ServerName*, followed by *Domain**User*. If *ServerName* is Server1, *Domain* is Reskit, and *User* is Admin, the command looks like the following:

   ```
   server1,reskit\admin
   ```

   For security reasons, the password is not included. When the text file is read by ReplMon, a dialog box prompts the user for the password.

4. Save the file with an .ini extension in the \Program Files\Active Directory Replication Monitor folder.

Another way to create a text file of servers is by using ReplMon. If ReplMon is open and monitoring servers, the current list of monitored servers can be saved for future use.

**To create a list of monitored servers from ReplMon**

1. While monitoring a list of servers in ReplMon, on the **File** menu, click **Save Monitored List As**.

2. Save the file with an .ini extension in the \Program Files\Active Directory Replication Monitor folder.

Populating the Monitored Servers List in ReplMon

**To load an .ini file to create a list of monitored servers for ReplMon**

1. With ReplMon open, on the **File** menu, click **Open Script**.

A list of .ini files from \Program Files\Active Directory Replication Monitor is displayed.

2. Select the .ini file that contains the list of servers to be monitored.

ReplMon and RepAdmin Similarities

Many of the tasks that can be performed with ReplMon can be performed from the command line by using Repadmin Overview.

**See Also**
**Concepts**
Replmon Overview
Replmon UI
Replmon Examples
Alphabetical List of Tools
Search Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Replmon Examples

Updated: April 18, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ReplMon Examples**
The Replmon graphical user interface tool was removed from Windows Server 2008 and later. Repadmin is still available for troubleshooting replication. For more information on using Repadmin, see Monitoring and Troubleshooting Active Directory Replication Using Repadmin and Getting Over Replmon.

Example 1: Generate a status report for a server

In this example, you use ReplMon to create a log file that records replication events for an individual server. This data is used to troubleshoot possible replication errors in the forest or domain.

**To create the log file**

1. In ReplMon, select the server in the **Monitored Servers** section of the window.
2. On the **Action** menu, select **Server**, and then click **Generate Status Report**. The **Save As** dialog box appears.
3. In the **Save As** dialog box, type a name for the log file to generate, (optionally) select a new path for the file, and then click **Save**.

   The **Report Options** dialog box appears.
4. In the **Report Options** dialog box, select the type of information to be logged.
5. To save the log file and return to ReplMon, click **OK**.

**Notes**

- If a log file is given the same name as an existing log file, ReplMon replaces the existing log file without a warning. To avoid this, incorporate the current date in your naming scheme.
- To set the default path, on the **View** menu, select **Options**, and then use the **General** tab of the **Options** dialog box.

Example 2: Display the replication topology for a server

In this example, you use ReplMon to display replication partners of a server, which graphically represents the replication topology.

**To display replication partners**

1. In ReplMon, select the server in the **Monitored Servers** section of the window.
2. On the **Action** menu, select **Server**, and then click **Show Replication Topologies**.

   The **View Replication Topology** window appears.
3. In the **View Replication Topology** window, on the **View** menu, click **Replication Topologies**.
4. In ReplMon on the **Options** menu, click **Report**.

Example 3: Display pending replication changes

Each domain controller keeps a record of the last changes received from other domain controllers by using the Update Sequence Number (USN) of the replication partner. In this example changes have occurred on the replication partner and the

Active Directory replication process has not propagated those changes to the monitored server. You can use ReplMon to show what objects have changed and therefore need to be replicated to the monitored server. This is done on a per-directory partition basis.

**To show what objects have changed**

1. In ReplMon, expand a server to display the list of directory partitions.

2. Expand and select the replication partner within the directory partition for which you want to view unreplicated changes. Only objects that have changed within the directory partition are displayed.

3. On the **Action** menu, click **Replication Partner**, and then select **Check Current USN and Unreplicated Objects**.

   A separate window opens. Detected changes are displayed and can be saved to a text file. If no changes are detected, a message box informs you of the current USN on the replication partner and that all changes have replicated to the monitored server.

Example 4: Synchronize directory partitions

In this example, you use ReplMon to force a replication event between two directory partitions. By default, directory partitions on servers periodically synchronize with each other. Manual synchronization is necessary only if the network or a server is down for an extended period of time.

**To manually synchronize two directory partitions**

1. In the **Monitored Servers** pane of ReplMon, select the server that needs to be brought up to date.

2. Expand a directory partition.

   A list of replication partners for that directory partition appears.

3. Select the directory partition that needs to be synchronized.

4. On the **Action** menu, click **Synchronize with this Replication Partner**.

5. In the dialog box that appears, select any additional parameters for synchronization, then click **OK** to continue with the synchronization.

**Notes**

- This method works only with direct replication partners.

- This method synchronizes only one directory partition at a time. To synchronize all directory partitions at once, on the **Action** menu, click **Server**, and then click **Synchronize Each Directory Partition with All Servers**.

**See Also**
**Concepts**
Replmon Overview
Replmon Remarks
Replmon UI
Alphabetical List of Tools
Search Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Replmon UI

Updated: April 18, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Replmon graphical user interface tool was removed from Windows Server 2008 and later. Repadmin is still available for troubleshooting replication. For more information on using Repadmin, see Monitoring and Troubleshooting Active Directory Replication Using Repadmin and Getting Over Replmon.

**Replmon UI**

ReplMon has the following menus:

File Menu

**File** menu options:

| Option | Description |
|---|---|
| **Open Script** | Opens a dialog box that populates the Active Directory Replication Monitor with monitored servers. These files have an .ini extension and can be created by using a text editor or by using the **Save Monitored List As** command in the **File** menu. |
| **Save Monitored List As** | Opens a dialog box that saves a list of monitored servers in an .ini file. This list can be opened later by using the **Open Script** command on the **File** menu. If the current user does not have permission to monitor a server listed, then a user name must be supplied after *ServerName*, using the *Domain**User* format. Do not supply the password in the .ini file. The password is requested as each server is added to the monitored list when the .ini file is read. |
| **Open Log** | Opens a previously created .log file. Log files are created by default for every monitored server, and statistics are tracked on a per-replication partner and per-directory partition basis. |
| **Exit** | Closes ReplMon and discontinues monitoring servers. |

Edit Menu

**Edit** menu options:

| Option | Description |
|---|---|
| **Add Monitored Server** | Opens a dialog box that adds a single server to the list of servers monitored by ReplMon. By default, ReplMon does not monitor any servers. If you know the name of the server, click **Add the server explicitly by name**, and then type the name of the server. If you need to browse for the server, click **Search the directory for the server to add**, and then click **Next**. The entire forest can be searched if necessary. Enter a user name and password by selecting the **Use Alternate Credentials** checkbox. |

Action Menu

The **Action** menu includes the **Domain**, **Server**, **Site**, **Naming Context**, and **Replication Partner** submenus.

**Domain** submenu option:

| Option | Description |
|---|---|
| **Domain** | Opens the **Search Domain Controllers for Replication Errors** submenu option. |

**Server** submenu options:

| Option | Description |
|---|---|
| **Update Status (only for this server)** | Triggers an immediate status check for the selected server. The time of the updated status is logged and displayed. |

| | |
|---|---|
| **Check Replication Topology** | Triggers the Knowledge Consistency Checker (KCC) to recalculate the replication topology for the selected server. |
| **Synchronize Each Directory Partition with All Servers** | Triggers replication for the selected server's directory partitions with each partition's replication partners. |
| **Generate Status Report** | Creates and saves a status report for the selected server. When you select this option, the **Save As** dialog box appears. In the **Save As** dialog box, you can name the status report. To select the information that should be on the status report, select **Report Options**. |
| **Show Domain Controllers in Domain** | Opens a dialog box that lists all of the domain controllers that belong to the selected server's domain. To close this dialog box, click **OK**. |
| **Show Replication Topologies** | Creates a graphical view of all replication partners for all directory partitions. |
| **Show Group Policy Object Status** | Lists Group Policy objects for a specified domain controller and the object's current version number. Also lists version numbers of the Group Policy object in SysVol. |
| **Show Current Performance Data** | Displays real-time performance monitor counters that are specific to Active Directory replication. |
| **Show Global Catalog Servers in Enterprise** | Opens a dialog box that lists a list of Global Catalog servers in the selected server's forest. To close this dialog box, click **OK**. |
| **Show Bridgehead Servers** | Displays a list of Bridgehead servers for the selected server's site or enterprise. To close this dialog box, click **OK**. |
| **Show Trust Relationships** | Displays a list of trust relationships for the selected server's enterprise. To close this dialog box, click **OK**. |
| **Show Attribute Meta-Data for Active Directory Object** | Displays the attribute metadata for a specified object. Enter the Distinguished Name (DN) of the object in the **Object** dialog box and then click **OK**. |
| **Clear Log** | Clears the contents of the log file generated by ReplMon. |
| **Delete** | Removes the selected server from the list of monitored servers in ReplMon. |
| **Properties** | Displays properties of the selected server. |

**Site** submenu option:

| Option | Description |
|---|---|
| **Site** | Starts the Add Monitored Server wizard. |

**Naming Context** submenu options:

| Option | Description |
|---|---|
| **Synchronize This Directory Partition with All Servers** | Synchronizes a selected directory partition with all domain controllers on the site. |
| **Show Change Notifications To Replication Partners** | Lists the replication messages sent to direct replication partners. |

**Replication Partner** submenu options:

| Option | Description |
|---|---|

| | |
|---|---|
| **Check Current USN and Unreplicated objects** | Displays information on a selected server's replication partner. Displays unreplicated changes on selected replication partners. |
| **Synchronize Each Directory Partition with All Servers** | Manually synchronizes all of the directory partitions on a server with its replication partners. To synchronize an individual replication partner instead, right-click that replication partner and select **Synchronize with this Replication Partner**. |

View Menu

**View** menu options:

| Option | Description |
|---|---|
| **Details View** | Opens a window that shows a log of server status updates. |
| **Refresh** | Updates all listed servers. Refreshes every five minutes. To change the refresh interval, click the **Update Automatically** button and enter a value between 1 and 60. |

**Options** dialog box **General** tab:

| Option | Description |
|---|---|
| **Show Retired Replication Partners** | Shows servers that are no longer replication partners in the graphical replication topology display. |
| **Show Transitive Replication and Extended Data** | When checked, logs replication events from transitive partners. (By default, only replication events from direct replication partners are entered in the log.) |
| **Notify when replication fails after this number of attempts:** | Enter a number of attempts after which ReplMon notes that replication is failing. Specify whether to enter this event in the Application log or to send an e-mail message. |
| **Auto Generate File Names for Log Files** | By default, ReplMon automatically creates a name for log files. |
| **Default Path For Replication Status Logs** | Specify a path to which ReplMon automatically saves log files. |
| **Enable Debug Logging** | Creates a separate log that can be used to troubleshoot errors in ReplMon. |

**Options** dialog box **Status Logging** tab:

| Option | Description |
|---|---|
| **Group Policy Objects** | Enables Group Policy object logging. |
| **Performance Statistics** | Enables the addition of performance objects and counters to the ReplMon log. |
| **Display Changed Attributes When Replication Occurs** | Records in the ReplMon log each attribute changed, and the Update Sequence Number (USN) associated with the change. |

**Options** dialog box **Cache** tab:

| Option | Description |
|---|---|
| **Reset Button** | Flushes the ReplMon cache. |

**See Also**
**Concepts**
Replmon Overview
Replmon Remarks
Replmon Examples
Alphabetical List of Tools

Search Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Sdcheck Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Sdcheck command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Sdcheck.exe: Security Descriptor Check Utility**
This command-line tool displays the security descriptor for any object stored in Active Directory. The security descriptor contains the access control lists (ACLs) defining the permissions that users have on objects stored in Active Directory.

### Notes

- You must run SDCheck from a command window.
- To enable administrators to determine the effective access controls on an object, SDCheck also displays the object hierarchy and any ACLs that are inherited by the object from its parent.
- As changes are made to the ACLs of an object or its parent, these changes are propagated automatically by Active Directory. SDCheck displays the security descriptor propagation metadata, so that administrators can monitor these changes with respect to propagation of inherited ACLs, as well as replication of ACLs from other domain controllers.
- As a complement to the replication monitoring tools (Repadmin.exe and Replmon.exe), SDCheck can be used to ensure that domain controllers are up-to-date with one another.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Every container and object on the network or in Active Directory has a set of access control information attached to it, known as a security descriptor. A security descriptor contains the security information associated with a securable object. In Windows Server 2003 operating systems, the security descriptor is automatically created when an object is created.

**System Requirements**
The following are the system requirements for SDCheck:

- Windows Server 2003 or Windows XP Professional

**File Required**

- Sdcheck.exe

**See Also**
**Concepts**
Sdcheck Remarks
Sdcheck Syntax
Sdcheck Examples
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview

Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Sdcheck Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SDCheck Remarks**
**Security Descriptors**
The security descriptor is the structure of the list of permissions for an object.

The descriptor contains a header that includes a revision number and control flags. It also contains the Owner's SID (Security Identifier), and the SID of the object's primary group. A security descriptor can also contain two other lists: A discretionary access control list (DACL) and a system access control list (SACL).

The DACL is used to determine what access to an object is allowed. It contains Access Control Entries (ACEs) that define which SIDs have what privileges. An ACE has three parts: header, SID, and mask. The header describes whether it allows or denies access. The SID is the security identifier of the group or user that the ACE pertains to. The mask tells what access type the ACE is referring to.

When an access request is received, the system scans the DACL to see if any of the SIDs match any of the SIDs of the requester. If no matches are found or a Deny is found, then access is denied. All the SIDs are scanned, because a Deny Access always has a higher priority than any Allow. The only exception to this rule is that the owner of an object always has full control of the object.

An SACL is similar to a DACL, except that it is used for auditing purposes. If the user's SID is listed in the SACL's ACE, the audit event will be recorded in the event log.

**See Also**
**Concepts**
Sdcheck Overview
Sdcheck Syntax
Sdcheck Examples
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Sdcheck Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SDCheck Examples**

**Example 1: Display the Summary Security Information for an Object**

In this example, no changes have been applied to the user object (Someone), or the parent organization unit (Sales), since their initial creation. Note the metadata version number.

Type the following at the command line:

```
sdcheck sprocket someone@example.com
```

The following output appears:

```
Microsoft(R) Windows (R) 2000 Operating System
Security Descriptor Check Utility - build(2015)

Input:  someone@example.com
Object: CN=Someone,OU=Sales,DC=example,DC=com
Domain: example.com
Domain: DC=example,DC=com
Server: sprocket

*** Warning: No values returned for dSCorePropagationData on
CN=Someone,OU=Sales,DC=example,DC=com
*** Warning: No values returned for dSCorePropagationData on OU=Sales,DC=example,DC=com
*** Warning: No values returned for dSCorePropagationData on DC=example,DC=com

Object:   CN=Someone,OU=Sales,DC=example,DC=com
Classes:  top person organizationalPerson user
SD:       1012 bytes
Metadata: 04/15/1999 14:53:04 @ SPROCKET.example.com ver: 1

  Object:   OU=Sales,DC=example,DC=com
  Classes:  top organizationalUnit
  SD:       424 bytes
  Metadata: 04/15/1999 14:54:13 @ SPROCKET.example.com ver: 1

    Object:   DC=example,DC=com
    Classes:  top domain domainDNS
    SD:       496 bytes
    Metadata: 04/15/1999 14:51:32 @ SPROCKET.example.com ver: 4

Checking ACL inheritance ...
Parent: 2 - DC=example,DC=com
Child:  1 - OU=Sales,DC=example,DC=com

Checking ACL inheritance ...
Parent: 1 - OU=Sales,DC=example,DC=com
Child:  0 - CN=Someone,OU=Sales,DC=example,DC=com
```

**Example 2: Display the Security Descriptor for an Object**

In this example, two additional access control lists (ACLs) have been added. This can be identified by the version number increase for the security metadata. One ACL denies Read/Write access to members of the Accounts Payable group, and the other ACL audits Read/Write attempts by the Administrator.

Type the following at the command line:

```
sdcheck sprocket someone@example.com -dumpSD
```

Expand to see SDCheck output

```
Microsoft(R) Windows (R) 2000 Operating System
Security Descriptor Check Utility - build(2015)

Input:  someone@example.com
Object: CN=Someone,OU=Sales,DC=example,DC=com
```

```
Domain: example.com
Domain: DC=example,DC=com
Server: sprocket

*** Warning: No values returned for dSCorePropagationData on OU=Sales,DC=example,DC=com
*** Warning: No values returned for dSCorePropagationData on DC=example,DC=com

Object:   CN=Someone,OU=Sales,DC=example,DC=com
Classes:  top person organizationalPerson user
SD:       1072 bytes
Metadata: 04/15/1999 14:59:08 @ SPROCKET.example.com ver: 3
History:  04/15/1999 14:59:00 flags(0x1) SD propagation
          04/15/1999 14:59:08 flags(0x1) SD propagation

  Object:   OU=Sales,DC=example,DC=com
  Classes:  top organizationalUnit
  SD:       424 bytes
  Metadata: 04/15/1999 14:54:13 @ SPROCKET.example.com ver: 1

    Object:   DC=example,DC=com
    Classes:  top domain domainDNS
    SD:       496 bytes
    Metadata: 04/15/1999 14:51:32 @ SPROCKET.example.com ver: 4

Checking ACL inheritance ...
Parent: 2 - DC=example,DC=com
Child:  1 - OU=Sales,DC=example,DC=com

Checking ACL inheritance ...
Parent: 1 - OU=Sales,DC=example,DC=com
Child:  0 - CN=Someone,OU=Sales,DC=example,DC=com


SD for CN=Someone,OU=Sales,DC=example,DC=com
SD Revision: 1
SD Control:  0x8c14
SE_DACL_PRESENT
SE_SACL_PRESENT
SE_DACL_AUTO_INHERITED
SE_SACL_AUTO_INHERITED
SE_SELF_RELATIVE
Owner: microsoft\Domain Admins S-1-5-21-640924683-4221571012-3872390550-512
Group: microsoft\Domain Users S-1-5-21-640924683-4221571012-3872390550-513
DACL:
Revision      4
Size:         944 bytes
# Aces:       24
Ace[0]
Ace Type:  0x1 - ACCESS_DENIED_ACE_TYPE
Ace Size:  36 bytes
Ace Flags: 0x0
Ace Mask:  0x000200bc
READ_CONTROL
ACTRL_DS_LIST
ACTRL_DS_SELF
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
ACTRL_DS_LIST_OBJECT
Ace Sid:   microsoft\Accounts Payable S-1-5-21-640924683-4221571012-3872390550-1130
Ace[1]
Ace Type:  0x0 - ACCESS_ALLOWED_ACE_TYPE
Ace Size:  24 bytes
Ace Flags: 0x0
Ace Mask:  0x000f01ff
DELETE
READ_CONTROL
WRITE_DAC
WRITE_OWNER
ACTRL_DS_CREATE_CHILD
ACTRL_DS_DELETE_CHILD
ACTRL_DS_LIST
ACTRL_DS_SELF
```

```
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
ACTRL_DS_DELETE_TREE
ACTRL_DS_LIST_OBJECT
ACTRL_DS_CONTROL_ACCESS
Ace Sid:   BUILTIN\Account Operators S-1-5-32-548
Ace[2]
Ace Type:  0x0 - ACCESS_ALLOWED_ACE_TYPE
Ace Size:  20 bytes
Ace Flags: 0x0
Ace Mask:  0x00020000
READ_CONTROL
Ace Sid:   NT AUTHORITY\Authenticated Users S-1-5-11
Ace[3]
Ace Type:  0x0 - ACCESS_ALLOWED_ACE_TYPE
Ace Size:  36 bytes
Ace Flags: 0x0
Ace Mask:  0x000f01ff
DELETE
READ_CONTROL
WRITE_DAC
WRITE_OWNER
ACTRL_DS_CREATE_CHILD
ACTRL_DS_DELETE_CHILD
ACTRL_DS_LIST
ACTRL_DS_SELF
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
ACTRL_DS_DELETE_TREE
ACTRL_DS_LIST_OBJECT
ACTRL_DS_CONTROL_ACCESS
Ace Sid:   microsoft\Domain Admins S-1-5-21-640924683-4221571012-3872390550-512
Ace[4]
Ace Type:  0x0 - ACCESS_ALLOWED_ACE_TYPE
Ace Size:  20 bytes
Ace Flags: 0x0
Ace Mask:  0x00020094
READ_CONTROL
ACTRL_DS_LIST
ACTRL_DS_READ_PROP
ACTRL_DS_LIST_OBJECT
Ace Sid:   NT AUTHORITY\SELF S-1-5-10
Ace[5]
Ace Type:  0x0 - ACCESS_ALLOWED_ACE_TYPE
Ace Size:  20 bytes
Ace Flags: 0x0
Ace Mask:  0x000f01ff
DELETE
READ_CONTROL
WRITE_DAC
WRITE_OWNER
ACTRL_DS_CREATE_CHILD
ACTRL_DS_DELETE_CHILD
ACTRL_DS_LIST
ACTRL_DS_SELF
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
ACTRL_DS_DELETE_TREE
ACTRL_DS_LIST_OBJECT
ACTRL_DS_CONTROL_ACCESS
Ace Sid:   NT AUTHORITY\SYSTEM S-1-5-18
Ace[6]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000010
ACTRL_DS_READ_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr set - Public Information
Object Ace Sid:   NT AUTHORITY\Authenticated Users S-1-5-11
Ace[7]
```

```
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000010
ACTRL_DS_READ_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr set - Web Information
Object Ace Sid:   NT AUTHORITY\Authenticated Users S-1-5-11
Ace[8]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000010
ACTRL_DS_READ_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr set - Personal Information
Object Ace Sid:   NT AUTHORITY\Authenticated Users S-1-5-11
Ace[9]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000010
ACTRL_DS_READ_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr set - General Information
Object Ace Sid:   NT AUTHORITY\Authenticated Users S-1-5-11
Ace[10]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  56 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000030
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr - userCertificate
Object Ace Sid:   microsoft\Cert Publishers S-1-5-21-640924683-4221571012-3872390550-
517
Ace[11]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000100
ACTRL_DS_CONTROL_ACCESS
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Control right - Change Password
Object Ace Sid:   Everyone S-1-1-0
Ace[12]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  56 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000010
ACTRL_DS_READ_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr set - Logon Information
Object Ace Sid:   microsoft\RAS and IAS Servers S-1-5-21-640924683-4221571012-
3872390550-553
Ace[13]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  56 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000010
ACTRL_DS_READ_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Control right - Modify Group Membership
Object Ace Sid:   microsoft\RAS and IAS Servers S-1-5-21-640924683-4221571012-
```

3872390550-553
Ace[14]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  56 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000010
ACTRL_DS_READ_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr set - Account Restrictions
Object Ace Sid:   microsoft\RAS and IAS Servers S-1-5-21-640924683-4221571012-
3872390550-553
Ace[15]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  56 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000010
ACTRL_DS_READ_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr set - Modify Remote Access Information
Object Ace Sid:   microsoft\RAS and IAS Servers S-1-5-21-640924683-4221571012-
3872390550-553
Ace[16]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000030
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr set - Web Information
Object Ace Sid:   NT AUTHORITY\SELF S-1-5-10
Ace[17]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000030
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Control right - Phone and Mail Options
Object Ace Sid:   NT AUTHORITY\SELF S-1-5-10
Ace[18]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000030
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Attr set - Personal Information
Object Ace Sid:   NT AUTHORITY\SELF S-1-5-10
Ace[19]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000100
ACTRL_DS_CONTROL_ACCESS
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Control right - Receive As
Object Ace Sid:   NT AUTHORITY\SELF S-1-5-10
Ace[20]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000100
ACTRL_DS_CONTROL_ACCESS

```
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Control right - Send As
Object Ace Sid:   NT AUTHORITY\SELF S-1-5-10
Ace[21]
Ace Type:  0x5 - ACCESS_ALLOWED_OBJECT_ACE_TYPE
Ace Size:  40 bytes
Ace Flags: 0x0
Object Ace Mask:  0x00000100
ACTRL_DS_CONTROL_ACCESS
Object Ace Flags: 0x1
ACE_OBJECT_TYPE_PRESENT
Object Ace Type:  Control right - Change Password
Object Ace Sid:   NT AUTHORITY\SELF S-1-5-10
Ace[22]
Ace Type:  0x0 - ACCESS_ALLOWED_ACE_TYPE
Ace Size:  24 bytes
Ace Flags: 0x12
CONTAINER_INHERIT_ACE
INHERITED_ACE
Ace Mask:  0x000f01bd
DELETE
READ_CONTROL
WRITE_DAC
WRITE_OWNER
ACTRL_DS_CREATE_CHILD
ACTRL_DS_LIST
ACTRL_DS_SELF
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
ACTRL_DS_LIST_OBJECT
ACTRL_DS_CONTROL_ACCESS
Ace Sid:   BUILTIN\Administrators S-1-5-32-544
Ace[23]
Ace Type:  0x0 - ACCESS_ALLOWED_ACE_TYPE
Ace Size:  36 bytes
Ace Flags: 0x12
CONTAINER_INHERIT_ACE
INHERITED_ACE
Ace Mask:  0x000f01ff
DELETE
READ_CONTROL
WRITE_DAC
WRITE_OWNER
ACTRL_DS_CREATE_CHILD
ACTRL_DS_DELETE_CHILD
ACTRL_DS_LIST
ACTRL_DS_SELF
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
ACTRL_DS_DELETE_TREE
ACTRL_DS_LIST_OBJECT
ACTRL_DS_CONTROL_ACCESS
Ace Sid:   microsoft\Enterprise Admins S-1-5-21-640924683-4221571012-3872390550-519
SACL:
Revision      2
Size:         52 bytes
# Aces:       2
Ace[0]
Ace Type:  0x2 - SYSTEM_AUDIT_ACE_TYPE
Ace Size:  24 bytes
Ace Flags: 0x82
CONTAINER_INHERIT_ACE
Ace Mask:  0x00000030
ACTRL_DS_READ_PROP
ACTRL_DS_WRITE_PROP
Ace Sid:   BUILTIN\Administrators S-1-5-32-544
Ace[1]
Ace Type:  0x2 - SYSTEM_AUDIT_ACE_TYPE
Ace Size:  20 bytes
Ace Flags: 0xd2
CONTAINER_INHERIT_ACE
```

```
INHERITED_ACE
Ace Mask:  0x000d016b
DELETE
WRITE_DAC
WRITE_OWNER
ACTRL_DS_CREATE_CHILD
ACTRL_DS_DELETE_CHILD
ACTRL_DS_SELF
ACTRL_DS_WRITE_PROP
ACTRL_DS_DELETE_TREE
ACTRL_DS_CONTROL_ACCESS
Ace Sid:   Everyone S-1-1-0
```

**Example 3: Determine Whether the Security Descriptor is Being Inherited Correctly**

In this example, an inheritable ACL was added to the domain object dc=microsoft,dc=com, denying Read and Write access to the members of the Finance group. Note that the metadata version number for the domain object dc=microsoft,dc=com has been incremented; however, the ACL has yet to propagate to the Sales object.

Type the following at the command line:

```
sdcheck sprocket someone@example.com
```

The following output appears:

```
Microsoft(R) Windows (R) 2000 Operating System
Security Descriptor Check Utility - build(2015)

Input:  someone@example.com
Object: CN=Someone,OU=Sales,DC=example,DC=com
Domain: example.com
Domain: DC=example,DC=com
Server: sprocket

*** Warning: No values returned for dSCorePropagationData on DC=example,DC=com

Object:   CN=Someone,OU=Sales,DC=example,DC=com
Classes:  top person organizationalPerson user
SD:       1108 bytes
Metadata: 04/15/1999 14:59:08 @ SPROCKET.example.com ver: 3
History:  04/15/1999 14:59:00 flags(0x1) SD propagation
          04/15/1999 14:59:08 flags(0x1) SD propagation
          04/15/1999 15:13:22 flags(0x1) SD propagation

  Object:   OU=Sales,DC=example,DC=com
  Classes:  top organizationalUnit
  SD:       460 bytes
  Metadata: 04/15/1999 14:54:13 @ SPROCKET.example.com ver: 1
  History:  04/15/1999 15:13:21 flags(0x1) SD propagation

    Object:   DC=example,DC=com
    Classes:  top domain domainDNS
    SD:       532 bytes
    Metadata: 04/15/1999 15:13:21 @ SPROCKET.example.com ver: 5

Checking ACL inheritance ...
Parent: 2 - DC=example,DC=com
Child:  1 - OU=Sales,DC=example,DC=com
*** Error: Parent ACE [0] specific Mask [0x4] not found1 in child
*** Error: Parent ACE [0] specific Mask [0x8] not found1 in child
*** Error: Parent ACE [0] specific Mask [0x10] not found1 in child
*** Error: Parent ACE [0] specific Mask [0x20] not found1 in child
*** Error: Parent ACE [0] specific Mask [0x80] not found1 in child
*** Error: Parent ACE [0] specific Mask [0x20000] not found1 in child


Checking ACL inheritance ...
Parent: 1 - OU=Sales,DC=example,DC=com
Child:  0 - CN=Someone,OU=Sales,DC=example,DC=com
```

**Example 4: Determine Whether Changes to the Security Descriptor are Being Replicated from One Domain Controller to Another**

In this example, changes to the security descriptor were performed against the same object on a different domain controller. Note in the security metadata that the version number has been incremented, and that the name of the domain controller that

originated the updated security descriptor is different (in this example: wombat.example.com).

Type the following at the command line:

```
sdcheck sprocket someone@example.com
```

The following output appears:

```
Microsoft(R) Windows (R) 2000 Operating System
Security Descriptor Check Utility - build(2015)

Input:  someone@example.com
Object: CN=Someone,OU=Sales,DC=example,DC=com
Domain: example.com
Domain: DC=example,DC=com
Server: sprocket

*** Warning: No values returned for dSCorePropagationData on DC=example,DC=com

Object:   CN=Someone,OU=Sales,DC=example,DC=com
Classes:  top person organizationalPerson user
SD:       1144 bytes
Metadata: 04/15/1999 17:26:52 @ WOMBAT.example.com ver: 4
History:  04/15/1999 14:59:00 flags(0x1) SD propagation
          04/15/1999 14:59:08 flags(0x1) SD propagation
          04/15/1999 15:13:22 flags(0x1) SD propagation

  Object:   OU=Sales,DC=example,DC=com
  Classes:  top organizationalUnit
  SD:       460 bytes
  Metadata: 04/15/1999 14:54:13 @ SPROCKET.example.com ver: 1
  History:  04/15/1999 15:13:21 flags(0x1) SD propagation

    Object:   DC=example,DC=com
    Classes:  top domain domainDNS
    SD:       532 bytes
    Metadata: 04/15/1999 15:13:21 @ SPROCKET.example.com ver: 5

Checking ACL inheritance ...
Parent: 2 - DC=example,DC=com
Child:  1 - OU=Sales,DC=example,DC=com

Checking ACL inheritance ...
Parent: 1 - OU=Sales,DC=example,DC=com
Child:  0 - CN=Someone,OU=Sales,DC=example,DC=com
```

**See Also**
**Concepts**
Sdcheck Overview
Sdcheck Remarks
Sdcheck Syntax
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)

# Sdcheck Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SDCheck Syntax**

SDCheck uses the following syntax:

**sdcheck** *Server Object* [**/dumpsd**] [**/dumpall**] [**/debug**] [[**/domain:***DomainName*] **/user:** *UserName* **/password:** *Password*]

**Parameters**

***Server***

Specifies the name of the domain controller on which to perform the query. The name may be either the DNS name or the NetBIOS name of the domain controller.

***Object***

Specifies the name of the object on which you want to determine the access control list (ACL) entries. The name can be either the logon name (for example, EXAMPLE\someone), the fully distinguished name (for example, cn=Someone, ou=Sales,dc=example,dc=com), or the user principal name (for example, someone@example.com).

**/dumpsd**

Displays only the security descriptor for the selected object.

**/dumpall**

Displays the security descriptor for the selected object along with all security descriptors inherited from parent containers. Administrators can use this parameter to determine the inheritance of security descriptors from parent objects through the object hierarchy.

**/debug**

Displays internal debugging information for the Security Descriptor Check Utility.

**/domain: *DomainName***

Specifies the domain name of *UserName* in either DNS or NetBIOS format. Used to specify credentials other than the current user's.

**/user: *UserName***

Specifies the name of a user other than the current user, in the form of either the logon name or user principal name. If a user principal name is specified, then a domain name is not required.

**/password: *Password***

Specifies the password of *UserName*. Used to specify credentials other than the current user's.

**See Also**

**Concepts**

Sdcheck Overview
Sdcheck Remarks
Sdcheck Examples
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview

Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Search Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Search command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools from the product CD, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Search.vbs: Active Directory Search Tool**
This command-line tool performs searches against a Lightweight Directory Access Protocol (LDAP) server. It can be used by an Administrator or any user who wants to get information from Active Directory.

All search criteria are taken from the command line. However, certain properties of the search, such as the page size, can only be changed by modifying variable values in the script.

**Corresponding UI**
To search Active Directory, open Active Directory Users and Computers.

**Concepts**
For a conceptual overview of LDAP, see the Understanding LDAP white paper on the MSDN web site.

For more information about Security Descriptors, see the Security Descriptor Definition Language topic on the MSDN Web site.

For more information about Active Directory, see Deploying Active Directory.

**System Requirements**
The following are the system requirements for this script:

- A connection to an LDAP server.

**File Required**

- Search.vbs

**See Also**
**Concepts**
Search Remarks
Search Syntax
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Search Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Search Remarks**
**Script Requirements**
Search.vbs works only with a Lightweight Directory Access Protocol (LDAP) provider, and requires Windows Script Host (WSH).

### Note

- WSH is already present on Windows 2000 Server and Windows Server 2003 operating systems. For WSH downloads, updates, and information, visit the Microsoft Windows Script Technologies Web site.

**Finding Required and Optional Values for an Attribute**
The schema defines objects as well as the attributes and permissible values for each. You can view schema classes that contain attribute information about objects. To view this information, search cn=Schema,cn=Configuration,dc=ForestRootDomain for classSchema objects.

**See Also**
**Concepts**
Search Overview
Search Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Search Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Search Examples**

```
search.vbs LDAP://DC=DomainDomain,DC=Domain,DC=Com /C:ObjectClass=User /S:SubTree /Q
```

**Additional details**

```
ADsPath1 = LDAP://CN=User1,CN=Users,DC=DomainDomain,DC=Domain,DC=Com

ADsPath2 = LDAP://CN=User2,CN=Users,DC=DomainDomain,DC=Domain,DC=Com

ADsPath3 = LDAP://CN=User3,CN=Users,DC=DomainDomain,DC=Domain,DC=Com

ADsPath4 = LDAP://CN=User4,CN=Users,DC=DomainDomain,DC=Domain,DC=Com
```

**See Also**

**Concepts**

Search Overview
Search Remarks
Search Syntax
Alphabetical List of Tools
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Search Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Search Syntax**

**search.vbs** *adspath* [**/C:***criteria*] [**/P:***properties*] [**/S:***scope*] [**/O**: *outputfile*] [**/U:***username*] [**/W:***password*] [**/Q**]

**Parameters**

***adspath***
Specifies the AdsPath of the root of the search (mandatory).

**/C: *criteria***
Specifies the criteria for the search (Default = "ObjectCategory=*").

**/P: *properties***
Specifies the properties to be retrieved (default = AdsPath).

**/S: *scope***
Specifies the scope of the search (default = OneLevel).

**/U: *username*and /W:*password***
Specifies the user name and password of the current user.

**/Q**
Suppresses all output messages.

**/O: *outputfile***
Specifies the search result output file name.

## Notes

Only the *adspath* parameter is mandatory. The default criteria is ObjectCategory =* (all classes), the default property to be retrieved is AdsPath, and the default scope is OneLevel. The default of *username* and *password* are those of the current user.

**See Also**
**Concepts**
Search Overview
Search Remarks
Alphabetical List of Tools
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview

# Setspn Overview

Updated: April 23, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008

This content has been moved to Setspn. Be sure to use a version of Setspn.exe that is included with Windows Server 2008 or later because it includes the –S parameter for adding an SPN. Although you can use Setspn –A to add an SPN, you should use Setspn -S instead because -S will verify that there are no duplicate SPNs.

If you are using Windows Server 2003 or earlier, you will not be able to use the -S switch because it is not available for that platform. You can apply Microsoft Knowledge Base 970536 (http://support.microsoft.com/kb/970536) to Windows Server 2003 to get an updated version of SetSPN that does support the –S switch. In the case where you cannot use -S, you should manually verify that there are no duplicate SPNs by first running Setspn -L.

For more information about service principal names (SPNs) and using Setspn.exe to troubleshoot problems with SPNs, see Service Principal Names (SPNs) SetSPN Syntax (Setspn.exe).

**See Also**

**Concepts**

Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc)
Acldiag Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

**Other Resources**

Service Principal Names (SPNs) on the TechNet Wiki
Service Principal Names on MSDN

# Disk and Data Management Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Disk and File System Management Tools**
Tools in this category are used to manage the following Windows features:

- Disk Defrag
- Disk Quotas
- Distribute File System (DFS)
- Encrypting File System (EFS)
- FAT
- Formatting disks
- NTFS
- Recovering disks
- Remote Storage
- Removable Storage

**Notes**

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Connstat Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connstat.cmd: FRS Connection Status Report**
This command-line tool processes the output generated when the NtfrsUtl tool is run with the **sets** parameter. Using the NtfrsUtl **sets** output as its data source, ConnStat creates a report that summarizes the status of the FRS replication connections to and from a specified computer.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
The File Replication Service (FRS), first introduced in Windows 2000, is an essential element of the Microsoft Active Directory technologies. The Active Directory service uses FRS for SYSVOL replication. FRS can also be used with the Microsoft Distributed File System (DFS) technologies for automatic replication of content between alternating targets. For more information on FRS and Active Directory replication please see the Windows Server 2003 Technical Reference.

A replica set consists of two or more computers configured to replicate the contents of a file folder. The individual computers are referred to as replica members.

FRS maintains logs of "change orders" that are either received from replication partners (the inbound log), or are to be sent to outbound replication partners (the outbound log).

The inbound log stores pending change orders to be processed. As entries are processed, FRS sends acknowledgment to the inbound partners. Data stored in the inbound log includes the globally unique identifier (GUID) of the change order, the file name, the object ID, the parent object ID, the version number, and the event time.

The outbound log stores pending change orders to be sent to outbound partners. Change orders remain in the outbound log until all outbound partners receive and acknowledge the change. Data stored in the outbound log is the same as that stored in the inbound log. Also in the outbound log is the leading (next change) and trailing (last acknowledged) index for each partner.

The outbound logs can become quite large, particularly when replica members are unavailable, links between replicas are slow, replication hours are restricted, or a large number of changes occur. For example, even when a member for a given replica set is unavailable, outbound log entries are maintained until the member becomes available. When the changes are then sent, the inbound partner sends all changes in log file order.

**System Requirements**

- Windows 2000 Server or Windows Server 2003
- Perl 5.005 521e or 522e

**Files Required**

- Connstat.cmd
- Ntfrsutl.exe

**See Also**
**Concepts**
Connstat Remarks
Connstat Syntax
Connstat Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview

Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Connstat Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ConnStat Syntax**

ConnStat uses the following syntax:

**connstat** [**-sort=***Option* [**-sort=***Option*]...] *DataFile*

**Parameters**

**-sort= *Option***

Specifies how the NtfrsUtl **sets** output is sorted. The following values are valid for the **-sort** parameter:

| Option | Description |
|---|---|
| **send** | Sorts outbound connections in ascending order by the Send Delta column of the ConnStat report. |
| **name** | Sorts outbound connections in ascending order by the Server Name column of the ConnStat report. |
| **clean** | Sorts outbound connections in ascending order by the Cleanup Delta column of the ConnStat report. |
| **lmt** | Sorts outbound connections in ascending order by the LMT (Leading minus Trailing) column of the ConnStat report. |
| **lastjointime** | Sorts outbound connections in ascending order by the LastJoinTime column of the ConnStat report. |
| **lastvvjoin** | Sorts outbound connections in ascending order by the Last VVJoin column of the ConnStat report. |

***DataFile***

Specifies the file that was generated by running NtfrsUtl with the **sets** parameter.

**Sample Usage**

```
connstat ntfrs_sets.txt

connstat sort=lastvvjoin ntfrs_sets.txt
```

**See Also**

**Concepts**

Connstat Overview
Connstat Remarks
Connstat Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Cabarc Overview
Bitsadmin Overview

# Connstat Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ConnStat Remarks**
**Before Using ConnStat**
If Perl is not installed on your system, if you are connected to the Internet, you can download a free version from the ActiveState Web site.

**Using the ConnStat Report**
The Connstat report consists of three main sections, the Report Header section, the Inbound Connections section, and the Outbound Connections section. Familiarity with FRS concepts and terminology is required in order to interpret these reports.

**The Report Header Section**
The first line of the Report Header lists the file specified in the ConnStat **datafile** parameter as well as the date that the file was last modified. The remainder of the Report Header gives information about the FRS service on the host where the input file was generated. The following table describes this information:

| Report Element | Description | Comment |
|---|---|---|
| **Replica** | Displays the replica set name and GUID. | SYSVOL replica sets are always called "DOMAIN SYSTEM VOLUME (SYSVOL SHARE)." Replica sets are uniquely identified by the GUID. |
| **Member** | Displays the member name. Individual computers in a replica set are referred to as replica members. | Note that for DFS replica sets this name is often a GUID. |
| **ServiceState** | Shows the state of the FRS service. Valid return values are between 0 and 7.<br>• **0** - Allocated: indicates an initializing state<br>• **1** - Initializing: indicates an initializing state<br>• **2** - Starting: indicates an initializing state<br>• **3** - Active: the replica set is functioning<br>• **4** - Stopped: the replica set is stopped with an error<br>• **5** - Error: the replica set is stopped with an error<br>• **6** - Jrnl_Wrap_Error: the replica set is stopped with an error<br>• **7** - Replica_Deleted: indicates the replica set is marked for deletion | The friendly name is displayed in parentheses next to the return value, for example 3 (ACTIVE). |

| | | |
|---|---|---|
| **OutLogSecNum** | Displays the current outbound log sequence number. | |
| **OutLogCleanup** | Displays the outbound log sequence number where the next clean-up will begin. | |
| **Delta** | Represents the difference between OutLogSecNum and OutLogCleanup. This is a good estimation of the number of change orders currently present in the outbound log. | |
| **Config Flags** | Displays state information about the replica set. Possible values are: **Multimaster**, **Primary**, or **Seeding**. | In multi-master replica sets, any domain controller accepts and replicates directory changes to any other domain controller. A primary flag means the host was a primary member when the replica set was created. The seeding flag, used for SYSVOL only, indicates the replica set is in the initial, seeding state. |
| **Root Path** | The root path used by the replica member. | |
| **Staging Path** | The staging path used by the replica member. | |
| **File Filter** | Lists file extensions that were excluded during replication, if any. | |
| **Directory Filter** | Lists directories that were excluded during replication, if any. | |

**The Inbound Connections Section**

This section gives state information for all inbound connections for the given member machine. The following table lists columns and descriptions for this section:

| Column Heading | Description | Comments |
|---|---|---|
| **Partner** | Displays the name of the immediate inbound partner for this replica member (member machine). | To participate in a valid FRS replica set, each member must have at least one inbound partner. |
| **I/O** | Specifies whether the connection is inbound (In) or outbound (Out). | This column appears in both the Inbound and Outbound report sections. |
| **State** | Specifies the current inbound connection state. | See the **Connection State** table in the **Connection State Definitions** remark for a list of valid Inbound connection states and their meanings. |

| | | |
|---|---|---|
| **Rev** | Represents operating system version information for the communication protocol of the inbound partner. Valid values include the following:<br>• **0** - The connection is not currently joined.<br>• **3** - The partner is running either Windows 2000 v. 2195, Windows 2000 v. 2195 SP1, or Windows 2000 v. 2195 SP2 with compression disabled.<br>• **4** - The partner is running Windows 2000 v. 2195 SP2 with compression enabled.<br>• **7** - The partner is running Windows 2000 v. 2195 SP3.<br>• **8** - The partner is running a member of the Windows Server 2003. | |
| LastJoinTime | Displays a timestamp showing when the replica member last successfully joined with this partner. | This data is available only for partners running Windows 2000 SP2 or later. |

**The Outbound Connections Section**

This section gives state information for all outbound connections for the given member machine. The following table lists columns and descriptions for this section.

Note

- The first several columns are identical to the Inbound connection report, with some additional information in the State column (see the **Connection State** table in the **Connection State Definitions** remark for more information). The columns starting with **OLog State** contains outbound log-related state information for the connection.

| **Column Heading** | **Description** | **Comments** |
|---|---|---|
| **Partner** | Displays the name of the immediate outbound partner for this replica member. | |
| **I/O** | Specifies whether the connection is inbound (In) or outbound (Out). | |
| **State** | Specifies the current outbound connection state.<br><br>See the **Connection State** table in the **Connection State Definitions** remark for a list of valid Outbound connection states and their meaning. | In the outbound connections report, the appearance of the tag "-vv" means that this connection is performing a Version Vector-based join operation. Typically this is done when the outbound partner is first added to the replica set or if it is in the process of doing a non-authoritative restore (a D2). Once this initial sync is complete, the connection leaves the VVJoin state. |

| | | |
|---|---|---|
| **Rev** | Represents operating system version information for the outbound partner's communication protocol.<br><br>See the **Inbound Connection** table in the **Inbound Connections Section** remark for definitions of the operating system versions. | |
| **LastJoinTime** | Displays a timestamp showing when the replica member last successfully joined with this partner. | This data is available only for partners running Windows 2000 SP2 or later. |
| **OLog State** | Specifies the state of outbound log processing for this connection.<br><br>• **OLP_UNJOINED** - The partner is not joined.<br><br>• **OLP_ELIGIBLE** - The partner can accept further change orders (Connection joined, and "change orders out" is less than the max limit).<br><br>• **OLP_STANDBY** - The partner is ready to join the eligible list.<br><br>• **OLP_AT_QUOTA** - The partner is at max quota for outstanding change orders.<br><br>• **OLP_INACTIVE** - The partner is not accepting change orders.<br><br>Note<br><br>• OLP_AT_QUOTA is a normal condition when FRS is actively replicating because the service always tries to keep eight change orders outstanding on each joined outbound connection. | |
| **Leadx** | Specifies the index into the outbound log for the next change order to be processed for this connection. | |

| | | |
|---|---|---|
| **Send Delta** | Specifies the difference between **OutLogSeqNum** (index of the most recent change order in the outlog) and the **Leadx** value for this connection. In other words, **Send Delta** specifies the number of change orders that remain to be processed by this connection.<br><br>A **Send Delta** of zero means that this connection is fully synchronized at this time. Unless there is very little activity in the replica set or the connection schedule is *always on*, you are unlikely to find connections with a **Send Delta** of zero. | |
| **Trailx** | Specifies the index into the outbound log of the oldest unacknowledged outbound change order. Each change order sent to the partner for this connection must eventually be acknowledged. | |
| **Cleanup Delta** | Specifies the difference between the **Trailx** value for this connection and **OutLogCleanup** index, the latter being the minimum **Trailx** value for all outbound connections. | Partners having connections with small cleanup deltas should be examined carefully to verify their operational status, since they are often the ones who are furthest behind in replicating content from this member. |
| **LMT** | Specifies the difference between the values in the **Leadx** and **Trailx** columns. This maximum value is limited by the size of the **AckVector**, typically 128.<br><br>A value of 127 in this column indicates that the connection may require investigation. | If the value of **LMT** isnbsp;127, wait for a while to see if the condition changes. If it does not, first try to delete and recreate the related connection object in the directory service. If the condition persists, then you may need to run a non-authoritative restore on the outbound partner. |
| **COs Out** | Specifies the number of active unacknowledged change orders pending at the partner. By default, FRS allows up to 8 change orders to be outstanding at a time on each outbound connection. | This limit is controlled by a registry parameter, which you may want to increase in an environment with high latency communication links. Note that FRS on the partner will try to initiate fetch requests for staging files for all eight change orders concurrently, so increasing this registry parameter will increase the load on the upstream member. See the Microsoft Windows Server 2003 Technical Reference for more details. |
| **Last VVJoin** | Specifies the last date and time that a Version Vector-based join was done on the connection. | |

**Connection State Definitions**

The following table lists Connection State values and their meaning:

| Connection State | Description |
|---|---|
| **Init** | Specifies that the connection is newly allocated. |
| **Unjoined** | Specifies that the connection is not yet joined to a partner. |
| **Start** | Specifies that the inbound partner has requested a join (StartJoin). |
| **Starting** | Specifies that the join has started. |
| **Scanning** | Specifies that FRS is scanning the inbound log for the replica set. |

| | |
|---|---|
| **SendJoin** | Specifies that the scan was completed, and the join request has been forwarded to the partner. |
| **WaitJoin** | Specifies that FRS is waiting for a reply from the partner. |
| **Joined** | Specifies that the join request succeeded. |
| **Unjoining** | Specifies that the join request failed after a series of retries. |
| **Deleted** | Specifies that the connection has been deleted. |

**See Also**

**Concepts**

Connstat Overview
Connstat Syntax
Connstat Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Cabarc Overview
Bitsadmin Overview

# Connstat Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ConnStat Examples**

**Example 1: Generating the NTFRS SETS Output**

ConnStat organizes the output of the NtfrsUtl tool, specifically the output generated when NtfrsUtl is run with the **sets** parameter. This example creates output that will be used by ConnStat. Type the following at the command line:

**ntfrsutl sets >ntfrs_sets.txt**

This command generates an output file called Ntfrs_sets.txt in the current directory. The command must be run on a computer that is a replica set member.

**Example 2: Using ConnStat to Parse an NtfrsUtl sets Datafile**

This example parses the data file generated in Example 1: Generating the NTFRS SETS Output and then displays the results to the screen. Type the following at the command line to save the ConnStat report to a file in the current directory:

**connstat.cmd ntfrs_sets.txt >connstat_out.txt**

This command runs the ConnStat script against the Ntfrs_sets.txt file, which was generated by the NtfrsUtl **sets** command and saves the resulting output in a file named Connstat_out.txt in the current directory.

**See Also**

**Concepts**

Connstat Overview
Connstat Remarks
Connstat Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Cabarc Overview
Bitsadmin Overview

# Devcon Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devcon.exe: Device Console Utility**
This command-line tool brings many of the functions of Device Manager to the command window. Devcon enables, disables, installs, configures, and removes devices on the local computer and displays detailed information about devices on local and remote computers.

- **Display.** Display the following properties of devices on local and remote computers:
  - Hardware IDs, compatible IDs, and instance IDs of devices
  - Device setup classes
  - The devices in a device setup class
  - Driver files
  - Driver nodes
  - Hardware resources
  - Status
  - Expected driver stack
- **Search.** Search for installed and uninstalled devices on a local or remote computer by hardware ID, instance ID, or device setup class.
- **Change.** Changes the status or configuration of Plug and Play devices on the local computer.
  - Enable
  - Disable
  - Update
  - Install
  - Remove
  - Rescan for Plug and Play devices
  - Change the upper and lower filter drivers for a device setup class
- **Restart.** Restart a local device, restart (reboot) the local system on demand, or restart the local system, when required, for another Devcon operation.

**Notes**

- To display help information for this tool, type the following at the command line:

  **devcon /?**

- To display detailed help information for an operation, type the following at the command line:

  **devcon /?operation**

**Corresponding UI**
You can also use Device Manager, a Microsoft Management Console snap-in, to view and manage devices. Devcon and Device Manager perform many of the same functions and the tools can be used together and interchangeably.

To open Device Manager

- Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Computer Management**.
- Expand the **System Tools** node and then click **Device Manager**.

or

- Click **Start** and then click **Run**.
- In the **Run** dialog box, type **devmgmt.msc**.

Open Device Manager now.

**Concepts**
A *device* is a piece of equipment that is connected to the computer and controlled by its microprocessors. This category includes equipment built into the computer by the manufacturer and equipment added later. The operating system interacts with devices through customized software, such as device drivers and filters. The device environment also includes supporting resources, such as memory, DMA channels, and port addresses.

Most devices manufactured since 1995 are Plug and Play devices. Plug and Play devices are automatically installed and configured for use by plugging them in and are automatically uninstalled by unplugging them. The operating system can control almost all aspects of a Plug and Play device.

With the advent of Plug and Play technology, the tasks of installing, configuring and uninstalling devices are simplified and most user configuration is eliminated. However, developers, testers and support engineers often need to view and manipulate devices and their supporting software and resources.

Developers typically use Device Manager, the Windows Management Interface (WMI) procedures, or functions in the standard Windows APIs to view and configure devices. Devcon provides a relatively complete set of basic device handling functions available from the command line.

For more information on device drivers and device configuration, see the Windows Driver Development Kit.

**System Requirements**
The following are the system requirements for this tool:

- Windows 2000 Professional, Windows 2000 Server, Windows XP Professional, or Windows Server 2003
- To change the properties of device drivers, users must have the following permissions set. Administrators have these permissions.
  - Load\Unload Drivers (set in System Policy Editor, Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment).
  - Permission to copy files to the System32\Drivers directory (set in the file system directory).
  - Permission to write settings to the registry (set in Regedit.exe).
- To install drivers, users must be logged on to an account in the Administrators group.

### Note

- When users do not have the required permissions, Devcon displays a generic "devcon failed" message with no further explanation.

### File Required

- Devcon.exe

**See Also**
**Concepts**
Devcon Remarks
Devcon Syntax
Devcon Examples
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Devcon Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devcon Remarks**

**Identifiers**
Devcon uses hardware IDs, compatible IDs, and instance IDs to identify devices.

- A *hardware ID* is a vendor-defined string that the Plug and Play Manager uses to find the INF file for a device. A device can have one or more hardware IDs.

- A *compatible ID* is a vendor-defined string that is more generic than the hardware ID. Devices that use the same INF file can have the same compatible ID. The Plug and Play Manager uses the compatible ID to find an INF file for a device only when it cannot find the hardware ID. However, because the compatible ID is not as reliable as a hardware ID, the system prompts the user for confirmation before installing a device based on a compatible ID match.

- An *instance ID* is a string that uniquely distinguishes a device from other devices on the computer. The instance ID of a device is unique only on a computer and can vary between computers and operating systems. The instance ID of most devices is also its registry path in the Enum registry subkey (HKLM\System\CurrentControlSet\Enum).

  In Devcon commands, instance IDs must be prefixed with the at character (@) to distinguish them from hardware IDs and compatible IDs. For example, **devcon status root*** searches for hardware IDs or compatible IDs that begin with "root\". However, **devcon hwids @root*** searches for instance IDs (device registry paths) that begin with "root\."

**Devcon Search Logic**
Devcon identifies devices by their computer name, hardware ID, compatible ID, instance ID, and device setup class.

When an ID or ID pattern (an ID with wildcard characters (*)) is preceded by the at character (@), Devcon searches only the instance IDs of devices for a match. Otherwise, it searches both the hardware IDs and compatible IDs of devices for a match.

When a command includes more than one ID or ID pattern, Devcon returns devices whose IDs match any of the IDs or ID patterns. That is, it assumes an OR between the ID arguments.

For example, **devcon hwids *pnp* *mou*** returns devices that include either "pnp" or "mou" in their hardware ID or compatible ID.

When a command includes a device setup class, Devcon first limits the search to the setup class and then applies the ID criteria. Devcon returns devices in the class that match any of the ID patterns. That is, it assumes an AND between the class and the IDs and an OR between each of the ID arguments.

For example, **devcon hwids =media *pnp* *microsoft*** returns devices in the media device setup class that include either "pnp" or "microsoft" in their hardware ID or compatible ID.

**Operational Details**

- **Classes**. Classes are listed in the order that they appear in the registry (alphanumeric order by GUID).

- **Stack**. The stack operation displays the expected driver stack for a device. Although the actual driver stack typically matches the expected stack, variations are possible.

- **Classfilter**. Devcon will not add a class filter to a class unless the driver is installed as a service, that is, there must be a subkey for the driver in the Services subkey (HKLM\System\CurrentControlSet\Services). This safeguard prevents you from accidentally adding a filter driver that does not exist and, thereby, rendering the system unbootable.

- **Findall**. The Findall operation differs from the Find operation, and from other Devcon device enumerations, in that it omits the DIGCF_PRESENT flag when calling SetupDiGetClassDevsEx function. As a result, the Findall display includes *phantom devices* that are not currently attached to the computer. For more information, see SetupDiGetClassDevsEx in

the Windows DDK.

- **Rescan**. Rescanning can cause the Plug and Play Manager to install device drivers without warning when it detects new devices. Rescanning can also detect some non-Plug and Play devices, particularly those that cannot notify the system when they are installed, such as parallel-port devices and serial-port devices. As a result, Administrator privileges are required to run rescan commands.

**Restart, Reboot, and /r**

Devcon provides a method to restart devices, and two methods to restart the operating system.

- The **/r** parameter is a conditional restart that restarts the operating system only when a restart is required to make the accompanying operation effective. This parameter is valid only in commands that include a Devcon operation. It can restart the system on a local or remote computer.

  If a system restart is required, Devcon displays a message explaining that you must restart the system to complete the operation. In response, you can use the **reboot** operation command.

- The **reboot** operation forces the operating system to restart. It is valid only on a local computer and it cannot be combined with other operations.

- The **restart** operation restarts the specified devices. It is valid only on a local computer and it cannot be combined with other operations. Some devices cannot be started without a reboot. The **/r** parameter can be used to reboot the operating system if the devices cannot be started without one.

**Devcon Return Codes**

Devcon returns a numeric (integer) code that can be used in programs and scripts to determine the success of a Devcon command. (For example, **return = devcon hwids ***). The following table lists and describes the return codes.

| Return code | Meaning |
|---|---|
| 0 | Success |
| 1 | Requires reboot |
| 2 | Failure |
| 3 | Syntax error |

**See Also**

**Concepts**

Devcon Overview
Devcon Syntax
Devcon Examples
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Apmstat Overview

# Devcon Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devcon Examples**

**Display operations**

hwids

classes

listclass

driverfiles

drivernodes

resources

stack

status

**hwids**

The hwids operation displays the hardware IDs, compatible IDs, and instance IDs of devices on the system.

**Example 1: Find all hardware IDs**

Because Devcon operations use IDs and ID patterns to identify devices, a common first step in using Devcon is to create a hardware ID reference file for devices on the computer.

The following command uses the hwids operation, which returns the IDs and the device description. It uses the wildcard character (*) to represent all devices on the local computer.

**devcon hwids ***

Because the output is lengthy and used repeatedly, save the output in a text file for reference.

The following command uses the wildcard character (*) in the hwids operation command to represent all devices on the computer. It uses the redirection character (>) to save the command output in the hwids.txt file.

**devcon hwids * > hwids.txt**

The following command finds the hardware IDs of devices on a remote computer, Server01. It uses the **/m** parameter to specify the name of the remote computer. The command redirects the output to the server01_hwids.txt file for later reference.

**Note**

- This command fails unless the user has the required permissions on the remote computer.

**devcon /m:\\server01 hwids * > server01_hwids.txt**

**Example 2: Find hardware IDs by using a pattern**

To find the hardware IDs of a particular device, you must enter the hardware ID or pattern, the compatible ID or pattern, the instance ID or pattern, or the name of the device setup class.

The following command uses a pattern to find the actual hardware ID of the floppy disk drive on the computer. (The user assumes that the pattern appears in the one of the device identifiers.) The command uses the wildcard character (*) to represent all characters that might precede or follow the word, "floppy," in any of the IDs.

**devcon hwids *floppy***

In response, Devcon displays the instance ID, hardware ID and compatible ID of the floppy disk drive on the computer. These IDs can be used in subsequent Devcon commands.

```
FDC\GENERIC_FLOPPY_DRIVE\5&39194F6D&0&0
    Name: Floppy disk drive
    Hardware ID's:
        FDC\GENERIC_FLOPPY_DRIVE
    Compatible ID's:
        GenFloppyDisk
1 matching device(s) found.
```

In this case, the phrase "floppy" occurs in the hardware ID or compatible ID of only one device on the computer. If it occurs in the ID of more than one device, all devices with "floppy" in their IDs appear in the output.

**Example 3: Find hardware IDs by using a class**
The following command uses a device setup class to find the hardware IDs of all devices in the Ports device setup class. The equal sign (=) preceding the class name indicates that it is a class, not an ID.

**devcon hwids =ports**

In response, Devcon displays the hardware ID and compatible ID of all three devices in the Ports setup class.

```
ACPI\PNP0401\4&B4063F4&0
    Name: ECP Printer Port (LPT1)
    Hardware ID's:
        ACPI\PNP0401
        *PNP0401
ACPI\PNP0501\1
    Name: Communications Port (COM1)
    Hardware ID's:
        ACPI\PNP0501
        *PNP0501
ACPI\PNP0501\2
    Name: Communications Port (COM2)
    Hardware ID's:
        ACPI\PNP0501
        *PNP0501
3 matching device(s) found.
```

**classes**
The classes operation lists the device setup class of devices in the system.

**Example 4: List classes on the local computer**
Because Devcon operations can use the device setup class to identify devices, it is useful to create a reference file of the device setup classes of devices on the computer.

The following command uses classes operation, which returns a list and description of all classes on the computer. The classes operation does not take parameter arguments that specify a device. It only lists all classes on the computer.

**devcon classes**

Because the output is lengthy and used repeatedly, save the output in a text file for reference.

The following command displays all device classes on the computer. It uses the redirection character (>) to save the command output in the classes.txt file.

**devcon classes > classes.txt**

**Example 5: List classes on the remote computer**
The following command lists the device setup classes on a remote computer, Server01

**devcon /m:\\server01 classes**

Because the output is lengthy and used repeatedly, save the output in a text file for reference.

The following command uses the redirection character (>) to save the command output in the server01_classes.txt file.

**devcon /m:\\server01 classes > server01_classes.txt**

**listclass**
The listclasses operation lists the devices in a device setup class.

**Example 6: List the devices in a device setup class**

The following command lists the devices in Net, the device setup class for network adapters. This command is usually preceded by a command to list all of the classes on the computer (**devcon classes**), to get the correct name of the setup class.

**devcon listclass net**

In response, Devcon displays the instance ID and description of each device in the Net setup class.

```
Listing 6 device(s) for setup class "Net" (Network adapters).
PCI\VEN_10B7&DEV_9200&SUBSYS_00BE1028&REV_78\4&BB7B4AE&0&60F0: 3Com 3C920 Integrated
Fast Ethernet Controller (3C905C-TX Compatible)
ROOT\MS_L2TPMINIPORT\0000                                    : WAN Miniport (L2TP)
ROOT\MS_NDISWANIP\0000                                       : WAN Miniport (IP)
ROOT\MS_PPPOEMINIPORT\0000                                   : WAN Miniport (PPPOE)
ROOT\MS_PPTPMINIPORT\0000                                    : WAN Miniport (PPTP)
ROOT\MS_PTIMINIPORT\0000                                     : Direct Parallel
```

This display, while interesting, does not provide the hardware IDs of the devices in the Net setup class. The following command uses the Devcon Syntax operation to list the devices in the Net setup class. In the hwids operation, the class name is preceded by an equal sign (=) to indicate that it is a class, not an ID.

**devcon hwids =net**

The resulting display lists the devices in the Net class and includes the instance ID, hardware IDs and compatible IDs of devices in the class.

```
PCI\VEN_10B7&DEV_9200&SUBSYS_00BE1028&REV_78\4&BB7B4AE&0&60F0
    Name: 3Com 3C920 Integrated Fast Ethernet Controller (3C905C-TX Compatible)
    Hardware ID's:
        PCI\VEN_10B7&DEV_9200&SUBSYS_00BE1028&REV_78
        PCI\VEN_10B7&DEV_9200&SUBSYS_00BE1028
        PCI\VEN_10B7&DEV_9200&CC_020000
        PCI\VEN_10B7&DEV_9200&CC_0200
    Compatible ID's:
        PCI\VEN_10B7&DEV_9200&REV_78
        PCI\VEN_10B7&DEV_9200
        PCI\VEN_10B7&CC_020000
        PCI\VEN_10B7&CC_0200
        PCI\VEN_10B7
        PCI\CC_020000
        PCI\CC_0200
ROOT\MS_L2TPMINIPORT\0000
    Name: WAN Miniport (L2TP)
    Hardware ID's:
        ms_l2tpminiport
ROOT\MS_NDISWANIP\0000
    Name: WAN Miniport (IP)
    Hardware ID's:
        ms_ndiswanip
ROOT\MS_PPPOEMINIPORT\0000
    Name: WAN Miniport (PPPOE)
    Hardware ID's:
        ms_pppoeminiport
ROOT\MS_PPTPMINIPORT\0000
    Name: WAN Miniport (PPTP)
    Hardware ID's:
        ms_pptpminiport
ROOT\MS_PTIMINIPORT\0000
    Name: Direct Parallel
    Hardware ID's:
        ms_ptiminiport
6 matching device(s) found.
```

**Example 7: List the devices in a multiples classes on a remote computer**

The following command lists the devices in the DiskDrive, CDROM, and TapeDrive classes on Server01, a remote computer.

**devcon /m:\\server01 listclass diskdrive cdrom tapedrive**

In response, Devcon displays the devices in those classes on the remote computer.

```
Listing 1 device(s) for setup class "DiskDrive" (Disk drives) on \\server01.
IDE\DISKWDC_WD204BA_____________________________16.13M16\4457572D414D373032313633393820
3120202020: WDC WD204BA
Listing 1 device(s) for setup class "CDROM" (DVD/CD-ROM drives) on \\server01.
IDE\CDROMSAMSUNG_DVD-ROM_SD-608__________________2.2_____\4&13B4AFD&0&0.0.0: SAMSUNG
DVD-ROM SD-608
No devices for setup class "TapeDrive" (Tape drives) on \\server01.
```

**driverfiles**

The driverfiles operation lists the file names of device drivers used by devices on the system.

**Example 8: List all driver files**

The following command lists the file names of drivers used by devices on the system. The command uses the wildcard character (*) to indicate all devices on the system. Because the output is extensive, the command uses the redirection character (>) to redirect the output to a reference file, driverfiles.txt.

### devcon driverfiles * > driverfiles.txt

**Example 9: List the driver files of a particular device.**

The following command searches for the device driver used by the mouse device on the local computer. It identifies the device by one of its hardware IDs, HID\Vid_045e&Pid_0039&Rev_0121. (To find the hardware ID of a device, use the Devcon Syntax operation.) The hardware ID is enclosed in quotation marks because it includes the ampersand character (&).

### devcon driverfiles "HID\Vid_045e&Pid_0039&Rev_0121"

In response, Devcon displays the two device drivers that support the mouse device.

```
HID\VID_045E&PID_0039\6&DC36FDE&0&0000
    Name: Microsoft USB IntelliMouse Optical
    Driver installed from c:\windows\inf\msmouse.inf [HID_Mouse_Inst]. 2 file(s)
 used by driver:
        C:\WINDOWS\System32\DRIVERS\mouhid.sys
        C:\WINDOWS\System32\DRIVERS\mouclass.sys
1 matching device(s) found.
```

**drivernodes**

The drivernodes operation lists all driver packages that are compatible with the device, along with their version and ranking. Valid only on the local computer.

**Example 10: List driver packages by hardware ID pattern**

The following command lists the driver nodes of all devices whose hardware ID begins with "sw" (software devices). Patterns are useful for finding information about similar devices that might not be in the same setup class.

The drivernodes operation is particularly useful for troubleshooting setup problems. For example, you can use the drivernodes operation to find out whether an INF file included in Windows was used for a device instead of a customized third-party INF file.

### devcon drivernodes sw*

In response, Devcon displays the driver nodes of devices whose hardware IDs or compatible IDs begin with "sw."

```
SW\{A7C7A5B0-5AF3-11D1-9CED-00A024BF0407}\{9B365890-165F-11D0-A195-0020AFD156E4}

    Name: Microsoft Kernel System Audio Device
DriverNode #0:
    Inf file is c:\windows\inf\wdmaudio.inf
    Inf section is WDM_SYSAUDIO
    Driver description is Microsoft Kernel System Audio Device
    Manufacturer name is Microsoft
    Provider name is Microsoft
    Driver date is 7/1/2001
    Driver version is 5.1.2535.0
    Driver node rank is 0
    Driver node flags are 00002244
```

```
        Inf is digitally signed
SW\{B7EAFDC0-A680-11D0-96D8-00AA0051E51D}\{9B365890-165F-11D0-A195-0020AFD156E4}

    Name: Microsoft Kernel Wave Audio Mixer
DriverNode #0:
    Inf file is c:\windows\inf\wdmaudio.inf
    Inf section is WDM_KMIXER
    Driver description is Microsoft Kernel Wave Audio Mixer
    Manufacturer name is Microsoft
    Provider name is Microsoft
    Driver date is 7/1/2001
    Driver version is 5.1.2535.0
    Driver node rank is 0
    Driver node flags are 00002244
        Inf is digitally signed
SW\{CD171DE3-69E5-11D2-B56D-0000F8754380}\{9B365890-165F-11D0-A195-0020AFD156E4}

    Name: Microsoft WINMM WDM Audio Compatibility Driver
DriverNode #0:
    Inf file is c:\windows\inf\wdmaudio.inf
    Inf section is WDM_WDMAUD
    Driver description is Microsoft WINMM WDM Audio Compatibility Driver
    Manufacturer name is Microsoft
    Provider name is Microsoft
    Driver date is 7/1/2001
    Driver version is 5.1.2535.0
    Driver node rank is 0
    Driver node flags are 00002244
        Inf is digitally signed
3 matching device(s) found.
```

**Example 11: List driver packages by instance ID pattern**

The following command lists the driver nodes of all devices whose instance IDs begin with ROOT\MEDIA, that is, devices in the Enum\Root\Media registry subkey. The command uses the at character (@) to indicate that the phrase is in the instance ID.

**devcon drivernodes @ROOT\MEDIA***

In response, Devcon displays the driver nodes of devices whose instance ID begins with "ROOT\MEDIA"

```
ROOT\MEDIA\MS_MMACM
    Name: Audio Codecs
DriverNode #0:
    Inf file is c:\windows\inf\wave.inf
    Inf section is MS_MMACM
    Driver description is Audio Codecs
    Manufacturer name is (Standard system devices)
    Provider name is Microsoft
    Driver date is 7/1/2001
    Driver version is 5.1.2535.0
    Driver node rank is 0
    Driver node flags are 00002240
        Inf is digitally signed
ROOT\MEDIA\MS_MMDRV
    Name: Legacy Audio Drivers
DriverNode #0:
    Inf file is c:\windows\inf\wave.inf
    Inf section is MS_MMDRV
    Driver description is Legacy Audio Drivers
    Manufacturer name is (Standard system devices)
    Provider name is Microsoft
    Driver date is 7/1/2001
    Driver version is 5.1.2535.0
    Driver node rank is 0
    Driver node flags are 00002240
        Inf is digitally signed
ROOT\MEDIA\MS_MMMCI
    Name: Media Control Devices
DriverNode #0:
    Inf file is c:\windows\inf\wave.inf
    Inf section is MS_MMMCI
```

```
    Driver description is Media Control Devices
    Manufacturer name is (Standard system devices)
    Provider name is Microsoft
    Driver date is 7/1/2001
    Driver version is 5.1.2535.0
    Driver node rank is 0
    Driver node flags are 00002240
        Inf is digitally signed
ROOT\MEDIA\MS_MMVCD
    Name: Legacy Video Capture Devices
DriverNode #0:
    Inf file is c:\windows\inf\wave.inf
    Inf section is MS_MMVCD
    Driver description is Legacy Video Capture Devices
    Manufacturer name is (Standard system devices)
    Provider name is Microsoft
    Driver date is 7/1/2001
    Driver version is 5.1.2535.0
    Driver node rank is 0
    Driver node flags are 00002240
        Inf is digitally signed
ROOT\MEDIA\MS_MMVID
    Name: Video Codecs
DriverNode #0:
    Inf file is c:\windows\inf\wave.inf
    Inf section is MS_MMVID
    Driver description is Video Codecs
    Manufacturer name is (Standard system devices)
    Provider name is Microsoft
    Driver date is 7/1/2001
    Driver version is 5.1.2535.0
    Driver node rank is 0
    Driver node flags are 00002240
        Inf is digitally signed
5 matching device(s) found.
```

**resources**

The resources operation lists the resources allocated to each device in the system.

**Example 12: List resources of a class of devices**

The following command displays the resources allocated to devices in the hdc class. This class includes IDE controllers. The equal sign (=) is prepended to "hdc" to indicate that it is a class and not an ID.

### devcon resources =hdc

In response, Devcon lists the resources allocated to IDE controllers on the local computer.

```
PCI\VEN_8086&DEV_244B&SUBSYS_00000000&REV_02\3&29E81982&0&F9
    Name: Intel(r) 82801BA Bus Master IDE Controller
    Device is currently using the following resources:
        IO  : ffa0-ffaf
PCIIDE\IDECHANNEL\4&37E53584&0&0
    Name: Primary IDE Channel
    Device is currently using the following resources:
        IO  : 01f0-01f7
        IO  : 03f6-03f6
        IRQ : 14
PCIIDE\IDECHANNEL\4&37E53584&0&1
    Name: Secondary IDE Channel
    Device is currently using the following resources:
        IO  : 0170-0177
        IO  : 0376-0376
        IRQ : 15
3 matching device(s) found.
```

**Example 13: List resources of device on a remote computer by ID**

The following command lists the resources allocated to the system timer on Server01, a remote computer. The command uses the hardware ID of the system timer, ACPI\PNP0100, to specify the device.

**devcon /m:\\Server01 resources *PNP0100**

In response, Devcon displays the resources of the Server01 system timer.

```
ROOT\*PNP0100\PNPBIOS_8
    Name: System timer
    Device has the following resources reserved:
        IO  : 0040-005f
        IRQ : 0
1 matching device(s) found on \\server01.
```

The following command uses the instance ID of the remote system timer in the Devcon resources command. The at character (@) indicates that the string is an instance ID, not a hardware ID or compatible ID.

**devcon /m:\\Server01 resources @ACPI\PNP0100\4&b4063f4&0**

**stack**

The stack operation displays the expected driver stack for devices on the computer.

**Example 14: Display the driver stack for storage devices**

The following command searches for devices in the volume setup class and displays the expected driver stack for those devices. The equal sign (=) indicates that the string is a class name.

**devcon stack =Volume**

In response, Devcon displays the expected stack for the devices. The data returned includes the instance ID and description of each device, the GUID and name of the device setup class, the names of upper and lower filter drivers, and controlling services, if any.

```
STORAGE\VOLUME\1&30A96598&0&SIGNATURE32323533OFFSET271167600LENGTH6E00D0C00
    Name: Generic volume
    Setup Class: {71A27CDD-812A-11D0-BEC7-08002BE2092F} Volume
    Class upper filters:
        VolSnap
    Controlling service:
        (none)
STORAGE\VOLUME\1&30A96598&0&SIGNATURE32323533OFFSET7E00LENGTH27115F800
    Name: Generic volume
    Setup Class: {71A27CDD-812A-11D0-BEC7-08002BE2092F} Volume
    Class upper filters:
        VolSnap
    Controlling service:
        (none)
2 matching device(s) found.
```

**Example 15: Find the setup class of a device**

The stack operation returns the setup class of a device in addition to the upper and lower filter drivers. The following commands find the setup class of the printer port interface by finding its instance ID, and then using the instance ID to find its setup class.

The following command finds the instance ID of the printer port interface by using the Devcon Syntax operation and "LPT," a phrase in the printer port hardware ID.

**devcon hwids *lpt***

In response, Devcon returns the instance ID (displayed in bold text), as well as the hardware ID of the printer port interface.

```
            LPTENUM\MICROSOFTRAWPORT\5&CA97D7E&0&LPT1
    Name: Printer Port Logical Interface
    Hardware ID's:
        LPTENUM\MicrosoftRawPort958A
        MicrosoftRawPort958A
1 matching device(s) found.
```

The next command uses the stack operation to find the class of the device represented by the instance ID. An at character (@)

identifies the ID as an instance ID. The ID is enclosed in quotation marks because it includes ampersand characters.

**devcon stack "@LPTENUM\MICROSOFTRAWPORT\5&CA97D7E&0&LPT1"**

In response, Devcon displays the driver stack for the printer port interface, including the class. The display reveals that the printer port is in the System class.

```
LPTENUM\MICROSOFTRAWPORT\5&CA97D7E&0&LPT1
    Name: Printer Port Logical Interface
    Setup Class: {4D36E97D-E325-11CE-BFC1-08002BE10318} System
    Controlling service:
        (none)
1 matching device(s) found.
```

**Example 16: Display the stack for related devices on a remote computer**
The following command displays the expected stack for miniport devices on Server01, a remote computer. It searches for devices in the Net setup class that have "miniport" in their hardware ID or compatible ID. Note that this command first limits the search to the Net setup class and then find the "miniport" string. It does not find devices other than those in the Net setup class.

**devcon /m:\\server01 stack =net *miniport***

In response, Devcon displays the expected stack for miniports on Server01.

```
ROOT\MS_L2TPMINIPORT\0000
    Name: WAN Miniport (L2TP)
    Setup Class: {4D36E972-E325-11CE-BFC1-08002BE10318} Net
    Controlling service:
        Rasl2tp
ROOT\MS_PPPOEMINIPORT\0000
    Name: WAN Miniport (PPPOE)
    Setup Class: {4D36E972-E325-11CE-BFC1-08002BE10318} Net
    Controlling service:
        RasPppoe
    Lower filters:
        NdisTapi
ROOT\MS_PPTPMINIPORT\0000
    Name: WAN Miniport (PPTP)
    Setup Class: {4D36E972-E325-11CE-BFC1-08002BE10318} Net
    Controlling service:
        PptpMiniport
    Lower filters:
        NdisTapi
ROOT\MS_PTIMINIPORT\0000
    Name: Direct Parallel
    Setup Class: {4D36E972-E325-11CE-BFC1-08002BE10318} Net
    Controlling service:
        Raspti
    Lower filters:
        PtiLink
4 matching device(s) found on \\Server01.
```

**status**
The classes operation lists the operational status of devices on the computer.

**Example 17: Display the status of all devices on the local computer.**
The following command finds the status of all devices on the local computer and saves the status in the status.txt file for logging or later review. The command uses the wildcard character (*) to represent all devices and the redirection character (>) to redirect the output to the status.txt file.

**devcon status * > status.txt**

**Example 18: Display the status of a particular device by instance ID**
The most reliable way to find the status of a particular device is to use the instance ID of the device.

The following command uses the instance ID to find the status of the I/O controller on the local computer. The command includes the instance ID of the device, PCI\VEN_8086&DEV_1130&SUBSYS_00000000&REV_02\3&29E81982&0&00. The at

character (@) prefixed to the ID identifies the string is an instance ID. The ID must be enclosed in quotation marks because it includes ampersand characters.

**devcon status "@PCI\VEN_8086&DEV_1130&SUBSYS_00000000&REV_02\3&29E81982&0&00"**

In response, Devcon displays the status of the I/O controller.

```
PCI\VEN_8086&DEV_1130&SUBSYS_00000000&REV_02\3&29E81982&0&00
    Name: Intel(R) 82815 Processor to I/O Controller - 1130
    Driver is running.
1 matching device(s) found.
```

**Example 19: Display the status of related devices on a remote computer.**
The following command displays the status of particular storage-related devices on Server01, a remote computer. It searches for the following devices:

- Disk drive, GenDisk
- CD-ROM drive, GenCdRom
- Floppy disk drive, FDC\GENERIC_FLOPPY_DRIVE
- Volumes, STORAGE\Volume
- Logical disk manager, ROOT\DMIO
- Volume manager, ROOT\FTDISK
- Floppy disk controller, ACPI\PNP0700

In the command, each ID is separated from the others by spaces. Note that GenDisk and GenCdRom are compatible IDs, while the other IDs are hardware IDs.

**devcon /m:\\server01 status GenDisk GenCdRom FDC\GENERIC_FLOPPY_DRIVE STORAGE\Volume ROOT\DMIO ROOT\FTDISK ACPI\PNP0700**

In response, Devcon displays the status of each device.

```
FDC\GENERIC_FLOPPY_DRIVE\1&3A2146F1&0&0
    Name: Floppy disk drive
    Driver is running.
IDE\CDROMSAMSUNG_DVD-ROM_SD-608_________________2.2_____\4&13B4AFD&0&0.0.0
    Name: SAMSUNG DVD-ROM SD-608
    Driver is running.
IDE\DISKWDC_WD204BA_____________________________16.13M16\4457572D414D373032313633393820
312
0202020
    Name: WDC WD204BA
    Driver is running.
ROOT\DMIO\0000
    Name: Logical Disk Manager
    Driver is running.
ROOT\FLOPPYDISK\0000
    Device has a problem: 28.
ROOT\FLOPPYDISK\0002
    Device has a problem: 01.
ROOT\FLOPPYDISK\0003
    Device has a problem: 01.
ROOT\FLOPPYDISK\0004
    Device is currently stopped.
ROOT\FTDISK\0000
    Name: Volume Manager
    Driver is running.
```

```
STORAGE\VOLUME\1&30A96598&0&SIGNATUREEA1AA9C7OFFSET1770DF800LENGTH3494AEA00
    Name: Generic volume
    Driver is running.
STORAGE\VOLUME\1&30A96598&0&SIGNATUREEA1AA9C7OFFSET7E00LENGTH1770CFC00
    Name: Generic volume
    Driver is running.
11 matching device(s) found on \\Server01.
```

**Search Operations**

find

findall

**find**

The find operation displays the instance ID and device description of devices installed on the computer that match the specified criteria.

**Example 20: Find devices by hardware ID pattern.**

The following command searches for mouse devices on Server01, a remote computer. Specifically, the command searches the Server01 computer for devices whose hardware ID or compatible ID includes "mou."

**devcon /m:\\Server01 find *mou***

In this case, Devcon found both two mouse devices.

```
ROOT\*PNP0F03\1_0_21_0_31_0                                    : Microsoft PS/2 Mouse
ROOT\RDP_MOU\0000                                              : Terminal Server Mouse
Driver
```

Because all of the Devcon display operations also find hardware IDs, you can use any display operation to search for hardware IDs. Select the operation based on the content that you need in the output. For example, to find the device drivers that mouse-related devices on a local computer use, submit the following command.

**devcon driverfiles *mou***

In response, Devcon finds the devices and lists their drivers.

```
HID\VID_045E&PID_0039\6&DC36FDE&0&0000
    Name: Microsoft USB IntelliMouse Optical
    Driver installed from c:\windows\inf\msmouse.inf [HID_Mouse_Inst]. 2 file(s) used
by d
river:
        C:\WINDOWS\System32\DRIVERS\mouhid.sys
        C:\WINDOWS\System32\DRIVERS\mouclass.sys
ROOT\RDP_MOU\0000
    Name: Terminal Server Mouse Driver
    Driver installed from c:\windows\inf\machine.inf [RDP_MOU]. 2 file(s) used by
driver:
        C:\WINDOWS\System32\DRIVERS\termdd.sys
        C:\WINDOWS\System32\DRIVERS\mouclass.sys
2 matching device(s) found.
```

**Example 21: Find devices by instance ID or class.**

The following commands display all legacy drivers on the local computer. Because legacy drivers do not have a hardware ID, you must search for them by their instance ID (registry path), ROOT\LEGACY, or their setup class, LegacyDriver.

The first command finds legacy drivers by an instance ID pattern. The ID pattern is prefaced by the at (@) symbol to indicate an instance ID and followed by the wildcard character (*) to find all devices in the ROOT\Legacy subkey.

**devcon find @root\legacy***

The second command finds legacy devices by searching for all devices in the legacydriver class.

**devcon find =legacydriver**

Both commands produce the same output, in this case, finding the same 27 legacy devices.

```
ROOT\LEGACY_AFD\0000                                        : AFD Networking Support
Environment
ROOT\LEGACY_BEEP\0000                                       : Beep
ROOT\LEGACY_DMBOOT\0000                                     : dmboot
ROOT\LEGACY_DMLOAD\0000                                     : dmload
ROOT\LEGACY_FIPS\0000                                       : Fips
ROOT\LEGACY_GPC\0000                                        : Generic Packet Classifier
ROOT\LEGACY_IPSEC\0000                                      : ipsec
ROOT\LEGACY_KSECDD\0000                                     : ksecdd
ROOT\LEGACY_MNMDD\0000                                      : mnmdd
ROOT\LEGACY_MOUNTMGR\0000                                   : mountmgr
ROOT\LEGACY_NDIS\0000                                       : ndis
ROOT\LEGACY_NDISTAPI\0000                                   : Remote Access NDIS TAPI
Driver
ROOT\LEGACY_NDISUIO\0000                                    : NDIS Usermode I/O
Protocol
ROOT\LEGACY_NDPROXY\0000                                    : NDProxy
ROOT\LEGACY_NETBT\0000                                      : netbt
ROOT\LEGACY_NULL\0000                                       : Null
ROOT\LEGACY_PARTMGR\0000                                    : PartMgr
ROOT\LEGACY_PARVDM\0000                                     : ParVdm
ROOT\LEGACY_RASACD\0000                                     : Remote Access Auto
Connection Driver
ROOT\LEGACY_RDPCDD\0000                                     : RDPCDD
ROOT\LEGACY_RDPWD\0000                                      : RDPWD
ROOT\LEGACY_TCPIP\0000                                      : tcpip
ROOT\LEGACY_TDPIPE\0000                                     : TDPIPE
ROOT\LEGACY_TDTCP\0000                                      : TDTCP
ROOT\LEGACY_VGASAVE\0000                                    : VgaSave
ROOT\LEGACY_VOLSNAP\0000                                    : VolSnap
ROOT\LEGACY_WANARP\0000                                     : Remote Access IP ARP
Driver
27 matching device(s) found.
```

**findall**

The findall operation finds all devices, including those that are represented in the registry, but are not present on the computer, because they have been moved or removed.

**Example 22: Find (and find all) devices in a setup class**

The following command finds all devices on the computer in the Net setup class. The equal sign (=) indicates that Net is a setup class and not an ID.

### devcon findall =net

In response, Devcon lists the following seven devices in the Net setup class. The first six are standard miniport devices. The seventh device, the RAS async adapter, is a software-enumerated device (SW*) that is not installed until it is needed.

```
PCI\VEN_10B7&DEV_9200&SUBSYS_00BE1028&REV_78\4&BB7B4AE&0&60F0: 3Com 3C920 Integrated
Fast
Ethernet Controller (3C905C-TX Compatible)
ROOT\MS_L2TPMINIPORT\0000                                   : WAN Miniport (L2TP)
ROOT\MS_NDISWANIP\0000                                      : WAN Miniport (IP)
ROOT\MS_PPPOEMINIPORT\0000                                  : WAN Miniport (PPPOE)
ROOT\MS_PPTPMINIPORT\0000                                   : WAN Miniport (PPTP)
ROOT\MS_PTIMINIPORT\0000                                    : Direct Parallel
SW\{EEAB7790-C514-11D1-B42B-00805FC1270E}\ASYNCMAC          : RAS Async Adapter
7 matching device(s) found.
```

The following command compares the find and findall commands by running a find command with the same parameters as the previous findall command.

### devcon find =net

In response, Devcon lists the following six devices in the Net setup class.

```
PCI\VEN_10B7&DEV_9200&SUBSYS_00BE1028&REV_78\4&BB7B4AE&0&60F0: 3Com 3C920 Integrated
Fast
Ethernet Controller (3C905C-TX Compatible)
```

```
ROOT\MS_L2TPMINIPORT\0000                                    : WAN Miniport (L2TP)
ROOT\MS_NDISWANIP\0000                                       : WAN Miniport (IP)
ROOT\MS_PPPOEMINIPORT\0000                                   : WAN Miniport (PPPOE)
ROOT\MS_PPTPMINIPORT\0000                                    : WAN Miniport (PPTP)
ROOT\MS_PTIMINIPORT\0000                                     : Direct Parallel
6 matching device(s) found.
```

The find command, which returns only currently present devices, does not list the software-enumerated device because it is not actually being used and software-enumerated devices are considered not-present if they are not being used.

**Change Operations**

classfilter

enable

disable

install

remove

restart

reboot

**classfilter**

Adds, deletes, displays, and changes the order of upper and lower filter drivers for a device setup class.

**Example 23: Display the filter drivers for a setup class.**

The following command displays the upper filter drivers for the DiskDrive setup class. Because this command does not include any classfilter operators, Devcon displays, but does not change, the filter drivers for the class.

**devcon classfilter DiskDrive upper**

In response, Devcon displays the new upper filter drivers for the DiskDrive class and confirms that it did not change them. In this case, the display shows that devices in the DiskDrive setup class use the PartMgr.sys (C:\Windows\System32\PartMgr.sys) upper filter driver.

```
Class filters unchanged.
    PartMgr
```

**Example 24: Add a filter driver to a setup class.**

The following command adds a fictitious filter, Disklog.sys, to the list of upper filter drivers for the DiskDrive setup class.

This command uses the add-after (**+**) classfilter operator to load the Disklog driver after the PartMgr driver so that it receives data that has already been processed by PartMgr.sys.

When the command starts, the virtual cursor is positioned before the first filter driver. Because it is not positioned on a particular driver, Devcon adds the Disklog driver to end of the filter driver list.

The command also uses the **/r** parameter, which reboots the system if it is necessary to make the class filter change effective.

**devcon /r classfilter DiskDrive upper +Disklog**

In response, Devcon displays the current upper filter drivers for the DiskDrive class.

```
Class filters changed. Class devices must be restarted for changes to take effect.
    PartMgr
    Disklog
```

If you misspell the driver name, or try to add a driver that is not installed on the system, the command fails. Devcon will not add a driver unless the driver is registered as a service, that is, unless the driver has an subkey in the Services registry subkey (HKLM\System\CurrentControlSet\Services).

The following command tests this safeguard feature. It adds "Disklgg," instead of Disklog, to the list of upper filters for the

DiskDrive class. The output demonstrates that the command fails.

**devcon /r classfilter DiskDrive upper +Disklgg**

```
devcon failed.
```

**Example 25: Insert a filter driver in the class list.**
The following command adds a fictitious filter driver, MyFilter.sys, to the list of upper filter drivers for the DiskDrive setup class. The command places MyFilter.sys between PartMgr.sys and Disklog.sys in the load order.

The following list shows the filter drivers for the DiskDrive class before the command is submitted.

```
    PartMgr
    Disklog
```

The first subcommand, **@Disklog**, uses the positioning operator (**@**) to put the virtual cursor on the Disklog filter driver. The second subcommand, **-MyFilter**, uses the add-before operator (**-**) to add MyFilter.sys before Disklog.sys.

The command also uses the **/r** parameter, which reboots the system if it is necessary to make the class filter change effective.

The positioning operator is essential in this example. Before Devcon processes any classfilter subcommands, the virtual cursor is at the beginning of the list and is not positioned on any filter drivers. If you use the add-before operator when the cursor is not on positioned on a driver, Devcon adds the driver to the beginning of the list. If you use the add-after operator when the cursor is not positioned on a driver, it adds the driver to the end of the list.

**devcon /r classfilter DiskDrive upper @Disklog -MyFilter**

In response, Devcon displays the current upper filter drivers for the DiskDrive class. ??

```
Class filters changed. Class devices must be restarted for changes to take effect.
    PartMgr
    MyFilter
    Disklog
```

You can also use the following command to add the MyFilter driver between PartMgr and Disklog. In this example, the first subcommand, **@PartMgr**, positions the virtual cursor on the PartMgr filter driver. The second subcommand, **+MyFilter**, uses the add after operator (+) to add MyFilter.sys after PartMgr.

**devcon /r classfilter DiskDrive upper @PartMgr +MyFilter**

**Example 26: Replace a filter driver**
The following command replaces the original copy of MyFilter.sys with a new and improved version, MyNewFilter.sys, in the list of filter drivers for the DiskDrive setup class.

The following list shows the filter drivers for the DiskDrive class before the command is submitted.

```
    PartMgr
    MyFilter
    Disklog
```

The first subcommand uses the delete operator (**!**) to delete MyFilter.sys from the list. This command deletes MyFilter from the list of upper filter drivers for the DiskDrive class. It does not affect the MyFilter.sys file in the C:\Windows\System32\Drivers directory.

The second subcommand uses the add-after operator (**+**) to place the new filter driver in the position that the deleted driver occupied. (The delete operator leaves the cursor on the position that the deleted filter occupied. Therefore, add-before (**-**) and add-after (**+**) have the same effect.)

The command also uses the **/r** parameter, which reboots the system if it is necessary to make the class filter change effective.

**devcon /r classfilter DiskDrive upper !MyFilter +MyNewFilter**

In response, Devcon shows the new class filter configuration for the DiskDrive class.

```
Class filters changed. Class devices must be restarted for changes to take effect.
    PartMgr
    MyNewFilter
    Disklog
```

**Example 27: Change the order of filter drivers**
The following command changes the order of filter drivers for the DiskDrive setup class. Specifically, it reverses the order of the second and third filter drivers.

The following list shows the filter drivers for the DiskDrive class before the command is submitted and the intended result of the command.

| Before | After |
|---|---|
| PartMgr | PartMgr |
| MyNewFilter | Disklog |
| Disklog | MyNewFilter |

The first subcommand uses the delete operator (**!**) to delete Disklog from the list. The second subcommand uses the start operator (**=**) to move the virtual cursor back to the starting position, and then uses the positioning operator (**@**) to put the cursor on the PartMgr driver. The start operator is necessary because the virtual cursor only moves forward through the list. The final subcommand uses the add-after operator (**+**) to add Disklog after PartMgr.

**devcon /r classfilter DiskDrive upper !Disklog =@PartMgr +Disklog**

In response, Devcon shows the new class filter configuration for the DiskDrive class.

```
Class filters changed. Class devices must be restarted for changes to take effect.
    PartMgr
    Disklog
    MyNewFilter
```

**enable**
The enable operation enables disabled devices on the system.

**Example 28: Enable a particular device.**
The following command enables a programmable interrupt controller that had been disabled to correct a system problem. Because the controller hardware ID, *PNP0000, includes an asterisk, the command uses the single quote character (') to direct Devcon to find the hardware ID precisely as specified in the command. Otherwise, the asterisk would be interpreted as a wildcard character.

**devcon enable '*PNP0000**

In response, Devcon displays the instance ID of the device, and explains that you must restart the system to enable the device.

```
ACPI\PNP0000\4&B4063F4&0                                    : Enabled on reboot
Not all of 1 device(s) enabled, at least one requires reboot to complete the operation.
```

You can respond by restarting the system, either manually, or by using the Devcon Syntax operation.

The following command adds the **/r** parameter to the previous command. The **/r** parameter reboots the system only when rebooting is required to complete an operation.

**devcon /r enable '*PNP0000**

In response, Devcon enables the device and then restarts the system to make the enabling effective.

When the system restarts, use a Devcon status command to confirm that the device is enabled.

**devcon status '*PNP0000**

```
ACPI\PNP0000\4&B4063F4&0
    Name: Programmable interrupt controller
    Driver is running.
```

**Example 29: Enable devices by class.**
The following command enables all printer devices on the computer by specifying the Printer setup class in the Devcon command. The command includes the **/r** parameter, which reboots the system if it is necessary to make the enabling effective.

**devcon /r enable =Printer**

In response, Devcon displays the instance ID of the printer that it found in the Printer class and reports that it is enabled. Although the command included the **/r** parameter, the system did not reboot because a reboot was not required to enable the printer.

```
LPTENUM\HEWLETT-PACKARDDESKJET_1120C\1&7530F08&0&LPT1.4        : Enabled
1 device(s) enabled.
```

**disable**
The disable operation disables the specified devices.

**Caution**

- If you have a USB keyboard or mouse installed on your computer, this test will disable or remove the device. To recover, you will need to use a PS/2 keyboard.

**Example 30: Disable devices by an ID pattern.**
The following command disables the USB devices on the local computer. It identifies the device by a hardware ID pattern, USB*. This pattern will match an device whose hardware ID or compatible ID begins with "USB." The command includes the **/r** parameter, which reboots the system if it is necessary to make the disabling effective.

**Caution**

- Before using an ID pattern to disable a device, determine which devices will be affected. To do so, use the pattern in a display command, such as **devcon status USB*** or **devcon hwids USB***.

**devcon /r disable USB***

In response, Devcon displays the instance IDs of the USB devices and reports that they are disabled. Although the command included the **/r** parameter, the system did not reboot because a reboot was not required to disable the devices.

```
USB\ROOT_HUB\4&2A40B465&0
: Disabled
USB\ROOT_HUB\4&7EFA360&0
: Disabled
USB\VID_045E&PID_0039\5&29F428A4&0&2
: Disabled
3 device(s) disabled.
```

**Example 31: Disable devices by instance ID.**
The following command disables the USB devices on the local computer. This command identifies the devices by their instance IDs (as indicated by the at character (@) that precedes each ID). Each instance ID is separated from the others by a space. Also, because the instance IDs include the ampersand character (&), they are enclosed in quotation marks. The command includes the **/r** parameter, which reboots the system if it is necessary to make the disabling effective.

**devcon /r disable "@USB\ROOT_HUB\4&2A40B465&0" "@USB\ROOT_HUB\4&7EFA360&0" "@USB\VID_045E&PID_0039\5&29F428A4&0&2"**

In response, Devcon displays the instance IDs of the USB devices and reports that they are disabled. Although the command included the **/r** parameter, the system did not reboot because a reboot was not required to disable the devices.

```
USB\ROOT_HUB\4&2A40B465&0
: Disabled
USB\ROOT_HUB\4&7EFA360&0
: Disabled
USB\VID_045E&PID_0039\5&29F428A4&0&2
: Disabled
3 device(s) disabled.
```

update

The update operation replaces the current device driver with the driver in the specified INF file.

**Example 32: Update the driver for communication ports**
The following command replaces the current device driver for communication ports on the system with the most suitable driver in the test.inf file. The command affects only only devices whose entire hardware ID is *PNP0501 (including the asterisk).

You can use this command to replace signed drivers on the system with alternate drivers for testing or troubleshooting, or to associate the devices with the newest version of the same drivers.

**Caution**

- Before replacing the driver for any device, determine which devices will be affected. To do so, use the pattern in a display command, such as **devcon hwids '*PNP0501** or **devcon driverfiles '*PNP0501**. This is especially important in an update operation, because Devcon does not list the devices affected by an update command.

**Note**

- Commands using the update operation only affect devices that have a hardware ID that is identical to the hardware ID specified in the command, not a pattern or partial match.

**devcon update test.inf *PNP0501**

In response, Devcon displays the following success message.

```
Updating drivers for *PNP0501 from test.inf.
```

To verify that the update was successful, use the Devcon Syntax operation, which lists the INF file and device driver files associated with the device.

The following command lists the driver files used by devices with a hardware ID of *PNP0501. In this command, which can use ID pattern matches, the string must be preceded by a single quote character (') to specify a literal match. Otherwise, the asterisk in the hardware ID is interpreted as a wildcard character.

**devcon driverfiles '*PNP0501**

In response, Devcon displays the driver files for the communication ports. The display shows that the devices are now associated with the devices specified in the test.inf file.

```
ROOT\*PNP0501\PNPBIOS_1
    Name: Communications Port (COM1)
    Driver installed from d:\test.inf [ComPort]. 2 file(s) used b
y driver:
        C:\WINDOWS\system32\DRIVERS\test.sys
        C:\WINDOWS\system32\DRIVERS\serenum.sys
ROOT\*PNP0501\PNPBIOS_2
    Name: Communications Port (COM2)
    Driver installed from d:\test.inf [ComPort]. 2 file(s) used b
y driver:
        C:\WINDOWS\system32\DRIVERS\test.sys
        C:\WINDOWS\system32\DRIVERS\serenum.sys
2 matching device(s) found.
```

**install**

The install operation installs the supporting software for a device.

**Example 33: Install a device**
The following command uses Devcon to install a Loopback device on the local computer. The command includes the full path to the INF file for the device (netloop.inf) and a hardware ID for the device, which typically is provided by the manufacturer (*MSLOOP).

### devcon /r c:\windows\inf\keyboard.inf *PNP030b

In response, Devcon reports that it installed the device by updating the driver files.

### Updating drivers for *PNPO30b from c:\windows\inf\keyboard.inf

**Example 34: Install a device using unattended setup**
The following example demonstrates how to install the MS Loopback Adapter during an unattended installation of Windows XP.

1. Add the following files to a floppy disk: Devcon.exe, Netset.exe, a tool in the Windows Server 2003 Resource Kit, and Netloop.inf (C:\Windows\inf\netloop.inf).

2. To the [GUIRunOnce] section of the unattended setup file, add the following Devcon command:

   **a:\devcon install a:\Netloop.inf *MSLOOP**

   This command installs the loopback adapter by using Netloop.inf and the hardware ID, *MSLOOP. The single quote character preceding "*MSLOOP" tells Devcon to interpret the asterisk as part of the hardware ID, not as a wildcard character.

3. Also, to the [GUIRunOnce] section of the unattended setup file, add a Netset command that specifies the desired network configuration settings.

4. Run the unattended installation.

**remove**
The remove operation uninstalls devices from the computer by deleting the device entry from the registry.

### Caution

- If you have a USB keyboard or mouse installed on your computer, this test will disable or remove the device. To recover, you will need to use a PS/2 keyboard.

**Example 35: Remove devices by instance ID pattern**
The following command removes all USB devices from the computer. It identifies the devices by an instance ID pattern that matches any instance ID (registry path) that begins with "USB\." The at character (@) distinguishes the instance ID from a hardware ID or compatible ID. The command also includes the **/r** parameter which reboots the system if it is required to make the remove procedure effective.

### Caution

- Before removing any devices by using a pattern, determine which devices are affected. To do so, use the pattern in a display command, such as **devcon status @usb*** or **devcon hwids @usb***.

### devcon /r remove @usb*

In response, Devcon displays the instance ID of the devices that it removed

```
USB\ROOT_HUB\4&2A40B465&0                                : Removed
USB\ROOT_HUB\4&7EFA360&0                                 : Removed
USB\VID_045E&PID_0039\5&29F428A4&0&2                     : Removed
3 device(s) removed.
```

**Example 36: Remove a particular network device**
The following command removes (uninstalls) the NDISWAN miniport from the local computer. The command specifies the Net class, and then refines the search by specifying devices in the class whose hardware ID or compatible ID include "ndiswan." The command also includes the **/r** parameter, which reboots the system if rebooting is required to make the remove procedure effective.

### Caution

- Before removing any devices by using a pattern, determine which devices will be affected. To do so, use the pattern in a display command, such as **devcon status =net *ndiswan*** or **devcon hwids =net *ndiswan***.

### devcon /r remove =net *ndiswan*

In response, Devcon displays the instance ID of the device that it removed.

ROOT\MS_NDISWANIP\0000 : Removed 1 device(s) removed.

rescan

The rescan operation scans the computer for new Plug and Play devices.

**Example 37: Scan the computer for new devices**
The following command scans the local computer for new devices.

### devcon rescan

In response, Devcon reports that it is scanning, but does not report finding any new devices.

```
Scanning for new hardware.
```

You can also use rescan on a remote computer. The following command runs the rescan operation on Server01, a remote computer, by adding the **/m** parameter to the command.

### devcon /m:\\server01 rescan

**restart**
The restart operation stops and then restarts devices on the local computer.

**Example 38: Restart a particular device**
The following command restarts the loopback adapter on the local computer. The command limits the search to the Net setup class and, within that class, specifies the instance ID of the loopback adapter, ROOT*MSLOOP\0000. The at character (@) identifies the string as an instance ID. The single quote character ('), which requests a literal search, prevents Devcon from interpreting the asterisk in the ID as a wildcard character.

### devcon restart =net @'ROOT*MSLOOP\0000

In response, Devcon displays the instance ID of the device and reports the result.

```
ROOT\*MSLOOP\0000                                                    : Restarted
1 device(s) restarted.
```

**reboot**
The reboot operation restarts the system on the local computer.

**Example 39: Restart the local computer**
The following command restarts the operating system on the local computer and associates the action to a hardware installation. Unlike the **/r** parameter, the **reboot** operation does not depend on the return code from another operation.

You can include the command in scripts and batch files that require the system to reboot.

### devcon reboot

In response, Devcon displays a message indicating that it is restarting the computer (**Rebooting local machine**.)

Devcon uses the standard ExitWindowsEx function to restart. If the user has open files on the computer or a program will not close, the system does not restart until the user has responded to system prompts to close the files or end the process.

**See Also**
**Concepts**
Devcon Overview
Devcon Remarks
Devcon Syntax
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Apmstat Overview

# Devcon Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devcon Syntax**

Devcon uses the following syntax:

**devcon** [**/m:***computer*] [**/r**] *operation* [*arguments*] | **devcon** {**/?** [*operation*] | **help** [*operation*]}

**Important**

- The parameters in a Devcon command must appear in the order shown in the syntax line. Otherwise, they are ignored without warning.

**Parameters**

**/m:\\ *computer***

Runs the command on the specified remote computer. The backslashes are required.

**/r**

Conditional restart. This parameter restarts the system after completing an operation only when a restart is required to make a change effective.

This parameter differs from the reboot operation, which forces the system to restart. Instead, the **/r** parameter determines whether a restart is required based on the return code from the accompanying operation.

**[ *operation*] /?**

Displays help. If you specify an operation, Devcon displays detailed help for the operation.

The parameters must appear in the specified order. For example, to display help for the status operation, type **devcon /? status** (or **devcon help status**), not **devcon status /?**.

**Operations**

There are three types of Devcon operations, those that display information, those that search for data, and those that manipulate the device or change its configuration.

The following operations *list and display* information about devices on the computer.

hwids

classes

listclass

driverfiles

drivernodes

resources

stack

status

find

findall

classfilter

enable

disable

update

install

remove

rescan

restart

reboot

**hwids**
Displays the hardware IDs, compatible IDs, and instance IDs of devices on the system. Valid on local and remote computers.

**Syntax**
**devcon** [**/m:***computer*] **hwids** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**
***/m:\\ computer***
Runs the command on the specified remote computer. The backslashes are required.

***
Represents all devices on the computer.

***ID***
Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

***= class***
Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- The **/m** parameter must precede the operation name (hwids). Otherwise, Devcon ignores the **/m** parameter and displays the hardware IDs of devices on the local computer without returning a syntax error.

- Because many Devcon operations require a hardware ID of a device, run a **devcon hwids *** command first to get the list of hardware IDs for subsequent operations.

**Sample Usage**
```
devcon hwids *

devcon /m:\\server01 hwids acpi*

devcon hwids acpi* *port*

devcon hwids =usb
```

**classes**
Lists all device setup classes, including classes that are not used by devices on the system. Valid on local and remote computers.

**Syntax**
**devcon** [**/m:***computer*]**classes**

**Parameters**

**/m:\\ *computer***
Runs the command on the specified remote computer. The backslashes are required.

**Remarks**

- The **/m** parameter must precede the operation name (classes). Otherwise, Devcon ignores the **/m** parameter and displays the classes on the local computer without returning a syntax error.
- To find the devices in a setup class, use the listclass operation. To find the setup class of a particular device, use the stack operation.

**Sample Usage**

```
devcon classes
```

```
devcon classes > setupclasses.txt
```

```
devcon /m:\\Server01 classes
```

**listclass**
Lists all of the devices in the specified device setup classes. Valid on local and remote computers.

**Syntax**
▸ **devcon**[**/m:***computer*]**listclass***class* [*class*...]

**Parameters**
**/m:\\ *computer***
Runs the command on the specified remote computer. The backslashes are required.

***class***
Specifies a device setup class. To find the device setup classes on the system, use the classes operation.

**Remarks**

- The **/m** parameter must precede the operation name (listclass). Otherwise, Devcon ignores the **/m** parameter and displays the devices on the local computer without returning a syntax error.
- The class name is not preceded by an equal sign (=) in the listclass operation syntax.
- Each entry in a setup class list represents one device. The entry consists of the unique instance name and a description of the device in *instance* **:** *description* format.
- To find the setup class of a particular device, use the stack operation.

**Sample Usage**

```
devcon listclass printers ports
```

```
devcon /m:\\Server01 listclass SmartCardReader
```

**driverfiles**
Displays the full path and file name of installed INF files and device drivers files for the specified devices. Valid only on the local computer.

**Syntax**
▸ **devcon driverfiles** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**
*****
Represents all devices on the computer.

***ID***
Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

***= class***

Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- The driverfiles operation runs only on the local computer.

**Sample Usage**

```
devcon driverfiles *

devcon driverfiles FDC\GENERIC_FLOPPY_DRIVE pci*

devcon driverfiles =media

devcon driverfiles =media isapnp*
```

**drivernodes**

The drivernodes operation lists all driver packages that are compatible with the device, along with their version and ranking. Valid only on the local computer.

**Syntax**

**devcon drivernodes** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**

***

Represents all devices on the computer.

***ID***

Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

***= class***

Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- The drivernodes operation runs only on the local computer.

**Sample Usage**

```
devcon drivernodes *

devcon drivernodes *miniport*

devcon drivernodes =netclient
```

```
devcon drivernodes =usb pci* usb*
```

**resources**
Lists the resources allocated to the specified devices. Resources are assignable, addressable bus paths, such as DMA channels, I/O ports, IRQ, and memory addresses. Valid on local and remote computers.

**Syntax**
▸ **devcon** [**/m:***computer*]**resources** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID*...]]}

**Parameters**
***/m:\\ computer***
Runs the command on the specified remote computer. The backslashes are required.

***
Represents all devices on the computer.

***ID***
Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

***= class***
Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- The **/m** parameter must precede the operation name (resources). Otherwise, Devcon ignores the **/m** parameter and displays the resources allocated to devices on the local computer without returning a syntax error.

**Sample Usage**
```
devcon resources *

devcon /m:\\server01 resources =media

devcon resources acpi* *port*
```

`devcon resources =class port*` (by class and hardware ID)

`devcon resources =class @port*`(by class and instance ID)

**stack**
Displays the expected driver stack for the specified devices. The stack operation also returns the GUID and the name of the device setup class of a device. Valid on local and remote computers.

A *driver stack* consists of the layered chain of drivers that support a device, such as upper and lower filters and controlling services. The display also includes the setup class of the device.

**Syntax**
▸ **devcon**[**/m:***computer*]**stack** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID*...]]}

**Parameters**
***/m:\\ computer***
Runs the command on the specified remote computer. The backslashes are required.

***
Represents all devices on the computer.

***ID***

Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

**= *class***

Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- The **/m** parameter must precede the operation name (stack). Otherwise, Devcon ignores the **/m** parameter and displays the stack of device drivers on the local computer without returning a syntax error.
- To investigate a device problem, compare the expected driver stack from the stack operation with the drivers installed for the device as displayed by the driverfiles operation, and the actual drivers used by the device available through the kernel debugger.

**Sample Usage**

```
devcon /m:\\Server01 stack * > Server01Stack.txt

devcon stack ISAPNP\ReadDataPort

devcon /m:\\Server01 stack pci*

devcon stack =multifunction
```

**status**

Displays the status (running, stopped, disabled) of the driver for devices on the computer. Valid on local and remote computers.

**Syntax**

**devcon**[**/m:***computer*]**status** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**

***/m:\\ computer***

Runs the command on the specified remote computer. The backslashes are required.

***

Represents all devices on the computer.

***ID***

Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

***= class***

Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- The **/m** parameter must precede the operation name (status). Otherwise, Devcon ignores the **/m** parameter and displays the status of device drivers on the local computer without returning a syntax error.

**Sample Usage**

```
devcon /m:\\Server01 status *

devcon status pci*

devcon status "PCI\VEN_115D&DEV_0003&SUBSYS_0181115D"

devcon status =printer
```

The following operations *search* for information about devices on the computer.

**find**

Finds all devices that are currently attached to the computer. Displays the instance ID and device description. Valid on local and remote computers.

**Syntax**

**devcon**[**/m:***computer*]**find** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**

***/m:\\ computer***

Runs the command on the specified remote computer. The backslashes are required.

*******

Represents all devices on the computer.

***ID***

Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

***= class***

Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- The **/m** parameter must precede the operation name (find). Otherwise, Devcon ignores the **/m** parameter and searches the local computer without returning a syntax error.

- The find operation with a single class argument is the same as a listclass operation.

**Sample Usage**

```
devcon find *

devcon find =media *pnp*
```

```
devcon /m:\\Server01 find *mou*

devcon find @*hub*
```

**findall**
Finds all devices on the computer, including devices that were once attached to the computer, but have been detached or moved (known as nonpresent or phantom devices). Findall also find devices that are enumerated differently as a result of a BIOS change. Valid on local and remote computers.

**Syntax**
**devcon** [**/m:***computer*]**findall** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**
***/m:\\ computer***
Runs the command on the specified remote computer. The backslashes are required.

***
Represents all devices on the computer.

***ID***
Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

***= class***
Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- The **/m** parameter must precede the operation name (findall). Otherwise, Devcon ignores the **/m** parameter and searches the local computer without returning a syntax error.

- To find devices that could be installed, but are not installed compare the output of find and findall operations.

**Sample Usage**
```
devcon resources STORAGE\Volume

devcon resources =ports lpt*

devcon resources @pci*
```

The following operations manipulate the device or change its configuration.

**classfilter**
Adds, deletes, displays, and changes the order of filter drivers for a device setup class. Valid only on the local computer.

**Syntax**
**devcon classfilter** *class*{**upper** | **lower**}[{**=** | **@** | **-** | **+** | **!**}*driver*[{**=** | **@** | **-** | **+** | **!**}*driver* ...]]

## Notes

- A classfilter command includes one or more subcommands consisting of an operator (=, @, -, +, !) and a filter driver name. Devcon executes the subcommands in the order that they appear in the command.

- The classfilter operation uses a virtual cursor to move through the list of drivers in the filter. The cursor starts at the

beginning of the list, before the first driver in the list. Unless returned to the starting position, the cursor always moves forward through the driver list as Devcon executes the subcommands.

- Without operators, "Devcon classfilter *class* {**upper | lower**}" displays the upper or lower filter drivers in the specified class.

**Parameters**

***class***
Specifies the device setup class.

**upper| lower**
Determines whether upper or lower filter drivers are displayed.

**Operators**

**=**
Moves the cursor to the beginning of the filter list (before the first driver).

**@ *driver***
Positions the cursor on the next instance of the specified driver.

**- *driver***
Add before. Inserts the specified driver before the driver on which the cursor is positioned.

If the cursor is not positioned on a driver, Devcon inserts the specified driver at the beginning of the list. When the subcommand completes, the cursor is positioned on the newly added driver.

**+ *driver***
Add after. Inserts the specified driver after the driver on which the cursor is positioned.

If the cursor is not positioned on a driver, Devcon inserts the specified driver at the end of the list. When the subcommand completes, the cursor is positioned on the newly added driver.

**! *driver***
Deletes the next occurrence of specified driver from the filter.

When the subcommand completes, the cursor occupies the position of the deleted driver. Subsequent + or - subcommands will insert a new driver at the cursor position.

**Remarks**

- Because filter driver changes require that the devices be restarted, include the **/r** parameter in the classfilter command.

- To list the device setup classes on the computer, use the classes operation. To list the devices in a device setup class, use the listclass operation. To list the class filters for a device, use the stack operation.

**Sample Usage**

**devcon classfilter mouse upper**

**devcon /r classfilter mouse upper !mouclass +newmou**

**devcon /r classfilter net lower @netfltr -testfltr**

**devcon /r classfilter volume upper !volsnap =!volsnap2**

**enable**
Enables devices on the computer. Valid only on the local computer.

When you enable a device, its driver is loaded into memory and the device is ready for use. Devcon reports a failure when it cannot find or cannot enable a device. It reports a success even when the device is already enabled. Before and after enabling a device, use the **devcon status** operation to verify the device specification and status.

**Syntax**

▸ **devcon enable** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**

***

Represents all devices on the computer.

***ID***

Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

**= *class***

Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- You might need to restart the device or the system to make the command effective. If so, add the restart parameter **/r** to the command.
- An **enable** command does not fail when the device is already enabled.
- Before and after enabling a device, use the status operation to verify the device status.

**Sample Usage**

`devcon enable *` (not recommended)

```
devcon /r enable *DVD-ROM*
```

```
devcon /r enable =printer
```

**disable**

Disables devices on the computer. Valid only on the local computer.

When you disable a device, the device remains physically connected to the computer, but its driver is unloaded from memory and its resources are freed, so the device cannot be used. Devcon reports a failure when it cannot find or cannot disable a device. It reports a success even when the device is already disabled. Before and after disabling a device, use the **devcon status** operation to verify the device specification and status.

**Syntax**

**devcon disable** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**

***

Represents all devices on the computer.

***ID***

Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |

| | |
|---|---|
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

**= *class***

Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- You might need to restart the system to make the command effective. If so, add the restart parameter **/r** to the command.
- A **disable** command does not fail when the device is already disabled.
- Before and after disabling a device, use the status operation to verify the device status.

**Sample Usage**

`devcon disable *` (not recommended)

```
devcon /r disable *DVD-ROM*
```

```
devcon /r disable =printer
```

**update**

Replaces the current device driver with the driver specified in the INF file. Valid only on the local computer. If Devcon cannot find the device, or if it cannot update it, it reports a failure. To isolate the problem, use a Devcon **find** command.

**Syntax**

**devcon update** *INFfile HardwareID*

**Parameters**

**INFfile**

Specifies the full path and file name of the INF (information) file for the device.

**HardwareID**

Specifies a hardware identifier for the device.

**Remarks**

- The hardware ID specified must exactly match the hardware ID of the device. Patterns are not valid. Do not type a single quote character (') before hardware IDs that include an asterisk.
- Because you might need to restart the system to make the command effective, add the restart parameter **/r** to the command.

**Sample Usage**

```
devcon update c:\windows\inf\newdvc.inf ISAPNP\CSC4324\0
```

```
devcon /r update c:\windows\inf\newdvc.inf *PNP030b
```

**install**

Installs the supporting software for a device. Valid only on the local computer. **ls/** can only be used to install devices that have no hardware (eg MS Loopback adapter) or non-PnP devices (eg non PnP ISA sound card).

**Syntax**

**devcon install** *INFfile HardwareID*

**Parameters**

**INFfile**

Specifies the full path and file name of the INF (information) file for the device. INF files contain the information that the system needs to install the software for a device.

**HardwareID**

Specifies a hardware identifier for the device.

For the syntax rules for hardward IDs, see "Hardware ID" in the Windows DDK Glossary.

**Remarks**

- You might need to restart the device or the system to make the command effective.
- The success message for the install operation says that Devcon has updated the drivers for the device.

**Sample Usage**

```
devcon install c:\windows\inf\newdvc.inf ISAPNP\CSC4324\0

devcon /r install c:\windows\inf\newdvc.inf ISAPNP\CSC4324\0
```

**remove**

Uninstalls the specified devices. Valid only on the local computer.

**Syntax**

▸ **devcon remove** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**

*****

Represents all devices on the computer.

***ID***

Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

**= *class***

Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- Because you might need to restart the system to make the command effective, add the restart parameter **/r** to the command.
- The **rescan** operation can also be used to reinstall a device that was removed.

**Sample Usage**

```
devcon /r remove "PCI\VEN_8086&DEV_7110"

devcon /r remove =printer

devcon /r remove =printer *deskj*
```

**rescan**

Directs Windows Plug and Play to search the computer for newly installed Plug and Play devices. Valid on local and remote computers.

**Syntax**

▸ **devcon** [**/m:***computer*]**rescan**

**Parameters**

**/m:\\ *computer***

Runs the command on the specified remote computer. The backslashes are required.

**Remarks**

- The **/m** parameter must precede the operation name (rescan). Otherwise, Devcon ignores the **/m** parameter and scans the local computer without returning a syntax error.

**Sample Usage**

```
devcon rescan

devcon /m:\\Server01 rescan
```

**restart**

Stops and restarts the specified devices. Valid only on the local computer. If Devcon cannot find a device, or if it cannot restart it, it reports a failure. To isolate the problem, use a Devcon **find** command.

**Syntax**

▸ **devcon restart** {***** | *ID* [*ID* ...] | **=***class* [*ID* [*ID* ...]]}

**Parameters**

***

Represents all devices on the computer.

***ID***

Specifies all or part of a hardware ID, compatible ID, or instance ID of a device. When specifying multiple IDs, separate each ID with a space. IDs that include an ampersand character (&) must be enclosed in quotation marks.

The following special characters modify the ID parameter.

| | |
|---|---|
| * | Matches any character or characters. Use the wildcard character (*) to create an ID pattern. For example, *disk*. |
| @ | Indicates an instance ID. For example, @ROOT\FTDISK\0000. |
| ' | Matches the string literally (exactly as it appears). Use it to indicate that an asterisk is part of the ID name, and is not a wildcard character. For example, '*PNP0600, where *PNP0600 (including the asterisk) is the hardware ID. |

**= *class***

Specifies the device setup class of the devices. To find the setup class of a device, use the classes operation.

You can also specify hardware IDs, compatible IDs, or instance IDs or ID patterns following the class name. Separate each ID or pattern with a space. Devcon finds devices in the class that match the specified IDs.

**Remarks**

- Because you might need to restart the system to make the command effective, add the restart parameter **/r** to the command.

**Sample Usage**

```
devcon restart *

devcon restart pci*

devcon restart "PCI\VEN_115D&DEV_0003&SUBSYS_0181115D"

devcon restart =printer

devcon restart =printer *desk*
```

**reboot**

Stops and then starts the operating system. Valid only on the local computer.

**Syntax**

▸ **devcon reboot**

**Remarks**

- Unlike the **/r** parameter, which restarts the system only when required to make a change effective, the reboot operation restarts the system without determining whether a restart is required.

- The reboot operation does not force a restart. If the system cannot close a program or if a user file is not saved, the restart cannot complete.

**Sample Usage**

```
devcon reboot
```

**See Also**

**Concepts**

Devcon Overview
Devcon Remarks
Devcon Examples
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Apmstat Overview

# Dfsutil Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dfsutil.exe: Distributed File System Utility**
This tool enables administrators to query and troubleshoot the Microsoft Distributed File System (DFS) from the command prompt. It can be used to perform maintenance of a DFS root and to clean up metadata left behind when domain-based DFS roots are orphaned or abandoned.

### Note

- You must run DFSUtil from the command window.

**Corresponding UI**
DFSUtil provides command-line functionality that parallels much of the Distributed File System Manager snap-in.

**Concepts**
Microsoft Distributed File System (DFS), a feature of Windows Server 2003, makes it easier to access and manage files that are physically distributed across a network. By using DFS, you can make files that are distributed across multiple servers appear to users as if they reside in one place on the network. Users and applications no longer need to know and specify the actual physical location of files in order to gain access to them: DFS makes the physical storage of data transparent by mapping it to a logical representation.

DFS does for servers and shares what file systems do for hard disks. File systems provide uniform named access to collections of sectors on disks; DFS provides a uniform naming convention and mapping for collections of servers, shares, and files. Thus, by using DFS you can organize file servers and their shares into a logical hierarchy, which makes it considerably easier to manage and use extensive information resources. In addition, DFS is not limited to a single file protocol. It can support the mapping of servers, shares, and files, regardless of the file client being used, provided that the client supports the native server and share.

In the past, with the Universal Naming Convention (UNC), to access file information a user or application was required to specify the physical server and share; that is, the user or application had to specify *Server**Share**Path**FileName*).

**System Requirements**
The following are the system requirements for DFSUtil:

- Windows Server 2003

**File Required**

- Dfsutil.exe

**See Also**
**Concepts**
Dfsutil Remarks
Dfsutil Syntax
Dfsutil Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview

Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Dfsutil Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DfsUtil Remarks**

DFSUtil Parameter Usage (Administration, Debugging/Monitoring, Troubleshooting/Disaster Recovery)

The following is an overview of when an administrator would typically use a specific DfsUtil parameters, as well as the parameters' general area of impact.

| Parameter | Use for Routine Administration | Use for Debugging And Monitoring | Use for Troubleshooting And Disaster Recovery | Impact: View (Passive) | Impact: Server Modification (Active) | Impact: Client Modification (Active) |
|---|---|---|---|---|---|---|
| AddStdRoot | yes | - | - | - | DFS Namespace modification | - |
| CheckBlob | - | yes | - | Verify blob stored in a file | - | - |
| Clean | - | - | yes | - | DFS Namespace modification | - |
| DisplayBlob | - | yes | - | Display blob contents to standard output | - | - |
| Export | yes | yes | yes | Save DFS namespace to a file | - | - |
| ExportBlob | - | yes | - | Export blob to a file | - | - |
| Import | yes | - | yes | - | DFS Namespace modification | - |
| ImportRoot | - | - | yes | - | DFS Namespace modification | - |
| Insite | yes | - | - | - | DFS Operations, Client behavior | - |
| PKTFlush | - | - | yes | - | - | Client behavior |
| PKTInfo | - | yes | - | Display target referrals and active target for a link (Client) | - | - |
| PurgeMupCache | - | - | yes | - | - | Client behavior |

| | | | | | | |
|---|---|---|---|---|---|---|
| PurgeWin2kStaticSiteTable | - | - | yes | - | DFS metadata site information | - |
| RemFtRoot | yes | - | - | - | DFS Namespace modification | - |
| RemoveReparse | - | - | - | - | DFS related directories on the server | - |
| AddFTRoot | - | - | yes | - | DFS related directories on the server | - |
| RemStdRoot | yes | - | - | - | DFS Namespace modification | - |
| RenameFtRoot | yes | - | - | - | DFS Namespace modification | - |
| RootScalability | yes | - | - | - | DFS Sensitivity | - |
| ShowWin2kStaticSiteTable | yes | yes | yes | Display Windows 2000 Server DFS site information when in mixed mode | - | - |
| SiteCosting | yes | - | - | - | DFS Operations, Client behavior | - |
| SiteName | yes | yes | - | View Site name of a computer | - | - |
| SPCFlush | - | - | yes | - | - | Client behavior |
| SPCInfo | - | yes | - | Display trusted domains and its DCs known on client | - | - |
| UnmapFtRoot | - | - | yes | - | DFS Namespace modification | - |
| UpdateWin2kStaticSiteTable | yes | - | yes | - | DFS metadata site information | - |
| View | - | yes | - | View DFS Namespace | - | - |

| | | | | | | |
|---|---|---|---|---|---|---|
| ViewDfsDirs | - | yes | yes | Displays DFS related directories in a volume | Displays DFS related directories in a volume | - |

**Output Description for /PktInfo**

Current Partition Knowledge Table (PKT) information can be displayed by using the /pktinfo parameter. The following describes the type of output produced when /pktinfo is used. The description numbers correspond to the numbers in the sample output.

1. The initial entry specifies how many entries are included in the output.

2. Every entry has both a full (**Entry**) and a short (**ShortEntry**) name.

3. **Expires** or Time To Live (TTL) is the length of time a DFS client stores referral information from the PKT in its local cache. The DFS client requests new information from the PKT if the TTL expires or the client is restarted. The TTL is reset if the shared folder is accessed before it expires.

4. **UseCount** is the number of times it's used since got this referral. **Type** specifies the type of referral being made.

5. Each entry contains a sequentially numbered target list. A referral is a list of targets a DFS client receives from the PKT when a user attempts to access a root or a link in the namespace. Referrals are cached locally on the client for a time specified by time indicated in **Expires**. The referral information can tell you which target is active for your DFS client. The number after each target name indicates the state of the target. Attempts to access an inactive target will be referred to the currently active target.

- State 0x21:inactive or not responding
- State 0x31:responding and an active target

6. An **ACTIVE** target status indicates it is currently being accessed by the client. If access to this target fails, the client will try to access the next target in the list. In this entry, the client would try to access the 12th target, \dom-dc-12.dom.company.com\SysVol.

**dfsutil /pktinfo**

```
Microsoft(R) Windows(TM) Dfs Utility Version 4.0
Copyright (C) Microsoft Corporation 1991-2001. All Rights Reserved.

--mup.sys--
```

**1.**

```
4 entries...
```

**2.**

```
Entry: \dom.company.com\SysVol
ShortEntry: \dom.company.com\SysVol
```

**3.**

```
Expires in 3014020 seconds
```

**4.**

```
UseCount:  Type:0x1 ( DFS )
```

**5.**

```
0:[\dom-dc-04.dom.company.com\SysVol] State:0x21 ( )
1:[\dom-dc-08.dom.company.com\SysVol] State:0x21 ( )
2:[\dom-dc-10.dom.company.com\SysVol] State:0x21 ( )
3:[\dom-dc-07.dom.company.com\SysVol] State:0x21 ( )
4:[\dom-dc-03.dom.company.com\SysVol] State:0x21 ( )
5:[\dom-dc-02.dom.company.com\SysVol] State:0x21 ( )
6:[\dom-dc-05.dom.company.com\SysVol] State:0x21 ( )
7:[\dom-dc-09.dom.company.com\SysVol] State:0x21 ( )
8:[\dom-dc-06.dom.company.com\SysVol] State:0x21 ( )
```

```
9:[\dom-dc-13.dom.company.com\SysVol] State:0x21 ( )
10:[\dom-dc-11.dom.company.com\SysVol] State:0x21 ( )
```

**6.**

```
11:[\dom-dc-01.dom.company.com\SysVol] State: 0x31 ( ACTIVE )
12:[\dom-dc-12.dom.company.com\SysVol] State:0x21 ( )

Entry: \dom.company.com\netlogon
ShortEntry: \dom.company.com\netlogon
Expires in 0 seconds
UseCount: 0 Type:0x1 ( DFS )
0:[\dom-dc-08.dom.company.com\netlogon] State:0x31 ( ACTIVE )
1:[\dom-dc-05.dom.company.com\netlogon] State:0x21 ( )
2:[\dom-dc-07.dom.company.com\netlogon] State:0x21 ( )
3:[\dom-dc-03.dom.company.com\netlogon] State:0x21 ( )
4:[\dom-dc-09.dom.company.com\netlogon] State:0x21 ( )
5:[\dom-dc-10.dom.company.com\netlogon] State:0x21 ( )
6:[\dom-dc-04.dom.company.com\netlogon] State:0x21 ( )
7:[\dom-dc-11.dom.company.com\netlogon] State:0x21 ( )
8:[\dom-dc-02.dom.company.com\netlogon] State:0x21 ( )
9:[\dom-dc-13.dom.company.com\netlogon] State:0x21 ( )
10:[\dom-dc-06.dom.company.com\netlogon] State:0x21 ( )

Entry: \dom\public
ShortEntry: \dom\public
Expires in 0 seconds
UseCount: 0 Type:0x81 ( REFERRAL_SVC DFS )
0:[\ntdsdfs2\Public] State:0x19 ( ACTIVE )
1:[\ntdsdfs1\Public] State:0x09 ( )

Entry: \dom\public\Desktop\Office2000
ShortEntry: \dom\public\Desktop\Office2000
Expires in 0 secondsUseCount: 0 Type:0x1 ( DFS )
0:[\imself-ia64-02\office2000] State:0x31 ( ACTIVE )
1:[\apzawsrv2ß\office2000] State:0x21 ( )
2:[\apzawsrv3\office2000] State:0x21 ( )

DfsUtil command completed successfully.
```

**Output Description for /SpcInfo**

SPC table information can be displayed by using the /spcinfo parameter. The following describes the type of output produced when /spcinfo is used. The description numbers correspond to the numbers in the sample output.

1. Entries beginning with [*] are not expanded, and are local to the domain.

2. Entries beginning with [+] are expanded into the DC list due to access.

3. Indented entries are expanded entries of the domain entry of the domain1.companydom.com domain.

4. Unexpanded domains which have not been accessed or required to expand by the client. The local client cache contains no further information about these domains. Only trusted domains are listed.

### dfsutil /spcinfo

```
Microsoft(R) Windows(TM) Dfs Utility Version 4.0
Copyright (C) Microsoft Corporation 1991-2001. All Rights Reserved.
```

**1.**

```
[*] [domain1-dc-08.domain1.companydom.com]
[*][DOMAIN1]
[*][domain1.companydom.com]
```

**2.**

```
[+] [domain1.companydom.com]
```

**3.**

```
[-domain1-dc-05.domain1.companydom.com]
```

```
[-domain1-dc-13.domain1.companydom.com]
[-domain1-dc-04.domain1.companydom.com]
[-domain1-dc-rn.ntwksta.domain1.companydom.com]
[-domain1-dc-02.domain1.companydom.com]

[+][DOMAIN1]
[-DOMAIN1-DC-07]
[-NTDEV-DC-11]
[-NTDEV-DC-05]
[+NTDEV-DC-08]
```

**4.**

```
[-][stressdsdom.nttest.companydom.com]
[-][STRESSDSDOM]
[-][vbl03stressdom.nttest.companydom.com]
[-][VBL03STRESSDOM]

DfsUtil command completed successfully.
```

**See Also**

**Concepts**

Dfsutil Overview
Dfsutil Syntax
Dfsutil Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Dfsutil Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DFSUtil Examples**
**Example 1: Control a DFS Client's Ability to Link to Sites**
**Enabling** the insite setting of a DFS server is useful when:

- You don't want the DFS clients to connect outside the site.
- You don't want the DFS client to connect to a site other than the site it is in, and hence avoid using expensive WAN links.

```
dfsutil /insite:\\example.com\dfsroot /enable
```

After using this command statement, clients will not get any referral for a replica outside the dfsroot site. This means that if the Replica sets in the client Site are down, the client will not do a failover to a Replica set in another site.

**Disabling** the insite setting of a DFS server is useful when:

- You want to enable outside site referrals.

**Note**

- If you want your DFS clients to be able to link to outside sites when no local server is available, and the DFS clients never seem to link outside the site, it may be because connectivity has been limited to an internal site using the **/insiteenable** setting. Disabling this setting will restore the ability of clients to link outside the site.

To reset your site preferences, type the following at the command prompt:

```
dfsutil /insite:\\example.com\Sales /disable
```

After using this command statement, clients will also get referrals for servers outside the example.com site. This means that, if all the servers in the client Site are down, the client will do a failover to a server in another site. Be aware that this may require traversing through expensive WAN links.

**Example 2: Configure a DFS Server to be Site Cost Aware**
You want DFS clients to be able to connect outside the internal site, but you want clients to connect to the closest site first, saving the expensive network bandwidth. You want to maintain *high availability* as a priority, but obviously you want DFS clients to connect to closer sites rather than farther sites when the former are reachable and up. To configure the server to be site cost aware, type either of the following statements at the command line:

```
dfsutil /sitecosting:\\example.com\sales /enable
```

```
dfsutil /root:\\example.com\sales /sitecosting /enable
```

Now the server sends the referral list composed of the randomly ranked targets in the same site as the client, followed by the targets in the next closest site from the site in which the client resides, followed by targets in the second closest site, and then the third and so on.

Notes

- If you do not know if a root is site costing aware, you can check its status by substituting the **/display** parameter for the **/sitecosting** parameter.
- If you want to disable site costing, substitute the **/disable** parameter for the **/enable** parameter.
- If you want additional insight into the processing behind this command, append the **/verbose** parameter to the end of either command statement.
- Do not use the **/sitecosting** parameter in conjunction with the **/api** parameter. Site costing cannot be done in the API

mode.

**Example 3: Back up the DFS Namespace**
You want to back up the DFS namespace for a specified root so that you can restore it later in case of system crash and loss of namespace from the system. Backing up namespace information is especially important when you have large namespaces. Using a single command statement per root, you can back up the namespaces into simple files. The files are in an XML format. To back up a namespace, type the following at the command line:

```
dfsutil /root:\\example.com\sales /export:c:\NameSpaceBackups\DateToday\saRoot20.txt
```

Note

- The output of the export file is in XML in Windows Server 2003 DFS. This means that, if your current DFS is a prior version and the output file is coming from a prior version, it should be converted to the XML format used by the **/import** parameter.

**Example 4: Restore the DFS Namespace from a Back Up**
Your system has crashed and you have lost your namespace data. In order to restore the namespace, type the following at the command line:

```
dfsutil /root:\\example.com\sales /import:c:\NameSpaceBackups\DateToday\saRoot20.txt /set
```

Notes

- It is often advisable to use the **/verify** option before doing the restore to assess the changes that will be made by the import using either the **merge** or **set** parameters. To use the **/verify** parameter, substitute it for the **/set** parameter.

- If you backed up the namespaces of the specified root from an earlier version and not Windows Server 2003, then this restore will not work.

**Example 5: Copy a Namespace**
You want to set up a backup server with the same namespace configuration as your main DFS server. This means that even if you have to take your main DFS server offline or if it goes down, the backup server can act as the DFS server. To copy one namespace to another, type the following at the command line:

```
dfsutil /root:\\example.com\sales /importroot:\\example1.com\sales1 /mirror
```

To verify that the namespace configuration has been mirrored on the backup server, type the following at the command line:

```
dfsutil /root:\\example.com\sales /importroot:\\example1.com\sales1 /compare /verbose
```

### Note

- By default, DFSUtil works in direct mode to get or update the DFS metadata. If you want DFSUtil to work in the API mode, use the **/api** parameter.

**Example 6: Force a DFS Client to Get a Target Referral List from the Server**
A client benefits from *high availability* in DFS by getting a list of link target referrals within the same site as well as targets in farther sites. In some cases targets in the closer sites may be inaccessible at the beginning of the client's use, causing the client to successfully failover to a target at a farther site. Once a closer and less expensive target is available, you would like the client to use it. If you do not want to reboot the client to cause a closer site to be selected, type the following at the command line:

```
dfsutil /pktflush
```

This command statement flushes the local partition knowledge table (PKT). This forces the client to get the referral list of the targets from the server again. Some of the entries in the PKT may not get flushed, especially if DFS is in the process of using the referrals. Once the PKT is flushed from the client cache, the client gets a new list of referrals from the server and it surely will try accessing the closer targets.

Notes

- There is no guarantee that the client select a desired target in the same site when there are other equivalent candidates available in the same site.

- Rebooting the client can replace the need of using the **/pktflush** parameter.

**Example 7: Use Windows 2000 DFS with Windows Server 2003 DFS**
You have a domain-based DFS with at least one member server that is a Windows 2000 DFS and at least one that is a Windows Server 2003 DFS. This could happen as a result of using site costing in mixed-mode DFS. It could also result from upgrading from a D-DFS with all member servers as Windows 2000 to one with all member servers as Windows Server 2003. To view what site information is currently kept on the Active Directory (blob), type the following at the command line:

```
dfsutil /root:\\example.com\Sales /showwin2kstaticsitetable
```

Now if you create links or roots on this DFS using one of the Windows Server 2003 tools, it never writes the Site Information to the DFS blob because Windows Server 2003 DFS discovers it dynamically. For Windows 2000 to function correctly, however, it needs to have this information on the blob it reads. To write links and roots to the blob which will allow Windows 2000 DFS to function properly, type the following at the command line:

```
dfsutil /root:\\example.com\Sales /updatewin2kstaticsitetable
```

This command statement updates the site information needed for Windows 2000 DFS to function properly. Repeat this command statement after every addition, deletion, and modification in a link or root on the DFS.

### Notes

- The Windows 2000 DFS uses an Active Directory object to learn about the site information in the DFS blob which in turn is written by the DFS administration tools when the roots and links were created. Windows Server 2003 does not read site information; it determines site information dynamically.

- The information kept on the blob is static in nature; it is not updated automatically on any changes in DFS.

**Example 8: Optimize Active Directory Metadata for Windows Server 2003 DFS**
Once you have completed your upgrade of DFS from Windows to Windows Server 2003 and all your DFS roots become Windows Server 2003 servers, the site information details stored in Active Directory are useless. It is good practice to keep the DFS blob optimal for best results and scalability of DFS. To remove this site information from the blob, type the following at the command line:

```
dfsutil /root:\\example.com\Sales /purgewin2kstaticsitetable
```

This ensures that the DFS blob will not contain any unnecessary site information.

**Example 9: Clean Up a Root-Server Entry in Active Directory**
If, for some reason, you have removed a root server and the metadata in Active Directory has not been updated, you can manually cause the root-server entry to be removed from the blob. This parameter is used as a problem/corruption repair setting. Do not confuse it with parameters like **/remftroot** or **/clean**. Unless absolutely necessary, use of this parameter should be avoided. To eliminate the resultant obsolete root target in the DFS namespace, type the following at the command line:

```
dfsutil /unmapftroot /root:\\ex.com\salesdata /server:\\mainshare1
```

### Caution

- Using this option deletes an entire namespace. This option is used **only** when you already have removed a domain-based root server. When this command statement is executed, the actual server is never contacted before it unmapped (because it is assumed to have been removed and, as such, would not exist). As a result, a functioning root can get removed accidentally.

Notes

- The root-server that has been removed must be domain-based for this method to work.

- DFSUtil works in direct mode to get or update the DFS metadata. This command statement will fail if the root-server is not domain-based or if direct mode cannot be initialized.

- Back up your namespace before using the **/unmapftroot** parameter.

**See Also**

**Concepts**

Dfsutil Remarks
Dfsutil Syntax
Dfsutil Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Dfsutil Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DFSUtil Syntax**

DFSUtil uses the following syntax:

**dfsutil** [**/?**] [*parameters*]

**Parameters**

/api

/verbose

/? | /command /?

/root:

/server:

/share:

/domain:dfsDomainName

/newDomain:dfsDomainName

/oldDomain:dfsDomainName

/comment:

/AddFTRoot:

/AddStdRoot:

/RemFTRoot

/RemStdRoot

/RenameFTRoot

/RootScalability

/SiteName

/InSite

/SiteCosting

/Export

/Import

/UpdateWin2kStaticSiteTable

/ShowWin2kStaticSiteTable

/PurgeWin2kStaticSiteTable

/View

/PktInfo

/SpcInfo

/ExportBlob:filepathname

/CheckBlob

/DisplayBlob:filepathname

/ViewDfsDirs:volume

/ImportRoot:SrcDfsName /{Mirror|Compare}

/UnMapFtRoot

/Clean

/PktFlush

/SpcFlush

/PurgeMupCache

/RemoveReparse:absolute_directory_path

**/api**
Instructs DFSUtil to go through regular server APIs instead of directly connecting to the Active Directory.

DFSUtil normally operates on domain-based roots directly, without communicating with the DFS Service on the server hosting a particular root. This default mode is known as Direct Mode. In Direct Mode, commands execute faster, especially when importing large namespaces. However, for debugging purposes, you may decide to communicate with the DFS server by using its APIs instead.

**Syntax**
- **dfsutil** *option* **/api** [**/verbose**]

**Parameters**
***option***
The option that is to be modified by the **/api** parameter.

**[ /verbose]**
Provides verbose output to the screen. This mode also notifies the user what mode a command executes in.

**Remarks**
Direct mode is essential when using parameters such as **/RenameFtRoot**, **/UnmapFtRoot**, and **/Insite**. These parameters will fail if Direct Mode does not initialize properly. Parameters that do not require a well-initialized Direct Mode (such as **/View**, **/Import**, and **/export**) will attempt to execute in the **/api** mode if the default Direct Mode fails. Specifying **/Verbose** displays the following information:

**Sample Usage**
```
dfsutil /view /root:\\dfsname\dfsshare /api
```

**/verbose**
Displays more detailed information when used with an Dfsutil Remarks.

**Syntax**
- **dfsutil** *option* **/verbose**

**Parameters**
**[ *option*]**
The option to be modified by **/verbose**.

**Sample Usage**
```
dfsutil /view /root:\\dfsname\dfsshare /verbose
```

**/? | /command /?**
Displays command-line help. The ***/command /?*** parameter displays more specific help about that command.

**Syntax**
- **dfsutil** [*option*] {**/?** | **/help**}

**Parameters**
**[ *option*]**
The option that is to be modified by the **/?** or **/help** parameter.

**Remarks**
Some parameters provide additional help text when used with **/?** or **/help**.

**Sample Usage**

```
dfsutil /?

dfsutil /help

dfsutil /export /?
```

**/root:**

Specifies the DFS root context where the specified command is to be performed.

**Syntax**

- **dfsutil** *option* **/root:** *dfs_root_name*

**Parameters**

***option***

The option that is to be modified by the **/root:** parameter.

**/root:** ***dfs_root_name***

The UNC path to the DFS root context.

**Sample Usage**

To display the DFS configuration data of a specific DFS root:

```
dfsutil /view /root:\\dfsname\dfsshare
```

**/server:**

Specifies the DFS namespace server on which the designated command is performed.

**Syntax**

- **dfsutil** *option* **/server:** *dfs_server_name*

**Parameters**

**[** ***option*** **]**

The option that is to be modified by the **/server:** parameter.

**/server:** ***dfs_server_name***

Specifies a DFS name space server.

**Sample Usage**

```
dfsutil /view /server:mycomputer.mydomain.com
```

**/share:**

Specifies the DFS or SMB share on which the designated command is performed.

**Syntax**

- **dfsutil** [*option*] **/share:***dfs_share_name*

**Parameters**

**[** ***option*** **]**

The option that is to be modified by the **/share:** parameter.

**/share:** ***dfs_share_name***

The name of a DFS or SMB share.

**Sample Usage**

```
dfsutil /clean /server:aseanl-test1 /share:dfsroot
```

**/domain:dfsDomainName**

Specifies the DFS domain on which the designated command is performed.

**Syntax**

- **dfsutil** *option* **/domain:** *dfs_domain_name*

**Parameters**

**[** ***option*** **]**

The option that is to be modified by the **/domain:** parameter.

***dfsDomainName***

The netbios name or the fully qualified name of a domain.

**/newDomain:dfsDomainName**

Specifies the DFS domain on which the designated command is performed.

**Syntax**

▸ **dfsutil** *option* **/newdomain:** *dfs_domain_name*

**Parameters**

**[ *option*]**

The option that is to be modified by the **/newdomain:** parameter.

***dfsDomainName***

The netbios name or the full qualified name of a domain.

**/oldDomain:dfsDomainName**

Specifies the DFS domain on which the designated command is performed.

**Syntax**

▸ **dfsutil** *option* **/olddomain:** *dfs_domain_name*

**Parameters**

**[ *option*]**

The option that is to be modified by the **/olddomain:** parameter.

***dfsDomainName***

The netbios name or the fully qualified name of a domain.

**/comment:**

Adds a custom text comment to the blob when creating or deleting a Root, Link, or Target.

**Syntax**

▸ **dfsutil** *option* **/comment:** *text*

**Parameters**

**[ *option*]**

The option that is to be modified by the **/comment** parameter.

***text***

The text to be added.

**Remarks**

Because they consume blob space, keep comments short.

### Important

- Avoid or reduce the usage of comments if you have a fairly high number of links and targets for a given root. Doing so will save blob space and improve overall DFS performance and scalability.

**Sample Usage**

```
dfsutil /AddFtRoot:MyDfs /server:MyServer /share:MyShare /comment:"A root of my dfs."
```

**/AddFTRoot:**

Creates a new DFS root in an existing domain-based DFS namespace. This command can also be used to add a computer as a root target to an existing namespace.

**Syntax**

▸ **dfsutil /addftroot:** *new_dfs_root_name* **/server:** *dfs_server_name* **/share:** *dfs_share_name* [**/comment:** *text*]

**Parameters**

***newDfsRootName***

Name of the DFS to which the root will belong.

**/server: *dfsServerName***

Name of the server on which to create the root.

**/share: *dfsShareName***

Name of the share that will back the root.

**/comment: *text***

Optional custom text added to the blob related to this action.

**Remarks**
Because they consume blob space, keep comments short.

**Sample Usage**
In a domain-based DFS named \\DomainName\MyDfs, add \\MyServer\MyShare as a root of that DFS.

```
dfsutil /AddFtRoot:MyDfs /server:MyServer /share:MyShare /comment:"A root of my dfs."
```

**/AddStdRoot:**
Creates a new DFS root in an existing standalone DFS namespace. This command can also be used to add a computer as a root target to an existing namespace.

**Syntax**
▸ **dfsutil /addstdroot:** *dfsRootName* **/server:** *dfsServerName* **/share:** *dfsShareName* [**/comment:** *text*] [**/api**]

**Parameters**
***dfsRootName***
Name of the root to remove.

**/server: *dfsServerName***
Name of the server on which to delete the root.

**/share: *dfsShareName***
Name of the root backup share.

**/comment: '' *text* ''**
Optional custom text added to the blob related to this action.

**Remarks**
Because they consume blob space, keep comments short.

**Sample Usage**
In a domain-based DFS named \\DomainName\MyDfs, add \\MyServer\MyShare as a root of that DFS.

```
dfsutil /AddStdRoot:MyDfs /server:MyServer /share:MyShare /comment:"A root of my dfs."
```

**/RemFTRoot**
Removes a root target from an existing domain-based DFS root.

**Syntax**
▸ **dfsutil /RemFTRoot:** *dfsRootName* **/server:** *dfsServerName* **/share:** *dfsShareName* [**/api**]

**Parameters**
***dfsRootName***
Name of root to remove.

**/server: *dfsServerName***
Name of the server on which to delete the root.

**/share: *dfsShareName***
Name of the root backup share.

**Remarks**
Prior to using this option, use **/View** to see how many replicas exist for this namespace.

**Caution**

- Removing the last replica of a DFS namespace causes the namespace to be wiped out.

To back up a namespace, use **/export**. To reload a namespace which has been backed up with **/export**, use **/import**.

**Sample Usage**
To remove the root target of MyDfs which currently is backed up by \\MyServer\MyShare:

```
dfsutil /RemFtRoot:MyDfs /server:MyServer /share:MyShare
```

**/RemStdRoot**
Removes a root target from an existing stand-alone DFS root.

**Syntax**

**dfsutil /RemStdRoot:** *dfsRootName* **/server:** dfsServerName **/share:** *dfsShareName* [**/api**]

**Parameters**

***dfsRootName***
Name of the root to remove.

**/server:** ***dfs_server_name***
Name of the server on which to delete the root.

**/share:** ***dfsShareName***
Name of the root backup share.

**Remarks**

Prior to using this option, use **/View** to see how many replicas exist for this namespace.

**Caution**

- Removing the last replica of a DFS namespace causes the namespace to be wiped out.

To back up a namespace, use **/export**. To reload a namespace which has been backed up with **/export**, use **/import**.

**Sample Usage**

To remove a root target of MyDfs that is backed up by \\MyServer\MyShare:

```
dfsutil /RemFtRoot:MyDfs /server:MyServer /share:MyShare
```

**/RenameFTRoot**

Given a DFS namespace, **/RenameFtRoot** renames references to an obsolete domain within that namespace.

This command can also be used to rename any link or target in a namespace.

**Syntax**

**dfsutil /renameftroot root:** *dfs_root_name* **/olddomain:** *domain_name* **/newdomain:** *domain_name* [**/verbose**]

**Parameters**

**/root:** ***dfs_root_name***
The UNC path to the DFS root context.

**/olddomain:** ***domain_name***
The name that you want to change your DFS namespace from. The same name that the **/root:** parameter references. This is a required parameter.

**/newdomain:** ***domain_name***
The name that you want to change your DFS namespace to. This is a required parameter.

**[ /verbose]**
Displays verbose information in the output.

**Remarks**

Prior to changing the root, it is a good practice to discover the existing state. The **/verbose** parameter is useful with this command as it lists all changes which have been made to the namespace.

**Important**

- After running this parameter, you must restart the DFS service on the servers where this command was run. Restarting the service causes the servers to recognize the new domain name in the namespace.

This command is typically run as a part of a domain rename. It is expected that the domain has already been renamed and the affected systems have been rebooted by the time this command is run. DFS servers that are running Windows 2000 Server or earlier are supported as root target servers; **/RenameFtRoot** will in fact change the registry entries of those systems appropriately.

**/RenameFtRoot** does not run in the **/Api** Mode. If Direct Mode initialization fails on the root specified (typically because the root is not Domain-based), then this command will fail.

**Sample Usage**

```
dfsutil /renameftroot /root: machname.OLDDOMAIN.company.com /olddomain: machname.OLDDOMAIN.company.com
```

```
/newdomain: machname.NEWDOMAIN.company.com
```

**/RootScalability**
An advanced command which increases the performance of a large scale deployment of DFS namespaces.

**Syntax**
▸ **dfsutil /rootscalability /root:** *dfs_root_name* {**/Enable**|**/Disable**|**/Display**} [**/verbose**]

**Parameters**
***/root: dfs_root_name***
The UNC path to the DFS root context.

**/rootscalability { /enable| disable| display}**
Select the appropriate root scaleability action.

| Value | Description |
|---|---|
| **enable** | Enables root scalability among DFS root servers. When enabled, network traffic among DFS root servers is kept to a minimum. In addition, there will be less traffic between the primary domain controller and DFS servers. This option is a tradeoff between users needing up-to-date information from DFS servers, and administrators needing minimal bandwidth consumptions by the DFS operations, (inter-DFS-server and DFS-to-domain-controller communication). |
| **disable** | Default behavior. This option results in tighter synchronization between DFS servers and Domain Controllers. |
| **display** | Display the current root scalability mode (enabled or disabled). |

**[ /verbose]**
Displays whether the root-scalability feature is currently enabled.

**Remarks**
Use this option in scenarios where the DFS metadata is comparably less dynamic and the network bandwidth has priority over up-to-date information or accessibility to up-to-date targets. Enabling this option may cause users to see outdated information from DFS servers at times, in terms of the target links and related information.

**Notes**

- This command does not work in **/Api** mode.
- The path specified must be a root, not a link.
- This has no effect on Standalone roots.
- It is a good practice to find out the current state before changing it.
- When **/rootscalability** is enabled, it is not uncommon to see an event log message such as, "DFS could not access its private data from the DS...". While this error may still indicate a problem in DS connectivity, typically this occurs because the nearest domain controller has outdated DFS information (expected behavior when **/rootscalability** is enabled).

**/SiteName**
Displays the site name that a certain computer belongs to.

**Syntax**
▸ **dfsutil /sitename:** *dfsSiteName*

**Parameters**
***dfsSiteName***
The name of the computer on which site(s) reside. Valid formats in which to specify the computer name include:

- computer name (local only)
- UNC name
- IP address

**Sample Usage**

```
dfsutil /sitename:machine4

dfsutil /sitename:mailserver.dom.com

dfsutil /sitename:georgeLaptop.domain.com

dfsutil /sitename:12.230.25.24
```

**/InSite**

Set or reset the insite setting on a root or a link, and view the current insite setting of a link or a root.

The insite setting on a link or root in a DFS defines how the DFS redirects the client access to that link or root, depending on the sites in which the client and the targets reside. The target could be a link target or a root target. This becomes important when the administrator wants to control the client accessing the root targets or link targets outside its own site. By using **/InSite**, the administrator can force the server to provide, in the referral list it returns to the client, only those targets which are in the same site as that of the client, when the client requests access to a certain root or link.

The **/InSite** setting has two modes: enabled and disabled.

**Syntax**

- **dfsutil /root:** *dfs_root_name* **/insite** /{**Enable** | **Disable** | **Display**} [**/verbose**]

**/root:** ***dfs_root_name***

The root name or the entire link name in UNC.

**/insite { /enable| disable| display}**

Select the appropriate insite action.

| Value | Description |
|---|---|
| **enable** | Clients will only access insite targets. This mode saves bandwidth by prohibiting client access to targets outside its site. However, this also reduces the availability of the DFS root on the link. |
| **disable** | Client will also access targets outside its site if necessary. |
| **display** | Displays the current insite mode (enabled or disabled) for the link or root. |

**[ /verbose]**

Displays verbose output to the screen.

**Remarks**

If you do not want to enable **/InSite**, see the **/SiteCosting** parameter. **/SiteCosting** provides more control over the DFS target referrals sent by the DFS server.

**Sample Usage**

```
dfsutil /root:\\dom.com\dfsroot /insite /display
```

Displays the current insite setting of the root dfsroot on domain dom.com.

```
Dfsutil /root:\\dom.com\DfsName\LinkA\B\C /insite /enable /verbose
```

This forces the DFS server to provide only insite target lists in the referrals for the link LinkA\B\C under the root \\dom.com\DfsName. As a result, the client accessing the link \\dom.com\DfsName\LinkA\B\C will not get a referral outside its site and will fail if there are NO targets responding in its site.

```
dfsutil /root:\\dom.com\DfsName\LinkA\B\C /insite /disable /verbose
```

This instructs the DFS server to provide all the link targets under link LinkA\B\C under the root \\dom.com\DfsName. As a result, the client accessing the link \\dom.com\DfsName\LinkA\B\C will get a list of referrals with targets in its site and outside its site as well. In fact, the targets of its site will be listed first (randomly), followed by the targets outside its site (also randomly).

**/SiteCosting**

Enables the DFS to issue referrals to the clients. The order in which the target referrals are listed is based on the following rules:

1. The targets of the same site as that of the client are listed first. The order is randomized to provide load balancing.

2. The targets from the next closest site are listed next (randomized again). This repeats until #3 occurs.

3. The targets without any site information are taken as the farthest site, listed at the end of the referrals and randomized.

**Syntax**

- **dfsutil /root:** *dfs_root_name* /sitecosting {/Enable | /Disable | /Display} [**/verbose**]

**Parameters**

***/root: dfs_root_name***

The UNC path to the DFS root context. This is a required parameter.

**/sitecosting { /enable| disable| display}**

Select the appropriate sitecosting action.

| Value | Description |
|---|---|
| **enable** | Enable sitecosting. |
| **disable** | Disable sitecosting. (Default.) |
| **display** | Display the current sitecosting mode (enabled or disabled). |

**Remarks**

To illustrate how SiteCosting works, the Root \\dom5.com\fsRoot has a link named linkA\B\C, which, in turn, has 10 targets in different sites as follows:

| | Entity | Site | Cost from client site |
|---|---|---|---|
| **Client** | \\myLaptop.dom1.com | NY-MANHATTAN-01 | 0 |
| Targets for link LinkA\B\C | \\nyServ1\Folder | NY-MANHATTAN-01 | 0 |
| | \\Asterix\universal | -- | Infinite |
| | \\laOrange\ShareData | LA-ORANGE-21 | 5 |
| | \\Bromley02\Kohinoor | LONDON-UK-10 | 22 |
| | \\Champs1\Elysee | PARIS-Douxieme-01 | 19 |
| | \\Eiffel4\Distributed | PARIS-Douxieme-01 | 19 |
| | \\LondonBridge01\Data | LONDON-UK-06 | 21 |
| | \\ny43rdFloor\Data | NY-MANHATTAN-01 | 0 |
| | \\Submarine\Air | -- | Infinite |
| | \\laOrange5\Replicated | LA-ORANGE-21 | 5 |

When SiteCosting is disabled, the DFS server issues the referral as follows:

1. Targets \\nyServ1\Folder, \\ny43rdFloor\Data randomly sequenced.

2. Remaining targets randomly sequenced irrespective of their site.

When SiteCosting is enabled, the Referral list looks like this:

1. Targets from the site same as the client randomly sequenced: \\ny43rdFloor\Data, \\nyServ1\Folder

2. Targets from the next least expensive Site randomly sequenced: here site LA-ORANGE-21. Targets:

\\laOrange\ShareData, \\laOrange5\Replicated.

3. Targets from 3rd least expensive site randomly sequenced: here the site PARIS-Douxieme-01. Targets: \\Eiffel4\Distributed, \\Champs1\Elysee.

4. Targets from 4th least expensive Site randomly sequenced: here the site LONDON-UK-06. Targets: \\LondonBridge01\Data.

5. Targets from 5th least expensive Site randomly sequenced: here the site LONDON-UK-10. Targets: \\Bromley02\Kohinoor.

6. Finally, targets without any site information known, believed to be the most expensive, are randomly sequenced. Targets: \\Asterix\universal, \\Submarine\Air.

**Sample Usage**

```
dfsutil /root:\\dom5.com\fsroot /sitecosting /enable /verbose
```

**/Export**

The **/export** parameter converts the configuration data into a reusable script file that describes an entire namespace. You specify the DFS namespace to export, and a file in which to save the exported configuration data. This file typically is the input for the **/import** parameter described later.

**Syntax**

▸ **dfsutil /root:** dfs_root **/export:** *file_name*

**Parameters**

***/root: dfs_root_name***

The UNC path to the DFS root context. This is a required parameter.

***/export: file_name***

The name of the file to export DFS configuration data to.

**Remarks**

The output file of the export is an XML file in DFS on a Windows Server 2003 system. Hence, it is important to note that if your current DFS is of a prior version, or the output files are coming from a prior version, it should be converted to the XML format to be reloaded by using the **/import** command option.

Warning

- If the file you specify to export already exists it will be overwritten without any warning. You will lose any data in the original file.

**Sample Usage**

```
dfsutil /root:\\dfsname\dfsshare /export:dfsoutput.txt
```

**/Import**

The **/import** parameter creates or updates a DFS namespace from a DFS metadata stored in a specified file.

**Syntax**

▸ **dfsutil /Root:** *dfsRootName* **/import:** *filename* **/{ Verify** | **Set** | **Merge**} [**/Verbose**]

**Parameters**

***dfsRootName***

An existing DFS root to which the *filename* file is imported. Only the namespace specified by the root argument is affected.

***filename***

An XML-formatted text file containing the DFS namespace. This file has usually been created by running **/export**.

**/{ Verify| Set| Merge}**

Select the appropriate namespace action. This is a required parameter.

| Value | Description |
|---|---|

| | |
|---|---|
| **Verify** | Displays changes that may happen during the import. |
| **Set** | Replaces the old namespace with the new namespace. All links and targets specified in the import file but not present in the new namespace are deleted. |
| **Merge** | Merge options decide the fate of the namespace that existed before the import command. The Merge option updates not only the links and targets that are common to both namespaces but also the new links and targets scripted in the input file. This is, in effect, a union of the existing namespace with the importing file-based namespace. If there is a conflict between the two, the file-based configuration is considered final. |

**Remarks**

Use **/Verify** prior to the actual import (that is, using **/Set** or **/Merge**), to assess changes that may happen during the import.

The output file of the export is in XML format for Windows Server 2003 DFS. If your current DFS is of a prior version, or the output files are coming from a prior version, it should be converted to the XML format to be reloaded in the import command option with this version 4.0, that is, this version of DFSUtil.exe.

**/UpdateWin2kStaticSiteTable**

This is a mixed mode-only (Windows 2000 *and* Windows Server 2003) option.

In case of a DFS environment with servers running both Windows 2000 and Windows Server 2003, this option runs on the server running Windows Server 2003 to update the site information useful to the server running Windows 2000. When a link target is being added to a link on a domain-based DFS from a server running Windows Server 2003, the site information for that link target is not added to the DFS blob on the Active Directory. (This is different from a Windows 2000 domain-based DFS, where it would be added to the blob. The reason for this difference in behavior is that the server running Windows Server 2003 determines the site information dynamically (on an as-needed basis), so there is no need to add it to the blob in order to save the blob space. In the mixed mode scenario, where there is a domain-based DFS with at least one member server running Windows Server 2003 and at least one member server running Windows 2000, site information of any link target added by the server running Windows Server 2003 will not be added to the blob. So the member server running Windows 2000 will never know its site information.

**Syntax**

- **dfsutil /root:** *dfsRootName* **/UpdateWin2kStaticSiteTable**

**Parameters**

***/root:h dfsRootName***

The DFS root on which this command is performed. Valid names are constructed as: \\dfsname\dfsshare.

**Remarks**

This command updates the Windows 2000 Site table.

The problem of mixed-mode interoperability is solved until more link targets are added from any server running Windows Server 2003. It is a good practice to update this site information table, as it is static, and an up-to-date table keeps the blob size optimal.

Use **/ShowWin2kStaticSiteTable** to display the state of the existing namespace. This can be especially useful in determining if the actual DFS namespace matches your namespace design.

**/ShowWin2kStaticSiteTable**

This is a mixed mode-only (Windows 2000 *and* Windows Server 2003) option.

Displays the current Target-to-Site mapping information stored in the blob of a DFS server running Windows 2000.

In a domain-based DFS with at least one member server running Windows Server 2003, and at least one member server is running Windows 2000, site information of any link target added by the server running Windows Server 2003 will not be added to the blob. So the member server running Windows 2000 will never know its site information.

To verify what the server running Windows 2000 knows about the sites of various targets, this option displays the Target-to-Site mapping.

**Syntax**

- **dfsutil /Root:** *dfsRootName* **/ShowWin2kStaticSiteTable**

**Parameters**

***dfsRootName***

The DFS root on which this command is performed. Valid names are constructed as: \\dfsname\dfsshare.

**Remarks**
This helps the Administrator verify if member server running Windows 2000 has all the necessary site information.

**/PurgeWin2kStaticSiteTable**
This is a mixed mode-only (Windows 2000 *and* Windows Server 2003) option.

This parameter flushes site information from the blob. In a domain-based DFS with at least one member server is running Windows Server 2003, and at least one member server is running Windows 2000, site information of any link target added by the server running Windows Server 2003 will not be added to the blob. So the member server running Windows 2000 will never know its site information.

**Syntax**
**dfsutil /Root:** *dfsRootName* **/PurgeWin2kStaticSiteTable**

**Parameters**
***dfsRootName***
The DFS root on which this command is performed. Valid names are constructed as: \\dfsname\dfsshare.

**/View**
Displays the DFS configuration data of a DFS root. Included in the configuration data is a list of the roots of the DFS, a list of all of the links of the DFS, and site information for the DFS.

**Syntax**
**dfsutil /view** {**/Root:** *dfsRootName*|**/server:** *dfs_root*|**/domain:** *dfs_root*}

**Parameters**
***dfsRootName***
The DFS root on which this command is performed. Valid names are constructed as: \\dfsname\dfsshare.

**/server:** ***dfs_root***
Specifies the DFS namespace server on which this command is performed.

**/domain:** ***dfs_root***
Specifies the domain on which this command is performed.

**Sample Usage**
```
dfsutil /view
```

**/PktInfo**
Displays the following client information currently in the Partition Knowledge Table (PKT):

- Parts of the DFS namespace cached by the client
- Names of the servers participating in the DFS share
- Clients' randomization order of the participating servers
- Current "go to" server

See Dfsutil Remarks for a description of the output produced by this parameter.

**Syntax**
**dfsutil /pktinfo** [**level:**{**0**|**1**}]

**Parameters**
**/level:{ 0|1}**
Optional report detail argument which takes a value of either 0 (low detail) or 1 (high detail).

**Remarks**
This is a DFS client-side operation.

Each entry will report a name, both full and short, the remaining time that the referral and information contained in it will remain in the referral list, followed by the referral list itself. Active alternate DFS client targets are indicated by the number "0X31" after the alternate name entry.

**Note**

- Flushing the Partition Knowledge Table (PKT) clears existing information and force the DFS client to fetch new information for the next series of access attempts. See \PKTFlush.

DFS servers store the DFS namespace and the servers backing up the namespace in a BLOB format in the registry (stand-alone DFS) or Active Directory (domain DFS). The DFS BLOB is formally known as the Partition Knowledge Table (PKT). DFS clients retrieve and cache each piece of the DFS namespace for a period of time determined by the server.

For DFS servers running Windows NT 4.0, the servers send the list of servers backing a piece of the DFS namespace in the same order. Clients randomize the list of servers to provide a degree of load balancing.

For DFS servers running Windows 2000, Windows XP, or Windows Server 2003 that are accessed by Windows 2000 DFS clients, the scrambling and site preference takes place on the server.

**Sample Usage**
```
dfsutil /pktinfo
```

**/SpcInfo**
Displays the content of SPC. The SPC contains the list of trusted domains; under each trusted domain there may be one or more domain controllers.

See Dfsutil Remarks for a description of the output produced by this parameter.

**Syntax**
▸ **dfsutil /spcinfo**

**Remarks**
This is a DFS client-side operation.

This cache can be flushed if you feel the client has inaccurate information or is not aligned with the design of a DFS environment. See the SPCflush option.

**Sample Usage**
```
dfsutil /spcinfo
```

**/ExportBlob:filepathname**
Exports the blob of a DFS namespace which already exists (that is, specific to a root). This option saves the blob of the namespace specified, without any format changes, directly to a file.

**Syntax**
▸ **dfsutil** *dfsRootName* **/exportblob:** *filename*

**Parameters**
***dfsRootName***
The DFS root on which this command is performed. Valid names are constructed as: \\dfsname\dfsshare.

***filename***
An absolute or relative pathname or working directory file.

**Sample Usage**
```
dfsutil /Root:\\dom.com\homedirroot /ExportBlob:e:\Backups\DFSbackups\Homedirroot\today.dfs
```

**/CheckBlob**
Checks the blob of a DFS namespace which already exists (that is, specific to a root). This is a basic consistency check.

**Syntax**
▸ **dfsutil /Root:** *dfsRootName* **/checkblob**

**Parameters**
***dfsRootName***
The DFS root on which this command is performed. Valid names are constructed as: \\dfsname\dfsshare.

**Sample Usage**
```
dfsutil /Root:\\dom.com\homedirroot /checkBlob
```

**/DisplayBlob:filepathname**
Displays the blob stored in a blob file. The blob file is generally exported with /ExportBlob. This option is used to view the blob of a specific root.

**Syntax**

**dfsutil /displayblob:** *filepathname*

**Parameters**

***filepathname***

Displays an existing file containing blob information.

**/ViewDfsDirs:volume**

Scans the specified volume and lists all DFS reparse directories. Reparse directories can also be removed when this option is used with the RemoveReparse option.

**Syntax**

**dfsutil /viewdfsdirs:** *volume* [**/removereparse**: *path*] [**/verbose**]

**Parameters**

***volume***

Name of the volume to scan.

**/removereparse:** ***path***

Removes all DFS reparse directories under a given absolute path directory/ hierarchy.

**/verbose**

Displays more detailed information.

**Remarks**

The **/RemoveReparse** parameter can be used separately over a directory tree to remove all the DFS reparse directories under it.

**Sample Usage**

The following syntax lists all DFS reparse directories existing on volume (drive) e:`dfsutil /viewdfsdirs:e:`

The following syntax views and deletes (removes) all DFS reparse directories on volume (drive) c:. The **/verbose** parameter includes more descriptive details.`dfsutil /ViewDfsDirs:c: /removereparse /verbose`

**/ImportRoot:SrcDfsName /{Mirror|Compare}**

Copies an entire namespace under a DFS root to another DFS root, which, in turn, enables you to perform other operations such as Namespace Backup, DFS Server Maintenance, Create DFS mirror etc.

**Syntax**

dfsutil **/Root:***dfsRootName***/ImportRoot**:*SrcDfsName* {**/Mirror** | **/Compare**} [**/Api**]

**Parameters**

***SrcDfsName***

An existing DFS namespace (root).

**/{Mirror| Compare}**

Select the appropriate namespace action.

| Value | Description |
|---|---|
| **Mirror** | Back up the *SrcDfsName* namespace. |
| **Compare** | Compare the *SrcDfsName* namespace with the new namespace. |

**/ API**

Execute in API mode.

**Remarks**

Back up the existing namespace using the **/export** parameter before performing this operation.

Warning

- When mirroring the DFS namespace, be careful to provide the **/Root** and **/ImportRoot** parameters in the correct order. Specifying these parameters in reverse order could be disastrous and irrecoverable.

**Sample Usage**

The namespace \\dom1\existing is compared to \\dom2\bin, where \\dom\bin can also exist.`dfsutil /Root:\\dom2\bin /ImportRoot:\\dom1\existing /Compare` The namespace \\dom1\existing is to be mirrored to \\dom2\bin, where \\dom\bin

can also exist but this will overwrite the existing namespace.

```
dfsutil /Root:\\dom2\bin /ImportRoot:\\dom1\existing /Mirror
```

**/UnMapFtRoot**
Manually removes references to an obsolete domain-based root target. At times it may not be feasible for the administrator to remove a root target server from the namespace gracefully. This command eliminates the resultant obsolete root target in the DFS namespace. Since the target host server is unlikely to exist at this point, DFSUtil does not attempt to contact it.

**Syntax**
**/unmapftroot /server:** *root_target_server* **/share:** *root_target_share*

**Parameters**
**/server: *root_target_server***
This option specifies the root target server. This is a required parameter.

**/share: *root_target_share***
This option specifies the target share server.

**Remarks**
This advanced command should be used only to delete references to a fault-tolerant root target that does not exist. It works only on domain-based roots.

### Caution

- This command is essentially a problem/corruption repair tool. Do not confuse it with commands such as **/RemFtRoot** or **/Clean**.

The command will fail if Direct Mode was unable to initialize.

**Sample Usage**
```
dfsutil /UnmapFtRoot /root:\\mydomain.com\namespaceroot /server:\\otherservername /share:\\myserver\share
```

**/Clean**
A troubleshooting option which removes a reference to an obsolete root.

**Syntax**
**dfsutil /clean /server:** *hostServer* **/share:** *rootname* [**/verbose**]

**Parameters**
**/server: *hostServer***
The server where the obsolete root existed.

**/share: *rootname***
The root reference to remove.

**Remarks**
If a root was not removed properly (by using the **/RemFtRoot** command), other root target servers may still try to reference the obsolete root. This option fixes the incorrect references.

This command contacts the registry of the specified system to delete obsolete references. It does not affect the Domain Controller or the Active Directory. This option should not be confused with commands such as **/UnmapFtRoot** and **/RemFtRoot**.

**/PktFlush**
Flushes the client's locally cached target referral list provided by the DFS server. When the client tries to access the corresponding link, this forces the client to refresh the referral list of targets from the DFS server. Some of the entries in the PKT may not get flushed, especially if DFS is in the process of using the referrals.

**Syntax**
**dfsutil /pktflush**

**Remarks**
This is a DFS client-side operation.

### Important

- This option should be avoided unless absolutely necessary. Under normal DFS operating conditions, there is no reason to

flush the existing referral list.

DFS clients cache the portions of the DFS namespace (PKT) for the duration of time specified in the Distributed File System Manager and referred to as the Time to Live (TTL). For DFS version 4.1, the PKT was cached for a hard-coded seven (7) days. Servers running Windows 2000 use 1800 seconds (30 minutes) as the default TTL. All Microsoft DFS clients delete the PKT for a given folder in the DFS tree if not accessed by the TTL expiration or when the client is rebooted, whichever occurs first.

Computers running Windows 2000 and Windows XP operating systems support the DFS tab in the Windows Explorer property sheet, which uses the **NetDFSGetClientInfo()** and **NetDFSSetClientInfo()** APIs to show the list of alternative paths residing in the volume, and which allows choosing another alternative (refresh PKT) without having to reboot or wait for TTL to expire.

DFS clients that run an operating system other than Windows 2000 or Windows XP must reboot or wait for the TTL for a cached portion of the DFS tree to expire.

**/SpcFlush**
Forces the client to fetch the latest information on the trusted domains and their domain controllers, when the client programs try to access a corresponding link or root. By using this option, the client flushes out its cached knowledge about the trusted domains and the domain controllers of these domains.

**Syntax**
**dfsutil /spcflush**

**Remarks**
This is a DFS client-side operation.

**/PurgeMupCache**
Resets the client's knowledge about the various sites' information. This is a troubleshooting option which should only be run on the client.

Clears the client MUP cache, preventing confusion about the current provider when such names conflict. Except for a temporary performance hit, this command has no other adverse effects. This command does not affect any DFS metadata. If this command is not run, and the namespace is not accessed, the obsolete cache entry eventually expires.

**Syntax**
**dfsutil /purgemupcache**

**Remarks**
This is a DFS client-side operation.

The MUP cache keeps information about both DFS and other shares on the client system. You create a non-DFS (for example, SMB) share and access it before mapping a DFS namespace by the same name on it. The MUP cache might still think that the share is served by the previous provider (SMB) although DFS is the correct provider now.

**/RemoveReparse:absolute_directory_path**
Removes all DFS reparse directories under a given absolute path directory/hierarchy.

**Syntax**
**dfsutil /removereparse:** *absolute_directory_path* [**/viewdfsdirs**]

**Parameters**
***absolute_directory_path***
The AbsoluteDirectoryPath is the path of the directory tree under which your reparse directories exist, all of which you intend to remove.

**[ /viewdfsdirs]**
See **/viewdfsdirs** for more information

**Remarks**
This option can be used as an additional parameter to **/viewdfsdirs**. In this context, it removes all the DFS reparse directories under a specified volume.

**Sample Usage**
The following syntax removes all DFS reparse directories under the specified directory. `dfsutil /RemoveReparse: c:\dfspath\dfsshares`

**See Also**
**Concepts**

Dfsutil Overview
Dfsutil Remarks
Dfsutil Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Dmdiag Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dmdiag.exe: Disk Manager Diagnostics**
This command-line tool displays system state and configuration information describing disk storage.

DmDiag displays the following information for the computer on which it is run:

- Computer name and operating system version
- Physical disk to disk type
- Mount points
- LDM file versions
- Drive letter usage, GetLogicalDrives(), GetDriveType()
- \Device
- Symbolic links
- ldmsize
- Kernel list
- Disk partition information

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
For more information, see Disk Management overview.

**System Requirements**
The following are the system requirements for DmDiag:

- Windows Server 2003 or Windows XP Professional

**File Required**

- Dmdiag.exe

**See Also**
**Concepts**
Dmdiag Syntax
Dmdiag Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview

Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Dmdiag Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DmDiag Examples**

**Example 1: Write disk configuration to file**

If you want to write disk configuration information to a file named Diskconf.txt, type the following at the command line:

```
dmdiag -f c:\data\diskconf.txt
```

**Example 2: Display disk configuration**

If you want to display disk configuration information in the command window, type the following at the command line:

```
dmdiag
```

**Sample DmDiag output**

```
---------- Computer Name and OS Version ----------

       Computer name: server2
            NT build: 3675
            CPU Type: x86
      DMDIAG Version: 5.2.3678.0 shp

---------- LDM File Versions ----------

                5.2.3675.0 shp - C:\WINNT\system32\dmadmin.exe
                5.2.3675.0 shp - C:\WINNT\system32\dmconfig.dll
                5.2.3675.0 shp - C:\WINNT\system32\dmdlgs.dll
                5.2.3675.0 shp - C:\WINNT\system32\dmdskmgr.dll
                5.2.3675.0 shp - C:\WINNT\system32\dmdskres.dll
                5.2.3675.0 shp - C:\WINNT\system32\dmintf.dll
                5.2.3675.0 shp - C:\WINNT\system32\dmremote.exe
                5.2.3675.0 shp - C:\WINNT\system32\dmserver.dll
                5.2.3675.0 shp - C:\WINNT\system32\dmutil.dll
                5.2.3675.0 shp - C:\WINNT\system32\dmview.ocx
                5.2.3675.0 shp - C:\WINNT\system32\drivers\dmboot.sys
                5.2.3675.0 shp - C:\WINNT\system32\drivers\dmio.sys
                5.2.3675.0 shp - C:\WINNT\system32\drivers\dmload.sys

---------- Mount Points ----------


---------- Drive Letter Usage, Drive Type ----------

A: = \Device\Floppy0 [Removable]
C: = \Device\HarddiskVolume1 [Fixed]
D: = \Device\HarddiskVolume2 [Fixed]
E: = \Device\CdRom0 [CDRom]

---------- Consolidated LDM Configuration Data ----------

ERROR: scan operation failed:
        A format error was found in the private region of the disk
ERROR: scan operation failed:
        A format error was found in the private region of the disk

---------- \Device\Harddisk0 ----------

\Device\Harddisk0\DP(1)0x7e00-0x7d040000+1 (Device)
\Device\Harddisk0\DP(2)0x7d047e00-0x845f1c00+2 (Device)
\Device\Harddisk0\DR0            (Device)
\Device\Harddisk0\Partition0     (SymbolicLink) -> \Device\Harddisk0\DR0
\Device\Harddisk0\Partition1     (SymbolicLink) -> \Device\HarddiskVolume1
\Device\Harddisk0\Partition2     (SymbolicLink) -> \Device\HarddiskVolume2

---------- Partition Table Info Disk 0 ----------
```

```
             525 Cylinders
             255 Tracks/Cylinder
              63 Sectors/Track
             512 Bytes/Sector
              12 MediaType
       8,434,125 Sectors (total)
   4,318,272,000 Bytes (total)
       4,217,063 KB
           4,118 MB
             4.0 GB

               0 StartingOffset
   4,321,787,904 PartitionLength
               0 HiddenSectors
               0 PartitionNumber
               0 PartitionType
               0 BootIndicator
               0 RecognizedPartition
               0 RewritePartition

             MBR PartitionStyle
               4 PartitionCount
        da0dda0d Signature

        Starting        Partition     Hidden       Total  Partition    Partition
       Boot Recognized    Rewrite
  Offset (bytes)   Length (bytes)    Sectors     Sectors     Number   Type (HEX)
  Indicator  Partition  Partition

         32,256    2,097,414,144         63   4,096,512          0         0x07
         1          1          0
   2,097,446,400    2,220,825,600  4,096,575   4,337,550          1         0x07
         0          1          0
              0                0          0           0          2         0x00
         0          0          0
              0                0          0           0          3         0x00
         0          0          0

   4,318,272,000 Bytes ( 8434125 sectors) Geometric size
   4,321,787,904 Bytes ( 8440992 sectors) True size (measured)
   4,321,787,904 Bytes ( 8440992 sectors) Reported size (Partition0)
               0 Bytes (       0 sectors) missing/wasted

---------- DMIO Kernel List ----------

Basic-info: version=21 maxspecio=512 maxio=512 maxkio=2048
Ktrans-state: dg=0.0 pendtid=0.0 state=UNKNOWN
#
# Kmem-regions:
#
id 0: (size=100)
id 1: (size=0)
#
# Disks: (cnt: 1)
#
Disk Harddisk0: dm= dgiid=0.0 darid=0.1026 dmrid=0.0
  kflag=()
  sflag=()
  vflag=(autoconfig|offline)
  publen=9223372036854775807 privlen=9223372036854775807 maxiosize=64
  type=ntdisk info=
#
# Devices: (cnt: 0)
#
#
# Groups: (cnt: 1)
#
Group NULLDG: iid=0.0 id=
  diskset=00000000-0000-0000-0000-000000000000
  configtid=0.0 logsize=0
  kflag=()
```

```
  vflag=(created|notempdb|noautoreimport|secondary)
# Group-Objects: (cnt: 0)
# End-group: NULLDG

---------- LDM Disk Header Harddisk0 ----------

ERROR: scan operation failed:
        A format error was found in the private region of the disk

---------- LDM Disk Config Harddisk0 ----------

ERROR: scan operation failed:
        A format error was found in the private region of the disk

---------- LDM Disk KLOG Harddisk0 ----------

ERROR: scan operation failed:
        A format error was found in the private region of the disk

---------- DMAdmin Simple Query ----------

dmdiag: ERROR: No disk groups loaded

---------- DMAdmin Verbose Query ----------

dmdiag: ERROR: No disk groups loaded
```

### Note

- If you want to see your output one screen at a time, type the following at the command line:

```
dmdiag |more
```

**See Also**
**Concepts**
Dmdiag Overview
Dmdiag Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Dmdiag Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DmDiag Syntax**

DmDiag uses the following syntax:

**dmdiag** [**/v**] [**/f** *FileName*]

**Parameters**

**/v**

Outputs verbose information. Verbose output includes the following fields: Computer Name and OS Version, LDM File Versions, Mount Points, Drive Letter Usage, Drive Type, Dynamic Disk Information, LDM Volume Information, Device Objects, DMIO Kernel List, LDM Disk Header HarddiskX, LDM Disk Config HarddiskX, LDM Disk KLOG HarddiskX, DMAdmin Simple Query, and DMAdmin Verbose Query.

***/f FileName***

Specifies the file name for the output file. If omitted, the file name defaults to Dmdiag.txt.

**Note**

- By default, DmDiag output includes the following:

| Field | Description |
|---|---|
| Dynamic Disk Information | Displays the disk group name and ID as well as the list of dynamic disks belonging to the group. For each disk the disk extents and the volumes to which these extents belong are displayed. |
| LDM Volume Information | Displays the list of dynamic volumes. For each volume the disk extents are displayed. |

**See Also**

**Concepts**

Dmdiag Overview
Dmdiag Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Dskprobe Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dskprobe.exe: DiskProbe**
This GUI tool is a sector editor for drives on computers running Windows XP and Windows Server 2003. A user with local Administrator rights can use DiskProbe to directly edit, save, and copy data on a physical hard drive that is not accessible in any other way.

This tool can help you prepare for disk-based problems by saving critical disk structures before problems arise. Documenting and preserving these disk structures, such as the Master Boot Record (MBR) and boot sector, provides a fallback in case of disk corruption. DiskProbe can also be used to resolve problems that you might encounter. With it, you can edit and repair disk sectors on a byte-by-byte basis if corruption occurs.

You can use DiskProbe to replace the Master Boot Record, repair damaged partition table information, and repair or replace damaged Partition Boot Sectors or other file system data. The tool can also save Master Boot Records and Partition Boot Sectors as files. They can then be replaced if the sectors become damaged at a later time. These on-disk data structures are not accessible through the file system, and so are not saved by any backup programs currently available.

### Caution

- Before you make changes with a low-level tool such as DiskProbe, make sure that a reliable, complete backup of the drive is available. Misuse of DiskProbe might make all data on the drive or volume permanently inaccessible.

  DiskProbe and other sector editors function at a level "below" the file system, so the normal checks for maintaining disk consistency are not enforced. This tool gives the user direct access to every byte on the physical disk without regard to access privilege, which makes it possible to damage or permanently overwrite critical on-disk data structures.

**Concepts**
The Master Boot Record (MBR) is created when the disk is partitioned. The MBR contains a small amount of executable code called the master boot code, the disk signature, and the partition table for the disk. At the end of the MBR is a 2-byte structure called a signature word or end of sector marker, which is always set to 0x55AA. A signature word also marks the end of an extended boot record (EBR) and the boot sector.

When you start a computer from the hard disk, the system BIOS code identifies the startup disk and reads the MBR. The master boot code in the MBR searches for the active, primary partition on the hard disk. Occasionally the MBR can become corrupted. This can be caused by human error, hardware problems, power fluctuations, viruses, and other factors.

### Note

- You can also use the Recovery Console to rewrite the MBR to resolve a corrupted MBR on a startup disk.

**Corresponding UI**
DiskProbe provides its own user interface.

**System Requirements**
The following are the system requirements for DiskProbe:

- Windows Server 2003

**Files Required**

- Dskprobe.exe
- Dskprobe.chm
- LoIOSrv.dll

**See Also**

**Concepts**

Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Ftonline Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ftonline.exe: Fault Tolerant Disk Mounter**
This command-line tool enables an administrator to mount and recover files from a fault tolerant disk created in a previous version of Windows. This tool is useful if you failed to upgrade your disks to dynamic disks, or failed to back the up your data before installing Windows Server 2003.

### Notes

- You must run FTOnline from the command window.
- FTOnline only mounts the disks as read-only to allow recovery of data. You can then back up your data, recreate the fault tolerant volumes, and restore your data.
- Disks mounted with FTOnline are not persistent. The disks will not be remounted when the computer is restarted.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Fault tolerant volumes on basic disks are no longer supported in Windows Server 2003.

Fault tolerant volumes are disks that use some type of RAID configuration to increase either performance or reliability. Windows Server 2003 supports disk striping (RAID-0), disk mirroring (RAID-1), and disk striping with parity (RAID-5).

A basic disk is a physical disk with primary and extended partitions. These disks are accessible from earlier versions of Windows.

**System Requirements**
The following are the system requirements for FTOnline:

- Windows Server 2003

**File Required**

- Ftonline.exe

**See Also**
**Concepts**
Ftonline Remarks
Ftonline Syntax
Ftonline Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview

Cabarc Overview
Bitsadmin Overview

# Ftonline Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FTOnline Remarks**

**FTOnline Usage**
This is a recovery tool for use in an emergency situation, and should not be used on a regular basis. It is used for emergency recovery of data when the computer running Windows Server 2003 does not have its data backed up.

Because Windows 2000 supports dynamic disks, you should upgrade your fault tolerant volumes to dynamic disks before installing Windows Server 2003. However, because Windows NT 4.0 does not support dynamic disks, you will need to back up your data before upgrading from Windows NT 4.0.

**See Also**
**Concepts**
Ftonline Overview
Ftonline Syntax
Ftonline Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Ftonline Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FTOnline Examples**

**Example: Mount an Inaccessible Fault Tolerant Disk**
In this example, you install a product from Windows Server 2003 and discover that a mirrored partition containing company data is assigned a letter, but is inaccessible. You look in Disk Administrator and see that it is labeled **e:**. To recover the data, type the following:

```
ftonline e:
```

This mounts the disk as a read-only volume, enabling you to back up the data. Once you back up the information, delete the volumes and recreate them as dynamic volumes, then restore the data from the backup.

**See Also**
**Concepts**
Ftonline Overview
Ftonline Remarks
Ftonline Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Ftonline Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FTOnline Syntax**
FTOnline uses the following syntax:

**ftonline** *DriveLetter*

**Parameters**
***DriveLetter***
Specifies the drive letter of the volume that needs to be mounted.

**Example**
See Ftonline Examples.

**See Also**
**Concepts**
Ftonline Overview
Ftonline Remarks
Ftonline Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Health_chk Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Health_chk.cmd: FRS Health Check**
This command-line tool retrieves FRS-related state information from the specified domain controller or server. It collects and provides administrators with in-depth data that can be used to monitor or troubleshoot an FRS system.

Health_Chk calls a series of tools for the output they produce. It can be run against a remote server, however, running the script directly on the server to be analyzed enables the tools to gather additional useful data not available when the script is run remotely.

- For monitoring, Health_Chk should be run regularly against hub servers and a small selection of spoke servers. Building a simple script that invokes these health checks and collects them to a single server for analysis is very valuable - especially when scheduled to run on a daily basis.

- For troubleshooting, Health_Chk should be run against the server that shows the symptoms of replication failures, but also, ideally, against all other related servers. It may not be possible to collect data for all servers in a large enterprise, but including direct replication partners or, if this is still too much, a good sample set of them, will provide more complete information about your topology.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
The File Replication Service (FRS) is a multi-master file replication service for NTFS that allows member machines to be interconnected in an arbitrary topology and uses last-writer-wins conflict resolution. It uses authenticated Remote Procedure Call (RPC) with encryption as the communication protocol between members, and can also support disconnected operation. FRS supports multiple replication sets per computer, each with potentially distinct topology.

**System Requirements**

- Windows 2000 Server or Windows Server 2003

- File Replication Service running on the target server(s)

- Administrator priveleges on the account running Health_Chk.

**Files Required**

- Health_chk.cmd

- Reg (Provided with the system)

- Eventquery.vbs (Provided with the system)

- Ntfrsutl Overview

- Iologsum Overview

- Connstat Overview

- Topchk.cmd

- Repadmin Overview

- Dcdiag Overview
- Netdiag Overview

**See Also**
**Concepts**
Health_chk Remarks
Health_chk Syntax
Health_chk Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Health_chk Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Health_Chk Remarks**
**Logs and Output Files**
Health_Chk retrieves FRS-related state information from the specified domain controller or server and stores this information in text files. These text files provide administrators with in-depth data that can be used to monitor or troubleshoot an FRS system.

Health_Chk can take quite a while to run. The script may collect huge amounts of data. This is especially true if you encounter problems with full staging areas, though staging area problems are not common on computers running Windows 2000 SP3 and later.

The text files are written to a specified directory using the name of the target machine.

**Note**

- Organizations typically have well-established tools for collecting event logs from Windows systems, and these should also continue to be used. The event log data collected by Health_Chk contains only a subset of the full event information.

The following lists the specific text files that Health_Chk creates, with a description of the data they contain:

| File | Description |
|---|---|
| ds_showconn.txt | The output from repadmin /**showconn**. |
| ds_showreps.txt | The output from repadmin /**showreps**. |
| evl_application.txt | Application event logs in text format. |
| evl_dns.txt | DNS server event logs in text format. |
| evl_ds.txt | Active Directory event logs in text format. |
| evl_ntfrs.txt | FRS event logs in text format. |
| evl_system.txt | System event logs in text format. |
| ntfrs_config.txt | Output from ntfrsutl **configtable**: Lists all replica sets. |
| ntfrs_ds.txt | The output from ntfrsutl **ds**: Displays the machine (FRS) topology definition in the local Domain Controller, from the perspective of the machine. |
| ntfrs_errscan.txt | A scan of error messages from the FRS debug logs. |
| ntfrs_inlog.txt | The output from ntfrsutl **inlog**: Shows inbound Change Orders (CO's). |
| ntfrs_machine.txt | A listing of free disk space on all drives. |
| ntfrs_outlog.txt | The output from ntfrsutl **outlog**: Shows outbound Change Orders (CO's). |
| Ntfrs_reg.txt | The registry data from HKLM\system\currentcontrolset\services\NtFrs and below. |
| Ntfrs_sets.txt | The output from ntfrsutl **sets**: A list of inbound and outbound connections and their status. |
| Ntfrs_sysvol.txt | The output from 'dir /s *.*' of the SYSVOL share. |
| Ntfrs_version.txt | The output from ntfrsutl **version**: Shows FRS version information. |

**Configuration Settings**
Health_Chk does not modify any configuration settings either on the target computer or in the Active Directory. The script only

reads configuration and status information.

**See Also**
**Concepts**
Health_chk Overview
Health_chk Syntax
Health_chk Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Health_chk Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Health_Chk Examples**

In this example, the remote computer, DfsRepA, is part of a DFS replica set. It is not a domain controller, so the warnings about no SYSVOL are reasonable.

```
C:\tools>health_chk logdir DfsRepA
```

When given this command, Health_Chk runs a number of tools and scripts, generates text output files in the logdir folder, and then exits.

```
Please WAIT....

 NTFRSUTL checks are running ...

 Dumping FRS inbound and outbound logs ..

 Dumping FRS registry parameters ...

SYSVOL check is running

The network path was not found.

SMB/DFS check is running

repadmin /showreps frsdoga is running

repadmin /showconn frsdoga is running

dcdiag /v /s:frsdoga is running

 Scanning eventlogs ...

 Scanning FRS debug logs for error/warning info ...

 Done ...
```

After Health_Chk completes, twenty text log files are created. The data in these generated text files can be used to monitor and troubleshoot the FRS system. For a description of the contents of the text files, see the Logs and Output Files section on the Health_chk Remarks page.

```
C:\tools>dir logdir\DfsRepA


 Volume in drive C is OS_VOL

 Volume Serial Number is 4D59-4447

 Directory of C:\Program Files\Support Tools\logdir\DfsRepA

07/21/2002  09:07 PM    <DIR>          .
07/21/2002  09:07 PM    <DIR>          ..
07/21/2002  09:06 PM               389 dcdiag.txt
07/21/2002  09:06 PM                92 ds_showconn.txt
07/21/2002  09:06 PM                91 ds_showreps.txt
07/21/2002  09:06 PM            16,216 evl_application.txt
07/21/2002  09:07 PM               161 evl_dns.txt
07/21/2002  09:07 PM               175 evl_ds.txt
07/21/2002  09:06 PM             3,998 evl_ntfrs.txt
07/21/2002  09:07 PM            18,398 evl_system.txt
07/21/2002  09:06 PM               766 net_view.txt
07/21/2002  09:06 PM            13,270 ntfrs_config.txt
07/21/2002  09:06 PM            21,562 ntfrs_ds.txt
07/21/2002  09:07 PM           731,201 ntfrs_errscan.txt
```

```
07/21/2002  09:06 PM                94 ntfrs_inlog.txt
07/21/2002  09:06 PM               337 ntfrs_machine.txt
07/21/2002  09:06 PM                95 ntfrs_outlog.txt
07/21/2002  09:06 PM             4,913 ntfrs_reg.txt
07/21/2002  09:06 PM            24,244 ntfrs_sets.txt
07/21/2002  09:06 PM                82 ntfrs_stage.txt
07/21/2002  09:06 PM                 0 ntfrs_sysvol.txt
07/21/2002  09:06 PM               465 ntfrs_version.txt
              20 File(s)        836,549 bytes
               2 Dir(s)   7,297,257,472 bytes free
```

**See Also**

**Concepts**

Health_chk Overview
Health_chk Remarks
Health_chk Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Health_chk Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Health_Chk Syntax**

Health_chk uses the following syntax:

**health_chk.cmd** *result_dir* [*target_computer*]

**Parameters**

***result_dir***

Created if it does not exist. No trailing backslash.

***target_computer***

Default is current computer. Valid computer names include the NetBIOS name with no leading slashes, or a full DNS name, xxx.yyy.zzz.com

**See Also**

**Concepts**

Health_chk Overview
Health_chk Remarks
Health_chk Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Iologsum Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Iologsum.cmd: FRS Inbound and Outbound Logs Summary Report Tool**
This command-line tool processes the output generated when the NtfrsUtl tool is run with the **inlog**, **outlog**, or **idtable** parameters.

Running NtfrsUtl **inlog** or NtfrsUtl **outlog** creates a summary inbound or outbound NtfrsUtl log report showing the pending change orders for the replica member. Running NtfrsUtl **idtable** creates a report with version information for each file in the replica tree; this version information tracks the last change order accepted for that particular file or directory. Running IOLogSum then sorts and formats this NtfrsUtl output.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
The File Replication Service (FRS), first introduced in Windows 2000, is an essential element of the Microsoft Active Directory technologies. The Active Directory service uses FRS for SYSVOL replication. FRS can also be used with the Microsoft Distributed File System (DFS) technologies for automatic replication of content between alternating targets. For more information on FRS and Active Directory replication please see the Windows Server 2003 Technical Reference.

A replica set consists of two or more computers configured to replicate the contents of a file folder. The individual computers are referred to as replica members.

FRS maintains logs of *change orders* that are either received from replication partners (the inbound log), or are to be sent to outbound replication partners (the outbound log).

The inbound log stores pending change orders to be processed. As entries are processed, FRS sends acknowledgment to the inbound partners. Data stored in the inbound log includes the change order's globally unique identifier (GUID), file name, object ID, parent object ID, version number, and event time.

The outbound log stores pending change orders to be sent to outbound partners. Change orders remain in the outbound log until all outbound partners receive and acknowledge the change. Data stored in the outbound log is the same as that stored in the inbound log. Also in the outbound log is the leading (next change) and trailing (last acknowledged) index for each partner.

The outbound logs can become quite large, particularly when replica members are unavailable, links between replicas are slow, replication hours are restricted, or a large number of changes occur. For example, even when a member for a given replica set is unavailable, outbound log entries are maintained until the member becomes available. When the changes are then sent, the inbound partner sends all changes in log file order.

**System Requirements**

- Windows 2000 Server or Windows Server 2003
- Perl 5.005 521e or 522e

**Files Required**

- Iologsum.cmd
- Ntfrsutl.exe

**See Also**
**Concepts**
Iologsum Remarks
Iologsum Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview

Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Iologsum Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IOLogSum Remarks**
**Before Using IOLogSum**
If Perl is not installed on your system, if you are connected to the Internet, you can download a free version from the ActiveState Web site.

**IOLogSum Report**
When used with NtfrsUtl **inlog** and NtfrsUtl **outlog**, the IOLogSum report can be used to determine which files are involved in FRS replication, and what has caused them to replicate. When used with NtfrsUtl **idtable**, IOLogSum can generate a report describing the state of all files and folders known to an FRS member.

Familiarity with FRS concepts and terminology is required in order to interpret IOLogSum reports.

Each IOLogSum report starts by identifying the name of the file that it processed, and the sort options used to generate the report. Then, for every replica set in which the computer participates, the report lists one row per change order.

The following table lists and describes each column in the report.

| Column Heading | Description |
|---|---|
| **SeqNumber** | Specifies the unique sequence number for the change order. For NtfrsUtl **idtable** reports, this is always reported as 00000000. |
| **EventTime** | Specifies the USN Journal Entry Timestamp; in other words, the time the file was changed on the originating machine. |
| **FileVersNum** | Specifies the file version number, incremented each time the file is closed. |
| **FileUsn** | Specifies the USN that the file must match on the Fetch request. |
| **FileSize** | Specifies the file size (in KB). |
| **FileGuid** | Specifies the object ID of the file or directory. For brevity, only the last 8 digits are printed. |
| **OriginatorGuid** | Specifies the GUID of the originating member. For brevity, only the last 8 digits are printed. The originator GUID can be mapped back to the server name by searching the NtfrsUtl **sets** output. |
| **Cxtion** | Specifies the GUID of the remote connection – for brevity, only the last 8 digits are printed |
| **FileName** | Specifies the Unicode file name. |

| | |
|---|---|
| **Flags** | Specifies a valid Change Order Flag. Valid flags are shown in the following list.<br>• **Abort** - Set when CO is being aborted.<br>• **VVAct** - Set when VV activate request is made.<br>• **Content** - Valid content command.<br>• **Locn** - Valid location command.<br>• **LclCo** - CO is locally generated.<br>• **Retry** - CO needs to retry.<br>• **InstallInc** - Local install not completed.<br>• **Refresh** - CO is an upstream-originated file refresh request.<br>• **OofOrd** - Do not check/update version vector.<br>• **NewFile** - If CO fails, delete IDTable entry.<br>• **DirectedCo** - This CO is directed to a single connection.<br>• **DemandRef** - CO is a downstream demand for refresh.<br>• **VVjoinToOri** - CO is from vvjoin to originator.<br>• **MorphGen** - CO generated as part of name morph resolution.<br>• **MoveinGen** - This CO was generated as part of a sub-dir MOVEIN.<br>• **OidReset** - This CO reset the Object ID (OID) back to the FRS defined value.<br>• **CmpresStage** - The stage file for this CO is compressed.<br>• **MD5** - The MD5 hash of the file is included.<br>• **DELETED** - The file has been deleted. On **idtable** reports, this indicates that the record is a tombstone, ensuring the file remains deleted throughout the replica set. |

| | |
|---|---|
| **LocationCmd** | Specifies the Change Order Location command. Valid Location Commands are shown in the following list.<br>• **CREATE** - Create a new file or directory.<br>• **DELETE** - Delete a file or directory.<br>• **MOVEIN** - Rename a file or directory from a point outside the Replica Tree, into a Replica Tree on the same volume.<br>• **MOVEOUT** - Rename a file or dir from a point inside the Replica Tree to a point outside the Replica Tree on the same volume.<br>• **MOVERS** - Rename a file or directory from one Replica Tree to another Replica Tree on the same volume (currently unsupported).<br>• **MOVEDIR** - Rename a file or directory from one directory to another in the same Replica Tree. |

| | |
|---|---|
| **ContentCmd** | Specifies the Change Order content changes from the USN reason. Flags set in the NTFS change log describing modifications to the file.<br>• **Close** - Change log close record.<br>• **Create** - File or directory was created.<br>• **Delete** - File or directory was deleted.<br>• **RenNew** - File or directory was renamed.<br>• **DatOvrWrt** - Main file data stream was overwritten.<br>• **DatExt** - Main file data stream was extended.<br>• **DatTrunc** - Main file data stream was truncated.<br>• **Info** - Basic info change (attrib, last write time, etc.).<br>• **Oid** - Object ID change.<br>• **StreamNam** - Alternate data stream name change.<br>• **StrmOvrWrt** - Alternate data stream was overwritten.<br>• **StrmExt** - Alternate data stream was extended.<br>• **StrmTrunc** - Alternate data stream was truncated.<br>• **EAChg** - Extended file attribute was changed.<br>• **Security** - File access permissions changed.<br>• **IndexableChg** - File change requires re-indexing.<br>• **HLink** - Hard link change.<br>• **CompressChg** - File compression attribute changed.<br>• **EncryptChg** - File encryption changed.<br>• **Reparse** - Reparse point changed. |

By default, the records are sorted by sequence number, but it is possible to sort by other columns.

- Sorting by File GUID makes it easy to detect repetitive changes to single files and the times at which they occurred (for example, once per hour).
- Sorting by Event Time is useful for identifying older change orders that may indicate a particular partner has been unavailable for a while. The IOLogSum report, coupled with the report that Connstat Overview generates, can tell you what change orders are pending for the delinquent partner.

**See Also**
**Concepts**
Iologsum Overview

Iologsum Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Iologsum Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IOLogSum Syntax**

IOLogSum uses the following syntax:

**iologsum** [**-sort=***Option*] *DataFile*

**Parameters**

**-sort= *Option***

Specifies a sort order for the IOLogSum report. The following table lists are valid values for *Option*:

| Option | Description |
|---|---|
| **seqnum** | Sorts output by sequence number. This is the default sort order. |
| **version** | Sorts output by file version number. |
| **filename** | Sorts output by file name. |
| **size** | Sorts output by file size. |
| **fileguid** | Sorts output by file GUID. |
| **origguid** | Sorts output by originator GUID. |
| **cxtion** | Sorts output by cxtion name/GUID. |
| **eventtime** | Sorts output by event time. |

***DataFile***

Specifies a file that was created by running NtfrsUtl with the **inlog**, **outlog** or **idtable** parameters. This is a required parameter.

Note

- The Perl interpreter should be installed and in the path for this command to function. See Iologsum Remarks for details on obtaining Perl.

**Sample Usage**

```
iologsum ntfrs_inlog.txt > in.txtiologsum -sort=fileguid ntfrs_outlog.txt > out.txt
```

**See Also**

**Concepts**

Iologsum Overview
Iologsum Remarks
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview

Cabarc Overview
Bitsadmin Overview

# Rsdiag Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Rsdiag.exe: Remote Storage Diagnostic Utility**
This command-line tool examines Remote Storage (HSM) databases and displays diagnostic information in text format about migration jobs, managed volumes of the version of NTFS file system used in Windows Server 2003, and physical media, as well as other Remote Storage information useful for system analysis.

Remote Storage stores its data in a series of databases and collection files in the *Systemroot*\System32\RemoteStorage directory. The data enables the Remote Storage service to associate file placeholders and file data, and to track managed and unmanaged volumes, files on managed volumes, jobs, and physical media. RsDiag lets you display this data in easy-to-read format on the console and in text files.

You can use RsDiag to display the following information:

- The contents of the Remote Storage database (EngDb), the FSA database (FsaDb: premigrated files), the FSA temporary unmanaged database, and the collection files that store the configuration options for each volume (RsFsa.col) and for the installation (RsEng.col).
- Version information for the configuration collection files.
- Volumes that are eligible for Remote Storage management.
- Properties of managed volumes.
- Properties of physical media used for Remote Storage.
- The physical location of the file data for premigrated and truncated files.
- Remote Storage jobs scheduled on the system, including migration, validation, free space, managed, and unmanaged jobs. You can also use RsDiag to stop a job in progress.
- The internal recall queue. You can use RsDiag to delete a job from the recall queue, preventing a file from being recalled.
- Descriptions of error codes generated by Rss.exe and found in the Service and Driver logs in Event Viewer.

RsDiag works with the Rsdir Overview to help diagnose problems in the Remote Storage databases and, if necessary, recover lost data.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Remote Storage extends the disk space on the computer by copying files to remote storage based on criteria that you establish, including the location and size of the files, and the amount of free space on the volume. Files are recalled from remote storage automatically when they are accessed.

Files stored remotely appear unchanged in Microsoft Windows Explorer. Users open them by the same methods, but the files do not consume any local disk space. Instead, Remote Storage uses NTFS Reparse Points as placeholders for the files on the local volumes. The Reparse Points store information about the file, including its actual location in remote storage.

**System Requirements**
The following are the system requirements for RsDiag:

- This version of RsDiag runs only on Windows Server 2003. A version of RsDiag for Windows 2000 is included in the Support Tools directory of the Windows 2000 Server operating system.

- The user must be a member of the Administrators group on the computer.
- RsDiag is a local computer tool. You cannot use it to display information about volumes on a remote computer.

**Files Required**

- Rsdiag.exe
- Rstools.dll

The Rstools.dll library must be registered on the system. The Support Tools installation process registers the DLL. However, if you copy these files to another computer, you must register the DLL manually by typing the following at the command prompt:

```
regsvr32 RsTools.dll
```

### For More Information

For more information about Remote Storage Server, download the "Enterprise Class Storage in the Next Generation of Windows NT" document from the Microsoft Windows 2000 Server Web site.

**See Also**
**Concepts**
Rsdiag Syntax
Rsdiag Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Rsdiag Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RsDiag Examples**
**Example 1: Display the engine database**
The following command displays diagnostic information about the Remote Storage Engine database (EngDb).

```
rsdiag /d e
```

The Engine database consists of the following three tables:

- The **Segment Table** associates the placeholder for a file to the location in remote storage where the file data is stored. This is done by mapping the BAG (migration group) identifier and location to the media storage.
- The **Media Table** associates the media identifier to an external media name that is used by the media management subsystem.
- The **BAG Information Table** contains global information about each BAG (migration group). There is one entry for each BAG.

**Segment Table**

| Short Name | Long Name | Description |
|---|---|---|
| **BagId** | BAG ID | A GUID that uniquely identifies the migration group of the file. A BAG is a group of files migrated in the same job and stored in a single data set (MTF/BKF) on the remote media. Remote Storage uses the BAG ID to locate and recall files. |
| **StartLoc** | BAG Segment Start Location | The starting position of this segment within the BAG. |
| **SegLen** | BAG Segment Length | The length of this segment of the BAG. This length can extend to the entire BAG or only a portion of the BAG. |
| **LocType** | Location Type | Indicates whether this location is a media location or a logical redirection to another BAG and location. |
| **PrimPos** | Primary Location Position Specifier | For media locations, this is a media identifier. For logical redirections, this is a BAG identifier. |
| **SecPos** | Secondary Location Position Specifier | For media locations, a starting location on media. For logical redirections, a BAG start location. |
| **IDB SeqNum** | | The database sequence number, indicating the sequence in which this record was originally added to the database |

**Segment Table dump example**

```
00000  BagId = {b010419f-fb1b-11d1-a1ba-00a0c9190447}, StartLoc = 3072, SegLen = 52224,
LocType = 0, PrimPos = {b0104199-fb1b-11d1-a1ba-00a0c9190447}, SecPos = 0  (IDB SeqNum
=      1)
```

**Media Table**

| Short Name | Long Name | Description |
|---|---|---|

| | | |
|---|---|---|
| **Id** | Media ID | |
| **NtmsId** | | The media handle of the media identifier for the Media Service. This is a Removable Storage Management (RSM) Media ID. |
| **StoragePoolId** | Storage Pool ID | A unique storage pool identifier. |
| **IDB SeqNum** | | The database sequence number, indicating the sequence in which this record was originally added to the database. |

**Media Table dump example**

```
00000  id = {b0104199-fb1b-11d1-a1ba-00a0c9190447} ntmsId = {da96022d-fb19-11d1-a1ba-
00a0c9190447} StoragePoolId = {e5390a42-fb19-11d1-a1ba-00a0c9190447} (IDB SeqNum =
1)
```

**BAG Information Table**

| Short Name | Long Name | Description |
|---|---|---|
| **Type** | | The record type. |
| **BagId** | BAG ID | A GUID created when the BAG is initially created. |
| **BirthDate** | BAG Birthdate | The date and time the BAG was created. |
| **Len** | Original BAG Length | The actual length of the BAG at create time. |
| **Type** | BAG Type | Not currently used. |
| **VolId** | Volume ID | The managed volume from which the data in this BAG was originally taken. |
| **DeletedAmount** | Deleted BAG amount | The amount of space within the BAG that has been deleted. Not currently used. |
| **RemoteDataSet** | Remote Data Set Number | The unique ID used on media to identify this session. |
| **IDB SeqNum** | | The database sequence number, indicating the sequence in which this record was originally added to the database. |

**Bag Information dump example**

```
00000 , Type = 2 BagId = {b010419f-fb1b-11d1-a1ba-00a0c9190447}, BirthDate = 06/03/1998
20:42:20, Len = 47051200, Type = 0, VolId = {9db6f302-fb1b-11d1-a1ba-00a0c9190447},
DeletedAmount = 0, RemoteDataSet = 1 (IDB SeqNum =      1)
```

**Example 2: Display the FSA database**

The following command displays the File System Agent (FSA) database. This database contains an entry for each premigrated file.

```
rsdiag /d f
```

The following is a sample display of an FSA database generated by using the **/d f** parameter:

```
Dump of DB: E:\WINDOWS\system32\RemoteStorage\FsaDb\{C0835F1E-01C1-804D-80B9-
FB1600000000}
  version = 1, # record types = 2
RecType =        1, Flags = 00000001, MaxSize =    65760, # keys =    3
  KeyType =        2, Size =       33, Flags = 00000001
  KeyType =        1, Size =       33, Flags = 00000002
  KeyType =        3, Size =       33, Flags = 00000001
RecType =        2, Flags = 00000001, MaxSize =    65708, # keys =    1
  KeyType =        1, Size =      254, Flags = 00000002
```

```
*** Dump of records of Type = 1 ***
00000  AccessTime = 12/17/2001 21:27:33 BagId = {7A4CBCE5-A0DB-4A74-9B0A-7B9AF561850F},
BagOffset = 1061888, FileId = 36028797018964002 Waiting for close = FALSE, Offset = 0,
Path = \test2.bkf, Size = 4002816 RecallTime =   01/01/1601 00:00:00, File USN = 1168
(IDB SeqNum =      2)
00001  AccessTime = 12/20/2001 17:05:56 BagId = {7A4CBCE5-A0DB-4A74-9B0A-7B9AF561850F},
BagOffset = 3072, FileId = 7921831744544702504 Waiting for close = FALSE, Offset = 0,
Path = \blslsslsl.txt, Size = 1057792 RecallTime = 01/01/1601 00:00:00, File USN = 920
(IDB SeqNum =      1)
00002  AccessTime = 12/20/2001 17:05:56 BagId = {7A4CBCE5-A0DB-4A74-9B0A-7B9AF561850F},
BagOffset = 5065728, FileId = 281474976710697 Waiting for close = FALSE, Offset = 0,
Path = \vrdsgsre.txt, Size = 2884608 RecallTime =  01/01/1601 00:00:00, File USN = 1424
(IDB SeqNum =      3)

*** File System Agent Database Dump Complete ***
```

The following table lists and describes the data fields in the record dump.

| Property | Description |
|---|---|
| **AccessTime** | The date and time that the file was last accessed, in Coordinated Universal Time (UTC). Migrating and recalling files does not change this value. |
| **BagID** | A GUID that uniquely identifies the migration group of the file. A BAG is a group of files migrated in the same job and stored in a single data set (MTF/BKF) on the remote media. Remote Storage uses the BAG ID to locate and recall files. |
| **BagOffset** | The offset of the file into the BAG. |
| **FileId** | The remote storage identifier for the file. |
| **Waiting for Close** | This field is not currently used. |
| **Offset** | This field is not currently used. |
| **Path** | The path for this file. This information may be compressed. |
| **Size** | The logical size of the premigrated data. |
| **RecallTime** | The date and time that the system most recently attempted to retrieve this file from remote storage, in Coordinated Universal Time (UTC). Files that have never been recalled have a recall date of 01/01/1601. |
| **File USN** | The Update Sequence Number (USN) for this file. |

**Example 3: Cancel a validation job**

The following series of commands finds a running validation job and then cancels the job before it completes. This procedure can be used on all types of Remote Storage jobs, such as migration jobs and jobs that force truncation to create free space on the volume.

The first command lists the Remote Storage jobs scheduled on the system:

```
rsdiag /j
```

In response, RsDiag displays the job list:

```
Name                                          JobState   IsUsers Done   DoneOk
Manage                                        Idle       False   Ok     Ok
Remote Storage Automatic Validate D           Active     Ok      False  Ok
Remote Storage Remove D                       Idle       Ok      Ok     Ok
Remote Storage Automatic Validate E           Idle       Ok      Ok     Ok
Remote Storage Validate E                     Idle       Ok      Ok     Ok
Remote Storage Copy Files E                   Idle       Ok      Ok     Ok
Remote Storage Free Disk Space E              Idle       Ok      Ok     Ok
```

The display reveals that the "Remote Storage Automatic Validate D" job is running.

The following command displays more detailed information about this particular job:

```
rsdiag /j "Remote Storage Automatic Validate D"


  Name              : Remote Storage Automatic Validate D
  Identifier        : {e65a104f-2753-49e2-bf8d-41573e1bbc27}
  Hsm Identifier    : {4d6e0baf-6d9c-4048-af3f-ca10d94cc8f9}
  Run Identifier    : 1
  Max Sessions      : 2
  Is User Defined   : Ok
  Is Active         : True
  Did Finish        : False
  Did Finish Ok     : False

    Resource Name     : SERVER01\NTFS\D:
    State             : Done
    Starting Path     : \
    Current Path      : \
    Start Time        : 03/15/2001 20:22:00
    Stop Time         : 03/15/2001 20:22:00
    Sub-Run Id        : 0

      Phase Name        : Scan
      State             : Done
      Media State       : Idle
      Priority          : Normal
      Start Time        : 03/15/2001 20:22:00
      Stop Time         : False
      Elapsed Time      : 0 days, 00:00:03
      Items             : 0
      Size              : 0
      Skipped Items     : 0
      Skipped Size      : 0
      Error Items       : 0
      Error Size        : 0
```

The following command stops the job:

```
rsdiag /c "Remote Storage Automatic Validate D"
```

Remember that stopping a job does not remove the job from the schedule. Also, stopping a job does not reverse actions already completed by the job. For example, if you cancel a free-disk-space job in progress, it does not restore data to files that were truncated earlier in the job.

**Example 4: Determine the tape on which a file is stored**

The following series of commands finds the data for a file on a physical tape. The first command uses Rsdir Overview, a tool that enumerates the files in a directory and adds Remote Storage properties to the attributes typically displayed.

```
C:\Inside Win2K>rsdir


Showing information for <\\?\C:\Inside Win2k\*.*>

Status        Attributes    Physical     Logical        Migrate Time File Name
------        ----------    --------     -------        ------------ ---------
non-HSM      --D--------           0           0                  NA .
non-HSM      --D--------           0           0                  NA ..
non-HSM      A-------R--          47          47                  NA Autorun.inf
non-HSM      --D-----R--           0           0                  NA Debuggers
non-HSM      --D-----R--           0           0                  NA Ebook
truncated    A-----OPR-Z           0       10004 08\11\2001 09:07:27 Readme.txt
non-HSM      --D-----R--           0           0                  NA Setup
non-HSM      --D-----R--           0           0                  NA StartCD
truncated    A-----OPR-Z           0       28672 04\27\2001 17:27:12 StartCD.exe

                3 File(s)              47 bytes
                4 Dir(s)            38676 bytes required to recall offline files
```

This display shows that the Readme.txt file is truncated and was migrated on 8/11/2001.

The next command displays the location of the Readme.txt file on a remote physical storage medium. This command uses the **/w** parameter to display the location of the file in secondary storage. It includes the complete path to the file and, because the directory name includes spaces, it encloses the path in quotation marks.

```
rsdiag /w "c:\inside win2k\readme.txt"
```

```
        File <c:\inside win2k\readme.txt> is truncated.
        Located on media <000302> in dataset 28
        Bag ID: <{7eae2dea-6634-4ae5-982a-c987c01ffa38}>
```

The output displays the BAG ID, the internal media name (00302), and the dataset in which data for this file is stored (28).

The final command displays all media used by Remote Storage on this system:

```
rsdiag /s
```

```
Showing:  <RS-SERVER01-4>
        Description     :  RS-SERVER01-4
        Name            :  000302
        Media ID        :  {4464192f-8164-422d-8541-43577071cb24}
        RSM ID          :  {7cb7bd09-888b-4021-bfda-400cfeb2e792}
        Media type      :  4
        Last update     :  12/20/2001 10:04:03
        Last error      :  Ok
        Recall only     :  FALSE
        Capacity        :  19036958720
        Free Bytes      :  3211743232
        Next Data Set   :  81
        Recreate        :  FALSE

Showing:  <RS-SERVER01-3>
        Description     :  RS-SERVER01-3
        Name            :  000301
        Media ID        :  {f926fccb-aa62-4116-8d4e-1d5b7eb8a4f3}
        RSM ID          :  {85e3bc66-d196-49f0-abf0-eb5c15ba2609}
        Media type      :  4
        Last update     :  08/15/2001 09:06:21
        Last error      :  Ok
        Recall only     :  FALSE
        Capacity        :  19036958720
        Free Bytes      :  20130557952
        Next Data Set   :  4
        Recreate        :  FALSE
```

The display reveals that the media known internally by the media name 00302 is a tape (media type 4) commonly known as RS-Server01-4. Optical media are listed as media type 3. Therefore, the data for the Readme.exe file is stored on the RS-SERVER01-4 tape.

**Example 5: Display the version information for a collection file**

The following command displays version information about Remote Storage components:

```
rsdiag /i
```

The following fields are included in the output:

| Name | Description |
|---|---|
| **Title** | Component title |
| **Registry Name** | Component title in registry |
| **Identifier** | Unique identifier of this resource |
| **DB Version** | Component's database version |
| **RS Version** | Remote Storage build number |

| | |
|---|---|
| **Windows Ver** | Windows version |
| **Windows Build** | Windows build number |
| **DB Path** | Path to the component collection file |

The command yields the following output:

```
RsTools
  Build        :  601

FSA
  RegistryName   :  Remote_Storage_Server
  Identifier     :  {4d6e0baf-6d9c-4048-af3f-ca10d94cc8f9}
  Db Version     :  1
  RS Version     :  611
  Windows Ver    :  5.1
  Windows Build  :  3590
  DbPath         :  C:\WINDOWS\system32\RemoteStorage\RsFsa.col


HSM
  RegistryName   :  Remote_Storage_Server
  Identifier     :  {4d6e0baf-6d9c-4048-af3f-ca10d94cc8f9}
  Db Version     :  3
  RS Version     :  611
  Windows Ver    :  5.1
  Windows Build  :  3590

  DbPath         :  C:\WINDOWS\system32\RemoteStorage\RsEng.col
```

**See Also**

**Concepts**

Rsdiag Overview

Rsdiag Syntax

Rsdir Overview

# Rsdiag Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RsDiag Syntax**

**rsdiag** [**/d** *DBType OutputFile*] [**/i**] [**/m**] [**/v** [*DriveLetter*]] [**/s**] [**/w** *Path**FileName*] [**/j** [*JobName*]] [**/c** *JobName*] [**/r** [**/f**]] [**/x** *RecallJob*] [**/e** *ErrorCode*] [**/t**] [**/?**]

**Parameters**

**Note**

- Except as noted, you can use only one RsDiag parameter with each command.

**/d *DBType OutputFile***

Writes the specified Remote Storage database to the specified text file. Typically, support engineers use these files to identify or resolve metadata errors.

*DBtype* identifies the database. This variable is required with **/d**. The following table lists valid values for *DBType*.

| DBType | Database type | Description |
|---|---|---|
| **e** | Engine database (EngDb) | The primary remote storage database for the installation. |
| **f** | FSA database (FsaDb) | The FSA database stores a list of premigrated files on each volume.<br><br>In *premigrated files*, the file data has been copied to remote storage, but it also remains on the local volume as long as space is available. Files that were migrated and then recalled are considered to be premigrated. |
| **u** | FSA temporary unmanaged database | A temporary database that stores FSA data for a volume during the unmanaged process. When the process is complete, this database is deleted. |
| **a** | FSA persistent (RsFsa.col) | A collection file that stores the configuration options set for each volume. |
| **n** | Engine persistent (RsEng.col) | A collection file that stores configuration options set for the Remote Storage installation. |

*OutputFile* is the full path and file name of the text file that RsDiag creates to store this display; for example, C:\RSS\myrssdb.txt. The full path is required, even when the file is in the local directory. All directories in the path must exist before you run the command. Use quotation marks to enclose names that contain spaces.

**/i**

Lists and displays version information for the collection (.col) files on this system. The collection files store configuration information for Remote Storage. To display the collection files, use the **/d a** or **/d n** parameter.

**/m**

Lists volumes on this system that are eligible for Remote Storage management.

**/v [*DriveLetter*]**

Lists the managed volumes on this system. With a drive letter, **/v** displays detailed information about the specified managed volume.

**/s**

Lists and describes the physical storage media used on this system.

**/w *Path* \ *FileName***

Finds the specified managed file in remote storage. This parameter displays the physical location of any file on a managed

volume. The display includes the physical storage medium on which the most current version of the specified file resides, and the data set of which the file is a member.

Specify the full path and file name of the file. The full path is required even when the file is in the local directory. Use quotation marks to enclose names that contain spaces.

To see which files are migrated, or to determine the Remote Storage status of a file, use Rsdir Overview.

**/j [*JobName*]**
Lists the Remote Storage jobs scheduled on this system. With a job name, **/j** displays detailed information about the specified job. Use quotation marks to enclose job names that include spaces.

This display lets you predict the duration of a job and compare the run time of different jobs.

**/c *JobName***
Stops a Remote Storage job in progress. This parameter stops the current run of the specified job. It does not affect the job schedule. To display the jobs scheduled for the system, use the **/j** parameter.

**Caution**

- This command stops the job, but does not cancel or reverse actions already completed by the job. Files processed before the job was stopped are not restored to their original state.

**/r [/f]**
Displays the jobs in the Remote Storage server's internal recall queue. The display includes the name and logical size of each recalled file. The **/f** (full) parameter adds the state of the recall, the user who initiated the recall, and the file offset into the media to the display.

**/x *RecallJob***
Deletes the specified recall job from the Remote Storage server's internal recall queue. Only pending recalls can be cancelled. To display the recall queue, use **/r**.

*RecallJob* is the name of the job that recalls the file. It consists of the full path to the file being recalled, including the server name. This value appears in the **Queued Recall** field of the **/r** display.

**/e *ErrorCode***
Explains HRESULT and WIN32 error codes. Type an *ErrorCode* value in decimal or hexadecimal format. The default is decimal. For hexadecimal error codes, prepend **0x** to the error code, for example, **0x81000005**.

**/t**
Reloads the trace files. This setting is for internal use only.

**/?**
Displays help information for the tool.

**See Also**
**Concepts**
Rsdiag Overview
Rsdiag Examples
Rsdir Overview

# Rsdir Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Rsdir.exe: Remote Storage File Analysis Utility**
This command-line tool displays the Remote Storage status of files on local NTFS file system volumes of servers running Windows Server 2003. It is designed to be used on volumes managed by the Remote Storage service. RSDir supplements the enumeration feature of Dir.exe with information about the Remote Storage properties of the objects, such as whether the files reside in local or remote storage, and details of their migration and location.

RSDir displays the following information about files on volumes formatted with NTFS:

- File name and path
- Remote Storage status (Non-Hierarchical Storage Management (HSM), Premigrated, or Truncated)
- File attributes
- Size of files on local and remote media, including the size of the unnamed data stream (file content) and the size of file metadata, such as Reparse Points
- Time stamps: Create, Access, Modify, Migrate, Recall
- IDs: File ID, HSM ID, BAG ID
- Cyclic Redundancy Check (CRC) data: Algorithm type and value for each file and for each migrated data set (Microsoft Tape Format (MTF) or .BKF files)
- Offset of files in remote media

RSDir compiles its display from information stored in Reparse Points on the local volume. You can use this information to monitor the use of Remote Storage and to detect problems with Remote Storage management on the system.

You can use RSDir with the Rsdiag Overview to identify problems in the Remote Storage databases and recover lost data. RSDir and Rsdiag complement Rss.exe, the command-line version of the Remote Storage snap-in.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Remote Storage extends the disk space on the computer by copying files to remote storage based on criteria that you establish, including the location and size of the files and the amount of free space on the volume. Files are recalled from remote storage automatically when they are accessed.

Files stored remotely appear unchanged in Windows Explorer and users use the same methods to open them, but they do not consume any local disk space. Instead, Remote Storage uses NTFS Reparse Points as placeholders for the files on the local volumes. The Reparse Points store information about the file, including its actual location in remote storage.

RSDir displays information stored in the Reparse Points in a formatted display at the command line.

**System Requirements**
The following are the system requirements for RSDir:

- This version of RSDir runs only on Windows Server 2003 operating systems. A version of RSDir for Windows 2000 is included in the Support Tools directory of Windows 2000 Server.
- The user must be a member of the Administrators group on the computer.

- RSDir is a local computer tool. You cannot use it to display information about volumes on a remote computer.

**File Required**

- RSDir.exe

**See Also**
**Concepts**
Rsdir Remarks
Rsdir Syntax
Rsdir Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Rsdir Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RSDir Remarks**
**General Information**

- RSDir displays information about the local computer only. You cannot use it to display information about files or directories on other computers.
- When you run RSDir without a path, or with an invalid or improperly specified path, RSDir displays the Remote Storage status of files in the local directory.
- RSDir displays time stamps in Coordinated Universal Time (UTC), not the time of the local computer.
- RSDir uses the standard set of NTFS file attributes. Truncated (remote) and premigrated files have the Archived (A), Reparse points (P), and Sparse (Z) attributes set. The Sparse attribute, which activates an efficient storage method for zero-filled files, is set for premigrated files even though they are not (yet) zero-filled. Only truncated files are set to Offline (O).

**RSDir Basic Display**
The following table lists and describes the fields in the basic RSDir display.

| Property | Description |
|---|---|
| Status | **Non-HSM**: The file resides only on the local volume, either because the volume is not managed by Remote Storage (Hierarchical Storage Management), or the file has not been migrated.<br><br>**Premigrated**: The file data has been copied to remote storage, but remains on the local volume as long as space is available. Files that were migrated and then recalled are considered to be premigrated.<br><br>**Truncated**: The file data resides in remote storage. A sparse placeholder file remains on the local volume with a Reparse Point that stores the actual location of the file in remote storage. |
| Attributes | **A** Archive<br><br>**C** Compressed<br><br>**E** Encrypted<br><br>**D** Directory<br><br>**H** Hidden<br><br>**O** Offline<br><br>**P** Reparse Points<br><br>**R** Read-only<br><br>**S** System File<br><br>**Z** Sparse |
| Physical | The physical file size, which is the number of bytes that the unnamed data stream (file content) occupies on the local volume. For files migrated to remote storage, this value is 0. |
| Logical | The logical file size, in bytes, which is the actual size of the unnamed data stream (file content), regardless of its physical location. For local files (non-HSM and premigrated), the physical size is equal to the logical size. |

| | |
|---|---|
| Migrate Time | The date and time that the file was most recently copied to remote storage, in Coordinated Universal Time (UTC). Recalling files does not change this value. |
| File Name | The name of the file. |
| Bytes | The combined size of files stored on the local volume, that is, non-HSM and premigrated files. |
| Bytes required to recall offline files | The combined size of files stored remotely, that is, truncated files. |

**RSDir Extended Display (/f)**

The **/f** parameter adds the following fields to the RSDir display.

| Property | Description |
|---|---|
| Create Time | The date and time that the file was created, in Coordinated Universal Time (UTC). |
| Last Access Time | The date and time that the file was last accessed, in Coordinated Universal Time (UTC). Migrating and recalling files does not change this value. |
| Last Modify Time | The date and time that the file was last changed, in Coordinated Universal Time (UTC). Migrating and recalling files does not change this value. |
| RC | (Recall Count) The number of times the file was recalled from remote storage. |
| Last Recall Time | The date and time that the file was most recently retrieved from remote storage, in Coordinated Universal Time (UTC). Files that have never been recalled have a recall date of 01/01/1601. |
| HSM ID | A GUID that uniquely identifies the Remote Storage installation. This ID is used to distinguish installations and resolve conflicts when Remote Storage has been reinstalled as part of a system recovery or reinstallation. |
| Bag ID | A GUID that uniquely identifies the migration group of the file. A *Bag* is a group of files migrated in the same job and stored in a single data set (MTF/BKF) on the remote media. Remote Storage uses the BAG ID to locate and recall files. |
| File ID | This field is not used. Files that are premigrated or truncated have a File ID of 0. |
| File Version ID | A numeric key used to track changes to premigrated and truncated files. |
| Data Stream | A unique identifier for the named data stream. All unnamed data streams (file content) have a Data Stream identifier of 0. |
| CRC Type | The algorithm used to perform the Cyclic Redundancy Check (CRC) on the data stream. Before truncating files, Remote Storage verifies that there are no differences between the local and remote copies of premigrated files. |
| Data Stream CRC | The value that the CRC generated for the data stream. |
| DS Size | The size of the migrated data stream, in bytes. Because Remote Storage migrates only the unnamed data stream (file content), this value is equal to the Data Size. |
| DS Start | The position of the unnamed data stream in remote storage. This value is the offset of the unnamed data stream (file content) from the beginning of the migrated data for this file. Because only the unnamed data stream is migrated, this value is always zero (0). |

| | |
|---|---|
| Data Size | The size that the unnamed data stream (file content) occupies in remote storage. Because Remote Storage always migrates file content as a unit, this value is either **N/A**, meaning zero (for local files), or it is equal to the logical file size (for truncated files). |
| Data Start | The position of the unnamed data stream in remote storage. This value is the offset of the unnamed data stream (file content) from the beginning of the file. Although only the unnamed data stream is migrated, this value is not zero because the data stream is preceded by MTF block headers. |
| File Size | The space that the file and its metadata occupy in remote storage, in bytes, also known as *size on disk*. In addition to the unnamed data stream (file content), this value includes the size of file system metadata, such as Reparse Points, and named streams. |
| File Start | Position of the file in remote storage. This value is the offset of this file, including file metadata, from the beginning of the remote data set for its *Bag* (migration group). |
| Ver Type | The algorithm used to perform a Cyclic Redundancy Check (CRC) on the entire migrated data set (the MTF/BKF, not just a particular data stream). |
| Ver Info | The value that the CRC generated for the entire migrated data set. |

**See Also**
**Concepts**
Rsdir Overview
Rsdir Syntax
Rsdir Examples
Alphabetical List of Tools
Topchk.cmd
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Rsdir Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RSDir Examples**

**Example 1: Display the status of files in the local directory**

The following command requests a basic display of the Remote Storage status of files in the directory from which the command is issued. Because the command uses the default value for all parameters, no parameters are needed.

```
C:\Inside Win2k\StartCD>RSDir
```

In response, RSDir displays the following information. For descriptions of the fields in the basic display, see Rsdir Remarks.

```
Showing information for <\\?\C:\Inside Win2k\StartCD\*.*>

Status        Attributes    Physical     Logical        Migrate Time File Name
------        ----------    --------     -------        ------------ ---------
non-HSM      --D-----R--           0           0                  NA .
non-HSM      --D-----R--           0           0                  NA ..
premigrated  A------PR-Z      104600      104600 04\27\2001 17:27:12 610215f.bmp
premigrated  A------PR-Z       93656       93656 04\27\2001 17:27:12 banner4.bmp
truncated    A-----OPR-Z           0      204800 04\27\2001 17:27:12 StartCD.exe
truncated    A-----OPR-Z           0        3747 04\27\2001 17:27:12 StartCD.ini

                4 File(s)          198256 bytes
                0 Dir(s)           208547 bytes required to recall offline files
```

This display lists the four files in the local directory and reveals that two are premigrated and two are migrated. The Migrate Time indicates that the files belong to the same BAG (migration group), meaning that they were copied to remote storage at the same time and are stored in the same data set on the remote media.

The two summary lines at the bottom of the display show the size that the file content consumes on the local disk and on remote media. Note that the local size is the sum of the physical size of all local files, including premigrated files. The "required to recall" size is the size of truncated files only.

This display does not reveal whether the premigrated files have ever been truncated, because the premigrated status is also used for files that are truncated and then recalled. That information can be inferred from data in the expanded (**/f**) display.

Finally, the local computer is in Eastern Standard Time, so the file content was actually moved at 12:27:12 local time.

**Example 2: Display the detailed status of a file**

The following command uses the **/f** parameter to request an expanded (full) display of the Remote Storage status of an executable file, Scan.exe. The command uses the redirection character (**>**) to direct the output to a text file, rss-scan.txt, for easier viewing. Also, because the path includes spaces, it is enclosed in quotation marks.

```
RSDir /f "c:\program files\scanco\scan.exe" > rss-scan.txt
```

In response, RSDir writes the following data to the rss-scan.txt file. (The display was edited in Notepad and arranged vertically for readability.)

For descriptions of the fields in the extended display, see Rsdir Remarks.

```
Status:             truncated
Attributes:         A-----OPR-Z
Physical:           0
Logical             2358435
Create Time         04\27\2001 17:08:19
Last Access Time    04\27\2001 17:08:20
Last Modify Time    09\30\1999 10:38:32
Migrate Time        04\27\2001 17:27:23
RC                  0
Last Recall Time    01\01\1601 00:00:00
HSM ID              {4d6e0baf-6d9c-4048-af3f-ca10d94cc8f9}
Bag ID              {21821bd8-dc91-4d13-ba8c-22f9e44b9e2f}
```

```
File ID             {00000000-0000-0000-0000-000000000000}
File Version ID     125831615120000000
Data Stream         0
CRC Type            1
Data Stream CRC     4003735577
DS Size             2358435
DS Start            0
Data Size           2358435
Data Start          322
File Size           2359296
File Start          327766016
Ver Type            1
Ver Info            59026
File Name           C:\Program Files\Scanco\scan.exe

        1 File(s)               0 bytes
        0 Dir(s)          2358435 bytes required to recall offline files
```

**Example 3: Display the status of files in a directory**
The following command displays the Remote Storage status of the Win2K subdirectory:

```
RSDir "c:\inside win2k\symbols\win2k"
```

In response, RSDir displays the Remote Storage status of the subdirectory, not of the files it contains. Because Remote Storage migrates files, not directories, this information is not particularly enlightening.

```
Status         Attributes    Physical     Logical        Migrate Time File Name
------         ----------    --------     -------        ------------ ---------
non-HSM       --D-----R--           0           0                  NA Win2k

                0 File(s)               0 bytes
                1 Dir(s)                0 bytes required to recall offline files
```

The following command adds wildcard characters to request a display of the files in the specified directory.

```
RSDir "c:\inside win2k\symbols\win2k\*.*"
```

In response, RSDir displays the Remote Storage status of files in the Win2k subdirectory, revealing that all files are truncated, even though the subdirectory status is "non-HSM."

```
Showing information for <*.*>

Status         Attributes    Physical     Logical        Migrate Time File Name
------         ----------    --------     -------        ------------ ---------
non-HSM       --D-----R--           0           0                  NA .
non-HSM       --D-----R--           0           0                  NA ..
truncated     A-----OPR-Z           0       99120 04\27\2001 17:27:12 ntkrnlmp.dbg
truncated     A-----OPR-Z           0      738304 04\27\2001 17:27:12 ntkrnlmp.pdb
truncated     A-----OPR-Z           0       98720 04\27\2001 17:27:12 ntkrnlpa.dbg
truncated     A-----OPR-Z           0      738304 04\27\2001 17:27:12 ntkrnlpa.pdb
truncated     A-----OPR-Z           0       99312 04\27\2001 17:27:12 ntkrpamp.dbg
truncated     A-----OPR-Z           0      746496 04\27\2001 17:27:12 ntkrpamp.pdb
truncated     A-----OPR-Z           0     1805372 04\27\2001 17:27:12 ntoskrnl.dbg
truncated     A-----OPR-Z           0      738304 04\27\2001 17:27:12 ntoskrnl.pdb

                8 File(s)               0 bytes
                0 Dir(s)          5063932 bytes required to recall offline files
```

Note that the path is not displayed in the output because RSDir considers the search field to be "*.*". To display the path, use the **/f** parameter, or run RSDir from the local subdirectory, as shown in Example 1: Display the status of files in the local directory.

**Example 4: Display the status of all files on a volume**
The following command lists the Remote Storage status of all files on the E drive of the local computer, and stores the list in a file, RSS-E.txt. The command is run from the root of the E drive and uses the **/s** parameter to request a recursive display. It also uses the redirection character (>) to direct the output to the RSS-E.txt file.

```
E:\>RSDir /s > RSS-E.txt
```

In response, RSDir writes the Remote Storage status of every file on the E drive to the RSS-E.txt file.

The command is run locally, that is, from the root of the E drive, so that correct file paths appear in the basic display output. When you run the command from another directory and specify the files on the E drive, such as `RSDir e:\*.*`, the phrase `*.*` is displayed in the output, instead of path statements. (This is not necessary if you are requesting an extended (full) display.)

Also, RSDir does not display the Remote Storage status of drives. The command `RSDir e:` returns blank or inaccurate output.

**Example 5: List all truncated files on a volume**
The following command creates a list of all truncated files on the C drive of the local computer and stores the list in a text file.

The command uses the **/f** parameter to specify an extended display, which includes the file path, and the **/s** parameter to make the command recursive. Then, the command uses the pipe symbol (`|`) to send the output directly to Find, a tool included in Windows Server 2003. For information about using Find, see "Find" in Help for Windows Server 2003. The Find command finds and displays only lines containing the word "truncated." The redirection character (`>`) directs the output of the Find command to the Trunc-c.txt file.

```
RSDir c:\*.* /f /s | find "truncated" > trunc-c.txt
```

You can use this command format to create logs of the truncated or premigrated files on your volumes. You can also save the command in a batch file and use Scheduled Tasks in Control Panel or use Schtasks to run the command periodically. For information about using Schtasks, see "Schtasks" in Help for Windows Server 2003.

**See Also**
**Concepts**
Rsdir Overview
Rsdir Remarks
Rsdir Syntax
Topchk.cmd
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Rsdir Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RSDir Syntax**

**RSDir** [*Path*[*****_File_]] [**/s**] [**/f**]

Without parameters, RSDir displays the Remote Storage status of the files in the local directory. The status tells whether the object resides in local or remote storage and provides details of its movement and location.

**Parameters**

***Path* [*File*]**

Displays the Remote Storage status of a specified file or the files on a volume or in a directory. If you omit this parameter, RSDir displays the Remote Storage status of files in the local directory.

To display the status of the files on a specified volume or in a specified directory, use wildcard characters, for example, d:*.* or c:\folder*.log.

**/s**

Recursively displays the Remote Storage status of all files and folders within the specified directory.

**/f**

Displays extended ("full") Remote Storage status information.

**Note**

- To accommodate the extended display, increase the width of the command window buffer to at least 500 characters, or redirect the output to a file.

**See Also**

**Concepts**

Rsdir Overview
Rsdir Remarks
Rsdir Examples
Alphabetical List of Tools
Topchk.cmd
Rsdiag Overview
lologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Topchk.cmd

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Topchk.cmd: DFS and SYSVOL Replication Topology Analysis Tool**
This command-line tool processes the output generated when the NtfrsUtl tool is run with the **ds** parameter. Using the NtfrsUtl **ds** output as its data source, TopChk creates a report that summarizes the replication topology of one or more replica sets. This report includes the following information:

- Topology
- Server Reference Table
- Server Host Table
- Server Inbound/Outbound partners
- Summary

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
The File Replication Service (FRS), first introduced in Windows 2000, is an essential element of the Microsoft Active Directory technologies. The Active Directory service uses FRS for SYSVOL replication. FRS can also be used with the Microsoft Distributed File System (DFS) technologies for automatic replication of content between alternating targets. For more information on FRS and Active Directory replication, see the Microsoft Windows Server 2003 Resource Kit.

Running NtrfsUtl **ds** produces output listing the raw definition of the FRS replication topology stored in Active Directory. A replica set consists of two or more computers configured to replicate the contents of a file folder. The individual computers are referred to as replica members.

**System Requirements**
The following are the system requirements for TopChk:

- Windows 2000 Server or Windows Server 2003
- PERL 5.005 521e or 522e

**Files Required**

- Topchk.cmd
- Ntfrsutl.exe

**See Also**
**Concepts**
Topchk Remarks
Topchk Syntax
Topchk Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview

Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Topchk Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TopChk Remarks**
**Before Using TopChk**
If Perl is not installed on your system, if you are connected to the Internet, you can download a free version from the ActiveState Web site.

**TopChk Report**
The TopChk report is used to verify that configuration information for FRS replica sets has been defined and stored correctly within the Active Directory. It is useful both as a regular health-check tool to confirm that replication is correctly configured, and as a troubleshooting tool for investigating possible causes of FRS replication issues.

Familiarity with FRS concepts and terminology is required in order to interpret TopChk reports.

The report consists of two main sections: The Subscriber List Summary followed by a Topology Analysis for each replica set. There is a section that summarizes the report following these sections.

**Subscriber List Summary**
This part of the report displays one row of information for each replica set that this member participates in. For example:

```
SUBSCRIBER: DOC|SALES      Root: f:\docs\sales                    Stage: e:\frs-stg
SUBSCRIBER: DOC|SUPPORT    Root: f:\docs\support                  Stage: e:\frs-stg
```

Each row shows the replica set name, the location on this computer where the replica set data exists, and the location of the FRS staging directory for this replica set.

**Replica Set Topology Analysis**
This part of the report contains a separate topology section for each replica set that this member participates in. The beginning of each section is indicated by the heading:

SET: <*replica set name*>

There may follow any of several possible sections, depending on the topology of the given replica set. The following sections can be included in the TopChk report:

SERVERS MISSING INBOUND CONNECTIONS

SERVERS MISSING OUTBOUND CONNECTIONS

MISSING NTDS SETTINGS REFERENCES

MEMBERS MISSING COMPUTER REFERENCE

MEMBERS MISSING CONNECTION OBJECTS

MEMBERS WITH SELF-REFERENCE CONNECTION OBJECTS

MEMBERS WITH 168 HOUR CONNECTIONS

**SERVERS MISSING INBOUND CONNECTIONS**
This part of the topology report appears if any FRS member servers have outbound replication partners but no inbound connection objects. There may be several reasons for this condition:

- There are no connection objects under the NTDS Settings object for this server. This is an error.
- The ServerReference Attribute for this server is null. This is an error.
- The FRS member object may be missing. This is an error.
- The server may be in a different domain, therefore no FRS member object can be identified for it. Domain controllers that hold a copy of the Global Catalog and are in a different domain are also listed this way. Such servers very likely do have inbound connections to replicate parts of the local domain naming context, but they are not members of the SYSVOL

replica set. As a result, these servers are listed but they are not in an error state.

**SERVERS MISSING OUTBOUND CONNECTIONS**
This part of the topology report appears if any FRS member servers have inbound replication partners but no outbound connection objects.

**MISSING NTDS SETTINGS REFERENCES**
This part of the topology report appears if any FRS member objects have no Server Reference to an NTDS Settings object. For more information, refer to the Knowledge Base article Q312862: "Recovering Missing FRS Objects and FRS Attributes in AD."

**MEMBERS MISSING COMPUTER REFERENCE**
This part of the topology report appears if any FRS member objects have no computer reference. For more information, refer to the Knowledge Base article Q312862: "Recovering Missing FRS Objects and FRS Attributes in AD."

**MEMBERS MISSING CONNECTION OBJECTS**
This part of the topology report appears if any FRS member objects have no inbound connection objects. This is most commonly caused by an Administrator manually defining a replication topology, and not creating a connection object.

In this situation, an Administrator should check for NTDS connection objects. If none exists, the Administrator can create one by using Active Directory Sites & Services. For more information, refer to Knowledge Base article Q257338, "Troubleshooting Missing SYSVOL and NETLOGON Shares."

**MEMBERS WITH SELF-REFERENCE CONNECTION OBJECTS**
This part of the topology report appears if any FRS member objects have connection objects that refer back to themselves. This is most commonly caused by an Administrator manually defining a replication topology and mistakenly creating this condition. In this situation, the topology must be manually corrected.

**MEMBERS WITH 168 HOUR CONNECTIONS**
This part of the topology report appears if the FRS member object has a connection object with a 168-hour (that is, 24x7) replication schedule. This section intentionally lists connections that are continuously enabled; it does not indicate an error.

**Summary Report Section**
This section provides a statistical overview of the other reports. Each potential entry for this section is described here.

Member objects with no NTDS Settings reference

This entry represents the number of FRS objects that did not replicate their SYSVOL share and need to be further investigated.

Member objects with no connection objects

This entry represents the number of NTDS Settings objects with no connection objects, meaning that the servers shown here are not replicating inbound. Note that each SYSVOL member must have one inbound and one outbound connection object. DFS replica sets are permitted to have inbound connection objects with no outbound connection object when using *hub and spoke* or custom topologies.

Connection objects set to disabled

This entry represents the number of disabled replication connections. The Administrator should investigate these to see which connections have been disabled, and determine if this is intended. The section titled **Connection Summary** printed as part of the Summary Report can be used for this purpose.

Monday Schedule

This section appears if TopChk finds a connection object scheduled to replicate at least daily (Monday in this case is assumed to be representative of the schedule for the business week). This provides some insight into whether the schedules have been staggered across connections, which is a best practice for large configurations.

Number connections with per-week active replication hours

This section displays replication hours for each connection object. It begins with a number of rows showing (a) number of servers, then (b) the number of hours the schedule is enabled for those servers.

### Note

- This is approximate. Any hour in which the schedule byte is non-zero is counted as an active replication hour.

It is possible that objects used for intra-site communication may show high replication hours (in the range of 160 or over); however, in most cases this level of network use is not advised.

Connection Summary

This section lists each connection, along with its Monday schedule, replication hours, and enabled/disabled status in order to facilitate further investigation of schedules.

**Note**

- A connection with a GUID, for example "cxtion: CC9BA0F4-3B12-4ADE-86BA-C405D86D4D59" indicates a connection that has been auto-generated by the Active Directory Knowledge Consistency Checker (KCC) or that has been created by using the DFS MMC snap-in utility. When the connection is associated with a name such as "cxtion: DC22," this indicates the connection was manually generated.

**Server Inbound/Outbound Partner Report**

The Server Inbound/Outbound Partner report (the final part of the Summary) lists the total number of inbound and outbound connections for a given server. Each server shows two entries: one with '<<' signifying an inbound connection list, and one with '>>' signifying an outbound connection list.

This part of the report is beneficial for investigating two aspects of the topology definition:

- It shows whether SYSVOL members have the required minimum of one inbound and one outbound connection object.

- It helps determine if the topology is reasonable with respect to the number of direct partners for each member, the amount of content being replicated (size * change), the link speed, and so forth.

An excessive number of connections to or from any one partner should be evaluated in order to see if the topology is balancing replication load appropriately.

**See Also**

**Concepts**

Topchk.cmd
Topchk Syntax
Topchk Examples
Alphabetical List of Tools
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Topchk Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TopChk Examples**

**Example 1: Generating the NtfrsUtl DS Data File**

TopChk organizes the output of NtfrsUtl ds. Therefore, when running the **topchk** command you must specify a file that contains this output. Type the following at the command line to create the NtfrsUtl output file.

```
ntfrsutl ds >ntfrsout.txt
```

This command statement generates an output file called Ntfrsout.txt in the current directory.

**Example 2: Using Topchk to Parse a Data File**

This example parses the data file generated in Example 1: Generating the NtfrsUtl DS Data File and then displays the results to the screen. Type the following at the command line to save the TopChk report to a file in the current directory:

```
topchk.cmd ntfrsout.txt >topchkout.txt
```

This command runs the TopChk script against the Ntfrsout.txt file, which was generated by the NtfrsUtl ds command and saves the resulting output in a file named Topchkout.txt in the current directory.

**See Also**

**Concepts**

Topchk.cmd
Topchk Remarks
Topchk Syntax
Alphabetical List of Tools
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Topchk Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TopChk Syntax**

TopChk uses the following syntax:

**topchk** *DataFile*

**Parameters**

***DataFile***

Specifies a file that was created by running NtfrsUtl with the **ds** parameter. This is a required parameter.

Note

- The Perl interpreter should be installed and in the path for this command to function. See Topchk Remarks for details on obtaining Perl.

**Sample Usage**

```
topchk ntfrs_ds.txt > top.txt
```

**See Also**

**Concepts**

Topchk.cmd
Topchk Remarks
Topchk Examples
Alphabetical List of Tools
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# File and Folder Management Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: File and Folder Management Tools**
Tools in this category are used to manage the following Windows features:

- File compression
- File server
- Indexing service
- Offline files
- Shared folders
- Shared folder shadow copy
- Synchronization Manager

**Notes**

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Bindiff Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Bindiff.exe: Binary File Difference Finder**
This command-line tool shows the differences between two binaries. BinDiff performs the following types of compares:

- A compare between two binary files that excludes time stamp differences. This is the default comparison.
- A compare between two binary files that excludes both time stamp and version information.
- A straight compare between two binary files that does not exclude time stamp or version information. This can be used to verify a tree copy or a CD burn.
- A compare between the files in one directory and the files in a separate directory.

### Note

- If you want to compare two ASCII text files, see the Windiff Overview Support Tool.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Application code is compiled into files that use binary format, and large applications are made up of a number of binary files. Over the programming life of that application, developers, testers, or writers might want to know what files, if any, have changed between the current and previous builds. While sometimes looking at the size of the file alerts the interested party to a change, it is also possible that an executable file could have had changes made to it, without the overall size of the file reflecting such changes. Being able to find files that have changed helps testers and writers direct their attention efficiently to the affected application areas. Developers can ensure that the proper versions of updated files are included in the installation directory for the product.

**System Requirements**
The following are the system requirements for this tool:

- Windows XP Professional or Windows Server 2003

### Note

- BinDiff can compare either compare 32-bit executables or 64-bit executables.

BinDiff also supports:

- Windows 95, Windows 98 and Windows ME
- Windows NT 4.0

**File Required**

- Bindiff.exe

**See Also**
**Concepts**
Bindiff Remarks
Bindiff Syntax

Bindiff Examples
Alphabetical List of Tools
Iasparse Overview
Bindiff Overview

# Bindiff Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BinDiff Remarks**
**BinDiff Limitations**
BinDiff excludes known time stamp fields and version information that are documented in the Windows module format (known as the PE specification). It also locates date and time string constants (such as __DATE__ and __TIME__) in the source code, and excludes them. This catches the majority of differing regions between two executables, but it is possible to generate files that contain differing regions not excluded by BinDiff. In these cases, BinDiff will report the files as "Different" when the user might have expected the files to be "Near Identical." The following list gives some examples that illustrate this behavior:

- Debug information comes in many formats, each of which has various time stamps scattered among the debug data. BinDiff does not try to parse the debug information. Instead, it excludes all known debug regions in the files. Thus, if debug data is stored outside of the specified debug regions, then BinDiff does not know to exclude it.

- There are many post-processing tools that change modules after they are built. These tools, while optimizing page layout, load order, and resources, might strip out information, such as debug information or user names. Localization tools, for example, also change resource strings. Many of these post-processing tools leave data behind, add data, or neglect to clean up dead space and padding areas. This can lead to unexpected differing regions in the two files, which then causes BinDiff to mark the files as "Different."

- Some modules embed other modules in their resource data. If an embedded module is rebuilt, the time stamps in that module change. Running BinDiff on the embedded module itself should result in a "Near Identical" match. When that module is embedded in another module, however, it becomes part of the resource data of the main module. When BinDiff compares two builds of the main module, it will result in "Different" because the differing regions are now a part of the resource data .

**Understanding FileSpec1 and FileSpec2**
*FileSpec1* and *FileSpec2* can represent either files or directories. If two files are specified, then the files will be compared with each other. If two directories are specified, then all the files within the first directory will be compared with files with the same name in the second directory.

**Regions Excluded By BinDiff**
The following regions are excluded by default if found in a module :

- The IMAGE_FILE_HEADER.TimeDateStamp field

- The IMAGE_OPTIONAL_HEADER.CheckSum field

- The DOS/Win16 stub code

- The IMAGE_EXPORT_DIRECTORY.TimeDateStamp field

- The IMAGE_IMPORT_DESCRIPTOR.TimeDateStamp field

- The entire IMAGE_DEBUG_DIRECTORY array

- All debug sections

- The IMAGE_LOAD_CONFIG_DIRECTORY.TimeDateStamp field

- The IMAGE_BOUND_IMPORT_DESCRIPTOR.TimeDateStamp field

- All IMAGE_BOUND_FORWARDER_REF.TimeDateStamp fields

- The IMAGE_RESOURCE_DIRECTORY.TimeDateStamp field
- All padding between resources
- All ANSI __TIMESTAMP__ string constants used in code
- All ANSI __DATE__ string constants used in code
- All ANSI __TIME__ string constants used in code
- All UNICODE __TIMESTAMP__ string constants used in code
- All UNICODE __DATE__ string constants used in code
- All UNICODE __TIME__ string constants used in code

If the **/v** parameter is specified, the following regions are excluded in addition to the default regions:

- The IMAGE_OPTIONAL_HEADER.MajorLinkerVersion field
- The IMAGE_OPTIONAL_HEADER.MinorLinkerVersion field
- The IMAGE_OPTIONAL_HEADER.MajorOperatingSystemVersion field
- The IMAGE_OPTIONAL_HEADER.MinorOperatingSystemVersion field
- The IMAGE_OPTIONAL_HEADER.MajorImageVersion field
- The IMAGE_OPTIONAL_HEADER.MinorImageVersion field
- The IMAGE_OPTIONAL_HEADER.MajorSubsystemVersion field
- The IMAGE_OPTIONAL_HEADER.MinorSubsystemVersion field
- The IMAGE_OPTIONAL_HEADER.Win32VersionValue field
- All VERSION_INFO resources

**Comparing Executable Files**
By default, BinDiff uses a special compare routine for all 32-bit and 64-bit Windows executable files that masks various build time stamp fields in both files when performing the compare. This allows two executable files to be marked as "Near Identical" when the files are truly identical, except the time they were built. You can override this behavior with the **/b** parameter, which forces BinDiff to perform a straight binary compare of all files.

**See Also**
**Concepts**
Bindiff Overview
Bindiff Syntax
Bindiff Examples
Alphabetical List of Tools
Iasparse Overview

# Bindiff Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BinDiff Examples**

**Example 1: Compare Two Versions of a File**

You have two versions of the file dcpromo.dll, and you want to see if there are any differences between the two files. Using BinDiff with no parameters other than *FileSpec1* and *FileSpec2* performs a basic compare. To check for different versions of dcpromo.dll, type the following at the command line:

**bindiff \\heapsort\htmdrop\binold\dcpromo.dll \\heapsort\htmdrop\binnew\dcpromo.dll**

Press Enter . The following output is displayed:

```
Different, Right is newer
```

In this case, the two versions are different: the second file (the one specified on the right) (\\heapsort\htmdrop\binnew\dcpromo.dll) is newer than the first file.

**Example 2: Compare Two Versions of a File But Eliminate Version and Time Stamp Differences**

You have two versions of a file that you want to compare, and you want to make sure that versioning and time stamp information are not the only differences between the two files. To compare the two files while ignoring versioning and time stamp information, type the following at the command line:

**bindiff /v \\heapsort\htmdrop\binold\dcpromo.dll \\heapsort\htmdrop\binnew\dcpromo.dll**

Press Enter. The following output is displayed:

```
Different, Right is newer
```

In this case, the two versions are different and the second file (the file on the right) (\\heapsort\htmdrop\binnew\dcpromo.dll) is newer than the first file (the file on the left). Versioning and time stamp information are not the only differences between the two files. If this were the case, the files would be listed as `Near Identical`.

**Example 3: Compare Differing Byte Counts for Two Versions of a File**

You have two versions of a file that you know differ, but you want to see how different they really are. To compare the two files and get a differing byte count, type the following at the command line:

**bindiff /c \\heapsort\htmdrop\binold\dcpromo.dll \\heapsort\htmdrop\binnew\dcpromo.dll**

Press Enter. The following output is displayed:

```
Different, Right is newer  16,294 bytes differ

File Count Summary
   Identical:      0 files
   Near Identical: 0 files
   Different:      1 files
   Left Only:      0 files
   Right Only:     0 files
   Errors:         0 files
   Total:          1 files

Byte Count Summary
   Matched:    16,294 bytes differ
   Left Only:       0 bytes
   Right Only:      0 bytes
   Total:      16,294 bytes
```

Both versions of the file are 29 KB, and 16 KB of code differs between the two files . The 16 KB difference results from deleting code in the original file, then adding code back in.

**Example 4: Compare Files in Two Directories**

You want to compare the files in one build folder to those in another folder. To compare the files in the two directories, discounting any differences due to versioning and time stamp information, type the following at the command line:

**bindiff /s /v \\heapsort\htmdrop\binold \\heapsort\htmdrop\binnew**

Press Enter. The following output is displayed:

```
appsec.dll   Right-only File
appsec.exe   Right-only File
APPSEC.HLP   Right-only File
dcpromo.csv  Identical
dcpromo.dll  Different, Right is newer
dcpromo.exe  Different, Right is newer
```

In this case, three files have been added to the new build directory (in ...\BinNew). One file, (dcpromo.csv), is unchanged between builds, and two files, (dcpromo.dll and dcpromo.exe), have had changes in addition to versioning and timestamp changes made.

**See Also**
**Concepts**
Bindiff Overview
Bindiff Remarks
Bindiff Syntax
Alphabetical List of Tools
Iasparse Overview

# Bindiff Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BinDiff Syntax**

BinDiff uses the following syntax:

**bindiff** [**/b**] [**/c**] [**/d**[*Bytes*]] [**/s**] [**/v**] [*FileSpec1 FileSpec2*]

**Parameters**

**/b**

Performs a straight binary compare. This parameter will not mask any time stamp or version information.

**/c**

Displays differing byte counts and summary information for the files specified in *FileSpec1* and *FileSpec2*.

**/d[ *Bytes*]**

Displays the byte differences between the files specified in *FileSpec1* and *FileSpec2*. Optionally specifies the number of non-differing bytes to show on each side of the differing regions in the files. The default value of *Bytes* is 8 bytes. Set *Bytes* to 0 to show only the differing bytes, or to -1 to dump both entire files.

**/s**

Compares all files in the subdirectory specified in *FileSpec1* to the files in the subdirectory specified in *FileSpec2*.

**/v**

Excludes all version information in the executable files when performing the compare. This allows two executable files to be marked as "Near Identical" when the files are truly identical, differing only in build time or version number information. This parameter is ignored if used in conjunction with **/b.**

***FileSpec1***

Specifies the directory and name of a file, the name of a file, or a subdirectory to be compared to the specification listed in *FileSpec2*. The name can use the Universal Naming Convention (UNC) format ("\\myserver") or reference a local drive. If the file is located in the current directory *FileSpec1* can also be the name of the file.

**Note**

- The type of information listed in *FileSpec1* must be the same type of information listed in *FileSpec2*. For example, if *FileSpec1* is set to c:\myfile.dll, then *FileSpec2* must also contain the name of a file, and not the name of a subdirectory (c:\new\myfile.dll, not c:\new\).

***FileSpec2***

Specifies the directory and name of a file, the name of a file, or a subdirectory to be compared to the specification listed in *FileSpec1*. The name can use the Universal Naming Convention (UNC) format ("\\myserver") or reference a local drive. If the file is located in the current directory *FileSpec2* can also be the name of the file.

**Note**

- The type of information listed in *FileSpec2* must be the same type of information listed in *FileSpec1*. For example, if *FileSpec2* is set to c:\myfile.dll, then *FileSpec1* must also contain the name of a file, and not the name of a subdirectory (c:\new\myfile.dll, not c:\new\).

**/?**

Displays command-line usage.

**Examples**

See Bindiff Examples, Bindiff Examples, Bindiff Examples, or Bindiff Examples.

**See Also**

**Concepts**

Bindiff Overview

Bindiff Remarks

Bindiff Examples
Alphabetical List of Tools
Iasparse Overview

# Bitsadmin Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Bitsadmin.exe: BITS Administration Utility**
This command-line tool allows administrators to manage the Background Intelligent Transfer Service (BITS), a background file transfer service and queue manager in Windows Server 2003 and Windows XP. Requests to BITS are submitted by an application and the files can be transferred in a throttled manner such that the interactive user is not affected by the bandwidth consumed. Requests are placed in a queue until the files are transferred, at which time the requesting application is invoked and notified of the completion.

**Corresponding UI**
The corresponding user interface for Background Intelligent Transfer Service can be found in the Windows Server 2003 Services snap-in, Services.msc. To open Services, click **Start**, click **Control Panel**, click **Performance and Maintenance**, click **Administrative Tools**, double-click **Services**, and then double-click **Background Intelligent Transfer Services**.

**Concepts**
Background Intelligent Transfer Service (BITS) is a file-transfer technology similar in concept to other standard file-transfer protocols, such as FTP. BITS is used by various applications to manage file transfer using standard Hypertext Transfer Protocol (HTTP).

BITS requires Internet Information Services (IIS) with BITS extensions to be running on the source server if you want to use BITS to upload files to your server.

### Note

- Internet Information Services is not required for downloads, only for uploads.

BITS is used by programs (such as Windows AutoUpdate) to download files by using spare bandwidth. BITS can maintain file transfers through network disconnections and computer restarts.

**System Requirements**
The following are the system requirements for this tool:

- Windows Server 2003, Windows XP Professional, or Windows XP Home Edition
- Internet Information Services (with the BITS Server Extension) is required to receive BITS uploads.
- Administrator rights on the target computer (the computer receiving transfers) is required to modify the jobs of other users.

### Note

- BITS requires a Web server running on the source server for downloading files. For file uploads, BITS requires IIS and a corresponding server component to be running on the destination server.

**File Required**

- Bitsadmin.exe

**See Also**
**Concepts**
Bitsadmin Remarks
Bitsadmin Syntax
Bitsadmin Examples
Alphabetical List of Tools
Topchk.cmd

Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Bitsadmin Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BitsAdmin Remarks**

**Windows Updates**

BITS is often erroneously confused with Windows Update service because Windows Update service uses Background Intelligent Transfer Service by default.

**Network Throttling**

BITS regulates the transfer rate to minimize impact on user interactivity, such as a job sent to a network printer or web pages viewed through Internet Explorer.

**Customized APIs**

BITS is exposed to programmers through a set of Application Program Interfaces (APIs). For more information on using BITS programmatic interfaces, see the Microsoft Developer Network.

**See Also**

**Concepts**

Bitsadmin Overview
Bitsadmin Syntax
Bitsadmin Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview

# Bitsadmin Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BitsAdmin Examples**

**Example 1: Create BITS Job**

You want to create a job. Type the following at the command line:

**bitsadmin /create myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

Created job {2DC7527D-C7A8-444C-84A9-D772E79D4B37}.
```

**Example 2: Get a Job's Type**

You want to get the type of job for **myjob**. Type the following at the command line:

**bitsadmin /gettype myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

DOWNLOAD
```

**Example 3: Get a Job's Size**

You want to know the size of a job. Type the following at the command line:

**bitsadmin /getbytestotal myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

236
```

**Example 4: Get the Number of Bytes Transferred for a Job**

You want to know how much of the job has already been transferred. Type the following at the command line:

**bitsadmin /getbytestransferred myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

106
```

**Example 5: Get the Number of Files in a Job**

You want to know how many files are in a job. Type the following at the command line:

**bitsadmin /getfilestotal myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.
```

```
3
```

**Example 6: Get the Number of Files Transferred for a Job**
You want to know how many files have been transferred from a job. Type the following at the command line:

**bitsadmin /getfilestransferred myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

1
```

**Example 7: Get a Job's Creation Time**
You want to know when a job was created. Type the following at the command line:

**bitsadmin /getcreationtime myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

6/3/2002 10:49:02 AM
```

**Example 8: Get a Job's Last Modification Time**
You want to know when a job was last modified. Type the following at the command line:

**bitsadmin /getmodificationtime myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

6/3/2002 10:49:02 AM
```

**Example 9: Get a Job's Completion Time**
You want to know if and when a job has completed. Type the following at the command line:

**bitsadmin /getcompletiontime myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

WORKING
```

The **myjob** job has not yet completed.

**Example 10: Get a Job's State**
You want to know a job's state. Type the following at the command line:

**bitsadmin /getstate myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

SUSPENDED
```

**Example 11: Get a Job's Owner**
You want to know who owns a job. Type the following at the command line:

**bitsadmin /getowner myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

DOMAINMAIN\smithj
```

**Example 12: Get a Job's Display Name**
You want to know the display name for a job. Type the following at the command line:

**bitsadmin /getdisplayname myjob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

myjob
```

### Note

- Using DisplayName as the argument for the job will work only if there is one job in the queue with this display name. Otherwise, Bitsadmin will return an error in which case, the GUID must be used.

**Example 13: Add a file to a specified job**
You want to add a file to a specified job. Type the following at the command line:

**bitsadmin /addfile myJob http://myserver/myfile.ext c:\myFile.ext**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3652.0 ]

BITS administration utility.

(C) Copyright 2000-2002 Microsoft Corp.

Added http://myserver/myfile.ext -> c:\myFile.ext to job.
```

### Note

- Instead of using myJob as the parameter the previous example, you could enter the GUID of the job. You can get the GUID of the job when you do the bitsadmin /create. You can also do a bitsadmin /list to get the GUID of the job. The GUIDS are unique, but the display names are not.

**Example 14: Resume a job**
You want to resume a job. Type the following at the command line:

**bitsadmin /resume myJob**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3652.0 ]

BITS administration utility.

(C) Copyright 2000-2002 Microsoft Corp.

Job resumed
```

**Example 15: Delete all jobs in the manager**
You want to delete all jobs in the manager. Type the following at the command line:

**bitsadmin /reset**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

{A305C29C-2247-42A6-918E-C575280CE885} canceled.
{6F2D447C-712E-4674-9FAE-42DBFAD071A8} canceled.
{0B8068F6-2572-4932-AE12-CD58DA78912E} canceled.
3 out of 3 jobs canceled.
```

**Example 16: List transfer jobs**

You want to list the transfer jobs. Type the following at the command line:

**bitsadmin /list**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

{A305C29C-2247-42A6-918E-C575280CE885} job1 SUSPENDED 0 / 0 0 / 0
{6F2D447C-712E-4674-9FAE-42DBFAD071A8} job2 SUSPENDED 0 / 0 0 / 0
{0B8068F6-2572-4932-AE12-CD58DA78912E} job3 SUSPENDED 0 / 0 0 / 0
Listed 3 job(s).
```

**Example 17: Delete a specified job**

You want to delete a specified job. Type the following at the command line:

**bitsadmin /cancel job1**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

Job canceled
```

**Example 18: Perform a series of operations simultaneously**

You want to run a series of operations such as: /create /addfile /resume /list /complete. Type the following at the command line:

**bitsadmin /create /addfile /resume /list /complete**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

Created job {9CFB8267-72BA-4725-A487-8CB1A9390637}.

bitsadmin /addfile job1 http://bitsnet/downloads/50mb.zip c:\temp\50mb.zip

BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

Added http://bitsnet/downloads/50mb.zip -> c:\temp\50mb.zip to job.

bitsadmin /resume job1

BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.
```

```
Job resumed.

bitsadmin /list

BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

{9CFB8267-72BA-4725-A487-8CB1A9390637} job1 TRANSFERRING 0 / 1 2000 / 50189685
Listed 1 job(s).

bitsadmin /list

BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

{9CFB8267-72BA-4725-A487-8CB1A9390637} job1 TRANSFERRED 1 / 1 50189685 / 501896
Listed 1 job(s).

bitsadmin /complete

BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

Invalid number of arguments.

bitsadmin /complete job1

BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

Job completed.

bitsadmin /list

BITSADMIN version 1.5 [ 5.1.3597.0 ]
BITS administration utility.
(C) Copyright 2000-2002 Microsoft Corp.

Listed 0 job(s).
```

**Example 19: Run /list /allusers /verbose**

You want to run /list /allusers /verbose. Type the following at the command line:

**bitsadmin /list /allusers /verbose**

Press **Enter**. The following output is displayed:

```
BITSADMIN version 1.5 [ 5.2.3639.0 ]

BITS administration utility.

(C) Copyright 2000-2002 Microsoft Corp.


GUID: {134E18A6-FC4B-46AC-A7E9-B7A7D59C84BE} DISPLAY: test

TYPE: DOWNLOAD STATE: SUSPENDED OWNER: DOMAIN1\joe

PRIORITY: NORMAL FILES: 0 / 1 BYTES: 0 / UNKNOWN

CREATION TIME: 9/13/2002 12:47:10 PM MODIFICATION TIME: 9/13/2002 12:47:13 PM

COMPLETION TIME: UNKNOWN

NOTIFY INTERFACE: UNREGISTERED NOTIFICATION FLAGS: 3
```

```
RETRY DELAY: 600 NO PROGRESS TIMEOUT: 1209600 ERROR COUNT: 0

PROXY USAGE: PRECONFIG PROXY LIST: NULL PROXY BYPASS LIST: NULL

DESCRIPTION:

JOB FILES:

        0 / UNKNOWN WORKING http://www.microsoft.com -> c:\ms

NOTIFICATION COMMAND LINE: none


Listed 1 job(s).


BITSADMIN version 1.5 [ 5.2.3639.0 ]

BITS administration utility.

(C) Copyright 2000-2002 Microsoft Corp.


{134E18A6-FC4B-46AC-A7E9-B7A7D59C84BE} canceled.

1 out of 1 jobs canceled.
```

**See Also**
**Concepts**

Bitsadmin Overview
Bitsadmin Remarks
Bitsadmin Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview

# Bitsadmin Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BitsAdmin Syntax**

BitsAdmin uses the following syntax:

**bitsadmin** [**/rawreturn**] [**/wrap**] [**/nowrap** *Job*] [**/help**|**/?**] [**/list** [**/allusers**] [**/verbose**]]
[**/monitor** [**/allusers**] [**/refresh** *Seconds*]] [**/reset**] [**/create** *DisplayName*] [**/info** *Job* [**/verbose**]
[**/addfile** *Job RemoteURL LocalName*] [**/listfiles** *Job*] [**/suspend** *Job*] [**/resume** *Job*] [**/cancel** *Job*] [**/complete** *Job*]
[**/gettype** *Job*] [**/getbytestotal** *Job*] [**/getbytestransferred** *Job*] [**/getfilestotal** *Job*] [**/getfilestransferred** *Job*]
[**/getcreationtime** *Job*] [**/getmodificationtime** *Job*] [**/getcompletiontime** *Job*] [**/getstate** *Job*] [**/geterror** *Job*]
[**/getowner** *Job*] [**/getdisplayname** *Job*] [**/setdisplayname** *Job DisplayName*] [**/getdescription** *Job*]
[**/setdescription** *Job Description*] [**/getpriority** *Job*] [**/setpriority** *Job Priority*] [**/getnotifyflags** *Job*]
[**/setnotifyflags** *Job NotifyFlags*] [**/getnotifyinterface** *Job*] [**/getminretrydelay** *Job*] [**/setminretrydelay** *Job RetryDelay*]
[**/getnoprogresstimeout** *Job*] [**/setnoprogresstimeout** *Job Timeout*] [**/geterrorcount** *Job*] [**/getproxyusage** *Job*]
[**/getproxylist** *Job*] [**/getproxybypasslist** *Job*] [**/setproxysettings** *Job* **use**, *List*, *Bypass*] [**/takeownership** *Job*]
[**/SetNotifycmdline** *Job command-line*] [**/GetNotifycmdline** *Job*]
[**/SetCredentials** *Job Target Scheme Username Password*] [**/RemoveCredentials** *Job Target Scheme*]
[**Getreplyfilename** *Job*] [**Setreplyfilename** *Job Path*] [**/GetReplyProgress** *Job*] [**/GetReplyData** *Job*] [**/GetReplyData** *Job*]

**Parameters**

**/help| /?**
Displays command line usage.

**/list[ /allusers] [ /verbose]**
Lists transfer jobs. The **/allusers** parameter lists jobs for all users on the system. The **/verbose** parameter provides detailed information about jobs.

**/monitor[ /allusers] [ /refresh*Seconds*]**
Monitors the copy manager. The **/allusers** parameter monitors the copy manager for all users on the system. The **/refresh** parameter reacquires copy manager data in the specified time interval in seconds.

**/reset**
Deletes all jobs in the manager. If executed by an administrator, all jobs are deleted. If executed by a normal user, only the jobs owned by the user are deleted.

**/create [type] *DisplayName***
Creates a job and assigns *DisplayName* to it. The **type** can be **/Download**, **/Upload**, or **/Upload-Reply.**

**/info *Job*[ /verbose]**
Displays information about the specified job. The **/verbose** parameter provides detailed information about the job.

**/addfile *Job RemoteURL LocalName***
Adds a file to the specified job. *RemoteURL* specifies the file's source location. *LocalName* specifies the file's name on the local target computer.

**/listfiles *Job***
Lists the files in the specified job.

**/suspend *Job***
Suspends the specified job. The job will not be scheduled to run again until the **/resume** parameter is executed.

**/resume *Job***
Queues the job to the list of jobs enabled for transfer.

**/cancel *Job***
Deletes the specified job.

**/complete *Job***
Completes the job and makes the files available for the destination directory. This is normally executed after the job moves to the transferred state.

**/gettype *Job***
Retrieves the job type of the specified job.

**/getbytestotal *Job***
Retrieves the size of the specified job.

**/getbytestransferred *Job***
Retrieves the number of bytes transferred for the specified job.

**/getfilestotal *Job***
Retrieves the number of files in the specified job.

**/getfilestransferred *Job***
Retrieves the number of files transferred for the specified job.

**/getcreationtime *Job***
Retrieves the job creation time for the specified job.

**/getmodificationtime *Job***
Retrieves the job modification time.

**/getcompletiontime *Job***
Retrieves the job completion time for the specified job.

**/getstate *Job***
Retrieves the job state for the specified job.

**/geterror *Job***
Retrieves detailed error information for the specified job

**/getowner *Job***
Retrieves the job owner for the specified job.

**/getdisplayname *Job***
Retrieves the display name for the specified job.

**/setdisplayname *Job DisplayName***
Sets the *DisplayName* for the specified job.

**/getdescription *Job***
Retrieves the job description of the specified job.

**/setdescription *Job Description***
Sets the description for the specified job.

**/getpriority *Job***
Retrieves the priority of the specified job.

**/setpriority *Job Priority***
Sets the priority for the specified job.

**/getnotifyflags *Job***
Retrieves the notify flags for the specified job.

**/setnotifyflags *Job NotifyFlags***
Sets the notify flags for the specified job.

**/getnotifyinterface *Job***
Determines if notify interface is registered for the specified job.

**/getminretrydelay *Job***
Retrieves the retry delay, in seconds, for the specified job.

**/setminretrydelay *Job RetryDelay***
Sets the *RetryDelay*, in seconds, for the specified job.

**/getnoprogresstimeout *Job***
Retrieves the no progress timeout, in seconds.

**/setnoprogresstimeout *Job Timeout***
Sets the no progress *Timeout*, in seconds, for the specified job.

**/geterrorcount *Job***
Retrieves an error count for the specified job.

**/getproxyusage *Job***
Retrieves the proxy usage setting for the specified job.

**/getproxylist *Job***
Retrieves the proxy list for the specified job.

**/getproxybypasslist *Job***
Retrieves the proxy bypass list for the specified job.

**/setproxysettings *Job* use, *List*, *Bypass***
Sets the proxy Settings for the specified job.

**/takeownership *Job***
Takes ownership of the specified job.

**/SetNotifycmdline *Job command-line***
Sets a command line for job notification.

**/GetNotifycmdline *Job***
Returns the job's notification command line.

**/SetCredentials *Job Target Scheme Username Password***
Adds credentials to a job. The *Target* can be either SERVER or PROXY. The *scheme* can be BASIC, DIGEST, NTLM, NEGOTIATE, or PASSPORT.

**/RemoveCredentials *Job Target Scheme***
Removes credentials from a job. The *Target* can be either SERVER or PROXY. The *scheme* can be BASIC, DIGEST, NTLM, NEGOTIATE, or PASSPORT.

The following options are valid for **Upload-reply** jobs only:

- **Getreplyfilename** *Job*

  **Gets the path of the file containing the server reply.**

- **Setreplyfilename *Job***

  ***Path* Sets the path of the file containing the server reply.**

- **/GetReplyProgress *Job***

  **Returns the size and progress of the server reply.**

- **/GetReplyData *Job***

  **Retrieves the server's reply data in hexadecimal format.**

**The following options can be placed before the command:**

- **/rawreturn**

  **Returns data suitable for parsing. Strips new line characters and formatting from the output. This parameter can be used in conjunction with the /create and /get* parameters.**

- **/wrap**

  **Wraps output to fit in a command window.**

- **/nowrap *Job***

**Does not wrap output to fit in a command window.**

**Note**

- **Commands that take a job parameter will accept either a job name or a job-ID GUID inside braces. BitsAdmin reports an error if the name is ambiguous.**

**See Also**
**Concepts**
**Bitsadmin Overview**
**Bitsadmin Remarks**
**Bitsadmin Examples**
**Alphabetical List of Tools**
**Topchk.cmd**
**Rsdir Overview**
**Rsdiag Overview**
**Iologsum Overview**
**Health_chk Overview**
**Ftonline Overview**
**Filever Overview**
**Efsinfo Overview**
**Dmdiag Overview**
**Dskprobe Overview**
**Diruse Overview**
**Dfsutil Overview**
**Connstat Overview**
**Cabarc Overview**

# Cabarc Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cabarc.exe: Cabinet Tool**
This command-line tool allows users to create, query and extract Windows cabinet (CAB) files.

### Note

- CABArc does not allow you to add files to an existing CAB (.cab) file.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
A CAB file or cabinet file is a single file created to hold a number of compressed files and used for the installation of software products. When a program is installed, the files within the CAB file are decompressed and copied to their appropriate directory locations. Windows CAB files use the extension .cab.

**System Requirements**
The following are the system requirements for this tool:

- Windows Server 2003, or Windows XP Professional

**File Required**

- Cabarc.exe

**See Also**
**Concepts**
Cabarc Remarks
Cabarc Syntax
Cabarc Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Cabarc Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CABArc Remarks**

**Using CABArc to View the Contents of a CAB File**
To view the contents of a CAB file with CABArc:

1. Open a command window and navigate to the folder containing the .cab file that has the contents you want to view.

2. Type the following at the command prompt:

   **cabarc l FileName.cab**

   *FileName* is the name of the .cab file.

**See Also**
**Concepts**
Cabarc Overview
Cabarc Syntax
Cabarc Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Bitsadmin Overview

# Cabarc Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CABArc Examples**

**Example 1: List Files in a Cabinet**

You want to list the files in support.cab. Type the following at the command prompt:

**cabarc l c:\builds\support.cab**

Press **Enter**. You will see the following output:

```
Microsoft (R) Cabinet Tool - Version 5.1.2600.0
Copyright (c) Microsoft Corporation. All rights reserved..

Listing of cabinet file 'c:\builds\support.cab' (size 4091527):
   89 file(s), 4 folder(s), set ID 32099, cabinet #0

File name                       File size     Date      Time   Attrs
-----------------------------  ---------- ---------- -------- -----
   acldiag.exe                      84992 2002/05/24 07:41:56  ----
   adprop.dll                      735232 2002/05/24 07:41:42  ----
   addiag.exe                      216091 2002/05/24 04:28:48  ----
   adsiedit.dll                    330240 2002/05/24 07:41:10  ----
   adsiedit.msc                     42986 2001/10/15 17:40:54  ----
   bitsadmin.exe                    48640 2002/05/24 07:41:16  ----
   bindiff.exe                      17408 2002/05/24 08:28:54  ----
   browstat.exe                     41472 2002/05/24 07:53:38  ----
```

**Example 2: Extract Files from Cabinet**

You want to extract all of the .dll files in the support.cab file. Type the following at the command prompt:

**cabarc x c:\builds\support.cab *.dll**

Press **Enter**. The .dll files will be extracted to the current directory. The following output is displayed:

```
Microsoft (R) Cabinet Tool - Version 5.1.2600.0
Copyright (c) Microsoft Corporation. All rights reserved..

Extracting file(s) from cabinet 'c:\builds\support.cab':
   extracting: adprop.dll
   -- File 'adsiedit.dll' already exists; overwrite? (Yes/No/All/Quit):
```

Type the following to make CABArc overwrite all preexisting files without continuing to prompt:

```
a
```

The following output is displayed:

```
   extracting: adsiedit.dll
   extracting: clonepr.dll
   extracting: depends.dll
   extracting: iadstools.dll
   extracting: movetree.dll
   extracting: rstools.dll
   extracting: sidwkr.dll

Operation successful
```

**Example 3: Create a New Cabinet**

You want to create a new cabinet file. Type the following at the command prompt:

**cabarc n test1.cab *.dll**

Press **Enter**. The following output is displayed:

```
Microsoft (R) Cabinet Tool - Version 5.1.2600.0
Copyright (c) Microsoft Corporation. All rights reserved..

Creating new cabinet 'test1.cab' with compression 'MSZIP':
  -- adding adprop.dll
  -- adding adsiedit.dll
  -- adding apidll.dll
  -- adding clonepr.dll
  -- adding depends.dll
  -- adding gutils.dll
  -- adding iadstools.dll
  -- adding movetree.dll
  -- adding ntfrsapi.dll
  -- adding rstools.dll
  -- adding sidwkr.dll
  -- adding traceprt.dll

Completed successfully
```

The cabinet file test1.cab is created in the current directory with MSZIP compression.

**See Also**

**Concepts**

Cabarc Overview
Cabarc Remarks
Cabarc Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Bitsadmin Overview

# Cabarc Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CABArc Syntax**
CABArc uses the following general syntax:

**cabarc** *parameters operation arguments*

**Operations**
**L**
Lists files contained within a CAB file.

**Syntax**
**cabarc l** *CabFileName* **.cab**

**Parameters**
***CabFileName* .cab**
Specifies the name of the CAB file for which the contents are to be listed. The *CabFileName* may include a path (for example, c:\test\mycab.cab).

**N**
Creates a CAB file.

**Syntax**
**cabarc** [**/m {lzx:*Compression*|mszip|none}**] [**/p**] [**/P** *Pathname*] [**/r**] [**/s**] [**/I** *SetID*] **n** *NewCabFileName*[*****]**.cab** *FileNameOrPattern* [*FileNameOrPattern* [...]]

**Parameters**
**/m {lzx: *Compression* |mszip|none}**
Sets compression type. If no compression type is specified, MSZIP format is used.

| Value | Description |
|---|---|
| LZX:*Compression* | Uses LZX compression. Valid *Compression* values are 15 to 21. |
| MSZIP | Uses MSZip format. |
| NONE | Uses no compression. |

**/p**
Preserves path names. Absolute paths are not allowed.

**/P *Pathname***
Strips the specified *Pathname* from files when they are added to the CAB file.

**Note**

- The **/P** option can be used multiple times to strip out multiple paths.

**/r**
Recurses into subdirectories when adding files.

**/s**
Reserves space in the CAB file for signing. For example, `/s 6144` reserves 6K bytes.

**/i *SetID***
Sets the CAB set ID when creating CAB files. The default set ID is 0.

**Note**

- When using the **/d** parameter, the *NewCabFileName* should contain a single "*" character where the cabinet number will be inserted.

***NewCabFileName*[*] .cab**
Specifies the name of the new cab file. The file must have a .cab file type.

**Note**

- When using the **/d** parameter, the *NewCabFileName* should contain a single "*" character where the cabinet number will be inserted.

***FileNameOrPattern***
Specifies the file or pattern of files to include in the CAB file. For example, `myfile.ini mfiles.*` will include the file myfile.ini and all files with a file name of mfiles regardless of type.

**Note**

- When creating a CAB file, the plus sign (+) may be used as a filename to force a folder boundary. For example, `cabarc n test.cab *.c test.h + *.bmp`.

**X**
Extracts files from a CAB file.

**Syntax**
**cabarc** [**/c**] [**/o**] **x** *CabFileNameFileNameOrPattern* [*FileNameOrPattern* [...]] [*Pathname*]

**Parameters**
**/c**
Confirms which files will be extracted.

**/o**
Specifies that existing files with the same name as those being extracted will be overwritten without asking for confirmation.

***CabFileName* .cab**
Specifies the name of the CAB file that contains the files to be extracted. The *CabFileName* may include a path (for example, c:\test\mycab.cab).

***FileNameOrPattern***
Specifies the file or pattern of files to extract from the CAB file. For example, `myfile.ini mfiles.*` will extract the file myfile.ini and all files with a file name of mfiles regardless of type.

***Pathname***
Specifies that path to which the files should be extracted. The path must end with a back slash example, `c:\myfiles\`. If no path is specified, the current path is used.

**Note**

- When run without parameters, operations, or arguments, CABArc displays command line usage.

**See Also**
**Concepts**
Cabarc Overview
Cabarc Remarks
Cabarc Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview

Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Bitsadmin Overview

# Depends Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Depends.exe: Dependency Walker**
This GUI tool scans any Windows module (including .exe, .dll, .ocx, and .sys, among others) and builds a hierarchical tree diagram of all dependent modules. For each module found, it lists all the functions that are exported by that module, and identifies which of those functions are actually being called by other modules. Another view displays the minimum set of required files, along with detailed information about each file, including a full path to the file, base address, version numbers, computer type, debug information, and more.

**Corresponding UI**
Depends provides its own user interface.

Depends is very useful for troubleshooting system errors related to module load problems. During the module scan, this tool detects dozens of common application problems such as missing modules, invalid modules, import/export mismatches, circular dependency errors, and mismatched computer types of modules.

Depends is useful for developers, testers, system administrators, creators of setup scripts, and anyone else who needs to examine what dependencies are required to make a specific module load successfully.

Depends runs on Windows 95, Windows 98, Windows NT 3.51, Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003. It can process any Windows modules, including those designed for Windows CE. It can run as a graphical application or as a console application. New features in version 2.0 include C++ function name undecoration, detection of delay-load and dynamic dependencies, application profiling, user-definable search paths, and the ability to save results to several file formats.

Details on running Depends can be found in Dependency Walker 2.0 Help (Depends.hlp), a separate documentation file.

**System Requirements**
The following are the system requirements for Depends:

- Windows 95, Windows 98, Windows NT 3.51, Windows NT 4.0, Windows 2000, Windows XP, or Windows Server 2003

**Files Required**

- Depends.cnt
- Depends.exe
- Depends.hlp
- Depends.dll

**See Also**
**Concepts**
Depends Remarks
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Depends Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Depends Remarks**

**Configuring Depends**

When you first run Depends, it registers itself with your system as a handler for all Win32 binary files. This causes a **View Dependencies** item to be added to the system's context menu. You must configure which file extensions Depends handles by opening the **Options** menu and adding file extensions in the **Handled File Extensions** dialog box. Then, when you right-click a Win32 module file in Windows Explorer, **View Dependencies** is displayed as one of the options.

**Opening a File in Depends**

To open a file in Depends, use one of four methods:

1. Right-click a module file in Windows Explorer and choose **View Dependencies** from the context menu. This requires that you configure Depends to handle certain file extensions. To do this, in Depends on the **View** menu, click **Handled File Extensions**.

2. In Depends on the **File** menu, click **Open**.

3. Drag and drop a file on top of the Depends main window.

4. At the command prompt, type **depends** followed by the name of the file to be viewed.

**Getting Help on Depends**

To get help on Depends while running the tool, open the table of contents by using one of two methods:

- In Depends on the **Help** menu, click **Help Topics**.

- To get help on the view that currently has focus, press F1.

**See Also**

**Concepts**

Depends Overview
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Diruse Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Diruse.exe: Directory Disk Usage**
This command-line tool displays directory size information, including compression information for NTFS volumes. You can use DirUse to determine the actual usage of space for compressed files and directories.

You can also specify a maximum folder size. DirUse then marks any folders that exceed the specified limit and, if you choose, alerts you to the problem.

If you run DirUse while logged on as a member of the Administrators group, you can check the use of disk space in directories even if you do not have access rights to those directories.

### Note

- DirUse is similar to du used in UNIX.

**Corresponding UI**
You can use Windows Explorer to determine the size of the first level of directories on a drive, but not to determine the overall size of directories and their subdirectories.

There is no other corresponding user interface for DirUse.

**Concepts**
For more information on NTFS, see NTFS.

**System Requirements**
The following are the system requirements for DirUse:

- Windows XP Professional or Windows Server 2003

**File Required**

- Diruse.exe

**See Also**
**Concepts**
Diruse Remarks
Diruse Syntax
Diruse Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Diruse Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DirUse Remarks**

**Using DirUse to Determine Compressed File Size**
The parameter to use for compressed directories and files is **/c**, which causes DirUse to display the actual size of a compressed file or directory instead of apparent uncompressed size. For example, if your D drive is an NTFS volume, and you want to determine the disk space used (in megabytes) and the number of files in each directory, at the command line, type:

```
diruse /s /m /c d:
```

To see compression information for a file, open Windows Explorer, right-click the file, and then click **Properties**.

**See Also**
**Concepts**
Diruse Overview
Diruse Syntax
Diruse Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Diruse Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DirUse Examples**

**Example 1: List current file usage for a shared volume**

In your company, users store their files on a shared volume called \\Shared. The \\Shared\Data folder is broken down into folders for each department. Each department folder is broken down into folders for users from that department as well as one or more folders that are shared by all users in that department. The \\Shared\Data folder also has one folder — \\Shared\Data\Common — that is used to share information between departments.

You have been asked to generate a report that shows the current usage for the network volume where users in your company store their data. This report must show the usage by department and list any individual user who has more than 50 MB of data on the volume. To generate the data for the first part of your report, type the following at the command line:

```
diruse /* /m \\shared\data > c:\deptdata.txt
```

This generates a list of usage in megabytes of the top-level folders contained in the \\Shared\Data folder. The output is sent to the Deptdata.txt file on your C drive.

Looking at Deptdata.txt, you see the following output:

```
    Size (mb)  Files  Directory
       248.60   3739  SUB-TOTAL: \\SHARED\DATA\Admin
        31.27     36  SUB-TOTAL: \\SHARED\DATA\Common
      1448.14  23654  SUB-TOTAL: \\SHARED\DATA\ConsultServ
       107.64    750  SUB-TOTAL: \\SHARED\DATA\CorpServ
      1053.28   6145  SUB-TOTAL: \\SHARED\DATA\Design
       147.11    258  SUB-TOTAL: \\SHARED\DATA\Engrg
        70.56    745  SUB-TOTAL: \\SHARED\DATA\Finance
        18.78   1188  SUB-TOTAL: \\SHARED\DATA\Hardware
        12.18     50  SUB-TOTAL: \\SHARED\DATA\HR
         2.22     23  SUB-TOTAL: \\SHARED\DATA\InfoTech
         5.44     46  SUB-TOTAL: \\SHARED\DATA\Legal
        18.53     56  SUB-TOTAL: \\SHARED\DATA\ProdDev
       470.43   2016  SUB-TOTAL: \\SHARED\DATA\ProdMktg
       107.92   1445  SUB-TOTAL: \\SHARED\DATA\Research
        31.10     43  SUB-TOTAL: \\SHARED\DATA\Sales
         9.98     14  SUB-TOTAL: \\SHARED\DATA\TechSupp
       114.82   1269  SUB-TOTAL: \\SHARED\DATA\Training
      3898.00  41477  TOTAL
```

You can see in this output that only the Admin, ConsultServ, CorpServ, Design, Engrg, Finance, ProdMktg, Research, and Training departments have more than 50 MB of data each.

To see what users have more than 50 MB of data stored on \\Shared\Data, type the following at the command line:

```
diruse /* /m /q:50 /d \\shared\data\admin, \\shared\data\consultserv, \\shared\data\corpserv,
\\shared\data\design, \\shared\data\engrg, \\shared\data\finance, \\shared\data\prodmktg,
\\shared\data\research, \\shared\data\training > c:\userdata.txt
```

This generates a list of the all top-level folders contained in any of the listed folders that use more than 50 MB of disk space. The output is sent to the Userdata.txt file on your C drive.

Looking at Userdata.txt, you see the following output:

```
Searching for directories that have exceeded 50 megabytes
    Size (mb)  Files  Directory
!       68.30    697  SUB-TOTAL: \\SHARED\DATA\Finance\Jack
        70.55    744  TOTAL
```

You can see in this output that only one user, Jack, is marked as using more than 50 MB of data.

**Example 2: Determine compressed file and folder sizes**

You want to compare the actual size of folders containing compressed files and the apparent uncompressed size in your Windows folders. To generate a list of top-level folders contained in the Windows folder with their associated compressed

sizes, type the following at the command line:

```
diruse /c /* /m c:\Windows > c:\data\compressed.txt
```

The Compressed.txt file in the C:\Data folder contains the output from the command-line statement. The file contains the following output:

```
    Size (mb)  Files  Directory
         0.00      0  SUB-TOTAL: C:\WINDOWS\addins
         0.19     64  SUB-TOTAL: C:\WINDOWS\Application Compatibility Scripts
         1.56      7  SUB-TOTAL: C:\WINDOWS\AppPatch
         2.52     16  SUB-TOTAL: C:\WINDOWS\Cluster
         0.00      0  SUB-TOTAL: C:\WINDOWS\Config
         0.00      0  SUB-TOTAL: C:\WINDOWS\Connection Wizard
         0.29    184  SUB-TOTAL: C:\WINDOWS\Cursors
         0.03      5  SUB-TOTAL: C:\WINDOWS\Debug
         0.00      3  SUB-TOTAL: C:\WINDOWS\Downloaded Program Files
        86.94      1  SUB-TOTAL: C:\WINDOWS\Driver Cache
        12.04    177  SUB-TOTAL: C:\WINDOWS\Fonts
        40.71   1536  SUB-TOTAL: C:\WINDOWS\Help
         0.00      0  SUB-TOTAL: C:\WINDOWS\IIS Temporary Compressed Files
         0.60      3  SUB-TOTAL: C:\WINDOWS\ime
         0.00      0  SUB-TOTAL: C:\WINDOWS\ime (x86)
        41.86   1323  SUB-TOTAL: C:\WINDOWS\inf
         1.71     29  SUB-TOTAL: C:\WINDOWS\Installer
        10.96     35  SUB-TOTAL: C:\WINDOWS\java
         2.43     40  SUB-TOTAL: C:\WINDOWS\Media
        29.36    137  SUB-TOTAL: C:\WINDOWS\Microsoft
         3.13     30  SUB-TOTAL: C:\WINDOWS\msagent
         0.00      0  SUB-TOTAL: C:\WINDOWS\msapps
         0.00      0  SUB-TOTAL: C:\WINDOWS\mui
         0.00      1  SUB-TOTAL: C:\WINDOWS\Offline Web Pages
        11.75    801  SUB-TOTAL: C:\WINDOWS\PCHEALTH
         1.07      5  SUB-TOTAL: C:\WINDOWS\Registration
        11.62     10  SUB-TOTAL: C:\WINDOWS\repair
         0.00      1  SUB-TOTAL: C:\WINDOWS\Resource Pro Kit
         2.01      4  SUB-TOTAL: C:\WINDOWS\Resources
         9.48     17  SUB-TOTAL: C:\WINDOWS\security
         0.87     25  SUB-TOTAL: C:\WINDOWS\system
       677.42   5524  SUB-TOTAL: C:\WINDOWS\system32
         0.00      1  SUB-TOTAL: C:\WINDOWS\TAPI
         0.00      3  SUB-TOTAL: C:\WINDOWS\Tasks
         0.02      1  SUB-TOTAL: C:\WINDOWS\Temp
         0.00      1  SUB-TOTAL: C:\WINDOWS\twain_32
         2.44     90  SUB-TOTAL: C:\WINDOWS\Web
         5.03     16  SUB-TOTAL: C:\WINDOWS\WinSxS
       956.01  10090  TOTAL
```

To generate a list of top-level folders contained in the Windows folder with their apparent sizes, type the following at the command line:

```
diruse /* /m c:\Windows > c:\data\uncompressed.txt
```

The Uncompressed.txt file in the C:\Data folder contains the output from the command-line statement. The file contains the following output:

```
    Size (mb)  Files  Directory
         0.00      0  SUB-TOTAL: C:\WINDOWS\addins
         0.19     64  SUB-TOTAL: C:\WINDOWS\Application Compatibility Scripts
         1.56      7  SUB-TOTAL: C:\WINDOWS\AppPatch
         2.52     16  SUB-TOTAL: C:\WINDOWS\Cluster
         0.00      0  SUB-TOTAL: C:\WINDOWS\Config
         0.00      0  SUB-TOTAL: C:\WINDOWS\Connection Wizard
         0.29    184  SUB-TOTAL: C:\WINDOWS\Cursors
         0.03      5  SUB-TOTAL: C:\WINDOWS\Debug
         0.00      3  SUB-TOTAL: C:\WINDOWS\Downloaded Program Files
        86.94      1  SUB-TOTAL: C:\WINDOWS\Driver Cache
        12.04    177  SUB-TOTAL: C:\WINDOWS\Fonts
        40.71   1536  SUB-TOTAL: C:\WINDOWS\Help
         0.00      0  SUB-TOTAL: C:\WINDOWS\IIS Temporary Compressed Files
         0.60      3  SUB-TOTAL: C:\WINDOWS\ime
         0.00      0  SUB-TOTAL: C:\WINDOWS\ime (x86)
```

```
    41.86   1323  SUB-TOTAL: C:\WINDOWS\inf
     1.71     29  SUB-TOTAL: C:\WINDOWS\Installer
    10.96     35  SUB-TOTAL: C:\WINDOWS\java
     2.43     40  SUB-TOTAL: C:\WINDOWS\Media
    29.36    137  SUB-TOTAL: C:\WINDOWS\Microsoft
     3.13     30  SUB-TOTAL: C:\WINDOWS\msagent
     0.00      0  SUB-TOTAL: C:\WINDOWS\msapps
     0.00      0  SUB-TOTAL: C:\WINDOWS\mui
     0.00      1  SUB-TOTAL: C:\WINDOWS\Offline Web Pages
    11.75    801  SUB-TOTAL: C:\WINDOWS\PCHEALTH
     1.07      5  SUB-TOTAL: C:\WINDOWS\Registration
    11.62     10  SUB-TOTAL: C:\WINDOWS\repair
     0.00      1  SUB-TOTAL: C:\WINDOWS\Resource Pro Kit
     2.01      4  SUB-TOTAL: C:\WINDOWS\Resources
     9.48     17  SUB-TOTAL: C:\WINDOWS\security
     0.87     25  SUB-TOTAL: C:\WINDOWS\system
   811.63   5524  SUB-TOTAL: C:\WINDOWS\system32
     0.00      1  SUB-TOTAL: C:\WINDOWS\TAPI
     0.00      3  SUB-TOTAL: C:\WINDOWS\Tasks
     0.02      1  SUB-TOTAL: C:\WINDOWS\Temp
     0.00      1  SUB-TOTAL: C:\WINDOWS\twain_32
     2.44     90  SUB-TOTAL: C:\WINDOWS\Web
     5.03     16  SUB-TOTAL: C:\WINDOWS\WinSxS
  1090.21  10090  TOTAL
```

When you compare the output, you see that the only folder that contains compressed files is C:\Windows\System32.

**Example 3: Use DirUse in a batch file**
You want to gather directory and subdirectory information on drive X. You use a batch file so that the data generated by DirUse can be manipulated in Microsoft Excel. The sample batch file contains the following lines of code:

```
X:
cd\
diruse /s x:\ > drvdata.txt
c:
"\program files\microsoft office\office\excel.exe" x:\drvdata.txt
```

The **/s** parameter specifies that you want information about all subdirectories on drive X. The redirect symbol (>) redirects the output to a file called Drvdata.txt.

After DirUse writes the directory information to the text file, the batch file switches back to drive C and starts Microsoft Excel. The command-line argument of the file name causes Microsoft Excel to open Drvdata.txt when the application starts. If you plan to use the batch file frequently, you can create a Microsoft Excel macro to sort, delete, or format the data.

To use this batch file, save the sample code in a file called Drvdata.bat and type the following at the command line:

```
drvdata
```

**See Also**
**Concepts**
Diruse Overview
Diruse Remarks
Diruse Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Diruse Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DirUse Syntax**

DirUse uses the following syntax:

**diruse** [**/s**|**/v**] [**/m**|**/k**|**/b**] [**/c**] [**/,**] [**/q:** *Size*] [**/l**] [**/a**] [**/d**] [**/o**] [**/***] *DirList* [**/?** | **/h** | **/help**]

**Parameters**

**/s**
Includes subdirectories of the specified directories in the output.

**/v**
Writes progress reports while scanning subdirectories. The **/v** parameter is ignored if **/s** is specified.

**/m**
Displays disk usage in megabytes.

**/k**
Displays disk usage in kilobytes.

**/b**
Displays disk usage in bytes (default).

**/c**
Uses compressed file size instead of apparent file size.

See Diruse Examples.

**/,**
Displays the thousands separator (comma or period) in file sizes.

**/q:** ***Size***
Marks, with an exclamation point (!), directories that exceed the specified size. If **/m** (megabytes) or **/k** (kilobytes) are not specified, the size is assumed to be in bytes. If **/q** is specified and any directory is found that exceeds the specified size, the return code is ONE. Otherwise the return code is ZERO.

**/l**
Writes overflows to the log file Diruse.log in the current directory.

**/a**
Specifies that an alert is generated if sizes specified by the **/q:** parameter are exceeded. The Alerter service must be running, and the alert appears only when you are using DirUse.

**/d**
Displays only directories that exceed specified sizes.

**/o**
Specifies that subdirectories are not checked to see if they exceed the specified size.

**/***
Uses the top-level directories residing in the directories specified in *DirList*.

***DirList***
Specifies a list of directories to check. *DirList* is required if you are doing anything other than getting command-line usage. Use semicolons, commas, or spaces to separate multiple directories.

See Diruse Examples for examples where *DirList* is a folder on a network volume and where *DirList* contains a list of folders.

**/?| /h| /help**
Displays command-line usage.

**Note**

- When run without parameters, DirUse displays command-line usage.

**See Also**
**Concepts**
Diruse Overview
Diruse Remarks
Diruse Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Efsinfo Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Efsinfo.exe: Encrypting File System Information**
This command-line tool displays information about files that are encrypted with Encrypting File System (EFS) on NTFS partitions.

The user can also encrypt or decrypt a file by using the command-line tool Cipher, included with the Windows XP Professional or Windows Server 2003 operating system CD.

Efsinfo.exe info is useful for finding certificate thumbprints.

**Corresponding UI**
The **Encryption Details Dialog Box** lists the user accounts with the associated certificate thumbprint who have access to a particular file. It provides an interface to manage access to an encrypted file and displays a list of Data Recovery Agents.

To access the **Encryption Details Dialog Box** right click an encrypted file and click **Properties**. In the **File Properties Dialog Box**, click **Advanced**. In the **Advanced Attributes Dialog Box** click **Details** to open the **Encryption Details Dialog Box**.

**Concepts**
EFS is a feature of Windows that enables users to encrypt and decrypt files. Using EFS, users can protect files from unauthorized access to their sensitive data, including from a stolen laptop or external disk drive. Folders themselves are not encrypted, but are marked so that files that are created within them will be encrypted.

Accessing an EFS folder is controlled by an ACL, not the certificate key. Only the file access is controlled by certificate keys.

In EFS, users work with encrypted files just as they do with any other files, and the encryption is transparent. If the EFS user is the same person who encrypted the file or folder, then the system decrypts the file or folder when the user accesses it later. However, unauthorized users are prevented from access to encrypted files or folders.

For more information on EFS, see "Encrypting File System overview" in Help for Windows Server 2003. For more information on NTFS, see "NTFS" in Help for Windows Server 2003.

**System Requirements**
The following are the system requirements for EFSInfo:

- Windows XP Professional or Windows Server 2003

**File Required**

- Efsinfo.exe

**See Also**
**Concepts**
Efsinfo Syntax
Efsinfo Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview

Connstat Overview
Cabarc Overview
Bitsadmin Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Efsinfo Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**EFSInfo Examples**

**Example: Display current EFS certificate**

If you want to display the current EFS certificate, type the following in the command window:

```
efsinfo /y
```

Information similar to the following output is displayed in the command window:

```
Your current EFS certificate thumbnail information about the computer named A-LOLAG2
is:

  BAAF 99CB 2B76 C49C E49C 6B17 C404 1030 E6BD BF3D
```

**See Also**

**Concepts**

Efsinfo Overview
Efsinfo Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Efsinfo Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**EFSInfo Syntax**

EFSInfo uses the following syntax:

- **efsinfo** [**/u**] [**/r**] [**/c**] [**/i**] [**/y**] [**/k**] [**/s:***Dir*] [*Path*[**,***Path*...]] [**/?**]
- **efsinfo /t:** *dir*

**Parameters**

**Notes**

- When EFSInfo is run without parameters, it displays encryption information about the files and folders in the current folder.
- If no encrypted files or folders have been created on this computer, then EFSInfo returns no information.

**/u**
Displays encryption information about the files and folders in the current folder. This is the default.

**/r**
Displays recovery agent information.

**/c**
Displays certificate thumbnail information.

**/i**
Continues performing the specified operation even after errors have occurred. By default, EFSInfo stops when an error is encountered.

**/y**
Displays the current Encrypting File System (EFS) certificate thumbnail on the local computer. The files specified might not be on this computer. If no items are returned, there are no encrypted files on the system.

**/s: *Dir***
Performs the specified operation on directories in the given directory and all subdirectories.

***Path*[ ,*Path*...]**
Specifies the path of one or more files or directories for which to display encryption information.

**/k**
Displays key and algorithm information used to encrypt the files.

**/t: *dir***
Used to update the EFS key information on the encrypted files in a given directory and all the subdirectories.

**/?**
Displays command-line help.

**See Also**

**Concepts**

Efsinfo Overview
Efsinfo Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview

Filever Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Filever Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Filever.exe: File Version**
This command-line tool examines the version resource structure of a file or a directory of files on either a local or a remote computer. It displays information on the versions of executable files such as .exe files and dynamic-link libraries (DLLs).

FileVer is useful in determining the exact binary build of an installed component. Executable files typically have their file dates and times modified to make them easy to identify as part of a major release. Hot fixes and service packs do not always follow this convention.

By identifying the actual version information, it is possible to track where a fix came from. This can be useful for customers who are requested by support personnel to identify hot fixes installed on their computers.

### Note

- FileVer might not work against localized binaries.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
For more information about comparing binary builds, see the Knowledge Base article Comparing Binary Images of Same Project Builds or Using Binary Version Compatibility.

**System Requirements**
The following is the system requirement for FileVer:

- Windows XP Professional or Windows Server 2003

**File Required**

- Filever.exe

**See Also**
**Concepts**
Filever Syntax
Filever Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
lologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Filever Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FileVer Examples**

**Example: Determine the version of Internet Explorer**
Some applications are not able to correctly determine the version number of various types of upgraded files. FileVer determines the version numbers of applications by examining the header information in the upgraded file. In this example, to determine the version number of Microsoft Internet Explorer, first append the path environment variable to include the Support Tools directory:

```
path= %PATH%;c:\program files\support tools
```

Change directories to the Internet Explorer directory:

```
cd c:\program files\internet explorer
```

Then, run FileVer by typing the following:

```
filever \v iexplore.exe
```

FileVer displays version information about Microsoft Internet Explorer similar to the following:

```
--a-- W32i   APP ENU      6.0.2484.0 shp     91,136 05-29-2001 iexplore.exe
        Language        0x0409 (English (United States))
        CharSet         0x04b0 Unicode
        OleSelfRegister Disabled
        CompanyName     Microsoft Corporation
        FileDescription Internet Explorer
        InternalName    iexplore
        OriginalFilenam IEXPLORE.EXE
        ProductName     Microsoft« Windows« Operating System
        ProductVersion  6.00.2484.0000
        FileVersion     6.00.2484.0000 (main.010529-2005)
        LegalCopyright  ¬ Microsoft Corporation. All rights reserved.

        VS_FIXEDFILEINFO:
        Signature:      feef04bd
        Struc Ver:      00010000
        FileVer:        00060000:09b40000 (6.0:2484.0)
        ProdVer:        00060000:09b40000 (6.0:2484.0)
        FlagMask:       0000003f
        Flags:          00000000
        OS:             00040004 NT Win32
        FileType:       00000001 App
        SubType:        00000000
        FileDate:       00000000:00000000
```

**See Also**
**Concepts**
Filever Overview
Filever Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview

Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Filever Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**FileVer Syntax**

FileVer uses the following syntax:

**filever** [**/s**] [**/v**] [**/e**] [**/x**] [**/b**] [**/a**] [**/d**] [[[*Drive***:**]*Path*]*FileName*]

**Parameters**

**/s**
Displays files in specified directory and all subdirectories.

**/v**
Lists verbose version information if available.

**/e**
Lists executables only.

**/x**
Displays short names generated for non-8.3 file names.

**/b**
Uses bare format (no dir listing).

**/a**
Specifies not to display file attributes.

**/d**
Specifies not to display file date and time.

**[[ *Drive*:] *Path*] *FileName***
Specifies the file using the name and optionally the drive and path.

**See Also**

**Concepts**

Filever Overview
Filever Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# This Topic Is No Longer Available

We're sorry—the topic that you requested is no longer available. Use the search box below to find related information.

# Msicuu Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Msicuu Remarks**

**Msicuu Log File**

Msicuu records the installations that a user deletes in a log file, Msicu.log. It creates a log for each user who runs Msicuu and stores the log in the user profile. The log is cumulative and is updated each time Msicuu runs.

The Msicuu log file, Msicu.log, is located in the *SystemDrive*\Documents and Settings*User*\Local Settings\Temp directory. Because the Local Settings directory is hidden by default, to view the file you must set your Windows Explorer folder options to show hidden files and folders.

The Msicuu log is not updated when you run Msizap.

**Msicuu and Msizap**

Msicuu uses the same programmatic functions as Msizap, and it runs Msizap (with the **T** parameter) to delete registry entries. (You can see the Msizap command window briefly when you click **Remove**.)

Unlike Msizap, Msicuu does not remove rollback data or the InProgress registry subkey, and it does not change permissions to Windows Installer objects. Also, Msicuu removes user-specific installation data only from the profile of the current user. Installation data stored in other user profiles remains in the system.

**Msicuu Products List**

Msicuu displays a list of products that Windows Installer installed on the system. It compiles this list by calling Windows Installer functions that search the registry. Msicuu or Msizap cannot be used on products that do not appear in this list.

Products that are installed on the system, but that do not appear in the list, might have been installed by other technologies, such as those used before the advent of Windows Installer. You must use other methods to clean up these installations.

**See Also**

**Concepts**

Msicuu Overview
Msicuu UI
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Msicuu UI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Msicuu UI**

**To use Msicuu**

- To run Msicuu, in the **Run** dialog box or at the command prompt, type `msicuu`.
- To quit Msicuu, click **Exit**.
- To delete installation data, select one or more product names from the **Installed Products** box, and then click **Remove**.
- To select all products, click **Select All**, or click **Clear All** to clear selection of all products.
- Msicuu prompts you for confirmation. Its actions are not reversible.

**See Also**
**Concepts**
Msicuu Overview
Msicuu Remarks
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Msizap Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Msizap.exe: Windows Installer Zapper**
This command-line tool deletes configuration data that Windows Installer maintains for products that it installs, including the directories, files, registry subkeys, and registry entries in which Windows Installer stores configuration data. You can use Msizap to remove remnants of corrupted Windows Installer installations that prevent you from installing a product correctly.

You can also use Msizap to give administrators Full Control permissions to Windows Installer folders, files, and registry subkeys.

**Caution**

- Msizap can delete installation data that is required by the system. Products installed by Windows Installer might not operate properly after using this tool.

  Msizap is designed as a *last resort* tool to be used only when conventional methods of repairing or removing an installation have failed. Before using Msizap, use Add or Remove Programs to remove or repair the installation. If that fails, use the Msicuu Overview (Msicuu.exe), which runs Msizap, but prevents you from causing unnecessary harm to the system.

You can use Msizap to correct the following installation failures.

- A Windows Installer installation, uninstallation, or rollback is interrupted or fails unpredictably. Because the installation did not complete successfully, Windows Installer will not uninstall or reinstall the program, or install any another programs.

- Registry entries for a product have been deleted or changed, damaging the operation of the products and preventing you from uninstalling or reinstalling the product.

- Multiple instances of a setup program are running simultaneously, or an instance of a setup program cannot run.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Windows Installer manages the installation and removal of applications on Windows XP Professional and Windows Server 2003 operating systems. Typically, these installations conclude successfully, or are successfully reversed ("rolled back") or repaired. However, in rare situations, an installation, a rollback, or an uninstallation fails, leaving installation files and registry entries that would otherwise be removed. As long as these objects remain on the system, Windows Installer considers the installation to be active and incomplete.

Because Windows Installer conducts only one installation at a time, it will not start another installation until all indicators of the active installation are removed.

You can use Msizap to remove particular installation remnants, such as the InProgress registry subkey and the rollback registry data. You can also delete all configuration data for a specified installation, or all Windows Installer installations on the system.

**System Requirements**

- This version of Msizap runs on Windows XP Professional and Windows Server 2003 operating systems. (Windows 2000 includes a version of Msizap for that operating system.)

- The user must be a member of the Administrators group on the computer affected.

- Msizap affects only the local computer. You cannot use it to manage a remote computer.

**File Required**

- Msizap.exe

**See Also**
**Concepts**
Msizap Remarks
Msizap Syntax
Msizap Examples
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Msizap Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Msizap Remarks**
**General Information**

- If you type `msizap` without any parameters, or with an incorrect parameter combination, command-line help displays.
- When you run Msizap, it displays a complete transaction log, including a list of the folders, files, and registry subkeys that it deleted or changed.
- Msizap affects only products installed by Windows Installer. It does not affect products that were installed by other technologies, including older Windows applications that were installed before the advent of Windows Installer.
- Msizap does not uninstall programs or delete objects that programs install during their setup procedures. It only deletes configuration data that Windows Installer stores for its own use.
- Programs removed by using Msizap might no longer appear in **Add or Remove Programs** in Control Panel.

**Troubleshooting Windows Installer**
Msizap is designed to be used as a last resort to remove the remnants of a corrupted installation. It is a blunt instrument that can delete installation data that is required by the system.

Before using Msizap, try the following less risky methods for repairing a failed installation:

- Use Add or Remove Programs in Control Panel to remove or repair the installation.
- Use Msiexec (command-line options), the command-line version of Windows Installer, with parameters to repair the installation and recache the local package. For example, use the following syntax to repair the MyPackage installation:

  **msiexec /fvomus package.msi**

**Troubleshooting Msizap**
When Msizap completes successfully, it displays a log of the objects that it searched, deleted, and changed. To determine whether Msizap ran to completion, look for one of the following statements at the end of the log.

| Final Log Statement | Description |
|---|---|
| Searching for files and folders in the %WINDIR%\Installer folder. | Msizap successfully completed a command using the **T** or **A** parameter. If the Installer folder is empty, then Msizap records a final action, in which it deletes the Installer folder. |
| No product data was found. | Msizap searched successfully, but did not find any of the specified objects on the system. Typically, this happens when an **msizap** command is repeated. |
| FAILED to clear all data. | Msizap cannot delete all of the data requested by the command. Output that ends with this statement also includes a failure statement and an error code at the point in the log at which the failure occurred. Failures are typically caused by problems with permissions, or with folders that do not reside on the local computer. |

**Stopping the Windows Installer Service**
The *** ALLPRODUCTS** and **T** parameters stop the Windows Installer service so that Msizap can obtain exclusive access to the files, folders, and registry subkeys that it needs to change or delete. Because the Windows Installer service starts automatically when needed and stops automatically when idle, you do not need to stop or restart this service manually.

**Identifying Package Files**
All Windows Installer Package files end with the file name extension .msi. However, the file names are sometimes obscure and are often numeric, making it hard to determine which program the package file installs. Fortunately, this data is included in the summary information associated with each package file.

To display the summary information for a package file, use one of the following methods:

- In Windows Explorer, place the mouse pointer over the name of an .msi file. (However, do not click.)
- In Windows Explorer, right-click the name of an .msi file, click **Properties**, and then click the **Summary** tab.

**Folders Deleted by Msizap**
Msizap attempts to delete the following directories when it is run with the ***** or **T** parameter. The list of objects actually deleted appears in the command window when you run Msizap.

- *SystemDrive*\Documents and Settings*UserName*\Application Data\Microsoft\Installer*ProductCode*
- *Systemroot*\MSI
- *Systemroot*\Installer*ProductCode* (a hidden system folder)
- *Drive*:\config.msi (for each local drive)

**Active (Incomplete) Installations**
When an installation does not complete or roll back successfully, some remnants of the installation might remain in the system. Windows Installer considers the incomplete installation to be active and will not install new products or reinstall existing products until the installation is resolved.

Because Windows Installer runs only one installation at a time, only one installation is active at any given time and there is only one instance of the indicator objects on any system.

The following objects indicate an active, incomplete installation on the system. When an installation completes successfully, Windows Installer automatically deletes these objects from the system.

- The InProgress registry subkey:

  **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\InProgress**

- Rollback data, which includes rollback scripts and registry entries for the following Rollback registry subkey:

  **HKLM\Software\Policies\Microsoft\Windows\Installer\Rollback**

To search for the InProgress subkey on a system, use Regedit.exe or Reg.exe, tools included in Windows XP Professional and Windows Server 2003 operating systems. Enter a command that uses the following syntax:

**reg /query** [*****ComputerName*] **HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Installer\InProgress**

To search for the Rollback subkey on a system, use Regedit or Reg, tools included in Windows XP Professional and Windows Server 2003 operating systems. Enter a **reg** command that uses the following syntax:

**reg /query** [*****ComputerName*] **HKLM\Software\Policies\Microsoft\Windows\Installer\Rollback**

**Registry Subkeys Deleted by Msizap**
Msizap attempts to delete the following registry subkeys when it is run with the ***** or **T** parameter. The list of objects actually deleted appears in the command window when you run Msizap.

- **HKCU\Software\Classes\Installer\Features***ProductCode*
- **HKCU\Software\Classes\Installer\Products***ProductCode*
- **HKCU\Software\Microsoft\Installer\Features***ProductCode*

- **HKCU\Software\Microsoft\Installer\Products***ProductCode*
- **HKCU\Software\Microsoft\Installer\UpgradeCodes***ProductCode*
- **HKCU\Software\Microsoft\Windows\CurrentVersion\Installer\UpgradeCodes***ProductCode*
- **HKLM\Software\Classes\Installer\Features***ProductCode*
- **HKLM\Software\Classes\Installer\Products***ProductCode*
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Features***ProductCode*
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Managed***SID***\Installer\Products***ProductCode*
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Managed***SID***\Installer\Features***ProductCode*
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Products***ProductCode*
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UpgradeCodes***ProductCode*
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData***SID***\Components**
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData***SID***\Products**
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Uninstall***ProductCode* -- if there are no longer any installations with the specified product code.
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\{usersid}\Components**
- **HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\{usersid}\Products**

**See Also**
**Concepts**
Msizap Overview
Msizap Syntax
Msizap Examples
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Msizap Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Msizap Examples**

**Example 1: Delete remnants of an active installation**

The following command deletes the InProgress subkey and the indicators of an Msizap Remarks on the computer. Because the command does not specify an installation, Msizap searches all installations until it finds these objects.

In this command, the **P** parameter deletes the InProgress registry subkey and the **S** parameter deletes rollback data:

```
msizap PS
```

In response, Msizap lists the actions it took to comply with the command:

```
MsiZapInfo: Performing operations for user S-1-5-21-2127521184-1604012920-1887927527-
00000

Searching for the Windows Installer InProgress key. . .
   Removed  \Software\Microsoft\Windows\CurrentVersion\Installer\InProgress
Searching for Installer files and folders associated with the product {ALL PRODUCTS}. .
.
  Searching for rollback folders. . .
Searching for the Windows Installer Rollback key. . .
   Removed  \Software\Microsoft\Windows\CurrentVersion\Installer\Rollback
```

The list begins by identifying the user profile from which user-specific data is deleted. Unless the **W** parameter is used, only the SID of the current user appears.

The remaining elements in the list report objects deleted or changed by the command. Note that Msizap searches for all products because the command did not specify a product by name or product code.

**Example 2: Remove installation data for a program**

The following command deletes all Windows Installer configuration data for MyApp.exe, a fictitious program installed by the MyApp.msi package file.

The command identifies the program by using the fully qualified path to the Windows Installer package file for the program. It uses the **T** parameter to delete the folders, files, and registry subkeys that Windows Installer created for the package, and the **P** and **S** parameters to remove rollback data and the InProgress subkey, in case they remain in the system.

Finally, the command redirects the output to a text file (Msizap.log), so that the detailed output is recorded in the file:

```
msizap TPS "C:\Program Files\My Apps\MyApp.msi" > msizap.log
```

The resulting Msizap.log file begins by displaying the SID of the user profile from which Msizap deleted user-specific objects (in this case, the current user profile), followed by a complete log of the objects that it searched, deleted, and changed.

As expected for a command using the **T** parameter, the log ends with a search of the %WINDIR%\Installer folder. If this folder no longer contains any directories, then Msizap deletes the folder. In this case, the folder still contains data for other programs, so Msizap does not delete it.

**Sample output**

```
C:\myapp>msizap tps myapp.msi
MsiZapInfo: Performing operations for user S-1-5-21-2000000004-1000000000-1000000000-
00000

Searching for the Windows Installer InProgress key. . .
Searching for Installer files and folders associated with the product {00000001-0001-
0000-
0000-624474736554}. . .
  Searching for rollback folders. . .
Searching for the Windows Installer Rollback key. . .
Searching for the product {00000001-0001-0000-0000-624474736554} cached package. . .
```

```
   Removed file: C:\WINDOWS\Installer\56cb1ce.msi
Searching for install property data for product {00000001-0001-0000-0000-624474736554}.
..
   Removed  \10000000100000000000264447375645\InstallProperties
Searching for product {00000001-0001-0000-0000-624474736554} data in the
HKLM\Software\Microsoft\Windows\CurrentVersion\Uninstall key. . .
   Removed  \{00000001-0001-0000-0000-624474736554}
Searching user's global config location for product {00000001-0001-0000-0000-
624474736554}
 data. . .
  Searching for product 10000000100000000000264447375645 upgrade codes in
Software\Microsoft\Windows\CurrentVersion\Installer\UpgradeCodes...
   Removed upgrade code '10000000100000000000264447375645' at
HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UpgradeCodes
  Searching for patches for product 10000000100000000000264447375645 in
Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-21-2000000004-
1000000000-1000000000-00000
\Products\10000000100000000000264447375645\Patches
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-21-
2000000004-1000000000-1000000000-00000\Products\10000000100000000000264447375645 for
product
data. . .
   Removed  \Features
   Removed  \Usage
   Removed  \Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-21-
2000000004-1000000000-1000000000-00000\Products\10000000100000000000264447375645
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-21-
2000000004-1000000000-1000000000-00000\Features\10000000100000000000264447375645 for
product
feature data. . .
Searching per-machine global config location for product {00000001-0001-0000-0000-
62447473
6554} data. . .
  Searching for product 10000000100000000000264447375645 upgrade codes in
Software\Microsoft\Windows\CurrentVersion\Installer\UpgradeCodes...
  Searching for patches for product 10000000100000000000264447375645 in
Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-
18\Products\100000001000000000000264447375
645\Patches
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-
18\Pro
ducts\10000000100000000000264447375645 for product data. . .
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-
18\Fea
tures\10000000100000000000264447375645 for product feature data. . .
Searching old global config location for product {00000001-0001-0000-0000-624474736554}
da
ta. . .
  Searching for product 10000000100000000000264447375645 upgrade codes in
Software\Microsoft\Windows\CurrentVersion\Installer\UpgradeCodes...
  Searching for patches for product 10000000100000000000264447375645 in
Software\Microsoft\Windows\CurrentVersion\Installer\Products\100000001000000000000264447
375645\Patches
  Searching
HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Products\100000001000
00000000264447375645 for product data. . .
  Searching
HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Features\100000001000
00000000264447375645 for product feature data. . .
Searching per-machine location for product {00000001-0001-0000-0000-624474736554} data.
..
  Searching for product 10000000100000000000264447375645 upgrade codes in
Software\Classes
\Installer\UpgradeCodes...
  Searching for patches for product 10000000100000000000264447375645 in
Software\Classes\I
nstaller\Products\10000000100000000000264447375645\Patches
  Searching Software\Classes\Installer\Components for published component data for the
pro
duct {00000001-0001-0000-0000-624474736554}. . .
  Searching HKLM\Software\Classes\Installer\Products\10000000100000000000264447375645
for
```

```
product data. . .
  Searching HKLM\Software\Classes\Installer\Features\10000000100000000000264447375645
for
product feature data. . .
Searching old per-user location for product {00000001-0001-0000-0000-624474736554}
data. ..
  Searching for product 10000000100000000000264447375645 upgrade codes in
Software\Classes
\Installer\UpgradeCodes...
  Searching for patches for product 10000000100000000000264447375645 in
Software\Classes\I
nstaller\Products\10000000100000000000264447375645\Patches
  Searching Software\Classes\Installer\Components for published component data for the
product {00000001-0001-0000-0000-624474736554}. . .
  Searching HKCU\Software\Classes\Installer\Products\10000000100000000000264447375645
for
product data. . .
  Searching HKCU\Software\Classes\Installer\Features\10000000100000000000264447375645
for
product feature data. . .
Searching per-user location for product {00000001-0001-0000-0000-624474736554} data. .
.
  Searching for product 10000000100000000000264447375645 upgrade codes in
Software\Microsoft\Installer\UpgradeCodes...
   Removed upgrade code '10000000100000000000264447375645' at
HKCU\Software\Microsoft\Installer\UpgradeCodes
  Searching for patches for product 10000000100000000000264447375645 in
Software\Microsoft\Installer\Products\10000000100000000000264447375645\Patches
  Searching Software\Microsoft\Installer\Components for published component data for
the p
roduct {00000001-0001-0000-0000-624474736554}. . .
   Removed product's published component qualifier value Convertor for published
component
 10000000400000000000264447375645
   Removed published component qualifier value Convertor
   Removed
HKCU\Software\Microsoft\Installer\Components\10000000400000000000264447375645
  Searching HKCU\Software\Microsoft\Installer\Products\10000000100000000000264447375645
fo
r product data. . .
   Removed  \Media
   Removed  \Net
   Removed  \SourceList
   Removed  \Software\Microsoft\Installer\Products\10000000100000000000264447375645
  Searching HKCU\Software\Microsoft\Installer\Features\10000000100000000000264447375645
fo
r product feature data. . .
   Removed  \Software\Microsoft\Installer\Features\10000000100000000000264447375645
Searching for product {00000001-0001-0000-0000-624474736554} in per-user managed
location.
  Searching for product 10000000100000000000264447375645 upgrade codes in
Software\Microsoft\Windows\CurrentVersion\Installer\Managed\S-1-5-21-2000000004-
1000000000-1000000000-00000\Installer\UpgradeCodes...
  Searching for patches for product 10000000100000000000264447375645 in
Software\Microsoft\Windows\CurrentVersion\Installer\Managed\S-1-5-21-2000000004-
1000000000-1000000000-00000\
Installer\Products\10000000100000000000264447375645\Patches
  Searching Software\Microsoft\Windows\CurrentVersion\Installer\Managed\S-1-5-21-
2000000004-1000000000-1000000000-00000\Installer\Components for published component
data for the product {00000001-0001-0000-0000-624474736554}. . .
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Managed\S-1-5-21-
2000000004-1000000000-1000000000-
00000\Installer\Products\10000000100000000000264447375645 for
 product data. . .
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Managed\S-1-5-21-
2000000004-1000000000-1000000000-
00000\Installer\Features\10000000100000000000264447375645 for
 product feature data. . .
Searching for shared DLL counts for components tied to the product
100000001000000000000264
447375645. . .
Searching for shared DLL counts for components tied to the product
```

```
1000000010000000000264
447375645. . .
   Reduced shared DLL count to 1 for: C:\Test\myapp.EXE
  Searching for product 10000000100000000000264447375645 client info data. . .
   Removed client of component 30000000300000000000264447375645
   Removed client of component 40000000300000000000264447375645
   Removed client of component 50000000300000000000264447375645
   Removed client of component 51000000300000000000264447375645
   Removed client of component 71000000300000000000264447375645
   Removed client of component A0000000300000000000264447375645
  Searching for product 10000000100000000000264447375645 client info data. . .
Searching for Installer files and folders associated with the product {00000001-0001-
0000-
0000-624474736554}. . .
  Searching for files and folders in the user's profile. . .
   Removed file: C:\Documents and Settings\user38.RESKIT\Application
Data\Microsoft\Installer\{00000001-0001-0000-0000-624474736554}\advt.ico
   Removed file: C:\Documents and Settings\user38.RESKIT\Application
Data\Microsoft\Installer\{00000001-0001-0000-0000-624474736554}\CD.ico
   Removed folder: C:\Documents and Settings\user38.RESKIT\Application
Data\Microsoft\Installer\{00000001-0001-0000-0000-624474736554}
  Searching for files and folders in the %WINDIR%\Installer folder
```

**Example 3: Change permissions to installation objects**

The following command grants members of the Administrators group on the computer Full Control permission to the folders, files, and registry subkeys for the MyApp.msi installation. The command changes permissions only; it does not delete any installation data.

The command uses the **A** parameter to specify that it is changing permissions, and the **T** and **W** parameters to specify which permissions it is changing. The command changes permissions on objects that the **T** parameter would otherwise delete, that is, the permissions to all files, folders, and registry subkeys that Windows Installer maintains for the installation. The **W** parameter extends the permission changes to objects in all user profiles in the registry.

In this command, the MyApp.msi package is identified by its product code, enclosed in the required braces:

```
msizap ATW {00000001-0001-0000-0000-624474736554}
```

The resulting Msizap.log file is divided into sections for each user profile that Msizap changes. The log lists all the objects for which Msizap searched, and those on which it changed permissions.

As expected for a command using the **A** parameter, each section of the log ends with a search of the %WINDIR%\Installer folder.

**Sample output**

```
   ***** Adjusting ACLs on data for user S-1-5-18 for product
{00000001-0001-0000-0000-624474736554} *****

MsiZapInfo: Performing operations for user S-1-5-18
Searching for the product {00000001-0001-0000-0000-624474736554} cached package. . .
Searching for install property data for product {00000001-0001-0000-0000-624474736554}.
. .
Searching user's global config location for product {00000001-0001-0000-0000-
624474736554} data. . .
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-
18\Prod
ucts\10000000100000000000264447375645 for product data. . .
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-
18\Feat
ures\10000000100000000000264447375645 for product feature data. . .
Searching per-machine global config location for product {00000001-0001-0000-0000-
624474736554} data. . .
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-
18\Prod
ucts\10000000100000000000264447375645 for product data. . .
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-
18\Feat
ures\10000000100000000000264447375645 for product feature data. . .
Searching old global config location for product {00000001-0001-0000-0000-624474736554}
```

```
data. . .
  Searching
HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Products\1000000010000
0000000264447375645 for product data. . .
  Searching
HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Features\1000000010000
0000000264447375645 for product feature data. . .
Searching per-machine location for product {00000001-0001-0000-0000-624474736554} data.
. .
  Searching HKLM\Software\Classes\Installer\Products\10000000100000000000264447375645
for
product data. . .
  Searching HKLM\Software\Classes\Installer\Features\10000000100000000000264447375645
for
product feature data. . .
Searching for product {00000001-0001-0000-0000-624474736554} in per-user managed
location.
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Managed\S-1-5-
18\Insta
ller\Products\10000000100000000000264447375645 for product data. . .
  Searching HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Managed\S-1-5-
18\Insta
ller\Features\10000000100000000000264447375645 for product feature data. . .
  Searching for product 10000000100000000000264447375645 client info data. . .
  Searching for product 10000000100000000000264447375645 client info data. . .
Searching for Installer files and folders associated with the product {00000001-0001-
0000-0000-624474736554}. . .
  Searching for files and folders in the user's profile. . .
   Removed ACL on file: C:\Documents and Settings\user38.RESKIT\Application
Data\Microsoft\Installer\{00000001-0001-0000-0000-624474736554}\advt.ico
   Removed ACL on file: C:\Documents and Settings\user38.RESKIT\Application
Data\Microsoft\Installer\{00000001-0001-0000-0000-624474736554}\CD.ico
   Removed ACL on folder: C:\Documents and Settings\user38.RESKIT\Application
Data\Microsoft\Installer\{00000001-0001-0000-0000-624474736554}
  Searching for files and folders in the %WINDIR%\Installer folder

***** Adjusting ACLs on data for user S-1-5-18 for product
{00000001-0001-0000-0000-624474736554} *****
... (Data truncated)

***** Adjusting ACLs on data for user S-1-5-21-2127521184-1604012920-1887927527-31305
for product {00000001-0001-0000-0000-624474736554} *****
... (Data truncated)

***** Adjusting ACLs on data for user S-1-5-21-2127521184-1604012920-1887927527-31305
for product {00000001-0001-0000-0000-624474736554} *****
... (Data truncated)
  Searching for files and folders in the %WINDIR%\Installer folder
```

**Example 4: Delete configuration data for another user**

The following command deletes Windows Installer configuration data for MyApp.msi for all users of the computer. The command is intended to remove installer data from the profile of User21, a user who is not a member of the Administrators group on the computer. To assure that the user's profile is loaded in the registry, the command must run while User21 is logged on to the system. However, because Msizap runs only with Administrator permissions, an administrator must run this command. In this example, the administrator uses the Runas command extension to provide user credentials. For information about using the Runas command, see "Runas" in Help for Windows Server 2003.

The command uses the **T** parameter to specify that all Windows Installer files, folders, and registry subkeys for MyApp.msi are deleted. It also uses the **W** parameter to make certain that the profile of User21 is affected, because Runas can temporarily make the administrator the current user of the computer.

The command is run by the administrator while User21 is logged on to the computer.

```
runas /env /user:Computer21\Administrator msizap TW MyApp.msi
```

In response, Runas prompts the administrator for a password, and then Msizap responds to the command. All MyApp.exe installer configuration data for all users, including User21, is deleted.

**Example 5: Delete all configuration data from the system**

The following command deletes all Windows Installer configuration data for all programs on the system. Because this

command can cause serious and irreparable harm, it is very rarely used.

The command uses the *** ALLPRODUCTS** parameters to extend the command to all products that Windows Installer installed on the system. It uses the **W** parameter to delete data from all user profiles. The only Windows Installer data that might remain in the system when the command completes is data that is stored in user profiles that are not present in the registry when the command runs.

```
msizap *W ALLPRODUCTS
```

**See Also**
**Concepts**
Msizap Overview
Msizap Remarks
Msizap Syntax
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Msizap Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Msizap Syntax**

Msizap uses the following syntax combinations:

To remove remnants of an active (incomplete) installation

**msizap**[**P**][**S**]

To remove data for a specified installation

**msizap**[**T**][**P**][**S**][**W**][**G**] *PackageName* | **{***ProductCode***}**

To remove all installation data

**msizap ***[**W**][**!**]**ALLPRODUCTS**

To give Administrators Full Control permission to installation data

**msizap A**[*****][**T**][**W**][**P**][**S**] [*PackageName* | **{***ProductCode***}** | **ALLPRODUCTS**]

To display help

**msizap /?**

### Notes

- When deleting objects that are unique in the system (such as the InProgress registry subkey or rollback data) or to change permissions, you do not need to specify an installation package.
- The package name, product code, or **ALLPRODUCTS**, if included, must be the last parameter in the command. Other parameters may be listed in any order.
- When using **T**, you must specify an installation by package name or product code. When using *****, you must specify **ALLPRODUCTS**.

**Parameters**

*****

Deletes all Windows Installer configuration data on the computer. **ALLPRODUCTS** is required when using *****.

This parameter deletes the folders, files, registry subkeys, and registry entries in which Windows Installer stores configuration data, including rollback data and the InProgress registry subkeys for each installation. It also stops the Windows Installer service, deletes published icons (icons that offer to install a program), and adjusts shared DLL counts.

**!**

Suppresses all confirmation messages.

**A**

Gives administrators Full Control permission to the folders, files, and registry subkeys specified by other parameters in the command. When combined with **A**, other parameters specify which object DACLs are changed; they do not direct Msizap to delete installation configuration data.

**P**

Deletes the InProgress registry subkey.

This parameter deletes the subkey that indicates to Windows Installer that an installation is in progress:

**HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\InProgress**

While this subkey remains in the registry, Windows Installer will not install another program. Because Windows Installer deletes the InProgress subkey when an installation completes, this subkey is present only for an active installation.

**S**

Deletes rollback data.

This parameter deletes the Rollback registry subkey (below), rollback scripts, and backup copies that enable Windows Installer to restore the computer to its original state when an installation stops before it is complete.

**HKLM\Software\Microsoft\Windows\CurrentVersion\Installer\Rollback**

Because Windows Installer automatically deletes rollback data when an installation completes, rollback data is present only for an active installation.

**T**

Deletes Windows Installer configuration data for a specified installation.

This parameter deletes the folders, files, registry subkeys, and registry entries in which Windows Installer stores configuration data for the installation (except for rollback data and the InProgress registry subkey, which are specified separately). It also stops the Windows Installer service, deletes published icons (icons that offer to install a program), and adjusts shared DLL counts.

**W**

Extends the command to all user profiles. By default, Msizap deletes and changes user-specific data only for the current user.

**G**

Removes orphaned cached Windows Installer data files for all users.

**?**

Displays help.

***PackageName* | {*ProductCode*}**

Specifies an installation as the target of the **msizap** command. Type the product code or the path and file name of the Windows Installer package file, such as:

```
c:\ptree.msi
```

If you do not specify an installation, then **msizap** commands apply to all Windows Installer installations on the computer.

**ALLPRODUCTS**

Indicates all Windows Installer installations on the computer. **ALLPRODUCTS** is required when using *****.

**Sample Usage**

```
msizap *! ALLPRODUCTS
msizap PS
msizap AT ptree.msi
msizap WTP MyPackage.msi
msizap TW {00000001-0001-0000-0000-624474736554}
```

**See Also**

**Concepts**

Msizap Overview
Msizap Remarks
Msizap Examples
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Showaccs Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Showaccs.exe: Show Access Control Lists**
This command-line tool enables users to examine the access control lists (ACLs) on the registry, file system, and file and print shares, and to look at membership in local groups. ShowAccs command options support looking at specific areas of Windows 2000 and Windows XP (such as specific file-system directories), or examining the entire system.

The output file of ShowAccs is a comma-separated text file, called the access-profile file, that shows all the object-specific access rights on the system. The access-profile file can be reviewed and analyzed by using a text editor, or it can be imported into a spreadsheet or database management program.

ShowAccs also creates a mapping file that lists the security principals (users and groups) that show up somewhere in ACLs. The mapping file can be input to the Security Migration Editor Microsoft Management Console (MMC) snap-in for account mapping.

### Note

- This tool is part of the Sidwalker Security Administration Tools.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Access control is implemented by access control lists (ACLs). Every file in the NTFS file system and every registry key has a unique ACL, granting access rights to file resources to users and groups, and defining what specific access rights each is granted. Each user and group is identified in the ACL by a security identifier (SID).

**System Requirements**
ShowAccs can be run as scheduled batch jobs. It must be run locally on the computer where the access permissions are to be changed. ShowAccs uses Windows NT security APIs (supported on Windows NT 4.0, Windows 2000, and Windows XP) to examine every access control list (ACL) on every object on the system. It is a CPU and I/O intensive program; plan to run it when system use is otherwise very light and expect significant resource use.

In Windows NT 4.0, ShowAccs requires Msvcrt50.dll, which is found in the \System32 directory of any Windows 2000 installation.

**Files Required**

- Showaccs.exe
- Msvcrt50.dll (found in the \System32 directory)

**See Also**
**Concepts**
Showaccs Remarks
Showaccs Syntax
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Showaccs Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ShowAccs Remarks**

**Using ShowAccs**

ShowAccs can examine the registry in approximately five minutes. Walking through the entire registry or file system, however, might require many hours, depending on the size of the file system and the scope of analysis. Examining every ACL on every file and converting security identifiers (SIDs) on a large file system of roughly 50 gigabytes (GB) can take two or more hours with 100 percent CPU utilization.

**See Also**

**Concepts**

Showaccs Overview
Showaccs Syntax
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Showaccs Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ShowAccs Syntax**

ShowAccs uses the following syntax:

**showaccs** *AccessProfileFile*[**/f** [*Path*]][**/r**][**/s**][**/p**][**/g**][**/m** [*MapFilePath*]][**/nobuiltins**][**/?**]

**Parameters**

**Note**

- When run without parameters, ShowAccs displays help.

***AccessProfileFile***

Specifies the path of the .csv file to be generated.

**/f[ *Path*]**

Shows access rights for all NTFS files or, optionally, for those in the directory tree of the specified path.

**/r**

Shows access rights for the registry.

**/s**

Shows access rights for file shares.

**/p**

Shows access rights for printer shares.

**/g**

Shows access rights for local groups.

**/m[ *MapFilePath*]**

Generates a map file, optionally in the directory specified by *MapFilePath*.

**/nobuiltins**

Omits access rights for built-in groups.

**/?**

Displays help on Showaccs.

**See Also**

**Concepts**

Showaccs Overview
Showaccs Remarks
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Spcheck Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Spcheck.exe: Service Pack Checker**
This command-line tool is used by administrators to determine which service pack is installed on a Windows Server 2003, Windows XP Professional, Windows 2000, or Windows NT 4.0 system. SPCheck checks which executable files, dynamic-link libraries, and drivers are installed on the target system. This helps to identify when service packs or hot fixes should be reinstalled, and ensures that all installed services are at the appropriate service pack level.

### Notes

- You must run SPCheck from the command window.
- SPCheck must run under a security context equivalent to the local administrator account.
- To obtain .ini files for SPCheck, search the Microsoft Web site and/or contact Microsoft Product Support Services.

**Concepts**
A *service pack* is a consolidated release of updates for a specific piece of software.

You can find service packs to download at the Windows Update Web site.

**System Requirements**
The following are the system requirements for SPCheck:

- Windows XP Professional or Windows Server 2003, or
- Windows 2000 Professional or Server, or
- Windows NT 4.0

**Files Required**

- Spcheck.exe
- Dbghelp.dll (located in the *Windir*\System32 folder)
- Spcheck.ini

**See Also**
**Concepts**
Spcheck Remarks
Spcheck Syntax
Spcheck Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Spcheck Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SPCheck Remarks**

**Understanding SPCheck**

SPCheck examines target components one file at a time. This means that SPCheck must run under a security context equivalent to the local administrator account. If the tool is not run under this security context, files may be reported as missing even though they are installed on the system. The tool compares the size and date of each file to those listed in Spcheck.ini. SPCheck's search algorithm uses system environment variables in combination with other methods to find and report on target files located on the target system.

**Interpreting the Results**

SPCheck creates a report detailing the results of its scan in the same directory from which SPCheck was run. It has the following format:

```
[WINS]
C:\WINNT\System32\Jet500.dll SP0-3
C:\WINNT\System32\Jet.dll SP0-6a
C:\WINNT\System32\Wins.exe Missing
C:\WINNT\System32\Winsadmin.exe Unknown
C:\WINNT\System32\Winsctrl.dll SP5
```

"Missing" indicates that the file was not found on the target system, and "Unknown" indicates that the file was found on the target system but did not match any of the files recorded in the .ini file.

Not all files are updated for each service pack release. For example, in the report excerpted here, SP0-3 means that Jet500.dll did not change in any of the first three service packs.

**See Also**

**Concepts**

Spcheck Overview
Spcheck Syntax
Spcheck Examples
Alphabetical List of Tools
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Spcheck Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SPCheck Examples**

**Example: Verify that the RRAS services are at the required service pack level**

Your organization has purchased some new RRAS management software made by another company. After it is installed, the RRAS services no longer function properly. You are concerned that the software might have replaced some of the critical RRAS files with previous versions. After searching on the Microsoft Web site, you find an Rras.ini file for use with SPCheck.

To verify that the RRAS files are current, type the following at the command line:

```
spcheck.exe /i rras.ini /r rras.txt
```

In this example, SPCheck used an .ini file named Rras.ini and created a report named Rras.txt. The following output is an excerpt from Rras.txt:

```
SPCHECK Report File
System Date:Tue May 28 14:38:13 2002

The last Windows service pack to be installed on this system was:
Version 5.0 Service Pack 1 (Build 2195)

Check the system analysis section of this report,
to see if all files are Service Pack 1 files

System Analysis
====================
 [RRAS]
  C:\WINNT\System32\Rascfg.dll SP0-3
  C:\WINNT\System32\Rasfil32.dll SP0-2
  C:\WINNT\System32\Rasadmin.exe SP0-3
  C:\WINNT\System32\Rassapi.dll SP0-1
  C:\WINNT\System32\Rasphone.exe SP0-2
```

You can see by this report that if the target system has Service Pack 1 (SP1) installed on it, each file is within the necessary range for the service pack: The target system contains RRAS-related files that are consistent with SP1. This indicates that the issue is not file-related, but probably a configuration problem. You can find more tips on interpreting the report's format at the end of each report.

**See Also**

**Concepts**

Spcheck Overview
Spcheck Remarks
Spcheck Syntax
Alphabetical List of Tools
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Spcheck Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SPCheck Syntax**

SPCheck uses the following syntax:

**spcheck** [**/i***INIFileName*] [**/r***RepFileName*]

**Parameters**

***/i INIFileName***

Specifies the name of the .ini file to be used. If this parameter is omitted, Spcheck.ini is used.

**Note**

- To obtain .ini files for SPCheck, search the Microsoft Web site and/or contact Microsoft Product Support Services.

***/r RepFileName***

Specifies the name of the report file that is generated when the tool is run. If this parameter is omitted, the default report name Spcheck.rpt is used.

**Examples**

See Spcheck Examples.

**See Also**

**Concepts**

Spcheck Remarks
Spcheck Remarks
Spcheck Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Windiff Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windiff.exe: WinDiff File and Directory Comparison**
This GUI tool shows the differences between specified ASCII text files, or folders containing ASCII text files. This is particularly useful for program source code.

The display shows either a summary of the comparison status of a list of files (outline mode) or a detailed line-by-line comparison of one of the files (expanded mode).

You can run WinDiff as a Windows-based application, using menu commands to choose files to compare and perform other actions. Or you can run WinDiff from the command line by using parameters described in the tool's Help.

**Files Required**

- Windiff.exe
- Windiff.hlp
- Gutils.dll

### For More Information

For more information on WinDiff, click Help on the tool's menu. Alternatively, you can run Windiff.hlp from the command line to see the WinDiff help file.

**See Also**
**Concepts**
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Hardware Management Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Hardware Management Tools**
Tools in this category are used to manage the following Windows features:

- Device Manager
- Driver signing
- Hardware devices
- UPS

### Notes

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Internet Services Management Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Internet Services Management Tools**
Tools in this category are used to manage the following Windows features:

- Internet Explorer
- Internet Information Services
- POP3
- Windows SharePoint Services
- Windows Media Services
- UDDI

**Notes**

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Httpcfg Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Httpcfg.exe: HTTP Configuration Utility**
The HTTP Configuration Utility can be used to manage the HTTP application programming interface (API). It may be used to perform many of the tasks associated with web server administration.

**Concepts**
The HTTP API enables applications to communicate over HTTP without using Microsoft Internet Information Services (IIS). Applications can register to receive HTTP requests for particular URLs, receive HTTP requests, and send HTTP responses. The HTTP API includes SSL support so applications can also exchange data over secure HTTP connections without depending on IIS. It is also designed to work with I/O completion ports.

Use this tool to configure a web server in any of the following ways:

- Add, query and delete SSL certificate meta-information. Such meta-information is maintained by the HTTP API in a metastore, and is used to locate certificates for certificate exchange in HTTPS sessions.
- Add, query and delete namespace reservations. The HTTP API allows administrators to reserve URI namespaces and protect them with Access Control Lists (ACLs), so that only specified HTTP API clients can use them.
- Add, query, and delete Internet Protocol (IP) addresses in the IP Listen List. If this list of addresses is specified, the HTTP API listens only to addresses on the list. Otherwise it listens on all IPs on the machine.

**System Requirements**
The following are the system requirements for this tool:

- Windows Server 2003

**Files Required**

- Httpcfg.exe

**See Also**
**Concepts**
Httpcfg Remarks
Httpcfg Syntax
Httpcfg Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Httpcfg Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Httpcfg Remarks**
**The /a Parameter**
**Security Descriptor Definition Language**

The /**a** parameter takes a string in the form of the Security Descriptor Definition Language (SDDL). The SDDL string defines the format that is used by the '**httpcfg** *set urlacl* /**a**' command. The SDDL also defines string elements for describing information in the components of a security descriptor.

For more information, see the Security Descriptor Definition Language topic on the MSDN website (http://msdn.microsoft.com).

**Httpcfg.exe and the HTTP API**
For more information about how Httpcfg corresponds to the HTTP API, see Using the HTTP API Configuration Tool.

**Httpcfg.exe return codes**
The httpcfg.exe commands return a standard WIN32 API error code. A return value of 0 means that the command completed successfully. For more information on WIN32 API error codes, see the Event Logging topic on the MSDN website (http://msdn.microsoft.com).

**See Also**
**Concepts**
Httpcfg Overview
Httpcfg Syntax
Httpcfg Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Httpcfg Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Httpcfg Examples**


[httpcfg ssl](#)

[httpcfg urlacl](#)


**httpcfg ssl**

This group of examples shows how to add, show, and delete a certificate in the SSL store.

**Adding an SSL Certificate to the Store**

In the following example, the user types the **set ssl** command, using the **/i**, **/h**, and **/g** parameters to specify the IP address, Thumbprint hash, and GUID, respectively, for the certificate being added.

```
httpcfg set ssl /i 10.0.0.1:80 /h 2c8bfddf59a4a51a2a5b6186c22473108295624d /g "{2bb50d9c-7f6a-4d6f-873d-5aee7fb43290}"
```

After running the command, Httpcfg displays the following text on the screen to confirm the command completed without an error (error code of 0).

```
HttpSetServiceConfiguration completed with 0.
```

**Viewing Certificates in the SSL Store.**

In this example, the user first types the **query ssl** command with the **/i** parameter, specifying an IP address in order to view the meta-information for a particular certificate. Following that, the user types **query ssl** without the **/i** paremeter, to view all certificates in the store.

```
httpcfg query ssl


     IP                      : 10.0.0.13:80
    Hash                     : 2c8bfddf59a4a51a2a5b6186c22473108295624d
    Guid                     : {2bb50d9c-7f6a-4d6f-873d-5aee7fb43290}
    CertStoreName            : (null)
    CertCheckMode            : 0
    RevocationFreshnessTime  : 0
    UrlRetrievalTimeout      : 0
    SslCtlIdentifier         : (null)
    SslCtlStoreName          : (null)
    Flags                    : 0
-------------------------------------------------------------------------------
    IP                       : 10.0.0.1:80
    Hash                     : 2c8bfddf59a4a51a2a5b6186c22473108295624d
    Guid                     : {2bb50d9c-7f6a-4d6f-873d-5aee7fb43290}
    CertStoreName            : (null)
    CertCheckMode            : 0
    RevocationFreshnessTime  : 0
    UrlRetrievalTimeout      : 0
    SslCtlIdentifier         : (null)
    SslCtlStoreName          : (null)
    Flags                    : 0
-------------------------------------------------------------------------------
```

**Deleting a Certificate from the SSL Store.**

In this example, the user types **delete ssl** with the required **/i** parameter to delete the associated certificate record from the SSL store.

```
httpcfg delete ssl /i 10.0.0.1:80
```

Httpcgf then displays the following text to the screen, verifying that the command completed successfully (error code of 0).

```
HttpDeleteServiceConfiguration completed with 0.
```

**httpcfg urlacl**

This group of examples show how to add, remove, and view URL ACL combinations in the urlacl store.

**Adding a URL ACL combination to the urlacl store**

```
httpcfg set urlacl /u http://woodgrovebank.com:443/ /a "O:DAG:DAD:(A;;GRGX;;;DA)(A;;GA;;;BA)"

HttpSetServiceConfiguration completed with 0.
```

**Viewing all of the URLs that have been assigned an ACL.**

```
httpcfg query urlacl


    URL : http://woodgrovebank.com:443/
    ACL : O:DAG:DAD:(A;;GXGR;;;DA)(A;;GA;;;BA)
-------------------------------------------------------------------------------
    URL : http://woodgrovebank.com:80/
    ACL : O:DAG:DAD:(A;;CCDC;;;SY)(A;;CCDC;;;DA)(OA;;CCDC;bf967aba-0de6-11d0-a285-
00aa003049e2;;AO)(OA;;CCDC;bf967a9c-0d
e6-11d0-a285-00aa003049e2;;AO)(OA;;CCDC;6da8a4ff-0e52-11d0-a286-00aa003049e2;;AO)
(OA;;CCDC;bf967aa8-0de6-11d0-a285-00aa0
03049e2;;PO)(A;;;;;AU)S:(AU;SAFA;CCDCSWWPSDWDWO;;;WD)
-------------------------------------------------------------------------------
```

**Deleting a URL ACL combination from the urlacl store.**

```
httpcfg delete urlacl /u http://woodgrovebank.com:80/

HttpDeleteServiceConfiguration completed with 0.
```

httpcfg iplisten

This group of examples show how to add, delete, and list IP address in the iplisten store.

Adding an IP address to the iplisten store.

```
httpcfg set iplisten -i 10.0.0.1:80

HttpSetServiceConfiguration completed with 0.
```

Viewing all off the IP addresses that the HTTP API is listening on.

```
httpcfg query iplisten

    IP                      : 10.0.0.1:80
-------------------------------------------------------------------------------
    IP                      : 10.0.0.13:80
-------------------------------------------------------------------------------
```

Shows the result of deleting a record from the iplisten store.

```
httpcfg delete iplisten -i 10.0.0.1:80

HttpDeleteServiceConfiguration completed with 0.
```

**See Also**

**Concepts**

Httpcfg Overview
Httpcfg Remarks
Httpcfg Syntax
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Httpcfg Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Httpcfg helps you manage the HTTP API.

#### Httpcfg Syntax

Httpcfg uses the following syntax:

**httpcfg** {**set** | **query** | **delete**} {**ssl** | **query** | **iplisten**} [**/i** *Ip:Port*] [**/h** *SSL Hash*] [**/g"{***GUID***}"**] [**/c** *StoreName*] [**/m** *CheckMode*] [**/r** *RevocationFreshness*] [**/x** *UrlRetrievalTimeout*] [**/t** *SslCtlIdentifier*] [**/n** *SslCtlStoreName*] [**/f** *Flags*] [**/u** {**http://***URL***:***Port***/** | **https://***URL*:*Port***/**}] [**/a** *ACL*]

#### Parameters

"**Action**" commands are: **set**, **query** and **delete**. They are followed by a set of arguments, **ssl**, **urlacl**, and **iplisten**, which are known as "**store**" arguments. Depending on the value of **action** and **store**, different parameters are then available. For example, the **set ssl** command may take a different set of parameters than the **query ssl** command.

- **httpcfg** {**action**} {**store**} [*parameter*]

#### Action commands

**set**

Creates a configuration record containing the values specified by the **ssl**, **urlacl**, or **iplisten** argument. This record is then added to the HTTP API configuration store. The call fails if a record with the specified values already exists. To change a given configuration record, you must first **delete** it, and then recreate it by using **set** with the updated value(s).

**query**

Retrieves one or more HTTP API configuration records.

**delete**

Deletes the specified information, such as IP addresses or SSL certificates, from the HTTP API configuration store, one record at a time.

#### Store arguments

**ssl**

Depending on the **Action** command used, adds (**set**), queries or deletes SSL certificate meta-information. Such meta-information is maintained by the HTTP API in a metastore, and is used to locate certificates for certificate exchange during HTTPS sessions.

**urlacl**

Depending on the **Action** command used, adds (**set**), queries, or deletes namespace reservations. The HTTP API allows administrators to reserve URI namespaces and protect them with Access Control Lists (ACLs), so that only specified HTTP API clients can use them.

**iplisten**

Depending on the **Action** command used, adds (**set**), queries, or deletes Internet Protocol (IP) addresses in the IP Listen List. If this list is present, the HTTP API listens only to addresses on the list.

#### httpcfg set ssl

**/i *Ip:Port***

The **/i** parameter takes a string that specifies the IP-Address:port combination. This serves as the record key identifying the SSL certificate being added. When using **set ssl**, the **/i** parameter is required.

**/h *SSL Hash***

The **/h** parameter takes a string of hexadecimal digits specifying the Thumbprint hash of the certificate being added. This is not a required parameter, however the SSL connection will fail if the hash is invalid.

**/g "{ *GUID* }"**

The **/g** parameter takes a string representing a Globally Unique Identifier (GUID) that identifies the application that added the certificate. In the case of Httpcfg, the GUID must be generated by the user. The enclosing quotation marks and curly braces are required; the **/g** parameter will not work without them. For more information on generating GUIDs, see Generating Interface UUIDs on the Microsoft Developer Network (MSDN).

**/c *StoreName***

The **/c** parameter takes a string that specifies the name of the store where the certificate being added resides. If no string is specified, the name "MY" is used by default.

**/m *CheckMode***

The **/m** parameter takes a string containing one or more numbers representing flags that determine the default mode for checking the certificate. The numbers may consist of one or more of the following flag values:

- 1 - Client certificate will not be verified for revocation.
- 2 - Use cached client certificate revocation.
- 4 - Enable revocation freshness time.
- 65536 - No usage check.

**/r *RevocationFreshness***

The **/r** parameter takes a string of numbers that specify the revocation freshness time for the certificate. Revocation freshness represents the number of seconds after which to check for an updated certificate revocation list (CRL). If this value is absent or zero, the new CRL is updated only when the previous one expires.

**/x *UrlRetrievalTimeout***

The **/x** parameter takes a string of numbers representing the timeout interval, in milliseconds, for retrieving a certificate revocation list from the remote URL.

**/t *SslCtlIdentifier***

The **/t** parameter takes a string that specifies an SSL control identifier, which restricts the group of certificate issuers to be trusted. This group must be a subset of the certificate issuers trusted by the machine being administered.

**/n *SslCtlStoreName***

The **/n** parameter takes a string containing the name of the store in which to look up the control identifier specified by the **/t** parameter.

**/f *Flags***

The **/f** parameter takes a string containing a number that controls how client certificates are handled. This number may consist of one or more of the following values:

- 1 - Use DS Mapper.
- 2 - Negotiate client certificate.
- 4 - Do not route to raw ISAPI filters.

**httpcfg query ssl**

**/i *Ip:Port***

The **/i** parameter takes a string that specifies the IP-Address:port combination. This serves as the record key identifying the SSL certificate being queried. If this parameter is omitted, then the query returns all records in the SSL store.

**httpcfg delete ssl**

**/i *Ip:Port***

The **/i** parameter takes a string specifying the IP-Address:port combination. This serves as the record key identifying the SSL certificate to be deleted. When using **delete ssl**, the **/i** parameter is required.

**httpcfg set urlacl**

**/u {http://*URL:Port*/ | https://*URL:Port*/}**

The **/u** parameter takes a string containing a fully qualified URL that will serve as the record key for the reservation being made. When using **set urlacl**, the **/u** parameter is required.

**/a *ACL***

The **/a** parameter takes a string containing an Access Control List in the form of a Security Descriptor Definition Language (SDDL) string. For more information see Httpcfg Remarks. When using **set urlacl**, the **/a** parameter is required.

**httpcfg query urlacl**

**/u *URL***

The **/u** parameter takes a string containing a fully qualified URL that identifies the reservation being queried. If no string is specified, the query returns all reservations in the store.

**httpcfg delete urlacl**

**/u *URL***

The **/u** parameter takes a string containing a fully qualified URL that identifies the reservation to be deleted. When using **delete urlacl**, the /**u** parameter is required.

**httpcfg set iplisten**

**/i *Ip:Port***

The **/i** parameter takes a string specifying the IP address to be added to the IP-Listen List. This can be either an IPv4 or IPv6 address. When using **set iplisten**, the /**i** parameter is required.

**httpcfg query iplisten**

**/i *Ip:Port***

The **/i** parameter takes a string specifying the exact IP address to be queried. If absent, the query returns all addresses in the store.

**httpcfg delete iplisten**

**/i *Ip:Port***

The **/i** parameter takes a string specifying the IP address to be deleted from the IP-Listen List. This can be either an IPv4 or IPv6 address. When using **delete iplisten**, the /**i** parameter is required.

**See Also**

**Concepts**

Httpcfg Overview
Httpcfg Remarks
Httpcfg Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Network Services Management Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Network Management Tools**
Tools in this category are used to manage the following Windows features:

- Connection Manager
- DHCP
- DNS
- Internet Authentication Service (IAS)
- Internet Protocol (IP)
- Internet Protocol Security (IPSEC)
- Modems
- Novell NetWare Integration
- Routing and Remote Access
- Quality of Service (QoS)
- Services for Macintosh
- Session Initiation Protocol (SIP)
- TCP/IP
- Telephony
- Virtual Private Networks (VPN)
- Windows ATM services
- Wireless networking

**Notes**

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Browstat Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Browstat.exe: Browser Status Utility**
This command-line tool is a general purpose character-based browser diagnostic tool. Use BrowStat to view the status of running browsers in Windows Server 2003 domains. This tool also provides information about the state of the browser, including the name of the master browser.

You can also find out which transports are on a computer. In this context, "transport" refers to the Windows XP or Windows Server 2003 network device name (case-insensitive) for a transport.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Microsoft Windows Server 2003 and Microsoft Windows XP Professional continue to support the browser service. The browser service identifies which resources, domains, and servers are available to a node on the network. It also provides interoperability with domains and computers that are not enabled to use the Active Directory directory service.

**System Requirements**
The following are the system requirements for BrowStat:

- Windows Server 2003

**File Required**

- Browstat.exe

**See Also**
**Concepts**
Browstat Remarks
Browstat Syntax
Browstat Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Browstat Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BrowStat Remarks**

**BrowStat Stats Information**

| Title | Description |
|---|---|
| **NumberOfServerEnumerations** | Number of browse requests for servers |
| **NumberOfDomainEnumerations** | Number of browse requests for domains |
| **NumberOfOtherEnumerations** | Number of "other" browse requests |
| **NumberOfMailslotWrites** | Number of mailslot writes received |
| **NumberOfServerAnnouncements** | Number of server announcements received |
| **NumberOfDomainAnnouncements** | Number of domain (work group) announcements received |
| **NumberOfElectionPackets** | Number of browser election packets received |
| **NumberOfGetBrowserServerListRequests** | Number of GetBrowserServerList requests received |
| **NumberOfMissedGetBrowserServerListRequests** | Number of GetBrowserServerList requests missed |
| **NumberOfDroppedServerAnnouncements** | Number of server announcements dropped |
| **NumberOfDroppedMailslotDatagrams** | Number of mailslot writes dropped because of memory |
| **NumberOfFailedMailslotReceives** | Number of mailslot writes dropped because of transport |
| **NumberOfMasterAnnouncements** | Number of WAN master browser announcements |
| **NumberOfIllegalDatagrams** | Number of illegal datagrams received |

**See Also**
**Concepts**

Browstat Overview
Browstat Syntax
Browstat Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview

# Browstat Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BrowStat Examples**

**Example 1: Display statistics for a browser**

To display statistics for the current browser, type the following at the command line:

**browstat stats**

Your results will look similar to the following output:

```
Browser statistics since 16:7:27.753 on 3/7/2001
NumberOfServerEnumerations:                     5
NumberOfDomainEnumerations:                     5
NumberOfOtherEnumerations:                      2
NumberOfMailslotWrites:                         0
NumberOfServerAnnouncements:                    18
NumberOfDomainAnnouncements:                    9
NumberOfElectionPackets:                        4
NumberOfGetBrowserServerListRequests:           1
NumberOfMissedGetBrowserServerListRequests:     0
NumberOfDroppedServerAnnouncements:             0
NumberOfDroppedMailslotDatagrams:               0
NumberOfFailedMailslotReceives:                 0
NumberOfMasterAnnouncements:                    0
NumberOfIllegalDatagrams:                       0
```

See Browstat Remarks for more information on interpreting this data.

**Example 2: Display status for a domain**

If you want to display the status of the **reskit** domain, type the following at the command line:

**browstat status -v reskit**

Your results will look similar to the following output:

```
Status for domain reskit on transport \Device\NetBT_Tcpip_{D61DB2EB-6C31-4516
-963B-2AACF0E80516}
    Browsing is active on domain.
    Master browser name is: TRINA
        Master browser is running build 2195
    1 backup servers retrieved from master TRINA
        \\TRINA
    There are 3 servers in domain reskit on transport \Device\NetBT_Tcpip_{D6
1DB2EB-6C31-4516-963B-2AACF0E80516}
        Number of NT Advanced Servers:                  0
        Number of OS/2 Domain controllers:              0
        Number of Windows For Workgroups machines:      2
        Number of Os/2 machines:                        0
        Number of NT machines:                          1

        Number of active WfW browsers:                  2
        Number of browsers:                             3
        Number of backup browsers:                      1
        Number of master browsers:                      1
    There are 1 domains in domain reskit on transport \Device\NetBT_Tcpip_{D61DB2EB-
6C31-4516
-963B-2AACF0E80516}
```

**See Also**

**Concepts**

Browstat Overview

Browstat Remarks
Browstat Syntax
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview

# Browstat Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BrowStat Syntax**

Browstat uses the following general syntax:

**browstat** *Operation* [*Options*]

To view the command syntax, click an operation.

Browstat dumpnet

Browstat elect

Browstat getblist

Browstat getmaster

Browstat getpdc

Browstat listwfw

Browstat stats

Browstat status

Browstat tickle

Browstat view

**Browstat dumpnet**

Displays the list of transports bound to browser.

**Syntax**

**browstat dumpnet**

**Sample Usage**

```
browstat dumpnet
```

**Browstat elect**

Forces election on a specified domain.

**Syntax**

**browstat elect** *Transport Domain* [**/help**]

**Parameters**

***Transport***

Enter a network device name (case-insensitive) for a transport in one of the following formats:

- **\device***Transport_***{MAC address}*** (for example, **\device\Nbf_{elnk1601}**)
- *****Transport_{MAC address}*
- *Transport*

***Domain***

Enter the domain in which the election will be forced.

**Browstat getblist**

Gets backup list for domain.

**Syntax**

**browstat getblist** *Transport* [[*Domain*] **refresh**] [**/help**]

**Parameters**

***Transport***

Enter a network device name (case-insensitive) for a transport in one of the following formats:

- **\device***Transport_***{MAC address}*** (for example, **\device\Nbf_{elnk1601}**)
- *****Transport_{MAC address}*
- *Transport*

***Domain***

Enter the name of the domain for which the backup list will be displayed.

**refresh**

Generates an updated browser for a segment.

**Browstat getmaster**

Displays the Master Browser name (using NetBIOS) for a specified domain.

**Syntax**

**browstat getmaster** *Transport Domain*

**Parameters**

***Transport***

Enter a network device name (case-insensitive) for a transport in one of the following formats:

- **\device***Transport_***{MAC address}*** (for example, **\device\Nbf_{elnk1601}**)
- *****Transport_{MAC address}*
- *Transport*

***Domain***

Enter the name of the domain for which the Master Browser name will be displayed.

**Browstat getpdc**

Displays the PDC name (using NetBIOS)for a specified domain.

**Syntax**

**browstat getpdc** *Transport Domain*

**Parameters**

***Transport***

Enter a network device name (case-insensitive) for a transport in one of the following formats:

- **\device***Transport_***{MAC address}*** (for example, **\device\Nbf_{elnk1601}**)
- *****Transport_{MAC address}*
- *Transport*

***Domain***

Enter the domain for which the PDC name is retrieved.

**Browstat listwfw**

Lists Windows for Workgroups computers that are acting as a browser.

**Syntax**

**browstat listwfw** *Domain*

**Parameters**

***Domain***

Enter the domain to which the Windows for Workgroups browser belong.

**Browstat stats**
Displays browser statistics.

**Syntax**
**browstat stats** [*****Computer* [**reset**]]

**Parameters**
**\\ *Computer***
Enter the name of the browser from which statistics are retrieved.

**reset**
Resets the statistics to zero.

**Examples**
See Browstat Examples.

**Browstat status**
Displays status about a domain.

**Syntax**
**browstat status** [**-v**] [*Domain*]

**Parameters**
**-v**
Displays a detailed output.

***Domain***
Enter the domain to query.

**Examples**
See Browstat Examples.

**Browstat tickle**
Forces the master browser to stop.

**Syntax**
**browstat tickle** *Transport* {*Domain*|*****Server*} [*EmulatedDomain*]

**Parameters**
***Transport***
Enter a network device name (case-insensitive) for a transport in one of the following formats:

- **\device***Transport_***{MAC address}*** (for example, **\device\Nbf_{elnk1601}**)
- *****Transport_{MAC address}*
- *Transport*

***Domain***
Enter the name of the domain to query.

**\\ *Server***
Enter the name of the server that receives the query.

***EmulatedDomain***
Enter an alternate domain the specified server should use when announcing its presence.

**Browstat view**
Lists the browsers and their browser roles for a server or domain.

**Syntax**
**browstat view** *Transport* [*Domain*|*****Server* [*Flags*|**/domain** [*DomainToQuery*]]]

**Parameters**
***Transport***

Enter the network device name (case-insensitive) for a transport in one of the following formats:

- **\device***Transport_***{MAC address}*** (for example, **\device\Nbf_{elnk1601}**)
- *****Transport_{MAC address}*
- *Transport*

***Domain***
Specifies the domain to query.

**\\ *Server***
Specifies the server that will receive the query.

***Flags***
Filters the list to show certain browser or server roles.

| Flag | Description |
|---|---|
| **AFP** | AFP server |
| **BBR** | Backup browser |
| **BDC** | Backup domain controller |
| **DFS** | Distributed file system |
| **DL** | Dial-in server |
| **DMB** | Domain master browser |
| **MBR** | Master browser |
| **MDC** | Member domain controller |
| **MFPN** | Microsoft File and Print for NetWare |
| **NT** | Windows NT |
| **NV** | Novell |
| **OSF** | OSF server |
| **PBR** | Potential browser |
| **PDC** | Primary domain controller |
| **PQ** | Print server |
| **S** | Server |
| **SQL** | SQL server |
| **SS** | Windows NT member server |
| **TS** | Time source |
| **VMS** | Vax VMS Server |
| **W** | Workstation |
| **W95** | Windows 95 |
| **WFW** | Windows for Workgroups |

| | |
|---|---|
| **XN** | Xenix |

**/domain [*DomainToQuery*]**

Passes the 0x80000000 flag to the query.

***DomainToQuery***

Enter name of the domain to query.

**See Also**

**Concepts**

Browstat Overview
Browstat Remarks
Browstat Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview

# Dcdiag Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tool Location**
The Dcdiag command-line tool is included when you install Windows Server 2003 Support Tools from the product CD or from the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=100114). For more information about how to install Windows Support Tools, see Install Windows Support Tools (http://go.microsoft.com/fwlink/?LinkId=62270).

**Dcdiag.exe: Domain Controller Diagnostic Tool**
This command-line tool analyzes the state of domain controllers in a forest or enterprise and reports any problems to assist in troubleshooting. As an end-user reporting program, Dcdiag encapsulates detailed knowledge of how to identify abnormal behavior in the system. Dcdiag displays command output at the command line.

Dcdiag consists of a framework for executing tests and a series of tests to verify different functional areas of the system. This framework selects which domain controllers are tested according to scope directives from the user, such as enterprise, site, or single server.

**Corresponding UI**
There is no corresponding user interface (UI) for this tool.

**Concepts**
All domain controllers in the same domain are peers of one another and any domain controller can make directory updates.

However, given the way in which directory updates are replicated from one domain controller to another, it is possible that difficulties can arise. For example, if the necessary domain controllers are not connected by a replication topology, the appropriate domain controllers do not receive directory updates when replication occurs.

Also, in order for the (Domain Controller) Locator to find a domain controller, it must have accurate information so that it can properly locate the resource. If a domain controller is incorrectly advertised, the Locator is unable to find it.

**System Requirements**
The following are the system requirements for DCDiag:

- Windows Server 2003 or Windows XP Professional

| **Note** |
| --- |
| The version of Dcdiag that ships with Windows Server 2003 Service Pack 1 (SP1) is enhanced with a new replication security test (CheckSecurityError). You can run the enhanced version of Dcdiag on computers running Windows Server 2003 with no service pack installed and on computers running Windows XP Professional. However, to run the new replication security test (**dcdiag /test:CheckSecurityError**), you must run Dcdiag on a domain controller running Windows Server 2003 with SP1. |

**File Required**

- Dcdiag.exe

**See Also**
**Concepts**
Dcdiag Remarks
Dcdiag Syntax
Dcdiag Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview

Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Dcdiag Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DCDiag Remarks**

Two DCDiag tests are new in Windows Server 2003 Service Pack 1 (SP1):

- DNS: Checks the health of Domain Name System (DNS) settings for the enterprise.
- CheckSecurityError: Locates security errors or errors that might be related to security problems, and performs initial diagnosis of the problems.

CheckSecurityError must be run on a domain controller running Windows Server 2003 with SP1. Both tests can be run against domain controllers running Windows 2000 Server with Service Pack 3 (SP3) or a later service pack, Windows Server 2003 with no service pack installed, and Windows Server 2003 with SP1.

**DCDiag Tests**

The following are tests that can be run using DCDiag. The tests are divided into three categories: Domain Controller tests that cannot be skipped, Domain Controller tests that can be skipped, and Non-Domain Controller tests. The tests that can be skipped are further divided into those that run by default and those that do not.

**Domain Controller tests that cannot be skipped**

**Connectivity**

Tests whether domain controllers are DNS registered, can be pinged, and have LDAP/RPC connectivity.

**Domain Contoller tests that can be skipped**

**Tests run by default**

**Replications**

Checks for timely replication and any replication errors between domain controllers.

**NCSecDesc**

Checks that the security descriptors on the naming context heads have appropriate permissions for replication.

**NetLogons**

Checks that the appropriate logon privileges exist to allow replication to proceed.

**Advertising**

Checks whether each domain controller is advertising itself in the roles it should be capable of. This test fails if the Netlogon Service has stopped or failed to start.

**KnowsOfRoleHolders**

Checks whether the domain controller can contact the servers that hold the five operations master roles (also know as flexible single master operations or FSMO roles).

**Intersite**

Checks for failures that would prevent or temporarily hold up intersite replication and tries to predict how long it will take before the KCC is able to recover.

| ⚠Caution |
| --- |
| Results of this test are often not valid, especially in atypical site or KCC configurations or at the Windows Server 2003 forest functional level. |

**FSMOCheck**

Checks that the domain controller can contact a KDC, a Time Server, a Preferred Time Server, a PDC, and a Global Catalog server. This test does not test any of the servers for operations master roles.

**RidManager**

Checks whether the RID master is accessible and to see if it contains the proper information.

**MachineAccount**
Checks whether the machine account has properly registered and the services are advertised. Use **/RecreateMachineAccount** to attempt a repair if the local machine account is missing. Use **/FixMachineAccount** if the machine account flags are incorrect.

**Services**
Checks whether the appropriate domain controller services are running.

**OutboundSecureChannels**
Checks that secure channels exist from all of the domain controllers in the domain to the domains specified by **/testdomain**. The **/nositerestriction** parameter prevents the test from being limited to the domain controllers in the site.

**ObjectsReplicated**
Checks that Machine Account and DSA objects have replicated. Use **/objectdn:***dn* with **/n:***nc* to specify an additional object to check.

**frssysvol**
This test checks that the file replication system (FRS) SYSVOL is ready.

**frsevent**
This test checks to see if there are errors in the file replication system. Failing replication of the SYSVOL share can cause policy problems.

**kccevent**
This test checks that the Knowledge Consistency Checker is completing without errors.

**systemlog**
This test checks that the system is running without errors.

**CheckSDRefDom**
This test checks that all application directory partitions have appropriate security descriptor reference domains.

**VerifyReplicas**
This test verifies that all application directory partitions are fully instantiated on all replica servers.

**CrossRefValidation**
This test verifies the validity of cross-references.

**VerifyReferences**
This test verifies that certain system references are intact for the FRS and Replication infrastructure.

**VerifyEnterpriseReferences**
This test verifies that certain system references are intact for the FRS and Replication infrastructure across all objects in the enterprise on each domain controller.

**/skip: *Test***
Skips the specified test. Should not be run in the same command with **/test**. The only test that cannot be skipped is **Connectivity**.

**Tests not run by default**

**Topology**
Checks that the KCC has generated a fully connected topology for all domain controllers.

**CheckSecurityError**
On domain controllers running Windows Server 2003 with SP1, reports on the overall health of replication with respect to Active Directory security. May be performed against one or all domain controllers in an enterprise. When the test has completed, DCDiag presents a summary of the results, along with detailed information for each domain controller tested and the diagnosis of security errors that are encountered.

The following argument is optional:

**/ReplSource:***SourceDomainController*

to check the ability to create a replication link between a real or potential source domain controller (*SourceDomainController*) and the local domain controller.

**CutoffServers**

Checks for any server that is not receiving replications because its partners are down.

**DNS**

New in Windows Server 2003 SP1. Includes six optional DNS-related tests, as well as the **/connectivity test**, which runs by default. The tests can be run individually or all at once. The tests include the following:

- **/DnsBasic** to confirm that essential services are running and available, necessary resource records are registered, and domain and root zones are present.
- **/DnsForwarders** to determine whether recursion is enabled and that any configured forwarders or root hints are functioning.
- **/DnsDelegation** to confirm that the delegated name server is function and to check for broken delegations.
- **/DnsDymanicUpdate** to verify that the Active Directory domain zone is configured for secure dynamic updates and to perform registration of a test record.
- **/DnsRecordRegistration** to test the registration of all essential DC Locator records.
- **/DnsResolveExtName** to verify basic resolution of either an intranet or Internet name.

**OutboundSecureChannels**

Checks that secure channels exist from all of the domain controllers in the domain to the domains specified by **/testdomain**. The **/nositerestriction** parameter prevents the test from being limited to the domain controllers in the site.

**VerifyReplicas**

This test verifies that all application directory partitions are fully instantiated on all replica servers.

**VerifyEnterpriseReferences**

This test verifies that certain system references are intact for the FRS and Replication infrastructure across all objects in the enterprise on each domain controller.

| **Note** |
| --- |
| Text (for example, naming context names and server names) containing international or unicode characters will be displayed correctly only if appropriate fonts and language support are loaded on the test computer. |

**Non-Domain Controller tests**

**DcPromo**

Tests the existing DNS infrastructure for promotion to domain controller. If the infrastructure is sufficient, the computer can be promoted to domain controller in a domain specified in **/DnsDomain:***Active_Directory_Domain_DNS_Name*. It reports whether any modifications to the existing DNS infrastructure are required. Required argument: **/DnsDomain:***Active_Directory_Domain_DNS_Name* One of the following arguments is required:

**/NewForest**

**/NewTree**

**/ChildDomain**

**/ReplicaDC**

If NewTree is specified, the ForestRoot argument is required: **/ForestRoot:***Forest_Root_Domain_DNS_Name*

**RegisterInDNS**

Tests whether this domain controller can register the Domain Controller Locator DNS records. These records must be present in DNS in order for other computers to locate this domain controller for the *Active_Directory_Domain_DNS_Name* domain. It reports whether any modifications to the existing DNS infrastructure are required. Required argument: **/DnsDomain:***Active_Directory_Domain_DNS_Name*

| **Note** |
| --- |
| All tests except DcPromo and RegisterInDNS must be run on computers after they have been promoted to domain controller. |

**See Also**

**Concepts**

Dcdiag Overview
Dcdiag Syntax
Dcdiag Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Browstat Overview

# Dcdiag Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DCDiag Syntax**

DCDiag uses the following syntax:

**dcdiag /s:** *DomainController* [**/n:***NamingContext*] [**/u:***Domain**UserName* **/p:**{***** | *Password* | **""**}] [{**/a** | **/e**}] [{**/q** | **/v**}] [**/i**] [**/f:***LogFile*] [**/ferr:***ErrLog*] [**/c** [**/skip:***Test*]] [**/test:***Test*] [**/fix**] [{**/h** | **/?**}] [**/ReplSource:***SourceDomainController*]

**Parameters**

**/s: *DomainController***

Uses *DomainController* as the home server. This parameter is required. It is ignored for DcPromo and RegisterInDns tests which can only be run locally.

**/n: *NamingContext***

Uses *NamingContext* as the naming context to test. Domains may be specified in NetBIOS, DNS or distinguished name format.

**/u: *Domain\UserName* /p:{* | *Password* | ""}**

Uses *Domain**UserName*DCDiag uses the process's or users default credentials. If alternate credentials are needed, use the following options to provide those credentials for binding with *Password* as the password. Use **""** for an empty or null password, or the wildcard character (*****) to prompt for the password.

**/a**

Tests all the servers on this site.

**/e**

Tests all the servers in the entire enterprise. Overrides **/a**.

**/q**

Quiet. Prints only error messages.

**/v**

Verbose. Prints extended information.

**/i**

Ignores superfluous error messages.

**/fix**

Only affects the *MachineAccount* test. It causes the test to fix the SPNs (Service Principal Names) on the domain controller's Machine Account Object.

**/f: *LogFile***

Redirects all output to *LogFile*. The **/f** parameter operates independently of **/ferr**.

**/ferr: *ErrLog***

Redirects fatal error output to a separate file *ErrLog*. The **/ferr** parameter operates independently of **/f**.

**/c**

Comprehensive. Runs all tests except **DCPromo** and **RegisterInDNS**, including non-default tests. Optionally, can be used with **/skip** to skip specified tests. The following tests are not run by default:

**Topology**

**CutoffServers**

**OutboundSecureChannels**

**{ /h | /?}**

Displays a syntax screen at the command prompt.

**/test: *Test***

Runs only this test. The nonskippable test **Connectivity** is also run. Should not be run in the same command with **/skip**.

**Note**

- All tests except **DcPromo** and **RegisterInDNS** must be run on computers that have been promoted to domain controller.

- The test **CheckSecurityError** is available only in the version of Dcdiag that is included with Windows Support Tools in Windows Server 2003 Service Pack 1 (SP1) and must be run on a domain controller that is running Windows Server 2003 with SP1.

**/ReplSource: *SourceDomainController***
Option for **/test:CheckSecurityError**. Tests the connection between the domain controller on which you run the command and the source domain controller. *SourceDomainController* is the DNS name, NetBIOS name, or distinguished name of a real or potential "from" server that is represented by a real or potential connection object.

**DNS Syntax**
The new DNS tests in Windows Server 2003 SP1 use the following syntax:

**dcdiag /test:DNS** [**/DnsBasic** | **/DnsForwarders** | **/DnsDelegation** | **/DnsDynamicUpdate** | **/DnsRecordRegistration** | **/DnsResolveExtName** [**/DnsInternetName:***InternetName*] | **/DnsAll**] [**/f:***LogFile*] [**/ferr:***ErrLog*] **/s:***DomainController* [**/e**] [**/v**]

**Parameters**
**/test:DNS [*DNS test*]**
Performs the specified DNS test. If no test is specified, defaults to **/DnsAll**.

**/DnsBasic**
Performs basic DNS tests, including network connectivity, DNS client configuration, service availability, and zone existence.

**/DnsForwarders**
Performs the **/DnsBasic** tests, and also checks the configuration of forwarders.

**/DnsDelegation**
Performs the **/DnsBasic** tests, and also checks for proper delegations.

**/DnsDynamicUpdate**
Performs **/DnsBasic** tests, and also determines if dynamic update is enabled in the Active Directory zone.

**/DnsRecordRegistration**
Performs the **/DnsBasic** tests, and also checks if the address (A), canonical name (CNAME) and well-known service (SRV) resource records are registered. In addition, creates an inventory report based on the test results.

**/DnsResolveExtName [/DnsInternetName:*InternetName*]**
Performs the **/DnsBasic** tests, and also attempts to resolve *InternetName*. If **/DnsInternetName** is not specified, attempts to resolve the name www.microsoft.com. If **/DnsInternetName** is specified, attempts to resolve the Internet name supplied by the user.

**/DnsAll**
Performs all tests, except for the **DnsResolveExtName** test, and generates a report.

**/f: *LogFile***
Redirects all output to *LogFile*. The **/f** parameter operates independently of **/ferr**.

**/ferr: *ErrLog***
Redirects fatal error output to a separate file *ErrLog*. The **/ferr** parameter operates independently of **/f**.

**/s: *DomainController***
Runs the tests against *DomainController*.

**/e**
Runs all tests specified by **/test:DNS** against all domain controllers in the Active Directory forest.

**/v**
Verbose. Presents extended information about successful test results, in addition to information about errors and warnings. When the **/v** parameter is not used, provides only error and warning information. Use the **/v** switch when errors or warnings are reported in the summary table.

**See Also**

**Concepts**

Dcdiag Overview
Dcdiag Remarks
Dcdiag Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Browstat Overview

# Dcdiag Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DCDiag Examples**

**Example 1: A normal DC**

In this example, you want to examine the domain controller so you can verify that it is healthy and functioning properly. Type the following command at the command prompt:

```
C:\Program Files\Support Tools>dcdiag /s:reskit-DC1 /u:<username> /p:*
```

Substitute the actual account user name for `<username>`. If you want to specify the domain of the user, use `<domainname>\<username>`. For example, if you want to specify a user account named AlanS from the Contoso domain, type `/u:Contoso\AlanS` to specify the user account. If you use an asterisk (*) for the password, the tool will prompt you for a password. If you want to type the password, you can substitute the actual user account password for the asterisk. However, the password will appear on the screen when you type it.

Output similar to the following appears:

```
Domain Controller Diagnosis

Performing initial setup:
   Done gathering initial info.

Doing initial required tests

   Testing server: Default-First-Site-Name\RESKIT-DC1
      Starting test: Connectivity
         ......................... RESKIT-DC1 passed test Connectivity

Doing primary tests

   Testing server: Default-First-Site-Name\RESKIT-DC1
      Starting test: Replications
         ......................... RESKIT-DC1 passed test Replications
      Starting test: NCSecDesc
         ......................... RESKIT-DC1 passed test NCSecDesc
      Starting test: NetLogons
         ......................... RESKIT-DC1 passed test NetLogons
      Starting test: Advertising
         ......................... RESKIT-DC1 passed test Advertising
      Starting test: KnowsOfRoleHolders
         ......................... RESKIT-DC1 passed test KnowsOfRoleHolders
      Starting test: RidManager
         ......................... RESKIT-DC1 passed test RidManager
      Starting test: MachineAccount
         ......................... RESKIT-DC1 passed test MachineAccount
      Starting test: Services
         ......................... RESKIT-DC1 passed test Services
      Starting test: ObjectsReplicated
         ......................... RESKIT-DC1 passed test ObjectsReplicated
      Starting test: frssysvol
         ......................... RESKIT-DC1 passed test frssysvol
      Starting test: kccevent
         ......................... RESKIT-DC1 passed test kccevent
      Starting test: systemlog
         ......................... RESKIT-DC1 passed test systemlog

   Running partition tests on : Schema
      Starting test: DeadCRTest
         ......................... Schema passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... Schema passed test CheckSDRefDom
```

```
   Running partition tests on : Configuration
      Starting test: DeadCRTest
         ......................... Configuration passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... Configuration passed test CheckSDRefDom

   Running partition tests on : RESKIT-DOM
      Starting test: DeadCRTest
         ......................... RESKIT-DOM passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... RESKIT-DOM passed test CheckSDRefDom

   Running enterprise tests on : RESKIT-DOM.reskit.com
      Starting test: Intersite
         ......................... RESKIT-DOM.reskit.com passed test Intersite
      Starting test: FsmoCheck
         ......................... RESKIT-DOM.reskit.com passed test FsmoCheck
```

**Example 2: Failed DNS registration**

In this example, you have noticed that one of the domain controllers is not replicating properly. After verifying that the domain controller is operational and can be pinged by IP address, use DCDiag to do an enterprise check. Type the following command at the command prompt:

```
C:\Program Files\Support Tools>dcdiag /s:reskit-DC1 /u:<username> /p:* /e
```

Output similar to the following displays:

```
Domain Controller Diagnosis

Performing initial setup:
   Done gathering initial info.

Doing initial required tests

   Testing server: Default-First-Site-Name\RESKIT-DC1
      Starting test: Connectivity
         ......................... RESKIT-DC1 passed test Connectivity

   Testing server: Default-First-Site-Name\RESKIT-DC2
      Starting test: Connectivity
         The host 7594898c-8ba4-4496-a01a-b0f2cadd28a6._msdcs.RESKIT-DOM.reskit.com
could not be resolved to an
         IP address.  Check the DNS server, DHCP, server name, etc
         Although the Guid DNS name
         (7594898c-8ba4-4496-a01a-b0f2cadd28a6._msdcs.RESKIT-DOM.reskit.com)
          couldn't be resolved, the server name
         (reskit-DC2.reskit-sib.RESKIT-DOM.reskit.com) resolved
         to the IP address (172.26.220.34) and was pingable.  Check that the IP
         address is registered correctly with the DNS server.
         ......................... RESKIT-DC2 failed test Connectivity

Doing primary tests

   Testing server: Default-First-Site-Name\RESKIT-DC1
      Starting test: Replications
         [Replications Check,RESKIT-DC1] A recent replication attempt failed:
            From RESKIT-DC2 to RESKIT-DC1
            Naming Context: CN=Configuration,DC=RESKIT-DOM,DC=reskit,DC=com
            The replication generated an error (1722):
            The RPC server is unavailable.
            The failure occurred at 2001-12-21 02:19:04.
            The last success occurred at 2001-12-21 01:57:43.
            1 failures have occurred since the last success.
            The source remains down. Please check the machine.
         ......................... RESKIT-DC1 passed test Replications
      Starting test: NCSecDesc
         ......................... RESKIT-DC1 passed test NCSecDesc
      Starting test: NetLogons
         ......................... RESKIT-DC1 passed test NetLogons
      Starting test: Advertising
```

```
         ......................... RESKIT-DC1 passed test Advertising
      Starting test: KnowsOfRoleHolders
         ......................... RESKIT-DC1 passed test KnowsOfRoleHolders
      Starting test: RidManager
         ......................... RESKIT-DC1 passed test RidManager
      Starting test: MachineAccount
         ......................... RESKIT-DC1 passed test MachineAccount
      Starting test: Services
         ......................... RESKIT-DC1 passed test Services
      Starting test: ObjectsReplicated
         ......................... RESKIT-DC1 passed test ObjectsReplicated
      Starting test: frssysvol
         ......................... RESKIT-DC1 passed test frssysvol
      Starting test: kccevent
         ......................... RESKIT-DC1 passed test kccevent
      Starting test: systemlog
         ......................... RESKIT-DC1 passed test systemlog

   Testing server: Default-First-Site-Name\RESKIT-DC2
      Skipping all tests, because server RESKIT-DC2 is
      not responding to directory service requests

   Running partition tests on : Schema
      Starting test: DeadCRTest
         ......................... Schema passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... Schema passed test CheckSDRefDom

   Running partition tests on : Configuration
      Starting test: DeadCRTest
         ......................... Configuration passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... Configuration passed test CheckSDRefDom

   Running partition tests on : RESKIT-DOM
      Starting test: DeadCRTest
         ......................... RESKIT-DOM passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... RESKIT-DOM passed test CheckSDRefDom

   Running partition tests on : reskit-sib
      Starting test: DeadCRTest
         ......................... reskit-sib passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... reskit-sib passed test CheckSDRefDom

   Running enterprise tests on : RESKIT-DOM.reskit.com
      Starting test: Intersite
         ......................... RESKIT-DOM.reskit.com passed test Intersite
      Starting test: FsmoCheck
         ......................... RESKIT-DOM.reskit.com passed test FsmoCheck
```

**Note**

If you receive output from diagnostic testing similar to the previous example, there is probably a DNS registration issue. For additional diagnostic procedures and potential resolutions, see Troubleshooting DNS (http://go.microsoft.com/fwlink/?LinkID=48893). Also, see Troubleshooting Active Directory—Related DNS Problems (http://go.microsoft.com/fwlink/?LinkId=130597).

**Example 3: Failed Netlogon Service**

In this example, the Netlogon Service has failed on one of the domain controllers. To troubleshoot, type the following command:

```
C:\Program Files\Support Tools>dcdiag /s:reskit-DC1 /u:<username> /p:*
```

Output similar to the following displays:

```
Domain Controller Diagnosis
```

```
Performing initial setup:
   Done gathering initial info.

Doing initial required tests

   Testing server: Default-First-Site-Name\RESKIT-DC1
      Starting test: Connectivity
         ......................... RESKIT-DC1 passed test Connectivity

Doing primary tests

   Testing server: Default-First-Site-Name\RESKIT-DC1
      Starting test: Replications
         ......................... RESKIT-DC1 passed test Replications
      Starting test: NCSecDesc
         ......................... RESKIT-DC1 passed test NCSecDesc
      Starting test: NetLogons
         ......................... RESKIT-DC1 passed test NetLogons
      Starting test: Advertising
         Fatal Error:DsGetDcName (RESKIT-DC1) call failed, error 1722
         The Locator could not find the server.
         ......................... RESKIT-DC1 failed test Advertising
      Starting test: KnowsOfRoleHolders
         ......................... RESKIT-DC1 passed test KnowsOfRoleHolders
      Starting test: RidManager
         ......................... RESKIT-DC1 passed test RidManager
      Starting test: MachineAccount
         ......................... RESKIT-DC1 passed test MachineAccount
      Starting test: Services
            NETLOGON Service is stopped on [RESKIT-DC1]
         ......................... RESKIT-DC1 failed test Services
      Starting test: ObjectsReplicated
         ......................... RESKIT-DC1 passed test ObjectsReplicated
      Starting test: frssysvol
         ......................... RESKIT-DC1 passed test frssysvol
      Starting test: kccevent
         ......................... RESKIT-DC1 passed test kccevent
      Starting test: systemlog
         An Error Event occured.  EventID: 0xC25A001D
            Time Generated: 12/21/2001   01:28:25
            Event String: The time provider NtpClient is configured to
         An Error Event occured.  EventID: 0xC25A001D
            Time Generated: 12/21/2001   01:40:30
            Event String: The time provider NtpClient is configured to
         An Error Event occured.  EventID: 0xC25A001D
            Time Generated: 12/21/2001   01:43:30
            Event String: The time provider NtpClient is configured to
         An Error Event occured.  EventID: 0xC25A001D
            Time Generated: 12/21/2001   01:58:46
            Event String: The time provider NtpClient is configured to
         An Error Event occured.  EventID: 0xC25A001D
            Time Generated: 12/21/2001   02:02:11
            Event String: The time provider NtpClient is configured to
         An Error Event occured.  EventID: 0xC25A001D
            Time Generated: 12/21/2001   02:05:11
            Event String: The time provider NtpClient is configured to
         An Error Event occured.  EventID: 0xC25A001D
            Time Generated: 12/21/2001   02:10:51
            Event String: The time provider NtpClient is configured to
         An Error Event occured.  EventID: 0xC25A001D
            Time Generated: 12/21/2001   02:13:51
            Event String: The time provider NtpClient is configured to
         ......................... RESKIT-DC1 failed test systemlog

   Running partition tests on : Schema
      Starting test: DeadCRTest
         ......................... Schema passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... Schema passed test CheckSDRefDom

   Running partition tests on : Configuration
      Starting test: DeadCRTest
```

```
         ......................... Configuration passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... Configuration passed test CheckSDRefDom

   Running partition tests on : RESKIT-DOM
      Starting test: DeadCRTest
         ......................... RESKIT-DOM passed test DeadCRTest
      Starting test: CheckSDRefDom
         ......................... RESKIT-DOM passed test CheckSDRefDom

   Running enterprise tests on : RESKIT-DOM.reskit.com
      Starting test: Intersite
         ......................... RESKIT-DOM.reskit.com passed test Intersite
      Starting test: FsmoCheck
         Warning: DcGetDcName(GC_SERVER_REQUIRED) call failed, error 1717
         A Global Catalog Server could not be located - All GC's are down.
         Warning: DcGetDcName(PDC_REQUIRED) call failed, error 1717
         A Primary Domain Controller could not be located.
         The server holding the PDC role is down.
         Warning: DcGetDcName(TIME_SERVER) call failed, error 1717
         A Time Server could not be located.
         The server holding the PDC role is down.
         Warning: DcGetDcName(GOOD_TIME_SERVER_PREFERRED) call failed, error 1717
         A Good Time Server could not be located.
         Warning: DcGetDcName(KDC_REQUIRED) call failed, error 1717
         A KDC could not be located - All the KDCs are down.
         ......................... RESKIT-DOM.reskit.com failed test FsmoCheck
```

**Note**

If you receive output from diagnostic testing similar to the previous example, see the list of related Knowledge Base articles on the Windows Operating System Event ID 5179 Netlogon page (http://go.microsoft.com/fwlink/?LinkId=130599) for additional troubleshooting information. Also, see Troubleshooting Active Directory—Related DNS Problems (http://go.microsoft.com/fwlink/?LinkID=130597).

**Example 4: Unresponsive or inaccessible server**

In this example, you have noticed replication problems. To resolve the issue, type the following at the command line:

```
C:\Program Files\Support Tools>dcdiag /s:reskit-DC1 /u:<username> /p:* /e
```

Output similar to the following displays:

```
Domain Controller Diagnosis

Performing initial setup:
   Done gathering initial info.

Doing initial required tests

   Testing server: Default-First-Site-Name\RESKIT-DC1
      Starting test: Connectivity
         ......................... RESKIT-DC1 passed test Connectivity

   Testing server: Default-First-Site-Name\RESKIT-DC2
      Starting test: Connectivity
         Server RESKIT-DC2 resolved to this IP address 172.26.220.34,
         but the address couldn't be reached(pinged), so check the network.
         The error returned was: Error due to lack of resources.
         This error more often means that the targeted server is
         shutdown or disconnected from the network
         ......................... RESKIT-DC2 failed test Connectivity

Doing primary tests

   Testing server: Default-First-Site-Name\RESKIT-DC1
      Starting test: Replications
         [Replications Check,RESKIT-DC1] A recent replication attempt failed:
            From RESKIT-DC2 to RESKIT-DC1
            Naming Context: CN=Configuration,DC=RESKIT-DOM,DC=reskit,DC=com
```

```
         The replication generated an error (1722):
         The RPC server is unavailable.
         The failure occurred at 2001-12-21 02:19:04.
         The last success occurred at 2001-12-21 01:57:43.
         1 failures have occurred since the last success.
         The source remains down. Please check the machine.
      ......................... RESKIT-DC1 passed test Replications
   Starting test: NCSecDesc
      ......................... RESKIT-DC1 passed test NCSecDesc
   Starting test: NetLogons
      ......................... RESKIT-DC1 passed test NetLogons
   Starting test: Advertising
      ......................... RESKIT-DC1 passed test Advertising
   Starting test: KnowsOfRoleHolders
      ......................... RESKIT-DC1 passed test KnowsOfRoleHolders
   Starting test: RidManager
      ......................... RESKIT-DC1 passed test RidManager
   Starting test: MachineAccount
      ......................... RESKIT-DC1 passed test MachineAccount
   Starting test: Services
      ......................... RESKIT-DC1 passed test Services
   Starting test: ObjectsReplicated
      ......................... RESKIT-DC1 passed test ObjectsReplicated
   Starting test: frssysvol
      ......................... RESKIT-DC1 passed test frssysvol
   Starting test: kccevent
      ......................... RESKIT-DC1 passed test kccevent
   Starting test: systemlog
      An Error Event occured.  EventID: 0xC25A001D
         Time Generated: 12/21/2001   01:28:25
         Event String: The time provider NtpClient is configured to
      An Error Event occured.  EventID: 0xC25A001D
         Time Generated: 12/21/2001   01:40:30
         Event String: The time provider NtpClient is configured to
      An Error Event occured.  EventID: 0xC25A001D
         Time Generated: 12/21/2001   01:43:30
         Event String: The time provider NtpClient is configured to
      An Error Event occured.  EventID: 0xC25A001D
         Time Generated: 12/21/2001   01:58:46
         Event String: The time provider NtpClient is configured to
      An Error Event occured.  EventID: 0xC25A001D
         Time Generated: 12/21/2001   02:02:11
         Event String: The time provider NtpClient is configured to
      An Error Event occured.  EventID: 0xC25A001D
         Time Generated: 12/21/2001   02:05:11
         Event String: The time provider NtpClient is configured to
      An Error Event occured.  EventID: 0xC25A001D
         Time Generated: 12/21/2001   02:10:51
         Event String: The time provider NtpClient is configured to
      An Error Event occured.  EventID: 0xC25A001D
         Time Generated: 12/21/2001   02:13:51
         Event String: The time provider NtpClient is configured to
      An Error Event occured.  EventID: 0xC25A001D
         Time Generated: 12/21/2001   02:18:58
         Event String: The time provider NtpClient is configured to
      ......................... RESKIT-DC1 failed test systemlog

Testing server: Default-First-Site-Name\RESKIT-DC2
   Skipping all tests, because server RESKIT-DC2 is
   not responding to directory service requests

Running partition tests on : Schema
   Starting test: DeadCRTest
      ......................... Schema passed test DeadCRTest
   Starting test: CheckSDRefDom
      ......................... Schema passed test CheckSDRefDom

Running partition tests on : Configuration
   Starting test: DeadCRTest
      ......................... Configuration passed test DeadCRTest
   Starting test: CheckSDRefDom
      ......................... Configuration passed test CheckSDRefDom
```

```
Running partition tests on : RESKIT-DOM
   Starting test: DeadCRTest
      ......................... RESKIT-DOM passed test DeadCRTest
   Starting test: CheckSDRefDom
      ......................... RESKIT-DOM passed test CheckSDRefDom

Running partition tests on : reskit-sib
   Starting test: DeadCRTest
      ......................... reskit-sib passed test DeadCRTest
   Starting test: CheckSDRefDom
      ......................... reskit-sib passed test CheckSDRefDom

Running enterprise tests on : RESKIT-DOM.reskit.com
   Starting test: Intersite
      ......................... RESKIT-DOM.reskit.com passed test Intersite
   Starting test: FsmoCheck
      ......................... RESKIT-DOM.reskit.com passed test FsmoCheck
```

**Note**

If you receive output from diagnostic testing similar to the previous example, see article 325487 in the Microsoft Knowledge base (http://go.microsoft.com/fwlink/?LinkID=46059) for additional troubleshooting steps.

**See Also**

**Concepts**

Dcdiag Overview
Dcdiag Remarks
Dcdiag Syntax
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Browstat Overview

# Dhcploc Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dhcploc.exe: DHCP Server Locator Utility**
This command-line tool displays the DHCP servers active on the subnet. If it detects any unauthorized DHCP servers, it beeps and sends out alert messages. It also displays packets that it detects from DHCP servers; you can specify whether to display packets from all DHCP servers or only those from unauthorized servers.

You can also use this tool to determine which DHCP servers are available to a DHCP client and to detect unauthorized DHCP servers on a subnet.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Dynamic Host Configuration Protocol (DHCP) is a TCP/IP standard for simplifying management of host IP configuration. The DHCP standards provide for the use of DHCP servers as a way to manage dynamic allocation of IP addresses and other related configuration details to DHCP-enabled clients on your network.

For more information on DHCP, see DHCP Concepts.

**System Requirements**
The following are the system requirements for DHCPLoc:

- Windows XP Professional or Windows Server 2003

**File Required**

- Dhcploc.exe

**See Also**
**Concepts**
Dhcploc Syntax
Dhcploc Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Dhcploc Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCPLoc Examples**

**Example: Sample DHCPLoc Output**
The output of DHCPLoc is formatted as follows, where "***" indicates a packet from an unauthorized server:

*Time* **(IP)** *IPAddress OfferedPacketType* **(S)** *ServerIPAddress* [*******]

Sample output:

```
17:34:58 (IP)0.0.0.0         NACK      (S)11.11.31.84          ***

17:36:38 (IP)11.101.190.130 OFFER     (S)11.101.12.226       ***

17:36:38 (IP)11.101.196.231 ACK       (S)11.101.13.53

17:36:53 (IP)11.101.196.231 ACK       (S)11.101.13.53

17:37:05 (IP)11.101.196.234 OFFER     (S)11.101.13.53

17:37:05 (IP)11.101.193.232 OFFER     (S)11.101.12.198

17:37:06 (IP)11.101.190.132 OFFER     (S)11.101.12.226        ***
```

**See Also**
**Concepts**
Dhcploc Overview
Dhcploc Syntax
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dcdiag Overview
Browstat Overview

# Dhcploc Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCPLoc Syntax**

DHCPLoc uses the following syntax:

**dhcploc** [**/p**] [**/a:"***AlertNameList***"**] [**/i:***AlertInterval*] *ComputerIPAddress* [*ValidDHCPServerList*]

**Parameters**

**/p**

Suppresses display of detected packets from any of the authorized DHCP servers specified in *ValidDHCPServerList*.

**/a:" *AlertNameList* "**

Sends alert messages to the names in *AlertNameList* if any unauthorized DHCP servers are found.

**/i: *AlertInterval***

Specifies the alert frequency in seconds.

***ComputerIPAddress***

Specifies the IP address of the computer from which you are running DHCPLoc. If the computer has multiple adapters, you must specify the IP address of the adapter that is connected to the subnet you want to test.

***ValidDHCPServerList***

Specifies the IP addresses of any number of authorized DHCP servers. The tool does not send alerts when it detects packets from the servers in this list; however, it displays those packets unless you use the **/p** parameter.

**See Also**

**Concepts**

Dhcploc Overview
Dhcploc Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dcdiag Overview
Browstat Overview

# Dnscmd Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dnscmd.exe: DNS Server Troubleshooting Tool**
This command-line tool assists administrators in Domain Name System (DNS) management.

DNSCmd displays and changes the properties of DNS servers, zones, and resource records. It manually modifies these properties, creates and deletes zones and resource records, and forces replication events between DNS server physical memory and DNS databases and data files. Some operations of this tool work at the DNS server level while others work at the zone level.

### Note

- DNSCmd enhances the functionality of and replaces Dnsstat.exe, a tool included in some versions of the *Windows NT Resource Kit*.

**Corresponding UI**
To manually view and manage DNS by using the DNS Server snap-in in Windows

1. Click **Start**, point to **Programs**, and then point to **Administrative Tools**.
2. Click **DNS**.
3. For information about how to use DNS, right-click **DNS**, and then click **Help**.

Open DNS Server snap-in now.

**Concepts**
DNS resolves computer names to IP addresses so that users can refer to computers by name rather than by a series of numbers.

For more information about DNS, see "DNS" and "DNS Concepts" in Help for Windows Server 2003.

**System Requirements**
DNSCmd runs on a source computer and acts on a target computer. The target computer can be the same computer as the source computer, or it can be a different computer.

**Source Computer Requirements**
The following are the system requirements for the source computer:

- Dnscmd.exe
- User's membership in the Administrators or Server Operators group on the target computer. Both the user account and the server computer must be members of the same domain or reside within trusted domains.

**Target Computer Requirements**
The DNS service must be installed on the target computer.

**File Required**

- Dnscmd.exe

**See Also**
**Concepts**
Dnscmd Syntax
Dnscmd Examples
Alphabetical List of Tools

Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Dnscmd Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNSCmd Examples**

**Example 1: Set the current time on a time stamp to resource records**

To set the current time on a time stamp to resource records, type:

```
dnscmd reskit.com /ageallrecords test.reskit.com
```

The following question is displayed:

```
Are you sure you want to force aging on node?
```

Type `y` to apply the time. Output similar to the following displays:

```
DNS Server reskit.com forced aging on records at root of zone test.reskit.com:
     Status = 0 (0x00000000)
Command completed successfully.
```

**Example 2: Display a complete list of zones on a DNS server**

To see a complete list of zones on your DNS server, type:

```
dnscmd reskit.com /enumzones
```

Output similar to the following displays:

```
Enumerated zone list:
     Zone Count = 6.

.                               0 DS            Up=0
0.0.10.in-addr.arpa             0 DS   Rev      Up=1
0.in-addr.arpa                  1 file Rev Auto Up=0
255.in-addr.arpa                1 file Rev Auto Up=0
reskit.com                      1 DS            Up=0
test.reskit.com                 1 file          Up=1

Command completed successfully.
```

**Example 3: Display a list of autocreated zones on a DNS server**

To display a list of autocreated zones that are also reverse lookup zones on the DNS server, type:

```
dnscmd reskit.com /enumzones /auto-created /reverse
```

Output similar to the following displays:

```
Enumerated zone list:
     Zone Count = 2.
0.in-addr.arpa                  1 file Rev Auto Up=0
255.in-addr.arpa                1 file Rev Auto Up=0

Command completed successfully.
```

**Example 4: Display the IsSlave setting from a DNS server**

To display the IsSlave setting from a DNS server, type:

```
dnscmd reskit.com /info isslave
```

Output similar to the following displays:

```
Query result:
Dword: 0 (00000000)
```

```
Command completed successfully.
```

**Example 5: Display the RecursionTimeout setting from a DNS server**
To display the RecursionTimeout setting from a DNS server, type:

```
dnscmd reskit.com /info recursiontimeout
```

Output similar to the following displays:

```
Query result:
Dword: 15 (0000000F)
Command completed successfully.
```

**Example 6: Delete the records from a node**
To delete the records in a node, type:

```
dnscmd reskit.com /nodedelete test.reskit.com node /tree
```

The following question displays:

```
Are you sure you want to delete node's subtree? y
```

Type `y` to delete the subtree.

Output similar to the following displays:

```
DNS Server reskit.com deleted node at node.test.reskit.com:
     Status = 0 (0x00000000)
Command completed successfully.
```

To delete the records in a node using the host, type:

```
dnscmd reskit.com /NodeDelete test.reskit.com host /F
```

Output similar to the following displays:

```
DNS Server reskit.com deleted node at host.test.reskit.com:
     Status = 0 (0x00000000)
Command completed successfully.
```

**Example 7: Display time statistics for a DNS server**
To display time statistics for a DNS server, type:

```
dnscmd reskit.com /statistics 00000001
```

Output similar to the following displays:

```
DNS Server reskit.com statistics:

DNS Server Time Statistics:
Server start time 3/7/2001 1:07:31 PM
Seconds since start 8482
Stats last cleared 3/7/2001 1:07:31 PM
Seconds since start 8482

Command completed successfully.
```

**Example 8: Display NbstatMem statistics for a DNS server**
To display NbstatMem statistics for a DNS server, type:

```
DNSCmd reskit.com /Statistics 00200000
```

Output similar to the following displays:

```
DNS Server reskit.com statistics:

NBStat Memory Usage:
NBStat Buffers:
```

```
     Alloc = 0
     Free = 0
     NetAllocs = 0
     Memory = 0
     Used = 0
     Returned = 0
     InUse = 0
     InFreeList = 0

Command completed successfully.
```

**Example 9: Delete a zone from a DNS server**
To delete the test.reskit.com zone from a server, type:

```
dnscmd reskit.com /zonedelete test.reskit.com
```

The following question is displayed:

```
Are you sure you want to ZoneDelete? (y/n)
```

Type `y` to delete the zone.

Output similar to the following displays:

```
DNS Server reskit.com deleted zone test.reskit.com:
        Status = 0 (0x00000000)
Command completed successfully.
```

**Example 10: Export zone resource records list to a file**
To export the resource record list from the test.reskit.com zone on the reskit.com DNS server, type:

```
dnscmd reskit.com /zoneexport test.reskit.com test.reskit.com.dns
```

Output similar to the following displays:

```
DNS Server reskit.com exported zone
     test.reskit.com to file
C:\Windows\System32\DNS\test.reskit.com.dns
Command completed successfully.
```

**Example 11: Display RefreshInterval setting from the registry**
To display the values in the RefreshInterval entry in the registry, type:

```
dnscmd reskit.com /zoneinfo test.reskit.com refreshinterval
```

Output similar to the following displays:

```
Zone query result:
Dword: 168 (000000A8)
Command completed successfully.
```

**Example 12: Display Aging setting from the registry**
To display the values in the Aging entry in the registry, type:

```
dnscmd reskit.com /zoneinfo test.reskit.com aging
```

Output similar to the following displays:

```
Zone query result:
Dword: 0 (00000000)
Command completed successfully.
```

**See Also**
**Concepts**
Dnscmd Overview

Dnscmd Syntax
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Dnscmd Syntax

Updated: November 5, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNSCmd Syntax**

DNSCmd uses the following general syntax:

**dnscmd** *ServerName Command* [*Command Parameters*]

**Operations**

Dnscmd ageallrecords

Dnscmd clearcache

Dnscmd config

Dnscmd createbuiltindirectorypartitions

Dnscmd createdirectorypartition

Dnscmd deletedirectorypartition

Dnscmd directorypartitioninfo

Dnscmd enlistdirectorypartition

Dnscmd enumdirectorypartitions

Dnscmd enumrecords

Dnscmd enumzones

Dnscmd info

Dnscmd nodedelete

Dnscmd recordadd

Dnscmd recorddelete

Dnscmd resetforwarders

Dnscmd resetlistenaddresses

Dnscmd startscavenging

Dnscmd statistics

Dnscmd unenlistdirectorypartition

Dnscmd writebackfiles

Dnscmd zoneadd

Dnscmd zonechangedirectorypartition

Dnscmd zonedelete

Dnscmd zoneexport

Dnscmd zoneinfo

Dnscmd zonepause

Dnscmd zoneprint

Dnscmd zoneresettype

Dnscmd zonerefresh

Dnscmd zonereload

Dnscmd zoneresetmasters

Dnscmd zoneresetscavengeservers

Dnscmd zoneresetsecondaries

Dnscmd zoneresume

Dnscmd zoneupdatefromds

Dnscmd zonewriteback

**Dnscmd ageallrecords**

Sets the current time on a time stamp to resource records at a specified zone or node on a DNS server.

**Syntax**

**dnscmd** [*ServerName*] **/ageallrecords** *ZoneName*[*NodeName*] | [**/tree**]|[**/f**]

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Specifies the FQDN name of the zone.

***NodeName***

Specifies a specific node or subtree in the zone.

**Note**

- NodeName specifies the node or subtree in the zone using:
  - **@** for root zone or FQDN.
  - FQDN of a node (name with a '.' at the end).
  - A single label for the name relative to zone root.

**/tree**

Specifies that all child nodes also receive the time stamp. The /tree switch for /ageallrecords will often be desired. Without it the command pertains only to the single node. This may be the desired behavior, but it is often desirable to execute this command on the entire zone or on an entire sub-tree of the zone.

**/f**

Executes the command without asking for confirmation.

**Remarks**

- The **ageallrecords** operation is for backward compatibility between the current version of DNS and previous releases in which aging and scavenging were not supported. It adds a time stamp with the current time to records that do not have one. It will not change the timestamp on records that already have a time stamp.
- Scavenging of records does not occur unless the records are timestamped. NS (name server), SOA (Start of Authority), and WINS records are not included in the scavenging process and are not timestamped even when the **ageallrecords** operation is run.
- This command fails unless scavenging is enabled for the DNS server and the zone. For information about how to enable scavenging for the zone, see the **aging** parameter under Zone-Level Syntax in the Dnscmd config operation in this document.

- The addition of a time stamp to DNS records makes them incompatible with DNS servers that run on operating systems other than Windows 2000, Windows XP, or Windows Server 2003. A time stamp you add by using the **ageallrecords** operation cannot be reversed.

- If none of the optional fields are specified, the command returns all records at the specified node. If a value is specified for at least one of the optional fields, then DNSCmd enumerates only records corresponding to the value or values specified in the optional field or fields.

**Example**
See Dnscmd Examples.

**Dnscmd clearcache**
Clears the DNS cache memory of resource records in the specified DNS server.

**Syntax**
**dnscmd** [*ServerName*] **/clearcache**

**Parameter**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

**Sample Usage**
```
dnscmd dnssvr1.contoso.com /clearcache
```

**Dnscmd config**
Allows the user to change values in the registry for the DNS server and individual zones. Accepts server-level and zone-level settings.

**Caution**

- Do not edit the registry directly unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can degrade performance, damage your system, or even require you to reinstall Windows. You can safely alter most registry settings by using the programs in Control Panel or Microsoft Management Console (MMC). If you must edit the registry directly, back it up first. Read the Registry Editor Help for more information.

**Server-Level Syntax**
**dnscmd** [*ServerName*] **/config** *Parameters*

**Parameters**
***ServerName***
Specifies the DNS server the administrator is planning to manage, represented by local computer syntax, IP address, FQDN, or Host name. If omitted, the local server is used.

***Parameters***
Specify a setting and, optionally, a value. *Parameters* values use this syntax: *Parameter* [*Value*]

The following *Parameters* values are documented in the remainder of this section:

- **/addressanswerlimit**
- **/bindsecondaries**
- **/bootmethod**
- **/defaultagingstate**
- **/defaultnorefreshinterval**
- **/defaultrefreshinterval**
- **/disableautoreversezones**

- **/disablensrecordsautocreation**
- **/dspollinginterval**
- **/dstombstoneinterval**
- **/ednscachetimeout**
- **/enablednsprobes**
- **/enablednssec**
- **/eventloglevel**
- **/forwarddelegations**
- **/forwardingtimeout**
- **/isslave**
- **/localnetpriority**
- **/logfilemaxsize**
- **/logfilepath**
- **/logipfilterlist**
- **/loglevel**
- **/maxcachesize**
- **/maxcachettl**
- **/namecheckflag**
- **/notcp**
- **/norecursion**
- **/recursionretry**
- **/recursiontimeout**
- **/roundrobin**
- **/rpcprotocol**
- **/scavenginginterval**
- **/secureresponses**
- **/sendport**

- **/strictfileparsing**
- **/updateoptions**
- **/writeauthorityns**
- **/xfrconnecttimeout**

**/addressanswerlimit [0|5-28]**
Specifies the maximum number of host records that a DNS server can send in response to a query. The value can be zero (0) or set in the range of 5-28 records. The default value is zero (0).

**/bindsecondaries [0|1]**
Changes the format of the zone transfer so that it can achieve maximum compression and efficiency. However, this efficient format is incompatible with earlier versions of BIND.

**0**
Uses maximum compression. Compatible with BIND versions 4.9.4 and later only.

**1**
Sends only one resource record per message to non-Microsoft DNS servers. Compatible with BIND versions earlier than 4.9.4. This is the default setting.

**/bootmethod [0|1|2|3]**
Determines the source from which the DNS server gets its configuration information.

**0**
No source.

**1**
Loads from the BIND file that is located in the DNS directory; by default, %systemroot%\System32\DNS..

**2**
Loads from the registry.

**3**
Loads from Active Directory and the registry. This is the default setting.

**/defaultagingstate [0|1]**
Determines whether the DNS scavenging feature is enabled by default on newly created zones.

**0**
Disables scavenging. This is the default setting.

**1**
Enables scavenging.

**/defaultnorefreshinterval [0x1-0xFFFFFFFF|0xA8]**
Sets a period of time in which no refreshes are accepted for dynamically updated records. This value is inherited automatically by zones on the server. To change the default, enter a value in the range of 0x1-0xFFFFFFFF. The default value from the server is 0xA8.

**/defaultrefreshinterval [0x1-0xFFFFFFFF|0xA8]**
Sets a period of time that is allowed for dynamic updates to DNS records. This value is inherited automatically by zones on the server. To change the default, enter a value in the range of 0x1-0xFFFFFFFF. The default value from the server is 0xA8.

**/disableautoreversezones [0|1]**
Enables or disables the automatic creation of reverse lookup zones.

**0**
Enables creation of autoreverse zones. This is the default setting.

**1**
Disables creation of autoreverse zones.

**/disablensrecordsautocreation {0|1}**
Specifies whether the DNS server automatically creates NS records for zones that it hosts.

**0**
Automatically creates NS records for zones that it hosts.

**1**
Does not automatically create NS records for zones that it hosts.

**/dspollinginterval 0-3600**
Specifies how often the DNS server polls Active Directory for changes in Active Directory-integrated zones.

| **Note** |
| --- |
| If the DNS server is running Windows Server 2008 or above, setting a value of 0 for dspollinginterval will result in the default interval of 180 seconds being configured, and values of 1-29 are not allowed. |

**/dstombstoneinterval [1-30]**
Amount of time in seconds to keep tombstoned records in Active Directory alive.

**/ednscachetimeout [3600-15724800]**
Specifies the number of seconds that EDns information is cached. The default value is 604,800 seconds (one week).

**/enableednsprobes {0|1}**
Enables or disables EDnsProbes.

**0**
Disables active support for EDnsProbes.

**1**
Enables active support for EDnsProbes.

**/enablednssec {0|1}**
Enables or disables DNSSEC.

**0**
Disables DNSSEC.

**1**
Enables DNSSEC.

**/eventloglevel [0|1|2|4]**
Determines which events are logged in the DNS server log in Event Viewer.

**0**
Logs no events.

**1**
Logs only errors.

**2**
Logs only errors and warnings.

**4**
Logs errors, warnings, and informational events. This is the default setting.

**/forwarddelegations [0|1]**
Determines how a query for a delegated subzone is handled by the DNS server. These queries can either be sent to the subzone referred to in the query or be sent to the list of forwarders named for the DNS server. Entries in the setting are used only when forwarding is enabled.

**0**
Automatically sends queries referring to delegated subzones to the appropriate subzone. This is the default setting.

**1**
Forwards queries referring to the delegated subzone to the existing forwarders.

**/forwardingtimeout [0x1-0xFFFFFFFF|0x5]**
Determines how many seconds (0x1-0xFFFFFFFF) a DNS server waits for a forwarder to respond before trying another. The default value is 0x5, which is 5 seconds.

**/isslave [0|1]**
Determines how the DNS server responds when forwarded queries receive no response.

**0**
Not a subordinate (also known as a *slave*). If the forwarder does not respond, the server attempts to resolve the query itself. This is the default setting.

**1**
A subordinate. If the forwarder does not respond, the server terminates the search and sends a failure to the resolver.

**/localnetpriority [0|1]**
Determines the order in which host records are returned when the server has multiple host records for the same name.

**0**
Returns the records in the order in which they are listed in the DNS database.

**1**
Returns the records that have similar IP network addresses first. This is the default setting.

**/logfilemaxsize [0x10000-0xFFFFFFFF|0x400000]**
Specifies the maximum size in bytes (0x10000-0xFFFFFFFF) of the Dns.log file. When the file reaches its maximum, DNS overwrites the oldest events. The default size is 0x400000, which is 4 MB.

**/logfilepath [*Path+LogFileName*]**
Specifies the path of the Dns.log file. The default path is %systemroot%\System32\Dns\Dns.log. A different path can be specified by using the format *Path+LogFileName*.

**/logipfilterlist *IPAddress* [,*IPAddress*...]**
Specifies which packets are logged in the debug log file. The entries are a list of IP addresses. Only packets going to and from the IP addresses in the list are logged.

**/loglevel [*EventType*]**
Determines which types of events are recorded in the Dns.log file. Each type is represented by a hexadecimal number. If you want more than one event in the log, use hexadecimal addition to add the values, then enter the sum.

**0x0**
The DNS server does not create a log. This is the default entry.

**0x10**
Logs queries.

**0x10**
Logs notifications.

**0x20**
Logs updates.

**0xFE**
Logs non-query.

**0x100**
Logs question transactions.

**0x200**
Logs answers.

**0x1000**
Logs send packets.

**0x2000**
Logs receive packets.

**0x4000**
Logs UDP packets.

**0x8000**
Logs TCP packets.

**0xFFFF**
Logs all packets.

**0x10000**
Logs Active Directory write transactions.

**0x20000**
Logs Active Directory update transactions.

**0x1000000**
Logs full packets.

**0x80000000**
Logs write-through transactions.

**/maxcachesize**
Specifies the maximum size, in kilobytes, of the DNS server's memory cache.

**/maxcachettl [0x0-0xFFFFFFFF|0x15180]**
Determines how many seconds (0x0-0xFFFFFFFF) a record is saved in cache. If the 0x0 setting is used, then the DNS server does not cache records. The default setting is 0x15180 (86,400 seconds or 1 day).

**/maxnegativecachettl [0x1-0xFFFFFFFF|0x384]**
Specifies how many seconds (0x1-0xFFFFFFFF) an entry that records a negative answer to a query remains stored in the DNS cache. The default setting is 0x384 (900) seconds.

**/namecheckflag [0|1|2|3]**
Specifies which character standard is used when allowing DNS names.

**0**
Uses ANSI characters that comply with IETF Requests For Comment (RFCs).

**1**
Uses ANSI characters that do not necessarily comply with IETF RFCs.

**2**
Uses multibyte UTF8 characters. This is the default setting.

**3**
Uses all characters.

**/norecursion [0|1]**
Determines whether a DNS server performs recursive name resolution.

**0**
DNS server performs recursive resolution if it is requested in a query. This is the default setting.

**1**
DNS server does not perform recursive resolution.

**/notcp**
This parameter is obsolete and has no effect in Windows Server 2003.

**/recursionretry [0x1-0xFFFFFFFF|0x3]**
Determines the number of seconds (0x1-0xFFFFFFFF) a DNS server waits before again trying to contact a remote server. The default setting is 0x3 (3 seconds). This value should be increased when recursion occurs over a slow WAN link.

**/recursiontimeout [0x1-0xFFFFFFFF|0xF]**
Determines the number of seconds (0x1-0xFFFFFFFF) a DNS server waits before discontinuing attempts to contact a remote server. The settings range from 0x1 through 0xFFFFFFFF. The default setting is 0xF (15 seconds). This value should be increased when recursion occurs over a slow WAN link.

**/roundrobin [0|1]**
Determines the order in which host records are returned when a server has multiple host records for the same name.

**0**
DNS server does not use round robin; instead, it returns the first record to every query.

**1**
DNS server rotates among the records that it returns from the top to the bottom of the list of matching records. This is the default setting.

**/rpcprotocol [0x0|0x1|0x2|0x4|0xFFFFFFFF]**
Specifies the protocol that RPC uses when making a connection from the DNS server.

**0x0**
Disables RPC for DNS.

**0x1**
Uses TCP/IP.

**0x2**
Uses named pipes.

**0x4**
Uses LPC.

**0xFFFFFFFF**
All protocols. This is the default setting.

**/scavenginginterval [0x0-0xFFFFFFFF|0x0]**
Determines whether the scavenging feature for the DNS server is enabled, and sets the number of hours (0x0-0xFFFFFFFF) between scavenging cycles. The default setting is 0x0, which disables scavenging for the DNS server. A setting greater than 0x0 enables scavenging for the server and sets the number of hours between scavenging cycles.

**/secureresponses [0|1]**
Determines whether DNS filters records that are saved in a cache.

**0**
Saves all responses to name queries to a cache. This is the default setting.

**1**
Saves only the records that belong to the same DNS subtree to a cache.

**/sendport [0x0-0xFFFFFFFF|0x0]**
Specifies the port number (0x0-0xFFFFFFFF) that DNS uses to send recursive queries to other DNS servers. The default setting is 0x0, which means the port number is randomly selected.

**/serverlevelplugindll[*DllPath*]**
Specifies the path of a custom plug-in. When *DllPath* specifies the fully qualified path name of a valid DNS server plug-in, the DNS server calls functions in the plug-in to resolve name queries that are outside the scope of all locally hosted zones. If a queried name is out of the scope of the plug-in, the DNS server performs name resolution using forwarding or recursion, as configured. If *DllPath* is not specified, the DNS server ceases to use a custom plug-in if a custom plug-in was previously configured.

**/strictfileparsing [0|1]**
Determines a DNS server's behavior when it encounters an erroneous record while loading a zone.

**0**
Continues to load even if the server encounters an erroneous record. The error is recorded in the DNS log. This is the default setting.

**1**
Stops loading and records the error in the DNS log.

**/updateoptions *RecordValue***
Prohibits dynamic updates of specified types of records. If you want more than one record type prohibited in the log, use hexadecimal addition to add the values, then enter the sum.

**0x0**
Does not restrict any record types.

**0x1**
Excludes SOA (Start of Authority) records.

**0x2**
Excludes NS (name server) records.

**0x4**
Excludes delegation NS records.

**0x8**
Excludes server host records.

**0x100**
On secure dynamic update, excludes SOA records.

**0x200**
On secure dynamic update, excludes root NS records.

**0x30F**
On standard dynamic update, excludes NS, SOA, and server host records. On secure dynamic update, excludes root NS and SOA records. Allows delegations and server host updates.

**0x400**
On secure dynamic update, excludes delegation NS records.

**0x800**
On secure dynamic update, excludes server host records.

**0x1000000**
Excludes DS records.

**0x80000000**
Disables DNS dynamic update.

**/writeauthorityns [0|1]**
Determines when the DNS server writes name server (NS) records in the Authority section of a response.

**0**
Writes NS records in the Authority section of referrals only. This setting complies with RFC 1034, Domain names—concepts and facilities, and with RFC 2181, Clarifications to the DNS Specification. This is the default setting.

**1**
Writes NS records in the Authority section of all successful authoritative responses.

**/xfrconnecttimeout [0x0-0xFFFFFFFF|0x1E]**
Determines the number of seconds (0x0-0xFFFFFFFF) a primary DNS server waits for a transfer response from its secondary. The default value is 0x1E (30 seconds). After the time-out value expires, the connection is terminated.

**Zone-Level Syntax**

**dnscmd /config** *Parameters*

**Parameters**

***Parameters***
Specify a setting, a zone name, and, optionally, a value. *Parameters* values use this syntax: *ZoneName Parameter* [*Value*]

The following *Parameters* values are documented in the remainder of this section:

- **/aging**
- **/allownsrecordsautocreation**
- **/allowupdate**
- **/forwarderslave**
- **/forwardertimeout**

- **/norefreshinterval**
- **/refreshinterval**
- **/securesecondaries**

**/aging *ZoneName***
Enables or disables scavenging in a specific zone.

**/allownsrecordsautocreation *ZoneName* [*Value*]**
Overrides the DNS server's NS record autocreation setting. NS resource records that were previously registered for this zone are not affected and so must be manually removed if not wanted.

**/allowupdate *ZoneName***
Determines whether the specified zone accepts dynamic updates.

**/forwarderslave *ZoneName***
Overrides the DNS server **/isslave** setting.

**/forwardertimeout *ZoneName***
Determines how many seconds a DNS zone waits for a forwarder to respond before trying another. This value overrides the value set at the server level.

**/norefreshinterval *ZoneName***
Sets a time interval for a zone during which no refreshes can dynamically update DNS records in a specified zone.

**/refreshinterval *ZoneName***
Sets a time interval for a zone during which refreshes can dynamically update DNS records in a specified zone.

**/securesecondaries *ZoneName***
Determines which secondary servers can receive zone updates from the master for this zone.

**Remarks**

- The zone name needs to be specified only for zone-level parameters.

**Dnscmd createbuiltindirectorypartitions**
Creates a DNS application directory partition. When DNS is installed, an application directory partition for the service is created at the forest and domain levels. Use this operation to create DNS application directory partitions that were deleted or never created. With no argument, this operation creates a built-in DNS directory partition for the domain.

**Syntax**
▸ **dnscmd** [*ServerName*] **/createbuiltindirectorypartitions** [**/forest**] [**/alldomains**]

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

**/forest**
Creates a DNS directory partition for the forest.

**/alldomains**
Creates DNS partitions for all domains in the forest.

**Dnscmd createdirectorypartition**
Creates a DNS application directory partition. When DNS is installed, an application directory partition for the service is created at the forest and domain levels. This operation creates additional DNS application directory partitions.

**Syntax**
▸ **dnscmd** [*ServerName*] **/createdirectorypartition** *PartitionFQDN*

**Parameters**
***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***PartitionFQDN***
The fully qualified domain name of the DNS application directory partition that will be created.

**Dnscmd deletedirectorypartition**
Removes an existing DNS application directory partition.

**Syntax**
- **dnscmd** [*ServerName*] **/deletedirectorypartition** *PartitionFQDN*

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***PartitionFQDN***
The fully qualified domain name of the DNS application directory partition that will be removed.

**Dnscmd directorypartitioninfo**
Lists information about a specified DNS application directory partition.

**Syntax**
- **dnscmd** [*ServerName*] **/directorypartitioninfo** *PartitionFQDN* [**/detail**]

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***PartitionFQDN***
The fully qualified domain name of the DNS application directory partition.

**/detail**
Lists all information about the records.

**Dnscmd enlistdirectorypartition**
Adds the DNS server to the specified directory partition's replica set.

**Syntax**
- **dnscmd** [*ServerName*] **/enlistdirectorypartition** *PartitionFQDN*

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***PartitionFQDN***
The fully qualified domain name of the DNS application directory partition.

**Dnscmd enumdirectorypartitions**
Lists the DNS application directory partitions for the specified server.

**Syntax**
- **dnscmd** [*ServerName*] **/enumdirectorypartitions** [**/custom**]

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

**/custom**
Lists only user-created directory partitions.

**Dnscmd enumrecords**
Lists the resource records of a specified node in a DNS zone.

**Syntax**

**dnscmd** [*ServerName*] **/enumrecords** *ZoneName NodeName*[**/type** *RRType RRData*] [**/authority**] [**/glue**] [**/additional**] [**/node** |**/child** |**/startchild***ChildName*] [**/continue** |**/detail**]

**Parameters**

***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Specifies the name of the zone to which the record belongs.

***NodeName***
Specifies the name of the node of the record.

**/type *RRType RRData***
Specifies the type of record to be listed and the type of data expected.

***RRType***
Specifies the type of record to be listed.

***RRData***
Specifies the type of data that is expected when using a certain data type.

**/authority**
Includes authoritative data.

**/glue**
Includes glue data.

**/additional**
Includes all additional information about the records listed.

**{/node | /child | /startchild *ChildName*}**
Filters or adds information to the record display.

**/node**
Lists only the records of the specified node.

**/child**
Lists only the records of a specified child.

**/startchild*ChildName***
Begins the list at the specified child.

**/continue| /detail**
Specifies how the returned data is displayed.

**/continue**
Lists only the records with their type and data.

**/detail**
Lists all information about the records.

**Sample Usage**

```
dnscmd /enumrecords test.contoso.com test /additional
```

**Dnscmd enumzones**

Lists zones that exist on the specified DNS server.

**Syntax**

**dnscmd** [*ServerName*] **/enumzones** [**/primary** | **/secondary** | **/forwarder** | **/stub** | **/cache** | **/auto-created**] [**/forward** | **/reverse** | **/ds** | **/file**] [**/domaindirectorypartition** | **/forestdirectorypartition** | **/customdirectorypartition** | **/legacydirectorypartition** | **/directorypartition** *PartitionFQDN*]

**Parameters**

***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

**/primary | /secondary | /forwarder | /stub | /cache | /auto-created**
Filters the types of zones to display.

**/primary**
Lists all zones that are either standard primary or Active Directory-integrated.

**/secondary**
Lists all standard secondary zones.

**/forwarder**
Lists zones that forward unresolved queries to another DNS server.

**/stub**
Lists all stub zones.

**/cache**
Lists only the zones that are loaded into cache.

**/auto-created**
Lists the zones that were created automatically during the DNS server installation.

**/forward | /reverse | /Ds | /file**
Specifies an additional filter of the types of zones to display.

**/forward**
Lists forward lookup zones.

**/reverse**
Lists reverse lookup zones.

**/ds**
Lists directory-integrated zones.

**/file**
Lists zones that are backed by files.

**/domaindirectorypartition**
Lists zones that are stored in the domain directory partition.

**/forestdirectorypartition**
Lists zones that are stored in the forest DNS application directory partition.

**/customdirectorypartition**
Lists all zones stored in a user-defined application directory partitions.

**/legacydirectorypartition**
Lists all zones stored in the domain directory partition.

**/directorypartition*PartitionFQDN***
Lists all zones stored in the specified directory partition.

**Remarks**

- The **enumzones** operation parameters act as filters on the list of zones. If no filters are specified, then a complete list of zones is returned. When a filter is used, only the zones that meet that filter's criteria are included in the returned list of zones.

**Example**
See Dnscmd Examples or Dnscmd Examples.

**Dnscmd info**
Displays settings from the DNS section of the registry of the specified server:
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters**

**Syntax**
**dnscmd** [*ServerName*] **/info**[*Setting*]

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***Setting***

Any setting returned by the **info** operation can be individually specified. If a setting is not specified, then all settings are returned.

**Remarks**

- This operation displays registry settings that are at the DNS server level. To display zone-level registry settings, use the Dnscmd zoneinfo operation. To see a list of settings that can be displayed with this operation, refer to the Dnscmd config operation Help.

- If you do not specify a value for the parameter when you use the Dnscmd config operation, the parameter's current value is reset to the default value.

**Example**

See Dnscmd Examples or Dnscmd Examples.

**Dnscmd nodedelete**

Deletes all records for a specified host.

**Syntax**

**dnscmd** [*ServerName*] **/nodedelete** *ZoneName NodeName* [**/tree**] [**/f**]

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Specifies the name of the zone.

***NodeName***

Specifies the host name of the node to delete.

**/tree**

Specifies to delete all of the child records.

**/f**

Executes the command without asking for confirmation.

**Example**

See Dnscmd Examples.

**Dnscmd recordadd**

Adds a record to a specified zone in a DNS server.

**Syntax**

**dnscmd** [*ServerName*] **/recordadd** *ZoneName NodeName RRType RRData*

**Parameters**

***ServerName***

Specifies the DNS server the administrator is planning to manage, represented by local computer syntax, IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Specifies the zone in which the record resides.

***NodeName***

Specifies a specific node in the zone.

***RRType***

Specifies the type of record to be added.

***RRData***
Specifies the type of data that is expected when using a certain data type.

**Note**

- Resource records added with this command are configured to be static unless the /Aging switch is used. See the examples below.

- When you add a record, make sure you use the correct data type and data format. For a list of resource record types and the appropriate data types, see Resource records reference.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /recordadd contoso.com test A 10.0.0.5

dnscmd dnssvr1.contoso.com /recordadd contoso.com test2 /Aging A 10.0.0.6

dnscmd /recordadd test.contoso.com test MX 10 mailserver.test.contoso.com
```

**Dnscmd recorddelete**
Deletes a resource record from a specified zone.

**Syntax**
**dnscmd** *ServerName* **/recorddelete** *ZoneName NodeName RRType RRData*[**/f**]

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Specifies the zone in which the record resides.

***NodeName***
Specifies the name of the host.

***RRType***
Specifies the type of record to be deleted.

***RRData***
Specifies the type of data that is expected when using a certain data type.

***/f***
Note

Executes the command without asking for confirmation.

- Because nodes can have more than one resource record, this command requires you to be very specific about the type of record that you want to delete.

  If you specify a data type and do not specify a type of resource record data, then all records with that specific data type for the specified node are deleted. For a list of resource record types and the appropriate data types, see Resource records reference.

**Sample Usage**

```
dnscmd /recorddelete test.contoso.com test MX 10 mailserver.test.contoso.com
```

**Dnscmd resetforwarders**
Selects or resets IP address(es) to which the DNS server forwards DNS queries when it cannot solve them locally.

**Syntax**
**dnscmd** [*ServerName*] **/resetforwarders** [*IPAddress* [,*IPAddress*]...][**/timeout** *TimeOut*] [**/slave**|**/noslave**]

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***IPAddress***
Lists the IP address(es) to which the DNS server forwards unresolved queries.

**/timeout *TimeOut***
Sets the number of seconds that the DNS server waits for a response from the forwarder. By default, this value is 5 seconds.

**/slave|/noslave**
Determines whether the DNS server performs its own iterative queries if the forwarder fails to resolve a query.

**/slave**
Prevents the DNS server from performing its own iterative queries if the forwarder fails to resolve a query.

**/noslave**
Allows the DNS server to perform its own iterative queries if the forwarder fails to resolve a query. This is the default setting.

**Remarks**

- By default, a DNS server performs iterative queries when it cannot resolve a query.
- Setting IP addresses by using the **resetforwarders** operation causes the DNS server to perform recursive queries to the DNS servers at the specified IP addresses. If the forwarders do not resolve the query, the DNS server can then perform its own iterative queries.
- If the **/slave** parameter is used, the DNS server does not perform its own iterative queries. This means the DNS server forwards unresolved queries only to the DNS servers in the list and does not try iterative queries if the forwarders do not resolve it. It is more efficient to set one IP address as a forwarder for a DNS server. The **resetforwarders** operation can be used for internal servers in a network to forward their unresolved queries to one DNS server that has an external connection.
- Listing a forwarder's IP address twice causes the DNS server to try to forward to that server twice.

**Sample Usage**
```
dnscmd dnssvr1.contoso.com /resetforwarders 10.0.0.1 /timeout 7 /slave

dnscmd dnssvr1.contoso.com /resetforwarders /noslave
```

**Dnscmd resetlistenaddresses**
Specifies the IP addresses on a server that listen for DNS client requests.

**Syntax**
**dnscmd** [*ServerName*] **/resetlistenaddresses** [*ListenAddress*]

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ListenAddress***
Specifies an IP address on the DNS server that listens for DNS client requests. If no listen address is specified, then all IP addresses on the server listen for client requests.

**Remarks**

- By default, all IP addresses on a DNS server listen for client DNS requests.

**Sample Usage**
```
dnscmd dnssvr1.contoso.com /resetlistenaddresses 10.0.0.1
```

**Dnscmd startscavenging**
Notifies a DNS server to attempt an immediate search for stale resource records in a specified DNS server.

**Syntax**
**dnscmd** [*ServerName*] **/startscavenging**

**Parameter**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

**Remarks**

- Successful completion of this command triggers an immediate start to a scavenge.
- Although the command to start the scavenge appears to complete successfully, the scavenge does not start unless the following preconditions are met:
  - Scavenging is enabled for both the server and the zone.
  - The zone is started.
  - The resource records have a time stamp.
- For information about how to enable scavenging for the server, see the **scavenginginterval** parameter under Server-Level Syntax in the Dnscmd config operation in this document.
- For information about how to enable scavenging for the zone, see the **aging** parameter under Zone-Level Syntax in the Dnscmd config operation in this document.
- For information about how to start a zone that is paused, see the Dnscmd zoneresume operation in this document.
- For information about how to check resource records for a time stamp, see the Dnscmd ageallrecords operation in this document.
- If the scavenge fails, no warning message displays.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /startscavenging
```

**Dnscmd statistics**

Displays or clears data for a specified DNS server.

**Syntax**

**dnscmd** [*ServerName*] **/statistics** [*StatID*] [**/clear**]

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***StatID***

Specifies which statistic or combination of statistics displays. An identification number is used to specify a statistic. If no statistic ID is specified, all statistics display.

The following table lists numbers that can be specified and the corresponding statistic that lists:

**00000001**
Time

**00000002**
Query

**00000004**
Query2

**00000008**
Recurse

**00000010**
Master

**00000020**
Secondary

**00000040**
WINS

**00000100**
Update

**00000200**
SkwanSec

**00000400**
Ds

**00010000**
Memory

**00100000**
PacketMem

**00040000**
Dbase

**00080000**
Records

**00200000**
NbstatMem

**/clear**
Resets the specified statistics counter to zero.

**Remarks**

- The **statistics** operation displays counters that begin on the DNS server when it is started or resumed.

**Examples**
See Dnscmd Examples or Dnscmd Examples.

**Dnscmd unenlistdirectorypartition**
Removes the DNS server from the specified directory partition's replica set.

**Syntax**
**dnscmd** [*ServerName*] **/unenlistdirectorypartition** *PartitionFQDN*

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***PartitionFQDN***
The fully qualified domain name of the DNS application directory partition that will be removed.

**Dnscmd writebackfiles**
Checks DNS server memory for changes and writes them to persistent storage.

**Syntax**
**dnscmd** [*ServerName*] **/writebackfiles** [*ZoneName*]

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Specifies the name of the zone to be updated.

**Remarks**

- The **writebackfiles** operation updates all dirty zones or a specified zone. A zone is dirty when there are changes in memory that have not yet been written to persistent storage. This is a server-level operation that checks all zones. One zone can be specified in this operation or the Dnscmd zonewriteback operation can be used.

**Sample Usage**
```
dnscmd dnssvr1.contoso.com /writebackfiles
```

**Dnscmd zoneadd**
Adds a zone to the DNS server.

**Syntax**
**dnscmd** [*ServerName*] **/zoneadd** *ZoneName ZoneType* [**/dp** *FQDN*| {**/domain**|**/enterprise**|**/legacy**}] **/load**

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Specifies the name of the zone.

***ZoneType***
Specifies the type of zone to create. Each type has different required parameters.

**/dsprimary**
Creates an Active Directory-integrated zone.

**/primary /file *FileName***
Creates a standard primary zone and specifies the name of the file that will store the zone information.

**/secondary *MasterIPAddress* [*MasterIPAddress*...] /file *FileName***
Creates a file-backed secondary zone. If the /file option is specified, a file is created in the **%windir%\system32\dns** directory.

**/stub *MasterIPAddress* [*MasterIPAddress*...] /file *FileName***
Creates a file-backed stub zone.

**/dsstub *MasterIPAddress* [*MasterIPAddress*...]**
Creates an Active Directory-integrated stub zone.

**/forwarder *MasterIPAddress* [*MasterIPAddress*]... /file *FileName***
Specifies that the created zone forwards unresolved queries to another DNS server.

**/dsforwarder**
Specifies that the created Active Directory-integrated zone forwards unresolved queries to another DNS server.

**/dp *FQDN* {/domain | /enterprise | /legacy}**
Specifies the directory partition on which to store the zone.

***FQDN***
Specifies fully qualified domain name of the directory partition.

**/domain**
Stores the zone on the domain directory partition.

**/enterprise**
Stores the zone on the enterprise directory partition.

**/legacy**
Stores the zone on a legacy directory partition.

***/load***

Required if the new zone is loaded from a file.

**Remarks**

- Specifying a zone type of **/forwarder** or **/dsforwarder** creates a zone that performs conditional forwarding.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zoneadd test.contoso.com /dsprimary

dnscmd dnssvr1.contoso.com /zoneadd secondtest.contoso.com /secondary 10.0.0.2

dnscmd dnssvr1.contoso.com /zoneadd secondtest.contoso.com /secondary 10.0.0.2 /file
secondtest.contoso.com.dns
```

**Dnscmd zonechangedirectorypartition**

Changes the directory partition on which the specified zone resides.

**Syntax**

**dnscmd** [*ServerName*] **/zonechangedirectorypartition** *ZoneName*] {[*NewPartitionName*] | [*ZoneType*] }

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***/ZoneName***

FQDN of the current directory partition on which the zone resides.

***NewPartitionName***

FQDN of the directory partition that the zone will be moved to.

***ZoneType***

Specifies the type of directory partition that the zone will be moved to.

**/domain**

Move the zone to the built-in domain directory partition.

**/forest**

Move the zone to the built-in forest directory partition.

**/legacy**

Move the zone to the directory partition created for pre-Active Directory domain controllers. These directory partitions are not needed when in native mode.

**Dnscmd zonedelete**

Deletes a specified zone.

**Syntax**

**dnscmd** [*ServerName*] **/zonedelete** *ZoneName* [**/dsdel**] [**/f**]

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Specifies the name of the zone to be deleted.

**/dsdel**

Deletes the zone from Active Directory.

**/f**

Executes the command without asking for confirmation.

**Example**

See Dnscmd Examples.

**Dnscmd zoneexport**

Creates a text file that lists the resource records of a specified zone.

**Syntax**

**dnscmd** [*ServerName*] **/zoneexport** *ZoneName ZoneExportFile*

**Parameters**

***ServerName***

Specifies the DNS server the administrator is planning to manage, represented by local computer syntax, IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Specifies the name of the zone.

***ZoneExportFile***

Specifies the name of the file to create.

**Remarks**

- The **zoneexport** operation creates a file of resource records for an Active Directory-integrated zone for troubleshooting purposes. By default, the file created by this operation is placed in the DNS directory, which is by default the %systemroot%/System32/Dns directory.

**Example**

See Dnscmd Examples.

**Dnscmd zoneinfo**

Displays settings from the section of the registry of the specified zone:
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters\Zones***ZoneName*

**Syntax**

**dnscmd** [*ServerName*] **/zoneinfo** *ZoneName* [*Setting*]

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Specifies the name of the zone.

***Setting***

Any setting returned with the **zoneinfo** operation can be individually specified. If a setting is not specified, then all settings are returned.

**Remarks**

- The **zoneinfo** operation displays registry settings that are at the DNS zone level found at **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters\Zones***ZoneName*.
- To display server-level registry settings, use the Dnscmd info operation.
- To see a list of settings that can be displayed with this operation, refer to the Dnscmd config operation.

**Example**

See Dnscmd Examples or Dnscmd Examples.

**Dnscmd zonepause**

Pauses the specified zone, which then ignores query requests.

**Syntax**

**dnscmd** [*ServerName*] **/zonepause** *ZoneName*

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Specifies the name of the zone to be paused.

**Remarks**

- To resume a zone and make it available after it has been paused, use the Dnscmd zoneresume operation.

**Sample Usage**
```
dnscmd dnssvr1.contoso.com /zonepause test.contoso.com
```

**Dnscmd zoneprint**
Lists the records in a zone.

**Syntax**
**dnscmd** [*ServerName*] **/zoneprint** *ZoneName*

**Parameters**
***ServerName***
Specifies the DNS server the administrator is planning to manage, represented by local computer syntax, IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Identifies the zone on which the type will be changed.

**Dnscmd zoneresettype**
Changes the type of the zone.

**Syntax**
**dnscmd** [*ServerName*] **/zoneresettype** *ZoneName ZoneType* [**/overwrite_mem** | **/overwrite_ds**]

**Parameters**
***ServerName***
Specifies the DNS server the administrator is planning to manage, represented by local computer syntax, IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Identifies the zone on which the type will be changed.

***ZoneType***
Specifies the type of zone to create. Each type has different required parameters.

**/dsprimary**
Creates an Active Directory-integrated zone.

**/primary /file *FileName***
Creates a standard primary zone.

**/secondary *MasterIPAddress* [,*MasterIPAddress*...]**
Creates a standard secondary zone.

**/stub *MasterIPAddress* [,*MasterIPAddress*...] /file *FileName***
Creates a file-backed stub zone.

**/dsstub *MasterIPAddress* [,*MasterIPAddress*...]**
Creates an Active Directory-integrated stub zone.

**/forwarder *MasterIPAddress* [,*MasterIPAddress*]... /file *FileName***
Specifies that the created zone forwards unresolved queries to another DNS server.

**/dsforwarder**
Specifies that the created Active Directory-integrated zone forwards unresolved queries to another DNS server.

**/overwrite_mem | /overwrite_ds**
Specifies how to overwrite existing data.

**/overwrite_mem**
Overwrites DNS data from data in Active Directory.

**/overwrite_ds**
Overwrites existing data in Active Directory.

**Remarks**

- Setting the zone type as **/dsforwarder** creates a zone that performs conditional forwarding.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zoneresettype test.contoso.com /primary /file test.contoso.com.dns

dnscmd dnssvr1.contoso.com /zoneresettype second.contoso.com /secondary 10.0.0.2
```

**Dnscmd zonerefresh**

Forces a secondary DNS zone to update from the master.

**Syntax**

**dnscmd** *ServerName* **/zonerefresh** *ZoneName*

**Parameters**

***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Specifies the name of the zone to be refreshed.

**Remarks**

- The **zonerefresh** operation forces a check of the version number in the master's SOA record. If the version number on the master is higher than the secondary's version number, then a zone transfer is initiated, updating the secondary server. If the version number is the same, no zone transfer occurs.

- The forced check occurs by default every 15 minutes. To change the default, use the **dnscmd config refreshinterval** operation.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zonerefresh test.contoso.com
```

**Dnscmd zonereload**

Copies zone information from its source.

**Syntax**

**dnscmd** *ServerName* **/zonereload** *ZoneName*

**Parameters**

***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Specifies the name of the zone to be reloaded.

**Remarks**

- If the zone is Active Directory-integrated, then it reloads from Active Directory.

- If the zone is a standard file-backed zone, then it reloads from a file.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zonereload test.contoso.com
```

**Dnscmd zoneresetmasters**

Resets the IP addresses of the master server that provides zone transfer information to a secondary zone.

**Syntax**

**dnscmd** *ServerName* **/zoneresetmasters** *ZoneName* [**/local**] [*ServerIPs*]

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Specifies the name of the zone to be reloaded.

**/local**

Sets a local master list. Used for Active Directory-integrated zones.

***ServerIPs***

The IP addresses of the master server(s) of the secondary zone. Enter this value as *IPAddress* [*IPAddress*]...

**Remarks**

- This value is originally set when the secondary zone is created. Use the **zoneresetmasters** operation on the secondary server. It has no effect if it is set on the master DNS server.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zoneresetmasters test.contoso.com 10.0.0.1

dnscmd dnssvr1.contoso.com /zoneresetmasters test.contoso.com /local
```

**Dnscmd zoneresetscavengeservers**

Changes the IP address(es) of the server(s) that can scavenge the specified zone.

**Syntax**

**dnscmd** [*ServerName*] **/zoneresetscavengeservers** *ZoneName* [*ServerIPs*]

**Parameters**

***ServerName***

Specifies the DNS server the administrator is planning to manage, represented by local computer syntax, IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Identifies the zone to scavenge.

***ServerIPs***

Lists the IP address(es) of the server(s) that can perform the scavenge. If this parameter is omitted, then all servers hosting this zone can scavenge it.

**Remarks**

- By default, all servers hosting a zone can scavenge that zone.
- If a zone is hosted on more than one DNS server, this operation can be used to reduce the number of times a zone is scavenged.
- Scavenging must be enabled on the DNS server and zone affected by this operation.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zoneresetscavengeservers test.contoso.com 10.0.0.1 10.0.0.2
```

**Dnscmd zoneresetsecondaries**

Specifies a list of IP addresses to which a master server responds when asked for a zone transfer.

**Syntax**

**dnscmd** [*ServerName*] **/zoneresetsecondaries** *ZoneName* {**/noxfr** | **/nonsecure** | **/securens** | **/securelist** *SecurityIPAddresses*} {**/nonotify** | **/notify** | **/notifylist** *NotifyIPAddresses*}

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Specifies the name of the zone to have its secondaries reset.

**/noxfr | /nonsecure | /securens | /securelist *SecurityIPAddresses***
Specifies whether only some or all of the secondary servers requesting an update get one.

**/noxfr**
Specifies that no zone transfers are allowed.

**/nonsecure**
Specifies that all zone transfer requests are granted.

**/securens**
Specifies that only the server listed in the NS record for the zone is granted a transfer.

**/securelist**
Specifies that zone transfers are granted only to the list of servers. This parameter must be followed by an IP address or addresses that the master server uses.

***SecurityIPAddresses***
Lists the IP addresses that receive zone transfers from the master server. This parameter is used only with the **/securelist** parameter.

**/nonotify | /notify | /notifylist *NotifyIPAddresses***
Specifies a change notification is sent only to certain secondary servers.

**/nonotify**
Specifies that no change notifications are sent to secondary servers.

**/notify**
Specifies that change notifications are sent to all secondary servers.

**/notifylist**
Specifies that change notifications are sent to only the list of servers. This command must be followed by an IP address or addresses that the master server uses.

***NotifyIPAddresses***
Specifies the IP address(es) of the secondary server(s) to which change notifications are sent. This list is used only with the **/notifylist** parameter.

**Remarks**

- Use the **zoneresetsecondaries** operation on the master server to set how it responds to zone transfer requests from secondary servers.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zoneresetsecondaries test.contoso.com /noxfr /nonotify

dnscmd dnssvr1.contoso.com /zoneresetsecondaries test.contoso.com /securelist 11.0.0.2
```

**Dnscmd zoneresume**
Starts a specified zone that was previously paused.

**Syntax**
**dnscmd** *ServerName* **/zoneresume** *ZoneName*

**Parameters**
***ServerName***
Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***
Specifies the name of the zone to resume.

**Remarks**

- This operation is used to undo the Dnscmd zonepause operation.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zoneresume test.contoso.com
```

**Dnscmd zoneupdatefromds**

Updates the specified Active Directory-integrated zone from Active Directory.

**Syntax**

▸ **dnscmd** *ServerName* **/zoneupdatefromds** *ZoneName*

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Specifies the name of the zone to update.

**Remarks**

- Active Directory-integrated zones perform this update by default every 180 seconds (3 minutes). To change this parameter, use the **dnscmd /config /dspollinginterval** operation.

  To view the current value, use the **dnscmd /info /dspollinginterval** operation. The value is displayed in seconds.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zoneupdatefromds
```

**Dnscmd zonewriteback**

Checks DNS server memory for changes relevant to a specified zone and writes them to persistent storage.

**Syntax**

▸ **dnscmd** *ServerName* **/zonewriteback** *ZoneName*

**Parameters**

***ServerName***

Specifies the DNS server the administrator plans to manage, represented by IP address, FQDN, or Host name. If omitted, the local server is used.

***ZoneName***

Specifies the name of the zone to update.

**Remarks**

- This is a zone-level operation. All zones on a DNS server can be updated with the Dnscmd writebackfiles operation.

**Sample Usage**

```
dnscmd dnssvr1.contoso.com /zonewriteback test.contoso.com
```

**See Also**

**Concepts**

Dnscmd Overview
Dnscmd Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Dnslint Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dnslint.exe: DNS Lint**
This command-line tool allows you to verify Domain Name System (DNS) records for a specified domain name. Use DNSLint to help diagnose potential causes of incorrect delegation and other common DNS problems.

**Concepts**
**Incorrect delegation**
System Administrators use DNS delegations to assign the authority for a DNS subdomain to a set of DNS servers. These delegations help administrators implement a DNS infrastructure that is scalable, flexible, and easy to administer.

Incorrect delegation is a common problem with DNS structures. This can occur when the authority for a DNS subdomain has been delegated to a particular DNS server that either does not exist, or does not have authority over the subdomain.

Incorrect delegation problems are compounded when other DNS issues are present in a DNS infrastructure. For example, if DNS domain information is not identical on all authoritative DNS servers, data returned from the servers may differ, leading to intermittent name resolution difficulties. In another scenario, if a Mail Exchange (MX) record exists but the corresponding glue record is missing, then e-mail delivery to that domain may be erratic.

Prior to the development of DNSLint, the nslookup utility was frequently used to diagnose and troubleshoot these types of DNS issues. Nslookup allows users to manually traverse a DNS infrastructure, inspecting the various DNS records. However, DNSLint offers quicker verification that record-level data is correct.

For more information about DNS, see DNS or Concepts in Microsoft Management Console Help.

For more information about troubleshooting DNS, see the chapter Troubleshooting DNS in the Networking Guide of the Windows Server 2003 Resource Kit.

**Common Active Directory DNS Problems**
Verifying a particular set of DNS records on multiple DNS servers can help diagnose and fix problems caused by missing or incorrect DNS records.

For example, when clients are experiencing problems logging on to the domain, verifying that the SRV records that clients use to find LDAP and Kerberos servers are available and accurate, can help determine if DNS is a cause of the problem.

Another scenario example is where you receive reports that customers are having problems accessing your Web site on the Internet. It would be nice to have a tool that quickly checks all the DNS records that are involved with the Web farm on all of the DNS servers that are supposed to have these records. You could quickly determine if there are missing or incorrect DNS records that may be related to the problem.

In a third example, you could be experiencing problems with e-mail delivery. You can send e-mail, but you are not receiving e-mail. Name resolution could be the problem? To confirm this theory or eliminate it as a possibility, you need to check all of the DNS records on all of the DNS servers that are used to resolve the IP address of the e-mail server.

**System Requirements**
DNSLint runs on a source computer and acts on a target computer. The target computer can be the same computer as the source computer, or it can be a different computer.

**Source Computer Requirements**

- Windows 2000, Windows XP, or Windows Server 2003

**Target Computer Requirements**

- The target computer's domain must be registered with an accredited domain registrar, or running DNS in a private namespace.

**File Required**

- Dnslint.exe

**See Also**
**Concepts**
Dnslint Remarks
Dnslint Syntax
Dnslint Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Dnslint Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNSLint Remarks**
**DNSLint Information**
The DNSLint tool streamlines the process of troubleshooting and diagnosing DNS issues on the Internet and private networks. Typically, it traverses a DNS network in under a minute, and then generates a report in HTML format.

DNSLint verifies that all authoritative DNS servers for a domain are responding to DNS queries, and have synchronized zone data. It verifies domain Host (A) records, Name Server (NS) records, Mail Exchange (MX) records, and the glue records associated with the delegation

**DNSLint and E-Mail**
Use the **/c** (connectivity test) parameter to have DNSLint test well-known e-mail ports on all e-mail servers that it finds while inspecting DNS servers for the specified domain name. The tool supports the following protocols:

- The Simple Mail Transfer Protocol (SMTP)
- Post Office Protocol (POP) version 3
- Internet Message Access Protocol (IMAP) version 4

By default, when the **/c** parameter is specified, DNSLint attempts to connect to all three ports on each e-mail server it finds. (for example, TCP port 25 for SMTP, TCP port 110 for POP3, and TCP port 143 for IMAP).

DNSLint reports the following port status: listening, not listening, or no response. In the case of a listening port, DNSLint returns any response the port might have. The following example demonstrates the response DNSLint might return for a listening SMTP port :

```
220 mailsrv.reskit.com Microsoft ESMTP MAIL Service, Version: 5.0.2195.3705 ready at Mon, 13 May 2002
17:08:36 -0700
```

**Using DNSLint /ql to diagnose Active Directory DNS**
DNSLint's /ql (Query List) option provides this functionality. DNSLint reads instructions from the text file specified using the /ql option. Once it has verified that the file is a valid DNSLint input file, it runs the queries that are specified within the file and reports the results in an easy to read HTML report (and optionally in a text report). This input file allows administrators to customize which DNS servers to query and exactly which DNS records to look for on each server. The format of the input file is as follows:

```
DNSLint
[dns~server] 169.254.46.138
www.fourthcoffee.com,a,r
169.254.197.1,ptr,r
[dns~server] 169.254.46.200
fourthcoffee.com,cname,r
fourthcoffee.com,mx,r
_kerberos._tcp.dc._msdcs.fourthcoffee.com,srv,r
```

The file must start with the word "dnslint" at the top of it. This is the first thing DNSLint looks for when the input file is opened. If it is not the first word read when the file is opened, the specified input file is rejected and an error is generated.

```
[dns~server] 169.254.46.138
```

This line specifies the IP address of a DNS server to send queries to. [dns~server] must be specified followed by a valid IP address. If either of these two components is missing, an error is generated and the specified input file is rejected.

Subsequent lines indicate the queries to send to the specified DNS server:

```
www.fourthcoffee.com,a,r
10.1.10.1,ptr,r
```

Format of the queries:

The first field in the line is the name to query. For example www.fourthcoffee.com. The name is then immediately followed by a comma. No spaces are allowed on either side of the comma.

The second field follows the comma immediately after the name to query. The second field is the type of record to query for. Valid types are as follows:

- a = Host
- ptr = Pointer
- cname = Alias
- mx = Mail Exchange
- srv = Service Location

The type of record is then immediately followed by a comma. No spaces are allowed on either side of the comma.

The third field is the type of query. This field immediately follows the comma after the type of record. Valid query types are as follows:

- r = recursive
- i = iterative

Nothing else is required to follow the third field. All three fields are required, and no spaces are allowed anywhere within the query line. A fourth field is optional. Appending ",tcp" to the third field will make DNSLint send the specified query using the TCP protocol instead of the default UDP protocol. Again, no spaces are allowed and nothing should follow this field if it is used.

**See Also**
**Concepts**
Dnslint Overview
Dnslint Syntax
Dnslint Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Dnslint Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNSLint Syntax**

DNSLint uses the following syntax:

**dnslint /d** *domain_name* | **/ad** | **/ql** {*text_file* | **autocreate**} [**/v**] [**/r** *report_name*] [**/y**] [**/no_open**] [**/s** *alternate_DNS_server_ip_address*] [**/c** *SMTP*|*POP*|*IMAP*] [**/t**] [**/test_tcp**]

**Parameters**

**/d *domain***

Specifies the domain to test. This is a required parameter.

**/ad**

Tests the DNS records used by Active Directory for replication. This is a required parameter. The **/s** parameter is required with this parameter.

**/ql {*text_file* | autocreate}**

Sends DNS queries based on input from a text file containing a list of DNS names. The command **dnslint /ql autocreate** creates an example textfile. The **/ql** parameter supports A, PTR, CNAME, SRV, and MX DNS record types. This is a required parameter. The **/ql** parameter will accept either a text file name or the **autocreate** option, not both.

**/v**

Turns on verbose output.

**/r *report_name***

Specifies a name for the DNSLint report file. The default report name is DNSLint Report.htm.

**Note**

- DNSLint always appends an .htm extension to the file, even if you specify an extension. For example, if you specify *report_name*=my_file.rpt, the resulting file will be called my_file.rpt.htm.

**/y**

Overwrites any existing report file without prompting the user.

**/no_open**

Prevents the DNSLint report from opening after it is generated.

**/s *alternate_DNS_server_ip_address***

This parameter supplies an alternate DNS server, forcing DNSLint to bypass InterNIC. This parameter is useful for testing domain structures that are not availible from the InterNIC database.

**/c *SMTP|POP|IMAP***

Used to perform connectivity tests on well-known e-mail ports when e-mail servers are found while inspecting DNS servers. See "Remarks" for more information.

**/t**

Used to request a plain text file with the same name as the DNSLint report file. (The default filename is DNLint Report.htm, unless the **/r** parameter is used to specify a different name.)

**/test_tcp**

Used to test TCP port 53 as well as UDP port 53. Unless this parameter is specified, DNSLint test the UDP port only.

**Notes**

- You must always specify one of the following parameters: **/d**, **/ad**, and **/ql**. However, they cannot be used together.
- The **/c** parameter cannot be used with the **/ad** or **/ql** parameters.
- When using the **/ad** parameter, **/s** must also be used.

- The **/test_tcp** parameter cannot be used with the **/ql** parameter.

**Sample Usage**

```
dnslint /d fourthcoffee.com

dnslint /ad /s 10.193.36.201

dnslint /ql my_dns_list.txt

dnslint /ql autocreate

dnslint /v /d fourthcoffee.com

dnslint /r newfile /d fourthcoffee.com

dnslint /y /d fourthcoffee.com

dnslint /no_open /d fourthcoffee.com

dnslint /s 10.193.36.201 /d labs.fourthcoffee.com
```

**See Also**

**Concepts**

Dnslint Overview
Dnslint Remarks
Dnslint Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Dnslint Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNSLint Examples**

**Example 1: Sample DNSLint Output with No Errors**

The following command instructs DNSLint to traverse the woodgrovebank.com domain, and return all relevant data about its DNS structure. This domain has no structural errors.

```
dnslint /d woodgrovebank.com
```

The link below shows a sample DNSLint report generated in response to this command.

DNSLint Report Example 1

**Example 2: Sample DNSLint Output with Errors**

The following command instructs DNSLint to traverse the lab.fourthcoffee.com domain. The command uses the /s parameter because lab.fourthcoffee.com is a private domain. As it turns out, this domain contains errors in its structure.

```
dnslint /d lab.fourthcoffee.com /s 10.46.138.20
```

The link below shows a sample DNSLint report generated in response to this command.

DNSLint Report Example 2

**See Also**

**Concepts**

Dnslint Overview
Dnslint Remarks
Dnslint Syntax
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Getsid Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Getsid.exe: Get Security ID**
This command-line tool compares the user security identifiers (SIDs) of two accounts.

You can use it to compare account SIDs between a primary domain controller and backup domain controller when you suspect user database corruption.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
For more information about security in Windows XP and Windows Server 2003, see *Security Model* or *Security and Active Directory* under Security in Help and Support Services.

**System Requirements**
The following are the system requirements for GetSID:

- Windows XP Professional or Windows Server 2003
- Administrative rights

**File Required**

- Getsid.exe

**See Also**
**Concepts**
Getsid Syntax
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc) Overview
Acldiag Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Getsid Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**GetSID Syntax**

GetSID uses the following syntax:

**getsid** *\\Server1 Account \\Server2 Account*

**Parameters**

***\\Server1***

Specifies the name of the server on which the accounts are found. Enter a netbios or FQDN. *Account* is the account SID.

***Account***

Specifies the name of a user or group account. Getsid will display the SID of this account. If the account name has a space in it, you must use quotation marks.

***\\Server2***

Specifies the name of the server on which the accounts are found. Enter a netbios or FQDN.

**Sample Usage**

- getsid \\reskitdc2 administrator \\reskitdc1 administrator
- getsid \\reskitdc2 "domain admins" \\reskitdc1 "domain admins"

**See Also**

**Concepts**

Getsid Overview
Alphabetical List of Tools
Search Overview
Replmon Overview
Repadmin Overview
Movetree.exe
Ldp Overview
Dsastat Overview
Clonepr Overview
ADSI Edit (adsiedit.msc) Overview
Acldiag Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Addiag.exe

# Iasparse Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Iasparse.exe: IAS Parse Tool**
This command-line tool parses Internet Authentication Service (IAS) and remote access server logs and converts them into a readable format.

IASParse is useful in these situations:

- When calculating billing information.
- For other diagnostic purposes, such as finding out the Remote Access Dial-In User Service (RADIUS) attributes received and sent.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
IAS can create a log file based on the authentication and accounting requests received from the network access servers when using RADIUS authentication. Remote access server can also log such information when using Windows Server 2003 authentication.

Setting up and using such logs helps simplify administration of the service. The logs can be useful for tracking accounting information, such as logon and logoff records for billing purposes. IASParse presents the information in them in a user-friendly way.

**System Requirements**
The following are the system requirements for this tool:

- Windows XP Professional or Windows Server 2003 operating systems

**File Required**

- Iasparse.exe

**See Also**
**Concepts**
Iasparse Remarks
Iasparse Syntax
Iasparse Examples
Alphabetical List of Tools
Iasparse Overview
Bindiff Overview

# Iasparse Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IASParse Remarks**

**IAS Logs**
IAS records the authentication and accounting information that it receives in a log file, which can be in either of two formats:

- **IAS Format** In the Microsoft® Windows® NT version 4.0 version of IAS, the log was formatted using a method by which attributes were logged as attribute-value pairs. This formatting is supported in IAS in Windows 2000, and the logs that use this format are referred to as IAS-formatted log files.

- **ODBC-Compatible Format** The ODBC-compatible format (Open Database Connectivity format) for log files is available starting in Windows 2000. Unlike IAS-formatted log files, ODBC-compatible log files present the data in a standard sequence and using a structure that is identical, regardless of the network access server (NAS) sending the data. This consistent sequence and structure helps simplify accounting and authentication record-keeping. Data can be easily exported to an ODBC database, which supports filtering and analysis by a variety of third-party programs.

The user interface of IAS allows the user to choose the format of the log file. This format can also be changed dynamically, which can produce intermixed records of both formats in the log file. Thus there is a need for a tool that can understand the format of the records logged into the log file dynamically, parse it, and produce an output that helps the user to understand the information logged into the file.

IAS uses the information in the file %windir%\system32\ias\dnary.mdb to find out the ODBC log sequence and also whether to log a particular attribute while logging in IAS format. The tool also reads the information from the same file before parsing the log file. Hence, hard coding these values inside the tool is eliminated.

IAS can also send the output to a pipe instead of a file. IASParse uses this feature to read the input from the pipe and send the output directly to the screen. In effect, this causes the tool to dynamically display the parsed output. This feature can be enabled by choosing the **-p** option. For this functionality to work, the user must configure the IAS to output to a pipe. This can be done by setting the log file directory from the IAS snap-in to \\.\pipe.

**See Also**
**Concepts**
Iasparse Overview
Iasparse Syntax
Iasparse Examples
Alphabetical List of Tools
Bindiff Overview

# Iasparse Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IASParse Examples**

**Example 1: IAS Format**

The following command starts IASparse and parses the iaslog.log file. It then records the authentication and accounting information that it receives in a log file in IAS format.

```
iasparse -f:iaslog.log
```

```
The line logged into the file: 172.31.230.187,rajeshp,02/09/2000,23:18:00,IAS,RA
JESHP2,6,2,7,1,5,11,61,5,64,1,65,1,31,172.31.225.108,66,172.31.225.108,4108,172.
31.230.187,4116,9,4128,rajeshp3,4147,311,4148,MSRASV5.00,4129,RAJESHP2\rajeshp,4136,1,4
142,0

 NAS_IP_Address      :    172.31.230.187
 User_Name           :     rajeshp
 Record_Date         :     02/09/2000
 Record_Time         :     23:18:00
 Service_Name        :     IAS
 Computer_Name       :     RAJESHP2
 Service-Type        :     Framed
 Framed-Protocol     :     PPP
 NAS-Port            :     11
 NAS-Port-Type       :     Virtual
 Tunnel-Type         :     PPTP
 Tunnel-Medium-Type  :     IP
 Calling-Station-Id  :     172.31.225.108
 Tunnel-Client-Endpt :     172.31.225.108
 Client-IP-Address   :     172.31.230.187
 Client-Vendor       :    VENDOR
 Client-Friendly-Name:     rajeshp3
 MS-RAS-Vendor       :    Microsoft
 MS-RAS-Version      :     MSRASV5.00
 SAM-Account-Name    :     RAJESHP2\rajeshp
 Packet-Type         :     Access-Request
 Reason-Code         :    The operation completed successfully.
```

The line logged into the file: 172.31.230.187,rajeshp,02/09/2000,23:18:00,IAS,RA
JESHP2,4130,RAJESHP2\rajeshp,4129,RAJESHP2\rajeshp,4128,rajeshp3,4116,9,4108,172.31.230.187,4136,3,4142,16

```
 NAS_IP_Address          :      172.31.230.187
 User_Name               :      rajeshp
 Record_Date             :      02/09/2000
 Record_Time             :      23:18:00
 Service_Name            :      IAS
 Computer_Name           :      RAJESHP2
 Fully-Qualifed-User-Name:      RAJESHP2\rajeshp
 SAM-Account-Name        :      RAJESHP2\rajeshp
 Client-Friendly-Name    :      rajeshp3
 Client-Vendor           :      VENDOR
 Client-IP-Address       :      172.31.230.187
 Packet-Type             :      Access-Reject
 Reason-Code             :      Authentication failure: unknown user name or bad
password
```

**Example 2: ODBC-Compatible Format**

The following command starts IASparse and parses the iaslog.log file. It then records the authentication and accounting information that it receives in a log file in ODBC-compatible format.

```
IASParse
```

The line logged into the file: "RAJESHP3","IAS",05/17/1999,15:31:11,4,"rajeshp"

,,"172.31.225.108",,"20.20.20.3",,"172.31.230.187",11,9,"172.31.230.187","rajes p3",4294967295,5,,1,2,,,0,"311 1 172.31.230.187 05/17/1999 17:13:09 3",,1200,,,
,0,318,370,"14",1,0,13,13,,"37","0x00000001",,1,1,"172.31.225.108",,,,,,,"MSRASV5.00",311,,,,

```
Computer_Name           :       "RAJESHP3"
Service_Name            :       "IAS"
Record_Date             :       05/17/1999
Record_Time             :       15:31:11
Packet-Type             :       Accounting-Request
User-Name               :       "rajeshp"
Fully-Qualifed-User-Name:
Called-Station-Id       :
Calling-Station-Id      :       "172.31.225.108"
Callback-Number         :
Framed-IP-Address       :       "20.20.20.3"
NAS-Identifier       :
NAS-IP-Address          :       "172.31.230.187"
NAS-Port                :       11
Client-Vendor           :       VENDOR
Client-IP-Address       :       "172.31.230.187"
Client-Friendly-Name    :       "rajeshp3"
Port-Limit              :       4294967295
NAS-Port-Type           :       Virtual
Connect-Info            :
Framed-Protocol         :       PPP
Service-Type            :       Framed
Authentication-Type     :       Unknown
NP-Policy-Name          :
Reason-Code             :       The operation completed successfully.
Class                   :       "311 1 172.31.230.187 05/17/1999 17:13:09 3"
Session-Timeout         :
Idle-Timeout            :       1200
Termination-Action      :
EAP-Friendly-Name       :
Acct-Status-Type        :       Start
Acct-Delay-Time         :       0
Acct-Input-Octets       :       318
Acct-Output-Octets      :       370
Acct-Session-Id         :       "14"
Acct-Authentic          :       Radius
Acct-Session-Time       :       0
Acct-Input-Packets      :       13
Acct-Output-Packets     :       13
Acct-Terminate-Cause    :       Unknown
Acct-Multi-Ssn-Id       :       "37"
Acct-Link-Count         :       "0x00000001"
Acct-Interim-Interval   :
Tunnel-Type             :       PPTP
Tunnel-Medium-Type      :       IP
Tunnel-Client-Endpt     :       "172.31.225.108"
Tunnel-Server-Endpt     :
Acct-Tunnel-Conn        :
Tunnel-Pvt-Group-ID     :
Tunnel-Assignment-ID    :
Tunnel-Preference       :
MS-Acct-Auth-Type       :
MS-Acct-EAP-Type        :
MS-RAS-Version          :       "MSRASV5.00"
MS-RAS-Vendor           :       Microsoft
MS-CHAP-Error           :
MS-CHAP-Domain          :
MS-MPPE-Encryption-Types:
MS-MPPE-Encryption-Policy:
```

**Example 3: Displaying the description associated with attributes in the log file.**

The following command displays the description associated with attributes in the log file.

```
-v
```

```
For example, for this attribute:

Client-Vendor       : VENDOR

it shows this description:

Manufacturer of RADIUS proxy or NAS. (IAS only)
```

**See Also**
**Concepts**
Iasparse Overview
Iasparse Remarks
Iasparse Syntax
Alphabetical List of Tools
Bindiff Overview

# Iasparse Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IASParse Syntax**

**IASParse** [**-f:***filename*] [**-p**] [**-v**] [**-?**]

**Parameters**

**-f: *filename***

Parses the file *filename*. By default, IASParse parses the file %windir%\system32\logfiles\iaslog.log

**-p**

Sends output directly to display. Sets the log file directory in the IAS snap-in to \\.\pipe.

**-v**

Displays the description associated with attributes in the log file.

**-?**

Displays a syntax screen at the command prompt.

**See Also**

**Concepts**

Iasparse Overview
Iasparse Remarks
Iasparse Examples
Alphabetical List of Tools
Bindiff Overview

# Netcap Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netcap.exe: Network Monitor Capture Utility**
This command-line tool allows a system administrator to monitor packets on a LAN and write the information to a log file. NetCap uses the Network Monitor Driver to sniff packets on local network segments.

### Notes

- You must run NetCap from the command window.
- If the Network Monitor Driver is not installed, NetCap installs it the first time the tool is run. To remove the driver, use **netcap /remove**.

**Corresponding UI**
This tool provides a command-line interface to some of the capture functionality of Netmon.

**Concepts**
NetCap captures frames directly from the network traffic data stream so they can be examined. You can use it to create capture files for support personnel.

Frames are packages of information transmitted as a single unit over a network. Every frame follows the same basic organization and contains the following:

- Control information such as synchronizing characters
- Source and destination addresses
- Protocol information
- An error-checking value
- A variable amount of data

**System Requirements**
NetCap requires one of the following operating systems:

- Windows Server 2003
- Windows XP Professional
- Windows 2000

**File Required**

- Netcap.exe

**See Also**
**Concepts**
Netcap Remarks
Netcap Syntax
Netcap Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview

Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Netcap Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetCap Remarks**

How the Network Monitor Driver Works

Network Monitor tracks the network data stream, which consists of all of the information transferred over a network at any given time. Before transmission, the networking software divides this information into smaller segments, called frames. Each frame contains the following information:

- The source address of the computer that sent the message.
- The destination address of the computer that receives the frame.
- Header information of each protocol used to send the frame.
- The data (or a portion of it) being sent to the destination computer.

All frames on a network segment pass through every computer connected to that segment. However, the network card typically passes only the frames addressed to it as the destination computer, to the networking software. The Windows XP Professional version of Network Monitor can copy frames originating from or sent to the local computer to a temporary capture file. The process by which Network Monitor copies frames is referred to as data capture.

The amount of information that the Network Monitor Driver can capture is limited only by disk space availability, up to one gigabyte. However, you usually need to capture only a small subset of the frames traveling on your network. To isolate a subset of frames, you can design a capture filter, which functions like a database query to isolate the information that you specify. You can filter frames on the basis of source and destination addresses, protocols, protocol properties, and pattern offset.

If you want a running capture to respond to specific conditions as soon as the Network Monitor detects them, you can design a capture trigger. When the Network Monitor detects a particular set of conditions on the network, this capture trigger performs a specified action, such as starting an executable file.

**Note**

- When used with a network adapter that supports promiscuous mode, the full version of Network Monitor (available with Microsoft Systems Management Server version 2.0) can capture all the frames it detects.

The Network Data Stream

Network Monitor monitors the network data stream, which consists of all information transferred over a network at any given time. Prior to transmission, this information is divided by the network software into smaller pieces, called frames or packets.

Frames, whether broadcast, multicast, or directed, are made up of several different pieces that can be analyzed separately. Some of these pieces contain data that Network Monitor can use to troubleshoot networking problems. For example, by examining the destination address, it can be determined whether the frame was a broadcast frame, indicating all hosts had to receive and process this frame, or a directed frame sent to a specific host. By analyzing frames, you can determine the exact cause of the frame, which helps determine whether the service generating these types of frames can be optimized.

**See Also**

**Concepts**

Netcap Overview
Netcap Syntax
Netcap Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview

Netdiag Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Netcap Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetCap Examples**

Example 1: Capture network packets for 3 minutes

In this example, you monitor your network port for three minutes and capture the information to the C:\Temp folder. Type the following at the command prompt:

```
netcap /c:c:\temp /l:00:03:00
```

NetCap displays the following result:

```
Total frames captured: 229

Capture file name: c:\temp\7D16111E39FA.cap
```

Example 2: Monitor with a trigger

In this example, you want to monitor with a trigger defining when to stop. You want to look for a certain pattern, and when found, keep listening until the trigger is halfway into the buffer. You have previously defined a filter restricting the monitoring to a suspicious host. Type the following at the command prompt:

```
NetCap /B:20 /N:2 /T BP 50 0a ff1f /F:d:\IPFilter.CF
```

NetCap produces the following output:

```
Total frames captured: 85220

Capture file name: c:\temp\7D234C233A3B.cap
```

Example 3: Remove the Network Monitor Driver

In this example, you remove the Network Monitor Driver that NetCap has installed. Type the following at the command prompt:

```
netcap /remove
```

NetCap produces the following output:

```
The NetCap.exe instance of the Network Monitor driver successfully removed.
```

**See Also**

**Concepts**

Netcap Overview
Netcap Remarks
Netcap Syntax
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Netcap Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetCap Syntax**

NetCap uses the following syntax:

**netcap** [**/b:***Number*] [**/t** *Type Buffer HexadecimalOffset HexadecimalPattern*] [**/f:** *FilterFile*] [**/c:** *CapturePath*] [**/n:** *Number*] [**/l:** *HH:MM:SS*] [**/tcf:** *FolderName*] [**/remove**]

**Parameters**

**/b: *Number***

Specifies the buffer size in megabytes. *Number* can be from 1 to 1000. The default is 1 (MB).

**/t *Type Buffer HexadecimalOffset HexadecimalPattern***

Specifies the use of a trigger to determine when to stop capturing. If the trigger is omitted, NetCap captures until the buffer is full, then stops. The **/t N** option continues to capture until the spacebar is pressed, using the buffer as a queue. If the buffer becomes full Netcap will overwrite the oldest entries.

***Type***

**B** = Buffer

**P = Pattern**

**BP = Buffer then pattern**

**PB = Pattern then buffer**

**N = No trigger**

***Buffer***

**Percent buffer size. Options are 25, 50, 75, and 100. It is used with B, BP, PB, but not P.**

***HexadecimalOffset***

**Hexadecimal offset from start of frame. Used with P, BP, PB, but not B.**

***HexadecimalPattern***

**Hexadecimal pattern to match. Used with P, BP, PB, but not B. The pattern must be an even number of hexadecimal digits.**

**/f: *FilterFile***

**Specifies a filter to use during the capture. A filter file is created using Network Monitor in versions 2.x and after. A filter file has a .cf extension.**

**/c: *CapturePath***

**Specifies a location where Netcap will move the temporary capture files. This can be any valid local or remote path. If /c is not specified, the capture path remains in the default temporary capture folder.**

**/n: *Number***

**Specifies which network adapter to capture from, based on the network adapter index number. Run Netcap with the /? parameter to assign index numbers to the network adapters on the computer.**

**/l: *HH:MM:SS***

**Specifies how long the capture will last in hours, minutes, and seconds (maximum = 99:99:99).**

**Note**

- **This option overrides the default 100 percent trigger, unless /t *Type* is also specified.**

**/tcf: *FolderName***

**Sets the location where the temporary capture files are stored. The path must be on a fixed local hard drive. Once set, you only need to use the option again to change the directory.**

**/remove**

**Removes the NetCap instance of the Network Monitor Driver.**

**Examples**

**See Netcap Examples**

**See Also**

**Concepts**

**Netcap Overview**
**Netcap Remarks**
**Netcap Examples**
**Alphabetical List of Tools**
**Spcheck Overview**
**Nltest Overview**
**Netdom Overview**
**Netdiag Overview**
**Httpcfg Overview**
**Dnslint Overview**
**Dnscmd Overview**
**Dhcploc Overview**
**Dcdiag Overview**
**Browstat Overview**

# Netdiag Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netdiag.exe: Network Connectivity Tester**
This command-line diagnostic tool helps to isolate networking and connectivity problems by performing a series of tests to determine the state of your network client. These tests and the key network status information they expose give network administrators and support personnel a more direct means of identifying and isolating network problems. Moreover, because this tool does not require parameters or switches to be specified, support personnel and network administrators can focus on analyzing the output rather than on training users how to use the tool.

### Notes

- You must run NetDiag from the command window.
- TCP/IP must be bound to one or more adapters before using this tool. In most cases, this binding already exists.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
The TCP/IP protocol stack is the primary communications protocol used in networks today. To access the Internet, you need to be using TCP/IP.

This tool provides diagnostic features that are useful for troubleshooting Windows Server 2003 operating system configurations. However, because this tool has been designed primarily for system administrators and support personnel, its output may be difficult to analyze. In this case, you should review the appropriate troubleshooting sections in the Windows Deployment and Resource Kits. For more information, search for network connectivity in the Windows Deployment and Resource Kits.

**System Requirements**
The following are the system requirements for NetDiag:

- Windows XP Professional or Windows Server 2003

**File Required**

- Netdiag.exe

**See Also**
**Concepts**
Netdiag Remarks
Netdiag Syntax
Netdiag Examples
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Alphabetical List of Tools

# Netdiag Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetDiag Remarks**
**NetDiag Features**
The NetDiag tool:

- Gathers static network information and tests the network driver, protocol driver, send/receive capability, and well-known target accessibility.
- Can be used by network administrators in conjunction with the Scheduler Service, to generate reports at regularly scheduled intervals.

**See Also**
**Concepts**
Netdiag Overview
Netdiag Syntax
Netdiag Examples
Spcheck Overview
Nltest Overview
Netdom Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Alphabetical List of Tools

# Netdiag Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetDiag Examples**
**Sample Scenario**
In this scenario, you have a problem connecting to a network resource. You receive only a "Network path not found" error message and no further assistance. You have no idea where to start troubleshooting the problem.

Assuming that the reason for the failure in this case is that the DNS server is down, the solution can be found with NetDiag.

At the command line, type:

```
netdiag
```

When executed in this scenario, NetDiag performs tests on each network adapter, and a set of global tests. The tests on the network adapters are performed in the following order:

- Netcard queries test
- IpConfig test
- Autoconfiguration test (APIPA)
- Default gateway test
- NetBT name test
- WINS Service test

Next in this example NetDiag performs a set of global tests in the following order:

- Domain membership test
- NetBT transports test
- Autonet address test (APIPA)
- IP loopback ping test
- Default gateway test
- NetBT name test
- Winsock test
- DNS test
- Redir and Browser test
- DC discovery test
- DC list test
- Trust relationship test

- Kerberos test
- LDAP test
- Bindings test
- WAN configuration test
- Modem configuration test
- IP Security test

The results of these tests show that the network adapter protocol, bindings, and IP address tests succeed. The DNS ping test fails and reports that the DNS server cannot be contacted.

With this information, the administrator knows that either the DNS server address is incorrect, or the DNS server is not responding. Because the DNS address is also displayed as output, you can easily verify whether it is correct.

After the problem is isolated, the administrator can perform additional troubleshooting to determine why the DNS server is down.

For more information on interpreting output from the NetDiag tool see the Interpreting NetDiag Output topic on the Netdiag Remarks page.

**Sample NetDiag Output**

```
.....................................

    Computer Name: RKSRVR-2
    DNS Host Name: rksrvr-2.reskita.microsoft.com
    System info : Windows 2000 Server (Build 2467)
    Processor : x86 Family 6 Model 6 Stepping 0, GenuineIntel
    List of installed hotfixes :
        Q147222


Netcard queries test . . . . . : Passed
    [WARNING] The net card 'Intel(R) PRO/100+ Management Adapter' may not be working.



Per interface results:

    Adapter : Local Area Connection

        Netcard queries test . . . : Passed

        Host Name. . . . . : rksrvr-2
        IP Address . . . . . : 10.10.1.51
        Subnet Mask. . . . . : 255.255.255.0
        Default Gateway. . . . . :
        Dns Servers. . . . . : 10.10.1.77


        AutoConfiguration results. . . . . : Passed

        Default gateway test . . . : Skipped
            [WARNING] No gateways defined for this adapter.

        NetBT name test. . . . . : Passed

        WINS service test. . . . . : Skipped
            There are no WINS servers configured for this interface.

    Adapter : Local Area Connection 2
```

```
        Netcard queries test . . . : Failed
        NetCard Status:          DISCONNECTED
            Some tests will be skipped on this interface.

        Host Name. . . . . . . . . : rksrvr-2
        Autoconfiguration IP Address : 169.254.74.217
        Subnet Mask. . . . . . . . : 255.255.0.0
        Default Gateway. . . . . . :
        Dns Servers. . . . . . . . :



Global results:


Domain membership test . . . . . . . : Passed


NetBT transports test. . . . . . . : Passed
    List of NetBt transports currently configured:
        NetBT_Tcpip_{A2D04C22-3BB8-4FA0-B7DA-414DC1DD08A7}
        NetBT_Tcpip_{56079E37-8246-4712-8B36-F503FF6F9873}
    2 NetBt transports currently configured.


Autonet address test . . . . . . . : Passed


IP loopback ping test. . . . . . . : Passed


Default gateway test . . . . . . . : Failed

    [FATAL] NO GATEWAYS ARE REACHABLE.
    You have no connectivity to other network segments.
    If you configured the IP protocol manually then
    you need to add at least one valid gateway.


NetBT name test. . . . . . . . . . : Passed


Winsock test . . . . . . . . . . . : Passed


DNS test . . . . . . . . . . . . . : Failed
          [WARNING] Cannot find a primary authoritative DNS server for the name
            'bdover.reskita.microsoft.com.'. [ERROR_TIMEOUT]
            The name 'bdover.reskita.microsoft.com.' may not be registered in DNS.

       [WARNING] The DNS entries for this DC cannot be verified right now on DNS
 server 10.10.1.77, ERROR_TIMEOUT.
    [FATAL] No DNS servers have the DNS records for this DC registered.


Redir and Browser test . . . . . . . : Passed
    List of NetBt transports currently bound to the Redir
        NetBT_Tcpip_{A2D04C22-3BB8-4FA0-B7DA-414DC1DD08A7}
        NetBT_Tcpip_{56079E37-8246-4712-8B36-F503FF6F9873}
    The redir is bound to 2 NetBt transports.

    List of NetBt transports currently bound to the browser
        NetBT_Tcpip_{A2D04C22-3BB8-4FA0-B7DA-414DC1DD08A7}
        NetBT_Tcpip_{56079E37-8246-4712-8B36-F503FF6F9873}
    The browser is bound to 2 NetBt transports.


DC discovery test. . . . . . . . . : Passed


DC list test . . . . . . . . . . . : Passed
```

```
Trust relationship test. . . . . : Failed
    Secure channel for domain 'RESKITA' is to '\\a-dcp.reskita.microsoft.com'.
    [FATAL] Cannot set secure channel for domain 'RESKITA' to PDC emulator. [ERR
OR_NO_LOGON_SERVERS]


Kerberos test. . . . . : Passed


LDAP test. . . . . : Passed
    [WARNING] Failed to query SPN registration on DC 'a-dcp.reskita.microsoft.com'.
    [WARNING] Failed to query SPN registration on DC 'a-dc1.reskita.microsoft.com'.
    [WARNING] Failed to query SPN registration on DC 'a-dc3.reskita.microsoft.com'.


Bindings test. . . . . : Passed


WAN configuration test . . . . . : Skipped
    No active remote access connections.


Modem diagnostics test . . . . . : Passed

IP Security test . . . . . : Passed
    Service status  is: Started
    Service startup is: Automatic
    IPSec service is available, but no policy is assigned or active

    Note: run "ipseccmd /?" for more detailed information

The command completed successfully
```

**See Also**

**Concepts**

Netdiag Overview
Netdiag Remarks
Netdiag Syntax
Spcheck Overview
Nltest Overview
Netdom Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Alphabetical List of Tools

# Netdiag Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetDiag Syntax**

NetDiag uses the following syntax:

**netdiag** [/**q**] [**/v**] [**/l**] [**/debug**] [**/d:***DomainName*] [**/fix**] [**/DcAccountEnum**] [**/test:***TestName*] [**/skip:***TestName*] [**/?**]

**Parameters**

**/q**
Specifies quiet output (errors only).

**/v**
Specifies verbose output.

**/l**
Sends output to Netdiag.log. This log file is created in the same directory where Netdiag.exe was run.

**/debug**
Specifies even more verbose output. With this parameter, NetDiag takes a few minutes to complete.

**/d: *DomainName***
Finds a domain controller in the specified domain.

**/fix**
Fixes minor problems.

**/DcAccountEnum**
Enumerates domain controller computer accounts.

**/test: *TestName***
Runs only the listed test(s). TCP/IP must be bound to one or more adapters before running any of the tests. Nonskippable tests are still run. A detailed explanation of the switches can be found on the Netdiag Remarks page.

Valid *TestName* values are:

| Value | Description |
|---|---|
| **Autonet** | Automatic Private IP Addressing (APIPA) address test. Tests whether APIPA is in use for the network adapters. |
| **Bindings** | Bindings test. Lists all bindings, including interface name, lower module name, upper module name, whether the binding is currently enabled, and the owner of the binding. |
| **Browser** | Redirector and Browser test. Lists the protocols bound to the Browser service and the redirector. |
| **DcList** | Domain controller list test. Obtains a list of domain controllers for the domain. |
| **DefGw** | Default gateway test. Attempts to contact each configured default gateway. |
| **DNS** | DNS test. Tests the availability of the configured DNS servers and verifies the current client's DNS registrations. |
| **DsGetDc** | Domain controller discovery test. First finds a generic domain controller from directory service, then finds the primary domain controller. Then, finds a Windows 2000 domain controller (DC). If the tested domain is the primary domain, checks whether the domain GUID stored in Local Security Authority (LSA) is the same as the domain GUID stored in the DC. If not, the test returns a fatal error; if the **/fix** option is on, DsGetDC tries to fix the GUID in LSA. |
| **IpConfig** | IP address configuration test. Enumerates the TCP/IP configuration information for each network adapter. |

| | |
|---|---|
| **IpLoopBk** | IP address loopback ping test. Pings the IP loopback address of 127.0.0.1 for each adapter. |
| **IPSec** | IP Security test. Tests whether IP Security is enabled and displays a list of active IP Security policies for the computer. |
| **IPX** | IPX test. Lists statistics for the IPX protocol installed on the computer. |
| **Kerberos** | Kerberos test. Checks whether the Kerberos package information is up-to-date. |
| **Ldap** | Lightweight Directory Access Protocl (LDAP) test. Contacts all available domain controllers and determines which LDAP authentication protocol is in use. |
| **Member** | Domain membership test. Checks to confirm details of the primary domain, including computer role, domain name, and domain GUID. Checks to see if NetLogon service is started, adds the primary domain to the domain list, and queries the primary domain security identifier (SID). |
| **Modem** | Modem diagnostics test. Lists configuration information for each modem found. |
| **NbtNm** | NetBT name test. Similar to the **nbtstat -n** command. It checks that the workstation service name `<00>` is equal to the computer name. It also checks that the messenger service name `<03>`, and server service name `<20>` are present on all interfaces and that none of these names are in conflict. |
| **Ndis** | Netcard queries test. Lists the network adapter configuration details, including the adapter name, configuration, media, globally unique identifier (GUID), and statistics. If this test shows an unresponsive network adapter, the remaining tests are aborted. |
| **NetBTTransports** | NetBT transports test. Lists the transport protocols that are bound to NetBT. |
| **Netstat** | Netstat information test. Lists protocol statistics and current TCP/IP connections. |
| **Netware** | Netware test. Queries the nearest Netware server for current login information. |
| **Route** | Routing table test. Lists static routes and whether they are persistent. |
| **Trust** | Trust relationship test. Tests trust relationships to the primary domain only if the computer is a member workstation, member server, or a Backup Domain Controller (BDC) domain controller that is not a PDC emulator. Checks that the primary domain security identifier (SID) is correct. Contacts an active DC. Connects to the SAM server on the DC. Uses the domain SID to open the domain to verify whether the domain SID is correct. Queries info of the secure channel for the primary domain. If the computer is a BDCDC, reconnects to the PDC emulator. If the computer is a member workstation or server, sets secure channel to each DC on the DC list for this domain. |
| **WAN** | Wide Area Network (WAN) configuration test. Lists settings and status on each COM port currently in use. |
| **WINS** | Windows Internet Name Service (WINS) service test. Tests the availability of the configured WINS server and the validity of the client registrations. |
| **Winsock** | Winsock test. Lists protocols and ports available to the WinSock service. |

**/skip:** ***TestName***

Skips the test specified by *TestName*. Nonskippable tests will still run. Valid *TestName* values are:

| Value | Description |
|---|---|
| **Autonet** | Automatic Private IP Addressing (APIPA) address test. |
| **Bindings** | Bindings test. |
| **Browser** | Redir and Browser test. |
| **DcList** | Domain controller list test. |

| | |
|---|---|
| **DefGw** | Default gateway test. |
| **DNS** | DNS test. |
| **DsGetDc** | Domain controller discovery test. |
| **IpConfig** | IP address configuration test. |
| **IpLoopBk** | IP address loopback ping test. |
| **IPX** | IPX test. |
| **Kerberos** | Kerberos test. |
| **Ldap** | Lightweight Directory Access Protocol (LDAP) test. |
| **Modem** | Modem diagnostics test. |
| **NbtNm** | NetBT name test. |
| **Netstat** | Netstat information test. |
| **Netware** | Netware test. |
| **Route** | Routing table test. |
| **Trust** | Trust relationship test. |
| **WAN** | Wide Area Network (WAN) configuration test. |
| **WINS** | Windows Internet Name Service (WINS) test. |
| **Winsock** | Winsock test. |

**See Also**
**Concepts**
Netdiag Overview
Netdiag Remarks
Netdiag Examples
Spcheck Overview
Nltest Overview
Netdom Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Alphabetical List of Tools

# Netdom Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netdom.exe: Windows Domain Manager**
This command-line tool enables administrators to manage Windows Server 2003 and Windows 2000 domains and trust relationships from the command line.

You can use NetDom to:

- Join a Windows XP Professional-based computer to a Windows Server 2003 or Windows 2000 or Windows NT 4.0 domain.
    - Provide an option to specify the organizational unit for the computer account.
    - Generate a random computer password for initial join.
- Manage computer accounts for domain member workstations and member servers. Management operations include:
    - Add, Remove, Query.
    - An option to specify the organizational unit for the computer account.
    - An option to move an existing computer account for a member workstation from one domain to another while maintaining the security descriptor on the computer account.
- Establish one-way or two-way trust relationships between domains, including the following kinds of trust relationships:
    - From a Windows 2000 or Windows Server 2003 domain to a Windows NT 4.0 domain.
    - From a Windows 2000 or Windows Server 2003 domain to a Windows 2000 or Windows Server 2003 domain in another enterprise (an "uplevel" external trust).
    - Between two Windows 2000 or Windows Server 2003 domains in an enterprise (a shortcut trust).
    - The Windows Server 2003 or Windows 2000 Server half of an interoperable Kerberos realm.
- Verify and/or reset the secure channel for the following configurations:
    - Member workstations and servers.
    - BDCs in a Windows NT 4.0 domain.
    - Specific Windows Server 2003 or Windows 2000 replicas.
- Manage trust relationships between domains, including the following operations:
    - Enumerate trust relationships (direct and indirect).
    - View and change some attributes on a trust.

**Note**

- You must run NetDom from the command window.

Corresponding UI

Much of the functionality of NetDom can be accessed from the Microsoft Management Console Active Directory Users and Computers snap-in, which is part of Windows Server 2003. More functionality can be accessed from the Active Directory: Domains and Trust snap-in which is the corresponding UI for NetDom's Trust Management functionality.

Concepts

A trust relationship is a defined affiliation between domains that enables pass-through authentication.

A one-way trust relationship between two domains means that one domain (the trusting domain) allows users who have accounts on the other domain (the trusted domain), access to its resources.

The one-way trust relationship described here is helpful in master domain models, but it is not the only kind of trust relationship. When two one-way trusts are established between domains, it is known as a two-way trust. In two-way trusts, each domain treats the users from the trusted (and trusting) domain as its own users.

System Requirements

The following are the system requirements for NetDom:

- Windows Server 2003 or Windows XP Professional

  **Note**

  - The Windows Server 2003 Administration Tools Pack must be installed on computers running Windows XP Professional for Netdom.exe to work. This tools pack is installed by default for computers running Windows Server 2003.

**File Required**

- Netdom.exe

**See Also**
**Concepts**
Netdom Remarks
Netdom Syntax
Netdom Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Netdom Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetDom Remarks**

**Using the /Verbose Parameter**

By default, only the result of an operation is reported. For example, if you use the **join** operation you see output similar to the following:

```
success: mywksta joined to mycompany domain
```

If you specify the **/verbose** parameter, the output lists the success or failure of each transaction necessary to perform the operation. For example, this time when you use the **join** operation, you see output similar to the following:

```
success: adding machine account for mywksta to mycompany domain
success: configuring lsa on mywksta
success: mywksta joined to mycompany domain
```

**Using the /Reboot Parameter**

The **/reboot** parameter specifies that the computer being acted upon by the specified NetDom operation is shut down and automatically rebooted after the completion of the operation. When you specify the **/reboot** parameter, the following message and a countdown timer display on the workstation screen, prior to the restart:

```
The system is shutting down. Please save
all work in progress and logoff. Any unsaved changes
will be lost. This shutdown was initiated because the
domain which this machine belongs to was changed by
nnn.
```

The name of the administrator entered by using the **/uo** parameter is substituted for *nnn*.

The default delay before computer restart is 20 seconds.

**See Also**

**Concepts**

Netdom Overview
Netdom Syntax
Netdom Examples
Spcheck Overview
Nltest Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Alphabetical List of Tools

# Netdom Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetDom Examples**

**Example 1: Add a Workstation or Member Server to a Windows NT 4.0 Domain**

To add the workstation **mywksta** to the Windows NT 4.0 domain**reskita**, type the following at the command line:

```
netdom add /d:reskita mywksta /ud:mydomain\admin /pd:password
```

**Example 2: Add a Workstation or Member Server to a Windows Server 2003 Domain**

To add the workstation mywksta to the Windows Server 2003 domain devgroup.example.com in the organizational unit (OU) Dsys/workstations, type the following at the command prompt:

```
netdom add/d:devgroup.example.com mywksta /OU:OU=Dsys,OU=Workstations,DC=microsoft,DC=com
```

### Note

- If the **/ou** parameter is not specified, the account is created in the Computers container.

**Example 3: Move a Windows NT 4.0 BDC to a new domain**

To join myBDC to the Windows NT 4.0 domain reskita type the following at the command prompt:

```
netdom mybdc moveNT4BDC /domain:reskita
```

**Example 4: Add an alternate name for a Windows Server 2003 domain controller**

To give an alternate name for the domain controller DC in the example.com domain, use the following syntax:

```
netdom computername dc /add:altDC.example.com
```

A name must first exist as an alternate before it can be made the primary name of a computer.

**Example 5: Rename a domain controller in a Windows Server 2003 domain**

To rename the domain controller DC to altDC in the example.com domain use the following syntax:

```
netdom computername dc /makeprimary:altdc.example.com
```

To rename a member server you must choose one of the existing alternate names for the computer and make it the new primary name.

**Example 6: Rename a Member Server**

To rename the member server member to member1, type the following at the command prompt:

```
netdom renamecomputer member /newname:member1.example.com /userd:administrator
```

**Example 7: Join a Workstation or Member Server to a Domain**

To join mywksta to the devgroup.example.com domain in the Dsys/workstations organizational unit, type the following at the command prompt:

```
netdom join /d:devgroup.example.com mywksta /OU:OU=Dsys,OU=Workstations,DC=microsoft,DC=com
```

Besides adding the computer account to the domain, the workstation is modified to contain the appropriate shared secret to complete the **join** operation.

**Example 8: Remove a Workstation or Member Server from a Domain**

To remove mywksta from the mydomain domain and make the workstation a part of a workgroup, type the following at the command prompt:

```
netdom remove /d:mydomain mywksta /ud:mydomain\admin /pd:password
```

**Example 9: Move a Workstation or Member Server from One Domain to Another**

To move mywksta from its current domain into the mydomain domain, type the following at the command prompt:

```
netdom move /d:mydomain mywksta /ud:mydomain\admin /pd:password
```

If the destination is a Windows 2000 domain, the Security ID history (SIDHistory) for the workstation is updated, retaining the security permissions that the computer account had previously.

**Example 10: Reset the secure channel for a workstation, member server, or Windows NT 4.0 BDC**
To reset the secure channel secret maintained between mywksta and devgroup.example.com (regardless of OU), type the following at the command prompt:

```
netdom reset /d:devgroup.example.com mywksta
```

To reset the secure channel between the Windows NT 4.0 PDC for Northamerica and the backup domain controller NABDC, type the following at the command prompt:

```
netdom reset /d:Northamerica NABDC
```

**Example 11: Force a Secure Channel Session Between a Member and a Specific Domain Controller**
Members often establish secure channel sessions with non-local domain controllers. To force a secure channel session between a member and a specific domain controller by using the **/server** parameter with the **reset** operation, type the following at the command prompt:

```
netdom reset /d:devgroup.example.com mywksta /Server:mylocalbdc
```

**Example 12: Verify a Workstation or Member Server Secure Channel**
To verify the secure channel secret is maintained between mywksta and devgroup.example.com, type the following at the command prompt:

```
netdom verify /d:devgroup.example.com mywksta
```

**Example 13: Establish a One-Way Trust Relationship**
When used with the **trust** operation, the **/d:***Domain* parameter always refers to the trusted domain.

To set the Windows NT 4.0 resource domain USA-Chicago to trust the Windows NT 4.0 account domain Northamerica, type the following at the command prompt:

```
netdom trust /d:Northamerica USA-Chicago /add /Ud:Northamerica\admin /Pd:* /Uo:USA-Chicago\admin /Po:*
```

Press Enter and the following prompt is displayed:

```
Password for Northamerica\admin:
```

Enter the password for Northamerica\admin and press Enter. The following prompt is displayed:

```
Password for USA-Chicago\admin:
```

Enter the password for USA-Chicago\admin and press Enter.

The user must have credentials for both domains. The **/pd** parameter can be used to specify the password for Northamerica\admin and the **/po** parameter can be used to specify the password for USA-Chicago\admin. If passwords are not provided on the command line, the user is prompted for both.

If you then want to specify a two-way trust, type the following at the command prompt

```
netdom trust /d:marketing.example.com engineering.example.com /add /twoway /Uo:admin@engineering.example.com
/Ud:admin@marketing.example.com:
```

**Example 14: Establish a One-Way Trust Relationship from a Windows Domain to a Non-Windows Kerberos Realm**
To establish a one-way trust where Northamerica trusts the non-Windows Kerberos realm ATHENA, type the following at the command prompt:

```
netdom trust /d:ATHENA Northamerica /add /PT:password /realm
```

The **/d** parameter specifies the trusted domain and the **/realm** parameter indicates that this is a non-Windows Kerberos realm. The order of the domains is not important. Credentials to the Windows 2000 domain can be supplied if needed.

### Note

- Verifying a specific trust relationship requires credentials unless the user has domain administrator privileges on both domains.

If you want to set the Kerberos realm ATHENA to trust the Northamerica domain, type the following at the command prompt:

```
netdom trust /d:Northamerica ATHENA /add
```

### Note

- To make the trust two-way, you can specify the **/twoway** parameter.

Non-Windows Kerberos trusts are created as non-transitive. If you want to change the trust from ATHENA to Northamerica as transitive, type the following at the command prompt:

```
netdom trust Northamerica /d:ATHENA /trans:yes
```

To display the transitive state, type the following at the command prompt:

```
netdom trust Northamerica /d:ATHENA /trans
```

The order of the two domains above is not important. Either can be the non-Windows Kerberos domain.

**Example 15: Break a One-Way Trust Relationship**
To undo the trust that USA-Chicago has for Northamerica, type the following at the command prompt:

```
netdom trust /d:Northamerica USA-Chicago /remove
```

**Example 16: Break a Two-Way Trust Relationship**
To break a two-way trust relationship, type the following at the command prompt:

```
netdom trust /d:marketing.example.com Engineering.example.com /remove /twoway
/Uo:admin@engineering.example.com /Ud:admin@marketing.example.com
```

**Example 17: Verify a Specific Trust Relationship**
To verify the one-way trust that USA-Chicago has for Northamerica, type the following at the command prompt:

```
netdom trust /d:Northamerica USA-Chicago /verify
```

To verify a two-way trust between the Northamerica and Europe domains, type the following at the command prompt:

```
netdom trust /d:Northamerica EUROPE /verify /twoway
```

The **/verify** parameter checks that the appropriate shared secrets are synchronized between the two items involved in the trust.

**Example 18: Reset a Specific Trust Relationship**
To reset the secure channel for the one-way trust between Northamerica and USA-Chicago, type the following at the command prompt:

```
netdom trust /d:Northamerica USA-Chicago /Ud:Northamerica\admin /reset
```

The **/reset** parameter synchronizes the appropriate shared secrets if they are not already synchronized.

**Example 19: Verify Kerberos Functionality**
To verify Kerberos authentication between a workstation and a service located in the domain devgroup.example.com, type the following at the command prompt:

```
netdom trust /d:devgroup.example.com /verify /KERBEROS
```

When you use the **NetDom trust** operation with the **/verify /kerberos** parameters, it seeks a session ticket for the Kerberos Admin service in the target domain. If successful, you can conclude that all Kerberos operations (for example KDC referrals) are operating correctly between the workstation and the target domain.

### Note

- This operation cannot be executed remotely. It must be run on the workstation being tested.

**Example 20: View All Workstation Members in a Domain**
To list all the workstations in the domain Northamerica, type the following at the command prompt:

```
netdom query /d:Northamerica WORKSTATION
```

**Example 21: View All Server Members in a Domain**
To list all of the servers in Northamerica, type the following at the command prompt:

```
netdom query /d:Northamerica SERVER
```

**Example 22: View All Domain Controller Members in a Domain**
To list all the domain controllers in the domain Northamerica, type the following at the command prompt:

```
netdom query /d:Northamerica DC
```

**Example 23: View All Organizational Unit Members in a Domain**
To list all of the OUs in devgroup.example.com, type the following at the command prompt:

```
netdom query /d:devgroup.example.com OU
```

**Example 24: List the Primary Domain Controller Member in a Domain**
To list the PDC for Northamerica, type the following at the command prompt:

```
netdom query /d:Northamerica PDC
```

**Example 25: List the Primary Domain Controller Emulator in a Domain**
To list the current PDC emulator for devgroup.example.com, type the following at the command prompt:

```
netdom query /d:devgroup.example.com FSMO
```

**Example 26: Perform Secure Channel Batch Repair**
You can use the **query** operation with the **/verify** and **/reset** parameters to perform these operations together. You can pipe the output of the **query** operation to the **NetDom verify** or **NetDom reset** operation.

To list all servers and verify secure channel secret, type the following at the command prompt:

```
netdom query /d:Northamerica SERVER /verify
```

To list all workstations and reset any unsynchronized secure channel secrets, type the following at the command prompt:

```
netdom query /d:Northamerica WORKSTATION /reset
```

**Example 27: View Domain Trusts**
To view all the direct trust relationships for the domain Northamerica, type the following at the command prompt:

```
netdom query /d:Northamerica /Ud:Northamerica\admin DOMAIN /Direct
```

To view all the direct and indirect trust relationships for the domain Northamerica, type the following at the command prompt:

```
netdom query /d:Northamerica /Ud:Northamerica\admin DOMAIN
```

To view all trust relationships and check their status, type the following at the command prompt:

```
netdom query /d:devgroup.example.com DOMAIN /verify
```

**Example 28: List the routed name suffixes for the trust between my TestDomain and the trustpartnerdomain**
**Note**

- The **/d** parameter is not needed for this operation which is an exception from other trust operations.

To list the routed name suffixes for the trust between my TestDomain abd the trustpartnerdomain, type the following at the command prompt:

```
netdom trust myTestDomain /namesuffixes:trustpartnerdomain
```

This will list out all the routed name suffixes for the trust between myTestDomain and the trustpartnerdomain. The trust **must** be either a Cross Forest trust or Non-Windows Realm Trust with the Forest Transitive attribute set.

**Sample Output**

```
C:\nt\ds\netapi\netdom\obj\i386>netdom trust shasandom2 /namesuffixes:powermatic

   Name, Type, Status, Notes

1. *.flotsam.org, Name Suffix, Enabled

2. *.powermatic.nttest.microsoft.com, Name Suffix, Enabled

3. *.jetsam.com, Name Suffix, Enabled

4. *.blah.com, Name Suffix, Conflicting, With shasandom2.nttest.microsoft.com

5. unisaw.powermatic.nttest.microsoft.com, Domain DNS name, Enabled

6. UNISAW, Domain NetBIOS name, Enabled, For unisaw.powermatic.nttest.microsoft.
```

```
com
7. s-1-5-21-1550512861-723516995-420396236, Domain SID, Enabled, For unisaw.powe
rmatic.nttest.microsoft.com
8. powermatic.nttest.microsoft.com, Domain DNS name, Enabled
9. POWERMATIC, Domain NetBIOS name, Enabled, For powermatic.nttest.microsoft.com

10. s-1-5-21-1390067357-1757981266-527237240, Domain SID, Enabled, For powermati
c.nttest.microsoft.com

The command completed successfully.
```

**Example 29: Enable/disable the first routed name suffix in the list generated by the previous command**

**Note**

- **/ToggleSuffix** must be used with **/NameSuffixes**. Use **/NameSuffixes** immediately before using **/ToggleSuffix** because the order in which the name suffixes are listed may change.

To enable/disable the first routed name suffix in the list generated by the previous command, type the following at the command prompt:

```
netdom trust myTestDomain /namesuffixes:foresttrustpartnerdomain /togglesuffix:1
```

**Sample Output**

**Note**

- The output generated reflects the routed name suffix list after the Toggling.

```
C:\nt\ds\netapi\netdom\obj\i386>netdom trust shasandom2 /ns:powermatic /ts:1
   Name, Type, Status, Notes
1. *.flotsam.org, Name Suffix, Admin-Disabled
2. *.powermatic.nttest.microsoft.com, Name Suffix, Enabled
3. *.jetsam.com, Name Suffix, Enabled
4. *.blah.com, Name Suffix, Conflicting, With shasandom2.nttest.microsoft.com
5. unisaw.powermatic.nttest.microsoft.com, Domain DNS name, Enabled
6. UNISAW, Domain NetBIOS name, Enabled, For unisaw.powermatic.nttest.microsoft.
com
7. s-1-5-21-1550512861-723516995-420396236, Domain SID, Enabled, For unisaw.powe
rmatic.nttest.microsoft.com
8. powermatic.nttest.microsoft.com, Domain DNS name, Enabled
9. POWERMATIC, Domain NetBIOS name, Enabled, For powermatic.nttest.microsoft.com

10. s-1-5-21-1390067357-1757981266-527237240, Domain SID, Enabled, For powermati
c.nttest.microsoft.com
```

```
The command completed successfully.
```

**Example 30: Adding/Removing TLN Suffixes**

To add the DNS name suffix blah.com to the Forest Trust Info with trustpartnerdomain, type the following at the command prompt:

```
Netdom trust myTestDomain /d:trustPartnerDomain /AddTln:blah.com
```

### Note

- This is only allowed for a trust with a Forest Transitive, Non-Windows Realm Trust. This is also true for:
    - Netdom trust myTestDomain /d:trustPartnerDomain /RemoteTln:blah.com
    - Netdom trust myTestDomain /d:trustPartnerDomain /AddTLNEx:something.blah.com (must have a TLN entry present for the parent Naming context, in this case blah.com, otherwise operation will be disallowed)
    - Netdom trust myTestDomain /d:trustPartnerDomain /RemoveTLNEx:something.blah.com

**List the name suffixes on a Non-Windows Realm Trust**

```
C:\nt\ds\netapi\netdom\obj\i386>netdom trust shasandom2 /ns:blah.mit.org

   Name, Type, Status, Notes

1. *.bowwow.com, Name Suffix, Enabled

2. *.meow.com, Name Suffix, Enabled


The command completed successfully.
```

**Add another TLN**

```
C:\nt\ds\netapi\netdom\obj\i386>netdom trust shasandom2 /d:blah.mit.org /addtln:

dude.com

The TLN or Exclusion was successfully added to the Forest Trust Info.

The command completed successfully.
```

**Add an invalid TLN exclusion**

```
C:\nt\ds\netapi\netdom\obj\i386>netdom trust shasandom2 /d:blah.mit.org /addtlne

x:dude.com

The Forest Trust Info for the specified trust could not be stored.

The parameter is incorrect.


Try "NetDom HELP" for more information.
```

**Add an valid TLN exclusion**

```
C:\nt\ds\netapi\netdom\obj\i386>netdom trust shasandom2 /d:blah.mit.org /addtlne

x:cool.dude.com
```

```
The TLN or Exclusion was successfully added to the Forest Trust Info.

The command completed successfully.
```

**View the result of previous operations**

```
C:\nt\ds\netapi\netdom\obj\i386>netdom trust shasandom2 /ns:blah.mit.org

   Name, Type, Status, Notes

1. *.cool.dude.com, Exclusion

2. *.bowwow.com, Name Suffix, Enabled

3. *.meow.com, Name Suffix, Enabled

4. *.dude.com, Name Suffix, Enabled


The command completed successfully.
```

**See Also**
**Concepts**
Netdom Overview
Netdom Remarks
Netdom Syntax
Spcheck Overview
Nltest Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Alphabetical List of Tools

# Netdom Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetDom Syntax**
NetDom uses the following general syntaxes:

- **NetDom** *Operation* [*Computer*] [{/**d: | /domain:**}*Domain*] [*Options*]
- **NetDom help** *Operation*

**NetDom add**
Adds a workstation or server account to the domain.

**Syntax**
- **NetDom add** *Computer* {**/d: | /domain:**}*Domain* [{**/ud: | /userd:**}[*Domain*\]*User* {**/pd: | /passwordd:**}{*Password*|*}] [{**/s: | /server:**}*Server*] [**/ou:***OUPath*] [**/dc**] [**/help | /?**]

**Parameters**
***Computer***
Specifies the name of the computer to be added.

**{/d: | /domain:} *Domain***
Specifies the domain in which to create the account. If this parameter is omitted, then the domain that the current computer belongs to is used.

**{/ud: | /userd:}[ *Domain*\] *User***
Specifies the user account that makes the connection with the domain that is specified in the **/d** or **/domain** parameter. If this parameter is omitted, the current user account is used.

**{/pd: | /passwordd:}{ *Password*|*}**
Specifies the password of the user account that is specified in the **/ud** or **/userd** parameter. Use the wildcard character (*****) to be prompted for the password.

**{/s: | /server:} *Server***
Specifies the name of a domain controller that performs the add.

**/ou: *OUPath***
Specifies the organizational unit (OU) under which to create the account. This must be the full RFC 1779 distinguished name of the OU. If omitted, the account is created under the default OU for machine objects for that domain.

**/dc**
Specifies that a domain controller's machine account is to be created. This allows the computer accounts for new Windows 2000, Windows Server 2003 domain controllers, and new Windows NT 4.0 backup domain controllers (BDCs) to be pre-created. If installing a new Windows NT 4.0 BDC into an existing Windows 2000 or Windows Server 2003 domain, the computer account must be pre-created. This parameter cannot be used with the **/ou** parameter.

**{/help | /?}**
Displays command-line help for the **add** operation.

**Sample Usage**
```
netdom add /d:reskit.ms.com mywksta

netdom add /d:reskit.ms.com mynt4dc /dc
```

**NetDom computername**
Manages the primary and alternate names for a computer. This command can safely rename Windows Server 2003 and Windows 2000 domain controllers as well as member servers.

**Syntax**
- **NetDom computername** *Computer* [**usero:***User* [**/passwordo:**[*Password*|*]] [**userd:***User* [**/passwordd:**[*Password*|*]] {**/add:***NewAltDNSName* | **/remove:***AltDNSName* | **/makeprimary:***ComputerDNSNam*

**Parameters**
***Computer***
Specifies the name of the computer to be added.

**/usero:[*domain*\]*UserName***
Specifies the user account to be used for the originating domain.

**/passwordo:[*Password* | *]**
Specifies the password to be used for the originating domain. If the wildcard character (*****) is used, then the user will be prompted for the password.

**/userd:[*Domain*\]*UserName***
Specifies the user account to be used for the destination domain.

**/passwordd:[*Password* | *]**
Specifies the password to be used for the destination domain. If the wildcard character (*****) is used, then the user will be prompted for the password.

**/add: *NewAltDNSName***
Specifies that a new alternate name should be created. The name must be a fully qualified domain name (computer name followed by primary DNS suffix, such as comp1.example.com).

**/remove: *AltDNSName***
Specifies that an existing alternate name should be deleted. The name must be a fully qualified domain name (computer name followed by primary DNS suffix, such as comp1.example.com).

**/makeprimary: *ComputerDNSName***
Specifies that an existing alternate name should be made into the primary name. The name must be a fully qualified domain name (computer name followed by primary DNS suffix, such as comp1.example.com).

**/enumerate[:{ALTERNATENAMES | PRIMARYNAME | ALLNAMES}]**
Lists the primary and/or any alternate names. The following valid values can be specified:

| Value | Description |
|---|---|
| **ALTERNATENAMES** | Lists the alternate names only. |
| **PRIMARYNAME** | Lists the primary name only. |
| **ALLNAMES** | Lists the primary and any alternate names. This is the default. |

**/verify**
Checks if there is a DNS A record and an Service Principal Name (SPN) for each computer name.

**{/help | /?}**
Displays command-line help for the **computername** operation.

**Sample Usage**
```
netdom computername reskit /enumerate:primaryname
```

**NetDom join**
Joins a workstation or member server to a domain. The act of joining a computer to a domain creates an account for the computer on the domain, if it does not already exist.

**Syntax**

**NetDom join** *Computer* {**/d: | /domain:**}*Domain* [**/ou:***OUPath*] [{**/ud: | /userd:**}[*Domain*\]*User* [{**/pd: | /passwordd:**}{*Password*|*}]] [{**/uo: | /usero**}*User* [{**/po: | /passwordo**} {*Password*|*}] [**/reboot**[**:***Delay*]] [**/help | /?**]

**Parameters**

**Note**

- When joining a computer running Windows NT 4.0 or earlier to the domain, the operation is not transacted. This means that a failure during the operation might leave the computer in an undetermined state with respect to the domain to which it was meant to join.

***Computer***
Specifies the name of the computer to be joined.

**{/d: | /domain:}*Domain***
Specifies the domain to which the account is joined. If this parameter is omitted, then the domain that the current computer belongs to is used.

**/ou: *OUPath***
Specifies the organizational unit (OU) under which to create the account. This must be the full RFC 1779 distinguished name of the OU. If omitted, the account is created under the default OU for machine objects for that domain.

**{/ud: | /userd:}[*Domain*\]*User***
Specifies the user account that makes the connection with the domain specified in the **/d** or **/domain** parameter. If this parameter is omitted, the current user account is used.

**{/pd: | /passwordd:}{*Password*|*}**
Specifies the password of the user account that is specified in the **/ud** or **/userd** parameter. Use the wildcard character (*****) to be prompted for the password.

**{/uo: | /usero}*User***
Specifies the user account that makes the connection with the computer to be joined. If this parameter is omitted, the current user account is used.

**{/po: | /passwordo}{*Password*|*}**
Specifies the password of the user account that is specified in the **/uo** or **/usero** parameter. Use the wildcard character (*****) to be prompted for the password.

**/reboot[:*Delay*]**
Specifies that the computer shuts down and automatically reboots after the join has completed. The *Delay* value is the number of seconds before automatic shutdown occurs. The default *Delay* value is 20 seconds.

**{/help | /?}**
Displays command-line usage for the **join** operation.

**Sample Usage**
```
netdom join /d:reskit.ms.com mywksta
```

**NetDom move**
Moves a workstation or member server to a new domain. The act of moving a computer to a new domain creates an account for the computer on the domain, if it does not already exist.

**Syntax**
**NetDom move** *Computer* {**/d: | /domain:**}*Domain* [**/ou:***OUPath*] [{**/ud: | /userd**}[*Domain*\]*User* [{**/pd: | /passwordd**}{*Password*|*}] [{**/uo: | /usero**}[*Domain*\]*User* [{**/po: | /passwordo**} {*Password*|*}]] [{**/uf: | /userf**}[*Domain*\]*User* [{**/pf: | /passwordf**}{*Password*|*}]] [**/reboot**[**:***Delay*]] [{**/help | /?**}]

**Parameters**

**Notes**

- When moving a computer running Windows NT 4.0 or earlier to the domain, the operation is not transacted. This means that a failure during the operation might leave the computer in an undetermined state with respect to the domain of its intended move.
- When moving a computer to a new domain, the old computer account in the previous domain is not deleted. If credentials are supplied for the former domain, the old computer account is disabled.
- The act of moving a computer to a new domain will create an account for the computer on the domain if it does not already exist.

***Computer***
Specifies the name of the computer to be moved.

**{/d: | /domain:}*Domain***
Specifies the domain to which the account is moved. If the parameter is omitted, then the domain that the current computer belongs to is used.

**/ou: *OUPath***
Specifies the organizational unit (OU) under which to create the account. This must be the full RFC 1779 distinguished name of the OU. If omitted, the account is created under the default OU for machine objects for that domain.

**{/ud: | /userd}[*Domain*\]*User***
Specifies the user account that makes the connection with the domain specified in the **/d** or **/domain** parameter. If this parameter is omitted, the current user account is used.

**{/pd: | /passwordd}{*Password*|*}**
Specifies the password of the user account that is specified in the **/ud** or **/userd** parameter. Use an asterisk (*) to be prompted for the password.

**{/uo: | /usero}*User***
Specifies the user account to make the connection with the computer to be moved. If this parameter is omitted, the current user account is used.

**{/po: | /passwordo}{*Password*|*}**
Specifies the password of the user account that is specified in the **/uo** or **/usero** parameter. Use the wildcard character (*****) to be prompted for the password.

**{/uf: | /userf}*User***
Specifies the user account to make the connection with the computer's former domain (of which the computer had been a member prior to the move). This parameter is used to disable the old computer account.

**{/pf: | /passwordf}{*Password*|*}**
Specifies the password of the user account that is specified in the **/uf** or **/userf** parameter. Use the wildcard character (*****) to be prompted for the password.

**/reboot[:*Delay*]**
Specifies that the computer shuts down and automatically reboots after the move has completed. The *Delay* value is the number of seconds before automatic shutdown occurs. The default *Delay* value is 20 seconds.

**{/help | /?}**
Displays command-line usage for the **move** operation.

**Sample Usage**
```
netdom move /d:newdomain mywksta
```

**NetDom query**
Queries the domain for information such as membership and trust.

**Syntax**
**NetDom query** {**/d: | /domain:**}*Domain* [{**/s: | /server:**}*Server*] [{**/ud: | /userd:**}[*Domain*\]*User* {**/pd: | /passwordd**}{*Password*|*}] [**/verify**] [**/reset**] [**/direct**] {**WORKSTATION**|**SERVER**|**DC**|**OU**|**PDC**|**FSMO**|**TRUST**} [{**/help | /?**}]

**Parameters**
**{/d: | /domain:}*Domain***

Specifies the domain to query for the information. If this parameter is omitted, then the domain that the current computer belongs to is used.

**{/s: | /server:}*Server***
Specifies the name of a domain controller that performs the query.

**{/ud: | /userd:}[*Domain*\]*User***
Specifies the user account that makes the connection with the domain in the **/d** or **/domain** parameter. If this parameter is omitted, the current user account is used.

**{/pd: | /passwordd}{*Password*|*}**
Specifies the password of the user account that is specified in the **/ud** or **/userd** parameter. Use the wildcard character (*****) to be prompted for the password.

**/verify**
Specifies verification of the secure channel secrets for all enumerated memberships or trusts, and displays them. Unless the user is an enterprise-level administrator, it will not be possible to verify all secure channel secrets.

**/reset**
Specifies resynchronization ofthe secure channel secrets for all enumerated memberships or trusts which are currently broken. The **/reset** parameter implies the **/verify** parameter. Unless the user is an enterprise-level administrator, it might not be possible to reset all enumerated trusts or memberships.

**/direct**
Indicates that the query for trust relationships returns only direct trust relationships, rather than direct and indirect relationships. This parameter is valid only when *Domain* is specified with the **/d** parameter.

**WORKSTATION|SERVER|DC|OU|PDC|FSMO|TRUST**
Specifies the type of list to generate.

| Object | Description |
|---|---|
| **WORKSTATION** | Queries the domain for the list of workstations. |
| **SERVER** | Queries the domain for the list of servers. |
| **DC** | Queries the domain for the list of domain controllers. |
| **OU** | Queries the domain for the list of OUs under which the specified user can create a machine object. |
| **PDC** | Queries the domain for the current primary domain controller. |
| **FSMO** | Queries the domain for the current list of operations master (also know as flexible single master operations or FSMO) owners. |
| **TRUST** | Queries the domain for the list of its trusts. |

**{/help | /?}**
Displays command-line usage for the **query** operation.

**Sample Usage**
```
netdom query /d:reskit.ms.com DC
```

**NetDom remove**
Removes a workstation or server from the domain.

**Syntax**
▸ **NetDom remove** *Computer* {**/d: | /domain:**}*Domain* [{**/ud: | /userd:**}[*Domain*\]*User* [{**/pd: | /passwordd**}{*Password*|*}]] [{**/uo: | /usero**}*User* [{**/po: |/passwordo**}{*Password*|*}]] [**/reboot**[**:***Delay*]] [{**/help | /?**}]

**Parameters**
***Computer***
Specifies the name of the computer to be removed.

**{/d: | /domain:}*Domain***
Specifies the domain from which the account is to be removed. If this parameter is omitted, then the domain that the current computer belongs to is used.

**{/ud: | /userd:}[ *Domain*\]*User***
Specifies the user account that makes the connection with the domain in the **/d** or **/domain** parameter. If this parameter is omitted, then the current user account is used.

**{/pd: | /passwordd}{*Password*|*}**
Specifies the password of the user account that is specified in the **/ud** or **/userd** parameter. Use the wildcard character (*****) to be prompted for the password.

**{/uo: | /usero}*User***
Specifies the user account to make the connection with the computer to be removed. If this parameter is omitted, then the current user account is used.

**{/po: |/passwordo}{*Password*|*}**
Specifies the password of the user account that is specified in the **/uo**or **/usero** parameter. Use the wildcard character (*****) to be prompted for the password.

**/reboot[:*Delay*]**
Specifies that the computer shuts down and automatically restarts after the **remove** operation has completed. The *Delay* value is the number of seconds before automatic shutdown occurs. The default *Delay* value is 20 seconds.

**{/help | /?}**
Displays command-line usage for the **remove** operation.

**Sample Usage**
```
netdom remove /d:reskit.ms.com mywksta
```

**NetDom movent4bdc**
Renames a Windows NT 4.0 backup domain controller to reflect a domain name change. This can assist in Windows NT 4.0 domain renaming efforts.

**Syntax**
▸ **netdom movent4bdc** *Computer* [{**/d: | /domain:**}*Domain*] [**/reboot**[**:***Delay*]] [{**/help | /?**}]

**Parameters**
***Computer***
Specifies the name of the backup domain controller to rename.

**{/d: | /domain:}*Domain***
Specifies the new name of the domain.

**/reboot[:*Delay*]**
Specifies that the computer shuts down and automatically reboots after the **rename** operation has completed. The *Delay* value is the number of seconds before automatic shutdown occurs. The default *Delay* value is 20 seconds.

**{/help | /?}**
Displays command-line usage for the **movent4bdc** operation.

**Sample Usage**
```
netdom movent4bdc /d:newdomain BDC51
```

**NetDom renamecomputer**
Renames a domain computer and its corresponding domain account. Use this command to rename domain workstations and member servers only. To rename domain controllers, use the NetDom

computername command.

**Syntax**

**NetDom renamecomputer** *Computer* **/newname:***NewComputerName* **/userd:**[*Domain*\]*UserName* [**/passwordd:**[*Password* | *]] **/usero:**[*Domain*\]*UserName* [**/passwordo:**[*Password* | *]] [**/reboot[:***Delc*

**Parameters**

***Computer***
Specifies the name of the computer to rename.

**/newname: *NewComputerName***
Specifies the new name of the computer.

**/userd:[*domain*\]*UserName***
Specifies the user account to be used for the destination domain.

**/passwordd:[*password* | *]**
Specifies the password of the user account that is specified in the **/ud** or **/userd** parameter. If the wildcard character (*****) is used, then the user will be prompted for the password.

**/usero:[*domain*\]*UserName***
Specifies the user account to be used for the originating domain.

**/passwordo:[*password*| *]**
Specifies the password of the user account that is specified in the **/uo** or **/usero** parameter. If the wildcard character (*****) is used, then the user will be prompted for the password.

**/force[:*Delay*]**
The user will be prompted for confirmation unless the **/force** parameter is specified.

**/reboot[:*Delay*]**
Specifies that the computer shuts down and automatically reboots after the **rename** operation has completed. The *Delay* value is the number of seconds before automatic shutdown occurs. The default *Delay* value is 20 seconds.

**{/help | /?}**
Displays command-line usage for the **renamecomputer** operation.

Note

- Do not use **renamecomputer** to rename Windows Server 2003 or Windows 2000 domain controllers. Using the **renamecomputer** operation to rename a domain controller may result in the domain controller no longer functioning as a domain controller on the network. To rename Windows Server 2003 and Windows 2000 servers and domain controllers, use the **computername** operation.

**NetDom reset**
Resets the secure connection between a workstation and a domain controller.

**Syntax**

**NetDom reset** *Computer* {**/d: | /domain:**}*Domain* [{**/s: | /server:**}*Server*] [{**/uo: | /usero:**}*User* {**/po: | /passwordo**}{*Password*|*}] [{**/help | /?**}]

**Parameters**

***Computer***
Specifies the name of the computer for which the connection is to be reset.

**{/d: | /domain:}*Domain***
Specifies the domain with which to establish the secure connection. If this parameter is omitted, then the domain that the current computer belongs to is used.

**{/s: | /server:}*Server***
Specifies the name of the domain controller to use to establish the secure connection.

**{/uo: | /usero:}*User***
Specifies the user account that makes the connection with the computer to be reset. If this parameter is omitted, then the current user account is used.

**{/po: | /passwordo}{*Password*|*}**
Specifies the password of the user account that is specified in the **/uo** or **/usero** parameter. Use the wildcard character (*****) to be prompted for the password.

**{/help | /?}**
Displays command-line usage for the **reset** operation.

**Sample Usage**

```
netdom reset /d:reskit.ms.com redmond
```

**NetDom resetpwd**
Resets the machine account password for a Domain Controller.

**Syntax**

**NetDom resetpwd** {**/s: | /server:**}*Server* {**/ud: | /userd:**}[*Domain*\]*User* {**/pd: | /passwordd:**}{*Password*|*}] [{**/help | /?**}]

**Parameters**

**{/s: | /server:}*Server***
Specifies the name of the domain controller to use for setting the machine account password.

**{/ud: | /userd:}[*Domain*\]*User***
Specifies the user account that makes the connection with the domain specified in the **/s** parameter. This must be in *Domain**User* format. If this parameter is omitted, then the current user account is used.

**{/pd: | /passwordd:}{*Password*|*}**
Specifies the password of the user account that is specified in the **/ud** parameter. Use the wildcard character (*****) to be prompted for the password.

**{/help | /?}**
Displays command-line usage for the **resetpwd** operation.

**NetDom trust**
Establishes, verifies, or resets a trust relationship between domains.

**Syntax**

**NetDom trust** *TrustingDomainName* {**/d: | /domain:**} *TrustedDomainName* [{**/ud: | /userd:**}[*Domain*\]*User* [{**/pd: | /passwordd:**}{*Password*|*}] [{**/uo: | /usero:**}*User*] [{**/po: | /passwordo:**} {*Password*|*}] [**/verify**] [**/reset**] [**/passwordt:***NewRealmTrustPassword*] [**/add** [**/realm**]] [**/remove** [**/force**]] [**/twoway**] [**/kerberos**] [**/transitive**[**:**{**YES**|**NO**}]] [**/oneside:**{**TRUSTED** | **TRUSTING**}] [**/force**] [**/**
[**/togglesuffix:**#]] [**/EnableSIDHistory**] [**/ForestTRANsitive**] [**/SelectiveAUTH**][**/AddTLN**][**/AddTLNEX**][**/RemoveTLN**] [**/RemoveTLNEX**][{**/help | /?**}]

**Parameters**

***TrustingDomainName***
Specifies the name of the trusting domain.

**{/d: | /domain:}*TrustedDomainName***
Specifies the name of the trusted domain. If this parameter is omitted, then the domain that the current computer belongs to is used.

**{/ud: | /userd:}[*Domain*\]*User*]**
Specifies the user account that makes the connection with the domain specified in the **/d** or **/domain** parameter. If this parameter is omitted, then the current user account is used.

**/pd:{*Password*|*}**
Specifies the password of the user account that is specified in the **/ud** or **/userd:** parameter. Use the wildcard character (*****) to be prompted for the password.

**{/uo: | /usero:}*User***
Specifies the user account that makes the connection with the trusting domain. If this parameter is omitted, then the current user account is used.

**{/po: | /passwordo:}{ *Password*|*}**
Specifies the password of the user account that is specified in the **/uo** or **/usero** parameter. Use the wildcard character (*****) to be prompted for the password.

**/verify**
Verifies the secure channel secrets upon which a specific trust is based.

**/reset**
Resets the trust secret between trusted domains or between the domain controller and the workstation.

**/passwordt: *NewRealmTrustPassword***
Specifies a new trust password. This parameter is valid only with the **/add** parameter, and only if one of the domains specified is a non-Windows Kerberos realm. The trust password is set on the Windows domain only, which means that credentials are not needed for the non-Windows domain.

**/add**
Specifies to create a trust.

**/realm**
Indicates that the trust is created to a non-Windows Kerberos realm. The **/realm** parameter is valid only with the **/add** and **/passwordt** parameters.

**/remove**
Specifies to break a trust.

**/force**
Removes both the trusted domain object and the cross-reference object for the specified domain from the forest. Use this option to clean up decommissioned domains that are no longer in use, and cannot be removed by using the Active Directory Installation wizard. This problem can occur if the domain controller for that domain was disabled or damaged and there were no domain controllers, or if it was not possible to recover the domain controller from backup media. This parameter is valid only when the **/remove** parameter is specified.

**/twoway**
Specifies establishment ofa two-way trust relationship rather than a one-way trust relationship.

**/kerberos**
Specifies exercisingthe Kerberos protocol between a workstation and a target domain. This parameter is valid only when the **/verify** parameter is specified.

**/transitive[:{YES|NO}]**
Specifies whether to set a transitive or non-transitive trust. This parameter is valid only for a non-Windows Kerberos realm. Non-Windows Kerberos trusts are created as non-transitive. If no value is specified, then the current transitivity state is displayed.

| Value | Description |
|---|---|
| **YES** | Sets the realm to a transitive trust. |
| **NO** | Sets the realm to a non-transitive trust. |

**/oneside:{TRUSTED| TRUSTING}**
Denotes that the trust object should only be created or removed on one domain.

| Value | Description |
|---|---|
| **TRUSTED** | Indicates that the trust object is created or removed on the trusted domain specified by the **/d** or **/domain** parameter). |
| **TRUSTING** | Indicates that the trust object is to be created or removed on the trusting domain. Valid only with the **/add** or /REMove parameter. The **/passwordt** parameter is required when using with the /ADD or /REMove option |

**/quarantine[:{YES | NO}]**
Sets or clears the domain quarantine attribute. If no value is specified then the current quarantine state is displayed.

| Value | Description |
|---|---|
| **YES** | Specifies that only SIDs from the directly trusted domain will be accepted for authorization data returned during authentication. SIDS from any other domains will be removed. |
| **NO** | Specifies that any SID will be accepted for authorization data returned during authentication. This is the default value. |

**/namesuffixes: *TrustName***
Lists the routed name suffixes for *TrustName* on the domain named by *TrustingDomainName*. The **/usero** and **/passwordo** parameters can be used for authentication. The **/domain** parameter is not required.

**/togglesuffix:#**
Changes the status of a name suffix. Used with the **/namesuffixes** parameter. The number of the name entry specified by the **/namesuffixes** parameter must be provided to indicate which name will have its status changed. Names that are in conflict cannot have their status changed until the name in the conflicting trust is disabled. Always precede this command with the **/namesuffixes** parameter because LSA will not always return the names in the same order.

**/EnableSIDistory**
Specifying yes allows users who migrate to the trusted forest from any other forest to use SID history to access resources in this forest. Valid only for an outbound forest trust. This should be done only if the trusted forest administrators can be trusted enough to specify SIDs of this forest in the SID history attribute of their users appropriately. Specifying no would disable the ability of the migrated users in the trusted forest to use SID history to access resources in this forest. Specifying /EnableSIDHistory without yes or no will display the current state.

**/ForestTRANsitive**
Specifying yes marks this trust as forest transitive. Specifying no marks this trust as not forest transitive. Specifying **/ForestTRANsitive** without yes or no will display the current state of this trust attribute. Valid only for non-Windows real trusts and can only be performed on the root domain for a forest.

**/SelectiveAUTH**
Specifying no disables selective authentification across this trust. Specifying **/SelectiveAUTH** without yes or no displays the currrent state of this trust attribute. Specifying yes enables selective authentification across this trust. Valid only on outbound forest and external trusts.

**/AddTLN**
Adds the specified top level name (DNS name suffix) to the forest trust info for the specified trust. Valid only for a forest transitive non-Windows realm trust and can only be performed on the root domain for a forest. Refer to the **/NameSuffixes** operation for a list of name suffixes.

**/AddTLNEX**
Adds the specified top level name exclusion(DNS name suffix) to the forest trust info for the specified trust. Valid only for a forest transitive non-Windows realm trust and can only be performed on the root domain for a forest. Refer to the **/NameSuffixes** operation for a list of name suffixes.

**/RemoveTLN**
Removes the specified top level name (DNS name suffix) from the forest trust info from the specified trust. Valid only for a forest transitive non-Windows realm trust and can only be performed on the root domain for a forest. Refer to the **/NameSuffixes** operation for a list of name suffixes.

**/RemoveTLNEX**
Removes the specified top level name exclusion (DNS Name Suffix) from the forest trust info from the specified trust. Valid only for a forest transitive non-Windows realm trust and can only be performed on the root domain for a forest. Refer to the **/NameSuffixes** operation for a list of name suffixes.

**{/help | /?}**

Displays command-line usage for the **trust** operation.

**Sample Usage**

```
netdom trust /d:masterdom resourcedom
```

**NetDom verify**

Verifies the secure connection between a workstation and a domain controller.

**Syntax**

▸ **NetDom verify** *Computer* {**/d: | /domain:**}*Domain* [{**/uo: | usero**}*User* {**/po: | /passwordo**}{*Password*|*}] [{**/help | /?**}]

**Parameters**

***Computer***

Specifies the name of the computer whose secure connection is verified.

**{/d: | /domain:}*Domain***

Specifies the domain with which to verify the secure connection. If this parameter is omitted, then the domain that the current computer belongs to is used.

**{/uo: | usero}*User***

Specifies the domain with which to verify the secure connection. If this parameter is omitted, then the current user account is used.

**{/po: | /passwordo}{*Password*|*}**

Specifies the password of the user account that is specified in the **/uo** or **/usero** parameter. Use the wildcard character (*****) to be prompted for the password.

**{/help | /?}**

Displays command-line usage for the **verify** operation.

**Sample Usage**

```
netdom verify /d:resourcedom yourwksta
```

**See Also**

**Concepts**

Netdom Overview
Netdom Remarks
Netdom Examples
Spcheck Overview
Nltest Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Alphabetical List of Tools

# Nltest Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Nltest.exe: NLTest**
This command-line tool helps perform network administrative tasks. You can use NLTest to:

- Get a list of domain controllers.
- Force a remote shutdown.
- Query the status of trust.
- Test trust relationships and the state of domain controller replication in a Windows domain.
- Force a user-account database to synchronize on Windows NT version 4.0 or earlier domain controllers.

NLTest can test and reset the secure channel established by the NetLogon service. This secure channel is established between clients and the domain controller that logs them on. NLTest does not work for clients using Kerberos for authentication since this secure channel is not used with Kerberos.

### Note

- You must run NLTest from the command window.

**Corresponding UI**
There is no corresponding user interface for this tool.

Concepts

A discrete communication channel, known as the secure channel, exists between trusted domains in a Windows NT 4.0 environment and parent domains and their immediate children in an Active Directory environment. In a Windows NT 4.0 environment these channels are also used to authenticate user accounts when a remote user connects to a network resource and the user account exists in a trusted domain. This is called pass-through authentication.

This tool provides diagnostic features that are useful for troubleshooting Windows Server 2003 operating system configurations. However, because this tool has been designed primarily for system administrators and support personnel, its output may be difficult to analyze. In this case, you should review the appropriate troubleshooting sections in the Windows Deployment and Resource Kits. Search for any of the keywords from the bulleted list in the tool description above.

System Requirements

The following are the system requirements for NLTest:

- Windows Server 2003 or Windows XP Professional

**File Required**

- Nltest.exe

**See Also**
**Concepts**
Nltest Syntax
Nltest Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview

Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Nltest Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NLTest Examples**

Example 1: Verify DCs in a domain

In this example, the **/dclist** parameter is used to create a list of domain controllers of the domain fourthcoffee.com

```
nltest /dclist:fourthcoffee
```

Output displays similar to the following:

```
Get list of DCs in domain 'ntdev' from '\\fourthcoffee-dc-01'.
    fourthcoffee-dc-01.forthcoffee.com       [DS] Site: Rome
    fourthcoffee-dc-03.forthcoffee.com       [DS] Site: LasVegas
    fourthcoffee-dc-04.forthcoffee.com       [DS] Site: LA
    fourthcoffee-dc-09.forthcoffee.com       [DS] Site: NYC
    fourthcoffee-dc-12.forthcoffee.com       [DS] Site: Paris
    fourthcoffee-dc-24.forthcoffee.com       [DS] Site: Chattaroy
    fourthcoffee-dc-32.forthcoffee.com       [DS] Site: Haifa
    fourthcoffee-dc-99.forthcoffee.com       [DS] Site: Redmond
    fourthcoffee-dc-63.forthcoffee.com [PDC] [DS] Site: London
The command completed successfully
```

Example 2: Advanced information about users

In this example, you want to find out detailed information about a certain user. At the command prompt, type:

```
nltest /user:"TestAdmin"
```

Output displays similar to the following:

```
User: User1
Rid: 0x3eb
Version: 0x10002
LastLogon: 2ee61c9a 01c0e947 = 5/30/2001 13:29:10
PasswordLastSet: 9dad5428 01c0e577 = 5/25/2001 17:05:47
AccountExpires: ffffffff 7fffffff = 9/13/30828 19:48:05
PrimaryGroupId: 0x201
UserAccountControl: 0x210
CountryCode: 0x0
CodePage: 0x0
BadPasswordCount: 0x0
LogonCount: 0x33
AdminCount: 0x1
SecurityDescriptor: 80140001 0000009c 000000ac 00000014 00000044 00300002 000000
02 0014c002 01050045 00000101 01000000 00000000 0014c002 000f07ff 00000101 05000
000 00000007 00580012 00000003 00240000 00020044 00000501 05000000 00000015 22cd
b7b4 7112b3f1 2b3be507 000003eb 00180000 000f07ff 00000201 05000000 00000020 000
00220 00140000 0002035b 00000101 01000000 00000000 00000201 05000000 00000020 00
000220 00000201 05000000 00000020 00000220
AccountName: User1
Groups: 00000201 00000007
LmOwfPassword: fb890c9c 5c7e7e09 ee58593b d959c681
NtOwfPassword: d82759cc 81a342ac df600c37 4e58a478
NtPasswordHistory: 00011001
LmPasswordHistory: 00010011
The command completed successfully
```

The detailed information provided can be used to troubleshoot many issues.

Example 3: Verify trust relationship with a specific server

In this example, you want to verify that the server a-dc1 has a valid trust relationship with the domain. At the command prompt, type:

```
nltest.exe /server:fourthcoffee-dc-01 /sc_query:fourthcoffee
```

Output displays similar to the following:

```
Flags: 30 HAS_IP  HAS_TIMESERV
Trusted DC Name \\fourthcoffee-dc-01.forthcoffee.com
Trusted DC Connection Status Status = 0 0x0 NERR_Success
The command completed successfully
```

Example 4: Determine the PDC emulator for a domain

In this example, you want to determine which DC in your domain the Windows NT 4.0–based computers are looking to as the PDC. At the command prompt, type:

```
nltest /dcname:fourthcoffee
```

Output displays similar to the following:

```
PDC for Domain fourthcoffee is \\fourthcoffee-dc-01
The command completed successfully
```

You can see that a-dcp is the PDC emulator for your domain.

Example 5: Show trust relationships for a domain

In this example, you want to view the established trust relationships for your domain. At the command prompt, type:

```
nltest /domain_trusts
```

Output displays similar to the following:

```
List of domain trusts:
    0: forthcoffee forthcoffee.com (NT 5) (Forest Tree Root) (Primary Domain)
The command completed successfully
```

This example shows that one domain is trusting itself and no others.

**See Also**
**Concepts**
Nltest Overview
Nltest Syntax
Alphabetical List of Tools
Spcheck Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Nltest Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NLTest Syntax**

NLTest uses the following syntax:

- **nltest** [**/server:***servername*] [*operation*[*parameter*] ...

**/server: *ServerName***

Runs NLTest at the specified remote domain controller. If this parameter is not specified, the command is executed on the local computer (domain controller).

**Operations**

**/query**

Reports on the state of the secure channel the last time it was used. This is the secure channel established by the NetLogon service.

**/repl**

Forces a synchronization with the PDC. Only changes not yet replicated to the BDC are synchronized. This command is for NT 4.0 BDCs only and is not for Active Directory replication. Administrative rights are required to perform this command.

**/sync**

Forces an immediate synchronization with the PDC of the entire SAM database. This command is for NT 4.0 BDCs only and is not for Active Directory replication. Administrative rights are required to perform this command.

**/pdc_repl**

Forces the PDC to send a synchronize notification to all BDCs. This command is for NT 4.0 PDCs only and is not for Active Directory replication. Administrative rights are required to perform this command.

**/sc_query: *DomainName***

Reports on the state of the secure channel the last time it was used. This is the secure channel established by the NetLogon service. Also, lists the name of the domain controller that was queried on the secure channel.

**/sc_reset:[ *DomainName*]**

Removes and then rebuilds the secure channel established by the NetLogon service. Administrative rights are required to perform this command.

**/sc_verify:[ *DomainName*]**

Checks the status of the secure channel established by the NetLogon service. If the secure channel is not working, this operation removes the existing channel and builds a new one. Administrative rights are required to perform this command. This operation is only valid on Windows 2000 with Service Pack 2 and Windows Server 2003 domain controllers.

**/sc_change_pwd:[ *DomainName*]**

Changes the password for the trust account of the specified domain. If this command is run on a domain controller, and an explicit trust relationship exists, then the password for the interdomain trust account is reset. Otherwise, the computer account password for the specified domain is changed. This command is only for computers that are Windows 2000, Windows XP, and Windows Server 2003.

**/dclist:[ *DomainName*]**

Lists all domain controllers in the domain. In an NT 4.0 domain environment, this command uses the Browser service to retrieve the list of domains. In an Active Directory environment, this command first queries Active Directory for a list of domain controllers. If this is unsuccessful the Browser service is used.

**/dcname:[ *DomainName*]**

Lists the primary domain controller or the primary domain controller emulator for *DomainName*.

**/dsgetdc:[ *DomainName*]**

Queries the DNS server for a list of domain controllers and their corresponding IP addresses. Contacts each domain controller to check for connectivity.

Use the following flags to filter the list of domain controllers or specify alternate names types in the syntax.

- **/PDC** Returns only the PDC (NT 4.0) or domain controller designated as the PDC emulator (Windows 2000 or Windows Server 2003).
- **/DS** Returns only those domain controllers that are Windows 2000 or Windows Server 2003 servers.
- **/DSP** Requests that Windows 2000 or Windows Server 2003 domain controllers be returned. If no Windows 2000 or Windows Server 2003 server is found, then this operation returns NT domain controllers.
- **/GC** Returns only those domain controllers designated as Global Catalog servers.
- **/KDC** Returns only those domain controllers designated as Kerberos key distribution centers.
- **/TIMESERV** Returns only those domain controllers designated as time servers.
- **/GTTIMESERV** Returns only those domain controllers designated as master time servers.
- **/NetBIOS** Use this command when specifying computer names in the syntax as NetBIOS names.
- **/DNS** Use this command when specifying computer names in the syntax as FQDNs.
- **/IP** Returns only domain controllers that have IP addresses. Domain controllers not using TCP/IP as their protocol stack are not returned.
- **/FORCE** Forces the computer to run the command against the DNS server instead of looking in cache for the information.

**/dnsgetdc: *DomainName***

Queries the DNS server for a list of domain controllers and their corresponding IP addresses.

Use the following flags to filter the list of domain controllers.

- **/PDC** Returns only only those domain controllers that are PDCs (NT 4.0) or designated as PDC emulators.
- **/GC** Returns only those domain controllers designated as Global Catalogs.
- **/KDC** Returns only those domain controllers designated as Kerberos key distribution centers.
- **/WRITABLE** Returns only those domain controllers that can accept changes to the directory database. All Active Directory domain controllers will be returned. Only NT 4.0 BDCs will not be returned with this command.
- **/LDAPONLY** Returns servers that are running an LDAP application. With this command LDAP servers are returned that are not necessarily DCs.
- **/FORCE** Forces the computer to run the command against the DNS server instead of looking in cache for the information.
- **/SITE:***Sitename* Sorts the returned records so that the ones pertaining to the site are listed first.
- **/SITESPEC** Filters the returned records so only those pertaining to the site are displayed. This operation can only be used with the **/SITE** operation.

**/dsgetfti: *DomainName*[ /UpdateTDO]**

Returns information about interforests trust(s). This operation is only for a Windows Server 2003 domain controller that is in the root of the forest. If no interforest trusts exist, this operation will return an error.

**/UpdateTDO** Updates the locally stored information on the interforest trust.

**/dsgetsite**
Returns the name of the site in which the domain controller resides.

**/dsgetsitecov**
Returns the name of the site that the domain controller covers. A domain controller can cover a site that has no local domain controller of its own.

**/parentdomain**
Returns the name of the parent domain of the server.

**/dsregdns**
Refreshes the registration of all domain controller-specific DNS records.

**/dsderegdns: *DnsHostName***
Deregisters DNS host records from DNS for the host specified with the *DnsHostName* parameter./DOM: /DOMGUID: /DSAGUID:.

Use the following flags to specify which records will be deregistered.

- **/DOM:** Specifies a DNS domain name for the host to use when searching for records in the DNS server. If not specified, the suffix of the DnsHostName is assumed to be the DNS domain name
- **/DOMGUID:** Deletes DNS records that are GUID based.
- **/DSAGUID** Deletes DSA records that are GUID based.

**/whowill: *Domain/ User***
Finds the domain controller that has the specified user account. Use this command to determine whether the account information has been replicated to other domain controllers.

**/finduser: *User***
Finds which directly trusted domain to which the specified user account belongs. Use this operation to troubleshoot logon issues of older client operating systems.

**/transport_notify**
Flushes the negative cache to force the discovery of a domain controller. Use this operation on NT 4.0 domain controllers only. This operation is done automatically when clients log on to Windows 2000 and Windows Server 2003 domain controllers.

**/dbflag: *HexadecimalFlags***
Sets a new debug flag. For most purposes, use 0x2000FFFF as the value for *HexadecimalFlags*. The entry in the Windows Server 2003registry for debug flags is HKLM\System\CurrentControlSet\Services\Netlogon\Parameters\DBFlag.

**/user: *UserName***
Displays many of the attributes for the specified user account that are maintained in the SAM account database. This operation will not work for user accounts stored in Active Directory.

**/time: *HexadecimalLSL HexadecimalMSL***
Converts Windows NT GMT time to ASCII. *HexadecimalLSL* is a hexadecimal value for least significant longword. *HexadecimalMSL* is a hexadecimal value for most significant longword.

**/logon_query**
Queries the cumulative number of NTLM logon attempts at the console or over the network.

**/domain_trusts**
Returns a list of trusted domains. /Primary /Forest /Direct_Out /Direct_In /All_Trusts /v.

Use the following flags to filter the list of domains.

- **/Primary** Returns only the domain to which the computer account belongs.
- **/Forest** Returns only those domains that are in the same forest as the primary domain.
- **/Direct_Out** Returns only the domains that are explicitly trusted with the primary domain.

- **/Direct_In** Returns only the domains that explicitly trust the primary domain.
- **/All_Trusts Returns all trusted domains.**
- **/v** Displays verbose output including domain SIDs and GUIDs if available.

**/dsquerydns**
Queries for the status of the last update for all DC-specific DNS records.

**/bdc_query: *DomainName***
Queries for a list of backup domain controllers in *DomainName* and displays their state of synchronization and replication status. This operation is only for NT 4.0 domain controllers.

**/sim_sync: *DomainName ServerName***
Simulates full synchronization replication. This operation is useful in test environments.

**/list_deltas: *FileName***
Displays the contents of the change log file *FileName*, which lists changes to the user account database. Netlogon.chg is the default name. This log file resides only on NT 4.0 BDCs.

**/cdigest: *Message* /domain: *DomainName***
Displays the current digest (calculation derived from the password) used by the client for the secure channel. Also, displays the digest based on the previous password. The secure channel is used for logons between client computers and a domain controller, or DC to DC for directory service replication. Use this operation in conjunction with the **/sdigest** operation to check trust account password synchronization.

**/sdigest: *Message* /rid: *RID_In_Hexadecimal***
Displays the current digest (calculation derived from the password) that the server is using for the secure channel. Also, displays the digest for the previous password. If the digest from the server matches the digest from the client (retrieved using the **/cdigest** operation), then the passwords used for the secure channel are synchronized. If the digests do not match, then a password change may not have replicated yet.

**/shutdown: *Reason*[ *Seconds*]**
Performs a remote shutdown of the *ServerName* for *Reason*, a string, after *Seconds*, an integer. For a complete description, see the Platform SDK documentation for **InitiateSystemShutdown**.

**/shutdown_abort**
Terminates a system shutdown.

**See Also**
**Concepts**
Nltest Overview
Nltest Examples
Alphabetical List of Tools
Spcheck Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview

# Portqry Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Portqry.exe: Port Query**
This command-line tool reports the status of TCP and UDP ports on a target computer. PortQry is used to troubleshoot TCP/IP connectivity issues. It provides an additional level of detail on port status not provided by other port scanning tools. You can use PortQry to query a single port, an ordered list of ports, or a sequential range of ports.

Note

- PortQry is a troubleshooting tool, not a security assessment tool.

**Concepts**
A TCP or UDP port is a mechanism that allows multiple sessions to use one IP address.

Port scanning tools search for and identify processes engaged at the port level. These processes may originate from, for example, operating system services such as DNS or WINS, or third-party applications or services. Generally, port scanning tools report only whether a port is busy or not, that is, whether there is a process on it. PortQry provides specific status information by analyzing the response from a process on a target port.

See the appendix from the Windows Server 2003 Internetworking Guide for information on TCP and UDP Ports and Protocol Numbers.

**System Requirements**

- Windows 2000, Windows XP, or Windows Server 2003
- A TCP/IP network connection

**File Required**

- Portqry.exe

**See Also**
**Concepts**
Portqry Remarks
Portqry Syntax
Portqry Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview

Ksetup Overview
Getsid Overview
Addiag.exe

# Portqry Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PortQry Remarks**

**Special Features**
PortQry improves on typical port scanning utilities by giving additional UDP-port status detail, and by providing protocol-specific support.

Determining the status of UDP ports depends on a response from the process that is listening on each UDP port. If an unformatted zero-length or fixed-length message is sent to a target UDP port, the port may not respond.

If the target port does respond, it is characterized as listening. If an "ICMP destination unreachable" message is returned from the port, the port is characterized as not listening. However, some port scanning utilities report that the port is listening simply if an "ICMP destination unreachable" message is *not* returned from a target port. This may, in fact, be inaccurate, because no response to a directed datagram may also indicate that the target port is being *filtered*. Usually only a properly-formatted message using the session layer or application layer protocol that the listening service or application understands, elicits a response from the target port.

When you are troubleshooting a connectivity problem, especially in an environment with firewalls, it is useful to know whether a port is first being filtered, or if it is actually listening. PortQry pinpoints this distinction on selected ports.

If there is no response from a target UDP port, PortQry reports that the port is "Listening or Filtered." Then, PortQry sends a properly-formatted message, using the session layer or application layer protocol that the listening service or application understands, to determine if the port is in fact listening.

PortQry supports the following session/application layer protocols

**NetBIOS Name Service**
By default, this service listens on UDP port 137. If PortQry receives no response from port 137, it reports that the port is "Listening or Filtered," and then determines whether the port is actually listening. If NetBIOS is available on the computer that is running PortQry, PortQry sends a NetBIOS adapter status query to the target computer. If the target computer responds to the query, PortQry reports the port status as Listening and returns the target computer's Media Access Control (MAC) address to the user.

**DNS**
PortQry can also send a properly formatted DNS query, by using UDP and TCP. PortQry first sends a DNS query, then waits for a response from the server. Whether the DNS response to the query is negative or positive is irrelevant, because any response indicates to PortQry that the port is listening.

**RPC**
PortQry can send a query to the RPC endpoint mapper (by using UDP and TCP) and interpret the response. This query dumps all of the endpoints that are currently registered with the RPC endpoint mapper. The response from the endpoint mapper is parsed, formatted, and returned to the user. Portqry Examples.

**LDAP**
PortQry can send an LDAP query (by using UDP and TCP) and interpret an LDAP server's response to the query. The response from the LDAP server is parsed, formatted, and returned to the user.

To query LDAP, PortQry automatically resolves UDP port 389 by using the %SystemRoot%\System32\Drivers\Etc\Services file that every Windows Server 2003, Windows XP, and Windows 2000-based computer has by default. If PortQry resolves the port to the LDAP service, it sends an unformatted user datagram to UDP port 389 on the target computer. However, PortQry does not receive a response from the port because the LDAP service responds only to a properly-formatted LDAP query. So, again, PortQry initially reports that the port is "Listening or Filtered."

PortQry then sends a properly-formatted LDAP query to UDP port 389. If PortQry receives a response to this query, it returns the entire response to the user, and reports that the port is Listening. If PortQry does not receive a response to the query, it reports that the port is Filtered. Portqry Examples.

**Port Status Reporting**
PortQry characterizes port status as one of the following three states:

- **Listening**

  A process is listening on the target port. PortQry received a response from the port.

- **Not Listening**

  No process is listening on the target port. PortQry received an Internet Control Message Protocol (ICMP) "Destination Unreachable - Port Unreachable" message back from the target UDP port. Or, if the target port is a TCP port, PortQry received a TCP Acknowledgement packet with the Reset flag set.

- **Filtered**

  The target port on the target computer is being filtered. PortQry dides not receive a response from the target port. A process may or may not be listening on the port. By default, TCP ports are queried three times; UDP ports are queried once before reporting that the target port is Filtered.

While PortQry supports querying a range of ports, its performance has not been optimized for security assessment. In particular, querying large ranges of TCP ports can be time-consuming and is not recommended.

**PortQry Errorlevels**
When using the **/q** parameter with PortQry from within a batch file, PortQry returns errorlevels based on the status of the target port. PortQry returns an errorlevel 0 if the port is Listening, 1 if the port is Not Listening, and 2 if the port is Listening or Filtered.

**The Services File**
In order to send protocol-specific queries, PortQry must first determine what protocol to use. On computers running Windows 2000 or later, this information is stored in the Services file, located in the %SystemRoot%\system32\drivers\etc directory. PortQry also uses the Services file to resolve the port name.

**Using PortQry in a Batch File**
You can use the **/q** parameter when running PortQry from a batch file, to report errorlevels based the status of the target port. PortQry returns an errorlevel 0 if the port is listening, 1 if the port is not listening, and 2 if the port is listening or filtered.

**Querying Filtered TCP Ports**
PortQry was developed as a trouble-shooting tool, rather than a security assessment tool. While PortQry supports querying a range of ports (by using the **/r** parameter), its performance has not been optimized to handle querying large ranges of TCP ports that may be filtered.

**See Also**
**Concepts**
Portqry Overview
Portqry Syntax
Portqry Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview

Ksetup Overview
Getsid Overview
Addiag.exe

# Portqry Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PortQry Examples**

**Example 1: PortQry Default Behavior**

To see the results of PortQry executed in its simplest form, type:

```
portqry /n 10.193.36.210
```

This command instructs PortQry to query TCP port 80, and determine whether it is listening on the target host (in this case 10.193.36.210).

PortQry displays results similar to the following output:

```
Querying target system called:

 10.193.36.210

Attempting to resolve IP address to a name...

IP address resolved to MYSERVER


TCP port 80 (http service): LISTENING
```

**Example 2: Querying RPC Ports**

To see the results of PortQry executed on an RPC port, type:

```
portqry /n 10.193.36.210 /p udp /e 135
```

PortQry displays results similar to the following output:

```
Querying target system called:

 10.193.36.210

Attempting to resolve IP address to a name...

IP address resolved to RKTLABDC2


UDP port 135 (epmap service): LISTENING or FILTERED
Querying Endpoint Mapper Database...
Server's response:

UUID: a00c021c-2be2-11d2-b678-0000f87a8f8e PERFMON SERVICE
ncacn_np:\\\\RKTLABDC2[\\pipe\\000003b8.000]

...

UUID: d049b186-814f-11d1-9a3c-00c04fc9b232 NtFrs API
ncacn_np:\\\\RKTLABDC2[\\pipe\\000003b8.000]

Total endpoints found: 69



==== End of RPC Endpoint Mapper query response ====

UDP port 135 is LISTENING
```

**Example 3: Querying an LDAP Port**
To see the results of PortQry executed on an LDAP port, type:

```
portqry /n 10.193.36.210 /p udp /e 389
```

For more information on how PortQry works on an LDAP port, see Portqry Remarks.

PortQry displays results similar to the following output:

```
Querying target system called:

 10.193.36.210

Attempting to resolve IP address to a name...

IP address resolved to RKTLABDC2


UDP port 389 (unknown service): LISTENING or FILTERED

Sending LDAP query to UDP port 389...

LDAP query response:


currentdate: 04/19/2002 01:32:58 (unadjusted GMT)
subschemaSubentry: CN=Aggregate,CN=Schema,CN=Configuration,DC=mydom,DC=com
dsServiceName: CN=NTDS Settings,CN=MYSERVER,CN=Servers,CN=Default-First-Site-
Name,CN=Sites,CN=Configuration,DC=mydom,DC=com
namingContexts: CN=Schema,CN=Configuration,DC=rktlabdom,DC=com
defaultNamingContext: DC=mydom,DC=com
schemaNamingContext: CN=Schema,CN=Configuration,DC=rktlabdom,DC=com
configurationNamingContext: CN=Configuration,DC=rktlabdom,DC=com
rootDomainNamingContext: DC=mydom,DC=com
supportedControl: 1.2.840.113556.1.4.319
supportedLDAPVersion: 3
supportedLDAPPolicies: MaxPoolThreads
highestCommittedUSN: 11019
supportedSASLMechanisms: GSSAPI
dnsHostName: myserver.mydom.com
ldapServiceName: mydom.com:myserver$@MYDOM.COM
serverName: CN=MYSERVER,CN=Servers,CN=Default-First-Site-
Name,CN=Sites,CN=Configuration,DC=mydom,DC=com
supportedCapabilities: 1.2.840.113556.1.4.800
isSynchronized: TRUE
isGlobalCatalogReady: TRUE


======== End of LDAP query response ========

UDP port 389 is LISTENING
```

**See Also**
**Concepts**
Portqry Overview
Portqry Remarks
Portqry Syntax
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview
Httpcfg Overview
Dnslint Overview

Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Portqry Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Portqry Syntax**

PortQry uses the following syntax:

- **portqry** /**n** name [/**p** protocol] [/**e** || /**r** || /**o** endpoint(s)] [/**l** logfile] [/**s**] [/**q**] [/**i**]

**Parameters**

**/n name**

A required parameter used to specify the target system by host name or IP address. No embedded spaces are allowed in either. If PortQry attempts to resolve an IP address from a host name and cannot, then it reports an error and quits. If PortQry attempts to resolve a host name from an IP address and cannot, then it reports an error, but continues processing.

**/p *protocol***

An optional parameter that specifies the type of port or protocol used to connect to the target port. Valid entries for protocol are *tcp*, *udp*, or *both*. The default is TCP. If you set the option to BOTH, you must also specify the port number by using the **/e** parameter. PortQry then sends queries to both the TCP and UDP port specified by the **/e** parameter.

**/e *port***

An optional parameter. Use **/e** to query a single port within the valid range of 1-65535. The **/e** parameter cannot be used with **/r** or **/o**.

**/r *start port:end port***

An optional parameter. Use **/r** to specify a range of port numbers to be queried in sequential order within the valid range of 1-65535. No spaces are allowed between the port numbers and the colon. The **/r** parameter cannot be used with **/o** or **/e**.

**/o *port,port,port...***

An optional parameter. Use **/o** to query a series of port(s) in a specific order (not sequential), within the valid range of 1-65535. No spaces are allowed between port numbers and commas. The **/o** parameter cannot be used with **/r** or **/e**.

**/l *filename***

An optional parameter. Use **/l** to copy the PortQry output to a text (log) file in the directory where you are running PortQry. Filename is the name (with extension) of the log file. If a log file with the same name already exists, PortQry prompts you to specify whether you want the existing file overwritten.

**/s**

An optional parameter. Use **/s** to request that Portqry wait longer for responses to UDP queries.

**/q**

An optional parameter. Use **/q** to suppress all output to the screen other than error messages. The **/q** parameter can be used with **/e**, but not with **/o**, **/r**, or **/p** with the BOTH option. For more information about errorlevel reporting, see Portqry Remarks.

**/i**

Bypasses default IP address-to-host-name lookup. This parameter is ignored unless an IP address is specified using the /**n** parameter. When an IP address is specified using /**n**, PortQry attempts to resolve the IP address to a host name. This process can be time consuming, particularly if the IP address cannot be successfully resolved. Specifying /**i** in this case allows you to bypass the default lookup, and instead, query the target port(s) immediately.

**See Also**

**Concepts**

Portqry Overview
Portqry Remarks
Portqry Examples
Alphabetical List of Tools
Spcheck Overview
Nltest Overview
Netdom Overview
Netdiag Overview
Netcap Overview

Httpcfg Overview
Dnslint Overview
Dnscmd Overview
Dhcploc Overview
Dcdiag Overview
Browstat Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Performance Monitoring Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Monitoring and Status Tools**
Tools in this category are used to manage the following Windows features:

- Device Manager
- Event Viewer
- Network monitoring
- Performance
- SNMP

**Notes**

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Exctrlst Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Exctrlst.exe: Extensible Performance Counter List**
This GUI tool lists all services and applications that provide registry-based performance counters on local and remote Windows 2000, Windows XP Professional, and Windows Server 2003 systems. It displays the names and locations of the dynamic link libraries (DLLs) that support performance counters, indicates disabled counters, and lets you enable and disable the performance counters for a service.

Unlike most performance tools, which load performance counters in order to collect data, Exctrlst displays information directly from the Performance subkeys in the registry (HKLM\SYSTEM\CurrentControlSet\Services*ServiceName*\Performance). As a result, Exctrlst lists all registered performance DLLs, even those that do not load or operate properly.

Performance counter developers can use this tool to find the DLLs, basic functions, and indexes of performance counters provided by a service. Administrators can use this tool to determine which services have installed performance counters on the system and to identify counters that are disabled or do not load properly.

Exctrlst displays the following information about performance counters:

- Services and applications that provide performance counters
- Names and locations of performance counter DLLs
- Enabled and disabled performance counters
- Indexes of performance objects, counter names, and help text provided by the service
- Names of the open, collect, and close procedure of each counter DLL
- The highest index on the system used for performance counters and their help text.

**Corresponding UI**
Exctrlst provides its own user interface.

**Concepts**
Windows collects data about system resources, such as disks, memory, processors, and network components. In addition, applications and services that you might be running on your system can also perform data collection. By default, the operating system obtains performance data for system resources using the registry.

When you use performance tools to access registry functions for performance data, the system collects the data from the appropriate system object managers, such as the Memory Manager, the input/output (I/O) subsystem, and so forth.

**System Requirements**
The following are the system requirements for Exctrlst:

- Windows Server 2003 or Windows XP Professional

**For More Information**
For information about counter types and for descriptions of the performance counters installed on Windows Server 2003, see the Performance Counters Reference of the Windows Server 2003 Resource Kit.

For information about the run-time reliability tests for performance counters, and to customize the tests for your system, search for "Configuration Flags" in the Registry Reference of the Windows Server 2003 Resource Kit.

For instructions on writing performance counter DLLs and loading performance counter data in the registry, see Adding Performance Counters in the Platform Software Development Kit (SDK).

For a list of performance counters that load and operate correctly, use Typeperf, System Monitor, or Showperf.

To see a graphic display of performance counter data, use System Monitor. To run System Monitor, at the command prompt or in the **Run** dialog box, type **Perfmon**.

To list the performance counters on a local or remote system from the command line, use Typeperf, a tool included in Windows XP Professional and Windows Server 2003. Typeperf lists only the counters that are enabled and load and operate properly. For more information about Typeperf, see "Typeperf" in Help for Windows Server 2003.

**File Required**

- Exctrlst.exe

**See Also**
**Concepts**
Exctrlst Remarks
Exctrlst UI
Exctrlst Examples
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Exctrlst Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Exctrlst Remarks**

**Exctrlst Data Source**

Exctrlst displays data directly from the Performance subkeys in the registry (HKLM\SYSTEM\CurrentControlSet\Services*ServiceName*\Performance). A Performance subkey is created when a service, application, or device registers a performance DLL.

Exctrlst provides a complete list of all performance counter DLLs registered on the system, including counter DLLs that do not load or operate properly, and disabled counter DLLs. Exctrlst is often used with tools that attempt to load and run the performance counters, such as Typeperf or System Monitor. By comparing the list of all counter DLLs in Exctrlst with the list of operating counter DLLs in System Monitor or Typeperf, you can detect counters that are not loading or operating properly. Typeperf is especially useful because it lists counters in the order in which they load (alphabetical order by service name), making it easier to detect which counter failed to load.

**Extensible Performance Counters**

In Windows NT and earlier systems, Exctrlst only displayed information about *extensible* performance counters, that is, custom counters that are not built into the operating system. However, in Windows 2000, Windows XP Professional, and Windows Server 2003, system counters and extensible counters are loaded and stored in the registry in the same way, and Exctrlst displays both types.

**Performance Counters for a Service**

Each performance counter DLL can support counters for one or more components and services. In the registry and, therefore, in Exctrlst, each DLL is listed in the subkey of the service that installed it.

**System Counters**

*System counters* are the performance counters that are built into the operating system. These counters appear in Exctrlst. You can use Exctrlst to enable and disable the system counters. However, because system counter names and explain text are not stored and accessed through the registry, their indexes appear as 0x0 or N/A in Exctrlst.

In Windows XP Professional and Windows Server 2003, system counters are defined in the following libraries and drivers:

| Service | Library | Objects |
|---|---|---|
| PerfOS | PerfOs.dll | Cache, Memory, Object, Paging File, Processor, System |
| PerfDisk | PerfDisk.dll | LogicalDisk, PhysicalDisk |
| PerfNet | PerfNet.dll | Browser, Redirector, Server |
| PerfProc | PerfProc.dll | Job Object, Job Object Details, Process, Thread, Thread Details |
| Spooler | Winspool.drv | Print Queue |
| TapiSrv | TapiPerf.dll | Telephony |
| Tcpip | Perfctrs.dll | IP, ICMP, NBT Connection, Network Interface, TCP, UDP |

**WMI Performance Counters**

Exctrlst lists all registry-based performance counter DLLs, including counter DLLs that are loaded by using the Windows Management Instrumentation High-Performance provider, Wbemperf.dll. However, because native WMI performance counters do not create Performance subkeys in the registry, Exctrlst cannot display information about these counters.

**Monitoring Remote Computers**

You can use Exctrlst to view the performance counters installed on remote computers. However, you must have permission to read the required files, you must have permission to read from the registry remotely, and the Remote Registry Service must be running on the remote computer. To enable or disable performance counters, you must also have permission to set values in the registry. For detailed instructions on configuring remote permissions, see *Monitoring Remote Computers* in the Windows XP Professional Resource Kit.

If Exctrlst cannot connect to a remote computer (for any reason), then it continues to monitor the computer that was

monitoring before the new computer name was entered. The name in the **Machine Name** box reverts to the name of the previously monitored computer, but Exctrlst does not display an error message.

**Disabled Performance Objects**
You can use Exctrlst to determine whether the performance counters for a service are enabled or disabled, and to enable and disable the performance counters for a service.

The system disables performance counters automatically when the performance counter DLL for a service does not pass the reliability tests that the system runs when it loads a counter in System Monitor or another performance monitoring tool.

Typically, users disable counters during testing or to prevent a malfunctioning counter DLL from loading. After replacing or repairing a counter DLL, you can use Exctrlst to re-enable the counters.

To disable the performance objects for a service, the system adds the **Disable Performance Counters** registry entry to the Performance subkey for the service (HKLM\SYSTEM\CurrentControlSet\Services*ServiceName*\Performance) and sets the value of the entry to 1.

To re-enable performance objects after the DLL is replaced or repaired, use Exctrlst, or set the value of the **Disable Performance Counters** registry entry to 0 or delete the entry.

For more information about the **Disable Performance Counters** registry entry, see the Registry Reference in the Windows XP Professional Resource Kit.

**Tip**
Event Viewer is one of the most valuable tools for diagnosing problems with performance counters. When the system detects a serious error in a performance counter DLL, it disables the counter and records Event 1016 in the Application Log in Event Viewer. If a performance counter does not appear in System Monitor, or if it does not operate properly, check the Application log for events generated by Perflib.

**See Also**
**Concepts**
Exctrlst Overview
Exctrlst UI
Exctrlst Examples
Alphabetical List of Tools
Pviewer Overview
Devcon Overview
Apmstat Overview

# Exctrlst Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Exctrlst Examples**

**Example 1: Find the performance objects installed on a remote computer**

The following procedure uses Exctrlst to find the performance counter objects installed on a remote computer, RK-USER-16. You can use this procedure to determine whether the performance features are installed correctly on a user's system, and you can use the information as an indirect indication of whether a service is installed on the user's system.

In this case, a user is complaining that he cannot monitor the Terminal Services performance counters. You have verified that the counters do not appear in System Monitor.

1. Start Event Viewer, connect to the RK-USER-16 computer, and search the Application Log for events recorded by Perflib. In this case, none of the events report a problem with the Terminal Services counters.

2. Start Exctrlst.

3. In the **Machine Name** box, type `RK-USER-16`, and then click **Refresh**.

4. In the **Sort Order** section, click **Service**.

The **Extensible Performance Counter** box now lists all of the performance counters installed on the remote computer. In this case, the Terminal Services service, TermService service and its DLL, Perfts.dll, do not appear in the list. This indicates that Terminal Services performance counters are not installed on the user's computer.

Note

- You cannot copy or export a list from Exctrlst. However, you can use the **query** operation of Reg.exe, a command-line tool in %windir%\system32, to list all services that have installed performance counters on a system.

  At the command prompt, type:

  ```
  reg query HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services /v Performance /s > PerfServices.txt
  ```

  This command creates a file called PerfServices.txt with the names of all services that have Performance subkeys in the registry. It also lists occurrences of the **Disable Performance Counters** registry entry and its value.

**Example 2: Find a missing performance object**

The following procedure determines why the performance object for a fictitious service, Reskit, does not appear in System Monitor. This procedure compares performance object lists from Exctrlst and Typeperf, a tool included in Windows XP Professional and Windows Server 2003. Because Typeperf only displays objects that load and run properly, objects that appear in Exctrlst but not in Typeperf are likely to have errors.

1. Check the Application Log in Event Viewer for events recorded by Perflib. When a counter DLL fails the reliability tests of the system or does not install properly, Perflib records errors in the log.

   In this case, there are no events pertaining to the Reskit service.

2. Use Exctrlst to verify that the counter is installed on the system.

   If a counter does not appear in the Exctrlst display, then it is not installed or properly registered in the system. In this case, the Reskit service is listed in Exctrlst, and its performance counter DLL, Reskit.dll, is listed in its proper location.

3. Use Exctrlst to determine whether the counter DLL for the Reskit service is disabled.

   In this case, when you click the Reskit service item, the **Performance Counters Enabled** check box is checked, indicating

that the Reskit service performance DLL passed the run-time reliability tests of the system.

4. Use Typeperf to load and display all properly loading performance counters on the system. Then, look for the **Reskit** performance object in the Typeperf list.

   The following command writes a list of all properly loading performance counters to the C:\Counters.txt file:

   ```
   typeperf -q -o C:\Counters.txt
   ```

   In this case, the Reskit object does not appear in the list. This indicates that the Reskit performance DLL passed the basic reliability tests, but does not load properly in Typeperf. This rare situation occurs when a DLL has a fault that is not evident in basic testing but appears later, such as when the service is stressed by high use or retrieves extreme data values not anticipated in the program design.

Counter DLLs might also be enabled but not load properly in Typeperf because one or more of the run-time reliability tests have been turned off, or because the system is prevented from disabling performance counters that fail its tests. To check for this occurrence, find the value of the **Configuration Flags** registry entry (HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib\Configuration Flags). The default value, which permits the system to disable faulty counters, is 0x10.

For information about the **Configuration Flags** registry entry, search the Registry Reference of the Windows Server 2003 Resource Kit.

**Example 3: Re-enable performance counters**
The following procedure re-enables the performance counter DLL for Reskit, a fictitious service.

To protect the performance monitoring capabilities of the system, the Performance Library (Perflib) tests all performance DLLs when they load and disables performance DLLs that fail its run-time reliability tests. In this case, when the system found that the Reskit performance DLL, Rkperf.dll, generated an exception, it disabled the DLL. The developer repaired Rkperf.dll, but it is still disabled in the system.

The following procedure uses Exctrlst to re-enable the performance counters for the Reskit service:

1. Start Exctrlst.

2. In the **Extensible Performance Counters** list, click the Reskit service.

3. Click to select the **Performance Counters Enabled** check box.

4. Click **Refresh**.

5. Use Typeperf to verify that the counter DLL loads and runs properly.

Typeperf attempts to load and run all installed performance counter DLLs. It lists the performance objects that successfully loaded in the order in which they load (alphabetical order by service name). In this case, the performance objects created by Rkperf.dll appear in Typeperf, indicating that the DLL repair was successful and the Reskit performance objects and counters are enabled and ready for use.

Note

- You cannot copy or export a list of disabled performance DLLs from Exctrlst. However, you can use Regfind, a command-line tool in the Windows 2000 Server Resource Kit, to list all instances of the **Disable Performance Counters** registry entry in the Services subkeys. When the value of this entry is 1, the performance DLL for the service, and the objects and counters it creates, are disabled.

  At the command prompt, type:

  ```
  regfind "Disable Performance Counters" -n -p HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services >
  Disabled.txt
  ```

  This command creates a file called Disabled.txt with the names of all services that have the registry entry. You can then

search the file for entries with a value of 1.

**Example 4: Find the performance objects for a service**
The following procedure comprises one method for determining which performance objects are associated with a service:

1. To list all performance counters that are enabled and operating properly on the system, use Typeperf, a tool included in Windows XP Professional and Windows Server 2003. For information about this tool, see "Typeperf" in Help for Windows Server 2003. At the command prompt, type:

   ```
   typeperf /q > all.txt
   ```

   This command saves the performance counter list in the All.txt file.

2. To disable the performance counters for the service, use Exctrlst.

3. Run Typeperf again. At the command prompt, type:

   ```
   typeperf /q > no_svc.txt
   ```

   This command saves the revised performance counter list in the No_Svc.txt file. This time, the disabled performance counters do not appear in the list.

4. Use a comparison tool, such as Ec, to compare the output from the Typeperf runs. For information about using Ec, see "Ec" in Help for Windows Server 2003. The counters that are missing from the No_Svc.txt file are those associated with the service.

5. Use Exctrlst to re-enable the counters for use.

**See Also**
**Concepts**
Exctrlst Overview
Exctrlst Remarks
Exctrlst UI
Alphabetical List of Tools
Pviewer Overview
Devcon Overview
Apmstat Overview

# Exctrlst UI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Exctrlst UI**

Exctrlst displays all registry-based performance counters installed on the system.

**Exctrlst dialog box**




**To run Exctrlst**

- At the command prompt or in the **Run** dialog box, type **Exctrlst**.

**To configure Exctrlst**

1. Select a computer.

   The **Machine Name** box displays the name of the computer being monitored. The default is the local computer.

   To monitor a remote computer, delete the name in the **Machine Name** box, type the name of the remote computer (preceded by backslashes), and click **Refresh**.

2. Select a sort order.

   Click a button in the **Sort Order** section to sort the items in the **Extensible Performance Counters** box.

   - **Library File**. Lists the performance DLLs installed on the system and the service that supports each DLL. The list is sorted in alphabetical order by service name.
   - **Service**. Lists the services that have installed performance DLLs and the name and location of the DLL for each service. The list is sorted in alphabetical order by the fully qualified path to the DLL.
   - **Counter ID**. Lists the services that have installed performance DLLs and the starting index for the performance

object installed by the service. The list is sorted in numerical order by index.

This starting index is the location of the performance object in the performance registry (HKEY_PERFORMANCE_DATA). Each object is assigned a range of indices within which it stores the counter names and explain text for the object.

**Note**

Performance objects that the system installs do not store their counter names in the performance registry. The index for these objects appears as 0x0 or N/A.

1. Select a service or DLL.

   The **Extensible Performance Counters** box displays a list of all registry-based performance DLLs on the system.

   Clicking an item in the **Extensible Performance Counters** displays information about the item in the fields below the box.

2. View, select, or clear the **Performance Counters Enabled** check box.

   The **Performance Counters Enabled** check box indicates whether the performance counters for the service are enabled or disabled.

   - To disable all performance counters for a service, clear the check box, and then click **Refresh**.
   - To enable all performance counters for a service, select the check box, and then click **Refresh**.

   The system automatically disables all performance counters for a service if the performance DLL fail the system's run-time reliability tests for performance DLLs. For more information, see Exctrlst Remarks in *Exctrlst Notes*.

**To interpret the Exctrlst display**

The following tables list and describe the fields in the Exctrlst dialog box.

- The system data in the upper section of the dialog box comes from the Perflib subkey in the registry.

| Item | Description |
|---|---|
| Last System Counter ID | Displays the highest index used to store a counter name, as recorded by the Performance Library (Perflib). Must match the value of **Last Counter Text ID.**<br><br>This field displays the value of the **Last Counter** registry entry (HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib\Last Counter). |
| Last Counter Text ID | Displays the highest index actually used to store a counter name.<br><br>This value is derived from the highest index used in the **Counter** registry entry (HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib*LanguageCode*\Counter). |
| Last System Help ID | Displays the highest index used to store explain text for a counter, as recorded by the Performance Library (Perflib). Must match the value of **Last Help Text ID.**<br><br>This field displays the value of the **Last Help** registry entry (HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib\Last Help). |
| Last Help Text ID | Displays the highest index actually used to store explain text.<br><br>This value is derived from the highest index used in the **Help** registry entry (HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib*LanguageCode*\Help). |

- The DLL-specific data in the lower section of the dialog box comes from the Performance subkeys in the registry. This data changes when you select an item from the **Extensible Performance Counters** dialog box.

| Item | Description |
|---|---|
| DLL Name | Displays the name and fully qualified path of the DLL that provides performance data for this service.<br><br>If the system cannot find the DLL file in this path, the registry records an error in the Application Log in Event Viewer. |
| Open Procedure | Displays the name of the DLL function that initializes the performance collection features of the service.<br><br>The open procedure must be exported by the DLL and it must return a successful status. Otherwise, the registry cannot collect performance data from this service. If this function is not found in the DLL or if it returns an error, the registry records an error in the Application Log in Event Viewer. |
| Collect Procedure | Displays the name of the DLL function that collects the performance data for the service.<br><br>The registry calls the collect function when a performance tool, such as System Monitor, requests performance data. If this function cannot be found in the DLL, the registry records an error in the Application Log in Event Viewer. |
| Close Procedure | Displays the name of the DLL function that terminates and completes the performance data collection features of the service.<br><br>The registry calls the close function when a performance tool, such as System Monitor, closes the registry interface. If this function cannot be found in the DLL, or it has a fault or causes an error, the registry records an error in the Application Log in Event Viewer. |
| Counter ID Range | Displays the range of index values that were assigned to the service to store counter names.<br><br>The values shown here must match the indices actually used, as shown in the value of the **Counter** registry entry (HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib*LanguageCode*\Counter).<br><br>"N/A" is displayed for system counters, which do not store counter names in the registry. If the value is **Not Found**, then the counters were not installed properly and will not appear correctly in tools such as System Monitor. |
| Help ID Range | Displays the range of index values that were assigned to the service to store explain text for its counters.<br><br>The values shown here must match the indices actually used, as shown in the value of the **Help** registry entry (HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib*LanguageCode*\Help).<br><br>"N/A" is displayed for system counters, which do not store explain text in the registry. If the value is **Not Found**, then the counters were not installed properly and will not appear correctly in tools such as System Monitor. |

**See Also**
**Concepts**
Exctrlst Overview
Exctrlst Remarks
Exctrlst Examples
Alphabetical List of Tools
Pviewer Overview
Devcon Overview
Apmstat Overview

# Gflags.exe

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Gflags.exe: Global Flags Editor (with Page Heap Verification)**
This tool allows you to enable and disable advanced internal system diagnostics and troubleshooting features on computers running Windows 2000, Windows XP Professional, and Windows Server 2003. Gflags is designed as a debugging tool for application developers. It is most often used to turn on indicators that other tools track, count, and log. You can run Gflags from the command line or use its graphical user interface dialog box.

Use Gflags to activate the following debugging features:

- Registry. Set systemwide debugging features for all Windows processes. These settings are stored in the registry. They take effect when you restart the system and remain effective until you change them.
- Kernel mode settings. Set debugging features for this session. These settings are effective immediately, but are lost when the system stops. The settings affect processes started after this command completes.
- Image file settings. Set debugging features for a particular program. These settings are stored in the registry. They take effect when you restart the program and remain effective until you change them.
- Debugger. Specify that a particular program always runs in a debugger. This setting is stored in the registry. It is effective immediately and remains effective until you change it. (This feature is available only from the Gflags dialog box.)
- Launch. Run a program with the specified debugging settings. The debugging settings are effective until the program stops. (This feature is available only from the Gflags dialog box.)
- Page heap verification. Enable, disable, and configure page heap verification for a program. When enabled, page heap monitors dynamic heap memory operations, including allocations and frees, and causes a debugger break when it detects a heap error.

### Caution

- Use Gflags only in consultation with Microsoft Product Support Services. It is an optional component provided solely for troubleshooting purposes. Incorrect use of this tool can degrade system performance or prevent the system from starting, requiring you to reinstall the operating system.

### Important

- Pool tagging is permanently enabled on all versions of Windows Server 2003. On these systems, the **Enable pool tagging** check box on the Gflags dialog box is dimmed and commands to enable or disable pool tagging fail.

**Corresponding UI**
This version of Gflags incorporates the functions of Pageheap.exe, a heap allocation monitoring tool formerly included in the Support Tools folder.

**System Requirements**
The following are the system requirements for Gflags:

- This version of Gflags runs only on Windows XP Professional and Windows Server 2003. A version of Gflags for Windows 2000 is included in the Support Tools directory of Windows 2000 Server.
- The user must be a member of the Administrators group on the computer.
- Gflags is a local computer tool. You cannot use it to set flags on a remote computer.

**File Required**

- Gflags.exe

**See Also**
**Concepts**
Gflags Remarks
Gflags Syntax
Gflags UI
Gflags Examples
Gflags Flag Reference
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Gflags Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Gflags Remarks**
**General Information**

- To display the Gflags dialog box, at the command line, type `Gflags` (with no parameters).
- To set flags in the registry or in kernel mode or to enable page heap verification, you must be a member of the Administrators group on the computer. However, users with at least Guest account access can launch a program from the Gflags dialog box.
- Gflags system-level registry settings appear in the registry immediately, but do not take effect until you restart the system.
- Gflags image file registry settings appear in the registry immediately, but do not take effect until you restart the process.
- The debugger and launch features in the Gflags dialog box are program specific. You can only set them on one image file at a time.
- For information about kernel special pool, see the Windows Driver Development Kit (DDK).
- For information about debuggers and to obtain debuggers and documentation, see Microsoft Debugging Tools.

**Flag Details**

- To clear all flags, set the flag to -FFFFFFFF. Setting the flag to 0 adds 0 to the current flag value.
- When you set the flags for an image file to FFFFFFFF (0xFFFFFFFF), Windows clears all flags and deletes the **GlobalFlag** entry in the image file subkey. The subkey is retained.

**Dialog Box and Command Line**
You can run Gflags by using its handy dialog box or from the command line. Most features are available in both forms, with the following exceptions.

**Dialog box only**

- Debugger. Specify that a particular program always runs in a debugger.
- Launch. Run a program with the specified debugging settings.
- Kernel Special Pool Tag. Detects kernel pool errors.

**Command line only**

- Set the size of the user mode stack trace database (**/tracedb**).
- Set page heap verification options.

**Gflags and Page Heap**
This version of Gflags includes the functionality of Pageheap.exe, a heap allocation monitoring tool formerly included in the Support Tools folder. The page heap feature reserves memory at the boundary of each allocation to trap access beyond the allocation that might otherwise go undetected.

The page heap options in Gflags let you select *standard heap verification*, which writes fill patterns at the end of each heap allocation and examines the patterns when the allocations are freed, or *full-page heap verification*, which places an inaccessible page at the end of each allocation so that the program stops immediately if it accesses memory beyond the allocation. Because full heap verification uses a full page of memory for each allocation, its widespread use can cause system memory shortages.

Use the following syntax variations to enable page heap verification:

- To enable standard page heap verification for all processes, use **gflags /r +hpa** or **gflags /k +hpa**.
- To enable standard page heap verification for one process, use **gflags /p /enable***ImageFile*.
- To enable full page heap verification for one process, **gflags /i***ImageFile***+hpa** or **gflags /p /enable***ImageFile***/full**.

All page heap settings, except **/k**, are stored in the registry and remain effective until you change them.

Use care in interpreting the **Enable page heap** check box for an image file in the Gflags dialog box. It indicates that page heap verification is enabled for an image file, but it does not indicate whether it is full or standard page heap. If the check results from selecting the check box, then full page heap verification is enabled for the image file. However, if the check results from use of the command-line interface, then the check can represent the enabling of either full or standard page heap verification for the image file.

To determine whether a full or standard page heap verification is enabled for a program, at the command line, type **gflags /p**. In the resulting display, **traces** indicates that standard page heap verification is enabled for the program and **full traces** indicates that full page heap verification is enabled for the program.

**Registry Information**
Debugging settings saved between sessions are stored in the registry. You can use the registry APIs, Regedit.exe, or Reg to query or change these values. For information about using Reg see "Reg" in Help for Windows Server 2003. The following table lists the types of settings and where they are stored in the registry.

| | |
|---|---|
| Systemwide debugging values. ("Registry") | HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager**GlobalFlag** |
| Program-specific debugging values. ("Image file") | HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options*ImageFile***GlobalFlag** |
| Page heap options for an image file | HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options*ImageFile***PageHeapFlags** |
| User mode stack trace database size (**tracedb**) | HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options*ImageFile***StackTraceDatabaseSizeInMb** |
| Create user mode stack trace database (ust, 0x1000) for an image file | Windows adds the image file name to the value of the **USTEnabled** registry entry (HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options**USTEnabled**). |
| Kernel Special Pool Tags | HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management**PoolTag** |
| Verify Start / Verify End | HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management**PoolTagOverruns**. The **Verify Start** option sets the value to 0. The **Verify End** option sets the value to 1. |

**See Also**
**Concepts**
Gflags.exe
Gflags Syntax
Gflags UI
Gflags Examples
Gflags Flag Reference
Alphabetical List of Tools

Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Gflags Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Gflags Examples**
**Example 1: Display flags**
The following command displays the current value of the system-level **GlobalFlag** registry entry (HKLM\System\CurrentControlSet\Control\Session Manager\GlobalFlag). It uses the **/r** parameter to specify the system-level registry entry.

```
gflags /r
```

In response, Gflags displays a single hexadecimal value representing the sum of all flags set and a list of the flags set.

```
Current Boot Registry Settings are: 40001400
    ptg - Enable pool tagging
    ust - Create user mode stack trace database
    bhd - Enable bad handles detection
```

In this example, the results show that there are three tags set, with a combined value of 0x40001400.

- Enable pool tagging (ptg) = 0x400
- Create user mode stack trace database (ust) = 0x1000
- Enable bad handles detection (bhd) = 0x40000000

Remember that the value in the registry might not be the current value. Changes to the registry are not effective until the process is restarted.

You can use similar commands to display the debugging values for the current session, known as *kernel mode*, and debugging values for a particular image file.

The following command displays the debugging values for the current session. It uses the **/k** parameter to indicate kernel mode.

```
gflags /k
```

The following command displays the debugging values set in the registry for Notepad.exe. This value is stored in a **GlobalFlag** registry entry located in the program-specific subkey, HKLM\System\CurrentControlSet\Control\Session Manager\Image File Execution Options\Notepad.exe. It uses the **/i** parameter to indicate image file mode and specifies the image file.

```
gflags /i notepad.exe
```

Remember that the value in the registry might not be the current value. Changes to the registry are not effective until the process is restarted.

**Example 2: Set a flag by using a flag abbreviation**
The following command sets the Show loader snaps flag in the **GlobalFlag** registry entry for the Notepad.exe image file. Show loader snaps takes snapshots of the load process, capturing in detail the loading and unloading of executable images and their supporting library modules.

The command uses the **/i** parameter to indicate image file mode and specifies the name of the image file, Notepad.exe. To identify the flag, the command uses **sls**, the abbrevation for Show loader snaps and it precedes the abbreviation with a plus sign (+) to indicate that the flag is set. Without the plus sign, the command has no effect.

```
gflags /i notepad.exe +sls
```

In response, Gflags displays the current flags for Notepad.exe. The display indicates that the command is successful. The Show loader snaps flag is enabled for all new instances of the Notepad.exe process.

```
Current Registry Settings for notepad.exe executable are: 00000002
    sls - Show Loader Snaps
```

**Example 3: Set a flag by using its hexadecimal value**
The following command sets the Enable page heap flag in the systemwide **GlobalFlag** registry entry. Enable page heap adds a guard page and other tracking features to each heap allocation.

The command uses the **/r** parameter to indicate system registry mode. To identify the flag, the command uses **2000000**, which represents 0x2000000, the hexadecimal value for Enable page heap.

Although the command sets a flag, it omits the plus sign. When using hexadecimal values, the sign is optional and add (+) is the default.

```
gflags /r 2000000
```

In response, Gflags displays the current value of **GlobalFlag**. The display indicates that the command is successful. The Enable page heap flag is set when you restart the system.

```
Current Boot Registry Settings are: 02000000
    hpa - Enable page heap
```

**Example 4: Set multiple flags**
The following command sets the following three debugging flags for the current session:

- Enable heap free checking (hfc, 0x20)
- Enable heap parameter checking (hpc, 0x40)
- Enable heap validation on call (hvc, 0x80)

This command uses the **/k** parameter to specify kernel mode (session only). It sets the value for kernel mode to **E0** (0xE0), the sum of the hexadecimal values of the flags.

```
gflags /k e0
```

In response, Gflags displays the current value of **GlobalFlag**. The display indicates that the command is successful and that the correct three flags are, indeed, set.

```
Current Running Kernel Settings are: 000000e0
    hfc - Enable heap free checking
    hpc - Enable heap parameter checking
    hvc - Enable heap validation on call
```

Please note that each of the following commands has the same effect as the command used in this example and can be used interchangeably:

```
gflags /k +20 +40 +80
```

```
gflags /k +E0
```

```
gflags /k +hfc +hpc +hvc
```

Flags set in kernel mode are effective immediately and remain effective until the system stops.

**Example 5: Clear a flag**
The following command clears the Enable page heap flag from the systemwide **GlobalFlag** registry entry. The command uses the **/r** parameter to indicate the systemwide registry and **hpa**, the abbreviation for the Enable page heap flag. The minus sign (-) specifies that the flag is to be cleared.

```
gflags /r -hpa
```

In response, Gflags displays the current value of **GlobalFlag**. The display indicates that the command is successful and that there are no longer any flags set in the registry value.

```
Current Boot Registry Settings are: 00000000
```

Please note that the following command, which uses the hexadecimal value of the Enable page heap flag, has the same effect as the command used in this example and can be used interchangeably:

```
gflags /r -02000000
```

**Example 6: Clear all flags**
There are two different ways to clear all flags set for a particular debugging mode:

- Clear the current flags.
- Clear all flags.

### Note

- The methods demonstrated by this example clear flags only. They do not reset the maximum stack trace size or kernel special pool tag to the default values.

**Clear the current flags**
The following command uses the first method for clearing flags. It clears the current value of **GlobalFlag** by clearing the flags set. In this example, the current value is 0xE0.

```
gflags /r -E0
```

In response, Gflags displays the current value of **GlobalFlag**. The display indicates that the command is successful and that there are no longer any flags set in the registry value.

```
Current Boot Registry Settings are: 00000000
```

Please note that the following commands have the same effect as the command used in this example and can be used interchangeably:

```
gflags /r -20 -40 -80
```

```
gflags /r -hfc -hpc -hvc
```

**Clear all flags**
The following command uses the second method for clearing flags. It clears all possible values set for **GlobalFlag**.

```
gflags /r -ffffffff
```

In response, Gflags displays the current value of **GlobalFlag**. The display indicates that the command is successful and that there are no longer any flags set in the registry value.

```
Current Boot Registry Settings are: 00000000
```

### Note

- Type this command into Notepad, then save the file as ClearFlag.bat. Thereafter, to clear all flags, just type `ClearFlag`.

Finally, the following example demonstrates that the intuitive method of clearing all flags does not work.

The following command should set the value of the systemwide **GlobalFlag** registry entry to 0 intuitively but it does not. In this example, the current value of **GlobalFlag** is 0xE0.

```
gflags /r 0
```

In response, Gflags displays the value of **GlobalFlag** after the command completes:

```
Current Boot Registry Settings are: 000000e0
    hfc - Enable heap free checking
    hpc - Enable heap parameter checking
    hvc - Enable heap validation on call
```

The command has no effect because it adds the value 0 to the current value of **GlobalFlag**.

**Example 7: Clear all flags for an image file**
The following command clears all flags and image debugger options for an image file. The command adds high-values (0xFFFFFFFF) to the current flag value. Gflags responds by deleting the **GlobalFlag** entry for the image file, thereby deleting all of the values it stores.

This method does not clear flags set in the systemwide registry entry or in kernel mode.

```
gflags /i notepad.exe ffffffff
```

In response, Gflags displays a message indicating that there are no flags set for the image file:

```
No Registry Settings for notepad.exe executable
```

To determine which image files have flags set, use the **query** operation of Reg.exe, a tool included in %windir%\system32, to search the registry subkeys in the Image file Execution Options subkey.

The following command searches for **GlobalFlag** in the specified path. The **/v** parameter directs Reg query to search for the value **GlobalFlag**. The /s parameter specifies to search all subkeys and values recursively.

```
reg query "HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Image File Execution Options" /v
GlobalFlag /s
```

In response, Reg query displays all image file subkeys with **GlobalFlag** registry entries. It also displays the value of the entry. Because the **gflags** command deleted the entry for Notepad.exe, it does not appear in the Reg query list.

```
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Image File Execu
tion Options
    f32main.exe
        GlobalFlag = 0x000000c0
    dgsetup.exe
        GlobalFlag = 0x00000080
```

### Note

- Type the **reg query** command into Notepad, then save the file as Images.bat. Thereafter, to find image files for which flags have been set, just type `Images`.

**Example 8: Enlarge the user mode stack trace database**
The following command increases the maximum size of the user mode stack trace database for the MyApp.exe process from 8 MB to 24 MB.

The command uses the **/i** parameter to specify the image file. It uses the **/tracedb** parameter to set the maximum stack trace database size and the value **24** to indicate the desired MB size in decimal units. (Hexadecimal units are not valid.)

```
            gflags /i MyApp.exe /tracedb 24
```

As the following error message indicates, this command fails because the "Create user mode stack trace database" (+ust) flag is not set for the process. You cannot set the size of a trace database until you create one.

```
Failed to set the trace database size for `MyApp.exe'
The user mode stack trace database flag (+ust) must be enabled before setting /tracedb.
```

The following commands fix the error. The first command creates a trace database for MyApp.exe and the second command sets the maximum size of the trace database to 24 MB. These commands cannot be combined into a single command. The listing below shows the commands and the success message from Gflags.

```
            gflags /i MyApp.exe +ust

Current Registry Settings for MyApp.exe executable are: 00001000
    ust - Create user mode stack trace database

gflags /i MyApp.exe /tracedb 24
```

```
Trace database size for `MyApp.exe' set to 24 Mb.
```

Gflags can change the size of the user mode stack trace database, but it does not display it. To display the trace database size, use registry APIs, Regedit.exe, or Reg, to check the value of the **StackTraceDatabaseSizeInMB** registry entry (HKLM\Software\Microsoft\Windows NT\CurrentVersion\Image File Execution Options*ImageFileName***StackTraceDatabaseSizeInMB**).

The following command uses Reg.exe to query the value of **StackTraceDatabaseSizeInMB**:

```
reg query "HKLM\Software\Microsoft\Windows NT\CurrentVersion\Image File Execution Options\MyApp.exe" /v
StackTraceDatabaseSizeInMB
```

In response, Reg.exe returns the value of **StackTraceDatabaseSizeInMB**, which confirms that the new value, 24 (0x18), was set. This value becomes effective when you restart the MyApp.exe.

```
! REG.EXE VERSION 3.0

HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Image File Execu
tion Options\MyApp.exe
    StackTraceDatabaseSizeInMB  REG_DWORD       0x18
```

### Notes

- Type the **reg query** command into Notepad, then save the file as TraceDb.bat. Thereafter, to display the value of **StackTraceDatabaseSizeInMB**, just type `TraceDb`.

- To see the registry locations of values set by Gflags, see Gflags Remarks.

#### Example 9: Detect a pool memory leak

The following example uses Gflags to set the **Enable pool tagging** flag in the registry entry for the system. Then, it uses Poolmon Overview, a Windows Support Tool, to display the size of the memory pools. Poolmon uses pool tags to monitor the bytes in the paged and nonpaged memory pools. By running Poolmon periodically, you can identify pools that expand continuously over time. This pattern often indicates a memory leak.

### Notes

- Pool tagging is permanently enabled on all versions of Windows Server 2003. On these systems, the **Enable pool tagging** check box on the Gflags dialog box is dimmed and commands to enable or disable pool tagging fail.

- If pool tagging is not enabled, Poolmon fails and diplays the following message: "Query pooltags Failed c0000002."

1. Set the Enable pool tagging flag in the systemwide registry entry.

   Use one of the following methods. These methods have the same effect and can be used interchangeably.

   - In the Gflags dialog box, click the **Registry** button, click the **Enable pool tagging** check box, and then click **Apply**.

   - At the command line, use the **/r** parameter to specify the systemwide registry entry, and the **ptg** abbreviation to specify the flag.

     ```
     gflags /r +ptg
     ```

   - At the command line, use the **/r** parameter to specify the systemwide registry entry, and the **0x400** hexadecimal value to specify the flag. The plus sign (+) is optional in this format.

     ```
     gflags /r +400
     ```

2. Restart the computer to make the registry change effective.

3. Run Poolmon periodically by using the following command. In this command, the **/b** parameter sorts the pools in descending size order.

   ```
   poolmon /b
   ```

   In response, Poolmon displays a table of memory pool size, including the number of allocs and frees, and the amount of memory remaining in the pool (in the Bytes column).

   ```
   Memory: 16224K Avail: 4564K PageFlts: 31 InRam Krnl: 684K P: 680K
    Commit: 24140K Limit: 24952K Peak: 24932K  Pool N: 744K P: 2180K

    Tag  Type    Allocs          Frees         Diff   Bytes      Per Alloc
   ---------------------------------------------------------------------

    CM   Paged    1283 (   0)    1002 (   0)    281 1377312 (     0) 4901
   Strg  Paged   10385 (  10)    6658 (   4)   3727  317952 (   512)   85
    Fat  Paged    6662 (   8)    4971 (   6)   1691  174560 (   128)  103
   MmSt  Paged     614 (   0)     441 (   0)    173   83456 (     0)  482
   ```

   If the value of Bytes in a pool expands continuously for no obvious reason, there might be a memory leak in that pool.

4. Turn off Enable pool tagging.

   The following command line uses the flag abbreviation method, but you can identify the flag by its hexadecimal value or use the **Global Flags** dialog box:

   ```
   gflags /r -ptg
   ```

5. Restart the system to make the registry change effective.

### Note

- Use the append symbol (>>) to redirect the Poolmon output to a log file. Later, you can examine the log file for pool size trends. For example:

  ```
  poolmon.exe /b >> poolmon.log
  ```

#### Example 10: Get a complete dump of active heaps in a process

This example uses Gflags and Umdh.exe (User Mode Dump Heap), a tool available from Microsoft Product Support Services.

To detect a leak in heap memory in Notepad.exe:

1. Set the Create user mode stack trace database flag for the Notepad.exe image file.

   The following command uses Gflags to set the Create user mode stack trace database flag. It uses the **/i** parameter to identify the image file and the **ust** abbreviation for the flag.

   ```
   gflags /i Notepad.exe +ust
   ```

   As a result of this command, a user mode stack trace is created for all new instances of the Notepad process.

2. Set the symbol file path.

   The following command creates an environment variable that stores the path to the directory of symbol files:

   ```
   set _NT_SYMBOL_PATH=C:\Windows\symbols
   ```

3. Start Notepad.

4. Find the process identifier (PID) of the Notepad process.

   You can find the PID of any running process from Task Manager or Tasklist, a tool included in Windows XP Professional and Windows Server 2003 operating systems. In this example, the Notepad PID is 1228.

5. Run Umdh.exe

   The following command runs Umdh.exe. It uses the **-p:** parameter to specify the PID which, in this example, is 1228. It uses the **/f:** parameter to specify the name and location of the output file for the heap dump, Notepad.dmp.

   ```
   umdh -p:1228 -f:notepad.dmp
   ```

   In response, Umdh writes a complete dump of all active heaps to the Notepad.dmp file.

**Example 11: Enable page heap verification**
The following command enables standard page heap verification for the MyApp.exe program.

```
gflags /p /enable myapp.exe
```

The following commands enable full page heap verification for the MyApp.exe program. Although these commands create different settings in the registry, they are all functionally equivalent to selecting the **Enable page heap** check box for the MyApp.exe image file in the Gflags dialog box. These methods can be used interchangeably.

```
gflags /p /enable myapp.exe /fullgflags /i myapp.exe +hpagflags /i myapp.exe +02000000
```

The following commands disable full or standard page heap verification for the Myapp.exe program, regardless of the command or dialog box method used to enable page heap verification.

```
gflags /p /disable myapp.exegflags /i myapp.exe -hpagflags /i myapp.exe -02000000
```

### Note

- When using the **/debug** or **/kdebug** parameters, use the **/p /disable** parameters to turn off the page heap verification (not the **/i -hpa** parameters). Otherwise, entries read by the debugger are not deleted from the registry.

**Example 12: Use page heap verification to find a bug**
This series of commands demonstrates how to use the page heap verification features of Gflags and the Windows NT Symbolic Debugger (NTSD) to detect an error in heap memory use. In this example, the programmer suspects that a fictitious application, Pheap-buggy.exe, has a heap error, and proceeds through a series of tests to identify the error.

**Step 1: Enable standard page heap verification**
The following command enables standard page heap verification for Pheap-buggy.exe:

```
gflags /p /enable pheap-buggy.exe
```

**Step 2: Verify that page heap is enabled**
The following command lists the programs on the system for which page heap verification is enabled:

```
gflags /p
```

In response, Gflags displays the following list of programs. In this display, **traces** indicates standard page heap verification and **full traces** indicates full page heap verification. In this case, Pheap-buggy.exe is listed with **traces**, indicating that standard page heap verification is enabled, as intended.

```
pheap-buggy.exe: page heap enabled with flags (traces )
```

**Step 3: Run the debugger**
The following command runs the CorruptAfterEnd function of Pheap-buggy.exe in NTSD with the **-g** (ignore initial breakpoint) and **-x** (set second-chance break on access violation exceptions) parameters set:

```
ntsd -g -x pheap-buggy /CorruptAfterEnd
```

When the application fails, NTSD generates the following display, which indicates that it detected an error in Pheap-buggy.exe:

```
============================================================
VERIFIER STOP 00000008: pid 0xAA0: corrupted suffix pattern

        00C81000 : Heap handle
        00D81EB0 : Heap block
        00000100 : Block size
        00000000 :
============================================================

Break instruction exception - code 80000003 (first chance)
eax=00000000 ebx=00d81eb0 ecx=77f7e257 edx=0006fa18 esi=00000008 edi=00c81000
eip=77f7e098 esp=0006fc48 ebp=0006fc5c iopl=0         nv up ei pl zr na po nc
cs=001b  ss=0023  ds=0023  es=0023  fs=0038  gs=0000             efl=00000246
ntdll!DbgBreakPoint:
77f7e098 cc              int     3
```

The header information includes the address of the heap with the corrupted block (00C81000 : Heap handle), the address of the corrupted block (00D81EB0 : Heap block) and the size of the allocation (00000100 : Block size).

The "corrupted suffix pattern" message indicates that the application violated the data integrity pattern that Gflags inserted after the end of the Pheap-buggy.exe heap allocation.

**Step 4: Display the call stack**
The next step is to use the addresses detected to locate the function that caused the error. The next two commands turn on line number dumping in the debugger and display the call stack with line numbers.

```
C:\>.linesLine number information will be loaded

C:\>kb
```

```
ChildEBP RetAddr  Args to Child
WARNING: Stack unwind information not available. Following frames may be wrong.
0006fc5c 77f9e6dd 00000008 77f9e3e8 00c81000 ntdll!DbgBreakPoint
0006fcd8 77f9f3c8 00c81000 00000004 00d81eb0 ntdll!RtlpNtEnumerateSubKey+0x2879
0006fcfc 77f9f5bb 00c81000 01001002 00000010 ntdll!RtlpNtEnumerateSubKey+0x3564
0006fd4c 77fa261e 00c80000 01001002 00d81eb0 ntdll!RtlpNtEnumerateSubKey+0x3757
0006fdc0 77fc0dc2 00c80000 01001002 00d81eb0 ntdll!RtlpNtEnumerateSubKey+0x67ba
0006fe78 77fbd87b 00c80000 01001002 00d81eb0 ntdll!RtlSizeHeap+0x16a8
0006ff24 010013a4 00c80000 01001002 00d81eb0 ntdll!RtlFreeHeap+0x69
0006ff3c 01001450 00000000 00000001 0006ffc0 pheap-buggy!TestCorruptAfterEnd+0x2b
[d:\nttest\base\testsrc\kernel\rtl\pageheap\pheap-buggy.cxx @ 185]
0006ff4c 0100157f 00000002 00c65a68 00c631d8 pheap-buggy!main+0xa9
[d:\nttest\base\testsrc\kernel\rtl\pageheap\pheap-buggy.cxx @ 69]
0006ffc0 77de43fe 00000000 00000001 7ffdf000 pheap-buggy!mainCRTStartup+0xe3 [crtexe.c
@ 349]
0006fff0 00000000 0100149c 00000000 78746341 kernel32!DosPathToSessionPathA+0x204
```

The debugger displays the call stack for Pheap-buggy.exe with line numbers. The call stack display shows that the error occurred when the TestCorruptAfterEnd function in Pheap-buggy.exe tried to free an allocation at 0x00c80000 by calling HeapFree, a redirect to RtlFreeHeap.

The most likely cause of this error is that the program wrote past the end of the buffer that it allocated in this function.

**Step 5: Enable full page heap verification**
Unlike standard mode, full page heap verification can catch the misuse of this heap buffer as soon as it occurs. The following command enables full page heap verification for Pheap-buggy.exe:

```
gflags /p /enable pheap-buggy.exe /full
```

**Step 6: Verify that full page heap is enabled**
The following command lists the programs on the system for which page heap verification is enabled:

```
gflags /p
```

In response, Gflags displays the following list of programs. In this display, **traces** indicates standard page heap verification and **full traces** indicates full page heap verification. In this case, Pheap-buggy.exe is listed with **full traces**, indicating that full page heap verification is enabled, as intended.

```
pheap-buggy.exe: page heap enabled with flags (full traces )
```

**Step 7: Run the debugger again**

The following command runs the CorruptAfterEnd function of Pheap-buggy.exe in the NTSD debugger with the **-g** (ignore initial breakpoint) and **-x** (set second-chance break on access violation exceptions) parameters set:

```
ntsd -g -x pheap-buggy /CorruptAfterEnd
```

When the application fails, NTSD generates the following display, which indicates that it detected an error in Pheap-buggy.exe:

```
Page heap: process 0x5BC created heap @ 00880000 (00980000, flags 0x3)
ModLoad: 77db0000 77e8c000   kernel32.dll
ModLoad: 78000000 78046000   MSVCRT.dll
Page heap: process 0x5BC created heap @ 00B60000 (00C60000, flags 0x3)
Page heap: process 0x5BC created heap @ 00C80000 (00D80000, flags 0x3)
Access violation - code c0000005 (first chance)
Access violation - code c0000005 (!!! second chance !!!)
eax=00c86f00 ebx=00000000 ecx=77fbd80f edx=00c85000 esi=00c80000 edi=00c16fd0
eip=01001398 esp=0006ff2c ebp=0006ff4c iopl=0         nv up ei pl nz na po nc
cs=001b  ss=0023  ds=0023  es=0023  fs=0038  gs=0000             efl=00000206
pheap-buggy!TestCorruptAfterEnd+1f:
01001398 889801010000     mov     [eax+0x101],bl          ds:0023:00c87001=??
```

With full page heap verification enabled, the debugger breaks at an access violation. To find the precise location of the access violation, turn on line number dumping and display the call stack trace.

The numbered call stack trace appears as follows:

```
ChildEBP RetAddr  Args to Child
0006ff3c 01001450 00000000 00000001 0006ffc0 pheap-buggy!TestCorruptAfterEnd+0x1f
[d:\nttest\base\testsrc\kernel\rtl\pageheap\pheap-buggy.cxx @ 184]
0006ff4c 0100157f 00000002 00c16fd0 00b70eb0 pheap-buggy!main+0xa9
[d:\nttest\base\testsrc\kernel\rtl\pageheap\pheap-buggy.cxx @ 69]
0006ffc0 77de43fe 00000000 00000001 7ffdf000 pheap-buggy!mainCRTStartup+0xe3 [crtexe.c
@ 349]
WARNING: Stack unwind information not available. Following frames may be wrong.
0006fff0 00000000 0100149c 00000000 78746341 kernel32!DosPathToSessionPathA+0x204
```

The stack trace shows that the problem occurs in line 184 of Pheap-buggy.exe ([d:\nttest\base\testsrc\kernel\rtl\pageheap\pheap-buggy.cxx @ 184]). Because full page heap verification is enabled, the call stack ends in the program code, not in a system DLL. As a result, the violation was caught where it happened, instead of when the heap block was freed.

**Step 8: Locate the error in the code**

A quick inspection reveals the source of the problem: The program tries to write to the 101st byte of a 100-byte buffer, a common off-by-one error.

```
*((PCHAR)Block + 0x100 + 1) = 0;
```

**See Also**

**Concepts**

Gflags.exe
Gflags Remarks
Gflags Syntax
Gflags UI
Gflags Flag Reference
Alphabetical List of Tools
Windiff Overview

Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Gflags Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Gflags Syntax**

- **gflags**[**/r** [*Flag* [*Flag*...]] | **/k** [*Flag* [*Flag*...]] | **/i** *ImageFile* [*Flag* [*Flag*...]] | **/i** *ImageFile* **/tracedb** *SizeInMB* | **/?**]
- **gflags /p** [**/disable** *ImageFile*] [**/?**]
- **gflags /p /enable** *ImageFile* [**/full** [**/backwards**] | [**/random** *Probability*] | [**/size** *SizeStartSizeEnd*] | [**/address** *AddressStart AddressEnd*] | [**/dlls** *DLL* [*DLL*...]] [**/debug** | **/kdebug**] [**/unaligned**] [**/decommit**] [**/notraces**] [**/fault** *Rate* [*TimeOut*]] [**/leaks**] [**/protect**] [**/no_sync**] [**/no_lock_checks**]

**Parameters**

- Typing **gflags** without parameters opens the Gflags dialog box.
- Typing **gflags /p** without additional parameters displays a list of programs that have page heap verification enabled.
- The **/full**, **/random**, **/size**, **/address**, and **/dlls** parameters for the page heap **/enable** operation determine which allocations are subject to page heap verification and the verification method used. You can use only one of these parameters in each command. The default is standard page heap verification of all allocations of the process. The remaining parameters set options for page heap verification.

***Flag***

A three-letter abbreviation (*FlagAbbr*) or hexadecimal value (*FlagHex*) that represents a debugging feature. The abbreviations and hexadecimal values are listed in the Global Flags Table.

Use one of the following flag formats:

| Format | Description |
|---|---|
| {**+** \| **-**}*FlagAbbr* | Sets (**+**) or clears (**-**) the flag represented by the flag abbreviation. Without a plus (**+**) or minus (**-**) symbol, the command has no effect. |
| [**+** \| **-**]*FlagHex* | Adds (**+**) or subtracts (**-**) the hexadecimal value of a flag. A flag is set when its value is included in the sum. Add (**+**) is the default. Enter a hexadecimal value (without **0x**) that represents a single flag or enter the sum of hexadecimal values for multiple flags. |

***ImageFile***

The name of an executable file (.exe or .dll), including the file name extension. For example, Notepad.exe or Mydll.dll.

**Note**

- The ImageFile name is case insensitive.

**/r[ *Flag*[ *Flag*...]]**

Registry. View and change systemwide debugging flags stored in the registry. These settings take effect when you restart the system and remain effective until you change them.

Without flags, **gflags /r** displays the systemwide flags set in the registry.

**/k[ *Flag*[ *Flag*...]**

Kernel mode settings. View and change systemwide debugging flags for this session. These settings are effective immediately, but are lost when the system stops. The settings affect processes started after this command completes.

Without flags, **gflags /k** displays systemwide flags set for the current session.

**/i *ImageFile*[ *Flag*[ *Flag*...]]**

Image file settings. View and change debugging flags for a particular process. These settings are stored in the registry. They

are effective for all new instances of this process and they remain effective until you change them.

Without flags, **gflags /i***ImageFile* displays the flags set for the specified process.

**/tracedb *SizeInMb***

Sets the maximum size of the user mode stack trace database for the process. To use this parameter, the Gflags Flag Reference (ust) flag must be set for the process.

*SizeInMB* is a whole number representing the number of MB in decimal units. The default value is the minimum size, 8 MB; there is no maximum size. To revert to the default size, set *SizeInMB* to 0 or 8.

**/p**

Sets page heap verification options for a process. Without additional parameters, **gflags /p** displays a list of programs for which page heap verification is enabled.

Page heap verification monitors dynamic heap memory operations, including allocations and frees, and causes a debugger break when it detects a heap error.

**Note**

- The page heap features in Gflags only monitor heap memory allocations that use the standard Windows heap manager functions (HeapAlloc, GlobalAlloc, LocalAlloc, malloc, operator new, operator new[], or their corresponding deallocation functions) or those that use custom operations that call the standard heap manager functions.

**/disable *ImageFile***

Turns off page heap verification (standard or full) for the process.

This command is equivalent to turning off the Gflags Flag Reference flag for a process (**gflags /i***ImageFile*-**hpa**). You can use the commands interchangeably.

**/enable *ImageFile***

Turns on standard page heap verification for the process. Standard page heap verification examines fill patterns when a heap block is freed.

This parameter sets the Gflags Flag Reference flag for the process in the registry. However, the **/enable** parameter sets standard heap verification for the process, unlike the **/i** parameter, which sets full heap verification for a process.

**/full**

Turns on full page heap verification for the process. Full page heap verification places an inaccessible page at the end of each allocation. The default is standard page heap verification, which places fill patterns at the end of an allocation and examines the patterns when a heap block is freed.

Using this parameter is equivalent to turning on the Gflags Flag Reference flag for a process (**gflags /i***ImageFile*+**hpa**). You can use the commands interchangeably.

**Caution**

- Because full page heap verification uses at least one page of committed memory for each allocation, system resources can be quickly exhausted. For memory intensive processes, use the **/size** or **/dlls** parameter.

**/backwards**

Places the inaccessible page at the beginning of an allocation, rather than the end. As a result, the debugger traps overruns at the beginning of the buffer, instead of those at the end of the buffer. Valid only with the **/full** parameter.

**/random *Probability***

Chooses full or standard page heap verification for each allocation based on the specified probability. *Probability* is a decimal integer from 0 through 100 representing the probability of full page heap verification. A probability of 100 is the same as using the **/full** parameter. A probability of 0 is the same as using standard page heap.

**/size *SizeStart SizeEnd***

Enables full page heap verification for allocations within the specified size range and enables standard page heap verification for all other allocations by the program. *SizeStart* and *SizeEnd* are decimal integers. The default is standard page heap verification for all allocations.

**/address *AddressStart AddressEnd***

Enables full page heap verification for memory allocated while a return address in the specified address range is on the runtime call stack. It enables standard page heap verification for all other allocations by the process. To use this feature, specify the address of a function that calls the function with the suspect allocation. The address of the calling function will be on the call stack when the suspect allocation occurs.

*AddressStart* and *AddressEnd* specify the address range searched in allocation stack traces. The addresses are specified in hexadecimal "C" format, such as, 0xAABBCCDD.

The **/address** parameter is valid only on *x*86 architectures. It is not reliable and is rarely used.

**/dlls *DLL*[ ,*DLL*...]**

Enables full page heap verification for allocations requested by the specified DLLs and standard page heap verification for all other allocations by the process.

*DLL* is the name of a binary file, including its file name extension. The specified file must be a function library that the process loads during execution.

**/debug**

Automatically launches the process specified by the **/enable** parameter under the NTSD debugger with page heap enabled and with the following options: **ntsd -g -G -x**.

- **-g** ignores initial breakpoint in debugger
- **-G** ignores final breakpoint at process termination
- **-x** sets second-chance break on AV exceptions

This option is used for programs that are difficult to start from a command prompt and those that are started by other processes.

**Note**

- When using the **/debug** or **/kdebug** parameters, use the **/p /disable** parameters to turn off the page heap functions (not the **/i -hpa** parameters). Otherwise, entries read by the debugger are not deleted from the registry.

**/kdebug**

Automatically launches the process specified by the **/enable** parameter under the NTSD debugger with page heap enabled and the following options: **ntsd -g -G -d -x**. It also redirects the debugger output to the kernel debugger if one is attached.

- **-g** ignores initial breakpoint in debugger
- **-G** ignores final breakpoint at process termination
- **-d** sends all debugger output to kernel debugger via DbgPrint
- **-x** sets second-chance break on AV exceptions

This parameter is not commonly used.

**Note**

- When using the **/debug** or **/kdebug** parameters, use the **/p /disable** parameters to turn off the page heap functions (not the **/i -hpa** parameters). Otherwise, entries read by the debugger are not deleted from the registry.

**/unaligned**

Aligns the end of each allocation at an end-of-page boundary, even if doing so means that the starting address is not aligned on an 8-byte block. By default, the heap manager guarantees that the starting address of an allocation is aligned on an 8-byte block, but cannot always end at a page boundary.

This option is used to to detect off-by-one-byte errors. When this parameter is used with the **/full** parameter, the inaccessible page begins at the last byte of the allocation and an immediate fault occurs when a process tries to read or write

even one byte beyond the allocation.

Note

- Some programs, including Microsoft Internet Explorer, assume 8-byte alignment and do not operate correctly with the **/unaligned** parameter.

**/decommit**

This option is no longer valid. It is accepted, but ignored. In previous versions of page heap, the **/decommit** parameter placed a zone of reserved virtual space after an allocation instead of an inaccessible page. In this version, the **/decommit** optimization is always implemented in full page heap verification.

**/notraces**

Specifies that runtime stack traces are not saved.

This option improves performance slightly, but it makes debugging much more difficult. This parameter is valid, but its use is not recommended.

**/fault *Rate*[ *TimeOut*]**

Forces the program's memory allocations to fail at the specified rate and after the specified timeout.

This parameter inserts heap allocation errors into the image file being tested (a practice known as "fault injection") so that some memory allocations fail, as might occur when the program runs in low memory conditions. This test helps to detect errors in handling allocation failure, such as failing to release resources.

*Rate* is a decimal integer from 1 through 10000 representing the probability that an allocation will fail. *TimeOut* is the number of seconds of delay between the start of the program and the start of the fault injection routines. The default is 5 (five seconds).

**/leaks**

Checks for heap leaks when a process ends.

**Caution**

- The **/leaks** parameter disables full page heap. When **/leaks** is used, the **/full** parameter and parameters that modify the **/full** parameter, such as **/backwards**, are ignored, and Gflags performs standard page heap verification with a leak check.

**/protect**

Protects heap internal structures. This test is used to detect random heap corruptions. It can make execution significantly slower.

**/no_sync**

Checks for unsynchronized access. This parameter causes a break if it detects that a heap created with the HEAP_NO_SERIALIZE flag is accessed by different threads.

Do not use this flag to debug a program that includes a customized heap manager. Functions that synchronize heap access cause the page heap verifier to report synchronization faults that do not exist.

**/no_lock_checks**

Disables the critical section verifier.

**/?**

Displays help for Gflags.

**/p /?**

Displays help for page heap verification options in Gflags.

**Global Flags Table**

The following table lists the flags that Gflags changes, the hexadecimal value and abbreviation for each flag, and the destination (R for registry, K for kernel, I for image file) in which the flag is valid.

For a detailed description of each flag, see the Gflags Flag Reference.

**Caution**

- The symbolic name of each flag is provided for reference only. Because symbolic names change, they are not a reliable identifier of a global flag.

**Important**

- Pool tagging is permanently enabled in all versions of Windows Server 2003. On these systems, the **Enable pool tagging** check box on the Gflags dialog box is dimmed and commands to enable or disable pool tagging fail.

| Description | Symbolic name | Hexadecimal Value | Abbreviation | Destination |
|---|---|---|---|---|
| Gflags Flag Reference | FLG_DISABLE_DBGPRINT | 0x08000000 | ddp | R,K |
| Gflags Flag Reference | FLG_KERNEL_STACK_TRACE_DB | 0x2000 | kst | R |
| Gflags Flag Reference | FLG_USER_STACK_TRACE_DB | 0x1000 | ust | R,K,I |
| Gflags Flag Reference | FLG_DEBUG_INITIAL_COMMAND | 0x4 | dic | R,K |
| Gflags Flag Reference | FLG_DEBUG_INITIAL_COMMAND_EX | 0x04000000 | dwl | R, K |
| Gflags Flag Reference | FLG_HEAP_DISABLE_COALESCING | 0x00200000 | dhc | R,K,I |
| Gflags Flag Reference | FLG_DISABLE_PAGE_KERNEL_STACKS | 0x80000 | dps | R,K |
| Gflags Flag Reference | FLG_DISABLE_PROTDLLS | 0x80000000 | dpd | R,K,I |
| Gflags Flag Reference | FLG_DISABLE_STACK_EXTENSION | 0x10000 | dse | I |
| Gflags Flag Reference | FLG_CRITSEC_EVENT_CREATION | 0x10000000 | cse | R,K,I |
| Gflags Flag Reference | FLG_APPLICATION_VERIFIER | 0x100 | vrf | R,K,I |
| Gflags Flag Reference | FLG_ENABLE_HANDLE_EXCEPTIONS | 0x40000000 | bhd | R,K |
| Gflags Flag Reference | FLG_ENABLE_CLOSE_EXCEPTIONS | 0x00400000 | ece | R,K |
| Gflags Flag Reference | FLG_ENABLE_CSRDEBUG | 0x20000 | d32 | R,K |
| Gflags Flag Reference | FLG_ENABLE_EXCEPTION_LOGGING | 0x00800000 | eel | R,K |
| Gflags Flag Reference | FLG_HEAP_ENABLE_FREE_CHECK | 0x20 | hfc | R,K,I |
| Gflags Flag Reference | FLG_HEAP_VALIDATE_PARAMETERS | 0x40 | hpc | R,K,I |
| Gflags Flag Reference | FLG_HEAP_ENABLE_TAGGING | 0x800 | htg | R,K,I |
| Gflags Flag Reference | FLG_HEAP_ENABLE_TAG_BY_DLL | 0x8000 | htd | R,K,I |
| Gflags Flag Reference | FLG_HEAP_ENABLE_TAIL_CHECK | 0x10 | htc | R,K,I |
| Gflags Flag Reference | FLG_HEAP_VALIDATE_ALL | 0x80 | hvc | R,K,I |
| Gflags Flag Reference | FLG_ENABLE_KDEBUG_SYMBOL_LOAD | 0x40000 | ksl | R,K |
| Gflags Flag Reference | FLG_ENABLE_HANDLE_TYPE_TAGGING | 0x01000000 | eot | R,K |
| Gflags Flag Reference | FLG_HEAP_PAGE_ALLOCS | 0x02000000 | hpa | R,K,I |
| Gflags Flag Reference (Windows 2000 and Windows XP only) | FLG_POOL_ENABLE_TAGGING | 0x400 | ptg | R |
| Gflags Flag Reference | FLG_ENABLE_SYSTEM_CRIT_BREAKS | 0x100000 | scb | R, K, I |

| | | | | |
|---|---|---|---|---|
| Gflags Flag Reference | FLG_LDR_TOP_DOWN | 0x20000000 | ltd | R, K, I |
| Gflags Flag Reference | FLG_MAINTAIN_OBJECT_TYPELIST | 0x4000 | otl | R,K |
| Gflags Flag Reference | FLG_SHOW_LDR_SNAPS | 0x2 | sls | R,K,I |
| Gflags Flag Reference | FLG_STOP_ON_EXCEPTION | 0x1 | soe | R,K,I |
| Gflags Flag Reference | FLG_STOP_ON_HUNG_GUI | 0x8 | shg | K |

**See Also**
**Concepts**
Gflags.exe
Gflags Remarks
Gflags UI
Gflags Examples
Gflags Flag Reference
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Gflags UI

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Gflags UI**
By using the Gflags dialog box, you can set and clear global flags from a user interface that lists all flags by name. There is no need to look up flag abbreviations or hexadecimal values.

Also, the dialog box provides access to the following features that are not available from the command line:

- Debugger. Specify that a particular program always runs in a debugger.
- Launch. Run a program with the specified debugging settings.
- Kernel Special Pool Tag. Detects kernel pool errors.

**Important**

- Pool tagging is permanently enabled on all versions of Windows Server 2003. On these systems, the **Enable pool tagging** check box on the Gflags dialog box is dimmed and commands to enable or disable pool tagging fail.

This section includes the following topics.

**Gflags dialog box**
The Gflags dialog box in System Registry mode.




**To open the dialog box**
Use one of the following methods:

- At the command line, type `gflags` (without parameters).

- In the **Run** dialog box, type `gflags` (without parameters).

**To set and clear systemwide registry flags**
Systemwide registry settings affect the entire system. They are saved in the registry and remain effective until you change them.

1. In the **Destination** area, click **System Registry**.
2. Set or clear a flag by selecting or clearing the check box associated with the flag.
3. When you have selected or cleared all the flags you want, click **Apply**.
4. Restart the system to make the changes effective.

**To set and clear kernel mode flags**
Kernel mode settings affect the entire system, but are effective only for the current session.

1. In the **Destination** area, click **Kernel Mode**.
2. Set or clear a flag by selecting or clearing the check box associated with the flag.
3. When you have selected or cleared all the flags you want, click **Apply**.

**To set and clear image file registry flags**
Image file settings affect instances of the specified program that start after the command completes. They are saved in the registry and remain effective until you change them.

1. In the **Destination** area, in the **Image File Name** box, type the name of the executable file or DLL, including the file name extension.
2. Click **Image File Options**.
3. Set or clear a flag by selecting or clearing the check box associated with the flag.
4. When you have selected or cleared all the flags you want, click **Apply**.

**To launch a program with flags**
This feature runs a program once with the specified flags. These settings affect only the instance of the program launched by the instruction in the **Command Line** box. They are not saved in the registry.

1. In the **Destination** area, in the **Command Line** box, type a command to start a program, including program parameters.
2. Set or clear a flag by selecting or clearing the check box associated with the flag.
3. When you have selected or cleared all the flags you want, click the **Launch** button.

### Notes

- Flags set in the registry do not affect the instance of the program that is launched.
- Flags set in the dialog box are used for the launched instance even when they are not image file flags.

**To run a program in a debugger**
This feature configures the program so that it always runs in a debugger with the specified options. This setting is saved in the registry. It affects all new instances of the program and remains effective until you change it.

The flags settings in the dialog box do not affect the debugger run.

1. In the **Destination** area, in the **Image File Name** box, type the name of the executable file, including the file name extension.

2. Click **Image File Options**.

3. In the **Image Debugger Options** area, click the **Debugger** check box.

4. In the **Debugger** box, type the command to run the debugger, including the program name and parameters. For example, **ntsd -d**.

5. Click **Apply**.

**To configure kernel special pool**

The kernel special pool feature uses the special pool for all allocations that have a specified tag or are within a specified size range. This feature helps you to detect and identify the source of errors in kernel pool use, such as writing beyond the allocated memory space, or referring to memory that has already been freed.

For information on using the special pool, see the *Windows Driver Development Kit*.

Kernel special pool settings are saved in the registry. They become effective when you restart the system and remain effective until you change them.

1. In the **Destination** area, click **System Registry**.

2. Choose one of the following methods:

   - To specify a pool tag in text format, click **Text**.

     Type a four-character pattern for the tag. The tag can include the **?** (substitute for any single character) and ***** (substitute for multiple characters) wildcard characters. For example, Fat* or Av?4.

     Gflags converts the ASCII characters to their hexadecimal format, placing the leftmost character in the lowest order position. For example, if you type **Tag1**, Gflags sets the tag to 0x31676154.

   - To specify a pool tag in hexadecimal format, click **Hex**.

     Type a hexadecimal value for the tag.

     Gflags sets the tag to the same numeric value. For example, if you type 12345678, Gflags stores pool tag 0x12345678.

     **Note**

Enter a hexadecimal value that is greater than the page size of the system. Otherwise, the value is interpreted as indicating a size range.

- To track kernel pool allocations of a specified size, click **Hex**.

  Allocations are organized into a range of pool block sizes. In 32-bit systems, pool blocks are divided into 8-byte ranges. In 64-bit systems, the range is 16 bytes.

  To indicate a range, type a hexadecimal number equal to the upper limit of the range plus 8 bytes for the pool header. For example, on a 32-bit system, to specify a pool block of between 9 and 16 bytes, type **18** to indicate 0x18.

  The following table shows the pattern of pool block ranges and the value you type to select that range for 32-bit systems.

| Size range (in bytes) | Value |
|---|---|
| 0 - 8 | 0x10 (16) |
| 9 - 16 | 0x18 (24) |
| 17 - 24 | 0x20 (32) |
| 25 - 32 | 0x28 (40) |
| ... | |

The following table shows the pattern of pool block ranges and the value you type to select that range for 64-bit systems.

| Size range (in bytes) | Value |
|---|---|
| 0 - 16 | 0x18 (24) |
| 17 - 32 | 0x28 (40) |
| 33 - 48 | 0x38 (56) |
| 49 - 64 | 0x46 (70) |
| ... | |

**Note**

The value for a size range must be equal to a multiple of 8 plus 8 (8x + 8) on 32-bit systems, or a multiple of 16 + 8 (16x + 8) on 64-bit systems. Also, the value must be smaller than the page size for the system. Otherwise, the value is interpreted as a hexadecimal tag.

1. Click **Verify Start** or **Verify End**.

   - **Verify Start** causes a system stop when a program tries to access tagged memory preceding its memory allocation.

   - **Verify End** causes a system stop when a program tries to access tagged memory beyond its memory allocation.

2. Click **Apply**.

3. To disable kernel special pool, delete the text or hexadecimal value from the **Kernel Special Pool Tag** box, and then click **Apply**.

**Note**

- The special pool can also be used for the memory pool allocations requested by a specific driver. This feature is activated by Driver Verifier (Verifier.exe), a tool included in the Windows\System32 directory in Windows 2000, Windows XP, and Windows Server 2003. For details, see the Driver Verifier documentation in the *Windows Driver Development Kit* (DDK).

**See Also**
**Concepts**
Gflags.exe
Gflags Remarks
Gflags Syntax
Gflags Examples
Gflags Flag Reference
Alphabetical List of Tools

Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Gflags Flag Reference

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Gflags Flag Reference**

**Caution**

- The symbolic name of each flag is provided for reference only. Because symbolic names change, they are not a reliable identifier of a global flag.

**Buffer DbgPrint output**

| **Abbreviation** | **ddp** |
|---|---|
| **Hexadecimal value** | 0x08000000 |
| **Symbolic Name** | FLG_DISABLE_DBGPRINT |
| **Destination** | Systemwide registry entry, kernel mode. |

Suppresses debugger output from DbgPrint(), DbgPrintEx(), KdPrint(), and KdPrintEx() calls. When this output is suppressed, it does not automatically appear in the kernel debugger. However, it can still be accessed by using the !dbgprint debugger extension.

**Create kernel mode stack trace database**

| **Abbreviation** | **kst** |
|---|---|
| **Hexadecimal value** | 0x2000 |
| **Symbolic Name** | FLG_KERNEL_STACK_TRACE_DB |
| **Destination** | Systemwide registry entry. |

Creates a run-time stack trace database of kernel operations, such as resource objects and object management operations. This feature works only when using a "checked build," that is, an internal debugging build of the operating system.

**Note**

- This flag is available as a kernel mode setting, but it is not effective in kernel mode because the kernel is already started.

**Create user mode stack trace database**

| **Abbreviation** | **ust** |
|---|---|
| **Hexadecimal value** | 0x1000 |
| **Symbolic Name** | FLG_USER_STACK_TRACE_DB |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Creates a run-time stack trace database in the address space of a particular process (image file mode) or all processes (systemwide).

**Debug initial command**

| **Abbreviation** | **dic** |
|---|---|
| **Hexadecimal value** | 0x4 |
| **Symbolic Name** | FLG_DEBUG_INITIAL_COMMAND |
| **Destination** | Systemwide registry entry, kernel mode. |

Runs Winlogon in the Windows Symbolic Debugger (Ntsd.exe) with the **-d** parameter, which directs its output to the kernel debugger console.

This flag is used for debugging the Client Server Run-time Subsystem (CSRSS), Winlogon.exe, and the Winlogon child processes.

See also: Debug Winlogon, Enable debugging of Win32 subsystem.

**Debug Winlogon**

| **Abbreviation** | **dwl** |
|---|---|
| **Hexadecimal value** | 0x04000000 |
| **Symbolic Name** | FLG_DEBUG_INITIAL_COMMAND_EX |
| **Destination** | Systemwide registry entry, kernel mode. |

Runs Winlogon in the Windows Symbolic Debugger (Ntsd.exe) with the following options:

- **-d** Send output to kernel debugger console.
- **-g** Ignore first breakpoint.
- **-x** Ignore first chance access violations.

See also: Debug Initial Command, Enable debugging of Win32 subsystem.

**Disable heap coalesce on free**

| **Abbreviation** | **dhc** |
|---|---|
| **Hexadecimal value** | 0x00200000 |
| **Symbolic Name** | FLG_HEAP_DISABLE_COALESCING |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Leaves adjacent blocks of heap memory separate when they are freed. By default, the system combines ("coalesces") newly freed adjacent blocks into a single block. Combining the blocks takes time, but reduces fragmentation that might force the heap to allocate additional memory when it can't find contiguous memory.

This flag is used to test optimizing strategies. It is not used for debugging.

**Disable paging of kernel stacks**

| **Abbreviation** | **dps** |
|---|---|
| **Hexadecimal value** | 0x80000 |
| **Symbolic Name** | FLG_DISABLE_PAGE_KERNEL_STACKS |
| **Destination** | Systemwide registry entry, kernel mode. |

Prevents paging of the kernel mode stacks of inactive threads. Generally, the kernel mode stack cannot be paged; it is guaranteed to be resident in memory. However, the system occasionally pages the kernel stacks of inactive threads. This flag prevents these occurrences.

The kernel debugger can provide information about a thread only when its stack is in physical memory. This flag is particularly important when debugging deadlocks and in other cases when every thread must be tracked.

**Disable protected DLL verification**

| **Abbreviation** | **dpd** |
|---|---|
| **Hexadecimal value** | 0x80000000 |

| | |
|---|---|
| **Symbolic Name** | FLG_DISABLE_PROTDLLS |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

This flag appears in Gflags, but it has no effect on the operating system.

**Disable stack extension**

| **Abbreviation** | **dse** |
|---|---|
| **Hexadecimal value** | 0x10000 |
| **Symbolic Name** | FLG_DISABLE_STACK_EXTENSION |
| **Destination** | Image file registry entry. |

Prevents the kernel from extending the stacks of the threads in the process beyond the initial memory committed. This is used to simulate low memory conditions (where stack extensions fail) and to test the strategic system processes that are expected to run well even with low memory.

**Early critical section event creation**

| **Abbreviation** | **cse** |
|---|---|
| **Hexadecimal value** | 0x10000000 |
| **Symbolic Name** | FLG_CRITSEC_EVENT_CREATION |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Creates event handles when a critical section is initialized, rather than waiting until the event is needed. When the system cannot create an event, it generates the exception during initialization and the calls to enter and leave the critical section do not fail.

Because this flag uses a significant amount of nonpaged pool memory, use it only on very reliable systems that have sufficient memory.

**Enable application verifier**

| **Abbreviation** | **vrf** |
|---|---|
| **Hexadecimal value** | 0x100 |
| **Symbolic Name** | FLG_APPLICATION_VERIFIER |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

**Enable bad handles detection**

| **Abbreviation** | **bhd** |
|---|---|
| **Hexadecimal value** | 0x40000000 |
| **Symbolic Name** | FLG_ENABLE_HANDLE_EXCEPTIONS |
| **Destination** | Systemwide registry entry, kernel mode. |

Raises a user mode exception (STATUS_INVALID_HANDLE) whenever a user mode process passes an invalid handle to the Object Manager.

**Enable close exception**

| **Abbreviation** | **ece** |
|---|---|
| **Hexadecimal value** | 0x00400000 |
| **Symbolic Name** | FLG_ENABLE_CLOSE_EXCEPTIONS |

| | |
|---|---|
| **Destination** | Systemwide registry entry, kernel mode. |

Raises a user mode exception whenever an invalid handle is passed to the CloseHandle() interface or related interfaces, such as SetEvent(), that take handles as arguments.

**Note**

- This flag is still supported, but it has been replaced by Enable bad handles detection, which performs a more comprehensive check of handle use.

#### Enable debugging of Win32 subsystem

| | |
|---|---|
| **Abbreviation** | **d32** |
| **Hexadecimal value** | 0x20000 |
| **Symbolic Name** | FLG_ENABLE_CSRDEBUG |
| **Destination** | Systemwide registry entry, kernel mode. |

Runs the Client Server Run-time Subsystem (Csrss.exe) in the Windows Symbolic Debugger (Ntsd.exe) with the following parameters:

- **-d** Send output to kernel debugger console.
- **-p 1** Attach the debugger to already running Csrss process, as identified by its process ID, 1.

This flag is effective only when the Debug Initial Command (dic) or Debug Winlogon (dwl) flag is also set.

#### Enable exception logging

| | |
|---|---|
| **Abbreviation** | **eel** |
| **Hexadecimal value** | 0x00800000 |
| **Symbolic Name** | FLG_ENABLE_EXCEPTION_LOGGING |
| **Destination** | Systemwide registry entry, kernel mode. |

Creates a log of exception records in the kernel run-time library. You can access the log from the kernel debugger.

#### Enable heap free checking

| | |
|---|---|
| **Abbreviation** | **hfc** |
| **Hexadecimal value** | 0x20 |
| **Symbolic Name** | FLG_HEAP_ENABLE_FREE_CHECK |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Validates the heap when it is freed.

See also: Enable heap tail checking, Enable heap parameter checking.

#### Enable heap parameter checking

| | |
|---|---|
| **Abbreviation** | **hpc** |
| **Hexadecimal value** | 0x40 |
| **Symbolic Name** | FLG_HEAP_VALIDATE_PARAMETERS |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Verifies some aspects of the heap whenever a heap API is called.

See also: Enable heap validation on call.

**Enable heap tagging**

| **Abbreviation** | **htg** |
|---|---|
| **Hexadecimal value** | 0x800 |
| **Symbolic Name** | FLG_HEAP_ENABLE_TAGGING |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Assigns unique tags to heap allocations. You can display the tag by using the !heap debugger extension with the **-t** parameter.

For information about the !heap extension, see the Microsoft Windows 2000 Debugging Tools Kit, which is available from the Microsoft Windows Driver Development Kits (DDK) Web site.

See also: Enable heap tagging by DLL.

**Enable heap tagging by DLL**

| **Abbreviation** | **htd** |
|---|---|
| **Hexadecimal value** | 0x8000 |
| **Symbolic Name** | FLG_HEAP_ENABLE_TAG_BY_DLL |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Assigns a unique tag to heap allocations created by the same DLL. You can display the tag by using the !heap debugger extension with the -t parameter.

For information about the !heap extension, see the Microsoft Windows 2000 Debugging Tools Kit, which is available from the Microsoft Windows Driver Development Kits (DDK) Web site.

See also: Enable heap tagging.

**Enable heap tail checking**

| **Abbreviation** | **htc** |
|---|---|
| **Hexadecimal value** | 0x10 |
| **Symbolic Name** | FLG_HEAP_ENABLE_TAIL_CHECK |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Checks for buffer overruns when the heap is freed. This flag adds a short pattern to the end of each allocation. The Windows heap manager detects the pattern when the block is freed and, if the block was modified, the heap manager breaks into the debugger.

See also: Enable heap free checking, Enable heap parameter checking.

**Enable heap validation on call**

| **Abbreviation** | **hvc** |
|---|---|
| **Hexadecimal value** | 0x80 |
| **Symbolic Name** | FLG_HEAP_VALIDATE_ALL |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Validates the entire heap each time a heap API is called.

To avoid the high overhead resulting from this flag, use the HeapValidate() API at critical junctures, such as when the heap is destroyed. However, this flag is useful for detecting random corruption in a pool.

See also: Enable heap parameter checking.

#### Enable loading of kernel debugger symbols

| Abbreviation | ksl |
|---|---|
| Hexadecimal value | 0x40000 |
| Symbolic Name | FLG_ENABLE_KDEBUG_SYMBOL_LOAD |
| Destination | Systemwide registry entry, kernel mode. |

Loads kernel symbols into the kernel memory space the next time the system starts. The kernel symbols are used in kernel profiling and by advanced kernel debugging tools.

#### Enable object handle type tagging

| Abbreviation | eot |
|---|---|
| Hexadecimal value | 0x01000000 |
| Symbolic Name | FLG_ENABLE_HANDLE_TYPE_TAGGING |
| Destination | Systemwide registry entry, kernel mode. |

This flag appears in Gflags, but it has no effect on the operating system.

#### Enable page heap

| Abbreviation | hpa |
|---|---|
| Hexadecimal value | 0x02000000 |
| Symbolic Name | FLG_HEAP_PAGE_ALLOCS |
| Destination | Systemwide registry entry, kernel mode, image file registry entry. |

Turns on page heap debugging, which verifies dynamic heap memory operations, including allocations and frees, and causes a debugger break when it detects a heap error.

This option enables full page heap debugging when set for image files and standard page heap debugging when set in system registry or kernel mode.

- *Full page heap debugging* (for **/i**) places an inaccessible page at the end of an allocation.
- *Standard page heap debugging* (for **/r** or **/k**) examines allocations as they are freed.

Setting this flag for an image file is the same as typing **gflags /p enable /full** for the image file at the command line.

#### Enable pool tagging

| Abbreviation | ptg |
|---|---|
| Hexadecimal value | 0x400 |
| Symbolic Name | FLG_POOL_ENABLE_TAGGING |
| Destination | Systemwide registry entry. |

**Important**

- This flag is permanently set in Windows Server 2003. On these systems, the **Enable pool tagging** check box on the Gflags dialog box is dimmed and commands to enable or disable pool tagging fail.

Collects data and calculates statistics about pool memory allocations. The data is grouped by pool tag value. Several tools that diagnose memory leaks and other kernel pool errors use the resulting data.

Use ExAllocatePoolWithTag or ExAllocatePoolWithQuotaTag to set the tag value. When no tag value is specified (ExAllocatePool, ExAllocatePoolWithQuota), the system creates a tag with the default value of "None." Because data for all

allocations with a "None" tag is combined, you cannot distinguish the data for a specific allocation.

In Windows XP and earlier systems, this flag also directs the system to attach a pool tag even when the pool memory is allocated by using ExAllocatePoolWithQuotaTag. Otherwise, the tag bytes are used to store the quota values. In Windows Server 2003, tag values and quota values are stored in separate fields that are attached to every pool memory allocation.

The **Enable pool tagging** flag is permanently set on all versions of Windows Server 2003. As a result, the system always collects data about pool memory allocations.

**Notes**

- For information on allocating pool memory, see the Windows Driver Development Kit.
- To display the data that the system collects about a tagged allocation, use Poolmon.exe, a tool in the Support Tools folder on Windows XP and Windows Server 2003.
- The description of the **Enable Pool Tagging** flag in Windows XP Support Tools is incomplete. This flag directs the system to collect and process data by tag value.

**Enable system critical breaks**

| **Abbreviation** | **scb** |
|---|---|
| **Hexadecimal value** | 0x100000 |
| **Symbolic Name** | FLG_ENABLE_SYSTEM_CRIT_BREAKS |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

- For per-process (image file) only: Forces a system breakpoint into the debugger whenever the specified process stops abnormally. This flag is effective only when the process calls the RtlSetProcessBreakOnExit() and RtlSetThreadBreakOnExit() interfaces.
- For systemwide (registry or kernel mode): Forces a system breakpoint to the debugger whenever processes that have called the RtlSetProcessBreakOnExit() and RtlSetThreadBreakOnExit() interfaces stop abnormally.

**Load DLLs top-down**

| **Abbreviation** | **ltd** |
|---|---|
| **Hexadecimal value** | 0x20000000 |
| **Symbolic Name** | FLG_LDR_TOP_DOWN |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Loads DLLs at the highest possible address. This flag is used to test 64-bit code for pointer truncation errors, because the most significant 32 bits of the pointers are not zeroes. It is designed for code running on the 64-bit versions of the Windows Server 2003.

**Maintain a list of objects for each type**

| **Abbreviation** | **otl** |
|---|---|
| **Hexadecimal value** | 0x4000 |
| **Symbolic Name** | FLG_MAINTAIN_OBJECT_TYPELIST |
| **Destination** | Systemwide registry entry, kernel mode. |

Collects and maintains a list of active objects by object type (for example, event, mutex, and semaphore). You can use Open Handles (Oh.exe), a tool included in the Windows 2000 Resource Kit, to display the object type lists. Because Oh.exe automatically sets the OTL flag, but does not clear it, you should use Gflags (-otl) to clear the flag.

**Notes**

- The linked lists created when you set this flag use eight bytes of overhead for each object. Remember to clear this flag when your analysis is complete.
- This flag is available as a kernel mode setting, but it is not effective in kernel mode because the kernel is already started.

**Show loader snaps**

| **Abbreviation** | **sls** |
|---|---|
| **Hexadecimal value** | 0x2 |
| **Symbolic Name** | FLG_SHOW_LDR_SNAPS |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

Captures detailed information about the loading and unloading of executable images and their supporting library modules.

- For systemwide (registry or kernel mode): Whenever a driver is loaded, this flag writes the kernel contents to the kernel debugger console. This feature works only when using a "checked build," that is, an internal debugging build of the operating system.
- For per-process (image file): Whenever a DLL is loaded, this flag writes the loader contents (and data related to DLL loading) to the program debugger console.

**Stop on exception**

| **Abbreviation** | **soe** |
|---|---|
| **Hexadecimal value** | 0x1 |
| **Symbolic Name** | FLG_STOP_ON_EXCEPTION |
| **Destination** | Systemwide registry entry, kernel mode, image file registry entry. |

The kernel breaks into the kernel debugger whenever a kernel mode exception occurs. The system passes all first chance exceptions (except for STATUS_PORT_DISCONNECT) with a severity of Warning or Error to the debugger before passing them to a local exception handler.

**Stop on hung GUI**

| **Abbreviation** | **shg** |
|---|---|
| **Hexadecimal value** | 0x8 |
| **Symbolic Name** | FLG_STOP_ON_HUNG_GUI |
| **Destination** | Kernel mode |

This flag appears in Gflags, but it has no effect on the operating system.

**See Also**
**Concepts**
Gflags.exe
Gflags Remarks
Gflags Syntax
Gflags UI
Gflags Examples
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview

Msicuu Overview
Memsnap Overview
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Pviewer Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Pviewer.exe: Process Viewer**
This GUI tool displays information about a running process and allows you to stop processes and change process priority.

### Note

- Process Viewer is similar to Pview.exe, but it can view processes on remote computers.

**Corresponding UI**
Process Viewer provides its own user interface.

**Concepts**
A process is a running instance of an application or executable file, along with all the system resources that have been allocated to that instance. In general, a process is equivalent to a single application or service; for example, Microsoft Word runs in the Winword.exe process.

Each process is composed of a set of threads. A thread is a unit of work that runs simultaneously with other units of work on the computer; each process must have at least one thread.

Threads represent the basic unit of execution in Windows Server 2003. When a thread starts, the Windows Server 2003 Memory Manager allocates enough physical memory and page file space to allow the thread to run. While a thread is running, it can request additional memory to enable it to complete its task. When a thread ends, it releases the memory it was using back to the Memory Manager for reallocation.

Memory is also allocated based on the number of threads spawned by each process. For example, consider a computer running two processes. Process A has 5 threads and Process B has 15 threads. Assuming the two processes have equal priority, Process A is allocated 25 percent (5 / 20) of the memory being used and Process B is allocated 75 percent (15 / 20) of the memory being used.

This method of thread allocation helps ensure that applications receive the memory they need. However, this same method can also result in memory leaks. A memory leak occurs when an application receives an increasing amount of memory but does not relinquish that memory back to the Memory Manager. This typically happens when a process creates threads but does not destroy those threads when they have finished their work.

**System Requirements**
The following are the system requirements for Process Viewer:

- Windows XP Professional or Windows Server 2003

### File required

- Pviewer.exe

**See Also**
**Concepts**
Pviewer Remarks
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Pviewer Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PViewer Remarks**

The **Process Viewer** dialog box contains the following elements:

**Memory Detail**

Click this button to see details about memory use for the selected process.

**Kill Process**

Click this button to stop the selected process.

**Caution**

- Do not attempt to stop processes required for running Windows 2000 or Windows NT. Make sure you understand which program owns a process before attempting to stop it.

**Computer**

Shows the name of the computer whose processes are currently displayed (default=local computer). To view the processes on a remote machine, type the UNC path to the computer in the space provided and then click the **Connect** button.

**Process**

The name of the process followed, in parentheses, by the process ID number in hexadecimal format.

**Privileged Time**

Privileged Time is the percentage of elapsed time that the threads of the process have been executing code in Privileged Mode, a processing mode that allows direct access to hardware and all memory. Privileged Mode is designed for operating system components, but Windows system services often run in Privileged Mode to gain access to data that is not accessible to User Mode threads. This value does not include Privileged Time that Windows components spend on behalf of the process.

**User Time**

User Time is the percentage of elapsed time that the threads of the process execute code in User Mode. User Mode is a restrictive processing mode designed for applications and environment subsystems. Application and system service threads spend most of their time executing in User Mode, but the operating system switches threads to Privileged Mode to gain access to operating system services.

**Process Memory Used**

The current number of bytes in the Working Set of this process. The Working Set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a certain threshold, pages are left in the Working Set of a process even if they are not in use. When free memory falls below the threshold, pages are trimmed from Working Sets. If they are needed, they will then be soft-faulted back into the Working Set before they leave main memory.

**Process Priority**

The current base priority of this process. Threads within a process can raise and lower their own base priority relative to the base priority of the process.

**Threads**

Thread State is the current state of the thread. It is 0 for Initialized, 1 for Ready, 2 for Running, 3 for Standby, 4 for Terminated, 5 for Wait, 6 for Transition, 7 for Unknown. A Running thread is using a processor; a Standby thread is about to use one. A Ready thread is waiting for a processor because none are free. A thread in Transition is waiting for a resource in order to execute, such as waiting for its execution stack to be paged in from disk. A Waiting thread has no use for the processor because it is waiting for a peripheral operation to complete, or a resource to become free.

**Thread Priority**

The current base priority of this thread. The system might raise the thread's dynamic priority above the base priority if the thread is handling user input, the priority can be lowered agian when a user is no longer interacting with the thread. Windows will never lower the thread below the base priority level.

**Thread Information**

**Context Switches** is the rate of switches from one thread to another in a second. Thread switches can occur either inside of a single process or across processes. A thread switch can be caused either by one thread asking another for information, or by a thread being preempted by another, higher-priority thread becoming ready to run.

**See Also**

**Concepts**

Pviewer Overview
Alphabetical List of Tools
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Printer and Fax Management Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Printer and Fax Management Tools**
Tools in this category are used to manage the following Windows features:

- Printers
- Print Services

### Notes

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Process and Service Management Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Process and Service Management Tools**
Tools in this category are used to manage the following Windows features:

- System or Network Monitoring
- Performance Counters
- Event Logs

**Notes**

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Memsnap Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Memsnap.exe: Memory Profiling Tool**
This command-line tool takes a snapshot of the memory resources being consumed by all running processes and writes this information to a log file.

MemSnap logs system memory usage to a log file with the default name of Memsnap.log. However, any file name can be chosen by specifying it at the command line when starting the tool.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Memory leaks occur when programs allocate memory but neglect to free it when the memory is no longer being used.

For more information about memory management and memory leaks, search the Microsoft Developer Network.

**System Requirements**
The following are the system requirements for MemSnap:

- Windows XP Professional or Windows Server 2003

**File Required**

- Memsnap.exe

**See Also**
**Concepts**
Memsnap Syntax
Memsnap Examples
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Memsnap Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MemSnap Examples**

**Example 1: Generate snapshot of memory resources**

To generate a snapshot of memory resources being consumed, type the following at the command line:

```
memsnap /m memsnap.log
```

This command sends the memory data to the memsnap.log file in the in the current directory.

**A sample memsnap log**

```
Process ID         Proc.Name Wrkng.Set PagedPool  NonPgdPl  Pagefile    Commit
Handles   Threads      User       Gdi
!LogType=memsnap
!ComputerName=ASEANL-TEST2
!buildnumber=3675
!buildtype=retail
!SystemTime=09\03\2002 15:39:31.0781 (GMT)
!TickCount=67179359
00000000              (null)     16384         0         0         0         0
0         1         0         0
00000004               System    233472         0         0         0     32768
738        50         0         0
00000124             smss.exe    499712      5500       600    172032    172032
24         2         0         0
00000170            csrss.exe   8376320     45408      4712   5165056   5165056
470        11         0         0
00000188         winlogon.exe   8220672     58864     46728   9465856   9465856
559        22        14        34
000001B4         services.exe   5763072    144696     10888   4149248   4149248
346        18         0         4
000001C0            lsass.exe  31318016     80252     71552  24829952  24829952
972        55         1         9
00000254          svchost.exe   3801088     20720     21528    995328    995328
201        10         0         4
0000028C          svchost.exe   7356416     33776      7680   2719744   2719744
228        27         0         5
0000034C          svchost.exe   5394432     31604      5544   3457024   3457024
140         9         0         0
00000374          svchost.exe   3031040     15252      2352    671744    671744
79         5         0         0
00000380          svchost.exe  20504576     72600     24392  12144640  12144640
1219        43         2         5
00000464          spoolsv.exe   9592832     38184      7232   5914624   5914624
222        15         0         5
000004CC            msdtc.exe   5820416     22640      5968   1617920   1617920
162        15         0         0
0000054C          certsrv.exe  10260480     39832     11640   6791168   6791168
247        13         2         8
00000580           dfssvc.exe   6217728     32072      5272   1732608   1732608
110        10         0         5
000005A0              dns.exe   6836224     35348     13256   6807552   6807552
166        12         0         5
000005D4          svchost.exe   2670592     15932      1952    471040    471040
53         2         0         4
00000610          ismserv.exe   4464640     28700      6088   1691648   1691648
119        10         0         5
00000630           llssrv.exe   6144000     30104      5088   1986560   1986560
118        10         0         0
00000684            ntfrs.exe   1015808     24308     11400   8994816   8994816
226        19         0         4
000006E0          svchost.exe   1622016      9908      1352    299008    299008
32         2         0         0
000006F0          locator.exe   2928640     25696      2272    643072    643072
41         3         0         0
```

```
00000860          dllhost.exe   8396800     25212      5488   2535424   2535424
234        19         1          4
00000AFC            csrss.exe   3133440     27420      3520    737280    737280
82         9          0          0
00000B14         winlogon.exe   5578752     36652      4056   1347584   1347584
138        10         0          0
00000C18          rdpclip.exe   3571712     27276      2080    741376    741376
56         3          3          5
00000C90         explorer.exe  16855040     51588      7488   5308416   5308416
285         9         70        158
0000073C              cmd.exe   1826816     15756      1336   1462272   1462272
23         1          0          4
00000AC8            logon.scr   1896448     16544      1400    364544    364544
13         1          0          0
00000B58          memsnap.exe   1372160     14272      1232    286720    286720
12         1          0          4
```

**Example 2: Generate snapshot of kernel memory pool information**

To generate a snapshot of the kernel memory pool type the following at the command line:

```
memsnap /p memsnap.log
```

This command sends the memory data to the memsnap.log located in the current directory.

**Sample memsnap /p output**

```
 Tag  Type     Allocs     Frees      Diff   Bytes  Per Alloc

!LogType=poolsnap
!ComputerName=RKTCLUST3
!buildnumber=3675
!buildtype=retail
!SystemTime=09\04\2002 18:38:42.0990 (GMT)
!TickCount=163617279
  TSQ Nonp         19        18         1      56         56
  waR Nonp          4         4         0       0          0
 1MEM Nonp          1         0         1    3400       3400
 2MEM Nonp          1         0         1    3944       3944
 3MEM Nonp          5         0         5     368         73
 8042 Nonp          4         0         4    3944        986
 8042 Paged        10        10         0       0          0
 ACPI Nonp         19        19         0       0          0
 AGP  Nonp          1         0         1     200        200
 AGP  Paged         7         5         2     448        224
 AcdN Nonp          2         0         2    1072        536
 AcpA Nonp         32        27         5     304         60
 AcpA Paged         1         0         1     504        504
 AcpB Nonp          2         2         0       0          0
 AcpB Paged        42        38         4     576        144
 AcpD Nonp         74        45        29    9320        321
 AcpF Nonp        206       205         1      40         40
 AcpI Nonp          3         3         0       0          0
 AcpM Nonp        949       949         0       0          0
 AcpM Paged         6         4         2     152         76
 AcpO Nonp          4         0         4     208         52
 AcpP Nonp         23        19         4     352         88
 AcpR Nonp          9         6         3    1616        538
 AcpR Paged       167       160         7    1040        148
 AcpS Nonp        185       151        34     864         25
 AcpS Paged        76        59        17     704         41
 Acpg Nonp          1         0         1      48         48
 Acpi Nonp         17        16         1     136        136
 Acpt Nonp          2         1         1      48         48
 Adap Nonp          9         0         9    1240        137
 AfdB Nonp         20         8        12   11864        988
 AfdC Nonp        376       298        78   12480        160
 AfdE Nonp        642       601        41   11480        280
 AfdI Nonp        340       340         0       0          0
 AfdL Nonp        535       507        28    1088         38
 AfdP Nonp       1659      1659         0       0          0
 AfdR Nonp          6         0         6     192         32
 AfdS Paged         2         2         0       0          0
```

| | | | | | | |
|---|---|---|---|---|---|---|
| AfdT | Nonp | 4 | 0 | 4 | 440 | 110 |
| AfdX | Paged | 619 | 586 | 33 | 9240 | 280 |
| Afdh | Nonp | 5 | 3 | 2 | 96 | 48 |
| Afdr | Nonp | 4 | 0 | 4 | 824 | 206 |
| Afdt | Paged | 5 | 0 | 5 | 792 | 158 |
| AgFM | Nonp | 1 | 0 | 1 | 40 | 40 |
| AgpG | Paged | 1 | 0 | 1 | 128 | 128 |
| AgpR | Paged | 1 | 1 | 0 | 0 | 0 |
| AmlC | Nonp | 40 | 36 | 4 | 32768 | 8192 |
| AmlH | Nonp | 1 | 0 | 1 | 65536 | 65536 |
| AmlI | Nonp | 2 | 0 | 2 | 392 | 196 |
| AmlI | Paged | 1 | 1 | 0 | 0 | 0 |
| AmlR | Nonp | 1 | 0 | 1 | 32 | 32 |
| AmlT | Nonp | 97 | 97 | 0 | 0 | 0 |
| ApcS | Nonp | 1 | 0 | 1 | 24 | 24 |
| ApcS | Paged | 2 | 2 | 0 | 0 | 0 |
| ArbA | Paged | 12 | 0 | 12 | 49152 | 4096 |
| ArbL | Paged | 50 | 24 | 26 | 6864 | 264 |
| ArbM | Nonp | 12 | 0 | 12 | 288 | 24 |
| ArbM | Paged | 36 | 36 | 0 | 0 | 0 |
| ArbR | Paged | 24 | 0 | 24 | 768 | 32 |
| AtmA | Paged | 784 | 180 | 604 | 37440 | 61 |
| AtmT | Paged | 9 | 0 | 9 | 5544 | 616 |
| BLes | Paged | 2 | 2 | 0 | 0 | 0 |
| Bmfd | Paged | 67 | 1 | 66 | 116688 | 1768 |
| CBRe | Nonp | 5 | 0 | 5 | 200 | 40 |
| CM | Nonp | 99 | 71 | 28 | 1120 | 40 |
| CM | Paged | 4046 | 4046 | 0 | 0 | 0 |
| CM1 | Paged | 30 | 30 | 0 | 0 | 0 |
| CM10 | Paged | 34 | 20 | 14 | 19488 | 1392 |
| CM11 | Paged | 18 | 16 | 2 | 8192 | 4096 |
| CM13 | Paged | 1 | 0 | 1 | 4096 | 4096 |
| CM16 | Paged | 480 | 423 | 57 | 237568 | 4167 |
| CM17 | Paged | 27 | 16 | 11 | 90112 | 8192 |
| CM18 | Paged | 1 | 0 | 1 | 4096 | 4096 |
| CM19 | Paged | 343 | 340 | 3 | 3480 | 1160 |
| CM2 | Paged | 30 | 30 | 0 | 0 | 0 |
| CM22 | Paged | 1 | 1 | 0 | 0 | 0 |
| CM23 | Paged | 1 | 0 | 1 | 8192 | 8192 |
| CM25 | Paged | 311 | 311 | 0 | 0 | 0 |
| CM26 | Paged | 14 | 4 | 10 | 81920 | 8192 |
| CM27 | Paged | 15 | 6 | 9 | 376 | 41 |
| CM28 | Paged | 1 | 0 | 1 | 4096 | 4096 |
| CM29 | Paged | 7 | 0 | 7 | 57344 | 8192 |
| CM3 | Paged | 30 | 30 | 0 | 0 | 0 |
| CM32 | Paged | 15 | 4 | 11 | 45056 | 4096 |
| CM35 | Paged | 6 | 0 | 6 | 118784 | 19797 |
| CM38 | Paged | 2881 | 2881 | 0 | 0 | 0 |
| CM39 | Paged | 1032 | 480 | 552 | 151872 | 275 |
| CM4 | Paged | 30 | 30 | 0 | 0 | 0 |
| CM44 | Nonp | 997 | 894 | 103 | 3264 | 31 |
| CM45 | Paged | 3 | 2 | 1 | 24 | 24 |
| CM7 | Paged | 15 | 15 | 0 | 0 | 0 |
| CM8 | Paged | 5108 | 5108 | 0 | 0 | 0 |
| CMAl | Paged | 1676 | 1638 | 38 | 155648 | 4096 |
| CMCa | Paged | 2 | 0 | 2 | 16384 | 8192 |
| CMDa | Paged | 18270 | 17278 | 992 | 226288 | 228 |
| CMDc | Paged | 1 | 0 | 1 | 24576 | 24576 |
| CMIn | Paged | 6400 | 6217 | 183 | 5136 | 28 |
| CMNb | Paged | 151054 | 150187 | 867 | 40984 | 47 |
| CMSb | Paged | 1 | 0 | 1 | 4096 | 4096 |
| CMSc | Paged | 275 | 83 | 192 | 47264 | 246 |
| CMVI | Paged | 15234 | 14546 | 688 | 23336 | 33 |
| CMVa | Paged | 99871 | 98025 | 1846 | 91936 | 49 |
| CMVw | Paged | 690 | 269 | 421 | 20208 | 48 |
| CMnb | Paged | 509 | 389 | 120 | 6720 | 56 |
| CMpa | Nonp | 991 | 892 | 99 | 6336 | 64 |
| CMpb | Paged | 1000 | 896 | 104 | 4992 | 48 |
| COMX | Nonp | 8 | 6 | 2 | 48 | 24 |
| COMX | Paged | 63 | 52 | 11 | 1032 | 93 |
| CPnp | Nonp | 91 | 0 | 91 | 22568 | 248 |
| CSdk | Nonp | 14 | 12 | 2 | 4248 | 2124 |

| | | | | | | |
|---|---|---|---|---|---|---|
| CSdk | Paged | 2 | 2 | 0 | 0 | 0 |
| CSnt | Nonp | 3411 | 3369 | 42 | 53176 | 1266 |
| CSnt | Paged | 2 | 2 | 0 | 0 | 0 |
| Call | Nonp | 5 | 0 | 5 | 360 | 72 |
| CcBc | Nonp | 40872 | 40856 | 16 | 2176 | 136 |
| CcPc | Nonp | 432 | 429 | 3 | 288 | 96 |
| CcPw | Nonp | 1 | 1 | 0 | 0 | 0 |
| CcSc | Nonp | 3905 | 3783 | 122 | 38064 | 312 |
| CcVa | Nonp | 1 | 0 | 1 | 139264 | 139264 |
| CcVl | Nonp | 6 | 0 | 6 | 3168 | 528 |
| CcVp | Nonp | 42 | 32 | 10 | 2336 | 233 |
| CcWK | Nonp | 1 | 0 | 1 | 136 | 136 |
| CcWk | Nonp | 292 | 284 | 8 | 192 | 24 |
| CcWq | Nonp | 8 | 0 | 8 | 192 | 24 |
| Cdic | Nonp | 1 | 0 | 1 | 72 | 72 |
| CmVn | Paged | 10553 | 10553 | 0 | 0 | 0 |
| Cont | Nonp | 2 | 0 | 2 | 144 | 72 |
| DG | Paged | 66 | 66 | 0 | 0 | 0 |
| DOPE | Nonp | 9 | 0 | 9 | 792 | 88 |
| Dacl | Paged | 9 | 8 | 1 | 96 | 96 |
| Ddk | Nonp | 7 | 7 | 0 | 0 | 0 |
| Ddk | Paged | 10 | 4 | 6 | 336 | 56 |
| Desk | Nonp | 12 | 0 | 12 | 2304 | 192 |
| Devi | Nonp | 252 | 34 | 218 | 187368 | 859 |
| DfsP | Nonp | 1 | 0 | 1 | 64 | 64 |
| DfsP | Paged | 1 | 0 | 1 | 2352 | 2352 |
| Dire | Paged | 57 | 18 | 39 | 8112 | 208 |
| Dmio | Nonp | 131 | 0 | 131 | 525720 | 4013 |
| Dmio | Paged | 4 | 3 | 1 | 128 | 128 |
| Dmip | Nonp | 131 | 0 | 131 | 17816 | 136 |
| Dmiq | Nonp | 1 | 0 | 1 | 264 | 264 |
| Dmld | Nonp | 1 | 1 | 0 | 0 | 0 |
| Dmst | Nonp | 2 | 2 | 0 | 0 | 0 |
| Dmst | Paged | 24 | 24 | 0 | 0 | 0 |
| Dnod | Nonp | 132 | 43 | 89 | 25632 | 288 |
| DrDr | Nonp | 81 | 75 | 6 | 344 | 57 |
| DrDr | Paged | 45 | 44 | 1 | 12288 | 12288 |
| DrEx | Nonp | 2 | 1 | 1 | 40 | 40 |
| DrIC | Nonp | 2 | 1 | 1 | 56 | 56 |
| Driv | Nonp | 102 | 4 | 98 | 23792 | 242 |
| Dump | Nonp | 6 | 0 | 6 | 37264 | 6210 |
| Dump | Paged | 1 | 1 | 0 | 0 | 0 |
| Dvgr | Paged | 2 | 2 | 0 | 0 | 0 |
| Envr | Nonp | 2 | 2 | 0 | 0 | 0 |
| Evel | Nonp | 1 | 0 | 1 | 24 | 24 |
| Even | Nonp | 142302 | 140326 | 1976 | 96672 | 48 |
| FLex | Nonp | 9 | 2 | 7 | 448 | 64 |
| FLfl | Paged | 58 | 45 | 13 | 936 | 72 |
| FLli | Nonp | 62 | 47 | 15 | 600 | 40 |
| FLln | Nonp | 1 | 0 | 1 | 40 | 40 |
| FLsh | Nonp | 1 | 0 | 1 | 56 | 56 |
| FSfm | Nonp | 1070 | 501 | 569 | 22760 | 40 |
| FSim | Paged | 1033 | 499 | 534 | 68352 | 128 |
| FSmg | Nonp | 3 | 0 | 3 | 1056 | 352 |
| FSrN | Nonp | 1249 | 1238 | 11 | 528 | 48 |
| FSrN | Paged | 1426 | 1308 | 118 | 11328 | 96 |
| FSrm | Nonp | 33 | 12 | 21 | 2688 | 128 |
| FSrm | Paged | 44 | 30 | 14 | 13072 | 933 |
| FSro | Nonp | 46 | 42 | 4 | 160 | 40 |
| FSro | Paged | 168 | 164 | 4 | 160 | 40 |
| FSrt | Nonp | 1 | 0 | 1 | 904 | 904 |
| FSun | Nonp | 2 | 2 | 0 | 0 | 0 |
| FatI | Nonp | 1 | 0 | 1 | 80 | 80 |
| File | Nonp | 225953 | 224319 | 1634 | 249680 | 152 |
| Flop | Nonp | 3 | 1 | 2 | 80 | 40 |
| Flop | Paged | 10 | 9 | 1 | 40 | 40 |
| Fsrc | Nonp | 3 | 2 | 1 | 24 | 24 |
| FstB | Nonp | 22 | 21 | 1 | 1208 | 1208 |
| Fstb | Nonp | 47 | 47 | 0 | 0 | 0 |
| Fstb | Paged | 6 | 6 | 0 | 0 | 0 |
| GDev | Paged | 10 | 0 | 10 | 33760 | 3376 |
| GFil | Paged | 9 | 5 | 4 | 328 | 82 |

| | | | | | | |
|---|---|---|---|---|---|---|
| GTmp | Paged | 9882 | 9882 | 0 | 0 | 0 |
| Galp | Paged | 235 | 235 | 0 | 0 | 0 |
| Gbaf | Paged | 197 | 172 | 25 | 17600 | 704 |
| Gcac | Paged | 703 | 659 | 44 | 246192 | 5595 |
| Gcsl | Paged | 2 | 0 | 2 | 976 | 488 |
| Gdbr | Paged | 102 | 100 | 2 | 160 | 80 |
| Gdcf | Paged | 879 | 853 | 26 | 1456 | 56 |
| Gdd | Paged | 5 | 5 | 0 | 0 | 0 |
| Gdde | Nonp | 2 | 0 | 2 | 80 | 40 |
| Gddf | Paged | 1 | 0 | 1 | 40 | 40 |
| Gddp | Paged | 3 | 2 | 1 | 96 | 96 |
| Gdev | Paged | 21 | 19 | 2 | 464 | 232 |
| Gdrs | Paged | 12613 | 12583 | 30 | 69768 | 2325 |
| Gdwd | Nonp | 12 | 0 | 12 | 1040 | 86 |
| Gdwd | Paged | 2 | 0 | 2 | 48 | 24 |
| Gebr | Paged | 22 | 15 | 7 | 2480 | 354 |
| Gedd | Paged | 6 | 6 | 0 | 0 | 0 |
| Geto | Paged | 21790 | 21790 | 0 | 0 | 0 |
| Gffv | Paged | 150 | 0 | 150 | 10800 | 72 |
| Gfil | Paged | 2 | 2 | 0 | 0 | 0 |
| Gfnt | Paged | 723 | 661 | 62 | 47616 | 768 |
| Gfsb | Paged | 42 | 40 | 2 | 16384 | 8192 |
| Gfsm | Nonp | 2 | 0 | 2 | 80 | 40 |
| Ggb | Paged | 692 | 656 | 36 | 59616 | 1656 |
| Ggdv | Paged | 8 | 0 | 8 | 976 | 122 |
| Gglb | Paged | 2 | 0 | 2 | 8192 | 4096 |
| Ggls | Paged | 366 | 132 | 234 | 22864 | 97 |
| Gh04 | Paged | 1 | 0 | 1 | 376 | 376 |
| Gh05 | Paged | 110 | 77 | 33 | 446088 | 13517 |
| Gh08 | Paged | 20 | 8 | 12 | 11552 | 962 |
| Gh09 | Paged | 2 | 0 | 2 | 1232 | 616 |
| Gh0< | Paged | 75 | 0 | 75 | 2400 | 32 |
| Gh0F | Paged | 1 | 1 | 0 | 0 | 0 |
| Gh14 | Paged | 1817 | 1814 | 3 | 1768 | 589 |
| Gh15 | Paged | 13674 | 13539 | 135 | 595296 | 4409 |
| Gh17 | Paged | 10 | 10 | 0 | 0 | 0 |
| Gh18 | Paged | 705 | 697 | 8 | 8960 | 1120 |
| Gh1< | Paged | 75 | 0 | 75 | 2400 | 32 |
| Gh1F | Paged | 6 | 6 | 0 | 0 | 0 |
| Ghab | Paged | 288 | 0 | 288 | 32256 | 112 |
| Ghas | Paged | 6 | 0 | 6 | 2640 | 440 |
| Ghtc | Paged | 4 | 4 | 0 | 0 | 0 |
| Gini | Nonp | 6 | 0 | 6 | 128 | 21 |
| Gla1 | Nonp | 2 | 0 | 2 | 224 | 112 |
| Gla1 | Paged | 304 | 156 | 148 | 236800 | 1600 |
| Gla4 | Nonp | 2 | 0 | 2 | 224 | 112 |
| Gla4 | Paged | 4648 | 4417 | 231 | 40656 | 176 |
| Gla5 | Nonp | 2 | 0 | 2 | 224 | 112 |
| Gla5 | Paged | 537 | 328 | 209 | 81928 | 392 |
| Gla8 | Nonp | 2 | 0 | 2 | 224 | 112 |
| Gla8 | Paged | 80 | 23 | 57 | 9120 | 160 |
| Gla: | Nonp | 2 | 0 | 2 | 224 | 112 |
| Gla: | Paged | 269 | 133 | 136 | 89216 | 656 |
| Gla; | Nonp | 2 | 0 | 2 | 224 | 112 |
| Gla@ | Nonp | 2 | 0 | 2 | 224 | 112 |
| Gla@ | Paged | 211 | 32 | 179 | 22912 | 128 |
| Gldv | Paged | 22 | 6 | 16 | 13888 | 868 |
| Glid | Paged | 2 | 2 | 0 | 0 | 0 |
| Glnk | Paged | 492 | 0 | 492 | 7872 | 16 |
| Gmap | Paged | 2 | 0 | 2 | 608 | 304 |
| Gpat | Paged | 1 | 0 | 1 | 4040 | 4040 |
| Gpfe | Paged | 150 | 0 | 150 | 25584 | 170 |
| Gpff | Paged | 150 | 0 | 150 | 28448 | 189 |
| Gpft | Paged | 4 | 0 | 4 | 960 | 240 |
| GreA | Paged | 2 | 0 | 2 | 64 | 32 |
| Gsem | Nonp | 1506 | 1324 | 182 | 14496 | 79 |
| Gsp | Paged | 75 | 59 | 16 | 2944 | 184 |
| Gspr | Paged | 5 | 3 | 2 | 880 | 440 |
| Gsth | Paged | 217 | 217 | 0 | 0 | 0 |
| Gtmp | Paged | 11028 | 11028 | 0 | 0 | 0 |
| Gtmw | Paged | 55 | 7 | 48 | 3840 | 80 |
| Gxlt | Paged | 15781 | 15765 | 16 | 1280 | 80 |

| | | | | | | |
|---|---|---|---|---|---|---|
| HT1H | Paged | 1 | 0 | 1 | 57344 | 57344 |
| HT40 | Paged | 2 | 0 | 2 | 16384 | 8192 |
| HT50 | Paged | 4 | 4 | 0 | 0 | 0 |
| HT52 | Paged | 1 | 1 | 0 | 0 | 0 |
| HT53 | Paged | 8 | 8 | 0 | 0 | 0 |
| Hal | Nonp | 113251 | 113248 | 3 | 9080 | 3026 |
| Hal | Paged | 600 | 600 | 0 | 0 | 0 |
| IPmm | Nonp | 1 | 1 | 0 | 0 | 0 |
| Ica | Nonp | 7440 | 7379 | 61 | 58088 | 952 |
| Ica | Paged | 2971 | 2971 | 0 | 0 | 0 |
| Icp | Nonp | 28 | 20 | 8 | 320 | 40 |
| IdeP | Nonp | 163690 | 163666 | 24 | 6816 | 284 |
| IdeP | Paged | 34 | 34 | 0 | 0 | 0 |
| IdeX | Nonp | 12 | 10 | 2 | 128 | 64 |
| IdeX | Paged | 22 | 22 | 0 | 0 | 0 |
| Ifs | Nonp | 266 | 264 | 2 | 128 | 64 |
| Ifs | Paged | 46 | 45 | 1 | 112 | 112 |
| ImRP | Paged | 1 | 1 | 0 | 0 | 0 |
| Info | Nonp | 4247 | 4246 | 1 | 32768 | 32768 |
| Info | Paged | 2 | 2 | 0 | 0 | 0 |
| Io | Nonp | 1 | 0 | 1 | 520 | 520 |
| Io | Nonp | 452234 | 452121 | 113 | 87600 | 775 |
| Io | Paged | 678 | 591 | 87 | 3568 | 41 |
| IoCc | Paged | 1689 | 1573 | 116 | 1856 | 16 |
| IoCo | Nonp | 172 | 108 | 64 | 4608 | 72 |
| IoEa | Nonp | 12130 | 12130 | 0 | 0 | 0 |
| IoEr | Nonp | 10 | 4 | 6 | 800 | 133 |
| IoFc | Paged | 38 | 38 | 0 | 0 | 0 |
| IoFs | Paged | 1 | 0 | 1 | 24 | 24 |
| IoFu | Paged | 1 | 1 | 0 | 0 | 0 |
| IoKB | Paged | 400 | 400 | 0 | 0 | 0 |
| IoNm | Paged | 299389 | 298490 | 899 | 115752 | 128 |
| IoOp | Nonp | 283619 | 283619 | 0 | 0 | 0 |
| IoRN | Paged | 50 | 50 | 0 | 0 | 0 |
| IoRi | Nonp | 13 | 13 | 0 | 0 | 0 |
| IoSe | Paged | 63 | 63 | 0 | 0 | 0 |
| IoSh | Nonp | 21 | 4 | 17 | 408 | 24 |
| IoTi | Nonp | 11 | 0 | 11 | 352 | 32 |
| IoUs | Nonp | 4 | 2 | 2 | 48 | 24 |
| Ioin | Nonp | 13 | 2 | 11 | 6864 | 624 |
| IpEQ | Nonp | 1 | 0 | 1 | 1808 | 1808 |
| IpIO | Nonp | 1 | 1 | 0 | 0 | 0 |
| IpLA | Nonp | 1 | 0 | 1 | 248 | 248 |
| IpSI | Nonp | 1027 | 0 | 1027 | 53272 | 51 |
| IpTI | Nonp | 3 | 0 | 3 | 5400 | 1800 |
| Irp | Nonp | 988 | 604 | 384 | 139016 | 362 |
| Isap | Nonp | 1 | 0 | 1 | 16 | 16 |
| Isap | Paged | 58 | 56 | 2 | 192 | 96 |
| Job | Nonp | 57 | 55 | 2 | 864 | 432 |
| KAPC | Nonp | 1 | 1 | 0 | 0 | 0 |
| KbdC | Nonp | 3 | 0 | 3 | 2544 | 848 |
| KbdC | Paged | 10 | 10 | 0 | 0 | 0 |
| Ke | Nonp | 54 | 53 | 1 | 296 | 296 |
| Ke | Paged | 1 | 0 | 1 | 8192 | 8192 |
| Key | Paged | 219499 | 218875 | 624 | 64848 | 103 |
| Keye | Paged | 1 | 0 | 1 | 64 | 64 |
| KrbC | Nonp | 890 | 886 | 4 | 360 | 90 |
| Ksec | Nonp | 6 | 0 | 6 | 5352 | 892 |
| Ksec | Paged | 4121 | 4119 | 2 | 64 | 32 |
| L2T1 | Nonp | 1 | 0 | 1 | 1064 | 1064 |
| L2T4 | Nonp | 1 | 0 | 1 | 4016 | 4016 |
| L2Ti | Nonp | 5 | 3 | 2 | 96 | 48 |
| LBan | Paged | 4 | 4 | 0 | 0 | 0 |
| LBbl | Nonp | 2 | 2 | 0 | 0 | 0 |
| LBbn | Paged | 3 | 0 | 3 | 192 | 64 |
| LBbs | Paged | 6 | 0 | 6 | 192 | 32 |
| LBci | Nonp | 14 | 14 | 0 | 0 | 0 |
| LBdi | Nonp | 1 | 0 | 1 | 168 | 168 |
| LBds | Nonp | 14 | 14 | 0 | 0 | 0 |
| LBea | Paged | 6 | 6 | 0 | 0 | 0 |
| LBma | Nonp | 2966 | 2966 | 0 | 0 | 0 |
| LBmb | Nonp | 9 | 6 | 3 | 1752 | 584 |

| | | | | | | |
|---|---|---|---|---|---|---|
| LBmh | Nonp | 14 | 14 | 0 | 0 | 0 |
| LBne | Paged | 41 | 41 | 0 | 0 | 0 |
| LBpn | Paged | 6 | 0 | 6 | 240 | 40 |
| LBpt | Paged | 2 | 0 | 2 | 816 | 408 |
| LBsl | Paged | 2 | 0 | 2 | 40 | 20 |
| LBtn | Nonp | 6 | 0 | 6 | 432 | 72 |
| LBwi | Nonp | 5444 | 5444 | 0 | 0 | 0 |
| LBxp | Nonp | 2 | 0 | 2 | 1152 | 576 |
| LSac | Paged | 4 | 0 | 4 | 168 | 42 |
| LSas | Nonp | 8 | 8 | 0 | 0 | 0 |
| LSbf | Paged | 21 | 21 | 0 | 0 | 0 |
| LScn | Nonp | 38 | 0 | 38 | 26144 | 688 |
| LSdb | Nonp | 3031 | 3028 | 3 | 312 | 104 |
| LSdb | Paged | 41 | 36 | 5 | 608 | 121 |
| LSep | Nonp | 8 | 0 | 8 | 3120 | 390 |
| LSfn | Paged | 2 | 0 | 2 | 64 | 32 |
| LSlf | Paged | 50 | 47 | 3 | 360 | 120 |
| LSlr | Paged | 2 | 2 | 0 | 0 | 0 |
| LSmf | Paged | 62 | 57 | 5 | 1432 | 286 |
| LSmi | Paged | 10350 | 10349 | 1 | 168 | 168 |
| LSnh | Nonp | 5498 | 5483 | 15 | 480 | 32 |
| LSni | Paged | 1 | 0 | 1 | 144 | 144 |
| LSpc | Paged | 36 | 0 | 36 | 6336 | 176 |
| LSpm | Nonp | 30 | 27 | 3 | 288 | 96 |
| LSpr | Paged | 23 | 22 | 1 | 64 | 64 |
| LSrf | Nonp | 45 | 41 | 4 | 992 | 248 |
| LSsd | Paged | 1 | 0 | 1 | 136 | 136 |
| LSsh | Nonp | 10 | 0 | 10 | 720 | 72 |
| LSsh | Paged | 5 | 0 | 5 | 2040 | 408 |
| LSsr | Paged | 6 | 1 | 5 | 208 | 41 |
| LSss | Paged | 37 | 35 | 2 | 216 | 108 |
| LStb | Nonp | 46 | 1 | 45 | 5984 | 132 |
| LStb | Paged | 112 | 0 | 112 | 7744 | 69 |
| LStc | Paged | 12 | 11 | 1 | 112 | 112 |
| LStr | Paged | 41281 | 41264 | 17 | 3840 | 225 |
| LSwi | Nonp | 1 | 0 | 1 | 995328 | 995328 |
| LSwq | Nonp | 2 | 0 | 2 | 1248 | 624 |
| LSwr | Nonp | 64 | 0 | 64 | 203264 | 3176 |
| Lfs | Paged | 9822 | 9769 | 53 | 5016 | 94 |
| LfsI | Paged | 2 | 0 | 2 | 131072 | 65536 |
| LfsS | Nonp | 2 | 0 | 2 | 208 | 104 |
| LfsS | Paged | 2 | 0 | 2 | 736 | 368 |
| Lfsr | Nonp | 2 | 0 | 2 | 32768 | 16384 |
| List | Nonp | 15 | 0 | 15 | 360 | 24 |
| LpcM | Paged | 19 | 15 | 4 | 1440 | 360 |
| MRXx | Paged | 11 | 11 | 0 | 0 | 0 |
| Mapr | Nonp | 12 | 11 | 1 | 2056 | 2056 |
| Mapr | Paged | 2 | 2 | 0 | 0 | 0 |
| Mdl | Nonp | 1522 | 674 | 848 | 108544 | 128 |
| Mdp | Nonp | 8 | 0 | 8 | 1536 | 192 |
| Mm | Nonp | 15 | 0 | 15 | 28424 | 1894 |
| Mm | Paged | 57 | 51 | 6 | 864 | 144 |
| MmCa | Nonp | 62939 | 62478 | 461 | 45208 | 98 |
| MmCi | Nonp | 387 | 114 | 273 | 59224 | 216 |
| MmCm | Nonp | 18 | 7 | 11 | 311840 | 28349 |
| MmCx | Nonp | 3 | 3 | 0 | 0 | 0 |
| MmDT | Nonp | 20 | 17 | 3 | 1272 | 424 |
| MmDT | Paged | 193 | 82 | 111 | 7928 | 71 |
| MmDb | Nonp | 1 | 0 | 1 | 144 | 144 |
| MmDm | Nonp | 8 | 0 | 8 | 640 | 80 |
| MmEx | Nonp | 1 | 0 | 1 | 32 | 32 |
| MmHi | Nonp | 9 | 3 | 6 | 240 | 40 |
| MmHt | Nonp | 2 | 0 | 2 | 1040 | 520 |
| MmIn | Nonp | 7 | 0 | 7 | 1120 | 160 |
| MmLd | Nonp | 105 | 7 | 98 | 11072 | 112 |
| MmLn | Nonp | 86 | 86 | 0 | 0 | 0 |
| MmMl | Nonp | 7 | 6 | 1 | 32 | 32 |
| MmPd | Nonp | 1 | 0 | 1 | 4096 | 4096 |
| MmRl | Nonp | 30 | 30 | 0 | 0 | 0 |
| MmSP | Nonp | 1 | 0 | 1 | 12288 | 12288 |
| MmSc | Nonp | 1 | 0 | 1 | 56 | 56 |
| MmSd | Nonp | 57 | 45 | 12 | 672 | 56 |

| | | | | | | |
|---|---|---|---|---|---|---|
| MmSe | Nonp | 247 | 245 | 2 | 64 | 32 |
| MmSg | Paged | 1 | 0 | 1 | 72 | 72 |
| MmSm | Paged | 1678 | 1386 | 292 | 16352 | 56 |
| MmSt | Paged | 2927 | 2184 | 743 | 535632 | 720 |
| MmSy | Nonp | 42 | 27 | 15 | 840 | 56 |
| MmWe | Nonp | 20 | 0 | 20 | 3360 | 168 |
| Mmdl | Nonp | 8 | 8 | 0 | 0 | 0 |
| Mmpr | Nonp | 2 | 0 | 2 | 80 | 40 |
| Mmpv | Nonp | 10 | 7 | 3 | 120 | 40 |
| MntA | Nonp | 16 | 10 | 6 | 2696 | 449 |
| MntA | Paged | 1477 | 1443 | 34 | 3016 | 88 |
| MntB | Paged | 6 | 6 | 0 | 0 | 0 |
| MsFD | Nonp | 1 | 0 | 1 | 264 | 264 |
| MsFD | Paged | 1 | 0 | 1 | 16 | 16 |
| MsFN | Paged | 65 | 65 | 0 | 0 | 0 |
| MsFc | Nonp | 15 | 15 | 0 | 0 | 0 |
| MsFd | Paged | 15 | 15 | 0 | 0 | 0 |
| MsFf | Nonp | 65 | 65 | 0 | 0 | 0 |
| MsFg | Nonp | 1 | 0 | 1 | 64 | 64 |
| MsvC | Nonp | 2038 | 2016 | 22 | 9112 | 414 |
| MsvC | Paged | 28 | 28 | 0 | 0 | 0 |
| MuoC | Nonp | 3 | 0 | 3 | 4944 | 1648 |
| MuoC | Paged | 8 | 8 | 0 | 0 | 0 |
| Mup | Nonp | 11184 | 11181 | 3 | 760 | 253 |
| Mup | Paged | 17686 | 17399 | 287 | 25464 | 88 |
| MupI | Nonp | 3 | 0 | 3 | 144 | 48 |
| Muta | Nonp | 888 | 733 | 155 | 10496 | 67 |
| NBSf | Nonp | 2 | 0 | 2 | 2280 | 1140 |
| NBSn | Nonp | 2 | 0 | 2 | 256 | 128 |
| NBSr | Paged | 2 | 2 | 0 | 0 | 0 |
| NBqh | Nonp | 5 | 0 | 5 | 160 | 32 |
| ND | Nonp | 334 | 334 | 0 | 0 | 0 |
| NDPt | Nonp | 2 | 0 | 2 | 4176 | 2088 |
| NDPw | Nonp | 2 | 0 | 2 | 4176 | 2088 |
| NDam | Nonp | 10530 | 10530 | 0 | 0 | 0 |
| NDan | Nonp | 20 | 2 | 18 | 1752 | 97 |
| NDar | Nonp | 2 | 0 | 2 | 480 | 240 |
| NDca | Paged | 3 | 3 | 0 | 0 | 0 |
| NDch | Nonp | 9 | 8 | 1 | 208 | 208 |
| NDco | Nonp | 24 | 12 | 12 | 1216 | 101 |
| NDcs | Nonp | 14 | 9 | 5 | 120 | 24 |
| NDcs | Paged | 3 | 3 | 0 | 0 | 0 |
| NDda | Paged | 3 | 3 | 0 | 0 | 0 |
| NDdb | Nonp | 713 | 703 | 10 | 2352 | 235 |
| NDfa | Nonp | 12 | 6 | 6 | 656 | 109 |
| NDfm | Nonp | 7 | 0 | 7 | 784 | 112 |
| NDlp | Nonp | 8497 | 8497 | 0 | 0 | 0 |
| NDmo | Nonp | 13 | 0 | 13 | 2992 | 230 |
| NDoa | Nonp | 3 | 0 | 3 | 1840 | 613 |
| NDoc | Nonp | 5275 | 5275 | 0 | 0 | 0 |
| NDpb | Nonp | 6 | 0 | 6 | 1344 | 224 |
| NDpc | Nonp | 4 | 3 | 1 | 192 | 192 |
| NDpf | Nonp | 21 | 0 | 21 | 5928 | 282 |
| NDpk | Nonp | 8 | 0 | 8 | 680 | 85 |
| NDpn | Nonp | 36 | 36 | 0 | 0 | 0 |
| NDpp | Nonp | 33 | 0 | 33 | 95008 | 2879 |
| NDrd | Nonp | 79 | 79 | 0 | 0 | 0 |
| NDsd | Paged | 3 | 3 | 0 | 0 | 0 |
| NDse | Nonp | 18 | 18 | 0 | 0 | 0 |
| NDse | Paged | 5 | 0 | 5 | 928 | 185 |
| NDst | Nonp | 20 | 17 | 3 | 88 | 29 |
| NDw1 | Nonp | 9 | 1 | 8 | 10832 | 1354 |
| NDw2 | Nonp | 8 | 8 | 0 | 0 | 0 |
| NDw3 | Nonp | 23 | 23 | 0 | 0 | 0 |
| NDwh | Nonp | 7 | 0 | 7 | 952 | 136 |
| NDwi | Nonp | 10992 | 10992 | 0 | 0 | 0 |
| NPLL | Nonp | 2 | 0 | 2 | 208 | 104 |
| Nb01 | Nonp | 2 | 0 | 2 | 2208 | 1104 |
| Nb02 | Nonp | 249 | 248 | 1 | 40 | 40 |
| Nb03 | Nonp | 437 | 436 | 1 | 16 | 16 |
| Nb05 | Nonp | 20 | 0 | 20 | 3200 | 160 |
| Nb07 | Nonp | 14 | 0 | 14 | 4256 | 304 |

| | | | | | | |
|---|---|---|---|---|---|---|
| Nb16 | Nonp | 14 | 7 | 7 | 1120 | 160 |
| Nb17 | Nonp | 3 | 2 | 1 | 120 | 120 |
| Nb22 | Nonp | 3233 | 3233 | 0 | 0 | 0 |
| Nb25 | Nonp | 8 | 8 | 0 | 0 | 0 |
| Nb26 | Nonp | 14 | 14 | 0 | 0 | 0 |
| Nb27 | Nonp | 14 | 12 | 2 | 1040 | 520 |
| Nb28 | Nonp | 8 | 8 | 0 | 0 | 0 |
| Nb29 | Nonp | 43 | 43 | 0 | 0 | 0 |
| Nb30 | Nonp | 385 | 385 | 0 | 0 | 0 |
| Nb32 | Nonp | 7 | 7 | 0 | 0 | 0 |
| Nb33 | Nonp | 59 | 59 | 0 | 0 | 0 |
| Nb36 | Nonp | 8 | 7 | 1 | 32 | 32 |
| Nb38 | Nonp | 16 | 16 | 0 | 0 | 0 |
| Nb39 | Nonp | 4 | 0 | 4 | 320 | 80 |
| NbL0 | Nonp | 69 | 48 | 21 | 3696 | 176 |
| NbL1 | Nonp | 41 | 20 | 21 | 3696 | 176 |
| NbL2 | Nonp | 1166 | 1040 | 126 | 22176 | 176 |
| NbL3 | Nonp | 1127 | 1118 | 9 | 1584 | 176 |
| NbL4 | Nonp | 150 | 105 | 45 | 7920 | 176 |
| NbTA | Nonp | 2875 | 2875 | 0 | 0 | 0 |
| Nba9 | Nonp | 8 | 8 | 0 | 0 | 0 |
| Nbt | Paged | 1 | 1 | 0 | 0 | 0 |
| Nbt0 | Nonp | 445 | 436 | 9 | 1368 | 152 |
| Nbt2 | Nonp | 3019 | 3019 | 0 | 0 | 0 |
| Nbt4 | Nonp | 18 | 18 | 0 | 0 | 0 |
| Nbt8 | Nonp | 697 | 696 | 1 | 16 | 16 |
| NbtC | Nonp | 9 | 0 | 9 | 1080 | 120 |
| NbtD | Nonp | 1166 | 1120 | 46 | 6992 | 152 |
| NbtF | Nonp | 2455 | 2455 | 0 | 0 | 0 |
| NbtG | Nonp | 1158 | 1158 | 0 | 0 | 0 |
| NbtH | Nonp | 743 | 743 | 0 | 0 | 0 |
| NbtK | Nonp | 10308 | 10308 | 0 | 0 | 0 |
| NbtL | Nonp | 10317 | 10317 | 0 | 0 | 0 |
| NbtM | Nonp | 1 | 1 | 0 | 0 | 0 |
| NbtN | Nonp | 240 | 240 | 0 | 0 | 0 |
| NbtO | Nonp | 65304 | 65304 | 0 | 0 | 0 |
| NbtR | Nonp | 513 | 512 | 1 | 152 | 152 |
| NbtS | Nonp | 45 | 45 | 0 | 0 | 0 |
| NbtV | Nonp | 40 | 40 | 0 | 0 | 0 |
| NbtX | Nonp | 1476 | 1476 | 0 | 0 | 0 |
| NbtY | Nonp | 1 | 1 | 0 | 0 | 0 |
| NbtZ | Nonp | 5 | 5 | 0 | 0 | 0 |
| Nbta | Nonp | 1 | 0 | 1 | 40 | 40 |
| Nbtc | Nonp | 780 | 780 | 0 | 0 | 0 |
| Nbte | Nonp | 1128 | 1128 | 0 | 0 | 0 |
| Nbtg | Nonp | 2 | 0 | 2 | 32 | 16 |
| Nbti | Nonp | 68 | 68 | 0 | 0 | 0 |
| Nbtj | Nonp | 1138 | 1138 | 0 | 0 | 0 |
| Nbtk | Nonp | 1138 | 1138 | 0 | 0 | 0 |
| Nbtl | Nonp | 2553 | 2331 | 222 | 31968 | 144 |
| Nbtm | Nonp | 2553 | 2553 | 0 | 0 | 0 |
| Nbtn | Nonp | 1157 | 1157 | 0 | 0 | 0 |
| Nbtq | Nonp | 2 | 2 | 0 | 0 | 0 |
| Nbtr | Nonp | 2 | 2 | 0 | 0 | 0 |
| Nbtt | Nonp | 18 | 18 | 0 | 0 | 0 |
| Nbtw | Nonp | 5 | 0 | 5 | 920 | 184 |
| Nls | Nonp | 1 | 1 | 0 | 0 | 0 |
| NpFC | Paged | 19 | 17 | 2 | 48 | 24 |
| NpFD | Paged | 1 | 0 | 1 | 112 | 112 |
| NpFR | Nonp | 53136 | 53136 | 0 | 0 | 0 |
| NpFW | Paged | 17 | 17 | 0 | 0 | 0 |
| NpFc | Nonp | 225 | 138 | 87 | 6960 | 80 |
| NpFc | Paged | 225 | 138 | 87 | 11832 | 136 |
| NpFi | Paged | 20 | 20 | 0 | 0 | 0 |
| NpFn | Nonp | 58 | 18 | 40 | 2168 | 54 |
| NpFn | Paged | 32 | 32 | 0 | 0 | 0 |
| NpFr | Nonp | 64969 | 64928 | 41 | 1640 | 40 |
| NpFs | Paged | 31767 | 31727 | 40 | 2880 | 72 |
| NpFt | Nonp | 1255 | 1255 | 0 | 0 | 0 |
| NpFw | Nonp | 1 | 1 | 0 | 0 | 0 |
| NpFw | Paged | 126 | 126 | 0 | 0 | 0 |
| NpfF | Paged | 58 | 18 | 40 | 4800 | 120 |

| | | | | | | |
|---|---|---|---|---|---|---|
| Npfs | Nonp | 1 | 0 | 1 | 272 | 272 |
| NtFA | Nonp | 591 | 591 | 0 | 0 | 0 |
| NtFA | Paged | 59980 | 59980 | 0 | 0 | 0 |
| NtFB | Paged | 473 | 463 | 10 | 28680 | 2868 |
| NtFC | Paged | 27415 | 27375 | 40 | 3096 | 77 |
| NtFE | Paged | 38 | 38 | 0 | 0 | 0 |
| NtFF | Paged | 92 | 92 | 0 | 0 | 0 |
| NtFI | Nonp | 4 | 4 | 0 | 0 | 0 |
| NtFI | Paged | 32383 | 32367 | 16 | 2064 | 129 |
| NtFL | Nonp | 21649 | 21641 | 8 | 8816 | 1102 |
| NtFL | Paged | 97 | 26 | 71 | 2840 | 40 |
| NtFS | Paged | 331 | 163 | 168 | 52744 | 313 |
| NtFd | Paged | 3626 | 3604 | 22 | 2016 | 91 |
| NtFf | Nonp | 1 | 0 | 1 | 40 | 40 |
| NtFf | Paged | 10 | 3 | 7 | 131264 | 18752 |
| NtFs | Nonp | 2366 | 1580 | 786 | 32016 | 40 |
| NtFs | Paged | 1739 | 889 | 850 | 44896 | 52 |
| NtFv | Paged | 72 | 72 | 0 | 0 | 0 |
| Ntf0 | Nonp | 3 | 0 | 3 | 196608 | 65536 |
| Ntf0 | Paged | 10627 | 10013 | 614 | 26952 | 43 |
| Ntf9 | Nonp | 54056 | 54056 | 0 | 0 | 0 |
| NtfC | Paged | 202 | 87 | 115 | 13800 | 120 |
| NtfD | Paged | 1 | 0 | 1 | 168 | 168 |
| NtfE | Paged | 78 | 53 | 25 | 5000 | 200 |
| NtfF | Paged | 376 | 109 | 267 | 254184 | 952 |
| NtfI | Nonp | 175 | 171 | 4 | 160 | 40 |
| NtfR | Nonp | 1 | 0 | 1 | 24 | 24 |
| NtfS | Paged | 6 | 0 | 6 | 2544 | 424 |
| NtfT | Nonp | 16 | 12 | 4 | 352 | 88 |
| NtfV | Nonp | 2 | 0 | 2 | 192 | 96 |
| Ntfc | Nonp | 1 | 0 | 1 | 72 | 72 |
| Ntfc | Paged | 389 | 252 | 137 | 9864 | 72 |
| Ntff | Nonp | 9 | 0 | 9 | 1872 | 208 |
| Ntff | Paged | 1190 | 693 | 497 | 413504 | 832 |
| Ntfi | Nonp | 1076 | 852 | 224 | 60928 | 272 |
| Ntfl | Paged | 12 | 0 | 12 | 1920 | 160 |
| Ntfn | Nonp | 1616 | 838 | 778 | 31432 | 40 |
| Ntfo | Paged | 515 | 262 | 253 | 24536 | 96 |
| Ntfr | Nonp | 1387 | 0 | 1387 | 89736 | 64 |
| Ntfs | Paged | 2 | 0 | 2 | 704 | 352 |
| Ntft | Nonp | 18 | 4 | 14 | 5216 | 372 |
| Nuio | Nonp | 4 | 0 | 4 | 560 | 140 |
| Ob | Paged | 1 | 1 | 0 | 0 | 0 |
| ObCI | Nonp | 1 | 0 | 1 | 136 | 136 |
| ObCi | Nonp | 7 | 0 | 7 | 392 | 56 |
| ObDi | Paged | 1646 | 860 | 786 | 12576 | 16 |
| ObDm | Paged | 15 | 10 | 5 | 280 | 56 |
| ObHd | Paged | 1191 | 1113 | 78 | 3072 | 39 |
| ObNM | Nonp | 1 | 0 | 1 | 136 | 136 |
| ObNm | Paged | 42629 | 41828 | 801 | 38280 | 47 |
| ObSc | Paged | 2009 | 1810 | 199 | 30176 | 151 |
| ObSq | Paged | 219595 | 219595 | 0 | 0 | 0 |
| ObSt | Nonp | 842084 | 842084 | 0 | 0 | 0 |
| ObjT | Nonp | 29 | 0 | 29 | 13456 | 464 |
| Obtb | Paged | 441 | 355 | 86 | 199648 | 2321 |
| PSwt | Nonp | 1 | 0 | 1 | 32 | 32 |
| PTIa | Nonp | 1 | 0 | 1 | 1024 | 1024 |
| PTIw | Nonp | 1 | 0 | 1 | 48 | 48 |
| PTPa | Nonp | 1 | 0 | 1 | 544 | 544 |
| PTPc | Nonp | 1 | 0 | 1 | 4008 | 4008 |
| PX2 | Nonp | 1 | 0 | 1 | 2008 | 2008 |
| PX3 | Nonp | 2 | 0 | 2 | 1200 | 600 |
| PX5 | Nonp | 2 | 0 | 2 | 144 | 72 |
| PX9 | Nonp | 2 | 0 | 2 | 224 | 112 |
| PXa | Nonp | 2 | 2 | 0 | 0 | 0 |
| PXb | Nonp | 2 | 0 | 2 | 1280 | 640 |
| PXc | Nonp | 2 | 2 | 0 | 0 | 0 |
| PXd | Nonp | 2 | 0 | 2 | 688 | 344 |
| PXg | Nonp | 2 | 0 | 2 | 144 | 72 |
| PXh | Nonp | 2 | 0 | 2 | 256 | 128 |
| PXl | Nonp | 1 | 0 | 1 | 2008 | 2008 |
| ParP | Nonp | 4 | 1 | 3 | 88 | 29 |

| | | | | | | |
|---|---|---|---|---|---|---|
| ParP | Paged | 16 | 10 | 6 | 288 | 48 |
| PciB | Nonp | 20 | 11 | 9 | 10640 | 1182 |
| PciB | Paged | 128 | 112 | 16 | 2520 | 157 |
| PlBd | Nonp | 2 | 0 | 2 | 4192 | 2096 |
| PnPi | Nonp | 1 | 0 | 1 | 80 | 80 |
| PoSL | Paged | 2 | 0 | 2 | 32 | 16 |
| PooL | Nonp | 1 | 0 | 1 | 8192 | 8192 |
| Pool | Nonp | 3 | 0 | 3 | 114688 | 38229 |
| Port | Nonp | 1331 | 1156 | 175 | 5600 | 32 |
| Port | Paged | 2593 | 2290 | 303 | 56408 | 186 |
| Pp | Nonp | 8 | 8 | 0 | 0 | 0 |
| Pp | Paged | 14558 | 14235 | 323 | 31576 | 97 |
| PpEB | Nonp | 1 | 0 | 1 | 192 | 192 |
| PpEE | Paged | 155 | 155 | 0 | 0 | 0 |
| PpEL | Nonp | 1 | 0 | 1 | 88 | 88 |
| PpUB | Nonp | 2 | 1 | 1 | 144 | 144 |
| PpWI | Nonp | 72 | 72 | 0 | 0 | 0 |
| Ppdd | Paged | 114 | 114 | 0 | 0 | 0 |
| Ppen | Nonp | 76 | 74 | 2 | 48 | 24 |
| Ppen | Paged | 174 | 136 | 38 | 3528 | 92 |
| Ppin | Paged | 40 | 39 | 1 | 40 | 40 |
| Ppio | Paged | 4137 | 4134 | 3 | 288 | 96 |
| Ppre | Paged | 878 | 794 | 84 | 9120 | 108 |
| Ppsu | Paged | 717 | 541 | 176 | 8808 | 50 |
| Prcr | Nonp | 4 | 3 | 1 | 32 | 32 |
| Prcr | Paged | 23 | 21 | 2 | 184 | 92 |
| Proc | Nonp | 203 | 171 | 32 | 21248 | 664 |
| PsEx | Nonp | 202 | 202 | 0 | 0 | 0 |
| PsIm | Paged | 440 | 345 | 95 | 2280 | 24 |
| PsQb | Nonp | 7 | 0 | 7 | 504 | 72 |
| PsTp | Paged | 3730 | 3466 | 264 | 6336 | 24 |
| Psap | Nonp | 31940 | 31940 | 0 | 0 | 0 |
| Psta | Nonp | 2 | 0 | 2 | 32 | 16 |
| Qpcf | Nonp | 3 | 0 | 3 | 264 | 88 |
| Qpct | Nonp | 5 | 0 | 5 | 560 | 112 |
| Qpfd | Nonp | 2 | 0 | 2 | 48 | 24 |
| Qphf | Nonp | 4 | 0 | 4 | 304 | 76 |
| Qppd | Nonp | 6 | 0 | 6 | 176 | 29 |
| Qpph | Nonp | 12 | 0 | 12 | 416 | 34 |
| Qppi | Nonp | 1 | 0 | 1 | 40 | 40 |
| Qppt | Nonp | 1 | 0 | 1 | 392 | 392 |
| Qprz | Nonp | 10 | 0 | 10 | 240 | 24 |
| RAGE | Paged | 24 | 3 | 21 | 59672 | 2841 |
| RBd1 | Nonp | 1 | 1 | 0 | 0 | 0 |
| RBd2 | Nonp | 1 | 1 | 0 | 0 | 0 |
| RBde | Nonp | 1 | 1 | 0 | 0 | 0 |
| RBev | Nonp | 1 | 0 | 1 | 88 | 88 |
| RBrg | Nonp | 1 | 0 | 1 | 120 | 120 |
| RBtw | Nonp | 1 | 1 | 0 | 0 | 0 |
| REMP | Nonp | 2 | 0 | 2 | 176 | 88 |
| RRle | Paged | 7112 | 6419 | 693 | 33264 | 48 |
| ReEv | Nonp | 199 | 16 | 183 | 4392 | 24 |
| ReSe | Nonp | 83 | 9 | 74 | 2368 | 32 |
| ReTa | Nonp | 208 | 7 | 201 | 6592 | 32 |
| Rqrv | Paged | 5525 | 5525 | 0 | 0 | 0 |
| RsfD | Nonp | 1 | 0 | 1 | 64 | 64 |
| Rx?? | Paged | 618 | 618 | 0 | 0 | 0 |
| RxBm | Nonp | 33 | 32 | 1 | 24 | 24 |
| RxCa | Nonp | 3 | 0 | 3 | 96 | 32 |
| RxCc | Nonp | 5496 | 5466 | 30 | 6320 | 210 |
| RxCd | Nonp | 2887 | 2887 | 0 | 0 | 0 |
| RxCd | Paged | 1131 | 1131 | 0 | 0 | 0 |
| RxCt | Nonp | 3 | 0 | 3 | 304 | 101 |
| RxCt | Paged | 6 | 0 | 6 | 264 | 44 |
| RxFc | Paged | 3983 | 3971 | 12 | 14336 | 1194 |
| RxIr | Nonp | 106144 | 106126 | 18 | 16128 | 896 |
| RxMs | Nonp | 123 | 122 | 1 | 1096 | 1096 |
| RxMs | Paged | 12004 | 12002 | 2 | 72 | 36 |
| RxMx | Paged | 4 | 1 | 3 | 264 | 88 |
| RxNc | Paged | 1641 | 1131 | 510 | 37080 | 72 |
| RxNf | Nonp | 3899 | 3887 | 12 | 2496 | 208 |
| RxNm | Nonp | 32 | 32 | 0 | 0 | 0 |

| | | | | | | |
|---|---|---|---|---|---|---|
| RxNr | Nonp | 1268 | 1253 | 15 | 9016 | 601 |
| RxSC | Nonp | 2543 | 2543 | 0 | 0 | 0 |
| RxSc | Nonp | 14890 | 14866 | 24 | 3512 | 146 |
| RxVn | Nonp | 1279 | 1264 | 15 | 3608 | 240 |
| RxWq | Nonp | 2276 | 2276 | 0 | 0 | 0 |
| Rxbw | Nonp | 125 | 125 | 0 | 0 | 0 |
| SWbr | Paged | 12 | 0 | 12 | 1888 | 157 |
| SWda | Paged | 22 | 0 | 22 | 880 | 40 |
| SWdr | Paged | 13 | 0 | 13 | 2232 | 171 |
| SWfd | Nonp | 1 | 0 | 1 | 216 | 216 |
| SWid | Paged | 12 | 0 | 12 | 1056 | 88 |
| SWki | Paged | 25 | 25 | 0 | 0 | 0 |
| SWre | Nonp | 1 | 1 | 0 | 0 | 0 |
| Sc?? | Nonp | 1 | 1 | 0 | 0 | 0 |
| ScB1 | Nonp | 2 | 2 | 0 | 0 | 0 |
| ScB2 | Nonp | 2 | 2 | 0 | 0 | 0 |
| ScB4 | Nonp | 4 | 0 | 4 | 384 | 96 |
| ScB5 | Paged | 4 | 4 | 0 | 0 | 0 |
| ScC1 | Paged | 3 | 1 | 2 | 240 | 120 |
| ScC7 | Nonp | 20359 | 20359 | 0 | 0 | 0 |
| ScC8 | Paged | 1 | 1 | 0 | 0 | 0 |
| ScC9 | Nonp | 32 | 32 | 0 | 0 | 0 |
| ScCE | Nonp | 89 | 89 | 0 | 0 | 0 |
| ScCI | Nonp | 125 | 124 | 1 | 32 | 32 |
| ScCM | Nonp | 2 | 2 | 0 | 0 | 0 |
| ScCP | Nonp | 8 | 8 | 0 | 0 | 0 |
| ScCS | Nonp | 124 | 124 | 0 | 0 | 0 |
| ScCs | Paged | 1 | 0 | 1 | 48 | 48 |
| ScDC | Nonp | 1 | 1 | 0 | 0 | 0 |
| ScDM | Nonp | 2 | 2 | 0 | 0 | 0 |
| ScDN | Paged | 3 | 0 | 3 | 144 | 48 |
| ScDS | Nonp | 180 | 180 | 0 | 0 | 0 |
| ScDa | Nonp | 45 | 45 | 0 | 0 | 0 |
| ScDg | Paged | 3 | 1 | 2 | 2096 | 1048 |
| ScDs | Nonp | 1 | 0 | 1 | 32 | 32 |
| ScFi | Nonp | 2 | 2 | 0 | 0 | 0 |
| ScFs | Nonp | 2 | 0 | 2 | 208 | 104 |
| ScFt | Nonp | 5 | 0 | 5 | 4280 | 856 |
| ScFt | Paged | 29 | 24 | 5 | 544 | 108 |
| ScIO | Nonp | 9 | 9 | 0 | 0 | 0 |
| ScLA | Nonp | 1 | 1 | 0 | 0 | 0 |
| ScLF | Nonp | 75 | 75 | 0 | 0 | 0 |
| ScLM | Nonp | 2 | 0 | 2 | 192 | 96 |
| ScLf | Nonp | 1 | 0 | 1 | 40 | 40 |
| ScLq | Nonp | 2 | 0 | 2 | 320 | 160 |
| ScPC | Nonp | 2 | 0 | 2 | 3888 | 1944 |
| ScRV | Nonp | 2 | 0 | 2 | 80 | 40 |
| ScRV | Paged | 2 | 0 | 2 | 256 | 128 |
| ScRi | Nonp | 12 | 12 | 0 | 0 | 0 |
| ScRp | Nonp | 5 | 0 | 5 | 448 | 89 |
| ScRp | Paged | 2 | 0 | 2 | 168 | 84 |
| ScRt | Paged | 3 | 2 | 1 | 40 | 40 |
| ScRw | Nonp | 1 | 1 | 0 | 0 | 0 |
| Scs$ | Nonp | 1 | 0 | 1 | 72 | 72 |
| Sdba | Paged | 784 | 632 | 152 | 10760 | 70 |
| Se | Nonp | 1 | 1 | 0 | 0 | 0 |
| Se | Paged | 29199 | 29199 | 0 | 0 | 0 |
| SeAc | Nonp | 64 | 64 | 0 | 0 | 0 |
| SeAc | Paged | 108761 | 108754 | 7 | 688 | 98 |
| SeCL | Nonp | 6 | 0 | 6 | 480 | 80 |
| SeDb | Paged | 6 | 6 | 0 | 0 | 0 |
| SeFS | Paged | 1 | 0 | 1 | 16 | 16 |
| SeHa | Paged | 6 | 6 | 0 | 0 | 0 |
| SeLS | Paged | 1 | 0 | 1 | 72 | 72 |
| SeLs | Paged | 65 | 55 | 10 | 320 | 32 |
| SeLu | Paged | 6692 | 6692 | 0 | 0 | 0 |
| SeLw | Paged | 49 | 49 | 0 | 0 | 0 |
| SeOt | Paged | 28 | 28 | 0 | 0 | 0 |
| SePa | Nonp | 73068 | 73038 | 30 | 3728 | 124 |
| SePa | Paged | 40 | 39 | 1 | 16 | 16 |
| SePr | Paged | 407 | 404 | 3 | 104 | 34 |
| SeSa | Paged | 139 | 139 | 0 | 0 | 0 |

| | | | | | | |
|---|---|---|---|---|---|---|
| SeSc | Paged | 57961 | 57961 | 0 | 0 | 0 |
| SeSd | Nonp | 98 | 89 | 9 | 1008 | 112 |
| SeSd | Paged | 107915 | 107900 | 15 | 1848 | 123 |
| SeSi | Paged | 416 | 388 | 28 | 664 | 23 |
| SeTI | Paged | 1 | 1 | 0 | 0 | 0 |
| SeTa | Paged | 139 | 139 | 0 | 0 | 0 |
| SeTd | Nonp | 29701 | 29591 | 110 | 7040 | 64 |
| SeTd | Paged | 29701 | 29591 | 110 | 8816 | 80 |
| SeUs | Paged | 53375 | 53375 | 0 | 0 | 0 |
| Sect | Paged | 65816 | 65527 | 289 | 26136 | 90 |
| Sema | Nonp | 1719 | 1216 | 503 | 28456 | 56 |
| Senm | Nonp | 10 | 10 | 0 | 0 | 0 |
| Senm | Paged | 4 | 4 | 0 | 0 | 0 |
| SmAd | Nonp | 1540 | 1532 | 8 | 888 | 111 |
| SmAd | Paged | 1 | 0 | 1 | 208 | 208 |
| SmDO | Nonp | 2447 | 2443 | 4 | 360 | 90 |
| SmDc | Paged | 344 | 344 | 0 | 0 | 0 |
| SmEc | Nonp | 1 | 0 | 1 | 64 | 64 |
| SmFc | Nonp | 1426 | 1415 | 11 | 1056 | 96 |
| SmKs | Paged | 2215 | 2215 | 0 | 0 | 0 |
| SmMa | Nonp | 746 | 738 | 8 | 2176 | 272 |
| SmMm | Nonp | 3451 | 3399 | 52 | 31760 | 610 |
| SmMs | Nonp | 916 | 916 | 0 | 0 | 0 |
| SmMs | Paged | 8 | 3 | 5 | 696 | 139 |
| SmMt | Nonp | 1498 | 1482 | 16 | 832 | 52 |
| SmNr | Paged | 2500 | 2491 | 9 | 656 | 72 |
| SmQP | Paged | 4672 | 4672 | 0 | 0 | 0 |
| SmRw | Nonp | 352 | 352 | 0 | 0 | 0 |
| SmRx | Nonp | 2809 | 2809 | 0 | 0 | 0 |
| SmSe | Nonp | 129 | 114 | 15 | 528 | 35 |
| SmSe | Paged | 4 | 4 | 0 | 0 | 0 |
| SmSr | Nonp | 1021 | 1009 | 12 | 448 | 37 |
| SmSr | Paged | 319 | 319 | 0 | 0 | 0 |
| SmTp | Nonp | 774 | 768 | 6 | 408 | 68 |
| SmTp | Paged | 6 | 6 | 0 | 0 | 0 |
| SmTr | Paged | 91256 | 91256 | 0 | 0 | 0 |
| SmVc | Nonp | 2996 | 2988 | 8 | 1024 | 128 |
| SmVr | Nonp | 3703 | 3694 | 9 | 2160 | 240 |
| Strg | Nonp | 10 | 10 | 0 | 0 | 0 |
| Strg | Paged | 11607 | 11570 | 37 | 2696 | 72 |
| Symb | Paged | 160 | 19 | 141 | 11280 | 80 |
| Symt | Paged | 156 | 15 | 141 | 7784 | 55 |
| TAPI | Nonp | 2 | 0 | 2 | 176 | 88 |
| TAPI | Paged | 1 | 1 | 0 | 0 | 0 |
| TC2r | Nonp | 8 | 6 | 2 | 48 | 24 |
| TCI1 | Nonp | 43466 | 43466 | 0 | 0 | 0 |
| TCI4 | Nonp | 2 | 0 | 2 | 816 | 408 |
| TCI5 | Nonp | 2 | 0 | 2 | 528 | 264 |
| TCI8 | Nonp | 2 | 2 | 0 | 0 | 0 |
| TCI9 | Nonp | 2 | 2 | 0 | 0 | 0 |
| TCIV | Nonp | 4 | 0 | 4 | 544 | 136 |
| TCIY | Nonp | 677 | 677 | 0 | 0 | 0 |
| TCIZ | Nonp | 39 | 0 | 39 | 2112 | 54 |
| TCIa | Nonp | 2 | 0 | 2 | 208 | 104 |
| TCIi | Nonp | 699 | 699 | 0 | 0 | 0 |
| TCIj | Nonp | 699 | 699 | 0 | 0 | 0 |
| TCIk | Nonp | 3 | 3 | 0 | 0 | 0 |
| TCIl | Nonp | 3 | 3 | 0 | 0 | 0 |
| TCIm | Nonp | 1 | 0 | 1 | 64 | 64 |
| TCIo | Nonp | 1 | 1 | 0 | 0 | 0 |
| TCIp | Nonp | 4 | 4 | 0 | 0 | 0 |
| TCIq | Nonp | 1 | 1 | 0 | 0 | 0 |
| TCIs | Nonp | 705 | 705 | 0 | 0 | 0 |
| TCIt | Nonp | 3 | 0 | 3 | 4552 | 1517 |
| TCIu | Nonp | 3 | 3 | 0 | 0 | 0 |
| TCIv | Nonp | 3 | 0 | 3 | 600 | 200 |
| TCIw | Nonp | 2 | 0 | 2 | 1360 | 680 |
| TCIy | Nonp | 4 | 0 | 4 | 192 | 48 |
| TCPA | Nonp | 1663 | 1616 | 47 | 16544 | 352 |
| TCPB | Nonp | 9 | 5 | 4 | 4624 | 1156 |
| TCPC | Nonp | 539 | 225 | 314 | 27080 | 86 |
| TCPT | Nonp | 100 | 85 | 15 | 5800 | 386 |

| | | | | | | |
|---|---|---|---|---|---|---|
| TCPY | Nonp | 1 | 0 | 1 | 88 | 88 |
| TCPa | Nonp | 217 | 212 | 5 | 160 | 32 |
| TCPc | Nonp | 4693 | 4307 | 386 | 18528 | 48 |
| TCPi | Nonp | 27 | 26 | 1 | 40 | 40 |
| TCPr | Nonp | 327 | 316 | 11 | 728 | 66 |
| TCPt | Nonp | 162073 | 162038 | 35 | 84544 | 2415 |
| TCPt | Paged | 1 | 1 | 0 | 0 | 0 |
| TCPu | Nonp | 14 | 14 | 0 | 0 | 0 |
| TChc | Nonp | 1 | 0 | 1 | 4096 | 4096 |
| TChe | Nonp | 1 | 0 | 1 | 4096 | 4096 |
| TChi | Nonp | 1 | 0 | 1 | 4096 | 4096 |
| TCht | Nonp | 1 | 0 | 1 | 4096 | 4096 |
| TChu | Nonp | 1 | 0 | 1 | 4096 | 4096 |
| TCi4 | Nonp | 39 | 39 | 0 | 0 | 0 |
| TCiA | Nonp | 13525 | 13515 | 10 | 880 | 88 |
| TCiJ | Nonp | 1 | 0 | 1 | 40 | 40 |
| TCiK | Nonp | 1 | 0 | 1 | 40 | 40 |
| TCiN | Nonp | 21941 | 21941 | 0 | 0 | 0 |
| TCiQ | Nonp | 271 | 267 | 4 | 288 | 72 |
| TCiR | Nonp | 2 | 2 | 0 | 0 | 0 |
| TCiS | Nonp | 2 | 0 | 2 | 48 | 24 |
| TCiT | Nonp | 1 | 0 | 1 | 32 | 32 |
| TCiW | Nonp | 3 | 3 | 0 | 0 | 0 |
| TCiX | Nonp | 5 | 5 | 0 | 0 | 0 |
| TCiY | Nonp | 3 | 3 | 0 | 0 | 0 |
| TCiZ | Nonp | 1 | 1 | 0 | 0 | 0 |
| TCid | Nonp | 34 | 34 | 0 | 0 | 0 |
| TCif | Nonp | 71 | 71 | 0 | 0 | 0 |
| TCij | Nonp | 1 | 0 | 1 | 48 | 48 |
| TD | Nonp | 45 | 22 | 23 | 4376 | 190 |
| TDIa | Nonp | 34 | 27 | 7 | 616 | 88 |
| TDIc | Nonp | 11 | 0 | 11 | 880 | 80 |
| TDId | Nonp | 11 | 0 | 11 | 792 | 72 |
| TDIe | Nonp | 5 | 0 | 5 | 480 | 96 |
| TDIf | Nonp | 7 | 0 | 7 | 504 | 72 |
| TDIg | Nonp | 6 | 0 | 6 | 216 | 36 |
| TDIj | Nonp | 4 | 4 | 0 | 0 | 0 |
| TDIk | Nonp | 4 | 0 | 4 | 2048 | 512 |
| TDIu | Nonp | 5 | 0 | 5 | 616 | 123 |
| TDIv | Nonp | 5 | 0 | 5 | 80 | 16 |
| TOBJ | Nonp | 15 | 0 | 15 | 1928 | 128 |
| TSdd | Paged | 1348 | 1320 | 28 | 357112 | 12754 |
| TSlc | Paged | 1 | 0 | 1 | 32 | 32 |
| TSmc | Paged | 11 | 0 | 11 | 6184 | 562 |
| TSwd | Nonp | 3 | 0 | 3 | 200 | 66 |
| TSwd | Paged | 13 | 3 | 10 | 417552 | 41755 |
| TTFC | Paged | 9 | 9 | 0 | 0 | 0 |
| Thre | Nonp | 18362 | 17939 | 423 | 270720 | 640 |
| Time | Nonp | 79 | 33 | 46 | 8512 | 185 |
| Toke | Paged | 29701 | 29591 | 110 | 74888 | 680 |
| Ttfd | Paged | 19893 | 19613 | 280 | 588432 | 2101 |
| Tun4 | Paged | 2 | 2 | 0 | 0 | 0 |
| TunP | Paged | 248 | 247 | 1 | 240 | 240 |
| UHUB | Nonp | 33 | 29 | 4 | 408 | 102 |
| UHUB | Paged | 4 | 2 | 2 | 152 | 76 |
| USBD | Nonp | 1 | 1 | 0 | 0 | 0 |
| USBD | Paged | 1 | 1 | 0 | 0 | 0 |
| Umrx | Nonp | 1 | 0 | 1 | 200 | 200 |
| Urdr | Paged | 5 | 2 | 3 | 120 | 40 |
| UsDI | Paged | 26 | 20 | 6 | 1776 | 296 |
| Usac | Paged | 17 | 4 | 13 | 1568 | 120 |
| Usai | Paged | 10 | 9 | 1 | 24 | 24 |
| Usbg | Nonp | 4 | 0 | 4 | 160 | 40 |
| Uscb | Paged | 14 | 11 | 3 | 992 | 330 |
| Uscc | Paged | 11674 | 11674 | 0 | 0 | 0 |
| Usci | Paged | 5 | 0 | 5 | 760 | 152 |
| Uscl | Paged | 40 | 0 | 40 | 2640 | 66 |
| Usct | Paged | 14 | 9 | 5 | 240 | 48 |
| Uscu | Paged | 3431 | 3329 | 102 | 9792 | 96 |
| Usd9 | Paged | 1 | 0 | 1 | 48 | 48 |
| Usdc | Paged | 74 | 59 | 15 | 720 | 48 |
| User | Nonp | 6 | 0 | 6 | 384 | 64 |

| | | | | | | |
|---|---|---|---|---|---|---|
| Usfd | Nonp | 1 | 0 | 1 | 8192 | 8192 |
| Ushk | Paged | 27 | 0 | 27 | 864 | 32 |
| Usih | Paged | 72 | 54 | 18 | 576 | 32 |
| Uskb | Paged | 27 | 25 | 2 | 160 | 80 |
| Uske | Nonp | 1 | 1 | 0 | 0 | 0 |
| Uskf | Paged | 2 | 0 | 2 | 208 | 104 |
| Usks | Paged | 14 | 14 | 0 | 0 | 0 |
| Uskt | Paged | 2 | 0 | 2 | 5360 | 2680 |
| Usla | Nonp | 6 | 0 | 6 | 672 | 112 |
| Usmr | Paged | 1 | 1 | 0 | 0 | 0 |
| Usms | Paged | 4 | 4 | 0 | 0 | 0 |
| Usny | Paged | 3 | 3 | 0 | 0 | 0 |
| Uspi | Paged | 617 | 529 | 88 | 12360 | 140 |
| Uspm | Paged | 38 | 38 | 0 | 0 | 0 |
| Uspn | Paged | 15 | 15 | 0 | 0 | 0 |
| Uspo | Nonp | 4 | 2 | 2 | 80 | 40 |
| Uspp | Nonp | 2 | 0 | 2 | 80 | 40 |
| Uspp | Paged | 25 | 20 | 5 | 864 | 172 |
| Uspr | Paged | 1617 | 1617 | 0 | 0 | 0 |
| Usqm | Paged | 358 | 350 | 8 | 448 | 56 |
| Usqu | Paged | 215 | 132 | 83 | 21912 | 264 |
| Ussa | Paged | 8 | 8 | 0 | 0 | 0 |
| Ussb | Paged | 78 | 78 | 0 | 0 | 0 |
| Ussc | Paged | 11236 | 11236 | 0 | 0 | 0 |
| Usse | Paged | 88 | 87 | 1 | 32 | 32 |
| Ussm | Paged | 37 | 30 | 7 | 504 | 72 |
| Usss | Paged | 13 | 13 | 0 | 0 | 0 |
| Usst | Paged | 1 | 1 | 0 | 0 | 0 |
| Ussw | Paged | 1255 | 1253 | 2 | 1072 | 536 |
| Ussy | Nonp | 62 | 22 | 40 | 1456 | 36 |
| Ussy | Paged | 18 | 11 | 7 | 544 | 77 |
| Usti | Paged | 524 | 447 | 77 | 27104 | 352 |
| Ustm | Paged | 969 | 955 | 14 | 784 | 56 |
| Usto | Paged | 4 | 4 | 0 | 0 | 0 |
| Usts | Paged | 335 | 271 | 64 | 3960 | 61 |
| Ustx | Paged | 5 | 0 | 5 | 200 | 40 |
| Usty | Paged | 294 | 294 | 0 | 0 | 0 |
| Usvl | Paged | 16686 | 16685 | 1 | 32 | 32 |
| Uswd | Paged | 3 | 0 | 3 | 528 | 176 |
| Uswe | Paged | 1 | 1 | 0 | 0 | 0 |
| Uswl | Paged | 17152 | 17150 | 2 | 408 | 204 |
| Vad | Nonp | 91751 | 90264 | 1487 | 71376 | 48 |
| VadF | Nonp | 1 | 0 | 1 | 32 | 32 |
| VadS | Nonp | 504709 | 504011 | 698 | 22336 | 32 |
| Vadl | Nonp | 15296 | 14813 | 483 | 30912 | 64 |
| Vga | Paged | 2 | 1 | 1 | 3272 | 3272 |
| VoSb | Nonp | 2 | 0 | 2 | 32768 | 16384 |
| VoSc | Nonp | 5 | 0 | 5 | 456 | 91 |
| VoSr | Paged | 2 | 2 | 0 | 0 | 0 |
| VoSs | Nonp | 4 | 4 | 0 | 0 | 0 |
| VoSs | Paged | 3 | 2 | 1 | 128 | 128 |
| VoSt | Nonp | 1 | 0 | 1 | 104 | 104 |
| VolS | Paged | 8 | 8 | 0 | 0 | 0 |
| Vpb | Nonp | 8 | 0 | 8 | 768 | 96 |
| Vprt | Nonp | 13 | 4 | 9 | 360 | 40 |
| Vprt | Paged | 1 | 0 | 1 | 120 | 120 |
| VraP | Nonp | 1 | 0 | 1 | 40 | 40 |
| VraP | Paged | 3 | 1 | 2 | 112 | 56 |
| Vtfd | Paged | 8 | 0 | 8 | 2176 | 272 |
| Wait | Nonp | 3818334 | 3818302 | 32 | 8512 | 266 |
| WanI | Nonp | 2 | 0 | 2 | 224 | 112 |
| WanJ | Nonp | 4 | 0 | 4 | 1344 | 336 |
| WanM | Nonp | 6 | 0 | 6 | 512 | 85 |
| WanN | Nonp | 1 | 0 | 1 | 680 | 680 |
| WanO | Nonp | 1 | 0 | 1 | 2064 | 2064 |
| WanS | Nonp | 2 | 0 | 2 | 144 | 72 |
| Wind | Nonp | 7 | 0 | 7 | 1008 | 144 |
| WmDS | Paged | 58 | 57 | 1 | 8192 | 8192 |
| WmGE | Paged | 1 | 0 | 1 | 49152 | 49152 |
| WmIS | Paged | 1 | 0 | 1 | 118784 | 118784 |
| WmMR | Paged | 1 | 0 | 1 | 544 | 544 |
| WmiD | Nonp | 48 | 48 | 0 | 0 | 0 |

```
 WmiG Nonp        563       532        31    6944        224
 WmiI Paged         1         1         0       0          0
 WmiR Nonp        197       128        69    4416         64
 Wmip Nonp        290       155       135    3240         24
 Wmip Paged      3071      2875       196   21104        107
 Wmit Nonp          1         0         1    4096       4096
 Wmit Paged         7         2         5   33096       6619
 Wmiw Nonp         23        23         0       0          0
 Wrpa Nonp          2         0         2     720        360
 Wrpc Nonp          1         0         1      72         72
 Wrpi Nonp          1         0         1     120        120
 Wrpr Nonp          1         1         0       0          0
 Wrps Nonp          6         4         2     128         64
 Xsmb Paged       129       129         0       0          0
 aEoP Nonp          1         0         1     672        672
 cbuf Paged         1         0         1      24         24
 cevt Nonp          3         0         3      96         32
 e100 Nonp         11         1        10   79536       7953
 iidg Paged         2         2         0       0          0
   rx Nonp          2         0         2     832        416
 sidg Nonp          4         0         4      96         24
 thdd Paged         2         0         2   40960      20480
 usbp Nonp         23        13        10   38960       3896
 usbp Paged        19        18         1      48         48
 vPrt Nonp          3         2         1      24         24
 vPrt Paged       171       152        19    3120        164
 xSMB Paged      5170      5170         0       0          0
```

**Example 3: Analyzing memsnap logs**

The /a parameter generate an analysis of memsnap logs created with the /m parameter.

```
memsnap /a memsnap.log
```

This command sends the memory data to the Memsnap.txt file in the C:\Data directory.

**Sample memsnap /t output**

```
Name            Probability Object          Change     Start       End Percent Rate/hour
net1.exe        Definite    PagedPool          260     16484     16744       1       318
net1.exe        Definite    NonPagedPool        40      1640      1680       2        49
net1.exe        Definite    Commit          217088    438272    655360      49    266112
net1.exe        Definite    Handles             11        22        33      50        13
memsnap.exe     Definite    PagedPool         1648     14272     15920      11      2020
memsnap.exe     Definite    NonPagedPool       176      1200      1376      14       215
memsnap.exe     Definite    Commit           40960    282624    323584      14     50209
memsnap.exe     Definite    Handles              9        12        21      75        11
net.exe         Definite    Commit           12288    323584    335872       3     15062
net.exe         Definite    Handles              1        21        22       4         1
net.exe         Definite    PagedPool        10792      5128     15920     210     13229
net.exe         Definite    NonPagedPool      1016       360      1376     282      1245
net.exe         Definite    Commit          217088    118784    335872     182    266112
net.exe         Definite    Handles             21         1        22    2100        25
```

**See Also**

**Concepts**

Memsnap Overview
Memsnap Syntax
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Gflags.exe
Devcon Overview
Pviewer Overview
Exctrlst Overview

Devcon Overview
Apmstat Overview

# Memsnap Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MemSnap Syntax**

MemSnap uses the following syntax:

- **memsnap** [**/m**] [**/p**] [**/a**] [**/ah**] *logfile* [**/?**]

**Parameters**

***/m logfile***

This parameter displays a snapshot of process information in the form of a log file. The default file name for the log file is memsnap.log in the current directory. If the file does not exist, it is created; if the file exists, the information is appended to it.

***/p logfile***

This parameter displays a snapshot kernel pool memory information. If the file does not exist, it is created; if the file exists, the information is appended to it.

***/a logfile***

This parameter analyzes a *logfile* for leaks and prints the output on the screen. The logfile to analyze is required with this parameter.

**/ah**

Same as '**/a**' but generates HTML tables.

**/?**

Displays command-line help.

**See Also**

**Concepts**

Memsnap Overview
Memsnap Examples
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Gflags.exe
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Poolmon Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Poolmon.exe**
This command-line tool displays data that the operating system collects about memory allocations from the system paged and non-paged kernel pools, and the memory pools used for Terminal Services sessions. The data is grouped by pool allocation tag. Poolmon is most often used to detect memory leaks.

Poolmon displays the following data for allocations grouped by tag.

- The number of allocations and frees (and allocations minus frees).
- The change in the number of allocations and frees between updates.
- The total size of allocations by tag, in bytes used, and the average allocation size.
- The change in bytes used between updates.
- The drivers that assign the tag value.

Poolmon also displays general memory information, including total and available memory, page faults, kernel physical memory, committed memory and the commit limit, peak memory, and the size of the paged and non-paged pools.

Using Poolmon, you can also:

- Sort and reconfigure the Poolmon display while it is running.
- Save configured data to a file.
- Generate a file of the tags used by drivers on the local system (32-bit system only).

### Note

- To use Poolmon on Windows XP and earlier systems, you must enable *pool tagging*, the feature that collects and processes data about memory allocations. On Windows Server 2003, pool tagging is permanently enabled. For more information, see "About Pool Tagging" in Poolmon Remarks.

**Corresponding UI**
This tool does not have a corresponding user interface.

**Concepts**
Memory leaks occur when programs allocate memory but do not free it when the memory is no longer being used.

To track kernel pool memory, the system attaches a label called a *tag* to each allocation and lets the allocating driver set the value of the tag. When pool tagging is enabled, the system collects data about the allocations, sorted by the tag value. You can use Poolmon to display the data that the system collected for all allocations, or for allocations with a particular tag.

In addition to the system memory pools, there is a paged pool for each Terminal Services session on the system. When a system is configured as a Terminal Server, (formerly known as Terminal Services in Application Server mode), the kernel-mode portion of the Win32 subsystem allocates memory from the session pool for data structures used in the session. Poolmon also displays those allocations from the Terminal Services session pools, individually and collectively, sorted by tag value.

For more information about memory management and memory leaks, search the Microsoft Developer Network.

**System Requirements**

- Windows XP Professional or Windows Server 2003

## Notes

- The system allocates memory from Terminal Services session pools only when the system is configured as a Terminal Server. Terminal Server is available on all Windows Server 2003 operating systems, except for Windows Server 2003, Web Edition.

- Poolmon cannot generate a Localtag.txt file on 64-bit versions of Windows Server 2003. As a result, the **/c** parameter and its functionality are available only on 32-bit versions of Windows.

### Files required

- Poolmon.exe

- msdis130.dll

- msvcp70.dll

- msvcr70.dll

- pooltag.txt

**See Also**
**Concepts**
Poolmon Remarks
Poolmon Syntax
Poolmon Examples
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Poolmon Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Poolmon Remarks**
**Enabling Pool Tagging**
Before running Poolmon on Windows XP and earlier systems, you must use Gflags.exe to enable pool tagging.

### Important

- Pool tagging is permanently enabled on Windows Server 2003.

**To enable pool tagging by using a dialog box:**

1. At the command line, or in the Run dialog box, type **Gflags**.
2. Check the **Enable Pool Tagging** check box.
3. Restart the computer.

**To enable pool tagging by using the command line:**

1. Type the following at the command line: `gflags /r +ptg`
2. Restart the computer.

**About Pool Tagging**
The pool tagging feature collects and calculates statistics about pool memory sorted by the tag value of the allocation.

To set a tag value, use **ExAllocatePoolWithTag** or **ExAllocatePoolWithQuotaTag** to allocate pool memory. If you do not assign a tag value (**ExAllocatePool**, **ExAllocatePoolWithQuota**), the system still creates a tag, but it assigns the default tag value, "None". As a result, you cannot distinguish the statistics for your allocation from that of other pool memory allocations.

For more information on pool tagging, see Gflags Flag Reference. For more information about pool tags, see the Windows Driver Development Kit.

**About Poolmon**
### Command-line Syntax and Running Syntax

You can configure the Poolmon display at the command line when you start Poolmon or while Poolmon is running. The command line parameters are described in the "Command-line Syntax" section of Poolmon Syntax. The keys that you press to configure the Poolmon display while it is running are described in the "Running Syntax" section of Poolmon Syntax.

### Accumulated Values

The data that Poolmon displays is collected and calculated by the system whenever pool tagging is enabled. The values for allocations, frees, and bytes used accumulate from the time that the system starts, and increase monotonically until the system is restarted. If a driver or component is started independently of the system, the values are accumulated from the last time that the driver or component started and reset only when the driver or system are restarted.

### Update Rate

Poolmon updates its display every five seconds. You cannot change the update rate.

### Saving Poolmon Output

You can add the **/n** parameter to a command to save a snapshot of the Poolmon output to a file. Because the data saved to the file is static, the columns that show the change in values between updates are not included in the file.

### Terminal Server Session Support

Poolmon can display allocations from the Terminal Services session pools, as well as the system pools. When the computer is configured as a Terminal Server, the kernel-mode portions of the Win32 subsystem allocate memory from the session pools. Otherwise, the system allocates pool memory for Terminal Services from the system pool.

Terminal Server, previously known as Terminal Services in Application Server mode, is available on all Windows Server 2003 operating systems, except for Windows Server 2003, Web Edition.

Poolmon can display allocations from a particular session pool, or from all session pools collectively. To display only the pool for a particular session, you specify the session ID. To find a Terminal Server session ID, use Task Manager. After Task Manager opens, click the Users tab to see the session IDs.

To display session IDs in Task Manager:

- From the **View** menu, click **Select Column**, and then select the **Session ID** check box.

### Viewing Driver Names

Poolmon has two features that display the names of drivers that assign each pool tag.

- **Local drivers**. The **/c** parameter displays the names of drivers on the local computer that assign the pool tag. The data for this feature comes from Localtag.txt, a file that Poolmon generates automatically when you use the **/c** parameter. Poolmon generates a new Localtag.txt file when it cannot find a Localtag.txt file in the current directory and you do not specify an alternate file.

- **Windows components and other drivers**. The **/g** parameter displays the names of Windows components and commonly used drivers that assign the pool tag. The data for this feature comes from Pooltag.txt, a file provided by Microsoft that is installed in the Support Tools folder with Poolmon.exe.

To display the driver name, Poolmon adds a column called "Mapped_Driver" to the right edge of the display and populates the column with data from Localtag.txt (for **/c**) or Pooltag.txt (for **/g**). The "Mapped_Driver" column appears only when you use the **/c** or **/g** parameters. You can use both the **/c** and **/g** parameters in a single command. In response, Poolmon displays the data from both files in the "Mapped_Driver" column.

You can also use **/c** and **/g** to display data from alternate files in the "Mapped_Driver" column by specifying a file name and location with either parameter (**poolmon /c***filename* | **poolmon /c***filename*). In that case, the **/c** and **/g** parameters behave identically and can be used interchangeably.

### Note

- Poolmon cannot generate a Localtag.txt file on 64-bit versions of Windows Server 2003.

### Updating Pooltag.txt

Poolmon can display the name of the Windows component or commonly used driver that assigns each pool tag. This feature uses data from Pooltag.txt, a file installed with Poolmon and with Windows debuggers. Occasionally, Microsoft updates this file. To check for updates, go to the Microsoft Product Support Services Web page, http://support.microsoft.com and search for "Pooltag.txt."

#### About the Poolmon Display

Poolmon displays columns of data about pool memory allocations in a command window. Use the arrow keys, page up and page down keys to scroll through the data. Poolmon updates the display every five seconds. You cannot change the update rate.

### Note

- To see the entire Poolmon display, the command window screen must be at least 80 characters wide (width=80) and at least 53 rows high (height=53); and the command window screen buffer must be at least 500 characters wide (width=500) and at least 2000 rows high (height=2000). Otherwise, the display might be truncated.

The following table describes the columns in the Poolmon display

| Column name | Description |
|---|---|
| Tag | The 4-byte tag assigned to the pool allocation. |
| Type | Whether the memory allocations are in paged or non-paged bytes. |
| Allocs | The number of allocations. |
| ( ) | The change in the number of allocations since the last update. |
| Frees | The number of frees. |
| ( ) | The change in the number of allocations since the last update. |
| Diff | The number of allocations minus the number of frees. |
| Bytes | The size of the allocations, in bytes used. |
| ( ) | The change in the allocation size since the last update. |
| Per Alloc | The value of Bytes divided by the value of Allocs. |
| Mapped_Driver | The local drivers (**/c**) and other commonly used drivers and system components (**/g**) that assign the pool tag value. This column appears only when you use the **/c** or **/g** parameters. |

The following sample Poolmon output is sorted by number of allocations. (To sort your display this way, start Poolmon with the /a parameter.)

```
Memory:  260620K Avail:   96364K  PageFlts:     0   InRam Krnl: 1916K P:17856K
Commit: 203500K Limit: 640916K Peak: 260632K            Pool N: 8332K P:27220K
System pool information
Tag  Type     Allocs            Frees            Diff   Bytes       Per Alloc

Wait Nonp    3971107 (    0)   3971077 (    0)      30    8456 (     0)    281
ObSt Nonp    2791258 (    0)   2791258 (    0)       0       0 (     0)      0
Gxlt Paged   1161638 (    0)   1161630 (    0)       8     864 (     0)    108
Ustm Paged   1088342 (    0)   1088298 (    0)      44    2464 (     0)     56
Io   Nonp    1021112 (    1)   1020985 (    1)     127   91912 (     0)    723
ObSq Paged    967615 (    0)    967615 (    0)       0       0 (     0)      0
Key  Paged    954821 (    0)    953979 (    0)     842   87528 (     0)    103
SePa Nonp     680348 (    0)    680321 (    0)      27    3656 (     0)    135
```

**See Also**

**Concepts**

Poolmon Overview
Poolmon Syntax
Poolmon Examples
Alphabetical List of Tools
Windiff Overview
Remote Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Poolmon Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Poolmon Examples**
Before running any of these examples on Windows XP and earlier systems, use Gflags.exe to enable pool tagging and then restart the system to make the change effective.

**Example 1: Display and Sort Poolmon Output**
This example describes various ways to configure the Poolmon display. By default, Poolmon displays all kernel memory allocations in alphabetical order by tag value. You can modify the sort order of the display at the command line or while Poolmon is running.

The following command starts Poolmon.

**poolmon**

The following command starts poolmon and sorts the display by number of frees.

**poolmon /f**

While poolmon is running, you can use the parameters in the running syntax to change the display. For example, to sort the display by number of bytes used, press **b** or **u**. To sort by bytes per allocation, press **m**.

The following command starts Poolmon and displays only allocations from the non-paged pool.

**poolmon /p**

While Poolmon is running, you can use press **p** to toggle between allocations from the paged pool, the non-paged pool, or both.

To display data for allocations with a particular tag, use the **/i** parameter. The following command displays allocations with the AfdB tag (the tag used by Afd.sys for data buffers).

**poolmon /iAfdB**

To exclude allocations with a particular tag, use the **/x** parameter. The following command displays all allocations that do not have the AfdB tag.

**poolmon /xAfdB**

You can use the wildcard character (*) and question mark (?) to specify groups of tags. The following command lists all allocations with tags beginning with Afd, the tag used by Afd.sys.

**poolmon /iAfd***

You can also sort the Poolmon display by the change in a value between updates. The **/(** parameter places Poolmon in "sort-by-change" mode.

The following command displays allocations with tags beginning with Afd, and sorts by the change (**/)**) in allocations (**/a**).

**poolmon /iAfd* /( /a**

The **/(** parameter is a toggle. When Poolmon is in sort-by-change mode, it continues to interpret all sort commands as commands to sort by the change in the value, until you press **(** again. If, for example, after using the previous command to start Poolmon, you then pressed **f**, Poolmon would sort the display by the *change in* the number of frees. To revert to sorting by the primary value while Poolmon is running, press **(** again.

**Example 2: Display Driver Names**
You can use the /g parameter when starting Poolmon to display the names of Windows components and commonly used drivers that assign each tag. If you find a problem in allocations with a particular tag, this feature helps you identify the offending component or driver. The components and drivers are listed in the Mapped_Driver column, on the right-most side of the display. The data for the Mapped_Driver column comes from Pooltag.txt, a file installed with Poolmon.

The following command displays memory allocated with tags that begin with NtF. (It uses the question mark character (?) as a

wildcard.) Using the **/g** parameter adds the Mapped_Driver column.

**poolmon /iNtF? /g**

The resulting display lists allocations with tags beginning in NtF. The right-most column in the display, Mapped_Driver, shows that the memory was allocated by Ntfs.sys, the driver for the NTFS file system. In this case, the display is even more specific, because Pooltag.txt includes the source files for Ntfs allocations.

```
 Memory:  260620K Avail:   65152K  PageFlts:    85   InRam Krnl: 2116K P:19560K
 Commit: 237688K Limit: 640916K Peak: 260632K            Pool N: 8500K P:33024K
 System pool information
 Tag  Type     Allocs            Frees            Diff   Bytes      Per Alloc
Mapped_Driver

 NtFA Nonp       9112 (   0)      9112 (   0)     0       0 (     0)      0 [ntfs.sys
-  AttrSup.c]
 NtFB Paged      3996 (   0)      3986 (   0)    10  252088 (     0)  25208 [ntfs.sys
-  BitmpSup.c]
 NtFC Paged   1579279 (   0)   1579269 (   0)    10     640 (     0)     64 [ntfs.sys
-  Create.c]
 NtFD Nonp         13 (   0)        13 (   0)     0       0 (     0)      0 [ntfs.sys
-  DevioSup.c]
 NtFF Paged      1128 (   0)      1128 (   0)     0       0 (     0)      0 [ntfs.sys
-  FileInfo.c]
 NtFI Nonp        152 (   0)       152 (   0)     0       0 (     0)      0 [ntfs.sys
-  IndexSup.c]
 NtFL Nonp      68398 (   0)     68390 (   0)     8   27280 (     0)   3410 [ntfs.sys
-  LogSup.c]
 NtFS Paged      2915 (   0)      2614 (   0)   301   80192 (     0)    266 [ntfs.sys
-  SecurSup.c]
 NtFa Paged       838 (   0)       829 (   0)     9     288 (     0)     32 [ntfs.sys
-  AllocSup.c]
 NtFd Paged    137696 (   0)    137688 (   0)     8     720 (     0)     90 [ntfs.sys
-  DirCtrl.c]
 NtFf Nonp          2 (   0)         1 (   0)     1      40 (     0)     40 [ntfs.sys
-  FsCtrl.c]
 NtFs Nonp      48825 (   0)     47226 (   0)  1599   64536 (     0)     40 [ntfs.sys
-  StrucSup.c]
 NtFv Paged       551 (   0)       551 (   0)     0       0 (     0)      0 [ntfs.sys
-  ViewSup.c]
```

Pooltag.txt is extensive, but it is not a complete list of all tags that can be used on the system. When a tag used in the system is not included in Pooltag.txt, Poolmon displays "Unknown driver" in the Mapped_Driver column for the tag. When this occurs, you can use the **/c** parameter to search the drivers on the local system and determine whether they use assign the tag.

The following examples demonstrate this method.

The following command uses the **/i** tag to list allocations with tags that end in MEM. The **/g** parameter adds the driver name to the display from the Pooltag.txt file.

**poolmon /i?MEM /g**

The resulting display lists the allocations with tags ending in MEM. However, because the MEM tags are not included in Pooltag.txt, "Unknown Driver" appears in the Mapped_Driver column in place of the driver name.

```
 Tag  Type        Allocs          Frees      Diff   Bytes      Per Alloc
Mapped_Driver

 1MEM Nonp       1 (   0)         0 (   0)     1    3344 (     0)   3344   Unknown
Driver
 2MEM Nonp       1 (   0)         0 (   0)     1    3944 (     0)   3944   Unknown
Driver
 3MEM Nonp       3 (   0)         0 (   0)     3     248 (     0)     82   Unknown
Driver
```

In this case, you can use the **/c** parameter to compile a list of local drivers and the tags they assign and then display the names of local drivers in the Mapped_Driver column.

The following command starts Poolmon. It uses the **/i** parameter to list allocations with tags ending in MEM, and the **/c** parameter to display the local drivers that assign the tags.

**poolmon /i?MEM /c**

If you do not specify a local tag file and Poolmon cannot find a Localtag.txt file on a 32-bit system, it creates one, as shown in the following screen messages. (Poolmon cannot generate a local tag file on 64-bit systems.)

```
d:\tools\Poolmon>poolmon /?MEM /c
Poolmon: No localtag.txt in current directory
Poolmon: Creating localtag.txt in current directory......
```

The resulting display, which uses the content from the newly created localtag.txt file, shows the local driver names in the Mapped_Driver column.

```
 Memory:  260620K Avail:   57840K  PageFlts:   162   InRam Krnl: 2116K P:19448K
 Commit: 244580K Limit: 640916K Peak: 265416K            Pool N: 8496K P:32904K
 System pool information
 Tag  Type     Allocs            Frees            Diff   Bytes      Per Alloc
Mapped_Driver

 1MEM Nonp          1 (   0)         0 (   0)        1    3344 (     0)   3344
[el90xbc5]
 2MEM Nonp          1 (   0)         0 (   0)        1    3944 (     0)   3944
[el90xbc5]
 3MEM Nonp          3 (   0)         0 (   0)        3     248 (     0)     82
[el90xbc5]
```

For a comprehensive driver name display, you can combine the **/c** and **/g** parameters in a command. (The order of parameters does not change the output.) The following command lists allocations for tags beginning with Ip. It uses the **/c** parameter, which adds contents of the Localtag.txt file to the Mapped_Driver column, and the **/g** parameter, which adds the contents of the Pooltag.txt to the Mapped_Driver column.

**poolmon /iIp** /c /g**

In the resulting display, the Mapped_Driver column contains data from both the localtag.txt and pooltag.txt files.

```
 Memory:  130616K Avail:   23692K  PageFlts:   146   InRam Krnl: 2108K P: 9532K
 Commit: 187940K Limit: 318628K Peak: 192000K            Pool N: 8372K P:13384K
 System pool information
 Tag  Type     Allocs            Frees            Diff   Bytes      Per Alloc
Mapped_Driver

 IpEQ Nonp          1 (   0)         0 (   0)        1    1808 (     0)   1808 [ipsec]
[ipsec.sys    -  event queue]
 IpFI Nonp         26 (   0)         0 (   0)       26    7408 (     0)    284 [ipsec]
[ipsec.sys    -  Filter blocks]
 IpHP Nonp          1 (   0)         1 (   0)        0       0 (     0)      0
[ipsec.sys    - IP Security]
 IpIO Nonp          1 (   0)         1 (   0)        0       0 (     0)      0 [ipsec]
 IpLA Nonp          1 (   0)         0 (   0)        1     248 (     0)    248 [ipsec]
[ipsec.sys    -  lookaside lists]
 IpSH Nonp          1 (   0)         1 (   0)        0       0 (     0)      0
[ipsec.sys    - IP Security]
 IpSI Nonp       1027 (   0)         0 (   0)     1027   53272 (     0)     51 [ipsec]
[ipsec.sys    - initial allocations]
IpTI Nonp          3 (   0)         0 (   0)        3    5400 (     0)   1800 [ipsec]
[ipsec.sys    -  timers]
```

**Example 3: Detect Memory Leakage**

This example outlines a procedure for using Poolmon to detect a memory leak.

1. Start Poolmon in default mode (no additional parameters).

2. Press P twice to display allocations from only the paged pool. (The P key toggles the display between paged, non-paged, and both.)

3. Press B to sort the Bytes column in descending order.

4. Let Poolmon run for a few hours. Because starting Poolmon changes the data, you must let it run until it reaches a steady state before the data is reliable.

5. Save the information generated by Poolmon, either as a screenshot, or by copying it from the command window and pasting it into Notepad.

6. Returning to Poolmon, press P twice again, this time to display only allocations from the non-paged pool.

Repeat steps 3, 5 and 6 approximately every half-hour for at least two hours.

When data collection is complete, examine the **Diff** (allocations minus frees) and **Bytes** (number of bytes allocated minus number of bytes freed) values for each tag, and note any that continually increase. Next, stop Poolmon, wait for a few hours, and then restart Poolmon. Examine the allocations that were increasing, and determine whether the bytes are now freed. Allocations that have still not been freed, or have continued to increase in size are the likely culprits.

**Example 4: Examine a Pool Memory Leak**
The following example demonstrates using Poolmon to investigate a pool memory leak from a suspected printer driver. In the example, Poolmon displays data that the system collects about memory allocations with the Dsrd tag, the tag that the printer drivers assign.

The Drsd tag is assigned to pool memory allocated by the **Drvalloc()** function in the Rasdd.dll and Hprasdd.dll files. Printer drivers use this function when allocating Graphical Device Interface (GDI) objects and associated memory. If a printer driver has an object leak, the memory allocated with the Drsd tag will also leak.

### Note

- Before running the steps in this example, ensure that the printer you are using will not be interrupted (if, for example, it is a network printer), until you are finished, otherwise the results may be invalid.

1. At the command line, type the following:

   ```
   poolmon /iDrsd
   ```

   This command directs Poolmon to display information for allocations with the Dsrd tag. (Pool tags are case-sensitive, so be sure to type the command exactly as shown.)

2. Record the values in the **Diff** and **Bytes** columns. In the following sample display, the value of **Diff** is 21 and the number of bytes is 17472.

   ```
   Memory:  130480K Avail:   91856K  PageFlts:  1220   InRam Krnl: 2484K P: 7988K
   Commit:  30104K Limit: 248432K Peak:  34028K            Pool N: 2224K P: 8004K
   Tag  Type        Allocs           Frees           Diff  Bytes           Per Alloc

   Drsd Paged       560 ( 177)       539 ( 171)       21   17472 (  4992)    832
   ```

3. Send a job to the printer, wait briefly for the system to return to normal, and then record the values for the **Diff** and **Bytes** columns.

   ```
   Memory:  130480K Avail:   91808K  PageFlts:  1240   InRam Krnl: 2488K P: 7996K
   Commit:  30152K Limit: 248432K Peak:  34052K            Pool N: 2224K P: 8012K
   ```

```
Tag  Type        Allocs           Frees           Diff  Bytes          Per Alloc

Drsd Paged       737 (   0)       710 (   0)       27   22464 (     0)    832
```

When the memory management for the printer driver is working properly, the value of **Diff** should return to its original value of 21 after printing. However, as the preceding output illustrates, the value of **Diff** rose to 27, and the number of bytes rose to 22464. The difference between the initial and subsequent output means that six Drsd blocks, with a total of 4992 bytes, leaked during printing.

**For More Information**

If you believe you have identified a leaking driver, you may be able to find more information by visiting the Product Support Services Web site and searching the Knowledge Base for relevant articles.

**Example 5: Monitor a Terminal Services Session**

This example shows several ways to display allocations from the Terminal Services session pools on the system. It demonstrates the use of the **/s** command-line parameter, and the **s**, TSSessionID, and **i** running parameters.

The following command displays allocations from all of the Terminal Services session pools. In this example, the local computer, which is configured as a Terminal Server, is hosting the sessions, and client computers are using the Remote Desktop feature to connect to the host.

**poolmon /s**

In response, Poolmon displays allocations from all session pools. Note the **All sessions pool information** title in the header.

```
Memory:  523572K Avail:  233036K  PageFlts:   344   InRam Krnl: 1828K P:18380K
Commit: 193632K Limit:1279764K Peak: 987356K            Pool N:14332K P:18644K
All sessions pool information
 Tag  Type     Allocs            Frees            Diff   Bytes       Per Alloc

 Bmfd Paged       361 (   0)       336 (   0)       25   57832 (     0)   2313
 DDfb Paged        34 (   0)        22 (   0)       12     720 (     0)     60
 Dddp Paged         8 (   0)         6 (   0)        2     272 (     0)    136
 Dh 1 Paged        24 (   0)        24 (   0)        0       0 (     0)      0
 Dh 2 Paged       344 (   0)       344 (   0)        0       0 (     0)      0
 Dvgr Paged         2 (   0)         2 (   0)        0       0 (     0)      0
 GDev Paged       108 (   0)       102 (   0)        6   20272 (     0)   3378
 GFil Paged        29 (   0)        27 (   0)        2     160 (     0)     80
 GPal Paged        11 (   0)         8 (   0)        3     816 (     0)    272
 GTmp Paged     88876 (   1)     88876 (   1)        0       0 (     0)      0
 GUma Paged         2 (   0)         2 (   0)        0       0 (     0)      0
 Galp Paged      3250 (   0)      3250 (   0)        0       0 (     0)      0
 Gbaf Paged      9829 (   0)      9801 (   0)       28   19712 (     0)    704
 Gcac Paged      3761 (   0)      3706 (   0)       55  288968 (     0)   5253
 Gcsl Paged         1 (   0)         0 (   0)        1     488 (     0)    488
 Gdbr Paged      6277 (   0)      6271 (   0)        6    1872 (     0)    312
 ...
```

To see allocations from a particular session pool, type the session ID immediately after the **/s** parameter, as shown in the following command. This command displays session pool allocations for Terminal Services session 0.

**poolmon /s0**

In response, Poolmon displays allocations from the session pool for Terminal Services session 0. Note the **Session 0 pool information** title in the header.

```
Memory:  523572K Avail:  233024K  PageFlts:   525   InRam Krnl: 1828K P:18384K
 Commit: 193760K Limit:1279764K Peak: 987356K            Pool N:14340K P:18644K
 Session 0 pool information
 Tag  Type     Allocs            Frees            Diff   Bytes       Per Alloc

 Bmfd Paged       361 (   0)       336 (   0)       25   57832 (     0)   2313
 DDfb Paged        34 (   0)        22 (   0)       12     720 (     0)     60
 Dddp Paged         8 (   0)         6 (   0)        2     272 (     0)    136
 Dh 1 Paged        24 (   0)        24 (   0)        0       0 (     0)      0
```

```
 Dh 2 Paged       344 (   0)       344 (   0)        0       0 (     0)      0
 Dvgr Paged         2 (   0)         2 (   0)        0       0 (     0)      0
 GDev Paged       108 (   0)       102 (   0)        6   20272 (     0)   3378
 GFil Paged        29 (   0)        27 (   0)        2     160 (     0)     80
 GPal Paged        11 (   0)         8 (   0)        3     816 (     0)    272
 GTmp Paged     89079 (  99)     89079 (  99)        0       0 (     0)      0
 GUma Paged         2 (   0)         2 (   0)        0       0 (     0)      0
 Galp Paged      3250 (   0)      3250 (   0)        0       0 (     0)      0
 Gbaf Paged      9830 (   0)      9802 (   0)       28   19712 (     0)    704
 Gcac Paged      3762 (   0)      3707 (   0)       55  283632 (     0)   5156
 Gcsl Paged         1 (   0)         0 (   0)        1     488 (     0)    488
 Gdbr Paged      6280 (   0)      6274 (   0)        6    1872 (     0)    312
 ...
```

To help determine which drivers and components are allocating memory from the session pool, add the **/g** parameter, as shown in the following command. The **/g** parameter adds a Mapped_Driver column listing the Windows components and drivers that assign each tag.

**poolmon /s0 /g**

```
Memory:  523572K Avail:  235876K  PageFlts:    43   InRam Krnl: 1900K P:18860K
Commit: 185040K Limit:1279764K Peak: 987356K            Pool N:14684K P:19124K
Session 0 pool information
Tag  Type     Allocs            Frees            Diff   Bytes      Per Alloc
Mapped_Driver

Bmfd Paged       421 (   0)       396 (   0)       25   57832 (     0)   2313 [Font
related stuff]
DDfb Paged        34 (   0)        22 (   0)       12     720 (     0)     60 Unknown
Driver
Dddp Paged        11 (   0)         6 (   0)        5     392 (     0)     78 Unknown
Driver
Dh 1 Paged        37 (   0)        35 (   0)        2     224 (     0)    112 Unknown
Driver
Dh 2 Paged       367 (   0)       364 (   0)        3     912 (     0)    304 Unknown
Driver
Dvgr Paged         2 (   0)         2 (   0)        0       0 (     0)      0 [vga for
risc video driver]
GDev Paged       119 (   0)       113 (   0)        6   20272 (     0)   3378 [Gdi
pdev]
GFil Paged        29 (   0)        27 (   0)        2     160 (     0)     80 [Gdi
engine descriptor list]
GPal Paged        11 (   0)         8 (   0)        3     816 (     0)    272 [Gdi
Objects]
GTmp Paged     98626 (   1)     98626 (   1)        0       0 (     0)      0 [Gdi
Objects]
GUma Paged         2 (   0)         2 (   0)        0       0 (     0)      0 [Gdi
Objects]
Galp Paged      3250 (   0)      3250 (   0)        0       0 (     0)      0 [Gdi
Objects]
Gbaf Paged     10331 (   0)     10305 (   0)       26   18304 (     0)    704 [Gdi
Objects]
Gcac Paged      4722 (   0)      4666 (   0)       56  305400 (     0)   5453 [Gdi
glyph cache]
Gcsl Paged         1 (   0)         0 (   0)        1     488 (     0)    488 [Gdi
string resource script names]
Gdbr Paged      6972 (   0)      6965 (   0)        7    2184 (     0)    312 [Gdi
driver brush realization]
```

You can also configure the Terminal Services session pool display while Poolmon is running. The following table shows a series of running commands, in the order in which they are typed, and the resulting Poolmon display.

The series begins with a command to start Poolmon. All other parameters are typed while Poolmon is running.

**poolmon**

| Key | Result | Description |
|---|---|---|

| | | |
|---|---|---|
| s | Displays all session pools. | |
| s | Displays the system pools. | The **s** parameter toggles the display between the system pools and the Terminal Services session pools. |
| 0 | Displays session 0 pool. | You can type a session ID while displaying the system pools. |
| 7 | Displays session 7 pool. | |
| a | Displays pool allocations for session 7, sorted by number of allocations | All standard running parameters are valid for session pool displays. |
| 0 | Displays allocations for session 0, sorted by number of allocations. | Session and sorting options are retained until changed. |
| s | Displays the system pools. | |
| s | Displays allocations for session 0, sorted by number of allocations. | Session option is retained. |
| 10 | Displays Session 1 allocations, and then displays Session 0 allocations. | Without **i**, you can enter only session IDs 0 - 9. |
| i | Prompts for a Terminal Services session ID. | |
| 10 | Displays Session 10 allocations. | |
| i | Prompts for a Terminal Services session ID. | To display all session pools, press **i** and then press **Enter**. |
| Enter | Displays all session pools. | |

Only systems configured as a Terminal Server allocate memory from the session pool. If you use Poolmon to display the session pool on a computer that is not a Terminal Server, or if you type a session ID that does not exist on the system, Poolmon does not display any allocations. Instead, it displays only the headings with general memory data.

The following command displays allocations from all Terminal Services session pools.

**poolmon /s**

The following figure shows the Poolmon display that would result, if the **/s** command were submitted to a computer running Windows XP Professional, that could not be configured as a Terminal Server.

```
Memory:  260620K Avail:   44956K  PageFlts:   308   InRam Krnl: 2744K P:20444K
Commit: 185452K Limit: 640872K Peak: 192472K            Pool N: 8112K P:20648K
All sessions pool information
Tag  Type     Allocs            Frees            Diff   Bytes       Per Alloc
```

**See Also**
**Concepts**
Poolmon Overview
Poolmon Remarks
Poolmon Syntax
Alphabetical List of Tools
Windiff Overview
Remote Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Poolmon Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Poolmon Syntax**
Poolmon has two separate, but related, sets of parameters. You use one parameter set at the command line when starting Poolmon; you use the other parameter set while running Poolmon.

**Command-line Syntax**
Poolmon uses the following syntax at the command line:

**poolmon**[**/i***Tag*][**/x***Tag*][**/c** [*LocalTagFile*]][**/g** [*PoolTagFile*]][**/s**[*TSSessionID*]][[**/p** [**/p**]][**/e**][**/(** | **/)**][**/t** | **/a**| **/f**| **/d** | **/b**| **/m**][**/l**] [**/n** [*File*]][**/?** | **/h**]

**Parameters**

**/i *Tag***

Displays only the allocations with the specified tag. Tags are case sensitive. Do not type a space between the **/i** and the tag.

You can use the wildcard character (*) to represent zero or more instances of any character and the question mark character (?) to represent one instance of any character. Do not begin a tag with the wildcard character.

**/x *Tag***

Displays all allocations *except* those matching the specified tag. Tags are case sensitive. Do not type a space between the **/x** and the tag.

You can use the wildcard character (*) to represent zero or more instances of any character and the question mark character (?) to represent one instance of any character. Do not begin a tag with the wildcard character.

**/c[ *LocalTagFile*]**

Adds a column to the display (Mapped_Driver) listing the drivers on the local computer that assign each tag.

*LocalTagFile* is the name and location of a formatted text file that lists the drivers on the local computer and the tag values that they assign. This file is the data source for the Mapped_Driver column when you use the **/c** parameter. The default is Localtag.txt.

Poolmon can generate a local tag file for you. If you use the **/c** parameter and do not specify an alternate local tag file, and Poolmon does not find a Localtag.txt file in the current directory, it scans the drivers on the local computer (C:\Windows\System32\Drivers*.sys) and generates a Localtag.txt file.

**Notes**

- Poolmon cannot generate a Localtag.txt file on 64-bit versions of Windows Server 2003. As a result, the **/c** parameter and its functionality are available only on 32-vit versions of Windows.
- You can use the **/c** and **/g** parameters in the same command. If you do, the Mapped_Driver column displays data from both the local tag and pool tag files.

**/g[ *PoolTagFile*]**

Adds a column to the display (Mapped_Driver) listing Windows components and commonly used drivers that assign each tag.

*PoolTagFile* is the name and location of a formatted text file that lists the names of Windows components and commonly used drivers and the tag values they assign. This file is the data source for the Mapped_Driver column that appears when you use the **/g** parameter. The default is pooltag.txt, a file provided by Microsoft. Pooltag.txt is installed when you install Support Tools or a Windows kernel debugger.

**Note**

- You can use the **/c** and **/g** parameters in the same command. If you do, the Mapped_Driver column displays data from both the local tag and pool tag files.

**/s[ *TSSessionID*]**
Displays allocations from the Terminal Services session pools. If you specify a Terminal Service session ID, Poolmon displays the allocations from the memory pool for that session. Do not type a space between the **/s** parameter and the session ID.

**Note**

- The system allocates memory from Terminal Services session pools only when the computer is configured as a terminal server. For more information, see "Terminal Services Session Support" in Poolmon Remarks.

**/p**
Displays only allocations from the non-paged pool.

**/p /p**
Displays only allocations from the paged pool.

**/e**
Displays pool totals. The totals appear at the bottom of the display.

**/(or /)**
Turns on the sort-by-change mode. With **/(** or **/)**, Poolmon sorts by the change in a value (allocations, frees, bytes), instead of the value. The change in each value is displayed in a parentheses after the value.

Use with **/a**, **/f**, **/b** or **/m**. For example, **poolmon /a** sorts the display by number of allocations, while **poolmon /( /a** sorts the display by the change in the number of allocations.

**/t**
Sorts alphabetically by tag name. **/t** is the default.

**/a**
Sorts tags by the number of allocations.

**/f**
Sorts tags by the number of frees.

**/d**
Sorts tags by size difference between allocations and frees.

**/b**
Sorts tags by bytes used.

**/m**
Sorts tags by bytes per allocation.

**/l**
Turns highlighting off. By default, Poolmon highlights values that have changed since the last update.

**/n[ *File*]**
Saves a snapshot of the Poolmon output to a file, instead of displaying it in a command window. You can include other command-line parameters to configure the output.

*File* specifies the name and location of the snapshot file. The default is poolsnap.log.

**/?or /?**
Displays command-line syntax.

**Running Syntax**
While Poolmon runs, you can use the following keys to change the display:

▸[**p**][**(** | **)**][**s**][*TSSessionID*][**i**][**l**][**e**][**t**][**a**][**f**][**d**][**b** | **u**][**m**][**h** | **?**]

**p**
Toggles the display between paged allocations, non-paged allocations, and both.

**(or )**
Toggles display between sorting by value (allocations, frees, bytes) and sorting by change in value.

**s**

Toggles the display between the system pools and the Terminal Services session pools.

***TSSessionID***
Displays allocations from the specified Terminal Services session pool. *TSSessionID* represents the session ID of a Terminal Services session. It must be an integer from 0 - 9. To display all session pools or to enter session IDs greater than 9, use the **i** parameter.

**i**
Prompts you for the session ID of a Terminal Server session.

To respond to the prompt:

- To display allocations from all Terminal Services session pools, press Enter.
- To display allocations from a particular Terminal Services session pool, type a session ID.

**l**
Toggles highlighting of changed lines on and off.

**e**
Toggles pool totals on and off. Totals appear at the bottom of the display.

**t**
Sorts by tag name. **t** is the default.

**a**
Sorts by number of allocations. When used with **(**, **a** sorts by the change in allocations.

**f**
Sorts by number of frees. When used with **(**, **f** sorts by the change in frees.

**d**
Sorts by difference between allocations and frees.

**bor u**
Sorts by bytes used. When used with **(**, **b** and **u** sort by the change in bytes used.

**m**
Sorts by bytes per allocation. When used with **(**, **m** sorts by the change in bytes per allocations.

**hor ?**
Displays help including command line syntax, running syntax, and a description of the output. To close help, press the ESC key.

**qor ESC**
Quits Poolmon.

**See Also**
**Concepts**
Poolmon Overview
Poolmon Remarks
Poolmon Examples
Alphabetical List of Tools
Iasparse Overview
Bindiff Overview

# Remote Administration Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Remote Administration Tools**
Tools in this category are used to manage the following Windows features:

- Emergency Management Services
- Remote Desktop
- Telnet
- Terminal Services

**Note**

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Remote Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote.exe: Remote Command Line**
This command-line tool allows you to run command-line programs on remote computers. You can also use Remote to distribute the processing requirements for a specified task across several computers.

Remote.exe provides remote network access by means of named pipes to applications that use STDIN and STDOUT for input and output. Users at other computers on a network or users connected by a direct-dial modem connection can either view the remote session or enter commands themselves.

### Notes

- You must run Remote from the command window.
- Remote cannot be used to control graphical Windows-based applications.

### Security Remark

- Remote performs no security authorization. It permits anyone running Remote.exe Client to connect to your remote server. Because of this, the account under which the remote server was run is open to anyone who connects.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Remote administration is a method of managing one or more remote computers from a single location. In a large organization, you can use remote administration to centrally manage hundreds or even thousands of computers located in other buildings or even in other cities. In a small organization, you can use remote administration to manage a single server located in an adjacent office.

**System Requirements**
The following are the system requirements for Remote:

- Windows XP Professional or Windows Server 2003

**File Required**

- Remote.exe

**See Also**
**Concepts**
Remote Remarks
Remote Syntax
Remote Examples
Alphabetical List of Tools
Windiff Overview
Remote Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Remote Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Remarks**
Remote Limitations

In addition to exposing your server to unauthorized users, Remote has the following limitations:

- Only non–MS-DOS-based programs that use STDIN/STDOUT/STDERR for input/output will function when redirected across the network.
- Programs that use the Win32 Console API do not work correctly.
- MS-DOS-based programs do not work.
- Generally, any programs that use on-screen character mode windows or cursor manipulation (such as screen-mode editors) do not work.

Using Remote

**To use Remote**

1. Start the server by running **remote /s** on the computer where you want to run the selected program.
2. Connect to the server from another computer by running **remote /c**.

**See Also**
**Concepts**
Remote Overview
Remote Syntax
Remote Examples
Alphabetical List of Tools
Windiff Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Remote Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Examples**

**Example 1: Interact with a batch program running on a remote server**

In this example, you must first create a batch program on the server (Remoteserver3 in this case). To create the batch program, you can use any text editor, such as Notepad, or the `copy con` command at the command prompt as follows:

```
copy con testprog.bat
pause
schtasks /?
pause
^Z
```

Press CTRL+Z (`^Z`) to end the `copy con` session and save the Testprog.bat file in the default directory.

**On the server**

To start the remote session at the server, type the following on the server at the command line:

```
remote /s "testprog.bat" testsession1
```

You will see output similar to the following, confirming that the remote server received the command to continue executing the batch file Testprog.bat:

```
C:\Program Files\Support Tools>remote /s "testprog.bat" "testsession1"
****************************************
************     REMOTE    ************
************     SERVER    ************
****************************************
To Connect: Remote /C REMOTESERVER3 "testsession1"


C:\Program Files\Support Tools>pause
Press any key to continue . . .
```

This indicates that Testprog.bat started, then paused, waiting for any key to be pressed before continuing as instructed in the batch program.

**On the client**

On the client computer, open a command window and type the following at the command prompt:

```
remote /c remoteserver3 "testsession1"
```

On the client computer, you will see output confirming that the remote server received the command to continue executing the batch file Testprog.bat, as follows:

```
C:\Program Files\Support Tools>remote /c remoteserver3 "testsession1"
****************************************
************     REMOTE    ************
************     CLIENT    ************
****************************************
Connected...


C:\Program Files\Support Tools>pause
Press any key to continue . . .
**Remote: Connected to REMOTESERVER3 administrator1 [Sun 3:27 PM]
```

**Client-server interaction**

After establishing the client-server interactive session, you can press any key to continue running the program, (in this case we

use the Enter key). When Testprog.bat continues on the server, you will see output similar to that shown below.

Server output looks like this:

```
**Remote: Connected to REMOTESERVER3 administrator1 [Sun 3:27 PM]
^M                      [CLIENTCOMPUTER7 3:27 PM]


C:\Program Files\Support Tools>schtasks /?

SCHTASKS /parameter [arguments]

Description:
    Enables an administrator to create, delete, query CHange, run and
    end scheduled tasks on a local or remote system. Replaces AT.exe.

Parameter List:
    /Create         Creates a new scheduled task.

    /Delete         Deletes the scheduled task(s).

    /Query          Displays all scheduled tasks.

    /Change         Changes the properties of scheduled task.

    /Run            Runs the scheduled task immediately.

    /End            Stops the currently running scheduled task.

    /?              Displays this help/usage.

Examples:
    SCHTASKS
    SCHTASKS /?
    SCHTASKS /Run /?
    SCHTASKS /End /?
    SCHTASKS /Create /?
    SCHTASKS /Delete /?
    SCHTASKS /Query  /?
    SCHTASKS /Change /?

Remote exiting. Child (testprog.bat) exit code was 0.
```

Client output (Clientcomputer7 in this case) is similar to output from the server Remoteserver3, as shown here:

```
C:\Program Files\Support Tools>schtasks /?

SCHTASKS /parameter [arguments]

Description:
    Enables an administrator to create, delete, query CHange, run and
    end scheduled tasks on a local or remote system. Replaces AT.exe.

Parameter List:
    /Create         Creates a new scheduled task.

    /Delete         Deletes the scheduled task(s).

    /Query          Displays all scheduled tasks.

    /Change         Changes the properties of scheduled task.

    /Run            Runs the scheduled task immediately.

    /End            Stops the currently running scheduled task.

    /?              Displays this help/usage.
```

```
*** SESSION OVER ***
```

### Note

- The session on both the client and the server computers ends the moment the **remote** command completes. In some cases, the output to the client computer might be truncated if the remote command terminates before all the output is sent to the client. In this example, the remote command (`schtasks /?`) is delayed from terminating by the `Pause` statement that occurs after the `schtsks /?` command.

**Example 2: Example of a keyword/color file**
The following command creates a keyword-highlighting scheme on the client command console using the text file Colorscheme.txt:

```
remote /c server3 remotesession2 /k colorscheme.txt
```

In this example, the format of the Colorscheme.txt text file is as follows:

```
Error
red
Continue
lblue, white
WaRnInG
white, purple
```

When used in conjunction with the **/K** parameter, this file causes all instances of the word "error" to be highlighted in red, "continue" to be highlighted in light blue with a white background, and "warning" to be highlighted in white with a purple background. Keywords are case-insensitive.

**See Also**
**Concepts**
Remote Overview
Remote Remarks
Remote Syntax
Alphabetical List of Tools
Windiff Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview
Gflags.exe
Devcon Overview

# Remote Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Syntax**
Remote uses the following general syntax:

**remote** *Operation*

**Operations**
Server

Client

Query

**Server**
Starts the remote service on the server where the command is to be run.

**Syntax**
**remote /s** *Command UniqueID*[**/U**[*DomainName*****]*UserName*] [**/UD** *Name*] [**/F** *Color*] [**/B** *Color*] [**/V**][**/-V**] [**@K**]

**Parameters**
***Command***
Specifies the command to be run on the server, with concurrent input from and output to a client that will connect later.

***UniqueID***
Specifies a unique name to distinguish this instance of the remote session from others running on the server. If the session name contains spaces it must be enclosed with quotation marks.

**/U[*DomainName*\]*UserName***
Specifies the user name or group names that are permitted to connect to a specified session running on the server. The use of *DomainName* is optional. This parameter can be specified more than once to enable permissions for a broader set of users or groups.

**Security Remark**

- Users and groups are not authenticated when connecting to remote sessions on servers. This can create potential security vulnerability on the server.

**/UD *Name***
Specifies the user name or group names that are to be denied access to a specified session running on the server.

**/F *Color***
Specifies the foreground color for the session on the server. Acceptable values are: **yellow**, **black**, **blue**, **lblue**, **white**, **red**, **green**, **cyan**, and **purple**. Remote does not discern color conflicts between background and foreground. If the same color is specified for both background and foreground, text in the command console will not be visible.

**/B *Color***
Specifies the background color for the session on the server. Acceptable values are: **yellow**, **black**, **blue**, **lblue**, **white**, **red**, **green**, **cyan**, and **purple**. Remote does not discern color conflicts between background and foreground. If the same color is specified for both background and foreground, text in the command console will not be visible.

**/V**
Makes the specified session visible to remote users logging in with **/q** parameter.

**/-V**
Makes the specified session invisible to remote users logging in with **/q** parameter. By default, if the *command* looks like a debugger, the session is visible to the remote client.

**@K**
Causes the server session to terminate.

**Client**
Starts the remote service on the client and connects to a specified server session based on a server name and unique identifier.

**Syntax**
**remote /c** *ServerName UniqueID*[**/F** *Color*] [**/B** *Color*] [**/L** *Number*] [**/K**][**@Q**]

**Parameters**
***ServerName***
Specifies the server name to connect to from the client.

***UniqueID***
Specifies a unique name to connect to on a specified server. If the session name contains spaces it must be enclosed in quotation marks.

**/F *Color***
Specifies the foreground color for the session on the server. Acceptable values are: **yellow**, **black**, **blue**, **lblue**, **white**, **red**, **green**, **cyan**, and **purple**. Remote does not discern color conflicts between background and foreground. If the same color is specified for both background and foreground, text in the command console will not be visible.

**/B *Color***
Specifies the background color for the session on the server. Acceptable values are: **yellow**, **black**, **blue**, **lblue**, **white**, **red**, **green**, **cyan**, and **purple**. Remote does not discern color conflicts between background and foreground. If the same color is specified for both background and foreground, text in the command console will not be visible.

**/L *Number***
Specifies the number of lines of output to obtain from a specified session running on a remote server.

**Note**

- Setting a small number of lines can truncate output from the session running on the server.

**/K**
Assigns colors to highlight keywords based on information in a text file of the following format:

*Keyword ForegroundColor*[**,** *BackgroundColor*]

See "Example 2: Example of a keyword/color file" in Remote Examples.

Acceptable values for *ForegroundColor* and *BackgroundColor* are: **yellow**, **black**, **blue**, **lblue**, **white**, **red**, **green**, **cyan**, and **purple**. Remote does not discern color conflicts between background and foreground. If the same color is specified for both background and foreground, text in the command console will not be visible.

**@Q**
Terminates the client session, leaving the remote server session running.

**Query**
Queries and retrieves the available unique identifiers of visible sessions on the specified computer.

**Syntax**
**remote /Q** *ServerName*

**Parameters**
***ServerName***
Specifies the server name to query.

**See Also**
**Concepts**
Remote Overview
Remote Remarks
Remote Examples
Alphabetical List of Tools
Windiff Overview
Poolmon Overview
Msizap Overview
Msicuu Overview
Memsnap Overview

Gflags.exe
Devcon Overview

# Security Management Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Security Management Tools**
Tools in this category are used to manage the following Windows features:

- Access control
- Authentication
- Encrypting File System
- Public Key Infrastructure
- Software restriction policies
- Windows Time service

**Notes**

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Ksetup Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ksetup.exe: Kerberos Setup**
This command-line tool configures a client connected to a server running Windows Server 2003 to use a server running Kerberos V5. The client then uses a Kerberos V5 realm instead of a Windows Server 2003 domain. This provides a single sign-on to the Key Distribution Center (KDC) and a local client account connected to a computer running Windows Server 2003.

KSetup is part of a group of tools, including Ktpass Overview, that is used to configure Windows for Kerberos V5 interoperability.

Administrators can use KSetup to:

- Set up a realm entry for a Kerberos V5 realm.
    - Set up a list of KDCs for that realm.
    - Set up a kpasswd server for that realm.
- Set up local account to Kerberos V5 account mappings. This is necessary to inform the operating system how to authorize a specific security principal. This links authorization data to an identity. Windows domains do not need this data because they get it through other means.
- Set the computer's password in the Kerberos realm. A Kerberos realm (and Windows domain) tracks computers joined to them using a shared secret in the form of the computer's password. When joining a Kerberos realm, a host principal must be created. The password used by the realm to create the host key must be entered using this command so that Windows can decode host tickets presented to it.
- Change a user's password in a Kerberos V5 realm.

Note

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Kerberos is an authentication system designed to enable two parties to exchange private information across an otherwise open network. It assigns a unique key, called a ticket, to each user who logs on to the network. The ticket is used to secure messages and identify the sender of the message.

**System Requirements**
The following are the system requirements for KSetup:

- Windows Server 2003
- User's membership in the Administrators or Server Operators group on the target computer. Also, both the user account and the server computer must be members of the same domain or reside within trusted domains.

**File Required**

- KSetup.exe

**See Also**
**Concepts**
Ksetup Remarks
Ksetup Syntax

Ksetup Examples
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Ksetup Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**KSetup Remarks**
**Understanding KSetup**
KSetup is used to change the computer settings for locating non-Microsoft Kerberos realms. In non-Microsoft Kerberos–based implementations, this information is usually kept in the Krb5.conf file. In Windows Server 2003 systems, it is kept in the registry. You can use this tool to modify these settings. These settings are used by workstations to locate non-Windows Kerberos realms and by domain controllers to locate non-Microsoft Kerberos realms for cross-realm trust relationships.

KSetup initializes registry keys that the Kerberos Security Provider uses to locate an non-Microsoft Kerberos Key Distribution Center (KDC) if the computer running a Windows Server 2003 operating system is not a member of a Windows domain. After configuration, the user of a computer running Windows XP Professional can log on to accounts in the Kerberos realm.

The Kerberos V5 authentication protocol is the default for network authentication on computers running Windows XP Professional. Computers running Windows XP Professional can participate in Kerberos realms that are not Windows Server 2003 domains. In this case, the KDC will not be on a domain controller running Windows Server 2003, so the DNS names for KDC servers must be stored in the registry of the client computer. The Kerberos Security Support Provider (SSP) looks in the registry for the DNS domain name of the user's realm, and then resolves the name to an IP address by querying a DNS server. Kerberos can use DNS to locate KDCs, using only the realm name but must be specially configured to do so.

**See Also**
**Concepts**
Ksetup Overview
Ksetup Syntax
Ksetup Examples
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Getsid Overview
Addiag.exe

# Ksetup Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**KSetup Examples**

**Example 1: Configure a UNIX KDC server and workstation realm**

Run the KSetup tool on a Windows XP Professional–based workstation to configure the UNIX KDC server and the realm that the workstation should use. Type the following at the command line to set the UNIX realm:

```
ksetup /addkdc DOMAIN.MICROSOFT.COM mitkdc.microsoft.com
```

Finally, type the following at the command line to set the local computer account password:

```
ksetup /setcomputerpassword p@ssword1!
```

Restart the computer for the changes to take effect.

**Example 2: Authenticate a realm's principals**

Type the following at the command line to authenticate the principals of the given realm if they present valid Kerberos tickets:

```
ksetup /mapuser user@DOMAIN.MICROSOFT.COM guest
```

### Note

- The computer will authenticate the principals of the given realm if they present valid Kerberos tickets. Any principal from the given Kerberos realm will be treated as Guest for local authorization purposes.

**Example 3: Configure a single sign-on to local workstation accounts**

Type the following at the command line to authenticate the principals of the given realm if they present valid Kerberos tickets:

```
ksetup /mapuser user@DOMAIN.MICROSOFT.COM guest
```

### Note

- In this case, any principal from the given Kerberos realm will be treated as Guest for local authorization purposes. For example, accessing a file or registry key will be granted or denied as though the Guest had attempted to access the file or registry key.

**Example 4: List the available realm flags that ksetup knows**

Type the following at the command line to map clients to local accounts of the same name:

```
ksetup /ListRealmFlags
```

```
Ksetup knows the following realm flags:
0x00 None         No Realm Flags
0x01 SendAddress  Include IP numbers within tickets.
                  Useful for solving SOME compatibility issues.
0x02 TcpSupported Indicates that this realm supports TCP.
                  (as opposed to just UDP)
0x04 Delegate     Everyone in this realm is trusted for delegation
0x08 NcSupported  This realm supports Name Canonicalization...
```

### Note

- To list the available realm flags that KSetup does not know, type either of the following commands at the command line:
    - /SetRealmFlags REALM SendAddress TcpSupported Delegate NcSupported
    - /SetRealmFlags REALM 0xF

**See Also**

**Concepts**

Ksetup Overview
Ksetup Remarks
Ksetup Syntax
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Getsid Overview
Addiag.exe

# Ksetup Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**KSetup Syntax**

KSetup uses the following syntax:

**ksetup** [**/SetRealm***DnsDomainName*] [**/MapUser***PrincipalAccount*] [**/Addkdc***RealmNameKdcName*] [**/Delkdc***RealmNameKdcName*] [**/Addkpasswd***RealmNameKpasswdName*] [**/Delkpasswd***Realm*
*[flag][...]*] [**/AddRealmFlags***Realmflag[flag][flag][...]*] [**/DelRealmFlags***Realmflag[flag][flag][...]*] [**/DumpState***No Arguments*] [**/?**] [**/help**]

**Parameters**

**/SetRealm *DnsDomainName***
Sets the name of an RFC1510 Kerberos realm.

**/MapUser *Principal Account***
Maps the name of a Kerberos principal to an account (* = any/all). If account name is omitted, mapping is deleted for the specified principal.

**/AddKdc *RealmName KdcName***
Adds an additional KDC address for the given realm. If KdcName is omitted, DNS may be used to locate KDCs.

**/DelKdc *RealmName KdcName***
Deletes instance(s) of KDC address for the realm.

**/AddKpasswd *RealmName KpasswdName***
Adds Kpasswd server address for a realm.

**/DelKpasswd *RealmName KpasswdName***
Deletes Kpasswd server address for a realm.

**/Server *ServerName***
Specifies the name of a Windows 2000 computer to target the changes.

**/SetComputerPassword *Password***
Sets the password for the local computer.

**/RemoveRealm *RealmName***
Deletes all information for this realm from the registry.

**/Domain *DomainName***
Uses this domain (if DomainName is unspecified, detects it).

**/ChangePassword *OldPasswd NewPasswd***
Uses Kpasswd to change the password of the logged-on user. Uses '*' to be prompted for passwords.

**/ListRealmFlags *No arguments***
Lists the available realm flags that ksetup knows.

**/SetRealmFlags *realm flag [flag] [flag] [...]***
Sets RealmFlags for a specific realm. /SetRealmFlags requires one or more arguments that are either numerical (0x5) or labels (from the /ListRealmFlags call).

**/AddRealmFlags *realm flag[flag][flag][...]***
Adds additional RealmFlags to a realm.

**/DelRealmFlags *realm flag [flag] [flag] [...]***
Deletes RealmFlags from a realm.

**/DumpState *No arguments***
Analyzes the kerberos configuration on the given machine.

**/? or /help**
Displays the usage screen.

**Note**

- The arguments for KSetup are evaluated in the order entered. This is irrelevant for most commands except /Server.

**Examples**

See Ksetup Examples, Ksetup Examples or Ksetup Examples.

**See Also**

**Concepts**

Ksetup Overview
Ksetup Remarks
Ksetup Examples
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Getsid Overview
Addiag.exe

# Ktpass Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ktpass.exe: Kerberos Keytab Setup**
This command-line tool enables an administrator to configure a non-Windows Server 2003 Kerberos service as a security principal in the Windows Server 2003 Active Directory. KtPass configures the server principal name for the host or service in Active Directory and generates an MIT-style Kerberos "keytab" file containing the shared secret key of the service. The tool allows UNIX-based services that support Kerberos authentication to use the interoperability features provided by the Windows Server 2003 Kerberos KDC service.

### Note

- The keytab output file is used to replace or merge with the MIT Kerberos /Etc/Krb5.keytab file. MIT Kerberos-based services (noninteractive) use the keytab to log on and use Kerberos services.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Kerberos is an authentication system designed to enable two parties to exchange private information across an otherwise open network. It assigns a unique key, called a ticket, to each user that logs on to the network. The ticket is embedded in messages to identify the sender of the message.

**System Requirements**
The following are the system requirements for KtPass:

- Windows Server 2003
- User's membership in the Administrators or Server Operators group on the target computer. Also, both the user account and the server computer must be members of the same domain or reside within trusted domains.

**File Required**

- Ktpass.exe

**See Also**
**Concepts**
Ktpass Remarks
Ktpass Syntax
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Ktpass Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**KtPass Remarks**
**Using KtPass**
Services running on UNIX systems can be configured with service instance accounts in Active Directory. This allows full interoperability. MIT Kerberos clients and servers on UNIX systems can authenticate by using the Windows Server 2003 Kerberos server, and clients connected to servers running Windows Server 2003 can authenticate to Kerberos services that support GSS API.

Unlike Kerberos principal names, Windows Server 2003 account names do not have multiple parts. For this reason, it is not possible to directly create an account of the name Sample/Unix1.microsoft.com. Such a principal instance is created by using the service principal name mappings.

To generate a UNIX host keytab file, map the principal to the account, and set the host principal password:

1. Use the Active Directory User and Computers snap-in to create a user account for the UNIX service. For example, create an account with the name SampleUnix1.

2. Set up an identity mapping for the user account using KtPass by typing a command line that uses this syntax:

   **ktpass /princ *ServiceInstance@REALM* /mapuser *AccountName* /pass *Password* /out *Unixmachine.Keytab***

   **Note**

   - You cannot map multiple service instances to the same user account.

3. Merge the keytab file with the /Etc/Krb5.keytab file on the UNIX host.

**See Also**
**Concepts**
Ktpass Overview
Ktpass Syntax
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Ktpass Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**KtPass Syntax**

**ktpass** [**/out***FileName*] [**/princ***UserName*] [**/pass** {*Password*|*}] [**/mapuser***LocalUserName*] [**/mapop** {**add**|**set**}] [{-|+}**desonly**] [**/in***FileName*] [**/crpyto** {**DES-CBC-CRC**|**DES-CBC-MD5**}] [**/ptype** {**KRB5_NT_PRINCIPAL**|**KRB5_NT_SRV_INST**|**KRB5_NT_SRV_HST**}] [**/kvno***KeyVersionNum*] [**/ktvno***KeyTabVersion*] [**/rawsalt**] [{-|+}**dumpsalt**] [{-|+}**setupn**] [{-|+}**setpass***Password*] [**/**

**Parameters**

**/out *FileName***
Specifies the name of the Kerberos 5 keytable file to generate.

**Note**

- This is the keytable you transfer to the UNIX computer and replace or merge with your existing keytable file, /Etc/Krb5.keytab.

**/princ *UserName***
Specifies the principal name in the form *user@REALM*.

**/pass{ *Password*| *}**
Specifies a password for the principal user name specified in the **/princ** parameter. Use "*" to prompt for a password.

**/mapuser *LocalUserName***
Maps the name of the Kerberos principal specified by the **/princ** parameter to the specified local user name. By default, Kerberos principals are not mapped to local user names.

**/mapop{ add| set}**
Specifies how the mapping attribute is set.

| Value | Description |
|---|---|
| **add** | Adds value. This is the default. |
| **set** | Sets value. |

**{ -| +} desonly**
Sets (+) or unsets (-) an account for des-only encryption. Des-only encryption is set by default.

**/in *FileName***
Specifies the keytab to read or digest.

**/crypto{ DES-CBC-CRC| DES-CBC-MD5]**
Sets the encryption type to use.

| Value | Description |
|---|---|
| **DES-CBC-CRC** | Default |
| **DES-CBC-MD5** | More MIT-like |

**/ptype{ KRB5_NT_PRINCIPAL| KRB5_NT_SRV_INST| KRB5_NT_SRV_HST}**
Specifies the principal type.

| Value | Description |
|---|---|
| **KRB5_NT_PRINCIPAL** | General ptype (recommended). |
| **KRB5_NT_SRV_INST** | User service instance. |
| **KRB5_NT_SRV_HST** | Host service instance. |

**/kvno *KeyVersionNum***
Specifies the Key version number. The default value is 1.

**/?| /h| /help**
Displays command-line help for KtPass.

**See Also**

**Concepts**

Ktpass Overview
Ktpass Remarks
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Sidwalk Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sidwalk.exe: Security ID Administration**
This command-line tool takes a mapping file as input and scans all the Access Control Lists (ACLs) in the registry, file system, file and print shares, and local group membership. SIDWalk uses the mapping information in the mapping file to either delete or replace every occurrence of an old security identifier (SID) with the corresponding new SID. The same mapping file can be used for SIDWalk conversion on multiple computers.

### Note

- This tool is part of the Sidwalker Security Administration Tools.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
Access control is implemented by access control lists (ACLs). Every file in the NTFS file system and every registry key has a unique ACL, granting access rights to file resources to users and groups, and defining what specific access rights each is granted. Each user and group is identified in the ACL by a security identifier (SID).

**System Requirements**
SIDWalk can be run as scheduled batch jobs. It must be run locally on the computer where the access permissions are to be changed. SIDWalk uses Windows NT security APIs (supported on Windows NT 4.0, Windows 2000, and Windows XP) to examine every ACL on every object on the system. It is a CPU and I/O intensive program; plan to run it when system use is otherwise very light and expect significant resource use.

**Files Required**

- Sidwalk.exe
- Msvcrt50.dll (found in the \System32 directory)

**See Also**
**Concepts**
Sidwalk Syntax
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Sidwalk Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SIDWalk Syntax**

SIDWalk uses the following syntax:

**sidwalk** *ProfileFile*[*ProfileFile* ...][**/t**][**/f** [*Path*]][**/r**][**/s**][**/p**][**/g**][**/l***File*][**/?**]

**Parameters**

***ProfileFile***

Specifies the path and file name of one or more .csv mapping files to be used for input.

**/t**

Conducts a test/dry run.

**/f[ *Path*]**

Scans all NTFS file system files or, optionally, only those on the directory tree of the specified path.

**/r**

Scans the registry.

**/s**

Scans file shares.

**/p**

Scans printer shares.

**/g**

Scans local groups.

**/l *File***

Generates a Converter log file with *File* as the file name and path.

**/?**

Displays help on SIDwalk.

**Remarks**

- When run without parameters, SIDWalk displays help.

**See Also**

**Concepts**

Sidwalk Overview
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Sidwalker Security Administration Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SIDWalker Security Administration Tools**
This set of programs helps system administrators manage access-control policies on systems running Windows Server 2003, Windows XP, Windows 2000, and Windows NT. The SIDWalker tools are designed to support a three-phase approach to modifying access control. Each phase of changing access control can take a day or longer to complete and verify, and requires significant amounts of system resources and administrator time. The phases are:

- **Planning.** The administrator needs to determine what users have been granted access to resources (file shares, print shares, NTFS files, registry keys, and local group membership) on a particular computer.
- **Mapping Accounts.** Based on who has access to what resources on the system, the administrator can choose to delete old, unused security identifiers (SIDs), or replace them with corresponding new SIDs, such as new security groups.
- **Converting ACLs.** Using the information from the planning and mapping phases, the third phase is the conversion of security identifiers found anywhere on a system to corresponding new SIDs.

**SIDWalker Components**

- Showaccs Overview
- Sidwalk Overview
- Sidwkr.dll

**Usage Scenario**
The following example is a scenario requiring migration of access control policies:

All of the files for a project on a file server have an access control list (ACL) that grants the group ProjectTeam Read/Write access to the project files.

ProjectTeam is a security group that has a unique security identity in the Windows XP account directory. If the project changes and ProjectTeam merges into a larger project team, the administrator might want to change the access permissions on the project files. The Read/Write access-control rights granted to the original ProjectTeam can be replaced by access rights granted to a NewProjectTeam group. NewProjectTeam can be defined in a different domain and project file server. The computer on which the project file server runs can move from the current domain to a new domain for the larger project. In changing situations like this, SIDWalker can facilitate managing and updating security access policies on the servers involved.

**See Also**
**Concepts**
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Sidwkr.dll

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sidwkr.dll: Security Migration Editor Snap-In**
The Security Migration Editor is a Microsoft Management Console (MMC) snap-in that provides a graphical user interface (GUI) for mapping old security identifiers (SIDs) (listed in a mapping file created by ShowAccs), to new security identifiers. The "migration" referred to is from old SIDs to corresponding new SIDs, which are mapped by selecting from the list of user and group accounts in target domains or computers. The snap-in also gives options to ignore or delete these SIDs as they show up in access-control lists (ACLs). The Security Migration Editor updates the mapping file, which is then used as input to SIDWalk.

### Note

- This tool is part of the Sidwalker Security Administration Tools.

**Corresponding UI**
Security Migration Editor provides its own user interface. See the MMC snap-in for more information.

**Concepts**
Access control is implemented by access control lists (ACLs). Every file in the NTFS file system and every registry key has a unique ACL, granting access rights to file resources to users and groups, and defining what specific access rights each is granted. Each user and group is identified in the ACL by a security identifier (SID).

**System Requirements**
Security Migration Editor does not have to run on the same computer on which ShowAccs or SIDWalk runs. An administrator can schedule ShowAccs to run remotely on one server in the network, can transfer the resulting files to a local workstation for analysis and update using Security Migration Editor, and then can copy the mapping file back to one or more target systems and run SIDWalk on them.

**Files Required**

- Sidwalk.msc
- Msvcrt50.dll (found in the \System32 directory)
- Sidwkr.dll (the snap-in file for Security Migration Editor)

**See Also**
**Concepts**
Sidwkr Remarks
Alphabetical List of Tools
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Sidwkr Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Migration Editor Remarks**

**Installing Security Migration Editor**

Before running the Security Migration Editor snap-in, it must be initialized from the command prompt with the command:

```
regsvr32 sidwkr.dll
```

This initializes the class libraries used by the snap-in.

**Using Security Migration Editor**

To use the snap-in, add it to the MMC. Security Migration Editor can also be started from the command prompt with the command:

```
sidwalk.msc
```

Further instructions on using snap-ins can be found in Help for the MMC.

**See Also**

**Concepts**

Sidwkr.dll
Alphabetical List of Tools
Xcacls Overview
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Xcacls Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Xcacls.exe**
This command-line tool sets all file-system security options that are accessible in Windows Explorer. XcAcls does this by displaying and modifying the access control lists (ACLs) of files.

XcAcls is especially useful in unattended installations of Microsoft Windows XP. Using this tool, you can set the initial access rights for folders in which the operating system resides. When you distribute software to servers or workstations, XcAcls also offers one-step protection against deletion of directories or files by users.

**Corresponding UI**
File-system security options can be set manually from Windows Explorer.

**To set file-security options from Windows Explorer:**

1. Right-click the file or folder for which you want to set security.
2. Select **Sharing** or **Properties**.
3. Select the **Security** tab.
4. Click **Advanced**.

**Concepts**
For more information on ACLs, see "Access control overview" or "Understanding Access Control" in Help for Windows Server 2003.

**System Requirements**
The following are the system requirements for XcAcls:

- Windows XP Professional or Windows Server 2003

**File required**

- Xcacls.exe

  To download the Xcacls tool, see Free Tool Downloads on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=18748).

**See Also**
**Concepts**
Xcacls Syntax
Xcacls Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview

Connstat Overview
Cabarc Overview
Bitsadmin Overview
Xcacls Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Xcacls Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**XcAcls Examples**

**Example 1: Replace ACLs of All Files and Directories in the Current Directory**

You want to replace the existing ACLs of all files and directories in the current directory with Read and Write access for the administrator, suppressing confirmation. Type the following at the command line:

```
xcacls *.* /g administrator:rw /y
```

Notice that you are not asked to confirm the change. You see output similar to the following:

```
processed file:  C:\data\compressed.txt
processed file:  C:\data\deptdata.txt
processed file:  C:\data\dirafter.txt
processed file:  C:\data\temp.txt
processed file:  C:\data\uncompressed.txt
processed file:  C:\data\userdata.txt
```

You can check to see that the command was executed by typing the following at the command line:

```
xcacls *.*
```

You see output similar to the following, confirming that the access rights have been set for the administrator:

```
C:\data\compressed.txt MYCOMPUTER\Administrator:(special access:)

                                          READ_CONTROL
                                          SYNCHRONIZE
                                          FILE_GENERIC_READ
                                          FILE_GENERIC_WRITE
                                          FILE_GENERIC_EXECUTE
                                          FILE_READ_DATA
                                          FILE_WRITE_DATA
                                          FILE_APPEND_DATA
                                          FILE_READ_EA
                                          FILE_WRITE_EA
                                          FILE_EXECUTE
                                          FILE_READ_ATTRIBUTES
                                          FILE_WRITE_ATTRIBUTES


C:\data\deptdata.txt MYCOMPUTER\Administrator:(special access:)

                                         READ_CONTROL
                                         SYNCHRONIZE
                                         FILE_GENERIC_READ
                                         FILE_GENERIC_WRITE
                                         FILE_GENERIC_EXECUTE
                                         FILE_READ_DATA
                                         FILE_WRITE_DATA
                                         FILE_APPEND_DATA
                                         FILE_READ_EA
                                         FILE_WRITE_EA
                                         FILE_EXECUTE
                                         FILE_READ_ATTRIBUTES
                                         FILE_WRITE_ATTRIBUTES


C:\data\dirafter.txt MYCOMPUTER\Administrator:(special access:)

                                         READ_CONTROL
                                         SYNCHRONIZE
                                         FILE_GENERIC_READ
                                         FILE_GENERIC_WRITE
                                         FILE_GENERIC_EXECUTE
                                         FILE_READ_DATA
```

```
                                        FILE_WRITE_DATA
                                        FILE_APPEND_DATA
                                        FILE_READ_EA
                                        FILE_WRITE_EA
                                        FILE_EXECUTE
                                        FILE_READ_ATTRIBUTES
                                        FILE_WRITE_ATTRIBUTES


C:\data\temp.txt MYCOMPUTER\Administrator:(special access:)

                                    READ_CONTROL
                                    SYNCHRONIZE
                                    FILE_GENERIC_READ
                                    FILE_GENERIC_WRITE
                                    FILE_GENERIC_EXECUTE
                                    FILE_READ_DATA
                                    FILE_WRITE_DATA
                                    FILE_APPEND_DATA
                                    FILE_READ_EA
                                    FILE_WRITE_EA
                                    FILE_EXECUTE
                                    FILE_READ_ATTRIBUTES
                                    FILE_WRITE_ATTRIBUTES


C:\data\uncompressed.txt MYCOMPUTER\Administrator:(special access:)

                                        READ_CONTROL
                                        SYNCHRONIZE
                                        FILE_GENERIC_READ
                                        FILE_GENERIC_WRITE
                                        FILE_GENERIC_EXECUTE
                                        FILE_READ_DATA
                                        FILE_WRITE_DATA
                                        FILE_APPEND_DATA
                                        FILE_READ_EA
                                        FILE_WRITE_EA
                                        FILE_EXECUTE
                                        FILE_READ_ATTRIBUTES
                                        FILE_WRITE_ATTRIBUTES


C:\data\userdata.txt MYCOMPUTER\Administrator:(special access:)

                                        READ_CONTROL
                                        SYNCHRONIZE
                                        FILE_GENERIC_READ
                                        FILE_GENERIC_WRITE
                                        FILE_GENERIC_EXECUTE
                                        FILE_READ_DATA
                                        FILE_WRITE_DATA
                                        FILE_APPEND_DATA
                                        FILE_READ_EA
                                        FILE_WRITE_EA
                                        FILE_EXECUTE
                                        FILE_READ_ATTRIBUTES
                                        FILE_WRITE_ATTRIBUTES
```

**Example 2: Edit the ACLs of the Current Directory**

You want to give TestUser Read, Write, Run, and Delete rights on all new files created in this directory, but only Read and Write permissions on the directory itself. Type the following at the command line:

**xcacls *.* /g TestUser:rwed;rw /e**

You see output similar to the following:

```
processed file:  C:\test\compressed.txt
processed file:  C:\test\deptdata.txt
processed file:  C:\test\dirafter.txt
processed file:  C:\test\temp.txt
processed file:  C:\test\uncompressed.txt
processed file:  C:\test\userdata.txt
```

```
C:\data\compressed.txt Everyone:(special access:)
                               READ_CONTROL
                               SYNCHRONIZE
                               FILE_GENERIC_READ
                               FILE_GENERIC_WRITE
                               FILE_GENERIC_EXECUTE
                               FILE_READ_DATA
                               FILE_WRITE_DATA
                               FILE_APPEND_DATA
                               FILE_READ_EA
                               FILE_WRITE_EA
                               FILE_EXECUTE
                               FILE_READ_ATTRIBUTES
                               FILE_WRITE_ATTRIBUTES

                       MYCOMPUTER\TestUser:C

C:\data\deptdata.txt Everyone:(special access:)
                             READ_CONTROL
                             SYNCHRONIZE
                             FILE_GENERIC_READ
                             FILE_GENERIC_WRITE
                             FILE_GENERIC_EXECUTE
                             FILE_READ_DATA
                             FILE_WRITE_DATA
                             FILE_APPEND_DATA
                             FILE_READ_EA
                             FILE_WRITE_EA
                             FILE_EXECUTE
                             FILE_READ_ATTRIBUTES
                             FILE_WRITE_ATTRIBUTES

                     MYCOMPUTER\TestUser:C

C:\data\dirafter.txt Everyone:(special access:)
                             READ_CONTROL
                             SYNCHRONIZE
                             FILE_GENERIC_READ
                             FILE_GENERIC_WRITE
                             FILE_GENERIC_EXECUTE
                             FILE_READ_DATA
                             FILE_WRITE_DATA
                             FILE_APPEND_DATA
                             FILE_READ_EA
                             FILE_WRITE_EA
                             FILE_EXECUTE
                             FILE_READ_ATTRIBUTES
                             FILE_WRITE_ATTRIBUTES

                     MYCOMPUTER\TestUser:C

C:\data\temp.txt Everyone:(special access:)
                          READ_CONTROL
                          SYNCHRONIZE
                          FILE_GENERIC_READ
                          FILE_GENERIC_WRITE
                          FILE_GENERIC_EXECUTE
                          FILE_READ_DATA
                          FILE_WRITE_DATA
                          FILE_APPEND_DATA
                          FILE_READ_EA
                          FILE_WRITE_EA
                          FILE_EXECUTE
                          FILE_READ_ATTRIBUTES
                          FILE_WRITE_ATTRIBUTES

                  MYCOMPUTER\TestUser:C

C:\data\uncompressed.txt Everyone:(special access:)
                                 READ_CONTROL
                                 SYNCHRONIZE
                                 FILE_GENERIC_READ
```

```
                                    FILE_GENERIC_WRITE
                                    FILE_GENERIC_EXECUTE
                                    FILE_READ_DATA
                                    FILE_WRITE_DATA
                                    FILE_APPEND_DATA
                                    FILE_READ_EA
                                    FILE_WRITE_EA
                                    FILE_EXECUTE
                                    FILE_READ_ATTRIBUTES
                                    FILE_WRITE_ATTRIBUTES

                           MYCOMPUTER\TestUser:C

C:\data\userdata.txt Everyone:(special access:)
                                READ_CONTROL
                                SYNCHRONIZE
                                FILE_GENERIC_READ
                                FILE_GENERIC_WRITE
                                FILE_GENERIC_EXECUTE
                                FILE_READ_DATA
                                FILE_WRITE_DATA
                                FILE_APPEND_DATA
                                FILE_READ_EA
                                FILE_WRITE_EA
                                FILE_EXECUTE
                                FILE_READ_ATTRIBUTES
                                FILE_WRITE_ATTRIBUTES

                       MYCOMPUTER\TestUser:C
```

The command edited the ACL of a file or a directory, but its effect on a directory was different. The ACE added to the directory is also an inherit ACE for new files created in this directory.

**Example 3: Edit Permissions on a Directory Without Creating an Inherit for New Files**

You want to grant Read and Write permissions on a directory for TestUser. You do not want to create an inherit entry for new files, but grant only Read access to existing files. Type the following at the command line:

**xcacls *.* /g TestUser:r;trw /e**

You see output similar to the following:

```
C:\data\compressed.txt Everyone:(special access:)
                                 READ_CONTROL
                                 SYNCHRONIZE
                                 FILE_GENERIC_READ
                                 FILE_GENERIC_WRITE
                                 FILE_GENERIC_EXECUTE
                                 FILE_READ_DATA
                                 FILE_WRITE_DATA
                                 FILE_APPEND_DATA
                                 FILE_READ_EA
                                 FILE_WRITE_EA
                                 FILE_EXECUTE
                                 FILE_READ_ATTRIBUTES
                                 FILE_WRITE_ATTRIBUTES

                          MYCOMPUTER\TestUser:C

C:\data\deptdata.txt Everyone:(special access:)
                                READ_CONTROL
                                SYNCHRONIZE
                                FILE_GENERIC_READ
                                FILE_GENERIC_WRITE
                                FILE_GENERIC_EXECUTE
                                FILE_READ_DATA
                                FILE_WRITE_DATA
                                FILE_APPEND_DATA
                                FILE_READ_EA
                                FILE_WRITE_EA
                                FILE_EXECUTE
                                FILE_READ_ATTRIBUTES
                                FILE_WRITE_ATTRIBUTES
```

```
                    MYCOMPUTER\TestUser:C

C:\data\dirafter.txt Everyone:(special access:)
                             READ_CONTROL
                             SYNCHRONIZE
                             FILE_GENERIC_READ
                             FILE_GENERIC_WRITE
                             FILE_GENERIC_EXECUTE
                             FILE_READ_DATA
                             FILE_WRITE_DATA
                             FILE_APPEND_DATA
                             FILE_READ_EA
                             FILE_WRITE_EA
                             FILE_EXECUTE
                             FILE_READ_ATTRIBUTES
                             FILE_WRITE_ATTRIBUTES

                    MYCOMPUTER\TestUser:C

C:\data\temp.txt Everyone:(special access:)
                         READ_CONTROL
                         SYNCHRONIZE
                         FILE_GENERIC_READ
                         FILE_GENERIC_WRITE
                         FILE_GENERIC_EXECUTE
                         FILE_READ_DATA
                         FILE_WRITE_DATA
                         FILE_APPEND_DATA
                         FILE_READ_EA
                         FILE_WRITE_EA
                         FILE_EXECUTE
                         FILE_READ_ATTRIBUTES
                         FILE_WRITE_ATTRIBUTES

                 MYCOMPUTER\TestUser:C

C:\data\uncompressed.txt Everyone:(special access:)
                                 READ_CONTROL
                                 SYNCHRONIZE
                                 FILE_GENERIC_READ
                                 FILE_GENERIC_WRITE
                                 FILE_GENERIC_EXECUTE
                                 FILE_READ_DATA
                                 FILE_WRITE_DATA
                                 FILE_APPEND_DATA
                                 FILE_READ_EA
                                 FILE_WRITE_EA
                                 FILE_EXECUTE
                                 FILE_READ_ATTRIBUTES
                                 FILE_WRITE_ATTRIBUTES

                           MYCOMPUTER\TestUser:C

C:\data\userdata.txt Everyone:(special access:)
                             READ_CONTROL
                             SYNCHRONIZE
                             FILE_GENERIC_READ
                             FILE_GENERIC_WRITE
                             FILE_GENERIC_EXECUTE
                             FILE_READ_DATA
                             FILE_WRITE_DATA
                             FILE_APPEND_DATA
                             FILE_READ_EA
                             FILE_WRITE_EA
                             FILE_EXECUTE
                             FILE_READ_ATTRIBUTES
                             FILE_WRITE_ATTRIBUTES

                    MYCOMPUTER\TestUser:C
```

**See Also**
**Concepts**

Xcacls Overview
Xcacls Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Xcacls Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**XcAcls Syntax**

XcAcls uses the following syntax:

**xcacls** *FileName*[**/t**][**/e**][**/x**][**/c**][**/g***User*:*Permissions*;*Spec*][**/r***User*][**/p***User*:*Permissions*;*Spec*] [...]][**/d***User* [...]][**/y**][**/?**|**/h**]

**Parameters**

***FileName***

Indicates the name of the file or directory to which the access control list (ACL) or access control entry (ACE) is typically applied. All standard wildcard characters can be used.

**/t**

Recursively walks through the current directory and all of its subdirectories, applying the chosen access rights to the matching files or directories.

**/e**

Edits the ACL instead of replacing it.

**/x**

Edits the ACL instead of replacing it; affects only the ACEs that the specified users already own.

**/c**

Causes XcAcls to continue if an "access denied" error occurs. If **/c** is omitted, XcAcls stops on this error.

**/g *User*: *Permissions*; *Spec***

Grants *User* access to the matching file or directory. *Permissions* apply the specified access right to files. *Spec* can be the same as *Permissions* and only applies to a directory. In this case, *Permissions* is used for file inheritance in this directory. *Permissions* and *Spec* accept the following values:

| Value | Description |
|---|---|
| **R** | Read |
| **C** | Change (write) |
| **F** | Full Control |
| **P** | Change Permissions (special access) |
| **O** | Take Ownership (special access) |
| **X** | Execute (special access) |
| **E** | Read (Special access) |
| **W** | Write (Special access) |
| **D** | Delete (Special access) |
| **T** | Used only by *Spec*. ACE not specified. Sets an ACE for the directory itself without specifying an ACE that is applied to new files created in that directory. At least one access right must follow. *Spec* entries that precede T (in other words, those between the semicolon and the T) are ignored. (Special value) |

**Notes**

- The access options for files (special file and directory access options for directories) are identical. For detailed explanations of these options, see Permissions for files and folders on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48679) or Share permissions on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48680).

- All other options, which can also be set in Windows Explorer, are subsets of all possible combinations of the basic access rights. Therefore, there are no special options for directory access rights, such as LIST or READ.

***/r User***
Revokes all access rights for the specified user.

***/p User: Permissions; Spec***
Replaces access rights for *User*. *Permissions* applies the specified access right to files. *Spec* can be the same as *Permissions* and only applies to a directory. In this case, *Permissions* is used for file inheritance in this directory. *Permissions* and *Spec* accept the following values:

| Value | Description |
|---|---|
| **R** | Read |
| **C** | Change (write) |
| **F** | Full Control |
| **P** | Change Permissions (special access) |
| **O** | Take Ownership (special access) |
| **X** | Execute (special access) |
| **E** | Read (special access) |
| **W** | Write (special access) |
| **D** | Delete (special access) |
| **T** | Used only by *Spec*. ACE not specified. Sets an ACE for the directory itself without specifying an ACE that is applied to new files created in that directory. At least one access right must follow. *Spec* entries that precede T (in other words, those between the semicolon and the T) are ignored. (Special value.) |

**Notes**

- The access options for files (special file and directory access options for directories) are identical. For detailed explanations of these options, see Permissions for files and folders on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48679) or Share permissions on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=48680).

- All other options, which can also be set in Windows Explorer, are subsets of all possible combinations of the basic access rights. Therefore, there are no special options for directory access rights, such as LIST or READ.

***/d User***
Denies the specified user access to the file or directory.

***/y***
Disables the confirmation when replacing user access rights. By default, XcAcls asks for confirmation. Because of this feature, when XcAcls is used in a batch routine, the routine stops responding until the right answer is entered. Use the **/y** parameter when you use XcAcls in batch mode.

**See Also**
**Concepts**
Xcacls Overview
Xcacls Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview

Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview
Sidwkr.dll
Sidwalker Security Administration Tools
Sidwalk Overview
Showaccs Overview
Sdcheck Overview
Ktpass Overview
Ksetup Overview
Getsid Overview
Addiag.exe

# Software and System Deployment Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: Software and System Deployment Tools**
Tools in this category are used to manage the following Windows features:

- RIS
- Terminal Services

**Notes**

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Addiag.exe

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Addiag.exe: Application Deployment Diagnosis**
This command-line tool provides information on the current state of software either installed or available for installation on a computer managed by IntelliMirror Software Installation and Maintenance.

AdDiag provides the following information about software on a computer:

- The current user, including logon credentials and security identifier (SID), and the platform (processor, locale) that could affect the managed software.
- Whether or not Terminal Server is running on the computer.
- Installed or advertised information about software, derived from the registry.
- Advertised information about available software stored in Active Directory.
- Windows Installer information.

**Corresponding UI**
There is no corresponding user interface for this tool.

**Concepts**
IntelliMirror is a set of features that enable desktop change and configuration management technology. These features include: user data management, software installation and maintenance, user settings management, and remote installation services. For more information, see IntelliMirror.

**System Requirements**
The following are the system requirements for AdDiag:

- Windows XP Professional or Windows Server 2003

**Files Required**

- Addiag.exe

**See Also**
**Concepts**
Addiag Remarks
Addiag Syntax
Addiag Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview

Cabarc Overview
Bitsadmin Overview

# Addiag Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AdDiag Remarks**

**Understanding AdDiag Report Output**
AdDiag can generate a report containing the following sections of information.

**General Info**
Provides user configuration details.

**TS Info**
Provides information about whether or not Microsoft Terminal Server is running.

**Managed Apps (Local List)**
Lists all applications with information stored in the registry. The applications are listed by GUID. The GUID is set from the Software Installation snap-in, so in a case where the friendly name does not appear correct, the GUID is the information of most use in troubleshooting.

Assign count is the count of full advertisements executed for the application. This should make a difference only in Roaming User Profile (RUP) cases, the only cases when additional full advertisements are made for the same user.

The following table lists values returned for UI level and their descriptions.

| Level | Description |
|---|---|
| 0x0 | UI level is unchanged. |
| 0x1 | Default UI. |
| 0x2 | Completely silent installation. |
| 0x3 | Progress bars and error handling. |
| 0x4 | Authored UI (wizard dialog boxes suppressed). |
| 0x5 | Authored UI (wizards, progress bars, and error handling). |
| 0x40 | Display only progress bars. |
| 0x80 | Display success or failure dialog box on completion. |

The following table lists values returned for Application State and their descriptions.

| State | Description |
|---|---|
| 0x1 | Assigned. |
| 0x2 | Published. |
| 0x4 | Remove unmanaged version before assignment. |
| 0x8 | Leave application unmanaged when Group Policy no longer applies. |
| 0x10 | Remove application when Group Policy no longer applies. |
| 0x20 | Leave application unmanaged. |
| 0x40 | Remove application. |
| 0x18 | Assigned, and leave application unmanaged when Group Policy no longer applies. |

**Deployed Apps (Server)**

Deployed Apps (Server) is a full list of applications from the Group Policy objects applied to the given user (or computer). In the case of upgrade relationships, the base application is still listed, even if there is an upgrade for it. This feature has not yet been implemented.

**MSI Apps**
MSI Apps lists all applications for which there is Windows Installer information. Managed applications come from Software Installation and Maintenance, unmanaged applications come from alternative distribution, such as Microsoft Systems Management Server.

**See Also**
**Concepts**
Addiag.exe
Addiag Syntax
Addiag Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Addiag Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AdDiag Examples**

**Example: Display diagnostic information for the local computer**

If you want to display diagnostic information for the local computer, type the following at the command line:

```
addiag
```

**Sample AdDiag output**

```
Microsoft (R) Software Distribution Diagnostics. Version 2.00
Copyright (C) Microsoft Corp 1998-2001. All rights reserved.

Collecting info...

Initializing Remote DS Data...
Initializing Local AppMgmt Registry Data...
Initializing Local AppMgmt File Data...
Initializing Local Windows Installer Data...
Initializing Local Shell Data...
Initializing Local Event Data...

========================= General Info =========================

User -- NameSamCompatible: TEST\vtwuster
User -- NameFullyQualifiedDN: CN=Vee T Wuster,
CN=Users,DC=test,DC=corp,DC=mstest,DC=com
User -- Logon Server: \\RED-DC-24
User -- SID: S-1-5-21-2127521184-1604012920-1887927527-340073
User -- Profile Type: LOCAL
User -- Locale: 1033
Processor Architecture: x86
System Locale: 1033

========================= TS Info =========================

Running Remote Admin TS

========================= Managed Apps (Local List) =========================

No Managed applications were found.

========================= Managed Applications (Server) =========================

User dump for test.corp.mstest.com
Dumping GPO list (3 items)...
GPO GUID: {12A7A27A-31AE-4DDC-AAC1-5744AB20EC41}
Name: EU-Office-MachinesDeployment-DuongN
GPO GUID: {6E810346-A86F-43CB-B6C6-BFA45220367A}
Name: EU-MngDskRecep-ITGIntel
cardkey
Object GUID: {31DEF562-37F4-4388-8DBC-77C0321F606D}
Package Flags:
PostBeta3
UserInstall
OnDemandInstall
Assigned
OrphanOnPolicyRemoval
MinimalInstallUI
FullInstallUI
ProductCode: {00AC0A39-F2C0-4605-93E7-B8FA37EB3C1D}
UI Level: Basic
Guest99
Object GUID: {0B294F89-AAD0-4AF2-BFF8-AF0A9891663D}
Package Flags:
```

```
PostBeta3
UserInstall
OnDemandInstall
Assigned
OrphanOnPolicyRemoval
MinimalInstallUI
FullInstallUI
ProductCode: {96F41056-72F2-4633-97FC-066C74570A70}
UI Level: Basic
Multi-Calendar Viewer
Object GUID: {818A3226-C828-4A1F-971F-4912D9007688}
Package Flags:
PostBeta3
UserInstall
OnDemandInstall
Assigned
OrphanOnPolicyRemoval
MinimalInstallUI
FullInstallUI
ProductCode: {057EA138-1893-49CC-B0DE-A39B9C6AAE5B}
UI Level: Basic
NT-PTouch
Object GUID: {B28338C5-9EFB-4F06-B94B-099C902B5AAB}
Package Flags:
PostBeta3
UserInstall
OnDemandInstall
Assigned
OrphanOnPolicyRemoval
MinimalInstallUI
FullInstallUI
ProductCode: {8B76F0D0-4DDD-4010-8A58-20F97DA1396B}
UI Level: Basic
GPO GUID: {C52B5368-2A6F-11D3-BB3C-00805FC792AD}
Name: EU-PublicApplicationPublishing -ITGINTEL
No Apps were found.

========================= Windows Installer Apps =========================

Found 3 MSI application(s)
WebFldrs
Windows 2000 Support Tools
Unkown Product

========================= Local GPO List =========================

Found 1 GPO(s)
{C52B5368-2A6F-11D3-BB3C-00805FC792AD}
Found 1 LDAP GPO path(s)
LDAP://CN=ClassStore,CN=User,CN={C52B5368-2A6F-11D3-BB3C-
00805FC792AD},CN=Policies,CN=System,DC=redmond,DC=corp,DC=mstest,DC=com

========================= Local script list =========================

No scripts were found.

========================= Local AD History =========================

Found 1 Applied GPO(s) in the history
EU-PublicApplicationPublishing -ITGINTEL
GPO GUID: {C52B5368-2A6F-11D3-BB3C-00805FC792AD}
Version: 0x8530853

========================= Local MSI Features =========================

Found 340 Component(s)
{72E6B800-9F98-41DA-BA77-2422BBDF7C20}
{D3D3AA10-5724-4452-A145-722E9A8A641E}
{F806B220-8624-4536-8FB5-E30AA3BA60E5}
{1987D340-49FF-4662-A030-B0B2EC218390}
{0498F050-4147-4D26-B060-7439CFA89A8F}
{4BF89850-1076-46A8-B1D9-868B48936900}
```

{4C628B60-32D1-4A1F-BE89-8A3D45972856}
{DC6FE870-F9BF-4ACB-B556-C31BE47F8111}
{3884E680-7F23-4C41-AEE2-4EF6B8BE53EA}
{2E2A8A80-A8B3-4B6F-B584-12303DC0338E}
{4C631E80-2C47-469A-BD6E-E536BC45BD6E}
{FA72DD90-8D2C-43B2-9D34-B24F69167CC2}
{713DC0A0-1FD8-4997-AD46-F00A51D4DD12}
{7A2AE1B0-C105-4104-B1FD-E00F177C6026}
{A20821D0-DBCC-4BF1-8C8E-B87AC1A5F77B}
{31DE2AD0-EF89-4384-AEE5-68C42112E842}
{B1F3F1E0-7C3E-4E7D-B277-68FF27E587FA}
{9A775CF0-A8CC-42CB-84DC-2AA630CF7C90}
{3A22F401-CCAD-4C56-8E39-BAFBAD767E88}
{9272CD01-F9F3-49B2-967D-6A5A987DC130}
{CF2C1711-621C-4178-88E3-E6F1D5DD4B69}
{184D3A11-6EAF-49DD-B6A0-E3C6161368C9}
{360A2321-32E9-4500-8931-2922A93E9416}
{0D472051-D355-4396-964C-8680C19CDCEC}
{46C48451-D488-11D1-825C-00A0C9056598}
{E3C05A71-7190-4C15-9FFC-67B9B02EFA40}
{346EDF71-BB1D-436A-ADE8-7312F16992E9}
{E1B4C781-9795-4D53-9DC4-C101038C0A93}
{74AE9A91-630D-43CF-B2F6-E55EC3C978FB}
{712E65A1-4D0A-4E4D-8091-5C8040EAA7BF}
{BA8AF7A1-A934-4DBC-AD54-212D055956C0}
{10B1A0B1-BD14-4174-BA2B-A83FDF95AD53}
{6D9584B1-A9F2-4CEF-8D86-64FB12CC4958}
{75A5C2C1-F1ED-4407-9364-8873C77F47F6}
{82D57FC1-AE20-4052-B6E7-1E18B05C19C7}
{8EC173D1-28D8-4A82-9955-DADC3A822E64}
{471019D1-8080-40CE-ACA1-7E68AD2528A9}
{F0F94AE1-A0CE-483F-BF9E-55706E3C370A}
{9114AEE1-35FA-11D2-AA89-00A0C90F57B0}
{FA5D46F1-904A-4FF6-ACB0-D532667AAA91}
{0C8FB9F1-5A98-408F-9E46-46FBA5208F7F}
{6FC80602-4D83-4EF0-9169-A6E839705AFF}
{23A31322-CDB0-40EB-8A87-BECBDD4B0AAB}
{24C80252-E892-4BDE-A1CE-10CF7CCBA983}
{F3514452-ADCC-408E-B787-8C2ABC21916E}
{E9BF2062-9DA7-4978-AD3F-A657422F2996}
{C28B2462-CF89-4733-AB67-5A49EF20CDF3}
{95144762-C2B8-11D1-998F-00A0C90A43B3}
{6A688D62-9CCC-4CE3-8327-5AD19A38C89D}
{AC248182-4629-4926-B1B5-83D3F9399D63}
{3FEC8482-CB99-4851-95F7-C956AACCA1D0}
{7EA3A482-87A3-48BF-8F83-977EE0C5A7A9}
{8B6AC782-C4CB-429D-8882-C20E47864CDA}
{5D871C92-7C90-4A64-96AA-329F08BA7664}
{E6E308A2-6C1D-41C3-A550-47ED6D83C6D3}
{CD7A2AE2-9A72-42B4-B3DC-3E4040039691}
{651C6013-AC86-4766-857B-7513F0D161F7}
{404D6423-5263-4344-8E1F-EC308A3E4E5A}
{1BC9CE23-D092-424a-9AC4-1D8E0C805FAD}
{FABD3633-4A5D-4BE0-9BC6-FCE695E8A154}
{F08F8E33-249A-4354-B545-AB1575F466C8}
{1A36D663-3D41-4FB7-B607-E958BABF6270}
{10FA5B63-BCB7-47fc-A528-12C6C5B4F2CE}
{95CFE083-EC87-43AB-B565-69042E50225D}
{BD036593-4F01-43B6-9484-FE5D472199AA}
{169BC993-212B-47FD-ABD7-93BA9783DFDA}
{A842C1B3-6B02-4DD9-8B2E-43675F4952F2}
{BD704BB3-4825-44BC-9528-2851AF45BFF1}
{98BBCDB3-E716-4EF1-8ED4-B1560943E13A}
{15B557D3-E69D-47F4-A635-E1489B672A58}
{89A01AD3-0024-4d10-B4F9-DE0118F0B7BF}
{51C3FFD3-056C-475F-92E7-F4F2FFFB0039}
{5DEAADE3-8ADD-4C32-A744-7CFF136E82F6}
{FC0675F3-40D3-4DA2-8EFC-8BADC80F9D54}
{71C6C8F3-0C32-439F-8D6A-08DA2080FEFE}
{6BD06DF3-30D4-4AE2-B76F-81B8506A891C}
{E091A004-FB86-4B76-87E6-B025AB4C0D0A}
{1EDF1324-66BE-495E-92E6-AC5FE0E75210}

{2A5CD034-623D-4EF3-AC31-E2550603ABAA}
{17856334-1F20-4698-9807-3BB947834DA5}
{74BA4744-69FC-4A11-A875-C1BDE850356B}
{54BB3944-D72D-4761-B393-D3EE0B22E4AA}
{22889D54-D591-46C8-87C3-6B82F9AF137E}
{5B23D074-05FE-4109-A6DE-6CBBFC532DBB}
{E274C084-3DCD-407B-BEFC-E359BD510738}
{1F9F2684-E80A-4F8A-87C9-5BB828BAA6E3}
{24BC3D84-F636-4335-AB13-73F6409B5582}
{E8152194-5020-4FC4-BBFC-CD8E999B628D}
{BE74C294-69B9-438A-9828-C48821348127}
{6E4CCBA4-270E-4A64-BAC8-0CCF47CD8B57}
{B6C009B4-66C9-4D1D-BBB7-F0936447EECE}
{930314C4-4CC2-431D-B14A-E81CFB69B13E}
{4BD778C4-80A6-4AA4-891E-D34867D3CED0}
{B1E881E4-615E-4A8A-B50A-69B4D524D3E4}
{27E29DE4-7EF0-48A7-ADE3-3FC856767FF8}
{5ACF21F4-2DE4-4F29-9A88-8A9425DDCE02}
{CE609FF4-5D85-402E-A850-ECC8E0582427}
{030CCA25-3E6D-41C6-B03D-6B33EB99545B}
{74321545-DA27-4D49-87EE-2F6EE362E9F7}
{3DAA2F45-86F4-4230-9CAB-D4CD5CE00DF1}
{EA2BA555-6C5B-4CC7-BAE4-80BD53B19E67}
{2425AB55-E917-4EA6-8E75-097012A82A9C}
{07463865-6A74-43D6-B877-CEECAC94B4E4}
{3E440985-B7A9-4BFE-BAD6-5E3A4DA6AD69}
{BE8F3B85-7B71-4C01-B528-6D2D0CB137B4}
{5CB1EE85-0D14-445E-A0E7-A96E7407D28D}
{26E90895-A003-466D-9C7F-8249D1397026}
{1F1E0B95-0267-494C-87FF-818C21924169}
{33A08F95-84AD-4AD9-9A87-CE7198B512AC}
{FB5D81A5-F73C-4A0A-9ED9-4231B5BA7E24}
{1A5459C5-8FAF-11D2-829D-00A0C9056598}
{E9CFEBC5-0288-433A-BBB7-743671008D51}
{CD1567E5-91CC-49AD-850A-4922105C1832}
{9E706BE5-3DFA-494A-88CE-837A42CF5B4B}
{A3129DE5-904C-4CEC-A2CD-0432A75069E2}
{F8EAF3F5-8509-43CC-BF9B-DEE445AB0A1E}
{B07BB5F5-5745-42E2-9AB3-BE784671FCD7}
{5B6E9FF5-80CC-43BC-B57E-13179681A27D}
{BE434906-5FE4-4951-834B-6F821F91E9FC}
{8BBF5A06-7325-4609-BC13-F85073BE84B8}
{01437F16-0EF8-4C81-9C3C-7B891BB02531}
{6F401036-34EB-4386-8E63-2FD33EA780C9}
{6ED00436-20C0-4AEE-851D-F82E39ECC3F1}
{284CEB36-F58C-4416-8649-455CDAC30A49}
{92713E36-B9B8-4113-ABB8-A17B8FE5C53A}
{E0384656-4B1C-448A-AC99-E0D9FBC30D59}
{035B4B56-D2BB-494B-969D-B90569E112D7}
{18863E76-7BE2-432B-A0C7-032478EFC06D}
{B842C486-AB4D-448B-8721-1212A47384BC}
{FE181F86-9B82-42B4-9303-AA6D26CDF83C}
{1B845896-EE51-4385-B8AE-05461C9B7869}
{B3213B96-D0B2-48E7-9F17-CC340389E8E7}
{308E3E96-711C-411D-BE17-7DF130C775E4}
{771992A6-DC57-4EA9-A607-A0727071E943}
{CBA468B6-4246-48E9-8684-B86AFEE1C8D2}
{35FC98B6-5591-4E75-91AF-739DADB46940}
{102834C6-CC7D-413A-A234-0AD5DC2D6176}
{502925D6-95DB-4E67-A564-D7BE7EC9BC0C}
{05C78DD6-825C-4BEA-A5BE-E8E79E59FEDD}
{3687B7E6-BAB8-40D7-A4A5-E2F12EEACE00}
{52A368E6-7168-4D76-A880-C733AB375EA7}
{79CBD0F6-6646-41EB-886D-FFA000F87351}
{406BA307-E1D6-43AE-A9B1-8895085C7BBD}
{957CD707-DDFA-4709-B022-D9928F852781}
{4C541F07-895C-4C20-A17A-F106175E0889}
{CC9CD617-7BF0-49B4-92B5-A0668E2F4C1D}
{0B504737-C5AD-4BB3-863C-DA065AF006A8}
{122E5B37-F6FB-474E-97D9-73054698B2DA}
{C25E2047-80BA-497B-8894-7DB2D49A9EA8}
{4B08E247-C041-46CA-AF0A-4B42B7636269}

{DFB89957-9903-41FD-8800-681B58A932F9}
{11342067-5685-41F3-9E76-AC9ED2EF8E11}
{D8DCC477-EDA4-434B-A186-9D3A2F61F1F7}
{7F98BA77-204E-45AE-812A-5CC79AA99F01}
{5B941587-24D1-478D-949B-31787DCD910D}
{2B11D297-40F4-4EDA-BD49-C7F3DD9B2FC1}
{01B44897-9AC9-4054-8485-6650AF6D5528}
{F71152A7-F2ED-47F9-AC82-C274B5E86617}
{6A78CFA7-B53F-40E0-B6F6-D8A36EE0824A}
{114C79B7-FF9F-4BB6-A70E-38DA0659C36A}
{F823DAB7-A6D5-44B8-898E-21610586B657}
{C1873EB7-5294-404E-9F62-6B46CF18117C}
{8478CEB7-85A1-4F27-8584-56B795DBCB32}
{185370C7-5F91-4E48-BB7C-40F65C4CF976}
{0A7443C7-6E9A-11D2-AAA2-00A0C90F57B0}
{ECB9ABC7-AF3F-4C55-936C-9EEC2F1330EE}
{5888ECC7-706A-4E54-B9B3-A113616BB479}
{B93B20D7-5936-447E-B004-4597505A2E82}
{D58182D7-8347-462F-8A5F-083C5292E7B7}
{E50863D7-51CA-47F9-ACA2-1743318B3A77}
{94301CD7-82BF-404B-BBE3-4353EF37648D}
{D50CA1E7-6E6B-458B-8326-157DA6EBCC61}
{1980C2E7-BF6B-4B0A-892C-98DA35E4ABD1}
{74B422F7-98FA-4F94-B274-FD10BCD1106F}
{86398EF7-6409-4E85-9B1C-F5F63C395081}
{1581CFF7-E8C7-4C67-B929-27451328C232}
{E960DE08-80D0-4FEF-8F3C-77A32DBADB7C}
{6B383418-CE9C-4708-8422-1D870C20DAF8}
{897CE818-7E34-4FB5-8F2A-B64901D9AA92}
{B0B6BC28-DE3E-485D-A655-B3E7DF087236}
{782D1038-EF7E-4608-91D0-A3A9E7D5760C}
{C78D7438-A0EC-4143-8BFF-DB858497D7C9}
{03C6D938-F8FC-450E-9A70-321617135A9A}
{4373BA38-08EF-4B30-B1CB-7D0F1CA9A4D4}
{E77E8848-3B66-4276-A412-887C510CBC0D}
{BC417A68-FDCF-4BC6-A777-D4FF7E60746A}
{E54DFF68-AC2B-4EAE-9948-0881F6E36B77}
{52C39188-F40A-43BA-B452-B8EAD2B4CCD2}
{64CAD688-3B52-4DBC-A082-B2B2DA2B8465}
{8E126B98-48E2-4614-956D-EF3FF0876828}
{8125BC98-17F1-42D4-98E0-DA14FBB2C1E0}
{163EC9A8-2904-4014-8CCF-51B94F9D9884}
{C9345EA8-A3DE-4414-9F0F-8C2AECD34D75}
{FA532FB8-A80A-4F9E-82F1-24F6A49F82D0}
{E0917FB8-5112-4347-A393-ADC78287B2F2}
{2A3EEBC8-5006-48C3-85A5-B97D72F65C1F}
{E1F122D8-004C-4D60-87CB-5D06FE96DD12}
{B13E84E8-374D-4291-B74E-92CDD7E1F24F}
{D63721F8-60DF-4074-A32B-90E023EDD96D}
{FBB37FF8-6953-4E69-989E-05B7EFF938A7}
{D691D719-0419-4FAF-B7E2-4108BF135B0B}
{7B524139-EC13-4D18-8DE3-BDD249A837F6}
{4E063539-580B-4586-A5A2-A34909B9ABE1}
{D13A7F39-2C18-48DF-A918-3D23C4A01CF7}
{A0CB8659-733B-4493-B7D6-0CFE15B27F67}
{8EAC9369-6387-4D07-8BC7-F4EA41430ACE}
{17ADF869-F24A-4C06-9684-7E7A6783E3C5}
{5911D979-E2B6-42B9-832F-9850A4F8B7DA}
{2867F979-F7AA-4B58-A6D4-54696BCC92BC}
{BD967C89-589C-4BBF-939F-643C95791F3E}
{2E047899-40C1-41DC-A0F9-080318FBCBDC}
{761C52A9-38A1-4ADE-AC32-B99D529F5356}
{22DA28A9-904D-4FB8-81F5-86D1E91A1E5D}
{27D471B9-6CB9-476C-9291-6AFC8E4D7D3B}
{0A7443C9-6E9A-11D2-AAA2-00A0C90F57B0}
{EBF07AC9-492C-479B-BDF7-F40F1DBEA6FE}
{CAF4DDD9-5F1F-4A8D-A352-59047B04FB7D}
{F1E95DE9-51EC-4647-B512-77A236207A57}
{32488DE9-DC7B-47CC-B919-3696CFE2107B}
{2D2ADA1A-E007-496E-9D33-E83B8957EFB6}
{4083532A-7E43-45C4-85A7-44986F983879}
{CAE1A82A-3CCC-4835-AC7F-8A097356FC8F}

{42D3475A-5646-49DD-B66E-47F8A51A89AC}
{8D8EA85A-45BD-475D-B1D0-FCC63723CB30}
{AC80795A-2F81-4D99-AE1E-EFCEE333677F}
{64B3086A-DBDD-43D7-A87D-62FD77029CFB}
{47F24F6A-1CC4-4696-A562-208A940E3DB6}
{437267AA-5501-4C7F-95F3-6FAD197BDC93}
{E20B84CA-838A-4CCB-8694-801B099EE47C}
{1C50B9CA-9DA1-4f89-B349-C6E7529D2B7D}
{6A1513DA-F97E-4F5C-B07B-D8622836EB6B}
{5F1629DA-6DB0-4CC6-852D-A5E984C102E6}
{6FB091FA-F69C-4936-8761-89ADC45C00BA}
{D42364FA-0ADA-49F1-8EF1-BEDC5D400525}
{EAEFDEFA-1BF5-4C2F-AB79-30024340A330}
{973E2E0B-06BB-4494-A6FC-5D6C43B1A471}
{4DBA7F0B-BD14-4116-91E1-2F37707BB289}
{0BF3C91B-AEC9-49D2-B82F-C7DD32DE7CA6}
{CBBA222B-297F-4B64-9EB7-D4957637F2CD}
{A1A9962B-3DBC-4FB8-BB2F-9AAFD29DD2E8}
{C0B30C2B-FB0C-4C8A-8B25-270E350A8E74}
{6053B13B-16EB-495B-81CB-5BE84F9BA6EA}
{9A72D63B-CFC8-4B4E-8BE9-AA775F0C4157}
{A3B3A93B-EC56-4500-AC14-2B65BCD5AB32}
{571E314B-4FDB-42B7-B9B4-56A51773E6B9}
{4302F55B-D821-484C-8210-4A4C44F783F8}
{CC29EA5B-7BC2-11D1-A921-00A0C91E2AA2}
{BB0A2A6B-D6C0-4543-9444-C796D5298D31}
{87372D6B-9102-473E-ADE0-64B2C71F7806}
{4F6B228B-735E-4DD9-882A-DCD4066FDDB6}
{AEFB398B-7387-4591-ACAF-169C1372A97E}
{738F7D9B-DBE1-43FB-A13F-EE4DD579866D}
{C9FB9E9B-7C03-4B16-955B-25FB900EF175}
{8BAC97BB-7E6F-41D8-92CE-411BA78DADD3}
{629E6DBB-863B-4083-9BA8-7EE69552F881}
{83106FBB-A17A-4E4C-8D26-C9BD609D4CC9}
{F4FD59CB-BB61-4C74-B4CE-0EE1F3A1BC40}
{B50830EB-3225-432e-8522-B634646CA7A4}
{89AAD3EB-B4DC-45DD-BDC1-F42F274388FE}
{7AE8D9EB-1A80-4E81-957F-4F435ED69637}
{9DB393FB-ACD8-4241-B29D-57C09B3244CF}
{7922340C-35CE-4894-9C98-F80FD34C0ADA}
{9E20F71C-4391-4496-B8C3-10A968B10D9B}
{6D73902C-5DF6-4993-A1E8-6099FD61A79C}
{34B2783C-FC60-4A7D-9691-E7E1F0D4465A}
{324A9D5C-603C-4724-9E93-4A7090F46659}
{E52A036C-2C42-44E6-A066-AD61977929D8}
{AAD1266C-4458-41A7-BD0A-D6954DEA7CE3}
{61D14E6C-49AA-49E3-989C-CDBC094C07F2}
{11BA698C-A95C-47A3-9D61-A08B390BDFF0}
{4486DB8C-A07A-4657-B442-1D159D68320D}
{8598FB8C-887C-4200-8FE6-690410E1C0F3}
{EF377F8C-A278-4F14-B099-994B3E2286BC}
{A42635BC-D3A4-4595-99C5-234E39DD8AE3}
{07BA08BC-A57D-49EA-ADBC-206C2DB7DF5B}
{3A2422CC-0A87-4EF8-8CF3-348EE6C4D52A}
{8F2419CC-A782-4BE6-9C63-CDFBCA5B5009}
{9D8B72DC-9435-4B8F-8179-B3A283857F55}
{0CCB94DC-230C-4F7E-B23D-D7E033E19EF0}
{77EFFADC-8D9F-42A2-9F05-13EC3A54687F}
{01FD82EC-1D8F-4138-9222-5FB566141D6D}
{F2960CEC-4082-4D6E-BE59-FAB36D9C1F5A}
{E1DBEBFC-6D48-4246-8401-9F2DF633F058}
{EA7ECF0D-840B-47BD-BD98-9A76959C3859}
{FCA5D11D-A8B7-4148-9090-31BA09783621}
{48D57A1D-70F4-4648-8B57-9938D6E1C30F}
{A6A2305D-E507-4960-8DFF-7293EA55B842}
{E214C19D-F60F-4F57-B530-B6B302F91CBC}
{D134139D-F7FC-4D17-A756-239310D42372}
{5E16639D-876D-437A-8FF0-CBCEAAA40341}
{0A525DCD-8395-4C95-99AB-32A2B65C9BCD}
{737FD0FD-99ED-4ACB-AC0C-55835ED57F07}
{D33D9DFD-BBB6-4131-9F71-3746AA98F31D}
{99EE222E-9D4F-4B13-B3E7-B13173E7F3E4}

{6599A45E-2A0F-4BED-A749-71C3A71923A3}
{BE25E95E-0FF9-4DDC-BB63-3445A0DB0CD4}
{61985B5E-B5BB-42D9-A54B-52F7D50CFEAB}
{31AC8E6E-7DE7-41BC-8420-C25B346AB8F7}
{BAA0FF6E-627A-4053-B447-886277AE7065}
{59A5B37E-B467-4016-990D-15CAA3925AF0}
{EF6B339E-7B24-4B92-9EC2-717A2E3D90F7}
{9638249E-88CF-445E-9C81-5A5E9D7A2019}
{9A4B469E-6774-4E51-81CC-060AF3E16E41}
{B9E7BC9E-18BF-4A4A-B9F0-73ED5259D3B0}
{D7793F9E-BF88-42D9-B6B8-55941C6A3AB2}
{3D2623AE-0E58-4DFC-9C23-3594EA94E4F3}
{DF7535AE-2F1B-43B7-8334-70F9A8F5976E}
{050C37AE-681F-46FA-8E3D-8D1CC27946C7}
{6CF858BE-92FE-4316-B0D0-2B9A950BE66A}
{108F2FBE-98E4-46F6-B3ED-F98E7D34489D}
{A460BECE-653E-4576-987A-47AF3643A8B8}
{FC94F9DE-1746-47CD-A641-D714B5CA40B5}
{161F5BDE-C122-4CE3-A3A8-3518D91C298B}
{9F7C23EE-1F13-42EF-9B59-A2B0F31CC4C1}
{B7E2EFFE-60B8-4532-BABD-DEE4EB6B4BCF}
{E705E00F-FA0A-4C27-9A7F-79255E45E7A1}
{AB9DF31F-E1C9-4F91-95D5-77013C5EA649}
{3989D12F-8BE9-4045-8B6F-5A21A3AA648B}
{423E5D3F-2CC2-4CB6-ACDD-43538CA73A98}
{5CEE7A4F-B606-4E5A-999F-0E1D671029F7}
{E490E65F-8F7B-4424-8EB3-1246974680BD}
{56DA575F-51D1-4E93-868B-B2743CCEBD23}
{CC29E95F-7BC2-11D1-A921-00A0C91E2AA2}
{37533E5F-F873-427C-AA7B-F4B25FDE4DCF}
{3950916F-CA46-4447-A7D6-8864F5F4D4B0}
{A3B53D6F-8EA0-41A2-925D-9E0A452E5E1D}
{BD67287F-FF6C-453F-A78E-31F29F839199}
{5F64EB7F-CF97-4D85-AA99-D748DA4F1431}
{BCDE229F-9364-431A-B21B-34AF07F2F10B}
{45DF439F-7E93-40F0-83E4-56461B8E647D}
{2D1633AF-12E6-4FE5-B526-5797DADA99AD}
{2ADFC8AF-FB0E-4C04-9242-F22E54D76633}
{87B70FAF-1116-415D-A7BD-2659AA7E9301}
{13A049BF-9FF9-44E5-84D4-9A06B02ACE55}
{EC366DBF-7A35-4E37-BDD4-90D3CE67D297}
{DC8587CF-DEF3-46EB-B3C8-BF35ED27D847}
{920BADCF-2C91-4E68-8514-B44BD3BF23A0}
{CA113DDF-6D42-4368-8A12-6A6F4BD7DE2E}
{BFF929EF-A1B3-4ECD-A286-CCC3CA826E75}
{1E2A9BFF-8E55-46A3-BAE0-F7561E1554F3}

========== Application log events ==========

EventID: 11925
Type: ERROR
Date: 12:25:44.0000 - 02/06/2001
User: N/A
Computer: RK-TEST-01
Source: MsiInstaller
Description: Product: Windows 2000 Support Tools -- Error 1925. You do not have sufficient privileges to complete this installation for all users of the machine. Log on as administrator and then retry this installation.
Data:

EventID: 11925
Type: ERROR
Date: 12:20:44.0000 - 02/06/2001
User: N/A
Computer: RK-TEST-01
Source: MsiInstaller
Description: Product: Windows 2000 Support Tools -- Error 1925. You do not have sufficient privileges to complete this installation for all users of the machine. Log on as administrator and then retry this installation.
Data:

```
========================= MSI Shortcuts in the profile =========================

No shortcuts were found.

========================= AD Integrity Checks =========================

Running "Compare GPO Names with LDAP Path"...
Check "Compare GPO Names with LDAP Path" succeeded.

Running "Check Darwin In Progress"...
Check "Check Darwin In Progress" succeeded.

Running "Compare Darwin Components vs Products"...
Check "Compare Darwin Components vs Products" succeeded.

Running "Check Darwin vs AppMgmt"...
Check "Check Darwin vs AppMgmt" succeeded.

Running "Check Darwin vs Shell"...
Check "Check Darwin vs Shell" succeeded.

Running "Compare Local Data vs Server Data"...

*********************************************************************
This check compares the local App Management data with the data found
on the domain controller.
Please note that inconsistency between them is common when the policy
has not been recently applied.
To ensure good results please reapply the policy before this test.
Last refresh: 9:34:55 on 3/13/2001.
*********************************************************************

Check "Compare Local Data vs Server Data" succeeded.
```

**See Also**

**Concepts**

Addiag.exe
Addiag Remarks
Addiag Syntax
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# Addiag Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AdDiag Syntax**

AdDiag uses the following syntax:

**addiag** [**/verbose:**{**TRUE**|**FALSE**}] [**/user:**{**TRUE**|**FALSE**}] [**/trace:***String*] [**/test:** *StringList*] [**/debug:** *Ulong*] [**/?:**{**TRUE**|**FALSE**}]

**Parameters**

**/verbose: {TRUE|FALSE}**

Specifies the level of detail returned.

| Value | Description |
|---|---|
| **TRUE** | Runs diagnostics in verbose mode. |
| **FALSE** | Runs diagnostics in concise mode. This is the default. |

**/user: {TRUE|FALSE}**

Specifies which diagnostics to run.

| Value | Description |
|---|---|
| **TRUE** | Runs diagnostics for the user. This is the default. |
| **FALSE** | Runs diagnostics for the computer. |

**/trace: *String***

Specifies in a string what system tracing to turn on or off. The string must be between 1 and 260 characters in length.

| Value | Description |
|---|---|
| **AppMgmtOn** | Turns on verbose logging for AppMgmt. The log file is generated at %WINDIR%\debug\usermode\appmgmt.log. |
| **AppMgmtOff** | Turns off verbose logging for AppMgmt. |
| **MSIOn** | Turns on verbose logging for Windows Installer. The log file is generated either at %WINDIR%\temp\msi*.log or %temp%\msi*.log. |
| **MSIOff** | Turns off verbose logging for Windows Installer. |
| **UserEnvOff** | Turns on verbose logging for the policy engine. The log file is generated at %WINDIR%\debug\usermode\userenv.log. |
| **UserEnvOn** | Turns off verbose logging for the policy engine. |

**/test: *StringList***

Specifies the list of diagnostics to be run. The string must be between 1 and 260 characters, and multiple diagnostics must be separated by comma (,) or semicolon (;). The default *StringList* is "**Info**,**TS**,**LocalApps**,**ServerApps**,**MSIApps**,**ADHistory**,**EventDump**,**All**".

| Value | Description |
|---|---|
| **Info** | Dumps Addiag Remarks information. |
| **TS** | Dumps Addiag Remarks information. |
| **LocalApps** | Dump Addiag Remarks List. |

| | |
|---|---|
| **ServerApps** | Dumps Addiag Remarks List. |
| **MSIApps** | Dumps Addiag Remarks List. |
| **GPOList** | Dumps Local Group Policy Object List. |
| **ScriptList** | Dumps Local Script Application List. |
| **ADHistory** | Dumps Local Active Directory policy history. |
| **MSIFeatures** | Dumps Local MSI Features List. |
| **MSILinks** | Dumps MSI shortcuts in the profile. |
| **EventDump** | Dumps Application log related events. |
| **Check** | Performs Active Directory Integrity check. |
| **all** | Performs all tests. |

**/debug: *Ulong***

Specifies a debug level of *Ulong* where *Ulong* is an unsigned long (debug only).

**/?:{TRUE|FALSE}**

Specifies whether command-line help is displayed.

| Value | Description |
|---|---|
| **TRUE** | Displays command-line usage. |
| **FALSE** | Does not display command-line usage. This is the default. |

**Examples**

See Addiag Examples.

**See Also**

**Concepts**

Addiag.exe
Addiag Remarks
Addiag Examples
Alphabetical List of Tools
Topchk.cmd
Rsdir Overview
Rsdiag Overview
Iologsum Overview
Health_chk Overview
Ftonline Overview
Filever Overview
Efsinfo Overview
Dmdiag Overview
Dskprobe Overview
Diruse Overview
Dfsutil Overview
Connstat Overview
Cabarc Overview
Bitsadmin Overview

# System Management Support Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Category: System Management Tools**
Tools in this category are used to manage the following Windows features:

- Automation tools
- Group Policy
- Logon scripts
- Message Queuing (MSMQ)
- Task Scheduler

### Notes

- There may not be a tool for every Windows feature listed here.
- Some tools may appear in more than one category.

# Apmstat Overview

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Apmstat.exe: Advanced Power Management Status**
This command-line tool provides status information on Advanced Power Management (APM) features.

The output of ApmStat can be helpful in troubleshooting APM problems. The tool reports on APM-related registry entries and on APM BIOS problems.

ApmStat tells you whether a computer has multiple processors and whether it is running in ACPI mode.

**Corresponding UI**

- The control panel includes a Power Options program. If APM is installed (enabled or disabled) there will be an **APM** tab in this program. You can turn APM on and off by selecting the box in this tab. This is the only recommended and supported way to turn APM off. Turning APM on is Plug and Play and on the fly, turning it off requires a reboot. If the tab is absent, the computer is either an ACPI computer or an APM-disabled computer.

- If APM is turned on, there is a **Stand by** option under **Shutdown** in the logon dialog box invoked with CTRL-ALT-DELETE. There might also be a hibernate entry, but that is not related. (Hibernate can work even if neither APM nor ACPI are present.) Standby under APM has the same use as under ACPI.

- If the battery display is turned on in the Power Options program of Control Panel, a battery status icon appears, and it works similarly to ACPI. Note that APM computers always report a single composite battery, regardless of how many are present. (The unified/composite number is used because this is thought to be more reliable on a wide range of APM BIOSs, and it is simpler.)

- On most APM computers, pressing the power button or sleep button can put the computer on standby (suspend). Most computers require pressing the power button to resume, although some resume with a keyboard touch. Windows 2000 APM does not support custom power buttons.

**Concepts**
Advanced Power Management (APM) is an older power management scheme based on an SMI/SMM (hidden hardware) BIOS approach. In APM, most of the functionality is in a computer-specific BIOS that is hidden from the operating system. APM was first widely supported on Windows 95.

Advanced Configuration and Power Interface (ACPI) is the new scheme for power management supported on Windows 98 and Windows 2000. Many computers support both APM and ACPI.

APM on Windows 2000 is intended primarily to support older notebook computers, although it can run on desktop computers. In general, it is preferable to disable APM on desktop computers.

On Windows 2000, APM supports only battery status, suspend, resume, and hibernate. Timer wakeup, wake on LAN, wake on ring, and other APM features are not supported.

ACPI should always be used if a computer has a working ACPI BIOS. Some useful features (timed wakeups, for example) are only available through ACPI on Windows 2000.

If no ACPI BIOS is available or if it does not work well, it might be better to use APM.

ACPI should always be used if a computer has a working ACPI BIOS. Some useful features (timed wakeups, for example) are only available through ACPI on Windows 2000.

If no ACPI BIOS is available or if it does not work well, it might be better to use APM.

**System Requirements**

- Windows Server 2003

- In Windows 2000 Server, APM is disabled.
- If Windows 2000 does not find a battery installed at startup, as is generally the case with desktop computers, stand-by status is not available. Only power-off status is enabled.

**File Required**

- Apmstat.exe

**See Also**
**Concepts**
Apmstat Remarks
Apmstat Syntax
Apmstat Examples
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Devcon Overview
Apmstat Overview

# Apmstat Remarks

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**APMstat Remarks**
**APM on desktop computers**
In general, APM should be avoided on desktop computers because third-party video cards or other add-in cards might cause problems. There are some computers on which the APM BIOS causes instability all by itself. Both APM support (Ntapm.sys) in the Windows 2000 operating system and the APM BIOS itself should be turned off on such computers. Other computers, though, may be more reliable under Windows 2000 with APM support turned on, for example, those that have an APM BIOS that cannot be disabled.

### Caution

- APM should not be run on any computer that could be made unstable by it, because an unstable APM BIOS might cause loss of data.

**Key components of APM**

- **APM BIOS:** Placed in ROM or flash by the OEM. Does most of the work of powering things on and off, but is invisible to the operating system. The APM BIOS must be turned on for APM to work. If it is unstable, it might be necessary to turn the APM BIOS off in a hardware configuration screen, in addition to turning off Windows 2000 APM support. Timeouts in the APM BIOS should be off, or set as long as possible, so that the operating system rather than the APM BIOS can time things out. The APM BIOS is replaced by ACPI.

- **NtLdr:** Does a restart of APM on resume from hibernate, if APM was active before the hibernate.

- **Ntapm.sys:** The device driver that hooks the system and the APM BIOS together. Some system operations end up here for dispatch to the APM BIOS. APM BIOS events and status are polled from here. Note that when the APM BIOS presents an event (for example to suspend or power off) Ntapm.sys catches this, and then does an NtInitiatePowerAction call, which tells the operating system to perform an operation. At the very end, the power manager calls into the Hardware Abstraction Layer (HAL) which calls back into Ntapm.sys which, in turn, calls the APM BIOS. So, almost all the power code in the operating system and driver stacks is the same between APM and ACPI.

- **Hal.dll**

- **Apmbatt.sys:** Simulates a battery unit so the system battery status code can work.

- **Power Options control panel program:** Allows the user to enable or disable APM support on his or her computer. This is the preferred way (and the only supported way) to turn operating system APM support on or off.

- **Biosinfo.inf:** Lists computers on the auto-enable list and the disable list, and the BIOS detection sequences used to match them.

**Troubleshooting APM**

1. If the computer is a desktop computer, it is unlikely that it will run APM well. But you can get most of the value of APM by using hibernate, which requires neither APM nor ACPI to work.

2. If the computer is a laptop or a desktop that must run APM, make sure the user has APM turned on, hibernate turned on, and the battery icon turned on.

3. Make sure that the APM BIOS screen blanking is off. This will often make the system seem to behave badly when all that is wrong is that the screen is black. (Use a screen saver of BLANK to get the right effect.)

4. Run Apmstat.exe. Its output will often reveal a common problem, such as trying to run APM on an multi-processor

computer, for example.

5. If Windows 2000 APM has been turned off, and the computer is still unstable, try turning off the APM BIOS itself in the computer's setup screen. Some computers work better with Windows 2000 APM turned on, since this puts them more in sync with their APM BIOSs.

While troubleshooting, keep in mind the following considerations:

- Virtually all APM computers have some sort of computer setup/configuration screen, usually invoked at boot by pressing ESC, DEL, F1, F2, or F10. If you can get access to these screens, you can often improve system behavior by changing or disabling APM options.

- Some APM BIOSs turn themselves off if you disable all their timers; in this case, it is best to set them all to some long timeout value.

- Some APM BIOSs refuse to suspend if the kernel debugger is active on one of their COM ports.

- Event Log If an APM_SETPOWER call fails (that is, a standby/suspend fails at the APM BIOS interface), Ntapm.sys writes a record into the system event log, and produces data that reports what happened in addition to the error code the APM BIOS returned. This information is sometimes helpful in determining if the computer refused to suspend due to activity (that the user is not aware of) on a port.

**See Also**
**Concepts**
Apmstat Overview
Apmstat Syntax
Apmstat Examples
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Devcon Overview

# Apmstat Examples

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**APMstat Examples**

**Example 1: Provide status information on Advanced Power Management (APM)**

To provide status information on Advanced Power Management (APM) features, type:

**apmstat**

Your results will look similar to the following output:

```
This machine has an APM bios present that looks OK, and it is
not on the list of machines known to have APM problems.
Check the power applet in the control panel to see if APM is enabled.
```

**Example 2: Display a verbose report**

To display a verbose report, type:

**apmstat /v**

Your results will look similar to the following output:

```
This computer appears to have an APM legal HAL.
This machine has an APM bios present that looks OK, and it is
not on the list of machines known to have APM problems.
Check the power applet in the control panel to see if APM is enabled.
APM Registry Data Dump
Major = 0001  Minor = 0002
InstallFlags = 0003
Code16Segment = f000  Code16Offset = f10c  DataSeg = f000
Signature = APM
Valid = 0001
Detection Log Data:
44 45 54 4c 4f 47 31 00 00 00 00 00 00 00 00 00
 D  E  T  L  O  G  1
```

**See Also**

**Concepts**

Apmstat Overview
Apmstat Remarks
Apmstat Syntax
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Devcon Overview

# Apmstat Syntax

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**APMstat Syntax**

**apmstat** [**-v**] [**-?**]

**Parameters**

**apmstat**

Without arguments, displays a basic report on Advanced Power Management status on the local computer.

**-v**

Displays a verbose report with more data about APM versions, address, and setup codes from the system description of the hardware tree in the registry.

**-?**

Displays this syntax listing.

**See Also**

**Concepts**

Apmstat Overview
Apmstat Remarks
Apmstat Examples
Alphabetical List of Tools
Pviewer Overview
Exctrlst Overview
Devcon Overview

# Support Tools Glossary

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Glossary**

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**A**

**access control**
A security mechanism that determines which operations a user, group, service, or computer is authorized to perform on a computer or on a particular object, such as a file, printer, registry subkey, or directory service object.

See also: permission; registry; service

**access control entry (ACE)**
An entry in an object's discretionary access control list (DACL) that grants permissions to a user or group. An ACE is also an entry in an object's system access control list (SACL) that specifies the security events to be audited for a user or group.

See also: access control; access control list (ACL); permission; security descriptor

**access control list (ACL)**
A list of security protections that apply to an entire object, a set of the object's properties, or an individual property of an object. There are two types of access control lists: discretionary and system.

See also: access control entry (ACE); security descriptor

**ACL**
See definition for: access control list (ACL)

**Active Directory**
The Windows-based directory service. Active Directory stores information about objects on a network and makes this information available to users and network administrators. Active Directory gives network users access to permitted resources anywhere on the network using a single logon process. It provides network administrators with an intuitive, hierarchical view of the network and a single point of administration for all network objects.

See also: directory service; domain; forest; replication

**active volume**
The volume from which the computer starts up. The active volume must be a simple volume on a dynamic disk. You cannot mark an existing dynamic volume as the active volume, but you can upgrade a basic disk containing the active partition to a dynamic disk. After the disk is upgraded to dynamic, the partition becomes a simple volume that is active.

See also: basic disk; dynamic disk; simple volume

**ActiveX**
A set of technologies that allows software components to interact with one another in a networked environment, regardless of the language in which the components were created.

**American Standard Code for Information Interchange (ASCII)**
A standard single-byte character encoding scheme used for text-based data. ASCII uses designated 7-bit or 8-bit number combinations to represent either 128 or 256 possible characters. Standard ASCII uses 7 bits to represent all uppercase and lowercase letters, the numbers 0 through 9, punctuation marks, and special control characters used in U.S. English. Most current x86-based systems support the use of extended (or "high") ASCII. Extended ASCII allows the eighth bit of each character to identify an additional 128 special symbol characters, foreign-language letters, and graphic symbols.

See also: Unicode

**application programming interface (API)**
A set of routines that an application uses to request and carry out lower-level services performed by a computer's operating system. These routines usually carry out maintenance tasks such as managing files and displaying information.

**ASCII (American Standard Code for Information Interchange)**

See definition for: American Standard Code for Information Interchange (ASCII)

**B**

**base priority**
A precedence ranking that determines the order in which the threads of a process are scheduled for the processor. Use Task Manager to view and change base priorities.

For Message Queuing, a property that specifies the queue's priority in a public queue. You can set the base priority from -32,768 to 32,767; the default priority is 0. Private queues do not support base priority. Message Queuing routes and delivers messages first by base priority, then by message priority.

**basic disk**
A physical disk that can be accessed by MS-DOS and all Windows-based operating systems. Basic disks can contain up to four primary partitions, or three primary partitions and an extended partition with multiple logical drives. If you want to create partitions that span multiple disks, you must first convert the basic disk to a dynamic disk by using Disk Management or the Diskpart.exe command-line tool.

See also: dynamic disk

**basic input/output system (BIOS)**
On x86-based computers, the set of essential software routines that test hardware at startup, start the operating system, and support the transfer of data among hardware devices. The BIOS is stored in read-only memory (ROM) so that it can be executed when you turn on the computer. Although critical to performance, the BIOS is usually invisible to computer users.

See also: Extensible Firmware Interface (EFI)

**batch program**
An ASCII (unformatted text) file that contains one or more operating system commands. A batch program's file name has a .cmd or .bat extension. When you type the file name at the command prompt, or when the batch program is run from another program, its commands are processed sequentially. Also called *batch files*.

See also: American Standard Code for Information Interchange (ASCII)

**BIOS**
See definition for: basic input/output system (BIOS)

**broadcast**
The transmission of packets by an Internet Protocol version 4 (IPv4) host to all computers on the subnet.

See also: packet

**browser**
Software that interprets the markup of files in HTML, formats them into Web pages, and displays them to the end user. Some browsers also permit end users to send and receive e-mail, read newsgroups, and play sound or video files embedded in Web documents.

**C**

**capture**
To assign a port to a printer. Documents that you print are sent to the printer through the captured port.

For Network Monitor, the process by which frames are copied.

See also: frame

**capture trigger**
A set of conditions that, when met, stop the capture or execute a program or command file to perform another type of action related to the capture.

**CCP**
See definition for: Compression Control Protocol (CCP)

**certification authority (CA)**
An entity responsible for establishing and vouching for the authenticity of public keys belonging to subjects (usually users or computers) or other certification authorities. Activities of a certification authority can include binding public keys to distinguished names through signed certificates, managing certificate serial numbers, and certificate revocation.

**CGI**
See definition for: common gateway interface (CGI)

**checked build**
A version of Windows that consists of binaries that provide additional error checking, argument verification, and system debugging code. Much of the extra code in the checked binaries is in the form of ASSERT macros that test an expression. If the expression evaluates to FALSE, the macro generates a kernel debugger error message and breaks into the debugger. This lets you immediately determine the cause and location of the error. Because of the additional error checking code and debugging information, the checked binaries are larger and run slower than the free binaries. This can conceal synchronization or other timing-related problems, such as race conditions, that become apparent only in the free build. If you run the checked build of Windows NT without having enabled kernel debugging, unexpected system shutdowns can occur. This is because the additional checks in the checked build increase the likelihood of encountering a breakpoint.

**cluster**
In data storage, the smallest amount of disk space that can be allocated to hold a file. All file systems used by Windows organize hard disks based on clusters, which consist of one or more contiguous sectors. The smaller the cluster size, the more efficiently a disk stores information. If no cluster size is specified during formatting, Windows picks defaults based on the size of the volume. These defaults are selected to reduce the amount of space that is lost and the amount of fragmentation on the volume. Also called an *allocation unit*.

In computer networking, a group of independent computers that work together to provide a common set of services and present a single-system image to clients. The use of a cluster enhances the availability of the services and the scalability and manageability of the operating system that provides the services.

See also: volume

**cluster storage**
Storage where one or more attached disks hold data used either by server applications running on the cluster or by applications for managing the cluster. Each disk on the cluster storage is owned by only one node of the cluster. The ownership of disks moves from one node to another when the disk group fails over or moves to the other node.

See also: cluster

**code page**
A means of providing support for character sets and keyboard layouts for different countries or regions. A code page is a table that relates the binary character codes used by a program to keys on the keyboard or to characters on the display.

**COM**
See definition for: Component Object Model (COM)

**common gateway interface (CGI)**
A server-side interface for initiating software services. For example a set of interfaces that describe how a Web server communicates with software on the same computer. Any software can be a CGI program if it handles input and output according to the CGI standard.

**Component Object Model (COM)**
An object-based programming model designed to promote software interoperability; it allows two or more applications or components to easily cooperate with one another, even if they were written by different vendors, at different times, in different programming languages, or if they are running on different computers running different operating systems. OLE technology and ActiveX are both built on top of COM.

See also: ActiveX

**Compression Control Protocol (CCP)**
A protocol used in the negotiation process in a Point-to-Point Protocol (PPP) connection. CCP is one type of Network Control Protocol. Network Control Protocols are used to establish and configure different network protocol parameters for Internet Protocol (IP), Internetwork Packet Exchange (IPX), and NetBIOS Extended User Interface (NetBEUI).

See also: NetBIOS Extended User Interface (NetBEUI)

**console tree**
The left pane in Microsoft Management Console (MMC) that displays the items contained in the console. The items in the console tree and their hierarchical organization determine the capabilities of a console.

See also: Microsoft Management Console (MMC)

**credentials**
A set of information that includes identification and proof of identification that is used to gain access to local and network resources. Examples of credentials are user names and passwords, smart cards, and certificates.

**cyclic redundancy check (CRC)**
A procedure used in checking for errors in data transmission. CRC error checking uses a complex calculation to generate a number based on the data transmitted. The sending device performs the calculation before transmission and sends its result to the receiving device. The receiving device repeats the same calculation after transmission. If both devices obtain the same result, it is assumed that the transmission was error-free. The procedure is known as a redundancy check because each transmission includes not only data but extra (redundant) error-checking values. Communications protocols such as XMODEM and Kermit use cyclical redundancy checking.

**D**

**daemon**
A networking program, usually associated with UNIX systems, that runs in the background performing tool functions such as housekeeping or maintenance without user intervention or awareness. Pronounced "demon."

**DFS**
See definition for: Distributed File System (DFS)

**DHCP**
See definition for: Dynamic Host Configuration Protocol (DHCP)

**DHCP client**
Any network-enabled device that supports the ability to communicate with a DHCP server for the purpose of obtaining dynamic leased IP configuration and related optional parameters information.

See also: DHCP server; Dynamic Host Configuration Protocol (DHCP)

**DHCP server**
A computer running the Microsoft DHCP service that offers dynamic configuration of IP addresses and related information to DHCP-enabled clients.

See also: Dynamic Host Configuration Protocol (DHCP); IP address

**dial-up connection**
The connection to your network if you use a device that uses the telephone network. This includes modems with a standard telephone line, ISDN cards with high-speed ISDN lines, or X.25 networks.

If you are a typical user, you might have one or two dial-up connections, for example, to the Internet and to your corporate network. In a more complex server situation, multiple network modem connections might be used to implement advanced routing.

**directory service**
Both the directory information source and the service that makes the information available and usable. A directory service enables the user to find an object when given any one of its attributes.

See also: Active Directory

**disk mirroring**
A set of software processes that maintains a backup copy of a volume at all times. Each mirror of a volume resides on a different disk; ideally, each disk has its own controller. If one mirror becomes unavailable (due to a disk failure, for example), you can use the other mirror to gain access to the volume's data.

See also: mirror; volume

**distinguished name**
A name that uniquely identifies an object by using the relative distinguished name for the object, plus the names of container objects and domains that contain the object. The distinguished name identifies the object as well as its location in a tree. Every object in Active Directory has a distinguished name. A typical distinguished name might be

CN=MyName,CN=Users,DC=Microsoft,DC=Com

This identifies the MyName user object in the microsoft.com domain.

See also: Active Directory; domain

**Distributed File System (DFS)**
A service that allows system administrators to organize distributed network shares into a logical namespace, enabling users to access files without specifying their physical location and providing load sharing across network shares.

See also: service

**DLL**
See definition for: dynamic-link library (DLL)

**DNS**
See definition for: Domain Name System (DNS)

**DNS server**
A server that maintains information about a portion of the DNS database and that responds to and resolves DNS queries.

See also: Domain Name System (DNS)

**domain**
In Active Directory, a collection of computer, user, and group objects defined by the administrator. These objects share a common directory database, security policies, and security relationships with other domains.

In DNS, any tree or subtree within the DNS namespace. Although the names for DNS domains often correspond to Active Directory domains, DNS domains should not be confused with Active Directory domains.

See also: Active Directory; Domain Name System (DNS)

**domain controller**
In an Active Directory forest, a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. Administrators can manage user accounts, network access, shared resources, site topology, and other directory objects from any domain controller in the forest.

See also: Active Directory; forest; shared resource

**domain DFS**
An implementation of Distributed File System (DFS) in which DFS topological information is stored in Active Directory. Because this information is made available on multiple domain controllers in the domain, domain DFS provides fault tolerance for any distributed file system in the domain.

See also: Active Directory; Distributed File System (DFS); fault tolerance

**Domain Name System (DNS)**
A hierarchical, distributed database that contains mappings of DNS domain names to various types of data, such as IP addresses. DNS enables the location of computers and services by user-friendly names, and it also enables the discovery of other information stored in the database.

See also: IP address; service; Transmission Control Protocol/Internet Protocol (TCP/IP)

**dynamic disk**
A physical disk that provides features that basic disks do not, such as support for volumes that span multiple disks. Dynamic disks use a hidden database to track information about dynamic volumes on the disk and other dynamic disks in the computer. You convert basic disks to dynamic by using the Disk Management snap-in or the DiskPart command-line tool. When you convert a basic disk to dynamic, all existing basic volumes become dynamic volumes.

See also: active volume; basic disk; partition; volume

**Dynamic Host Configuration Protocol (DHCP)**
A TCP/IP service protocol that offers dynamic leased configuration of host IP addresses and distributes other configuration parameters to eligible network clients. DHCP provides safe, reliable, and simple TCP/IP network configuration, prevents address conflicts, and helps conserve the use of client IP addresses on the network.

DHCP uses a client/server model where the DHCP server maintains centralized management of IP addresses that are used on the network. DHCP-supporting clients can then request and obtain lease of an IP address from a DHCP server as part of their network boot process.

See also: DHCP client; DHCP server; IP address; service; Transmission Control Protocol/Internet Protocol (TCP/IP)

**dynamic-link library (DLL)**

An operating system feature that allows executable routines (generally serving a specific function or set of functions) to be stored separately as files with .dll extensions. These routines are loaded only when needed by the program that calls them.

See also: Resource DLL

**E**

**EFI**
See definition for: Extensible Firmware Interface (EFI)

**Encrypting File System (EFS)**
A feature in this version of Windows that enables users to encrypt files and folders on an NTFS volume disk to keep them safe from access by intruders.

See also: NTFS file system

**environment variable**
A string consisting of environment information, such as a drive, path, or file name, associated with a symbolic name that can be used by Windows. You use System in Control Panel or the **set** command from the command prompt to define environment variables.

See also: variable

**Extensible Firmware Interface (EFI)**
In computers with the Intel Itanium processor, the interface between a computer's firmware, hardware, and the operating system. EFI defines a new partition style called GUID partition table (GPT). EFI serves the same purpose for Itanium-based computers as the basic input/output system (BIOS) found in x86-based computers. However, it has expanded capabilities that provide a consistent way to start any compatible operating system and an easy way to add EFI drivers for new bootable devices without the need to update the computer's firmware.

See also: basic input/output system (BIOS)

**F**

**FAT**
See definition for: file allocation table (FAT)

**FAT32**
A derivative of the file allocation table (FAT) file system. FAT32 supports smaller cluster sizes and larger volumes than FAT, which results in more efficient space allocation on FAT32 volumes.

See also: file allocation table (FAT)

**fault tolerance**
The ability of computer hardware or software to ensure data integrity when hardware failures occur. Fault-tolerant features appear in many server operating systems and include mirrored volumes, RAID-5 volumes, and server clusters.

See also: cluster; RAID-5 volume

**file allocation table (FAT)**
A file system used by MS-DOS and other Windows operating systems to organize and manage files. The file allocation table is a data structure that Windows creates when you format a volume by using FAT or FAT32 file systems. Windows stores information about each file in the file allocation table so that it can retrieve the file later.

See also: FAT32; NTFS file system

**File Share resource**
A file share accessible by a network path that is supported as a cluster resource by a Resource DLL.

See also: Resource DLL

**File Transfer Protocol (FTP)**
A member of the TCP/IP suite of protocols, used to copy files between two computers on the Internet. Both computers must support their respective FTP roles: one must be an FTP client and the other an FTP server.

See also: Transmission Control Protocol/Internet Protocol (TCP/IP)

**filter**
For Indexing Service, software that extracts content and property values from a document to index them.

For Internet Protocol security (IPSec), a specification of Internet Protocol (IP) traffic that provides the ability to trigger security negotiations for a communication based on the source, destination, and type of IP traffic.

For Internet Information Services (IIS), a feature of Internet Server Application Programming Interface (ISAPI) that allows preprocessing of requests and postprocessing of responses, permitting site-specific handling of Hypertext Transfer Protocol (HTTP) requests and responses.

In IP and Internetwork Packet Exchange (IPX) packet filtering, a definition in a series of definitions that indicates to the router the type of traffic allowed or disallowed on each interface.

See also: Internet Server Application Programming Interface (ISAPI)

**flexible single-master operations (FSMO)**
See definition for: operations master

**flow**
A stream of data sent or received by a host. Also known as *network traffic*.

**forest**
One or more Active Directory domains that share the same class and attribute definitions (schema), site and replication information (configuration), and forest-wide search capabilities (global catalog). Domains in the same forest are linked with two-way, transitive trust relationships.

See also: Active Directory; domain; global catalog; schema; two-way trust

**frame**
In synchronous communication, a package of information transmitted as a single unit from one device to another.

See also: capture

**FSMO**
See definition for: operations master

**FTP**
See definition for: File Transfer Protocol (FTP)

**G**

**gateway**
A dedicated device (or a set of services running on a dedicated computer) that routes network traffic and enables communication between different networking protocols. A gateway is a multiprotocol Internet Protocol (IP) router that translates between different transport protocols or data formats.

**global catalog**
A directory database that applications and clients can query to locate any object in a forest. The global catalog is hosted on one or more domain controllers in the forest. It contains a partial replica of every domain directory partition in the forest. These partial replicas include replicas of every object in the forest, as follows: the attributes most frequently used in search operations and the attributes required to locate a full replica of the object.

In Microsoft Provisioning System, the Exchange server maintains a list of global catalogs, and it maintains a load balance across global catalogs.

See also: Active Directory; domain controller; forest; replication

**globally unique identifier (GUID)**
A 16-byte value generated from the unique identifier on a device, the current date and time, and a sequence number. A GUID is used to identify a particular device or component.

**GPO**
See definition for: Group Policy object (GPO)

**graphical user interface (GUI)**
A display format, like that of Windows, that represents a program's functions with graphic images such as buttons and icons. GUIs enable a user to perform operations and make choices by pointing and clicking with a mouse.

**Group Policy**
The infrastructure within Active Directory directory service that enables directory-based change and configuration management of user and computer settings, including security and user data. You use Group Policy to define configurations for groups of users and computers. With Group Policy, you can specify policy settings for registry-based policies, security, software installation, scripts, folder redirection, remote installation services, and Internet Explorer maintenance. The Group Policy settings that you create are contained in a Group Policy object (GPO). By associating a GPO with selected Active Directory system containers—sites, domains, and organizational units—you can apply the GPO's policy settings to the users and computers in those Active Directory containers. To create an individual GPO, use the Group Policy Object Editor. To manage Group Policy objects across an enterprise, you can use the Group Policy Management console.

See also: Active Directory; Group Policy object (GPO)

**Group Policy object (GPO)**
A collection of Group Policy settings. GPOs are essentially the documents created by the Group Policy Object Editor. GPOs are stored at the domain level, and they affect users and computers that are contained in sites, domains, and organizational units. In addition, each computer has exactly one group of policy settings stored locally, called the *local Group Policy object*.

See also: Group Policy

**GUID**
See definition for: globally unique identifier (GUID)

**H**

**handle**
In the user interface, an interface added to an object that facilitates moving, sizing, reshaping, or other functions pertaining to an object. In programming, a pointer to a pointer, that is, a token that lets a program access an identified resource.

**hardware abstraction layer (HAL)**
A thin layer of software provided by the hardware manufacturer that hides, or abstracts, hardware differences from higher layers of the operating system. By means of the filter provided by the HAL, different types of hardware look alike to the rest of the operating system. This enables the operating system to be portable from one hardware platform to another. The HAL also provides routines that enable a single device driver to support the same device on all platforms.

**high performance file system (HPFS)**
The file system designed for the OS/2 version 1.2 operating system.

**hive**
A file in which the system stores a portion of the registry (named for their resemblance to the cellular structure of a beehive). A hive is backed by a single file and a .log file, which are in either the *systemroot*\System32\Config folder or the *systemroot*\Profiles*username* folder.

By default, most hive files (Default, SAM, Security, and System) are stored in the *systemroot*\System32\Config folder. The *systemroot*\Profiles folder contains the user profile for each user of the computer. Because a hive is a file, it can be moved from one system to another. However, you must use the Registry Editor, Regedit.exe, to edit the file.

See also: key; registry

**Hosts file**
A local text file in the same format as the 4.3 Berkeley Software Distribution (BSD) UNIX /etc/hosts file. This file maps host names to IP addresses, and it is stored in the \%Systemroot%\System32\Drivers\Etc folder.

**HTML**
See definition for: Hypertext Markup Language (HTML)

**HTTP**
See definition for: Hypertext Transfer Protocol (HTTP)

**Hypertext Markup Language (HTML)**
A simple markup language used to create hypertext documents that are portable from one platform to another. HTML files are simple ASCII text files with codes embedded (indicated by markup tags) to denote formatting and hypertext links.

See also: American Standard Code for Information Interchange (ASCII)

**Hypertext Transfer Protocol (HTTP)**
The protocol used to transfer information on the World Wide Web. An HTTP address (one kind of Uniform Resource Locator

(URL)) takes the following form: http://www.microsoft.com.

See also: protocol; Uniform Resource Locator (URL)

**I**

**incorrect delegation**
A condition in which the resource records establishing a Domain Name System (DNS) zone delegation contain incorrect data or indicate a DNS server that does not have authority for the delegated zone. An improper delegation causes a DNS resolver to direct queries to an incorrect server and the query fails.

See also: Domain Name System (DNS)

**infrastructure master**
A domain controller that holds the infrastructure operations master role in Active Directory. The infrastructure master updates the group-to-user reference whenever group memberships change and replicates these changes across the domain. At any time, the infrastructure master role can be assigned to only one domain controller in each domain.

See also: Active Directory; domain controller; operations master

**IntelliMirror**
A set of change and configuration management features based on Active Directory that enables management of user and computer data and settings, including security data. IntelliMirror also provides limited ability to deploy software to Windows 2000 and later workstations or servers.

See also: Active Directory

**Internet Server Application Programming Interface (ISAPI)**
An application programming interface (API) that resides on a server computer for initiating software services tuned for Windows operating systems.

In Microsoft Provisioning System, ISAPI resides on the Web server.

See also: application programming interface (API)

**Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX)**
Transport protocols used in Novell NetWare networks, which together correspond to the combination of TCP and IP in the TCP/IP protocol suite. Windows implements IPX through NWLink.

See also: NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink); Transmission Control Protocol/Internet Protocol (TCP/IP)

**IP address**
For Internet Protocol version 4 (IPv4), a 32-bit address used to identify a node on an IPv4 internetwork. Each node on the IP internetwork must be assigned a unique IPv4 address, which is made up of the network ID, plus a unique host ID. This address is typically represented with the decimal value of each octet separated by a period (for example, 192.168.7.27). You can configure the IP address statically or dynamically by using Dynamic Host Configuration Protocol (DHCP).

For Internet Protocol version 6 (IPv6), an identifier that is assigned at the IPv6 layer to an interface or set of interfaces and that can be used as the source or destination of IPv6 packets.

See also: Dynamic Host Configuration Protocol (DHCP)

**IPX/SPX**
See definition for: Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX)

**ISAPI**
See definition for: Internet Server Application Programming Interface (ISAPI)

**J**
There are no glossary terms that begin with this letter.

**K**

**KDC**
See definition for: Key Distribution Center (KDC)

**Kerberos V5 authentication protocol**
An authentication mechanism used to verify user or host identity. The Kerberos V5 authentication protocol is the default authentication service. Internet Protocol security (IPSec) can use the Kerberos protocol for authentication.

See also: Key Distribution Center (KDC); ticket-granting service (TGS); ticket-granting ticket (TGT)

**kernel**
The core of layered architecture that manages the most basic operations of the operating system and the computer's processor. The kernel schedules different blocks of executing code, called threads, for the processor to keep it as busy as possible and coordinates multiple processors to optimize performance. The kernel also synchronizes activities among Executive-level subcomponents, such as I/O Manager and Process Manager, and handles hardware exceptions and other hardware-dependent functions. The kernel works closely with the hardware abstraction layer.

**kernel mode**
A highly privileged mode of operation where program code has direct access to all memory, including the address spaces of all user-mode processes and applications, and to hardware. Also known as *supervisor mode*, *protected mode*, or *Ring 0*.

**key**
In Registry Editor, a folder that appears in the left pane of the Registry Editor window. A key can contain subkeys and entries. For example, Environment is a key of **HKEY_CURRENT_USER**.

In IP security (IPSec), a value used in combination with an algorithm to encrypt or decrypt data. Key settings for IPSec are configurable to provide greater security.

See also: registry

**Key Distribution Center (KDC)**
A network service that supplies session tickets and temporary session keys used in the Kerberos V5 authentication protocol.

See also: Kerberos V5 authentication protocol; ticket-granting service (TGS); ticket-granting ticket (TGT)

**Knowledge Consistency Checker (KCC)**
A built-in process that runs on all domain controllers and generates the replication topology for the Active Directory forest. At specified intervals, the KCC reviews and makes modifications to the replication topology to ensure propagation of data either directly or transitively.

See also: Active Directory; domain controller; forest; replication topology

**Korn shell (ksh)**
A command shell that provides the following functionality: file input and output redirection; command-line editing by using **vi**; command history; integer arithmetic; pattern matching and variable substitution; command name abbreviation (aliasing); and built-in commands for writing shell programs.

**ksh**
See definition for: Korn shell (ksh)

**L**

**lame delegation**
See definition for: incorrect delegation

**LAN**
See definition for: local area network (LAN)

**Layer Two Tunneling Protocol (L2TP)**
An industry-standard Internet tunneling protocol that provides encapsulation for sending Point-to-Point Protocol (PPP) frames across packet-oriented media. For IP networks, L2TP traffic is sent as User Datagram Protocol (UDP) messages. In Microsoft operating systems, L2TP is used in conjunction with Internet Protocol security (IPSec) as a virtual private network (VPN) technology to provide remote access or router-to-router VPN connections. L2TP is described in RFC 2661.

See also: frame

**LDAP**
See definition for: Lightweight Directory Access Protocol (LDAP)

**Lightweight Directory Access Protocol (LDAP)**
The primary access protocol for Active Directory. LDAP is an industry-standard protocol, established by the Internet

Engineering Task Force (IETF), that allows users to query and update information in a directory service. Active Directory supports both LDAP version 2 and LDAP version 3.

See also: Active Directory; directory service; protocol

**local area network (LAN)**
A communications network connecting a group of computers, printers, and other devices located within a relatively limited area (for example, a building). A LAN enables any connected device to interact with any other on the network.

See also: NetBIOS Extended User Interface (NetBEUI); network basic input/output system (NetBIOS); virtual local area network (VLAN)

**Local Security Authority (LSA)**
A protected subsystem that authenticates and logs users on to the local computer. In addition, the LSA maintains information about all aspects of local security on a computer (collectively known as the *local security policy*), and it provides various services for translation between names and identifiers.

**M**

**MAC**
See definition for: media access control (MAC)

**majority node set server cluster**
A cluster configuration that has two or more nodes and that is configured so that the nodes may or may not be attached to one or more cluster storage devices. The cluster configuration data is stored on multiple disks across the cluster, and the Cluster service makes sure that this data is kept consistent across the different disks. There are advantages and limitations for each cluster configuration (single node server cluster, single quorum device server cluster, and majority node set server cluster).

See also: cluster; cluster storage; single quorum device server cluster

**management information base (MIB)**
A set of objects that represent various types of information about a device, used by Simple Network Management Protocol (SNMP) to manage the device. Because different network management services are used for different types of devices and protocols, each service has its own set of objects.

See also: service; Simple Network Management Protocol (SNMP)

**master boot record (MBR)**
The first sector on a hard disk, which begins the process of starting the computer. The MBR contains the partition table for the disk and a small amount of executable code called the *master boot code*.

**MBR**
See definition for: master boot record (MBR)

**media access control (MAC)**
A sublayer of the IEEE 802 specifications that defines network access methods and framing.

**Message Authentication Code (MAC)**
An algorithm that ensures the quality of a block of data.

**MIB**
See definition for: management information base (MIB)

**Microsoft Management Console (MMC)**
A framework for hosting administrative tools called *snap-ins*. A console might contain tools, folders or other containers, World Wide Web pages, and other administrative items. These items are displayed in the left pane of the console, called a *console tree*. A console has one or more windows that can provide views of the console tree. The main MMC window provides commands and tools for authoring consoles. The authoring features of MMC and the console tree itself might be hidden when a console is in User Mode.

See also: console tree; snap-in

**mirror**
One of the two volumes that make up a mirrored volume. Each mirror of a mirrored volume resides on a different disk. If one mirror becomes unavailable (due to a disk failure, for example), Windows can use the remaining mirror to gain access to the volume's data.

See also: fault tolerance; volume

**MMC**
See definition for: Microsoft Management Console (MMC)

**multicasting**
The process of sending a message simultaneously to more than one destination on a network.

**N**

**native mode**
In Windows 2000 domains, the domain mode in which all domain controllers in a domain are running Windows 2000 and a domain administrator has switched the domain operation mode from mixed mode to native mode. Native mode supports universal groups and nesting of groups. In native mode, domain controllers running Windows NT 4.0 or earlier are not supported.

In Windows Server 2003 domains, native mode is referred to as *Windows 2000 native*, and it is one of three domain functional levels available.

See also: Active Directory

**NetBEUI**
See definition for: NetBIOS Extended User Interface (NetBEUI)

**NetBIOS**
See definition for: network basic input/output system (NetBIOS)

**NetBIOS Extended User Interface (NetBEUI)**
A network protocol native to Microsoft Networking. It is usually used in small, department-size local area networks (LANs) of 1 to 200 clients. NetBEUI can use Token Ring source routing as its only method of routing. NetBEUI is the Microsoft implementation of the NetBIOS standard.

See also: Compression Control Protocol (CCP); local area network (LAN); network basic input/output system (NetBIOS); protocol

**Netsh**
A command-line and scripting tool for networking components for local or shared computers running Windows 2000, Windows XP Professional, or Windows Server 2003.

**network basic input/output system (NetBIOS)**
An application programming interface (API) that can be used by programs on a local area network (LAN). NetBIOS provides programs with a uniform set of commands for requesting the lower-level services required to manage names, conduct sessions, and send datagrams between nodes on a network.

See also: application programming interface (API); basic input/output system (BIOS); local area network (LAN); service

**nonpaged pool**
Operating system memory that is never paged to disk. Paging is the moving of infrequently used parts of a program's working memory from RAM to another storage medium, usually the hard disk. In Task Manager, the amount of memory used by a process, in kilobytes.

**notification area**
The area on the taskbar to the right of the taskbar buttons. The notification area displays the time and can also contain shortcuts that provide quick access to programs, such as Volume Control and Power Options. Other shortcuts can appear temporarily, providing information about the status of activities. For example, the printer shortcut icon appears after a document has been sent to the printer and disappears when printing is complete.

**NTFS file system**
An advanced file system that provides performance, security, reliability, and advanced features that are not found in any version of file allocation table (FAT). For example, NTFS guarantees volume consistency by using standard transaction logging and recovery techniques. If a system fails, NTFS uses its log file and checkpoint information to restore the consistency of the file system. NTFS also provides advanced features, such as file and folder permissions, encryption, disk quotas, and compression.

See also: FAT32; file allocation table (FAT)

**NTLM authentication protocol**
A challenge/response authentication protocol. The NTLM authentication protocol is supported in Windows 2000, Windows XP, and the Windows Server 2003 family, but it is not the default. It was the default authentication protocol for earlier versions of Windows.

See also: protocol

**NWLink**
See definition for: NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)

**NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)**
The Microsoft implementation of the Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) protocol used on NetWare networks. NWLink allows connectivity between Windows-based computers and NetWare networks running IPX/SPX. NWLink also provides network basic input/output system (NetBIOS) functionality and the Routing Information Protocol (RIP).

See also: Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX); network basic input/output system (NetBIOS); Routing Information Protocol over IPX (RIPX)

**O**

**object linking and embedding (OLE)**
A method for sharing information among applications. Linking an object, such as a graphic, from one document to another inserts a reference to the object into the second document. Any changes you make in the object in the first document will also be made in the second document. Embedding an object inserts a copy of an object from one document into another document. Changes you make in the object in the first document will not be updated in the second unless the embedded object is explicitly updated.

**ODBC**
See definition for: Open Database Connectivity (ODBC)

**one-way trust**
A trust relationship between two domains in which only one of the two domains trusts the other domain. For example, domain A trusts domain B, and domain B does not trust domain A. One-way trusts are often used to enable authenticated access to resource domains.

See also: domain; trust relationship; two-way trust

**Open Database Connectivity (ODBC)**
An application programming interface (API) that enables database applications to access data from a variety of existing data sources.

See also: application programming interface (API)

**operations master**
A domain controller that has been assigned one or more special roles in an Active Directory domain. The domain controllers assigned these roles perform operations that are single-master (not permitted to occur at different places on the network at the same time). Examples of these operations include resource identifier allocation, schema modification, PDC emulation, adding and removing domains to and from the forest, and tracking changes to security principals across all domains in a forest.

See also: Active Directory; domain controller; infrastructure master

**organizational unit**
An Active Directory container object used within domains. An organizational unit is a logical container into which users, groups, computers, and other organizational units are placed. It can contain objects only from its parent domain. An organizational unit is the smallest scope to which a Group Policy object (GPO) can be linked, or over which administrative authority can be delegated.

See also: Active Directory; Group Policy object (GPO)

**orphan**
A member of a mirrored volume or a RAID-5 volume that has failed due to a severe cause, such as a loss of power or a complete hard-disk head failure. When this happens, the fault-tolerant driver determines that it can no longer use the orphaned member and directs all new reads and writes to the remaining members of the fault-tolerant volume.

See also: fault tolerance; RAID-5 volume

## P

**packet**
An Open Systems Interconnection (OSI) network layer transmission unit that consists of binary information representing both data and a header containing an identification number, source and destination addresses, and error-control data.

**packet header**
In network protocol communications, a specially reserved field of a defined bit length that is attached to the front of a packet for carry and transfer of control information. When the packet arrives at its destination, the field is then detached and discarded as the packet is processed and disassembled in a corresponding reverse order for each protocol layer.

See also: packet

**page fault**
The interrupt that occurs when software attempts to read from or write to a virtual memory location that is marked *not present*.

In Task Manager, page fault is the number of times data has to be retrieved from disk for a process because it was not found in memory. The page fault value accumulates from the time the process started.

**paged pool**
The system-allocated virtual memory that has been charged to a process and that can be paged. Paging is the moving of infrequently used parts of a program's working memory from random access memory (RAM) to another storage medium, usually the hard disk.

In Task Manager, the amount of system-allocated virtual memory, in kilobytes, used by a process.

**parity**
A calculated value that is used to reconstruct data after a failure. RAID-5 volumes stripe data and parity intermittently across a set of disks. When a disk fails, some server operating systems use the parity information together with the data on good disks to recreate the data on the failed disk.

See also: fault tolerance; RAID-5 volume; striped volume

**parity bit**
In asynchronous communications, an extra bit used in checking for errors in groups of data bits transferred within or between computer systems. In modem-to-modem communications, a parity bit is often used to check the accuracy with which each character is transmitted.

See also: parity

**partition**
A portion of a physical disk that functions as though it were a physically separate disk. After you create a partition, you must format it and assign it a drive letter before you can store data on it.

On basic disks, partitions are known as basic volumes, which include primary partitions and logical drives. On dynamic disks, partitions are known as dynamic volumes, which include simple, striped, spanned, mirrored, and RAID-5 volumes.

See also: basic disk

**partition table**
On a hard disk, the data structure that stores the offset (location) and size of each primary partition on the disk. On MBR disks, the partition table is located in the master boot record. On GPT disks, the partition table is located in the GUID partition entry array.

See also: globally unique identifier (GUID); master boot record (MBR); partition

**path**
A sequence of directory (or folder) names that specifies the location of a directory, file, or folder within the Windows directory tree. Each directory name and file name within the path must be preceded by a backslash (\). For example, to specify the path of a file named Readme.doc located in the Windows directory on drive C, type **C:\Windows\Readme.doc**.

**PCMCIA device**
A removable device, approximately the size of a credit card, that can be plugged into a PCMCIA slot in a portable computer. PCMCIA devices can include modems, network adapters, and hard disk drives.

Some PCMCIA cards can be connected to and disconnected from your computer without restarting it. Before you remove the

PCMCIA card, however, you should use the Add Hardware Wizard to notify Windows that you are doing so. Windows will then notify you when you can remove the device.

**PDC**
See definition for: primary domain controller (PDC)

**Performance Monitor**
A Windows NT administrative tool that monitors performance on local or remote computers. Performance Monitor was replaced by the Performance console in Windows 2000.

**permission**
A rule associated with an object to regulate which users can gain access to the object and in what manner. Permissions are assigned or denied by the object's owner.

**POP**
See definition for: Post Office Protocol

**Portable Operating System Interface for UNIX (POSIX)**
An Institute of Electrical and Electronics Engineers (IEEE) standard that defines a set of operating-system services. Programs that adhere to the POSIX standard can be easily ported from one system to another. POSIX was based on UNIX system services, but it was created in a way that allows it to be implemented by other operating systems.

See also: service

**POSIX**
See definition for: Portable Operating System Interface for UNIX (POSIX)

**Post Office Protocol**
A maildrop service that allows a client to retrieve mail that the server is holding for it. The most recent implementation is Version 3, or POP3.

**primary domain controller (PDC)**
In a Windows NT domain, a domain controller running Windows NT Server 4.0 or earlier that authenticates domain logon attempts and updates user, computer, and group accounts in a domain. The PDC contains the master read-write copy of the directory database for the domain. A domain has only one PDC.

In a Windows 2000 or Windows Server 2003 domain, the PDC emulator master supports compatibility with client computers that are not running Windows 2000 or Windows XP Professional.

See also: Active Directory

**printer**
A device that puts text or images on paper or other print media. Examples are laser printers or dot-matrix printers.

**process**
An operating system object that consists of an executable program, a set of virtual memory addresses, and one or more threads. When a program runs, a process is created.

**promiscuous mode**
A mode in which a network device listens to all traffic present on the local network segment. A device operating in promiscuous mode reads every frame it receives, whether the frames are broadcast, multicast, or unicast. Certain devices typically operate in promiscuous mode, such as bridges, which must listen to all traffic in order to build their media access control (MAC) address tables, and network troubleshooting devices, which capture and analyze all traffic on a particular local area network (LAN) segment.

See also: broadcast; local area network (LAN)

**protocol**
A set of rules and conventions for sending information over a network. These rules govern the content, format, timing, sequencing, and error control of messages exchanged among network devices.

See also: Transmission Control Protocol/Internet Protocol (TCP/IP)

**Q**

**QoS**
See definition for: Quality of Service (QoS)

**Quality of Service (QoS)**
A set of quality assurance standards and mechanisms for data transmission, implemented in this version of Windows.

**R**

**RAID-5 volume**
A fault-tolerant volume with data and parity striped intermittently across three or more physical disks. Parity is a calculated value that is used to reconstruct data after a failure. If a portion of a physical disk fails, Windows recreates the data that was on the failed portion from the remaining data and parity. You can create RAID-5 volumes only on dynamic disks on computers running the Windows 2000 Server or Windows Server 2003 families of operating systems. You cannot mirror or extend RAID-5 volumes. In Windows NT 4.0, a RAID-5 volume was known as a *striped set with parity*.

See also: dynamic disk; fault tolerance; parity; volume

**RAS**
See definition for: Remote Access Service (RAS)

**reduced instruction set computing**
A type of microprocessor design that focuses on rapid and efficient processing of a relatively small set of instructions. RISC architecture limits the number of instructions that are built into the microprocessor, but optimizes each so it can be carried out very rapidly{bmct emdash.bmp}usually within a single clock cycle.

**Redundant Array of Independent Disks (RAID)**
A method used to standardize and categorize fault-tolerant disk systems. RAID levels provide various mixes of performance, reliability, and cost. Some servers provide three of the RAID levels: Level 0 (striping), Level 1 (mirroring), and Level 5 (RAID-5).

See also: fault tolerance; RAID-5 volume

**registry**
A database repository for information about a computer's configuration. The registry contains information that Windows continually references during operation, such as:

- Profiles for each user
- The programs installed on the computer and the types of documents that each can create
- Property settings for folders and program icons
- What hardware exists on the system
- Which ports are being used

The registry is organized hierarchically as a tree, and it is made up of keys and their subkeys, hives, and entries.

See also: hive; key

**registry key**
An identifier for a record or group of records in the registry.

**relative ID (RID)**
The part of a security ID (SID) that uniquely identifies an account or group within a domain.

See also: domain; forest; security ID (SID)

**remote access**
Part of the integrated Routing and Remote Access service that provides remote networking for telecommuters, mobile workers, and system administrators who monitor and manage servers at multiple branch offices. Users can use Network Connections to dial in to remotely access their networks for services such as file and printer sharing, electronic mail, scheduling, and SQL database access.

See also: service

**Remote Access Service (RAS)**
A Windows NT 4.0 service that provides remote networking for telecommuters, mobile workers, and system administrators

who monitor and manage servers at multiple offices.

**remote procedure call (RPC)**
A message-passing facility that allows a distributed application to call services that are available on various computers on a network. Used during remote administration of computers.

See also: service

**reparse points**
NTFS file system objects that have a definable attribute containing user-controlled data and that are used to extend functionality in the input/output (I/O) subsystem.

See also: NTFS file system

**replication**
The process of copying updated data from a data store or file system on a source computer to a matching data store or file system on one or more destination computers to synchronize the data.

In Active Directory, replication synchronizes schema, configuration, application, and domain directory partitions between domain controllers.

In Distributed File System (DFS), replication synchronizes files and folders between DFS roots and root targets.

See also: Active Directory; Distributed File System (DFS); topology

**replication topology**
In Active Directory replication, the set of physical connections that domain controllers use to replicate directory updates among domain controllers within sites and between sites.

In the File Replication service (FRS), the interconnections between replica set members. These interconnections determine the path that data takes as it replicates to all replica set members.

See also: Active Directory; Distributed File System (DFS); domain controller; replication

**Request for Comments (RFC)**
An official document of the Internet Engineering Task Force (IETF) that specifies the details for protocols included in the TCP/IP family.

See also: protocol; Transmission Control Protocol/Internet Protocol (TCP/IP)

**resource**
Generally, any part of a computer system or network, such as a disk drive, printer, or memory, that can be allotted to a running program or a process.

For Device Manager, any of four system components that control how the devices on a computer work. These four system resources are interrupt request (IRQ) lines, direct memory access (DMA) channels, input/output (I/O) ports, and memory addresses.

For server clusters, a physical or logical entity that is capable of being managed by a cluster, brought online and taken offline, and moved between nodes. A resource can be owned only by a single node at any point in time.

**Resource DLL**
A dynamic-link library (DLL) containing an implementation of the Resource application programming interface (API) for a specific type of resource. The Resource DLL is loaded into the address space of its Resource Monitor.

See also: application programming interface (API); dynamic-link library (DLL)

**RID**
See definition for: relative ID (RID)

**ring transition**
A ring transition or kernel-mode transition occurs when user-mode processes use application program interfaces (APIs) to switch their threads from user mode to kernel mode.

**RIPX**
See definition for: Routing Information Protocol over IPX (RIPX)

**RISC**

See definition for: reduced instruction set computing

**router**
Hardware that helps local area networks (LANs) and wide area networks (WANs) achieve interoperability and connectivity and that can link LANs that have different network topologies (such as Ethernet and Token Ring). Routers match packet headers to a LAN segment and choose the best path for the packet, optimizing network performance.

See also: local area network (LAN); packet header; Routing Information Protocol over IPX (RIPX); static routes; wide area network (WAN)

**Routing Information Protocol over IPX (RIPX)**
A protocol used by routers to exchange information between routers on an Internetwork Packet Exchange (IPX) network and by hosts to determine the best router to use when forwarding IPX traffic to a remote IPX network.

See also: NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink); protocol; router

**RPC**
See definition for: remote procedure call (RPC)

**S**

**SAM**
See definition for: Security Accounts Manager (SAM)

**schema**
The set of definitions for the universe of objects that can be stored in a directory. For each object class, the schema defines which attributes an instance of the class must have, which additional attributes it can have, and which other object classes can be its parent object class.

**script**
A type of program consisting of a set of instructions to an application or tool program. A script usually expresses instructions by using the application's or tool's rules and syntax, combined with simple control structures such as loops and if/then expressions. "Batch program" is often used interchangeably with "script" in the Windows environment.

**SCSI**
See definition for: small computer system interface (SCSI)

**Security Accounts Manager (SAM)**
A Windows service used during the logon process. SAM maintains user account information, including groups to which a user belongs.

See also: service

**security descriptor**
A data structure that contains security information associated with a protected object. Security descriptors include information about who owns the object, who can access it and in what way, and what types of access are audited.

See also: permission

**security ID (SID)**
A data structure of variable length that identifies user, group, and computer accounts. Every account on a network is issued a unique SID when the account is first created. Internal processes in Windows refer to an account's SID rather than the account's user or group name.

**Server Message Block (SMB)**
A file-sharing protocol designed to allow networked computers to transparently access files that reside on remote systems over a variety of networks. The SMB protocol defines a series of commands that pass information between computers. SMB uses four message types: session control, file, printer, and message.

See also: protocol

**service**
A program, routine, or process that performs a specific system function to support other programs, particularly at a low (close to the hardware) level. When services are provided over a network, they can be published in Active Directory, facilitating service-centric administration and usage. Some examples of services are the Security Accounts Manager service, File Replication service, and Routing and Remote Access service.

See also: Security Accounts Manager (SAM)

**sh**
See definition for: shell (sh)

**share**
To make resources, such as folders and printers, available to others.

See also: resource

**shared resource**
Any device, data, or program that is used by more than one program or one other device. For Windows, *shared resource* refers to any resource that is made available to network users, such as folders, files, printers, and named pipes. *Shared resource* can also refer to a resource on a server that is available to network users.

See also: resource; share

**shell (sh)**
A command-line POSIX tool that contains a large subset of Korn shell functionality (though no variable arrays).

**SID**
See definition for: security ID (SID)

**Simple Mail Transfer Protocol (SMTP)**
A member of the TCP/IP suite of protocols that governs the exchange of electronic mail between message transfer agents.

See also: protocol; Transmission Control Protocol/Internet Protocol (TCP/IP)

**Simple Network Management Protocol (SNMP)**
A network protocol used to manage TCP/IP networks. In Windows, the SNMP service is used to provide status information about a host on a TCP/IP network.

See also: protocol; service; Transmission Control Protocol/Internet Protocol (TCP/IP)

**simple volume**
A dynamic volume made up of disk space from a single dynamic disk. A simple volume can consist of a single region on a disk or multiple regions of the same disk that are linked together. If the simple volume is not a system volume or boot volume, you can extend it within the same disk or onto additional disks. If you extend a simple volume across multiple disks, it becomes a spanned volume. You can create simple volumes only on dynamic disks. Simple volumes are not fault tolerant, but you can mirror them to create mirrored volumes on computers running the Windows 2000 Server or Windows Server 2003 families of operating systems.

See also: dynamic disk; fault tolerance; spanned volume; volume

**single quorum device server cluster**
A cluster configuration that has two or more nodes and that is configured so that every node is attached to one or more cluster storage device. The cluster configuration data is stored on a single cluster storage device. There are advantages and limitations for each cluster configuration (single node server cluster, single quorum device server cluster, and majority node set server cluster).

See also: cluster; cluster storage; majority node set server cluster

**small computer system interface (SCSI)**
A standard high-speed parallel interface defined by the American National Standards Institute (ANSI). A SCSI interface is used for connecting microcomputers to peripheral devices, such as hard disks and printers, and to other computers and local area networks (LANs).

See also: local area network (LAN)

**SMB**
See definition for: Server Message Block (SMB)

**SMTP**
See definition for: Simple Mail Transfer Protocol (SMTP)

**snap-in**
A type of tool that you can add to a console supported by Microsoft Management Console (MMC). A stand-alone snap-in can

be added by itself; an extension snap-in can be added only to extend the function of another snap-in.

See also: Microsoft Management Console (MMC)

**spanned volume**
A dynamic volume consisting of disk space on more than one physical disk. You can increase the size of a spanned volume by extending it onto additional dynamic disks. You can create spanned volumes only on dynamic disks. Spanned volumes are not fault tolerant and cannot be mirrored.

See also: dynamic disk; fault tolerance; simple volume; volume

**SQL**
See definition for: structured query language (SQL)

**standard error (STDERR)**
In UNIX, the defined receiver of error messages about a process. By default, the standard error goes to the terminal.

**standard input (STDIN)**
In UNIX, the defined source of input for a process. By default, standard input comes from the terminal.

**standard output (STDOUT)**
In UNIX, the defined receiver for output from a process. By default, the standard output goes to the terminal.

**static routes**
Routes in a routing table that are permanent until changed by a network administrator or by an automatically scheduled auto-static update.

See also: router

**striped volume**
A dynamic volume that stores data in stripes on two or more physical disks. Data in a striped volume is allocated alternately and evenly (in stripes) across the disks. Striped volumes offer the best performance of all the volumes that are available in Windows, but they do not provide fault tolerance. If a disk in a striped volume fails, the data in the entire volume is lost. You can create striped volumes only on dynamic disks. Striped volumes cannot be mirrored or extended.

See also: dynamic disk; fault tolerance; volume

**structured query language (SQL)**
A widely accepted standard database sublanguage used in querying, updating, and managing relational databases.

**subnet mask**
A 32-bit value that enables the recipient of Internet Protocol version 4 (IPv4) packets to distinguish the network ID and host ID portions of the IPv4 address. Typically, subnet masks use the format 255.*x.x.x*. IPv6 uses network prefix notations rather than subnet masks.

See also: IP address

**system tray**
See definition for: notification area

**T**

**TAPI**
See definition for: Telephony API (TAPI)

**TCP**
See definition for: Transmission Control Protocol/Internet Protocol (TCP/IP)

**TCP/IP**
See definition for: Transmission Control Protocol/Internet Protocol (TCP/IP)

**TDI**
See definition for: Transport Driver Interface (TDI)

**Telephony API (TAPI)**
An application programming interface (API) used by communications programs to work with telephony and network services. Communications programs like HyperTerminal and Phone Dialer use TAPI to dial, answer, and route telephone calls on

conventional telephony devices, including PBXs, modems, and fax machines. TAPI 3.0 also provides Internet Protocol (IP) telephony support, which Phone Dialer and other programs use to transmit, route, and control real-time audio and video signals over IP-based networks such as the Internet.

See also: application programming interface (API); service

**Telnet**
A protocol that enables an Internet user to log on to and enter commands on a remote computer linked to the Internet, as if the user were using a text-based terminal directly attached to that computer. Telnet is part of the TCP/IP suite of protocols. The term *telnet* also refers to the software (client or server component) that implements this protocol.

See also: protocol; Transmission Control Protocol/Internet Protocol (TCP/IP)

**TGS**
See definition for: ticket-granting service (TGS)

**TGT**
See definition for: ticket-granting ticket (TGT)

**thread**
A type of object within a process that runs program instructions. Using multiple threads allows concurrent operations within a process and enables one process to run different parts of its program on different processors simultaneously. A thread has its own set of registers, its own kernel stack, a thread environment block, and a user stack in the address space of its process.

See also: kernel

**ticket**
A set of identification data for a security principal, issued by a domain controller for purposes of user authentication. Two forms of tickets in Windows are ticket-granting tickets (TGTs) and service tickets.

See also: domain controller; ticket-granting ticket (TGT)

**ticket-granting service (TGS)**
A Kerberos V5 service provided by the Kerberos V5 Key Distribution Center (KDC) service that issues service tickets that allow users to authenticate to services in a domain.

See also: Kerberos V5 authentication protocol; Key Distribution Center (KDC); ticket-granting ticket (TGT)

**ticket-granting ticket (TGT)**
A credential issued to a user by the Kerberos Key Distribution Center (KDC) when the user logs on. The user must present the TGT to the KDC when requesting session tickets for services. Because a TGT is normally valid for the life of the user's logon session, it is sometimes called a *user ticket*.

See also: Kerberos V5 authentication protocol; Key Distribution Center (KDC); ticket-granting service (TGS)

**topology**
The physical layout of computers, cables, switches, routers, and other components of a network. *Topology* also refers to the underlying network architecture, such as Ethernet or Token Ring.

In Active Directory replication, the set of connections that domain controllers use to replicate information among themselves.

See also: domain controller

**touch**
Command-line POSIX tool which changes the modification date of files.

**Transmission Control Protocol/Internet Protocol (TCP/IP)**
A set of networking protocols widely used on the Internet that provides communications across interconnected networks of computers with diverse hardware architectures and various operating systems. TCP/IP includes standards for how computers communicate and conventions for connecting networks and routing traffic.

See also: protocol

**Transport Driver Interface (TDI)**
A common set of routines for network layer components that communicate with the session layer of the Open Systems Interconnection (OSI) model. These routines allow software components above and below the transport layer to be mixed and matched without reprogramming.

**transport protocol**
A protocol that defines how data should be presented to the next receiving layer in the Windows NT and Windows 2000 networking model and packages the data accordingly. The transport protocol passes data to the network adapter driver through the Network Driver Interface Specification (NDIS) interface and to the redirector through the Transport Driver Interface (TDI).

**trust relationship**
A logical relationship established between domains to allow pass-through authentication, in which a trusting domain honors the logon authentications of a trusted domain. User accounts and global groups defined in a trusted domain can be given rights and permissions in a trusting domain, even though the user accounts or groups don't exist in the trusting domain's directory.

See also: domain; permission

**tunnel server**
A server or router that terminates tunnels and forwards traffic to the hosts on the target network.

See also: router

**two-way trust**
A trust relationship between two domains in which both domains trust each other. For example, domain A trusts domain B, and domain B trusts domain A. All parent-child trusts are two-way.

See also: domain; one-way trust; trust relationship

**U**

**UCS**
See definition for: Universal Character Set (UCS)

**UNC**
See definition for: Universal Naming Convention (UNC)

**UNC (Universal Naming Convention) name**
The full name of a resource on a network. It conforms to the *servername**sharename* syntax, where *servername* is the name of the server and *sharename* is the name of the shared resource. UNC names of directories or files can also include the directory path under the share name, with the following syntax:

\\servername\sharename\directory\filename

See also: resource

**Unicode**
A character encoding standard developed by the Unicode Consortium that represents almost all of the written languages of the world. The Unicode character repertoire has multiple representation forms, including UTF-8, UTF-8, and UTF-32. Most Windows interfaces use the UTF-16 form.

See also: American Standard Code for Information Interchange (ASCII); Universal Character Set (UCS)

**Uniform Resource Locator (URL)**
An address that uniquely identifies a location on the Internet. A URL for a World Wide Web site is preceded by *http://*, as in the fictitious URL *http://www.example.microsoft.com*. A URL can contain more detail, such as the name of a page of hypertext, usually identified by the file name extension .html or .htm.

**uninterruptible power supply (UPS)**
A device that connects a computer and a power source to ensure that electrical flow is not interrupted. UPS devices use batteries to keep the computer running for a period of time after a power failure. UPS devices usually provide protection against power surges and brownouts as well.

**Universal Character Set (UCS)**
An international standard character set reference that is part of the Unicode standard. The most widely held existing version of the UCS standard is UCS-2, which specifies 16-bit character values currently accepted and recognized for use to encode most of the world's languages.

See also: American Standard Code for Information Interchange (ASCII); Unicode

**Universal Naming Convention (UNC)**

A convention for naming files and other resources beginning with two backslashes (\), indicating that the resource exists on a network computer. UNC names conform to the *servername**sharename* syntax, where *servername* is the server's name and *sharename* is the name of the shared resource. The UNC name of a directory or file can also include the directory path after the share name, by using the following syntax: *servername**sharename**directory**filename*.

**universal serial bus (USB)**
An external bus that supports Plug and Play installation. Using USB, you can connect and disconnect devices without shutting down or restarting your computer. You can use a single USB port to connect up to 127 peripheral devices, including speakers, telephones, CD-ROM drives, joysticks, tape drives, keyboards, scanners, and cameras. A USB port is usually located on the back of your computer near the serial port or parallel port.

**UNIX**
A powerful, multiuser, multitasking operating system initially developed at AT&T Bell Laboratories in 1969 for use on minicomputers. UNIX is considered more portable, that is, less computer-specific, than other operating systems because it is written in C language. Newer versions of UNIX have been developed at the University of California at Berkeley and by AT&T.

**UPS**
See definition for: uninterruptible power supply (UPS)

**URL**
See definition for: Uniform Resource Locator (URL)

**USB**
See definition for: universal serial bus (USB)

**user profile**
A file that contains configuration information for a specific user, such as desktop settings, persistent network connections, and application settings. Each user's preferences are saved to a user profile that Windows uses to configure the desktop each time a user logs on.

**UUEncoding/Decoding**
A form of encoding and decoding that allows binary files to be sent as e-mail. A binary file, such as a program or image, cannot be sent directly over the Internet by e-mail. To send such a file, you must first encode it to hide the control characters. This encoding can be accomplished by converting the binary file to a text file. UUEncoding carries out this conversion, and UUDecoding regenerates the original file.

**V**

**variable**
In programming, a named storage location capable of containing a certain type of data that can be modified during program execution.

System environment variables are defined by the operating systems in the Windows Server 2003 family, and they are the same no matter who is logged on to the computer. Administrator group members can add new variables or change the values, however.

User environment variables can be different for each user of a particular computer. They include any environment variables you want to define or variables defined by your applications, such as the path where application files are located.

See also: environment variable

**virtual local area network (VLAN)**
A logical grouping of hosts on one or more local area networks (LANs) that allows communication to occur between hosts as if they were on the same physical LAN.

See also: local area network (LAN)

**VLAN**
See definition for: virtual local area network (VLAN)

**VoIP (Voice over Internet Protocol)**
A method for sending voice over a local area network (LAN), a wide area network (WAN), or the Internet using TCP/IP packets.

See also: local area network (LAN); Transmission Control Protocol/Internet Protocol (TCP/IP); wide area network (WAN)

**volume**

An area of storage on a hard disk. A volume is formatted by using a file system, such as file allocation table (FAT) or NTFS, and has a drive letter assigned to it. You can view the contents of a volume by clicking its icon in Windows Explorer or in My Computer. A single hard disk can have multiple volumes, and volumes can also span multiple disks.

See also: file allocation table (FAT); NTFS file system; simple volume; spanned volume

**volume set**
A volume that consists of disk space on one or more physical disks. A volume set is created by using basic disks and is supported only in Windows NT 4.0 or earlier. Volume sets were replaced by spanned volumes, which use dynamic disks.

See also: basic disk; dynamic disk; partition; spanned volume; volume

**voluntary tunnel**
A tunnel that is initiated by the client. A voluntary tunnel tunnels Point-to-Point Protocol (PPP) over Internet Protocol (IP) from the client to the tunnel server, and then the data is forwarded to the target host by the tunnel server.

See also: tunnel server

**W**

**WAN**
See definition for: wide area network (WAN)

**wc**
Word count. Command-line POSIX tool that returns the number of bytes, words, and lines in files.

**wide area network (WAN)**
A communications network connecting geographically separated computers, printers, and other devices. A WAN enables any connected device to interact with any other on the network.

See also: local area network (LAN)

**Windows Internet Name Service (WINS)**
A software service that dynamically maps IP addresses to computer names (network basic input/output system (NetBIOS) names). This enables users to access resources by name instead of requiring them to use IP addresses that are difficult to recognize and remember.

See also: IP address; network basic input/output system (NetBIOS); resource; service; WINS resource

**Windows Support Tools**
Tools that administrators, developers, or support personnel can use to diagnose and troubleshoot operating system configuration problems. Although Windows Support Tools are included on the operating system CD, they are not guaranteed or supported by Microsoft, and they must be installed separately from the operating system.

**WINS**
See definition for: Windows Internet Name Service (WINS)

**WINS resource**
A resource type that provides Windows Internet Name Service (WINS) from a cluster.

See also: cluster; Windows Internet Name Service (WINS)

**working set**
For a process, the amount of physical memory assigned to the process by the operating system.

**X**

**X.500**
A set of standards defining a distributed directory service, developed by the International Standards Organization (ISO).

See also: directory service

**Y**
There are no glossary terms that begin with this letter.

**Z**
There are no glossary terms that begin with this letter.

# Security Policy Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Policy Settings**

**In this subject**


- Account Policy Settings
- Local Policy Settings
- Event Log Policy Settings
- Restricted Groups Policy Settings
- System Services
- Registry
- File System
- Wireless Network (IEEE 802.11) Policies
- Public Key Policy Settings
- Software Restriction Policy Settings
- IP Security Policies on Active Directory


Security policy settings are among the settings contained in Group Policy objects (GPOs). When you edit security settings in a local GPO, only the security policy settings on that computer are affected. When you edit security policy settings in a GPO in Active Directory directory service, the policy settings affect sites, domains, and organizational units (OUs) to which the GPO is linked. Some settings, like password policy settings, are only operative at the domain level.

This subject documents security policy settings in Microsoft Windows Server 2003.

# Account Policy Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Account Policies**

**In this section**

- Password Policy
- Account Lockout Policy
- Kerberos Policy

The policy settings under Account Policies are implemented at the domain level. A Windows Server 2003 domain must have a single password policy, account lockout policy, and Kerberos version 5 authentication protocol policy for the domain. Configuring these policy settings at any other level in Active Directory will only affect local accounts on member servers. If there are groups that require separate password policies, they should be segmented into another domain or forest, based on any additional requirements.

For domain accounts, there can be only one account policy per domain. The account policy must be defined in the Default Domain Policy or in a new policy that is linked to the root of the domain and given precedence over the Default Domain Policy, which is enforced by the domain controllers that make up the domain. A domain controller always pulls the account policy from a Group Policy object (GPO)linked to the domain, which by default is the Default Domain Policy GPO. This behavior occurs even if there is a different account policy applied to the organizational unit (OU) that contains the domain controller.

By default, workstations and servers joined to a domain — the domain member computers — also receive the same account policy for their local accounts. However, local account policies for member computers can be differentiated from the domain account policy by defining an account policy for the OU that contains the member computers.

Account Policies contains three subsets:

- **Password Policy**. These policy settings are used for domain or local user accounts. They determine settings for passwords, such as enforcement and lifetimes.
- **Account Lockout Policy**. These policy settings are used for domain or local user accounts. They determine the circumstances and length of time that an account will be locked out of the system.
- **Kerberos Policy**. These policy settings are used for domain user accounts. They determine Kerberos-related settings, such as ticket lifetimes and enforcement. Kerberos policy settings do not exist in local computer policy.

Account Policies policy settings can be configured in the following location in Group Policy Object Editor:

Computer Configuration\Windows Settings\Security Settings\Account Policies\

**Password Policy**

In Windows and many other operating systems, the most common method for authenticating a user's identity is to use a secret passphrase or password. Securing your network environment requires that strong passwords be used by all users. This helps avoid the threat of a malicious user guessing a weak password, whether through manual methods or by using tools, to acquire the credentials of a compromised user account. This is especially true for administrative accounts. When you change a complex password regularly, it reduces the likelihood of a successful password attack. Password policy settings control the complexity and lifetime of passwords. Each specific password policy account setting is discussed in this section. Password Policy settings can be configured in the following location in Group Policy Object Editor:

Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

**Enforce password history**

The **Enforce password history** policy setting determines the number of unique new passwords that must be associated with a user account before an old password can be reused.

The possible values for this Group Policy setting are:

- A user-defined number from 0 through 24.
- Not defined.

**Discussion**

Password reuse is an important concern in any organization. Many users want to reuse the same password for their account over a long period of time. The longer the same password is used for a particular account, the greater the chance that an attacker will be able to determine the password through brute force attacks. If users are required to change their password, but nothing prevents them from using the old password or continually reusing a small number of passwords, the effectiveness of a good password policy is greatly reduced.

Specifying a low number for **Enforce password history** allows users to continually use the same small number of passwords repeatedly. If you do not also set **Minimum password age**, users can change their password as many times in a row as necessary in order to reuse their original password.

If you set **Enforce password history** to a number greater than zero, users must come up with a new password every time they are required to change their old one. This improves security, but it can increase the risk that users will write down their passwords so they do not forget them.

If you set the value to the maximum of **24**, it helps to ensure that vulnerabilities caused by password reuse are kept to a minimum.

For this policy setting to be effective in your organization, configure **Minimum password age** so that you do not allow passwords to be changed immediately. **Enforce password history** should be set at the level that combines a reasonable maximum password age with a reasonable password change interval requirement for users.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | 24 passwords remembered |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 0 passwords remembered |
| DC Effective Default Settings | 24 passwords remembered |
| Member Server Effective Default Settings | 24 passwords remembered |

**Maximum password age**

The **Maximum password age** policy setting determines the number of days that a password can be used before the system requires the user to change it.

The possible values for this Group Policy setting are:

- A user-defined number of days from 0 through 999.
- Not defined.

**Discussion**

If your users change their passwords frequently, it helps to reduce the risk that a valid password might be cracked, and it mitigates the risk that someone may use a password that has been wrongfully acquired. **Maximum password age** can be configured so that users are never required to change their passwords, but this results in a major security risk.

Many cautious administrators choose to set **Maximum password age** to a value between **30** and **60** days. You can set **Maximum password age** to never expire by setting the number of days to **0**.

Setting **Maximum password age** very low requires users to change their passwords too often. This might actually reduce

security in the organization, because it can increase the possibility of that users will write their passwords down to avoid forgetting them. If you set the value very high, it reduces the level of security within an organization, because it gives a potential attacker more time to guess a user's password.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | 42 days |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 42 days |
| DC Effective Default Settings | 42 days |
| Member Server Effective Default Settings | 42 days |

**Minimum password age**
The **Minimum password age** policy setting determines the number of days that a password must be used before the user is allowed to change it. **Minimum password age** must be less than **Maximum password age**.

Configure this to be greater than **0** if you want the **Enforce password history** policy setting to be effective. If **Enforce password history** is set to **0**, the user does not have to choose a new password. If **Enforce password history** is set to a number greater than zero, users must enter a new unique password when they change their password.

The possible values for this Group Policy setting are:

- A user-defined number of days from 0 through 999.
- Not defined.

**Discussion**
Forcing users to change passwords regularly is ineffective if they can cycle through passwords until they get to an old favorite. Use **Minimum password age** with **Enforce password history** to prevent this. For example, if **Enforce password history** is set to 12 (which ensures that users cannot reuse any of their last 12 passwords) and **Minimum password age** is 0, users could change their password 13 times in a row to return to the password they were originally using. Therefore, **Minimum password age** must be set to a value greater than 0 for **Enforce password history** to be effective.

It is advisable to set **Minimum password age** to a value of 2 days. Setting the number of days to 0 allows immediate password changes, which is not recommended.

There is one minor issue associated with setting **Minimum password age** to a number greater than 0. If an administrator sets a password for a user, and then would like for that user to change the administrator-defined password, the administrator must select the **User must change password at next logon** check box. Otherwise, the user will not be able to change the password until the number of days specified by **Minimum password age**.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | 1 day |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 0 days |
| DC Effective Default Settings | 1 day |

| | |
|---|---|
| Member Server Effective Default Settings | 1 day |

**Minimum password length**

The **Minimum password length** policy setting determines the least number of characters that can make up a password for a user account. There are many different theories about determining the best password length for an organization. One possibility is to use *pass phrases* rather than passwords. In Microsoft Windows 2000, Windows XP, and Windows Server 2003, pass phrases can be quite long and they can include spaces. A valid pass phrase such as "I want to drink a $5 milkshake" is considerably stronger than a ten-character string of random numbers and letters, yet is easier to remember.

The possible values for this Group Policy setting are:

- A user-defined number of characters from 0 through 14.
- Not defined.

**Discussion**

There are several types of password attacks that can be performed to obtain the password for a particular user account. These include dictionary attacks — which attempt to use common words and phrases — and brute-force attacks, which try every possible combination of characters. In addition, an attack can be performed by obtaining the account database and using tools to crack the accounts and passwords.

Set **Minimum password length** to at least a value of 8. If the number of characters is set to 0, no password is required. In most environments, an eight-character password is recommended because it is long enough to provide adequate security and still short enough for users to easily remember. This value will help to provide adequate defense against a brute-force attack. Adding complexity requirements will help reduce the possibility of a dictionary attack. Complexity requirements are discussed in the description of the **Passwords must meet complexity requirements** policy setting.

Permitting short passwords reduces security, because short passwords can be easily broken with tools that perform either dictionary or brute-force attacks against the passwords. Requiring very long passwords can result in mistyped passwords that might cause an account lockout and subsequently increase the volume of Help desk calls.

In addition, requiring extremely long passwords can actually decrease the security of an organization because users might be more likely to write their passwords down to avoid forgetting them. However, if users are taught they can use pass phrases as described above, they should be much more likely to remember.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | 7 characters |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 0 characters |
| DC Effective Default Settings | 7 characters |
| Member Server Effective Default Settings | 7 characters |

**Passwords must meet complexity requirements**

The **Passwords must meet complexity requirements** policy setting determines whether passwords must meet a series of guidelines that are considered important for a strong password. Enabling this policy setting requires passwords to meet the following requirements:

- The password is at least six characters long.
- The password contains characters from three of the following four categories:
    - English uppercase characters (from A through Z)
    - English lowercase characters (from a through z)

    - Base 10 digits (from 0 through 9)
    - Non-alphanumeric characters (for example: !, $, #, or %)

- The password does not contain three or more characters from the user's account name. If the account name is less than three characters long, this check is not performed because the rate at which passwords would be rejected would be too high. When checking against the user's full name, several characters are treated as delimiters that separate the name into individual tokens: commas, periods, dashes, hyphens, underscores, spaces, number signs (#), and tab characters. Each token that is three or more characters long is searched for in the password, and if it is present, the password change is rejected. For example, the name "Erin M. Hagens" would be split into three tokens: "Erin," "M," and "Hagens." Because the second token is only one character long, it would be ignored. Therefore this user could not have a password that included either "erin" or "hagens" as a substring anywhere in the password. None of these checks are case-sensitive.

These complexity requirements are enforced when passwords are changed or new passwords created.

The rules that are included in the Windows Server 2003 password complexity requirements are part of Passfilt.dll and cannot be directly modified. The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Passwords that contain only alphanumeric characters are easy to crack by using publicly available tools. To prevent this, passwords should contain additional characters and meet complexity requirements.

Set **Passwords must meet complexity requirements** to **Enabled**. This policy setting, combined with a minimum password length of 8, ensures that there are at least 218,340,105,584,896 different possibilities for a single password. This makes the use of a brute-force attack difficult, but still not impossible. An attacker who had enough processing power to test 1 million passwords per second could determine such a password in about seven-and-a-half days or less.

Enabling the default Passfilt.dll may cause some additional Help desk calls for locked-out accounts because users might not be used to having passwords that contain characters other than those found in the alphabet. However, this policy setting is liberal enough that all users should be able to abide by the requirements with a minor learning curve.

Additional settings that can be included in a custom Passfilt.dll are the use of non–upper-row characters. Upper-row characters are those that are typed by holding down the SHIFT key and typing any of the digits from 1 through 10.

Also, the use of ALT key character combinations can greatly enhance the complexity of a password. However, requiring all users in an organization to adhere to such stringent password requirements can result in unhappy users and an extremely busy Help desk. Consider implementing a requirement in your organization to use ALT characters in the range from 0128 through 0159 as part of all administrator passwords. (ALT characters outside of this range can represent standard alphanumeric characters that do not add additional complexity to the password.)

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Enabled |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Enabled |

| | |
|---|---|
| Member Server Effective Default Settings | Enabled |

**Store password using reversible encryption for all users in the domain**

The **Store password using reversible encryption for all users in the domain** policy setting determines whether Windows Server 2003, Windows 2000, and Windows XP Professional store passwords by using reversible encryption.

This policy setting provides support for applications that use protocols that require knowledge of the user's password for authentication purposes. By definition, storing encrypted passwords in a way that is reversible means that the encrypted passwords can be decrypted. A knowledgeable attacker who is able to break this encryption can then log on to network resources by using the compromised account. For this reason, never enable **Store password using reversible encryption for all users in the domain** unless application requirements outweigh the need to protect password information.

If you use Challenge Handshake Authentication Protocol (CHAP) authentication through remote access or Internet Authentication Service (IAS) services, you must enable this policy setting. CHAP is an authentication protocol used by Microsoft remote access and Network Connections. Digest Authentication in Internet Information Services (IIS) also requires that you enable this policy setting.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

This policy setting determines whether Windows Server 2003 will store passwords in a weaker format that is much more susceptible to brute force attacks.

It is advisable to set the value for **Store password using reversible encryption for all users in the domain** to **Disabled**. If you use either Digest Authentication in IIS or the CHAP authentication protocol through remote access or IAS, you must set this value to **Enabled**. This is an extremely dangerous policy setting to apply by using Group Policy on a user-by-user basis, because it requires opening the appropriate user account object in Active Directory Users and Computers.

### Note

- *Never* enable this policy setting unless business requirements outweigh the need to protect password information.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Disabled |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Account Lockout Policy**

Someone who attempts to use more than a few unsuccessful passwords while trying to log on to your system might be a malicious user attempting to determine an account password by trial and error. Windows Server 2003 keeps track of logon attempts, and it can be configured to respond to this type of potential attack by disabling the account for a preset period of time. Account Lockout Policy settings control the threshold for this response and the actions to be taken after the threshold is reached. The Account Lockout Policy settings can be configured in the following location in Group Policy Object Editor:

Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy\

**Account lockout duration**

The **Account lockout duration** policy setting determines the number of minutes a locked-out account remains locked out before automatically becoming unlocked. The available range is from 1 through 99,999 minutes. A value of 0 specifies that the account will be locked out until an administrator explicitly unlocks it. If **Account lockout threshol**d is set to a number greater than zero, **Account lockout duration** must be greater than or equal to the value of **Reset account lockout counter after**.

The possible values for this Group Policy setting are:

- A user-defined number of minutes from 0 through 99,999.
- Not defined.

**Discussion**

A denial-of-service attack can occur if an attacker abuses the **Account lockout threshold** and repeatedly attempts to log on to an account. If the **Account lockout threshold** is configured, after the specified number of failed attempts the account will be locked out. If the **Account lockout duration** is set to 0, the account will remain locked out until an administrator unlocks it manually.

It is advisable to set **Account lockout duration** to approximately 30 minutes. To specify that the account will never be locked out, set the value to 0. To configure the value for this policy setting so that it never automatically unlocks the account might seem like a good idea, doing so can increase the number of requests that your organization's Help desk receives to unlock accounts that were locked by mistake.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not applicable |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Account lockout threshold**

The **Account lockout threshold** policy setting determines the number of failed logon attempts that will cause a user account to be locked out. A locked-out account cannot be used until it is reset by an administrator or until the number of minutes specified by **Account lockout duration** expires. You can set a value from 1 through 999 failed logon attempts, or you can specify that the account will never be locked out by setting the value to 0. If **Account lockout threshold** is set to a number greater than zero, **Account lockout duration** must be greater than or equal to the value of **Reset account lockout counter after**.

Failed password attempts on workstations or member servers that have been locked by using either CTRL+ALT+DELETE or password-protected screen savers do not count as failed logon attempts unless **Interactive logon: Require Domain Controller authentication to unlock workstation** is set to **Enabled**. If **Interactive logon: Require Domain Controller authentication to unlock workstation** is enabled, repeated failed password attempts to unlock the workstation will count against the account lockout threshold.

The possible values for this Group Policy setting are:

- A user-defined number from 0 through 999.
- Not defined.

**Discussion**

Brute force password attacks can be automated to try thousands or even millions of password combinations for any or all user accounts. Limiting the number of failed logons that can be performed nearly eliminates the effectiveness of such attacks.

However, it is important to note that, in contrast, a denial-of-service attack could be performed on a domain that has an account lockout threshold configured. A malicious user could programmatically attempt a series of password attacks against all users in the organization. If the number of attempts is greater than the value of **Account lockout threshold**, the attacker could potentially lock out every account.

Because vulnerabilities can exist both when this value is configured and when it is not, any organization should weigh their identified threats and the risks that they are trying to mitigate. There are two options to consider for this policy setting:

- Set **Account lockout threshold** to 0. This ensures that accounts will not be locked out. This setting will prevent a denial-of-service attack that intentionally locks out all or some accounts. In addition, this setting helps reduce Help desk calls because users cannot accidentally lock themselves out of their accounts.

  Because it will not prevent a brute force attack, a value of 0 should only be chosen if both of the following criteria are explicitly met:

  - Password Policy settings force all users to have complex passwords made up of eight or more characters.
  - A robust auditing mechanism is in place to alert administrators when a series of failed logons are occurring in the environment.

  If these criteria cannot be met, set **Account lockout threshold** to a high enough value that users can accidentally mistype their password several times before they are locked out of their account, but ensure that a brute-force password attack would still lock out the account. It is advisable to specify a value of 50 invalid logon attempts. Keep in mind, however, that although this setting can reduce the number of Help desk calls by reducing the number of user lockouts, it cannot prevent a denial-of-service attack.

Setting **Account lockout threshold** to a number greater than zero locks out user accounts until they are reset by an administrator, or until the number of minutes specified in **Account lockout duration** expires. This setting will likely generate a number of additional Help desk calls. In fact, in many organizations, locked accounts generate the most calls to the Help desk.

By setting **Account lockout threshold** to 0, you leave your organization open to the possibility that you will not detect a brute force attack.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | 0 invalid login attempts |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 0 invalid login attempts |
| DC Effective Default Settings | 0 invalid login attempts |
| Member Server Effective Default Settings | 0 invalid login attempts |

**Reset account lockout counter after**
The **Reset account lockout counter after** policy setting determines the number of minutes that must elapse from the time a user fails to log on before the failed logon attempt counter is reset to 0 bad logon attempts. If **Account lockout threshold** is set to a number greater than zero, this reset time must be less than or equal to the value of **Account lockout duration**.

The possible values for this Group Policy setting are:

- A user-defined number of minutes from 1 through 99,999.
- Not defined.

**Discussion**

A disadvantage to setting this too high is that users lock themselves out for an inconveniently long period, if they exceed the account lockout threshold through login errors. Users may make excessive help desk calls.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not applicable |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Kerberos Policy**
In Windows Server 2003, the Kerberos version 5 authentication protocol provides the default mechanism for authentication services and the authorization data necessary for a user to access a resource and perform a task on that resource. By reducing the lifetime of Kerberos tickets, you reduce the risk of a legitimate user's credentials' being stolen and successfully used by an attacker; however, you increase authorization overhead. In most environments, these settings should not need to be changed. In a default installation of a Windows 2000 or Windows Server 2003 Active Directory domain, the default values of these policy settings, which are applied at the domain level, are configured in the Default Domain Policy Group Policy object (GPO). The Kerberos Policy settings can be configured in the following location in Group Policy Object Editor:

Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

**Enforce user logon restrictions**
The **Enforce user logon restrictions** policy setting determines whether the Kerberos V5 Key Distribution Center (KDC) validates every request for a session ticket against the user rights policy of the user account. Validating each request for a session ticket is optional because the extra step takes time, and that can slow network access to services.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
If this policy setting is disabled, users might be granted session tickets for services they do not have the right to use.

It is advisable to set **Enforce user logon restrictions** to **Enabled**.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Enabled |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not applicable |
| DC Effective Default Settings | Enabled |

| | |
|---|---|
| Member Server Effective Default Settings | Not applicable |

**Maximum lifetime for service ticket**

The **Maximum lifetime for service ticket** policy setting determines the maximum number of minutes that a granted session ticket can be used to access a particular service. The value must be 10 minutes or greater, and it must be less than or equal to the value of the **Maximum lifetime for user ticket** policy setting.

The possible values for this Group Policy setting are:

- A user-defined number of minutes from 10 through 99,999, or 0, in which case service tickets do not expire.
- Not defined.

**Discussion**

If a client presents an expired session ticket when it requests a connection to a server, the server returns an error message. The client must request a new session ticket from the Kerberos V5 KDC. After a connection is authenticated, however, it no longer matters whether the session ticket remains valid. Session tickets are used only to authenticate new connections with servers. Ongoing operations are not interrupted if the session ticket that authenticated the connection expires during the connection.

If the value for this policy setting is too high, users might be able to access network resources outside of their logon hours, or users whose accounts have been disabled might be able to continue accessing network services by using valid service tickets that were issued before their account was disabled. If the value is set to 0, service tickets never expire.

It is advisable to set **Maximum lifetime for service ticket** to **600** minutes.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | 600 minutes |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not applicable |
| DC Effective Default Settings | 600 minutes |
| Member Server Effective Default Settings | Not applicable |

**Maximum lifetime for user ticket**

The **Maximum lifetime for user ticket** policy setting determines the maximum amount of time (in hours) that a user's ticket-granting ticket can be used. When a user's ticket-granting ticket expires, a new one must be requested or the existing one must be renewed.

The possible values for this Group Policy setting are:

- A user-defined number of hours from 0 through 99,999.
- Not defined.

**Discussion**

If the value for this policy setting is too high, users might be able to access network resources outside of their logon hours, or users whose accounts have been disabled might be able to continue to access network services by using valid service tickets that were issued before their account was disabled. If the value is set to 0, ticket-granting tickets never expire.

It is advisable to set **Maximum lifetime for user ticket** to 10 hours.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | 10 hours |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not applicable |
| DC Effective Default Settings | 10 hours |
| Member Server Effective Default Settings | Not applicable |

**Maximum lifetime for user ticket renewal**

The **Maximum lifetime for user ticket renewal** policy setting determines the period of time (in days) during which a user's ticket-granting ticket can be renewed.

The possible values for this Group Policy setting are:

- A user-defined number of days from 0 through 99,999.
- Not defined.

**Discussion**

If the value for this policy setting is too high, users may be able to renew very old user ticket-granting tickets. If the value is 0, ticket-granting tickets never expire.

It is advisable to set **Maximum lifetime for user ticket renewal** to **7** days.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | 7 days |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not applicable |
| DC Effective Default Settings | 7 days |
| Member Server Effective Default Settings | Not applicable |

**Maximum tolerance for computer clock synchronization**

The **Maximum tolerance for computer clock synchronization** policy setting determines the maximum time difference (in minutes) that Kerberos V5 tolerates between the time on the client clock and the time on the domain controller running Windows Server 2003 that provides Kerberos authentication.

The possible values for this Group Policy setting are:

- A user-defined number of minutes from 1 through 99,999.
- Not defined.

**Discussion**

To prevent replay attacks, Kerberos V5 uses time stamps as part of its protocol definition. For time stamps to work properly, the clocks of the client and the domain controller need to be synchronized as closely as possible. Because the clocks of two computers are often out of sync, administrators can use this policy setting to establish the maximum acceptable difference to Kerberos V5 between a client clock and domain controller clock. If the difference between a client clock and the domain controller clock is less than **Maximum tolerance for computer clock synchronization**, any time stamp that is used in a session between the two computers is considered authentic.

It is advisable to set **Maximum tolerance for computer clock synchronization** to a value of 5 minutes.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | 5 minutes |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not applicable |
| DC Effective Default Settings | 5 minutes |
| Member Server Effective Default Settings | Not applicable |

# Local Policy Settings

Updated: October 8, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local Policies**

**In this section**

- Audit Policy
- User Rights Assignment
- Security Options

These policy settings apply to a computer and contain these subsets:

- **Audit policy**. Determines whether security events are written to the security log in Event Viewer on the computer. Also determines whether to log successful attempts, failed attempts, or both.
- **User rights assignment**. Determines which users or groups have logon rights or privileges on the computer.
- **Security options**. Enables or disables security policy settings for the computer, such as digital signing of data, Administrator and Guest account names, floppy disk drive and CD drive access, driver installation, and logon prompts.

Because a computer can have more than one policy setting applied to it, security policy settings can conflict with each other. The order of precedence from highest to lowest is: OU, domain, and local computer.

**Audit Policy**

An audit log records an entry whenever users perform certain actions that you specify. For example, the modification of a file or policy setting can trigger an audit entry. The audit entry shows the action performed, the associated user account, and the date and time of the action. You can audit both successful and failed attempts at actions.

The state of the operating system and applications on a computer is dynamic. For example, an administrator might need to change security levels temporarily to resolve an administrative or network issue, but all too often, such changes are forgotten about and never undone. In a case like this, the computer might no longer meet the requirements for enterprise security.

Regular analysis enables an administrator to track and ensure an adequate level of security on each computer as part of an enterprise risk-management program. Analysis focuses on highly specific information about all system aspects related to security. This enables an administrator to tune the security levels and—most importantly—detect any security flaws that might occur in the system over time.

Security auditing is extremely important for any enterprise system, because audit logs sometimes give the only indication that a security breach has occurred. If the breach is discovered some other way, proper audit policy settings generate an audit log that contains important information about the breach.

Often, failure logs are much more informative than success logs, because failures typically indicate an error. For example, if a user successfully logs on to the system, this is considered normal. However, if a user unsuccessfully tries to log on to the system multiple times, it can indicate that someone is trying to break into the system by using someone else's user ID.

All computers in your organization should have sensible auditing policy settings so that legitimate users can be held accountable for their actions and unauthorized activity can be detected and tracked. If no auditing is configured, or if the auditing is set too low on the computers in your organization, you will not have sufficient evidence to analyze after security incidents take place. On the other hand, if too much auditing is enabled, the security log will fill up with meaningless entries. The Audit Policy settings can be configured in the following location in Group Policy Object Editor:

Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy

**Audit account logon events**

The **Audit account logon events** policy setting determines whether to audit each instance of a user logging on to or logging off from another computer, in which the computer recording the audit event is used to validate the account. If you configure

this policy setting, you can specify whether to audit successful logons, audit failed logons, or not audit logon events at all.

The possible values for this Group Policy setting are:

- Success.
- Failure.
- No auditing.

**Discussion**

If **Audit account logon events** is set to **Success**, an audit entry is recorded in the security log whenever an account logon attempt succeeds. This can be useful information for accounting purposes and for post-incident forensics when you want to determine who successfully logged on to which computer. If **Audit account logon events** is set to **Failure**, an audit entry is recorded whenever an account logon attempt fails. This can be useful for intrusion detection, but this setting can also leave your enterprise vulnerable to a denial-of-service condition, because a malicious user could generate millions of logon failures and fill your security log.

If the value of this policy setting on a domain controller is **Success**, an entry is recorded for each user who is validated against that domain controller, even though the user is actually logging on to a workstation that is joined to the domain.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Success |
| Stand-Alone Server Default Settings | Success |
| DC Effective Default Settings | Success |
| Member Server Effective Default Settings | Success |

**Audit account management**

The **Audit account management** policy setting determines whether to audit each event of account management on a computer.

The possible values for this Group Policy setting are:

- Success.
- Failure.
- No auditing.

**Discussion**

Examples of account management events include the following:

- A user account or group is created, changed, or deleted.
- A user account is renamed, disabled, or enabled.
- A password is set or changed.

If you configure this policy setting, you can specify whether to audit successes, audit failures, or not audit account management events at all. When you set the value of **Audit account management** to **Success**, an audit entry is generated when any

account management event succeeds. It is advisable to set the value to **Success** on all computers in your enterprise. To respond to security incidents, it is critical that organizations be able to track who created, changed, or deleted an account. When the value is set to **Failure**, an audit entry is generated when any account management event fails.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Success |
| Stand-Alone Server Default Settings | No auditing |
| DC Effective Default Settings | Success |
| Member Server Effective Default Settings | No auditing |

**Audit directory service access**
The **Audit directory service access** policy setting determines whether to audit the event of a user accessing an Active Directory object that has specified its own system access control list (SACL). An SACL is a list of users and groups for which actions on an object are to be audited on a Windows 2000–based network. If you configure this policy setting, you can specify whether to audit successes, audit failures, or not audit directory service access at all. Success audits generate an audit entry when a user successfully accesses an Active Directory object that has an SACL specified. Failure audits generate an audit entry when a user unsuccessfully attempts to access an Active Directory object that has a SACL specified.

The possible values for this Group Policy setting are:

- Success.
- Failure.
- No auditing.

**Discussion**
Auditing directory service access and configuring SACLs on directory objects can generate a large volume of entries in the security logs on domain controllers. You should only configure these policy settings if you actually intend to use the information.

Note that you can set an SACL on an object in Active Directory by using the **Security** tab in that object's properties. This is analogous to the **Audit object access** policy setting, except that it applies only to Active Directory objects and not to file system and registry objects.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Success |
| Stand-Alone Server Default Settings | No auditing |
| DC Effective Default Settings | Success |
| Member Server Effective Default Settings | No auditing |

**Audit logon events**
The **Audit logon events** policy setting determines whether to audit each instance of a user logging on to, logging off from, or

making a network connection to the computer that is recording the audit event.

The possible values for this Group Policy setting are:

- Success.
- Failure.
- No auditing.

**Discussion**
If you are logging successful account logon audit events on a domain controller, workstation logon attempts do not generate logon audits. Only interactive and network logon attempts to the domain controller itself generate logon events. In short, account logon events are generated where the account is located; logon events are generated where the logon attempt occurs.

If you configure this policy setting, you can specify whether to audit successes, audit failures, or not audit the logon events at all. Success audits generate an audit entry when a logon attempt succeeds, which is useful information for accounting purposes and for post-incident forensics so that you can determine who successfully logged on to which computer. Failure audits generate an audit entry when a logon attempt fails, which is useful for intrusion detection, but this setting also creates a potential denial-of-service condition because a malicious user could generate millions of logon failures and fill your security event log.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Success |
| Stand-Alone Server Default Settings | Success |
| DC Effective Default Settings | Success |
| Member Server Effective Default Settings | Success |

**Audit object access**
The **Audit object access** policy setting determines whether to audit the event of a user accessing an object—for example, a file, folder, registry key, or printer—that has specified its own SACL.

The possible values for this Group Policy setting are:

- Success.
- Failure.
- No auditing.

**Discussion**
If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit object access at all. Success audits generate an audit entry when a user successfully accesses an object that has specified an SACL. Failure audits generate an audit entry when a user unsuccessfully attempts to access an object that has specified an SACL; some failure events are to be expected in the course of normal system operations. For example, many applications, such as Microsoft Word, always attempt to open files with both read and write user rights; if they are unable to do so, they try to open them with read-only user rights. When this happens, a failure event will be recorded if you have enabled both failure auditing and the appropriate SACL on that file.

Enabling auditing of object access and configuring SACLs on objects can generate a large volume of entries in the security logs on systems in your enterprise; therefore, you should only enable these policy settings if you actually intend to use the information.

Enabling the capability to audit an object—such as a file, folder, printer, or registry key—is a two-step process in Windows Server 2003. After enabling the **Audit object access** policy setting, you must determine the objects to which you want to monitor access, and then modify their SACLs accordingly. For example, if you want to audit any attempts by users to open a particular file, you can use Group Policy or Windows Explorer to set a **Success** or **Failure** attribute directly on the file that you want to monitor for that particular event.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No auditing |
| Stand-Alone Server Default Settings | No auditing |
| DC Effective Default Settings | No auditing |
| Member Server Effective Default Settings | No auditing |

**Audit policy change**
The **Audit policy change** policy setting determines whether to audit every incidence of a change to user rights assignment policy settings, audit policy settings, or trust policy settings.

The possible values for this Group Policy setting are:

- Success.
- Failure.
- No auditing.

**Discussion**
If you configure this policy setting, you can specify whether to audit successes, audit failures, or not audit the policy settings changes at all. Success audits generate an audit entry when a change to user rights assignment, audit, or trust policy settings is successful, which is useful information for accounting purposes and for post-incident forensics so that you can determine who successfully modified policy settings in the domain or on individual computers. Failure audits generate an audit entry when a change to user rights assignment, audit, or trust policy settings fails.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Success |
| Stand-Alone Server Default Settings | No auditing |
| DC Effective Default Settings | Success |
| Member Server Effective Default Settings | No auditing |

**Audit privilege use**
The **Audit privilege use** policy setting determines whether to audit each instance of a user exercising a user right.

The possible values for this Group Policy setting are:

- Success.

- Failure.
- No auditing.

**Discussion**
If you configure this policy setting, you can specify whether to audit successes, audit failures, or not audit privilege use at all. Success audits generate an audit entry when the exercise of a user right succeeds. Failure audits generate an audit entry when the exercise of a user right fails. The volume of events generated by enabling these policy settings is very high and difficult to sort through. You should only enable these policy settings if you have a plan for how you will use the information that is generated.

By default, audit events are not generated for use of the following user rights, even if success audits or failure audits are specified for **Audit privilege use**:

- Bypass traverse checking
- Debug programs
- Create a token object
- Replace process level token
- Generate security audits
- Backup files and directories
- Restore files and directories

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No auditing |
| Stand-Alone Server Default Settings | No auditing |
| DC Effective Default Settings | No auditing |
| Member Server Effective Default Settings | No auditing |

**Audit process tracking**
The **Audit process tracking** policy setting determines whether to audit detailed tracking information for events such as program activation, process exit, handle duplication, and indirect object access.

The possible values for this Group Policy setting are:

- Success.
- Failure.
- No auditing.

**Discussion**
If you configure this policy setting, you can specify whether to audit successes, audit failures, or not audit process tracking at

all. Success audits generate an audit entry when the process being tracked succeeds. Failure audits generate an audit entry when the process being tracked fails.

Setting the value of **Audit process tracking** to **Success** or **Failure** will generate a large number of events, so typically it is set to **No auditing**. However, process tracking can be helpful in post-incident forensics because it provides a detailed log of which processes were started and the time they were started.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No auditing |
| Stand-Alone Server Default Settings | No auditing |
| DC Effective Default Settings | No auditing |
| Member Server Effective Default Settings | No auditing |

**Audit system events**
The **Audit system events** policy setting determines whether to audit when a user restarts or shuts down their computer, or when an event occurs that affects either system security or the security log.

The possible values for this Group Policy setting are:

- Success.
- Failure.
- No auditing.

**Discussion**
If you configure this policy setting, you can specify whether to audit successes, audit failures, or not audit the system events at all. Success audits generate an audit entry when a system event is executed successfully. Failure audits generate an audit entry when a system event is attempted unsuccessfully. Because only a few additional events will be recorded by setting **Audit system events** to both **Failure** and **Success**, and all of those events will be very significant, it is recommended that you set these values on all computers in your organization.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Success |
| Stand-Alone Server Default Settings | No auditing |
| DC Effective Default Settings | Success |
| Member Server Effective Default Settings | No auditing |

**User Rights Assignment**
User rights are tasks that a user is permitted to perform on a computer system or domain. There are two types of user rights: logon rights and privileges. Logon rights control who is authorized to log on to a computer and how they can log on. Privileges control access to system-wide resources on a computer and can override permissions set on specific objects. An example of a

logon right is the right to log on to a computer locally. An example of a privilege is the right to shut down the system. Both types of user rights are assigned by administrators to individual users or groups as part of the security policy settings for the computer.

The User Rights Assignment policy settings can be configured in the following location in Group Policy Object Editor:

Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Access this computer from the network**

The **Access this computer from the network** policy setting allows a user to connect to the computer from the network. This user right is required by a number of network protocols, including server message block (SMB)–based protocols, network BIOS (NetBIOS), Common Internet File System (CIFS), and Component Object Model Plus (COM+).

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

Users who can connect from their computer to the network can access resources on those computers for which they have permission. For example, this right is required for a user to connect to shared printers and folders. If this right is assigned to the Everyone group, and some shared folders have both share and NTFS file system permissions configured so that the Everyone group has read access, anyone will be able to view files in those shared folders. This is an unlikely situation for fresh installations of Windows Server 2003, because the default share and NTFS permissions in Windows Server 2003 do not include the Everyone group. For systems upgraded from Windows NT 4.0 or Windows 2000, this policy setting can have a higher level of risk because the default permissions for these operating systems are not as restrictive as the default permissions in Windows Server 2003.

It is advisable to assign the **Access this computer from the network** user right only to those users who require access to the server. For example, assigning this to the Administrators and Users groups will allow a user who logged on to the domain to access servers in the domain.

Removing this user right on domain controllers from all users will prevent anyone from logging on to the domain or using network resources. Removing this right on member servers will prevent users from connecting to those servers by way of the network. For these reasons, it is important to verify that authorized users have this right for the computers they need to access on the network.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Everyone, Administrators, Authenticated Users, ENTERPRISE DOMAIN CONTROLLERS, Pre-Windows 2000 Compatible Access |
| Stand-Alone Server Default Settings | Everyone, Administrators, Users, Power Users, Backup Operators |
| DC Effective Default Settings | Everyone, Administrators, Authenticated Users, ENTERPRISE DOMAIN CONTROLLERS, Pre-Windows 2000 Compatible Access |
| Member Server Effective Default Settings | Backup Operators, Power Users, Users, Administrators, Everyone |

**Act as part of the operating system**

The **Act as part of the operating system** user right allows a process to assume the identity of any user and thus gain access to the resources that the user is authorized to access. Typically, only low-level authentication services require this user right. Note that potential access is not limited to what is associated with the user by default; the calling process might request that arbitrary additional user rights be added to the access token. The calling process might also build an access token that does not provide a primary identity for events recorded in the security log.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
This user right is extremely powerful; anyone with this right can take complete control of the computer and erase virtually any evidence of their activities.

Limit the **Act as part of the operating system** user right to as few accounts as possible—it should not even be assigned to the Administrators group under normal circumstances. When a service requires this user right, configure the service to log on by using the local System account, which has this user right inherently. Do not create a separate account and assign this user right to it.

This user right is rarely needed by any accounts other than the local System account.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No one |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | No one |
| Member Server Effective Default Settings | Not defined |

**Add workstations to domain**
The **Add workstations to domain** user right allows the user to add a computer to a specific domain. For the user right to take effect, it must be assigned to the user as part of the Default Domain Controllers Policy for the domain. A user who has this user right can add up to 10 workstations to the domain. Users can also join a computer to a domain if they have been assigned the **Create Computer Objects** permission for an OU or for the Computers container in Active Directory. Users who have the **Create Computer Objects** permission can add an unlimited number of computers to the domain regardless of whether they have been assigned the **Add workstations to domain** user right.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
This user right has a moderate risk. Users with this right can add a computer to the domain that has been configured in a way that violates corporate security policies. For example, even if your organization does not want its users to have administrative user rights on their computer, a user could install Windows on their computer and then add the computer to the domain. They would know the password for the local Administrator account, and they could log on by using that account and then add their domain account to the local Administrators group.

Configure the **Add workstations to domain** user right so that only authorized members of the IT team are allowed to add computers to the domain. For organizations that have never allowed users to set up their own systems or add them to the

domain, this countermeasure will have no impact. For those that have allowed some or all users to configure their own computers, this countermeasure will force the organization to establish a formal process for these procedures.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Authenticated Users |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Authenticated Users |
| Member Server Effective Default Settings | Not defined |

**Adjust memory quotas for a process**
The **Adjust memory quotas for a process** user right allows a user to adjust the maximum memory available to a process. This user right is useful for system tuning, but it can be abused. In the wrong hands, it could be used to start a denial-of-service attack.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
A user with this right can reduce the amount of memory available to any process. This could cause business-critical network applications to slow down or even to fail.

Restrict the **Adjust memory quotas for a process** user right to users who require it to perform their job, such as application administrators who are responsible for maintaining database management systems or domain administrators responsible for managing the corporate directory and its supporting infrastructure.

Only organizations that have been lax with regard to restricting users to roles with limited user rights will find imposing this countermeasure difficult. For most organizations, implementing this restriction should have no impact.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | LOCAL SERVICE, NETWORK SERVICE, Administrators |
| Stand-Alone Server Default Settings | LOCAL SERVICE, NETWORK SERVICE, Administrators |
| DC Effective Default Settings | LOCAL SERVICE, NETWORK SERVICE, Administrators |
| Member Server Effective Default Settings | Administrators, NETWORK SERVICE, LOCAL SERVICE |

**Allow log on locally**
**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The **Allow log on locally** user right allows a user to start an interactive session on the computer. Users who do not have this

right can still start a remote interactive session on the computer if they have the **Allow log on through Terminal Services** user right.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Any account with the right to log on locally can be used to log on to the console of the computer. Not restricting this user right to legitimate users who need to be able to log on to the console of the system could result in unauthorized users downloading and executing malicious code to elevate their user rights.

For domain controllers, only assign the **Allow log on locally** user right to the Administrators group. For servers that fill other roles, add Backup Operators and Power Users. For end-user computers, assign this right to the Users group also.

Alternatively, you can add groups such as Account Operators, Server Operators, and Guests to the **Deny log on locally** user right. Denying access to these groups could limit the abilities of users assigned to specific administrative roles in your environment. Confirm that delegated activities will not be adversely affected.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators, Backup Operators, Account Operators, Server Operators, Print Operators |
| Stand-Alone Server Default Settings | Administrators, Users, Power Users, Backup Operators |
| DC Effective Default Settings | Administrators, Backup Operators, Account Operators, Server Operators, Print Operators |
| Member Server Effective Default Settings | Backup Operators, Power Users, Users, Administrators |

**Allow log on through Terminal Services**
The **Allow log on through Terminal Services** user right allows a user to log on to the computer by using a Remote Desktop connection. You should not assign this right to individual users or groups. Instead, it is a best practice to control who can open a Remote Desktop connection to the computer by adding users to or removing them from the Remote Desktop Users group.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Any account with the right to log on via Terminal Services can be used to log on to the remote console of the computer. Not restricting this user right to legitimate users who need to log on to the console of the system might allow unauthorized users to download and execute malicious code to elevate their user rights.

For domain controllers, only assign the **Allow log on through Terminal Services** user right to the Administrators group. For servers that fill other roles and end-user computers, add the Remote Desktop Users group. For all server roles other than Terminal Servers, ensure that only authorized members of the IT team who need to be able to manage the systems remotely belong to either of these groups.

For Terminal Servers, ensure that only users who require access to the server have accounts that belong to the Remote Desktop Users group, because this built-in group by default has this logon right.

Alternatively, you can assign the **Deny log on through Terminal Services** user right to groups such as Account Operators, Server Operators, and Guests. However, be careful when implementing this method because you might block access to legitimate administrators who also happen to belong to a group that has the **Deny log on through Terminal Services** logon right.

Changing membership in these default groups, or removing this user right from other groups, can limit the abilities of users assigned to specific administrative roles in your environment. Confirm that delegated activities will not be adversely affected.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Administrators, Remote Desktop Users |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Remote Desktop Users, Administrators |

**Back up files and directories**
The **Back up files and directories** user right allows the user to circumvent file and directory permissions to back up the system. This user right is selected only when an application attempts access via the NTFS backup application programming interface (API) by using, for example, Ntbackup.exe. Otherwise, normal file and directory permissions apply.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Users who are able to back up data from a computer might take the backup media to a computer that does not belong to the domain where they have administrative user rights and restore the data. They could then take ownership of the files and view anything contained within the backup set.

Limit the **Back up files and directories** user right to members of the IT team who need to be able to back up corporate data as part of their day-to-day responsibilities.

Changing membership of the groups that have this user right might limit the abilities of users assigned to specific administrative roles in your environment. Confirm that authorized backup administrators are still able to perform backup operations.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators, Backup Operators, Server Operators |
| Stand-Alone Server Default Settings | Administrators, Backup Operators |
| DC Effective Default Settings | Administrators, Backup Operators, Server Operators |
| Member Server Effective Default Settings | Backup Operators, Administrators |

**Bypass traverse checking**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The **Bypass traverse checking** user right allows a user who is navigating an object path in NTFS or the registry to pass through folders without checking for the special access permission Traverse Folder. This user right does not allow the user to list the contents of a folder, it only allows the user to traverse the folder's directories.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.

- Not defined.

**Discussion**
The default setting for this user right is to allow anyone to bypass traverse checking. For experienced Windows system administrators, this is the expected behavior, and they configure file-system access control lists (SACLs) accordingly. The only scenario where the default configuration might lead to a mishap is if the administrator who is configuring permissions expects that users who are unable to access a folder will be unable to access the contents of any child folders. This is an unlikely situation, and therefore the risk is very low.

Organizations that are extremely concerned about security might want to remove the Everyone group, or perhaps even the Users group, from the list of groups that have the **Bypass traverse checking** user right.

Windows and many applications have been designed with the expectation that anyone who can legitimately access the computer will have this user right. Therefore, removing the Everyone group from the list of users who have this user right by default might lead to operating-system instability or application failure. It is recommended that you leave this policy setting at its default.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Everyone, Administrators, Authenticated Users, Pre-Windows 2000 Compatible Access |
| Stand-Alone Server Default Settings | Everyone, Administrators, Users, Power Users, Backup Operators |
| DC Effective Default Settings | Everyone, Administrators, Authenticated Users, Pre-Windows 2000 Compatible Access |
| Member Server Effective Default Settings | Backup Operators, Power Users, Users, Administrators, Everyone |

**Change the system time**
The **Change the system time** user right allows the user to adjust the time on the computer's internal clock. This user right is not required to change the time zone or other display characteristics of the system time.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.

- Not defined.

**Discussion**
Users who are able to set a computer to the wrong time can cause several problems: Time stamps on event log entries will no longer be accurate; time stamps that indicate when files and folders are created or modified will not be correct; and, if the computer belongs to a domain, it might be unable to authenticate itself or users who try to log on to the domain from it. The

risk of this happening is mitigated on most domain controllers, member servers, and end-user computers because the Windows Time service automatically synchronizes time with domain controllers in the following ways:

- All client desktop computers nominate the authenticating domain controller as their inbound time partner.
- All member servers follow the same process that client desktop computers follow.
- All domain controllers in a domain nominate the primary domain controller (PDC) operations master as their inbound time partner.
- All PDC operations masters follow the hierarchy of domains in the selection of their inbound time partner.
- The PDC operations master at the root of the domain is authoritative for the organization; therefore, it is recommended that you synchronize it with a reliable external time server.

This issue becomes much more serious if a malicious user is able to change the system time and then stop the Windows Time service, or reconfigure it to be synchronized with a time server that is not accurate.

Limit the **Change the system time** user right to users who have a legitimate need to change the system time, such as members of the IT team.

There should be no impact for most organizations because time synchronization will be fully automated for all computers that belong to the domain. Computers that do not belong to the domain should be synchronized with an external source.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators, Server Operators |
| Stand-Alone Server Default Settings | Administrators, Power Users |
| DC Effective Default Settings | Administrators, Server Operators |
| Member Server Effective Default Settings | Power Users, Administrators |

**Create a page file**
The **Create a page file** user right allows the user to create and change the size of a paging file. This is done by clicking **Performance** on the **Advanced** tab of the **System Properties** property sheet and specifying a paging file size for a particular drive in the **Performance Options** dialog box.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Users who can change the paging file size can make it extremely small or move the file to a highly fragmented storage volume, which might lead to reduced system performance.

You should only assign the **Create a page file** user right to members of the Administrators group.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators |

**Create a token object**

The **Create a token object** user right allows a process to create a token. If the process uses NtCreateToken() or another token-creation API, it can use the token to access any local resources. It is advisable not to assign this user right to any users.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

The operating system determines the level of a user's privileges by examining the user's access token. Access tokens are built when users log on to the local computer or connect to a remote computer over a network. When you revoke a user right, the change is immediately recorded in the registry, but the change is not reflected in the user's access token until the next time the user logs on or connects. A user with the ability to create or modify tokens can change the level of access for any currently logged-on account. Such a user can escalate their own user rights or create a denial-of-service condition.

Do not assign the **Create a token object** right to any users. Processes that require this user right should use the local System account, which already includes this user right.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No one |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | No one |
| Member Server Effective Default Settings | Not defined |

**Create global objects**

The **Create global objects** user right is required for a user account to create global objects that are available to all sessions. Users can still create objects specific to their own session without being assigned this user right.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

Users who can create global objects can affect processes that run under other users' sessions. This could lead to a variety of problems, such as application failure or data corruption.

Only assign the **Create global objects** user right to members of the local Administrators and Service groups.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Administrators, SERVICE |
| DC Effective Default Settings | Administrators, SERVICE |
| Member Server Effective Default Settings | Administrators, SERVICE |

**Create permanent shared objects**
The **Create permanent shared objects** user right allows a user to create a directory object in the object manager. This means that users with this user right can create shared folders, printers, and other objects. This user right is useful for kernel-mode components that extend the object namespace. Components that run in kernel mode have this user right inherently; therefore, it is normally not necessary to specifically assign this user right.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Users with this user right can expose sensitive data to the network by creating a new shared object.

Do not assign the **Create permanent shared objects** right to any users. Processes that require this user right should use the local System account, which already includes this user right. Users who belong to the local Administrators group will still be able to create, modify, and delete shared folders and printers.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No one |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | No one |
| Member Server Effective Default Settings | Not defined |

**Debug programs**
The **Debug programs** user right allows the user to attach a debugger to any process. This user right provides access to sensitive and critical operating system components.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

The **Debug programs** user right can be exploited to capture sensitive system information from system memory. Some attack tools exploit the debug program's user right to extract hashed passwords and other private security information. The risk of attackers' being able to exploit this user right is mitigated by the fact that the debug program's user right is by default assigned only to administrators.

Revoke the **Debug programs** user right from all users and groups. Revoking this user right will mean that no one except administrators will be able to debug programs. However, under normal circumstances, debugging is rarely necessary on production systems. If a problem arises that requires debugging an application on a production server temporarily, move the server to a different OU and assign the **Debug programs** user right to the appropriate account.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators |

**Deny access to this computer from the network**
The **Deny access to this computer from the network** user right prohibits a user from connecting to the computer from the network.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Users who can log onto the computer over the network can enumerate lists of account names, group names, and shared resources. Users with permission to access shared folders and files can connect over the network and possibly view or modify data. Explicitly denying this logon right to high-risk accounts, such as the local Guest account and other accounts that have no business reason for accessing the computer over the network, provides an additional layer of protection.

Assign the **Deny access to this computer from the network** user right to the following accounts:

- ANONYMOUS LOGON
- The built-in local Administrator account
- The local Guest account
- The built-in Support account (the Support_388945a0 account is the logon account for Remote Assistance)
- All service accounts

  There is an important exception to this list: do not assign the logon right to any service accounts that are used to start services that need to connect to the computer over the network. For example, if you have configured a shared folder for Web servers to access, and you present content within that folder through a Web site, you might need to allow the account under which IIS is running to log on to the server with the shared folder via the network.

Configuring this logon right for other groups might limit the abilities of users assigned to specific administrative roles in your environment. Verify that delegated tasks will not be negatively affected.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | SUPPORT_388945a0 |
| Stand-Alone Server Default Settings | SUPPORT_388945a0 |
| DC Effective Default Settings | SUPPORT_388945a0 |
| Member Server Effective Default Settings | SUPPORT_388945a0 |

**Deny log on as a batch job**
The **Deny log on as a batch job** user right prohibits a user from logging on by using a batch-queue facility. Batch-queue in Windows Server 2003 is used for scheduling jobs that will be automatically started one or more times in the future. This user right is needed for any accounts that are used for scheduling jobs via the Task Scheduler service.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Accounts with this logon right might be used to schedule jobs that consume excessive system resources, leading to a denial-of-service condition.

Assign the **Deny log on as a batch job** user right to the built-in Support account — the Support_388945a0 account is the logon account for Remote Assistance — and the local Guest account. Configuring this logon right for other accounts might prevent users who are assigned to specific administrative roles from performing their required job activities. Confirm that delegated tasks will not be affected adversely.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No one |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | No one |
| Member Server Effective Default Settings | Not defined |

**Deny log on as a service**
The **Deny log on as a service** logon right prohibits a user from logging on as a service.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

Accounts that can log on as a service might be used to configure and start new unauthorized services such as a Trojan horse or backdoor (a backdoor is a hidden entrance to an operating system that can be used to download system information). The benefit of configuring this countermeasure is somewhat mitigated by the fact that only users with administrative user rights can install and configure services, and a malicious user who has already attained that level of access can configure the service to run under the local System account.

It is advisable to keep this user right in its default state; that is, do not assign the **Deny log on as a service** logon right to any accounts. Organizations that are extremely concerned about security might want to assign this logon right to groups and accounts that they are certain will never need to log on as a service. Assigning this user right to specific accounts can prevent services from being started, which might lead to a denial-of-service condition.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No one |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | No one |
| Member Server Effective Default Settings | Not defined |

**Deny log on locally**
The **Deny log on locally** logon right prohibits a user from logging on directly at the computer's keyboard.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Any account with the right to log on locally can be used to log on at the console of the computer. If this user right is not restricted to legitimate users who need to log on to the console of the system, unauthorized users can download and execute malicious code that elevates their user rights.

Assign the **Deny log on locally** logon right to the built-in Support account. The Support_388945a0 account is the logon account for Remote Assistance.

The Support_388945a0 account enables Help and Support Center interoperability with signed scripts. This account is primarily used to control access to signed scripts that are accessible from within Help and Support Center. Administrators can use this account to delegate the ability for an ordinary user, who does not have administrative access to a computer, to run signed scripts from links embedded within Help and Support Center. These scripts can be programmed to use the Support_388945a0 account credentials instead of the user's credentials to perform specific administrative operations on the local computer, operations that otherwise would not be supported by the ordinary user's account. When the delegated user clicks a link in Help and Support Center, the script runs under the security context of the Support_388945a0 account. This account has limited access to the computer and is disabled by default.

Assigning this logon right to additional accounts might limit the abilities of users assigned to specific administrative roles in your environment. Confirm that delegated activities will not be adversely affected.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|

| | |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | SUPPORT_388945a0 |
| Stand-Alone Server Default Settings | SUPPORT_388945a0 |
| DC Effective Default Settings | SUPPORT_388945a0 |
| Member Server Effective Default Settings | SUPPORT_388945a0 |

**Deny log on through Terminal Services**

The **Deny log on through Terminal Services** logon right prohibits a user from logging on to the computer by using a Remote Desktop connection.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

Any account with the right to log on using Terminal Services can be used to log on to the remote console of the computer. If this user right is not restricted to legitimate users who need to log on to the console of the system, unauthorized users might download and execute malicious code that elevates their user rights.

Assign the **Deny log on through Terminal Services** logon right to the built-in local Administrator account and all service accounts. Assigning this logon right to other groups might limit the abilities of users assigned to specific administrative roles in your environment. Confirm that delegated tasks will not be negatively affected.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Enable computer and user accounts to be trusted for delegation**

The **Enable computer and user accounts to be trusted for delegation** user right allows the user to change the **Trusted for Delegation** policy setting on a user or computer object in Active Directory. The user or computer that is assigned this user right must also have write access to the account control flags on the object.

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

Delegation of authentication is used by multitier client/server applications. It allows a front-end service to use the credentials of

a client while authenticating to a back-end service. For this to be possible, both client and server must be running under accounts that are trusted for delegation.

Misuse of this user right can lead to unauthorized users' impersonating other users' on the network. A malicious user can exploit this user right to gain access to network resources while appearing to be a different user, which could make post–security incident forensics more difficult.

It is advisable to assign the **Enable computer and user accounts to be trusted for delegation** user right only to the Administrators group on domain controllers. There is no reason to assign it to anyone on member servers or workstations that belong to a domain, because it has no meaning in those contexts; this user right is only relevant to domain controllers and stand-alone systems.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Not defined |

**Force shutdown from a remote system**
The **Force shutdown from a remote system** user right allows a user to shut down a computer from a remote location on the network.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Any user who can shut down a computer can cause a denial-of-service condition; therefore, this user right should be tightly restricted.

Only assign the **Force shutdown from a remote system** user right to members of the Administrators group. Removing the Server Operators group might limit the abilities of users assigned to specific administrative roles in your environment. Confirm that delegated activities will not be adversely affected.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators, Server Operators |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators, Server Operators |
| Member Server Effective Default Settings | Administrators |

**Generate security audits**
The **Generate security audits** user right allows a process to generate audit records in the security log. The security log can be

used to trace unauthorized system access.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Accounts that are able to write to the security log can be used by a malicious user to fill that log with meaningless events. If the computer is configured to overwrite event logs as they fill up, the attacker can use this method to remove evidence of their unauthorized activities. If the computer is configured to shut down when it is unable to write to the security log, this method can be used to create a denial-of-service condition.

Ensure that only the Local Service and Network Service groups have the **Generate security audits** user right assigned to them.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | LOCAL SERVICE, NETWORK SERVICE |
| Stand-Alone Server Default Settings | LOCAL SERVICE, NETWORK SERVICE |
| DC Effective Default Settings | LOCAL SERVICE, NETWORK SERVICE |
| Member Server Effective Default Settings | LOCAL SERVICE, NETWORK SERVICE |

**Impersonate a client after authentication**
The **Impersonate a client after authentication** user right allows programs running on behalf of a user to impersonate a client. Requiring this user right for this kind of impersonation prevents an unauthorized user from convincing a client to connect — for example, by remote procedure call (RPC) or named pipes — to a service that they have created and then impersonating that client, which can elevate the unauthorized user's credentials to administrative or system levels.

Services that are started by the Service Control Manager have the built-in Service group added by default to their access tokens. COM servers started by the COM infrastructure that are configured to run under a specific account also have the Service group added to their access tokens. As a result, these services get this user right when they are started.

In addition, a user can also impersonate an access token if any of the following conditions exist.

- The access token that is being impersonated is for this user.
- The user, in this logon session, created the access token by logging on to the network with explicit credentials.
- The requested level is less than Impersonate, such as Anonymous or Identify.

Because of these factors, users do not usually need this user right.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
A user with this user right can trick a client into connecting to a service that they have created and then impersonating that

client. The user can then elevate their level of access to that of the other client.

Ensure that only the Administrators and Service groups have the **Impersonate a client after authentication** user right assigned to them.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Administrators, SERVICE |
| DC Effective Default Settings | Administrators, SERVICE |
| Member Server Effective Default Settings | Administrators, SERVICE |

**Increase scheduling priority**
The **Increase scheduling priority** user right allows a user to increase the base priority class of a process. Increasing relative priority within a priority class is not a privileged operation. This user right is not required by administrative tools supplied with the operating system, but might be required by software development tools.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
A user with this user right can increase the scheduling priority of a process to Real-Time. This might leave little processing time for all other processes, which could lead to a denial-of-service condition.

Only Administrators should have the **Increase scheduling priority** user right assigned to them.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators |

**Load and unload device drivers**
The **Load and unload device drivers** user right determines which users can dynamically load and unload device drivers. This user right is not required if a signed driver for the new hardware already exists in the Driver.cab file on the computer.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Device drivers run as highly privileged code. A user who has the **Load and unload device drivers** user right might unintentionally install malicious code that masquerades as a device driver. It is assumed that administrators will exercise greater care and only install drivers with verified digital signatures.

You must have this user right, and also be a member of either Administrators or Power Users, to install a new driver for a local printer or manage a local printer by setting defaults for options such as duplex printing. This requirement — to have the user right and also be a member of the Administrators or Power Users group — is new in the Windows XP and Windows Server 2003 operating systems.

Do not assign the **Load and unload device drivers** user right to any user or group other than Administrators. Removing this user right from the Print Operators group or other accounts might limit the abilities of users assigned to specific administrative roles in your environment. Ensure that delegated tasks will not be negatively affected.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators, Print Operators |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators, Print Operators |
| Member Server Effective Default Settings | Administrators |

**Lock pages in memory**
The **Lock pages in memory** user right allows a process to keep data in physical memory, which prevents the system from paging the data to virtual memory on disk. Assigning this user right can result in significant degradation of system performance.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Users with this user right can assign physical memory to several processes, leaving little or no RAM for other processes. This could lead to a denial-of-service condition.

Do not assign the **Lock pages in memory** user right to any accounts.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No one |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | No one |
| Member Server Effective Default Settings | Not defined |

**Log on as a batch job**

The **Log on as a batch job** user right allows a user to log on by using a batch-queue facility such as the Task Scheduler service. When an administrator uses the Add Scheduled Task wizard to schedule a task to run under a particular user name and password, that user automatically receives the **Log on as a batch job** user right. When the scheduled time arrives, the Task Scheduler service logs the user on as a batch job rather than as an interactive user, and the task runs in the user's security context.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

For most organizations, the default values are sufficient. Allow the system to manage this user right automatically if you want to allow tasks to be scheduled to run for specific user accounts. If you do not want to use the Task Scheduler in this manner, configure the **Log on as a batch job** user right for only the Local Service group and the local support account (Support_388945a0). For Internet Information Services (IIS) servers, configure this policy locally, rather than through domain-based Group Policy, so you can ensure that the local IUSR_computername and IWAM_computername accounts have this logon right.

If you specify this policy setting through domain-based Group Policy, the system will not be able to assign this user right to accounts used for scheduled jobs in the Task Scheduler, and the IUSR_computername and IWAM_computername accounts will not have this logon right. This will cause IIS to be unable to run some COM objects that are necessary for proper functioning.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | LOCAL SERVICE, SUPPORT_388945a0 |
| Stand-Alone Server Default Settings | LOCAL SERVICE, SUPPORT_388945a0 |
| DC Effective Default Settings | LOCAL SERVICE, SUPPORT_388945a0 |
| Member Server Effective Default Settings | LOCAL SERVICE, SUPPORT_388945a0 |

**Log on as a service**

The **Log on as a service** logon right allows a security principal to log on as a service. Services can be configured to run under the Local System, Local Service, or Network Service accounts, which have a built-in right to log on as a service. Any service that runs under a separate user account must be assigned this right.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

This is a powerful logon right that allows accounts to start network services. The risk is mitigated by the fact that only users with administrative user rights can install and configure services. A malicious user who has already attained that level of access can configure the service to run under the local System account.

The default security principal with the **Log on as a service** logon right is Network Service, a built-in local group. Service accounts should be added to this local group, and you should monitor membership for this group and this user right to ensure that unexpected changes are not being made.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | NETWORK SERVICE |
| Stand-Alone Server Default Settings | NETWORK SERVICE |
| DC Effective Default Settings | NETWORK SERVICE |
| Member Server Effective Default Settings | NETWORK SERVICE |

**Manage auditing and security log**

The **Manage auditing and security log** user right allows a user to specify object access auditing options for individual resources such as files, Active Directory objects, and registry entries. Object access auditing is not performed unless you enable it by using Group Policy Object Editor and configuring Audit Policy in the following location:

Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy

A user who has this user right can also view and clear the security log from Event Viewer.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

The right to manage the security event log is a powerful user right, and it should be closely guarded. Anyone with this user right can clear the security log, possibly erasing important evidence of unauthorized activity.

Ensure that only the local Administrators group has the **Manage auditing and security log** user right.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators |

**Modify firmware environment values**

The **Modify firmware environment values** user right allows system environment variables to be modified, either by a process through an application programming interface (API) or by a user through **System Properties**.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

Anyone with this user right can configure the settings of a hardware component in such a way that it fails. This can lead to data

corruption or a denial-of-service condition.

Ensure that only the local Administratorsgroup has the **Modify firmware environment values** user right.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators |

**Perform volume maintenance tasks**
The **Perform volume maintenance tasks** user right allows a non- administrative or remote user to manage volumes or disks. Windows Server 2003 checks for the user right in a user's access token when a process running in the user's security context calls **SetFileValidData()**.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
A user with this user right can delete a volume, leading to the loss of data or a denial-of-service condition.

Ensure that only the local Administrators group has the **Perform volume maintenance tasks** user right.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators |

**Profile single process**
The **Profile single process** user right allows a user to sample the performance of an application process. Ordinarily, you do not need this user right to use the Performance snap-in; however, you do need the user right if System Monitor is configured to collect data by using Windows Management Instrumentation (WMI).

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
A malicious user with this user right can monitor a computer's performance to help identify critical processes that they might want to attack directly. The attacker can also determine what processes are running on the system and identify countermeasures that they should avoid — such as antivirus software or an intrusion-detection system — or identify which other users are logged on to the system.

Ensure that only the local Administrators group has the **Profile single process** user right. Removing this user right from the Power Users group or other accounts might limit the abilities of users who are assigned to specific administrative roles in your environment. Ensure that delegated tasks will not be negatively affected.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Administrators, Power Users |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators, Power Users |

**Profile system performance**
The **Profile system performance** user right allows a user to sample the performance of system processes. Ordinarily, you do not need this user right to use the Performance snap-in; however, you do need the user right if System Monitor is configured to collect data by using WMI.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
This is a moderate risk. A malicious user with this user right can monitor a computer's performance to help identify critical processes that they might want to attack directly. The attacker might also be able to determine what processes are running on the system so they can identify countermeasures that they might need to avoid, such as antivirus software or an intrusion-detection system.

Ensure that only the local Administrators group has the **Profile system performance** user right.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators |

**Remove computer from docking station**
The **Remove computer from docking station** user right allows the user of a portable computer to undock the computer by

clicking **Eject PC** on the **Start** menu.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Anyone who has this user right can remove a portable computer that has been started from its docking station. The value of employing this countermeasure is reduced by several factors: If a malicious user can restart the computer, they can remove it from the docking station after the BIOS has started but before the operating system has started; or a malicious user might steal the computer and the docking station together. This setting will not affect servers, because they are not typically installed in a docking station

Ensure that only the local Administrators and Power Users groups have the **Remove computer from docking station** user right.

This is the default configuration, so it should have little impact; however, if your organization's end users are not members of the Power Users or Administrators groups, they will be unable to remove their own portable computers from their docking stations without shutting down the system. Therefore, on portable computers, you might want to assign this user right to the local Users group.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Administrators, Power Users |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators, Power Users |

**Replace a process level token**
The **Replace a process level token** user right allows a parent process to replace the access token that is associated with a child process.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
If a user has this right and the **Adjust memory quotas for a process** user right, they can start processes as other users. They could use this method to hide the fact that they are performing unauthorized actions on the computer.

Ensure that only the Local Service and Network Service groups have the **Replace a process level token** user right.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |

| | |
|---|---|
| Default Domain Controller Policy | LOCAL SERVICE, NETWORK SERVICE |
| Stand-Alone Server Default Settings | LOCAL SERVICE, NETWORK SERVICE |
| DC Effective Default Settings | LOCAL SERVICE, NETWORK SERVICE |
| Member Server Effective Default Settings | LOCAL SERVICE, NETWORK SERVICE |

**Restore files and directories**

The **Restore files and directories** user right allows a user to circumvent file and directory permissions when restoring backed-up files and directories, and to set any valid security principal as the owner of an object.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

A malicious user with this user right can restore sensitive corporate data to a system, overwriting data that is more recent. This could lead to loss of important data, data corruption, or a denial-of-service condition.

This countermeasure will not prevent a malicious user from restoring data to an unmanaged system, so it is critical that organizations carefully protect the media they use for backing up data.

Ensure that only the local Administrators group has the **Restore files and directories** user right. Removing this user right from the Backup Operators group and other accounts might limit the abilities of users assigned to specific administrative roles in your environment. Ensure that delegated tasks will not be negatively affected.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators, Backup Operators, Server Operators |
| Stand-Alone Server Default Settings | Administrators, Backup Operators |
| DC Effective Default Settings | Administrators, Backup Operators, Server Operators |
| Member Server Effective Default Settings | Administrators, Backup Operators |

**Shut down the system**

The **Shut down the system** user right allows a user to shut down the local computer.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**

The ability to shut down domain controllers should be limited to a very small number of trusted administrators. Even though a system shutdown requires the ability to log on to the server, you should be very careful about the accounts and groups that you allow to shut down a domain controller.

Shutting down domain controllers makes them unavailable to perform such functions as processing logons, serving Group Policy, and answering Lightweight Directory Access Protocol (LDAP) queries. Shutting down domain controllers that have been

assigned operations master roles (also known as flexible single master operations or FSMO) roles can disable key domain functionality, such as processing logons for new passwords, which is performed by the PDC emulator operations master.

Ensure that only Administrators and Backup Operators have the **Shut down the system** user right on member servers, and that only Administrators have the user right on domain controllers. Removing these default groups might limit the abilities of users assigned to specific administrative roles in your environment. Ensure that delegated tasks will not be negatively affected.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators, Backup Operators, Server Operators, Print Operators |
| Stand-Alone Server Default Settings | Administrators, Power Users, Backup Operators, Users |
| DC Effective Default Settings | Administrators, Backup Operators, Server Operators, Print Operators |
| Member Server Effective Default Settings | Administrators, Power Users, Backup Operators, Users |

**Synchronize directory service data**
The **Synchronize directory service data** user right allows a process to read all objects and properties in the directory, regardless of the protection on the objects and properties. This user right is required in order to use LDAP directory synchronization (Dirsync) services.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
This user right affects domain controllers; only domain controllers should be able to synchronize directory service data. They have this right inherently because the synchronization process runs in the context of the System account on domain controllers. A malicious user who has this user right can view all information stored within the directory. They can then use some of that information to facilitate additional attacks or expose sensitive corporate data, such as direct telephone numbers or physical addresses.

Ensure that no accounts have the **Synchronize directory service data** user right.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | No one |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | No one |
| Member Server Effective Default Settings | Not defined |

**Take ownership of files or other objects**
The **Take ownership of files or other objects** user right allows a user to take ownership of any securable object in the system, including Active Directory objects, NTFS files and folders, printers, registry keys, services, processes, and threads.

The possible values for this Group Policy setting are:

- A user-defined list of accounts.
- Not defined.

**Discussion**
Any user with this ability can take control of any object regardless of the permissions on that object and then make any changes they want to that object. This can result in exposure of data, corruption of data, or a denial-of-service condition.

Ensure that only the local Administrators group has the **Take ownership of files or other objects** user right.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Administrators |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators |

**Security Options**
The Security Options section of Group Policy enables or disables computer security policy settings, such as digital data signing, Administrator and Guest account names, access to floppy disk and CD drives, driver installation behavior, and logon prompts. The Security Options policy settings can be configured in the following location in Group Policy Object Editor:

Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Accounts: Administrator account status**
The **Accounts: Administrator account status** policy setting enables or disables the Administrator account under normal operation. When a computer is booted in Safe Mode, the Administrator account is always enabled, regardless of this policy setting.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
For some organizations, periodically changing the password for local accounts can be a daunting management challenge, therefore they might want to disable the built-in Administrator account. Another reason to consider disabling this built-in account is that by default it cannot be locked out, no matter how many failed logons it accrues. This makes it a prime target for brute-force password-guessing attacks. Additionally, this account has a well-known security identifier (SID); some non-Microsoft tools allow you to authenticate over the network by specifying the SID rather than the account name. This means that even if you have renamed the Administrator account, a malicious user could start a brute-force attack by using the SID.

Disabling the Administrator account can become a maintenance issue under certain circumstances. For example, in a domain environment, if the secure channel that constitutes your connection fails for any reason, and there is no other local Administrator account, you must restart the computer in Safe Mode to fix the problem that broke your connection status.

If the current Administrator password does not meet the password requirements, you will not be able to enable the Administrator account again after it has been disabled. In this case, another member of the Administrators group must set the password on the Administrator account by using the Local Users and Groups tool.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Accounts: Guest account status**
The **Accounts: Guest account status** policy setting determines whether the Guest account is enabled or disabled.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
This account allows unauthenticated network users to gain access to the system by logging in as Guest with no password. Unauthorized users can access any resources that are accessible to the Guest account over the network. This means that any network shares with permissions allowing access to the Guest account, the Guests group, or the Everyone group will be accessible over the network. This can lead to the exposure or corruption of data.

Set **Accounts: Guest account status** to **Disabled** so that the built-in Guest account is no longer usable. All network users will have to authenticate before they can access shared resources on the system. If the Guest account is disabled and **Network access: Sharing and security model for local accounts** is set to **Guest only**, network logons—such as those performed by the SMB Service—will fail. This should have little impact on most organizations because it is the default value for this policy setting in Windows 2000, Windows XP, and Windows Server 2003.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Accounts: Limit local account use of blank passwords to console logon only**
The **Accounts: Limit local account use of blank passwords to console logon only** policy setting determines whether remote interactive logons by network services such as Terminal Services, Telnet, and File Transfer Protocol (FTP) are allowed for local accounts that have blank passwords. If this policy setting is enabled, a local account must have a nonblank password to be used to perform an interactive or network logon from a remote client.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
This policy setting does not affect interactive logons that are performed physically at the console, or logons that use domain accounts. It is possible for non-Microsoft applications that use remote interactive logons to bypass this policy setting.

Blank passwords are a serious threat to computer security and they should be forbidden through both corporate policy and suitable technical measures. In fact, the default policy settings for Windows Server 2003 Active Directory domains require complex passwords of at least seven characters. Nevertheless, if a user with the ability to create new accounts creates one that has bypassed your domain-based password policy settings, that account might have a blank password. For example, a user could build a stand-alone system, create one or more accounts with blank passwords, and then join the computer to the domain. The local accounts with blank passwords would still function. Anyone who knows the account name can then use accounts with blank passwords to log on to systems.

It is advisable to set **Accounts: Limit local account use of blank passwords to console logon only** to **Enabled**.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Accounts: Rename administrator account**
The **Accounts: Rename administrator account** policy setting determines whether a different account name is associated with the SID for the Administrator account.

The possible values for this Group Policy setting are:

- User-defined text.
- Not defined.

**Discussion**
Because the Administrator account exists on all Windows 2000, Windows Server 2003, and Windows XP Professional computers, renaming the account makes it slightly more difficult for attackers to guess this user name and password combination. By default, the built-in Administrator account cannot be locked out no matter how many times a malicious user might use a bad password. This makes the Administrator account a popular target for brute-force password-guessing attacks. The value of this countermeasure is lessened because this account has a well-known SID and there are non-Microsoft tools that allow you to initiate a brute-force attack over the network by specifying the SID rather than the account name. This means that even if you have renamed the Administrator account, a malicious user could start a brute-force attacking by using the SID.

Rename the Administrator account by specifying a value for the **Accounts: Rename administrator account** policy setting.

This policy setting is not configured in the security templates, nor is a new user name for the account suggested here. This is so that numerous organizations reading this guide will not implement the same new user name in their environments.

Be sure to inform users who are authorized to use this account of the new account name.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Administrator |
| DC Effective Default Settings | Administrator |
| Member Server Effective Default Settings | Administrator |

**Accounts: Rename guest account**

The **Accounts: Rename guest account** policy setting determines whether a different account name is associated with the SID for the Guest account.

The possible values for this Group Policy setting are:

- User-defined text.
- Not defined.

**Discussion**

Because the Guest account exists on all Windows 2000, Windows Server 2003, and Windows XP Professional computers, renaming the Guest account makes it slightly more difficult for unauthorized persons to guess this privileged user name and password combination.

Rename the Guest account by specifying a value for the **Accounts: Rename guest account** policy setting.

This policy setting is not configured in the security templates.

The Guest account is disabled by default in Windows 2000, Windows XP, and Windows Server 2003.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Guest |
| DC Effective Default Settings | Guest |
| Member Server Effective Default Settings | Guest |

**Audit: Audit the access of global system objects**

Enabling the **Audit: Audit the access of global system objects** policy setting creates system objects—such as mutexes, events, semaphores, and MS-DOS devices—with a default SACL. If the **Audit object access** policy setting is also enabled, access to these system objects is audited.

Global system objects, also known as *base system objects* or *base named objects*, are ephemeral kernel objects that have had names assigned to them by the application or system component that created them. These objects are most commonly used to synchronize multiple applications or multiple parts of a complex application. By virtue of having a name, these objects are global in scope, and therefore visible to all processes on the system. These objects all have a security descriptor, but typically have a NULL SACL. Enabling this policy setting at boot time causes the kernel to assign an SACL to these objects when they are created.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
The threat is that a globally visible named object, if incorrectly secured, might be acted on by a malicious program that knows the name of the object. For instance, if a synchronization object such as a mutex has a poorly chosen discretionary access control list (DACL), a malicious program can access that mutex by name and cause the program that created it to malfunction; however, the risk of this occurring is very low.

Enabling this policy setting can generate a large number of security events, especially on busy domain controllers and application servers. This might cause servers to respond slowly and force the security log to record numerous events of little significance. Auditing for access to global system objects is an all-or-nothing affair; there is no way to filter which events get recorded and which do not. Even if an organization has the resources to analyze events generated when this policy setting is enabled, it is unlikely to have the source code or a description of what each named object is used for; therefore, it is unlikely that many organizations could benefit from enabling this policy setting.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Audit: Audit the use of Backup and Restore privilege**
The **Audit: Audit the use of Backup and Restore privilege** policy setting determines whether to audit the use of all user rights, including Backup and Restore, when the **Audit privilege use** policy setting is configured. Enabling both policy settings generates an audit event for every file that is backed up or restored.

Enabling this policy setting in conjunction with the **Audit privilege use** policy setting records any instance of user rights being exercised in the security log. If **Audit privilege use** is enabled but **Audit: Audit the use of Backup and Restore privilege** is disabled, when users use Backup or Restore user rights, those events will not be audited.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Enabling this policy setting when the **Audit privilege use** policy setting is also enabled generates an audit event for every file that is backed up or restored. This can help you to track down an administrator who is accidentally or maliciously restoring data in an unauthorized manner.

It is advisable to set **Audit use of Backup and Restore privilege** to **Disabled**. Enabling this policy setting can generate a large number of security events, which might cause servers to respond slowly and force the security log to record numerous events of little significance.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Audit: Shut down system immediately if unable to log security audits**
**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The **Audit: Shut down system immediately if unable to log security audits** policy setting determines whether the system shuts down if it is unable to log security events. This policy setting is a requirement for Trusted Computer System Evaluation Criteria (TCSEC)-C2 and Common Criteria certification to prevent auditable events from occurring if the audit system is unable to log those events. Microsoft has chosen to meet this requirement by halting the system and displaying a Stop message in the case of a failure of the auditing system. Enabling this policy setting stops the system if a security audit cannot be logged for any reason. Typically, an event fails to be logged when the security audit log is full and the value of **Retention method for security log** is either **Do not overwrite events (clear log manually)** or **Overwrite events by days**.

With **Audit: Shut down system immediately if unable to log security audits** set to **Enabled**, if the security log is full and an existing entry cannot be overwritten, the following Stop message appears:

STOP: C0000244 {Audit Failed}

An attempt to generate a security audit failed.

To recover, an administrator must log on, archive the log (optional), clear the log, and reset this option as desired.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
If the computer is unable to record events to the security log, critical evidence or important troubleshooting information might not be available for review after a security incident.

The administrative burden of enabling this policy setting can be very high, especially if you also set the **Retention methodfor security log** to **Do not overwrite events (clear log manually)**. This setting turns a repudiation threat (a backup operator could deny that they backed up or restored data) into a denial-of-service threat, because a server can be forced to shut down if it is overwhelmed with logon events and other security events that are written to the security log. Additionally, because the shutdown is not graceful, it is possible that irreparable damage to the operating system, applications, or data could result. Although NTFS will guarantee that the file system's integrity will be maintained during a sudden system shutdown, it cannot guarantee that every data file for every application will still be in a usable form when the system is restarted.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Devices: Allow undock without having to log on**

The **Devices: Allow undock without having to log on** policy setting determines whether a user must log on to request permission to remove a portable computer from a docking station. Enabling this policy setting allows a user to undock a computer by pressing the portable computer's physical eject button. Disabling this policy setting requires that the user log on to get undocking permission. Only users who have the **Remove computer from docking station** user right can be assigned this permission.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

This policy setting should only be disabled for portable computers that cannot be mechanically undocked. Otherwise, enabling the policy setting allows a user to physically remove the computer when they cannot log in and the eject button does not work.

Enabling this policy setting means that anyone with physical access to a computer that has been placed in its docking station can remove the computer and possibly tamper with it. For computers that do not have docking stations, this policy setting has no impact. However, for users with a mobile computer that is normally docked while they are in the office, this policy setting will help lower the risk of equipment theft or a malicious user gaining physical access to these computers.

It is advisable to disable the **Devices: Allow undock without having to log on** policy setting. Users who have docked their computers will have to log on to the local console before they can undock their systems.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Devices: Allowed to format and eject removable media**

The **Devices: Allowed to format and eject removable media** policy setting determines who is allowed to format and eject removable media.

The possible values for this Group Policy setting are:

- Administrators.

- Administrators and Power Users.
- Administrators and Interactive Users.
- Not defined.

**Discussion**
Users can move removable disks to a different computer where they have administrative user rights and then take ownership of any file, assign themselves full control, and view or modify any file. The advantage of configuring this policy setting is diminished by the fact that most removable storage devices will eject media with the press of a button.

It is advisable to set **Allowed to format and eject removable media** to **Administrators**. Only administrators will be able to eject NTFS-formatted removable media.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Administrators |
| DC Effective Default Settings | Administrators |
| Member Server Effective Default Settings | Administrators |

**Devices: Prevent users from installing printer drivers**
For a computer to print to a network printer, the driver for that network printer must be installed on the local computer. The **Devices: Prevent users from installing printer drivers** policy setting determines who can install a printer driver as part of adding a network printer. When you set the value to **Enabled**, only Administrators and Power Users can install a printer driver as part of adding a network printer. Setting the value to **Disabled** allows any user to install a printer driver as part of adding a network printer. This setting prevents unprivileged users from downloading and installing an untrusted printer driver.

This setting has no impact if an administrator has configured a trusted path for downloading drivers. When using trusted paths, the print subsystem attempts to use the trusted path to download the driver. If the trusted path download succeeds, the driver is installed on behalf of any user. If the trusted path download fails, the driver is not installed and the network printer is not added.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Although it might be appropriate in some organizations to allow users to install printer drivers on their own workstations, this is not suitable for servers. Installing a printer driver on a server can cause the system to become less stable. Only administrators should have this user right on servers. A malicious user might deliberately try to damage the system by installing inappropriate printer drivers.

It is advisable to set **Devices: Prevent users from installing printer drivers** to **Enabled**. Only users in the Administrative, Power User, or Server Operator groups will be able to install printers on servers. If this policy setting is enabled, but the driver for a network printer already exists on the local computer, users can still add the network printer. This policy setting does not affect a user's ability to add a local printer.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Devices: Restrict CD-ROM access to locally logged-on user only**

The **Devices: Restrict CD-ROM access to locally logged-on user only** policy setting determines whether a CD-ROM will be accessible to both local and remote users simultaneously. Enabling this policy setting allows only the interactively logged-on user to access removable CD-ROM media. If this policy setting is enabled and no one is logged on interactively, the CD-ROM can be accessed over the network.

The possible values for these Group Policy settings are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

The security benefit of enabling this policy setting is small because it only prevents network users from accessing the drive when someone is logged on to the local console of the system at the same time. Additionally, CD drives are not automatically made available as network shares; administrators must deliberately choose to share the drive. This is important when administrators are installing software or copying data from a CD-ROM and they do not want network users to be able to execute the applications or view the data.

If this policy setting is enabled, users connecting to the server over the network will not be able to use any CD drives installed on the server whenever anyone is logged on to the local console of the server. Enabling this policy setting is not suitable for a system that serves as a CD jukebox for network users.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Devices: Restrict floppy access to locally logged-on user only**

The **Devices: Restrict floppy access to locally logged-on user only** policy setting determines whether removable floppy disks are accessible to both local and remote users simultaneously. Enabling this policy setting allows only the interactively logged-on user to access removable floppy disks. If this policy setting is enabled and no one is logged on interactively, the floppy disk can be accessed over the network.

The possible values for these Group Policy settings are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
The security benefit of enabling this policy setting is small because it only prevents network users from accessing the floppy disk drive when someone is logged on to the local console of the system at the same time. Additionally, floppy disk drives are not automatically made available as network shares; administrators must deliberately choose to share the drive. This becomes important when administrators are installing software or copying data from a floppy disk and they do not want network users to be able to execute the applications or view the data.

If this policy setting is enabled, users connecting to the server over the network will not be able to use any floppy disk drives installed on the server whenever anyone is logged on to the local console of the server.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Devices: Unsigned driver installation behavior**
The **Devices: Unsigned driver installation behavior** policy setting determines what happens when a user attempts to install a device driver that has not been certified by the Windows Hardware Quality Lab (WHQL) by means of Setup API.

The possible values for this Group Policy setting are:

- Silently succeed.
- Warn but allow installation.
- Do not allow installation.
- Not defined.

**Discussion**
This policy setting can prevent the installation of unsigned drivers or warn the administrator that an unsigned driver software is about to be installed. It can prevent Setup API from installing drivers that have not been certified to run on Windows Server 2003. This policy setting will not prevent a method used by some attack tools in which malicious .sys files are copied and registered to start as system services.

It is advisable to set **Unsigned driver installation behavior** to **Warn but allow installation**. This differs from the default value for this policy setting in Windows Server 2003, which is **Not defined**. A potential problem with setting **Unsigned driver installation behavior** to **Warn but allow installation** is that users who have sufficient user rights to install device drivers will still be able to install unsigned device drivers, which can result in stability problems for system servers. Another potential problem with this setting is that unattended installation scripts will fail if they attempt to install unsigned drivers.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Warn but allow installation |
| DC Effective Default Settings | Warn but allow installation |
| Member Server Effective Default Settings | Warn but allow installation |

**Domain controller: Allow server operators to schedule tasks**

The **Domain controller: Allow server operators to schedule tasks** policy setting determines whether users in the Server Operators group are allowed to submit jobs by means of the AT schedule service. This security option only affects the AT schedule facility, it does not affect the Task Scheduler service.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

Enabling this policy setting means that jobs created by server operators via the AT service will be executed in the context of the account that is running that service—by default, that is the local System account. This means that server operators can perform tasks that the System account is able to do but server operators would normally not be able to do, such as add their account to the local Administrators group.

The impact of enabling this policy setting should be small for most organizations. Users, including those in the Server Operators group, will still be able to create jobs by using the Task Scheduler Wizard, but those jobs will run in the context of the account the user authenticated with when they set up the job.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Domain controller: LDAP server signing requirements**

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The **Domain controller: LDAP server signing requirements** policy setting determines whether the LDAP server requires LDAP clients to negotiate data signing.

The possible values for this Group Policy setting are:

- None.

Data signing is not required in order to bind with the server. If the client requests data signing, the server supports it.

- Require signature.

  The LDAP data-signing option must be negotiated unless Transport Layer Security/Secure Sockets Layer (TLS/SSL) is being used.

- Not defined.

**Discussion**
Unsigned network traffic is susceptible to man-in-the-middle attacks, where an intruder captures packets between the server and the client and modifies them before forwarding them to the client. In the case of an LDAP server, this means that a malicious user can cause a client to make decisions based on false records from the LDAP directory. You can lower the risk of a malicious user accomplishing this in a corporate network by implementing strong physical security measures to protect the network infrastructure. Furthermore, implementing Internet Protocol security (IPsec) Authentication Header (AH) mode, which provides mutual authentication and packet integrity for IP traffic, can make all types of man-in-the-middle attacks extremely difficult.

It is advisable to set **Domain controller: LDAP server signing requirements** to **Require signature**. Clients that do not support LDAP signing will be unable to execute LDAP queries against the domain controllers. This means that all computers in your organization that are managed from Windows Server 2003 or Windows XP-based computers using NTLM authentication and the Windows Server 2003 Administration Tools Pack must either have Windows 2000 Service Pack 3 (SP3) or later installed.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | None |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | None |
| Member Server Effective Default Settings | Not defined |

**Domain controller: Refuse machine account password changes**
The **Domain controller: Refuse machine account password changes** policy setting determines whether or not a domain controller will accept password change requests for computer accounts.

The possible values for this Group Policy setting are:

- Enabled.

- Disabled.

- Not defined.

**Discussion**
Enabling this policy setting on all domain controllers in a domain prevents domain members from changing their computer account passwords. This, in turn, leaves those passwords susceptible to attack.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Domain member: Digitally encrypt or sign secure channel data (multiple related policy settings)**

**Note**

- Misuse of these policy settings is a common error that can cause data loss or problems with data access or security.

The following policy settings determine whether a secure channel can be established with a domain controller that is not capable of signing or encrypting secure channel traffic:

- Domain member: Digitally encrypt or sign secure channel data (always)
- Domain member: Digitally encrypt secure channel data (when possible)
- Domain member: Digitally sign secure channel data (when possible)

Setting **Domain member: Digitally encrypt or sign secure channel data (always)** to **Enabled** prevents establishing a secure channel with any domain controller that cannot sign or encrypt all secure channel data.

To protect authentication traffic from man-in-the-middle, replay, and other types of network attacks, Windows-based computers create a communication channel through NetLogon called secure channels. These channels authenticate computer accounts. They also authenticate user accounts when a remote user connects to a network resource and the user account exists in a trusted domain. This is called pass-through authentication, and it allows a computer running Windows that has joined a domain to have access to the user account database in its domain and in any trusted domains.

To enable the **Domain member: Digitally encrypt or sign secure channel data (always)** policy setting on a member workstation or server, all domain controllers in the domain that the member belongs to must be capable of signing or encrypting all secure-channel data. This means that all such domain controllers must be running Windows NT 4.0 Service Pack 6a (SP6a) or later.

Enabling the **Domain member: Digitally encrypt or sign secure channel data (always)** policy setting automatically enables the **Domain member: Digitally sign secure channel data (when possible)** policy setting.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

When a Windows Server 2003 system joins a domain, a computer account is created. After joining the domain, the computer uses the password for that account to create a secure channel with the domain controller for its domain every time that it restarts. Requests sent on the secure channel are authenticated—and sensitive information such as passwords are encrypted—but the integrity of the channel is not checked, and not all information is encrypted. If a system is set to always encrypt or sign secure channel data, a secure channel cannot be established with a domain controller that is not capable of signing or encrypting all secure channel traffic. If the computer is configured to encrypt or sign secure channel data when possible, a secure channel can be established, but the level of encryption and signing is negotiated.

The following values are advised:

- Set **Domain member: Digitally encrypt or sign secure channel data (always)** to **Enabled**.

- Set **Domain member: Digitally encrypt secure channel data (when possible)** to **Enabled**.
- Set **Domain member: Digitally sign secure channel data (when possible)** to **Enabled**.

Digital encryption and signing of the secure channel is a good idea where it is supported. The secure channel protects domain credentials as they are sent to the domain controller. However, only Windows NT 4.0 SP6a or later supports digital encryption and signing of the secure channel. Windows 98 Second Edition clients do not support it (unless they have the dsclient installed). Thus, you cannot enable the **Domain member: Digitally encrypt or sign secure channel data (always)** policy setting on domain controllers that support Windows 98 clients as members of the domain. The potential impact of such a setting can include making it impossible to create or delete trust relationships with earlier versions, disabling logons from earlier-version clients, and not being able to authenticate other domains' users from a trusted domain that is running an earlier version of Windows.

You can enable this policy setting after you eliminate all Windows 95, Windows 98, and Windows Millennium Edition clients from the domain and upgrade all Windows NT 4.0 servers and domain controllers from trusted or trusting domains to Windows NT 4.0 SP6a. You can enable the other two policy settings, **Domain member: Digitally encrypt secure channel data (when possible)** and **Domain member: Digitally sign secure channel data (when possible)**, on all computers in the domain that support these policy settings without affecting earlier-version clients and applications.

**Domain member: Digitally encrypt or sign secure channel data (always)**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Enabled |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Domain member: Digitally encrypt secure channel data (when possible)**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |

| | |
|---|---|
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Domain member: Digitally sign secure channel data (when possible)**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Domain member: Disable machine account password changes**

The **Domain member: Disable machine account password changes** policy setting determines whether a domain member periodically changes its computer account password. Setting its value to **Enabled** prevents the domain member from changing its computer account password. Setting it to **Disabled** allows the domain member to change its computer account password as specified by the value of the **Domain member: Maximum machine account password age** policy setting, which is every 30 days by default.

Do not enable this policy setting. Computer account passwords are used to establish secure channel communications between members and domain controllers and—within the domain—between the domain controllers themselves. After it is established, the secure channel transmits sensitive information that is necessary for making authentication and authorization decisions.

Do not use this policy setting in an attempt to support dual-boot scenarios that use the same computer account. If you want to dual-boot two installations that are joined to the same domain, give the two installations different computer names. This policy setting was added to Windows to make it easier for organizations that stockpile prebuilt computers that are put into production months later; those computers do not have to be rejoined to the domain.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

The default configuration for computers running Windows Server 2003 that belong to a domain is that they are automatically required to change the passwords for their accounts every 30 days. Disabling this feature causes computers running Windows Server 2003 to retain the same passwords as their computer accounts. Computers that are no longer able to automatically change their account password are at risk of a malicious user determining the password for the system's domain account.

Verify that the **Domain member: Disable machine account password changes** option is set to **Disabled**.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Domain member: Maximum machine account password age**

The **Domain member: Maximum machine account password age** policy setting determines the maximum allowable age for a computer account password. This policy setting also applies to computers running Windows 2000, but it is not available through the Security Configuration Manager tools on those computers.

The possible values for this Group Policy setting are:

- A user-defined number of days from 0 through 999.
- Not defined.

**Discussion**

In Active Directory–based domains, each computer has an account and password just like every user. By default, the domain members automatically change their domain password every 30 days. Increasing this interval significantly, or setting it to **0** so that the computers no longer change their passwords, gives a malicious user more time to undertake a brute-force password-guessing attack against one of the computer accounts.

It is often advisable to set **Domain member: Maximum machine account password age** to about **30** days.

Some organizations prebuild computers and then store them for later use or ship them to remote locations. If the computer's account has expired, it will no longer be able to authenticate with the domain. Computers that cannot authenticate with the domain must be removed from the domain and rejoined to it. For this reason, some organizations might want to create a special OU for computers that are prebuilt and configure the value for this policy setting to a larger number of days.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 30 days |
| DC Effective Default Settings | 30 days |
| Member Server Effective Default Settings | 30 days |

**Domain member: Require strong (Windows 2000 or later) session key**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The **Domain member: Require strong (Windows 2000 or later) session key** policy setting determines whether a secure channel can be established with a domain controller that is not capable of encrypting secure channel traffic with a strong, 128-bit session key. Enabling this policy setting prevents establishing a secure channel with any domain controller that cannot encrypt secure channel data with a strong key. Disabling this policy setting allows 64-bit session keys.

**Note**

- To enable this policy setting on a member workstation or server, all domain controllers in the domain that the member belongs to must be capable of encrypting secure channel data with a strong, 128-bit key. This means that all such domain controllers must be running Windows 2000 or Windows Server 2003.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Session keys used to establish secure channel communications between domain controllers and member computers are much stronger in Windows 2000 and Windows Server 2003 than they were in previous Microsoft operating systems.

Whenever possible, you should take advantage of these stronger session keys to help protect secure channel communications from eavesdropping and session-hijacking network attacks. Eavesdropping is a form of hacking in which network data is read or altered in transit. The data can be modified to hide or change the name of the sender, or it can be redirected.

It is advisable to set **Domain member: Require strong (Windows 2000 or later) session key** to **Enabled**. Enabling this policy setting ensures that all outgoing secure channel traffic will require a strong, Windows 2000 or Windows Server 2003 encryption key. Disabling this policy setting requires that key strength be negotiated. Only enable this option if the domain controllers in all trusted domains support strong keys. By default, this value is disabled.

You will not be able to join computers with this policy setting enabled to Windows NT 4.0 domains, nor will you be able to join computers that do not support this policy setting to domains where the domain controllers have this policy setting enabled.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Interactive logon: Do not display last user name**
The **Interactive logon: Do not display last user name** policy setting determines whether the **Log on to Windows** dialog box displays the name of the last user who logged on to the computer. If the value is **Enabled**, the **Log on to Windows** dialog box does not display the name of the last user who logged on. If it is **Disabled**, the dialog box does display the name of the last user to log on.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
A malicious user with access to the console—for example, someone with physical access or someone who is able to connect to the server via Terminal Services—can view the name of the last user who logged on to the server. The attacker might then

attempt to log on to the server by guessing the password, using a dictionary, or by brute-force attack.

It is advisable to set **Do not display last user name in logon screen** to **Enabled**. Users will always have to type their user names when they log on to the servers.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Interactive logon: Do not require CTRL+ALT+DEL**
The **Interactive logon: Do not require CTRL+ALT+DEL** policy setting determines whether pressing CTRL+ALT+DELETE is required before a user can log on. Enabling this policy setting on a computer allows a user to log on without pressing CTRL+ALT+DELETE. Disabling this policy setting requires users to press CTRL+ALT+DELETE before logging on, unless they are using a smart card for Windows logon. A smart card is a tamper-proof device that stores security information.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Microsoft developed this feature to make it easier for users with certain types of physical impairments to log on to computers running Windows; however, not having to press the CTRL+ALT+DELETE key combination leaves users susceptible to attacks that attempt to intercept their passwords. Requiring CTRL+ALT+DELETE before users log on ensures that users are communicating by means of a trusted path when entering their passwords.

A malicious user might install a Trojan horse program that looks like the standard Windows logon dialog box and capture a user's password. The attacker can then log on to the compromised account with whatever level of user rights that user has.

It is advisable to set **Disable CTRL+ALT+DEL requirement for logon** to **Disabled**. Unless they are using a smart card to log on, users will have to simultaneously press three keys before the logon dialog box appears.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Interactive logon: Message text for users attempting to log on**

The **Interactive logon: Message text for users attempting to log on** and **Interactive logon: Message title for users attempting to log on** policy settings are closely related. **Interactive logon: Message text for users attempting to log on** specifies a text message to be displayed to users when they log on. **Interactive logon: Message title for users attempting to log on** specifies a title to appear in the title bar of the window that contains the text message. This text is often used for legal reasons — for example, to warn users about the ramifications of misusing company information, or to warn them that their actions might be audited.

Windows XP Professional adds support for configuring logon banners that can exceed 512 characters in length and can also contain carriage return/line feed sequences. However, Windows 2000 clients cannot interpret and display message text that is created by Windows XP Professional computers. You must use a Windows 2000 computer to create a logon message that applies to Windows 2000 computers. If you inadvertently create a logon message by using a Windows XP Professional computer, and you discover that it is not displayed properly on Windows 2000 computers, do the following:

- Undefine the setting.
- Redefine the setting by using a Windows 2000 computer.

Simply changing a Windows XP Professional–defined logon message setting by using a Windows 2000 computer does not work. The setting must be undefined first.

The possible values for these Group Policy settings are:

- User-defined text.
- Not defined.

**Discussion**

Not using this warning-message policy setting leaves your organization legally vulnerable to trespassers who unlawfully penetrate your network. Legal precedents have established that organizations that display warnings to users who connect to their servers over a network have a higher rate of successfully prosecuting trespassers.

It is advisable to set **Message text for users attempting to log on** to the following value:

This system is restricted to authorized users. Individuals attempting unauthorized access will be prosecuted. If unauthorized, terminate access now! Clicking OK indicates your acceptance of the information in the background.

It is advisable to set **Message title for users attempting to log on** to a value similar to the following:

IT IS AN OFFENSE TO CONTINUE WITHOUT PROPER AUTHORIZATION.

Any warning that you display should first be approved by representatives from your organization's legal and human resources departments.

When these policy settings are configured, users will see a dialog box before they can log on to the server console.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Interactive logon: Number of previous logons to cache (in case domain controller is not available)**

The **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** policy setting determines whether a user can log on to a Windows domain by using cached account information. Logon information for domain accounts can be cached locally so that, if a domain controller cannot be contacted on subsequent logons, a user can still log on. This policy setting determines the number of unique users whose logon information is cached locally.

If a domain controller is unavailable and a user's logon information is cached, the user is prompted with the following message:

A domain controller for your domain could not be contacted. You have been logged on using cached account information. Changes to your profile since you last logged on might not be available.

If a domain controller is unavailable and a user's logon information is not cached, the user is prompted with this message:

The system cannot log you on now because the domain <DOMAIN_NAME> is not available.

The possible values for this Group Policy setting are:

- A user-defined number from 0 through 50.
- Not defined.

**Discussion**

The value of this policy setting indicates the number of users whose logon information the server caches locally. If the value is **10**, the server caches logon information for 10 users. When an eleventh user logs on to the computer, the server overwrites the oldest cached logon session.

Users who access the server console will have their logon credentials cached on that server. A malicious user who is able to access the file system of the server can locate this cached information and use a brute-force attack to determine user passwords. Windows mitigates this type of attack by encrypting the information and keeping the cached credentials in the system's registries, which are spread across numerous physical locations.

It is advisable to set **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** to **0**. Setting this value to 0 disables the local caching of logon information. Additional countermeasures include enforcing strong password policies and physically securing the computers. If the value is set to **0**, users will be unable to log on to any computers if there is no domain controller available to authenticate them. Organizations might want to set **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** to **2** for end-user systems, especially for mobile users. Setting this value to 2 means that the user's logon information will still be in the cache even if a member of the IT department has recently logged on to their computer to perform system maintenance. This way, those users will be able to log on to their computers when they are not connected to the corporate network.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 10 logons |
| DC Effective Default Settings | 10 logons |
| Member Server Effective Default Settings | 10 logons |

**Interactive logon: Prompt user to change password before expiration**

The **Interactive logon: Prompt user to change password before expiration** policy setting determines how many days in advance users are warned that their passwords are about to expire. With this advance warning, the user has time to construct a password that is sufficiently strong.

The possible values for this Group Policy setting are:

- A user-defined number of days from 1 through 999.

- Not defined.

**Discussion**
It is recommended that user passwords be configured to expire periodically. Users will need warning their passwords are going to expire, or they might inadvertently get locked out of the system. This could lead to confusion for users who access the network locally, or make it impossible for users who access the network via dial-up or virtual private network (VPN) connections to log on.

It is advisable to set **Prompt user to change password before expiration** to **14** days. When their password expiration date is 14 or fewer days away, users will see a dialog box each time they log on to the domain.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 14 days |
| DC Effective Default Settings | 14 days |
| Member Server Effective Default Settings | 14 days |

**Interactive logon: Require Domain Controller authentication to unlock workstation**
Unlocking a locked computer requires logon information. For domain accounts, the **Interactive logon: Require Domain Controller authentication to unlock workstation** policy setting determines whether it is necessary to contact a domain controller to unlock a computer. Enabling this policy setting requires a domain controller to authenticate the domain account that is being used to unlock the computer. Disabling this policy setting allows a user to unlock the computer without the computer verifying the logon information with a domain controller. However, if **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** is set to a value greater than zero, the user's cached credentials will be used to unlock the system.

This policy setting also applies to computers running Windows 2000, but it is not available through the Security Configuration Manager on those computers.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
The computer caches the credentials of any users that have been authenticated locally in memory. The computer uses these cached credentials to authenticate anyone who attempts to unlock the console.

When cached credentials are used, any changes that have recently been made to the account — such as user rights assignments, account lockout, or the account being disabled — are not considered or applied after this authentication process. This means not only that user rights are not updated, but more importantly that disabled accounts are still able to unlock the console of the system.

It is advisable to set **Interactive logon: Require Domain Controller authentication to unlock workstation** to **Enabled** and set **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** to **0**. When the console of a computer is locked, either by a user or automatically by a screen saver time-out, the console can only be unlocked if the user is able to reauthenticate to the domain controller. If no domain controller is available, users cannot unlock their workstations.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Interactive logon: Require smart card**
The **Interactive logon: Require smart card** policy setting requires users to log on to a computer by using a smart card.

This policy setting also applies to Windows 2000 computers, but it is not available through the Security Configuration Manager on those computers.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Requiring users to use long, complex passwords for authentication enhances network security, especially if the users must change their passwords regularly. This reduces the chance that a malicious user will be able to guess a user's password via a brute-force attack. Using smart cards rather than passwords for authentication dramatically increases security because, with today's technology, it is nearly impossible for a malicious user to impersonate another user. Smart cards that require personal identification numbers (PINs) provide two-factor authentication: the user who attempts to log on must possess the smart card and know its PIN. A malicious user who captures the authentication traffic between the user's computer and the domain controller will find it extremely difficult to decrypt the traffic: even if they do, the next time the user logs on to the network, a new session key will be generated for encrypting traffic between the user and the domain controller.

It is advisable to set **Interactive logon: Require smart card** to **Enabled**. All users will have to use smart cards to log on to the network. This means that the organization must have a reliable public key infrastructure (PKI) in place, and provide smart cards and smart card readers for all users. These are significant challenges, because planning for and deploying these technologies requires expertise and resources; however, Windows Server 2003 includes Certificate Services, a highly advanced server for implementing and managing certificates that—when combined with Windows XP—includes features such as automatic user and computer enrollment and renewal.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Interactive logon: Smart card removal behavior**

The **Interactive logon: Smart card removal behavior** policy setting determines what happens when the smart card for a logged-on user is removed from the smart card reader.

The possible values for this Group Policy setting are:

- No Action.
- Lock Workstation.
- Force Logoff.
- Not defined.

**Discussion**

If smart cards are used for authentication, the computer should automatically lock itself when the card is removed—that way, if users forget to manually lock their workstations when they are away from them, malicious users cannot gain access.

It is advisable to set **Interactive logon: Smart card removal behavior** to **Lock Workstation**. If you select **Lock Workstation** in the property sheet for this policy setting, the workstation is locked when the smart card is removed. This allows users to leave the area, take their smart card with them, and still maintain a protected session.

If you select **Force Logoff** in the property sheet for this policy setting, the user is automatically logged off when the smart card is removed. Users will have to reinsert their smart cards and reenter their PINs when they return to their workstations.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | No Action |
| DC Effective Default Settings | No Action |
| Member Server Effective Default Settings | No Action |

**Microsoft network client and server: Digitally sign communications (four related policy settings)**

**Note**

- Misuse of these policy settings is a common error that can cause data loss or problems with data access or security.

There are four separate policy settings relating to digitally signing SMB communications: **Microsoft network client: Digitally sign communications (always)**, **Microsoft network server: Digitally sign communications (always)**, **Microsoft network client: Digitally sign communications (if server agrees)**, and **Microsoft network server: Digitally sign communications (if client agrees**).

The possible values for these Group Policy settings are:

- Enabled.
- Disabled.
- Not defined.

Implementing digital signing in high-security networks helps to prevent the impersonation of clients and servers. This type of impersonation is known as "session hijacking," in which a malicious user who has access to the same network as the client or

server interrupts, ends, or steals a session in progress. Attackers can potentially intercept and modify unsigned SMB packets, then modify the traffic and forward it so the server performs undesirable actions. Alternatively, the attacker can pose as the server or client after a legitimate authentication and gain unauthorized access to data.

SMB is the resource-sharing protocol supported by many Microsoft operating systems. It is the basis of NetBIOS and many other protocols. SMB signing authenticates both the user and the server hosting the data. If either side fails the authentication process, data transmission will not take place.

An alternative countermeasure that can protect all network traffic is to implement digital signatures by using IPsec. There are hardware-based accelerators for IPsec encryption and signing that can be used to minimize the performance impact on the servers' CPUs. There are no such accelerators available for SMB signing.

The following values are advisable

- Set **Microsoft network client: Digitally sign communications (always)** to **Disabled**.
- Set **Microsoft network server: Digitally sign communications (always)** to **Disabled**.
- Set **Microsoft network client: Digitally sign communications (if server agrees)** to **Enabled**.
- Set **Microsoft network server: Digitally sign communications (if client agrees)** to **Enabled**.

Some sources recommend setting all of these policy settings to **Enabled**, but enabling them can cause slower performance on client computers and prevent them from communicating with legacy SMB applications and operating systems.

The Windows 2000, Windows Server 2003, and Windows XP Professional file and print sharing protocol SMB supports mutual authentication, which closes session-hijacking attacks and supports message authentication, thus preventing man-in-the-middle attacks. SMB signing provides this authentication by placing a digital signature into each SMB, which is then verified by both the client and the server. Implementing SMB signing can cause a noticeable performance impact as a result of having to sign and verify each packet between the servers. Configuring computers to ignore all unsigned SMB communications prevents legacy applications and operating systems from connecting. Completely disabling all SMB signing leaves the computers vulnerable to session-hijacking attacks.

#### Microsoft network client: Digitally sign communications (always)

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

#### Location

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

#### Microsoft network client: Digitally sign communications (if server agrees)

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

#### Location

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Microsoft network server: Digitally sign communications (always)**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Enabled |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Not defined |

**Microsoft network server: Digitally sign communications (if client agrees)**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Enabled |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Not defined |

**Microsoft network client: Send unencrypted password to third-party SMB servers**

Enabling the **Microsoft network client: Send unencrypted password to third-party SMB servers** policy setting allows the

SMB redirector to send plaintext passwords to non-Microsoft SMB servers that do not support password encryption during authentication.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Enabling this policy setting allows the server to transmit passwords in plaintext across the network to other systems that offer SMB services. These others systems might not use any of the SMB security mechanisms included with Windows Server 2003.

It is advisable to set **Microsoft network client: Send unencrypted password to connect to third - party SMB servers** to **Disabled**. Some very old applications and operating systems such as MS-DOS, Windows for Workgroups 3.11, and Windows 95 might not be able to communicate with the servers in your organization via the SMB protocol.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Microsoft network server: Amount of idle time required before suspending session**
The **Microsoft network server: Amount of idle time required before suspending session** policy setting determines the amount of continuous idle time that must pass in an SMB session before the session is suspended due to inactivity. Administrators can use this policy setting to control when a computer suspends an inactive SMB session. The session is automatically reestablished when client activity resumes.

For this policy setting, a value of **0** means to disconnect an idle session as quickly as is reasonably possible. The maximum value is 99,999, which is 208 days—in effect, this disables the policy setting.

The possible values for this Group Policy setting are:

- A user-defined number of minutes from 0 through 99,999.
- Not defined.

**Discussion**
Each SMB session consumes server resources. Establishing numerous null sessions will slow or possibly crash the server. A malicious user might repeatedly establish SMB sessions until the server stops responding; at this point, SMB services will become slow or unresponsive.

It is advisable to set **Microsoft network server: Amount of idle time required before disconnecting session** to **15** minutes. There will be little impact because SMB sessions will be reestablished automatically if the client resumes activity.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 15 minutes |
| DC Effective Default Settings | 15 minutes |
| Member Server Effective Default Settings | 15 minutes |

**Microsoft network server: Disconnect clients when logon hours expire**

The **Microsoft network server: Disconnect clients when logon hours expire** policy setting determines whether to disconnect users who are connected to the local computer outside their user account's valid logon hours. This policy setting affects the SMB component. Enabling this policy setting causes client sessions with the SMB service to be forcibly disconnected when the client's logon hours expire. Disabling this policy setting maintains an established client session after the client's logon hours have expired. When you enable this policy setting, you should also enable **Network security: Force logoff when logon hours expire**.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

If your organization has configured logon hours for users, it makes sense to enable this policy setting; if it is disabled, users who are assumed to be unable to access network resources outside their logon hours might actually be able to continue to use those resources with sessions that were established during allowed hours.

Enable the **Microsoft network server: Disconnect clients when logon hours expire** policy setting. If logon hours are not used in your organization, enabling this policy setting will have no impact. If logon hours are used, existing user sessions will be forcibly terminated when their logon hours expire.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Network access: Allow anonymous SID/Name translation**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The **Network access: Allow anonymous SID/Name translation** policy setting determines whether an anonymous user can request SID attributes for another user.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
If this policy setting is enabled, a user might use the well-known Administrators SID to get the real name of the built-in Administrator account, even if the account has been renamed. That person might then use the account name to initiate a brute-force password-guessing attack.

It is advisable to set **Network access: Allow anonymous SID/Name translation** to **Disabled**. This is the default value on member computers; therefore, it will have no impact on them. The default value for domain controllers is **Enabled**. Disabling this policy setting on domain controllers means that legacy systems might be unable to communicate with Windows Server 2003-based domains. For example, the following systems might not work:

- Windows NT 4.0–based Remote Access Service servers.
- Microsoft SQL Server running on Windows NT 3.51–based or Windows NT 4.0–based computers.
- Remote Access Service or Microsoft SQL Server running on Windows 2000–based computers that are located in Windows NT 3.51 domains or Windows NT 4.0 domains.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Disabled |

**Network access: Do not allow anonymous enumeration of SAM accounts**
### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The **Network access: Do not allow anonymous enumeration of SAM accounts** policy setting determines which additional permissions will be assigned for anonymous connections to the computer. Windows allows anonymous users to perform certain activities, such as enumerating the names of domain accounts and network shares. This is convenient, for example, when an administrator wants to give access to users in a trusted domain that does not maintain a reciprocal trust. However, even with this policy setting enabled, anonymous users will still have access to any resources that have permissions that explicitly include the special built-in group ANONYMOUS LOGON.

This policy setting has no impact on domain controllers.

In Windows 2000, a similar policy setting named **Additional Restrictions for Anonymous Connections** managed a registry entry **RestrictAnonymous**, located in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA key. In Windows Server 2003, the policy settings **Network access: Do not allow anonymous enumeration of SAM accounts** and **Network access: Do not allow anonymous enumeration of SAM accounts and shares** replace the Windows 2000 policy setting. They manage the registry entries **RestrictAnonymousSAM** and **RestrictAnonymous**, respectively, both located in the HKLM\System\CurrentControlSet\Control\Lsa\ key.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

A malicious user can anonymously list account names and use the information to attempt to guess passwords or perform social-engineering attacks. "Social engineering" is a term for tricking people into revealing their password or some other form of security information.

If you set **Network access: Do not allow anonymous enumeration of SAM accounts** to **Enabled**, it will be impossible to establish trusts with Windows NT 4.0–based domains. This value will also cause problems with earlier-version clients such as Windows NT 3.51 and Windows 95 that are trying to use resources on the server.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Network access: Do not allow anonymous enumeration of SAM accounts and shares**

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The **Network access: Do not allow anonymous enumeration of SAM accounts and shares** policy setting determines which additional permissions will be assigned for anonymous connections to the computer. Windows allows anonymous users to perform certain activities, such as enumerating the names of domain accounts and network shares. This is convenient, for example, when an administrator wants to give access to users in a trusted domain that does not maintain a reciprocal trust. However, even with this policy setting enabled, anonymous users will still have access to any resources that have permissions that explicitly include the special built-in group ANONYMOUS LOGON.

This policy setting has no impact on domain controllers.

In Windows 2000, a similar policy setting named **Additional Restrictions for Anonymous Connections** managed a registry value **RestrictAnonymous**, located in the HKLM\SYSTEM\CurrentControlSet\Control\LSA key. In Windows Server 2003, the policy settings **Network access: Do not allow anonymous enumeration of SAM accounts** and **Network access: Do not allow anonymous enumeration of SAM accounts and shares** replace the Windows 2000 policy setting. They manage the registry entries **RestrictAnonymousSAM** and **RestrictAnonymous**, respectively, both located in the HKLM\System\CurrentControlSet\Control\Lsa\ key.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
If this policy setting is disabled, an unauthorized user could anonymously list account names and use the information to attempt to guess passwords or perform social-engineering attacks.

It is advisable to set **Network access: Do not allow anonymous enumeration of SAM accounts and shares** to **Enabled**. Note that when this is enabled, it will be impossible to establish trusts with Windows NT 4.0–based domains. This policy setting will also cause problems with earlier-version clients such as Windows NT 3.51 and Windows 95 that are trying to use resources on the server.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Network access: Do not allow storage of credentials or .NET Passports for network authentication**
The **Network access: Do not allow storage of credentials or .NET Passports for network authentication** policy setting determines whether Stored User Names and Passwords saves passwords or credentials for later use when it gains domain authentication. Enabling this policy setting prevents Stored User Names and Passwords from storing passwords and credentials.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Passwords cached by Stored User Names and Passwords can be accessed by a user when logged on to the computer. A problem can arise if the user unknowingly executes hostile code that reads the passwords and forwards them to another, malicious user.

It is advisable to set **Network access: Do not allow storage of credentials or .NET Passports for network authentication** to **Enabled**. Users will be forced to enter passwords whenever they log on to their Passport account or other network resources that are not accessible to their domain account. This will not have an impact when the user visits Web sites that require him to log on, because enabling this policy setting does not prevent passwords from being saved when the user requests that they be. Enabling this policy setting should have no impact on users when they access network resources that have been configured to allow access with their Active Directory–based domain account.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |

| | |
|---|---|
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Network access: Let Everyone permissions apply to anonymous users**

The **Network access: Let Everyone permissions apply to anonymous users** policy setting determines which additional permissions are assigned for anonymous connections to the computer. Enabling this policy setting allows anonymous users to perform certain activities, such as enumerating the names of domain accounts and network shares. This is convenient, for example, when an administrator wants to give access to users in a trusted domain that does not maintain a reciprocal trust. By default, the token created for anonymous connections does not include the Everyone SID; therefore, permissions assigned to the Everyone group do not apply to anonymous users. Enabling this policy setting adds the Everyone SID to the token created for anonymous connections, which allows anonymous users to access any resource for which the Everyone group has been given permissions.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

A malicious user might anonymously list account names and shared resources and use the information to attempt to guess passwords or perform social-engineering attacks. It is advisable to set **Network access: Let Everyone permissions apply to anonymous users** to **Disabled**.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Network access: Named Pipes that can be accessed anonymously**

The **Network access: Named Pipes that can be accessed anonymously** policy setting determines which communication sessions, or pipes, will have attributes and permissions that allow anonymous access.

The possible values for this Group Policy setting are:

- A user-defined list of shares.
- Not defined.

**Discussion**

Restricting access over named pipes such as COMNAP and LOCATOR helps prevent unauthorized access to the network. The following table lists default named pipes and their purpose.

### Default Named Pipes That Are Accessible Anonymously

| Named Pipe | Purpose |
|---|---|

| | |
|---|---|
| COMNAP | SnaBase named pipe. Systems Network Architecture (SNA) is a collection of networking protocols originally developed for IBM mainframe computers. |
| COMNODE | SNA Server named pipe. |
| SQL\QUERY | Default named pipe for SQL Server. |
| SPOOLSS | Named pipe for the Print Spooler service. |
| EPMAPPER | End Point Mapper named pipe. |
| LOCATOR | RPC Locator service named pipe. |
| TrkWks | Distributed Link Tracking Client named pipe. |
| TrkSvr | Distributed Link Tracking Server named pipe. |

It is advisable to set **Network access: Named Pipes that can be accessed anonymously** to a null value; that is, enable the policy setting, but do not enter named pipes in the text box. This will disable null session access over named pipes, and applications that rely on this feature or on unauthenticated access to named pipes will no longer function. For example, with Microsoft Commercial Internet System 1.0, the Internet Mail Service runs under Inetinfo.exe. Inetinfo.exe starts in the context of the System account. When Internet Mail Service needs to query the Microsoft SQL Server database, it uses the System account, which uses null credentials to access a SQL pipe on the computer running SQL Server.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | COMNAP,COMNODE, SQL\QUERY, SPOOLSS, EPMAPPER, LOCATOR,TrkWks,TrkSvr |
| DC Effective Default Settings | COMNAP,COMNODE, SQL\QUERY, SPOOLSS, EPMAPPER, LOCATOR,TrkWks,TrkSvr |
| Member Server Effective Default Settings | COMNAP,COMNODE, SQL\QUERY, SPOOLSS, EPMAPPER, LOCATOR,TrkWks,TrkSvr |

**Network access: Remotely accessible registry paths**
The **Network access: Remotely accessible registry paths** policy setting determines which registry paths will be accessible after referencing the WinReg key to determine access permissions to those paths.

The possible values for this Group Policy setting are:

- A user-defined list of paths.
- Not defined.

**Discussion**
The registry is a database for computer configuration information, much of which is sensitive. A malicious user can use the registry to facilitate unauthorized activities. To reduce the risk of this happening, suitable access control lists (ACLs) are assigned throughout the registry to help protect it from access by unauthorized users.

It is advisable to set **Network access: Remotely accessible registry paths** to a null value; that is, enable the policy setting but do not enter any paths in the text box. Remote management tools, such as the Microsoft Baseline Security Analyzer and Microsoft Systems Management Server (SMS), require remote access to the registry. Removing the default registry paths from the list of accessible paths might cause these and other management tools to fail.

If you do want to allow remote access, you must also enable the Remote Registry service.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | A (see below) |
| DC Effective Default Settings | A (see below) |
| Member Server Effective Default Settings | A (see below) |

A = System\CurrentControlSet\Control\ProductOptions, System\CurrentControlSet\Control\Server Applications, Software\Microsoft\Windows NT\CurrentVersion

**Network access: Remotely accessible registry paths and sub-paths**

The **Network access: Remotely accessible registry paths and sub-paths** policy setting determines which registry paths and subpaths will be accessible after referencing the WinReg key to determine access permissions to those paths.

The possible values for this Group Policy setting are:

- A user-defined list of paths.
- Not defined.

**Discussion**

The registry is a database for computer configuration information, much of which is sensitive. A malicious user can use it to facilitate unauthorized activities. The chance of this happening is reduced by the fact that the default ACLs assigned throughout the registry are fairly restrictive, and they help protect it from access by unauthorized users.

It is advisable to set **Networkaccess: Remotely accessible registry paths and sub-paths** to a null value; that is, enable the policy setting but do not enter any paths in the text box. Remote management tools such as the Microsoft Baseline Security Analyzer and SMS require remote access to the registry. Removing the default registry paths from the list of accessible paths might cause these and other management tools to fail.

If you do want to allow remote access, you must also enable the Remote Registry service.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | A (see below) |
| DC Effective Default Settings | A (see below) |
| Member Server Effective Default Settings | A (see below) |

A = System\CurrentControlSet\Control\Print\Printers, System\CurrentControlSet\Services\Eventlog, Software\Microsoft\OLAP Server, Software\Microsoft\Windows NT\CurrentVersion\Print, Software\Microsoft\Windows NT\CurrentVersion\Windows, System\CurrentControlSet\Control\ContentIndex, System\CurrentControlSet\Control\Terminal Server, System\CurrentControlSet\Control\Terminal Server\UserConfig, System\CurrentControlSet\Control\Terminal Server\DefaultUserConfiguration, Software\Microsoft\Windows NT\CurrentVersion\Perflib, System\CurrentControlSet\Services\SysmonLog

**Network access: Restrict anonymous access to Named Pipes and Shares**

When enabled, the **Network access: Restrict anonymous access to Named Pipes and Shares** policy setting restricts

anonymous access to shares and pipes to the settings for **Network access: Named pipes that can be accessed anonymously** and **Network access: Shares that can be accessed anonymously**. This setting controls null session access to shares on your computers by adding the registry entry **RestrictNullSessAccess** with the value **1** to the registry key HKLM\System\CurrentControlSet\Services\LanManServer\Parameters. This registry entry toggles null session shares on or off to determine whether the Server service restricts access to clients that have logged on to the System account without user name and password authentication.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Null sessions are a weakness that can be exploited through the various shares on the computers in your environment.

It is advisable to set **Network access: Restrict anonymous access to Named Pipes and Shares** to **Enabled**. Enabling this policy setting restricts null session access to unauthenticated users to all server pipes and shares except those listed in the **NullSessionPipes** and **NullSessionShares** registry entries.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Network access: Shares that can be accessed anonymously**
The **Network access: Shares that can be accessed anonymously** policy setting determines which network shares anonymous users can access.

The possible values for this Group Policy setting are:

- A user-defined list of shares.
- Not defined.

**Discussion**
Enabling this policy setting is very dangerous. Any shares listed can be accessed by any network user. This might lead to the exposure or corruption of sensitive corporate data.

It is advisable to set **Network access: Shares that can be accessed anonymously** to a null value. There should be little impact because this is the default value. All users will have to be authenticated before they can access shared resources on the server.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|

| | |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | COMCFG,DFS$ |
| DC Effective Default Settings | COMCFG,DFS$ |
| Member Server Effective Default Settings | COMCFG,DFS$ |

**Network access: Sharing and security model for local accounts**

The **Network access: Sharing and security model for local accounts** policy setting determines how network logons that use local accounts are authenticated. If it is set to **Classic - local users authenticate as themselves**, network logons that use local account credentials authenticate by using those credentials. If it is set to **Guest only - local users authenticate as Guest**, network logons that use local accounts are automatically mapped to the Guest account. The Classic model allows fine control over access to resources. By using the Classic model, you can assign different types of access to different users for the same resource. Using the Guest-only model treats all users equally. All users authenticate as Guest, and they all receive the same level of access to a given resource, which can be either Read Only or Modify.

The default value for this policy setting on stand-alone Windows XP Professional is **Guest only**. The default for Windows XP systems joined to a domain and Windows Server 2003 systems is **Classic**. This policy setting does not affect network logons that use domain accounts, nor does it affect interactive logons that are performed remotely by using services such as Telnet or Terminal Services.

When the computer is not joined to a domain, this policy setting also tailors the **Sharing** and **Security** tabs in Windows Explorer to correspond to the sharing and security model that is being used. This policy setting has no effect on Windows 2000 computers.

The possible values for this Group Policy setting are:

- Classic - local users authenticate as themselves.
- Guest only - local users authenticate as Guest.
- Not defined.

**Discussion**

When the value of this policy setting is **Guest only - local users authenticate as Guest**, any user who can access your computer over the network does so with Guest user rights. This means that they will probably be unable to write to those shares. Although this does increase security, it makes it impossible for authorized users to access shared resources on those systems. When the value is **Classic - local users authenticate as themselves**, local accounts must be password-protected; otherwise, anyone can use those user accounts to access shared system resources.

For network servers, set **Network access: Sharing and security model for local accounts** to **Classic - local users authenticate as themselves**. On end-user systems, set it to **Guest only - local users authenticate as Guest**.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Classic (local users authenticate as themselves) |
| DC Effective Default Settings | Classic (local users authenticate as themselves) |
| Member Server Effective Default Settings | Classic (local users authenticate as themselves) |

**Network security: Do not store LAN Manager hash value on next password change**

The **Network security: Do not store LAN Manager hash value on next password change** policy setting determines whether, at the next password change, LAN Manager will be prevented from storing hash values for the new password.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
By attacking the Security Accounts Manager file, attackers can potentially gain access to user names and password hashes. Attackers can use a password-cracking tool to determine what the password is. After they have access to this information, they can use it to gain access to resources on your network by impersonating users. Enabling this policy setting will not prevent these types of attacks, but it will make them much more difficult.

It is advisable to set **Network security: Do not store LAN Manager hash value on next password change** to **Enabled**. Require all users to set new passwords the next time they log on to the domain so that LAN Manager hashes are removed. Legacy operating systems—such as Windows 95, Windows 98, and Windows Millennium Edition—will fail, as will some non-Microsoft applications.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Network security: Force logoff when logon hours expire**
The **Force logoff when logon hours expire** policy setting determines whether to disconnect users who are connected to the local computer outside their user account's valid logon hours. This policy setting affects the SMB component. Enabling this policy setting forcibly disconnects client sessions with the SMB Service when the client's logon hours expire. Disabling this policy setting maintains an established client session after the client's logon hours have expired.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
If this policy setting is disabled, a user can remain connected to the system outside of their allotted logon hours. It is advisable to set **Network security: Force logoff when logon hours expire** to **Enabled**. SMB sessions will be terminated on member servers when a user's logon time expires, and the user will be unable to log on to the system until their next scheduled access time begins.

This policy setting does not apply to Administrator accounts.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Disabled |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Network security: LAN Manager authentication level**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

LAN Manager (LM) is a family of early Microsoft client/server software that allows users to link personal computer systems together on a single network. Networking capabilities include transparent file and print sharing, user security features, and network administration tools. In Active Directory domains, Kerberos is the default for authentication, but if Kerberos is not negotiated for some reason, Active Directory will use LM, NTLM, or NTLM version 2 (NTLMv2).

LM authentication—including the LM, NTLM, and NTLMv2 variants—is the protocol used to authenticate all Windows clients for such operations as:

- Joining a domain.
- Authenticating between Active Directory forests.
- Authenticating to earlier-version domains.
- Authenticating to systems that are not based on Windows 2000, Windows Server 2003, or Windows XP.
- Authenticating to systems that are not in the domain.

The **Network security: LAN Manager authentication level** policy setting determines which challenge/response authentication protocol is used for network logons. This choice affects the authentication protocol level clients use, the session security level the systems negotiate, and the authentication level servers accept, as follows:

- **Send LM & NTLM responses.** Clients use LM and NTLM authentication and never use NTLMv2 session security; domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Send LM & NTLM - use NTLMv2 session security if negotiated.** Clients use LM and NTLM authentication and use NTLMv2 session security if the server supports it. Domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Send NTLM response only.** Clients use NTLM authentication only and use NTLMv2 session security if the server supports it. Domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Send NTLMv2 response only.** Clients use NTLMv2 authentication only and use NTLMv2 session security if the server supports it. Domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Send NTLMv2 response only\refuse LM.** Clients use NTLMv2 authentication only and use NTLMv2 session security if the server supports it. Domain controllers refuse LM (accept only NTLM and NTLMv2 authentication).
- **Send NTLMv2 response only\refuse LM & NTLM.** Clients use NTLMv2 authentication only and use NTLMv2 session security if the server supports it. Domain controllers refuse LM and NTLM (accept only NTLMv2 authentication).

The possible values for this Group Policy setting are:

- Send LM & NTLM responses.
- Send LM & NTLM - use NTLMv2 session security if negotiated.
- Send NTLM responses only.
- Send NTLMv2 responses only.
- Send NTLMv2 responses only\refuse LM.
- Send NTLMv2 responses only\refuse LM & NTLM.
- Not defined.

These settings correspond to the levels discussed in other Microsoft documents, as follows:

- **Level 0–Send LM and NTLM response; never use NTLMv2 session security.** Clients use LM and NTLM authentication, and never use NTLMv2 session security. Domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Level 1–Use NTLMv2 session security if negotiated.** Clients use LM and NTLM authentication, and use NTLMv2 session security if the server supports it. Domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Level 2–Send NTLM response only.** Clients use only NTLM authentication, and use NTLMv2 session security if the server supports it. Domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Level 3–Send NTLMv2 response only.** Clients use NTLMv2 authentication, and use NTLMv2 session security if the server supports it. Domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Level 4–Domain controllers refuse LM responses**. Clients use NTLM authentication, and use NTLMv2 session security if the server supports it. Domain controllers refuse LM authentication, that is, they accept NTLM and NTLMv2.
- **Level 5–Domain controllers refuse LM and NTLM responses (accept only NTLMv2).** Clients use NTLMv2 authentication, use and NTLMv2 session security if the server supports it. Domain controllers refuse NTLM and LM authentication (they accept only NTLMv2).

**Discussion**

Windows 2000, Windows Server 2003, and Windows XP are configured by default to send LM and NTLM authentication responses. Windows 95, Windows 98, and Windows Millennium Edition clients only send LM. The default server setting allows all clients to authenticate with servers and use their resources; however, this means that LM responses—the weakest form of authentication response—are sent over the network, and it is potentially possible for attackers to sniff that traffic to more easily reproduce a user's password.

The Windows 95, Windows 98, Windows Me, and Windows NT operating systems cannot use the Kerberos version 5 protocol for authentication. For this reason, in a Windows Server 2003 domain, these systems authenticate by default with both the LM and NTLM protocols for network authentication. You can enforce a more secure authentication protocol for Windows 95, Windows 98, Windows Me, and Windows NT by using NTLMv2. For the logon process, NTLMv2 uses a secure channel to protect the authentication process. Even if you do use NTLMv2 for legacy clients and servers, Windows-based clients and servers that are members of the domain will authenticate with Windows Server 2003 domain controllers by using the Kerberos protocol.

Windows NT 4.0 requires Service Pack 4 (SP4) to support NTLMv2, and Windows 95, Windows 98, and Windows Me need the directory service client installed in order to support NTLMv2.

It is advisable to set **LAN Manager Authentication Level** to **Send NTLMv2 responsesonly**. When all clients support NTLMv2, this level of authentication is strongly recommended by Microsoft and a number of independent organizations. Clients that do not support NTLMv2 authentication will not be able to authenticate in the domain or access domain resources

by using LM and NTLM.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Send NTLM response only |
| Stand-Alone Server Default Settings | Send NTLM response only |
| DC Effective Default Settings | Send NTLM response only |
| Member Server Effective Default Settings | Send NTLM response only |

**Network security: LDAP client signing requirements**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

The **Network security: LDAP client signing requirements** policy setting determines the level of data signing that is requested on behalf of clients that issue LDAP BIND requests, as follows:

- **None.** The LDAP BIND request is issued with the caller-specified options.
- **Negotiate signing.** If TLS/SSL has not been started, the LDAP BIND request is initiated with the LDAP data signing option set in addition to the caller-specified options. If TLS/SSL has been started, the LDAP BIND request is initiated with the caller-specified options.
- **Require signature.** This is the same as **Negotiate signing**. However, if the LDAP server's intermediate saslBindInProgress response does not indicate that LDAP traffic signing is required, the caller is told that the LDAP BIND command request failed.

This policy setting does not have any impact on ldap_simple_bind or ldap_simple_bind_s. No Microsoft LDAP clients that are included with Windows XP Professional use ldap_simple_bind or ldap_simple_bind_s to communicate with a domain controller.

The possible values for this Group Policy setting are:

- None.
- Negotiate signing.
- Require signature.
- Not defined.

**Discussion**
Unsigned network traffic is susceptible to man-in-the-middle attacks, where an intruder captures the packets between the client and server and modifies them before forwarding them to the server. In the case of an LDAP server, this means that a malicious user might cause a server to make decisions based on false queries from the LDAP client. You can lower this risk in a corporate network by implementing strong physical security measures to protect the network infrastructure. Furthermore, all types of man-in-the-middle attacks can be made extremely difficult by requiring digital signatures on all network packets via IPsec Authentication Headers.

It is advisable to set **Domain controller: LDAP server signing requirements** to **Require signature**. If you set the server to require LDAP signatures, you must also set the client to do so. Not setting the client will prevent the client from communicating

with the server—this can cause many features to fail, including user authentication, Group Policy, and logon scripts.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Negotiate signing |
| DC Effective Default Settings | Negotiate signing |
| Member Server Effective Default Settings | Negotiate signing |

**Network security: Minimum session security for NTLM SSP based (including secure RPC) clients**
The **Network security: LDAP client signing requirements** security policy setting determines the level of data signing that is requested on behalf of clients issuing LDAP BIND requests, as follows:

- **None:** The LDAP BIND request is issued with the caller-specified options.
- **Negotiate signing:** If TLS/SSL has not been started, the LDAP BIND request is initiated with the LDAP data-signing option set in addition to the caller-specified options. If TLS/SSL has been started, the LDAP BIND request is initiated with the caller-specified options.
- **Require signature:** This is the same as **Negotiate signing**. However, if the LDAP server's intermediate saslBindInProgress response does not indicate that LDAP traffic signing is required, the caller is told that the LDAP BIND command request failed.

This policy setting does not have any impact on ldap_simple_bind or ldap_simple_bind_s. No Microsoft LDAP clients that are included with Windows XP Professional use ldap_simple_bind or ldap_simple_bind_s to communicate with a domain controller.

The possible values for this Group Policy setting are:

- None.
- Negotiate signing.
- Require signature.
- Not defined.

**Discussion**
Unsigned network traffic is susceptible to man-in-the-middle attacks, where an intruder captures the packets between the client and server and modifies them before forwarding them to the server. In the case of an LDAP server, this means that a malicious user could cause a server to make decisions based on false queries from the LDAP client. You can lower this risk in a corporate network by implementing strong physical security measures to protect the network infrastructure. Furthermore, all types of man-in-the-middle attacks can be made extremely difficult by requiring digital signatures on all network packets by way of IPsec Authentication Headers.

It is advisable to set **Domain controller: LDAP server signing requirements** to **Require signature**. If you set the server to require LDAP signatures, you must also set the client. Not setting the client will prevent the client from communicating with the server — this can cause many features to fail, including user authentication, Group Policy, and logon scripts.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | No minimum |
| DC Effective Default Settings | No minimum |
| Member Server Effective Default Settings | No minimum |

**Network security: Minimum session security for NTLM SSP based (including secure RPC) servers**

The **Network security: Minimum session security for NTLM SSP based (including secure RPC) servers** policy setting allows a server to require the negotiation of message confidentiality (encryption), message integrity, 128-bit encryption, or NTLMv2 session security. These values are dependent on the value of the **Network security: LAN Manager authentication level** policy setting.

The possible values for this Group Policy setting are:

- Require message confidentiality.

  The connection will fail if encryption is not negotiated. Encryption converts data into a form that is not readable by anyone until it is decrypted.

- Require message integrity.

  The connection will fail if message integrity is not negotiated. The integrity of a message can be assessed through message signing. Message signing proves that the message has not been tampered with by attaching a cryptographic signature, which identifies the sender and contains a numeric representation of the contents of the message.

- Require 128-bit encryption.

  The connection will fail if strong encryption (128-bit) is not negotiated.

- Require NTLMv2 session security.

  The connection will fail if the NTLMv2 protocol is not negotiated.

- Not defined.

**Discussion**

Setting all of these values for this policy setting will help protect network traffic that uses the NTLM Security Support Provider (NTLM SSP) from being exposed or tampered with by a malicious user who has gained access to the same network. That is, these settings help protect against man-in-the-middle attacks.

Enable all four values that are available for the **Network security: Minimum session security for NTLM SSP based (including secure RPC) servers** object. Legacy clients that do not support these policy settings will be unable to communicate with the computer.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |

| | |
|---|---|
| Stand-Alone Server Default Settings | No minimum |
| DC Effective Default Settings | No minimum |
| Member Server Effective Default Settings | No minimum |

**Recovery console: Allow automatic administrative logon**

The **Recovery console: Allow automatic administrative logon** policy setting determines whether a user must give the Administrator account password before getting access to the system. Enabling this policy setting allows any user to automatically log on to the system without giving a password at the Recovery Console.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

The Recovery Console can be very useful when troubleshooting and repairing systems that cannot be restarted; however, enabling this policy setting so a user can automatically log on to the console is dangerous. Anyone can walk up to the server, shut it down by disconnecting the power, reboot it, select **Recovery Console** from the **Restart** menu, and then assume full control of the server.

It is advisable to set **Recovery Console: Allow automatic administrative logon** to **Disabled**. This requires a user to enter a user name and password to access the Recovery Console account.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Recovery console: Allow floppy copy and access to all drives and all folders**

Enabling the **Recovery console: Allow floppy copy and access to all drives and all folders** policy setting makes the Recovery Console SET command available, which allows you to set the following Recovery Console environment variables.

- **AllowWildCards.** Enables wildcard support for some commands (such as the DEL command).
- **AllowAllPaths.** Allows access to all files and folders on the computer.
- **AllowRemovableMedia.** Allows files to be copied to removable media, such as a floppy disk.
- **NoCopyPrompt.** Do not prompt when overwriting an existing file.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.

- Not defined.

**Discussion**
An authorized administrator might forget to remove a CD or floppy disk with sensitive data or applications that a malicious user could then steal. An authorized administrator could accidentally leave a startup disk in the computer after using the Recovery Console. If the computer is restarted for any reason and the BIOS has been configured to boot from the CD-ROM or floppy disk drive before the hard disk, the server will start from the removable disk. This causes the server's network services to be unavailable.

It is advisable to set **Recovery Console: Allow floppy copy and access to drives and folders** to **Disabled**. Users who have started a server by using the Recovery Console and logged in with the built-in Administrator account will not be able to copy files and folders to a floppy disk.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Shutdown: Allow system to be shut down without having to log on**
The **Shutdown: Allow system to be shut down without having to log on** policy setting determines whether a computer can be shut down by a user who has not logged on to Windows. Enabling this policy setting makes the **Shut Down** command available on the Windows logon screen. Disabling this policy setting removes the **Shut Down** command from the Windows logon screen. When the setting is disabled, users must be able to log on to the computer successfully and have the **Shut down the system** user right before they can perform a system shutdown.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
Users who can access the console locally can shut the system down. Attackers or misguided users can connect to the server using Terminal Services and shut it down or restart it without having to identify themselves. A malicious user might also cause a temporary denial-of-service condition by walking up to the local console and restarting the server, or shutting down the server and thus rendering unavailable all its applications and services.

It is advisable to set the **Allow system to be shut down without having to log on** user right to **Disabled**. Administrators must log on to servers to shut them down or restart them.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |

| | |
|---|---|
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**Shutdown: Clear virtual memory page file**

The **Shutdown: Clear virtual memory page file** policy setting determines whether the virtual memory paging file is cleared when the system is shut down. Virtual memory support uses a system paging file to swap pages of memory to disk when they are not used. On a running system, this paging file is opened exclusively by the operating system, and it is well-protected. However, systems that have been configured to allow starting other operating systems might have to make sure that the system paging file is wiped clean when the running system shuts down. This will ensure that sensitive information from process memory that might have been written to the paging file is not available to an unauthorized user who might manage to access the paging file directly.

When this policy setting is enabled, it clears the system paging file when the system shuts down properly. On a portable computer system on which hibernation has been disabled, the system will also clear the hibernation file, Hiberfil.sys.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

Important information kept in real memory might be written periodically to the paging file. This helps Windows Server 2003 handle multitasking functions. A malicious user who has physical access to a server that has been shut down can view the contents of the paging file. The attacker can move the system volume into a different computer and then analyze the contents of the paging file. This is a time-consuming process, but it can expose data cached from RAM to the paging file. A malicious user who has physical access to the server can bypass this countermeasure by simply unplugging the server from its power source.

It is advisable to set **Shutdown: Clear virtual memory page file** to **Enabled**. This causes Windows Server 2003 to clear the paging file when the system is shut down. Depending on the size of the paging file, this process might take several minutes before the system completely shuts down. This delay in shutting down the server is especially noticeable on servers with large paging files. For a server with 2 gigabytes (GB) of RAM and a 2-GB paging file, this setting can add 20 to 30 minutes, or more, to the shutdown process. For some organizations, this downtime violates their internal service level agreements. Use caution when implementing this countermeasure in your environment.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**System cryptography: Force strong key protection for user keys stored on the computer**

The **System cryptography: Force strong key protection for user keys stored on the computer** policy setting determines

whether users can use private keys, such as their Secure/Multipurpose Internet Mail Extensions (S/MIME) key, without a password.

The possible values for this Group Policy setting are:

- User input is not required when new keys are stored and used.
- User is prompted when the key is first used.
- User must enter a password each time they use a key.
- Not defined.

**Discussion**
Configuring this policy setting so that users must provide a password every time they use a key — in addition to their domain password — makes it more difficult for a malicious user to access locally-stored user keys, even if the attacker takes control of the user's computer and determines their logon password.

It is advisable to set **System cryptography: Force strong key protection for user keys stored on the computer** to **User must enter a password each time they use a key**. Users must enter their password every time they access a key that is stored on their computer. For example, if users use an S/MIME certificate to digitally sign their e-mail, they will be forced to enter the password for that certificate every time they send a signed e-mail message. For some organizations, the overhead caused by using this value might be too high, but even they should set the value at a minimum to **User is prompted when the key is first used**.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing**
The **System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing** policy setting determines whether the TLS/SSL security provider will only support the strong cipher suite known as TLS_RSA_WITH_3DES_EDE_CBC_SHA. In effect, this means that the provider only supports the TLS protocol as a client and as a server, if applicable. It uses only the Triple Data Encryption Standard (3DES) encryption algorithm for the TLS traffic encryption, only the Rivest - Shamir-Adleman (RSA) public key algorithm for the TLS key exchange and authentication, and only the Secure Hash Algorithm version 1 (SHA-1) hashing algorithm for the TLS hashing requirements.

For Encrypting File System (EFS), it supports only the 3DES encryption algorithm for encrypting file data supported by NTFS. By default, EFS uses the DESX algorithm (a variation of the DES algorithm) for encrypting file data.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

Enabling this policy setting ensures that the computer will use the most powerful algorithms available for digital encryption, hashing, and signing. This minimizes the risk of a malicious user compromising digitally encrypted or signed data.

It is advisable to set **System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing** to **Enabled**. Clients with this policy setting enabled will be unable to communicate via digitally encrypted or signed protocols with servers that do not support these algorithms. Network clients that do not support these algorithms will not be able to use servers that require them for network communications. If you enable this policy setting, you must also configure Internet Explorer to use TLS.

**To enable Internet Explorer to use TLS**

1. In Internet Explorer, on the **Tools** menu, click **Internet Options**.
2. Click the **Advanced** tab.
3. Select the **Use TLS 1.0** check box.

It is also possible to configure this setting by using Group Policy or the Internet Explorer Administrators Kit.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

**System objects: Default owner for objects created by members of the Administrators group**
The **System objects: Default owner for objects created by members of the Administrators group** policy setting determines whether the Administrators group or the object creator is the default owner of any system objects that are created.

The possible values for this Group Policy setting are:

- Administrators group.
- Object creator.
- Not defined.

**Discussion**
Setting the value to **Administrators group** makes it impossible to hold individuals accountable for creating new system objects.

It is advisable to set **System objects: Default owner for objects created by members of the Administrators group** to **Object creator**. When a system object is created, its ownership reflects which specific account created the object, rather than the more generic Administrators group.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|

| | |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Administrators group |
| DC Effective Default Settings | Administrators group |
| Member Server Effective Default Settings | Administrators group |

**System objects: Require case insensitivity for non-Windows subsystems**

The **System objects: Require case insensitivity for non-Windows subsystems** policy setting determines whether case-insensitivity is enforced for all subsystems. The Microsoft Win32 subsystem is not case-sensitive; however, the kernel supports case-sensitivity for other subsystems, such as Portable Operating System Interface for UNIX (POSIX). Enabling this policy setting enforces case-insensitivity for all directory objects, symbolic links, and input/output (I/O) objects, including file objects. Disabling this policy setting does not allow the Win32 subsystem to become case-sensitive.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

Because Windows is case-insensitive but the POSIX subsystem will support case-sensitivity, if this policy setting is not enforced it will be possible for a user of that subsystem to create a file with the same name as another file but with a different mix of capital letters. That might confuse users when they try to access these files by using normal Win32 tools, because only one of the files will be available.

It is advisable to set **System objects: Require case insensitivity for non-Windows subsystems** to **Enabled**. All subsystems will be forced to observe case-insensitivity. This might confuse users who are familiar with one of the UNIX-based operating systems and are used to a case-sensitive operating system.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**System objects: Strengthen default permissions of internal system objects (e.g. Symbolic Links)**

The **System objects: Strengthen default permissions of internal system objects (e.g. Symbolic Links)** policy setting determines the strength of the default discretionary access control list (DACL) for objects. Windows maintains a global list of shared system resources, such as MS-DOS device names, mutexes, and semaphores. By using this list, processes can locate and share objects.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.

- Not defined.

**Discussion**
This policy setting determines the strength of the default DACL for objects. Windows Server 2003 maintains a global list of shared system resources such as MS-DOS device names, mutexes, and semaphores. By using this list, processes can locate and share objects. Each type of object is created with a default DACL that specifies who can access the objects with what permissions. Enabling this policy setting strengthens the default DACL and allows users who are not administrators to read, but not to modify, shared objects that they did not create.

It is advisable to set **Strengthen default permissions of global system objects (for example, Symbolic Links)** to **Enabled**.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Enabled |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**System settings: Optional subsystems**
The **System settings: Optional subsystems** policy setting determines which subsystems support your applications. With this policy setting, you can specify as many support subsystems as your environment demands.

The possible values for this Group Policy setting are:

- A user-defined list of subsystems.

- Not defined.

**Discussion**
The POSIX subsystem is an Institute of Electrical and Electronic Engineers (IEEE) standard that defines a set of operating system services. The POSIX subsystem is required if the server supports applications that use that subsystem.

The subsystem introduces a security risk relating to processes that can potentially persist across logons. If a user starts a process and then logs out, the next user who logs on to the system might access the process the previous user started. This is dangerous, because the process started by the first user can retain that user's system user rights; therefore, anything the second user does using that process is performed with the user rights of the first user. This makes it difficult to trace who creates processes and objects, which is essential for post-security incident forensics.

It is advisable to set the **System settings: Optional subsystems** policy setting to a null value. The default value is **POSIX**, so applications that rely on the POSIX subsystem will no longer run. For example, Microsoft Services for UNIX 3.0 installs an updated version of the POSIX subsystem. Reset this policy setting in Group Policy for any servers that use Services for UNIX 3.0.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |

| | |
|---|---|
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Posix |
| DC Effective Default Settings | Posix |
| Member Server Effective Default Settings | Posix |

**System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**

The **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies** policy setting determines whether digital certificates are processed when a user or process attempts to run software with an .exe file name extension. This policy setting enables or disables certificate rules (a type of Software Restriction Policies rule). By using policy settings under Software Restriction Policies, you can create a certificate rule that allows or disallows the running of Authenticode-signed software, based on the digital certificate that is associated with the software. In order for certificate rules to take effect, you must enable this policy setting.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**

Software restriction policies help to prevent users and computers from executing unauthorized code, such as viruses or Trojan horses.

It is advisable to set **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies** to **Enabled**. Enabling certificate rules results in software restriction policies checking a certificate revocation list (CRL) to make sure the software's certificate and signature are valid. When you start signed programs, this setting can decrease system performance. You can disable CRLs by editing Software Restriction Policies in the desired GPO. In the **Trusted Publishers Properties** dialog box, clear the **Publisher** and **Timestamp** check boxes.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Disabled |
| DC Effective Default Settings | Disabled |
| Member Server Effective Default Settings | Disabled |

# Event Log Policy Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Event Log**

The Event Log service records events on the system by writing to one of three default logs that you can read in Event Viewer: the security, application, and system logs. The security log records audit events. You use the settings under Event Log to specify attributes of the security, application, and system logs, such as maximum log size, access rights for each log, and retention settings and methods.

Event Log policy settings can be configured in the following location in Group Policy Object Editor:

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

**Maximum event log size (settings for application, security and system logs)**

**Note**

- Misuse of these policy settings is a common error that can cause data loss or problems with data access or security.

These policy settings specify the maximum size of the application, security, and system logs. Although you can specify values as large as 4 GB in Group Policy Object Editor and Event Viewer, there are factors that make the effective maximum size for these logs much smaller.

The Event Log service uses memory-mapped files, and it runs as Eventlog.dll, one of the services under the Services.exe process. When files are loaded in this way, the entire file is loaded into system memory. All of the current versions of Windows have an architectural limitation regarding memory-mapped files: no process can have more than 1 GB of memory-mapped files in total. This means that all of the services running under the Services.exe process must share the 1-GB pool. The memory is assigned as contiguous 64-KB chunks of memory. If the system is unable to assign additional memory needed to expand memory-mapped files, problems will arise.

For the Event Log service, this means that regardless of how large the log has been configured to be, events might no longer be written to the log. Error messages will not be displayed. The events will simply not appear in the event log, or they might overwrite other events that have been recorded previously. Fragmentation of the log files in memory has also been shown to lead to significant performance problems on busy systems.

Due to these limitations — even though the theoretical limit for memory-mapped files suggests that you should be able to configure up to 1 GB for all the event logs, and you can actually specify as much as 4 GB per log — Microsoft has verified that the practical size limit for all event logs combined is around 300 megabytes (MB) on most servers. On Windows XP, member servers, and stand-alone servers, the combined size of the application, security, and system event logs should not exceed 300 MB. On domain controllers, the combined size of these three logs — plus the Directory Service, File Replication Service, and DNS Server logs — should not exceed 300 MB.

These limitations have caused problems for some Microsoft customers, but addressing them will require fundamental changes to the architecture for recording system events. Microsoft plans to resolve these problems in the next version of Windows by rewriting the event logging system from the ground up.

Although you will need to use trial and error to determine the best log size for a particular server, you can use the fact that the average event writes about 500 bytes to each log. Estimate the average number of events generated each day for each type of log in your enterprise. Remember that log file sizes must be a multiple of 64 KB.

For example, if your file server generates 5,000 events per day in its security log and you want to ensure that you have at least four weeks of data available at all times, you want to set the size of that log to about 70 MB (500 bytes 5,000 events/day 28 days = 70,000,000 bytes.) Then, check the servers occasionally over the following four weeks to verify that the logs are retaining enough events. Event log size and log wrapping should be configured to match the business and security requirements you determined when you designed your enterprise security plan.

The possible values for this Group Policy setting are:

- A user-defined number of kilobytes from 64 through 4,194,240; however, it must be a multiple of 64.

If you significantly increase the number of objects to audit in your organization, you run the risk of filling the security log to

capacity and thus forcing the system to shut down. If this occurs, the system will be unusable until an administrator clears the security log. To prevent this, disable the **Audit: Shut down system immediately if unable to log security audits** policy setting and increase the security log size.

All computers in your organization should have sensible log size policy settings enabled so that legitimate users can be held accountable for their actions, unauthorized activity can be detected and tracked, and system problems can be detected and diagnosed.

When event logs fill to capacity, entries can no longer be written to them unless the retention method for each is set so that the system will overwrite the oldest entries with the most recent ones. The risk of the event logs' not containing recent entries can be mitigated by setting the retention method so that older events are overwritten as needed. The consequence of this action is that older events will be removed from the logs. Attackers can use this to their advantage by generating a large number of extraneous events to overwrite any evidence of their attack.

Ideally, all specifically monitored events will be sent to a server by using Microsoft Operations Manager (MOM) or some other automated monitoring tool. This is particularly important because an attacker who successfully compromises a server could clear the security log. If all events are sent to a monitoring server, you will be able to gather post-incident forensic information about the attacker's activities.

**Maximum application log size**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 16384 KB |
| DC Effective Default Settings | 16384 KB |
| Member Server Effective Default Settings | 16384 KB |

**Maximum security log size**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 16384 KB |
| DC Effective Default Settings | 16384 KB |
| Member Server Effective Default Settings | 16384 KB |

**Maximum system log size**

**Note**

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | 16384 KB |
| DC Effective Default Settings | 16384 KB |
| Member Server Effective Default Settings | 16384 KB |

**Prevent local guests group from accessing application log**
The **Prevent local guests group from accessing application log** policy setting determines whether guests are prevented from accessing the application log.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
An attacker who has successfully logged on to a computer with guest user rights can learn important information about the system by viewing the application log. The attacker might then use this information to launch additional attacks.

It is advisable to enable this policy setting.

This policy setting does not appear in the Local Computer Policy object. It only affects computers running Windows 2000, Windows XP, and Windows Server 2003.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Prevent local guests group from accessing security log**
The **Prevent local guests group from accessing security log** policy setting determines whether guests are prevented from accessing the security log.

The possible values for this Group Policy setting are:

- Enabled.

- Disabled.
- Not defined.

**Discussion**
An attacker who has successfully logged on to a computer with guest user rights can learn important information about the system by viewing the security log. The attacker might then use this information to launch additional attacks.

It is advisable to enable this policy setting.

This policy setting does not appear in the Local Computer Policy object. It only affects computers running Windows 2000,Windows XP, and Windows Server 2003.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Prevent local guests group from accessing system log**
The **Prevent local guests group from accessing system log** policy setting determines whether guests are prevented from accessing the system log.

The possible values for this Group Policy setting are:

- Enabled.
- Disabled.
- Not defined.

**Discussion**
An attacker who has successfully logged on to a computer with guest user rights can learn important information about the system by viewing the system log. The attacker might then use this information to launch additional attacks.

It is advisable to enable this policy setting.

This policy setting does not appear in the Local Computer Policy object. It only affects computers running Windows 2000,Windows XP, and Windows Server 2003.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |

| | |
|---|---|
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Enabled |
| Member Server Effective Default Settings | Enabled |

**Retain application log**

The **Retain application log** policy setting determines how many days' worth of events to retain, if the retention method specified for the log is **By Days**.

The possible values for this Group Policy setting are:

- A user-defined number of days from 1 through 365.
- Not defined.

**Discussion**

If you archive the log at scheduled intervals, in the property sheet for this policy setting, specify the appropriate number of days and select **Overwrite events by days** for the event log retention method. Also, ensure that the maximum log size is large enough to accommodate the amount of information you want to gather during the archive interval.

It is advisable to set the value of the **Retain event logs** policy setting for all three event logs to **Not defined**.

This policy setting does not appear in the Local Computer Policy object. A user must have been assigned the **Manage auditing and security log** user right to access the security log.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Retain security log**

The **Retain security log** policy setting determines how many days' worth of events to be retained, if the retention method specified for the log is **By Days**.

The possible values for this Group Policy setting are:

- A user-defined number of days from 1 through 365.
- Not defined.

**Discussion**

If you archive the log at scheduled intervals, in the property sheet for this policy setting, specify the appropriate number of days and select **Overwrite events by days** for the event log retention method. Also, ensure that the maximum log size is large enough to accommodate the amount of information you want to gather during the archive interval.

It is advisable to set the value of the **Retain event logs** policy setting for all three event logs to **Not defined**.

This policy setting does not appear in the Local Computer Policy object. A user must have been assigned the **Manage auditing and security log** user right to access the security log.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Retain system log**

The **Retain system log** policy setting determines how many days' worth of events to be retained, if the retention method specified for the log is **By Days**.

The possible values for this Group Policy setting are:

- A user-defined number of days from 1 through 365.
- Not defined.

**Discussion**

If you archive the log at scheduled intervals, in the property sheet for this policy setting, specify the appropriate number of days and select **Overwrite events by days** for the event log retention method. Also, ensure that the maximum log size is large enough to accommodate the amount of information you want to gather during the archive interval.

It is advisable to set the value of the **Retain event logs** policy setting for all three event logs to **Not defined**.

This policy setting does not appear in the Local Computer Policy object. A user must have been assigned the **Manage auditing and security log** user right to access the security log.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Retention method for event log (separate policy settings for application, security and system logs)**

**Note**

- Misuse of these policy settings is a common error that can cause data loss or problems with data access or security.

The **Retention method for event log** policy settings determine the wrapping method for the application, security, and system logs.

If you do not want to archive the logs, in the property sheet for this policy setting, select the **Define this policy setting** check box, and then click **Overwrite events as needed**.

If you want to archive the log at scheduled intervals, in the property sheet for this policy setting, select the **Define this policy setting** check box, click **Overwrite events by days**, and then specify the appropriate number of days in the **Retain**

**application log** policy setting. Ensure that the maximum log size is large enough to accommodate the amount of information you expect to gather during the archive interval.

If you must retain all the events in the log, in the property sheet for this policy setting, select the **Define this policy setting** check box, and then click **Do not overwrite events (clear log manually)**. This value requires that the log be cleared manually. When the maximum log size is reached, new events are not written to the log, they are discarded.

The possible values for these Group Policy settings are:

- Overwrite events by days.
- Overwrite events as needed.
- Do not overwrite events (clear log manually).
- Not defined.

These policy settings do not appear in the Local Computer Policy object.

If you significantly increase the number of objects you audit in your organization, you run the risk of filling the security log to capacity and thus forcing the system to shut down. If this occurs, the system will be unusable until an administrator clears the security log. To prevent this, disable the **Audit: Shut down system immediately if unable to log security audits** policy setting and then increase the security log size.

If you set the value of **Event log retention method** to **Do not overwrite events (clear log manually)** or **Overwrite events by days**, important recent events might not be recorded, or you might suffer a denial-of-service attack.

It is advisable to set **Event log retention method** for all three event logs to the value **Overwrite events as needed**. Some resources recommend configuring this policy setting to **Do not overwrite events (clear log manually)**; however, the administrative burden that this setting imposes is too high for most organizations. Ideally, all significant events will be sent to a monitoring server by using MOM or some other automated monitoring tool.

#### Retention method for application log

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

#### Location

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

### Default Values

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Overwrite as needed |
| DC Effective Default Settings | Overwrite as needed |
| Member Server Effective Default Settings | Overwrite as needed |

#### Retention method for security log

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

#### Location

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

### Default Values

| Server Type or GPO | Default Value |
| --- | --- |
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Overwrite as needed |
| DC Effective Default Settings | Overwrite as needed |
| Member Server Effective Default Settings | Overwrite as needed |

**Retention method for system log**

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**

*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Event Log\

### Default Values

| Server Type or GPO | Default Value |
| --- | --- |
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Overwrite as needed |
| DC Effective Default Settings | Overwrite as needed |
| Member Server Effective Default Settings | Overwrite as needed |

# Restricted Groups Policy Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restricted Groups**
This security setting allows an administrator to define two properties for security-sensitive groups (or *restricted* groups). The two properties are **Members** and **Member Of**.

The **Members** list defines exactly who belongs and who does not belong to the restricted group. Both inclusion on the list and exclusion from the list, are enforced.

The **Member Of** list specifies groups in which the restricted group is included. Only inclusion in the **Member Of** list is enforced, not exclusion: If you remove a group from the **Member Of** list, the restricted group is allowed to remain a member of the removed group.

### Note

- Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Restricted Groups\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Discussion**
The Restricted Groups folder is available only in Group Policy objects associated with domains, organizational units, and sites. The Restricted Groups folder does not appear in the local Group Policy object.

If a restricted group is defined so that it has no members (that is, the **Members** list is empty), all members of the group are removed when the policy is enforced on the computer. If the **Member Of** list is empty, no changes are made to any groups to which the restricted group belongs.

# System Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Services**
This security setting allows you to define the startup mode (manual, automatic, or disabled) and access permissions for all services (**Start**, **Stop**, or **Pause**). You can also edit other security properties of the service, such as which user or group accounts have permission to read, write, delete, or execute inheritance settings or auditing and ownership permission by clicking **Edit Security**.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\System Services\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined for any service |
| Default Domain Controller Policy | Not defined for any service |
| Stand-Alone Server Default Settings | Not defined for any service |
| DC Effective Default Settings | Not defined for any service |
| Member Server Effective Default Settings | Not defined for any service |

**Discussion**
System services are processes that normally run in the background. An administrator uses the Services snap-in to manage system services directly. To manage system services through security policy, administrators can use the System Services security setting.

This security setting does not appear in the local Group Policy object.

If you choose to set system service startup to **Automatic**, perform adequate testing to verify that the services can start without user intervention.

For performance optimization, set unnecessary or unused services to **Manual**.

To find descriptions of each individual service (including what happens when you stop it from running), see the “System Services chapter of Threats and Countermeasures: Security Settings” in Windows Server 2003 and Windows XP, which can be downloaded from the Microsoft Web site.

# Registry

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Registry**
This security setting allows you to define access permissions and audit settings, and to take or assign ownership of individual registry keys.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Registry\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined for any key |
| Default Domain Controller Policy | Not defined for any key |
| Stand-Alone Server Default Settings | Not defined for any key |
| DC Effective Default Settings | Not defined for any key |
| Member Server Effective Default Settings | Not defined for any key |

**Discussion**
The default permissions on registry keys allow full control by administrators and the system. Users have read-only permission, and the creator/owner can take or assign ownership of the key. This security setting can be used to change these default permissions, or to set specific permissions for other users and groups.

This security setting does not appear in the local Group Policy object.

# File System

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File System**
This security setting allows you to define access permissions and audit settings, and to take or assign ownership of files and folders.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\File System\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Discussion**
The default permissions on files and folders allow full control by administrators and the system. Users have read-only permission, and the creator/owner can take or assign ownership of them. This security setting can be used to change these default permissions, or to set specific permissions for other users and groups.

This security setting does not appear in the local Group Policy object.

# Wireless Network (IEEE 802.11) Policies

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless Network (IEEE 802.11) Policies**
This security setting allows you to create or change a wireless network policy.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Wireless Network (IEEE 802.11) Policies\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Discussion**
When you create a wireless network policy, you can restrict computers from accessing some kinds of wireless networks.

You can choose to permit:

1. **Wireless access to any available network.** With this choice, the policy client computer will preferentially connect to infrastructure in the form of fixed Institute of Electrical and Electronic Engineers (IEEE) 802.11 transmitter/receivers. The policy client will make ad-hoc connections to wireless-enabled computers if there is no access point available.

2. **Access point (infrastructure) networks only.**

3. **Computer-to-computer (ad hoc) networks only.**

A wireless network policy can also designate preferred networks, and key types for encryption and authentication. You can employ Smart Cards or certificates, if you enable access control through IEEE 802.1x.

You can create at most one wireless network policy per Group Policy object (GPO). It is recommended that you use a domain-linked GPO.

This security setting does not appear in the local GPO.

# Public Key Policy Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Public Key Policies**
Use this security setting to specify how to manage certificates and certificate life cycles.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Public Key Policies\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Discussion**
This security setting specifies whether certificates are automatically enrolled, renewed when they are expired, and removed when they are revoked.

You can specify certificates for data recovery, and you can specify Trusted Root Certification Authorities, and Enterprise Trust lists.

This security setting does not appear in the local Group Policy object.

# Software Restriction Policy Settings

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Software Restriction Policies**
This security setting allows you create or change a software restriction policy.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\Software Restriction Policies\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Discussion**
When you create a software restriction policy, you can prohibit users from running administratively specified software on their computers. More restrictively, you can prevent the use of all except administratively specified software on their computers. Optionally, local administrators can be exempted from software restriction policy.

You can exempt software from restriction if the software publisher is trusted by the group that you specify: End users, local computer administrators, or Enterprise administrators.

You can create rules to override the default security level. These rules can be based on certificates, hashes, Internet zones, and paths.

You can create at most one wireless network policy per Group Policy object (GPO).

The wireless network policies security setting does not appear in the local GPO.

# IP Security Policies on Active Directory

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP Security Policies on Active Directory**
This security setting allows you to permit, block, or negotiate security for TCP/IP traffic.

**Location**
*GPO_name*\Computer Configuration\Windows Settings\Security Settings\IP Security Policies on Active Directory\

**Default Values**

| Server Type or GPO | Default Value |
|---|---|
| Default Domain Policy | Not defined |
| Default Domain Controller Policy | Not defined |
| Stand-Alone Server Default Settings | Not defined |
| DC Effective Default Settings | Not defined |
| Member Server Effective Default Settings | Not defined |

**Discussion**
Use this setting to create filter lists and corresponding actions. Filter lists are based on ports, protocols, and whether the traffic is incoming or outgoing. Actions you specify are carried out when traffic matches a filter in the list. Actions include blocking, permitting, and requiring or requesting that traffic be secured through encryption.

Specify authentication methods. Choices include Kerberos V5, certificates, and (for testing only) preshared keys.

IP Security Policy can also be set in the local Group Policy object. IP Security Policy that is set in Active Directory takes precedence.

# Windows Firewall Settings

Updated: July 10, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Firewall Settings**

**In this topic:**


- Windows Firewall Settings: Optional Components
- Windows Firewall Settings: Remote Administration Tools
- Windows Firewall Settings: Server Roles
- Windows Firewall Settings: Services


This technical reference provides a summary of the Windows Firewall settings you need to configure for each Microsoft® Windows Server™ 2003 optional component, remote administration tool, server role, and service.

Windows Firewall supports configuring its settings in two profiles: domain and standard. Only one profile is active at a time. If all of the network connections on a computer allow the computer to reach a domain controller for the domain to which it is joined, then the domain profile is used. If ANY network connection is to a network through which a domain controller cannot be reached, then the standard profile is used. For more information about profiles, see Help: Understanding Windows Firewall Profiles (http://go.microsoft.com/fwlink/?linkid=157699).

For more information about Windows Firewall, see Windows Firewall Operations Guide (http://go.microsoft.com/fwlink/?LinkId=43262) and Help: Windows Firewall (http://go.microsoft.com/fwlink/?LinkId=43263).

# Windows Firewall Settings: Optional Components

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Firewall is a stateful host firewall that blocks all unsolicited incoming TCP/IP traffic, including Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic. If you enable Windows Firewall on a server and you install and use an optional component that listens for unsolicited incoming traffic, you must usually configure Windows Firewall settings so that the optional component can receive unsolicited traffic. Although you can configure Windows Firewall settings manually, the recommended method is to use the Security Configuration Wizard (SCW).

To use an optional component with Windows Firewall, you usually need to add a program or port to the Windows Firewall exceptions list. When you add a program or port to the exceptions list, you instruct Windows Firewall to allow unsolicited incoming traffic to reach the specified program or pass through the specified port. In some cases, you might need to change a registry setting or enable one of the preconfigured Windows Firewall exceptions, such as the File and Printer Sharing exception or the Remote Administration exception.

| **Caution** |
| --- |
| The Remote Administration exception allows traffic through numerous ports, which can make your computer more accessible to attack. Be sure to read the Windows Firewall documentation so that you understand the risks of using the Remote Administration exception. Incorrectly editing the registry can severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

For more information about configuring Windows Firewall exceptions, see Managing Program, Port, and Service Exceptions (http://go.microsoft.com/fwlink/?LinkId=43261) and Help: Understanding Windows Firewall Exceptions (http://go.microsoft.com/fwlink/?LinkId=43260). For more information about SCW, see Security Configuration Wizard Overview on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43259).

Use the A-Z list to find out how to configure Windows Firewall for use with the following optional components.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**

Access Control system (Windows Firewall: Access Control system)

Accessibility Wizard (Windows Firewall: Accessibility Wizard)

ASP.NET (Windows Firewall: ASP.NET session state server)

**B**

No entries

**C**

Certificate Services (Windows Firewall: Certificate Services)

Certificate Services Web Enrollment Support (Windows Firewall: Certificate Services Web Enrollment Support)

Chat (Windows Firewall: Chat)

Connection Manager Administration Kit (Windows Firewall: Connection Manager Administration Kit)

Connection Point Services (Windows Firewall: Connection Point Services)

**D**

No entries

**E**

No entries

**F**

No entries

**G**
No entries

**H**
HyperTerminal (Windows Firewall: HyperTerminal)

**I**
IIS: Active Server Pages (Windows Firewall: IIS: Active Server Pages)

IIS: BITS server extension ISAPI (Windows Firewall: IIS: BITS server extension ISAPI)

IIS: FrontPage 2002 Server Extensions (Windows Firewall: IIS: FrontPage 2002 Server Extensions)

IIS: Internet Data Connector (Windows Firewall: IIS: Internet Data Connector)

IIS: Internet Printing (Windows Firewall: IIS: Internet Printing)

IIS: Remote Administration (HTML) (Windows Firewall: IIS: Remote Administration (HTML))

IIS: Remote Desktop Web Connection (Windows Firewall: IIS: Remote Desktop Web Connection)

IIS: Server Side Includes (Windows Firewall: IIS: Server Side Includes)

IIS: WebDAV Publishing (Windows Firewall: IIS: WebDAV Publishing)

Internet Gateway Device Discovery and Control Client (Windows Firewall: Internet Gateway Device Discovery and Control Client)

**J**
No entries

**K**
No entries

**L**
No entries

**M**
Message Queuing: Active Directory Integration (Windows Firewall: Message Queuing: Active Directory Integration)

Message Queuing: Core Functionality (Windows Firewall: Message Queuing: Core Functionality)

Message Queuing: Downlevel Client Support (Windows Firewall: Message Queuing: Downlevel Client Support)

Message Queuing: Local Storage (Windows Firewall: Message Queuing: Local Storage)

Message Queuing: MSMQ HTTP Support (Windows Firewall: Message Queuing: MSMQ HTTP Support)

Message Queuing: Routing Support (Windows Firewall: Message Queuing: Routing Support)

Message Queuing: Triggers (Windows Firewall: Message Queuing: Triggers)

**N**
Network COM+ access (Windows Firewall: Network COM access)

Network DTC access (Windows Firewall: Network DTC access)

Network Monitor Tools (Windows Firewall: Network Monitor Tools)

**O**
Outlook Express (Windows Firewall: Outlook Express)

**P**
POP3 Service Web Administration (Windows Firewall: POP3 Service Web Administration)

**Q**
No entries

**R**
Remote Access Quarantine Service (Windows Firewall: Remote Access Quarantine Service)

RPC over HTTP Proxy (Windows Firewall: RPC over HTTP Proxy)

**S**
Security auditing (Windows Firewall: Security auditing)

Simple Network Management Protocol (Windows Firewall: Simple Network Management Protocol)

**T**
No entries

**U**
UDDI Services Database Components (Windows Firewall: UDDI Services Database Components)

UDDI Services Web Server Components (Windows Firewall: UDDI Services Web Server Components)

Update Root Certificates (Windows Firewall: Update Root Certificates)

**V**
No entries

**W**
Windows Media Services Administrator for the Web (Windows Firewall: Windows Media Services Administrator for the Web)

WMI SNMP Provider (Windows Firewall: WMI SNMP Provider)

WMI Windows Installer Provider (Windows Firewall: WMI Windows Installer Provider)

**X**
No entries

**Y**
No entries

**Z**
No entries

# Windows Firewall Settings: Remote Administration Tools

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Firewall is a stateful host firewall that blocks all unsolicited incoming TCP/IP traffic, including Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic. If you enable Windows Firewall on a server and you want to manage the server with a remote administration tool, usually you must configure Windows Firewall settings on the server so that the server can receive unsolicited traffic from the remote administration tool. You might also have to configure Windows Firewall settings on the computer that is running the remote administration tool if Windows Firewall is enabled on that computer. Although you can configure Windows Firewall settings manually, the recommended method is to use the Security Configuration Wizard (SCW).

To use a remote administration tool with Windows Firewall, you usually need to add a program or port to the Windows Firewall exceptions list. When you add a program or port to the exceptions list, you instruct Windows Firewall to allow unsolicited incoming traffic to reach the specified program or pass through the specified port. In some cases, you might need to configure a registry setting or enable one of the preconfigured Windows Firewall exceptions, such as the File and Printer Sharing exception or the Remote Administration exception.

| **Caution** |
| --- |
| The Remote Administration exception allows traffic through numerous ports, which can make your computer more accessible to attack. Be sure to read the Windows Firewall documentation so that you understand the risks of using the Remote Administration exception. Incorrectly editing the registry can severely damage your system. Before you make changes to the registry, you should back up any valued data on the computer. |

For more information about configuring Windows Firewall exceptions, see Managing Program, Port, and Service Exceptions (http://go.microsoft.com/fwlink/?LinkId=43261) and Help: Understanding Windows Firewall Exceptions (http://go.microsoft.com/fwlink/?LinkId=43260). For more information about SCW, see Security Configuration Wizard Overview on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43259).

Use the A-Z list to find out how to configure Windows Firewall for use with the following remote administration tools.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**

Active Directory Domains and Trusts (Windows Firewall: domain)

Active Directory Management (Windows Firewall: admgmt)

Active Directory Schema Management (Windows Firewall: schmmgmt)

Active Directory Sites and Services (Windows Firewall: dssite)

Active Directory Users and Computers (Windows Firewall: dsa)

Authorization Manager (Windows Firewall: azman)

**B**

No entries

**C**

Certificate Templates (Windows Firewall: certtmpl)

Certificates (Windows Firewall: certmgr)

Certification Authority (Windows Firewall: certsrv)

Certutil command (Windows Firewall: certutil)

Cluster Administrator (Windows Firewall: cluadmin)

Cluster command (Windows Firewall: cluster)

Component Services (Windows Firewall: comexp)

Computer Management (Windows Firewall: compmgmt)

Connection Manager Administration Kit Binaries (Windows Firewall: cmbins)

Connection Manager Administration Kit Wizard (Windows Firewall: cmak)

**D**

Device Manager (Windows Firewall: devmgr)

Dfscmd command (Windows Firewall: dfscmd)

DHCP (Windows Firewall: dhcpmgmt)

Directory Service Utilities (Windows Firewall: ntdsutil)

Disk Defragmenter (Windows Firewall: dfrg)

Disk Management (Windows Firewall: diskmgmt)

Distributed File System (Windows Firewall: dfsgui)

DNS Management (Windows Firewall: dnsmgmt)

Dsadd command (Windows Firewall: dsadd)

Dsget command (Windows Firewall: dsget)

Dsmod command (Windows Firewall: dsmod)

Dsmove command (Windows Firewall: dsmove)

Dsquery command (Windows Firewall: dsquery)

Dsrm command (Windows Firewall: dsrm)

**E**

Event Viewer (Windows Firewall: eventvwr)

**F**

Fax client console (Windows Firewall: fxsclnt)

Fax Service Manager (Windows Firewall: fxsadmin)

File Server Management (Windows Firewall: filesvr)

**G**

Group Policy Object Editor (Windows Firewall: gpedit)

**H**

No entries

**I**

IIS Application Management script (Windows Firewall: iisapp)

IIS Backup script (Windows Firewall: iisback)

IIS Configuration script (Windows Firewall: iiscnfg)

IIS FTP script (Windows Firewall: iisftp)

IIS FTP Virtual Directory script (Windows Firewall: iisftpdr)

IIS Help script (Windows Firewall: iisschlp)

IIS Service Extension script (Windows Firewall: iisext)

IIS Virtual Directory script (Windows Firewall: iisvdir)

IIS Web Management script (Windows Firewall: iisweb)

Indexing Service (Windows Firewall: ciadv)

Internet Authentication Service (Windows Firewall: iasmsc)

Internet Information Services (IIS) Manager (Windows Firewall: iis)

IP Security Monitor (Windows Firewall: ipsecmon)

IP Security Policies (Windows Firewall: ipsecpol)

**J**
No entries

**K**
No entries

**L**
Local Security Settings (Windows Firewall: secpol)

Local Users and Groups (Windows Firewall: lusrmgr)

**M**
No entries

**N**
Network Load Balancing Manager (Windows Firewall: nlbmgr)

Network Monitor tools (Windows Firewall: netmon)

**O**
No entries

**P**
Performance (Windows Firewall: perfmon)

POP3 Service (Windows Firewall: p3server)

Public Key Management (Windows Firewall: pkmgmt)

**Q**
No entries

**R**
Remote Desktops (Windows Firewall: tsmmc)

Remote Storage (Windows Firewall: rsadmin)

Removable Storage Operator Requests (Windows Firewall: ntmsoprq)

Removable Storage (Windows Firewall: ntmsmgr)

Resultant Set of Policy (Windows Firewall: rsop)

Routing and Remote Access (Windows Firewall: rrasmgmt)

**S**
Security Configuration and Analysis (Windows Firewall: sca)

Services (Windows Firewall: services)

Shared Folders (Windows Firewall: fsmgmt)

**T**
Telephony (Windows Firewall: tapimgmt)

Terminal Services Configuration (Windows Firewall: tscc)

Terminal Services Manager (Windows Firewall: tsadmin)

**U**

UDDI Services Console (Windows Firewall: uddi)

**V**
No entries

**W**
Windows Management Infrastructure (Windows Firewall: wmimgmt)

Windows Media Services (Windows Firewall: wmsadmin)

Windows Server 2003 Administration Tools Pack (Windows Firewall: adminpak)

WINS (Windows Firewall: winsmgmt)

Wireless Monitor (Windows Firewall: wiremon)

**X**
No entries

**Y**
No entries

**Z**
No entries

# Windows Firewall Settings: Server Roles

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Firewall is a stateful host firewall that blocks all unsolicited incoming TCP/IP traffic, including Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic. If you enable Windows Firewall on a server and you install or configure a server role that listens for unsolicited incoming traffic, you must usually configure Windows Firewall settings so that the server role can receive unsolicited traffic. Although you can configure Windows Firewall settings manually, the recommended method is to use the Security Configuration Wizard (SCW).

To use a server role work with Windows Firewall, you usually need to add one or more programs or ports to the Windows Firewall exceptions list. When you add a program or port to the exceptions list, you instruct Windows Firewall to allow unsolicited incoming traffic to reach the specified program or pass through the specified port. In some cases, you might need to configure a registry setting or enable one of the preconfigured Windows Firewall exceptions, such as the File and Printer Sharing exception or the Remote Administration exception.

| **Caution** |
| --- |
| The Remote Administration exception allows traffic through numerous ports, which can make your computer more accessible to attack. Be sure to read the Windows Firewall documentation so that you understand the risks of using the Remote Administration exception. Incorrectly editing the registry can severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

For more information about configuring Windows Firewall exceptions, see Managing Program, Port, and Service Exceptions (http://go.microsoft.com/fwlink/?LinkId=43261) and Help: Understanding Windows Firewall Exceptions (http://go.microsoft.com/fwlink/?LinkId=43260). For more information about SCW, see Security Configuration Wizard Overview on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43259).

Use the A-Z list to find out how to configure Windows Firewall for use with the following server roles.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**

Application Server (Windows Firewall: Application Server)

ASP.NET session state server (Windows Firewall: aspnetstate)

**B**

No entries

**C**

Certificate Server (Windows Firewall: Certificate Server)

Cluster server (Windows Firewall: Cluster server)

**D**

DFS server (Windows Firewall: DFS server)

DHCP server (Windows Firewall: DHCP server)

DNS server (Windows Firewall: DNS server)

Domain controller (Windows Firewall: Domain controller)

**E**

No entries

**F**

Fax server (Windows Firewall: Fax server)

File and Print server for NetWare (Windows Firewall: File and Print server for NetWare)

File server (Windows Firewall: File server)

File server for Macintosh (Windows Firewall: macfile)

FTP server (Windows Firewall: FTP server)

**G**
No entries

**H**
No entries

**I**
IAS server (Windows Firewall: IAS server)

Internet Connection Sharing server (Windows Firewall: Internet Connection Sharing server)

**J**
No entries

**K**
No entries

**L**
No entries

**M**
Message Queuing server (Windows Firewall: Message Queuing server)

Microsoft Identity Integration Server 2003 (Windows Firewall: Microsoft Identity Integration Server 2003)

Microsoft Operations Manager 2005 server (Windows Firewall: Microsoft Operations Manager 2005 server)

Microsoft Operations Manager 2005 Workgroup Edition server (Windows Firewall: Microsoft Operations Manager 2005 Workgroup Edition server)

Middle-tier application server (COM+) (Windows Firewall: Middle-tier application server)

MSDE database server for UDDI (Windows Firewall: MSDE database server for UDDI)

**N**
NNTP (Usenet news) server (Windows Firewall: NNTP (Usenet news) server)

**O**
No entries

**P**
POP3 server (Windows Firewall: POP3 server)

Print server (Windows Firewall: Print server)

Print server for Macintosh (Windows Firewall: Print server for Macintosh)

Print server for Unix (LPD) (Windows Firewall: Print server for Unix (LPD))

**Q**
No entries

**R**
Remote access/VPN server (Windows Firewall: Remote access/VPN server)

Remote installation server (Windows Firewall: Remote installation server)

Remote Storage server (Windows Firewall: remotestorageserver)

**S**
SMTP server (Windows Firewall: SMTP server)

SNMP server (Windows Firewall: SNMP server)

SNMP trap server (Windows Firewall: SNMP trap server)

**T**
Telnet server (Windows Firewall: Telnet server)

Terminal server (Windows Firewall: Terminal server)

**U**
UDDI Web server (Windows Firewall: UDDI Web server)

**V**
No entries

**W**
Web server (Windows Firewall: Web server)

Windows Media server (Windows Firewall: Windows Media server)

WINS server (Windows Firewall: WINS server)

**X**
No entries

**Y**
No entries

**Z**
No entries

# Windows Firewall Settings: Services

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Firewall is a stateful host firewall that blocks all unsolicited incoming TCP/IP traffic, including Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic. If you enable Windows Firewall on a server and the server uses a service that listens for unsolicited incoming traffic, you must usually configure Windows Firewall settings so that the service can receive unsolicited traffic. Although you can configure Windows Firewall settings manually, the recommended method is to use the Security Configuration Wizard (SCW).

To use a service with Windows Firewall, you usually need to add a program or port to the Windows Firewall exceptions list. When you add a program or port to the exceptions list, you instruct Windows Firewall to allow unsolicited incoming traffic to reach the specified program or pass through the specified port. In some cases, you might need to change a registry setting or enable one of the preconfigured Windows Firewall exceptions, such as the File and Printer Sharing exception or the Remote Administration exception.

| **Caution** |
| --- |
| The Remote Administration exception allows traffic through numerous ports, which can make your computer more accessible to attack. Be sure to read the Windows Firewall documentation so that you understand the risks of using the Remote Administration exception. Incorrectly editing the registry can severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

For more information about configuring Windows Firewall exceptions, see Managing Program, Port, and Service Exceptions (http://go.microsoft.com/fwlink/?LinkId=43261) and Help: Understanding Windows Firewall Exceptions (http://go.microsoft.com/fwlink/?LinkId=43260). For more information about SCW, see Security Configuration Wizard Overview on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=43259).

Use the A-Z list to find out how to configure Windows Firewall for use with the following services.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**

Alerter (Windows Firewall: alerter)

Application Layer Gateway Service (Windows Firewall: alg)

Application Management (Windows Firewall: appmgmt)

ASP.NET State Service (Windows Firewall: aspnetstate)

Automatic Updates (Windows Firewall: wuausrv)

**B**

No entries

**C**

Certificate Services (Windows Firewall: certsvc)

Client Service for NetWare (Windows Firewall: nwcworkstation)

ClipBook (Windows Firewall: clipsrv)

Cluster Service (Windows Firewall: clussvc)

COM+ Event System (Windows Firewall: eventsystem)

COM+ System Application (Windows Firewall: comsysapp)

Computer Browser (Windows Firewall: browser)

Cryptographic Services (Windows Firewall: cryptsvc)

**D**

DCOM Server Process Launcher (Windows Firewall: dcomlaunch)

DHCP Client (Windows Firewall: dhcp)

DHCP Server (Windows Firewall: dhcpserver)

Distributed File System (Windows Firewall: dfs)

Distributed Link Tracking Client (Windows Firewall: trkwks)

Distributed Link Tracking Server (Windows Firewall: trksvr)

Distributed Transaction Coordinator (Windows Firewall: msdtc)

DNS Client (Windows Firewall: dnscache)

DNS Server (Windows Firewall: DNS server)

**E**

Error Reporting Service (Windows Firewall: ersvc)

Event Log (Windows Firewall: eventlog)

**F**

Fax (Windows Firewall: fax)

File and Print Services for Netware (Windows Firewall: fpsvcnetware)

File Replication (Windows Firewall: ntfrs)

File Server Macintosh (Windows Firewall: File Server Macintosh)

FTP Publishing Service (Windows Firewall: msftpsvc)

**G**

No entries

**H**

Help and Support (Windows Firewall: helpsvc)

Human Interface Device Access (Windows Firewall: hidserv)

**I**

IAS Jet Database Access (Windows Firewall: iasjet)

IASJet (used by RRAS) (Windows Firewall: iasjetrras)

IIS Admin Service (Windows Firewall: iisadmin)

IMAPI CD-Burning COM Service (Windows Firewall: imapiservice)

Indexing Service (Windows Firewall: cisvc)

Infrared Monitor (Windows Firewall: irmon)

Internet Authentication Service (Windows Firewall: ias)

Intersite Messaging (Windows Firewall: ismserv)

IPSEC Services (Windows Firewall: policyagent)

IPv6 Helper Service (Windows Firewall: 6to4)

**J**

No entries

**K**

Kerberos Key Distribution Center (Windows Firewall: kdc)

**L**

License Logging (Windows Firewall: licenseservice)

Logical Disk Manager (Windows Firewall: dmserver)

Logical Disk Manager Administrative Service (Windows Firewall: dmadmin)

**M**

Message Queuing (Windows Firewall: msmq)

Message Queuing Downlevel Client (Windows Firewall: mqds)

Message Queuing Triggers (Windows Firewall: msmqtriggers)

Messenger (Windows Firewall: messenger)

Microsoft Operations Manager (Windows Firewall: mom)

Microsoft POP3 Service (Windows Firewall: pop3svc)

Microsoft Software Shadow Copy Provider (Windows Firewall: swprv)

**N**

Net Logon (Windows Firewall: netlogon)

NetMeeting Remote Desktop Sharing (Windows Firewall: mnmsrvc)

Network Connections (Windows Firewall: netman)

Network DDE (Windows Firewall: netdde)

Network DDE DSDM (Windows Firewall: netddedsdm)

Network Location Awareness (NLA) (Windows Firewall: nla)

Network News Transport Protocol (Windows Firewall: nntpsvc)

Network Provisioning Service (Windows Firewall: xmlprov)

NTLM Security Support Provider (Windows Firewall: ntlmssp)

**O**

No entries

**P**

Performance Logs and Alerts (Windows Firewall: sysmonlog)

Plug and Play (Windows Firewall: plugplay)

Portable Media Serial Number Service (Windows Firewall: wmdmpmsn)

Print Server for Macintosh (Windows Firewall: macprint)

Print Spooler (Windows Firewall: spooler)

Protected Storage (Windows Firewall: protectedstorage)

**Q**

No entries

**R**

Remote Access Connection Manager (Windows Firewall: rasman)

Remote Administration Service (Windows Firewall: srvcsurg)

Remote Desktop Help Session Manager (Windows Firewall: rdsessmgr)

Remote Installation (Windows Firewall: binlsvc)

Remote Procedure Call (RPC) (Windows Firewall: rpcss)

Remote Registry (Windows Firewall: remoteregistry)

Remote Server Manager (Windows Firewall: appmgr)

Remote Storage Notification (Windows Firewall: remotestorageuserlink)

Remote Storage Server (Windows Firewall: remotestorageserver)

Removable Storage (Windows Firewall: ntmssvc)

Resultant Set of Policy Provider (Windows Firewall: rsopprov)

Routing and Remote Access (Windows Firewall: remoteaccess)

**S**

SAP Agent (Windows Firewall: nwsapagent)

Secondary Logon (Windows Firewall: seclogon)

Security Accounts Manager (Windows Firewall: samss)

Server (Windows Firewall: lanmanserver)

Shell Hardware Detection (Windows Firewall: shellhwddetection)

Simple Mail Transfer Protocol (SMTP) (Windows Firewall: smptsvc)

Simple TCP/IP Services (Windows Firewall: simptcp)

Single Instance Storage Groveler (Windows Firewall: groveler)

Smart Card (Windows Firewall: scardsvr)

SNMP Service (Windows Firewall: snmp)

SNMP Trap Service (Windows Firewall: snmptrap)

Special Administration Console Helper (Windows Firewall: sacsvr)

SSL for World Wide Web Publishing (Windows Firewall: sslwww)

System Event Notification (Windows Firewall: sens)

**T**

Task Scheduler (Windows Firewall: schedule)

TCP/IP NetBIOS Helper (Windows Firewall: lmhosts)

TCP/IP Print Server (Windows Firewall: lpdsvc)

Telephony (Windows Firewall: tapisrv)

Telnet (Windows Firewall: tlntsvr)

Terminal Server Licensing (Windows Firewall: termservlicensing)

Terminal Services (Windows Firewall: termservice)

Terminal Services Session Directory (Windows Firewall: tssdis)

Trivial FTP Daemon (Windows Firewall: tftpd)

**U**

Uninterruptible Power Supply (Windows Firewall: ups)

**V**

Virtual Disk Service (Windows Firewall: vds)

Volume Shadow Copy (Windows Firewall: vss)

**W**

Web Usage Logging (Windows Firewall: webusage)

WebClient (Windows Firewall: webclient)

Windows Audio (Windows Firewall: audiosrv)

Windows Firewall/Internet Connection Sharing (Windows Firewall: sharedaccess)

Windows Image Acquisition (WIA) (Windows Firewall: stisvc)

Windows Installer (Windows Firewall: msiserver)

Windows Internet Name Service (WINS) (Windows Firewall: wins)

Windows Management Instrumentation (Windows Firewall: winmgmt)

Windows Management Instrumentation Driver Extensions (Windows Firewall: wmi)

Windows Media Services (Windows Firewall: wmserver)

Windows System Resource Manager (Windows Firewall: winsysresmgr)

Windows Time (Windows Firewall: w32time)

Windows User Mode Driver Framework (Windows Firewall: umwdf)

WinHTTP Web Proxy Auto-Discovery Service (Windows Firewall: winhttpautoproxysvc)

Wireless Configuration (Windows Firewall: wzcsvc)

WMI Performance Adapter (Windows Firewall: wmiapsrv)

Workstation (Windows Firewall: lanmanworkstation)

World Wide Web Publishing (Windows Firewall: w3svc)

**X**
No entries

**Y**
No entries

**Z**
No entries

# IIS 6.0 Technical Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Welcome to Internet Information Services (IIS) 6.0 Technical Reference from Microsoft.

Microsoft® Windows® Server™ 2003 and Internet Information Services (IIS) 6.0 provide the services to support a secure, available, and scalable Web server on which to run your Web sites and applications. Whether you manage a single Web server or many, Internet Information Services (IIS) 6.0 Technical Reference will help you effectively plan, deploy, operate, and troubleshoot your IIS 6.0 solution. This comprehensive technical resource delivers an in-depth description of the new IIS 6.0 architecture, as well as reference information about IIS 6.0 features and services. It also includes practical information and tools to help you accomplish everyday administrative tasks.

This guide contains the following sections:


Introducing IIS 6.0

IIS 6.0 Architecture

Running IIS 6.0 as an Application Server

Working with the Metabase

Managing a Secure IIS 6.0 Solution

Optimizing IIS 6.0 Performance

Web Server Scalability

Configuring Internet Sites and Services

IIS 6.0 Administration Scripts, Tips, and Tricks

Analyzing Log Files

Troubleshooting IIS 6.0

Common IIS 6.0 Administrative Tasks

Using FrontPage 2002 Server Extensions with IIS 6.0

IPv6 and IIS 6.0

Running 32-bit Applications on 64-bit Windows

IIS 6.0 Administration Reference

# Introducing IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0 is a significant new release. Completely redesigned — from its core components and architecture to its default locked-down installation — IIS 6.0 offers organizations significant advantages over previous versions for reliability, manageability, security, scalability, and performance. Whether you manage a single Web server or are a system administrator at a large Internet service provider, *Internet Information Services (IIS) 6.0 Technical Reference* provides the conceptual and practical information you need to work with IIS 6.0. You will find information about how IIS 6.0 is different from previous versions; how IIS 6.0 works and integrates with the Microsoft® Windows® Server 2003 operating system; and how to manage, maintain, and troubleshoot IIS.

In This Section

Internet Information Services 6.0

What's Changed in IIS 6.0

IIS 6.0 Resource Kit Tools

Additional Resources for IIS 6.0

# Internet Information Services 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 is tightly integrated with the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Web Edition; and Windows® Server 2003, Datacenter Edition operating systems. Along with Windows Server 2003, IIS 6.0 provides Web server capabilities over an intranet, the Internet, or an extranet. A versatile tool for creating a reliable communications platform of dynamic network applications, IIS 6.0 can be deployed in many different environments. Small businesses and organizations might use one server running IIS 6.0 to host a single Web site and a small intranet. Large organizations, including Internet service providers (ISPs), corporations, and educational institutions, can manage many application servers, each hosting thousands of Web sites and applications.

Organizations also use IIS 6.0 to host File Transfer Protocol (FTP) sites and route news or e-mail by using the Network News Transfer Protocol (NNTP) and the Simple Mail Transfer Protocol (SMTP). Because IIS 6.0 supports the latest Web standards, such as ASP.NET, XML, and Simple Object Access Protocol (SOAP), it is a robust platform for the development, implementation, and management of new Web sites and applications.

Installing IIS as an Application Server

To help protect your systems from attackers who target unused or unmonitored services, IIS 6.0 is not installed on Windows Server 2003 by default. You must install IIS 6.0 manually when you are ready to deploy your Web solutions.

To install IIS 6.0, you can use the Configure Your Server wizard, which is included in Windows Server 2003. The wizard allows you to install or remove Windows Server 2003 services that are grouped into different *server roles*. These server roles include file server, print server, terminal server, DNS server, DHCP server, and application server. The application server role can include the following services:

- Internet Information Services (IIS) 6.0
- COM+ for remote transactions
- Distributed Transaction Coordinator (DTC) for remote access
- ASP.NET
- FrontPage Server Extensions

For more information about the application server role, see Running IIS 6.0 as an Application Server. For information about installing IIS, see Internet Information Services (IIS) 6.0 overview in Help and Support Center for Windows Server 2003.

Enabling Web Service Extensions to Serve Dynamic Content

When you initially install IIS, Web service extensions are locked down, which means that IIS 6.0 only serves static content until you enable specific extensions. To protect your server, select and enable features like ASP, ASP.NET, Server-Side Includes, WebDAV publishing, and FrontPage® 2002 Server Extensions from Microsoft *only if you need them*. If you do not enable these features before you try to serve dynamic content, IIS returns an HTTP 404 level error message (File or directory not found).

For information about enabling Web service extensions, see Enabling Dynamic Content in IIS 6.0 Operations Guide, which is accessible from IIS Manager.

# What's Changed in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 core components and architecture are redesigned, which makes this version of IIS significantly different from previous versions. IIS 6.0 meets organizations demands for the following improvements:

- **Reliability.** Greater reliability by providing a way to isolate one application from another.
- **Manageability.** Better manageability by providing a configurable XML-based metabase.
- **Security.** Tighter security by providing a locked-down security strategy.
- **Performance and scalability.** Improved performance by providing isolation in-process. Built-in scalability by offering an architecture that allows you to consolidate sites and applications on fewer servers.

# Reliability Improvements

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 promotes greater reliability and reduces the chance of defective applications compromising overall server performance and application reliability. Applications that do not function properly can adversely affect any Web server, regardless of the reason for the applications faulty functionality. For example, one application can consume CPU resources, causing other applications to perform slowly. In addition, a Web application that fails can potentially bring down a Web server, including all its sites and applications. Web server administrators must sometimes restart the entire Web server, bringing down all its applications and sites, in an attempt to update or fix one faulty application. IIS 6.0 offers solutions for all these scenarios.

The IIS 6.0 redesigned architecture, which includes two new request processing models, worker process isolation mode and IIS 5.0 isolation mode, helps you avoid the problems caused by faulty applications. Worker process isolation mode, which is the default application isolation mode for IIS 6.0, separates key components of the World Wide Web Publishing Service (WWW service) from the effects of malfunctioning applications by using isolated worker processes. In IIS 5.0 isolation mode, requests are processed similarly to how they were processed in IIS 5.0. This process model is intended for existing applications that do not initially work when you test them during deployment testing in worker process isolation mode. For more information about IIS architecture changes, see IIS 6.0 Architecture.

Worker process isolation mode is the default process model in IIS 6.0 because it enables the greatest number of benefits, many of which are achieved by using application pools. An *application pool* is a grouping of one or more URLs that are served by a worker process or set of worker processes. You can configure specific settings for each application pool and the worker processes that serve it.

If your server is running applications that contain imperfect code, using worker process isolation mode allows you to isolate a defective application in an application pool. You can then use health monitoring (a feature that makes an application pool self-healing), worker process recycling (periodic restarts of unstable applications or ones that leak memory), and rapid-fail protection (auto-restart of failed applications) to specifically address the applications faulty functionality, which results in fewer or no system restarts. For more information about how IIS improves reliability of applications, see Running IIS 6.0 as an Application Server.

# Manageability Improvements

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Whether you are responsible for one Web server or for many clusters of Web servers, you can find new features and improvements in IIS 6.0 that will make your Web server easier to manage than with previous versions of IIS. For example, changes in the IIS metabase, additions to programmatic administration, logging enhancements, and debugging improvements can help you manage your Web servers more efficiently. From propagating configuration changes across multiple servers to pinpointing the problem in a slow-to-respond Web site, you can administer your Web servers more quickly and easily with IIS 6.0.

Metabase Improvements

The metabase configuration and schema for IIS 4.0 and IIS 5.0 were stored in a binary file, which was difficult to read or edit. IIS 6.0 replaces the single binary file (MetaBase.bin), with two XML files, which are named MetaBase.xml and MBSchema.xml. In IIS 6.0, administrators who want to manage IIS by editing the metabase directly can use a simple text editor, such as Microsoft Notepad, to edit MetaBase.xml (the metabase configuration file). Editing the metabase configuration file directly is a good solution for administrators who do not want to use scripts or code to administer IIS, and it is a faster administration solution than using IIS Manager when you are working remotely over a slow network connection. For more information about the IIS metabase, see Working with the Metabase.

Programmatic Administration Additions

Many administrators prefer to manage IIS programmatically by using scripts and batch files, rather than by using IIS Manager, which provides a graphical user interface. Administrators working in an ISP environment with many servers — both local and remote — and thousands of Web sites, can save time by managing IIS programmatically.

IIS 6.0 includes a Windows Management Instrumentation (WMI) provider that allows administrators to control services and applications programmatically. In addition, you can still administer IIS programmatically by using Active Directory Services Interfaces (ADSI) and other custom scripts and batch files. For more information about ADSI and WMI, see IIS 6.0 Administration Scripts, Tips, and Tricks.

| **Note** |
| --- |
| In previous releases of IIS, IIS Manager was called the Internet Service Manager. For information about using IISManager to administer IIS, see Common Administrative Tasks. |

Logging Enhancements

You can manage and troubleshoot your Web server more effectively if you understand how to read and analyze log files. In IIS 6.0, logging differs from previous versions in several ways, so be sure to familiarize yourself with the changes before you configure logging on your server.

In IIS 5.0, Inetinfo.exe (a user-mode process) performed all logging. In the redesigned architecture of IIS 6.0, if a request is cached, the response is served from the HTTP.sys kernel-mode cache without passing through Inetinfo.exe or another user-mode component; therefore, HTTP.sys handles most logging in IIS 6.0 (Open Database Connectivity [ODBC] and custom logging are the exceptions) to ensure that responses are recorded.

IIS 6.0 provides the following additional enhancements to logging:

- Substatus error code logging allows you to record substatus error codes for requests that fail (substatus error codes are not returned to the client).
- Centralized binary logging provides a way to record detailed information about all your Web sites but to use a minimum of system resources. For example, you can use a single log file to record data for all your sites.
- HTTP.sys error logging ensures that all errors generated by HTTP.sys can be recorded, including error reasons.
- Remote logging allows you to write log data to a remote share over a network by using a Universal Naming Convention (UNC) path.

For more information about IIS logging, see Analyzing Log Files.

Troubleshooting

Troubleshooting and debugging can be time-consuming tasks for an administrator. To help you troubleshoot effectively, be sure you are aware of the changes and new features in IIS 6.0 and Windows Server 2003 that can make troubleshooting easier. For example, Windows Server 2003 provides a new framework, Enterprise Tracing for Windows (ETW), which you can use for debugging and capacity planning. IIS 6.0 (through ETW) implements several different tracing providers to trace key transitions for an HTTP request, following its progress through the IIS service framework. These tracing providers include HTTP.sys, IIS Web Server, IIS ISAPI Filter, IIS ISAPI Extension, and IIS ASP traces. The information generated by these providers can help you determine where, in the lifetime of a request, a problem is occurring. For more information about troubleshooting, including how to use new and existing features, tools, and techniques with IIS 6.0, see Troubleshooting IIS 6.0.

# Security Improvements

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Attackers and malicious users can take advantage of a system with services that are running but not being used. Such an attack is possible because if administrators do not use a service, they might forget to maintain it with current hotfixes, service packs, and patches. To reduce this security risk, IIS 6.0 is locked down by default — only request handling for static Web pages is enabled, and only the WWW service is installed. None of the applications that run on IIS — including ASP, ASP.NET, Common Gateway Interface (CGI) scripting, FrontPage 2002 Server Extensions, and WebDAV — are turned on by default.

IIS 6.0 improves security in other ways as well. Many aspects of IIS 6.0, including default functionality and settings, perform different than they did in previous versions. These changes can result in existing applications or sites performing in unexpected ways. To save troubleshooting time, familiarize yourself with the security features in IIS 6.0. For more information about IIS 6.0 security, see Managing a Secure IIS 6.0 Solution.

# Performance and Scalability Improvements

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you want to host many applications or Web sites on the same server or if you have applications that consume large amounts of system resources, such as bandwidth or CPU processing, you can host those applications on IIS 6.0 and still obtain good performance. IIS 6.0 architecture provides the foundation for a number of performance enhancements in the Web server. For example, Web gardens, processor affinity, idle timeout, bandwidth throttling, HTTP compression, and CPU monitoring are a few of the features that help you conserve system resources while optimizing performance.

When your Web server is running in worker process isolation mode, you can create *Web gardens*, which enable multiple worker processes to serve one application pool. Web gardens can improve performance because if one worker process in an application pool is busy processing a request, other worker processes can also accept and process requests for that application pool.

IIS 6.0 enables you to configure application pools to establish affinity between worker processes and specific processors in a multiprocessor computer. Processor affinity, when set, enables IIS to efficiently use the caches on multiprocessor servers. Also, you can configure an application pool to have its worker processes request a shutdown if the worker processes are idle for a configurable amount of time. Shutting down idle worker processes frees up unused resources, allowing CPU resources to be used where needed. For more information about idle timeout, see Running IIS 6.0 as an Application Server.

If the network or Internet connection that your Web server uses is also used by other services, such as e-mail or news, you can monitor (and limit) the CPU usage of individual worker processes so that a portion of CPU resources are available for those other services. If your Web server hosts more than one Web site, you can individually throttle the bandwidth used by each site. In turn, you can make best use of available bandwidth by enabling HTTP compression, which allows faster transmission of pages between the Web server and compression-enabled clients. Also, IIS 6.0 supports HTTP Keep-Alives, an HTTP specification that increases server efficiency by reducing server activity and resource consumption.

If your server hosts applications that have workloads requiring large amounts of cached data, you can configure IIS 6.0 to cache up to 64 gigabytes (GB) on an x86-based (or compatible processor) system. For more information about these features and how to improve Web server performance, see Optimizing IIS 6.0 Performance.

IIS supports Windows scalability solutions, such as clustering and network load balancing; however, IIS 6.0 also provides a number of specific changes that can improve Web server scalability. Scalability is particularly important for organizations that want to reduce overhead and maintenance costs by consolidating sites and applications on fewer servers.

Another way to extend a system is to store content in a remote location, such as on a UNC–based drive. IIS 6.0 provides caching, authentication, and administration changes that make it easier for you to create, secure, and manage centralized content storage.

For more information about these and other scalability features in IIS 6.0, see Web Server Scalability.

# IIS 6.0 Resource Kit Tools

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 Resource Kit Tools includes 14 tools to help you implement and administer IIS 6.0 more effectively. You can install the IIS 6.0 Resource Kit Tools, which includes Help documentation, by running IIS60rkt.exe from the Microsoft Download Center. When you perform a default installation of the Resource Kit Tools, most of the files are installed in the *LocalDrive*:\Program Files\IIS Resources folder. A few tools are installed in other directories. Read the Help documentation that is included with the tools if you want to know the installation directory for a specific tool.

After installing the tools, you can access the tools and their documentation from the **Start** menu (point to **All Programs** and then **IIS Resources**).

The IIS 6.0 Resource Kit Tools provides the following tools for migrating to IIS 6.0, troubleshooting IIS, testing and improving IIS performance, applying IIS security settings, and administering IIS.

| **Important** |
| --- |
| These tools are provided as-is without a warranty of any kind. Use these tools at your own risk. Microsoft disclaims any implied warranty of merchantability and fitness for a particular purpose. |

Migration Tools

The IIS 6.0 Resource Kit Tools includes two migration tools, the IIS 6.0 Migration Tool and the Apache to IIS 6.0 Migration Tool.

**IIS 6.0 Migration Tool (IISmt.exe)**

Use the IIS 6.0 Migration Tool to migrate Web site content and configuration settings from a server running IIS 4.0, IIS 5.0, or IIS 6.0 to a server running IIS 6.0. The IIS 6.0 Migration tool saves time and reduces the chance of errors by automating many of the repetitive processes involved with moving sites manually. This command-line tool is not documented in Resource Kit Tools Help, but a separate user guide is included in the IIS 6.0 Migration Tool folder and is also available from the Start menu.

**Apache to IIS 6.0 Migration Tool (Wsmk_targ_virdir.exe and Wsmk_windows.exe)**

Use this tool to migrate Web sites — for example, FrontPage Server Extensions, .htaccess files, and home directories — from Linux-based servers that run Apache Web applications to servers running IIS 6.0 on Windows Server 2003. Use this Perl-based command-line tool to migrate both Web content and configuration settings. This migration tool supports a variety of both Linux and Apache versions. The tool is not documented in Resource Kit Tools Help, but a separate user guide is included in the Apache to IIS 6.0 Migration Tool folder and is also available from the Start menu.

Performance and Troubleshooting Tools

The IIS 6.0 Resource Kit Tools provides several tools to help you improve performance and to troubleshoot an application server that is running IIS 6.0. These tools include the Web Capacity Analysis Tool (WCAT), Log Parser, IISState, Wfetch, and TinyGet.

**WCAT (Wcclient.exe and Wcctl.exe)**

Use the two WCAT tools, Wcclient.exe and Wcctl.exe, as a stress client to simulate clients and as a controller to initiate and monitor the testing. WCAT can test how your IIS 6.0-based server and your network configuration respond to a variety of client requests for content, data, or Web pages. Then use the test results to determine the optimal server and network configuration for your environment.

**Log Parser (LogParser.exe or LogParser.dll)**

Use Log Parser to perform many tasks related to log files, such as filtering entries, converting log files to other formats, and performing data mining. Log Parser supports many different input formats, including all the IIS log file formats and multiple output formats, such as text files and database tables. You can use Log Parser to obtain logging information for WCAT and TinyGet. Log Parser is available as a command-line tool (LogParser.exe) and as a set of COM objects that support scripting (LogParser.dll).

**IISState (IISState.exe and IISState.log)**

Use IISState to troubleshoot slow performance and application failures for applications running on IIS 6.0. IISState can also help you identify the causes of ASP0115 error messages and the causes of errors detected by Dr. Watson. For example, IISState can analyze an application like Inetinfo.exe as it executes, record all running threads, and display and log the analysis. IISState

also identifies each thread, and displays the raw thread stack and the time that each thread spent executing in kernel and user modes. IISState is a command-line tool.

**WFetch**

Use WFetch to troubleshoot, test, and log HTTP connections between a server and client. For example, you can use Wfetch to test the performance of new Web sites or Web sites that contain new elements, such as ASP or wireless protocols. Because Wfetch provides only HTTP connection information, it provides a faster way to analyze and troubleshoot HTTP performance than by running a Network Monitor trace. Wfetch provides a graphical user interface (GUI).

**TinyGet (TinyGet.exe)**

Use TinyGet to troubleshoot, test, and log HTTP connections between a server and client. With TinyGet.exe, you can customize each test request by configuring many different factors, including the authentication method, HTTP version, and output format. You can also use scripts that specify looping and multithreading. TinyGet.exe is a command-line tool.

Security Tools

The IIS 6.0 Resource Kit Tools provides several tools to help secure your Web server and troubleshoot security issues, including CustomAuth, Permissions Verifier, SelfSSL, and IISCertDeploy.vbs.

**CustomAuth (CustomAuth.dll and CustomAuth.ini)**

Use CustomAuth as a Web client authentication alternative to Basic, NTLM, and other IIS authentication methods. With CustomAuth, you can create a custom logon form for collecting user credentials, or you can use or modify the logon form that is included with the tool. As an ISAPI filter that must also be installed as a wildcard application map, CustomAuth runs at the beginning of every request, regardless of the extension of the requested file. Because CustomAuth allows users to log off after they complete a session, it is a good solution for a public kiosk.

**Permissions Verifier (Permissions Verifier.js)**

Use Permissions Verifier and its accompanying XML sample to verify that IIS is configured with the minimum permissions required to run. The tool checks the access control lists (ACLs) set for users and groups to help you determine whether permissions issues are causing any Web server problems.

**SelfSSL (SelfSSL.exe)**

Use SelfSSL to generate and install a self-signed Secure Sockets Layer (SSL) certificate. Because SelfSSL.exe generates a self-signed certificate that does not originate from a commonly trusted source, use this tool only when you need to troubleshoot third-party certificate problems or when you need to create a secure private channel between your server and a limited, known group of users, such as exists in a software test environment. SelfSSL.exe is a command-line tool.

**IISCertDeploy.vbs**

Use IISCertDeploy.vbs to deploy and back up SSL certificates on servers running IIS 6.0.

Administration Tools

The IIS 6.0 Resource Kit Tools provides several tools to help you administer your Web server and manage your sites, including IIS Host Helper Service, RemapURL, and Metabase Explorer.

**IIS Host Helper Service (IISHostsvc.exe)**

Use IIS Host Helper Service to register IIS host header strings with name resolution services. You can use IIS Host Helper Service in corporate intranet settings when you want to allow users to access Web sites by using host header names, but you do not want to register each sites name statically with Windows Internet Name Service (WINS) or a DNS server. When you install the IIS 6.0 Resource Kit Tools, IISHostsvc.exe is placed in the *systemroot*\System32\Inetsrv directory, and an IIS Host Helper Service log file directory is created at *systemroot*\System32\LogFiles\Iishostsvc.

**RemapURL (RemapUrl.ini and RemapUrl.dll)**

Use RemapURL to redirect client HTTP requests for URLs. RemapURL is a sample ISAPI filter that redirects HTTP requests by using an ISAPI PREPROC_HEADERS notification. This method enables you to redirect client GET or POST requests in a way that is completely transparent to the user. You can also use RemapURL if you want to display custom error messages.

**Metabase Explorer**

Use Metabase Explorer for viewing and editing the IIS 6.0 metabase. You can also use this tool to edit security settings for keys, export and import keys and subkeys, copy and paste keys and subkeys, and compare records. Metabase Explorer provides a GUI.

# Additional Resources for IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this section.

Related Information

- Analyzing Log Files for information about improvements to IIS logging, including centralized binary logging.
- Common IIS 6.0 Administrative Tasks for information about how to perform frequently used administrative tasks.
- Configuring Internet Sites and Services for information about configuring Web, FTP, NNTP, and SMTP sites and services.
- IIS 6.0 Administration Scripts, Tips, and Tricks for information about configuring IIS programmatically. Includes steps for creating and configuring a site from a batch file.
- IIS 6.0 Architecture for information about IIS 6.0 architecture as it pertains to the application isolation modes. Includes diagrams of how request processing occurs in worker process isolation mode and in IIS 5.0 isolation mode.
- Managing a Secure IIS 6.0 Solution for information about improvements to IIS security.
- Optimizing IIS 6.0 Performance for information about specific performance counters and tools for monitoring and tuning performance. Includes information about managing system resources, such as setting connection timeouts, configuring HTTP compression, and enabling bandwidth throttling.
- Running IIS 6.0 as an Application Server for information about installing IIS 6.0, choosing and configuring an application isolation mode, using health monitoring features and application pools, and managing Web server resources.
- Troubleshooting IIS 6.0 for information about troubleshooting IIS, including HTTP status codes.
- Web Server Scalability for information about improving scalability by optimizing IIS caches.
- Working with the Metabase for information about editing, backing up, and importing and exporting the IIS 6.0 metabase.
- IIS 6.0 Deployment Guide for information about deploying IIS 6.0.

Related IIS 6.0 Operations Guide Topics

- Common Administrative Tasks in IIS 6.0 for detailed information about how to complete many of the tasks described in this section.
- Enabling and Disabling Dynamic Content in IIS 6.0, which is accessible from IIS Manager, for information about enabling and disabling features in the Web Service Extensions list.
- Frequently Asked Questions about IIS 6.0 for basic information about IIS 6.0.
- Metabase Property Reference for reference information about metabase properties.
- Server Administration Guide for information about administering IIS 6.0.
- Site Setup for information about setting up IIS 6.0.
- Web Application Guide for information about using ASP and ASP.NET with IIS 6.0.

Related Windows Server 2003 Help Topics

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the**Search in title only** check box.

- Internet Information Services (IIS) 6.0 overview in Help and Support Center for Windows Server 2003 for information about installing IIS 6.0.

Related Tools

- Adsutil.vbs

  Obtain syntax, parameters, and examples for using Adsutil.vbs to configure and manage IIS from IIS 6.0 Administration Scripts, Tips, and Tricks.

- IIS 6.0 Resource Kit tools (IIS60rkt.exe)

- Windows Support Tools

  Obtain the Windows Support Tools by installing them from the \Support\Tools folder of the Windows Server 2003 operating system CD. See the Support Tools Help file for information about using these tools.

# IIS 6.0 Architecture

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) version 6.0, which runs on all editions of the Microsoft® Windows® Server 2003 operating system, provides a new architecture that offers flexibility in the choice of two application isolation modes. The new architecture helps you run a faster Web service that is more reliable and secure. Read this chapter to obtain a detailed review of the new architecture so that you can understand how IIS works and can effectively deploy, manage, and troubleshoot IIS 6.0.

**In This Section**

Overview of IIS 6.0 Architecture

IIS 6.0 Services

IIS 6.0 Core Components

Application Isolation Modes

How ASP.NET Works with IIS 6.0

Additional Resources for IIS 6.0 Architecture

**Related Information**

- For more information about running IIS as an application server, see Running IIS 6.0 as an Application Server.
- For more information about what has changed since IIS 5.0, see Introducing IIS 6.0.

# Overview of IIS 6.0 Architecture

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 provides a redesigned World Wide Web Publishing Service (WWW service) architecture that can help you achieve better performance, reliability, scalability, and security for your Web sites, whether they run on a single server running IIS or on multiple servers.

IIS 6.0 runs a server in one of two distinct request processing models, called application isolation modes. *Application isolation* is the separation of applications by process boundaries that prevents one application or Web site from affecting another and reduces the time that you spend restarting services to correct problems related to applications.

In IIS 6.0, application isolation is configured differently for each of the two IIS application isolation modes. Both modes rely on the *HTTP protocol stack* (also referred to as *HTTP.sys*) to receive Hypertext Transfer Protocol (HTTP) requests from the Internet and return responses. HTTP.sys resides in *kernel mode*, where operating system code, such as device drivers, runs. HTTP.sys listens for, and queues, HTTP requests. For more information about HTTP.sys, see HTTP Protocol Stack.

The new request-processing architecture and application isolation environment enables individual Web applications, which always run in *user mode*, to function within a self-contained worker process. A *worker process* is user-mode code whose role is to process requests, such as returning a static page or invoking an Internet Server API (ISAPI) extension or filter. Worker processes use HTTP.sys to receive requests and send responses over HTTP. For more information about worker processes, see Worker Processes.

**IIS 6.0 Request Processing Models**

*Worker process isolation mode* is the new IIS request processing model. In this application isolation mode, you can group Web applications into *application pools*, through which you can apply configuration settings to the worker processes that service those applications. An application pool corresponds to one request routing queue within HTTP.sys and one or more worker processes.

Worker process isolation mode enables you to completely separate an application in its own process, with no dependence on a central process such as Inetinfo.exe to load and execute the application. All requests are handled by worker processes that are isolated from the Web server itself. Process boundaries separate each application pool so that when an application is routed to one application pool, applications in other application pools do not affect that application. By using application pools, you can run all application code in an isolated environment without incurring a performance penalty. For more information about application pools, see How Application Pools Work.

For a visual representation of worker process isolation mode architecture, see Figure 2.1.




Worker process isolation mode delivers all the benefits of the new IIS 6.0 architecture, including multiple application pools, health monitoring and recycling, increased security and performance, improved scalability, and processor affinity. For example, the new health monitoring features can help you discover and prevent application failures, and can also help protect your Web server from imperfect applications.

*IIS 5.0 isolation mode* provides compatibility for applications that were designed to run in earlier versions of IIS. When IIS 6.0 is running in IIS 5.0 isolation mode, request processing is almost identical to the request processing in IIS 5.0. When a server is working in IIS 5.0 isolation mode, application pools, recycling, and health monitoring features are unavailable.

For a visual representation of IIS 5.0 isolation mode architecture, see Figure 2.2. The dashed line in Figure 2.2 indicates the dependency of the worker process on the WWW service, which manages the worker process. For more information about the role of the WWW service, see WWW Service Administration and Monitoring.




Use IIS 5.0 isolation mode only if components or applications do not function in worker process isolation mode. The latter mode is designed to provide an environment in which most existing applications or sites function correctly.

For more information about IIS 6.0 isolation modes, see Application Isolation Modes.

**IIS 6.0 Services**

IIS 6.0 provides four Internet services: the World Wide Web Publishing Service (WWW service) for hosting Internet and intranet content; the File Transfer Protocol (FTP) service for hosting sites where users can upload and download files; the Network News Transfer Protocol (NNTP) service for hosting discussion groups; and the Simple Mail Transfer Protocol (SMTP) service for sending and receiving e-mail messages.

After installing these services, you can create sites or virtual servers, configure properties and security settings, and set up components to further customize your system. For more information about IIS services, see IIS 6.0 Services later in this section and Configuring Internet Sites and Services.

WWW Service Administration and Monitoring, which is a new set of features that were added to the WWW service in IIS 6.0, manages worker processes, supports the new request processing model, and is responsible for health management and maintenance, including application pool health monitoring, recycling worker processes, and rapid-fail protection.

IIS Admin service is a service of the Microsoft® Windows® Server 2003, Standard Edition; Microsoft® Windows® Server 2003, Enterprise Edition; Microsoft® Windows® Server 2003, Web Edition; and Microsoft® Windows® Server 2003, Datacenter Edition operating systems. The IIS Admin service manages the IIS metabase, which stores IIS configuration data. The IIS Admin service makes metabase data available to applications and the core components of IIS. For more information about the IIS metabase, see The IIS Metabase.

**Comparing Basic Functionality Between IIS 6.0 and Earlier IIS Versions**

For a quick preview of IIS architecture, comparing IIS 6.0 architecture with the architecture for earlier versions of IIS, see Table 2.1, which summarizes the basic functionality for different versions of IIS.

**Table 2.1 Basic Functionality in Different Versions of IIS**

| Functionality | IIS 4.0 | IIS 5.0 | IIS 5.1 | IIS 6.0 |
|---|---|---|---|---|
| Platform | Microsoft® Windows NT® 4.0 operating system | Microsoft® Windows® 2000 Server operating system | Microsoft® Windows® XP Professional operating system | Windows Server 2003 |
| Architecture | 32-bit | 32-bit | 32-bit and 64-bit | 32-bit and 64-bit |
| Network subsystem | TCP/IP kernel | TCP/IP kernel | TCP/IP kernel | HTTP.sys kernel |
| Application request processing model | **MTX.exe:** Multiple DLL hosts in High application isolation.<br>**Inetinfo.exe**: In-process as DLLs with Low isolation. | **DLLHost.exe:** Multiple DLL hosts in Medium or High application isolation.<br>**Inetinfo.exe:** In-process as DLLs with Low isolation. | **DLLHost.exe:** Multiple DLL hosts in Medium or High application isolation.<br>**Inetinfo.exe:** In-process as DLLs with Low isolation. | **W3wp.exe:** In worker process isolation mode (multiple worker processes).<br>**Inetinfo.exe:** Only in IIS 5.0 isolation mode with *in-process* applications.<br>**DLLHost.exe**: Only in IIS 5.0 isolation mode with *out-of-process* applications. |
| Metabase configuration | Binary | Binary | Binary | XML |
| Security | Windows authentication<br>Secure Sockets Layer (SSL) | Windows authentication<br>SSL<br>Kerberos | Windows authentication<br>SSL<br>Kerberos | Windows authentication<br>SSL<br>Kerberos<br>.NET Passport support |
| Remote administration | HTMLA[1] | HTMLA | No HTMLA<br>Terminal Services | Remote Administration Tool (HTML)<br>Terminal Services |

[1] Hypertext Markup Language Administration or HTML-based administration. HTMLA-based administration allows

administrators to manage IIS Web servers by using a standard Web browser, such as Microsoft® Internet Explorer, as the client-side interface.

# IIS 6.0 Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 offers four Internet services that you can use to create sites or virtual servers, configure properties and security settings, and set up components to customize your system.

When you install IIS 6.0 on a computer that does not contain an earlier version of IIS, IIS 6.0 automatically installs the following two services:

- The WWW service, which hosts Internet and intranet content.
- The IIS Admin service, which manages the IIS metabase.

You can also choose to install one or more of the following services:

- The FTP service for hosting sites from which users can upload and download files.
- The NNTP service for hosting discussion groups.
- The SMTP service for sending and receiving e-mail messages.

Table 2.2 lists the IIS services, their core components, and their service hosts.

**Table 2.2 Basic Services Provided by IIS 6.0**

| Service Name | Description | Service Short Name | Core Component | Host |
|---|---|---|---|---|
| World Wide Web Publishing Service (WWW service) | Delivers Web publishing services. | W3SVC | Iisw3adm.dll | Svchost.exe |
| File Transfer Protocol (FTP) | Allows file uploads and downloads from remote systems. | MSFTPSVC | Ftpsvc2.dll | Inetinfo.exe |
| Simple Mail Transfer Protocol (SMTP) | Sends and receives electronic messages (e-mail). | SMTPSVC | Smtpsvc.dll | Inetinfo.exe |
| Network News Transfer Protocol (NNTP) | Distributes network news messages. | NNTPSVC | Nntpsvc.dll | Inetinfo.exe |
| IIS Admin Service | Manages the metabase. | IISADMIN | Iisadmin.dll | Inetinfo.exe |

**World Wide Web Publishing Service**

The World Wide Web Publishing Service (WWW service) provides Web publishing for IIS, connecting client HTTP requests to Web sites running on an IIS-based Web server.

The WWW service manages and configures the IIS core components that process HTTP requests. These core components include the HTTP protocol stack (HTTP.sys) and the worker processes.

The WWW service includes these subcomponents: Active Server Pages (ASP), Internet Data Connector, Remote Administration (HTML), Remote Desktop Web Connection, server-side includes (SSI), Web Distributed Authoring and Versioning (WebDAV) publishing, and ASP.NET.

**FTP Service**

IIS provides an FTP service, which you can use to allow users on remote computer systems to copy files to and from your server on a network that uses TCP/IP. The IIS FTP service is an implementation of the File Transfer Protocol, RFC 959, File Transfer Protocol (FTP), and several extensions. The FTP protocol is implemented on top of TCP, which ensures that file transfers are complete and that data transfer is accurate.

You can deploy the FTP service across an arbitrary number of front-end and back-end servers, which increases reliability and availability. By adding virtual directories and servers, you can easily scale FTP without affecting end users.

In IIS 6.0, the FTP service allows you to isolate users at the site level, a feature known as *FTP user isolation*, to help administrators secure and commercialize their Internet sites. Because of the easy availability and wide adoption of FTP, Internet service providers (ISPs) and application service providers (ASPs) traditionally have used FTP to upload their Web content. IIS 6.0 allows the isolation of users into their own directory, thus preventing users from viewing or overwriting other users Web content. The user's top-level directory appears as the root of the FTP site, thus restricting access by not allowing users to navigate farther up the directory tree or across to other users home directories. Within the users own site, the user can create, modify, or delete files and folders.

**SMTP Service**

The SMTP service in IIS processes messages by using the Simple Mail Transfer Protocol (SMTP), which is a TCP/IP protocol that is used to send and receive messages from one computer to another on a network. This protocol is used in intranets and on the Internet to route e-mail.

SMTP is the Internet standard for transporting and delivering electronic messages. Based on specifications in RFCs 2821 and 2822, Microsoft SMTP service is included in the Windows Server 2003 operating system. In Windows Server 2003, the SMTP service is actually a component of IIS and runs as part of Inetinfo.exe. Windows Server 2003 uses SMTP as its native transport protocol to route all messages internally and externally. SMTP is also the default transport for Microsoft® Exchange 2000 Server.

The SMTP component of IIS can send or receive SMTP e-mail messages. You can program the server to automatically send messages in response to events — for example, to confirm a successful form submission by a user. You can also use SMTP to receive messages — for example, to collect feedback from Web site customers.

SMTP does not provide a complete e-mail service. To obtain complete e-mail services for your users, use Microsoft® Exchange Server. For more information about Exchange Server, see the Microsoft Exchange Server TechCenter.

**NNTP Service**

IIS provides a Network News Transfer Protocol (NNTP) service, which you can use to distribute network news messages to NNTP servers and to NNTP clients (news readers) on the Internet. NNTP provides for the distribution, inquiry, retrieval, and posting of news articles by using a reliable stream-based transmission of news on the Internet. With NNTP, news articles are stored on a server in a central database from which users select specific items to read. Indexing, cross-referencing, and expiration of aged messages also are provided.

You can host NNTP local discussion groups on a single computer. Because this feature complies fully with the NNTP protocol, users can use any NNTP client to participate in the newsgroup discussions. The IIS NNTP service does not support news feeds or replication. To use news feeds or to replicate a newsgroup across multiple computers, use Microsoft Exchange Server.

| **Note** |
| --- |
| For more information about these four IIS Internet services, see Configuring Internet Sites and Services. |

**IIS Admin Service**

IIS Admin service is a Windows Server 2003 service that manages the IIS metabase. The metabase stores IIS configuration data in a plaintext XML file that you can read and edit by using common text editors. IIS Admin service makes metabase data available to other applications, including the core components of IIS, applications built on IIS, and applications that are independent of IIS, such as management or monitoring tools. For more information about the IIS metabase, see The IIS Metabase and Working with the Metabase.

# IIS 6.0 Core Components

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 contains several core components that perform important functions in the new IIS architecture. The more that you know about these components, the better you can use IIS 6.0 to your advantage and effectively troubleshoot any issues that might arise.

Most of these core components were briefly described in Overview of IIS 6.0 Architecture. This section provides additional details, including how these core components interact with each other or how they function in each of the two application isolation modes. For example, you can read about the following topics in this section:

- **HTTP Protocol Stack (HTTP.sys).** How the HTTP protocol stack protects the operating system kernel from the effects of imperfect application code, and how it handles kernel-mode queuing.
- **Worker processes.** How the worker process is controlled by the WWW service in both application isolation modes, but in different ways.
- **WWW Service Administration and Monitoring.** How the WWW Service Administration and Monitoring component manages the lifetime of the worker process.
- **Inetinfo.exe.** How the role of Inetinfo.exe has changed since IIS 5.0.
- **IIS Metabase.** Benefits of the new XML metabase, especially in diagnosing possible metabase corruption or in reading and editing the metabase configuration.

# HTTP Protocol Stack

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

The HTTP listener is implemented as a kernel-mode device driver called the HTTP protocol stack (HTTP.sys). IIS 6.0 uses HTTP.sys, which is part of the networking subsystem of the Windows operating system, as a core component.

Earlier versions of IIS use Windows Sockets API (Winsock), which is a user-mode component, to receive HTTP requests. By using HTTP.sys to process requests, IIS 6.0 delivers the following performance enhancements:

- Kernel-mode caching. Requests for cached responses are served without switching to user mode.
- Kernel-mode request queuing. Requests cause less overhead in context switching, because the kernel forwards requests directly to the correct worker process. If no worker process is available to accept a request, the kernel-mode request queue holds the request until a worker process picks it up.

Using HTTP.sys and the new WWW service architecture provides the following benefits:

- When a worker process fails, service is not interrupted; the failure is undetectable by the user because the kernel queues the requests while the WWW service starts a new worker process for that application pool.
- Requests are processed faster because they are routed directly from the kernel to the appropriate user-mode worker process instead of being routed between two user-mode processes.

For information about improvements made to HTTP.sys for Windows Server 2003 SP1, see the HTTP API in Windows Server 2003 SP1 white paper.

**How HTTP.sys Works**

When you create a Web site, IIS registers the site with HTTP.sys, which then receives any HTTP requests for the site. HTTP.sys functions like a forwarder, sending the Web requests it receives to the request queue for the user-mode process that runs the Web site or Web application. HTTP.sys also sends responses back to the client.

Other than retrieving a stored response from its internal cache, HTTP.sys does not process the requests that it receives. Therefore, no application-specific code is ever loaded into kernel mode. As a result, bugs in application-specific code cannot affect the kernel or lead to system failures.

HTTP.sys provides the following services in IIS 6.0:

- Routing HTTP requests to the correct request queue.
- Caching of responses in kernel mode.
- Performing all text-based logging for the WWW service.
- Implementing Quality of Service (QoS) functionality, which includes connection limits, connection timeouts, queue-length limits, and bandwidth throttling.

**How HTTP.sys Handles Kernel-Mode Queuing**

When IIS 6.0 runs in worker process isolation mode, HTTP.sys listens for requests and queues those requests in the appropriate queue. Each request queue corresponds to one application pool. An *application pool* corresponds to one request queue within HTTP.sys and one or more worker processes.

When IIS 6.0 runs in IIS 5.0 isolation mode, HTTP.sys runs like it runs in worker process isolation mode, except that it routes requests to a single request queue.

If a defective application causes the user-mode worker process to terminate unexpectedly, HTTP.sys continues to accept and queue requests, provided that the WWW service is still running, queues are still available, and space remains in the queues.

When the WWW service identifies an unhealthy worker process, it starts a new worker process if outstanding requests are waiting to be serviced. Thus, although a temporary disruption occurs in user-mode request processing, an end user does not experience the failure because TCP/IP connections are maintained, and requests continue to be queued and processed.

# Worker Processes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A *worker process* is user-mode code whose role is to process requests, such as processing requests to return a static page, invoking an ISAPI extension or filter, or running a Common Gateway Interface (CGI) handler.

In both application isolation modes, the worker process is controlled by the WWW service. However, in worker process isolation mode, a worker process runs as an executable file named W3wp.exe, and in IIS 5.0 isolation mode, a worker process is hosted by Inetinfo.exe. Figure 2.2, which depicts the architecture for IIS 5.0 isolation mode, uses a dashed line to suggest the relationship between the worker process and the WWW service.

Worker processes use HTTP.sys to receive requests and to send responses by using HTTP. Worker processes also run application code, such as ASP.NET applications and XML Web services. You can configure IIS to run multiple worker processes that serve different application pools concurrently. This design separates applications by process boundaries and helps achieve maximum Web server reliability.

By default, worker processes in worker process isolation mode run under the Network Service account, which has the strongest security (least access) compatible with the functionality that is required. For more information about Network Service accounts, including Windows Server 2003 built-in security accounts, see Running IIS 6.0 as an Application Server, and IIS and Built-in Accounts.

# WWW Service Administration and Monitoring

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The WWW Service Administration and Monitoring component, which is hosted in Svchost.exe, is a set of new features contained in the WWW service. Like the worker process, this new component runs in user mode, and like HTTP.sys, this component contains critical IIS 6.0 functionality that never loads external code. IIS 6.0 completely isolates external application code from this component by running only the external code in worker processes.

| **Note** |
| --- |
| Because the WWW Service Administration and Monitoring component implements and controls a set of new features contained in the WWW service, to simplify the text, the term *WWW service* is used when discussing these new features instead of the term *WWW Service Administration and Monitoring component*. |

The WWW service is responsible for two roles in IIS: HTTP administration and worker process management.

**HTTP administration**
The WWW service interacts with the metabase to obtain the configuration data that is needed to set up HTTP.sys and that is used to manage the worker process.

When you start a server that is running IIS, and the WWW service is started, this service reads metabase information and initializes the HTTP.sys namespace routing table with one entry for each application. HTTP.sys then uses the routing table data to determine which application pool responds to requests from what parts of the namespace. When HTTP.sys receives a request, it signals the WWW service to start a worker process for the application pool.

When you add or delete an application pool, the WWW service processes the configuration changes, which includes adding or deleting the application pool queue from HTTP.sys. The WWW service also listens for and processes configuration changes that occur as a result of application pool recycling.

**Managing the lifetime of the worker process**
The WWW service is responsible for managing the worker processes, which includes starting the worker processes and maintaining information about the running worker processes. This service also determines when to start a worker process, when to recycle a worker process, and when to restart a worker process if it becomes blocked and is unable to process any more requests.

In worker process isolation mode, the WWW service is responsible for health management and maintenance, including application pool health monitoring, recycling worker processes, and rapid-fail protection. For more information about these new features, see Features of Worker Process Isolation Mode, and Running IIS 6.0 as an Application Server.

The WWW service runs under the Local System account. For more information about Local System accounts, including Windows Server 2003 built-in security accounts, see Running IIS 6.0 as an Application Server and Managing a Secure IIS 6.0 Solution.

# Inetinfo.exe

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When IIS 6.0 runs in worker process isolation mode, Inetinfo.exe is a user-mode component that hosts the IIS metabase and that also hosts the non-Web services of IIS 6.0, including the FTP service, the SMTP service, and the NNTP service. Inetinfo.exe depends on IIS Admin service to host the metabase.

When IIS 6.0 runs in IIS 5.0 isolation mode, Inetinfo.exe functions much as it did in IIS 5.0. In IIS 5.0 isolation mode, however, Inetinfo.exe hosts the worker process, which runs ISAPI filters, Low-isolation ISAPI extensions, and other Web applications.

**Important**

In IIS6.0, regardless of the application isolation mode used, the services that run in Inetinfo.exe run as dynamic-link libraries (DLLs) under the Local System account. Because a Local System account allows users access to every resource on the local computer, use caution when you run services in Inetinfo.exe, especially on computers that provide services on the Internet.

**Note**

When you run services in Inetinfo.exe that can be accessed from the Internet, be sure to protect your system by using a firewall, such as Microsoft® Internet Security and Acceleration (ISA) Server, or by disabling services that you do not use.

# The IIS Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS *metabase* is a plaintext, XML data store that contains most IIS configuration information. Although most of the IIS configuration settings are stored in the IIS metabase, a few settings are maintained in the Windows registry. If a configuration setting is one that you might need to configure or change, or if you can access the setting in the IIS 6.0 user interface, then the setting is typically stored in the IIS metabase.

The new XML metabase allows administrators to directly read and edit the configuration file without using scripts or code, by using a common text editor such as Notepad. The XML metabase makes it easier to diagnose potential metabase corruption, and to back up and restore your metabase files.

If you upgrade to IIS 6.0, the existing binary metabase is automatically converted to the new XML metabase. The metabase is 100 percent compatible with IIS 5.*x* metabase application programming interfaces (APIs) and Active Directory® directory service Service Interfaces (ADSI), ensuring that existing scripts and code continue to work.

For more information about the metabase, see Working with the Metabase.

**Metabase Elements**

IIS 6.0 replaces the single binary file (MetaBase.bin) of earlier IIS versions with two XML files: MetaBase.xml and MBSchema.xml. IIS stores these files in the *systemroot*\System32\Inetsrv folder of your computer. To view and modify these files, you must be a member of the Administrators group.

The metabase consists of the following elements:

- **MetaBase.xml file.** This file stores IIS configuration information that is specific to an installation of IIS.
- **MBSchema.xml file.** This file contains the metabase schema. The MBSchema.xmlfileis a master configuration file that defines default attributes for all metabase properties and enforces rules for constructing and placing metabase entries within the metabase.
- **In-memory metabase.** The in-memory metabase contains the most current metabase and metabase schema configuration. The in-memory metabase accepts changes to the metabase configuration and schema, storing them in RAM, and periodically writing changes to the on-disk metabase and metabase schema files.

When IIS starts, the MetaBase.xml and MBSchema.xml files are read by the IIS storage layer and copied to the in-memory metabase. While IIS is running, any changes that you make to the in-memory metabase are periodically written to disk. IIS also saves the in-memory metabase to disk when you stop IIS.

The IIS Admin service makes the metabase available (by means of the Admin Base Object) to other applications, including the core components of IIS, applications built on IIS, and applications that are independent of IIS, such as management or monitoring tools.

**Automatic Versioning and History**

IIS 6.0 also provides automatic versioning and history. This metabase history feature automatically tracks any changes to the metabase that are written to disk. When the metabase is written to disk, IIS marks the new MetaBase.xml file with a version number and saves a copy of the file in the history folder. Each history file is marked with a unique version number, which is then available, if needed, for metabase rollback or restore. The metabase history feature is enabled by default.

For more information about the metabase, see Working with the Metabase.

# Application Isolation Modes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 runs a server in one of two distinct request processing models, called *application isolation modes*. In each isolation mode, IIS functions differently, although both application isolation modes rely on HTTP.sys as the HTTP listener.

**Worker process isolation mode**
Worker process isolation mode takes advantage of the redesigned architecture for IIS 6.0. In this application isolation mode, you can group Web applications into application pools, through which you can apply specific configuration settings to groups of applications and to the worker processes servicing those applications. By using application pools, you can run all application code in an isolated environment; however, unlike earlier versions of IIS, IIS 6.0 provides isolation without a performance penalty, because there are no added process hops.

Worker process isolation mode provides application (or site) compatibility for most existing applications. This application isolation mode is specified by a **false** value for the metabase property **IIs5IsolationModeEnabled**.

**IIS 5.0 isolation mode**
IIS 5.0 isolation mode is provided for applications that depend upon specific features and behaviors of IIS 5.0. Use this mode only if an application has a compatibility issue when it runs in worker process isolation mode and you cannot resolve the problem.

In this application isolation mode, you can isolate applications as you did in IIS 5.0: in Low isolation (in-process), Medium isolation (pooled out-of-process), or High isolation (out-of-process). This application isolation mode is specified by a **true** value for the **IIs5IsolationModeEnabled** metabase property.

Whenever possible, run your server in worker process isolation mode, which is more secure and reliable than IIS 5.0 isolation mode and is also the default application isolation mode for a clean installation of IIS 6.0. For more information about IIS 6.0 application isolation mode defaults, see Application Isolation Mode Defaults.

# Evaluating Application Isolation Modes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS application isolation mode that you choose affects performance, reliability, security, and feature availability. Worker process isolation mode is the recommended mode of operation for IIS 6.0 because it provides a more robust platform for applications and the ability to isolate applications in separate application pools without incurring a performance penalty. Also, in the default configuration, worker process isolation mode provides a higher level of security than does IIS 5.0 isolation mode because its default host process identity is Network Service.

Some applications might have compatibility problems when they run in worker process isolation mode. If compatibility problems occur, you might need to use IIS 5.0 isolation mode. For applications that run in IIS 5.0 isolation mode, Local System is the only available process identity when you run an application in-process in Inetinfo.exe. For more information about IIS security, see Managing a Secure IIS 6.0 Solution. Also see Configuring Worker Process Identities and IIS and Built-in Accounts.

**Important**

IIS6.0 cannot run both application isolation modes simultaneously on the same computer. Therefore, on a single computer running IIS6.0, you cannot run some Web applications in worker process isolation mode and others in IIS5.0 isolation mode. If your applications require different modes, you must run them on separate computers.

When you are deciding which application isolation mode to use, follow these guidelines:

- Use worker process isolation mode unless you need to use IIS 5.0 isolation mode because of a specific compatibility problem. For example, an application that checks to see if an application is running in a process named Inetinfo.exe is incompatible with worker process isolation mode.
- For Web sites that contain static content and simple ASP applications, run the applications in worker process isolation mode because these Web sites usually require little or no modification.
- For applications that run correctly on IIS 5.0, try running them first in worker process isolation mode. If they do not work correctly, change them back to IIS 5.0 isolation mode until you can diagnose and fix the incompatibility.

**Applications that cannot run in worker process isolation mode**

Applications that run in worker process isolation mode can take advantage of all the worker process behaviors, which include multi-instancing, recycling, and configurable process identity. Certain application characteristics conflict with the architecture of worker process isolation mode and must be run in IIS 5.0 isolation mode until you can modify them.

The following applications or application characteristics conflict with worker process isolation mode:

- **Read raw data filters.** This type of filter can be loaded only in IIS 5.0 isolation mode.
- **Dependency on Inetinfo.exe.** Applications that must run in an Inetinfo.exe worker process can run only in IIS 5.0 isolation mode because applications cannot run in Inetinfo.exe in worker process isolation mode.
- **Requires Dllhost.exe.** Web applications that must run in a Dllhost.exe environment can run only in IIS 5.0 isolation mode because running Web applications in Dllhost.exe is not an option in worker process isolation mode.

**Applications that use SF_NOTIFY_READ_RAW_DATA filters**

ISAPI filters that use SF_NOTIFY_READ_RAW_DATA notifications work differently depending on the isolation mode in which IIS 6.0 is running. To understand the differences, remember that a global filter is installed for the entire IIS service and sees requests for all Web sites, whereas a site filter is installed at the site level and sees requests only for the site where it is installed.

**Note**

The only raw data filter that works differently in IIS 6.0 is SF_NOTIFY_READ_RAW_DATA. SF_NOTIFY_SEND_RAW_DATA works as expected and is supported in both application isolation modes.

For IIS 6.0 running in IIS 5.0 isolation mode, the behavior is the same as it was in IIS 5.0. Because IIS must read the host header information to route the request to the correct site, and because the READ_RAW_DATA notification is sent before the host header arrives from the client, it is impossible to know which site must be notified. For this reason, site fiters cannot register for SF_NOTIFY_READ_RAW_DATA notifications. Global filters, however, can register for SF_NOTIFY_READ_RAW_DATA notifications.

IIS 6.0 running in worker process isolation mode does not support SF_NOTIFY_READ_RAW_DATA notifications. Worker process isolation mode is based on application pools, which do not correspond directly to sites. Http.sys routes a request based on server bindings and the URL namespace. Because a READ_RAW_DATA notification must be sent before IIS has received enough information from the client request to determine the server bindings and URL, IIS cannot route a READ_RAW_DATA notification to the correct worker process. This is not an issue in IIS 6.0 running in IIS 5.0 isolation mode, because Inetinfo.exe is always the only worker process.

For information about how the SF_NOTIFY_READ_RAW_DATA notification works with Secure Sockets Layer (SSL), see SSL and SF_NOTIFY_READ_RAW_DATA filters in Managing a Secure IIS 6.0 Solution.

**Applications that require special handling to run in worker process isolation mode**

Applications that you cannot load into multiple processes simultaneously have the following limitations:

- You cannot run them in application pools configured to perform overlapped recycling.
- You cannot run them in application pools configured as Web gardens.

Instead, run this type of application in an application pool that is not configured for overlapped recycling and that is not configured to run as a Web garden. For more information about overlapped recycling, see Features of Worker Process Isolation Mode. For information about configuring overlapped recycling, see Recycling Worker Processes with IIS 6.0.

If the IIS 6.0 WWW service is running in worker process isolation mode (the IIS 6.0 default mode), and you must run applications that require a functionality that is available only in IIS 5.0 isolation mode, change to IIS 5.0 isolation mode until you can correct the incompatibility. When running your server in IIS 5.0 isolation mode, you cannot take advantage of worker process isolation and the other features of worker process isolation mode.

# Application Isolation Mode Defaults

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you install IIS 6.0 on a computer that does not contain an earlier version of IIS, IIS 6.0 automatically sets the application isolation mode to worker process isolation mode. If you upgrade from an earlier version of IIS, IIS 6.0 sets the application isolation mode to IIS 5.0 isolation mode.

Table 2.3 lists the IIS 6.0 default settings for the two application isolation modes.

**Table 2.3 Default IIS 6.0 Installation Settings for Each Application Isolation Mode**

| Type of Installation | Application Isolation Mode |
|---|---|
| New installation of IIS 6.0 (clean install) | Worker process isolation mode |
| Upgrade from an earlier version of IIS 6.0, such as a beta release | No change in isolation mode |
| Upgrade from IIS 5.0 | IIS 5.0 isolation mode |
| Upgrade from IIS 4.0 | IIS 5.0 isolation mode |

Although IIS 5.0 isolation mode is the default application isolation mode when you upgrade from earlier versions of IIS, be sure to test your applications for compatibility issues so that, if possible, you can upgrade to worker process isolation mode.

# Application Isolation Mode Functions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use Table 2.4 to compare how certain IIS 6.0 functions run in worker process isolation mode with how they run in IIS 5.0 isolation mode. Knowing the host component for each application isolation mode can help you decide which application isolation mode to use. For example, if an application must run in DLLhost.exe, and you cannot change the application to remove this requirement, this table shows you the two available options, which are to run the application in Medium or High isolation in IIS 5.0 isolation mode.

**Subhead**

**Table 2.4 Comparison of IIS 6.0 Functions in the Two Application Isolation Modes**

| IIS 6.0 Function | Worker Process Isolation Mode Host/Component | IIS 5.0 Isolation Mode Host/Component |
|---|---|---|
| Worker process management | Svchost.exe/WWW service | Svchost.exe/WWW service |
| Worker process | W3wp.exe | Inetinfo.exe |
| Running Low-isolation ISAPI applications | W3wp.exe | Inetinfo.exe |
| Running Medium-isolation ISAPI applications | Not applicable (All ISAPI extensions are in-process.) | DLLHost.exe (1) |
| Running High-isolation ISAPI applications | Not applicable (All ISAPI extensions are in-process.) | DLLHost.exe |
| Running ISAPI filters | W3wp.exe | Inetinfo.exe |
| HTTP.sys configuration | Svchost.exe/WWW service | Svchost.exe/WWW service |
| HTTP protocol support | Windows kernel/HTTP.sys | Windows kernel/HTTP.sys |
| HTTP SSL (HTTPFilter) | Lsass.exe | Inetinfo.exe |
| IIS metabase | Inetinfo.exe | Inetinfo.exe |
| FTP | Inetinfo.exe | Inetinfo.exe |
| NNTP | Inetinfo.exe | Inetinfo.exe |
| SMTP | Inetinfo.exe | Inetinfo.exe |

**HTTP SSL (HTTPFilter)**

When HTTP.sys receives requests that are encrypted by using Secure Sockets Layer (SSL), the kernel-mode HTTP service cannot decrypt the requests or encrypt the responses. Instead, HTTP SSL (HTTPFilter), which is a user-mode service, implements SSL for the HTTP service. HTTP SSL runs in the IIS 6.0 isolation modes as follows:

- In worker process isolation mode, Lsass.exe hosts HTTPFilter.
- In IIS 5.0 isolation mode, HTTP filter runs in Inetinfo.exe.

| **Important** |
|---|
| Lsass.exe cannot be stopped except during computer shutdowns — either planned or unexpected. |

**IIS 6.0 Identity vs. IIS 5.0 Identity**

In IIS 6.0, worker processes use a different process identity than in IIS 5.0. *Process identity* is an operating system term used to

denote the account that a process runs under. Every process that is running on a Windows NT operating system has a process identity that is used to control access to resources on the system.

The different process identity used by worker processes in IIS 6.0 can affect applications migrated from IIS 5.0 if the application expects the process identity to run as a specific account, such as IWAM_*ComputerName*. In addition, for Medium-isolation or High-isolation Web applications created in IIS 5.0 or run on IIS 6.0 in IIS 5.0 isolation mode, identity is configured by Component Services; however, when IIS 6.0 is running in worker process isolation mode, identity is configured in IIS Manager.

# Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 worker process isolation mode takes the concept of application isolation, which was introduced in IIS 4.0, one step further. In IIS 6.0, you can isolate one application from another so that an error in one application does not affect another application running in a different process. This application isolation mode provides better isolation while not incurring a performance penalty for isolation.

Worker process isolation mode loads application code — for example, ASP and ASP.NET applications — into the worker process only. By isolating application code in the worker process, this application isolation mode ensures a reliable environment for an application server: the WWW service, IIS Admin service, and HTTP.sys can run continuously despite any service interruptions that occur in a worker process. Also, Web sites running in the worker processes are not affected by failures in other worker processes because they are isolated from each other through operating system process boundaries.

Worker process isolation mode uses all of the new IIS 6.0 core components and supports application pooling, recycling, and health monitoring features.

# Benefits of Using Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

One benefit of isolating applications or sites into separate worker processes is to simplify management tasks, such as taking an application offline or putting an application back online (independent of all other sites or applications running on the system), changing a component that an application uses, debugging an application, and throttling (limiting) resources that an application uses.

IIS 6.0 worker process isolation mode improves on earlier versions of IIS in the following areas:

- **It is more robust.** Prevents Web applications and Web sites in an application pool from disrupting or harming applications or Web sites in another application pool.

- **It requires fewer or no reboots.** Does not force you to reboot the server or to shut down the WWW service. You can perform common tasks, such as upgrading site content or components, debugging Web applications, or dealing with faulty Web applications, without affecting service to other sites or applications on the server.

- **It is self-healing.** Supports autorestart of failed applications and periodic restart of unstable applications or applications that leak memory.

- **It is more scalable and reliable.** Supports scaling for ISPs, which sometimes run thousands of sites on a server, and offers better multiprocessor scalability.

- **It is application-oriented.** Makes the application the unit of administration and robustness by enabling application isolation and also by enabling resource throttling (limiting the amount of resources that an application can use) and scaling based on the application.

The result is a Web server that is more reliable and always available, even if applications cause their hosting worker processes to terminate.

# Features of Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Worker process isolation mode includes several features that you can use to improve IIS robustness, scalability, and reliability without reducing performance. These features include multiple application pools, application pool health monitoring, worker process recycling, and rapid-fail protection. When you run your Web server in worker process isolation mode, you can configure application pools to take advantage of these features, which can only be used when a Web server is running in worker process isolation mode.

**Multiple Application Pools**
In IIS 5.0, applications can be pooled together out-of-process, but only in one application pool (typically referred to as *a pooled out-of-process* in IIS 5.0, rather than an *application pool*). In IIS 6.0, worker process isolation mode allows customers to create multiple application pools, each with a different configuration, such as a unique recycling configuration or a unique process ID. In addition, all HTTP application run-time features, such as the services that support ISAPI extensions, are available in all application pools. This design prevents a malfunctioning Web application or Web site from disrupting other Web applications or Web sites on that server.

**Application Pool Health Monitoring**
You can configure an application pool to monitor not only the health of its worker process, but the health of the entire pool. The WWW service can detect an unhealthy application pool in these situations:

- If a worker process in the application pool terminates abnormally.
- If all of the available IIS threads in the worker process that is assigned to the application pool are blocked.
- If the application notifies IIS that a problem exists.

Monitoring the health of a worker process includes detecting that the worker process is not able to serve requests and taking corrective action. For example, if a worker process does not respond to a ping from the WWW service, the worker process probably does not have threads available for processing incoming requests. When this happens, the WWW service either terminates the worker process or releases the worker process and leaves it running. (You can configure IIS to attach a debugger to an unhealthy worker process when the WWW server releases it.) If requests are waiting in the request queue, the WWW service starts a new worker process to replace the one that was terminated or released.

For more information about configuring IIS for application pool health monitoring, see Configuring Worker Process Health Monitoring in IIS 6.0 Help and Running IIS 6.0 as an Application Server.

**Worker Process Recycling**
Web applications sometimes leak memory, are poorly coded, or have other unidentified problems. In earlier versions of IIS, these faulty applications might force you to restart the Web server. In worker process isolation mode, you can use the recycling feature to periodically restart the worker processes in an application pool in order to manage faulty applications. You can schedule worker processes to restart based on several options, such as elapsed time or the number of requests served.

You can configure the WWW service so that it recycles a worker process in one of two ways:

- By default, the WWW service establishes an *overlapped recycle*, creating a new worker process to handle any new incoming requests *before* signaling the old worker process to finish processing its existing requests and shut down.
- Optionally, the WWW service first signals the old worker process to shut down and, after it shuts down, starts a new worker process.

In both types of recycling, the WWW service allows a configured amount of time for the old worker process to handle the requests that it has already taken from HTTP.sys before it shuts itself down.

In the default configuration, when a worker process requests a recycle, the WWW service initiates an overlapped recycle, creating a new worker process to replace the old process. While the new worker process is starting, the old process continues to serve requests. As soon as the new worker process starts and initializes successfully, the WWW service signals the old worker process to shut down. The old worker process then stops pulling new requests and begins to shut down. The WWW

service gives the old worker process a configured time period in which to finish processing its requests and shut down, and it terminates the worker process if it does not shut down within the specified time.

For more information about worker process recycling, including when to use it and how to configure it, see Recycling Worker Processes with IIS 6.0.

**Rapid-Fail Protection**

Rapid-fail protection stops an application pool when too many of the worker processes assigned to it are found to be unhealthy within a specified period of time. When an application pool is stopped, HTTP.sys either returns an out-of-service message (503: Service Unavailable) or resets a connection based on the configuration of the **LoadBalancerCapabilities** property of the application pool. Also, when an application pool stops automatically, such as when rapid-fail protection engages, you can configure an action (a debugging action, for example) to notify the administrator that the application pool has stopped.

For more information about rapid-fail protection, including how to configure this feature, see Running IIS 6.0 as an Application Server.

**Orphaning Worker Processes**

You can configure worker process isolation mode to orphan a worker process that the WWW service deems to be failing. The WWW service usually terminates a failing worker process and replaces it. If you enable orphaning, the WWW service allows a failing worker process to continue running, but separates it from the application pool (making it an orphan) and starts a new worker process in its place. You can configure the WWW service to run a command on the orphaned worker process — for example, launching a debugger.

For more information about orphaning worker processes, including how to configure this feature, see Running IIS 6.0 as an Application Server in this book and OrphanWorkerProcess Metabase Property metabase property.

**Web Gardens**

*Web gardens* are application pools that are configured to run multiple worker processes. By default, each application pool in IIS 6.0 has only one worker process. However, you can configure an application pool to have a set of equivalent worker processes that share the work, thus creating a *Web garden*. In a Web garden, HTTP.sys distributes the requests among the set of worker processes in the application pool.

| **Note** |
| --- |
| Do not confuse a Web garden with a *Web farm*, which is a Web site or Web hosting service that runs off of more than one server and that uses clustering and load balancing. |

The benefits of Web gardens are that they can handle more load, and that if one worker process stops responding, other worker processes are available to accept and process requests. Hence, Web gardens can enhance performance and reliability of the applications in the application pool.

The **MaxProcesses** metabase property determines the maximum number of worker processes that an application pool allows to service its requests. For more information about Web gardens, including how to create them, see Running IIS 6.0 as an Application Server, and Configuring Web Gardens with IIS 6.0.

**Processor Affinity**

*Processor affinity* is an application pool property that forces worker processes to run on a specific set of microprocessors or CPUs. Processor affinity applies to all worker processes that serve a particular application pool and allows you to take advantage of more frequent CPU caching. If a Web garden is running on a multiprocessor computer, you can use processor affinity to dedicate certain application pools to specific clusters of CPUs.

To set up processor affinity, configure the **SMPAffinitized** and **SMPProcessorAffinityMask** metabase properties. For more information about assigning processor affinity, including how to configure these metabase properties, see Running IIS 6.0 as an Application Server.

**Idle Timeout**

You can configure an application pool to have its worker processes request a shutdown if they are idle for a certain period of time. Configuring idle timeout can help you to free unused resources. New worker processes are started when demand exists for that application pool.

Use either IIS Manager or the **IdleTimeout** metabase property to configure this feature. For information about setting the **IdleTimeout** metabase property, see IdleTimeout in IIS 6.0 Help. For more information about configuring idle timeout for a worker process, see Running IIS 6.0 as an Application Server.

**Application Pool Identity**

*Application pool identity* is the user account that the worker processes servicing the application pool use as their process identity. *Process identity* is the operating system term that is used to denote the account that a process runs under. Every running process on a Windows NT operating system has a process identity that is used to control access to resources on the system.

You can assign a predefined account for the application pool identity or create your own account for this purpose. For security reasons, choose an application pool identity with the minimum permissions that your application requires.

For more information about configuring process identities, see Configuring Worker Process Identities.

# How Application Pools Work

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you run IIS 6.0 in worker process isolation mode, you can separate different Web applications and Web sites into groups known as *application pools*. An *application pool* is a group of one or more URLs that are served by a worker process or set of worker processes. Any Web directory or virtual directory can be assigned to an application pool.

Every application within an application pool shares the same worker process. Because each worker process operates as a separate instance of the worker process executable, W3wp.exe, the worker process that services one application pool is separated from the worker process that services another. Each separate worker process provides a process boundary so that when an application is assigned to one application pool, problems in other application pools do not affect the application. This ensures that if a worker process fails, it does not affect the applications running in other application pools.

Use multiple application pools when you want to help ensure that applications and Web sites are confidential and secure. For example, an enterprise organization might place its human resources Web site and its finance Web site on the same server, but in different application pools. Likewise, an ISP that hosts Web sites and applications for competing companies might run each company's Web services on the same server, but in different application pools. Using different application pools to isolate applications helps prevent one customer from accessing, changing, or using confidential information from another customer's site.

In HTTP.sys, an application pool is represented by a *request queue*, from which the user-mode worker processes that service an application pool collect the requests. Each pool can manage requests for one or more unique Web applications, which you assign to the application pool based on their URLs. Application pools, then, are essentially worker process configurations that service groups of namespaces.

Multiple application pools can operate at the same time. An application, as defined by its URL, can only be served by one application pool at any time. While one application pool is servicing a request, you cannot route the request to another application pool. However, you can assign applications to another application pool while the server is running.

# Request Processing in Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Figure 2.3 shows the standard request process for IIS 6.0 when it is running in worker process isolation mode.




The HTTP protocol stack (HTTP.sys), which runs in kernel mode, receives client requests and routes them to the appropriate request queue. Worker processes, which run in user mode, pull the requests directly from their own kernel request queues, eliminating the process hops that occur in IIS 5.0 (and that also occur in IIS 5.0 isolation mode) when the Web server sends a request to a High-isolation, out-of-process application. Because these extra process hops are eliminated in worker process isolation mode, IIS can provide application isolation without sacrificing performance.

| **Note** |
| --- |
| Before routing a request to a queue, HTTP.sys checks to see if the response is located in its kernel-mode cache; if a response is cached, it is returned immediately. |

To see a visual representation of the request process for the two application isolation modes, compare Figure 2.3 (worker process isolation mode) with Figure 2.6 (High isolation in IIS 5.0 isolation mode).

| **Note** |
| --- |
| In Figure2.3, and also in Figure2.4, 2.5, 2.6, and 2.7, the numbered arrows indicate the request process flow. |

When HTTP.sys receives requests that are encrypted by using SSL, the kernel-mode HTTP service cannot decrypt the requests or encrypt the responses. Instead, HTTP SSL (HTTPFilter), which is a user-mode service, implements SSL for the HTTP service. In worker process isolation mode, the default host for HTTPFilter is Lsass.exe.

Figure 2.4 illustrates the process flow for an SSL request when IIS 6.0 is running in worker process isolation mode. The HTTP protocol stack receives the SSL request from the client and places it in the SSL queue. The HTTPFilter listens for requests to enter the queue, picks up any new request, uses SSL to decrypt the request, and then returns the decrypted request to HTTP.sys. HTTP.sys then places the decrypted request in the appropriate request queue. A worker process pulls the request directly from its own kernel request queue.




On the return trip, the worker process returns the response to HTTP.sys, which places it in the appropriate queue from which the HTTPFilter picks it up and performs SSL encryption. After encryption, the HTTPFilter again returns the response to HTTP.sys, which sends the encrypted response back to the client.

Keep in mind that SSL requires more resources than does standard HTTP. If you use SSL, ensure that the information being sent, such as credit card numbers or medical information, is sensitive enough to justify the added performance expense. For example, your home page or a search results page does not necessarily need to be run through SSL. However, you need to make sure that a checkout or purchase page is SSL-secured.

When you use SSL, the initial connection takes significantly longer than reconnecting. Reconnections use security information in the SSL session cache. The default timeout for the SSL session cache is changed to 5 minutes in Windows 2000 (10 hours in Windows 2000 with Service Pack 2) and 10 hours in Windows Server 2003. After cached data is deleted, the client and server must establish a completely new connection. You can change the default timeout by configuring the **ServerCacheTime** registry entry. For more information about configuring this registry setting, see Optimizing IIS 6.0 Performance.

| **Caution** |
|---|
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first. |

For more information about balancing performance and security, see Performance and Security in IIS 6.0 Help. For more information about encryption, see About Encryption in IIS 6.0 Help. For more information about the security features available in IIS 6.0, see Managing a Secure IIS 6.0 Solution.

Table 2.5 provides additional details about how worker process isolation mode processes HTTP (standard) and HTTPS (SSL) requests.

**Table 2.5 HTTP and HTTPS Request Processing**

| Step | Request Type | Description |
|---|---|---|

| | | |
|---|---|---|
| I | HTTP & HTTPS | HTTP.sys receives the request from the client and determines the type of request (HTTP or HTTPS). (Corresponds to #1 in Figures 2.3 and 2.4.)<br><br>If the request is an HTTPS request, HTTP.sys places it in the SSL queue. |
| | HTTPS (SSL) only | For an SSL request, the HTTPFilter performs these steps:<br>• Listens for requests to be added to the SSL queue (see #2 in Figure 2.4).<br>• Picks up any requests that are queued.<br>• Uses SSL to decrypt the request.<br>• Returns the decrypted request to HTTP.sys for processing (see #3 in Figure 2.4). |
| II | HTTP & HTTPS | HTTP.sys determines if the request is valid:<br>• If the request is invalid, HTTP.sys sends a 400-class error message for an invalid request back to the client. 1<br>• If the request is valid, HTTP.sys checks to see if the response is located in its kernel-mode cache. |
| III | HTTP & HTTPS | HTTP.sys checks the kernel-mode cache for a response:<br>• If the response is in the kernel-mode cache, HTTP.sys returns the response immediately, without transitioning to user mode.<br>• If the response is not in the cache, HTTP.sys determines the correct request queue and places the request in that queue.<br>• If the target request queue is out of service or if the queue is full, HTTP.sys returns a 503 Service Unavailable message to the client. |
| IV | HTTP & HTTPS | If no worker process is available and listening in the request queue and if HTTP.sys is able to request that a new worker process be started, HTTP.sys signals the WWW service to start a worker process. |
| V | HTTP & HTTPS | The worker process pulls the request from the queue and processes the request. (Corresponds to #2 in Figure 2.3 and #4 in Figure 2.4.) |
| VI | HTTP & HTTPS | The worker process sends the response back to HTTP.sys and instructs HTTP.sys whether to cache the response. (Corresponds to #3 in Figure 2.3 and #5 in Figure 2.4.) |
| VII | HTTP & HTTPS | HTTP.sys determines the type of response (HTTP or HTTPS).<br><br>If the request is an HTTPS response, HTTP.sys places it in the SSL queue. |

| | | |
|---|---|---|
| | HTTPS (SSL) only | For an SSL response, the HTTPFilter performs the following steps:<br>• Listens for responses to enter the SSL queue (see #6 in Figure 2.4).<br>• Picks up any responses that enter the queue.<br>• Uses SSL to encrypt the response.<br>• Returns the encrypted response to HTTP.sys (see #7 in Figure 2.4). |
| VIII | HTTP & HTTPS | HTTP.sys sends the response back to the client and logs the request if logging is enabled. (Corresponds to #4 in Figure 2.3 and #8 in Figure 2.4.) |

For information about 400-class error messages, including keys to the many substatus codes, see Troubleshooting IIS 6.0.

# IIS 5.0 Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 5.0 isolation mode provides as much compatibility as possible for applications that were developed for IIS 5.0. For example, the request processing of IIS 6.0 when it is running in IIS 5.0 isolation mode is nearly identical to the request processing of IIS 5.0. Hence, in IIS 5.0 isolation mode, you can run read raw data filters and can run applications in Inetinfo.exe or in DLLHost.exe — all of which is not accommodated in worker process isolation mode.

Before you decide to run an application in IIS 5.0 isolation mode, carefully evaluate whether you can substitute other filters and applications for the ones that require IIS 5.0 architecture because when a server is running in IIS 5.0 isolation mode, application pools, recycling, and health monitoring features are unavailable.

For more information about evaluating application isolation modes, see Evaluating Application Isolation Modes.

# Application Isolation in IIS 5.0 Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When IIS 6.0 is running in IIS 5.0 isolation mode, your options for isolation, which are similar to your options when you run an application on IIS 5.0, are as follows:

- **Low isolation (in-process).** Applications run in-process (on the Web server) as DLLs in Inetinfo.exe and are not protected from other applications running in-process. The process identity of the application (the account under which the applications process runs) is Local System.

- **Medium isolation (pooled out-of-process).** All pooled applications run as DLLs in one instance of DLLHost.exe and are protected from the effects of applications running in High isolation. However, because all pooled applications run in the same process, they are not protected from each other. The default application identity is IWAM_*ComputerName*.

- **High isolation (out-of-process).** Applications run as DLLs in DLLHost.exe and are protected from the effects of other applications running in Medium or High isolation. In addition, all other applications that run on the same computer are protected from the effects of applications that run in High isolation. The default application identity is IWAM_*ComputerName*.

**Security Considerations**

When IIS 6.0 runs in IIS 5.0, Web applications that are set to Low isolation run with the Local System account. The Local System account has access to all of the resources on the computer. This means that when an attack by a malicious user takes over a Web application that is set to Low isolation, all of the assets of the local computer are open to the attacker.

When you set Web applications to Medium or High isolation, they run with IWAM_*ComputerName* as the default account; however, you can configure these applications to run with other identities.

**Performance Considerations**

When you run IIS 6.0 in IIS 5.0 isolation mode, applications that run in Medium or High isolation incur a performance penalty. The penalty occurs because the Component Object Model (COM) uses a remote procedure call (RPC) during these parts of the cycle:

- When a request is sent to an application that is running in Medium or High isolation and the response is sent back.

- When a Web application that is running in Medium or High isolation communicates with the Web server during the processing of a request.

In contrast, when IIS 6.0 runs in worker process isolation mode, applications are loaded in-process to W3wp.exe. Therefore, no performance penalty occurs.

# Request Processing in IIS 5.0 Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 5.0 isolation mode, HTTP.sys is used in almost the same way that it is in worker process isolation mode. The difference is that HTTP.sys routes requests to only a single request queue, which is maintained by the WWW Service Administration and Monitoring component of IIS. Depending on how you configure the IIS 5.0 isolation mode — as in-process or out-of-process — IIS processes the request in either Inetinfo.exe or in DLLHost.exe.

Figure 2.5 illustrates the processing of an SSL request for a Low-isolation application that is running in IIS 5.0 isolation mode. When HTTP.sys receives requests that are encrypted by using SSL, the kernel-mode HTTP service cannot decrypt the requests or encrypt the responses. Instead, HTTP SSL (HTTPFilter), which is a user-mode service, implements SSL for the HTTP service. In IIS 5.0 isolation mode, Inetinfo.exe, rather than Lsass.exe, hosts HTTPFilter.




In Figure 2.5, if you eliminate the decrypting and encrypting steps that are required for SSL requestsand follow only arrows 1, 4, 5, and 8, you see the standard (non-SSL) request process for a Low-isolation application that is running in IIS 5.0 isolation mode. For example, an HTTP request comes in to the HTTP protocol stack from the Internet, enters a single request-processing queue, and is picked up by a worker process that delivers the request to the Low-isolation application and returns it directly to the HTTP protocol stack.

Figure 2.6 illustrates request processing for an HTTP (non-SSL) request when IIS 6.0 is running in IIS 5.0 isolation mode. Note the extra process hops (represented by arrows 3 and 4) that occur when the request hops from the worker process, which is hosted by Inetinfo.exe and managed by the WWW service, to the High-isolation out-of-process application, which is hosted by DLLHost.exe.




If you compare Figure 2.6 with Figure 2.3, you can see the advantage of using worker process isolation mode, in which application isolation occurs without a performance penalty (the additional process hops are not needed) and in which health monitoring and worker process recycling promote the stability of your application server and its applications.

| **Note** |
| --- |
| In IIS5.0 isolation mode, request processing is almost identical for High-isolation and Medium-isolation applications. The only differences between these two isolation methods are that multiple applications can run in the pool for Medium isolation and there can only be one Medium-isolation process; however, only a single application can run in the pool for High isolation and you can have multiple High-isolation applications on one computer. |

Figure 2.7 illustrates the processing of an SSL request for a High-isolation application that is running in IIS 5.0 isolation mode. The HTTP protocol stack receives the SSL request from the client and places it in the SSL queue. The HTTPFilter listens for requests to be added to the SSL queue, picks up any requests that are queued, uses SSL to decrypt the request, and then returns the decrypted request to HTTP.sys. HTTP.sys then places the decrypted request in the worker process request queue (in this application isolation mode, there is only a single request queue), and a worker process pulls the request directly from the queue.




On the return trip, the worker process returns the response to HTTP.sys, which places it in the SSL queue, from which the HTTPFilter picks it up and performs SSL encryption. After encryption, the HTTPFilter returns the response to HTTP.sys, which sends the encrypted response back to the client.

If performance is a concern, and you are running applications in High isolation or Medium isolation on a computer that is running in IIS 5.0 isolation mode, keep in mind that using SSL requires more resources to process a request. In addition, the extra process hops required for High isolation and Medium isolation add to the resource requirements.

For more information about balancing performance and security, see Balancing Performance and Security. For more information about configuring IIS 5.0 isolation mode, see Running IIS 6.0 as an Application Server.

# How ASP.NET Works with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you run your ASP.NET applications in IIS 6.0, you obtain a significant advantage over running your ASP.NET applications in IIS 5.0. For example, in IIS 5.0, you cannot isolate individual applications into their own worker processes; hence, if one ASP.NET application fails, it can affect all the other applications. IIS 6.0, however, provides a new architecture that allows you to isolate your ASP.NET applications without incurring a performance penalty.

Before you run ASP.NET applications on IIS 6.0, be sure that you understand how ASP.NET works with the two IIS 6.0 isolation modes. When running in IIS 5.0 isolation mode, ASP.NET uses its own request processing model and configuration settings in the Machine.config file — just as it does when ASP.NET runs on IIS 5.*x*. When running in worker process isolation mode, ASP.NET disables its own request processing model and uses the new worker process architecture of IIS 6.0.

**Important**

The version of ASP.NET that is included with WindowsServer2003 is not available with the Microsoft® Windows® XP 64-Bit Edition and the 64-bit versions of Windows Server 2003. For more information, see Features unavailable on 64-bit versions of the Windows Server 2003 family in Help and Support Center for Windows Server 2003.

**ASP.NET Request Processing Model**

When ASP.NET is enabled on a server running Windows Server 2003, the ASP.NET request processing model depends on the application isolation mode in which IIS 6.0 is running. The preferred application isolation mode for IIS 6.0 is worker process isolation mode, which enables you to use application pools, worker process recycling, and application pool health monitoring. When IIS 6.0 is running in worker process isolation mode, ASP.NET uses the IIS 6.0 request processing model settings.

In almost all situations, if your ASP.NET applications are running on Windows Server 2003, configure them to run in worker process isolation mode. However, if your application has compatibility issues in this mode — for example, if your application depends upon read raw ISAPI data filters — you can change your configuration to run IIS 6.0 in IIS 5.0 isolation mode.

If you configure IIS 6.0 to run in IIS 5.0 isolation mode, ASP.NET runs in its own request processing model, Aspnet_wp.exe, and uses its own configuration settings for process configuration, which are stored in the Machine.config configuration file. When your server is running in IIS 5.0 isolation mode, the ASP.NET ISAPI extension implements a request processing model that is similar to worker process isolation mode and contains worker process management capabilities similar to those provided by the WWW service. ASP.NET also provides recycling, health monitoring, and processor affinity.

When ASP.NET is installed on servers running Windows XP or Windows 2000 Server, ASP.NET runs in IIS 5.1 or IIS 5.0, respectively, and runs in the Aspnet_wp.exe process by default. The ASP.NET request processing model provides process recycling, health detection, and support for affinities between worker processes and particular CPUs on a server with multiple CPUs. For more information about running ASP.NET in its own request processing model, see Using the ASP.NET Process Model.

**When Different Versions of ASP.NET Share the Same Application Pool**

When multiple versions of the .NET Framework are installed on a computer that uses IIS 6.0, you might encounter the following error message in the Application Event log:

**Note**

It is not possible to run different versions of ASP.NET in the same IIS process. Please use the IIS Administration Tool to reconfigure your server to run the application in a separate process.

This error occurs when more than one version of ASP.NET is configured to run in the same process. Different versions of the .NET Framework and run time cannot coexist side by side within the same process. Therefore, an ASP.NET application that uses one version of the run time must not share a process with an application that uses a different version. This error commonly occurs when two or more applications are mapped to different versions of ASP.NET but share the same application pool.

For more information about using ASP.NET with IIS 6.0 application pools, see Configuring an ASP.NET Application for an ASP.NET Version. For more information Deploying ASP.NET Applications, see Deploying ASP.NET Applications in IIS 6.0. For more information about ASP.NET, see About ASP.NET.

# Additional Resources for IIS 6.0 Architecture

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information related to this chapter.

Related Information

- Configuring Internet Sites and Services for information about the IIS FTP, NNTP, and SMTP services.
- Managing a Secure IIS 6.0 Solution for information about the security features available in IIS 6.0.
- Running IIS 6.0 as an Application Server for information about configuring application pool health monitoring, orphaning worker processes, creating Web gardens, configuring idle timeout, configuring worker process isolation mode and IIS 5.0 isolation mode, and using the built-in Network Service account.
- Deploying ASP.NET Applications in IIS 6.0 for information about deploying ASP.NET applications on IIS 6.0.
- Encryption with IIS 6.0 for information about how to use encryption to help reduce the attack surface of your Web server.
- Metabase Structure in IIS 6.0 for information about the architecture of the IIS metabase.
- About ASP.NET for information about ASP.NET.
- Configuring an ASP.NET Application for an ASP.NET Version for information about using ASP.NET with IIS 6.0.
- Configuring Web Gardens in IIS 6.0 for information about creating Web gardens.
- Configuring Health Monitoring for information about configuring IIS 6.0 for application pool health monitoring.
- Configuring Worker Process Identities and IIS and Built-in Accounts for information about IIS security, including the recommended use of the built-in Network Service account.
- IdleTimeout Metabase Property metabase property for information about using this property to specify how long a worker process runs idle if no new requests are received.
- OrphanWorkerProcess Metabase Property for information about using this property to separate (instead of terminate) a failing worker process from its application pool.
- Balancing Performance and Security for information about successfully balancing performance and security.
- Recycling Worker Processes with IIS 6.0 for information about worker process recycling, including when to use it, how to configure it, and how to configure overlapped recycling.
- Using the ASP.NET Process Model for information about running ASP.NET in its own request processing model.

# Application Pool Configuration in IIS 6.0

Published: February 19, 2010

Updated: February 19, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic provides some best practices recommendations for configuring IIS application pools on your IIS 6.0 Web server. Application pools enable specific configuration settings to be applied to groups of applications, and the worker processes that service those applications. All Web sites and virtual directories must be assigned to an application pool. Configuring application pools appropriately can increase service availability on the server as application pools isolate sites and applications from other sites and applications.

**Planning Application Pool Configuration**

You can configure application pools to optimize performance for your server and organization. The ideal configuration for your environment depends on many factors. These include hardware and network configuration, the number of Web sites that are running on the server, and the applications on those sites. The guidance provided here is intended to highlight points to consider instead of the specific settings for your particular environment. For example, you may want to isolate sites or applications by technology, such as ASP.NET, or by business segment, such as Human Resources or Finance. You may also want to isolate sites and applications by business priority or by specific application. You can configure application pools using one or more of these categories, or define categories to support your specific business needs.

**Application Pools and Service Availability**

How you configure application pools on your server directly affects the availability of the WWW service on your IIS server. For example, if there are two Web sites that are running in a single application pool, the same worker process (W3wp.exe) handles requests for both the sites. If one of the sites experiences a fatal error, the process stops processing requests for both the sites in the application pool. To avoid this, configure a separate application pool for each site so that errors in one are isolated from the other sites on the server.

As a best practice, you should configure application pools to isolate mission critical applications from other applications on the server. This can also result in better performance for critical applications as the computer continues to process requests for the critical application, even if there are problems in other application pools. Requests for other applications or sites are handled by the process in a separate, isolated application pool.

In some cases, administrators prefer to configure each site on the server to run in its own application pool. This is a supported configuration. However, exceeding about 60 application pools that use unique identities on a server requires special configuration, and may cause a decrease in performance because of high utilization of hardware resources. If it is a priority to configure more than 60 application pools but hardware limitations prevent doing this, you could also add another IIS server to your environment and then run some sites and applications on the additional server.

**Advanced Configurations**

If you are configuring your application pools with unique identities, depending on the applications and memory resources on your server, you will reach a limit of about 60 application pools. There are finite limits to some system resources that get allocated with each new logon session. This means that 60 processes can run concurrently as distinct accounts. IIS 6.0 supports running these processes in a single shared workstation and desktop, at a cost of sharing a single encapsulation of a user session among all parties.

To scale beyond 60 application pools and to share a single desktop, change the value of the registry **HKey_Local_Machine\System\CurrentControlSet\Services\W3SVC\Parameters\UseSharedWPDesktop** subkey to **1** (the data type is DWORD). After changing this registry entry, you can scale to hundreds of application pools and hundreds of concurrently running worker processes.

| **Warning** |
| --- |
| Before you change a registry key or subkey, you should export or make a backup copy of the key or subkey. You can save the backup copy to a location that you specify, such as a folder on your hard disk or a removable storage device. If you make changes that you want to undo, you can import the backup copy. |

If IIS starts more worker processes than can be supported, or if the server is running out of resources, you can set a maximum number of concurrent worker processes. By using the IIS 6.0 and HTTP.sys architecture, and with a request load distributed across many application pools, IIS will try to start more worker processes than the server can sustain without running out of memory or CPU resources. The **DemandStartThreshold Metabase** property (also named *concurrent process gating*), when

set to a value less than the default, applies a hard limit to the number of IIS worker processes that can run concurrently. If the hard limit is exceeded, HTTP.sys returns an HTTP 503 error (Service Unavailable).

# Running IIS 6.0 as an Application Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0, running on Microsoft® Windows® Server 2003, helps provide a secure, reliable, and easily managed application server on which you can host sites over an intranet, the Internet, or an extranet. IIS helps you create a platform of dynamic network applications, allowing you to use the latest Web standards to develop, implement, and manage your Web applications.

To match the needs of your applications to your server, IIS offers an adjustable architecture that allows you to isolate applications within a self-contained worker process. The new process model helps prevent one application or Web site from stopping another and reduces the time that you spend restarting services to maintain the health of your server. If health issues arise, IIS helps you manage them, usually without restarting your Web server or affecting your users.

**In This Section**


Introduction to Running an Application Server

Configuring an Application Server Role

Enabling Dynamic Content

Configuring an Application Isolation Mode

Creating Application Pools

Ensuring Application Health in Worker Process Isolation Mode

Managing Resources in Worker Process Isolation Mode

Running Web Applications

Additional Resources for Running IIS 6.0 as an Application Server

# Introduction to Running an Application Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Establishing an application server begins with configuring the application server role, which includes deciding what Web-related technologies you want to use and then enabling those technologies on your server. Attackers often target unused or unmonitored services; so if you enable only what you need, your IIS 6.0 application server provides a reduced attack surface when it is connected to the Internet.

In IIS 6.0, deciding which request processing model to use is your next decision. Helping you make this decision is IIS 6.0 Architecture, the section that introduces the two IIS 6.0 request processing models, known as *application isolation modes*. Where the earlier section provides architectural diagrams and descriptions of the two application isolation modes, this section offers guidelines for deciding where and when to use each mode and provides procedures for configuring the isolation modes and their features. For information about choosing and configuring an application isolation mode, see Configuring an Application Isolation Mode.

Your application server runs in one of two mutually exclusive application isolation modes: worker process isolation mode or IIS 5.0 isolation mode. Worker process isolation mode, which is the default request processing model, supports new features such as application pools, health monitoring, worker process recycling, and rapid-fail protection. These new features help provide more robust performance. For example, you can recycle faulty applications to prevent service disruption and can use application pools to isolate one application from another without incurring a performance penalty. For more information about using these new features, see Creating Application Pools and Ensuring Application Health in Worker Process Isolation Mode.

Worker process isolation mode also provides better default security for your Web applications — in particular, better balance between security and functionality. To help you evaluate the security differences between the two application isolation modes, see Security Considerations When Choosing an Application Isolation Mode.

Worker process isolation mode provides several features to help you manage system resources, including Web gardens, idle timeout, CPU monitoring, processor affinity, and startup and shutdown limits. For information about configuring these features, see Managing Resources in Worker Process Isolation Mode.

To help you migrate from earlier versions of IIS, IIS 6.0 provides IIS 5.0 isolation mode, which is a processing model that offers compatibility for applications that depend on certain request processing characteristics of IIS 5.0. For example, if an application must run within Inetinfo.exe, it can only run in IIS 5.0 isolation mode unless you modify the application, making it compatible with worker process isolation mode. For information about which applications might require a processing model like IIS 5.0 isolation mode, see Evaluating Application Requirements.

You can configure most IIS 6.0 features either by using a graphical user interface (GUI) or by using a programmatic interface. For the procedures included here, the method that uses the GUI (usually implemented through IIS Manager) appears first, and an optional method for configuring metabase properties programmatically (by using the administrative utility Adsutil.vbs) appears second.

# Configuring an Application Server Role

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use the Configure Your Server wizard to install or remove many services on a computer running the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; or Windows® Server 2003, Web Edition operating system. The server roles that you can install by using this wizard include the following: file server, print server, application server (IIS, ASP.NET), terminal server, DNS server, and DHCP server.

The application server role can include IIS 6.0, COM+, ASP.NET, and FrontPage® Server Extensions from Microsoft. Before you install IIS 6.0 on your server, review the technologies offered in the application server role and decide which technologies you want to install and enable during initial installation, and which technologies you will install later as needed. Then prepare for the IIS 6.0 installation by reviewing security considerations.

# Technologies Offered in the Application Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you install IIS 6.0 as an application server by using the Configure Your Server wizard, IIS, COM+, and Microsoft® .NET Framework are automatically installed. To help ensure the security of your system after a default installation, ASP.NET and FrontPage 2002 Server Extensions are installed but not enabled in a default installation; the wizard gives you the option of enabling these services.

Table 3.1 describes the technologies included in the IIS 6.0 application server role, noting which are installed and enabled by default when you use the Configure Your Server wizard to install an application server role on a computer running Windows Server 2003.

**Table 3.1 Technologies That Make Up the Application Server Role**

| Technology | Enabled by Default |
|---|---|
| **Internet Information Services (IIS).** Gives you a full-featured Web server that provides the infrastructure for .NET and existing Web applications and Web services. IIS 6.0 is optimized to run Web applications and XML Web services in a hosting environment. | Yes |
| **COM+ for remote transactions.** Provides an extension to COM. COM+ builds on COMs integrated services and features, making it easier for developers to create and use software components in any language by using any tool. | Yes |
| **Distributed Transaction Coordinator (DTC) for remote access**. Coordinates COM+ transactions across multiple computers. | Yes |
| **ASP.NET.** Offers a unified Web development platform that provides the services necessary for developers to build enterprise-class Web applications.[1] | |
| **FrontPage Server Extensions.** Enables users on a client computer to publish and administer Web sites on a server remotely over a network.[1] | |

[1] To ensure a more secure server during default installations, ASP.NET and FrontPage Server Extensions are automatically installed but in a disabled state. You must enable each of these technologies before you can use them. You can optionally enable either ASP.NET or FrontPage Server Extensions (or both) when using the Configure Your Server wizard by selecting the appropriate check box.

| **Important** |
|---|
| The version of ASP.NET that is included with Windows Server 2003 is not available in the Microsoft® Windows XP 64-Bit Edition and the 64-bit versions of Windows Server 2003. For more information, see Features unavailable on 64-bit versions of the Windows Server 2003 family in Help and Support Center for Windows Server 2003. |

If you are unsure whether you need either of the optional technologies, you can install IIS by using the default installation, which installs them in a disabled state, and then enable the optional technologies when you need them by using **Add or Remove Programs** in Control Panel. For more information about installing IIS 6.0 by using Control Panel, see Installing IIS.

As with every functionality that you enable on your server, use caution and take appropriate steps to secure each new technology that you add. For more information about securing your application server, see Managing a Secure IIS 6.0 Solution.

# Benefits of Enabling ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP.NET is an application environment that provides the services necessary to build and deploy enterprise-class Web applications. ASP.NET is installed but not enabled during a default installation; you must enable it before you can use your application server to host ASP.NET Web pages. ASP.NET is part of the .NET Framework, which simplifies application development in the highly distributed environment of the Internet.

ASP.NET offers several important benefits:

- **Manageability.** ASP.NET uses a text-based, hierarchical configuration system that makes it easier for you to apply settings to your server environment and Web applications. Because ASP.NET stores configuration information as plain text, you can apply new settings without using local administration tools.

- **Security.** ASP.NET provides default authorization and authentication schemes for Web applications. Developers can easily add to, remove, or replace these schemes.

- **Easy deployment.** You can deploy an ASP.NET application to a server by merely copying the necessary files to the server, without restarting the server.

- **Enhanced performance.** ASP.NET is compiled code that runs on the server. ASP.NET takes advantage of early binding, just-in-time (JIT) compilation, native optimization, and caching services out of the box.

- **Flexible output caching.** ASP.NET can cache page data, portions of a page, or entire pages. When ASP.NET is running on Windows Server 2003 with IIS 6.0, ASP.NET can cache entire pages in the high-performance IIS kernel cache.

- **Mobile device support.** ASP.NET supports any browser on any device.

- **Scalability and availability.** ASP.NET includes features specifically designed to improve performance in clustered and multiprocessor environments.

ASP.NET and ASP pages can run compatibly side by side on an IIS application server because they are serviced by different ISAPI extensions that do not interfere with each other. Installing ASP.NET does not corrupt any ASP applications; the ASP engine continues to process files with an .asp extension. You can isolate these two processing environments in separate application pools or run both types of application in the same application pool.

The security configurations for ASP.NET and IIS are complementary and in part interdependent. For example, when IIS is running in worker process isolation mode, the ASP.NET processes are secured by means of restrictions on worker process identities; when IIS is running in IIS 5.0 isolation mode, the ASP.NET processes are secured by means of the identity assigned within the Machine.config file. Additional IIS security features are configured in the IIS metabase, and additional ASP.NET security features are configured in XML-based Web.config files. To secure an ASP.NET Web application, you might need to configure security in both places.

For more information about ASP.NET, see Overview of ASP.NET.

# Benefits of Installing FrontPage Server Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage Server Extensions provides Web-based and command-line administration for extending virtual servers. Extending virtual servers with FrontPage Server Extensions enables the site owner to author the site in FrontPage and delegate site ownership and administrative credentials.

If you build and maintain Web sites for customers and coordinate these efforts with multiple authors, the features included in FrontPage Server Extensions can help you accomplish the following tasks:

- Manage Web sites either on the local server or remotely by using HTML Administration pages or a command-line interface.
- Secure Web sites, and grant authoring, browsing, site management, or other user rights to authorized users.
- Analyze site usage to find out which users are visiting the site and how often.
- Monitor the health of FrontPage Server Extensions and detect errors.

FrontPage Server Extensions is typically used on development, test, or staging servers but not on application servers in a production environment. Use appropriate caution if you install a publishing system like FrontPage Server Extensions on a production server.

If you are unsure whether you need FrontPage Server Extensions, you can install them after installing IIS 6.0 by using **Add or Remove Programs** in Control Panel. For more information about installing FrontPage 2002 Server Extensions, see Using FrontPage 2002 Server Extensions with IIS 6.0.

# Preparing to Install IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you install IIS and configure it as an application server, review the security considerations outlined in this section:

- Avoid installing IIS on a domain controller.
- Verify that all disk volumes on your application servers use the NTFS file system, not FAT or FAT32 file systems.
- Ensure that your computer has network connectivity and a static or dynamic IP address. For more information, see Before Configuring IIS and Securing Sites with IP Address Restrictions.

**Avoid installing IIS on a domain controller**

Whenever possible, do not install IIS on a domain controller for the following reasons:

- The networking and processor load generated by authentication and other domain controller processes can degrade IIS performance.
- Adding users to a group that can log on locally to the domain controller can increase the attack surface that is exposed by a Web application. If security is compromised on the domain controller, security is compromised on the entire domain.

**Verify that all disk volumes on your application servers use NTFS**

To provide a more secure file system, use NTFS on your application servers. NTFS is a more powerful and secure file system than the FAT or FAT32 file system. When you use NTFS on your application servers, you can limit access to the files and directories on the server and configure the access level that you grant to a particular user or group. An added benefit of NTFS is that it allows larger volume sizes than FAT.

Table 3.2 compares the benefits of using the NTFS and FAT file systems.

**Table 3.2 Comparison of Support Offered by NTFS and FAT File Systems**

| Type of Support | NTFS File System | FAT File System |
|---|---|---|
| Directory and file security | Enables more secure directories and files through the use of NTFS permissions, which you can set at both the file and directory levels. | Does not allow you to secure directories and files. |
| The Active Directory® directory service and domain-based security | Supports Active Directory and domain-based security. | No support. |
| File encryption | Supports file encryption, which enhances file security. | No support. |

To determine the type of file system used by a disk volume, view the properties of the volume (open **My Computer**, right-click the disk volume, and then click **Properties**). Windows provides a conversion utility (Convert.exe) that you can use to convert an existing FAT or FAT32 volume to NTFS without losing data. For more information about converting FAT or FAT32 volumes to NTFS, see "Reformatting or converting a partition to use NTFS" in Help and Support Center for Windows Server 2003.

# Installing IIS 6.0 Technical Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To help protect the system from malicious users and attackers, IIS is not installed by default when you install Windows Server 2003. In addition, when you initially install IIS, the service is installed in a secure, locked mode.

In Windows Server 2003, you can install IIS by using either Control Panel or the Configure Your Server wizard. These options offer slightly different default installations:

| **Note** |
| --- |
| You can use the Web Platform Installer (Web PI) to easily install IIS, and applications that run on IIS. The Web PI is a free, lightweight tool that lets you install IIS and related technologies such as ASP.NET, SQL Server Express, Visual Web Developer, other popular Web applications, and more. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI. |

- If you install IIS by using **Add or Remove Programs** in Control Panel, only IIS is installed and enabled, so the default installation serves only static content. Features such as ASP, ASP.NET, server-side includes (SSI), Web Distributed Authoring and Versioning (WebDAV) publishing, and FrontPage Server Extensions are either not installed or not enabled, and therefore do not work.

- If you install IIS by using the Configure Your Server wizard, ASP.NET and ASP are automatically installed with IIS. In addition, you have the option of enabling ASP.NET and installing FrontPage Server Extensions by using the wizard.

For more information about installing IIS by using **Add or Remove Programs** in Control Panel, or by using the Configure Your Server wizard, see Installing IIS.

If you want to serve dynamic content, you can unlock applications such as ASP.NET, ASP, and FrontPage Server Extensions through the **Web Service Extensions** node in IIS Manager. For more information, see Enabling Dynamic Content.

| **Note** |
| --- |
| When you upgrade from a previous version of IIS, most dynamic content features remain enabled. The installation program notes the exceptions. |

Be sure to install IIS on an NTFS-formatted drive. NTFS is a more powerful and secure file system than FAT or FAT32. For more information about the advantages of using NTFS, see Preparing to Install IIS 6.0.

# Enabling Dynamic Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The default installation of IIS 6.0 protects your system from malicious users and attackers by delivering IIS in a locked mode, in which IIS can only serve static HTTP content. After a default installation, the request handlers that process dynamic content are disabled, which means that features like ASP, ASP.NET, SSI, WebDAV publishing, and FrontPage Server Extensions do not work.

You can configure the request handlers (ISAPI extensions or Common Gateway Interface [CGI] programs), known as *Web service extensions*, by using either the **Web Service Extensions** node in IIS Manager or the command-line script Iisext.vbs, which is located in *systemroot*\System32.

Iisext.vbs can be used to perform the same operations that are available in IIS Manager. You can use either tool to administer dynamic content, performing the following tasks:

- Allow or prohibit Web service extensions on the local computer.
- Add new Web service extensions.
- Allow all the Web service extensions that an application requires.
- Prohibit the running of all Web service extensions on the local computer.

For more information about using IIS Manager to administer dynamic content, including instructions for adding new Web service extensions, allowing an application to call a Web service extension, or disabling all Web service extensions, see Enabling and Disabling Dynamic Content in IIS 6.0.

To enable a Web service extension from the command line, use the Iisext.vbs command-line script. Before enabling a Web service extension or application type, obtain a list of all the defined Web service extensions on the computer. (To view a list in IIS Manager, expand the **Web Service Extensions** node.)

**Important** You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /Profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters).

**To view a list of a Web service extensions on the local computer by using Iisext.vbs**

1. In the **Run** dialog box, type **cmd**, and then click **OK**.
2. At the command prompt, type:

   **cscript %SystemRoot%\system32\iisext.vbs /listapp**

   This command displays a list of all Web service extensions that are located on the local computer. For example, using the */listapp* parameter might return the following:

   ```
   Connecting to server ...Done.
   Active Server Pages
   Internet Data Connector
   Server Side Includes
   WebDAV
   FrontPage Server Extensions
   ASP.NET v1.1.4322
   ```

If you want to enable ASP pages, type the command-line syntax that is shown in the following procedure. If you want to enable another type of application, replace Active Server Pages with the name of that application. If the application name contains spaces, be sure to enclose the name with quotation marks.

**To enable ASP pages (or other Web service extension)**

1. At the command prompt, type:

   **cscript %systemroot%\system32\iisext.vbs /EnApp "Active Server Pages"**

For more information about using the lisext.vbs command-line script to administer dynamic content, such as enabling, disabling, or listing Web service extensions and their files, see Managing Applications and Web Service Extensions Using Scripts.

# Configuring an Application Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

So that you can match the needs of your server configuration to the requirements of your applications, and to provide more flexibility if you are upgrading from an earlier version of IIS, IIS 6.0 offers an adjustable architecture for isolating your applications. IIS 6.0 can run an application server in one of two process models, called application isolation modes:

- Worker process isolation mode, which is the default application isolation mode for IIS 6.0, isolates key components of the World Wide Web Publishing Service (WWW service) from the effects of defective applications and protects applications from each other by using isolated worker processes.

- IIS 5.0 isolation mode, which processes requests in nearly the same way that IIS 5.0 does, is provided for applications that depend on specific features and behaviors of IIS 5.0 or earlier. Use this isolation mode only if an application has a compatibility problem, which you cannot resolve, when it runs in worker process isolation mode.

Both application isolation modes use the HTTP protocol stack, HTTP.sys, in almost identical ways. The primary difference is that in worker process isolation mode, HTTP.sys can route requests to multiple queues; in IIS 5.0 isolation mode, HTTP.sys routes requests to a single queue. For more information about IIS 6.0 request processing, see IIS 6.0 Architecture.

# Choosing an Application Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

One of the most important decisions that you will make when you configure your application server is which application isolation mode to use. Although worker process isolation mode is recommended if your environment supports it, you must evaluate the relative capabilities of the two application isolation modes, considering their compatibility with the applications that the server will make available, the level of security that your organization requires, and performance requirements.

# Evaluating the Capabilities of the Two Application Isolation Modes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you choose an application isolation mode for your application server, review the benefits of each mode and then review your ISAPI (Internet Server API) filters and applications to determine whether any conflicts exist with worker process isolation mode. Keep in mind that when an application server is running in IIS 5.0 isolation mode, many of the benefits of the new request-processing architecture — such as application pools, recycling, and health monitoring — are unavailable. If a filter or application cannot run in worker process isolation mode, evaluate whether you can substitute another filter or application that does not conflict with worker process isolation mode.

Benefits of Worker Process Isolation Mode

In worker process isolation mode, you can isolate an application within its own process, where it does not depend on a central process such as Inetinfo.exe. When you run IIS in this isolation mode, you can group Web applications into *application pools*, through which you can apply specific configuration settings to groups of applications and to the worker processes that service those application groups.

An application pool corresponds to one request routing queue within HTTP.sys and one or more worker processes. Process boundaries separate each application pool from the others so that when an application is routed to one application pool, applications in other application pools do not affect that application. Application pools allow you to run all application code in an isolated environment without incurring a performance penalty.

Worker processes operate independently; if one fails, it does not affect other worker processes. The pooling of applications protects applications from the effects of worker processes that support other application pools. In this way, applications are protected from each other.

In worker process isolation mode, a kernel-mode device driver called the HTTP protocol stack (HTTP.sys) routes HTTP requests directly to a kernel-mode request queue that is assigned to a specific application pool. Worker processes that serve the application pool pull incoming requests directly from the request queue, avoiding the overhead of user-mode process switching.

To protect the WWW service, IIS 6.0 isolates critical components, such as HTTP.sys and WWW Service Administration and Monitoring, from external application code that runs in worker processes. HTTP.sys receives and queues requests for Internet services.

When a worker process enters an unhealthy state and stops processing requests, HTTP.sys continues to queue requests. The WWW service detects that the worker process is unhealthy and shuts it down, starting a new worker process if requests are waiting in the queue or if new requests arrive. Because HTTP.sys continues to accept incoming requests even when worker processes fail, end users are shielded from the failure; the client can still connect to the WWW service while a replacement worker process is initialized.

IIS 6.0 worker process isolation mode delivers the following improvements over earlier versions of IIS:

- **Robust performance.** Isolation prevents Web applications and Web sites in an application pool from disrupting applications or Web sites in other application pools. Restarts of the operating system and WWW service are avoided.
- **Self-healing.** Automated management provides auto-restart of failed applications and periodic restart of unstable applications or applications that leak memory.
- **Scalability.** IIS 6.0 supports hundreds (even thousands) of sites on a server through features such as start on demand and idle timeout for a worker process. In addition, IIS 6.0 offers improved multiprocessor scalability.
- **Processor affinity.** You can bind worker processes to specific processors on multiprocessor servers.
- **Automated debugging.** You can configure debugging tools to start automatically when worker processes fail.
- **CPU monitoring.** You can limit the amount of CPU resources that an application pool can consume within a configured amount of time.

For more information about worker process isolation mode, including architectural diagrams, see IIS 6.0 Architecture. For more information about automated debugging and CPU monitoring, see Enabling Debugging and Enabling CPU Monitoring.

Before you decide which application isolation mode to run on a server, read Evaluating Application Requirements.

Isolation Options in IIS 5.0 Isolation Mode

When IIS 6.0 is running in IIS 5.0 isolation mode, the following options for application isolation, which are similar to those in IIS 5.0, are available:

- **Low isolation (in-process).** Applications run in-process (on the Web server) as DLLs in Inetinfo.exe and are not protected from other applications that run in-process. The process identity (the security context of the process running the application) is LocalSystem.

- **Medium isolation (pooled out-of-process).** All pooled applications run as DLLs in the same single instance of DLLHost.exe and are protected from the effects of applications running in High isolation. However, because all pooled applications run in the same process, they are not protected from each other. The default application identity is IWAM_*ComputerName*.

- **High isolation (out-of-process).** Applications run as DLLs in their own DLLHost.exe and are protected from the effects of other applications running in Medium or High isolation. Also, all other applications that run on the same computer are protected from the effects of applications that run in High isolation. The default application identity is IWAM_*ComputerName*.

| **Note** |
| --- |
| Only the parts of applications that are served by ISAPI extensions can be run out-of-process — that is, in a memory space that is separate from the Web server. Thus, if an ASP page that is set to High-isolation includes static files, those static files are served by Inetinfo.exe. |

For more information about IIS 5.0 isolation mode, including architectural diagrams, see IIS 6.0 Architecture.

# Evaluating Application Requirements

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 cannot run both application isolation modes simultaneously on the same computer. Therefore, on a single computer running IIS 6.0, you cannot run some Web applications in worker process isolation mode and others in IIS 5.0 isolation mode. Applications that require different application isolation modes must be run on separate computers.

**Guidelines for Optimizing Application Performance**

When you are deciding which application isolation mode to use, follow these guidelines:

- Use worker process isolation mode unless you need to use IIS 5.0 isolation mode because of a specific compatibility issue. For example, read raw data filters only work in IIS 5.0 isolation mode.
- Run Web sites that contain static content or simple ASP applications in worker process isolation mode. These Web sites usually require little or no modification to run in worker process isolation mode.
- For applications that run correctly in IIS 5.0, try running them first in worker process isolation mode. If they do not work correctly, attempt to diagnose and resolve any compatibility issues. You can also turn off or scale down some of the features that worker process isolation mode provides. For example, you might turn off recycling or change the recycling settings so that the server does not perform an overlapped recycle.

If a server is running in IIS 5.0 isolation mode and you upgrade the servers applications so that they are compatible with worker process isolation mode, be sure to change the isolation mode on that server to worker process isolation mode. If, after changing the server from IIS 5.0 isolation mode to worker process isolation mode, you find that you cannot easily maintain a healthy server, change back to IIS 5.0 isolation mode while you diagnose and fix the compatibility issues. Then reconfigure the server to run in worker process isolation mode.

| **Important** |
| --- |
| Before you change the application isolation mode, consider that each time you reconfigure a server to use a different mode, you must restart IIS, which temporarily interrupts the service and can cause loss of application state data. |

**Applications That Cannot Run in Worker Process Isolation Mode**

Applications that run in worker process isolation mode can take advantage of all the worker process behaviors, which include multi-instancing, recycling, and configurable process identity. Certain application characteristics conflict with the architecture of worker process isolation mode. Applications with those characteristics must be run in IIS 5.0 isolation mode until you can modify them.

Applications with the following characteristics cannot run in worker process isolation mode:

- **Read raw data filters.** This type of ISAPI filter only works in IIS 5.0 isolation mode.
- **Dependency on Inetinfo.exe.** Applications that must run within an Inetinfo.exe worker process can only run in IIS 5.0 isolation mode.
- **Requires Dllhost.exe.** ISAPI applications that must run in a Dllhost.exe environment can only run in IIS 5.0 isolation mode because Dllhost.exe is not available for ISAPI applications in worker process isolation mode. (Dllhost.exe can still be used to host out-of-process COM+ applications, regardless of the mode in which IIS is running.)

If IIS 6.0 is running in worker process isolation mode (the default mode in IIS 6.0), and you must run applications that require a functionality that is only available in IIS 5.0 isolation mode, reconfigure the server to run in IIS 5.0 isolation mode until you can correct the incompatibility. Running your server in IIS 5.0 isolation mode means that you cannot take advantage of the reliability features in worker process isolation mode.

# Security Considerations When Choosing an Application Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you choose an application isolation mode, keep these security considerations in mind:

- **Worker process isolation mode.** Worker process isolation mode provides better default security for running Web applications than does IIS 5.0 isolation mode. In worker process isolation mode, worker processes run by default as Network Service, which is the account that provides a better balance between security and functionality.

- **IIS 5.0 isolation mode.** Web applications that are set to Low isolation, such as low-isolation ISAPI extensions, run in a process that runs as LocalSystem. The LocalSystem account can read, execute, and change all of the resources on the computer. Thus, if an attack by a malicious user takes over a Web application that runs in Low isolation, many assets of the local computer are open to the attacker. Also, if you set Web applications to Medium or High isolation, they run with IWAM_*ComputerName* as the default identity, which has fewer rights than LocalSystem.

Windows Server 2003 provides the following built-in system accounts, which you can use to provide the security context for worker processes:

- **Local Service.** Has limited rights on the local computer and limited access (Anonymous) to network resources. Use the Local Service account if the worker process does not require authenticated access to network resources.

- **Network Service.** Has limited rights on the local computer and authenticated access (as the computer account) to network resources.

- **LocalSystem.** Has full access to the system because it belongs to the Administrators group.

Table 3.3 provides additional details about the built-in system accounts that are available in IIS 6.0, and Table 3.4 shows the default security account for both application isolation modes**.**

**Table 3.3 Built-in System Accounts in IIS 6.0**

| Built-in Account | Privilege Level | Group or Account Used on the Local Machine | Group or Account Used on the Network |
|---|---|---|---|
| Local Service | Least privileged | Users group | Anonymous access account |
| Network Service | More privileged | Users group | Computer account |
| LocalSystem | Most privileged | Administrator with full access account | Computer account |

**Table 3.4 Default Security Accounts for Each Application Isolation Mode**

| Mode or Options Within a Mode | Default Account |
|---|---|
| Worker process isolation mode | Network Service |
| IIS 5.0 isolation mode | |
| ISAPI extensions set to Low isolation that run in Inetinfo.exe | LocalSystem |
| ISAPI extensions set to Medium-isolation that run in pooled out-of-process application hosts | IWAM_*ComputerName* |
| ISAPI extensions set to High-isolation that run in out-of-process application hosts | IWAM_*ComputerName* |
| ASP.NET worker processes | ASPNET |

For more information about built-in security accounts in IIS 6.0 and Windows Server 2003, including IWAM_*ComputerName*, see Managing a Secure IIS 6.0 Solution and IIS and Built-in Accounts.

# Performance Considerations When Choosing an Application Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you run IIS 6.0 in IIS 5.0 isolation mode, applications that run in Medium or High isolation incur a performance penalty when COM uses a remote procedure call (RPC) during these parts of the request processing cycle:

- When a request is sent to an application that is running in Medium or High isolation and the response is sent back.
- When a Web application that is running in Medium or High isolation communicates with the Web server during the processing of a request.

In contrast, when IIS 6.0 runs in worker process isolation mode, applications are loaded in-process to W3wp.exe and remain there. Therefore, no performance penalty is incurred.

| **Note** |
| --- |
| You cannot run SSI and Internet Database Connector (IDC) applications out-of-process. You need to run these applications in-process. |

For more information about performance issues, see Performance Tuning and Optimizing IIS 6.0 Performance.

# Setting an Application Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After configuring the IIS application server role and enabling dynamic content, configure the application isolation mode for the server. IIS 6.0 can run in one of two application isolation modes: worker process isolation mode, which is the default mode for IIS 6.0, or IIS 5.0 isolation mode. To configure IIS 6.0 to run in a specific application isolation mode, use either IIS Manager or the Adsutil.vbs administrative utility.

If you are configuring IIS to run in IIS 5.0 isolation mode, after setting the application isolation mode, you must perform the additional task of setting the level of isolation for each application hosted on the server.

For information about choosing the appropriate application isolation mode for your environment, see Choosing an Application Isolation Mode. For more information about configuring the application mode for a server, see Configuring Isolation Modes.

# Setting Isolation for Applications Running in IIS 5.0 Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you configured the server to run in IIS 5.0 isolation mode, you must now set the level of application protection, or isolation, for each application that runs on the server. Table 3.5 summarizes the security and performance considerations to take into account when you choose an application isolation mode. If you do not set an isolation level, by default, IIS 5.0 isolation mode runs an application in Medium isolation pooled out-of-process.

**Table 3.5 Isolation Levels in IIS 5.0 Isolation Mode**

| Isolation Level/Host | Protection | Default Security Account | Performance Penalty? |
|---|---|---|---|
| Low/Inetinfo.exe | **In-process.** Not protected from other applications that run in-process. | LocalSystem | |
| Medium/ DLLHost.exe | **Pooled out-of-process.** Protected from the effects of applications that run in High isolation; however, the applications are not protected from each other. | IWAM_*ComputerName* | Yes |
| High/ DLLHost.exe | **Out-of-process.** Protected from the effects of other applications that run in Medium or High isolation; cannot affect other applications running on the same computer. | IWAM_*ComputerName* | Yes |

For more information about the security and performance implications of running applications at each of these application levels, see Security Considerations When Choosing an Application Isolation Mode and Performance Considerations When Choosing an Application Isolation Mode. For more information about setting or changing the level of application protection for applications running in IIS 5.0 isolation mode, see Isolating Applications in IIS 6.0.

# Creating Application Pools

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you run IIS 6.0 in worker process isolation mode, you can isolate different Web applications or Web sites in pools, which are called *application pools*. An application pool is a group of URLs that are routed to one or more worker processes that share the same configuration. The URLs that you assign to an application pool can be for an application, a Web site, a Web directory, or a virtual directory.

In an application pool, process boundaries separate each worker process from other worker processes so that when an application is routed to one application pool, applications in other application pools do not affect that application.

By using an application pool, you can assign specific configuration settings to a worker process (or, in the case of a Web garden, to a set of worker processes) that services a group of applications. For example, you can configure worker process recycling, which offers several configuration options to match the needs of each application. If, for example, you suspect that an application has a memory leak, you might configure the application pools worker process to recycle when its memory use reaches a certain threshold. If another application fails because of the volume of requests that it receives, you can set the application pools worker process to recycle when the application exceeds a specified number of requests.

By creating new application pools and assigning Web sites and applications to them, you can make your server more efficient, reliable, and secure, and ensure that your applications remain available even when a worker process serving an application pool is recycled because of a faulty application.

For more information about the application pool architecture in IIS 6.0, including how application pools work, see IIS 6.0 Architecture. For more information about creating application pools and configuring worker processes, see Configuring Application Pools in IIS 6.0 and Configuring Application Pool Identity with IIS 6.0.

# Guidelines for Creating Application Pools

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Consider the following guidelines when you create application pools:

- To isolate Web applications on a Web site from Web applications on other sites running on the same computer, create an individual application pool for each Web site.

- To improve application security, you can specify a unique user account (process identity) for each application pool. If you specify a unique user account, be sure to add that account to the IIS_WPG group. When you use a built-in account, use Network Service if possible because it offers a better balance between security and functionality.

  | **Note** |
  | --- |
  | If you specify a large number of unique user accounts, your Web server might reach a limit to the number of worker processes you can start under separate identities, in which case, all subsequent application pools fail with a message of Service unavailable. |

  The use of unique user accounts also improves security auditing capabilities, because you can more easily trace security events to the corresponding applications that use the security context — that is, the process identity —of the account listed in the security events. For information about auditing security events, see Auditing in IIS 6.0.

- If you want to configure an application to run with its own unique set of properties, create a unique application pool for that application.

- If you test an application on the same server on which the application is running, create separate application pools and be sure to use different virtual directories for the test and production versions of the application. Using a test application pool for the test virtual directory helps to isolate the test version of the application.

  By running the test version in a separate application pool, you reduce the risk of the test version crashing or corrupting the production version because the applications are separated by process boundaries. However, separation by application pool does not protect any shared resources that both versions use, such as shared COM components or common data sources (for example, flat files, databases, or registry keys). You can improve the isolation of COM components by using side-by-side assemblies, but you must segregate data sources by some other means to completely isolate the applications. For information about using side-by-side assemblies, see Configuring Applications to Use COM+ Services.

# Ensuring Application Health in Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Worker process isolation mode provides several features that can help you identify and correct faulty applications in your application pools before they affect your systems or your users. These features often do the work for you, ensuring that you have robust performance and healthy systems:

- **Application pool health monitoring.** Provides the basis for application pool health and health monitoring, including worker process health monitoring and debugging.
- **Recycling worker processes.** Periodically restarts worker processes in an application pool.
- **Rapid-fail protection.** Protects system resources when a worker process that is serving an application pool fails repeatedly within a short period of time.

IIS 5.0 isolation mode does not offer these features.

# Monitoring Application Pool Health

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The WWW service monitors the health of worker processes by periodically pinging them to determine their responsiveness. If a worker process fails to respond to a ping (for example, the worker process might not have available threads to process the incoming request), the WWW service either terminates the worker process and creates another worker process to replace it, or releases the worker process but lets it continue to run. You can configure IIS to perform a specified action when an unhealthy worker process is released — for example, you can specify that IIS attach a debugger to a released worker process.

Additionally, the WWW service maintains a communication channel to each worker process and can detect a drop in the communication channel, which indicates a worker process failure. When a worker process fails of its own accord, the WWW service starts another worker process in its place, provided the application pool does not go into rapid-fail protection.

In addition to monitoring the health of worker processes, the WWW service can detect ongoing problems with the entire application pool by using rapid-fail protection. For example, if worker processes are terminating abnormally every few seconds, the WWW service can determine that the application pool is unhealthy and disable it, preventing the unhealthy application pool from affecting applications in other pools.

The following conditions cause an application pool to start or stop:

- Initiation of rapid-fail protection causes an application pool to stop. For more information about rapid-fail protection, see Configuring Rapid-Fail Protection.

- CPU monitoring shuts down an application pool after it reaches the configured Maximum CPU use limit. When the limit is reached, the application pool shuts down and remains down until the **CPUResetInterval** time limit expires, at which time the application pool restarts. For more information about CPU monitoring, see Enabling CPU Monitoring.

- A configuration error occurs — such as when an incorrect user name or password, or a nonexistent or misconfigured account identity, is used — causing an application pool to stop. For more information about configuring worker process account identities, see Managing a Secure IIS 6.0 Solution and Configuring Worker Process Identities.

- A Windows administrator performs a demand stop or start on an application pool. For more information about performing demand stops and starts on an application pool with the IIsApplicationPool (WMI) class and the IIsApplicationPool (ADSI) class, see the IIS Programmatic Administration Reference in the Platform SDK on MSDN.

You can enable health monitoring for an application pool by using IIS Manager or by using the Adsutil.vbs utility to update metabase properties from the command line. For more information about configuring health monitoring, see Configuring Health Monitoring.

# Using an ISAPI Extension That Declares Itself Unhealthy

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use an ISAPI extension application that is built to programmatically signal the WWW service that it needs to be recycled. ISAPI developers can configure an ISAPI extension to signal the WWW service by using the **ServerSupportFunction** callback function together with the parameter *HSE_REQ_REPORT_UNHEALTHY*.

The ISAPI extension requires an internal mechanism for determining that it is unhealthy, such as monitoring the status of its internal thread pool. During the ping operation, the WWW service detects whether a worker process is hosting an ISAPI extension that is unhealthy and needs recycling; therefore, for this function to work, you must enable worker process pinging on the application pool that is running the ISAPI extension.

When you use this type of ISAPI extension and pinging is enabled, consider that it will cause a worker process recycle in which the ISAPI extension is running. This affects all applications that are running in that worker process.

The ASP and ASP.NET ISAPI extensions apply logic that takes advantage of this feature as it monitors the status of its internal thread pool. If too many of its threads enter a blocking state, it reports itself as unhealthy, which triggers a worker process recycle if pinging is enabled for the application pool.

For more information about using the **ServerSupportFunction** callback function with *HSE_REQ_REPORT_UNHEALTHY*, see HSE_REQ_REPORT_UNHEALTHY in the IIS Platform SDK on MSDN.

# Enabling Debugging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By default, when a worker process is determined to be unhealthy, the WWW service terminates the unhealthy process and starts a new worker process to serve new or queued client requests. Optionally, you can allow an application pool to keep the worker process running, but release, or *orphan*, the worker process from the application pool for purposes of debugging while at the same time initializing a new worker process to handle incoming client requests. An orphaned worker process ceases to take new requests from HTTP.sys.

If you enable debugging on an application pool, the WWW service allows an unhealthy worker process to keep running when the WWW service releases the worker process from serving an application pool. When enabling debugging, you can also configure the WWW service to start an executable application or script for a released worker process. For example, the WWW service might start an application that sends you an e-mail notifying you of the failure.

Be sure to closely monitor any worker process that is released from its application pool but continues to run. The worker process can still terminate because the WWW service leaves it running but no longer monitors it, which effectively tells the worker process to shut down without forcing it to do so. If the worker process recovers from its unhealthy state, it detects that it has no relationship to the WWW service and self-terminates. When an orphaned worker process self-terminates, you might find an event log entry stating that a worker process was released but find no evidence of the worker process ever having run.

If you enable debugging on unhealthy worker processes, be sure to monitor these released worker processes, because IIS does not remove them from memory. If you do not correctly handle worker processes that you keep alive for debugging purposes, you can have large numbers of failed worker processes running on your computer. These worker processes can tie up resources needed by other processes. You might need to terminate them quickly to free those resources. In some cases, these worker processes can block metabase access and cause problems with other worker processes or with the WWW service itself.

For more information about setting a metabase property to configure the WWW service to release rather than terminate an unhealthy worker process, see the OrphanWorkerProcess Metabase Property. For more information about configuring IIS to start an executable application or script when a worker process is orphaned, see Troubleshooting IIS 6.0, the OrphanActionParams Metabase Property and the OrphanActionExe Metabase Property.

# Ensuring Application Pool Health

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Health monitoring cannot detect application failures that do not cause the worker process to fail or do not block the available threads in the worker process. For example, an application can return an error response code, such as an HTTP 404 error (Not Found), but still respond normally to a ping from the WWW service. (The exception is an application like a custom ISAPI extension that implements specific code to indicate its unhealthy state.)

At a minimum, monitoring application pool health and taking corrective actions requires that you take the following steps:

- Allow the WWW service to detect unhealthy worker processes by enabling the WWW service to ping worker processes.
- Configure rapid-fail protection to disable application pools when the worker processes assigned to them fail a set number of times within a specified period.

For more information about rapid-fail protection, see Configuring Rapid-Fail Protection.

# Recycling Worker Processes in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In worker process isolation mode, you can configure IIS to periodically restart worker processes so that you can recycle faulty Web applications. Recycling ensures that applications in those pools remain healthy and that system resources are recovered.

You can configure worker processes to restart based on one of several options, including elapsed time, the number of requests served, scheduled times, and memory usage.

IIS 5.0 isolation mode does not offer these features.

# How Worker Process Recycling Works

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The WWW service can recycle a worker process in one of two ways, both of which do not interrupt service:

- By default, the WWW service establishes an *overlapped recycle*, in which the worker process that is to be shut down is kept running until after a new worker process is started.

- If the applications in an application pool cannot tolerate multi-instancing, you can optionally configure the WWW service to shut down a worker process and then start a new worker process. Use this type of recycling only if the applications in the application pool cannot tolerate multi-instancing. To turn off overlapped recycling, set the **DisallowOverlappingRotation** metabase property to **true**.

When a worker process requests a recycle, the WWW service initiates an overlapped recycle, creating a new worker process to replace the old one. While the new worker process is starting, the old process continues to serve requests. After the new process starts and initializes successfully, the WWW service instructs the old worker process to shut down. At this point, the old worker process stops accepting new requests from HTTP.sys and begins to shut down. The WWW service allows the old worker process a configured time period in which to finish processing its requests before the worker process is shut down. The WWW service terminates the worker process if it fails to shut down within the configured time.

With overlapped recycling, the old worker process remains in communication with HTTP.sys to handle requests until the new worker process is running. HTTP.sys establishes and maintains TCP/IP connections. When the WWW service recycles a worker process, it does not disconnect the existing TCP/IP connection. However, because the shutdown timeout value of a shutdown or startup is configurable, the worker process can be terminated while it is still serving requests if it does not finish servicing existing requests within the time limit.

If the application pool is a Web garden, which uses a set of equivalent worker processes, and you configure worker process recycling based on elapsed time, all the worker processes are recycled during the elapsed time set, but they are not recycled all at once. Worker process recycling in a Web garden is spread over the time set to reduce interruptions of service to client requests and to reduce the impact of restarting all processes at the same time.

Similarly, when you configure a Web garden to recycle applications based on the number of requests processed, recycling takes place at intervals in order to spread out the overhead that is associated with the worker process recycling.

For more information about Web gardens and how to configure them, see Configuring Web Gardens in IIS 6.0. For more information about how you can recycle a Web garden based on length of time, including examples of Web garden recycling, see Deploying ASP.NET Applications in IIS 6.0.

# When and How to Use Worker Process Recycling

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Consider the following guidelines when you decide whether to enable worker process recycling. Worker process recycling is not intended as a replacement for fixing a faulty application; however, recoding is not always an option, especially in cases when other application code is running that cannot be modified.

Consider using recycling in these situations:

- Your Web server hosts faulty applications that you cannot fix.
- Your Web server is experiencing undetermined or intermittent problems.
- Performance monitoring suggests that an application is leaking memory.
- You have implemented a temporary reset solution, such as a scheduled execution of the IISReset utility or a scheduled restart.

You might not need to use recycling in these situations:

- You are hosting a Web site that contains static content only and no custom ISAPI applications.
- You are hosting an application that is thoroughly tested and does not have memory or resource allocation issues.

To effectively use recycling, see Table 3.6, which presents the criteria on which recycling can be based.

**Table 3.6 When to Use Each Recycling Option**

| Recycle Trigger | Timing of Recycling | Conditions for Use |
|---|---|---|
| ISAPI unhealthy request | Recycles based on a request from an ISAPI extension in an application pool. | An ISAPI extension is built to declare itself unhealthy and pinging is enabled. |
| Elapsed time | Recycles after a specified number of minutes. | Applications cannot run for extended time periods. |
| Number of requests | Recycles when the number of HTTP requests exceeds a specified number of requests. | Applications fail because of the quantity of requests that they receive. |
| Scheduled times | Recycles at specified times within a 24-hour period. | Applications cannot run for an extended time period, or you want to periodically refresh applications during times of low use, such as on weekends or during late evening and early morning. |
| Virtual memory | Recycles when the worker process virtual memory (reserved plus used) reaches a specified maximum. | The memory heap becomes fragmented, which is caused by applications reserving memory multiple times. A steady increase in the use of virtual memory can indicate a fragmented memory heap. |
| Used memory | Also called *private memory*. Recycles when the memory that is used by the worker process reaches a specified maximum. | Some applications have memory leaks. |

| | | |
|---|---|---|
| User demand | Recycles when you use IIS Manager to manually recycle all the worker processes in an application pool, or when you run a script that calls the Recycle() method for the application pool. | A single application is causing problems while other sites are running, and you need to act quickly. Consider recycling the application pool instead of resetting the entire WWW service. |

For additional information about configuring worker processes for recycling, see Recycling Worker Processes with IIS 6.0, which provides detailed instructions for recycling worker processes immediately or based on the following criteria: elapsed time, number of processing requests, scheduled time, and maximum memory used.

# Logging Worker Process Recycling Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By setting the metabase property **LogEventOnRecycle**, you can direct the WWW service to log worker process recycling events in the Windows Server 2003 system log. You can use **LogEventOnRecycle** parameters to monitor and log recycle events based on the following criteria:

- Time consumed
- Requests processed
- Scheduled recycles
- On-demand recycles initiated by an administrator
- Memory consumed (both virtual and physical)
- Changes in application pool properties
- ISAPI unhealthy requests

For more information about logging worker process recycling events, see the LogEventOnRecycle Metabase Property.

# Configuring Rapid-Fail Protection

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When a worker process fails, the WWW service detects the failure and takes action, such as logging the event and restarting the process. You can also configure IIS so that if a worker process fails multiple times within a specified time period, its application pool is disabled to protect other application pools that are running on the same computer. This feature, known as *rapid-fail protection*, stops application pools from a continuous cycle of failing and restarting, which uses up system resources.

The WWW service shuts down an application pool whenever a worker process in the application pool fails often enough to equal or exceed the Rapid-Fail Protection (RFP) interval time window (for example: five failures in five minutes). The WWW service detects failure whenever:

- A worker process does not start within the startup time limit.
- A worker process does not shut down within the shutdown time limit.
- A worker process shuts itself down because of a fatal error and sends the WWW service an error code.
- A worker process fails to respond to a ping message.
- The WWW service detects that a worker process is sending non-standard communications (the worker process may have been taken over).

Rapid-fail protection places the disabled application pool in out-of-service mode, and HTTP.sys returns a 503 error (Service Unavailable) message for requests to that portion of the namespace, including requests that have already been queued for the disabled application pool. In addition, you can configure IIS so that when an application pool is stopped automatically, a specified action occurs, such as running a debugger or notifying the administrator that the application pool has stopped.

| **Note** |
| --- |
| Alternatively, HTTP.sys can do a connection reset based on the configuration of the **LoadBalancerCapabilities** property of the application pool. However, if a single computer hosts multiple applications, be sure to configure load balancers or switching hardware to reroute only the traffic intended for a failed application pool. Do not route requests away from healthy application pools, which can still receive and process requests. |

The rapid-fail protection feature is unavailable when IIS 6.0 is running in IIS 5.0 isolation mode. For more information about configuring rapid-fail protection, see Configuring Rapid-Fail Protection in IIS 6.0.

# Managing Resources in Worker Process Isolation Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0, when running in worker process isolation mode, offers several features that you can use to better manage your systems resources and to help maintain healthy and robust systems:

- Configure Web gardens.
- Configure idle timeout for a worker process.
- Configure startup and shutdown time limits.
- Enable CPU monitoring.
- Assign processor affinity.

# Configuring Web Gardens in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In worker process isolation mode, you can configure an application pool to be supported by multiple worker processes. An application pool that uses more than one worker process is called a *Web garden*. (Web gardens are to be distinguished from *Web farms*, which use multiple servers for a Web site.)

| **Important** |
| --- |
| Web gardens can reduce contention for system resources, such as threads, but they can also cause resource contention. For example, if application resources are designed for exclusive access by a single process, such as a flat file used for temporary storage, resource contention can occur between the worker processes within a Web garden. Also, there are documented application characteristics (singletons) for which Web gardens might cause increased contention. |

Creating a Web garden for an application pool enhances performance and reliability in the following ways:

- Robust processing of requests. When a worker process in an application pool is tied up (for example, when a script engine stops responding), other worker processes can accept and process requests for the application pool.

- Reduced contention for resources. When a Web garden reaches a steady state, each new TCP/IP connection is assigned, according to a round-robin scheme, to a worker process in the Web garden. The round-robin scheme smoothes out workloads and reduces contention for resources that are bound to a worker process.

For more information about configuring Web gardens, see Configuring Web Gardens with IIS 6.0.

# Configuring Idle Timeout for a Worker Process

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can conserve system resources by configuring worker processes to shut down after a specified period of inactivity. This feature, known as *idle timeout*, can be used to better manage system resources when the processing load is heavy, when identified worker processes are consistently idle, or when new processing space is not available. If an application pool is shut down due to inactivity and HTTP.sys receives a request for the idle application pool, the WWW service starts a new worker process to replace the one that timed out.

For more information about configuring idle timeout, see Configuring Worker Processes for Idle Timeout.

# Configuring Shutdown and Startup Time Limits

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For a worker process that has a requested recycle, you can set the maximum time that the worker process continues to process requests before it shuts down. When you set the shutdown time limit (the default is 90 seconds), consider that a worker process shuts down gracefully when the configured shutdown time is long enough to allow it to finish processing its queued requests. Be sure you tune the shutdown time period to fit the load characteristics of the applications that are being processed. If the worker process cannot shut down within the configured time limit, a termination event appears in the event log even when the application shuts down gracefully.

You can also configure a startup time limit (the default is 90 seconds), which specifies how long the WWW service waits for a worker process to start before identifying the worker process as unhealthy and in need of being shut down.

For more information about configuring shutdown and startup time limits, see Configuring Startup and Shutdown Time Limits.

# Enabling CPU Monitoring

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By using CPU monitoring, you can monitor worker processes for CPU usage and optionally shut down the worker processes that consume large amounts of CPU time. CPU monitoring is only available in worker process isolation mode.

To configure CPU monitoring for individual application pools, choose between two monitoring options: error event logging or the shutdown of all worker processes in the application pool that is consuming more CPU time than the specified limit.

Event logging

When the CPU usage for a specific application pool reaches a set limit, the WWW service writes an error to the Windows Server 2003 system log, which can be viewed in Event Viewer. The entry includes the worker process identifier and the name of the application pool that exceeded the CPU limit.

To view the event log, open Control Panel and access **Event Viewer** through Administrative Tools. Then double-click the system event log.

Shut down the unhealthy application pool (also includes event logging)

When you choose this option, IIS logs an error in the system log and shuts down all worker processes in the application pool. After the application pool shuts down, HTTP.sys returns an out-of-service message (503: Service Unavailable) for any request that is received during the interval when the application pool is stopped. The application pool remains shut down until the CPU reset interval (set by the **CPUResetInterval** metabase property) elapses. IIS then restarts the application pool, and CPU monitoring starts over again.

Obtaining more detailed CPU monitoring

If your Web server is running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, you can use Windows System Resource Manager to obtain more detailed CPU monitoring. Using Windows System Resource Manager, you can control how CPU resources are allocated to applications, services, and processes. For more information, see Windows System Resource Manager Frequently Asked Questions.

Enabling CPU Monitoring

You can enable CPU monitoring by using either IIS Manager or from the command line.

Enabling CPU monitoring from the command line can take as many as three steps:

- Set the **CPULimit** metabase property, which limits the worker processes in a specified application pool to a percentage of CPU time.
- Set the **CPUResetInterval** metabase property, which specifies the time interval for CPU monitoring.
- Set the **CPUAction** metabase property, which specifies the type of action you want IIS to take, such as write to the event log or shut down worker processes that exceed the CPU limit. Table 3.7 lists the possible values for **CPUAction**. Use only one value from the table for each application pool. If you do not set the **CPUAction** metabase property, IIS assigns the default value, which is zero (event logging only).

**Table 3.7 Values for the CPUAction metabase property**

| Value | Description |
|---|---|
| 0 | Writes a warning to the system log in Event Viewer when the worker processes for the application pool exceed the CPU limit, but takes no other action. |
| 1 | Writes a warning to the system log, and shuts down worker processes in the application pool, when the worker processes exceed the CPU limit. |

For more information about enabling CPU monitoring by using IIS Manager and from the command line, see Enabling CPU Monitoring in IIS 6.0.

# Assigning Processor Affinity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To efficiently use CPU caches on multiprocessor servers, you can configure application pools to establish affinity between worker processes and the processors, a feature known as *processor affinity*. Processor affinity forces worker processes to run on specific CPUs, and applies to all worker processes serving a particular application pool. You can also use processor affinity with Web gardens that run on multiprocessor computers where specific application pools are dedicated to specific clusters of CPUs.

Processor affinity typically yields the greatest performance benefits when you configure affinity on a computer that has eight or more processors. You can set processor affinity from the command line by using the **set** command with Adsutil.vbs. To set processor affinity, configure the **SMPAffinitized** and **SMPProcessorAffinityMask** metabase properties.

Table 3.8 shows which processors, on a computer with eight processors, are made available to worker processes that serve an application pool when various processor affinity masks are assigned.

**Table 3.8 Sample Processor Affinity Mask Values for Binding an Application Pool to Specific Processors on an Eight-Processor Computer**

| Hexadecimal Mask (Binary Representation) | Processor7 | Processor6 | Processor5 | Processor4 | Processor3 | Processor2 | Processor1 | Processor10 |
|---|---|---|---|---|---|---|---|---|
| 0x2 (0010) | | | | | | | Yes | |
| 0x5 (0101) | | | | | | Yes | | Yes |
| 0xD (1101) | | | | | Yes | Yes | | Yes |
| 0xF (1111) | | | | | Yes | Yes | Yes | Yes |
| 0x11 (00010001) | | | | | Yes | | | Yes |
| 0x1E | | | | Yes | Yes | Yes | Yes | |
| 0xF0 (11110000) | Yes | Yes | Yes | Yes | | | | |
| 0xFE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | |
| 0xFF | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

[1] Processor numbering begins with 0.

If you extend the table, a hexadecimal value of 0xFFFFFFFF (maximum value) allows all processors up to 32 to be available to worker processes that serve an application pool.

For more information about processor affinity mask settings, see the SMPProcessorAffinityMask Metabase Property and the SMPAffinitized Metabase Property. For more information about using the **set** command with Adsutil.vbs, see Processor Affinity.

# Running Web Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

With IIS 6.0 as your application server, you can use the latest Web application platforms and standards, such as ASP.NET, XML, and Simple Object Access Protocol (SOAP), to develop, implement, and manage your Web applications. While continuing to build on the ISAPI programming model, IIS 6.0 provides several new features, including custom errors, server-side redirection, streamlined file and buffer response, and improved Unicode support.

# Enhancements to ISAPI Support

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 supports both ISAPI extensions and ISAPI filters:

- *ISAPI extensions* are DLLs that handle specific requests.
- *ISAPI filters* are DLLs that you can register with IIS to modify the behavior of the server.

ISAPI extensions and filters load the first time a client requests resources that are processed by the ISAPIs and are kept in memory to handle future requests. ISAPI extensions handle specific requests from the client.

Like ISAPI extensions, ISAPI filters are programs that respond when the application server receives an HTTP request. Unlike ISAPI extensions, ISAPI filters are driven by Web server events instead of client requests. For example, you can write an ISAPI filter to control which files are mapped to a URL, modify the response sent by the server, and perform other actions to modify the behavior of the server.

IIS 6.0 supports these programmatic enhancements to ISAPI extensions:

- **HSE_REQ_EXEC_URL**
- **HSE_REQ_VECTOR_SEND**
- **HSE_REQ_REPORT_UNHEALTHY**
- **HSE_REQ_SEND_CUSTOM_ERROR**

These enhancements are new **ServerSupportFunction** callback functions, known as *support functions*. The **ServerSupportFunction** function, which returns data for several auxiliary functions that are not covered by other callback functions, is supplied in the EXTENSION_CONTROL_BLOCK structure that is associated with HTTP requests. IIS and the ISAPI extensions use the EXTENSION_CONTROL_BLOCK structure to exchange information.

For more information about ISAPI filters and extensions, including the **ServerSupportFunction** callback functions, see the IIS Web Application Technology Reference on MSDN.

Using HSE_REQ_EXEC_URL to Call Other URLS on the Server

The **HSE_REQ_EXEC_URL** support function (also known as **ExecuteURL**) allows ISAPI extensions to rewrite any part of the request, including the entity body, and have IIS reprocess this modified request. As long as the rewritten URL is in the same application pool and uses the same authentication, IIS can successfully reprocess the modified request.

This new support function gives ISAPI developers greater flexibility in transparently rewriting URLs on the server without the clients knowledge, allows an arbitrary number of URLs to be executed in a chain, and supports custom authentication. When used with wildcard application mapping, this support function can replace most read raw data filters by offering easy examination and modification of request entity body prior to request reprocessing. For information about wildcard application mapping, see Setting Application Mappings.

Using HSE_REQ_VECTOR_SEND to Consolidate Multiple Buffers

The **HSE_REQ_VECTOR_SEND** support function (also known as **VectorSend**) allows developers to group and send, in order, a list of buffers and file handles. This list is handed to IIS 6.0 and passed on to HTTP.sys, which compiles all the buffers and file handles within the kernel and then directly sends the group in one kernel-mode transition. This new functionality frees ISAPI from building large memory buffers or making multiple **WriteClient** calls.

In earlier versions of IIS, ISAPI developers have only two possibilities if they have multiple buffers that make up a response: they can either call **WriteClient** multiple times or assemble the response in one big memory buffer. The first approach causes a performance bottleneck because there is one kernel-mode transition for each **WriteClient** function call. The second approach also affects performance and uses additional memory. **VectorSend** is a more all-purpose transmit function for IIS 6.0 customers.

Using HSE_REQ_REPORT_UNHEALTHY to Recycle a Worker Process

The **HSE_REQ_REPORT_UNHEALTHY** support function (also known as **ReportUnhealthy**) allows an ISAPI extension to request that its worker process be recycled. Developers can use this new function to request a recycle if their ISAPI application becomes unstable or enters an unknown state. Note that in order to cause a recycle after an ISAPI calls **ReportUnhealthy**, health monitoring must be enabled for the ISAPI extensions application pool. When calling **ReportUnhealthy**, the developer can also pass in a string representing the reason why the ISAPI extension is calling **ReportUnhealthy**. This string is then added to the event that the worker process writes to the application log, which you can view in Windows Server 2003 Event Viewer.

Using HSE_REQ_SEND_CUSTOM_ERROR to Send Custom Error Messages

The **HSE_REQ_SEND_CUSTOM_ERROR** support function (also known as **SendCustomErrors**) allows ISAPI developers to send the client a specific IIS-configured custom error message, including all sub-error codes. This support function allows an ISAPI extension to better integrate its error messages with a Web sites custom error configuration.

Using the New Unicode ServerSupportFunctions

In addition to the preceding enhancements, IIS 6.0 allows ISAPI extensions to retrieve server variable values that are encoded as Unicode characters. In particular, ISAPI extensions can retrieve the request URL in Unicode, which can benefit international companies that have multilanguage Web sites for which the URL is not in the same code page as the server.

The following are new **ServerSupportFunction** callback functions that support values encoded as Unicode characters:

**HSE_REQ_EXEC_UNICODE_URL**

The **HSE_REQ_EXEC_UNICODE_URL** support function is functionally identical to **HSE_REQ_EXEC_URL** except that it expects the URL to be in Unicode instead of ANSI. This support function, when used with the new Unicode **GetServerVariable**, easily substitutes for **HSE_REQ_EXEC_URL**.

**HSE_REQ_MAP_UNICODE_URL_TO_PATH**

Allows your ISAPI extension to map a logical URL path to a physical path.

**HSE_REQ_MAP_UNICODE_URL_TO_PATH_EX**

Allows your ISAPI extension to map a logical URL path to a physical path, as well as gather several types of attributes that are associated with that physical path, such as access-control or cache-control flags.

# Using ASP in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP is a server-side scripting environment that you can use to create dynamic and interactive Web pages, and powerful Web applications. When the server receives a request for an ASP file, it processes the server-side script code contained in the ASP file to build a Hypertext Transfer Protocol (HTML) Web page, which is sent to the browser.

In addition to server-side script code, ASP files can contain HTML (including related client-side scripts) and calls to COM components that perform a variety of tasks, such as connecting to a database or processing business logic.

Because ASP scripts run on the server rather than on the client, your Web server does all the work involved in generating the HTML pages sent to browsers. Because only the result of the script is returned to the browser, the client cannot access the server-side scripts; therefore, users cannot view the script commands that created the page they are viewing.

For more information about ASP pages, see the IIS Web Development in the IIS Software Development Kit (SDK) on MSDN. For more information about writing server-side scripts with ASP and creating ASP applications, see the MSDN Library.

# Improvements to the ASP Programming Environment

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 6.0, you can make use of the following enhancements to the ASP programming environment:

- Improved international and UTF-8 support. UTF-8 support has been expanded to include all ASP built-in object properties and methods. The previous version of ASP supported UTF-8 only for Response.Write.
- Improved POST support. ASP can now read chunked-encoded POST data from a client.
- UNC enhancements. If your ASP scripts reside on a Universal Naming Convention (UNC) share, ASP has improved performance in executing those scripts. For information about configuring IIS 6.0 with remotely stored content on UNC servers, see Deploying and Configuring Internet Information Services (IIS) 6.0 with Remotely Stored Content on UNC Servers and NAS Devices.
- Optional script source access in ASP pages. *Script source access*, which allows access to the source code of scripts in ASP pages and other scripts, is new and is disabled by default. To enable script source access, you must select either Read or Write permission.
- COM+ services in ASP. You can now configure your ASP applications to use the following COM+ services without having to create a COM component:
    - Apartment model selection
    - Side-by-side assemblies
    - COM+ partitions
    - COM+ tracker

**COM+ Services in ASP**

In previous versions of IIS, if you wanted to use COM+ services from an ASP application, you had to create a COM component to call the methods of those services. In IIS 6.0, COM+ services are separate from the components. This separation allows ASP applications to use a set of COM+ services without instantiation by any COM components. Because of this separation, you can configure your ASP application to use the following COM+ services without having to create a COM component.

| **Note** |
| --- |
| Batch Logon is required for COM+ server-activated application identities. COM+ will add this automatically when you enter a new account for the identity of a COM+ application using the UI. Batch Logon gives your application the ability to execute code across servers, but does not allow the account to logon interactively to home server. |

Since only side-by-side assemblies can be enabled using IIS Manager, example code for enabling each service on an ASP application is provided below.

This topic includes the following information:

- Apartment Model Selection
- Side-by-Side Assemblies
- COM+ Partitions
- COM+ Tracker

**Apartment model selection**

ASP is now capable of running all of its threads in a multithreaded apartment (MTA). If your COM components are primarily free-threaded or both-threaded, running the ASP threads as MTA can improve performance significantly.

To enable an ASP application to run in an MTA, set the **AspExecuteInMTA** metabase property at the application level. This enables you to have one application running on ASP MTA threads and a second application running on ASP single-threaded apartment (STA) threads. By default, ASP runs applications in the STA.

| **Important** |
| --- |
| When you switch an ASP application from running in STA to MTA (or from MTA to STA), the impersonation token becomes obsolete, which can cause the application to run with no impersonation. Without impersonation, the application runs with the identity of the process, which might allow access to other resources. If you must switch threading models, disable and unload the application before you make the change. |

The following example sets the Default Web Site application (*W3SVC/1/ROOT*) to execute in MTA:

```
On Error Resume Next
set providerObj = GetObject("winmgmts://MyMachine/root/MicrosoftIISv2")
' Get a reference to the ASP application called Default Web Site
set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
WScript.Echo "Before: AspExecuteInMTA = " & IIsWebVirtualDirSettingObj.AspExecuteInMTA
' Set the ASP application to execute in MTA
IIsWebVirtualDirSettingObj.AspExecuteInMTA = 1
IIsWebVirtualDirSettingObj.Put_()
WScript.Echo "After: AspExecuteInMTA = " & IIsWebVirtualDirSettingObj.AspExecuteInMTA
```

**Side-by-side assemblies**

Side-by-side (SxS, in metabase property names) assemblies allow ASP applications to specify which version of a system DLL or shared COM component to use. For example, if your ASP application relies on MSXML version 2.0, you can ensure that your application still uses MSXML 2.0 even if a later version is installed on the computer for use by other applications.

Configuring SxS assemblies requires that you know the path to the DLL, and that the COM+ manifest file exists in every virtual directory that requires the use of the DLL. The COM+ manifest is an XML file that contains information about where a DLL is installed. IIS does not verify that the manifest exists. A manifest looks like the following file excerpt:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<assembly xmlns="urn:schemas-microsoft-com:asm.v1" manifestVersion="1.0">
<assemblyIdentity publicKeyToken="75e377300ab7b886" type="win32" name="Test4Dir"
version="1.0.0.0" processorArchitecture="x86"/>
<file name="DirComp.dll" hash="35ca6f27b11ed948ac6e50b75566355f0991d5d9"
hashalg="SHA1">
<comClass clsid="{6C6CC20E-0F85-49C0-A14D-D09102BD7CDC}" progid="DirComp.PathInfo"
threadingModel="apartment"/>
<typelib tlbid="{AA56D6B8-9ADB-415D-9E10-16DD68447319}" version="1.0" helpdir=""/>
</file>
</assembly>
```

You can enable side-by-side assemblies either programmatically or by using IIS Manager. To programmatically enable side-by-side assemblies, set the **AspEnableSxs** flag of the **AspAppServiceFlags** metabase property to **true**. Also, set the **AspSxsName** metabase property to the name of the COM+ manifest. Set both metabase properties at the application level.

| **Important** |
| --- |
| Even though this feature is configurable at the application level, you must not use more than one version of a system DLL in any IIS6.0 application pool. If you load more than one version of a system DLL in an application pool, the application that is loaded first has its version loaded and the other applications are forced to use that DLL until the applications are unloaded. |

The following example enables side-by-side assemblies on the Default Web Site application (*W3SVC/1/ROOT*). Notice that after setting only the **AspEnableTracker** property, the **AspAppServiceFlags** property changes as well:

```
On Error Resume Next
set providerObj = GetObject("winmgmts://MyMachine/root/MicrosoftIISv2")
' Get a reference to the ASP application called Default Web Site
set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
```

```
WScript.Echo "Before: AspEnableSxs = " & IIsWebVirtualDirSettingObj.AspEnableSxs
WScript.Echo "        AspSxsName = " & IIsWebVirtualDirSettingObj.AspSxsName
WScript.Echo "        AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
' Set the ASP application to enable COM+ side-by-side assemblies
IIsWebVirtualDirSettingObj.AspEnableSxs = 1
' Set the AspSxsName property
IIsWebVirtualDirSettingObj.AspSxsName = "VersionInfo"
' Save the values to the IIS metabase
IIsWebVirtualDirSettingObj.Put_()
' Get the reference again in order to refresh the AspAppServiceFlags property.
set IIsWebVirtualDirSettingObj = Nothing
set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
WScript.Echo "After: AspEnableSxs = " & IIsWebVirtualDirSettingObj.AspEnableSxs
WScript.Echo "       AspSxsName = " & IIsWebVirtualDirSettingObj.AspSxsName
WScript.Echo "       AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
```

**COM+ partitions**

You can use COM+ partitions to isolate Web applications in their own COM+ partitions. Isolation is useful to prevent one Web application from accessing the private COM+ applications, configuration information, and data of another Web application.

COM+ partitions can hold different versions of your own custom COM components. For example, if you host Web sites for two competing companies that both use COM+ in their Web applications, you can use COM+ partitions to ensure that one company's Web application cannot access the COM+ components in the other company's Web applications. If one of those companies asks you to change certain features in a COM+ application that both companies use, you can isolate the new version of that COM+ application in the partition that is linked to that company's Web application.

To enable COM+ partitions in IIS, set the **AspUsePartition** flag of the **AspAppServiceFlags** metabase property at the application level. The partition is identified by a globally unique identifier (GUID) (created by using the Component Services Manager snap-in), which is set by the **AspPartitionID** metabase property. If no partition is specified, the default system partition is used. For more information, please see "Creating and Configuring COM+ Partitions" in the COM+ SDK or online at COM+ (Component Services).

| **Important** |
| --- |
| Even though you can configure this feature at the application level, use only one version of a COM+ component in any application pool. For example, if application App1 uses version1.0 of a custom COM+ application called Shop.dll, and application App2 uses version2.0 of Shop.dll, App1 and App2 cannot be in the same application pool. If you put them in the same application pool, the application that you load first has its version of Shop.dll loaded, and the other application is forced to use that version until the applications are unloaded. |

The following example enables partitions on the Default Web Site application (*W3SVC/1/ROOT*). Notice that after setting only the **AspEnableTracker** property, the **AspAppServiceFlags** property changes as well.

```
On Error Resume Next
set providerObj = GetObject("winmgmts://MyMachine/root/MicrosoftIISv2")
' Get a reference to the ASP application called Default Web Site
set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
WScript.Echo "Before: AspUsePartition = " & IIsWebVirtualDirSettingObj.AspUsePartition
WScript.Echo "        AspPartitionID = " & IIsWebVirtualDirSettingObj.AspPartitionID
WScript.Echo "        AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
' Set the ASP application to enable COM+ partitioning
IIsWebVirtualDirSettingObj.AspUsePartition = 1
' Set the AspPartitionID property to the GUID configured in Component Services Manager
' when you created the COM+ partition
IIsWebVirtualDirSettingObj.AspPartitionID = "{00000000-0000-0000-0000-000000000000}"
' Save the values to the IIS metabase
IIsWebVirtualDirSettingObj.Put_()
' Get the reference again in order to refresh the AspAppServiceFlags property.
set IIsWebVirtualDirSettingObj = Nothing
set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
WScript.Echo "After: AspUsePartition = " & IIsWebVirtualDirSettingObj.AspUsePartition
WScript.Echo "       AspPartitionID = " & IIsWebVirtualDirSettingObj.AspPartitionID
WScript.Echo "       AspAppServiceFlags = " &
```

```
IIsWebVirtualDirSettingObj.AspAppServiceFlags
```

**COM+ tracker**

COM+ tracker allows you to debug ASP applications. For example, if you need to diagnose and fix a faulty Web application, you can enable COM+ tracker to determine when ASP pages are being loaded, when COM components are loaded, and when threads leave a page. After debugging your application, you can disable COM+ tracker to return your application to normal performance speed.

To enable COM+ tracker on the IIS side, set the **AspEnableTracker** flag of the **AspAppServiceFlags** metabase property at the application level.

The following example enables tracking on the Default Web Site application (*W3SVC/1/ROOT*). Notice that after setting only the **AspEnableTracker** property, the **AspAppServiceFlags** property changes as well:

```
On Error Resume Next
set providerObj = GetObject("winmgmts://MyMachine/root/MicrosoftIISv2")
' Get a reference to the ASP application called Default Web Site
set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
WScript.Echo "Before: AspEnableTracker = " &
IIsWebVirtualDirSettingObj.AspEnableTracker
WScript.Echo "        AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
' Set the ASP application to enable COM+ tracking
IIsWebVirtualDirSettingObj.AspEnableTracker = 1
IIsWebVirtualDirSettingObj.Put_()
' Get the reference again in order to refresh the AspAppServiceFlags property.
set IIsWebVirtualDirSettingObj = Nothing
set IIsWebVirtualDirSettingObj =
providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'")
WScript.Echo "After: AspEnableTracker = " & IIsWebVirtualDirSettingObj.AspEnableTracker
WScript.Echo "       AspAppServiceFlags = " &
IIsWebVirtualDirSettingObj.AspAppServiceFlags
```

For more information about these enhancements to COM+ services in ASP, including application examples of each, see What's New in ASP.

**New Metabase Properties for ASP**

Table 3.9 describes the metabase properties for ASP that are new in IIS 6.0.

**Table 3.9 New ASP Metabase Properties in IIS 6.0**

| Property | Purpose |
|---|---|
| **AspAppServiceFlags** | Contains the flags that must be set to enable COM+ services on your IIS applications. |
| **AspDiskTemplateCacheDirectory** | Specifies the location of the ASP disk cache. |
| **AspExecuteInMTA** | Enables ASP threads to execute in an MTA. |
| **AspKeepSessionIDSecure** | Sends the ASP session cookie to a browser securely. |
| **AspMaxDiskTemplateCacheFiles** | Specifies the maximum number of compiled ASP templates that can be stored. |
| **AspPartitionID** | Sets the COM+ partition to use for the application. |
| **AspRunOnEndAnonymously** | Enables ASP to run the **Global.asa Application_OnEnd** and **Session_OnEnd** events anonymously. |
| **AspBufferingLimit** | Specifies the buffer size limit. |
| **AspMaxRequestEntityAllowed** | Specifies the maximum number of bytes allowed in the entity body of an ASP request. |
| **AspSxsName** | Allows ASP applications to specify which version of a system DLL or COM component to use. |

For more information about how to use these new metabase properties, see the Metabase Property Reference.

**Related Information**

- For more information about using IIS Manager to enable side-by-side assemblies, see Enabling Side-by-Side Assemblies

using IIS Manager.

- For more information about configuring launch permissions for distributed Web applications, see Configuring Launch Permissions for COM+ Server Objects.

# Security and Performance Enhancements in ASP

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 6.0, ASP has undergone several important changes to improve security and performance:

- Default disabling of ASP pages during clean installations of IIS 6.0. In a default installation, the IIS service is installed in a secure mode, and only serves static content. You must enable features such as ASP, ASP.NET, CGI, ISAPI, and WebDAV if you need them. For more information about these features, including how to enable them, see Enabling Dynamic Content.

- Secure ASP built-in functions. All ASP built-in functions always run as the low-privileged account, IUSR_*ComputerName*, or under an authenticated user account if one is selected and has valid credentials.

- Default disabling of parent paths. As a security precaution, the metabase property **AspEnableParentPaths** is now set to zero by default. This new default setting affects Web pages that contain or use the .. notation to refer to a parent directory in the #include SSI directive. Unless you explicitly set this property to **true**, such Web pages do not respond to a client request (they instead generate a Disallowed Parent Path message).

- ASP hang detection. When an IIS Web site is busy, there might be instances when the maximum number of ASP threads are spawned and some of the ASP threads hang, resulting in degraded performance. (Threads are considered to be hung if they do not respond to a timeout.) If sufficient ASP requests hang so that ASP cannot service requests quickly, ASP detects this unhealthy state and requests that the worker process recycle itself.

- Monitoring of disk-based caching. The Active Server Pages performance object now provides the Templates Cached counter to monitor disk-based caching in addition to in-memory caching. This counter counts the number of cached ASP files, both on disk and in memory. A new ASP counter, In Memory Templates Cached, reports the number of ASP pages in the memory cache. The default location of the ASP disk cache is *systemroot*\System32\Inetsrv\ASP Compiled Templates. You can change the cache location by setting the **AspDiskTemplateCacheDirectory** metabase property.

- Anonymous user access for Global.asa events. Earlier versions of ASP executed events in the security context (or user identity) of the host process because during these events there was not necessarily a user context. This sometimes caused access denied errors when writing to a file in the Application_OnEnd or Session_OnEnd event. By default, ASP now runs the Global.asa events, Application_OnEnd and Session_OnEnd, anonymously. Setting the **AspRunOnEndAnonymously** metabase property to **true** allows these global ASP functions to run as anonymous user.

- Performance enhancements. To limit the amount of memory allocated to ASP pages, IIS 6.0 sets the default value of the **AspScriptFileCacheSize** metabase property to 500 ASP pages and the default value of the **AspScriptEngineCacheMax** metabase property to 250 script engines. For sites that include a large set of frequently requested ASP pages, you can increase the number of ASP pages allowed by the **AspScriptFileCacheSize** property, which improves performance because ASP page compilation is substantially slower than retrieving pages from cache. On a site with only a small number of frequently requested ASP pages, you can save memory by setting a smaller value.

ASP Metabase Properties No Longer in Use in IIS 6.0

The following metabase properties are no longer used by ASP, even though they are still in the metabase for compatibility with existing administration scripts:

- **AspAllowOutOfProcComponents**

- **AspTrackThreadingModel**

- **AspExceptionCatchEnable**

In IIS 6.0, all properties that relate to ASP thread gating, such as the **ASPThreadGatingModel** property, were removed from

ASP and the metabase. If you are using scripts to upgrade from IIS 5.0, you might need to make manual adjustments for the properties that were removed.

# Setting Application Mappings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because you can choose from many programming and scripting languages when you develop Web applications, IIS uses the file name extension of a requested resource on your Web site to determine which ISAPI or CGI program to run to process a request. For example, a request for a file with an .aspx extension causes the Web server to call the ASP.NET ISAPI (Aspnet_isapi.dll) to process the request. The association of a file name extension with an ISAPI or CGI program is called *application mapping* (or sometimes *script mapping* or *application extension mapping*).

IIS 6.0 is preconfigured to support common application mappings. You can add or remove mappings for all of the applications on a Web site or for an individual application. Use IIS Manager to map a file name extension to an application or remove an application mapping. For more information about setting application mappings, see Setting Application Mappings in IIS 6.0.

Benefits of Using Wildcard Application Mappings Instead of ISAPI Filters

You can configure a Web site or virtual directory to run any number of ISAPI applications at the beginning of every request to that Web site or virtual directory, regardless of the file name extension of the requested file. This type of configuration, which can only be an ISAPI application, is called *wildcard application mapping*.

When an IIS Web server receives a request from a client, the Web server looks at the file name extension for the requested file to determine whether an ISAPI or CGI application handles that file type. For example, if a request comes in for a Web page called Default.asp, IIS knows from its default application mappings that ASP files are handled by the Asp.dll ISAPI extension.

Prior to wildcard application mapping, if you wanted certain tasks to be executed for every client request before the requested page was sent to its mapped application, you had to write an ISAPI filter that intercepted requests as they arrived. Wildcard application mapping offers an alternative to ISAPI filters. Although ISAPI filters are excellent for the functions that they perform, they have a more narrow use than do ISAPI applications.

For example, when you use wildcard application mapping, your ISAPI application:

- Can easily access the entity body of a request.
- Can be configured for virtual directories.
- Can contain long-running operations.
- Offers greater flexibility because the ISAPI extension API, which is used to write ISAPI applications, has more functionality than does the ISAPI filter API.

For more information about ISAPI filters, see Installing ISAPI Filters.

Request Handling with Wildcard Application Mappings

Highly secure Web applications, such as banking applications, require that every user who makes a request be authenticated against a database of user names, passwords, and account numbers. After the user is authenticated, certain rules are applied that specify what information the user can access. If the wildcard application map determines that the user has no rights, it can prevent the request from going any farther, or it can redirect the request to a new Web page.

You can implement a custom authentication and authorization scheme like this in an ISAPI application that you install as a wildcard application mapping, and all users must go through authentication regardless of the file type of their requests. For example, when you install wildcard application mappings, IIS processes a client request in this way:

1. An incoming request is handled by any existing ISAPI filters that are installed at the global and Web site levels. IIS executes the ISAPI filters in priority order from global to Web site.

- The request is sent to any existing wildcard application mappings, which IIS also executes in order.
    - If the incoming request for a Web file triggers wildcard application mappings at multiple levels — such as Web site, virtual directory, or nested virtual directories — IIS only executes the wildcard application mappings installed in the virtual directory closest to the Web file. That is, if wildcard application mappings exist at the Web site and

virtual directory level, a request for a Web file in the virtual directory only triggers the wildcard application mappings for the virtual directory.

- If no custom application mappings are installed in a virtual directory, that virtual directory inherits the effective application mappings, including wildcard application mappings, from its parent.

- The IIS server looks at the application mappings table to determine which ISAPI application or CGI script specifically processes the requested resource.

For information about developing ISAPI extensions, see Creating ISAPI Extensions. For information about developing ISAPI filters, see Creating ISAPI Filters.

For more information about application mapping in general, see Setting Application Mappings in IIS 6.0 and Installing Wildcard Application Mappings.

# Additional Resources for Running IIS 6.0 as an Application Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this section.

Related Information

- Configuring Internet Sites and Services for information about configuring the File Transfer Protocol (FTP), Network News Transfer Protocol (NNTP), and Simple Mail Transfer Protocol (SMTP) services.
- IIS 6.0 Architecture for information about the IIS 6.0 application isolation modes.
- Using FrontPage 2002 Server Extensions with IIS 6.0 for information about installing and using FrontPage 2002 Server Extensions.
- Deploying ASP.NET Applications in IIS 6.0 in *Deploying Internet Information Services (IIS) 6.0* of the *Microsoft® Windows® Server 2003 Deployment Kit* for information about how you can recycle a Web garden based on length of time, including examples of Web garden recycling.
- ServerSupportFunction Extension Function on the MSDN Web site for information about the **ServerSupportFunction** callback function.

Related IIS 6.0 Operations Guide Topics

- Recycling Worker Processes with IIS 6.0 for information about worker process isolation mode features.
- IIS and Built-in Accounts and Configuring Application Pool Identity with IIS 6.0 for security-related information that pertains to running your application server.
- Installing IIS and Enabling and Disabling Dynamic Content in IIS 6.0 for information about installing IIS 6.0 and enabling the features and services necessary for your applications to run properly.
- Installing ISAPI Filters for information about ISAPI filters.
- The Metabase Property Reference and Code Examples to Configure Metabase Properties for information about configuring metabase properties.
- For information about performing demand stops and starts on an application pool, see IIsApplicationPool (ADSI) and IIsApplicationPool (WMI) in the IIS Programmatic Administration Reference.
- Setting Application Mappings in IIS 6.0 for information about setting application mappings, including step-by-step instructions.
- What's New in ASP and Overview of ASP.NET for information about using ASP and ASP.NET with IIS 6.0.

Related Windows Server 2003 Help Topics

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the**Search in title only** check box.

- Features unavailable on 64-bit versions of the Windows Server 2003 family in Help and Support Center for Windows Server 2003 for information about features, such as ASP.NET, that are not supported on 64-bit versions of Windows Server 2003.

- Reformatting or converting a partition to use NTFS in Help and Support Center for Windows Server 2003 for information about converting FAT or FAT32 volumes to NTFS.

Related Tools

- Adsutil.vbs

You can use the Adsutil.vbs command-line tool, which is an unsupported IIS administrative utility that is provided as a learning tool, to configure your application server. For more information about this tool, including a description of its syntax, commands, and parameters, see Using the Adsutil.vbs Administration Script.

# Working with the Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The metabase is a hierarchical store of configuration information and schema that are used to configure Internet Information Services (IIS) 6.0. Read this section for an overview of the metabase structure and terminology, as well as detailed information about effectively managing the IIS 6.0 metabase by using the graphical user interface (GUI)–based and command-line utilities that are available with IIS 6.0.

**In This Section**


Overview of the IIS 6.0 Metabase

Metabase Structure in IIS 6.0

Metabase Security and Auditing

Backing Up and Restoring the Metabase

Editing the Metabase

Metabase Import and Export

Additional Resources for the IIS 6.0 Metabase

# Overview of the IIS 6.0 Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Internet Information Services (IIS) 6.0 metabase comprises two plaintext XML files that store IIS configuration values and schema. The metabase can be edited manually, by using IIS Manager, or it can be edited programmatically. In addition, the IIS 6.0 metabase is extensible in a highly efficient manner — as your IIS deployment grows, so does the metabase. By using an inheritance model, explicit declarations of duplicate values are avoided, which reduces overhead when IIS reads configuration values from the metabase.

The metabase is a hierarchical store of configuration information and schema that are used to configure IIS. The metabase configuration and schema for IIS 4.0 and IIS 5.0 were stored in a single binary file named MetaBase.bin, which was not easy to read or edit. IIS 6.0 replaces MetaBase.bin with two XML files named MetaBase.xml and MBSchema.xml. These files are stored on your computer in the *systemroot*\System32\Inetsrv folder. Only users who are members of the Administrators group can view and modify these files.

The metabase configuration file, MetaBase.xml, stores most of the IIS configuration. Some IIS configuration values are also stored in the Windows registry. If a configuration value is one that you might need to configure or change, or if you can access the setting in the IIS 6.0 user interface, then the setting is typically stored in the IIS metabase.

The MetaBase.xml file is organized in a hierarchical structure, which can vary depending on the choices that are made when IIS is installed and, if needed, reconfigured. When IIS is started or restarted, the configuration settings are read from MetaBase.xml and copied into the IIS cache in-memory, which is referred to as the *in-memory metabase*. When IIS Manager or programmatic interfaces are used to change IIS configuration, the changes are made to the in-memory metabase and persisted to MetaBase.xml after a previously configured number of changes, or according to regular time intervals.

In physical terms, the metabase is a combination of the MetaBase.xml and MBSchema.xml files and the in-memory metabase. The IIS configuration information is stored in the MetaBase.xml file, while the metabase schema is stored in the MBSchema.xml file. When IIS is started, these files are read by the storage layer and then written to the in-memory metabase through Admin Base Objects (ABOs), as shown in Figure 4.1.




The metabase configuration and schema are stored in separate nodes in the in-memory metabase. The in-memory metabase itself resides in the IIS file cache in your computer's RAM. While IIS is running, changes that are made to the in-memory metabase are periodically saved to disk. The in-memory metabase is also saved to disk when IIS is stopped.

# IIS 6.0 Metabase Features

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS 6.0 metabase includes the features described below.

**The metabase schema can be modified by an administrator**
Active Directory® Service Interfaces (ADSI) and Windows Management Instrumentation (WMI) use the metabase schema to enforce which properties can be written at a particular key in the metabase. The schema also enforces the data types that can be used for a particular property attribute. You can modify the metabase schema by using ADSI or WMI to customize this enforcement, or to allow out-of -schema properties to be written in the metabase. For information about extending the metabase schema, see The Metabase Schema.

**The metabase is 100 percent compatible with IIS 5.0 metabase APIs**
Because the XML metabase in IIS 6.0 is fully compatible with IIS 5.0 metabase application programming interfaces (APIs), existing scripts and code will continue to work.

**Metabase performance and scalability are improved**
The IIS 6.0 metabase improves performance and scalability over earlier versions of IIS in the following ways:

- Comparable or better disk footprint size compared to the IIS 5.0 metabase
- Faster read times on Web server startup than the IIS 5.0 binary metabase
- Equivalent write performance to the IIS 5.0 binary metabase

**The metabase incorporates rich functionality**
The metabase incorporates rich functionality with the following features:

- Property inheritance
- Data typing of property values
- Change notification when metabase properties change
- Security
- Simplified backup and restore of your IIS configuration
- Easier import and export of individual nodes to other computers

**Benefits of Plaintext XML Files**
Plaintext XML metabase files offer the benefits described below.

**Metabase files can be edited directly by using common text editors**
The metabase configuration file, MetaBase.xml, can be easily read and edited by using common text editors, such as Notepad, which is provided as part of the Microsoft® Windows® operating system. Editing the metabase configuration file directly will be of interest primarily for administrators who do not want to use scripts or code to administer IIS. In addition, editing the metabase configuration file directly is faster than using IIS Manager when administering IIS over a slow network connection. It is recommended that you become familiar with the structure of MetaBase.xml before you attempt to configure IIS by editing the MetaBase.xml file.

| **Important** |
| --- |
| Editing the MBSchema.xml file directly is not supported. Changes made to the metabase schema must be made programmatically by using the IIS ADSI provider. |

**Improved metabase corruption recovery and troubleshooting**

The metabase makes it easier to diagnose potential metabase corruption because the metabase is stored in a plaintext file, and it can be analyzed by using tools such as the WinDiff utility. The WinDiff utility is located in the Deploy.cab file in the \Support\Tools folder of your Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; or Windows® Server 2003, Web Edition operating system CD. As the metabase is edited and saved to disk, IIS makes copies of each version of the MetaBase.xml and MBSchema.xml files in the History folder. In other words, the metabase history feature is like an automatic backup mechanism. You can use history files to compare changes that have been made to the MetaBase.xml and MBSchema.xml files.

**Improved metabase backup and restore capabilities**

By using Backup/Restore Configuration, administrators can back up the metabase and encrypt it with any password. If a critical failure occurs, the metabase can be restored on a different computer or on a different installation of the operating system. Additionally, the metabase can be restored by using an earlier version of the MetaBase.xml and MBSchema.xml files from the History folder.

# Metabase Terminology

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

MetaBase.xml is structured by using keys. *Keys* are analogous to file folders in a file system or registry keys in the registry. Each key contains a **Location** attribute, which specifies its hierarchical location in the metabase. Properties, which define the values for configuring IIS, are contained in keys. Because MetaBase.xml is structured hierarchically, a particular property can be configured differently at different keys. Properties contain attributes, such as the **Value** attribute.

Table 4.1 summarizes the metabase terminology that is used to describe the structure of the IIS metabase.

**Table 4.1 Metabase Terminology**

| Term | Description |
|---|---|
| Key | A hierarchical container in MetaBase.xml that is used to contain properties. Each key contains the **Location** attribute within the start tag of the key, specifying the hierarchical position of the key in MetaBase.xml. A metabase key is analogous to a file folder or registry key. |
| Location | Each key contains a **Location** attribute — a sequence of location names separated by one or more forward slashes (/) — that uniquely identifies the hierarchical position of the key in MetaBase.xml. For example /LM is the location of the IIsComputer key, and /LM/W3SVC is the location of the IIsWebService key. The latter is a child key of the LM location. A child key is a key that is contained by another key. |
| Property or Custom Property | A property or custom property defines values that are used to configure IIS. Properties contain attributes; some properties, such as **AccessFlags**, also contain flags. Properties are contained by keys and can be written, in MetaBase.xml, as schema-defined or custom. |
| Attributes | Metabase properties contain several attributes that describe some characteristics of the property. Keys also contain attributes, such as the **Location** attribute or the **DefaultDoc** attribute found within the IIsWebService key.<br>Do not confuse the term *attributes* as used here to describe IIS metabase properties with XML attributes, which describe general XML structure. |
| Flag | A flag is a bit within the DWORD value of a property that contains flags. Each flag within a property is used to configure particular functionality of the property, such as read or write access. The sum of the flag values for a property becomes the sum of the value attribute for that property. |

**Keys and Locations**

Metabase keys are analogous to registry keys or a directory in a file system. Each metabase key is identified by its name and location, which are defined within the start tag of the key. The following example shows two keys, IIsConfigObject and IIsLogModules, each having a unique location as defined by the value of its respective **Location** attribute:

Most key names are the same as the value of the **KeyType** property, also known as the ABO class name, which relates to the type of key. The IIsLogModules key, seen in the example above, is named this way. If you change the **KeyType** value of a metabase key programmatically, the name of the key changes in MetaBase.xml when the in-memory metabase is saved. This **KeyType** value is important because the schema-defined collection that is named by using the same **KeyType** enforces which properties the key can contain.

Keys named IIsConfigObject are special because they are not named by using a **KeyType** value. Because all properties are schema-defined in the IIsConfigObjectcollection, keys in MetaBase.xml that are named IIsConfigObject can contain any property. You can specify a **KeyType** value for a key named IIsConfigObject by writing the **KeyType** property in the key as a custom property. This is not a recommended method, however, because it does not work with ADSI. Preferably, use a real object, not IIsConfigObject.

Some key names describe the type of data at a unique location in the metabase, and thus might appear at multiple locations. For example, because you can have more than one virtual directory, IIsWebVirtualDir is a key that can appear at multiple locations. The **Location** attribute defines the hierarchical position of a key in MetaBase.xml. The value of the **Location** attribute, with only two exceptions, is a sequence of location names separated by one or more forward slashes (/) that uniquely identifies the hierarchical position of the key in MetaBase.xml.

The exceptions are the IIS_Global key, which has the highest hierarchical position in MetaBase.xml (**Location** = .), and the IIS_ROOT key, which has the next-highest hierarchical position in MetaBase.xml (**Location** = /). The next-highest hierarchical position is the IIsComputer key (Location = /LM). All remaining keys in MetaBase.xml are contained by the IIsComputer key; therefore, their location always begins with /LM/*LocationName*, where *LocationName* is the name of the location.

The hierarchical position of any metabase key can be determined by looking at its **Location** attribute. For example. /LM is the location of the IIsComputer key, and /LM/W3SVC is the location of the IIsWebService key. Because the /LM/W3SVC location begins with /LM followed by a forward slash character (/) and another location name (W3SVC), the IIsWebService key at location /LM/W3SVC is a child of the IIsComputer key. A *child key* is a key that is contained by another key. A key that contains another key is also called a *parent key*.

# XML Terminology Related to IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because MetaBase.xml is an XML file constructed with XML elements, it helps to be familiar with XML terminology and metabase terminology to understand how MetaBase.xml is structured. Table 4.2 identifies and defines the XML terms and relates them to their use in MetaBase.xml.

**Table 4.2 XML Terminology and How It Relates to the IIS Metabase**

| Term | XML Description | Relation to MetaBase.xml |
|---|---|---|
| XML declaration | A processing instruction at the start of an XML document, such as `<?xml version = 1.0?>`, that declares the file to be XML code. | The first line in MetaBase.xml is an XML declaration. |
| Element | A logical unit of information. Elements are surrounded by <*ElementName*> and </*ElementName*> tags, where *ElementName* is the name of the element. Elements can contain attributes. | Keys and custom properties are written by using XML elements. |
| Root element | A single element that contains all the other elements and character data that make up an XML document; also known as a *document element*. | MetaBase.xml contains a root element named <configuration>. |
| Start tag | A tag that marks the start of an element, such as <*ElementName*>, where *ElementName* is the name of the element. | In MetaBase.xml, start tags are used to identify the start of a key or custom property. |
| End tag | A tag that marks the end of an element, such as </*ElementName*>, where *ElementName* is the name of the element. | In MetaBase.xml, end tags are used to identify the end of a key or custom property. |
| Element type | Indicated by the name that occurs in its start tag and end tag. | In MetaBase.xml, the element type for metabase keys is actually the value of the **KeyType** property. For example, IIsWebService is the **KeyType** that is used to identify the World Wide Web Publishing Service (WWW service) key. Custom properties use the name **Custom** as the element type. |
| Attribute | A name and value pair that is associated with an XML element and that provides more information about the content of that element.<br><br>Do not confuse the term **attribute** as used here to describe XML structure with an IIS metabase property attribute that describes an IIS property. | A metabase property contains several attributes that describe the characteristics of the property, such as the data type or value. Keys also contain attributes, such as the **Location** attribute or the **DefaultDoc** attribute found within the IIsWebService key. |
| Comment | Text that is not to be treated as part of the document. Comments are delimited by <!-- and --> tags. | In MetaBase.xml, comments are delimited by <!-- and --> tags, and must be written above the key to which they pertain. If you put a comment inside an XML node, it is removed by IIS. In addition, no comments are allowed above the IIS_Global node. |

# Metabase Structure in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The XML declaration `<?xml version = "1.0"?>` occupies the first line in MetaBase.xml and indicates that the file is formatted as XML. The root element, named <configuration>, contains all other elements.

The IIS_Global key is the highest-level key in the metabase that contains properties. Consider the IIS_Global key to be read-only; do not configure its properties.

The IIS_ROOT key (Location = /) is the highest-level key in which you can configure properties. Properties that are configured at the IIS_ROOT key affect the general operation of IIS.

The IIsComputer key (Location = /LM) contains all remaining metabase locations. Table 4.3 lists, in hierarchical order, child locations that can exist within the namespace of the /LM location. The File Transfer Protocol (FTP), Network News Transfer Protocol (NNTP), and Simple Mail Transfer Protocol (SMTP) components are not installed with IIS by default. Therefore, the /LM/MSFTPSVC, /LM/NNTPSVC, and /LM/SmtpSvc locations do not exist in MetaBase.xml unless their respective components have been installed.

**The Local Machine Namespace**

The local machine (LM) namespace is the parent location in the hierarchical structure of keys where all services and sites are organized. These keys, each of which contains their own unique location, are organized in the following format:

/LM/*Service*/*Site*/*ROOT*/*VirtualDirectory*/*Directory*/*File*

Where the replaceable parameters are:

- *Service* = MSFTPSVC, NNTPSVC, SmtpSvc, or W3SVC
- *Site* = unique number identifier for the FTP, NNTP, SMTP, or Web, site
- *ROOT* = root virtual directory of the site
- *VirtualDirectory* = virtual directory
- *Directory* = physical directory
- *File* = file name

Each site is a server instance, and is referred to by the number used in its namespace. For example, /LM/W3SVC/1 specifies the location of the key that contains the first Web site, and /LM/MSFTPSVC/3 specifies the location of the key that contains the third FTP site.

| **Important** |
| --- |
| Beginning with IIS6.0, the default Web site is numbered 1, and when a new Web site is created, a Web site identification number is randomly generated based on the name of the Web site. You can enable sequential numbering of Web sites, which was the method of generating Web site identification numbers in IIS5.1 and earlier, by setting the IncrementalSiteIDCreation registry key to 1. |

Table 4.3 lists the locations of the parent keys that are contained in the LM namespace. Because their child keys are too numerous, they are not listed here.

**Table 4.3 Local Machine Locations**

| Location | Keys at this location |
| --- | --- |
| /LM/EventManager | Keys at the /LM/EventManager location and its child locations configure Microsoft Exchange Server. You should not configure keys at the /LM/EventManager location or its child keys by editing the MetaBase.xml file. |

| | |
|---|---|
| /LM/IISADMIN | Keys at the /LM/IISADMIN location and its child locations record DCOM extensions to IIS. Because the /LM/IISADMIN location is internally configured by IIS, you should not configure keys at /LM/IISADMIN or its child locations. |
| /LM/Logging | Keys at the /LM/Logging location and its child locations configure IIS logging. |
| /LM/MimeMap | Keys at the /LM/MimeMap location configure the **MimeMap** property. |
| /LM/MSFTPSVC | Keys at the /LM/MSFTPSVC location and its child locations configure the FTP service and sites. |
| /LM/NNTPSVC | Keys at the /LM/NNTPSVC location and its child locations configure the NNTP service and sites. |
| /LM/SmtpSvc | Keys at the /LM/SmtpSvc location and its child locations configure the SMTP service and sites. |
| /LM/W3SVC | Keys at the /LM/W3SVC location configure properties that are global to the WWW service. |
| /LM/W3SVC/1 | Keys at the /LM/W3SVC/1 location and its child locations configure properties that are specific to the first Web site. |
| /LM/W3SVC/AppPools | Keys at the /LM/W3SVC/AppPools location configure properties that are global to all application pools. Keys at child locations, such as /LM/W3SVC/AppPools/DefaultAppPool, configure properties that are specific to a particular application pool. |
| /LM/W3SVC/Filters | Keys at the /LM/W3SVC/Filters location and its child locations configure Internet Server API (ISAPI) filters, and global compression schemes. |
| /LM/W3SVC/Info | Keys at the /LM/W3SVC/Info location configure properties that are global to all Web sites. |
| /LM/W3SVC/Info/Templates | The /LM/W3SVC/Info/Templates location and its child locations were used by the IIS Permissions Wizard in earlier versions of IIS. It remains solely for legacy compatibility. You should not configure keys at the /LM/W3SVC/Info/Templates location by editing the MetaBase.xml file. |

# Property Inheritance in the IIS 6.0 Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Property inheritance configures IIS by using as few property settings as possible. The benefits of using a minimal number of property settings to configure IIS are as follows:

- Faster IIS cache performance
- Reduced memory consumption
- Reduced time required for IIS administration

Metabase properties are used to provide detailed control of IIS functionality. All metabase properties contain attributes, such as the **DefaultValue** attribute, which specify the configuration of the properties. Most metabase properties are inheritable, and they are specified as inheritable in MBSchema.xml by using the string value INHERIT as part of the value of the **Attributes** attribute. The following example shows the default attributes for the **HttpErrors** property in MBSchema.xml:

```
<Property
   InternalName = "HttpErrors"
   ID = "6008"
   Type = "MULTISZ"
   UserType = "IIS_MD_UT_FILE"
   Attributes = "INHERIT"
   MetaFlagsEx = "CACHE_PROPERTY_MODIFIED"
   DefaultValue = ""/
>
```

Hierarchical Relationship of Keys

The metabase is designed so that properties can be set differently at different keys. To understand how property inheritance works, it is necessary to first understand the hierarchical relationship of keys in the MetaBase.xml file. In relation to other keys, a key can be located in MetaBase.xml at a higher level, lower level, or the same level. The **Location** attribute of each key determines the level in which a key exists in the hierarchical structure.

The following example illustrates the relationship between the IIsWebService, IIsWebServer, and IIsWebVirtualDir keys. The example uses indentation to illustrate the hierarchical relationship of the keys.

```
<IIsWebService Location = "/LM/W3SVC"
   <IIsWebServer Location = "/LM/W3SVC/1">
      <IIsWebVirtualDir Location = "/LM/W3SVC/1/ROOT">
   <IIsWebServer Location = "/LM/W3SVC/2">
      <IIsWebVirtualDir Location = "/LM/W3SVC/2/ROOT">
```

In this example, the IIsWebService key is the highest-level key. The IIsWebServer keys are at a lower level than IIsWebService, and the IIsWebVirtualDir keys are at a lower level than the IIsWebServer keys. Keys that are at the same level cannot inherit properties from each other. Therefore, because the IIsWebServer keys are at the same hierarchical level in the metabase, they cannot inherit properties from each other. The IIsWebVirtualDir keys are also at the same level in relation to each other, and therefore they cannot inherit properties from each other.

| **Note** |
| --- |
| The hierarchical structure of the metabase is determined by the **Location** attribute of the keys themselves, not the class name of the keys. |

Inheritance

Inheritable properties that are not explicitly set at a lower-level key are automatically inherited from the key at the next-higher level in which the property is configured. The IIS_ROOT key is the highest key in MetaBase.xml from which a property can be inherited.

The following example illustrates properties that are configured differently in different keys in MetaBase.xml, as well as properties that are inherited.

```
<IIsWebService Location = "/LM/W3SVC"
<!-- The DefaultDoc and ContentIndexed properties are specifically set at this key. -->
= "Default.htm,Default.asp,index.htm,Default.aspx"
ContentIndexed = "TRUE"
   <IIsWebServer Location = "/LM/W3SVC/1">
   <!-- The DefaultDoc property is specifically set at this key. -->
   = "Default.htm,Default.asp,index.htm,iisstart.htm,Default.aspx"
   <!-- The ContentIndexed property is inherited from the IIsWebService key and
   does not appear here -->
   </IIsWebServer>

<IIsWebVirtualDir Location = "/LM/W3SVC/1/ROOT"
<!-- The DefaultDoc and ContentIndexed properties are inherited from the /LM/W3SVC/1
location and do not appear here. -->
</IIsWebVirtualDir>
>
   <IIsWebServer Location = "/LM/W3SVC/2"
   <!-- The DefaultDoc and ContentIndexed properties are inherited from the
   IIsWebService key. -->
   </IIsWebServer>
   >
   <IIsWebVirtualDir Location = "/LM/W3SVC/2/ROOT"
   <!-- The DefaultDoc property is inherited from the /LM/W3SVC/2 location. -->
   <!-- The ContentIndexed property is specifically set at this key. -->
   ContentIndexed = "TRUE"
   >
</IIsWebVirtualDir>
</IIsWebService>
```

# The Metabase Schema

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The metabase schema defines what properties can exist in the metabase configuration file. The metabase schema also enforces which properties can be set at which keys. The metabase schema provides this enforcement only for *schema-defined* properties, which are properties written to the metabase configuration file that adhere to the rules of metabase schema. Properties that are not schema-defined can be written in the metabase configuration file as *custom properties*. Custom properties can be used in the metabase configuration file to override schema enforcement or to configure a property with values other than the default values that are defined by the schema.

# Metabase Schema Collections

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The metabase schema is organized into containers called *collections*. Collections are defined in MBSchema.xml by using <Collection> start and end tags, and they are written in the following format:

```
<Collection InternalName = "CollectionName" MetaFlagName = "MetaFlagValue" >
   <Property PropertyName = "PropertyValue" >
      <Flag FlagName = "FlagValue" >
   .
   .
</Collection>
```

In the example above, *CollectionName* is the name of the entire collection. *MetaFlagName* and its associated *MetaFlagValue* apply to the entire collection. *PropertyName* and its associated *PropertyValue* are individual settings within the collection. *FlagName* and its associated *FlagValue* are individual settings within a property.

All collections exist at the same hierarchical level in MBSchema.xml. Therefore, no collection is a parent of another. Reading the MBSchema.xml file from top to bottom, the first collection, named **MetabaseBaseClass**, defines global settings in the schema.

The following code example is taken from the metabase schema file to illustrate how collections are defined.

```
<Collection InternalName = "IIsConfigObject" MetaFlagsEx = "NOTABLESCHEMAHEAPENTRY"
MetaFlags = "HASUNKNOWNSIZES | HIDDEN">
   <Property InternalName = "AccessFlags" ID = "6016" Type = "DWORD"
   UserType = "IIS_MD_UT_FILE" Attributes = "INHERIT"
   MetaFlagsEx = "CACHE_PROPERTY_MODIFIED" DefaultValue = "AccessRead"
   EndingNumber = "0">
      <Flag InternalName = "AccessExecute" Value = "4" ID = "6209"/>
</Collection>
```

The collection directly below **MetabaseBaseClass** is named **IIsConfigObject**. The **IIsConfigObject** collection is unique, because it is where all properties are defined. Only properties that are defined in the **IIsConfigObject** collection can be written to the metabase configuration as schema-defined properties.

# Metabase Schema Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Properties* are configuration objects that are associated with IIS Admin Objects and their corresponding ABOs. All properties that are defined in the **IIsConfigObject** collection contain attributes. The default values of these attributes for each property are also defined.

| **Note** |
| --- |
| The default values that are assigned to property attributes in the **IIsConfigObject** collection are not necessarily the values that are written to the metabase configuration when IIS is installed. The default values of attributes, which are defined by the **IIsConfigObject** collection, are used when you write a schema-defined property to the metabase configuration file after IIS is installed. |

Each property that is defined in the **IIsConfigObject** collection contains the following attributes: **ID**, **Type**, **UserType**, **Attributes**, **InternalName**, and **DefaultValue**. Some properties also contain the following attributes: **MetaFlags**, **MetaFlagsEx**, **StartingNumber**, and **EndingNumber**. Property attributes are contained in the <Property> start tag. The following example from the **IIsConfigObject** collection illustrates how a property is defined.

```
<Property
   InternalName = "ServerComment"
   ID = "1015"
   Type = "STRING"
   UserType = "IIS_MD_UT_SERVER"
   Attributes = "INHERIT"
   MetaFlagsEx = "CACHE_PROPERTY_MODIFIED | WAS_NOTIFICATION"
   DefaultValue = ""/>
```

# Flags

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to attributes, some properties also contain *flags*. Each flag configures a unique functionality of the property, such as read or write access. Each flag is identified by the ABO that uses the unique *bitmask identifier* of the flag.

Flags contain attributes named **Value** and **ID**. Flags are defined in the property's start and end tags, and they are identified by using <Flag> tags. The sum value of all flags that are set for a property determines the value of the property's **DefaultValue** attribute, which is also called the default value of the property. The following example illustrates a property that contains flags.

```
<Property InternalName = "PasswordChangeFlags"
   ID = "2068"
   Type = "DWORD"
   UserType = "IIS_MD_UT_SERVER"
   Attributes = "INHERIT"
   MetaFlagsEx = "CACHE_PROPERTY_MODIFIED"
   DefaultValue = "0"
   EndingNumber = "0">
      <Flag InternalName = "AuthChangeUnsecure" Value = "1" ID = "6292"/>
      <Flag InternalName = "AuthChangeDisable" Value = "2" ID = "6293"/>
      <Flag InternalName = "AuthAdvNotifyDisable" Value = "4" ID = "6294"/>
</Property>
```

| **Note** |
| --- |
| Flags and their attributes are written in the start and end tags of a property. A property that does not contain a flag is encapsulated in the start tag of the property, and it does not use end tags. |

# Remaining Collections

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

All remaining collections, which are below the **IIsConfigObject** collection, are organized by key type. These collections enforce at which keys in the metabase configuration a schema-defined property can be written. For example, if a collection named **IIsWebServer** contains the **ServerComment** property, IIS allows the **ServerComment** property to be written to any key in the metabase configuration that uses the IIsWebServer key type. Therefore, only properties that are contained in these remaining collections can be written (as schema-defined properties) in a metabase key of the same key type in MetaBase.xml.

The exception to this rule is the IIsConfigObject key type, described earlier. Because all properties are contained in the **IIsConfigObject** collection, all properties can be written to any key in the metabase configuration of the IIsConfigObject key type.

Properties are written in these remaining collections in the following format:

```
<Property InheritsPropertiesFrom = "IIsConfigObject:PropertyName"/>
```

where *PropertyName* is the name of the property.

All of the properties that are contained by these remaining collections inherit their default values from the **IIsConfigObject** collection, because the **IIsConfigObject** collection is where all properties are defined.

You cannot use ADSI or WMI, which are not defined in metabase schema, to update properties in the metabase configuration. Changes written directly to the MetaBase.xml file that do not comply with metabase schema rules might not be written to the in-memory metabase. If this occurs, an error or warning is sent to the event log.

# Metabase Security and Auditing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A default installation of IIS 6.0 ensures metabase security by setting strict access control entries (ACEs) on the metabase files and by encrypting sensitive data within the files. If you maintain this level of security, perform regular backups, use a strong administrator password, and limit the number of users who have administrative credentials, you are taking the proper precautions to protect your metabase files.

# File-Level Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

As described in Table 4.4, IIS installs the metabase files with strict ACEs set to prevent anyone but administrators from viewing your configuration data. An access control list (ACL) is a container for ACEs.

**Table 4.4 Metabase files, purpose, and permissions**

| File | Purpose | ACL |
|---|---|---|
| *systemroot*\System32\Inetsrv\MetaBase.xml | Stores configuration data for the IIS services. | NT AUTHORITY\SYSTEM : Full control<br>BUILTIN\Administrators : Full control |
| *systemroot*\System32\Inetsrv\MBSchema.xml | Stores the schema for the configuration file. The schema defines what IIS properties can be set at certain metabase keys. | NT AUTHORITY\SYSTEM : Full control<br>BUILTIN\Administrators : Full control |
| *systemroot*\System32\Inetsrv\ History*HistoryFile* | Stores the metabase history files that are created automatically by IIS. | NT AUTHORITY\SYSTEM : Full control<br>BUILTIN\Administrators : Full control |
| *systemroot*\System32\Inetsrv\MetaBack*BackupFile* | Stores the metabase backup files that are created on demand by using Backup/Restore Configuration. | NT AUTHORITY\SYSTEM : Full control<br>BUILTIN\Administrators : Full control |

# Encrypted Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS encrypts sensitive data in the metabase configuration file, MetaBase.xml, so that it cannot be viewed even if an unauthorized user gains access to the file.

| **Important** |
| --- |
| Do not manually change encrypted properties in MetaBase.xml. There is no way to encrypt your data before inserting it by using Notepad or another text editor. Only WMI, ADSI, or ABOs can be used to change the data in encrypted properties. |

Metabase properties are marked for encryption by an attribute that is set on the property in the metabase schema file. The following example from the metabase schema file shows how the **AnonymousUserPass** property is marked for encryption with a **SECURE** attribute:

```
<Property InternalName = "AnonymousUserPass" ID = "6021" Type = "STRING" UserType =
"IIS_MD_UT_FILE"
   Attributes = "INHERIT|SECURE" MetaFlagsEx = "CACHE_PROPERTY_MODIFIED"
   DefaultValue = ""/>
```

| **Caution** |
| --- |
| You cannot use WMI, ADSI, or ABOs to add an attribute to an existing property in the metabase schema. For example, you cannot add the **SECURE** attribute to an existing property. If you attempt to use WMI, ADSI, or ABOs to remove a property in the metabase schema and then create it again with the **SECURE** attribute, the metabase configuration data for that property is lost. You can, however, use ADSI to create new properties in the metabase schema with the **SECURE** attribute. Manual changes made to the metabase schema are not supported and might cause an error. |

The following metabase properties are encrypted:

- **ADConnectionsPassword**
- **AdminACL**
- **AdminACLBin**
- **AnonymousUserPass**
- **ImapDsPassword**
- **LogOdbcPassword**
- **Pop3DsPassword**
- **RoutePassword**
- **SmtpDsPassword**
- **UNCPassword**
- **FeedPassword**
- **WAMUserPass**

# Checklist: Metabase Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table 4.5 is a list of recommended steps to ensure the security of the metabase.

**Table 4.5 Steps to ensure metabase security**

| Complete | Step | Reference |
|---|---|---|
| | Make sure that the passwords for all accounts on your computer are strong and not written anywhere near the computer or within view. | See "Passwords" in Help and Support Center for Windows Server 2003. |
| | Maintain strict access control permissions on the metabase files that are listed in Table 4.4. Only the LocalSystem account and the Administrators group should be listed in the ACLs as having full control of the metabase files, including the history files and backups. IIS sets this access control by default. | See "Access control" in Help and Support Center for Windows Server 2003. |
| | Employ the concept of *least privilege* on your computer. Allow only one person to know the Administrator password, and limit the number of people whose accounts are in the Administrators group. Set ACLs on all folders and files so that the fewest users have the lowest-level privileges needed to run required tasks. Set IIS security on all IIS applications and virtual directories so that the least permission necessary exists to allow the appropriate clients to view your content. | See "Access control" and "Security" in Help and Support Center for Windows Server 2003. |
| | Do not run the **cipher** command or use Encrypting File System (EFS) on the metabase files. Sensitive data, such as passwords that are stored in MetaBase.xml, is already encrypted. Encrypting the entire MetaBase.xml file slows the performance of your IIS server and might cause errors. | See the "Cipher" and "Encrypting File System" topics in Help and Support Center for Windows Server 2003. |
| | When you edit the MetaBase.xml file manually, copy the MetaBase.xml file first, and then work on the copy. When you have checked your changes, replace MetaBase.xml with your copy. | See Editing the MetaBase.xml File While IIS Is Running. |
| | Create backups of your metabase files using Backup/Restore Configuration whenever you make a significant change in the metabase. If you allow other people to configure the metabase, make it a policy that they create a backup with their name in the title whenever they make a change. IIS periodically makes its own backup files, called history files, but history files might not be created for every change to the metabase. | See Backing Up the Metabase |
| | Do not use the metabase import and export feature to create regular backup files. This feature is meant to transport sections of the metabase to other computers. However, it is recommended that you create an export file for your entire metabase periodically in the event that your computer hardware fails and you must move your Web sites to another computer. | See Metabase Import and Export. |

| | | |
|---|---|---|
| | Remember that backing up or exporting the metabase does not back up your content (for example, .htm files, .asp files, and components). The metabase holds only configuration data, such as the location where your content is stored. To back up your content, use Windows Backup. | See "Backing up and restoring data" in Help and Support Center for Windows Server 2003. |
| | Monitor your event log for IIS event messages. For example, if you attempt to make changes to the metabase when there is a write-lock on the in-memory database, a metabase history file is created and an error is written to the event log. If you successfully create an FTP virtual directory, but IIS cannot enumerate the physical directory, an error is written to the event log. | See "Event Viewer" in Help and Support Center for Windows Server 2003. |
| | Set up file auditing on the various files that affect the metabase. This way, if the metabase becomes corrupted, you can identify the account of the last user who made a change. Select your auditing choices carefully for the files listed below, and monitor the audit logs for a period of time to determine if the settings are meeting your needs.<br>`Systemroot\System32\Inetsrv\MetaBas.xml….`<br>If you choose to audit the MetaBase.xml file, only users who edit the metabase manually are identified in the audit logs by their own account name. Users who edit the metabase by using IIS Manager or other tools are logged as accessing the metabase under the LocalSystem account. Because IIS itself is constantly accessing the metabase, audit logs can grow very quickly unless you limit auditing to specific users, such as the local Administrators group, and specific access attempts, such as writing to a file. Auditing MetaBase.xml is still valuable for determining when changes were made, but you might have to audit the other files listed here to find the exact identity of the user.<br>Custom tools and script files that use WMI, ADSI, or ABOs to edit the metabase:<br>Auditing should be set on these tools to identify who executes them. If you audit MetaBase.xml and want to find out who made a change at 3:01 P.M. on a certain day, you can look in the audit logs for your scripts and tools to see if someone ran one of them at 3:01 P.M.<br>`Systemroot\System32\Inetsrv\Inetmgr.exe….`<br>Audit this file for the same reason you audit custom tools and script files, above. | See "Auditing overview" in Help and Support Center for Windows Server 2003. |

# Metabase Auditing

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

Beginning with the Windows Server 2003 Service Pack 1 (SP1) release, IIS 6.0 includes a metabase auditing feature that allows tracking of each change that is made to the metabase. Metabase auditing is enabled by enabling an audit access control entry (ACE) on a node in the metabase. After ACE is enabled, whenever a metabase change takes place on that node, an audit event is published in the NT Security event log.

The following information is recorded in the NT Security event log:

- What was changed (metabase node, property, and old and new values).
- When the change was made (date and time).
- Who made the change (domain and user name).
- Success or failure of the change attempt (HRESULT).
- When a change is made remotely (client IP number).

Without metabase auditing, you cannot determine who made changes to the metabase and when those changes were made. IIS 6.0 does feature a metabase history, with up to 10 (configurable) Metabase.xml files stored in the Inetsrv\History folder. However, history files do not reveal what changed, who made the change, and when the change was made.

The auditing requirements for changes to metabase properties differ from changes to metabase keys. For a change to a key, it is necessary to log only the key data, because changes to a key affect everything beneath it (child operations are not audited). Changing a property requires more information, including both the key and the property name.

The following list specifies the audit event content for operations on both metabase keys and metabase properties:

- Delete key: name of node being deleted
- Delete property: name of key and name of property being deleted
- Move key: old node location and new node location
- Copy key: source node location and new node location
- Add key: name of node being added
- Add property: name of key, name of property, and value of property
- Rename key: names of old node and new node
- Change property: name of key, name of property, old and new values of property

| **Note** |
| --- |
| To avoid disclosing sensitive information, such as passwords, values of secure properties will not appear in audit event log entries. |

**Enabling and Disabling Metabase Auditing**

Metabase auditing is enabled when the following conditions are met:

- In the Windows Group Policy Editor (gpedit.msc), audit-object access is enabled and auditing is turned on for both

successes and failures.

- ACEs are enabled on the metabase nodes that you want to audit.
- The command script iiscnfg.vbs was used to enable auditing.

Metabase auditing is disabled when the following conditions are met:

- In the Windows Group Policy Editor (gpedit.msc), audit-object access is disabled.
- ACEs are disabled on the metabase nodes that you no longer want to audit.
- The command script iiscnfg.vbs was used to disable auditing.

For information about metabase audit event log messages, see Metabase Auditing Event Messages.

# Backing Up and Restoring the Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Creating backups of the metabase is a vital part of maintaining metabase reliability. The ability to create backup files and to restore the metabase by using backup files has been enhanced in IIS 6.0. You can now restore the backup on other computers if portable backup has been chosen. History file copies of the metabase are created automatically by IIS, and an IIS administrator using IIS Manager can create them on demand. All backups, regardless of how they were created, are stored together and displayed together in IIS Manager.

Restoring a backup from an earlier version of IIS is not supported. After you install or upgrade to IIS 6.0, it is recommended that you perform a metabase backup as soon as possible to preserve your configuration data.

Backing up the metabase should be supplemented with proper metabase security practices. For more information about metabase security, see Metabase Security and Auditing.

# Backing Up the Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Using Backup/Restore Configuration, you can create complete backup copies of the IIS metabase. The backup files are either portable or nonportable copies of both the metabase configuration file (MetaBase.xml) and the matching metabase schema file (MBSchema.xml). Metabase backup files provide a way to restore your metabase configuration and schema data in the event that the metabase becomes corrupted. Portable backup files can be restored either to the computer on which the backup was made or to other installations of Windows Server 2003, while non-portable backup files can only be restored to the same system from which it was made.

When a backup is created, a set of two files is created and named in the following way:

```
Name.mdx and Name.scx
```

In the syntax above, *Name* is the name that the administrator uses to name the backup set, and *x* is the version number of the backup set, starting with 0 (zero). The version number is increased by one for each backup set that uses the same name, up to 999. For example, if you use the name *MyBackup* as the name of your backup set twice, MyBackup.md0 and MyBackup.sc0 are created on the first backup; on the second backup, the files are named MyBackup.md1 and MyBackup.sc1. By default, backup files are stored in the *systemroot*\System32\Inetsrv\MetaBack folder.

Backup files contain only configuration data; they do not include your content. To back up your content, use Windows Backup. For more information about using Windows Backup, see "Backing up and restoring data" in Help and Support Center for Microsoft® Windows® Server 2003.

| **Note** |
| --- |
| The metabase is locked while the backup is in progress. |

**Portable Backup**

When creating a portable backup, the administrator provides a password that is used by IIS to encrypt the backup files. The password is encrypted and stored in the header of the backup file. Only the administrator password and secure properties within the backup files are encrypted; all other information within the backup files is plain text. After the backup file is encrypted, the password within the backup file cannot be changed.

**Non-portable Backup**

When creating a non-portable backup, the administrator does not supply a password. Therefore, non-portable backup files are encrypted with a blank password, which allows any member of the Administrators group to restore the metabase. Only the blank password and secure properties are encrypted; all other information within the backup file is plain text.

# The Metabase History Feature

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS automatically maintains a record of metabase changes in history files that are saved to disk. You can also configure the number of metabase history files to save. By using the metabase history feature, you can revert the metabase through any number of changes to restore a particular configuration, or to see what has changed between revisions.

By marking each new metabase file with a unique version number and saving a copy of the file in the History folder, the metabase history feature automatically keeps track of metabase changes that are saved to disk. Each history file is then available for restoring the MetaBase.xml and MBSchema.xml files from the History folder, editing the MetaBase.xml file while IIS is running, and troubleshooting event log errors. The metabase history feature is enabled by default.

**History Folder**

The History folder stores versioned copies of the MetaBase.xml and MBSchema.xml files. These copies can only be viewed by members of the Administrators group. The location of the History folder is *systemroot*\System32\Inetsrv\History.

A *history file pair* consists of a MetaBase.xml and MBSchema.xml file of the same major and minor version numbers. The default setting allows for a maximum of 10 versioned history file pairs to be stored in the History folder, but you can configure this number. However, before increasing the number of versioned sets of files stored in the History folder, you should ensure that you have adequate hard disk space on the drive volume where the History folder is located. (You can estimate the amount of hard disk space that is required by multiplying the average combined file size of MetaBase.xml and MBSchema.xml files by the number of versions that you want to store in the History folder.) You should not decrease the value of the **MaxHistoryFiles** property to below 10.

When the in-memory metabase is written to disk, IIS checks to determine whether the number of history file pairs that are contained in the History folder exceeds the value of the **MaxHistoryFiles** property. If the number of history file pairs exceeds the value of the **MaxHistoryFiles** property, the oldest history file pair is deleted.

For information about backing up your entire operating system, including special steps that must be performed to include IIS in that backup, see Backing up and restoring data in Help and Support Center for Windows Server 2003. For information about the event log, see Event log in Help and Support Center for Windows Server 2003.

# Configuring the Metabase History Feature

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By default, the metabase history feature is enabled and the number of history file pairs that are saved within the History folder is defined by the default value of the **MaxHistoryFiles** property. A history file pair is the combination of a copy of a MetaBase.xml file and an MBSchema.xml file of the same version in the History folder. If the **EnableHistory** property is not written in the MetaBase.xml file, IIS interprets the absence of the property to mean that the history feature is enabled, and the **MaxHistoryFiles** property is set to the default value.

**To change the metabase history default settings**

1. Add the **EnableHistory** and **MaxHistoryFiles** metabase properties to the /LM level of the metabase.

**To disable metabase history**

1. Set the value of the **EnableHistory** property to **0** (zero).

**To enable metabase history**

1. Set the value of the **EnableHistory** property to **1**.

**To change the number of history files that are stored in the History folder**

1. Change the value of the **MaxHistoryFiles**property.

| ⚠**Caution** |
| --- |
| Disabling history files can have serious consequences if the metabase becomes corrupted by improper editing. Without backup files, the IIS server might have to be reconfigured from an empty state. However, keeping backup files can also be a security risk if you do not ensure that the files are protected by ACLs. By default, only members of the Administrators group can view the history files, but if you copy history files to a new location, you must duplicate those settings. For information about configuring ACLs, see "Access control" in Help and Support Center for Windows Server 2003. |

# Naming the Metabase History Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To keep track of metabase history, IIS uses the following versioning scheme to enumerate the copies of MetaBase.xml that are stored in the History folder.

**Metabase Version Numbers**

The value of the **HistoryMajorVersionNumber** metabase property, also referred to as the major version number, increases when the in-memory metabase is written to disk.

The minor version number increases when the edit-while-running feature is enabled and an administrator changes and saves the MetaBase.xml file directly with an application such as Notepad.

**HistoryMajorVersionNumber Property**

The **HistoryMajorVersionNumber** property is located under the IIS_Global node in the metabase, as shown in the following example:

```
<?xml version = "1.0"?>
<configuration xmlns = "urn:microsoft-catalog:XML_Metabase_V54_0">
<MBProperty>
<IIS_Global Location = "."
   BINSchemaTimeStamp = "9676376defb8c201"
   ChangeNumber = "1023"
   HistoryMajorVersionNumber = "154"
   SessionKey = "4585thge754u4509853743hg0948054ug6t095"
   XMLSchemaTimeStamp = "444e436defb8c201"
   >
</IIS_Global>
```

| **Important** |
| --- |
| IIS uses the **HistoryMajorVersionNumber** property to track the versions of the metabase. Thus, manually changing the value of the **HistoryMajorVersionNumber** property is not recommended. |

**Minor Version Number**

The minor version number is not a property that is stored in the metabase. IIS calculates the minor version number when the MetaBase.xml file is edited and saved directly, and if edit-while-running is enabled, as follows:

IIS looks in the History folder to determine the last minor version number that was used for a history file with the same HistoryMajorVersionNumber number as MetaBase.xml. IIS then uses the next sequential minor version number when naming the new history file.

When a history file is created because the in-memory metabase is written to disk, either through the regular timed history algorithm, or by the user explicitly saving the current configuration to disk, the major version number is incremented by 1, and the minor version number is reset to zero.

**Naming the History Files**

When a copy of MetaBase.xml is written to the History folder, the value of the **HistoryMajorVersionNumber** property is incremented and added to the names of the file pair. If edit-while-running is enabled, and the user makes a manual change to the metabase, the minor version number is incremented and added to the names of the file pair.

The file names of versioned MetaBase.xml files and MBSchema.xml files are created in the following format, respectively, where both the value of the **HistoryMajorVersionNumber** property and the minor version number are 10-digit numbers:

```
MetaBase_HistoryMajorVersionNumber_MinorVersionNumber.xml
MBSchema_HistoryMajorVersionNumber_MinorVersionNumber.xml
```

When either the HistoryMajorVersionNumber or the minor version number value is less than 10 digits, the number is padded with zeros in the file name to create a 10-digit number. This keeps the file names aligned and in numerical order when you sort by file name, making the contents of the History folder easier to read.

The following is an example of file names in the History folder:

```
MetaBase_0000000001_0000000000.xml
MetaBase_0000000001_0000000001.xml
MetaBase_0000000002_0000000000.xml
MetaBase_0000000002_0000000001.xml
MetaBase_0000000002_0000000002.xml
MetaBaseError_0000000000.xml
MetaBaseError_0000000001.xml
MBSchema_0000000001_0000000000.xml
MBSchema_0000000001_0000000001.xml
MBSchema_0000000002_0000000000.xml
MBSchema_0000000002_0000000001.xml
MBSchema_0000000002_0000000002.xml
```

For each instance of a MetaBase_*MajorVersion_MinorVersion*.xml file within the History folder, there is an MBSchema_*MajorVersion_MinorVersion*.xml file of the same version that is used with the metabase configuration.

# Metabase Error Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

File names that contain the word *error* are created in the History folder under the following circumstances:

- Edit-while-running is enabled.
- You edit MetaBase.xml directly by using an application, such as Notepad.
- An error, such as (but not limited to) the following occurs:
    - An XML tag is omitted in MetaBase.xml.
    - A property name is misspelled in MetaBase.xml.
    - The corresponding history file with the same HistoryMajorVersionNumber value as MetaBase.xml is not found.

Error files are named in the following format, where *TenDigitNumber* is the version number of the error file, starting with version number 0000000000:

```
MetaBaseError_TenDigitNumber.xml
```

The version number increases by one for each additional error file that is created.

# Restoring the Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Restoring the metabase from history files is similar to restoring it from backup files. For a comparison of the two methods of restoring the metabase, see Table 4.6.

**Table 4.6 Comparison of Metabase Restore Methods**

| Using Backup Files | Using History Files |
|---|---|
| If a password is provided by the user to encrypt secure properties, the metabase can be restored to a different installation of Windows, even on a different computer. | The metabase can be restored only on the same installation of Windows from which the history files were created. This is because data is encrypted with the machine key in the backups. |
| Backup files can be created on demand, scheduled, or created by using a programmatic administration script. | History files can be created automatically (when the IIS history feature is enabled) or by using a programmatic administration script. |
| Backup files can be restored from IIS Manager, by using APIs or by using programmatic administration scripts. | History files can be restored manually or by using IIS Manager. |

# Restoring the Metabase from History Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use automatically generated history files to restore your IIS configuration from IIS Manager. However, you may find it necessary to restore the metabase without using IIS Manager. For example, suppose a configuration change is made incorrectly, or if it is made using incorrect XML syntax. In this case, IIS Manager would be unable to enumerate the metabase.

For more information about restoring the metabase from history files, with or without IIS Manager, see Restoring the Metabase with History Files.

# Restoring the Metabase from Backup Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can restore a metabase backup to the same computer on which the backup was made or, if the backup was created with a password, onto a different computer. Because a metabase backup does not include content files, if you restore a metabase backup to a different server you must copy your content files as well. The content files must exist in the same location and on the same drive used on the original server. Restoring a backup from an earlier version of IIS is not supported.

For more information about restoring the IIS metabase from backup files, see Restoring the Metabase with Backup Files.

**SSL Certificates**

If you use a Secure Sockets Layer (SSL) certificate on your Web server and you restore the metabase backup to a different installation of the operating system, an error occurs if the certificate is not available or installed. If a site's SSL certificate is missing, connections are dropped, and the browser displays a page unavailable message. The solution is to reinstall the certificate into the same certificate store as on the original computer, and restart IIS.

# Restoring a Section of the Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can make fixes to specific sections of the metabase in two ways:

- You can use the import feature without the merge option to import metabase elements over damaged data.
- If the damaged section prevents IIS from loading the metabase, you can manually edit the MetaBase.xml file in any text editor by copying sections from an export or backup file and pasting them into the MetaBase.xml file.

  For information about importing and exporting metabase elements, see Metabase Import and Export.

You can be prepared to make fixes to the metabase by creating a template of your entire metabase, as well as by creating export files of new sections of your metabase file at various stages of its life, so that you have something to work from if a section becomes corrupted.

To create a template of your entire metabase, first create a backup file. Remove any machine-specific and system-specific information from your backup file, and then remove extraneous information and configuration information that is specific to applications, sites, and virtual directories.

What you are left with is a template of top-level configuration information that might not be recoverable using metabase import/export. Together with the export files of specific sections that you created, you are prepared for any loss of data or corruption. One way to identify where XML syntax corruption exists is to open the MetaBase.xml file with Internet Explorer. Internet Explorer will identify the line at which it encountered invalid xml syntax.

# Editing the Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Editing the MetaBase.xml file directly by using common text-editing tools is an efficient way to configure the metabase over low-bandwidth connections. Editing MetaBase.xml with Notepad is more efficient than using IIS Manager, because IIS Manager is relatively slow over low-bandwidth lines, such as a 56 KB modem. Additionally, editing the metabase file directly is faster than through user interfaces (UIs), such as IIS Manager, when hundreds or thousands of sites are running on the same computer. Editing the MetaBase.xml file directly is particularly useful for administrators who do not want to use programmatic methods to administer IIS but want an alternative to the UI methods. Editing the Metabase directly over a remote connection is not recommended because a dropped connection can leave the Metabase in an indeterminate state.

You can make changes to the MetaBase.xml file while IIS is running only if the edit-while-running feature is enabled. If edit-while-running is not enabled, you must stop IIS before you edit the MetaBase.xml file. This is inefficient and not a recommended practice.

# Best Practices for Maintaining Metabase Integrity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following practices can greatly reduce the chance that the metabase will become corrupted.

- Enable edit-while-running. This ensures that changes made to the metabase configuration while IIS is running are written to the on-disk metabase.
- Always back up the metabase before making configuration changes. Make the right kind of backup. A password-protected backup is cleanly transportable to other systems, while a backup made without a password is tied to the system on which the backup was made.
- When editing the metabase with a text editor, use a text editor that write-locks the open file. This prevents the possibility of more than one person editing the metabase at a time.
- After editing the metabase, open the file in Internet Explorer 6 or above. Any errors in the XML syntax will be displayed when Internet Explorer opens the file.
- Never allow changes to the metabase by more than one entity at a time. For example, do not edit the metabase by hand at the same time an administrative script that changes metabase settings is running.
- If you make changes to the metabase and it does not load properly, use a file comparison tool to compare the backup you made before editing the metabase with the edited metabase. This can help pinpoint any unintended changes.

Note the following about editing the metabase:

- Recoverable metabase schema errors occur when text editing yields well-formed XML but one or more properties are invalid, according the metabase schema.
- Editing the metabase schema is not supported.
- If the metabase contains incorrectly formatted XML, IIS cannot read the metabase and will not run. You can restore the last good history file to correct this problem.
- If an invalid configuration is created in MetaBase.xml, the error cannot read property is logged to Event Viewer, but the rest of IIS keeps running.

# Configuring the Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS metabase is a plaintext XML file containing most IIS configuration values. The metabase can be configured manually, by using IIS Manager or a text editor, or programmatically, using the IIS WMI and ADSI providers. Deciding the method you use to configure the metabase depends on the task you want to accomplish. For a one-time, one-step configuration task, editing the metabase file directly in a text editor might be easiest; for repetitive tasks that require a number of steps, a programmatic approach is most efficient. Another important aspect of configuring the metabase is knowing how to back up and restore the metabase if changes result in an error.

IIS includes several methods for changing values in the metabase. As shown in Figure 4,2, you can set or view metabase values programmatically by using any of the following resources:

- IIS Manager. The user interface (UI) that you can use to update the in-memory metabase. This is the safest way to update the metabase. However, it does not expose all available metabase properties.
- IIS ADSI provider. A set of scripting interfaces that you can run from a script.
- IIS WMI provider. A set of scripting interfaces that you can run from a script.
- Admin Base Objects (ABOs). A set of binary COM interfaces.

All configuration settings are written to the metabase through ABOs, as illustrated in Figure 4.2.




**About the IIS ABOs**

As shown in Figure 4.2, the ABOs are central to updating the in-memory metabase. An ABO is a set of binary COM objects that are used to write IIS configuration values that are stored in the IIS metabase. Each object has properties, and each property is stored in the metabase. You can manipulate IIS configuration by using ABOs to add new virtual directories, new Web sites, and so forth.

IIS Manager and the other resources in the illustration that are used to update the in-memory metabase write directly to the ABOs. The ABOs then write the changes to the in-memory metabase. You can also access ABOs directly by using C or Microsoft® Visual C++® development system, Microsoft® Visual Basic® development system, Delphi, or any COM-compliant language to programmatically update the in-memory metabase. You cannot use scripting language code to programmatically update the in-memory metabase.

**Metabase Storage Layer**

The in-memory metabase is periodically written to disk through the metabase storage layer. The metabase storage layer is the part of IIS that reads the metabase files into the in-memory metabase and also writes the in-memory metabase to the metabase files. The storage layer writes the in-memory metabase directly to the metabase files when a write-to-disk event occurs or when IIS is shut down.

| **Important** |
| --- |
| You cannot edit the MBSchema.xml file directly by using edit-while-running. Changes to MBSchema.xml must be made programmatically by using ADSI. |

**Adding Comments to the Metabase**

A metabase comment must be stored in the in-memory metabase in the correct XML format and associated with a particular node within the metabase, so the context of the comment is preserved during subsequent writes of the metabase to disk. Put node-specific comments right above the node. If you delete the node, the comments about the node are deleted as well.

When loading the metabase, the configuration system loads comments as a generic property into the metabase, along with other metabase data. Any comments found at the bottom of the file (below all nodes) will always be preserved at the bottom of the file.

Comments are added in the following format:

```
<!-- your comment goes here -->
```

The following is a list of rules concerning metabase comments:

- Comments that are added above the MBProperty node are not picked up by edit-while-running.
- Comments that are added between the MBProperty node and the IIS_Globalnode cause an error file to be created.
- Comments that are added between IIS_Global and IIS_ROOT are picked up by edit-while-running.
- Adding comments between start and end XML tags results in an error file and the in-memory metabase being flushed to disk.

# Editing Metabase Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can write metabase properties in MetaBase.xml as either schema-defined or custom properties. Writing schema-defined properties results in an efficient use of the IIS cache, which is where the in-memory metabase and schema are stored, because the attributes are inherited from the schema rather than rewritten for each instance of a property. Using custom properties is not recommended, because they might not show up in the UI and ADSI cannot access them after they have been created.

A schema-defined property is contained in the start tag of a key, and only the property name and value are defined. Attributes for a schema-defined property come from the metabase schema. In the following example, the **EnableEditWhileRunning** property is a schema-defined property with a value of zero, and it is written in the IIs_Computer key:

```
<IIs_Computer Location = "/LM"
   EnableEditWhileRunning = "0">
```

A custom property defines some or all of the attributes of a property. Unlike a schema-defined property, a custom property is written between the XML start tag and the end tag of a metabase key. Custom properties are easy to identify because the name of the element that contains the custom property is **Custom**. The name and attributes of a custom property are contained in an XML element named **Custom**.

At a minimum, the **Name** and **ID** attributes must be defined in a custom property. Other than the name of the property and the **ID** attribute, all remaining attributes for a custom property come from the schema, if they are undefined. You need to define the other attributes only when you want them to contain a different value than the schema-defined attribute values. In the following example, the **Win32Error** property is defined as a custom property with a value of zero, and it is written within the IIsFtpVirtualDir key.

```
<IIsFtpVirtualDir Location = "/LM/MSFTPSVC/1/ROOT">
   <Custom
      Name = "Win32Error"
      ID = "1099"
      Value = "0"
      Attributes = "INHERIT"/>
</IIsFtpVirtualDir>
```

| **Important** |
| --- |
| Each property that is defined in the schema (MBSchema.xml) has a unique **ID** attribute. When you define a known property as a custom property, you must specify the unique ID number for that property. |

#### General Notes for Editing the Metabase

Keep the following points in mind when editing the metabase:

- Properties with the **SECURE** attribute set are stored encrypted.
- All text in MetaBase.xml is case sensitive.
- The **KeyType** property contained in a node is used to determine the element or node type that is applied to the metabase path.
- Nodes that do not contain a **KeyType** property are stored as the generic object IISConfigObject, which can contain any property, and any node can be contained within any other node.
- Custom properties are allowed for any node type.
- The IIS_Global node must not be edited.
- A schema-defined property that violates the schema — for example, a misspelled property name — is logged and the property is ignored.

- Custom properties require a metabase **ID** attribute.
- The **Multi-SZ** attribute is replaced with the multistring value.
- Flags contain bitmask identifiers.
- The internal metabase implementation is Unicode. To support international and non-ASCII characters, the MetaBase.xml file is encoded in UTF-8 text format.

**Changing Schema-Defined or Out-of-Schema Properties**
You can use custom properties for the following purposes:

- To use properties that are not schema-defined.
- To use properties with different attribute values than what is available in the schema-defined properties.

When you write a property by using different attribute values than what is schema-defined, you must write the property as a custom property. Only the attributes that are different from the schema should be defined in MetaBase.xml.

The following example shows a schema-defined property:

```
AccessFlags = AccessWrite
```

The following example shows a custom property:

```
<Custom Name = "AccessFlags"
   ID = "6016"
   Type = "DWORD"
   Value = "1"
   UserType = "IIS_MD_UT_FILE"
   Attributes = "INHERIT"/>
```

**Starting or Stopping IIS**
When IIS is started or restarted, the metabase configuration that is stored in MetaBase.xml and the schema information that is stored in MBSchema.xml is read by the storage layer and then written to the in-memory metabase through ABOs, as shown in Figure 4.2.

If the MetaBase.xml file cannot be parsed when IIS is started or restarted, an error is sent to the event log and IIS does not start. The reasons that MetaBase.xml might not be parsed correctly include — but are not limited to — missing XML tags, misspelled property names, or corruption of the MetaBase.xml file. To determine the cause of failure, you can use the last history file to determine what has changed in the MetaBase.xml file. If the metabase cannot be parsed, you can also restore it to an earlier version by using the last manual backup or an automatically generated history file.

When IIS is stopped, it checks to determine whether the configuration node or the schema node have changed in the in-memory metabase since the last time the in-memory metabase was written to disk. The configuration node contains the metabase configuration, which is stored in the MetaBase.xml file. The schema node contains the metabase schema, which is stored in the MBSchema.xml file.

To minimize the amount of time that is required to stop IIS, the MetaBase.xml file is not written to disk if the configuration node in the in-memory metabase has not changed. If the configuration node in the in-memory metabase has changed since the last time the in-memory metabase was saved to disk, the MetaBase.xml file is overwritten with the contents of the configuration node, and a new history file is created. For more information about history files, see Naming the Metabase History Files.

# Editing the MetaBase.xml File While IIS Is Running

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This section contains information about the metabase edit-while-running feature, which is also known as Direct Metabase Edit. The ability to edit the metabase while IIS is running makes administration simpler and avoids site down time, because you do not have to stop IIS when you need to change the metabase configuration.

# The Edit-While-Running Process

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS uses Windows file change notification to determine when the MetaBase.xml file has been saved. When IIS detects that the metabase configuration file has been saved, a series of checks are made, as illustrated in Figure 4.3. Several possible scenarios are presented, based on the results of the steps shown in the illustration, to demonstrate the behavior of edit-while-running.

| **Important** |
| --- |
| The MBSchema.xml file cannot be edited directly by using edit-while-running. Changes made to metabase schema must be made programmatically. |




**Explanation of the Edit-While-Running Process**

**Step 1: IIS receives a file change notification that MetaBase.xml has been saved.**
IIS receives a file change notification from Windows when the metabase configuration file is saved. If MetaBase.xml is write-locked and you attempt to save MetaBase.xml directly with an application such as Notepad, an error is generated and a file change notification does not occur.

**Step 2: IIS looks in MetaBase.xml for the HistoryMajorVersionNumber value.**
IIS reads the value of the **HistoryMajorVersionNumber** property in MetaBase.xml to determine the major version number.

**Step 3: IIS looks in the History folder for the corresponding history file.**
The corresponding history file is named with the highest minor version number with the same HistoryMajorVersionNumber value found in MetaBase.xml in step 2. If the corresponding history file is found, IIS proceeds to step 4.

**Step 4: Parse MetaBase.xml.**
IIS parses MetaBase.xml to determine whether there are fatal XML errors, such as missing XML end tags. If MetaBase.xml can be parsed successfully, IIS proceeds to step 5.

**Step 5: Compare MetaBase.xml to the corresponding history file.**
IIS compares MetaBase.xml to the corresponding history file of the same major version within the History folder. IIS determines the differences between the two files and does one of the following:

- If no changes were made to MetaBase.xml, a history file is not created and no further action is taken.
- If changes were made to MetaBase.xml, IIS proceeds to step 6.

**Step 6: Check that the level exists in the in-memory metabase that changes were made to in MetaBase.xml.**
IIS determines whether the level/node exists within the in-memory metabase that changes were made to in the MetaBase.xml file. If the level/key exists within the in-memory metabase, IIS proceeds to step 7.

**Step 7: Write the changes to the in-memory metabase through ABO.**
The changes are written to the in-memory metabase through ABOs.

A change that was made to MetaBase.xml might not be sent to the in-memory metabase for any of the following reasons:

- The change violates metabase schema; for example, a property name is misspelled.
- The in-memory metabase is write-locked by someone making a simultaneous change by using ADSI, ABO, or WMI to modify the same metabase node or property.
- The key in the in-memory metabase that changes were made to in MetaBase.xml does not exist. This could occur if an administrator deletes the key in the in-memory metabase, by using a programmatic interface or IIS Manager, before MetaBase.xml is changed and saved to disk.

If a change is not written to the in-memory metabase for any of the reasons listed above, an error or warning is sent to the event log.

**Step 8: IIS creates a new history file that contains the contents of the corresponding history file and the changes that were written to the in-memory metabase.**
A history file is created in the History folder. This new file includes the combination of contents of the corresponding history file plus the changes that were written to the in-memory metabase in step 7. The new history file is named with the major version number that was found within MetaBase.xml (the value of the **HistoryMajorVersionNumber** property that was determined in step 2) and the next highest minor version number. IIS looks within the History folder to find the highest minor version number that was used for the same major version number, and it uses the next sequential minor version number to name the new history file. This naming convention permits succeeding edits to the same major version of the file to work correctly.

# Simultaneous Updates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS supports simultaneous updates between programmatic interfaces such as ADSI and WMI. However, under some circumstances, errors can occur as described in step 7 above. The rule is that the last write always wins.

IIS does not support simultaneous updates between the programmatic interfaces and an administrator who directly edits MetaBase.xml. Simultaneous writes to the same metabase node or property can result in an error because of the resulting race condition. You can protect against this type of error by allowing only one administrator to edit the metabase at a time.

# Enabling Edit-While-Running

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Edit-while-running can be enabled by using either IIS Manager or a command prompt. Before you can enable edit-while-running, the metabase history feature must be enabled and set to a value of at least 10 (the default metabase history setting).

| **Important** |
| --- |
| All text within MetaBase.xml is case sensitive. |

For more information about enabling edit-while-running from IIS Manager or a command prompt, see Enabling Edit-While-Running in IIS 6.0.

# Writing the In-Memory Metabase and Schema to Disk

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The metabase is written to disk either when MetaBase.xml is edited directly and manually saved to disk, or when the in-memory metabase is saved to disk.

| **Important** |
| --- |
| MetaBase.xml can be edited by hand and saved to disk only when edit-while-running is enabled or when the IIS Admin service is stopped before the MetaBase.xml file is edited and saved. |

The in-memory metabase is written to disk only under the following circumstances:

- When the IIS Admin service is stopped, and there are changes pending in the in-memory metabase.
- The in-memory metabase is written to disk 60 seconds after a change is made to the metabase; however, if more than 30 changes occur during those 60 seconds, the write-to-disk event is postponed for another 60 seconds. The in-memory metabase write-to-disk event is postponed a maximum of five times, after which it is written to disk.
- When saved programmatically, as when using a resource such as ADSI, WMI, or ABO.

Figure 4.4 and the following discussion describe how the in-memory metabase is written to disk. The verification checks can influence how IIS saves the in-memory metabase to disk, as illustrated by the possible scenarios.

# Writing the In-Memory Metabase to Disk

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Figure 4.4 shows the process IIS uses to determine when to write the in-memory metabase to disk.




Step 0: A change is made to the metabase, which schedules a write-to-disk event 60 seconds in the future.

**Explanation of How the In-Memory-Metabase Is Written to Disk**

**Step 1: Are there changes pending in the in-memory metabase?**

When a write-to-disk event occurs, IIS first checks for pending changes that occurred after the first change in the in-memory metabase. Changes that are written to the MetaBase.xml file when edit-while-running is enabled are counted as pending changes in the in-memory metabase. If there are no pending changes in the in-memory metabase, no further action is taken and the in-memory metabase is written to disk; otherwise, IIS proceeds to step 2.

**Step 2: IIS locks the in-memory metabase to prevent external writes.**

The in-memory metabase is write-locked, meaning that no writes can occur but the metabase can still be read; this prevents writes while the in-memory metabase is written to disk and ensures that a valid copy of the in-memory metabase is saved.

**Step 3: IIS determines the major version number.**
IIS increases the value of the **HistoryMajorVersionNumber** property in the in-memory metabase by one. IIS then verifies whether a file in the History folder is named with the same major version number. If a file that uses the same major version number exists in the History folder, IIS increases the value of the **HistoryMajorVersionNumber** property in the in-memory metabase again by one, and checks the History folder to determine whether the major version number is available. This process is repeated until an available major version number is found.

This method of determining the major version number prevents the possibility of overwriting a file in the History folder. Although unlikely under normal circumstances, the possibility exists that a file in the History folder has a nonsequential version number for one of the following reasons:

- An administrator changes the value of the **HistoryMajorVersionNumber** property, which causes the value of the **HistoryMajorVersionNumber** property to roll over to 1 or to the lowest available major version number.

- An administrator copies an old history file into the History folder.

**Step 4: IIS creates a temporary file that contains the in-memory metabase configuration.**
When the major version number is determined, a temporary file named Xml.tmp is created in the same folder as MetaBase.xml. The major version number that was determined in step 3 is written in the temporary file.

**Step 5: IIS unlocks the in-memory metabase to allow writes.**
After the temporary file is created, the in-memory metabase is unlocked to allow writes through ABOs.

**Step 6: IIS copies the temporary file to the History folder and renames the history file by using the major and minor version numbers.**
The temporary file is copied to the History folder and renamed in the following format, where *HistoryMajorVersionNumber*is the major version number that was determined in step 3 and the minor version number is reset to 0.

MetaBase_*HistoryMajorVersionNumber_MinorVersionNumber*.xml

**Step 7: Is MetaBase.xml write-locked?**
The MetaBase.xml file can become write-locked if edit-while-running is not enabled and an administrator uses an application to open MetaBase.xml for editing, and the editing application enforces a write-lock. An application that enforces a write-lock releases the lock when the file is closed. However, in most circumstances, MetaBase.xml is not write-locked.

If MetaBase.xml is write-locked by an application, IIS cannot write changes to the file. Instead, a history file is created and an error is sent to the event log.

**Step 8: Is MetaBase.xml newer than the temporary file?**
IIS checks to determine whether the MetaBase.xml file is newer than the temporary file. The MetaBase.xml file can be newer than the temporary file when the in-memory metabase is written to disk at about the same time that an administrator saves changes to the MetaBase.xml file when edit-while-running is enabled.

If the MetaBase.xml file is newer than the temporary file, the MetaBase.xml file is not overwritten with the temporary file. Instead, an error is sent to the event log.

**Step 9: The temporary file copies over MetaBase.xml, and is then deleted.**
If all verification tests are passed, the MetaBase.xml file is overwritten with the temporary file and the temporary file is deleted.

**Step 10: Delete oldest history file?**
IIS checks to determine whether the number of history file pairs in the History folder exceeds the value of the **MaxHistoryFiles** property. (A history file pair is a MetaBase.xml file and an MBSchema.xml history file that are named by using the same major and minor version numbers.) If the number of history file pairs exceeds the value of the **MaxHistoryFiles** property, the oldest history file pair is deleted.

# Writing the In-Memory Schema to Disk

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS can distinguish changes that are made to the in-memory metabase schema separately from changes that are made to the in-memory metabase configuration. The process of writing the in-memory metabase schema to disk is similar to writing the in-memory metabase configuration to disk, as follows:

- If there are no pending changes within the in-memory metabase schema when a write-to-disk event occurs, the in-memory metabase schema node is not written to disk. Instead, the MBSchema.xml file is copied to the History folder and renamed with the same major and minor version numbers that are used to name the metabase configuration history file in step 6 above.

- If changes are made to the in-memory metabase schema, the following occurs:

    - The MBSchema.xml file is renamed to a temporary schema file.

    - The in-memory schema is written to MBSchema.xml.

    - The temporary schema file is copied to the History folder and renamed with the same major and minor version numbers that are used in step 6 above.

# Metabase Import and Export

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The metabase import and export feature enables administrators to create an export file, named MetaBase.xml by default, which contains specifically selected elements of a metabase configuration file. This export file can later be imported into the metabase configuration file of the same computer, or to another computer running Windows Server 2003. For more information about procedures for performing metabase import and export, see Metabase Import and Export in IIS 6.0.

Additionally, you can use the metabase export feature to create a metabase template, which can be used to configure multiple computers with the same IIS configuration settings, excluding machine-specific and system-specific information. For more information about metabase templates, see Restoring a Section of the Metabase.

Metabase export does not replace the functionality of metabase backup. You use metabase backup to create a backup file of the entire metabase configuration and schema, which can be restored only in its entirety (including encrypted properties). Metabase export does not export the metabase schema.

| **Note** |
| --- |
| Neither metabase backup nor export can be used to save or copy content files, scripts, or binaries. To back up your entire server — including content files, scripts, and binaries — use Windows Backup. To simply copy these files to a new location, use Windows Explorer or the Windows **xcopy** command. Remember to reregister any components that you copy to a new computer. |

**Specifying Metabase Elements for Import or Export**

You can import and export the following metabase elements:

- All Web site configurations or a single site configuration
- All FTP site configurations or a single FTP site configuration
- All SMTP site configurations or a single SMTP site configuration
- All NNTP site configurations or a single NNTP site configuration
- An IIS application configuration
- A directory, virtual directory, or file configuration properties
- All application pool configurations or a single application pool configuration

**Importing and Exporting a Metabase Element**

To import or export a metabase element, you can use any of the following:

- IIS Manager
- The Iiscnfg.vbs tool that is included with IIS
- A custom script that uses the **Export** method of the WMI provider for IIS or the **Export** method of the ADSI provider for IIS
- A custom application that uses ABO

Regardless of whether you use a programmatic interface, a script, or IIS Manager to import and export metabase elements, there are specific choices you must make. The following list outlines both the required and optional steps needed to import and export a metabase element. The inheritance and merge options are discussed after the list. When a metabase element is

exported, IIS creates an XML file that contains the exported element.

**Before exporting a metabase element**

1. Select a metabase element to export.
2. Specify the name and location of the XML file to contain the exported element.
3. Optionally choose to export inherited properties.
4. Optionally choose to export lower-level child keys recursively (recommended).
5. Optionally encrypt the export file with a password that you provide (recommended).

**Before importing a metabase element**

1. Specify a metabase export file to import.
2. Specify an element within the export file to import.
3. Optionally choose to import inherited properties.
4. The export file must contain inherited properties for this option to be available.
5. Optionally choose to import lower-level child keys recursively.
6. The export file must contain lower-level child keys for this option to be available.
7. Specify the password, if any, that was used to encrypt the export file.
8. Optionally choose to merge the settings from the export file into the metabase configuration.
9. If you intend to import a metabase element to another computer, you must first change the machine-specific and system-specific information in the export file.

**Setting Inherit and Merge Flags**

If you set the inherit flag during a metabase export, inheritable properties that are applicable to the metabase element are saved as part of the export file. If the inherit flag is set during an import operation, these inheritable properties are imported into the metabase.

The behavior of importing inheritable properties differs, depending on whether the merge flag is used along with the inherit flag.

For example:

- If you import a metabase element by using the inherit flag without the merge flag, inheritable properties are imported into the metabase, unless the property already exists at the same location in MetaBase.xml. (If the value of an imported property differs between the export file and the MetaBase.xml file, the value of the property in the MetaBase.xml file remains unchanged.)
- If you import a metabase element by using the inherit flag with the merge flag, inherited properties are imported into the metabase and the value of the properties that already exist at the same location in MetaBase.xml are overwritten. (If the value of an imported property differs between the export file and the MetaBase.xml file, the value of the property in the MetaBase.xml file is overwritten with the value from the export file.)

**Important**

Inherited properties in the destination metabase affect the configuration of all sites. Use caution when you import a metabase element without using the inherit flag. To avoid unintended changes to other sites, import to a unique location that has no child nodes.

**Exporting Lower-Level Child keys**

When specifying a metabase element to export, you can optionally export lower-level child keys. If an element is exported by using this option, you can optionally import the element and all recursive child keys.

# Importing and Exporting Metabase Elements with IIS Manager

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use the metabase import/export feature to create an export file that contains specifically selected elements of a metabase configuration file. This exported file can later be imported into the metabase configuration file on the same computer or another computer running Windows Server 2003.

The simplest way to use the metabase import/export feature is through IIS Manager. However, if you need to use this feature on numerous machines quickly, you can use one of the following:

- The Iiscnfg.vbs tool that is included with IIS 6.0
- A custom script that uses the **Export** method of the WMI provider for IIS or the **Export** method of the ADSI provider for IIS
- A custom application that uses ABO

For more information about importing and exporting metabase elements using IIS Manager, see Exporting Metabase Elements with IIS Manager and Importing Metabase Elements with IIS Manager.

# Moving Sites and Applications to Another Computer

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use the metabase export feature to configure multiple computers with the same basic IIS configuration settings. The procedure below illustrates how to create a metabase template for the WWW service — though any service, site, virtual directory, or application can be moved.

For more information about moving sites and applications to other computers, see Moving Sites and Applications to Another Computer in IIS 6.0.

# Metabase Import/Export Behavior

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When using the metabase import/export feature, you need to be aware of the following behaviors with which the feature operates:

- If you attempt to import a property to a computer whose metabase schema file does not include the property, the import process continues, but errors are logged in Event Viewer. If you are using a script or application to perform the import, the script or application issues a warning.
- If the key type of any imported metabase element does not match the destination element, an error occurs.
- If you use a script or application that uses the inherited option, inherited settings are imported and applied to the top element. If you do not use the inherited option, only the elements that are explicitly set are imported.
- If you import configuration data by using a script or application and you use the merge flag, existing elements are not overwritten, and new ones are added. If you do not use the merge flag, existing elements are deleted or overwritten.

# Machine-Specific and System-Specific Information

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When importing a metabase element to another computer, you must change or remove machine-specific and system-specific information in the import file before you use it; otherwise IIS cannot use these values.

| **Important** |
| --- |
| Never include password properties in import files. If you use an import file with a blank password property, the new computer subsequently stores encrypted properties as plain text. |

Metabase properties that contain paths to resources in the file system might not match the paths on other computers. For example, the following properties might be written in the import file in the following way:

```
LogFileDirectory = ""
Path = "c:\inetpub\ftproot"
HttpErrors = "400,*,FILE,
   more HTTP errors
InProcessIsapiApps = "
   more ISAPI applications
ScriptMaps = ".asp,,1,GET,HEAD,POST,TRACE
   more script maps
```

Change the paths in bold above to the path used by the computer to which you want to import the metabase element. Search for other paths in your import file before using it.

Metabase properties that specify the IWAM and IUSR accounts might not match the user accounts used by other computers. If your import file contains these properties, delete or alter them so that their values do not overwrite the existing valid properties. These properties might be written in the import file in the following way:

```
<IIsWebService
   Location = "/LM/MSFTPSVC"
   WAMUserName = "AccountName"
   AnonymousUserName = "AccountName">
```

The **ServerBindings** property, which is used to configure the port number for a site, might not be set the same way on another computer. However, you might want to import these settings to the new computer so that the computers will be configured the same way. The **ServerBindings**property is set at the site level, and can be written in the import file in the following way:

```
<IIsFtpServer
   Location = "/LM/MSFTPSVC/1"
   ServerBindings = ":21:" >
<IIsWebServer
   Location = "/LM/W3SVC/1"
   ServerBindings = ":80:">
```

Metabase properties that contain ACLs cannot be moved to another computer. If your import file contains these properties, delete them so that their values do not overwrite the existing valid properties. These properties might be written in the import file in the following way:

```
AdminACL = "LongNumber"
```

Metabase properties that contain passwords cannot be moved to another computer. If your import file contains properties with passwords, delete them so that their values do not overwrite the existing valid properties. These properties might be written in the import file in the following way:

```
<IIsWebService
   Location = "/LM/W3SVC"
   WAMUserPass = "LongNumber"
   AnonymousUserPass = "LongNumber"
   LogOdbcPassword = "LongNumber">
```

# Additional Resources for the IIS 6.0 Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to the IIS 6.0 metabase.

**Related Information**

- Common IIS 6.0 Administrative Tasks for step-by-step instructions to help you complete many metabase-related tasks.
- IIS 6.0 Architecture for information about the role of the metabase in the overall architecture of IIS 6.0.

**Related topics in the IIS 6.0 Operations Guide**

- IIS 6.0 Registry Reference for information about the IncrementalSiteIDCreation registry key.
- Metabase Property Reference for information about the IIS 6.0 metabase properties.
- The Metabase Schema for information about the IIS 6.0 metabase schema.

**Related Windows Server 2003 Help Topics**
For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the**Search in title only** check box.

- "Access control" in Help and Support Center for Windows Server 2003 for information about how to configure ACLs.
- "Backing up and restoring data" in Help and Support Center for Windows Server 2003 for information about how to back up your entire operating system that include special steps that must be performed to include IIS in that backup.
- "Passwords" in Help and Support Center for Windows Server 2003 for information about strong passwords.

**Related Tools**

- The IIS MetaEdit 2.2 Utility

  MetaEdit is a metabase editing tool available for download from Microsoft. For information about how to obtain MetaEdit, see Knowledge Base article 232068.

# Managing a Secure IIS 6.0 Solution

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

Appropriate security safeguards on your Web server can help reduce or eliminate various security threats. These threats can come from malicious users as well as from well-intentioned users who accidentally gain access to restricted information or inadvertently alter important files. Internet Information Services (IIS) 6.0 provides the elements for a secure IIS solution, including a new approach to server configuration that emphasizes installing only the services that you know you will run on your server.

**In This Section**


Security Changes in IIS 6.0

Automatic Installation of Security Updates

Authentication Methods Supported in IIS 6.0

Access Control in IIS 6.0

Encryption in IIS 6.0

TCP/IP Port Filtering

Secure Code

MIME Types

Auditing Server Activity

Standards Implemented

Additional Resources for IIS 6.0 Security

# Security Changes in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To improve the security of your Web server, many aspects of IIS 6.0, including default behavior and settings, function differently than in earlier versions of IIS. Some of the most notable changes were made to take a more proactive stance against malicious users and attackers. A significant change is that IIS is not installed by default on Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition operating systems, and many services and features of IIS are not installed or enabled by default when you install IIS. Other security changes in IIS 6.0 affect components of Active Server Pages (ASP), authentication, and access control methods. As a result of these changes, some existing applications and sites might require you to enable services, change settings, or make other adjustments before they run as expected. However, if you change default settings, you should do so carefully to maintain the most secure solution possible.

The most significant security-related changes are as follows:

- IIS installs in a locked-down mode.
- Restrictive Multipurpose Internet Mail Extensions (MIME) types reduce the attack surface of IIS.
- Multiple worker processes affect Internet Server API (ISAPI) filter status display.
- ASP and ASP.NET functionality are disabled by default.
- Parent paths are disabled by default.
- Global.asa events are run as anonymous user.
- Anonymous password synchronization is disabled by default.
- Advanced Digest authentication requires Windows Server 2003.
- Microsoft® .NET Passport authentication requires LocalSystem user account rights.
- Kerberos authentication requires service principal names (SPNs) for multiple worker processes.
- Access is restricted for executables.
- Access is restricted for non-default identities for Common Gateway Interface (CGI) processes.

The following sections describe the security-related changes and provide information about how to customize your IIS work environment.

# IIS Installs in a Locked-Down Mode

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you install IIS 6.0, it is locked down — only request handling for static Web pages is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. None of the features that sit on top of IIS are turned on by default, including ASP, ASP.NET, CGI scripting, FrontPage® 2002 Server Extensions from Microsoft, and Web Distributed Authoring and Versioning (WebDAV). This locked-down state minimizes the attack surface that is available to intruders, who sometimes target computers by attacking services that are running but that are unused. These attacks can happen if, for example, an administrator forgets to turn off an unused service, and then to maintain it with current hotfixes, service packs, and security updates. Over time, the service might become increasingly vulnerable to attackers.

IIS 6.0 installation and service-enabling features simplify administration and management of IIS services for security purposes. For example:

- You can install and enable services when you need them.
- If you need additional services, you can enable them through the Web Service Extensions node in IIS Manager.
- When enabled features are no longer required, you can disable them through the Web Service Extensions node in IIS Manager.

| **Important** |
| --- |
| It is possible to use the Web Service Extensions node to allow all unknown ISAPI and CGI extensions, but this presents a security risk and is not recommended. This option is available for use in development and test environments, but it should never be enabled on a production server. |

**IIS Lockdown Tool Capabilities Are Built-in**

You cannot install the IIS Lockdown Tool in IIS 6.0. However, to prevent the WWW service from being disabled after an upgrade, you can run the IIS Lockdown Tool prior to upgrading from an earlier version of IIS. When you upgrade from a server running Microsoft Windows® 2000 Server and IIS 5.0, IIS 6.0 runs in IIS 5.0 isolation mode by default, and IIS implements the following security measures:

- UrlScan functionality is left unchanged. UrlScan is a tool that reduces the attack surface of Web servers running earlier versions of IIS. By default, IIS 6.0 has features that significantly improve security by reducing the attack surface of the Web server. UrlScan provides flexible configuration for advanced administrators, while maintaining the improved security in IIS 6.0. When you need this flexibility in configuring your Web server, you can run UrlScan on IIS 6.0.
- Web service extensions that are mapped to 404.dll are mapped to their original location and added to the Web service extensions list with a prohibited status.
- Access control lists (ACLs) for WebDAV are reset correctly, and WebDAV is added to the Web service extensions list with a prohibited status.
- Indexing Service is left disabled, and existing ACLs are unchanged.

# Restrictive MIME Types Reduce the Attack Surface of IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To help prevent attackers from sending files that could compromise the security of your Web server, the IIS static file handler serves only the files that are explicitly listed in the list of Multipurpose Internet Mail Extensions (MIME) types. IIS is preconfigured to recognize a default set of global MIME types that are recognized by all of the Web sites that you create with IIS.

If the default list of registered MIME types is too restrictive and you want to allow IIS to serve other file types, you can add allowable MIME types on a global basis or for a specific Web site or directory.

For more information about MIME types, see Working with MIME Types.

**Allowing All Extensions**

In IIS 5.0, if a request fell through the static file handler, IIS served the file, regardless of its extension. IIS 6.0 does not automatically serve file requests that fall through the static file handler. Although it is not generally recommended as a best security practice, you might want to allow all file extensions in some scenarios — for example, if you want to serve files that have ascending extensions such as .001, .002, and so on. If you want to serve all files, you can add a wildcard character (*) to the MIME types list. Files that do not have extensions can be served only if the wildcard character is added to the list.

You can add a wildcard character through IIS Manager, or you can add one programmatically by using the Adsutil.vbs command-line script. For information about adding a wildcard character through IIS Manager, see Creating Global MIME Types and Creating MIME Types for a Web site or directory.

To add a wildcard programmatically by using Adsutil.vbs, use the customizable text file provided in Listing 5.1.

**Important** You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters).

**Listing 5.1 Sample Script for Adding Wildcards**

```
c:\inetpub\adminscripts\adsutil.vbs set /w3svc/MimeMap ".*,common/type"
Sample
' VBScript source code
Const ADS_PROPERTY_UPDATE = 2
Set MimeMapObj = GetObject("IIS://LocalHost/MimeMap")
aMimeMap = MimeMapObj.GetEx("MimeMap")
' Add a new mapping.
i = UBound(aMimeMap) + 1
Redim Preserve aMimeMap(i)
Set aMimeMap(i) = CreateObject("MimeMap")
aMimeMap(i).Extension = ".*"
aMimeMap(i).MimeType = "common/type"
MimeMapObj.PutEx ADS_PROPERTY_UPDATE, "MimeMap", aMimeMap
MimeMapObj.SetInfo
```

For more information about using adsutil.vbs, see Using the Adsutil.vbs Administration Script.

# Multiple Worker Processes Affect ISAPI Filter Status Display

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

As a routine security practice, you should monitor the status of your ISAPI filters to ensure that they are functioning correctly. In IIS 6.0, you can use IIS Manager to check the ISAPI filter status. Because multiple worker processes are supported in IIS 6.0, be careful when you interpret the status as it is displayed. For example, if a red arrow appears, that indicates that the last ISAPI filter failed to load; it does not indicate that all of the ISAPI filters for all of the worker processes failed.

For more information about checking the status of ISAPI filters, see Checking ISAPI Filter Status.

# ASP-Related Security Changes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

There are several important changes to the way IIS interacts with ASP that improve security and reduce the attack surface of IIS.

**ASP and ASP.NET Are Disabled by Default**
In IIS 6.0, ASP and ASP.NET are disabled by default. If you want to serve ASP or ASP.NET content from IIS 6.0, you must first enable ASP or ASP.NET in the Web Service Extensions node.

| **Note** |
| --- |
| If your content requires ISAPI, CGI, WebDAV, or other extensions, you must enable those extensions as well. |

For more information about enabling ASP or ASP.NET, see Enabling ASP Pages in IIS 6.0 and Enabling ASP.NET.

**Parent Paths Are Disabled by Default**
When parent paths are enabled, ASP pages allow you to include files with paths that are relative to the current directory (using the ..\ notation). This constitutes a potential security risk because a server-side include (SSI) path can provide access to critical or confidential files that are outside the root directory of the application. When parent paths are enabled, a Web server is more susceptible to *traversal attacks*, which allow a malicious user to execute commands or view data that is outside the intended target path.

In IIS 6.0, because parent paths are disabled by default, an existing application that has a Web page that contains the #include SSI directive and uses ..\ notation to refer to a parent directory will not work unless you enable parent paths. The parent paths correspond to the **AspEnableParentPaths** metabase property. For more information about enabling parent paths, see Configuring ASP Applications. For more information about SSI directives, see Using Server-Side Include Directives.

**Global.asa Events Are Run As Anonymous User**
In earlier versions of IIS, ASP executes events in the security context (or user identity) of the host process because during ASP-executed events there is no user context. Global.asa executes events in earlier versions of IIS as follows:

- **Application_OnStart** Impersonated user
- **Application_OnEnd** Host identity
- **Session_OnStart** Impersonated user
- **Session_OnEnd** Host identity

In IIS 6.0, ASP executes the Global.asa events **Application_OnEnd** and **Session_OnEnd** anonymously by default. This change prevents user code from running as the base identity of the worker process. For example, if **Application_OnEnd** and **Session_OnEnd** run under the host identity context, and an Internet service provider (ISP) allows users to upload Global.asa files, a malicious user might be able to use the **FileSystemObject** to access files that are otherwise inaccessible in the anonymous context.

If you want to change this default setting, use the **AspRunOnEndAnonymously** metabase property. For more information about this setting, see the AspRunOnEndAnonymously Metabase Property.

# Security-Related Changes in Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 introduces a number of changes to user authentication. These changes include new default settings and requirements for different authentication methods.

**Anonymous Password Synchronization is Disabled by Default**

- Anonymous password synchronization requires the IIS subauthentication component, Iissuba.dll, which was installed by default in earlier versions of IIS. The subauthentication component allows IIS to authenticate an anonymous user even if the password in the IIS metabase does not match the password in the Windows local account database. Because this involves a security risk, the subauthentication component is not installed by default in IIS 6.0. However, some administrators need to use anonymous password synchronization, in, for example, large ISP settings where thousands of anonymous IIS accounts are used to isolate Web sites.

- For more information about enabling subauthentication and anonymous password synchronization, see Configuring Subauthentication on a New Installation of IIS 6.0.

**Advanced Digest Authentication Requires Windows Server 2003**
To use Advanced Digest authentication, the domain infrastructure and the server running IIS must both be running Windows Server 2003. If you upgraded to IIS 6.0, the default digest authentication is Digest authentication. If you performed a clean installation of IIS 6.0, the default digest authentication is Advanced Digest authentication. For more information about Advanced Digest authentication, see Advanced Digest Authentication.

**Intermediate Certificates Are Not Downloaded to Complete a Certificate Chain**
When IIS receives an incoming SSL request from a client, IIS attempts to build its certificate chain before sending its certificate information back to the client. In IIS 6.0, missing intermediate certificates are not automatically downloaded by default, which was the behavior for previous versions of IIS. To ensure that client requests for SSL-enabled content can be served, you must manually add the intermediate certificate to the Intermediate Certification Authorities folder on the local computer. In addition, every time a certificate in the external chain is renewed or modified, you must manually update the certificate store. For information about manually updating the certificate store, see Client receives 403.16 error when IIS cannot process a complete certificate chain.

This change affects the following scenarios:

- IIS servers in workgroups, because no Group Policy is available to populate the cache.

- Cross Certificates, because the certificates might not be published into the Directory Service so Group Policy can populate the cache.

- IIS servers that are set to explicitly trust a foreign Certification Authority (CA).

**.NET Passport Authentication Requires LocalSystem User Account Rights**
Microsoft .NET Passport authentication works only when the account that is used for the IIS worker process has *trusted computing base*(TCB) rights, which means that the account acts as part of the operating system. Only the LocalSystem user account has TCB rights by default. You can add TCB rights to an account by using the Local Security Settings snap-in. However, because a worker process running with TCB rights acts as part of the operating system, this account has a very high level of access, which presents a security risk.

**Kerberos Authentication Requires SPNs for Multiple Worker Processes**
You can isolate Web sites that are in different worker processes and running under different identities; however, unexpected IIS behavior might occur if you use Integrated Windows authentication. Integrated Windows authentication attempts to use Kerberos authentication, which might not work, depending upon the identity of the worker process. To use Kerberos authentication, a service must register its Service Principal Name (SPN) in the account in the Microsoft Active Directory® directory service under which the service is running. By default, Active Directory registers the NetBIOS name or computer name and allows the Network Service or LocalSystem user accounts to use Kerberos. However, for Kerberos to work with the following configurations, you must first register an SPN:

- If your site is referenced by a Windows Internet Name Service (WINS) or Domain Name System (DNS) name that is different from the computer name of the server running IIS (a Web farm, for example), authentication defaults to NTLM (which is also known as Microsoft Windows NT® Challenge/Response authentication).

- If you configure the worker process for the Web site to run as a domain account, and you did not use a WINS or DNS name, authentication will fail. This is because the SPN can be found, but it is registered for the Network Service and LocalSystem user accounts, not for the domain account under which the worker process is running.

You can use the Setspn.exe command-line tool to register an SPN on the account under which the worker process is running. You must be a domain administrator to set an SPN. The Setspn.exe tool is included in Resource Kit Tools for Windows Server 2003, available on the *Windows Server 2003 Deployment Kit* companion CD, or on the Web at http://www.microsoft.com/reskit.

**Forcing NTLM**

When you isolate Web sites on a virtual directory level by configuring worker process identities as different domain accounts, Kerberos fails. If you are using Integrated Windows authentication, are not using a WINS or DNS name for the server running IIS, and you want to use a local user account or the Local Service user account as a worker process identity, Kerberos authentication fails because Active Directory does not trust the accounts.

To work around these issues, you can force IIS to use NTLM authentication by setting the NTAuthenticationProviders Metabase Property to NTLM.

# Security-Related Changes in Access Control

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 improves security by changing access control requirements for executables and CGI processes. These changes can affect existing applications that might require less restrictive permissions.

**Access Is Restricted for Executables**
Windows Server 2003 requires a user to be a member of the Administrators group to run most command-line executables in the System folder, so remote access is limited to administrators. This might affect you if your application uses command-line tools from a Web page. To work around this issue, change the ACL for the executable by using an ACL editing tool, such as the **cacls** command. For more information about ACLs and ACL editing tools, see Access Control in IIS 6.0.

**Access Is Restricted for Nondefault Identities for CGI Processes**
IIS 6.0 worker processes use the **CreateProcessAsUser** API to start CGI processes. The **CreateProcessAsUser** API must have the SE_ASSIGNPRIMARYTOKEN_NAME and SE_INCREASE_QUOTA_NAME user rights to succeed. The Network Service, Local Service, and LocalSystem user accounts have these user rights. If you change the identity of a worker process and want to run CGI processes, ensure that the new identity has these two user rights. You can assign user rights using the Local Security Settings snap-in in Windows Server 2003.

For more information about configuring user rights for identities that run CGI processes, see Configuring User Rights for Nondefault Identities to Run CGI Processes.

# Automatic Installation of Security Updates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

As a security best practice, keep IIS up-to-date with the current security updates that Microsoft releases. With Windows Server 2003 Automatic Updates, you can install updates automatically so that the most critical code is updated on your computer, based on a schedule that you specify. You do not need to be logged on to the computer for the installation to occur. By minimizing the time that you spend logged on to your computer as a member of the Administrators group, your computer is less vulnerable to attack.

The only IIS updates that are available for scheduled automatic installation are critical security-related hotfixes. Because these hotfixes might require restarting the operating system, you can delay or change the time of the scheduled installation. You must be logged on as an administrator to change the scheduled installation time.

When an update is available for automatic installation, two events occur:

- At the scheduled installation time, the Priority 0 Critical security hotfix is installed.
- After the scheduled installation takes place, an event is logged notifying you that the hotfix has been installed successfully.

For information about scheduling automatic updates, see To schedule automatic updates in Microsoft Windows Server 2003 documentation.

# Authentication Methods Supported in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Using authentication with IIS is critical if your Web server contains resources that are not meant for anonymous or public access, but approved users must be able to access the Web server over the Internet. For example, if you run Microsoft® Outlook® Web Access (OWA) or the Microsoft Terminal Services Advanced Client on your Web server, you must configure an appropriate authentication method to control access to those applications. You can require users to confirm their identity by providing a valid Microsoft® Windows® name and password before they access any information on your Web server. This identification process is called *authentication*, and, like many features of IIS, authentication can be set at the Web site, directory, or file level.

IIS 6.0 supports the following authentication methods:

- Anonymous authentication. This authentication method allows everyone access to the public areas of your Web sites, without asking for a user name or password.
- Basic authentication. This authentication method prompts the user for a user name and a password, also called *credentials*, which are sent unencrypted over the network.
- Digest authentication. This authentication method operates much like Basic authentication, except that passwords are sent across the network as a hash value for additional security. Digest authentication is available only on domains with domain controllers running Windows server operating systems.
- Advanced Digest authentication. This authentication method is identical to Digest authentication, except Advanced Digest authentication stores the client credentials as a Message Digest (MD5) hash in Active Directory on your domain controller running Windows Server 2003.
- Integrated Windows authentication. This authentication method uses hashing technology to scramble user names and passwords before sending them over the network.
- UNC authentication. This authentication method passes users' credentials through to the computer with the Universal Naming Convention (UNC) share.
- .NET Passport Authentication. This authentication method is a user-authentication service that lets Web site users create a single sign-in name and password for access to all .NET Passport–enabled Web sites and services. .NET Passport–enabled sites rely on the .NET Passport central server to authenticate users; .NET Passport–enabled sites do not host and maintain their own proprietary authentication systems.
- Certificate authentication. This authentication method uses Secure Sockets Layer (SSL) certificates to authenticate servers and clients.

Table 5.2 compares the Web site authentication methods that are discussed in this chapter.

**Table 5.3 Comparison of Web Site Authentication Methods**

| Method | Security Level | How Passwords Are Sent | Crosses Proxy Servers and Firewalls | Client Requirements |
|---|---|---|---|---|
| Anonymous authentication | None | N/A | Yes | Any browser |
| Basic authentication | Low | Base64 encoded clear text | Yes, but sending passwords across a proxy server or firewall in clear text is a security risk because Base64 encoded clear text is not encrypted. | Most browsers |

| | | | | |
|---|---|---|---|---|
| Digest authentication | Medium | Hashed | Yes | Internet Explorer 5 or later |
| Advanced Digest authentication | Medium | Hashed | Yes | Internet Explorer 5 or later |
| Integrated Windows authentication | High | Hashed when NTLM is used; Kerberos ticket when Kerberos is used. | No, unless used over a PPTP connection | Internet Explorer 2.0 or later for NTLM; Windows 2000 or later with internet Explorer 5 or later for Kerberos |
| Certificate authentication | High | N/A | Yes, using an SSL connection | Internet Explorer and Netscape |
| .NET Passport authentication | High | Encrypted | Yes, using an SSL connection | Internet Explorer and Netscape |

If you configure multiple authentication methods, IIS attempts to negotiate the most restrictive method first, and then it works down the list of available authentication protocols until finding an authentication method that is supported by both the client and the server.

You can also provide security for FTP sites by using one of the following authentication methods:

- Anonymous FTP authentication. This authentication method gives users access to the public areas of your Web site without prompting them for a user name and password.

- Basic FTP authentication. This authentication method requires users to log on with a user name and password corresponding to a valid Windows user account.

You can use these authentication methods to grant access to public areas of your site while minimizing the risk of unauthorized access to your private files and directories. Each authentication method is covered in detail in the following sections.

# Anonymous Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Anonymous authentication gives users access to the public areas of your Web or File Transfer Protocol (FTP) site without prompting the users for a user name or password. When a user attempts to connect to your public Web or FTP site, your Web server assigns the connection to the Windows user account that is specified for anonymous connections. By default, this account is named IUSR_*ComputerName*, where *ComputerName* is the name of the computer on which IIS is running, and it is included in the Windows user group, Guests. This group has security restrictions, which are imposed by NTFS file system permissions, that designate the level of access and the type of content that is available to public users.

| **Important** |
| --- |
| Do not confuse the anonymous user context that Windows Server2003 provides with the anonymous user account that IIS provides. The IIS anonymous user account is configurable; the default is IUSR_*ComputerName*. The operating system also provides a notion of anonymous access across the network. |

When you use Anonymous authentication, the IUSR_*ComputerName* account is added to the Guests group on the computer running IIS during setup. When IIS receives a request, it impersonates the IUSR_*ComputerName* account before executing any code. IIS can impersonate the IUSR_*ComputerName* account because it recognizes the user name and password for this account. Before returning a page to the client, IIS checks NTFS and directory permissions to see whether the IUSR_*ComputerName* account is allowed access to the file. If access is allowed, the access process, which is also called authorization, completes, and the resources are made available to the user. If access is not allowed, IIS returns a 401.3 error message to the client.

You can change the account that is used for Anonymous authentication in IIS Manager, either at the Web server service level or for individual virtual directories and files. However, be aware that when you change the IUSR_*ComputerName* account, the changes affect every anonymous HTTP request that a Web server services, so use caution if you modify this account.

If you enable Anonymous authentication, IIS always attempts to authenticate the user with Anonymous authentication first, even if you enable additional authentication methods. If the anonymous user account does not have permission to access a specific file or resource, your Web server will not establish an anonymous connection for that resource.

In IIS 6.0, NETWORK_CLEARTEXT is the default logon type for Anonymous authentication. As a result, Anonymous authentication no longer requires the Allow log on locallyuser right. For more information about the NETWORK_CLEARTEXT logon type, see Basic Authentication.

# Sub-Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In earlier versions of IIS, sub-authentication was enabled by default. The sub-authentication component, Iissuba.dll, enables IIS to manage passwords on anonymous accounts. Because using this component involves a security risk, sub-authentication is not enabled by default when you install Windows Server 2003 with IIS 6.0. However, you can use sub-authentication to manage passwords for anonymous accounts by meeting the following requirements:

- For applications you grant anonymous access, the worker process runs as LocalSystem.
- The sub-authentication component, Iissuba.dll, is registered.
- The AnonymousPasswordSync Metabase Property is enabled (set to **true**).

The actions taken to meet the above requirements are different for clean installs of IIS 6.0 and upgrades to IIS 6.0 from installations of IIS with sub-authentication configured.

**Configuring Sub-Authentication on a New Installation of IIS 6.0**
By default, on a new installation, IIS 6.0 runs in worker process isolation mode, and sub-authentication is disabled (**AnonymousPasswordSync** is set to **false**).

If after a new installation of IIS 6.0, you switch to IIS 5.0 isolation mode, then by default, the worker process assigned to in-process applications will run as LocalSystem.

**Configuring Sub-Authentication on an Upgrade to IIS 6.0**
When you upgrade a server to IIS 6.0 from an earlier version of IIS that uses sub-authentication to manage passwords on anonymous accounts, sub-authentication is enabled by default (**AnonymousPasswordSync** is set to **true**). However, sub-authentication will not work because two configuration tasks have not been accomplished: Iissuba.dll is not registered and the worker processes using Anonymous authentication are not running as LocalSystem. Your event log should have entries on this.

# Basic Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Basic authentication method is a widely used industry-standard method for collecting user name and password information. When you use Basic authentication, the browser displays a dialog box into which users are required to enter a previously assigned Windows account user name — which includes a Windows domain name, for example, Domain1\User1 — and password, which are also known as *credentials*. The browser then attempts to establish a connection to a server using the user's credentials. The plaintext password is Base64-encoded before it is sent over the network.

One advantage of Basic authentication is that it is part of the HTTP specification and is supported by most browsers. However, Base64 encoding is not encryption. Web browsers that use Basic authentication transmit passwords in an unencrypted form. If a Base64-encoded password is intercepted over the network by a network sniffer, unauthorized users can easily decode and reuse the password. Therefore, Basic authentication is not recommended unless you are confident that the connection between the user and your Web server has been secured, for example, with a dedicated line or an SSL connection. For more information about SSL, see SSL and Certificates.

If a user's credentials do not correspond to a valid Windows user account, Internet Explorer displays a dialog box into which the user re-enters his or her credentials. Internet Explorer — not IIS — allows the user three connection attempts before it terminates the connection and reports an error to the user. If a user's credentials correspond to a valid Windows user account, a connection is established.

Basic authentication does not automatically configure your Web server to authenticate users; to require authentication, you must also disable anonymous access. As a security best practice, you should create Windows user accounts with NTFS permissions correctly set. For step-by-step instructions for configuring Basic authentication, see Basic Authentication in IIS 6.0.

**Logon Types**

In IIS 6.0, the default logon type for Basic authentication (and Anonymous authentication) is NETWORK_CLEARTEXT. This is a change from earlier versions of IIS, in which the default logon type is INTERACTIVE. As a result of this change, when you use the default logon type for Basic authentication, users no longer need interactive logon rights. Basic authentication works with domain controllers, and the NETWORK and NETWORK_CLEARTEXT settings no longer require logon rights. Table 5.4 lists the logon types that are available for Basic authentication and Anonymous authentication.

**Table 5.4 Logon Types for Basic Authentication and Anonymous Authentication**

| Logon Type | LogonMethod Setting | Logon Right Required | Security Identifier (SID) Added to Access Token | Outbound Credentials |
|---|---|---|---|---|
| NETWORK_ CLEARTEXT (default) | 3 - MD_LOGON_NETWORK_ CLEARTEXT | Network | NT AUTHORITY\NETWORK_ CLEARTEXT | Yes |
| NETWORK | 2 - MD_LOGON_NETWORK | Network | NT AUTHORITY\NETWORK | No |
| BATCH | 1 - MD_LOGON_BATCH | Batch | NT AUTHORITY\BATCH | Yes |
| Interactive | 0 - MD_LOGON_INTERACTIVE | Interactive | NT AUTHORITY\INTERACTIVE | Yes |

For more information, see the LogonMethod Metabase Property.

**Token Cache Security Considerations**

When you use Basic authentication, user tokens are cached in the token cache. By default, tokens remain in the cache for 15 minutes. If you log on using Basic authentication with an account that has a high level of user logon rights, a successful attacker could use the account to gain access to the resources on your computer. There are several ways to help minimize this threat:

- Do not log on, or allow anyone to log on, using Basic authentication with an account that has a high level of user logon rights.
- Disable user token caching by setting the global registry entry, UserTokenTTL, to zero (0).

**Caution**

Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference.

- Change UserTokenTTL to less than the default of 15 minutes. Note that time is expressed in seconds for this metabase property. For more information about the UserTokenTTL global registry entry, see Global Registry Entries.

**Note**

If you use Basic authentication and disable the token cache, IIS must log on for every request made — potentially thousands for a requested Web page. This will have a negative effect on performance. For more information, see Balancing Performance and Security.

**Related Topics**

For information about editing the registry, see the Registry Reference.

# Digest Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Digest authentication offers the same functionality as Basic authentication; however, Digest authentication provides a security improvement because a user's credentials are not sent across the network in plaintext. Digest authentication sends credentials across the network as a Message Digest 5 (MD5) hash, which is also known as the MD5 message digest, in which the credentials cannot be deciphered from the hash.

If your server is running an earlier version of IIS with Digest authentication enabled and you upgrade to IIS 6.0, Digest authentication remains the default authentication method. In all other instances, you must enable Digest authentication.

**Client Authentication Process for Digest Authentication**

Figure 5.1 shows how a client is authenticated using Digest authentication. The steps that follow the figure describe the process in more detail.




1. The client requests a file from the server running IIS.
2. The server running IIS denies the request and sends the client the following information:
   a. Digest authentication is being used.
   b. The realm name.
3. Internet Explorer prompts the user for credentials (user name and password). Internet Explorer then combines these credentials with the name of the realm to create an MD5 hash and resubmits the request for the file from the server running IIS, this time sending the MD5 hash.
4. The server running IIS receives the hash and sends the client's hash to the domain controller for verification.
5. The domain controller informs the server running IIS of the authentication results.
6. If the client is authenticated, IIS sends the requested document or data to the client.

| **Important** |
| --- |
| Digest authentication completes only if the domain controller has a reversibly encrypted (clear text) copy of the requesting user's password stored in Active Directory. To allow passwords to be stored in clear text, you need to activate the **Store password using reversible encryption** setting on the **Account** tab of the user in Active Directory. Alternatively, you can set a Group Policy to enable this capability. After making this setting you need to set a new password to activate this feature because the old password cannot be determined. For more information about setting user properties in Active Directory, see Help and Support Center for Windows Server 2003. |

**Requirements for Digest Authentication**

You do not need to install additional client software to use Digest authentication, but Digest authentication relies on the

HTTP 1.1 protocol, as defined in RFC 2617 HTTP Authentication: Basic and Digest Access Authentication, and not all browsers support that protocol. If a non–HTTP 1.1–compliant browser requests a file from a server that uses Digest authentication, the server requests that the client provide Digest authentication credentials. The non–HTTP 1.1–compliant browser rejects the request because the client cannot support Digest authentication. For more information about HTTP authentication, see RFC 2617, HTTP Authentication: Basic and Digest Access Authentication on the Internet Engineering Task Force Web site.

Before you enable Digest authentication on your server, ensure that:

- All of the clients that need access to a resource that is secured with Digest authentication are using Internet Explorer 5.*x* or Internet Explorer 6.
- The user and the server running IIS are members of, or are trusted by, the same domain.
- Users have a valid Windows user account that is stored in Active Directory on the domain controller, with reversible encryption enabled.
- The domain has a domain controller running the Microsoft® Windows® 2000 Server operating system or Windows Server 2003.

  **Important**

  If the authenticating domain controller is on a computer running Windows2000 Server, IIS6.0 requires subauthentication for Digest authentication to work. For more information about subauthentication, see Using Sub-Authentication with Digest Authentication later in this topic.

- The server running IIS 6.0 is running Windows 2000 Server or Windows Server 2003.
- If the server operates in worker process isolation mode, the application pool that contains an application that uses Digest authentication must run with LocalSystem as the worker process identity. This requirement applies because a user password is not stored outside the hash that the client sends to IIS, and a worker process that is associated with the application cannot impersonate the user by using the user name and password.
- The Iissuba.dll file is registered on the server running IIS.

**Note**

Only domain administrators can verify that the domain controller requirements are met. Check with your domain administrator if you are unsure about whether your domain controller meets the preceding requirements.

Enabling Digest authentication on a server running IIS requires the following two tasks:

- Enabling Digest authentication for Windows domain servers. For more information about enabling Digest authentication for Windows domain servers, see Help and Support Center for Windows Server 2003.
- Configuring the realm name at any level of the metabase by using IIS Manager or by scripting.

**Note**

If Basic authentication is enabled for the site, virtual directory, or folder you are configuring, the **Default domain** box will also be available. However, only **Realm** is meaningful to Digest authentication.

**Configuring the Realm Name**

Table 5.4 lists and describes the levels of the metabase where the realm name can be configured. If a child key is not specifically configured, it inherits its configuration from the next configured level up.

**Table 5.5 Metabase Levels Where the Realm Name Can Be Configured**

| Metabase Level | Description |
|---|---|

| | |
|---|---|
| W3SVC | The W3SVC level, also known as the **IISWebService** level, is the highest level in the metabase where Digest authentication can be configured. Configurations set at this level are inherited by lower levels that do not have specific configuration settings. |
| W3SVC/*n* | The W3SVC/*n* level, also known as the **IISWebService** level, is a specific Web site, where *n* is the number of the site. Sites are numbered starting at 1. The default Web site is 1. |
| W3SVC/*n*/root | The W3SVC/*n*/Root level, known as the **IISWebVirtualDir** level, is the starting point for a Web Site, where *n* is the number of the site. |
| W3SVC/*n*/root/vdir | The W3SVC/*n*/ROOT/WebVirtualDir level, known as the **IISWebVirtualDir** level, is a virtual directory within a Web Site, where *n* is the number of the site. |
| W3SVC/*n*/root/vdir/webdir | The W3SVC/*n*/Root/WebVirtualDir/WebDir level, also known as the **IISWebDirectory** level, is a physical directory within a virtual directory within a Web site, where *n* is the number of the site. |
| W3SVC/*n*/root/vdir/file | The W3SVC/*n*/root/vdir/file level is an individual file within the W3SVC/*n*/ROOT/WebVirtualDir level, where *n* is the number of the site. |
| W3SVC/*n*/root/webdir/file | The W3SVC/*n*/root/vdir/file level is an individual file within the W3SVC/*n*/Root/WebVirtualDir/WebDir level, where *n* is the number of the site. |

You can configure one or multiple realm names on a server running IIS. For example, you might want to configure multiple realm names to allow members of Domain1 to access the Sales directory, and members of Domain2 to access the Engineering directory. This is particularly useful if no trust relationship exists between Domain1 and Domain2. If you configure multiple realm names, they must be configured at different levels of the metabase. For more information about domains, see Help and Support Center for Windows Server 2003.

If a child key in the metabase is not configured with a realm name, that child key inherits the realm name from the next parent key that has the realm name configured. If the realm name is not configured, IIS sends its own computer name as the realm name. If IIS sends its own name as the realm name and IIS is not running on a Windows Server 2003 domain controller with Active Directory, Digest authentication will fail.

For step-by-step instructions for configuring Digest authentication, see Digest Authentication in IIS 6.0.

**Using Sub-Authentication with Digest Authentication**

To use Digest authentication in IIS 6.0 when the domain controller is running Windows 2000 Server, you must enable sub-authentication. The sub-authentication file, Iissuba.dll, is automatically copied to the %Windir%\System32 folder when you install Windows 2000 Server and Windows Server 2003. However, by default, sub-authentication is not enabled in IIS 6.0. To enable sub-authentication, you must perform the following tasks:

- Install the sub-authentication component, iisuba.dll.
- Set the **UseDigestSSP** metabase property to false.

| **Note** |
|---|
| The W3SVC service must be restarted before changes to this property take effect. |

- Set the identity of the application pool to LocalSystem.

For more information about how to configure sub-authentication, see Configuring Subauthentication.

Logging on with a user principal name (UPN) in the form username@domain.com will still fail at this point, even though your IIS 6.0 server is a member of a Windows 2000 domain and is correctly configured for Digest authentication as well as subauthentication.. This failure occurs because Digest authentication accepts only domain logons in one of the following formats:

- If a user logs on to a server in the same domain where the user's account resides, the username alone is sufficient for authentication.

- For logons with a user account from another domain, use the domain\username format.

**Understanding the MD5 Hash**

An MD5 hash is used for sending encrypted user credentials across a network within an HTTP header. An MD5 hash, also known as the MD5 message-digest, is created by an HTTP 1.1–compliant browser such as Internet Explorer 5 and above, using the MD5 message-digest algorithm as defined in RFC 1321, The MD5 Message-Digest Algorithm.

| **Note** |
| --- |
| The MD5 hash is a security improvement over base64 encoded clear text passwords. An unauthorized person can easily intercept base64 encoded passwords by using a network sniffer, and decoding the passwords is a trivial matter. A user name and password that is encrypted using the MD5 message-digest algorithm cannot feasibly be decrypted from the hash. |

An MD5 hash contains a user name, password, and the name of the realm. The realm is the domain that will authenticate or reject the users credential. The users credential is the password that is encrypted within the MD5 hash.

**MD5 Hash Properties**

An MD5 hash consists of a small amount of binary data, typically no more than 160 bits, and is sent by across the network within an HTTP header. All hash values share the following properties:

- **Hash length.** The length of the hash value is determined by the type of algorithm used, and its length does not vary with the size of the message. The most common hash value lengths are either 128 or 160 bits.

- **Non-discoverability.** Every pair of nonidentical messages will translate into a completely different hash value, even if the two messages differ only by a single bit. Using today's technology, it is not feasible to discover a pair of messages that translate to the same hash value.

- **Repeatability.** Each time a particular message is hashed using the same algorithm, the exact same hash value will be produced.

- **Irreversibility.** All hashing algorithms are one-way. Given a hash value, it is not feasible to discover the password. In fact, none of the properties of the original message can be determined given the hash value alone.

For more information about the MD5 message-digest algorithm, see RFC 1321, The MD5 Message-Digest Algorithm, on the Internet Engineering Task Force Web site.

# Advanced Digest Authentication

Updated: August 14, 2013

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Advanced Digest authentication is available only on Windows Server 2003. When Advanced Digest authentication is enabled, user credentials are stored on the domain controller as an MD5 hash. Advanced Digest authentication does not require that credentials are stored using reversible encryption. Instead, Advanced Digest authentication stores a few precalculated hashes in Active Directory, so user passwords cannot feasibly be discovered by anyone with access to the domain controller, including the domain administrator.

When you perform a clean installation of IIS 6.0, Advanced Digest authentication is enabled as the default authentication method, and it is recommended over Digest authentication for the following reasons:

- Subauthentication is not required for Advanced Digest authentication, because Advanced Digest authentication uses the Windows Security Support Provider Interface (SSPI) traditional implementation for security.
- Reversible password encryption is not required for Advanced Digest authentication. In Windows Server 2003, the Active Directory extended schema properties ensure that every newly created user account automatically has the Advanced Digest authentication hashed and stored as a field in the AltSecId property of the user object.
- A worker process with an application that uses Advanced Digest authentication does not have to run with LocalSystem as the process identity, because subauthentication is not required.

**Client Authentication Process for Advanced Digest Authentication**

Figure 5.2 shows how a client is authenticated using Advanced Digest authentication, and the steps that follow describe the process in more detail.




1. The client browser requests a file from the server running IIS.
2. The server running IIS denies the initial request and sends the client the following information:
   - Digest authentication is being used, because the same Digest authentication algorithm is used between the server running IIS and the client for both Digest and Advanced Digest authentication.
   - The realm name.
3. Internet Explorer prompts the user for credentials (user name and password). Internet Explorer then combines these credentials with the name of the realm to create an MD5 hash and resubmits the request for the file to the server running IIS, this time also sending the MD5 hash in the header of the HTTP request.
4. The server running IIS receives the client's hash and sends it to the domain controller for verification.
5. The domain controller compares the client's hash to the copy stored in Active Directory. If the hash values match, the domain controller informs the server running IIS that the client is authenticated.

6. The server running IIS sends the requested file to the client.

**Requirements for Advanced Digest Authentication**

As with Digest authentication, you do not need to install additional client software to use Advanced Digest authentication. Because Advanced Digest authentication relies on the HTTP 1.1 protocol as defined the RFC 2617, HTTP Authentication: Basic and Digest Access Authentication, your browsers should be HTTP 1.1 compliant. If a non-HTTP 1.1–compliant browser requests a file from a server using Advanced Digest authentication, the server will request the client to provide Advanced Digest authentication credentials. The non-HTTP 1.1–compliant client rejects the request because Advanced Digest authentication is not supported by the client.

Advanced Digest authentication, like Digest authentication, will not work unless the following minimum requirements are met:

- All of the clients that need access to a resource that is secured with Advanced Digest authentication are using Internet Explorer 5 or Internet Explorer 6.
- The user and the server running IIS are members of the same domain, or the user is a member of a trusted domain within the same forest.
- Users have a valid Windows user account that is stored in Active Directory on the domain controller.
- The domain controller and the server running IIS are both running Windows Server 2003. If either your domain controller or the server running IIS is running Windows 2000 Server or earlier, IIS defaults to Digest authentication and does not notify you of this action.

| **Note** |
| --- |
| Only domain administrators can verify that the domain controller requirements are met. Check with your domain administrator if you are unsure about whether your domain controller meets these requirements. |

**Configuring Advanced Digest Authentication**

Enabling Advanced Digest authentication on the server running IIS requires the following three tasks:

- Enable Digest authentication for Windows domain servers. For more information about enabling Digest authentication for Windows domain servers, see Configuring Digest Authentication in IIS 6.0.
- Configure the realm name. For more information about configuring the realm name, see "Configuring the Realm Name" in Configuring Digest Authentication.
- Set the UseDigestSSP Metabase Property to **true**. The W3SVC service must be restarted before changes to this property take effect.

| **Important** |
| --- |
| If you follow the first two procedures but do not configure the **UseDigestSSP** metabase key, you will be using Digest authentication, not Advanced Digest authentication. |

Advanced Digest authentication uses the **UseDigestSSP** metabase property to switch between Digest and Advanced Digest security support provider interface (SSPI) code. If this property is set to **false**, Digest authentication is used. In all other cases (true, empty, or not set), IIS uses Advanced Digest authentication. If you configure **UseDigestSSP** at the W3SVC level of the metabase, all child keys inherit their configuration from that level.

| **Important** |
| --- |
| If you use Digest authentication and select worker process isolation mode, you must use the LocalSystem user account as the application pool identity. |

Additionally, IIS has two registry keys that specify timeout periods for Advanced Digest authentication: DigestPartialContextCacheTTL and DigestContextCacheTTL. These keys are not present by default, but you can add them to the

registry and configure them if the default values used by IIS are not sufficient.

The value of DigestPartialContextCacheTTL specifies the timeout period for partially-formed security contexts, during which a user is initially challenged for credentials to authenticate. IIS uses a default value of 30 seconds if the registry key is not present on the Web server.

| **Important** |
| --- |
| If you choose to configure this key, understand that keeping the value low will minimize the risk of Denial of Service (DoS) attacks. |

The value of DigestContextCacheTTL specifies the timeout period for fully-formed security contexts, during which a user remains authenticated after providing valid credentials on the initial challenge response. IIS uses a default value of 300 seconds if the registry key is not present on the Web server.

For more information about DigestPartialContextCacheTTL and DigestContextCacheTTL, see Global Registry Entries.

# Integrated Windows Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

With Integrated Windows authentication (formerly called NTLM, and also known as Windows NT Challenge/Response authentication), the user name and password (credentials) are hashed before being sent across the network. When you enable Integrated Windows authentication, the client browser proves its knowledge of the password through a cryptographic exchange with your Web server, involving hashing.

| **Important** |
| --- |
| Integrated Windows authentication is disabled by default if you install Windows Server 2003 Service Pack 1 (SP1) as part of a slipstream installation of a Windows Server 2003 operating system. If you install Windows Server 2003 SP1 by itself, as an upgrade to a Windows Server 2003 operating system, the setting for Integrated Windows authentication is unchanged from its Windows Server 2003 setting. Integrated Windows authentication is enabled by default for Windows Server 2003 operating systems. |

Integrated Windows authentication uses Kerberos v5 authentication and NTLM authentication. Kerberos is an industry-standard authentication protocol that is used to verify user or host identity. If Active Directory is installed on a domain controller running Windows 2000 Server or Windows Server 2003, and the client browser supports the Kerberos v5 authentication protocol, Kerberos v5 authentication is used; otherwise, NTLM authentication is used.

Integrated Windows authentication includes the Negotiate, Kerberos, and NTLM authentication methods. Negotiate, a wrapper for Kerberos v5 and NTLM, allows the client application to select the most appropriate security support provider for the situation. Kerberos v5 and NTLM have the following restrictions:

- NTLM can get past a firewall, but is generally stopped by proxies because NTLM is connection-based, and proxies do not necessarily keep connections established.
- Kerberos v5 requires that the client have a direct connection to Active Directory, which is generally not the case in Internet scenarios.
- Kerberos v5 requires that both the client and the server have a trusted connection to a Key Distribution Center (KDC) and be Active Directory–compatible.
- Kerberos v5 requires SPNs for multiple worker processes. If your Web site uses multiple worker processes, you can use Kerberos authentication, but you must manually register service names. For more information about Kerberos and service registration, see Kerberos and service registration later in this topic.

Integrated Windows authentication has the following limitations:

- Integrated Windows authentication is supported by only Internet Explorer 2 and later.
- Integrated Windows authentication does not work over HTTP proxy connections.

Therefore, Integrated Windows authentication is best suited for an intranet environment, where both user and Web server computers are in the same domain and where administrators can ensure that every user has Internet Explorer 2 or later.

**Client Authentication Process**

Unlike Basic authentication, Integrated Windows authentication does not initially prompt for a user name and password. The current Windows user information that is on the client is used for Integrated Windows authentication. If the authentication exchange initially fails to authorize the user, Internet Explorer prompts the user for a Windows account user name and password, which it processes using Integrated Windows authentication. Internet Explorer prompts the user for the correct user name and password up to three times. If, however, the user has logged on to the local computer as a domain user, then no authentication is required when the user accesses a network computer in that domain.

| **Note** |
| --- |

Microsoft Internet Explorer versions 4.0 and later can be configured to initially prompt for user name and password. For more information, see Internet Explorer Help.

**Kerberos and Service Registration**

Kerberos is the default method of network authentication for services in Windows Server 2003. A service can use Kerberos only if that service first registers a service principal name (SPN) in Active Directory. By default, Active Directory registers the NetBIOS, or computer, name of the server, but if the SPN is different, you must manually register the service. Before Kerberos can authenticate a service, the service must be registered on only one account object. If the logon account name of a service instance changes, the service must be reregistered under the new account. Therefore, only one application pool that has the service registered can authenticate with Kerberos.

IIS uses the NTAuthenticationProviders Metabase Property to pass the Negotiate header when Integrated Windows authentication is used to authenticate client requests. The Negotiate header permits clients to choose either Kerberos or NTLM. Negotiate uses Kerberos, unless Kerberos cannot be used by one of the systems involved in the authentication, or unless the calling application does not provide sufficient information to use Kerberos.

By default, Active Directory allows the Network Service and Local System accounts to use Kerberos. If **NtAuthenticationProviders** is set to **Negotiate,NTLM** for an application pool, Kerberos works correctly with Network Service as the worker process identity, because this is the default worker process identity for IIS 6.0 in worker process isolation mode. However, suppose you isolate applications on a Web site by configuring each application to run in a separate application pool under a separate process identity that is a domain account. If you attempt to use Integrated Windows authentication with the isolated applications, because you want to use Kerberos authentication, Kerberos fails.

Unless Kerberos authentication is configured correctly, it fails with a 401.1 error in the following situations:

- The Fully Qualified Domain Name (FQDN) is different from the NetBIOS name. For example, the IIS server might host a Web site called www.contoso.com on a server named web01.
- The process—for example, Dllhost.exe or W3wp.exe—that is performing the underlying authentication runs under an identity other than System, and no SPN is registered for that identity.
- Applications are hosted across multiple servers that use the same computer name. Only one user can be registered per computer name.
- If all servers in a Web farm use one computer name, but load balancing distributes requests to multiple servers, and no server in the Web farm has a unique SPN.

For Kerberos authentication to work correctly, you must configure isolation at the site level, not at the application pool level. All application pools in your Web site that use Kerberos must run with the same application pool process identity. You must then register that process identity as an SPN with Kerberos by using the Setspn.exe command-line tool. The Setspn.exe tool is included in Resource Kit Tools for Windows Server 2003 Deployment Kit companion CD, or on the Web at http://www.microsoft.com/reskit. You must be a domain administrator to set an SPN. For more information about registering a service principal name by using Setspn.exe, see Configuring Constrained Delegation for Kerberos.

For more information about this limitation and using Setspn.exe, see Security-Related Changes in Authentication.

# UNC Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The UNC authentication method, which is also known as UNC Passthrough authentication, determines the credentials to use for gaining access to a UNC share on a remote computer. UNC Passthrough authentication works only with authentication methods that can perform delegation, for example, Basic authentication and Kerberos authentication. If you use IIS Manager to create a Web site or virtual directory and you specify a UNC path for the content, IIS Manager prompts you to type a user name and a password for the network resource. The user name must be specified as *domain\username*. If you specify a user name and password, both the UNCUserName Metabase Propertyy and the UNCPassword Metabase Property are set.

In IIS 6.0, UNC authentication works by looking at the requesting user and the credentials that are stored in the **UNCUserName** and **UNCPassword** properties of the metabase to determine the credentials to pass through to the computer with the UNC share. If **UNCUserName**is specified (not empty) and **UNCPassword** is valid, the metabase user credentials are sent as the user identity for access to the remote share. If **UNCUserName** is specified (not empty) and **UNCPassword** is not valid, a "500 Internal Server Error: Invalid User Name or Password" message is sent to the client. If **UNCUserName** is empty, the credentials of the request user, which can be either an authenticated set of credentials for authenticated requests or IUSR_*computername* credentials for anonymous requests, are sent as the user identity for access to the remote share.

Table 5.6 shows the behavior resulting from the values of **UNCUserName** and **UNCPassword**.

**Table 5.6 Values and Behavior of UNC Passthrough Properties in the IIS Metabase**

| UNCUserName Value | UNCPassword Value | Resulting Behavior |
|---|---|---|
| Specified (not empty) | Valid value | Metabase user credentials are sent as the user identity for access to the remote share. |
| Specified (not empty) | Invalid value | 500 Internal Server Error: Invalid Username or Password message is sent to the client. |
| Empty | Not applicable | Credentials of the request user — either an authenticated set of credentials for authenticated requests or IUSR_*computername* credentials for anonymous requests — are sent as the user identity for access to the remote share. |

| **Note** |
|---|
| The **UNCAuthenticationPassthrough** metabase property is no longer used for UNC authentication. |

**Constrained Delegation**

*Delegation* is the act of allowing a service to impersonate a user account or computer account to access resources throughout the network. When a service is trusted for delegation, that service can impersonate a user to use other network services. *Constrained delegation* is a new option for Windows Server 2003 that allows you to specify the SPNs to which an account can delegate. A service can be trusted for delegation, but the domain administrator can limit that trust to a select group of explicitly specified services. By allowing delegation only to specific services, you can control the specific network resources the service or computer can use.

Constrained delegation is particularly useful in scenarios in which a site that requires authentication — a site that does not allow anonymous access — contains content that is housed on a remote UNC file server. With constrained delegation, you can enable Integrated Windows authentication, which can use NTLM authentication or send credentials across the network as a Kerberos token. For more information about Integrated Windows authentication, see Integrated Windows Authentication.

Figure 5.3 shows how constrained delegation works in a simple configuration.




**Deciding Whether to Use Constrained Delegation**

Constrained delegation is best implemented in a Windows Server 2003 domain. Although you can use constrained delegation in a Windows 2000 Server mixed domain, there are significant limitations. In a Windows 2000 Server mixed domain, constrained delegation works only for users that are authenticated using Basic or Kerberos authentication. In addition, you cannot restrict delegation to a specific set of services on a computer running Windows 2000 Server, which makes it difficult to implement this feature and reduce the attack surface of your server.

If you do not use constrained delegation but you enable Integrated Windows authentication, the token that the Web server obtains from the security infrastructure of Windows does not have sufficient permissions to access another computer, such as your file server. However, with constrained delegation and Integrated Windows authentication, the token received by the Web server from the security infrastructure of Windows is a Kerberos-based token with permission to access other computers, including the file server. Essentially, constrained delegation allows an NTLM-based token to be upgraded to a Kerberos-based token.

If you use Basic authentication, you probably do not need constrained delegation because the tokens cached by Basic authentication have sufficient permissions to access other computers in the domain, including the file server. However, in some situations you do not need to allow clients to authenticate with Integrated Windows authentication, but you might need constrained delegation. For example, you might have applications that need to access a file server using computer accounts, such as an enterprise intranet application that impersonates the client to gain access to information on various servers in the enterprise and then presents the consolidated data to the client over HTTP.

If you are running IIS 6.0 in IIS 5.0 isolation mode and you have an out-of-process application, this process normally runs as the local IWAM_*computername* account, which prevents constrained delegation from working correctly. One example of an application that might fail to perform as expected in this scenario is an ASP application that accesses a remote computer running Microsoft® SQL Server™ on which you intend for the authenticated users credentials to be used. To fix this, switch IIS to run in worker process isolation mode.

# .NET Passport Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Microsoft .NET Passport is a user-authentication service and a component of the Microsoft .NET Framework. The .NET Passport single sign-in service and express purchase service enable organizations to deliver a fast and convenient way for consumers to sign in and make transactions on a Web site. With the .NET Passport single sign-in service, you can map sign-in names to information in your databases, which enables you to offer .NET Passport members a personal Web experience through targeted content, advertisements, and promotions. The .NET Passport single sign-in service is similar to the forms-based authentication model that is commonly used on the Web today. The .NET Passport network extends this model in the following ways:

- The sign-in, sign-out, and registration pages are centrally hosted rather than being specific to each individual site.
- .NET Passport can be extensively co-branded to match the look and feel of your site. Co-branding material is served directly from your site and is included in the centrally hosted pages when those pages are displayed on a client browser.
- Centrally hosted pages that require additional security to exchange credentials or other information are always served using SSL.
- All .NET Passport sign-in and core profile cookies are strongly encrypted. Each participating Web site receives a unique encryption key that ensures information privacy.
- The central .NET Passport servers return encrypted sign-in and profile information to your site, which can then be used to write local cookies and thus avoid redirects back to the central .NET Passport servers on subsequent page views.
- Members do not need to retype their sign-in name and password when moving from site to site. .NET Passport–enabled sites issue a set of encrypted cookies in the .NET Passport central servers' domain to facilitate silent and seamless sign-in across sites. However, sites might still choose to always force members to be redirected to .NET Passport sign-in and to authenticate upon first viewing of their site.
- Participating Web sites never receive a member's password. The authentication cookie is actually a pair of encrypted time stamps asserting when the member signed in. When members choose to sign out by clicking the .NET Passport sign-out link, they are redirected to a central page that enables deletion of all .NET Passport cookies from all of the sites that the member visited during the session.
- There is no real-time, server-to-server communication between participating Web sites and the central .NET Passport servers. All information exchange occurs through the client's browser using HTTP redirects, encrypted information on the query string, and cookies. Server-to-server communication occurs only when a centrally hosted XML configuration file is periodically downloaded and cached locally by .NET Passport's server-side object (supplied in the .NET Passport SDK); this XML file contains current URLs for all the .NET Passport servers and the current profile schema.

Microsoft .NET Passport uses standard Web technologies, such as SSL, HTTP redirects, cookies, Microsoft® JScript® development software, and strong symmetric key encryption to deliver the single sign-in service. In addition, .NET Passport is compatible with Internet Explorer 4 and later, Netscape Navigator version 4.0 and later, and some versions of UNIX.

.NET Passport–enabled sites rely on the .NET Passport central server to authenticate users. However, the .NET Passport central server does not authorize or deny the access of a user to individual .NET Passport–enabled sites. It is the responsibility of the Web site to control user access rights.

After .NET Passport authentication is verified, a Windows Server 2003 Passport user can be mapped to a user of Active Directory through the users Windows Server 2003 Passport identification, if such a mapping exists. A token is created by the Local Security Authority (LSA) for the user, and is set by IIS for the HTTP request. Application developers and Web site administrators can use this security model for authorization based on users of Active Directory. You can also delegate these credentials through the constrained delegation feature supported by Windows Server 2003. For more information about constrained delegation, see Constrained Delegation for UNC File Content. When using Microsoft .NET Passport authentication, you can store account information and user credentials in Active Directory. If account information is not mapped in Active

Directory, IIS automatically creates an Anonymous authentication session with anonymous user account credentials. In most cases, anonymous user account credentials equal IUSR_*ComputerName*.

The only difference between .NET Passport authentication IUSR credentials and standard Anonymous authentication IUSR credentials is that .NET Passport authentication creates an IIS authentication token for the anonymous account that contains a security identifier, known as the .NET Passport User ID (PUID). Active Directory returns a user authentication token if the account information is mapped in Active Directory, and the corresponding user is logged on to the local computer.

Users can register at a .NET Passport–enabled site, and their user information is then stored in encrypted profiles on .NET Passport servers. When.NET Passport users register at a participating site, their personal information can be shared with the site. This sharing can make the registration process faster. When the .NET Passport users sign in to that site again, their .NET Passport profiles can enable access to personalized accounts or services on that site.

For more information about .NET Passport authorization, see the Microsoft .NET Passport homepage.

**Conforming to Passport Privacy Guidelines**

When you implement .NET Passport for Web sites running on Windows Server 2003, you must ensure that you conform to the Passport privacy guidelines. Use one of the following configurations for your Web site:

- Configure IIS 6.0 to use .NET Passport authentication and Anonymous authentication. Your Web site will exhibit typical Passport behavior, but this configuration requires that you store all Passport-enabled Web pages on the Web site itself.

- Create a virtual directory for an IIS 6.0 Web site, and configure it for Anonymous authentication. Store a default Passport-enabled Web page in this virtual directory. This page should display the Passport sign-in logo, which results from the call to Manager.LogoTag(), and should also contain a return URL value that points to a second virtual directory on the same Web site.

  Configure the second virtual directory to use .NET Passport authentication only. You can set the Passport Manager Administration Utility **ForceSignin** and **TimeWindow** values as you choose. Store the rest of your Passport-enabled Web pages in this second virtual directory.

# FTP Site Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Based on your security requirements, you can select one of two IIS authentication methods to validate users requesting access to your FTP pages. Available authentication settings must be set at the site level for FTP sites.

| **Note** |
| --- |
| Digest and Integrated Windows authentication cannot be used with FTP sites. |

- Anonymous FTP authentication. If you select Anonymous FTP authentication for a resource, all requests for that resource are accepted without prompting the user for a user name or password. This is possible because IIS automatically creates a Windows user account called IUSR_*ComputerName*, where *ComputerName* is the name of the server on which IIS is running. This is very similar to Web-based Anonymous authentication. If Anonymous FTP authentication is enabled, IIS always tries to use that method first, even if you enable Basic FTP authentication.

- Basic FTP authentication. To establish an FTP connection with your Web server by using the Basic FTP authentication method, users must log on with a user name and password corresponding to a valid Windows user account. If the FTP server cannot verify a user's identity, the server returns an error message. Basic FTP authentication provides only low security because the user transmits the user name and password across the network in an unencrypted form.

Table 5.6 compares Anonymous and Basic FTP authentication.

**Table 5.6 Comparison of Anonymous and Basic FTP Authentication Methods**

| Method | Security Level | How Passwords Are Sent | Crosses Proxy Servers and Firewalls | Client Requirements |
| --- | --- | --- | --- | --- |
| Anonymous FTP authentication | None | N/A | Yes | Any FTP client |
| Basic FTP authentication | Low | Plaintext | Yes | Any FTP client |

# Access Control in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can control which users and computers can access your Web server and its resources. There are two main types of access control: NTFS permissions, which are controlled by the operating system, and Web site permissions, which are controlled by IIS. You can also use IP address restrictions to limit access by specific computers. To help secure your Web server, you should use a defense-in-depth strategy that combines suitable authentication methods with tools such as firewalls and the appropriate NTFS permissions and Web site permissions. IIS built-in accounts and features, such as URL authorization, help you control access to resources on multiple levels, from entire Web and FTP sites to individual files or specific URLs.

- Securing your files with NTFS. When you use the NTFS file system, you can limit access to your Web server's files and directories. You can configure the file and directory permissions that set the access level assigned to a particular user account or user group. For example, you can configure your Web server to enable a specific user to view and execute a file, while excluding all other users from accessing that file.

- Securing your site with Web site permissions. Web permissions, unlike NTFS permissions, apply to all users accessing your Web sites. NTFS permissions apply only to a specific user or group of users with a valid Windows account. NTFS controls access to physical directories on your server, whereas Web permissions control access to virtual directories on your Web site.

- Securing your site with IP address restrictions. With IP address restrictions, you can configure your Web server to assign or deny specific computers, groups of computers, or domains access to Web sites, directories, or files. For example, if your intranet server is connected to the Internet, you can prevent Internet users from accessing your Web server by assigning access only to members of your intranet, and explicitly denying access to outside users.

How Access Control Works

Access control involves both user rights and permissions. *User rights* are rights that allow a user to perform specific actions on a computer or network. *Permissions* are rules that are associated with an object, such as a file or folder, to regulate which accounts can gain access to that object. After you configure your Windows file system and Web server security features, when a user attempts to access your Web server, the server carries out several access-control processes to identify the user and determine the allowed level of access.

Figure 5.4 shows how the access-control process works.




The following steps outline the access-control process:

1. The client requests a resource on the server.

2. The IP address of the client is checked against any IP address restrictions in IIS. If the IP address is denied access, the request fails, and a 403 Access Forbidden message is returned to the user.

3. The server, if configured to require it, requests authentication information from the client. The browser either prompts the user for a user name and password or offers this information automatically.

4. IIS checks whether the user has a valid Windows user account. If the user does not, the request fails, and a 401 Access is denied message is returned to the user.

5. IIS checks whether the user has Web permissions for the requested resource. If the user does not, then the request fails, and a 403 Access Forbidden message is returned to the user.

6. Any security modules, such as ASP.NET impersonation, are added.

7. IIS checks the NTFS permissions on static files, ASP, and Common CGI files for the resource. If the user does not have NTFS permissions for the resource, then the request fails, and a 401 Access is denied message is returned to the user.

8. If the user has NTFS permissions, the request is fulfilled.

# NTFS Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

With NTFS access permissions as the foundation of your Web servers security, you can define the level of file and directory access that is granted to Windows users and groups. For example, if a business decides to publish its catalog on your Web server, you need to create a Windows user account for that business and then assign permissions for the specific Web site, directory, or file. The permissions should permit only the server administrator and the owner of the business to update the contents of the Web site. Public users should be allowed to view the Web site, but not to alter its contents.

To control access to directories and files in this manner, you must be using drives formatted in NTFS, not FAT32. Using FAT32 is strongly discouraged because it allows users access to every file on your hard drive. NTFS is a more powerful and restrictive file system than FAT and FAT32. With NTFS, you can limit access to your Web server's files and directories. You can also configure the access level granted to a particular user or group for the files and directories on your server.

Table 5.7 compares the security options of NTFS and FAT file systems.

**Table 5.7 Comparison of NTFS and FAT File Systems**

| NTFS | FAT |
|---|---|
| Allows administrators to control access to directories and files using NTFS permissions. Permissions can be set on the file level, as well as the directory level. | Does not allow administrators to control access to directories and files. |
| Supports file encryption, which greatly enhances file security. | Does not support file encryption. |
| Allows administrators to enable Web Distributed Authoring and Versioning (WebDAV) properties. | Does not allow administrators to enable WebDAV properties. |
| Supports Active Directory and domain-based security. | Does not support Active Directory and domain-based security. |

The NTFS permission levels are as follows:

- Full Control. Users can do anything to the file, including taking ownership of it. It is recommended that you grant this level of access only to administrators.
- Modify. Users can view and modify files and file properties, including deleting and adding files to a directory or file properties to a file. Users cannot take ownership or change permissions on the file.
- Read & Execute. Users can run executable files, including scripts.
- List Folder Contents. Users can view a list of a folder's contents.
- Read. Users can view files and file properties.
- Write. Users can write to a file.

**Important**

If you do not assign Read & Execute permissions to the IUSR_*ComputerName* account for a resource, anonymous users will be denied access to that resource. The default permissions assigned when you add a user or group to the Access Control List (ACL) are Read & Execute, List Folder Contents, and Read.

For more details about setting NTFS permissions, see Setting NTFS Permissions for Directories or Files.

# Access Control Lists

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

NTFS permissions are set and managed by using access control lists (ACLs). An *ACL* is a list that indicates which users or groups are able to access or modify a particular file. An ACL is made up of access control entries (ACEs).

**ACL Tools**
Many tools, such as the **cacls** command and the Xcacls.exe command-line tool, are available to help you set and manage ACLs.

**Cacls**
**Cacls** is a Windows Server 2003 command that allows you to display or modify discretionary access control lists (DACLs). For more information about the **cacls** command, see Cacls in Help and Support Center for Windows Server 2003.

**Xcacls.exe**
Xcacls.exe is a Windows Server 2003 command-line tool that you can use to set all file-system security options that are accessible in Windows Explorer from the command line. Xcacls.exe does this by displaying and modifying the ACLs of files. For more information, see Xcacls.exe.

**Default ACLs and Privileges for the IIS_WPG Group**
When you install IIS 6.0, some ACLs and privileges are set automatically for the IIS_WPG group. Table 5.8, Table 5.9, and Table 5.10 list the default ACLs that are set during the installation process.

**Table 5.8 Default File System ACLs for the IIS_WPG Group**

| Location | Setting | Description |
|---|---|---|
| %windir%\Help\iishelp\common | Read, Execute | Access to default error page |
| %windir%\IIS Temporary Compressed Files | Full | Access to the compression directory |
| %windir%\system32\MicrosoftPassport | Read, Execute, Write, Delete | Passport configuration |
| %windir%\inetsrv\ASP Compiled Templates | Full | Access to ASP compiled templates that are saved to disk |
| %windir%\inetpub\wwwroot | Read, Execute | Access to the content root directory |

**Table 5.9 Default Metabase ACLs for the IIS_WPG Group**

| Location | Setting | Description |
|---|---|---|
| / (IIS_ROOT key) | Read, Unsecure Read, Enumerate Keys | Give worker process access to read metabase data |
| /LM/W3SVC/AppPools | Unsecure Read | Give worker process access to read unsecured metabase properties |
| /LM/W3SVC/Filters | Read, Unsecure Read, Enumerate Keys, Write | Allow IIS to update filter state |
| /LM/W3SVC/*X*/Filters | Read, Unsecure Read, Enumerate Keys, Write | Allow IIS to update filter state (for new sites, where *X* is the Web site number) |

**Table 5.10 Default Registry Key ACLs for the IIS_WPG Group**

| Location | Setting | Description |
|---|---|---|
| HKLM\System\CurrentControlSet\Services\W3SVC | Read | Read W3SVC configured registry keys |
| HKLM\System\CurrentControlSet\Services\InetInfo | Read | Read InetInfo configured registry keys |
| HKLM\System\CurrentControlSet\Services\HTTP | Read | Read HTTP configured registry keys |
| HKLM\System\CurrentControlSet\Services\ASP | Read | Read ASP configured registry keys |

All worker process identities must be members of the IIS_WPG group. If the worker process identity account is not in the IIS_WPG group and does not have the appropriate privileges and permissions, the worker process will not start. In addition, if the worker process identity accounts are not members of the IIS_WPG group, they might not be able to launch a worker process after a modification to the operating system, such as upgrading to the next version or installing a service pack.

In addition, IIS sets restrictive ACLs on log files. For more information about log file permissions, see Analyzing Log Files.

**Metabase ACLs**

IIS installs the metabase files with strict ACLs set to prevent anyone but administrators from viewing your configuration data. You can use the MetaACL.vbs command-line tool to change the ACLs and grant granular permissions on a site-by-site or application-by-application basis. For more information about MetaACL.vbs, see Knowledge Base article 267904, Metaacl.exe modifying metabase permissions for IIS Admin Objects. For more information about metabase security, see Working with the Metabase.

| **Important** |
| --- |
| Before you edit the metabase, verify that you have a backup copy that you can restore if a problem occurs. For information about how to back up the metabase, see Working with the Metabase. |

# Web Site Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When configuring access control for your servers, it is important to understand the distinction between Web site permissions and NTFS permissions. Unlike NTFS permissions, Web site permissions apply to all of the users that access your Web sites. NTFS permissions apply only to a specific user or group of users with a valid Windows account. NTFS permissions control access to physical directories on your server, whereas Web site permissions control access to virtual directories on your Web site. For example, you can use Web site permissions to control whether users visiting your Web site are allowed to view a particular page, upload information, or run scripts on the site.

| **Note** |
| --- |
| If Web site permissions conflict with NTFS permissions for a directory or file, the more restrictive settings are applied. |

IIS 6.0 supports the following Web site permissions:

- Read. Users can view file content and properties. This permission is set by default.
- Write. Users can change content and properties of directories or files

  | **Important** |
  | --- |
  | Allowing users to have Read and Write access to source code can compromise the security of your server. |

- Script source access. Users can access the source code for files, such as the scripts in an ASP application. This option is available only if Read or Write permissions are assigned. If the Read permission is assigned, users can view the source code. If the Write permission is assigned, users can view and modify the source code.

  | **Note** |
  | --- |
  | When you use this option, users might be able to view sensitive information, such as a user name and a password, from scripts in an ASP program. They might also be able to change source code that runs on your server. This can seriously affect the security and the performance of your server. You might want to control access to this type of information and to these functions by using individual Windows accounts and a higher-level authentication method, such as Integrated Windows authentication. |

- Directory browsing. Users can view file lists and collections.
- Log visits. A log entry is created for each visit to the Web site.
- Index this resource. Allows Indexing Service to index this resource.

Whereas Web site permissions control whether users can access sites, code, or pages, another set of permissions, which are known as *Execute permissions*, control whether scripts and other executables can be run.

The following Execute permissions are supported by IIS 6.0:

- None. Prevents all scripts and executables from running.
- Scripts only. Enables applications that are mapped to a script engine to run in this directory, even if the applications are not assigned permissions for executables.
- Scripts and Executables. Enables any application to run in this directory, including applications that are mapped to script engines and Windows binaries (.dll and .exe files).

The Scripts only permission is more secure than the Scripts and Executables permission because the Scripts only permission allows you to limit the applications that can be run in the directory. For information about how to configure Web site permissions, see Securing Sites with Web Site Permissions.

# IIS and Built-in Accounts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS uses a number of built-in Windows accounts, as well as accounts that are specific to IIS. For security reasons, you should be aware of the different accounts and their default user privileges. It can be a security risk to change the identity of a worker process so that it runs as an account with a high level of access, such as the LocalSystem user account.

**LocalSystem**
The built-in LocalSystem user account has a high level of access privileges; it is part of the Administrators group. If a worker process identity runs as the LocalSystem user account, that worker process has full access to the entire system. When IIS 6.0 is running in IIS 5.0 isolation mode, this is the default user account for worker process identities. LocalSystem has one default user right, Full access.

**Network Service**
The built-in Network Service user account has fewer access privileges on the system than the LocalSystem user account, but the Network Service user account is still able to interact throughout the network with the credentials of the computer account. For IIS 6.0, it is recommended that the worker process identity that is defined for application pools run as the Network Service user account, which is the default setting. The following table shows the default user privileges for the Network Service account, along with how each privilege is derived.

| Privilege | Source |
|---|---|
| Replace a process-level token (SeAssignPrimaryTokenPrivilege) | Explicit assignment |
| Adjust memory quotas for a process (SeIncreaseQuotaPrivilege) | Explicit assignment |
| Generate security audits (SeAuditPrivilege) | Explicit assignment |
| Bypass traverse checking (SeChangeNotifyPrivilege) | Through membership in the Everyone group |
| • Access this computer from the network (SeNetworkLogonRight) | Through membership in the Everyone group |
| • Log on as a batch job (SeBatchLogonRight) | Through membership in the IIS_WPG group |
| • Log on as a service (SeInteractiveLogonRight) | Explicit assignment |
| • Impersonate a client after authentication | Through membership in the IIS_WPG group |

**Local Service**
The built-in Local Service user account has fewer access privileges on the computer than the Network Service user account, and those user privileges are limited to the local computer. Use the Local Service user account if the worker process does not require access outside the server on which it is running. The following table shows the default user privileges for the Local Service account, along with how each privilege is derived.

| Privilege | Source |
|---|---|
| Replace a process-level token (SeAssignPrimaryTokenPrivilege) | Explicit assignment |
| Adjust memory quotas for a process (SeIncreaseQuotaPrivilege) | Explicit assignment |
| Generate security audits (SeAuditPrivilege) | Explicit assignment |
| Bypass traverse checking (SeChangeNotifyPrivilege) | Through membership in the Everyone group |

| | |
|---|---|
| • Access this computer from the network (SeNetworkLogonRight) | Through membership in the Everyone group |
| • Log on as a batch job (SeBatchLogonRight) | Explicit assignment |

**IIS_WPG**

The IIS IIS_WPG group account has the minimum permissions and user privileges that are necessary to start and run a worker process on a Web server. Application pool identities must be members of this group so the application pool can register with Http.sys. The following table shows the default user privileges for the IIS_WPG account, along with how each privilege is derived.

| Privilege | Source |
|---|---|
| Access this computer from the network (SeNetworkLogonRight) | Through membership in the Everyone group |
| Bypass traverse checking (SeChangeNotifyPrivilege) | Through membership in the Everyone group |
| Impersonate a client after authentication (SeImpersonatePrivilege) | Explicit assignment |
| • Log on as a batch job (SeBatchLogonRight) | Explicit assignment |

**IUSR_ComputerName**

The IIS IUSR_*ComputerName* user account is for anonymous access to IIS. By default, when a user accesses a Web site that uses Anonymous authentication, that user is mapped to the IUSR_*ComputerName* account. The following table shows the default user privileges for the IUSR_*ComputerName* account, along with how each privilege is derived.

| Privilege | Source |
|---|---|
| Access this computer from the network (SeNetworkLogonRight) | Explicit assignment |
| • Allow log on locally (SeInteractiveLogonRight) | Explicit assignment |
| Bypass traverse checking (SeChangeNotifyPrivilege) | Through membership in the Everyone group |
| • Log on as a batch job (SeBatchLogonRight) | Explicit assignment |

**IWAM_ComputerName**

The IIS IWAM_*ComputerName* user account is for starting out-of-process applications in IIS 5.0 isolation mode. The following table shows the default user privileges for the IWAM_*ComputerName* account, along with how each privilege is derived.

| Privilege | Source |
|---|---|
| Access this computer from the network (SeNetworkLogonRight) | Explicit assignment |
| • Adjust memory quotas for a process (SeIncreaseQuotaPrivilege) | Explicit assignment |
| Bypass traverse checking (SeChangeNotifyPrivilege) | Through membership in the Everyone group |
| • Log on as a batch job (SeBatchLogonRight) | Explicit assignment |
| • Replace a process-level token (SeAssignPrimaryTokenPrivilege) | Explicit assignment |

**ASPNET**

The built-in ASPNET user account is for running the ASP.NET worker process in IIS 5.0 isolation mode. The following table

shows the default user privileges for the ASPNET account, along with how each privilege is derived.

| **Privilege** | **Source** |
|---|---|
| Access this computer from the network (SeNetworkLogonRight) | Explicit assignment |
| • Allow logon locally (SeInteractiveLogonRight) | Through membership in the Users group |
| Bypass traverse checking (SeChangeNotifyPrivilege) | Through membership in the Users group |
| • Deny logon locally (SeDenyInteractiveLogonRight) | Explicit assignment |
| • Log on as a batch job (SeBatchLogonRight) | Explicit assignment |
| • Log on as a service (SeInteractiveLogonRight) | Explicit assignment |

# Anonymous Access

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Anonymous access, the most common Web site access control method, allows anyone to visit the public areas of your Web sites. In IIS 6.0, anonymous users are assigned to the IUSR_*ComputerName* account, which is created when IIS is installed. This account is a valid Windows account and is a member of the **Guests** group. For example, if your computer name is SalesDept1, the anonymous account name is IUSR_SalesDept1. Each Web site on your server can use either the same or different anonymous user logon accounts. The IUSR_*ComputerName* account can be defined on a computer or on a domain.

In IIS 6.0, the IUSR_*ComputerName* account is denied Write access to Web content by default. In addition, this account is explicitly denied access to all executables (such as cmd.exe) in the %windir%\System32 folder. For security reasons, to run most executables in this folder, you must be a member of the Administrators group or of the LocalSystem, Interactive, or Service account. Therefore, an anonymous user cannot run these executables (for example, Common Gateway Interface (CGI) code that calls external programs such as PHP).

When the Web sites and applications running on your Web server require anonymous access, you must configure a user account specifically for anonymous access. This user account can be stored in the local account database on the Web server or in a domain.

| **Important** |
| --- |
| If you are upgrading to IIS6.0, be aware that the upgrade process automatically configures IIS to use the default anonymous user account IUSR_*ComputerName*, where *ComputerName* is the local computer. If you have a domain-based anonymous user account, you must configure the anonymous user identity to use this same domain-based account after the upgrade. |

# Configuring Application Pool Identity in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you run IIS in worker process isolation mode, applications can be configured to run in separate application pools. An application pool is a configuration that links one or more applications to a set of one or more worker processes. The *application pool identity* is the name of the account under which the worker process of the application pool runs.

| **Important** |
| --- |
| This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. |

In earlier versions of IIS, worker processes ran as the LocalSystem account. Because the LocalSystem account has access to almost all resources of the operating system, this had serious security implications. IIS 6.0 strengthens security because the worker process runs under the new built-in Network Service account by default. IIS 6.0 also allows you to configure the account under which worker processes will run. You have the option of using one of the following three predefined accounts, or creating your own account:

- Network Service. By default, application pools run under the Network Service user account, which has low-level user access rights, so it provides better security against attackers or malicious users who might attempt to take over the computer on which the WWW service is running.

- Local Service. The Local Service user account also has low-level access rights, and is useful when you want to configure application pools to run as the Local Service user account in situations that do not require access to resources on remote computers.

- LocalSystem. You can configure application pools to run as the LocalSystem user account, which has more user rights than the Network Service or Local Service user account. However, running an application pool under an account with increased user rights presents a high security risk.

For example, suppose an ISP wants to allow customers to upload CGI applications and then add them to an application pool. The ISP can run CGI-enabled applications in a separate application pool under the Network Service user account, which has lower user rights, to reduce the risk that these applications will be used to attack the server.

IIS_WPG Group

The IIS_WPG user group is new in IIS 6.0. It provides the minimum set of privileges and permissions required to start and run a worker process on a Web server. On a clean installation of IIS 6.0, the IIS_WPG group contains the Network Service, Local Service, LocalSystem, and IWAM_*ComputerName* accounts.

To designate a specific user account as a worker process identity for a Web site, make the user account a member of the IIS_WPG group. This is a convenient way to configure a custom identity for a worker process without manually assigning privileges and permissions.

All worker process identities must be members of the IIS_WPG group. If the worker process identity is not in the IIS_WPG group and does not have the appropriate privileges and permissions, the associated worker process will not start. In addition, the worker process identity might not be able to launch a worker process after a modification to the operating system, such as an upgrade or service pack.

| **Note** |
| --- |
| By default, the IIS_WPG group does not have the right to start CGI processes. If you create a configurable account and add it to the IIS_WPG group as a worker process identity, you must also assign this new account the following two user rights to start CGI processes: Adjust memory quotas for a process and Replace a process level token. |

When you install IIS 6.0, the setup process adds an ACL to the wwwroot directory of the default Web site. This ACL gives the IIS_WPG group permission to access that directory, which means that all worker process identities in the IIS_WPG group can access the contents of this directory on the default Web site. However, you should not give this group access by ACLs to content for a specific Web site. Instead, to isolate applications, do one of the following:

- Remove the IIS_WPG group from ACLs on site content.
- Allow access to a Web site's content directories and files by providing ACLs for the worker process identity associated with the Web site, rather than for the IIS_WPG group.
- If you create a configurable account to use as the worker process identity, add that account to any ACLs on the site's content directories and files.

For information about changing the account under which an application pool runs by using IIS Manager, see Configuring Application Pool Identity with IIS 6.0.

You can change the worker process identity programmatically by using the Microsoft® Visual Basic® Scripting Edition (VBScript) development system sample script provided in Listing 5.2. Listing 5.2 provides a customizable VBScript-based file that creates an application that runs in an application pool under the supplied identity. Then the script starts Task Manager to show the worker process, and it attempts to access the newly created application using Internet Explorer. When the script is complete, you should see, in Task Manager, a new worker process that serves the request and that runs under the supplied identity.

**Listing 5.2 Sample Script for Configuring Worker Process Identity**

```
' VBScript source code
option explicit

if WScript.Arguments.Count < 4 OR WScript.Arguments.Count > 5 then
   WScript.Echo "Usage:"
   WScript.Echo WScript.ScriptFullName & " username password AppPoolName vdirname
<physical path for vdir>"
   WScript.Quit
end if
dim basepath         :   basepath = "IIS://localhost/w3svc/1/root"
dim username         :   username = WScript.Arguments(0)
dim password         :   password = WScript.Arguments(1)
dim AppPoolName      :   AppPoolName = WScript.Arguments(2)
dim vdirname         :   vdirname = WScript.Arguments(3)
dim VDirPhysicalPath

if WScript.Arguments.Count < 5 then
   VDirPhysicalPath = "c:\inetpub\wwwroot"
else
   VDirPhysicalPath = WScript.Arguments(4)
end if

if CreateNewUser(username, password) = false then
   WScript.Echo "Error creating user " & username & ". "
   WScript.Quit
end if

if AddUserToIIS_WPG(username) = false then
   WScript.Echo "Error adding user " & username & " IIS_WPG. "
   WScript.Quit
end if

if CreateNewAppPool(username, password, AppPoolName) = false then
   WScript.Echo "Error creating AppPool " & AppPoolName & " running as user " &
username & ". "
   WScript.Quit
end if

if CreateNewVDir(vdirname, vdirphysicalpath) = false then
   WScript.Echo "Error creating new virtual directory " & vdirname & ". "
   WScript.Quit
end if

if AssignAppPool(vdirname, AppPoolName) = false then
   WScript.Echo "Error assigning Application Pool " & AppPoolName & " to vdir " &
vdirname & ". "
   WScript.Quit
end if
```

```
dim wsh : set wsh = CreateObject("WScript.Shell")
WScript.Echo "Waiting for iisreset to return"
wsh.Run "iisreset /restart",,true
wsh.Run "http://localhost/" & vdirname
wsh.Run "taskmgr"

function AssignAppPool(vdirname, AppPoolName)
   dim vdirobj

   set vdirobj = GetObject(basepath & "/" & vdirname)
   vdirobj.AppPoolId = AppPoolName
   vdirobj.SetInfo
   if err = 0 then
      AssignAppPool = true
   else
      AssignAppPool = false
   end if
end function

function CreateNewVDir (vdirname, vdirphysicalpath)
   dim VdirObj, VRoot
   set VRoot = GetObject(basepath)
   set VdirObj = VRoot.Create("IIsWebVirtualDir", vdirname)
   VdirObj.AccessRead = True
   VdirObj.AccessScript = True
   VdirObj.Put "Path", vdirphysicalpath
   VdirObj.SetInfo

   vdirobj.AppCreate true
   vdirobj.SetInfo
   if err = 0 then
      CreateNewVDir = True
   else
      CreateNewVdir = false
   end if
end function

function CreateNewAppPool(username, password, AppPoolName)
   dim AppPools, newAppPool
   set AppPools = GetObject("IIS://localhost/w3svc/AppPools")
   set newAppPool = AppPools.Create("IIsApplicationPool", AppPoolName)
   newAppPool.WamUserName = username
   newAppPool.WamUserPass = password
      newAppPool.LogonMethod =1
   newAppPool.AppPoolIdentityType=3
   newAppPool.SetInfo
   if err = 0 then
      CreateNewAppPool = True
   else
      CreateNewAppPool = false
   end if
end function

function AddUserToIIS_WPG (username)
   dim nw : set nw = CreateObject("WScript.Network")
   dim computername : computername = nw.Computername
   dim MyUser
   dim IIS_WPG

   set IIS_WPG = GetObject("WinNT://" & computername & "/IIS_WPG,group")
   set MyUser = GetObject("WinNT://" & computername & "/" & username & ",user")
   if not IIS_WPG.IsMember(MyUser.AdsPath) then
      IIS_WPG.Add(MyUser.AdsPath)
   end if
   if err = 0 then
      AddUserToIIS_WPG = True
   else
      AddUserToIIS_WPG = false
   end if

end function
```

```
function CreateNewUser(username, password)
   dim nw             :   set nw = CreateObject("WScript.Network")
   dim computername   :   computername = nw.Computername
   dim container      :   set container = GetObject("WinNT://" & computername)
   dim newuser        :   set newuser = container.Create( "user", username)
   newuser.SetPassword(password)
   newuser.SetInfo
   if err = 0 then
      CreateNewUser = True
   else
      WScript.Echo "Error: " & err.Description & " - " & err.number
      CreateNewUser = false
   end if

end function
```

# URL Authorization

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A new authorization framework is available with Windows Server 2003. Implemented through the Windows Server 2003 Authorization Manager, role-based access control is designed to make it easier to manage access control for Web-based or line-of-business applications in enterprise environments.

For example, for an expense-reporting or Web-based shopping application, the authorization decisions often do not involve determining access to persistent objects. Instead, the applications might need to verify a workflow, or they might involve multiple distinct operations, such as querying a database and sending mail. Access decisions often are made based not only on token group membership but also on business logic, such as the amount that is submitted in an expense application or verification of workflow completion. For applications such as these, which do not have well-defined persistent objects, object-oriented ACLs are not effective.

IIS extends the use of role-based authentication by providing gatekeeper authorization to specific URLs. IIS URL authorization allows you to authorize user access to the URLs that make up a Web application. When a user requests access to a URL, URL authorization validates the users access based on the roles of that user, which can be defined in Lightweight Directory Access Protocol (LDAP) queries, custom user roles, and the Authorization Manager rules (scripts that are written in VBScript or JScript, which you can include in role definitions and task definitions). This allows you to control all user access to URLs instead of controlling access per ACL, on each resource.

Additionally, Web applications can use URL authorization together with Authorization Manager to control — from within the same policy store — access to the URLs that make up a Web application and the application-specific tasks and operations. Maintaining the policy in the same policy store allows you to manage access to the URLs and application features from a single point of administration while leveraging the store-level application groups and user-programmable business rules.

URL authorization is task-based, allowing you to create flexible authorization schemes that selectively allow or deny access to parts of the URL namespace for certain sets, users, or roles. The **URLAuthorizationModule** class maps users and roles to pieces of the URL namespace. This module implements both positive and negative authorization assertions and is available for use anytime. You need only to list users or roles in the <allow> or <deny> elements of the <authorization> section of a configuration file.

URL authorization is implemented as an ISAPI interceptor. When an application, virtual directory, or URL is configured to use URL authorization, each request to a URL is routed to the URL authorization ISAPI interceptor. The URL authorization wildcard script map uses Authorization Manager to authorize access to the requested URL. The URL must be associated with an Authorization Manager policy store that contains the authorization policy for the URL. After the client has been authorized to access the URL, the Execute URL feature of the URL authorization ISAPI passes the request to the appropriate handler for the URL, such as ASP.dll, another ISAPI, or the static file handler. Figure 5.5 shows this process. In this figure, the wildcard script map is named URL Authz ISAPI, Authorization Manager is named .NET Authz Framework, and the Execute URL feature is named ExecURL.



By using URL authorization, you can control access based on information that is available only at run time. For example, if you have a Web page that should be available only to employees in a given cost center or to employees of a certain age, you can assign roles to the correct users based on LDAP queries that check the cost center or age attributes on a user object. If employees can access only certain pages on certain days of the week or during a certain time of day, you can create a BizRule that grants access to the URL based on these values or any value that can be asserted at run time, including IIS server variables.

For more information about Authorization Manager, see "Authorization Manager" in Help and Support Center for Windows Server 2003.

# Using URL Authorization

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 does not provide graphical user interface (GUI) support for URL authorization. To use URL authorization in IIS 6.0, you must enable the wildcard script map, Urlauth.dll. In addition, you must set the following metabase properties on the application, virtual directory, or URL (Web site):

- AzEnable. This property enables URL authorization for the virtual directory, application, or URL that corresponds to the entry in the metabase.
- AzStoreName. This property associates an Authorization Manager store with the virtual directory, application, or URL.
- AzScopeName. This property associates the virtual directory, application, or URL with a scope. This scope will be the name of a scope in the IIS 6.0 URL authorization application in the Authorization Manager policy store referred to in the AzStoreName property. If no scope or an empty string is specified, the default scope of the URL authorization is used.
- AzImpersonationLevel. This property determines the impersonation behavior for the application, allowing you to configure the Web application to impersonate the client user, the IIS worker process, or the IUSR_*ComputerName* account for the worker process. Each setting significantly changes the environment and implied design of the Web application.

For more information about these metabase properties, see URL Authorization in IIS 6.0.

Although URL authorization controls access to other forms of authorization, such as settings for IIS directory security or ACL permissions, the application context still requires the correct IIS directory security and ACL permissions. URL authorization allows the IIS directory security and ACL permissions to be more easily maintained.

When URL authorization is configured, the AzStoreName property in the IIS metabase entry for the application, virtual directory, or URL identifies an Authorization Manager policy store. To manage the authorization policy, you must run Authorization Manager in Developer mode and use the **Open Policy Store** command from the **Action** menu. URL authorization is an application in this store. The AzScopeName property in the metabase entry is an Authorization Manager scope in URL authorization. Use this scope to manage access to the corresponding URL. When you configure an application, virtual directory, or URL for URL authorization, you must create a scope in the *authorization policy store*, a database that stores the authorization manager policies and that has the same name as the name that is specified in the AzScopeName property of the corresponding metabase entry.

Listing 5.3 provides a customizable text file, written in VBScript, that you can use to mark the root of the first site as a URL in MyAZScope, which is defined in the MyAZStore.xml file. Users with URLAccess rights in this scope will be able to access the site.

**Listing 5.3 Sample Script for URL Authorization**

```
var objVDir = GetObject("IIS://localhost/w3svc/1/root")
objVDir.AzEnable = true
objVDir.AZStoreName = "MSXML://d:\MyAZStore.xml"
objVDir.AzScopeName = "MyAZScope"
objVDir.AZImpersonationLevel = 0
objVDir.SetInfo()
```

# Securing FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Although you can configure the FTP service to require a valid user name and password combination, neither the credentials that are specified at logon nor the data itself is encrypted or encoded in any way. The credentials are sent across the network in plaintext and can be intercepted and analyzed by any station on any network between the FTP client and the FTP server. If your plaintext credentials are intercepted and analyzed by a malicious user, someone other than the intended user can log on to your FTP site and download the files that you placed there or gain access to other network resources, even network resources that require Windows account authentication. For more information about setting the authentication method for your FTP site, see Configuring FTP Site Properties.

If you intend to place sensitive data on your FTP site or if communication between clients and your FTP server is important, consider using FTP over an encrypted channel, such as a virtual private network (VPN) secured with Point-to-Point Tunneling Protocol (PPTP) or Internet Protocol security (IPSec). To learn how to set up a private network over the Internet using PPTP or how to set up IP communications with a VPN and IPSec, see the PPTP, VPN, and IPSec topics in Help and Support Center for Windows Server 2003.

If you need to authorize users or host sensitive content, WebDAV is a good alternative to FTP because WebDAV can take advantage of SSL encryption for communications. For more information about WebDAV, see Securing WebDAV Publishing Directories.

# Securing WebDAV Publishing Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

WebDAV is an extension of the HTTP 1.1 protocol that facilitates file and directory manipulation over an HTTP connection. Through the use of WebDAV *verbs*, or commands, properties can be added to and read from files and directories. Files and directories can also be remotely created, deleted, moved, or copied. Additional access control can be configured through both Web server permissions and NTFS.

Just as with a standard Web site, you can use the **Directory Security** tab in the **Web Site Properties** dialog box to configure WebDAV authentication, Web site access control, and IP address and domain name restrictions, and to enable secure communications. Furthermore, for security reasons and to enable WebDAV custom search properties, your WebDAV publishing directory must reside on an NTFS file system partition. For more information about NTFS partitions, see "Choosing a file system: NTFS, FAT, or FAT32" in Help and Support Center for Windows Server 2003.

**Setting the Client Authentication Method**

Because WebDAV is an HTTP-based protocol, the same authentication methods are available for your WebDAV publishing directory as for your Web sites. When you create a WebDAV virtual directory, the authentication settings will be set by default to Anonymous authentication and Integrated Windows authentication.

| **Important** |
| --- |
| Change the default authentication settings for your WebDAV site. Anonymous authentication and Integrated Windows authentication are appropriate for Web sites, but not for resources on a WebDAV publishing directory. |

Use the following best practices as guidelines when configuring client authentication for WebDAV. For more information about authentication, see Authentication in IIS 6.0.

**Anonymous authentication**

Turn off Anonymous authentication to your WebDAV publishing directory. Without controlling access, your directory might be vandalized by unknown clients.

**Basic authentication**

Turn on Basic authentication only if you encrypt passwords through SSL. Basic authentication sends passwords over the connection in plaintext, so passwords can be intercepted and read. For more information about SSL, see Configuring SSL on a Web Server or Web Site.

**Digest authentication**

Turn on Digest authentication if users access your WebDAV publishing directory over the Internet and through firewalls because Digest authentication sends the passwords over the network as an MD5 hash. However, passwords are stored as plaintext in the Active Directory® directory service.

**Advanced Digest authentication**

Turn on Advanced Digest authentication if users access your WebDAV publishing directory over the Internet and through firewalls. Advanced Digest authentication is the best choice because, in addition to sending passwords over the network as an MD5 hash, the passwords are also stored in Active Directory as an MD5 hash rather than as plaintext.

**Integrated Windows authentication**

Turn off Integrated Windows authentication unless users access your WebDAV publishing directory only from an intranet.

**.NET Passport authentication**

Turn off .NET Passport authentication. Microsoft® .NET Passport authentication uses cookies to validate user credentials against a central .NET Passport server. However, the central .NET Passport server does not authorize or deny the access of a specific user to individual .NET Passport–enabled sites.

**Configuring Access Control**

You can control access to your WebDAV publishing directory by coordinating permissions in IIS and Windows Server 2003. The following information can help you set the correct permissions depending on what kind of content your WebDAV publishing directory contains and what you want to allow users to do with that content. For step-by-step instructions for setting permissions, see Securing Sites with Web Site Permissions.

**Read, Write, Directory browsing enabled**

Set these permissions if you want clients to see a list of resources, modify them (except for resources that do not have Write

permission), publish their own resources, or manipulate files.

**Write enabled; Read and Directory browsing disabled**
Set these permissions if you want clients to publish private information on the directory, but you do not want others to see what has been published. This configuration works well if clients are submitting ballots or performance reviews.

**Read and Write enabled; Directory browsing disabled**
Set these permissions only if you want to obscure file names as a security method. However, be aware that security by obscurity is a low-level security precaution because an attacker or a malicious user can guess file names by trial and error.

**Index this resource enabled**
Enable Indexing Service if you plan to let clients search directory resources.

You can also control access to your WebDAV publishing directory through discretionary access control lists (DACLs). When you set up a WebDAV publishing directory on an NTFS file system disk, the server running Windows Server 2003 is configured with a secure set of defaults that can be too restrictive for Web publishing. If you change the defaults, assign Write permission only to the specific users or groups that absolutely need Write permission. (Whenever possible, assign permissions to groups rather than users.) By default, the Users group is assigned only Read permission for most portions of the file system.

If you have script files in your WebDAV publishing directory that you do not want users to access, you can deny access to these files by verifying that the Script Source Access permission is not assigned. The extensions of your script files are listed in the Applications Mapping list. All other executable files, including files with .exe extensions, are treated as static HTML files unless the Scripts and Executables permission is assigned to the directory.

You can prevent users from downloading and viewing .exe files as HTML files but still allow clients to run them by assigning the Scripts and Executables permission. To do so, on the **Virtual Directory** tab of the publishing directory, on the **Execute permissions** list, click **Scripts and Executables**. When you assign the Scripts and Executables permission to the directory, all of the executable files are subject to the Script Source Access permission setting. You can assign the Script Source Access permission on the **Home Directory** tab of the **Web Site Properties** dialog box of the Web site that you are configuring. When the Script Source Access permission is assigned, clients with Read permission can see all the executables, and clients with Write permission can edit and run all the executables.

If you want clients to be able to write to an executable file that does not appear in the Application Mapping list, assign the Write permission and the Scripts Only Execute permission to the directory.

If you want clients to be able to write to any executable file, regardless of whether it appears in the Application Mapping list, assign the Script Source Access permission and the Scripts and Executables Execute permission to the directory.

# Controlling Access to a Database on a Web Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For Web applications that communicate with a database, controlling access to the database is a critical element of Web server security. There are several ways to control access to a database that is published on an IIS 6.0 Web server:

- Use the built-in security permissions of the database application or the database management system. Many of the integrated user authentication methods in your database provide control access at a fine level of granularity. For example, the Database Results Wizard in Microsoft FrontPage 2002 allows you to create a Web page that can access a database, and to use password protection on the database connection. For more information about Microsoft SQL Server security, see the Help topics in SQL Server Enterprise Manager. For more information about Microsoft Access 2002 Security, see the Help topics that are accessible from the Help menu.

- Use the data source name (DSN). The DSN that you create on the Web server is used by an external program or an ASP page to refer to the database that you want to publish on your Web site. For more information about how to determine the DSN password setting for a database or database management system, consult the application's help.

- Use NTFS permissions. Windows Server 2003 NTFS permissions can restrict access to specific folders and files on your Web site. For more information about NTFS, see Help and Support Center for Windows Server 2003.

- Use subwebs to restrict access to a section of the Web site. For example, if you use FrontPage to create a Web site, you can create security boundaries by using subwebs. Each subweb can maintain separate security settings, so you can store database results pages or ASP pages that reference the database in a subweb with unique permissions. For more information about creating subwebs and assigning unique permissions, see Knowledge Base article 301432, How to Create a Subweb and Add Permissions Using FrontPage 2000.

- Use an appropriate Web server authentication method. Database user authentication can depend on the protocol that is used for the database connection. For example, if you decide to use the Named Pipes default connection protocol for SQL Server, Windows user account credentials might be authenticated with SQL Server authentication.

# Encryption in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Encryption is the process of scrambling information by applying a mathematical function in such a way that it is extremely difficult for anyone other than an intended recipient to retrieve the original information. Central to this process is a mathematical value, called a *key*, which is used by the function to scramble the information in a unique and complex way.

Your Web server uses essentially the same encryption process to secure communication links with users. After establishing a link, a special *session key* is used by both your Web server and the user's Web browser to both encrypt and decrypt information. For example, when an authenticated user attempts to download a file from a Web site requiring a SSL–secured channel, your Web server uses a session key to encrypt the file and related HTTP headers. After receiving the encrypted file, the Web browser then uses a copy of the same session key to recover the file.

This method of encryption has an inherent drawback: During the process of creating a link, a copy of the session key might be transmitted across an unsecured network. This means that a malicious user can compromise the link by intercepting and stealing the session key. To safeguard against this possibility, however, your Web server implements an additional method of encryption.

# About Secure Sockets Layer

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 uses the Secure Sockets Layer (SSL) protocol, which provides a way of establishing an encrypted communication link with users. SSL is a public key-based security protocol that is used by Internet services and clients to authenticate each other and to establish message integrity and confidentiality. SSL uses certificates for authentication, and it uses encryption for message integrity and confidentiality. To establish encrypted communications using SSL, you must have a valid server certificate installed.

SSL confirms the authenticity of your Web site and, optionally, confirms the identity of users that are accessing restricted Web sites. You can use the SSL encryption features of your Web server to cryptographically encode information transmitted and received by your Web server. When you enable SSL encryption, you can prevent unauthorized individuals from decoding the original content of your transmissions. For example, you can use encryption to help protect the private information that clients exchange with your server, such as credit card numbers or phone numbers.

When you set security properties for a specific Web site, you automatically set the same security properties for directories and files belonging to that site, unless the security properties of the individual directories and files have been previously set.

Your Web server will prompt you for permission to reset the properties of individual directories and files when you attempt to set security properties for your Web site. If you choose to reset these properties, your previous security settings will be replaced by the new settings. The same condition applies when you set security properties for a directory containing subdirectories or files with previously set security properties.

To maintain the performance level of your Web server, consider using SSL encryption only for sensitive information, such as financial transactions. Encrypted transmissions can significantly reduce transmission rates and server performance.

# Session Keys and Browsers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

During the exchange of information that establishes a SSL–secured connection, the server creates a *session key*, also called an encryption key. IIS and the browser use the session key to encrypt and decrypt information that is exchanged during a session. The required encryption strength is measured in bits (for example, 40-bit encryption or 128-bit encryption). The encryption strength provides the degree of confidentiality that is deemed necessary for the application. A large number of bits indicates greater encryption strength because attackers who do not have the session key find it difficult to decrypt larger encrypted strings. Although greater encryption key strengths offer greater security, they also require more server resources to implement. The session key for your Web server is typically 128 bits long. The session key is not the same as an SSL key pair, which is used to negotiate and establish a communication link.

You can configure your Web server to require a 128-bit minimum session-key strength, the default for Microsoft Windows .NET, for all SSL–secured communication sessions. If you set a minimum 128-bit key strength, however, users attempting to establish a SSL–secured communications channel with your server must use a browser capable of communicating with a 128-bit session key.

You can require users to establish an encrypted channel (https:// rather than http://) with your server before accessing a restricted Web site, directory, or file. The use of an encrypted channel, however, requires that the user's Web browser and your Web server both support the encryption scheme used to secure the channel. Specifically, when you enable your Web server's default communication security settings, you require the user's Web browser to support a session key strength of 40 bits, or greater.

# Session Key Encryption Strength

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A session key's *strength* is proportional to the number of binary *bits* that comprise the session key file. This means that session keys with a greater number of bits have a greater degree of security, and are considerably more difficult to forcibly decode.

When a user attempts to establish an encrypted communication channel with your Web server, the user's browser must negotiate the strongest possible level of encryption—that is, session key strength—that can be used to secure communications over that channel. This means that both your Web server and the user's browser must be equipped with compatible session key encryption and decryption capabilities. For example, when you configure your Web server to require a session key with a minimum of 40-bit (default) encryption strength, a user attempting to secure a connection must have a Web browser capable of processing information with a 40-bit session key. If you require 128-bit key encryption, your users must use Web browsers that support 128-bit encryption.

# Public Key Encryption

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Your Web server's SSL security feature uses *public key* encryption to shield the session key from interception during transmission. Public key encryption involves the use of two additional keys: a *private* and a *public* key.

Figure 5.6 shows how public key encryption works, and the steps that follow describe the process in more detail.




1. The user's Web browser establishes a communication link with your Web server by using the https:// protocol.

2. The user's Web browser and your server engage in negotiation to determine the degree of encryption to use for securing communications.

3. Your server sends the browser its certificate, which includes a public key.

4. The user's Web browser generates a session key, encrypts it with the server's public key, and sends it to your server.

5. Using the private key, your server decrypts the message and retrieves the session key sent by the Web browser.

6. Your Web server and the browser both use the session key to encrypt and decrypt transmitted data.

The private key plays an important role in ensuring that your communication link remains secured. You should take every reasonable precaution to protect the private key from loss or theft. If you suspect that your private key has been compromised, notify your certification authority and then use the Web Server Certificate Wizard to create a new certificate request so you can obtain a new server certificate.

# Kernel-Mode SSL

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

In Windows Server 2003 Service Pack 1 (SP1) and later, you can run SSL in *kernel mode* instead of the default *user mode*. Kernel mode operation means that components or processes run in the core address space of the operating system.

Running SSL in kernel mode improves SSL performance by moving encryption and decryption operations to the kernel, thereby reducing the number of transitions between kernel mode and user mode.

The following restrictions apply to running SSL in kernel mode:

- Client certificates are not supported.
- RC2 ciphers are not supported.
- The Private Communications Technology (PCT) 1.0 protocol is not supported.
- Configuration changes to server certificates require a restart of the HTTP service.
- ISAPI GetServerVariable calls for certificate information do not work.
- Bulk encryption offload is not supported.
- ISAPI filters do not get READ RAW notifications regardless of the security of the connection. This restriction affects only IIS 5.0 compatibility mode, because IIS 6.0 worker process isolation mode does not support READ RAW DATA filters regardless of the setting for kernel-mode SSL.

For information about enabling kernel-mode SSL, see Enabling Kernel Mode SSL.

# SSL and SF_NOTIFY_READ_RAW_DATA Filters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 6.0, SSL communication must use the SF_NOTIFY_READ_RAW_DATA notification in worker process isolation mode. However, these notifications are not supported in IIS 6.0 in worker process isolation mode. For this reason, IIS creates a new SSL service that is associated with worker process isolation mode. This service runs entirely outside of any IIS process; currently, it runs as lsass.exe, although it can be configured to run as an instance of Svchost.exe. For raw, unencrypted SSL data, Http.sys uses this new SSL service in much the same way it uses Inetinfo.exe in IIS 5.0 isolation mode for raw data filter support. However, in IIS 5.0 isolation mode, the SSL code runs in Inetinfo.exe instead of in its own service.

For more information about how SF_NOTIFY_READ_RAW_DATA notifications are handled in IIS 6.0 isolation modes, see Application Isolation Modes.

# SSL and Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Certificates* are digital identification documents that allow both servers and clients to authenticate each other. If you want the server and client browser to set up an SSL connection over which encrypted information can be sent, certificates are required. Certificates include keys that are used to establish an SSL-encrypted connection. A *key* is a unique value that is used to authenticate the server and the client in establishing an SSL connection. A public key and a private key form an *SSLkey pair*. The World Wide Web Publishing Service (WWW service) on your Web server uses this key pair to negotiate an encrypted connection with the client browser.

The certificate-based SSL features in IIS consist of a server certificate, a client certificate, and various digital keys. You can use Certificate Services to create certificates, or you can obtain certificates from a mutually trusted third-party organization called a certification authority (CA). For more information about obtaining, installing, and managing certificates, see Certificates_IIS_SP1_Ops.

*Server certificates* provide a way for users to confirm the identity of your Web site before they transmit personal information, such as a credit card number. A server certificate contains detailed identification information, such as the name of the organization that is affiliated with the server content, the name of the organization that issued the certificate, and a public key that is used to establish an encrypted connection. This information helps to assure users of the authenticity of Web server content and the integrity of the SSL–secured connection.

With SSL, your Web server also has the option of authenticating users by checking the contents of their *client certificates*. A typical client certificate contains detailed identification information about a user and the organization that issued the certificate, as well as a public key. You can use client certificate authentication, along with SSL encryption, to implement a method for verifying the identity of your users.

Use the Web Server Certificate Wizard to either generate a certificate request file (NewKeyRq.txt, by default) that you can send to a certification authority, or to generate a request for an online certification authority, such as Microsoft® Certificate Services. For more information, see Obtaining Server Certificates.

If you are not using Microsoft Certificate Services to issue your own server certificates, a third-party certification authority must approve your request and issue your server certificate. Depending on the level of identification assurance offered by your server certificate, you can expect to wait anywhere from several days to several months for the certification authority to approve your request and send you a certificate file. You can have only one server certificate per Web site.

After you receive a server certificate file, use the wizard to install it. The installation process attaches, or *binds*, your certificate to a Web site. For more information about installing certificates, see Requesting a New Server Certificate from an Online CA.

**Certification Authorities**

You can obtain a certificate from a mutually trusted, third-party commercial organization called a *certification authority* (CA). The primary responsibility of the CA is to confirm the identity of the party that is seeking a certificate, thus ensuring the validity of the identification information that is contained in the certificate.

Before issuing a certificate, the CA requires you to provide identification information, such as name, address, and organization. The extent of this information can vary with the identification assurance requirements of the certificate. If you need a certificate to provide absolute assurance about your identity, the CA will require substantial information from you; gathering this information might require a personal interview with the CA as well as the endorsement of a notary.

Ultimately, the success of certification authentication depends on whether the party receiving a certificate trusts the authority who issued the certificate, and that the authority correctly verified the certificate owner's identity. However, beyond this trust, certificates do not provide conclusive proof of the identity, trustworthiness, or intentions of the user or server.

**Certificate Trust Lists**

You can configure computers running Windows Server 2003 with IIS 6.0 to accept certificates from a predefined list of certification authorities (CAs). For example, you might create a different list of trusted certification authorities for each department on your network. IIS would accept only client certificates supplied by certification authorities in the CTL for that department. Each Web site can be configured to accept certificates from a different list by using certificate trust lists (CTLs). You can automatically verify client certificates against the CTL. You can use the Certificate Trust List Wizard to create and edit CTLs and to add new root certificates to your CTLs.

**Certificate Revocation Lists**

Certification authorities (CAs) cannot physically revoke your users' client certificates. However, they might publish Certificate

Revocation Lists (CRLs) that are copied onto your computer, where they can be searched for client certificates that are in revoked status. This process, called CRL checking, allows IIS to reject client certificates that do not meet the CA criteria for validity. CRLs pair subscribers and revocation status, along with a description of the reason for the revocation. Most CRL services also include a prospective date for the next CRL to be published. Frequent updates help you stay current with subscribers' status with the CA. When determining the schedule for CRL publishing, there is a trade-off between performance and security. The more often CRLs are published, the more current clients can be with the list of revoked CRLs and the less likely they are to accept recently-revoked certificates. However, frequent publishing can negatively impact network and client performance. To help mitigate this trade off, delta CRLs can be used if the environment supports them. For more information about certificate revocation and CRLs, see Revoking Certificates and Publishing CRLs.

CRL checking is managed by setting metabase properties. The metabase properties that control CRL checking can be set or viewed using a COM object, or WMI scripts, or ADSI scripts.

**Enabling and Disabling CRL Checking**

You can enable and disable CRL checking by setting the CertCheckMode Metabase Property. CRL checking is enabled by default. The CRL will be refreshed by the CA whenever a new CRL is issued, unless you set a different CRL refresh interval by specifying a value other than 0 for the CertCheckMode metabase property.

**Configuring CRL Refreshes**

Optionally, you can refresh the CRL on your computer with the CRL on the CA server even when the cached CRL on your computer is valid, by changing the setting of the RevocationFreshnessTime Metabase Property.

You can change the length of time allocated for downloading a CRL to a time, in seconds, by setting the RevocationURLRetrievalTimeout Metabase Property.

# Server Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To activate your Web server's SSL security features, you must obtain and install a valid server certificate. Server certificates are digital identifications containing information about your Web server and the organization that is sponsoring the server's Web content. A server certificate enables users to authenticate your server, check the validity of the Web content, and establish an encrypted connection. The server certificate also contains a public key, which is used in creating an encrypted connection between the client and server.

The success of a server certificate as a means of identification depends on whether the user trusts the validity of the information that is contained in the certificate. For example, a user logging on to your company's Web site might be hesitant to provide credit card information, despite having viewed the contents of your company's server certificate. This might be especially true if your company is new and not well known. If this is the case, consider obtaining a server certificate from a certification authority.

Alternatively, depending on your organization's relationship with its Web site users, you can issue your own server certificates. For example, in the case of a large corporate intranet handling employee payroll and benefits information, corporate management might decide to run a certificate service assuming responsibility for validating identification information and issuing server certificates. For more information see Obtaining and Installing Server Certificates.

# Using Certificate Wizards

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 includes the following wizards that simplify the tasks necessary to establish and maintain communications for a Web site using server certificates:

- Web Server Certificate Wizard. Obtain and manage server certificates, which are used in negotiating a link between your server and a user's browser.
- Certificate Trust List Wizard. Manage certificate trust lists (CTLs), which are lists of certification authorities that are trusted by each Web site or virtual directory.

**The Web Server Certificate Wizard**

In IIS, you can obtain, configure, and renew server certificates using the Web Server Certificate Wizard. The wizard detects whether a server certificate has already been installed and whether it is about to expire. Use the wizard to perform the following tasks:

- Create a certificate request.
- Replace the server certificate with another one from a certification authority (CA), from an online CA, such as Microsoft Certificate Services, or from a file previously obtained in Key Manager.
- Reassign a certificate from one Web site to another Web site.
- View certificates.

When creating a new certificate, the Web Server Certificate Wizard allows you to choose the strength of encryption, the type of certificate, and a cryptographic service provider for your certificate.

Online requests for server certificates can be made only to local and remote Enterprise Certificate Services and remote Stand-alone Certificate Services. The IIS Web Server Certificate Wizard does not recognize a stand-alone installation of Certificate Services on the same computer when requesting a certificate. To get around this, use the offline certificate request to save the request to a file and then process as an offline request. For more information about offline requests, see "Certificate Services Help" in Help and Support Center for Windows Server 2003.

If you are not using an online certification authority, you will need to save the request file generated by the Web Server Certificate Wizard to disk and send it to the CA. When the response is received, you can start the wizard and it will begin where it left off. If you are replacing a certificate, IIS will continue to use the old certificate until the new request is completed.

**The Certificate Trust List Wizard**

Use the Certificate Trust List Wizard to obtain and manage certificate trust lists (CTLs). A CTL is a list of trusted certification authorities (CA) for a particular Web site. By configuring your CTL, you can allow client certificates issued by one CA to be used but not from another CA. CTLs are especially useful for Internet Service Providers (ISPs) who have several Web sites on their server and who need to have a different list of approved certification authorities for authenticating clients at each site. CTLs are available only at the Web site level and are not available for FTP sites.

# Obtaining and Installing Server Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can obtain server certificates from an outside certification authority (CA), or you can issue your own server certificates by using Microsoft Certificate Services.

Use Microsoft Certificate Services to create a customizable service for issuing and managing certificates. You can create server certificates for the Internet or for corporate intranets, giving your organization complete control over certificate management policies.

**Obtaining Server Certificates from a certification authority**

If you are replacing your current server certificate, IIS will continue to use the old certificate until the new request has been completed.

When you are choosing a certification authority (CA), consider the following questions:

- Will the CA be able to issue a certificate that is compatible with all of the browsers used to access my server?
- Is the CA a recognized and trusted organization?
- How will the CA provide verification of my identity?
- Does the CA have a system for receiving online certificate requests, such as requests generated by the IIS Web Server Certificate Wizard?
- How much will the certificate cost initially, and how much will renewal or other services cost?
- Is the CA familiar with my organization or my company's business interests?

**Issuing Your Own Server Certificates**

When deciding whether to issue your own server certificates, consider the following:

- Microsoft Certificate Services accommodates different certificate formats and provides for auditing and logging of certificate-related activity.
- Compare the cost of issuing your own certificates against the cost of buying a certificate from a certification authority.
- Remember that your organization will require an initial adjustment period to learn, implement, and integrate Certificate Services with existing security systems and policies.
- Assess the willingness of your connecting clients to trust your organization as a certificate supplier.

**Assigning Resources to Server Certificates**

Use the following guidelines when assigning IP addresses, Web sites, and SSL ports to your server certificates:

- Each Web site can have only one server certificate assigned to it.
- One certificate can be assigned to multiple Web sites.
- You can assign multiple IP addresses per Web site.
- You can assign multiple SSL ports per Web site.

**Installing Server Certificates**

After obtaining a server certificate from a CA, or after issuing your own server certificate using Certificate Services, use the Web Server Certificate Wizard to install it. When you use the Server Certificate Wizard to obtain and install a server certificate, the process is referred to as *creating* and *assigning* a server certificate. For instructions on how to install your server certificate, see Requesting a New Server Certificate from an Online CA.

**Backing Up Server Certificates**

The Web Server Certificate Wizard is used to back up server certificates. Because IIS works closely with Windows, you can also use Certificate Manager, which is called "Certificates" in Microsoft Management Console (MMC), to export and back up your server certificates.

# Managing Certificates Programmatically

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 includes the new certificate object, **IISCertObj**, which allows you to use scripts to process SSL certificates on your local computer or on a remote computer. You can use scripts in conjunction with **IISCertObj**, a COM object, to perform the following tasks:

- Import copies of a certificate from a .pfx file (a file containing a PFX-encrypted certificate and a private key) from one server onto multiple servers.
- Save backup copies of certificates on a central site.
- Copy a certificate from one server onto multiple servers.

You can use **IISCertObj** to perform the following tasks:

- Import copies of a server certificate from a .pfx file (a file containing a PFX-encrypted server certificate and a private key) from a server to multiple servers.
- Save backup copies of server certificates on a central site.
- Copy a server certificate from one server to multiple servers.

When using **IISCertObj**, you must set properties for the methods that you call. The following sections outline the necessary properties and the methods that you need to manage server certificates.

**Setting IISCertObj Properties**

Because a method will return an error message if the required properties for that method are not available when the method is called, you should set properties before you call methods. Table 5.11 lists the **IISCertObj** properties that you need to generate server certificate requests.

**Table 5.11 IISCertObj Properties Required for Server Certificate Requests**

| Property Attribute | Name Attribute | Value | Description |
|---|---|---|---|
| ServerName | Data Type | String, write only | Contains the name of the computer on which server certificate operations are executed. |
| UserName | Data Type | String, write only | Specifies the user name that is used to log on to the remote computer. If UserName is empty, the credentials of the user that is currently logged on are used. |
| UserPassword | Data Type | String, write only | Specifies the password that is used for logging onto the remote computer. |
| InstanceName | Data Type | String, write only | Identifies the targeted metabase instance. W3svc/1 is an example. |

**Applying IISCertObj Methods**

You can apply the following methods to the **IISCertObj** object: **Copy**, **Export**, **GetCertInfo**, **Import**, **ImportToCertStore**,

**IsExportable**, **IsInstalled**, **Move**, and **RemoveCert**.

**Copy**

The **Copy** method allows you to copy a server certificate from one instance of IIS to another and from one computer to another. Use the following syntax with the **Copy** method:

```
Copy bAllowExport, bOverwriteExisting, DestinationServerName,
DestinationServerInstance, DestinationServerUserName, DestinationServerPassword
```

Table 5.12 lists the parameters that the **Copy** method uses.

**Table 5.12 Copy Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| *bAllowExport* | Boolean | Indicates whether the imported server certificate will be exportable (**true**) or not exportable (**false**). |
| *bOverwriteExisting* | Boolean | Indicates whether, if an identical server certificate already exists, it is overwritten (**true**) or not overwritten (**false**). |
| *DestinationServerName* | String | Specifies the computer name of the server to which the server certificate will be moved. |
| *DestinationServerInstance* | String | Specifies the IIS metabase instance at which the server certificate will be pointed. For example, w3svc/1. |
| *DestinationServerUserName* | String | Specifies the user name for the destination server. If empty, the user name of the user who is currently logged on will be used. |
| *DestinationServerPassword* | String | Specifies the password for the destination server when the *DestinationServerUserName* parameter is used. |

**Export**

The **Export** method exports a copy of a server certificate to a file. The target file can be on a local or remote computer. Use the following syntax with the **Export** method:

```
Export FileName, Password, bPrivateKey, bCertChain, bRemoveCert
```

Table 5.13 lists the parameters that the **Export** method uses.

**Table 5.13 Export Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| *FileName* | String | Specifies the name of the targeted file. For example, C:\Mydir\Mycert.pfx. |
| *Password* | String | Specifies the password that is used to secure the file specified as *FileName*. |
| *bPrivateKey* | Boolean | Indicates whether the private key is exported (**true**) or not exported (**false**). |
| *bCertChain* | Boolean | Indicates whether the certificate trust chain is exported (**true**) or not exported (**false**). |
| *bRemoveCert* | Boolean | Indicates whether the binding of a server certificate is removed (**true**) or not removed (**false**). |

**GetCertInfo**

The **GetCertInfo** method allows you to acquire information on a server certificate that is installed on the computer specified by the **IISCertObj** properties, **ServerName** and **InstanceName**.

The **GetCertInfo** method uses no parameters.

**Import**

The **Import** method imports a copy of a server certificate from a file to a local or remote computer. Use the following syntax with the **Import** method:

```
Import FileName, Password, bAllowExport, bOverwriteExisting
```

Table 5.14 lists the parameters that the **Import** method uses.

**Table 5.14 Import Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| *FileName* | String | Specifies the name of the server certificate file to import. |
| *Password* | String | Specifies the password that is used to secure the file specified as *FileName*. |
| *bAllowExport* | Boolean | Indicates whether the imported server certificate is exportable (**true**) or not exportable (**false**). |
| *bOverwriteExisting* | Boolean | Indicates whether, if an identical server certificate already exists, it is overwritten (**true**) or not overwritten (**false**). |

**ImportToCertStore**

The **ImportToCertStore** method imports a copy of a server certificate from a file to the certificate store on a local or remote computer. Use the following syntax with the **ImportToCertStore** method:

```
ImportToCertStoreFileNamePasswordbAllowExportbOverwriteExisting
```

Table 5.15 lists the parameters that the **ImportToCertStore** method uses.

**Table 5.15 ImportToCertStore Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| *FileName* | String | Specifies the name of the server certificate file to import. |
| *Password* | String | Specifies the password that is used to secure the file specified as *FileName*. |
| *bAllowExport* | Boolean | Indicates whether the imported server certificate is exportable (**true**) or not exportable (**false**). |
| *bOverwriteExisting* | Boolean | Indicates whether, if an identical server certificate already exists, it is overwritten (**true**) or not overwritten (**false**). |

**IsExportable**

The **IsExportable** method allows you to determine whether a server certificate is exportable from a computer that is specified by the IIS **ServerName** and **InstanceName** properties.

The **IsExportable** method uses no parameters.

**IsInstalled**

The **IsInstalled** method allows you to determine whether a server certificate is installed on a computer that is specified by the **IISCertObj** properties, **ServerName** and **InstanceName**.

The **IsInstalled** method uses no parameters.

**Move**

The **Move** method allows you to move a server certificate from one instance of IIS to another, and from one computer to another. Use the following syntax with the **Move** method:

```
Move bAllowExport, bOverWriteExisting, DestinationServerName,
DestinationServerInstance, DestinationServerUserName, DestinationServerPassword
```

Table 5.16 lists the parameters that the **Move** method uses.

**Table 5.16 Move Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| *bAllowExport* | Boolean | Indicates whether the imported server certificate is exportable (**true**) or not exportable (**false**). |
| *bOverWriteExisting* | Boolean | Indicates whether, if an identical server certificate already exists, it is overwritten (**true**) or not overwritten (**false**). |
| *DestinationServerName* | String | Specifies the computer name of the server to which the server certificate will be moved. |

| | | |
|---|---|---|
| *DestinationServerInstance* | String | Specifies the IIS metabase instance at which the server certificate will be pointed. For example, w3svc/1. |
| *DestinationServerUserName* | String | Specifies the user name for the destination server. If no user name is entered, the user name of the user that is currently logged on will be used. |
| *DestinationServerPassword* | String | Specifies the password for the destination server when the *DestinationServerUserName* parameter is used. |

**RemoveCert**

The **RemoveCert** method allows you to delete a server certificate from a certificate store. Use the following syntax with the **RemoveCert** method:

```
Copy bRemoveFromCertStore, bPrivateKey
```

Table 5.17 lists the parameters that the **RemoveCert** method uses.

**Table 5.17 RemoveCert Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| *bRemoveFromCertStore* | Boolean | Indicates whether the server certificate is removed (**true**) or not removed (**false**). |
| *bPrivateKey* | Boolean | Indicates whether the private key of the server certificate is removed (**true**) or not removed (**false**). |

**Sample Server Certificate Management Scripts**

In IIS, if you use Windows Script Host or ASP, you can execute scripts that communicate with **IISCertObj** programmatically to import, archive, or export server certificates.

**Importing server certificates to multiple servers**

Large sites often need to use one server certificate for user logon to multiple servers. It would be time-consuming to use the Microsoft Management Console (MMC) to add copies of the server certificate to each server. By using the **Import** method of **IISCertObj**, you can install copies of the same server certificate on all of the targeted servers.

To import a server certificate from a certificate store on one server to other servers, first save the following script in your text editor as CertImport.vbs.

**Important**

You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyMachine******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters).

```
Option Explicit
Dim iiscertobj, pfxfile, pfxfilepassword, InstanceName, WebFarmServers, IISServer
Set iiscertobj = WScript.CreateObject("IIS.CertObj")
pfxfile = WScript.Arguments(0)
pfxfilepassword = WScript.Arguments(1)
InstanceName = WScript.Arguments(2)
WebFarmServers = split(WScript.Arguments(3), ",")
iiscertobj.UserName = WScript.Arguments(4)
iiscertobj.UserPassword = WScript.Arguments(5)
For Each IISServer in WebFarmServers
  iiscertobj.ServerName = IISServer
  iiscertobj.InstanceName = InstanceName
  iiscertobj.Import pfxfile, pfxfilepassword, true, true
Next
```

After you save the script as CertImport.vbs, run this command:

```
Certimport.vbs cert.pfx pfxPassword w3svc/1 iisserver1,iisserver2,iisserver3
Administrator aal34290
```

**Saving backup copies of server certificates in a central archive**

Using the **Export** method of **IISCertObj**, you can archive a backup of each server certificate on your server farm to a central site.

To export copies of server certificates to the central site, first save the following script in your text editor as Save_all_certs.vbs.

```
Option Explicit
Dim iiscertobj, targetServer, targetServers, pfxbasename, pfxpassword, InstanceName
Set iiscertobj = WScript.CreateObject("IIS.CertObj")
pfxbasename = WScript.Arguments(0)
pfxpassword = WScript.Arguments(1)
InstanceName = WScript.Arguments(2)
targetServers = split(WScript.Arguments(3), ",")
iiscertobj.UserName = WScript.Arguments(4)
iiscertobj.UserPassword = WScript.Arguments(5)
iiscertobj.InstanceName = InstanceName
For Each targetServer in targetServers
  iiscertobj.ServerName = targetServer
  iiscertobj.Export pfxbasename + targetServer + ".pfx", pfxpassword, true, false,
false
Next
```

After you save the script as Save_all_certs.vbs, run this command:

```
Save_all_certs.vbs C:\certbackup\ pfxpassword w3svc/1 iisserver2,iisserver3,iisserver4
Administrator aal34290
```

**Copying a server certificate from an existing server to a new server**

Using the **Copy** method of **IISCertObj**, you can copy a server certificate to a new server after you add the server to a server farm.

To copy a server certificate from an existing server to a new server, first save the following script in your text editor as Certcopy.vbs.

```
Option Explicit
Dim iiscertobj, targetServer, targetServers, targetInstance
Set iiscertobj = WScript.CreateObject("IIS.CertObj")
iiscertobj.ServerName = WScript.Arguments(0)
iiscertobj.Instancename = WScript.Arguments(1)
targetServers = split(WScript.Arguments(2), ",")
targetInstance = WScript.Arguments(3)
iiscertobj.UserName = WScript.Arguments(4)
iiscertobj.UserPassword = WScript.Arguments(5)
For Each targetServer in targetServers
  iiscertobj.Copy true, true, targetServer, targetInstance
Next
```

After you save this script as Certcopy.vbs, run this command:

```
Certcopy.vbs iisServer1 w3svc/1 iisServer2 w3svc/1 Administrator asdf-0324
```

Listing 5.4 provides a customizable text file that you can use to import a certificate from a .pfx file and assign it to the desired site.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyMachine******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Listing 5.4 Sample Script for Importing Certificates**

```
' certimp.vbs
'
' author: uladm   05/03/2002
'
' certimp.vbs for usage info
'


REM
REM LOCALIZATION
REM

L_SWITCH_PFX      = "-c"
L_SWITCH_PFX_PASSWORD   = "-p"
L_SWITCH_INSTANCE_ID   = "-i"
L_SWITCH_SERVER      = "-s"
L_SWITCH_USER      = "-u"
L_SWITCH_PASSWORD   = "-pwd"

L_DESC_PROGRAM   = "certimp - imports a certificate from a pfx file and assigns it to
the desired site"

L_DESC_PFX      = "Specify pfx file name"
L_DESC_PFX_PASSWORD   = "pfx password"
L_DESC_INSTANCE_ID   = "<site> Specify server instance (site) name"
L_DESC_SERVER      = "<server> Specify computer to configure"
L_DESC_USER      = "User name"
L_DESC_PASSWORD      = "User password"

L_DESC_EXAMPLES         = "Examples:"
L_DESC_EXAMPLE1         = "certimp.vbs -c cert.pfx -p pfxpassword -i w3svc/1 -s
iisserver1 -u Administrator -pwd aal34290"

L_ERR_IMPORT_ERROR   = "Error importing certificate: "
L_ERR_IMPORT_SUCCESS   = "Import succeeded."

REM
REM END LOCALIZATION
REM

REM
REM --- Globals ---
REM

dim g_dictParms
set g_dictParms = CreateObject ( "Scripting.Dictionary" )

REM
REM --- Set argument defaults ---
REM

g_dictParms(L_SWITCH_PFX)      = ""
g_dictParms(L_SWITCH_PFX_PASSWORD)   = ""
g_dictParms(L_SWITCH_INSTANCE_ID)   = "w3svc/1"
g_dictParms(L_SWITCH_SERVER)      = ""
g_dictParms(L_SWITCH_USER)      = ""
g_dictParms(L_SWITCH_PASSWORD)      = ""

REM
REM --- Begin Main Program ---
REM

if NOT ParseCommandLine ( g_dictParms, WScript.Arguments ) then
   usage
   WScript.Quit ( 0 )
end if

pfxfile = g_dictParms(L_SWITCH_PFX)
pfxfilepassword = g_dictParms(L_SWITCH_PFX_PASSWORD)
instance = g_dictParms(L_SWITCH_INSTANCE_ID)
server = g_dictParms(L_SWITCH_SERVER)
user = g_dictParms(L_SWITCH_USER)
```

```
userpassword = g_dictParms(L_SWITCH_PASSWORD)

REM
REM   Check arguments
REM
if ( pfxfile = "" ) then
   usage
   WScript.Quit 0
end if

REM
REM   Import certificate
REM
dim iiscertobj
set iiscertobj = WScript.CreateObject( "IIS.CertObj" )
if (user <> "") then
   iiscertobj.UserName = user
   if (userpassword <> "") then
      iiscertobj.UserPassword = userpassword
   end if
end if
if (server <> "") Then
   iiscertobj.ServerName = server
end if
iiscertobj.InstanceName = instance
iiscertobj.Import pfxfile, pfxfilepassword, true, true

REM
REM   Show Result
REM
if (err.number <> 0) then
   WScript.echo L_ERR_IMPORT_ERROR
   WScript.echo Err.Description & " (Error 0x" & Hex(Err.Number) & ")"
else
   WScript.echo L_ERR_IMPORT_SUCCESS
end if
WScript.Quit 0

REM
REM
REM --- End Main Program ---
REM
REM

REM
REM ParseCommandLine ( dictParameters, cmdline )
REM     Parses the command line parameters into the given dictionary
REM
REM Arguments:
REM     dictParameters  - A dictionary containing the global parameters
REM     cmdline - Collection of command line arguments
REM
REM Returns - Success code
REM

Function ParseCommandLine ( dictParameters, cmdline )
    dim     fRet
    dim     cArgs
    dim     i
    dim     strSwitch
    dim     strArgument

    fRet    = TRUE
    cArgs   = cmdline.Count
    i       = 0

    do while (i < cArgs)

        REM
        REM Parse the switch and its argument
        REM
```

```
        if i + 1 >= cArgs then
            REM
            REM Not enough command line arguments - Fail
            REM

            fRet = FALSE
            exit do
        end if

        strSwitch = cmdline(i)
        i = i + 1

        strArgument = cmdline(i)
        i = i + 1

        REM
        REM Add the switch,argument pair to the dictionary
        REM

        if NOT dictParameters.Exists ( strSwitch ) then
            REM
            REM Bad switch - Fail
            REM

            fRet = FALSE
            exit do
        end if

        dictParameters(strSwitch) = strArgument

    loop

    ParseCommandLine = fRet
end function

REM
REM Usage ()
REM     prints out the description of the command line arguments
REM

Sub Usage
   WScript.Echo L_DESC_PROGRAM

   WScript.Echo vbTab & L_SWITCH_PFX & " " & L_DESC_PFX
   WScript.Echo vbTab & L_SWITCH_PFX_PASSWORD & " " & L_DESC_PFX_PASSWORD
   WScript.Echo vbTab & L_SWITCH_INSTANCE_ID & " " & L_DESC_INSTANCE_ID
   WScript.Echo vbTab & L_SWITCH_SERVER & " " & L_DESC_SERVER
   WScript.Echo vbTab & L_SWITCH_USER & " " & L_DESC_USER
   WScript.Echo vbTab & L_SWITCH_PASSWORD & " " & L_DESC_PASSWORD

   WScript.Echo L_DESC_EXAMPLES
   WScript.Echo L_DESC_EXAMPLE1
end sub
```

# Server-Gated Cryptography

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Server-Gated Cryptography (SGC) offers financial institutions a solution for worldwide financial transactions by using 128-bit encryption. SGC is an extension of SSL that enables financial institutions with export versions of IIS to use strong encryption.

SGC does not require an application to run on the client browser and will enable export clients to connect with 128-bit strength. Although SGC capabilities are built into IIS 4.0, 5.0, and 6.0, an SGC certificate is required to use SGC. Contact your CA for availability information.

If you want to support clients that run software without strong encryption support, you need only SGC certificates. On January 14, 2000, the United States government eliminated the remaining export controls for most computer hardware and software products that incorporate *strong encryption*, which was defined as products that implemented symmetric key encryption with key lengths of over 64 bits. Windows 2000 Server was the first platform to be shipped internationally under the new regulations with 128-bit+ encryption. For information about United States export regulations, see United States Commerce Department Bureau of Industry and Security. For more information about exporting Microsoft products, see Exporting Microsoft Products.

| **Note** |
| --- |
| If you are running Windows 2000 and you open your SGC certificate, you might receive the following notice on the **General** tab: **"The certificate has failed to verify for all its intended purposes."** This notice is issued because of the way SGC certificates interact with Windows 2000 Server and does not necessarily indicate that the certificate does not work correctly. |

# Cryptographic Service Provider

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Cryptographic service provider (CSP) allows you to select a Microsoft or third-party cryptographic provider to handle cryptography and certificate management. Each cryptographic provider can create a public and private key to encrypt the data that is sent to and from the Web server. The private key is stored at the server on hardware, on a Peripheral Component Interconnect (PCI) card, on a smart card, or in the registry, as it is for the two default providers that Microsoft installs: Microsoft DH SChannel Cryptographic Provider and Microsoft RSA SChannel Cryptographic Provider.

IIS certificate storage is now integrated with the Microsoft Cryptographic API (Crypto API). The Crypto API contains identical methods and properties for every cryptographic provider. This allows you to switch between cryptographic providers without having to rewrite code.

# Client Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Client certificates are electronic documents that contain information about clients. These certificates, like server certificates, contain not only this information but also public encryption keys that form part of the SSL security feature of IIS. The public keys, or encryption codes, from the server and the client certificates facilitate encryption and decryption of transmitted data over an open network, such as the Internet.

The typical client certificate contains several items of information: the identity of the user, the identity of the certification authority, a public key that is used for establishing encrypted communications, and validation information, such as an expiration date and serial number. Certification authorities offer different types of client certificates, which contain differing amounts of information, depending on the level of authentication that is required.

Your Web server cannot process client certificates unless you have previously installed a server certificate and enabled your server's communication securityfeatures.

# Obtaining Client Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can obtain a client certificate from a trustworthy, third-party organization, called a *certification authority* (CA). Before issuing a certificate, the CA might need to verify your identity. The identification information required by the CA varies, depending on the type of certificate that you want to obtain. How these certificates are obtained and implemented depends upon the browser that is being used.

When you choose a CA to issue client certificates, consider the following questions:

- Is the CA a well-known and trusted organization?
- What kind of information will the CA require from you to verify your identity? Most CAs require you to provide detailed identifying information, such as your name, address, organization name, and e-mail address.
- Does the CA have a Web site you can use to request your client certificate?
- Are the certificates issued by the CA compatible with your browser and server?

# Enabling Client Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can require users attempting to access your Web site to log on with a client certificate. Requiring a client certificate, however, does not reliably protect your content from unauthorized access. Any user with a valid and trusted client certificate can establish a connection and access your resource. To protect your Web content from unauthorized access you must do one of the following:

- Use Basic, Digest, or Integrated Windows authentication, in addition to requiring a client certificate.
- Create a Windows account mapping for client certificates. For more information, see Mapping Client Certificates to User Accounts in IIS 6.0.

When you attempt to set properties for a specific Web site, your Web server prompts you for permission to reset the properties of individual directories and files in the Web site. If you choose to reset these properties, your previous settings are replaced by the new settings. This is also true when you are setting properties for a directory containing subdirectories or files with previously set security properties.

# Client Certificate Mapping

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can authenticate users who log on with a client certificate by creating *mappings*, which relate the certificate information to a Windows user account. After you create and enable a certificate mapping, each time a user logs on with a client certificate, your Web server automatically associates that user with the appropriate Windows user account. This way, you can automatically authenticate users who log on with client certificates, without requiring the use of other supported authentication methods such as Basic, Digest, or Integrated Windows authentication. There are three ways to map client certificates: Directory Service (DS) mapping, one-to-one mapping, and many-to-one mapping.

**Directory Service Mapping**
*Directory Service (DS) certificate mapping* uses native Windows Active Directory features to authenticate users with client certificates. There are advantages and disadvantages to using DS mapping. For example, an advantage is that the client certificate information is shared across many servers. A disadvantage, however, is that wildcard matching is not as advanced as it is in the IIS mapper.

You can enable DS mapping only at the **Master** properties level, and only if you are a member of a Windows domain. Activating DS mapping will exclude the use of one-to-one and many-to-one mapping for the entire World Wide Web Publishing Service (WWW service).

**One-to-One Mapping**
*One-to-one mapping* maps individual client certificates to accounts. The server compares the copy of the client certificate it holds with the client certificate that is sent by the browser. The two must be absolutely identical for the mapping to proceed. If a client gets another certificate containing all of the same user information, it must be mapped again.

**Many-to-One Mapping**
*Many-to-one mapping* uses wildcard matching rules that verify whether a client certificate contains specific information, such as issuer or subject. This mapping does not compare the actual client certificate, but rather accepts all client certificates that fulfill the specific criteria. If a client gets another certificate containing all of the same user information, the existing mapping will work.

# Mapping Strategies

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Client certificate mapping is very flexible because any of the three mapping methods can be used to map client certificates to user accounts. You can map a client certificate to any number of user accounts. Likewise, you can map any number of client certificates to a single user account.

Certificate mapping can be used in several situations, including the following:

- Large Networks. Networks with a large number of client certificates can use many-to-one or Directory Service (DS) mapping. You can create one or more matching rules to map certificates to one or more Windows user accounts.

- Small Networks. Networks with very few users can use one-to-one mapping to provide greater control of certificate usage and revocation, or use many-to-one mapping to facilitate easier administration.

- Additional Security. For resources that have few users and that require additional security, you can use one-to-one mapping. In this way, you can be sure that only selected certificates are used. This allows more stringent certificate revocation policies to be enforced.

- Internet. Internet sites that use certificate authentication can use many-to-one mapping by accepting a wide range of certificates and mapping them all to an account with rights that are similar to the IUSR_*ComputerName* account.

- By certification authority. To map all users who log on with a client certificate that was issued by a particular organization, you can use many-to-one mapping. Then you can define a matching rule that automatically maps any certificate issued by that organization to a user account.

If you are using mapping to integrate your Web sites into a Windows domain, the Windows Directory Service mapper will best suit your purpose. For more information, see "Mapping certificates to user accounts" in Help and Support Center for Windows Server 2003.

# TCP/IP Port Filtering

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*TCP/IP port filtering* is the practice of selectively enabling or disabling Transmission Control Protocol (TCP) ports and User Datagram Protocol (UDP) ports on computers or network devices. When used in conjunction with other security practices, such as deploying firewall software at your Internet access point, applying port filters to intranet and Internet servers insulates those servers from many TCP/IP-based security attacks, including internal attacks by malicious users.

An Internet or intranet host, such as a computer or network device on a TCP/IP-based network, uses a combination of an IP address and port number to communicate with an application or service running on another Internet or intranet host. Together, an IP address and port number make up a *socket*. Because TCP/IP hosts are assigned a unique IP address, and standard TCP/IP-based applications and services typically use a specific TCP or UDP port number, sockets can direct communications between specific applications or services running on specific hosts.

A port number is identified in a TCP or UDP packet header and represents the transport protocol address of a specific application and service that uses TCP or UDP. For example, HTTP services use TCP port 80 by default, Telnet uses TCP port 23 by default, and Simple Network Management Protocol (SNMP) uses UDP port 161 by default.

The Internet Assigned Numbers Authority (IANA) categorizes TCP and UDP ports into three categories. Table 5.18 lists these categories.

**Table 5.18 IANA Categories of TCP and UDP Ports**

| Port Category | Port Number | Range Description |
|---|---|---|
| Well-known ports | 0–1023 | Typically used by standard system processes or programs that are executed by users with administrator credentials. Assigned by IANA. |
| Registered ports | 1024–49151 | Used by ordinary user processes or programs that are executed by ordinary users. IANA does not assign these ports, but registers use of them as a convenience for the TCP/IP community. |
| Dynamic or private ports | 49152–65535 | Unassigned and unregistered ports used for private applications, client-side processes, or other processes that dynamically allocate port numbers. |

Typically, the server side of a TCP or UDP process listens to the associated well-known port number. The client side of the process uses either the well-known port number or, more commonly, a dynamically allocated port number that is assigned only for the duration of the process.

To enable communications with the applications and services that your servers use, you must ensure that the associated ports are enabled. However, because malicious users on your internal network can attempt to exploit enabled ports to attack your servers, you should disable the TCP and UDP ports on your servers that are not used. This reduces the avenues of attack to your servers and improves the security of hosts that connect to your servers.

| **Important** |
|---|
| Server-based port filtering is not the only method you should use to secure your servers and network from TCP/IP-based security attacks. To provide a more complete network security solution, you should also deploy network firewall software at your Internet access point. |

Table 5.19 lists some of the default TCP port numbers for processes that are commonly used with Internet services.

**Table 5.19 TCP Ports and Associated Services**

| Default TCP Port Number | Internet Service |
|---|---|
| 20 | FTP Data Channel |
| 21 | FTP Control Channel |
| 23 | Telnet (enabled on some intranet or Internet servers) |

| | |
|---|---|
| 25 | Simple Mail Transfer Protocol (SMTP) |
| 80 | HTTP for World Wide Web |
| 119 | Network News Transfer Protocol (NNTP) |
| 443 | Hypertext Transfer Protocol over TLS/SSL (HTTPS) for secure World Wide Web |
| 563 | Network News Transfer Protocol over TLS/SSL (NNTPS) |

Notes

- Windows Media Services can provide streaming media services over unicast or multicast IP through a variety of static or dynamic UDP and TCP ports, or through a single port, depending on the configuration. You can also configure Windows Media Services to provide HTTP streaming media services through the default HTTP port, which is TCP port 80. For more information about Windows Media Services, see Help and Support Center for Windows Server 2003, and search for Using Windows Media Services.

- SSL 3.0 and Transport Layer Security (TLS) provide a way for clients and servers to exchange encrypted information. SSL 3.0 and TLS also provide a way for the server to verify the identity of the client before a user logs on to a server.

Table 5.20 lists the well-known UDP port numbers for the processes that are commonly used with Internet services.

**Table 5.20 UDP Ports and Associated Services**

| UDP Port Number | Service |
|---|---|
| 53 | DNS name queries (supports some Internet services) |
| 161 | SNMP |

For a list of TCP and UDP port numbers that are used by Windows Server 2003, see the services.txt file in the *systemroot*\System32\Drivers\Etc folder.

**Related Information**

For more information about network firewalls, see Microsoft Internet Security and Accelerations Server.

# Using Isolation to Secure Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Network Service built-in user account is the default worker process identity for application pools in IIS 6.0. You can have more than one application pool running with this process identity. However, if you use the Network Service account to configure ACLs on application pool resources, every application pool will then have access to every other application pool's resources.

**Configuring an Application Pool Identity to Isolate Applications**
To isolate applications effectively, each application pool should have its own account identity. That account identity must be added to the IIS_WPG local group so the account can launch the worker process associated with the application pool. However, you should not use the IIS_WPG group itself to configure ACLs on content files and directories. In addition, if you configure any accounts to be used either for anonymous access or for site authors or owners, you should not add them to the IIS_WPG group. Doing so would give those accounts access to directories that are shared among application pools, such as the compression cache and the ASP template cache. For more information about configuring an application pool identity, see Configuring Application Pool Identity in IIS 6.0.

To configure applications with a high level of isolation, create a unique user account for anonymous access, and then make this account the anonymous user account for your Web site. For step-by-step instructions, see Changing the Account Used for Anonymous Authentication. A unique anonymous user identity combined with a designated account for the application pool identity provides two of the essential elements needed for constructing an effective security boundary for the application. In addition, having a unique anonymous user identity allows you to configure authorization so that applications launched by the anonymous user can access only the required resources.

**Configuring Authorization to Isolate Applications**
Authorization uses the authenticated identity of the user, including the anonymous user, to limit access to resources. In this context, an application pool identity can be considered a user. You can use the application pool identity to configure ACLs to file system resources, the IIS metabase, and registry keys.

Use an ACL to permit or deny specific users access to an application's directories and files. When you configure ACLs on content, use groups to contain application pool identities and assign ACLs to those groups. Your groups can represent various user functions or actions, such as authoring or anonymous access, or specific permissions such as Browse or Read. Then, assign NTFS permissions to file system resources to allow permissions where the application pool identity requires access, for example, logging directories or script caches.

Use the Metabase Explorer tool in the IIS 6.0 Resource Kit Tools or the Meta-Data Edit Tool to view and set ACLs on metabase keys. Default ACLs on application pools give the IIS_WPG, Network Service, and Local System accounts the ability to query metabase properties for all application pools. You can increase isolation by:

- Removing metabase entries for IIS_WPG.
- Assigning the application pool identity Read permission.
- Assigning any Web site author or administrators Read permission.
- Giving the Administrators group and any other system administrators Full Control.

To configure ACLs on registry keys, you can use templates from the Windows Server 2003 Security Guide to restrict permissions on the registry. In addition, you can specifically restrict access to portions of the registry that contains information about Component Object Model (COM) objects used by your applications.

For more information about isolating applications, see Configuring Application Isolation using Windows Server 2003 and IIS 6.0.

# Secure Code

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

No matter how strong the security of your server is, if the applications that it hosts are not programmed according to best security practices, your network might be vulnerable to attacks. As part of a defense-in-depth strategy, IIS administrators should work with developers to ensure that the code hosted on the server running IIS is as secure as possible. For example, developers can reduce the risk of certain types of attacks, such as cross-site scripting and SQL injection, by validating user input.

*Cross-site scripting* occurs when an attacker sends a link in e-mail to a user or otherwise points the user to a Web site, and the link actually contains malicious script code, which can be VBScript or JScript. As a result of cross-site scripting, an attacker potentially can perform such actions as reading cookies, instantiating browser plug-ins or native code and scripting it with untrusted data, and intercepting user input. Potential targets are any browsers that support scripting and any Web servers that support HTML forms.

In a *SQL-injection attack*, an attacker finds a way to modify a SQL command that is used in a dynamic Web page. The attacker typically crafts user input in a way that modifies the logic of the SQL command, so that the attacker is able to obtain information that would otherwise be unavailable.

Developers who write the applications that are hosted on your Web site must know how to include safeguards against these types of attacks. It is paramount that all user input be validated thoroughly. The book *Writing Secure Code* covers these and other Web application security issues in depth. For more information about *Writing Secure Code*, see Additional Resources for IIS 6.0 Security.

# MIME Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To reduce the available attack surface, IIS serves only a small number of default static file types unless you add an approved extension to the MIME types list in IIS Manager. MIME types instruct a browser or mail application about how to handle the files received from a server. For example, when a browser requests an item on a server, the browser also requests the MIME type of the object. When IIS delivers a Web page to a client browser, IIS also sends the MIME type of the data it is sending. If there is an attached or embedded file in a specific format, IIS also tells the client application the MIME type of the embedded or attached file. The client application then knows how to process or display the data being received from IIS.

A MIME type consists of two required parts (the media-type and sub-type) and one optional part (the parameter-list) in the following format:

```
media-type/sub-type;parameter-list
```

| **Important** |
| --- |
| IIS does not support the optional parameter-list. Instead, IIS attempts to match the entire string that follows the media-type/. |

Known extensions are handled somewhat differently for static pages and for dynamic content, but you can add, change, or remove allowable extensions for both static and dynamic content. For more information about enabling dynamic content, see Enabling Dynamic Content.

If a user requests a file with an extension that is not permitted, the user receives a 404 error (File not found), rather than a 403 error (Access is denied). This reduces the amount of information that is supplied to a would-be attacker. For more information about interpreting substatus error codes, see Analyzing Log Files.

**Inheritance and Merge Behavior**

IIS merges the MIME type mappings that it finds on different levels into one in-memory representation. For example, when IIS receives a request, it searches for the **MimeMap** property for the level, such as lm/w3svc/1/root/MimeMap, and the closest inherited level. If the MIME type for an extension is not in the resulting array, IIS looks for the MIME type in the registry (HKEY_CLASSES_ROOT*extension*). If no matching MIME type is found, IIS can use a wildcard mapping if one is configured. Otherwise, the user receives a 404 error. The default wildcard for IIS is .*, application/octet-stream.

| **Important** |
| --- |
| Using a wildcard mapping can pose a security risk, and is recommended for troubleshooting only. By default, wildcard mapping is not configured in IIS, and only extensions that are defined at the **MimeMap** node in the IIS metabase are served. |

If you use IIS Manager to remove a MIME type mapping, you must also verify that the MIME type does not exist in the registry. If the MIME type continues to exist in the registry, files with that extension continue to be served. Be aware, however, that removing MIME type mappings from the registry may affect other applications, because the registry database is shared by multiple applications.

**Merge considerations**

When you configure MIME types, be aware of the restrictive MIME merge behavior of IIS 6.0. Values that are set on nodes are inherited by default by child nodes, but values that are set on child nodes do not propagate back up to parent nodes. For example, if you allow the .ini MIME type on *LocalComputer*/Directory, but you do not allow this MIME type on the root directory, users might not receive the files that you expect. In this situation, if a user requests http://server/directory/default.ini, the user receives the file. However, if the user requests http://server/default.ini, the user does not receive the file because IIS uses the **MimeMap** property for *LocalComputer*/Directory and does not merge it with entries in *LocalComputer*/Root. Instead, the user receives a 404 error.

# Auditing Server Activity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use security auditing techniques to monitor a broad range of user and Web server security activity. You should routinely audit your server configuration to detect areas where resources might be susceptible to unauthorized access and tampering. For example, you can audit the following user activities:

- Logon attempts, both successful and unsuccessful
- Attempts to access restricted accounts
- Attempts to execute restricted commands

You can use the integrated Windows utilities, or the logging features built into IIS, or use Active Server Pages (ASP) applications to create your own auditing logs.

# Standards Implemented

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Many of the security features in IIS implement Internet community standards. These standards help to facilitate uniformity and cross-platform utilization of applications and information. Microsoft is committed to working with the Internet and computer communities both in assisting to structure good standards, but also in its implementation of those standards. The following standards are implemented by IIS security features:

- Secure Sockets Layer (SSL) 3.0. SSL 3.0 is a public key-based security protocol implemented by the Secure Channel (Schannel) security provider. SSL security protocols are used widely by Internet browsers and servers for authentication, message integrity, and confidentiality.

- Basic authentication. This authentication method is a part of the HTTP 1.1 protocol, as defined in RFC 2617 HTTP Authentication: Basic and Digest Access Authentication, and sends passwords over networks in Base64 encoded format. Most browsers support this specification.

- Digest authentication. This authentication method uses the MD5 message-digest algorithm as defined in RFC 1321, The MD5 Message-Digest Algorithm, and sends authentication information over networks as a *hash* and is compatible with proxy servers.

- Public-Key Cryptography Standard (PKCS) #7. This security standard describes the format of encrypted data such as digital signatures or digital envelopes that contain information. Both of these are involved in the certificate features of IIS.

- PKCS #10. This security standard describes the format of requests for certificates that are submitted to certification authorities.

# Additional Resources for IIS 6.0 Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to Managing a Secure IIS 6.0 Solution.

**Related Information**

- Analyzing Log Files for information about interpreting substatus error codes.
- Configuring Internet Sites and Services for information about configuring host header–based sites.
- Optimizing IIS 6.0 Performance for information about managing security while maintaining server performance.
- Working with the Metabase for information about metabase security.
- Microsoft Knowledge Base article 301432, How to Create a Subweb and Add Permissions for information about creating subwebs and assigning unique permissions.
- The Microsoft Internet Security & Acceleration Server Web site for information about network firewalls.
- The Microsoft .NET Passport Web site for information about .NET Passport authorization.
- MetaACL.exe Sample for information about MetaACL.vbs.
- Cryptography Essentials for information about CSP and managing installed third-party cryptographic providers.
- The RFC-Editor Web page for information about RFC 1321 and RFC 2617.
- *Writing Secure Code, 2nd Edition,* by Michael Howard and David LeBlanc, 2003, Redmond: Microsoft Press.

**Related IIS 6.0 Operations Guide Topics**

- The AspRunOnEndAnonymously Metabase Property for information about ASP OnEnd routines.
- Basic Authentication in IIS 6.0 for information about the Basic authentication method.
- Digest Authentication in IIS 6.0 for information about the Digest authentication method.
- IIS 6.0 Encryption for information about SSL encryption features.
- Global Registry Entries for information about the UserTokenTTL property.
- Securing Sites with Web Site Permissions for information about Web site permissions.

**Related Windows Server 2003 Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the**Search in title only** check box.

- Add a member to a local group for information about adding users to groups.
- Authorization Manager for information about Authorization Manager.
- Cacls for information about the **cacls** command.

- Domain and forest functionality for information about domain functional levels.
- Managing IPSec from the command line for information about Netsh commands for IPSec.
- Using Windows Automatic Updates for information about scheduled installation of security updates.

**Related Tools**

- The Xcacls.exe command-line tool

  See the Xcacls.exe Web site for information about Xcacls.exe, including instructions for downloading this command-line tool.

# Optimizing IIS 6.0 Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you monitor your application servers to maintain a baseline of performance data, you can spot performance trends as they develop, take steps to prevent unsatisfactory performance, decide how to best tune or upgrade your servers, and determine whether your changes are beneficial. By tuning your Internet Information Services (IIS) version 6.0 application servers, you improve the client experience, help avoid bottlenecks, and can extend the interval between hardware upgrades.

**In This Section**


Overview of Performance Monitoring and Tuning

Controlling Memory Usage

Preventing Processor Bottlenecks

Improving Application Performance

Balancing Performance and Security

Optimizing Performance Through Design

Additional Resources for IIS 6.0 Performance

# Overview of Performance Monitoring and Tuning

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Performance monitoring* is the process of capturing and analyzing performance data from different areas of your server environment, which include applications, memory, processors, hardware, and your network. You obtain performance data to help you recognize trends as they develop, prevent unsatisfactory performance, and optimize the use of your system resources. Monitoring also helps you decide when to upgrade your hardware and whether upgrades are actually improving your servers performance.

Although some performance problems and their solutions are immediate and obvious, others develop over time and require careful monitoring and tuning. First, monitor to establish a performance baseline against which to judge and compare the performance of your server; without a baseline, your tuning efforts might not give you optimal performance.

By monitoring performance and analyzing performance data, you can identify performance patterns to help you locate bottlenecks and to identify underused or overused resources. After locating a bottleneck, you can make changes to the component to improve performance. Bottlenecks can occur anywhere in your server environment at any time, so you must regularly monitor performance to capture baseline information about your system.

To get started with performance monitoring, familiarize yourself with the tools used, which include System Monitor, Performance Logs and Alerts, and Network Monitor; with the counters that are available for monitoring performance objects; and with the basics of setting up monitoring in order to collect useful data.

# Using Performance Tools to Obtain a Baseline

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Using the appropriate tools, you can monitor your application server to collect data on performance, make specific changes in order to tune your server and its Web applications, evaluate the results of changes, and plan additional changes to help your application server run optimally. Which tools you should use, and when you should use them, depends on the data you need and your purpose in collecting it.

The Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; or Windows® Server 2003, Web Edition operating system provides the Performance console, which includes System Monitor, Performance Logs and Alerts, and Task Manager. System Monitor and Performance Logs and Alerts are Microsoft Management Console (MMC) snap-ins. In addition, you can access these tools through the **Run** dialog box by using the **perfmon** command.

From the Performance console, you can examine the output of *performance counters*, which monitor the activity of *performance objects*. Each performance counter is named based on the object from which it collects data (for example, a processor, process, or thread) and the type of data that it collects (for example, the Processor\% Processor Time performance counter reports the average percentage of processor time in use for the Processor performance object).

Windows Server 2003 provides several command-line tools that you can use to monitor performance. This section cites two command-line tools for performing event tracing. For more information about using command-line tools to monitor performance, see "Monitoring performance from the command line" in Help and Support Center for Windows Server 2003.

# System Monitor

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

System Monitor enables you to collect and display real-time performance data for a local computer or remote computers according to criteria that you define. System Monitor can also display data that is collected in counter logs.

You can use System Monitor to monitor your server's activity and summarize its performance at selected intervals. You can display performance data in real-time charts or reports, collect data in files, and generate alerts when critical events occur. Use the data to determine the cause of system bottlenecks and to fine-tune system and application performance.

By default, the System Monitor graph is plotted on a scale of 0 through 100. For counters that register small values, you might need to change the scale. When interpreting the counters, remember that most of them show the most recently observed value, not a long-term average. Use Performance Logs and Alerts to log data over time and determine average values.

**Related Information**

For more information about using System Monitor, see Setting Up Monitoring, and "System Monitor overview" in Help and Support Center for Windows Server 2003.

# Performance Logs and Alerts_IIS_SP1_TechRef

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

Use Performance Logs and Alerts to automatically collect performance data in logs and to send alerts based on criteria that you set:

- Create custom log files that automatically log the values for specified counters based on the start and stop times that you provide.
- Set an alert on a counter that defines what action you want the system to take (for example, send an alert, run a program, make an entry in a log file, or initiate logging) when the alert condition is met.
- View counter data during or after collection. After collecting data in logs, use System Monitor to view the data, or export the data to a spreadsheet or database application.

A binary log file format is provided for circular logging or for logging instances, such as threads or processes, that begin after the system starts logging data. (*Circular logging* is the process of continuously logging data to a single file, overwriting previous data with new data.)

| **Note** |
| --- |
| You can also monitor your server by examining IIS logs. For information about IIS logs, see Analyzing Log Files. |

# Task Manager

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use Task Manager to monitor key indicators of your computers performance.

The **Performance** tab displays a dynamic overview of your computers performance, including usage history for CPU and memory; the total number of handles, threads, and processes currently running on the computer; and physical memory, kernel memory, and commit memory usage.

The **Processes** tab shows information about the processes running on your computer. You can select columns to display. For example, display information about CPU and memory usage, page faults, handle count, and a number of other parameters.

# Event Tracing with Log Manager and Trace Report

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A Web request goes through many layers of the operating system, most of which can be tuned. By tracing IIS and kernel events, you can pinpoint Web bottlenecks on the server, and can often determine where to tune your server for improved performance. You can use the trace data with performance monitors, capacity planning tools, and applications that analyze information relating to system resource usage.

Use the Log Manager and Trace Report command-line tools in Windows Server 2003 to collect and report trace data relating to IIS and kernel events:

- Use the Log Manager tool (Logman.exe) to collect trace data for specified IIS and kernel events.
- Then use the Trace Report tool (Tracerpt.exe) to process the event trace log and to generate trace analysis reports and comma-delimited files for the events.

For information about Logman.exe and Tracerpt.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click **Command-line reference A-Z**. For information about using Log Manager and Trace Report with IIS 6.0, including syntax examples and analysis of output reports, see Troubleshooting IIS 6.0 with Tracing.

# Network Monitor_IIS_SP1_TechRef

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

Network Monitor is a useful tool for conducting a network analysis of your Web applications. Use this tool to document the number of roundtrips that are required for a client to download a specific Web page from your server. Make adjustments to your applications, such as the changes recommended in Optimizing Application Design and Administration, and then rerun Network Monitor to verify that the adjustments improved performance.

Network Monitor collects data about the network traffic that flows to and from the network adapter of the computer on which it is installed. By capturing and analyzing that data, you not only can improve performance, you can prevent, diagnose, and solve many types of network problems.

You can configure Network Monitor to provide only the information that is most relevant to you. For example, you can configure *capture triggers*, which cause Network Monitor to start or stop capturing data when a circumstance or set of circumstances occurs, and you can configure *capture filters* to control the types of data that Network Monitor captures or displays.

Network Monitor monitors the frames in which data is sent over the network, which, in addition to the data, contain the addresses of the source and destination computers and the protocols in use.

Because using Network Monitor can affect server performance, you might want to run this tool when usage is low or for short periods during peak usage. Use capture filters so that you collect only the data that you need, then repeat each network capture one or more times to verify that the data you obtained accurately reflects network activity.

# Monitoring with Performance Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Both System Monitor and Performance Logs and Alerts allow you to examine the output of *performance counters*, which monitor the activity of specific *performance objects*. Each performance counter is named for the object for which it collects data and the type of data that it collects. For example, the Processor\% ProcessorTime counter is the primary indicator of processor activity, displaying the average percentage of busy time observed during the sample interval.

A performance object often offers a choice of more than one *instance*. Using the Processor object as an example, if you are monitoring a dual-processor computer, the first processor instance is 0 and the second processor instance is 1. A commonly used instance is the _Total instance, which monitors the sum of the values of a specific counter's instances.

**Components of a counter path**
The counter path contains the computer name, and the performance object, instance, and counter.It is typically represented in the following way:

```
\\ComputerName\Object(Instance)\Counter
```

The computer name is optional. If you do not designate a computer name, the counter monitors the local computer.

Again, using the Processor object as an example, if you want to monitor the percentage of processor time that your #0 processor uses, your selection is represented in the following way:

```
Processor (#0)\% ProcessorTime
```

**Sampling methods for counters**
Performance counters collect and display their data in one of two ways:

- Instantaneous counters. These counters display the most recent measurement of an aspect of resource use. For example, the Process\Thread Count counter shows the number of threads for a particular process as of the last time it was measured. An instantaneous counter might have a name containing the word Current. Unless the server has a steady workload, these counters might not provide meaningful data.

- Averaging counters. Based on the previous two measurements, averaging counters derive an average value for the interval between the measurements. These counters typically calculate a percentage or the number of occurrences per second. For example, the Memory\Pages/sec counter shows the number of memory pages read over the sample interval, divided by the number of seconds in the interval. When you start one of these counters, no value is displayed until the second measurement is taken.

# Counters Provided by Windows and by IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Windows Server 2003 provides hundreds of counters. Frequently used performance objects include cache, memory, process, and server. For more information about Windows performance objects and counters, see "Performance objects and counters" in Help and Support Center for Windows Server 2003.

When you install IIS, you automatically install the IIS performance counters, which include the following:

- Web Service counters to monitor the World Wide Web Publishing Service (WWW service).
- Web Service Cache counters to monitor the WWW service cache.
- FTP Service counters to monitor the File Transfer Protocol (FTP) service.
- Active Server Pages counters to monitor applications that run as Active Server Pages (ASP pages).

If you use Simple Network Management Protocol (SNMP) services to monitor your Web services, IIS makes available a set of counters by means of SNMP. They include SNMP FTP service counters and SNMP HTTP service counters, which you can use to monitor an SNMP WWW service. However, you cannot view the SNMP counters in System Monitor or Performance Logs and Alerts; you must use Windows Management Instrumentation (WMI) or a management information base (MIB) browser tool instead. For more information about specific IIS counters, including how to locate and view the SNMP counters, see IIS 6.0 Performance Counters.

# Suggested Performance Counters to Watch

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table 6.1 and Table 6.2 list some of the counters that are typically used for monitoring application servers. The tables include the preferred value of each counter and suggest values that might indicate performance problems.

Most of the values provided in this section are relative: Optimal values vary with the Web application, system, and network architecture. For example, data collected by the PhysicalDisk\Avg Disk Bytes/Transfer counter is different for different kinds of controllers and drives. Some counters may not be relevant to your applications. For example, data collected by the Active Server Pages\Transactions/sec counter is not relevant to a Web site with only static content.

In Table 6.1 and Table 6.2, *ComputerName* is omitted to conserve space. For example, the first counter in Table 6.1 is actually *ComputerName*\Memory\Pages/sec.

**Table 6.1 Preferred Values for Frequently Used Performance Counters**

| Object\Counter | Preferred or Ideal Value |
|---|---|
| Memory\Pages/sec | 0–20. (Unhealthy if greater than 80; probably indicates not enough RAM.) |
| Memory\Available Bytes | 10% of physical memory. |
| Memory\Committed Bytes | No more than 75 percent of physical memory. |
| Memory\Pool Nonpaged Bytes | A steady value. (A slow rise might indicate a memory leak.) |
| Processor\% Processor Time | Less than 75 percent. |
| Processor\Interrupts/sec | Depends on the processor, and on network hardware and drivers. Up to 3,500 for a 90–megahertz (MHz) Pentium; more than 19,000 for a 500–MHz Pentium or more than 58,000 for a 1.5–gigahertz (GHz) Pentium. Lower is better. If the value is too high, try moving some hardware devices to a different server. |
| Processor\System Processor Queue Length | 4 or less. |
| LogicalDisk\% Disk Time<br><br>PhysicalDisk\% Disk Time | As low as possible. |
| LogicalDisk\Avg. Disk Queue Length<br><br>PhysicalDisk\Avg. Disk Queue Length | Less than 4. |
| LogicalDisk\Avg. Disk Bytes/Transfer<br><br>PhysicalDisk\Avg. Disk Bytes/Transfer | As high as possible. |
| System\Context Switches/sec | Compare this value with the value of Web Service\Total Method Requests/sec. Context switches per request (Context Switches/sec divided by Total Method Requests/sec) should be low. |
| System\System Calls/sec | As low as possible. |

| | |
|---|---|
| Web Service\Bytes Total/sec | As high as possible. |
| Web Service\Total Method Requests/sec | As high as possible. |
| Web Service\Current Connections | As high as possible. |
| Web Service Cache\File Cache Hits % | As high as possible for static content.<br><br>**Note**: This value might be low if the Kernel URI cache hits percentage is high. |
| Web Service Cache\Kernel:URI Cache Flushes | As low as possible, relative to the number of requests.<br><br>**Note**: This number increases every time a file is flushed from the HTTP.sys response cache (HTTP.sys is the kernel-mode device driver in IIS 6.0), which means that the content has not been accessed in the past 2–4 minutes. The only way to decrease this number is to flush the cache less often, although frequent flushing can cause HTTP.sys to use more memory for content that is not being accessed. |
| Web Service Cache\Kernel:URI Cache Misses | As low as possible. (Each request for dynamic content increases the value of the counter by one.) |
| Web Service Cache\Kernel:URI Cache Hits % | As high as possible. (Applies to static unauthenticated content and dynamic content that is marked as cacheable.) |
| Active Server Pages\Request Wait Time | As low as possible. |
| Active Server Pages\Requests Queued | As low as possible. |
| Active Server Pages\Transactions/sec | As high as possible.<br><br>**Note**: ASP transactions degrade overall server performance because each transaction requires interaction with a database. If you are concerned about server performance, use ASP transactions sparingly. |

The counters listed in Table 6.2 are useful for monitoring your FTP servers.

**Table 6.2 Preferred Values for Useful FTP Performance Counters**

| Object\Counter | Preferred or Ideal Value |
|---|---|
| FTP Service\Bytes Sent/sec | As high as possible. |
| FTP Service\Bytes Received/sec | As high as possible. |
| FTP Service\Bytes Total/sec | As high as possible. |

In addition to the counters in Table 6.1 and Table 6.2, the following system counters are useful for monitoring your servers use of system resources:

- System\Threads
- System\Processes
- System\Context Switches/sec

For a discussion of system-related counters and how their data relates to the performance of your Web server, see Identifying Processor Bottlenecks.

# Setting Up Monitoring

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use the counters listed in Table 6.1 and Table 6.2 to generate a baseline log that typifies your server's performance under ordinary conditions. Be sure to obtain multiple logs so that your baseline data contains enough information to give you a good sense of your server's behavior. For example, you might log the recommended counters for one week of typical operation. Shorter logs, which you need to periodically create, can help you monitor how changes in network activity are affecting your system performance.

For information about selecting performance objects and counters to view in System Monitor, see Viewing Counter Data in the Performance Console.

# Using the Predefined System Overview Log

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To get started quickly, use the predefined System Overview log that Windows Server 2003 provides under Counter Logs. The System Overview log is configured to create a binary log in the *systemroot*\Perflogs folder that, after manual startup, is updated every 15 seconds until the log reaches a maximum size. To minimize the system resources used during logging, the System Overview log provides *circular logging*, which overwrites older data with new data.

By default, the log includes three counters: Memory\Pages/sec, PhysicalDisk(_Total)\Avg. Disk Queue Length, and Processor(_Total)\% Processor Time.

Interpreting the data from a binary log requires the use of a parser, such as the Log Parser tool (LogParser.exe), to extract the data and convert it to formatted text. You can obtain the Log Parser tool by downloading the IIS 6.0 Resource Kit Tools from the Microsoft Download Center.

For information about creating counter logs, including baseline IIS logs, see Microsoft Knowledge Base article 313064, Monitor Web Server Performance by Using Counter Logs in System Monitor in IIS.

For information about using Performance Logs and Alerts to create custom logs, see the "Performance Logs and Alerts overview" in Help and Support Center for Windows Server 2003. For information about analyzing log data and a list of acceptable values for counters, see "Analyzing performance data" in Help and Support Center for Windows Server 2003.

# Troubleshooting System Monitor when IIS Host Thousands of Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you add IIS Web service counters to a Performance Monitor log, either by right-clicking the System Monitor details pane and clicking **Add Counters** or by clicking the plus sign (+) on the tool bar, Performance Monitor may take a long time to enumerate the list of available instances on the server. If you are hosting a very large number of Web sites, for example more than 20,000, Performance Monitor may appear to stop responding (hang).

To work around this behavior, use Typeperf.exe to display all counters currently available on a particular local or remote computer and then use Logman.exe to create, start, stop, and schedule performance counter trace logs and collections on local and remote systems. For more information about these command line tools, see Help and Support Center.

# Collecting Useful Data

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you set up performance monitoring, keep the following guidelines in mind:

- Be selective when choosing which counters to run, because several counters running at the same time can cause a small decrease in performance and can consume disk space.

- For routine monitoring, start by sampling data every 15 minutes.

- If you are monitoring for a specific problem, try varying the length of the interval until you find the optimal setting for obtaining the data that you need.

  For example, if you are monitoring the activity of a specific process at a specific time, set a frequent update interval; however, if you are monitoring a problem that reveals itself slowly, such as a memory leak, use a longer interval.

- When deciding the logging interval, consider the overall length of time over which you want to monitor.

  For example, you might sample data every 15 seconds for monitoring that is 4 hours or less. If you plan to monitor for 8 hours or more, set an interval longer than 300 seconds.

When you decide to change your system in any way, such as by tuning settings, adding hardware, or upgrading your system, follow these guidelines:

- Always make changes one at a time. Unsatisfactory performance that appears to relate to a single component might be the result of bottlenecks involving multiple components. However, if you make multiple changes simultaneously, it is often impossible to assess the impact of each change. For this reason, even when you are troubleshooting performance problems that involve multiple components, you might need to address each problem individually, making changes one at a time.

- Repeat monitoring after every change. Because tuning one resource can affect other resources, it is recommended that you monitor after every change that you make, being sure to monitor not only the area you changed but also more globally. Unless you monitor after you make a change, you cannot determine whether your change improved performance.

- Review event logs. Sometimes poor or unsatisfactory performance generates output in event logs. Use Event Viewer to monitor events that appear in event logs, such as application, system, and security logs. In addition, IIS logs can provide you with information about your changes. For more information about Event Viewer, see Event Viewer in Help and Support Center for Windows Server 2003. For more information about IIS logs, see Analyzing Log Files.

- Test your changes in a lab. When your changes are significant, test them in a performance lab before making the changes in a production environment.

- Periodically, obtain a new performance baseline. After making several changes, monitor again to update your performance baseline. Compare the new baseline to previous logs to observe the actual effects your changes are having on performance and capacity.

# Managing Network Activity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Creating and maintaining a Web site involves using hardware, software, and network bandwidth to manage network traffic. Servers send out pages in response to requests. In order to issue a request, a browser first establishes a Transmission Control Protocol (TCP) connection with the server and then sends the request through the connection. Network traffic, as the term applies to Web servers, is the mixture of incoming requests and outgoing responses.

Network traffic typically occurs in bursts and clumps that are only partly predictable. For example, many intranet sites have activity peaks at the beginning and end of the day, and around lunchtime. However, the exact size of these peaks varies from day to day, and the actual traffic load changes from moment to moment. There is a direct relationship between the amount of traffic and the network bandwidth needed. *Network bandwidth* is the capacity of the transmission medium stated in bits per second (bps). On computer networks, greater bandwidth means faster data-transfer capability. The more visitors your site has and the larger the pages the site provides, the more network bandwidth your server requires.

# Estimating Bandwidth Requirements and Connection Speed

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Data traveling on a network is split into packets. In addition to the data that it carries, each packet includes about 20 bytes of header information and other network protocol information (this extra information is known as *overhead*). The amount of data in a packet is not fixed, and thus the ratio of overhead to data can vary. Most incoming HTTP requests are small. A typical request (for example, `GET http://www.microsoft.com/default.asp`), including the TCP/IP overhead, consists of no more than a few hundred bytes as it travels across the network.

Overhead can become an important consideration when you are estimating your sites bandwidth requirements and deciding how fast a connection you need. If current usage is close to the limits of your connections capacity, an extra 20 percent for overhead might mean that you need the next faster type of connection.

Think of a server that displays static Hypertext Markup Language (HTML) text-only pages that average 5 kilobytes (KB), which is nearly equivalent to a full page of printed text. The server is connected to the Internet through a DS1/T1 line, which can transmit data at 1.536 megabits per second. Inherent overhead makes it impossible to use the full-rated T1 bandwidth of 1.544 megabits per second. For this 5-KB file, protocol overhead is significant, amounting to about 30 percent of the files size. For larger files, the overhead accounts for a smaller percentage of network traffic.

Table 6.3 shows the traffic generated by a typical request for a 5-KB page. Note that all the figures for overhead are estimates. The precise number of bytes sent varies with each request.

**Table 6.3 Network Traffic Generated by a Request for a 5-KB Page**

| Traffic Type | Bytes Sent |
|---|---|
| TCP connection | 180 (approx.) |
| GET request | 256 (approx.) |
| 5-KB file | 5,120 |
| Protocol overhead | 1,364 (approx.) |
| **Total** | **6,920** |

To calculate the number of bits for this 5-KB page, multiply the total number of bytes sent by 8 bits per byte:

6,920 8 = 55,360

As stated previously, a T1 line can transmit 1.536 megabits per second. Dividing bits per second by bits per page (1,536,000/55,360) indicates a maximum transmission rate of just under 28 pages per second. (Because modems add a start bit and a stop bit to each byte, the transmission rate through a modem is slower than the raw numbers appear to indicate.) For this small text-only page, Table 6.4 shows the relative speed for connections to several types of network interfaces and an estimated page delivery rate for each connection type.

**Table 6.4 Relative Network Interface Speeds to Request a 5-KB Page**

| Connection Type | Connection Speed | 5-KB Pages Sent per Second |
|---|---|---|
| Dedicated PPP/SLIP via modem | 28.8 kilobits per second (Kbps) | About half of 1 page |
| Frame Relay or fast modem | 56 Kbps | Almost 1 page |
| Integrated Services Digital Network (ISDN) | 128 Kbps | Just over 2 pages |
| Typical DSL | 640 Kbps | Almost 11 pages |
| DS1/T1 | 1.536 megabits per second (Mbps) | 26 pages |

| | | |
|---|---|---|
| 10-megabit Ethernet | 8 Mbps (best case) | (Up to) 136 pages |
| DS3/T3 | 44.736 Mbps | 760 pages |
| OC1 | 51.844 Mbps | 880 pages |
| 100-megabit Ethernet | 80 Mbps (best case) | (Up to) 1,360 pages |
| OC3 | 155.532 Mbps | 2,650 pages |
| OC12 | 622.128 Mbps | 10,580 pages |
| 1-gigabit/sec Ethernet | 800 Mbps (best case) | (Up to) 13,600 pages |

If you add a small graphic to the 5-KB page, the results are considerably different. An image, in the form of a .jpg file that appears on-screen as a 1-inch by 2-inch rectangle (the actual physical size depends on monitor settings), takes up about as much disk space as the original text file. Adding one such image file to each page nearly doubles the average page size and also the number of requests that are needed. This increased size reduces the number of pages that the server can send to the Internet on a DS1/T1 line to a maximum of about 15 pages per second, regardless of how fast the computer itself runs.

If a page contains several images, if the images are relatively large, or if the page contains multimedia content, the page can take much longer to download. You can often find simple solutions for improving download time, such as consolidating, removing, sizing down, or compressing images or converting images to a more compact graphics format. You might also connect to the network by using a faster interface (this change can improve performance for the server, but not necessarily for the client, as is discussed later in this section).

A site that serves primarily static HTML pages, especially pages with a simple structure, is likely to run out of network bandwidth before it runs out of processing power. In contrast, a site that performs a lot of dynamic page generation, or that acts as a transaction or database server, uses more processor cycles and can create bottlenecks in its processor, memory, disk, or network.

For more information about how to create more efficient Web pages, see Creating a More Efficient Web Site.

**Browser Download Time**

Up to this point, the bandwidth discussion has focused on the number of pages that a *server* can send in a specified amount of time. The second part of this discussion focuses on how long it takes a *browser* to download a page.

Consider how much time a browser needs to download a page that, including overhead, amounts to 90 KB, which equals about 720 kilobits. If you ignore *network latency*, which is the amount of time required for a packet or signal to travel from one point on a network to another (latency typically adds a few seconds before any of the data arrives), and if everything is working perfectly, a 720-kilobit page downloads over a 28.8-Kbps connection in about 25 seconds. The download takes longer if any blocking or bottlenecks occur at the server, if the network is overloaded and slow, or if the users connection is slower than 28.8 Kbps (due to poor line quality, for example).

If the client computer has a higher bandwidth connection on an intranet, the download time is usually much shorter. If your Web site is on the Internet, however, you cannot count on a majority of users having faster connections.

**Server-Side Request Processing**

It takes about 52 connections at 28.8 Kbps to saturate a DS1/T1 line. Not accounting for delays (which are fairly typical), if no more than 52 clients simultaneously request the page used in the preceding example, and if the server can keep up with the requests, the clients all receive the 90-KB page in the 25 seconds calculated in the example.

If 100 clients simultaneously request that same page, however, the total number of bits to be transferred is 100 737,280 (720 kilobits). It takes between 47 and 48 seconds for that many bits to travel over a DS1/T1 line. When that many bits are involved, the network connection for the server, not the client, becomes the limiting factor.

A DS3/T3 line carries nearly 45 Mbps, about 30 times the capacity of a DS1/T1 line, and it takes more than 1,500 clients at 28.8 Kbps to saturate its bandwidth. Moreover, the increase in download time for each new client is much smaller on a DS3/T3 line. When there are 2,000 simultaneous 28.8-Kbps connections, for example, it still takes less than 33 seconds for a client to download the page.

This example assumes that the server is capable of performing the relevant processing and handling of 2,000 simultaneous connections, which is not the same as handling 2,000 simultaneous users: users occasionally stop to read or think, and typically

spend only a small percentage of their time downloading, except while receiving streaming multimedia content. Because of this difference between users and connections, IIS 6.0 can support more users than the figures appear to indicate. A Web server on a DS1/T1 line can typically handle several hundred users connecting at 28.8 Kbps, and with a DS3/T3 line the number typically climbs to 5,000 or more. Although these numbers are derived from actual servers, you can expect performance to vary with differing content types and user needs, and with the number and type of services that a particular computer is performing.

These network performance differences scale linearly, and the scaling continues at higher data-transfer rates. Two DS3/T3 lines, for example, can serve approximately twice as many clients as one line can, provided that processor power is sufficient to keep up with user demand and no bottlenecks prevent your servers from operating at maximum processing power. Connecting to the network by using a faster interface often resolves any performance problems for the server, but not necessarily for the client, as is shown later in this section.

**Perceived Delay, Acceptable Delay, and Site Performance**

The amount of time in which a user perceives that a Web page appears, known as *perceived delay*, is not identical to the actual time that is required to fully display the page. If the first thing that the user sees when opening a page is a set of buttons that allow further navigation, the user might never know or care that the rest of the page takes longer than a minute to download. If, on the other hand, the page takes longer than a minute to download, and the navigation buttons do not appear until after the rest of the page, users might not wait for the download.

The length of an *acceptable delay* depends partly on the kind of information that the page provides, but no more than 30 seconds is usually considered acceptable. If the information is important, users are more likely, but still reluctant, to wait.

As a best practice, set minimum performance goals that specify the acceptable performance for each of your sites. For example, the managers of certain Microsoft e-commerce Internet sites test their site performance against the following minimum performance goals:

- All pages must load within 10 seconds.
- Average CPU usage for each server must be less than 70 percent.
- Available memory must remain stable throughout the test period.
- The site must maintain 10 customer checkouts per second.

# Monitoring Network Activity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The primary functions of IIS 6.0 are to establish connections for clients, to receive and interpret requests, and to deliver files — all as quickly as possible. The pace at which these vital functions are performed depends, in large part, on two factors: the effective bandwidth of the link between the server and the network, and the capacity of this link and the server to support network resources.

The speed of the network interfaces also affects the pace. Some servers have two or more network interfaces, which are frequently called front-end and back-end servers. *Front-end servers* are client-accessible Web servers that run application server software, such as IIS, to handle traffic coming from the Internet. Front-end servers can add a layer of protection for your *back-end servers*, which include database servers, file servers, domain controllers, and WINS servers. Different interfaces do not necessarily run at the same speed. This is the case, for example, if a Web server is connected to a database server by means of a private network.

If more bandwidth is needed, the network must be upgraded or — in the case of shared-resource networks such as Ethernet — the network must be broken into subnets.

**Bandwidth and Capacity**

The main purpose of most Web servers is to manage input/output (I/O): Requests come in, and pages go out. Handling I/O requires a certain amount of bandwidth and other server resources as well. In addition to IIS 6.0, network I/O involves TCP/IP, which is implemented by Windows Server 2003 TCP/IP.

Network capacity is measured, in part, by the number of connections that the server establishes and maintains. Bandwidth is measured in several ways:

- By the rate at which bytes are transferred to and from the server.
- By the rate at which the server sends data packages, which include frames, packets, segments, and datagrams.
- By the rate at which the server sends and receives files.

Effective bandwidth varies widely and depends upon the transmission capacity of the link, the server configuration, and the server workload. The values for a single server also change as it operates in response to demand and to competition for shared network resources.

To ensure that your network has sufficient bandwidth and capacity for the network activity it must support, monitor the following performance indicators:

- Data transmission rates at the different Open Systems Interconnection (OSI) layers, because the components that transmit data reside in different layers.
- File transfer rates, because a Web page often requires multiple file transfers.
- TCP connections, because a plateau in connections established, or increases in connection failures and connection resets, can indicate insufficient bandwidth.

# Monitoring Data Transmission Rates at Different OSI Layers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The simplest measure of the effective bandwidth of a server is the rate at which the server sends and receives data. System Monitor displays counts of data transmissions that are collected by many components of the server computer.

The components that collect data reside in different layers of the OSI reference model:

- Counters on the Web, FTP, and Simple Mail Transfer Protocol (SMTP) services performance objects measure data transmitted at the OSI application layer.
- Counters on the TCP object measure data transmitted at the transport layer.
- Counters on the IP object measure data at the network layer.
- Counters on the Network Interface object measure data at the data-link layer.

As a result of their different positions in the OSI stack, the counters collect different data. For example, the counters at the application layer measure data in the form in which the application sends it, counting the bytes sent before the data is divided into packets and prefixed with protocol headers and control packets. Counters at the application layer do not include retransmitted data.

In addition, the counters display the data in units native to the component being measured. For example, the Web Service object displays data in bytes, and the TCP object displays data in segments.

For more information about TCP/IP and about the OSI reference model, see Additional Resources for IIS 6.0 Performance.

Tables 6.5 through 6.8 list and describe some of the counters that can be used to measure transmission rates. These counters display the transmission rates observed during the last sample interval; they do not display a rolling or cumulative average of the rate.

**Counters for Monitoring Data Transmission Rates at the Application Layer**

Table 6.5 lists counters at the application layer. As with all the Web Service counters, these collect data that shows the rate at which the WWW service (a user-mode service) is sending and receiving bits. These counters do not measure the transmittal rates for HTTP.sys.

**Table 6.5 Counters for Measuring Data Transmission Rates at the Application Layer**

| Object\Counter | Value |
|---|---|
| Web Service\Bytes Sent/sec | The rate at which the WWW service is sending data, in bytes per second. |
| Web Service\Bytes Received/sec | The rate at which the WWW service is receiving data, in bytes per second. |
| Web Service\Bytes Total/sec | The rate at which the WWW service is sending and receiving data, in bytes per second (the sum of Web Service\Bytes Sent/sec and Web Service\Bytes Received/sec). |
| FTP Service\Bytes Sent/sec | The rate at which the FTP service is sending data, in bytes per second. |
| FTP Service\Bytes Received/sec | The rate at which the FTP service is receiving data, in bytes per second. |
| FTP Service\Bytes Total/sec | The rate at which the FTP service is sending and receiving data, in bytes per second (the sum of FTP Service\Bytes Sent/sec and FTP Service\Bytes Received/sec). |

| | |
|---|---|
| SMTP Server\Bytes Sent/sec | The rate at which the SMTP server is sending data, in bytes per second. |
| SMTP Server\Bytes Received/sec | The rate at which the SMTP server is receiving data, in bytes per second. |
| SMTP Server\Bytes Total/sec | The rate at which the SMTP server is sending and receiving data, in bytes per second (the sum of SMTP Service\Bytes Sent/sec and SMTP Service\Bytes Received/sec). |

**Analyzing the data**

The IIS 6.0 service counters listed in Table 6.5 display the number of bytes transmitted on behalf of each service that runs on the server. To calculate the total number of bytes sent or received by all IIS 6.0 services, calculate the sum of the values for each service. To determine the proportion of bytes transmitted by each service, compute the ratio of bytes for one service to the sum of bytes for all services or for the network.

Data collected by the IIS 6.0 service counters underestimates the total number of bytes that the IIS 6.0 services actually transmit to the network. Because these values are collected at the application layer, they measure data only. They do not measure protocol headers, control packets, or retransmitted bytes. In general, the bytes counted by the services represent 60 to 70 percent of the total number of bytes transmitted by the services on the network. If the sum of bytes for all services accounts for two-thirds or more of total network bandwidth, you can assume that your network is running at or near the total capacity of its communications link.

If you are using bandwidth throttling to limit the amount of bandwidth that the WWW service or an individual Web site uses, or you want to evaluate whether bandwidth throttling might be useful, monitor the Web Service\Bytes Total/sec and the Web Service\Bytes Sent/sec counters. These counters can indicate whether your Web server has sufficient bandwidth to handle its load. If it does not, consider using bandwidth throttling to redistribute network availability.

For information about using bandwidth throttling, see Throttling Bandwidth to Manage Service Availability.

**Counters for Monitoring Data Transmission Rates at the Transport and Network Layers**

Table 6.6 lists the counters on the TCP object. These counters monitor *TCP segments* — the unit of data that TCP sends down the protocol stack to IP. Windows Server 2003 provides two TCP objects, TCPv4 and TCPv6. Choose the counter that monitors the version of TCP that your server uses.

**Table 6.6 Counters for Measuring Data Transmission Rates at the Transport Layer**

| Object\Counter | Value |
|---|---|
| TCPv4\Segments Sent/sec<br><br>TCPv6\Segments Sent/sec | The rate at which TCP segments are sent by using the TCP protocol. |
| TCPv4\Segments Received/sec<br><br>TCPv6\Segments Received/sec | The rate at which TCP segments are received by using the TCP protocol. |
| TCPv4\Segments/sec<br><br>TCPv6\Segments/sec | The rate at which TCP segments are sent and received by using the TCP protocol (the sum of Segments Sent/sec and Segments Received/sec). |
| TCPv4\Segments Retransmitted/sec<br><br>TCPv6\Segments Retransmitted/sec | The rate at which segments are transmitted that contain one or more bytes that TCP recognizes as having been transmitted before. Segments Retransmitted/sec is a subset of Segments Sent/sec and Segments/sec. To determine the proportion of transmissions caused by failed transmission attempts, divide Segments Retransmitted/sec by Segments Sent/sec. |

Table 6.7 describes lists some of the datagram counters for the IP object. These counters monitor *IP datagrams*, which are the units of data that IP sends down the protocol stack to the network interface, such as a network adapter. The sum of IP\Datagrams/sec and IP\Datagrams Forwarded/sec represents the rate at which the server handles all IP datagrams.

**Table 6.7 Counters for Measuring Data Transmission Rates at the Network Layer**

| Object\Counter | Value |
|---|---|

| | |
|---|---|
| IPv4\Datagrams Sent/sec<br><br>IPv6\Datagrams Sent/sec | The rate at which IP datagrams are sent to the network interfaces. This counter does not include datagrams forwarded to another server. |
| IPv4\Datagrams Received/sec<br><br>IPv6\Datagrams Received/sec | The rate at which IP datagrams are received from the network interfaces. This counter does not include datagrams forwarded to another server. |
| IPv4\Datagrams/sec<br><br>IPv6\Datagrams/sec | The overall transmission rate for IP datagrams being sent and received over the network interfaces (the sum of IP\Datagrams Sent/sec and IP\Datagrams Received/sec). |
| IPv4\Datagrams Forwarded/sec<br><br>IPv6\Datagrams Forwarded/sec | The rate, in incidents per second, at which the server attempts to find routes over which to forward IP datagrams to their final destination. |

**Analyzing the data**

Counters on the TCP and IP performance objects display the rate at which data is sent and received over a TCP/IP connection at the transport and network layers, but the rate is not expressed in bytes. Counters on the IP performance object display data in datagrams, and counters on the TCP performance object display data in segments. It is difficult to convert segments to bytes because bytes per segment can vary from 8 KB to 64 KB (the size can increase to 1 gigabyte [GB] if a window scaling option is in effect), depending upon the size of the TCP/IP receive window and the maximum segment size negotiated when each connection is established.

**Counters for Monitoring Data Transmission Rates at the Data-Link Layer**

Table 6.8 lists several counters on the Network Interface performance object that might be useful for obtaining data about the network adapters on the server. The first instance of the Network Interface object that you see in System Monitor represents the *loopback* — a local path through the protocol driver and the network adapter; all other instances represent installed network adapters.

**Table 6.8 Counters for Measuring Data Transmission Rates at the Data-Link Layer**

| Object\Counter | Value |
|---|---|
| Network Interface (*Adapter ID*)\Bytes Sent/sec | The rate, in seconds, at which bytes are sent over this network adapter. The counted bytes include framing characters. This counter is a subset of Network Interface\Bytes Total/sec. |
| Network Interface (*Adapter ID*)\Bytes Received/sec | The rate, in seconds, at which bytes are received over this network adapter. The counted bytes include framing characters. This counter is a subset of Network Interface\Bytes Total/sec. |
| Network Interface (*Adapter ID*)\Bytes Total/sec | The rate, in bytes per second, at which bytes are sent and received over this network adapter (the sum of Network Interface\Bytes Sent/sec and Network Interface\Bytes Received/sec). |
| Network Interface\Packets Sent/sec | The rate, in seconds, at which packets are sent over the network adapter. |
| Network Interface\Packets Received/sec | The rate, in seconds, at which packets are received over the network adapter. |

**Analyzing the data**

Counters for the Network Interface performance object display the rate at which bytes are transmitted over a TCP/IP connection by monitoring the counters on the network adapter at the data-link layer. The values of these Network Interface counters include all prepended frame header bytes and bytes that have been retransmitted. They provide a relatively accurate estimate of the number of bytes transmitted over the network, but do not measure the bytes transmitted by a specific IIS 6.0 service.

It is useful to compare the data produced by these counters with other performance measures, such as the total number of connections served at a given bandwidth or processor use at different throughput rates.

# Monitoring File and Message Transfers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Most static Web pages include multiple files, such as a file of text and one or more graphics files. File counters are available for each IIS 6.0 service. The counters for the WWW and FTP services count the total number of files that the services send and receive. The counters for the SMTP service tally total messages sent and total messages received, but not total messages sent and received.

Table 6.9 lists the counters for monitoring file and message transfers.

**Table 6.9 Counters for Monitoring IIS 6.0 File and Message Transfers**

| Object\Counter | Value |
|---|---|
| Web Service\Total Files Sent<br>FTP Service\Total Files Sent<br>SMTP Server\Messages Sent Total | For the WWW and FTP services, the total number of files sent by the service since service startup.<br>For the SMTP service, the total number of outbound messages sent since service startup. |
| Web Service\Files Received<br>FTP Service\Total Files Received<br>SMTP Server\Messages Received Total | For the WWW and FTP services, the total number of files received by the service since service startup.<br>For the SMTP service, the total number of inbound messages received since service startup. |
| Web Service\Total Files Transferred<br>FTP Service\Total Files Transferred | For the WWW and FTP services, the sum of Total Files Sent and Total Files Received by the service since service startup.<br>The SMTP service does not provide a sum total counter. |

# Analyzing the Data

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The file counters for a particular service are an indicator of the network activity generated by that service. This data can also be associated with other performance measures to determine the effects of high and low file activity on server components.

The file counters in Table 6.9 display a cumulative total for all traffic since the service was started, regardless of when System Monitor was started. These counters do not display the current value or the rate at which files are transmitted. In IIS 6.0, the Web Service object (but not the FTP Service object) provides counters to measure the rate, in seconds, at which files are sent and received. The SMTP service provides similar counters that measure the rate at which messages are received and sent.

To calculate file transmission rates for the FTP service, you can use Performance Logs and Alerts to log the values for the file counters and the times at which measurements are taken. To derive the transmission rates, you can export the log files to a spreadsheet that associates the time of the measurement with the file count.

# Monitoring TCP Connections

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If the bandwidth of your server is insufficient to handle its workload, clients usually become aware of this before the server does. Client requests to the server might be rejected or time out or the response might be delayed. On the server side, the indicators are less clear because the server continues to establish connections, receive requests, and transmit data.

Bandwidth shortages are not uncommon. You can detect one on your server (perhaps even before clients do) by monitoring the success and failure of connections established and rejected by TCP. With ample bandwidth, the server can establish and serve connections before they time out. If bandwidth is insufficient, the connections fail.

Table 6.10 lists the counters that monitor the success and failure of connections to TCP. Windows Server 2003 provides two TCP objects, TCPv4 and TCPv6; choose the counter that monitors the version of TCP that your server is using.

**Table 6.10 Counters for Monitoring the Success or Failure of TCP Connections**

| Object\Counter | Value |
|---|---|
| TCPv4\Connections Established<br><br>TCPv6\Connections Established | The number of simultaneous connections supported by TCP. This counter displays the number of connections last observed to be in the ESTABLISHED or CLOSE-WAIT state. This counter displays the last observed value (indicating the current state); its value is not an average. |
| TCPv4\Connection Failures<br><br>TCPv6\Connection Failures | The number of connections that have failed since the service was started (regardless of when you started System Monitor). TCP counts a connection as having failed when it goes directly from sending (SYN-SENT) or receiving (SYN-RCVD) to CLOSED, or from receiving (SYN-RCVD) to listening (LISTEN). |
| TCPv4\Connections Reset<br><br>TCPv6\Connections Reset | The number of connections reset since the service was started (regardless of when you started System Monitor).<br><br>TCP counts a connection as having been reset when it goes directly from ESTABLISHED or CLOSE-WAIT to CLOSED. |

The counters on the TCPv4 and TCPv6 objects are the best indicators of the success of connection requests. The counters on the Web Service and FTP Service performance objects monitor connections maintained by each IIS 6.0 service, but display only successful connection requests, not failed attempts. Like all counters at the application layer, they do not have information about connections until the connections are established.

For a discussion of performance counters that display the number of simultaneous connections maintained by IIS 6.0, see Preventing Processor Bottlenecks.

**Analyzing the Data**

If your server does not support current or increasing demand, look for the following patterns when you analyze the data from the TCP\Connections counters:

- Look for a plateau in the Connections Established counter value. If the counter value of TCP\Connections Established often reaches, but rarely exceeds, a maximum value (that is, the line in the graph rises and then reaches a plateau), the peak value is likely to indicate the maximum number of connections that can be established with the current bandwidth and application workload. If you observe such a pattern, the server probably cannot support any greater demand.

- Look for consistent increases of the Connection Failures and Connections Reset counter values. An increasing number of failures and resets, or a consistently increasing rate of failures and resets, can indicate a bandwidth shortage. The counters that monitor failures and resets show cumulative values, but you can use Performance Logs and Alerts to set alerts on the values or to log values over time. You can then use a spreadsheet to calculate the rates at which connections are rejected and reset.

**Note**

Be cautious when interpreting the number of reset connections shown by the TCP\Connections Reset counter. Resets do not always indicate dropped or failed connections. To minimize connection overhead, many browsers routinely close connections by sending a TCP reset (RST) packet rather than by using a normal close operation. The TCP\Connections Reset counter does not distinguish between connections that are reset because they are dropped and those that are reset in order to be closed abruptly.

# Administering Network Resources

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 provides a set of Quality of Service (QoS) features to help you maintain acceptable service levels of data transmission on your network. The goal of QoS is to ensure that particular sites or applications do not monopolize server resources, such as memory or CPU cycles, which can adversely affect performance. QoS helps administrators control how IIS components — such as sites, application pools, or the WWW service as a whole — use resources.

If the bandwidth on your server is not sufficient to support demand, you can increase overall server bandwidth. You can also increase the effective bandwidth of existing communication links. Some suggestions on how to do so follow; many involve setting parameters that can only be modified by editing the Windows registry or the IIS 6.0 metabase.

It is frequently possible to reduce your use of bandwidth by optimizing Web application scripts or content. This option is worth looking into as an interim solution, partly because it can also improve response time and the user experience; however, if your client base is growing, you might eventually have to increase the bandwidth of your network connection anyway.

You can sometimes effectively increase existing bandwidth by limiting connections, by increasing the length of the connection queues, or by enabling HTTP Keep-Alives.

QoS in IIS provides the following methods for managing and maintaining network performance:

- Limiting Connections to Manage Resources: Set limits on the number of connections allowed to a Web or FTP server or to specific sites. Use this feature to ensure that all your services remain active during periods of peak use and to protect against malicious attack.

- Enabling HTTP Keep-Alives to Keep Connections Open: To maintain an open connection when browsers send multiple requests to download a Web page, be sure that HTTP Keep-Alives are enabled.

- Setting Connection Timeouts to Save Resources: To reduce the loss of processing resources consumed by idle connections, set connection timeouts on your Web server.

- Using HTTP Compression for Faster Downloads: To provide faster transmission time between IIS and compression-enabled browsers, use HTTP compression.

- Throttling Bandwidth to Manage Service Availability: Limit the bandwidth used by your Web server so that the server remains available for other services, such as e-mail or news, even during periods of peak use.

- Using Other Features to Enhance Performance: In addition to the QoS features, use other IIS features, including Web gardens, idle timeout, CPU monitoring, and processor affinity, to help you manage system resources and enhance performance.

# Limiting Connections to Manage Resources

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Connection limits restrict the number of simultaneous connections to your Web sites and your Web server. If the number of connections reaches the maximum allotted, all subsequent connection attempts return an error and then are disconnected.

By limiting connections, you can conserve bandwidth for other uses, such as e-mail servers, news servers, or another Web site running on the same installation. Limiting connections also conserves memory and protects against malicious attacks designed to overload your Web server with thousands of client requests.

To determine whether you need to limit connections, use System Monitor to log the Current Connections, Maximum Connections, and Total Connection Attempts counters on the Web Service and FTP Service objects. Continue logging until you have a good sense of the normal range; typically, logging can take several days to a week or more and must be repeated at regular intervals. For more information about performance logging, see Enabling Logging and Analyzing Log Files.

You can establish global WWW service or FTP service connection limits for all Web and FTP sites, or you can establish connection limits on individual Web or FTP sites. IIS checks for a global connection limit before checking for connection limits on individual sites. If the global connection limit is exceeded, IIS returns error message 403.9 (Forbidden - Too Many Users) regardless of the connection limit set for individual sites. (You can customize the connection limit error message in IIS 5.0; however, IIS 6.0 does not offer that option.) For information about limiting connections for the WWW service, the FTP service, a Web site, or an FTP site, see Limiting Connections.

Another way to limit connections to your Web server is through bandwidth throttling. For information about using this feature, see Throttling Bandwidth to Manage Service Availability. For a list of counters related to performance, see IIS 6.0 Performance Counters. For information about connection timeouts, see Setting Connection Timeouts to Save Resources.

# Enabling HTTP Keep-Alives to Keep Connections Open

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A browser typically makes multiple requests in order to download an entire Web page. To enhance server performance, most Web browsers request that the server keep the connection open across these multiple requests, which is a feature known as *HTTP Keep-Alives*.

Without HTTP Keep-Alives, a browser that makes numerous requests for a page containing multiple elements, such as graphics, might require a separate connection for each element. These additional requests and connections require extra server activity and resources, decreasing server efficiency. The additional connections also make a browser much slower and less responsive, especially across a slow connection.

HTTP Keep-Alives are enabled by default in IIS 6.0, which complies with the HTTP/1.1 specification for HTTP Keep-Alives. IIS holds open an inactive connection for as long as the **ConnectionTimeout** metabase property specifies (the default value is 120 seconds).

If you disable HTTP Keep-Alives, the server ignores a clients request to keep the connection open. Therefore, disable HTTP Keep-Alives only for a specific reason and if you clearly understand how this change affects your server.

HTTP Keep-Alives are required for integrated security or connection-based authentication services, such as Integrated Windows authentication. If you disable HTTP Keep-Alives for Web sites that use Integrated Windows authentication, requests to the Web site may fail.

To verify that your server is running with HTTP Keep-Alives enabled, or to enable HTTP Keep-Alives if this feature is disabled, see Enabling HTTP Keep-Alives.

# Setting Connection Timeouts to Save Resources

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Connection timeouts help reduce the amount of memory resources that are consumed by idle connections. Time-out settings also allow you to specify how long server resources are allocated to specific tasks or clients. When you enable connection timeouts, IIS 6.0 enforces the following types of connection timeouts at the connection level:

- A *connection timeout*, in which the client has sent data to the server, but the client is now idle. Use the **ConnectionTimeout** metabase property to set a connection timeout limit for the WWW, FTP, Network News Transfer Protocol (NNTP), and SMTP services.

- A *request timeout*, which prevents clients from issuing unreasonably slow requests to the server (for example, 1 bit per second). Use the **HeaderWaitTimeout** metabase property to set a request timeout for the WWW service.

- A *response timeout*, which prevents malicious or malfunctioning clients from consuming resources by holding a connection open with minimal data. Use the **MinFileBytesPerSec** metabase property to set a response timeout for the WWW service.

**Monitoring with Counters to Evaluate Connection Limits**

In IIS 6.0, the default connection timeout settings are more restrictive than in earlier versions of IIS, which helps prevent denial of service attacks on the server. To determine whether you can improve performance by changing a default connection timeout setting or by adding an optional setting, begin by obtaining a baseline of how your server performs with the current connection limits. For example, use System Monitor to log the Current Connections, Maximum Connections, and Total Connection Attempts counters on the Web Service and FTP Service objects. Continue logging until you have a good sense of the normal range; typically, logging can take several days to a week or more and must be repeated at regular intervals.

After obtaining baseline performance data for the default configuration, make incremental changes to the connection timeout settings, and then collect additional performance data by using these same counters. Compare the results to determine if changing the connection limits improves performance, keeping in mind that more aggressive limits can increase protection against malicious attacks.

**Setting Connection Timeouts by Using IIS Manager**

You can set global connection timeouts for the WWW or FTP service, or for individual Web sites and FTP sites. You can also set global connection timeouts on SMTP and NNTP servers. For more information about setting connection timeouts, see Setting Connection Timeouts.

**Setting Connection Timeouts by Editing the Metabase**

IIS 6.0 provides three metabase properties, **ConnectionTimeout**, **HeaderWaitTimeout**, and **MinFileBytesPerSec**, which you can use to set different types of connection timeouts. In IIS 6.0, these properties replace the **ServerListenTimeout** metabase property, which is no longer used for the WWW service but can be used for the FTP, SMTP, and NNTP services.

**Setting connection timeouts**

The **ConnectionTimeout** metabase property specifies the amount of time (in seconds) that the server waits before disconnecting an inactive connection. IIS applies this timeout limit after the client sends the first request to the server and the client is idle. The default value is 120 seconds for the WWW and FTP services (global settings); 120 seconds for individual Web and FTP sites; and 10 minutes for the SMTP and NNTP services. (In IIS Manager, when you change the value of the **ConnectionTimeout** property, you change this setting.)

For security reasons, the **ConnectionTimeout** property cannot be disabled. Thus, if you try to set the **ConnectionTimeout** property to 0, the property retains its previous setting.

**Setting request timeouts**

The **HeaderWaitTimeout** metabase property specifies the amount of time (in seconds) that the server waits for the client computer to send all HTTP headers for a request (indicated by a double carriage return) before HTTP.sys resets the connection. The purpose of this property is to help impede a type of denial of service attack that attempts to exhaust connection limits and keep those connections connected. You can apply this connection timeout only at the WWW service level.

For security reasons, the **HeaderWaitTimeout** property cannot be disabled. Thus, if you try to set the **HeaderWaitTimeout** property to 0, the property retains its previous setting.

**Setting response timeouts**

The**MinFileBytesPerSec**metabase property determines the length of time that the client has to receive the server's entire response to its request. If the client computer does not receive the entire HTTP response within the interval set by the time-out value (by default, 240 bytes per second), HTTP.sys terminates the connection. You can apply this connection timeout only at the WWW service level.

Configuring the **MinFileBytesPerSec** metabase property prevents a client computer from sending a request for a large response (such as a file download) and then receiving the response at a maliciously slow rate that is meant to consume resources on the server and potentially interrupt service for other client computers.

The time-out period is calculated by dividing the size of the entire response (including headers) by the value of the **MinFileBytesPerSec** property to obtain a maximum allowable response time, in seconds. For example, a 2-KB response (2,048 bytes) is allowed 8.5 seconds to complete if **MinFileBytesPerSec** has the default value of 240 bytes per second.

To accommodate very slow applications, you can disable the **MinFileBytesPerSec** property by setting the value to 0.

**Reference to Default Time-out Settings**

Additional IIS 6.0 metabase properties set time-out values for ASP, Common Gateway Interface (CGI) scripts, and Internet database connection pooling. Table 6.11 gives a summary of the metabase properties for setting timeouts and the default time-out limit for each property. For information about configuration options, see Code Examples to Configure Metabase Properties. The final column of the table indicates which properties can alternatively be updated in IIS Manager.

**Table 6.11 Default Time-out Values for IIS 6.0**

| Metabase Property | Default Time-Out Value | Configured in IIS Manager |
|---|---|---|
| **AspQueueTimeout** | Unlimited | |
| **AspScriptTimeout** | 90 seconds | Yes |
| **AspSessionTimeout** | 20 minutes | Yes |
| **CGITimeout** | 300 seconds | Yes |
| **ConnectionTimeout** | 120 seconds (Web and FTP);<br>10 minutes (SMTP and NNTP) | Yes |
| **HeaderWaitTimeout** | None (Turned off by default.) | |
| **MinFileBytesPerSec** [1] | 240 bytes per second | |
| **PoolIdcTimeout** | None (Turned off by default.) | |

[1] This metabase property cannot be modified in IIS Manager, but it can be modified by adding the **MinFileBytesPerSec** entry to the Windows registry.

For more information about ASP–related properties and counters, see Monitoring ASP Performance. For information about the registry path for the **MinFileBytesPerSec** entry, see Global Registry Entries.

Another way to limit connections to your Web server is to use bandwidth throttling. For information, see Throttling Bandwidth to Manage Service Availability. A related way to manage resources is to limit the number of simultaneous connections to your sites and server. For information about limiting connections, see Limiting Connections to Manage Resources.

# Using HTTP Compression for Faster Downloads

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If your Web sites use large amounts of bandwidth or if you want to use bandwidth more effectively, consider enabling HTTP compression, which provides faster transmission times between IIS and compression-enabled browsers regardless of whether your content is served from local storage or a UNC resource. If your network bandwidth is restricted, HTTP compression can be beneficial unless your processor usage is already very high.

IIS provides the following compression options:

- Static files only.
- Dynamic application responses only.
- Both static files and dynamic application responses.

Dynamic processing can affect CPU resources because IIS does not cache compressed versions of dynamic output. If compression is enabled for dynamic responses and IIS receives a request for a file that contains dynamic content, the response that IIS sends is compressed every time it is requested. Because dynamic compression consumes considerable CPU time and memory resources, use it only on servers that have slow network connections but CPU time to spare.

Compressed static responses can be cached and therefore do not affect CPU resources like dynamic responses do.

**How HTTP Compression Works**

When IIS receives a request, it checks whether the browser that sent the request is compression-enabled. (Recent versions of Microsoft® Internet Explorer and most other browsers typically send the following header if they are compression-enabled: Accept-Encoding: gzip, deflate.) IIS then determines whether the request is for a static file or for dynamic content.

If the content of the file is static, IIS checks whether the file has previously been requested and is already stored in a compressed format in the temporary compression directory. If a compressed version of the requested file is not found, IIS sends an uncompressed version of the requested file to the client browser while a background thread compresses the requested file. The newly compressed file is then stored in the compression directory, and subsequent requests for that file are serviced directly from the compression directory. In other words, an uncompressed version of the file is returned to the client unless a compressed version of the file already exists in the compression directory.

If the file contains dynamic content, IIS compresses the response as it is generated and sends the compressed response to the browser. No copy of the file is cached by the Web server.

The performance cost of compressing a static file is modest and is typically incurred only once, because the file is then stored in the temporary compression directory. The cost of compressing dynamically generated files is somewhat higher because the files are not cached and must be regenerated with each request. The cost of expanding the file at the browser is minimal. Compressed files download faster, which makes them particularly beneficial to the performance of any browser that uses a network connection with restricted bandwidth (a modem connection, for example).

When you enable HTTP compression, compressed files are given a default expiration date of Jan. 1, 1997. This expiration date prevents proxy servers from serving cached copies of compressed files to browsers that are not compression-enabled. This expiration date also forces browsers to return to the server for a fresh copy of the file when the user makes a new request instead of displaying a cached copy of the file. When you enable HTTP compression, the default settings for the **HcExpiresHeader**, **HcCacheControlHeader**, and **HcSendCacheHeaders** metabase properties ensure that older clients and proxy servers do not attempt to cache compressed files. Before you change these settings, see Metabase Property Reference for information about these and related metabase properties.

**Determining Whether HTTP Compression Will Improve Performance**

If your server generates a large volume of dynamic content, consider whether the additional processing cost of HTTP compression is one that you can reasonably afford. If the % Processor Time counter is already 80 percent or higher, enabling HTTP compression is not recommended.

To evaluate how much of your processor is typically being used, follow these steps:

1. Establish a baseline for your processor usage by using System Monitor to log the following counters over several days. If

you use Performance Logs and Alerts, you can log the data to a database and then query the data, examining the results in detail.

- Processor\% Processor Time. This counter has a total instance and a separate instance for each processor in the system. If your server has more than one processor, you need to watch the individual processors as well as the total to discover any imbalance in the workload.
- Network Interface\Bytes Sent/sec. Counters for the Network Interface performance object display the rate at which bytes are transmitted over a TCP/IP connection by monitoring the counters on the network adapter. For information about additional counters to monitor for this object, see Table 6.8.

2. Enable compression, and continue to log the values for these counters for an extended period — preferably for several days — so you have a good basis for comparison.

   Collect a broad sample to determine how compression affects various aspects of performance. Conduct the following tests:

   - Enable static compression only, dynamic compressions only, and both.
   - Change the list of files that you use for compression testing for both static and dynamic content.
   - Vary the compression level. Try this on all content types.

3. Compare the data from monitoring with and without compression, and with different types of compression, different compression levels, and different files.

| **Important** |
|---|
| If you see signs of blocking or bottlenecking during the preceding testing, promptly stop the test. A significant drop in the value of either counter indicates that performance with compression enabled has decreased relative to performance without compression enabled. |

**Choosing Compression Options**

When you configure HTTP compression by using IIS Manager, IIS 6.0 automatically applies the default settings shown in Table 6.12. You can configure custom settings for each compression option by editing the metabase or alternatively — in the case of the compression directory — by using IIS Manager.

**Table 6.12 Default Settings When HTTP Compression Is Enabled**

| Compression Options | File Type | Default Configuration | Configuration Method |
|---|---|---|---|
| File types compressed | Static | .txt, .htm, and .html | Metabase |
| | Dynamic | .exe, .dll, and .asp | Metabase |
| Compression schemes | Static | Both gzip and deflate | Metabase |
| | Dynamic | Both gzip and deflate | Metabase |
| Compression level | Static | 10 | Metabase |
| | Dynamic | 0 | Metabase |
| Compression directory | Static | Size: 95 MB<br>Location[1]: %Windir%\IIS Temporary Compressed Files | Metabase or IIS Manager |
| | Dynamic | No directory[2] | N.A. |

[1] The compression directory must reside on an NTFS file system volume.

[2] When dynamic files are compressed, they are not cached.

If the default HTTP compression configuration does not meet the needs of your organization, you can make the following changes to the configuration:

- Specify additional file types to compress. Edit the **HcFileExtensions** metabase property (for static files) and the **HcScriptFileExtensions** metabase property (for dynamic files) to apply compression to additional file types. Add only compressible files types to your compression scheme because trying to compress an already compressed file — such as .jpg, .mp3, or .zip files — wastes processing time.

- Enable or disable the gzip or deflate compression scheme. When you use IIS Manager to enable compression, IIS enables both compression schemes (gzip and deflate), so that IIS can apply whichever compression scheme the client requests. You can edit the metabase to specify that only a particular compression scheme is used. If you disable the only compression scheme that a client supports, the server replies with an uncompressed file.

  | **Note** |
  | --- |
  | Disabling either the gzip or deflate compression schemes can have unintended consequences, such as not compressing responses to a particular browser. For maximum compatibility, it is recommended that you leave both compression schemes enabled. |

- Change the level of compression for static or dynamic files. Compression levels range from 0 through 10. Higher compression levels produce smaller compressed files but use more CPU and memory. Lower compression levels produce slightly larger compressed files, but with less impact on CPU and memory usage. To configure a compression level other than 10 for static files and other than zero for dynamic files, edit the **HcOnDemandCompLevel** and **HcDynamicCompressionLevel** metabase properties.

  For dynamic compression, increasing the compression level can significantly increase CPU usage. The default compression level of zero uses the least amount of CPU resources and can increase performance if network bandwidth is adequate. Raise the dynamic compression level above zero only if you need more network bandwidth and have sufficient CPU capacity to handle the extra load.

- Change the location or size of the compression directory. To change the location and size of the compression directory for static files, edit the **HcCompressionDirectory**, **HcDoDiskSpaceLimiting**, and **HcMaxDiskSpaceUsage** metabase properties.

**Compression Strategies**

When the CPU of your server is not heavily loaded, the simplest compression strategy is to enable static and dynamic compression for all of the sites and site elements (directories and files) on the server. This is known as global HTTP compression. However, when the CPU load of your server is high, you might not want to enable compression for all of the sites and site elements on the server.

Depending on CPU usage, there are two strategies for fine tuning which sites and site elements are compressed:

- Enable compression globally, and then disable compression selectively for specific sites or site elements.

- Leave global compression disabled, and then enable compression selectively for specific sites or site elements.

**Enabling Global HTTP Compression**

When you enable HTTP compression by using IIS Manager, IIS applies compression according to the default settings shown in Table 6.12. With the exception of the location and size of the compression directory, you cannot change the default settings by using IIS Manager. Instead, you must edit the metabase to change the default settings.

| **Important** |
| --- |

After enabling HTTP compression (either by using IIS Manager or by editing the metabase), you must restart IIS for the change to take effect.

For more information about enabling global HTTP compression by using IIS Manager and through Adsutil.vbs, see Enabling HTTP Compression.

**Configuring HTTP Compression for Individual Sites and Site Elements**

All of the compression configurations discussed so far are located under W3SVC/Filters/Compression in the metabase and are used to apply global HTTP compression, which is applied by using IIS Manager or by using one of the metabase properties that are used for global compression (beginning with Hc). IIS 6.0 offers two new metabase properties, **DoStaticCompression** and **DoDynamicCompression**(beginning with Do), which can be used to enable or disable compression at the individual site and site element levels. Use these two new properties to apply compression to individual namespaces. Hence, you can enable global compression by using IIS Manager or the Hc global properties and then disable one of the Do properties (or vice versa) to obtain compression at a specific location in the IIS metabase.

For more information about enabling compression for individual Sites and Site elements, see Enabling HTTP Compression.

**Specifying File Types and Compression Schemes by Editing the Metabase**

After enabling HTTP compression, you can use the Adsutil.vbs command-line utility (located in the *systemroot*\Inetpub\AdminScripts folder) to edit the metabase, specifying which static or dynamic file types to compress and which compression scheme to apply.

Using Adsutil.vbs, you can customize HTTP compression in the following ways:

- Apply the gzip or deflate compression scheme, or both. To apply both schemes, you must execute a separate **set** command for each compression scheme.

- Customize the file types to which both static and dynamic compression are applied. You must execute a separate **set** command for each compression type. If you use IIS Manager to enable compression and you then use Adsutil.vbs to specify specific static or dynamic file types, you replace any default values that were applied by IIS Manager.

**Note**

For compression to be performed, the **HcDoStaticCompression** metabase property must be set to **true** (for static compression) and the **HcDoDynamicCompression** metabase property must be set to **true** (for dynamic compression). You can set these by using IIS Manager or Adsutil.vbs. For instructions, see the procedures in the preceding section.

The following syntax enables compression at the W3SVC level (which means it is applied to all Web sites and virtual directories on the Web server):

```
cscript.exe adsutil.vbs set W3svc/Filters/Compression/{GZIP|DEFLATE}/
{HcFileExtensions "StaticFileType" ...}|{HcScriptFileExtensions "DynamicFileType" ...}
```

You must restart the WWW service before any changes take effect. For information about restarting a service, see Common IIS 6.0 Administrative Tasks. Table 6.13describes the parameters for customizing file types and compression schemes.

**Table 6.13 Parameters for Setting Compression Schemes and File Types for HTTP Compression**

| Parameter | Description |
|---|---|
| `GZIP\|DEFLATE` | Specifies which compression scheme (either gzip or deflate) to apply to the specified file types. To enable both compression schemes, you must run this command once for each compression scheme. For maximum compatibility, leave both compression schemes enabled. |
| `HcFileExtensions "StaticFileType" ...` | Specifies which static file types to apply the compression scheme to — for example, .txt, .js, .css, .doc, or .xls files. To specify a file type, enclose the file name extension (without the leading period) in quotation marks. Separate multiple file types with a space. This parameter cannot be executed together with the `HcScriptFileExtensions "`*DynamicFileType*`"` parameter. |

| | |
|---|---|
| `HcScriptFileExtensions "DynamicFileType" ...` | Specifies which dynamic file types to apply the compression scheme to — for example, .exe, .dll, or .asp files. To specify a file type, enclose the file name extension (without the leading period) in quotation marks. Separate multiple file types with a space. This parameter cannot be executed together with the `HcFileExtensions "StaticFileType"` parameter.<br><br>**Important** Setting the `HcScriptFileExtensions` parameter to an empty string can adversely affect your servers performance. If this parameter is empty, all dynamic responses are sent compressed. In addition, any static file type not specified in `HcFileExtensions` is dynamically compressed and therefore not cached. |

For example, the following command uses a deflate compression scheme to apply static compression at the W3SVC level to text files (.txt) and cascading style sheets (.css):

```
cscript.exe adsutil.vbs set w3svc/Filters/Compression/DEFLATE/HcFileExtensions "txt"
"css"
```

#### Specifying the Location and Size of the Compression Directory for Static Compression

Three metabase properties enable you to change the location of the default directory for static compression and to limit its size.

#### Changing the location of the compression directory

The **HcCompressionDirectory** metabase property specifies the folder where compressed versions of static files are temporarily cached. By default, the compression directory is at %Windir%\IIS Temporary Compressed Files.

Due to a number of differences between NTFS and FAT volumes — such as differences in time stamping and access control mechanisms, and the efficiency with which NTFS handles directories that contain large numbers of files — the compression directory must reside on an uncompressed and unshared NTFS volume. In addition, the access control lists (ACLs) for the folder must include Full Control access to the identity of the application pool or to the IIS_WPG group. If a requested file resides on an NTFS volume and the compression directory is on a FAT volume, IIS does not perform HTTP compression on that file.

#### Setting a size limit for the compression directory

In addition, you can configure the **HcDoDiskSpaceLimiting** and **HcMaxDiskSpaceUsage** metabase properties to limit the amount of disk space that all files in the compression directory can occupy. **HcDoDiskSpaceLimiting** enables disk space limiting. If it is set to **true**, IIS limits compressed files to no more than the number of bytes specified by **HcMaxDiskSpaceUsage**. After the limit is reached, compressed files are removed from the compression directory on a least recently used basis. If **HcDoDiskSpaceLimiting** is set to **false**, no disk space limit is enforced for the compression directory.

**HcDoDiskSpaceLimiting** is relevant only if you have set the property **HcDoOnDemandCompression** to **true**. For more information about these metabase properties, see Metabase Property Reference.

| **Important** |
|---|
| Limiting the use of disk space has a significant impact on performance because additional overhead is required to check for the limit. Use this feature only when you lack adequate disk space to cache all content files on your Web server. |

#### Disabling HTTP Compression for Requests That Come Through Certain Proxy Servers

Certain HTTP proxy servers, including some advertised as HTTP 1.1–compliant, do not handle the caching of compressed objects correctly. The **HcNoCompressionForProxies** metabase property allows you to disable the HTTP 1.1 response for compression requests that come through proxy servers.

In addition, you can use the **HcNoCompressionForHttp10** metabase property to disable compression for requests that contain an HTTP 1.0 version number. HTTP 1.0, as described in RFC 1945, Hypertext Transfer Protocol — HTTP/1.0, provides a minimal level of support for certain types of compression. However, some confusion exists concerning HTTP 1.0 compression, especially with regard to proxy servers. To minimize the chance of inappropriately returning a compressed file to a client that cannot decompress it, you can use the **HcNoCompressionForHttp10** metabase property to disable compression in questionable scenarios.

For more information about the many metabase properties that you can set when configuring HTTP 1.1 compression, including a complete list of all the properties you can set, see IIsCompressionSchemes (ADSI) in the IIS 6.0 SDK.

# Throttling Bandwidth to Manage Service Availability

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If the network or Internet connection used by your Web server is also used by other services, such as e-mail or news, you can limit the bandwidth used by your Web server so that the server is available for those other services. If your Web server hosts more than one Web site, you can individually throttle the bandwidth used by each site. However, keep in mind that the IPv6 network traffic handled by your Web sites is not throttled. Bandwidth throttling uses Packet Scheduler to manage when data packets are sent.

**Determining When and Where to Use Bandwidth Throttling**
Before enabling bandwidth throttling, determine your Web server's typical load by using System Monitor to log the following counters over several days:

- Bytes Total/sec or Current Bandwidth counter on the Network Interface object. To compare incoming and outgoing traffic, examine both Bytes Sent/sec and Bytes Received/sec.

- Compare the values from the Network Interface object counters with the total bandwidth of your network connection.

If you use Performance Logs and Alerts, you can log the data to a database and then query the data, examining the results in detail.

For a typical load, your server should use no more than 50 percent of its total available bandwidth. If your server is subject to large peaks in use, keep your typical load lower than 50 percent. The remaining bandwidth can be used during peak periods.

If testing of your servers' typical loads shows that a server or a particular site is consistently using more than 50 percent of the available bandwidth, use the procedures that follow to throttle bandwidth.

In addition, before you set global bandwidth throttling, be sure that you understand the following:

- Individual sites throttle bandwidth according to their established maximum; the global setting limits the total network bandwidth that is available for all unthrottled Web sites on a server. Setting a global WWW service maximum bandwidth does not override established bandwidth maximums for individual Web sites on the server.

- Global bandwidth throttling settings for the WWW service do not affect FTP sites or the FTP service as a whole.

**Related Information**

- For more information about bandwidth throttling, including enabling packet scheduler, and enabling and disabling bandwidth throttling, see Throttling Bandwidth.

# Using Other Features to Enhance Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0, when running in worker process isolation mode, offers several features that you can use in addition to the QoS feature to better manage your systems resources:

- Use Web gardens to enhance performance.
- Conserve processor resources by using idle timeout.
- Enable CPU monitoring to conserve processing time.
- Assign processor affinity to better use CPU caching.

**Using Web Gardens to Enhance Performance**

If the server is running in worker process isolation mode, you can configure an application pool to be supported by multiple worker processes. An application pool that uses more than one worker process is called a *Web garden*. Creating a Web garden for an application pool enhances performance and reliability by providing the following enhancements:

- Robust processing of requests. When a worker process in an application pool is tied up (for example, when a script engine stops responding), other worker processes can accept and process requests for the application pool.
- Reduced contention for resources. When a Web garden reaches a steady state, each new TCP/IP connection is assigned, according to a round-robin scheme, to a worker process in the Web garden. The round-robin scheme smoothes out workloads and reduces contention for resources that are bound to a worker process.

For more information about Web gardens, including how to configure them, see Running IIS 6.0 as an Application Server.

**Conserving Processor Resources by Using Idle Timeout**

If the server is running in worker process isolation mode, you can conserve system resources by terminating idle worker processes; that is, you can configure an application pools worker process to shut down after a specified period of inactivity. Use this feature to better manage your resources when the processing load is heavy, when identified worker processes are consistently idle, or when new processing space is not available. If an application pool times out because the inactivity of its worker process exceeds a configured limit for idle timeout, and if HTTP.sys then receives a request for the idle application pool, the WWW service starts a new worker process to replace the one that timed out.

For more information about configuring idle timeout for a worker process, see Running IIS 6.0 as an Application Server.

**Enabling CPU Monitoring to Conserve Processing Time**

Through CPU monitoring of application pools, you can monitor the CPU usage of worker processes and, if you configure it to do so, automatically shut down worker processes that consume large amounts of CPU processing time. CPU monitoring logs and throttles only the applications that are isolated in a process that is separate from IIS; hence, this feature is only available when your server is running in worker process isolation mode.

When you configure CPU monitoring for an application pool, you can choose from two monitoring options:

- Perform only event logging. When the CPU usage for an application pool exceeds the Maximum CPU use limit, the WWW service writes an error to the Windows event log. The log entry identifies the specific worker process and the application pool that exceeded the CPU limit.
- Log events, and also shut down the unhealthy application pool. After an application pool reaches the configured Maximum CPU use limit, CPU monitoring writes an error to the event log and shuts down the application pool. The application pool remains down until the **CPUResetInterval** time limit expires, at which time the application pool restarts. If any requests are received while the application pool is stopped, HTTP.sys returns a 503 error (Server Too Busy).

For more information about CPU monitoring, including how to configure it, see Running IIS 6.0 as an Application Server.

**Obtaining more detailed CPU monitoring**

If your Web server is running the Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, operating system, you can use the Windows System Resource Manager to monitor different types of CPU activity more closely. With this feature, you can control how CPU and memory resources are allocated to applications, services, and processes. However, Windows System Resource Manager cannot manage system resources correctly if the computer is already being managed by another resource manager, such as IIS CPU monitoring. For best results, if you use Windows System Resource Manager, do not use any other resource manager on that computer.

**Related Topics**

For more information about Windows System Resource Manager, see Windows System Resource Manager — Fast Facts.

# Controlling Memory Usage

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

With memory, more is almost always better. However, just adding more RAM to your server does not optimize performance. You must first determine whether you are making good use of the memory that the server has.

The biggest performance gains with IIS occur when you resolve issues related to insufficient memory. Before changing your hardware configuration, rule out memory problems. Monitor memory first to verify that your server has enough memory, and then move on to other parts of your server environment. Problems caused by memory shortages can often appear to be problems in other parts of the system.

**Minimum Memory Requirements for IIS 6.0**
To run IIS 6.0 with the Windows Server 2003 operating system, a dedicated Web server must have a minimum of 128 MB of RAM. However, your server might require more RAM to handle the resource needs of custom applications. If your custom applications require large amounts of memory, provide 256 MB or more of RAM. Additional memory is particularly beneficial for e-commerce sites, sites with a lot of content, and sites with a high volume of traffic.

**Memory Management and Memory Allocation in Windows Server 2003**
In order to use memory most efficiently, the Windows Server 2003 memory system is largely self-tuning. Similarly, IIS 6.0 regulates the size of the IIS object cache. Therefore, the primary purpose of monitoring memory on a Windows Server 2003–based server running IIS 6.0 is not to adjust the size of each memory component, but rather to make sure that the system as a whole has enough memory and is using it efficiently.

The largest chunk of activity in Windows Server 2003 is a *process*. The physical RAM available to a process is known as its *working set*. If the needs of a process exceed its working set — that is, if the process is not able to store all its code and frequently used data in RAM — some of the information is stored elsewhere, usually on disk as virtual memory.

Storing information on disk, instead of in memory, increases disk activity, which slows down the server. Some memory contains information that is frequently used, and sending that data to disk instead of retaining it in memory can degrade performance.

In addition, a section of system memory, known as *pool memory*, is reserved for use by the operating system. Pool memory is either *paged* or *nonpaged*. Nonpaged pool is the section of pool memory that is not pageable. Data in a nonpaged pool is always stored in physical memory and never saved to the disk. This section usually reserved for device drivers or other critical system components that cannot wait for memory to be paged. Paged pool is the section of pool memory that can be managed in the same way as other virtual memory - that is, it is pageable in the inactive memory can reside on the disk, and the current or active memory can reside in physical memory.

Within each process, Windows Server 2003 uses threads to accomplish particular tasks. The memory occupied by threads is part of the working set of the process, except for some special threads that run in nonpageable memory. Threads that service connections in IIS 6.0, for example, cannot be paged.

This section tells how to monitor various aspects of overall server memory — available memory, paging, use of the file system cache, the size of paging files, the IIS working set, and disk I/O — and, based on your observations, how to optimize memory usage.

# Monitoring Overall Server Memory

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use System Monitor to check the values of the counters described in the following sections. This will help you determine if the current amount of memory on your server is sufficient for your needs.

To evaluate memory usage on your Web server, look at available memory, paging, the size of the file system cache, the size of the paging files, the size of the components of the IIS 6.0 working set, and disk I/O.

# Monitoring Available Memory

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Monitor the counter in Table 6.14 to determine whether the server has performance bottlenecks associated with memory.

**Table 6.14 Counter for Monitoring Available Memory**

| Object\Counter | Value |
|---|---|
| Memory\Available Bytes | The amount of physical memory, in bytes, available to processes running on the computer. This counter displays the last observed value only; it is not an average. |

Compare the total physical memory that is available to Windows Server 2003 with the available memory remaining when all server services are running. To gather reliable data, log this value over time, being sure to include periods of peak activity. Try to reserve at least 10 percent of available memory for peak use.

# Monitoring Paging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If a process requests a page in memory and the system cannot find the page at the requested location, a page fault occurs. If the page is elsewhere in memory, the fault is a *soft page fault*. If the page must be retrieved from disk, the fault is a *hard page fault*. Most processors can handle large numbers of soft page faults without consequence. However, hard page faults can cause significant delays. Continuous high rates of disk paging indicate a memory shortage.

Table 6.15 lists the counters that are available for monitoring hard and soft page faults.

**Table 6.15 Counters for Monitoring Page Faults, Pages Input, and Page Reads**

| Object\Counter | Value |
|---|---|
| Memory\Page Faults/sec | The overall rate at which the processor handles both hard and soft page faults. |
| Memory\Pages Input/sec | The total number of pages read from disk to resolve hard page faults. |
| Memory\Page Reads/sec | The number of times that the disk was read to resolve hard page faults. |
| Memory\Transition Faults/sec | The rate at which page faults are resolved by recovering pages without incurring additional disk activity. *Transition faults*, which measure soft page faults, are counted in numbers of faults because only one page is faulted in each operation; the number of transition faults is equal to the number of pages faulted. |

Hard page faults increment the Page Faults/sec counter (which reports the rate of hard and soft page faults combined) in addition to the Pages Input/sec and Page Reads/sec counters (which report the rate of hard page faults only). The Transition Faults/sec counter is only incremented by soft page faults.

The value of the Memory\Pages Input/sec counter (the number of pages read from disk as a result of hard page faults) will be greater than or equal to the value of Page Reads/sec (the number of times the disk was read as a result of hard page faults), and can give you a good idea of your hard page fault rate. If these numbers are low, your server is responding to requests quickly. If these numbers are high, investigate whether you have dedicated too much memory to the caches, leaving too little memory for the rest of the system. If reducing cache sizes does not improve system performance, you might need to increase the amount of RAM on your server.

# Monitoring the File System Cache

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Monitor the counters in Table 6.16 to determine whether a server has a performance bottleneck associated with the size of the file system cache. These counters display the last observed value only; they do not display an average value.

**Table 6.16 Counters for Monitoring the File System Cache**

| Object\Counter | Value |
|---|---|
| Memory\Cache Bytes | The current size, in bytes, of the file system cache. By default, the cache uses up to 50 percent of available physical memory. The counter value is the sum of Memory\System Cache Resident Bytes, Memory\System Driver Resident Bytes, Memory\System Code Resident Bytes, and Memory\Pool Paged Resident Bytes. |
| Memory\System Cache Resident Bytes | The current size, in bytes, of the pageable operating system code in the file system cache. This value includes only current physical pages and excludes any virtual memory pages not currently resident. |
| Memory\System Driver Resident Bytes | The current size, in bytes, of the pageable physical memory in use by device drivers. This represents the *working set* (physical memory area) of the drivers. This value is a component of Memory\System Driver Total Bytes. |
| Memory\System Code Resident Bytes | The current size, in bytes, of the operating system code currently in physical memory that can be written to disk when it is not in use. This value is a component of Memory\System Code Total Bytes. |
| Memory\Pool Paged Resident Bytes | The current size, in bytes, of the paged pool. The paged pool is an area of *system memory* (physical memory used by the operating system) reserved for objects that can be written to disk when they are not in use. The space that is used by the paged and nonpaged pools is taken from physical memory; thus, a pool that is too large denies memory space to processes. |

The value of the Memory\System Cache Resident Bytes counter equals the System Cache value shown on the **Performance** tab in Task Manager. As a result, this value can be smaller than the actual amount of virtual memory in use by the file system cache.

Windows Server 2003 automatically trims the cache when memory is scarce and enlarges the cache when memory is ample. Use the Memory\Cache Bytes counter to monitor the size of the file system cache. A log of cache size readings reveals how the size of the file system cache changes over time. Compare the log data to a measure of general memory availability, such as data from the Memory\Available Bytes counter. By monitoring these counters and watching the trends, you can expose memory shortages in a server.

When examining the logs, note when the cache is smallest. Track how small the cache gets and how often its size is reduced. Also, note how much system memory is available when the cache size is reduced. By watching these trends, you can begin to associate the size of the cache with the servers performance. For example, if performance degrades when the cache is large, consider adding RAM. If performance degrades when the cache is small, consider increasing the time interval during which objects remain in the cache before being flushed. For more information about monitoring and tuning IIS file cache, see Web Server Scalability.

# Monitoring the Size of the Paging Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The counters in Table 6.17 are useful in examining paging file usage. *Paging files* store pages of memory used by a process that are not contained in other files. Paging files are shared by all processes, and the lack of space in paging files can prevent processes from allocating memory.

**Table 6.17 Counters for Monitoring Page Files**

| Object\Counter | Value |
|---|---|
| Process\Page File Bytes | The current amount of virtual memory, in bytes, that this process has reserved for use in the paging files. |
| Process\Page File Bytes Peak | The maximum amount of virtual memory, in bytes, that this process has reserved for use in the paging files. |
| Paging File\% Usage | The percentage of the Page File instance that is in use. |
| Paging File\% Usage Peak | The peak usage of the Page File instance, expressed as a percentage of total file size. |

Windows Server 2003 creates a paging file (with the file name Page File) on the system drive. You can create a paging file on each logical disk. By striping a paging file across separate physical drives, you can improve paging file performance by using drives that do not contain your sites content or log files.

You can change the sizes of existing paging files. The larger the paging file, the more memory the system commits to it. Be sure that the paging file on the system drive is at least twice the size of the physical memory, so that the system can write the entire contents of RAM to disk if the system unexpectedly locks up or shuts down.

# Monitoring the IIS 6.0 Working Set

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If free memory on the computer is above a threshold, pages are left in the working set of a process even if they are not in use. When free memory falls below that threshold, pages are trimmed from working sets. If the trimmed pages are later needed, they are soft-faulted back into the working set before leaving main memory.

When available memory falls below about 4 MB, the system attempts to make more memory available by taking memory away from the working sets of processes. This strategy increases the rate of page faults, because each process must now retrieve data that was once in its working set either from disk or from elsewhere in memory. When the rate of page faults for a particular process rises, however, the system attempts to expand the working set of that process to lower the page fault rate. As the system tries to maintain this balance, the size of the working set of each process fluctuates accordingly.

The Process\Working Set counters measure the number of memory pages that each process uses. If the system has sufficient memory, it can maintain enough space in the working set so that IIS 6.0 rarely must perform disk operations. One indicator of memory sufficiency is how much the size of the process working set fluctuates in response to general memory availability on the server. Significant fluctuation can indicate a lack of available memory.

The systems memory pools hold objects created and used by applications and the operating system. The contents of the memory pools are accessible only in privileged mode — that is, only the kernel of the operating system can directly use the memory pools; user processes cannot.

Table 6.18 describes the counters that are available for monitoring the IIS 6.0 working set.

**Table 6.18 Counters for Monitoring the IIS 6.0 Working Set**

| Object\Counter | Value |
|---|---|
| Memory\Pool Paged Bytes<br>Memory\Pool Non-paged Bytes | Pool Paged Bytes and Pool Non-paged Bytes monitor the pool space for all processes on the server. |
| Process(W3wp1)\Virtual Bytes<br>Process(Inetinfo)\Virtual Bytes | The amount of virtual address space, in bytes, reserved directly by IIS, either by the Inetinfo.exe process (where IIS runs when set to IIS 5.0 isolation mode) or by the W3wp.exe process (where isolated or pooled applications run when IIS is set to worker process isolation mode) instantiated on your server. |
| Process(W3wp1)\Working Set<br>Process(Inetinfo)\Working Set | Size, in bytes, of the working set of the W3wp.exe or Inetinfo.exe process. These counters display the last observed value, not an average over time. |

Monitor the W3wp instance if you are running your Web server in worker process isolation mode. If you are running your server in IIS 5.0 isolation mode, monitor either the Inetinfo instance (for in-process applications) or the Dllhost instance (for out-of-process applications).

**Important**

You can run IIS6.0 in one of two request processing models: *worker process isolation mode* and *IIS 5.0 isolation mode*. IIS typically performs better while running in worker process isolation mode, although results depend on a variety of factors, including the design of your applications. Also note that each mode uses different instances of a counter to monitor performance. For more information about the two request processing models, see Running IIS 6.0 as an Application Server.

Be sure that you monitor Process\Working Set counters for all instances of W3wp.exe (the component that hosts the worker process when IIS runs in worker process isolation mode) on your server; otherwise, you might obtain an inaccurate reading of the virtual memory used by IIS.

You can use the Memory\Available Bytes counter (in Table 6.14) as an indicator of memory availability and the

Process\Working Set counters (in Table 6.18) as an indicator of the size of the IIS 6.0 working set. These counters display the last value observed, rather than an average, so be sure to examine data collected over time.

**Monitoring Disk I/O**

IIS 6.0 writes its logs to disk, so there is almost always some activity, even when requests are processed out of the cache 100 percent of the time. Under ordinary circumstances, disk activity other than that generated by logging might indicate a problem in other areas. For example, if your server has insufficient RAM, you see lots of disk activity because many hard page faults are occurring.

Under certain conditions, disk activity is normal. Examples include a large site that serves more content than your server can cache and any site that serves large amounts of static content that is not cacheable. You are also likely to notice lots of disk activity if your server hosts a database or if your users request many different Web pages.

Typically, you use the Physical Disk counters in Table 6.1 to watch for spikes in the number of disk reads when the server is busy. If the server has enough RAM, most connections result in cache requests, unless the server also hosts a database and clients are making dissimilar queries, which precludes caching.

If the server does not host any application or service that is obviously disk-intensive, but you see lots of disk activity anyway, you need to check RAM use immediately to make sure that the server has enough memory.

# Optimizing Memory Usage

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Servers running IIS 6.0 benefit from ample physical memory. Generally, the more memory that you add, the more the servers use and the better they perform. IIS 6.0 requires a minimum of 128 MB of memory; at least 256 MB is recommended. If you are running memory-intensive applications, your server might require much more memory to run optimally — more than the recommended 256 MB of memory.

Adding RAM to your system is not the only option, however. The following methods can be used to optimize memory performance without adding memory.

**Improve Data Organization**
To improve the performance of the file system cache, keep related Web files on the same logical partitions of a disk.

Use Disk Defragmenter to consolidate fragmented files and folders on the server's hard disk so that each occupies a single, contiguous space on the volume. As a result, the system can gain access to files and folders and save new ones more efficiently.

**Increase the Efficiency of High-Traffic Web Sites**
Evaluate your high-traffic Web sites to find ways to increase their efficiency. For more information about creating efficient Web sites, see Creating a More Efficient Web Site.

**Try Disk Mirroring or Striping**
The optimum configuration is to have enough physical memory to hold all static Web pages. However, if pages must be retrieved from disk, use mirroring or striping to make reading from disk sets faster. In some cases, a caching disk controller is helpful. For more information about mirrored or striped volumes, see "Using mirrored volumes" or "Using striped volumes" in Help and Support Center for Windows Server 2003.

**Add or Enlarge Paging Files**
Add paging files, or increase the size of existing ones. Windows Server 2003 creates one paging file on the system disk, but you can also create a new paging file on each logical partition of each disk. By striping a paging file across separate physical drives, you can improve paging file performance by using drives that do not contain your sites content or log files.

You can change the sizes of existing paging files. The larger the paging file, the more memory the system commits to it. Be sure that the paging file on the system drive is at least twice the size of the physical memory, so that the system can write the entire contents of RAM to disk if the system unexpectedly locks up or shuts down.

For information about resizing paging files, see "Change the size of the virtual memory paging file" in Help and Support Center for Windows Server 2003.

**Retime the IIS Object Cache**
Consider shortening the period of time that an unused object remains in the cache, or lengthening the time that a used object remains in the cache, by adding the ObjectCacheTTL entry to the registry.

| ⚠Caution |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see IIS 6.0 Registry Reference. |

The value of ObjectCacheTTL controls the static file cache by specifying the Time To Live (TTL), which sets the length of time that objects are held in cached memory. If an object in the memory cache has not been referenced for the defined period, that object is phased out of the cache.

The default value is 30 seconds. If system memory is limited or if the contents of the server are dynamic, you can use a lower TTL to prevent system memory from being used to cache a large number of volatile objects. Setting the ObjectCacheTTL entry to unlimited disables the object-cache scavenger and allows cached objects to remain in the cache until the cached object changes. Disabling the cache scavenger is useful if your server has ample system memory and your data is relatively static.

For more information about setting the time limit for the object cache, see Setting the IIS Object Cache Time Period.

**Change the Balance of the File System Cache to the IIS 6.0 Working Set**
By default, servers running Windows Server 2003 are configured to give preference to the file system cache over the working sets of processes when allocating memory (through the server property **Maximize data throughput for file sharing**). Although IIS 6.0–based servers benefit from a large file system cache, giving the file system cache preference often causes the IIS 6.0 pageable code to be written to disk, which results in lengthy processing delays. To avoid these processing delays, you can set server properties to maximize data throughput for network applications. For more information about setting the server properties to maximize throughput, see Maximizing Data Throughput for Network Applications.

**Limit Connections**
If your server does not have enough memory, limiting the number of connections on the server might help alleviate the shortage because some physical memory (about 10 KB per connection) is consumed by the data structures that the system uses to keep track of connections. For more information about limiting connections to save memory, see Limiting Connections to Manage Resources.

**Limit the Queue Length for Application Pools**
If you are running your server in worker process isolation mode, you can limit application pool queue lengths to prevent large numbers of requests from queuing up and overloading your server. When the queue length limits feature is enabled (it is enabled by default), IIS monitors the number of requests in a designated application pool queue before queuing a new request. Use IIS Manager to change the default limit of 1,000 requests for application pool request queues.

If adding a new request to the queue exceeds the maximum queue length, the server rejects the request and sends a 503 error (Server Too Busy) to the client. However, requests that are already queued remain queued even if the limit is changed to a value that is less than the current queue length.

For information about changing the queue length limit for an application pool by using IIS Manager, see Configuring Application Pool Queue-Length Limits. For information about configuring application pool queue length limits by using metabase properties, see the **AppPoolQueueLength** property in the Metabase Property Reference.

**Adjust Resource Allocation in Windows**
System processing is managed by Windows Server 2003, which can allocate processor and memory resources among tasks. If your server needs a temporary performance boost, you can temporarily adjust resource allocation in Windows to allocate the resources where you need them:

- To obtain a faster response time, you can set Windows to allocate more processor time to the program that you are currently running.

- If you want to run background programs, such as printing or disk backup, while you work, you can have Windows share processor resources equally between background and foreground programs.

For more information about changing the way that Windows allocates processor resources, see Change the performance of foreground and background programs in Help and Support Center for Windows Server 2003.

**Limit Performance Logging**
Always keep in mind that performance logging uses system resources. When you are not actively checking performance, disable performance-related logging to squeeze a bit more performance from your server.

# Preventing Processor Bottlenecks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Servers running IIS 6.0 rely on the speed and efficiency of their processors. IIS 6.0 code is multithreaded for efficient scaling on uniprocessor and multiprocessor computers, and is largely self-tuning. Even so, processor bottlenecks can occur on very active servers.

A *processor bottleneck* occurs when one or more processes consume the majority of the processor time. This forces threads that are ready to be executed to wait in a queue for processor time. Processor bottlenecks can occur on multiprocessor computers even when only a single processor is fully loaded, if the work in the queue cannot be — or is not — distributed to the other processors. Because a processor bottleneck is not centered in other components, such as memory, disks, or network connections, upgrading or adding to those components does not fix this performance problem, and can sometimes make it worse.

The processors on a server running IIS 6.0 must support the operating system and processes unrelated to Internet services, in addition to IIS 6.0 processes. The processors must also support applications related to Internet services, such as those that assemble data from Microsoft® SQL Server™ databases or generate dynamic Web pages.

To identify processor bottlenecks on the server, you need to monitor process activity, connections, and threads. After evaluating the data, determine the appropriate measures for improving processor performance through hardware upgrades, software configuration changes (limiting connections or allowing a higher thread count), or the redesigning of Web sites to reduce the processor workload.

# Identifying Processor Bottlenecks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The relationships among threads, connections, and requests are complex, especially in IIS 6.0, which uses the worker thread model rather than the simpler, and less efficient, thread-per-client model. Because of the complexity of the model, it is important to obtain a baseline from which to evaluate your system if a bottleneck occurs.

To identify performance bottlenecks, perform the following types of monitoring activity:

- Monitor processor activity by using performance counters and other tools.
- Monitor connections to determine how your server responds to managing different numbers of connections.
- Monitor threads, including context switches, interrupts, and deferred procedure calls.

# Monitoring Processor Activity

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Several tools are available for monitoring processor performance:

- Use the Web Capacity Analysis Tool (WCAT) to perform stress tests. For information about using WCAT, see Using WCAT as a Stress Client.
- To monitor process performance on servers running Windows Server 2003, you can use Task Manager, System Monitor, and Performance Logs and Alerts.

Remember that all monitoring tools use system resources. Monitor the processor use of the process in which the tool runs; then, before you analyze your data, subtract the processor time of the tool process from the data.

To collect data about processor activity, monitor the counters that are listed in Table 6.19. The % Processor Time counter is the primary indicator of processor activity, displaying the average percentage of busy time observed during the sample interval.

**Table 6.19 Counters for Monitoring Processor Activity**

| Object\Counter | Value |
|---|---|
| System\Processor Queue Length | The number of threads in the processor queue. Shows ready threads only, not threads that are running. Even multiprocessor computers have a single queue for processor time; thus, for multiprocessors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than two threads per processor is normally acceptable, depending upon the workload. |
| System\Context Switches/sec | The combined rate at which all processors on the computer are switched from one thread to another. |
| Processor (_Total)\% Processor Time | The average percentage of time that all processors are active. This counter is the primary indicator of average processor activity. This value is derived by calculating the percentage of time during the sample interval that all processors spent in executing their idle thread (which consumes cycles on a processor when no other threads are ready to run), and then subtracting the result from 100 percent. |
| Processor\% Processor Time | The average percentage of processor use for each processor (#0, #1, and so on). |
| Processor\% Privileged Time | The average percentage of processor time that is spent in privileged mode. In Windows Server 2003, only privileged mode code has direct access to hardware and memory. Application threads can be switched to privileged mode to run operating system services. |
| Processor\% User Time | The average percentage of processor time that is spent in user mode. *User mode* is a restricted processing mode designed for applications, environment subsystems, and integral subsystems. The alternative, *privileged mode*, is designed for operating system components, allowing direct access to hardware and memory. The operating system switches application threads to privileged mode to access operating system services. |
| Process\% Processor Time | The average percentage of processor use attributable to each processor, either for a particular process or for the total for all processes (shown in the list of instances). |

**Tip**

A memory bottleneck can sometimes look like a processor or disk bottleneck. If the system does not have enough physical memory to store the code and data that are needed, the processor spends substantial time paging. Before adding or upgrading processors or disks, you need to monitor the memory on your server. For more information about monitoring memory, see Monitoring Overall Server Memory.

**Analyzing Processor Activity Data**

Of the counters listed in Table 6.19, the System\Processor Queue Length counter is probably the most important for analyzing processor activity data. This counter displays the number of threads waiting to be processed in the single queue shared by all processors. Sustained high rates of processor activity, which leave little excess capacity to handle peak loads, are often associated with processor bottlenecks. Processor activity itself only indicates that the resource is being used, not that maximum use of the resource is a problem. However, a long, sustained queue indicates that threads are being kept waiting because a processor cannot handle the load assigned to it.

A sustained processor queue length of two or more threads (as indicated by the System\Processor Queue Length counter) typically indicates a processor bottleneck. You can set an alert in Performance Logs and Alerts to notify administrators when the processor queue length is unacceptably high.

The Processor\% Processor Time counter is most often used as a general measure of processor activity on both uniprocessor and multiprocessor computers. On a multiprocessor server, use this counter to reveal unequal distribution of processor load. You might have a processor bottleneck if % Processor Time numbers are high while usage of the network adapter and disk I/O remain well below capacity.

On multiprocessor computers, use the total instance of the Processor\% Processor Time counter — for example, Processor(_Total)\% Processor Time — to monitor system-wide processor use. For multiprocessor use, this counter represents the active time of all processors divided by the number of processors. If the server workload is shared equally among all processors, the total instance time is an excellent measure of processor activity.

Windows Server 2003 is designed for efficient scaling and includes several strategies for balancing processor load. An application, however, can create an imbalance by setting a processor affinity, which binds a process to a single processor. To identify processor bottlenecks related to processor affinity, chart Processor\% Processor Time for each processor on the computer.

| **Note** |
| --- |
| Though processor affinity can overload a single processor, it can improve performance by reducing the number of processor cache flushes as threads move from one processor to another. You will need to assess, through monitoring, whether the trade-off is acceptable. For information about assigning processor affinity to an application pool or Web garden, see Improving Processor Performance. |

**Optimizing Processor Activity**

If data shows a large processor queue with all processors busy, view a histogram of Process\% Processor Time for each process. To do this, click the **View Histogram** button on the toolbar in System Monitor. The histogram shows the processor time consumed by each process.

Based on the histogram readings, consider making the following adjustments to optimize performance:

- If a single bar in the histogram rises above all of the others, the process represented by the bar might be consuming a disproportionate share of processor time and causing a bottleneck. Consider replacing the application running in that process, or moving the process to another server.

- If the processors are being shared equally by several processes, consider upgrading or adding processors. (Multithreaded processes benefit most from additional processors.)

**Related Topics**

For more information about processor use by applications related to IIS, see Improving Application Performance.

# Monitoring Connections

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

It is important to determine how your server responds when it is managing different numbers of connections. After collecting data on connection trends, you can associate data about general server performance with the number of connections being served.

Each connection that an IIS 6.0 service establishes consumes processor time. The network adapter interrupts the processor to signal that a client has requested a connection. Further processing is required to establish and maintain the connection, to fulfill client requests sent over the connection, and — when the connection is closed — to delete the structures that serviced the connection. Each time a connection is established, the load on the server increases.

IIS 6.0 includes several features to optimize its handling of connections. For example, the HTTP Keep-Alives feature, which is enabled by default, maintains an open connection between the server and browser across multiple requests until an entire Web page is downloaded. This feature significantly improves bandwidth performance on most servers and improves response times for clients. For more information about HTTP Keep-Alives, see Enabling HTTP Keep-Alives to Keep Connections Open.

You can use IIS 6.0 logging to monitor the number of connections that your server makes and to track patterns of client demand for your server. Table 6.20 lists the counters that you can use with Performance Logs and Alerts to monitor connections to IIS 6.0.

Table 6.20 Performance Counters for IIS 6.0 service connections/p>

| Object\Counter | Value |
|---|---|
| Web Service\Current Connections FTP Service\Current Connections | The number of connections maintained by the service during the most recent sample interval. |
| Web Service\Maximum Connections FTP Service\Maximum Connections | The maximum level of simultaneous connections since the server started up. |

Because these counters display the last observed value, and not an average, you must log their values over time to get a reliable sample. These counters can exaggerate the number of simultaneous connections because, at any given moment, some entries may not have been deleted even though the connections on which they are based have been closed.

For more information about configuring and interpreting IIS 6.0 logs, see Analyzing Log Files.

**Analyzing Connection Data**

Identify patterns of client demand for your server by monitoring the number of connections. Log connection data at regular time intervals, and look at the number of connections served during each interval. Observe the size of the processor queue and the usage on each processor during periods of small, moderate, and large numbers of connections. This data shows how your configuration responds to each load level.

You can identify a processor bottleneck during a given time interval by observing one of the following:

- A long, sustained processor queue (with more than two threads).
- High usage rates on one or more processors.
- A curve in the graph of the Current Connections counter on any IIS 6.0 service performance object that reaches a high value and then plateaus. This pattern often indicates that additional connections are being blocked or rejected.

The Active Server Pages, Web Service, FTP Service, and SMTP Server counters collect data at the OSI application layer. If any connections are blocked, rejected, or reset between the transport and application layers, counts of TCP/IP connections might not equal the sum of HTTP, FTP, and SMTP connections. For information about monitoring connections at lower layers, see Monitoring Network Activity.

**Optimizing Connections**

To prevent processor bottlenecks, make certain that a lengthy processor queue is not forming when you serve large numbers of connections. You can usually avoid a bottleneck during peak usage by setting the connection limit to twice the average value of the Current Connections counter.

If the processor regularly develops a bottleneck when servicing large numbers of connections, consider upgrading or adding processors, or limit the maximum number of connections on the server. Limiting connections can cause the server to block or reject connections, but it helps to ensure that accepted connections are processed promptly.

For information about limiting connections, including how to configure this setting, see Limiting Connections to Manage Resources.

# Monitoring Threads

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 runs in a set of multithreaded processes designed for efficient scaling on uniprocessor and multiprocessor systems. *Threads* are the sequences of execution in each process that run the process code on the processor. In the IIS 6.0 process, there is no simple association between threads and connections or between threads and requests, nor is there an easily quantifiable relationship between the optimum number of threads in the process and the number of files served, the number of requests filled, or the number of connections maintained.

Because IIS 6.0 uses the worker thread model (rather than the simpler, thread-per-client model, which dedicates a thread to each connection or request), a pool of threads, known as the *worker threads*, is dedicated to accepting and monitoring all connections. This frees other threads to do the remaining work of the application, such as authenticating users; parsing client requests; locating, opening, and transmitting files; and managing internal data structures.

Even though you cannot associate individual threads with connections or requests, you can:

- Count the number of threads in the IIS 6.0 process.
- Measure the amount of processor time that each thread gets.
- Associate the number of threads (and processor activity) with the number of current connections, number of files served, and other measures of server activity and performance.

Several tools — including Performance Logs and Alerts, Task Manager, and System Monitor — monitor or enumerate the threads in a process. Individual threads are difficult to monitor, especially if they frequently start and stop. Monitoring threads is also costly. Be sure to monitor the overhead of the process in which your tool runs (by monitoring the Process\% Processor Time counter), and subtract this from the overall processor time for the data that you collect.

**Monitoring Threads and Context Switches**

Table 6.21 lists the counters that monitor threads and context switches. You can add to this list any counters that you use to associate numbers of threads with performance, such as Web Service\Current Connections, Web Service\Bytes/sec, or Server\Logon/sec.

**Table 6.21 Counters for Monitoring IIS 6.0 Threads**

| Object\Counter | Value |
|---|---|
| Process (W3wp1)\Thread Count | The number of threads created by the process. This counter does not indicate which threads are busy and which are idle. It displays the last observed value, not an average. |
| Thread (W3wp1:*Thread #*)\% Processor Time | The percentage of processor time that each thread of the worker process is using. |
| Thread (W3wp1/*Thread #*)\Context Switches/sec | The rate of switches from one thread to another. Thread switches can occur either inside of a single process or across processes. |
| System\Context Switches/sec | The combined rate at which all processors on the computer are switched from one thread to another. |

1 Monitor the W3wp instance if you are running your Web server in worker process isolation mode. If you are running your server in IIS 5.0 isolation mode, monitor either the Inetinfo instance (for in-process applications) or the Dllhost instance (for out-of-process applications).

Balancing threads against overall performance as measured by connections and requests can be difficult. Any time that you tune threads, follow up with overall performance monitoring to see whether performance increases or decreases.

You can chart the Process\Thread Count: W3wp value over time to see how many threads the worker process creates and how

the number of threads varies. Then, observe the processor time for each thread in the process (Thread\% Processor Time: W3wp/*Thread #*) during periods of high, medium, and low server activity (as indicated by the other performance measures).

To determine if you need to adjust the thread count, compare the number of threads and the processor time for each thread in the process to the total processor time. If the threads are constantly busy, but are not fully using the processor time, you might improve performance by allowing more threads. However, if all the threads are busy and the processors are operating at close to their maximum capacity, it is best to distribute the load across more servers instead of increasing the number of threads.

Context switch counters are provided for the System and Thread objects. If you decide to increase the maximum size of any of the thread pools, it is important to monitor the counters for context switches. *Context switches* occur when a running thread voluntarily relinquishes the processor, is preempted by a higher priority ready thread, or switches between user mode and kernel mode to use an Executive or subsystem service. Switching to the subsystem process causes one context switch in the application thread; switching back causes another context switch in the subsystem thread.

Increasing the number of threads can increase the number of context switches to the point that performance decreases rather than increases. Ten thousand or more context switches per second is high. If you see numbers in that range, consider reducing the maximum thread pool size.

You should also observe the patterns of context switches over time, which indicate that the kernel has switched the processor from one thread to another. A large number of threads is likely to increase the number of context switches. Although they allow multiple threads to share the processor, context switches also interrupt the processor and might interfere with its performance, especially on multiprocessor computers.

Controlling processor affinity can improve performance by reducing the number of processor cache flushes as threads move from one processor to another. For example, you might be able to reduce cache flushes for dedicated servers by assigning processor affinity. However be aware that dedicating a program to a particular processor might prevent other program threads from migrating to the least-busy processor.

**Monitoring Interrupts and Deferred Procedure Calls**

You also need to monitor the counters for interrupts and deferred procedure calls (DPCs). Table 6.22 lists the counters to use.

**Table 6.22 Counters to Monitor Interrupts and Deferred Procedure Calls**

| Object\Counter | Value |
|---|---|
| Processor\Interrupts/sec | The rate, in average number of interrupts in incidents per second, at which the processor received and serviced hardware interrupts. It does not include deferred procedure calls, which are counted separately. |
| Processor\% DPC Time | The percentage of time that the processor spent receiving and servicing deferred procedure calls during the sample interval. Deferred procedure calls are interrupts that run at a lower priority than standard interrupts. |

The value of the Processor\Interrupts/sec counter is an indirect indicator of the activity of devices that generate interrupts, such as the system clock, the mouse, disk drivers, data communication lines, network adapters, and other peripheral devices. These devices normally interrupt the processor when they have completed a task or require attention, suspending normal thread execution. The system clock typically interrupts the processor every 10 milliseconds, creating a background of interrupt activity.

Use the Processor\Interrupts/sec and Processor\% DPC Time counters to determine how much time the processor is spending on interrupts and deferred procedure calls. These two factors can be another source of load on the processor. Client requests can be a major cause of interrupts and deferred procedure calls. Some new network adapters include interrupt moderation, which accumulates interrupts in a buffer when the level of interrupts becomes too high.

**Optimizing Thread Values**

By default, the IIS 6.0 process creates as many as four threads per processor. IIS 6.0 continually adjusts the number of threads in its process in response to server activity. For most systems, this tuning is sufficient to maintain the optimum number of threads, but you can change the maximum number of threads per processor if your system requires this. If the threads in the IIS 6.0 process appear to be overworked or underutilized, consider these tuning strategies:

- If nearly all of the threads of the IIS 6.0 process are busy nearly all of the time, and the processors are at or near their maximum capacity, consider distributing the workload among more servers. You can also add processors.

- If nearly all threads appear busy, but the processors are not always active, consider increasing the maximum number of threads per processor. Do not increase the maximum number of threads unless you have processors with excess capacity. More threads on the same number of processors cause more interrupts and context switches, and result in less processor time per thread.

To adjust the maximum number of threads in the IIS 6.0 service process, use a registry editor to add the **MaxPoolThreads** entry to the following registry path:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Inetinfo\Parameters

| ⚠**Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see Registry Reference. |

The **MaxPoolThreads** entry is calculated in units of threads per processor. Do not set this value below 2 or above 20. You can try setting the **MaxPoolThreads** entry to a higher value if back-end latency is higher; set the entry to a lower value for applications that are CPU-intensive. Continue monitoring the system carefully to make sure that changing the number of threads allowed achieves the effect that you want.

# Improving Processor Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If your data on processor performance indicates that processor queues are developing regularly or while a server is handling large numbers of connections, start by monitoring the server's memory. Rule out a memory bottleneck or add more memory before (or in addition to) adding or upgrading processors.

In addition to the methods for optimizing thread values (discussed in Monitoring Threads), consider making the following adjustments.

**Upgrade the L2 cache**
When you add or upgrade processors, choose processors with a large secondary (L2) cache. Server applications, such as IIS, benefit from a large processor cache because their instruction paths involve many different components. A large processor cache (2 MB or more if external, up to the maximum available if on the CPU chip) is recommended to improve performance on active servers running IIS 6.0.

**Add network adapters**
If you are administering a multiprocessor system that does not distribute interrupts symmetrically, you can improve the distribution of the processor workload by adding network adapters so that there is one adapter for every processor. Generally, you only add adapters when you need to improve the throughput of your system. Network adapters, like any additional hardware, have some intrinsic overhead. However, if one of the processors is nearly always active (that is, if Processor\% Processor Time = 100) and more than half of its time is spent servicing deferred procedure calls (if Processor\% DPC Time > 50), then adding an adapter is likely to improve system performance, as long as the available network bandwidth is not already saturated.

**Upgrade network adapters**
If you are adding or upgrading network adapters, choose those with drivers that support interrupt moderation cards, which can perform transfers in the background without processor help. Interrupt moderation prevents the processor from being overwhelmed by bursts of interrupts. Consult the driver manufacturer for details. Also, if at all possible, use adapters that support TCP checksum offloading.

**Upgrade to a multiprocessor computer**
A server-application environment, such as e-commerce sites running COM+ applications, is much more processor-intensive than a typical file or print server environment because much more processing is taking place at the server. For that reason, high-end multiprocessor computers can increase performance in server-application environments. IIS 6.0 introduces a new kernel-mode driver for HTTP parsing and caching that is tuned to increase Web server throughput and scalability of multiprocessor computers.

**Assign processor affinity to better use CPU caching**
To efficiently use CPU caches on multiprocessor servers, you can configure application pools to establish affinity between worker processes and the processors. *Processor affinity* forces worker processes to run on a specific microprocessor or CPU, and applies to all worker processes serving a particular application pool. You can also use processor affinity with Web gardens that run on multiprocessor computers where clusters of CPUs are dedicated to specific application pools.

Controlling processor affinity can improve performance by reducing the number of processor cache flushes as threads move from one processor to another. For example, you might be able to reduce cache flushes for dedicated servers by assigning processor affinity. However, the performance trade-off is that dedicating a program to a particular processor might prevent other program threads from migrating to the least-busy processor.

If you want to assign a particular process or program to a single processor to improve its performance — possibly at the expense of other processes — configure the SMPAffinitized Metabase Property and SMPProcessorAffinityMask Metabase Property. For more information about assigning processor affinity, including how to configure these metabase properties, see Running IIS 6.0 as an Application Server.

**Limit connections**
If you cannot upgrade or add processors, consider reducing the maximum number of connections that each IIS 6.0 service accepts. Limiting connections can result in connections being blocked or rejected, but it helps ensure that accepted connections are processed promptly.

**Adjust the maximum number of threads**
IIS 6.0 tunes the number of threads in its process dynamically. The dynamic values are usually optimal; in extreme cases of very

active or underused processors, however, it might help to adjust the maximum number of threads. If you change the maximum number of threads, continue careful testing to make sure that the change has improved performance. The difference is usually quite subtle.

**Redesign the Web site**

You can improve performance and reduce the processor workload by optimizing database use, calling compiled components from scripts in ASP pages, using Internet Server API (ISAPI) instead of ASP (and ASP or ISAPI instead of CGI applications), and eliminating large bitmapped images or optimizing them to reduce file size. For more information about optimizing Web site performance, see Improving Application Performance.

# Improving Application Performance

Updated: November 5, 2009

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Web applications continue to increase in popularity, constituting an ever larger proportion of the average Web server file base. The challenge for administrators is to preserve speed and efficiency while serving complex applications. By monitoring and tuning different types of applications, in addition to static HTML pages, you can estimate application overhead during periods of varying activity to ensure that your server is prepared for the workload.

Use the tools and counters that are included with IIS 6.0, such as those discussed here, to help you optimize application performance:

- Testing Applications with Stress Tools. Use a stress client like the Web Capacity Analysis Tool (WCAT) to test your applications under a variety of workloads.

- Monitoring and Tuning Applications in IIS 6.0. When you enable ASP.NET and ASP pages on a Web server, IIS provides application-specific counters for monitoring and tuning their performance.

# Testing Applications with Stress Tools

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If Web applications are an important part of your site, the performance of those applications is critical in determining the capacity of the site. In turn, testing is the only way to find out the capacity and performance of a Web application.

If you are unfamiliar with stress testing or need a refresher, look at the following sections:

- Measuring performance by using a stress client. Provides an introduction to performing stress testing.
- Using WCAT as a stress client. Details how to obtain the WCAT tool and its documentation. WCAT is a stress client provided in the IIS 6.0 Resource Kit Tools.
- Estimating baseline performance for applications. Compares the relative performance capabilities of specific application types to help you proactively develop applications that provide optimal performance.

# Measuring Performance by Using a Stress Client

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use a stress client to test different server and network configurations under a variety of workloads. For example, use a stress client to test how your server responds to different workloads or to test the same workload with varying configurations of the server. Some stress clients include prepared test workloads; optionally, you can create your own workloads.

A stress client simulates clients and servers communicating over a network, and usually requires at least four primary subsystems:

- At least one computer that simulates a client, which runs one or more virtual clients.
- One computer that acts as a server.
- One computer, called a *controller*, which initiates and monitors the test.
- A network.

Both the client and controller components can run on one computer. (You can also run the server component on the same computer, but doing so might provide inaccurate results.)

To produce a more realistic test of your server, link it to client computers that are powerful enough to drive the server. The processors in the client computers must be fast enough to prevent a client bottleneck when accessing the server computer. If the client processors are slower than the server, you might need to link more client computers to each server.

Stress clients usually work best if the network that connects the computers has little or no traffic that is unrelated to the test. If possible, use a link dedicated only to the test, and run the tests on a 100 megabits per second — or faster — network.

Test different performance features to observe the effect they have on your server. To test a feature, first run a stress client test with the feature, then run the same test without it. Be sure to run the with feature and without feature versions of the test on varying workloads and multiple clients. For example, use a stress client to test the following features:

- Set different connection limits, such as HTTP Keep-Alives and connection timeouts.
- Test several compression combinations. For example, compress only static files, only dynamic files, or both. Compress at different compression levels for both file types. Also, check to see that the correct set of compressible file types is included in your compression configuration.
- Test bandwidth throttling at different times and on different services.
- Test different hardware combinations, including multiprocessor machines with and without processor affinity.
- Compare the benefits of running different health-monitoring features, such as rapid-fail protection and worker process recycling.
- Test security-related features, such as SSL.

Repeat each test several times, and average the results to eliminate the effects of unintended variations in the test conditions. Then compare the results of the with feature and without feature tests. Consistent differences in the results of the tests are likely to indicate the overhead associated with the feature. You can use these results for planning configuration changes.

# Using WCAT as a Stress Client

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Web Capacity Analysis Tool (WCAT) is a stress client available to IIS users. It is a script-driven command-line tool. WCAT includes a stress client tool (Wcclient.exe) to simulate the clients and a controller tool (Wcctl.exe) to initiate and monitor the testing.

Downloading the IIS 6.0 Resource Kit Tools performs a default installation of the WCAT client and controller files in the *LocalDrive*:\Program Files\IIS Resources folder.

A WCAT test begins with the controller locating and interpreting input files. A controller input file includes three required sections (Configuration, Script, and Distribution) and one optional section (Performance, for specifying which performance counters to monitor). The controller uses the input files to send instructions to the clients, including which tests to run. WCAT includes a prepared input file (Sample.ubr, in the *LocalDrive*:\Program Files\IIS Resources\WCAT Controller\Scenarios folder) that contains a test for stressing the default IIS Web page.

When a WCAT test is complete, the controller collects the test results and writes them to output files. Output files include a log file, which you can optionally configure as an XML-formatted file, and an optional performance results file, which contains data collected from performance counters on the server.

If you specify performance counters in the Performance section of the .ubr file, WCAT logs the counter data in its usual report but also produces a performance log file in a tab-separated format, so that you can import the logged counter data into a spreadsheet or data-processing application. The log of performance data that WCAT generates is stored, by default, in the WCAT Controller\Scenarios folder. You can open this log, which is a text file, by using any word processor.

# Estimating Baseline Performance for Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you develop a new application, it is useful to understand the performance capabilities of different types of Web applications. Table 6.23 provides a guide to the relative performance of different types of applications, where *n* represents the fastest application, and all other application types are rated relative to its performance.

**Table 6.23 Relative Performance of Different Application Types**

| Application Type | Relative Cost | Comments |
|---|---|---|
| Static page (kernel mode) | *n* | Fastest and simplest Web page. |
| Static page (user mode) | 2*n* | |
| ISAPI | 2*n* | Relative speed for a simple ISAPI application. |
| ASP.NET | 5*n* | Relative speed for a simple ASP.NET application. |
| ASP | 10*n* | Relative speed for a simple ASP application. |
| CGI | 100*n* | Slowest. |

Microsoft® Visual Studio® .NET provides Application Center Test (ACT) to help you test and monitor the performance of your Web application. ACT simulates a large group of users by opening multiple connections to a Web server and rapidly sending HTTP requests. ACT is designed to stress-test Web servers and to analyze performance and scalability problems with Web applications. For more information about ACT, see the Microsoft Application Center Test 1.0, Visual Studio .NET Edition documentation on MSDN.

**Related Topics**

For more information about Web application development, including ASP.NET and ASP applications, see Web Application Guide.

# Monitoring and Tuning Applications in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

To ensure that all of your Web applications run reliably and efficiently during peak hours as well as normal usage periods, you should monitor and tune your applications, obtaining baseline data to compare the performance of different types of applications. If you uncover applications or groups of applications that appear to run inefficiently, review the suggestions for improving application performance in Optimizing Application Design and Administration.

You can also improve application performance by using the connection timeout and HTTP compression features that IIS provides. Certain applications, such as ASP pages, provide metabase properties for customizing these features. For more information about using these features to improve application performance, see Administering Network Resources.

When you make changes to your applications or to the hardware on which your applications run, be sure to test the changes in a nonproduction environment before implementing them on production servers.

IIS provides application-specific counters that you can monitor to obtain performance data about different types of applications that run on your servers.

- Monitor applications that use the WWW service by observing counters that monitor ISAPI extensions, CGI requests, and so forth.
- Monitor ASP pages, and tune metabase properties for ASP.
- Monitor ASP.NET applications.

| **Note** |
| --- |
| If after installing Windows Server 2003 SP1 you see a decrease in performance, and investigation reveals a decrease in the CGI Requests/sec processed from what was typical before installing SP1, see Configuring the Maximum Number of Concurrent CGI Requests for information on increasing the value of MaxConcurrentCgisExecuting. |

# Monitoring Applications That Use the IIS 6.0 WWW Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Monitor the counters in Table 6.24 to determine how the WWW service is processing ISAPI, CGI, POST, and GET requests.

**Table 6.24 Counters That Provide Data About WWW Service Applications**

| Object\Counter | Value |
|---|---|
| Web Service\ISAPI Extension Requests/sec<br><br>Web Service\CGI Requests/sec | The rates at which the server is processing ISAPI and CGI application requests. If these values decrease due to increasing loads, you might need to redesign the applications.<br><br>**Note**: ASP is an ISAPI extension; therefore ASP requests are included in the first counter. |
| Web Service\Get Requests/sec<br><br>Web Service\Post Requests/sec | The rate at which these two common HTTP request types are made to the server. POST requests are generally used for forms and are sent to ISAPIs (including ASP) or CGIs. GET requests make up almost all other requests from browsers and include requests for static files, ASPs and other ISAPIs, and CGI requests. |

# Monitoring ASP Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you are running ASP pages on your server, the counters listed in Table 6.25 can help you determine how well the applications are performing. These counters provide the sum of the ASP performance counters for all processes hosting ASP; the data cannot be broken down by process. For information about monitoring the performance of your ASP.NET applications, see Monitoring ASP.NET Performance.

**Table 6.25 Counters Used to Monitor ASP Performance**

| Object\Counter | Value |
|---|---|
| Active Server Pages\Requests/sec | The number of ASP responses during the past second. This counter excludes requests for static files or other dynamic content, and can fluctuate considerably based on the complexity of the ASP pages and the capacity of your Web server. |
| Active Server Pages\Requests Executing | The number of ASP requests currently executing (for example, the number of active worker threads). |
| Active Server Pages\Request Wait Time | The number of milliseconds that the most recent ASP request was waiting in the queue. |
| Active Server Pages\Request Execution Time | The number of milliseconds that the most recent ASP request took to execute. This value can be somewhat misleading because it is not an average. |
| Active Server Pages\Requests Queued | The number of queued ASP requests that are waiting to be processed. The maximum number for this counter is determined by the metabase property **AspRequestQueueMax**. |

**Analyzing the Data**

The Active Server Pages\Requests/sec counter excludes requests for static files or other dynamic content, and can fluctuate considerably based on the complexity of the ASP pages and the capacity of your Web server. If this counter is low during spikes in traffic on your server, your applications might be causing a bottleneck.

It is difficult to give an ideal value for the Active Server Pages\Requests Executing counter. Some factors to consider include:

- If pages execute quickly and do not wait for I/O (for example, loading a file or making a database query), the number of ASP requests that are currently executing is likely to be low — little more than the number of processors when the computer is busy.

- If pages must wait for I/O, the number of pages executing is likely to be higher —up to the maximum thread count allowed for an ASP processor (set by the **AspProcessorThreadMax** metabase property) multiplied by the number of processors, multiplied by the number of processes hosting ASP. (The default **AspProcessorThreadMax** value is 25.)

- If the Requests Executing counter value equals the value of the **AspProcessorThreadMax** property, if Requests Queued is large, and if CPU usage is low, you might need to increase the value of the **AspProcessorThreadMax** metabase property.

| **Important** |
|---|
| The value of the AspProcessorThreadMax Metabase Property can significantly influence the scalability of your Web applications and the performance of your server. Before changing the setting, be sure to read about this property and its use. |

The Request Execution Time counter can be misleading because it is not an average. For example, if you regularly receive 30 requests for a page that executes in 10 milliseconds (ms) to every one request for a page that executes in 500 ms, the counter is likely to indicate 10 ms even though the average execution time is over 25 ms.

Ideally, the Requests Queued and Request Wait Time counters remain close to zero; however, they can vary under varying loads (depending on the application) while the server still provides good performance. If these numbers vary too much from their expected range, you might need to rewrite your ASP applications. If the queue size and wait time increase, or if the queue limit is reached, browsers display the message HTTP 500/Server Too Busy.

**Related Topics**

For more information about ASP performance counters, including additional ASP counters, see IIS 6.0 Performance Counters.

# Tuning ASP Metabase Settings

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If your Web server is running ASP applications, you can help optimize server performance by tuning the metabase properties described in this section. You can retrieve and change these properties by using WMI or ADSI, or by directly editing the metabase with a text editor. Changes to most of these properties do not take effect until you restart the WWW service.

**AppAllowDebugging**
The AppAllowDebugging Metabase Property specifies whether ASP debugging is enabled on the server. When this property is enabled, IIS application threads are serialized; only one thread can execute for each application at one time. Serialized threads adversely affect Web server performance. You should retain the default setting for this property, which is **false,** on all production servers. Set this property to **true** only on a test server when you are debugging an application during its deployment phase.

**AspBufferingOn**
The AspBufferingOn Metabase Property turns on buffering in ASP, so that all output from an application is collected in the ASP output buffer before the buffer is flushed to the client browser (default behavior). If this property is set to **false**, output from ASP scripts is written to the client browser as the client browser becomes available. Set this property to **true** on all production Web servers.

**AspProcessorThreadMax**
The AspProcessorThreadMax Metabase Property specifies the maximum number of worker threads per processor that IIS can create. To find out the maximum number of worker threads that IIS allows per ASP process, multiply this value by the number of processors on your server. If you decrease this value, monitor performance to make sure that the lower thread limit does not degrade performance. If it does, increase the value again.

**AspQueueConnectionTestTime**
The AspQueueConnectionTestTime Metabase Property specifies the maximum time, in seconds, that a request remains queued before ASP checks to determine whether a client is still connected. If a request is in the queue longer than the queue connection test time, the server checks whether the client is still connected before beginning execution.

This feature can help you handle impatient users who fill up the request queue with numerous attempts at the same page. Often, users do not wait more than a few seconds for ASP pages to process. Although the maximum waiting time varies from user to user, the generally accepted maximum is about 10 seconds; so set a small value — for example, 3 seconds.

This property is useful for making ASP processing efficient only up to the point at which ASP begins to process the script. After the script is running, your application should continue to check for client connection at appropriate times, by using the **IsClientConnected** method of the ASP built-in **Response** object. For long-running pages, make cautious use of **Response.Flush** to ensure that users perceive that the page is still alive and doing productive work.

**AspRequestQueueMax**
The AspRequestQueueMax Metabase Property specifies the maximum number of ASP requests that are permitted in a queue. The default value is 3,000, but the optimal setting depends on the behavior of the application. If the execution time of the request is very short and the time in the queue is short, it is reasonable to decrease this value.

**AspScriptEngineCacheMax**
The AspScriptEngineCacheMax Metabase Property specifies the maximum number of scripting engines that ASP pages keep cached in memory (excludes currently running script engines). The default value is 250. Adjust the value according to the type of content in the application. If you are serving thousands of unique pages, server performance might improve if you increase the cache size so that the most frequently requested pages are accessed easily. A hit in the script engine cache means that you can avoid recompiling the template into byte code.

**AspScriptFileCacheSize**
The AspScriptFileCacheSize Metabase Property specifies the number of precompiled script files that can be stored in the ASP template cache. If you set the property to 4,294,967,295 (unlimited), all requested script files are cached. The default value is 500. Increase the value if you have many ASP pages. Do not set the value of this property to zero, which turns off all ASP caching and can severely impair your servers performance.

**AspDiskTemplateCacheDirectory**
The AspDiskTemplateCacheDirectory Metabase Property specifies the name of the directory where ASP stores compiled ASP templates to disk when the in-memory cache overflows. If you change the default directory, you must ensure that the identity

of the processes running Asp.dll is assigned Full Control permission to the new directory. Typically, the identity of the worker processes running Asp.dll is the LocalSystem or Network Service account. You must specify a local path; ASP template caching fails if this property is set to a Universal Naming Convention (UNC) or mapped network path.

**AspSessionTimeout and AspSessionMax**

The AspSessionTimeout Metabase Property specifies, in minutes, the default amount of time that a session object is maintained after the last request associated with the object is received. The AspSessionMax Metabase Property specifies the maximum number of concurrent sessions that IIS permits. For applications that take advantage of sessions, you might want to decrease the session timeout in order to reduce server overhead. If decreasing the session timeout causes the number of concurrent sessions to increase too much, you might need to set a lower session maximum.

**AspTrackThreadingModel**

The AspTrackThreadingModel Metabase Property specifies whether IIS checks the threading model of any components that your application creates. ASP no longer uses this metabase property, although it is still in the metabase to provide compatibility with existing administrative scripts. It is recommended that you retain this property's default value, which is **false,** to avoid the overhead incurred by the ASP track threading model and to help improve performance of your ASP application.

**CacheISAPI**

The CacheISAPI Metabase Property indicates whether ISAPI extensions are cached in memory after first use. If the value of this property is **true** (the default setting), ISAPI extensions, after they are loaded, remain in the cache until the server is stopped. If the value is **false**, extensions are unloaded from memory after the ISAPI extension is no longer in use.

ISAPI extensions are cached or not cached based on the value of this property when they are loaded into memory. Thus, if this property is changed after an extension is loaded and cached, the change has no effect on that extension until the Web server is restarted.

Although you might temporarily set this property to **false** to help with debugging, you need to keep this value set to **true** on all production Web servers. Because the ASP.dll is an ISAPI extension, disabling this property — which requires that every request be reloaded — degrades performance.

# Monitoring ASP.NET Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP.NET includes a number of performance counters that you can use to track the execution of your Web applications. When you do steady-state running analysis, you might want to ignore the first request and also any one-time initialization costs for objects. For example, when an ASP.NET page is requested for the first time, there is always the cost of compiling an instance of the **Page** class. After the initial cost is incurred, performance improves for subsequent requests.

| **Note** |
| --- |
| You can also use the built-in ASP.NET tracing feature to track code execution for a page or an application. |

Two performance objects are defined for ASP.NET, supporting two types of performance counters:

- The ASP.NET performance object contains system counters that monitor events on the state server. These global counters are not bound to a particular application instance.
- The ASP.NET Applications performance object contains application counters.

For example, there is a significant difference between the State Server Sessions counters for the ASP.NET performance object, which apply only to the computer on which the state server is running, and the Sessions counters for the ASP.NET Applications performance object, which apply only to user sessions that occur in process.

Unlike the ASP performance counters discussed earlier, which are exposed globally for the entire server computer, most of the ASP.NET performance counters are exposed per application.

Use System Monitor to view the counters for ASP.NET. If your Web server serves multiple applications, you need to specify a particular application instance when selecting a counter to monitor. In addition, System Monitor provides a **__Total__** application instance, which aggregates the counter values for all applications on a server.

**Useful Counters for Monitoring ASP.NET Applications**

When monitoring the performance of your ASP.NET Web applications, try tracking the performance counters described in Table 6.26.

**Table 6.26 ASP.NET Performance Counters**

| Object\Counter | Value |
| --- | --- |
| ASP.NET\Application Restarts | The aggregate number of restarts for all ASP.NET applications during the Web servers lifetime. |
| ASP.NET\Requests Queued | The number of ASP.NET requests waiting to be processed. |
| ASP.NET\Worker Process Restarts | The number of times that a worker process has restarted on the computer. |
| ASP.NET Applications\Errors Total | The total number of errors that have occurred in ASP.NET applications. |
| ASP.NET Applications\Requests/Sec | The number of requests executed per second for ASP.NET applications. |

Application restarts, as reported by the ASP.NET\Application Restarts counter, are incremented with each **Application_OnEnd** event. An application restart can occur because changes were made to the Web.config file or to assemblies stored in the applications \Bin directory, or because too many changes occurred in Web Forms pages. A sudden increase in this counter can mean that your Web application is shutting down. If an unexpected increase occurs, be sure to investigate it promptly. This value is reset every time that IIS is restarted.

| **Note** |
| --- |
| The value of each ASP.NET performance counter is updated every 400milliseconds. |

**Useful Counters for Troubleshooting ASP.NET Performance Problems**

Table 6.27 describes some additional performance counters that you can use to troubleshoot performance problems with your Web applications.

**Table 6.27 Performance Counters for Troubleshooting ASP.NET Applications**

| Object\Counter | Indicates |
|---|---|
| ASP.NET Applications\Pipeline Instance Count | The number of active pipeline instances. |
| .NET CLR Exceptions\# of Exceps Thrown | The total number of exceptions thrown since the start of the application. Collects data both on .NET exceptions and on unmanaged exceptions that are converted to .NET exceptions. For example, a null pointer reference exception in unmanaged code is rethrown in managed code as a .NET SystemNullReferenceException. Exceptions that are rethrown are recounted. This counter collects data from both handled and unhandled exceptions. |
| System\Context Switches/sec | The rate at which thread contexts are switched by all CPUs in the Web server. A high number usually indicates either high contention for locks or many switches between user and kernel mode by the thread. Be sure to investigate these problems by using sampling profiles and other tools. |

Although exceptions occur infrequently, you might want to monitor the # of Exceps Thrown counter, because exceptions can have performance implications. Some code paths, however, rely on exceptions for their proper functioning. For example, the **HttpResponse.Redirect** method always throws an exception, **ThreadAbortException**. Therefore, it can be more useful to track the number of exceptions thrown by using the Errors Total counter to see if the exception generated an error for the application.

**Related Topics**

For more information about the ASP.NET performance counters discussed in this section, as well as other ASP.NET performance counters, see IIS 6.0 Performance Counters.

# Balancing Performance and Security

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Balancing performance with concerns about the security of your Web applications can be an important issue, particularly if you run an e-commerce Web site. Secure Web communications require more resources than do Web communications that are not secure, so you need to know when to use security techniques such as SSL certificates or any of the Windows authentication methods. Because SSL uses complex encryption, and encryption requires considerable processor resources, it takes much longer to retrieve and send data from SSL-enabled directories. Therefore, you should place only those files that contain or receive sensitive information in your SSL-enabled directory, keeping pages free of resource-consuming elements such as images.

If you use SSL, be aware that establishing the initial connection takes five times longer than reconnecting by using security information in the SSL session cache. The default time-out interval for the SSL session cache is 10 hours in Windows Server 2003. After secure data is deleted from the cache, the client and server must establish a new connection.

# Configuring ServerCacheTime for SSL Sessions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you plan to support long SSL sessions, consider lengthening the SSL cache time-out interval by adding the ServerCacheTime entry to the registry. If you expect thousands of users to connect to your site by using SSL, estimate how long you expect SSL sessions to last, and then set the value of the ServerCacheTime entry to a number slightly higher than your estimate. Do not set the value much higher than your estimate, because the resulting time-out interval might cause your server to retain stale data in the cache.

| **Caution** |
| --- |
| The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference. |

One reason for changing the default value for the SSL session cache is to force the client to authenticate more often. More frequent caching is sometimes useful, for example, if you know that the client is using a smart card and you want the Web page to be accessible only when the user inserts the smart card in the reader.

Before changing the SSL cache time-out interval, make sure that HTTP Keep-Alives are enabled (HTTP Keep-Alives are enabled by default). SSL sessions do not expire when you use them with HTTP Keep-Alives except when the browser closes the connection.

For information about adding the ServerCacheTime entry to the registry, see Setting the SSL Cache Time-out Interval. For more information about configuring SSL session caching, including how to set the ClientCacheTime registry setting, see Knowledge Base article 247658, How to Configure Secure Sockets Layer Server and Client Cache Elements.

**Related Topics**

For more information about editing the registry, see the Registry Reference.

# Testing How Security Features Affect Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The security services in Windows Server 2003 and IIS 6.0 are integrated into a number of operating system services. This means that you cannot monitor security features separately from other aspects of those services. Instead, the most common way to measure security overhead is to run tests that compare server performance with and without a security feature. Run the tests with fixed workloads and a fixed server configuration so that the security feature is the only variable.

In your tests of the effects of security services on performance, measure the following:

- Processor activity and the processor queue. Security features like authentication, certificates, encryption, and the SSL protocol require significant processing; so it is important to balance the performance costs associated with these security features with your performance objectives.

  When using security features, you are likely to see increased processor activity, in both privileged mode and user mode, and an increase in the rate of context switches and interrupts. If the processors on the server are not sufficient to handle the increased load, queues are likely to form.

  If the SSL protocol is being used, Lsass.exe might consume an unexpected amount of CPU resources because SSL processing occurs in the Lsass.exe process. Custom hardware, such as cryptographic accelerators that perform encryption, can help. When you create certificates by using the certificate wizards provided in IIS, you can select a cryptographic service provider that uses cryptographic accelerators.

- Physical memory used. Security features require that the system store and retrieve more user information.

- Network traffic. You are likely to see an increase in traffic between the IIS-based server and the domain controller that is used for authenticating logon passwords and verifying IP addresses.

- Latency and delays. The most obvious performance degradation resulting from complex security features like SSL is the time and effort involved in encryption and decryption, both of which use many processor cycles. Downloading files from servers by using the SSL protocol can be 10 to 100 times slower than downloading from servers that are not using SSL.

| **Important** |
| --- |
| It is recommended that you do not use a domain controller as a Web server. If a domain controller is running IIS, the proportion of processor use, memory, and network and disk activity consumed by domain services is likely to increase significantly. The increased activity can be enough to prevent IIS services from running efficiently. |

**Related Topics**

For more information about IIS certificate wizards, see Using Certificate Wizards.

# Measuring Security Overhead by Using a Stress Client

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To test a security feature like SSL, first run a stress client test with the security feature; then run the same test without it. Be sure to run the "with feature" and "without feature" versions of the test on varying workloads and multiple clients.

To eliminate the effects of unintended variations of the test conditions, repeat each test several times and average the results. Then, compare the results of the "with feature" and "without feature" tests. Consistent differences in the results of the tests are likely to indicate the overhead associated with the security feature. You can use these results for planning configuration changes to handle the security overhead.

For more information about testing with a stress client, see Testing Applications with Stress Tools.

WCAT is a particularly useful tool for determining security overhead. Run IIS 6.0 logging together with WCAT to count the number of logons and file accesses. In addition, you can use System Monitor, which includes a set of counters for monitoring user authentication.

**Related Topics**

For more information about obtaining and using WCAT, see Using WCAT as a Stress Client.

# Tracking Anonymous and Nonanonymous Requests

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Web Service and FTP Service performance objects include counters that report the number of anonymous and nonanonymous requests that are sent to each of these IIS 6.0 services. (The term *nonanonymous* is used instead of *authenticated* to account for custom authentication schemes that require data from the client other than, or in addition to, the user name and password, as is also the case with authentication.)

By themselves, these counters help you determine the number and proportion of each type of request. You can also use the counter values to estimate the effect of changing how you handle anonymous and nonanonymous users. For example, if most connections are anonymous, prohibiting anonymous requests has a more significant impact than if most requests are nonanonymous.

Analysis of data from these counters, together with general measures of server performance — such as the use and allocation of processor time, the size of processor queues, memory use, the number of disk reads and writes, and throughput — is even more useful. By analyzing the combined data, you can associate the effects of varying numbers and proportions of anonymous and nonanonymous users with changes in the performance of system components.

The counters that report the number of anonymous and nonanonymous requests are called Current Anonymous Users and Current NonAnonymous Users. In IIS 5.0, these counters actually displayed connections. In IIS 6.0, these counters count each request as a unique user.

Table 6.29 describes the counters for monitoring anonymous and nonanonymous requests. These counters are associated with the Web Service and FTP Service performance objects. The Anonymous Users and Nonanonymous Users counters display the number of anonymous and nonanonymous requests to the IIS 6.0 service when the values were last observed. They do not report averages or rates. These counters can exaggerate the number of requests because closed requests might not yet have been deleted when the counter is displayed.

**Table 6.29 Counters for Anonymous and Nonanonymous Users**

| Object\Counter | Indicates |
|---|---|
| Web Service\Anonymous Users/sec<br>Web Service\NonAnonymous Users/sec | The number of requests per second that anonymous or nonanonymous users sent to the IIS 6.0 service. (These counters provide an instantaneous value rather than an average.) Note that IIS 6.0 counts each request as a new user. |
| Web Service\Current Anonymous Users<br>Web Service\Current NonAnonymous Users<br>FTP Service\Current Anonymous Users<br>FTP Service\Current NonAnonymous Users | The number of requests from anonymous or nonanonymous users that the server is currently processing. Note that IIS 6.0 counts each request as a new user. |
| Web Service\Maximum Anonymous Users<br>Web Service\Maximum NonAnonymous Users<br>FTP Service\Maximum Anonymous Users<br>FTP Service\Maximum NonAnonymous Users | The maximum number of anonymous or nonanonymous users that made concurrent requests to the service since the service was started. Note that IIS 6.0 counts each request as a new user. |

| | |
|---|---|
| Web Service\Total Anonymous Users<br><br>Web Service\Total NonAnonymous Users<br><br>FTP Service\Total Anonymous Users<br><br>FTP Service\Total NonAnonymous Users | The total number of requests that anonymous or nonanonymous users sent to the service since the service was started. Note that IIS 6.0 counts each request as a new user. |
| Web Service\Total Logon Attempts<br><br>FTP Service\Total Logon Attempts | The total number of requests that both anonymous and nonanonymous users sent to the service since the service was started. |

**Related Topics**

If you prefer to use SNMP services to monitor the WWW and FTP services, IIS provides a set of SNMP counters that monitor anonymous and nonanonymous users. For more information about the SNMP Web counters, see IIS 6.0 Performance Counters.

# Watching 404 Not-Found Errors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Web Service performance object includes a counter that displays not-found errors. Not-found errors result from client requests that were not satisfied because they included a reference to a Web page or a file that does not exist. These errors are identified by their HTTP status code number, which is 404. Many not-found errors occur because Web pages and files have been deleted or moved to another location.

You can use the Web Service\Not Found Errors/sec counter, described in Table 6.30, to track the rate at which not-found errors are occurring on your server. Alternatively, set an alert in Performance Logs and Alerts to notify the administrator when the rate of not-found errors exceeds a threshold.

**Table 6.30 Counters for Not-Found Errors**

| Object\Counter | Indicates |
|---|---|
| Web Service\Not Found Errors/sec | The number of client read requests that were not satisfied because the URL did not point to a valid file. |

An increase in not-found errors might indicate that a file has been moved without its link being updated. However, it can also indicate failed attempts to access protected documents, such as user lists and file directories.

An increase in not-found errors might also be related to IIS 6.0 security measures. When you perform a clean installation of IIS 6.0, the default settings ensure that IIS is installed in a secure, locked mode. This default installation serves only static HTTP content, such as .htm files; the request handlers for dynamic content, such as ISAPI and CGI, are disabled.

IIS 6.0 provides new substatus error codes that can help you determine why an HTTP 404 error (file or directory not found) occurred. For example, if a request is denied because an ISAPI extension or CGI is locked, a 404.2 error is returned, and a 404.3 substatus code indicates a MIME map restriction policy.

**Related Topics**

For more information about analyzing substatus codes, see Troubleshooting IIS 6.0. For more information about unlocking dynamic content, see Common IIS 6.0 Administrative Tasks.

# Analyzing Security Data and Planning Upgrades

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you collect data to assess the effects of adding security features to your server configuration, use the results of your analysis to plan configuration changes that can help you handle the additional workload that is required to support security features. Consider the following options.

**Upgrade or add processors**
Security features in general and the SSL protocol in particular are usually processor-intensive. Because Windows Server 2003 security features are multithreaded, they can run simultaneously on multiple processors. Thus, adding processors can improve performance significantly and prevent the processors from becoming a bottleneck.

For best results, choose processors with large secondary (L2) cache space (up to 2 MB). When a security feature like SSL encrypts and decrypts data, much of the processors time goes to reading and writing small units of data to and from main memory. If your server can store this data in the processor cache, the data is retrieved much faster.

**Add memory**
If security features cause increased paging or shortages in virtual memory, adding more memory can help. The physical memory used to support the security service consumes space that can be used to cache files. To accommodate peak use, allow for twice as much memory as required during times of average use while still maintaining 10 MB of available memory.

| **Important** |
| --- |
| Avoid adding disk space for security-related activity. Do not add disk space solely to support security features. Any increased disk activity associated with security features is likely to result from a shortage of physical memory, not from an actual need for more disk space. Security features such as the SSL protocol rely primarily on processors and physical memory, as opposed to disk space. |

**Use custom hardware**
Custom hardware, such as cryptoaccelerator cards, can make a significant difference in the way that your server handles security overhead. Be sure to thoroughly test any nonstandard hardware to make sure that it is compatible with the other hardware and software that you use.

**Related Topics**

- For more information about IIS security features, see Configuring Advanced Security and Customization.
- For more information about setting up certificates and keys, see Configuring SSL on a Web Server or Web Site.
- For more information about enabling HTTP keep-alives, see Enabling HTTP Keep-Alives to Keep Connections Open.

# Optimizing Performance Through Design

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to monitoring and tuning your server, you can optimize performance by improving the design and administration of your Web applications, by optimizing Web page design to create efficient Web pages, and by partnering with other groups in your organization who share the same network.

# Optimizing Application Design and Administration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can often improve the performance of your Web server by paying attention to the design and administration of your Web applications and by upgrading hardware components to support fast application response times.

**Upgrading Hardware to Improve Application Performance**
Several types of hardware upgrade can significantly improve application response times on your Web servers:

- Upgrade or add processors. Web applications benefit from faster processors.

- Add memory. Adding memory can improve application response times if applications are running within their own processes.

**Optimizing Application Design**
Check this list for general application design guidelines that can optimize performance.

- Use static files wherever possible. Typically, static files require much less processor time than do files with dynamic content.

  If you are generating pages dynamically to satisfy user preferences, consider substituting as many as 10 or 20 static variations for a single dynamic page. If you are generating dynamic pages to provide frequently updated data, redesign your application so that it generates a single static page on a fixed schedule and stores that page for retrieval until the next update.

- Use managed code (C# or Microsoft® Visual Basic® .NET), Microsoft ASP.NET, or ISAPI applications. C# and Visual Basic .NET applications usually run faster than ASP applications because the run-time environment executes compiled code. All of these types of applications typically run faster than equivalent CGI applications.

- Enable ISAPI and ASP.NET output caching where possible, which allows HTTP.sys to cache the response and serve requests for the applications from the kernel instead of transitioning to user mode. For more information about tuning ISAPI output caching, see Tuning ASP Metabase Settings.

- Avoid transitions from managed code to unmanaged code in your applications, if possible.

- Enhance database-driven performance by using Microsoft SQL Server. Create and maintain good indexes to minimize I/O during database queries. Also, use stored procedures, which take much less time to execute and are easier to write than ASP scripts designed to do the same task.

- Design your site so that applications push the processing load onto the client.

  Client-side checks reduce roundtrips and help to avoid network latencies, conserve server-side resources, and allow changes to appear instantaneously to the user. For example, add client-side code to validate that forms are filled out with good data, such as checking that credit-card numbers have a correct checksum. However, the server must still validate the data, because malicious clients can send malformed data.

**Optimizing the Design of ASP Applications**
Check this list for ASP application design guidelines that can optimize performance.

- Turn off ASP debugging, and verify that both the **AppAllowDebugging** and **AppAllowClientDebug** metabase properties are set to false.

- Avoid storing apartment-threaded components in ASP application or session state. Remove Visual Basic objects from ASP Session state if they are apartment threaded.

- Cache output from ASP if possible; if not, cache inputs to ASP if possible.

- Convert scripts in ASP pages into COM objects. Scripts in ASP pages are interpreted; if a script involves one or more loops, converting it to a COM object can improve its performance considerably.

- ASP is now able to run all of its threads in a multithreaded apartment (MTA). If your COM components are primarily free-threaded or both-threaded, running the ASP threads as MTA can improve performance significantly. To enable all ASP applications to run in an MTA, you can use the metabase property **AspExecuteInMTA** at the application level. For more information about designing ASP to run its threads in MTA, see Configuring Applications to Use COM+ Services.

- If an ASP script is more than 100 lines (counting lines of code in included files), consider creating a COM+ component to provide the same function. If COM+ components are written efficiently and debugged properly, they can offer 20 to 30 times the processing speed of a script for the same dynamic page.

| **Tip** |
| --- |
| The easiest way to measure the size of an ASP script with **#include** directives is to change the file name extension of the page from .asp to .stm, and then open the .stm file by using your browser. Use your browser's **View Source** command to display the .asp file and lines of code from the included files. |

**Optimizing Network Administration and Design**

Check this list for network administration and design guidelines that can optimize performance.

- Set HTTP expire headers for all images and for HTML so that proxy servers and browsers make fewer calls to the Web server. For information about setting an expiration date for Web site content, see Using Content Expiration.

- Use SSL certificates only when needed and only for content that requires security. Because SSL uses complex encryption that consumes considerable processor resources, it takes much longer to retrieve and send data from SSL-enabled directories. Therefore, only secure the files that contain or receive sensitive information, such as credit-card numbers or medical information. Also, keep these pages free of other elements that consume resources, such as images. For more information about implementing an acceptable level of security while preserving performance, see Balancing Performance and Security.

- Minimize initialization processing wherever possible. For example, enable caching, when appropriate; on multiprocessor computers, assign processor affinity to better use CPU caching; and use HTTP compression for faster downloads. For information about caching, see Monitoring Overall Server Memory. For information about using compression on different file types, see Using HTTP Compression for Faster Downloads.

- Run IIS in worker process isolation mode, which allows you to periodically restart worker processes assigned to an application pool (either manually or automatically by using preconfigured settings). Recycling worker processes can help minimize problems attributed to faulty Web applications.

- Establish Web gardens, which are application pools that can run multiple worker processes.

- Stress-test applications in your lab before you place them on a production server.

- If you have a large installation with many FTP sites, each attached to a different IP port, your server might create too many backlog monitor threads (the system creates a thread for every 64 sites). The workaround for this is to disable the backlog monitor. To disable the backlog monitor, set the value of the DisableBacklogMonitor registry entry to 1 and then restart the FTP services. The DisableBacklogMonitor entry is located in the following registry subkey:

  HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Inetinfo\Parameters

| **Caution** |
| --- |
| The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference. |

- If you expect a large spike in traffic on a particular page or set of pages, you can change that page or set it to static HTML for a few days. Switching from a dynamic page to static HTML can involve some redesign, but the effort is often worthwhile. For example, when many thousands of people, during a short period, traverse three or four scripted pages to get to a particular download page, your server uses a lot more resources than it needs to. To save time for your users, you can instead provide a temporary static front page that has simple navigation.

**Related Topics**

For information about editing the registry, see the Registry Reference.

# Creating a More Efficient Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Creating a page-load performance profile provides a systematic way of looking at a Web page and recording its performance characteristics across a wide area network (WAN). A page-load performance profile incorporates the distance, congestion, and limited bandwidth that 99 percent of your users experience when accessing your site.

Page-load performance profiles can help you determine the following:

- The impact of not setting expiration dates.
- Which files are taking the longest to load.
- Whether you have unneeded file content.
- Where to reduce and consolidate files.
- How long your pages take to load (on first access and second access).
- Whether you need to enable or adjust your compression configuration.

To perform a page-load test, install the appropriate tools to collect data so that you can develop performance profiles by analyzing HTTP traffic. Network Monitor is a protocol analyzer that looks into packet-level detail. For more information about obtaining and using this tool, see Network Monitor overview in Help and Support Center for Windows Server 2003.

After obtaining page-load performance profiles for your Web pages, you can use the performance data to identify changes that you can make in order to optimize the performance of the Web pages.

# Obtaining a Page-Load Performance Profile

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To obtain an accurate picture of a Web pages performance during normal use, you must obtain page-load performance data for the second access to a Web page as well as the first access. The first-access load time is usually much longer in duration than second-access load times. When a first-time user visits a Web page, the user's computer does not hold cached files from previous visits; hence, the user's browser must load all the files that are a part of the Web page. The procedure for measuring a Web page's first-access load time differs slightly from the procedure for measuring second-access load times.

For information about measuring first-access and second-access page-load times for a Web page, see Measuring Load Time for a Web Page.

After you complete your own page-load performance profiles, perform a page-load test against your competitor's sites, compare the results with your own performance data, and ask:

- Which site is faster?
- Does speed contribute to ratings?
- Which sites will users choose to come back to?

# Optimizing a Page-Load Performance Profile

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By optimizing page-load performance profiles, you can provide more efficient Web sites for your users. Begin by making sure that HTTP Keep-Alives are enabled, enabling HTTP compression if appropriate for your Web sites, and setting expiration dates on the files associated with your Web pages. Then identify your slow-loading files, reviewing them for unneeded file content. Next measure your file quantity and size. Finally, determine port usage and minimize unneeded screen refreshes. After reviewing all these factors and optimizing your Web sites, obtain a new page-load profile to determine whether the changes significantly improve the efficiency of your Web site.

For additional tips for optimizing your Web sites, see Best Practices for Creating Efficient Sites. For instructions telling how to obtain a page-load profile, see Obtaining a Page-Load Performance Profile.

**Verify That HTTP Keep-Alives Are Enabled**

Another item to verify is the connection = keep-alive state for each TCP connection. If keep-alive connections are turned off, every file requires a new TCP connection. For small files, enabling HTTP Keep-Alives in IIS effectively doubles the number of roundtrips. For more information about HTTP Keep-Alives, see Enabling HTTP Keep-Alives to Keep Connections Open.

**Set HTTP Compression**

If you have Web sites that use large amounts of bandwidth, consider enabling HTTP compression to obtain faster transmission times between your server and compression-enabled browsers. The performance cost of compressing a static file is small and is typically incurred only once because the compressed file is stored in a compression cache.

The cost of expanding the file at the browser is minimal. Compressed files download faster, which makes them especially beneficial to the performance of any browser that uses a network connection with restricted bandwidth, such as a modem connection.

For more information about enabling HTTP compression, see Using HTTP Compression for Faster Downloads.

**Set Expiration Dates on Files**

A common performance mistake is forgetting to set the expiration dates on the files associated with Web pages. When a user returns to a Web page, Internet Explorer already has most of the files for the page in its cache, but it does not use these files if the expiration dates are in the past. Instead, Internet Explorer sends a GET request to the server for the file, indicating the date of the file in the cache. If the file has not changed, the server sends a Not Modified message. This GET/Not-Modified sequence costs the client a roundtrip. For small files of 3 KB or less, this roundtrip is as expensive as getting the original file.

For more information about setting expiration dates, see Using Content Expiration.

**Identify Slow-Loading Files**

Identifying which files are taking the longest to load often provides the best clues about what you need to improve. Causes of very long load times can include server capacity issues, network congestion, and large file content. Use the data that you collect from Network Monitor captures to identify how long each file takes to load. Visit a Web page several times and observe which files are regularly loaded most slowly. Multiple visits to a page help improve statistical veracity.

**Remove Unneeded File Content**

Files often contain clutter, such as tabs, spaces, new lines, and comments, which consumes up to 15 percent or more of a page. For production Web sites that are visited by millions of users, remove every bit of file content that is not absolutely necessary for the rendering of the page.

View the packet contents of your files with Network Monitor. Look for repetitive text such as <Option Value = data> statements that are used to insert text into drop-down lists. It takes much less space to add a few lines of code that sends an array (that reads the drop-down lists) to the client.

**Measure File Quantity and Size**

The number and size of files that you load on your Web site have a large impact on page-load times. Fewer large files load faster than many small files, as follows:

- Number of files. For each file, HTTP must perform a GET-File. Each GET-File is at least a single roundtrip. If the files are on a new server that the user has not previously connected to, HTTP must also make another roundtrip to create a new connection. Roundtrips across a modem cost 250 milliseconds or more.

The optimal minimum number of files for a Web site is about four. Consider consolidating like files; for example, have only one .js or .css file instead of two or three.

- Size of files. The TCP slow-start algorithm favors larger files. When a new connection is first opened, TCP sends a small number of packets in response to a GET-File. Thereafter, for each acknowledgement received by the server, TCP increases the number of response packets sent to the client.

  Small files of 3 KB or less are complete in one roundtrip. Files larger than 20 KB receive only about 16 KB of data per roundtrip. The size of a file can make the difference between a very small bit rate of 1 to 10 kilobits per second (Kbps) and a high bit rate of 40 to hundreds of Kbps.

**Determine Your Port Usage**

Using HTTP 1.1, Internet Explorer can open a maximum of two ports to any one server and a maximum of four ports total. The best performance occurs when four ports are opened and are kept open to be reused by all the files. Reusing a port is a good thing; when TCP slow-start increases the window size for a file, the next file that uses the same port can also use the increased window size.

View your Network Monitor capture summary to determine how many ports were used during your page-load tests, and by which files. Keep in mind that measuring Web pages through a proxy server hides the true number of ports being used. The best tests from your office are conducted by using a direct connection to the Internet without a proxy.

**Control Screen Refreshes**

In a refresh request, Internet Explorer (the client) tells the server the date stamp of the version of a file that it has in its cache by using the If-Modified-Since header. IIS 6.0 then determines whether the file has been modified since that time. If the file is unmodified, IIS 6.0 replies with a brief Not Modified response.

In routine file updates for your Web site, do not copy files that have not been modified. A static HTML page is not retrieved during a screen refresh if it has not been updated. Some publication processes copy files that have not been modified, which gives the files new time stamps and thus updates them as far as the system is concerned. These files are downloaded even though they are actually unchanged. When you set up your publication process, make every effort to avoid this waste of resources.

# Best Practices for Creating Efficient Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following best practices can help you improve the efficiency of important, often-used Web pages.

**Optimize settings**

- Set file expiration dates. By setting future expiration dates on each file at the IIS server, you can reduce second-access load times by 50 to 70 percent.
- Enable HTTP Keep-Alives in IIS.
- Author hyperlinks consistently because the Internet Explorer cache is case sensitive.
- Use cache-control extensions. Internet Explorer version 5 introduced support for two extensions to the cache-control HTTP response header: pre-check and post-check.
- Enable HTTP compression to increase efficiency of bandwidth use.

**Use design elements judiciously**

- Standardize images. For example, if your site uses two images that are slightly different from each other, either in size, color, or design, standardize to a single image.
- Reduce Graphics Interchange Format (GIF) image color depth or consider using Portable Network Graphics (PNG) format.
- Eliminate unnecessary images that appear below the scroll bar. Users view these images less frequently; so weigh their value against their performance cost. For example, use text rather than images for arrows, lines, and bullets. Incorporate several small images into one large graphic (or remove some of the images). Substitute smaller images for large ones.
- Use the most compact graphics format that supports the quality of image that you need. Images saved as a file type that is inherently large, such as a Tagged Image File Format (.tif), can increase download times; consider substituting a more compact format such as a Graphics Interchange Format (.gif) or a Joint Photographic Experts Group (.jpg) file. You might also consider compressing existing images.

**Ensure that developers design scripts and code that load quickly**

- Reduce clutter. Space characters, tabs, and remarks often consume more than 20 percent of an .htm or an .asp file. Delete this clutter before you load the files on production servers.
- Remove redundant tags. Consider using arrays to fill drop-down lists instead of repeating [<Option Value = data> data</Option>] again and again.
- Use JavaScript (.js) files to consolidate commonly used scripts.
- Reduce the number of small files. Build whole images instead of multiple component images. For tabs, use a single image map.
- Consolidate multiple images and adjacent images into a single file.
- Do not use nested <table> tags.

**Test load times for your pages**

You can use the following tools to test the load times for your Web pages:

- WAN simulators
- Network Monitor
- Load simulators for the Internet environment
- Reporting tools

# Partnering with Other Groups to Improve Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In enterprise organizations, where corporate intranets might be accessed by employees in far-reaching parts of the country or the world, IT administrators in areas that share the intranet must work together to ensure optimal performance for all users. Take time to find out how the performance of other areas is affecting the performance of your applications and systems; then work with IT administrators in other areas to coordinate efforts to improve performance.

If the systems and applications that you manage depend on the performance of other systems and applications, use this guide to help you collaborate with your IT partners:

1. Identify your partners:

   - Start early in the IT cycle to contact partners and find performance allies in other groups and on your own team. Obtain agreement to work together to maintain and improve performance, including upgrading, adding to, or changing your network components.

   - Determine how the performance of other systems and applications affects your systems and applications, and vice versa; then document the dependencies and share the data with your partners.

2. Before you begin your design, obtain information about your users. You must know your users in order to set realistic goals for response times and system preferences. If possible, conduct usability tests to discover the actual needs, behaviors, and preferences of real users.

3. Set performance goals for your group, and establish a timeline:

   - Set performance goals and test plans for your group.

   - Discuss key milestones in your timeline.

   - Identify competing priorities, especially conflicting priorities of partners outside your immediate group.

4. Coordinate your plans and goals with your partners plans and goals:

   - Share your goals, timeline, and milestones with your partners; then synchronize goals and timelines with them. Define terminology to ensure that everyone understands each goal and service level. Include a plan for communicating with your partners.

   - Write contracts, if needed, to set service level agreements for performance and to balance competing priorities so that other groups priorities do not adversely affect yours (and vice versa).

   - Establish baseline requirements for operating systems and hardware.

   - Establish guidelines for how to monitor any mutual performance dependencies. Establish minimum testing standards. Establish a system for logging bugs against performance problems.

5. Design the systems and applications for your group to minimize their impact on other components:

   - Design your client systems so they do not put unnecessary load on servers. For example, ensure that clients' browsers are compression-enabled, and enable HTTP compression of static files on your Web server to provide faster transmission times.

- Test to identify and minimize unnecessary client requests, such as requests from impatient users who make numerous attempts to access the same page. It is recommended that you set the AspQueueConnectionTestTime Metabase Property so that your server does not process abandoned requests, enable HTTP Keep-Alives, and set up caching to minimize requests from clients Web browsers.

# Additional Resources for IIS 6.0 Performance

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this optimizing IIS 6.0 performance.

**Related Information**

- Analyzing Log Files for information about using IIS logs to monitor and tune performance.
- Common IIS 6.0 Administrative Tasks for step-by-step instructions to help you complete many of the tasks described in the performance topics.
- Troubleshooting IIS 6.0 for information about troubleshooting performance.
- Web Server Scalability for information about IIS 6.0 capacity planning and scalability.
- Article 247658, Configure Secure Sockets Layer Server and Client Cache Elements, in the Microsoft Knowledge Base for information about configuring SSL session caching, including how to set the ClientCacheTime registry setting.
- Article 313064, Monitor Web Server Performance by Using Counter Logs in System Monitor in IIS, in the Microsoft Knowledge Base for information about creating counter logs, including baseline IIS logs.
- The Windows System Resource Manager — Fast Facts site for information about using Windows System Resource Manager.
- Microsoft Application Consulting and Engineering (ACE) Team, *Performance Testing Microsoft .NET Web Applications*. Redmond, WA: Microsoft Press, 2003. Read this source for information about how to perform stress testing by using Microsoft Application Center Test (ACT), how to use System Monitor and Network Monitor, and how to conduct an application network analysis.

**Related IIS 6.0 Operations Guide Topics**

- Enabling Logging, which is accessible from IIS Manager, for information about performance logging.
- Limiting Connections for information about setting global and individual site connection limits.
- Administering Network Resources for information about the QoS features in IIS 6.0.
- Analyzing Security Data and Planning Upgrades for information about IIS 6.0 security features.
- Web Application Guide for information about Web application development, including ASP.NET and ASP applications.

**Related Windows Server 2003 Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- Analyzing performance data in Help and Support Center for Windows Server 2003 for information about analyzing log data and a list of acceptable values for counters.
- Performance console overview in Help and Support Center for Windows Server 2003 for information about Windows performance monitoring tools.
- Monitoring performance from the command line in Help and Support Center for Windows Server 2003 for information about command-line tools that you can use to monitor performance.

- The TCP/IP model in Help and Support Center for Windows Server 2003 for information about how the OSI model relates to the TCP/IP reference model.

**Related Tools**

- The **logman** command

  Use the **logman** command to schedule performance counter and event trace log collections on local and remote systems. For more information about the **logman** command, see Logman in Help and Support Center for Windows Server 2003.

- Log Parser tool (LogParser.exe)

  Use the Log Parser tool to extract data from the raw file. The Log Parser tool and its accompanying user documentation are included in the IIS 6.0 Resource Kit Tools.

- The **start netmon** command

  Use the **start netmon** command to start the Network Monitor user interface from the command line. You can also specify parameters for capturing and filtering frames. For more information about using this command-line option, see Managing Network Monitor from the command line in Help and Support Center for Windows Server 2003.

- Web Capacity Analysis Tool (WCAT), which contains a stress client tool (Wcclient.exe) and a controller tool (Wcctl.exe).

  Use WCAT to perform stress testing. WCAT and its accompanying user documentation are included in the IIS 6.0 Resource Kit Tools.

# Web Server Scalability

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0, running on the Microsoft® Windows® Server 2003 operating system, includes a new architecture and new features to help your application server scale. *Scaling* is the ability of a system to handle increasing demands at an acceptable performance level. If an application server and the underlying infrastructure scale, the application can handle steady traffic growth, in addition to occasional peaks in traffic volumes, while maintaining good performance. If an application server and the underlying infrastructure do not scale, clients connecting to that application can experience poor response times, errors, or other problems that might ultimately frustrate end users.

**In This Section**

Scaling IIS 6.0

Improving IIS 6.0 Scalability and Availability with Network Load Balancing

Improving Scalability by Optimizing IIS 6.0 Queues

Improving Scalability by Optimizing IIS 6.0 Caches

ISP Scaling -- Strategies for Hosting Thousands of Sites

Improving Scalability Through UNC-Based Centralized Content Storage

Case Study: Scaling an ASP.NET Web Application on IIS 6.0

Additional Resources for IIS 6.0 Scalability

**Related Information**

For information about performance monitoring and tuning, see Optimizing IIS 6.0 Performance.

# Scaling IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Changes in the Internet Information Services (IIS) 6.0 architecture have led to dramatic improvements in how IIS scales. During preliminary testing, IIS 6.0 scaled better than earlier versions of IIS. Those preliminary results showed that on a default installation, while running in worker process isolation mode, IIS 6.0 was capable of the following:

- Hosting as many as 20,000 static-content sites.
- Hosting as many as 2,000 application pools.
- Processing 25,000 requests per second for a 15-KB cached static file with Hypertext Transfer Protocol (HTTP) Keep-Alives enabled.
- Running hundreds — or potentially thousands — of simultaneous worker processes.

| **Note** |
| --- |
| Performance and scalability results vary depending on hardware and software configurations. Microsoft does not make any warranties or guarantees about these preliminary results. |

Additionally, processing throughput has greatly improved for installations with up to eight processors. Default thread pool, queue, and cache settings have all been tuned to increase throughput. Startup and shutdown times are faster than earlier versions. Simply put, your IIS Web server installation should scale better on a default installation of IIS 6.0 than on any earlier version of IIS. However, there are many variables in a Web server installation that can affect how IIS 6.0 scales. Scalability is dependent on the least scalable part in your server installation. Factors such as network latency, slow back-end processing, a shortage of RAM on your servers, or poorly designed and untested applications can all undermine how IIS 6.0 scales. If one part does not scale, the whole installation suffers.

This section provides a brief overview of the many changes and features that help IIS scale, and also describes some of the changes and features in the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Web Edition; and Windows® Server 2003, Datacenter Edition operating systems that can increase the scalability of your server installation. Be aware that scaling an application server is a process much like tuning the performance of your server. Performance and scalability are intricately linked, and both goals require long-term planning, testing, and monitoring. For more information about performance monitoring and tuning, see Optimizing IIS 6.0 Performance.

# Scalability Features in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This section describes features in IIS 6.0 that can improve the scalability of your application server.

| **Important** |
| --- |
| Do not enable the features described in this chapter on a production server before you establish a performance baseline in a test environment. If enabling any of these features degrades performance, continue testing or disable the feature and return your server to the pre-change configuration. |

**Worker process isolation mode**

IIS 6.0 introduces worker process isolation mode, which runs all Web applications in an isolated environment. When you run IIS in worker process isolation mode, you can configure applications to run in separate application pools. An *application pool* is a grouping of URLs routed to one or more worker processes. An application pool allows specific configuration settings to apply to these worker processes.

Worker processes operate independently: if one fails, it does not affect other worker processes. The pooling of applications protects applications from the effects of worker processes that support other application pools. In this way, applications are protected from each other.

When a worker process enters an unhealthy state and stops processing requests, the HTTP protocol stack (HTTP.sys) continues to queue requests. The World Wide Web Publishing Service (WWW service) detects that the worker process is unhealthy and shuts it down, starting a new worker process if requests are waiting in the queue or if new requests arrive. Hence, even when a worker process fails, the WWW service, by means of HTTP.sys, continues to accept requests and shields the user from loss of service because the client can still connect to the WWW service.

For more information about IIS 6.0 application isolation modes, including how to enable and configure worker process isolation mode, see Running IIS 6.0 as an Application Server. For more information about how IIS 6.0 processes requests, see IIS 6.0 Architecture.

**On-demand process start**

IIS 6.0 does not allocate resources at initialization time; instead, it allocates resources as needed, which increases site scalability. In particular, IIS 6.0 application pools only start processes when the first request for a URL is served by the application pool.

**HTTP Keep-Alives**

IIS 6.0 supports HTTP Keep-Alives, an HTTP specification that can significantly improve server performance. Most Web browsers request that the server keep the connection open across multiple requests; this is called a *Keep-Alive*. Without this feature, a browser must make numerous connection requests for a page that contains multiple elements, such as graphics. A separate connection might need to be made for each element. HTTP Keep-Alives increase server efficiency by reducing server activity and resource consumption. HTTP Keep-Alives can also increase browser responsiveness across a slow connection. For more information about HTTP Keep-Alives, see Optimizing IIS 6.0 Performance.

**Web gardens**

When running in worker process isolation mode, IIS can be configured to have multiple worker processes servicing requests for an application pool. This configuration is known as a Web garden. When a worker process in an application pool is busy processing a request, other worker processes can accept and process requests for the application pool. For more information about Web gardens, including how to enable or configure this feature, see Running IIS 6.0 as an Application Server.

**Processor affinity**

To efficiently use resources on large multiprocessor servers, IIS 6.0 enables you to configure application pools to establish affinity between worker processes and individual processors. When set, processor affinity forces worker processes to run on specific microprocessors or CPUs. This affinity applies to all worker processes that serve the Web sites and applications of an application pool. For more information about processor affinity, including how to enable or configure this feature, see Optimizing IIS 6.0 Performance.

**Idle timeout for worker processes**

An IIS 6.0 application pool can be configured so that its worker processes request a shutdown if they are idle for a configurable amount of time. Shutting down worker processes frees up unused resources and reduces CPU load. For more information about idle timeout, including how to enable or configure this feature, see Running IIS 6.0 as an Application Server.

**Bandwidth throttling**
If the network or Internet connection used by your Web server is also used by other services, such as e-mail or news, you can limit the bandwidth used by your Web server so it is available for those other services. If your Web server hosts more than one Web site, you can individually throttle the bandwidth used by each site. For more information about bandwidth throttling, including how to enable or configure this feature, see Optimizing IIS 6.0 Performance.

**HTTP compression**
HTTP compression allows faster transmission of pages between the Web server and compression-enabled clients. In addition, HTTP compression makes the best use of available bandwidth and can significantly increase site performance. For more information about HTTP compression, including how to enable or configure this feature, see Optimizing IIS 6.0 Performance.

**64 GB cache capacity**
For workloads that require large amounts of cached data, the IIS 6.0 kernel-mode driver, HTTP.sys, can be configured to cache up to 64 gigabytes (GB) on an x86-based or compatible processor–based system. For more information about IIS caches, see Improving Scalability by Optimizing IIS 6.0 Caches.

**Centralized binary logging**
IIS 6.0 supports centralized binary logging, where multiple Web sites send binary, unformatted log data to a single log file. When many Web sites are hosted on the same server, the process of creating hundreds or thousands of formatted log files and writing the log data to disk can quickly consume valuable CPU and memory resources, thereby creating performance and scalability problems. Centralized binary logging in IIS 6.0 minimizes the server resources that are used for logging, while providing detailed log data for organizations that require it. For more information about centralized binary logging, including how to enable or configure this feature, see Analyzing Log Files.

# Scalability Features in Windows Server 2003

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This section describes features in Windows Server 2003 that can improve the scalability of your application server.

| **Important** |
| --- |
| Do not enable features described here on a production server before you establish a performance baseline in a test environment. If enabling any of these features degrades performance, continue testing or disable the feature and return your server to the pre-change configuration. |

**Kernel-mode driver**
Windows Server 2003 introduces a new kernel-mode driver, HTTP.sys, for HTTP parsing and caching. IIS 6.0 is built on top of HTTP.sys. HTTP.sys is specifically tuned to increase Web server throughput by directly processing requests in the kernel (in specific circumstances) or by efficiently routing requests to user-mode worker processes. The combination of kernel request processing and efficient user-mode routing have dramatically improved how IIS 6.0 scales and performs. For more information about HTTP.sys, see IIS 6.0 Architecture.

**64-bit support**
The complete Windows Server 2003 code base is compiled for both 32-bit and 64-bit platforms. Organizations that need highly scalable applications can take advantage of an operating system that runs — and is supported — on these two platforms.

**WSRM**
Windows System Resource Manager (WSRM) is a feature of Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.

With WSRM, you can control how CPU and memory resources are allocated to applications, services, and processes. Managing resources in this way improves system performance and reduces the chance that applications, services, or processes will interfere with the rest of the system; it also creates a more consistent and predictable experience for users of applications and services running on the computer.

WSRM can be installed from the Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition operating system CD. For more information about WSRM, see Help in the WSRM snap-in.

**DFS**
Distributed File System (DFS) unifies files on different computers into a single namespace, making it easy to build a single, hierarchical view of multiple file servers and file server shares on a network. To use DFS as the file system for IIS, select the root for the Web site as a DFS root. You can then move resources within a DFS tree without affecting any HTML links. For more information about DFS, see Distributed File System and File Replication Services.

**FRS**
The File Replication service (FRS) provides multimaster file replication for designated directory trees between designated servers running Windows Server 2003. The designated directory trees must be on disk partitions formatted with the version of the NTFS file system used with Windows Server 2003. FRS must be used with DFS. DFS uses FRS to automatically synchronize content between assigned replicas. The combination of DFS and FRS can also work with the Active Directory® directory service to automatically synchronize the content of system volume information across domain controllers. For more information about FRS, see Distributed File System and File Replication Services.

# Improving IIS 6.0 Scalability and Availability with Network Load Balancing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The best way to guarantee the scalability and availability of your Internet services is to host your site with more than one computer. Microsoft Windows Server 2003 includes the Network Load Balancing service. Network Load Balancing enhances the availability and scalability of Internet server applications such as those used on Web, File Transfer Protocol (FTP), firewall, proxy, and virtual private network (VPN) servers, in addition to other mission-critical servers. A single computer running Windows can provide a limited level of server reliability and scalable performance. However, by combining the resources of two or more computers running Windows Server 2003 into a single cluster, Network Load Balancing can deliver the reliability and performance that Web servers and other mission-critical servers need. Figure 7.1 shows two connected Network Load Balancing clusters. The first cluster consists of two hosts and the second cluster consists of four hosts.




Each host runs separate copies of the server applications, such as those for a Web, FTP, or Telnet server. Network Load Balancing distributes incoming client requests across the hosts in the cluster. The load weight to be handled by each host can be configured as necessary. You can also dynamically add hosts to the cluster to handle the increased load. In addition, Network Load Balancing can direct all traffic to a designated single host, called the default host.

Network Load Balancing allows all of the computers in the cluster to be addressed by the same cluster IP address (but also maintains their existing unique, dedicated IP addresses). For more information about Network Load Balancing, see Help and Support Center for Windows Server 2003.

| **Note** |
| --- |
| FrontPage® Server Extensions from Microsoft do not work in a Network Load Balancing environment. Also, FrontPage Server Extensions do not work side-by-side with SharePoint™ Team Services from Microsoft on the same virtual server or Web site. |

**Related Topics**

For more information about upgrading your FrontPage-extended Web sites, see SharePoint Team Services.

# IIS Responses to Load-Balanced Application Pool Behaviors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 is designed to work with a variety of hardware and software network load balancers. However, the introduction of application pools and health-monitoring features such as rapid-fail protection in IIS 6.0 has the potential to confuse load balancers when an application pool fails to respond to a request, or when IIS health-monitoring features act on a faulty application pool. This section describes how HTTP.sys and IIS respond to load balancers for various application pool failures or problems.

For Layer 3 and Layer 4 load balancers (that is, basic load balancers such as Network Load Balancing), when an application pool fails, IIS causes the load balancer to run a TCP reset command. For Layer 7 load balancers (more advanced load balancers that balance the load based on higher-level application data, such as URLs, session data, and cookies) HTTP.sys sends a 503 error and attempts to connect to the application pool again after a specified period of time. IIS can also be configured to shut down the faulty application pool altogether. These responses can remedy application failures at the application pool level instead of forcing the load balancer to fail over the entire server. These responses also allow healthy application pools to continue processing requests without forcing the load balancer to fail over the entire server.

Table 7.1 summarizes how HTTP.sys and IIS respond to load balancers for various application pool failures.

**Table 7.1 IIS Responses to Load-Balanced Application Pool Behaviors**

| Condition | Auto-Shutdown of Application Pool? | Layer 7 Load Balancer Behavior | Layer 3 and Layer 4 Load Balancer Behavior |
|---|---|---|---|
| The application pool entered rapid-fail protection. | Yes | HTTP 503 error | IIS resets the TCP connection. |
| The application pool CPU threshold was exceeded. | Yes | HTTP 503 error | IIS resets the TCP connection. |
| An administrator disabled the application pool or shut down the WWW service. | No | HTTP 503 error | IIS resets the TCP connection. |
| The kernel-mode request queue is full. | No | HTTP 503 error | IIS resets the TCP connection. |

By default, the IIS metabase property **LoadBalancerCapabilities** is configured for Layer 7 responses (**LoadBalancerCapabilities** = 2). If your organization uses Layer 3 or Layer 4 load balancers, you will need to change this property in the IIS metabase to a value of 1.

| **Important** |
|---|
| It is recommended that you back up the IIS metabase before you change the metabase. For additional information, see Working with the Metabase. |

For more information about configuring load balancing, see the LoadBalancerCapabilities Metabase Property. For more information about the IIS 6.0 metabase, including how to change a metabase property, see Working with the Metabase.

**IIS Response When a Load-Balanced Application Pool Enters Rapid-Fail Protection**

When an application pool enters the rapid-fail protection state, HTTP.sys responds to requests with a 503 error. If IIS is configured to shut down the application pool when the application pool enters the rapid-fail protection state, IIS executes the auto-shutdown application pool command. For more information about rapid-fail protection, see Running IIS 6.0 as an Application Server. For load-balanced applications, HTTP.sys responds to requests in one of the following ways:

- Layer 7 response. When HTTP.sys receives requests for an application pool that is in the rapid-fail protection state, HTTP.sys sends a 503 error. If IIS is configured to shut down the application pool when it enters the rapid-fail protection state, IIS executes the auto-shutdown application pool command. When the error has been resolved, the administrator must restart the application pool.

- Layer 3 and Layer 4 response. When HTTP.sys receives requests for an application pool that is in the rapid-fail protection

state, HTTP.sys resets the TCP connection that carried the request. If IIS is configured to shut down the application pool when it enters the rapid-fail protection state, IIS executes the auto-shutdown application pool command. When the error has been resolved, the administrator must restart the application pool.

For information about rapid-fail protection, see Running IIS 6.0 as an Application Server.

**IIS Response When a Load-Balanced Application Pool Exceeds a Configured CPU Threshold**
The IIS CPU monitoring feature allows you to configure a CPU threshold for a specific application pool. The IIS CPU monitoring feature includes the option to shut down an application pool if the CPU exceeds a configured maximum. For load-balanced applications, HTTP.sys responds to requests in one of the following ways:

- Layer 7 response. When HTTP.sys receives subsequent requests for an application pool that has exceeded the configured CPU maximum, HTTP.sys sends a 503 error. If IIS is configured to shut down the application pool when it exceeds its CPU threshold, IIS executes the auto-shutdown application pool command. When the error has been resolved, the administrator must restart the application pool.

- Layer 3 and Layer 4 response. When HTTP.sys receives requests for an application pool that has exceeded the configured CPU maximum, HTTP.sys resets the TCP connection that carried the request. If IIS is configured to shut down the application pool when it exceeds its CPU threshold, IIS executes the auto-shutdown application pool command. When the error has been resolved, the administrator must restart the application pool.

For information about CPU monitoring, see Running IIS 6.0 as an Application Server.

**IIS Response When a Load-Balanced Application Pool Is Disabled by an Administrator or an Administrator Shuts Down the WWW Service**
When an application pool is shut down, the requests that have already been queued in kernel mode are lost. A shutdown of this nature indicates serious application failure and presents a potential threat to the rest of the system if the application continues processing. Such an application failure can affect the integrity of data in custom applications. For this reason, all requests to that application pool should be stopped.

- Layer 7 response. When HTTP.sys receives requests for a disabled application pool, or if the administrator shuts down the WWW service, HTTP.sys sends a 503 error. When the error has been resolved, the administrator must restart the application pool.

- Layer 3 and Layer 4 response. When HTTP.sys receives requests for a disabled application pool, or if the administrator shuts down the WWW service, HTTP.sys resets the TCP connection that carried the request.

**IIS Response When the Kernel-Mode Request Queue Is Full**
A full kernel-mode request queue is a temporary state. After application queues finish processing other requests, the kernel-mode request queue will begin accepting new requests.

- Layer 7 response. When HTTP.sys receives requests for an application pool but the request queue is full, HTTP.sys sends a 503 error.

- Layer 3 and Layer 4 response. When HTTP.sys receives requests for an application pool but the request queue is full, HTTP.sys resets the TCP connection that carried the request.

# Preserving Session State in Network Load Balancing Web Server Clusters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS enables a Web application to maintain user session data (also called *session state*) across multiple requests. For example, an online ordering application might maintain the user's shopping cart as part of the user's session. IIS uses *session cookies*, which are text files supplied by the user's browser and stored on the user's computer for each request to keep track of which session belongs to which user.

| **Note** |
| --- |
| Worker processes lose session state if the worker process is recycled. The ideal configuration for storing session state with IIS 6.0 is to store user session data in a back-end database. |

Ideally, Network Load Balancing always routes connections from a user who has an active session to the server where that session is stored. However, Network Load Balancing and other Layer 4 network load balancers do not use sessions or cookies to decide which Web server should handle an incoming request. For this reason, sessions can be lost between requests in a Network Load Balancing server cluster.

There are several ways to avoid losing sessions:

- Design the Web site application to store session state on the client. For more information, see Preserving Session State with ASP Web Applications in Network Load Balancing and Preserving Session State with ASP.NET Web Applications in Network Load Balancing.

- Store session state on a central server-side store, such as a database.

- Use the client affinity feature. When client affinity is enabled, Network Load Balancing directs all TCP connections to the same cluster host. This allows session state to be maintained in host memory. You can enable client affinity in the **Add/Edit Port Rules** dialog box in Network Load Balancing Manager. Choose either **Single** or **Class C** affinity to ensure that only one cluster host will handle all connections that are part of the same client session. This is important if the server application running on the cluster host maintains session state (such as server cookies) between connections. For more information about Network Load Balancing affinity, see Help in the Network Load Balancing snap-in.

# Preserving Session State with ASP Web Applications in Network Load Balancing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can enable session state persistence for Active Server Pages (ASP) applications. If you enable session state, the server creates a new Session object for each connection, session state is accessible, session storage is allowed, **Session_OnStart**and **Session_OnEnd**occurs, and the **ASPSessionID** cookie is sent to the client. If you do not enable session state, state access and storage are not allowed, events are not processed, and no cookie is sent. By default, session state is enabled with the session timeout set to 20 minutes. The corresponding metabase property is **AspAllowSessionState**.

- For information about enabling session state for an ASP application, see Enabling Session State Persistence.

# Preserving Session State with ASP.NET Web Applications in Network Load Balancing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you use **SqlServer** or **StateServer** session state mode, session state might be lost when you run an ASP.NET Web application in a Network Load Balancing server cluster. To maintain session state across different Web servers in the cluster, the application path of the Web site (for example, \LM\W3SVC\2) in the IIS metabase must be the same for all of the Web servers in the cluster.

On one Web server, the instance ID of the Web site where the ASP.NET application is hosted might be 1 (where the application path is \LM\W3SVC\1). On another Web server, the instance ID of the Web site might be 5 (where the application path is \LM\W3SVC\5).

| **Note** |
| --- |
| Web sites created on a clean installation of IIS 6.0 will maintain their individual site numbers, as determined by the **ServerComment** metabase property, across the load-balanced installation. For more information about the **ServerComment** metabase property, see the ServerComment Metabase Property.<br><br>For information about changing the application path of a Web site, see Synchronizing the Application Path of a Web Site. |

# Improving Scalability by Optimizing IIS 6.0 Queues

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A request queue holds a collection of one or more requests waiting to be executed. Request queues can develop when your site experiences high traffic volumes, or if your site blocks to wait for database processing. Large queues, or queues that are not cleared quickly, impact scalability by limiting the number of requests the server can process and by reducing throughput. More importantly, if users are forced to wait for a response from your server, they might perceive that the site is not available and click the **Stop** button on their Web browser. In this situation, the queuing problem grows because clicking the **Stop** button can fill a queue with requests that are no longer connected to a client's browser. If your site routinely develops queues, you might dramatically improve scalability and availability by adjusting the number of requests that can wait in a queue. You can set limits for the following request queues:

- TCP/IP queue
- Kernel request queue
- ASP request queue
- ASP.NET request queues

# TCP/IP Queue

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

At the TCP/IP layer, high-priority threads respond to client requests by validating that the server is or is not accepting and processing requests. If the server is accepting requests, the thread maintains the customer connection and attempts to hand off the request to HTTP.sys. The only reason you might need to adjust the TCP/IP queue is if users cannot connect to your site on the Web. Monitor the TCPv4\Connection Failures and TCPv4\Connections Reset counters (for TCPv6, use the same counters for the TCPv6 object). If the Connection Failures and Connections Reset counters increase steadily, examine your TCP/IP settings. For more information about TCP/IP, see Help and Support Center for Windows Server 2003.

# Kernel Request Queue

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After HTTP.sys detects the TCP/IP thread, it either begins serving the request with content stored in the HTTP.sys cache, if available (this is called a *cache hit*), or puts the request in a kernel request queue for the designated application pool. Because the request is still in the kernel, but is waiting in a queue for a specific application pool, this queue is sometimes called the *application pool queue*. In IIS Manager, this queue is called the *kernel request queue*.

When IIS 6.0 runs in worker process isolation mode, by default HTTP.sys checks the number of requests for a designated application pool queue before it adds a new request to the queue. If adding the new request to the queue exceeds the queue limit, the server rejects the request, logs a QueueFull error in the HTTP error log, and sends a 503 error (Service Unavailable) to the client. This 503 error response cannot be customized. During times of high traffic, if you experience this 503 error, try adjusting the number of requests that wait in the queue. Increase the number to a large value such as 16,000 and test to see if the 503 errors persist.

Requests in the application pool queue remain in the queue if you lower the queue maximum to a number less than 1,000 by using the following procedure. For information about adjusting the number of requests allowed in the kernel request queue, see Configuring Application Pool Queue-Length Limits.

# ASP Request Queue

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For requests to ASP pages, after the worker process picks up the request in the application pool queue, the request is sent to the Internet Server API (ISAPI) handler, which passes the request along to Asp.dll. If a thread is available, the ASP request is processed. If a thread is not available, the request waits in the ASP request queue. Requests are processed in the order in which they are received. If the ASP request queue fills to 3,000 requests (the default metabase value of **AspRequestQueueMax**) the server rejects the request, logs the error in the IIS log, and sends a 500 error (Server Too Busy) to the client.

If the number of queued requests is under 3,000, users with the last requests are forced to wait (called queue latency) as other requests are pulled from the queue and executed. This is an acceptable condition if the queue is cleared within a few seconds. After a few seconds, most users believe the server is not responding and click the **Stop** button on their Web browser. High queue latency results in poor performance for all ASP requests for the entire time the queue is saturated, because all ASP requests in a worker process share the same queue. Additionally, CPU utilization tends to spike when the queue is saturated, which results in slower response times.

You can improve ASP queuing by adjusting the default values for the **AspProcessorThreadMax** and **AspRequestQueueMax** metabase properties. **AspProcessorThreadMax** represents the maximum number of threads per processor, and **AspRequestQueueMax** represents the maximum size of the request queue. For more information about these metabase properties, see the AspProcessorThreadMax Metabase Property and AspRequestQueueMax Metabase Property.

Table 7.2 describes performance counters and their ideal values for ASP queuing. Monitor these counters in System Monitor by using a chart with one-second intervals before you change any metabase properties on your server. For more information about IIS 6.0 counters, see Performance Counters Reference.

**Table 7.2 Preferred Values for Active Server Pages Queuing Performance Counters**

| Object\Counter | Preferred or Ideal Value |
|---|---|
| Active Server Pages\Request Wait Time | As low as possible. This counter determines how long clients connecting to your site must wait before the page starts executing. If this number is high, users connected to your site might be frustrated by slow response times. |
| Active Server Pages\Requests Queued | As low as possible. If large queues develop, or if this number fluctuates, you might need to adjust the metabase properties that affect ASP queuing, as described in Tuning AspRequestQueueMax. This counter is not as important as the Request Wait Time counter. Queues can clear quickly, but long wait times can frustrate users connected to your site. |
| Active Server Pages\Requests Rejected | As low as possible. If this number is high and Requests Queued is high, you might be reaching the ASP request queue default maximum of 3,000. You might need to adjust the metabase properties that affect ASP queuing, as described later in Tuning AspRequestQueueMax. |
| Active Server Pages\Requests/second | Application-specific. You want this number to be as high as possible. If Requests Queued is high, this number will typically be lower. Adjusting the metabase properties that affect ASP queuing can increase it. |
| Processor\% Processor Time (for each processor) | Application-specific. If Requests Queued never increases and % Processor Time is low, it is likely that ASP queuing is not affecting the performance of your site. |

# Tuning AspProcessorThreadMax

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If the number of requests in the ASP queue is fluctuating and your CPUs are running below 50 percent, your requests might be blocking. *Blocking* occurs when a request that is being processed is forced to wait for an available thread to perform component processing or back-end processing. For example, an ASP page using ActiveX® Data Objects (ADO) to access a computer running Microsoft® SQL Server™ located on a separate computer will typically be blocked. If the SQL request executes a query that takes five seconds, and if dozens of requests arrive every second to execute this same query, soon all available threads will be allocated and each thread will wait five seconds for the query to be completed. You can reduce blocking by increasing the number of ASP threads.

| **Important** |
| --- |
| It is recommended that you back up the IIS metabase before you change the metabase. For additional information, see Working with the Metabase. |

Increasing the number of threads increases concurrency, which means that the server has more threads available to process requests (especially nonblocking requests). This might reduce queuing and improve response times. On a site with minor blocking (that is, the queue only occasionally reaches the value of the AspRequestQueueMax Metabase Property), increasing the threads can have a dramatic effect; in fact, it might eliminate all ASP queuing, because there are always enough threads to handle the blocking and nonblocking requests.

Typically, you can reduce blocking by increasing the number of threads in the AspProcessorThreadMax Metabase Property from 25 (the default) to 50 threads.

| **Note** |
| --- |
| After changing the value of **ASPProcessorThreadMax**, you will need to restart the WWW service. |

After increasing the number of ASP threads, you should expect to see a slight increase in CPU utilization, and the queue should fluctuate more quickly. If, after increasing the number of threads, Requests Queued increases and stays high and % Processor Time decreases, a significant percentage of your requests are being blocked. You might want to reexamine your application code and determine whether the source of the blocking can be mitigated or worked around. Or you might want to simply add more servers to your Web farm to reduce contention.

If your queue decreases and CPU utilization increases, continue increasing **AspProcessorThreadMax** until you hit your target CPU utilization (do not exceed 70 percent, and keep **AspProcessorThreadMax** below 100). Ultimately, you might discover that the network input/output (I/O) or another factor is placing a limit on your server's performance. You might not be able to improve this limit.

For more information about the IIS Metabase, see Working with the Metabase. For more information about ASPProcessorThreadMax and ASPRequestQueueMax, see the Metabase Property Reference.

# Tuning AspRequestQueueMax

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The goal of tuning the AspRequestQueueMax Metabase Property is to ensure good response time while minimizing how often the server sends the HTTP 503 error (Server Too Busy) to clients when the ASP request queue is full. If the value of **AspRequestQueueMax** is too low, the server will send the HTTP 503 error too often. If the value of **AspRequestQueueMax** is too large, users might perceive that the server is not responding when in fact their request is waiting in the queue. By watching the queue during periods of high traffic, you should see a pattern of ups and downs. Make note of the peak value, and set the value of **AspRequestQueueMax** just above the peak value. Use the queue to handle short-term spikes, ensure response time, and throttle the system to avoid overload when sustained, unexpected spikes occur. If you do not have data for adjusting **AspRequestQueueMax**, a good initial setting seems to be a one-to-one ratio of queue to total threads. For example, if the AspProcessorThreadMax Metabase Propertyy is set to 25 and you have four processors (4  25 = 100 threads), set **AspRequestQueueMax** to 100 and tune from there.

| **Note** |
| --- |
| After changing the value of **AspRequestQueueMax**, you will need to restart the WWW service.<br><br>For more information about the IIS Metabase, see Working with the Metabase. |

# Additional Metabase Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following metabase properties also affect queuing.

**AspQueueTimeout**
The AspQueueTimeout Metabase Property specifies the amount of time (in seconds) that an ASP script request is allowed to wait in the queue. When requests are pulled from the queue, they are checked to see if they have expired (that is, they have waited longer than the value of this property). Expired requests are rejected with a 503 error (Service Unavailable). By default, ASP does not reject requests from valid connections, no matter how long they have been in the queue.

**AppPoolQueueLength**
The AppPoolQueueLength Metabase Property indicates to HTTP.sys how many requests to queue up for an application pool before it rejects future requests. When the limit for this property is exceeded, IIS rejects the additional requests with a 503 error. Setting the **AppPoolQueueLength** value to 0 means there is no maximum request queue length. This is not recommended.

For more information about these properties, see the Metabase Property Reference.

# ASP.NET Request Queues

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP.NET uses two request queues:

- The *global queue*, which is managed by the process that runs ASP.NET (Aspnet_wp). The global queue is configured in the Machine.config file by the **<processModel requestQueueLimit>** property.

- The *application queue*, or *virtual directory queue*, which is managed by the HttpRuntime class. The HttpRuntime class provides run-time services for ASP.NET applications. There is one queue for each virtual directory. The application queue is configured in Machine.config by the **<httpRuntime appRequestQueueLimit>** property.

When either queue exceeds its default limit, the request is rejected with a 503 error (Service Unavailable).

Worker processes that handle ASP.NET requests might run more threads than needed for ASP requests. Also, ASP.NET might queue requests more readily than ASP queues requests, but this should not be regarded as a scalability or performance issue. ASP.NET requests are executed more quickly than ASP requests. This means that with ASP.NET it is actually better to have fewer concurrent requests and more serial requests. Because ASP.NET maximizes the overall work on the CPU, this leads to fewer context switches.

Overall, ASP.NET is very self-tuning. For the global queue (Aspnet_wp process), the default ASP.NET thread numbers on four-processor and eight-processor servers should not require additional tuning. For servers with more than eight processors, you might want to decrease the default thread numbers by modifying the values of the **<processModel maxworkerthreads>** and **<processModel maxiothreads>** properties in Machine.config.

# ASP.NET Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use the following counters to monitor queuing for all ASP.NET applications on your server.

**Request Execution Time**
The value of this counter is the number of milliseconds taken to execute the latest request. In the Microsoft .NET Framework version 1.0, request execution time begins when the worker process receives the request and stops when the ASP.NET ISAPI sends HSE_REQ_DONE_WITH_SESSION to IIS. This counter includes the time taken to write the response to the client for IIS 5.0, but not for IIS 6.0. Consequently, for IIS 5.0, a client with a slow network connection increases the value of this counter considerably.

In .NET Framework version 1.1, request execution time begins when the HttpContext for the request is created and stops before the response is sent to IIS. Assuming that user code does not call HttpResponse.Flush, this means that request execution time stops before any bytes are sent to IIS or to the client.

There is no threshold for this counter. The value of this counter should be stable. Experience will help you set a threshold for a particular site.

**Requests Queued**
The value of this counter is the number of requests currently in the queue. It includes requests in all ASP.NET queues. The Aspnet_wp process queue is a named pipe through which the request is sent from one process to the other. The number of requests in this queue increases if there is a shortage of available I/O threads in the Aspnet_wp process. When the limit specified by **<processModel requestQueueLimit=/>** is exceeded, requests will be rejected with a 503 error (Service Unavailable). Note that if a request is rejected for this reason, it will never reach managed code, and error handlers will not be notified. Normally this only occurs when the server is under a very heavy load.

Virtual directory queues maintain the availability of worker and I/O threads. The number of requests in these queues increases if the number of available worker threads or available I/O threads falls below the limit specified by **<httpRuntime minFreeThreads=/>**. When the limit specified by **<httpRuntime appRequestQueueLimit=/>** is exceeded, the request is rejected with a 503 error and the client is sent an HttpException with the message Service Unavailable.

**Requests Rejected**
The value of this counter is the number of rejected requests. Requests are rejected when one of the queue limits is exceeded. Back-end latency, such as that caused by a slow computer running SQL Server, is often preceded by a sudden increase in the number of pipeline instances and a decrease in % Processor Time and Requests/second. A server might be overwhelmed during times of heavy load due to processor or memory constraints that ultimately result in the rejection of requests.

The value of this counter should be 0. Values greater than 0 should be investigated.

**Request Wait Time**
The value of this counter is the number of milliseconds that the most recent request spent waiting in the global queue. This does not include any time the request spent waiting in the application queues.

The threshold for this counter is 1,000. The average request should spend 0 milliseconds waiting in the queue.

# ASP.NET Application Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use the following counters to monitor queuing for individual ASP.NET applications.

**Requests Executing**
The number of requests currently being executed. This counter is incremented when the ASP.NET request pipeline begins to process the request and is decremented after the ASP.NET request pipeline finishes the request.

**Requests In Application Queue**
The number of requests in the application request queue. In addition to Requests Executing, Requests In Application Queue provides a warning when requests will be rejected. If there are only a few virtual directories, increasing the default value of the **appRequestQueueLimit** property in the Machine.config file to 200 or 300 might be suitable, especially for slow applications under heavy load.

**Pipeline Instance Count**
The number of active pipeline instances. Only one thread of execution can run within a pipeline instance, so this number gives the maximum number of concurrent requests that are being processed for a given application. The number of pipeline instances should not fluctuate. Sudden increases are indicative of back-end latency.

# Registry Entries for Thread Pool Queuing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following entries in the Windows registry determine the number of threads in the IIS thread pool. Tuning thread pools can produce dramatic improvement in IIS performance and scalability.

**MaxPoolThreads**
Specifies the maximum number of pool threads per processor. This is not an absolute limit; that is, IIS can increase the number of threads if the current load would benefit from having more threads available. Pool threads monitor the network for requests and process incoming requests. The value of **MaxPoolThreads** does not include threads that are consumed by ISAPI applications.

**PoolThreadLimit**
Specifies the maximum number of pool threads that can be created in the system. Pool threads monitor the network for requests and process incoming requests. **PoolThreadLimit** is an absolute limit that includes all core Web service threads. It is always greater than or equal to the value of **MaxPoolThreads**.

For more information about these registry entries, see the Global Registry Entries reference, and Optimizing IIS 6.0 Performance.

# Improving Scalability by Optimizing IIS 6.0 Caches

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A cache is a special subsystem that stores frequently requested data. Without caches or caching, IIS can potentially impose three time-consuming processes on each request, which can cause poor response times. These processes include:

- Reading metadata for each request, which generates a time-consuming remote procedure call (RPC) to the IIS Admin service.
- Retrieving the authentication token for the IUSR_*computername* account with a time-consuming RPC to Lsass.exe.
- Opening and reading the file.

To mitigate these time-consuming processes, HTTP.sys and IIS use caches. These caches are all enabled by default and should not require much tuning by administrators. If, while monitoring Web Service Cache counters in System Monitor, you begin to see rising numbers of cache misses or very low cache hit rates, you can adjust settings for the following:

- URI cache (user-mode URI cache)
- Token cache
- File cache
- HTTP.sys response cache (kernel URI cache)
- ASP caches, including the ASP template cache and the ASP script engine cache
- Global cache registry entries

For information about Universal Naming Convention (UNC) caching, see, Caching UNC-Based Files. For information about IIS cache counters, see the Performance Counters Reference.

# URI Cache

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The user-mode URI cache is known as the URI cache. The URI cache stores metadata about a URL. Metadata can include a variety of configuration data, including data about headers and authentication. This configuration data is stored in the IIS Admin service. The URI cache stores the metadata as an object. For each request that reaches IIS user mode, the worker process checks the URI cache to see if the object for the requested URL has been cached. If the object has not been cached, the worker process must call the IIS Admin service for the appropriate metadata. This call is expensive from a performance and scalability standpoint. If the URI cache has a stored object for the URL, the request pulls in the object and continues processing.

The IIS Admin service is responsible for URI cache change notification. If, for example, Default.htm changes, the IIS Admin service tells the worker process to flush the URI cache for that specific URL. Subsequent requests for that URL are cached in the URI cache. Table 7.3 describes the preferred or ideal values for URI cache counters. For more information about IIS 6.0 counters, see the Performance Counters Reference.

**Table 7.3 Preferred Values for URI Cache Counters**

| Object\Counter | Preferred or Ideal Value |
|---|---|
| Web Service Cache\URI Cache Hits | Depends on content. If content cannot be cached in the HTTP.sys response cache, the value of this counter should be as high as possible. If your content can be stored in the HTTP.sys response cache, the value of this counter should be low. |
| Web Service Cache\URI Cache Hits% | Depends on content. If content cannot be cached in the HTTP.sys response cache, the value of this counter should be as high as possible. If your content can be stored in the HTTP.sys response cache, the value of this counter should be low. |
| Web Service Cache\URI Cache Misses | Low is better. The value of this counter increases when IIS cannot locate requested content in either the HTTP.sys response cache or the IIS file cache, and must locate the requested content on the hard disk. If the value of this counter is low, responses are being sent from either the IIS file cache or the HTTP.sys response cache. |

# Token Cache

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When a request is made to the server, the security credentials for the request (or the configured anonymous user) are used to create a user token on the server. The server impersonates this user token when accessing files or other system resources. The token is cached in what is commonly called the *token cache*, so that the Windows logon only takes place the first time the user accesses the system or after the user's token has been flushed from the cache. If a token does not exist in the cache for an incoming request, IIS must call the Lsass.exe process to get the token. This call is expensive from a performance and scalability standpoint.

| **Note** |
| --- |
| Integrated Windows authentication tokens are not cached. |

The IIS worker process is responsible for flushing the token cache. The worker process monitors the **UserTokenTTL** registry entry for change notification (*TTL* stands for time to live). If the token has expired (the default time to live is 15 minutes) or if the token has changed in any way, IIS flushes the token cache. Currently, there are no performance counters that monitor the token cache.

# File Cache

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS file cache stores file contents in memory. If the file is larger than the 256 KB default maximum file size, IIS creates a file handle, passes the handle to HTTP.sys, and then closes the handle. The maximum file size is controlled by the MaxCachedFileSize registry key. For more information about the MaxCachedFileSize registry key, see Global Registry Entries.

The IIS 6.0 file cache uses an algorithm that attempts to cache only frequently requested files. The algorithm requires that a file be requested at least twice in a 10-second activity period, or the file will not be cached. In the case of a request for a Default.htm file, the activity-period algorithm works in this way:

1. The worker process receives the request for Default.htm.
2. The worker process checks to see if the file is in the IIS file cache table. At this point, the file has not been cached yet (that is, no cache entry is found).
3. The IIS file cache checks the activity table to see if Default.htm has been requested within the last 10 seconds. Because the file has not been requested, the file cache enters Default.htm in the activity table and serves out the file handle.
4. Three seconds later, the worker process receives another request for Default.htm.
5. The file cache does not find the cache entry in the cache table, but the activity table contains an entry that is within the 10-second activity period.
6. The file cache again serves out the file handle, but this time the file cache gets the cache entry from the worker process and stores the cache entry in the cache table (the cache table flushes all unused entries every 30 seconds).
7. A third request for Default.htm is served directly from the file cache.

You can configure the activity period (which is 10 seconds by default) by modifying the **ActivityPeriod** registry entry. If you set the value of **ActivityPeriod** to 0, IIS always caches files. The **ActivityPeriod** registry entry is located in the following subkey:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Inetinfo\Parameters

| **Caution** |
|---|
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference. |

Windows provides file cache change notification at the directory level. This means Windows monitors the parent file directory and informs IIS when a file in the directory or any subdirectory has changed. When IIS gets the change notification, the worker process flushes the file from the IIS file cache.

Table 7.4 describes the preferred or ideal values for file cache counters. For more information about IIS 6.0 counters, see the Performance Counters Reference.

**Table 7.4 Preferred Values for File Cache Counters**

| Object\Counter | Preferred or Ideal Value |
|---|---|
| Web Service Cache\Current Files Cached | Depends on content. A low value can indicate that HTTP.sys is caching a majority of static files, which is ideal. If the content files cannot be cached by HTTP.sys, a high value is good. |

| | |
|---|---|
| Web Service Cache\File Cache Flushes | Note that this counter value represents the number of times the cache was flushed, and not the number of files flushed. This value will continue to increase over time. If you edit files or change files often, this value will increase more rapidly. |
| Web Service Cache\File Cache Hits | Depends on content. If the content files cannot be cached by HTTP.sys, a high value is good. |
| Web Service Cache\File Cache Hits% | If the Kernel: Cache Hits % counter is low, this value should be high. If Kernel: Cache Hits % is low and this value is low, examine your file set and determine why your files are not being cached. Note that this counter does not include dynamic content, only static files. If your static files are not being cached, you might want to lengthen the activity period for this cache. |
| Web Service Cache\File Cache Misses | A *file cache miss* means the request for the file must go to the hard disk. File cache misses negatively impact performance and scalability. This value should be as low as possible. Also note that kernel cache hits will cause this value to be low. |
| Web Service Cache\Total Flushed Files | As low as possible. This value will increase if you edit or change content files, which results in a cache flush. |
| Web Service Cache\Maximum File Cache Memory Usage | The value of this counter and the value of Web Service Cache\Current File Cache Memory Usage should be the same. If the numbers are different, the cache was flushed. If this counter is significantly higher than Web Service Cache\Current File Cache Memory Usage, you might consider recycling the worker process because the application might have a memory leak. |
| Web Service Cache\Total Files Cached | Ideally, this value will be the same as Web Service Cache\Current Files Cached. If you edit or change your content files and subsequently flush the file cache, this value will decrease. |

Related Topics

For more information about editing the registry, see the Registry Reference.

# Http.sys Response Cache

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Kernel caching with the Http.sys response cache can be one of the most effective means of scaling and improving Web server performance. Cached responses are served from the kernel, which greatly improves response times and increases the number of requests per second that IIS can serve because requests for cached content never enter IIS user mode.

| **Note** |
| --- |
| The HTTP response cache is known as the kernel URI cache in System Monitor. |

For each request, the IIS worker process tells Http.sys whether or not to cache a response based on the activity-period cache algorithm. If a file is requested twice within 10 seconds (the default value for the **ActivityPeriod** registry entry) the IIS worker process tells Http.sys to cache the full response by URI. All subsequent requests for the cached response will be served from the cache. For more information about the activity-period cache algorithm, see File Cache.

Every 120 seconds, the Http.sys response cache runs a flush algorithm, which flushes cached files that have not been requested within the 120-second interval. The flush algorithm is also called when IIS receives change notification for a file — that is, when the file has been edited or changed in some way.

# HTTP.sys Cache Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table 7.5 describes the preferred or ideal values for HTTP.sys response cache counters. For more information about IIS 6.0 counters, see the Performance Counters Reference.

**Table 7.5 Preferred Values for HTTP.sys Response Cache Counters**

| Object\Counter | Preferred or Ideal Value |
|---|---|
| Web Service Cache\Kernel: Current URIs Cached | The value of this counter should nearly equal the total number of cacheable responses on the server (presuming that all cacheable responses have been cached). If this value is considerably lower than the total number of cacheable responses on the server, make sure the corresponding files are not larger than the default size, 256 KB. For information about the criteria that determine which files cannot be cached in the HTTP.sys response cache, see Events and Conditions That Disable HTTP.sys Response Caching. |
| Web Service Cache\Kernel: Total Flushed URIs | As low as possible. The value of this counter represents the number of times a response was flushed. Note that responses are flushed from the cache if the associated file is edited or changed in any way, or if the file is not requested within 120 seconds. |
| Web Service Cache\Kernel: Total URIs Cached | Depends on content, but the value of this counter should nearly equal the total number of cacheable responses on the server (presuming all cacheable responses have been cached). If this value is significantly larger than the total number of cacheable responses on the server, the HTTP.sys response cache is caching responses, flushing responses, and then caching responses again, which can impact performance and scalability. Investigate why responses are being flushed from the cache. If the value of this counter is significantly lower than the total number of cacheable responses on your sites, determine why some of the corresponding content files are not being cached. |
| Web Service Cache\Kernel: URI Cache Flushes | As low as possible. This counter denotes the number of times the flush algorithm has been called. Note that a response will be flushed from the cache if the corresponding file is edited or changed in any way, or if the file was not requested within 120 seconds. |
| Web Service Cache\Kernel: Cache Hits | As high as possible. This counter is incremented continually. If the value of this counter is very low, investigate why requests are not finding the cached response. |
| Web Service Cache\Kernel: Cache Hits % | Depends on content. The value of this counter should reflect the ratio of requests for cacheable content to requests for content that is not cacheable. If, for example, roughly half the requests to your sites are for cacheable content, the value of this counter should be close to 50 percent. |
| Web Service Cache\Kernel: URI Cache Hits/sec | As high as possible. |
| Web Service Cache\Kernel: URI Cache Misses | Depends on content. All requests for content that cannot be cached result in a cache miss. If the value of this counter is high and the value of Cache Hits is low, investigate why your responses are not being cached. |

# Events and Conditions That Disable HTTP.sys Response Caching

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

The HTTP.sys response cache will cache any response that has the appropriate flag in the request header; however, HTTP.sys will not cache the response if one of the following criteria applies:

| **Note** |
| --- |
| This cache is disabled on a per-request basis. |

- The request is not an anonymous request.
- The request required authentication (that is, there was an "Authorization:" header present in the request).
- The site is configured to use a footer.
- Dynamic compression is enabled and is used for the response. (Static compression can be used for HTTP.sys caching.)
- There is a query string in the request.
- The request verb was not GET.
- The cache is disabled (the **MD_VR_NO_CACHE** metabase property equals 1).
- The request has an entity body.
- Certificate mapping is enabled for the URL in question.
- Custom logging is enabled for the site in question.
- The request HTTP version was neither 1.1 nor 1.0.
- There was a "Translate: f" header in the request (that is, it was a Web Distributed Authoring and Versioning [WebDAV] request).
- There was an "Expect:" header that did not contain exactly "100 continue."
- There was either an "If-Range" or "Range" header in the request. HTTP.sys processes only whole responses; it does not attempt to send ranged responses.
- The response spanned multiple SendResponse/SendResponseEntityBody calls. A cacheable response must come down in a single, vectored SendResponse call.
- The total response size was greater than the per-response maximum size, which is set by using the **UriMaxUriBytes** registry key. The default value is 256 KB.
- The size of the response header was greater than the per-response maximum header size. The default maximum value is 1024 bytes.
- The cache is already full. The default size of the cache is proportional to the physical memory in the computer.
- The response is zero length.

- There is an ISAPI filter installed that is not cache-aware. By default, ISAPI filters are not cache-aware, but you can set the **FilterEnableCache** metabase property for the filter to make it cache-aware. All filters that are included with IIS 6.0, including FrontPage and ASP.NET filters, are cache-aware.

- A static file is accessed as a default document. For example, Default.htm exists in the root directory. Accessing the specific file by name (Default.htm) causes HTTP.sys to cache the file. Accessing the site by requesting the root folder (/), for example http://contoso.com/ with no file specified results in an uncached response. Note, however, that a request to a URL such as http://contoso.com/ will be cached by user-mode file caching in the worker process, as usual.

- The Web site has an IP Restriction list (applies only to Windows 2003 Server Service Pack 1 (SP1)).

# ASP Caching

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The process of creating compiled ASP script files (called *templates*) and script engines each time an ASP page is requested can consume memory and CPU resources, which in turn can impact performance and scalability. The following section describes how you can improve ASP performance and scalability by caching ASP templates and ASP script engines.

# ASP Template Caching

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP processes the templates or template files that contain ASP scripts. ASP stores these templates in a template cache and then serves the cached templates for subsequent client requests. Caching ASP templates enhances performance and scalability, because cached templates are not compiled each time they are called.

The ASP template cache uses a least-recently-used algorithm for determining which templates are cached. This means that if the cache is full, the template that has been in the cache and has not been requested for the longest amount of time is replaced by the next template to enter the cache. Every time a new template is added to the cache, the template goes to the beginning of the cache. The ASP template cache stores up to 500 templates by default (as determined by the **AspScriptFileCacheSize** metabase property). Templates are moved to the beginning of the cache when they are requested. In the case where the 501st template is bumped from the cache, the template is written to disk. By default, the ASP disk cache can hold 2,000 templates (as determined by the **AspMaxDiskTemplateCacheFiles** metabase property). If a template in the disk cache is referenced, the template is moved back to the ASP template cache.

In worker process isolation mode, **AspScriptFileCacheSize** and **AspMaxDiskTemplateCacheFiles** are global settings, which means they apply to all ASP pages on the Web server. (They are process-level settings in IIS 5.0 isolation mode.) ASP pages can contain #include files and these #includefiles can make the ASP files very large. The default 500-file limit established by the ASP template cache applies only to ASP pages and not #include files. However, ASP files that contain #include files can grow large enough to cause memory shortages and errors. ASP will flush the ASP template cache if IIS runs out of memory. If the ASP template cache is flushed, each subsequent request to an ASP page will need to be recompiled before it is cached again. In these situations, you might see ASP queuing. If you display the Active Server Pages\In Memory Templates Cached counter in System Monitor and see spikes followed by a severe reduction in the number of cached templates, your ASP pages might be too big. This could cause your system to run out of memory, flush the cache, recompile, and then run out of memory again. To remedy this situation, reduce the size of the ASP template cache to a number less than 500. Table 7.6 lists the preferred or ideal values for ASP template cache counters. For more information about IIS 6.0 counters, see the Performance Counters Reference.

**Table 7.6 Preferred Values for ASP Template Cache Counters**

| Object\Counter | Preferred or Ideal Value |
|---|---|
| Active Server Pages\In Memory Templates Cached | The value of this counter should reflect the number of frequently requested ASP files on your site (also known as your *hot set*). |
| Active Server Pages\In Memory Template Cache Hit Rate | As high as possible. |
| Active Server Pages\Templates Cached | The value of this counter should be close to the value of the **AspMaxDiskTemplateCacheFiles** metabase property; however, this value represents the total number of templates cached in memory and cached to disk, so this value could be higher than **AspMaxDiskTemplateCacheFiles**. |
| Active Server Pages\Template Cache Hits | As high as possible. |
| Active Server Pages\Template Cache Hit Rate | If this value is low and the Active Server Pages\In Memory Template Cache Hit Rate is high, a majority of user requests are being served from the memory cache. Performance and scalability should show a positive impact by not having to read the template from disk. If the value of Active Server Pages\Templates Cached equals the value of the **AspMaxDiskTemplateCacheFiles** metabase property, the cache has reached its maximum value. When the cache has reached its maximum value, this value should be high, otherwise the cache is being flushed. |

The ASP disk cache begins caching files into a directory when the ASP memory cache is full. Furthermore, a file has to be requested at least twice before it becomes a candidate for the ASP disk cache. The corresponding metabase property, **AspMaxDiskTemplateCacheFiles**, specifies the maximum number of compiled ASP templates that can be stored. **AspDiskTemplateCacheDirectory** contains the name of the directory that ASP uses to store compiled ASP templates to disk after overflow of the in-memory cache.

| **Important** |
| --- |
| In order for disk caching to work, ASP needs a properly configured disk cache directory in which to store the compiled ASP files. This means that the Administrators and IIS_WPG groups have Read/Write and Delete permission on the disk cache directory, and the disk cache directory is a local directory on the Web server. |

The default location of the disk cache is *systemroot*\System32\inetsrv\ASP Compiled Templates. In worker process isolation mode, subdirectories under the disk cache directory are created for each application pool. In IIS 5.0 isolation mode, the default disk cache directory is *windir*\system32\inetsrv\ASP Compiled Templates.

For information about changing the settings for caching ASP templates and creating an ASP disk cache directory, see Configuring ASP Template Caching.

# Script Engine Caching

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP code is compiled and executed by Microsoft® Visual Basic® Scripting Edition (VBScript) by default. To improve performance and scalability of ASP code, ASP creates script engines. ASP script engines contain the object and marshalling data that VBScript.dll needs to execute the request. Each entry in the ASP script file cache can point to one or more entries in the ASP script engine cache.

Table 7.7 describes the preferred or ideal values for the ASP Script Engine Cache counter. For more information about IIS 6.0 counters, see the Performance Counters Reference. For information about changing the settings for caching ASP script engines, see Configuring ASP Script Engine Caching.

**Table 7.7 Preferred Value for ASP Script Engine Cache Counter**

| Object\Counter | Preferred or Ideal Value |
|---|---|
| Active Server Pages\Script Engines Cached | The value of this counter should be greater than or equal to the value of Active Server Pages\Requests Executing. If this number is low, the server might not have enough memory to store the script engines in the cache. Performance and scalability will suffer because the server must create script engines often. |

# Setting the ASP Cache Metabase Property and Registry Entry

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following metabase property and Windows registry entry affect ASP caching on your server. In most cases, these settings should not require tuning or adjusting. If you adjust these settings, you should monitor performance and scalability after making changes.

**AspScriptFileCacheSize**
The AspScriptFileCacheSize Metabase Property specifies the number of precompiled script files to cache. If set to 0, no script files are cached. If set to 4,294,967,295, all script files requested are cached. This property is used to tune performance, depending on the amount of available memory and the number of script file cache hits.

The metabase represents unlimited as the DWORD value of 4,294,967,295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Earlier versions of IIS represented unlimited as -1.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference. |

**DisableLazyContentPropagation**
Windows registry path: HKLM\SYSTEM\CurrentControlSet\Services\ASP\Parameters

Data type: REG_DWORD

Default value: 0 (lazy propagation enabled)

Range: 0 - 1

*Lazy propagation* refers to the action that IIS takes when a large amount of content is updated at one time. IIS has an internal limit on the amount of content that can be updated in the in-memory template cache. If the size of the updated content exceeds that limit, IIS marks each of the files in the in-memory template cache as invalid. When it receives the first request to an invalid file, IIS begins to compile a new template, but it serves the expired template until the new template is compiled.

If you set the value of the registry entry DisableLazyContentPropagation to 1, IIS behaves as it does for IIS 5.0 and IIS 5.1 when a large amount of content is updated at one time. IIS flushes the in-memory template cache, and performance can slow to a standstill as each new request to the server forces IIS to compile new templates.

If you are developing Web pages and making changes to a few files at a time, you do not have to refresh your pages twice to see new changes because lazy propagation only applies when a large amount of content is updated at one time.

For information about the DisableLazyContentPropagation registry entry, see ASP Registry Entries.

**Related Topics**
For more information about editing the registry, see the Registry Reference.

# Global IIS Caching Registry Entries

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following registry entries affect caching on your server. In most cases, these settings should not require tuning or adjusting. If you find it necessary to adjust these settings, use caution and monitor performance and scalability after making changes.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference. |

**DisableMemoryCache**
Windows registry path: HKLM\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters

Data type: REG_DWORD

Default value: 0 (disabled)

Range: 0 - 1

Disables server static file caching, which negatively impacts performance and scalability. It is not recommended that you disable caching unless you are troubleshooting an issue or debugging application code.

**MemCacheSize**
Windows registry path: HKLM\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters

Data type: REG_DWORD

Default value: approximately one-half of available physical memory, in megabytes (MB)

Range: 0 - 2500 MB

Controls the size of the static file cache. The default value is dynamically adjusted every 60 seconds.

**ObjectCacheTTL**
Windows registry path: HKLM\CurrentControlSet\Services\InetInfo\Parameters

Data type: REG_DWORD

Default value: 30 (seconds)

Range: 0 - 4,294,967,295 (unlimited)

Controls the TTL setting, which defines the length of time that objects are held in the cache. If an object in the memory cache has not been referenced for the specified period, that object will be phased out of the cache. If system memory is limited, or the contents of the server are dynamic, you can use a lower TTL to prevent system memory from being used to cache a large number of volatile objects. Setting **ObjectCacheTTL** to 0xFFFFFFFF disables the *object-cache scavenger* (the process that routinely reads the items in the cache to determine which items should stay in the cache and which items should be flushed) and allows cached objects to remain in the cache until they are overwritten. Disabling the cache scavenger is useful if your server has ample system memory and your data is relatively static.

**Related Topics**
For more information about these registry entries, see Global Registry Entries.

For more information about editing the registry, see the Registry Reference.

# ISP Scaling -- Strategies for Hosting Thousands of Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The ability for IIS to scale up to thousands of sites on one server depends on the type of content the site hosts and the number of application pools on the server. Your business model and service agreements with customers will also determine how many sites you can put on one server. There are three common Internet service provider (ISP) installation strategies you can use when IIS is running in worker process isolation mode:

- Shared static hosting
- Shared static and dynamic hosting
- Dedicated hosting

# Shared Static Hosting

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The *shared static hosting* strategy involves collecting large numbers of static sites together on one server. These static sites are typically personal home pages or business pages that experience very low traffic volumes on a daily basis, so performance and scalability issues are rarely an issue. Before provisioning your server, consider the following guidelines.

**Administer IIS by using command-line administration scripts or batch files**
Using IIS Manager to perform tasks such as provisioning thousands of new sites or configuring sites is extremely inefficient. If your server hosts a few thousand sites, IIS Manager can take several minutes to refresh as it reads configuration data from the metabase. In this case, it is more efficient to administer IIS by using command-line tools. For information about command-line administration tools in IIS 6.0, see IIS 6.0 Administration Scripts, Tips, and Tricks, or see Using Command-Line Administration Scripts. For general information about Web site and virtual directory configuration, see Configuring Internet Sites and Services.

**Create a logical folder structure**
A logical folder structure, for example Domain/username, will alleviate some of the administration overhead in locating sites when they require configuration changes.

**Use host header names to create multiple sites on one server**
Static IP addresses are in limited supply, and obtaining thousands of static IP addresses can be problematic. Static IP addresses are only necessary when a site requires Secure Sockets Layer (SSL). Also, static IP addresses have performance overhead costs. The WWW service must manage an endpoint for every site identified by a unique IP address, and this consumes memory from the nonpaged pool.

**Put all static sites in one application pool**
Static sites, if they are truly static, pose little or no crash risk for an application pool because static sites do not execute code. By putting the sites in one application pool, you limit the resource overhead that multiple application pools impose. You also limit the administrative overhead that comes with multiple application pools.

**Run all static sites with the same anonymous user account**
For static, public Web sites, the anonymous user account provides ample security. Before returning a page to the client, IIS checks NTFS file and directory permissions to see whether the IUSR_*computername* account is allowed access to the file. If access is not allowed, IIS attempts to use another authentication method. If none is selected, IIS returns an HTTP 403 error (Access Denied) to the browser.

**Recycle the application pool worker process (optional)**
You will need to monitor site traffic to determine a regular recycle schedule, if a regular recycle is needed. Ideally, you should recycle the worker process during low-traffic times, for example 3:00 A.M. local time. For more information about worker process recycling, see Running IIS 6.0 as an Application Server.

**Monitor startup times**
Although you should not need to restart your server very often, monitor startup times when you do restart the server. When an IIS server hosts thousands of Web sites, the size of the metabase grows significantly and the task of reading configuration data from the metabase while IIS is starting can strain your CPU resources. Your startup times might be a determining factor for when you need to add a new server to your installation.

**Monitor performance**
As you provision sites, regularly monitor site traffic and resource consumption. In this way, you will be able to see how the system scales and proactively determine when it is time to add a new server to your installation.

**Enable centralized binary logging**
Centralized binary logging writes all log file data in binary format to one centralized file, and preserves memory and CPU resources by not creating and writing to thousands of individual log files. For more information about centralized binary log files, see Analyzing Log Files.

**Partition log files on multiple disks**
If you choose not to use centralized binary logging, partition log files on multiple disks to minimize I/O. With multiple disk controllers writing log data to different disks, you will minimize I/O and preserve memory resources for other needs.

# Shared Static and Dynamic Hosting

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The *shared static and dynamic hosting* strategy mixes static and dynamic sites on the same server. If your customer base requires you to host a mix of static and dynamic sites on the same server, the following guidelines can help you maximize the number of sites and application pools you can host on one server. Most of the guidelines in Shared Static Hosting apply to this strategy as well.

**Separate static and dynamic content sites in different application pools**
If you need to host static and dynamic sites on the same server, put the static sites in one or more application pools and the dynamic sites in one or more different application pools. By isolating the static sites from the dynamic sites, you mitigate the possibility of a poorly developed application crashing the static content application pool.

**Set content quotas for sites**
Your business model and customer service agreement should define content quotas, for example 200 MB for each site. Set disk quotas in the properties of each physical hard disk by using the Quota service available in Windows Server 2003.

**Consider enabling Quality of Service (QoS) features for sites that are resource-intensive**
Features such as CPU monitoring, bandwidth throttling, and idle timeout will ensure that one site does not consume a majority of the resources and cause errors for other sites on the same server. For more information about these features, including how to enable them, see Optimizing IIS 6.0 Performance.

**Configure one unique anonymous user per site**
Configuring one unique anonymous user per site adds a security layer by ensuring that users cannot access content on sites to which they are not permitted access.

**Develop processes for handling service packs, security updates, and hotfixes**
If your customers will be using IIS 6.0 to administer their own dynamic sites, set up a process for handling service packs, security updates, and hotfixes to minimize downtime for your customers and to help ensure the security of your installation. Windows Server 2003 has improved management of updates with Auto Update version 1.0, which provides you with three options:

- Notify you the moment an update is available.
- Download the update and notify you of its availability.
- Scheduled installation (this option enables the update to be downloaded and automatically installed at a time you choose).

**For ASP content, allow ASP pages to run on MTA instead of STA**
Because of the overhead memory and CPU cost of running dynamic applications, your server will reach a limit in the number of application pools it can host. With ASP content, you can mitigate this limitation on application pools by enabling ASP pages to run on multithreaded apartment (MTA) threads instead of single-threaded apartment (STA) threads. ASP is capable of running all of its threads in an MTA. If your COM components are primarily free-threaded or both-threaded, running the ASP threads as MTA can improve performance significantly. By default, the **AspExecuteInMTA** metabase property is set to 0, which means that ASP does not execute in MTA. Set this property to 1 at the application level to enable ASP to run in MTA.

| **Note** |
| --- |
| MTA has some limitations with respect to components that can be instantiated. Objects that are strictly Apartment might show a performance regression because marshalling will occur and the potential for STA thread contention will increase. |

**If you need to set up more than 60 application pools where each pool is running as a unique identity on a server, change the UseSharedWPDesktop registry key**
If you are setting up your application pools with unique identities, depending on the applications and memory resources on your server, you will reach a limit of about 60 application pools. There are finite limits to some system resources that get allocated with each new logon session. This means that 60 processes can run concurrently as distinct accounts. IIS 6.0 supports

running these processes in a single shared workstation and desktop, at a cost of sharing a single encapsulation of a user session among all parties.

| ⚠Caution |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference. |

To scale beyond 60 application pools and to share a single desktop, change the value of the registry entry UseSharedWPDesktop, of data type DWORD, to 1. The registry entry is located in the subkey HKLM\System\CurrentControlSet\Services\W3SVC\Parameters\UseSharedWPDesktop. After changing this registry entry, you should be able to scale to hundreds of application pools and hundreds of concurrently running worker processes.

**If IIS is starting too many worker processes and your server is running out of resources, set a maximum number of concurrent worker processes**

With the IIS 6.0/HTTP.sys architecture and with a request load distributed across many application pools, IIS might attempt to start more worker processes than the server can sustain without running out of memory or CPU resources. The DemandStartThreshold Metabase Property (also called *concurrent process gating*), when set to a value less than the default, applies a hard limit to the number of IIS worker processes that can run concurrently. If the hard limit is exceeded, HTTP.sys returns an HTTP 503 error (Service Unavailable).

**Locate misbehaving applications and isolate them in their own application pool**

You can use Event Viewer and IIS logs to locate misbehaving applications, or you can use Log Manager and Trace Report, which are general-purpose tools supplied with Windows Server 2003. Log Manager creates trace data, which is then processed by the Windows Trace Report utility. The Trace Report utility creates a formatted, detailed report that is useful in assessing IIS and IIS-related activity in the operating system. Once you have located a misbehaving application, isolate the application in its own application pool so it doesn't affect other applications. Recycle the application pool regularly.

**Related Topics**

For more information about editing the registry, see the Registry Reference.

For more information about Log Manager and Trace Report, see The IIS Request-Based Tracing Process.

For more information about viewing and interpreting log files, see Analyzing Log Files.

# Dedicated Hosting

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Dedicated hosting is typically reserved for those customers who have the greatest availability, security, and performance requirements for their dynamic Web sites. In some cases, you might put only one or two sites on an IIS server to ensure there are enough memory and CPU resources to meet customer requirements. In other cases, dedicated hosting means the customer pays to have their own server or servers. Dedicated hosting can be configured in the following ways.

**Configure one unique application pool per Web site**
Process boundaries separate each worker process so that problems in one application pool do not affect other application pools on the server. By configuring each Web site in a different application pool, you ensure greater reliability, availability, and security. For information about how to configure application pools, see Running IIS 6.0 as an Application Server.

**Configure one unique anonymous user per site**
Configuring one unique anonymous user per site adds a security layer by ensuring that users cannot access content on sites to which they are not permitted access.

**Configure one unique user to run as the process identity of the application pool**
Configuring one unique user to run as the process identity of the application pool adds an additional layer of security by ensuring that one site cannot access content on another site. If you specify a unique user account, be sure to add that account to the IIS_WPG group. When you use a built-in account, try to use Network Service because it offers the best balance between security and functionality.

If you experience scalability issues on a dedicated hosting computer, see Shared Static and Dynamic Hosting.

# Improving Scalability Through UNC-Based Centralized Content Storage

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 delivers several technical changes that make it easier for ISPs or other organizations to create UNC-based centralized storage solutions. Centralized content storage can make it much easier to scale your installation to thousands of Web sites. Centralized storage can also help minimize time commitments for server management and administration. This section describes the following:

- Caching Changes. For UNC content, IIS 6.0 provides the new last-modified caching algorithm for both static and ASP files. In addition, IIS 6.0 improves the performance of change-notification-based caching for ASP files.

- UNC Authentication. IIS 6.0 now performs UNC authentication, formerly called Passthrough authentication, by default.

- Constrained Delegation. With constrained delegation, you can allow delegation only to certain services, so you can control the specific network resources that a service or computer can use.

# Caching UNC-based Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you can create UNC-based virtual directories, the IIS caches must treat each UNC path as a separate directory structure, and therefore ASP and static file caches must monitor each directory individually. In a centralized-content architecture with many different UNC shares, this has the potential to exhaust system resources.

The Web server and the file server are connected with a Server Message Block (SMB) connection. Only one SMB connection is created between a Web server and a file server; however, within that SMB connection, there are a finite number of work items. Work items can be consumed in a variety of ways and for varying amounts of time. For example, performing an operation — such as a **CreateFile** or **GetFileAttributes** command — on a file consumes an I/O work item for a short amount of time, but asking for a change notification on a directory structure consumes a work item for as long as the connection is intact. Each change-notification instance uses a work item until the connection is lost, and each request from a Web server to the file server uses a work item only temporarily. However, the scalability impact of the SMB connection is reduced by using last-modified caching.

IIS 6.0 has two methods for caching static and ASP files, called last-modified caching and change-notification caching. For information about configuring change-notification tracking or the time to live (TTL) for static and ASP files, see Configuring Last-Modified Caching and Configuring ASP Change-Notification Caching.

# UNC-based Caching Considerations

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You should decide which caching method to use based on your network architecture and your scalability and performance requirements. In general, last-modified caching is the most reliable method. It is recommended that you use this caching method when you require scalability for a system with a Web server that points to different file structures. You can also experience performance gains by using change-notification caching if your system uses a file server that reliably reports change notifications and you have only a few sites or virtual directories. For example, if you have a single Web site with one or two applications, and all of your content is stored on a Windows Server 2003–based file server, use change-notification caching.

To configure caching for optimum scalability and performance, you must edit the registry entries **MaxMpxCt**, **MaxWorkItems**, and **MaxCmds**. On the server running IIS, these entries are stored in the subkey HKLM\System\CurrentControlSet\Services\LanmanWorkstation\Parameters. On the remote file server running Windows Server 2003, these entries are stored in the subkey HKLM\System\CurrentControlSet\Services\LanmanServer\Parameters.

**MaxMpxCt**

The **MaxMpxCt** registry entry permits a server to provide a specified maximum number of simultaneous client requests to itself, and then enforces that limit. Because each client connection generates multiple instances (for example, Windows shell and Explorer.exe), with multiple clients connecting to the server, the number of client requests can be greater than the default **MaxMpxCt** value of 50. As soon as this limit is reached for that server, additional requests are refused until the number of requests drops below the limit again.

**MaxWorkItems**

The value of the registry entry **MaxWorkItems** should be at least four times as large as that for **MaxMpxCt**. For example, if **MaxMpxCt** has a value of 1,024, **MaxWorkItems** should have a value of at least 4,096.

Some resources, including Microsoft Knowledge Base articles, recommend setting the value of **MaxWorkItems** to 16,000. However, this could consume up to 250 MB of nonpaged pool memory, which could starve the system. To increase your systems reliability, it is recommended that you set the value of **MaxWorkItems** lower than 16,000.

To determine the number to set for the value of **MaxWorkItems**, monitor the Server object performance counters. The Work Item Shortages counter will indicate whether you need to make these values larger or increase the load capability of your file server by adding memory or clustering the file server. You must allow enough work items for two activities: change notifications and normal I/O requests.

**MaxCmds**

The **MaxCmds** registry entry sets the limit for the number of concurrent outstanding network requests between the SMB client and server that the client can support. The general formula for calculating the value of **MaxCmds** is as follows:

(number of distinct physical directories that IIS needs to monitor for change notifications) (1 (if static files exist) + 1 (if ASP content exists) + 1 (if ASP.NET content exists) + 50 (for concurrent default/regular file I/O)

This formula is merely a starting point for tuning your servers. You might need to increase your settings depending on concurrent traffic. For better scaling, you typically need to modify the values of the **MaxCmds**, **MaxMpxCt**, and **MaxWorkItems** registry entries on both the server running IIS and the remote file server. Before you modify the entries, however, watch the Server object performance counters while you stress-test your system. Carefully monitor the Work Item Shortages counter to see if you need to increase the values of the entries or increase the load capability of your file server by adding memory or clustering the file server. If you are using the non-default change-notification option for caching, you must allow enough work items for two activities: change notifications and normal I/O requests for the file.

# UNC-based Caching Scenarios

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following three scenarios illustrate the different UNC–based caching options in more detail and provide recommendations for setting maximum values:

- Wide content, low traffic. This scenario is an example of a typical low-end shared hosting server.
- Wide content, high traffic. This scenario represents an enterprise Web site that hosts hundreds of small applications and receives a high volume of traffic consisting mainly of anonymous users.
- Narrow content, high traffic. This scenario represents a dedicated hosted Web site with fewer than 10 applications, but with a very high volume of traffic consisting only of anonymous users.

When evaluating the settings in the tables in the following sections, consider the number of simultaneous work items that are opened from the server running IIS to the remote file server, but reflected (inversely) by the value of the **MaxCmds** entry. You should calculate all other remote file server settings based on that result.

| **Important** |
| --- |
| The settings in the following scenarios are intended as guidelines only. Changing the registry settings listed here will produce different results, depending on your specific system configuration. You should test your application with the default settings, and then increase the maximum settings accordingly. |

# Scenario 1: Wide Content, Low Traffic

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In this scenario, each site has its own directory structure, and there are no virtual directories below each site. Each server running IIS might have up to 20,000 sites, but only approximately 1,000 sites are active at any given time. Each site runs in its own process. Users are authenticated if they are publishing on the Web; otherwise, they are anonymous. The content is primarily ASP and static files, with a few ASP.NET applications.

If you enable file caching with the change-notification algorithm, you might use the following equation to calculate **MaxCmds**:

20,000 physical directories (1(static) + 1(ASP) +1(ASP.NET) + 50 = approximately 60,000 for **MaxCmds**

However, this setting would jeopardize server performance and scalability because it would consume hundreds of megabytes of nonpaged pool memory.

If you use caching with the default last-modified algorithm, no work items are used for change notification. Depending on active traffic for file I/O and your analysis of the performance counters, you might still need to increase the values of **MaxCmds**, **MaxWorkItems**, and **MaxMpxCt** from their default values. Table 7.8 lists the recommended settings for the wide-content, low-traffic scenario that uses the default last-modified caching algorithm.

**Table 7.8 Recommended Registry Settings for Wide Content, Low Traffic**

| Registry Entry | Server Running IIS | Remote File Server |
|---|---|---|
| MaxCmds | 3,072 (0xC00). This is the maximum number of work items allowed by this computer. | Default setting. |
| MaxMpxCt | Default setting. | 3,072 (0xC00). This is the maximum number of work items allowed per client computer. The client computer in this case is the server running IIS. |
| MaxWorkItems | Default setting. | 12,288 (0x3000). This value must be four times MaxMpxCt. |

# Scenario 2: Wide Content, High Traffic

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In this scenario, a large Web site has 300 applications. The site contains static files, ASP content and ASP.NET content. Most users access the site anonymously.

If you enable file caching with the change-notification algorithm, this site will consume approximately three work items per cache, per application, for a total of 900 work items.

If you enable file caching with the last-modified algorithm, the site will not consume any work items that would be consumed for long periods, but there will be temporary work items and heavy traffic between the Web server and the file server. The caches are all frequently used.

Because there are a manageable number of applications in this scenario that use less than 1,000 work items per client, you might attempt to use the change-notification caching algorithm. This would reduce the back-end network traffic.

Table 7.9 lists the recommended settings for the wide-content, high-traffic scenario that uses the change-notification algorithm.

**Table 7.9 Recommended Registry Settings for Wide Content, High Traffic**

| Registry Entry | Server Running IIS | Remote File Server |
|---|---|---|
| MaxCmds | 2,048 (0x800). This is the maximum number of work items allowed by this computer. | Default setting. |
| MaxMpxCt | Default setting. | 2,048 (0x800). This is the maximum number of work items allowed per client computer. The client computer in this case is the server running IIS. |
| MaxWorkItems | Default setting. | 8,192 (0x2000). This value must be four times MaxMpxCt. |

**Note**

The value listed in Table 7.9 for MaxCmds is higher than the value that results when you use the general formula for calculating MaxCmds because the value in the table was rounded.

If you have too many applications to use the change-notification algorithm, you can use the last-modified algorithm. You should extend the cache TTL value to approximately 30 seconds to reduce the traffic between the Web server and the file server. This scenario is the most difficult to tune properly.

# Scenario 3: Narrow Content, High Traffic

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In this scenario, a small site contains 10 applications and static, ASP, and ASP.NET files that receive a very high number of hits.

If you enable file caching with the change-notification algorithm, this site will consume approximately three work items per cache, per application, for a total of 30 work items.

If you enable file caching with the last-modified algorithm, the site will not consume any work items, but there will be heavy traffic between the Web server and the file server. The caches are all frequently used.

Because this scenario requires only a few work items to cache the content reliably, change notification is the recommended caching algorithm. You do not need to increase the default maximum settings by very much. Table 7.10 lists the recommended settings for the narrow-content, high-traffic scenario that uses the change-notification algorithm.

**Table 7.10 Recommended Registry Settings for Narrow Content, High Traffic**

| Registry Entry | Server Running IIS | Remote File Server |
|---|---|---|
| MaxCmds | 512 (0x200). This is the maximum number of work items allowed by this computer. | Default setting. |
| MaxMpxCt | Default setting. | 512 (0x200). This is the maximum number of work items allowed per client computer. The client computer in this case is the server running IIS. |
| MaxWorkItems | Default setting. | 2,048 (0x800). This value must be four times MaxMpxCt. |

**Note**

The value listed in Table 7.10 for MaxCmds is higher than the value that results when you use the general formula for calculating MaxCmds because the value in the table was rounded.

# Authentication for UNC Shares

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The UNC authentication method, also called UNC Passthrough authentication, determines the credentials to be used for gaining access to a UNC share on a remote computer. In IIS 6.0, UNC authentication works by looking at the request user and the credentials stored in the UNCUserName Metabase Property and UNCPassword Metabase Property to determine the credentials to pass through to the computer with the UNC share.

If you use IIS Manager to create a Web site or virtual directory and specify a UNC path for the content, IIS Manager prompts you to type a user name and password for the network resource. If you specify a user name and password, the **UNCUserName** and **UNCPassword**metabase properties are set.

If **UNCUserName** is specified (not empty) and **UNCPassword** is valid, the metabase user credentials are sent as the user identity for access to the remote share. If **UNCUserName** is specified (not empty) and **UNCPassword** is not valid, a 500.13 error (Internal Server Error: Invalid Username or Password) is sent to the client.

If **UNCUserName** is empty, the credentials of the request user — which can be either an authenticated set of credentials for authenticated requests or IUSR_*computername* credentials for anonymous requests — are sent as the user identity for access to the remote share. In other words, by default in IIS 6.0, you do not need to specify user name and password credentials.

| **Note** |
| --- |
| The **UNCAuthenticationPassthrough** metabase property is no longer used for UNC authentication. |

For more information about enabling UNC authentication, see Securing Virtual Directories.

# Constrained Delegation for UNC File Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Delegation* is the act of allowing a service to impersonate a user account or computer account in order to access resources throughout the network. When a service is trusted for delegation, that service can impersonate a user to use other network services. *Constrained delegation* is a new option for Windows Server 2003, and it is recommended that you use it with a Windows Server 2003 domain. With this option, you can specify the Service Principal Names (SPNs) to which an account can delegate. With this option, a service can be trusted for delegation, but the domain administrator can limit that trust to a select group of explicitly specified services. By only allowing delegation to specific services, you can control the specific network resources the service or computer can use.

Constrained delegation is particularly useful in scenarios where a site that requires authentication — that is, a site that does not allow anonymous access — contains content that is housed on a remote UNC file server. With constrained delegation, you can enable Windows Integrated authentication, which can use NTLM authentication or send credentials across the network as a Kerberos-based token.

If you do not use constrained delegation, but you enable Windows Integrated authentication and NTLM, the token that the Web server obtains from the Windows security infrastructure does not have sufficient permissions to access another computer, such as your file server. However, with constrained delegation and Windows Integrated authentication, the token received by the Web server from the Windows security infrastructure is a Kerberos-based token with permission to access other computers, including the file server. Essentially, constrained delegation allows an NTLM-based token to be upgraded to a Kerberos-based token. Be aware, however, that Kerberos-based authentication can degrade performance because each access check occurs on the file server.

For more information about constrained delegation, see Managing a Secure IIS 6.0 Solution.

When you set up a UNC-based virtual directory and specify a user name and password for that directory, you might receive an error message when you attempt to access the directory by using IIS Manager. This error occurs because IIS Manager sends the current Windows user credentials when it accesses the virtual directory. The credentials of the current Windows user might not be the same as those you specified for the UNC-based virtual directory. However, if you access your Web server over the Internet by using IIS Manager on a remote computer, the contents of the virtual directory will be displayed without error. Because using IIS Manager is not a reliable method of verifying whether your customers can access the content stored on the remote computer, you should test your virtual directories by making requests with a Web browser.

# Case Study: Scaling an ASP.NET Web Application on IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following case study illustrates how to monitor and tune IIS 6.0 for improved performance and scalability. In this study, network administrators in a fictional company, Contoso, were getting ready to roll out a new Web application for the purpose of selling books online. The application code had been extensively tested, but performance and scalability statistics were still unknown. To gather performance and scalability metrics, the administrators used the Microsoft Web Capacity Analysis Tool (WCAT) to stress-test the site.

Stress testing helps administrators locate performance and scalability bottlenecks in their Web application installation. By using a tool such as WCAT or the Microsoft Application Center Test (ACT) in a test environment to simulate the server load they can expect in a production environment, administrators can isolate and fix potential bottlenecks that might undermine the customer experience. Administrators can also plan the installation capacity.

The administrators in this study were preparing to roll out the Duwamish 7.0 ASP.NET application, a fictitious online bookseller application. Their rollout consisted of the following phases:

- Capacity planning
- Tuning production servers
- Scaling up
- Scaling out

| **Note** |
| --- |
| This case study does not provide information about how to create a WCAT script or how to perform stress testing by using WCAT. See the WCAT documentation (Wcatguide.doc) for information about how to create a WCAT script and how to use WCAT for stress testing. |

Related Topics

For more information about WCAT, see the IIS 6.0 Resource Kit Tools.

# Capacity Planning_IIS_SP1_TechRef

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

The administrators initially wanted to see how their application scaled on a server with four 500-MHz processors running IIS 6.0 in worker process isolation mode. Their goal was to stress the site to see how many requests per second the installation processed at (or close to) 100 percent CPU saturation. By understanding how the application performed when the CPU was saturated, they could begin to understand how the site would scale under a heavy load, and they could begin to plan the capacity of the installation based on traffic estimates. Table 7.11 describes the software and hardware used for this case study.

**Table 7.11 Case Study Software and Hardware**

| Software | Hardware |
|---|---|
| Web Servers:<br>Windows Server 2003 (RC1)<br>Default installation of IIS 6.0 (running in worker process isolation mode)<br>ASP.NET<br>Duwamish.NET 7.0 Web application | Web Servers:<br>Four 550-MHz Intel Pentium 3 processors<br>Compaq Proliant 6400R<br>512 MB RAM<br>Single 3COM Gigabyte Ethernet Server network adapter<br>Eight 900-MHz Pentium 3 processors<br>Compaq Proliant 6400R<br>8 GB RAM<br>Intel Pro Gigabit Optic Fiber Server network adapter |
| Database:<br>SQL Server 2000 | Database:<br>Two 1000-MHz Pentium 3 processors<br>Compaq DL360<br>264 MB RAM |

# Preparing to Test

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The administrators set up their WCAT configuration file to run 40 virtual clients (four client computers running 10 threads each). They set up their WCAT distribution file to imitate real-world customer use of their site. Table 7.12 shows the breakdown of the WCAT distribution file.

**Table 7.12 Description of WCAT Distribution File Classes for Contoso Stress Tests**

| Class Number | Percentage of Requests | Class Description |
|---|---|---|
| <class> 1 | 60 | User browses for books. The administrators assigned the highest percentage of requests to this action to reflect what they expect to be the highest percentage of real-world traffic to their site. |
| <class> 2 | 10 | User creates an account and buys a book. |
| <class> 3 | 20 | User logs on to an existing account, puts five books in the shopping cart, takes two books out of the shopping cart, and then completes the purchase. |
| <class> 4 | 10 | Invalid searches and other user mistakes. |

The administrators ran the stress test for 10 minutes. If they had wanted to locate bugs or bottlenecks in their Web application, they would have run the test for much longer; however, because they were not concerned about bugs in the application, the 10-minute test run was sufficient.

While the test ran, the administrators monitored the network to make sure the connection (with a capacity of 1 GB) was not saturated. In Windows Task Manager, on the **Networking** tab, the administrators noted that network traffic was steady at about 4 percent, which meant that the network was not a bottleneck.

The administrators then opened System Monitor and monitored the Processor(_Total)\% Processor Time counter. The counter showed that the CPU was saturated at 100 percent. The administrators ran the test several more times with fewer virtual clients and the CPU did not reach 100 percent saturation, so they concluded that the data from the test that ran 40 virtual clients would suit their needs for capacity planning, provided the WCAT log did not reveal test errors.

The log revealed no errors, so the administrators concluded the bottleneck is the CPU. The server installation could not scale beyond the 40 virtual clients because the server did not have sufficient processing power. If the administrators had failed to saturate the CPU, they would have looked elsewhere in the installation for the bottleneck — for example, disk I/O (reading and writing from disk) or database activity. Also, a badly written application can be the source of a bottleneck in cases where the CPU is not saturated during stress testing.

# Interpreting Test Data

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In a situation where the CPU constrains maximum performance, the WCAT results can be used to perform a cycles-per-request calculation for capacity planning. The WCAT log showed 893 requests per second on a server with four 550-MHz processors. Test results can be summarized as follows:

- 4 x 550 megacycles/second = 2,200 megacycles/second.
- 2,200,000,000 cycles/second divided by 893 requests/second = 2,463,606 cycles/request (on average). Rounded up, the server used 2.5 million cycles/request while the CPU was at 100 percent saturation.
- At this point, the administrators spoke with the business managers and marketing managers to learn about expected traffic loads per day. The managers explained that after extensive research, they expected 200,000 customers per day. Each customer will send roughly 60 requests per session, so the total requests/day is 60 200,000, or 12,000,000 total requests/day (estimated).
- 12,000,000 divided by 24 60 minutes 60 seconds = 139 anticipated requests per second.
- Rounding to 150 requests/second 2.5 megacycles per request = 375 MHz needed to process the expected traffic at any given time (on average).

The administrators did not want their Web server to run at 100 percent CPU utilization. Furthermore, the administrators knew from experience that all Web sites have a pattern to their traffic. In order to accurately assess their capacity needs, they would need to know how much CPU power would be needed to handle traffic spikes. After researching other online booksellers, the administrators determined that their site would be busiest in the evening hours from 6 P.M. to midnight (local time) and that their peaks would be around 300 requests per second. The administrators also knew that according to their company's policy, maximum CPU utilization during peak loads should rarely exceed 30 percent. This maximum helps to assure quality of service during times of heavy traffic.

With a goal of 300 requests per second and a maximum CPU utilization of 30 percent, the administrators adjusted their WCAT distribution and configuration files and then ran the stress test again. The tests showed that at 275 requests per second, CPU utilization approached 30 percent. In this way, the administrators determined that their current capacity on the server with four 550-MHz processors is sufficient for expected regular traffic and peak loads. In other words, the installation scales well for their current needs, and the installation should have sufficient resources available to handle peak loads.

# Tuning Production Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Contoso best practices mandate that each server be configured for optimum performance and scalability before the server is put into production. These best practices specifically target disk setup, services running on the server by default, network setup, and caches.

**Use RAID**

To begin with, the administrators set up the site content on a redundant array of independent disks (RAID). RAID is a data storage method in which data — along with information used for error correction, such as parity bits — is distributed among two or more hard disk drives to improve performance and reliability. For information about setting up RAID, see Help and Support Center for Windows Server 2003.

If using RAID is not an option, consider moving the site content to a disk other than the system disk (that is, the disk on which the operating system is installed). By storing your site content on a disk other than the system disk, you reduce contention for disk access between the operating system and the Web sites. This further increases performance.

**Delete unused default documents**

The administrators knew that if a user tried to access the Web site by using a trailing slash — for example, http://www.contoso.com/website/ — the server would read each item listed in the default document list (such as Default.htm, Default.asp, and Index.htm) and perform a search to see if each document existed. As a best practice, the administrators deleted all the documents in the default document list that do not exist. They left only the default document that does exist — for Contoso, this is Default.aspx. You can delete unnecessary default documents by using IIS Manager to modify the **Documents** tab on the **Web Sites** property pages.

**Shut down unused services**

To conserve memory and CPU resources for the Web service, the administrators shut down unnecessary services. The server will host only the Contoso Web application, so the administrators stopped other IIS services such as the FTP, Simple Mail Transfer Protocol (SMTP), and Network News Transfer Protocol (NNTP) services. Using the Services snap-in, the administrators stopped the Automatic Updates, Task Scheduler, Wireless Configuration, and Print Spooler services (because they would not be printing from that server).

| **Important** |
| --- |
| Refer to Help and Support Center for Windows Server2003 to ensure that you understand the role of each service before you stop it. Monitor performance after stopping each service to ensure that the change does not negatively affect performance. |

**Change file and printer sharing defaults**

The administrators changed the default File and Printer Sharing for Network properties from **Maximize data throughput for file sharing** to **Maximize data throughput for network applications**. This small change resulted in a 10 percent increase in throughput. For information about how to change these properties, see To configure File and Printer Sharing for Microsoft Networks.

**Related Topics**

For more information about performance tuning, see Optimizing IIS 6.0 Performance.

# Scaling Up

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Scaling up* means adding hardware, such as RAM or CPUs, to a Web server to increase the number of sites a server can host. Scaling up can be an inexpensive way to boost performance, especially if you host static-content Web sites.

Shortly after moving the server with four 550-MHz processors into production, the Contoso Marketing department finalized plans for a major television advertisement campaign. The marketing director informed the administrators that site traffic would likely double in the next few weeks. The administrators needed to determine whether the server with four 550-MHz processors could handle sustained traffic loads of 600 requests per second while staying within the 30 percent CPU restriction. The cycles-per-request analysis is as follows:

- 4 x 550 megacycles/second = 2,200 megacycles/second.
- 2,200,000,000 cycles/second divided by 893 requests/second = 2,463,606 cycles/request (on average). Rounded up, the server performed 2.5 million cycles/request while the CPU was at 100 percent saturation.
- 600 requests/second 2.5 million cycles/request = 1500 MHz.
- 1500 MHz divided by 30 percent CPU saturation maximum (1500/.3) = 5000 megacycles/second.

As a result of this analysis, the administrators realized they needed a server that could perform 5,000 megacycles/second in order to use 1,500 megacycles/second at 30 percent CPU. The server with four 550-MHz processors could only achieve 2,200 megacycles/second. Consequently, they would need to scale up, scale out, or both.

To determine the results of scaling up, the administrators tested a server with eight 900-MHz processors. This new server was capable of 7,200 megacycles/second, which covered their current needs.

The administrators migrated the application and performed their WCAT stress test again. In an effort to saturate the CPU at 100 percent, the administrators ran the test with 48, then 72, and finally 120 virtual clients. CPU utilization never exceeded 70 percent. The administrators checked Task Manager and found that the network hovered around 11 percent. The administrators determined at this point that they had a bottleneck somewhere in the system, but for the purpose of their tests, the result of 1,904 requests/second with an average response time of 630 milliseconds was sufficient test data.

The server with four 550-MHz processors performed 893 requests/second. The server with eight 900-MHz processors performed 1,900 requests/second. The installation scaled at a ratio of 1 to 1.6 [(7200/1900)/(2200/893)] when moving from four to eight processors. The administrators understood this to be an acceptable scaling ratio. Realistically, the overhead in system resources and issues such as lock contention do not allow a system to scale linearly; and so they did not expect to see a scale ratio of 1 to 2 when moving from four to eight processors. According to their calculations, the administrators determined that the server with eight 900-MHz processors, barring any unforeseen problems, should scale and perform satisfactorily for their application.

# Scaling Out

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Scaling out* is the process of adding servers to an existing server environment to improve performance and to increase the number of Web sites that the system can host or publish. Scaling out reduces bottlenecks and lock contention because requests coming into the system do not share resources. The request load is balanced among servers.

The Contoso administrators understood that although scaling up to an eight-processor server would meet their performance and scalability needs, they would potentially gain more throughput and increased reliability by hosting their Web application on multiple servers. They decided to use two four-processor servers (one with four 550-MHz processors and one with four 900-MHz processors) with Network Load Balancing for load balancing and failover. After installing and configuring Network Load Balancing on both servers, the administrators ran their WCAT stress test using 500 virtual clients (five client computers running 100 threads each).

As one Network Load Balancing node, the two servers processed (on average) 1,900 requests per second when the Processor(_Total)\% Processor Time counter approached 100 percent. The administrators were initially concerned that the requests per second had not increased beyond the results of the eight-processor server by itself. After looking at the actual megacycles per request, as shown in Table 7.13, the administrators realized they actually had an 18 percent increase in throughput, because the number of megacycles per request had decreased in the Network Load Balancing test.

**Table 7.13 Megacycles/Request Calculations**

| Server | Megacycles/Second (# of Processors/MHz) | Requests/Second | Megacycles/Request |
|---|---|---|---|
| 4 x 550 MHz | 2,200 | 900 | 2.45 |
| 8 x 900 MHz | 7,200 | 1,900 | 3.8 |
| 4 x 550 MHz + 4 x 900 MHz | 5,800 | 1,900 | 3.1 |

With the 18 percent increase in throughput and the added reliability of hosting the Web application on two servers, the administrators asked themselves, Are we done? Is this good enough? Have we scaled this installation as much as we can scale it? They realized that the only truly accurate benchmark they could use for their application was the application itself. In other words, they had scaled the installation out to the point where they could readily process the expected traffic loads while staying within the 30 percent CPU utilization guideline. They had increased reliability by hosting the site on two servers, and they had executed their stress testing with no errors. Unless they began to experience errors from clients, or they noticed — during regular performance monitoring — that traffic loads were increasing and CPU usage was surpassing the 30 percent restriction, they had achieved their objective.

# Additional Resources for IIS 6.0 Scalability

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to IIS scalability.

**Related Information**

- IIS 6.0 Administration Scripts, Tips, and Tricks for information about command-line administration tools in IIS 6.0.
- IIS 6.0 Architecture for information about IIS 6.0 request processing modes.
- Optimizing IIS 6.0 Performance for information about performance tuning.
- Running IIS 6.0 as an Application Server for information about IIS 6.0 application isolation modes, including how to enable and configure worker process isolation mode.
- Working with the Metabase for information about the IIS 6.0 metabase, including how to change a metabase property.
- *Performance Testing Microsoft .NET Web Applications* by Microsoft Application Consulting and Engineering (ACE) Team, 2003, Redmond: Microsoft Press.

**Related IIS 6.0 Operations Guide Topics**

- Performance Counters Reference for information about IIS cache counters.
- Metabase Property Reference for information about specific IIS 6.0 properties.
- Using Command-Line Administration Scripts for information about administering IIS 6.0 programmatically.
- Performance Counters Reference for information about IIS 6.0 counters.

**Related Tools**

- WCAT

  The Microsoft Web Capacity Analysis Tool (WCAT) runs simulated workloads on client/server configurations. Using WCAT, you can test how IIS and network configurations respond to a variety of different client requests for content or data. The results of these tests can be used to determine the optimal server and network configuration for your server. WCAT and its documentation (Wcatguide.doc) are available in the IIS 6.0 Resource Kit Tools.

# Configuring Internet Sites and Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) provides four Internet services: the World Wide Web Publishing Service (WWW service) for hosting Internet and intranet content, the File Transfer Protocol (FTP) service for hosting sites from which users can upload and download files, the Network News Transfer Protocol (NNTP) service for hosting discussion groups, and the Simple Mail Transfer Protocol (SMTP) service for sending and receiving e-mail messages. The WWW service is installed by default. Otherwise, install only the services that you need. After you install IIS–based Internet services, you can create sites or virtual servers, configure properties and security settings, and set up components to further customize your server.

**In This Section**


DNS Overview_IIS_SP1_TechRef

Configuring Web Sites in IIS 6.0

Configuring FTP Sites

Configuring the NNTP Service

Configuring the SMTP Service

Additional Resources for Configuring Sites and Services in IIS 6.0


**Related Information**

- For information about the WWW service and IIS core components, see IIS 6.0 Architecture.
- For information about creating a Web site, see Common IIS 6.0 Administrative Tasks.
- For information about IIS security, see Managing a Secure IIS 6.0 Solution.
- For information about creating Web sites programmatically, see IIS 6.0 Administration Scripts, Tips, and Tricks.

# DNS Overview_IIS_SP1_TechRef

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

The Internet employs Transmission Control Protocol/Internet Protocol (TCP/IP) to locate and connect to network hosts, which are computers or other network devices. Each network host is identified by a unique IP address. For users to reach your IIS-based Internet site or service, you must have a unique IP address.

IP addresses are expressed in dotted-decimal notation and segmented into four 8-bit octets, for example, 192.168.1.42. It is difficult to remember and use numeric IP addresses, and many people prefer to use alphanumeric names. To address this issue, Domain Name System (DNS) provides a standard naming convention for locating IP-based computers. When you use DNS, clients can connect to an Internet site or service by typing a DNS domain name (such as www.microsoft.com) in their browsers, rather than an IP address (such as 192.168.1.42). A DNS domain name such as www.microsoft.com is easier to remember and type correctly.

The process of mapping a DNS domain name to an IP address is called *name resolution*. DNS name resolution is performed by a DNS service, such as the DNS service provided with the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; and Windows® Server 2003, Datacenter Edition, operating systems. For more information about Windows Server 2003 DNS, see DNS in Help and Support Center for Windows Server 2003.

Before you can set up an IIS-based Internet site or service, you must obtain an IP address, and if you want users to be able to use a DNS domain name to connect to your Internet site or service, you must also register a unique DNS domain name. Before you request a DNS domain name to register, decide on a plan for your DNS namespace.

**Namespace Planning for DNS**

DNS uses a namespace that is hierarchically structured and based on the concept of a tree of named domains. The tree has multiple levels:

- The root domain. The root domain is the top-level organizational domain. Root domain names include com, org, and gov.
- Second-level domains. Second-level domains provide an identifying name. For example, if your company is named Contoso Pharmaceuticals, then the second-level domain can be *contoso* in contoso.com.
- Subdomains. Subdomains are arranged according to departments, geography, or other logical divisions to help users find information as quickly as possible. For example, you can create a subdomain of research.contoso.com.

The DNS namespace includes globally unique identifiers — Internet domain names, IP address numbers, protocol parameter and port numbers, and others — that are assigned to the organizations that request them. These identifiers must be unique for the Internet to function.

To effectively plan your DNS namespace, consider the long-term organization of your network architecture, as well as how users might look for information. For more information about namespace planning, see Namespace planning for DNS in Help and Support Center for Windows Server 2003.

**Registering a DNS Domain Name**

Your Internet service provider (ISP) should assign you an IP address for your server that you will use when you configure IIS to host your Internet site or sites. To register a DNS domain name, you can work directly with a naming authority. Many ISPs can register a DNS domain name for you. The InterNIC Web site maintains a list of registrars that are currently taking registrations for all available top-level domains. The InterNIC Web site organizes and lists these registrars alphabetically, by name, by location, and by supported languages. After you register your DNS domain name, your ISP or registrar can host the domain, or you can host the domain yourself (if you are a more advanced user). For more information about domain name registration services, see the InterNIC Web site.

# Configuring Web Sites in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS WWW service allows you to publish content on the Internet or an intranet. You can publish a static Web site or a dynamic application. The primary function of the WWW service is to connect client Hypertext Transfer Protocol (HTTP) requests to your Web site or sites. The WWW service is installed at *SystemDrive*:\Windows\System32by default when you start IIS.

| **Note** |
| --- |
| If you want to install the WWW service somewhere other than at the default location, you must use an unattended installation answer file and specify the alternative path using the **PathWWWRoot** parameter. For more information, see Appendix E: Unattended Setup. |

Related Information

For more information about the WWW service and IIS core components, see IIS 6.0 Architecture.

# Creating and Setting Up Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you use the Manage Your Server Wizard to install the application server role, the default Web site is created automatically. You can use this Web site as it is, change it, or create and use new Web sites. For step-by-step instructions for creating your first Web site, see Common Administrative Tasks or see Creating Web or FTP Sites.

Whether you plan to publish your site on an intranet or the Internet, the principles of providing content are the same. You place your Web files in directories on your server and configure IIS to serve those directories. Then users establish an HTTP connection and view the files with a browser. The first step in creating a Web site, therefore, is to design your server, directory, and file structure.

The two main types of directories in IIS are physical directories and virtual directories. Each Web site must have one home directory, which is the central location for the published pages of the site. In addition, you can create Web Distributed Authoring and Versioning (WebDAV) publishing directories, where users can publish, lock, and manage resources.

# Home Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The *home directory* is the central location for your published pages. It contains the pages or applications that you will publish. When you install the application server role, IIS creates a default home directory for the default Web site. The home directory for the default Web site is *SystemDrive*:\Inetpub\Wwwroot. You can change the home directory to another local directory or a shared directory on another computer, or you can redirect requests to a URL if your content is located on a different server on the Internet. For step-by-step instructions for setting the home directory, see Setting Home Directories.

If you set your home directory as a shared directory on another computer, then users might need to enter a user name and password to access the resource. If you set the home directory as a shared directory on another computer, it is important to use an account that has only the minimum credentials that are required to run the application. If you use an account with administrative credentials when you create this directory, then clients can gain access to server operations, which seriously jeopardizes the security of your network. For more information about security credentials, see Managing a Secure IIS 6.0 Solution.

If you have a simple Web site, you can set up subdirectories under the home directory. You can create subdirectories for the different elements used on your Web site, such as images, scripts, and server-side includes. Alternatively, in a simple Web site setting, you can set up virtual directories instead.

# Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To publish from any directory not contained within your home directory, you must create a *virtual directory*. A *virtual directory* is an *alias*, either for a physical directory on your server hard disk that does not reside in the home directory or for the home directory on another computer. An *alias* is a name that Web browsers use to access the virtual directory. Because an alias is usually shorter than the path name of the directory, it is more convenient for users to type. A client simply adds the alias to the URL of the Web site to browse the Web content in that virtual directory.

Using virtual directories offers several benefits. Not only is an alias usually easier for users to type and remember, but also it is a good security practice to use virtual directories because an alias hides the physical location of your files from unauthorized users. In addition, virtual directories make it easier for you to change and manage the directories on your site. If you change the name or location of a directory, you can simply change the mapping between the alias and the physical directory, rather than change the URL of the Web site.

In addition, virtual directories allow administrators to store content in different locations. Users can access the directories using simple URLs. For example, if a default Web site is located at C:\Inetpub\WWWRoot, an administrator can create a virtual directory as a subdirectory of the default Web site, but the virtual directory can have a physical directory in a different location, such as D:\Data. If the virtual directory has an alias of *Data*, and the server name is *MyWeb*, users can access the default Web site and the virtual directory by typing **http://MyWeb** and **http://MyWeb/Data**, respectively.

If you have a simple Web site housed on one directory, using virtual directories is an optional, although recommended, practice. However, if you have a complex site that spans multiple directories, drives, or computers, or if you want to specify different URLs for different parts of the site, you must use virtual directories. To set up a virtual directory that maps to a directory on another computer, you must specify the Universal Naming Convention (UNC) name of the directory and provide a user name and password.

For step-by-step instructions for creating and deleting virtual directories, see Using Virtual Directories.

# WebDAV Publishing Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

WebDAV is an industry-standard set of extensions to HTTP 1.1. WebDAV is a good alternative to FTP if you need a secure solution for publishing resources because it enables the same strong authentication, encryption, proxy support, and caching capabilities as any other HTTP-based Web site. In addition, with WebDAV, you can send multiple file transfers through a single TCP connection, whereas FTP requires a new connection for each file transferred. For more information about the WebDAV standard, see WebDAV Resources.

WebDAV is integrated into IIS, but it is disabled by default. For step-by-step instructions for enabling WebDAV in IIS, see Web Authoring with WebDAV. For information about securing WebDAV, see Securing WebDAV Publishing Directories later in this section.

With WebDAV, you can create Web publishing directories where users can publish, lock, and manage resources. To set up a WebDAV publishing directory, enable WebDAV, and then set up a virtual directory using IIS Manager in the same way that you set up a virtual directory for a standard Web site. For security purposes and to configure DAV custom properties, which are special properties that users can search on, a WebDAV publishing directory must reside on an NTFS file system partition. For step-by-step instructions for setting up a WebDAV publishing directory, see Creating Publishing Directories.

**Making the WebDAV Publishing Directory Available to Users**

After you enable WebDAV and set up a publishing directory, users can use the directory on your server as a collaborative work environment. Users with the correct permissions can copy and move files, and they can search for and modify the content and the properties of files. Multiple users can read a file concurrently, but the WebDAV lock and unlock capabilities allow only one person to modify the file at a time.

Users can access a WebDAV publishing directory through clients running Microsoft® Windows® 2000 or Microsoft® Windows® XP Professional operating systems, Microsoft® Internet Explorer 5 and 6, Microsoft® Office 2000 or Office XP products, and through any other client that supports the industry-standard WebDAV protocol.

For more information about WebDAV clients, including how they can access your WebDAV publishing directory, see About WebDAV.

**Enabling WebDAV Redirector**

WebDAV Redirector is a remote file system that allows clients running Windows XP Professional or Windows Server 2003 to connect to your WebDAV publishing directory through the command line. Once the client is connected, a user can create, modify, and save files on the WebDAV publishing directory using any existing applications that are written for a Windows 32-bit operating system. WebDAV Redirector enables users to access files on the Web as though they are on a mapped network drive. When enabling the WebDAV Redirector, you must ensure that the WebClient service is started.

For the WebDAV Redirector to work, the site that contains the virtual directory that is mapped to your WebDAV publishing directory must be set up exclusively as a WebDAV site. Therefore, you must delete all custom headers on that site, including custom headers for the Microsoft® FrontPage® 2002 Server Extensions. If your WebDAV Web site is set up as both a WebDAV site and a site running FrontPage 2002 Server Extensions, IIS will send the FrontPage custom header and the WebDAV Redirector will assume that the server is not a proper WebDAV site.

# Configuring Web Site Properties in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Every Web site is configured with default properties. You can change the default properties at the WWW service level (which affects all the Web sites on a server), the site level, the directory level, or the file level. IIS employs property inheritance to reduce Web site administration tasks. Lower-level properties are automatically inherited from higher-level properties. For example, the default **Connection Timeout** property is 900, which means that every file on every Web site on the server will time out after 900 seconds. If you want to reduce the time-out period for a certain Web site, you can change the setting at the site level and the change will be replicated to all the files and directories on that site. You can edit settings at lower levels individually to override inherited settings from the next level up. If you change a setting at a lower level, and then you later change a setting at a higher level that conflicts with the lower-level setting, you will be prompted to choose whether you want to change the lower-level setting to match the new higher-level setting. For step-by-step instructions for changing default settings using IIS Manager, see Changing Default Web Site Settings.

# Modifying the Default Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you install IIS, it creates the default Web site automatically. This Web site is configured with default settings that are not optimal for every environment. You can do one of three things with the default Web site:

- **Use it as is, without changing any of the settings**. This practice is not recommended for a production Internet or intranet site. It is recommended that you use the default Web site only in test or demonstration situations.

- **Delete the default Web site, and create new sites**. Many ISPs will delete the default Web site and then use custom scripts to create new Web sites as needed. For more information about creating Web sites programmatically, see IIS 6.0 Administration Scripts, Tips, and Tricks.

- **Modify the default Web site to work in your environment**. If you choose to modify the default Web site, then, as a best practice, change the properties immediately, before you create any directories or virtual directories. Directories and virtual directories inherit the properties of the default Web site.

  You can modify properties of the default Web site as appropriate for your environment.

**Modifying Default Web Site Identity Settings**
You can use the Web Site tab of the **Web Site Properties** dialog box to configure several key Web site identity settings.

**Description**
Change the description of the Web site, using a name that is appropriate for your environment. This name will appear only in IIS Manager.

**IP address**
Enter an IP address if you are creating a Web site bound to a specific IP address on the computer and you do not want to use the All Unassigned default setting, or if you recently changed the IP address of the server.

**TCP and SSL port**
Change the value for TCP, User Datagram Protocol (UDP), or both TCP and UDP ports only if you have special requirements. It is recommended that you leave the default TCP and SSL (Secure Sockets Layer) port settings (80 and 443, respectively) for production servers.

**Host header**
Enter a host header value if you plan to host multiple Web sites on the same server, or if you plan to publish the site so that users can access it by using a DNS name rather than the computer name. To create a host-header value, click **Advanced**, and then in the **Advanced Web Site Configuration** dialog box, click **Add to type in a host header value**. Any subsequent Web sites that you create cannot use the host header value that you enter for the default Web site. For more information about using host headers to host multiple Web sites, see Hosting Multiple Web Sites on the Same Server.

**Modifying the Default Authentication Method**
You can use the **Directory Security** tab of the **Web Site Properties** dialog box to change the authentication method. The authentication method determines whether users are identified, and how users must be identified to access your site. The authentication method you select varies, depending on the kind of site you are creating and the purpose of the site. For a complete discussion of authentication, see Managing a Secure IIS 6.0 Solution.

**Modifying the Default Physical File Location**
You can use the **Home Directory** tab of the **Web Site Properties** dialog box to specify the path to the home directory for your Web site and whether the home directory resides on the local computer or on a network computer. The default Web site home directory is *SystemDrive*:\Inetpub\Wwwroot.

**Enabling and Configuring Log File Formats**
You can use the **Web Site** tab of the **Web Site Properties** dialog box to enable logging and to specify the log format that you want to use. For more information about IIS logs, see Analyzing Log Files.

**Specifying the Application Pool**
You can use the **Home Directory** tab of the **Web Site Properties** dialog box to specify the application pool. For more

information about application pools, see IIS 6.0 Architecture and Running IIS 6.0 as an Application Server.

**Extending Web Sites with FrontPage 2002 Server Extensions**

You can install and enable FrontPage 2002 Server Extensions for Web sites. FrontPage 2002 Server Extensions provide Web-based and command-line administration that allow you to *extend* your virtual servers, which means that the site owner can author the site in FrontPage and delegate site ownership and administrative credentials for Web sites. FrontPage 2002 Server Extensions are not installed by default. For more information about FrontPage 2002 Server Extensions, including information about installing and enabling, see Using FrontPage 2002 Server Extensions with IIS 6.0.

# Redirecting Requests

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Redirecting* refers to the process of configuring the Web server to issue an HTTP 302 redirect message to the client, which instructs the client to resubmit the request for a new location. This process helps ensure that users always receive the Web page that they want, even if you changed the location of the page and did not update all of the links.

It is useful to redirect requests when the content is not hosted on the Web server or when you are updating your Web site and want to make a portion of the site temporarily unavailable. In addition, you can use a redirect when you have changed the name of a virtual directory and you want users to be able to access files using the old URL.

You can use redirect wildcards to perform basic rewriting operations for the redirected URL. For example, to redirect all requests for /Scripts/Filename.stm to a single file named Default.stm, and to redirect all requests for /Scripts/Filename.htm to a single file named Default.htm, you would type the following in the **Redirect to** box for the /Scripts virtual directory: ***;Filename.stm;/Default.stm;Filename.htm;/Default.htm**

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **Run as** command to run IIS Manager as an administrator. At a command prompt, type **runas /***User***:***Administrative_AccountName* **mmc %systemroot%\system32\inetsrv\iis.msc**. |

**To use redirect wildcards**

1. In IIS Manager, click the local computer, right-click the Web site or directory that you want to redirect, and then click **Properties**.
2. In the **Web Site Properties** dialog box, click the **Home Directory** tab, the **Virtual Directory** tab,or the **Directory** tab.
3. In **The content for this resource should come from**, click **A redirection to a URL**.
4. In the **Redirect to** box, type the wildcard character (*), a semicolon (;), and then the URL. Separate pairs of wildcard characters and destination URLs with a semicolon.
5. In **The client will be sent to**, select the **The exact URL entered above** check box.

**Redirecting to Dynamic Content**

You can also redirect requests to dynamic content, which is useful for sending users to Common Gateway Interface (CGI) scripts, Active Server Pages (ASP), and .aspx pages. If you want to redirect to dynamic content, make sure that dynamic content is enabled on your server running IIS. To redirect to dynamic content, use redirect variables to pass parameters to the page. For example, to redirect all requests for scripts in the Scripts folder to a logging program named Logger.exe that records the requested URL and any parameters that are passed with the URL, type: **/Scripts/Logger.exe?URL=$V+PARAMS=$P**. In this example, $V and $P are redirect variables. For more complete information about redirect variables, see Redirect Reference.

**See Also**
**Concepts**
Redirecting Requests to Files, Directories, or Programs

# Securing Web Sites in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To secure your Web sites, implement a defense-in-depth strategy that combines many different security measures. Use appropriate settings for authentication methods, Web site access control, IP address and domain name restrictions, and secure communications, including certificates and encryption. In addition, implement tools such as firewalls, and lock the rooms that contain hardware. You can use the **Directory Security** tab in the **Web Site Properties** dialog box to configure authentication, Web site access control, and IP address and domain name restrictions, and to enable secure communications. For a thorough discussion of IIS security, including how to configure security settings, see Managing a Secure IIS 6.0 Solution.

There are a number of considerations for configuring authentication and access-control settings for WebDAV-based remote publishing. These considerations are outlined in this section.

# Securing WebDAV Publishing Directories in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Just as with a standard Web site, you can use the **Directory Security** tab in the **Web Site Properties** dialog box to configure WebDAV authentication, Web site access control, and IP address and domain name restrictions, and to enable secure communications. Furthermore, for security reasons and to enable WebDAV custom search properties, your WebDAV publishing directory must reside on an NTFS file system partition. For more information about NTFS partitions, see Choosing a file system: NTFS, FAT, or FAT32 in Help and Support Center for Windows Server 2003.

Setting the Client Authentication Method

Because WebDAV is an HTTP-based protocol, the same authentication methods are available for your WebDAV publishing directory as for your Web sites. When you create a WebDAV virtual directory, the authentication settings will be set by default to Anonymous authentication and Integrated Windows authentication.

| **Important** |
| --- |
| Change the default authentication settings for your WebDAV site. Anonymous authentication and Integrated Windows authentication are appropriate for Web sites, but not for resources on a WebDAV publishing directory. |

Use the following best practices as guidelines when configuring client authentication for WebDAV. For more information about authentication, see Setting Web Site Authentication.

**Anonymous authentication**

Turn off Anonymous authentication to your WebDAV publishing directory. Without controlling access, your directory might be vandalized by unknown clients.

**Basic authentication**

Turn on Basic authentication only if you encrypt passwords through SSL. Basic authentication sends passwords over the connection in plaintext, so passwords can be intercepted and read. For more information about SSL, see Configuring SSL on a Web Server or Web Site.

**Digest authentication**

Turn on Digest authentication if users access your WebDAV publishing directory over the Internet and through firewalls because Digest authentication sends the passwords over the network as an MD5 hash. However, passwords are stored as plaintext in the Active Directory® directory service.

**Advanced Digest authentication**

Turn on Advanced Digest authentication if users access your WebDAV publishing directory over the Internet and through firewalls. Advanced Digest authentication is the best choice because, in addition to sending passwords over the network as an MD5 hash, the passwords are also stored in Active Directory as an MD5 hash rather than as plaintext.

**Integrated Windows authentication**

Turn off Integrated Windows authentication unless users access your WebDAV publishing directory only from an intranet.

**.NET Passport authentication**

Turn off .NET Passport authentication. Microsoft® .NET Passport authentication uses cookies to validate user credentials against a central .NET Passport server. However, the central .NET Passport server does not authorize or deny the access of a specific user to individual .NET Passport–enabled sites.

Configuring Access Control

You can control access to your WebDAV publishing directory by coordinating permissions in IIS and Windows Server 2003. The following information can help you set the correct permissions depending on what kind of content your WebDAV publishing directory contains and what you want to allow users to do with that content. For step-by-step instructions for setting permissions, see Securing Sites with Web Site Permissions.

**Read, Write, Directory browsing enabled**

Set these permissions if you want clients to see a list of resources, modify them (except for resources that do not have Write

permission), publish their own resources, or manipulate files.

**Write enabled; Read and Directory browsing disabled**

Set these permissions if you want clients to publish private information on the directory, but you do not want others to see what has been published. This configuration works well if clients are submitting ballots or performance reviews.

**Read and Write enabled; Directory browsing disabled**

Set these permissions only if you want to obscure file names as a security method. However, be aware that security by obscurity is a low-level security precaution because an attacker or a malicious user can guess file names by trial and error.

**Index this resource enabled**

Enable Indexing Service if you plan to let clients search directory resources.

You can also control access to your WebDAV publishing directory through discretionary access control lists (DACLs). When you set up a WebDAV publishing directory on an NTFS file system disk, the server running Windows Server 2003 is configured with a secure set of defaults that can be too restrictive for Web publishing. If you change the defaults, assign Write permission only to the specific users or groups that absolutely need Write permission. (Whenever possible, assign permissions to groups rather than users.) By default, the Users group is assigned only Read permission for most portions of the file system.

If you have script files in your WebDAV publishing directory that you do not want users to access, you can deny access to these files by verifying that the Script Source Access permission is not assigned. The extensions of your script files are listed in the Applications Mapping list. All other executable files, including files with .exe extensions, are treated as static HTML files unless the Scripts and Executables permission is assigned to the directory.

You can prevent users from downloading and viewing .exe files as HTML files but still allow clients to run them by assigning the Scripts and Executables permission. To do so, on the **Virtual Directory** tab of the publishing directory, on the **Execute permissions** list, click **Scripts and Executables**. When you assign the Scripts and Executables permission to the directory, all of the executable files are subject to the Script Source Access permission setting. You can assign the Script Source Access permission on the **Home Directory** tab of the **Web Site Properties** dialog box of the Web site that you are configuring. When the Script Source Access permission is assigned, clients with Read permission can see all the executables, and clients with Write permission can edit and run all the executables.

If you want clients to be able to write to an executable file that does not appear in the Application Mapping list, assign the Write permission and the Scripts Only Execute permission to the directory.

If you want clients to be able to write to any executable file, regardless of whether it appears in the Application Mapping list, assign the Script Source Access permission and the Scripts and Executables Execute permission to the directory.

# Hosting Multiple Web Sites on the Same Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can host and manage multiple Web sites on one server running Windows Server 2003 and IIS 6.0. To host multiple Web sites on a single Web server, you must configure a unique identity for each Web site by using one of three methods:

- Port-Based Routing
- IP-Based Routing
- Host-Header Routing

By using one of these three methods, you can create unique identities for multiple Web sites without installing a dedicated Web server for each site. You can also create a unique home directory for each site and store the content on the local server or on remote network shares. By doing so, you configure each Web site as a separate entity or virtual server.

For each server hosting multiple Web sites, choose one of the three routing methods listed earlier. When you use one method per server, performance improves because cache and routing lookups are optimized. Conversely, when you use a combination of nonstandard port numbers, unique IP addresses, and host headers to identify multiple Web sites on the same server, performance is degraded for all of the Web sites on the server. For more information about Web site performance, see Optimizing IIS 6.0 Performance.

You can add Web sites to a server by using IIS Manager or by using the lisweb.vbs command-line script. For step-by-step instructions for adding Web sites, see Adding Web Sites to an IIS 6.0 Server.

# Port-Based Routing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use a nonstandard TCP port number for each Web site. Some administrators use ports other than port 80 for development and testing purposes or for running other server applications that require exclusive access to port 80.

When using port-based routing to host multiple Web sites on the same server, consider the following:

- A server can listen on a maximum of 64 ports at one time. The actual number of ports a server can listen on varies depending on the available RAM on the computer. Therefore, you cannot use port-based routing to route a large number of sites.
- Port-based routing consumes a vast amount of kernel resources and can degrade server performance.
- If you use a nonstandard TCP port number to identify a Web site, users cannot reach the site by using a standard name or URL. Instead, users must know the nonstandard TCP port number assigned to the Web site and add it to the name or IP address of the Web site in their browser. For example, to reach a private Web site named Contoso with an IP address of 10.0.0.02 and HTTP assigned to the nonstandard TCP port 8080, users would have to type **http://contoso:8080** or **http://10.0.0.02:8080** in their browser to reach the site.

If you choose to use a nonstandard port for a nonproduction server, use a port number higher than 1023 so that it does not conflict with the well-known port numbers that the Internet Assigned Numbers Authority (IANA) assigns.

# IP-Based Routing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

With IIS 6.0, you can configure Web sites to listen on all available IP addresses or on an individual IP address. If you use multiple IP addresses to distinguish between different Web sites on the same server, you must configure IIS to assign each site a unique, static IP address.

Although it is simple to configure IP–based routing, this strategy typically is viable only for large corporations and ISPs. Unique IP addresses are in short supply, so it can be difficult and expensive to obtain a block of static IP addresses.

One advantage of IP–based routing is that it enables you to host multiple Web sites that require SSL on the same server. However, hosting a large number of sites that require SSL on the same server can degrade performance because so many requests have to be encrypted and decrypted.

# Host-Header Routing

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

You can use one IP address (or the **All Unassigned** setting) for multiple Web sites and distinguish the Web sites with unique host-header names. This option is easy to configure by using IIS Manager. To use host-header names, your computer or network must be using a name resolution system such as DNS.

When a client sends a request for an HTTP connection to the server running IIS, the client request includes a field in the HTTP header called HOST. This field contains the host name of the requested Web server. For example, when you type **http://www.microsoft.com** into your browser, your browser sends a packet with an HTTP header that includes the following field: HOST: www.microsoft.com. Because the name of the field is *HOST* and it is in the HTTP header of the client, it is referred to as the *host header*. For more information on the HTTP 1.1 standard, see the RFC 2616, Hypertext Transfer Protocol -- HTTP/1.1 on the Internet Engineering Tasks Force Web site.

With host-header routing, when a client request reaches the server, IIS uses the host name, rather than only the IP address as it does with IP–based routing, to determine which Web site the client is requesting. To set up multiple Web sites on the same server to take advantage of host-header routing, the host-header names must be DNS domain names that are publicly available and registered with an authorized Internet naming authority. For more information about obtaining a DNS name, see DNS Overview_IIS_SP1_TechRef.

| **Note** |
| --- |
| To set up an intranet site for host-header routing, the host-header name can be an intranet site name. |

**Host-Header Routing and SSL**

Beginning with Windows Server 2003 Service Pack 1 (SP1), IIS supports the use of SSL with host header routing. For more information about setting up host headers for an SSL-enabled Web site, see Configuring SSL Host Headers.

# Configuring FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS includes the FTP service for publishing and managing files. FTP uses TCP at the transport layer, so all the data that is exchanged between the server and the client is guaranteed not only to be delivered but also to be delivered intact, as originally sent.

The IIS Admin service, which is physically part of the Inetinfo.exe process, manages the protocols in IIS that are not related to the Web, including FTP. By default, the FTP service is not installed on a server running Windows Server 2003. Once it is installed, the FTP sites start automatically the next time you restart the IIS services.

# Installing the FTP Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you can set up an FTP site, you must install the FTP service from Add or Remove Programs in Control Panel. For information about how to install the FTP service, see FTP Site Setup.

When you install the FTP service, a default FTP site directory is created with the home directory *SystemDrive*:\Inetpub\Ftproot.

There are two kinds of directories for FTP site content: user home directories and shared virtual directories. Each FTP site, except for those in Active Directory Isolation mode, must have a home directory that serves as the initial directory for users connecting to the site. In most configurations, each user has an individual directory under this site home directory. Create a virtual directory to hold additional content that is accessible to every user that logs on to the site, or if you want to use different URLs for different parts of the site. This virtual directory is different than the anonymous user home directory that is accessible to all anonymous users. Home and virtual directories for FTP sites function similarly to Web site directories. For more information about directories, see Creating and Setting Up Web Sites.

| **Note** |
| --- |
| Virtual directory names do not show up when listing the home directory of a user. Instead, virtual directories are used as if they were subdirectories of the user's home directory. A site administrator can publish these virtual directories in the site welcome message. |

# Configuring FTP Site Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The FTP service uses inherited default property settings, much like the WWW service does. Settings on lower levels automatically inherit higher-level settings, and you can edit lower-level settings individually to override inherited settings from the next level up. If you use IIS Manager to change a setting at a lower level and then later change a setting at a higher level that conflicts with the lower-level setting, you are prompted to choose whether you want to change the lower-level setting to match the new higher-level setting. For step-by-step instructions for changing settings, see Changing Default FTP Site Settings.

**Configuring Site Messages**

When you set up an FTP site, you can configure IIS to send messages when a user accesses the FTP site. You can create the following messages:

- **Banner**. This message is sent to users as they connect but before they log on to an FTP site. A banner message usually informs users about what kind of information is available on the FTP site.
- **Welcome**.This message welcomes users to an FTP site after they have logged on.
- **Exit**. This message appears when users log off from an FTP site.
- **Maximum connections**. This message is sent to users when the maximum number of connections for the FTP site has been reached.

**Setting the Directory Output Style**

You can set the directory output style of your FTP site. The directory output style can be either Microsoft® MS-DOS® operating system-based or UNIX-based, and the output style defines the format that is used when an FTP client sends a **list** command.

By default, the MS-DOS directory output style displays the year in a two-digit format, for example, 03. You can change this setting to display the year in a four-digit format, for example, 2003, by enabling the **FtpDirBrowseShowLongDate** metabase property. For more information, see FtpDirBrowseShowLongDate Metabase Property.

The UNIX directory output style displays the year in a four-digit format when the date of the file is different from the year of the FTP server. The year is not returned in cases where the file date is the same as the year of the FTP server.

**Setting the Authentication Method**

Two types of users can connect to an FTP site: anonymous users and authenticated users. Set the authentication method to allow only anonymous users, authenticated users, or both anonymous users and authenticated users. If you want to allow only anonymous users, in the **FTP Sites Properties** dialog box, on the**Security Accounts** tab, select the **Allow anonymous connections** check box and the **Allow only anonymous connections** check box. If you want to allow only authenticated users, clear the **Allow anonymous connections** check box. If you want to allow anonymous users and authenticated users, select only the **Allow anonymous connections** check box. Anonymous access to FTP resources works similarly to anonymous access to Web sites. If anonymous access is allowed, all client connections with the user name *anonymous* or *ftp* are accepted and authenticated as the anonymous user. By default, the user account that is configured as the anonymous user is IUSR_*ComputerName*. As an Internet convention for anonymous FTP user identification, the anonymous user is prompted for a password, which is used as a user name.

If a site uses FTP User Isolation and a user connects with a user name and password, then the user is authenticated with his or her Windows account and mapped to individual home directories. This practice is not secure because credentials are transmitted over the network in plaintext and can therefore be intercepted and read. For step-by-step instructions for setting the FTP authentication method, see FTP Site Authentication in IIS 6.0.

| **Note** |
| --- |
| You cannot set the authentication method and the anonymous user properties using IIS Manager for FTP sites that are configured for the **Isolate users using Active Directory** mode. The **Isolate users using Active Directory**mode is explained in more detail later in this section. |

**Setting the Passive Connection Port Range**

FTP clients connect to FTP servers through a well-known service port, TCP port 21, and establish a control channel for sending requests and receiving replies. For actual file data transfers, a second connection is established: a *data channel*. By default, the server connects back to the client to establish the data connection, which is an active-mode connection. However, to support clients that are behind firewalls and cannot accept incoming connections, the FTP protocol allows the client to connect to the server to establish the data channel, which is a passive-mode connection. Passive-mode connections are established through the following process: the client requests a passive mode (PASV), the server replies with a unique port number, and the client connects to that port.

By default, the FTP server allocates ports for passive-mode connections from the WinSock dynamic range, 1024 to 5000. If the server itself is behind a firewall, the 1024 to 5000 port range must be open for incoming connections at the firewall. This requirement can expose the server to potential attack because many applications share the WinSock dynamic port range. To reduce the attack surface, you can configure the FTP service to allocate ports for passive-mode connections from a port range above 5000, which allows you to assign a port range to be used exclusively by FTP passive-mode connections and to create firewall and router policies that open that range for incoming TCP connections.

To set the port range for passive-mode connections, edit the **PassivePortRange**property in the metabase. When you set the port range, consider the number of anticipated concurrent file transfers because each client might need a distinct port. For performance reasons, the port range should be at least two times the number of anticipated concurrent file transfers.

For more information about setting the port range used for FTP passive-mode connections, see PassivePortRange Metabase Property.

# Securing FTP Sites in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FTP is not a secure means of transferring data. Although you can configure the FTP service to require a valid user name and password combination, neither the credentials that are specified at log on nor the data itself is encrypted or encoded in any way. The credentials are sent across the network in plaintext and can be intercepted and analyzed by any station on any network between the FTP client and the FTP server. If your plaintext credentials are intercepted and analyzed by a malicious user, then someone other than the intended user can log on to your FTP site and download the files that you placed there or gain access to other network resources, even network resources that require Windows account authentication. For more information about setting the authentication method for your FTP site, see Configuring FTP Site Properties.

If you intend to place sensitive data on your FTP site or if secure communication between clients and your FTP server is important, consider using FTP over an encrypted channel, such as a virtual private network (VPN) secured with Point-to-Point Tunneling Protocol (PPTP) or Internet Protocol security (IPsec). To learn how to set up a private network over the Internet using PPTP or how to set up secure IP communications with a VPN and IPsec, see the PPTP, VPN, and IPsec topics in Help and Support Center for Windows Server 2003. For more information about IPsec and IIS, see Managing a Secure IIS 6.0 Solution.

If you need to authorize users or host sensitive content, then WebDAV is a good alternative to FTP because WebDAV can take advantage of SSL for secure communications. For more information about WebDAV, see WebDAV Publishing Directories.

# Creating Additional FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you install the FTP service, the default FTP site is created with a default directory that is located at *SystemDrive*:\Inetpub\Ftproot. If you want to add additional FTP sites, then you can use the FTP Site Creation Wizard or the Iisftp.vbs command-line script, which is stored in the *SystemDrive*:\Windows\System32 directory. Iisftp.vbs sets up the home directory and configuration in the metabase for the new FTP site. For step-by-step instructions for creating new FTP sites with the wizard or the command-line script, see Adding FTP Sites to Your Server.

Each FTP site must have a unique IP and port number combination, so each site must have a unique IP address or a unique TCP port number (other than 21, which is the default port number for FTP connections). It is recommended that you create multiple sites with unique IP addresses instead of multiple sites with unique port numbers because when you use a port other than the default port 21, users must type the port number, in addition to the computer name, to locate the site.

If you use a name resolution system, such as Windows Server 2003 DNS, and you use port 21 for FTP, then visitors can type **ftp://** followed by your computer name in their browsers to reach your site. If you use a port other than port 21, then users have to type **ftp://**followed by the port number and the name of your computer. If you have multiple sites with multiple IP addresses, then each address needs to have a corresponding registered domain name, which users must type instead of your computer name. For more information about registering a domain name, see DNS Overview_IIS_SP1_TechRef.

# Hosting Multiple FTP Sites with FTP User Isolation

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can host multiple FTP sites on the same server running IIS. If you are an ISP or an application service provider in a multi-site Internet hosting scenario and you want to ensure that your customers cannot access the FTP directories of one another, then you can enable FTP User Isolation. When you enable FTP User Isolation, the user's top-level directory appears as the root of the FTP site, so other users cannot view or overwrite content. Within the user's specific site, the user can create, modify, or delete files and folders.

FTP User Isolation supports three isolation modes: **Do not isolate users**, **Isolate users**, and **Isolate users using Active Directory**. You can select the isolation mode during FTP site setup using the FTP Site Creation Wizard. You can use lisftp.vbs to configure FTP User Isolation, using the */isolation*parameter. When you use the */isolation* parameter, specify either **AD**, for Active Directory isolation, or **Local**, for local isolation. If you do not include the */isolation* parameter, the site will not isolate users.

You can select a different isolation mode for each FTP site.

| **Important** |
| --- |
| After you set the FTP User Isolation mode and finish the FTP Site Creation Wizard or create the site using lisftp.vbs, do not change the isolation setting manually. |

#### Do Not Isolate Users Mode

The **Do not isolate users** mode does not enforce FTP User Isolation and is designed to work like earlier versions of the FTP service in IIS. Because isolation is not enforced among the different users that log on to your FTP server, this mode is ideal for a site that offers only download capabilities for shared content or for sites that do not require protection of data between users.

In this mode, all user directories are in one level as a subdirectory of the FTP site directory. The site directory can reside either on the local computer or on a network share. When a user accesses the FTP site, the home directory is determined as follows:

- By default, the initial user directory is set to the physical path configured as the FTP site directory. If this directory does not exist, the user connection is denied.
- If a user name is supplied, the home directory is derived from the user name. For anonymous users, the derived name is *anonymous*. For local computer users, the derived name is the user name. For domain users, the derived name is the user name without the domain name.
- If a directory with the derived name exists in the site directory, then that directory becomes the initial log-on directory for the user.

| **Important** |
| --- |
| The user is not restricted to this subdirectory and can navigate up to the site directory and into other subdirectories unless you have set ACLs to prevent users from accessing them. |

#### Isolate Users Mode

The **Isolate users**mode determines a unique home directory for each user derived from the user name. The home directory of the user is treated as a root directory for the user, and the user cannot navigate or access the physical file system outside of the root directory. If users need access to dedicated shared folders, then you can establish a virtual directory.

User home directories are located in a two-level directory structure under the FTP site directory. The site directory can reside either on the local computer or on a network share. When a user accesses the FTP site, the home directory is determined in one of three ways:

- For anonymous users, the home directory is **LocalUser\Public** under the FTP root directory.
- For local users, the home directory is **LocalUser***UserName* under the FTP root directory.

- For users that log on with *Domain\UserName*, the home directory is *Domain\UserName* under the FTP root directory.

The user home directory must be created before the user logs on. If the directory does not exist when the user attempts to connect, the connection is denied.

For information about creating a new FTP site or converting an existing FTP site to Isolate users mode, see Creating a New FTP Site with Isolate Users Mode and Converting an Existing FTP Site to Isolate Users Mode.

**Isolate Users Using Active Directory Mode**

In the **Isolate users using Active Directory** mode, the FTP service is integrated with Active Directory to retrieve home directory information for users. To accomplish this integration, the Active Directory user object is extended with two properties: **msIIS-FTPRoot** and **msIIS-FTPDir.** The **msIIS-FTPRoot** property stores the file server share, and the **msIIS-FTPDir** property stores the relative physical path to the home directory for each user. You can use Iisftp.vbs to get and set these properties.

Information that is frequently retrieved from Active Directory is cached by the FTP service. The default caching interval is 10 minutes. You can adjust the caching interval by editing the registry entry DsCacheRefreshSecs.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference on the Microsoft® Windows® Server 2003 Resource Kit companion CD or at the Windows Server 2003 Resource Kit web site. |

If the FTP service successfully accesses the path, it becomes the home directory for the user. The user cannot access the file system outside this directory. The user is denied access if either the **msIIS-FTPRoot** or **msIIS-FTPDir** property does not exist, or if the resulting home directory cannot be accessed.

This configuration option provides maximum flexibility and control over user home directories in an ISP environment. For example, John Doe connects to his FTP site at ftp.example.com. The example.com network load-balancing server (NLBS) resolves this request to the server FTPS3. John Doe enters the user name JohnDoe. FTPS3 is configured with the default domain name Domain4, so the user Domain4\JohnDoe is authenticated and the home directory information for John Doe is retrieved from Active Directory as \\FS1\Share2\Users\u2\johndoe_dir. From now on, FTPS3 will access the data under this home directory for all the FTP requests from John Doe. This example is illustrated in Figure 8.1.




**User Home Directories**

In the **Isolate users using Active Directory** mode, each user's home directory resides on an arbitrary network path, which provides you with the flexibility to distribute user home directories across multiple servers, volumes, and directories. In addition, you can move user's home directories transparently from one location to another, which affects the service for the users only during the duration of the move. You can also set the **msIIS-FTPRoot** and **msIIS-FTPDir** properties so that the users home directory maps to a local folder on the FTP server; for example, **msIIS-FTPRoot** is set to **E:\FTPUsers**.

**Configuring Isolate Users Using Active Directory Mode**

There are three main steps to configuring the **Isolate users using Active Directory** mode:

- Configure the file servers.
- Configure Active Directory.
- Create and configure the FTP sites to be isolated.

When you configure the file servers, you must create the shares and user directories for all the users that are permitted to connect to the FTP service, including the user configured to impersonate anonymous users. Before you complete this step, consider factors such as expected disk space usage, storage management, and network traffic.

To configure Active Directory, you need a server running Windows Server 2003, Standard Edition, and Active Directory. Configure the user object in Active Directory for each user, including the user configured to impersonate the anonymous user, by setting the **msIIS-FTPRoot** and **msIIS-FTPDir** properties to point to the home directories that you previously created. You can do this by using the lisftp.vbs command-line utility with the **/SetADProp**command. To learn more about setting up Active Directory, see Active Directory in Help and Support Center for Windows Server 2003.

For information about working with FTP sites with **Isolate users using Active Directory** mode, see Creating a New FTP Site with Isolate Users Using Active Directory Mode, Enabling Anonymous Access for an FTP Site Configured with Isolate Users Using Active Directory Mode, and Converting an Existing FTP Site to Isolate Users Using Active Directory Mode.

# Configuring the NNTP Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS NNTP service enables you to create a news server to host discussion groups. NNTP is the protocol for distributing, posting, and retrieving news items between a news server and client.

You can use the NNTP service to host local discussion groups, called *newsgroups*, on a single computer. For example, you can create departmental newsgroups where company department members can discuss projects and issues. Or you can create a newsgroup and allow customers to access it over the Internet for technical or customer support. Because the NNTP service is in complete compliance with the NNTP protocol, users can use any newsreader client, such as the Microsoft® Outlook® Express messaging and collaboration client, to participate in newsgroup discussions.

You can also configure the NNTP service for *newsfeeds*. Newsfeeds enable users of different newsgroups to read and post articles to other newsgroups, as if the users were using a single group. An article that is posted to one newsgroup gets sent to other newsgroups, where it can be read.

The NNTP service is hosted by Inetinfo.exe. Clients connect to the NNTP service through TCP port 119 by default. If you enable encrypted SSL connections, the default port is 563.

# Installing the NNTP Service and Setting Up NNTP Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The NNTP service is not installed by default. You must install the NNTP service using Control Panel, which creates a default NNTP configuration that you can customize using IIS Manager. For step-by-step instructions for installing the NNTP service, see NNTP Server Setup.

The NNTP service stores newsgroup articles on virtual servers with one or more directory hierarchies. Each virtual server has its own virtual directory or directories, which contain newsgroup articles, or messages, stored as files. In addition, the NNTP service involves other file types, which include internal files as well as the public newsgroup message files. When you install the NNTP service, a default NNTP virtual server configuration is created with a message store in *SystemDrive*:\Inetpub\Nntpfile\Root.

**Virtual Servers**
You can have multiple NNTP virtual servers running on IIS. Each virtual server has its own properties, expiration policies, and virtual directories. Create as many NNTP virtual servers as you need. Consider creating multiple virtual servers to host different types of newsgroups on the same server. For example, you might want to host internal departmental newsgroups on the same server with Internet-based customer service newsgroups. Because those newsgroups will have different settings, it is easiest to set them up on different virtual servers. However, if you have many different departmental newsgroups that have the same property settings, consider hosting all of them on the same virtual server.

Each virtual server should have a unique IP address/TCP port combination. Although it is possible to host multiple NNTP virtual servers on a single IP address by using a different TCP port for each one, this configuration is not recommended because users will have to configure their client software specifically to use the nonstandard TCP port.

**File and Message Types**
The NNTP service involves several file and message types, including:

- **Message files**. Message files are messages that are posted to newsgroups, receive an .nws extension, and are stored in the directory that the NNTP service creates for them.

- **Internal files**. Users never see or use internal files, which are files that the NNTP service creates that list the subjects of the stored articles. These internal files have an .xix extension, and the NNTP service creates one subject file for every 128 articles in a newsgroup. The NNTP service also uses .tag files to mark folders as newsgroups or virtual roots. In addition, the NNTP service maintains a number of internal data structure files with .hsh, .hdr, .lst, and .txt extensions. These files are essential for the NNTP service; do not delete them. If you accidentally delete the internal files or if they become corrupted, you can restore them by rebuilding the NNTP service. For more information about rebuilding the NNTP service, see Rebuilding the NNTP Service.

- **Control messages**. Control messages are specially formatted messages that direct the NNTP service to delete a specific article from a newsgroup or to add or delete an entire newsgroup. Control messages are stored in directories that the NNTP service creates. For more information about control messages, see Using Control Messages.

When you create an NNTP virtual server, you are prompted to enter two directory paths: a path for internal files and a path for newsgroup message files. If possible, put these two paths on different disk partitions because the internal files can be corrupted if the disk partition runs out of disk space. Also, if you expect a high volume of traffic on your NNTP virtual server, you can place the files on a drive in a redundant array of independent disks (RAID) to improve performance.

**Virtual Directories**
A virtual directory is a directory on a local or remote disk drive that acts as the root of a part of the newsgroup hierarchy. Virtual directories enable you to store newsgroup files on multiple disk drives. Using more than one disk drive can improve the performance of a heavily used disk drive and can provide more storage. Virtual directories also enable you to change the physical location of the directory without changing the name of the newsgroup.

Virtual directories can be located on a local disk drive or on a network drive. However, if you store content in a remote location, then you must configure the remote computer with at least Read and Write permissions for the user name that you provide. Make sure that the remote location is secure and maintained by a trusted administrator. There is no advantage to creating a virtual directory on the same physical drive as the home directory. For step-by-step instructions for creating virtual directories, see Creating Virtual Directories in IIS 6.0.

# Configuring NNTP Virtual Server Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NNTP virtual server properties** dialog box in IIS Manager allows you to configure many different settings for your virtual server, including access control, logging, newsgroup limits, and security. Use the **General** tab to set the IP address and, optionally, the port number for the NNTP virtual server. You can also set connection limits and time-outs, specify a path header, and enable logging. Auditing server activity with log files is a good way to detect if unauthorized external users are attempting to access your virtual server or if internal users are trying to access resources that they do not have authority to access. Use the **Settings** tab for management tasks, such as enabling moderated newsgroups and control messages. For more information, see Creating and Managing Newsgroups. Use the **Access**, **Settings**, and **Security** tabs to configure security-related settings. For more information about setting properties that affect security, see Securing NNTP Virtual Servers.

For a detailed description of each property on every **Properties** tab and for more information about logging, see Configuring NNTP Virtual Servers.

# Securing NNTP Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By configuring the properties of your NNTP virtual server correctly, you can increase the security of your NNTP server and your network. The **Access**, **Settings**, and **Security** tabs each have security-related properties that you can use to configure access-related settings, set post and connection limits, and assign administrative credentials. You can also increase security through the use of access control lists (ACLs).

**Configuring Access-Related Settings**
The **Access** tab has three important security-related settings: Authentication, Certificate, and Connection.

**Authentication**
The Authentication option allows you to select from the following authentication methods for newsgroup users:

**Anonymous access**. Anonymous access is enabled by default. If you do not disable Anonymous access, everyone is able to access all of the newsgroups on this virtual server.

- **Basic authentication**. If you enable Basic authentication and you disable anonymous access, users must enter a valid Windows user name and an authorized password to access the newsgroup. However, those credentials are not encrypted when they are sent across the network (credentials are sent in plaintext). Basic authentication leaves your server vulnerable to a dictionary attack. Use this option only if your newsgroups are internal.

- **Integrated Windows authentication**. If you enable Integrated Windows authentication and you disable anonymous access, users must enter a valid Windows user name and an authorized password to access the newsgroup, and those credentials are encrypted when they are sent over the network. Integrated Windows authentication can also leave your Web server vulnerable to a dictionary attack. Outlook Express supports this protocol.

- **SSL client authentication**. SSL provides a secure, encrypted connection between the NNTP service and the client. SSL support requires that an SSL certificate is installed on the computer running Windows Server 2003 and that the client software supports SSL. (Outlook Express supports SSL.) If a server certificate is installed, then the NNTP service uses SSL whenever a client requests it. You have the option to require SSL for all newsgroups or for newsgroups located in a virtual directory.

**Certificate**
The Certificate option starts the Web Server Certificate Wizard, which you can use to create and administer server certificates used in secure Web communications, such as those that require SSL.

**Connection**
The Connection option allows you to restrict access to the virtual server based on the IP address of the client. By default, all IP addresses have access to the NNTP service. You can either allow or deny access to a specific list of IP addresses, and you can specify the IP addresses individually or as a group using a subnet mask. You can also specify IP addresses using a domain name, but doing so adds the overhead of a DNS lookup for each connection. For more information about domain names and DNS, see DNS Overview.

**Setting Post and Connection Limits**
There are several options on the **Settings** tab that affect the security of your NNTP virtual server. You can set the post size, in kilobytes, of the largest single article a user can post. You can also set the connection size, in megabytes, of the maximum amount of data that a user can post to a newsgroup during a connection. When you set limits on the size of posts and connections, you help prevent a malicious individual from trying to overload your server resources by posting extremely large articles.

**Assigning Administrative Credentials**
You can use the **Security** tab to grant administrative credentials to Windows user accounts. By default, the Windows Administrator account is granted administrative credentials. For more information about security and access control, see Managing a Secure IIS 6.0 Solution.

**Creating ACLs**
After you set the authentication requirements for your virtual server, you can further restrict access to newsgroups with ACLs. For example, in the case of departmental newsgroups, ACLs can prevent non-executives from reading the executive

newsgroup. You can control access to individual newsgroups or sets of newsgroups by assigning permissions on your computer running Windows Server 2003 for the directories that contain those newsgroups. You can set permissions for an individual directory or for a set of directories. For more information about ACLs, see Access control in Help and Support Center for Windows Server 2003.

# Creating and Managing Newsgroups

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A *newsgroup* is the organizational structure for a group of related news articles. For example, if you are creating departmental newsgroups, then you can create the accounting, human resources, and executive newsgroups. If you expect a lot of newsgroup activity, then you can break down the newsgroups into narrower categories. However, because articles within each newsgroup are organized by subject, you do not need to create a separate newsgroup for each subject.

# Creating Newsgroups

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

There are three ways to create a newsgroup:

- Use the New Newsgroup Wizard in IIS Manager. For step-by-step instructions for creating a newsgroup with the wizard, see Creating Newsgroups in IIS 6.0.
- Use the Rgroup.vbs command-line script.
- Use a control message. For more information about using control messages, see Using Control Messages.

When you create a newsgroup, the NNTP service creates the directory for the newsgroup automatically.

**Creating Newsgroups Using Rgroup.vbs**

You can use Rgroup.vbs to create newsgroups. You can also use Rgroup.vbs to delete and modify newsgroups. Table 8.1 lists and describes the parameters that you can use with Rgroup.vbs.

| **Important** |
|---|
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **Run as** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyMachine******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Table 8.1 Rgroup.vbs Parameters**

| Parameter | Description |
|---|---|
| -t *Operation*, where *Operation* is one of the following arguments:<br>• **f**<br>• **a**<br>• **d**<br>• **g**<br>• **l** | Finds a group.<br>Adds a group.<br>Deletes a group.<br>Gets the properties of a newsgroup.<br>Loads newsgroups from an active file. |
| **-s** *Server* | Specifies the name of the server that is affected by this change. |
| **-v** *VirtualServerID* | Specifies the ID of the virtual server. |
| **-c** *CreationDate* | Specifies the date on which the newsgroup is created. |
| **-g** *Group* | Specifies the group to be modified or deleted. The wildcard character (*) can be used during searches (**Rgroup.vbs –t f**). |
| **-n** *NumberOfResults* | Specifies the maximum number of newsgroups to return for a search (required with **Rgroup.vbs –t f**). |

| | |
|---|---|
| **-m** *Moderator* | Specifies the full e-mail address of the moderator. |
| **-d** *Description* | Provides a description of the newsgroup. |
| **-p** *Description* | Specifies the response to the**List PrettyName**s command. |
| **-r true** \| **false** | Indicates whether the newsgroup is read-only. |
| **-u true** \| **false** | Indicates whether the group is moderated by the default moderator. |
| **-a** *FileName* | Specifies the names of the active file. |

Table 8.2 lists examples of how to use Rgroup.vbs.

**Table 8.2 Rgroup.vbs Script Examples**

| Command-line Script | Function |
|---|---|
| **Rgroup.vbs -t f -g alt.*** | Finds all newsgroups beginning with *alt*. |
| **Rgroup.vbs -t d -g my.old.group** | Deletes the newsgroup my.old.group. |
| **Rgroup.vbs -t a -g my.new.group** | Adds the newsgroup my.new.group. |

# Moderating Newsgroups

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you want to control which articles are posted to a newsgroup or if you want to designate someone to control the articles, then create a moderated newsgroup. Articles submitted to a moderated newsgroup are not posted until they are approved by the moderator that you choose for that newsgroup. The moderator can filter inappropriate material and can detect if someone is trying to compromise information on the server or the server itself. To use moderated newsgroups, you must specify either an SMTP mail server to send articles directly to the moderator or a directory where the articles are stored for the moderator.

When a user submits an article to a moderated newsgroup, the NNTP service sends the article to the newsgroup moderator. The moderator can either approve the article and send it back to the NNTP service to be posted, or discard the article. For discarded articles, the moderator can send a message to the author of the article to explain why the article was rejected. Any NNTP client can be used to post messages to a moderated newsgroup. However, the moderator needs to use an NNTP client that supports moderating functions.

Alternatively, you can specify a default moderator domain. If you choose to use the default moderator for a newsgroup, articles are sent to *Newsgroup_Name@DefaultDomain*, where *Newsgroup_Name* is the name of the newsgroup. (The periods are replaced by underscores.) For example, articles posted to the newsgroup rec.sports.tennis, with a default domain of contoso.com, are sent to rec_sports_tennis@contoso.com.

Use the **Settings** tab in the **NNTP virtual server properties** dialog box to enable moderated newsgroups. For step-by-step instructions for enabling and setting up moderated newsgroups, see Moderating Newsgroups in IIS 6.0.

# Managing Newsgroup Messages

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can set expiration policies to delete newsgroup messages automatically after a certain period of time. You can also delete messages that are no longer needed.

#### Setting Expiration Policies

You can specify how long articles are kept in an NNTP service newsgroup by defining expiration policies. An expiration policy can apply to a single newsgroup or to any number of newsgroups. You can define as many policies as you need. In each case, the oldest articles are deleted first.

If you are hosting a newsgroup that is intended for internal project tracking or to record organizational information, then do not set an expiration policy. Manually delete and save articles. If you are hosting a busy newsgroup that receives many articles, set a short expiration policy to avoid overloading your server resources.

If you do not specify an expiration policy for a newsgroup, then delete articles manually when they are no longer needed. You can set and modify expiration policies using IIS Manager, or you can use Rexpire.vbs. To set expiration policies in IIS Manager, right-click **Expiration Policies**, point to **New**, and click **Expiration Policy**.

#### Setting Expiration Policies Using Rexpire.vbs

You can use Rexpire.vbs to set expiration policies. You can also use this command to delete and modify expiration policies. Table 8.3 lists and describes the parameters you can use with Rexpire.vbs.

**Table 8.3 Rexpire.vbs Parameters**

| Parameter | Description |
|---|---|
| -t *Operation*, where *Operation* is one of the following arguments:<br>• **a**<br>• **d**<br>• **g**<br>• **s**<br>• **e** | Adds the expiration policy.<br>Deletes the expiration policy.<br>Gets the expiration policy.<br>Sets the expiration policy.<br>Enumerates the expiration policies. |
| **-s** *Server* | Specifies the name of the server to be configured. |
| **-v** *VirtualServerID* | Specifies the ID of the virtual server. |
| **-I** *ExpireID* | Specifies the ID of the expiration policy. |
| **-h** *ExpireTime* | Specifies the number of hours until articles are discarded. |
| **-d** *MegabytesExpireSize* | Specifies the limit, in megabytes, of the size of the data in the virtual directory at which articles begin to be deleted. The default is 500 MB. |
| **-n** *Newsgroups* | Specifies the newsgroups to which the command is applied. NOTE: You can use strings that contain the wildcard character (*). |
| **-p** *PolicyName* | Specifies the name of the expiration policy. |
| **o- true** \| **false** | Indicates whether this is a one-time expiration policy. |

Table 8.4 lists examples of how you can use Rexpire.vbs.

**Table 8.4 Rexpire.vbs Script Examples**

| Command Script | Function |
|---|---|
| **Rexpire.vbs -t e -v 1** | Enumerates the expiration policies for NNTP site 1. |
| **Rexpire.vbs -t a -n alt.binaries.* -h 24** | Creates an expiration policy that expires articles on alt.binaries that are older than 24 hours. |
| **Rexpire.vbs -t s -i 1 -h 24** | Sets policy 1, which expires any articles older than 24 hours. |
| **Rexpire.vbs -t d -i 1** | Deletes expiration policy 1. |

**Deleting Messages Using Rcancel.vbs**

You cannot delete messages in IIS Manager, but you can delete messages using Rcancel.vbs or control messages. For more information about control messages, see Using Control Messages.

Table 8.5 lists and describes the parameters you can use with Rcancel.vbs. You must supply the message ID to cancel the message.

**Table 8.5 Rcancel.vbs Parameters**

| Parameter | Description |
|---|---|
| **-m** *MessageID* | Specifies the ID of message to cancel. |
| **-s** *Server* | Specifies the computer to configure. |
| **-I** *VirtualServerID* | Specifies the ID of the virtual server. |

The following example shows how you would use Rcancel.vbs:

```
Rcancel.vbs -m "<ugjWStxCBHA.1240@example.microsoft.com>"
```

# Limiting Newsgroup Enumeration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you host many newsgroups on one server, then you might want to limit or filter the newsgroups that are displayed in IIS Manager. You can specify how many newsgroups to display or you can filter the display to show only newsgroups of a certain type. For example, if you have many different departmental newsgroups that are broken down into subgroups, then you might want to filter the display to show only the top-level departmental newsgroups. To limit or filter newsgroup enumeration, in IIS Manager, right-click **Newsgroups**, and then click **Limit groups enumeration**. For more information, see Limiting Newsgroup Enumeration in IIS 6.0.

# Managing Sessions_IIS_SP1_TechRef

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

There might be times when you need to manage or monitor sessions as an administrator. You can terminate a session or all current sessions, or you can enumerate users. For example, if a specific user is posting inappropriate messages to an unmoderated newsgroup, you can terminate the session for that user.

You can manage sessions in IIS Manager or by using the Rsess.vbs command-line script. To manage sessions in IIS Manager, right-click the session, and then click the appropriate action.

**Managing Sessions Using Rsess.vbs**

Table 8.6 lists and describes the parameters that you can use with Rsess.vbs.

**Table 8.6 Rsess.vbs Parameters**

| Parameter | Description |
|---|---|
| **-t** *operation*, where *Operation* is one of the following arguments:<br>• **e**<br>• **d**<br>• **a** | Enumerates current sessions.<br>Deletes the session. (You must also specify **-u** *UserName*or **-I** *IPAddress*.)<br>Deletes all sessions and terminates the client connection. |
| **-u** *UserName* | Specifies the user name to delete. |
| **-I** *IPAddress* | Specifies the IP or host name to delete. |
| **-s** *Server* | Specifies the name of the server to configure. |
| **-v** *VirtualServerID* | Specifies the ID of the virtual server (an integer). |

Table 8.7 lists examples of how you can use Rsess.vbs.

**Table 8.7 Rsess.vbs Script Examples**

| Command Script | Function |
|---|---|
| `Rsess.vbs -t e -v 1` | Lists the current sessions for NNTP site 1. |
| `Rsess.vbs -t d -u john_doe` | Terminates the session for the user john_doe. |
| `Rsess.vbs -t a` | Terminates all sessions. |

| Note |
|---|
| **Rsess.vbs -t e** returns a list of the local feed (that is, the inbound push feed) as an anonymous connection. |

# Using Control Messages

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*Control messages* are specially formatted newsgroup messages that can create newsgroups, delete newsgroups, and delete posts. You can enable or disable control messages for newsgroups and for newsfeeds. If you enable control messages, then they are enabled for every user. If you want to restrict access to control messages to administrators only, then set appropriate ACLs for the control groups and enable either Basic authentication or Integrated Windows authentication.

In IIS Manager, use the **Settings** tab in the **NNTP virtual server properties** dialog box to enable or disable control messages for newsgroups. Use the Rfeed.vbs command-line script to enable or disable control messages for newsfeeds. Control messages are disabled by default in IIS 6.0. However, if you have upgraded from the Microsoft® Windows® 2000 operating system and you enabled control messages in that operating system, control messages are enabled in IIS 6.0.

#### Generating Control Messages

Control messages are generated automatically by NNTP client software, NNTP servers, and special utility programs. You can also generate control messages manually by connecting to the NNTP server using Telnet and entering the control messages directly.

The NNTP service accepts the following basic control messages:

**Cancel**: Deletes a specified article.

**Newgroup**: Creates a newsgroup.

**Rmgroup**: Deletes an existing newsgroup.

The Cancel control message is generated by client programs. The Newgroup and Rmgroup messages are typically generated by NNTP servers.

Control messages look like other newsgroup messages, but control messages have the following headers:

**From**: The name of the user or service from which this message originated.

**Subject**: A description of the subject.

**Newsgroups**: The name of the newsgroup affected by the control message, such as control.cancel|newgroup|rmgroup or any other valid newsgroup.

**Approved**: The user or administrator who created the message. If the newsgroup is moderated, the name of the moderator must appear. If the group is not moderated, any name can be used.

**Control**: The control header. Control messages must contain one of the following control headers:

- **Newgroup** <*name of new group*>. Creates a new group, for example, `newgroup Test.One`.
- **Rmgroup** <*name of group to delete*>. Deletes an existing newsgroup, for example, `rmgroup Test.One`.
- **cancel** <*MessageID*>. Deletes the message that is referenced by the message ID, for example, `cancel <BMO2mVOECHA.2772@server.domain.microsoft.com>`.

A blank line separates the headers from the message body. The message body can contain any text, or it can be blank. The following appears at the end of the control message:

```
{CRLF}.{CRLF}
```

#### Restricting Access to Control Messages

When the NNTP service receives Cancel, Newgroup, and Rmgroup control messages, it posts them into the following newsgroups, respectively: control.cancel, control.newgroup, and control.rmgroup. These newsgroups are created automatically. Control messages are sent to the control newsgroups regardless of where they were originally posted. The NNTP service also automatically creates a control virtual directory.

Restrict access to control messages to prevent unauthorized users from changing your newsgroups or deleting messages. You can restrict access by configuring the control newsgroups as moderated newsgroups or by restricting access to a limited group of users. For more information about moderating and restricting access, see Moderating Newsgroups and Configuring NNTP Virtual Server Properties.

In addition, any ACLs that you have in place for your existing newsgroups apply to control messages. For example, if you have a newsgroup named Microsoft.Software and you have secured the \Inetpub\Nntpfile\Root\Microsoft\Software directory with ACLs that restrict anonymous users from creating subfolders, then when an anonymous user posts a newsgroup control message in Microsoft.Software, the control message is not processed.

# Rebuilding the NNTP Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can fix several common problems with the NNTP service by rebuilding it. Rebuild the NNTP service when:

- The virtual server does not start.
- You delete internal newsgroup files accidentally.
- You recover the system after a disk failure occurs.
- Users cannot gain access to articles that they are authorized to read.
- An error message that is written to the Event Viewer system log tells you to rebuild the server.

For more information about rebuilding the NNTP service, see Rebuilding the IIS NNTP Service.

# Backing Up and Restoring an NNTP Virtual Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can back up your locally hosted NNTP virtual servers so that you can restore them in case of a serious system failure.

To back up an NNTP virtual server, use the **ntbackup** command or the Backup or Restore wizard to back up everything under *LocalDrive*:\Inetpub\Nntpfile. For more information about backing up and restoring Windows Server 2003–based files and directories, see Backing up and restoring data in Help and Support Center for Windows Server 2003.

# Managing Newsfeeds

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Newsfeeds send articles between the IIS NNTP service and other NNTP servers. With newsfeeds, the articles that are posted to one server can be read by users of the other servers. There are several different kinds of newsfeeds:

- Usenet: You can create an external newsfeed to Usenet to enable users of your news site to access the articles that are posted to any Usenet news site. A Usenet provider uses one or more NNTP servers to push new articles to the NNTP service. The NNTP service pushes the articles that are posted by your users to your Usenet provider.

- Peer newsfeed: If your organization has more than one news site, you can share articles between the sites. A peer newsfeed is ideal for private newsgroups that are not appropriate for public distribution on Usenet.

- Master/slave newsfeed: If the volume of traffic on your news site exceeds the capacity of a single computer, you can use multiple computers and share articles among them. A single master news site pushes articles to and from multiple subordinate sites. All clients connect to the subordinate sites.

  The master site controls the article numbers and keeps all the subordinate sites synchronized. Because all the subordinate sites have the same data, a master/slave newsfeed also provides a backup in case one of the computers fails.

Newsfeeds are not enabled by default. For information about enabling newsfeeds, see Enabling Newsfeeds.

**Adding, Deleting, and Modifying Newsfeeds**

Use Rfeed.vbs to add, delete, and modify newsfeeds. Table 8.8 lists and describes the parameters that you can use with Rfeed.vbs.

**Table 8.8 Rfeed.vbs Parameters**

| Parameter | Description |
|---|---|
| **-t** *Operation*, where *Operation* is one of the following:<br>• **a**<br>• **d**<br>• **g**<br>• **s**<br>• **e** | Adds a feed.<br>Deletes a feed.<br>Gets information about a feed.<br>Sets feed information.<br>Enumerates feeds. |
| **-s***Serve* | Specifies the computer name of the server to configure. |
| **-v***VirtualServerID* | Specifies the ID of the virtual server. |
| **-fpeer \|master\| slave** | Indicates the feed type. |
| **-r** *FeedServer* | Specifies the server that you push newsfeeds to or pull newsfeeds from. |
| **-u***UucpName* | Specifies the name that you put in the path header. (The default is the host name of the server.) |
| **-a** *AuthInfoAccount* | Specifies the user name to use when you connect. |
| **-b** *AuthInfoPassword* | Specifies the password to use when you connect. |

| | |
|---|---|
| **-I***FeedID* | Specifies the ID of the specified feed. |
| Parameters for inbound feed settings:<br>• **-ia pull** \| **accept**\|**none**<br>• **-in NewsgroupPatterns**<br>• **-ix***#Minute*<br>• **-iz true**\| **false**<br>• **-im** *#Connection*<br>• **-io** *Date*<br>• **-ic true** \|**false**<br>• **-ip** *PortNumber*<br>• **-it***Directory* | Indicates whether to enable or disable pull feeds.<br>Specifies newsgroup patterns.<br>Specifies the amount of time (in minutes) between feed runs.<br>Indicates whether to process control messages.<br>Specifies the maximum number of connection attempts that are allowed.<br>Gets articles after this date (pull feed only).<br>Indicates whether to create newsgroups automatically during pull feeds.<br>Specifies the outbound IP port for pull feeds.<br>Specifies the temporary files directory. |
| Parameters for outbound feed settings:<br>• **-oa push** \| **none**<br>• **-on Patterns**<br>• **-ox** *#Minutes*<br>• **-oz true**\| **false**<br>• **-om** *#Connections*<br>• **-op***PortNumber*<br>• **-ot***Directory* | Indicates whether to enable a push feed.<br>Specifies newsgroup patterns.<br>Specifies the amount of time (in minutes) between feed runs.<br>Indicates whether to process control messages.<br>Specifies the maximum number of connection attempts allowed.<br>Specifies the outbound IP port for push feeds.<br>Specifies the temporary files directory. |

Table 8.9 lists examples of how you can use Rfeed.vbs.

**Table 8.9 Rfeed.vbs Script Examples**

| Command Script | Function |
|---|---|
| **Rfeed.vbs -t e** | Enumerates feeds. |
| **Rfeed.vbs -t a -r FeedServer.microsoft.com -f peer -ia pull -in *** | Adds a pull feed from peer server FeedServer.microsoft.com, from which the articles from all the newsgroups are to be pulled. |
| **Rfeed.vbs -t d -i 1** | Deletes feed 1. |
| **Rfeed.vbs -t -s -i 1 -ix 120** | Sets the wait time between feed runs to 120 minutes. |

# Configuring the SMTP Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS SMTP service is a simple component for delivering and receiving e-mail messages. A message is transferred to a designated SMTP server to initiate the delivery. The originating SMTP server initiates communications with a DNS server, based on the domain name of the recipient e-mail address. The DNS server looks up and then returns the host name of the destination SMTP server for that domain.

Next, the originating SMTP server communicates with the destination SMTP server directly through TCP/IP port 25. If the user name of the recipient e-mail address matches one of the authorized user accounts on the destination SMTP server, the message is transferred to the destination SMTP sever, where it waits for the recipient to pick up the message through a client program.

The SMTP service also can transfer messages through one or more intermediate relay SMTP servers. A *relay SMTP server* receives the original message and then delivers it to the destination SMTP server or redirects it to another relay server. This process is repeated until the message is delivered or a designated time-out period passes.

The IIS SMTP service is commonly used to create a *smart host*, which is an SMTP server through which all outgoing messages are routed. For example, you can set up a smart host as a stand-alone SMTP server that sits between the Microsoft® Exchange server at your organization and the Internet. In this scenario, the smart host routes all e-mail from the Exchange server to the Internet and provides an added layer of protection between the Internet and the internal network. A smart host can also operate within a network.

Another common use for the SMTP service is to enable e-mail messages to be delivered from a Web site. For example, a company can provide an e-mail link for sending feedback messages or for requesting information. The SMTP service receives and routes those messages to the appropriate mailbox or to another SMTP server. Collaboration Data Objects (CDO) provides a simple object model that applications can leverage to submit mail to the SMTP server. For more information about CDO, see Overview of CDO on MSDN.

# Installing the SMTP Service and Creating an SMTP Virtual Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS SMTP service is not installed by default. Use Control Panel to install the SMTP service, and use IIS Manager to change the default SMTP configuration. For step-by-step instructions for installing the SMTP service, see SMTP Server Setup.

When you install the SMTP service, IIS creates the default SMTP virtual server automatically. In most cases, you need only one SMTP virtual server so that you can modify the default SMTP virtual server to work in your environment. However, if you are hosting multiple domains and want to have more than one default domain, for example, then you can create multiple SMTP virtual servers. To an end user, each SMTP virtual server appears as a separate server with a unique IP address/TCP port combination. To create a new virtual server, right-click **Default SMTP Virtual Server**, point to **New**, and click **Virtual Server.**

**Default Settings**

The default SMTP virtual server has the following default settings.

**Name**: The name of the virtual server that appears in IIS Manager. You can change the name of the virtual server to anything that you want by right-clicking the virtual server in IIS Manager and then clicking Rename.

**IP address/TCP port**: All unassigned/25. You can use the **General** tab in the **SMTP virtual server properties** dialog box to change this setting. If you change this setting, you must specify an IP address/TCP port combination that is not being used by another SMTP virtual server. TCP port 25 is both the default TCP port and the recommended TCP port. More than one virtual server can use the same TCP port, provided that they are configured with different IP addresses. If you do not set a unique IP address/TCP port combination, the SMTP virtual server will not start.

**Default domain**: The domain name that is listed on the **Computer Name** tab in **System Properties**. The default domain is used to stamp messages from addresses that do not have a domain. An SMTP virtual server can have only one default domain, and it cannot be deleted. To change the name of the default domain, double-click the virtual server, and double-click **Domains**. Right-click the local (default) domain, and click **Rename**. For more information about domains, see Organizing Messages Using SMTP Domains.

**Home directory**: *LocalDrive*:\Inetpub\Mailroot. The home directory is the root of your SMTP content directories, and it must be local to the computer on which the SMTP service runs.

If you create a new virtual server, you can configure default settings using the New Virtual Server Wizard.

# Configuring an SMTP Virtual Server as a Smart Host

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This section provides information about configuring an SMTP virtual server as a smart host. Only the properties for which there are special considerations for the smart host scenario are covered. For more complete information about configuring every virtual server property, see SMTP Administration.

**Setting Delivery Properties**

Delivery properties govern how the SMTP service handles outbound e-mail messages and are especially important when you configure a smart host.

To designate a smart host for delivering messages, in the **SMTP virtual server properties** dialog box, click the **Delivery** tab, and then click **Advanced**. In the **Advanced Delivery** dialog box, type the name of the virtual server that will act as the smart host. If the server you are currently configuring will be the smart host, then type its name.

You can also designate a *masquerade domain* in the **Advanced Delivery** dialog box. In a masquerade domain, e-mail that is sent from one domain looks as though it were sent from another. For example, if you have one SMTP virtual server named mail.contoso.com and you want all e-mail to look as though it were sent from contoso.com, then configure a masquerade domain.

If you configure a smart host that sits outside your firewall, click **Outbound Connections** to set the authentication type and provide the necessary credentials for the SMTP server to which your smart host will connect. For more information about authentication methods, see Securing SMTP Virtual Servers.

# Securing SMTP Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

There are many steps you can take to make your SMTP virtual server more secure, including configuring access-related settings such as the authentication method, limiting administrative access, and setting connection limits.

**Configuring Access-Related Settings**
You can restrict access to your SMTP server by requiring authentication, limiting access according to IP address, or both. Restrict access by using the **Access** tab of the **SMTP virtual server properties** dialog box. The **Access** tab has five important security-related settings: Authentication, Certificate, Communication, Connection, and Relay.

**Authentication**
The Authentication option allows you to select from the following methods for authenticating users who attempt to connect to your SMTP server:

- **Anonymous access**. Anonymous access does not require users to enter a user name and password. This option is intended for servers that accept mail from outside the network.

  For example, if you are creating a stand-alone smart host that sits outside of your firewall and in front of your Exchange server, then you might select Anonymous access. Or, if you are creating a server to receive e-mail from a public Web site, then you might select Anonymous access.

- **Basic authentication**. Basic authentication requires users to enter a valid user name and password; however, the credentials are sent across the network unencrypted. If you select Basic authentication, then select **Requires TLS encryption**, which encrypts the user credentials. To require TLS encryption, you must have a valid SSL certificate installed on the server.

- **Integrated Windows authentication**. Integrated Windows authentication requires users to enter a valid Windows user name and password to connect to your SMTP server. Credentials are sent across the network encrypted. You can select Integrated Windows authentication if you are setting up a smart host to relay messages within your network or if you are setting up a server to receive e-mail from internal sites, such as a company intranet.

**Certificate and Communication**
The Certificate and Communication options allow you to secure communication by installing security certificates and requiring encryption. To install a certificate, click **Certificate** to start the New Certificate Wizard. Then, if you want to require the SMTP service to use SSL to encrypt every message, click **Communication**.

**Connection**
The Connection option allows you to restrict access to your server based on IP address. For example, if you are setting up a smart host that works inside your network, then you can use the **Only the list below** option to restrict access to the range of IP addresses for your network. If you are setting up a smart host that relays external e-mail or if you are setting up a server to accept e-mail from a public Web site, then connection control can be an effective way to restrict unsolicited commercial e-mail. After you identify the IP address or addresses of computers that send bulk e-mail, you can use the **All except the list below**option to prevent computers with specific IP addresses from connecting to your server. Although you can restrict access based solely on the domain name, this practice is not recommended because of the resources required to perform a reverse DNS lookup to identify the IP address of the computer attempting to connect.

**Relay**
The Relay option allows you to send e-mail to an SMTP server, which then sends it to the destination server or to another SMTP relay server. By default, the SMTP service blocks all computers from relaying unwanted mail through the virtual server, except those that meet the authentication requirements you designated in the **Authentication** dialog box. If you are setting up a smart host, enable only authenticated users to relay messages. If you are setting up a server to receive e-mail from the Internet, do not allow relaying, which can make you vulnerable to outside users who attempt to send unsolicited commercial e-mail through your SMTP server.

**Limiting Administrative Access**
You can restrict access to your server to only trusted administrators and make sure that administrators are granted only the access that they need to accomplish their work. Limit administrator access by using the **Security** tab of the **SMTP virtual**

**server properties** dialog box.

**Setting Messages Limits**

You can set message limits to prevent large, incoming messages from affecting the performance of your server. Set message limits by using the **Messages** tab of the **SMTP virtual server properties** dialog box.

There are two message limit settings. The first limits the message size. If a mail client sends a message that exceeds the limit, then the mail client receives an error. The second limits the session size, the maximum amount of data that is accepted during the total connection. This setting applies to the sum of all the messages that are sent during the connection and to only the message body.

To help protect against users who send unsolicited, commercial e-mail to your SMTP server, you can limit the maximum number of recipients for a single message sent in one connection. For more information about setting message and recipient limits, see Managing Messages_IIS_SP1_Ops.

# Organizing Messages Using SMTP Domains

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

SMTP domains are used for organizing messages for delivery. There are two main types of domains: local and remote.

**Organizing Messages Using Local Domains**
A *local domain* is a DNS domain that is serviced by the local SMTP server. Any message with a local domain name that arrives at an SMTP server must be delivered to a local mailbox or drop directory, forwarded to another local address, or returned to the sender with a nondelivery report.

You can specify local domains to be either default or alias. The SMTP service uses the default domain to stamp message headers that lack a domain specification. An alias domain is an alias for the default domain. For example, if Contoso Pharmaceuticals merges with Contoso Corporate, then customers might not know which domain to use. If ContosoPharmaceuticals.com is the default domain, then you can set ContosoCorporate.com as an alias domain, which allows users to send mail to either the ContosoPharmaceuticals.com domain or the ContosoCorporate.com domain so that the mail is still delivered to the correct recipient.

**Organizing Messages Using Remote Domains**
A *remote domain* is a domain that is not served by the local SMTP server. When mail is routed to a remote domain and the address is found, the SMTP service sends the mail to another server. Therefore, if you are setting up a smart host as a stand-alone server in front of your Exchange server, then you need to create a remote domain using the domain of your Exchange server.

**Related Information**

- For information about creating alias and remote domains, see Creating SMTP Domains.

# Additional Resources for Configuring Sites and Services in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this section.

**Related Information**

- Analyzing Log Files for information about IIS logs.
- Common IIS 6.0 Administrative Tasks for information about creating a Web site.
- IIS 6.0 Administration Scripts, Tips, and Tricks for information about creating Web sites programmatically.
- IIS 6.0 Architecture for information about the WWW service and IIS core components.
- Managing a Secure IIS 6.0 Solution for information about IIS security.
- Running IIS 6.0 as an Application Server for information about application pools.
- Appendix E: Unattended Setup for information about unattended installation.
- Working with the Metabase for information about the IIS metabase.

**Related Windows Server 2003 Help Topics**
For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** checkbox.

- Access control in Help and Support Center for Windows Server 2003.
- Active Directory in Help and Support Center for Windows Server 2003 for information about using Active Directory.
- Backing up and restoring data in Help and Support Center for Windows Server 2003 for information about how to back up servers running Windows Server 2003.
- Choosing a file system: NTFS, FAT, or FAT32 in Help and Support Center for Windows Server 2003 for information to help you decide which file system is the best solution for your server.
- DNS in Help and Support Center for Windows Server 2003 for information about the Windows Server 2003 DNS feature.
- Namespace planning for DNS in Help and Support Center for Windows Server 2003 for information about planning namespaces for DNS.

# IIS 6.0 Administration Scripts, Tips, and Tricks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Day-to-day administration of Internet Information Services (IIS) 6.0 includes tasks such as creating and configuring Web and File Transfer Protocol (FTP) sites, managing server certificates, and backing up IIS–based content and configuration. To work more efficiently, you can perform most of these administration tasks programmatically and remotely, whether you are administering IIS in a single-site or multi-site environment.

**In This Section**


Creating and Configuring a Site Programmatically

Creating a Web Site from a Template

Comparing Programmatic Administration Methods for IIS 6.0

Managing Server Certificates Programmatically

Managing IIS 6.0 Remotely

Scheduling IIS 6.0 Backups

Restarting and Alternatives to Restarting IIS 6.0

Additional Resources for IIS 6.0 Administration


**Related Sections**

- For information about specific commands, including command syntax, see Command-line reference A-Z in Help and Support Center for Microsoft® Windows® Server 2003.
- For information about configuring Web sites as well as the File Transfer Protocol (FTP), Network News Transfer Protocol (NNTP), and Simple Mail Transfer Protocol (SMTP) services, see Configuring Internet Sites and Services.
- For information about installing and using FrontPage® 2002 Server Extensions from Microsoft, see Using FrontPage 2002 Server Extensions with IIS 6.0.

# Creating and Configuring a Site Programmatically

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

One of the most common tasks for Internet service providers (ISPs) and administrators who manage complex intranet sites hosted on servers running Internet Information Services (IIS) 6.0 is creating and configuring new Web sites and File Transfer Protocol (FTP) sites. Although you can create these sites by using IIS Manager, if you create sites often or create many sites at once, it can be faster and easier to create sites programmatically. For more information about creating Web sites and FTP sites by using IIS Manager, see Configuring Internet Sites and Services.

You can use a combination of scripts and tools to create a batch file that will automate all of the creating and configuring tasks. The steps for creating the batch file are as follows:

- Set up and configure DNS.
- Create directories.
- Add temporary content.
- Create Web sites.
- Set up FrontPage® 2002 Server Extensions from Microsoft.
- Create FTP sites.
- Set access control lists (ACLs).
- Include verification and error checking.

The following sections provide details, including sample scripts, for each step of the process for creating a batch file. In the final section, a complete batch file that you can customize for your environment is provided.

**Variables**
Some of the sample scripts in the following sections contain one or more of the following variables:

**%1** =User

**%2** = Domain

**%3** = .com (extension)

**%4** = y or n (indicates whether FrontPage 2002 Server Extensions should be installed)

If you type **newsite.bat user domain.com y** at a command prompt, after you save the complete batch file in a text editor, each entry becomes a variable in the script. The variable **%1** represents User; the variable **%2%3** represents Domain.com; and the variable **%2** represents Domain.

# Setting Up and Configuring DNS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Domain Name System (DNS) is an important part of creating new sites, especially if you host many sites on a single server with one Internet Protocol (IP) address. For more information about DNS and IP addresses, see Configuring Internet Sites and Services.

One tool that can be helpful in managing DNS is the **Dnscmd.exe** command-line tool that is provided with the Support Tools for Microsoft® Windows® Server 2003, Standard Edition; Microsoft® Windows® Server 2003, Enterprise Edition; and Microsoft® Windows® Server 2003, Datacenter Edition operating systems. You can use **Dnscmd.exe** to view, manage, and update existing DNS server configurations or to set up and configure new DNS servers on your network. Before you use Dnscmd.exe, decide how best to configure DNS zones for your network, or, if you have a designated DNS administrator for the network, ask him or her which parameters to use. For more information about using Dnscmd.exe, see Server administration using Dnscmd in Help and Support Center for Microsoft® Windows® Server 2003.

Listing 9.1 shows how you can use Dnscmd.exe to set up a DNS zone for contoso.com, create entries for contoso.com and www.contoso.com that point to IP address 10.1.1.1, and then set up a canonical name (CNAME) for www.contoso.com that instructs the DNS server to use the same IP address that is configured for contoso.com.

| **Important** |
| --- |
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /User:***MyMachine******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

**Listing 9.1 Sample Script for Setting Up and Configuring DNS**

```
dnscmd PRIMARY /zoneadd contoso.COM /primary /file contoso.COM.dns
//sets up the overall zone contoso.COM on the primary name server

dnscmd SECONDARY /zoneadd contoso.COM /secondary /file contoso.COM.dns
//sets up the zone on the secondary

dnscmd PRIMARY /recordadd contoso.COM @ A 1.1.1.1
//sets up the record on the primary server for contoso.COM pointing to the IP 10.1.1.1

dnscmd PRIMARY /recordadd contoso.COM www CNAME contoso.COM
//sets up a CNAME record for WWW.contoso.com to point to the same IP as
// contoso.COM
```

# Creating Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you set up and configure DNS, you should create the physical directory for your sites. You can use the **mkdir** command to create a directory programmatically. For more information about **mkdir**, including command syntax, see Mkdir in the command-line reference section of Help and Support Center for Windows Server 2003.

Although it is simple to create a directory, the task requires advanced planning to create logical structures. Listing 9.2 uses C:\Web\UserName\Domain, where a users content is in a directory that has the same name as the user name. Each users domain is a subdirectory of the user name directory. This structure allows users to use FTP to download files to their folders and access other directories as necessary.

**Listing 9.2 Sample Script for Setting Up Directories**

```
mkdir C:\WEB\USERNAME\
mkdir C:\WEB\USERNAME\DOMAIN\
```

To prevent mistakes, such as unintentionally overwriting an existing directory, include error handling in your script. You can use the **if [not] exist** command in batch programs to check for and respond to certain conditions, as seen in Listing 9.3. For more information about the **if** commands, including command syntax, see If in the command-line reference section of Help and Support Center for Windows Server 2003.

**Listing 9.3 Sample Script for Error Handling**

```
If NOT exist C:\WEB\USERNAME\ goto errornofolder1
If NOT exist C:\WEB\USERNAME\DOMAIN\ goto errornofolder2
```

**Creating Directories Remotely**

| **Important** |
| --- |
| As a best security practice, use the **net use** command to authenticate with the remote computer before you use **mkdir** remotely. For more information about the **net use**command, including command syntax, see Net services commands in the command-line reference section of Help and Support Center for Windows Server 2003. |

To create directories on a remote computer, use the following command-line syntax:

```
\\Server\c$\Web\UserName\
```

# Adding Temporary Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you create sites, you can add temporary content files that users can replace with their own content when they are ready. An example of a temporary content file is a home page that says Coming Soon or Under Construction. Use the Adsutil.vbs command-line script to add temporary content files.

Adsutil.vbs is a flexible tool that uses the Microsoft® Visual Basic® development system, Scripting Edition (VBScript), with Microsoft® Active Directory® Service Interfaces (ADSI) to manipulate the IIS configuration. Note that Adsutil.vbs is not a Microsoft-supported product. The documentation for, and the format of, Adsutil.vbs can change at any time. For more information about Adsutil.vbs, including syntax and parameters, see Using the Adsutil.vbs Administration Script.

You must use the CScript.exe command-line tool, which is installed with Windows Script Host, to run Adsutil.vbs. You can set CScript.exe as the default run-time program by typing **CScript //h:cscript //nologo //s** at the command prompt.The optional **//nologo** parameter disables the display of the CScript copyright banner. Alternatively, you can type **Cscript.exe** at the beginning of the Adsutil.vbs script. For more information about using Adsutil.vbs, including commands and parameters, see Using the Adsutil.vbs Administration Script.

Listing 9.4 uses Adsutil.vbs to add a file named temp.htm to the list of default documents. Although you do not have to use the name temp.htm, do not name the file default.htm because default.htm would take precedence if the user added only index.htm as the home page. The sample script adds the temp.htm file as the last file on the **Documents** tab in the Web site properties so that temp.htm can never override a users home page, even if the user never deletes temp.htm. You must type the complete Adsutil.vbs script on one line.

**Listing 9.4 Sample Script for Setting Up Temporary Files**

```
Adsutil.vbs set w3svc/defaultdoc
"default.htm,default.html,default.asp,default.aspx,index.htm,index.html,index.
asp,index.aspx,temp.htm"
```

After you add a temporary file, use the **copy** command to copy the file to the users directory. For example, to copy temp.htm to the users directory, at the command prompt, type **copy temp.htm C:***web******username******domain***\temp.htm**.

**Adding Temporary Content Remotely**

You can run any of the Adsutil.vbs commands remotely. To add temporary content to a remote computer, add **-s:***Server* to the end of the command, where *Server* is the name of the target server.

# Creating Web Sites in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you configure DNS, create the directories, and populate the directories with temporary content, you need to create the Web sites. You can use IIS Manager to create Web sites, or you can use the Iisweb.vbs command-line script, which is stored in the *systemroot*\System32 folder.

When you use Iisweb.vbs to create a Web site, you specify only the basic properties that are needed to create the site and identify its contents. Iisweb.vbs uses the same default properties that IIS Manager uses when creating new Web sites, and it adheres to the same rules for inheriting properties. For more information about default Web site properties, see Configuring Internet Sites and Services. For more information about Iisweb.vbs, including command syntax, see Creating Web Sites Using Iisweb.vbs. For example, to create a Web site on the local computer using the directory specified in the preceding sample scripts, at a command prompt, type **iisweb /create** *c:\web\user\domain contoso.com* **/d** *www.contoso.com*.

**Creating Web Sites Remotely**

To create a Web site on a remote computer, type **/s** *Server*after the **/create** parameter, where *Server* is the name of the target computer. To specify the IP address of the target computer, type **/i** *IPAddress* after the **/s** parameter, where*IPAddress* is the IP address.

# Providing Additional Host Headers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Many users request additional host headers, such as contoso.com, without the www prefix. You can use Adsutil.vbs to set host headers, but you need to specify the site identification (ID). Although you can find out the site ID by parsing it from Adsutil.vbs output, an easier way is to integrate the Translate.js command-line script, which uses the IIS Windows Management Instrumentation (WMI) provider, into your Adsutil.vbs command. For more information about the IIS WMI provider, see IIS WMI Provider on MSDN.

The Translate.js script appears below. Copy it exactly as it is shown, and then save it in the same location as your other scripts. If you do not specify a location, it will be saved to C:\Windows\System32 by default. You need to use CScript.exe to run Translate.js. For information about using CScript.exe, see Adding Temporary Content.

```
TRANSLATE.JS
var serverComment = WScript.Arguments(0);
var query = "SELECT Name, ServerComment FROM IIsWebServerSetting WHERE ServerComment='"
+ serverComment + "'";
var providerObj = GetObject("winmgmts:/root/MicrosoftIISv2");
var sites = providerObj.ExecQuery(query);

if (sites.Count != 0)
{
   for(e = new Enumerator(sites); ! e.atEnd(); e.moveNext())
   {
      var item = e.item();
      // Print the site ID only by leaving the "W3SVC/" prefix out.
      WScript.Echo(item.Name.substr(6));
   }
}
else
{
   WScript.Echo("No sites found.");
}
```

After you save Translate.js, you can integrate it with Adsutil.vbs to provide additional host headers. Listing 9.5 uses 80:contoso.com to designate the site IP binding as All Unassigned. If you want to designate a specific IP address, type the IP address followed by a colon (:) in front of this string, as in **IPAddress:80:contoso.com**.

**Listing 9.5 Sample Script for Setting Up Host Headers**

```
cscript //nologo translate.js "CONTOSO.COM" > siteid.txt

for /f %%I in (siteid.txt) do SET SITEID=%%I

Adsutil set w3svc/%SITEID%/serverbindings ":80:WWW.CONTOSO.COM" ":80:CONTOSO.COM"

Del siteid.txt
```

**Providing Additional Host Headers Remotely**

| **Important** |
| --- |
| To provide additional host headers remotely while logged in with a domain account, you must have administrative credentials on both computers. When logged in with a local user account, you must have an account with the same user name and password on the remote computer, as well as administrative credentials on both computers. |

To modify Translate.js to access a remote server, add the server name when you get the object by changing **GetObject(winmgmts:/root/MicrosoftIISv2);** to **GetObject(winmgmts://***Server***/root/MicrosoftIISv2);**, where *Server*is the name of the remote server.

# Installing and Setting Up FrontPage 2002 Server Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage 2002 Server Extensions provide Web-based and command-line administration for extending virtual servers. *Extending*a virtual server means enabling various FrontPage 2002 Server Extensions features to improve the way that you manage the content development and security of your site. Extending virtual servers with FrontPage 2002 Server Extensions enables the site owner to author the site in FrontPage and delegate site ownership and administration credentials.

FrontPage 2002 Server Extensions include new security features, such as roles and rights, and new features for monitoring server health and Web site use. For more information about improvements in FrontPage 2002 Server Extensions, see Using FrontPage 2002 Server Extensions with IIS 6.0.

Listing 9.6 enables both FrontPage 2002 Server Extensions and FTP on the same site. The sample script sets up FTP by default and checks the *%4*variable. If the *%4* variable is **y**, the script goes to the FrontPage 2002 Server Extensions setup section. Otherwise, FrontPage 2002 Server Extensions setup is skipped.

**Listing 9.6 Sample Script for Setting up FTP and FrontPage 2002 Server Extensions**

```
:iffp
if "%4"=="y" goto fpse
goto skipfp
```

There are two options for setting up FrontPage 2002 Server Extensions. You can add the full FrontPage 2002 Server Extensions path to System Properties, or you can include the path in the batch file each time you run it. The locations are as follows:

- For FrontPage 2002 Server Extensions: *Drive*:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\50\Bin

- For FrontPage 2000 Server Extensions: *Drive*:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\40\Bin

**To add the FrontPage 2002 Server Extensions path to System Properties**

1. From the **Start**menu, right-click **My Computer**, and then click **Properties**.
2. In the **System Properties**dialog box, click the **Advanced** tab, and then click **Environment Variables**.
3. In the **Environment Variables**dialog box, in **System variables**, click the path, and then click **Edit**.
4. In the **Edit System Variable** dialog box, in the **Variable value** box, type a semicolon (**;**) at the end of the path, and then do one of the following:
   - For FrontPage 2002 Server Extensions, type: *Drive***:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\50\bin**
   - For FrontPage 2000 Server Extensions, type: *Drive***:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\40\bin**

To include the FrontPage 2002 Server Extensions path in the batch file each time you run it, run a script like Listing 9.7. This sample script does not add the path to System Properties.

**Listing 9.7 Sample Script for Installing FrontPage 2002 Server Extensions Batch Files**

```
@set path=%path%;%SYSTEMDRIVE%\inetpub\adminscripts;%SystemDrive%\Program Files\Common
Files\Microsoft Shared\Web Server Extensions\50\bin
```

After you set up FrontPage 2002 Server Extensions, you can use SharePoint™ Team Services from Microsoft command-line tools to install and administer FrontPage 2002 Server Extensions from the local server or from a remote computer. The Owsadm.exe command-line tool runs on only the server, and the Owsrmadm.exe command-line tool can be run from a remote

computer. These tools are available at *systemroot*\Program Files\Common Files\Microsoft Shared\Web Server Extensions\50\Bin.

**Installing FrontPage 2002 Server Extensions Locally**

To view the command-line options for Owsadm.exe, at the command prompt, change the directory to the *systemroot*\Program Files\Common Files\Microsoft Shared\Web Server Extensions\50\Bin folder, and then type **owsadm.exe -h**.

| **Note** |
| --- |
| In the following sample scripts, CONTOSO.COM is case sensitive. |

Listing 9.8 uses Owsadm.exe to install FrontPage 2002 Server Extensions on the local server.

**Listing 9.8 Sample Script for Installing FrontPage 2002 Extensions System Properties**

```
:fpse
owsadm -o install -p 80 -m www.CONTOSO.COM -t msiis -u USER
```

To install FrontPage 2000 Server Extensions, use Listing 9.9.

**Listing 9.9 Sample Script for Installing FrontPage 2000 Extensions System Properties**

```
:fpse
fpsrvadm -o install -p 80 -m www.CONTOSO.COM -t msiis -u USER
```

**Security Considerations for Installing FrontPage 2002 Server Extensions Remotely**

Be aware of security considerations when using Owsrmadm.exe. When you administer a server remotely, a wider community of users is given greater access to the Web server from the Internet, which can create a security risk. During remote administration of Windows SharePoint Team Services and FrontPage 2002 Server Extensions, an unauthorized person can gain access to Web sites, based on the Windows SharePoint Team Services or FrontPage 2002 Server Extensions on your server, and modify Web site settings, and even delete Web sites. To increase security when administering FrontPage 2002 Server Extensions remotely, you should use a Secure Sockets Layer (SSL) connection, and you should use a nonstandard administration port to access Fpadmdll.dll.

To require SSL security for authoring and site administration tasks, use the **Require SSL for authoring and administration** option on the **Set Installation Defaults** tab in **Server Administration**, or set the **RequireSSL** property by using the **setproperty** command.

When you install FrontPage 2002 Server Extensions, a nonstandard administration port is created. Use this administration port to access Fpadmdll.dll, which, on the Web server, acts as the form handler for any of the FrontPage 2002 Server Extensions HTML–based administration pages.

For more information about using Owsrmadm.exe, see Using the Administration Tools Remotely in the SharePoint Team Services Administrator's Guide.

# Creating FTP Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The directory structures for the sample scripts in this section include an overall user folder that contains a folder for each of the user's domains. For example, if the user has two domains, the structure would be *C:\Web\User\Domain1* and *C:\Web\User\Domain2*. For the purposes of uploading content through FTP, you need to set up an FTP virtual directory that points to C:\Web\User so that the user can upload to either domain easily. Instructions for setting up a virtual directory are provided later in this section.

You can have only one FTP site per IP address on TCP port 21, so you need to use virtual directories and ACLs to ensure that users are sent to the correct directory. You can create and host multiple FTP sites with the same IP address, as long as they have different ports. This configuration can be difficult to manage. It requires making sure that the port number is explicitly defined in the URL, and you might also need to open the port at the firewall.

If you want to configure multiple users to a single FTP site but also ensure that users cannot access and download from the directories of other users, you can create FTP sites by using FTP isolation. Some of the benefits of FTP isolation include added security and easier migration of users to different servers.

FTP user isolation supports two isolation modes: Isolate users and Isolate users using Active Directory. You can also create FTP sites with no isolation. You can configure FTP user isolation with the Iisftp.vbs command-line script by using the **/isolation** parameter with either the **AD** argument for Active Directory isolation, or the **Local** argument, for local isolation. If you do not include the **/isolation** parameter, the site will not isolate users. For more detailed information about creating multiple FTP sites and setting up FTP user isolation, see Configuring Internet Sites and Services.

# Setting Up FTP Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can set up an FTP virtual directory by using lisftpdr.vbs. The directory structure in these sample scripts has a domain beneath each user folder. If you set up your directory this way, when your users upload content they can change directories into any domain they own. For more information about lisftpdr.vbs, including command syntax, see Creating FTP Virtual Directories Using lisftpdr.vbs.

For example, to use one FTP site that points to the same IP address as the Web sites and creates a virtual directory for a user named USER, at the command prompt, type **iisftpdr /create** *ftpsitename* **USER** *c:\web\user*. In Listing 9.10, Translateftp.js gets the site ID and uses an Adsutil.vbs command to assign write permissions for the user on the new virtual directory. The sample script creates Translateftp.js by copying Translate.js, renaming it Translateftp.js, and changing two elements.

The first element that is changed is the IISWebServerSetting class. Translate.js queries the IISWebServerSetting class by using WMI. For FTP, you must switch to the IISFtpServerSetting class. This class corresponds to the IIsFtpServer IIS Admin object, and it contains the read/write properties for the object. For more information about the IISWebServerSetting element, see see the IIS Programmatic Administration Reference ion MSDN.

The second element that is changed is the JavaScript *item.Name***.substr()** call. The **substr()** method returns a specified part of a string. By changing the setting from six (6) to nine (9) for Translateftp.js, the script will return the FTP site ID as a return string.

**Listing 9.10 Sample Script for Setting UP FTP Virtual Directories**

```
TRANSLATEFTP.JS
var serverComment = WScript.Arguments(0);
var query = "SELECT Name, ServerComment FROM IIsFTPServerSetting WHERE ServerComment='"
+ serverComment + "'";
var providerObj = GetObject("winmgmts:/root/MicrosoftIISv2");
var sites = providerObj.ExecQuery(query);

if (sites.Count != 0)
{
   for(e = new Enumerator(sites); ! e.atEnd(); e.moveNext())
   {
      var item = e.item();
      //Print the site ID only by leaving the "MSFTPSVC/" prefix out.
      WScript.Echo(item.Name.substr(9));
   }
}
else
{
   WScript.Echo("No sites found.");
}

The ADSUTIL command and integration with translateftp.js.

cscript //nologo translate.js "FTPSITEID" > ftpsiteid.txt

for /f %%I in (ftpsiteid.txt) do SET FTPSITEID=%%I

adsutil.vbs set msftpsvc/%FTPSITEID%/root/USER/accesswrite "true"

del ftpsiteid.txt
```

**Setting Up FTP Remotely**

| **Important** |
| --- |
| To set up FTP remotely while logged in with a domain account, you must have administrative credentials on both computers. When logged in with a local user account, you must have an account with the same user name and password on the remote computer, as well as administrative credentials on both computers. |

To modify Translateftp.js to access a remote server, add the server name when you get the object by changing **GetObject(winmgmts:/root/MicrosoftIISv2);** to **GetObject(winmgmts://***Server***/root/MicrosoftIISv2);**, where *Server* is

the name of the remote server.

# Setting ACLs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Depending on the requirements of your Web sites and services, you might want to explicitly set ACLs for groups and accounts in IIS. For example, you might need to ensure that ASP.NET can run on a Web site. You can use one of several tools to programmatically display and modify ACLs for your physical folders. For a list of the default IIS ACLs, and a list of popular ACL tools, see Managing a Secure IIS 6.0 Solution.

| **Important** |
|---|
| If you use FrontPage2002 Server Extensions and FTP on the same site (which is not recommended), you should set ACLs after you install FrontPage2002 Server Extensions to prevent FrontPage2002 Server Extensions from changing the permissions. |

For example, to edit the ACLs instead of replacing them, so that the defaults will be inherited from c:\web, and the FrontPage 2002 Server Extensions permissions will be left intact, at the command prompt, type **cacls** *c:\web\user* **/t /e /g user:c networkservice:r <***y.txt*.

You will need to use a text file to simulate pressing ENTER+Y+ENTER, which confirms the operation. To do so, create a .text file named Y.txt, type the following, and save the file:

**<enter>**

**Y**

**<enter>**

Because this sample script uses cacls, the account name is specified as networkservice instead of Network Service, but both names refer to the same account.

#### Recommended and Inherited ACLs

Table 9.1 lists several accounts or groups with ACLs that you might need to change, as well as an explanation for the recommendation.

**Table 9.1 Recommended ACLs**

| Account or Group Name | Recommended ACL | Explanation |
|---|---|---|
| IUSR_*ComputerName* | Read | Anonymous browsing account (used by default). |
| Network Service | Read | ASP.NET requires read access on the Network Service account.<br>**Note**<br>IIS 6.0 runs under this account by default when IIS is running in worker process isolation mode. In IIS 5 compatibility mode, the default account is ASPNET. |
| USER | Change | Account for the site administrator. |

Table 9.2 lists the ACLs that are inherited from the folder structure, but that can be added in the script.

**Table 9.2 Inherited ACLs**

| Account or Group Name | Inherited ACL | Explanation |
|---|---|---|
| Administrators | Full | Account for users in the Administrators group. |
| System | Full | Account for users who need system access. |

#### Setting ACLs Remotely

To run cacls remotely, use the following command-line syntax, where *Server* is the name of the target server, and *c$\web* is the path:

**cacls ** *Server\c$\web*

# Including Verification and Error Checking

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

As a best practice, you should place error-checking and verification code at the start of and throughout your batch file. If you plan to be present while you run the batch file, you do not need extensive verification.

One way to have your batch file check for errors is to include a **net user** command that verifies that the user name is in the domain and then parses the command output for any text that says the user does not exist. For example, use the following command-line syntax to verify that the user name is in the domain:

**net user** *user/domain* **>** *user.txt*

Then, parse the command output for the text string The user name could not be found, which would indicate that the user is not in the domain. You can use the **findstr** command to parse a file for a text string.

Another way to have your batch file check for errors is to include the **if not exist** command and other commands that ensure that the previous portion of the script succeeded. Ideally, you should run this script after each series of commands to check, for example, that the DNS entries were created, and then to check that the directories were created, and so on. You should also include error-handling code after the error-checking code, such as a loop that waits a few minutes to attempt to complete a series of commands again, or a more complex series of commands, while the batch file parses a specific error.

Listing 9.11 provides simple error checking for the existence of all variables. In addition, this sample script allows another administrator, who might not be familiar with your script, to use the **/?** parameter to get information on how to use the script.

**Listing 9.11 Sample Script for Error Checking**

```
if "%1" == "?" goto usage
if "%1" == "/?" goto usage
if "%1" == "-?" goto usage


if "%1" == "" goto errormissingvariable
if "%2" == "" goto errormissingvariable
if "%3" == "" goto errormissingvariable
if "%4" == "" goto errormissingvariable
```

# The Complete Batch File

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Listing 9.12 provides the complete batch file named Newsite.bat. The complete batch file includes the series of commands that are documented in the preceding sections, variables for elements such as the server IP address and DNS servers, and the @ symbol in front of certain lines to hide them from the console window.

After you save Newsite.bat, type the following at the command prompt to run it:

**newsite.bat user domain.com n**

**%1 = user**

**%2 = domain**

**%3 = .com** (extension)

**%4 = y|n** (to indicate whether to install FrontPage 2002 Server Extensions)

**Listing 9.12 NEWSITE.BAT: Sample Batch File**

```
@SET PRIMARYNS=dns.contoso.com
@SET SECONDARYNS=dns1.contoso.com
@SET IP=1.1.1.1
@SET LOC=c:\web
@SET FTPSITE=ftp_computername
@SET ANONACCT=IUSR_computername
@SET IIS_ASPNET_ACCT=networkservice

@if "%1" == "?" goto usage
@if "%1" == "/?" goto usage

@if "%1" == "" goto errormissingvariable
@if "%2" == "" goto errormissingvariable
@if "%3" == "" goto errormissingvariable
@if "%4" == "" goto errormissingvariable

@set path=%path%;%SYSTEMDRIVE%\inetpub\adminscripts;%SystemDrive%\Program Files\Common
Files\Microsoft Shared\Web Server Extensions\50\bin

:dns
dnscmd %PRIMARYNS% /zoneadd %2%3 /primary /file %2%3.dns
dnscmd %SECONDARYNS% /zoneadd %2%3 /secondary /file %2%3.dns
dnscmd %PRIMARYNS% /recordadd %2%3 @ NS %SECONDARYNS%
dnscmd %PRIMARYNS% /recordadd %2%3 @ A %IP%
dnscmd %PRIMARYNS% /recordadd %2%3 www CNAME %2%3

:dir
mkdir %LOC%\%1
mkdir %LOC%\%1\%2
copy temp.htm %LOC%\%1\%2\temp.htm

:web
iisweb /create %LOC%\%1\%2 "%2%3" /d www.%2%3

cscript //nologo translate.js "%2%3" > siteid.txt
for /f %%I in (siteid.txt) do SET SITEID=%%I
Adsutil set w3svc/%SITEID%/serverbindings ":80:WWW.%2%3" ":80:%2%3"
Del siteid.txt

:iffp
if "%4"=="y" goto fpse
goto skipfp

:fpse
owsadm -o install -p 80 -m www.%2%3 -t msiis -u %1

:skipfp
```

```
:ftp
Iisftpdr /create "%FTPSITE%" %1 %LOC%\%1

cscript //nologo translateftp.js "%FTPSITE%" > ftpsiteid.txt
for /f %%I in (ftpsiteid.txt) do SET FTPSITEID=%%I
adsutil.vbs set msftpsvc/%FTPSITEID%/root/%1/accesswrite "true"
del ftpsiteid.txt

:acls
cacls %LOC%\%1 /t /e /g %1:c %IIS_ASPNET_ACCT %:r <y.txt
@goto end

:errormissingvariable
@echo You are missing a variable.
@goto usage

:usage
@echo Creates a domain on the LOCAL server.
@echo var1=username, var2=domain (without extension!), var3=extension, var4=y (for
FrontPage 2002 Server Extensions, n  for none.)
@echo EXAMPLE myusername mydomain .com y
@goto end

:end
```

# Comparing Programmatic Administration Methods for IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 can be configured programmatically from standalone scripts or applications. Table 9.3 describes these methods of configuring IIS, and lists some of the advantages and disadvantages of using each method.

**Table 9.3 Comparison of Programmatic Administration Methods for IIS 6.0**

| Configuration Method | Advantages | Disadvantages |
|---|---|---|
| Command-line tools included with IIS | • Provides solutions to the most common administration tasks. | • Some tools are not available on earlier versions of IIS. |
| Active Directory Service Interfaces (ADSI) | • Can configure large sites and multiple servers.<br>• Supports scripting.<br>• Applies changes immediately without restarting the server.<br>• Supports IIS 4.0 through IIS 6.0. | • Requires knowledge of which supporting properties need to be created and set.<br>• Some functionality is missing when using the IIS ADSI provider from System.DirectoryServices. |
| Windows Management Instrumentation (WMI) | • Can configure large sites and multiple servers<br>• Supports scripting.<br>• Applies changes immediately without restarting the server. | • Requires knowledge of which supporting properties need to be created and set.<br>• Cannot be used when you are configuring a server running IIS 5.1 or earlier. |
| Admin Base Objects (ABO) | • Provides a faster method of configuring IIS than ADSI or WMI. | • Requires knowledge of which supporting properties need to be created and set.<br>• Does not support scripting.<br>• Only supports applications written in C++ or Visual Basic 6.0.<br>• Is more difficult to use than ADSI or WMI.<br>• Provides no safeguards to prevent configuration of invalid settings. |

For more detailed information about the programmatic administration methods described in Table 9.3, see the IIS Software Development Kit on MSDN.

# Managing Server Certificates Programmatically

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

There are many ways to manage server certificates programmatically. This section provides a description of two different methods, including sample scripts for performing management tasks related to server certificates. One method uses the **IISCertObj** COM object, and the other method uses a *CAPICOM*-based script. CAPICOM is a Cryptographic interface (CryptoAPI).

# Managing Server Certificates by Using IISCertObj

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can use **IISCertObj** to perform the following tasks:

- Import copies of a server certificate from a .pfx file (a file containing a PFX-encrypted server certificate and a private key) from a server to multiple servers.
- Save backup copies of server certificates on a central site.
- Copy a server certificate from one server to multiple servers.

When using **IISCertObj**, you must set properties for the methods that you call. The following sections outline the necessary properties and the methods that you need to manage server certificates.

**Setting IISCertObj Properties**

Because a method will return an error message if the required properties for that method are not available when the method is called, you should set properties before you call methods. Table 9.4 lists the **IISCertObj** properties that you need to generate server certificate requests.

**Table 9.4 Required IISCertObj Properties**

| Property Attribute | Name Attribute | Value | Description |
|---|---|---|---|
| `ServerName` | Data Type | String, write only | Contains the name of the computer on which server certificate operations are executed. |
| `UserName` | Data Type | String, write only | Specifies the user name that is used to log on to the remote computer. If UserName is empty, the credentials of the user that is currently logged on are used. |
| `UserPassword` | Data Type | String, write only | Specifies the password that is used for logging onto the remote computer. |
| `InstanceName` | Data Type | String, write only | Identifies the targeted metabase instance. W3svc/1 is an example. |

**Applying IISCertObj Methods**

You can apply the following methods to the **IISCertObj** object: **Copy**, **Export**, **GetCertInfo**, **Import**, **ImportToCertStore**, **IsExportable**, **IsInstalled**, **Move**, and **RemoveCert**.

**Copy**

The **Copy** method allows you to copy a server certificate from one instance of IIS to another and from one computer to another. Use the following syntax with the Copy method:

```
            bAllowExport, bOverwriteExisting, DestinationServerName,
DestinationServerInstance, DestinationServerUserName, DestinationServerPassword
```

Table 9.5 lists the parameters that the Copy method uses.

**Table 9.5 Copy Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|

| | | |
|---|---|---|
| `bAllowExport` | Boolean | Indicates whether the imported server certificate will be exportable (true) or not exportable (false). |
| `bOverwriteExisting` | Boolean | Indicates whether, if an identical server certificate already exists, it is overwritten (true) or not overwritten (false). |
| `DestinationServerName` | String | Specifies the computer name of the server to which the server certificate will be moved. |
| `DestinationServerInstance` | String | Specifies the IIS metabase instance at which the server certificate will be pointed. For example, w3svc/1. |
| `DestinationServerUserName` | String | Specifies the user name for the destination server. If empty, the user name of the user who is currently logged on will be used. |
| `DestinationServerPassword` | String | Specifies the password for the destination server when the *DestinationServerUserName* parameter is used. |

**Export**

The **Export** method exports a copy of a server certificate to a file. The target file can be on a local or remote computer. Use the following syntax with the Export method:

```
Export FileName Password bPrivateKey bCertChain bRemoveCert
```

Table 9.6 lists the parameters that the Export method uses.

**Table 9.6 Export Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| `FileName` | String | Specifies the name of the targeted file. For example, C:\Mydir\Mycert.pfx. |
| `Password` | String | Specifies the password that is used to secure the file specified as *FileName*. |
| `bPrivateKey` | Boolean | Indicates whether the private key is exported (true) or not exported (false). |
| `bCertChain` | Boolean | Indicates whether the certificate trust chain is exported (true) or not exported (false). |
| `bRemoveCert` | Boolean | Indicates whether the binding of a server certificate is removed (true) or not removed (false). |

**GetCertInfo**

The **GetCertInfo** method allows you to acquire information on a server certificate that is installed on the computer specified by the **IISCertObj** properties, **ServerName** and **InstanceName**.

The GetCertInfo method uses no parameters.

**Import**

The **Import** method imports a copy of a server certificate from a file to a local or remote computer. Use the following syntax with the Import method:

```
FileName Password bAllowExport bOverWriteExisting
```

Table 9.7 lists the parameters that the Import method uses.

**Table 9.7 Import Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| `FileName` | String | Specifies the name of the server certificate file to import. |
| `Password` | String | Specifies the password that is used to secure the file specified as *FileName*. |
| `bAllowExport` | Boolean | Indicates whether the imported server certificate is exportable (true) or not exportable (false). |
| `bOverwriteExisting` | Boolean | Indicates whether, if an identical server certificate already exists, it is overwritten (true) or not overwritten (false). |

**ImportToCertStore**

The **ImportToCertStore** method imports a copy of a server certificate from a file to the certificate store on a local or remote computer. Use the following syntax with the ImportToCertStore method:

```
ImportToCertStore FileName Password bAllowExport bOverwriteExisting
```

Table 9.8 lists the parameters that the ImportToCertStore method uses.

**Table 9.8 ImportToCertStore Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| `FileName` | String | Specifies the name of the server certificate file to import. |
| `Password` | String | Specifies the password that is used to secure the file specified as *FileName*. |
| `bAllowExport` | Boolean | Indicates whether the imported server certificate is exportable (true) or not exportable (false). |
| `bOverwriteExisting` | Boolean | Indicates whether, if an identical server certificate already exists, it is overwritten (true) or not overwritten (false). |

**IsExportable**

The **IsExportable** method allows you to determine whether a server certificate is exportable from a computer that is specified by the IIS **ServerName** and **InstanceName** properties.

The IsExportable method uses no parameters.

**IsInstalled**

The **IsInstalled** method allows you to determine whether a server certificate is installed on a computer that is specified by the IIS **ServerName** and **InstanceName** properties.

The IsInstalled method uses no parameters.

**Move**

The **Move** method allows you to move a server certificate from one instance of IIS to another, and from one computer to another. Use the following syntax with the Move method:

```
Move bAllowExport bOverwriteExisting DestinationServerName DestinationServerInstance
DestinationServerUserName DestinationServerPassword
```

Table 9.9 lists the parameters that the Move method uses.

**Table 9.9 Move Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| bAllowExport | Boolean | Indicates whether the imported server certificate is exportable (true) or not exportable (false). |
| bOverWriteExisting | Boolean | Indicates whether, if an identical server certificate already exists, it is overwritten (true) or not overwritten (false). |
| DestinationServerName | String | Specifies the computer name of the server to which the server certificate will be moved. |
| DestinationServerInstance | String | Specifies the IIS metabase instance at which the server certificate will be pointed. For example, w3svc/1. |
| DestinationServerUserName | String | Specifies the user name for the destination server. If no user name is entered, the user name of the user that is currently logged on will be used. |
| DestinationServerPassword | String | Specifies the password for the destination server when the *DestinationServerUserName* parameter is used. |

**RemoveCert**

The **RemoveCert** method allows you to delete a server certificate from a certificate store. Use the following syntax with the RemoveCert method:

```
RemoveCert bRemoveFromCertStore bPrivateKey
```

Table 9.10 lists the parameters that the RemoveCert method uses.

**Table 9.10 RemoveCert Method Parameters**

| Parameter | Data Type | Description |
|---|---|---|
| bRemoveFromCertStore | Boolean | Indicates whether the server certificate is removed (true) or not removed (false). |
| bPrivateKey | Boolean | Indicates whether the private key of the server certificate is removed (true) or not removed (false). |

**Sample Server Certificate Management Scripts**

In IIS, if you use Windows Script Host or ASP, you can execute scripts that communicate with **IISCertObj** programmatically to import, archive, or export server certificates.

**Importing server certificates to multiple servers**

Large sites often need to use one server certificate for secure user logon to multiple servers. It would be time-consuming to use the Microsoft Management Console (MMC) to add copies of the server certificate to each server. By using the **Import** method, you can install copies of the same server certificate on all of the targeted servers.

To import a server certificate from a certificate store on one server to other servers, first save the following script in your text editor as CertImport.vbs.

| **Important** |
|---|
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyMachine******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

```
Option Explicit
```

```
Dim iiscertobj, pfxfile, pfxfilepassword, InstanceName, WebFarmServers, IISServer
Set iiscertobj = WScript.CreateObject("IIS.CertObj")
pfxfile = WScript.Arguments(0)
pfxfilepassword = WScript.Arguments(1)
InstanceName = WScript.Arguments(2)
WebFarmServers = split(WScript.Arguments(3), ",")
iiscertobj.UserName = WScript.Arguments(4)
iiscertobj.UserPassword = WScript.Arguments(5)
For Each IISServer in WebFarmServers
  iiscertobj.ServerName = IISServer
  iiscertobj.InstanceName = InstanceName
  iiscertobj.Import pfxfile, pfxfilepassword, true, true
Next
```

After you save the script as CertImport.vbs, run this command:

**Certimport.vbs cert.pfx pfxPassword w3svc/1 iisserver1,iisserver2,iisserver3 Administrator aal34290**

**Saving backup copies of server certificates in a central archive**

Using the **Export** method, you can archive a backup of each server certificate on your server farm to a central site.

To export copies of server certificates to the central site, first save the following script in your text editor as Save_all_certs.vbs.

```
Option Explicit
Dim iiscertobj, targetServer, targetServers, pfxbasename, pfxpassword, InstanceName
Set iiscertobj = WScript.CreateObject("IIS.CertObj")
pfxbasename = WScript.Arguments(0)
pfxpassword = WScript.Arguments(1)
InstanceName = WScript.Arguments(2)
targetServers = split(WScript.Arguments(3), ",")
iiscertobj.UserName = WScript.Arguments(4)
iiscertobj.UserPassword = WScript.Arguments(5)
iiscertobj.InstanceName = InstanceName
For Each targetServer in targetServers
  iiscertobj.ServerName = targetServer
  iiscertobj.Export pfxbasename + targetServer + ".pfx", pfxpassword, true, false,
false
Next
```

After you save the script as Save_all_certs.vbs, run this command:

**Save_all_certs.vbs C:\certbackup\ pfxpassword w3svc/1 iisserver1,iisserver2,iisserver3 Administrator aal34290**

**Copying a server certificate from an existing server to a new server**

Using the **Copy** method, you can copy a server certificate to a new server after you add the server to a server farm.

To copy a server certificate from an existing server to a new server, first save the following script in your text editor as Certcopy.vbs.

```
Option Explicit
Dim iiscertobj, targetServer, targetServers, targetInstance
Set iiscertobj = WScript.CreateObject("IIS.CertObj")
iiscertobj.ServerName = WScript.Arguments(0)
iiscertobj.Instancename = WScript.Arguments(1)
targetServers = split(WScript.Arguments(2), ",")
targetInstance = WScript.Arguments(3)
iiscertobj.UserName = WScript.Arguments(4)
iiscertobj.UserPassword = WScript.Arguments(5)
For Each targetServer in targetServers
  iiscertobj.Copy true, true, targetServer, targetInstance
Next
```

After you save this script as Certcopy.vbs, run this command:

**Certcopy.vbs iisServer1 w3svc/1 iisServer2 w3svc/1 Administrator asdf-0324**

# Extracting Server Certificate Information with a CAPICOM-Based Script

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

CAPICOM is an easy-to-use interface into CryptoAPI. CAPICOM allows you to incorporate Public Key Infrastructure (PKI) functionality into your applications by providing a simple interface for searching the certificate store and using server certificates to sign or encrypt data. Applications using CAPICOM objects call interfaces into CAPICOM.dll. CAPICOM.dll must be present and registered at run time to use CAPICOM objects. If you use CAPICOM objects in a Visual Basic project, you should add CAPICOM.dll to the project references. CAPICOM.dll is available with the Platform SDK. For more information about CAPICOM, see the CAPICOM Reference on MSDN.

**Registering CAPICOM.dll**
To register CAPICOM.dll, at the command prompt, change directories to the directory where CAPICOM.dll is stored, and then type **regsvr32 capicom.dll**.

You can use CAPICOM to create any number of server certificate–related management scripts. Listing 9.15 uses CAPICOM to extract SSL script information from IIS.

**Listing 9.15 Sample Script for Extracting CAPICOM SSL Script Information**

```
REM
Option Explicit
on error resume next
Const CAPICOM_MY_STORE = "My"
Const CAPICOM_LOCAL_MACHINE_STORE  = 1
Const CAPICOM_CURRENT_USER_STORE  = 2
Const CAPICOM_STORE_OPEN_READ_ONLY = 0
Const CAPICOM_EKU_CLIENT_AUTH = 2
Const CAPICOM_EKU_CODE_SIGNING = 3
Const CAPICOM_EKU_EMAIL_PROTECTION = 4
Const CAPICOM_EKU_SERVER_AUTH = 1
Const CAPICOM_EKU_OTHER = 0
Const CR_DISP_ISSUED  = &H3
Const CR_OUT_CHAIN = &H100
Const CR_OUT_BASE64 = &H1
Const CERT_SYSTEM_STORE_LOCAL_MACHINE = &H20000
Const CR_IN_BASE64  = &H1
Const CR_IN_PKCS10  = &H100
Dim oCert, oStore
Set oStore = CreateObject ("CAPICOM.Store")
if Err.Number <> 0 Then
   wscript.echo "CAPICOM NOT detected"
   Wscript.Quit(1)
End if
oStore.Open CAPICOM_LOCAL_MACHINE_STORE, CAPICOM_MY_STORE, CAPICOM_STORE_OPEN_READ_ONLY
For Each oCert in oStore.Certificates
   WScript.Echo "  Subject Name: " & oCert.SubjectName
        WScript.Echo "  Issuer Name: " & oCert.IssuerName
        WScript.Echo "  SHA-1 Thumbprint: " & oCert.Thumbprint
        WScript.Echo "  Serial Number: " & oCert.SerialNumber
        WScript.Echo "  Version: " & oCert.Version
        WScript.Echo "  Valid From: " & oCert.ValidFromDate
        WScript.Echo "  Valid To: " & oCert.ValidToDate
Next
```

# Managing IIS 6.0 Remotely

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can remotely administer your server running IIS on an intranet or the Internet by using WMI or ADSI scripts, or you can use one of the following tools to manage IIS remotely:

- IIS Manager. You can use IIS Manager on your server to remotely connect to and administer an intranet server running IIS 5.0, IIS 5.1, and IIS 6.0. (IIS 3.0 and IIS 4.0 are not supported.)

- Terminal Services. Terminal Services does not require you to install IIS Manager on the remote client computer because once you are connected to the server running IIS, you use IIS Manager on the Web server as if you are logged on locally.

- Remote Administration (HTML) tool. You can use the Remote Administration (HTML) tool to administer your Web server from any Web browser on your intranet. This version of the Remote Administration (HTML) tool runs only on servers running IIS 6.0.

- For more information about using these tools to manage IIS remotely, see Administering Servers Remotely in IIS 6.0.

# Scheduling IIS 6.0 Backups

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Using available tools, you can create a simple script to back up your IIS content and configuration as often as is required in your environment. After you create this script, you can type one additional command to schedule the backup to occur regularly.

**Backing Up IIS Content**

You can include the **xcopy** command in a simple script for copying IIS and backing up content. **Xcopy** offers useful parameters for copying permissions, auditing settings, and copying hidden files.

| **Important** |
|---|
| You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /***User:MyMachine******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any parameters). |

To back up content with **xcopy**, use the following command-line syntax:

**xcopy** *Source Destination* **[/o] [/x] [/e] [/h] [/y] [c]**

Table 9.11 lists relevant **xcopy** parameters.

**Table 9.11 Xcopy Parameters**

| Parameter | Description |
|---|---|
| *Source* | Specifies the name of the source server. |
| *Destination* | Specifies the name of the destination server. |
| **o** | Copies ACLs. |
| **x** | Copies auditing settings (used with /**o**). |
| **e** | Copies the empty directories. |
| **h** | Copies the hidden directories. |
| **y** | Overwrites without prompting. |
| **c** | Continues to copy even if errors occur. |

**Backing Up IIS Configuration**

To back up the metabase, you can run the Iisback.vbs command-line script with the **/backup** parameter, which places a backup file in *systemroot*\System32\Inetsrv\Metaback. After you run Iisback.vbs, you should copy the backup file to your main backup location. These examples use the backup location \\Mybackupserver\Share$.

If you want to create a batch file, you can set the name with dynamic variables, such as **%date%**, and static, user-defined variables, such as **%server%**.

Next, copy the backup files that are created in the *systemroot*\System32\Inetsrv\Metaback folder. The files that are created are named *Name*.MD# and *Name*.SC#, where *Name* is the name you specify and # is the version number of the backup set, starting with zero (0). If you use the same name for subsequent backups, the version number will be incremented by 1 for each backup set that uses the same name. For more information about backing up the metabase, see Working with the Metabase.

Listing 9.16 is an example of a backup batch file that copies both content and configuration.

**Listing 9.16 Backup.bat: Sample Backup Batch File**

```
set server=myserver
```

```
set name=%date%-%server%
iisback /backup /b %name% /e %password%
xcopy %windir%\system32\inetsrv\metaback\%name%.* \\mybackupserver\share$\%server%
xcopy /o /x /e /h /y /c c:\web \\mybackupserver\share$\%server%
```

This batch file copies the backup files and all of the Web content — one or more Web sites that are subfolders of C:\WEB\ — to the backup location that is a subfolder of the folder that has the same name as the server. The defined variables **%server%** and **%name%** allow you to change the server name quickly without having to go through the batch file.

To schedule this backup process as a nightly task at 9 pm, type **at 9:00p /every:m,t,w,th,f,s,su c:\backup.bat** at the command prompt.

For more information about backing up a computer running Windows Server 2003, see How to Back up Data on the Windows Server System 2003 Web site.

# Restarting and Alternatives to Restarting IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In earlier versions of IIS, if an application became unavailable, restarting IIS by using the IISReset.exe command-line tool was a common recovery practice. However, restarting or stopping IIS, or rebooting your Web server is a severe action. When you restart IIS, all of the sessions that are connected to your Web server, including the World Wide Web Publishing service (WWW service), FTP, Simple Mail Transfer Protocol (SMTP), and Network News Transfer Protocol (NNTP), are dropped. Any data held in Web applications is lost, and all of the Internet sites are unavailable until the Internet services are restarted. To avoid this, IIS 6.0 requires fewer restarts than earlier versions of IIS.

For more information about restarting IIS and alternatives to restarting IIS, see Restarting IIS.

| **Important** |
| --- |
| Although using IISReset.exe is not recommended, if you do use IISReset.exe, always use the **/noforce** parameter. The **/noforce** parameter is recommended as a safeguard against data loss in case the IIS services cannot all be stopped within the one-minute time-out period. |

**Automatic Restart**

In IIS 6.0, the WWW service is hosted by the service host, Svchost.exe. The FTP, NNTP, and SMTP services and the IIS metabase, known as the IIS Admin service (IISAdmin.dll), live in Inetinfo.exe. If the IIS Admin service terminates abnormally, IIS restarts automatically. This feature is known as Automatic Restart.

In IIS 5.0, if the IIS Admin service terminated abnormally, both the WWW service and the IIS Admin service had to be restarted because they shared the same application space. In IIS 6.0, if the IIS Admin service terminates abnormally, the WWW service does not go down because the IIS Admin service and the WWW service run in separate process spaces. In this case, the WWW service acknowledges that the metabase has terminated abnormally and checks to see whether IISReset.exe is configured on the IIS Admin service. If IISReset.exe is configured on the IIS Admin service, IIS waits for the IIS Admin service to start again and then reconnects the WWW service.

**Worker Process Isolation Mode**

Worker process isolation mode is the default service mode in IIS; it is designed to ensure that your applications do not interfere with other services and cause the need to restart IIS. If you are running IIS in worker process isolation mode, you can configure it to periodically restart worker processes in the application pool. You can also recycle application pools manually. In addition, you can configure worker process isolation mode so that if a particular application pool experiences consecutive multiple failures, the application can be automatically disabled, allowing for rapid-fail protection. Finally, in worker process isolation mode you can use CPU monitoring, which is a tool for monitoring and automatically shutting down worker processes that consume large amounts of CPU time.

# Additional Resources for IIS 6.0 Administration

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this section.

**Related Information**

- Configuring Internet Sites and Services for information about configuring the File Transfer Protocol (FTP), Network News Transfer Protocol (NNTP), and Simple Mail Transfer Protocol (SMTP) services.
- IIS 6.0 Architecture for information about the IIS 6.0 application isolation modes.
- Managing a Secure IIS 6.0 Solution for information about setting ACLs.
- Using FrontPage 2002 Server Extensions with IIS 6.0 for information about installing and using FrontPage 2002 Server Extensions.
- Working with the Metabase for information about metabase terminology, including definitions of metabase keys, values, and KeyTypes.
- The CAPICOM site for information about using CAPICOM.
- The Microsoft Windows Management Instrumentation (WMI) SDK for information about the WMI provider, including the WMI Object Browser.

**Related IIS 6.0 Operations Guide Topics**

- Creating Web Sites Using lisweb.vbs for information about creating Web sites by using the lisweb.vbs command-line script.
- Creating FTP Virtual Directories Using lisftpdr.vbs for information about setting up FTP virtual directories.
- Restarting IIS for information about restarting IIS 6.0 and alternatives to restarting IIS 6.0.

# Analyzing Log Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0 offers a number of ways to record the activity of your Web sites, File Transfer Protocol (FTP) sites, Network News Transfer Protocol (NNTP) service, and Simple Mail Transfer Protocol (SMTP) service and allows you to choose the log file format that works best for your environment. IIS logging is designed to be more detailed than the event logging or performance monitoring features of the Microsoft® Windows® Server 2003, Standard Edition, Windows® Server 2003, Enterprise Edition, and Windows® Server 2003, Datacenter Edition, operating systems. IIS log files can include information such as who has visited your site, what was viewed, and when the information was last viewed. You can monitor attempts to access your sites, virtual folders, or files and determine whether attempts were made to read or write to your files. IIS log file formats allow you to record events independently for any site, virtual folder, or file.

**In This Section**


What's New in Logging

Log File Formats in IIS

Centralized Binary Log File Format

HTTP.sys Error Log Files

Custom Logging Modules

Additional Resources for IIS 6.0 Logging


**Related Information**

- For information about the IIS architecture and core components, see IIS 6.0 Architecture.

# What's New in Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Logging in IIS 6.0 differs from logging in previous versions of IIS in several ways. The redesigned IIS architecture, which allows for multiple worker processes, affects how the IIS core components handle logging. In addition, IIS 6.0 provides new logging features and log file formats.

**Events Are Logged by HTTP.sys**
In IIS 5.0, all logging is done by the Inetinfo.exe component and is accomplished using Component Object Model (COM) modules that are written for logging. This system is effective in IIS 5.0 because there is only one server process model, Inetinfo.exe. However, in IIS 6.0, logging is done by the HTTP protocol stack (HTTP.sys). IIS 6.0 passes user-mode events to HTTP.sys through application programming interfaces (APIs), and then the user-mode events are logged by HTTP.sys.

HTTP.sys handles logging for two main reasons. When IIS 6.0 is running in worker process isolation mode, each worker process can perform its own logging and a sites applications can spread across multiple application pools. If Inetinfo.exe performed logging in this environment, multi-instance or synchronization problems could occur. Such problems can be avoided because HTTP.sys performs logging. In addition, in IIS 6.0, requests can be served completely from the HTTP.sys kernel-mode cache without ever passing through Inetinfo.exe, so logging from HTTP.sys ensures that cached responses are recorded.

The log file format for which HTTP.sys does not perform logging is only Open Database Connectivity (ODBC) logging. For more information about how ODBC logging is handled, see ODBC Logging. For more information about the IIS 6.0 architecture and core components, see IIS 6.0 Architecture.

**W3C Centralized Logging**
W3C centralized logging is a global configuration on the server where all Web sites write data to a single log file. This single log file uses the W3C Extended log file format and can be viewed in a text editor, unlike centralized binary logging, which writes data in a binary format and requires a parsing tool to view the data. W3C centralized logging is available in Windows Server 2003 Service Pack 1 or later. For more information about this logging option, see W3C Centralized Logging.

**Substatus Error Code Logging**
To reduce the possibility that a malicious user might use the information provided by substatus error codes to attack the Web server, substatus error codes are not returned to the client in IIS 6.0. Instead, administrators using the World Wide Web Consortium (W3C) Extended log file format can record substatus error codes when requests fail. Substatus error code logging is turned on by default for the W3C Extended log file format. For more information about substatus error code logging, see Substatus Error Codes.

**IIS Centralized Binary Logging**
Centralized binary logging is a process in which IIS creates a single log file that contains binary, unformatted log data for all of the Web sites hosted on a server. In contrast, the other available IIS logging methods create one formatted, human-readable log file per Web site. Centralized binary logging provides organizations with a way to record detailed information about all of their Web sites, using a minimum of system resources. For more information about centralized binary logging, see IIS Centralized Binary Logging.

**HTTP.sys Error Logging**
HTTP.sys sometimes generates its own errors that, because of the IIS 6.0 architecture, are not recorded with the worker-process-driven events. To account for these errors, IIS 6.0 supports HTTP.sys error log files. One example of an event that would trigger a log file entry in the HTTP.sys error log file is a connection time-out. For more information about HTTP.sys error log files, see HTTP.sys Error Log Files.

**Remote Logging**
You can write log data to a remote share over a network using a full, Universal Naming Convention (UNC) path. For example, you can specify *Servername**LogFiles* as the storage directory for your log files, where *Servername*represents the name of the remote server, and *LogFiles* represents the name of the directory where the log files are stored.

For more information about remote logging, see Remote Logging.

# Log File Formats in IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS provides six different log file formats that you can use to track and analyze information about your IIS-based sites and services. In addition to the six available formats, you can create your own custom log file format. For more information about custom logging, see Custom Logging Modules.

The following log file formats and logging options are available in IIS:

- W3C Extended Log File Format Text-based, customizable format for a single site. This is the default format.
- W3C Centralized Logging All data from all Web sites is recorded in a single log file in the W3C log file format.
- NCSA Common Log File Format Text-based, fixed format for a single site.
- IIS Log File Format Text-based, fixed format for a single site.
- ODBC Logging Fixed format for a single site. Data is recorded in an ODBC-compliant database.
- Centralized Binary Logging Binary-based, unformatted data that is not customizable. Data is recorded from multiple Web sites and sent to a single log file. To interpret the data, you need a special parser.
- HTTP.sys Error Log Files Fixed format for HTTP.sys-generated errors.

You can read text-based log files using a text editor such as Notepad, which is included with Windows, but administrators often import the files into a report-generating software tool for further analysis.

For step-by-step instructions on how to enable logging, see Enabling Logging.

**UTF-8 Format for Non-English Languages and Security**
*UTF-8* is a method of character encoding that allows for both single and multibyte characters in one string. If you want to read text-based (W3C Extended, IIS, and NCSA Common) log files in a language other than English, you can enable the UTF-8 format. You also need the UTF-8 format if you serve URLs in a language other than the native server code page, for example, if you serve Chinese URLs on a Japanese system. In addition, because IIS 6.0 tries to serve URLs in code pages other than the server default code page by default, you might want to enable the UTF-8 format for security reasons. It is a good security practice to enable the UTF-8 format in case of an attack based on URLs in UTF-8 that might not translate correctly to the default code page.

| **Note** |
| --- |
| IIS does not support the UTF-8 format for FTP site log files. |

For information about enabling UTF-8, see Enabling UTF-8 Format for Non-English Languages and Security.

**Log File Locations and ACLs**
If you create a log file directory of C:\LogFilesfor HTTP.sys-generated logging (W3C Extended log file format, NCSA Common log file format, IIS log file format, centralized binary logging, or HTTP.sys error logging), then HTTP.sys generates the following subdirectories, and the log files are created under these subdirectories:

- For the W3C Extended, NCSA Common, and IIS log file formats, HTTP.sys generates the subdirectory C:\LogFiles\W3SVC#, where # is the site ID.
- For centralized binary logging, HTTP.sys generates the subdirectory C:\LogFiles\W3SVC.
- For HTTP.sys error logging, HTTP.sys generates the subdirectory C:\WINDOWS\System32\LogFiles\HTTPErr.

By default, the log file directory has the following access control lists (ACLs):

- NT Authority\System: Full access
- Built-in\Administrators: Full access
- Everyone: No access (Although the No access permission is the effective permission, this setting is not explicitly set by HTTP.sys.)

Individual log files in the log file directory have the following controls:

- NT Authority\System: Full access
- Built-in\Administrators: Read and delete access
- Everyone: No access (Although the No access permission is the effective permission, this setting is not explicitly set by HTTP.sys.)

| **Important** |
| --- |
| The default log file directory ACLs (DACLs) are set for optimum security. If you choose to create log file directories with less restrictive ACLs, your system might be more vulnerable to attack. |

If you are using the W3C Extended, NCSA Common, or IIS log file formats and it is necessary to do so, you can create a subdirectory called W3SVC# of the C:\LogFiles directory and set ACLs for that directory that are different from the defaults. For example, if you need to allow a site administrator to read the log files, create a W3SVC# subdirectory of the C:\LogFiles directory and set Read ACLs for the site's administrator.

| **Important** |
| --- |
| The owner of the W3SVC# subdirectory must be a member of the Local Administrators group. If not, HTTP.sys writes an error to the NT Event log and disables logging. |

For more information about file and directory security and access control, see Managing a Secure IIS 6.0 Solution.

# W3C Extended Log File Format

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The W3C Extended log file format is the default log file format for IIS. It is a customizable ASCII text-based format. You can use IIS Manager to select which fields to include in the log file, which enables you to keep log files as small as possible. Because HTTP.sys handles the W3C Extended log file format, this format records HTTP.sys kernel-mode cache hits.

Table 10.1 lists and describes the available fields. Default fields are noted.

**Table 10.1 W3C Extended Log File Fields**

| Field | Appears As | Description | Default Y/N |
|---|---|---|---|
| Date | date | The date on which the activity occurred. | Y |
| Time | time | The time, in coordinated universal time (UTC), at which the activity occurred. | Y |
| Client IP Address | c-ip | The IP address of the client that made the request. | Y |
| User Name | cs-username | The name of the authenticated user who accessed your server. Anonymous users are indicated by a hyphen. | Y |
| Service Name and Instance Number | s-sitename | The Internet service name and instance number that was running on the client. | N |
| Server Name | s-computername | The name of the server on which the log file entry was generated. | N |
| Server IP Address | s-ip | The IP address of the server on which the log file entry was generated. | Y |
| Server Port | s-port | The server port number that is configured for the service. | Y |
| Method | cs-method | The requested action, for example, a GET method. | Y |
| URI Stem | cs-uri-stem | The target of the action, for example, Default.htm. | Y |
| URI Query | cs-uri-query | The query, if any, that the client was trying to perform. A Universal Resource Identifier (URI) query is necessary only for dynamic pages. | Y |
| HTTP Status | sc-status | The HTTP status code. | Y |
| Win32 Status | sc-win32-status | The Windows status code. | N |
| Bytes Sent | sc-bytes | The number of bytes that the server sent. | N |
| Bytes Received | cs-bytes | The number of bytes that the server received. | N |
| Time Taken | time-taken | The length of time that the action took, in milliseconds. | N |
| Protocol Version | cs-version | The protocol version see commentHTTP or FTP see commentthat the client used. | N |
| Host | cs-host | The host header name, if any. | N |

| | | | |
|---|---|---|---|
| User Agent | cs(User-Agent) | The browser type that the client used. | Y |
| Cookie | cs(Cookie) | The content of the cookie sent or received, if any. | N |
| Referrer | cs(Referrer) | The site that the user last visited. This site provided a link to the current site. | N |
| Protocol Substatus | sc-substatus | The substatus error code. | Y |

You can select as many of the W3C Extended log file fields as you want. Some fields may not contain information. When a select field does not contain information, a hyphen (-) appears as a placeholder. If a field contains a nonprintable character, HTTP.sys replaces it with a plus sign (+) to preserve the log file format. Nonprintable characters can appear as a result of virus attacks such as when a malicious user sends carriage returns and line feeds that could break the log file format. You can avoid this problem if you replace nonprintable characters with the plus sign (+).

Fields are separated by spaces. Field prefixes have the following meanings:

- s- Server actions
- c- Client actions
- cs- Client-to-server actions
- sc- Server-to-client actions

**Note**

FTP log files do not record the following fields:

- cs-uri-query
- cs-host
- cs(User-Agent)
- cs(Cookie)
- cs(Referrer)
- sc-substatus

**Note**

For the **Time Taken** field, the client-request time stamp is initialized when HTTP.sys starts parsing the request. If the response size is less than or equal to 2 kilobytes and from memory, or the **force buffer** flag is set, the **Time Taken** value usually does not include network time. Otherwise, the **TimeTaken** value generally includes network time. For more information about the **Time Taken** field in IIS, see Description of the time-taken field in IIS 6.0 and IIS 7.0 HTTP logging.

For information about status codes, see IIS Status Codes.

For more information about the W3C Extended log file format, see Extended Log File Format.

For information about HTTP request buffering that uses the HTTP Server API Version 1, see HttpSendHttpResponse Function.

# W3C Extended Log File Examples

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

**Example of an Intranet Site**

The following example shows a W3C Extended log file entry from an intranet site, as viewed in a text editor. The log file includes only the default fields.

```
#Software: Microsoft Internet Information Services 6.0
#Version: 1.0
#Date: 2002-05-02 17:42:15
#Fields: date time c-ip cs-username s-ip s-port cs-method cs-uri-stem cs-uri-query sc-status cs(User-Agent)
2002-05-02 17:42:15 172.22.255.255 - 172.30.255.255 80 GET /images/picture.jpg - 200 Mozilla/4.0+(compatible;MSIE+5.5;+Windows+2000+Server)
```

You can interpret the preceding log file entry as follows:

#Software: Microsoft Internet Information Services 6.0

This indicates the version of IIS that is running.

#Version: 1.0

This indicates the log file format.

#Date: 2002-05-02 17:42:15

This indicates when the first log file entry was recorded, which is when the entire log file was created.

#Fields: date time c-ip cs-username s-ip s-port cs-method cs-uri-stem ...

This indicates the fields that are recorded in the log file entry and the order in which they are recorded

Table 10.2 lists and describes the fields that are recorded in this example.

| **Note** |
|---|
| Use the order shown in the Fields log file entry to parse all entries until the next log file entry begins. The field order is subject to change and is different in IIS 6.0 than it was in IIS 5.0. |

**Table 10.2 Example of a W3C Extended Log File Entry from an Intranet Site**

| Field | Appears As | Description |
|---|---|---|
| date | 2002-05-02 | This log file entry was recorded on May 2, 2002. |
| time | 17:42:15 | This log file entry was recorded at 5:42 P.M. UTC. Entries are recorded to the log file when the send completion for the last IIS send occurs. |
| c-ip | 172.22.255.255 | The IP address of the client. |
| cs-username | - | The user was anonymous. |
| s-ip | 172.30.255.255 | The IP address of the server. |
| s-port | 80 | The server port. |
| cs-method | GET | The user issued a **GET**, or download, command. |

| | | |
|---|---|---|
| cs-uri-stem | /images/picture.jpg | The user wanted to download the picture.jpg file from the Images folder. |
| cs-uri-query | - | The URI query did not occur. (URI queries are necessary only for dynamic pages, such as ASP pages, so this field usually contains a hyphen for static pages.) |
| sc-status | 200 | The request was fulfilled with no errors. |
| cs(User-Agent) | Mozilla/4.0+ (compatible;MSIE+5.5;+Windows+2000+Server) | The type of browser that the client used, as represented by the browser. |

**Example of an Internet Site**

The following example shows a W3C Extended log file entry from an Internet site, as viewed in a text editor. The log file contains several fields in addition to the default fields and is typical of the kind of log file you would run to monitor an Internet site.

```
#Software: Microsoft Internet Information Services 6.0
#Version: 1.0
#Date: 2002-05-24 20:18:01
#Fields: date time c-ip cs-username s-ip s-port cs-method cs-uri-stem cs-uri-query sc-
status sc-bytes cs-bytes time-taken cs(User-Agent) cs(Referrer)
2002-05-24 20:18:01 172.224.24.114 - 206.73.118.24 80 GET /Default.htm - 200 7930 248
31 Mozilla/4.0+(compatible;+MSIE+5.01;+Windows+2000+Server) http://64.224.24.114/
```

You can interpret the preceding log file entry as follows:

#Software: Microsoft Internet Information Services 6.0

This indicates the version of IIS that is running.

#Version: 1.0

This indicates the log file format.

#Date: 2002-05-24 20:18:01

This indicates when the first log file entry was recorded, which is when the entire log file was created.

#Fields: date time c-ip cs-username s-ip s-port cs-method cs-uri-stem ...

This indicates the fields that are recorded in the log file entry and the order in which they are recorded. Table 10.3 lists and describes the fields recorded in this example.

**Table 10.3 Example of a W3C Extended Log File Entry from an Internet Site**

| Field | Appears As | Description |
|---|---|---|
| date | 2002-05-24 | This log file entry was recorded on May 24, 2002. |
| time | 20:18:01 | This log file entry was recorded at 8:18 P.M. UTC. |
| c-ip | 172.224.24.114 | The IP address of the client. |
| cs-username | - | The user was anonymous. |
| s-ip | 206.73.118.24 | The IP address of the server. |
| s-port | 80 | The server port. |
| cs-method | GET | The user issued a **GET**, or download, command. |
| cs-uri-stem | /Default.htm | The user wanted to download the contents of Default.htm. |
| cs-uri-query | - | The URI query did not occur. |

| | | |
|---|---|---|
| sc-status | 200 | The request was fulfilled without error. |
| sc-bytes | 7930 | The number of bytes that the server sent to the client. |
| cs-bytes | 248 | The number of bytes that the client sent to the server. |
| time-taken | 31 | The action was completed in 31 milliseconds. |
| cs(User-Agent) | Mozilla/4.0+ (compatible;+MSIE+5.01;+Windows+2000+Server) | The type of browser that the client used, as represented by the browser. |
| cs(Referrer) | http://62.224.24.114/ | The Web page that provided the link to the Web site. |

**Example of an FTP Site**

The following example shows a W3C Extended log file entry from an Internet FTP site, as viewed in a text editor. The example does not include all of the default fields and is typical of the kind of log file that you would run to monitor an FTP site.

```
#Software: Microsoft Internet Information Services 6.0
#Version: 1.0
#Date: 2002-06-04 16:40:23
#Fields: time c-ip cs-method cs-uri-stem sc-status
16:40:23 10.152.10.200 [6994]USER anonymous 331
16:40:25 10.152.10.200 [6994]PASS anonymous@example.net 530
```

You can interpret the preceding log file example as follows:

#Software: Microsoft Internet Information Services 6.0

This indicates the version of IIS that is running.

#Version: 1.0

This indicates the log file format.

#Date: 2002-06-04 16:40:23

This indicates when the first log file entry was recorded, which is when the log was created.

#Fields: time c-ip cs-method cs-uri-stem sc-status

This indicates the fields that are recorded in the log file entry and the order in which they are recorded. Table 10.4 lists and describes the fields recorded in this example.

**Table 10.4 Example of a W3C Extended Log File Entry from an FTP Site**

| Field | Appears As | Description |
|---|---|---|
| time | 16:40:23 | This log file entry was recorded at 4:40 P.M. UTC. |
| c-ip | 10.152.10.200 | The IP address of the client. |
| cs-method | [6994]USER | The **USER FTP** command was used, which requests a user name and is always followed by a **PASS FTP** command. 6994 is the connection number corresponding to an anonymous user. |
| cs-uri-stem | anonymous | The user (the target of the **USER** command) was anonymous. |
| sc-status | 331 | The user name was accepted. |
| time | 16:40:25 | The next recorded action occurred at 4:40 P.M. UTC. |
| c-ip | 10.152.10.200 | The IP address of the client. |
| cs-method | [6994]PASS | The **PASS FTP** command was used, which supplies a password for the user name and is always preceded by a **USER** command. |

| | | |
|---|---|---|
| cs-uri-stem | anonymous@example.net | The password (the target of the **PASS** command) supplied. |
| sc-status | 530 | The user was not logged on. |

# Substatus Error Codes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In an effort to reduce the attack surface of IIS, in IIS 6.0, error messages do not return specific error message content, including the substatus error code, to clients. However, to allow administrators to track errors and debug failed requests, IIS provides the ability to record substatus error codes in the W3C Extended log file format. Substatus error code logging is enabled by default. For more information about enabling and disabling substatus error code logging, see Substatus Error Codes in Log Files.

If an error message contains too much information about the server and an explanation of why a particular request failed to execute, malicious users can use the information from the error message to attack the server. For example, an error code such as 404.2 indicates that a file or directory is not returned because the server lockdown policy restricts it. In IIS 6.0, a 404.2 error is returned to the client as a 404 error message. The simple 404 error message does not provide a malicious user with any details about the configuration of the server.

When substatus error code logging is enabled by default, only members of the Administrators group and users with LocalSystem user accounts can access log files that contain substatus error codes. To analyze an error, the administrator locates the error and substatus error code in the log file and checks the code against the IIS 6.0 HTTP.sys error code list. For example, if a client request to an Active Server Pages (ASP) page returned a 403 error, the administrator can determine that the actual error was 403.9 by viewing the log file. The administrator checks the IIS 6.0 HTTP.sys error code list and learns that, in this situation, the file or directory was not found because too many clients were trying to connect to the server at once. The administrator can remedy the situation quickly and easily by changing the maximum number of connections setting to unlimited. To see the IIS 6.0 HTTP.sys error code list, see HTTP Status Codes in IIS 6.0.

# W3C Centralized Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

W3C centralized logging is a global configuration on the server where all Web sites write data to a single log file. Data is stored in the log file using the W3C Extended log file format. The log file can be viewed in a text editor, unlike IIS Centralized Binary Logging which writes data in binary format and requires a parsing tool to view the data. W3C centralized logging is available in Windows Server 2003 Service Pack 1 or later.

| **Important** |
| --- |
| FTP, NNTP, and SMTP do not support W3C centralized logging. |

By default, IIS Web sites write data to individual log files (one log file per Web site). On servers hosting large numbers of sites (hundreds or thousands), the process of maintaining large numbers of open file handles to log files can negatively impact how the server scales. W3C centralized logging improves server scalability on servers hosting large numbers of sites because IIS requires only one open file handle for logging.

W3C centralized logging is a server property, not a site property, so when you enable this feature all Web sites on that server are configured to write log data to the central log file.

| **Note** |
| --- |
| W3C centralized logging uses the W3C Extended log format, which includes the following four fields: **HostHeader**, **Cookie**, **UserAgent**, and **Referrer**. These fields are not available with centralized binary logging. |

For the procedure to enable W3C centralized logging, see Enabling W3C Centralized Logging.

**See Also**
**Concepts**
CentralW3CLoggingEnabled Metabase Property

# NCSA Common Log File Format

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The NCSA Common log file format is a fixed ASCII text-based format, so you cannot customize it. The NCSA Common log file format is available for Web sites and for SMTP and NNTP services, but it is not available for FTP sites. Because HTTP.sys handles the NCSA Common log file format, this format records HTTP.sys kernel-mode cache hits.

The NCSA Common log file format records the following data:

- Remote host address
- Remote log name (This value is always a hyphen.)
- User name
- Date, time, and Greenwich mean time (GMT) offset
- Request and protocol version
- Service status code (A value of 200 indicates that the request was fulfilled successfully.)
- Bytes sent

Not all fields will contain information. For fields for which there is no information, a hyphen (-) appears as a placeholder. If a field contains a nonprintable character, HTTP.sys replaces it with a plus sign (+) to preserve the log file format. This typically occurs with virus attacks, when, for example, a malicious user sends carriage returns and line feeds that, if not replaced with the plus sign (+), would break the log file format.

Fields are separated by spaces, and time is recorded as local time with the GMT offset.

# NCSA Common Log File Example

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following example shows an NCSA Common log file entry, as viewed in a text editor.

```
172.21.13.45 - Microsoft\JohnDoe [08/Apr/2001:17:39:04 -0800] "GET
/scripts/iisadmin/ism.dll?http/serv HTTP/1.0" 200 3401
```

Table 10.5 lists and describes the fields recorded in this example.

**Table 10.5 Example of an NCSA Common Log File Entry**

| Field | Appears As | Description |
|---|---|---|
| Remote host address | 172.21.13.45 | The IP address of the client. |
| Remote log name | - | The name is unavailable. |
| User name | Microsoft\JohnDoe | The name of the user is JohnDoe. |
| Date, time, and GMT offset | [08/Apr/2001:17:39:04 -0800] | The log file entry was created on April 8, 2001 at 5:39 P.M. The difference between the local time and the GMT is eight hours. |
| Request and protocol version | GET /scripts/iisadmin/ism.dll?http/serv HTTP/1.0 | The client issued a **GET** command for the Ism.dll file using HTTP version 1.0. |
| Service status code | 200 | The request was fulfilled successfully. |
| Bytes sent | 3401 | The number of bytes sent. |

# IIS Log File Format

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS log file format is a fixed ASCII text-based format, so you cannot customize it. Because HTTP.sys handles the IIS log file format, this format records HTTP.sys kernel-mode cache hits.

The IIS log file format records the following data:

- Client IP address
- User name
- Date
- Time
- Service and instance
- Server name
- Server IP address
- Time taken
- Client bytes sent
- Server bytes sent
- Service status code (A value of 200 indicates that the request was fulfilled successfully.)
- Windows status code (A value of 0 indicates that the request was fulfilled successfully.)
- Request type
- Target of operation
- Parameters (the parameters that are passed to a script)

Not all fields will contain information. For fields for which there is no information, a hyphen (-) appears as a placeholder. If a field contains a nonprintable character, HTTP.sys replaces it with a plus sign (+) to preserve the log file format. This typically occurs with virus attacks, when, for example, a malicious user sends carriage returns and line feeds that, if not replaced with the plus sign (+), would break the log file format.

Fields are separated by commas, making the format easier to read than the other ASCII formats, which use spaces for separators. The time is recorded as local time. A Web site instance is displayed as W3SVC#, and an FTP site instance is displayed as MSFTPSVC#, where # is the instance of the site. Time taken is recorded in milliseconds. For more information about the time taken field, see W3C Extended Log File Format.

# IIS Log File Example

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following example shows an IIS log file entry, as viewed in a text editor:

```
192.168.114.201, -, 03/20/01, 7:55:20, W3SVC2, SERVER, 172.21.13.45, 4502, 163, 3223,
200, 0, GET, /DeptLogo.gif, -,
```

Table 10.6 lists and describes the fields recorded in this example.

**Table 10.6 Example of an IIS Log File Entry**

| Field | Appears As | Description |
|---|---|---|
| Client IP address | 192.168.114.201 | The IP address of the client. |
| User name | - | The user is anonymous. |
| Date | 03/20/01 | This log file entry was made on March 20, 2001. |
| Time | 7:55:20 | This log file entry was recorded at 7:55 A.M. |
| Service and instance | W3SVC2 | This is a Web site, and the site instance is 2. |
| Server name | SERVER | The name of the server. |
| Server IP | 172.21.13.45 | The IP address of the server. |
| Time taken | 4502 | This action took 4,502 milliseconds. |
| Client bytes sent | 163 | The number of bytes sent from the client to the server. |
| Server bytes sent | 3223 | The number of bytes sent from the server to the client. |
| Service status code | 200 | The request was fulfilled successfully. |
| Windows status code | 0 | The request was fulfilled successfully. |
| Request type | GET | The user issued a **GET**, or download, command. |
| Target of operation | /DeptLogo.gif | The user wanted to download the DeptLogo.gif file. |
| Parameters | - | There were no parameters passed. |

# ODBC Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ODBC logging records a fixed set of data fields in an ODBC-compliant database, such as Microsoft® Access or Microsoft® SQL Server™. With ODBC logging, you must set up a database to receive the data, and then specify this database as the database to which log files should be recorded.

When ODBC logging is enabled, IIS disables the HTTP.sys kernel-mode cache. For this reason, implementing ODBC logging will degrade overall server performance. ODBC logging is not recommended unless your database limits you to this option.

IIS includes a Structured Query Language (SQL) template file that you can run in a SQL database to create a table that accepts log file entries from IIS. The file is named Logtemp.sql, and it is located in the *systemroot*\System32\Inetsrv directory by default.

**Important**

For security reasons, do not use the well-known SQL Server administrator account name SA for ODBC logging. If a malicious user accessed the worker process, that user could use the SA account to access the SQL Server. Remove the SA account, and create a new account with the least necessary rights. For more information, see SQL Server 2000 Operations Guide: Security Administration.

Table 10.7 lists the fields, data types, and descriptions of the fields recorded in the ODBC logging format.

**Table 10.7 ODBC Logging Fields**

| Field | Data Type | Description |
|---|---|---|
| ClientHost | varchar(255) | The IP address of the client. |
| UserName | varchar(255) | The domain name of the client. |
| LogTime | datetime | The connection date and time. |
| Service | varchar(255) | The IIS service. |
| Machine | varchar(255) | The name of the computer. |
| ServerIP | varchar(50) | The IP address of the server. |
| ProcessingTime | integer | The processing time, in milliseconds. |
| BytesRecvd | integer | The number of bytes that the server received. |
| BytesSent | integer | The number of bytes that the server sent. |
| ServiceStatus | integer | The service status or reply code. A value of 200 indicates that the request was fulfilled successfully. |
| Win32Status | integer | The Windows status or error code. A value of 0 indicates that the request was fulfilled successfully. |
| Operation | varchar(255) | The command that the user requested. |
| Target | varchar(255) | The recipient. |
| Parameters | varchar(255) | The request parameters that the client sent. |

# IIS Centralized Binary Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Centralized binary logging is a process where multiple Web sites send binary, unformatted log data to a single log file. In contrast, the other IIS logging methods create one log file per Web site.

| **Important** |
| --- |
| FTP, NNTP, and SMTP do not support centralized binary logging. |

When a server running IIS hosts many Web sites, the process of creating hundreds or thousands of formatted log files and writing the log data to a disk can quickly consume valuable CPU and memory resources, thereby creating performance and scalability problems. Centralized binary logging minimizes the amount of system resources that are used for logging, while at the same time providing detailed log data for organizations that require it.

Centralized binary logging is a server property, not a site property, so when you enable centralized binary logging on a server, all of the Web sites on that server are configured to write log data to the central log file. When you enable centralized binary logging, you cannot record data from individual Web sites in a different format. The centralized binary logging log file has an Internet binary log (.ibl) file name extension. This file name extension ensures that text utilities do not try to open and read the central binary logging log file.

When to Use Centralized Binary Logging

Centralized binary logging is particularly useful when many Web sites are hosted on the same server or where server resources are at a premium. With centralized binary logging, an administrator can maximize the number of Web sites that a server can host and record activity for because centralized binary logging reduces the amount of system resources that are consumed during logging and decreases code paths in IIS, thus increasing performance and scalability.

Centralized binary logging can also reduce administration burdens for Internet service providers (ISPs) for whom collecting and storing logged data is imperative. For example, if an ISP has six servers with 10,000 Web sites per server, the ISP would have to manage 10,000 log files per day per server running IIS. With centralized binary logging, the ISP would have to manage and store only one file per server, per day.

Log files can be vital to troubleshooting applications. However, an organization might not want to take up system resources for log files that are used only in debugging situations. For example, a Web application team at an enterprise organization uses ASP and Microsoft® Visual Basic® development system to develop COM components that access a SQL database. The application team needs to maximize the response time of individual requests and maximize the amount of concurrent work that a group of servers can handle. Central binary logging can help the team achieve these goals.

For information about enabling centralized binary logging for all Web sites hosted on a server running IIS 6.0, see Centralized Binary Logging in IIS 6.0.

When you are ready to extract data from a raw log file, you can do one of the following:

- Create a tool that locates and extracts the data that you want from the raw file and converts the data into formatted text. You can view a header file and log file format descriptions in the IIS 6.0 Software Development Kit on MSDN.

- Use the Log Parser tool to extract data from the raw file. The Log Parser tool and its accompanying user documentation are included in the IIS 6.0 Resource Kit Tools.

Centralized Binary Logging File Format

The raw centralized binary logging log file is made up of fixed-length records or index records that contain string identifiers. The index records appear because, in an effort to record as much information as possible, variable-length string fields are replaced by numeric identifiers — indexes — that map the variable-length string to the logged identifier.

The raw log file is not human-readable, and it cannot be read using most available log analyzers. To extract data from a raw log file, you can create a tool that locates and extracts the data and then converts it into formatted text. For more information about the header file and log file format descriptions that you need to create such a tool, see Centralized Binary Log File Format.

Centralized binary logging records the following information, which is similar, but not identical, to the W3C Extended log file

format:

- Date
- Time
- Client IP address
- User name
- Site ID
- Server name
- Server IP address
- Server port
- Method
- URI stem
- URI query
- Protocol status
- Windows status
- Bytes sent
- Bytes received
- Time taken
- Protocol version
- Protocol substatus

For information about the **Time Taken** field, see W3C Extended Log File Format.

| **Note** |
| --- |
| The following fields are reported in W3C Extended log files, but they are not recorded in centralized binary logging log files: |

- Host. The host header.
- User agent. The browser type of the client; this string is too large to be practical for the binary format.
- Cookie. The content of the cookie that was sent.
- Referrer. The site that the user last visited.

# Centralized Binary Log File Format

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Centralized binary log files require a specialized parser to interpret the log data. You can use the file format provided in this appendix to create a custom parser.

See Also

**Concepts**

Analyzing Log Files

# Centralized Binary Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Centralized binary logging was designed to log as much information as possible while using the least amount of system resources. Therefore, numeric identifiers, which are called *indexes*, are substituted for variable-length strings in key fields. The indexes map the variable-length string to the logged identifier. When centralized binary logging is enabled, every site hosted on a server running Internet Information Services (IIS) sends raw log data to a single file.

# IIS Binary Log File Names

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The binary log file has an Internet binary log, or .ibl, file name extension. This file name extension was created so that text utilities that recognize the .log format would not attempt to open the binary log file.

Depending on the recycling option you choose, the binary log file will be named as follows:

- **Date-based recycling**: rayymmddhh.ibl. The time in the name is expressed in Universal Coordinated Time (UTC).
- **Size recycling**: raw#.ibl. IIS assigns a number to each log file, in increments of one. The # represents the number of the log file.

# Record Order

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The typical central binary log file has the records shown in Figure F.1, and they appear in the order shown. However, the exact order depends on the arrival order of cache-hit requests and cache-miss requests.



# Record Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

All binary log record files have a two-byte **RecordType** field defined as the first element. Parsers should read the **USHORT** field first to determine the record type before parsing the whole record. Table F.1 lists the numbers that represent the binary log file record types.

Table F.1 Record Types

| Record Type | Number |
|---|---|
| #define HTTP_RAW_RECORD_HEADER_TYPE | (0) |
| #define HTTP_RAW_RECORD_INDEX_DATA_TYPE | (2) |
| #define HTTP_RAW_RECORD_HIT_LOG_DATA_TYPE | (3) |
| #define HTTP_RAW_RECORD_MISS_LOG_DATA_TYPE | (4) |
| #define HTTP_RAW_RECORD_CACHE_NOTIFICATION_DATA_TYPE | (5) |

Basic data types, such as **ULONG**, shown in the records in the rest of this appendix are written to the binary log file in host order, unless otherwise specified.

| **Note** |
|---|
| You must use the Windows.H header file to make most of the following data types available. For more information about the Windows.H header file, search for Windows.H on MSDN. |

Header Record

The header record is the first log file record. It contains fields that identify which server generated the log file, and the date and time that the file was created. This information enables you to validate the file, see the statement of capabilities, and move the logs from one computer to another for post-processing, all without losing the source of the information.

A header record is always included at the beginning of each log file, and if the server is restarted during the logging period, header records are also written into the middle of each log file. For example, if the server is configured to start a new log file every day and the server is then restarted in the middle of the day, a second header is placed in the middle of that days log file.

Example F.1 contains a description of the header record with inline comments.

Example F.1 Header Record

| **Note** |
|---|
| Please be aware that Microsoft might change this data format at any time and is likely to do so in future Windows releases. Developers who use this data structure definition will need to alter program logic, recompile, and then re-release their applications for future releases of Windows (or service packs). |

```
#define MAX_COMPUTER_NAME_LEN           (256)

typedef struct _RAW_FILE_HEADER
{
    // Must be RAW_RECORD_HEADER_TYPE, i.e. value is 0
    //
    USHORT          RecordType;

    // Identifies the version of the Internet Binary Log File.
    //
    union {
```

```
        struct {
            UCHAR MajorVersion; // MAJOR_RAW_LOG_FILE_VERSION == 1
            UCHAR MinorVersion; // MINOR_RAW_LOG_FILE_VERSION == 0
        };
        USHORT Version;
    };

    // Shows the alignment padding size. sizeof(PVOID). Important for    // parsers to
be able to determine the padding size (depending on    // the processor architecture of
the computer that generated the    // log field). ie. X86 == 4 bytes ia64 == 8 bytes.
    //
    ULONG           AlignmentSize;

// UTC based system timestamp for the raw file creation/opening. (Can// be used  to
populate a FILETIME structure.)
//
    LARGE_INTEGER   DateTime;

    // Name of the server which created/opened the raw file. Maximum    // computer
name length is set to 256, and it is null-terminated if    // the length is not exactly
256 characters. This is important to    // identify the source computer for the log. In
either case, the    // parser should skip the remaining bytes to reach the next record.
    //
    WCHAR           ComputerName[MAX_COMPUTER_NAME_LEN];

} RAW_FILE_HEADER, *PRAW_FILE_HEADER;
```

File Footers

The binary log file format does not write a footer. You can detect the next file based on the timestamps in the header and the log file records written in the previous log file.

Inline Records

Each request record contains the following inline records:

- The RAWLOGID, which is a unique identifier for URLs.
- An IP address type flag, which tells you whether the IP address is IPv4 or IPv6.

**Unique identifier for URLs**

Centralized binary logging contains a 64-bit URL identifier. To conserve space, frequent URLs are written once and then given a random identifier called the RAWLOGID. Example F.2 lists a RAWLOGID.

Example F.2 URL Identifier

| **Note** |
| --- |
| Please be aware that Microsoft might change this data format at any time and is likely to do so in future Windows releases. Developers who use this data structure definition will need to alter program logic, recompile, and then re-release their applications for future releases of Windows (or service packs). |

```
typedef struct _RAWLOGID
{
    //
    // The identifier is a 64-bit field and is a random number.
    //

    ULONG   UIDLowPart;
    ULONG   UIDHighPart;

} RAWLOGID, *PRAWLOGID;
```

**IP address type**

Each request record has either an IPv4 or IPv6 address encoded within it. Additionally, a flag in the first part of the log entry

indicates the IP version of the address field. Example F.3 lists an IPv4 address, and Example F.4 lists an IPv6 address.

Example F.3 IPv4 Address

| **Note** |
| --- |
| Please be aware that Microsoft might change this data format at any time and is likely to do so in future Windows releases. Developers who use this data structure definition will need to alter program logic, recompile, and then re-release their applications for future releases of Windows (or service packs). |

```
typedef struct _RAWLOG_IPV4_ADDRESSES {

    // The IPV4 client address or source of the request
    ULONG client;

    // The IPV4 server address or address of the server that received    // the request
ULONG server;

} RAWLOG_IPV4_ADDRESSES, *PRAWLOG_IPV4_ADDRESSES;
```

Example F.4 IPv6 Address

| **Note** |
| --- |
| Please be aware that Microsoft might change this data format at any time and is likely to do so in future Windows releases. Developers who use this data structure definition will need to alter program logic, recompile, and then re-release their applications for future releases of Windows (or service packs). |

```
typedef struct _RAWLOG_IPV6_ADDRESSES {

    // The IPV6 client address or source of the request
   USHORT client[8];

    // The IPV6 server address or address of the server that received    // the request
   USHORT server[8];


} RAWLOG_IPV6_ADDRESSES, *PRAWLOG_IPV6_ADDRESSES;
```

Cache or User-Mode Request

For every HTTP request sent to IIS 6.0, the HTTP protocol stack (HTTP.sys) checks the cache entry for a response that matches the incoming request. If the request matches, then HTTP.sys fulfills the request from the cache. If the request does not match, then requests are sent to the user-mode IIS worker process for further processing. The binary log file contains one record type for cache hits and a different record type for cache misses, which are requests sent to the user-mode IIS worker process.

Centralized binary logging records data for all of the fields that are recorded by default in W3C Extended logging format, except for the User-Agent field. This field is too large to record because the string comes from the client browser, and if an index is used, then it would imply the comparison and hashing of this field for every record that is logged.

Cache-Hit Record

Example F.5 applies to requests that are served directly from the HTTP.sys cache. There is one record per request.

Example F.5 Cache-Hit Record

| **Note** |
| --- |
| Please be aware that Microsoft might change this data format at any time and is likely to do so in future Windows releases. Developers who use this data structure definition will need to alter program logic, recompile, and then re-release their applications for future releases of Windows (or service packs). |

```
typedef struct _RAW_FILE_HIT_LOG_DATA
{
    // Type must be RAW_RECORD_HIT_LOG_DATA_TYPE i.e. value is 3.
```

```
    //
    USHORT          RecordType;

    // Optional flags.
    //
    union
    {
        struct
        {
            USHORT IPv6:1;              // Are IPs for this request IPv6 or not?
            USHORT ProtocolVersion:3;   // See protocol version table below.
            USHORT Method:6;            // See HTTP Verb Table below.
            USHORT Reserved:6;
        };
        USHORT Value;
    } Flags;

    // Site ID. Represents which site owns this log record as per the
    // metabase numeric metabase SiteID.
    //
    ULONG           SiteID;

    // UTC Timestamp for the Log Hit.
    //
    LARGE_INTEGER   DateTime;

    // The IP Port the request came in on
    //
    USHORT          ServerPort;

    // ProtocolStatus - 200, 404, 503 etc - value should never be
    // greater than 999
    //
    USHORT          ProtocolStatus;

    // Other fields of the response
    //
    ULONG           Win32Status;
    ULONGLONG       TimeTaken;      // in milliseconds
    ULONGLONG       BytesSent;   // in bytes
    ULONGLONG       BytesReceived; // in bytes

    // For cache hits, there will always be at least one Uri Index    // record written
prior to this record. (see _RAW_INDEX_FIELD_DATA    // record)
    //
    RAWLOGID   UriId;

    // Either IPv4 or IPv6 address structure will follow.
    //
    // _RAWLOG_IPV4_ADDRESSES or
    // _RAWLOG_IPV6_ADDRESSES

} RAW_FILE_HIT_LOG_DATA, *PRAW_FILE_HIT_LOG_DATA;
```

Table F.2 lists the values that correspond to each protocol version.

Table F.2 Protocol Versions

| Value | Protocol Version |
|---|---|
| 0 | Unknown |
| 1 | HTTP 0.9 |
| 2 | HTTP 1.0 |
| 3 | HTTP 1.1 |

Table F.3 lists the HTTP methods used for the binary log file format. The mapping of the method field is essentially the same as the HTTP_VERB type definition in the Http.h HTTP API header file in the Microsoft Platform SDK. For more information about the Http.h header file, see HTTP API.

**Table F.3 HTTP Methods**

| Value | HTTP Method |
|---|---|
| 0 | Unparsed |
| 1 | Unknown (for new verbs) |
| 2 | Invalid |
| 3 | OPTIONS |
| 4 | GET |
| 5 | HEAD |
| 6 | POST |
| 7 | PUT |
| 8 | DELETE |
| 9 | TRACE |
| 10 | CONNECT |
| 11 | TRACK |
| 12 | MOVE |
| 13 | COPY |
| 14 | PROPFIND |
| 15 | PROPPATCH |
| 16 | MKCOL |
| 17 | LOCK |
| 18 | UNLOCK |
| 19 | SEARCH |

Cache-Miss Record

Example F.6 applies to requests that are served from the user-mode worker process. There is one record per request.

**Example F.6 Cache-Miss Record**

| **Note** |
|---|
| Please be aware that Microsoft might change this data format at any time and is likely to do so in future Windows releases. Developers who use this data structure definition will need to alter program logic, recompile, and then re-release their applications for future releases of Windows (or service packs). |

```
typedef struct _RAW_FILE_MISS_LOG_DATA
{
    // Type must be RAW_RECORD_MISS_LOG_DATA_TYPE i.e. value is 4
    //
    USHORT          RecordType;

    // Optional flags.
    //
    union
    {
        struct
        {
            USHORT IPv6:1;              // Are IPs for this request IPv6 or not?
            USHORT ProtocolVersion:3;   // See protocol version table above.
            USHORT Method:6;            // See HTTP Verb Table above.
            USHORT Reserved:6;
        };
        USHORT Value;
    } Flags;

    // Site ID. Represents which site owns this log record as per the
    // metabase numeric metabase SiteID.
    //
    ULONG           SiteID;

    // UTC Timestamp for the Log Hit.
    //
    LARGE_INTEGER   DateTime;

    // The IP Port the request came in on.
    //
    USHORT          ServerPort;

    // ProtocolStatus - 200, 404, 503 etc - value should never be
    // greater than 999
    //
    USHORT          ProtocolStatus;

    // Other fields of the response
    //
    ULONG           Win32Status;
    ULONGLONG       TimeTaken;         // in milliseconds
    ULONGLONG       BytesSent;      // in bytes
    ULONGLONG       BytesReceived;      // in bytes

    // Substatus is only relevant for requests served from user mode.
    // All cache-hit responses will have an implied Substatus of 0.
    //
    USHORT          SubStatus;

    // Variable-length fields follow the structure.
    //
    USHORT          UriStemSize;
    USHORT          UriQuerySize;
    USHORT          UserNameSize;

    // Either IPv4 or IPv6 address structure will follow. The client    // address
appears first, followed by the server address.
    //
    // _RAWLOG_IPV4_ADDRESSES or
    // _RAWLOG_IPV6_ADDRESSES

    // Then the URIStem - UriStemSize bytes
    // Then the URIQuery - UriQuerySize bytes
    // Then the UserName - Align the UserNameSize bytes up to the size of
    // AlignmentSize (in the header record) multiple

} RAW_FILE_MISS_LOG_DATA, *PRAW_FILE_MISS_LOG_DATA;
```

After writing the RAW_FILE_MISS_LOG_DATA record to the file that is already aligned, HTTP.sys appends the variable size strings to the end of the structures. However, if the total length of the bytes written after the structure is not aligned, HTTP.sys

provides padding to ensure proper alignment of the next record.

Parsers should be aware of this padding and must skip it after reading the last variable size string that follows the RAW_FILE_MISS_LOG_DATA record.

For example, if a cache-miss record has a URIStem of 8 bytes, URIQuery of 0 bytes, and UserName of 4 bytes, then HTTP.sys adds the following padding:

ALIGN_UP ((8 + 0 + 4), AlignmentSize) - (8 + 0 + 4). This would be 4 bytes for a binary log file generated on an IA64, and 0 bytes for a binary log file generated for I386.

After reading the last variable size string, the parser should skip the padding bytes.

Log File Index Mapping Record

A log file index mapping record maps a specific URL to a RAWLOGID. This record always appears before the RAWLOGID is referenced in a cache-hit record. If a subsequent record for the same URL appears in a log file, then it receives a new RAWLOGID and subsequent cache-hit records refer to the new ID. This occurs because HTTP.sys periodically flushes cached items and recaches after a period of time, and a new RAWLOGID is generated for the recached item. Example F.7 lists the log file index mapping record.

**Example F.7 Log File Index Mapping Record**

| **Note** |
| --- |
| Please be aware that Microsoft might change this data format at any time and is likely to do so in future Windows releases. Developers who use this data structure definition will need to alter program logic, recompile, and then re-release their applications for future releases of Windows (or service packs). |

```
#define URI_BYTES_INLINED       (4)
#define URI_WCHARS_INLINED      (URI_BYTES_INLINED/sizeof(WCHAR))

typedef struct _RAW_INDEX_FIELD_DATA
{
    // RAW_RECORD_INDEX_DATA_TYPE == i.e. 2
    //
    USHORT          RecordType;

    // Size of the variable size string (in bytes).
    // When reading and writing need to align the DIFF(Size - 4)
    // up to AlignmentSize.
    //
    USHORT          Size;

    // Unique ID for the URI.
    //
    RAWLOGID   Id;

    // Variable size string follows immediately after the structure.
    // Array of 4 bytes is defined to be able to make it PVOID aligned
    // on ia64. Typically URIs will be bigger than 4 bytes.
    //
    WCHAR           Str[URI_WCHARS_INLINED];

} RAW_INDEX_FIELD_DATA, *PRAW_INDEX_FIELD_DATA;
```

Cache Flushes

The HTTP.sys kernel-mode cache sometimes flushes completely. Significant configuration changes, such as changes to bandwidth throttling settings, can trigger a full cache flush. In the rare event that the cache flushes, the binary log file writes a record that indicates that you should dump all unique identifiers in memory and restart the process of recording the URL to new and unique ID mappings.

These records are rare and typically occur on servers that are reconfigured frequently. Example F.8 lists a cache flush record.

**Example F.8 Cache Flush Record**

| **Note** |
| --- |

Please be aware that Microsoft might change this data format at any time and is likely to do so in future Windows releases. Developers who use this data structure definition will need to alter program logic, recompile, and then re-release their applications for future releases of Windows (or service packs).

```
typedef struct _RAW_FILE_CACHE_NOTIFICATION
{
    // Must be RAW_RECORD_CACHE_NOTIFICATION_DATA_TYPE i.e. 5
    //
    USHORT          RecordType;

    // Reserved for alignment
    //
    USHORT          Reserved[3];

} RAW_FILE_CACHE_NOTIFICATION, *PRAW_FILE_CACHE_NOTIFICATION;
```

# HTTP.sys Error Log Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For the W3C Extended, IIS, and NCSA Common log file formats and for centralized binary logging, the log data is generated from the worker processes, passed to HTTP.sys through APIs, and then recorded in log files. However, because of the separation of the HTTP protocol stack in kernel mode and the worker processes in user mode, HTTP.sys-initiated error responses are not recognized by the user mode and are, therefore, not passed back to HTTP.sys for logging. To record these error responses, HTTP.sys needs its own error logging scheme that comes directly from the kernel mode. This is accomplished with HTTP.sys error logging. By default, HTTP.sys error logging is turned on.

**Common HTTP.sys Error Log File Entries**

There are a number of cases where HTTP.sys will generate its own errors. For example, when HTTP.sys sends an error response back to a client and a site is not yet associated with the request, an error is recorded in the HTTP.sys error log files. This error usually happens when there is a parse error in the request that was last received on the connection and HTTP.sys disconnects the connection after sending the response.

Another example of an HTTP.sys-generated error is a connection time-out. If a time-out occurs during a request, the request is recorded in the HTTP.sys error log file to provide more information about the connection.

Also, if a worker process in user mode unexpectedly terminates or closes its application pool handle while there are still outstanding requests that were queued and routed to this worker process, the outstanding requests are recorded in the HTTP.sys error log files. In this case, you can use the site ID field entry to identify the problematic site.

HTTP.sys error log files will sometimes contain connection errors that occur when IIS does not immediately destroy connections that the client terminates before a response for the last request on these connections is complete. HTTP.sys cannot resolve these connections (also known as *zombie* connections) until IIS completes its send or passes the last send call containing the logging information. Typically, unresolved connections are recorded in normal IIS log files. However, if IIS never sends the logging information back to HTTP.sys, unresolved connections are dropped after 30 seconds to one minute. HTTP.sys records these dropped connections in the HTTP.sys error log file to alert the administrator to their presence.

# Configuring HTTP.sys Error Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To change the configuration of HTTP.sys error logging, you must edit the registry. HTTP.sys reads the configuration only once during startup. If you change the configuration, HTTP.sys can pick up the change only at the next startup. The HTTP.sys error logging configuration is global.

You must have administrative credentials to edit the HTTP.sys configuration in the registry.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Windows Server 2003 Resource Kit Registry Reference. |

**Registry Parameters**

Three parameters control HTTP.sys error logging. They are located under the following registry key: **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\HTTP\Parameters**.

There are three entries that can be set:

- **EnableErrorLogging (DWORD)**. The default value is TRUE.
- **ErrorLogFileTruncateSize (DWORD)**. The default value is 1 * 1024 * 1024 bytes. The value cannot be smaller than this, and it must be specified in bytes.
- **ErrorLoggingDir (String)**. The default value is *systemroot*\System32\LogFiles. When specified, this value must be a fully qualified directory string, but you can use *systemroot*. For example:

  C:\LogFiles

  — or —

  *systemroot*\LogFiles

Although you can specify the directory, HTTP.sys will create a subdirectory named HTTPERR under the given log directory and put the error log files in this subdirectory. By default, members of the Administrators group and users with a LocalSystem user account have full access to the log files and directories. No other users have access rights.

The log file naming convention for HTTP.sys error logging is Httperr.#, where # is a sequence number. HTTP.sys error log files have a .log extension. For example, Httperr1.log. The error log files are recycled, based on their sizes. The default recycle size for error log files is 1 megabyte (MB). You can configure the value in the registry, but it cannot be smaller than 1 MB.

When the following events occur, HTTP.sys records them in the Windows Event log:

- HTTP.sys cannot create an error log file.
- The administrator sets an unacceptable configuration in the registry.
- There is not enough disk space.
- There is any other write failure.

# Log File Format for HTTP.sys Error Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

HTTP.sys error log files contain a fixed set of fields. Table 10.8 lists the fields that are recorded in HTTP.sys error logging log files and describes them.

**Table 10.8 HTTP.sys Error Logging Fields**

| Field | Description |
|---|---|
| Date | The date, in UTC time. This entry is always 10 characters long, for example, `2000-01-31`. |
| Time | The time, in UTC time. This entry is always eight characters long, for example, `00:12:23`. |
| Client IP | The IP address of the client. The version of the IP address can be either IPv4 or IPv6. If it is IPv6, the Scope ID field also appears. |
| Client Port | The port number of the client. |
| Server IP | The IP address of the server. The server IP address can be either IPv4 or IPv6. If it is IPv6, the Scope ID field also appears. |
| Server Port | The port number of the server. |
| Protocol version | The protocol version, if the last request on the connection has been parsed enough to identify the protocol version. If either the major or the minor version is greater than or equal to 10, the driver records the version as HTTP/?.?. |
| Verb | The verb, if the last request that was parsed passed the verb state. Unknown verbs are also recorded. HTTP.sys enforces a length limit of 255 bytes for the verb; anything longer is truncated. |
| CookedURL and query | The URL and its query, if both exist. A question mark (?) separates the URL from the query. If the URL of the request is completely processed (also known as *cooked*), then the processed URL is recorded with a local code page conversion and is treated as a Unicode field.<br><br>If only the unprocessed (raw) URL was present at the time of logging, then it is recorded as is, without a local code page conversion. HTTP.sys enforces a length limit of 4096 bytes for the URL; anything longer is truncated. |
| Protocol status | The protocol status of the response for the request, if it is available. The value cannot be greater than 999. |
| Site ID | The site ID, as a numeric value. For example, instead of recording W3SVC1, the field contains 1. There is no maximum value for the site ID. (This value can be as large as a MAX_ULONG.) |
| Reason phrase | Detailed information about why the error occurred, depending on the error type. This field can never be empty. For more information about reason phrases, see HTTP.sys Reason Phrases later in this section. |

For fields for which information is unavailable, a hyphen (-) appears as a placeholder. Fields are separated by spaces. If a field contains a nonprintable character, HTTP.sys replaces it with a plus sign (+) to preserve the log file format.

# HTTP.sys Reason Phrases

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following reason phrases can appear in an HTTP.sys error log file to describe the error that caused the log file entry:

**Connection_Abandoned_By_AppPool**
This reason phrase indicates that the error was caused by a worker process stopping unexpectedly and orphaning requests.

**Connection_Dropped**
This reason phrase indicates that a *zombie* connection was dropped by IIS and not resolved correctly.

**Various connection time-out errors**
These reason phrases include the following various connection time-out errors:

- Timer_ConnectionIdle. There has been no new data sent over the connection since the last send from the client to server, and the connection timed out.
- Timer_HeaderWait. A connection was initiated with the server, but the headers for the request were not received in a timely manner.
- Timer_MinBytesPerSecond. The minimum throughput rate was not maintained.
- Timer_EntityBody. The connection expired while waiting for the entity body to arrive.
- Timer_AppPool. The connection expired because the request waited too long in an application pool queue.

**Various errors**
These reason phrases include the following errors, most of which are parse errors:

- BadRequest
- Verb
- URL
- Header
- Hostname
- Invalid_CR/LF
- Number
- FieldLength
- RequestLength
- Forbidden
- LengthRequired
- Precondition
- EntityTooLarge

- URL_Length
- Internal
- N/I

**Internal**
This reason phrase indicates that an HTTP Error 500, internal server error, occurred.

**N/I**
This reason phrase indicates that an HTTP Error 501, not implemented, occurred.

**All 503 errors**
The 503 errors are service unavailable errors. These reason phrases include the following 503 errors:

- N/A. The service is unavailable.
- ConnLimit. The site connection limit has been reached.
- AppOffline. Because of rapid fail protection, the application was taken offline by IIS.
- QueueFull. The application request queue is full.
- Disabled. The administrator has taken the application offline.
- AppShutdown. The application was automatically shut down because of an administrator policy.
- AppPoolTimer. The application pool process is too busy to handle the request.

**Version_N/S**
This reason phrase indicates that an HTTP error 505, HTTP version not supported, occurred.

# HTTP.sys Error Log File Examples

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following example shows an HTTP.sys error log file entry, as viewed in a text editor:

```
2002-07-05 18:45:09 172.31.77.6 2094 172.31.77.6 80 HTTP/1.1 GET /qos/1kbfile.txt 503 -
ConnLimit
```

Table 10.9 lists and describes the fields that are recorded in this example.

**Table 10.9 Example of an HTTP.sys Error Log File Entry**

| Field | Appears As | Description |
|---|---|---|
| Date | 2002-07-05 | This log file entry was recorded on July 5, 2002. |
| Time | 18:45:09 | This log file entry was recorded at 6:45 P.M. |
| Client IP | 172.31.77.6 | The IP address of the client. |
| Client port | 2094 | The client port number. |
| Server IP | 172.31.77.6 | The IP address of the server. |
| Server port | 80 | The server port number. |
| Protocol version | HTTP/1.1 | The protocol version. |
| Verb | GET | The user requested a **GET**, which requests a file. |
| Cooked URL and query | /qos/1kbfile.txt | The query was unavailable, and the URL of the request was completely processed (also known as cooked), so /qos/1kbfile.txt is the local code page conversion. |
| Protocol status | 503 | The service was unavailable. |
| Site ID | – | The site ID was unavailable. |
| Reason Phrase | ConnLimit | The connection limit was reached. |

# Custom Logging Modules

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS supports custom logging modules that are COM objects that must expose a base set of methods for IIS to use. These methods are defined in the ILogPlugin, ILogPluginEx, and ILogUIPlugin interfaces. You can implement custom logging modules when you create your own log file formats or when you manipulate log data. For more information about the methods for the ILogPlugin, ILogPluginEx, and ILogUIPlugin interfaces, see Custom Logging Module Interfaces on MSDN. For information about creating custom logging modules, see Creating a Custom Logging Module for IIS on MSDN.

**Custom Logging Considerations**

Before you create and use custom logging modules, it is important to consider the following:

- If you want to implement custom logging for a site, consider synchronization issues. For example, different instances of a logging module can record information for different applications within the same site — but from different worker processes. To avoid synchronization problems, you can write a per-application pool log instead of a per-site log or you can use Microsoft® Win32® application programming interface synchronization primitives. Alternatively, for existing custom logging modules that you cannot change, all of the site's applications must be assigned to the same application pool to avoid the possibility of multiple worker processes writing entries to the same log file. In addition, the application pool must have the following configuration:

    - **MaxProcesses** = 1

    - **DisallowOverlappingRotation** = True

- Because the worker process, not HTTP.sys, handles custom logging modules, the HTTP.sys kernel-mode cache is disabled for the site for which a custom logging module is specified.

- In IIS 6.0, HTTP.sys uses the LocalSystem user account when writing log data to the currently selected LogFiles directory. Custom logging modules, however, use the accounts of the application pool in which they are running. The default account for application pools is the Network Service user account. For this reason, custom logging modules can cause write problems or errors when writing to the default LogFiles folder. You can correct this problem by configuring your custom logging modules to write to a different folder.

- If you choose to have custom logging modules write to a different folder, be sure to set a proper ACL on the new folder. Consider giving write permissions to the IIS_WPG group, which is a built-in user group that has the minimum set of permissions that IIS requires. If the identity account is not in the IIS_WPG group and has not been assigned appropriate rights, then the worker process fails to start. For more information about ACLs, see Managing a Secure IIS 6.0 Solution.

| **Important** |
| --- |
| When a custom logging module is loaded, IIS disables the HTTP.sys kernel-mode cache. For this reason, implementing a custom logging module can degrade overall server performance. |

Once you have created a custom logging module, you can add your format to the Active Log Format list in IIS Manager. To add the custom format to the Active Log Format list, you must edit the metabase with the Adsutil.vbs command-line script.

# Additional Resources for IIS 6.0 Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to this section.

**Related Information**

- Centralized Binary Log File Format for information about centralized binary logging.
- IIS 6.0 Architecture for information about the IIS 6.0 architecture and core components.
- Managing a Secure IIS 6.0 Solution for information about file and directory security and access control.
- The W3C Extended Log File Format site for information about the W3C Extended log file format.

**Related Help Topics**

- Enabling Logging for information about how to enable logging.
- Substatus Error Codes in Log Files for information about enabling and disabling substatus error code logging.

**Related Tools**

- Log Parser tool (LogParser.exe)

  Use the Log Parser tool to extract data from the raw file. The Log Parser tool and its accompanying user documentation are included in the IIS 6.0 Resource Kit Tools.

# Troubleshooting IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When administering Web applications, you will have problems to solve — whether fixing a service interruption, analyzing a configuration problem, or working around system resource contention problems. Whether the application is an internal corporate Web-presence or a widely available, scaled-out, data-intensive Web application, any problems that you encounter can seem nearly impossible to resolve, and attempting to do so consumes valuable time and resources.

This section describes how to troubleshoot Web applications on production computers running the Microsoft® Windows® Server 2003 operating system and Internet Information Services (IIS) 6.0. It explains how to use troubleshooting tools to administer IIS and discusses in detail service and program areas that might require troubleshooting, such as security, service startup, the IIS metabase, HTTP, and ASP.

**In This Section**


Overview of Troubleshooting IIS 6.0

Tools for Troubleshooting IIS 6.0

IIS Fundamentals

Common IIS Troubleshooting Tasks

HTTP Status Codes in IIS 6.0

Troubleshooting IIS Configuration Problems

Troubleshooting HTTP.sys Problems

Troubleshooting Miscellaneous IIS Problems

Advanced IIS Troubleshooting with Tracing and Debugging

Additional Resources for Troubleshooting IIS 6.0


**Related Information**

- For more information about troubleshooting, see Troubleshooting in IIS 6.0.

# Overview of Troubleshooting IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Problems that occur when administering IIS 6.0 might be difficult to solve for any of the following reasons:

- Problem reports from end users might be vague or incomplete, lacking such information as the content of error messages, how frequently a problem occurs, or the specific conditions under which the problem occurred.
- Logs do not record information that reveals the cause of the problem.
- The problem occurs in a complex, widely distributed application, which is difficult to track comprehensively.

The goal of troubleshooting is to isolate and determine the causes of problems that interfere with normal operation. This section describes a troubleshooting methodology and some of the tools that you can use to troubleshoot a server running IIS 6.0 and the Web applications that it hosts.

# Troubleshooting Methodology

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To efficiently troubleshoot Web applications, you need to apply and consistently use a troubleshooting framework, or methodology. Doing so helps you to streamline your troubleshooting activities and approach problems with confidence.

The fundamental troubleshooting methodology has five phases:

- **Phase 1: Discovery.** Gather information about the problem.
- **Phase 2: Planning.** Create a plan of action.
- **Phase 3: Problem Reproduction.** Reproduce the problem, or determine that you cannot reproduce it. If you cannot reproduce the problem, then you might not have enough information to confirm that there is a problem.
- **Phase 4: Problem Isolation.** Isolate the variables that relate directly to the problem.
- **Phase 5: Analysis.** Analyze your findings to determine the cause of the problem.

**Phase 1: Discovery**

Most problems are discovered by the end user. Problems might also be revealed by the operating system or by application logging and event tracking utilities. When a problem is reported, begin the troubleshooting process by gathering information to confirm that a problem exists.

**Interview the user who reported the problem**

If an end user reported the problem, ask that user for detailed information about what occurred. If possible, get the exact text of any error messages. If the problem recurs, ask the user to provide screen shots for error-message information or other symptoms that appeared on a screen. If the problem is behavioral and does not generate an error message, ask very specific, closed-ended questions. For example, you might ask the following questions:

- What time did the problem occur? (You can use this information to locate proximate events a large log file or to correlate the problem with external causes.)
- Which button(s) or link(s) did you click just before the problem occurred?
- Did you refresh or re-request the Web page to recover from the problem?

**Examine the records**

The Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Web Edition; and Windows® Server 2003, Datacenter Edition operating systems provide logging and event tracking. Check the logs for entries around the time that the problem occurred. Additionally, the application itself might log events. For example, the IIS logs and the HTTP error log record this information. For more information, see Analyzing Log Files.

**Phase 2: Planning**

Creating a plan of action is quite possibly the most important phase in the troubleshooting process. In this phase, you list the steps that you will follow as you proceed through the other troubleshooting phases. Stick to this plan and refer to it often. When you are bogged down in the tasks of the problem isolation phase, referring to your plan will help you remember where you are in the troubleshooting process and avoid getting sidetracked. Revise the plan as you progress — in most cases one step will dictate later steps.

When troubleshooting IIS–related problems, your plan might consist of configuring IIS logging to log extra details and then setting up a Performance Monitor log to run during a specific period of time. Based on the results of these actions, the next step might involve making a configuration change.

**Phase 3: Problem Reproduction**

Determine whether the problem is readily reproducible. The ability to reproduce the problem on demand is fundamental to properly and efficiently troubleshooting the problem. Use the information that you gathered in Phase 1 to reproduce the

problem, most often by repeating the actions taken by the end user before the problem occurred. If you can readily reproduce the problem, the isolation phase is more manageable. When you have determined a set of steps or events that trigger the problem, move to the isolation phase.

If you cannot readily reproduce the problem, the problem isolation phase can be tedious or isolating the problem can be impossible. If you cannot reproduce the problem, prepare to gather the right kind of information the next time it happens. Consider doing any or all of the following:

- Enable more detailed event tracking.
- Ask users to watch for the problem and pay close attention to what they are doing if it occurs.
- Write additional code in the application that looks for and highlights the problem if it happens again.

For example, when dealing with a Web application problem, you might add extra tracing code to your ASP pages or you might configure Windows security auditing for specific failures. The goal is to gather enough information so that the next time the problem occurs, the information that you obtain will be sufficient to correctly diagnose the problem.

**Phase 4: Problem Isolation**
In the problem isolation phase, you reproduce the problem as efficiently as possible by using repetitive steps. In this phase, you eliminate variables — such as settings, file actions, component starts/stops, or any change in the execution — that do not cause the problem, narrowing the variables down to those that are responsible for the problem.

You have succeeded in isolating the problem when you achieve the following conditions:

- The problem can be reproduced consistently when you take a fixed series of actions.
- The problem cannot be reproduced when you omit any of those actions.

For example, when troubleshooting IIS–related problems, you might do the following:

- Configure a Web application to be hosted in its own worker process or configure a component to execute in an isolated host process like a COM+ Dllhost.exe in order to determine which part of an application is consuming CPU time or leaking memory.
- Configure security settings on a URL in varying degrees through repetitive tests to isolate authentication problems.
- Capture performance data to determine whether a problem lies in the core Web server or in a Web service extension–based application, such as ASP or ASP.NET.

**Phase 5: Analysis**
Depending on the nature of the problem, the analysis phase can be the most difficult phase of the troubleshooting process. In the analysis phase, you use everything that you learned in the previous phases to do the following:

- Determine the cause of the problem.
- Explore how the problem affects the application.
- Determine the best way to solve the problem.

  This step might be a simple formality if, after Phase 3 is complete, you know what the problem is and how to fix it. If not, you must perform additional analysis of the data you have gathered. You might find that you must spend more time in the isolation phase.

  You might find that there are several ways to fix the problem. Take what you have learned about the problem and decide which actions to take. Ask yourself these kinds of questions:

  - What is the impact of the problem?

    - Do the benefits of fixing the problem outweigh the costs of fixing it?
    - Is there an acceptable workaround?

- Record your progress and the data that you have gathered.

  As you obtain information about a problem, you will make decisions about the steps that you need to take that will often alter your plan of action. It is important to record the information that you collect so that you have a way to account for your decisions as you move through the process of troubleshooting. If you find yourself taking the wrong path, having this information will allow you to backtrack to the point where you made the wrong decision. In addition, your ability to explain the changes or improvements in the problem is often important to other stakeholders.

# Tools for Troubleshooting IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Good tools are an indispensable part of any troubleshooting effort. This section identifies tools that you can use to help troubleshoot problems that you encounter on applications running on IIS 6.0. It also provides tips for getting the most benefit from each tool. Specific tool usage scenarios are also discussed in this section.

The following tools are useful for troubleshooting problems that occur with IIS 6.0:

- **WFetch (WFetch.exe).** This tool allows you to generate customized HTTP requests and view request and response data.
- **File Monitor (FileMon.exe).** This tool allows you to view and capture real-time file system activity.
- **Registry Monitor (RegMon.exe).** This tool allows you to view and capture real-time system registry activity.
- **IIS Request Monitor.** This tool captures information about HTTP requests in IIS worker processes, which can help you isolate and understand problems when worker processes become slow or unresponsive.
- **Secure Socket Layer (SSL) Diagnostic Tool (SSLDiag).** This tool helps you identify configuration problems in the IIS metabase, certificates, or certificate stores when running Web sites that use SSL.
- **Authentication and Access Control Diagnostics (authdiag.exe).** This tool allows you to review, test, and correct problems with authentication and authorization.
- **IIS 6.0 Enterprise Tracing for Windows.** This tool allows you to trace HTTP requests as they move through various components in your server architecture. IIS components provide informational events and error events (if applicable) to the ETW architecture as requests enter and exit each component. You can view the event data to see exactly where in the IIS architecture a request failed or became unresponsive.
- **Network Monitor.** This network tracing utility allows you to view activity in the network stack. Network Monitor is included with Windows Server 2003.
- **System Monitor.** This tool allows you to view and collect system performance data. System Monitor is included with Windows Server 2003.
- **HRPlus.** This tool provides error lookup functionality.
- **Microsoft Debugging Tools for Windows.** This tool allows you to debug and diagnose application problems.

# WFetch

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A simple HTTP request that is issued by typing a URL in the address bar of a browser might produce multiple network requests to the target server running IIS. However, the user does not see all of the resulting response data because the browser does not display it. WFetch allows you to fully customize an HTTP request and send it to a Web server so that you can see the raw HTTP request and response data. WFetch is included in the Internet Information Services (IIS) 6.0 Resource Kit Tools in the Microsoft Download Center.

For more information about using WFetch in a testing environment, see Knowledge Base article 284285, How To: Wfetch.exe to Troubleshoot HTTP Connections.

# File Monitor and Registry Monitor

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When troubleshooting IIS 6.0, you can use File Monitor and Registry Monitor to obtain two specific types of information:

- Verification that a specific file (File Monitor) or registry key (Registry Monitor) is being accessed by a given process.
- Information about whether that access results in an Access Denied error.

Both File Monitor and Registry Monitor can be downloaded from the System Internals Web site.

To install File Monitor and Registry Monitor, unzip and copy each .exe file to a desired location. To use one of these tools, double-click the appropriate executable file (FileMon.exe or RegMon.exe). Both utilities start in capture mode, which means that they display all file system (File Monitor) or registry (Registry Monitor) activity from all processes as it happens. Before you begin troubleshooting a problem, stop the capture and clear the entries that have been written.

**To stop or start a capture in FileMon.exe or RegMon.exe**

- On the **Options** menu, click **Capture Events**.

**To clear the current capture in FileMon.exe or RegMon.exe**

- On the **Edit** menu, click **Clear Display**.

**Tips for Using File Monitor and Registry Monitor for Troubleshooting**

Use the following tips to streamline the troubleshooting process when using File Monitor and Registry Monitor:

- **Isolate your task as much as possible.** The output from a capture comes very quickly and at a high volume. The larger the capture, the harder it is to find the information that you want. When you are ready, perform only the tasks that will produce the file or registry access that you are trying to verify.
- **Avoid unnecessary clicking or activity.** Performing extraneous operations while File Monitor or Registry Monitor is running will produce unrelated and unnecessary data in the output.
- **Start a capture as closely as possible to when you perform the task.** Separate your utility window from the workspace that contains the process that you are using to troubleshoot so that you can easily switch between programs. If possible, press ALT+TAB to quickly switch programs.
- **Create a capture filter for the process name that you are working with.** Registry Monitor and File Monitor both have capture filter functionality, which you can use to remove extraneous activity from the capture. You can filter on any string in an entry. For example, a common filter is the IIS 6.0 worker process name, W3wp.exe.

**To set a capture filter for W3wp.exe**

1. On the **Edit** menu of the appropriate tool, click **Filter/Highlight**.
2. In the **Include** box, type **W3wp.exe** and then click **Apply**.

When you have finished the task, stop the capture and scan the results for specific output. Usually, your objective will be to verify whether the problem is being caused by or is related to an access denial. In either tool, you can see whether this is the case by looking for the text ACCESS DENIED in the Result column.

If you need to compare the time that an event occurred with some other activity on your network, change the time format of the capture from the elapsed time, which is the default, to the literal time of day.

# IIS Request Monitor

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS Request Monitor is a command-line tool that captures information about HTTP requests in IIS worker processes. Administrators and developers can use Request Monitor to understand which HTTP requests are executing in a worker process when the worker process has become unresponsive or very slow. Request Monitor works with IIS 6.0 running on Windows Server 2003, Service Pack 1 or later.

For more information about the IIS Request Monitor, see Currently Executing Requests Tracing.

# SSL Diagnostics Version 1.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

SSL Diagnostics Version 1.0 helps you identify configuration problems in the IIS metabase, certificates, or certificate stores when running Web sites that use SSL. This tool allows you to review SSL configuration information, simulate the SSL handshake to locate errors, or "hot swap" certificates for testing purposes.

For information about the x86 version of SSL Diagnostics Version 1.0, see SSL Diagnostics Version 1.0 (x86).

For information about the 64-bit version of SSL Diagnostics Version 1.0, see SSL Diagnostics Version 1.0 (ia64).

# Authentication and Access Control Diagnostics Version 1.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Authentication and Access Control Diagnostics (Authdiag.exe) Version 1.0 allows you to review, test, and correct problems with authentication and authorization. You can use Authdiag.exe to check settings on Web sites, FTP sites, virtual directories, Web directories, and files. Authdiag.exe can help you troubleshoot the following types of issues:

- ACCESS DENIED errors caused by the following situations:
    - NTFS authentication failures
    - Required privileges missing
    - Invalid permissions on registry keys
    - Incorrect configuration for IIS metabase properties.
- Authentication failures based on invalid configuration.
- UNC configuration issues.
- Failure of the system default permissions for IIS, based on incorrect permissions to directories such as *%windir%*\system32\inetsrv.
- Failure within an Active Server Page (ASP) based on failed permissions.
- Failure to access the Home Directory or FTP Root because the Anonymous User or Authenticated User does not have the appropriate permissions.

Authdiag.exe provides both user interface (UI) and command-line modes of operation. Use the UI mode for troubleshooting authentication and authorization problems on a local server. Use the command-line mode for troubleshooting on a local server.

# IIS 6.0 Enterprise Tracing for Windows

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Windows Server 2003 includes Enterprise Tracing for Windows (ETW), a framework for implementing tracing providers that can be used for debugging and capacity planning. IIS 6.0 implements a number of providers to trace key transitions and the progress that an HTTP request makes as it moves through the IIS service framework. These providers include the following:

- IIS: WWW Server. This provider traces the transition of a request from HTTP.sys through various IIS subcomponents.
- IIS: SSL Filter. This provider traces calls into and out of the SSL Filter, which is used for SSL traffic.
- IIS: ISAPI Extension. This provider traces the transition of a request into and out of Internet Server API (ISAPI) extension processing.
- IIS: ASP. Traces the start and completion of the execution of an ASP request.
- IIS: IIS Admin Tracing. Traces the IIS Admin service during startup and shutdown.

For more detailed information about tracing IIS requests, including an introduction to tracing, an explanation of the different tracing providers in IIS, the different types of tracing, and how to troubleshoot IIS with trace data, see Troubleshooting IIS 6.0 with Tracing.

For the procedures to trace IIS requests and view trace data, see Advanced Troubleshooting with IIS Tracing.

For an explanation of the tracing events returned by IIS, see Tracing Events Reference.

# Network Monitor Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When troubleshooting IIS–related problems, you can use Network Monitor to start a capture, reproduce a problem, and view the capture. Network Monitor is included with Windows Server 2003.

**Procedures**

**To install Network Monitor**

1. From the **Start** menu, click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. In the left pane, click **Add/Remove Windows Components**.
4. Click **Management and Monitoring Tools** and then click **Details**.
5. Select the **Network Monitor Tools** check box and then click **OK**.
6. Click **Next**.
7. When installation is complete, click **Finish**.

   Network Monitor is now available on the Administrative Tools menu.

| **Note** |
| --- |
| Network Monitor, unlike WFetch, is a good tool for analyzing TCP connection management or three-way handshake information. For more information about using Network Monitor in this way, see Knowledge Base article 172983, Explanation of the Three-Way Handshake via TCP/IP. |

**Tips for Using Network Monitor for Troubleshooting**

Consider the following tips when using Network Monitor for troubleshooting:

- As with any utility that captures large volumes of data, it is important to plan your Network Monitor capture and to time the beginning and end of captures so that they minimize extraneous network traffic.
- Network Monitor, like File Monitor and Registry Monitor, can display the timestamp instead of the elapsed time in the capture. Displaying the timestamp along with synchronizing system clocks on participating computers can help you accurately isolate problems.
- Increase the capture buffer size from its default of 1 MB to 20 MB.

**To adjust Network Monitor capture settings**

1. On the **Capture** menu, click **Buffer Settings**.
2. In the **Buffer Size (MB)** box, type **20**.
3. Click **OK**.

Applying filters to captures allows you to record only relevant data. You can also apply a filter after the capture has completed to filter the data that is displayed. Filtering by protocol and by network addresses are the two most common ways to quickly filter Network Monitor capture data.

**To filter capture data by HTTP protocol**

1. While viewing a saved capture file, on the **Display** menu, click **Filter**.
2. In the **Display Filter** dialog box, click **Protocol == Any**, and then click **Edit Expression**.

3. In the **Expression** dialog box, click **Disable All**.
4. In the **Disabled Protocols** list box, click **HTTP**, and then click **Enable**.
5. Click **OK** twice.

   Network Monitor filters and then displays only HTTP traffic.

| To filter capture data by network addresses |
| --- |
| 1. While viewing a saved capture file, on the **Display** menu, click **Filter**.<br>2. In the **Display Filter** dialog box, click **ANY<-->ANY**, and then click **Edit Expression**.<br>3. In the **Expression** dialog box, choose the source and destination network. If you don't see the one you want, type the address that you want in the **Edit Addresses box**. |

By default, Network Monitor only parses traffic on port 80 as HTTP traffic. You can customize Network Monitor to parse traffic on other ports, for example, on port 8080, as HTTP traffic.

**To configure Network Monitor to parse traffic on non-standard HTTP ports**

1. Open Notepad.
2. Open the **Tcpip.ini file**, which is located in the **Parsers** folder under the Network Monitor installation point (by default, *systemroot*\System32\Netmon).
3. Locate the **[TCP_HandoffSet]** section.
4. Below the **80=HTTP** line, type the following new line: **8080=HTTP**
5. Save and then close the Tcpip.ini file.
6. Open the **Http.ini file**, which is located in the same folder as Tcpip.ini.
7. Locate the **[Ports]** section.
8. Immediately after the text **SrcPorts=80**, type the new port, preceded by a comma. The line now reads **SrcPorts=80,8080**.
9. Save and then close the Http.ini file.

To start any new parsing activity, close and reopen your captures.

For more information about HTTP traffic and Network Monitor, see Knowledge Base article 252876, How to View HTTP Data Frames Using Network Monitor. For more information about capturing network traffic, see Knowledge Base article 148942, How to Capture Network Traffic with Network Monitor.

# System Monitor Tool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

System Monitor is part of the Performance snap-in, which is accessible through Administrative Tools. When generating a performance log for troubleshooting, add, at a minimum, the following counter objects:

- .NET CLR (Common Language Runtime)
- Active Server Pages (ASP)
- ASP.NET
- Memory
- Process
- Processor
- Thread
- Web Service

Alternatively, you can observe real-time data on the display, if doing so meets your troubleshooting needs.

For information about starting System Monitor, and selecting performance objects and counters, see Viewing Counter Data in the Performance Console.

# HRPlus

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

HRESULTS are a type of 32-bit error code returned from various parts of the Windows Server 2003 operating system and its components. HRPlus is a free tool that is used primarily for decoding HRESULT values returned by COM applications, but it also can be used to generate custom HRESULT values for use in your own applications. Additionally, you can use HRPlus to decode Win32 error codes.

To obtain a copy of HRPlus, search for "hrplus" on the Simtel Web site.

# Microsoft Debugging Tools for Windows

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Microsoft Debugging Tools for Windows toolkit provides a robust suite of utilities for troubleshooting and debugging software. For more information about Microsoft Debugging Tools for Windows, see Debugging Tools for Windows - Overview, and Advanced IIS Troubleshooting with Tracing and Debugging.

# IIS Fundamentals

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS is an HTTP protocol server, which means that it accepts and replies to HTTP requests. IIS processes three types of HTTP requests:

- Static content (like HTM document) requests
- CGI requests
- ISAPI requests

By understanding the fundamentals of IIS architecture and the related functionality, you can discover and isolate problems much more quickly. You will also be better equipped to differentiate between real problems and expected behavior.

This section explores the following functionality, which any IIS troubleshooter must understand:

- HTTP Protocol Basics: Explains the way Internet browsers and servers running IIS communicate with each other.
- IIS Service Startup: Sometimes problems arise even before IIS starts servicing requests. This section describes the startup process and typical startup behavior, which can help you determine the causes of problems.
- HTTP Request Walkthrough: Explains the various phases that a typical request passes through as it moves through the service.

# HTTP Protocol Basics

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Troubleshooting IIS often requires that you examine the communications between a Web client and the server running IIS. By understanding the fundamentals of this communication, you can easily interpret network monitor traces and similar records.

Hypertext Transfer Protocol (HTTP) is the communication protocol used to exchange information between a client system and a Web server across a TCP/IP connection. The interchange is generally referred to as an HTTP transaction. In an HTTP transaction, the client system opens a TCP connection with the Web server and submits an HTTP request. The Web server, in turn, issues an HTTP response, completing one HTTP transaction. HTTP is stateless, in that there is no provision in the protocols design for information about any single request persisting beyond one transaction.

The following code segment is an example of an HTTP GET request:

```
GET /default.htm HTTP/1.1\r\n
Host: rhynoruv\r\n
Accept: */*\r\n
\r\n
```

The request line contains an HTTP method, a Uniform Resource Identifier (URI), and the HTTP version number, followed by carriage return and line feed characters. Typically, the request line is followed by a series of HTTP headers. Depending on the type of request, an HTTP request can have a variety of headers. With HTTP 1.1, the simplest GET request consists of the request line and only one HTTP header — the Host header. Two carriage returns/line feeds designate the end of the HTTP request header. In a simple GET request, these indicate the end of the request. Some requests include data, which is called the *entity body,* in addition to the HTTP headers. In the HTTP request, the division between the headers and the entity body is designated by two carriage returns/line feeds.

The following code sample shows an HTTP response, which is similar to a request:

```
HTTP/1.1 200 OK\r\n
Content-Length: 50\r\n
Content-Type: text/html\r\n
Last-Modified: Sun, 20 Oct 2002 22:52:16 GMT\r\n
Accept-Ranges: bytes\r\n
ETag: "255591568b78c21:5cd"\r\n
Server: Microsoft-IIS/6.0\r\n
MicrosoftOfficeWebServer: 5.0_Pub\r\n
X-Powered-By: ASP.NET\r\n
Date: Sun, 20 Oct 2002 22:52:35 GMT\r\n
\r\n

<HTML>\r\n
<BODY>\r\n
\r\n
Hello World.\r\n
\r\n
</BODY>\r\n
</HTML>
```

The first line is the status line, which consists of the HTTP version, a numeric HTTP status code, and then a reason phrase that briefly explains the status code. The status code is often the first indicator of a problem. For information about status codes, see HTTP Status Codes in IIS 6.0.

There are two other important aspects of the HTTP protocol that you must understand before you attempt troubleshooting: the provisions for managing TCP connections and how credential management (authentication) occurs.

# HTTP Connection Management

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because the TCP/IP protocol is so robust, establishing a new TCP connection for every HTTP request is relatively expensive in terms of network roundtrips. The two versions of HTTP (1.0 and 1.1) avoid incurring this overhead differently.

HTTP 1.0 transactions typically involved a client opening a TCP connection with the Web server and issuing an HTTP request. The server would then send its response and close the TCP connection. The client could optionally ask that the TCP connection remain open beyond the first transaction with the HTTP header **Connection: keep-alive**.

Starting with HTTP 1.1, persistent connections are implicit; in other words, the Web client and Web server keep the underlying TCP connection open until either sends an indication that the connection can be closed with the **Connection: close** header. Because the connection remains open, a way to designate the end of the entity-body is required. The **Content-Length** header performs this function by specifying the size of the entity-body, in bytes.

# HTTP Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

HTTP contains provisions for negotiating access to a secured resource, a process called *authentication*. Typically, an initial request from a client will not contain any authentication information; such a request is considered an *anonymous* request. If the requested resource is secured, the Web server can respond with a 401-Access Denied error and, optionally, can include one or more **WWW-Authenticate** HTTP headers. Each **WWW-Authenticate** header indicates a type of authentication scheme the server uses. The client can then choose to engage in one of these authentication schemes; doing so sometimes requires several more HTTP roundtrips.

For more information about authentication schemes, see 401.1 and 401.2-Authentication Problems. For more information about IIS authentication types, see Managing a Secure IIS 6.0 Solution.

For more detailed information about the HTTP protocol, see RFC 2616, Hypertext Transfer Protocol — HTTP/1.1 at the Internet Engineering Task Force Web site.

# IIS Service Startup

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

On rare occasions, you might encounter problems related to starting the IIS services infrastructure. The IIS services are generally robust enough to handle problems and recover from them gracefully. If a problem does occur and if it is appropriate, an event is logged to the application/system event log explaining the failure. Typically, the startup process never needs troubleshooting beyond examining messages written to the event log and taking the required corrective action. However, because it is helpful to understand the inner workings of an application in order to recognize problematic behavior, this section describes the major stages of the IIS services infrastructure startup process.

For information about the components of the services infrastructure, see IIS 6.0 Architecture.

The Windows Service Control Manager (SCM) performs Windows operating system service management operations. Among other things, SCM is responsible for signaling services to begin the startup process and for managing service dependencies. Some services are configured to *auto-start*, which means that SCM asks those services to start when the system is restarted. Other services are configured to manually start (or *demand-start*), which means that SCM starts those services only when prompted by an external request or when a service dependency demands it.

IIS services startup begins when the computer restarts. The IIS Admin service, which is hosted by Inetinfo.exe, is configured as an auto-start service. This is the first service in the IIS service infrastructure that is asked by SCM to start. The IIS Admin service depends upon the Remote Procedure Call (RPCSS) service and the Security Accounts Manager (SAMSS) service, which means that SCM guarantees that these are started before the IIS Admin service.

The next auto-start service in the IIS services infrastructure that starts is the World Wide Web Publishing Service (WWW service). It is hosted by Svchost.exe. In addition to the IIS Admin and RPCSS services, the WWW service has dependencies on the HTTP SSL (HTTPFilter) service. The IIS Admin service and RPCSS were started earlier, but before SCM starts the WWW service, it needs to demand-start the HTTPFilter service.

The HTTPFilter service has dependencies on the IIS Admin service and the HTTP service. It is this dependency on HTTP that causes SCM to start the HTTP.sys kernel-mode driver. Consequently, SCM starts the HTTPFilter service, which runs in Lsass.exe (or Inetinfo.exe if you are running in IIS 5.0 isolation mode), and then starts the WWW service.

The following sections describe what happens when these services start.

**HTTP.sys**

In IIS 6.0, unlike in earlier versions of IIS, the HTTP.sys kernel-mode device driver is not part of the core IIS services. In Windows Server 2003, the HTTP.sys driver actually exposes an API set that any application or service, not just IIS, can leverage. However, since IIS 6.0 is the major client, so to speak, of HTTP.sys, it is beneficial to understand how HTTP.sys relates to IIS 6.0.

The primary responsibility of HTTP.sys is to listen for and respond to HTTP requests on the network endpoints on which it is requested to do so. When HTTP.sys starts, by default it listens for HTTP requests on all IP addresses configured on the system. In some cases, though, it is beneficial for HTTP.sys to listen only on specific addresses — for example, if Windows Media Services needs to listen on port 80 of a specific IP address, outside of the HTTP.sys driver. To accommodate this need, HTTP.sys has an IP inclusion list. This list, when it contains entries, tells HTTP.sys which IP addresses to listen on, and, thereby, which IP addresses not to listen on. At startup, HTTP.sys examines this list and configures itself accordingly. As a result, changes to this list take effect only when the service is restarted.

| **Note** |
| --- |
| You can use the HttpCfg.exe utility to configure the IP inclusion list. For information about HttpCfg.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click **Windows Support Tools**. |

**IIS Admin Service**

The IIS Admin service primarily is responsible for managing the IIS metabase. During startup, it reads the MetaBase.xml file and creates its in-memory representation. Because the IIS Admin service creates this in-memory snapshot, if the metabase file is corrupt or if the XML code in the file is not syntactically well-formed, the IIS Admin service will fail to start. The IIS Admin service is responsible for synchronizing access to the metabase through the various control interfaces, even when manual updates are made directly to the MetaBase.xml file. During startup, the IIS Admin service sets up the control interfaces, like IMSAdmin, and registers for necessary file change notifications from the Windows Server 2003 operating system in order to monitor changes to the XML file.

**WWW Service**
The WWW service recognizes the following entities:

- URLs. The combined application and site binding information.
- Applications. Unique URLs that, when merged with a sites bindings and endpoints, provide full mapping information that enables HTTP.sys to determine which application pool requests that application should be directed to, and the metadata that should be used to support them.
- Sites. The specific bindings and endpoints that HTTP.sys should listen to. Sites contain one or more applications.
- Application pools. Sets of applications that are routed to the same set of worker processes. Application pools are also a common configuration for the worker processes that serve the applications in the pool.
- Application root. Also known as the Web application starting point, an application root is designated when you create and configure an application on a node in a Web site directory structure. You can have more than one application per Web site. The root of a newly created Web site, including the default Web site that is created when you install IIS, is an application starting point. Every file and directory beneath the starting-point directory in your Web site is considered part of the application until another starting-point directory is found. Thus, directory boundaries define the scope of an application.

When the WWW service starts, it performs the following operations:

1. Reads WWW Service configuration information from the IIS Admin service. This information includes global data, application pool data, Web site data, and application data.
2. Validates this data and logs any errors to the event log, choosing to ignore a record or to alter properties in the record to the defaults.
3. Configures HTTP.sys for all sites, applications, and application pools that are registered, and starts the ones that are set to start automatically.
4. Enables HTTP.sys configuration all at once, allowing IIS to accept requests.
5. If IIS 5.0 isolation mode is enabled, starts a worker process in Inetinfo.exe and waits for the worker process to start.

At this point, IIS can accept incoming HTTP requests. The first HTTP request to an application pool launches the worker process (W3wp.exe). In other words, an application pool can be started even though a worker process is not running.

**Worker Process Startup**
The fundamental aspects of a worker process startup routine are fairly straightforward. The worker process typically does the following:

1. Starts the metabase listener and listens for change notifications.
2. Initializes the IIS thread pool.
3. Obtains the application pool, which was specified in the command-line parameters, and then sets up the communication mechanism that it will use to receive requests from HTTP.sys.
4. Sets up communication with the WWW service to enable pinging and recycling.
5. Initializes other internal caches and structures that are used to process HTTP requests.
6. Initializes the Web service extension restriction list.
7. Begins accepting requests from HTTP.sys.

# HTTP Request Walkthrough

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

An HTTP request takes only one of three basic paths through the WWW service depending on whether the request is for a static file, an ISAPI extension, or for CGI processing. This section describes these paths. Understanding the request flow is invaluable during the problem isolation phase of the troubleshooting process.

# Common Request Operations

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

All HTTP requests arrive from the network and are passed up to the HTTP.sys listener from the TCP protocol stack. HTTP.sys parses the HTTP headers and makes sure that the request is well formed and valid. In most cases, HTTP.sys then determines into which application request queue to place the request.

Worker processes are signaled by the presence of requests in the HTTP request queue to start processing the requests. Using the HTTP headers, the worker process determines first which kind of authentication needs to be performed. It is here that the HTTP request goes through the process of getting a user context. All HTTP requests are authenticated (except in the case of anonymous requests served from the HTTP.sys kernel-mode cache) and have a Windows security context, which becomes the *user context of the request*. Even a request for content that has anonymous access enabled acquires the security context of the configured anonymous account; it is not anonymous in that there is no Windows user whatsoever. Through the authentication process, the worker process obtains a security token for the user by means of a Windows logon API. The user context and the associated Windows security token are fundamental and have a wide-reaching effect on how the rest of the request proceeds. Failures in this area and in areas related to security typically manifest themselves as HTTP 401 errors. For information about troubleshooting this type of error, see HTTP Status Codes in IIS 6.0.

After a user context has been obtained, the worker process determines from the HTTP headers the kind of content the request is requesting, and then chooses the appropriate handler. At this point, the three request types become relevant.

# Requests for Static Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A request for static content is the most basic kind of request that IIS handles. For this request type, IIS validates that the user has access to the file and then returns the file.

This request type also involves validation against the Multipurpose Internet Mail Exchange (MIME) type restriction feature. If the request is for a restricted file type, IIS returns a 404.3-MIME Map Policy Prevents This Request error. For more information, see Requests for static files return 404 error.

IIS can also cause a response to be cached in the HTTP.sys kernel-mode cache. For more information, see the Http.sys Response Cache.

# Requests Handled by ISAPI Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Active Server Pages (ASP) and ASP.NET requests are handled by ISAPI extensions. A request is designated an ISAPI request when it requests a file whose extension is mapped to an ISAPI extension DLL in the **ScriptMaps** metabase property. A request is also designated an ISAPI request when it requests an ISAPI DLL directly in the URL or when a wildcard **ScriptMaps** entry exists, whereby all requests go through the mapped ISAPI DLL. IIS calls a well-known entry point in the ISAPI DLL and passes a structure containing all of the information required about a request. When IIS is configured for IIS 5.0 isolation mode, ISAPI requests also can be passed *out-of-process* to a Dllhost.exe worker process. Before an ISAPI request is executed, IIS checks the **WebSvcExtRestrictionList** metabase property to ensure that the ISAPI is allowed. If the ISAPI is not allowed, IIS returns a 404.2 error.

In this kind of request, the ISAPI extension DLL is responsible for processing the HTTP request and then sending the response. When an ISAPI extension DLL processes a request, any number of things can happen. In the case of ASP and ASP.NET, the user context obtained during authentication is typically impersonated before the ScriptMap file is opened. These kinds of requests, where distributed and dynamic operations are performed, often involve custom-built components and backend data stores.

# CGI Requests

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

*CGI requests* directly request a CGI executable. They are similar to ISAPI requests in that the CGI performs all of the execution work and also has to be enabled in the **WebSvcExtRestrictionList** metabase property. Unlike ISAPI requests, CGI requests involve launching a new process for each request. However, this design does not perform or scale very well — especially not on heavily trafficked Web sites. Consequently, CGI requests are typically used less often than other kinds of requests.

# Common IIS Troubleshooting Tasks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Several basic IIS troubleshooting tasks are helpful in many different troubleshooting situations. These tasks include: disabling HTTP friendly error messages in Internet Explorer, generating a custom HTTP request, checking basic request functionality, browsing with different host name styles, using substatus and Win32 errors in W3C extended logging, configuring the WWW service to log worker process recycling events, checking Windows NT system and application event logs, checking the HTTP error and IIS logs, restarting the IIS service, and identifying the process ID of the worker process. Perform these tasks during the initial steps of the discovery phase, or when initially diagnosing a problem.

# Generating an HTTP Request

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To observe the raw HTTP traffic between a browser and IIS, use WFetch. Commonly, you might use WFetch when a user receives an HTTP 403-Forbidden error. In such a situation, you might begin troubleshooting by checking the access control list (ACL) configuration on the requested content. If that does not reveal the cause, you most likely would start troubleshooting the problem as an authentication problem. However, the 403 message indicates that these are the wrong actions to take. Disabling friendly HTTP error messages in Internet Explorer and re-trying the request with WFetch would have yielded a much more meaningful error: Directory Listing Denied. This virtual directory does not allow contents to be listed. You would then immediately realize that Directory Browsing is disallowed for that virtual directory and that either Default Documents is disabled or it is enabled, but none of the configured documents are present in that directory.

**To generate a simple HTTP request using WFetch**

1. Start WFetch.
2. From the **Verb** list box, click **GET**.
3. In the **Host** box, type a host name — for example, *www.contoso.com*.
4. In the **Path** box, type the following: **/**
5. In the **Auth** list box, click **Anonymous**.

   The **Port**, **Ver** (HTTP Version), and **Connect** settings can be left at their defaults.
6. Click **Go**.

WFetch sends an HTTP request to the local instance of IIS and then receives the response. In the Log Output window, you can view the HTTP request and response data. Notice the individual headers in both the request and response, and the entity body in the response. Figure 11.1 WFetch Startup Screen below shows WFetch in the default startup state.



# Checking Basic Functionality with Test Request Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Occasionally, you might need to assess the core functional health of the various request types in IIS. Determining which kinds of requests work and which do not is a fundamental step in the isolation phase. The easiest way to accomplish this is to create several test files that can be placed in various application roots in your various Web sites.

Making these test requests from different locations is as important as making requests of different types. Consider that a directory location determines which application, URL, and application pool a request will be routed through. These locations can indicate at which point a particular kind of request is failing. The following examples are test request files for static content, ASP ISAPI, and ASP.NET ISAPI requests, respectively.

**Static Content Test Request File**
To create a static content test request file, create an HTML file with the following HTML code:

```
<HTML>
<BODY>
<H3>Hello World</H3>
</BODY>
</HTML>
```

Save this code in a file called Simple_Request.htm and use it to troubleshoot static content problems.

**ASP Test Request File**
For an ASP test file, you need include only dynamic ASP code inside of <% %> delimiters. However, output that returns the current time, user, and application pool is useful, too. To create an ASP request file that returns this information, create an HTML file with the following code:

```
<%@ Language=VBScript %>
<html>
<body>
<%
Response.Write ("The time is now: " & Now() & "<p>")
Response.Write ("The authenticated user is (blank if anonymous) : " & _
         Request.ServerVariables("LOGON_USER") & _
         "<p>")
Response.Write ("The application pool that served this request is: " & _
         Request.ServerVariables("APP_POOL_ID") & _
         "<p>")
%>
</body>
</html>
```

Save this code in a file called Simple_Request.asp.

**ASP.NET Test Request File**
The code for an ASP.NET test request is similar to that of an ASP test request:

```
<%@ Page Language="VB" Debug="True" Trace="True" %>
<%@ Import Namespace=System.Web%>

<HTML>
   <script language="VB" runat="server">

      Sub Page_Load(Src As Object, E As EventArgs)

      dim ctx as HttpContext = HttpContext.Current

      lblTime.Text = Now()
        lblUserIdentity.Text = ctx.Request.ServerVariables("LOGON_USER")
        lblAppPool.Text = ctx.Request.ServerVariables("APP_POOL_ID")

     End Sub

   </script>
   <body style="FONT: 8pt verdana">
```

```
      <form runat="server">
       The time is now:
      <asp:Label id="lblTime" runat="server" /><p>
       The authenticated user is (blank if anonymous):
      <asp:Label id="lblUserIdentity" runat="server" /><p>
   The application pool that served this request is:
       <asp:Label id="lblAppPool" runat="server" /><p>
      </form>
   </body>
</HTML>
```

Save this code in a file called Simple_Request.aspx.

**Testing Simple Object Creation**

A simple COM or .NET object that returns a basic hello world string is a useful troubleshooting tool. Placing that object locally on the server running IIS or remotely on your application tier, and then testing from ASP or from a simple script file or compiled client can help you determine where problems exist in your distributed applications architecture.

**Testing Application Code Outside of IIS**

Another way to test basic functionality is to test your applications completely outside of IIS. For example, if you are testing COM object creation failures, you can create a simple Microsoft Visual Basic® Scripting Edition (VBScript) file that contains your object creation code and then run the script when you log on to the IIS Web server. Variations might involve creating other custom applications by using other development tools.

# Browsing with Different Host Name Styles

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Connectivity and authentication problems can be related to Domain Name System (DNS) issues, the configuration of host headers, and general network anomalies. To understand the nature of a problem, it is sometimes helpful to vary the way that you type the host name in the requested Uniform Resource Locator (URL). Use any of the following variations:

- Fully Qualified Domain Name (FQDN). A fully qualified DNS name that uniquely identifies the computer. This name is a combination of the DNS host name and the DNS domain name, in the format *Hostname,DomainName*. For example, if the DNS host name is host1 and the DNS domain name is example.contoso.com, the fully qualified DNS name is host1.example.contoso.com.

- Network BIOS (NetBIOS) name. The computer name recognized by Windows Internet Name Service (WINS) in the format *HostName*, which maps the name to an IP address. For example, the NetBIOS name for host1.example.contoso.com is host1.

- IP address. The address used to identify a unique node on an IP network. This address is typically represented with the decimal value of each octet, separated by a period (for example, 192.168.7.27).

| **Note** |
| --- |
| IIS6.0 also supports IPv6 addressing. For more information, see IPv6 and IIS 6.0. |

Browsing with different host name styles is useful for the following reasons:

- Certificate authentication requires matching the host name from the URL to the Common Name (CN) imprinted in the server certificate.

- In Kerberos authentication, the requested Kerberos service ticket is keyed off of the host name entered in the URL. If the wrong host name is used, a service ticket might not be obtainable.

- The host header configuration in the **ServerBindings** metabase property dictates how an HTTP request is routed into the service. This relies upon matching the requested host name with an appropriate host header configuration.

- The simple name resolution process can differ when you use a NetBIOS name instead of an FQDN. Using a NetBIOS name in the URL is similar to pinging to test name resolution of the NetBIOS name to the IP address.

- Internet Explorer can unexpectedly prompt for credentials during NTLM authentication (also referred to as Windows NT Challenge/Response authentication) when the host name portion of the requested URL contains a period. Using the NetBIOS name rather than an FQDN or IP address might indicate the cause. For more information, see Knowledge Base article 258063, Internet Explorer May Prompt You for a Password.

- Browsing by IP address is a foolproof way to rule out name resolution problems. However, the host header (set with the **ServerBindings** property) must be appropriately configured.

# Using Substatus and Win32 Errors in W3C Extended Logging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To reduce the attack surface of IIS 6.0, custom error messages do not return specific error message information, including the substatus code, to remote client computers. If a custom error message contains too much information about the core Web server and an explanation of why a particular request failed to execute, malicious users can use the information to attack the Web server. As an example, an error code such as 404.2-Web Service Extension Lockdown Policy Prevents This Request is returned to a Web client as 404-File or Directory Not Found, and, thereby, gives the Web client no indication of why the request failed.

Although this change to error message reporting reduces the possibility of an attack by a malicious user, it makes it difficult to debug a failed request. For this reason, you can enable substatus error codes logging to IIS log files when appropriate. Note that on clean installs, substatus code logging for W3C Extended format files is enabled by default. Doing so is particularly useful when the logging of Win32 error codes is enabled. Then, to obtain the information that you need, you simply locate the specific request in the log file and check the Win32 error code and substatus code. This feature is called W3C extended logging. Figure 11.2 Extended Logging Properties Sheet shows substatus codes as they appear on the extended logging properties sheet.




For more information about logging substatus error codes, see Substatus Error Codes in Log Files.

# Configuring the WWW Service to Log Worker Process Recycling Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Proper configuration of IIS is always essential to successful troubleshooting. By setting the **LogEventOnRecycle** metabase property, you can direct the WWW service to log worker process recycling events in the event log. To accomplish this, you can set the **LogEventOnRecycle** property directly as a bitmask or you can set individual **AppPoolRecycle** property settings to **true**. The **LogEventOnRecycle** property is set to **AppPoolRecycleMemory** by default.

This is important as there could be occasions when the problem reported happened at or near the same time as a worker process recycle for the application. Having this information makes it easier to troubleshoot exactly what happened at the time of the problem to the entire server.

Table 11.1 summarizes the event IDs that are logged, and the related metabase properties and their bitmasks.

**Table 11.1 LogEventOnRecycle Configuration**

| Recycle Condition | Event ID | Metabase Property | Bitmask |
|---|---|---|---|
| Elapsed time | 1074 | **AppPoolRecycleTime** | 1 (0x1) |
| Number of requests processed | 1075 | **AppPoolRecycleRequests** | 2 (0x2) |
| Scheduled recycle | 1076 | **AppPoolRecycleSchedule** | 4 (0x4) |
| Virtual memory consumed | 1077 | **AppPoolRecycleMemory** | 8 (0x8) |
| ISAPI reported itself unhealthy | 1078 | **AppPoolRecycleIsapiUnhealthy** | 16 (0x10) |
| Recycled on demand | 1079 | **AppPoolRecycleOnDemand** | 32 (0x20) |
| Recycled due to configuration change | 1080 | **AppPoolRecycleConfigChange** | 64 (0x40) |
| Private memory consumed | 1117 | **AppPoolRecyclePrivateMemory** | 128 (0x80) |

# Checking NT System and Application Event Logs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When looking for information about a problem, always consult the event logs for the application. Logging begins automatically when Windows is started. There are three types of logs: system, security, and application. For each type of log, you can filter the events to be viewed, designate the number of entries to view, specify how long to save entries, and specify whether to automatically overwrite existing events when the log becomes full.

For information about opening Event Viewer to see the event logs, see About Events.

# Checking the HTTP Error Log

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 6.0, HTTP.sys logs errors in the HTTP error log. The HTTP error log is the first place you should look when IIS returns an HTTP 503–Service Unavailable error, for example. The HTTP error log is located in *systemroot*\System32\LogFiles\HTTPERR.

For information about HTTP errors, see HTTP Status Codes in IIS 6.0.

# Checking the IIS Logs

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The IIS logs provide a great deal of information about the activity of a Web application. You can find the IIS logs in *systemroot*\System32\LogFiles\W3SVC*number*, where *number* is the site ID for the Web site.

# Restarting IIS Services

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

During troubleshooting, you might need to restart the IIS services. You can use IIS Manager to start or to stop any Web site on your server running IIS, to stop or recycle an application pool, or to restart the WWW service, which restarts the complete IIS service base.

# Restarting a Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Stopping a Web site causes the WWW service to remove the URLs for that site from the HTTP.sys configuration. The WWW service then marks the Web site as stopped in the metabase and directs the Web site not to start again automatically when the WWW service is restarted. Stopping a Web site does not cause loaded segments of the applications to unload, nor does it affect the worker processes that might be running when you stop the site. Additionally, requests that are waiting in the application pool queue in HTTP.sys will still be processed. However, new requests will either generate HTTP 404-Not Found errors, or, if no other Web site is listening on the port, will send the client a socket connection error.

Web sites can be stopped and started by using either IIS Manager or the IisWeb.vbs command-line script. For information about stopping and restarting a Web site by using either IIS Manager or IisWeb.vbs, see Starting and Stopping Web Sites.

# Recycling an Application Pool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Application pools are made up of a listening and routing structure in HTTP.sys and one or more ready-to-start worker processes that are waiting to process incoming requests. In worker process isolation mode, you can configure IIS to periodically restart worker processes under certain conditions. This automatic restarting of worker processes is called *recycling*.

Recycling an application pool causes the WWW service to shut down all running worker processes that are serving the application pool, and then start new worker processes. Whether the WWW service starts the new worker processes before it stops the existing one depends on the **DisallowOverlappingRotation** property in the metabase. Recycling an application pool does not alter any state in HTTP.sys or change any configuration in the metabase.

| **Note** |
| --- |
| When an application pool is serviced by more than one worker process, the recycle operation is staggered in some scenarios. Only when a worker process is recycled on demand or when all of the worker processes hit their memory limits at the same time would they be restarted simultaneously.<br><br>For information about recycling an application pool on demand, see Configuring Worker Processes for Recycling. |

# Stopping an Application Pool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Stopping an application pool causes the WWW service to shut down all running worker processes serving that application pool. HTTP.sys is then directed to return HTTP 503-Service Unavailable errors for any requests currently queued for this pool, and to return HTTP 503 errors for all future requests. The application pool is marked as stopped in the metabase and is configured not to start automatically when the WWW service is restarted. You must manually start the application pool to cause it to serve requests again.

# Restarting the WWW Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Sometimes, as a last resort, you need to restart the WWW service. For information about restarting the WWW service, see Starting and Stopping Services.

# HTTP Status Codes in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

HTTP status codes provide the best information about what is happening with IIS. On an active Web server, HTTP status codes provide a continuous feed of up-to-the-moment information about the activity and health of the server.

Most HTTP status codes have three digits, for example, 401. Some status codes have three digits followed by a decimal point, and one or two more digits (for example, 300.12). In such a case, the number that follows the decimal point is called the substatus code.

Codes within each hundred-number range provide specific types of information, as summarized in Table 11.2 HTTP Status Codes.

**Table 11.2 HTTP Status Codes**

| Status Code Range | Type of Code |
|---|---|
| 100 | Informational |
| 200 | Successful |
| 300 | Redirection |
| 400 | Client Error |
| 500 | Server Error |

As an administrator, at some point you will inevitably need to determine which HTTP status codes are being returned by IIS. There are a number of ways to do this. Your choice of method depends on the amount of information that you need and the type of problem that you are troubleshooting. If you have a problem that is limited to one or a few clients, viewing the status codes from the clients Web browser or issuing a reproduced HTTP request with WFetch might best serve your needs. However, if you need to look for an individual request or for a collection of multiple requests over a period of time, you should examine the IIS logs. Viewing IIS logs allows you to see not only the main HTTP status code, but the Win32 error code and the HTTP substatus code (if logging is configured to do so). The Win32 error codes and the HTTP substatus codes often contain information that is critical to your troubleshooting effort.

The following sections describe the HTTP status codes and provide troubleshooting suggestions.

# HTTP 1xx-2xx -- Informational and Success Codes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

HTTP status codes in the 1*xx* are typically informational. Status codes in the 2*xx* range indicate that the transaction was successful. Because the codes in both ranges are self-explanatory, they do not require further discussion.

# HTTP 3xx -- Redirection Codes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

HTTP status codes in the 3*xx* range pertain to *redirection* (telling the browser to issue a new request at a different location). These codes do not indicate that a problem has occurred.

Table 11.3 HTTP 300-Class Error Codes Returned by IIS lists the HTTP redirection codes that IIS returns.

**Table 11.3 HTTP 300-Class Error Codes Returned by IIS**

| Status Code | Condition |
|---|---|
| 301 | Permanent Redirect. |
| 302 | Object Moved |
| 304 | Not Modified. |
| 307 | Temporary redirect. |

# 301-Permanent Redirect

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This redirection status code tells the Web client that the location for the requested resource has permanently changed. The new location is specified in the **Location** HTTP header. This is primarily useful for Web clients that keep record of HTTP URL links. Upon receiving this status code, the Web client can update the stored URL with the new location.

# 302-Object Moved

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This standard redirection status code tells the Web client to issue a new request to the location specified in the **Location** HTTP header.

# 304-Not Modified

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

RFC 2616, "Hypertext Transfer Protocol — HTTP/1.1", provides functionality for conditional GET requests. For example, an HTTP client can specify an **If-Modified-Since** HTTP header that consists of a time and date stamp, which tells the server that it has a cached copy of the content being requested. If the request meets the condition specified in the **If-Modified-Since** header, the server will not return a body, but will return the 304-Not Modified result.

For more information about RFC 2616, see the W3C World Wide Web Consortium Web site.

# 307-Temporary Redirect

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The temporary redirect option, which is new in IIS 6.0, is used to prevent a Web browser from losing data when the browser issues an HTTP POST request. Normally, when a Web browser issues a POST request and then receives a 302-Object Moved redirect message from the Web server, the browser issues a GET request for the new location and loses the data in the **POST** request. With a 307 redirect, the browser reissues the POST request with the original data to the new location.

For more information about temporary HTTP redirects, see the HttpRedirect Metabase Property.

# Courtesy Redirects

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When a browser makes an HTTP request, the request sometimes takes the form http://example.contoso.com/vdir1. Notice that there is no trailing / at the end of the URL. In this case, vdir1 is a directory, instead of a file. As a *courtesy* to the Web browser, IIS responds to this request with a 302-Object Moved redirect message with the HTTP header: **Location: /***vdir1***/**. This is known as a courtesy redirect.

Courtesy redirection can produce unexpected behavior when the Web server is behind a Network Address Translation (NAT) device —a proxy server, for example. Suppose, for example, that the Web site is configured to listen on port 8080 and the NAT device is configured to forward HTTP requests on port 80 to the Web server on port 8080. In this configuration, the NAT device accepts the Web browsers request, but when the courtesy redirect occurs, IIS appends the port number to the host name. The NAT device needs to be aware of this and alter the response appropriately before returning it to the client.

For more information, see Knowledge Base article 298408, IIS Generates Courtesy Redirect When Folder Without Trailing Slash Is Requested, and Knowledge Base article 313074, How To Redirect Browser Requests in Internet Information Services 5.0.

# HTTP 40x -- Client Error Codes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

HTTP status codes in the 4*xx* range indicate that a problem occurred with the request. Request problems can range from the request not meeting authentication requirements to the request being malformed. RFC 2616, "Hypertext Transfer Protocol — HTTP/1.1", refers to request problems as client errors because the requests originate from the client. For more information about RFC 2616, see the W3C World Wide Web Consortium Web site.

Table 11.4 HTTP 400-Class Client Error Codes Returned by IIS lists the client error codes that IIS returns. The descriptions for most codes are self-explanatory. When additional information about a code is required, it is provided in one of the following sections.

**Table 11.4 HTTP 400-Class Client Error Codes Returned by IIS**

| Status Code | Condition |
|---|---|
| 400 | Cannot resolve the request. |
| 401.*x* | Unauthorized. |
| 403.*x* | Forbidden. |
| 404.*x* | File or directory not found. |
| 405 | HTTP verb used to access this page is not allowed. |
| 406 | Client browser does not accept the MIME type of the requested page. |
| 407 | Initial proxy authentication required by the Web server. |
| 412 | Precondition set by the client failed when evaluated on the Web server. |
| 413 | Request entity too large. |
| 414 | Request URL is too large and therefore unacceptable on the Web server. |
| 415 | Unsupported media type. |
| 416 | Requested range not satisfiable. |
| 417 | Expectation failed. |
| 423 | Locked error. |

# HTTP 400-Cannot Resolve the Request

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS returns a 400 error code when the request is malformed. Though there are a variety causes for this error, the most common reason is that the request does not comply with RFC 2616. For example, if the request is made in HTTP 1.1, RFC 2616 requires that the request include a **Host** header. The simplest HTTP 1.1 request would look similar to the following:

```
GET / HTTP/1.1
Host: example.contoso.com
```

You can use WFetch to issue such a custom-built request. Most HTTP clients adhere to HTTP RFC 2616 and issue well-formed requests. However, if you are developing a custom HTTP client and are experiencing HTTP 400 problems, use WFetch to make test requests.

**To issue a simple HTTP 1.1 request using WFetch**

1. Start **WFetch**.
2. From the **Advanced Request** list box, click **Raw Request**.
3. In the box, type the following text, where *ComputerName* is the name of your server running IIS:

   ```
   GET / HTTP/1.1\r\n
   Host: ComputerName\r\n
   \r\n
   ```

4. Ensure that the **Host** box contains the same *ComputerName* as noted above.
5. Click **Go**.

This sends the most basic HTTP 1.1 GET request to the server.

HTTP.sys records some errors in its own log file. Table 11.5 HTTP.sys Logged Reason Codes for HTTP 400 Errors lists the HTTP.sys reason codes and explanations for each of the 400 errors. These errors are not recorded in the IIS logs.

**Table 11.5 HTTP.sys Logged Reason Codes for HTTP 400 Errors**

| Reason Code | Condition |
|---|---|

| | |
|---|---|
| BadRequest | The request could not be understood by the server because the syntax was incorrect.<br>In the following instances, HTTP.sys, the kernel-mode driver, will reject a request as a BadRequest. HTTP.sys returns an HTTP 400 error (Bad Request) to the client. This HTTP 400 error is not customizable:<br>• The verb used in the request is invalid.<br>• The requested URL is invalid.<br>• The requested header name is invalid.<br>• The requested hostname is invalid.<br>• The request contains an invalid carriage return and/or line feed as a terminator.<br>• The request contains an invalid content length header or invalid entity chunk length.<br>• The cumulative length of the headers in the request is longer than the maximum allowed.<br>• The header field in the request is too long.<br>• The request header is too long. |
| Verb | Invalid verb. |
| URL | Invalid URL. |
| Header | Invalid header name. |
| Hostname | Invalid hostname. |
| Invalid_CR/LF | Invalid carriage return or line feed. |
| Number | Invalid number. |
| FieldLength | A header field in the request was too long. |
| RequestLength | The request length was too long. |

# HTTP 401.x-Unauthorized

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Authentication is one of the first operations performed when an HTTP request is issued. *Authentication* is the process whereby IIS creates a user context for an HTTP request, typically by obtaining credentials from the HTTP client using a preconfigured authentication method, and then calling a Windows logon API using those credentials. The API then returns a security token for the request, which represents the user context.

After the authentication process concludes, IIS determines where the HTTP request goes next based on the resource that is being requested. Regardless of this choice, IIS issues an authorization check against the requested resource. IIS checks to ensure that the user context associated with this request is allowed to make the request. Usually, IIS performs a file ACL check to authenticate the request.

When IIS cannot authenticate a request, it returns a 401.*x*-Unauthorized code. The substatus codes provide detailed information about why the request failed, as shown in Table 11.1 LogEventOnRecycle Configuration. The descriptions for most substatus codes are self-explanatory. When additional information about a substatus code is required, it is provided in one of the following sections.

**Table 11.6 HTTP 401-Substatus Codes**

| 401 Substatus Code | Condition |
|---|---|
| 1 | Access is denied due to invalid credentials. |
| 2 | Access is denied due to server configuration favoring an alternate authentication method. |
| 3 | Access is denied due to an ACL set on the requested resource. |
| 4 | Authorization failed by a filter installed on the Web server. |
| 5 | Authorization failed by an ISAPI/CGI application. |
| 7 | Access denied by URL authorization policy on the Web server. |

# 401.1 and 401.2-Authentication Problems

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Typically, HTTP 401.1 and 401.2 errors are encountered when the authentication process fails in some way — either because the mechanism that IIS used to obtain the credentials has failed or because the credentials themselves are invalid. This is an important distinction that is a critical to the problem isolation process. These errors tell you that, before performing any other function, IIS has failed to authenticate the user.

For a brief overview of the authentication options available in IIS, see Managing a Secure IIS 6.0 Solution.

**Initial Isolation**
When you know that authentication has failed, gather the following information:

- The URL that is being requested
- The authentication method that is configured for that node in the metabase

When you know the requested URL, you can use IIS Manager to verify the authentication methods that are enabled. Alternatively, you can use the Adsutil.vbs command-line script to check the **AuthFlags** metabase property. For information about the settings for this property, see AuthFlags Metabase Property.

**Multiple Authentication Methods**
Typically, initial HTTP requests to a server are anonymous. In most cases, a Web browser does not know to send authentication information, and, therefore, it issues a request without this information. If Anonymous authentication is enabled for the requested resource, IIS does not attempt to use other authentication protocols, but instead attempts to obtain a security token for the IUSR_*ComputerName* account by calling one of the Windows logon APIs. If this fails, then IIS looks to see if other authentication methods are enabled for the URL and sends back a 401.2 response to the browser that includes a list of **WWW-Authenticate** HTTP headers that specify the authentication schemes IIS is configured to accept for that URL.

At this stage, typically it is up to the HTTP client to choose to continue the HTTP transaction, to pick a protocol from the list, and to proceed to authenticate with that protocol. If this process is successful, IIS then calls a logon API and obtains a security token that it can use to associate a user context with the HTTP request.

**LogonMethod Default Has Changed**
The **LogonMethod** metabase property designates the kind of logon token that IIS asks for when calling the Windows logon APIs for Anonymous or Basic authentication. In earlier versions of IIS, the default was set to **0**, which caused IIS to request an INTERACTIVE token. This sometimes caused problems when the IUSR_*ComputerName*, or, in the case of Basic authentication, the user credentials did not have the **Log on locally** user right. In such a case, IIS returns a 401.1 error.

In IIS 6.0, the default value of **LogonMethod** has changed to **3**, which evaluates to a NETWORK_CLEARTEXT token. This token is very similar to an INTERACTIVE token; however, it does not allow the user to physically log on to the system console and does not require the Log on locally user right.

For more information, see Securing Sites with Web Site Permissions.

**No Authentication Methods Selected (AuthFlags=0)**
IIS returns a 401.2 error when no authentication methods are enabled. If you receive this error, ensure that an authentication method is enabled.

**Anonymous Accounts (IUSR_ComputerName) Attempting Sub-authentication Logon Receive the 401.1 Error: Logon Failed**
By default, the sub-authentication component, Iissuba.dll, is not enabled in IIS 6.0. In earlier versions, Iissuba.dll allowed IIS to manage passwords on anonymous accounts, which created a potential security risk. In IIS 6.0, you can use sub-authentication to manage passwords for anonymous accounts. To do so, your configuration must meet the following requirements:

- For applications for which you grant anonymous access, the worker process must run as LocalSystem.
- The sub-authentication component, Iissuba.dll, must be registered.
- The **AnonymousPasswordSync** metabase property must be enabled (set to **true**).

The actions taken to meet these requirements are different for clean installations of IIS 6.0 and upgrades to IIS 6.0 than they are for installations of IIS with sub-authentication configured. For more information about configuring sub-authentication, see Anonymous Authentication in IIS 6.0.

Additionally, when you encounter a 401.1 error, ensure that the account configured for anonymous authentication has not expired or is not locked out.

**Non-Anonymous authentication**
Initially, a Web browser issues an anonymous request. If Anonymous authentication is not enabled or if Anonymous authentication fails, IIS sends a list of supported authentication protocols as a group of **WWW-Authenticate** HTTP headers. At this point, the browser chooses the protocol to use, and then continues processing the request.

**Basic authentication**
Basic authentication is the simplest authentication protocol. Because of its simplicity, it is a good authentication method to enable temporarily when you are troubleshooting authentication problems. Use this authentication protocol if you want to quickly verify that a password is correct or that communication with a domain controller is working.

If the user name is a domain account, IIS contacts a domain controller to verify the credentials. If this fails, IIS returns a 401.2 error. If communication with the domain controller is successful, but the user name or password is invalid, IIS returns a 401.1 error.

Internet Explorer — not IIS — allows the user three connection attempts before rejecting the connection request and reporting an error to the user. When the Web server verifies that the user name and password correspond to a valid Windows Server 2003 user account, a connection is established.

**Integrated Windows authentication**
Integrated Windows authentication actually encompasses two different authentication protocols: Kerberos and Windows NT Challenge/Response, also known as NTLM. When this authentication protocol is enabled, IIS sends the following HTTP headers to the browser:

- **WWW-Authenticate: Negotiate**
- **WWW-Authenticate: NTLM**

The **Negotiate** header indicates that IIS can negotiate which of the two Windows authentication protocols to use. Internet Explorer 5, or later, recognizes this authentication header and determines which of the two authentication protocols to use. The **NTLM** authentication header is used as a fallback for browsers that do not support the **Negotiate** authentication header.

**Kerberos protocol authentication**
Authentication with the Kerberos protocol occurs as follows:

1. IIS requests authentication.
2. If the Windows user has not logged on to a domain, a dialog box requesting credentials appears in Internet Explorer, and then IIS contacts the Ticket Granting Server or Key Distribution Center (KDC) to request and receive a Ticket Granting Ticket. It then sends the Ticket Granting Ticket along with information about the server running IIS to the KDC.
    a. If the Windows user has already successfully logged on to the domain and received a Ticket Granting Ticket, IIS sends this ticket along with information about the IIS server to the KDC.
    b. The KDC issues the client a Resource Ticket.
    c. The client passes this ticket to the server running IIS.

    When troubleshooting Kerberos, ensure that the browser can make a trusted connection to the KDC by meeting the requirements listed in Microsoft Knowledge Base article 326985, How To: Troubleshoot Kerberos-related Issues in IIS.

Additionally, consider potential problems with the following:

- Service Principal Names. Kerberos uses a service for authentication. When the client passes the Ticket Granting Ticket and the information about IIS to the KDC, part of this information includes the server name from the URL that the browser is attempting to reach. By default, upon installation a systems NetBIOS name is registered in Active Directory® directory service. This is the name that a KDC uses to recognize IIS. It is called the Service Principal Name (SPN). If the browser is accessing the Web server by using some other host name — either the FQDN or another host header name — that name must be registered with Active Directory. You can do this by using the Setspn.exe command-line utility. For more information about Setspn.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click **Windows Support Tools**.

- Kerberos authentication. If you suspect that you are having a problem with the SPN, test Kerberos authentication by using the NetBIOS computer name. If the test is successful, you probably need to use the Setspn.exe command-line utility.

  To run Setspn.exe, at a command prompt, type the following:

  **Setspn -A HOST/***ServicePrincipalNameServiceAccount*

  In the above syntax, *ServicePrincipalName* is the new host name that you want to register and *ServiceAccount* is the identity under which the process that is hosting the service will be running. For computer accounts, like NetworkService or LocalSystem, use the NetBIOS name of the computer.

| **Note** |
| --- |
| The LocalService account, along with any other local accounts on an IIS server, cannot be authenticated by a KDC. Therefore, you cannot use the Kerberos protocol for authentication if the worker process identity is set to LocalService or to another account on the local computer. |

- Trust for delegation configuration. When troubleshooting Kerberos authentication, also ensure that the IIS computer is configured to be trusted for delegation in Active Directory. If the server running IIS is a member of the domain and is not a domain controller, then it must be trusted for delegation in order for Kerberos to work properly.

**To configure a computer to be trusted for delegation**

1. On the Domain controller, click **Start**, point to **Administrative Tools**, and click **Active Directory Users and Computers**.
2. Under the domain, double-click the **Computers** folder, right-click the server that is running IIS, and then click **Properties**.
3. Click the **General** tab.
4. Select the **Trusted for delegation** check box.
5. Click **OK**.

   The computer account is now trusted for delegation.

   - Security auditing. By enabling security auditing on the server running IIS, you can use the events that are logged to discover whether problems are occurring.

     For information about using the security log to troubleshoot Kerberos, see article 326985, How To: Troubleshoot Kerberos-related Issues in IIS.

**Windows NT Challenge/Response protocol**

Windows NT Challenge/Response (NTCR) protocol differs from Kerberos in that the server presents the HTTP client with a "challenge" and the client responds with its response. This way, the client's password is never sent over the network. Authentication with the NTCR protocol occurs as follows:

1. Typically, the client issues an initial anonymous request. When the anonymous request is rejected, IIS returns a 401.2

error and the **WWW-Authenticate** headers.

2. If the client fails or does not support Kerberos, the **Negotiate** and **NTLM** header values initiate an NTCR authentication exchange. The client closes the TCP connection, opens a new one, and sends a request that includes an **Authorization: NTLM** header. This header also includes encoded text that represents the users *UserName*, *ComputerName*, and *Domain*. This text is used by the Windows Security Support Provider Interface (SSPI) to generate the challenge. If the user account is not a local Windows account on the IIS server, the data is passed on to an appropriate domain controller, which then generates the challenge.

3. The challenge is sent to the client and IIS returns another 401.2 error.

4. The client uses its password and the challenge to create a mathematical hash. The client sends the hash back to the server in another **Authorization: NTLM** header.

5. The server accepts the response, and the local security provider or the appropriate domain controller recreates the same hash and compares the two. If they match, the user is successfully authenticated.

When troubleshooting NTCR authentication, note the following:

- Point-to-point connections. NTCR authentication requires point-to-point connection between the client and the server that is running IIS. Consequently, the NTRC protocol will not work with some European Organization for Nuclear Research (CERN) proxies and other Internet devices. As a result, NTCR is not suitable for use over the Internet. If a connection is broken during authentication, IIS returns a 401.2 error.

- HTTP keep-alives. A connection must be maintained through the entire authentication sequence. Therefore, HTTP keep-alives must be enabled on the server running IIS. If they are not enabled, IIS continues to return 401.2 errors until the browser fails the authentication sequence.

- Security privileges. The user account accessing the server must have the **Access this computer from the network** privilege. If this is not the case, IIS returns a 401.1 error.

- Domain controller communication. As with Basic authentication, if the credentials presented by the client are for a domain user account, the server running IIS needs to contact a domain controller.

**AuthPersistence Usage in IIS 6.0**

In earlier versions of IIS, the **AuthPersistence** metabase property had three possible settings. Two of the settings allowed administrators to enhance performance by specifying persistence based on the existence of a proxy server. Administrators could use either of those **AuthPersistence** settings to force IIS to negotiate one-time-per-client connections and then use those credentials for subsequent requests over the same connections. These two settings have been removed from IIS 6.0 for security reasons.

In IIS 6.0, the only valid setting for the **AuthPersistence** metabase property is **AuthPersistSingleRequest**, and NTLM is the only IIS 6.0 authentication protocol that honors this setting. The setting for **AuthPersistSingleRequest** is honored only in the following circumstances:

- Integrated Windows authentication is set to *NTLM*.

- Integrated Windows authentication is set to *Negotiate*, and NTLM authentication is used.

In either of these cases, **AuthPersistSingleRequest** is **False** — that is, not set — for backward compatibility with earlier versions of IIS. A value of **False** means that authentication persists for subsequent requests over the same connection.

In IIS 6.0, all other authentication protocols assume that the value of **AuthPersistSingleRequest** is **True** — that is, set — so authentication persists only for a single request over a connection. IIS 6.0 automatically resets authentication at the end of a request and forces each subsequent request over the same connection to authenticate.

**401.3-Unauthorized: Access Is Denied Due to an ACL Set on the Requested Resource**

A 401.3 error is a catchall "access denied" error. To troubleshoot this error, first check the authentication method that is

enabled on the requested URL. This will help you determine the security context under which the request is expected to execute: as an anonymous user or as an authenticated user. How you proceed depends on whether the user is authenticated or not.

Verify file and directory ACLs. When the requested URL is a directory, not a file, check the default document(s) that are configured for that URL. Sometimes, if the first document in the default document list is missing, the second document in the list might generate a 401.3 error.

After verifying the ACLs, you have a few options. You can enable security auditing for object access failures. This can be difficult, though, because you have to enable auditing on the individual objects (files, directories, and so on) that you want to monitor. Alternatively, you can use File Monitor and Registry Monitor to look for access denied problems in the registry or file system.

**Simplified UNC authentication model**

The Universal Naming Convention (UNC) authentication method, also called UNC Pass-through authentication, determines the credentials used for gaining access to a UNC share on a remote computer. Beginning with IIS 6.0, the UNC authentication method looks at the request user and the credentials stored in the **UNCUserName** and **UNCPassword** properties of the metabase to determine the credentials to pass through to the computer with the UNC share. This occurs in the following way:

1. If the **UNCUserName** is specified (not empty) and **UNCPassword** is valid, the requested user credentials are sent as the user identity for access to the remote share. If the **UNCUserName** is specified and the **UNCPassword** is not valid, a 500-Internal Server Error message is sent to the client. For information about troubleshooting this error, see HTTP 500.x -- Internal Server Error Codes.

2. If the **UNCUserName** is not specified, then the credentials of the request user (either an authenticated set of credentials for authenticated requests or IUSR_*ComputerName* credentials for anonymous requests) are sent as the user identity for access to the remote share.

| **Note** |
| --- |
| The **UNCAuthenticationPassthrough** metabase property is no longer used for UNC authentication. |

**Related Information**

For more Information about authentication and authorization, see the following articles in the Microsoft Knowledge Base:

- Article 264921, Info: How IIS Authentications Browser Clients
- Article 326985, How To: Troubleshoot Kerberos-related Issues in IIS
- Article 294382, Authentication May Fail with 401.3 Error If Host Header Differs from Servers NetBIOS Name
- Article 811889, How To: Troubleshoot the Cannot Generate SSPI Context Error Message
- Article 325358, How To: Configure User and Group Access on an Intranet in Windows 2000 or Windows NT 4.0

# HTTP 403.x-Forbidden

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Typically, 403 errors occur when an operation or request is disallowed because a requirement other than proper authentication credentials is not met. Commonly, this error occurs when a script request is made to a Web directory for which script access is not enabled. In such a case, first verify for which resource the request is being made. If the request is, indeed, for script code, such as ASP or ASP.NET, check the **execute permissions** for that directory in IIS Manager and ensure that at least **Script** is selected.

Table 11.7 HTTP 403 Substatus Codes lists the HTTP 403 error substatus codes.

**Table 11.7 HTTP 403 Substatus Codes**

| 403 Substatus Code | Description |
|---|---|
| None | Access is denied. |
| 1 | Execute access is denied. |
| 2 | Read access is denied. |
| 3 | Write access is denied. |
| 4 | SSL is required to view this resource. |
| 5 | SSL 128 is required to view this resource. |
| 6 | IP address of the client has been rejected. |
| 7 | SSL client certificate is required. |
| 8 | DNS name of the client is rejected. |
| 9 | Too many clients are trying to connect to the Web server. |
| 10 | Web server is configured to deny Execute access. |
| 11 | Password has been changed. |
| 12 | Client certificate is denied access by the server certificate mapper. |
| 13 | Client certificate has been revoked on the Web server. |
| 14 | Directory listing is denied on the Web server. |
| 15 | Client access licenses have exceeded limits on the Web server. |
| 16 | Client certificate is ill-formed or is not trusted by the Web server. |
| 17 | Client certificate has expired or is not yet valid. |
| 18 | Cannot execute requested URL in the current application pool. |
| 19 | Cannot execute CGIs for the client in this application pool. |

| | |
|---|---|
| 20 | Passport logon failed. |

# HTTP 404.x-File or Directory Not Found

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A HTTP 404 error indicates that the requested resource was not found. Table 11.8 HTTP 404 Substatus Codes lists the substatus codes for the 404 error. The descriptions for most substatus codes are self-explanatory. When additional information about a substatus code is required, it is provided in one of the following sections.

**Table 11.8 HTTP 404 Substatus Codes**

| 404 Substatus Code | Condition |
|---|---|
| None | File or directory not found. |
| 1 | Web site not accessible on the requested port. |
| 2 | Web service extension lockdown policy prevents this request. |
| 3 | MIME map policy prevents this request. |

Whether the requested resource is a file or a directory, begin by checking the following potential causes:

- Verify the file or directory. Examine the requested URL and check the physical path to which it maps.
- Check for unexpected host header routing. Problems with host headers occur when you have created more than one virtual Web server and rely on host header routing to serve content from the appropriate server. A 404 error results when you request a URL that exists on only one virtual server and, because of the configuration of the host header, the request is routed to the other virtual server.
- Verify that the request has been made on the correct port. The 404.1 error is returned when a request is made to a port on which the WWW service is not listening.
- Ensure that the requested file is not hidden. IIS sends a 404 error if the **Hidden** file attribute for a requested file is set.
- Validate (*ScriptMaps) Web service extension requests. 404 errors are sometimes returned when you have a *ScriptMap Web service extension and the Verify that file exists option is set to **true**. Because no file is actually associated with a *ScriptMap request, IIS cannot check for the file, so it returns a 404 error.

Both the Web service extension lockdown policy and **MimeMap** restriction features can cause a 404 error to be returned to the client making the request. If one of these features is the cause of an HTTP 404 error, do the following:

- Check the sc-substatus field for the request in the IIS W3C extended log. For information about enabling substatus logging, see Using Substatus and Win32 Errors in W3C Extended Logging.
- Check the Win32 error in the IIS W3C extended log. By default, IIS logs all Win32 errors associated with the request that are returned by the underlying operating system.
- Enable a modified custom error for the HTTP 404 error that you suspect. You can configure IIS to return a response to the client that contains the substatus code by editing the custom error page for the 404 error, and placing some custom text into the file. The custom error pages are located in *systemroot*\Help\IisHelp\Common.

| **Important** |
|---|
| Enabling a custom error that provides substatus information undermines the security strategy used by IIS and should be done temporarily and only for troubleshooting purposes. |

**404.1-Web Site Not Accessible on the Requested Port**
The 404.1 error can occur only on computers with multiple IP addresses. If a specific IP address/port combination receives a client request, and the IP address is not configured to listen on that particular port, IIS returns a 404.1 HTTP error.

For example, if a computer has two IP addresses and only one of those IP addresses is configured to listen on port 80, any requests received on the other IP address with port 80 causes IIS to return a 404.1 error. This error should be set only at the service level because it is returned to clients only when multiple IP addresses are used on the server.

**404.2-Lockdown Policy Prevents This Request**
If a request is denied because the associated ISAPI or CGI has not been unlocked, a 404.2 error is returned. When substatus logging is enabled, if you look at this request in the IIS logs, you will see an entry similar to the following:

```
2002-11-25 05:46:15 127.0.0.1 GET /default.asp - 80 - 127.0.0.1 - 404 2 1260
```

If IIS logging is not configured to log the substatus code, you can check the Win32 error to verify this condition. When a 404.2 error occurs, IIS logs the Win32 error 1260, which is an ERROR_ACCESS_DISABLED_BY_POLICY error.

**404.3-MIME Map Policy Prevents This Request**
If a request is denied because a MIME map restriction is in effect, a 404.3 error is returned. When substatus logging is enabled, if you look at this request in the IIS logs, you will see an entry similar to the following:

```
2002-11-25 05:46:27 127.0.0.1 GET /somefile.unkext - 80 - 127.0.0.1 - 404 3 50
```

As with a 404.2 error, a 404.3 error causes a Win32 error to be logged. In this case, the error is 50, which is an ERROR_NOT_SUPPORTED error.

Another quick way to verify that a MIME map restriction is in effect is to add a wildcard MIME type to the virtual directory in question. This enables all MIME types to be served without restriction. Then, repeat the failed HTTP request and verify that it is served successfully.

| **Note** |
| --- |
| Adding a wildcard MIME type to the virtual directory is a troubleshooting step, not a solution to the problem. Leaving the wildcard in place and allowing all MIME types to be served compromises the security of the server. |

For more information about adding a wildcard MIME type, see Working with MIME Types.

**405-HTTP Verb Used to Access This Page Is Not Allowed**
This HTTP code is returned when the client makes an HTTP request that contains a verb that is not allowed. This condition can occur when:

- A request for static content contains verbs other than GET or HEAD, and the request is made to a URL that did not end with a /. Instead performing a courtesy redirect, IIS sends the 405 error.
- An HTTP request for an ISAPI application contains a verb not listed in the **ScriptMaps**configuration for that ISAPI.

**407-Initial Proxy Authentication Required by the Web Server**
This error indicates that an intermediary proxy server between the HTTP client and the Web server requires some form of authentication. How you troubleshoot this kind of error depends upon the proxy server itself. Generally speaking, running a network trace with Network Monitor is helpful. If the Web client is a custom client, ask its developer to ensure that it is handling security appropriately.

**413-Request Entity Is Too Large**
For security reasons, you can limit the size of the *entity-body* of an HTTP request by modifying the **MaxRequestEntityAllowed** metabase property. When an entity-body of a client request exceeds the size that is specified in the **MaxRequestEntityAllowed** property, IIS returns a 413 error. If this error is logged for an individual request, an application on the Web server might have encountered an unexpected event and generated a request that is too large. If this error is logged for many requests, a malicious user might be attempting to compromise your Web server.

**414-Request URL Is Too Large and Therefore Unacceptable on the Web Server**
Just as the entity body of a request can be too large for IIS to process, a URL can be too long for IIS to process. IIS returns a 414 error if this occurs.

# HTTP 5xx Server Error Codes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Error codes in the 5*xx* range indicate server problems. Table 11.9 HTTP 500-Class Client Error Codes Returned by IIS lists the HTTP 5*xx*-class codes that IIS returns. The descriptions for most codes are self-explanatory. When additional information about a code is required, it is provided in one of the following sections.

**Table 11.9 HTTP 500-Class Client Error Codes Returned by IIS**

| Status Code | Condition |
|---|---|
| 500.*x* | Internal server error. |
| 501 | Header values specify a configuration that is not implemented. |
| 502 | Bad gateway. |
| 503 | Service unavailable. |
| 504 | Gateway timeout. |
| 505 | HTTP version not supported. |

# HTTP 500.x -- Internal Server Error Codes

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table 11.10 HTTP 500 Substatus Codes lists the most common 500 substatus codes returned by IIS. The descriptions for most substatus codes are self-explanatory. When additional information about a substatus code is required, it is provided in one of the following sections.

**Table 11.10 HTTP 500 Substatus Codes**

| 500 Substatus Code | Condition |
|---|---|
| 11 | Application is shutting down on the Web server. |
| 12 | Application is busy restarting on the Web server. |
| 13 | Web server is too busy. |
| 14 | Invalid application configuration on the server. |
| 15 | Direct requests for Global.asa are not allowed. |
| 16 | UNC authorization credentials are incorrect. |
| 17 | URL authorization store cannot be found. |
| 18 | URL authorization store cannot be opened. |
| 100 | Internal ASP error. |

**500.11 and 500.12-Application State Issues**

Users rarely encounter 500.11 and 500.12 errors in their browser. These errors indicate that the ASP application to which the requested ASP page belongs is in the process of shutting down, restarting, or both. Restarts and shutdowns can be triggered directly from code when the built-in APIs available to ASP developers are used. Occasionally, these errors occur when the Global.asa file for the application (if one exists) is updated or altered in such a way that a file change notification is issued. Because IIS monitors for file change notifications on ASP content, IIS receives this notification, and because the file is the Global.asa file, IIS restarts the application.

If users often encounter these errors in their browser and you are certain that the code in the Global.asa file is not being updated, virus scanning and backup software, which sometimes sweep the ASP content and trigger a notification, might be causing these errors. For more information, see Microsoft Knowledge Base article 248013, Err Msg: HTTP Error 500-12 Application Restarting.

**500.13-Web Server Too Busy**

ASP requests are executed by threads from a pool of threads whose size is limited to 25 threads per processor, by default. If more ASP requests arrive than there are threads available to execute them, ASP places the extra requests in a queue, where they wait until a thread becomes available. By default, this queue is limited to 3000 requests. If the queue becomes full, the next ASP request, instead of being queued, causes a 500.13-Web server too busy error. The number of requests that can wait in the queue is controlled by the **AspRequestQueueMax** metabase property. Change the **AspRequestQueueMax** property to produce the desired user experience — a smaller queue allows users to more quickly see the 500.13 error when ASP requests are backed up.

If users encounter this error, select the **Active Server Pages/Requests Queued** performance counter check box. If the queue fills up, but ultimately empties again, or if the queue remains partially filled but requests move through very quickly, there probably is not a problem, just a spike in traffic. However, if the queue remains partially or completely full for a sustained period and requests do not appear to be moving through quickly enough, there might be a problem.

Diagnosing a problem of this type can be a complex process. For information about diagnosing such problems, see Advanced IIS Troubleshooting with Tracing and Debugging.

**500.16-UNC Authorization Credentials Are Incorrect**

Problems with content stored on remote file servers can be more complex than problems with locally stored content. The 500.16 error indicates that there is a specific configuration problem with remotely stored content. When IIS is configured to use specific credentials to access the UNC path in the metabase, it uses those credentials to open the content on the remote file server. To do this, IIS uses a Windows logon API to obtain a security token that it can use to impersonate a security identity when accessing the remotely stored content. If the call to the Windows logon API fails, indicating that there is a problem with the user name or password, IIS returns an HTTP 500.16 error.

**500.17 and 500.18-IIS URL Authorization Issues**

IIS URL Authorization is a security feature that is new in IIS 6.0. Among other administrative tasks, IIS URL Authorization configures the Authorization Manager store by setting the **AzStoreName** metabase property. This store can be an XML file or an LDAP location. It is configured by setting the **AzStoreName** property to an appropriate string that represents the path to a store. For example: **msxml://c:\MyStore.xml** or **msldap://CN=MyAzStore, CN=Program Data, DC=MyDomain, DC=com**.

The 500.17 and 500.18 errors pertain to problems accessing the Authorization Manager store. A 500.17 error indicates that the store could not be found. To begin troubleshooting one of these errors, try to verify the path specified in the string by either opening the XML file or testing the LDAP path with a simple client, such as LDP.exe. For more information about LDP.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click **Windows Support Tools**.

A 500.18 error typically indicates that the store was found, but could not be opened. The error description might contain information such as an HRESULT and description. However, this error generally occurs because there is an ACL–related problem. The security context that is used to open the store is that of the worker process identity. To troubleshoot, verify that the configured worker process identity is a member of the built-in IIS_WPG group and that this group has access to the store. File Monitor is a good utility to use in such a case.

For more information about IIS URL Authorization, see URL Authorization in IIS 6.0.

# 503-Service Unavailable

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

HTTP 503 errors are returned more often in IIS 6.0 than in earlier versions of IIS. Typically, the HTTP 503-Service Unavailable error is returned directly from HTTP.sys when there is a problem with getting the request up to the worker process. However, this error can occur for a variety of reasons, as listed in Table 11.11 HTTP 503 Errors from HTTP.sys.

**Table 11.11 HTTP 503 Errors from HTTP.sys**

| Reason String | Text in HTTP.sys Error Log | Cause(s) of Error |
|---|---|---|
| Service unavailable. | N/A | • IIS cannot start any new worker processes because of limited system resources or because starting a new worker process would exceed the **DemandStartThreshold**property.<br>• Bandwidth throttling is enabled, but the filter addition fails.<br>• The control channel or internal configuration group for the URL is inactive.<br>• The send for a request that was serviced from the cache failed (typically under low memory conditions). |
| Too many users. | ConnLimit | The connection limit for the site or server as a whole has been reached. |
| Application taken offline. | AppOffline | The application pool has been put into Rapid Fail Protection and has been disabled automatically. |
| Application request queue full. | QueueFull | The application pools request queue has been exceeded and the request cannot be queued. |
| Administrator has taken application offline. | Disabled | The administrator has stopped the application pool. |
| Application automatically shut down due to administrator policy. | AppShutdown | The application pool has been shut down because it exceeded its CPU usage limit. |
| Request timed out in app pool queue. | AppPoolTimer | The application pool is too busy to handle the request. The request has, therefore, timed out in the queue and has been returned with a 503 error. |

For security purposes, the description that is sent back to the client (Service Unavailable) does not fully describe the error. HTTP.sys keeps its own error log. Requests that generate 503 errors never pass into user mode; therefore, the HTTP error log is the only place these problems are documented. To troubleshoot the cause of the 503 error, use the following procedure.

**To check for possible causes of HTTP 503 errors**

1. Check the HTTP.sys error log for information about the error being returned. See Table 11.11 HTTP 503 Errors from HTTP.sys for error descriptions.
2. Check the event logs. Typically, the event logs contain more detailed information about the cause of a 503 error.
3. Verify that health/recycling issues are not causing the error.
4. Check for server binding conflicts.
5. Check the IP inclusion list for errors.

# Other HTTP.sys Error Log Errors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In some cases, HTTP.sys generates entries in its error log that are not associated with a specific HTTP request and that do not have HTTP error codes. Table 11.12 Other HTTP.sys Errors lists these errors.

**Table 11.12 Other HTTP.sys Errors**

| Logged Reason Code | Condition |
|---|---|
| Timer_ConnectionIdle | The connection was reset as its idle time exceeded the **ConnectionTimeout** property. |
| Timer_HeaderWait | The headers for the request were not received within the time limit specified by the **HeaderWaitTimeout** property. |
| Timer_MinBytesPerSecond | The throughput rate of this connection has dropped below the limit specified in the **MinFileBytesPerSec** property. |
| Timer_EntityBody | The entity body was not received within the specified timeout. |
| Timer_Response | The connection was reset because the response was not processed within the specified time limit. |
| Timer_AppPool | The connection was reset because the request stayed in the request queue of the application pool longer than the specified time limit. |

For complete information about the errors logged in the HTTP error log, see Error Logging in the HTTP API. For more information about available configuration options in HTTP.sys, see Microsoft Knowledge Base article 820129, Inf: Http.sys Registry Settings for IIS.

# Troubleshooting IIS Configuration Problems

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Configuration problems present their own unique troubleshooting challenges. This section provides information about troubleshooting common configuration problems that occur with the integrity of XML in the IIS metabase and with UNC content.

# Preserving the Integrity of XML in the IIS Metabase

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

With the introduction of the XML-based IIS configuration store, administrators now have the ability to directly edit XML in a text editor. This feature is called edit-while-running. Because XML has very strict rules about syntax, the XML store needs to be well-formed at all times. Manually editing the XML store with a text editor can introduce syntax errors in the MetaBase.xml file. Therefore, when manually editing the MetaBase.xml file with a text editor, observe the following guidelines:

- If a syntax error in the XML file creates a problem, use the metabase backup/restore feature to restore the metabase from the backup files.
- Open the MetaBase.xml file in Internet Explorer, which has a built-in XML parser. If a syntax error occurs, Internet Explorer displays it and notes its location.

For information about the metabase edit-while-running feature of IIS, see Enabling Edit-While-Running in IIS 6.0. For information about the metabase backup/restore feature, see Restoring the Metabase from Backup Files.

# Troubleshooting Problems with UNC Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Traditionally, situations that involve Universal Naming Convention (UNC) content have experienced configuration and performance problems related to security, user delegation, and file caching. Improvements in IIS 6.0 affect how these problems are resolved. Figure 11.3 IIS 6.0 UNC Authentication illustrates the logic IIS 6.0 uses to manage and serve content from a UNC file share.



# Accessing and Executing Remote Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When IIS accesses content stored on a remote file server, user delegation becomes a factor because IIS uses the clients credentials to access content on the file server. In effect, the Web client passes the users credentials to IIS, which then delegates them to the file server for the authorization check. In IIS 6.0, user delegation is easier because, by default, IIS uses the credentials of the request user (the anonymous user identity for anonymous requests and the authenticated users identity for authenticated requests) to access the content on the remote file server. So, by default, IIS does *pass-through authentication*, which means that IIS passes the user credentials through to the remote file server. If you want IIS to use a specific user when accessing remote content, set the **UNCUsername** and **UNCPassword** properties in the IIS metabase. For information about setting these properties, see the UNCUserName Metabase Property and the UNCPassword Metabase Property.

**Logging On**

In Figure 11.3 IIS 6.0 UNC Authentication, the request user is the identity of the request coming into the Web server, such as a GET request for http://www.contoso.com/staticfile.txt. If anonymous access is enabled, this request uses the identity of the anonymous user as specified in the metabase (typically, IUSR_*ComputerName*). If anonymous access is disabled and a valid authenticated request comes in, IIS uses that users account to log on. However, in the case of UNC content, at this stage, IIS checks if **UNCUserName** is specified in the metabase. If so, IIS logs on as that user. If **UNCUserName** is not specified in the metabase or if it is set to an empty string, then IIS logs on as the user who originally made the request (that is, as an anonymous user or as an authenticated user).

| **Note** |
| --- |
| The **UNCPassthroughAuthentication** metabase property has been removed from IIS 6.0. In IIS 6.0, pass-through authentication is enabled by default. To disable pass-through authentication, use the **UNCUserName** and **UNCPassword** properties. |

**Accessing the Remote Server**

After IIS has a user identity, it must access the remote file server. The type of authentication (Anonymous, Basic, NTLM, Kerberos) enabled on the server can significantly affect how successful this step is. In almost all cases, Basic authentication enables IIS to access a remote file server (although you must always use SSL with Basic authentication since user names and passwords are sent in plain text). Using Basic authentication can help you verify that everything else is working properly.

**Executing the Remote Content**

So far, we have discussed how remote content is *accessed*. Remote content must also be *executed* in the context of this user account. In other words, if **UNCUserName** and **UNCUserPassword** are set, then all remote content will be accessed and executed in the context of **UNCUserName**. This might not be desirable for most applications, because with this approach you lose the identity of the originating Internet/Intranet user. If pass-through authentication is being used, then the request is processed under the identity of the request user.

The process illustrated in Figure 11.3 IIS 6.0 UNC Authentication applies to all content that is executed or retrieved by the server. Consequently, static content, such as Graphics Interchange Format (GIF) and Hypertext Markup Language(HTML) pages, Script-Mapped files, CGI executables, and ISAPI Extension DLLs (like ASP, ASP.Net, and so on) runs under this user context. (ISAPI filters run under the context of the worker process identity. By default, the worker process runs as NetworkService, which is a local account.)

**Working in a Domain Environment**

In most cases, you will want to run your Web servers and file server in a domain environment because it is much simpler to use a domain environment in which you can control one set of users and groups from the domain controller. Furthermore, if you want to use Windows Integrated Authentication in this scenario, then you need to use a domain environment to take advantage of the new constrained delegation feature of Windows Server 2003. Instructions on setting up constrained delegation follow. For more information about constrained delegation, see Delegating authentication in Help and Support Center for Windows Server 2003.

**Working in a Workgroup Environment**

Web servers and a file server used in a workgroup environment are difficult to manage and maintain. In a workgroup environment, because one user account has access to another computer only if that computer has a user with the same user name and password, you have added administrative overhead. To add a new user to your environment, you have to add the

user to both servers and you need to ensure that both accounts have the same password.

For an detailed explanation of using remote content on a UNC share, see Deploying and Configuring Internet Information Services (IIS) 6.0 with Remotely Stored Content on UNC Servers and NAS Devices.

# Troubleshooting File Caching Problems

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

There are two causes for problems with file caching of UNC content:

- Vendor's implementation and support for file change notifications. Usually non-Windows-powered Network Attached Storage (NAS) vendors provide fixes for known problems; however, as a workaround, caching is typically disabled to prevent those problems.

- The number of simultaneous distinct file structures that the ASP and static file caches can monitor for changes. Because each monitor takes up a work item indefinitely from the Server Message Block (SMB) connection with the file server, IIS can exhaust a system's SMB resources by trying to cache a wide and disparate set of directories.

To address these problems, several changes have been made to the way IIS caches content. These changes improve performance and reliability in most cases; however, these changes can have unexpected effects on UNC stored content.

**ASP and Static File Caching Changes**

IIS 6.0 has a new caching algorithm for the ASP script and static file cache when the content is being stored on a UNC share. In earlier versions, IIS registered for a file change notification from the operating system for each directory structure so that it would be notified when a file changed. In IIS 6.0, the Web server simply checks the Last-Modified date on the cached file instead of requesting these notifications. If the Last-Modified date has not changed, the cached version of the file is sent out. If the file is new, the cache is updated with the new version of the file and is sent to the client.

To further improve performance, the algorithm has a staleness aspect that causes IIS to check the Last-Modified date only if more than five seconds have passed since the last check. Otherwise, IIS assumes that the file has not changed, and serves the cached version. For more information about this algorithm, see Troubleshooting Performance Problems with UNC Content.

| **Note** |
| --- |
| The server-side includes file handler (ssinc.dll) makes use of the user-mode static-file cache, so the configuration you apply to the user-mode static-file cache also applies to .stm, .shtml, and any other files you have mapped to ssinc.dll. |

**Troubleshooting Problems with Stale Content**

When stale content is being served by the Web Server even though the content has been updated on the remote file server, disable the Web Server cache as the first step in isolating the problem.

| **Note** |
| --- |
| First ensure that your results are not being affected by the 5-second staleness factor. |

For information about disabling caching of ASP files and static files, see Configuring ASP Template Caching and Configuring Last-Modified Caching.

# Troubleshooting Performance Problems with UNC Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Using the Last-Modified time to check for changes (instead of monitoring each virtual directory for change notifications) permits greater scalability when many virtual directories are mapped to UNC shares. Checking the Last-Modified time on the remote file server every time IIS wants to send a file is relatively expensive in terms of processing overhead. IIS checks the Last-Modified time on a file every five seconds by default; otherwise, it assumes that nothing has changed and serves the content from the cache, even though it might be stale.

For some sites, a five-second window will be too large a window, and for others, too small. The sampling interval is configurable with the following registry keys, which are not present by default:

- **Static pages** Set the registry property, **FileAttributeCheckThreshold**, a DWORD value, located at **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Inetinfo\Parameters**.

- **ASP Scripts** Set the registry property, **FileMonitoringTimeoutSeconds**, a DWORD value, located at **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ASP\Parameters**.

Using the last-modified cache update mechanism will result in more reliable and secure delivery of content. Additionally, it will scale well on a wide application structure (where your Web server has lots of applications pointing at thousands of different directories). There is, however, a performance penalty that might be noticeable in regards to static content.

Even with these new file-caching algorithms, overall performance of UNC-based content can still be slightly slower than that of local content. However, depending on the content size, the difference is usually negligible.

If you think that the performance of your server has diminished due to these changes, you can enable change notification–based caching by adding the following registry keys, which are not present by default:

- **Static Pages** Set the registry property, **DoDirMonitoringForUNC**, a DWORD value, located at **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Inetinfo\Parameters**.

- **ASP Scripts** Set the registry property, **EnableChangeNotificationForUNC**, a DWORD value, located at **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ASP\Parameters**.

These entries do not exist in the registry by default. To add it to the registry, use the registry editor, Regedit.exe.

However, note that this option is not always recommended because it can cause the operating system to suffer from the scalability problems that were present in earlier version of IIS. Additionally, share-level permissions will no longer be honored by the Web server (NTFS ACLs are always honored). Consider modifying this registry entry only if you have a small directory structure, are not concerned about share-level permissions, and consider performance the most important metric.

When deciding which caching mechanism to use, consider the following tradeoffs. The most reliable and secure method is using the default: Last-Modified-based caching. It is reliable because it does not rely on the change-notification functionality of the file server and it scales well on a wide application structure (where your Web server has lots of applications pointing at thousands of different directories). It is more secure because share-level permission changes are honored (the cache accesses the file often, which automatically triggers the authorization check against the current file system ACL and the current share ACL).

However, if you are using a file server that reliably reports change notifications *and* if you have few sites or virtual directories, using change-notification is significantly faster. For example, if you have a single Web site with a couple of applications and all of your content is stored on a Windows Server 2003 file server, it is recommended that you use the change-notification method of caching.

**Related Information**

- For more information about using remote content on a UNC share, see Deploying and Configuring Internet Information Services (IIS) 6.0 with Remotely Stored Content on UNC Servers and NAS Devices.

- For more information about the FileAttributeCheckThreshold registry key, see Global Registry Entries.

- For more information about the FileMonitoringTimeoutSeconds and EnableChangeNotificationForUNC registry keys, see ASP Registry Entries.

# Troubleshooting IIS Manager and UNC Content Problems

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because of the way that IIS Manager displays the virtual directory, managing Web content that is stored on remote UNC shares can be confusing. IIS Manager uses the credentials of the logged-in user (the user actually running IIS Manager), who might not necessarily have the right credentials to access the remote file share. In this case, IIS Manager displays a stop-sign icon for the virtual directory. If this occurs, test UNC access by making requests with a Web browser.

You might want to create virtual directories that reference UNC-mapped drives. IIS does not support mapped drives because they are user-specific. For more information, see Microsoft Knowledge Base article 257174, Using Mapped Drives with IIS.

# Troubleshooting Errors That Occur When UNC Content Is Under High Load

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If the Web server and the remote file server are not configured properly for high stress conditions, errors might occur. The following configuration errors might be reported in the browser, in the event log on the IIS server or remote file server, or the IIS logs:

- `Windows cannot find the network path. Verify that the network path is correct and the destination computer is not busy or turned off. If Windows still cannot find the network path, contact your network administrator. Error code: 0x80070033`
- `The network location cannot be reached. For information about network troubleshooting, see Windows Help.Error code: 0x800704CF`
- `Not enough server storage is available to process this command. Error code: 0x8007046`

**Modifying the MaxUserPort Registry Entry**

The first two errors are returned when the server running IIS has a shortage of outbound TCP ports when it attempts to make the UNC connection. In this case, the Web server returns HTTP 500 errors for all requests until TCP ports are available again. To fix this problem, you can modify the MaxUserPort registry entry on the server running IIS to specify an increase in the number of available ports. This entry does not exist in the registry by default. To add it to the registry, use the registry editor, Regedit.exe.

| **Caution** |
| --- |
| Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference in the Microsoft Windows Server 2003 Deployment Kit. |

For example, if you have many concurrent authenticated users, a different SMB connection is opened from the computer running IIS to the remote file server for each individual user. These SMB connections use ports. By default, you are limited to ports from 1024 to 5000 for outbound TCP connections (slightly fewer than 4,000 ports are available), so you might need to increase this value to a higher number, such as 10000 or 20000, to provide additional ports for SMB connections. Be aware that this might cause problems with applications that create sockets on these higher ports. For more information, see Microsoft Windows 2000 TCP/IP Implementation Details.

**Checking the File Server Setting**

If you receive the `Not enough server storage is available to process this command. Error code: 0x8007046A` message, check your file server settings. This IIS error usually accompanies event log messages like the following:

- `The server was unable to allocate from the system nonpaged pool because the pool was empty.`
- `The server was unable to allocate a work item <n> times in the last 60 seconds.`

Normally, these messages indicate that your file server is set to use too much nonpaged system memory. This can be quickly verified by opening Task Manager, clicking the **Performance** tab and checking the **Nonpaged** value in the **Kernel Memory** frame. Limit the amount of kernel memory that the Server service can use by configuring the SMB settings as discussed later in this section.

When troubleshooting this kind of error, check the following performance counters on the computer running the Web server:

- Server\Files Open
- Server\Server Sessions

- Server\Work Item Shortages

**Modifying the Work Item Configuration**

When two computers communicate, one SMB connection is used for each user. However, within an SMB connection, there are a finite number of work items. Work items can be consumed in a variety of ways and for varying amounts of time. For example, performing an operation against a file (by using **CreateFile**, **GetFileAttributes**, or other similar Windows APIs) takes up an I/O work item for a short period of time, but asking for a change notification on a directory structure takes up a work item for as long as the connection is maintained.

`The network location cannot be reached. For information about network troubleshooting, see Windows Help.Error code: 0x800704CF and Not enough server storage is available to process this command. Error code: 0x8007046` errors usually can be fixed by changing the following registry entries, which are related to SMB configuration: MaxCmds, MaxMpxCt, and MaxWorkItems.

On the remote file server, you can update these registry entries in the following location: HKLM\System\CurrentControlSet\Services\LanmanServer\Parameters.

On the server running IIS, you can update these registry entries in the following location: HKLM\System\CurrentControlSet\Services\LanmanWorkstation\Parameters.

The MaxCmds, MaxMpxCt, and MaxWorkItems registry entries do not exist in the registry by default. To add them to the registry, use the registry editor, Regedit.exe.

| **Note** |
| --- |
| For these registry values to take effect, you must stop and start the Server service on the remote file server and the Workstation service on the server running IIS. You do not need to restart the computers. |

Before changing the MaxCmds, MaxMpxCt, and MaxWorkItems registry entries to troubleshoot work item shortages and system memory consumption, you need to understand how work items are consumed. Work item consumption occurs as follows:

- Only one SMB connection is created per Web server/file server connection per user.
- Each change notification instance uses a work item until the connection is ended.
- Each request from the Web server to the file server uses a work item temporarily.
- Typically, each work item takes up about 20K of non-paged pool memory. Although certain Microsoft Knowledge Base articles specify setting the value of the MaxWorkItems registry entry to 16000, doing so consumes approximately 320 MB of non-paged pool memory, which could overwhelm system resources. To increase the reliability of your system, always set this number lower than this recommended value.
- To calculate the value of the MaxCmds registry entry, use the following formula:

  (the number of distinct physical directories that IIS needs to monitor for change notifications) * (1 (if static files exist) + 1 (if ASP content exists) + 1 (if ASP.Net content exists) ) + 50 (for concurrent default/regular file IO)

| **Note** |
| --- |
| This formula approximates a recommended value. Depending on the amount of concurrent traffic, you might need to increase your values. |

Table 11.13 Values for SMB Settings lists the default and maximum values for the relevant SMB settings on Windows Server 2003. These SMB settings also apply to the Microsoft Windows® 2000 Server operating system.

**Table 11.13 Values for SMB Settings**

| Registry Entry | Default Value | Maximum Value |
| --- | --- | --- |

| | | |
|---|---|---|
| MaxCmds | 50 | 65535 |
| MaxMpxCt | 50 | 65535 |
| MaxWorkItems | Scaled on memory size, usually between 1KB and 4KB. | 65535 |

# Troubleshooting Logging Problems

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Historically, IIS administrators have rarely had to troubleshoot logging problems. In earlier versions of IIS, log configuration was simple and straightforward, and in most cases, the default settings were adequate.

In IIS 6.0, new logging functionality, such as centralized binary logging and UTF-8 logging, and the fact that logging is now performed by HTTP.sys instead of the WWW Service (W3SVC) present new configuration options and architectural elements as well as new troubleshooting challenges. This section explores some of the problems that can occur with HTTP.sys logging.

Table 11.14 HTTP.sys Logging-Related Event Log Errors lists and describes HTTP.sys logging-related Event Log errors, and provides troubleshooting recommendations for each.

**Table 11.14 HTTP.sys Logging-Related Event Log Errors**

| Event ID | Description | Recommendation |
|---|---|---|
| 15000 | Unable to create log file *%2*. | Ensure that the logging directory is correct and this computer account has write access to that directory. |
| 15001 | Unable to create the log file for site W3SVC*%2*. | Ensure that the logging directory for the site is correct and this computer has write access to that directory. |
| 15002 | Unable to write to the log file *%2* for site W3SVC*%3*. The disk might be full. | If this is a network path, ensure that network connectivity is not broken. Otherwise, the disk is full. |
| 15003 | Unable to create the centralized binary log file. | Ensure that the logging directory is correct and this computer has write access to that directory. |
| 15004 | Unable to write to the centralized binary log file *%2*. The disk might be full. . | If this is a network path, make sure that network connectivity is not broken. Otherwise, the disk is full. |
| 15006 | Owner of the log file or directory *%2* is invalid. Another user might have already created the log file or the directory. | Check for the existence of the logging directory and ensure that the owner is the local Administrators group. |

**Logging Security Configuration**

One of the fundamental aspects of IIS logging design involves the expected security configuration of the log file directory or UNC share when logging to UNC.

**Configuration process**

Upon first receiving an HTTP request, HTTP.sys begins logging. On this initial request, if the appropriate logging subdirectory and file does not exist, HTTP.sys creates them, giving System and local Administrators full control. For each, the owner is the Local Administrators account. If the directory and file already exist, HTTP.sys performs an access check for security purposes and verifies that the owner of the file and directory is the local Administrators group. If not, HTTP.sys logs event ID 15006 to the event log and disables logging for that site. If the ownership check passes, HTTP.sys then ensures that System and local Administrators have full control.

When configuring logging to a UNC share, HTTP.sys does not perform an ownership check. Instead, the administrator is responsible for ensuring that security on the log file share and directory is configured appropriately. HTTP.sys sets the ACLs on the remote share to allow the Web servers system account the appropriate access.

**Event ID 15000, 15001, and 15003**

These errors occur when the log file and directory do not have proper security settings. When the IIS logging configuration settings are changed, if the file and directory security settings are not correct (if the System and Administrators do not have full control), HTTP.sys writes event ID 15001 or event ID 15003 (if binary logging is enabled) to the Event Log. HTTP.sys writes an event ID 15000 to the event log if security settings change during operation, for example, if the file ACLs are altered manually.

**Event ID 15002 and Event ID 15004**

Event ID 15002 and Event ID 15004 errors are similar. Both occur when HTTP.sys is unable to write to the log file because the

disk is full. When binary logging is enabled, HTTP.sys writes an event ID 15004 error to the Event Log. An event ID 15004 also is written to the Event Log if a network problem prevents HTTP.sys from writing to the log file.

**Logging Headers**
In the past, the logging headers that were written to a log file typically indicated the restarting of the Web service. In IIS 6.0 this is not the case. After waiting for 15 minutes for a given site, HTTP.sys stops logging. When the next request for that site arrives, HTTP.sys restarts logging and writes a new set of headers to the log file.

**Present in Log Files**
In rare cases, you might see in an IIS log file entry. When a query string or a header value is longer than 4 KB, HTTP.sys writes the string as instead of writing the entire string.

**UTF-8 Logging Problems**
When using the UTF-8 logging feature, note the following:

- A log file logged in UTF-8 does not contain a Byte Order Mark (BOM). File editors use this mark to identify text as UTF-8 text. Therefore, if you attempt to open a log file that is logged in UTF-8 in Notepad by double-clicking the file or by using the **Open With** option, the file might not display correctly. To open the file in a way that displays it correctly, use the **Open** command on the **File** menu and then select **UTF-8** in the **Encoding** box.

- Because IIS 6.0 now supports UTF-8 URLs, you can now log those URL requests to an ASCII log file. UTF-8 is a double-byte character set standard. Because ASCII is a single-byte character set standard, logging UTF-8 information to an ASCII file presents a problem. In such a case, ? is logged for the characters that cannot be converted to the codepage of the server.

# Troubleshooting HTTP.sys Problems

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

HTTP.sys can provide information through three mechanisms: HTTP status codes, Windows Event Logs, and the HTTP error log. Preceding sections discussed the HTTP status codes and several event log errors that are related to logging. This section addresses the initial startup process of IP addresses binding and the mechanism other services use to ask HTTP to listen to specific URL endpoints through the HTTP API.

# Configuring the IP Inclusion List

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The HTTP.sys IP inclusion list allows administrators to limit the IP addresses to which HTTP.sys automatically binds when it starts. This provides functionality similar to the functionality provided by the **DisableSocketPooling** property that was available in IIS 5.0. At startup, HTTP.sys checks to determine if there are any entries in the IP inclusion list. If there are, HTTP.sys attempts to listen only on these IP addresses and only at initial startup. If the list has no entries, HTTP.sys listens on all available IP addresses.

To configure the IP inclusion list, use the HttpCnfg.exe utility. For information about HttpCfg.exe, see Httpcfg.exe: HTTP Configuration Utility.

Table 11.15 HTTP.sys IP Binding Event Log Errors lists and describes the HTTP.sys IP binding event log errors. Only the errors that require additional explanation are addressed below.

**Table 11.15 HTTP.sys IP Binding Event Log Errors**

| Event ID | Description |
|---|---|
| 15005 | Unable to bind to the underlying transport for *%2*. The IP Listen-Only list might contain a reference to an interface that might not exist on this computer. The data field contains the error number. |
| 15007 | Reservation for the namespace identified by URL prefix *%2* was successfully added. |
| 15008 | Reservation for namespace identified by URL prefix *%2* was successfully deleted. |
| 15009 | An error occurred while initializing namespace reservations. The error status code is contained within the returned data. |
| 15010 | An error occurred while initializing a namespace reservation identified by URL prefix *%2*. The error status code is contained within the returned data. |
| 15014 | Unable to convert IP Listen-Only list entry *%2*. The data field contains the error number. |
| 15015 | Unable to convert all entries on IP Listen-Only list. The driver will listen on all available interfaces. |

**Event ID 15014**

If the IP inclusion list contains entries, HTTP.sys iterates through the list and, for each entry, it first checks to ensure that the IP address is valid and is assigned to one of the systems network interfaces. If this check fails, HTTP.sys logs Event ID 15014 to the event log.

**Event ID 15015**

This error is logged to the event log when HTTP.sys parses the IP inclusion list and finds that all of the entries in the list are invalid. If this happens, as the description in Table 11.15 notes, HTTP.sys listens to all IP addresses.

Services that use the HTTP listening functionality — namely IIS — have to configure HTTP and, through available APIs, tell HTTP.sys which endpoints to listen to. If a service requests an endpoint where HTTP.sys is not listening, HTTP.sys logs an event to the event log that states this. In the case of IIS, the WWW service also makes note of these kinds of errors.

| **Note** |
|---|
| Once the IP inclusion list is set, HTTP.sys never reverts to listening on all IP addresses. |

# Troubleshooting HTTP.sys Communication Problems

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In rare cases, the WWW service might encounter problems when communicating with the HTTP service. If this occurs, the WWW service logs errors to the event log. These errors are summarized in Table 11.16

**Table 11.16 WWW Service to HTTP Communication Errors**

| Event ID | Description |
|---|---|
| 1020 | The World Wide Web Publishing Service encountered difficulty configuring the HTTP.sys control channel property *%1*. The data field contains the error number. |
| 1109 | The World Wide Web Publishing Service failed to set the control channels filter configuration. The data field contains the error number. |
| 1062 | The World Wide Web Publishing Service encountered a failure configuring the logging properties on the HTTP Control Channel. Logging Enabled is *%1*. Log File Directory is *%2*. Log Period is *%3*. Log Truncate Size is *%4*. The data field contains the error number. |
| 1037 | A failure occurred while opening the HTTP control channel for the World Wide Web Publishing Service. The data field contains the error number. |

# Troubleshooting Miscellaneous IIS Problems

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Additional problems that might require troubleshooting include unexpected metabase site IDs, CGI process difficulties, server connection problems, and client request timeout and error-out problems.

# Metabase Site IDs Are Unexpected Numbers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you create a new Web site in IIS Manager, the newly generated site ID in the metabase is a random number. In earlier versions of IIS, the site ID was incremented with each new site. If you have an application that in some way relies on the old behavior, you can change the current behavior by adding the registry entry IncrementalSiteIDCreation, a DWORD value, to the registry with a value of 1. Add IncrementalSiteIDCreation to the following subkey:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\InetMgr\Parameters

**Caution** Do not edit the registry unless you have no alternative. The registry editor bypasses standard safeguards, allowing settings that can damage your system, or even require you to reinstall Windows. If you must edit the registry, back it up first and see the Registry Reference in the Microsoft Windows Server 2003 Deployment Kit.

When generating a new site ID, IIS uses a hash of the Web site name. This new design ensures that all IIS 6.0 servers in a Web farm have a good chance of generating the same site identification numbers for sites with the same name. For more information, see Microsoft Knowledge Base article 332122, IIS 6.0: Web Site Identification Numbers Are Not Incremental.

# Anonymous Users Performance Counters in IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Current Anonymous Users performance counter is often used for capacity planning and performance troubleshooting. The counter is incremented differently in IIS 6.0 than it was in IIS 5.0.

In IIS 5.0**.** Current Anonymous Users is the number of current anonymous HTTP connections. This number decreases when the connection is closed (for various reasons, such as when the connection explicitly closes or times out).

In IIS 6.0**.** Current Anonymous Users is the number of current anonymous requests being processed. This number decreases when the request is completed. (HTTP.sys and IIS user mode both provide numbers for this counter.)

In short, IIS 6.0 counts each request as a user, whereas IIS 5.0 counts each connection as a user.

# CGI Processes Will Not Start

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In order for CGI processes to run, the account that starts the CGI processes (the worker process identity) must be assigned certain user rights. You can add the account as a member of the IIS_WPG group and assign the account these two necessary user rights:

- Adjust memory quotas for a process.
- Replace a process-level token.

For information about assigning user rights to an account on the local computer, see CGI processes do not start.

# Clients Cannot Connect to the Server

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Windows Server 2003 provides a software-based firewall to help prevent unauthorized connections to your server from remote computers. By default, the Internet Connection Firewall (ICF) is disabled. However, if you have enabled the firewall in its default configuration after installing Windows Server 2003 and before installing IIS 6.0, clients will not be able to connect to your server. Use the following procedure to configure the ICF to allow clients to initiate Web and other IIS–related connections to your server.

For information about configuring ICF for IIS, see Before Configuring IIS.

# Client Requests Error Out or Time Out

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 6.0, settings are set to aggressive and secure defaults to minimize attacks due to timeouts and limits that were previously too generous. IIS enforces the following timeouts at the connection level:

- Limits on response buffering. The metabase property **ASPBufferingLimit** is 4 MB, by default. If ASP scripts buffer more than this, they generate an error. Prior to IIS 6.0, there was no limit to buffering.

- Limits on posts. The metabase property **AspMaxRequestEntityAllowed** enforces a maximum ASP POST size of 204800 bytes, with each individual field limited to 100 KB. Prior to IIS 6.0, there was no limit to posts.

- The **ServerListenTimeout**property is no longer available**.** The metabase property **ServerListenTimeout** has been replaced by the following metabase properties:

    - **ConnectionTimeout**. This property specifies the amount of time, in seconds, that the server waits before disconnecting an inactive connection.

    - **MinFileBytesPerSec**. When IIS responds to a client request, the **MinFileBytesPerSec** property determines the length of time that the client has to receive the entire response. If the client computer takes too long to receive the entire response, the kernel-mode driver, HTTP.sys, terminates the connection according to the timeout value.

    - **HeaderWaitTimeout**. When a client connects to the Web server, the client computer is given a time limit to send in all headers for the request (demarked by a final double \r\n). If the complete header set for the request is not received within the time period indicated by **HeaderWaitTimeout**, HTTP.sys resets the connection. You can configure the value of **HeaderWaitTimeout**.

- **Header size limitation**. By default, HTTP.sys accepts only requests where the size of the request header is less than 16 KB. This means that if HTTP.sys does not receive the terminating carriage return/line feed (CRLF) pair sequence within 16 KB, HTTP.sys considers the request malicious and terminates the connection. You can change the header size limitation by adjusting the value of the registry entry MaxRequestBytes. Note that this entry does not exist in the registry by default. To add it to the registry, use the registry editor, Regedit.exe.

# Advanced IIS Troubleshooting with Tracing and Debugging

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Up to this point, this section has explored problems and issues that are specific to IIS, its interaction with various Windows operating system mechanisms (like the file system and security), and its implementation of the HTTP protocol. These types of problems are identified by definitive error messages and demonstrate specific faulty behaviors. In addition, these types of problems usually reveal themselves early — during either application development or initial configuration of the IIS application environment. For example, problems with permissions usually appear during security configuration, which usually occurs during application deployment. Problems like these display symptoms immediately, making it is easy to identify the problem.

Today's Web-based applications often have many components and systems that interact. This complexity provides greater opportunity for components to fail or to behave in unexpected ways. Further complicating matters, large and complex applications often exhibit subtle or non-specific symptoms. The effects of such problems are abstract and might seem random. For example, a problem with data access synchronization might cause an application to stop responding in one instance, but a crash in another. Such "quirks" can occur deep within an application's processing, often masking or distorting symptoms and ultimately hiding the root cause of the problem.

IIS is the primary runtime environment for the execution of application logic implemented in custom ISAPI applications, ASP, ASP.NET, or a Web service. These components often communicate with middle-tier business logic servers and database servers. Consequently, IIS is often the main execution environment for the core workings of an application and, therefore, IIS is the environment where problems begin or symptoms are propagated.

This section discusses tracing and debugging, advanced troubleshooting techniques for isolating problems in applications and IIS components.

# Troubleshooting IIS 6.0 with Tracing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Windows operating system includes the Event Tracing for Windows (ETW) infrastructure to help individuals troubleshoot problems in the operating system, including problems in HTTP components. IIS HTTP components include Active Server Pages (ASP), Internet Server Application Programming Interface (ISAPI) extensions, and the Secure Sockets Layer (SSL) Filter service, to name a few. If an HTTP request in IIS fails or becomes unresponsive while ETW is enabled, you can view ETW trace data, called events, to determine which component caused the failure.

In Windows Server 2003 Service Pack 1 or later, IIS includes the following tracing features. These features leverage the ETW infrastructure.

- IIS Currently-executing Requests Tracing: This tracing feature provides general statistics and details about all requests executing on the server at the moment tracing was started. If the CPU on your server is spiking or if requests become unresponsive, currently-executing requests tracing can help you understand which URLs are being requested, which application pool the requests reside in, and other similar details. This feature does not include the option to specify which components or URLs to trace. For more information, see Currently Executing Requests Tracing.

- IIS Request-Based Tracing: This tracing feature tracks HTTP requests as they move through IIS components. Request-based tracing can help you target the IIS component processing a request when the request failed or became unresponsive. Request-based tracing allows you to define and trace a specific URL or a group of URLs. For more information, see IIS Request-Based Tracing.

Windows Server 2003 Service Pack 1 or later also includes a provider for tracing the IIS Admin service during startup and shutdown. The IIS Admin service provides access to the in-memory configuration store and other dependent services. If you experience problems while IIS is starting up or shutting down, the IIS: IIS Admin Service provider can help you understand the nature of the problem. For more information about this provider, see IIS providers.

# Currently Executing Requests Tracing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use currently-executing requests tracing to find out which requests are executing in a worker process when the worker process has become slow or unresponsive. When enabled, currently-executing requests tracing prompts worker processes to report statistics and details about each request in each worker process. Currently-executing requests tracing is enabled via the command-line.

To understand how Request Monitor works, consider the following scenario:

You are the administrator for an e-commerce Web application. The Web application is hosted on IIS 6.0 running Microsoft Windows Server 2003 Service Pack 1.

Your developers make substantial changes to the application for a planned upgrade. While testing the upgraded application in a test environment, you notice in Task Manager that an IIS worker process (W3wp.exe) is periodically spiking to consume 100% of the CPU cycles on the server. You don't know which requests are executing in the worker process or which URLs are being requested. In short, you have no idea where to begin searching for the source of this problem.

Previously, administrators in this situation could terminate the process in Task Manager and then check the HTTPERR log for clues as to what happened. Or, administrators may call Microsoft Product Support Services for assistance in attaching a debugger to the process to get a memory dump of the activity inside the worker process. This is not an ideal solution because of the time and expense involved in such a call.

With currently-executing requests tracing, you can narrow down the source of the problem by executing a command from a command prompt. For example, the administrator types the following command at a command prompt and then presses ENTER:

**logman start W3wpTrace -p IIS: "Request Monitor" -ets**

The administrator lets the trace session run for a few moments and then types the following at a command prompt and presses ENTER:

**logman stop W3wpTrace -ets**

Currently-executing requests tracing returns details about requests executing at the moment the command was executed. For example:

```
<WorkerProcessProcessID=3420
                    AppPoolId=DefaultAppPool
                    RequestsService=126
                    TotalWorkerThreads=4
                    AvailableWorkerThreads=4
                    ProcessUpTime=0:26:34
                                 >
<WebRequestList>
              <WebRequest ContextId={00000000-0000-0000-8f00-0060000000fe}
                     SiteId=1
LocalIPAddr=<data>
LocalPort=80
HostName=Contoso3
ConnectionID=1073741965
                     ProcessID=3420
ClientIP=<data>
Verb=Get
URL=/asp/asptest.asp
ProcessingState=processing
TimeElapsed=29723TimeFromLastState=29723

/>

      <WebRequestList>
             <WebRequest ContextId={00000000-0000-0000-2900-0060000000fb}
                     SiteId=1
LocalIPAddr=<data>
LocalPort=80
HostName=Contoso3
```

```
ConnectionID=1073741859
                    ProcessID=3420
ClientIP=<data>
Verb=Get
URL=/asp/asptest.asp
ProcessingState=processing
TimeElapsed=28281TimeFromLastState=28281
/>

</WebRequest>
</WebRequestList>
</WorkerProcess>
```

With this information, the administrator determines there are two separate requests for asptest.asp, the processing state for each request is processing, and both of these requests have been processing for more than 28 seconds. With this information, the administrator can focus their troubleshooting efforts on the asptest.asp page. For example, the administrator could send the Request Monitor data to the developer of asptest.asp and ask the developer to look for code errors.

| **Note** |
| --- |
| If you expect Request Monitor to return hundreds of request details, you can filter the data output using other parameters. |

For more information about tracing currently-executing requests with Request Monitor, including the step-by-step procedure for tracing requests and a detailed explanation of the command-line switches, see How to Trace All Requests Currently Executing in IIS Worker Processes.

# IIS Request-Based Tracing

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use IIS request-based tracing to track requests in IIS components like ASP, SSL filters, and ISAPI extensions. Request-based tracing also allows you to trace a request to a specific URL or a group of URLs, which can further refine your troubleshooting efforts.

When an HTTP request is picked up by an IIS user mode worker process (meaning the request was not served by the HTTP.sys response cache), the request passes through multiple components before a response is generated and IIS sends the response back to the client computer. Depending on the request type, the request is processed by the static file request handler, the CGI request handler, or an ISAPI extension handler. Requests that enter an ISAPI extension handler are processed further by Active Server Pages (ASP) or ASP.NET pages. If a request fails or becomes unresponsive as it makes its way through these components, error-reporting channels like the Windows NT Event Log and the HTTP Error log might not provide enough detail to help you locate the source of the problem. Requested-based tracing will provide the details you need to help locate the source of the problem.

In the context of Event Tracing for Windows (ETW), IIS components are called **providers** because they provide trace data or events to ETW.

From a high level, request-based tracing works as shown here:




1. The administrator enables a tracing session on the IIS server from a command prompt.

2. ETW notifies IIS providers to begin reporting trace events.




3. A request enters the IIS worker process. As the request moves through IIS, the providers report events to the ETW infrastructure. In this example, the worker process reports a request start event, ASP_Start_Request, and then each of the IIS components report Start and End events. Finally, the worker process reports a request end event, ASP_End_Request. ETW writes these events to the trace log. If a request experienced a change while processing, for example the client computer disconnected while the request was on the ASP queue, ASP would report the ASP_CLIENT_DISCONNECTED event. If a request experienced an error, ASP would report an event that described the source of the problem, for example, ASP_FILE_NOT_FOUND or ASP_FILE_ACCESS_DENIED.




4. The administrator stops the trace session and reviews the trace log to locate the source of the problem. Problems are typically identified by error events or a START event that does not contain a corresponding END event.




You can specify which providers report events to ETW during a specific trace session. You can also customize how much trace data providers report to ETW. For example, you can configure a trace session so IIS providers report error events only, or you can configure the trace session so providers report informational events and error events. You can also determine which URLs to trace, thereby focusing your troubleshooting efforts on a specific application.

# The IIS Request-Based Tracing Process

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This section describes, from a high level, the process of setting up, executing, and handling the post-processing for a request-based tracing session in IIS. This section is intended to give the reader an overview of the tracing process, including the technical details of tracing. This section does not include the step-by-step procedure for executing a trace in IIS. For the step-by-step procedure to execute a trace in IIS, see Advanced Troubleshooting with IIS Tracing.

Tracing a request in IIS involves the following:

- Configuring the IIS Request-Based Tracing Session
- Starting the IIS Request-Based Tracing Session
- Reproducing the Problem
- Stopping the IIS Request-Based Tracing Session
- Post-processing and Viewing IIS Request-Based Tracing Data

The following section discusses tracing using a **provider** file. A provider file is a simple text file that contains a list of the providers you want to trace. When you execute tracing, you can identify which providers (components) to trace at the command-line using switches, or, you can identify which providers to trace using a provider file. The provider file method is described here because a provider file reduces the amount of information you need to type if you plan to use tracing regularly.

# Configuring the IIS Request-Based Tracing Session

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Setting up a request-based tracing session in IIS requires that you first decide what you want to trace. Your decision on what to trace will be determined by the nature of your Web applications. For example, if your Web applications use Active Server Pages (ASP), then you would list this provider in the provider file. If your Web applications do not use ASP, then you would not want to list this provider in the provider file. After you start tracing, the components defined in the provider file report trace data to the ETW infrastructure.

Your decision on what to trace will also be determined by the nature of the problem you are troubleshooting. If you know which component is causing the problem you can create your provider file to target the problematic component. If you are having troubles with a specific application, you can trace the specific URL.

If you are not certain which URL or component(s) are causing the problem, use IIS currently executing requests to determine which URL is being requested when a worker process becomes unresponsive. Knowing which URL is causing the problem will also help you determine which other providers to define in the provider file. For example, if you know the problem exists in a URL that utilizes an ISAPI filter, then you would want to define that provider in the provider file. By explicitly defining which components to trace, you reduce the amount of data in the trace log, which makes troubleshooting easier.

This section discusses the following concepts related to setting up an IIS request-based trace session:

- The provider file
- IIS providers
- Verbosity levels

# The provider file

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The provider file is a simple text file that resides on the local computer and tells ETW which providers should report trace data. This file also tells ETW how verbose the trace reporting should be; that is, how much data the providers should report, as well as what flags or areas to trace for each provider.

The contents of a provider file contain the following:

- The names of the providers that will report trace data to ETW
- Provider flags (for applicable providers) that define which areas of the provider to trace
- A number that tells ETW how much trace data to collect for the specified provider, called verbosity level

| **Note** |
| --- |
| Providers, flags, and verbosity levels are described later in this section. |

In the following example, the first provider is WWW Server. The flag 0xFFFFFFFE denotes that all subcomponents of the WWW Server component should be traced, including authentication, security, compression, etc. The number five (5) denotes that the provider should report all tracing events for that provider, the most detailed trace information possible. For the other two providers, ASP and ISAPI extensions, the provider name is followed by zero and then five. The zero tells ETW not to trace any flags for the stated providers. The number five denotes the verbosity level for the providers.

```
"IIS: WWW Server" 0xFFFFFFFE 5
"IIS: Active Server Pages (ASP)" 0 5
"IIS: WWW Isapi Extension" 0 5
```

# IIS providers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Components that provide tracing events to ETW, like ASP and ISAPI extensions, are called providers. In Windows Server 2003, Service Pack 1, the following IIS providers report trace data to ETW:

- IIS: WWW Server
- IIS: SSL Filter
- IIS: ISAPI Extension
- IIS: Active Server Pages (ASP)
- IIS: IISAdmin
- IIS: Request Monitor

When you configure the trace session by creating the provider file, you define which providers report data. Components not listed in the provider file do not report trace data. Each of the providers listed here is described in greater detail later in this section.

If you are actively troubleshooting a problem on a busy production server, try to limit the number of providers defined in the provider file or use the IIS: WWW Server **UrlFilter** flag to restrict tracing to a specific URL. Doing so will result in a more manageable set of trace data in the trace log which should make it easier to locate the source of the problem.

If you are new to request-based tracing, run a few practice trace sessions on a test computer with multiple (or all) providers defined in the provider file so you can see what kind of events are generated with “normal” traffic and how quickly the trace log can fill with data.

**IIS: WWW Server**

This provider traces the following core WWW service functional areas. If you wish to trace one or more of these functional areas, you must define these areas in the provider file using the designated flag or the friendly flag. If a flag or friendly flag is not designated in the provider file, the functional area will not be traced. The following table lists all of the functional areas, the flags to enable them, and the "friendly" flag name of each functional area. Each of these functional areas is described in detail later in this section.

| Functional Area | Flag | Friendly Flag Name |
|---|---|---|
| UrlFilter | 1 | UseURLFilter |
| Authentication | 2 | IISAuthentication |
| Security | 4 | IISSecurity |
| ISAPI Filter | 8 | IISFilter |
| Static File | 10 | IISStaticFile |
| CGI | 20 | IISCGI |
| Compression | 40 | IISCompression |
| Cache | 80 | IISCache |
| WWW Server All | 0xFFFFFFFE | IISAll |

**Note**

You can identify more than one functional area in a provider file by adding the numeric flag values. For example, if you want to trace using the **Urlfilter** flag (1), the **Authentication** flag (2), and the **Security** flag (4), you can enter a flag value of 7 in your provider file to trace all of these areas.

**UseUrlFilter**

The **UseUrlFilter** flag allows you to specify which URL(s) to trace. If you do not enable this flag, request-based tracing traces all URLs. If you enable this flag and configure the TraceUriPrefix Metabase Property for the URL's you want to trace, you can preserve system resources like disk space and CPU overhead while narrowing your troubleshooting efforts to a much smaller data set in the trace log.

**Important**

If you plan to trace a specific URL or group of URLs using the **UseUrlFilter** flag, you must configure the **TraceUriPrefix** property in the IIS metabase with the URL's you want to trace. If you enable the **UseUrlFilter** flag and do not enable the **TraceUriPrefix** property, ETW will not trace any URLs. For information on how to enable the **TraceUriPrefix** metabase property, see How to Trace Requests for a Specific URL or Set of URLs.

**Authentication**

The **IISAuthentication** flag traces authentication actions, including the name of the authenticated user, the authentication scheme (anonymous, basic, etc) and the results of the authentication attempt (successful, failed, error, etc). If users are having problems authenticating on your Web application, use the **IISAuthentication** flag to understand what is happening with authentication on the server.

**Security**

The **IISSecurity** flag generates trace events when requests are rejected by the IIS server for security-related reasons, like a client request was denied access to a resource. Use this flag to understand why requests are being rejected for security reasons.

**ISAPI Filter**

The **IISFilter** flag traces calls into and out of ISAPI filters. Use this flag to understand how requests have been changed after entering an ISAPI filter, as well as how long requests are spending in the filters.

**Static File**

The **IISStaticFile** flag generates trace events when a request is for a static file.

**CGI**

The **IISCGI** flag generates trace events when a request is for a CGI file.

**Compression**

The **IISCompression** flag generates trace events when a response is compressed.

**Cache**

The **IISCache** flag generates trace events for cache operations associated with the request.

**WWW Server All**

The **IISAll** flag forces all of the WWW Server functional areas to generate trace events.

**IIS: SSL Filter**

This provider traces calls into and out of the Secure Sockets Layer (SSL) Filter, which is used for encrypting and decrypting SSL requests.

**IIS: WWW ISAPI Extension**

This provider traces ISAPI extension requests.

**IIS: Active Server Pages (ASP)**

This provider traces requests for Active Server Pages (ASP) files, including calls into and out of COM+ objects. The events captured by this provider are useful for diagnosing hangs or crashes caused by COM objects used in ASP pages.

**IIS: IISAdmin Service**

This provider traces the IIS Admin service during startup and shutdown. The IIS Admin service reads the IIS metabase.xml file which stores all IIS configuration data. Use this provider if you notice that the IISAdmin service has trouble starting up or shutting down.

This provider includes the following flags:

### Startup

The Startup flag forces ETW to trace IIS Admin service startup events.

### Shutdown

The Shutdown flag forces ETW to trace IIS Admin service shutdown events.

For more information about this provider, see Tracing Startup and Shutdown Problems in the IIS Admin Service.

#### IIS: Request Monitor

This provider traces all requests currently executing in worker processes at the moment tracing was started. For more information about this provider, see Currently Executing Requests Tracing.

# Verbosity levels

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Once you have established which providers will report data to the trace log, you must determine how much data each provider should report. You establish how much data each provider reports by setting the *verbosity level* in the provider file. The verbosity level is a number in the provider file, ranging from 0 to 5, where 0 is the least information and 5 is the most information. If you are troubleshooting a specific error and only want to see tracing events related to that error, you would probably choose verbosity level 1 or 2. If you are troubleshooting a problem and you need tracing to help you locate the source of the problem, you would probably choose verbosity level 3, 4, or 5. Table 11.3 describes verbosity levels in more detail.

If you plan to execute tracing using a provider file, you can specify different verbosity levels for different providers. For example, you could configure the **IIS:Active Server Pages (ASP)** provider to log trace events at verbosity level 5 and **IIS:WWW Server** at verbosity level 2. If you plan to execute tracing solely from the command line, you cannot specify different verbosity levels. **IIS:Active Server Pages (ASP)** and **IIS:WWW Server** would need to have the same verbosity level to execute the trace.

**Table 11.3**

| Level | Name | Description |
|---|---|---|
| 0 | General | IIS providers log only those few events required to provide some context to the request activity. These events include a **General_Request_Start** event that logs the URL and verb for the request and a **General_Request_End** event. |
| 1 | Fatal Error | IIS providers report events for actions that can cause a process to exit or actions that are about to cause a process to exit. |
| 2 | Error | IIS providers report events when components cannot continue processing a request. The request is immediately rejected.<br>**Note**<br>These events typically indicate a problem on the server side rather than a client-induced problem. |
| 3 | Warning | IIS providers report events when components experience an unexpected result/error but can continue processing the request. |
| 4 | Information | IIS providers report general information about requests. |
| 5 | Verbose | IIS providers report detailed information about requests. |

Once you have determined the providers you want to trace, the flags (if any), and the amount of information you want to collect in the trace log (the verbosity level), you enter these items in a provider file, for example, MyProviders.txt.

# Starting the IIS Request-Based Tracing Session

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS request-based tracing uses Log Manager (Logman.exe) to start, stop, and update trace sessions. Log Manager is a command-line tool included with Windows Server 2003. Log Manager can schedule trace collections on a local machine or on a remote machine.

| **Note** |
| --- |
| This section does not contain the step-by-step procedure for starting and stopping a trace session. This section describes some of the essential concepts for executing a trace session. For the step-by-step procedure to execute a trace session in IIS, see Advanced Troubleshooting with IIS Tracing. |

When you enable tracing, you enter a command at the command prompt with the following:

1. A command to start Log Manager (**Logman start trace**).
2. A file name where the trace data should be collected.
3. The provider file switch (**-pf**)
4. The name of the provider file, including the path to the provider file, if necessary.
5. A command to start the trace immediately (**-ets**)

The following is an example of a command that starts tracing with Log Manager:
**Logman start trace ISAPItrace -pf MyProvidersInputFile -ets**

After you enable a trace session, Logman creates the session in the kernel. When Logman has enough session data, it creates a trace log file (.etl extension) in the directory where you executed the start trace command. Log Manager assigns the file the name you entered after the **start trace** command, for example ISAPItrace.etl (in the command above).

After you start the trace session, Logman prints to the screen a few general statistics and details about the session you just started, as seen below. Notice that Logman tells you the name of the trace session, the file name and path where the trace data should be collected, and the details of the provider file.

```
D:\Documents and Settings\TestMachine>logman start trace ISAPItrace -pf
d:\MyProviders.txt -ets
Name:ISAPItrace
Age Limit:15
Buffer Size:8
Buffers Written:1
Clock Type:System
Events Lost:0
Flush Timer:0
Buffers Free:2
Buffers Lost:0
File Mode:Sequential
File Name: D:\Documents and Settings\TestMachine\ISAPItrace.etl
Logger Id:3
Logger Thread Id:1468
Maximum Buffers:25
Maximum File Size:0
Minimum Buffers:3
Number of buffers:3
Real Time Buffers Lost:0

Provider                          Flags                    Level
-----------------------------------------------------------------------
* "IIS: WWW Server" (IISAuthentication, IISSecurity, IISFilter,
IISStaticFile,IISCGI,IISCompression,IISCache,IISAll)0x05
```

```
{3A2A4E84-4C21-4981-AE10-3FDA0D9B0F83}  0x00000ffe0x05
* "IIS: WWW Isapi Extension"             0x000000000x05
{A1C2040E-8840-4C31-BA11-9871031A19EA}  0x000000000x05

The command completed successfully.
```

For information on Log Manager command-line syntax and parameters, type **Logman /?** at a command prompt and press ENTER. For a list of providers or to see the exact syntax required in the provider file, type **logman query providers** at a command prompt and press ENTER.

# Reproducing the Problem

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Once you have created your provider file and started the trace session, you must reproduce the problem. Reproducing the problem may be as simple as making one request from a client computer to a troublesome application, or as complex as using multiple clients to create load on the server.

Once you have reproduced the problem, you will be able to check the trace log file for details about the error or problem. If, after reproducing the problem, you do not find helpful details in the trace log, you may need to enable other providers or increase the verbosity level of your providers in the provider file. In short, if you have reproduced the problem, the events captured in the trace log should provide you the details of the problem.

For more information about tracing events and what they mean, see Tracing Events Reference.

# Stopping the IIS Request-Based Tracing Session

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Stop the trace session after you have reproduced the problem. If you are unable reproduce the problem, stop the trace session after a few minutes and then review the trace log. Be aware that tracing consumes system resources. On a server processing hundreds of requests per second, the trace log file created by ETW can grow in size to several megabytes after a few seconds, depending on the number of providers defined in the provider file. Immediately disable tracing on production servers after you have gathered enough data to help you isolate and troubleshoot the problem.

To stop the trace session, execute the **logman stop trace** command from the command prompt. For example:
**Logman stop trace ISAPItrace -ets**

| **Note** |
| --- |
| ETW does not support running multiple trace sessions with the same provider at the same time. You must stop an existing trace session before starting a new one. For example, you can not run two trace sessions where both sessions use the **IIS: WWW Server** provider. |

# Post-processing and Viewing IIS Request-Based Tracing Data

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Event Tracing for Windows writes trace data to the trace log (.etl log file) in a binary format. To view the trace data, the trace log must be processed by a tool. The tools described below process .etl log files.

**Trace Report (tracerpt.exe)**
Trace Report ships with Windows Server 2003. To process the trace log file with Trace Report, type tracerpt filename and press ENTER. For example:

**tracerpt NewTrace.etl**

Trace Report converts the binary trace data into text and produces two files in the directory where you executed the **tracerpt** command:

- Summary.txt contains general details about the trace session, including which providers were used.

- DumpFile.csv contains the actual trace data in a text format. Open DumpFile.csv in a text editor like Notepad to view the trace data. The data in DumpFile.csv is readable (shown below), but the formatting is not as reader-friendly as the format produced by Log Parser and DumpTraceReqs (described later in this topic):

```
Event Name,       Type,        TID,           Clock-Time, Kernel(ms),   User(ms), User
Data
  EventTrace,     Header, 0x000000EC,   127311198897300433,         10,         60,
8192, 33620485,     3790,        1, 127311199293485921,   100144,        0, 0x00000000,
6,        1,        4,        0,      730, 0x00B82930, 0x00B82944,                0,
3579545, 127311198897300433, 0x00000002,        0, 0, 0
  IISGeneral, GENERAL_REQUEST_START, 0x00000A74,   127311199074261949,          0,
0, {00000000-0000-0000-0200-0060000000ed},        1, "DefaultAppPool",
17077649788062662657,                0, "/ASPtest", "GET", 0, 0
   IISFilter, FILTER_START, 0x00000A74,   127311199074261949,          0,          0,
{00000000-0000-0000-0200-0060000000ed}, "D:\IISDebugTools\IISCHAgent.dll", 0, 0
   IISFilter, FILTER_PREPROC_HEADERS_START, 0x00000A74,   127311199074261949,
0,          0, {00000000-0000-0000-0200-0060000000ed}, 0, 0
   IISFilter, FILTER_PREPROC_HEADERS_END, 0x00000A74,   127311199074261949,          0,
0, {00000000-0000-0000-0200-0060000000ed}, 0, 0
   IISFilter, FILTER_END, 0x00000A74,   127311199074261949,          0,          0,
{00000000-0000-0000-0200-0060000000ed}
```

**Log Parser 2.2 and the DumpTraceReqs.js script**
Log Parser is a free tool for parsing and formatting log files or other files that contain large amounts of unformatted data. The DumpTraceReqs.js script, which is included with Log Parser 2.2 on the Microsoft Download Center, was written by the IIS team for the purpose of processing and formatting IIS request-based tracing logs. DumpTraceReqs.js groups all events per request in chronological order, and displays them per request. You do not have to be familiar with Log Parser and the sample script to format your trace logs. You simply execute a command from the command line and the tool processes the log file for you. Whenever possible, use Log Parser and DumpTraceReqs.js to process and view trace logs because the output is much easier to read than the output from tracerpt.exe. For example:

```
Microsoft (R) Windows Script Host Version 5.6
Copyright (C) Microsoft Corporation 1996-2001. All rights reserved.
---------------------------------------------------------------------
Request n.1: /ASPtest {00000000-0000-0000-0200-0060000000ed} 2004-5-7 22:11:47

  IISGeneral: GENERAL_REQUEST_START - IIS starts processing a new request
    SiteId: 1
    AppPoolId: DefaultAppPool
    ConnId: 1073741825
    RawConnId: 0
    RequestURL: /ASPtest
    RequestVerb: GET
    ContextIDSeq: 1
    Timestamp: 22:11:47.426.194900
```

```
IISFilter: FILTER_START - W3 filter starts
  FilterName: D:\IISDebugTools\IISCHAgent.dll
  ContextIDSeq: 1
  Timestamp: 22:11:47.426.194900

IISFilter: FILTER_PREPROC_HEADERS_START - Filter starts PREPROC_HEADERS notification
  ContextIDSeq: 1
  Timestamp: 22:11:47.426.194900

IISFilter: FILTER_PREPROC_HEADERS_END - Filter ends PREPROC_HEADERS notification
  ContextIDSeq: 1
  Timestamp: 22:11:47.426.194900

IISFilter: FILTER_END - W3 filter ends
  ContextIDSeq: 1
  Timestamp: 22:11:47.426.194900
```

Once Log Parser and DumpTraceReqs.js have processed your .etl file, Log Parser prints the data to the screen. Viewing the data on the screen may not be an optimal viewing experience, so it is recommended that you open the processed .etl file in Notepad, as described in How to Process and View Trace Log Files

Once you have opened the processed .etl file in Notepad, you will see that each request contains a list of the trace events that were recorded during the trace session. You can scroll down the page to view all the subsequent requests and events.

For the procedure to configure Log Parser 2.2 and process an .etl file using DumpTraceReqs.js, see How to Process and View Trace Log Files.

| **Note** |
| --- |
| DumpTraceReqs.js was not designed to process data from the IISAdmin provider. If you need to process and view IISAdmin trace data, use the Trace Report tool included with Windows Server 2003. |

For more information about the tracing events in your trace log file, see the Tracing Events Reference.

# Tracing Startup and Shutdown Problems in the IIS Admin Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 6.0, the IIS Admin Service loads the primary configuration file for IIS sites, called the IIS MetaBase.xml file, and related in-memory configuration services. The IIS Admin service also loads the following:

- The File Transport Protocol (FTP) service
- The Network News Transport Protocol (NNTP) service
- The Simple Mail Transfer Protocol (SMTP) service

If the IIS Admin Service experiences a problem (hangs) while starting up or shutting down, diagnosing the problem can be extremely difficult because there are potentially four services involved. In the past, troubleshooting IISADMIN startup or shutdown hangs usually required calling Microsoft Product Support Services who would then help the user collect a process dump of the IIS Admin Service.

Currently, the best way to troubleshoot hangs in the IIS Admin service is to trace the service using the IIS Admin provider. Once you have collected trace data in the trace log, you can view the trace log to learn specific details about when the service stops responding. You can trace the IIS Admin service during startup or shutdown by choosing the **Startup** or the **Shutdown** flag. If necessary, you can trace using both flags at the same time.

In order to capture data in the trace log, you need to enable IIS Admin tracing and then reproduce the problem on your IIS server. If the IIS Admin service is currently hung, you have two options:

1. If the IIS Admin service hung on startup, enable tracing as described in How to Trace the IIS Admin Service. Then use the **iisreset /no force***computername* command to restart the IIS server. After **iisreset** terminates the process and IIS restarts, the IIS Admin service will report trace data in the trace log. For more information about the **iisreset** command, see Restarting IIS.

2. If the IIS Admin service hung on shutdown, use the **iisreset /no force***computername* command to restart the IIS server. Once the IIS server has restarted, enable tracing and then shutdown the IIS Admin service.

| **Note** |
| --- |
| The IIS Admin service can be configured to restart automatically, a feature known as Automatic Restart. If you are attempting to trace the IIS Admin service to troubleshoot hangs, you may want to disable this feature until you have remedied the problem. For more information about this feature, see "Automatic Restart" in Restarting IIS. |

For information about the events listed in the trace log, see Tracing Events Reference.

# Troubleshooting IIS with a Debugger

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Advanced troubleshooting activities will sometimes require a debugging tool. A debugging tool allows you to peer inside an executing Win32 process to observe what is happening. You can use the debugging tool to capture a record of the memory space of the process that is hosting the faulty application. In IIS 6.0, the process is usually W3WP.EXE. This memory dump is written to a DMP file.

The Microsoft Debugging Tools for Windows include a script file, named Adplus.vbs, which is a scripted front-end to the Microsoft Console Debugger (CDB). Adplus.vbs makes capturing dump files easy.

The key to using any debugger is to understand the kind of problem that you are troubleshooting. Once you identify the kind problem you are troubleshooting, you can decide what to do with the debugger and what kind of dump file to obtain. After you have obtained a dump file, it is recommended that you then have it debugged by a developer or a Microsoft support professional.

The IIS Resource Kit Tools contain a tool called IISState. This tool is like a debugger and can also be used to look at a running process. IISState contains some heuristic functionality that attempts to figure out what is happening in the process for you. For more information, see the readme.txt that comes with IISState.

For detailed information about troubleshooting with Adplus.vbs, see Microsoft Knowledge Base article 286350, How To: Use ADPlus to Troubleshoot "Hangs" and "Crashes".

# Troubleshooting Low CPU Hangs with a Debugger

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

An application that becomes unresponsive to any kind of request while the process hosting the application is consuming little or no system processor(s) CPU time is in a low CPU hang state. This type of problem can be caused by unrecoverable lock situations in program code (deadlocks) or by threads that are waiting on an external resource, such as a database.

To capture a memory dump of the offending process, at a command prompt, type the following: **C:\debuggers\> cscript.exe adplus.vbs -hang -p** *pid*

If you know that the hung process is an IIS process but do not know which process is causing the problem, instead of specifying a specific PID or process name, specify the **-iis** parameter. Specifying this parameter will cause Adplus.vbs to capture all IIS–related processes.

# Troubleshooting High CPU Hangs with a Debugger

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

An application that becomes unresponsive to any kind of request while the process hosting the application is consuming a great deal of system processor(s) CPU time, usually near 100%, is in a high CPU hang state. This problem can be caused by an endless loop in script code or by some other intensive kind of processing.

In such a case, it is useful to identify the thread that is consuming the CPU before you attempt to capture a memory dump. To identify the offending thread, you can use a tool like QSlice.exe. For information about QSlice.exe, in Help and Support Center for Windows Server 2003, click **Tools**, and then click **Windows Support Tools**. Alternatively, you can check the .log file that is generated by the Adplus.vbs script. This log file contains output generated by the **!runaway** debugger command. The **!runaway** command lists the CPU consumption times of the various running threads. For information about debugging and the debugger commands, see Debugging Tools for Windows Help.

To capture a memory dump, at a command prompt, type the following: **C:\debuggers\> cscript.exe adplus.vbs -hang -p** *PID*

# Troubleshooting Crashes with a Debugger

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

A crash can occur when code executing in the process hosting the application encounters an unexpected event that causes an exception and results in a process shutdown. Examples of such events include a stack overflow, encountering an invalid memory address (an access violation), or executing an invalid instruction. Because these situations are unexpected and, thus, unobservable, you need to run Adplus.vbs in monitoring, or *crash*, mode before the problem occurs again.

To run Adplus.vbs in crash mode, at the command prompt, type: **C:\debuggers\> cscript.exe adplus.vbs -crash -p** *PID*

If you do not know which process is crashing, use the **-iis** parameter, which causes Adplus.vbs to monitor all of the IIS–related processes for crashes. When a crash occurs, the dump files are written to disk.

# Troubleshooting Memory Problems with a Debugger

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In a typical memory problem, application code consumes too much physical system memory or too much of the virtual address space for the process in which the application is running. These problems are called memory leaks. Memory leaks can also cause the other crash and hang problems, as described earlier in this section. High CPU hangs can occur when the memory allocations of the application are intensive and consume a disproportionately great amount of CPU time. Crashes can occur when memory has been mismanaged and corruption has occurred.

# IIS 6.0 Debugging Features

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ISAPI application developers that take advantage of the **ServerSupportFunction** callback function with the **HSE_REQ_REPORT_UNHEALTHY** parameter can make debugging their ISAPI applications easier. When worker process pinging is enabled, you can also enable orphaning to aide in obtaining memory dumps. This is explained in the next section.

The ASP and ASP.NET ISAPI extensions apply logic that takes advantage of this feature as well. For example, ASP monitors the status of its internal thread pool. If too many of its threads enter a blocking state, it reports itself as unhealthy, which triggers a worker process recycle if pinging is enabled for the application pool. ASP.NET also has several algorithms to detect an unhealthy state.

For more information about the ServerSupportFunction callback function, see ServerSupportFunction Extension Function in the IIS Platform SDK on MSDN.

# Enabling Debugging for IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By default, when a worker process is determined to be unhealthy, the WWW service terminates it and starts a new worker process to serve new or queued client requests. Optionally, you can allow an application pool to keep the worker process running, but release, or *orphan*, the worker process from the application pool for purposes of debugging while at the same time initializing a new worker process to handle incoming client requests. An orphaned worker process ceases to take new requests from HTTP.sys.

If you enable debugging on an application pool, the WWW service allows an unhealthy worker process to keep running when the WWW service releases the worker process from serving an application pool. When enabling debugging, you can also configure the WWW service to start an executable application or script for a released worker process. For example, the WWW service might start an application that sends you an e-mail notifying you of the failure.

To enable debugging, set the **OrphanWorkerProcess** property to **true**. Use the **OrphanActionExe** and **OrphanActionParams** properties to configure the application that will run when the worker process is orphaned. For example, to use Adplus.vbs to obtain a memory dump, you could set the properties as follows:

**OrphanActionExe: %SystemRoot%\System32\CScript.exe**

**OrphanActionParams: C:\debuggers\adplus.vbs -quiet -hang -p %1%**

Be sure to closely monitor any worker process that is released from its application pool but continues to run. The worker process can terminate because the WWW service leaves it running but no longer monitors it, which effectively tells the worker process to shut down without forcing it to do so. If the worker process recovers from its unhealthy state, it detects that it has no relationship to the WWW service and self-terminates. When an orphaned worker process self-terminates, you might find an event log entry stating that a worker process was released, but find no evidence of the worker process ever having run.

If you enable debugging on unhealthy worker processes, be sure to monitor these released worker processes, because IIS does not remove them from memory. If you do not correctly handle worker processes that you keep alive for debugging purposes, you can have many failed worker processes running on your computer. These worker processes can tie up resources needed by other processes. In such a case, you might need to terminate those worker processes quickly to free those resources. In some cases, these worker processes can block metabase access and cause problems with other worker processes or with the WWW service itself.

The **OrphanActionExe** and **OrphanActionParams** property settings shown above will leave an orphaned worker process running after obtaining a dump file. As an alternative, you can use CDB directly from a .cmd script file. For example, paste the following text into a new text file and save the file as *LocalDrive*\debuggers\action.cmd:

```
@if "%_echo%"==="" echo off
setlocal

set WORKINGDIR=c:\debuggers
set FILENAME=%WORKINGDIR%\crash.dmp
set LOG=%WORKINGDIR%\log.txt
set COMMAND=c:\debuggers\cdb.exe -c ".dump /o /mhf /u %FILENAME%;q" -p %1

echo %COMMAND% > %LOG%
%COMMAND%

Endlocal
```

Then you can set the orphaning properties as follows:

**OrphanActionExe: C:\Debuggers\action.cmd**

**OrphanActionParams: %%1%%**

Using the **.dump** command combined with the **q** command with the CDB debugger will capture the dump file and terminate the worker process.

# Limitations of Health Detection

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Health detection cannot be used to determine application failures that do not cause the worker process to crash or to block the available threads in the worker process. As an example, an application that returns invalid response codes, such as HTTP 500 errors, but otherwise functions normally, will still respond to being pinged from the WWW service, unless the application is a custom ISAPI extension that implements specific code to indicate its unhealthy state.

# Additional Resources for Troubleshooting IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to troubleshooting IIS 6.0.

**Related Information**

- IIS 6.0 Administration Scripts, Tips, and Tricks for information about how to enable an application that has been locked down.
- IIS 6.0 Architecture for information about the components of the IIS services infrastructure.
- Managing a Secure IIS 6.0 Solution for a brief overview of the authentication options available in IIS 6.0.
- The Shareware.com Web site to download a copy of the HRPlus utility.
- Debugging Tools for Windows - Overview to download Microsoft Debugging Tools for Windows.
- The RFC-Editor Web site for detailed information about RFC 2616, Hypertext Transfer Protocol — HTTP/1.1.
- Deploying and Configuring Internet Information Services (IIS) 6.0 with Remotely Stored Content on UNC Servers and NAS Devices for a complete look at using remote content on a UNC share.
- The following Microsoft Knowledge Base articles:
    - Article 148942, How to Capture Network Traffic with Network Monitor for information about capturing network traffic.
    - Article 172983, Explanation of the Three-Way Handshake Via TCP/IP for information about using Network Monitor in this way.
    - Article 314611, Error Message: HTTP Error 500-12 (Application Restarting) for information about errors that occur when the ASP application to which the requested ASP page belongs is in the process of shutting down, restarting, or both.
    - Article 252876, How to View HTTP Data Frames Using Network Monitor for information about HTTP traffic and Network Monitor.
    - Article 257174, Using Mapped Drives with IIS.
    - Article 258063, Internet Explorer May Prompt You for a Password for information about using the NetBIOS name instead of an FQDN or IP address during NTLM authentication.
    - Article 264921, Info: How IIS Authenticates Browser Clients for information about authentication and authorization.
    - Article 284285, How To: Use Wfetch.exe to Troubleshoot HTTP Connections for information about using WFetch in a testing environment.
    - Article 294382, Authentication May Fail with 401.3 Error If Host Header Differs from Servers NetBIOS Name for information about authentication and authorization.
    - Article 298408, IIS Generates Courtesy Redirect When Folder Without Trailing Slash Is Requested for information about courtesy redirects.

- Article 313074, How To: Redirect Browser Request in Internet Information Services 5.0 for information about courtesy redirects.
- Article 325358, How To: Configure User and Group Access on an Intranet in Windows 2000 or Windows NT 4.0 for information about authentication and authorization.
- Article 326985, How To: Troubleshoot Kerberos-related Issues in IIS for information about using the security log to troubleshoot Kerberos authentication.
- Article 811889, How To: Troubleshoot the Cannot Generate SSPI Content Error Message for information about authentication and authorization.

**Related topics in the IIS 6.0 Operations Guide**

- AuthFlags Metabase Property for information about the settings for the **AuthFlags** property.
- Enabling Edit-While-Running in IIS 6.0 for information about the metabase edit-while-running feature in IIS 6.0.
- HttpRedirect Metabase Property for information about temporary HTTP redirects.
- Listing Running Web Applications Using Iisapp.vbs for information about using Iisapp.vbs.
- Recycling Worker Processes with IIS 6.0 for information about when to manually recycle an application pool or worker process.
- Requests for static files return 404 error for information about troubleshooting static content problems.
- Restoring the Metabase from Backup Files for information about the metabase backup/restore feature.
- Securing Sites with Web Site Permissions for information about securing the Web sites hosted on your IIS 6.0 server.
- Troubleshooting in IIS 6.0 for additional information about troubleshooting strategies for IIS 6.0.
- Working with MIME Types for information about adding a wildcard MIME type.

**Related Windows Server 2003 Help Topics**
For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the **Search in title only** check box.

- "Create and Configure Trace Logs" for information about setting flags and levels in IIS Enterprise Tracing for Windows.
- "Delegating authentication" for information about constrained delegation.

**Related Tools**

- IIS 6.0 Resource Kit tools (IIS60rkt.exe)

  Install the Resource Kit Tools by running IIS60rkt.exe. When you perform a default installation of the Resource Kit Tools, the files are installed in the *LocalDrive*:\Program Files\IIS Resources folder.

- Windows Support Tools

  Obtain the Windows Support Tools by installing them from the \Support\Tools folder of the Windows Server 2003 operating system CD. See the Support Tools Help file for information about using these tools.

# Common IIS 6.0 Administrative Tasks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you are already familiar with running Internet Information Services (IIS) 6.0 and need only a quick refresher to perform an administrative task, or if you are learning how to use IIS and want to know the most common tasks that follow an IIS installation, use this appendix as a quick reference. IIS 6.0 provides a wealth of tools and features, many of which are outlined in this appendix, to help you create a strong and secure communications platform of dynamic network applications.

**In This Appendix**

Overview of Common Administrative Tasks

Important First Tasks after Installing IIS 6.0

Tasks New to IIS 6.0

Security-Related Tasks

Tasks for Managing Servers and Applications

Tasks for Administering Servers

**Related Resources**

- For information about configuring Web sites as well as the FTP, NNTP, and SMTP services, see Configuring Internet Sites and Services.
- For information about configuring IIS 6.0 programmatically, see IIS 6.0 Administration Scripts, Tips, and Tricks.

# Overview of Common Administrative Tasks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Although this appendix is a reference for the common tasks you frequently perform when administering IIS as your Web server, most tasks in this appendix do not include guidelines for their use or detailed explanations. A few tasks include brief conceptual information, although they are the exception. However, all topics include cross-references to IIS Operations Guide topics, and to other topics and appendices in this reference so that you can quickly find additional information on a task.

Table A.1 outlines the tasks that are included in this appendix.

**Table A.1 Overview of the Tasks in This Appendix**

| Task | Describes How To |
|---|---|
| **Important First Tasks In IIS 6.0** | |
| Starting IIS Manager | Open IIS Manager (provides three options). |
| Starting and stopping services and sites | Start and stop IIS services and sites. |
| Enabling dynamic content services | Enable dynamic content services, such as Active Server Pages and ASP.NET (includes information about default installations). |
| Creating a Web FTP site | Create Web site configurations, and install and use File Transfer Protocol (FTP) services. |
| Creating virtual directories | Create virtual directories within Web and FTP sites. |
| **Tasks New to IIS 6.0** (tasks that are new or significantly different than they were in IIS 5.0) | |
| Creating and isolating applications | Create and manage applications. |
| Creating application pools | Group Web applications into application pools. |
| Configuring recycling | Periodically restart worker processes assigned to an application pool. |
| Backing up and restoring metabase configurations | Save metabase and application configurations, including portable backups. |
| Saving and copying site configurations | Back up your site configurations. |
| **Security-related Tasks** | |
| Setting authentication settings | Set up and use user authentication methods. |
| Obtaining and backing up server certificates | Set up and use Secure Sockets Layer (SSL) certification on your sites. |
| Controlling access to applications | Help reduce the attack surface of your applications with permissions and restrictions. |
| **Tasks for Managing Servers and Applications** | |
| Hosting multiple Web sites | Create and host multiple Web sites. |

| | |
|---|---|
| Redirecting Web sites | Automatically direct users to the correct page on your site. |
| Assigning resources to applications | Control the amount of resources an application uses. |
| **Tasks for Administering Servers** | |
| Administering servers from the command line | Use powerful scripting and programming tools to access and configure settings. |
| Administering servers remotely | Use tools to remotely administer your sites. |

When performing these tasks — especially security-related tasks — be sure that you are familiar with best practice guidelines and that you apply those guidelines to these tasks.

To perform most of the tasks described in this appendix, you must be a member of the Administrators group on the local computer, or you must be delegated the appropriate authority. If you log on to your computer as a member of the Administrators group, you might make your system vulnerable to malicious programs that could cause security risks. Instead, use the **Run as** feature to perform administrative tasks so that you do not need to log on to your computer as a member of the Administrators group. Using this feature, you can open and execute a program that uses a different account and security context than the one you logged on with.

You can use **Run as** through the user interface (UI) or as a command-line tool. The **Run as** feature that is built into the UI is a shortcut that you access by right-clicking some programs (files with the .exe file name extension), some Control Panel items (files with the .cpl file name extension), and Microsoft Management Console (MMC) (files with the .msc file name extension) snap-ins.

**To use the Run as feature to run IIS Manager as an administrator**

1. From the **Start** menu, point to **Administrative Tools**, right-click **Internet Information Services (IIS) Manager**, and then click **Run as**.

The **runas** command provides the same capabilities as the built-in **Run as** feature.

**To use the runas command to run IIS Manager as an administrator**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following:

   **runas /User:** *Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**

**To use the runas command to run a command-line script as an administrator**

1. From the **Start** menu, click **Run**.
2. In the **Open** box, type **cmd**, and then click **OK**.
3. At the command prompt, type the following:

   **runas /profile /User:** *MyMachine* **** *Administrator* **cmd**

   A new command window, which has administrative rights, opens.
4. In the new command window, type the following at the command prompt:

   **cscript.exe** *ScriptName* (include the script's full path and any parameters)

For more information about using the **Run as** feature or the **runas** command to perform procedures, see "Using Run as", "Runas", and "Create a shortcut using the runas command" in Help and Support Center for Microsoft® Windows® Server 2003.

# Important First Tasks after Installing IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you install IIS, if you want to quickly build a few Web sites and virtual directories, this section introduces the following first steps for these basic tasks:

- Starting IIS Manager: Describes three ways to open IIS Manager.
- Starting and Stopping Services and Sites: Describes how to start and stop IIS services.
- Enabling and Disabling Dynamic Content: Describes how to enable Web service extensions, such as Active Server Pages (ASP) and Microsoft® ASP.NET, so you can serve dynamic content. Includes information about default installations.
- Creating a Web Site or an FTP Site: Describes how to create Web site configurations, and how to install and use FTP services.
- Creating IIS 6.0 Virtual Directories: Describes how to create virtual directories within Web and FTP sites.

# Starting IIS Manager

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS Manager is a graphical user interface (GUI) for configuring your application pools or your Web, FTP, Simple Mail Transfer Protocol (SMTP), or Network News Transfer Protocol (NNTP) sites. You can use IIS Manager to configure IIS security, performance, and reliability features. For example, you can add or delete sites; start, stop, and pause sites; back up and restore server configurations; and create virtual directories for better content management.

In earlier versions of IIS, this tool was called the Internet Service Manager. For information about starting IIS Manager, see IIS Manager.

# Starting and Stopping Services and Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Infrequently, you might make configuration changes in IIS 6.0 that require you to restart IIS before the changes can take effect. For example, if you change the application isolation mode in which your server is running, such as when you change from worker process isolation mode to IIS 5.0 isolation mode or vice versa, you need to restart IIS. If you make this configuration change by using IIS Manager, you are prompted to restart IIS after you click **OK** to confirm the change. If you make this configuration change by using a command-line utility, such as Adsutil.vbs, you can use the **IISReset** command-line utility to complete the change. Both methods — using the Restart IIS command in IIS Manager or using a command-line utility — allow you to stop, start, and restart IIS Internet services, as well as restart your computer.

When you restart the Internet service, all sessions connected to your Web server (including Internet, FTP, SMTP, and NNTP) are dropped. Any data held in Web applications is lost. All Internet sites are unavailable until Internet services are restarted. For this reason, avoid restarting, stopping, or rebooting your server.

For a list of features designed to improve IIS reliability and remedy the need to restart IIS, see the Alternatives to Restarting IIS section in the Restarting IIS topic in the IIS 6.0 Operations Guide, which is accessible from IIS Manager.

For information about restarting IIS services, and starting, stopping, or pausing individual sites, see Starting and Stopping Services.

**Saving Your Configuration to Disk**

As a safeguard, if you must stop or restart IIS, save your configuration to disk before you perform the restart. Your configuration is automatically saved if you enable the edit-while-running feature (this feature is not enabled by default). For more information about the edit-while-running feature, see IIS Metabase.

Alternatively, you can manually save your configuration to disk. For information about manually saving your configuration to disk, see Saving Configurations.

**If You Receive an Error Stating That IISReset Is Disabled**

If the **IISReset** command-line utility is disabled, then the command-line or IIS Manager calls that require IISReset.exe fail and return an error stating that **IISReset** is disabled. Actions that fail include the **Restart IIS** command in IIS Manager and Service Control Manager (SCM) recovery configuration actions that use the **IISReset** command-line utility (for example, the default IIS Admin SCM recovery path). However, SCM recovery actions that do not use the **IISReset** command-line utility continue to function (for example, the default World Wide Web Publishing Service [WWW service] SCM recovery path that restarts the WWW service).

# Enabling and Disabling Dynamic Content

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To help reduce the attack surface of systems, IIS 6.0 is not installed by default on the Microsoft® Windows® Server 2003, Standard Edition; Microsoft® Windows® Server 2003, Enterprise Edition; and Microsoft® Windows® Server 2003, Datacenter Edition. After installing these products, administrators must manually install IIS 6.0.

When you perform a clean installation of IIS 6.0, the default settings help protect your system from malicious users and attackers. When you install IIS 6.0, it is locked down — only request handling for static Web pages (HTTP content, such as .htm and .html files ) is enabled, and only the World Wide Web Publishing Service (WWW service) is installed. The request handlers that process dynamic content are disabled, which means that features like ASP, ASP.NET, server-side includes (SSI), FrontPage® 2002 Server Extensions from Microsoft, and Web Distributed Authoring and Versioning (WebDAV) do not work by default.

You can configure the request handlers (for example, Internet Server API [ISAPI] extensions or Common Gateway Interface [CGI] programs), which are called Web service extensions, by using the Web Service Extensions node in IIS Manager or by using the command-line script lisext.vbs, which is stored in the *systemroot*\System32 folder. You can individually enable or disable a Web service extension if it is registered in the Web Service Extensions node in IIS Manager.

For more information about using IIS Manager to administer dynamic content, including step-by-step documentation for adding new Web service extensions, allowing an application to call a Web service extension, or disabling all Web service extensions, see Enabling and Disabling Dynamic Content in IIS 6.0.

For more information about using the lisext.vbs command-line script to administer dynamic content, such as enabling, disabling, or listing Web service extensions and their files, see Managing Applications and Web Service Extensions Using Scripts.

# Enabling ASP Pages

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Active Server Pages (ASP) is a server-side scripting environment that you can use to create dynamic and interactive Web pages, and build powerful Web applications. When the server receives a request for an ASP file, it processes server-side script code contained in the file to build the Hypertext Markup Language (HTML) Web page that is sent to the browser. In addition to server-side script code, ASP files can contain HTML (including related client-side scripts), as well as calls to Component Object Model (COM) components that perform a variety of tasks, such as connecting to a database or processing business logic.

For information about enabling ASP, see Enabling ASP Pages in IIS 6.0. For an introduction to ASP concepts, see About ASP. For information about creating ASP pages and developing Web applications, see the Web Application Guide.

# Installing and Enabling ASP.NET

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Microsoft ASP.NET is a unified Web development platform that provides the services necessary for developers to build enterprise-class Web applications. Although ASP.NET is largely syntax compatible with ASP, it also provides a new programming model and infrastructure for more secure, scalable, and stable applications. You can augment your existing ASP applications by incrementally adding ASP.NET functionality to them.

ASP.NET is a compiled, .NET-based environment; you can author applications in any .NET compatible language, including Microsoft® Visual Basic® .NET, Microsoft® Visual C#®, and Microsoft® JScript® .NET.

For information about enabling ASP.NET, and installing ASP.NET on a server running Windows Server 2003, see Installing ASP.NET.

# Creating a Web Site or an FTP Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS creates a default Web site on your computer during installation. You can use the *LocalDrive*:\Inetpub\Wwwroot directory to publish your Web content, or you can create any directory or virtual directory you choose. Because the FTP service is not installed by default, you must first install and start the File Transfer Protocol (FTP) service to create an FTP site.

Creating a Web or FTP site by using IIS Manager does not create content but merely creates a directory structure and configuration files from which to publish the content. For information about creating Web sites and FTP sites, and installing FTP services, see Creating Web or FTP Sites.

For more information about configuring Web and FTP sites, see Configuring Internet Sites and Services, Web Site Setup and FTP Site Setup.

# Creating IIS 6.0 Virtual Directories

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In most cases, the content you publish to your Web or FTP site is located in a root or home directory on your computer, such as *LocalDrive*:\Inetpub\Wwwroot\. However, you might need to publish content that is located elsewhere, such as on a remote computer.

To publish from a directory that is not contained within your home or root directory, you can create a virtual directory. A virtual directory is a directory that is not contained in the home directory but appears to client browsers as though it were.

For information about creating Web and FTP virtual directories, see Creating Virtual Directories in IIS 6.0.

# Tasks New to IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you upgrade from an earlier version of IIS, you need to learn the tasks that are new to IIS 6.0. This section introduces the following basic tasks that are new, or significantly different, in IIS 6.0:

- Creating and Isolating Applications: Describes how to create and manage applications.
- Creating an Application Pool: Describes how to group Web applications into application pools.
- Configuring Recycling: Describes how to periodically restart worker processes assigned to an application pool.
- Backing Up and Restoring the Metabase Using Backup Files: Describes how to save metabase and application configurations.
- Saving and Copying Site Configurations: Describes how to copy site configurations for use on new sites.

# Creating and Isolating Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

To create an application, you designate a directory as the starting point (application root) for the application. You can then set properties for the application. Each application can have a friendly name; this name appears in IIS Manager and gives you a way to distinguish between applications.

Web sites are root-level applications by default. When you create a Web site, a default application is created at the same time. You can use this root-level application, remove it, or replace it with a new application by removing it and creating a new application.

To isolate an application means that you configure it to run in a process (memory space) separate from the Web server and other applications. You can run IIS 6.0 in one of two application isolation modes: *worker process isolation mode* or *IIS 5.0 isolation mode*. IIS cannot run both application isolation modes simultaneously on the same computer. If you have applications that require different modes, you must run them on separate computers.

Worker process isolation mode is the default application isolation mode that the server runs in on a clean installation. In this mode, you can isolate an application by adding it to an application pool that includes isolation settings.

IIS 5.0 isolation mode is provided for applications that depend on specific features and behaviors of IIS 5.0. Use this mode only if an application has a compatibility issue when it runs in worker process isolation mode and you cannot resolve the problem.

For more information about creating and isolating applications, see Creating Applications in IIS 6.0.

# Creating an Application Pool

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you run IIS 6.0 in worker process isolation mode, you can group Web applications into application pools. An *application pool* is a grouping of URLs routed to one or more worker processes that share the same configuration. Application pools allow you to apply specific configuration settings to groups of applications and the worker processes servicing those applications. Any Web site, Web directory or virtual directory can be assigned to an application pool.

By creating new application pools and assigning Web sites and applications to them, you can make your server more efficient and reliable. In addition, your other applications are always available, even when the applications in the new application pool terminate.

For more information about creating application pools and configuring worker processes, see Creating Application Pools in IIS 6.0, Configuring Application Pool Identity with IIS 6.0, and Running IIS 6.0 as an Application Server.

# Configuring Recycling

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In worker process isolation mode, you can configure IIS to periodically restart the worker processes that are assigned to an application pool in order to manage faulty Web applications. Recycling keeps problematic applications running smoothly, especially when it is not practical to modify the application code. Recycling ensures that application pools remain healthy and that system resources are recovered.

You can configure worker processes to restart based on one of several options, including elapsed time, number of requests served, scheduled times, and memory usage; or you can configure worker processes to start on demand. In IIS 6.0, worker process recycling is available only when IIS is running in worker process isolation mode.

For more information about configuring recycling for worker processes, see Configuring Worker Processes for Recycling, and Running IIS 6.0 as an Application Server.

# Backing Up and Restoring the Metabase Using Backup Files

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Metabase backup files provide a way to restore your metabase configuration and schema data if your metabase becomes corrupted. You can create backup files by using IIS Manager or an administration script. The backup files are copies of the metabase configuration file (MetaBase.xml) and the matching metabase schema file (MBSchema.xml). Use the metabase configuration backup and restore feature to restore the metabase from backup files.

You can create two types of metabase backups:

- Portable backups. When you create a portable backup, you provide a password that is used by IIS to encrypt the backup files. IIS encrypts the password and stores it in the header of the backup file. Only the administrator password and secure properties within the backup files are encrypted; all other information within the backup files are plaintext. After the backup file is encrypted, you cannot change the password of the backup file.

- Non-portable backups. When you create a non-portable backup, you do not supply a password. Therefore, IIS encrypts non-portable backup files with a blank password, which allows any member of the Administrators group to restore the metabase by using non-portable backup files. Only the blank password and secure properties are encrypted; all other information within the backup file is plaintext.

You can restore a metabase backup to the computer on which the backup was made or to a different computer that is running Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; and Windows® Server 2003, Web Edition. However, before you restore a metabase backup from one computer to another, you must first delete the machine-specific information from the metabase file. For more information about restoring a metabase backup to a different computer, see Machine-Specific and System-Specific Information in the IIS 6.0 Operations Guide, which is accessible from IIS Manager.

Before you create a backup of the metabase, consider the following:

- A non-portable backup can only be restored to the computer on which the backup was made.

- The metabase is locked while the backup is in progress.

- Backup files contain only configuration data; they do not include your content.

To back up your content, use the Windows Backup feature. For more information about Windows Backup, see "Backing up and Restoring Data" in Help and Support Center for Windows Server 2003.

For more information about backing up and restoring the metabase, see Backing Up and Restoring the Metabase in IIS 6.0.

# Saving and Copying Site Configurations

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After your site and applications are running the way you want, you can save all or part of the configurations for a backup copy, or for import and export to other sites or computers. You can use the following procedure to save site configurations for Web or FTP sites, as well as Web or FTP virtual directories.

IIS automatically makes a backup copy of the metabase configuration and schema files each time you make changes to the metabase. Backup files contain only configuration data; they do not include your content (.asp files, .htm files, .dll files, and so on). You can also create backup files on demand, or create backup copies of individual site or application configurations and then export and import them to and from other sites or computers.

For more information about backing up the metabase, see Saving Configurations and IIS Metabase.

# Security-Related Tasks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 and Windows Server 2003 provide several ways to help secure your application servers and their content. This section provides information about the following security-related topics:

- Setting Web Site Authentication: Describes how to set up Web site authentication for your Web sites.
- Setting FTP Site Authentication: Describes how to set up FTP site authentication to validate users who request access to your FTP sites.
- Obtaining and Backing Up Server Certificates: Describes how to set up Secure Sockets Layer (SSL) certification on your sites. SSL certificates enable Web servers and users to authentication each other before establishing a connection.
- Controlling Access to Applications: Describes how to reduce the attack surface of your applications with permissions and restrictions, and how to control which users and computers are allowed to access your Web server and its resources.

Before you do any of these security-related tasks, be sure that you are familiar with best practice guidelines and that you apply those guidelines to these tasks. If you are new to using IIS or if you are unfamiliar with any of the following security-related tasks, be sure to read the cross-references to additional information that are provided in this section.

# Setting Web Site Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can require users to provide a valid Windows user account name and password before they access any information on your server. This identification process is called *authentication*. Authentication, like many of the features in IIS, can be set at the Web site, directory, or file level.

This section contains step-by-step procedures for configuring Web site authentication. For information about configuring FTP sites, see Setting FTP Site Authentication later in this appendix.

To set Web authentication, choose from the following authentication methods:

- Anonymous authentication. This authentication method gives users access to the public areas of your Web site without prompting them for a user name or password.
- Basic authentication. This authentication method requires a previously assigned Windows account user name and password, also known as *credentials*.
- Digest authentication. This authentication method offers the same functionality as Basic authentication, while providing an additional level of security because the user's credentials are not sent over the network in plaintext.
- Advanced Digest authentication. This authentication method offers similar functionality to Digest authentication; however, collects user credentials and stores them on the domain controller as an MD5 hash, or message digest. Advantages of this authentication are that the worker process does not need to run as local system and the user password is not stored as plaintext on the domain controller. This authentication method requires a Windows Server 2003 domain controller infrastructure.
- Integrated Windows authentication. This authentication method collects information through a method where the user name and password are hashed before being sent across the network.
- Certificate authentication. This authentication method adds SSL security through client or server certificates. For information about this type of authentication, see Obtaining and Backing Up Server Certificates.
- .NET Passport authentication. This authentication method provides a single sign-in service that is HTTP cookie-based.

# Configuring Anonymous Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Anonymous authentication gives users access to the public areas of your Web or FTP site without prompting users for a user name or password. When a user attempts to connect to your public Web or FTP site, your Web server assigns the connection to the Windows user account IUSR_*computername*, where *computername* is the name of the computer on which IIS is running. By default, the IUSR_*computername* account is included in the Users and Guests user groups.

For information about enabling Anonymous authentication and changing the account used for Anonymous authentication, see Anonymous Authentication in IIS 6.0.

# Configuring Basic Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Basic authentication method transmits user names and passwords across the network in an unencrypted form. You need to use SSL, in combination with Basic authentication, to encrypt user account information that is transmitted across the network.

Configuration settings made at the Web Sites folder level can be inherited by all Web sites unless you specify otherwise at the individual Web site level. Enabling Basic authentication does not automatically configure your Web server to authenticate users. You must also create Windows user accounts and set NTFS permissions.

For information about enabling Basic authentication, see Enabling Basic Authentication and Configuring the Realm Name.

# Configuring Digest Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Digest authentication offers the same functionality as Basic authentication; however, Digest authentication provides a means to help ensure that user credentials are not sent across the network in plaintext. Digest authentication transmits credentials across the network as an MD5 hash, or message digest, where the original user name and password cannot be deciphered from the hash. Digest authentication is available to WebDAV directories.

Digest authentication is enabled by default for upgrades from an earlier version of IIS. If you need to enable Digest authentication on a server running IIS 6.0, do the following:

1. Enable Digest authentication for Windows domain servers.

2. Configure the realm name.

**Configuring the Realm Name**

In addition to using IIS Manager to enable Digest authentication on a Windows domain server, you can use scripting to configure the realm name at any level of the metabase, as shown in the Table A.2.

If a child key in the metabase is not configured with a realm name, that child key inherits the realm name from the next parent key that has the realm name configured. If the realm name is not configured, IIS sends its own computer name as the realm name. If IIS sends its own name as the realm name and IIS is not running on a Windows Server 2003 domain controller with Active Directory® directory service, Digest authentication fails. As a best practice, avoid running IIS on a domain controller; whenever possible, physically separate a server that is running IIS from a domain controller.

**Table A.2 Configuring the Realm Name at Any Level of the Metabase**

| Metabase Level | Description |
|---|---|
| W3SVC | The W3SVC level, also known as the IISWebService level, is the highest level in the metabase where Digest authentication can be configured. Lower levels that do not have specific configuration settings inherit configurations set at this level. |
| W3SVC/*n* | The W3SVC/*n* level, also known as the IISWebService level, is a specific Web site, where *n* is the number of the site. Sites are numbered starting at 1. The default Web site is 1. |
| W3SVC/*n*/root | The W3SVC/*n*/Root level, known as the IISWebVirtualDir level, is the starting point for a Web Site, where *n* is the number of the site. |
| W3SVC/*n*/root/vdir | The W3SVC/*n*/Root/WebVirtualDir level, known as the IISWebVirtualDir level, is a virtual directory within a Web Site, where *n* is the number of the site. |
| W3SVC/*n*/root/vdir/webdir | The W3SVC/*n*/Root/WebVirtualDir/WebDir level, also known as the IISWebDirectory level, is a physical directory within a virtual directory within a Web site, where *n* is the number of the site. |
| W3SVC/*n*/root/vdir/file | The W3SVC/*n*/Root/Vdir/file level is an individual file within the W3SVC/*n*/Root/WebVirtualDir level, where *n* is the number of the site. |
| W3SVC/*n*/root/vdir/webdir/file | The W3SVC/*n*/Root/Vdir/file level is an individual file within the W3SVC/*n*/Root/WebVirtualDir/WebDir level, where *n* is the number of the site. |

You can configure either single or multiple realm names on a server running IIS. You might want to configure multiple realm names if the domains do not have a trusted relationship. If you configure multiple realm names, you must configure them at different levels of the metabase.

For information about enabling Digest authentication and configuring the realm name, see Configuring Digest Authentication in IIS 6.0.

# Configuring Advanced Digest Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Under Advanced Digest authentication, user credentials are stored on the domain controller as an MD5 hash. Because credentials are stored in Active Directory as an MD5 hash, user passwords cannot be discovered by anyone with access to the domain controller, not even by the domain administrator. Advanced Digest authentication is available to WebDAV directories.

In IIS 6.0, Advanced Digest authentication is preferred over Digest authentication; however, you can use Digest authentication where your systems do not meet the requirements for Advanced Digest authentication.

Advanced Digest authentication is enabled by default on a clean installation of IIS 6.0. If you need to enable Advanced Digest authentication on a server running IIS 6.0, do the following:

1. Enable Digest authentication for Windows domain servers.
2. Configure the realm name.
3. Set the **UseDigestSSP** metabase property to **TRUE**.

| **Important** |
| --- |
| If you perform the two procedures that precede this note but do not configure the **UseDigestSSP** metabase key, you are using Digest authentication, not Advanced Digest authentication.<br><br>For information about enabling Advanced Digest authentication, see Advanced Digest Authentication in IIS 6.0. |

# Configuring Integrated Windows Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Integrated Windows authentication (formerly called NTLM, and also referred to as Windows NT Challenge/Response authentication) is a form of authentication that hashes the user name and password before they are sent across the network. When you enable Integrated Windows authentication, the client submits the password through a cryptographic exchange with your Web server that involves hashing. Integrated Windows authentication is the default authentication method used in Windows Server 2003.

Integrated Windows authentication has the following limitations:

- Only Microsoft® Internet Explorer version 2.0 and later support this authentication method.
- This authentication method might not work over HTTP proxy connections.

Therefore, Integrated Windows authentication is best suited for an intranet environment, where both user and Web server computers are in the same domain and where administrators can ensure that every user has Internet Explorer version 2.0 or later.

For information about enabling Integrated Windows authentication, see Configuring Integrated Windows Authentication in IIS 6.0.

# Enabling .NET Passport Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can enable Microsoft® .NET Passport authentication on a Web site by using IIS Manager. When .NET Passport is enabled, requests coming into IIS must contain .NET Passport credentials on either the query string or within a cookie. The credentials also have to be valid, meaning the ticket has not expired. If IIS does not detect .NET Passport credentials, requests are redirected to the .NET Passport sign-in page.

.NET Passport uses cookies, which contain information that can be compromised. However, you can use .NET Passport authentication over a SSL connection, which reduces the potential of replay attacks.

For information about enabling .NET Passport authentication, see Enabling .NET Passport Authentication in IIS 6.0.

# Setting FTP Site Authentication

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Based on your security requirements, you can select an IIS authentication method to validate users who request access to your FTP sites. To set FTP site authentication, choose from the following authentication methods:

- Anonymous FTP authentication. This authentication method gives users access to the public areas of your FTP site without prompting them for a user name or password.

- Basic FTP authentication.This authentication method requires users to log on with a user name and password corresponding to a valid Windows user account.

You cannot use Digest or Integrated Windows authentication with FTP sites, and you must set available authentication settings at the site level for FTP sites.

**Enabling Anonymous FTP Authentication**
You can configure your FTP server to allow anonymous access to FTP resources. If you select Anonymous FTP authentication for a resource, all requests for that resource are accepted without prompting the user for a user name or password. This is possible because IIS automatically creates a Windows user account called IUSR_*computername*, where *computername* is the name of the server on which IIS is running. This is very similar to Web-based Anonymous authentication. If Anonymous FTP authentication is enabled, IIS always try to use it first, even if you enable Basic FTP authentication.

For information about enabling Anonymous FTP authentication, see Configuring Anonymous FTP Authentication. For more information about creating Windows user accounts and setting NTFS permissions, see Securing Files with NTFS Permissions and Setting NTFS Permissions for Directories or Files.

**Enabling Basic FTP Authentication**
To establish an FTP connection with your Web server by using Basic FTP authentication, users must log on with a user name and password corresponding to a valid Windows user account. If the FTP server cannot verify a user's identity, the server returns an error message. Basic FTP authentication transmits the user name and password across the network in an unencrypted form.

For information about enabling Basic FTP authentication, see Configuring Basic FTP Authentication.

# Obtaining and Backing Up Server Certificates

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Server certificates contain information used in establishing identities over a network, which is a process called *authentication*. Similar to conventional forms of identification, certificates enable Web servers and users to authenticate each other before establishing a SSL connection.

*Server certificates* contain information about the server that allows the client to positively identify the server before sharing sensitive information. *Client certificates* contain personal information about the clients requesting access to your site that allow you to positively identify them before allowing them access to the site.

Certificates include *keys* used in establishing an SSL secure connection. A *key* is a unique value used to authenticate the server and the client in establishing an SSL connection. A *public key* and a *private key* form an SSL *key pair*. Your Web server uses this key pair to negotiate a secure connection with the user's Web browser to determine the level of encryption required for securing communications.

You can obtain a server certificate in one of two ways: issue your own server certificate or obtain a server certificate from a certification authority.

To request and install your own server certificate, use the Web Server Certificate Wizard to create a customizable service for issuing and managing certificates. You can create server certificates for the Internet or for corporate intranets, giving your organization complete control over certificate management policies.

To obtain a server certificate from a certification authority, follow these steps:

1. Do either of the following:

   - Find a certification authority that provides services that meet your business needs and request a server certificate.

     –Or–

   - Use the Web Server Certificate Wizard to create a certificate request that you can send to the certification authority.

2. After the certificate is processed and returned to you, use the Web Server Certificate Wizard to install the certificate.

For more information about SSL certificates, including how to obtain, install, and back up server certificates, see Certificates_IIS_SP1_Ops. For more information about Microsoft Certificate Services, see "Certificate Services" in Help and Support Center for Windows Server 2003.

# Controlling Access to Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can control which users and computers are allowed to access your Web server and its resources. You can use both the NTFS file system and Web server security features to assign users specific permissions to directories and files. You can also use IP address restrictions to limit access by specific computers or groups of computers.

- Securing your files with NTFS permissions. With the NTFS file system, you can limit access to your Web server's files and directories. You can configure the file and directory permissions that set the access level assigned to a particular user account or user group. For example, you can configure your Web server to enable a specific user to view and execute a file, while excluding all other users from accessing that file.

- Securing your Web site with Web permissions. Web permissions, unlike NTFS permissions, apply to all users accessing your Web sites. NTFS permissions apply only to a specific user or group of users with a valid Windows account. NTFS controls access to physical directories on your server, whereas Web permissions control access to virtual directories on your Web site.

- Restricting access to your Web site by using IP addresses. With IP address restrictions, you can configure your Web server to assign or deny specific computers, groups of computers, or domains access to Web sites, directories, or files. For example, if your intranet server is connected to the Internet, you can prevent Internet users from accessing your Web server by assigning access only to members of your intranet, and explicitly denying access to outside users.

# Securing Your Files with NTFS Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can control access to your Web site's directories and files by setting NTFS access permissions. You can use NTFS permissions to define the level of access that you want to assign to specific users, and groups of users. Correct configuration of file and directory permissions is crucial for preventing unauthorized access to your resources. For information about securing a Web site by using NTFS permissions, see Securing Files with NTFS Permissions.

With NTFS permissions, you also have the choice of assigning special permissions to groups or users. Special permissions are permissions on a more detailed level. For better management, assign broad-level permissions to users or groups, where it is applicable. For descriptions of permissions, see "Permissions for Files or Folders" in Help and Support Center for Windows Server 2003.

# Securing Your Web Site with Web Permissions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can also use Web permissions to help secure your Web site. When you configure your Web site's access permissions for specific sites, directories, and files, keep in mind the following:

- Web permissions are not meant to be used in place of NTFS permissions; instead, use Web permissions with NTFS permissions.

  Although you can use Web permissions with both the NTFS and FAT file systems, use NTFS whenever possible.

- Unlike NTFS permissions, Web permissions affect everyone who tries to access your Web site.

- Disabling permissions restricts all users. For example, disabling the Read permission restricts all users from viewing a file, regardless of the NTFS permissions applied to those users' accounts. However, enabling the Read permission can allow all users to view that file, unless NTFS permissions that restrict access have also been applied.

- If Web permissions conflict with NTFS permissions for a directory or file, the more restrictive settings are applied. For example, if you set both IIS and NTFS permissions, the permissions that explicitly deny access take precedence over permissions that grant access.

- If you want to set permissions for a WebDAV directory, keep in mind the following:

  - You must enable WebDAV before you can publish in a WebDAV directory.

  - You must turn off Anonymous access to your WebDAV directory. Protect your WebDAV directory by using one of the following authentication methods in IIS 6.0, such as Integrated Windows authentication and the discretionary access control lists (DACL) in NTFS. For more information about WebDAV security, see Managing WebDAV Security in the IIS 6.0 Operations Guide, which is accessible from IIS Manager.

For information about setting permissions for Web content (including WebDAV), see Securing Sites with Web Site Permissions.

# Restricting Access to Your Web Site by Using IP Addresses

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

You can configure your Web site to grant or deny specific computers, groups of computers, or domains access to Web sites, directories, or files. For example, if your intranet server is connected to the Internet, you can prevent Internet users from accessing your Web server by assigning access only to members of your intranet, and explicitly denying access to outside users.

IP address restrictions apply only to Internet Protocol version 4 (IPv4) addresses.

For information about granting or denying access to computers and domains, see Granting or Denying Access to a Computer and Granting or Denying Access to a Domain.

# Tasks for Managing Servers and Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use the following tasks in this section to manage your servers and applications:

- Using Host Header Names to Host Multiple Web Sites: Describes how to create and host multiple Web sites.
- Redirecting Web Sites: Describes how to automatically direct users to the correct page on your site.
- Assigning Resources to Applications: Describes how to control the amount of resources an application uses.

# Using Host Header Names to Host Multiple Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS supports multiple Web sites on a single server. To create and host multiple Web sites, you must configure a unique identity for each site on the server. To assign a unique identity, distinguish each Web site with at least one of three unique identifiers: a host header name, an IP address, or a TCP port number.

One method for providing each site with a unique identifier is to use IIS Manager to assign multiple host header names. Browsers must comply with HTTP 1.1 to support the use of host header names. Microsoft® Internet Explorer 3.0, Netscape Navigator 2.0, and later versions of both browsers support host header names.

| **Important** |
| --- |
| Avoid assigning a host header name to the Default Web Site; instead, use an IP address of All Unassigned, a TCP port of 80, and no host header name. |

Table A.4 briefly describes and compares the three ways that you can uniquely identify your Web sites.

**Table A.4 Ways You Can Uniquely Identify Multiple Web Sites**

| Web Site Identifier | Description |
| --- | --- |
| Host header name | Recommended for most situations. Requires that your computer or network use a name resolution system. Organizations typically use DNS name resolution. |
| Unique IP Address | Used primarily for Internet services that host SSL on the local server. Typically, only large corporations and Internet service providers (ISPs) obtain and maintain multiple IP addresses. |
| Nonstandard TCP port number | Generally not recommended. Can be used for private Web site development and testing purposes but rarely used on production Web servers because users must add the port to the URL and the port must be opened on the firewall. |

Choose one method of uniquely identifying Web sites for each server. Using one method for each server improves performance by optimizing cache and routing lookups. Conversely, using any combination of host headers, unique IP addresses, or non-standard port numbers to identify multiple Web sites degrades the performance of all Web sites on the server.

If you use host header names to identify a new Web site, select a unique name as follows:

- **On the Internet.** The host header must be a publicly available DNS name, such as support.microsoft.com. Register a public DNS name with an authorized Internet name authority. For more information about using DNS names to identify your site, see Domain Name Resolution in the IIS 6.0 Operations Guide, which is accessible from IIS Manager.

- **On a private network.** The host header can be an intranet site name. To resolve host names to an IP address, register the host header name with your intranets DNS administrator. Alternatively, you can resolve host names to an IP address by using a locally stored database file called the Hosts file, which is located in the *systemroot*\System32\Drivers\Etc folder. The following is an example of the contents of a Hosts file:

```
#
# Table of IP addresses and host names
#
# IP Address     TCP Port   Host Header Name
#
192.168.0.100    80         www.example1.com
192.168.0.100    80         example1.com
```

Use Windows Explorer to create a home directory for your Web site content. Create subdirectories to store HTML pages, image files, and other content as needed. To organize home directories for multiple Web sites on the same server, you can create a top-level directory for storing all home directories, and then create subdirectories for each site.

**Important** You must be a member of the Administrators group on the local computer to perform the following procedure or procedures. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run IIS Manager as an administrator. At a command prompt, type **runas /user:***Administrative_AccountName* **"mmc %systemroot%\system32\inetsrv\iis.msc"**.

**To assign multiple host header names to a single Web site**

1. In IIS Manager, expand the local computer, click the **Web Sites** folder, right-click the Web site for which you want to assign a host header name, and then click **Properties.**
2. Click the **Web Site** tab.
3. In the **IP Address**list, select the IP address that you want assigned to this Web site.
4. Click **Advanced**.
5. Under **Multiple identities for this Web Site**, select the IP address, and then click **Edit**.
6. In the **Host Header Value** box, type the host header name (for example, www.contoso.com).
7. Click **OK** twice.

For more information about using host header names to host multiple Web sites, see Hosting Multiple Web Sites on a Single Server. Also see Knowledge Base article 324287, Use Host Header Names to Configure Multiple Web Sites in Internet Information Services 6.0.

# Redirecting Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When a browser requests a page or program on your Web site, the Web server locates the page identified by the URL and returns it to the browser. When you move a page on your Web site, you cannot always correct all the links that refer to the old URL. To make sure that browsers can find the page at the new URL, you can instruct the Web server to *redirect* the browser to the new URL.

You can redirect requests for files in one directory to a different directory, to a different Web site, or to another file in a different directory. When the browser requests the file at the original URL, the Web server instructs the browser to go to the new URL.

For information about redirecting requests to other Web sites or directories, redirecting requests to a single file, or redirecting requests to a program, see Redirecting Web Sites in IIS 6.0.

# Assigning Resources to Applications

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

There are several ways to assign, or limit, the amount of resources an application can use. You can use application property sheets to control performance, cache, and process options of individual applications, or you can use *bandwidth throttling* to limit resources for an entire site or application pool. A third alternative is to control access to your application through the NTFS file system and Web server security features.

For information about assigning resources by using property sheets, globally assigning resources by using bandwidth throttling, or assigning resources to an individual Web site by using bandwidth throttling, see Assigning Resources to Applications in IIS 6.0 and Throttling Bandwidth.

# Tasks for Administering Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 and Windows Server 2003 provide several tools to help you administer your servers. This section provides the following task-based information:

- Administering Servers from the Command Line: Describes how to use the following powerful scripting and programming tools to access and configure settings.
    - Supported command-line scripts. Supplied scripts for the IIS Windows Management Instrumentation (WMI) provider to manage and set IIS metabase configurations.
    - Adsutil.vbs IIS administration utility. Uses VBScript and Active Directory Service Interfaces (ADSI) as a learning tool to manipulate the IIS configuration.
- Administering Servers Remotely: Describes how to use the following tools to remotely administer your sites.
    - IIS Manager. Remotely connect to and administer an intranet server.
    - Terminal Services. After you use Terminal Services to connect to the server that is running IIS, you can use IIS Manger on the Web server as if you were logged on locally.
    - Remote Administration (HTML) tools. Administer your IIS Web server from Web browsers on your intranet.
    - Supported command-line scripts. Use with the IIS WMI provider to remotely manage an IIS machine. Each command-line script supports the **/s** parameter, which you can use to specify the remote server against which you want to perform the command.

# Administering Servers from the Command Line

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS provides powerful scripting and programming tools that you can use to access and configure settings from within a command-line script or compiled application. You can use these tools to create, delete, start, stop, pause, and list various sites and applications, as well as to copy, import, and export configurations.

**Important** You must be a member of the Administrators group on the local computer to run scripts and executables. As a security best practice, log on to your computer by using an account that is not in the Administrators group, and then use the **runas** command to run your script or executable as an administrator. At a command prompt, type **runas /profile /user:***MyComputer******Administrator* **cmd** to open a command window with administrator rights and then type **cscript.exe** *ScriptName* (include the script's full path and any known parameters).

IIS provides the following command-line tools:

- Supported command-line scripts. Use the supplied scripts for the IIS WMI provider to manage and set IIS metabase configurations.

- Adsutil.vbs IIS administration utility. Uses Microsoft® Visual Basic® Scripting Edition (VBScript) with ADSI to manipulate the IIS configuration.

# Supported Command-Line Scripts

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS 6.0 includes supported scripts that you can find in the *systemroot*\System32 directory. These VBScript scripts use the IIS WMI provider to manage configuration settings in the IIS metabase.

When you use these scripts to create a new site or virtual directory, you can specify the basic properties that are needed to create the site or directory and identify its contents. The scripts apply the same default properties that IIS Manager uses to create new sites or virtual directories, and they adhere to the same rules for inheriting properties.

To configure more advanced properties for a site or virtual directory, use IIS Manager. Alternatively, you can build an XML template that contains the properties you want to apply to a new Web site or virtual directory and then use the Iiscnfg.vbs script to apply this template to any Web site or virtual directory in the IIS metabase.

The computer that runs the command must be running Microsoft® Windows® XP Professional or Windows Server 2003. The computer that the command affects must be running Windows Server 2003 with IIS 6.0. You cannot use these scripts to manage IIS 6.0 on clients running Windows XP Professional because Windows XP runs with IIS 5.1.

Table A.5 lists the supported command-line scripts for IIS 6.0.

**Table A.5 Supported Command-Line Scripts for IIS 6.0**

| Script | For These Areas | Tasks | IIS Operations Guide Topics |
|---|---|---|---|
| Iisweb.vbs | Web sites | Create, delete, start, stop, pause, and query or list Web sites. | Managing Web Sites |
| Iisftp.vbs | FTP sites | Create, delete, start, stop, pause, and query or list FTP sites.<br>Query and set Active Directory properties for a users FTP home directory (use in FTP user isolation). | Managing FTP Sites<br>Setting Active Directory User Isolation |
| Iisvdir.vbs | Web virtual directories | Create, delete, or list the Web virtual directories of a given root. | Managing Web Virtual Directories |
| Iisftpdr.vbs | FTP virtual directories | Create, delete, or list the FTP virtual directories of a given root. | Managing FTP Virtual Directories |
| iisback.vbs | Back up and restore IIS configuration | Create, delete, restore, and list backup copies of IIS configuration. | Managing IIS Configurations |
| Iiscnfg.vbs | Export or import IIS configuration | Export or import an IIS configuration to or from an XML file; copy the metabase and schema to another computer; and save configuration changes to disk. | Managing IIS Configurations |
| Iisext.vbs | Applications and dynamic content services | Configure and manage applications, Web service extensions (like ASP and ASP.NET), and individual files. | Managing Applications and Web Service Extensions |
| Iisapp.vbs | Application pools and their worker processes | List the worker processes (W3wp.exe) currently running and the application pool each one serves. | Listing Running Web Applications |

# Adsutil.vbs IIS Administration Utility

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Adsutil.vbs is an IIS administration utility that uses VBScript with ADSI to manipulate the IIS configuration and modify the metabase. You must run this script by using CScript, which is installed with Windows Script Host.

Adsutil.vbs is a flexible and generic command-line tool. Because Adsutil.vbs is not supported (its documentation and format can change at any time), use Adsutil.vbs primarily as a learning tool.

For more information about how to use Adsutil.vbs, including Adsutil.vbs syntax, parameters, commands and examples, see IIS 6.0 Administration Scripts, Tips, and Tricks in this reference.

# Administering Servers Remotely

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you run IIS on an intranet or the Internet, you can administer your server remotely by using the following tools:

- IIS Manager. Use on your server to remotely connect to and administer an intranet server running IIS 5.*x* or IIS 6.0.
- Terminal Services. Does not require you to install IIS Manager on the remote client computer because, after you are connected to the server that is running IIS, you can use IIS Manager on the Web server as if you are logged on locally.
- Remote Administration (HTML) tool. Use to administer your IIS Web server from any Web browser on your intranet. This version of the tool runs only on servers running IIS 6.0.
- Supported command-line scripts. Use the IIS-supported command-line scripts with the IIS WMI provider to remotely manage a server running IIS. Each command-line script supports the */s* parameter, which you can use to specify the remote server against which you want to perform the command.

For information about enabling the Remote Administration (HTML) tool, and administering servers remotely through IIS Manager, Terminal Services, and the Remote Administration (HTML) tool, see Administering Servers Remotely in IIS 6.0.

For more information about remotely managing servers by using this tool, from the **Help** menu in the Remote Administration (HTML) tool, click **Help Topics**.

# Using FrontPage 2002 Server Extensions with IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage® 2002 Server Extensions from Microsoft® is a technology that works with Internet Information Services (IIS) to support administering, authoring, and browsing Web sites. FrontPage Server Extensions provides you with tools for managing site security, organizing your content into subwebs, and checking site usage. Users who are running the FrontPage client can create and manage professional-looking Web pages for corporate Internet or intranet sites.

**In This Appendix**

Overview of FrontPage 2002 Server Extensions

Preparing to Extend Web Sites

Extending and Configuring Web Sites

Configuring Advanced Security and Customization

**Related Information**

For information about configuring sites on IIS 6.0, see Configuring Internet Sites and Services.

# Overview of FrontPage 2002 Server Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage 2002 Server Extensions is a set of server-side applications that extend the functionality of a Web server and support authoring in Microsoft FrontPage. FrontPage 2002 Server Extensions enables you and your customers to use advanced features without writing complex *server-side* programs, which are applications and programs that run on the Web server rather than on the client computer. Using this technology allows you to deploy content over HTTP without requiring a virtual private network (VPN) or an unsecured method, such as File Transfer Protocol (FTP).

FrontPage 2002 Server Extensions includes new features, such as roles and user rights, and server health monitoring.

A FrontPage client (a computer that is running FrontPage) communicates with a Web server by using WinSock and TCP/IP and with the server extensions by using HTTP. FrontPage implements a remote procedure call mechanism on top of the HTTP POST request, so that the FrontPage client can request documents, add new authors, and so on.

A Web server with FrontPage 2002 Server Extensions recognizes POST requests addressed to the server extensions Internet Server API (ISAPI) or Common Gateway Interface (CGI) applications and directs those requests accordingly. FrontPage can communicate between clients and servers even through proxy servers (firewalls).

# Why Use FrontPage 2002 Server Extensions?

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you use FrontPage 2002 Server Extensions on your Web server, you can divide Web content on a single server into multiple content areas, called *subwebs*, each with its own set of authors and users. In addition, if you build and maintain Web sites for customers and you coordinate with multiple authors, the features included with FrontPage 2002 Server Extensions can help you do the following:

- Host Web sites that are created by using FrontPage-based Web authoring.
- Manage Web sites either on the local server or remotely by using HTML administration pages or a command-line interface.
- Grant authoring, browsing, site management, or other user rights to authenticated users.
- Create, delete, and rename FrontPage-extended subwebs.
- Analyze site usage to find out who is viewing the site and how often.

When you install FrontPage 2002 Server Extensions on your Web server, your users can do the following:

- Create Web pages and subwebs directly on your Web server.
- Enhance their sites with FrontPage components, such as hit counters, link bars, form handlers, full-text searches, and top 10 lists.
- Publish content directly to your Web servers by means of the Internet.
- Collaborate with other users on the same Web site and check files in and out of a FrontPage-extended subweb.
- Save directly to a Web server when using Microsoft Office programs.
- Use search functionality that provides a full-text index.

For more information about how you can use FrontPage 2002 Server Extensions in your organization, see the FrontPage 2002 Server Extensions Support Center.

# FrontPage Server Extensions Terminology

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage 2002 Server Extensions uses the term *virtual server* to describe what IIS documentation typically calls an extended Web site. In this appendix, the term *extended Web site,* rather than the term *virtual server,* is used to describe a Web site that is extended with FrontPage Server Extensions. However, where using IIS terminology might be confusing, such as when documenting procedures that use the FrontPage 2002 Server Extensions user interface, this appendix refers to an extended Web site as a virtual server. Just remember that an IIS extended Web site is the same as a FrontPage 2002 Server Extensions virtual server.

A FrontPage-extended Web site is also known as the *rootweb* because it is the top-level content directory of one or more subwebs. Hence, if you are using FrontPage terminology, every virtual server contains a rootweb, which can have many levels of subdirectories and subwebs that contain its content.

When you extend a Web site with FrontPage 2002 Server Extensions, each extended Web site can have multiple subwebs. A *subweb* is a subdirectory of the rootweb or of another subweb. Subwebs are the FrontPage mechanism for breaking up a Web site so that different areas can be owned and maintained by different people or groups. Each subweb can have many levels of subdirectories.

Several extended Web sites, or rootwebs, can reside on one server. Each extended Web site has its own name and can be accessed through its domain name, IP address, or host header. For more information about using host headers, see Common IIS 6.0 Administrative Tasks.

# Preparing to Extend Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage 2002 Server Extensions provides Web-based and command-line administration for extending Web sites, which enables site owners to author Web sites in FrontPage and delegate site ownership and administration credentials. *Extending* a Web site means enabling FrontPage 2002 Server Extensions features to augment how you manage the content development and security of your sites. Extending a Web site can add an additional level of security management to a Web site and enable a FrontPage client to open the site and author it in FrontPage.

Before you can extend a Web site by using server extensions, you need to install and enable FrontPage 2002 Server Extensions.

# Installing and Enabling FrontPage Server Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By default, IIS 6.0 is installed in a *locked* mode that serves only static HTTP content, thus protecting your systems from malicious users and attackers. When your server receives a default installation, the request handlers that process dynamic content are disabled, which means that features like FrontPage Server Extensions and ASP.NET do not work.

If you use the Configure Your Server Wizard to install IIS 6.0, you can elect to install and enable FrontPage Server Extensions at the time you configure your server role. If you do not install FrontPage Server Extensions at that time, you can add them to your application server at a later time.

**IIS 6.0 Application Isolation Modes**
An IIS 6.0 Web server always runs in one of two mutually exclusive *application isolation modes*. When you install FrontPage 2002 Server Extensions on an IIS 6.0 Web server, the server extensions run in the application isolation mode that you have chosen to run on the server. The application isolation mode you choose depends on what mode best fits your needs, and you can change the isolation mode after installation. At the time you install IIS 6.0, a default application isolation mode is assigned as follows:

- For clean installations, *worker process isolation* mode is the default mode.
- For upgrade installations, *IIS 5.0 isolation mode* is the default mode.

For more information about choosing an IIS 6.0 isolation mode, see Running IIS 6.0 as an Application Server.

**Installing FrontPage 2002 Server Extensions**
Before you install FrontPage 2002 Server Extensions, consider the following guidelines:

- You must run FrontPage 2002 Server Extensions on a computer that uses only the NTFS file system.

  Before you install FrontPage 2002 Server Extensions on an application server, verify that all disk volumes on the computer are formatted with the NTFS file system. Windows provides a conversion utility (Convert.exe) that you can use to convert an existing file allocation table (FAT) or FAT32 volume to NTFS without losing data. If you try to install FrontPage 2002 Server Extensions to a disk drive that is formatted with FAT, your default Web site is not extended.

  For more information about converting FAT or FAT32 volumes to NTFS, see Reformatting or converting a partition to use NTFS in Help and Support Center for Windows Server 2003.

- If you upgrade from an earlier version of FrontPage Server Extensions, you must manually upgrade each FrontPage-extended Web site.

  If you are running multiple extended Web sites and you upgrade from FrontPage 2000 Server Extensions running on the Microsoft® Windows® 2000 Server operating system to FrontPage 2002 Server Extensions running on Microsoft® Windows® Server 2003, Standard Edition; Microsoft® Windows® Server 2003, Enterprise Edition; Microsoft Windows Server 2003, Web Edition; or Microsoft® Windows® Server 2003, Datacenter Edition operating system, you must manually upgrade each FrontPage-extended Web site. If you have a single FrontPage-extended Web site and you upgrade to FrontPage 2002 Server Extensions running on Windows Server 2003, the extended Web site automatically upgrades to FrontPage 2002 Server Extensions.

If you have a Web server that runs on Windows Server 2003 and that has multiple extended Web sites, you must upgrade all your extended Web sites to FrontPage 2002 Server Extensions.

| **Note** |
| --- |
| In FrontPage2002 Server Extensions, the lightweight server extensions option is no longer supported. When you upgrade an extended Web site to FrontPage2002 Server Extensions, it receives a full (not a lightweight) implementation. |

For more information about upgrading your FrontPage-extended Web sites, see Upgrading from FrontPage Server Extensions and Office Server Extensions.

**To install and enable FrontPage 2002 Server Extensions**

1. From the **Start** menu, click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. In the Windows Components Wizard, double-click **Application Server**, double-click **Internet Information Services**, and then select the **FrontPage 2002 Server Extensions** check box.
5. Click **OK** twice, click **Next**, and then click **Finish**.

For more information about installing and setting up FrontPage 2002 Server Extensions, including sample scripts to automate installation, see IIS 6.0 Administration Scripts, Tips, and Tricks.

**Enabling Web Service Extensions**

If you have FrontPage 2002 Server Extensions installed on your Web server but they do not work, check that you have them enabled. For information about enabling Web service extensions, see Enabling Web Service Extensions. For more information about enabling dynamic content, such as FrontPage Server Extensions, see Enabling and Disabling Dynamic Content in IIS 6.0.

# Creating Web Sites to Extend as Virtual Servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Before you can use FrontPage 2002 Server Extensions with your Web sites, you must create one or more IIS Web sites to contain them. By default, IIS provides a working Web site called *Default Web Site*, which points to the content directory *systemroot*\Inetpub\Wwwroot. To create a new Web site, you must create a content directory.

Be sure that you do not nest Web sites that you want to extend as FrontPage virtual servers because FrontPage Server Extensions does not recognize Web sites nested within each other. For example, if you need to create two extended Web sites, you can locate the first content directory in *systemroot*\Inetpub*Test* and the second content directory in *systemroot*\Inetpub*Test2*. If you place the *Test2* directory within the *Test* directory, the permissions in the *Test* directory overwrite all the permissions in the *Test2* directory.

When you create new Web sites, assign them to the default IIS application pool (DefaultAppPool) or to another IIS application pool of your choice.

| **Important** |
| --- |
| Do not use, or make changes to, the MSSharePointAppPool, which is created by default in IIS Manager when you enable FrontPage2002 Server Extensions. MSSharePointAppPool holds special files that are used internally by FrontPage Server Extensions. |

FrontPage 2002 Server Extensions does not support Unicode content directory names or Web Distributed Authoring and Versioning (WebDAV).

# Securing FrontPage 2002 Server Extensions

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage 2002 Server Extensions relies on the security features of the Windows Server 2003 family to provide security for Web site content. FrontPage Server Extensions uses the following two elements of Windows security:

- **User authentication.** Validates the user account that is attempting to gain access to a Web site or network resource.
- **File system security**. Controls which users can access specific files and folders in the file system.

In addition to these elements of Windows security, FrontPage 2002 Server Extensions includes a new security feature: *user roles*. With user roles, you do not have to control the file and folder permissions separately or be concerned about keeping your local groups synchronized with your list of Web users. You use roles to give users permissions on your Web site, and use FrontPage 2002 Server Extensions administration tools to add new users directly. For more information about managing users and user roles, see Using Roles to Manage User Rights later in this appendix.

**User Authentication**

When you use FrontPage 2002 Server Extensions with IIS 6.0 on a computer running Windows Server 2003, user authentication is based on one or more of the following Windows authentication methods: Anonymous authentication, Basic authentication, Integrated Windows authentication, Digest authentication, Advanced Digest authentication, or Certificate authentication.

When you set up your Web server, choose the authentication method you want to use for FrontPage 2002 Server Extensions. With the exception of Anonymous authentication, you cannot change the authentication method by using the FrontPage 2002 Server Extensions administration tools; you must instead use IIS Manager to change the authentication method.

FrontPage 2002 Server Extensions supports connectivity through firewalls. Make sure your firewall is open for the standard HTTP ports 80 and 443 (which port you use depends on your configuration). If you use a firewall and want to use NTLM Integrated Windows authentication (formerly called NTLM and also referred to as Windows NT Challenge/Response authentication), make sure your firewall supports NTLM. Otherwise, configure your Web sites with Basic authentication, preferably with SSL.

For more information about these authentication methods, including brief descriptions of each, see Common IIS 6.0 Administrative Tasks. For additional information about security options for IIS 6.0, see Managing a Secure IIS 6.0 Solution.

**File System Security**

FrontPage 2002 Server Extensions relies in part on the Windows operating system to help secure the file system for your Web sites. The Windows Server 2003 family uses access control lists (ACLs) to help secure files and folders. Because ACLs require the NTFS file system, your Web servers must use NTFS when you run FrontPage 2002 Server Extensions on IIS 6.0. For more information about access control, see Understanding access control in Help and Support Center for Windows Server 2003. For more information about setting ACLs for FrontPage 2002 Server Extensions, see Windows Security Model.

| **Caution** |
| --- |
| When you build FrontPage-extended Web sites and subwebs, all the ACLs for those sites are maintained by FrontPage Server Extensions. Do not try to change the ACLs in Windows Server2003. Instead, in order to maintain security, always make permissions changes and other security-related changes by using the FrontPage2002 Server Administration tools. |

# Extending and Configuring Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After you install FrontPage 2002 Server Extensions, you can extend a Web site and then customize the Web site by configuring several options. To configure FrontPage Server Extensions, use one of the HTML administration pages, such as the Server Administration page, which you can access through Microsoft SharePoint Administrator. You can also use the FrontPage Server Extensions command-line tools. For more information about using the command-line tools, see Using Command-Line Tools later in this appendix.

The HTML administration pages, which are automatically created when you install FrontPage 2002 Server Extensions, replace the Microsoft Management Console (MMC) interface for administering the FrontPage 2000 Server Extensions. You must be a member of the Administrators group on the local computer to view and use these HTML administration pages. For more information about accessing the HTML administration pages, see Extending Web sites later in this appendix.

Microsoft SharePoint Administrator provides the following HTML administration pages to help you administer your FrontPage-extended Web sites:

- **Server Administration pages.** Allow you to extend a Web site (make it a FrontPage virtual server) and administer settings for your Web server. You can also view a list of virtual servers on your Web server and obtain information about each one.

  For example, you can specify default settings for extended Web sites. You can also reset user passwords and specify which rights are available to assign to roles and users.

- **Virtual Server Administration pages.** Allow you to globally control settings for any extended Web site (FrontPage virtual server) on your computer.

  A newly created extended Web site inherits settings from defaults set on the Server Administration pages. You can use the Virtual Server Administration pages to change these default settings and to specify what settings to use for subwebs of the extended Web site. For example, you can individually upgrade an extended Web site to FrontPage 2002 Server Extensions. You can also change configuration settings and set user account limits.

- **Site Administration pages.** Control settings for each individual Web site that is part of a FrontPage rootweb or subweb.

  For example, you can manage users and roles; perform usage analysis tasks; create, merge, or delete subwebs; and perform server health checks on virtual servers.

# Extending Web sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

After creating a Web site in IIS, you can use the Server Administration pages provided by FrontPage 2002 Server Extensions to extend the Web site and administer settings for your Web server.

For information about extending a Web site by using the Server Administration pages, see Using FrontPage Server Extensions to Extend Web Sites.

# Configuring Extended Web Sites

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Any changes that you make to an extended Web site at the rootweb level affect all sites and subwebs for that rootweb. You can set the following options at the rootweb level by using the Virtual Server Administration pages.

- **Authoring options.** Turn authoring on or off for an extended Web site. When authoring is turned on, users can add new pages, edit pages, and remove pages for a Web site. When authoring is turned off, FrontPage client users cannot perform any of these tasks.

- **Mail settings.** Specify the Simple Mail Transfer Protocol (SMTP) mail server, the mail encoding option, and the **From** and **Reply-to** address to use for e-mail messages related to FrontPage Server Extensions.

- **Performance tuning options.** Use default settings to manage the performance of your FrontPage-extended Web sites, or specify settings for several performance tuning options, including caching limits for documents, images, and files.

- **Client scripting options.** Control the types of scripts that authors can run on your site, including ISAPI, CGI, ASP, and others.

- **Security settings.** Change security options, track authoring data by creating logs, and specify whether users can store executable files on your server.

If you do not change these configuration settings, the extended Web site inherits the settings from the Set Installation Defaults administration page (in the Server Administration pages). Use the Change Configuration Settings page in the Virtual Server Administration pages to configure your extended Web site.

**To configure FrontPage-extended Web sites by using the Virtual Server Administration pages**

1. From the **Start** menu, point to **Administrative Tools**, and then click **Microsoft SharePoint Administrator**.
2. On the **Server Administration** page, next to the FrontPage-extended Web site name, click **Administration**.
3. Click **Change configuration settings**.
4. On the **Change Configuration Settings** page, click the options you want, and then click **Submit**.

# Configuring Web Sites under the Rootweb

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to configuring the settings for an extended Web site (FrontPage virtual server), you can configure the settings for individual Web sites that are part of a rootweb (these Web sites become extended because they are contained in, or in a subdirectory of, the rootweb). Use the Site Administration page or the command-line tools to configure the following features:

- **User and role settings**. Add or delete users and roles, and change which roles a user is assigned to. For more information about how FrontPage 2002 Server Extensions manages user rights, see Using Roles to Manage User Rights later in this appendix.

- **Usage analysis settings**. Configure usage analysis to track how many users visit your site; track the type and number of hits that your site receives, and track other site-usage information. The usage reports rely on a log file that is generated by your IIS application server.

- **Server health check settings**. Detect and repair certain functionalities relating to extended Web sites, such as checking role configurations, managing security, or checking anonymous access.

**To configure Web sites by using the Site Administration page**

1. If you are a server administrator

    a. From the **Start** menu, point to **Administrative Tools**, and then click **Microsoft SharePoint Administrator**.

    b. On the **Server Administration** page, click the name of the extended Web site (FrontPage virtual server) that you want to manage, which takes you to the Site Administration page for the rootweb of that extended Web site.

2. If you are a site administrator

    a. Open your Web site in the FrontPage client.

    b. On the **Tools** menu, point to **Server**, and then click **Administration Home**.

# Using Command-Line Tools

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage 2002 Server Extensions provides two new tools for command-line administration:

- **Owsadm.exe**. Administers FrontPage 2002 Server Extensions from the local server. You must have administrative credentials on the local server to use this tool.

- **Owsrmadm.exe**. Administers FrontPage 2002 Server Extensions from a remote computer. The Web server must be running before you can use this tool. You must have administrative credentials on the server to use this tool. In addition, you must have a copy of the tool on your local computer, which must be running Windows.

When you run Owsadm.exe, you supply an operation and a set of command-line parameters in the following form:

**-operation CommandName -parameter value**

Each parameter for the command line also has a short form that you can use instead of the full parameter name. For example, you can use the following command to add a specified user to the Web site on port 80:

**Owsadm.exe -o users -c add -u username -p 80**

To configure settings for an extended Web site (FrontPage virtual server) from the command line, use the **SetProperty** parameter with Owsadm.exe or Owsrmadm.exe. You can set authoring, performance tuning, client scripting, mail settings, and security settings properties by using the command line. For example, to turn on authoring for an extended Web site at port 1088, configure the **Authoring** property by using the following command:

**Owsadm.exe -o setproperty -p 1088 -pn Authoring -pv enabled**

For more information about using these command-line tools to extend and administer sites, including links to command-line parameters, operations, and properties, see "Command-line Administration" in the SharePoint Team Services Administrator's Guide. For more information about using FrontPage 2002 Server Extensions with IIS 6.0 and for additional links to FrontPage-related Web sites, see Using FrontPage Server Extensions to Extend Web Sites.

# Configuring Advanced Security and Customization

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to choosing a Windows authentication method for your extended Web sites, you need to manage user access by configuring user rights. FrontPage 2002 Server Extensions are designed to be scalable so that you can configure Web site permissions to be as simple or as complex as you need.

FrontPage 2002 Server Extensions also provides advanced customization, which Internet service providers (ISPs) and Web Presence Providers (WPPs) might find especially useful. For example, you can authenticate users separately for each extended Web site (FrontPage virtual server), which allows you to better manage anonymous user access to Web sites on separate virtual servers. In addition, you can create extended Web sites that have access to shared folders on your network by using the Universal Naming Convention (UNC) standards. Customizing FrontPage 2002 Server Extensions by using UNC network shares can be helpful if your Web server is hosting many large Web sites and subwebs.

All these advanced customizations are intended for experienced IT administrators who are familiar with Windows authentication methods and with configuring FrontPage 2002 Server Extensions.

| **Note** |
| --- |
| Registered WPPs for Microsoft® FrontPage® version2002 are companies that provide hosting services for FrontPage–based sites. Many WPPs also provide technical support, training and consulting, 24-hour monitoring, maintenance and traffic reporting, security management, or other Web-based services. |

# Using Roles to Manage User Rights

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

FrontPage 2002 Server Extensions uses roles to manage user rights. Each user is assigned at least one role with corresponding rights. You can edit or customize the rights assigned to a role, create a new role, or delete an unused role.

You can add users to a Web site without assigning them to a role. For example, if you are creating new user accounts for a Web site, you can create the user accounts and then assign or remove roles at any time.

Not only does FrontPage 2002 Server Extensions offer greater flexibility in assigning roles, it also provides user-friendly management tools. In earlier versions, you managed user roles from within the Microsoft FrontPage client by using the **Security** command on the **Tools** menu. With FrontPage 2002 Server Extensions, you can manage roles by using the HTML administration pages or the command-line tool.

You must be both a Web site administrator and an administrator on the local computer to perform administrative tasks that affect settings for all extended Web sites and subwebs in a rootweb. FrontPage 2002 Server Extensions includes the following roles by default:

- **Browser**. Has rights to view pages, view Web document discussions, and read lists.
- **Author**. Has Browser rights, plus rights to edit pages, directories, and lists.
- **Advanced Author**. Has Author rights, plus rights to define and apply themes and borders, link style sheets, and recalculate a Web site (check the links). This role appears as "advauthor" on the command line.
- **Administrator**. Has all Browser, Author, and Advanced Author rights, plus rights to configure roles, create local computer user accounts, manage source control, create subwebs, manage server health, and manage usage analysis. This role appears as "admin" on the command line and is a Web site administrator role, not a server administrator role.

For more information about user roles, including a complete list of user rights, their default settings, and how to configure them, see "Managing Roles," "Managing Users," and "User Rights" in the SharePoint Team Services Administrator's Guide.

# Authenticating Users Separately for Each Extended Web Site

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you use FrontPage 2002 Server Extensions on Windows Server 2003, you can take advantage of a new feature that allows you to better manage anonymous user access for Web sites on separate extended Web sites (FrontPage virtual servers). You can use group accounts to separate anonymous access for each extended Web site, so that only anonymous users and registered users on the same extended Web site have access to the source files. When you use this new feature, the users on one extended Web site can no longer read source files on another extended Web site.

This is an advanced feature and is intended for use by Internet service providers (ISPs) and Web Presence Providers (WPPs) who allow users to have direct access to the file system of their Web sites through methods such as FTP, ASP, and Perl.

To take advantage of this new feature, you must create a group account to track all user accounts that have access to a virtual server. To turn on this feature and to specify the name for the group account, you must edit the Windows registry. For more information about using this new feature, including step-by-step instructions for enabling group accounts, see "Authenticating Users Separately for Each Virtual Server" in the Windows® SharePoint® Services Administrator's Guide.

# Connecting Web Sites to UNC Network Shares

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you run FrontPage 2002 Server Extensions on Windows Server 2003, you can create FrontPage-extended Web sites that have access to shared folders on your network by using the Universal Naming Convention (UNC) standard. A configuration that uses UNC shares is often advantageous when you set up a server that hosts many large Web sites and subwebs. You can create Web sites on your Web server and have the files for the Web sites distributed across multiple file servers on your network. Connecting Web sites to UNC network shares requires that you configure the file server and the Web server.

You must be a local administrator of both servers to perform the steps needed for this advanced configuration. Both servers must be running Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Advanced Edition; or Windows® Server 2003 Enterprise Edition.

For more information about connecting a FrontPage extended Web site to a UNC network share, see "Connecting Web Sites to UNC Network Shares," in the Windows® Sharepoint® Services Administrator's Guide.

For more information about configuring Web sites that have content stored on a file server, see Deploying and Configuring Internet Information Services (IIS) 6.0 with Remotely Stored Content on UNC Servers and NAS Devices.

# IPv6 and IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0 provides Internet services to clients connecting over the next generation of Internet Protocol (IP) — known as IP version 6 (IPv6). When you use IPv6 with IIS on a server running the Microsoft® Windows® Server 2003 operating system, your sites can respond to both IPv6 requests and Internet Protocol version 4 (IPv4) requests. Although current demand for IPv6 is limited, the adoption of new Internet-connected devices, such as phones and handheld computers, is expected to quickly exhaust the remaining IPv4 address space and speed the transition to IPv6.

**In This Appendix**


Summary of Protocol Changes from IPv4 to IPv6

Comparing IPv4 and IPv6 Addresses

How IIS 6.0 Supports IPv6

Securing IPv6 Networks

Installing or Removing IPv6

Additional Resources for IPv6 and IIS 6.0


**Related Information**

- For more information about IIS 6.0 architecture, including HTTP.sys, see IIS 6.0 Architecture.

# Summary of Protocol Changes from IPv4 to IPv6

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In response to concern about the finite number of IPv4 addresses, the Internet Engineering Task Force (IETF) developed a suite of protocols and standards known as Internet Protocol version 6 (IPv6). IPv6 was designed to replace IPv4, which is more than 20 years old. Because IPv6 is expected to become increasingly important as the Internet continues to grow, IPv6 functionality is included in the Microsoft® Windows® XP and the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Datacenter Edition; and Windows® Server 2003, Web Edition operating systems.

The Internet Protocol (IP) is the network-layer protocol used by TCP/IP for addressing and routing packets of data between hosts. The current version of the IP, IPv4, has remained primarily unchanged since RFC 791 was published in 1981. IPv4 is robust, easily implemented, and interoperable, and it scales well. However, continued dramatic growth of the Internet is pushing the limits of IPv4 design. Concerns with IPv4 include the scarcity of public IPv4 addresses for use on the Internet, the size and complexity of its backbone routing tables, and the need for simpler, more automatic configuration.

IPv6, previously named IP Next Generation (IPng), was developed primarily to allow for larger (128-bit) IP addresses. Additional enhancements include the following:

- A more efficient routing infrastructure that allows backbone routers to maintain much smaller routing tables.
- A new header format that reduces header overhead and provides more efficient processing at intermediate routers.
- A simplified host configuration that uses both stateful and stateless address configuration.
- Built-in security provided by Internet Protocol security (IPsec). (In IPv4, IPsec is optional.)
- Better support for Quality of Service (QoS), which is the set of methods or processes that a service-based organization uses to maintain a specific level of quality.
- A way to efficiently manage the interaction of neighboring nodes by using multicast and unicast messaging.
- Extension headers for adding new features to IPv6.

Table G.1 lists pertinent IPv6 RFCs and serves as a useful reference to the source documents.

**Table G.1 Summary of Useful Source Documents for Migrating from IPv4 to IPv6**

| Version | Date | RFC | RFC Title |
|---|---|---|---|
| IPv4 | September, 1981 | RFC 791 | Internet Protocol<br>(Protocol specification) |
| IPv6[1] | January, 1995 | RFC 1752 | The Recommendation for the IP Next Generation Protocol<br>(Standards track for IPv6 General RFCs) |
| IPv6 | December, 1995 | RFC 1883 | Internet Protocol, Version 6 (IPv6) Specification<br>(Proposed standard — made obsolete by RFC 2460) |
| IPv6 | December, 1995 | RFC 1886 | DNS Extensions to support IP version 6<br>(IPv6 applications RFCs) |

| | | | |
|---|---|---|---|
| IPv6 | July, 1998 | RFC 2373<br>RFC 3513[2] | IP Version 6 Addressing Architecture<br>Internet Protocol Version 6 (IPv6) Addressing Architecture<br>(Standards track for IPv6 addressing RFCs) |
| IPv6 | December, 1998 | RFC 2460[3] | Internet Protocol, Version 6 (IPv6) Specification<br>(Standards track for the Network Layer RFCs and Internet drafts) |

[1] IPv5 was an experimental non-IP real-time stream protocol called ST, which was never widely used.

[2] RFC 3513 supercedes RFC 2373; however Windows Server 2003 does not implement its changes with regard to site-local addresses.

[3] In RFC 2460, Internet Protocol, Version 6 (IPv6) Specification, IPv6 is described as a connectionless, unreliable datagram protocol that is used primarily for addressing packets and routing them between hosts. *Connectionless* means that a session is not established before data exchange begins. *Unreliable* means that delivery is not guaranteed. IPv6, like IPv4, always makes a best-effort attempt to deliver a packet, but an IPv6 packet might be lost, delivered out of sequence, duplicated, or delayed. IPv6 itself does not attempt to recover from these types of errors; the acknowledgment of packet delivery and the recovery of lost packets is done by a higher-layer protocol, such as TCP. TCP performs reliably over both IPv4 and IPv6.

For more information about the design changes needed by IPv4 to accommodate the increasing demands of network traffic, see Introduction to IPv6 in Help and Support Center for Windows Server 2003. For more information about IPv6, see IPv6 features in Help and Support Center for Windows Server 2003.

For more information about RFCs, see RFC-Editor.org.

# Comparing IPv4 and IPv6 Addresses

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The size of an IPv6 address is 128 bits, which is four times larger than an IPv4 address. In theory, the 32-bit address space that IPv4 uses provides 4,294,967,296 possible addresses; however, previous and current allocation practices limit the number of public IPv4 addresses to a few hundred million. By contrast, the 128-bit address space that IPv6 uses provides $3.4 \times 10^{38}$ possible addresses.

The size of the IPv6 address allows for subdividing the address into a hierarchical routing structure that reflects the current topology of the Internet. This structure provides great flexibility for hierarchical addressing and routing, which the IPv4-based Internet lacks.

# Comparing Address Formats

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IPv4 addresses are represented in a *dotted-decimal format*, in which the 32-bit address is divided into four 8-bit sections. Each set of 8 bits is converted into its decimal equivalent and is separated from adjacent 8-bit decimal equivalents by periods. The following is an example of an IPv4 address:

```
131.107.16.200
```

In IPv6, the 128-bit address is divided into eight 16-bit blocks, each of which is converted to a 4-digit hexadecimal number that is separated from adjacent blocks by colons. The resulting representation is called *colon-hexadecimal format*.

The following is an IPv6 address in binary form:

```
0010000111011010000000001101001100000000000000000001011110011101 1
0000001010101010000000001111111111111110001010001001110001011010
```

First, the 128-bit address is divided into eight 16-bit blocks, as follows:

```
0010000111011010   0000000011010011   0000000000000000   0010111100111011
0000001010101010   0000000011111111   1111111000101000   1001110001011010
```

Then, each of the eight 16-bit blocks is converted to hexadecimal and delimited with colons. The result is the following:

```
21DA:00D3:0000:2F3B:02AA:00FF:FE28:9C5A
```

An IPv6 address can be further simplified by removing the leading zeros within each 16-bit block. However, each block must have at least a single digit. With leading zero suppression, the address used in this example becomes the following:

```
21DA:D3:0:2F3B:2AA:FF:FE28:9C5A
```

**Compressing Zeros in IPv6 Addresses to the Double-Colon Format**

IPv6 addressing conventions also allow you to simplify an address that contains long sequences of zeros. If an address contains consecutive groups of 16-bit blocks that are set to 0 in the colon-hexadecimal format, you can compress the consecutive blocks to :: (known as double-colon) to simplify the address. To avoid ambiguity, use zero compression only once within any one address. Otherwise, you cannot determine the number of 0 bits represented by each instance of a double-colon (::).

Table G.2 provides two examples of IP addresses and shows how zero compression changes each address.

**Table G.2 Effect of Zero Compression on Sample IP Addresses**

| Address Before Zero Compression | Address After Zero Compression |
|---|---|
| FE80:0:0:0:2AA:FF:FE9A:4CA2 | FE80::2AA:FF:FE9A:4CA2 |
| FF02:0:0:0:0:0:0:2 | FF02::2 |

To determine how many 0 bits are represented by the double-colon in a compressed address, count the number of blocks in the address, subtract this number from 8, and then multiply the result by 16. For example, the address FF02::2 contains two blocks (the FF02 block and the 2 block). Therefore, the number of bits expressed by the double-colon is 96 (96 = (8 - 2) 16).

**Understanding Prefixes**

IPv4 implementations commonly use a dotted-decimal representation of the network prefix length, which is called the subnet mask. IPv6 does not use subnet masks; it supports only prefix length notation.

The *prefix* is the part of an IP address where the bits have fixed values or are the bits of a route or subnet identifier. Prefixes for IPv6 routes and subnet identifiers are expressed in the same way as classless inter-domain routing (CIDR) notation is expressed for IPv4, that is, *address/prefix length*. The *prefix length* specifies the number of left-most bits that make up the subnet prefix. For example, an IPv6 prefix can be represented as follows:

```
3FFE:2900:D005:F28B::/64
```

In this example, the first 64 bits of the global unicast address are the prefix, and the remaining 64 bits (128 - 64 = 64) are the interface ID.

For more information about the architecture of IPv6 addresses, see RFC 2373, IP Version 6 Addressing Architecture.

# Comparing Address Types

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you install IPv6 on a computer that is not connected to a network in which an IPv6 router is present, the computer automatically configures a link-local IPv6 address, which is a type of address that allows you to communicate with computers on your subnet. If you connect your computer to a subnet in which an IPv6 router is present, the router assigns your computer an IPv6 global or site-local address. The *site-local address* allows your computer to communicate within your intranet. The *global address* allows your computer to communicate with computers on the IPv6 Internet.

The left-most bits of an IP address are called the *format prefix* (FP), which indicates the specific type of IPv6 address. IPv6 accommodates many address types, including the following:

- **Unicast addresses**. Provide point-to-point, directed communication between two hosts on a network.
- **Multicast addresses**. Provide a method for sending a single IP packet to multiple hosts in a group. A multicast address is used for one-to-many communication.
- **Anycast addresses**. Provide a method of delivering a packet to the nearest member of a group. Currently, anycast addresses are used only as destination addresses and are assigned only to routers. An anycast address is used for one-to-one-of-many communication.

Table G.3 compares some basic elements of IPv4 and IPv6 addressing.

**Table G.3 Comparison of IPv4 and IPv6 Addressing Elements**

| Address Space Element | IPv4 Address | IPv6 Address |
|---|---|---|
| Unspecified address | 0.0.0.0 | 0:0:0:0:0:0:0:0: or :: |
| Loopback address | 127.0.0.1 | 0:0:0:0:0:0:0:1: or ::1 |
| Address type | Public IPv4 addresses | Global addresses (aggregatable global unicast addresses) |
| • Address type | Private IPv4 addresses, such as:<br>• 10.0.0.0/8<br>• 172.16.0.0/12<br>192.168.0.0/16 | Site-local addresses, which always begin with FEC0::/48) |
| Address type | Automatic Private IP Addressing (APIPA), which uses the 169.254.0.0/16 prefix | Link-local addresses, which always begin with FE80::/64Text representation |
| Text representation | Dotted-decimal format | Colon-hexadecimal format with suppression of leading zeros and zero compression. IPv4-compatible addresses are expressed in dotted-decimal notation. |
| Network bits representation | Subnet mask in dotted-decimal format or prefix-length notation | Prefix-length notation only |

| | | |
|---|---|---|
| DNS name resolution | IPv4 host address (A) resource record[1] | IPv6 host address (AAAA) resource record[2] |

[1] An A resource record, which is stored on your DNS servers, enables mapping from a host name to an IPv4 32-bit address.

[2] AAAA (quad-A) resource records enable mapping from a host name to an IPv6 128-bit address.

**Unicast Addresses**
Unicast addresses identify a single interface within the scope of a particular type of unicast address. The scope of an address is the region of the IPv6 network over which the address is unique. With the appropriate unicast routing topology, packets addressed to a unicast address are delivered only to a single interface.

The following are types of unicast IPv6 addresses:

- **Aggregatable global unicast addresses**. Identified by the format prefix (FP) of 001, these addresses are equivalent to public IPv4 addresses.
- **Local-use unicast addresses**. Provide two types of addresses:
  - **Link-local addresses**. Identified by the FP of 1111 1110 10, these addresses are used by nodes when they are communicating with neighboring nodes on the same link.
  - **Site-local addresses**. Identified by the FP of 1111 1110 11, these addresses are equivalent to the IPv4 private address space. Use these addresses between nodes that communicate with other nodes in the same site.
- **Unspecified address**. Used only to indicate the absence of an address; this type of address cannot be assigned to a node. The IPv6 unspecified address, 0:0:0:0:0:0:0:0 or ::, is equivalent to the IPv4 unspecified address of 0.0.0.0.
- **Loopback address**. Used to identify a loopback interface, which enables a node to send packets to itself. The IPv6 loopback address, 0:0:0:0:0:0:0:1 or ::1, is equivalent to the IPv4 loopback address of 127.0.0.1.
- **Transition, or compatibility, addresses**. Provided to help you migrate from IPv4 to IPv6; these addresses allow both types of hosts to coexist on your network.

**Types of Transition IPv6 Addresses**
To help you transition from IPv4 to IPv6 and to facilitate the coexistence of both types of hosts, IPv6 defines the following transition IPv6addresses.

**IPv4-compatible addresses**
*IPv4-compatible addresses* are used by IPv6/IPv4 nodes that communicate with IPv6 over an IPv4 infrastructure. *IPv6/IPv4 nodes* are nodes that run both the IPv4 and IPv6 protocols. The format for an IPv4-compatible address is 0:0:0:0:0:0:w.x.y.z or ::w.x.y.z (where w.x.y.z is the dotted-decimal representation of a public IPv4 address). The IPv6 protocol for Windows XP and Windows Server 2003 provides support for IPv4-compatible addresses, but support is not enabled by default.

**IPv4-mapped addresses**
*IPv4-mapped addresses* are used to represent an IPv4-only node to an IPv6 node. The IPv4-mapped address is never used as a source or destination address for an IPv6 packet. It is used only for internal representation. The format for an IPv4-mapped address is 0:0:0:0:0:FFFF:w.x.y.z or ::FFFF:w.x.y.z. The IPv6 protocol for Windows XP and Windows Server 2003 does not support IPv4-mapped addresses.

**6to4 addresses**
*6to4 addresses* are used for communicating between two nodes that are running both IPv4 and IPv6 over the Internet. The 6to4 address is formed by combining the prefix 2002::/16 with the 32 bits of the public IPv4 address of the node or the site of the node, thus forming a 48-bit prefix. For example, for the IPv4 address of 131.107.0.1, the 6to4 address prefix is 2002:836B:1::/48 (where 836B:1 is the colon hexadecimal notation for 131.107.0.1). Support for 6to4 addresses is provided by the IPv6 Helper service (known as the *6to4 service*) that is included with the IPv6 protocol for Windows XP and Windows Server 2003.

**Global Addresses**
*Global addresses*, which are identified by an FP of 001, and which are also called aggregatable global unicast addresses, are

equivalent to public IPv4 addresses. Global addresses are globally routable and reachable on the IPv6 Internet.

As the name implies, you can aggregate, or summarize, global addresses to produce an efficient routing infrastructure. Unlike the current IPv4-based Internet, which has a mixture of both flat and hierarchical routing, the IPv6-based Internet is designed to support efficient hierarchical addressing and routing. The *scope* of a global address, which is the region of the IPv6 internetwork over which the address is unique, is the entire IPv6 Internet.

Figure G.1 illustrates the structure of an IPv6 global address.




**TLA ID field**
Indicates the Top Level Aggregation Identifier (TLA ID) for the address. The size of this field is 13 bits. The TLA ID identifies the highest level in the routing hierarchy. TLA IDs are administered by the Internet Assigned Numbers Authority (IANA) and allocated to local Internet registries that, in turn, allocate individual TLA IDs to large, global Internet service providers (ISPs). A 13-bit field allows up to 8,192 different TLA IDs. Routers in the highest level of the IPv6 Internet routing hierarchy (called default-free routers) do not have a default route — only routes with 16-bit prefixes corresponding to the allocated TLA IDs and additional entries for routes based on the TLA ID assigned to the routing region where the router is located.

**Res field**
Reserves space for future expansion of either the TLA ID or the NLA ID field. The size of this field is 8 bits.

**NLA ID field**
Indicates the Next Level Aggregation Identifier (NLA ID) for the address. The NLA ID identifies a specific customer site. The size of this field is 24 bits. The NLA ID allows an ISP to create multiple levels of addressing hierarchy to organize addressing and routing and to identify sites. The structure of the ISP's network is not visible to default-free routers.

**SLA ID field**
Indicates the Site Level Aggregation Identifier (SLA ID) for the address. The SLA ID is used by an individual organization to identify subnets within its site. The size of this field is 16 bits. An organization can use this field to create 65,536 subnets or multiple levels of addressing hierarchy. Being assigned an SLA Id is equivalent to being allocated an IPv4 Class A network ID (assuming that the last octet is used for identifying nodes on subnets). The structure of the customer's network is not visible to the ISP.

**Interface ID field**
Identifies the interface of a node on a specific subnet. The size of this field is 64 bits.

For more information about IPv6 addressing, see RFC 2373, IP Version 6 Addressing Architecture, at the Internet Engineering Task Force Web site.

# How IIS 6.0 Supports IPv6

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Except for the differences in functionality outlined in this section, IIS 6.0 provides the same Web services to clients that connect by using IPv6 as it does for clients that connect by using IPv4.

Differences in IIS Functionality Between IPv4 and IPv6

IIS 6.0 provides the same core functionality for users of IPv6 as it does for users of IPv4. However, only a subset of IIS 6.0 functionality is available for users of IPv6. The following are the most significant differences:

- **The IIS Manager user interface does not support IPv6 addresses**. IIS Manager does not display IPv6 addresses as it does IPv4 addresses; that is, the IIS user interface does not provide a way to work with or manipulate IPv6 addresses. However, manipulating literal hexadecimal IPv6 addresses is usually not recommended, so lack of this functionality is unlikely to prevent you from installing and using IPv6.

- **The IP Address Restrictions feature is not supported**. The IP Address Restrictions feature in IIS 6.0 does not support IPv6 addresses or IPv6 prefixes.

- **Bandwidth throttling is not supported for responses sent over IPv6 addresses**. If you change the **MaxBandwidth** and **MaxGlobalBandwidth** metabase properties, you do not affect IPv6 network traffic. However, connection limits and connection time-outs, which are related IIS 6.0 features, are supported.

- **The ServerBindings and SecureBindings metabase properties do not support IPv6 addresses**. Both properties specify strings that IIS 6.0 uses to determine which network endpoints are used by the server instance. The string format for the **ServerBindings** property is *IPAddress:Port:HostName*. Both the *IPAddress* and *HostName* parameters of the string are optional; however, the *IPAddress* component of the string is limited to storing an IPv4 address. Any unspecified parameters default to an all-inclusive wildcard.

  Because of the limitations in IPv6 functionality for these metabase properties, IIS 6.0 functionality for IPv6 is affected as follows:

  - **Site routing is limited to host headers only**. You cannot configure sites to route on an IPv6 address or on a combination of an IPv6 address and host header. This limitation also affects sites that are configured to route based on an IPv4 address. When you install IPv6, sites that are already specifically configured for IPv4 site-based routing do not respond to requests that come in over IPv6.

  - **The number of Secure Sockets Layer (SSL) sites is limited to one**. Due to the IP routing restriction for IPv6, IIS deployments designed for IPv6 addresses are limited to one SSL site per computer.

- **File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), and Network News Transfer Protocol (NNTP) services are not supported**. IPv6 is supported only for the WWW service. The FTP, SMTP, and NNTP services do not have IPv6 support in IIS 6.0.

- **If you install IPv6 on a computer that is running IIS 6.0, all sites on the computer respond to IPv6 clients**. You cannot configure individual sites or virtual directories to respond to IPv6 traffic while others on the same server respond to IPv4 traffic.

- **Logging tools must support IPv6 address formats in order to function correctly**. IIS writes IPv6 addresses to the log file when IPv6 is enabled and client computers connect to the server by using IPv6 addresses. Log parsing tools that are to be used with log files for IPv6 sites must support IPv6 address formats.

- **The EnableReverseDnsLookup property is not supported**. For IPv4, a value of **true** for the **EnableReverseDnsLookup** metabase property allows reverse DNS lookup to determine the DNS name of the client computer. For IPv6, however, this functionality is not enabled.

| Note |
| --- |
| Setting the **EnableReverseDnsLookup** metabase property to **true** does not cause the REMOTE_HOST server variable to return the DNS name of the client, as it does for IPv4. REMOTE_HOST always contains the IPv6 address, regardless of the **EnableReverseDNSLookup** property setting. |

Using the IPv6-Aware ISAPI Server Variables

When you install IPv6 on a server running IIS 6.0, Internet Server API (ISAPI) server variables provide support for IPv6. The ISAPI framework provides the appropriate local-host and remote-host server variables for IPv6 network addresses: LOCAL_ADDR and REMOTE_ADDR. When clients connect over IPv6, these variables store the IPv6 address.

It is important to note that IPv6 addresses can be longer than IPv4 addresses, so you need to take steps to prevent buffer overruns when you install IPv6.

| Important |
| --- |
| Buffer overruns are one of the most common causes of security breaches. Preventing buffer overruns helps protect your server from being attacked. |

To prevent buffer overruns, you must allocate more space to hold the string representation of IPv6 addresses. For example, the longest possible IPv4 string looks something like 123.123.123.123 (16 characters, including the trailing zero required to express the string in some programming languages), whereas the longest IPv6 string looks like 1111:2222:3333:4444:5555:6666:123.123.123.123%1234567890, which is 57-characters long, including the trailing zero. Therefore, when you allocate buffers, use 16 characters for IPv4 addresses and 57 characters for IPv6 addresses.

| Note |
| --- |
| The "%1234567890" portion of the string indicates the zone ID, which is an integer that specifies the scope, or zone, of the destination. The zone ID is needed when you are specifying a link-local destination address or a site-local destination address (if you are using multiple sites). For link-local addresses, the zone ID is typically equal to the interface index of the desired sending interface. For site-local addresses, the zone ID is equal to the site number. |

Both ISAPI server variables (LOCAL_ADDR and REMOTE_ADDR) use the typical IP address format for the applicable IP version (IPv4 or IPv6). For example, for an IPv6 request, both server variables use an IPv6 IP address in the colon-hexadecimal format; for an IPv4 request, both server variables use an IPv4 IP address in the dotted-decimal format. Note that the IPv6 address is simply a string variable, which can contain the % character. For example, fe80::1:1:1:1%4 is an acceptable IPv6 address.

How Installing IPv6 Affects Web Browser Use and Client Connectivity

When you install IPv6 on a computer in your network, two limitations of Internet Explorer 6 (the version provided with Windows Server 2003) can affect the computers ability to use the Internet Explorer Web browser:

- **Internet Explorer 6 does not connect when literal IPv6 addresses are used in the URL**. Internet Explorer 6 does not support the format for literal IPv6 addresses in URLs that is described in RFC 2732, Format for Literal IPv6 Addresses in URLs. However, Internet Explorer does support IPv6 by means of DNS name resolution. Therefore, you should create DNS AAAA resource records that resolve Web server names to an IPv6 address and then use Web server names in the URL.

- **Internet Explorer 6 cannot browse IPv6 Web sites by using a proxy server if the proxy server is not IPv6-enabled**. When Internet Explorer is configured to use a proxy server, all name resolution requests for Web sites are forwarded to the proxy server. If the proxy server is not IPv6-enabled, proxy-based requests for local or remote IPv6 Web pages are unsuccessful. For information about how to configure Internet Explorer to operate without the use of a proxy server, see Internet Explorer Help.

Table G.4 shows the network protocol that clients use to connect to different IP configurations.

**Table G.4 Client Connectivity Matrix for IPv6**

| Client IPv6 Awareness | Type of Client Operating System | IIS 6.0 without IPv6 Installed | IIS 6.0 with IPv6 Installed |
| --- | --- | --- | --- |

| | | | | |
|---|---|---|---|---|
| Client is IPv6 unaware | IPv4-only | Connect over IPv4 | Connect over IPv4 | |
| | | IPv6-only | No connectivity | No connectivity |
| | | IPv4/IPv6 | Connect over IPv4 | Connect over IPv4 |
| Client is IPv6 aware[1] | IPv4-only | Connect over IPv4 | Connect over IPv4 | |
| | | IPv6-only | No connectivity | Connect over IPv6 |
| | | IPv4/IPv6 | Connect over IPv4 | Connect over IPv6 or IPv4[2] |
| Client is IPv6-only | IPv4-only | No connectivity | No connectivity | |
| | | IPv6-only | No connectivity | Connect over IPv6 |
| | | IPv4/IPv6 | No connectivity | Connect over IPv6 |

[1] IPv6 aware means that the client can also communicate over IPv4.

[2] When the DNS AAAA server returns the IPv6 and IPv4 addresses to the client, the client uses the connection protocol that applies to the address that is listed first.

# Securing IPv6 Networks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Windows Server 2003 currently supports the use of IPv6 only when IPv4 is also installed. Because TCP/IP internetworks are susceptible to a variety of attacks, ranging from passive attacks, such as eavesdropping, to active attacks, such as denial-of-service attacks, be sure to follow best practices for security when using IPv6 on your network. A few general and IPsec-related IPv6 security suggestions follow.

**General Recommendations for Securing IPv6 Networks**

Be aware of the following known security risks for IPv6, and consider reconfiguring your system to meet the recommendation that is shown for each:

- **The installation of an unauthorized router can cause reconfiguration of clients and rerouting of IPv6 traffic**.

  To communicate with IPv6 nodes on other network segments, IPv6 must use a default router. A default router is automatically assigned based on the receipt of a router advertisement. Malicious users with physical access to the IPv6-enabled network can install an unauthorized IPv6 router on the network segment, enabling a denial of service attack on IPv6 hosts. The unauthorized IPv6 router can reconfigure IPv6 clients, set itself as the default router, reroute link traffic, and disrupt other network services.

  **Recommendation**: Ensure that unauthorized individuals do not have physical or wireless access to your network. For more information, see Best practices for security in Help and Support Center for Windows Server 2003.

- Internet Connection Firewall (ICF) and Basic Firewall cannot filter or block IPv6 traffic.

  - **ICF**, which is available in the 32-bit versions of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition, restricts the traffic that can enter your network from the Internet. Because ICF can filter only IPv4 traffic, IPv6 traffic might get through the firewall and enter your network.

  - **Basic Firewall**, which is a component of Routing and Remote Access, can be enabled for public interfaces. Because Basic Firewall filters only IPv4 traffic, IPv6 traffic might get through the firewall and enter your network.

  **Recommendation**: If you are running IPv6 on your network, use firewall software or hardware that can filter and block IPv6 traffic.

- On-link computers (computers on a link, or LAN segment) can take control of another IPv6 address, causing on-link devices to create an incorrect entry in their neighbor cache.

  Nodes on an IPv6 link use address resolution to resolve a neighboring nodes IPv6 address to its link-layer address in the same way that nodes on an IPv4 link resolve addresses in IPv4. The resolved link-layer IPv6 address becomes an entry in a nodes neighbor cache (equivalent to the ARP cache in IPv4). If an attacker causes an IPv6 node to maliciously use another nodes address, it can cause other computers on the link to add a false entry to their neighbor cache. All traffic that is intended for the original computer goes instead to the attacker's computer, and the attacker can appear to send traffic from the original computer.

  **Recommendation**: Ensure that unauthorized individuals do not have physical or wireless access to your network. For more information, see Best practices for security in Help and Support Center for Windows Server 2003.

- When native IPv6 connectivity is not present (that is, when an IPv6 router does not exist on the network segment), it is easier to *spoof*, or appear to use, off-link IPv6 source addresses.

  A common defense against IP source address spoofing involves using packet-filtering routers. However, because traffic between hosts on the same link does not cross a router, this protective filtering is not used and spoofing might go undetected.

  Although spoofing can occur in native IPv4 or IPv6 networks, where on-link hosts can spoof off-link addresses in

communication with other hosts on the link, the threat is greater outside the native IPv6 network. Outside the native IPv6 network, encapsulation technologies are used. Because the logical link used for packet encapsulation spans a large portion of an IPv4 network (often the entire IPv4 Internet), an attacker can be anywhere on that IPv4 network and still spoof off-link addresses.

**Recommendation**: Ensure that an IPv6 router is present on the network segment.

For more information about IPv6 security, see "Security information for IPv6" in Help and Support Center for Windows Server 2003.

**Securing with IPsec on IPv6 Networks**

The IPv6 protocol for Windows Server 2003 incorporates Internet Protocol security (IPsec), which protects IPv6 data as it is sent over the network. *IPsec* is a set of Internet standards that uses cryptographic security services to provide the following:

- **Confidentiality**. Captured IPsec traffic cannot be deciphered without the appropriate encryption key.
- **Data origin authentication**. IPsec traffic contains a cryptographic checksum that incorporates a shared encryption key so that the receiver can verify that it was actually sent by the apparent sender.
- **Data integrity**. The cryptographic checksum is also used by the receiver to verify that the packet was not modified in transit.

IPsec in IPv6 is separate from, and does not interoperate with, IPsec for TCP/IP. IPsec policies that are configured with the IP Security Policies or Group Policy snap-ins have no effect on IPv6 traffic. You need to manually configure IPsec policies, security associations (SAs), and keys for IPv6. For more information about IPsec for TCP/IP, see Internet Protocol security (IPsec) in Help and Support Center for Windows Server 2003. For an example of configuring IPsec for IPv6, see "Using IPsec between two local link hosts" in Help and Support Center for Windows Server 2003.

In addition to the general security recommendations listed in the preceding section, follow these guidelines when you use IPv6 with IPsec:

- **Do not use IPsec for IPv6 in a production environment**. The current implementation of IPsec for IPv6 is not recommended for use in a production environment because it relies on static keying, which means that it has no provisions for updating encryption keys when sequence numbers are reused.
- **Use random numbers to configure SPIs**. When you manually configure IPsec Security Parameters Indexes (SPIs) for IPv6, always use random numbers so that you do not compromise the security of your IPsec for IPv6 policies.
- **Use only supported Encapsulating Security Payload (ESP) (ESP with NULL encryption)**. The IPv6 protocol for Windows Server 2003 does not support the use of IPsec ESP data encryption. However, the use of ESP with NULL encryption is supported. Although NULL encryption uses the ESP header, only data origin authentication and data integrity services are provided.

For more information about IPv6 security features, see "Security features for IPv6" in Help and Support Center for Windows Server 2003.

# Installing or Removing IPv6

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

When you initially install IPv6 on a computer in your network, be sure to install it first in a test environment so that you can configure and test the IPv6 protocol before installing it in your production environment. For more information about setting up a test lab, including typical test lab tasks and a sample infrastructure, see Setting up an IPv6 test lab in Help and Support Center for Windows Server 2003.

**To install IPv6**

1. In Control Panel, double-click **Network Connections**.
2. Right-click any local area connection, and then click **Properties**.
3. Click **Install**.
4. In the **Select Network Component Type** dialog box, click **Protocol**, and then click **Add**.
5. In the **Select Network Protocol** dialog box, click **Microsoft TCP/IP version 6**, and then click **OK**.
6. Click **Close** to save changes to your network connection.

After you install the IPv6 protocol, IIS automatically supports IPv6 on your Web server. However, if a Web site is already running when you install IPv6, you must restart IIS before the site begins listening for IPv6 requests. Sites that you create after you enable IPv6 support automatically listen for IPv6.

For more information about how to install the IPv6 network protocol stack, see IP version 6 in Help and Support Center for Windows Server 2003. For information about restarting IIS, see Restarting IIS.

**To remove IPv6**

1. In Control Panel, double-click **Network Connections**.
2. Right-click any local area connection, and then click **Properties**.
3. Click **Microsoft TCP/IP version 6** in the list of installed components, and then click **Uninstall.**
4. In the **Uninstall Microsoft TCP/IP version 6** dialog box, click **Yes**.
5. Click **Close** to save changes to your network connection.
6. Click **Yes** to restart your computer.

   Windows prompts you to restart your computer so that the new settings can take effect.

# Additional Resources for IPv6 and IIS 6.0

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

These resources contain additional information and tools related to IPv6.

**Related Information**

- Davies, J. *Understanding IPv6*. Redmond, WA: Microsoft Press, 2003 for information about using IPv6 with Windows Server 2003 and Windows XP.
- Howard, Michael & LeBlanc, David. *Writing Secure Code*. Second Edition. Redmond, WA: Microsoft Press, 2003 for information about buffer overruns.
- Huitema, Christian. *IPv6: The New Internet Protocol*. Second edition. Upper Saddle River, NJ: Prentice Hall PTR, 1998 for information about IPv6.
- The Internet Engineering Task Force (IETF) for a repository of RFCs.
- Miller, M. *Implementing IPv6: Supporting the Next Generation of Protocols*. Second edition. Foster City, CA: M&T Books, 2000 for information about using IPv6.

**Related Windows Server 2003 Help Topics**

For best results in identifying Help topics by title, in Help and Support Center, under the **Search** box, click **Set search options**. Under **Help Topics**, select the**Search in title only** check box.

- Internet Protocol security (IPsec) in Help and Support Center for Windows Server 2003.
- Introduction to IPv6 in Help and Support Center for Windows Server 2003.
- IP version 6 in Help and Support Center for Windows Server 2003 for information about the IPv6 protocol and how you can use it with Windows Server 2003 and IIS 6.0.
- IPv6 features in Help and Support Center for Windows Server 2003.
- Security information for IPv6 in Help and Support Center for Windows Server 2003.
- Security features for IPv6 in Help and Support Center for Windows Server 2003.
- Setting up an IPv6 test lab in Help and Support Center for Windows Server 2003.

**Related Tools**

- Ipsec6.exe

  Use the Ipsec6.exe command-line tool to configure IPsec policies and security associations. For more information about Ipsec6.exe, see IPv6 utilities in Help and Support Center for Windows Server 2003.

- Netsh.exe

  Use the command-line tool Netsh.exe to configure interfaces, addresses, caches, and routes for IPv6 and to display them on your screen. For more information about Netsh.exe, see Windows Support Tools Help on the Windows Server 2003 operating system CD.

# Running 32-bit Applications on 64-bit Windows

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Windows Server 2003™, Service Pack 1 enables IIS 6.0 to run 32-bit Web applications on 64-bit Windows using the Windows-32-on-Windows-64 (WOW64) compatibility layer. IIS 6.0 using WOW64 is intended to run 32-bit personal productivity applications needed by software developers and administrators, including 32-bit Internet Information Services (IIS) Web applications.

On 64-bit Windows, 32-bit processes cannot load 64-bit DLLs, and 64-bit processes cannot load 32-bit DLLs. If you plan to run 32-bit applications on 64-bit Windows, you must configure IIS to create 32-bit worker processes. Once you have configured IIS to create 32-bit worker processes, you can run the following types of IIS applications on 64-bit Windows:

- Internet Server API (ISAPI) extensions
- ISAPI filters
- Active Server Page (ASP) applications (specifically, scripts calling COM objects where the COM object can be 32-bit or 64-bit)
- ASP.NET applications

IIS can, by default, launch Common Gateway Interface (CGI) applications on 64-bit Windows, because CGI applications run in a separate process.

**Configuring IIS to run 32-bit Web applications on 64-bit Windows**
Before you configure IIS to run 32-bit applications on 64-bit Windows, note the following:

- IIS only supports 32-bit worker processes in Worker Process Isolation mode on 64-bit Windows.
- On 64-bit Windows, the World Wide Web Publishing service can run 32-bit and 64-bit worker processes. Other IIS services like the IIS Admin service, the SMTP service, the NNTP service, and the FTP service run 64-bit processes only.
- On 64-bit Windows, the World Wide Web Publishing service does not support running 32-bit and 64-bit worker processes concurrently on the same server.

For the procedure to run 32-bit Web applications on 64-bit Windows, see Configuring IIS to Run 32-bit Applications on 64-bit Windows.

**After configuring IIS to run 32-bit Web applications on 64-bit Windows**
After you configure IIS 6.0 to run 32-bit Web applications, IIS stores 32-bit DLLs and ISAPIs in the *%windir%*\syswow64\inetsrv directory. All other IIS files, including the MetaBase.xml file, are stored in the *%windir%*\system32\inetsrv directory. File access to the System32 and sub directories are transparently redirected based on the bitness of the process making that file access (64-bit processes have full access, while 32-bit processes have access to System32 redirected to Syswow64). If your legacy applications have specific 32-bit file access needs and you notice application failures, see if the application needs to reference the new *%windir%*\syswow64\inetsrv to resolve the problem.

# IIS 6.0 Administration Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The reference section provides detailed information about the interfaces, objects, and data structures available for programming and administering IIS.

**In This Section**

This section contains the following topics:

- Metabase Property Reference: This section defines the metabase properties used for IIS and includes code examples that show how to access the properties.
- IIS 6.0 Registry Reference: This section defines the registry entries that pertain to IIS.
- Events Reference: This section describes the events that IIS lists in the Event Viewer and explains how to use them to monitor your Web server.
- IIS 6.0 Event Messages: This section describes the Error, Warning, and Information event messages generated by IIS 6.0, and the logs to use to monitor Web server performance.
- Performance Counters Reference: This section describes the counters that you can enable in the System Monitor MMC snap-in, allowing you to evaluate the performance of your Web server.
- IIS 6.0 Performance Counters: This section describes performance objects, the components that generate the data, and how you can use Windows Management Instrumentation (WMI) to view these counters.
- Timeouts Reference: This section defines the IIS time-out properties and explains how to individually set them to configure your Web server for different tasks.
- Redirect Reference: This section explains how to format a server redirection correctly. For reference topics related to ASP, see the IIS Web Application Technology Reference section of the IIS Software Developer Kit (SDK).
- Tracing Events Reference: This reference describes the events that can be recorded in the trace log when using request-based tracing or IIS Admin tracing.

# Metabase Property Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This section contains reference pages for each of the Internet Information Services (IIS) 6.0 metabase properties. Most of these properties can be changed using IIS Manager. You can also use the methods of the IIS Windows Management Instrumentation (WMI) provider or the IIS Active Directory Service Interfaces (ADSI) provider in command-line scripts to change these properties.

As an alternative to writing your own command-line scripts, IIS installs Command-Line Administration Scripts that provide enough functionality to help you administer your server. Before you use any of these methods, please review the IIS Metabase section.To view a particular property, click the property name from the table of contents, or from the table below.

**Key**
@Application-level property: Application-level properties, when set, apply to the entire application regardless of the isolation mode in which your IIS server is running. Setting an application-level property in one location does not affect the entire application pool (see individual property topics for exceptions).

#Process-level property: Process-level properties only have an effect on IIS behavior if they are set at the service level (for both isolation modes), or at the server level (for IIS 5.0 isolation mode, if the application is set to High isolation). The documentation might list other locations where IIS allows you to set these properties, but they have no effect on IIS behavior.

**A**

| | |
|---|---|
| AccessFlags Metabase Property | AccessSSLFlags Metabase Property |
| ADConnectionsPassword Metabase Property | ADConnectionsUserName Metabase Property |
| AddNoHeaders Metabase Property | AdminACL Metabase Property |
| AdminACLBin Metabase Property | AdminEmail Metabase Property |
| AdminName Metabase Property | AdminServer Metabase Property |
| AllowAnonymous Metabase Property | AllowClientPosts Metabase Property |
| AllowControlMsgs Metabase Property | AllowFeedPosts Metabase Property |
| AllowKeepAlive Metabase Property | AllowPathInfoForScriptMappings Metabase Property |
| AlwaysUseSsl Metabase Property | AnonymousOnly Metabase Property |
| AnonymousPasswordSync Metabase Property | AnonymousUserName Metabase Property |
| AnonymousUserPass Metabase Property | AppAllowClientDebug Metabase Property |
| AppAllowDebugging Metabase Property | AppFriendlyName Metabase Property |
| AppIsolated Metabase Property | AppOopRecoverLimit Metabase Property |
| AppPackageID Metabase Property | AppPackageName Metabase Property |
| AppPoolAutoStart Metabase Property | AppPoolCommand Metabase Property |
| AppPoolId Metabase Property | AppPoolIdentityType Metabase Property |
| AppPoolQueueLength Metabase Property | AppPoolState Metabase Property |
| AppRoot Metabase Property | AppWamClsid Metabase Property |
| ApplicationDependencies Metabase Property | ArticleTableFile Metabase Property |

| | |
|---|---|
| ArticleTimeLimit Metabase Property | AspAllowOutOfProcComponents Metabase Property |
| AspAllowSessionState Metabase Property | AspAppServiceFlags Metabase Property |
| AspBufferingLimit Metabase Property | AspBufferingOn Metabase Property |
| AspCalcLineNumber Metabase Property | AspCodepage Metabase Property |
| AspDiskTemplateCacheDirectory Metabase Property | AspEnableApplicationRestart Metabase Property |
| AspEnableAspHtmlFallback Metabase Property | AspEnableChunkedEncoding Metabase Property |
| AspEnableParentPaths Metabase Property | AspEnableTypelibCache Metabase Property |
| AspErrorsToNTLog Metabase Property | AspExceptionCatchEnable Metabase Property |
| AspExecuteInMTA Metabase Property | AspKeepSessionIDSecure Metabase Property |
| AspLCID Metabase Property | AspLogErrorRequests Metabase Property |
| AspMaxDiskTemplateCacheFiles Metabase Property | AspMaxRequestEntityAllowed Metabase Property |
| AspPartitionID Metabase Property | AspProcessorThreadMax Metabase Property |
| AspQueueConnectionTestTime Metabase Property | AspQueueTimeout Metabase Property |
| AspRequestQueueMax Metabase Property | AspRunOnEndAnonymously Metabase Property |
| AspScriptEngineCacheMax Metabase Property | AspScriptErrorMessage Metabase Property |
| AspScriptErrorSentToBrowser Metabase Property | AspScriptFileCacheSize Metabase Property |
| AspScriptLanguage Metabase Property | AspScriptTimeout Metabase Property |
| AspSessionMax Metabase Property | AspSessionTimeout Metabase Property |
| AspSxsName Metabase Property | AspTrackThreadingModel Metabase Property |
| AuthChangeURL Metabase Property | AuthExpiredUnsecureURL Metabase Property |
| AuthExpiredURL Metabase Property | AuthFlags Metabase Property |
| AuthNotifyPwdExpURL Metabase Property | AuthNotifyPwdExpUnsecureURL Metabase Property |
| AuthPersistence Metabase Property | AuthTurnList Metabase Property |
| AutoShutdownAppPoolExe Metabase Property | AutoShutdownAppPoolParams Metabase Property |
| AzEnable Metabase Property | AzImpersonationLevel Metabase Property |
| AzScopeName Metabase Property | AzStoreName Metabase Property |

**B**

| | |
|---|---|
| BadMailDirectory Metabase Property | BannerMessage Metabase Property |
| BinSchemaTimeStamp Metabase Property | |

**C**

| | |
|---|---|
| CacheControlCustom Metabase Property | CacheControlMaxAge Metabase Property |
| CacheControlNoCache Metabase Property | CacheISAPI Metabase Property |

| | |
|---|---|
| CentralBinaryLoggingEnabled Metabase Property | CertCheckMode Metabase Property |
| CGITimeout Metabase Property | ChangeNumber Metabase Property |
| ClientPostHardLimit Metabase Property | ClusterEnabled Metabase Property |
| CollectionComment Metabase Property | ClientPostSoftLimit Metabase Property |
| ConnectResponse Metabase Property | ConnectionTimeout Metabase Property |
| ContentIndexed Metabase Property | CPUAction Metabase Property |
| CPULimit Metabase Property | CPUResetInterval Metabase Property |
| CreateCGIWithNewConsole Metabase Property | CreateProcessAsUser Metabase Property |
| CSideEtrnDomains Metabase Property | CustomErrorDescriptions Metabase Property |

**D**

| | |
|---|---|
| DefaultDoc Metabase Property | DefaultDocFooter Metabase Property |
| DefaultDomain Metabase Property | DefaultLogonDomain Metabase Property |
| DefaultModeratorDomain Metabase Property | DemandStartThreshold Metabase Property |
| DirBrowseFlags Metabase Property | DirectoryLevelsToScan Metabase Property |
| DisableNewNews Metabase Property | DisableSocketPooling Metabase Property |
| DisableStaticFileCache Metabase Property | DisallowOverlappingRotation Metabase Property |
| DisallowRotationOnConfigChange Metabase Property | DoDynamicCompression Metabase Property |
| DomainRouting Metabase Property | DoMasquerade Metabase Property |
| DontLog Metabase Property | DoStaticCompression Metabase Property |
| DownlevelAdminInstance Metabase Property | DropDirectory Metabase Property |

**E**

| | |
|---|---|
| EnableDocFooter Metabase Property | EnableEditWhileRunning Metabase Property |
| EnableHistory Metabase Property | EnableReverseDns Metabase Property |
| EnableReverseDnsLookup Metabase Property | EtrnDays Metabase Property |
| EtrnSubdomains Metabase Property | ExitMessage Metabase Property |
| ExMdbGuid Metabase Property | ExpireNewsgroups Metabase Property |
| ExpirePolicyName Metabase Property | ExpireSpace Metabase Property |
| ExpireTime Metabase Property | |

**F**

| | |
|---|---|
| FeedAccountName Metabase Property | FeedAdminError Metabase Property |
| FeedAllowControlMsgs Metabase Property | FeedAuthenticationType Metabase Property |

| | |
|---|---|
| FeedConcurrentSessions Metabase Property | FeedCreateAutomatically Metabase Property |
| FeedDisabled Metabase Property | FeedDistribution Metabase Property |
| FeedErrParmMask Metabase Property | FeedFeedpairId Metabase Property |
| FeedHandshake Metabase Property | FeedInterval Metabase Property |
| FeedMaxConnectionAttempts Metabase Property | FeedNewsgroups Metabase Property |
| FeedNextPullHigh Metabase Property | FeedNextPullLow Metabase Property |
| FeedOutgoingPort Metabase Property | FeedPassword Metabase Property |
| FeedPeerGapSize Metabase Property | FeedPeerTempDirectory Metabase Property |
| FeedPostHardLimit Metabase Property | FeedPostSoftLimit Metabase Property |
| FeedReportPeriod Metabase Property | FeedSecurityType Metabase Property |
| FeedServerName Metabase Property | FeedStartTimeHigh Metabase Property |
| FeedStartTimeLow Metabase Property | FeedTempDirectory Metabase Property |
| FeedType Metabase Property | FeedUucpName Metabase Property |
| FilterDescription Metabase Property | FilterEnableCache Metabase Property |
| FilterEnabled Metabase Property | FilterFlags Metabase Property |
| FilterLoadOrder Metabase Property | FilterPath Metabase Property |
| FilterState Metabase Property | FrontPageWeb Metabase Property |
| FsPropertyPath Metabase Property | FtpDirBrowseShowLongDate Metabase Property |
| FtpLogInUtf8 Metabase Property | FullyQualifiedDomainName Metabase Property |

**G**

| | |
|---|---|
| GreetingMessage Metabase Property | GroupHelpFile Metabase Property |
| GroupListFile Metabase Property | GroupvarListFile Metabase Property |

**H**

| | |
|---|---|
| HcCacheControlHeader Metabase Property | HcCompressionBufferSize Metabase Property |
| HcCompressionDirectory Metabase Property | HcCompressionDll Metabase Property |
| HcCreateFlags Metabase Property | HcDoDiskSpaceLimiting Metabase Property |
| HcDoDynamicCompression Metabase Property | HcDoOnDemandCompression Metabase Property |
| HcDoStaticCompression Metabase Property | HcDynamicCompressionLevel Metabase Property |
| HcExpiresHeader Metabase Property | HcFileExtensions Metabase Property |
| HcFilesDeletedPerDiskFree Metabase Property | HcIoBufferSize Metabase Property |
| HcMaxDiskSpaceUsage Metabase Property | HcMaxQueueLength Metabase Property |
| HcMinFileSizeForComp Metabase Property | HcNoCompressionForHttp10 Metabase Property |

| | |
|---|---|
| HcNoCompressionForProxies Metabase Property | HcNoCompressionForRange Metabase Property |
| HcOnDemandCompLevel Metabase Property | HcPriority Metabase Property |
| HcScriptFileExtensions Metabase Property | HcSendCacheHeaders Metabase Property |
| HeaderWaitTimeout Metabase Property | HistoryExpiration Metabase Property |
| HistoryMajorVersionNumber Metabase Property | HistoryTableFile Metabase Property |
| HonorClientMsgIds Metabase Property | HopCount Metabase Property |
| HttpCustomHeaders Metabase Property | HttpErrors Metabase Property |
| HttpExpires Metabase Property | HttpPics Metabase Property |
| HttpRedirect Metabase Property | |

**I**

| | |
|---|---|
| IdleTimeout Metabase Property | Iis5IsolationModeEnabled Metabase Property |
| ImapClearTextProvider Metabase Property | ImapDefaultDomain Metabase Property |
| ImapDsAccount Metabase Property | ImapDsBindType Metabase Property |
| ImapDsDataDirectory Metabase Property | ImapDsDefaultMailRoot Metabase Property |
| ImapDsHost Metabase Property | ImapDsNamingContext Metabase Property |
| ImapDsPassword Metabase Property | ImapDsSchemaType Metabase Property |
| ImapExpireDelay Metabase Property | ImapExpireMail Metabase Property |
| ImapExpireStart Metabase Property | ImapMailExpirationTime Metabase Property |
| ImapRoutingDll Metabase Property | ImapRoutingSources Metabase Property |
| ImapRoutingTableType Metabase Property | ImapServiceVersion Metabase Property |
| InProcessIsapiApps Metabase Property | IPSecurity Metabase Property |
| IsapiRestrictionList Metabase Property | |

**J**

**K**

| |
|---|
| KeyType Metabase Property |

**L**

| | |
|---|---|
| LimitRemoteConnections Metabase Property | ListFile Metabase Property |
| LoadBalancerCapabilities Metabase Property | LocalDomains Metabase Property |
| LocalRetryAttempts Metabase Property | LocalRetryInterval Metabase Property |
| LogAnonymous Metabase Property | LogCustomPropertyDataType Metabase Property |
| LogCustomPropertyHeader Metabase Property | LogCustomPropertyID Metabase Property |
| LogCustomPropertyMask Metabase Property | LogCustomPropertyName Metabase Property |

| | |
|---|---|
| LogCustomPropertyNodeID Metabase Property | LogCustomPropertyServicesString Metabase Property |
| LogEventOnRecycle Metabase Property | LogExtFileFlags Metabase Property |
| LogFileDirectory Metabase Property | LogFileLocaltimeRollover Metabase Property |
| LogFilePeriod Metabase Property | LogFileTruncateSize Metabase Property |
| LogModuleId Metabase Property | LogModuleList Metabase Property |
| LogInUTF8 Metabase Property | LogModuleUiId Metabase Property |
| LogNonAnonymous Metabase Property | LogOdbcDataSource Metabase Property |
| LogOdbcPassword Metabase Property | LogOdbcTableName Metabase Property |
| LogOdbcUserName Metabase Property | LogonMethod Metabase Property |
| LogPluginClsid Metabase Property | LogType Metabase Property |

**M**

| | |
|---|---|
| MajorIlsVersionNumber Metabase Property | MasqueradeDomain Metabase Property |
| MaxBandwidth Metabase Property | MaxBandwidthBlocked Metabase Property |
| MaxBatchedMessages Metabase Property | MaxClientsMessage Metabase Property |
| MaxConnections Metabase Property | MaxDirChangeIOSize Metabase Property |
| MaxEndpointConnections Metabase Property | MaxErrorFiles Metabase Property |
| MaxGlobalBandwidth Metabase Property | MaxHistoryFiles Metabase Property |
| MaxMailObjects Metabase Property | MaxMessageSize Metabase Property |
| MaxOutConnections Metabase Property | MaxOutConnectionsPerDomain Metabase Property |
| MaxProcesses Metabase Property | MaxRecipients Metabase Property |
| MaxRequestEntityAllowed Metabase Property | MaxSearchResults Metabase Property |
| MaxSessionSize Metabase Property | MaxSmtpErrors Metabase Property |
| MaxSmtpLogonErrors Metabase Property | MD_0 Metabase Property |
| MD_1001 Metabase Property | MD_APP_LAST_OUTPROC_PID Metabase Property |
| MD_ETAG_CHANGENUMBER Metabase Property | MD_IISADMIN_EXTENSIONS Metabase Property |
| MD_ISM_ACCESS_CHECK Metabase Property | MD_METADATA_ID_REGISTRATION Metabase Property |
| MD_SERVER_CAPABILITIES Metabase Property | MD_SERVER_PLATFORM Metabase Property |
| MimeMap Metabase Property | MinFileBytesPerSec Metabase Property |
| MinorIlsVersionNumber Metabase Property | ModeratorFile Metabase Property |
| MSDOSDirOutput Metabase Property | |

**N**

| | |
|---|---|
| NameResolutionType Metabase Property | NewsCrawlerTime Metabase Property |

| | |
|---|---|
| NewsDropDirectory Metabase Property | NewsFailedPickupDirectory Metabase Property |
| NewPickupDirectory Metabase Property | NntpClearTextProvider Metabase Property |
| NntpCommandLogMask Metabase Property | NntpOrganization Metabase Property |
| NntpServiceVersion Metabase Property | NntpUucpName Metabase Property |
| NotDeletable Metabase Property | NTAuthenticationProviders Metabase Property |

**O**

| | |
|---|---|
| OrphanActionExe Metabase Property | OrphanActionParams Metabase Property |
| OrphanWorkerProcess Metabase Property | |

**P**

| | |
|---|---|
| PassivePortRange Metabase Property | PassportRequireADMapping Metabase Property |
| PasswordCacheTTL Metabase Property | PasswordChangeFlags Metabase Property |
| PasswordExpirePrenotifyDays Metabase Property | Path Metabase Property |
| PeriodicRestartMemory Metabase Property | PeriodicRestartPrivateMemory Metabase Property |
| PeriodicRestartRequests Metabase Property | PeriodicRestartSchedule Metabase Property |
| PeriodicRestartTime Metabase Property | PickupDirectory Metabase Property |
| PingingEnabled Metabase Property | PingInterval Metabase Property |
| PingResponseTime Metabase Property | PoolIdcTimeout Metabase Property |
| Pop3ClearTextProvider Metabase Property | Pop3DefaultDomain Metabase Property |
| Pop3DsAccount Metabase Property | Pop3DsBindType Metabase Property |
| Pop3DsDataDirectory Metabase Property | Pop3DsDefaultMailRoot Metabase Property |
| Pop3DsHost Metabase Property | Pop3DsNamingContext Metabase Property |
| Pop3DsPassword Metabase Property | Pop3DsSchemaType Metabase Property |
| Pop3ExpireDelay Metabase Property | Pop3ExpireMail Metabase Property |
| Pop3ExpireStart Metabase Property | Pop3MailExpirationTime Metabase Property |
| Pop3RoutingDll Metabase Property | Pop3RoutingSources Metabase Property |
| Pop3RoutingTableType Metabase Property | Pop3ServiceVersion Metabase Property |
| PostmasterEmail Metabase Property | PostmasterName Metabase Property |
| PrettyNamesFile Metabase Property | ProcessNTCRIfLoggedOn Metabase Property |

**Q**

| |
|---|
| QueueDirectory Metabase Property |

**R**

| | |
|---|---|
| RapidFailProtection Metabase Property | RapidFailProtectionInterval Metabase Property |
| RapidFailProtectionMaxCrashes Metabase Property | Realm Metabase Property |
| RedirectHeaders Metabase Property | RelayForAuth Metabase Property |
| RelayIpList Metabase Property | RemoteRetryAttempts Metabase Property |
| RemoteRetryInterval Metabase Property | RemoteSmtpPort Metabase Property |
| RemoteSmtpSecurePort Metabase Property | RemoteTimeout Metabase Property |
| RestrictionListCustomDesc Metabase Property | RevocationFreshnessTime Metabase Property |
| RevocationURLRetrievalTimeout Metabase Property | RouteAction Metabase Property |
| RouteActionString Metabase Property | RoutePassword Metabase Property |
| RouteUserName Metabase Property | RoutingDll Metabase Property |
| RoutingSources Metabase Property | |

**S**

| | |
|---|---|
| SaslLogonDomain Metabase Property | ScriptMaps Metabase Property |
| SecureBindings Metabase Property | SecuritySetupRequired Metabase Property |
| SendBadTo Metabase Property | SendNdrTo Metabase Property |
| ServerAutoStart Metabase Property | ServerBindings Metabase Property |
| ServerCommand Metabase Property | ServerComment Metabase Property |
| ServerConfigFlags Metabase Property | ServerID Metabase Property |
| ServerListenBacklog Metabase Property | ServerListenTimeout Metabase Property |
| ServerSize Metabase Property | ServerState Metabase Property |
| SessionKey Metabase Property | ShouldDeliver Metabase Property |
| ShouldPickupMail Metabase Property | ShouldPipelineIn Metabase Property |
| ShouldPipelineOut Metabase Property | ShutdownLatency Metabase Property |
| ShutdownTimeLimit Metabase Property | SmartHost Metabase Property |
| SmartHostType Metabase Property | SMPAffinitized Metabase Property |
| SMPProcessorAffinityMask Metabase Property | SmtpAdvQueueDll Metabase Property |
| SmtpAqueueWait Metabase Property | SmtpAuthTimeout Metabase Property |
| SmtpBdatTimeout Metabase Property | SmtpClearTextProvider Metabase Property |
| SmtpCommandLogMask Metabase Property | SmtpConnectTimeout Metabase Property |
| SmtpDataTimeout Metabase Property | SmtpDisableRelay Metabase Property |
| SmtpDomainValidationFlags Metabase Property | SmtpDotStuffPickupDirFiles Metabase Property |
| SmtpDsAccount Metabase Property | SmtpDsBindType Metabase Property |
| SmtpDsDataDirectory Metabase Property | SmtpDsDefaultMailRoot Metabase Property |

| | |
|---|---|
| SmtpDsDomain Metabase Property | SmtpDsFlags Metabase Property |
| SmtpDsHost Metabase Property | SmtpDsNamingContext Metabase Property |
| SmtpDSNLanguageID Metabase Property | SmtpDSNOptions Metabase Property |
| SmtpDsPassword Metabase Property | SmtpDsPort Metabase Property |
| SmtpDsSchemaType Metabase Property | SmtpDsUseCat Metabase Property |
| SmtpEventlogLevel Metabase Property | SmtpFlushMailFile Metabase Property |
| SmtpHeloNoDomain Metabase Property | SmtpHeloTimeout Metabase Property |
| SmtpInboundCommandSupportOptions Metabase Property | SmtpIpRestrictionFlag Metabase Property |
| SmtpLocalDelayExpireMinutes Metabase Property | SmtpLocalNDRExpireMinutes Metabase Property |
| SmtpMailFromTimeout Metabase Property | SmtpMailNoHelo Metabase Property |
| SmtpMaxRemoteQThreads Metabase Property | SmtpOutboundCommandSupportOptions Metabase Property |
| SmtpRcptToTimeout Metabase Property | SmtpRemoteDelayExpireMinutes Metabase Property |
| SmtpRemoteNDRExpireMinutes Metabase Property | SmtpRemoteProgressiveRetry Metabase Property |
| SmtpRemoteRetryThreshold Metabase Property | SmtpRoutingTableType Metabase Property |
| SmtpRsetTimeout Metabase Property | SmtpSaslTimeout Metabase Property |
| SmtpServer Metabase Property | SmtpServiceVersion Metabase Property |
| SmtpSSLCertHostnameValidation Metabase Property | SmtpSSLRequireTrustedCA Metabase Property |
| SmtpTurnTimeout Metabase Property | SmtpUseTcpDns Metabase Property |
| SSIExecDisable Metabase Property | SSLAlwaysNegoClientCert Metabase Property |
| SSLCertHash Metabase Property | SslCtlIdentifier Metabase Property |
| SslCtlStoreName Metabase Property | SSLStoreName Metabase Property |
| SslUseDsMapper Metabase Property | StartupTimeLimit Metabase Property |
| SuppressDefaultFTPBanner Metabase Property | |

**T**

**U**

| | |
|---|---|
| UNCPassword Metabase Property | UNCUserName Metabase Property |
| UpdatedDefaultDomain Metabase Property | UpdatedFQDN Metabase Property |
| UploadReadAheadSize Metabase Property | UseDigestSSP Metabase Property |
| UserIsolationMode Metabase Property | |

**V**

| | |
|---|---|
| VrDoExpire Metabase Property | VrDriverClsid Metabase Property |
| VrDriverProgid Metabase Property | VrOwnModerator Metabase Property |

| | |
|---|---|
| VrUseAccount Metabase Property | |

**W**

| | |
|---|---|
| WAMUserName Metabase Property | WAMUserPass Metabase Property |
| WebSvcExtRestrictionList Metabase Property | Win32Error Metabase Property |

**X**

| | |
|---|---|
| XMLSchemaTimeStamp Metabase Property | XoverTableFile Metabase Property |

**Y**

**Z**

# Code Examples to Configure Metabase Properties

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following examples demonstrate the syntax for setting a metabase property using different technologies, such as Windows Management Instrumentation (WMI) and Active Directory Service Interfaces (ADSI). The code examples are written in Microsoft® Visual Basic® Scripting Edition (VBScript), JScript®, and Perl.

The examples are simplified to show the essential code that is needed to change a metabase property. If you use these examples, add error checking code and security checking code before using them to configure your server.

**Implementation Guidelines**
**To implement one of the VBScript, JScript, or Perl examples in a command-line administration script**

- For VBScript: Save the script file with the file name extension of .vbs and call a script with "cscript //nologo *filename*.vbs".
- For JScript: Save the script file with the file name extension of .js and call a script with "cscript //nologo *filename*.js".
- For Perl: Save the script file with the file name extension of .pl and call a script with "perl *filename*.pl".

**To implement one of the VBScript, JScript, or Perl examples in an ASP page**

1. Set security on the Web application to disable anonymous access. Anonymous access to metabase properties is not allowed. The user requesting the ASP page should be an Administrator. In addition, the process executing the ASP page must have permission to modify the metabase.
2. Add the following lines of ASP code, where the first line is at the top of the file, and "*language*" is "Microsoft Visual Basic Scripting Edition (VBScript)", "JScript", or "PerlScript":

```
<%@ Language="language" %>
<%
insert the example code here
%>
```

3. Change instances of **WScript.Echo** to **Response.Write**, and **print** to **$Response->Write**.
4. Change the file name extension of your file to .asp.
5. Install ActivePerl with PerlScript to run the Perl examples. ActivePerl can be downloaded from the ActiveState Web site.

**Code Examples**
In the following examples, there are italicized words that must be replaced by a valid string for the example to execute properly. Replace *property_name_1* and *property_name_2* with metabase property names. Replace *numeric_or_Boolean_value* and *String_value* with a valid data types for that property.

These examples demonstrate how to set a metabase property to a string, number, or Boolean value. These examples do not cover WMI situations where you need to pass in an object.

**WMI**

| | |
|---|---|
| Microsoft Visual Basic Scripting Edition (VBScript) | `var providerObj=GetObject("winmgmts:/root/MicrosoftIISv2");`<br>`var vdirObj=providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'");`<br>`// Print out the current value of some properties:`<br>`WScript.Echo("Before: " + vdirObj.property_name_1 + ", " + vdirObj.property_name_2);`<br>`// Set some properties:`<br>`vdirObj.property_name_1=numeric_or_Boolean_value;`<br>`vdirObj.property_name_2="String_value";`<br>`// Save the property changes in the metabase:`<br>`vdirObj.Put_();`<br>`WScript.Echo("After: " + vdirObj.property_name_1 + ", " + vdirObj.property_name_2);` |
| JScript | `var providerObj=GetObject("winmgmts:/root/MicrosoftIISv2");`<br>`var vdirObj=providerObj.get("IIsWebVirtualDirSetting='W3SVC/1/ROOT'");`<br>`// Print out the current value of some properties:`<br>`WScript.Echo("Before: " + vdirObj.property_name_1 + ", " + vdirObj.property_name_2);`<br>`// Set some properties:`<br>`vdirObj.property_name_1=numeric_or_Boolean_value;`<br>`vdirObj.property_name_2="String_value";`<br>`// Save the property changes in the metabase:`<br>`vdirObj.Put_();`<br>`WScript.Echo("After: " + vdirObj.property_name_1 + ", " + vdirObj.property_name_2);` |
| Perl | `use Win32;`<br>`use Win32::OLE;`<br>`my $providerObj=Win32::OLE -> GetObject("winmgmts:/root/MicrosoftIISv2");`<br>`my $vdirObj=$providerObj->Get("IIsWebVirtualDirSetting='W3SVC/1'");`<br>`# Print out the current value of some properties:`<br>`print "Before: ", $vdirObj->{"property_name_1"}, ", ", $vdirObj->{"property_name_2"}, "\n";`<br>`# Set some properties:`<br>`$vdirObj->{"property_name_1"}=numeric_or_Boolean_value;`<br>`$vdirObj->{"property_name_2"}="String_value";`<br>`# Save the property changes in the metabase:`<br>`vdirObj->Put_();`<br>`print "After: ", $vdirObj->{"property_name_1"}, ", ", $vdirObj->{"property_name_2"}, "\n";` |

**ADSI**

| | |
|---|---|
| Microsoft Visual Basic Scripting Edition (VBScript) | `set vdirObj=GetObject("IIS://localhost/W3svc/1/Root")`<br>`' Print out the current value of some properties:`<br>`WScript.Echo "Before: " & vdirObj.property_name_1 & ", " & vdirObj.property_name_2`<br>`' Set some properties:`<br>`vdirObj.Put "property_name_1", numeric_or_Boolean_value`<br>`vdirObj.Put "property_name_2", "String_value"`<br>`' Save the property changes in the metabase:`<br>`vdirObj.SetInfo`<br>`WScript.Echo "After: " & vdirObj.property_name_1 & ", " & vdirObj.property_name_2` |
| JScript | `var vdirObj=GetObject("IIS://localhost/W3svc/1/Root");`<br>`// Print out the current value of some properties:`<br>`WScript.Echo("Before: " + vdirObj.property_name_1 + ", " + vdirObj.property_name_2);`<br>`// Set some properties:`<br>`vdirObj.Put("property_name_1", numeric_or_Boolean_value);`<br>`vdirObj.Put("property_name_2", "String_value");`<br>`// Save the property changes in the metabase:`<br>`vdirObj.SetInfo();`<br>`WScript.Echo("After: " + vdirObj.property_name_1 + ", " + vdirObj.property_name_2);` |

| | |
|---|---|
| Perl | `use Win32;`<br>`use Win32::OLE;`<br>`my $vdirObj=Win32::OLE ->`<br>`GetObject("IIS://localhost/W3SVC/1/Root");`<br>`# Print out the current value of some properties:`<br>`print "Before: ", $vdirObj->{"property_name_1"}, ", ",`<br>`$vdirObj->{"property_name_2"}, "\n";`<br>`# Set some properties:`<br>`$vdirObj->{"property_name_1"}=numeric_or_Boolean_value;`<br>`$vdirObj->{"property_name_2"}="String_value";`<br>`# Save the property changes in the metabase:`<br>`$vdirObj->SetInfo();`<br>`print "After: ", $vdirObj->{"property_name_1"}, ", ", $vdirObj-`<br>`>{"property_name_2"}, "\n";` |

You may encounter errors when running example code that you have altered. The following table lists possible errors and their causes.

| Error Code | Possible Cause |
|---|---|
| 0x80005006 | The property does not exist, or it does not exist at the access location where you are attempting to set it.<br><br>For example, you cannot set the AdminServer Metabase Property at the Web site level (W3SVC/1, W3SVC/1/ROOT) or the virtual directory level (W3SVC/1/ROOT/vdir); you can set it only at the World Wide Web Publishing Service (WWW service) level. You cannot set the AdminEmail Metabase Property on any Web, File Transfer Protocol (FTP), or Simple Mail Transfer Protocol (SMTP) metabase key; you can set it only on a Network News Transfer Protocol (NNTP) metabase key.<br><br>See the Metabase Property Reference for the metabase property that you are trying to set. Use the Access Locations table to verify that you are accessing the property at a valid location. Or, you can open the MetaBase.xml file on your server and search for the property to see what locations it is set on. |
| 0x8000500C | The data type of the property is different than the data type to which you are attempting to set it.<br><br>For example, you cannot set the AspCodepage Metabase Property property to a string ("English"); you can set it only to a number representing a code page (1252). |

**Related Information**

- For information about scripting in JScript or Microsoft Visual Basic Scripting Edition (VBScript), see the Microsoft Script Technologies SDK.
- For an article about scripting in PerlScript, see Active Scripting with PerlScript.
- For information about WMI and ADSI, including additional examples, see IIS Administration Technologies in the IIS SDK on MSDN.

# AccessFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AccessFlags** property contains flags for configuring file access permissions, which are defined in the Flags section. The last four flags augment the first four flags in that they relate to access granted to remote computers. You can enable local access without remote access, but you cannot enable remote access without local access.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ACCESS_PERM |
| Attributes | INHERIT |
| Default Value | accessread |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6016 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/MSFTPSVC/*n*/ROOT /LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

| | |
|---|---|
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/ROOT /LM/SMTPSVC/*n*/ROOT/*virtual_directory_name* | IIsSmtpVirtualDir |

**Flags**

| **Flag Name** | **AccessExecute** |
|---|---|
| Description | A value of **true** indicates that the file or the contents of the folder may be executed, regardless of file type. |
| Metabase Bitmask Identifier | MD_ACCESS_EXECUTE |
| Decimal Value | 4 |
| Hexadecimal Value | 0x00000004 |

| **Flag Name** | **AccessSource** |
|---|---|
| Description | A value of **true** indicates that users are allowed to access source code if either Read or Write permissions are set. Source code includes scripts in Microsoft Active Server Pages (ASP) applications. |
| Metabase Bitmask Identifier | MD_ACCESS_SOURCE |
| Decimal Value | 16 |
| Hexadecimal Value | 0x00000010 |

| **Flag Name** | **AccessRead** |
|---|---|
| Description | A value of **true** indicates that the file or the contents of the folder may be read through Microsoft Internet Explorer. |
| Metabase Bitmask Identifier | MD_ACCESS_READ |
| Decimal Value | 1 |
| Hexadecimal Value | 0x00000001 |

| **Flag Name** | **AccessWrite** |
|---|---|
| Description | A value of **true** indicates that users are allowed to upload files and their associated properties to the enabled directory on your server or to change content in a Write-enabled file. Write can be implemented only with a browser that supports the PUT feature of the HTTP 1.1protocol standard. |
| Metabase Bitmask Identifier | MD_ACCESS_WRITE |
| Decimal Value | 2 |
| Hexadecimal Value | 0x00000002 |

| Flag Name | AccessScript |
| --- | --- |
| Description | A value of **true** indicates that the file or the contents of the folder may be executed if they are script files or static content. A value of **false** only allows static files, such as HTML files, to be served. |
| Metabase Bitmask Identifier | MD_ACCESS_SCRIPT |
| Decimal Value | 512 |
| Hexadecimal Value | 0x00000200 |

| Flag Name | AccessNoRemoteExecute |
| --- | --- |
| Description | A value of **true** indicates that remote requests to execute applications are denied; only requests from the same computer as the IIS server succeed if the **AccessExecute** property is set to **true**. You cannot set **AccessNoRemoteExecute** to false to enable remote requests, and set **AccessExecute** to **false** to disable local requests. |
| Metabase Bitmask Identifier | MD_ACCESS_NO_REMOTE_EXECUTE |
| Decimal Value | 8192 |
| Hexadecimal Value | 0x00000400 |

| Flag Name | AccessNoRemoteRead |
| --- | --- |
| Description | A value of **true** indicates that remote requests to view files are denied; only requests from the same computer as the IIS server succeed if the **AccessRead** property is set to **true**. You cannot set **AccessNoRemoteRead** to **false** to enable remote requests, and set **AccessRead** to **false** to disable local requests. |
| Metabase Bitmask Identifier | MD_ACCESS_NO_REMOTE_READ |
| Decimal Value | 4096 |
| Hexadecimal Value | 0x00001000 |

| Flag Name | AccessNoRemoteWrite |
| --- | --- |
| Description | A value of **true** indicates that remote requests to create or change files are denied; only requests from the same computer as the IIS server succeed if the **AccessWrite** property is set to **true**. You cannot set **AccessNoRemoteWrite** to **false** to enable remote requests, and set **AccessWrite** to **false** to disable local requests. |
| Metabase Bitmask Identifier | MD_ACCESS_NO_REMOTE_WRITE |
| Decimal Value | 1024 |
| Hexadecimal Value | 0x00000400 |

| Flag Name | AccessNoRemoteScript |
| --- | --- |
| Description | A value of **true** indicates that remote requests to view dynamic content are denied; only requests from the same computer as the IIS server succeed if the **AccessScript** property is set to **true**. You cannot set **AccessNoRemoteScript** to **false** to enable remote requests, and set **AccessScript** to **false** to disable local requests. |
| Metabase Bitmask Identifier | MD_ACCESS_NO_REMOTE_SCRIPT |
| Decimal Value | 16384 |
| Hexadecimal Value | 0x00004000 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AccessSSLFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AccessSSLFlags** property contains the Secure Sockets Layer (SSL) file permission flags as seen in the Flags section. The default value of 0 means that no SSL permissions are set.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SSL_ACCESS_PERM |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6030 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/ROOT /LM/SMTPSVC/*n*/ROOT/*virtual_directory_name* | IIsSmtpVirtualDir |

**Flags**

| **Flag Name** | **AccessSSL** |
|---|---|
| Description | A value of **true** indicates that file access requires SSL file permission processing with or without a client certificate. |
| Metabase Bitmask Identifier | MD_ACCESS_SSL |
| Decimal Value | 8 |
| Hexadecimal Value | 0x00000008 |

| **Flag Name** | **AccessSSL128** |
|---|---|
| Description | A value of **true** indicates that file access requires SSL file permission processing with a minimum key size of 128 bits, with or without a client certificate. |
| Metabase Bitmask Identifier | MD_ACCESS_SSL128 |
| Decimal Value | 256 |
| Hexadecimal Value | 0x00000100 |

| **Flag Name** | **AccessSSLNegotiateCert** |
|---|---|
| Description | A value of **true** indicates that SSL file access processing requests a certificate from the client. A value of **false** indicates that access continues if the client does not have a certificate. Some versions of Internet Explorer will close the connection if the server requests a certificate and a certificate is not available (even if **AccessSSLRequireCert** is also set to **false**). |
| Metabase Bitmask Identifier | MD_ACCESS_NEGO_CERT |
| Decimal Value | 32 |
| Hexadecimal Value | 0x00000020 |

| **Flag Name** | **AccessSSLRequireCert** |
|---|---|
| Description | A value of **true** indicates that SSL file access processing requests a certificate from the client. If the client provides no certificate, the connection is closed. **AccessSSLNegotiateCert** must also beset to **true** when using **AccessSSLRequireCert**. |
| Metabase Bitmask Identifier | MD_ACCESS_REQUIRE_CERT |
| Decimal Value | 64 |
| Hexadecimal Value | 0x00000040 |

| **Flag Name** | **AccessSSLMapCert** |
|---|---|
| Description | A value of **true** indicates that SSL file permission processing maps a client certificate to a Microsoft Windows operating system user-account. The **AccessSSLNegotiateCert** property must also be set to **true** for the mapping to occur. |

| | |
|---|---|
| Metabase Bitmask Identifier | MD_ACCESS_MAP_CERT |
| Decimal Value | 128 |
| Hexadecimal Value | 0x00000080 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

AccessFlags Metabase Property

# ADConnectionsPassword Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ADConnectionsPassword** property is used when the FTP property, UserIsolationMode Metabase Property, is set to 2, which makes FTP user isolation dependent upon the Active Directory directory service. Set **ADConnectionsUserName** to the user name that has access to the domain controller, and set **ADConnectionsPassword** to the corresponding password.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_AD_CONNECTIONS_PASSWORD |
| Attributes | INHERIT \| SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5015 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

UserIsolationMode Metabase Property

# ADConnectionsUserName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ADConnectionsUserName** property is used when the FTP property, UserIsolationMode Metabase Property, is set to 2, which makes FTP user isolation dependent upon Active Directory. Set **ADConnectionsUserName** to the user name that has access to the domain controller, and set **ADConnectionsPassword** to the corresponding password.

**ADConnectionsUserName** must be in the *domain\user* format.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_AD_CONNECTIONS_PASSWORD |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5014 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- UserIsolationMode Metabase Property

# AddNoHeaders Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AddNoHeaders** property disables adding the Received header. Normally, a server adds a Received header to the top of a message when it receives the message.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37000 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/ SMTPSVC | IIsSmtpService |
| /LM/ SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AdminACL Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AdminACL** property contains a Windows discretionary access control list (DACL) that controls access to any metabase key. This property grants read access, restricted write access, or unrestricted write access. For information about the possible values that can be set for this property, see the Bitmask Values table.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | NTACL |
| WMI Data Type | NTACL |
| ADSI Data Type | NTACL |
| ABO Data Type | NTACL |
| ABO Metabase Identifier | MD_ADMIN_ACL |
| Attributes | INHERIT \| SECURE \| REFERENCE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 6027 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/MSFTPSVC/*n*/ROOT /LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |
| /LM/logging | IIsLogModules |
| /LM/W3SVC/Info | IIsWebInfo |

| | |
|---|---|
| /LM/MSFTPSVC/Info | IIsFtpInfo |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| / | IIS_ROOT |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Values**

| Constant | Value | Description |
|---|---|---|
| MD_ACR_READ | 0x00000001 | Enables read access to all properties. |
| MD_ACR_WRITE | 0x00000002 | Enables write access to all properties. |
| MD_ACR_ENUM_KEYS | 0x00000008 | Enables key enumeration. |
| MD_ACR_RESTRICTED_WRITE | 0x00000020 | See the Remarks section following this table. |
| MD_ACR_UNSECURE_PROPS_READ | 0x00000080 | Enables read access to properties that do not have the METADATA_SECURE attribute set. |
| MD_ACR_WRITE_DAC | 0x00040000 | Enables write access to **AdminACL** for security descriptor creator. |

**Remarks**

MD_ACR_RESTRICTED_WRITE enables restricted write access to the following properties.

| | |
|---|---|
| AdminACL Metabase Property | AppIsolated Metabase Property |
| Path Metabase Property | AccessFlags Metabase Property |
| AnonymousUserName Metabase Property | AnonymousUserPass Metabase Property |
| MaxBandwidth Metabase Property | MaxBandwidthBlocked Metabase Property |
| SecureBindings Metabase Property | ServerBindings Metabase Property |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AdminACLBin Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AdminACLBin** property is used by Microsoft Exchange Server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | binary |
| WMI Data Type | binary |
| ADSI Data Type | binary |
| ABO Data Type | binary |
| ABO Metabase Identifier | MD_VPROP_ADMIN_ACL_RAW_BINARY |
| Attributes | INHERIT \| SECURE \| REFERENCE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 6286 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/MSFTPSVC/*n*/ROOT<br>/LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |

| | |
|---|---|
| /LM/logging | IIsLogModules |
| /LM/W3SVC/Info | IIsWebInfo |
| /LM/MSFTPSVC/Info | IIsFtpInfo |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/W3SVC/AppPools/DefaultAppPool<br><br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AdminEmail Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AdminEmail** property corresponds to the "Administrator e-mail account" option, which is on the **Settings** tab of the Network News Transfer Protocol (NNTP) virtual server Properties.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45060 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AdminName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AdminName** property holds the name of the NNTP administrator.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45061 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AdminServer Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AdminServer** property identifies the Web server instance that is being administered by the remote administration program.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_ADMIN_INSTANCE |
| Attributes | INHERIT |
| Default Value | 1 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2115 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/Info | IIsWebInfo |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AllowAnonymous Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AllowAnonymous** property specifies whether the FTP server will allow anonymous access.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ALLOW_ANONYMOUS |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5005 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AllowClientPosts Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AllowClientPosts** property corresponds to **Allow client posting** on the **Settings** tab of NNTP virtual server Properties.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45062 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AllowControlMsgs Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AllowControlMsgs** property corresponds to **Allow control messages** on the **Settings** tab of NNTP virtual server Properties.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45064 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AllowFeedPosts Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AllowFeedPosts** property corresponds to **Allow feed posting** on the **Settings** tab of NNTP virtual server Properties.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45063 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AllowKeepAlive Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AllowKeepAlive** property specifies whether Keep-Alive processing is permitted.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ALLOW_KEEPALIVES |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 6038 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AllowPathInfoForScriptMappings Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AllowPathInfoForScriptMappings** property indicates whether clients are permitted to specify path information in script-mapped requests.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ALLOW_PATH_INFO_FOR_SCRIPT_MAPPINGS |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2095 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AlwaysUseSsl Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AlwaysUseSsl** property is used by the Microsoft Exchange Server messaging system.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_ SERVER |
| ID | 36924 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AnonymousOnly Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AnonymousOnly** property specifies whether the service or server restricts access to anonymous FTP users. If the value is set to **true**, only anonymous users will be allowed to log on to the FTP server. If set to **false**, users will be allowed to log on using a valid user name, as well as anonymously.

| **Note** |
| --- |
| If the value of **AnonymousOnly** is set to **true**, and the value of the AllowAnonymous Metabase Property is set to **false**, no users will be permitted to log on to the FTP server. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ANONYMOUS_ONLY |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_ SERVER |
| ID | 5006 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AnonymousPasswordSync Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AnonymousPasswordSync** property indicates whether IIS should handle the user password for anonymous users attempting to access resources. The following list details the behavior of this property:

- If **AnonymousPasswordSync** is set to **false**, the administrator must manually set the **AnonymousUserPass** property to the anonymous user password; otherwise anonymous access will not function properly.
- If **AnonymousPasswordSync** is set to **true**, the anonymous user password is set by IIS.
- If **AnonymousPasswordSync** is set to **true** and the value of the AllowAnonymous Metabase Property is set to **false**, no users will be permitted to log on to the FTP server.
- For anonymous password synchronization to work when the domain controller for the IIS server is running Microsoft Windows 2000, you must have subauthentication enabled. For more information, see Configuring Subauthentication in a Windows 2000 Domain.

| **Important** |
| --- |
| Setting **AnonymousPasswordSync** on a server running IIS 6.0 has no effect unless you run your application under the System identity and enable subauthentication. However, it is strongly recommended that you never run an application under the System identity because of the security risk. If your application contains a buffer-overrun, malicious users can do anything they want under the guise of the System identity. Also, subauthentication is not enabled by default on a new installation of IIS 6.0. (See Configuring Subauthentication.) This new default behavior of IIS 6.0 might break old applications that use Anonymous authentication. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ANONYMOUS_USE_SUBAUTH |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_ FILE |
| ID | 6022 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- AnonymousUserName Metabase Property
- AnonymousUserPass Metabase Property

# AnonymousUserName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AnonymousUserName** property specifies the name of the registered local user that is used to authenticate anonymous users. The server associates a user name and password with every server action.

**AnonymousUserName** must be in the *domain\user* format.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_ANONYMOUS_USER_NAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_ FILE |
| ID | 6020 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |

| | |
|---|---|
| /LM/SMTPSVC/*n* | /LM/SMTPSVC/*n* |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- AnonymousUserPass Metabase Property
- AnonymousPasswordSync Metabase Property

# AnonymousUserPass Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AnonymousUserPass** property specifies the password of the registered local user that is used to authenticate anonymous users. The server associates a user name and password with every server action.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_ANONYMOUS_PWD |
| Attributes | INHERIT \| SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6021 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- AnonymousUserName Metabase Property
- AnonymousPasswordSync Metabase Property

# AppAllowClientDebug Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppAllowClientDebug** property specifies whether ASP client-side debugging is enabled. This property is independent of AppAllowDebugging Metabase Property, which applies to server-side debugging.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ENABLECLIENTDEBUG |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7019 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppAllowDebugging Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppAllowDebugging** property specifies whether ASP debugging is enabled on the server. This property is independent of the AppAllowClientDebug Metabase Property, which applies to client-side debugging.

| **Note** |
| --- |
| IIS application threads are serialized when server-side debugging is enabled: Only one thread at a time is allowed to run for each application. This could adversely affect server performance on busy sites. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ENABLESERVERDEBUG |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7018 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppFriendlyName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppFriendlyName** property specifies the user-friendly name of the package or application.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APP_FRIENDLY_NAME (IIS_MD_HTTP_BASE+102) |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 2102 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

AppPackageName Metabase Property

# AppIsolated Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppIsolated** property indicates whether an application is to run in-process, out-of-process, or in a pooled-process. A value of 0 indicates that the application is to run in-process, a value of 1 indicates out-of-process, and a value of 2 indicates a pooled-process. Use the application management methods of the **IIsWebVirtualDir** and **IIsWebDirectory** objects to set the process space for your application.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APP_ISOLATED |
| Attributes | INHERIT |
| Default Value | 2 |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | iis_md_ut_wam |
| ID | 2104 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

InProcessIsapiApps Metabase Property

# ApplicationDependencies Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ApplicationDependencies** property contains the Application Dependencies list for the items in the Web Service Restriction List. If an application is contained in this property, the application has dependencies on an item (or multiple items) in the Web Service Restriction List.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_APP_DEPENDENCIES |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2167 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppOopRecoverLimit Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The AppOopRecoverLimit Metabase Property specifies the maximum number of times an out-of-process application will be restarted after a failure. The server will not service requests from components that have exceeded this limit. This property does not apply to applications or extensions that are run in-process.

| **Note** |
| --- |
| IIS must be restarted before changes to this property take effect. |

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APP_OOP_CRASH_LIMIT |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 2110 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

InProcessIsapiApps Metabase Property

# AppPackageID Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppPackageID** property provides the COM+ application identifier (ID) for a transaction. This ID is used in all transactions managed by Component Services.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APP_PACKAGE_ID |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | iis_md_ut_wam |
| ID | 2106 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

AppPackageName Metabase Property

# AppPackageName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppPackageName** property provides the COM+ application name for a transaction.

Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APP_PACKAGE_NAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 2107 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- AppPackageID Metabase Property
- AppFriendlyName Metabase Property

# AppPoolAutoStart Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppPoolAutoStart** property indicates to the World Wide Web Publishing Service (WWW service) to automatically start an application pool when the application pool is created or when IIS is started, if the value of this property is set to **true**.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_APPPOOL_AUTO_START |
| Attributes | NO_ATTRIBUTES |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9028 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppPoolCommand Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppPoolCommand** property allows an application to send commands to the application pool to start or stop the application pool. Applications can read the value of the AppPoolState Metabase Property to determine if the application pool has responded to the command issued by the use of this property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_COMMAND |
| Attributes | VOLATILE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9026 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |

**Values**

| Constant | Value | Description |
|---|---|---|
| MD_APPPOOL_COMMAND_START | 0x00000001 | Starts the application pool. |
| MD_APPPOOL_COMMAND_STOP | 0x00000002 | Stops the application pool. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppPoolId Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppPoolId** property specifies the application pool where the application is routed.

| **Note** |
| --- |
| If an **AppPoolId** is not specified for a particular application; the application inherits the **AppPoolId** from a parent key that has the **AppPoolId** property set. If none of the application's parents have this property set, the application pool inherits this setting from the site level, or from the World Wide Web Publishing Service (WWW service) level. If no valid **AppPoolId** is available through inheritance, the application pool fails and an error is written in the event log. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APP_APPPOOL_ID |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9101 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/n/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppPoolIdentityType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppPoolIdentityType** property allows application pools to run as LocalSystem, LocalService, NetworkService, or a specific user account. For the possible values to use in setting this property, see the Bitmask Values table below.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_IDENTITY_TYPE |
| Attributes | INHERIT |
| Default Value | 2 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9021 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Bitmask Values**

| Constant | Value | Description |
|---|---|---|
| MD_APPPOOL_IDENTITY_TYPE_LOCALSYSTEM | 0x00000000 | The application pool runs as LocalSystem. |
| MD_APPPOOL_IDENTITY_TYPE_LOCALSERVICE | 0x00000001 | The application pool runs as LocalService. |
| MD_APPPOOL_IDENTITY_TYPE_NETWORKSERVICE | 0x00000002 | The application pool runs as NetworkService. |
| MD_APPPOOL_IDENTITY_TYPE_SPECIFICUSER | 0x00000003 | The application pool runs as a specified user account. |

**Remarks**

MD_ACR_RESTRICTED_WRITE enables write access to the following properties:

| | |
|---|---|
| AdminACL Metabase Property | AppIsolated Metabase Property |

| | |
|---|---|
| Path Metabase Property | AccessFlags Metabase Property |
| AnonymousUserName Metabase Property | AnonymousUserPass Metabase Property |
| MaxBandwidth Metabase Property | MaxBandwidthBlocked Metabase Property |
| SecureBindings Metabase Property | ServerBindings Metabase Property |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppPoolQueueLength Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

The **AppPoolQueueLength** property indicates to HTTP.sys how many requests to queue up for an application pool before rejecting future requests. When the limit for this property is exceeded, IIS rejects the additional requests with a 503 error.

| **Note** |
| --- |
| For Windows Server 2003 with Service Pack 1 (SP1), the default value of this property has changed from 4000 to 1000. If you had IIS 6.0 installed before applying SP1 to your server, the value will not be changed; however, you should consider lowering the value for performance reasons. |

| **Note** |
| --- |
| If you set this property to a value outside the range, you receive a 1041 error in the event log. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_UL_APPPOOL_QUEUE_LENGTH |
| Attributes | INHERIT |
| Default Value | 1000 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9017 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppPoolState Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppPoolState** property value is used by applications to determine the state of an application pool. The possible states are listed below in the Bitmask Values table.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_STATE |
| Attributes | VOLATILE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9027 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |

**Bitmask Values**

| Constant | Value | Description |
| --- | --- | --- |
| MD_APPPOOL_STATE_STARTING | 0x00000001 | The application pool is starting. |
| MD_APPPOOL_STATE_STARTED | 0x00000002 | The application pool has started. |
| MD_APPPOOL_STATE_STOPPING | 0x00000003 | The application pool is stopping. |
| MD_APPPOOL_STATE_STOPPED | 0x00000004 | The application pool has stopped. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppRoot Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppRoot** property contains the metabase path to the application root. For example, the metabase path for the application root of your Default Web Site is **/LM/W3SVC/1/ROOT**. To configure a key as an application root, use the application management methods of the **IIsWebVirtualDir** and **IIsWebDirectory** objects.

| **Important** |
|---|
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APP_ROOT |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | IIS_MD_UT_FILE |
| ID | 2103 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AppWamClsid Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AppWamClsid** property provides the class ID for the application's Web Application Management (WAM) interface.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APP_WAM_CLSID |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 2105 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ArticleTableFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ArticleTableFile** property specifies the location on the file system of the article hash table.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45158 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNNtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ArticleTimeLimit Metabase Property

Updated: July 7, 2006

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ArticleTimeLimit** property corresponds to ArticleTimeLimit registry key described in Microsoft Knowledge Base article, XADM: Registry Setting to Allow the Pulling of Newsfeed Articles Older Than 30 Days at URL.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45056 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspAllowOutOfProcComponents Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In IIS 4.0, the **AspAllowOutOfProcComponents** property specifies whether an ASP script is allowed to call out-of-process components, which are executables that are launched from within an application. This property was made obsolete in IIS 5.0 and its value is ignored. However, scripts using this property will continue to function normally.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ALLOWOUTOFPROCCOMPONENTS |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7014 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspAllowSessionState Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspAllowSessionState** property enables session state persistence for the ASP application. If the value is set to true, the server creates a new Session object for each connection, session state is accessible, session storage is allowed, **Session_OnStart** and **Session_OnEnd** occurs, and the ASPSessionID cookie is sent to the client. If the value is set to false, state access and storage are not allowed, events are not processed, and no cookie is sent.

| **Note** |
| --- |
| The ASP page directive, **@ENABLESESSIONSTATE**, can be used to override this metabase property. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ALLOWSESSIONSTATE |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7011 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspAppServiceFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspAppServiceFlags**property contains the flags that must be set to enable COM+ services on your IIS applications. Once you enable the flags, you may also need to set other properties. For more information, see Configuring Applications to Use COM+ Services. Set this property at the application level.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_SERVICE_FLAGS |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7044 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Flags**

| Flag Name | AspEnableTracker |
|---|---|
| Description | A value of **true** enables COM+ tracker, which allows administrators or developers to debug ASP applications. For more information, see the COM+ Software Development Kit (SDK) and see Configuring Applications to Use COM+ Services. |
| Metabase Bitmask Identifier | MD_ASP_SERVICE_FLAG_TRACKER |

| | |
|---|---|
| Decimal Value | 1 |
| Hexadecimal Value | 0x00000001 |

| **Flag Name** | **AspEnableSxs** |
|---|---|
| Description | A value of **true** enables COM+ side-by-side assemblies, which allow ASP applications to specify which version of a system DLL or classic COM component to use, such as WinHTTP 5.1, Shell Common Controls version 6.0 (Comctl32.dll), GDI Plus version 1.0 (GDIplus.dll), and Visual C++ Run-time Libraries version 6.0. For more information, see the COM+ Software Development Kit (SDK) and see Configuring Applications to Use COM+ Services. |
| Metabase Bitmask Identifier | MD_ASP_SERVICE_FLAG_FUSION |
| Decimal Value | 2 |
| Hexadecimal Value | 0x00000002 |

| **Flag Name** | **AspUsePartition** |
|---|---|
| Description | A value of **true** enables COM+ partitioning, which can be used to isolate Web applications into their own COM+ partitions. COM+ partitions can hold different versions of your own custom COM components. If you set this flag, also set the AspPartitionID Metabase Property. For more information, see Configuring Applications to Use COM+ Services. |
| Metabase Bitmask Identifier | MD_ASP_SERVICE_FLAG_PARTITIONS |
| Decimal Value | 4 |
| Hexadecimal Value | 0x00000004 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspBufferingLimit Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspBufferingLimit** property sets the maximum size of the ASP buffer. If response buffering is turned on, this property controls the maximum number of bytes that an ASP page can write to the response buffer before a flush occurs.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_BUFFER_LIMIT |
| Attributes | INHERIT |
| Default Value | 4194304 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7052 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspBufferingOn Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspBufferingOn** property specifies whether output from an ASP application will be buffered. If the value is set to **true**, all output from the application is collected in the buffer before the buffer is flushed to the client browser. With buffering on, the ASP application has to completely process the ASP script before the client browser receives any output. For this reason, applications that have buffering enabled may seem slower to users than those applications that have buffering disabled, even though the total time taken for the buffered script is less. Therefore, if buffering is turned on, consider using the Flush method of the ASP built-in Response object to pass the user pieces of content as the script is being processed.

If this property is set to **false**, output from an ASP script is written to the client browser as it becomes available. If buffering is turned off, any methods that modify existing HTTP headers or generate new headers must be executed before the content body is sent to the client browser. If buffering is turned on, this restriction is unnecessary.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_BUFFERINGON |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_APP |
| ID | 7000 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspCalcLineNumber Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspCalcLineNumber** property specifies whether ASP should calculate and store the line number of each executed line of code in order to provide the number in an error report.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_CALCLINENUMBER |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_APP |
| ID | 7050 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspCodepage Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspCodepage** property specifies the default code page for an application. This setting can be overridden at the page level by using the ASP page directive, **@CODEPAGE**, for literal strings, or one of the **Codepage** properties of the ASP built-in Session and Response objects for dynamic data. If this property is set to 0 at any node, the system default code page (CP_ACP) is inferred.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_CODEPAGE |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_APP |
| ID | 7016 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspDiskTemplateCacheDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

**AspDiskTemplateCacheDirectory** contains the name of the directory that ASP uses to store compiled ASP templates to disk after overflow of the in-memory cache. If the default directory is changed, then the administrator must ensure that the identity of the processes running Asp.dll has been assigned full control rights to the new directory. Typically, the identity of the processes running Asp.dll is the IWAM_*USER* and System accounts.

**AspDiskTemplateCacheDirectory** can only be set to a local path. ASP template caching fails if this property is set to a UNC or mapped network path.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | EXPANDSZ |
| WMI Data Type | String |
| ADSI Data Type | EXPANDSZ |
| ABO Data Type | EXPANDSZ |
| ABO Metabase Identifier | MD_ASP_DISKTEMPLATECACHEDIRECTORY |
| Attributes | INHERIT |
| Default Value | %windir%\system32\inetsrv\asp compiledtemplates |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7036 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspEnableApplicationRestart Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspEnableApplicationRestart** determines whether an ASP application can be automatically restarted. When changes are made to Global.asa or metabase properties that affect an application, the application will not restart unless the **AspEnableApplicationRestart** property is set to **true**.

| **Important** |
| --- |
| When this property is changed from **false** to **true**, the application immediately restarts. |

IIS also provides an ADSI method, **IIsWebDirectory.AspAppRestart**, which can be used to restart the application. For more information about this method, see the IIS ADSI Provider reference section of the IIS Software Development Kit (SDK)

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ENABLEAPPLICATIONRESTART |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7027 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspEnableAspHtmlFallback Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspEnableAspHtmlFallback** property controls the behavior of ASP when a new request is to be rejected due to a full request queue. A **true** value for this property will cause an .htm file with a similar name as the requested .asp file, if it exists, to be sent instead of the .asp file. The naming convention for the .htm file is the name of the .asp file with _asp appended. For example, if the .asp file is hello.asp, then the .htm file should be hello_asp.htm

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ENABLEASPHTMLFALLBACK |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7021 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspEnableChunkedEncoding Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspEnableChunkedEncoding** property specifies whether HTTP 1.1 chunked transfer encoding is enabled for the World Wide Web Publishing Service (WWW service). The WWW service setting for this property is inheritable by all in-process and pooled out-of-process applications, at all levels below the key where this property is set. Metabase settings at the Web server level or lower are ignored for in-process and pooled out-of-process applications. However, settings at the Web server level or lower will be used if that key is an isolated out-of-process application.

| **Note** |
| --- |
| Even if **AspEnableChunkedEncoding** is set to **true**, chunked transfer encoding is used only by browsers that support and have HTTP 1.1 enabled. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ENABLECHUNKEDENCODING |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7022 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspEnableParentPaths Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspEnableParentPaths** property specifies whether an ASP page allows paths relative to the current directory (using the ..\ notation). If set to **true**, this property constitutes a potential security risk, because an include path could access critical or confidential files outside the root directory of the application.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ENABLEPARENTPATHS |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | SP_MD_UT_APP |
| ID | 7008 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspEnableTypelibCache Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspEnableTypelibCache** property specifies whether type libraries are cached on the server. The World Wide Web Publishing Service (WWW service) setting for this property is applicable to all in-process and pooled out-of-process application nodes, at all levels. Metabase settings at the Web server level or lower are ignored for in-process and pooled out-of-process applications. However, settings at the Web server level or lower are used if that node is an isolated out-of-process application.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ENABLETYPELIBCACHE |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_WAM |
| ID | 7023 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspErrorsToNTLog Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspErrorsToNTLog** property specifies which ASP errors are written to the Windows event Log. ASP errors are written to the client browser and to the IIS log files by default.

**AspErrorsToNTLog** is set to **FALSE** by default, and modifies the **AspLogErrorRequests** property in the following way:

- If **AspLogErrorRequests** is set to **FALSE**, then ASP errors are not written to the Windows event log, regardless of the value of **AspErrorsToNTLog**.
- If **AspLogErrorRequests** is set to **TRUE**, then if IIS fails to write an item to the IIS log file, the item is written to the Windows event log as a warning, regardless of the value of **AspErrorsToNTLog**.
- If **AspLogErrorRequests** is set to **TRUE** and **AspErrorsToNTLog** is set to **FALSE**, then only the most serious ASP errors are sent to the Windows event log. Serious ASP errors are numbers 100, 101, 102, 103, 104, 105, 106, 107, 115, 190, 191, 192, 193, 194, 240, 241, and 242.
- If **AspLogErrorRequests** is set to **TRUE** and **AspErrorsToNTLog** is set to **TRUE**, then all ASP errors are written to the Windows event log.

| **Note** |
| --- |
| This property does not control logging requests to the IIS log files. To accomplish this see the **LogType** and **DontLog** properties. Also, this property does not control IIS Event Messages, which are always sent to the Windows event log. |

This property is a process-level property.

If this property set at the Web service level, it is applicable to all in-process application keys. Metabase property settings at the Web site level or lower are ignored for in-process applications unless the Web site key is isolated as an out-of-process application.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ERRORSTONTLOG |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7024 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspExceptionCatchEnable Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspExceptionCatchEnable** property specifies whether ASP pages trap exceptions thrown by components. If set to **false** (or, **disabled**), the Microsoft Script Debugger tool does not catch exceptions sent by the component that you are debugging.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_EXCEPTIONCATCHENABLE |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7015 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspExecuteInMTA Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP is capable of running all of its threads in a multi-threaded apartment (MTA). If your COM components are primarily free-threaded or both-threaded, running the ASP threads as MTA can improve performance significantly.

By default, the **AspExecuteInMTA** property is set to 0, which means that ASP does not execute in MTA. Set this property to 1 at the application level to enable ASP to run in MTA.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_EXECUTEINMTA |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7041 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspKeepSessionIDSecure Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspKeepSessionIDSecure** property, when enabled, ensures that a SessionID is sent as a secure cookie if assigned over a secure channel.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_KEEPSESSIONIDSECURE |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| StartingNumber | Not applicable |
| EndingNumber | Not applicable |
| ID | 7043 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspLCID Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspLCID** property specifies the default locale identifier (LCID) for an application. LCIDs define how dates, times, and currencies are formatted. LCIDs are not the same for each geographical locale.

Override this setting at the page level by using the ASP page directive, @LCID for literal strings, or one of the **LCID** properties of the ASP built-in Response or Session objects for dynamic data. If this property is set to 0 at any node, the system default locale (LOCALE_SYSTEM_DEFAULT) is inferred. Setting **AspLCID** changes the locale for both the ASP application and the scripting engine. Using the Microsoft Visual Basic Scripting Edition (VBScript) function setLocale() only changes the locale for the scripting engine.

| **Note** |
| --- |
| The locale group has to be installed on the Web server before its locale identifier can be specified in a Web page or application. Use the **Regional and Language Options** item in Control Panel to install locale groups. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_LCID |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7042 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspLogErrorRequests Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspLogErrorRequests** property controls whether the Web server writes ASP errors to the application section of the Windows event log. ASP errors are written to the client browser and to the IIS log files by default.

**AspLogErrorRequests** is set to **TRUE** by default, and is modified by the **AspErrorsToNTLog** property in the following way:

- If **AspLogErrorRequests** is set to **FALSE**, then ASP errors are not written to the Windows event log, regardless of the value of the AspErrorsToNTLog Metabase Property.
- If **AspLogErrorRequests** is set to **TRUE**, then if IIS fails to write an item to the IIS log file, the item is written to the Windows event log as a warning, regardless of the value of the **AspErrorsToNTLog** property.
- If **AspLogErrorRequests** is set to **TRUE** and the **AspErrorsToNTLog** property is set to **FALSE**, then only the most serious ASP errors are sent to the Windows event log. Serious ASP errors are numbers 100, 101, 102, 103, 104, 105, 106, 107, 115, 190, 191, 192, 193, 194, 240, 241, and 242.
- If **AspLogErrorRequests** is set to **TRUE** and the **AspErrorsToNTLog** property is set to **TRUE**, then all ASP errors are written to the Windows event log.

| **Note** |
| --- |
| This property does not control logging requests to the IIS log files. To accomplish this see the **LogType** and **DontLog** properties. Also, this property does not control IIS Event Messages, which are always sent to the Windows event log. |

This property is a process-level property.

| ASP error code | ASP error | Description |
| --- | --- | --- |
| ASP 0100 | Out of memory | Unable to allocate required memory. |
| ASP 0101 | Unexpected error | The function returned (*exception name*). |
| ASP 0102 | Expecting string input | The function expects a string as input. |
| ASP 0103 | Expecting numeric input | The function expects a number as input. |
| ASP 0104 | Operation not Allowed | Operation not Allowed |
| ASP 0105 | Index out of range | An array index is out of range. |
| ASP 0106 | Type Mismatch | An unhandled data type was encountered. |
| ASP 0107 | Stack Overflow | The data being processed is over the allowed limit. |
| ASP 0115 | Unexpected error | A trappable error (*exception name*) occurred in an external object. The script cannot continue running. |
| ASP 0190 | Unexpected error | A trappable error occurred while releasing an external object. |
| ASP 0191 | Unexpected error | A trappable error occurred in the **OnStartPage** method of an external object. |

| | | |
|---|---|---|
| ASP 0192 | Unexpected error | A trappable error occurred in the **OnEndPage** method of an external object. |
| ASP 0193 | OnStartPage Failed | An error occurred in the **OnStartPage** method of an external object. |
| ASP 0194 | OnEndPage Failed | An error occurred in the **OnEndPage** method of an external object. |
| ASP 0240 | Script Engine Exception | A script engine threw exception (*exception name*) in (*object name*) from (*object name*). |
| ASP 0241 | CreateObject Exception | The **CreateObject** of (*object name*) caused exception (*exception name*). |
| ASP 0242 | Query OnStartPage Interface Exception | Querying object's (*object name)* **OnStartPage** or **OnEndPage** methods caused exception (*exception name*). |

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_LOGERRORREQUESTS |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7001 |

**Configurable Locations**
You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**
For general code examples, see Code Examples to Configure Metabase Properties.

# AspMaxDiskTemplateCacheFiles Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The value of the **AspMaxDiskTemplateCacheFiles** property specifies the maximum number of compiled ASP templates that can be stored. The directory that stores the compiled templates is configured by the AspDiskTemplateCacheDirectory Metabase Property property. This is a per process setting, so if ASP is hosted in more than one process, the number of files in this directory could exceed the individual maximum.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | Boolean |
| ADSI Data Type | SINT32 |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_MAXDISKTEMPLATECACHEFILES |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7040 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

AspDiskTemplateCacheDirectory Metabase Property

# AspMaxRequestEntityAllowed Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspMaxRequestEntityAllowed** property specifies the maximum number of bytes allowed in the entity body of an ASP request. If a Content-Length header is present and specifies an amount of data greater than the value of **AspMaxRequestEntityAllowed**, IIS returns a 403 error response. This property is related in function to **MaxRequestEntityAllowed**, but is specific to ASP request. Whereas you might set the **MaxRequestEntityAllowed** property to 1 MB at the general World Wide Web Publishing Service (WWW Service) level, you may choose to set **AspMaxRequestEntityAllowed** to a lower value, if you know that your specific ASP applications handle a smaller amount of data.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_MAX_REQUEST_ENTITY_ALLOWED |
| Attributes | INHERIT |
| Default Value | 200000 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7053 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspPartitionID Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

COM+ partitions are used to isolate Web applications into their own COM+ partitions. COM+ partitions hold different versions of your own custom COM components.

Set the **AspPartitionID** property to the Globally Unique Identifier (GUID) of the COM+ partition. Also, set the **AspUsePartition** flag of the **AspAppServiceFlags** metabase property. Set both properties at the application level.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_ASP_SERVICE_PARTITION_ID |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7048 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspProcessorThreadMax Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspProcessorThreadMax** property specifies the maximum number of worker threads per processor that IIS may create.

| **Note** |
| --- |
| This setting can dramatically influence the scalability of your Web applications, and the performance of your server in general. Because this property defines the maximum number of ASP requests that can execute simultaneously, this setting should remain at the default value unless your ASP applications are making extended calls to external components. In this case, you may increase the value of **AspProcessorThreadMax**. Doing so allows the server to create more threads to handle more concurrent requests. |

If the value of this property is changed, the Web server instance must be stopped and restarted for the change to take effect. The World Wide Web Publishing Service (WWW service) setting for this property is applicable to all in-process and pooled out-of-process application nodes, at all levels. Metabase settings at the Web server level or lower ignores in-process and pooled out-of-process applications. However, settings at the Web server level or lower are used if that node is an isolated out-of-process application.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_PROCESSORTHREADMAX |
| Attributes | INHERIT |
| Default Value | 10 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7025 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

| | |
|---|---|
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspQueueConnectionTestTime Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS places all ASP requests into a queue. If the request is queued longer than the number of seconds specified by the **AspQueueConnectionTestTime** property, ASP checks to determine whether the client is still connected before executing the request. If the client is no longer connected, the request is not processed and is deleted from the queue.

| **Note** |
| --- |
| Most often, users do not wait more than a few seconds for ASP pages to process. Although the maximum waiting time varies from user to user, the generally accepted maximum is approximately 10 seconds. You can use the **AspQueueConnectionTestTime** property to ensure that IIS does not waste time processing a request that has been abandoned by the user. This property is useful for making ASP processing efficient only up to the point at which ASP begins to process the script. Once the script is running, however, your application should continue to check for client connection at appropriate times, by using the **IsClientConnected** method of the ASP built-in **Response** object. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_QUEUECONNECTIONTESTTIME |
| Attributes | INHERIT |
| Default Value | 3 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7028 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspQueueTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspQueueTimeout** property specifies the amount of time (in seconds) that an ASP script request is allowed to wait in the queue. When requests are pulled from the queue, they are checked to see if they have expired (have waited longer than the value of this parameter). Expired requests are rejected with a message indicating the server is too busy.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_QUEUETIMEOUT |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7013 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspRequestQueueMax Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspRequestQueueMax** property specifies the maximum number of concurrent ASP requests that are permitted into the queue. Any client browser that attempts to request ASP files when the queue is full is given an HTTP 500 - Server Too Busy error.

In IIS 6.0 ASPRequestQueueMax can be set at lm/w3svc/AppPools and lm/w3svc/AppPools/<AppPoolName>.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_REQEUSTQUEUEMAX |
| Attributes | INHERIT |
| Default Value | 500 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_MD_UT_APP |
| ID | 7026 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspRunOnEndAnonymously Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspRunOnEndAnonymously** property specifies whether the **SessionOnEnd** and **ApplicationOnEnd** global ASP functions should be run as the anonymous user.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_RUN_ONEND_ANON |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7051 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspScriptEngineCacheMax Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The AspScriptEngineCacheMax Metabase Property specifies the maximum number of scripting engines that ASP pages will keep cached in memory.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_SCRIPTENGINECACHEMAX |
| Attributes | INHERIT |
| Default Value | 120 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7005 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspScriptErrorMessage Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspScriptErrorMessage** property specifies the error message to send to the browser if specific debugging errors are not sent to the client (if AspScriptErrorSentToBrowser Metabase Property is set to **false**).

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_ASP_SCRIPTERRORMESSAGE |
| Attributes | INHERIT |
| Default Value | An error occurred on the server when processing the URL. Please contact the system administrator. |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7003 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- AspScriptErrorSentToBrowser Metabase Property

# AspScriptErrorSentToBrowser Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspScriptErrorSentToBrowser** property specifies whether the Web server writes debugging specifics (file name, error, line number, description) to the client browser, in addition to logging them to the Windows Event Log. The **AspScriptErrorMessage** property provides the error message to be sent if this property is set to false.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_SCRIPTERRORSSENTTOBROWSER |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7002 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

AspScriptErrorMessage Metabase Property

# AspScriptFileCacheSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspScriptFileCacheSize** property specifies the number of precompiled script files to cache. If set to 0, no script files are cached. If set to 4294967295, all script files requested are cached. This property is used to tune performance, depending on the amount of available memory and the amount of script file traffic.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_SCRIPTFILECACHESIZE |
| Attributes | INHERIT |
| Default Value | 500 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7004 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspScriptLanguage Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspScriptLanguage** property specifies the default script language for all ASP applications running on the Web server. The scripting engine for this language will be used for ASP code within the script language delimiters (<% and %>), unless specifically overridden by the following directive which must appear as the first line in any ASP page.

**<%@ LANGUAGE = ScriptingLanguage %>**

| **Note** |
| --- |
| IIS issues an error if *ScriptingLanguage* name is not registered on the system. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_ASP_SCRIPTLANGUAGE |
| Attributes | INHERIT |
| Default Value | VBScript |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7012 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspScriptTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspScriptTimeout** property specifies (in seconds) the default length of time that ASP pages allow a script to run before terminating the script and writing an event to the Windows Event Log. ASP script can override this value by using the **ScriptTimeout** property of the ASP built-in **Session** object. The **ScriptTimeout** property allows your ASP application to set a higher script timeout value. For example, you can use this setting to adjust the timeout once a particular user establishes a valid session by logging in or ordering a product.

| **Note** |
| --- |
| The minimum value for this property is one second. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_SCRIPTTIMEOUT |
| Attributes | INHERIT |
| Default Value | 90 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7006 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

AspSessionTimeout Metabase Property

# AspSessionMax Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspSessionMax** property specifies the maximum number of concurrent sessions that IIS will permit. If a client attempts to establish a new session with IIS once this limit is reached, the client will receive an HTTP 500 - Server Too Busy error.

| **Note** |
| --- |
| ASP requires a certain amount of memory overhead for each session maintained. Although you could use **AspSessionMax** to limit this memory overhead, it is generally more appropriate to manage the lifetime of session objects within IIS by using the AspSessionTimeout Metabase Property property, so that client browsers will be rejected once the limit is reached. |

If an incoming request does not have an ASP session cookie, or has a session cookie that does not match an existing session, the incoming request is considered a new session request. Setting this property at the **IIsWebService** level applies to all in-process and pooled out-of-process application keys, at all levels. Metabase settings at the **IIsWebService** level or lower are ignored for in-process and pooled out-of-process applications. However, settings at the **IIsWebService** level or lower are used if that key is an isolated out-of-process application.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_SESSIONMAX |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7029 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

| | |
|---|---|
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

AspSessionTimeout Metabase Property

# AspSessionTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspSessionTimeout** property specifies (in minutes) the default amount of time that a Session object is maintained after the last request associated with the object is made. You can override this property in script by using the Timeout property of the ASP built-in Session object.

| **Note** |
| --- |
| You can use **AspSessionTimeout** to fine tune your ASP applications. Because Session objects consume memory resources, using this property to limit the lifetime of an individual Session makes your applications more scalable. |

Setting this property at the **IIsWebService** level applies to all in-process and pooled out-of-process application keys, at all levels. Metabase settings at the **IIsWebServer** level or lower are ignored for in-process and pooled out-of-process applications. However, settings at the **IIsWebServer** level or lower are used if that key is an isolated out-of-process application.

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ASP_SESSIONTIMEOUT |
| Attributes | INHERIT |
| Default Value | 10 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7007 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

AspScriptTimeout Metabase Property

# AspSxsName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspSxsName** property enables side-by-side (SxS) assemblies. SxS assemblies allow ASP applications to specify which version of a system DLL or classic COM component to use, such as WinHTTP 5.1, Shell Common Controls version 6.0 (Comctl32.dll), GDI Plus version 1.0 (GDIplus.dll), and Visual C++ Run-time Libraries version 6.0.

Set the **AspSxsName** metabase property to the name of the COM+ application. Also, set the **AspEnableSxs** flag of the AspAppServiceFlags Metabase Property. Set both metabase properties at the application level.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_ASP_SERVICE_SXS_NAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | ASP_MD_UT_APP |
| ID | 7049 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AspTrackThreadingModel Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AspTrackThreadingModel** property specifies whether IIS checks the threading model of any components that your application creates.

| **Important** |
| --- |
| This metabase property is usually best left at the default value (**false**), so that ASP will not use system resources to track the threading model and your ASP application provides the best possible performance. If this property is set to **false**, and you give application scope to components that you create, those components must aggregate the free-threaded marshaler. If you do not aggregate the free-threaded marshaler, ASP generates an error when you try to create an instance of the component. For more information about this topic, see the "Building Components for ASP" section in the Platform SDK. |

Another reason to leave this property at the default setting (**false**), is that any objects that are created in your ASP application without the **OnStartPage** or **OnEndPage** method are released earlier than they would be otherwise and there is degradation of your application's scalability. For more details about improving performance, see Performance Tuning. Also with this property set to **false** ASP is prevented from writing the response before destroying the component.

| **Note** |
| --- |
| In IIS 4.0, the default value for this property was **true**. |

Setting this property at the **IIsWebService** level applies to all in-process and pooled out-of-process application keys, at all levels. Metabase settings at the **IIsWebServer** level or lower are ignored for in-process and pooled out-of-process applications. However, settings at the **IIsWebServer** level or lower are used if that key is an isolated out-of-process application.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ASP_ID_LAST |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 7020 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |

| | |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AuthChangeURL Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AuthChangeURL** property contains the URL to be invoked when a user enters a new password.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_AUTH_CHANGE_URL |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 2060 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AuthExpiredURL Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AuthExpiredURL** property contains a URL that will be invoked when a user's password has expired. It will be invoked in a secure (HTTPS) manner.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_AUTH_EXPIRED_URL |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 2061 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AuthExpiredUnsecureURL Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AuthExpiredUnsecureURL** property contains a URL that will be invoked when a user's password has expired.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_AUTH_EXPIRED_UNSECUREURL |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 2067 |

**Configurable Locations**
You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**
For general code examples, see Code Examples to Configure Metabase Properties.

# AuthFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AuthFlags** property contains the settings for the Windows authentication schemes that are returned to clients as being available. This property contains the file access authentication flags, which are listed below in the Flags section.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_AUTHORIZATION |
| Attributes | INHERIT |
| Default Value | authanonymous |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6000 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/file_name /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Flags**

| Flag Name | AuthAnonymous |
|---|---|
| Description | Specifies Anonymous authentication as one of the possible Windows authenticationschemes returned to clients as being available. |
| Metabase Bitmask Identifier | MD_AUTH_ANONYMOUS |
| Decimal Value | 1 |
| Hexadecimal Value | 0x00000001 |

| Flag Name | AuthBasic |
|---|---|
| Description | Specifies Basic authentication as one of the possible Windows authenticationschemes returned to clients as being available. |
| Metabase Bitmask Identifier | MD_AUTH_BASIC |
| Decimal Value | 2 |
| Hexadecimal Value | 0x00000002 |

| Flag Name | AuthMD5 |
|---|---|
| Description | Specifies Digest authentication and Advanced Digest authentication as one of the possible Windows authentication schemes returned to clients as being available. |
| Metabase Bitmask Identifier | MD_AUTH_MD5 |
| Decimal Value | 16 |
| Hexadecimal Value | 0x00000010 |

| Flag Name | AuthNTLM |
|---|---|
| Description | Specifies Integrated Windows authentication (also known as Challenge/Response or NTLMauthentication) as one of the possible Windows authentication schemes returned to clients as being available. |
| Metabase Bitmask Identifier | MD_AUTH_NT |
| Decimal Value | 4 |
| Hexadecimal Value | 0x00000004 |

| Flag Name | AuthPassport |
|---|---|
| Description | A value of **true** indicates that Microsoft .NET Passport authentication is enabled. For more information, see .NET Passport Authentication in IIS 6.0. |
| Metabase Bitmask Identifier | MD_AUTH_PASSPORT |
| Decimal Value | 64 |
| Hexadecimal Value | 0x00000040 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AuthNotifyPwdExpUnsecureURL Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AuthNotifyPwdExpUnsecureURL** property contains a URL that will be invoked if the user's password expires sooner than the number of days specified in **PasswordExpirePreNotifyDays**.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_AUTH_NOTIFY_PWD_EXP_UNSECUREURL |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 2069 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AuthNotifyPwdExpURL Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AuthNotifyPwdExpURL** property contains a URL that will be invoked if the user's password expires sooner than the number of days specified in **PasswordExpirePreNotifyDays**. It will be invoked in a secure (HTTPS) manner.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_AUTH_NOTIFY_PWD_EXP_URL |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 2062 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AuthPersistence Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AuthPersistence** property specifies authentication persistence across requests on a connection using NTLM authentication. The default for NTLM is **false (0)**, for legacy reasons, while all other authentication schemes assume the single request to be **true (1)**. When the **AuthPersistSingleRequest** flag is set to **true** on this property, IIS 6.0 automatically reauthenticates every request, even those on the same connection. This behavior is different on IIS 5.0 and IIS 5.1, where there were two other flags for the **AuthPersistence** property specifying persistence based on the existence of a proxy server.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_AUTHORIZATION_PERSISTENCE |
| Attributes | INHERIT |
| Default Value | authpersistsinglerequest |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6031 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Flags**

| **Flag Name** | **AuthPersistSingleRequest** |
|---|---|
| Description | Setting this flag to **true** specifies that authentication persists only for a single request on a connection. IIS resets the authentication at the end of each request, and forces re-authentication on the next request of the session. |

| | |
|---|---|
| Metabase Bitmask Identifier | MD_AUTH_SINGLEREQUEST |
| Decimal Value | 64 |
| Hexadecimal Value | 0x00000040 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AuthTurnList Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AuthTurnList** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36966 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/Domain | IIsSmtpDomain |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AutoShutdownAppPoolExe Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AutoShutdownAppPoolExe** property specifies an executable to run when the World Wide Web Publishing Service (WWW service) shuts down an application pool for rapid fail protection. You can use the **AutoShutdownAppPoolParams** property to send parameters to the executable.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_AUTO_SHUTDOWN_APPPOOL_COMMAND |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9035 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AutoShutdownAppPoolParams Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AutoShutdownAppPoolParams** property specifies any command-line parameters for the executable that is specified in the AutoShutdownAppPoolExe Metabase Property. You can use these two properties in the following way to send e-mail, for example, when the World Wide Web Publishing Service (WWW service) shuts down an application pool for rapid fail protection:

AutoShutdownAppPoolExe = "C:\LogAndSendMail.bat"

AutoShutdownAppPoolParams = "-AppPoolName %1%"

where %1% represents the application pool name.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APPPOOL_STATE_STOPPED |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9036 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AzEnable Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AzEnable** property enables URL Authorization for the virtual directory, application, or URL corresponding to the entry in the metabase.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_ENABLE_URL_AUTHORIZATION |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6048 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/n/ROOT/*file_name* /LM/W3SVC/n/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AzImpersonationLevel Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AzImpersonationLevel** property determines the impersonation behavior for the application, allowing the Web application to be configured to impersonate the client user, the IIS worker process, or the IUSER_* account for the worker process.

| **Note** |
| --- |
| Take care in setting this value, as each setting significantly changes the environment and implied design of the Web application. |

| Impersonation Level | Numeric value of the Field |
| --- | --- |
| Client User | 0 |
| IIS Worker Process | 1 |
| IUSER_* | 2 |

| Attribute Name | Attribute Value |
| --- | --- |
| Attribute Name | Attribute Value |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_URL_AUTHORIZATION_IMPERSONATION_LEVEL |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6053 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |

| | |
|---|---|
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AzScopeName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AzScopeName** property associates the virtual directory, application, or URL with a scope. If no scope or an empty string is specified, the default scope of the IIS 6.0 URL Authorization will be used.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_URL_AUTHORIZATION_SCOPE_NAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6050 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# AzStoreName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **AzStoreName** property associates an Authorization Manager policy store with the virtual directory, application, or URL. This string value must specify the store type and path to the policy store (i.e. **"msxml://c:\mystore.xml"** or **"msldap://CN=MyAzStore, CN=Program Data, DC=mydomain, DC=com"**).

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_URL_AUTHORIZATION_STORE_NAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6049 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# BadMailDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **BadMailDirectory** property contains the directory where undeliverable mail messages are placed. This corresponds to "Badmail directory" on the **Message** tab.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36909 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# BannerMessage Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **BannerMessage** property specifies the message text an FTP server transmits to a client once the client connects to the FTP server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_BANNER_MESSAGE |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5011 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- ExitMessage Metabase Property
- GreetingMessage Metabase Property
- MaxClientsMessage Metabase Property

# BinSchemaTimeStamp Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **BINSchemaTimeStamp** property is used internally by IIS to keep the MBSchema.xml file in sync with the binary schema file.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | binary |
| WMI Data Type | binary |
| ADSI Data Type | binary |
| ABO Data Type | binary |
| ABO Metabase Identifier | MD_GLOBAL_BINSCHEMATIMESTAMP |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 9991 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| . | IIS_Global |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CacheControlCustom Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CacheControlCustom** property specifies custom HTTP 1.1 cache control directives. Multiple directives can be specified if they are separated by carriage return/line feed (CRLF) pairs. There should not be a trailing carriage return/line feed after the final directive.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_CC_OTHER |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6043 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CacheControlMaxAge Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CacheControlMaxAge** property specifies the HTTP 1.1 cache control maximum age value. The value of this property ranges from 0 to 4294967295 (or unlimited), measured in seconds.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_CC_MAX_AGE |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6042 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CacheControlNoCache Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Reference topic for IIS CacheControlNoCache metabase properties.

The **CacheControlNoCache** property specifies the HTTP 1.1 directive to prevent caching of content.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_CC_NO_CACHE |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6041 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CacheISAPI Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CacheISAPI** property indicates whether ISAPI extensions are cached in memory after first use. If the value of this property is set to **true**, ISAPI extensions, once loaded, remain in the cache until the server is stopped. If the value is set to **false**, extensions are unloaded from memory after the ISAPI extension is no longer in use.

ISAPI extensions are cached or not cached based on the value of this property at the time they were loaded into memory for use. Thus, if this property is changed after an extension has been loaded and cached, the change will have no effect on that extension until the Web server is restarted.

| **Important** |
|---|
| There are severe performance implications to setting the value of this property to **false**. Also, some extensions cannot work properly when this property is set to **false**. |

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_CACHE_EXTENSIONS |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6034 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CentralBinaryLoggingEnabled Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CentralBinaryLoggingEnabled** property is a global flag to denote that IIS will generate one centralized binary log file rather than a separate file per Web site.

| **Note** |
| --- |
| This parameter is not inheritable and only exists on the parent node. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_GLOBAL_BINARY_LOGGING_ENABLED |
| Attributes | NO_ATTRIBUTES |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4016 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CentralW3CLoggingEnabled Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CentralW3CLoggingEnabled** metabase property configures W3C centralized logging on the server, a global configuration where all Web sites write data to a single log file. This single log file uses the W3C Extended log file format and can be viewed in a text editor. In contrast, centralized binary logging writes data in a binary format and requires a parsing tool to view the data. W3C centralized logging is available in Windows Server 2003 Service Pack 1 or later.

| **Note** |
| --- |
| This parameter is not inheritable and only exists on the parent node. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | XX |
| Attributes | NO_ATTRIBUTES |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2119 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**See Also**

**Tasks**

Enabling W3C Centralized Logging

# CertCheckMode Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The CertCheckMode Metabase Property enables or disables Certificate Revocation List (CRL) checking. **CertCheckMode** is set to 0 by default (**CertCheckMode**=0), which means that IIS by default searches for an updated CRL.

This metabase property is closely related to the RevocationFreshnessTime Metabase Property and the RevocationURLRetrievalTimeout Metabase Property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_CERT_CHECK_MODE |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2160 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Values**

| Constant | Value | Description |
|---|---|---|
| MD_CERT_NO_REVOC_CHECK | 0x00000001 | Certificate revocation checking is not performed. |
| MD_CERT_CACHE_RETRIEVAL_ONLY | 0x00000002 | During certificate revocation verification, the CRL will not be updated from a remote location, such as a CRL at an external URL. In this case, the CRL that is cached on the client is used. If the CRL is expired, the certificate revocation verification fails. |

| | | |
|---|---|---|
| MD_CERT_CHECK_REVOCATION_FRESHNESS_TIME | 0x00000004 | The client CRL is replaced by the CRL at a remote location, even if the CRL that is cached on the client is valid. The value of the RevocationFreshnessTime Metabase Property determines the frequency of this action. |
| MD_CERT_NO_USAGE_CHECK | 0x00010000 | The certificate provided by the client is not verified as valid. |

**Important**

The MD_CERT_NO_REVOC_CHECK, MD_CERT_CACHE_RETRIEVAL_ONLY, and MD_CERT_CHECK_REVOCATION_FRESHNESS_TIME flags are mutually exclusive.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Topics**

- For more information about using SSL and certificates with IIS 6.0, see Encryption in IIS 6.0.
- For more information about certificate revocation and CRLs, see Revoking Certificates and Publishing CRLs in Windows Server 2003 documentation.

# CgiRestrictionList Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is controlled by IIS.

| Attribute Name | Attribute Value |
|---|---|
| ADSI/WMI Data Type | MULTISZ |
| ABO Data Type | MULTISZ_METADATA |
| Schema Default | **null** |
| Internal Default | 0 |
| Upper Bound | Not specified |
| Lower Bound | Not specified |
| Internal ID | 2164 |
| Friendly ID | MD_CGI_RESTRICTION_LIST |
| Property Attributes | NO_ATTRIBUTES |
| User Type | IIS_MD_UT_SERVER |

**Configurable Locations**
This property is not configurable at any location in the metabase.

**Flags**
There are no flags for this property.

# CGITimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CGITimeout** property specifies the timeout (in seconds) for CGI applications. The value range is 10 seconds to &H80000000 (2.1 million seconds).

| **Note** |
| --- |
| &H represents Microsoft Visual Basic Scripting Edition (VBScript) hexadecimal format. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SCRIPT_TIMEOUT |
| Attributes | INHERIT |
| Default Value | 300 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6033 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ChangeNumber Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ChangeNumber** property is used internally by IIS to keep track of changes to the metabase.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_GLOBAL_CHANGE_NUMBER |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 9997 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| . | IIS_Global |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ClientPostHardLimit Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ClientPostHardLimit** property contains the number of bytes the server will take before breaking the session.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45165 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ClientPostSoftLimit Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ClientPostSoftLimit** property contains the largest post the server will take.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45166 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ClusterEnabled Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ClusterEnabled** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_CLUSTER_ENABLED |
| Attributes | NO_ATTRIBUTES |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1025 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CollectionComment Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CollectionComment** property exposes comments in the metabase.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_COMMENTS |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9990 |

**Configurable Locations**
This property cannot be configured.

**Code Example**
For general code examples, see Code Examples to Configure Metabase Properties.

# ConnectResponse Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ConnectResponse** property changes the banner response issued at connection to Microsoft Exchange. You must restart the service for this to take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36907 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ConnectionTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ConnectionTimeout** property assumes the functional role of the IIS 5.0 **ServerListenTimeout** property.

In IIS 6.0, use this property in conjunction with the **HeaderWaitTimeout** property, and for Web service nodes, use the **MinFileBytesPerSec** property instead of the **ServerListenTimeout** property.

The **ConnectionTimeout** property specifies the time in seconds that IIS waits before it disconnects a connection considered inactive for any of the following reasons:

- Timer_ConnectionIdle

  The connection was closed because there was not activity and the Timer_ConnectionIdle timer expired.

- Timer_EntityBody

  The connection was closed because no data was received, and the Timer_EntityBody timer expired.

  If a request has an entity body, HTTP.sys begins a Timer_EntityBody timer, and sets the **ConnectionTimeout** value to match the timer. When HTTP.sys receives data on the request, it resets the Timer_EntityBody timer and **ConnectionTimeout**.

- Timer_AppPool

  The connection was closed because the request waited longer than the **ConnectionTimeout** value while sitting in an Application Pool queue, waiting to be processed. Increasing the **ConnectionTimeout** value increases the time a request can wait in an application pool queue before being disconnected.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_CONNECTION_TIMEOUT |
| Attributes | INHERIT |
| Default Value | 120 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1013 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ContentIndexed Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ContentIndexed** property specifies whether the installed content indexer should index content under this directory tree.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_IS_CONTENT_INDEXED |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6039 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CPUAction Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CPUAction** property configures the action(s) that IIS takes when Microsoft Windows NT job objects run. Only one Windows NT job object exists per application pool, therefore the **CPUAction** property is configured on a per application pool basis.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_CPU_ACTION |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9022 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Values**

| **Important** |
|---|
| Only one value from the following table should be used per application pool. |

| Constant | Value | Description |
|---|---|---|
| MD_CPU_NO_ACTION | 0x00000000 | No action is taken except that a warning is written to the event log when the CPU limit is exceeded. |
| MD_CPU_KILL_W3WP | 0x00000001 | Application pool worker processes that exceed their CPU limit will be forced to shut down. |
| MD_CPU_TRACE | 0x00000002 | Enables capacity tracing for the application pool when 80%of the CPU limit is reached. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CPULimit Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CPULimit** property configures the maximum percentage of CPU time (in 1/1000ths of a percent) that the worker processes in an application pool are allowed to consume over a period of time as indicated by the CPUResetInterval Metabase Property property. If the limit set by the **CPULimit** property is exceeded, an event is written to the event log and an optional set of events can be triggered as determined by the CPUAction Metabase Property property.

| **Note** |
| --- |
| Setting the value of this property to 0 disables limiting the worker processes to a percentage of CPU time. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_CPU_LIMIT |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9023 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CPUResetInterval Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CPUResetInterval** property specifies the reset period (in minutes) for CPU monitoring and throttling limits on the application pool. When the number of minutes elapsed since the last process accounting reset equals the number specified by this property, IIS will reset the CPU timers for both the logging and limit intervals. Setting the value of this property to 0 disables CPU monitoring.

| **Important** |
| --- |
| The value in **CpuResetInterval** must be greater than the time between logging operations, otherwise IIS will reset counters before logging has occurred, and process accounting will not occur. |

| **Note** |
| --- |
| Because process accounting in IIS uses Windows Job Objects to monitor CPU times for an entire process, process accounting will log and throttle only applications that are isolated in a separate process from IIS. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_CPU_RESET_INTERVAL |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2144 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CreateCGIWithNewConsole Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CreateCGIWithNewConsole** property indicates whether a CGI application runs in its own console. If the value is set to **true**, each CGI application creates a new console when started. A value of **false** indicates that CGI applications should run without a console.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_CREATE_PROC_new_CONSOLE |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6036 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CreateProcessAsUser Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CreateProcessAsUser** property specifies whether a CGI process is created as the application pool identity or the identity of the requesting user.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_CREATE_PROCESS_AS_USER |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6035 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CSideEtrnDomains Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **CSideEtrnDomains** property lists the domains that will be sent when this Simple Mail Transfer Protocol (SMTP) virtual server sends the **ETRN** command, acting as a client.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36968 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC/*n*/Domain | IIsSmtpDomain |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# CustomErrorDescriptions Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The value of the **CustomErrorDescriptions** property is a list of custom error messages. Each string in the list specifies an error code and subcode for the custom message, what the specific error text and subcode error text is, and whether the custom message is supported by file or URL, or by file only. Each string value is in the format *ErrorCode*, *ErrorSubcode*, *ErrorText*, *ErrorSubcodeText*, followed by another comma and either a 0 for file and URL support, or a 1 for file support only.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_CUSTOM_ERROR_DESC |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2120 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Info | IIsWebInfo |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HttpErrors Metabase Property

# DefaultDoc Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DefaultDoc** contains one or more file names of default documents that will be returned to the client if no file name is included in the client's request. The default document will be returned if the **EnableDefaultDoc** flag of the DirBrowseFlags Metabase Property is set to **true** for the directory. This property can contain a list of default document file names separated by a comma and a space, for example Default.htm, Default.asp.

| **Important** |
| --- |
| If the **EnableDefaultDoc** flag is not set in the DirBrowseFlags Metabase Property property, the server will not search for a default document. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_DEFAULT_LOAD_FILE |
| Attributes | INHERIT |
| Default Value | default.htm |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6006 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n*/ROOT<br><br>/LM/W3SVC/*n*/ROOT/virtual_directory_name | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/physical_directory_name<br><br>/LM/W3SVC/*n*/virtual_directory_name/physical_directory_name | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- DirBrowseFlags Metabase Property

# DefaultDocFooter Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DefaultDoc** contains one or more file names of default documents that will be returned to the client if no file name is included in the client's request. The default document will be returned if the **EnableDefaultDoc** flag of the DirBrowseFlags Metabase Property is set to **true** for the directory. This property can contain a list of default document file names separated by a comma and a space, for example Default.htm, Default.asp.

| **Important** |
|---|
| If the **EnableDefaultDoc** flag is not set in the **DirBrowseFlags** property, the server will not search for a default document. |

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_DEFAULT_LOAD_FILE |
| Attributes | INHERIT |
| Default Value | default.htm |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6006 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- DirBrowseFlags Metabase Property

# DefaultDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DefaultDomain** property specifies the default domain for logon to an SMTP server. If this value is not specified, the default domain is the domain name controlled by the computer, if the computer is a domain controller. If the computer is not a domain controller, the default domain is the computer name.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_DEFAULT_LOGON_DOMAIN |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36908 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DefaultLogonDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In Web, FTP, and NNTP servers, the **DefaultLogonDomain** property is used to specify the default domain that the server uses to authenticate users (in a Web hosting scenario where **UserIsolationMode** = 2). If the value of **DefaultLogonDomain** is not specified, the default domain will be the domain name controlled by the computer, if the computer is a domain controller. If the computer is not a domain controller, the default domain will be the computer name.

In the case of FTP, IIS uses this domain name to query the corresponding Active Directory for the two properties required for each user in the Web hosting scenario: msIIS-FTPRoot and msIIS-FTPDir.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_DEFAULT_LOGON_DOMAIN |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6012 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |

| | |
|---|---|
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br><br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DefaultModeratorDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DefaultModeratorDomain** property corresponds to "Default moderator domain" on the **Settings** tab of NNTP virtual server Properties.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45065 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DemandStartThreshold Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DemandStartThreshold** property specifies the maximum number of worker processes allowed to run concurrently on a machine.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_DEMAND_START_THRESHOLD |
| Attributes | NO_ATTRIBUTES |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9207 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/ | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DirBrowseFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DirBrowseFlags** metabase property contains flags, as seen in the Flags section, that control whether directory browsing is enabled, how much directory and file information is provided if browsing is enabled, and whether there is a default page in the directory.

| **Note** |
| --- |
| If a client accesses the directory and does not provide a file name, and the **EnableDefaultDoc** flag is set to **true**, the server will look for a default file specified in the DefaultDoc Metabase Property. If a default file does not exist in the directory, and the **EnableDirBrowsing** flag is set to **true**, then directory browsing is enabled. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_DIRECTORY_BROWSING |
| Attributes | INHERIT |
| Default Value | dirbrowseshowdate \| dirbrowseshowtime \|dirbrowseshowsize \| dirbrowseshowextension \| enabledefaultdoc |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6005 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT<br><br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br><br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Flags**

| Flag Name | EnableDirBrowsing |
|---|---|
| Description | When set to **true**, directory browsing is enabled. |
| Metabase Bitmask Identifier | MD_DIRBROW_ENABLED |
| Decimal Value | 42147483648 |
| Hexadecimal Value | 0x80000000 |

| Flag Name | DirBrowseShowDate |
|---|---|
| Description | When set to **true**, date information is displayed when browsing directories. |
| Metabase Bitmask Identifier | MD_DIRBROW_SHOW_DATE |
| Decimal Value | 2 |
| Hexadecimal Value | 0x00000002 |

| Flag Name | DirBrowseShowTime |
|---|---|
| Description | When set to **true**, file time information is displayed when displaying directories. |
| Metabase Bitmask Identifier | MD_DIRBROW_SHOW_TIME |
| Decimal Value | 4 |
| Hexadecimal Value | 0x00000004 |

| Flag Name | DirBrowseShowSize |
|---|---|
| Description | When set to **true**, file size information is displayed when browsing directories. |
| Metabase Bitmask Identifier | MD_DIRBROW_SHOW_SIZE |
| Decimal Value | 8 |
| Hexadecimal Value | 0x00000008 |

| Flag Name | DirBrowseShowExtension |
|---|---|
| Description | When set to **true**, file name extensions are displayed when browsing directories. |
| Metabase Bitmask Identifier | MD_DIRBROW_SHOW_EXTENSION |
| Decimal Value | 16 |
| Hexadecimal Value | 0x00000010 |

| Flag Name | DirBrowseShowLongDate |
|---|---|
| Description | When set to **true**, date information is displayed in extended format when displaying directories. |
| Metabase Bitmask Identifier | MD_DIRBROW_LONG_DATE |
| Decimal Value | 32 |
| Hexadecimal Value | 0x00000020 |

| Flag Name | EnableDefaultDoc |
|---|---|
| Description | When set to **true**, the default document (specified by the DefaultDoc Metabase Property) for a directory is loaded when the directory is browsed. |

| | |
|---|---|
| Metabase Bitmask Identifier | MD_DIRBROW_LOADDEFAULT |
| Decimal Value | 1073741824 |
| Hexadecimal Value | 0x40000000 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DirectoryLevelsToScan Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DirectoryLevelsToScan** property specifies the number of virtual directory levels to scan in the metabase from the root virtual directory for a server in order to build a list of directories to be monitored for file system changes. If a virtual directory is not included in the list, then the IIS cache will not be flushed when a file change occurs in the directory. This list is also used for down-level remote administration from earlier versions of IIS. The value of this property ranges from 0 to 4294967295 (or unlimited).

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, Microsoft Visual Basic Scripting Edition (VBScript) represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LEVELS_TO_SCAN |
| Attributes | NO_ATTRIBUTES |
| Default Value | 2 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1022 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DisableNewNews Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DisableNewNews** property is set to disable the NEWNEWS command. Administrators may want to disable NEWNEWS for performance reasons so that other servers are not pulling entire feeds.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45068 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DisableSocketPooling Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Because the **DisableSocketPooling** property is defined as a valid property in the IIS 6.0 metabase schema (MBSchema.xml), you can still set this property, but the setting has no effect. In IIS 6.0, the HTTP listener is in the new kernel-level driver, HTTP.sys. To configure HTTP.sys, you must modify an IP inclusion list by using Httpcfg.exe. For more information, see Microsoft Knowledge Base article Q813368, IIS 6.0: Setting Metabase Property DisableSocketPooling Has No Effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_DISABLE_SOCKET_POOLING |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1029 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DisableStaticFileCache Metabase Property

Updated: August 19, 2013

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DisableStaticFileCache** property specifies the caching of static files for a directory and is disabled on a per-request basis. If **DisableStaticFileCache** = 1, static files are not cached for the specified directory. Note: Active Server Pages (ASP) files and templates are cached independently of this setting. This property may be set on any valid virtual directory and applies to the virtual directory root and any subdirectories.

The **DisableStaticFileCache** property is used when multiple virtual directories are created on Universal Naming Convention (UNC) shares and are secured with individual identities and access control lists (ACLs) to fully secure files from access by unauthorized users. For this configuration scenario, not setting this property would potentially allow unauthorized access to a site's files in the cache through other security identities.

The **DisableStaticFileCache** property may also be set to permit case-sensitive requests for files that are stored in virtual directories located on network file system (NFS) volumes and accessed through an NFS client, such as one provided by Services For Unix (SFU).

The static file cache may alternatively be disabled at the machine level through the registry key HKLM\System\CurrentControlSet\Services\inetinfo\parameters\DisableMemoryCache.

Disabling the static file cache will, in most cases, reduce the Web server request handling performance. However, the effect may be negligible if requests are broadly and randomly distributed across a large number of files.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_VR_NO_CACHE |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 3007 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DisallowOverlappingRotation Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DisallowOverlappingRotation** property specifies whether or not the World Wide Web Publishing Service (WWW Service) should start up another worker process to replace the existing worker process while it is shutting down. The value of this property should be set to **true** if the worker process loads any application code that does not support multi-instance.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_APPPOOL_DISALLOW_OVERLAPPING_ROTATION |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9015 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br><br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DisallowRotationOnConfigChange Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DisallowRotationOnConfigChange** property specifies whether or not the World Wide Web Publishing Service (WWW Service) should rotate worker processes in an application pool when the configuration has changed. This means that the worker processes will not pick up application pool changes to values passed to the worker process, such as IdleTimeout Metabase Property and PeriodicRestartRequests Metabase Property. The value of **DisallowRotationOnConfigChange** should be set to **true** if the worker process loads any application code that does not support multi-instance.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_APPPOOL_DISALLOW_ROTATION_ON_CONFIG_CHANGE |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9018 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br><br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DoDynamicCompression Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DoDynamicCompression** property performs the same function as the global **HcDoDynamicCompression** property, but it is used for enabling and disabling static content compression at the individual site or site element level.

Before you can use **DoDynamicCompression** to customize compression at the site or site element level, at least one compression scheme (gzip or deflate) must be enabled in the **HcDoDynamicCompression** property at the **IIsCompressionScheme** level. This setting determines which compression schemes IIS will support for dynamic content: gzip, deflate, or both.

You can either enable compression of dynamic content server-wide and then disable it at specific locations, or you can disable compression of dynamic content server-wide and then enable it only at the locations you specify. For more information about these two approaches, see the HcDoDynamicCompression Metabase Property.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_DO_NAMESPACE_DYNAMIC_COMPRESSION |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 2255 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT<br><br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br><br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |

| | |
|---|---|
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br><br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcDoDynamicCompression Metabase Property

# DomainRouting Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DomainRouting** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36920 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DoMasquerade Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DoMasquerade** property toggles domain masquerading.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36913 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DontLog Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DontLog** property specifies whether client requests are written to the IIS log files. By default, logging is enabled by the **LogType** property on a service or server node. Set the **DontLog** property to **TRUE** on child directories or files to exempt them from logging.

ASP and ODBC errors do not get logged to the IIS log files.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_DONT_LOG |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6023 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

| | |
|---|---|
| /LM/MSFTPSVC/*n*/ROOT<br><br>/LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/NNTPSVC/*n*/ROOT<br><br>/LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/ROOT<br><br>/LM/SMTPSVC/*n*/ROOT/*virtual_directory_name* | IIsSmtpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- LogType Metabase Property

# DoStaticCompression Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DoStaticCompression** property performs the same function as the global **HcDoStaticCompression Metabase Property** property, but is used for enabling or disabling static content compression at the individual site or site element level.

Before you can use **DoStaticCompression** to customize compression at the site or site element level, at least one compression scheme (gzip or deflate) must be enabled in the **HcDoStaticCompression** property at the **IIsCompressionScheme** level. This setting determines which compression schemes IIS will support for static content: gzip, deflate, or both.

You can either enable compression of static content server-wide and then disable it at specific locations, or you can disable compression of static content server-wide and then enable it only at the locations you specify. For more information about these two approaches, see HcDoStaticCompression Metabase Property.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_DO_NAMESPACE_STATIC_COMPRESSION |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 2256 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT<br><br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br><br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |

| | |
|---|---|
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br><br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcDoStaticCompression Metabase Property

# DownlevelAdminInstance Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DownlevelAdminInstance** indicates the server instance for remote administration clients.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_DOWNLEVEL_ADMIN_INSTANCE |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1021 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# DropDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **DropDirectory** property lists the location of the Drop directory.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36882 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Enable32BitAppOnWin64 Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **Enable32BitAppOnWin64** property indicates whether IIS creates 32-bit worker processes or 64-bit worker processes on 64-bit Windows. This property is only relevant on 64-bit Windows with IIS 6.0 running in Worker Process Isolation mode.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_APPPOOL_32_BIT_APP_ON_WIN64 |
| Attributes | INHERIT |
| Default Value | **FALSE** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9040 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

# EnableDocFooter Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **EnableDocFooter** property enables or disables custom footers specified by the DefaultDocFooter Metabase Property. Setting **EnableDocFooter** to **false** disables custom footers, and results in higher performance.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_FOOTER_ENABLED |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6010 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- DefaultDocFooter Metabase Property

# EnableEditWhileRunning Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **EnableEditWhileRunning** property is used to enable the edit-while-running feature. The **EnableHistory** property must also be set to **true** for the edit-while-running feature to be enabled. For more information about the edit-while-running feature, see Enabling Edit-While-Running in IIS 6.0.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ROOT_ENABLE_EDIT_WHILE_RUNNING |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9998 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM | IIsComputer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- EnableHistory Metabase Property

# EnableHistory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **EnableHistory** property enables or disables the metabase history feature. Setting this property to **true** enables metabase history. Setting the MaxHistoryFiles property to 0 (zero) effectively disables the metabase history feature. For more information, see Configuring the Metabase History Feature in IIS 6.0.

| **Important** |
| --- |
| If this property is set to **false** when EnableEditWhileRunning Metabase Property is set to **true**, IIS will write a warning to the event log and reset this property to **true**. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ROOT_ENABLE_HISTORY |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9996 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM | IIsComputer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- MaxHistoryFiles Metabase Property
- EnableEditWhileRunning Metabase Property

# EnableReverseDns Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **EnableReverseDns** property enables or disables reverse Domain Name Server (DNS) lookups for the World Wide Web Publishing Service (WWW service). Reverse lookups involve looking up the domain name when the IP address is known. Reverse DNS lookups can use significant resources and time.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_DO_REVERSE_DNS |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6029 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# EnableReverseDnsLookup Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **EnableReverseDnsLookup** property corresponds to the **Perform reverse DNS lookup on incoming messages** check box in the **Advanced Delivery** dialog box.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_REVERSE_NAME_LOOKUP |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36884 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# EtrnDays Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **EtrnDays** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ETRN_DAYS |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36918 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# EtrnSubdomains Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **EtrnSubdomains** property indicates whether Etrn subdomains (the @ extension to Etrn) are allowed.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36938 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ExitMessage Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ExitMessage** property specifies the message text that an FTP server transmits to a client when the client sends a quit command.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_EXIT_MESSAGE |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5001 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- GreetingMessage Metabase Property
- BannerMessage Metabase Property
- MaxClientsMessage Metabase Property

# ExMdbGuid Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ExMdbGuid** property is used by the Microsoft Exchange driver.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 46162 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/ROOT<br>/LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ExpireNewsgroups Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ExpireNewsgroups** property lists newsgroups for which this expiration policy applies.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45658 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Expires/expire *n* | IIsNntpExpire |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ExpirePolicyName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ExpirePolicyName** property contains the name of the News Network Transport Protocol (NNTP) expiration policy.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45659 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC/*n*/Expires/expire *n* | IIsNntpExpire |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ExpireSpace Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ExpireSpace** property contains the size limit for expiration by size. This property is not currently supported.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45656 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Expires/expire *n* | IIsNntpExpire |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ExpireTime Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ExpireTime** property contains the amount of elapsed time before article expiration begins.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45657 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Expires/expire *n* | IIsNntpExpire |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedAccountName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedAccountName** property contains the user name when authentication is required for a feed.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45461 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedAdminError Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedAdminError** property contains the error text displayed after an administrative error.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45481 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedAllowControlMsgs Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedAllowControlMsgs** property allows control messages from a remote feed.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45466 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedAuthenticationType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedAuthenticationType** contains the feed authentication type (cleartext or none).

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45460 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedConcurrentSessions Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedConcurrentSessions** property contains the number of concurrent sessions allowed from feed server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45470 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedCreateAutomatically Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedCreateAutomatically** property controls whether groups are created automatically.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45467 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedDisabled Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedDisabled** property controls whether the feed is enabled or disabled.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45468 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedDistribution Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedDistribution** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45469 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedErrParmMask Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedErrParmMask** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45482 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedFeedpairId Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedFeedpairId** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45479 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedHandshake Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedHandshake** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45480 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedInterval Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedInterval** property corresponds to the "Run Interval" box on the **Schedule** tab of feed properties. The property defines how often the feed runs.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45465 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedMaxConnectionAttempts Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedMaxConnectionAttempts** property contains the number of times feed tries to connect without success before disabling the feed.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45471 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedNewsgroups Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedNewsgroups** property corresponds to the list of newsgroups that the feed accepts, asks for, or pushes.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45458 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedNextPullHigh Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedNextPullHigh** contains the high DWORD of the time of the next pull feed.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45474 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedNextPullLow Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedNextPullLow** contains the low DWORD of the time of the next pull feed.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45475 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedOutgoingPort Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedOutgoingPort** property corresponds to "Outbound IP Port" of the feed.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45478 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedPassword Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If authentication is required for feed, the **FeedPassword** property contains the password for the account.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45462 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedPeerGapSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedPeerGapSize** property contains the gap size to leave in the file before the article.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45477 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedPeerTempDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedPeerTempDirectory** property contains the temporary directory used for the feed.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45476 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds | IIsNntpFeeds |
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedPostHardLimit Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedPostHardLimit** property corresponds to **Limit connection size** for the feed on the **Settings** tab of NNTP virtual server Properties. Units are in kilobytes.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45167 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedPostSoftLimit Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedPostSoftLimit** property corresponds to **Limit post size** for the feed on the **Settings** tab of NNTP virtual server Properties. Units are in kilobytes.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45168 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedReportPeriod Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedReportPeriod** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45180 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedSecurityType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedSecurityType** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45459 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedServerName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedServerName** property corresponds to **Remote Server Name** on **General** tab of feed properties.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45456 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedStartTimeHigh Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedStartTimeHigh** property contains the high DWORD of the feed start time.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45463 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedStartTimeLow Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedStartTimeLow** property contains the low DWORD of the feed start time.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45464 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedTempDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedTempDirectory** contains the temporary directory used for feeds.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45473 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedType** property specifies whether the remote server is a master server, subordinate server, or peer server in relation to the local server.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45457 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FeedUucpName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FeedUucpName** property corresponds to **Unique path ID generated by remote server** on the **General** tab of feed properties.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45472 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC/*n*/Feeds/feed *n* | IIsNntpFeed |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FilterDescription Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The value of **FilterDescription** contains the description of the corresponding filter.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_FILTER_DESCRIPTION |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2045 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/*filter_name*<br><br>/LM/W3SVC/*n*/Filters/*filter_name* | IIsFilter |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FilterEnableCache Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FilterEnableCache** property can be set on individual filters to mark them cache friendly to HTTP.sys. *Cache friendly* means that the filter does not conditionally reroute URLs for /default.htm to either /default-1.htm or /default-2.htm, which would break the HTTP.sys cache. By default, filters are considered unfriendly to the HTTP.sys cache.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_FILTER_ENABLE_CACHE |
| Attributes | NO_ATTRIBUTES |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2046 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/*filter_name*<br><br>/LM/W3SVC/*n*/Filters/*filter_name* | IIsFilter |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FilterEnabled Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FilterEnabled** indicates whether the installed filter is enabled.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_FILTER_ENABLED |
| Attributes | NO_ATTRIBUTES |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2043 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/*filter_name*<br><br>/LM/W3SVC/*n*/Filters/*filter_name* | IIsFilter |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FilterFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FilterFlags** property contains flags that indicate the events that are registered to notify the installed ISAPI filter, which are defined in the Flags section.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

For more information, see "Developing ISAPI Filters" in the Platform SDK.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_FILTER_FLAGS |
| Attributes | NO_ATTRIBUTES |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2044 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/Filters/*filter_name*<br><br>/LM/W3SVC/*n*/Filters/*filter_name* | IIsFilter |

**Flags**

| **Flag Name** | **NotifySecurePort** |
| --- | --- |
| Description | Notify only if port secure. |
| Metabase Bitmask Identifier | MD_NOTIFY_SECURE_PORT |
| Decimal Value | 1 |
| Hexadecimal Value | 0x00000001 |

| Flag Name | NotifyNonSecurePort |
| --- | --- |
| Description | Notify only if port nonsecure. |
| Metabase Bitmask Identifier | MD_NOTIFY_NONSECURE_PORT |
| Decimal Value | 2 |
| Hexadecimal Value | 0x00000002 |

| Flag Name | NotifyReadRawData |
| --- | --- |
| Description | Notify if reading raw data. |
| Metabase Bitmask Identifier | MD_NOTIFY_READ_RAW_DATA |
| Decimal Value | 32768 |
| Hexadecimal Value | 0x00008000 |

| Flag Name | NotifyPreProcHeaders |
| --- | --- |
| Description | Notify if preprocessing headers. |
| Metabase Bitmask Identifier | MD_NOTIFY_PREPROC_HEADERS |
| Decimal Value | 16384 |
| Hexadecimal Value | 0x00004000 |

| Flag Name | NotifyAuthentication |
| --- | --- |
| Description | Notify if performing authentication. |
| Metabase Bitmask Identifier | MD_NOTIFY_AUTHENTICATION |
| Decimal Value | 8192 |
| Hexadecimal Value | 0x00002000 |

| Flag Name | NotifyUrlMap |
| --- | --- |
| Description | Notify if mapping URL to physical path. |
| Metabase Bitmask Identifier | MD_NOTIFY_URL_MAP |
| Decimal Value | 4096 |
| Hexadecimal Value | 0x00001000 |

| Flag Name | NotifyAccessDenied |
| --- | --- |
| Description | Notify if server sending HTTP error 401 to client. |
| Metabase Bitmask Identifier | MD_NOTIFY_ACCESS_DENIED |
| Decimal Value | 2048 |
| Hexadecimal Value | 0x00000800 |

| Flag Name | NotifySendResponse |
| --- | --- |
| Description | Notify if response being sent to client. |
| Metabase Bitmask Identifier | MD_NOTIFY_SEND_RESPONSE |

| | |
|---|---|
| Decimal Value | 64 |
| Hexadecimal Value | 0x00000040 |

| **Flag Name** | **NotifySendRawData** |
|---|---|
| Description | Notify if sending raw data. |
| Metabase Bitmask Identifier | MD_NOTIFY_SEND_RAW_DATA |
| Decimal Value | 1024 |
| Hexadecimal Value | 0x00000400 |

| **Flag Name** | **NotifyLog** |
|---|---|
| Description | Notify if logging. |
| Metabase Bitmask Identifier | MD_NOTIFY_LOG |
| Decimal Value | 512 |
| Hexadecimal Value | 0x00000200 |

| **Flag Name** | **NotifyEndOfRequest** |
|---|---|
| Description | Notify if request has ended. |
| Metabase Bitmask Identifier | MD_NOTIFY_END_OF_REQUEST |
| Decimal Value | 128 |
| Hexadecimal Value | 0x00000080 |

| **Flag Name** | **NotifyEndOfNetSession** |
|---|---|
| Description | Notify if network session is ending. |
| Metabase Bitmask Identifier | MD_NOTIFY_END_OF_NET_SESSION |
| Decimal Value | 256 |
| Hexadecimal Value | 0x00000100 |

| **Flag Name** | **NotifyOrderHigh** |
|---|---|
| Description | High priority for notifications. |
| Metabase Bitmask Identifier | MD_NOTIFY_ORDER_HIGH |
| Decimal Value | 524288 |
| Hexadecimal Value | 0x00080000 |

| **Flag Name** | **NotifyOrderMedium** |
|---|---|
| Description | Medium priority for notifications. |
| Metabase Bitmask Identifier | MD_NOTIFY_ORDER_MEDIUM |
| Decimal Value | 262144 |
| Hexadecimal Value | 0x00040000 |

| Flag Name | NotifyOrderLow |
| --- | --- |
| Description | Low priority for notifications. |
| Metabase Bitmask Identifier | MD_NOTIFY_ORDER_LOW |
| Decimal Value | 131072 |
| Hexadecimal Value | 0x00020000 |

| Flag Name | NotifyAuthComplete |
| --- | --- |
| Description | Notify if the client's identity has been negotiated with the client. Because of the timing of this notification, the AUTH_USER server variable can be used to reliably obtain the identity of the user. If the value of this property is 1, then the filter registers for the SF_NOTIFY_AUTH_COMPLETE notification. If the value is 0, the filter does not. |
| Metabase Bitmask Identifier | MD_NOTIFY_AUTH_COMPLETE |
| Decimal Value | 67108864 |
| Hexadecimal Value | 0x04000000 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FilterLoadOrder Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The value of **FilterLoadOrder**, is a comma-delimited list, which specifies the order that filters are to be loaded, such as Filter1, Filter2, and so on.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_FILTER_LOAD_ORDER |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2040 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters<br>/LM/W3SVC/*n*/Filters | IIsFilters |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FilterPath Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FilterPath** property specifies the full path to a filter DLL; for example C:\Wwwroot\Bin\Filter.dll.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_FILTER_IMAGE_PATH |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2041 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters<br><br>/LM/W3SVC/*n*/Filters | IIsFilters |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FilterState Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FilterState** property indicates whether the installed filter is loaded or unloaded, as defined in the Bitmask Values table.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_FILTER_STATE |
| Attributes | NO_ATTRIBUTES |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2042 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/Filters/*filter_name*<br><br>/LM/W3SVC/*n*/Filters/*filter_name* | IIsFilter |

**Values**

| Constant | Value | Description |
| --- | --- | --- |
| MD_FILTER_STATE_LOADED | 0x00000001 | Filter is loaded. |
| MD_FILTER_STATE_UNLOADED | 0x00000004 | Filter is unloaded. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FrontPageWeb Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FrontPageWeb** property indicates whether a server instance will be targeted by Microsoft FrontPage. Setting **FrontPageWeb** to **true** causes FrontPage Manager to create the files required for FrontPage Server Extensions. Setting **FrontPageWeb** to **false** causes these files to be deleted.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_FRONTPAGE_WEB |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2072 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br><br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| LM/W3SVC/*n*/ROOT/*physical_directory_name*<br><br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FsPropertyPath Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FsPropertyPath** property contains the path to driver property files.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 46158 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/ROOT<br>/LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FtpDirBrowseShowLongDate Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FtpDirBrowseShowLongDate** indicates whether to display two-digit or four-digit years when browsing directories.

| **Note** |
| --- |
| The **MSDOSDirOutput** property allows UNIX-style directory browsing, and if set to **false**, **FtpDirBrowseShowLongDate** will have no effect. **FtpDirBrowseShowLongDate** is valid only when **MSDOSDirOutput**=**true**, which sets the directory output to MS-DOS format. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_SHOW_4_DIGIT_YEAR |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 5010 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/MSFTPSVC/*n*/ROOT<br>/LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FtpLogInUtf8 Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FtpLogInUtf8** property is used by IIS.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_FTP_LOG_IN_UTF_8 |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5013 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# FullyQualifiedDomainName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **FullyQualifiedDomainName** contains the fully qualified domain name (FQDN) of the Simple Mail Transfer Protocol (SMTP) virtual server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_FQDN_VALUE |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36935 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# GreetingMessage Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **GreetingMessage** property specifies the message text an FTP server transmits to a client once the client logs in to the FTP server and is authenticated.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_GREETING_MESSAGE |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5002 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- ExitMessage Metabase Property
- BannerMessage Metabase Property
- MaxClientsMessage Metabase Property

# GroupHelpFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **GroupHelpFile** property contains the path to the descript.txt file, which is currently unused.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45156 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# GroupListFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **GroupListFile** property contains the path to the *group.lst* file.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45157 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# GroupvarListFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **GroupvarListFile** property contains the path to the *groupvar.lst* file.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45182 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcCacheControlHeader Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcCacheControlHeader** property specifies the directive that IIS adds to the Cache-Control header. This header, along with the Expires header specified by HcExpiresHeader Metabase Property, is included with every compressed file that is sent to a client browser.

The **HcCacheControlHeader** property specifies a header that overrides the HTTP Expires header, which is specified by **HcExpiresHeader**, and is included to ensure that older clients and proxy servers do not attempt to cache compressed files.

Neither the header specified by **HcCacheControlHeader** nor the header specified by **HcExpiresHeader** is sent with a response if the metabase property HcSendCacheHeaders Metabase Property is set to **false**.

The World Wide Web Publishing Service (WWW service) must be restarted before changes to this property take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_HC_CACHE_CONTROL_HEADER |
| Attributes | NO_ATTRIBUTES |
| Default Value | max-age=86400 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2211 |

**Configurable Locations**
You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**
For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcExpiresHeader Metabase Property
- HcSendCacheHeaders Metabase Property

# HcCompressionBufferSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcCompressionBufferSize** property specifies the size of the compression buffer (in bytes) that IIS uses for receiving compressed data. This buffer is used for on-demand compression of static content, as well as dynamic content compression. A larger buffer yields slightly faster compression performance, but at the cost of using additional memory.

The size limit specified by **HcCompressionBufferSize** is used by IIS, only if either HcDoOnDemandCompression Metabase Property or HcDoDynamicCompression Metabase Property are set to **true**. The maximum allowable size of the compression buffer is 1,048,576 bytes (1 MB).

The World Wide Web Publishing Service (WWW service) must be restarted before changes to this property take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_COMPRESSION_BUFFER_SIZE |
| Attributes | NO_ATTRIBUTES |
| Default Value | 8192 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2223 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcDoOnDemandCompression Metabase Property
- HcDoDynamicCompression Metabase Property

# HcCompressionDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcCompressionDirectory** property specifies the directory where compressed versions of static files are temporarily cached. If compression is enabled for IIS, every time IIS receives a request for a static file, such as an .htm or .txt file, the server searches the compression directory for an already compressed version of the requested static file.

To force IIS to recompress all outgoing static files, delete all files found in the compression directory specified by **HcCompressionDirectory**.

| **Important** |
| --- |
| It is strongly recommended that the directory specified by **HcCompressionDirectory** reside on an NTFS file system volume, not on a FAT volume. If the file being requested resides on an NTFS volume, and you have set **HcCompressionDirectory** to specify a FAT volume, IIS will not perform HTTP compression on that file. This behavior is due to a number of differences between FAT and NTFS volumes, such as time stamping and access control mechanisms. In addition, NTFS handles directories that contain a large number of files more efficiently than FAT volumes do. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | EXPANDSZ |
| WMI Data Type | String |
| ADSI Data Type | EXPANDSZ |
| ABO Data Type | EXPANDSZ |
| ABO Metabase Identifier | MD_HC_COMPRESSION_DIRECTORY |
| Attributes | NO_ATTRIBUTES |
| Default Value | %windir%\iis temporary compressed files |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2210 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcDoOnDemandCompression Metabase Property

# HcCompressionDll Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcCompressionDll** property specifies the fully qualified file system path and file name of the compression DLL associated with the compression scheme.

The World Wide Web Publishing Service (WWW service) must be restarted before changes to this property take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | EXPANDSZ |
| WMI Data Type | String |
| ADSI Data Type | EXPANDSZ |
| ABO Data Type | EXPANDSZ |
| ABO Metabase Identifier | MD_HC_COMPRESSION_DLL |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2237 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/gzip<br>/LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcCreateFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcCreateFlags** property contains a flag that is passed to the **HcCompressionDLL** property. This flag identifies the compression schema to be used (either "gzip" or "deflate").

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_CREATE_FLAGS |
| Attributes | NO_ATTRIBUTES |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2243 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/gzip /LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcDoDiskSpaceLimiting Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcDoDiskSpaceLimiting** allows you to limit the amount of disk space that all compressed files, which are stored in the compression directory specified by HcCompressionDirectory Metabase Property can occupy. If **HcDoDiskSpaceLimiting** is set to **true**, IIS will not allow compressed files to occupy more than the number of bytes specified by HcMaxDiskSpaceUsage Metabase Property. Once the limit is reached, compressed files are removed from the compression directory on an LRU (least-recently) basis. If **HcDoDiskSpaceLimiting** is set to **false**, no limit is enforced on the amount of disk space that the files in the compression directory can occupy.

| **Note** |
| --- |
| Limiting the use of disk space has a significant impact on performance because additional overhead is required to check for the limit. Use this feature only when you lack adequate disk space to cache all content files on your Web server. |

**HcDoDiskSpaceLimiting** is relevant only if you have set the property HcDoOnDemandCompression Metabase Property to **true**.

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_DO_DISK_SPACE_LIMITING |
| Attributes | NO_ATTRIBUTES |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2216 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcCompressionDirectory Metabase Property
- HcMaxDiskSpaceUsage Metabase Property

- HcDoOnDemandCompression Metabase Property

# HcDoDynamicCompression Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcDoDynamicCompression** property specifies whether responses to requests for dynamic content, such as scripts in ASP pages and ISAPI extensions are compressed. Setting **HcDoDynamicCompression** to **true**, indicates that dynamic content is compressed when requested by a client. If dynamic compression is enabled, each request for dynamic content that contains the "Accept-Encoding: gzip" or "Accept-Encoding: deflate" header causes the content to be compressed, as long as the compression scheme that is specified in the header is enabled in IIS. If the request does not contain this header, or the header specifies a compression scheme that is not enabled in IIS, the content is not compressed.

| **Note** |
| --- |
| Because dynamic content is always changing, IIS does not cache compressed versions of dynamic content. Dynamic compression consumes considerable CPU time and memory resources, and should only be used on servers that have slow network connections and CPU time to spare. |

At the individual compression scheme level (**IIsCompressionScheme**), the **HcDoDynamicCompression** property determines which compression schemes IIS will support for dynamic content: gzip, deflate, or both.

You can either enable compression of dynamic content server-wide and then disable it at specific locations, or you can disable compression of dynamic content server-wide and then enable it only at the locations you specify. To enable and disable dynamic compression by using the **HcDoDynamicCompression** property, do the following:

- To enable dynamic content compression server-wide, set the **HcDoDynamicCompression** property to **true** at the **IIsCompressionSchemes** level. You can then disable dynamic content compression settings for an individual Web site or site element, for example, a virtual directory or file, by using the DoDynamicCompression Metabase Property at the individual site or site element level.

- To disable dynamic content compression server-wide, set the **HcDoDynamicCompression** property to **false** at the **IIsCompressionSchemes** level. You can then enable dynamic content compression settings for an individual Web site or site element, for example, a virtual directory or file, by using the DoDynamicCompression Metabase Property at the individual site or site element level.

| **Important** |
| --- |
| If you change the value of the **HcDoDynamicCompression**property at the **IIsCompressionScheme** level, the World Wide Web Publishing Service (WWW service) must be restarted before the change takes effect. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_DO_DYNAMIC_COMPRESSION |
| Attributes | NO_ATTRIBUTES |
| Default Value | **false** |

| | |
|---|---|
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2213 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/gzip<br><br>/LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcDoOnDemandCompression Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcDoOnDemandCompression** property specifies whether static files, such as .htm and .txt files, are compressed if a compressed version of the file does not exist in the compression directory specified by HcCompressionDirectory Metabase Property.

For on-demand compression of static files to occur, **HcDoOnDemandCompression** must be set to **true** at the global (**IIsCompressionSchemes**) level and at least one compression scheme (**IIsCompressionScheme**) level. The supported compression schemes are gzip and deflate. For information about the levels that this property can be set, see the Access Locations table.

If **HcDoOnDemandCompression** is set to **true** at the **IIsCompressionScheme** level and at least one **IIsCompressionSchemes** level, and IIS receives a request for static content, IIS checks the compression directory. If a corresponding compressed version of the requested file is not found, IIS will send an uncompressed version of the requested file to the client browser, while a background thread compresses the requested file. The newly compressed file is then stored in the compression directory, and subsequent requests for that file are serviced directly from the compression directory. In other words, an uncompressed version of the file is returned to the client, unless a compressed version of the file already exists in the compression directory.

If **HcDoOnDemandCompression** is set to **false**, at the global level, IIS does not compress static content files when a corresponding compressed version is not found in the compression directory.

If you change the value of **HcDoOnDemandCompression** at the individual compression scheme level, the World Wide Web Publishing Service (WWW service) must be restarted for the change to take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_DO_ON_DEMAND_COMPRESSION |
| Attributes | NO_ATTRIBUTES |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2215 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/gzip /LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcCompressionDirectory Metabase Property

# HcDoStaticCompression Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcDoStaticCompression** property specifies whether responses to requests for static content, such as such as .htm and .txt files are compressed. Setting **HcDoStaticCompression** to **true**, indicates that static content is compressed when requested by a client. If static compression is enabled, each request for static content that contains the "Accept-Encoding: gzip" or "Accept-Encoding: deflate" header causes the content to be compressed, so long as the compression scheme specified in the header is enabled in IIS. If the request does not contain this header, or the header specifies a compression scheme that is not enabled in IIS, the content is not compressed.

If IIS determines that the response should be compressed, IIS checks the compression directory specified by the HcCompressionDirectory Metabase Property. If a fresh, compressed version of the static file is found in the directory, it is sent to the client browser. If a fresh, compressed version of the file is not found, and on-demand compression is enabled with HcDoOnDemandCompression Metabase Property, IIS sends the requested file in uncompressed form, and adds that file to the background compression queue.

At the individual compression scheme level (**IIsCompressionScheme**), the **HcDoStaticCompression** property determines which compression schemes IIS will support for static content: gzip, deflate, or both.

You can either enable compression of static content server-wide and then disable it at specific locations, or you can disable compression of static content server-wide and then enable it only at the locations you specify. To enable and disable static compression by using the **HcDoStaticCompression** property, do the following:

- To enable static compression server-wide, set the **HcDoStaticCompression** property to **true** at the **IIsCompressionSchemes** level. You can then disable static file compression settings for an individual Web site or site element -- for example, a virtual directory or file -- by using the DoStaticCompression property at the individual site or site element level.

- To disable static compression server-wide, set the **HcDoStaticCompression** property to **false** at the **IIsCompressionSchemes** level. You can then enable static file compression for an individual Web site or site element -- for example, a virtual directory or file -- by using the DoStaticCompression property at the individual site or site element level.

| **Important** |
| --- |
| If you change the value of the **HcDoStaticCompression** property at the **IIsCompressionScheme** level, the World Wide Web Publishing Service (WWW service) must be restarted before the change takes effect. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_DO_STATIC_COMPRESSION |
| Attributes | NO_ATTRIBUTES |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |

| | |
|---|---|
| User Type | IIS_MD_UT_SERVER |
| ID | 2214 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/gzip /LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcDynamicCompressionLevel Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcDynamicCompressionLevel** property specifies the compression level for the compression scheme, when the scheme is compressing dynamic content. Low compression levels produce slightly larger compressed files, but with lower overall impact on CPU and memory resources. Higher compression levels generally result in smaller compressed files, but with higher CPU and memory usage.

The World Wide Web Publishing Service (WWW service) must be restarted for any changes to this property take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_DYNAMIC_COMPRESSION_LEVEL |
| Attributes | NO_ATTRIBUTES |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2241 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/gzip /LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcExpiresHeader Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcExpiresHeader** property specifies the content of the HTTP Expires header that is sent with all requested compressed files, along with the Cache-Control header discussed in HcCacheControlHeader Metabase Property.

The combination of **HcExpiresHeader** and **HcCacheControlHeader** ensures that older clients and proxy servers do not attempt to cache compressed files.

Neither the header specified by **HcCacheControlHeader**, nor the header specified by **HcExpiresHeader**, is sent with the response if the metabase property HcSendCacheHeaders Metabase Property is set to **false**.

The World Wide Web Publishing Service (WWW service) must be restarted before any changes to this property take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_HC_EXPIRES_HEADER |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2212 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcCacheControlHeader Metabase Property
- HcSendCacheHeaders Metabase Property

# HcFileExtensions Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcFileExtensions** indicates which file name extensions are supported by the compression scheme. Only static files with the specified file extensions are compressed by IIS. If this setting is empty, no static files are compressed.

In addition to file name extension matches, HcDoOnDemandCompression Metabase Property and HcDoStaticCompression Metabase Property must be set to **true** for compression to be performed.

The World Wide Web Publishing Service (WWW service) must be restarted for any changes to this property take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_HC_FILE_EXTENSIONS |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2238 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/gzip<br><br>/LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcScriptFileExtensions Metabase Property

# HcFilesDeletedPerDiskFree Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcFilesDeletedPerDiskFree** property specifies the number of compressed files that are deleted from the compression directory when IIS determines that the compression directory has grown beyond its limits. IIS deletes a number of files equal to the setting of this property, starting with the LRU (least-recently used) files.

The **HcFilesDeletedPerDiskFree** property is relevant only if HcDoDiskSpaceLimiting Metabase Property is set to **true**, and HcMaxDiskSpaceUsage Metabase Property is set to something other than 0xFFFFFFFF (unlimited). The maximum allowable value of **HcFilesDeletedPerDiskFree** is 1024.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_FILES_DELETED_PER_DISK_FREE |
| Attributes | NO_ATTRIBUTES |
| Default Value | 256 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2225 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcIoBufferSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcIoBufferSize** property specifies the size of the buffer (in bytes) that IIS uses to read uncompressed files. A larger buffer yields slightly faster compression performance, but at the cost of additional memory usage.

**HcIoBufferSize** is used by IIS only if HcDoOnDemandCompression Metabase Property is set to **true**. The maximum allowable value is 1,048,576 bytes (1 MB).

The World Wide Web Publishing Service (WWW service) must be restarted before any changes to this property take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_IO_BUFFER_SIZE |
| Attributes | NO_ATTRIBUTES |
| Default Value | 8192 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2222 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcMaxDiskSpaceUsage Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcMaxDiskSpaceUsage** property specifies whether IIS limits the number of bytes of disk space occupied by compressed files in the compression directory.

Disk space limiting is performed, and **HcMaxDiskSpaceUsage** is checked by IIS, only if HcDoDiskSpaceLimiting Metabase Property is set to **true**. If **HcMaxDiskSpaceUsage** is set to 4294967295, there will be no limit on the number of bytes occupied by compressed files.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_MAX_DISK_SPACE_USAGE |
| Attributes | NO_ATTRIBUTES |
| Default Value | 1000000 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2221 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcMaxQueueLength Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS performs on-demand compression of static content in the background. When a static file is requested, and IIS has determined that a previously compressed version of the requested file does not already exist in the compression directory, the request is put into a simple FIFO (first in, first out) queue. The **HcMaxQueueLength** property specifies how many concurrent background compression requests can be in the queue. If the queue is full, additional compression requests are ignored until there is room in the queue.

If **HcMaxQueueLength** is set to 4294967295, there is no limit to the number of pending requests in the queue.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_MAX_QUEUE_LENGTH |
| Attributes | NO_ATTRIBUTES |
| Default Value | 1000 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2224 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcMinFileSizeForComp Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcMinFileSizeForComp** allows you to specify the minimum number of bytes a file must contain for it to be compressed using on-demand compression. Very small files do not compress well, and in some cases, compression may even make the file larger than it was initially.

If **HcMinFileSizeForComp** is set to 0 (zero), all files are compressed.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_MIN_FILE_SIZE_FOR_COMP |
| Attributes | NO_ATTRIBUTES |
| Default Value | 256 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2226 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcNoCompressionForHttp10 Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcNoCompressionForHttp10** property disables compression for requests containing an HTTP 1.0 version number. HTTP 1.0, as described in RFC 1945, provides a minimal level of support for certain types of compression. However, some confusion exists around HTTP 1.0 compression, especially with regard to proxy servers. To minimize the chance of inappropriately returning a compressed file to a client that cannot decompress it, you can use this metabase property to disable compression in questionable scenarios.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_NO_COMPRESSION_FOR_HTTP_10 |
| Attributes | NO_ATTRIBUTES |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2217 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcNoCompressionForProxies Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcNoCompressionForProxies** allows you to disable the HTTP 1.1 response for compression requests that have come through proxy servers. Certain HTTP proxy servers, including some advertised as HTTP 1.1 compliant, do not correctly handle the caching of compressed objects.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_NO_COMPRESSION_FOR_PROXIES |
| Attributes | NO_ATTRIBUTES |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2218 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcNoCompressionForRange Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcNoCompressionForRange** property specifies how IIS handles HTTP Range requests. The HTTP 1.1 Request for Comment, RFC 2068, is ambiguous about whether Range requests apply to the compressed or uncompressed version of a file. Some browsers could be unprepared for compressed Range responses. As a result, IIS, by default, does not send compressed responses to Range requests, to avoid problems with browser misbehavior.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_NO_COMPRESSION_FOR_RANGE |
| Attributes | NO_ATTRIBUTES |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2219 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcOnDemandCompLevel Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcOnDemandCompLevel** property specifies the compression level for the compression scheme, when the scheme is to compress static content on demand. Low compression levels produce slightly larger compressed files, but with lower overall impact on CPU and memory resources. Higher compression levels generally result in small compressed files but with higher CPU and memory usage.

Valid compression levels range from 1 to 10.

The World Wide Web Publishing Service (WWW service) must be restarted before any changes to this property take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_ON_DEMAND_COMP_LEVEL |
| Attributes | NO_ATTRIBUTES |
| Default Value | 10 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2242 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/gzip /LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcDynamicCompressionLevel Metabase Property

# HcPriority Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcPriority** property specifies the priority rating assigned to a particular compression scheme. If your installation of IIS has multiple compression schemes installed and a client browser indicates, in the Accept-Encoding header, that it can handle multiple compression schemes, IIS uses the priority numbers specified in **HcPriority** to determine which scheme to use for the request. In that case, IIS uses the matching scheme that has the highest priority number.

Valid priority numbers range from 1 to 10. In general, if you have multiple compression schemes installed with IIS, you should assign each compression scheme a different priority number.

The World Wide Web Publishing Service (WWW service) must be restarted before any changes to this property take effect.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HC_PRIORITY |
| Attributes | NO_ATTRIBUTES |
| Default Value | 1 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2240 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Filters/Compression/gzip /LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HcScriptFileExtensions Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcScriptFileExtensions** indicates which file name extensions are supported by the compression scheme. The output from dynamic files with the file extensions specified in this property are compressed by IIS.

| **Important** |
| --- |
| Setting this parameter to an empty string can have an adverse effect on your server's performance. |

If this parameter is empty, all dynamic responses are sent compressed. In addition, any static file type not specified in HcFileExtensions Metabase Property is dynamically compressed and therefore, not cached.

Also, HcDoDynamicCompression Metabase Property must be set to **true** for dynamic compression to be performed.

The World Wide Web Publishing Service (WWW service) must be restarted before any changes to this property take effect.

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_HC_SCRIPT_FILE_EXTENSIONS |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2244 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/Filters/Compression/gzip /LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- HcFileExtensions Metabase Property

# HcSendCacheHeaders Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HcSendCacheHeaders** property specifies whether the headers specified by HcCacheControlHeader Metabase Property and HcExpiresHeader Metabase Property are sent with each compressed response. If this property is set to **true**, Cache-Control and Expires headers are sent with all compressed responses. If this property is set to **false**, however, only the HTTP 1.1 Vary header is sent.

| **Important** |
| --- |
| In general, the Cache-Control and Expires HTTP headers are included with compressed responses to prevent inappropriate caching of compressed files, especially by proxy servers and older browsers. If you choose to disable these headers, you should also set HcNoCompressionForProxies Metabase Property and HcNoCompressionForHttp10 Metabase Property to **true**. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_HC_SEND_CACHE_HEADERS |
| Attributes | NO_ATTRIBUTES |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2220 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HeaderWaitTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HeaderWaitTimeout** property specifies the amount of time (in seconds) the server waits for all HTTP headers for the request (indicated by a double carriage return) to be received before disconnecting the client. The purpose of this property is to help prevent a common variant of the Denial of Service (DoS) attack that attempts to max out connection limits and keep those connections connected.

In IIS 6.0, use this property in conjunction with the **ConnectionTimeout** property, and for Web service nodes, use the **MinFileBytesPerSec** property instead of the **ServerListenTimeout** property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HEADER_WAIT_TIMEOUT |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9204 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HistoryExpiration Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HistoryExpiration** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45057 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HistoryMajorVersionNumber Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HistoryMajorVersionNumber** is used internally only by IIS to keep track of the major version number of the MetaBase.xml and MBSchema.xml files. The major version number is also used to version and name history files.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_GLOBAL_EDIT_WHILE_RUNNING_MAJOR_VERSION_NUMBER |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| StartingNumber | Not applicable |
| EndingNumber | Not applicable |
| ID | 9994 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| . | IIS_Global |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HistoryTableFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HistoryTableFile** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45159 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HonorClientMsgIds Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HonorClientMsgIds** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45058 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HopCount Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HopCount** property specifies the number of hops that are allowed before a message is determined to be in a looping condition and marking for non-delivery.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_HOP_COUNT |
| Attributes | INHERIT |
| Default Value | 10 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36886 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/n | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HttpCustomHeaders Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HttpCustomHeaders** property contains custom headers that are sent to the client, in addition to the default header of the HTML file. Each string in this property is formatted as a key name and value pair: *Keyname: Value*.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_HTTP_CUSTOM |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6004 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HttpErrors Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HttpErrors** property specifies the custom error string sent to clients in response to HTTP 1.1 errors. Each string in the list specifies the HTTP error code and subcode, indicates whether the handler will be a URL or a file, and specifies which URL or file the client will be sent. Each string can be in either format: a URL or a file.

To send a URL to the client, use *HTTPErrorCode*, *HTTPErrorSubcode*, URL, *HandlerURL*. To send a file, use *HTTPErrorCode*, *HTTPErrorSubcode*, FILE, *Filename*.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_CUSTOM_ERROR |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6008 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- CustomErrorDescriptions Metabase Property

# HttpExpires Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HttpExpires** property specifies the expiration of HTML document content by returning the value to the browser in the HTML file header. The browser compares the given value with the current date to determine whether to display a cached page or to request an updated page from the server. This property can be made relative, or dynamic, by using the format D, *#SecondsUntilExpiration*, with 0xFFFFFFFF indicating no expiration date. It can also be absolute, with the format *S,ValidGMTstring*.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_HTTP_EXPIRES |
| Attributes | INHERIT |
| Default Value | d, 0x15180 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6002 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HttpPics Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HttpPics** property specifies the World Wide Web Consortium (W3C) Platform for Internet Content Selection (PICS) rating header. For more information about PICS, see the World Wide Web Consortium Web site.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_HTTP_PICS |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6003 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# HttpRedirect Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **HttpRedirect** property specifies the directory or URL to which a client is redirected when attempting to access a specific resource. There are two general forms that the value for this property can take.

The simple format is *Destination,Flag*.

- Where Destination can specify either a URL or a virtual path to a file.
- Where *Flag* can have one of four values:
- EXACT_DESTINATION indicates that the value provided for *Destination* should be considered an absolute target location.
- CHILD_ONLY indicates that the value should be added to the beginning of the file name of the request to be redirected.
- PERMANENT indicates that this redirection is permanent for this resource.
- TEMPORARY indicates that this redirection is temporary for this resource.

The more complex form involves the use of wildcards. The format is *; *Wildcard1*; *Destination1*; *Wildcard2*; *Destination2*, *Flag*. Each *Wildcard*; *Destination* pair indicates that requests matching the wildcard are redirected to the specified destination. *Flag* can have the same values as the simple string format.

| **Important** |
| --- |
| Setting **HttpRedirect** at the site or service level has no affect. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_HTTP_REDIRECT |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6011 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# IdleTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **IdleTimeout** property specifies how long (in minutes) a worker process should run idle if no new requests are received and the worker process is not processing requests. After the allotted time passes, the worker process should request to be shut down by the World Wide Web Publishing Service (WWW Service).

| **Note** |
| --- |
| If the value of this property is less than the value of PeriodicRestartTime Metabase Property, IIS sets both properties to their default values. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_IDLE_TIMEOUT |
| Attributes | INHERIT |
| Default Value | 20 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9005 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Iis5IsolationModeEnabled Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **IIs5IsolationModeEnabled** property specifies whether to run in IIS 5.0 isolation mode. Setting this property to **true** causes IIS to run in IIS 5.0 isolation mode. Setting this property to **false** causes IIS to run in worker process isolation mode.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_GLOBAL_STANDARD_APP_MODE_ENABLED |
| Attributes | NO_ATTRIBUTES |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9203 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapClearTextProvider Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49375 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapDefaultDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49368 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapDsAccount Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49383 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapDsBindType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49379 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapDsDataDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49377 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapDsDefaultMailRoot Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49378 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapDsHost Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49381 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapDsNamingContext Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49382 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapDsPassword Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT \| SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49384 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapDsSchemaType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49380 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapExpireDelay Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 1440 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49370 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapExpireMail Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49373 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapExpireStart Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49371 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapMailExpirationTime Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49153 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapRoutingDll Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49369 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapRoutingSources Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | iis_md_ut_wam |
| ID | 49366 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapRoutingTableType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | ldap |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49376 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ImapServiceVersion Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 49652 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# InProcessIsapiApps Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **InProcessIsapiApps** property specifies a list of ISAPI filters and extensions that must be run in the Web server process. The string format is either *Fullpath\Appname*.dll, or just *Appname*.dll. ISAPI DLL entries without a qualifying path will match any ISAPI request that contains the DLL name, regardless of the physical path specified in the request.

For example, any of the following would be a valid list for this property: "D:\winnt\system32\inetsrv\ssinc.dll", "D:\winnt\system32\inetsrv\httpodbc.dll", "author.dll", "shtml.dll", "admin.dll".

| **Note** |
| --- |
| This property only indicates which ISAPI applications are required to run in the Web server process. The inclusion of an application in this list does not imply that the application is automatically run when starting up the server. There can be significant performance issues if too many applications are forced to run outside of the server process. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_IN_PROCESS_ISAPI_APPS |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2073 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- AppIsolated Metabase Property

# IPSecurity Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **IPSecurity** property specifies the IP access restrictions for a URL. This property can be used to assign or deny access to browsers, based on either their IP address or DNS host name. For more information about managing this property, see the **IIsIPSecurity** object in the IIS ADSI Provider reference section of the IIS Software Development Kit (SDK).

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | IPSECLIST |
| WMI Data Type | IPSECLIST |
| ADSI Data Type | IPSECLIST |
| ABO Data Type | IPSECLIST |
| ABO Metabase Identifier | MD_IP_SEC |
| Attributes | inherit \| reference |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6019 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/MSFTPSVC/*n*/ROOT /LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

| | |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# IsapiRestrictionList Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is controlled by IIS.

| Attribute Name | Attribute Value |
|---|---|
| ADSI/WMI Data Type | MULTISZ |
| ABO Data Type | MULTISZ_METADATA |
| Schema Default | **null** |
| Internal Default | Not specified. |
| Upper Bound | Not specified. |
| Lower Bound | Not specified. |
| Internal ID | 2163 |
| Friendly ID | MD_ISAPI_RESTRICTION_LIST |
| Property Attributes | NO_ATTRIBUTES |
| User Type | IIS_MD_UT_SERVER |

**Configurable Locations**
This property is not configurable at any location in the metabase.

**Flags**
There are no flags for this property.

# KeyType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **KeyType** property specifies the type of metabase key and the IIS Admin Object that relates to that key. The **KeyType** property specifies the class of the object associated with the key, such as **IIsWebServer** or **IIsWebVirtualDir**. The **KeyType** determines the set of properties accessible at the key.

| **Note** |
| --- |
| The **KeyType** property is accessible by all IIS Admin Objects. The value is the name of the IIS Admin Object accessed, and is the same as the ADSI Class property. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_KEY_TYPE |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1002 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM | IIsComputer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

| | |
|---|---|
| /LM/MSFTPSVC/*n*/ROOT /LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |
| /LM/W3SVC/Filters/*filter_name* /LM/W3SVC/*n*/Filters/*filter_name* | IIsFilter |
| /LM/W3SVC/Filters /LM/W3SVC/*n*/Filters | IIsFilters |
| /LM/W3SVC/Filters/Compression/gzip /LM/W3SVC/Filters/Compression/deflate | IIsCompressionScheme |
| /LM/W3SVC/Filters/Compression/Parameters | IIsCompressionSchemes |
| /LM/W3SVC/*n*/IIsCertMapper | IIsCertMapper |
| /LM/MimeMap | IIsMimeMap |
| /LM/logging | IIsLogModules |
| /LM/logging/Microsoft IIS Log File Format /LM/logging/NCSA Common Log File Format /LM/logging/ODBC Logging /LM/logging/W3C Extended Log File Format | IIsLogModule |
| /LM/logging/Custom Logging /LM/logging/Custom Logging/Date /LM/logging/Custom Logging/Time /LM/logging/Custom Logging/Extended Properties /LM/logging/Custom Logging/Extended Properties/*extended_property_name* | IIsCustomLogModule |
| /LM/W3SVC/Info | IIsWebInfo |
| /LM/MSFTPSVC/Info | IIsFtpInfo |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |
| /LM/NNTPSVC/Info | IIsNntpInfo |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/ROOT /LM/SMTPSVC/*n*/ROOT/*virtual_directory_name* | IIsSmtpVirtualDir |
| /LM/SMTPSVC/*n*/Domain | IIsSmtpDomain |
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |
| /LM/SMTPSVC/Info | IIsSmtpInfo |
| /LM/NNTPSVC/*n*/Rebuild | IIsNntpRebuild |
| /LM/NNTPSVC/*n*/Sessions | IIsNntpSessions |
| /LM/NNTPSVC/*n*/Feeds | IIsNntpFeeds |
| /LM/NNTPSVC/*n*/Feeds/feed*n* | IIsNntpFeed |
| /LM/NNTPSVC/*n*/Expires | IIsNntpExpiration |
| /LM/NNTPSVC/*n*/Expires/expire*n* | IIsNntpExpire |
| /LM/NNTPSVC/*n*/Groups | IIsNntpGroups |
| /LM/SMTPSVC/*n*/Sessions | IIsSmtpSessions |
| / | IIS_ROOT |

| | |
|---|---|
| . | IIS_Global |
| /LM/SMTPSVC/*n*/Alias | IIsSmtpAlias |
| /LM/SMTPSVC/*n*/DL | IIsSmtpDL |
| /LM/SMTPSVC/*n*/user | IIsSmtpUser |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LimitRemoteConnections Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LimitRemoteConnections** property specifies the maximum number of inbound remote connections allowed at one time for a given virtual server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36926 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ListFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ListFile** property contains the path to the *listfile.hsh* file.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45172 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LoadBalancerCapabilities Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LoadBalancerCapabilities** property specifies behavior when a service is unavailable. A setting of 1 terminates the connection. A setting of 2 sends error code 503.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOAD_BALANCER_CAPABILITIES |
| Attributes | INHERIT |
| Default Value | 2 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9034 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LocalDomains Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LocalDomains** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36919 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LocalRetryAttempts Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LocalRetryAttempts** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOCAL_RETRY_ATTEMPTS |
| Attributes | INHERIT |
| Default Value | 48 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36896 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LocalRetryInterval Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LocalRetryInterval** property specifies the number of minutes between retries for local delivery.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOCAL_RETRY_MINUTES |
| Attributes | INHERIT |
| Default Value | 60 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36897 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogAnonymous Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogAnonymous** property specifies whether anonymous requests are logged to the event log.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_LOG_ANONYMOUS |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5007 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogCustomPropertyDataType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogCustomPropertyDataType** property contains bitmask values that indicate the data type of a custom logging field. The bitmask values are listed in the Bitmask Values section.

For more information about custom logging fields, see "Setting Metabase Properties for Logging" in the Platform SDK.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOGCUSTOM_SERVICES_string |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4505 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/logging/Custom Logging<br>/LM/logging/Custom Logging/Date<br>/LM/logging/Custom Logging/Time<br>/LM/logging/Custom Logging/Extended Properties<br>/LM/logging/Custom Logging/Extended Properties/*extended_property_name* | IIsCustomLogModule |

**Values**

| Constant | Value | Description |
|---|---|---|
| MD_LOGCUSTOM_DATATYPE_INT | 0x00000000 | Datatype is int. |
| MD_LOGCUSTOM_DATATYPE_UINT | 0x00000001 | Datatype is unsigned int. |
| MD_LOGCUSTOM_DATATYPE_LONG | 0x00000002 | Datatype is long. |
| MD_LOGCUSTOM_DATATYPE_ULONG | 0x00000003 | Datatype is unsigned long. |

| | | |
|---|---|---|
| MD_LOGCUSTOM_DATATYPE_FLOAT | 0x00000004 | Datatype is float. |
| MD_LOGCUSTOM_DATATYPE_DOUBLE | 0x00000005 | Datatype is double. |
| MD_LOGCUSTOM_DATATYPE_LPSTR | 0x00000006 | Datatype is long pointer to string array. |
| MD_LOGCUSTOM_DATATYPE_LPWSTR | 0x00000007 | Datatype is lpwstr. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogCustomPropertyHeader Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogCustomPropertyHeader** property indicates which header string will be written to the log file for the custom logging field.

For more information about custom logging fields, see "Setting Metabase Properties for Logging" in the Platform SDK.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOGCUSTOM_PROPERTY_HEADER |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4502 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/logging/Custom Logging<br>/LM/logging/Custom Logging/Date<br>/LM/logging/Custom Logging/Time<br>/LM/logging/Custom Logging/Extended Properties<br>/LM/logging/Custom Logging/Extended Properties/*extended_property_name* | IIsCustomLogModule |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogCustomPropertyID Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogCustomPropertyID** property specifies the ABO Metabase identifier of the metabase property that determines if logging is enabled or disabled for the logging field in question. Consult the metabase property reference page, to obtain the ABO Metabase identifier for a particular property.

For instance, the ABO Metabase identifier for LogExtFileFlags Metabase Property is MD_LOGEXT_FIELD_MASK, and the ID number is 4013.

When a logging module is called to log custom information fields, the logging module accesses the property specified by **LogCustomPropertyID**, and checks the individual flag bits, specified by LogCustomPropertyMask Metabase Property, to determine if that particular field is currently enabled or disabled.

For more information about custom logging fields, see "Setting Metabase Properties for Logging" in the Platform SDK.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOGCUSTOM_PROPERTY_ID |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4503 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/logging/Custom Logging<br>/LM/logging/Custom Logging/Date<br>/LM/logging/Custom Logging/Time<br>/LM/logging/Custom Logging/Extended Properties<br>/LM/logging/Custom Logging/Extended Properties/*extended_property_name* | IIsCustomLogModule |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogCustomPropertyMask Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogCustomPropertyMask** property specifies the bitmask that can be used against the metabase property, specified by LogCustomPropertyID Metabase Property, to determine whether the particular logging field is currently enabled or disabled.

When a logging module logs custom information fields, the logging module accesses the property specified by **LogCustomPropertyID**, and checks the individual flag bits, specified by **LogCustomPropertyMask**, to determine if that particular field is currently enabled or disabled.

For more information about custom logging fields, see "Setting Metabase Properties for Logging" in the Platform SDK.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOGCUSTOM_PROPERTY_MASK |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4504 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/logging/Custom Logging<br>/LM/logging/Custom Logging/Date<br>/LM/logging/Custom Logging/Time<br>/LM/logging/Custom Logging/Extended Properties<br>/LM/logging/Custom Logging/Extended Properties/*extended_property_name* | IIsCustomLogModule |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogCustomPropertyName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogCustomPropertyName** property specifies the name of the custom logging information field. Typically, this name is the same as the name used in the logging module user interface.

For more information about custom logging fields, see "Setting Metabase Properties for Logging" in the Platform SDK.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOGCUSTOM_PROPERTY_NAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| StartingNumber | Not applicable |
| EndingNumber | Not applicable |
| ID | 4501 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/logging/Custom Logging<br>/LM/logging/Custom Logging/Date<br>/LM/logging/Custom Logging/Time<br>/LM/logging/Custom Logging/Extended Properties<br>/LM/logging/Custom Logging/Extended Properties/*extended_property_name* | IIsCustomLogModule |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogCustomPropertyNodeID Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogCustomPropertyNodeID** property contains the numeric ID of the logging information field for identification purposes. Each field has a corresponding flag in the **LogExtFileFlags** property which should be set if you want the information logged.

| Information Field | Numeric ID of the Field |
|---|---|
| Bytes Received | 17 |
| Bytes Send | 16 |
| Client IP Address | 5 |
| Cookie | 22 |
| Date | 1 |
| Extended Properties | 3 |
| Host | 20 |
| Method | 11 |
| Protocol Status | 14 |
| Protocol Version | 19 |
| Referer | 23 |
| Server IP Address | 9 |
| Server Name | 8 |
| Server Port | 10 |
| Service Name | 7 |
| Time | 2 |
| Time Taken | 18 |
| URI Query | 13 |
| URI Stem | 12 |
| User Agent | 21 |
| User Name | 6 |

| | |
|---|---|
| Win32 Status | 15 |

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_NOTIFY_ORDER_MASK |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4508 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/logging/Custom Logging<br>/LM/logging/Custom Logging/Date<br>/LM/logging/Custom Logging/Time<br>/LM/logging/Custom Logging/Extended Properties<br>/LM/logging/Custom Logging/Extended Properties/*extended_property_name* | IIsCustomLogModule |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogCustomPropertyServicesString Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogCustomPropertyServicesstring** property specifies to which services the custom logging fields apply.

For more information about custom logging fields, see "Setting Metabase Properties for Logging" in the Platform SDK.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_LOGCUSTOM_SERVICES_string |
| Attributes | INHERIT |
| Default Value | w3svc<br>msftpsvc<br>smtpsvc<br>nntpsvc |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4506 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/logging/Custom Logging<br><br>/LM/logging/Custom Logging/Date<br><br>/LM/logging/Custom Logging/Time<br><br>/LM/logging/Custom Logging/Extended Properties<br><br>/LM/logging/Custom Logging/Extended Properties/*extended_property_name* | IIsCustomLogModule |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogEventOnRecycle Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogEventOnRecycle** property specifies that IIS should log an event when an application pool is recycled. Application pools recycle for a variety of reasons. In order for IIS to log the event, the **LogEventOnRecycle** property must have a bit set corresponding to the reason for the recycle.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APP_POOL_LOG_EVENT_ON_RECYCLE |
| Attributes | INHERIT |
| Default Value | APPPOOLRECYCLEMEMORY |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9037 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Flags**

| Flag Name | AppPoolRecycleTime |
|---|---|
| Description | **Important** This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. For more information, see Recycling Worker Processes with IIS 6.0.<br><br>Setting this flag to **true** causes IIS to log the event (the worker process is recycled after a specified elapsed time). |
| Metabase Bitmask Identifier | MD_APP_POOL_RECYCLE_TIME |
| Decimal Value | 1 |

| | |
|---|---|
| Hexadecimal Value | 0x00000001 |

| **Flag Name** | **AppPoolRecycleRequests** |
|---|---|
| Description | **Important** Important This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. For more information, see Recycling Worker Processes with IIS 6.0.<br><br>Setting this flag to **true** causes IIS to log the event (the worker process is recycled after a specified number of requests). |
| Metabase Bitmask Identifier | MD_APP_POOL_RECYCLE_REQUESTS |
| Decimal Value | 2 |
| Hexadecimal Value | 0x00000002 |

| **Flag Name** | **AppPoolRecycleSchedule** |
|---|---|
| Description | **Important** Important This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. For more information, see Recycling Worker Processes with IIS 6.0.<br><br>Setting this flag to **true** causes IIS to log the event (the worker process is recycled at specified times). |
| Metabase Bitmask Identifier | MD_APP_POOL_RECYCLE_SCHEDULE |
| Decimal Value | 4 |
| Hexadecimal Value | 0x00000004 |

| **Flag Name** | **AppPoolRecycleMemory** |
|---|---|
| Description | **Important** Important This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. For more information, see Recycling Worker Processes with IIS 6.0.<br><br>Setting this flag to **true** causes IIS to log the event (the worker process is recycled once a specified amount of used or virtual memory, expressed in megabytes, is in use). |
| Metabase Bitmask Identifier | MD_APP_POOL_RECYCLE_MEMORY |
| Decimal Value | 8 |
| Hexadecimal Value | 0x00000008 |

| **Flag Name** | **AppPoolRecycleIsapiUnhealthy** |
|---|---|
| Description | **Important** This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. For more information, see Recycling Worker Processes with IIS 6.0.<br><br>Setting this flag to **true** causes IIS to log the event (the worker process is recycled if IIS finds that an ISAPI is unhealthy). |
| Metabase Bitmask Identifier | MD_APP_POOL_RECYCLE_ISAPI_UNHEALTHY |

| | |
|---|---|
| Decimal Value | 16 |
| Hexadecimal Value | 0x00000010 |

| **Flag Name** | **AppPoolRecycleOnDemand** |
|---|---|
| Description | **Important** This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. For more information, see Recycling Worker Processes with IIS 6.0.<br><br>Setting this flag to **true** causes IIS to log the event (the worker process is recycled on demand by an administrator). |
| Metabase Bitmask Identifier | MD_APP_POOL_RECYCLE_ON_DEMAND |
| Decimal Value | 32 |
| Hexadecimal Value | 0x00000020 |

| **Flag Name** | **AppPoolRecycleConfigChange** |
|---|---|
| Description | **Important** Important This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. For more information, see Recycling Worker Processes with IIS 6.0.<br><br>Setting this flag to **true** causes IIS to log the event (the worker process is recycled after configuration changes are made). |
| Metabase Bitmask Identifier | MD_APP_POOL_RECYCLE_CONFIG_CHANGE |
| Decimal Value | 64 |
| Hexadecimal Value | 0x00000040 |

| **Flag Name** | **ppPoolRecyclePrivateMemory** |
|---|---|
| Description | **Important** This feature of IIS 6.0 is available only when IIS is running in worker process isolation mode. For more information, see Recycling Worker Processes with IIS 6.0.<br><br>Setting this flag to **true** causes IIS to log the event (the worker process is recycled when private memory reaches a specified amount). |
| Metabase Bitmask Identifier | MD_APP_POOL_RECYCLE_PRIVATE_MEMORY |
| Decimal Value | 128 |
| Hexadecimal Value | 0x00000080 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogExtFileFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogExtFileFlags** property contains flags that determine which categories of information are written to either the log file or ODBC data source during logging events. The flags are described in the Flags section.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOGEXT_FIELD_MASK |
| Attributes | INHERIT |
| Default Value | logextfiletime \| logextfileclientip \|logextfilemethod \| logextfileuristem \| logextfilehttpstatus |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9013 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Flags**

| Flag Name | LogExtFileDate |
|---|---|
| Description | Log the date. |
| Metabase Bitmask Identifier | MD_EXTLOG_DATE |

| | |
|---|---|
| Decimal Value | 1 |
| Hexadecimal Value | 0x00000001 |

| **Flag Name** | **LogExtFileHost** |
|---|---|
| Description | Log the name of host server. |
| Metabase Bitmask Identifier | MD_EXTLOG_HOST |
| Decimal Value | 1048576 |
| Hexadecimal Value | 0x00100000 |

| **Flag Name** | **LogExtFileTime** |
|---|---|
| Description | Log the time. |
| Metabase Bitmask Identifier | MD_EXTLOG_TIME |
| Decimal Value | 2 |
| Hexadecimal Value | 0x00000002 |

| **Flag Name** | **LogExtFileClientIp** |
|---|---|
| Description | Log the client IP address. |
| Metabase Bitmask Identifier | MD_EXTLOG_CLIENT_IP |
| Decimal Value | 4 |
| Hexadecimal Value | 0x00000004 |

| **Flag Name** | **LogExtFileUserName** |
|---|---|
| Description | Log the username. |
| Metabase Bitmask Identifier | MD_EXTLOG_USERNAME |
| Decimal Value | 8 |
| Hexadecimal Value | 0x00000008 |

| **Flag Name** | **LogExtFileSiteName** |
|---|---|
| Description | Log the site name. |
| Metabase Bitmask Identifier | MD_EXTLOG_SITE_NAME |
| Decimal Value | 16 |
| Hexadecimal Value | 0x00000010 |

| **Flag Name** | **LogExtFileComputerName** |
|---|---|
| Description | Log the local computer name. |
| Metabase Bitmask Identifier | MD_EXTLOG_COMPUTER_NAME |
| Decimal Value | 32 |
| Hexadecimal Value | 0x00000020 |

| Flag Name | LogExtFileServerIp |
|---|---|
| Description | Log the server's own IP address. |
| Metabase Bitmask Identifier | MD_EXTLOG_SERVER_IP |
| Decimal Value | 64 |
| Hexadecimal Value | 0x00000040 |

| Flag Name | LogExtFileMethod |
|---|---|
| Description | Log the protocol method. |
| Metabase Bitmask Identifier | MD_EXTLOG_METHOD |
| Decimal Value | 128 |
| Hexadecimal Value | 0x00000080 |

| Flag Name | LogExtFileUriStem |
|---|---|
| Description | Log the Universal Resource Identifier (URI) stem information. The URI stem usually consists of the actual resource being requested. |
| Metabase Bitmask Identifier | MD_EXTLOG_URI_STEM |
| Decimal Value | 256 |
| Hexadecimal Value | 0x00000100 |

| Flag Name | LogExtFileUriQuery |
|---|---|
| Description | Log the Universal Resource Identifier (URI) query information. The URI query usually consists of parameters passed to the URL by using the URL?Parameters format. |
| Metabase Bitmask Identifier | MD_EXTLOG_URI_QUERY |
| Decimal Value | 512 |
| Hexadecimal Value | 0x00000200 |

| Flag Name | LogExtFileHttpStatus |
|---|---|
| Description | Log the HTTP status. |
| Metabase Bitmask Identifier | MD_EXTLOG_HTTP_STATUS |
| Decimal Value | 1024 |
| Hexadecimal Value | 0x00000400 |

| Flag Name | LogExtFileWin32Status |
|---|---|
| Description | Log the current Microsoft Win32 error status. |
| Metabase Bitmask Identifier | MD_EXTLOG_WIN32_STATUS |
| Decimal Value | 2048 |
| Hexadecimal Value | 0x00000800 |

| **Flag Name** | **LogExtFileBytesSent** |
| --- | --- |
| Description | Log the total bytes sent. |
| Metabase Bitmask Identifier | MD_EXTLOG_BYTES_SENT |
| Decimal Value | 4096 |
| Hexadecimal Value | 0x00001000 |

| **Flag Name** | **LogExtFileBytesRecv** |
| --- | --- |
| Description | Log the total bytes received. |
| Metabase Bitmask Identifier | MD_EXTLOG_BYTES_RECV |
| Decimal Value | 8192 |
| Hexadecimal Value | 0x00002000 |

| **Flag Name** | **LogExtFileTimeTaken** |
| --- | --- |
| Description | Log the total time taken for a request to be completed. |
| Metabase Bitmask Identifier | MD_EXTLOG_TIME_TAKEN |
| Decimal Value | 16384 |
| Hexadecimal Value | 0x00004000 |

| **Flag Name** | **LogExtFileServerPort** |
| --- | --- |
| Description | Log the active server port. |
| Metabase Bitmask Identifier | MD_EXTLOG_SERVER_PORT |
| Decimal Value | 32768 |
| Hexadecimal Value | 0x00008000 |

| **Flag Name** | **LogExtFileUserAgent** |
| --- | --- |
| Description | Log the contents of the user agent field sent by the client. |
| Metabase Bitmask Identifier | MD_EXTLOG_USER_AGENT |
| Decimal Value | 65536 |
| Hexadecimal Value | 0x00010000 |

| **Flag Name** | **LogExtFileCookie** |
| --- | --- |
| Description | Log information from the client cookie. |
| Metabase Bitmask Identifier | MD_EXTLOG_COOKIE |
| Decimal Value | 131072 |
| Hexadecimal Value | 0x00020000 |

| **Flag Name** | **LogExtFileReferer** |
| --- | --- |
| Description | Log the referrer field sent by the client. |

| | |
|---|---|
| Metabase Bitmask Identifier | MD_EXTLOG_REFERER |
| Decimal Value | 262144 |
| Hexadecimal Value | 0x00040000 |

| **Flag Name** | **LogExtFileProtocolVersion** |
|---|---|
| Description | Log the client server protocol version. |
| Metabase Bitmask Identifier | MD_EXTLOG_PROTOCOL_VERSION |
| Decimal Value | 64 |
| Hexadecimal Value | 0x00080000 |

| **Flag Name** | **LogExtFileHttpSubStatus** |
|---|---|
| Description | Log the sub-status code of the HTTP error. For example, for the 500.18 HTTP error, the status code is 500 and the sub-status code is 18. |
| Metabase Bitmask Identifier | na |
| Decimal Value | 2097152 |
| Hexadecimal Value | 0x00200000 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogFileDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogFileDirectory** property specifies the default logging directory, where the log file and logging-related support files are stored.

| **Important** |
| --- |
| The specified log-file directory must be on the local computer. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | EXPANDSZ |
| WMI Data Type | String |
| ADSI Data Type | EXPANDSZ |
| ABO Data Type | EXPANDSZ |
| ABO Metabase Identifier | MD_LOGFILE_DIRECTORY |
| Attributes | INHERIT |
| Default Value | %windir%\system32\logfiles |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4001 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogFileLocaltimeRollover Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogFileLocaltimeRollover** property specifies IIS log file rollover behavior for W3C Extended logging. If the property is set to **false**, the default rollover behavior will be in effect, meaning a new log file will be created based on Coordinated Universal Time (UTC), not local time. If this property is set to **true**, a new log file will be created based on local time, not UTC. UTC was formerly called Greenwich Mean Time, or GMT.

| **Note** |
| --- |
| Regardless of the setting, the time stamp for each W3C Extended Logging log record is UTC-based. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_DEFAULT_EXTLOG_FIELDS |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4015 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogFilePeriod Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogFilePeriod** property specifies how often IIS creates a new log file. This property can be set to the following values as seen in the Bitmask Values section: 1 (daily), 2 (weekly), 3 (monthly), or 4 (hourly). If this property is set to 0, a new file will be created when the maximum size designated in LogFileTruncateSize Metabase Property is reached.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOGFILE_PERIOD |
| Attributes | INHERIT |
| Default Value | 1 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4003 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Values**

| Constant | Value | Description |
|---|---|---|
| MD_LOGFILE_PERIOD_MAXSIZE | 0x00000000 | Create new log file after reaching the maximum size. |

| | | |
|---|---|---|
| MD_LOGFILE_PERIOD_DAILY | 0x00000001 | Create new log file daily. |
| MD_LOGFILE_PERIOD_WEEKLY | 0x00000002 | Create new log file weekly. |
| MD_LOGFILE_PERIOD_MONTHLY | 0x00000003 | Create new log file monthly. |
| MD_LOGFILE_PERIOD_HOURLY | 0x00000004 | Create a new log file hourly. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- LogFileTruncateSize Metabase Property

# LogFileTruncateSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogFileTruncateSize** property specifies the maximum size of the log file (in bytes). A new log file is created if a logging event causes the original log file to exceed the size specified in this property.

| **Note** |
| --- |
| The minimum file size allowed is 131,072 bytes. If this property is set to a value less than 131,072 bytes, the default value is implicitly assumed as 131,072 bytes. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOGFILE_TRUNCATE_SIZE |
| Attributes | INHERIT |
| Default Value | 2000000 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4004 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogInUTF8 Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is controlled by IIS.

| Attribute Name | Attribute Value |
|---|---|
| ADSI/WMI Data Type | BOOL |
| ABO Data Type | DWORD_METADATA |
| Schema Default | **FALSE** |
| Internal Default | **FALSE** |
| Upper Bound | Not specified |
| Lower Bound | Not specified |
| Internal ID | 9206 |
| Friendly ID | MD_GLOBAL_LOG_IN_UTF_8 |
| Property Attributes | NO_ATTRIBUTES |
| User Type | IIS_MD_UT_SERVER |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Flags**

There are no flags for this property.

# LogModuleId Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogModuleId** property specifies the CLSID for the logging module.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOG_PLUGIN_MOD_ID |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4005 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/logging/Microsoft IIS Log File Format<br>/LM/logging/NCSA Common Log File Format<br>/LM/logging/ODBC Logging<br>/LM/logging/W3C Extended Log File Format | IIsLogModule |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- LogModuleUiId Metabase Property

# LogModuleList Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogModuleList** property specifies a string that contains a comma-delimited list that names the available logging modules.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOG_PLUGINS_AVAILABLE |
| Attributes | INHERIT |
| Default Value | NCSA common log file format, IIS log fileformat, W3C extended log file format |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4012 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogModuleUiId Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogModuleUiId** property specifies the CLSID for the logging module's user interface.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOG_PLUGIN_UI_ID |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4006 |

**Configurable Locations**
You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/logging/Microsoft IIS Log File Format /LM/logging/NCSA Common Log File Format /LM/logging/ODBC Logging /LM/logging/W3C Extended Log File Format | IIsLogModule |

**Code Example**
For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- LogModuleId Metabase Property

# LogNonAnonymous Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogNonAnonymous** property specifies whether non-anonymous hits are logged to the event log.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_LOG_NONANONYMOUS |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5008 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogOdbcDataSource Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogOdbcDataSource** property specifies the name of the ODBC data source to which Microsoft Windows writes data during logging events.

| **Important** |
| --- |
| The logging module's CLSID determines whether Microsoft Windows logs to an ODBC data source or to the file system. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOGSQL_DATA_SOURCES |
| Attributes | INHERIT |
| Default Value | httplog |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4007 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogOdbcPassword Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogOdbcPassword** property specifies the ODBC database password to be used when writing information to the database during logging events.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOGSQL_PASSWORD |
| Attributes | INHERIT \| SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4010 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogOdbcTableName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogOdbcTableName** property specifies the name of the ODBC database table to which Microsoft Windows writes information during logging events.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOGSQL_TABLE_NAME |
| Attributes | INHERIT |
| Default Value | internetlog |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4008 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogOdbcUserName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogOdbcUserName** property specifies the ODBC database user name to be used when writing information to the database during logging events.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOGSQL_USER_NAME |
| Attributes | INHERIT |
| Default Value | internetadmin |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4009 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogonMethod Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogonMethod** property contains an integer that specifies the logon method for cleartext logons. Valid settings are:

- 0 for interactive logon (defined in iiscnfg.h as MD_LOGON_INTERACTIVE).
- 1 for batch logon (defined in iiscnfg.h as MD_LOGON_BATCH).
- 2 for network logon (defined in iiscnfg.h as MD_LOGON_NETWORK).
- 3 for cleartext logon (defined in iiscnfg.h as MD_LOGON_NETWORK_CLEARTEXT).

For more information about logon methods, see UNC Authentication in IIS 6.0.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOGON_METHOD |
| Attributes | INHERIT |
| Default Value | 3 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6013 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name*<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |

| | |
|---|---|
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br><br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |
| /LM/W3SVC/AppPools/DefaultAppPool<br><br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogPluginClsid Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogPluginClsid** property specifies the order of precedence for the logging modules.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_LOG_PLUGIN_ORDER |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4011 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# LogType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **LogType** property indicates whether logging to the IIS log files is enabled. IIS log files are stored in *systemroot*\system32\Logfiles\. This property is read-write.

If **LogType** is set to 0, then logging is disabled. If **LogType** is set to 1, then logging is enabled. To disable logging for child nodes under a parent node for whom **LogType** is set to 1, set the **DontLog** property to **TRUE** on the child nodes.

ASP and ODBC errors do not get logged to the IIS log files.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_LOG_TYPE |
| Attributes | INHERIT |
| Default Value | 1 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 4000 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- DontLog Metabase Property

# MajorIIsVersionNumber Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MajorIIsVersionNumber** property specifies the major version number of IIS.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SERVER_VERSION_MAJOR |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 1101 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/Info | IIsWebInfo |
| /LM/MSFTPSVC/Info | IIsFtpInfo |
| /LM/NNTPSVC/Info | IIsNntpInfo |
| /LM/SMTPSVC/Info | IIsSmtpInfo |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MasqueradeDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MasqueradeDomain** property specifies the domain name used for non-delivery reports.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_MASQUERADE_NAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36912 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxBandwidth Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxBandwidth** property specifies the maximum network bandwidth used for IIS. You can use this setting to help prevent overloading the network with IIS activity. This is not an inheritable property, but the value set at the machine level is globally available to all server instances. **MaxBandWidth** can be set individually, so that specific server instances are used instead of the global value, and can exceed the global setting established at the machine level.

The World Wide Web Publishing Service (WWW service) must be restarted before changes to this property take effect.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_BANDWIDTH |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1000 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM | IIsComputer |
| /LM/W3SVC | IIsWebService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxBandwidthBlocked Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxBandwidthBlocked** property specifies the maximum number of asynchronous requests to be queued in the outstanding idle queue when the maximum bandwidth of the server is exceeded. The value range for this property is 0 to 4294967295 (unlimited).

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| **Important** |
| --- |
| This is not an inheritable property and cannot be set at child nodes. **MaxBandwidthBlocked** is set at the machine level, and is globally available to all server instances. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_BANDWIDTH_BLOCKED |
| Attributes | NO_ATTRIBUTES |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1003 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM | IIsComputer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxBatchedMessages Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxBatchedMessages** property limits the number of messages per connection.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 36871 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxClientsMessage Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxClientsMessage** property contains a string that is displayed to an FTP client when there are too many FTP clients accessing the server. You can configure the maximum number of FTP clients by using the **MaxConnections** property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_MAX_CLIENTS_MESSAGE |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5003 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- GreetingMessage Metabase Property
- BannerMessage Metabase Property
- ExitMessage Metabase Property

# MaxConnections Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxConnections** property specifies the maximum number of simultaneous connections to a server. The valid range is 0 to 4294967295 (unlimited). The **MaxClientsMessage** property can be used to send a message to clients when this value has been exceeded.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, Microsoft Visual Basic Scripting Edition (VBScript) represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_CONNECTIONS |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1014 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- MaxClientsMessage Metabase Property

# MaxDirChangeIOSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxDirChangeIOSize** property determines the size of the pickup directory buffer. This property is used for debugging purposes.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36878 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxEndpointConnections Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxEndpointConnections** property specifies the maximum number of "listen" sockets that will be aggregated on a network endpoint. For example, if this value is set to 15, then a maximum of 15 total connections can be made to a single port, even if more than one domain is bound to the port.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_ENDPOINT_CONNECTIONS |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1024 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxErrorFiles Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxErrorFiles** property specifies the maximum number of error files. When a syntax error occurs as a result of an edit to Metabase.XML, an error file is generated in systemroot\system32\inetsrv\history. If more than **MaxErrorFiles** files need to be created, the oldest files are deleted to ensure that this limit is not exceeded.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_ERROR_FILES |
| Attributes | NO_ATTRIBUTES |
| Default Value | 10 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9988 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM | IIsComputer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxGlobalBandwidth Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxGlobalBandwidth** property specifies the maximum total bandwidth for the server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_GLOBAL_BANDWIDTH |
| Attributes | NO_ATTRIBUTES |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9201 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxHistoryFiles Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxHistoryFiles** property specifies the maximum number of paired history files that are saved in the history folder. A pair of history files consists of a MetaBase.xml and a MBSchema.xml file, named with the same major and minor version number, and stored in the history folder.

| **Important** |
| --- |
| If the EnableEditWhileRunning Metabase Property property is set to **true**, the **MaxHistoryFiles** property should not be configured to a value less than 10. Setting the value of the **MaxHistoryFiles** property to a value lower than 10 may cause edit-while-running to fail, whereas setting the value too high may have an impact on disk capacity. |

Before increasing the value of the **MaxHistoryFiles** property above the value of 10, ensure that you have adequate hard disk space on the drive volume that the history folder is located on. You can estimate the amount of hard disk space required by multiplying the average combined file size of the MetaBase.xml and MBSchema.xml files by the number of versions that you want to store in the history folder. For information about configuring metabase history, see Configuring the Metabase History Feature in IIS 6.0.

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ROOT_MAX_HISTORY_FILES |
| Attributes | NO_ATTRIBUTES |
| Default Value | 10 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9995 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM | IIsComputer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- EnableHistory Metabase Property

# MaxMailObjects Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxMailObjects** property specifies the number of inbound connections allowed at one time. This key is used for debugging purposes.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36873 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxMessageSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxMessageSize** property specifies the maximum message size the server accepts. **MaxMessageSize** is also the value sent by the SMTP service in response to the EHLO command.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_MSG_SIZE |
| Attributes | INHERIT |
| Default Value | 33554432 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36891 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxOutConnections Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxOutConnections** property specifies the maximum number of concurrent outbound connections.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_OUTBOUND_CONNECTION |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36894 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxOutConnectionsPerDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxOutConnectionsPerDomain**property specifies the maximum number of concurrent outbound connections per outbound domain.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_OUT_CONN_PER_DOMAIN |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36925 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxProcesses Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxProcesses** property determines the maximum number of worker processes an application pool allows to service requests for an application pool.

| **Note** |
| --- |
| A Web Garden is an application pool with a **MaxProcesses** setting higher than 1. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_MAX_PROCESS_COUNT |
| Attributes | INHERIT |
| Default Value | 1 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9003 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxRecipients Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxRecipients** property specifies the maximum number of recipients for a given message for a given virtual server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 100 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36895 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxRequestEntityAllowed Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxRequestEntityAllowed** property specifies the maximum number of bytes allowed in the entity body of a request. If a Content-Length header is present and specifies an amount of data greater than the value of **MaxRequestEntityAllowed**, IIS sends a 403 error response.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MAX_REQUEST_ENTITY_ALLOWED |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6051 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxSearchResults Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxSearchResults** property specifies the maximum number of results for a search command.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45181 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxSessionSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxSessionSize** property specifies the maximum amount of data per SMTP session for a given virtual server.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 268435456 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36892 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxSmtpErrors Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxSmtpErrors** property specifies the maximum number of SMTP errors that can occur in one session before the server drops the connection.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36887 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MaxSmtpLogonErrors Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MaxSmtpLogonErrors** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36945 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| | |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MD_0 Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MD_0** property is used by IIS.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | hidden |
| User Type | unknown_usertype |
| ID | 0 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| | |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MD_1001 Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MD_1001** property is used by IIS.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MEMORY_CACHE_SIZE |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | hidden |
| User Type | IIS_MD_UT_SERVER |
| ID | 1001 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM | IIsComputer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MD_APP_LAST_OUTPROC_PID Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MD_APP_LAST_OUTPROC_PID** property contains a value that is stored in an application package ID to denote that the application package has been moved from out-of-process to in-process. In some cases, when an application is moved from out-of-process to in-process, the application's corresponding package may not be deleted by IIS. In this case, IIS stores the application package ID with the value of the **MD_APP_LAST_OUTPROC_PID** property so that if the application is later moved out-of-process, IIS can use the existing package in MTS/COM+.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APP_LAST_OUTPROC_PID |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | hidden |
| User Type | iis_md_ut_wam |
| ID | 2108 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MD_ETAG_CHANGENUMBER Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS uses the value specified in the **MD_ETAG_CHANGENUMBER** property to generate a portion of the entity tags (ETags) that are sent to clients in server response headers. By setting a matching value for **MD_ETAG_CHANGENUMBER** on all the Web servers in a Network Load Balanced (NLB) cluster, the value of the entity tag (ETag) will become the same on all the servers, preventing clients from having to re-download content that already exists in the client cache when a conditional request is sent to a different Web server in the cluster.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ETAG_CHANGE_NUMBER |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | hidden |
| User Type | IIS_MD_UT_SERVER |
| ID | 2039 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MD_IISADMIN_EXTENSIONS Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MD_IISADMIN_EXTENSIONS** property contains a list of CLSIDs for all registered administration extensions installed on the server.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_SECURE_PORT |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | hidden |
| User Type | IIS_MD_UT_SERVER |
| ID | 1028 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| | |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MD_ISM_ACCESS_CHECK Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MD_ISM_ACCESS_CHECK** property is used by IIS.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ISM_ACCESS_CHECK |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | hidden |
| User Type | IIS_MD_UT_FILE |
| ID | 6269 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MD_METADATA_ID_REGISTRATION Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MD_METADATA_ID_REGISTRATION** property is used by IIS.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_METADATA_ID_REGISTRATION |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | hidden |
| User Type | IIS_MD_UT_SERVER |
| ID | 1030 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| | |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MD_SERVER_CAPABILITIES Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MD_SERVER_CAPABILITIES** property is used by IIS.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SERVER_CAPABILITIES |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | hidden |
| User Type | IIS_MD_UT_SERVER |
| ID | 1103 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/Info | IIsWebInfo |
| /LM/MSFTPSVC/Info | IIsFtpInfo |
| /LM/NNTPSVC/Info | IIsNntpInfo |
| /LM/SMTPSVC/Info | IIsSmtpInfo |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MD_SERVER_PLATFORM Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MD_SERVER_PLATFORM** property is used by IIS.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SERVER_PLATFORM |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | hidden |
| User Type | IIS_MD_UT_SERVER |
| ID | 1100 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/Info | IIsWebInfo |
| /LM/MSFTPSVC/Info | IIsFtpInfo |
| /LM/NNTPSVC/Info | IIsNntpInfo |
| /LM/SMTPSVC/Info | IIsSmtpInfo |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MimeMap Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MimeMap** property provides a list of the file name extensions for Multipurpose Internet Mail Extensions (MIME) mappings. The **MimeMap** key at the local machine level of the metabase establishes the default list, which can be overridden at the World Wide Web Publishing Service (WWW service) and Web server level keys. For more information about managing this property, see the **IIsMimeMap** object in the IIS ADSI Provider reference section of the IIS Software Development Kit (SDK).

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | MIMEMAP |
| WMI Data Type | MIMEMAP |
| ADSI Data Type | MIMEMAP |
| ABO Data Type | MIMEMAP |
| ABO Metabase Identifier | MD_MIME_MAP |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6015 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT<br>/LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM | IIsComputer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/MimeMap | IIsMimeMap |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MinFileBytesPerSec Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MinFileBytesPerSec** property is used to configure the minimum throughput rate that HTTP.sys enforces when sending data from the client to the server, and back from the server to the client. The **MinFileBytesPerSec** property prevents malicious or malfunctioning software clients from consuming resources by holding a connection open with minimal data. If the throughput rate is below the **MinFileBytesPerSec** setting, the connection is terminated.

In IIS 6.0, use this property in conjunction with the **ConnectionTimeout** property, and for Web service nodes, use the **HeaderWaitTimeout** property instead of the **ServerListenTimeout** property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_MIN_FILE_BYTES_PER_SEC |
| Attributes | INHERIT |
| Default Value | 240 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9205 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MinorIIsVersionNumber Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MinorIIsVersionNumber** property specifies the minor version number of IIS.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 1102 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/Info | IIsWebInfo |
| /LM/MSFTPSVC/Info | IIsFtpInfo |
| /LM/NNTPSVC/Info | IIsNntpInfo |
| /LM/SMTPSVC/Info | IIsSmtpInfo |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ModeratorFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ModeratorFile** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45160 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# MSDOSDirOutput Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **MSDOSDirOutput** property specifies the style of directory output for a list operation request from an FTP client. If the value is **true**, the format is in MS-DOS style, if **false**, it is produced in UNIX style.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_MSDOS_DIR_OUTPUT |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5004 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NameResolutionType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NameResolutionType** property tells SMTP how to resolve a host name. The default is 1. Valid values are:

1= UNCACHEDNS

2= GETHOSTBYNAME

3= DNS_GETHOSTBYNAME

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 1 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36885 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NewsCrawlerTime Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NewsCrawlerTime** property specifies the time between expiration runs, in minutes.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45069 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NewsDropDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NewsDropDirectory** property specifies the directory where NNTP puts mail that needs to go out.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45176 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NewsFailedPickupDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NewsFailedPickupDirectory** property specifies the directory where NNTP puts invalid posts from the pickup directory.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45174 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NewPickupDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NewsPickupDirectory** property specifies the directory where NNTP looks to pick up posts.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45173 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NntpClearTextProvider Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NntpClearTextProvider** property specifies the clear text authentication provider.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45179 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NntpCommandLogMask Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NntpCommandLogMask** property specifies the path to the moderatr.txt file, currently unused.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45067 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NntpOrganization Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NntpOrganization** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45171 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NntpServiceVersion Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NntpServiceVersion** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45175 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NntpUucpName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **NntpUucpName** property contains the NNTP hub name (UUCP name).

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45170 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NotDeletable Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is not valid in IIS 6.0. In earlier versions of IIS, the **NotDeletable** property indicates whether a server can be deleted by using IIS Manager. If this property is set to **true**, an attempt to delete the affected server instance using IIS Manager results in an error.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_NOT_DELETABLE |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2116 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# NTAuthenticationProviders Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

When Integrated Windows authentication is enabled, the **NTAuthenticationProviders** property contains a comma-delimited list of Windows authentication providers that you can attempt to use, including Integrated Windows authentication (also known as NTLM) and Kerberos.

| **Important** |
| --- |
| Integrated Windows authentication is disabled by default if you install Windows Server 2003 Service Pack 1 (SP1) as part of a slipstream installation of a Windows Server 2003 operating system. If you install Windows Server 2003 SP1 by itself as an upgrade to a Windows Server 2003 operating system, the setting for Integrated Windows authentication is unchanged from its Windows Server 2003 setting. Integrated Windows authentication is enabled by default for Windows Server 2003 operating systems. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_NTAUTHENTICATION_PROVIDERS |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6032 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC /LM/W3SVC/*n* /LM/*W3SVC/n/ROOT /LM/W3SVC/n/ROOT/MyApp* | IIsWebService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# OrphanActionExe Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **OrphanActionExe** property specifies an executable to run when the World Wide Web Publishing Service (WWW service) orphans a worker process. You can use the OrphanActionParams Metabase Property to send parameters to the executable.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APPPOOL_ORPHAN_ACTION_EXE |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9031 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# OrphanActionParams Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **OrphanActionParams** property specifies command-line parameters for the executable specified by the OrphanActionExe Metabase Property. To specify the process ID of the orphaned process, use %1%.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_APPPOOL_ORPHAN_ACTION_PARAMS |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9032 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- OrphanActionExe Metabase Property

# OrphanWorkerProcess Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **OrphanWorkerProcess** property, when set to true, notifies the World Wide Web Publishing Service (WWW Service) not to terminate a worker process that fails to respond to pings, but to instead orphan the worker process in the application pool if the worker process suffers fatal errors.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_APPPOOL_ORPHAN_PROCESSES_FOR_DEBUGGING |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9009 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PassivePortRange Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PassivePortRange** property specifies the range of data ports to be used by the FTP service in response to PASV commands.

PASV FTP requires the server to open a data port for the client to make a second connection. This is a separate connection than the typical port 21 that is used for the control channel. The second connection is used when data files are transferred back to the client. By configuring the port range, you can write firewall and router rules to allow external clients access only to the ports they need and reduce the attack surface available to malicious users. In other words, if you have applications other than FTP that are using the default port range of 1025-5000, and do not want to expose these ports through your firewall in order to enable PASV FTP, you can use this value to change the range that you must open through your firewall. If this value is not specified, or is set to an empty string, the default value of 1025-5000, as specified by Winsock, is used. If this property is specified, the valid range that FTP will validate is from 5001 to 65535 (see **StartingNumber** and **EndingNumber** below), and may be a range or a single number.

| **Important** |
| --- |
| This property can be set only at the service level. In order to make the changes effective, the service must be restarted. If the value is invalid, the service will invalidate it and will not restart. |

For more information, see Modes and Data Transmission.

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_PASSIVE_PORT_RANGE |
| Attributes | INHERIT |
| Default Value | 1025-5000 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5016 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PassportRequireADMapping Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PassportRequireADMapping** property specifies how IIS handles Microsoft .NET Passport authentication and Active Directory mapping.

By default, IIS 6.0 with .NET Passport authentication enabled tries to map the .NET Passport user to an account in Active Directory. This default behavior can create performance overhead for sites that do not accept .NET Passport Active Directory mapping. The mapping behavior can be controlled by the **PassportRequireADMapping** flag. The flag values are:

- **PassportRequireADMapping** set to 0: IIS does not attempt Active Directory mapping. The request is handled as Anonymous User.
- **PassportRequireADMapping** set to 1 (default setting): In this situation, the worker process must have TCB privileges (meaning, the worker process acts as part of the operating system). IIS then attempts to map to an Active Directory account (called LsaLogonUser). If this attempt fails, the request is handled as Anonymous User.
- **PassportRequireADMapping** set to 2: IIS enforces a mapping from the .NET Passport account to an Active Directory account before returning the requested Web page. If the mapping fails, IIS returns a 401 error.

For more information, see Setting Up .NET Passport in IIS 6.0. Also, the AuthFlags Metabase Property contains the available settings for Windows authentication schemes and the file access authentication flags.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_PASSPORT_REQUIRE_AD_MAPPING |
| Attributes | INHERIT |
| Default Value | 1 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6052 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |

| | |
|---|---|
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PasswordCacheTTL Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PasswordCacheTTL** property is used for password types of security authentication schemes. The value specifies the amount of time (in seconds) that an expired password will be held in the memory cache.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ADV_CACHE_TTL |
| Attributes | INHERIT |
| Default Value | 600 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2064 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.RapidFailProtectionInterval Metabase Property

**Related Information**

- PasswordExpirePrenotifyDays Metabase Property

# PasswordChangeFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PasswordChangeFlags** property contains flags, as seen in the Flags section, that control password expiration and password change processing between the server and client. A value of 0, which is the default, indicates that an SSL connection is required, 1 indicates that changing is allowed on nonsecure ports, 2 indicates that changing is disabled, and 4 indicates that password expiration notification is disabled.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_AUTH_CHANGE_FLAGS |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2068 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Flags**

| Flag Name | AuthChangeUnsecure |
|---|---|
| Description | Password changing allowed on nonsecure ports. |
| Metabase Bitmask Identifier | MD_AUTH_CHANGE_UNSECURE |
| Decimal Value | 1 |
| Hexadecimal Value | 0x00000001 |

| Flag Name | AuthChangeDisable |
|---|---|
| Description | Password changing disabled. |

| | |
|---|---|
| Metabase Bitmask Identifier | MD_AUTH_CHANGE_DISABLE |
| Decimal Value | 2 |
| Hexadecimal Value | 0x00000002 |

| **Flag Name** | **AuthAdvNotifyDisable** |
|---|---|
| Description | Advance notification of password expiration disabled. |
| Metabase Bitmask Identifier | MD_AUTH_ADVNOTIFY_DISABLE |
| Decimal Value | 4 |
| Hexadecimal Value | 0x00000004 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.RapidFailProtectionInterval Metabase Property

**Related Information**

PasswordExpirePrenotifyDays Metabase Property

# PasswordExpirePrenotifyDays Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PasswordExpirePrenotifyDays** property specifies the number of days remaining before the client's password expires, and indicates when a password prenotification message is sent.

| **Note** |
| --- |
| This property is dependant on the PasswordChangeFlags Metabase Property being enabled. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_ADV_NOTIFY_PWD_EXP_IN_DAYS |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2063 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Path Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **Path** property specifies the physical path associated with a virtual directory.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_VR_PATH |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 3001 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/MSFTPSVC/*n*/ROOT /LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |
| /LM/SMTPSVC/*n*/ROOT /LM/SMTPSVC/*n*/ROOT/*virtual_directory_name* | IIsSmtpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PeriodicRestartMemory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PeriodicRestartMemory** property specifies the amount of virtual memory (in KB) that a worker process can use before the worker process recycles. The maximum value supported for this property is 4,294,967 KB.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_PERIODIC_RESTART_MEMORY |
| Attributes | INHERIT |
| Default Value | 512000 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9024 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PeriodicRestartPrivateMemory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PeriodicRestartPrivateMemory** property specifies the amount of private memory (in KB) that a worker process can use before the worker process recycles. The maximum value supported for this property is 4,294,967 KB.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_PERIODIC_RESTART_PRIVATE_MEMORY |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9038 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PeriodicRestartRequests Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PeriodicRestartRequests** property indicates the number of requests the OOP application should process, after which it is recycled.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_PERIODIC_RESTART_REQUEST_COUNT |
| Attributes | INHERIT |
| Default Value | 10000 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9002 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PeriodicRestartSchedule Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PeriodicRestartSchedule** property specifies the time (in 24 hour format) that the application will be rotated. Each time is in local time and is specified in the following format:

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_APPPOOL_PERIODIC_RESTART_SCHEDULE |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9020 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PeriodicRestartTime Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PeriodicRestartTime** property specifies the period of time, in minutes, after which IIS rotates an isolated OOP application. Setting the value of this property to 0 disables the property. The maximum supported value for this property is 71,582.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_PERIODIC_RESTART_TIME |
| Attributes | INHERIT |
| Default Value | 60 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9001 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- IdleTimeout Metabase Property

# PickupDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PickupDirectory** property specifies the path to the directory that is used to pick up inbound mail.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_MAIL_PICKUP_DIR |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36880 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PingingEnabled Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PingingEnabled** property specifies whether the World Wide Web Publishing Service (WWW Service) should periodically monitor the health of a worker process. Setting the value of this property to true indicates to the WWW service to monitor the worker processes to ensure that they are running and healthy.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_APPPOOL_PINGING_ENABLED |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9004 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PingInterval Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PingInterval** property specifies the period of time (in seconds) between health-monitoring pings that the World Wide Web Publishing Service (WWW Service) sends to a worker process.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_PING_INTERVAL |
| Attributes | INHERIT |
| Default Value | 240 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9013 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PingResponseTime Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PingResponseTime** property specifies the amount of time (in seconds) that a worker process is given to respond to a health monitoring ping. After the time limit is exceeded, the World Wide Web Publishing Service (WWW Service) terminates the worker process.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_PING_RESPONSE_TIMELIMIT |
| Attributes | INHERIT |
| Default Value | 60 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9014 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PoolIdcTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PoolIdcTimeout** property specifies the timeout value (in seconds) for Internet database connection pooling. A value of 0 (zero) specifies that no pooling is implemented.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_POOL_IDC_TIMEOUT |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6037 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3ClearTextProvider Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41182 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3DefaultDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41176 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3DsAccount Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41190 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3DsBindType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41186 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3DsDataDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41184 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3DsDefaultMailRoot Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41185 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3DsHost Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41188 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3DsNamingContext Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41189 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3DsPassword Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT \| SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41191 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3DsSchemaType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41187 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3ExpireDelay Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 1440 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41178 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3ExpireMail Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41181 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3ExpireStart Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41179 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3MailExpirationTime Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 40961 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3RoutingDll Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41177 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3RoutingSources Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | iis_md_ut_wam |
| ID | 41174 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3RoutingTableType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | ldap |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41183 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Pop3ServiceVersion Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | IIS_MD_UT_SERVER |
| ID | 41460 |

**Configurable Locations**

Currently, this property is not configurable at any location in the IIS metabase.

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PostmasterEmail Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PostmasterEmail** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 36917 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.RapidFailProtectionInterval Metabase Property

# PostmasterName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PostmasterName** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 36916 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# PrettyNamesFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **PrettyNamesFile** property contains the path to the prettynm.txt file, which is currently unused.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45178 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ProcessNTCRIfLoggedOn Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ProcessNTCRIfLoggedOn** property enables processing of Integrated Windows authentication (NTLM), even if a user has already logged on using an alternate authentication scheme.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_PROCESS_NTCR_IF_LOGGED_ON |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2070 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# QueueDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **QueueDirectory** property designates the directory where mail queues are held.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36875 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties

# RapidFailProtection Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Setting the **RapidFailProtection** property to **true** instructs the World Wide Web Publishing Service (WWW service) to put all applications in an application pool out of service if the number of worker process crashes has reached the maximum specified by the RapidFailProtectionMaxCrashes Metabase Property, within the number of minutes specified by the RapidFailProtectionInterval Metabase Property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_APPPOOL_RAPID_FAIL_PROTECTION_ENABLED |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9006 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RapidFailProtectionInterval Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RapidFailProtectionInterval** property specifies the number of minutes before the failure count for a process is reset. See RapidFailProtection Metabase Property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_RAPID_FAIL_PROTECTION_INTERVAL |
| Attributes | INHERIT |
| Default Value | 5 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9029 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RapidFailProtectionMaxCrashes Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RapidFailProtectionMaxCrashes** property specifies the maximum number of failures allowed within the number of minutes specified by the RapidFailProtectionInterval Metabase Property. See RapidFailProtection Metabase Property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_RAPID_FAIL_PROTECTION_MAX_CRASHES |
| Attributes | INHERIT |
| Default Value | 5 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9030 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# Realm Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **Realm** property specifies the realm that IIS should use to authenticate a client that is attempting to access a resource secured with Digest authentication or Advanced Digest authentication.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_REALM |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6001 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RedirectHeaders Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RedirectHeaders** property specifies additional redirection headers. Each string in the string array represents a header and its value.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_REDIRECT_HEADERS |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6044 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RelayForAuth Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RelayForAuth** property corresponds to the **Allow all computers which successfully authenticate to relay, regardless of the list above** check box in the **Relay Restrictions** dialog box. This value is set to -1 if the check box is selected (that is, allow relay to all those who have authenticated) and is set to 0 if the check box is not selected.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36973 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/Domain | IIsSmtpDomain |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RelayIpList Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RelayIpList** property is an octet string containing a list of IPs allowed or denied for relay listed in the Relay Restrictions dialog box.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | IPSECLIST |
| WMI Data Type | IPSECLIST |
| ADSI Data Type | IPSECLIST |
| ABO Data Type | IPSECLIST |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | inherit \| reference |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 36970 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/Domain | IIsSmtpDomain |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RemoteRetryAttempts Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RemoteRetryAttempts** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 48 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36898 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RemoteRetryInterval Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RemoteRetryInterval** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 60 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36899 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RemoteSmtpPort Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RemoteSmtpPort** property determines which port is used when a virtual server connects to a remote server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 25 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36914 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RemoteSmtpSecurePort Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RemoteSmtpSecurePort** property is obsolete.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 465 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36929 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RemoteTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RemoteTimeout** property sets the timeout for outbound connections, which correspond to the timeout value in the **Outbound Connections** dialog box.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36921 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RestrictionListCustomDesc Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is controlled by IIS.

| Attribute Name | Attribute Value |
|---|---|
| ADSI/WMI Data Type | MULTISZ |
| ABO Data Type | MULTISZ_METADATA |
| Schema Default | **null** |
| Internal Default | Not specified |
| Upper Bound | Not specified |
| Lower Bound | Not specified |
| Internal ID | 2165 |
| Friendly ID | No friendly name. Use the Internal ID attribute listed above. |
| Property Attributes | NO_ATTRIBUTES |
| User Type | IIS_MD_UT_SERVER |

**Configurable Locations**
This property is not configurable at any location in the metabase.

**Flags**
There are no flags for this property.

# RevocationFreshnessTime Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If the value of the CertCheckMode Metabase Property is set to **MD_CERT_CHECK_REVOCATION_FRESHNESS_TIME**, the client's CRL is replaced by the CRL on the CA, even if the CRL that is cached on the client is valid. The value of this metabase property, **RevocationFreshnessTime**, in seconds, determines the frequency of this action.

The *freshness time* for a CRL is the period of time for which it is valid. If you set this property to a very short period of time, IIS makes many requests for an updated CRL, which can cause network delays. Setting a longer period of time for the **RevocationFreshnessTime** metabase property is recommended.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_REVOCATION_FRESHNESS_TIME |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2161 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- For more information about using SSL and certificates with IIS 6.0, see IIS 6.0 Encryption.

For more information about certificate revocation and CRLs, see Revoking Certificates and Publishing CRLs.

# RevocationURLRetrievalTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If the value of the CertCheckMode Metabase Property is set to **MD_CERT_CHECK_REVOCATION_FRESHNESS_TIME**, the client's CRL is replaced by the CRL on the CA, even if the CRL that is cached on the client is valid. The default timeout interval for downloading the CRL is 0 unless you use this metabase property, **RevocationURLRetrievalTimeout**, to specify a different timeout interval (in milliseconds). If the CRL is a large file, a higher value for this metabase property allows more time for downloading the file before a timeout occurs.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_REVOCATION_URL_RETRIEVAL_TIMEOUT |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2162 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- For more information about using SSL and certificates with IIS 6.0, see Encryption in IIS 6.0.
- For more information about certificate revocation and CRLs, see Revoking Certificates and Publishing CRLs.

# RouteAction Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RouteAction** property flags the various SMTP options. This gives properties to a connector.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36946 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/Domain | IIsSmtpDomain |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RouteActionString Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RouteActionstring** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36947 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/Domain | IIsSmtpDomain |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RoutePassword Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For outbound connections, the **RoutePassword** property sets the password used to connect to a remote server (this value is listed for a given connector, for example /smtpsvc/1/domain/username.com).

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT \| SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36949 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/Domain | IIsSmtpDomain |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RouteUserName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

For outbound connections, the **RouteUserName** property sets the account name used to connect to a remote server (this value is listed for a given connector, for example /smtpsvc/1/domain/username.com).

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36948 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/Domain | IIsSmtpDomain |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RoutingDll Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RoutingDll** property is obsolete.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36883 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# RoutingSources Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **RoutingSources** property is obsolete.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | iis_md_ut_wam |
| ID | 36910 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SaslLogonDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SaslLogonDomain** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36940 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ScriptMaps Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ScriptMaps** property specifies the file name extensions of applications used for script processor mappings. The list string is written in the following format:

```
"Extension, ScriptProcessor, Flags, IncludedVerbs"
```

Where Extension is the file name extension, such as .htm, *ScriptProcessor* is the full path to the DLL, *Flags* is the integer value corresponding to the requested behavior described in the Bitmask Table, and *IncludedVerbs* is a list of the verbs that a particular ISAPI DLL processes.

For example, to specify the file extension for the ISAPI "Test.dll" with a file extension of ".htm", you might provide the following list (string):

```
".htm,C:\Windows40\System32\Inetsrv\Test.dll,5,GET, HEAD, POST"
```

Note IIS allows an "*" in place of the file name extension, which effectively routes all requests though the designated *ScriptProcessor*. This feature is useful when applying a custom ISAPI extension that will handle requests from all file types.

In IIS version 4.0 and earlier, the syntax was to list excluded verbs rather than included verbs. In version 5.0 and later, if no verbs are listed, a value of "all verbs" is assumed. You should list the verbs you want your ISAPI filter or extension to handle.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_SCRIPT_MAPS |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6014 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

| | |
|---|---|
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Values**

| Constant | Value | Description |
|---|---|---|
| MD_SCRIPTMAPFLAG_SCRIPT | 0x00000001 | Scripts are allowed to run in directories given Script permission. If this value is not set, then the script can only be run in directories that are flagged for Execute permission. |
| MD_SCRIPTMAPFLAG_CHECK_PATH_INFO | 0x00000004 | The server attempts to access the PATH_INFO portion of the URL, as a file, before starting the scripting engine. If the file can't be opened, or doesn't exist, an error is returned to the client. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SecureBindings Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SecureBindings** property specifies a string that is used by IIS to determine which secure network endpoints are used by the server instance. The string format is IP:Port.

| **Note** |
| --- |
| The IP component of each string is optional. If it is not provided, IIS assumes that any IP address is allowable (wildcard). |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_SECURE_BINDINGS |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 2021 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.RapidFailProtectionInterval Metabase Property

**Related Information**

- ServerBindings Metabase Property

# SecuritySetupRequired Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

| **Important** |
| --- |
| The **SecuritySetupRequired** property is not supported in IIS 6.0. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 2166 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| | |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SendBadTo Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SendBadTo** property sets the alias to which bad mail is sent. The default is "Postmaster."

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36923 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SendNdrTo Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SendNdrTo** property corresponds to "Send copy of Non-Delivery Report to" on the **Messages** tab.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36922 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ServerAutoStart Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ServerAutoStart** property indicates if the server instance should start automatically when the service is started.

This property is reset automatically when a server instance is stopped or restarted, to maintain state across service restarts. For example, if a server is stopped, the value of this property is set to **false** so that if the service is stopped and restarted, the server will remain stopped. Likewise, if a server is started, this property is set to **true**, allowing the server to remain running whenever the service is running.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_SERVER_AUTOSTART |
| Attributes | NO_ATTRIBUTES |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1017 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ServerBindings Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ServerBindings** property specifies a string that IIS uses to determine which network endpoints are used by the server instance. The string format is IP:Port:Hostname.

| **Note** |
| --- |
| Both the IP and Hostname parameters of the string are optional. Any unspecified parameters default to an all-inclusive wildcard. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_SERVER_BINDINGS |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1023 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- SecureBindings Metabase Property

# ServerCommand Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ServerCommand** property can be set to start, stop, continue, or pause a site. The commands, and the corresponding codes, are listed in the Bitmask Values table.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SERVER_COMMAND |
| Attributes | VOLATILE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1012 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Bitmask Values**

| Value | Description |
|---|---|
| 1 | Start site. |
| 2 | Stop site. |
| 4 | Pause site. |
| 8 | Continue site. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ServerComment Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ServerComment** property specifies a comment for the server instance.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_SERVER_COMMENT |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1015 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ServerConfigFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ServerConfigFlags** property contains flags, as seen in the Flags section, that specify the server configuration established at startup.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SERVER_CONFIGURATION_INFO |
| Attributes | NO_ATTRIBUTES |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1027 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/Info | IIsWebInfo |

**Flags**

| Flag Name | ServerConfigSSL40 |
|---|---|
| Description | A value of **true** indicates that 40-bit Secure Sockets Layer (SSL) processing is installed on the server. |
| Metabase Bitmask Identifier | MD_SERVER_CONFIG_SSL_40 |
| Decimal Value | 1 |
| Hexadecimal Value | 0x00000001 |

| Flag Name | ServerConfigSSL128 |
|---|---|
| Description | A value of **true** indicates that 128-bit Secure Sockets Layer (SSL) processing is installed on the server. |

| | |
|---|---|
| Metabase Bitmask Identifier | MD_SERVER_CONFIG_SSL_128 |
| Decimal Value | 2 |
| Hexadecimal Value | 0x00000002 |

| **Flag Name** | **ServerConfigSSLAllowEncrypt** |
|---|---|
| Description | A value of **true** indicates that the locality allows encryption. |
| Metabase Bitmask Identifier | MD_SERVER_CONFIG_ALLOW_ENCRYPT |
| Decimal Value | 4 |
| Hexadecimal Value | 0x00000004 |

| **Flag Name** | **ServerConfigAutoPWSync** |
|---|---|
| Description | A value of **true** indicates that the server can perform automatic anonymous user password synchronization between the Web authentication manager (WAM) and the Microsoft Windows authentication manager. |
| Metabase Bitmask Identifier | MD_SERVER_CONFIG_AUTO_PW_SYNC |
| Decimal Value | 8 |
| Hexadecimal Value | 0x00000008 |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ServerID Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| ADSI/WMI Data Type | STRING |
| ABO Data Type | STRING_METADATA |
| Schema Default | **null** |
| Internal Default | Not specified |
| Upper Bound | Not specified |
| Lower Bound | Not specified |
| Internal ID | 130010 |
| Friendly ID | No friendly name. Use the Internal ID attribute listed above. |
| Property Attributes | NO_ATTRIBUTES |
| User Type | IIS_MD_UT_SERVER |

**Configurable Locations**
You can configure this property at the following locations in the IIS metabase.

**Flags**
There are no flags for this property.

# ServerListenBacklog Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ServerListenBacklog** property specifies the number of outstanding sockets that can be queued. The value is based on the *AcceptEx* operating system parameter, and the server size specified in the ServerSize Metabase Property. Valid values for this property range from 5 to 500.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SERVER_LISTEN_BACKLOG |
| Attributes | INHERIT |
| Default Value | 40 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1019 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- ServerSize Metabase Property

# ServerListenTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ServerListenTimeout** property specifies the amount of time (in seconds) the server waits for a client to send data. After the allotted amount of time passes, the server disconnects the client. This property is not used by the WWW service in IIS 6.0, but it remains available for the FTP, SMTP, and NNTP services.

In IIS 6.0, this property does not apply to the Web service (IIsWebService, IIsWebServer). Use ConnectionTimeout Metabase Property, HeaderWaitTimeout Metabase Property, and MinFileBytesPerSec Metabase Property instead.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SERVER_LISTEN_TIMEOUT |
| Attributes | INHERIT |
| Default Value | 120 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1020 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ServerSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ServerSize** property specifies the general size of the server, in terms of the number of client requests processed per day. A value of 0 (zero) indicates a small Web site that expects to receive fewer than 10,000 requests per day; a value of 1 indicates a medium site handling between 10,000 and 100,000 requests per day; and a value of 2 designates a large site, processing more than 100,000 requests a day. The value of this property is used in calculating the value for the ServerListenBacklog Metabase Property.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SERVER_SIZE |
| Attributes | INHERIT |
| Default Value | 1 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1018 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/W3SVC | IIsWebService |
| /LM/MSFTPSVC | IIsFtpService |

**Values**

| Constant | Value | Description |
|---|---|---|
| MD_SERVER_SIZE_SMALL | 0x00000000 | Anticipate fewer than 10,000 requests per day. |
| MD_SERVER_SIZE_MEDIUM | 0x00000001 | Anticipate 10,000 to 100,000 requests per day. |
| MD_SERVER_SIZE_LARGE | 0x00000002 | Anticipate more than 100,000 requests per day. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- ServerListenBacklog Metabase Property

# ServerState Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ServerState** property presents the current state of the server instance. The state codes, and the corresponding states, are as follows:

1 (starting)

2 (started)

3 (stopping)

4 (stopped)

5 (pausing)

6 (paused)

7 (continuing)

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_SERVER_STATE |
| Attributes | VOLATILE |
| Default Value | 4 |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1016 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Bitmask Values**

| Constant | Value | Description |
|---|---|---|

| | | |
|---|---|---|
| MD_SERVER_STATE_STARTING | 0x00000001 | Server starting. |
| MD_SERVER_STATE_STARTED | 0x00000002 | Server started. |
| MD_SERVER_STATE_STOPPING | 0x00000003 | Server stopping. |
| MD_SERVER_STATE_STOPPED | 0x00000004 | Server stopped. |
| MD_SERVER_STATE_PAUSING | 0x00000005 | Server pausing. |
| MD_SERVER_STATE_PAUSED | 0x00000006 | Server paused. |
| MD_SERVER_STATE_CONTINUING | 0x00000007 | Server continuing. |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SessionKey Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SessionKey** property is used for the encryption of secure metabase properties.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | binary |
| WMI Data Type | binary |
| ADSI Data Type | binary |
| ABO Data Type | binary |
| ABO Metabase Identifier | MD_GLOBAL_SESSIONKEY |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 9999 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| . | IIS_Global |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ShouldDeliver Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ShouldDeliver** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36881 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ShouldPickupMail Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ShouldPickupMail** property determines if mail is picked up from the Pickup directory.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36876 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ShouldPipelineIn Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ShouldPipelineOut** property allows pipelining of outgoing commands (Boolean value).

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36901 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ShouldPipelineOut Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ShouldPipelineIn** property allows pipelining of incoming commands (Boolean value).

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36902 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ShutdownLatency Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ShutdownLatency** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45070 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC | IIsNntpService |
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# ShutdownTimeLimit Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **ShutdownTimeLimit** property specifies the amount of time (in seconds) after a recycle threshold has been reached that IIS waits for all old requests to finish running in a worker process before terminating the worker process.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_SHUTDOWN_TIMELIMIT |
| Attributes | INHERIT |
| Default Value | 90 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9012 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n*/*virtual_directory_name*/*physical_directory_name* | IIsWebDirectory |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmartHost Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmartHost** property corresponds to "Smart host" in the Advanced Delivery dialog box (for a given virtual server, not in the general SmtpSvc container).

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36906 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmartHostType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmartHostType** property corresponds to the "Attempt direct delivery before sending to smart host" check box in the **Advanced Delivery** dialog box. If no smart host is specified in the smart host text box, **SmartHostType** will be set to 0. When the smart host is specified and the check box is not selected, **SmartHostType** will be set to 2 (always) and the smart host will be used to send mail without consulting the DNS server. When the smart host is specified and the check box is selected, **SmartHostType** will be set to 1 (that is, direct delivery through DNS lookup first; if DNS fails to look up the domains of the recipients, the smart host entry of the SMTP virtual server is used).

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36905 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SMPAffinitized Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Setting the **SMPAffinitized** property to true indicates that a particular worker process assigned to an application pool should be assigned to a given CPU. This property is used in conjunction with the SMPProcessorAffinityMask Metabase Property to configure a particular processor a worker process will be assigned to.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_APPPOOL_SMP_AFFINITIZED |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9007 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- SMPProcessorAffinityMask Metabase Property

# SMPProcessorAffinityMask Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SMPProcessorAffinityMask** property configures the hexadecimal processor mask. The hexadecimal processor mask indicates to which CPU the worker processes in an application pool should be bound. Before this property takes affect, the SMPAffinitized Metabase Property must be set to **true** for the application pool.

| **Note** |
| --- |
| Do not set this property to zero. Doing so causes no SMP affinity to be configured, creating an error condition. |

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_SMP_AFFINITIZED_PROCESSOR_MASK |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9008 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- SMPAffinitized Metabase Property

# SmtpAdvQueueDll Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpAdvQueueDll** property indicates the string containing the location of the advanced queuing DLL. The default is *systemroot*\Inetsrv\Aqueue.dll.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36987 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpAqueueWait Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpAqueueWait** property key is for debugging purposes only.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36990 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpAuthTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpAuthTimeout** property sets the amount of time the server waits after a client issues this command before hanging up because of inactivity.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37006 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpBdatTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpBdatTimeout** property sets the amount of time the server waits after a client issues this command before hanging up because of inactivity.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37005 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpClearTextProvider Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpClearTextProvider** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36942 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpCommandLogMask Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpCommandLogMask** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36864 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpConnectTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpConnectTimeout** property sets the amount of time, after a client makes a connection, that the server will wait before hanging up because of inactivity.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37001 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDataTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDataTimeout** property sets the amount of time, after a client issues this command, that the server will wait before hanging up because of inactivity.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37004 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDisableRelay Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDisableRelay** property controls whether mail may be relayed. Note that the Boolean flag value is reversed.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36974 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDomainValidationFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Setting the **SmtpDomainValidationFlags** property key to a nonzero value causes the domain passed to the EHLO command to be checked to determine if it is well formatted.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36992 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDotStuffPickupDirFiles Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDotStuffPickupDirFiles** property specifies if SMTP will process dotstuffed files in the Pickup directory.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37030 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsAccount Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsAccount** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36957 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsBindType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsBindType** property specifies whether to authenticate to a remote LDAP server or not.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36953 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsDataDirectory Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsDataDirectory** property is obsolete.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36951 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsDefaultMailRoot Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsDefaultMailRoot** property is obsolete.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36952 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsDomain** property sets the domain to use to authenticate to a remote LDAP server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36993 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsFlags Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsFlags** property determines special flags passed to set options on an LDAP connection.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37024 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsHost Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsHost** property sets the host name of a remote LDAP server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36955 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsNamingContext Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsNamingContext** property sets the LDAP base naming context to search in on a remote LDAP server.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36956 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDSNLanguageID Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDSNLanguageID** property indicates the language ID that the SMTP virtual server uses when sending Delivery Status Notifications.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37022 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDSNOptions Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDSNOptions** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37021 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsPassword Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsPassword** property sets the password to use to bind to a remote LDAP server.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT \| SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36958 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsPort Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsPort** property sets the port on which a remote LDAP server listens.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37023 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsSchemaType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsSchemaType** property specifies which directory schema to use.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36954 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpDsUseCat Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpDsUseCat** property enables the Categorizer.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36994 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpEventlogLevel Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpEventlogLevel** property sets the event logging level:

0=none

1=minimum

3=medium

5=maximum

7=field engineering

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37027 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpFlushMailFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpFlushMailFile** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36866 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpHeloNoDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpHeloNoDomain** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36981 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpHeloTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpHeloTimeout** property determines how long the server waits, after a client issues this command, before hanging up because of inactivity.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37012 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpInboundCommandSupportOptions Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This property is reserved for future use by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 7697601 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36998 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpIpRestrictionFlag Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpIpRestrictionFlag** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37031 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpLocalDelayExpireMinutes Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpLocalDelayExpireMinutes** property corresponds to where users set the local "Delay notification" time period on the **Delivery** tab. The value in the key is in minutes. This property is listed for both the \smtpsvc and \smtpsvc\1. Changing the setting changes the value of the key for the individual virtual server.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 60 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37017 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpLocalNDRExpireMinutes Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpLocalNDRExpireMinutes** property corresponds to where users set the local "Expiration timeout" time period on the **Delivery** tab. The value in the key is in minutes. This property is listed for both the \smtpsvc and \smtpsvc\1. Changing the setting changes the value of the key for the individual virtual server.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 2880 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37018 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpMailFromTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpMailFromTimeout** property determines how long the server will wait, after a client issues this command, before hanging up because of inactivity.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37002 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpMailNoHelo Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpMailNoHelo** property allows a "mail from" to be issued without issuing an EHLO (or HELO).

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36982 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpMaxRemoteQThreads Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpMaxRemoteQThreads** property is for debugging purposes only.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36964 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpOutboundCommandSupportOptions Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpOutboundCommandSupportOptions** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 7697601 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36999 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpRcptToTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpRcptToTimeout** property determines how long the server waits, after a client issues this command, before hanging up because of inactivity.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37003 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpRemoteDelayExpireMinutes Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpRemoteDelayExpireMinutes** property corresponds to where users set the remote "Delay notification" time period on the **Delivery** tab. The value in the key is in minutes. This property is listed for both the \smtpsvc and \smtpsvc\1. Changing the setting changes the value of the key for the individual virtual server.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 60 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37019 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpRemoteNDRExpireMinutes Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpRemoteNDRExpireMinutes** property corresponds to where users set the remote "Expiration timeout" time period on the **Delivery** tab. The value in the key is in minutes. This property is listed for both the \smtpsvc and \smtpsvc\1. Changing the setting changes the value of the key for the individual virtual server.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 2880 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37020 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpRemoteProgressiveRetry Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpRemoteProgressiveRetry** property corresponds to where users set "First retry interval," "Second retry interval," "Third retry interval," and "Subsequent retry interval" for an SMTP message on the **Delivery** tab. This property exists for \smtpsvc and \smtpsvc\1. Setting the value changes the metabase setting for the individual virtual server, not the whole smtpsvc.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36996 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpRemoteRetryThreshold Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpRemoteRetryThreshold** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 3 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36995 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpRoutingTableType Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpRoutingTableType** property is obsolete.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | ldap |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36950 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC/*n*/RoutingSources | IIsSmtpRoutingSource |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpRsetTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpRsetTimeout** property determines how long the server waits, after a client issues this command, before hanging up because of inactivity.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37010 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpSaslTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpSaslTimeout** property determines how long the server waits, after a client issues this command, before hanging up because of inactivity.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37007 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpServer Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpServer** property contains the address of the SMTP server for use when sending e-mail to a moderator.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45059 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpServiceVersion Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpServiceVersion** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36931 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpSSLCertHostnameValidation Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Certificates have a "common name" that matches the name of the computer that owns it. The **SmtpSSLCertHostnameValidation** property verifies that the "common name" of a certificate matches the name of the remote computer that sent the certificate.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37026 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpSSLRequireTrustedCA Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpSSLRequireTrustedCA** property enables verification of a certificate that a remote host gives you when you connect to it. It makes sure this certificate is in your trusted certificate root on your computer. If not, it ends the SSL session and the handshake is rejected.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37025 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SmtpTurnTimeout Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpTurnTimeout** property determines how long the server waits, after a client issues this command, before hanging up because of inactivity.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 37009 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase PropertiesCode Examples to Configure Metabase Properties.

# SmtpUseTcpDns Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SmtpUseTcpDns** property specifies if SMTP uses TCP for DNS. The value can be 0 (IE UDP only), 1 (TCP only), 2 (try UDP and then TCP), or others.

The metabase represents unlimited as the DWORD value of 4294967295 (0xFFFFFFFF); however, VBScript represents unlimited in hexadecimal format as &HFFFFFFFF. Previous versions of IIS represented unlimited as -1.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | Boolean |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | 4294967295 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 36997 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SSIExecDisable Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SSIExecDisable** property specifies whether server-side include (SSI) #exec directives are disabled under this path.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_SSI_EXEC_DISABLED |
| Attributes | INHERIT |
| Default Value | **FALSE** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6028 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name* /LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SSLAlwaysNegoClientCert Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SSLAlwaysNegoClientCert** property controls SSL client connection negotiations. If this property is set to **true**, any time SSL connections are negotiated, the server will immediately negotiate a client certificate, preventing an expensive renegotiation. Setting **SSLAlwaysNegoClientCert** also helps eliminate client certificate renegotiation deadlocks, which may occur when a client is blocked on sending a large request body when a renegotiation request is received.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | BOOL |
| WMI Data Type | BOOL |
| ADSI Data Type | BOOL |
| ABO Data Type | BOOL |
| ABO Metabase Identifier | MD_SSL_ALWAYS_NEGO_CLIENT_CERT |
| Attributes | INHERIT |
| Default Value | **FALSE** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5521 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SSLCertHash Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SSLCertHash** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | binary |
| WMI Data Type | binary |
| ADSI Data Type | binary |
| ABO Data Type | binary |
| ABO Metabase Identifier | MD_SSL_CERT_HASH |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5506 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SslCtlIdentifier Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SslCtlIdentifier** property contains a unique value that identifies a specific certificate trust list (CTL). It must be used with SslCtlStoreName Metabase Property to accurately reference a CTL.

| **Attribute Name** | **Attribute Value** |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5512 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SslCtlStoreName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SslCtlStoreName** property contains the name of the CryptoAPI store that contains certificate trust lists (CTL). It must be used with SslCtlIdentifier Metabase Property to accurately reference a CTL.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_SSL_CERT_STORE_NAME |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5517 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SSLStoreName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SSLStoreName** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_SSL_CERT_STORE_NAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5511 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SslUseDsMapper Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **SslUseDsMapper** property specifies whether IIS is to use the Windows Directory Service certificate mapper or IIS certificate mapper. If **SSLUseDSMapper** is set to **false**, IIS uses the IIS certificate mapper.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_SSL_USE_DS_MAPPER |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5019 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- AccessSSLFlags Metabase Property

# StartupTimeLimit Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The value of the **StartupTimeLimit** property specifies the amount of time (in seconds) that the World Wide Web Publishing Service (WWW Service) should wait for a worker process to finish starting up and reporting to the WWW Service.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_APPPOOL_STARTUP_TIMELIMIT |
| Attributes | INHERIT |
| Default Value | 90 |
| MetaFlagsEx | CACHE_PROPERTY_CLEARED |
| User Type | IIS_MD_UT_SERVER |
| ID | 9011 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# SuppressDefaultFTPBanner Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By default, the FTP service sends a "220 Microsoft FTP Service" banner when clients connect to the server. The **SuppressDefaultFTPBanner** property allows you to disable the display of the default FTP banner.

When **SuppressDefaultFTPBanner** is set to **true**, the message sent to the client will only include the number "220", not the "Microsoft FTP Service" part of the banner name.

| **Important** |
| --- |
| Configurable banner messages which follow this banner are not affected. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_SHOW_4_DIGIT_YEAR |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlags | CASEINSENSITIVE |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5017 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/MSFTPSVC/*n* | IIsFtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# TraceUriPrefix Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **TraceUriPrefix** metabase property defines a specific URL or set of URLs to trace for IIS 6.0 request-based tracing (available in Windows Server 2003 Service Pack 1 or later). When setting this property specify either a physical to a file or a relative path with a wild card.

For example: To trace a specific URL, you would type the following at a command prompt:

**adsutil.vbs set w3svc/** *n* **/TraceUriPrefix=** *path to URL*

and press ENTER. *N* is the number of the site as it is listed in the metabase, e.g. w3svc/1 is the default Web site (by default). Path to URL is the physical path to the file you want to trace, e.g. /Inetpub/wwwroot/asp/default.asp

To trace a set of URLs, for example all .asp pages or .aspx pages, you would type the following at a command prompt:

**adsutil.vbs set w3svc/** *n* **/TraceUriPrefix**=*path to URL*extension*

and press ENTER. *N* is the number of the site as it is listed in the metabase, e.g. w3svc/1 is the default Web site (by default). *Path to URL*extension* is the physical path to the group of files you want to trace, e.g. /Inetpub/wwwroot/asp/*.asp or /Inetpub/wwwroot/aspnet/*.aspx

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_APP_TRACE_URI_LIST |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 2118 |

**Configurable Locations**
You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IISWebServer |
| /LM/W3SVC/*n* | IISWebServer |

**Code Example**
For general code examples, see Code Examples to Configure Metabase Properties.

**See Also**

**Tasks**

How to Trace Requests for a Specific URL or Set of URLs

# UNCPassword Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **UNCPassword** property specifies the encrypted password used to gain access to UNC (Universal Naming Convention) virtual roots.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_VR_PASSWORD |
| Attributes | INHERIT \| SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 3003 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/MSFTPSVC/*n*/ROOT /LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |
| /LM/SMTPSVC/*n*/ROOT /LM/SMTPSVC/*n*/ROOT/*virtual_directory_name* | IIsSmtpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# UNCUserName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **UNCUserName** property specifies the user name for Universal Naming Convention (UNC) virtual roots. Use this property only as a flag for authentication schemes that support delegation.

| **Note** |
| --- |
| Integrated Windows authentication (NTLM) does not support delegation. |

When a request comes in to the IIS server, IIS looks at the **UNCUserName** property on the virtual directory. If the property exists and it is not an empty string, then the UNC user name and password are used for authentication. If authentication fails, a 401.2 error is returned to the client. If authentication passes, then the client's original credentials are passed through to the resources that were requested.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_VR_USERNAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 3002 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/MSFTPSVC/*n*/ROOT /LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |
| /LM/SMTPSVC/*n*/ROOT /LM/SMTPSVC/*n*/ROOT/*virtual_directory_name* | IIsSmtpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

- UNCPassword Metabase Property
- Path Metabase Property

# UpdatedDefaultDomain Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **UpdatedDefaultDomain** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 36937 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# UpdatedFQDN Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **UpdatedFQDN** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **true** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 36936 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/SMTPSVC | IIsSmtpService |
| /LM/SMTPSVC/*n* | IIsSmtpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# UploadReadAheadSize Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **UploadReadAheadSize** property establishes the number of bytes a Web server will read into a buffer and pass to an ISAPI extension. This occurs once per client request. The ISAPI extension receives any additional data directly from the client. The range is 0 to &HFFFFFFFF (4 GB).

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_UPLOAD_READAHEAD_SIZE |
| Attributes | INHERIT |
| Default Value | 49152 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6045 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name*<br>/LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# UseDigestSSP Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Advanced Digest authentication uses the **UseDigestSSP** metabase property to switch between Digest and Advanced Digest security support provider interface (SSPI) code. If this property is set to **false**, Digest authentication is used. In all other cases (**true**, empty, or not set), IIS uses Advanced Digest authentication. If you configure UseDigestSSP at the W3SVC level of the metabase, all child keys inherit their configuration from that level.

| **Important** |
| --- |
| If Digest SSPI code is used and dedicated application mode is selected, the Local System account must be used as the identity. |

The World Wide Web Publishing Service (WWW service) must be restarted before changes to this property take effect.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | MD_USE_DIGEST_SSP |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 6047 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n*/ROOT/*file_name* /LM/W3SVC/*n*/ROOT/*virtual_directory_name/file_name* | IIsWebFile |
| /LM/W3SVC/*n*/ROOT/*physical_directory_name* /LM/W3SVC/*n/virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# UserIsolationMode Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **UserIsolationMode** property defines the isolation type for the corresponding FTP site.

If **UserIsolationMode** is set to **0**, there is no user isolation. This setting is the default and is backward-compatible.

If the property is set to **1**, a client authenticates using local or domain accounts and is then sent to a folder under the root that matches the user name. This setting is called "Isolated (Locally)," and it supports users who do not want to use Active Directory.

If the property is set to **2**, user isolation is dependent on Active Directory. This setting is called "Isolated (Active Directory)," and it is primarily used by Internet service providers (ISPs) and other customers who want to set up large numbers of FTP accounts.

If you set this property to **2**, you should also set the DefaultLogonDomain Metabase Property, ADConnectionsPassword Metabase Property, and ADConnectionsUserName Metabase Property properties.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_USER_ISOLATION |
| Attributes | INHERIT |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 5012 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC | IIsFtpService |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# VrDoExpire Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **VrDoExpire** property controls whether the virtual root uses its own expiration policy.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **false** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 46161 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# VrDriverClsid Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **VrDriverClsid** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 46156 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# VrDriverProgid Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **VrDriverProgid** Metabase Property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 46157 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# VrOwnModerator Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **VrOwnModerator** property is used by Microsoft Exchange.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | Boolean |
| WMI Data Type | Boolean |
| ADSI Data Type | Boolean |
| ABO Data Type | Boolean |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 46163 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# VrUseAccount Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **VrUseAccount** property controls wither the virtual root account is used to log on.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 46160 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# WAMUserName Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **WAMUserName** property specifies the account user name that IIS uses by default as the COM+ application identity for newly created IIS out-of-process applications. The values of this property and its companion property, WAMUserPass Metabase Property, are set when IIS is installed, and match the user name and password values in the Microsoft Windows user account, which is established at the same time. Changing the value of this property is not recommended. If you do, change it to a valid Windows user account, and change **WAMUserPass** to the corresponding password for the new account.

| **Important** |
| --- |
| Changes to **WAMUserName** and **WAMUserPass** may disrupt the operation of existing IIS out-of-process applications. You can synchronize application identities using Component Services to edit the user name and password values, found on the **Identity** tab of the property sheet for each package. |

In-process applications are not affected by these property values.

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_WAM_USER_NAME |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 7501 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

WAMUserPass Metabase Property

# WAMUserPass Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **WAMUserPass** property specifies the password for the account that IIS uses by default as the COM+ application identity for newly created IIS out-of-process applications. The values of this property and its companion property, **WAMUserName**, are set when IIS is installed, and match the password and user name values in the Microsoft Windows user account (IWAM_*MachineName*, where *MachineName* is the name of the machine on which IIS is installed) established at the same time.

| **Important** |
| --- |
| Changing the value of this property is not recommended. If you do, you must change the Windows account password to the identical value. You must also synchronize existing IIS out-of-process application identities, using Component Services to edit the user name and password values, which are found on the **Identity** tab of the property sheet for each package. |

In-process applications are not affected by these property values.

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | MD_VR_PATH_DELIMITER_string |
| Attributes | INHERIT \| SECURE |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_FILE |
| ID | 7502 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

**Related Information**

WAMUserName Metabase Property

# WebDAVMaxAttributesPerElement Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003 with SP1

The **WebDAVMaxAttributesPerElement** metabase property specifies the maximum number of attributes allowed per XML element, including namespaces.

| **Note** |
| --- |
| For Windows Server 2003 with Service Pack 1 (SP1), this property will be added to the IIS schema, but it must be set manually. |

| **Attribute Name** | **Attribute Value** |
| --- | --- |
| **XML Data Type** | DWORD |
| **WMI Data Type** | DWORD |
| **ADSI Data Type** | DWORD |
| **ABO Data Type** | DWORD_METADATA |
| **ABO Metabase Identifier** | MD_WEBDAV_MAX_ATTRIBUTES_PER_ELEMENT |
| **Attributes** | INHERIT |
| **Default Value** | 32 |
| **MetaFlagsEx** | CACHE_PROPERTY_MODIFIED |
| **User Type** | IIS_MD_UT_SERVER |
| **ID** | 8501 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| **Metabase Path** | **IIS Admin Object Type** |
| --- | --- |
| /LM/W3SVC | IIsWebService |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/W3SVC/*n*/ROOT | IIsWebVirtualDir |
| /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n*/ROOT/*virtual_directory_name/physical_directory_name* | IIsWebDirectory |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# WebSvcExtRestrictionList Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **WebSvcExtRestrictionList** property contains a multi-string list of ISAPIs and CGIs used by applications on the Web server, including their location, descriptive information, and a flag for determining whether the Web service extension is enabled or disabled. The **WebSvcExtRestrictionList** property is used by the Web server to determine whether an ISAPI extension is permitted to run.

Each string element is written as:

```
AllowDenyFlag,ExtensionPhysicalPath,UIDeletableFlag,GroupID,Description
```

- AllowDenyFlag is "0" if the extension is not allowed or "1" if the extension is allowed.
- ExtensionPhysicalPath is the full path to the extension. If the last four characters of the physical path are ".dll" (case insensitive), then the extension is an ISAPI, otherwise it is a CGI. In the case of CGI script mappings, the physical path must contain the physical path from the mapping (i.e. "c:\windows\system32\cmd.exe %s %s").
- UIDeletableFlag is "1" if the UI is allowed to delete the extension from the list or "0" if the UI is not allowed to delete it.
- GroupID is a non-localizable string used to identify groups of extensions.
- Description is a localizable description of the extension.

**Example**

Some examples of elements in this property are shown below. The first two entries, for *.dll and *.exe are special cases, referring to "All Unlisted ISAPIs" and "All Unlisted CGIs," respectively. The Web server core enables or disables all ISAPIs or all CGIs depending on the setting, if these entries are present:

```
"0,*.dll"
"0,*.exe"
"0,c:\windows\system32\inetsrv\asp.dll,0,ASP,Active Server Pages"
"0,c:\windows\system32\inetsrv\httpodbc.dll,0,HTTPODBC,Internet Data Connector"
"0,c:\windows\system32\inetsrv\ssinc.dll,1,SSINC,Server Side Includes"
```

Configuring the **WebSvcExtRestrictionList** property using Admin Base Object (ABO) can be risky because the format of the property must be perfect. The property is represented as a buffer of strings that are delimited by \0 and composed of five fields separated by commas. A safer way to configure this property is by using the extension-related ADSI and WMI methods of the **IIsWebService** object. For a brief comparison of ADSI and WMI, see Administering IIS Programmatically.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | MULTISZ |
| WMI Data Type | MULTISZ |
| ADSI Data Type | MULTISZ |
| ABO Data Type | MULTISZ |
| ABO Metabase Identifier | MD_WEB_SVC_EXT_RESTRICTION_LIST |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |

| | |
|---|---|
| ID | 2168 |

**Configurable Locations**
You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC | IIsWebService |

**Code Example**
For general code examples, see Code Examples to Configure Metabase Properties.

# Win32Error Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **Win32Error** property indicates the Microsoft Win32 error status code. For more information about Microsoft Win32 error codes, see "Win32 Error Codes" on MSDN.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | DWORD |
| WMI Data Type | SINT32 |
| ADSI Data Type | DWORD |
| ABO Data Type | DWORD |
| ABO Metabase Identifier | MD_WIN32_ERROR |
| Attributes | VOLATILE |
| Default Value | 0 |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 1099 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/W3SVC/*n*/ROOT /LM/W3SVC/*n*/ROOT/*virtual_directory_name* | IIsWebVirtualDir |
| /LM/W3SVC/*n* | IIsWebServer |
| /LM/MSFTPSVC/*n* | IIsFtpServer |
| /LM/MSFTPSVC/*n*/ROOT /LM/MSFTPSVC/*n*/ROOT/*virtual_directory_name* | IIsFtpVirtualDir |
| /LM/NNTPSVC/*n* | IIsNntpServer |
| /LM/NNTPSVC/*n*/ROOT /LM/NNTPSVC/*n*/ROOT/*virtual_directory_name* | IIsNntpVirtualDir |
| /LM/SMTPSVC/*n* | IIsSmtpServer |
| /LM/SMTPSVC/*n*/ROOT /LM/SMTPSVC/*n*/ROOT/*virtual_directory_name* | IIsSmtpVirtualDir |
| /LM/W3SVC/AppPools/DefaultAppPool<br>/LM/W3SVC/AppPools/DefaultAppPool/*application_pool_name* | IIsApplicationPool |
| /LM/W3SVC/AppPools | IIsApplicationPools |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# XMLSchemaTimeStamp Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **XMLSchemaTimeStamp** property keeps the MBSchema.xml file in sync with the binary schema file.

| **Important** |
| --- |
| Because this property is internally configured by IIS, you should consider it to be read-only. Do not configure this property. |

| Attribute Name | Attribute Value |
| --- | --- |
| XML Data Type | binary |
| WMI Data Type | binary |
| ADSI Data Type | binary |
| ABO Data Type | binary |
| ABO Metabase Identifier | MD_GLOBAL_XMLSCHEMATIMESTAMP |
| Attributes | NO_ATTRIBUTES |
| Default Value | **null** |
| MetaFlagsEx | Not applicable |
| User Type | IIS_MD_UT_SERVER |
| ID | 9992 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
| --- | --- |
| . | IIS_Global |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# XoverTableFile Metabase Property

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The **XoverTableFile** property specifies the Xover hash file.

| Attribute Name | Attribute Value |
|---|---|
| XML Data Type | String |
| WMI Data Type | String |
| ADSI Data Type | String |
| ABO Data Type | String |
| ABO Metabase Identifier | No friendly name is defined for this identifier in iiscnfg.h. Search the *systemroot*\System32\Inetsrv\MBSchema.xml file for this property and use the integer from its ID attribute. |
| Attributes | INHERIT |
| Default Value | **null** |
| MetaFlagsEx | CACHE_PROPERTY_MODIFIED |
| User Type | IIS_MD_UT_SERVER |
| ID | 45161 |

**Configurable Locations**

You can configure this property at the following locations in the IIS metabase.

| Metabase Path | IIS Admin Object Type |
|---|---|
| /LM/NNTPSVC/*n* | IIsNntpServer |

**Code Example**

For general code examples, see Code Examples to Configure Metabase Properties.

# IIS 6.0 Registry Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The configuration registry stores entries that define the working environment for the Microsoft® Windows® operating system and any services that are installed on a computer running Windows. You can use graphical tools (such as Control Panel or Setup) to change these values, or you can use the Registry Editor (Regedit.exe) utility that is included in Windows to modify the configuration registry directly.

This section includes the following information:

- About the Registry: Provides conceptual topics related to the registry.
- Global Registry Entries: Specifies entries used for global control of Internet services.
- Service-Specific Registry Entries: Specifies entries stored in the registry by Internet services for service-specific behavior.
- ASP Registry Entries: Specifies entries for Active Server Pages (ASP).
- WWW Service Registry Entries: Specifies entries for the World Wide Web Publishing Service (WWW service) only.
- FTP Service Registry Entries: Specifies entries for FTP services only.

**Related Information**

- For information about the metabase, see IIS Metabase.
- For information about routine server administration tasks, see Server Administration Guide.

# About the Registry

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The registry is a repository for information about a computer's configuration, such as profiles for each user and the programs that are installed. The registry is organized hierarchically and is made up of keys and their subkeys, hives, and entries.

Many entries that existed in the registry for previous versions of Internet Information Services (IIS) were migrated to a fast-access, hierarchical database, called the *metabase*. Entries that were migrated to the metabase are known as *metabase properties*. Most of the remaining registry keys are used to initiate Internet services and enable backward compatibility with previous versions of IIS, are specific to a service, or are global to IIS and all services. Any new configuration for IIS should be written to the metabase.

Windows includes Registry Editor, which can be used to configure the registry. For information about using Registry Editor, see "Registry Editor" in Help and Support Center for Windows Server 2003.

| **Important** |
| --- |
| For a registry change to take effect, you must restart the service that is affected by the change. For global changes to take effect, you must restart all services. |

**Related Information**

- For more information about the registry, see "Registry structure" in Help and Support Center for Windows Server 2003.
- For information about how to restart IIS, see Restarting IIS.
- For more information about the IIS metabase, see IIS Metabase.

# Global Registry Entries

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following registry entries affect all of the IIS services, without having to be set for each individual service. On a default installation of IIS, you might not see some of these registry entries, which indicates that the default values are being used. To set these entries to values other than the default values, you must manually create the keys and set them.

| **Important** |
| --- |
| For a registry change to take effect, you must restart the service that is affected by the change. For global changes to take effect, you must restart all services. |

| **CacheSecurityDescriptor** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 1 (enabled) |
| Range: 0 - 1 |
| Specifies whether security descriptors are cached for file objects. When **CacheSecurityDescriptor** is enabled, IIS retrieves security permissions when caching a file object and does not need to gain access to the file object to check access rights for new users. The feature is useful only if you have more than one user account (Web server does not use only Anonymous authentication). |

| **CheckCertRevocation** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 0 (disabled) |
| Range: 0 - 1 |
| Specifies whether client certificates are checked for revocation by IIS. **CheckCertRevocation** is disabled by default because checking for common certificate issuers is typically done over the Internet and has severe performance impact. However, enabling **CheckCertRevocation** may be useful if you issue your own certificates and the revocation process is local. |

| **DataSetCacheSize** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 50 |

| |
|---|
| Each IIS worker process (w3wp.exe) contains a local metabase data cache of up to 50 entries. For each request, the required dataset is obtained from the local metabase cache if possible. If the dataset is not found in the cache, then an expensive LPC request is made to the Inetinfo.exe process for the dataset residing in the metabase. When there are more than 50 virtual directories, IIS disables the metabase data cache, forcing all requests to go via the expensive path (which negatively impacts throughput).<br><br>This registry entry allows administrators or IT managers to set the maximum number of dataset entries in the cache if the number of total installed virtual directories exceeds the default (50). This key is not created by default. |

| **DigestContextCacheTTL** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 300 (seconds) |
| Range: 0 - 4,294,967 |
| When using Advanced Digest Authentication, this value overrides the Time To Live (TTL) setting for fully formed security contexts after the first authentication request succeeds. |

| **DigestPartialContextCacheTTL** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 30 (seconds) |
| Range: 0 - 4,294,967 |
| When using Advanced Digest Authentication, this value overrides the Time To Live (TTL) setting for partially formed security contexts for the first authentication challenge request. |

| **DisableLazyContentPropagation** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters |
| Data Type: REG_DWORD |
| Default Value: 0 (lazy propagation enabled) |
| Range: 0 - 1 |

| |
|---|
| Lazy propagation refers to the action that IIS takes when a large amount of content is updated at one time.<br><br>IIS has an internal limit of how much content can be updated in the in-memory template cache. If the size of the updated content exceeds that limit, IIS marks each of the files in the in-memory template cache as invalid. On the first request to an invalid file, IIS begins to compile a new template but the expired template is served and is continued to be served until the new template is compiled.<br><br>If you set **DisableLazyContentPropagation** to 1, IIS behaves as it does for IIS 5.0 and IIS 5.1 when a large amount of content is updated at one time. IIS flushes the in-memory template cache, and performance can slow to a standstill as each new request to the server forces IIS to compile new templates.<br><br>If you are developing Web pages and making changes to a few files at a time, you do not have to refresh your pages twice to see new changes because lazy propagation only applies when a large amount of content is updated at one time. |

| **DisableMemoryCache** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 0 (disabled) |
| Range: 0 - 1 |
| Disables server caching. |

| **DoDirMonitoringForUNC** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Inetinfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 0 (disabled) |
| Range: 0 (disabled) - 1 (enabled) |
| This setting determines whether static files rely on the Change Notification mechanism for UNC to determine whether or not content is stale. The File Change Notification mechanism is not considered robust enough for tracking changes over UNC, and for this reason this setting is disabled. IIS and ASP use the Last Modified mechanism to track files instead. |

| **ErrorLoggingDir** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\HTTP\Parameters |
| Data Type: REG_STRING |
| Default Value: *systemroot*\system32\logfiles |
| Range: N/A |
| Specifies the location of the Httperr.log file. Httperr.log is a new feature in IIS 6.0. If HTTP.sys is unable to write an error to a Web site log because the request did not reach user mode, the error will be entered in Httperr.log. Errors that are written to Httperr.log include the following: 503 errors, idle connection time-outs, URL parsing errors, and the last 10 requests to a failed application pool. |

| **FileAttributeCheckThreshold** |
|---|

| |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 5 (seconds) |
| Range: 0 - 4,294,967 |
| When the Last Modified mechanism is used to track files, this key enables caching of static content from an UNC share in the HTTP.sys cache and specifies the amount of time for the content to remain in the cache. |

| **IncrementalSiteIDCreation** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\InetMgr\Parameters |
| Data Type: REG_DWORD |
| Default Value: Key is not created by default, and therefore site identifiers are large, unique, randomly generated numbers. |
| Range: 0 - 1 |
| When a new site is created on IIS 6.0, a site identification number is randomly generated using the name of the Web site. This ensures that all IIS 6.0 servers in a Web farm have a good chance of generating the same site identification number for sites of the same name. With IIS 5.1 and previous versions, site identification numbers were incremental. For example, because the default Web site is created first, its site identification number is 1, and the next site to be created is identified as 2.<br><br>If you have administration scripts that depend upon the IIS 5.1 method of generating site identification numbers, you can create the **IncrementalSiteIDCreation** key at the path specified above, and set it to 1 to force IIS to use incremental site identification numbers.<br><br>**Note**<br>When you remotely administer another IIS server, the value of the registry key of the local server is used to determine the generation of site identification numbers on the remote server.<br><br>To choose the identification number of your Web site, create the site programmatically using the IIsWebService.CreateNewSite or IIsFtpService.CreateNewSite WMI method, whose optional fourth parameter holds a user-chosen site identification number. |

| **LogBufferSize** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Http\Parameters\LogBufferSize |
| Data Type: REG_DWORD |
| Default Value: N/A |
| Range: 12,288 - 65,536 (12K - 64K) |
| Overrides the default logging buffer (per site) of 64K. Using this key, administrators can prevent excessive memory usage when running many thousands of sites on a server. This registry key does not exist by default and must be added to the registry. |

| **MaxCachedFileSize** |
|---|

| |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 256 KB (262,144 bytes) |
| Controls the maximum size of cache files. The default is 256 KB if no value is present. |

| **MaxConcurrency** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 0 |
| Range: 0 - 4,294,967,295 (unlimited) |
| Specifies the amount of concurrency that a system should provide. IIS uses completion ports to handle input-output (I/O). In general, have more than one thread running and conflicting on shared memory or locks. This key specifies how many threads per processor should be allowed to run simultaneously, if there is a pending I/O operation. A value of 0 allows the system to make the intelligent choice of the number of threads to use. Any nonzero value specifies that the system should allow that many threads per processor to run simultaneously. |

| **MaxPoolThreads** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 4 |
| Range: 0 - 4,294,967,295 (unlimited) |
| Specifies the number of pool threads to create per processor. Pool threads watch for the network for requests and process incoming requests. The **MaxPoolThreads** count does not include threads that are consumed by ISAPI applications. Generally, you should not create more than 20 threads per processor. |

| **MemCacheSize** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: approximately one half the available physical memory and calibrated in megabytes |
| Range: 0 - 2500MB |
| Controls the size of the cache. The default value dynamically adjusts every 60 seconds. |

| **MinFileBytesPerSec** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |

| |
|---|
| Data Type: REG_DWORD |
| Default Value: 240 |
| Range: 1 - 8192 |
| When a Web server sends a file to a client, a time-out is established for how long the server will allow the transfer to continue before ending the connection. The time-out chosen is the maximum of the connection time-out specified in IIS Manager plus the size of the file, and divided by the value specified for **MinFileBytesPerSec**. For example, a file size of 100 KB is given a time-out of 100 seconds, or the connection time-out if the latter is greater. |

| **ObjectCacheTTL** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 30 (seconds) |
| Range: 0 - 4,294,967,295 (unlimited) |
| Controls the Time To Live (TTL) setting, which defines the length of time that objects are held in cached memory. If an object in the memory cache has not been referenced for the defined period, that object will be phased out of the cache. If system memory is limited, or the server's contents are dynamic, you can use a lower TTL to prevent system memory from being used to cache a large number of volatile objects. Setting **ObjectCacheTTL** to 0xFFFFFFFF disables the object-cache scavenger and allows cached objects to remain in the cache until they are overwritten. Disabling the cache scavenger is useful if your server has ample system memory and your data is relatively static. |

| **PoolThreadLimit** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 2 * # MB |
| Range: 0 - 4,294,967,295 (unlimited) |
| Specifies the maximum number of pool threads that can be created in the system. Pool threads watch the network for requests and process incoming requests. **PoolThreadLimit** is a hard limit that includes all IIS threads and is always greater than or equal to **MaxPoolThreads**. |

| **ThreadPoolMaxCPU** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: -1 (CPU metric is turned off) |
| Range: 0 - 100 |
| ThreadPoolMaxCPU is the maximum CPU that the thread pool should throttle back at. |

| **ThreadPoolStartDelay** |
|---|

| |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 1000 ms |
| Range: 1000 - 10,000 |
| ThreadPoolStartDelay is how long the thread pool should set the timer for when zero threads are waiting on the completion port. |

| **ThreadTimeout** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 24 * 60 * 60 (24 hours, expressed in seconds) |
| Range: 0 - 4,294,967,295 (unlimited) |
| Specifies the amount of time an input-output (I/O) processing thread should be maintained, even if there is no I/O activity on the system. In general, when there is no I/O activity and no requests outstanding, the server remains idle and does not consume memory. But if the server continues and exceeds the **ThreadTimeout** interval, the I/O processing thread is stopped. Units are in seconds. |

| **UserTokenTTL** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters |
| Data Type: REG_DWORD |
| Default Value: 15 * 60 (15 Minutes, expressed in seconds) |
| Range: 0 - 0x7FFFFFFF |
| When a request is made to the server, the security credentials for the request (or the configured anonymous user) are used to create a user token on the server. The server impersonates this user token when accessing files or other system resources (see **CacheSecurityDescriptor**). The token is cached so that the Windows log on takes place only the first time the user accesses the system or after the user's token has fallen out of the cache. Integrated Windows authentication tokens are not cached. Units are in seconds. |

# Service-Specific Registry Entries

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

**Service-Specific Registry Entries**

The following registry key, LanguageEngines, is not installed by default. You must create this key in the following location in order to use it.

HKEY_LOCAL_MACHINE\SYSTEM

\CurrentControlSet

\Services

\W3SVC

\ASP

\LanguageEngines

\LanguageName

Value: Write REG_SZ: Response.WriteEquiv |

Value: WriteBlock REG_SZ: Response.WriteBlockEquiv |

| **Important** |
| --- |
| For a registry change to take effect, you must restart the service that is affected by the change. For global changes to take effect, you must restart all services. |

**LanguageEngines**

| | |
| --- | --- |
| Registry Path | (See above) |
| Data Type | REG_STRING |
| Default Value | Not applicable |
| Description | Specifies a scripting language that does not support the **Object.Method** syntax as the Active Server Pages (ASP) primary scripting language (the language used within the ASP script delimiters <% ... %>). *LanguageName* is the name of the chosen language. `Response.WriteEquiv` is the language's equivalent of **Response.Write**; and `Response.WriteBlockEquiv` is the language's equivalent of **Response.WriteBlock**. The pipe symbol (\|) is an insertion used by ASP to send expressions and HTML blocks that are normally processed with **Response.Write** and **Response.WriteBlock** methods. This can be done automatically when installing additional scripting languages. If your scripting language is sensitive to white space or newline characters, you may not be able to use it as the primary scripting language even when this registry value is set. As an alternative, you can manually write HTML blocks to the browser or write that language's functions within tagged script blocks (<SCRIPT> ... </SCRIPT> ) and call them from any other language. |

# ASP Registry Entries

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The following registry entries affect the behavior and performance of ASP on IIS. On a default installation of IIS, you will not see these registry entries, which means the default values are being used. To set these entries to their non-default values, you must manually create the keys and set them.

| **DisableLazyContentPropagation** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 0 (lazy propagation enabled) |
| Range: 0 - 1 |
| Lazy propagation refers to the action that IIS takes when a large amount of content is updated at one time.<br><br>IIS has an internal limit of how much content can be updated in the in-memory template cache. If the size of the updated content exceeds that limit, IIS marks each of the files in the in-memory template cache as invalid. On the first request to an invalid file, IIS begins to compile a new template but the expired template is served and is continued to be served until the new template is compiled.<br><br>If you set **DisableLazyContentPropagation** to 1, IIS behaves as it does for IIS 5.0 and IIS 5.1 when a large amount of content is updated at one time. IIS flushes the in-memory template cache, and performance can slow to a standstill as each new request to the server forces IIS to compile new templates. If you are developing Web pages and making changes to a few files at a time, you do not have to refresh your pages twice to see new changes because lazy propagation only applies when a large amount of content is updated at one time. |

| **ThreadMax** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 100 |
| Range: 0 - PoolThreadLimit |
| This value is a hard limit on the maximum number of ASP worker threads. The number of threads in ASP is generally determined by the number of processors in the system multiplied by the **AspProcessorThreadMax** metabase setting. If the product of the number of processors present and **AspProcessorThreadMax** exceeds the value of **ThreadMax**, the ASP engine will not spawn any additional threads than the value specified by **ThreadMax**.<br><br>**Note**<br>A large number of threads can reduce system performance because of the resulting high context switching rate. |

| **DisableComPlusCpuMetric** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |

| |
|---|
| Default Value: 0 (enabled) |
| Range: 0 (enabled) - 1 (disabled) |
| This setting determines whether COM+ controls the rate at which new threads are spawned. If **DisableComPlusCpuMetric** is set to 1 (disabled), COM+ will not control the rate of spawning of additional threads. If **DisableComPlusCpuMetric** is set to 0(the default), COM+ controls spawning of additional threads based on a number of factors; primarily on the CPU utilization and the Context Switch rate for the system. These values can also be adjusted using the registry. To specify the CPU utilization at which COM+ begins to throttle the thread creation, use the MaxCPU Registry key. To specify the context switch rate at which COM+ should throttle thread creation use the MaxCSR registry key.<br>**Important**<br>Changing these settings from the default value can have a severe impact on performance. It is recommended that you leave these values at their defaults. |

| **AspMaxResponseHeaderLength** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 64 Kilobytes |
| Range: 0 - ffffffff |
| This is a security setting that limits the size of the header generated using the response object by an ASP script file. If the size of the response header exceeds the value of **AspMaxResponseHeaderLength**, the system returns a 500 response to the client. |

| **AspMaxPropertyStringLength** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 16 Kilobytes |
| Range: 0 - ffffffff |
| This is a security setting that limits the string length of a single property name. If the string length of the property name exceeds the value of **AspMaxPropertyStringLength**, the system returns a 500 response to the client. |

| **DisableCachedResponseOnUNCAccessFailure** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 0 (enabled) |
| Range: 0 (enabled) - 1 (disabled) |

This setting determines whether ASP will serve a cached response in the event a network or other communication problem prevents ASP from validating the requested page against the back end system (mapped by way of a UNC path) containing the content to be served. By default, if communication with the back end system is not possible when ASP attempts to validate the requested page, ASP will serve the cached page. Note that the response is still checked against existing credentials stored with the file. An event log message will be logged to let the administrator know that the system took a corrective action. The administrator may want to look at the event log if customers complain of stale content. There might be a problem with the network connection with the back end. If **DisableCachedResponseOnUNCAccessFailure** is set to 1 (disabled), ASP returns a 500 error to the client. This setting can be useful for failure detection systems to kick in (e.g., if a load balancer attached to the system it can direct the request to another system.) This prevents serving stale content, and also notifies the administrator of the error.

**F5AttackDetectionEnabled**

Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\

Data Type: REG_DWORD

Default Value: 1 (enabled)

Range: 0 (disabled) - 1 (enabled)

This is a security setting that determines whether ASP scans the request queue at regular intervals to detect disconnected requests and remove them from the queue. This has the effect of not allowing disconnected requests to fill up the request queue. In addition, cleaning up the queue improves system efficiency.

**Note**

Because disabling this setting can make your system vulnerable to attack, it is strongly recommended that you leave this setting enabled.

**HangDetRequestThreshold**

Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\

Data Type: REG_DWORD

Default Value: 1000

Range: 0 - ffffffff

This setting determines whether the hang detection algorithm should be invoked based on the number of requests in the request queue. A value of 0 means that the hang detection algorithm is never invoked.

**HangDetThreadHungThreshold**

Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\

Data Type: REG_DWORD

Default Value: 50 (%)

Range: 1 - 100

This setting determines what percentage of the total number of ASP application threads must be unresponsive to fulfill one of the criteria for recycling the application.

| **HangDetConsecIllStatesThreshold** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 3 |
| Range: 0 - 3 |
| This setting helps prevent applications from being declared unhealthy and recycled because of transitory bursts in the application's request queue. If the number of times an application is consecutively detected to be in the "ill" state exceeds the **HangDetConsecIllStatesThreshold** value, AND if the application's request queue length has been increasing during the same period of time, ASP reports the application as being unhealthy, and IIS recycles the application. |

| **HangDetEnabled** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 1 (enabled) |
| Range: 0 (disabled) - 1 (enabled) |
| This setting determines whether or not ASP checks for and tracks unresponsive threads.<br><br>Once ASP hands over a page to a COM component for execution, it does not regain control unless the component returns a success or error. If a thread does not return within the timeout period specified on the **AspScriptTimeout** metabase property), it is considered unresponsive. Because there are a limited number of threads that can be executing at any one time on the system, each unresponsive thread means there is 1 less thread available to service a request queue. This results in a build up of the request queue unless the unresponsive threads are closed down.<br><br>ASP invokes the hang detection algorithm based on a combination of ASP registry and metabase request queue length settings. If the hang detection algorithm determines that the ratio of queued requests to the number of available threads high, ASP reports the application(s) with the unresponsive thread(s) as unhealthy, and IIS recycles the application(s). |

| **EnableChangeNotificationForUNC** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 0 (disabled) |
| Range: 0 (disabled) - 1 (enabled) |
| This setting determines whether ASP relies on the Change Notification mechanism for UNC to determine whether or not content is stale. The file change notification mechanism is not considered robust for tracking changes over UNC, and for this reason this setting is disabled. IIS and ASP use the Last Modified mechanism to track files instead. |

| **FileMonitoringEnabled** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |

| |
|---|
| Default Value: 1 (enabled) |
| Range: 0 (disabled) - 1 (enabled) |
| This setting determines whether files are registered for change notification. This setting is no longer used as the functionality has been encompassed in system logic and the **FileMonitoringEnabled** setting. Do not use this setting. |

| **FileMonitoringTimeoutSeconds** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 5 |
| Range: 0 - ffffffff |
| The last modified time stamp method of maintaining consistency in cache verifies if last modified time stamp of a file has changed since the file was last checked. This means that the check occurs on every request. This has a serious affect on performance. **FileMonitoringTimeoutSeconds** specfies the number of seconds that the system should return the cached copy since its last timestamp check.<br><br>Increasing the value of this key will improve performance, but can lead to returning stale content for the number of seconds specified in this key. This key applies to cases where the content is stored on a back end UNC or the local disk has had a huge content propagation and the change notification buffer has not been able to capture all the changes, thus pushing the system into a last modified algorithm for all files. You may want to tailor this key to suit your needs. |

| **MaxCSR** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: None. Determined dynamically. |
| Range: 0 - ffffffff |
| If COM+ is managing the CPU metric then this setting can be adjusted to tune COM+'s creation of the number of threads that ASP will use to process requests. COM+ tries to create threads on demand while staggering their creation so as not to put the system into an overloaded state. MaxCSR watches the context switch rate, and if the context switching in the system exceeds the rate specified in this key, COM+ stops spawning threads until the context switch rate falls below the value specified by this key. |

| **MaxCPU** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 100 |
| Range: 0 - 100 |

| |
|---|
| If COM+ is managing the CPU metric then this setting can be adjusted to tune COM+'s creation the number of threads that ASP uses to process requests, while staggering their creation so as not to put the system into an overloaded state. MaxCPU watches CPU utilization, and if the CPU utilization in the system exceeds the rate specified in this key, COM+ will stop spawning threads until CPU usage falls below the value specified by this key. |

| **DisableOOMRecycle** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ASP\Parameters\ |
| Data Type: REG_DWORD |
| Default Value: 0 (enabled) |
| Range: 0 (enabled) - 1 (disabled) |
| If ASP is not able to allocate memory for its internal processing, it assumes the system is out of memory and requests that the application be recycled, thereby freeing memory possibly held by a misbehaving application.<br><br>Allowing ASP to request recycling based on out of memory conditions must be considered carefully because of unintended consequences, such as loss of session state. |

# WWW Service Registry Entries

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to Service-Specific Registry Entries, the World Wide Web Publishing Service (WWW service) maintains the following entries.

| **Important** |
| --- |
| For a registry change to take effect, you must restart the service that is affected by the change. For global changes to take effect, you must restart all services. |

| **EnableTraceMethod** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters |
| Data Type: REG_DWORD |
| Default Value: 0 (disabled) |
| Range: 0 - 1 |
| Determines whether IIS recognizes the **HTTP TRACE** method. The **TRACE** method is used to invoke a remote, application-layer loop-back of a request message. **TRACE** allows a client to see what is being received at the other end of the request chain and use that data for testing and debugging information.<br>**Important**<br>Do not leave **EnableTraceMethod** enabled on a production system, as it can reveal back-end server address information to a malicious user. |

| **LogErrorRequests** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters |
| Data Type: REG_DWORD |
| Default Value: 1 (enabled) |
| Range: 0 - 1 |
| Determines whether to record failed activities in the log file. A value of 1 logs failed activities, and a value of 0 turns it off. |

| **LogSuccessfulRequests** |
| --- |
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters |
| Data Type: REG_DWORD |
| Default Value: 1 (enabled) |
| Range: 0 - 1 |

| |
|---|
| Determines whether to record successful activities in the log file. A value of 1 logs successful activities, and a value of 0 turns it off. |

| **MaxConcurrentCgisExecuting** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters |
| Data Type: REG_DWORD |
| Default Value: 256 |
| Range: 0 - 4294967295 |
| Determines the maximum number of concurrent CGIs allowed to run on the system. Additional CGI requests will be queued and then processed when the number of processing requests falls below the number set in this registry key. |

| **SSIEnableCmdDirective** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters |
| Data Type: REG_DWORD |
| Default Value: 0 (disabled) |
| Range: 0 - 1 |
| To execute shell commands, the **#exec cmd** directive of server-side includes is used. Security-conscious sites may want to disable the **#exec cmd** directive when untrusted parties are allowed to place files on the server by setting **SSIEnableCmdDirective** to 0 as an added security precaution. **SSIEnableCmdDirective** does not exist in the registry by default; to allow the **#exec cmd** directive to execute shell commands, you must first create **SSIEnableCmdDirective** and set it to 1. |

| **TryExceptDisable** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters |
| Data Type: REG_DWORD |
| Default Value: 0 (disabled) |
| Range: 0 - 1 |
| Disables exception caching when calling the **HttpExtensionproc** of an Internet Server API (ISAPI) application. Setting **TryExceptDisable** to 1 results in any exception that is caused by a misbehaving ISAPI application to stop the server, but it allows developers to do just-in-time debugging. Therefore, you should set **TryExceptDisable** to 1 only to debug ISAPI applications. |

| **UploadReadAhead** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters |
| Data Type: REG_DWORD |
| Default Value: 48KB |

| |
|---|
| Range: 0 - 0x80000000 |
| When a client posts data to the server, some data is read by the server and some by the application. **UploadReadAhead** sets the amount of data the server will read before passing control to the application. The application is then responsible for reading the rest of the data. If you increase this size, you should increase the amount of RAM on the server. |

| **UsePoolThreadForCGI** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters |
| Data Type: REG_DWORD |
| Default Value: 1 (enabled) |
| Range: 0 - 1 |
| IIS, by default, uses a server pool thread to do Common Gateway Interface (CGI) processing. This means that CGI requests that take an extended period of time can consume a server pool thread. |

| **UseSharedWPDesktop** |
|---|
| Registry Path: HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\W3SVC\Parameters |
| Data Type: REG_DWORD |
| Default Value: 0 |
| Range: 0 - 1 |
| If you are setting up your application pools with unique identities, then depending on the applications and memory resources on your server, you will reach a limit of about 60 application pools. There are finite limits of some system resources that get allocated with each new logon session. This means that 60 processes can run concurrently as distinct accounts. IIS 6.0 supports running these processes in a single shared workstation and desktop, at a cost of sharing a single encapsulation of a user session among all parties. To scale beyond 60 application pools and to share a single desktop, change the **UseSharedWPDesktop** to a DWORD value of 1. After changing this registry key, you should be able to scale to hundreds of application pools and hundreds of concurrently running worker processes. |

# FTP Service Registry Entries

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In addition to Service-Specific Registry Entries,the FTP service maintains the following value entries.

**Details**

| **Important** |
| --- |
| For a registry change to take effect, you must restart the service that is affected by the change. For global changes to take effect, you must restart all services. |

**AnnotateDirectories**

| | |
| --- | --- |
| Registry Path | HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSFTPSVC\Parameters |
| Data Type | REG_DWORD |
| Default Value | 0 (disabled) |
| Range | 0,1 |
| Description | The FTP service supports annotating a directory with custom messages. The annotation text is stored in a file named ~ftpsvc~.ckm in the directory to be annotated. If ~ftpsvc~.ckm exists in the target directory of a Change Directory (CWD - Change Working Directory) FTP operation, the FTP service responds with the contents of this file for the operation. This provides a way for administrators to add custom messages for directories under consideration. By default, the service is configured to not send annotation text. If you choose to add a custom message, you should also create the annotation file. Also, you should make the annotation file a hidden file so that the file does not appear on a directory listing. |

**AnonymousHomeDirInfoGrace**

| | |
| --- | --- |
| Registry Path | HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSFTPSVC\Parameters\ |
| Data Type | REG_DWORD |
| Default Value | 1 |
| Range | 0 - 4,294,967,295 |
| Description | This entry specifies the number of seconds (grace period) during which the cached value for anonymous home directory will still be used after the interval in AnonymousHomeDirInfoInterval expires. After the grace period expires, the entries in the cache are updated and the requests will receive the new updated entries. The first request made after that time will update the cached entry. Although the range of values allows up to the maximum for DWORD properties, it is recommended that you keep this value low. |

**AnonymousHomeDirInfoInterval**

| | |
| --- | --- |
| Registry Path | HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSFTPSVC\Parameters\ |
| Data Type | REG_DWORD |
| Default Value | 10 |

| | |
|---|---|
| Range | 0 - 4,294,967,295 |
| Description | This entry specifies the number of seconds during which the entry of the anonymous home directory is cached. The first anonymous login made after that time will update the cached entry. Although the range of values allows up to the maximum for DWORD properties, you should consider potential changes in the home directory path when you decide on the value set for this registry key. |

**CacheAnonymousHomeDirInfo**

| | |
|---|---|
| Registry Path | HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSFTPSVC\Parameters\ |
| Data Type | REG_DWORD |
| Default Value | 0 (disabled) |
| Range | 0 to 4,294,967,295. Any DWORD value other than 0 will cause the FTP server to cache the anonymous home directory. |
| Description | This entry enables caching of the anonymous home directory. Caching the anonymous home directory helps to avoid having the FTP server refuse anonymous connections when the home directory resides on an UNC share and too many users are already connected. In such cases, you may see FTP event message number 8: "FTP Server could not create a client worker thread for user at host %1. The connection to this user is terminated. The data is the error."<br><br>When this key is enabled, the cache will have a 10 second refresh interval and a 1 second grace period, both of which are configurable--respectively--in the registry keys AnonymousHomeDirInfoInterval and AnonymousHomeDirInfoGrace. |

**EnablePortAttack**

| | |
|---|---|
| Registry Path | HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSFTPSVC\Parameters |
| Data Type | REG_DWORD |
| Default Value | 0 (disabled) |
| Range | 0,1 |
| Description | This entry is disabled by default to prevent a security problem in the FTP protocol specification. Because EnablePortAttack is disabled by default, you should enable this flag if you want users to connect by using other ports as specified in the FTP RFC. The FTP service specification allows passive connections to be established based on the port address given by the client. This can allow malicious users to execute destructive commands in the FTP service. The problem occurs when the FTP service makes a connection to a client using a port other than FTP Data port (20) and that port number is less than IP_PORT_RESERVED (1024). EnablePortAttack controls whether or not such an attack should be allowed. By default, the service does not make any connections to port numbers lower than IP_PORT_RESERVED (other than 20). |

**LowercaseFiles**

| | |
|---|---|
| Registry Path | HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSFTPSVC\Parameters\ |
| Data Type | REG_DWORD |
| Default Value | 0 (disabled) |
| Range | 0,1 |

| | |
|---|---|
| Description | The FTP service uses uppercase for file names. In order for comparisons with case-sensitive file systems to work, you need to ensure that the exact file names are compared. Administrators can add LowercaseFiles to the registry to ensure that the FTP service also uses lowercase for such comparisons. |

**MaxPassiveAcceptEvents**

| | |
|---|---|
| Registry Path | HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSFTPSVC\Parameters\ |
| Data Type | REG_DWORD |
| Default Value | 60 |
| Range | 1 through 300. Any number is rounded up to the nearest multiple of 60 and capped off at 300. |
| Description | This registry value is intended for advanced users. When FTP clients initiate passive file transfer operations, the FTP server uses asynchronous network notification to wait for the actual client connection to a passive data socket. By default, the FTP server allows 60 passive file transfer operations to be queued simultaneously. Once the queue reaches 60, subsequent client requests receive error 451: Too many Passive Users. Client connections remain in this transient state for a short time, between issuing a file transfer command (e.g. RETR, STOR) and connecting the data channel to the passive socket. For this reason, this error can only occur in one of the following scenarios:<br>1. A DoS attack has issued concurrent file transfer commands without actually connecting to the data socket.<br>2. An extremely large number of concurrent users are transferring files in passive mode (which is possible on slow networks).<br>Consistent failure with the above error may indicate a DoS attack. However, sporadic failures at peak usage times may indicate valid usage of this registry value. |

# Events Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS generates events in Event Viewer so that you can verify the performance of IIS. Event Viewer tracks (in order of importance) error events, warning events, and informational events. The logs in Event Viewer provide an audited record of all services and processes in the Microsoft® Windows® Server 2003 operating system, such as logon and connection times.

**In This Section**

- About Events: Provides conceptual topics related to Windows Events.
- ASP Errors: Specifies errors generated by the World Wide Web Publishing Service (WWW service) when processing ASP applications that are configured incorrectly.
- ASP Events: Specifies events generated from Active Server Pages.
- FTP Service Events: Specifies events generated by the FTP service.
- ODBC Driver Errors Returned by IIS: Specifies events caused by Web pages that contain ODBC driver errors.
- Server-Side Includes Events: Specifies events generated by the part of the WWW service that processes server-side include directives in Web pages.
- Core WWW Service Events: Specifies events generated from Web services.
- WWW Performance Counter Events: Specifies events generated by the part of the WWW service that monitors and records performance counter data.
- WWW Service Events: Specifies events generated by the part of the WWW service that handles internal administration of the IIS server.

**Related Information**

- For information about error messages displayed in the browser of a client, see Error Messages in IIS 6.0.
- For information about logging client requests, see Logging Site Activity.
- For information about viewing and interpreting log files, see Analyzing Log Files.
- For information about performance monitoring and tuning, see Optimizing IIS 6.0 Performance.

# About Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Logging begins automatically when Windows is started. There are three types of logs: system, security, and application. For each type of log, you can filter the events to be viewed, designate the number of entries to view, specify how long to save entries, and specify whether to automatically overwrite existing events when the log becomes full.

You can open Event Viewer in any one of the following ways:

- From the **Start** menu, click **Run**. In the **Open** box, type **eventvwr**. Click **OK**.

- or -

- Type **start eventvwr** at the command prompt, and press ENTER.

- or -

- From the **Start** menu, point to **Administrative Tools**, and then click **Event Viewer**.

- or -

- From the **Start** menu, point to **Administrative Tools**, and then click **Computer Management**. **Event Viewer** is listed under the **System Tools** node. The benefit of this display is that IIS Manager is in the same window under the **Services and Applications** node.

**ASP Errors and the Windows Events Log**

For more information, see **AspLogErrorRequests** and **AspErrorsToNTLog**. **AspLogErrorRequests** can be set in IIS Manager, or it can be set programmatically using ADSI, WMI, or ABO. See Code Examples to Configure Metabase Properties for instructions on how to programmatically change metabase properties.

**AspErrorsToNTLog** cannot be set in IIS Manager, and must be set programmatically.

**Related Information**

- For instructions on how to programmatically change metabase properties, see Code Examples to Configure Metabase Properties.

- For descriptions of ASP Event errors, see ASP Events.

- For more information about the metabase and programmatic administration of IIS, see Administering IIS Programmatically.

# ASP-Related Errors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This topic lists the errors that are generated by the World Wide Web Publishing Service (WWW service) when processing ASP applications that are configured incorrectly or ASP pages that contain errors. These errors are logged to the Windows Event application log only if the AspLogErrorRequests metabase property is set to **true**. If the AspErrorsToNTLog metabase property is also set to **true**, then all of the errors get logged to the Windows Event application log; otherwise, only a subset are logged. The subset includes the error numbers 100, 101, 102, 103, 104, 105, 106, 107, 115, 190, 191, 192, 193, 194, 240, 241, and 242.

**Details**

| **Note** |
| --- |
| When messages appear in the event log or Web browser, the symbols that appear below (such as \|, %s, %X, %1, %2, %2!.20s!, and so on) are replaced with details about the error at the time that the error actually occurs. |

| Event ID | Description |
| --- | --- |
| 0100 | Out of memory. Unable to allocate required memory. |
| 0101 | Unexpected error. The function returned \|. |
| 0102 | Expecting string input. The function expects a string as input. |
| 0103 | Expecting numeric input. The function expects a number as input. |
| 0104 | Operation not Allowed. |
| 0105 | Index out of range. An array index is out of range. |
| 0106 | Type Mismatch. An unhandled data type was encountered. |
| 0107 | Data size too large. Size of data being sent in the request is over the allowed limit. |
| 0108 | Create object failed. An error occurred while creating object '%s'. |
| 0109 | Member not found. |
| 0110 | Unknown name. |
| 0111 | Unknown interface. |
| 0112 | Missing parameter. |
| 0113 | Script timed out. The maximum amount of time for a script to execute was exceeded. You can change this limit by specifying a new value for the property Server.ScriptTimeout or by changing the value in the IIS administration tools. |
| 0114 | Object not free threaded. The application object accepts only free threaded objects; object '%s' is not free threaded. |
| 0115 | Unexpected error. A trappable error (%X) occurred in an external object. The script cannot continue running. |
| 0116 | Missing close of script delimiter. The Script block lacks the close of script tag (%>). |
| 0117 | Missing close of script tag. The Script block lacks the close of script tag (</SCRIPT>) or close of tag symbol (>). |
| 0118 | Missing close of object tag. The Object block lacks the close of object tag (</OBJECT>) or close of tag symbol (>). |
| 0119 | Missing Classid or Progid attribute. The object instance '\|' requires a valid Classid or Progid in the object tag. |
| 0120 | Invalid Runat attribute. The Runat attribute of the Script tag or Object tag can only have the value 'Server'. |

| | |
|---|---|
| 0121 | Invalid Scope in object tag. The object instance '\|' cannot have Application or Session scope. To create the object instance with Session or Application scope, place the Object tag in the Global.asa file. |
| 0122 | Invalid Scope in object tag. The object instance '\|' must have Application or Session scope. This applies to all objects created in a Global.asa file. |
| 0123 | Missing Id attribute. The required Id attribute of the Object tag is missing. |
| 0124 | Missing Language attribute. The required Language attribute of the Script tag is missing. |
| 0125 | Missing close of attribute. The value of the '\|' attribute has no closing delimiter. |
| 0126 | Include file not found. The include file '\|' was not found. |
| 0127 | Missing close of HTML comment. The HTML comment or server-side include lacks the close tag (-->). |
| 0128 | Missing File or Virtual attribute. The Include file name must be specified using either the File or Virtual attribute. |
| 0129 | Unknown scripting language. The scripting language '\|' is not found on the server. |
| 0130 | Invalid File attribute. File attribute '\|' cannot start with forward slash or back slash. |
| 0131 | Disallowed Parent Path. The Include file '\|' cannot contain '..' to indicate the parent directory. |
| 0132 | Compilation Error. The Active Server Page '\|' could not be processed. |
| 0133 | Invalid ClassID attribute. The object tag has an invalid ClassID of '\|'. |
| 0134 | Invalid ProgID attribute. The object has an invalid ProgID of '\|'. |
| 0135 | Cyclic Include. The file '\|' is included by itself (perhaps indirectly). Please check include files for other Include statements. |
| 0136 | Invalid object instance name. The object instance '\|' is attempting to use a reserved name. This name is used by Active Server Pages intrinsic objects. |
| 0137 | Invalid Global Script. Script blocks must be one of the allowed Global.asa procedures. Script directives within <% ... %> are not allowed within the Global.asa file. The allowed procedure names are Application_OnStart, Application_OnEnd, Session_OnStart, or Session_OnEnd. |
| 0138 | Nested Script Block. A script block cannot be placed inside another script block. |
| 0139 | Nested Object. An object tag cannot be placed inside another object tag. |
| 0140 | Page Command Out Of Order. The @ command must be the first command within the Active Server Page. |
| 0141 | Page Command Repeated. The @ command can only be used once within the Active Server Page. |
| 0142 | Thread token error. A thread token failed to open. |
| 0143 | Invalid Application Name. A valid application name was not found. |
| 0144 | Initialization Error. The page level objects list failed during initialization. |
| 0145 | New Application Failed. The new Application could not be added. |
| 0146 | New Session Failed. The new Session could not be added |
| 0147 | 500 Server Error. |
| 0148 | Server Too Busy. |
| 0149 | Application Restarting. The request cannot be processed while the application is being restarted. |
| 0150 | Application Directory Error. The Application directory could not be opened. |

| | |
|---|---|
| 0151 | Change Notification Error. The change notification event could not be created. |
| 0152 | Security Error. An error occurred while processing a user's security credentials. |
| 0153 | Thread Error. A new thread request failed. |
| 0154 | Write HTTP Header Error. The HTTP headers could not be written to the client browser. |
| 0155 | Write Page Content Error. The page content could not be written to the client browser. |
| 0156 | Header Error. The HTTP headers are already written to the client browser. Any HTTP header modifications must be made before writing page content. |
| 0157 | Buffering On. Buffering cannot be turned off once it is already turned on. |
| 0158 | Missing URL. A URL is required. |
| 0159 | Buffering Off. Buffering must be on. |
| 0160 | Logging Failure. Failure to write entry to log. |
| 0161 | Data Type Error. The conversion of a Variant to a String variable failed. |
| 0162 | Cannot Modify Cookie. The cookie 'ASPSessionID' cannot be modified. It is a reserved cookie name. |
| 0163 | Invalid Comma Use. Commas cannot be used within a log entry. Please select another delimiter. |
| 0164 | Invalid TimeOut Value. An invalid TimeOut value was specified. |
| 0165 | SessionID Error. A SessionID string cannot be created. |
| 0166 | Uninitialized Object. An attempt was made to access an uninitialized object. |
| 0167 | Session Initialization Error. An error occurred while initializing the Session object. |
| 0168 | Disallowed object use. An intrinsic object cannot be stored within the Session object. |
| 0169 | Missing object information. An object with missing information cannot be stored in the Session object. The threading model information for an object is required. |
| 0170 | Delete Session Error. The Session did not delete properly. |
| 0171 | Missing Path. The Path parameter must be specified for the MapPath method. |
| 0172 | Invalid Path. The Path parameter for the MapPath method must be a virtual path. A physical path was used. |
| 0173 | Invalid Path Character. An invalid character was specified in the Path parameter for the MapPath method. |
| 0174 | Invalid Path Character(s). An invalid '/' or '\\' was found in the Path parameter for the MapPath method. |
| 0175 | Disallowed Path Characters. The '..' characters are not allowed in the Path parameter for the MapPath method. |
| 0176 | Path Not Found. The Path parameter for the MapPath method did not correspond to a known path. |
| 0177 | Server.CreateObject Failed. %s |
| 0178 | Server.CreateObject Access Error. The call to Server.CreateObject failed while checking permissions. Access is denied to this object. |
| 0179 | Application Initialization Error. An error occurred while initializing the Application object. |
| 0180 | Disallowed object use. An intrinsic object cannot be stored within the Application object. |
| 0181 | Invalid threading model. An object using the apartment threading model cannot be stored within the Application object. |

| | |
|---|---|
| 0182 | Missing object information. An object with missing information cannot be stored in the Application object. The threading model information for the object is required. |
| 0183 | Empty Cookie Key. A cookie with an empty key cannot be stored. |
| 0184 | Missing Cookie Name. A name must be specified for a cookie. |
| 0185 | Missing Default Property. A default property was not found for the object. |
| 0186 | Error parsing certificate |
| 0187 | Object addition conflict. Could not add object to application. Application was locked down by another request for adding an object. |
| 0188 | Disallowed object use. Cannot add objects created using object tags to the session intrinsic. |
| 0189 | Disallowed object use. Cannot add objects created using object tags to the application intrinsic. |
| 0190 | Unexpected error. A trappable error occurred while releasing an external object. |
| 0191 | Unexpected error. A trappable error occurred in the OnStartPage method of an external object. |
| 0192 | Unexpected error. A trappable error occurred in the OnEndPage method of an external object. |
| 0193 | OnStartPage Failed. An error occurred in the OnStartPage method of an external object. |
| 0194 | OnEndPage Failed. An error occurred in the OnEndPage method of an external object. |
| 0195 | Invalid Server Method Call. This method of the Server object cannot be called during Session_OnEnd and Application_OnEnd. |
| 0196 | Cannot launch out of process component. Only InProc server components should be used. If you want to use LocalServer components, you must set the AspAllowOutOfProcComponents metabase setting. Please consult the help file for important considerations. |
| 0197 | Disallowed object use. Cannot add object with apartment model behavior to the application intrinsic object. |
| 0198 | Server shutting down. Cannot process request |
| 0199 | Disallowed object use. Cannot add JScript objects to the session. |
| 0200 | Out of Range 'Expires' attribute. The date and time given for 'Expires' precedes Jan 1, 1980 or excedes Jan 19, 2038, 3:14:07 GMT. |
| 0201 | Invalid Default Script Language. The default script language specified for this application is invalid. |
| 0202 | Missing Code Page. The code page attribute is missing. |
| 0203 | Invalid Code Page. The specified code page attribute is invalid. |
| 0204 | Invalid CodePage Value. An invalid CodePage value was specified. |
| 0205 | Change Notification. Failed to create event for change notification. |
| 0206 | Cannot call BinaryRead. Cannot call BinaryRead after using Request.Form collection. |
| 0207 | Cannot use Request.Form. Cannot use Request.Form collection after calling BinaryRead. |
| 0208 | Cannot use generic Request collection. Cannot use the generic Request collection after calling BinaryRead. |
| 0209 | Illegal value for TRANSACTION property. The TRANSACTION property can only be REQUIRED, REQUIRES_new, SUPPORTED or NOT_SUPPORTED. |
| 0210 | Method not implemented. This method has not yet been implemented. |

| | |
|---|---|
| 0211 | Object out of scope. A built-in ASP object has been referenced, which is no longer valid. |
| 0212 | Cannot Clear Buffer. Response.Clear is not allowed after a Response.Flush while Client Debugging is Enabled. |
| 0214 | Invalid Path parameter. The Path parameter excedes the maximum length allowed. |
| 0215 | Illegal value for ENABLESESSIONSTATE property. The ENABLESESSIONSTATE property can only be TRUE or FALSE. |
| 0216 | MSDTC Service not running. Transactional web pages cannot be run if the MSDTC service is not running. |
| 0217 | Invalid Scope in object tag. Object scope must be Page, Session or Application. |
| 0218 | Missing LCID. The LCID attribute is missing. |
| 0219 | Invalid LCID. The specified LCID is not available. |
| 0220 | Requests for GLOBAL.ASA Not Allowed. Requests with the URL pointing to GLOBAL.ASA are not allowed. |
| 0221 | Invalid @ Command directive. The specified '\|' option is unknown or invalid. |
| 0222 | Invalid TypeLib Specification. METADATA tag contains an invalid Type Library specification. |
| 0223 | TypeLib Not Found. METADATA tag contains a Type Library specification that does not match any Registry entry. |
| 0224 | Cannot load TypeLib. Cannot load Type Library specified in the METADATA tag. |
| 0225 | Cannot wrap TypeLibs. Cannot create a Type Library Wrapper object from the Type Libraries specified in METADATA tags. |
| 0226 | Cannot modify StaticObjects. StaticObjects collection cannot be modified at run time. |
| 0227 | Server.Execute Failed. The call to Server.Execute failed |
| 0228 | Server.Execute Error. The call to Server.Execute failed while loading the page. |
| 0229 | Server.Transfer Failed. The call to Server.Transfer failed |
| 0230 | Server.Transfer Error. The call to Server.Transfer failed while loading the page. |
| 0231 | Server.Execute Error. Invalid URL form or fully-qualified absolute URL was used. Use relative URLs. |
| 0232 | Invalid Cookie Specification. METADATA tag contains an invalid cookie specification. |
| 0233 | Cannot load cookie script source. Cannot load cookie script source file specified in the METADATA tag. |
| 0234 | Invalid include directive. Server side include directives may not be present in script blocks. Please use the SRC= attribute of the <SCRIPT> tag. |
| 0235 | Server.Transfer Error. Invalid URL form or fully-qualified absolute URL was used. Use relative URLs. |
| 0236 | Invalid Cookie Specification. METADATA tag contains an invalid or missing SRC parameter. |
| 0237 | Invalid Cookie Specification. METADATA tag contains an invalid or missing NAME parameter. |
| 0238 | Missing attribute value. No value was specified for the '\|' attribute. |
| 0239 | Cannot process file. UNICODE ASP files are not supported. |
| 0240 | Script Engine Exception. A ScriptEngine threw exception '%X' in '%s' from '%s'. |
| 0241 | CreateObject Exception. The CreateObject of '%s' caused exception %X. |
| 0242 | Query OnStartPage Interface Exception. Querying Object '%s''s OnStartPage or OnEndPage methods caused exception %X. |

| | |
|---|---|
| 0243 | Invalid METADATA tag in Global.asa. Only METADATA TYPE="TypeLib" may be used in Global.asa. |
| 0244 | Cannot Enable Session State. Session state cannot be enabled when it has been disabled in the application. |
| 0245 | Mixed usage of Code Page values. The @CODEPAGE value specified differs from that of the including file's CODEPAGE or the file's saved format. |
| 0246 | Too many concurrent users. Please try again later.. |
| 0247 | Bad Argument to BinaryRead. The argument to BinaryRead must be non-negative. |
| 0248 | Script isn't transacted. This ASP file must be transacted in order to use the ObjectContext object. |
| 0249 | Cannot use IStream on Request. Cannot use IStream on Request object after using Request.Form collection or Request.BinaryRead. |
| 0250 | Invalid Default Code Page. The default code page specified for this application is invalid. |
| 0251 | Response Buffer Limit Exceeded. Execution of the ASP page caused the Response Buffer to exceed its configured limit. |

# ASP-Related Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This topic lists the events that are generated by the World Wide Web Publishing Service (WWW service) when processing ASP applications that are configured incorrectly or ASP pages that contain errors. These events are logged to the Windows Event application log regardless of the value of the AspLogErrorRequests Metabase Property.

| **Note** |
| --- |
| When messages appear in the event log or Web browser, the symbols that appear below (such as \|, %s, %X, %1, %2, %2!.20s!, and so on) are replaced with details about the error at the time that the error actually occurs. |

| Event ID | Description |
| --- | --- |
| 0x10 | Warning: %1, %2. |
| 0x11 | Warning: %1, %2, %3. |
| 0x12 | Warning: %1, %2, %3, %4. |
| 0x13 | Failed to create an ASP Session due to Access Denied when activating COM+ Partition %1. |
| 0x14 | Failed to create ASP Application %1 due to invalid or missing COM+ Partition ID. |
| 0x15 | Failed to retrieve the Anonymous User Token for ASP Application%1. Global.ASA OnEnd routines will not be executed. |
| 0x16 | Failed to impersonate the Anonymous User for ASP Application%1. Global.ASA OnEnd routines will not be executed. |
| 0x17 | No Default Script Language specified for Application %s. Using default: %s. |
| 0x1f3 | Invalid ProgID. |
| 0x5 | Error: %1. |
| 0x6 | Error: %1, %2. |
| 0x7 | Error: %1, %2, %3. |
| 0x8 | Error: %1, %2, %3, %4. |
| 0x9 | Warning: %1. |

# FTP Service-Related Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This topic lists the events that are generated by the FTP service.

| **Note** |
| --- |
| When messages appear in the event log or Web browser, the symbols that appear in the documentation below (such as \|, %s, %X, %1, %2, %2!. 20s!, and so on) are replaced with details about the error at the time that the error actually occurs. |

| Event ID | Description |
| --- | --- |
| 1 | FTP Server could not initialize its security. The data is the error. |
| 3 | FTP Server could not initialize the socket library. The data is the error. |
| 4 | FTP Server was unable to initialize due to a shortage of available memory. The data is the error. |
| 5 | FTP Server could not locate the FTP/TCP service. The data is the error. |
| 6 | FTP Server could not create the main connection socket. The data is the error. |
| 7 | FTP Server could not create the main connection thread. The data is the error. |
| 8 | FTP Server could not create a client worker thread for user at host %1. The connection to this user is terminated. The data is the error. |
| 9 | A call to a system service failed unexpectedly. The data is the error. |
| 10 | User %1 at host %2 has timed-out after %3 seconds of inactivity. |
| 11 | Anonymous logon request received from %1 at host %2. |
| 12 | User logon request received from %1 at host %2. |
| 13 | User %1 failed to log on, could not access the home directory%2. |
| 14 | User %1 denied access to the current directory %2 due to a security change. |
| 15 | Failed to read metabase property **PassivePortRange** under the FTP Service node %1. |
| 16 | The FTP Service configuration property **PassivePortRange** has invalid value: %1. |
| 14401 | The server failed to create a data channel connection to a client on local interface *ip_address*. The FTP data port *port_number* may be in use by another service or application. |

# ODBC Driver Errors Returned by IIS

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This topic lists the events that are caused by Web pages that contain ODBC driver errors.

| **Note** |
| --- |
| When errors appear in the IIS error log, the symbols that appear below (such as \|, %s, %X, %1, %2, %2!. 20s!, and so on) are replaced with details about the error at the time that the error actually occurs. |

| Event ID | Description |
| --- | --- |
| 2602 | The form you just filled in requires the %1 entry to be filled in. Please fill in that entry in the form and resubmit the form. |
| 2603 | The query file could not be opened. The file may not exist or you may have insufficient permission to open the file. |
| 2605 | The template file contains an **else** tag without a matching **if** tag. |
| 2606 | The template file contains an **endif** tag without a matching **if** tag. |
| 2607 | The template file contains an expression with parameters that are of different type. For example, an **if** tag comparing a string with a number would cause this error. |
| 2608 | The template file contains an expression using the CONTAINS operator where one or both parameters are not strings. CONTAINS is only valid with string values. |
| 2609 | The template file contains an expression that uses a quoted string but the close quote is missing. |
| 2610 | The template file contains an expression with an unrecognized operator. Valid operators are GT, LT, EQ and CONTAINS. |
| 2611 | The query file contains an unrecognized field %1. |
| 2612 | The query file must contain a valid, non-empty Datasource: field and SQLStatement: field. |
| 2613 | The library odbc32.dll could not be loaded to perform the query. Make sure ODBC has been properly installed on the server. |
| 2615 | Error performing query. %1 |
| 2616 | The maximum number of SQLStatement fields in the IDC file has been exceeded. |

# Server-Side Includes Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This topic lists the events that are generated by the part of the World Wide Web Publishing Service (WWW service) that processes server-side include directives in Web pages.

| **Note** |
| --- |
| When messages appear in the event log or Web browser, the symbols that appear below (such as \|, %s, %X, %1, %2, %2!.20s!, and so on) are replaced with details about the error at the time that the error actually occurs. |

| Event ID | Description |
| --- | --- |
| 2501 | Error processing SSI file. |
| 2502 | Invalid SSI tag. |
| 2503 | Cannot perform flastmod(): Win32 Error Code = %2!.20s! |
| 2504 | Cannot perform fsize(): Win32 Error Code = %2!.20s! |
| 2505 | Variable '%1!.40s!' cannot be found. |
| 2506 | Variable '%1!.40s!' cannot be evaluated. |
| 2507 | SSI feature not supported |
| 2508 | Error handling SSI File : Win32 Error Code = %2!.20s! |
| 2509 | Failed to process SSI file. |
| 2510 | Failed to execute CMD : Win32 Error Code = %2!.20s! |
| 2511 | Failed to execute CGI: Win32 Error Code = %2!.20s! |
| 2512 | Failed to execute ISAPI extension: Win32 Error Code =%2!.20s! |
| 2513 | Cannot #EXEC file due to lack of EXECUTE permission. |
| 2514 | Cannot process file due to access denied. |
| 2515 | A server side include file has included itself or the maximum depth of server side includes has been exceeded. |
| 2516 | The CMD option is not enabled for #EXEC calls. |
| 2517 | #EXEC calls have been disabled for this virtual path. |
| 2518 | There is a recursive #INCLUDE chain. |
| 2519 | Failed to execute script: Http Status Code = %2!.20s! |

# Core WWW Service Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This topic lists the events that are generated by the core processes of the World Wide Web Publishing Service (WWW service) such as authentication and authorization, application problems, memory monitoring, and so on.

| **Note** |
| --- |
| When messages appear in the event log or Web browser, the symbols that appear below (such as \|, %s, %X, %1, %2, %2!.20s!, and so on) are replaced with details about the error at the time that the error actually occurs. |

| Event ID | Description |
| --- | --- |
| 2201 | HTTP Server could not initialize its security. The data is the error. |
| 2203 | HTTP Server could not initialize the socket library. The data is the error. |
| 2204 | HTTP Server was unable to initialize due to a shortage of available memory. The data is the error. |
| 2206 | HTTP Server could not create the main connection socket. The data is the error. |
| 2208 | HTTP Server could not create a client connection object for user at host %1. The connection to this user is terminated. The data is the error. |
| 2214 | The HTTP Filter DLL %1 failed to load. The data is the error. |
| 2216 | The script started from the URL '%1' with parameters '%2' has not responded within the configured timeout period. The HTTP server is terminating the script. |
| 2218 | The HTTP server encountered an error processing the server side include file '%1'.The error was '%2'. |
| 2219 | The HTTP server encountered an unhandled exception while processing the ISAPI Application '%1'. |
| 2220 | The HTTP server was unable to load the ISAPI Application '%1'.The data is the error. |
| 2221 | A server side include file has included itself or the maximum depth of server side includes has been exceeded. |
| 2222 | An attempt was made to load filter '%1' on a server instance but it requires the SF_NOTIFY_READ_RAW_DATA filter notification so it must be loaded as a global filter. |
| 2223 | For compatibility with previous versions of IIS, the filter'%1' was loaded as a global filter from the registry. To control the filter with the Internet Service Manager, remove the filter from the registry and add it as a global filter with Internet Service Manager. Filters in the registry are stored at "HKLM\System\CurrentControlSet\Services\W3Svc\Parameters\FilterDLLs". |
| 2226 | The server was unable to read the file %1 due to a lack of access permissions. |
| 2227 | The server was unable to acquire a license for a SSL connection. |
| 2228 | The server stop serving requests for application '%1' because the number of Out of Process component crashes exceed a limit. |

| | |
|---|---|
| 2229 | The server failed to shut down application '%1'. The error was'%2'. |
| 2230 | The server was unable to read the file %1. The file does not exist. |
| 2231 | The server was unable to read the file %1. The Windows 32 error returned from the attempt is %2. |
| 2232 | The server was unable to read the file %1. The file exceeds the maximum allowable size of %2. |
| 2233 | The server was unable to allocate a buffer to read the file%1. |
| 2236 | The server failed to load application '%1'. The error was'%2'. |
| 2237 | Out of process application '%1' terminated unexpectedly. |
| 2238 | Job object query failed. The data is the error. |
| 2239 | Job object set limit failed. The data is the error. |
| 2240 | Schedule a work item for CPU Accounting or Limits failed. The data is the error. |
| 2241 | Site '%1' hit its CPU Limit. No action was taken. |
| 2242 | Site '%1' hit its CPU Limit. The priority of all process on that site has been lowered to idle class. |
| 2243 | Site '%1' hit its CPU Limit. All processes on that site have been terminated. |
| 2244 | Site '%1' hit its CPU Limit. The site has been paused. |
| 2245 | The server certificate for instance '%1' has expired or is not yet valid. |
| 2246 | The server certificate for instance '%1' has been revoked. |
| 2247 | The server certificate for instance '%1' does not chain up to a trusted root certificate. |
| 2248 | One of the certificates in the certificate chain of the server certificate for instance '%1' has an invalid signature. |
| 2249 | CPU Limits/Accounting failed to queue work item. The data is the error. |
| 2250 | The server certificate for instance '%1' had invalid metabase data associated with it and could not be retrieved; the error encountered was '%2'. |
| 2251 | A CryptoAPI error was encountered trying to retrieve the server certificate for instance '%1'; the error encountered was '%2'. |
| 2252 | The server certificate for instance '%1' could not be retrieved because it could not be found in a certificate store; the error encountered was '%2' |
| 2253 | The server certificate for instance '%1' could not be retrieved due to an internal error; the error encountered was '%2'. |
| 2254 | The server Certificate Trust List for instance '%1' had invalid metabase data associated with it and could not be retrieved; the error encountered was '%2'. |
| 2255 | A CryptoAPI error was encountered trying to retrieve the server Certificate Trust List for instance '%1'; the error encountered was'%2'. |

| | |
|---|---|
| 2256 | The server Certificate Trust List for instance '%1' could not be retrieved because it could not be found in a certificate store; the error encountered was '%2'. |
| 2257 | The server certificate for instance '%1' could not be retrieved due to an internal error; the error number is '%2'. |
| 2258 | The server failed to close client connections to the following URLs during shutdown: '%1'. |
| 2259 | The COM Application '%1' at '%2' failed to activate out of process. |
| 2261 | An attempt was made to load filter '%1' but it requires the SF_NOTIFY_READ_RAW_DATA filter notification and this notification is not supported in Worker Process Isolation Mode. |
| 2262 | ISAPI '%1' reported itself as unhealthy for the following reason: '%2'. |
| 2263 | ISAPI '%1' reported itself as unhealthy. No reason was given by the ISAPI. |
| 2264 | The directory specified for caching compressed content %1 is invalid. Static compression is being disabled. |
| 2265 | The registry key for IIS subauthenticator is not configured correctly on local machine, the anonymous password sync feature is disabled. |
| 2266 | The account that the current worker process is running under does not have SeTcbPrivilege privilege, the anonymous password sync feature and the Digest authentication feature are disabled. |
| 2267 | Raw ISAPI Filter could not be loaded due to configuration problem. HTTPFilter Service has to be hosted in Inetinfo.exe for Raw ISAPI filters to work correctly. The most likely reason is that IIS5IsolationModeEnabled setting was changed and HTTPFilter service was not restarted. Restart HTTPFilter for configuration changes to be taken into effect. |
| 2268 | Could not load all ISAPI filters for site/service. Therefore startup aborted. |
| 2269 | The worker process failed to initialize the HTTP.sys communication or the w3svc communication layer and therefore could not be started. The data field contains the error number. |
| 2270 | IIS Client certificate mapping configuration for site %1 failed to be loaded. The data field contains the error number. |
| 2271 | Could not initialize the logging module for site %1. The site will therefore be non-functional. |
| 2272 | The registry key for IIS subauthenticator is not configured correctly on domain controller, the Digest feature is disabled. |

# WWW Performance Counter Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This topic lists the events that are generated by the part of the World Wide Web Publishing Service (WWW service) that monitors and records performance counter data.

| **Note** |
| --- |
| When messages appear in the event log or Web browser, the symbols that appear below (such as \|, %s, %X, %1, %2, %2!.20s!, and so on) are replaced with details about the error at the time that the error actually occurs. |

| Event ID | Description |
| --- | --- |
| 2000 | Unable to query the W3SVC (HTTP) service performance data. The error code returned by the service is data DWORD 0. |
| 2001 | It has taken too long to refresh the W3SVC counters, the stale counters are being used instead. |
| 2002 | Setting up Web Service counters failed, please make sure your Web Service counters are registered correctly. |
| 2003 | It has taken too long to refresh the W3SVC counters, the stale counters are being used instead. This is the second message within the past %1 (hours, minutes, seconds). No more stale counter messages will be logged for this client session until the time limit expires. %n For additional information specific to this message please visit the Microsoft Help and Support site. |

# WWW Service-Related Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This topic lists the events that are generated by the part of the World Wide Web Publishing Service (WWW service) that handles internal administration of the IIS server.

| **Note** |
| --- |
| When messages appear in the event log or Web browser, the symbols that appear below (such as \|, %s, %X, %1, %2, %2!. 20s!, and so on) are replaced with details about the error at the time that the error actually occurs. |

| Event ID | Description |
| --- | --- |
| 1001 | The World Wide Web Publishing Service's Inetinfo.exe process monitor encountered an error and has shut down. This condition prevents the World Wide Web Publishing Service from detecting an unexpected termination of the Inetinfo.exe process and, therefore, may not detect changes to the metabase as expected. It is recommended that you restart the World Wide Web Publishing Service. |
| 1002 | Application pool '%1' is being automatically disabled due to a series of failures in the process (es) serving that application pool. |
| 1003 | Cannot register the URL prefix '%1' for site '%2'. The URL maybe invalid. The site has been deactivated. The data field contains the error number. |
| 1004 | Cannot register the URL prefix '%1' for site '%2'. The site has been deactivated. The data field contains the error number. |
| 1005 | The World Wide Web Publishing Service is exiting due to an error. The data field contains the error number. |
| 1006 | Not Used. |
| 1007 | Cannot register the URL prefix '%1' for site '%2'. The necessary network binding may already be in use. The site has been deactivated. The data field contains the error number. |
| 1008 | Not Used. |
| 1009 | A process serving application pool '%1' terminated unexpectedly. The process id was '%2'. The process exit code was'0x%3'. |
| 1010 | A process serving application pool '%1' failed to respond to aping. The process id was '%2'. |
| 1011 | A process serving application pool '%1' suffered a fatal communication error with the World Wide Web Publishing Service. The process id was '%2'. The data field contains the error number. |
| 1012 | A process serving application pool '%1' exceeded time limits during start up. The process id was '%2'. |
| 1013 | A process serving application pool '%1' exceeded time limits during shut down. The process id was '%2'. |
| 1014 | The World Wide Web Publishing Service encountered an internal error in its process management of worker process '%2' serving application pool '%1'. The data field contains the error number. |

| | |
|---|---|
| 1015 | A process serving application pool '%1' was orphaned but the specified orphan action %2 could not be executed. The data field contains the error number. |
| 1016 | Unable to configure logging for site '%1'. The data field contains the error number. |
| 1017 | The World Wide Web Publishing Service has terminated due to a Worker process failure. |
| 1018 | The World Wide Web Publishing Service was not able to initialize performance counters. The service will run without the performance counters. Please restart the World Wide Web Publishing Service to reinitialize the performance counters. The data field contains the error number. |
| 1019 | The World Wide Web Publishing Service encountered difficulty initializing performance counters. Performance counters will function, however, not all counter data may be accurate. The data field contains the error number. |
| 1020 | The World Wide Web Publishing Service encountered difficulty configuring the HTTP.sys control channel property '%1'. The data field contains the error number. |
| 1021 | The identity of application pool, '%1' is invalid. If it remains invalid when the first request for the application pool is processed, the application pool will be disabled. The data field contains the error number. |
| 1022 | The World Wide Web Publishing Service failed to create a worker process for the application pool '%1'. The data field contains the error number. |
| 1023 | Not Used. |
| 1024 | Virtual site '%1' is configured to truncate its log every '%2'bytes. Since this value must be at least 1048576 bytes (1 megabyte), 1048576 bytes will be used. |
| 1025 | Application pool '%1' exceeded its job limit settings. |
| 1026 | An error occurred securing the handle of application pool '%1'.Therefore, the application pool was not secured as expected. Change the identification information to cause the World Wide Web Publishing Service to secure the handle of the application pool again. The data field contains the error number. |
| 1027 | The World Wide Web Publishing Service's request of HTTP.sys tenable the application pool '%1' failed. The data field contains the error number. |
| 1028 | The World Wide Web Publishing Service failed to add the worker process '%1' to the job object representing application pool '%2'.The data field contains the error number. |
| 1029 | A failure occurred while configuring an application's bindings for site '%1'. The site has been deactivated. The data field contains the error number. |
| 1030 | Inetinfo.exe terminated unexpectedly and the system was not configured to restart IIS Admin. The World Wide Web Publishing Service has shut down. |
| 1031 | Not Used. |
| 1032 | A failure occurred while configuring the logging properties for the site '%1'. Probable cause: A mapped network path is being used for the site's log file directory path, which is not supported by IIS. Please use a UNC path instead. |
| 1033 | A failure occurred while configuring the logging properties for the site '%1'. Probable cause: The site's log file directory contains an invalid machine or share name. |
| 1034 | A failure occurred while configuring the logging properties for the site '%1'. The server does not have access permissions to the site's log file directory. |

| | |
|---|---|
| 1035 | A failure occurred while configuring the logging properties for the site '%1'. The site's log file directory is not fully qualified. Probable cause: It is missing the drive letter or other crucial information. |
| 1036 | A failure occurred while initializing the configuration manager for the World Wide Web Publishing Service. The data field contains the error number. |
| 1037 | A failure occurred while opening the HTTP control channel for the World Wide Web Publishing Service. The data field contains the error number. |
| 1038 | Not Used. |
| 1039 | A process serving application pool '%1' reported a failure. The process id was '%2'. The data field contains the error number. |
| 1040 | Not Used. |
| 1041 | The '%1' '%2' failed range validation for property '%3'. The configured value '%4' is outside of the range '%5' to '%6'. The value will default to '%7'. |
| 1042 | The AppPoolCommand property set on the application pool '%1' is not a valid value. It is '%2'. It must be eitherMD_APPPOOL_COMMAND_START = 1 or MD_APPPOOL_COMMAND_STOP = 2. |
| 1043 | The virtual site '%1' has been invalidated and will be ignored because valid site bindings could not be constructed, or no site bindings exist. |
| 1044 | The virtual site '%1' has been invalidated and will be ignored because the site's AppPoolId is null. |
| 1045 | Not Used. |
| 1046 | The server command property for virtual site '%1' is not a valid value. It is set to '%2'. It must be one of the following:MD_SERVER_COMMAND_START = 1 or MD_SERVER_COMMAND_STOP = 2 orMD_SERVER_COMMAND_PAUSE = 3 or MD_SERVER_COMMAND_CONTINUE = 4. |
| 1047 | The application '%2' belonging to site '%1' has an AppPoolIdset, but the property is empty. Therefore, the application will be ignored. |
| 1048 | The application '%2' belonging to site '%1' has an invalidAppPoolId '%3' set. Therefore, the application will be ignored. |
| 1049 | The application pool '%1' has an IdleTimeout '%2' greater than the PeriodicRestartTime '%3'. The defaults ( IdleTimeout = '%4'& PeriodicRestartTime '%5' ) will be used. |
| 1050 | The application pool id '%1' exceeds length limits. It is '%2'characters and can not be more than '%3' characters. The application pool will therefore be ignored. |
| 1051 | An application pool id was defined as a zero length string. An application pool id must not be a zero length string. The application pool will therefore be ignored. |
| 1052 | "The application '%2' belongs to an invalid site '%1'.Therefore, the application will be ignored." |
| 1053 | The World Wide Web Publishing Service received a change notification, but was unable to process it correctly. The data field contains the error number. |
| 1054 | The World Wide Web Publishing Service failed to run the auto stop action %2 for application pool '%1'. The data field contains the error number. |

| | |
|---|---|
| 1055 | The site '%1' was disabled because the application pool defined for the site '%2' is not a valid application pool. |
| 1056 | The site '%1' was disabled because the application pool '%2'defined for the site's root application is not a valid application pool. |
| 1057 | The identity of application pool '%1' is invalid, so the Worldwide Web Publishing Service can not create a worker process to serve the application pool. Therefore, the application pool has been disabled. |
| 1058 | The World Wide Web Publishing Service encountered a failure while setting the affinity mask of application pool '%1'. Probable cause: The mask does not contain any processors available on this machine. The data field contains the error number. |
| 1059 | A failure was encountered while launching the process serving application pool '%1'. The application pool has been disabled. |
| 1060 | The job object associated with application pool '%1'encountered an error during configuration. CPU monitoring may not function as expected. The data field contains the error number. |
| 1061 | The job object attached to application pool '%1' failed to start its timer. The application pool's CPU usage is being monitored and will eventually reach the limit and report a failure. The data field contains the error number. |
| 1062 | The World Wide Web Publishing Service encountered a failure configuring the logging properties on the HTTP Control Channel.Logging Enabled is '%1'. Log File Directory is '%2'. Log Period is'%3'. Log Truncate Size is '%4'. The data field contains the error number. |
| 1063 | The World Wide Web Publishing Service encountered a failure requesting metabase change notifications. The data field contains the error number. |
| 1064 | The World Wide Web Publishing Service encountered a failure requesting metabase change notifications during recovery fromInetinfo.exe terminating unexpectedly. While the World Wide Web Publishing Service will continue to run, it is highly probable that it is no longer using current metabase data. Please restart the World Wide Web Publishing Service to correct this condition. The data field contains the error number. |
| 1065 | Centralized logging is configured to truncate its log every'%1' bytes. Since this value must be at least 1048576 bytes ( 1megabyte ), 1048576 bytes will be used. |
| 1066 | The World Wide Web Publishing Service encountered an error attempting to configure centralized logging. It is not configured as expected. The data field contains the error number. |
| 1067 | The World Wide Web Publishing Service property '%1' failed range validation. The configured value, '%2' is outside of the range '%3' to '%4'. The default value, '%5', will be used. |
| 1068 | The World Wide Web Publishing Service failed to record the proper state '%2' and win32error '%3' of site '%1'in the metabase. To correct, start/stop the site or restart the World Wide Web Publishing Service. The data field contains the error number. |
| 1069 | The World Wide Web Publishing Service failed to record the proper state '%2' and win32error '%3' of application pool '%1' in the metabase. To correct, start/stop the application pool or restart the World Wide Web Publishing Service. The data field contains the error number. |
| 1070 | The World Wide Web Publishing Service failed to issue recycle requests to all worker processes of application pool '%1'. The data field contains the error number. |
| 1071 | The World Wide Web Publishing Service failed to enable bandwidth throttling on site '%1' because QoS Packet Scheduler is not installed. |

| | |
|---|---|
| 1072 | The World Wide Web Publishing Service failed to enable bandwidth throttling on site '%1'. The data field contains the error number. |
| 1073 | The World Wide Web Publishing Service failed to enable global bandwidth throttling because QoS Packet Scheduler is not installed. |
| 1074 | A worker process with process id of '%1' serving application pool '%2' has requested a recycle because the worker process reached its allowed processing time limit. |
| 1075 | A worker process with process id of '%1' serving application pool '%2' has requested a recycle because it reached its allowed request limit. |
| 1076 | A worker process with process id of '%1' serving application pool '%2' has requested a recycle because it reached its scheduled recycle time. |
| 1077 | A worker process with process id of '%1' serving application pool '%2' has requested a recycle because it reached its virtual memory limit. |
| 1078 | An ISAPI reported an unhealthy condition to its worker process. Therefore, the worker process with process id of '%1' serving application pool '%2' has requested a recycle. |
| 1079 | An administrator has requested a recycle of all worker processes in application pool '%1'. |
| 1080 | The worker processes serving application pool '%1' are being recycled due to 1 or more configuration changes in the application pool properties which necessitate a restart of the processes. |
| 1081 | The worker processes serving application pool '%1' are being recycled due to detected problems with the metabase that may make current cached metadata invalid. |
| 1082 | A worker process with pid '%2' that serves application pool'%1' has been determined to be unhealthy (see previous event log message), but because a debugger is attached to it, the World Wide Web Publishing Service will ignore the error. |
| 1083 | HTTP.sys provided inconsistent site performance counter data tithe World Wide Web Publishing Service. Therefore, the World Wide Web Publishing Service will ignore the data provided. |
| 1084 | The World Wide Web Publishing Service's request of HTTP.sys to disable the application pool '%1' failed. The data field contains the error number. |
| 1085 | The World Wide Web Publishing Service failed to apply a new configuration for application pool '%1'. The data field contains the error number. |
| 1086 | The World Wide Web Publishing Service failed to properly configure the load balancer capabilities on application pool '%1'.The data field contains the error number. |
| 1087 | The World Wide Web Publishing Service failed to properly configure the application pool queue length on app pool '%1'. The data field contains the error number. |
| 1088 | The World Wide Web Publishing Service failed to properly configure the job object for application pool '%1'. The data field contains the error number. |
| 1089 | The World Wide Web Publishing Service failed to issue a demand start to HTTP.sys for application pool '%1'. The data field contains the error number. |

| | |
|---|---|
| 1090 | The World Wide Web Publishing Service failed to update the AutoStart property for application pool '%1'. The data field contains the error number. |
| 1091 | The World Wide Web Publishing Service failed to run the auto stop action for application pool '%1'. The data field contains the error number. |
| 1092 | The World Wide Web Publishing Service failed to issue a request for the worker process '%2' of application pool '%1' to supply its performance counters. The data field contains the error number. |
| 1093 | The World Wide Web Publishing Service failed to overlap recycle for application pool '%1' worker process '%2'. The data field contains the error number. |
| 1094 | The World Wide Web Publishing Service failed to disassociate the application '%1' from the virtual site '%2'. The data field contains the error number. |
| 1095 | The World Wide Web Publishing Service failed to delete the config group for the application '%1' in site '%2'. The data field contains the error number. |
| 1096 | The World Wide Web Publishing Service failed to remove the urls for the virtual site '%1'. The data field contains the error number. |
| 1097 | The World Wide Web Publishing Service failed to set the application pool for the application '%1' in site '%2'. The data field contains the error number. |
| 1098 | The World Wide Web Publishing Service failed to set the max connections for the virtual site '%1'. The data field contains the error number. |
| 1099 | The World Wide Web Publishing Service failed to set the connection timeout for the virtual site '%1'. The data field contains the error number. |
| 1100 | The World Wide Web Publishing Service failed to create the application '%2' in site '%1'. The data field contains the error number. |
| 1101 | The World Wide Web Publishing Service failed to create app pool'%1'. The data field contains the error number. |
| 1102 | The World Wide Web Publishing Service failed to create site'%1'. The data field contains the error number. |
| 1103 | The World Wide Web Publishing Service failed to delete the application '%2' in site '%1'. The data field contains the error number. |
| 1104 | The World Wide Web Publishing Service failed to delete app pool'%1'. The data field contains the error number. |
| 1105 | The World Wide Web Publishing Service failed to delete site'%1'. The data field contains the error number. |
| 1106 | The World Wide Web Publishing Service failed to modify the application '%2' in site '%1'. The data field contains the error number. |
| 1107 | The World Wide Web Publishing Service failed to modify app pool'%1'. The data field contains the error number. |
| 1108 | The World Wide Web Publishing Service failed to modify site'%1'. The data field contains the error number. |
| 1109 | The World Wide Web Publishing Service failed to set the control channel's filter config. The data field contains the error number." |

| | |
|---|---|
| 1110 | The World Wide Web Publishing Service failed to write the AutoStart Property for site '%1'. The data field contains the error number. |
| 1111 | The World Wide Web Publishing Service failed to delete all IIS owned SSL configuration data. The data field contains the error number. |
| 1112 | The World Wide Web Publishing Service failed to query SSL configuration data for site '%1'. The data field contains the error number. |
| 1113 | One of the IP/Port combinations for site '%1' has already be configured to be used by another site. The other site's SSL configuration will be used. |
| 1114 | One of the IP/Port combinations for site '%1' has already been configured to be used by another program. The other program's SSL configuration will be used. |
| 1115 | The World Wide Web Publishing Service failed to set SSL configuration data for site '%1'. The data field contains the error number. |
| 1116 | The World Wide Web Publishing Service failed to delete SSL configuration data for site '%1'. The data field contains the error number. |
| 1117 | A worker process with process id of '%1' serving application pool '%2' has requested a recycle because it reached its private bytes memory limit. |
| 1118 | During recovery from an unexpected termination of theInetinfo.exe process, the World Wide Web Publishing Service failed to identify the appropriate records requiring deletion from its metadata cache. The data field contains the error number. |
| 1119 | During recovery from an unexpected termination of theInetinfo.exe process, the World Wide Web Publishing Service failed to queue action items necessary to handle the recovery. It is recommended that you restart the World Wide Web Publishing Service. The data field contains the error number. |
| 1120 | The World Wide Web Publishing Service failed to obtain cache counters from HTTP.sys. The reported performance counters do not include performance counters from HTTP.sys for this gathering. The data field contains the error number. |
| 1121 | The World Wide Web Publishing Service failed to obtain site performance counters from HTTP.sys. The reported performance counters do not include counters from HTTP.sys for this gathering. The data field contains the error number. |
| 1122 | The World Wide Web Publishing Service failed to publish the performance counters it gathered. The data field contains the error number. |
| 1123 | The World Wide Web Publishing Service failed to cancel the performance counter gathering timer. The data field contains the error number. |
| 1124 | The World Wide Web Publishing Service failed to cancel the performance counter timer. The data field contains the error number. |
| 1125 | The World Wide Web Publishing Service failed to start the performance counter gathering timer. The data field contains the error number. |
| 1126 | The World Wide Web Publishing Service failed to copy a change notification for processing. Therefore, the service may not be in sync with the current data in the metabase. The data field contains the error number. |

| | |
|---|---|
| 1127 | A worker process '%2' serving application pool '%1' is no longer trusted by the World Wide Web Publishing Service, based on ill-formed data the worker process sent to the service. |
| 1128 | The World Wide Web Publishing Service was not able to initialize asp performance counters. The service will therefore run without asp performance counters. Restart the World Wide Web Publishing Service to start asp performance counter gathering. The data field contains the error number. |
| 1129 | Cannot register the URL prefix '%1' for site '%2'. Either the network endpoint for the site's IP address could not be created, or the IP listen list for HTTP.sys did not contain any usable IP addresses. The site has been deactivated. The data field contains the error number. |
| 1130 | Cannot register the URL prefix '%1' for site '%2'. The IP address for the site is not in the HTTP.sys IP listen list. The site has been deactivated. The data field contains the error number. |
| 1131 | Cannot register the URL prefix '%1' for site '%2'. Too many listening ports have been configured in HTTP.sys. The site has been deactivated. The data field contains the error number. |

# IIS 6.0 Event Messages

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) 6.0 generates events in Event Viewer so that you can verify the performance of IIS. Event Viewer tracks the following: error events, warning events, and informational events. The logs in Event Viewer provide an audited record of all services and processes in the Microsoft® Windows® Server 2003 operating system, such as logon and connection times.

**In This Section**

Event Logging Overview

WWW Service Events

WWW Service Worker Process Events

ASP Events and Errors

FTP Service Events

WWW Service Performance Counter Events

**Related Sections**

- For information about viewing and interpreting log files, see Analyzing Log Files in this reference.
- For information about performance monitoring and tuning, see Optimizing IIS 6.0 Performance in this reference.

# Event Logging Overview

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

By default, a computer running the Microsoft® Windows® Server 2003, Standard Edition; Windows® Server 2003, Enterprise Edition; Windows® Server 2003, Web Edition; or Windows® Server 2003, Datacenter Edition operating system records events in three kinds of logs:

- **Application**. Contains events logged by applications or programs.
- **Security**. Records events such as valid and invalid logon attempts, as well as events related to resource use.
- **System**. Contains events logged by Windows system components.

Logging begins automatically when Windows is started. For each type of log, you can filter the events to be viewed, designate the number of entries to view, specify how long to save entries, and specify whether to automatically overwrite existing events when the log becomes full.

You view Windows event logs through Event Viewer. You can open Event Viewer by using any one of the following methods:

- From the **Start** menu, click **Run**. In the **Open** box, type **eventvwr**, and then click **OK**.
- At the command prompt, type **start eventvwr**, and then press ENTER.
- From the **Start** menu, point to **Administrative Tools**, and then click **Event Viewer**.
- From the **Start** menu, point to **Administrative Tools**, and then click **Computer Management**. Event Viewer is listed under the **System Tools** node. The benefit of this display is that IIS Manager is in the same window under the **Services and Applications** node.

For more information about event logs and Event Viewer, see "Event Log Overview" in Help and Support Center for Windows Server 2003.

Operating system events related to IIS are logged in the event logs. These logs are not the same as the IIS logs that record specific information about Web site usage and activity. For more information about IIS logs, see Analyzing Log Files in this reference.

| **Note** |
| --- |
| When messages appear in the event log or Web browser, the symbols that appear in the tables in this chapter (such as \|, %s, %X, %1, %2, %2!.20s!, and so on) are replaced with details about the error at the time that the error actually occurs. |

# WWW Service Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table E.1 lists the events that are generated by the part of the World Wide Web Publishing Service (WWW service) that handles internal administration of the W3SVC. The WWW service events are listed in the System log, with the source name W3SVC. For information about troubleshooting IIS, see Troubleshooting IIS 6.0

**Table E.1 WWW Service Events**

| Event ID | Description |
|---|---|
| 1001 | The WWW services InetInfo process monitor encountered an error and has shut down. This condition prevents the WWW service from detecting an unexpected termination of the InetInfo process and, therefore, it might not detect changes to the metabase as expected. Restart the WWW service. |
| 1002 | Application pool %1 is being automatically disabled because of a series of failures in the processes serving that application pool. |
| 1003 | Cannot register the URL prefix %1 for site %2. The URL might be invalid. The site has been deactivated. The data field contains the error number. |
| 1004 | Cannot register the URL prefix %1 for site %2. The site has been deactivated. The data field contains the error number. |
| 1005 | The WWW service is exiting because an error occurred. The data field contains the error number. |
| 1007 | Cannot register the URL prefix %1 for site %2. The necessary network binding might already be in use. The site has been deactivated. The data field contains the error number. |
| 1009 | A process serving application pool %1 was terminated unexpectedly. The process ID was %2. The process exit code was '0x%3. |
| 1010 | A process serving application pool %1 failed to respond to a ping. The process ID was %2. |
| 1011 | A process serving application pool %1 suffered a fatal communication error with the WWW service. The process ID was %2. The data field contains the error number. |
| 1012 | A process serving application pool %1 exceeded time limits during start up. The process ID was %2. |
| 1013 | A process serving application pool %1 exceeded time limits during shut down. The process ID was %2. |
| 1014 | The WWW service encountered an internal error in its process management of worker process %2 serving application pool %1. The data field contains the error number. |
| 1015 | A process serving application pool %1 was orphaned, but the specified orphan action %2 could not be executed. The data field contains the error number. |
| 1016 | Unable to configure logging for site %1. The data field contains the error number. |
| 1017 | The WWW service has terminated because of a worker process failure. |
| 1018 | The WWW service was not able to initialize performance counters. The service will run without the performance counters. Restart the WWW service to reinitialize the performance counters. The data field contains the error number. |

| | |
|---|---|
| 1019 | The WWW service encountered difficulty initializing performance counters. Performance counters will function; however, some counter data might be inaccurate. The data field contains the error number. |
| 1020 | The WWW service encountered difficulty configuring the HTTP.sys control channel property %1. The data field contains the error number. |
| 1021 | The identity of application pool %1 is invalid. If it remains invalid when the first request for the application pool is processed, the application pool will be disabled. The data field contains the error number. |
| 1022 | The WWW service failed to create a worker process for the application pool %1. The data field contains the error number. |
| 1024 | Virtual site %1 is configured to truncate its log every %2 bytes. Because this value must be at least 1,048,576 bytes (1 megabyte), 1,048,576 bytes will be used. |
| 1025 | Application pool %1 exceeded its job limit settings. |
| 1026 | An error occurred securing the handle of application pool %1'; therefore, the application pool was not secured as expected. Change the identification information to cause the WWW service to secure the handle of the application pool again. The data field contains the error number. |
| 1027 | The WWW service's request of HTTP.sys to enable the application pool %1 failed. The data field contains the error number. |
| 1028 | The WWW service failed to add the worker process %1 to the job object representing application pool %2. The data field contains the error number. |
| 1029 | A failure occurred while configuring an application's bindings for site %1. The site has been deactivated. The data field contains the error number. |
| 1030 | InetInfo was terminated unexpectedly, and the system was not configured to restart the IIS Admin service. The WWW service has shut down. |
| 1032 | A failure occurred while configuring the logging properties for the site %1. Probable cause: A mapped network path is being used for the site's log file directory path, which is not supported by IIS. Use a UNC path instead. |
| 1033 | A failure occurred while configuring the logging properties for the site %1. Probable cause: The site's log file directory contains an invalid machine or share name. |
| 1034 | A failure occurred while configuring the logging properties for the site %1. The server does not have access permissions to the site's log file directory. |
| 1035 | A failure occurred while configuring the logging properties for the site %1. The site's log file directory is not fully qualified. Probable cause: It is missing the drive letter or other crucial information. |
| 1036 | A failure occurred while initializing the configuration manager for the WWW service. The data field contains the error number. |
| 1037 | A failure occurred while opening the HTTP control channel for the WWW service. The data field contains the error number. |
| 1039 | A process serving application pool %1 reported a failure. The process ID was %2. The data field contains the error number. |

| | |
|---|---|
| 1041 | The %1 %2 failed range validation for property %3. The configured value %4 is outside of the range %5 to %6. The value will default to %7. |
| 1042 | The **AppPoolCommand** property set on the application pool %1 is not a valid value. It is %2. It must be either MD_APPPOOL_COMMAND_START = 1 or MD_APPPOOL_COMMAND_STOP = 2. |
| 1043 | The virtual site %1 has been invalidated and will be ignored because valid site bindings could not be constructed, or no site bindings exist. |
| 1044 | The virtual site %1 has been invalidated and will be ignored because the site's **AppPoolID** is null. |
| 1046 | The server command property for virtual site %1 is not a valid value. It is set to %2. It must be one of the following: MD_SERVER_COMMAND_START = 1, MD_SERVER_COMMAND_STOP = 2, MD_SERVER_COMMAND_PAUSE = 3, or MD_SERVER_COMMAND_CONTINUE = 4. |
| 1047 | The application %2 belonging to site %1 has an **AppPoolId** set, but the property is empty; therefore, the application will be ignored. |
| 1048 | The application %2 belonging to site %1 has an invalid **AppPoolId** %3 set; therefore, the application will be ignored. |
| 1049 | The application pool %1 has an **IdleTimeout** %2 greater than the **PeriodicRestartTime** %3. The defaults ( **IdleTimeout** = %4 & PeriodicRestartTime %5 ) will be used. |
| 1050 | The application pool ID %1 exceeds length limits. It is %2 characters and cannot be more than %3 characters. The application pool will, therefore, be ignored. |
| 1051 | An application pool ID was defined as a zero-length string. An application pool ID must not be a zero-length string. The application pool will, therefore, be ignored. |
| 1052 | The application %2 belongs to an invalid site %1'; therefore, the application will be ignored. |
| 1053 | The WWW service received a change notification, but was unable to process it correctly. The data field contains the error number. |
| 1054 | The WWW service failed to run the auto-stop action %2 for application pool %1. The data field contains the error number. |
| 1055 | The site %1 was disabled because the application pool defined for the site %2 is not a valid application pool. |
| 1056 | The site %1 was disabled because the application pool %2 defined for the site's root application is not a valid application pool. |
| 1057 | The identity of application pool %1 is invalid, so the WWW service cannot create a worker process to serve the application pool; therefore, the application pool has been disabled. |
| 1058 | The WWW service encountered a failure while setting the affinity mask of application pool %1. Probable cause: The mask does not contain any processors available on this machine. The data field contains the error number. |
| 1059 | A failure was encountered while launching the process serving application pool %1. The application pool has been disabled. |
| 1060 | The job object associated with application pool %1 encountered an error during configuration. CPU monitoring might not function as expected. The data field contains the error number. |

| | |
|---|---|
| 1061 | The job object attached to application pool %1 failed to start its timer. The application pool's CPU usage is being monitored and will eventually reach the limit and report a failure. The data field contains the error number. |
| 1062 | The WWW service encountered a failure configuring the logging properties on the HTTP Control Channel. **Logging Enabled** is %1. **Log File Directory** is %2. **Log Period** is %3. **Log Truncate Size** is %4. The data field contains the error number. |
| 1063 | The WWW service encountered a failure requesting metabase change notifications. The data field contains the error number. |
| 1064 | During recovery from an unexpected termination of the InetInfo process, the WWW service encountered a failure requesting metabase change notifications. Although the WWW service will continue to run, it is highly probable that it is no longer using current metabase data. Restart the WWW service to correct this condition. The data field contains the error number. |
| 1065 | Centralized logging is configured to truncate its log every %1 bytes. Because this value must be at least 1,048,576 bytes (1 megabyte), 1,048,576 bytes will be used. |
| 1066 | The WWW service encountered an error attempting to configure centralized logging. It is not configured as expected. The data field contains the error number. |
| 1067 | The WWW service property %1 failed range validation. The configured value %2 is outside of the range %3 to %4. The default value, %5, will be used. |
| 1068 | The WWW service failed to record the proper state %2 and **Win32Error** %3 of site %1 in the metabase. To correct, start/stop the site, or restart the WWW service. The data field contains the error number. |
| 1069 | The WWW service failed to record the proper state %2 and **Win32Error** %3 of application pool %1 in the metabase. To correct, start/stop the application pool, or restart the WWW service. The data field contains the error number. |
| 1070 | The WWW service failed to issue recycle requests to all worker processes of application pool %1. The data field contains the error number. |
| 1071 | The WWW service failed to enable bandwidth throttling on site %1 because QoS Packet Scheduler is not installed. |
| 1072 | The WWW service failed to enable bandwidth throttling on site %1. The data field contains the error number. |
| 1073 | The WWW service failed to enable global bandwidth throttling because QoS Packet Scheduler is not installed. |
| 1074 | A worker process with process ID of %1 serving application pool %2 has requested a recycle because the worker process reached its allowed processing time limit. |
| 1075 | A worker process with process ID of %1 serving application pool %2 has requested a recycle because it reached its allowed request limit. |
| 1076 | A worker process with process ID of %1 serving application pool %2 has requested a recycle because it reached its scheduled recycle time. |
| 1077 | A worker process with process ID of %1 serving application pool %2 has requested a recycle because it reached its virtual memory limit. |
| 1078 | An ISAPI reported an unhealthy condition to its worker process; therefore, the worker process with process ID of %1 serving application pool %2 has requested a recycle. |
| 1079 | An administrator has requested a recycle of all worker processes in application pool %1. |

| | |
|---|---|
| 1080 | The worker processes serving application pool %1 are being recycled, because one or more configuration changes in the properties of the application pool necessitated a restart of the processes. |
| 1081 | The worker processes serving application pool %1 are being recycled because of detected problems with the metabase that might make current cached metadata invalid. |
| 1082 | A worker process with process ID %2 that serves application pool %1 has been determined to be unhealthy (see previous event log message), but — because a debugger is attached to it — the WWW service will ignore the error. |
| 1083 | HTTP.sys provided inconsistent site performance counter data to the WWW service. Therefore, the WWW service will ignore the data provided. |
| 1084 | The WWW service's request of HTTP.sys to disable the application pool %1 failed. The data field contains the error number. |
| 1085 | The WWW service failed to apply a new configuration for application pool %1. The data field contains the error number. |
| 1086 | The WWW service failed to properly configure the load balancer capabilities on application pool %1. The data field contains the error number. |
| 1087 | The WWW service failed to properly configure the application pool queue length on application pool %1. The data field contains the error number. |
| 1088 | The WWW service failed to properly configure the job object for application pool %1. The data field contains the error number. |
| 1089 | The WWW service failed to issue a demand start to HTTP.sys for application pool %1. The data field contains the error number. |
| 1090 | The WWW service failed to update the **AutoStart** property for application pool %1. The data field contains the error number. |
| 1091 | The WWW service failed to run the auto-stop action for application pool %1. The data field contains the error number. |
| 1092 | The WWW service failed to issue a request for the worker process %2 of application pool %1 to supply its performance counters. The data field contains the error number. |
| 1093 | The WWW service failed to overlap recycle for application pool %1 worker process %2. The data field contains the error number. |
| 1094 | The WWW service failed to disassociate the application %1 from the virtual site %2. The data field contains the error number. |
| 1095 | The WWW service failed to delete the configuration group for the application %1 in site %2. The data field contains the error number. |
| 1096 | The WWW service failed to remove the URLs for the virtual site %1. The data field contains the error number. |
| 1097 | The WWW service failed to set the application pool for the application %1 in site %2. The data field contains the error number. |
| 1098 | The WWW service failed to set the maximum connections for the virtual site %1. The data field contains the error number. |
| 1099 | The WWW service failed to set the connection time-out for the virtual site %1. The data field contains the error number. |

| | |
|---|---|
| 1100 | The WWW service failed to create the application %2 in site %1. The data field contains the error number. |
| 1101 | The WWW service failed to create application pool %1. The data field contains the error number. |
| 1102 | The WWW service failed to create site %1. The data field contains the error number. |
| 1103 | The WWW service failed to delete the application %2 in site %1. The data field contains the error number. |
| 1104 | The WWW service failed to delete application pool %1. The data field contains the error number. |
| 1105 | The WWW service failed to delete site %1. The data field contains the error number. |
| 1106 | The WWW service failed to modify the application %2 in site %1. The data field contains the error number. |
| 1107 | The WWW service failed to modify application pool %1. The data field contains the error number. |
| 1108 | The WWW service failed to modify site %1. The data field contains the error number. |
| 1109 | The WWW service failed to set the control channel's filter configuration. The data field contains the error number. |
| 1110 | The WWW service failed to write the **AutoStart** property for site %1. The data field contains the error number. |
| 1111 | The WWW service failed to delete all IIS-owned Secure Sockets Layer (SSL) configuration data. The data field contains the error number. |
| 1112 | The WWW service failed to query SSL configuration data for site %1. The data field contains the error number. |
| 1113 | One of the IP/Port combinations for site %1 has already been configured to be used by another site. The other site's SSL configuration will be used. |
| 1114 | One of the IP/Port combinations for site %1 has already been configured to be used by another program. The other program's SSL configuration will be used. |
| 1115 | The WWW service failed to set SSL configuration data for site %1. The data field contains the error number. |
| 1116 | The WWW service failed to delete SSL configuration data for site %1. The data field contains the error number. |
| 1117 | A worker process with process ID of %1 serving application pool %2 has requested a recycle because it reached its private-bytes-memory limit. |
| 1118 | During recovery from an unexpected termination of the InetInfo process, the WWW service failed to identify the appropriate records requiring deletion from its metadata cache. The data field contains the error number. |
| 1119 | During recovery from an unexpected termination of the InetInfo process, the WWW service failed to queue action items necessary to handle the recovery. Restart the WWW service. The data field contains the error number. |
| 1120 | The WWW service failed to obtain cache counters from HTTP.sys. The reported performance counters do not include performance counters from HTTP.sys for this gathering. The data field contains the error number. |
| 1121 | The WWW service failed to obtain site performance counters from HTTP.sys. The reported performance counters do not include counters from HTTP.sys for this gathering. The data field contains the error number. |
| 1122 | The WWW service failed to publish the performance counters it gathered. The data field contains the error number. |

| | |
|---|---|
| 1123 | The WWW service failed to cancel the performance counter gathering timer. The data field contains the error number. |
| 1124 | The WWW service failed to cancel the performance counter timer. The data field contains the error number. |
| 1125 | The WWW service failed to start the performance counter gathering timer. The data field contains the error number. |
| 1126 | The WWW service failed to copy a change notification for processing. Therefore, the service might not be in sync with the current data in the metabase. The data field contains the error number. |
| 1127 | A worker process %2 serving application pool %1 is no longer trusted by the WWW service, based on ill-formed data the worker process sent to the service. |
| 1128 | The WWW service was not able to initialize Active Server Pages (ASP) performance counters. The service will, therefore, run without ASP performance counters. Restart the WWW service to start ASP performance-counter gathering. The data field contains the error number. |
| 1129 | Cannot register the URL prefix %1 for site %2. Either the network endpoint for the site's IP address could not be created or the IP listen list for HTTP.sys did not contain any usable IP addresses. The site has been deactivated. The data field contains the error number. |
| 1130 | Cannot register the URL prefix %1 for site %2. The IP address for the site is not in the HTTP.sys IP listen list. The site has been deactivated. The data field contains the error number. |
| 1131 | Cannot register the URL prefix %1 for site %2. Too many listening ports have been configured in HTTP.sys. The site has been deactivated. The data field contains the error number. |

# WWW Service Worker Process Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table E.2 lists the events that the WWW service worker process generates, such as authentication and authorization, application problems, memory monitoring, and so on. These events are listed in the Application log with the source name W3SVC-WP.

**Table E.2 Worker Process Events for the WWW Service**

| Event ID | Description |
|---|---|
| 2201 | The HTTP server could not initialize its security. The data is the error. |
| 2203 | The HTTP server could not initialize the socket library. The data is the error. |
| 2204 | The HTTP server was unable to initialize because of a shortage of available memory. The data is the error. |
| 2206 | The HTTP server could not create the main connection socket. The data is the error. |
| 2208 | The HTTP server could not create a client connection object for the user at host %1. The connection to this user is terminated. The data is the error. |
| 2214 | The HTTP filter DLL %1 failed to load. The data is the error. |
| 2216 | The script that was started from the URL %1 with parameters %2 has not responded within the configured time-out period. The HTTP server is terminating the script. |
| 2218 | The HTTP server encountered an error processing the server-side include file %1. The error was %2. |
| 2219 | The HTTP server encountered an unhandled exception while processing the ISAPI application %1. |
| 2220 | The HTTP server was unable to load the ISAPI application %1. The data is the error. |
| 2221 | A server-side include file has included itself, or the maximum depth of server-side includes has been exceeded. |
| 2222 | An attempt was made to load filter %1 on a server instance, but it requires the SF_NOTIFY_READ_RAW_DATA filter notification so it must be loaded as a global filter. |
| 2223 | For compatibility with previous versions of IIS, the filter %1 was loaded as a global filter from the registry. To control the filter with IIS Manager, remove the filter from the registry and add it as a global filter with IIS Manager. Filters in the registry are stored at HKLM\System\CurrentControlSet\Services\W3Svc\Parameters\Filter DLLs. |
| 2226 | The server was unable to read the file %1 because of a lack of access permissions. |
| 2227 | The server was unable to acquire a license for an SSL connection. |
| 2228 | The server stopped serving requests for application %1 because the number of crashes of out-of-process components exceeded a limit. |
| 2229 | The server failed to shut down application %1. The error was %2. |
| 2230 | The server was unable to read the file %1. The file does not exist. |

| | |
|---|---|
| 2231 | The server was unable to read the file %1. The Windows 32 error returned from the attempt is %2. |
| 2232 | The server was unable to read the file %1. The file exceeds the maximum allowable size of %2. |
| 2233 | The server was unable to allocate a buffer to read the file %1. |
| 2236 | The server failed to load application %1. The error was %2. |
| 2237 | Out-of-process application %1 terminated unexpectedly. |
| 2238 | The job object query failed. The data is the error. |
| 2239 | The job object set limit failed. The data is the error. |
| 2240 | Schedule a work item for CPU Accounting or Limits failed. The data is the error. |
| 2241 | Site %1 hit its CPU Limit. No action was taken. |
| 2242 | Site %1 hit its CPU Limit. The priority of all process on that site has been lowered to idle class. |
| 2243 | Site %1 hit its CPU Limit. All processes on that site have been terminated. |
| 2244 | Site %1 hit its CPU Limit. The site has been paused. |
| 2245 | The server certificate for instance %1 has expired or is not yet valid. |
| 2246 | The server certificate for instance %1 has been revoked. |
| 2247 | The server certificate for instance %1 does not chain up to a trusted root certificate. |
| 2248 | One of the certificates in the certificate chain of the server certificate for instance %1 has an invalid signature. |
| 2249 | CPU Limits/Accounting failed to queue work item. The data is the error. |
| 2250 | The server certificate for instance %1 had invalid metabase data associated with it and could not be retrieved. The error encountered was %2. |
| 2251 | A CryptoAPI error was encountered trying to retrieve the server certificate for instance %1. The error encountered was %2. |
| 2252 | The server certificate for instance %1 could not be retrieved because it could not be found in a certificate store. The error encountered was %2'. |
| 2253 | The server certificate for instance %1 could not be retrieved because of an internal error. The error encountered was %2. |
| 2254 | The server Certificate Trust List for instance %1 had invalid metabase data associated with it and could not be retrieved. The error encountered was %2. |
| 2255 | A CryptoAPI error was encountered trying to retrieve the server Certificate Trust List for instance %1. The error encountered was %2. |
| 2256 | The server Certificate Trust List for instance %1 could not be retrieved because it could not be found in a certificate store. The error encountered was %2. |

| | |
|---|---|
| 2257 | The server certificate for instance %1 could not be retrieved because of an internal error. The error number is %2. |
| 2258 | The server failed to close client connections to the following URLs during shutdown: %1. |
| 2259 | The COM application %1 at %2 failed to activate out of process. |
| 2261 | An attempt was made to load filter %1, but it requires the SF_NOTIFY_READ_RAW_DATA filter notification, which is not supported in worker process isolation mode. |
| 2262 | ISAPI %1 reported itself as unhealthy for the following reason: %2. |
| 2263 | ISAPI %1 reported itself as unhealthy. No reason was given by the ISAPI. |
| 2264 | The directory specified for caching compressed content %1 is invalid. Static compression is being disabled. |
| 2265 | The registry key for IIS subauthenticator is not configured correctly on the local machine; the anonymous password sync feature is disabled. |
| 2266 | The account that the current worker process is running under does not have SeTcbPrivilege privilege; the anonymous password sync feature and the Digest authentication feature are disabled. |
| 2267 | A raw ISAPI filter could not be loaded because of a configuration problem. HTTPFilter Service has to be hosted in InetInfo.exe for raw ISAPI filters to work correctly. The most likely reason is that the **IIS5IsolationModeEnabled** setting was changed and HTTPFilter service was not restarted. Restart HTTPFilter for configuration changes to be taken into effect. |
| 2268 | Could not load all ISAPI filters for site/service. Therefore, startup was aborted. |
| 2269 | The worker process failed to initialize the HTTP.sys communication or the W3svc.exe communication layer and, therefore, could not be started. The data field contains the error number. |
| 2270 | IIS client certificate mapping configuration for site %1 failed to be loaded. The data field contains the error number. |
| 2271 | Could not initialize the logging module for site %1. The site will, therefore, be nonfunctional. |
| 2272 | The registry key for IIS subauthenticator is not configured correctly on the domain controller; the Digest authentication feature is disabled. |
| 2273 | The worker process could not access the metabase because a disconnection error occurred. Marking process as unhealthy. |

# ASP Events and Errors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

In Event Viewer, you can open the property pages for each type of log and configure the filter options that generally apply to all events.

IIS increases the degree of control by providing two metabase keys that enable you to specify what is sent to the Windows Events log when ASP events and errors occur: **AspLogErrorRequests** and **AspErrorsToNTLog**.

**AspLogErrorRequests**
The **AspLogErrorRequests** property controls whether the Web server writes unsuccessful client requests to the Windows event log file. If **AspLogErrorRequests** is set to TRUE, a standard set of ASP error requests are logged.

If **AspLogErrorRequests** is set to TRUE and **AspErrorsToNTLog** is set to FALSE, then the ASP errors are sent to the IIS log. If these errors are serious or logging to the IIS log previously failed, then each error is also sent to the Windows event log.

If **AspLogErrorRequests** is set to FALSE, these errors are sent only to the IIS log and not to the Windows Event Log. If **AspLogErrorRequests** is set to TRUE and **AspErrorsToNTLog** is set to TRUE, then all ASP-related errors go to the Windows event log and the IIS log. For more information about setting the **AspLogErrorRequests** property, see About Events.

To completely prevent ASP errors from being logged in the IIS log file, you must disable logging altogether. For more information about disabling logging, see Analyzing Log Files.

**AspErrorsToNTLog**
If the **AspErrorsToNTLog** key is set to TRUE, all possible ASP errors are sent to the application event log. If this key is set to FALSE, only a subset of all ASP errors is sent. The subset includes the error numbers 100, 101, 102, 103, 104, 105, 106, 107, 115, 190, 191, 192, 193, 194, 240, 241, and 242. This key cannot be set in IIS Manager. It must be set programmatically, using scripts, for example.

# ASP Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table E.3 lists the events that the WWW service generates when processing ASP applications that are configured incorrectly or ASP pages that contain errors. These events are listed in the Application log regardless of the value of the **ASPLogErrorRequests** metabase property.

**Table E.3 ASP Events**

| Event ID | Description |
|---|---|
| 0x5 | Error: %1. |
| 0x6 | Error: %1, %2. |
| 0x7 | Error: %1, %2, %3. |
| 0x8 | Error: %1, %2, %3, %4. |
| 0x9 | Warning: %1. |
| 0x10 | Warning: %1, %2. |
| 0x11 | Warning: %1, %2, %3. |
| 0x12 | Warning: %1, %2, %3, %4. |
| 0x13 | Failed to create an ASP session because access was denied when activating COM+ partition %1. |
| 0x14 | Failed to create ASP application %1 because of invalid or missing COM+ partition ID. |
| 0x15 | Failed to retrieve the anonymous user token for ASP application %1. Global.asa OnEnd routines will not be executed. |
| 0x16 | Failed to impersonate the anonymous user for ASP application %1. Global.asa OnEnd routines will not be executed. |
| 0x17 | No default script language was specified for application %s. Using default: %s. |
| 0x1f3 | Invalid progID. |

# ASP Errors

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table E.4 lists the errors that the WWW service generates when processing ASP applications that are configured incorrectly or ASP pages that contain errors. These errors are listed in the Application log only if the **ASPLogErrorRequests** metabase property is set to TRUE. If the **ASPErrorsToNTLog** metabase property is also set to TRUE, all the errors get logged to the Windows event application log; otherwise, only a subset is logged. The subset includes the error numbers 100, 101, 102, 103, 104, 105, 106, 107, 115, 190, 191, 192, 193, 194, 240, 241, and 242. These error messages are listed in the Application log with the source name Active Server Pages.

**Table E.4 ASP Errors**

| Error ID | Description |
|---|---|
| 0100 | Out of memory. Unable to allocate required memory. |
| 0101 | Unexpected error. The function returned \|. |
| 0102 | Expecting string input. The function expects a string as input. |
| 0103 | Expecting numeric input. The function expects a number as input. |
| 0104 | Operation not allowed. |
| 0105 | Index out of range. An array index is out of range. |
| 0106 | Type mismatch. An unhandled data type was encountered. |
| 0107 | Data size too large. The size of the data being sent in the request is over the allowed limit. |
| 0108 | Create object failed. An error occurred while creating object %s. |
| 0109 | Member not found. |
| 0110 | Unknown name. |
| 0111 | Unknown interface. |
| 0112 | Missing parameter. |
| 0113 | Script timed out. The maximum amount of time for a script to be executed was exceeded. You can change this limit by specifying a new value for the property **Server.ScriptTimeout** or by changing the value in the IIS administration tools. |
| 0114 | Object not free-threaded. The application object accepts only free-threaded objects; object %s is not free-threaded. |
| 0115 | Unexpected error. A trappable error (%X) occurred in an external object. The script cannot continue running. |
| 0116 | Missing close of script delimiter. The script block lacks the close of the <script> tag (%>). |
| 0117 | Missing close of <script> tag. The script block lacks the close of the <script> tag (</SCRIPT>) or the close-of-tag symbol (>). |

| | |
|---|---|
| 0118 | Missing close of <object> tag. The object block lacks the close of the <object> tag (</OBJECT>) or the close-of-tag symbol (>). |
| 0119 | Missing **Classid** or **Progid** attribute. The object instance '\|' requires a valid **Classid** or **Progid** in the <object> tag. |
| 0120 | Invalid **Runat** attribute. The **Runat** attribute of the <script> tag or <object> tag can have only the value 'Server. |
| 0121 | Invalid scope in <object> tag. The object instance '\|' cannot have Application or Session scope. To create the object instance with Session or Application scope, place the <object> tag in the Global.asa file. |
| 0122 | Invalid scope in <object> tag. The object instance '\|' must have Application or Session scope. This applies to all objects created in a Global.asa file. |
| 0123 | Missing **ID** attribute. The required **ID** attribute of the <object> tag is missing. |
| 0124 | Missing **Language** attribute. The required **Language** attribute of the <script> tag is missing. |
| 0125 | Missing close of attribute. The value of the '\|' attribute has no closing delimiter. |
| 0126 | Include file not found. The include file '\|' was not found. |
| 0127 | Missing close of HTML comment. The HTML comment or server-side include lacks the close tag (-->). |
| 0128 | Missing **File** or **Virtual** attribute. The include file name must be specified by using either the **File** or **Virtual** attribute. |
| 0129 | Unknown scripting language. The scripting language '\|' is not found on the server. |
| 0130 | Invalid **File** attribute. **File** attribute '\|' cannot start with a forward slash or a backward slash. |
| 0131 | Disallowed parent path. The include file '\|' cannot contain '..' to indicate the parent directory. |
| 0132 | Compilation error. The ASP page '\|' could not be processed. |
| 0133 | Invalid **ClassID** attribute. The <object> tag has an invalid **ClassID** of '\|. |
| 0134 | Invalid **ProgID** attribute. The object has an invalid **ProgID** of '\|. |
| 0135 | Cyclic include. The file '\|' is included by itself (perhaps indirectly). Select the **Include files for other include statements** box. |
| 0136 | Invalid object instance name. The object instance '\|' is attempting to use a reserved name. This name is used by ASP pages intrinsic objects. |
| 0137 | Invalid global script. Script blocks must be one of the allowed Global.asa procedures. Script directives within <% ... %> are not allowed within the Global.asa file. The allowed procedure names are **Application_OnStart**, **Application_OnEnd**, **Session_OnStart**, and **Session_OnEnd**. |
| 0138 | Nested script block. A script block cannot be placed inside another script block. |
| 0139 | Nested object. An <object> tag cannot be placed inside another <object> tag. |
| 0140 | Page command out of order. The @ command must be the first command within the ASP page. |
| 0141 | Page command repeated. The @ command can be used only once within the ASP page. |

| | |
|---|---|
| 0142 | Thread token error. A thread token failed to open. |
| 0143 | Invalid application name. A valid application name was not found. |
| 0144 | Initialization error. The page-level objects list failed during initialization. |
| 0145 | New application failed. The new application could not be added. |
| 0146 | New session failed. The new session could not be added. |
| 0147 | 500 server error. |
| 0148 | Server too busy. |
| 0149 | Application restarting. The request cannot be processed while the application is being restarted. |
| 0150 | Application directory error. The application directory could not be opened. |
| 0151 | Change notification error. The change notification event could not be created. |
| 0152 | Security error. An error occurred while processing a user's security credentials. |
| 0153 | Thread error. A new thread request failed. |
| 0154 | Write HTTP header error. The HTTP headers could not be written to the client browser. |
| 0155 | Write page content error. The page content could not be written to the client browser. |
| 0156 | Header error. The HTTP headers are already written to the client browser. Any HTTP header modifications must be made before writing page content. |
| 0157 | Buffering on. Buffering cannot be turned off once it is turned on. |
| 0158 | Missing URL. A URL is required. |
| 0159 | Buffering off. Buffering must be on. |
| 0160 | Logging failure. Failure to write entry to the log. |
| 0161 | Data type error. The conversion of a **Variant** to a **String** variable failed. |
| 0162 | Cannot modify cookie. The cookie 'ASPSessionID' cannot be modified. It is a reserved cookie name. |
| 0163 | Invalid comma use. Commas cannot be used within a log entry. Select another delimiter. |
| 0164 | Invalid **TimeOut** value. An invalid **TimeOut** value was specified. |
| 0165 | **SessionID** error. A **SessionID** string cannot be created. |
| 0166 | Uninitialized object. An attempt was made to access an uninitialized object. |
| 0167 | Session initialization error. An error occurred while initializing the **Session** object. |

| | |
|---|---|
| 0168 | Disallowed object use. An intrinsic object cannot be stored within the **Session** object. |
| 0169 | Missing object information. An object with missing information cannot be stored in the **Session** object. The threading model information for an object is required. |
| 0170 | Delete session error. The session was not deleted properly. |
| 0171 | Missing path. The *Path* parameter must be specified for the **MapPath** method. |
| 0172 | Invalid path. The *Path* parameter for the **MapPath** method must be a virtual path. A physical path was used. |
| 0173 | Invalid path character. An invalid character was specified in the *Path* parameter for the **MapPath** method. |
| 0174 | Invalid path characters. An invalid '/' or '\\' was found in the *Path* parameter for the **MapPath** method. |
| 0175 | Disallowed path characters. The '..' characters are not allowed in the *Path* parameter for the **MapPath** method. |
| 0176 | Path not found. The *Path* parameter for the **MapPath** method did not correspond to a known path. |
| 0177 | **Server.CreateObject** failed. %s' |
| 0178 | **Server.CreateObject** access error. The call to **Server.CreateObject** failed while checking permissions. Access is denied to this object. |
| 0179 | Application initialization error. An error occurred while initializing the **Application** object. |
| 0180 | Disallowed object use. An intrinsic object cannot be stored within the **Application** object. |
| 0181 | Invalid threading model. An object using the apartment threading model cannot be stored within the **Application** object. |
| 0182 | Missing object information. An object with missing information cannot be stored in the **Application** object. The threading model information for the object is required. |
| 0183 | Empty cookie key. A cookie with an empty key cannot be stored. |
| 0184 | Missing cookie name. A name must be specified for a cookie. |
| 0185 | Missing default property. A default property was not found for the object. |
| 0186 | Error parsing certificate. |
| 0187 | Object addition conflict. Could not add object to application. Application was locked down by another request for adding an object. |
| 0188 | Disallowed object use. Cannot add objects created by using <object> tags to the session intrinsic. |
| 0189 | Disallowed object use. Cannot add objects created by using <object> tags to the application intrinsic. |
| 0190 | Unexpected error. A trappable error occurred while releasing an external object. |
| 0191 | Unexpected error. A trappable error occurred in the **OnStartPage** method of an external object. |
| 0192 | Unexpected error. A trappable error occurred in the **OnEndPage** method of an external object. |

| | |
|---|---|
| 0193 | **OnStartPage** failed. An error occurred in the **OnStartPage** method of an external object. |
| 0194 | **OnEndPage** failed. An error occurred in the **OnEndPage** method of an external object. |
| 0195 | Invalid server method call. This method of the **Server** object cannot be called during **Session_OnEnd** and **Application_OnEnd**. |
| 0196 | Cannot launch out-of-process component. Only InProc server components should be used. If you want to use LocalServer components, you must set the **AspAllowOutOfProcComponents** metabase setting. |
| 0197 | Disallowed object use. Cannot add object with apartment model behavior to the application intrinsic object. |
| 0198 | Server shutting down. Cannot process request. |
| 0199 | Disallowed object use. Cannot add **JScript** objects to the session. |
| 0200 | Out-of-range **Expires** attribute. The date and time given for the **Expires** attribute is earlier than Jan 1, 1980 or later than Jan 19, 2038, 3:14:07 GMT. |
| 0201 | Invalid default script language. The default script language specified for this application is invalid. |
| 0202 | Missing code page. The **CodePage** attribute is missing. |
| 0203 | Invalid code page. The specified **CodePage** attribute is invalid. |
| 0204 | Invalid **CodePage** value. An invalid **CodePage** value was specified. |
| 0205 | Change notification. Failed to create event for change notification. |
| 0206 | Cannot call BinaryRead. Cannot call BinaryRead after using Request.Form collection. |
| 0207 | Cannot use Request.Form. Cannot use Request.Form collection after calling BinaryRead. |
| 0208 | Cannot use generic Request collection. Cannot use the generic Request collection after calling BinaryRead. |
| 0209 | Illegal value for **Transaction** property. The **Transaction** property can only be REQUIRED, REQUIRES_NEW, SUPPORTED, or NOT_SUPPORTED. |
| 0210 | Method not implemented. This method has not yet been implemented. |
| 0211 | Object out of scope. A built-in ASP object, which is no longer valid, has been referenced. |
| 0212 | Cannot clear buffer. Response.Clear is not allowed after a Response.Flush while client debugging is enabled. |
| 0214 | Invalid *Path* parameter. The *Path* parameter exceeds the maximum length allowed. |
| 0215 | Illegal value for **EnableSessionState** property. The **EnableSessionState** property can be only TRUE or FALSE. |
| 0216 | MSDTC Service not running. Transactional Web pages cannot be run if the MSDTC service is not running. |
| 0217 | Invalid scope in <object> tag. Object scope must be Page, Session, or Application. |
| 0218 | Missing LCID. The LCID attribute is missing. |

| | |
|---|---|
| 0219 | Invalid LCID. The specified LCID is not available. |
| 0220 | Requests for Global.asa not allowed. Requests with the URL pointing to Global.asa are not allowed. |
| 0221 | Invalid @ Command directive. The specified '\|' option is unknown or invalid. |
| 0222 | Invalid TypeLib Specification. The <metadata> tag contains an invalid type library specification. |
| 0223 | TypeLib not found. The <metadata> tag contains a type library specification that does not match any registry entry. |
| 0224 | Cannot load TypeLib. Cannot load the type library specified in the <metadata> tag. |
| 0225 | Cannot wrap TypeLibs. Cannot create a type library wrapper object from the type libraries specified in <metadata> tags. |
| 0226 | Cannot modify StaticObjects. StaticObjects collection cannot be modified at run time. |
| 0227 | Server.Execute failed. The call to Server.Execute failed. |
| 0228 | Server.Execute error. The call to Server.Execute failed while loading the page. |
| 0229 | Server.Transfer failed. The call to Server.Transfer failed |
| 0230 | Server.Transfer error. The call to Server.Transfer failed while loading the page. |
| 0231 | Server.Execute error. Invalid URL form or fully qualified absolute URL was used. Use relative URLs. |
| 0232 | Invalid cookie specification. The <metadata> tag contains an invalid cookie specification. |
| 0233 | Cannot load cookie script source. Cannot load cookie script source file specified in the <metadata> tag. |
| 0234 | Invalid include directive. Server-side include directives cannot be present in script blocks. Use the **SRC=** attribute of the <script> tag. |
| 0235 | Server.Transfer error. Invalid URL form or fully qualified absolute URL was used. Use relative URLs. |
| 0236 | Invalid cookie specification. The <metadata> tag contains an invalid or missing *SRC* parameter. |
| 0237 | Invalid cookie specification. The <metadata> tag contains an invalid or missing *Name* parameter. |
| 0238 | Missing attribute value. No value was specified for the '\|' attribute. |
| 0239 | Cannot process file. Unicode ASP files are not supported. |
| 0240 | Script engine exception. A script engine threw exception %X in %s from %s. |
| 0241 | **CreateObject** exception. The **CreateObject** of %s caused exception %X. |
| 0242 | Query **OnStartPage** interface exception. Querying the **OnStartPage** or **OnEndPage** methods of object %s caused exception %X. |
| 0243 | Invalid <metadata> tag in Global.asa. Only <metadata> TYPE="TypeLib" can be used in Global.asa. |
| 0244 | Cannot enable session state. Session state cannot be enabled when it has been disabled in the application. |

| | |
|---|---|
| 0245 | Mixed usage of code page values. The **@CODEPAGE** value specified differs from that of the including file's code page or the file's saved format. |
| 0246 | Too many concurrent users. Try again later. |
| 0247 | Bad argument to **BinaryRead**. The argument to **BinaryRead** must be nonnegative. |
| 0248 | Script is not transacted. This ASP file must be transacted to use the **ObjectContext** object. |
| 0249 | Cannot use IStream on **Request**. Cannot use IStream on **Request** object after using Request.Form collection or Request.BinaryRead. |
| 0250 | Invalid default code page. The default code page specified for this application is invalid. |
| 0251 | Response buffer limit exceeded. Execution of the ASP page caused the response buffer to exceed its configured limit. |

# FTP Service Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table E.5 lists the events that are generated by the FTP service. The FTP service events are listed in the System log with the source name MSFTPSVC.

**Table E.5 FTP Service Events**

| Event ID | Description |
|---|---|
| 1 | FTP server could not initialize its security. The data is the error. |
| 3 | FTP server could not initialize the socket library. The data is the error. |
| 4 | FTP server was unable to be initialized because of a shortage of available memory. The data is the error. |
| 5 | FTP server could not locate the FTP/TCP service. The data is the error. |
| 6 | FTP server could not create the main connection socket. The data is the error. |
| 7 | FTP server could not create the main connection thread. The data is the error. |
| 8 | FTP server could not create a client worker thread for the user at host %1. The connection to this user is terminated. The data is the error. |
| 9 | A call to a system service failed unexpectedly. The data is the error. |
| 10 | User %1 at host %2 has timed out after %3 seconds of inactivity. |
| 11 | An anonymous logon request was received from %1 at host %2. |
| 12 | A user logon request was received from %1 at host %2. |
| 13 | User %1 failed to log on, could not access the home directory %2. |
| 14 | User %1 denied access to the current directory %2 because of a security change. |
| 15 | Failed to read metabase property **PassivePortRange** under the FTP service node %1. |
| 16 | The FTP service configuration property **PassivePortRange** has the invalid value: %1. |

# WWW Service Performance Counter Events

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Table E.6 lists the events that are generated by the part of the WWW service that monitors and records performance counter data. These events are listed in the Application log with the source name W3Ctrs.

**Table E.6 WWW Service Performance Counter Events**

| Event ID | Description |
|---|---|
| 2000 | Unable to query the WWW service performance data. The error code returned by the service is data DWORD 0. |
| 2001 | It has taken too long to refresh the WWW service counters; the stale counters are being used instead. |
| 2002 | Setting up Web service counters failed; make sure your Web service counters are registered correctly. |
| 2003 | It has taken too long to refresh the WWW service counters; the stale counters are being used instead. This is the second message within the past %1 (hours, minutes, seconds). No more stale counter messages will be logged for this client session until the time limit expires. |

# Metabase Auditing Event Messages

Applies To: Windows Server 2003 with SP1

When metabase auditing is enabled, IIS tracks and logs each change to the metabase keys or properties using event messages in the Security log. Each change type is associated with an event message ID in the Security log. Metabase auditing records similar information for each event message.

For information about the common output produced for metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

The Metabase Auditing Event Messages table shows the event IDs, names, and descriptions for IIS 6.0 metabase auditing event messages.

**Metabase Auditing Event Messages**

| Event ID | Event Name | Description | Event Message Output |
|---|---|---|---|
| 4500 | IISADMIN_AUDIT_ADD_KEY | Shows information about a change to add a metabase key | See Metabase Audit Add Key |
| 4501 | IISADMIN_AUDIT_DELETE_KEY | Shows information about a change to delete a metabase key | See Metabase Audit Delete Key |
| 4502 | IISADMIN_AUDIT_DELETE_CHILD_KEYS | Shows information about a change to delete child keys from a metabase property | See Metabase Audit Delete Child Keys |
| 4503 | IISADMIN_AUDIT_COPY_KEY | Shows information about a change to copy a metabase key | See Metabase Audit Copy Key |
| 4504 | IISADMIN_AUDIT_RENAME_KEY | Shows information about a change to rename a metabase key | See Metabase Audit Rename Key |
| 4505 | IISADMIN_AUDIT_SET_DATA | Shows information about a change to add or change data in a metabase property | See Metabase Audit Set Data |
| 4506 | IISADMIN_AUDIT_DELETE_DATA | Shows information about a change to delete data from a metabase property | See Metabase Audit Delete Data |
| 4507 | IISADMIN_AUDIT_DELETE_ALL_DATA | Shows information about a change to delete all data from the metabase properties | See Metabase Audit Delete All Data |
| 4508 | IISADMIN_AUDIT_COPY_DATA | Shows information about a change to copy data from the metabase properties | See Metabase Audit Copy Data |
| 4509 | IISADMIN_AUDIT_SET_LAST_CHANGETIME | Shows the last time a change was made to the metabase | See Metabase Audit Set Last Change Time |
| 4510 | IISADMIN_AUDIT_RESTORE | Shows information about a change to restore the metabase | See Metabase Audit Restore |
| 4511 | IIADMIN_AUDIT_DELETE_BACKUP | Shows information about a change to delete the metabase backup | See Metabase Audit Delete Backup |
| 4512 | IISADMIN_AUDIT_IMPORT | Shows information about a change to import data into the metabase | See Metabase Audit Import |

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Common Output in Metabase Auditing Event Messages

Applies To: Windows Server 2003 with SP1

All event messages produced when metabase auditing is enabled provide information specific to the event with content that is common to every event message. For example, the specific event message for a metabase audit set data event differs from the specific event message provided for a metabase audit copy data event. The set data event message includes the common output and information about the change to a specific property ID and property name, the old value and new value. While the event log message for a copy data metabase auditing event provides the common output and the destination and the source path of the copy action.

The Event Viewer details pane provides the following information about the metabase auditing event:

- Event type - Success or Failure
- When the event occurred - Date and time
- Source - IIS Metabase
- Event ID - specific to the metabase auditing event
- User - the user account
- Computer - short name for the computer

The metabase auditing event log message common output appears in the event properties view. The following table describes the event information common to each metabase auditing event log message.

**Message Output - Common**

| Log Output Name | Description |
|---|---|
| Primary User Name | The name of the user account that tried to change the metabase file manually when this information is possible to obtain, or System when a local system account was used. |
| Primary User Domain | The user domain when that information is possible to obtain. |
| Primary Logon ID | The ID used to log on to and change the metabase, when this information is possible to obtain. |
| Caller PID | The process ID of the application pool. |
| Caller Image Path | The absolute path of the resource. |
| Result | The variable representing the result. |

For information about using metabase auditing, see Metabase Auditing.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Add Key

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_ADD_KEY table shows the information about the attempt to add a metabase key, the name of the metabase key, and whether the attempt was successful.

**IISADMIN_AUDIT_ADD_KEY**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4500 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_ADD_KEY |
| Metabase Add Key | %n%n |
| Path | %4%n |
| Result | %7%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about using metabase auditing, see Metabase Auditing.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Delete Key

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_DELETE_KEY table shows information about the attempt to delete a metabase key, the name of the metabase key, and whether the attempt was successful.

**IISADMIN_AUDIT_DELETE_KEY**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4501 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_DELETE_KEY |
| Metabase Delete Key | %n%n |
| Path | %4%n |
| Result | %7%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Delete Child Keys

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_DELETE_CHILD_KEYS table shows information about the attempt to delete child metabase keys, the name of the keys, and whether the attempt was successful.

**IISADMIN_AUDIT_DELETE_CHILD_KEYS**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4502 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_DELETE_CHILD_KEYS |
| Metabase Delete Child Keys | %n%n |
| Path | %4%n |
| Result | %7%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Copy Key

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_COPY_KEYS table shows information about the attempt to copy data from a metabase key, the name of the key, and whether the attempt was successful.

**IISADMIN_AUDIT_COPY_KEYS**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4503 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_COPY_KEY |
| Metabase Copy Key | %n%n |
| Source Path | %4%n |
| Destination Path | %5%n |
| Result | %8%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Rename Key

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_RENAME_KEY table shows information about the attempt to rename a metabase key, the name of the metabase key, and whether the attempt was successful.

**IISADMIN_AUDIT_RENAME_KEYS**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4504 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_RENAME_KEY |
| Metabase Rename Key | %n%n |
| Source Path | %4%n |
| New Key Name | %5%n |
| Result | %8%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Set Data

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_SET_DATA_KEY table shows information about the attempt to add or change data in a metabase key, the name of the key, and whether the attempt was successful.

**IISADMIN_AUDIT_SET_DATA_KEY**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4505 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_SET_DATA_KEY |
| Metabase Set Data | %n%n |
| Path | %4%n |
| Property ID | %5%n |
| Property Name | %6%n |
| Old Value | %7%n |
| New Value | %8%n |
| Result | %11%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Delete Data

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_DELETE_DATA table shows information about the attempt to delete data from a property in a metabase key, the name of the key, and whether the attempt was successful.

**IISADMIN_AUDIT_DELETE_DATA**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4506 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_DELETE_DATA |
| Metabase Delete Data | %n%n |
| Path | %4%n |
| Property ID | %5%n |
| Property Name | %6%n |
| Old Value | %7%n |
| Result | %10%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Delete All Data

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_DELETE_ALL_DATA table shows information about the attempt to delete all data from the metabase, and whether the attempt was successful.

**IISADMIN_AUDIT_DELETE_ALL_DATA**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4507 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_DELETE_ALL_DATA |
| Metabase Delete All Data | %n%n |
| Path | %4%n |
| Result | %7%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Copy Data

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_COPY_DATA table shows information about the attempt to copy data from the metabase and whether the attempt was successful.

**IISADMIN_AUDIT_COPY_DATA**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4508 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_COPY_DATA |
| Metabase Copy Data | %n%n |
| Source Path | %4%n |
| Destination Path | %5%n |
| Result | %8%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Set Last Change Time

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_SET_LAST_CHANGETIME table shows information about the attempt to set the last change timestamp for the metabase and whether the attempt was successful.

**IISADMIN_AUDIT_SET_LAST_CHANGETIME**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4509 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_SET_LAST_CHANGETIME |
| Metabase Set Last Change Time | %n%n |
| Path | %4%n |
| Result | %7%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Restore

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_RESTORE table shows information about the attempt to restore a backup file to the metabase and whether the attempt was successful.

**IISADMIN_AUDIT_RESTORE**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4510 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_RESTORE |
| Metabase Restore | %n%n |
| Name of Backup | %4%n |
| Result | %7%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Delete Backup

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_DELETE_BACKUP table shows information about the attempt to delete the metabase backup and whether the attempt was successful.

**IISADMIN_AUDIT_DELETE_BACKUP**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4511 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_DELETE_BACKUP |
| Metabase Delete Backup | %n%n |
| Name of Backup | %4%n |
| Result | %7%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Metabase Audit Import

Applies To: Windows Server 2003 with SP1

The IISADMIN_AUDIT_IMPORT table shows information about the attempt to import a file into the metabase, and whether the attempt was successful.

**IISADMIN_AUDIT_IMPORT**

| Output | Description or Placeholder |
|---|---|
| Message ID | 4512 |
| Severity | Success |
| Symbolic Name | IISADMIN_AUDIT_IMPORT |
| Metabase Import | %n%n |
| Source Path | %4%n |
| Destination Path | %5%n |
| Filename | %6%n |
| Result | %9%n |

**Note**

The placeholders designated by percent sign, n or a number, and a following percent sign, such as in the metabase key, path, and result fields are replaced by the actual values when the metabase auditing event occurs.

For information about the common output produced for all metabase auditing event messages, see Common Output in Metabase Auditing Event Messages.

To view sample output for one metabase auditing event message, see Sample Metabase Auditing Event Log.

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Sample Metabase Auditing Event Log

Applies To: Windows Server 2003 with SP1

The Security log shows Metabase Audit Log Message Event ID 4505 indicating a set data metabase auditing event that changed the MaxConnections property ID 1014 from the old value of 4294967295 to a new value of 2. The Security log event header information is repeated in the detailed view of the event.

```
Event Type:Success Audit
Event Source:IIS-METABASE
Event Category:Object Access
Event ID:4505
Date:8/22/2006
Time:1:49:42 PM
User:NT AUTHORITY\SYSTEM
Computer:ATestMachine
Description:
Metabase Set Data
 Primary User Name: SYSTEM
 Primary User Domain: NT AUTHORITY
 Primary Logon ID: (0x0,0x3E7)
 Path: /LM/W3SVC
 Property ID: 1014
 Property Name: MaxConnections
 Old Value: 4294967295
 New Value: 2
 Caller PID: 2588
 Caller Image Path: \Device\HarddiskVolume1\WINNT\system32\inetsrv\inetinfo.exe - EWR
 Result: 0x0
```

For information about using metabase auditing, see Metabase Auditing.

For information about enabling metabase auditing, see Auditing the Metabase.

# Performance Counters Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Performance counters measure quantities of information or data, according to the number, size, duration, and rate of data being requested or received. You should not expect to get a list of details from a counter, such as a list of error messages. Instead, use performance counters to get quantities, such as the number of error message since startup or the rate at which error messages are being generated.

You can access counter data through the System Monitor (formerly known as Performance Monitor) or through the "Simple Network Management Protocol" (SNMP) service in the Help and Support for Windows 2003.

The names and descriptions of the following counters are included in this reference:

System Monitor Counters

- WWW Service Counters
- ASP System Monitor Counters
- WWW Service Cache Counters
- IIS Global Counters
- FTP Service Counters

SNMP Counters

- SNMP FTP Service Counters
- SNMP HTTP Service Counters

You must be a member of the Administrators group on the IIS server or a member of the Performance Logging group to retrieve valid data from the following counters. If you are not a member of one of these groups, or if the service you are monitoring is not running, the performance counters return a zero value.

**WWW Service Counters**
These counters help you determine how well the World Wide Web Publishing Service (WWW service) processes requests. These counters can be configured to monitor performance for all sites on a server (using the _Total instance) or for individual Web sites. Note that a few counters included with IIS 5.*x* are now obsolete and are marked as such. Obsolete counters return a zero value.

| Counter | Description |
|---|---|
| Total Bytes Sent | Number of data bytes that have been sent by the WWW service (since service startup). This counter is new in IIS 6.0. |
| Bytes Sent/Sec | Rate that data bytes are sent by the WWW service, per second. This counter was called Bytes Sent in IIS 5.*x*. |
| Total Bytes Received | Total bytes of data received by the WWW service since service startup. This counter is new in IIS 6.0. |
| Bytes Received/Sec | Rate at which data bytes are received by the WWW service, per second. This counter was called Bytes Received in IIS 5.*x*. |
| Total Bytes Transferred | Total bytes of data that were sent and received by the WWW service since service startup. This counter is new in IIS 6.0. |

| | |
|---|---|
| Bytes Total/Sec | Sum of bytes sent/sec and bytes received/sec. This counter was called Bytes Total in IIS 5.*x*. |
| Total Files Sent | Number of user-mode files that were sent by the WWW service since service startup (not including cache hits). Note that this counter does not increment when files are being served from the kernel-mode cache. For more information, see Kernel: URI Cache Hits under "WWW Service Cache Counters". Also note that in IIS 5.0, when a 404 error (File not found) was returned to the client, this counter was not incremented, even if a file was sent with a custom error message. In IIS 6.0, the counter is incremented whether a custom or standard error is returned to the browser. |
| Files Sent/Sec | Rate that files are sent, per second. |
| Total Files Received | Number of files received by the WWW service since service startup. |
| Files Received/Sec | Rate at which files are received by the WWW service, per second. |
| Total Files Transferred | Sum of files sent and received by the WWW service since service startup. Note that this counter does not increment when files are being served from the kernel-mode cache. For more information, see Kernel: URI Cache Hits under "WWW Service Cache Counters". |
| Files/Sec | Rate at which files are sent and received by the WWW service, per second. |
| Current Anonymous Users | Number of users who currently have an anonymous request using the WWW service. Note that IIS 6.0 counts each request as a new user. This is a change from IIS 5.*x*, which counted each connection as a new user. Current Users (Anonymous or Non Anonymous) are now the number of requests currently being worked on by the server. |
| Current Non Anonymous Users | Number of users who currently have a non anonymous request using the WWW service. Note that IIS 6.0 counts each request as a unique user. This is a change from IIS 5.*x*, which counted each connection as a new user. Current Users (Anonymous or Non Anonymous) are now the number of requests currently being worked on by the server. |
| Total Anonymous Users | Number of users who established an anonymous request since the WWW service started. Note that IIS 6.0 counts each request as a unique user. This counter does not increment when files are being served from the kernel cache. For more information, see Kernel: URI Cache Hits under "WWW Service Cache Counters". |
| Anonymous Users/Sec | Rate at which users are making anonymous requests to the WWW service. Note that IIS 6.0 counts each request as a new user. |
| Total NonAnonymous Users | Number of users who had nonanonymous requests to the WWW service since service startup. Note that IIS 6.0 counts each request as a unique user. This counter does not increment when files are being served from the kernel cache. For more information, see Kernel: URI Cache Hits under "WWW Service Cache Counters". |
| NonAnonymous Users/Sec | Rate at which users are making nonanonymous requests to the WWW service, per second. Note that IIS 6.0 counts each request as a new user. |
| Maximum NonAnonymous Users | Maximum number of users who made concurrent nonanonymous requests to the WWW service since service startup. Note that IIS 6.0 counts each request as a new user. |
| Current Connections | Current number of active connections to the WWW service. The behavior of this counter is unchanged from IIS 5.*x*. |
| Maximum Connections | Maximum number of simultaneous connections to the WWW service since service startup. |
| Total Connection Attempts | Number of attempted connections to the WWW service since service startup. |
| Connection Attempts/Sec | Rate of attempted connections to the WWW service since service startup, per second. |

| | |
|---|---|
| Total Logon Attempts | Number of logons attempted on the WWW service since service startup. |
| Logon Attempts/Sec | Rate at which logons to the WWW service are being attempted, per second. |
| Total Options Requests | Number of HTTP requests that were using the **OPTIONS** method since WWW service startup. |
| Options Requests/Sec | Rate at which HTTP requests using the **OPTIONS** method are made. |
| Total Get Requests | Number of HTTP requests that were using the **GET** method since WWW service startup. |
| Get Requests/Sec | Rate at which HTTP requests using the **GET** method are made to the WWW service, per second. |
| Total Post Requests | Number of HTTP requests that were using the **POST** method since WWW service startup. |
| Post Requests/Sec | Rate at which requests using the **POST** method are made to the WWW service, per second. |
| Total Head Requests | Number of HTTP requests that were made using the **HEAD** method since WWW service startup. |
| Head Requests/Sec | Rate at which HTTP requests using the **HEAD** method are made to the WWW service, per second. |
| Total Put Requests | Number of HTTP requests that were using the **PUT** method since WWW service startup. |
| Put Requests/Sec | Rate at which HTTP requests using the **PUT** method are made to the WWW service, per second. |
| Total Delete Requests | Number of HTTP requests that were using the **DELETE** method since WWW service startup. |
| Delete Requests/Sec | Rate at which HTTP requests using the **DELETE** method are made to the WWW service, per second. |
| Total Trace Requests | Number of HTTP requests that were using the **TRACE** method since WWW service startup. |
| Trace Requests/Sec | Rate at which HTTP requests using the **TRACE** method are made to the WWW service, per second. |
| Total Move Requests | Number of HTTP requests that were using the **MOVE** method since WWW service startup. |
| Move Requests/Sec | Rate at which HTTP requests using the **MOVE** method are made to the WWW service, per second. |
| Total Copy Requests | Number of HTTP requests that were using the **COPY** method since WWW service startup. |
| Copy Requests/Sec | Rate at which HTTP requests using the **COPY** method are made to the WWW service, per second. |
| Total Mkcol Requests | Number of HTTP requests that were using the **MKCOL** method since WWW service startup. |

| | |
|---|---|
| Mkcol Requests/Sec | Rate at which HTTP requests using the **MKCOL** method are made to the WWW service, per second. |
| Total Propfind Requests | Number of HTTP requests that were using the **PROPFIND** method since WWW service startup. |
| Propfind Requests/Sec | Rate at which HTTP requests using the **PROPFIND** method are made to the WWW service, per second. |
| Total Proppatch Requests | Number of HTTP requests that were using the **PROPPATCH** method since WWW service startup. |
| Proppatch Requests/Sec | Rate at which HTTP requests using the **PROPPATCH** method are made to the WWW service, per second. |
| Total Search Requests | Number of HTTP requests that were using the **SEARCH** method since WWW service startup. |
| Search Requests/Sec | Rate at which HTTP requests using the **SEARCH** method are made to the WWW service, per second. |
| Total Lock Requests | Number of HTTP requests that were using the **LOCK** method since WWW service startup. |
| Lock Requests/Sec | Rate at which HTTP requests using the **LOCK** method are made to the WWW service, per second. |
| Total Unlock Requests | Number of HTTP requests that were using the **UNLOCK** method since WWW service startup. |
| Unlock Requests/Sec | Rate at which HTTP requests using the **UNLOCK** method are made to the WWW service, per second. |
| Total Other Requests Methods | Number of HTTP requests that do not use the **OPTIONS**, **GET**, **HEAD**, **POST**, **PUT**, **DELETE**, **TRACE**, **MOVE**, **COPY**, **MKCOL**, **PROPFIND**,**PROPPATCH**, **SEARCH**, **LOCK**, or **UNLOCK** methods since WWW service startup. May include **LINK** or other methods supported by gateway applications. |
| Other Requests Methods/Sec | Rate at which HTTP requests are made to the WWW service that do not use the methods listed for Total Other Requests, per second. |
| Total Method Requests | Number of HTTP requests since WWW service startup. |
| Total Method Requests/Sec | Rate at which all HTTP requests are received, per second. |
| Total Not Found Errors | Number of requests that could not be satisfied by the server because the requested document could not be found (since service startup). Generally reported as HTTP error 404. |
| Not Found Errors/Sec | Rate of errors due to requests that couldn't be satisfied by the server because the requested document couldn't be found, per second. |
| Total Locked Errors | Number of requests that could not be satisfied by the server because the requested document was locked (since service startup).Generally reported as HTTP error 423. |
| Locked Errors/Sec | Rate at which requests couldn't be satisfied because the document requested was locked, per second. |
| Current CGI Requests | Current number of CGI requests that are being processed simultaneously by the WWW service. |
| Total CGI Requests | Number of all CGI requests that have been made since WWW service startup. |

| | |
|---|---|
| CGI Requests/Sec | Rate at which CGI requests are being processed simultaneously by the WWW service, per second. |
| Maximum CGI Requests | Maximum number of CGI requests simultaneously processed by the WWW service since service startup. |
| Current ISAPI Extension Requests | Current number of ISAPI extension requests that are being processed simultaneously by the WWW service. |
| Total ISAPI Extension Requests | Number of all ISAPI extension requests that have been made since WWW service startup. |
| ISAPI Extension Requests/Sec | Rate of ISAPI extension requests that are processed simultaneously by the WWW service, per second. |
| Maximum ISAPI Extension Requests | Maximum number of ISAPI extension requests that have been processed simultaneously by the WWW service. |
| Current CAL count for authenticated users | This counter is no longer valid. The value will always be zero. |
| Maximum CAL count for authenticated users | This counter is no longer valid. The value will always be zero. |
| Total count of failed CAL requests for authenticated users | This counter is no longer valid. The value will always be zero. |
| Current CAL count for SSL connections | This counter is no longer valid. The value will always be zero. |
| Maximum CAL count for SSL connections | This counter is no longer valid. The value will always be zero. |
| Total count of failed CAL requests for SSL connections | This counter is no longer valid. The value will always be zero. |
| Total Allowed Async I/O Requests | This counter is no longer valid. The value will always be zero. |
| Total Blocked Async I/O Requests | This counter is no longer valid. The value will always be zero. |
| Total Rejected Async I/O Requests | This counter is no longer valid. The value will always be zero. |

| | |
|---|---|
| Current Blocked Async I/O Requests | This counter is no longer valid. The value will always be zero. |
| Measured Async I/O Bandwidth Usage | This counter is no longer valid. The value will always be zero. |
| Total blocked bandwidth bytes | Total bytes blocked due to bandwidth throttling settings (since service startup). This counter is new in IIS 6.0. |
| Current blocked bandwidth bytes | Current bytes temporarily blocked due to bandwidth throttling settings (since service startup). This counter is new in IIS6.0. |
| Service Uptime | Uptime for WWW service or Web site. |
| MaxCalSsl | This counter is no longer valid. The value will always be zero. |

**ASP System Monitor Counters**

If you are running Active Server Pages (ASP) on your server, these counters can help you determine how well the pages are performing. The counters are designed to monitor server performance and cannot be configured for individual sites.

| Counter | Description |
|---|---|
| ASP Debugging Requests | Number of requests for debugging documents since WWW service startup. |
| Errors During Script Runtime | Number of requests failed due to runtime errors. |
| Errors From ASP Preprocessor | Number of requests failed due to preprocessor errors. |
| Errors From Script Compilers | Number of requests failed due to script compilation errors. |
| Errors/Sec | Number of errors per second. |
| Request Bytes In Total | Total size, in bytes, of all requests. |
| Request Bytes Out Total | Total size, in bytes, of responses sent to clients. This does not include standard HTTP response headers. |
| Request Execution Time | Number of milliseconds that it took to execute the most recent request. |
| Request Wait Time | Number of milliseconds the most recent request was waiting in the queue. |
| Requests Disconnected | Number of requests that were disconnected due to communication failure. |
| Requests Executing | Number of requests currently executing. |
| Requests Failed Total | Number of requests failed due to errors, authorization failure, and rejections. |
| Requests Not Authorized | Number of requests failed due to insufficient access rights. |
| Requests Not Found | Number of requests for files that were not found. |
| Requests Queued | Number of requests waiting for service from the queue. |
| Requests Rejected | Number of requests not executed because there were insufficient resources to process them. |
| Requests Succeeded | Number of requests that executed successfully. |
| Requests Timed Out | Number of requests that timed out. |

| | |
|---|---|
| Requests Total | Number of requests since the service was started. |
| Requests/Sec | Number of requests executed per second. |
| Script Engines Cached | Number of script engines in cache. |
| Session Duration | Number of milliseconds that the most recent session persisted. |
| Sessions Current | The current number of sessions being serviced. |
| Sessions Timed Out | Number of sessions timed out. |
| Sessions Total | Number of sessions since the service was started. |
| Templates Cached | Number of templates currently cached. |
| Template Cache Hit Rate | Percent of requests found in template cache. |
| Template Notifications | Number of templates invalidated in the cache due to change notification. |
| Transactions Aborted | Number of transactions aborted. |
| Transactions Committed | Number of transactions committed. |
| Transactions Pending | Number of transactions in progress. |
| Transactions Total | Number of transactions since the service was started. |
| Transactions/Sec | Transactions started per second. |
| In Memory Templates Cached | Number of compiled templates cached in memory. |
| In Memory Template Cache Hit Rate | Percent of requests found in cache in memory |
| Script Engine Cache Hit Rate | Percent of requests found in script engine cache. |
| Engine Flush Notifications | Number of engines invalidated in the cache due to change notification. |

**WWW Service Cache Counters**

The WWW service and FTP service do not share a common cache. The caches are split into two separate performance objects: one for FTP service (IIS global counters) and one for the WWW service (WWW service cache counters).

WWW service cache counters are designed to monitor server performance and cannot be configured to monitor individual sites.

| **Counter** | **Description** |
|---|---|
| Active Flushed Entries | Active Flushed Entries are file handlers cached in user mode that will be closed when all current transfers complete. This counter is new in IIS 6.0. |
| Current File Cache Memory Usage | Current number of bytes used for the user-mode file cache. This counter is new in IIS 6.0. |
| Current Files Cached | Current number of files whose content is in the user-mode cache. This counter is new in IIS 6.0. |
| Total Files Cached | Number of files whose content was ever added to the user-mode cache since WWW service startup. This counter is new in IIS 6.0. |
| File Cache Hits | Number of successful lookups in the user-mode file cache since WWW service startup. This counter is new in IIS 6.0. |

| | |
|---|---|
| File Cache Hits % | Ratio of user-mode file cache hits to total cache requests since WWW service startup. This counter is new in IIS 6.0. |
| File Cache Misses | Number of unsuccessful look ups in the user-mode file cache since WWW service startup. This counter is new in IIS 6.0. |
| File Cache Flushes | Number of files removed from the user-mode cache since WWW service startup. This counter is new in IIS 6.0. |
| Maximum File Cache Memory Usage | Maximum number of bytes used for user-mode file cache since service startup. This counter is new in IIS 6.0. |
| Total Flushed Files | Number of file handles that have been removed from the user-mode cache since WWW service startup. This counter is new in IIS 6.0. |
| Current URIs Cached | Number of URI information blocks that are currently in the user-mode cache. This counter is new in IIS 6.0. |
| Total URIs Cached | Number of URI information blocks that have been added to the user-mode cache since WWW service startup. This counter is new in IIS 6.0. |
| URI Cache Hits | Number of successful look ups in the user-mode URI cache since WWW service startup. This counter is new in IIS 6.0. |
| URI Cache Misses | Number of unsuccessful lookups in the user-mode URI cache since WWW service startup. This counter is new in IIS 6.0. |
| URI Cache Hit % | Ratio of URI Cache Hits to total cache requests since WWW service startup. This counter is new in IIS 6.0. |
| URI Cache Flushes | Total URI cache flushes since WWW service startup. This counter is new in IIS 6.0. |
| Total Flushed URIs | Number of URI information blocks that have been removed from the user-mode cache since WWW service startup. This counter is new in IIS 6.0. |
| Current Metadata Cached | Current number of metadata information blocks in the user-mode cache. This counter is new in IIS 6.0. |
| Total Metadata Cached | Number of metadata information blocks that have been added to the user-mode cache since service startup. This counter is new in IIS 6.0. |
| Metadata Cache Hits | Number of successful lookups in the user-mode metadata cache since WWW service startup. This counter is new in IIS 6.0. |
| Metadata Cache Misses | Number of unsuccessful look ups in the user-mode metadata cache since WWW service startup. This counter is new in IIS 6.0. |
| Metadata Cache Hits % | Ratio of successful lookups to total metadata cache requests. This counter is new in IIS 6.0. |
| Metadata Cache Flushes | Number of user-mode metadata cache flushes since WWW service startup. This counter is new in IIS 6.0. |
| Total Flushed Metadata | Number of metadata information blocks that have been removed from the user-mode cache since WWW service startup. This counter is new in IIS 6.0. |
| Kernel: Current URIs Cached | URI information blocks currently cached by the kernel. Kernel counters reflect all HTTP.sys activity, not just IIS activity. This counter is new in IIS 6.0. |
| Kernel: Total URIs Cached | Number of URI information blocks added to the kernel URI cache since WWW service startup. Kernel counters reflect all HTTP.sys activity, not just IIS activity. This counter is new in IIS 6.0. |
| Kernel: URI Cache Hits | Number of successful look ups in the kernel URI cache since WWW service startup. Kernel counters reflect all HTTP.sys activity, not just IIS activity. This counter is new in IIS 6.0. |

| | |
|---|---|
| Kernel: URI Cache Hits% | The ratio of kernel URI cache hits to total cache requests since WWW service startup. Kernel counters reflect all HTTP.sys activity, not just IIS activity. This counter is new in IIS 6.0. |
| Kernel: URI Cache Hits/Sec | Number of kernel URI cache hits, per second. Kernel counters reflect all HTTP.sys activity, not just IIS activity. This counter is new in IIS 6.0. |
| Kernel URI Cache Misses | Number of unsuccessful lookups in the kernel URI cache since WWW service startup. Kernel counters reflect all HTTP.sys activity, not just IIS activity. This counter is new in IIS 6.0. |
| Kernel: URI Cache Flushes | Number of kernel URI cache flushes since WWW service startup. Kernel counters reflect all HTTP.sys activity, not just IIS activity. This counter is new in IIS 6.0. |
| Kernel: Total Flushed URIs | Number of URI information blocks that have been removed from the kernel cache since WWW service startup. Kernel counters reflect all HTTP.sys activity, not just IIS activity. This counter is new in IIS 6.0. |

**IIS Global Counters**

These counters monitor FTP, SMTP, and NNTP services as a whole and cannot be configured to monitor individual sites. This is a change from IIS 5.x, in which global counters monitored WWW and FTP services. Additionally, some IIS 5.x counters are now obsolete and return a zero value.

| **Note** |
|---|
| If the service you wish to monitor (FTP, SMTP, or NNTP) is not installed or is not running, the counters return a zero value. |

| **Counter** | **Description** |
|---|---|
| Measured Async I/O Bandwidth Usage | This counter is no longer valid. The value will always be zero. |
| Total Blocked Async I/O Requests | This counter is no longer valid. The value will always be zero. |
| Total Rejected Async I/O Requests | This counter is no longer valid. The value will always be zero. |
| Current Blocked Async I/O Requests | This counter is no longer valid. The value will always be zero. |
| Measured Async I/O Bandwidth Usage | This counter is no longer valid. The value will always be zero. |
| Current Files Cached | Current number of files whose content is in the cache. |
| Total Files Cached | Number of files whose content was ever added to the cache. |
| File Cache Hits | Number of successful lookups in the file cache. |
| File Cache Misses | Number of unsuccessful lookups in the file cache. |
| File Cache Hits % | Ratio of File Cache Hits to total cache requests. |
| File Cache Flushes | Number of file cache flushes since server startup. |
| Current File Cache Memory Usage | Current number of bytes used for file cache. |
| Maximum File Cache Memory Usage | Maximum number of bytes used for file cache. |
| Active Flushed Entries | Number of cached file handles that will be closed when all current transfers complete. |
| Total Flushed Files | Number of file handles that have been removed from the cache since service start up. |
| Current URIs Cached | URI information blocks currently in the cache. |
| Total URIs Cached | Number of URI information blocks ever added to the cache. |

| | |
|---|---|
| URI Cache Hits | Number of successful lookups in the URI cache. |
| Uri Cache Misses | Number of unsuccessful lookups in the URI cache. |
| URI Cache Hit % | The ratio of URI Cache Hits to total cache requests. |
| URI Cache Flushes | URI Cache Flushes since server startup. |
| Total Flushed URIs | Number of URI information blocks that have been removed from the cache since service start up. |
| Current Blobs Cached | BLOB information blocks currently in the cache. |
| Total Blobs Cached | Number of BLOB information blocks ever added to the cache. |
| Blob Cache Hits | Number of successful lookups in the BLOB cache. |
| Blob Cache Misses | Number of unsuccessful lookups in the BLOB cache. |
| Blob Cache HitRatio | The ratio of BLOB Cache Hits to total cache requests. |
| Blob Cache Flushes | BLOB Cache Flushes since server startup. |
| Total Flushed Blobs | Number of BLOB information blocks that have been removed from the cache since service start up. |

**FTP Service Counters**

These counters are designed to monitor FTP service performance. They can be configured to monitor a server or individual FTP sites.

| **Counter** | **Description** |
|---|---|
| Bytes Sent/Sec | Rate that data bytes are sent by the FTP service. |
| Bytes Received/Sec | Rate that data bytes are received by the FTP service. |
| Bytes Total/Sec | Sum of Bytes Sent/sec and Bytes Received/sec. This is the total rate of bytes transferred by the FTP service. |
| Total Files Sent | Total number of files sent by the FTP service since service startup. |
| Total Files Received | Total number of files received by the FTP service since service startup. |
| Total Files Transferred | Sum of Files Sent and Files Received. This is the total number of files transferred by the FTP service since service startup. |
| Current Anonymous Users | Number of users who currently have an anonymous connection using the FTP service. |
| Current NonAnonymous Users | Number of users who currently have a nonanonymous connection using the FTP service. |
| Total Anonymous Users | Number of users who established an anonymous connection withthe FTP service (since service startup). |
| Total NonAnonymous Users | Number of users who had nonanonymous connections to the FTP service since FTP service startup. |
| Maximum Anonymous Users | Maximum number of users who established concurrent anonymous connections using the FTP service (since service startup). |
| Maximum NonAnonymous Users | Maximum number of users who established concurrent nonanonymous connections using the FTP service (since service startup). |

| | |
|---|---|
| Current Connections | Current number of connections established with the FTP service. |
| Maximum Connections | Maximum number of simultaneous connections established with the FTP service. |
| Total Connection Attempts (all instances) | Number of connections that have been attempted using the FTP service (since service startup). This counter is for all instances listed. |
| Total Logon Attempts | Number of logons that have been attempted using the FTP service (since service startup). |
| FTP Service Uptime | FTP Service and FTP site uptime, in seconds. |

**SNMP Counters**

If you prefer to use Simple Network Management Protocol (SNMP) services to monitor your Web servers, IIS exposes a set of counters via SNMP. These counters are defined in 3 .mib files in the %windir%\system32 folder. These files are named inetsrv.mib, http.mib and ftp.mib. Inetsrv.mib is a container object for the other two and does not contain any counters.

For information about installing and using the SNMP service see "SNMP" in Help and Support Center for Windows Server 2003.

**SNMP FTP Service Counters**

These SNMP counters are designed to monitor FTP service performance.

| Counter | Description |
|---|---|
| totalBytesSentHighWord | This is the high 32-bits of the total number of bytes sent by the FTP server. |
| totalBytesSentLowWord | This is the low 32-bits of the total number of bytes sent by the FTP server. |
| totalBytesReceivedHighWord | This is the high 32-bits of the total number of bytes received by the FTP server. |
| totalBytesReceivedLowWord | This is the low 32-bits of the total number of bytes received by the FTP server. |
| totalFilesSent | This is the total number of files sent by this FTP server. |
| totalFilesReceived | This is the total number of files received by this FTP server. |
| currentAnonymousUsers | This is the number of anonymous users currently connected to the FTP server. |
| currentNonAnonymousUsers | This is the number of nonanonymous users currently connected to the FTP server. |
| totalAnonymousUsers | This is the total number of anonymous users that have ever connected to the FTP server. |
| totalNonAnonymousUsers | This is the total number of nonanonymous users that have ever connected to the FTP server. |
| maxAnonymousUsers | This is the maximum number of anonymous users simultaneously connected to the FTP server. |
| maxNonAnonymousUsers | This is the maximum number of nonanonymous users simultaneously connected to the FTP server. |
| currentConnections | This is the current number of connections to the FTP server. |
| maxConnections | This is the maximum number of simultaneous connections to the FTP server. |
| connectionAttempts | This is the number of connection attempts that have been made to the FTP server. |
| logonAttempts | This is the number of logon attempts that have been made to the FTP server. |

**SNMP HTTP Service Counters**

These counters are designed to monitor World Wide Web Publishing (WWW) service performance.

| Counter | Description |
|---|---|

| | |
|---|---|
| totalBytesSentHighWord | This is the high 32-bits of the total number of bytes sent by the WWW service (since service startup). |
| totalBytesSentLowWord | This is the low 32-bits of the total number of bytes sent by the WWW service (since service startup). |
| totalBytesReceivedHighWord | This is the high 32-bits of the total number of bytes received by the WWW service (since service startup). |
| totalBytesReceivedLowWord | This is the low 32-bits of the total number of bytes received by the WWW service (since service startup). |
| totalFilesSent | This is the total number of files sent by this WWW service (since service startup). |
| totalFilesReceived | This is the total number of files received by the WWW service (since service startup). |
| currentAnonymousUsers | This is the number of users who currently have an anonymous connection using the WWW service. |
| currentNonAnonymousUsers | This is the number of users who currently have a nonanonymous connection using the WWW service. |
| totalAnonymousUsers | This is the total number of users who established an anonymous connection with the WWW service (since service startup). |
| totalNonAnonymousUsers | This is the total number of users who established a nonanonymous connection with the WWW service (since service startup). |
| maxAnonymousUsers | This is the maximum number of users who established concurrent anonymous connections using the WWW service (since service startup). |
| maxNonAnonymousUsers | This is the maximum number of concurrent nonanonymous connections to the WWW service (since service startup). |
| currentConnections | This is the current number of connections established with the WWW service. |
| maxConnections | This is the maximum number of concurrent connections established with the WWW service (since service startup). |
| connectionAttempts | This is the number of connections that have been attempted using the WWW service (since service startup). |
| logonAttempts | This is the number of logons attempts to the WWW service (since service startup). |
| totalOptions | This is the number of HTTP requests using the OPTIONS method (since service startup). |
| totalGets | This is the number of HTTP requests using the GET method (since service startup). Get requests are the most common HTTP request. |
| totalPosts | This is the number of HTTP requests using the POST method (since service startup). |
| totalHeads | This is the number of HTTP requests using the HEAD method (since service startup). Head requests generally indicate a client is querying the state of a document they already have to see if it needs to be refreshed. |
| totalPuts | This is the number of HTTP requests using the PUT method (since service startup). |
| totalDeletes | This is the number of HTTP requests using the DELETE method (since service startup). Delete requests are generally used for file removals. |
| totalTraces | This is the number of HTTP requests using the TRACE method (since service startup). Trace requests allow the client to see what is being received at the end of the request chain and use the information for diagnostic purposes. |

| | |
|---|---|
| totalMove | This is the number of HTTP requests using the MOVE method (since service startup). Move requests are used for moving files and directories. |
| totalCopy | This is the number of HTTP requests using the COPY method (since service startup). Copy requests are used for copying files and directories. |
| totalMkcol | This is the number of HTTP requests using the MKCOL method (since service startup). Mkcol requests are used to create directories on the server. |
| totalPropfind | This is the number of HTTP requests using the PROPFIND method (since service startup). Propfind requests retrieve property values on files and directories. |
| totalProppatch | This is the number of HTTP requests using the PROPPATCH method (since service startup). Proppatch requests set property values on files and directories. |
| totalSearch | This is the number of HTTP requests using the SEARCH method (since service startup). Search requests are used to query the server to find resources that match a set of conditions provided by the client. |
| totalLock | This is the number of HTTP requests using the LOCK method (since service startup). Lock requests are used to lock a file for one user so that only that user can modify the file. |
| totalUnlock | This is the number of HTTP requests using the UNLOCK method (since service startup). Unlock requests are used to remove locks from files. |
| totalOthers | This is the number of HTTP requests that are not OPTIONS, GET,HEAD, POST, PUT, DELETE, TRACE, MOVE, COPY, MKCOL, PROPFIND,PROPPATCH, SEARCH, LOCK or UNLOCK methods (since service startup). |
| currentCGIRequests | This is the current number of CGI requests being simultaneously processed by the WWW service. |
| currentBGIRequests | This is the current number of ISAPI requests being simultaneously processed by the WWW service. |
| totalCGIRequests | This is the total number of CGI requests (since service startup). |
| totalBGIRequests | This is the total ISAPI requests received (since service startup). |
| maxCGIRequests | This is the maximum number of CGI requests simultaneously processed by the WWW service (since service startup). |
| maxBGIRequests | This is the maximum number of ISAPI requests simultaneously processed by the WWW service (since service startup). |
| currentBlockedRequests | This is the current number of requests that have been temporarily blocked by the WWW service due to bandwidth throttling settings. |
| totalBlockedRequests | This is the total number of requests that have been temporarily blocked by the WWW service due to bandwidth throttling settings. |
| totalAllowedRequests | This is the total number of requests that the WWW service has allowed when the bandwidth throttling settings on the server have been enabled (since service startup). |
| totalRejectedRequests | This is the total number of requests rejected due to bandwidth throttling settings (since service startup). |
| totalNotFoundErrors | This is the number of requests that couldn't be satisfied by the server because the requested document could not be found (since service startup). These are generally reported as an HTTP 404 error code to the client. |
| totalLockedErrors | This is the number of requests that couldn't be satisfied by the server because the requested was locked (since service startup). These are generally reported as an HTTP 423 error code to the client. |

| | |
|---|---|
| measuredBandwidth | This is the I/O bandwidth used by the WWW service, averaged over a minute. |
| currentCALsforAuthenticatedUsers | This counter is no longer valid. Value will always be zero. |
| maxCALsforAuthenticatedUsers | This counter is no longer valid. Value will always be zero. |
| totalCALFailedAuthenticatedUser | This counter is no longer valid. Value will always be zero. |
| currentCALsforSecureConnections | This counter is no longer valid. Value will always be zero. |
| maxCALsforSecureConnections | This counter is no longer valid. Value will always be zero. |
| totalCALFailedSecureConnection | This counter is no longer valid. Value will always be zero. |

# IIS 6.0 Performance Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Microsoft® Windows® Server 2003 operating system obtains performance data from components in your computer as those components are utilized. This data is described as a performance object and is typically named for the component that is generating the data. For example, the Web Service performance object is a collection of performance data about Internet Information Services (IIS) 6.0. Each performance object provides performance counters that represent data on specific aspects of a system or service. When you install IIS 6.0, all of the IIS performance counters, with the exception of the SNMP counters, appear in System Monitor and in Performance Logs and Alerts. You can also use Windows Instrumentation Management (WMI) to view and use these counters and the SNMP counters.

**In This Section**


Web Service Counters for the WWW Service

Web Service Cache Counters for the WWW Service

FTP Service Counters

Internet Information Services Global Counters

SNMP Counters

Active Server Pages Performance Counters

ASP.NET Performance Counters


**Related Sections**

- For information about performance monitoring and tuning, see Optimizing IIS 6.0 Performance.

# Web Service Counters for the WWW Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The Web Service counters help you determine how well the World Wide Web Publishing Service (WWW service) processes requests. The WWW service is a user-mode service. These counters also reflect the processing that occurs in the kernel-mode driver, HTTP.sys.

Configure these counters to monitor performance for individual Web sites or for all sites on a server (by using the _Total instance). Note that some counters that were included with IIS 5.*x* are now obsolete and, therefore, return a zero value. Obsolete counters are listed in Table D.8 and Table D.9.

Table D.1 describes the Bytes counters for the Web Service performance object.

**Table D.1 Bytes (Sent, Received, and Transferred) Counters**

| Counter | Description |
|---|---|
| Total Bytes Sent | The number of data bytes that have been sent by the WWW service since the service started. This counter is new in IIS 6.0. |
| Bytes Sent/sec | The rate, in seconds, at which data bytes have been sent by the WWW service. |
| Total Bytes Received | The total bytes of data that have been received by the WWW service since the service started. This counter is new in IIS 6.0. |
| Bytes Received/sec | The rate, in seconds, at which data bytes have been received by the WWW service. |
| Total Bytes Transferred | The total number of bytes of data that have been sent and received by the WWW service since the service started. This counter is new in IIS 6.0. |
| Bytes Total/sec | The sum of Bytes Sent/sec and Bytes Received/sec. |

Table D.2 describes the Files counters for the Web Service performance object.

**Table D.2 Files (Sent, Received, and Transferred) Counters**

| Counter | Description |
|---|---|
| Total Files Sent | The number of user-mode files that have been sent by the WWW service since the service started. This counter does not include cache hits. Note that this counter does not increment when files are being served from the kernel-mode cache. For more information, see Kernel: URI Cache Hits in Table D.13. |
| Files Sent/sec | The rate, in seconds, at which files have been sent. |
| Total Files Received | The number of files that have been received by the WWW service since the service started. |
| Files Received/sec | The rate, in seconds, at which files have been received by the WWW service. |
| Total Files Transferred | The sum of Total Files Sent and Total Files Received by the WWW service since the service started. Note that this counter does not increment when files are being served from the kernel-mode cache. For more information, see Kernel: URI Cache Hits in Table D.13. |
| Files/sec | The rate, in seconds, at which files have been sent and received by the WWW service. |

Table D.3 describes the Anonymous Users and NonAnonymous Users counters for the Web Service performance object. In this

context, nonanonymous users are authenticated users.

| Note |
| --- |
| With Anonymous and NonAnonymous Users counters, IIS6.0 counts each request as a new user. This behavior differs from IIS5.*x*, which counted each connection as a new user. |

**Table D.3 Anonymous Users and NonAnonymous Users Counters**

| Counter | Description |
| --- | --- |
| Current Anonymous Users | The number of users who currently have an anonymous request pending with the WWW service. In IIS 6.0, Current Users (Anonymous or NonAnonymous) is the number of requests currently being worked on by the server. |
| Current NonAnonymous Users | The number of users who currently have a nonanonymous request pending with the WWW service. In IIS 6.0, Current Users (Anonymous or NonAnonymous) is the number of requests currently being worked on by the server. |
| Total Anonymous Users | The number of users who have established an anonymous request since the WWW service started. This counter does not increment when files are being served from the kernel cache. For more information, see Kernel: URI Cache Hits in Table D.13. |
| Anonymous Users/sec | The rate, in seconds, at which users have made anonymous requests to the WWW service. |
| Total NonAnonymous Users | The number of users who have made nonanonymous requests to the WWW service since the service started. For more information, see Kernel: URI Cache Hits in Table D.13. |
| NonAnonymous Users/sec | The rate, in seconds, at which users have made nonanonymous requests to the WWW service. |
| Maximum NonAnonymous Users | The maximum number of users who have made concurrent nonanonymous requests to the WWW service since the service started. |

Table D.4 describes the Connections and Attempts counters for the Web Service performance object.

**Table D.4 Connections and Attempts Counters**

| Counter | Description |
| --- | --- |
| Current Connections | The number of active connections to the WWW service. |
| Maximum Connections | The maximum number of simultaneous connections made to the WWW service since the service started. |
| Total Connection Attempts | The number of connections to the WWW service that have been attempted since the service started. |
| Connection Attempts/sec | The rate, in seconds, at which connections to the WWW service have been attempted since the service started. |
| Total Logon Attempts | The number of attempts to log on to the WWW service that have occurred since the service started. |
| Logon Attempts/sec | The rate, in seconds, at which attempts to log on to the WWW service have occurred. |

Table D.5 describes the Requests counters for the Web Service performance object.

**Table D.5 Requests Counters**

| Counter | Description |
| --- | --- |

| | |
|---|---|
| Total Options Requests | The number of HTTP requests that have used the **OPTIONS** method since the WWW service started. |
| Options Requests/sec | The rate, in seconds, at which HTTP requests that use the **OPTIONS** method have been made. |
| Total Get Requests | The number of HTTP requests that have used the **GET** method since the WWW service started. |
| Get Requests/sec | The rate, in seconds, at which HTTP requests that use the **GET** method have been made to the WWW service. |
| Total Post Requests | The number of HTTP requests that have used the **POST** method since the WWW service started. |
| Post Requests/sec | The rate, in seconds, at which requests that use the **POST** method have been made to the WWW service. |
| Total Head Requests | The number of HTTP requests that have used the **HEAD** method since the WWW service started. |
| Head Requests/sec | The rate, in seconds, at which HTTP requests that use the **HEAD** method have been made to the WWW service. |
| Total Put Requests | The number of HTTP requests that have used the **PUT** method since the WWW service started. |
| Put Requests/sec | The rate, in seconds, at which HTTP requests that use the **PUT** method have been made to the WWW service. |
| Total Delete Requests | The number of HTTP requests that have used the **DELETE** method since the WWW service started. |
| Delete Requests/sec | The rate, in seconds, at which HTTP requests that use the **DELETE** method have been made to the WWW service. |
| Total Trace Requests | The number of HTTP requests that have used the **TRACE** method since the WWW service started. |
| Trace Requests/sec | The rate, in seconds, at which HTTP requests that use the **TRACE** method have been made to the WWW service. |
| Total Move Requests | The number of HTTP requests that have used the **MOVE** method since the WWW service started. |
| Move Requests/sec | The rate, in seconds, at which HTTP requests that use the **MOVE** method have been made to the WWW service. |
| Total Copy Requests | The number of HTTP requests that have used the **COPY** method since the WWW service started. |
| Copy Requests/sec | The rate, in seconds, at which HTTP requests that use the **COPY** method have been made to the WWW service. |

| | |
|---|---|
| Total Mkcol Requests | The number of HTTP requests that have used the **MKCOL** method since the WWW service started. |
| Mkcol Requests/sec | The rate, in seconds, at which HTTP requests that use the **MKCOL** method have been made to the WWW service. |
| Total Propfind Requests | The number of HTTP requests that have used the **PROPFIND** method since the WWW service started. |
| Propfind Requests/sec | The rate, in seconds, at which HTTP requests that use the **PROPFIND** method have been made to the WWW service. |
| Total Proppatch Requests | The number of HTTP requests that have used the **PROPPATCH** method since the WWW service started. |
| Proppatch Requests/sec | The rate, in seconds, at which HTTP requests that use the **PROPPATCH** method have been made to the WWW service. |
| Total Search Requests | The number of HTTP requests that have used the **SEARCH** method since the WWW service started. |
| Search Requests/sec | The rate, in seconds, at which HTTP requests that use the **SEARCH** method have been made to the WWW service. |
| Total Lock Requests | The number of HTTP requests that have used the **LOCK** method since the WWW service started. |
| Lock Requests/sec | The rate, in seconds, at which HTTP requests that use the **LOCK** method have been made to the WWW service. |
| Total Unlock Requests | The number of HTTP requests that have used the **UNLOCK** method since the WWW service started. |
| Unlock Requests/sec | The rate, in seconds, at which HTTP requests that use the **UNLOCK** method have been made to the WWW service. |
| Total Other Request Methods | The number of HTTP requests that did not use the **OPTIONS**, **GET**, **HEAD**, **POST**, **PUT**, **DELETE**, **TRACE**, **MOVE**, **COPY**, **MKCOL**, **PROPFIND**, **PROPPATCH**, **SEARCH**, **LOCK**, or **UNLOCK** methods since the WWW service started. Can include **LINK** or other methods supported by gateway applications. |
| Other Request Methods/sec | The rate, in seconds, at which HTTP requests that do not use the methods listed for the Total Other Requests Methods counter have been made to the WWW service. |
| Total Method Requests | The number of HTTP requests that have been made since the WWW service started. |
| Total Method Requests/sec | The rate, in seconds, at which all HTTP requests have been received. |

Table D.6 describes the Errors counters for the Web Service performance object.

**Table D.6 Errors (Not Found and Locked) Counters**

| Counter | Description |
| --- | --- |
| Total Not Found Errors | The number of requests that have been made since the service started that were not satisfied by the server because the requested document was not found. Usually reported as HTTP error 404. |
| Not Found Errors/sec | The rate, in seconds, at which requests were not satisfied by the server because the requested document was not found. |
| Total Locked Errors | The number of requests that have been made since the service started that could not be satisfied by the server because the requested document was locked. Usually reported as HTTP error 423. |
| Locked Errors/sec | The rate, in seconds, at which requests were not satisfied because the requested document was locked. |

Table D.7 describes the Common Gateway Interface (CGI) Requests and ISAPI Extension Requests counters for the Web Service performance object.

**Table D.7 CGI Requests and ISAPI Extension Requests Counters**

| Counter | Description |
| --- | --- |
| Current CGI Requests | The number of CGI requests that are being processed simultaneously by the WWW service. |
| Total CGI Requests | The number of all CGI requests that have been made since the WWW service started. |
| CGI Requests/sec | The rate, in seconds, at which CGI requests are being processed simultaneously by the WWW service. |
| Maximum CGI Requests | The maximum number of CGI requests that have been processed simultaneously by the WWW service since the service started. |
| Current ISAPI Extension Requests | The number of ISAPI extension requests that are being processed simultaneously by the WWW service. |
| Total ISAPI Extension Requests | The number of ISAPI extension requests that have been made since the WWW service started. |
| ISAPI Extension Requests/sec | The rate, in seconds, at which ISAPI extension requests are being processed by the WWW service. |
| Maximum ISAPI Extension Requests | The maximum number of ISAPI extension requests that were processed simultaneously by the WWW service. |

Table D.8 lists the CAL (client access license) Count counters for the Web Service performance object, which are not valid in IIS 6.0. The value of these counters is always zero.

**Table D.8 CAL Count Counters**

| Counter |
| --- |
| Current CAL count for authenticated users |
| Maximum CAL count for authenticated users |
| Total count of failed CAL requests for authenticated users |
| Current CAL count for SSL connections |
| Maximum CAL count for SSL connections |

| |
|---|
| Total count of failed CAL requests for SSL connections |

Table D.9 describes other miscellaneous counters, some of which are no longer valid, for the Web Service performance object.

**Table D.9 Async I/O Requests, Bandwidth Bytes, and Miscellaneous Counters**

| Counter | Description |
|---|---|
| Total Allowed Async I/O Requests | This counter is no longer valid; its value is always zero. |
| Total Blocked Async I/O Requests | This counter is no longer valid; its value is always zero. |
| Total Rejected Async I/O Requests | This counter is no longer valid; its value is always zero. |
| Current Blocked Async I/O Requests | This counter is no longer valid; its value is always zero. |
| Measured Async I/O Bandwidth Usage | This counter is no longer valid; its value is always zero. |
| Total Blocked Bandwidth Bytes | This counter is no longer valid; its value is always zero. |
| Current Blocked Bandwidth Bytes | This counter is no longer valid; its value is always zero. |
| Service Uptime | The uptime for the WWW service or a Web site. |
| MaxCalSsl | This counter is no longer valid; its value is always zero. |

# Web Service Cache Counters for the WWW Service

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

The WWW service and FTP service do not share a common cache. Instead, the caches are split into two separate performance objects: one for FTP service and one for the WWW service. WWW service cache counters are designed to monitor server performance only; therefore, you cannot configure them to monitor individual sites.

For information about FTP service counters, see Internet Information Services Global Counters later in this appendix.

Table D.10 describes the File counters for the Web Service Cache performance object. These counters are new in IIS 6.0.

**Table D.10 File Counters for the Web Service Cache Object**

| Counter | Description |
|---|---|
| Active Flushed Entries | The number of user-mode cache entries that have been flushed, though memory is still allocated for these entries. The allocated memory will be released after all current transfers complete. |
| Current File Cache Memory Usage | The number of bytes currently used for the user-mode file cache. |
| Current Files Cached | The number of files whose content is currently in the user-mode cache. |
| Total Files Cached | The number of files whose content has been added to the user-mode cache since the WWW service started. |
| File Cache Hits | The number of successful lookups in the user-mode file cache that have occurred since the WWW service started. |
| File Cache Hits % | The ratio of user-mode file cache hits to total cache requests that have been made since the WWW service started up. |
| File Cache Misses | The number of unsuccessful lookups in the user-mode file cache that have been made since the WWW service started. |
| File Cache Flushes | The number of files that have been removed from the user-mode cache since the WWW service started. |
| Maximum File Cache Memory Usage | The maximum number of bytes that have been used for the user-mode file cache since the WWW service started. |
| Total Flushed Files | The number of file handles that have been removed from the user-mode cache since the WWW service started. |

Table D.11 describes the URI counters for the Web Service Cache performance object. These counters are new in IIS 6.0.

**Table D.11 URI Counters for the Web Service Cache Object**

| Counter | Description |
|---|---|
| Current URIs Cached | The number of URI information blocks that are currently stored in the user-mode cache. |
| Total URIs Cached | The number of URI information blocks that have been added to the user-mode cache since the WWW service started. |

| | |
|---|---|
| URI Cache Hits | The number of successful lookups that have been made in the user-mode URI cache since the WWW service started. |
| URI Cache Misses | The number of unsuccessful lookups that have been made in the user-mode URI cache since the WWW service started. |
| URI Cache Hits % | The ratio of URI Cache Hits to total cache requests that have occurred since the WWW service started. |
| URI Cache Flushes | The total number of URI cache flushes that have occurred since the WWW service started. |
| Total Flushed URIs | The number of URI information blocks that have been removed from the user-mode cache since the WWW service started. |

Table D.12 describes the Metadata counters for the Web Service Cache performance object. These counters are new in IIS 6.0.

**Table D.12 Metadata Counters for the Web Service Cache Object**

| Counter | Description |
|---|---|
| Current Metadata Cached | The current number of metadata information blocks in the user-mode cache. |
| Total Metadata Cached | The number of metadata information blocks that have been added to the user-mode cache since the WWW service started. |
| Metadata Cache Hits | The number of successful lookups in the user-mode metadata cache that have occurred since the WWW service started. |
| Metadata Cache Misses | The number of unsuccessful lookups in the user-mode metadata cache that have occurred since the WWW service started. |
| Metadata Cache Hits % | The ratio of successful lookups to total metadata cache requests. |
| Metadata Cache Flushes | The number of user-mode metadata cache flushes that have occurred since the WWW service started. |
| Total Flushed Metadata | The number of metadata information blocks that have been removed from the user-mode cache since the WWW service started. |

Table D.13 describes the Kernel counters for the Web Service Cache performance object. These counters are new in IIS 6.0.

| **Note** |
|---|
| Kernel counters reflect all HTTP.sys activity, not just IIS activity. |

**Table D.13 Kernel Counters for the Web Service Cache Object**

| Counter | Description |
|---|---|
| Kernel: Current URIs Cached | The number of URI information blocks currently cached by the kernel. |
| Kernel: Total URIs Cached | The number of URI information blocks that have been added to the kernel URI cache since the WWW service started. |
| Kernel: URI Cache Hits | The number of successful lookups in the kernel URI cache that have occurred since the WWW service started. |

| | |
|---|---|
| Kernel: URI Cache Hits% | The ratio of Kernel: URI Cache Hits to total cache requests since the WWW service started. |
| Kernel: URI Cache Hits/sec | The average number of kernel URI cache hits that are being made per second. |
| Kernel URI Cache Misses | The number of unsuccessful lookups in the kernel URI cache that have occurred since the WWW service started. |
| Kernel: URI Cache Flushes | The number of kernel URI cache flushes that have occurred since the WWW service started. |
| Kernel: Total Flushed URIs | The number of URI information blocks that have been removed from the kernel cache since the WWW service started. |

# FTP Service Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use the counters in the following tables to monitor FTP service performance. You can configure the counters to monitor a server or an individual FTP site.

Table D.14 describes Bytes and Total Files counters for the FTP Service performance object.

**Table D.14 Bytes and Total Files Counters**

| Counter | Description |
|---|---|
| Bytes Sent/sec | The rate at which data bytes are being sent by the FTP service. |
| Bytes Received/sec | The rate at which data bytes are being received by the FTP service. |
| Bytes Total/sec | The sum of Bytes Sent/sec and Bytes Received/sec. |
| Total Files Sent | The total number of files that have been sent by the FTP service since the service started. |
| Total Files Received | The total number of files that have been received by the FTP service since the service started. |
| Total Files Transferred | The sum of Total Files Sent and Total Files Received. This is the total number of files transferred by the FTP service since the service started. |

Table D.15 describes Anonymous and NonAnonymous Users counters for the FTP Service performance object.

**Table D.15 Anonymous and NonAnonymous Users Counters**

| Counter | Description |
|---|---|
| Current Anonymous Users | The number of users who currently have an anonymous connection that was made by using the FTP service. |
| Current NonAnonymous Users | The number of users who currently have a nonanonymous connection that was made by using the FTP service. |
| Total Anonymous Users | The number of users who have established an anonymous connection with the FTP service since the service started. |
| Total NonAnonymous Users | The number of users who have established nonanonymous connections with the FTP service since the service started. |
| Maximum Anonymous Users | The maximum number of users who have established concurrent anonymous connections using the FTP service since the service started. |
| Maximum NonAnonymous Users | The maximum number of users who have established concurrent nonanonymous connections using the FTP service since the service started. |

Table D.16 describes Connections, Attempts, and Uptime counters for the FTP Service performance object.

**Table D.16 Connections, Attempts, and Uptime Counters**

| Counter | Description |
|---|---|
| Current Connections | The current number of connections that have been established with the FTP service. |

| | |
|---|---|
| Maximum Connections | The maximum number of simultaneous connections that have been established with the FTP service. |
| Total Connection Attempts (all instances) | The number of connections that have been attempted by using the FTP service since the service started. This counter applies to all instances listed. |
| Total Logon Attempts | The number of logons that have been attempted by using the FTP service since the service started. |
| FTP Service Uptime | The amount of time, in seconds, that the FTP service has been running. |

# Internet Information Services Global Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Use the Internet Information Services Global counters described in Table D.17 through Table D.20 to monitor FTP, SMTP, and NNTP services as a whole. You cannot configure these counters to monitor individual sites. This differs from IIS 5.*x*, in which global counters monitored the WWW and FTP services.

If the service that you want to monitor (FTP, SMTP, or NNTP) is not installed or is not running, the counters return a zero value.

Table D.17 lists the Async I/O counters that are no longer valid for the Internet Information Services Global performance object. The value of these counters is always zero.

**Table D.17 Async I/O Counters**

| Counter |
|---|
| Measured Async I/O Bandwidth Usage |
| Total Blocked Async I/O Requests |
| Total Rejected Async I/O Requests |
| Current Blocked Async I/O Requests |
| Measured Async I/O Bandwidth Usage |

Table D.18 describes File Cache counters for the Internet Information Services Global performance object.

**Table D.18 File Cache Counters**

| Counter | Description |
|---|---|
| Current Files Cached | The number of files whose content is currently in the cache. |
| Total Files Cached | The number of files whose content has been added to the cache since the service started. |
| File Cache Hits | The number of successful lookups in the file cache. |
| File Cache Misses | The number of unsuccessful lookups in the file cache. |
| File Cache Hits % | The ratio of File Cache Hits to the total number of cache requests. |
| File Cache Flushes | The number of file cache flushes that have occurred since the service started. |
| Current File Cache Memory Usage | The current number of bytes used for the file cache. |
| Maximum File Cache Memory Usage | The maximum number of bytes used for the file cache. |

Table D.19 describes Flushed and URI Cached counters for the Internet Information Services Global performance object.

**Table D.19 Flushed and URI Cached Counters**

| Counter | Description |
|---|---|
| Active Flushed Entries | The number of cached file handles that will close when all current transfers are complete. |
| Total Flushed Files | The number of file handles that have been removed from the cache since the service started. |
| Current URIs Cached | The number of URI information blocks that are currently in the cache. |

| | |
|---|---|
| Total URIs Cached | The number of URI information blocks that have been added to the cache. |
| URI Cache Hits | The number of successful lookups in the URI cache. |
| Uri Cache Misses | The number of unsuccessful lookups in the URI cache. |
| URI Cache Hits % | The ratio of URI Cache Hits to the total number of cache requests. |
| URI Cache Flushes | The number of URI cache flushes that have occurred since the server started. |
| Total Flushed URIs | The number of URI information blocks that have been removed from the cache since the service started. |

Table D.20 describes the BLOBs counters for the Internet Information Services Global performance object. A *BLOB* is a binary large object.

**Table D.20 BLOBs Counters**

| Counter | Description |
|---|---|
| Current BLOBs Cached | The BLOB information blocks currently in the cache. |
| Total BLOBs Cached | The number of BLOB information blocks that have been added to the cache. |
| BLOB Cache Hits | The number of successful lookups in the BLOB cache. |
| BLOB Cache Misses | The number of unsuccessful lookups in the BLOB cache. |
| BLOB Cache Hit % | The ratio of BLOB Cache Hits to the total number of cache requests. |
| BLOB Cache Flushes | The number of BLOB cache flushes that have occurred since the service started. |
| Total Flushed BLOBs | The number of BLOB information blocks that have been removed from the cache since the service started. |

# SNMP Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

IIS exposes the following counters for Simple Network Management Protocol (SNMP) services. These counters are defined in more detail in three .mib files — Inetsrv.mib, Http.mib and Ftp.mib — which you can find in the *Windir*\System32 folder. Inetsrv.mib is a container object for the other two files and, therefore, does not contain counters. To view the contents of the .mib files, open the files by using Notepad.

You cannot view the SNMP counters in System Monitor or Performance Logs and Alerts. Use WMI or a management information base (MIB) Browser tool instead.

For information about installing and using the SNMP service, see Checklist: Implementing the SNMP service in Help and Support Center for Microsoft® Windows® Server 2003.

**SNMP FTP Service Counters**

Use the SNMP counters in Table D.21 through Table D.23 to monitor FTP service performance.

Table D.21 describes the totalBytes and totalFiles counters that you can use to monitor an SMTP FTP service.

**Table D.21 totalBytes and totalFiles Counters**

| Counter | Description |
|---|---|
| totalBytesSentHighWord | The high 32-bits of the total number of bytes sent by the FTP server. |
| totalBytesSentLowWord | The low 32-bits of the total number of bytes sent by the FTP server. |
| totalBytesReceivedHighWord | The high 32-bits of the total number of bytes received by the FTP server. |
| totalBytesReceivedLowWord | The low 32-bits of the total number of bytes received by the FTP server. |
| totalFilesSent | The total number of files sent by this FTP server. |
| totalFilesReceived | The total number of files received by this FTP server. |

Table D.22 describes the AnonymousUsers and NonAnonymousUsers counters that you can use to monitor an SNMP FTP service.

**Table D.22 AnonymousUsers and NonAnonymousUsers Counters**

| Counter | Description |
|---|---|
| currentAnonymousUsers | The number of anonymous users currently connected to the FTP server. |
| currentNonAnonymousUsers | The number of nonanonymous users currently connected to the FTP server. |
| totalAnonymousUsers | The total number of anonymous users that have been connected to the FTP server. |
| totalNonAnonymousUsers | The total number of nonanonymous users that have been connected to the FTP server. |
| maxAnonymousUsers | The maximum number of anonymous users that have been simultaneously connected to the FTP server. |
| maxNonAnonymousUsers | The maximum number of nonanonymous users that have been simultaneously connected to the FTP server. |

Table D.23 describes the Connections and Attempts counters that you can use to monitor an SNMP FTP service.

**Table D.23 Connections and Attempts Counters**

| Counter | Description |
|---|---|
| currentConnections | The current number of connections to the FTP server. |
| maxConnections | The maximum number of simultaneous connections that have been made to the FTP server. |

| | |
|---|---|
| connectionAttempts | The number of attempts that were made to connect to the FTP server. |
| logonAttempts | The number of attempts that were made to log on to the FTP server. |

**SNMP HTTP Service Counters**

Use the SNMP HTTP counters in Table D.24 through Table D.31 to monitor WWW service performance.

Table D.24 describes the totalBytes and totalFiles counters that you can use to monitor an SNMP WWW service.

**Table D.24 totalBytes and totalFiles Counters**

| Counter | Description |
|---|---|
| totalBytesSentHighWord | The high 32-bits of the total number of bytes that have been sent by the WWW service since the service started. |
| totalBytesSentLowWord | The low 32-bits of the total number of bytes that have been sent by the WWW service since the service started. |
| totalBytesReceivedHighWord | The high 32-bits of the total number of bytes that have been received by the WWW service since the service started. |
| totalBytesReceivedLowWord | The low 32-bits of the total number of bytes that have been received by the WWW service since the service started. |
| totalFilesSent | The total number of files that have been sent by the WWW service since the service started. |
| totalFilesReceived | The total number of files that have been received by the WWW service since the service started. |

Table D.25 describes the AnonymousUsers and NonAnonymousUsers counters that you can use to monitor an SNMP WWW service.

**Table D.25 AnonymousUsers and NonAnonymousUsers Counters**

| Counter | Description |
|---|---|
| currentAnonymousUsers | The number of users who currently have an anonymous connection by using the WWW service. |
| currentNonAnonymousUsers | The number of users who currently have a nonanonymous connection by using the WWW service. |
| totalAnonymousUsers | The total number of users who have made an anonymous connection by using the WWW service since the service started. |
| totalNonAnonymousUsers | The total number of users who have made a nonanonymous connection by using the WWW service since the service started. |
| maxAnonymousUsers | The maximum number of users who have made concurrent anonymous connections by using the WWW service since the service started. |
| maxNonAnonymousUsers | The maximum number of users who have made concurrent nonanonymous connections to the WWW service since the service started up. |

Table D.26 describes the Connections and Attempts counters that you can use to monitor an SNMP WWW service.

**Table D.26 Connections and Attempts Counters**

| Counter | Description |
|---|---|
| currentConnections | The current number of connections made with the WWW service. |
| maxConnections | The maximum number of concurrent connections that have been made with the WWW service since the service started. |
| connectionAttempts | The number of connections that have been attempted by using the WWW service since the service started. |

| | |
|---|---|
| logonAttempts | The number of attempts to log on to the WWW service that have been made since the service started. |

Table D.27 describes the total counters (totalOptions, totalGets, and others) that you can use to monitor an SNMP WWW service.

**Table D.27 total Counters**

| Counter | Description |
|---|---|
| totalOptions | The number of HTTP requests that have been made using the **OPTIONS** method since the service started. |
| totalGets | The number of HTTP requests that have been made using the **GET** method since the service started. |
| totalPosts | The number of HTTP requests that have been made using the **POST** method since the service started. |
| totalHeads | The number of HTTP requests that have been made using the **HEAD** method since the service started. Head requests generally indicate that a client is querying the state of a document it already has to see if it needs to be refreshed. |
| totalPuts | The number of HTTP requests that have been made using the **PUT** method since the service started. |
| totalDeletes | The number of HTTP requests that have been made using the **DELETE** method since the service started. Delete requests are typically used for file removal. |
| totalTraces | The number of HTTP requests that have been made using the **TRACE** method since the service started. Trace requests allow the client to see what is being received at the end of the request chain and to use this information for diagnostic purposes. |
| totalMove | The number of HTTP requests that have been made using the **MOVE** method since the service started. Move requests move files and directories. |
| totalCopy | The number of HTTP requests that have been made using the **COPY** method since the service started. Copy requests copy files and directories. |
| totalMkcol | The number of HTTP requests that have been made using the **MKCOL** method since the service started. Mkcol requests create directories on the server. |
| totalPropfind | The number of HTTP requests that have been made using the **PROPFIND** method since the service started. Propfind requests retrieve property values for files and directories. |
| totalProppatch | The number of HTTP requests that have been made using the **PROPPATCH** method since the service started. Proppatch requests set property values for files and directories. |
| totalSearch | The number of HTTP requests that have been made using the **SEARCH** method since the service started. Search requests query the server to find resources that match a set of conditions provided by the client. |
| totalLock | The number of HTTP requests that have been made using the **LOCK** method since the service started. Lock requests lock a file so that only the user who locked it can modify the file. |
| totalUnlock | The number of HTTP requests that have been made using the **UNLOCK** method since the service started. Unlock requests remove locks from files. |
| totalOthers | The number of HTTP requests that have been made by methods other than **OPTIONS**, **GET**, **HEAD**, **POST**, **PUT**, **DELETE**, **TRACE**, **MOVE**, **COPY**, **MKCOL**, **PROPFIND**, **PROPPATCH**, **SEARCH**, **LOCK**, or **UNLOCK** methods since the service started. |

Table D.28 describes the CGIRequests and BGIRequests counters that you can use to monitor an SNMP WWW service.

**Table D.28 CGIRequests and BGIRequests Counters**

| Counter | Description |
|---|---|
| currentCGIRequests | The current number of CGI requests being processed simultaneously by the WWW service. |
| currentBGIRequests | The current number of ISAPI requests being processed simultaneously by the WWW service. |
| totalCGIRequests | The total number of CGI requests that have been made since the service started. |

| | |
|---|---|
| totalBGIRequests | The total number of ISAPI requests that have been received since the service started. |
| maxCGIRequests | The maximum number of CGI requests that have been processed simultaneously by the WWW service since the service started. |
| maxBGIRequests | The maximum number of ISAPI requests that have been processed simultaneously by the WWW service since the service started. |

Table D.29 lists the Requests counters (Blocked, Allowed, and Rejected), which are no longer valid in IIS 6.0. The value for these counters is always zero.

**Table D.29 Requests Counters (Blocked, Allowed, and Rejected)**

| Counter |
|---|
| currentBlockedRequests |
| totalBlockedRequests |
| totalAllowedRequests |
| totalRejectedRequests |

Table D.30 describes the Errors and Bandwidth counters that you can use to monitor an SNMP WWW service.

**Table D.30 Errors and Bandwidth Counters**

| Counter | Description |
|---|---|
| totalNotFoundErrors | The number of requests made since the service started that the server did not satisfy because the requested document was not found. This type of request is usually reported to the client as an HTTP 404 error message. |
| totalLockedErrors | The number of requests made since the service started that the server did not satisfy because the requested document was locked. This type of request is usually reported to the client as an HTTP 423 error message. |
| measuredBandwidth | The I/O bandwidth that the WWW service used, averaged over a minute. |

Table D.31 lists the CAL-related counters that are no longer valid for the SNMP WWW service. The value of these counters is always zero.

**Table D.31 CAL-related Counters**

| Counter |
|---|
| currentCALsforAuthenticatedUsers |
| maxCALsforAuthenticatedUsers |
| totalCALFailedAuthenticatedUser |
| currentCALsforSecureConnections |
| maxCALsforSecureConnections |
| totalCALFailedSecureConnection |

# Active Server Pages Performance Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

If you are running Active Server Pages (ASP) on your server, the ASP counters can help you determine how well the server or site is responding to ASP requests. The ASP counters are designed to monitor server performance; you cannot monitor individual ASP applications because ASP counters collect global data across the entire WWW service.

Table D.32 describes the ASP Debugging and Errors counters for the Active Server Pages performance object.

**Table D.32 ASP Debugging and Errors Counters**

| Counter | Description |
|---|---|
| ASP Debugging Requests | The number of requests for debugging documents that have been made since the WWW service started. |
| Errors During Script Runtime | The number of requests that failed because run-time errors occurred. |
| Errors From ASP Preprocessor | The number of requests that failed because preprocessor errors occurred. |
| Errors From Script Compilers | The number of requests that failed because script compilation errors occurred. |
| Errors/sec | The average number of errors that occurred per second. |

Table D.33 describes the Requests counters for the Active Server Pages performance object.

**Table D.33 Requests Counters for ASP Pages**

| Counter | Description |
|---|---|
| Request Bytes In Total | The total size, in bytes, of all requests. |
| Request Bytes Out Total | The total size, in bytes, of responses sent to clients. This total does not include standard HTTP response headers. |
| Request Execution Time | The number of milliseconds that it took to execute the most recent request. |
| Request Wait Time | The number of milliseconds that the most recent request waited in the queue. |
| Requests Disconnected | The number of requests that were disconnected because communication failed. |
| Requests Executing | The number of requests that are currently executing. |
| Requests Failed Total | The number of requests that failed due to errors, authorization failure, and rejections. |
| Requests Not Authorized | The number of requests that failed because access rights were insufficient. |
| Requests Not Found | The number of requests that were made for files that were not found. |
| Requests Queued | The number of requests that are waiting in the queue for service. |

| | |
|---|---|
| Requests Rejected | The number of requests that were not executed because there were insufficient resources to process them. |
| Requests Succeeded | The number of requests that executed successfully. |
| Requests Timed Out | The number of requests that timed out. |
| Requests Total | The number of requests that have been made since the service was started. |
| Requests/sec | The average number of requests that were executed per second. |

Table D.34 describes miscellaneous counters for the Active Server Pages performance object.

**Table D.34 Script Engines, Sessions, Templates, and Transactions Counters**

| Counter | Description |
|---|---|
| Script Engines Cached | The number of script engines in the cache. |
| Script Engine Cache Hit Rate | The percentage of requests that were found in the script engine cache. |
| Engine Flush Notifications | The number of engines invalidated in the cache because change notification occurred. |
| Session Duration | The length of time that the most recent session lasted, in milliseconds. |
| Sessions Current | The number of sessions currently being serviced. |
| Sessions Timed Out | The number of sessions that have timed out. |
| Sessions Total | The number of sessions that have run since the service was started. |
| Templates Cached | The number of templates that are currently cached. |
| Template Cache Hit Rate | The percentage of requests that have been found in the template cache. |
| Template Notifications | The number of templates that have been invalidated in the cache because change notification occurred. |
| In Memory Templates Cached | The number of compiled templates that are cached in memory. |
| In Memory Template Cache Hit Rate | The percentage of requests that have been found in the memory cache. |
| Transactions Aborted | The number of transactions that have been aborted. |
| Transactions Committed | The number of transactions that have been committed. This counter increments after page execution if the transaction does not abort. |
| Transactions Pending | The number of transactions that are in progress. |
| Transactions Total | The number of transactions that have occurred since the service was started. |
| Transactions/sec | The average number of transactions that have been started, per second. |

| Note |
| --- |
| The Templates Cached counter, Template Cache Hit Rate counter, and Template Notifications counter provide information about the total number of templates cached on disk and in memory. To determine the number of requests cached on disk only, subtract the value of the In Memory Templates Cached counter from the value of the Templates Cached counter. |

# ASP.NET Performance Counters

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

ASP.NET supports the following ASP.NET system performance counters, which aggregate information for all ASP.NET applications on a Web server computer, or, alternatively, apply generally to a system of ASP.NET servers running the same applications.

**ASP.NET Counters**

Use the ASP.NET counters in Table D.35 through Table D.38 to monitor ASP.NET system performance.

Table D.35 describes the Application Restarts and Applications Running counters for the ASP.NET performance object.

**Table D.35 Application Restarts and Applications Running Counters**

| Counter | Description |
|---|---|
| Application Restarts | The number of times that an application has been restarted since the Web service started. Application restarts are incremented with each **Application_OnEnd** event. An application restart can occur because changes were made to the Web.config file or to assemblies stored in the application's \Bin directory, or because too many changes occurred in Web Forms pages. Sudden increases in this counter can mean that your Web application is shutting down. If an unexpected increase occurs, be sure to investigate it promptly. This value resets every time IIS is restarted. |
| Applications Running | The number of applications that are running on the server computer. |

Table D.36 describes the Requests counters for the ASP.NET performance object.

**Table D.36 Requests Counters**

| Counter | Description |
|---|---|
| Requests Disconnected | The number of requests that were disconnected because a communication failure occurred. |
| Requests Queued | The number of requests in the queue waiting to be serviced. If this number increases as the number of client requests increases, the Web server has reached the limit of concurrent requests that it can process. The default maximum for this counter is 5,000 requests. You can change this setting in the computer's Machine.config file. |
| Requests Rejected | The total number of requests that were not executed because insufficient server resources existed to process them. This counter represents the number of requests that return a 503 HTTP status code, which indicates that the server is too busy. |
| Request Wait Time | The number of milliseconds that the most recent request waited in the queue for processing. |

Table D.37 describes the State Server Sessions counters for the ASP.NET performance object. These counter are available only on a computer on which the state server service (aspnet_state) is running.

**Table D.37 State Server Sessions Counters**

| Counter | Description |
|---|---|
| State Server Sessions Abandoned | The number of user sessions that were explicitly abandoned. These are sessions that have been ended by specific user actions, such as closing the browser or navigating to another site. |
| State Server Sessions Active | The number of active user sessions. |
| State Server Sessions Timed Out | The number of user sessions that are now inactive. In this case, the user is inactive, not the server. |

| | |
|---|---|
| State Server Sessions Total | The number of sessions created during the lifetime of the process. This counter represents the cumulative value of State Server Sessions Active, State Server Sessions Abandoned, and State Server Sessions Timed Out counters. |

Table D.38 describes the worker process counters for the ASP.NET performance object.

**Table D.38 Worker Process Counters**

| Counter | Description |
|---|---|
| Worker Process Restarts | The number of times that a worker process restarted on the server computer. A worker process can be restarted if it fails unexpectedly or when it is intentionally recycled. If worker process restarts increase unexpectedly, investigate immediately. |
| Worker Processes Running | The number of worker processes that are running on the server computer. |

**ASP.NET Applications Performance Counters**

ASP.NET supports the following application performance counters, which you can use to monitor the performance of a single instance of an ASP.NET application. A unique instance of these counters, named __Total__, aggregates counters for all applications on a Web server. The __Total__ instance is always available. When no applications are running on the server, the counters display zero.

Table D.39 describes the Anonymous Requests counters for the ASP.NET Applications performance object.

**Table D.39 Anonymous Requests Counters**

| Counter | Description |
|---|---|
| Anonymous Requests | The number of requests that use anonymous authentication. |
| Anonymous Requests/sec | The average number of requests that have been made per second that use anonymous authentication. |

Table D.40 describes the Cache Total counters for the ASP.NET Applications performance object.

**Table D.40 Cache Total Counters**

| Counter | Description |
|---|---|
| Cache Total Entries | The total number of entries in the cache. This counter includes both internal use of the cache by the ASP.NET framework and external use of the cache through exposed APIs. |
| Cache Total Hits | The total number of responses served from the cache. This counter includes both internal use of the cache by the ASP.NET framework and external use of the cache through exposed APIs. |
| Cache Total Misses | The number of failed cache requests. This counter includes both internal use of the cache by ASP.NET and external use of the cache through exposed APIs. |
| Cache Total Hit Ratio | The ratio of cache hits to cache misses. This counter includes both internal use of the cache by ASP.NET and external use of the cache through exposed APIs. |
| Cache Total Turnover Rate | The number of additions to and removals from the cache per second. Use this counter to help determine how efficiently the cache is being used. If the turnover rate is high, the cache is not being used efficiently. |

Table D.41 describes the Cache API counters for the ASP.NET Applications performance object.

**Table D.41 Cache API Counters**

| Counter | Description |
|---|---|
| Cache API Entries | The total number of entries in the application cache. |

| | |
|---|---|
| Cache API Hits | The total number of requests served from the cache when it was accessed only through the external cache APIs. This counter does not track use of the cache internally by ASP.NET. |
| Cache API Misses | The total number of requests to the cache that failed when the cache was accessed through the external cache APIs. This counter does not track use of the cache internally by ASP.NET. |
| Cache API Hit Ratio | The cache hit-to-miss ratio when the cache was accessed through external cache APIs. This counter does not track use of the cache internally by ASP.NET. |
| Cache API Turnover Rate | The number of additions to and removals from the cache per second, when it is used through the external APIs (excluding internal use by the ASP.NET framework). This counter is useful for determining how effectively the cache is being used. If the turnover is great, then the cache is not being used effectively. |

Table D.42 describes the Errors counters for the ASP.NET Applications performance object.

**Table D.42 Errors Counters**

| Counter | Description |
|---|---|
| Errors During Preprocessing | The number of errors that occurred during parsing. Excludes compilation and run-time errors. |
| Errors During Compilation | The number of errors that occurred during dynamic compilation. Excludes parser and run-time errors. |
| Errors During Execution | The total number of errors that occurred during the execution of an HTTP request. Excludes parser and compilation errors. |
| Errors Unhandled During Execution | The total number of unhandled errors that occurred during the execution of HTTP requests.<br><br>An *unhandled error* is any uncaught run-time exception that escapes user code on the page and enters the ASP.NET internal error-handling logic. Exceptions occur in the following circumstances:<br><br>• When custom errors are enabled, an error page is defined, or both.<br><br>• When the **Page_Error** event is defined in user code and either the error is cleared (by using the **HttpServerUtility.ClearError** method) or a redirect is performed. |
| Errors Unhandled During Execution/sec | The number of unhandled exceptions that occurred per second during the execution of HTTP requests. |
| Errors Total | The total number of errors that occurred during the execution of HTTP requests. Includes parser, compilation, or run-time errors. This counter represents the sum of the Errors During Compilation, Errors During Preprocessing, and Errors During Execution counters. A well-functioning Web server should not generate errors. |
| Errors Total/sec | The average number of errors that occurred per second during the execution of HTTP requests. Includes any parser, compilation, or run-time errors. |

Table D.43 describes the Output Cache counters for the ASP.NET Applications performance object.

**Table D.43 Output Cache Counters**

| Counter | Description |
|---|---|
| Output Cache Entries | The total number of entries in the output cache. |
| Output Cache Hits | The total number of requests serviced from the output cache. |

| | |
|---|---|
| Output Cache Misses | The number of output-cache requests that failed per application. |
| Output Cache Hit Ratio | The percentage of total requests that were serviced from the output cache. |
| Output Cache Turnover Rate | The average number of additions to and removals from the output cache per second. If the turnover is great, the cache is not being used effectively. |

Table D.44 describes the Request Bytes counters for the ASP.NET Applications performance object.

**Table D.44 Request Bytes Counters**

| Counter | Description |
|---|---|
| Request Bytes In Total | The total size, in bytes, of all requests. |
| Request Bytes Out Total | The total size, in bytes, of responses sent to a client. This does not include standard HTTP response headers. |

Table D.45 describes the Requests counters for the ASP.NET Applications performance object.

**Table D.45 Requests Counters**

| Counter | Description |
|---|---|
| Requests Executing | The number of requests that are currently executing. |
| Requests Failed | The total number of failed requests. All status codes greater than or equal to 400 increment this counter.<br><br>**Note:** Requests that cause a 401 status code increment this counter and the Requests Not Authorized counter. Requests that cause a 404 or 414 status code increment this counter and the Requests Not Found counter. Requests that cause a 500 status code increment this counter and the Requests Timed Out counter. |
| Requests In Application Queue | The number of requests in the application request queue. |
| Requests Not Found | The number of requests that failed because resources were not found (status code 404, 414). |
| Requests Not Authorized | The number of requests that failed because of lack of authorization (status code 401). |
| Requests Succeeded | The number of requests that executed successfully (status code 200). |
| Requests Timed Out | The number of requests that timed out (status code 500). |
| Requests Total | The total number of requests that have been made since the service started. |
| Requests/sec | The average number of requests that have been executed per second. This counter represents the current throughput of the application. |

Table D.46 describes the Session State and SQL Server Connections counters for the ASP.NET Applications performance object.

**Table D.46 Session State and SQL Server Connections Counters**

| Counter | Description |
|---|---|
| Session State Server Connections Total | The total number of session-state connections that were made to a computer on which out-of-process session-state data is stored. |

| | |
|---|---|
| Session SQL Server Connections Total | The total number of session-state connections that were made to the Microsoft® SQL Server™ database in which session-state data is stored. |

Table D.47 describes the Sessions counters for the ASP.NET Applications performance object. These counters are supported only with in-memory session state.

**Table D.47 Sessions Counters**

| Counter | Description |
|---|---|
| Sessions Active | The number of sessions that are active. |
| Sessions Abandoned | The number of sessions that have been explicitly abandoned. |
| Sessions Timed Out | The number of sessions that timed out. |
| Sessions Total | The total number of sessions. |

Table D.48 describes the Transactions counters for the ASP.NET Applications performance object.

**Table D.48 COM+ Transactions Counters**

| Counter | Description |
|---|---|
| Transactions Aborted | The number of COM+ transactions that were aborted. |
| Transactions Committed | The number of COM+ transactions that were committed. This counter increments after page execution if the transaction does not abort. |
| Transactions Pending | The number of COM+ transactions that are in progress. |
| Transactions Total | The total number of COM+ transactions that have occurred since the service was started. |
| Transactions/sec | The average number of COM+ transactions that were started per second. |

Table D.49 describes a few miscellaneous counters for ASP.NET Applications.

**Table D.49 Miscellaneous Counters for ASP.NET Applications**

| Counter | Description |
|---|---|
| Compilations Total | The total number of times that the Web server process dynamically compiled requests for files with .aspx, .asmx, .ascx, or .ashx extensions (or a code-behind source file).<br><br>**Note:** This number initially climbs to a peak value as requests are made to all parts of an application. After compilation occurs, however, the resulting binary compilation is saved on disk, where it is reused until its source file changes. This means that, even when a process restarts, the counter can remain at zero (be inactive) until the application is modified or redeployed. |
| Debugging Requests | The number of requests that occurred while debugging was enabled. |
| Pipeline Instance Count | The number of active request pipeline instances for the specified ASP.NET application. Because only one execution thread can run within a pipeline instance, this number represents the maximum number of concurrent requests that are being processed for a specific application. In most circumstances, it is better for this number be low when the server is busy, because this means that the CPU is well used. |

# Timeouts Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Internet Information Services (IIS) has several time-out values that are set by default when you install a member of the Microsoft Windows Server-TM- 2003 family of operating systems. Time-out values allow the server to specify how long server resources are allocated to specific tasks or clients.

Time-out settings are configurable from IIS Manager as well. For more information about setting the various time-out values, see Setting Connection Timeouts.

The following table lists the time-out values in IIS and where you can set them.

| Property Name | Metabase Identifier | Available in IIS Manager |
|---|---|---|
| AspScriptTimeout Metabase Property | MD_ASP_SCRIPTTIMEOUT | Yes |
| AspSessionTimeout Metabase Property | MD_ASP_SESSIONTIMEOUT | Yes |
| ConnectionTimeout Metabase Property | MD_CONNECTION_TIMEOUT | Yes |
| CGITimeout Metabase Property | MD_SCRIPT_TIMEOUT | No |
| ServerListenTimeout Metabase Property | MD_SERVER_LISTEN_TIMEOUT | No |
| PoolIdcTimeout Metabase Property | MD_POOL_IDC_TIMEOUT | No |
| AspQueueTimeout Metabase Property | MD_ASP_QUEUETIMEOUT | No |
| HeaderWaitTimeout Metabase Property | MD_HEADER_WAIT_TIMEOUT | No |
| MinFileBytesPerSec Metabase Property | MD_MIN_FILE_BYTES_PER_SEC | No |

# Redirect Reference

Updated: June 29, 2009

Applies To: Windows Server 2003, Windows Server 2003 with SP1

Redirecting a client request is one way to ensure that users get the correct page, if your site is under construction or has changed identity. You can redirect client requests to a directory on the same Web server or to a different URL.

**Configuring URL and Wildcard Redirection**

You can configure redirection to a URL and wildcard redirection programmatically. Redirection variables and flags allow you to specify targets and behavior by changing the **HttpRedirect** metabase property.

The following example shows the format of a simple redirection string.

`http://` *DestinationURL* `[,` *Flag[s]* `]`

Where *DestinationURL* can include the following redirect variables to pass portions of the original URL with the destination URL. You can use more than one variable in a single redirection string.

| Variable | Description | Example |
|---|---|---|
| $S | Passes the suffix of the requested URL to the new URL. The suffix is the portion of the original URL that remains after the redirected URL is substituted.<br><br>If the EXACT_DESTINATION flag is not set, the resulting destination URL will have the name of the requested file appended as a folder name, as well as the file name itself. | If the virtual directory that is mapped to the following Web site:<br><br>`http://fabrikam/Scripts`<br><br>is configured to be redirected to the following exact URL:<br><br>`http://fabrikam/NewScripts$S`<br><br>A new request for the following URL:<br><br>`http://fabrikam/Scripts/Program.exe`<br><br>Would be redirected automatically to the following URL:<br><br>`http://fabrikam/NewScripts/Program.exe` |
| $P | Passes parameters (such as querystring parameters) in the original URL to the new URL, without the question mark (?). | If the virtual directory that is mapped to the following Web site:<br><br>`http://fabrikam/Scripts`<br><br>is configured to be redirected to the following URL:<br><br>`http://fabrikam/NewScripts$P`<br><br>A new request for the following URL:<br><br>`http://fabrikam/Scripts/File.asp?var1=5&var2=6`<br><br>Would be redirected automatically to the following URL:<br><br>`http://fabrikam/NewScriptsvar1=5&var2=6` |

| | | |
|---|---|---|
| $Q | Passes parameters (such as querystring parameters) in the original URL to the new URL, including the question mark (?). | If the virtual directory that is mapped to the following Web site:<br>`http://fabrikam/Scripts`<br>is configured to be redirected to the following exact URL:<br>`http://fabrikam/NewScripts$S$Q`<br>A new request for the following URL:<br>`http://fabrikam/Scripts/File.asp?var1=5&var2=6`<br>Would be redirected automatically to the following URL:<br>`http://fabrikam/NewScripts/File.asp?var1=5&var2=6` |
| $V | Passes the requested URL, without the server name and without any parameters. To include parameters, use the $P or $Q variable with the $V variable.<br>If the EXACT_DESTINATION flag is not set, the resulting destination URL will have the name of the requested file appended as a folder name, as well as the file name itself. | If the virtual directory that is mapped to the following Web site:<br>`http://fabrikam/Scripts`<br>is configured to be redirected to the following exact URL:<br>`http://contoso$V`<br>A new request for the following URL:<br>`http://fabrikam/Scripts/File.asp`<br>Would be redirected automatically to the following URL:<br>`http://contoso/Scripts/File.asp` |

**Redirection Using Wildcards**

You can use redirect wildcards to match any number of characters in the original URL. Begin the destination URL with an asterisk (*) and a semicolon (;), and separate pairs of wildcard characters and destination URLs with a semicolon. The following example shows the format of a redirection string that uses wildcards.

```
*; Wildcard1; Destination1[; Wildcard2; Destination2][, Flags]
```

Where *WildcardN* can include an asterisk (*) to match any number of characters in the original URL and pass them to *DestinationN* by using the $0 through $9 variables defined in the following table.

| "Variable" | "Description" | "Example" |
|---|---|---|

| | | |
|---|---|---|
| $0 through $9 | Passes the portion of the requested URL that matches the indicted wildcard character. Wildcards must start at $0 for each *Wildcard*;*Destination* set. The EXACT_DESTINATION flag must be appended to the redirection string. | If the virtual directory that is mapped to the following Web site:<br>`http://fabrikam/Scripts`<br>is configured with the following redirection string:<br>`*; /Scripts/*_in.asp; /NewScripts/$0.dll; /Scripts/*_out.asp; http://contoso/NewScripts/$0.dll`<br>A new request for the following URLs:<br>`http://fabrikam/Scripts/data_in.asp`<br>`http://fabrikam/Scripts/data_out.asp`<br>Would be redirected automatically to the following respective URLs:<br>`http://fabrikam/NewScripts/data.dll`<br>`http://contoso/NewScripts/data.dll` |
| ! | Does not allow a redirect.<br>This variable is useful when you want to exempt files or subdirectories from redirection when a parent directory is configured to be redirected. | If the virtual directory that is mapped to the following Web site:<br>`http://fabrikam/Scripts`<br>is configured to be redirected, but the following Web site:<br>`http://fabrikam/Scripts/Commerce/Accounts`<br>needs to be exempt, programmatically set the redirection string for the exempted folder to the following<br>`*;!` |

For example, to redirect all requests for /Scripts/Filename.stm to a single file called Default.stm, and to redirect all requests for /Scripts /Filename.htm to a single file called Default.htm, use the following redirect request for the /Scripts virtual directory:

```
*;Filename.stm;/Default.stm/Filename.htm;/Default.htm
```

**Using Redirection Flags**

The following flags augment redirect strings of either URL or wildcard format. More than one flag is allowed. Flags are appended programmatically to the redirection string in HttpRedirect, or configured by selecting check boxes in IIS Manager.

| Flag | User Interface Check Box | Description |
|---|---|---|
| EXACT_DESTINATION | **The exact URL entered above** | Indicates that the value provided for the destination should be considered an absolute target location. |
| CHILD_ONLY | **A directory below this one** | Alerts IIS that redirection should occur only once because the destination is in a subdirectory of the original URL. This flag avoids loops. Also, this flag instructs IIS to redirect requests only if they are to subfolders or files of the original URL. |
| PERMANENT | **A permanent redirection for this resource** | Indicates that this redirection is permanent for this resource. |
| TEMPORARY | Clear the **A permanent redirection for this resource** check box" | Indicates that this redirection is temporary for this resource. |

**Related Information**

- For information about how to use IIS Manager to configure URL redirection, see Redirecting Requests to Files, Directories, or Programs.

- For information about how to use the URL Rewrite module to configure URL redirection in IIS 7, see Using the URL Rewrite Module.

# Tracing Events Reference

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 with SP1

This reference describes the events that can be recorded in the trace log when using request-based tracing or IIS Admin tracing. When applicable, an event description may include troubleshooting steps. The events listed here are organized by tracing provider:

- Active Server Pages (ASP)
- IIS Admin Service
- IIS:SSL Filter
- World Wide Web Service Global Tracing Events
- IIS: WWW ISAPI Extensions
- IIS: WWW Server
    - IIS: CGI
    - IIS: Cache
    - IIS: Compression
    - IIS: Filter
    - IIS: General
    - IIS: ISAPI
    - IIS: Security

If you are searching for information on a specific event, press CTRL+F and search for the desired event description.

**Active Server Pages (ASP)**

| Event | Description |
|---|---|
| ASP_DEQUEUE_REQUEST | The ASP request was removed from the ASP request queue for processing/execution. |
| ASP_DONE_PROCESSING | The ASP engine finished executing the request, but ASP may still have to send request data back to the client. ASP is done with the request when the ASP_END_REQUEST event is triggered. |
| ASP_END_REQUEST | ASP finished processing the request and has sent all associated response data. |
| ASP_START_REQUEST | ASP received a new request from IIS. |
| ASP_APPLICATION_ONSTART_ERROR | ASP encountered an error when trying to execute the **Application_OnStart** function in the Global.asa file. The ASP execution was aborted. Look at the error response sent back to the client to help determine the source of the error. |
| ASP_CLIENT_DISCONNECTED | The client disconnected while the request was on the ASP queue. |

| | |
|---|---|
| ASP_COMPILE_FAILED | ASP could not compile the ASP file successfully. Look at other events between ASP_START_COMPILE and ASP_END_COMPILE in the trace log to determine why the compile failed. Or look at the error response sent back to the client to help determine why the compile failed. |
| ASP_FILE_ACCESS_DENIED | ASP tried to read the file but failed because of insufficient permissions. The user may not have the appropriate permissions. Check the NTFS permissions on the file. Enable the IIS authentication tracing provider, request the file, and check the trace log for more information about the problem. |
| ASP_FILE_NOT_FOUND | The requested ASP file does not exist at the specified location. Verify the URL of the path. Also look at the data in the ASP_READ_FILE event to determine the file that ASP is trying to read. You may need to run the trace again with verbosity set to level four. |
| ASP_GLOBAL_ASA_ACCESS_DENIED | ASP tried to read the Global.asa file but failed because of insufficient permissions. The user does have the appropriate permissions. Check the NTFS permissions on the Global.asa file for the specified application. Check the ASP_START_GLOBAL_ASA_CHECK event which specifies the location of the Global.asa file it is trying to access. If you do not see the ASP_START_GLOBAL_ASA_CHECK event run the trace again with verbosity set to level four. |
| ASP_INIT_FAILURE | ASP failed to initialize. There are four possible reasons why ASP failed to initialize:<br><br>1. IIS could not read the MetaBase.xml configuration file. This is a fatal error meaning IIS will not be able to process requests. This error may be the result of invalid metabase configuration or an invalid access control list (ACL) on the MetaBase.xml file. This is a level two error.<br><br>To remedy this problem:<br><br>1. Verify that the MetaBase.xml file exists in the %windir%\system32\inetsrv directory.<br><br>2. Check the permissions on the MetaBase.xml file and ensure that System has Full Control.<br><br>3. If necessary, restart IIS and all associated services using the **IISReset** command. From a command prompt, type **iisreset** and press ENTER.<br><br>4. If ASP fails to initialize after restarting all IIS services, contact Microsoft Product Support Services.<br><br>1. ASP could not initialize the ASP disk cache. This is not a fatal error meaning IIS and ASP will continue to process requests, but you may see significant performance degradation because ASP will not be able to use the disk cache.<br><br>To remedy this problem:<br><br>1. Ensure that the ASP Templates Cache directory exists. It is located in the %windir%\system32 directory by default.<br><br>2. Check access control lists (ACLs) on the directory and ensure that the IIS_WPG group is listed. The IIS_WPG group must have Read and Write permissions on the ASP Templates Cache directory. |

| | |
|---|---|
| | 3. Ensure that the user identity that the process is running as is a member of the IIS_WPG group. In the Computer Management console, expand **Local Users and Groups**, right-click the **IIS_WPG group**, and select **Properties**.<br><br>1. ASP could not initialize the Component Object Model (COM). This is a fatal error meaning IIS will not be able to process requests.<br><br>To remedy this problem, look at the Windows NT Event Log for COM errors. Debug any COM errors.<br><br>1. ASP could not initialize the metabase change notification listener. This is a fatal error meaning IIS will not be able to process requests.<br><br>To remedy this problem:<br><br>1. Recycle the application pool. Right-click the application pool and click **Recycle**.<br><br>2. If the problem persists, restart IIS and all associated services using the **IISReset** command. From a command prompt, type **iisreset** and press ENTER.<br><br>3. If the problem persists, contact Microsoft Product Support Services. |
| ASP_QUEUE_REQUEST_FAILED | ASP failed to put a request in the request queue. The queue may be full or another request may be blocking further queue processing. If this failure was preceded by the ASP_SERVER_TOO_BUSY_TOO_QUEUE event, then the queue is full.<br><br>In this situation, ASP should report itself as not healthy and the Windows NT Application Event log should contain an event from W3svc_WP. If Worker Processes Health Monitoring is enabled, the ping to the worker process should fail and the application pool should recycle once IIS detects that the ping failed.<br><br>To remedy this problem<br><br>1. Ensure that the AspRequestQueueMax Metabase Property was not changed from its default value (500). If **AspRequestQueueMax** was not changed, use Currently Executing Requests Tracing to determine which requests were executing when the failure occurred.<br><br>2. Once you have found which requests are executing, trace the requests to the indicated URL to find where the blockage is occurring. |

| | |
|---|---|
| ASP_REQUEST_TIMEOUT | The ASP request spent too much time on the queue. The queue may be full or another request may be blocking further queue processing. Recycle the application pool or use Currently Executing Requests Tracing to understand which requests were executing when the failure occurred.<br><br>To remedy this problem<br><br>1. Ensure that the AspQueueTimeout Metabase Property was not changed from its default value. The default value is an unlimited number, represented by 4294967295.<br><br>2. If **AspQueueTimeout** was not changed, use Currently-Executing Requests Tracing to determine which requests were executing when the failure occurred.<br><br>3. Once you have determined which requests were executing, trace the requests to the indicated URL to find where the blockage is occurring. |
| ASP_SERVER_TOO_BUSY_TO_QUEUE | ASP cannot queue the request. The server may be running out of memory.<br><br>To remedy this problem<br><br>1. Recycle the application pool. Right-click the application pool and click **Recycle**.<br><br>2. If the problem persists, restart IIS and all associated services using the **IISReset** command. From a command prompt, type **iisreset** and press ENTER.<br><br>3. If the problem persists, contact Microsoft Product Support Services. |
| ASP_SESSION_ONSTART_ERROR | ASP encountered an error when trying to execute the **Session_OnStart** function in the Global.asa file. Look at the error response sent back to the client to help determine the source of the error. The ASP execution was aborted. |
| ASP_TRANSACTION_ABORTED | The ASP transaction aborted. Transactions in this case are Microsoft Transaction System (MTS) transactions. Work with the developer of the ASP file to debug the application. |

| | |
|---|---|
| ASP_UNHEALTHY_FAILURE | ASP is not healthy and could not continue to execute requests. The **Reason** field for this event shows either DEADLOCK_DETECTED or OUT_OF_MEMORY. If you are running IIS 6.0 in worker process isolation mode and you see this event, the application pool hosting the application will be recycled if Worker Process Health Monitoring is enabled. If this failure happens repeatedly, your application may have a memory leak. If you are running IIS 6.0 in IIS 5.0 isolation mode, you will see an ASP error in the Windows NT Event Log, but no action will be taken. With IIS 6.0 running in IIS 5.0 isolation mode, you must restart the World Wide Web Publishing Service and then look at the Web site logs for all executed requests to understand why ASP is not healthy. Trace a request using Currently Executing Requests Tracing to understand which requests where executing when the failure occurred.<br><br>**Note**<br>If IIS 6.0 is running in worker process isolation mode and a worker process does not shut down within the time limit configured by the ShutdownTimeLimit Metabase Property, IIS will forcefully shut down the process. If IIS forcefully shuts down the process, all requests executing in that process are logged in the HTTPERR log. |
| ASP_GLOBAL_ASA_DOES_NOT_EXIST | ASP could not find the Global.asa file. If your application uses a Global.asa file, determine why the file does not exist in the specified location. If your application does not use a Global.asa file, then you can ignore this event. |
| ASP_APPLICATION_ONSTART_SUCCESS | ASP successfully executed the **Application_Onstart** function. |
| ASP_CURRENT_SESSION | This event tells you the session ID for the current session. |
| ASP_END_APPLICATION_ONSTART | ASP finished executing the **Application_Onstart** function. |
| ASP_END_CACHE_ACCESS | ASP finished checking the cache. Check the **Reason** field in the trace log for one of the following:<br><br>NOT_SERVED: The file was not found or not served. The file may not have been served because of improper permissions on the file. The file may not exist or there could be a larger problem with the server, for example, a hard disk crash.<br><br>SERVED_CACHE_MISS: The file did not exist in the cache. ASP retrieved the file from disk and tried to load it into the cache.<br><br>SERVED_CACHE_HIT_CHANGENOTIF: The file was found in the cache and served from the cache. Change notification was used to maintain cache consistency.<br><br>SERVED_CACHE_HIT_LASTMOD: The file was found in the cache and served from the cache. ASP polled the last-modified timestamp on the file to maintain cache consistency.<br><br>SERVED_CACHE_HIT_LASTMOD_NOACCESS: The connection between the IIS server and a Universal Naming Connection (UNC) file share is down. ASP cannot check the timestamp on the requested file. However, the requested file exists in the cache so it was served from the cache. The content may be stale. |
| ASP_END_COMPILE | ASP finished generating the ASP template. |
| ASP_END_GLOBAL_ASA_CHECK | ASP finished checking for the Global.asa file. |

| | |
|---|---|
| ASP_END_SCRIPTLESS_SEND | The requested ASP page did not contain any script tags. ASP did not need to compile the page before sending it.<br>**Note**<br>ASP files that do not contain scripts should be renamed with the .htm or .html extension. Script-less .asp files are not processed as efficiently as static .htm or .html files. Responses for static files can be returned by the HTTP.sys kernel cache, which can significantly increase performance. |
| ASP_END_SCRIPT_EXECUTION | ASP finished executing the script code. ASP may not be finished processing the request. |
| ASP_END_SESSION_ONSTART | ASP finished executing the **Session_Onstart** function. |
| ASP_FIRST_REQUEST_FOR_APPLICATION | ASP received the first request for this application. ASP_APPLICATION_ONSTART and GLOBAL.ASA_CHECK events should follow. |
| ASP_FIRST_REQUEST_FOR_ASP | The ASP.dll received its first request. Part of ASP initialization takes place on the first request. |
| ASP_GLOBAL_ASA_ACCESS_SUCCESS | ASP successfully read the Global.asa file. |
| ASP_NEW_SESSION_CREATED | ASP accepted the first request for a new session and created a session ID for this session. |
| ASP_QUEUE_REQUEST | ASP put the request on the queue. This event should be followed by the ASP_DEQUEUE_REQUEST event when the request is pulled off the queue for execution. |
| ASP_READ_FILE | ASP is trying to read the file. The trace log will show one of these events for each included file. |
| ASP_SECURE_SESSION_ID_SET | ASP made a transition to a secure state. This event marks the session ID of the secure session. |
| ASP_SESSION_ONSTART_SUCCESS | ASP successfully executed the **Session_Onstart** function. |
| ASP_START_APPLICATION_ONSTART | ASP is about to call the **Application_Onstart** function. |
| ASP_START_CACHE_ACCESS | ASP is about to check for the file template in the ASP File Template Cache. |
| ASP_START_COMPILE | ASP is about to compile the specified file. |
| ASP_START_GLOBAL_ASA_CHECK | ASP is about to check for the existence of a Global.asa at the specified location. |
| ASP_START_SCRIPTLESS_SEND | The requested page does not contain any script tags. ASP will not need to compile the page before sending it. |
| ASP_START_SCRIPT_EXECUTION | ASP started executing the script code. |
| ASP_START_SESSION_ONSTART | ASP started executing the **Session_Onstart** function. |
| ASP_TRANSACTION_COMMITTED | The ASP transaction was committed. Transactions in this case are Microsoft Transaction System (MTS) transactions. |
| ASP_SCRIPT_TRACE_COM_CALL_END | The ASP script engine finished calling into a COM object, other than the response object, as indicated by your script. |
| ASP_SCRIPT_TRACE_COM_CALL_START | The ASP script engine is making a call into a COM object other than the response object, as indicated by your script. |
| ASP_SCRIPT_TRACE_CREATE_OBJECT_END | The ASP script engine finished creating a COM object as indicated by your script. |

| | |
|---|---|
| ASP_SCRIPT_TRACE_CREATE_OBJECT_START | The ASP script engine is creating a COM object as indicated by your script. |
| ASP_SCRIPT_TRACE_END | The script engine has finished executing. This event shows where the script engine finished executing. For example, the execution could end in an include file. |
| ASP_SCRIPT_TRACE_GET_OBJECT_END | The script engine finished calling the **GetObject** function. |
| ASP_SCRIPT_TRACE_GET_OBJECT_START | The script engine started calling the **GetObject** function. |
| ASP_SCRIPT_TRACE_START | The script engine started executing. This event shows where the execution is starting. For example, the execution could start in an include file. |

**IIS Admin Service**

| Event | Description |
|---|---|
| IISADMIN_END | The IIS Admin service finished shutting down. |
| START | The IIS Admin service has started shutting down. The service may be shutting down because of administrator action, for example, the administrator may have executed the IISreset.exe command or the net stop IIS Admin command. |
| ABO_DISCONNECT_END | The IIS Admin service finished disconnecting metabase COM client connections (as indicated by the process ID). |
| ABO_DISCONNECT_START | The IIS Admin service is about to disconnect metabase COM client connections (as indicated by the process ID) before shutting down. |
| COUNINITIALIZE_END | The IIS Admin service finished calling the **CoUninitialize** function to shutdown all COM activity in the process. All COM activity in the process was successfully shutdown. |
| COUNINITIALIZE_START | The IIS Admin service is about to call the **CoUninitialize** function to shutdown all COM activity in the process. |
| EWR_SHUTDOWN_END | The IIS Admin service successfully shutdown the Edit-While-Running mechanism, which monitors the MetaBase.xml file for changes. |
| EWR_SHUTDOWN_START | The IIS Admin service is about to shutdown the Edit-While-Running mechanism, which monitors the MetaBase.xml file for changes. |
| LISTENER_SHUTDOWN_END | The IIS Admin service successfully shutdown the listener that monitors change notification in the metabase. This means the IIS Admin service has stopped publishing IIS metabase change notifications.<br>**Note**<br>All processes listening for change notification, as indicated by the process ID (PID), will no longer receive IIS metabase change notification. |
| LISTENER_SHUTDOWN_START | The IIS Admin service is about to shutdown the listener that monitors change notification in the metabase. This means the IIS Admin service will no longer publish IIS metabase change notifications.<br>**Note**<br>All processes listening for change notification, as indicated by the PID, no longer receive IIS metabase change notification. |
| METADATA_SHUTDOWN_END | The IIS Admin service successfully shutdown the in-memory metabase. |
| METADATA_SHUTDOWN_START | The IIS Admin service is about to shutdown the in-memory metabase. |

| | |
|---|---|
| SVCEXT_SHUTDOWN_END | The IIS Admin service successfully shut down the service extensions. |
| SVCEXT_SHUTDOWN_START | The IIS Admin service is about to shutdown the service extensions. |
| END | The IIS Admin service successfully started. |
| IISADMIN_START | The IIS Admin service is starting up. This should be the first published IIS Admin service event. If you do not see this event in the trace log, check the NT System event log for service control manager events with the IIS Admin service. |
| ABO_REGISTER_FAILURE | The IIS Admin service could not register the metabase ABO interface in COM, an indication that there is a problem in COM. Look at the error code published with this event and take the action described by the error code. If the problem persists, call Microsoft Product Support Services. |
| COM_SECURITY_FAILURE | The IIS Admin service could not set the permissions on the ABO interface in COM. Look at the error code published with this event and take the action described by the error code. If the problem persists, call Microsoft Product Support Services. |
| DCPROMO_FAILURE | The IIS Admin service could not register the IIS accounts in Active Directory after the machine was promoted to a domain controller. The problem probably resides in Active Directory. Once Active Directory is back online, stop and restart the IIS Admin service. By its very nature, this failure can only happen one time. Look at the error code published with this event and take the action described by the error code. If the problem persists, call Microsoft Product Support Services.<br><br>**Note**<br>When you see this event in the trace log, the IIS Admin service is still running. However, all IIS accounts may have shut down, including the default anonymous user account. |
| MBSCHEMA_PARSE_FAILURE | The IIS Admin service failed to parse the MBSchema.xml file. Look at the error code published with this event and take the action described by the error code. You can also look in the NT System event log for errors in the IIS Admin service. These errors may provide more information about this failure. If the errors do not provide you with the necessary information, role back to an earlier version of the MBSchema.xml file if you have one. You can also open the MBSchema.xml file in an XML editor to ensure that the file is properly formatted. If the problem persists, call Microsoft Product Support Services. |
| METABASE_PARSE_FAILURE_MB_XML | The IIS Admin service failed to parse the MetaBase.xml file. Look at the error code published with this event and take the action described by the error code. You can also check the NT System event log for errors in the IIS Admin service. If the errors do not provide you with the necessary information, role back to an earlier version of the MetaBase.xml file. If the problem persists, call Microsoft Product Support Services. |
| METABASE_PARSE_FAILURE_SESSION_KEY | The IIS Admin Service failed to decrypt the encrypted session key in the metabase, which means IIS cannot read the metabase. The machine key IIS has stored does not match the machine key used to encrypt the session key. This is a critical failure. As a result, IIS will not run. Stop and restart the IIS Admin service. If the problem persists, rollback to an earlier version of the MetaBase.xml file and MBSchema.xml file from the History folder and restart the IIS Admin service. If the problem persists, uninstall IIS, reinstall IIS, and restore a known-good backup from an IIS 6.0 server that was backed up with a password (the password is used in the encryption process). If the problem persists, contact Microsoft Product Support Services. |
| SVCEXT_STARTUP_FAILURE | The IIS Admin service attempted to start a service extension and the startup failed. Restart the IIS Admin service. If the problem persists, contact Microsoft Product Support Services. |

| | |
|---|---|
| COM_CATALOG_FAILURE | The IIS Admin service attempted to start the COM catalog and failed. There is probably an error in the COM catalog. |
| VSS_FAILURE | The IIS Admin service attempted to register with the Volume Snapshot Service (VSS), but failed. The IIS Admin service will continue to run, but if any system-wide backups are performed using the Windows Server backup utility, then IIS may not be able to get the MetaBase.xml file ready for backup. |
| DCPROMO_END | The IIS Admin service finished updating the user accounts in the case where the IIS server was promoted to a domain controller. |
| DCPROMO_START | The IIS Admin service is about to check if the machine has been promoted to a domain controller. If the machine has been updated to a domain controller, the IIS Admin service will begin updating the user accounts. |
| METABASE_PARSE_FAILURE_SETTING | The IIS Admin service failed to read one or more of the property settings in the metabase. The IIS Admin service will still run, but will ignore the settings that failed. Check the Windows NT System Event log to determine which properties were ignored. |
| ABO_REGISTER_END | The IIS Admin service finished registering the ABO interface with COM. |
| ABO_REGISTER_START | The IIS Admin service is about to register the ABO interface with COM. |
| COM_CATALOG_START | The IIS Admin service is about to start using the COM catalog. |
| COM_CATALOG_END | The IIS Admin service successfully started and is using the COM catalog. |
| COM_SECURITY_END | The IIS Admin service successfully setup the permissions in COM on the ABO interface. |
| COM_SECURITY_START | The IIS Admin service is about to setup the permissions in COM on the ABO interface. |
| MBSCHEMA_PARSE_END | The IIS Admin service successfully parsed the MBSchema.xml file. |
| MBSCHEMA_PARSE_START | The IIS admin service is about to parse the MBSchema.xml file. |
| METABASE_PARSE_END | The IIS Admin service successfully parsed the MetaBase.xml file. |
| METABASE_PARSE_START | The IIS Admin service is about to parse the MetaBase.xml file. |
| SVCEXT_STARTUP_END | The IIS Admin service successfully started a service extension. |
| SVCEXT_STARTUP_START | The IIS Admin service is about to start a service extension. |
| VSS_END | The IIS Admin service successfully registered with the Volume Snapshot Service (VSS). |
| VSS_START | The IIS Admin service is about to register with the Volume Snapshot Service (VSS). |

**IIS: SSL Filter**

| Event | Description |
|---|---|
| SSLERROR_DECRYPT | The HTTP Filter service failed while decrypting Secure Sockets Layer (SSL) data. This is a fatal error. The request was rejected as a result of this error. This error may be the result of a change in the encrypted data on the client side or on the transport layer. Because this error is the result of a change on the client side, this error does not require any action on the part of the administrator |

| | |
|---|---|
| SSLERROR_ENCRYPT | The HTTP Filter service failed while encrypting SSL data. This is a fatal error. This error could be the result of a low-memory condition on the server. The Windows NT Event logs should contain more information about this error. If the problem persists, contact Microsoft Product Support Services. |
| SSLERROR_HANDSHAKE | The HTTP Filter service failed during the SSL handshake. This is a fatal error. The request was rejected. This error could be the result of a low-memory condition on the server. Also, the server may not support the SSL protocol and/or cipher requested by the client. This error most likely resides with the client computer. The client computer may be using the SSL 2.0 protocol and the server may be configured to use client certificates, as indicated by following error in the trace log: SEC_E_ALGORITHM_MISMATCH. This error does not require any action on the part of the administrator. |
| SSLERROR_INTERNALINCONSISTENCY | The HTTP Filter service failed because the SSL filter detected an internal error. This is a fatal error. The request was rejected. This error can happen during the SSL handshake. This error does not require any action on the part of the administrator. |
| SSLERROR_INVALIDCONTEXT | The HTTP Filter service failed because the SSL filter detected an internal error. This is a fatal error. The request was rejected. This error can happen during the SSL handshake. This error does not require any action on the part of the administrator. |
| SSLERROR_RENEGOTIATION | The HTTP Filter service failed while the client certificate was being renegotiated. This is a fatal error. The request was rejected. The source of this problem may be in the secure channel (Schannel) call or an internal error. This error can happen during the SSL handshake. This error does not require any action on the part of the administrator. |
| SSLERROR_USERDATADURINGRENEGOTIATION | The HTTP Filter service failed during the SSL renegotiation handshake. This is a fatal error. The request was rejected. This is an internal error. This error does not require any action on the part of the administrator. |
| ENDCONNECTION | The HTTP Filter service closed the connection and all data has been sent. |
| SSLCLIENTCERTINFO | This event tells you the details of the client certificate that was negotiated during the SSL handshake. |
| SSLCREDENTIALINFO | This event tells you the credentials that will be used during the SSL handshake. |
| SSLDECRYPTEDDATA | The HTTP Filter service successfully decrypted the client data and wrote the data as clear text into the ETW log. The HTTP Filter service may limit the amount of decrypted data going into the ETW log to 32K. |
| SSLENDHANDSHAKE | The HTTP Filter service successfully completed the SSL handshake. |
| SSLSCHANNELCALL | The HTTP Filter service is monitoring calls into the Schannel, which implements the SSL protocol handling. |
| SSLSTARTHANDSHAKE | The HTTP Filter service started an SSL handshake on a secure connection. This event can be triggered in two situations:<br>• This is a new connection to the secure port.<br>• This is a renegotiation when client-certificate authentication is requested, but the initial SSL handshake did not include the request for client authentication. |

| | |
|---|---|
| SSLTOBEENCRYPTEDDATA | The HTTP Filter service received the clear-text data and is preparing to encrypt that data. |
| STARTCLOSECONNECTION | The HTTP Filter service has started to shut down the client connection. The connection may still contain data. |
| STARTCONNECTION | The HTTP Filter service was notified by HTTP.sys that a client has connected to a secure port. |

**World Wide Web Service Global Tracing Events**

These events describe WWW service actions but exclude request-processing events.

| Event | Description |
|---|---|
| FILTER_UNLOAD_END | A worker process finished unloading an Internet Server API (ISAPI) filter. A worker process will unload an ISAPI filter during normal shutdown. The name of the unloaded filter is identified by the **Filter Name** property in the trace log. |
| FILTER_UNLOAD_START | A worker process started unloading an ISAPI filter. A worker process will unload an ISAPI filter during normal shutdown. The name of the unloaded filter is identified by the **Filter Name** property in the trace log. |
| GLOBAL_PROCESS_SHUTDOWN_MODE | A worker process was told by the WWW Service to stop listening for new requests and to shutdown. This can happen for a variety of reasons, including a process recycle. This is not a crash. If a worker process was instructed to shutdown immediately, then the worker process also shutdown all Common Gateway Interface (CGI) scripts. This event is followed by the GLOBAL_W3_SERVER_SHUTDOWN_START event. |
| GLOBAL_W3_SERVER_SHUTDOWN_END | A worker process finished shutting down and finished cleaning up after its objects. |
| GLOBAL_W3_SERVER_SHUTDOWN_START | A worker process is preparing to shutdown and preparing to clean up after its objects. |
| GLOBAL_W3_SERVER_INIT_ERROR | This is an internal error. A worker process failed one of its initialization routines. The INITSTATUS event indicates how far the worker process got in its initialization routine. The ERROR CODE event indicates the Win32 error received when trying to execute the initialization routine. Recycle the application pool. If the problem persists, contact Microsoft Product Support Services. Tell the support engineer the value of the INITSTATUS and ERROR CODE properties. |
| FILTER_LOAD_END | A worker process finished loading an ISAPI filter. If the ERROR CODE event for this event is zero, the filter loaded successfully. If the ERROR CODE event is a numeric value other than zero, the filter did not load. The Windows NT Event log will contain more information explaining why the filter did not load. If the problem persists, contact Microsoft Product Support Services. |
| FILTER_LOAD_START | A worker process is attempting to load an ISAPI filter. The **Filter Name** property indicates the name of the filter the worker process is attempting to load. |
| GLOBAL_PROCESS_START_INFO | A worker process is starting up. This event records the parameters passed to the worker process during startup, including command-line arguments, which mode the worker process is running in, the identity under which the process is running, and the processor affinity mask for that process. |
| GLOBAL_ULATQ_START_LISTENING | A worker process started listening for requests from HTTP.sys. |
| GLOBAL_W3_SERVER_END | A worker process finished initializing. The worker process is ready to process requests. |
| GLOBAL_W3_SERVER_START | A worker process has started up and is beginning initialization. |

| | |
|---|---|
| GLOBAL_CACHE_START | The worker process caches are about to be initialized. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_CERTIFICATE_START | A worker process is about to initialize the certificate object wrapper. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_CGI_RESTRICTION_LIST_START | A worker process is about to initialize the CGI restriction list. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_COMPRESSION_START | A worker process is about to initialize HTTP compression. Failure to initialize HTTP compression does not result in a failure to initialize the worker process. The worker process will continue to run, but will not be able to compress files. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_FILTER_START | A worker process is about to initialize global ISAPI filters. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_IISUTIL_START | A worker process is about to initialize the IISUTIL (a set of common utility routines). If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_ISAPI_RESTRICTION_LIST_START | A worker process is about to initialize the ISAPI restriction list. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_LOGGING_START | A worker process is about to initialize logging. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_METABASE_LISTENER_START | A worker process is about to initialize the metabase listener, which monitors the metabase for changes. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_METABASE_START | A worker process is about to initialize the metabase. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_MIME_MAP_START | A worker process is about to initialize the Multipurpose Internet Mail Extensions (MIME) map. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_RAW_CONNECTION_START | A worker process running in IIS 5.0 isolation mode is about to initialize stream filter support. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_RESOURCE_DLL_START | A worker process is about to load its resource DLL. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_SERVER_VARIABLE_START | A worker process is about to initialize the global server variables. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_SITE_LIST_START | A worker process is about to initialize the site list. This means the objects that will store the list of sites have been initialized by the worker process. This does not mean the list of sites have been completely loaded into the worker process. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_ULATQ_START | A worker process is about to initialize communication to HTTP.sys and the World Wide Web Publishing Service (WWW service). If the worker process initializes communication, this event will be followed with a GLOBAL_FILTER_START event. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_W3_CONNECTION_START | A worker process is about to initialize the connection table. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |

| | |
|---|---|
| GLOBAL_W3_CONTEXT_START | A worker process is about to initialize its World Wide Web context. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_W3_REQUEST_START | A worker process is about to initialize the request object. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_W3_RESPONSE_START | A worker process is about to initialize the response object. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_W3_SITE_START | A worker process is about to initialize the World Wide Web site object. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |
| GLOBAL_W3_WINSOCK_START | A worker process is about to initialize Winsock for network address translation. In IIS 6.0, Winsock is not used as the transport layer. The kernel-mode driver, HTTP.sys, is the request/response transport layer. If the action fails, the GLOBAL_W3_SERVER_INIT_ERROR will be written to the trace log. |

**IIS: WWW ISAPI Extensions**

| Event | Description |
|---|---|
| CALL_ISAPI_EXTENSION | A worker process is about to pass a request to the ISAPI extension. |
| ISAPI_EXTENSION_DONE | The client computer received the final response from the ISAPI extension or disconnected. The ISAPI context for this request has been reclaimed. When tracing for capacity planning purposes, calculate the delta between the NOTIFY_ISAPI_COMPLETION event and this event to determine how long it took to send the last bit of the response. |
| NOTIFY_ISAPI_COMPLETION | The ISAPI extension sent the last segment of the response. The ISAPI extension is done with the request. |

**IIS: WWW Server**

| Event | Description |
|---|---|

| | |
|---|---|
| AUTH_ANON_PASSWD_CHANGE_NEEDED | The virtual directory is configured for anonymous authentication. The password for the anonymous user ID has expired or needs to be changed per policy. To resolve this issue, reset the anonymous user password in two places:<br>• The IIS metabase<br>• The Local Users and Groups node of the Computer Management MMC snap-in<br><br>To reset the anonymous user password in the IIS metabase:<br>1. Enter the following command at a command prompt: **cscript.exe adsutil.vbs set w3svc/site number/AnonymousUserPass** *new password*<br>2. Press ENTER.<br><br>or<br>1. Open IIS Manager, expand the Local Computer, right-click the Web site or virtual directory where you want to change the Anonymous Authentication password and click **Properties**.<br>2. Click the **Directory Security** tab.<br>3. In the **Authentication and access control** section, click **Edit**.<br>4. Change the password in the **Password** text box.<br>5. Click **OK**.<br><br>To reset the anonymous user password in the Local Users and Groups node of the Computer Management MMC snap-in<br>1. Right-click My Computer and click **Manage**.<br>2. Expand the Local Users and Groups node and click **Users**.<br>3. Right-click the anonymous user account (IUSR_machine name) and click **Set password**.<br>4. Click **OK**. |
| AUTH_INVALID_ANON_ACCOUNT | The virtual directory is configured for anonymous authentication. The anonymous user name for this virtual directory does not exist or is an invalid account. To resolve this issue, ensure that the user name stored in the AnonymousUserName Metabase Property exists on the computer or is available to the computer in the case of a domain account. |
| AUTH_BAD_BASIC_HEADER | The server computer requires Basic authentication to access the requested URL. The client computer supplied an invalid basic header or failed to supply a basic header. The request failed. This server sent an HTTP 401: Unauthorized response. |
| AUTH_BASIC_LOGON_FAILED | The client attempted to authenticate using basic authentication. The client supplied an invalid user name and/or password. The request failed. This server sent an HTTP 401: Unauthorized response. |

| | |
|---|---|
| AUTH_IISDIGEST_LOGON_FAILED | The client attempted to authenticate using digest authentication. The client supplied an invalid user name and/or password. The request failed. This server sent an HTTP 401: Unauthorized response. |
| AUTH_NTLM_NULL_SESSION | The virtual directory is configured for NTLM authentication. The client computer did not provide a user name and/or password during authentication. The request failed. This server sent an HTTP 401: Unauthorized response. |
| AUTH_PASSPORT_LOGON_FAILED | The virtual directory is configured for Passport authentication. The client computer provided an invalid user name and/or password during authentication. The request failed. This server sent an HTTP 401: Unauthorized response. |
| AUTH_PASSWD_CHANGE_DISABLED | The authenticated user's password is about to expire and the server is not configured for password change according to the **PasswordChangeFlags** metabase property. Configure the IIS 6.0 server for password change. For more information, see the PasswordChangeFlags Metabase Property. |
| AUTH_PASSWD_CHANGE_NEEDED | The user's password needs to be changed. IIS is about to change the user's password as configured by the **PasswordChangeFlags** metabase property. For more information, see the PasswordChangeFlags Metabase Property. |
| AUTH_SSPI_LOGON_FAILED | IIS attempted to logon to the user with the credentials supplied in the client request. The logon failed due to invalid credentials. Logon was denied. |
| AUTH_TYPE_NOT SUPPORTED | The client requested an authentication type that is not configured on the IIS server in the metabase. The server is not configured for the authentication type requested by the client. |
| AUTH_WDIGEST_LOGON_FAILED | IIS attempted to logon the user using Windows Digest Authentication (Advanced Digest Authentication). The authorization attempt failed. |
| AUTH_END | IIS finished the logon attempt. Authentication for this request is finished. |
| AUTH_REQUEST_AUTH_TYPE | IIS determined the authentication type requested by the client computer. This authentication type is included in the trace log with this event. |
| AUTH_SSPI_CONTINUE_NEEDED | The server needs to continue the authentication handshake between the client and the server. The server sent an HTTP 401: Bad Logon response. The response included the WWW authenticate header. |
| AUTH_START | IIS started authenticating the request. |
| AUTH_SUCCEEDED | The request was successfully authenticated. This event includes the following information:<br>• The requested authentication type<br>• The remote user's name<br>• The local user name<br>• The impersonation level |

**IIS: CGI**

| Event | Description |
|---|---|
| CGI_END | IIS finished processing the CGI. IIS may still be processing the request. |
| CGI_START | IIS started executing a CGI. |
| CGI_HEADERS_RECIEVED | IIS received the complete set of HTTP response headers generated by the CGI. |

| | |
|---|---|
| CGI_LAUNCH | The CGI executable process has been created and launched. |
| CGI_PREMATURE_TERMINATION | A CGI terminated unexpectedly. The server sent an HTTP 502: Premature Exit error code to the client. Investigate why the CGI failed or debug the CGI executable. |
| CGI_REQUEST_ENTITY_SENT | The CGI sent the complete request-entity body. |
| CGI_TIMEOUT | The CGI timed out. The server sent an HTTP 502: Timeout error to the client. Investigate why the CGI timed out or debug the CGI executable. |

**IIS: CACHE**

| Event | Description |
|---|---|
| FILE_CACHE_ACCESS_END | IIS finished checking the file cache for the requested file. The trace log for this event includes the following information:<br>• Whether the file was found in the cache<br>• Whether the file was added to the cache<br>• Whether the directory where the file resides is being monitored for changes<br>• Whether IIS encountered an error when checking the last-modified date stamp on the file<br>• Whether IIS ignored the error for the last-modified date stamp on the file |
| FILE_CACHE_ACCESS_START | IIS started to check the file cache for the requested file. |
| HTTPSYS_CACHEABLE | IIS checked to see if a response could be cached in the HTTP.sys response cache. For information about the HTTP.sys response cache and why a response was not cached, see Events and Conditions That Disable HTTP.sys Response Caching. |
| URL_CACHE_ACCESS_END | IIS finished checking the URL metadata configuration cache for the requested URL. The trace log details for this event indicate whether or not the URL was added to the cache. |
| URL_CACHE_ACCESS_START | IIS started checking the metadata cache for the URL metadata. |

**IIS: Compression**

| Event | Description |
|---|---|
| DYNAMIC_COMPRESSION_NOT_SUCCESS | IIS could not compress the dynamic response. Ensure that dynamic compression is enabled. |
| DYNAMIC_COMPRESSION_START | IIS attempted to compress a dynamic response (any request that is not static). |
| DYNAMIC_COMPRESSION_SUCCESS | IIS verified various criteria for dynamic compression, including whether the client can accept compressed responses and whether the server is configured for dynamic compression. IIS is ready to start compressing the entity body for the dynamic response. |
| STATIC_COMPRESSION_CREATE_END | IIS attempted to compress a static file and store the compressed response in the IIS Temporary Compressed Files directory (by default). The trace log for this event indicates the error code, if any. If the trace log for this event shows a zero, the response was successfully compressed. The trace log also indicates the original file name and size and the compressed file name and size. Use this information to determine how much the static file was compressed. |
| STATIC_COMPRESSION_CREATE_START | IIS started compressing a static file. IIS will attempt to store the compressed response in the IIS Temporary Compressed Files directory (by default). |

| | |
|---|---|
| STATIC_COMPRESSSION_NOT_SUCCESS | IIS attempted to compress a static file but failed. The trace log for this event contains details about the error. IIS will attempt to compress the file again on the next request. |
| STATIC_COMPRESSION_START | IIS started compressing the static file. |
| STATIC_COMPRESSION_SUCCESS | IIS found the requested file in the IIS Temporary Compressed Files directory (the default directory) and is preparing to send it. |
| DYNAMIC_COMPRESSION_DO | IIS dynamically compressed a portion of the response. |

**IIS: Filter**

| Event | Description |
|---|---|
| FILTER_END | An ISAPI filter finished processing a request. |
| FILTER_START | An ISAPI filter started processing a request. The filter name is indicated in the event. |
| FILTER_ERROR | An ISAPI filter returned an error after processing a request. The error number is indicated in the event. |
| FILTER_ACCESS_DENIED_END | An ISAPI filter that requested SF_NOTIFY_ACCESS_DENIED notification finished processing a request. |
| FILTER_ACCESS_DENIED_START | An ISAPI filter that registered for the SF_NOTIFY_ACCESS_DENIED notification denied access to a request. The trace log for this event contains the name of the requested URL, the physical path, and the reason why the request was denied access. |
| FILTER_ADD_REQ_HEADER | An ISAPI filter that has registered for the PREPROC_HEADER notification added additional headers to the request. This event is triggered each time the filter adds a header to the request. The trace log for this event contains the header name and value. |
| FILTER_ADD_RESP_HEADER | An ISAPI filter that registered for the SEND_RESPONSE notification added additional headers to the response. This event is triggered each time the filter adds a header to the response. The trace log for this event contains the header name and value. |
| FILTER_AUTHENTICATION_END | The SF_NOTIFY_AUTHENTICATION notification completed. The trace log for this event contains the final user name and whether or not the user name changed. |
| FILTER_AUTHENTICATION_START | IIS started the SF_NOTIFY_AUTHENTICATION notification. IIS logs this event only if a filter has registered for this notification The trace log for this event contains the original user name. |
| FILTER_AUTH_COMPLETE_END | The ISAPI filter finished reporting the SF_NOTIFY_AUTHENTICATION event. The trace log for this event includes the final user name (in case the filter changed the user name) and whether or not the password changed. |
| FILTER_AUTH_COMPLETE_START | The ISAPI filter registered for the SF_NOTIFY_AUTHENTICATION filter event notification and is about to be notified. |
| FILTER_END_OF_REQUEST_END | The ISAPI filter finished handling the SF_NOTIFY_END_OF_REQUEST filter event notification. |
| FILTER_END_OF_REQUEST_START | The ISAPI filter is about to handle the SF_NOTIFY_END_OF_REQUEST filter event. |
| FILTER_LOG_END | The ISAPI filter finished handling the SF_NOTIFY_LOG filter event notification. The trace log for this event contains the final log data parameters. |
| FILTER_LOG_START | The ISAPI filter is about to handle the SF_NOTIFY_LOG filter event notification. The event contains the original data to be logged. |

| | |
|---|---|
| FILTER_PREPROC_HEADERS_END | The ISAPI filter finished handling the SF_NOTIFY_PREPROC_HEADERS filter event notification. |
| FILTER_PREPROC_HEADERS_START | The ISAPI filter is about to handle the SF_NOTIFY_PREPROC_HEADERS filter event notification. |
| FILTER_SEND_RAW_DATA_END | The ISAPI filter handled the SF_NOTIFY_SEND_RAW_DATA filter event notification. |
| FILTER_SEND_RAW_DATA_START | The ISAPI filter is about to handle the SF_Nofity_Send_Raw_Data filter event notification. |
| FILTER_SEND_RESPONSE_END | The ISAPI filter handled the SF_NOTIFY_SEND_RESPONSE_START filter event notification. |
| FILTER_SEND_RESPONSE_START | The ISAPI filter is about to handle the SF_NOTIFY_SEND_RESPONSE_START filter event notification. |
| FILTER_SET_REQ_HEADER | An ISAPI filter successfully set an HTTP request header. This event is logged each time a request header is set. The trace log for this event contains the header name and header value. |
| FILTER_SET_RESP_HEADER | An ISAPI filter successfully set an HTTP response header. This event is logged each time a response header that is set. The trace log for this event contains the header name and header value. |
| FILTER_URL_MAP_END | The ISAPI filter handled the SF_NOTIFY_URL_MAP_START filter event notification. The trace log contains the following information for this event:<br>• The original URL<br>• The file path<br>• Access permission for the URL<br>• The number of matching characters in the physical path<br>• The number of matching characters in the URL<br>• The script map entry for this URL |
| FILTER_URL_MAP_START | The ISAPI filter is about to handle the SF_NOTIFY_URL_MAP_START filter event notification. The trace log contains the following information for this event:<br>• The original URL<br>• The file path<br>• Access permission for the URL<br>• The number of matching characters in the physical path<br>• The number of matching characters in the URL<br>• The script map entry for this URL |

**IIS: General**

| **Event** | **Description** |
|---|---|
| GENERAL_CGI_HANDLER | IIS determined that the request was for a CGI and started processing that CGI. |

| | |
|---|---|
| GENERAL_CHILD_REQUEST_END | The child execution completed and the child context was deleted. The trace log for this event contains the response status code, sub-status codes, and the number of bytes sent. |
| GENERAL_CHILD_REQUEST_START | IIS started to execute the child request.<br><br>**Note**<br>An ISAPI extension can call HSE_REQ_EXECUTE_URL to re-execute a request.<br><br>The trace log for this event contains the following:<br>• The site id<br>• The new URL to execute<br>• The new verb to execute<br>• The number of times the verb/URL has been re-executed |
| GENERAL_DAV_HANDLER | IIS determined that the incoming request was a Web Distributed Authoring and Versioning (WebDAV) request. |
| GENERAL_ISAPI_HANDLER | IIS set up the initial context for the ISAPI, but has yet to determine if the ISAPI extension is enabled. |
| GENERAL_MAP_HANDLER | IIS started processing a map file request. |
| GENERAL_OOP_ISAPI_HANDLER | IIS determined that the request was for an ISAPI extension that is Out of Process. This is only applicable in IIS 5.0 Isolation mode. |
| GENERAL_OPTIONS_HANDLER | IIS started processing the Options request. A client request can use an **Options** method to determine the options available once it connects to the server. |
| GENERAL_REDIRECTION_HANDLER | IIS started sending the configured redirect-response to the client. |
| GENERAL_REQUEST_END | IIS completely finished processing the request and the response. The response was sent to the client. |
| GENERAL_REQUEST_START | An IIS worker process picked up a request from HTTP.sys and successfully associated that request with a site on the server. This event is logged before any filter notifications are triggered. This event marks the beginning of the request in the worker process. |
| GENERAL_SEND_CUSTOM_ERROR | The response status and sub-status codes have been generated. The HTTP status code is greater than or equal to 400 (that is, the requested generated an error). The custom error will be sent to the client if one of the following is true:<br>• The request and all of its children allow an error body to be sent back to the client<br>• The response code is greater than 400.<br><br>**Note**<br>This event is logged before IIS checks to see if the customer (if configured as a file) exists. |
| GENERAL_STATIC_FILE_HANDLER | IIS started processing a request for a static file. |
| GENERAL_TRACE_HANDLER | IIS started processing a request whose method is **Trace**. |

| | |
|---|---|
| GENERAL_GET_URL_METADATA | IIS found the URL metadata in the cache, or IIS created a new URL metadata cache object for the request. |

**IIS: ISAPI**

| Event | Description |
|---|---|
| ISAPI_END | IIS started its cleanup routine after a request was processed by an ISAPI. |
| ISAPI_START | IIS is about to call into the ISAPI extension. |

**IIS: Security**

| Event | Description |
|---|---|
| SECURITY_DENIED_BY_CGI_RESTRICTION | A request was either script-mapped to a CGI or called a CGI directly and that CGI was not enabled in the Web Service Extensions. |
| SECURITY_DENIED_BY_ISAPI_RESTRICTION | A request was either script-mapped to an ISAPI or called an ISAPI directly and that ISAPI was not enabled in the Web Service Extensions. |
| SECURITY_DENIED_BY_MIMEMAP | IIS could not locate a MIME map for the requested static file. IIS sent an HTTP 404.3 status code to the client. |
| SECURITY_FILE_ACCESS_DENIED | A request attempted to load the default file for a virtual directory and the file was either not found (the path was not found) or the requested resource had an invalid name. The trace log for this event contains the requested file system path and the user name (including domain) used to access the requested file. Check the file system ACL on the requested resource and verify that the user name has permission to access the file. |
| SECURITY_ILLEGAL_SHORT_FILENAME | A requested resource contains a short file name which expands to something other than the actual resource indicated. For example: */~test expands to /testsite/user/filename* |
| SECURITY_REJECTED_HOSTNAME | The client computer's host name is not allowed. IIS rejected the request. For more information, see Securing Sites with IP Address Restrictions. |
| SECURITY_REJECTED_IP | The originating client IP address is listed in the Web server's IP address restriction list. IIS rejected the request. For more information, see Securing Sites with IP Address Restrictions. |
| SECURITY _REJECTED_REQUIRED_SSL_128 | The requested directory requires 128-bit SSL encryption. IIS rejected the request. For more information, see IIS 6.0 Encryption. |

# Windows Server 2003 Product Help

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Welcome**

- Roadmap to Help and Information
- Choosing the features that are most useful for your server
- Common Administrative Tasks
- Troubleshooting Strategies
- New ways to do familiar Windows tasks
- Checklists Provided in Windows Server 2003 Help
- Best practices: pointers to recommendations and tips
- Submit product feedback to Microsoft

© 1985-2005 Microsoft Corporation. All rights reserved.

# Roadmap to Help and Information

Updated: July 19, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Roadmap to help and information**
Select an option to view:

- Viewing Help for the Windows Server 2003 family from another computer
- Viewing Help using Remote Desktop
- Viewing classic Help
- Viewing Help for Windows 2000
- Using the Windows Deployment and Resource Kits
- Using TechNet
- Using MSDN
- Accessing documentation on the Web
- Accessing Microsoft white papers
- Accessing information about training and certification
- Accessing information about compatibility
- Accessing information about accessibility
- Copyright

# Viewing Help for the Windows Server 2003 family from another computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing Help for the Windows Server 2003 family from another computer**

You can view Help for the Windows Server 2003 family from another computer by using Help and Support Center. For more information about how to do this, see To install Help content from another Windows computer, CD, or disk image.

You can also use Remote Desktop for Administration. For more information about how to do this, see Connect to another computer using Remote Desktop Web Connection. For more information about Remote Desktop for Administration, see Remote Desktop for Administration.

# Viewing Help using Remote Desktop

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing Help using Remote Desktop**

You can view Help for Windows XP and the Windows Server 2003 family from any computer running Windows XP by using Remote Desktop for Administration. For more information about how to do this, see Connect to another computer using Remote Desktop Web Connection. For more information about Remote Desktop for Administration, see Remote Desktop for Administration.

# Viewing classic Help

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing classic Help**
Classic Help is Help in the format that shipped with Windows 2000. This is the same content that appears in Help and Support Center, but in a different format.

To access classic Help, open Windows Help.

### Notes

- To open the Windows Help file, click **Start**, click **Run**, and then type **hh windows.chm**.
- If you install Help content from another Windows computer, CD, or disk image, the Help file contains the Help content for the local computer, not the Help content from the other source.

# Viewing Help for Windows 2000

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing Help for Windows 2000**

You can access Help for all Windows 2000 Server operating systems at the Microsoft Web site:

- To access Help for Windows 2000 Server, see Windows 2000 Server Documentation at the Microsoft Web site.
- To access Help for Windows 2000 Advanced Server, see Windows 2000 Advanced Server Documentation at the Microsoft Web site.
- To access Help for Windows 2000 Datacenter Server, see Windows 2000 Datacenter Server Documentation at the Microsoft Web site.

# Viewing Help on the Web

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can find product Help for the Microsoft® Windows Server™ 2003 operating systems on the Microsoft Web site. However, the site does not offer a separate set of topics for each version and each edition in that set of products. To reduce duplication and improve search efficiency, each topic in the product Help contains information about how a feature works across all versions and editions. In addition, each topic specifies any differences between versions and editions. If a topic does not specify any differences, the functionality that it describes is the same across all Windows Server 2003 operating systems.

In addition to differences between versions and editions, other differences can affect whether your server appears and functions exactly as the topics describe. For example, the steps to open an item in Control Panel vary slightly based on how a server is configured. The following list explains in more detail why Help topics on the Web might differ from what you see on a particular server:

- **Versions of Windows Server 2003**

  Different versions of Windows Server 2003 are available for installation on computers with different microprocessors, and not all features are available on all versions. For example, the x64-based versions of Windows Server 2003 have different sets of features than do the 32-bit and Itanium-based versions of Windows Server 2003. You can find out whether a particular feature is available on a particular version by reading the topics about that feature. For more information, see Accessing information about compatibility and Considerations for 64-bit drivers and software.

- **Editions of Windows Server 2003**

  Windows Server 2003 comprises multiple editions so that you can choose which sets of features best address your needs. For example, the feature set of Windows Server 2003, Web Edition, is optimized for a particular purpose, and several features that do not relate to that purpose are not available in that edition. Procedure topics state when the functionality that they describe are not available on all editions. For more information about feature availability on specific editions of Windows Server 2003, see the Microsoft Web site.

- **Service packs**

  Service packs also affect how a server appears and functions, and topics describe the changes that result from installing Windows Server 2003 Service Pack 1 (SP1). For more information about SP1, see Technical Overview of Windows Server 2003 Service Pack 1 (SP1) and Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1.

- **Account permissions**

  You might not be able to perform certain tasks based on the permissions that have been granted to the account that you used to log on. For more information about account permissions, see User and computer accounts and Local user accounts.

- **User preference settings**

  Product Help for Windows Server 2003 documents functionality based on default settings for options such as menu styles and file locations. If these settings have been changed, you might need to follow a slightly different set of steps than those presented in Help topics. For more information about user preference settings, see Settings saved in a user profile and Setting environment variables.

**See Also**
**Concepts**
Roadmap to Help and Information
Access Windows Help
Server scalability

# Using the Windows Deployment and Resource Kits

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the Microsoft Windows Deployment and Resource Kits**
The Microsoft® Windows® Deployment and Resource Kits provide IT professionals with the technical information and tools they need to successfully deploy and operate Windows.

See the topics in this Help and Support site for the content in the Microsoft® Windows® Server 2003 Deployment Kit, the Microsoft® Windows® Server 2003 Resource Kit, and Microsoft® Windows® XP Professional Resource Kit Documentation. For more information about purchasing the kits, or to browse the documentation on the Web, go to the Windows Deployment and Resource Kits Web site. This Web site also contains other important documentation developed by the Windows Resource Kit teams, including the Windows Server 2003 System Administration Scripting Guide and Windows 2000 documentation.

**Important**

- Some of the links on this page are to English-only content on the Web, regardless of the language of your operating system. In addition, if you install a Windows Deployment or Resource Kit on a non-English operating system, you will see English content mixed with non-English content in the Help and Support Center table of contents.

**Using Windows Server 2003 Help and Support Center from Windows XP Professional**
Windows Server 2003 Help and Support Center contains technical content for the IT professional that you won't find in Windows XP Professional. If you administer servers running Windows Server 2003 from a computer running Windows XP Professional, you should install Windows Server 2003 Help and Support Center. When you do this, your installed Windows Deployment and Resource Kits can be viewed and searched from within the Windows Server 2003 Help and Support Center. This way you have easy access to a more complete documentation set for server deployment and administration.

To install Windows Server 2003 Help and Support Center on a computer running Windows XP Professional, follow the steps in the Help topic Install Help content from a Windows Server 2003 server operating system onto a computer running Windows XP Professional. Windows Server 2003 family Help content is available on the Windows Server 2003 Deployment Kit companion CD, the Windows Server 2003 Resource Kit companion CD, and the Windows Server 2003 operating system CD.

# Updated technical information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Updated technical information**

The Microsoft Product Support Services Web site contains an extensive collection of articles, troubleshooting wizards, files you can download, and other useful product information. The articles can help you understand more about the features included in the Windows Server 2003 family.

In addition, Microsoft Product Support Services contains other information you may find useful, such as frequently asked questions (FAQs), troubleshooting information, and information from newsgroups.

For more information, see Microsoft Product Support Services at the Microsoft Web site.

# Using TechNet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using TechNet**

TechNet is a service that contains resources for information technology (IT) professionals. This includes access to white papers and articles written by support professionals at Microsoft, career and industry information, discussion groups, technical feedback, and software. For more information, see Microsoft Technet at the Microsoft Web site.

# Using MSDN

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using MSDN**

The Microsoft Developer Network (MSDN) is a complete resource for developers. You can find downloads, guidelines for hardware and driver compatibility testing, and how-to guides to help you build device drivers for Windows operating systems. For more information, see MSDN at the Microsoft Web site.

# Accessing documentation on the Web

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing documentation on the Web**

To download the latest product documentation, see Windows Update at the Microsoft Web site. Windows Update is a catalog of items, such as drivers, patches, the latest Help files, and Internet products, that you can download to keep your computer up to date.

# Accessing Microsoft white papers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing Microsoft white papers**

Microsoft white papers, many of which are written by Microsoft system engineers and network specialists, provide in-depth information on a variety of subjects. For more information, see Technical White Papers at the Microsoft Web site.

# Accessing information about training and certification

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing information about training and certification**

The Training and Certification Web site provides information about Microsoft training, certification, and technical careers. For more information, see Microsoft Training and Certification at the Microsoft Web site.

# Accessing information about compatibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing information about compatibility**

There are several ways to assess the readiness of your system and to make sure that your hardware and software will run properly.

**Compatible Hardware and Software**

You can search the Compatible Hardware and Software area in Help and Support Center to find compatibility information about specific products. You can search by manufacturer or by product, or your can browse through the list of hardware and software devices to see if your device is compatible or if it needs an update. For more information, see Compatible Hardware and Software overview.

**Program Compatibility Wizard**

This wizard prompts you to test a program in different modes (environments) and with various settings. For example, if the program was originally designed to run on Windows 95, set the compatibility mode to Windows 95 and try running the program again. If this is successful, the program will start in that mode every time. You can also run the Program Compatibility Wizard on the setup file for a program.

**Compatibility Mode**

Compatibility Mode provides an environment for running applications that more closely reflects the behavior of Windows 95, Windows 98, Windows NT 4.0, or Windows 2000 operating systems. These modes resolve several of the most common issues that prohibit older applications from working correctly on the Windows Server 2003 family of products. Applications that experience problems after migration may benefit from being started in one of these compatibility environments. You can also use Compatibility Mode to provide solutions for your own applications.

**Microsoft Web site**

You can search for compatible computers, hardware devices, and software applications at the Microsoft Web site. For more information, see Hardware compatibility information in the Windows Catalog (search page) at the Microsoft Web site.

**Notes**

- To get the most current compatibility information when using Compatible Hardware and Software in Help and Support Center, you must be connected to the Internet.

- Compatible Hardware and Software in Help and Support Center is not necessarily a complete list of programs and devices that are compatible with this release. It is a collection of products for which test data has been submitted to Microsoft. If you do not find a specific product, you can contact the manufacturer and ask the manufacturer about registering its product in the Windows Logo Program.

# Accessing information about accessibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing information about accessibility**

The Windows Server 2003 family contains features to help people with disabilities use the product. In addition, Microsoft offers other products and services for people with disabilities. For more information, and for information about products and services offered by other organizations, see Accessibility for People with Disabilities.

# Copyright

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Copyright**
Information in this document, including URL and other Internet Web site references, is subject to change without notice. The entire risk of the use or the results of the use of this document remains with the user. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious and no association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

All public network IDs or addresses referenced in this document are for demonstration/example purposes only. Microsoft's inclusion of such IDs and addresses in this document does not grant you any express or implied right to test or otherwise use or access any such network ID or address.

The links in this area will let you leave Microsoft's product and/or Web site. The linked sites are not under the control of Microsoft and Microsoft is not responsible for the contents of any linked site or any link contained in a linked site, or any changes or updates to such sites. Microsoft is not responsible for webcasting or any other form of transmission received from any linked site. Microsoft is providing these links to you only as a convenience, and the inclusion of any link does not imply endorsement by Microsoft of the site.

For your convenience, Microsoft may make available on this service or in its software products, tools and utilities for use and/or download. Microsoft does not make any assurances with regard to the accuracy of the results or output that derives from such use of any such tools and utilities. Please respect the intellectual property rights of others when using the tools and utilities made available on this service or in Microsoft software products.

© 1985-2005 Microsoft Corporation. All rights reserved.

Active Desktop, Active Directory, ActiveMovie, ActiveX, Authenticode, BackOffice, Direct3D, DirectAnimation, DirectDraw, DirectInput, DirectMusic, DirectPlay, DirectShow, DirectSound, DirectX, DoubleSpace, DriveSpace, FrontPage, IntelliMirror, IntelliMouse, JScript, Microsoft, Microsoft Press, Microsoft QuickBasic, MSDN, MS-DOS, MSN, Natural, NetMeeting, OpenType, Outlook, PowerPoint, SideWinder, Visual Basic, Visual C++, Visual FoxPro, Win32, Windows, Windows Media, Windows NT, and Wingdings are either registered trademarks or trademarks of Microsoft Corporation in the U.S.A. and/or other regions and/or countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

**Acknowledgements**
Portions of this product are based in part on the work of Mark H. Colburn and sponsored by the USENIX Association. Copyright © 1989 Mark H. Colburn. All rights reserved.

This product includes software developed by the University of California, Berkeley and its contributors.

Portions of this product are based in part on the work of the Regents of the University of California, Berkeley and its contributors. Because Microsoft has included the Regents of the University of California, Berkeley, software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1985, 1988 Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Portions of this product are based in part on the work of Greg Roelofs. Because Microsoft has included the Greg Roelofs software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1998-1999 Greg Roelofs. All rights reserved.

This software is provided "as is," without warranty of any kind, express or implied. In no event shall the author or contributors be held liable for any damages arising in any way from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. Redistributions of source code must retain the above copyright notice, disclaimer, and this list of conditions.

2. Redistributions in binary form must reproduce the above copyright notice, disclaimer, and this list of conditions in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

This product includes software developed by Greg Roelofs and contributors for the book, *PNG: The Definitive Guide*, published by O'Reilly and Associates.

Portions of this software are based in part on the work of Hewlett-Packard Company. Because Microsoft has included the Hewlett-Packard Company software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1994 Hewlett-Packard Company

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Hewlett-Packard Company makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Portions of this software are based in part on the work of the University of Southern California. Because Microsoft has included the University of Southern California software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1996 by the University of Southern California. All rights reserved.

Permission to use, copy, modify, and distribute this software and its documentation in source and binary forms for any purpose and without fee is hereby granted, provided that both the above copyright notice and this permission notice appear in all copies - and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed in part by the University of Southern California, Information Sciences Institute. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

THE UNIVERSITY OF SOUTHERN CALIFORNIA makes no representations about the suitability of this software for any purpose. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Portions of this software are based in part on the work of Luigi Rizzo. Because Microsoft has included the Luigi Rizzo software in this product, Microsoft is required to include the following text that accompanied such software:

© 1997-98 Luigi Rizzo (luigi@iet.unipi.it)

Portions derived from code by Phil Karn (karn@ka9q.ampr.org), Robert Morelos-Zaragoza (robert@spectra.eng.hawaii.edu) and Hari Thirumoorthy (harit@spectra.eng.hawaii.edu), Aug 1995

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions of this software are based in part on the work of Massachusetts Institute of Technology. Because Microsoft has included the Massachusetts Institute of Technology software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1989,1990 by the Massachusetts Institute of Technology. All Rights Reserved.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Under U.S. law, this software may not be exported outside the US without license from the U.S. Commerce department.

Portions of this software are based in part on the work of Regents of The University of Michigan. Because Microsoft has included the Regents of The University of Michigan software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1995,1996 Regents of The University of Michigan. All Rights Reserved.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of The University of Michigan not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. This software is supplied as is without expressed or implied warranties of any kind.

# Getting Started

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Getting Started**
The Getting Started topics provide information about new features as well as information designed to help you install or upgrade a network with one to five servers and 100 or fewer clients. If you are an administrator of a larger collection of servers, also see the Windows Server 2003 Deployment Kit at the Microsoft Windows Server System Web site. For information about obtaining the Windows Deployment and Resource Kits, see Using the Windows Deployment and Resource Kits.

For additional information, see:

- Step-by-step guides for deployment
- New Features
- Installing and Upgrading the Operating System
- Installing and Upgrading on Cluster Nodes
- Configuring Roles for Your Server
- Windows Product Activation

# Step-by-step guides for deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Step-by-step guides for deployment**
The Microsoft Web site provides tutorials that you can use to learn about and deploy the operating system. These technical walkthroughs provide step-by-step instructions and illustrations for installing and configuring key features.

You can access the Step-by-Step Guides at the Microsoft Web site.

The following list includes some of the topics that you can find in these guides:

**Infrastructure**

- Building a common network infrastructure

**Active Directory**

- Setting up additional domain controllers
- Delegating control of objects
- Managing Active Directory
- Managing replication topology
- Using LDAP and VBScript to manage Active Directory
- Extending Active Directory using the Schema and display specifiers

**Management services**

- Understanding Group Policy
- Managing remote installation of the operating system
- Managing application deployments for groups of users and computers
- Managing user data and user settings

**Security services**

**Setting up and configuring certificates**

- Setting up a certification authority
- Configuring a domain to trust an external certification authority
- Managing Certificate Services
- Obtaining and managing certificates
- Configuring automatic certificate requests

**Using certificates**

- Requesting certificates using Web pages
- Mapping certificates to user accounts to be used on a Web site

**Additional security features**

- Configuring local security policy on individual computers
- Managing security in an enterprise network
- Encrypting files
- Using Kerberos V5
- Encrypting IP network traffic using Internet Protocol security (IPSec)

**File and print services**

- Creating logical views of directories and files

**Clustering**

- Installing a typical two-node cluster on Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition

# Considerations for 64-bit drivers and software

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Considerations for 64-bit drivers and software**

Device drivers that are 32-bit do not work on computers running 64-bit versions of the Windows operating systems. If you are attempting to install hardware that only has 32-bit drivers available, that hardware will not work correctly on 64-bit versions of Windows. Additionally, applications that attempt to install 32-bit device drivers on 64-bit versions of Windows will encounter errors.

For information about updating drivers and troubleshooting issues with device drivers for 64-bit versions of Windows, contact the manufacturer of the hardware or software.

# New Features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features**

Windows Server 2003, Standard Edition, gives you a high level of reliability and scalability, lowers your cost of computing with flexible management services, and provides the best foundation for running business applications.

Windows Server 2003, Web Edition, is an easy-to-deploy, easy-to-manage server operating system that is optimized to be a complete Web hosting package.

Windows Server 2003, Enterprise Edition, includes the full feature set of Windows Server 2003, Standard Edition, and it also offers support for greater than 4 GB of memory and multiple processors, as well as for clustering. For more information about clustering, see Windows Clustering.

Windows Server 2003, Datacenter Edition, includes the full feature set of Windows Server 2003, Standard Edition, and it adds the high availability, reliability, and scalability required for enterprise and larger departmental solutions.

For information about feature changes in Windows Server 2003 with Service Pack 1, see topics in Features for Management, Security, and Best Performance, Features for Directory, File, Print, and Collaboration Services, and

Features for Internet, Application, and Networking Services.

- For information about features and components that have changed names or locations, see New ways to do familiar Windows tasks.
- To view a table that compares features in this release with features in previous releases, see Feature comparison: Windows NT, Windows 2000, and the Windows Server 2003 family. This table is especially useful if you upgrade from Windows NT 4.0 or Windows 2000 and you want to examine the differences between features.
- To learn more about the features that are useful for a specific function or administrative task, see Choosing the features that are most useful for your server.
- For a concise list of features in alphabetical order, see Alphabetical list of features.
- For an overview of the features in these topics, see Overview of Features.

# New ways to do familiar Windows tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

Many components have changed names or locations since Windows NT 4.0 and Windows 2000. To find the new way to do a task, click the letter that the name of the component starts with, and then click the name of the component.


A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**

New ways to do familiar Active Directory tasks

Active Directory Domains & Trusts: New ways to do familiar tasks

New ways to do familiar Licensing tasks

Active Directory Users & Computers: New ways to do familiar tasks

New ways to do familiar Services for Macintosh tasks

**B**

Backup: New ways to do familiar tasks

Tape Devices: New ways to do familiar tasks

Backup: New ways to do familiar tasks

**C**

CD Player: New ways to do familiar tasks

New ways to do familiar Novell NetWare Integration tasks

New ways to do familiar Server Cluster tasks

New ways to do familiar command-line tasks

**D**

Devices: New ways to do familiar tasks

Devices: New ways to do familiar tasks

New ways to do DHCP tasks

New ways to do DHCP tasks

New ways to do Network Connections tasks

New ways to do familiar tasks for Disk Management

New ways to do familiar tasks for Disk Management

New ways to do familiar DNS tasks

New ways to do familiar DNS tasks

**F**

Find

Find computer

Find files or folders

Find printers

View Options: New ways to do familiar tasks

**H**
New ways to do familiar modem tasks

**I**
Inbox

Internet: New ways to do familiar tasks

Internet: New ways to do familiar tasks

**L**
New ways to do familiar Licensing tasks

New ways to do familiar performance monitoring tasks

**M**
New ways to do familiar Services for Macintosh tasks

New ways to do familiar Active Directory tasks

Media Player

Message Queuing (also known as MSMQ)

New ways to do familiar modem tasks

New ways to do familiar command-line tasks

MSMQ Explorer

Multimedia: New ways to do familiar tasks

My Briefcase: New ways to do familiar tasks

**N**
New ways to do Network Connections tasks

New ways to do familiar Network Load Balancing tasks

New ways to do familiar Network Monitor tasks

New ways to do familiar Novell NetWare Integration tasks

New ways to do familiar Novell NetWare Integration tasks

**O**
Outlook Express

**P**
New ways to do familiar performance monitoring tasks

New ways to do familiar modem tasks

Ports: New ways to do familiar tasks

New ways to do familiar Printing tasks

New ways to do familiar Printing tasks

New ways to do familiar tasks with TCP/IP

New ways to do VPN tasks

**R**
Regional Settings: New ways to do familiar tasks

Regional Settings: New ways to do familiar tasks

New ways to do familiar Remote Access tasks

New ways to do familiar Remote Desktop for Administration tasks

New ways to do familiar Routing tasks

New ways to do familiar Routing tasks

New ways to do familiar Remote Access tasks

**S**

SCSI Adapters: New ways to do familiar tasks

SCSI Adapters: New ways to do familiar tasks

New ways to do familiar Security tasks

Server: New ways to do familiar tasks

New ways to do familiar Server Cluster tasks

Active Directory Users & Computers: New ways to do familiar tasks

New ways to do familiar Services for Macintosh tasks

Server: New ways to do familiar tasks

New ways to do familiar SNMP tasks

Sounds

Multimedia: New ways to do familiar tasks

New ways to do familiar performance monitoring tasks

**T**

Tape Devices: New ways to do familiar tasks

New ways to do familiar Telephony tasks

New ways to do familiar tasks with TCP/IP

New ways to do familiar Telephony tasks

New ways to do familiar Terminal Server tasks

New ways to do familiar Remote Desktop for Administration tasks

**U**

Active Directory Users & Computers: New ways to do familiar tasks

New ways to do User Profiles tasks

**V**

View Options: New ways to do familiar tasks

New ways to do VPN tasks

**W**

CD Player: New ways to do familiar tasks

New ways to do WINS tasks

New ways to do familiar Network Load Balancing tasks

# Feature comparison: Windows NT, Windows 2000, and the Windows Server 2003 family

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Feature comparison: Windows NT, Windows 2000, and the Windows Server 2003 family**
The following tables provide a feature comparison based on your upgrade path. For example, if you upgrade from Windows NT 4.0, select the "Windows NT 4.0 to the Windows Server 2003 family" option below. For links to more information about the features in this release, see New Features.

The tables, which are categorized by feature area, display information about improved features that were introduced in earlier operating systems, as well as features that are new to this release.

Some features are not available on some of the products in the Windows Server 2003 family. For more information, click the links in the following tables. For information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

Click the option that describes your upgrade path:

- Windows NT 4.0 to the Windows Server 2003 family
- Windows 2000 to the Windows Server 2003 family

**Windows NT 4.0 to the Windows Server 2003 family**

| Feature area | Features that were added since Windows NT 4.0 and are updated in this release | Features that are new to the Windows Server 2003 family |
|---|---|---|
| Server reliability | • Improved memory management<br>• Robust system architecture<br>• Diagnostic tools<br>• Safe Mode start<br>• Recovery Console<br>• Windows File Protection<br>• Enhanced Backup utility | • Automated System Recovery (ASR)<br>• Compatibility Mode |
| Server availability | • Server clusters<br>• Network Load Balancing clusters<br>• Fewer server restarts<br>• Plug and Play<br>• Start from mirrored volumes | • User State Migration Tool<br>• Emergency Management Services |

| | | |
|---|---|---|
| Server scalability | • Hardware scalability<br>• I2O support<br>• Symmetric multiprocessing enhancements<br>• Network Load Balancing<br>• Server clusters<br>• Enterprise Memory Architecture | • Internet Information Services (IIS) 6.0 |
| Server manageability | • IntelliMirror<br>• Windows Management Instrumentation (WMI)<br>• Remote Installation Services (RIS)<br>• Manage Your Server | • Remote Desktop for Administration<br>• Remotely administered server support<br>• User State Migration Tool<br>• Emergency Management Services<br>• Resultant Set of Policy (RSOP)<br>• Authorization Manager |
| Performance and tuning features | • Winsock Direct<br>• I2O support<br>• CPU throttling for Web applications or sites<br>• Process Accounting for Web applications or sites<br>• Predictable, end-to-end Quality of Service (QoS)<br>• Network Load Balancing<br>• New performance tools<br>• Disk Defragmenter<br>• Event Tracing enhancements | • Performance Options<br>• New command-line utilities for System Monitor |

| | | |
|---|---|---|
| Authentication and smart card support | • Full support for the Kerberos V5 protocol<br>• Public key infrastructure (PKI), Certificate Services, and smart cards<br>• Virtual private networking<br>• Flexible, secure authentication and authorization<br>• Internet Authentication Service (IAS)<br>• Routing and Remote Access | • Stored User Names and Passwords<br>• Smart card access through Terminal Server |
| Additional security features | • Encrypting File System (EFS)<br>• Internet Protocol Security (IPSec) | • TCP/UDP port ownership<br>• Software restriction policies |
| Remote management | • Microsoft Management Console (MMC)<br>• Scripting with Windows Script Host<br>• Remote Installation Services (RIS) | • Remote Desktop for Administration<br>• Remote assistance<br>• Web Interface for Remote Administration (Windows Server 2003, Web Edition only) |
| Maintenance and product updates | • Windows Update | • Does not apply |
| Internationalization features | • Multilingual support<br>• Multiregional support | • New Phonetic Input Method Editor (IME) for Traditional Chinese |

| | | |
|---|---|---|
| Active Directory | • Simplified user and network-resource management<br>• Flexible, secure authentication and authorization<br>• Directory consolidation<br>• Directory-enabled applications and infrastructure<br>• Scalability without complexity<br>• Use of Internet standards<br>• A powerful development environment<br>• Replication and trust monitoring<br>• Message Queuing distribution lists | • For a list of the Active Directory features that are new in this release, see New features for Active Directory. |
| Disk and file management features | • Disk Management<br>• Disk quota support<br>• Distributed File System<br>• Distributed Link Tracking<br>• Distributed Authoring and Versioning (DAV)<br>• High-performance content indexing<br>• Encrypting File System (EFS)<br>• Removable Storage and Remote Storage | • Remote document sharing<br>• Command-line support for disk management<br>• GUID partition table (GPT). Available on all versions of Windows Server 2003 with Service Pack 1 (SP1) and on all versions of Windows Server 2003, x64 Edition.<br>• Shadow copies of shared folders |

| | | |
|---|---|---|
| Printer and protocol support | • Printers published in Active Directory<br>• Broad device and protocol support<br>• Plug and Play<br>• Internet Printing Protocol<br>• Printer management features<br>• Network printing features | • New command-line support for printing<br>• Easier printer driver installation on server clusters |
| Internet and e-mail services and features | • Active Server Pages (ASP)<br>• Windows Media Services<br>• Web Distributed Authoring and Versioning (WebDAV)<br>• Application recycling | • Microsoft .NET Framework components<br>• Internet Information Services (IIS) 6.0<br>• Universal Description, Discovery, and Integration (UDDI) Services<br>• E-mail services (POP3, SMTP)<br>• Windows Management Instrumentation (WMI) Provider |
| Features for collaboration | • Windows Media Services<br>• Predictable, end-to-end Quality of Service (QoS)<br>• Asynchronous transfer mode (ATM) | • Microsoft .NET Framework<br>• Audio Acceleration on the remote access server |
| Support for existing systems(see also Application support) | • Industry standards<br>• Support for existing server applications<br>• Integration with existing operating systems | • Compatibility Mode<br>• Program Compatibility Wizard |

| | | |
|---|---|---|
| Application support | • Widest range of turnkey applications<br>• Comprehensive Internet application services<br>• Microsoft Active Accessibility<br>• TAPI 3.1<br>• Transaction services<br>• Message Queuing<br>• Component Services<br>• Job objects<br>• Host data integration | • Compatibility Mode<br>• Web gardens in Internet Information Services (IIS) |
| Management and deployment | • Change and configuration management with Group Policy<br>• Directory consolidation in Active Directory<br>• IntelliMirror<br>• Remote Installation Services (RIS)<br>• Windows Management Instrumentation (WMI)<br>• Side-by-side component sharing<br>• Add or Remove Programs<br>• Windows Installer<br>• Application management | • Remotely administered server support<br>• Resultant Set of Policy (RSOP) |

| | | |
|---|---|---|
| Terminal Server features | • Centralized deployment of programs<br>• Remote access to applications<br>• Single application access<br>• Terminal Services Manager<br>• Remote control | • Audio redirection |
| Development support | • Visual Studio .NET<br>• Multilingual technology | • Automatic generation of link IDs<br>• Background Intelligent Transfer Service (BITS) |
| Network protocols and technologies | • Winsock Direct<br>• DHCP with DNS and Active Directory<br>• Internet Connection Sharing (ICS)<br>• Network Address Translation (NAT)<br>• Virtual private networking<br>• Routing and Remote Access<br>• Asynchronous transfer mode (ATM) | • Simple Object Access Protocol (SOAP)<br>• Fibre Channel support<br>• Automatic configuration for multiple networks<br>• Web Digital Authoring and Versioning (WebDAV) redirector<br>• Internet Authentication Service (IAS) as RADIUS proxy<br>• Wireless networking<br>• IP version 6 (IPv6) |

**Windows 2000 to the Windows Server 2003 family**

| Feature area | Features that were added in Windows 2000 and are updated in this release | Features that are new to the Windows Server 2003 family |
|---|---|---|
| Server reliability | • Does not apply. | • Automated System Recovery (ASR)<br>• Compatibility Mode |
| Server availability | • Storage area network (SAN) and network access server (NAS) support<br>• Network Load Balancing enhancements | • User State Migration Tool<br>• Emergency Management Services |
| Server scalability | • Network Load Balancing enhancements | • Internet Information Services (IIS) 6.0 |

| | | |
|---|---|---|
| Server manageability | • Windows Management Instrumentation (WMI) improvements<br>• Remote Installation Services (RIS) improvements | • Remote Desktop for Administration<br>• Remotely administered server support<br>• User State Migration Tool<br>• Emergency Management Services<br>• Resultant Set of Policy (RSoP)<br>• Authorization Manager |
| Performance and tuning features | • System Monitor enhancements<br>• Performance Logs and Alerts improvements | • Performance Options<br>• New command-line utilities for System Monitor |
| Authentication and smart card support | • Subordinate certification authority (CA) improvements | • Stored User Names and Passwords<br>• Smart card access through Terminal Server |
| Additional security features | • Encrypting File System (EFS) improvements<br>• Internet Protocol Security (IPSec) monitoring improvements | • TCP/UDP Port Ownership<br>• Software restriction policies |
| Remote management | • Remote Installation Services (RIS) improvements | • Remote Desktop for Administration<br>• Remote assistance<br>• Web Interface for Remote Administration (Windows Server 2003, Web Edition only) |
| Maintenance and product updates | • Windows Update enhancements | • Does not apply |
| Internationalization features | • Improved Japanese Input Method Editor (IME)<br>• Improved Korean IME<br>• Improved Simplified Chinese IME | • New Phonetic IME for Traditional Chinese |

| | | |
|---|---|---|
| Active Directory | • For a list of the Active Directory features that are new in this release, see New features for Active Directory. | • For a list of the Active Directory features that are new in this release, see New features for Active Directory. |
| Disk and file management features | • File management improvements | • Remote document sharing<br>• Command-line support for disk management<br>• GUID partition table (GPT). Available on all versions of Windows Server 2003 with Service Pack 1 (SP1) and all versions of Windows Server 2003, x64 Edition.<br>• Shadow copies of shared folders |
| Printer and protocol support | • Easier print driver installation on server clusters | • New command-line support for printing |
| Internet and e-mail services and features | • Does not apply | • Microsoft .NET Framework<br>• Internet Information Services (IIS) 6.0<br>• Universal Description, Discovery, and Integration (UDDI) Services<br>• E-mail services (POP3, SMTP)<br>• Windows Management Instrumentation (WMI) Provider for IIS |
| Features for collaboration | • NetMeeting enhancements | • Microsoft .NET Framework<br>• Audio Acceleration on the remote access server |
| Support for existing systems(see also Application support) | • Does not apply | • Compatibility Mode<br>• Program Compatibility Wizard |
| Application support | • Does not apply | • Compatibility Mode<br>• Web gardens in Internet Information Services (IIS) |
| Management and deployment | • Windows Management Instrumentation (WMI) improvements<br>• Scripting support<br>• Remote Installation Services (RIS) improvements | • Remotely administered server support<br>• Resultant Set of Policy (RSOP) |

| | | |
|---|---|---|
| Terminal Server features | • Resolution and color enhancements<br>• Group Policy integration | • Audio redirection |
| Development support | • Does not apply | • Automatic generation of link IDs<br>• Background Intelligent Transfer Service (BITS) |
| Network protocols and technologies | • Connection Manager improvements | • Simple Object Access Protocol (SOAP)<br>• Fibre Channel support<br>• Automatic configuration for multiple networks<br>• Web Digital Authoring and Versioning (WebDAV) redirector<br>• Internet Authentication Service (IAS) as RADIUS proxy<br>• Wireless networking<br>• IP version 6 (IPv6) |

# Choosing the features that are most useful for your server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing the features that are most useful for your server**

To learn more about the features that are most useful for your server according to its function, see the following table. For links to more information about the features in this release, see New Features.

| Server type | Description of the feature or type of support | Information on how to use the feature |
|---|---|---|
| File server | • Features for Active Directory<br>• Disk and file management features | • Administering Active Directory<br>• Disks and Data<br>• Client Computers |
| Print server | • Features for Active Directory<br>• Printer and protocol support | • Administering Active Directory<br>• Printing<br>• Printers and Faxes |
| Web server | • Overview of Windows Server 2003, Web Edition<br>• Internet and e-mail services and features<br>• Application support<br>**Related areas:**<br>Application server<br>Networking/communications<br>Server security | • Internet and E-mail Services<br>• Internet Information Services (IIS) 6.0<br>• Technologies for Application and Development Support |
| Multimedia server | • Features for collaboration | • Using Windows Media Services<br>• Quality of Service (QoS)<br>• Windows ATM Services |

| | | |
|---|---|---|
| Application server | • Microsoft .NET Framework<br>• Support for existing systems<br>• Terminal Server features<br>• Application support<br>• Management and deployment<br>• Performance and tuning features<br>• Development support | • Remote Desktop for Administration<br>• Group Policy (pre-GPMC)<br>• IntelliMirror<br>• Monitoring and Status Tools<br>• Quality of Service (QoS)<br>• Windows ATM Services<br>• Enterprise Memory Architecture (Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, not including 64-bit versions)<br>• Network Load Balancing Clusters<br>• Windows Clustering (Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition only)<br>• Technologies for Application and Development Support<br><br>See also the Microsoft Platform SDK. You can download the SDK at the Platform SDK Update Web site. |
| Networking/ communications | • Network protocols and technologies<br>• Support for existing systems<br>• Performance and tuning features | • Network Services<br>• Configuring Communication Links<br>• Monitoring and Status Tools<br>• Quality of Service (QoS)<br>• Windows ATM Services |
| Server reliability, availability, scalability, and manageability | • Server reliability<br>• Server availability<br>• Server scalability<br>• Server manageability | • Disks and Data<br>• Monitoring and Status Tools<br>• Plug and Play overview<br>• Network Load Balancing Clusters<br>• Windows Clustering (Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition only) |

| | | |
|---|---|---|
| Server security | • Authentication and smart card support<br>• Additional security features | • Security<br>• New features in security<br>• Authentication<br>• Internet Authentication Service (not available on Windows Server 2003, Web Edition)<br>• Internet Protocol Security (IPSec) |
| Configuration and management | • Features for Active Directory<br>• Terminal Server features<br>• Remote management<br>• Disk and file management features<br>• Internationalization features<br>**Related area:**<br>File server | • Client Computers<br>• Active Directory Overview<br>• IntelliMirror<br>• Remote Administration<br>• Administration and Scripting Tools<br>• Remote Installation Services (not available on Windows Server 2003, Web Edition) |

# Alphabetical list of features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Alphabetical list of features**
This is a partial list of the extensive set of features included in the Windows Server 2003 family. For links to more information about the features in this release, see New Features.

- Active Directory
- Asynchronous transfer mode
- Automated System Recovery
- Certificate Services
- Clustering
- Compatibility Mode
- Component Services
- Connection Manager
- DHCP with DNS and Active Directory
- Disk Management
- Disk quota support
- E-mail services (POP3, SMTP)
- Encrypting File System
- Group Policy (part of Active Directory)
- Indexing Service
- Internet Authentication Service
- Internet Connection Sharing
- Internet Information Services 6.0
- Internet protocol security support
- Kerberos V5 protocol support
- Layer Two Tunneling protocol support
- Lightweight Directory Access protocol support
- Message Queuing

- Microsoft Management Console
- Network address translation
- Network Load Balancing clusters
- Operating system migration, support, and integration
- Plug and Play
- Public key infrastructure and smart card infrastructure
- Quality of Service
- Recovery Console
- Remote Assistance
- Remote Desktop for Administration
- Remote Installation Services
- Remote Storage
- Removable Storage
- Routing and Remote Access
- Safe Mode
- Security Configuration Wizard
- Server clusters (Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition only)
- Smart card infrastructure
- Software restriction policies
- TAPI 3.1
- Terminal Server
- Universal Description, Discovery, and Integration Services
- Virtual private networking
- Windows Firewall
- Windows Media Services
- Windows Script Host
- Windows Update
- Winsock Direct

**Active Directory**
Active Directory is the directory service for Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. It stores information about objects on the network and makes this information easy for administrators and users to find and use. Active Directory uses a structured data store as the basis for a logical, hierarchical organization of directory information. For more information, see Active Directory.

**Asynchronous transfer mode**
Asynchronous transfer mode (ATM) is a high-speed, connection-oriented protocol designed to transport multiple types of traffic across a network. It is applicable to both local area networks (LANs) and wide area networks (WANs). Using ATM, your network can simultaneously transport a wide variety of network traffic: voice, data, image, and video.

**Automated System Recovery**
Automated System Recovery (ASR) helps you recover your system in the event that your hard disk fails or your system is seriously damaged. You can easily prepare the necessary backup materials for ASR through wizards that are part of Backup. By using ASR in combination with standard backups (data and application backups), you can restore a system to the state it was in when you backed it up. For more information about ASR, see Automated System Recovery (ASR) overview. For more information about other system recovery options, see Disaster Recovery.

**Certificate Services**
Using Certificate Services and certificate management tools, you can deploy your own public key infrastructure (PKI). With a PKI, you can implement standards-based technologies, such as smart card logon capabilities, client authentication (through Secure Sockets Layer (SSL) and Transport Layer Security (TLS)), secure e-mail, digital signatures, and secure connectivity (using Internet protocol security (IPSec)). Using Certificate Services, you can set up and manage certification authorities (CAs) that issue and revoke X.509 v3 certificates. This means that you do not have to depend on commercial client authentication services, although you can integrate commercial client authentication into your PKI if you choose. For more information, see Certificate Services.

**Clustering**
A cluster is a group of independent computers that work together to run a common set of applications and provide the image of a single system to the client and application. For more information, see "Network Load Balancing clusters" or "server clusters" in this list.

Server clusters are available only on Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.

**Compatibility Mode**
A new Compatibility Mode feature ensures out-of-box compatibility for many popular applications. Compatibility Mode technology provides an environment that more closely reflects the behavior of Windows 95, Windows 98, Windows NT 4.0, or Windows 2000 operating systems. These modes resolve several of the most common issues that prevent older applications from working correctly. Applications that experience problems after migration might benefit from being started in one of these compatibility environments. Advanced users can also take advantage of this technology to create solutions for their own applications.

**Component Services**
Component Services is a set of services based on extensions of Component Object Model (COM) and on Microsoft Transaction Server (an earlier release of a component-based transaction processing system). Component Services provides threading and security, transaction management, object pooling, queued components, and application administration and packaging.

To see the complete Help documentation for COM+ applications and Component Services, click **Start**, point to **All programs**, point to **Administrative Tools**, and then click **Component Services**. In Component Services, click **Help**, and then click **Help Topics**.

**Connection Manager**
Connection Manager provides support for local and remote connections to your service by using a network of access points, such as those available worldwide through Internet service providers (ISP). For more information, see Connection Manager Administration Kit.

**DHCP with DNS and Active Directory**
Dynamic Host Configuration protocol (DHCP) works with DNS and Active Directory on Internet protocol (IP) networks, helping to free you from assigning and tracking static IP addresses. DHCP dynamically assigns IP addresses to computers or other resources that are connected to an IP network. For more information, see Network Services.

**Disk Management**
Disk Management is a graphical tool for managing disk storage that includes such features as:

- Support for dynamic volumes, working outside the restriction of four partitions per disk.
- Online disk management (fewer restarts).
- Local and remote drive management.
- Volume Mount Points, which you can use to connect, or mount, a local drive at any empty folder on a local NTFS-formatted volume.

For more information, see Disk Management.

**Disk quota support**
You can use disk quotas on volumes that are formatted with the NTFS file system to monitor and limit the amount of disk space that is available to individual users. You can define the responses that result when users exceed your specified thresholds. For more information, see Disk Quotas.

**E-mail services (POP3, SMTP)**
The POP3 (Post Office protocol 3) service provides e-mail transfer and retrieval services. Administrators can use the POP3 service to store and manage e-mail accounts on the mail server. Simple Mail Transfer protocol (SMTP) supports the transfer of e-mail between servers. For more information, see POP3 Service and Simple Mail Transfer Protocol (SMTP) service.

**Encrypting File System**
Encrypting File System (EFS) complements other access controls and provides an additional level of protection for your data. EFS runs as an integrated system service, making it easy to manage, difficult to attack, and transparent to the user. For more information, see Encrypting File System.

**Group Policy (part of Active Directory)**
You can use policy to define the settings and allowed actions for users and computers. In contrast with local policy, you can use Group Policy to set policies that apply across a given site, domain, or organizational unit in Active Directory. Policy-based management simplifies such tasks as system update operation, application installation, user profiles, and desktop-system lockdown. For more information, see Group Policy (pre-GPMC).

**Indexing Service**
Indexing Service provides a fast, easy, and secure way for users to search for information locally or on the network. Users can search in files in different formats and languages, either through the **Search** command on the **Start** menu or through Hypertext Markup Language (HTML) pages that they view in a browser.

**Internet Authentication Service**
Internet Authentication Service (IAS) provides a central point for managing authentication, authorization, accounting, and auditing of dial-up or virtual private network (VPN) users. IAS uses the Internet Engineering Task Force (IETF) protocol called Remote Authentication Dial-In User Service (RADIUS). For more information, see Internet Authentication Service.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition. In Windows Server 2003, Standard Edition, you can configure IAS with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. For more information about these limits, see Internet Authentication Service.

**Internet Connection Sharing**
With the Internet Connection Sharing (ICS) feature of Network Connections, you can connect your home network or small office network to the Internet. For example, you might have a home network that connects to the Internet through a dial-up connection. By enabling Internet Connection Sharing on the computer that uses the dial-up connection, you can provide network address translation, addressing, and name resolution services for all the computers on your network. For more information, see Internet Connection Sharing and network address translation.

Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Internet Information Services 6.0**
Internet Information Services (IIS) 6.0 is a full-featured Web server that provides the foundation for the Windows Server 2003 family and existing Web applications and Web services. IIS 6.0 offers dedicated application mode, which runs all application code in an isolated environment. IIS 6.0 also supports Web gardens, in which a set of equivalent processes on a computer each receive a share of the requests that are normally served by a single process, achieving better multiprocessor scalability.

**Note**

- By default, IIS is not installed with new installations of Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. You can add it by using the Configure Your Server Wizard or by using Add or Remove programs.

**Internet protocol security support**
You can use Internet protocol security (IPSec) to secure communications within an intranet and to create secure virtual private network (VPN) solutions across the Internet. IPSec was designed by the Internet Engineering Task Force (IETF), and it is an industry standard for encrypting TCP/IP traffic. For more information, see Internet Protocol Security (IPSec).

**Kerberos V5 protocol support**
Kerberos V5 is a mature, industry-standard network authentication protocol. With Kerberos V5 support, a fast, single logon process gives users the access they need to enterprise resources, as well as to other environments that support this protocol. Support for Kerberos V5 includes additional benefits, such as mutual authentication (client and server must both provide authentication) and delegated authentication (the user's credentials are tracked end to end). For more information, see Kerberos V5 authentication.

**Layer Two Tunneling protocol support**
Layer Two Tunneling protocol (L2TP) is a more secure version of Point-to-Point Tunneling protocol (PPTP), and it is used for tunneling, address assignment, and authentication. For more information, see Layer Two Tunneling Protocol.

**Lightweight Directory Access protocol support**
Lightweight Directory Access protocol (LDAP), an industry standard, is the primary access protocol for Active Directory. LDAP version 3 was defined by the Internet Engineering Task Force (IETF). For more information, see Directory access protocol.

**Message Queuing**
The Message Queuing feature helps developers build and deploy applications that run more reliably over networks, including the Internet. These applications can interoperate with applications running on different platforms, such as mainframe computers and UNIX-based systems. For more information, see Using Message Queuing.

**Microsoft Management Console**
You can use Microsoft Management Console (MMC) to arrange administrative tools and processes within a single interface. You can also delegate tasks to specific users by creating preconfigured MMC consoles for them. The consoles provide the users with the tools that you select. For more information, see Microsoft Management Console.

**Network address translation**
Network address translation (NAT) hides internally managed IP addresses from external networks by translating private internal addresses to public external addresses. This reduces IP address registration costs by letting you use private IP addresses internally, with translation to a small number of registered IP addresses externally. It also hides the internal network structure, reducing the risk of attacks against internal systems. For more information, see Understanding Network Address Translation.

**Network Load Balancing clusters**
Previously known as Windows NT Load Balancing Service (WLBS), Network Load Balancing distributes incoming TCP/IP traffic among multiple servers. Your clustered applications, especially Web server applications, can handle more traffic and provide higher availability and faster response times. For more information, see Network Load Balancing Clusters.

**Operating system migration, support, and integration**
The Windows Server 2003 family of products integrates seamlessly with existing systems, and it contains support for earlier Windows operating systems, as well as other popular operating systems. The following support is included:

- Interoperability with Windows NT Server 3.51 and 4.0.

- Support for clients that run a variety of operating systems, including Windows 3.*x*, Windows 95, Windows 98, Windows 2000, and Windows NT Workstation 4.0.

- UNIX interoperability, provided through an add-on pack called Windows Services for UNIX. Telnet services (for remote access and administration) are included in the Windows Server 2003 family, as well as in the add-on pack.

- Mainframe and midrange connectivity, using S/390 and AS/400 transaction and queuing gateways through Host Integration Server 2000.

- File Server for Macintosh, which allows Macintosh clients to use the TCP/IP protocol (AFP over IP) to share files and to access shares on a server running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. For more information, see Services for Macintosh.

  This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Plug and Play**
With Plug and Play, a combination of hardware and software support, the server can recognize and adapt to hardware configuration changes automatically, without your intervention and without restarting. For more information, see Plug and Play overview.

**Public key infrastructure and smart card infrastructure**
Using Certificate Services and certificate management tools, you can deploy your own public key infrastructure (PKI). With a PKI, you can implement standards-based technologies, such as smart card logon capabilities, client authentication (through Secure Sockets Layer (SSL) and Transport Layer Security (TLS)), secure e-mail, digital signatures, and secure connectivity (using Internet protocol security (IPSec)). Using Certificate Services, you can set up and manage certification authorities (CAs) that issue and revoke X.509 v3 certificates. This means that you do not have to depend on commercial client authentication services, although you can integrate commercial client authentication into your PKI if you choose. For more information, see Certificate Services and Smart Cards.

**Quality of Service**
Quality of Service (QoS)-based services and protocols provide a guaranteed, end-to-end, express delivery system for information across the network. For more information, see Quality of Service (QoS).

**Recovery Console**
With Recovery Console, you can start a command-line console on a system on which a software problem (such as a problem with drivers or files) is preventing the system from starting. From Recovery Console, you can use basic commands to try to recover the system. For more information, see Recovery Console overview.

**Remote Assistance**
You can use Remote Assistance to remotely administer a computer. If you have an invitation, Remote Assistance is a convenient way for you to connect to a remote computer from a computer running Windows XP or any product in the Windows Server 2003 family. After you are connected, you can view the remote computer's screen and chat in real time. If the person requesting assistance allows you to, you can even use the mouse and keyboard to work on the remote computer. For more information, see Windows interface administrative tool reference A-Z: Remote Assistance.

This feature is not available on the Itanium-based versions of the Windows operating systems.

**Remote Desktop for Administration**
With Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode), you can administer a computer from virtually any computer on your network. Based on Terminal Services technology, Remote Desktop for Administration is specifically designed for server management. For more information about this feature, see Remote Desktop for Administration.

**Remote Installation Services**
With Remote Installation Services (RIS), you can create installation images of operating systems or even of complete computer configurations, including desktop settings and applications. You can then make these images available to users at client computers. The client computers must support remote booting with the Pre-Boot eXecution Environment (PXE) ROM, or else they must be started with a remote-startup floppy disk. For more information, see Remote Installation Services.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Remote Storage**
Remote Storage uses criteria that you specify to automatically copy little-used files to removable media. If hard-disk space drops below specified levels, Remote Storage removes the (cached) file content from the disk. If the file is needed later, the content is automatically recalled from storage. For more information, see Remote Storage.

Remote Storage is not available on Windows Server 2003, Web Edition or on Windows Server 2003, Standard Edition. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

**Removable Storage**
Removable Storage makes it easy to track your removable storage media (tapes and optical discs) and to manage the hardware libraries, such as changers and jukeboxes, that contain them. Because removable optical discs and tapes are less

expensive per megabyte (MB) than hard disks, Removable Storage and Remote Storage can decrease your costs. For more information, see Removable Storage.

**Routing and Remote Access**
Routing and Remote Access replaces the Routing and Remote Access Service (RRAS) and Remote Access Service (RAS) features from Windows NT 4.0. Routing and Remote Access is a single, integrated service that terminates connections from either dial-up or virtual private network (VPN) clients or that provides routing (IP, IPX, and Services for Macintosh), or both. With Routing and Remote Access, your server can function as a remote access server, a VPN server, a gateway, or a branch-office router. For more information, see Routing and Remote Access.

**Safe Mode**
With Safe Mode, you can start the Windows operating system with a minimal set of drivers and services and then view a log showing the sequence of events at startup. By using Safe Mode, you can diagnose problems with drivers and other components that might be preventing normal startup. For more information, see Startup options.

**Security Configuration Wizard**
The Security Configuration Wizard reduces the attack surface of computers that are running Windows Server 2003 with Service Pack 1 (SP1). The wizard asks the user a series of questions designed to determine the functional requirements of the server. Any functionality that is not required by the roles being performed by the server is then disabled. In addition to being a fundamental security best practice, reducing the attack surface increases the diversity of your Windows landscape, and reduces the number of systems that need to be immediately updated if a security issue arises. For more information, see Windows Server TechCenter.

**Server clusters (Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition only)**
Server clusters provide high availability, scalability, and manageability for important resources and applications. Multiple servers (nodes) in a cluster remain in constant communication. If one of the nodes in a cluster is unavailable as a result of failure or maintenance, another node immediately begins providing service (a process known as failover). Users who access the cluster are constantly connected to server-based resources.

With Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, you can configure server clusters in one of three available cluster models, depending on your requirements. For more information, see Server Clusters.

**Smart card infrastructure**
Using Certificate Services and certificate management tools, you can deploy your own public key infrastructure (PKI). With a PKI, you can implement standards-based technologies, such as smart card logon capabilities, client authentication (through Secure Sockets Layer (SSL) and Transport Layer Security (TLS)), secure e-mail, digital signatures, and secure connectivity (using Internet protocol security (IPSec)). For more information, see Certificate Services and Smart Cards.

**Software restriction policies**
With software restriction policies, you can protect your computer environment from untrusted code by identifying and specifying which applications are allowed to run. For more information, see Software Restriction Policies.

**TAPI 3.1**
TAPI 3.1 unifies Internet protocol (IP) and traditional telephony to enable developers to create computer telephony applications that work as effectively over the Internet or an intranet as they do over the traditional telephone network. For more information, see Telephony.

**Terminal Server**
By using Terminal Server, a user can access programs running on the server from a variety of older devices. For example, a user can access a virtual Windows XP Professional desktop and x86-based Windows applications from hardware that cannot run the software locally. Terminal Server provides this capability for both Windows client devices and other client devices. (Other devices require add-on software by Citrix Systems.) For a more detailed description, see Terminal Server features.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition. Note that Remote Desktop for Administration is available on Windows Server 2003, Web Edition.

**Universal Description, Discovery, and Integration Services**
Universal Description, Discovery, and Integration (UDDI) is an industry specification for publishing and locating information about Web services. Some products in the Windows Server 2003 family include UDDI Services, a Web service that provides UDDI capabilities for use within an enterprise or across organizations. For more information, see UDDI Services overview.

UDDI Services is not included with Windows Server 2003, Web Edition. In addition, Windows Server 2003, Standard Edition, supports only stand-alone installations of UDDI Services. Distributed installation support is available with Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. In a stand-alone installation of UDDI Services, both the UDDI Web server component and the UDDI database component are installed onto a single server. In a distributed

installation, UDDI components are distributed across multiple servers.

For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

**Virtual private networking**
You can give users ready access to your organization's network even when they are out of the office--and reduce the cost of this access--by implementing a virtual private network (VPN). The VPN connection creates a secure tunnel across the Internet into the private network.

There are two types of VPN technology in the Windows Server 2003 family:

- Point-to-Point Tunneling protocol (PPTP), which employs user-level Point-to-Point protocol (PPP) authentication methods and Microsoft Point-to-Point Encryption (MPPE) for data encryption.

- Layer Two Tunneling protocol (L2TP) with Internet protocol security (IPSec). L2TP employs user-level PPP authentication methods and computer-level certificates with IPSec for data encryption.

On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

For more information about virtual private networking, see Virtual Private Networks.

**Windows Firewall**
Windows Firewall is a host firewall technology. As a host firewall, Windows Firewall runs on each of your clients and servers. It provides protection from network attacks that pass through your perimeter network or originate inside your organization, such as Trojan horse attacks, port scanning attacks, and worms. Like many firewall technologies, Windows Firewall is a stateful firewall. As such, Windows Firewall inspects and filters all Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic. Unsolicited incoming traffic is dropped unless it is a response to a request by the host (solicited traffic), or it is specifically allowed (excepted traffic). You can specify excepted traffic according to port number, application name, or service name by configuring Windows Firewall settings. With the exception of some Internet Control Message Protocol (ICMP) messages, Windows Firewall allows all outgoing traffic.

Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

For more information, see Windows Server TechCenter.

**Windows Media Services**
Windows Media Services provides support for the distribution of streaming audio and video over corporate intranets and the internet. In Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, Windows Media Services delivers advanced streaming functionality, such as multicasting, wireless network support, internet authentication, server plug-ins, and Cache/proxy APIs.

Windows Media Services is not included with Windows Server 2003, Web Edition, or the Itanium-based versions of Windows Server 2003.

**Windows Script Host**
By using Windows Script Host (WSH), you can automate such actions as creating a shortcut and connecting to and disconnecting from a network server. WSH is language independent: you can write scripts in common scripting languages, such as Microsoft Visual Basic, Scripting Edition, and JScript. For more information, see Windows Script Host.

**Windows Update**
Windows Update is an online extension of Windows that detects and replaces outdated Windows system files with the most recent versions. Windows Update can significantly increase computer system stability over a long period. It can also reduce the load on administrators and support staff charged with maintaining a productive operating environment. Windows Update requires an internet connection. For more information, see Windows Update.

**Winsock Direct**
Winsock Direct, now available on all server editions, enables applications that use Winsock to perform faster and with less overhead when they communicate across a system area network (SAN). If there is a SAN in place, Winsock Direct has the effect of streamlining communications between distributed components.

# Overview of Features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Overview of features**
This section provides information about new and improved features in Microsoft® Windows Server 2003 family operating systems. For links to more information about the features in this release, see New Features.

This section covers:

- Features for Management, Security, and Best Performance
- Features for Directory, File, Print, and Collaboration Services
- Features for Internet, Application, and Networking Services

# Features for Management, Security, and Best Performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Features for management, security, and best performance**
The features in the Windows Server 2003 family ensure that data access, storage, and exchange comply with your needs while providing strong security and performance.

For links to more information about the features in this release, see New Features.

This section covers:

- Server reliability
- Server availability
- Server scalability
- Server manageability
- Performance and tuning features
- Authentication and smart card support
- Additional security features
- Remote management
- Maintenance and product updates
- Internationalization features

# Server reliability

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server reliability**

This topic contains a brief overview of the features that support server reliability. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

Reliability is the likelihood of a computer system or device continuing to function over a given period of time and under specified conditions. For descriptions of features that are related to availability, which is closely related to reliability, see Server availability.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT) that help provide increased levels of server reliability:

**Improved memory management**

Improvements to memory management in both Windows 2000 and the Windows Server 2003 family mean greater speed, reliability, and flexibility for your servers.

**Robust system architecture**

A robust system provides you with greater reliability, even if one application or service encounters problems. For example, with better handling of nonresponsive applications, you can move, resize, minimize, and close the window of a nonresponsive application. Also, in the event that an updated driver does not function properly, you can restore the previous drivers. The system's support infrastructure ensures that applications that worked on previous Windows operating systems will continue to work on the Windows Server 2003 family of products.

**Diagnostic tools**

You can monitor system states and prevent problems by using the improved diagnostic tools that are included with this release. For more information, see Monitoring and Status Tools.

**Safe Mode start**

With Safe Mode, you can start the operating system with a minimal set of drivers and services and then view a log showing the sequence of events at startup. By using Safe Mode, you can diagnose problems with drivers and other components that might be preventing normal startup. For more information, see Disaster Recovery.

**Recovery Console**

With Recovery Console, you can start a command-line console on a system on which a software problem (such as a problem with drivers or files) is preventing the system from starting. From Recovery Console, you can use basic commands to help recover the system. For more information, see Recovery Console overview.

**Windows File Protection**

Windows File Protection prevents the replacement of protected system files, such as .sys, .dll, .ocx, .ttf, .fon, and .exe files. Windows File Protection runs in the background and prevents other programs from changing these files, which are needed by the operating system. For more information, see Windows File Protection.

**Enhanced Backup utility**

You can back up data to a wide variety of storage media, such as tape drives, external hard disks, and Zip disks. For more information, see Backup.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of server reliability:

**Automated System Recovery**

Automated System Recovery (ASR) helps you recover your system in the event that a hard disk fails or the system is seriously damaged. You can easily prepare the necessary backup materials for ASR through wizards that are part of Backup. By using ASR in combination with standard backups (data and application backups), you can restore a system to the state it was in when you backed it up. For more information about ASR, see Automated System Recovery (ASR) overview. For more information about other system recovery options, see Disaster Recovery.

**Compatibility Mode**

A new Compatibility Mode feature ensures out-of-box compatibility for many popular applications. Compatibility Mode technology provides an environment that more closely reflects the behavior of Windows 95, Windows 98, Windows NT 4.0, or Windows 2000 operating systems. These modes resolve several of the most common issues that prevent older applications from working correctly. Applications that experience problems after migration may benefit from being started in one of these compatibility environments. Advanced users can also take advantage of this technology to create solutions for their own applications.

# Server availability

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server availability**

This topic contains a brief overview of the features that support server availability. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

Availability is a level of service provided by applications, services, or systems. Highly available systems have minimal downtime, whether planned or unplanned. The Windows Server 2003 family supports high levels of server availability for your critical business solutions. This operating system also gives you clustering, advanced fault tolerance, and file system recoverability features.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT) that help provide increased levels of server availability:

**Server clusters (Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition only)**

Server clusters provide high availability, scalability, and manageability for important resources and applications. Multiple servers (nodes) in a cluster remain in constant communication. If one of the nodes in a cluster is unavailable as a result of failure or maintenance, another node immediately begins providing service (a process known as failover). Users who access the cluster are constantly connected to server-based resources. For more information, see Windows Clustering.

**Network Load Balancing**

Network Load Balancing enhances the availability and scalability of Internet server programs on mission-critical servers. The two principal functions of Network Load Balancing are to provide high availability for Web server programs and to scale server performance. By combining the resources of two or more computers into a single cluster, Network Load Balancing delivers reliability and performance by distributing the incoming network traffic. For more information, see Network Load Balancing Overview.

**Fewer server restarts**

Configuring hardware and software is easier, because you can perform many activities without restarting the server. These activities include, but are not limited to:

- Extending a storage volume
- Configuring network protocols
- Managing storage dynamically (Dynamic Volume Management)
- Reconfiguring settings on Peripheral Component Interconnect (PCI) and other Plug and Play hardware

**Plug and Play**

With Plug and Play, a combination of hardware and software support, the server can recognize and adapt to hardware configuration changes automatically, without your intervention and without restarting. With Plug and Play, planned system restarts are reduced by more than 90 percent.

**Start from mirrored volumes**

With appropriate preparation, you can start your server from a redundant copy of data on a mirrored volume, which increases system availability during server recovery or planned maintenance. For more information, see Using mirrored volumes.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of server availability:

**User State Migration Tool**

The User State Migration Tool (USMT) aids deployment by capturing and restoring user settings, files, and documents. Users

do not have to reconfigure desktop settings for such things as e-mail servers, proxy servers, desktop color schemes, and desktop wallpaper.

**Emergency Management Services**

With Emergency Management Services, combined with the appropriate hardware, you can perform remote management and system recovery tasks, even when the server is not available through the standard remote administration tools and mechanisms. For more information, see Emergency Management Services.

**Storage area network and network-attached storage support**

Storage area network (SAN) and network-attached storage (NAS) support in the Windows Server 2003 family has a number of new system enhancements to increase server availability. These improvements include fault tolerant, network-attached scenarios, such as connecting to SAN and NAS servers. When connecting to SANs, the Windows Server 2003 family supports multipath failover, which enables redundant paths from the host server (running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or any product in the Windows Server 2003 family) to the SAN volume. The Windows Server 2003 family also supports connecting to NAS and being used as a NAS.

**Network Load Balancing enhancements**

Several improvements have been made to Network Load Balancing in this release:

- Network Load Balancing can now be bound to multiple network adapters, so that you can configure multiple independent clusters on each host. For more information about virtual clusters, see Virtual clusters.

- Network Load Balancing uses Internet Group Management Protocol (IGMP) support for limiting switch flooding, so that traffic intended for a Network Load Balancing cluster passes through only those ports that serve the cluster hosts and not through all switch ports.

- With Network Load Balancing Manager, you can create new Network Load Balancing clusters, and you can configure and manage clusters and all the cluster's hosts from a single remote or local computer.

# Server scalability

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server scalability**

This topic contains a brief overview of the features that support server scalability. It is divided into three sections: New and updated features since Windows Server 2003 (without SP1), New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

Scalability is the measure of how well a computer, service, or application can grow to meet increasing performance demands. For server clusters, it is the ability to incrementally add one or more systems to an existing cluster when the overall load of the cluster exceeds its capabilities. Products in the Windows Server 2003 family scale from small workgroups to enterprise data center deployments, and they support up to 64 processors and advanced I/O. The Windows Server 2003 family also integrates network load balancing and multiprocessor optimizations for your business applications.

**New and updated features since Windows Server 2003 (without SP1)**

Windows Server 2003 operating systems with Service Pack 1 (SP1) offer the following improvements (compared to Windows Server 2003 without SP1) that help provide increased levels of server scalability:

**Physical Address Extension (PAE)**

To support the addition of Data Execution Prevention (DEP), which is also known as no-execute page protection, the following changes have been made:

- PAE is automatically enabled on computers running Windows Server 2003 with Service Pack 1 and Windows XP with Service Pack 2 when DEP is also enabled on a computer with a processor that supports the no-execute page protection feature.

- When PAE mode is enabled on computers running Windows Server 2003, Standard Edition with Service Pack 1 and Windows XP with Service Pack 2, physical address space is limited to 4 gigabytes (GB). Limiting physical address space to 4 GB helps prevent driver compatibility issues with PAE mode.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT) that help provide increased levels of server scalability:

**Hardware scalability**

By choosing from products in the Windows Server 2003 family, you can take advantage of the growing number of competitively priced multiprocessor computers.

**I2O support**

I2O (also called Intelligent Input/Output) architecture provides for higher I/O performance on your servers by offloading certain I/O operations to a secondary processor. I2O improves I/O performance in high-bandwidth applications, such as networked video, groupware, and client/server processing.

**Symmetric multiprocessing enhancements**

The Windows Server 2003 family supports single or multiple CPUs that conform to the symmetric multiprocessing (SMP) standard. Using SMP, the operating system can run threads on any available processor, which makes it possible for applications to use multiple processors when additional processing power is required to increase the capability of a system. New features include SMP locking performance, improved registry performance, and increased Terminal server sessions.

**Network Load Balancing**

Previously known as Windows NT Load Balancing Service (WLBS), Network Load Balancing distributes incoming TCP/IP traffic between multiple servers. Your clustered applications, especially Web server applications, can handle more traffic and provide higher availability and faster response times.

**Server clusters (Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition only)**

Server clusters provide high availability, scalability, and manageability for important resources and applications. Multiple servers (nodes) in a cluster remain in constant communication. If one of the nodes in the cluster is unavailable due to failure

or maintenance, another node immediately begins providing service (a process known as failover). Users who access the cluster are constantly connected to server-based resources. For more information, see Windows Clustering.

**Enterprise Memory Architecture (Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, not including 64-bit versions)**

With Enterprise Memory Architecture, you can run applications that take advantage of large amounts of physical memory on Windows Server 2003, Enterprise Edition. Using applications that are written with the Address Windowing Extensions (AWE) API, you can map more physical memory into the applications' virtual address memory space for improved performance. For systems with 2 GB to 4 GB of memory, you can also use application memory tuning, also known as 4-gigabyte tuning (4GT). This enables you to provide up to 3 GB of your virtual address memory space to applications (by providing less memory to the operating system).

#### New and updated features since Windows 2000

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of server scalability:

**Internet Information Services 6.0**

Internet Information Services (IIS) 6.0 is a full-featured Web server that provides the foundation for the Windows Server 2003 family and existing Web applications and Web services. IIS 6.0 offers dedicated application mode, which runs all application code in an isolated environment. IIS 6.0 also supports Web gardens, in which equivalent processes on a computer each receive a share of the requests that are normally served by a single process, achieving better multiprocessor scalability.

**Note**

- By default, IIS is not installed with new installations of Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. You can add it by using the Configure Your server Wizard or by using Add or Remove Programs.

**Network Load Balancing enhancements**

Several improvements have been made to Network Load Balancing in this release:

- Network Load Balancing can now be bound to multiple network adapters, so that you can configure multiple independent clusters on each host. For more information about virtual clusters, see Virtual clusters.

- Network Load Balancing uses Internet Group Management Protocol (IGMP) support for limiting switch flooding, so that traffic intended for a Network Load Balancing cluster passes through only those ports that serve the cluster hosts and not through all switch ports.

- With Network Load Balancing Manager, you can create new Network Load Balancing clusters, and you can configure and manage clusters and all the cluster's hosts from a single remote or local computer.

# Server manageability

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server manageability**

This topic contains a brief overview of the features that support server manageability. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

Improved manageability in the Windows Server 2003 family makes it easy to deploy and maintain Windows applications and systems.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT) that help provide increased levels of server manageability:

### IntelliMirror

To help reduce costs, administrators need the highest levels of control over portable and desktop systems. IntelliMirror provides this control on client systems running Windows 2000 Professional or Windows XP Professional. You can use IntelliMirror to define policy settings based on business roles, group memberships, and locations. With these policy settings, Windows 2000 Professional desktops and Windows XP Professional desktops are automatically reconfigured to meet a specific user's requirements each time that user logs on to the network, regardless of where the user logs on. For more information about IntelliMirror, see IntelliMirror.

### Windows Management Instrumentation

Windows Management Instrumentation (WMI) is a scalable management infrastructure that is based on the Web-Based Enterprise Management (WBEM) initiative. Using WMI, which includes command-line and scripting capabilities, you can monitor, track, and control system events that are related to software applications, hardware components, and networks. WMI includes a uniform scripting application programming interface (API) that defines all managed objects under a common object framework based on the Common Information Model (CIM). WMI makes it possible for applications and scripts to discover what information is available for a system by enumerating the classes that are available. WMI can submit queries that filter requests for very specific information, and it can subscribe to WMI events based on your particular interests, rather than being limited to events that are predefined by the original developers. For more information about WMI, see WMI Overview.

### Active Directory Migration Tool

The Active Directory Migration Tool (ADMT) can help you migrate user accounts, groups, and computer accounts from Windows NT 4.0 domains to Active Directory domains (that is, domains in which the domain controllers run Windows 2000 or Windows Server 2003). For more information, see Active Directory support tools and *Designing and Deploying Directory and Security Services* at the Microsoft Windows Server System Web site.

### Remote Installation Services

With Remote Installation Services (RIS), you can create installation images of operating systems or even of complete computer configurations, including desktop settings and applications. You can then make these images available to users at client computers. The client computers must support remote booting with the Pre-Boot eXecution Environment (PXE) ROM, or else they must be started with a remote-startup floppy disk. For more information, see Remote Installation Services.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

### Manage Your Server

Manage Your Server provides a central location for you to set up various server roles and perform daily administrative tasks. For more information, see Configuring Roles for Your Server.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of server manageability:

### Remote Desktop for Administration

With Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode), you can administer a computer from virtually any computer on your network. Based on Terminal Services technology, Remote

Desktop for Administration is specifically designed for server management. For more information about this feature, see Remote Desktop for Administration.

**Remotely administered server support**

Remotely administered server support provides the ability to install and manage a computer that does not have a monitor, VGA display adaptor, keyboard, or mouse. This feature also includes Emergency Management Services, which provides a mechanism to manage the remotely administered server when the operating system is not functional, such as when the computer is restarted, when the server is unavailable on a network, or during use of Remote Installation Services (RIS).

**User State Migration Tool**

The User State Migration Tool (USMT) aids deployment by capturing and restoring user settings, files, and documents. Users do not have to reconfigure desktop settings for such things as e-mail servers, proxy servers, desktop color schemes, and desktop wallpaper.

**Emergency Management Services**

With Emergency Management Services, combined with the appropriate hardware, you can perform remote management and system recovery tasks, even when the server is not available through the standard remote administration tools and mechanisms. For more information, see Emergency Management Services.

**Resultant Set of Policy**

You can use Resultant Set of Policy (RSoP) to simulate and test policy settings that are applied to computers or users through Group Policy. RSoP is a query engine that polls existing policies--and also planned policies--and then reports the results of those queries. It polls existing policies based on site, domain, domain controller, and organizational unit. RSoP gathers this information from Windows Management Instrumentation (WMI). For more information, see Resultant Set of Policy.

**Windows Management Instrumentation improvements**

Windows Management Instrumentation (WMI) filtering enhances the Group Policy infrastructure by helping you to determine the effect of a Group Policy object (GPO).

**Remote Installation Services improvements**

Enhancements to Remote Installation Services (RIS) include support for both Windows Server 2003 family and Windows XP product installations, more control over answer file processing for an RIS installation, and access to network files in recovery modes.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Authorization Manager**

Authorization Manager is a new tool for authorization-related tasks. It is helpful for application developers who use authorization services in their applications, as well as for administrators who define groups and assign security principals to perform roles and tasks requiring appropriate user rights and permissions. For more information, see Authorization Manager.

# Performance and tuning features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Performance and tuning features**

This topic contains a brief overview of the performance and tuning features in the Windows Server 2003 family. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

New features in the Windows Server 2003 family provide for better performance and the ability to tune your server to your needs and environment.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of server performance:

**Winsock Direct**

Winsock Direct enables applications that use Winsock to perform faster and with less CPU overhead when they communicate across a system area network (SAN). Winsock Direct has the effect of streamlining communications between distributed components, if there is a SAN in place.

**I2O support**

I2O (also called Intelligent Input/Output) architecture allows for higher I/O performance on your servers by offloading certain I/O operations to a secondary processor. I2O improves I/O performance in high-bandwidth applications, such as networked video, groupware, and client/server processing.

**CPU throttling for Web applications or sites**

CPU throttling limits the amount of CPU time a Web application or site can use, which ensures that processor time is available to other Web sites or to non-Web applications.

**Process Accounting for Web applications or sites**

With Process Accounting, you can obtain information about how Web sites use CPU resources. This information can be crucial for making decisions about changing or improving hardware or applications. It can also be important when monitoring or diagnosing problems with scripts or applications.

**Predictable, end-to-end Quality of Service**

Quality of Service (QoS)-based services and protocols provide a guaranteed, end-to-end, express delivery system for traffic that is generated by QoS-enabled programs across a network or across multiple disparate networks. For more information, see Quality of Service (QoS).

**Network Load Balancing**

Previously known as Windows NT Load Balancing Service (WLBS), Network Load Balancing distributes incoming TCP/IP traffic between multiple servers. Your clustered applications, especially Web server applications, can handle more traffic and provide higher availability and faster response times.

**New performance tools**

Several new performance tools are available to help you diagnose and resolve memory, registry, driver, and other system problems.

**Disk Defragmenter**

Disk Defragmenter enhancements include Defrag.exe, a new command-line tool. Both Disk Defragmenter and Defrag.exe can now defragment the master file table (MFT) on NTFS volumes. For more information, see Disk Defragmenter.

**Event tracing enhancements**

Event tracing provides a mechanism for managing data that is useful to performance monitors, capacity planning tools, and other applications that analyze information relating to system resource utilization. New command-line tools control event trace log scheduling, collection, and analysis on local and remote computers.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of server performance:

**Performance Options**

Performance Options is a tool that consolidates processor, memory, and virtual memory settings, so that you can quickly and easily configure a computer for best performance.

**New command-line utilities for System Monitor**

New command-line tools are provided to control performance counter and event trace log scheduling, as well as collection and analysis on local and remote computers without using System Monitor (Sysmon). These command-line tools include **logman**, **relog**, **tracerpt**, and **typeperf**. For more information, see Logman, Relog, Tracerpt, and Typeperf.

**System Monitor enhancements**

The Windows Server 2003 family supports viewing data from multiple log files simultaneously and from data stored in a SQL database that has been collected using the Performance Logs and Alerts service. Selected data from a performance log file or SQL database can be saved to a new file for analysis at a later time. Display enhancements make it easy to select multiple counter objects and to display the data properties page for a selected counter directly from the list window.

**Performance Logs and Alerts improvements**

The Windows Server 2003 family supports log files greater than 1 GB in size and, with the new log file format, performance data can now be appended to an existing log file. You can log data directly to a SQL database through Open Database Connectivity (ODBC) data connections.

# Authentication and smart card support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication and smart card support**

This topic contains a brief overview of the authentication and smart card support features in this release. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The Windows Server 2003 family provides a wide variety of authentication technologies, including technologies for using smart cards. For more information about new security features, see New features in security.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for authentication and use of smart cards:

**Full support for the Kerberos V5 protocol**

With Kerberos V5 support, a fast, single logon process gives users the access they need to enterprise resources running on Windows 2000 or the Windows Server 2003 family of products. Support for Kerberos V5 includes additional benefits, such as mutual authentication (the client and server must both provide authentication) and delegated authentication (the user's credentials are tracked end to end). For more information, see Kerberos V5 authentication.

**Public key infrastructure, Certificate Services, and smart cards**

By using Certificate Services and certificate management tools, you can deploy your own public key infrastructure (PKI). With a PKI, you can implement standards-based technologies, such as smart card logon capabilities, client authentication (through Secure Sockets Layer (SSL) and Transport Layer Security (TSL)), secure e-mail, digital signatures, and secure connectivity (through Internet Protocol security (IPSec)). Using Certificate Services, you can set up and manage certification authorities (CAs) that issue and revoke X.509 v3 certificates. This means that you do not have to depend on commercial client authentication services, although you can integrate commercial client authentication into your public key infrastructure if you choose.

**Virtual private networking**

You can give users ready access to your organization's network even when they are out of the office, and reduce the cost of such access, by implementing a virtual private network (VPN). The VPN connection creates a secure tunnel across the Internet into the private network.

There are two types of VPN technology in the Windows Server 2003 family:

- Point-to-Point Tunneling Protocol (PPTP), which employs user-level Point-to-Point Protocol (PPP) authentication methods and Microsoft Point-to-Point Encryption (MPPE) for data encryption.
- Layer Two Tunneling Protocol (L2TP) with Internet Protocol security (IPSec). L2TP employs user-level PPP authentication methods and computer-level certificates with IPSec for data encryption.

On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Flexible, secure authentication and authorization**

Flexible and secure authentication and authorization services provide protection for data while minimizing barriers to doing business over the Internet. Active Directory supports multiple authentication protocols, such as Kerberos V5, Secure Sockets Layer (SSL) v3, and Transport Layer Security (TSL) using X.509 v3 certificates, and security groups that span domains efficiently.

**Internet Authentication Service**

Internet Authentication Service (IAS) provides you with a central point for managing authentication, authorization, accounting, and auditing of dial-up or virtual private network (VPN) users. IAS uses the Internet Engineering Task Force (IETF) protocol called Remote Authentication Dial-In User Service (RADIUS).

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition. In Windows Server 2003, Standard Edition, you can configure IAS with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. For more information about these limits, see Internet Authentication Service.

**Routing and Remote Access**

Routing and Remote Access replaces the Routing and Remote Access Service (RRAS) and Remote Access Service (RAS) features in Windows NT 4.0. Routing and Remote Access is a single, integrated service that terminates connections from either dial-up or virtual private network (VPN) clients or that provides routing (IP, IPX, and AppleTalk), or both. With Routing and Remote Access, your server can function as a remote access server, a VPN server, a gateway, or a branch-office router. For more information, see Routing and Remote Access.

#### New and updated features since Windows 2000

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for authentication and use of smart cards:

**Stored User Names and Passwords**

Stored User Names and Passwords stores user credentials, including passwords and X.509 v3 certificates, so that they can be associated with remote computers and automatically reused. This provides a consistent, single sign-on experience for users, including roaming users.

**Smart card access through Terminal Server**

With Terminal Server, you can log on from a Remote Desktop Connection to a Windows Server 2003 domain by using trusted X.509 v3 certificates that are stored on a smart card. For example, you can use a smart card to log on to Active Directory on a terminal server.

**Subordinate certification authority improvements**

Subordinate certification authorities (CA) now support qualified subordination, which allows public key infrastructure (PKI) administrators to use qualified subordinate CAs for more precise security and administrative purposes than standard CAs, such as having a qualified subordinate CA only issue certificates to subjects in certain domains that can only be used by specific software applications. For more information, see Qualified subordination.

# Additional security features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Additional security features**

This topic contains a brief overview of the additional security features in the Windows Server 2003 family. It is divided into three sections: New and updated features since Windows Server 2003 (without SP1), New and updated features since Windows NT 4.0, and New and updated features since Windows 2000.

For information about authentication and smart card support features, see Authentication and smart card support. For more information about new security features, see New features in security. For links to more information about the features in this release, see New Features.

The Windows Server 2003 family includes a variety of features for the creation of a strong, flexible security system, including file encryption and the security features that are built into Active Directory. For information about the kinds of authentication that are available, see Authentication and smart card support.

**New and updated features since Windows Server 2003 (without SP1)**

Windows Server 2003 operating systems with Service Pack 1 (SP1) offer the following improvements (compared to Windows Server 2003 without SP1) that help provide increased levels of support for security:

**Security Configuration Wizard**

The Security Configuration Wizard reduces the attack surface of computers that are running Windows Server 2003 with SP1. The wizard asks the user a series of questions designed to determine the functional requirements of the server. Any functionality that is not required by the roles being performed by the server is then disabled. In addition to being a fundamental security best practice, reducing the attack surface increases the diversity of your Windows landscape, and reduces the number of systems that need to be immediately updated if a security issue arises. For more information, see Windows Server TechCenter.

**Windows Firewall**

Windows Firewall is a host firewall technology. As a host firewall, Windows Firewall runs on each of your clients and servers. It provides protection from network attacks that pass through your perimeter network or originate inside your organization, such as Trojan horse attacks, port scanning attacks, and worms. Like many firewall technologies, Windows Firewall is a stateful firewall. As such, Windows Firewall inspects and filters all Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic. Unsolicited incoming traffic is dropped unless it is a response to a request by the host (solicited traffic), or it is specifically allowed (excepted traffic). You can specify excepted traffic according to port number, application name, or service name by configuring Windows Firewall settings. With the exception of some Internet Control Message Protocol (ICMP) messages, Windows Firewall allows all outgoing traffic.

Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

For more information, see Windows Server TechCenter.

**Digital Identity Management Service (DIMS) and Credential Roaming**

DIMS and credential roaming make it possible for users who log onto any domain-joined computer running Windows Server 2003 with SP1 to seamlessly and silently have all of their certificates and private keys available for applications and services. Administrators can use Group Policy to enable credential roaming, thereby automatically synchronizing and replicating user certificates and private keys with the Active Directory user object. For more information, see Windows Server TechCenter.

**Windows Server 2003 Privacy Information**

Windows Server 2003 with SP1 includes Group Policy settings that were not available in the original release of Windows Server 2003. You can find new settings designed to help protect your privacy in Group Policy under a key called Internet Communication Management. The settings under the Internet Communication Management key are designed to help you control the way Windows components communicate with the Internet. One of the settings new to Windows Server 2003 with SP1 is Restrict Internet communication, which you can use to control multiple privacy-related settings simultaneously. For more information about privacy, see *Using Windows Server 2003 in a Managed Environment* at the Microsoft TechNet Web site.

**Enabling Authentication for Terminal Services connections**

With Windows Server 2003 with SP1, you can configure Terminal Services connections to use Transport Layer Security (TLS) for server authentication and encryption. TLS is often referred to as Secure Sockets Layer/Transport Layer Security (SSL/TLS).

By default, TLS is not enabled when SP1 is installed. For more information, see *Terminal Services Security* at the Microsoft Web site.

#### New and updated features since Windows NT 4.0

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for security:

**Encrypting File System**

Encrypting File System (EFS) complements other access controls and provides an added level of protection for your data. EFS runs as an integrated system service, making it easy for you to manage, difficult to attack, and transparent to the user. For more information, see Encrypting File System.

**Internet Protocol security**

You can use Internet Protocol security (IPSec) to secure communications within an intranet and to create secure virtual private network (VPN) solutions across the Internet. IPSec was designed by the Internet Engineering Task Force (IETF), and it is an industry standard for encrypting TCP/IP traffic. For more information, see Internet Protocol Security (IPSec).

#### New and updated features since Windows 2000

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for security:

**TCP/UDP port ownership**

A new **netstat** command option displays the process that owns the Transmission Control Protocol/User Datagram Protocol (TCP/UDP) port. You can use this feature to configure secure servers, security audits, and performance improvements.

**Encrypting File System improvements**

Encrypting File System (EFS) has been improved in several ways. Stronger encryption algorithms are available with larger keys. Multiple users can be authorized to share encrypted files. Offline files can be encrypted through EFS, so that you can protect locally cached documents. For more information, see Encrypting File System.

**Software restriction policies**

With software restriction policies, you can protect your computer environment from untrusted code by identifying and specifying which applications are allowed to run. For more information, see Software Restriction Policies.

**Internet Protocol security monitoring improvements**

Internet Protocol security (IPSec) monitoring improvements include a Microsoft Management Console (MMC) snap-in that provides detailed information about IPSec policy configuration and active security states. This feature replaces the Ipsecmon.exe monitor program in Windows 2000. For more information, see Internet Protocol Security (IPSec).

# Remote management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote management**

This topic contains a brief overview of the remote management features available in the Windows Server 2003 family. It is divided into three sections: New and updated features since Windows Server 2003 (without SP1), New and updated features since Windows NT 4.0, and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The Windows Server 2003 family includes a variety of remote management features. For detailed information on how to use these features, see Remote Administration.

**New and updated features since Windows Server 2003 (without SP1)**

Windows Server 2003 operating systems with Service Pack 1 (SP1) offer the following improvements (compared to Windows Server 2003 without SP1) that help provide increased levels of support for remote management:

### Remote Installation Services

Users can now use the Remote Installation Preparation (RIPrep) Wizard to install an x64-based server image.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for remote management:

### Microsoft Management Console

You can use Microsoft Management Console (MMC) to arrange the administrative tools and processes that you need within a single interface. You can also delegate tasks to specific users by creating preconfigured MMCs that contain the tools you select. For more information, see Microsoft Management Console.

### Scripting with Windows Script Host

By using Windows Script Host (WSH), you can automate such actions as creating a shortcut and connecting to and disconnecting from a network server. WSH is language independent. You can write scripts in common scripting languages, such as Microsoft Visual Basic®, Scripting Edition, (VBScript) and JScript. For more information, see Windows Script Host.

### Remote Installation Services

With Remote Installation Services (RIS), you can create installation images of operating systems or even of complete computer configurations, including desktop settings and applications. You can then make these images available to users at client computers. The client computers must support remote booting with the Pre-Boot eXecution Environment (PXE) ROM, or else they must be started with a remote-startup floppy disk. For more information, see Remote Installation Services.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for remote management:

### Remote Desktop for Administration

With Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode), you can administer a computer from virtually any computer on your network. Based on Terminal Services technology, Remote Desktop for Administration is specifically designed for server management. You can use Remote Desktop for Administration to log on to the server remotely as though you were logged on to it locally. Additionally, Remote Desktop for Administration now supports greater color depth and screen resolution. For more information about this feature, see Remote Desktop for Administration.

### Remote Assistance

You can use Remote Assistance to remotely administer a computer. If you have an invitation, Remote Assistance is a convenient way for you to connect to a remote computer from a computer running Windows XP or any product in the Windows Server 2003 family. After you are connected, you can view the remote computer's screen and chat in real time. If the person requesting assistance allows you to, you can even use the mouse and keyboard to work on the remote computer. For more information, see Windows interface administrative tool reference A-Z: Remote Assistance.

This feature is not available on the Itanium-based versions of the Windows operating systems.

**Web Interface for Remote Administration (Windows Server 2003, Web Edition only)**

On Windows Server 2003, Web Edition, the Web Interface for Remote Administration is a Hypertext Markup Language (HTML)-based application that is used to configure and manage the server from a remote client. Individual servers, entire server farms, and multiple sites per server can be managed from a single remote workstation.

The Web Interface for Remote Administration provides a simple means of performing common Web server configuration tasks, including:

- Creating and deleting Web sites
- Configuring network settings
- Managing local user accounts
- Restarting the Web server

**Remote Installation Services improvements**

Enhancements to Remote Installation Services (RIS) include support for both Windows Server 2003 family and Windows XP product installations, more control over answer file processing for an RIS installation, and access to network files in recovery modes.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

# Maintenance and product updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maintenance and product updates**

This topic contains a brief overview of the maintenance and product updates features in the Windows Server 2003 family. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The Windows Server 2003 family offers several features that can help keep your system current with the latest system files and product updates.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help keep your system updated with the latest system files:

### Windows Update

Windows Update is an online extension of Windows that detects outdated Windows system files and replaces them with the most recent versions. Windows Update can significantly increase computer system stability over a long period. It can also reduce the workload on administrators and support staff who must maintain a productive operating environment. Windows Update requires an Internet connection. For more information, see Windows Update at the Microsoft Web site.

Windows Update includes Dynamic Update, which is available during Setup. If you have a working Internet connection on the computer on which you run Setup, you can use Dynamic Update to obtain the most up-to-date Setup files, including drivers and other files. For more information, see Using Dynamic Update for updated drivers and other setup files.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help keep your system updated with the latest system files:

### Windows Update enhancements

Windows Update now provides additional choices:

- With Windows Update Catalog, you can select updates that you plan to deploy later, which means you can control how and when the updates are deployed. For more information, see Windows Update Catalog at the Microsoft Web site.

- With Automatic Updates, which is built into Windows XP and products in the Windows Server 2003 family, the operating system can automatically download and install critical updates from the Windows Update site. For more information, see Automatic Updates.

- With Microsoft Windows Update Services, you can deploy a limited version of the Windows Update site inside your firewall. You can control how and when updates are deployed, and you can deliver updates to users through your intranet rather than having updates occur directly from the Internet. For more information, see Windows Update Services at the Microsoft Web site.

# Internationalization features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### Internationalization features

This topic contains a brief overview of the internationalization features in the Windows Server 2003 family of products. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

#### New and updated features since Windows NT 4.0

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for internationalization:

**Multilingual support**

The Windows Server 2003 family provides support for multiple languages on a single system. As a result of the system architecture, the operating system, utilities, and properly written applications can be translated and modified into local versions for hundreds of languages. Local administrators therefore do not need to be fluent in technical English.

**Multiregional support**

Support for additional locales and regional and language options is also included in this release, which makes it even easier to access relevant content.

#### New and updated features since Windows 2000

Many products in the Windows Server 2003 family offer the following improvements (in comparison to Windows 2000) that help provide increased support for internationalization. Some of these languages may not be available on certain products:

**New Phonetic Input Method Editor for Traditional Chinese**

The new Phonetic Input Method Editor (IME) is a new design, based on the latest state-of-the-art Chinese keyboard input technology. It is an update of the Microsoft Phonetic IME 98a CHT edition.

**Improved Japanese Input Method Editor**

The Japanese IME Version 8 receives alphabetical keyboard characters, converts them to Japanese characters, and then enters the result in the application. This new feature also supports the use of Japanese IME dictionaries over a network. IME saves user-defined words and conversion rules in a user dictionary, which can be edited on the user's local computer. (This dictionary was read-only in Windows 2000).

**Improved Korean Input Method Editor**

The Korean IME maps key sequences to Korean Jamo and then converts them into Korean syllables. The new Korean IME supports handwriting input for Hangul and Hanja.

**Improved Simplified Chinese Input Method Editor**

The MSPY3.0 IME provides an enlarged IME lexicon (approximately 100,000 words and phrases) and updated IME rules, an improved language model with higher IME accuracy and performance, and modeless typing (Chinese/English mixed input).

# Features for Directory, File, Print, and Collaboration Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Features for directory, file, print, and collaboration services**
The Windows Server 2003 family contains improved support for directory, file, print, and collaboration services. The topics in this section describe the features that support these server types.

Windows Server 2003, Web Edition, is not intended for use as a file or print server. Support for file and print services is limited on the Windows Server 2003, Web Edition, operating system.

For links to more information about the features in this release, see New Features.

This section covers:

- Features for Active Directory
- Features for collaboration
- Disk and file management features
- Printer and protocol support

# Features for Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Features for Active Directory**

This topic contains a brief overview of the features of the Active Directory directory service in the Windows Server 2003 family. It is divided into three sections: New and updated features since Windows Server 2003 (without SP1), New and updated features since Windows NT 4.0, and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features. For more information about Active Directory, see Active Directory. For a more detailed list of new Active Directory features, see New features for Active Directory.

Active Directory is an enterprise-class directory service that is scalable, built from the ground up using Internet-standard technologies, and fully integrated at the operating-system level. Active Directory simplifies administration and makes it easier for users to find resources. Active Directory provides a wide range of features and capabilities.

### Note

- Computers running the Windows Server 2003, Web Edition operating system cannot function as domain controllers. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

**New and updated features since Windows Server 2003 (without SP1)**

Windows Server 2003 operating systems with Service Pack 1 (SP1) offer the following improvements (compared to Windows Server 2003 without SP1) that help provide increased levels of support for Active Directory:

### Improved replication and DNS diagnostic testing capabilities

Active Directory® has been updated to provide automatic directory service backup reminders, improved protection against replication errors, improvements to Install from Media (to facilitate adding new domain controllers that are DNS servers), improved DNS diagnostic testing capabilities, and access to a new platform for running domain controllers in virtual machines under Microsoft® Virtual Server 2005. For more information about these and other new Active Directory features and enhancements, see New features for Active Directory.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for Active Directory:

### Simplified user and network-resource management

Using Active Directory, you can build hierarchical information structures that make it easier for you to control administrative credentials and other security settings and that make it easier for your users to locate network resources, such as files and printers.

### Flexible, secure authentication and authorization

Flexible and secure authentication and authorization services provide protection for data while minimizing barriers to doing business over the Internet. Active Directory supports multiple authentication protocols, such as the Kerberos V5 protocol, Secure Sockets Layer (SSL) v3, and Transport Layer Security (TLS) using X.509 v3 certificates, and security groups that span domains efficiently.

### Directory consolidation

You can organize and simplify the management of users, computers, applications, and devices, and make it easier for users to find the information they need. You can take advantage of synchronization support through Lightweight Directory Access Protocol (LDAP)-based interfaces, and you can work with directory consolidation requirements specific to your applications.

### Directory-enabled applications and infrastructure

Active Directory features make it easier for you to configure and manage applications and other directory-enabled network components.

### Scalability without complexity

Active Directory scales to millions of objects per domain and uses indexing technology and advanced replication techniques to speed performance.

### Use of Internet standards

Active Directory provides access through LDAP and uses a Domain Name System (DNS)-based namespace.

**A powerful development environment**

Active Directory provides a powerful development environment through Active Directory Service Interfaces (ADSI), which provides an object-oriented interface to Active Directory. ADSI makes it easy for programmers and administrators to create directory programs by using high-level tools such as Microsoft Visual Basic, Java, C, or Visual C++, without having to worry about the underlying differences between the different namespaces. For more information, see Programming interfaces.

**Replication and trust monitoring**

Active Directory provides Windows Management Instrumentation (WMI) classes to monitor whether domain controllers are successfully replicating Active Directory information and whether trusts are functioning properly.

**Message Queuing distribution lists**

Message Queuing (also known as MSMQ) enables you to send messages to distribution lists that are hosted in Active Directory.

#### New and updated features since Windows 2000

The Windows Server 2003 family offers several improvements (in comparison to Windows 2000) that help provide increased levels of support for and better management of Active Directory. For a list of the Active Directory features that are new in this release, see New features for Active Directory.

# Features for collaboration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Features for collaboration**

This topic contains a brief overview of the collaboration features, including multimedia support, in the Windows Server 2003 family. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The Windows Server 2003 family improves employee communication and collaboration with team collaboration tools, as well as secure real-time communications, mobile data, and streaming media on a wide range of connected devices. Because servers are continuously available, employees can access information anywhere, anytime, and on any device, improving productivity across your organization.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for collaboration:

### Windows Media Services

Windows Media Services provides support for the distribution of streaming audio and video over corporate intranets and the Internet. In Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, Windows Media Services delivers advanced streaming functionality, such as multicasting, wireless network support, internet authentication, server plug-ins, and Cache/Proxy APIs.

The core components of Windows Media Services were developed using Component Object Model (COM), creating a flexible architecture that is easily customized for specific applications. Many features are implemented through the use of plug-ins. With this platform, developers can easily extend Windows Media Services functionality and integrate with existing systems and solutions, such as storage systems, billing applications, and logging applications.

Windows Media Services is not included with Windows Server 2003, Web Edition or the Itanium-based versions of the Windows Server 2003 family.

### Predictable, end-to-end Quality of Service

Quality of Service (QoS)-based services and protocols provide a guaranteed, end-to-end, express delivery system for traffic that is generated by QoS-enabled programs across a network or across multiple disparate networks. For more information, see Quality of Service (QoS).

### Asynchronous transfer mode

Asynchronous transfer mode (ATM) is a high-speed, connection-oriented protocol designed to transport multiple types of traffic across a network. It is applicable to both local area networks (LANs) and wide area networks (WANs). Using ATM, your network can simultaneously transport a wide variety of network traffic: voice, data, image, and video.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for collaboration:

### Microsoft .NET Framework

Microsoft .NET Framework provides the most comprehensive and powerful means to build, deploy, and manage XML Web services. XML Web services allow applications to communicate and share data over the Internet, regardless of operating system or programming language. The Windows Server 2003 family is the foundation of the Microsoft .NET vision, delivering an unprecedented .NET computing foundation to organizations of all sizes. For more information, see Microsoft .NET Framework.

### Audio Acceleration on the remote access server

The Audio Acceleration feature prioritizes multimedia traffic over other data streams, which provides better performance, especially for lower-speed networks. Audio Acceleration on the remote access server ensures the quality of streaming audio and video communication.

### NetMeeting enhancements

You can use NetMeeting to communicate with another person by using Microsoft Exchange multicast conferencing, which

allows you to communicate and collaborate in real time over intranets and the Internet. NetMeeting also offers a platform for real-time communication, so that software vendors can integrate Internet conferencing into their applications and services. For more information, see Using NetMeeting.

This feature is not available on the Itanium-based versions of the Windows operating systems.

# Disk and file management features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk and file management features**

This topic contains a brief overview of the disk and file management features in the Windows Server 2003 family. It is divided into three sections: New and updated features since Windows Server 2003 (without SP1), New and updated features since Windows NT 4.0, and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The disk, volume, and file management features in this release make it a flexible platform for file servers and make it easier for users to locate and access information.

**New and updated features since Windows Server 2003 (without SP1)**

Windows Server 2003 operating systems with Service Pack 1 (SP1) offer the following improvements (compared to Windows Server 2003 without SP1) that help provide increased levels of support for disk and file management:

**Distributed File System (DFS)**

You can use the /move parameter in Dfscmd.exe to rename a DFS link.

**GUID partition table (GPT) disks**

The GUID partition table (GPT) disk partitioning style supports volumes up to 18 exabytes in size and up to 128 partitions per disk. In comparison, the master boot record (MBR) disk partitioning style supports volumes up to 2 terabytes in size and up to four primary partitions per disk (or three primary partitions, one extended partition, and unlimited logical drives). Unlike MBR partitioned disks, data that is critical to platform operation is located in partitions instead of unpartitioned or hidden sectors. In addition, GPT partitioned disks have redundant primary and backup partition tables for improved partition data structure integrity.

For more information about using GPT and MBR disks on Itanium-based computers, on x64-based computers, and on x86-based computers running Windows Server 2003 Service Pack 1 (SP1), see GUID partition table.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for disk and file management:

**Disk Management**

Disk Management is a system utility for managing hard disks and the volumes, or partitions, that they contain. With Disk Management, you can initialize disks; create volumes; format volumes with FAT, FAT32, or NTFS file systems; and create fault-tolerant disk systems. Disk Management enables you to perform most disk-related tasks without shutting down the system or interrupting users; most configuration changes take effect immediately. For more information, see Disk Management.

**Disk quota support**

You can use disk quotas on volumes formatted with the NTFS file system to monitor and limit the amount of disk space available to individual users. You can define the responses that result when users exceed your specified thresholds. For more information, see Disk Quotas.

**Distributed File System**

With Distributed File System (DFS), you can create a single directory tree that includes multiple file servers and file shares in a group, division, or enterprise. This allows users to easily find files or folders distributed across the network. DFS shares can also be published as Volume Objects in Active Directory. For more information, see Distributed File System (DFS).

**Distributed Link Tracking**

Using Distributed Link Tracking and NTFS, you can enable client applications to track linked sources that have been moved. For example, a client application will always be able to access a linked database, even if the database location changes.

**Distributed Authoring and Versioning**

Using Distributed Authoring and Versioning, remote authors can edit, move, or delete files, file properties, directories, and directory properties on your server over a Hypertext Transfer Protocol (HTTP) connection.

**High-performance content indexing**

You can use Indexing Service to provide a fast, easy, and secure way for users to search for information locally or on the

network. Users can search in files in different formats and languages, either through the **Search** command on the **Start** menu or through Hypertext Markup Language (HTML) pages that they view in a browser.

**Encrypting File System**

Encrypting File System (EFS) complements other access controls and provides an added level of protection for your data. EFS runs as an integrated system service, making it easy to manage, difficult to attack, and transparent to the user. For more information, see Encrypting File System.

**Removable Storage and Remote Storage**

Removable Storage makes it easy to track your removable storage media (tapes and optical discs) and to manage the hardware libraries, such as changers and jukeboxes, that contain them. Remote Storage uses criteria that you specify to automatically copy little-used files to removable media. If hard-disk space drops below specified levels, Remote Storage removes the (cached) file content from the disk. If the file is needed later, the content is automatically recalled from storage. Because removable optical discs and tapes are less expensive per megabyte (MB) than hard disks, Removable Storage and Remote Storage can decrease your costs. For more information about Removable Storage, see Removable Storage. For more information about Remote Storage, see Remote Storage.

Remote Storage is not available on Windows Server 2003, Web Edition or on Windows Server 2003, Standard Edition. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

#### New and updated features since Windows 2000

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for disk and file management:

**Remote document sharing**

A new Web Distributed Authoring and Versioning (WebDAV) redirector supports the WebDAV protocol for remote document sharing over HTTP. The WebDAV redirector supports the use of existing applications and allows file sharing across the Internet (through firewalls, routers, and so forth) to HTTP servers. With the WebDAV redirector, you can host Web repositories that can be accessed from anywhere on the Internet and with any application, not just those that interpret Web publishing protocols.

**Command-line support for disk management**

With the new **DiskPart** command, you can perform many disk management tasks from the command line. For more information, see Managing disks and volumes from the command line and DiskPart.

**GUID partition table (Itanium-based versions only)**

Windows XP 64-bit Edition (Itanium) and the Itanium-based versions of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, support a new disk partitioning style, the GUID partition table (GPT). Unlike master boot record (MBR) partitioned disks, data that is critical to platform operation is located in partitions instead of in unpartitioned or hidden sectors. In addition, GPT partitioned disks have redundant primary and backup partition tables for improved partition data structure integrity. For more information, see GUID partition table.

**Shadow copies of shared folders**

With shadow copies of shared folders, you can open, view, and save point-in-time versions of files that are stored in shared network folders. For example, if you accidentally delete a file from a network share, you can use the shadow copy feature to retrieve a previous version of the file. For more information, see Shadow Copies for Shared Folders.

**File management improvements**

With improved folder options, you can manage files more effectively. For example, you can customize your folder view to show common tasks for files and folders, which makes it easier to send files through e-mail or give a user access to a place on a Web server. Increased thumbnail size makes it easier to identify documents.

**Support for storage area networks (SANs)**

For better interoperability in storage area network (SAN) environments, on Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, volumes on new basic disks are not automatically mounted and assigned drive letters by default when added to the system. You must manually mount basic volumes and assign the drive letters using Disk Management or the DiskPart and mountvol commands. Dynamic disks, removable media devices (such as Zip or Jaz drives), and optical discs (such as CD-ROM or DVD-RAM) are always automatically mounted by the system when they are detected.

# Printer and protocol support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

### Printer and protocol support

This topic contains a brief overview of the printer and protocol support features in the Windows Server 2003 family. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

The Windows Server 2003 family provides a variety of printer and protocol support features for print servers. For links to more information about the features in this release, see New Features.

### New and updated features since Windows NT 4.0

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for print servers:

**Printers published in Active Directory**

Users can quickly search for and locate the most convenient printing resources through an intuitive user interface.

**Broad device and protocol support**

You can choose from among thousands of different printers.

**Plug and Play**

With Plug and Play, a combination of hardware and software support, the server can recognize and adapt to hardware configuration changes automatically, without your intervention and without restarting.

**Internet Printing Protocol**

With Internet Printing Protocol (IPP), users can print directly to a Uniform Resource Locator (URL) over an intranet or the Internet. In addition, servers that run products in the Windows Server 2003 family can automatically generate print job information in Hypertext Markup Language (HTML) format, so that users can view it in their browsers.

**Printer management features**

You can monitor the operation of local and remote printers using System Monitor. You can set counters for a variety of criteria, such as Bytes Printed/Second, Job Errors, and Total Pages Printed.

**Network printing features**

You can remotely administer a print server running Windows Server 2003 by using Microsoft Management Console (MMC) from anywhere on the network. In addition, if the client computer is running Windows 2000 or Windows XP, you do not have to install a printer driver for that client computer to use a printer. The driver is downloaded automatically when the client connects to the print server. This simplifies printing across a network.

### New and updated features since Windows 2000

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for print servers:

**New command-line support for printing**

With new command-line utilities, you can manage printers and print queues by using a Windows Management Instrumentation (WMI) provider. You can also develop scripts or batch files to manage printers more efficiently.

**Easier printer driver installation on server clusters (Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition only)**

In Windows 2000, each printer driver had to be installed manually on every node of a server cluster. With Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, when you install a printer driver on a cluster, the driver is automatically propagated to all nodes.

# Features for Internet, Application, and Networking Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Features for Internet, application, and networking services**
The Windows Server 2003 family includes a variety of Internet capabilities, support for applications and application development, and networking and communications capabilities.

For links to more information about the features in this release, see New Features.

This section covers:

- Microsoft .NET Framework
- Internet and e-mail services and features
- Application and Development Support
- Network protocols and technologies

# Microsoft .NET Framework

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft .NET Framework**

Microsoft .NET Framework provides the most comprehensive and powerful means to build, deploy, and manage the broad range of Web services that are made possible by XML and related technologies. For more information, see the .NET Framework page at the Microsoft Web site.

**XML Web services**

XML Web services allow applications to communicate and share data over the Internet, regardless of operating system or programming language. The Windows Server 2003 family is the foundation of the Microsoft .NET vision, delivering an unprecedented computing foundation to organizations of all sizes.

**.NET Framework tools**

Using the tools that are provided with .NET Framework, application developers can easily build Windows applications that blend XML Web services with traditional applications, improving communication, collaboration, and connectivity inside and outside their organizations.

**Other .NET technologies**

Other .NET technologies that are included in the Windows server 2003 family are the integration of Microsoft .NET Passport with Active Directory, native support for Simple Object Access Protocol (SOAP) as the foundation for XML message passing in Message Queuing, and integration of component architecture with XML Web services.

Universal Description, Discovery, and Integration (UDDI) Services is also included with some products in the Windows server 2003 family. UDDI Services is not included with Windows Server 2003, Web Edition. In addition, Windows Server 2003, Standard Edition, supports only stand-alone installations of UDDI Services. Distributed installation support is available with Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. In a stand-alone installation of UDDI Services, both the UDDI Web server component and the UDDI database component are installed onto a single server. In a distributed installation, UDDI components are distributed across multiple servers.

For more information about Message Queuing, see Using Message Queuing. For more information about SOAP, see Network protocols and technologies. For more information about Internet-related services, see Internet and e-mail services and features. For more information about UDDI Services, see UDDI Services overview.

**Built on industry standards**

The Windows Server 2003 family is built on industry standards so that you can extend existing applications and quickly develop new ones.

The Microsoft .NET Framework is built into all products in the Windows Server 2003 family except for the 64-bit versions.

The 32-bit version of .NET Framework 1.1 is supported by 64-bit versions of Windows Server operating systems and can be installed for 32-bit applications running on 64-bit versions of Windows operating systems (WOW64).

To install .NET Framework 1.1, see "Microsoft .NET Framework Version 1.1 Redistributable Package" at the Microsoft Download Center Web site. To install the service pack, see "Microsoft .NET Framework 1.1 Service Pack 1" at the Microsoft Download Center Web site.

**Important performance and scalability information**

If a 32-bit version of the .NET Framework is installed on the 64-bit Itanium-based systems, applications created for the 32-bit version of the Framework will bind and run against it. Therefore, applications created with the Microsoft .NET Framework Version 1.1 run only in a 32-bit process within the WOW64 compatibility environment. Due to the design of x86 emulation for the Itanium processor family and the WOW64 subsystem, applications are restricted to execution on one processor.

Single processor execution and x86 emulation reduce the performance and scalability of 32-bit .NET Framework applications running on Itanium-based systems. We recommend that applications using the .NET Framework Version 1.1 be used for interactive, client applications. With .NET Framework applications that demand high performance and scalability, such as high-load ASP.NET applications, we do not recommend this environment. We suggest that application developers instead evaluate the .NET Framework Version 2.0, which includes 64-bit support for Itanium-based systems, for increased performance and scalability.

# Internet and e-mail services and features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet and e-mail services and features**

This topic contains a brief overview of the Internet services and features in this release. It is divided into three sections: New and updated features since Windows Server 2003 (without SP1), New and updated features since Windows NT 4.0, and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The Windows Server 2003 family includes improved support for a wide variety of Internet capabilities.

**New and updated features since Windows Server 2003 (without SP1)**

Windows Server 2003 operating systems with Service Pack 1 (SP1) offer the following improvements (compared to Windows Server 2003 without SP1) that help provide increased levels of support for Internet services:

### Group Policy additions under "Internet Communication Management"

Windows Server 2003 with SP1 contains Group Policy settings that were not available in the original release of Windows Server 2003. Some of these settings can be found in Group Policy under the key called Internet Communication Management. This key is located within GPMC in Computer Configuration\Administrative Templates\System and in User Configuration\Administrative Templates\System.

The settings under the Internet Communication Management key are designed to help you control the way components in Windows Server 2003 with SP1 communicate with the Internet.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for Internet services:

### Active Server Pages

Active Server Pages (ASP) is a server-side scripting environment for creating dynamic, interactive Web server applications. With ASP, developers can combine Hypertext Markup Language (HTML) pages, script commands, and Component Object Model (COM) components as needed to create flexible Web-based applications.

### Windows Media Services

Using Windows Media Services, you can deliver high-quality streaming multimedia to users on the Internet and intranets.

Windows Media Services is not included with Windows Server 2003, Web Edition or the Itanium-based versions of the Windows Server 2003 family.

### Web Distributed Authoring and Versioning

Using Web Distributed Authoring and Versioning (WebDAV), remote authors can use a Hypertext Transfer Protocol (HTTP) connection to work with files. For more information, see Disk and file management features.

### Application recycling

Application recycling enhances the reliability and health of servers. You can refresh applications and prevent possible resource leaks, which increases the reliability and availability of your servers.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for Internet services:

### Microsoft .NET Framework components

Microsoft .NET Framework offers a wide range of possibilities for building, deploying, and managing Web services. For more information, see Microsoft .NET Framework.

### Internet Information Services 6.0

Internet Information Services (IIS) 6.0 is a full-featured Web server that has been optimized for running Web applications and Web services in a hosting environment. Many new features have been included in IIS 6.0 to enhance reliability, security, manageability, and performance.

**Note**

- By default, IIS is not installed with new installations of Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. You can add it by using the Configure Your Server Wizard or by using Add or Remove Programs.

**Universal Description, Discovery, and Integration Services**

Universal Description, Discovery, and Integration (UDDI) is an industry specification for publishing and locating information about Web services. Some products in the Windows Server 2003 family include UDDI Services, a Web service that provides UDDI capabilities for use within an enterprise or across organizations. UDDI Services is a standards-based XML Web service designed for enterprise developers to efficiently discover, share, and reuse Web services directly through their development tools.

UDDI Services is not included with Windows Server 2003, Web Edition. In addition, Windows Server 2003, Standard Edition, supports only stand-alone installations of UDDI Services. Distributed installation support is available with Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. In a stand-alone installation of UDDI Services, both the UDDI Web server component and the UDDI database component are installed onto a single server. In a distributed installation, UDDI components are distributed across multiple servers.

For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

**E-mail services (POP3, SMTP)**

The POP3 service provides e-mail transfer and retrieval services. Administrators can use the POP3 service to store and manage e-mail accounts on the mail server. Simple Mail Transfer Protocol (SMTP) supports the transfer of e-mail between servers. For more information, see POP3 Service and Simple Mail Transfer Protocol (SMTP) service.

**Configuration Database WMI Provider**

The Configuration Database WMI Provider for IIS facilitates cross-application administration of servers that are running IIS. Windows Management Instrumentation (WMI) has a rich schema and associative approach to modeling application configuration, which makes it possible to write scripts that manage all aspects of a Web application, including system configuration, application files, and databases.

# Application and Development Support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application and development support**
For a table that relates the types of support listed in this topic to other types of support in this release, see Choosing the features that are most useful for your server.

This section covers:

- Support for existing systems
- Application support
- Management and deployment
- Terminal Server features
- Development support

# Support for existing systems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Support for existing systems**

This topic contains a brief overview of the interoperability features in the Windows Server 2003 family. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The Windows Server 2003 family supports Windows NT 4.0, Windows 2000, UNIX, and Macintosh, as well as interoperability with various client operating systems. You can introduce products in the Windows Server 2003 family incrementally into a network environment, because migration paths are provided from any number of existing systems, devices, and applications.

This release includes the following design features to ensure that it is compatible with as many systems and applications as possible.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for existing systems:

### Industry standards

The Windows Server 2003 family supports the latest in industry standards, including:

- XML
- Simple Object Access Protocol (SOAP)
- Lightweight Directory Access Protocol (LDAP)
- DNS
- Active Directory Service Interfaces (ADSI)
- Kerberos V5 authentication protocol
- Universal Description, Discovery, and Integration (UDDI)

  UDDI Services is not included with Windows Server 2003, Web Edition. In addition, Windows Server 2003, Standard Edition, supports only stand-alone installations of UDDI Services. Distributed installation support is available with Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. In a stand-alone installation of UDDI Services, both the UDDI Web server component and the UDDI database component are installed onto a single server. In a distributed installation, UDDI components are distributed across multiple servers.

### Support for existing server applications

Existing server applications, such as Microsoft Exchange Server and Microsoft SQL Server, are supported on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. For full support, you may need to install the latest product updates. For more information, see the Windows Update site at the Microsoft Web site. Also, see the Application Center at the Microsoft Web site.

### Integration with existing operating systems

Products in the Windows Server 2003 family integrate seamlessly with existing operating systems and contain support for earlier Windows operating systems, as well as for other popular operating systems. The following features are included in this release:

- Interoperability with Windows NT Server 4.0 and Windows 2000 Server.
- Support for clients running a variety of operating systems, including Windows 3.x, Windows 95, Windows 98,

Windows Millennium Edition, Windows NT Workstation 4.0, Windows 2000 Professional, and Windows XP.

- UNIX interoperability, provided through an add-on pack called Windows Services for UNIX. The add-on pack provides support for Network File System (NFS), password synchronization, a UNIX command shell, and a collection of UNIX utilities. The add-on pack also supports Network Information Service (NIS). Telnet services (for remote access and administration) are also included.

- File Server for Macintosh, allowing Macintosh clients to use the TCP/IP protocol (AFP over IP) to share files and to access shares on a server running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition.

  This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for existing systems:

### Compatibility Mode

Compatibility Mode provides an environment for running applications that more closely reflects the behavior of Windows 95, Windows 98, Windows NT 4.0, or Windows 2000 operating systems. These modes resolve several of the most common issues that prevent older applications from working correctly. You can also use this technology to create solutions for your own applications.

### Program Compatibility Wizard

You can use the Program Compatibility Wizard to detect and test compatibility settings. For more information, see Troubleshooting compatibility issues.

# Application support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application support**

This topic contains a brief overview of the application support features in the Windows Server 2003 family. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The Windows Server 2003 family provides support for a wide range of turnkey independent software vendor (ISV) solutions and a comprehensive set of Internet application services for rapid development of custom applications.

**Note**

- The Windows Server 2003, Web Edition operating system is intended for Web application services and not as a general application server.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for applications:

**Widest range of turnkey applications**

You can run a wide and varied range of applications on the Windows Server 2003 family of products. For the most recent information on compatible applications, see Software compatibility information in the Windows Catalog at the Microsoft Web site.

**Comprehensive Internet application services**

Using the integrated Web, component management, transaction processing, and message queuing services in this release, programmers can rapidly develop and deploy scalable, component-based applications.

**Microsoft Active Accessibility**

Microsoft Active Accessibility is an application programming interface (API) that makes it easy for developers to build applications that can be customized for people with varying degrees of disabilities. Because Active Accessibility reduces the need for application-specific code, it reduces overall maintenance costs and allows software developers to innovate in their user interface without sacrificing compatibility with accessibility aids. For more information, see the Accessibility page at the Microsoft Web site.

**TAPI 3.1**

TAPI 3.1 unifies Internet Protocol (IP) and traditional telephony to enable developers to create computer telephony applications that work as effectively over the Internet or an intranet as over the traditional telephone network.

**Transaction services**

The integrated transaction services in the Windows Server 2003 family provide simple building blocks that can reliably and efficiently execute complex transactions across widespread distributed networks.

**Message queuing**

Integrated message queuing functionality helps developers build and deploy applications that run more reliably over networks, including the Internet. These applications can interoperate with applications running on different platforms, such as mainframes and UNIX-based systems.

**Component Services**

Component Services is a set of services based on extensions of the Component Object Model (COM) and on Microsoft Transaction Server (an earlier release of a component-based transaction processing system). Component Services provides threading and security, transaction management, object pooling, queued components, and application administration and packaging.

**Job objects**

Products in the Windows Server 2003 family contain an extension to the process model called a job. Job objects are nameable, securable, sharable objects that control attributes of the processes that are associated with them. A job object's basic function is to allow groups of processes to be managed and manipulated as a unit. For example, a Web application can

use a job object to account for CPU usage.

**Host data integration**

Products in the Windows Server 2003 family provide an excellent platform for running database applications by including native support for Open Database Connectivity (ODBC), Oracle, and Microsoft SQL Server.

#### New and updated features since Windows 2000

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for applications:

**Compatibility Mode**

Compatibility Mode provides an environment for running applications that more closely reflects the behavior of Windows 95, Windows 98, Windows NT 4.0, or Windows 2000 operating systems. These modes resolve several of the most common issues that prohibit older applications from working correctly. You can also use this technology to create solutions for your own applications. For more information, see Accessing information about compatibility.

**Web gardens in Internet Information Services**

In Internet Information Services (IIS), a Web garden is an application pool that has multiple worker processes serving the requests that are routed to that pool. A Web garden can be configured so that the processes serving it have affinity for particular processors on a multiprocessor computer. Using Web gardens, Web applications have increased scalability. If one worker process experiences a problem with an Active Server Page (ASP), the other worker processes can still serve content and take over for the locked worker process.

# Management and deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Management and deployment**

This topic contains a brief overview of the application management and deployment features in the Windows Server 2003 family. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The following features help simplify the deployment of applications on desktop and portable systems. Terminal Server provides additional options and features that simplify deployment. For more information, see Terminal Server features.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for management and deployment:

**Change and configuration management with Group Policy**

You can use policy settings to define the permitted actions and settings for users and computers. In contrast with local policy, Group Policy can be used to set policy settings that apply across a given site, domain, or organizational unit in Active Directory. For more information about Group Policy, see Group Policy (pre-GPMC). The Windows Server 2003 family includes more than 200 new policy settings.

**Directory consolidation in Active Directory**

With Active Directory, you can organize and simplify the management of users, computers, applications, and devices and make it easier for users to find the information they need. You can take advantage of synchronization support through Lightweight Directory Access protocol (LDAP)-based interfaces and work with directory consolidation requirements that are specific to your applications.

**IntelliMirror**

To help reduce costs, administrators need the highest levels of control over portable and desktop systems. IntelliMirror provides this control on client systems running Windows 2000 Professional or Windows XP Professional. You can use IntelliMirror to define policy settings based on business roles, group memberships, and locations. With these policy settings, Windows 2000 Professional desktops and Windows XP Professional desktops are automatically reconfigured to meet a specific user's requirements each time that user logs on to the network, regardless of where the user logs on. For more information about IntelliMirror, see IntelliMirror.

**Remote Installation Services**

With Remote Installation Services (RIS), you can create installation images of operating systems or even of complete computer configurations, including desktop settings and applications. You can then make these images available to users at client computers. The client computers must support remote booting with the Pre-Boot eXecution Environment (PXE) ROM, or else they must be started with a remote-startup floppy disk. For more information, see Remote Installation Services.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Windows Management Instrumentation**

Windows Management Instrumentation (WMI) is a scalable management infrastructure that is based on the Web-Based Enterprise Management (WBEM) initiative. Using WMI, which includes command-line and scripting capabilities, you can monitor, track, and control system events that are related to software applications, hardware components, and networks. WMI is both standards-based (coming out of the Common Information Model (CIM)) and flexible, allowing you to customize the way that information is gathered and filtered. For more information about WMI, see WMI Overview.

**Side-by-side component sharing**

Side-by-side sharing minimizes application fragility by enabling multiple versions of the same Component Object Model (COM) or Win32 component, such as a .dll file, to run at the same time in memory. This means that applications can use the specific components that they were designed for and tested with, even if another application requires a different version of the same component. This allows developers to build more reliable applications, because developers can choose the version of the component that they want to use for their application, independent of the other applications on the system.

**Add or Remove programs**

With Add or Remove programs, you can easily add or remove programs or components and keep track of the list of programs that are currently installed.

**Windows Installer**

With Windows Installer, you can customize your deployments of software components, as well as monitoring file resiliency and maintaining basic system recovery with rollbacks. For more information, see Windows Installer.

**Application management**

With application management, you can make changes in controlled ways to applications, without recompiling or rereleasing.

#### New and updated features since Windows 2000

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for management and deployment:

**Remotely administered server support**

Remotely administered server support provides the ability to install and manage a computer that does not have a monitor, VGA display adaptor, keyboard, or mouse. This feature also includes Emergency Management Services, which provides a mechanism to manage the remotely administered server when the operating system is not functional, such as when the computer is restarted, when the server is unavailable on a network, or during use of Remote Installation Services (RIS).

**Windows Management Instrumentation improvements**

Windows Management Instrumentation (WMI) filtering enhances the Group Policy infrastructure by helping you to determine the effect of a Group Policy object (GPO).

**Scripting support**

The Windows Server 2003 family has better scripting support, including new command-line tools and improvements to Windows Script Host. For a list of new command-line tools, see New command-line tools. For more information about the improvement to Windows Script Host, see Windows Script Host overview.

**Resultant Set of Policy**

You can use Resultant Set of Policy (RSoP) to simulate and test policy settings that are applied to computers or users through Group Policy. RSoP is a query engine that polls existing policy settings--and also planned policy settings--and then reports the results of those queries. It polls existing policy settings based on site, domain, domain controller, and organizational unit. RSoP gathers this information from Windows Management Instrumentation (WMI).

**Remote Installation Services improvements**

Enhancements to Remote Installation Services (RIS) include support for both Windows Server 2003 family and Windows XP product installations, more control over answer file processing for a Remote Installation Services installation, and access to network files in recovery modes.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

# Terminal Server features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Server features**

This topic contains a brief overview of the Terminal Server features in this release. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

Terminal Server provides remote computers with access to Windows-based programs running on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. With Terminal Server, you can provide a single point of installation that gives multiple users access to any computer that is running one of these products. Users can run programs, save files, and use network resources--all from a remote location--as if these resources were installed on their own computers.

### Notes

- Several Terminal Server and Remote Desktop features have been renamed or replaced in this release. For a table that lists the new names for features in Terminal Server, see New ways to do familiar Terminal Server tasks. For a table that lists the new names for Remote Desktop for Administration (formerly Terminal Services in Remote Administration mode), see New ways to do familiar Remote Desktop for Administration tasks.

- Remote Desktop for Administration is another part of the Terminal Services technology that is offered in this release. Remote administration of servers with Remote Desktop for Administration can be enabled on any computer running the Windows Server 2003 family of products, whether or not Terminal Server is installed. For more information, see Remote Desktop for Administration.

- Terminal Server (previously known as Terminal Services in Application Server mode) is not included on computers running the Windows Server 2003, Web Edition operating system. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT) that help provide increased levels of support for Terminal Server:

### Centralized deployment of programs

With Terminal Server, all program execution, data processing, and data storage occur on the server, which centralizes the deployment of programs. Terminal Server ensures that all clients can access the same version of a program. Software is installed only once on the server, rather than on every desktop throughout the organization, which reduces the costs that are associated with updating individual computers.

### Remote access to applications

Terminal Server efficiently sends the application display, rather than the underlying application data, from server to client. This means that you can give users access to data-heavy applications over lower-bandwidth connections (for example, a 28.8-Kbps dial-up connection) and provide higher performance than would be possible without Terminal Server.

### Single-application access

Terminal Server can provide users with access to a single published application when access to a full Windows XP desktop is not required.

### Terminal Services Manager

With Terminal Services Manager, you can view information about terminal servers in trusted domains. This information includes all sessions, users, and processes for each terminal server. You can also use this utility to perform various actions to manage the server.

### Remote control

Terminal Services provides built-in remote control for the Windows Server 2003 family.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for Terminal Server:

**Audio redirection**

Audio redirection enables sound reproduction on a client computer with any application that attempts to play a wave sound inside a Terminal Server session. This feature provides the capability for basic audio, for example, for attached .wav files in e-mail, documents, or simple streaming media.

**Group Policy integration**

Terminal Server is integrated with Group Policy, so that you can manage how a server that has Terminal Server enabled can be used by enforcing redirection capabilities, password access, and wallpaper settings.

**Resolution and color enhancements**

Terminal Services, the underlying technology for Terminal Server, now supports greater color depth and screen resolution. Using Remote Desktop Connection, color depth can be any setting from 256 colors to True Color, and resolution can be set from 640 × 480 to the highest level that is supported by the client video devices. Resolution and color limits must be enabled on both the server and the client.

# Development support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Development support**

This topic contains a brief overview of the development support features in the Windows Server 2003 family. It is divided into two sections: New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

Developers writing server-based applications can work with the tools that are available in Microsoft Visual Studio® and also with other development tools. The system architecture in the Windows Server 2003 family allows developers to design applications that are easily translated and modified into local versions.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide increased levels of support for developers:

### Visual Studio .NET

Visual Studio .NET is designed to address the major requirements of organizations creating modern enterprise solutions. The design themes for Visual Studio .NET include:

- A complete suite of tools for component-based development
- Enhanced life cycle productivity
- Enterprise database tools
- Enhanced support for teams
- Integration with application services

### Multilingual technology

Because of the system architecture in the Windows Server 2003 family, the operating system, utilities, and properly written applications can be translated and modified into local versions for hundreds of languages.

**New and updated features since Windows 2000**

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for developers:

### Automatic generation of link IDs

Link IDs are used by applications as they build new Active Directory structures (extensions to the Active Directory schema). If an application requests a new link ID that duplicates a link ID that is already in use, the schema extension fails to load. Automatic generation of link IDs avoids this problem by giving the application developer the option to request that the link ID be generated automatically.

### Background Intelligent Transfer Service

Background Intelligent Transfer Service (BITS) is a background file transfer mechanism. BITS prioritizes requests, and you can use it to specify that low-priority transfers occur only when the network is not at full capacity. The transfers continue even after network disconnections or computer restarts. You can write an application that downloads or uploads files in the background using BITS. For example, if a user has requested that Windows Update perform multiple downloads, Windows Update retrieves those updates in the background by submitting a request to BITS. The BITS service then downloads the files in the background and notifies the user when the files are downloaded.

# Network protocols and technologies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network protocols and technologies**

This topic contains a brief overview of the network protocols and technologies in the Windows Server 2003 family. It is divided into three sections: New and updated features since Windows Server 2003 (without SP1), New and updated features since Windows NT 4.0 and New and updated features since Windows 2000.

For links to more information about the features in this release, see New Features.

The Windows Server 2003 family includes a variety of network protocols and technologies to extend the capabilities of your server. For information about additional technologies related to networking, see Performance and tuning features.

### Notes

- The following network protocols have been removed from all products in the Windows Server 2003 family: the Data Link Control (DLC) protocol (related to certain models of printers) and the NetBIOS Extended User Interface (NetBEUI).
- The following network protocols have been removed from the Itanium-based versions of the Windows Server 2003 family: Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX), IPX-dependent services, and Open Shortest Path First (OSPF).
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition supports infrared (IR) networking.

**New and updated features since Windows Server 2003 (without SP1)**

Windows Server 2003 operating systems with Service Pack 1 (SP1) offer the following improvements (compared to Windows Server 2003 without SP1) that help provide increased levels of support for network protocols and technologies:

### Remote Access Quarantine Service, part of Network Access Quarantine Control

Remote Access Quarantine Service is an essential component of Network Access Quarantine Control for Windows Server 2003 with SP1.

Network administrators can enforce network access requirements on remote computers with Network Access Quarantine Control. Lack of access makes it difficult to enforce network requirements (such as the use of antivirus software) on remote computers. By deploying the components of Network Access Quarantine Control, including installing the optional Remote Access Quarantine Service (RQS.exe) component on Windows Server 2003 with SP1, and by including the RQC.exe component with the advanced customization features of Connection Manager, network administrators can create connections that check for required programs, registry settings, files, or combinations thereof. They can also quarantine a remote-access session until these checks have been performed.

Installing the Remote Access Quarantine Service does not automatically start the Remote Access Quarantine Agent service, another aspect of Network Access Quarantine Control. As the network administrator, you must decide the best time to start the service in relation to the configuration of the Routing and Remote Access service.

The Remote Access Quarantine Agent service depends on the Routing and Remote Access service. However, when the Routing and Remote Access service is restarted, the Remote Access Quarantine Agent service is not automatically restarted. You must manually restart the Remote Access Quarantine Agent service.

For more information, see the following:

*Network Access Quarantine Control in Windows Server 2003* at the Microsoft Web site.

*VPN Quarantine Sample Scripts for Verifying Client Health Configurations* at the Microsoft Web site.

### ValueAdd folder has been replaced by Support folder

The ValueAdd folder has been removed from Windows Server 2003 with SP1. To install support tools that were previously found in that folder, such as Phone Book Administrator, go to the Support folder on your Setup CD.

**New and updated features since Windows NT 4.0**

The Windows Server 2003 family offers the following improvements (in comparison to Windows NT 4.0) that help provide

increased levels of support for network protocols and technologies:

**Winsock Direct**

With Winsock Direct, applications that use Winsock can perform faster and with less CPU overhead when they communicate across a system area network (SAN). Winsock Direct has the effect of streamlining communications between distributed components if there is a SAN in place.

**DHCP with DNS and Active Directory**

Dynamic Host Configuration Protocol (DHCP) works with DNS and Active Directory on Internet Protocol (IP) networks, helping to free you from assigning and tracking static IP addresses. DHCP dynamically assigns IP addresses to computers or other resources connected to an IP network.

**Internet Connection Sharing**

With the Internet Connection Sharing feature of Network Connections, you can connect your home network or small office network to the Internet. For example, you might have a home network that connects to the Internet with a digital subscriber line (DSL) connection. By enabling Internet Connection Sharing on the computer that uses DSL, you can provide network address translation, addressing, and name resolution services for all computers on your home network. For more information, see Internet Connection Sharing and network address translation.

Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Network Address Translation**

Network Address Translation (NAT) hides internally managed IP addresses from external networks by translating private internal addresses to public external addresses. This reduces IP address registration costs by letting you use private IP addresses internally, with translation to a small number of registered IP addresses externally. It also hides the internal network structure, reducing the risk of attacks against internal systems.

**Virtual private networking**

You can give users ready access to your organization's network even when they are out of the office, and reduce the cost of such access, by implementing a virtual private network (VPN). The VPN connection creates a secure tunnel across the Internet into the private network.

There are two types of VPN technology in the Windows Server 2003 family:

- Point-to-Point Tunneling Protocol (PPTP), which employs user-level Point-to-Point Protocol (PPP) authentication methods and Microsoft Point-to-Point Encryption (MPPE) for data encryption.

- Layer Two Tunneling Protocol (L2TP) with Internet Protocol security (IPSec). L2TP employs user-level PPP authentication methods and computer-level certificates with IPSec for data encryption.

On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Routing and Remote Access**

Routing and Remote Access replaces the Routing and Remote Access Service (RRAS) and Remote Access Service (RAS) features in Windows NT 4.0. Routing and Remote Access is a single, integrated service that terminates connections from either dial-up or VPN clients or that provides routing (IP, IPX, and AppleTalk), or both. With Routing and Remote Access, your server can function as a remote access server, a VPN server, a gateway, or a branch-office router. For more information, see Routing and Remote Access.

**Asynchronous transfer mode**

Asynchronous transfer mode (ATM) is a high-speed, connection-oriented protocol designed to transport multiple types of traffic across a network. It is applicable to both local area networks (LANs) and wide area networks (WANs). Using ATM, your network can simultaneously transport a wide variety of network traffic: voice, data, image, and video.

#### New and updated features since Windows 2000

The Windows Server 2003 family offers the following improvements (in comparison to Windows 2000) that help provide increased levels of support for network protocols and technologies:

**Simple Object Access Protocol**

Simple Object Access Protocol (SOAP) is an XML-based protocol for exchanging information in a decentralized, distributed environment. It consists primarily of two core parts: an envelope for handling extensibility and modularity and an encoding mechanism for representing data types within the envelope. It does not require synchronous execution or request/response interaction. SOAP provides a simple protocol with an extensibility mechanism for adding functionality, such as message routing, security, and guaranteed delivery. By relying on the availability of ubiquitous protocols, including not only XML and Hypertext Transfer Protocol (HTTP) but Simple Mail Transfer Protocol (SMTP) and File Transfer Protocol (FTP), SOAP is available to the broadest assortment of platforms and clients.

**Fibre Channel support**

Fibre Channel is a technology for 1-gigabit-per-second data transfer that maps common transport protocols such as small computer system interface (SCSI) and Internet Protocol (IP), merging networking and high-speed I/O into a single connectivity technology. Fibre Channel technology gives you a way to address the distance and the address-space limitations of conventional channel technologies. For more information, see Fibre Channel hardware configuration.

**Automatic configuration for multiple networks**

Automatic configuration for multiple networks provides easy access to network devices and the Internet. It also provides the mobile computer user with seamless operations on both office and home networks without having to reconfigure TCP/IP settings manually. With automatic configuration for multiple networks, TCP/IP can use an alternative configuration if a DHCP server is not found.

**Connection Manager improvements**

Connection Manager offers several new features:

- **Automatic proxy configuration for Connection Manager clients.** Configure user proxy settings to ensure that the user has appropriate access to internal and external resources while connected to your service.
- **Connection Manager Favorites.** Eliminates repetitive configuration of Connection Manager properties when users switch among common dialing locations.
- **Client log files.** Troubleshoot problems with Connection Manager connections.
- **Connection Manager Administration Kit.** Customize Connection Manager for your needs using the Connection Manager Administration Kit (CMAK) Wizard.

**Web Distributed Authoring and Versioning redirector**

The Web Distributed Authoring and Versioning (WebDAV) redirector supports the WebDAV protocol for remote document sharing over HTTP. The WebDAV redirector supports the use of existing applications, and it supports file sharing across the Internet (through firewalls, routers, and so forth) to HTTP servers.

**Internet Authentication Service as RADIUS proxy**

This feature supports the ability of Internet Authentication Service (IAS) to forward Remote Authentication Dial-In User Service (RADIUS) authentication and accounting requests to another RADIUS server. This includes flexible, rule-based forwarding, load balancing, and failover between multiple IAS/RADIUS server and load balancing RADIUS-EAP requests; the ability to force the client into a compulsory tunnel with or without user authentication; and selective forwarding of authentication and accounting requests to different RADIUS servers.

This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition. In Windows Server 2003, Standard Edition, you can configure IAS with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. For more information about these limits, see Internet Authentication Service.

**Wireless networking**

Wireless networking functionality in the Windows Server 2003 family supports the IEEE 802.11 standard, and it minimizes the configuration that is required to access wireless networks. The automatic wireless network configuration feature makes it possible for users to roam across different wireless networks in different locations, without the need to reconfigure their network connection settings for each new location. Security options for 802.11 include authentication services and encryption services that are based on the Wired Equivalent Privacy (WEP) algorithm. For enhanced security, 802.1X authentication is enabled by default in the Windows Server 2003 family of products. IEEE 802.1X authentication provides authenticated access to 802.11 wireless networks and to wired Ethernet networks. IEEE 802.1X authentication provides support for Extensible Authentication Protocol (EAP) security types, so that users can use authentication methods such as certificates.

**IP version 6**

The Windows Server 2003 family supports IP version 6 (IPv6), which is a suite of standard protocols that is the next generation of network layer protocols for the Internet. For more information, see IP Version 6.

# Installing and Upgrading the Operating System

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing and upgrading the operating system**

These topics can help you plan for upgrading to or installing Windows Server 2003 operating systems in a network with one to five servers and 100 or fewer clients. For help with planning the upgrade or installation of a larger number of servers, see the Windows Server 2003 Deployment Kit at the Microsoft Windows Server System Web site. For information about obtaining the Windows Server 2003 Resource Kit, see Using the Windows Deployment and Resource Kits.

- Before running Setup, see Checklists for installing and upgrading the operating system.
- For tips about installing, see Best practices for installing and upgrading the operating system.
- For help with specific tasks, see How to install and upgrade the operating system.
- For general background information, see Concepts for installing and upgrading the operating system.
- For problem-solving instructions, see Troubleshooting for installing and upgrading the operating system.

# Checklists for installing and upgrading the operating system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

- Checklist: Performing an upgrade
- Checklist: Performing a new installation

# Checklist: Performing an upgrade

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Performing an upgrade**

| Step | Reference |
|---|---|
| Check system requirements. | System Requirements and Hardware Compatibility |
| Read Relnotes.htm and hardware compatibility information. | Important files to review<br><br>Hardware compatibility |
| Determine whether your current operating system is one from which you can upgrade to Windows Server 2003. | Upgrading or Installing |
| If you have server clusters, review the choices for upgrading on cluster nodes. | Upgrading to Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition on Cluster Nodes<br><br>Checklist: Preparation for upgrading a cluster |
| For upgrades in any domain where all domain controllers run Windows NT 4.0, review the upgrade plan. | Upgrades in a Windows NT 4.0 Domain |
| Decide whether to upgrade any FAT or FAT32 partitions to NTFS. | Choosing a File System for the Installation Partition |
| Check the system log for errors that could cause problems during the upgrade. | Preparing your system for an upgrade |
| Back up files. | Preparing your system for an upgrade |
| If upgrading from Windows NT 4.0, prepare mirror sets or other disk sets for upgrade. | Working with volume, mirror, or stripe sets or stripe sets with parity |
| Disconnect UPS devices. | Preparing your system for an upgrade |
| Review hardware and software. | Preparing your system for an upgrade |
| Start Setup, and follow the instructions provided. | Start Setup on an x86-based computer<br><br>Start Setup on an Itanium architecture-based computer<br><br>Start Setup on an x64-based computer |
| If domain controller upgrades are complete, review concepts about domain functional levels and, if appropriate, raise the functional level. | Domain and forest functionality<br><br>Upgrading from a Windows NT domain |

# Checklist: Performing a new installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Performing a new installation**

| Step | Reference |
|---|---|
| Check system requirements. | System Requirements and Hardware Compatibility |
| Read Relnotes.htm and hardware compatibility information. | Important files to review Hardware compatibility |
| Decide whether to upgrade or install. | Upgrading or Installing |
| If you have server clusters, review the process for installation on cluster nodes. | Checklist: Preparation for installing a cluster |
| Decide what licensing mode to use: Per Device or Per User or Per server. | Choosing a licensing mode |
| Decide whether you want to be able to choose between different operating systems each time you start the computer. | Deciding Whether a Computer Will Contain More Than One Operating System |
| Decide what file system to use: NTFS, FAT, or FAT32. | Choosing a File System for the Installation Partition |
| Decide what partition you will install on, or whether you will need a new partition for your installation. | Planning Disk Partitions or Volumes for New Installations |
| Decide how to handle networking, IP addresses, and TCP/IP name resolution. | Networks: TCP/IP, IP Addresses, and Name Resolution |
| Decide whether to create domains or workgroups for your servers. | Deciding Between Workgroups and Domains |
| If the computer already has a working operating system on it, back up files, and check the system log for errors. | Preparing your system for a new installation |
| Uncompress any DriveSpace or DoubleSpace volumes. | Preparing your system for a new installation |
| If the computer runs Windows NT 4.0 and uses mirror sets or other disk sets, prepare the disks. | Working with volume, mirror, or stripe sets or stripe sets with parity |
| Disconnect UPS devices. | Preparing your system for a new installation |
| Start Setup, and follow the instructions provided. | Start Setup on an x86-based computer<br><br>Start Setup on an Itanium architecture-based computer<br><br>Start Setup on an x64-based computer |

# Best practices for installing and upgrading the operating system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for installing and upgrading the operating system**

**Prepare the server before upgrading or installing.**

- If you are performing an upgrade, see:
  - Checklist: Performing an upgrade
  - Preparing your system for an upgrade
- If you are performing a new installation, see:
  - Checklist: Performing a new installation
  - Preparing your system for a new installation

**Confirm server settings after running Setup.**

- **Check Event Viewer.** Check Event Viewer for messages associated with the installation or startup process. For more information, see Event Viewer.
- **Set Event Viewer log size and wrap setting.** Define your event log size and wrap (overwrite) setting to match your business and security requirements. For information, see Set event logging options.
- **Check server optimization.** Adjust the server optimization setting to match the role the computer will play in your organization. For information about how to access server optimization settings, see To configure memory-related settings on your computer.
- **Verify settings such as IP, DNS, WINS, and default gateway settings.** Use the administrative tools described in Network Services to configure and verify your settings. You can also open a command prompt, type **ipconfig /all** and verify the displayed settings.
- **View the full computer name.** Open a command prompt and type **net config rdr** to view the full computer name. Compare the results against the Active Directory name to confirm that they match or vary as intended.
- **Validate server cluster settings.** Follow the steps in Checklist: Validating your clustering system to confirm the basic configuration settings for your cluster storage.
- **Specify options for paging files and memory dumps.** Set the paging file size and placement based on memory size and server usage. For information about how to do this, see Change virtual memory settings. Also, set the options for the operating system to use if the system stops unexpectedly (for example, the options for memory dumps). For more information, see Specify what happens when the system stops unexpectedly
- **Configure for remote administration.** Identify the remote administration methods you plan to use, and configure them as necessary. For more information, see Remote Administration.

**Confirm server security settings after running Setup.**

- **Review startup settings for services.** Open Computer Management. In the Services folder under Services and

Applications, change the **Startup Type** of services so that services necessary for your server, and only those services, start automatically. Also, confirm that all services that are set to start automatically can start without user intervention or multiple retries. For a list of topics about services, see Services Snap-in. For specific information about services that start by default after Setup (which are different than in Windows 2000), see Default settings for services.

- **Review Internet Explorer security settings.** In Windows server 2003 family of operating systems, the Internet Explorer Enhanced Security Configuration is enabled by default. The security settings in this configuration can help make your computer more secure by limiting its exposure to malicious Web sites. Therefore, with this enhanced level of security, you might find that some Web sites do not display correctly in Internet Explorer when you are browsing Internet and intranet Web sites. Also, you might be prompted to enter your credentials when accessing network resources, such as files in Universal Naming Convention (UNC) shared folders. You can easily change the enhanced security settings. For more information, see Internet Explorer Enhanced Security Configuration overview.

- **Review open network ports.** To help protect servers against attacks, close ports that are not necessary for the system to function properly. You can review ports with the **netstat** command. It is useful to also use an external port scanner and compare the results with the **netstat** results. Be sure to notify other system administrators or affected operations staff before running the port scanner, in case it affects system behavior or sets off your existing triggers for detecting intrusion. For more information about **netstat** and how you can determine which process is associated with each open port, see Netstat. For more information, also see port number assignments at the Internet Assigned Numbers Authority (IANA) Web site.

  **Note**

  - Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

- **Rename or disable the Administrator account; set up or review permissions and accounts that are relevant for this server.** Rename or disable the Administrator account created during Setup (it cannot be deleted). Set up the appropriate accounts and permissions for configuring, backing up, operating, and using this server. For both users and administrators, limit access to the minimum access that is necessary; for example, for someone who only performs backups, use the Backup Operators group. Also, set up policies to ensure that users maintain strong passwords. For more information, see:

  - Best practices for security and Best practices for Active Directory Users and Computers
  - Understanding Groups and Default groups
  - Strong passwords

- **Understand the security implications that apply to your server.** The questions to ask about security vary with the intended server use. For example, for a file server, who needs access to a given file or directory, what type of access is needed, and what are the possible consequences of unauthorized access? For a Web server, what Web applications do you want to use, what risks are inherent in those applications, and what settings will help mitigate the risks? For more information, see the links in the following table:

| Type of server | Suggested security topics |
|---|---|
| Any server | See Best practices for security and topics for the technologies you use in Best practices: pointers to recommendations and tips. |
| File server | If you are using Shared Folders, see Best practices for Shared Folders for Shared Folders. If you are using Active Directory, see Security overview. |
| Multimedia server | See security sections in the documentation for Windows Media Services, which is available when you install Windows Media Services. For more information, see Using Windows Media Services. |
| Print server | See Printing Best practices for printing. |

| | |
|---|---|
| Terminal server | See Terminal Server Best practices for Terminal server. |
| Web server | See the Security section of the server Administration Guide for Internet Information Services (IIS); this documentation is available when you install IIS. |

- **Configure auditing settings, set up system protections, and install a virus scanner.** Review and implement auditing settings, including those for security events, so that you can track all changes to the system and protect against unwanted changes to software. For more information, see:

  - Auditing Security Events

  - Topics about system information and the protection of system files and drivers:

    System Information

    Windows File Protection

    Driver signing for Windows

    System File Checker (sfc)

    Using File Signature Verification

    To set file signature verification options

  - Best practices for security

- **Ensure that an appropriate firewall is in place.** A firewall is a security system that acts as a protective boundary between a network and the outside world. For information about the firewall functionality that is available in Windows Server 2003 operating systems, see Help: Windows Firewall.

- **Consider using the Security Configuration Wizard.** The Security Configuration Wizard reduces the attack surface of servers that are running Windows Server 2003 with Service Pack 1 (SP1). The wizard asks the user a series of questions designed to determine the functional requirements of the server. Any functionality that is not required by the roles being performed by the server is then disabled. In addition to being a fundamental security best practice, reducing the attack surface increases the diversity of your Windows landscape, and reduces the number of systems that need to be immediately updated if a security issue arises. For more information, see Windows Server 2003 TechCenter.

**Begin routine operations such as backups.**

- **Analyze the hard disk for fragmentation.** After you run Setup and install programs, analyze the hard disk to see if it needs to be defragmented. For more information, see Analyze a volume.

- **Perform backups and install the Recovery Console.** Perform backups, install the Recovery Console so that it is readily available whenever you start the computer, and practice or simulate recovery operations (at times when demand on the server is low). For more information, see:

  - Backup

  - Install the Recovery Console as a startup option

  - Disaster Recovery

**Familiarize yourself with available information.**

- **Review the Best practices topics and checklists that relate to the features and components you will use.** For a

list of Best Practices topics, see Best practices: pointers to recommendations and tips. For a list of Checklists, see Checklists Provided in Windows Server 2003 Help.

- **Familiarize yourself with other sources of information and tools.** There are a variety of lists of resources in Help and Support, including the Resources for installing and upgrading the operating system topic for Getting Started. A variety of information is also available at the sources listed in Support resources.

# How to install and upgrade the operating system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Start Setup on an x86-based computer
- Start Setup on an Itanium architecture-based computer
- Start Setup on an x64-based computer
- Enter Server Settings for a New Installation

# Start Setup on an x86-based computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Start Setup on an x86-based computer**

- Start Setup from the CD on a computer running Windows
- Start Setup from a network
- Start Setup for a new installation from the CD on a computer running MS-DOS
- Start Setup for a new installation by starting the computer from the CD
- Start Setup and provide a mass storage driver or a HAL file

# Start Setup from the CD on a computer running Windows

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Setup from the CD on a computer running Windows**

1. Insert the CD in the drive.
2. To begin Setup, do one of the following:
   - For a computer running any version of Windows other than Windows 3.*x*, wait for Setup to display a dialog box.
   - For a computer running Windows 3.*x*, use File Manager to change to the CD-ROM drive and to change to the I386 directory, and then double-click **Winnt.exe**.
3. Follow the Setup instructions.

## Notes

- To run Setup for an upgrade, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- If you are running Setup on a computer running Windows 3.*x* or MS-DOS, for best efficiency, use disk caching. Otherwise, the Setup process (started from Winnt.exe) could take a long time. To enable disk caching on a computer running Windows 3.*x* or MS-DOS, you can use SMARTDrive. For information about SMARTDrive, see the documentation for Windows 3.*x* or MS-DOS.
- Setup works in several stages, prompting you for information, copying files, and restarting.
- For information about unattended Setup, see Related Topics.
- If you have a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for your hard disk, see Related Topics for information about how to ensure that the controller driver file is available to the Setup program.
- If you have specialized hardware requiring a custom hardware abstraction layer (HAL) file supplied by your computer manufacturer, see Related Topics for information about how to make the HAL file available during Setup.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Performing a new installation
Start Setup from a network
Start Setup for a new installation from the CD on a computer running MS-DOS
Start Setup for a new installation by starting the computer from the CD
Start Setup and provide a mass storage driver or a HAL file
Planning for unattended Setup

# Start Setup from a network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Setup from a network**

1. On a network server, share the installation files, either by inserting the CD and sharing the CD-ROM drive, or by copying the files from the I386 folder on the CD to a shared folder.

2. On the computer on which you want to install, connect to the shared Setup files:

   - If you are sharing the CD-ROM drive, connect to the shared drive and change to the I386 folder.

   - If you are sharing a folder, connect to that folder.

3. Find and run the appropriate file in the I386 directory of the CD or in the shared folder:

   - From a computer running MS-DOS or Windows 3.*x*, run Winnt.exe.

   - From a computer running Windows 95, Windows 98, Windows Millennium Edition, Windows NT with Service Pack 5 or later, Windows 2000, or Windows XP, run Winnt32.exe.

4. Follow the Setup instructions.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Setup works in several stages, prompting you for information, copying files, and restarting.

- For information about unattended Setup, see Related Topics.

- If you have a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for your hard disk, see Related Topics for information about how to ensure that the controller driver file is available to the Setup program.

- If you have specialized hardware requiring a custom hardware abstraction layer (HAL) file supplied by your computer manufacturer, see Related Topics for information about how to make the HAL file available during Setup.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Performing a new installation
Start Setup from the CD on a computer running Windows
Start Setup for a new installation from the CD on a computer running MS-DOS
Start Setup for a new installation by starting the computer from the CD
Start Setup and provide a mass storage driver or a HAL file

# Start Setup for a new installation from the CD on a computer running MS-DOS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Setup for a new installation from the CD on a computer running MS-DOS**

1. Insert the CD in the drive.
2. At the command prompt, type:

   ***d*:**

   where *d* is the drive letter of the CD-ROM drive.
3. Type:

   **cd i386**
4. Type:

   **winnt**
5. Follow the Setup instructions.

## Notes

- If you are running Setup on a computer running Windows 3.*x* or MS-DOS, for best efficiency, use disk caching. Otherwise, the Setup process (started from Winnt.exe) could take a long time. To enable disk caching on a computer running Windows 3.*x* or MS-DOS, you can use SMARTDrive. For information about SMARTDrive, see the documentation for Windows 3.*x* or MS-DOS.
- Setup works in several stages, prompting you for information, copying files, and restarting.
- For information about unattended setup, see Related Topics.
- If you have a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for your hard disk, see Related Topics for information about how to ensure that the controller driver file is available to the Setup program.
- If you have specialized hardware requiring a custom hardware abstraction layer (HAL) file supplied by your computer manufacturer, see Related Topics for information about how to make the HAL file available during Setup.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Performing a new installation
Start Setup from the CD on a computer running Windows
Start Setup from a network
Start Setup for a new installation by starting the computer from the CD
Start Setup and provide a mass storage driver or a HAL file

# Start Setup for a new installation by starting the computer from the CD

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Setup for a new installation by starting the computer from the CD**

1. Determine whether the computer on which you want to start Setup can be started from the CD-ROM drive, and whether you want to perform a new installation (not an upgrade). Continue only if both are true.

2. Insert the CD in the drive, and then restart the computer.

3. Follow the instructions for your operating system to boot the computer from the CD.

4. Wait for Setup to display a dialog box, and then follow the Setup instructions.

## Notes

- This method applies only if you want to perform a new installation, not an upgrade. Using this method, you can perform an installation on a computer that does not have an operating system, although you can also use this method on computers that have operating systems.

- Setup works in several stages, prompting for information, copying files, and restarting.

- For information about unattended Setup, see Related Topics.

- If you have a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for your hard disk, see Related Topics for information about how to ensure that the controller driver file is available to the Setup program.

- If you have specialized hardware requiring a custom hardware abstraction layer (HAL) file supplied by your computer manufacturer, see Related Topics for information about how to make the HAL file available during Setup.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Performing a new installation
Start Setup from the CD on a computer running Windows
Start Setup from a network
Start Setup for a new installation from the CD on a computer running MS-DOS
Start Setup and provide a mass storage driver or a HAL file
Planning for unattended Setup

# Start Setup and provide a mass storage driver or a HAL file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Setup and provide a mass storage driver or a HAL file**

1. Determine whether you need to supply Setup with a special file (a mass storage driver or a hardware abstraction layer (HAL) file) supplied by your hardware manufacturer:

   - If you have a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for your hard disk, check to see if the manufacturer has supplied a separate driver file for use with the Windows Server 2003 family.
   - If you have a highly specialized computer system, review your hardware documentation to see if a custom HAL file is supplied by your computer manufacturer.

2. If you determine that you need a driver or HAL file supplied by your hardware manufacturer, before beginning Setup, locate the floppy disk or other medium containing the file.

3. During the early part of Setup, watch the bottom of the screen for a line that prompts you to press F6. Press the appropriate key as follows:

   - To supply the driver for a mass storage controller to Setup, press F6 as instructed.
   - To supply a HAL file to Setup, press F5.

4. Follow further prompts for guidance in supplying the driver file to Setup so that it can gain access to the mass storage controller.

### Note

- If you are not sure whether you must obtain a separate driver file from the manufacturer of your mass storage controller, you can try running Setup. If the controller is not supported by the driver files on the Setup CD and therefore requires a driver file that is supplied by the hardware manufacturer, Setup stops and displays a message saying that no disk devices can be found, or it displays an incomplete list of controllers. After you obtain the necessary driver file, restart Setup, and press F6 when you are prompted.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Performing a new installation
Start Setup from the CD on a computer running Windows
Start Setup from a network
Start Setup for a new installation from the CD on a computer running MS-DOS
Start Setup for a new installation by starting the computer from the CD

# Start Setup on an Itanium architecture-based computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Setup on an Itanium architecture-based computer**

1. Immediately after turning on the computer, insert the Setup CD in the CD-ROM drive.
2. When the **EFI Boot Manager** menu appears, select the CD-ROM option.
3. When prompted, press any key to start from the CD-ROM.
4. Follow the on-screen Setup instructions.

### Important

- Itanium architecture-based computers require a minimum 100-MB FAT partition for the operation of the Extensible Firmware Interface (EFI). This EFI system partition is created automatically during Setup, and it stores programs and information files that the EFI uses to start the operating system. Do not delete or reformat this partition.
- Setup also creates a Microsoft Reserved (MSR) partition, which is required by the operating system. The size of this partition depends on the size of the hard disk. On drives less than 16 GB in size, the MSR is 32 MB. On drives greater than or equal to 16 GB, the MSR is 128 MB. Do not delete or reformat this partition.

### Notes

- Before installing, verify that you are using the correct version of the firmware. For more information, refer to the hardware manufacturer's documentation.
- The EFI defines a new partition style called GUID partition table (GPT). The installation partition on an Itanium architecture-based computer must be on a GPT disk. Setup automatically specifies GPT as the disk to be used for installation. You can still read and write to master boot record (MBR) volumes, which are common to most computers, but you cannot install on master boot record volumes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Extensible Firmware Interface
GUID partition table

# Start Setup on an x64-based computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Setup on an x64-based computer**

1. Determine whether the computer on which you want to start Setup can be started from the CD-ROM drive.
2. Insert the CD in the drive, and then restart the computer.
3. Follow the instructions for your operating system to boot the computer from the CD.
4. Wait for Setup to display a dialog box, and then follow the Setup instructions.

# Enter Server Settings for a New Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enter server settings for a new installation**

- Specify regional and language options and user's name during Setup
- Specify licensing mode during Setup
- Enter your computer name
- Set the Administrator account password
- Set the date and time
- Allow Setup to assign or obtain an IP address
- Specify a static IP address and settings needed for DNS and WINS
- Specify the workgroup or domain name

## Notes

- The topics in this section provide details about entering server settings for a new installation. If you are upgrading, you can skip these topics because Setup will use your previous settings.
- If you have a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for your hard disk, or if you have a specialized computer requiring a custom hardware abstraction layer (HAL) file supplied by your computer manufacturer, see Start Setup and provide a mass storage driver or a HAL file.
- After you start Setup, a process begins in which necessary Setup files are copied to the hard disk. During this process, Setup displays dialog boxes that you can use to select various options.

# Specify regional and language options and user's name during Setup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify regional and language options and user's name during Setup**

1. During Setup, in the **Regional and Language Options** dialog box, follow the instructions for customizing language and locale settings. You can use multiple regional and language options.

2. In the **Personalize Your Software** dialog box, type your name and, optionally, your organization.

3. Continue with Setup.

## Notes

- To open Regional and Language Options, click **Start**, click **Control Panel**, and then double-click **Regional and Language Options**.

- If you select a European country or region in the list of countries/regions, or if you live in a country or region where the euro has been introduced, it is a good idea to verify that the default currency settings in **Regional and Language Options** meet your needs.

- After you run Setup, you can modify regional and language options in Control Panel. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Choosing components to install

# Specify licensing mode during Setup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify licensing mode during Setup**

1. During Setup, in the **Licensing Modes** dialog box, select your client licensing mode:

    - **Per Server**
    - **Per Device or Per User**

2. Continue with Setup.

### Note

- If you are unsure which mode to use, select Per Server, because you can change once from Per Server to Per Device or Per User at no cost. For more information about licensing modes, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Choosing a licensing mode
Licensing

# Enter your computer name

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enter your computer name**

- During Setup, in the **Computer Name and Administrator Password** dialog box, follow the instructions for entering your computer name. The recommended length for most languages is 15 characters or less. For languages that require more storage space per character, such as Chinese, Japanese, and Korean, the recommended length is 7 characters or less.

## Notes

- It is recommended that you use only Internet-standard characters in the computer name. The standard characters are the numbers from 0 through 9, uppercase and lowercase letters from A through Z, and the hyphen (-) character. Computer names cannot consist entirely of numbers.

- If you are using DNS on your network, you can use a wider variety of characters, including Unicode characters and other nonstandard characters, such as the ampersand (&). Using nonstandard characters may affect the ability of non-Microsoft software to operate on your network. For more information about DNS, see Related Topics.

- The maximum length for a computer name is 63 bytes. If the name is longer than 15 bytes (15 characters in most languages, 7 characters in some), computers running Windows NT Server 4.0 and earlier will recognize this computer by the first 15 bytes of the name only. In addition, there are additional configuration steps for a name that is longer than 15 bytes. For more information, see Related Topics.

- If a computer is part of a domain, you must choose a computer name that is different from any other computer in the domain. To avoid name conflicts, the computer should be unique on the domain, workgroup, or network. If this computer is part of a domain, and it contains more than one operating system, you must use a unique computer name for each operating system that is installed. For example, if the computer name is FileServerNT when the computer is started with Windows NT Server 4.0, the computer must have a different name, perhaps FileServerNew, when it is started with a product in the Windows Server 2003 family. This requirement also applies to a computer that contains multiple installations of the same operating system.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Name resolution for TCP/IP
Workgroups compared with domains
Active Directory naming
Namespace planning for DNS

# Set the Administrator account password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the Administrator account password**

1. During Setup, in the **Computer Name and Administrator Password** dialog box, type a password of up to 127 characters in the **Administrator Password** box. For the strongest system security, use a password of at least 7 characters, and use a mixture of uppercase and lowercase letters, numbers, and other characters, such as *, ?, or $.

2. In the **Confirm Password** box, type the password again.

3. Continue with Setup.

### Important

- After Setup is completed, for best security, change the name of the Administrator account (it cannot be deleted) and keep a strong password on the account at all times. For more information, see Strong passwords.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
User and computer accounts
Predefined security templates
Accounts: Administrator account status

# Set the date and time

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the date and time**

1. During Setup, in the **Date and Time Settings** dialog box, set the date, time, and time zone.
2. If you want the system to automatically adjust for daylight saving time, select the **Automatically adjust clock for daylight saving changes** check box.
3. Continue with Setup.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Allow Setup to assign or obtain an IP address

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow Setup to assign or obtain an IP address**

- During Setup, in the **Networking Settings** dialog box, click **Typical settings**.

## Notes

- When you click **Typical settings** in the **Networking Settings** dialog box, Setup checks to see if there is a DHCP server on your network. If there is a DHCP server on your network, DHCP provides an IP address. If there is no DHCP server on your network, Setup will use a limited IP addressing option called Automatic Private IP Addressing (APIPA). On a server using APIPA, complete the network configuration after Setup, because a server using APIPA can communicate only with other computers using APIPA on the same network segment.

- For information about configuring automatic IP addressing after Setup, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Networks: TCP/IP, IP Addresses, and Name Resolution
Overview of networking and TCP/IP
Internet Protocol (IP) addresses
Specify a static IP address and settings needed for DNS and WINS
Configure TCP/IP for automatic addressing

# Specify a static IP address and settings needed for DNS and WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify a static IP address and settings needed for DNS and WINS**

1. During Setup, in the **Networking Settings** dialog box, click **Custom settings**, and then click **Next**.
2. In the **Networking Components** dialog box, click **Internet Protocol (TCP/IP)**.
3. Click **Properties**.
4. In the **Internet Protocol (TCP/IP) Properties** dialog box, click **Use the following IP address**.
5. In **IP address**, **Subnet mask**, and **Default gateway**, type the appropriate addresses.
6. Under **Use the following DNS server addresses**, type the address of a preferred DNS server and, optionally, an alternate DNS server.

   If the local server is the preferred or alternate DNS server, type the same IP address as assigned in the previous step.
7. If you will use a WINS server, click **Advanced**, and then click the **WINS** tab in the **Advanced TCP/IP Settings** dialog box to add the IP address of one or more WINS servers.
8. Click **OK** in each dialog box, and continue with Setup.

**Note**

- For information about configuring networking settings (including a static IP address) after Setup, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Networks: TCP/IP, IP Addresses, and Name Resolution
Overview of networking and TCP/IP
Internet Protocol (IP) addresses
DHCP
Allow Setup to assign or obtain an IP address
Configure TCP/IP for static addressing

# Specify the workgroup or domain name

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify the workgroup or domain name**

1. During Setup, in the **Workgroup or Computer Domain** dialog box, choose whether your server will be in a workgroup or in a domain.

2. Specify the workgroup or domain name.

3. Continue with Setup.

## Notes

- To avoid name conflicts, the computer should be unique on the domain, workgroup, or network. If this computer is part of a domain, and it contains more than one operating system, you must use a unique computer name with each of the operating systems that is installed. For example, if the computer name is FileserverNT when the computer is started with Windows NT Server 4.0, the computer must have a different name, perhaps FileserverNew, when it is started with Windows Server 2003, Standard Edition. This requirement also applies to a computer that contains multiple installations of the same operating system.

- For more information about the requirements for joining either a workgroup or a domain, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Deciding Between Workgroups and Domains
Workgroups compared with domains
Join a domain
Join a workgroup

# Concepts for installing and upgrading the operating system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts for installing and upgrading the operating system**
This section provides background for planning the installation of Windows Server 2003, Standard Edition, Windows Server 2003, Web Edition, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. For information about which topics apply if you are upgrading and which apply if you are performing a new installation, see the following:

- Checklist: Performing an upgrade
- Checklist: Performing a new installation

The following topics provide background information for the choices you make during the installation process.

- Important files to review
- Choosing a licensing mode
- Using Dynamic Update for updated drivers and other setup files
- Planning for unattended Setup
- System Requirements and Hardware Compatibility
- Upgrading or Installing
- Deciding Whether a Computer Will Contain More Than One Operating System
- Choosing a File System for the Installation Partition
- Planning Disk Partitions or Volumes for New Installations
- Networks: TCP/IP, IP Addresses, and Name Resolution
- Deciding Between Workgroups and Domains
- Preparing Your System Before Running Setup
- Resources for installing and upgrading the operating system

# Important files to review

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Important files to review**

At some point in the process of planning for a new installation or an upgrade, familiarize yourself with the **ReInotes.htm** file found in the \Docs folder on the CD for the Windows Server 2003 operating systems. This file contains important usage information about hardware, networking, applications, and printing. Also familiarize yourself with information about hardware compatibility for products in the Windows Server 2003 family. For more information, see Hardware compatibility.

For the most recent information on compatible applications for products in the Windows Server 2003 family, see the software compatibility section of the Microsoft Web site.

Most of the information contained in the Getting Started topics also appears in the \Docs folder on the Windows Server 2003 operating systems CDs, in text files (which use a file name extension of "TXT").

# Choosing a licensing mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing a licensing mode**

Products in the Windows Server 2003 family support two licensing modes:

- Per Device or Per User
- Per Server

The following diagram illustrates the two modes:







If you choose the **Per Device or Per User** mode, each device or user that accesses a server running a product in the Windows Server 2003 family requires a separate Client Access License (CAL). With one CAL, a particular device or user can connect to any number of servers running products in the Windows Server 2003 family. This is the most commonly used licensing method for companies with more than one server running products in the Windows Server 2003 family.

In contrast, Per Server licensing means that each concurrent connection to this server requires a separate CAL. In other words, this server can support a fixed number of connections at any one time. For example, if you select the **Per Server** client licensing mode with five licenses, this server could have five concurrent connections at any one time (if each client requires one connection, this is five clients at any one time). The clients using the connections do not need any additional licenses.

The Per Server licensing mode is often preferred by small companies with only one server. It is also useful for Internet or remote access servers where the client computers might not be licensed as network clients for products in the Windows Server 2003 family. You can specify a maximum number of concurrent server connections and reject any additional logon requests.

If you are unsure which mode to use, choose **Per Server**, because you can change once from Per Server mode to Per Device or Per User mode at no cost. For more information about licensing and licensing modes, see Licensing overview and the Licensing Checklists for licensing.

**Note**

- If you use Terminal Server, you need to take additional steps to set up and activate Terminal Server Licensing. For more information, see Terminal Server Licensing.

# Using Dynamic Update for updated drivers and other setup files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Dynamic Update for updated drivers and other setup files**

If you have a working Internet connection on the computer on which you run Setup, you can choose Dynamic Update during Setup and obtain the most up-to-date Setup files, including drivers and other files. Whenever an important update is made to any crucial Setup file, that update is made available through Dynamic Update functionality built into the Windows Update Web site. Some of the updated files will be replacements (for example, an updated driver or updated Setup file) and some will be additions (for example, a driver not available at the time that the Setup CD was created). It is recommended that you use Dynamic Update when running Setup.

Dynamic Update has been carefully designed so that it is reliable and easy to use:

- The files on the Dynamic Update section of the Windows Update Web site have been carefully tested and selected. Only files that are important in ensuring that Setup runs well are made available through Dynamic Update. Files with minor updates that will not significantly affect Setup are not part of Dynamic Update.

- Because Dynamic Update downloads only the files that are required for your computer, the Dynamic Update software briefly examines your computer hardware. No personal information is collected, and no information is saved. The only purpose is to select appropriate drivers for your hardware configuration. This keeps the download as short as possible and ensures that only necessary drivers are downloaded to your hard disk.

- You can use Dynamic Update when running a preinstallation compatibility check from the product CD, or when running Setup itself. Either way, you obtain the most up-to-date files for running Setup. For information about running the compatibility check, see Hardware compatibility.

- You can use Dynamic Update with unattended Setup. Preparing for this requires several steps. For a brief description of unattended Setup, see Planning for unattended Setup. For details about how to use Dynamic Update with unattended Setup (also called automated installation), see *Automating and Customizing Installations* at the Microsoft Windows Server System Web site. For information about the *Windows Server 2003 Resource Kit*, see Using the Windows Deployment and Resource Kits.

The Windows Update Web site offers a variety of updates that you can use after completing Setup. For more information, see Windows Update at the Microsoft Web site. For a description of the different ways you can use Windows Update, see Windows Update.

# Planning for unattended Setup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for unattended Setup**

This topic provides general information about unattended Setup. For detailed instructions on running unattended Setup (also called automated installation), see "Automating and Customizing Installations" at the Microsoft Windows Server System Web site. Also see the Windows XP Professional Resource Kit at the Microsoft Windows XP Resource Kits Web site. For information about the Windows Deployment and Resource Kits, see Using the Windows Deployment and Resource Kits.

To simplify the process of setting up a product in the Windows Server 2003 family on multiple computers, you can run Setup unattended. To do this, you create and use an *answer file*, a customized script that answers the Setup questions automatically. Then, you run Winnt32.exe or Winnt.exe with the appropriate options for unattended Setup. Choose the command according to the operating system that is running when you start unattended Setup:

- To start unattended Setup on a computer running MS-DOS or Windows 3.*x*, use Winnt.exe (with the appropriate options).
- To start unattended Setup on a computer running Windows 95, Windows 98, Windows Millennium Edition, Windows NT, Windows 2000, Windows XP, or a product in the Windows Server 2003 family, use Winnt32.exe (with the appropriate options). With Windows NT 4.0, before starting unattended Setup, apply Service Pack 5 or later.

To learn about unattended Setup, see the following sources:

- **For detailed instructions on running unattended Setup (also called automated installation):** See the Windows Deployment and Resource Kits listed in the first paragraph of this topic.
- **To view the command options available for Winnt.exe:** On a computer running Windows 3.*x* or MS-DOS, insert the Setup CD for a product in the Windows Server 2003 family in the CD-ROM drive and open the command prompt. Then, change to the CD-ROM drive, change to the I386 directory, and type:

  **winnt /?**
- **To use an x86-based or x64-based computer to view the command options available for Winnt32.exe:** On a computer running Windows 95, Windows 98, Windows Millennium Edition, Windows NT, Windows 2000, Windows XP, or a product in the Windows Server 2003 family, insert the Setup CD for Windows Server 2003, Standard Edition, in the CD-ROM drive, and open the command prompt (click **Start**, click **Run**, and then type **cmd**). Then, change to the CD-ROM drive, change to the I386 directory, and type:

  **winnt32 /?**
- **To use an Itanium-based computer to view the command options available for Winnt32.exe:** On an Itanium-based computer running Windows XP 64-Bit Edition; the Itanium-based version of Windows Server 2003, Enterprise Edition; or the Itanium-based version of Windows Server 2003, Datacenter Edition, insert the Setup CD for the product in the CD-ROM drive and open the command prompt (click **Start**, click **Run**, and then type **cmd**). Then, change to the CD-ROM drive, change to the IA64 directory, and type:

  **winnt32 /?**

For information about using Winnt.exe or Winnt32.exe for interactive setup (not unattended Setup), see How to install and upgrade the operating system.

# System Requirements and Hardware Compatibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System requirements and hardware compatibility**

- Windows Server 2003, Standard Edition: System requirements
- Windows Server 2003, Web Edition: System requirements
- Windows Server 2003, Enterprise Edition: System requirements
- Windows Server 2003, Datacenter Edition: System requirements
- Hardware overview
- Hardware compatibility
- Taking an inventory of devices that do not use Plug and Play
- Mass storage drivers and the setup process
- Understanding the ACPI BIOS for an x86-based computer

# Windows Server 2003, Standard Edition: System requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003, Standard Edition: System requirements**
To ensure adequate performance, make sure that computers on which you install or upgrade to Windows Server 2003, Standard Edition meet the following requirements:

- **For an x86-based computer**:
  - One or more processors with a recommended minimum speed of 550 megahertz (MHz). The minimum supported speed is 133 MHz. A maximum of four processors per computer is supported. Processors from the Intel Pentium/Celeron family, AMD K6/Athlon/Duron family, or compatible processors are recommended.
  - 256 megabytes (MB) of RAM (recommended minimum). 128 MB is the minimum supported, and 4 gigabytes (GB) is the maximum supported.

    For computers with more than 4 GB of RAM, be sure to confirm hardware compatibility by clicking the appropriate link in Support resources.
- **For an x64-based computer**:
  - One or more processors with a minimum speed of 1.4 gigahertz (GHz). A maximum of four processors per computer is supported.
  - 512 MB of RAM (minimum).

    For computers with more than 4 GB of RAM, be sure to confirm hardware compatibility by clicking the appropriate link in Support resources.
- A hard disk partition or volume with enough free space to accommodate the setup process. To ensure that you have flexibility in your later use of the operating system, we recommend that you allow considerably more space than the minimum required for running Setup, which is approximately 1.25 GB to 2 GB for x86-based versions of Windows Server 2003, and 4 GB for x64-based versions of Windows Server 2003. The larger amount of space is required if you are running Setup across a network instead of from a CD-ROM, or if you are installing on a FAT or FAT32 partition (NTFS is the recommended file system).

  In addition, a domain controller upgrade from Windows NT 4.0 could require much more space than other upgrades or installations, because the existing user accounts database can expand by as much as a factor of ten during the upgrade, as Active Directory functionality is added.

  **Note**

  - The setup process requires the free disk space described in the previous paragraphs. After Setup is finished, actual hard disk space used for the operating system will be more than the free space required for Setup, because of space needed for the paging file, for any optional components you install, and (on domain controllers) for user accounts and other Active Directory information. The usual size for the paging file is 1.5 times the size of the RAM. For more information, see the following:

    Change virtual memory settings

    Specify what happens when the system stops unexpectedly

    Choosing components to install

    Understanding Active Directory

Directory data store

For additional information about planning the sizes of disk partitions or volumes, including information about creating a separate partition for Remote Installation Services, see Considerations when planning disk partitions and Remote Installation Services system requirements.

- A VGA or higher-resolution monitor (Super VGA 800x600 or higher is recommended), keyboard, and (optionally) a mouse or other pointing device.

  As an alternative, for operation without a monitor or keyboard, you can choose a remote diagnostic and support processor that is designed for products in the Windows Server 2003 family. For more information about hardware compatibility, click the appropriate link in Support resources.

*For CD-ROM installation:*

- A CD-ROM or DVD drive.

*For network installation:*

- One or more network adapters and related cables that are designed for products in the Windows Server 2003 family. For more information about hardware compatibility, click the appropriate link in Support resources.

- A server from which to offer network access for the Setup files.

*For additional hardware-based functionality:*

- Appropriate hardware for the functionality you require. For example, if you require sound, you must have a compatible sound card with speakers or headphones. If you plan to support network clients, the servers and clients must have appropriate network adapters and cables. For more details about your hardware, click the appropriate link in Support resources.

# Windows Server 2003, Web Edition: System requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003, Web Edition: System requirements**
To ensure adequate performance, make sure that computers on which you install or upgrade to Windows Server 2003, Web Edition meet the following requirements:

- **For an x86-based computer**:

  - One or two processors with a recommended minimum speed of 550 megahertz (MHz). The minimum supported speed is 133 MHz. A maximum of two processors per computer is supported. Processors from the Intel Pentium/Celeron family, AMD K6/Athlon/Duron family, or compatible processors are recommended.

  - 256 megabytes (MB) of RAM (recommended minimum). 128 MB is the minimum supported and 2 gigabytes (GB) is the maximum supported.

  - A hard disk partition or volume with enough free space to accommodate the setup process. To ensure that you have flexibility in your later use of the operating system, we recommend that you allow considerably more space than the minimum required for running Setup, which is approximately 1.25 GB to 2 GB. The larger amount of space is required if you are running Setup across a network, instead of from a CD-ROM.

    You must use NTFS as the file system with Windows Server 2003, Web Edition.

    **Note** The setup process requires the free disk space described in the previous paragraphs. After Setup is finished, actual hard disk space used for the operating system will be more than the free space required for Setup, because of space needed for the paging file, for any optional components you install, and (on domain controllers) for user accounts and other Active Directory information. The usual size for the paging file is 1.5 times the size of the RAM. For more information, see the following:

    Change virtual memory settings

    Specify what happens when the system stops unexpectedly

    Choosing components to install

    Understanding Active Directory

    Directory data store

    For additional information about planning the sizes of disk partitions or volumes, including information about creating a separate partition for Remote Installation Services, see Considerations when planning disk partitions and Remote Installation Services system requirements.

  - A VGA or higher-resolution monitor (Super VGA 800x600 or higher is recommended), keyboard, and (optionally) a mouse or other pointing device.

  - As an alternative, for operation without a monitor or keyboard, you can choose a remote diagnostic and support processor that is designed for products in the Windows Server 2003 family. For more information about hardware compatibility, click the appropriate link in Support resources.

*For CD-ROM installation:*

  - A CD-ROM or DVD drive.

*For network installation:*

- One or more network adapters and related cables that are designed for products in the Windows Server 2003 family. For more information about hardware compatibility, click the appropriate link in Support resources.

- A server from which to offer network access for the Setup files.

*For additional hardware-based functionality:*

- Appropriate hardware for the functionality you require. For example, if you require sound, you must have a compatible sound card with speakers or headphones, and if you plan to support network clients, the servers and clients must have appropriate network adapters and cables. For more details about your hardware, click the appropriate link in Support resources.

# Windows Server 2003, Enterprise Edition: System requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003, Enterprise Edition: System requirements**
To ensure adequate performance, make sure that computers on which you install or upgrade to Windows Server 2003, Enterprise Edition meet the following requirements:

- **For an x86-based computer**:
  - One or more processors with a recommended minimum speed of 550 megahertz (MHz). The minimum supported speed is 133 MHz. A maximum of eight processors per computer is supported. Processors from the Intel Pentium/Celeron family, AMD K6/Athlon/Duron family, or compatible processors are recommended.
  - 256 megabytes (MB) of RAM (recommended minimum). 128 MB is the minimum supported, and 32 gigabytes (GB) is the maximum supported.

    For computers with more than 4 GB of RAM, be sure to confirm hardware compatibility by clicking the appropriate link in Support resources.
- **For an Itanium-based computer**:
  - One or more processors with a minimum speed of 733 MHz. A maximum of eight processors per computer is supported.
  - 1 GB of RAM minimum, 1024 GB (1 Terabyte) maximum.

    For computers with more than 4 GB of RAM, be sure to confirm hardware compatibility by clicking the appropriate link in Support resources.
- **For an x64-based computer**:
  - One or more processors with a minimum speed of 1.4 GHz. A maximum of eight processors per computer is supported.
  - 1 GB of RAM minimum, 1024 GB (1 Terabyte) maximum.

    For computers with more than 4 GB of RAM, be sure to confirm hardware compatibility by clicking the appropriate link in Support resources.
- A hard disk partition or volume with enough free space to accommodate the setup process. To ensure that you have flexibility in your later use of the operating system, we recommend that you allow considerably more space than the minimum required for running Setup, which is approximately 2 GB to 3 GB for x86-based versions of Windows Server 2003, and 4 GB for Itanium-based and x64-based versions of Windows Server 2003. The larger amount of space is required if you are running Setup across a network instead of from a CD-ROM, or if you are installing on a FAT or FAT32 partition in the case of x86-based or x64-based versions of Windows Server 2003. (NTFS is the recommended file system for both x86-based and x64-based versions of Windows Server 2003. Itanium-based versions of Windows Server 2003 support only NTFS).

  In addition, a domain controller upgrade from Windows NT 4.0 could require much more space than other upgrades or installations, because the existing user accounts database can expand by as much as a factor of ten during the upgrade, as Active Directory functionality is added.

  **Note**

- The setup process requires the free disk space described in the previous paragraphs. After Setup is finished, actual hard disk space used for the operating system will be more than the free space required for Setup, because of space needed for the paging file, for any optional components you install, and (on domain controllers) for user accounts and other Active Directory information. The usual size for the paging file is 1.5 times the size of the RAM. For more information, see the following:

  Change virtual memory settings

  Specify what happens when the system stops unexpectedly

  Choosing components to install

  Understanding Active Directory

  Directory data store

  For additional information about planning the sizes of disk partitions or volumes, including information about creating a separate partition for Remote Installation Services, see Considerations when planning disk partitions and Remote Installation Services system requirements.

- A VGA or higher-resolution monitor (Super VGA 800x600 or higher recommended), keyboard, and (optionally) a mouse or other pointing device.

  As an alternative, for operation without a monitor or keyboard, you can choose a remote diagnostic and support processor that is designed for products in the Windows Server 2003 family. For more information about hardware compatibility, click the appropriate link in Support resources.

*For CD-ROM installation:*

- A CD-ROM or DVD drive.

*For network installation:*

- One or more network adapters and related cables that are designed for products in the Windows Server 2003 family. For more information about hardware compatibility, click the appropriate link in Support resources.

- A server from which to offer network access for the Setup files.

*For additional hardware-based functionality:*

- Appropriate hardware for the functionality you require. For example, if you plan to support network clients, the servers and clients must have appropriate network adapters and cables. As another example, if you require a server cluster, the entire cluster solution must be listed as compatible with Windows Server 2003, Enterprise Edition, (for more details, see information about cluster hardware and drivers in Planning and preparing for cluster installation). For more details about your hardware, click the appropriate link in Support resources.

# Windows Server 2003, Datacenter Edition: System requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003, Datacenter Edition: System requirements**
Computers that run Windows Server 2003, Datacenter Edition must be on the compatibility list for Windows Server 2003, Datacenter Edition. For details, see the hardware compatibility information in the Windows Server Catalog at the Microsoft Web site.

The requirements for systems on the compatibility list for Windows Server 2003, Datacenter Edition include the following:

- **For an x86-based computer**:
  - Must be able to support eight or more processors with a minimum speed of 400 megahertz (MHz). For new installations, the minimum processor speed is 550 MHz. A maximum of 32 processors per computer is supported.
  - 512 megabytes (MB) of RAM minimum, 128 gigabytes (GB) maximum.
  - 2 GB of free hard-disk space (minimum) for Setup.

    For computers with more than 4 GB of RAM, be sure to confirm hardware compatibility by clicking the appropriate link in Support resources.
- **For an Itanium-based computer**:
  - Must be able to support eight or more processors with a minimum speed of 733 MHz. A maximum of 64 processors per computer is supported.
  - 1 GB of RAM minimum, 1024 GB (1 Terabyte) maximum.
  - 4 GB of free hard-disk space (minimum) for Setup.

    For computers with more than 4 GB of RAM, be sure to confirm hardware compatibility by clicking the appropriate link in Support resources.
- **For an x64-based computer**:
  - Must be able to support eight or more processors with a minimum speed of 1.4 GHz. A maximum of 64 processors per computer is supported.
  - 1 GB of RAM minimum, 1024 GB (1 Terabyte) maximum.
  - 4 GB of free hard-disk space (minimum) for Setup.

    For computers with more than 4 GB of RAM, be sure to confirm hardware compatibility by clicking the appropriate link in Support resources.
- A hard-disk partition or volume with enough free space to accommodate the setup process. To ensure that you have flexibility in your later use of the operating system, we recommend that you allow considerably more space than the minimum required for running Setup, which is approximately 2 to 3 GB on an x86-based computer, and 4 GB on an Itanium-based or x64-based computer. The larger amount of space is required if you are running Setup across a network instead of from a CD-ROM.

  **Note**

- The setup process requires the free disk space described in the previous paragraphs. After Setup is finished, the actual hard-disk space used for the operating system will be more than the free space required for Setup. This is because of the space that is needed for the paging file, for any optional components you install, and (on domain controllers) for user accounts and other Active Directory information. The usual size for the paging file is 1.5 times the size of the RAM. For more information, see the following:

  Change virtual memory settings

  Specify what happens when the system stops unexpectedly

  Choosing components to install

  Understanding Active Directory

  Directory data store

  For additional information about planning the sizes of disk partitions or volumes, including information about creating a separate partition for Remote Installation Services, see Considerations when planning disk partitions and Remote Installation Services system requirements.

- A VGA or higher-resolution monitor (Super VGA 800x600 or higher recommended), keyboard, and (optionally) a mouse or other pointing device.

  As an alternative, for operation without a monitor or keyboard, you can choose a remote diagnostic and support processor that is designed for products in the Windows Server 2003 family. For more information about hardware compatibility, click the appropriate link in Support resources.

*For CD-ROM installation:*

- A CD-ROM or DVD drive.

*For network installation:*

- One or more network adapters and related cables that are designed for products in the Windows Server 2003 family. For more information about hardware compatibility, click the appropriate link in Support resources.

- A server from which to offer network access for the Setup files.

*For additional hardware-based functionality:*

- Appropriate hardware for the functionality you require. For example, if you plan to support network clients, the servers and clients must have appropriate network adapters and cables. As another example, if you require a server cluster, the entire cluster solution must be listed as compatible with Windows Server 2003, Datacenter Edition, (for more details, see information about cluster hardware and drivers in Planning and preparing for cluster installation). For more details about your hardware, click the appropriate link in Support resources.

# Hardware overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware overview**

Hardware includes any device that is connected to your computer and is controlled by your computer's microprocessor. This includes equipment that was connected to your computer when it was manufactured, as well as peripheral equipment that you added later. The following are examples of devices:

- Universal Serial bus (USB)
- IEEE 1394
- Modems
- Disk drives
- CD-ROM drives
- Digital video disc (DVD) drives
- Printers
- Network adapters
- Keyboards
- Video adapter cards
- Monitors
- Game controllers

Devices (Plug and Play and non-Plug and Play) can be connected to your computer in several ways. Some devices, such as network adapters and sound cards, are connected to expansion slots inside your computer. Other devices, such as printers and scanners, are connected to ports on the outside of your computer. Some devices, known as PC Cards, connect only to PC Card slots on a portable computer.

For a device to work properly with Windows, software known as a device driver must be installed on the computer. Each device is supported by one or more device drivers, which are typically supplied by the device manufacturer. However, some device drivers are included with Windows. If the device is Plug and Play, Windows can automatically detect it and install the appropriate device drivers.

If the device is not automatically installed by Windows, the Found New Hardware Wizard will appear and ask you to insert any media (such as compact discs or floppy disks) that were provided with the device. Non-Plug and Play devices are installed using the Add Hardware Wizard in Control Panel. If you want to manually install device drivers, you must use Device Manager. Before manually installing device drivers, you should consult the device documentation provided by the manufacturer.

**Important**

- You must be logged on as an administrator or as a member of the Administrators group in order to install or configure a device if:
- The device driver does not have the Designed for Windows Logo or a digital signature.

- Further action is required to install the device, requiring Windows to display a user interface.
- The device driver is not already on your computer.
- You need to configure a device using Device Manager.
- Your computer is connected to a network; network policy settings may also prevent you from installing hardware.

**Note**

- For the latest updates, including security fixes, service packs, new Help files, and device drivers, see Windows Update Web site.

**See Also**
**Concepts**
Hardware profiles overview
Using Internet Explorer
**Other Resources**
Hardware types overview
Installing devices overview

# Hardware compatibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware compatibility**

One of the most important steps to take before running Setup on a server is to confirm that your hardware is compatible with products in the Windows Server 2003 family. You can do this by running a preinstallation compatibility check from the Setup CD or by checking the hardware compatibility information available from Microsoft. Also, as part of confirming hardware compatibility, check to see that you have obtained updated hardware device drivers and an updated system BIOS (or for an Itanium architecture-based computer, the correct firmware).

Regardless of whether you run a preinstallation compatibility check, Setup checks hardware and software compatibility at the beginning of an upgrade or new installation and displays a report if there are incompatibilities.

**Running a preinstallation compatibility check from the Setup CD**

You can run a hardware and software compatibility check from the Setup CD. The compatibility check does not require you to actually begin an upgrade or an installation. To run the check, insert the Setup CD in the CD-ROM drive and, when a display appears, follow the prompts for checking system compatibility. You will be offered the option to download the latest Setup files (through Dynamic Update) when you run the check. If you have Internet connectivity, it is recommended that you allow the download.

For more information about downloading the latest Setup files, see Using Dynamic Update for updated drivers and other setup files.

Another way to run the compatibility check is to insert the Setup CD in the CD-ROM drive, open a command prompt, and type

***d*:\i386\winnt32 /checkupgradeonly**

where *d* represents the CD-ROM drive.

**Checking hardware and software compatibility information available from Microsoft**

Another way to confirm that your hardware and software are designed for products in the Windows Server 2003 family is to click the appropriate links for compatibility information in Support resources.

Note that one type of hardware no longer supported is the microchannel bus.

**Checking drivers and system BIOS or firmware**

Check that you have obtained updated drivers for your hardware devices and that you have the latest system BIOS (for an x86-based or x64-based computer) or the correct firmware (for an Itanium architecture-based computer). The device manufacturers can help you obtain these items. For information about the Advanced Configuration and Power Interface (ACPI) standard for the BIOS, see Understanding the ACPI BIOS for an x86-based computer.

Finally, if you have devices that do not use Plug and Play, or you are aware that your Plug and Play devices are not implemented exactly to the standards, consider taking a device inventory of the hardware devices in your computer. For more information, see Taking an inventory of devices that do not use Plug and Play.

# Taking an inventory of devices that do not use Plug and Play

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Taking an inventory of devices that do not use Plug and Play**

This section describes steps you can take if the devices in your computer do not use Plug and Play technology. For important information about hardware compatibility (including the compatibility of devices), see:

- Windows Server 2003, Enterprise Edition: System requirements
- Hardware compatibility

For specific information about using a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) with a driver that was supplied by the manufacturer, see Mass storage drivers and the setup process.

Products in the Windows Server 2003 family include Plug and Play technology so that devices (for example, video and network adapters) can be automatically recognized by the operating system, configuration conflicts are avoided, and you do not have to specify each device's settings by hand. However, if you have devices that do not use Plug and Play, or you are aware that your Plug and Play devices are not implemented exactly to the standards, you might want to take steps to avoid device configuration conflicts. This section describes steps you can take, if you choose, to understand your device configuration before running Setup.

To take an inventory of your devices, if your computer has an existing operating system, use it to obtain the current settings, such as memory address and interrupt request (IRQ), used with your devices. For example, with Windows NT 4.0, you can use Control Panel to view settings (on the **Start** menu, point to **Settings**, click **Control Panel**, and then double-click icons such as **Network** and **Ports**). You might also choose to view system BIOS information. To do this, watch the screen while starting the computer, and then press the appropriate key when prompted.

At the beginning of an installation or upgrade, the Setup program automatically takes a device inventory as well. For devices that do not use Plug and Play, or that are not implemented exactly to Plug and Play standards, taking your own inventory helps prevent the following difficulties:

- If two or more adapters share IRQ settings or memory addresses, the Setup program might not be able to resolve the conflict. To prevent this, you can take one of two approaches.

  You can remove one of the adapters before running Setup and reinstall it afterward.

  As an alternative, you can modify one adapter's IRQ settings and memory addresses before running Setup, so that each adapter's settings are unique.

- If adapters do not respond in a standard way to the attempts by Setup to detect or enumerate them, Setup might receive indecipherable or inaccurate information. In this case, you might need to remove these devices before running Setup, and reinstall and configure them afterward.

The following table lists the kinds of information to gather if you have devices that do not use Plug and Play, and you decide to take a device inventory before starting Setup.

| Adapter | Information to gather |
|---|---|
| Video | Adapter or chip set type and how many video adapters |
| Network | IRQ, I/O address, DMA (if used) connector type (for example, BNC or twisted pair), and bus type |
| SCSI controller | Adapter model or chip set, IRQ, and bus type |
| Mouse | Mouse type and port (COM1, COM2, or PS/2) or USB |

| | |
|---|---|
| I/O port | IRQ, I/O address, and DMA (if used) for each I/O port |
| Sound adapter | IRQ, I/O address, and DMA |
| Universal serial bus (USB) | Which devices and hubs are attached |
| PC card | Which adapters are inserted and in which slots |
| Plug and Play | Whether enabled or disabled in BIOS |
| BIOS settings | BIOS revision and date |
| External modem | COM port connections (COM1, COM2, and so on) |
| Internal modem | COM port connections; for nonstandard configurations, IRQ and I/0 address |
| Advanced Configuration and Power Interface (ACPI); Power Options | Enabled or disabled; current setting |
| PCI | Which PCI adapters are inserted and in which slots |

# Mass storage drivers and the setup process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Mass storage drivers and the setup process**

If you have a mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for your hard disk, confirm that the controller is designed for products in the Windows Server 2003 family by clicking the appropriate link in Support resources.

If your controller is compatible, but you are aware that the manufacturer has supplied a separate driver file for use with products in the Windows Server 2003 family, obtain the file (on a floppy disk) before you begin Setup. During the early part of Setup, a line at the bottom of the screen will prompt you to press F6. Further prompts will guide you in supplying the driver file to Setup so that it can gain access to the mass storage controller.

If you are not sure whether you must obtain a separate driver file from the manufacturer of your mass storage controller, you can try running Setup. If the controller is not supported by the driver files on the Setup CD, and therefore requires a driver file that is supplied by the hardware manufacturer, Setup stops and displays a message saying that no disk devices can be found, or it displays an incomplete list of controllers. After you obtain the necessary driver file, restart Setup, and press F6 when you are prompted.

# Understanding the ACPI BIOS for an x86-based computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding the ACPI BIOS for an x86-based computer**
For an x86-based computer, the basic input/output system (BIOS) is a set of software through which the operating system (or Setup) communicates with the computer's hardware devices. The Advanced Configuration and Power Interface (ACPI) is the current standard for the way the BIOS works. Products in the Windows Server 2003 family support not only ACPI-compliant BIOS versions, but also some BIOS versions based on older advanced power management (APM) and Plug and Play designs.

Some ACPI-based BIOS versions are not compliant with the standard. The more recent the version of an ACPI BIOS, the more likely that it is compliant. An ACPI-based BIOS that is not compliant with the ACPI standard might not support workable communication between the operating system (or Setup) and your hardware. If workable communication is not supported, Setup stops and displays instructions for contacting your hardware manufacturer and taking other steps to solve the problem. If this happens, follow the instructions provided.

To learn more about the ACPI compliance of your BIOS:

- For information about your BIOS version, before running Setup, restart the computer and watch the text on the screen. Pay particular attention to blocks of text containing the words "BIOS" or "ACPI BIOS."

- For information about BIOS versions for your hardware, check your hardware documentation and contact your hardware manufacturer.

- If you think your BIOS might be a noncompliant ACPI BIOS, see the section about ACPI BIOS in the "System Recovery and Troubleshooting" appendix of the printed *Getting Started* manual.

# Upgrading or Installing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrading or installing**
This section provides information that will help you decide whether to upgrade or to install. It provides details about upgrading and installing, such as a list of the operating systems from which you can upgrade. In addition, it provides information about the steps to take when upgrading a small domain.

- Windows Server 2003, Standard Edition: Upgrade or new installation?
- Windows Server 2003, Enterprise Edition: Upgrade or new installation?
- Windows Server 2003, Datacenter Edition: Upgrade or new installation?
- Upgrades in a domain containing Windows 2000 domain controllers
- Upgrades in a Windows NT 4.0 Domain

# Windows Server 2003, Standard Edition: Upgrade or new installation?

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003, Standard Edition: Upgrade or new installation?**
After you have reviewed System Requirements and Hardware Compatibility and Important files to review, and before you run Setup, you will need to determine whether to upgrade or to perform a new installation. *Upgrading* is either replacing Windows NT 4.0 (with Service Pack 5 or later) or Windows 2000 with a product in the Windows Server 2003 family. In contrast, *installing* means completely removing the previous operating system, or installing a product in the Windows Server 2003 family on a disk or disk partition with no previous operating system.

If you have already decided whether to upgrade or to perform a new installation, see Checklist: Performing an upgrade or Checklist: Performing a new installation.

The following lists describe points to consider when deciding between an upgrade and a new installation.

**Points to consider for an upgrade**

- With an upgrade, configuration is simpler and your existing users, settings, groups, rights, and permissions are retained.
- With an upgrade, you do not need to reinstall files and applications. As with any major changes to the hard disk, however, we recommend that you back up the disk before beginning an upgrade.
- Before planning to perform an upgrade, see "Operating systems from which you can upgrade" later in this topic.
- If you are upgrading in a domain that includes domain controllers running Windows 2000, be sure to read Upgrades in a domain containing Windows 2000 domain controllers.
- If you are upgrading in a domain where all domain controllers are running Windows NT 4.0, be sure to read Upgrades in a Windows NT 4.0 Domain.
- If you want to upgrade and then use the same applications as before, be sure to review applications information in Relnotes.htm (in the \Docs folder on the product CD). Also, for the most recent information on compatible applications for products in the Windows Server 2003 family, see the software compatibility section of the Microsoft Web site.

**Points to consider for a new installation**

- If you reformat your hard disk and then perform a new installation, the efficiency of your disk might improve (compared to not reformatting it). Reformatting also gives you the opportunity to modify the size or number of disk partitions, to make them match your requirements more closely.
- If you want to practice careful configuration management, for example for a server where high availability is important, you might want to perform a new installation on that server instead of an upgrade. This is especially true on servers on which the operating system has been upgraded several times in the past.
- It is possible to install Windows Server 2003, Standard Edition, and also allow the computer to sometimes run another operating system. Setting up the computer this way, however, presents complexities because of file system issues. For more information, see Deciding Whether a Computer Will Contain More Than One Operating System.

**Note**

- If you used Windows NT 4.0 to create a volume set, mirror set, stripe set, or stripe set with parity, and you want to run Setup for Windows Server 2003, Standard Edition on that computer, you must prepare the disk set first. For details, see Working with volume, mirror, or stripe sets or stripe sets with parity.

**Operating systems from which you can upgrade**
If you upgrade, Setup automatically installs Windows Server 2003, Standard Edition into the same folder as the currently installed operating system. You can upgrade to Windows Server 2003, Standard Edition from the following versions of Windows:

- Windows NT Server 4.0 with Service Pack 5 or later
- Windows NT Server 4.0, Terminal Server Edition with Service Pack 5 or later
- Windows 2000 Server

### Note

- You cannot upgrade from an earlier version of a Windows server operating system to an Itanium-based or x64-based version of a Windows server operating system.

# Windows Server 2003, Enterprise Edition: Upgrade or new installation?

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003, Enterprise Edition: Upgrade or new installation?**
After you have reviewed System Requirements and Hardware Compatibility and Important files to review, and before you run Setup, you will need to determine whether to upgrade or to perform a new installation. *Upgrading* is either replacing Windows NT 4.0 (with Service Pack 5 or later) or Windows 2000 with a product in the Windows Server 2003 family. In contrast, *installing* means completely removing the previous operating system, or installing a product in the Windows Server 2003 family on a disk or disk partition with no previous operating system.

If you have already decided whether to upgrade or to perform a new installation, see Checklist: Performing an upgrade or Checklist: Performing a new installation.

The following lists describe points to consider when deciding between an upgrade and a new installation.

**Points to consider for an upgrade**

- With an upgrade, configuration is simpler and your existing users, settings, groups, rights, and permissions are retained.
- With an upgrade, you do not need to reinstall files and applications. As with any major changes to the hard disk, however, we recommend that you back up the disk before beginning an upgrade.
- Before planning to perform an upgrade, see "Operating systems from which you can upgrade" later in this topic.
- If you are upgrading in a domain that includes domain controllers running Windows 2000, be sure to read Upgrades in a domain containing Windows 2000 domain controllers.
- If you are upgrading in a domain where all domain controllers run Windows NT 4.0, be sure to read Upgrades in a Windows NT 4.0 Domain.
- If you want to upgrade and then use the same applications as before, be sure to review applications information in Relnotes.htm (in the \Docs folder on the product CD). Also, for the most recent information on compatible applications for products in the Windows Server 2003 family, see the software compatibility section of the Microsoft Web site.

**Points to consider for a new installation**

- If you reformat your hard disk and then perform a new installation, the efficiency of your disk might improve (compared to not reformatting it). Reformatting also gives you the opportunity to modify the size or number of disk partitions, to make them match your requirements more closely.
- If you want to practice careful configuration management, for example for a server where high availability is important, you might want to perform a new installation on that server instead of an upgrade. This is especially true on servers on which the operating system has been upgraded several times in the past.
- It is possible to install more than one operating system. Setting up the computer this way, however, presents complexities because of file system issues. For more information, see Deciding Whether a Computer Will Contain More Than One Operating System.

  **Note**

  - Do not install Windows Server 2003 on a compressed drive unless the drive was compressed with the NTFS file system compression utility. Uncompress a DriveSpace or DoubleSpace volume before running Setup on it.
- If you used Windows NT 4.0 to create a volume set, mirror set, stripe set, or stripe set with parity, and you want to run

Setup for Windows Server 2003 on that computer, you must prepare the disk set first. For details, see Working with volume, mirror, or stripe sets or stripe sets with parity.

**Operating systems from which you can upgrade**

If you upgrade, Setup automatically installs Windows Server 2003 into the same folder as the currently installed operating system. You can upgrade to Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition, from the following versions of Windows:

- Windows NT Server 4.0 with Service Pack 5 or later
- Windows NT Server 4.0, Terminal Server Edition with Service Pack 5 or later
- Windows NT Server 4.0, Enterprise Edition with Service Pack 5 or later
- Windows 2000 Server
- Windows 2000 Advanced Server

**Notes**

- You can also upgrade from Windows Server 2003, Standard Edition to Windows Server 2003, Enterprise Edition.
- If you have a version of Windows NT earlier than Windows NT 4.0, you cannot upgrade directly to a product in the Windows Server 2003 family. You must first upgrade to Windows NT 4.0 and apply Service Pack 5 before upgrading to a product in the Windows Server 2003 family.
- If you have servers or client computers that run Windows NT 3.51, we recommend that you install or upgrade to a newer operating system on all these computers, or retire them from operation. If you have more than one domain, you must upgrade domain controllers running Windows NT 3.51 for reliable logon validation. In any case, upgrading or retiring computers running Windows NT 3.51 strengthens security and reduces the number of version differences between computers, simplifying management and troubleshooting.
- You cannot upgrade from an earlier version of a Windows server operating system to an Itanium-based or x64-based version of a Windows server operating system.

# Windows Server 2003, Datacenter Edition: Upgrade or new installation?

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003, Datacenter Edition: Upgrade or new installation?**
After you have reviewed System Requirements and Hardware Compatibility and Important files to review, and before you run Setup, you will need to determine whether to upgrade or to perform a new installation. *Upgrading* is replacing Windows 2000 Datacenter Server with Windows Server 2003, Datacenter Edition. In contrast, *installing* means completely removing the previous operating system, or installing a product in the Windows Server 2003 family on a disk or disk partition with no previous operating system.

If you have already decided whether to upgrade or to perform a new installation, see Checklist: Performing an upgrade or Checklist: Performing a new installation.

The following lists describe points to consider when deciding between an upgrade and a new installation.

**Points to consider for an upgrade**

- With an upgrade, configuration is simpler and your existing users, settings, groups, rights, and permissions are retained.
- With an upgrade, you do not need to reinstall files and applications. As with any major changes to the hard disk, however, we recommend that you back up the disk before beginning an upgrade.
- If you want to upgrade and then use the same applications as before, be sure to review applications information in Relnotes.htm (in the \Docs folder on the product CD). Also, for the most recent information on compatible applications for products in the Windows Server 2003 family, see the software compatibility section of the Microsoft Web site.

**Points to consider for a new installation**

- If you reformat your hard disk and then perform a new installation, the efficiency of your disk might improve (compared to not reformatting it). Reformatting also gives you the opportunity to modify the size or number of disk partitions, to make them match your requirements more closely.
- If you want to practice careful configuration management, for example, for a server where high availability is important, you might want to perform a new installation on that server instead of an upgrade. This is especially true on servers on which the operating system has been upgraded several times in the past.

# Upgrades in a domain containing Windows 2000 domain controllers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrades in a domain containing Windows 2000 domain controllers**

If you are upgrading servers in a domain containing Windows 2000 domain controllers, there are a few points to keep in mind. The following points are the most important:

- The first step in planning the upgrade of a particular server is to check the compatibility of the hardware with products in the Windows Server 2003 family. For more information, see Hardware compatibility.

- Before you upgrade the first domain controller in a domain where one or more domain controllers run Windows 2000, or add a domain controller running a product in the Windows Server 2003 family, you must prepare the domain (and the forest in which it is located) using a simple tool on the Setup CD for Windows Server 2003.

  Before preparing the domain (and the forest in which it is located), it is recommended that you apply Service Pack 2 or later to all domain controllers running Windows 2000.

- If you have Internet Information Services (IIS) 5.0 installed on a server, when you upgrade the operating system, IIS is upgraded to IIS 6.0. However, for application compatibility, IIS runs in IIS 5.0 isolation mode after the upgrade. For more information about IIS isolation modes, see the IIS 6.0 Help.

- The Internet Explorer Enhanced Security Configuration is enabled by default when you upgrade. The security settings in this configuration can help make your computer more secure by limiting its exposure to malicious Web sites. Therefore, with this enhanced level of security, you might find that some Web sites do not display correctly in Internet Explorer when you are browsing Internet and intranet Web sites. Also, you might be prompted to enter your credentials when accessing network resources, such as files in Universal Naming Convention (UNC) shared folders. You can easily change the enhanced security settings. For more information, see Internet Explorer Enhanced Security Configuration overview.

- As you complete domain controller upgrades and reduce the number of different operating system versions running on them, you can adjust the domain and forest functional levels appropriately. For more information about functional levels, see Domain and forest functionality and Raising domain and forest functional levels. For general information about Active Directory, see Active Directory.

If you have servers running Windows NT, also review the following points:

- Before running Setup, it is recommended that you review the file systems and partitions that exist on the server. You must have at least one NTFS partition on domain controllers. It is recommended that you use NTFS on all partitions on all the servers in the domain, because any FAT or FAT32 partition lacks many security features. For example, on FAT or FAT32 partitions, a shared folder can be protected only by the permissions set on the shared folder, not on individual files, and there is no software protection against local access to the partition. For more information, see Reformatting or converting a partition to use NTFS and Convert.

- If some of your domain controllers run Windows 2000 and some run Windows NT, it is recommended that you upgrade the Windows NT domain controllers as soon as is practical, to reduce the number of version differences between computers, simplify management and troubleshooting, and strengthen security.

- Before you begin an upgrade from Windows NT 4.0, you must apply Service Pack 5 or later.

- If you have servers or client computers that run Windows NT 3.51, it is recommended that you install or upgrade to a newer operating system on all these computers, or retire them from operation. If you have more than one domain, you must upgrade domain controllers running Windows NT 3.51 for reliable logon validation. In any case, upgrading or retiring computers running Windows NT 3.51 strengthens security and reduces the number of version differences between computers, simplifying management and troubleshooting.

After the domain and forest are prepared, there are two stages in the upgrade of a domain controller. First you run Setup to upgrade the operating system and then, when Setup is complete, you run the Active Directory Installation Wizard (which appears on the screen the first time you log on).

# Upgrades in a Windows NT 4.0 Domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrades in a Windows NT 4.0 domain**
This section provides information about upgrading servers in a Windows NT 4.0 domain, that is, a domain where all domain controllers run Windows NT 4.0. The guidelines provided are for a network with one to five servers and 100 or fewer clients (the size of network discussed in these Getting Started topics). For information about upgrading servers in a domain where one or more domain controllers run Windows 2000, see Upgrades in a domain containing Windows 2000 domain controllers.

The following list describes the topics about upgrading servers in a Windows NT 4.0 domain:

- For information about initial preparations such as checking hardware compatibility and reviewing the size of partitions on server hard disks, see Basic preparations for the upgrade of servers in a Windows NT domain.
- For more information about planning steps that help you tailor the upgrade to your specific situation, see the following topics:
  - Planning server roles and the order of server upgrades from Windows NT
  - Planning the order of server upgrades in a Windows NT 4.0 domain
  - Planning DNS for the upgrade of domain controllers running Windows NT 4.0
- For more information about the final preparations to make just before beginning an upgrade, and guidelines about choices to make during the upgrade of domain controllers, see Performing final preparations and carrying out upgrades of domain controllers running Windows NT 4.0.

**Important**

- If you have servers or client computers that run Windows NT 3.51, it is recommended that you install or upgrade to a newer operating system on all these computers, or retire them from operation. If you have more than one domain, you must upgrade domain controllers running Windows NT 3.51 for reliable logon validation. In any case, upgrading or retiring computers running Windows NT 3.51 strengthens security and reduces the number of version differences between computers, simplifying management and troubleshooting.

**Note**

- The Internet Explorer Enhanced Security Configuration is enabled by default when you upgrade. The security settings in this configuration can help make your computer more secure by limiting its exposure to malicious Web sites. Therefore, with this enhanced level of security, you might find that some Web sites do not display correctly in Internet Explorer when you are browsing Internet and intranet Web sites. Also, you might be prompted to enter your credentials when accessing network resources, such as files in Universal Naming Convention (UNC) shared folders. You can easily change the enhanced security settings. For more information, see Internet Explorer Enhanced Security Configuration overview.

For additional information, see:

- For more information about Active Directory in relation to an upgrade from Windows NT, see Upgrading from a Windows NT domain and Domain and forest functionality.
- For information about deployments and upgrades in multiple domains or in domains larger than five servers, see the *Windows Server 2003 Deployment Kit* at the Microsoft Windows Server System Web site. For information about obtaining the *Windows Server 2003 Resource Kit*, see Using the Windows Deployment and Resource Kits.

# Basic preparations for the upgrade of servers in a Windows NT domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Basic preparations for the upgrade of servers in a Windows NT domain**

There are a small number of basic preparations you must carry out to lay the groundwork for upgrading servers in a domain where all the domain controllers run Windows NT 4.0. For information about how these preparations fit into the overall process of the upgrade, see Upgrades in a Windows NT 4.0 Domain.

**Important**

- If you have servers or client computers that run Windows NT 3.51, it is recommended that you install or upgrade to a newer operating system on all these computers, or retire them from operation. If you have more than one domain, you must upgrade domain controllers running Windows NT 3.51 for reliable logon validation. In any case, upgrading or retiring computers running Windows NT 3.51 strengthens security and reduces the number of version differences between computers, simplifying management and troubleshooting.

**Compatibility check**

Check to see whether the applications and hardware on your systems are compatible with products in the Windows Server 2003 family, and then update or replace them as necessary. For important details about how to do this, see Hardware compatibility and Important files to review.

**Review partition sizes, especially on domain controllers**

Review the partitions on the servers to make sure there is plenty of free space available for the Windows Server 2003 family operating system to run. It is especially important to allow plenty of free space on domain controllers, because the existing user accounts database can expand by as much as a factor of ten when Active Directory functionality is added. For information about the disk requirements for the operating system, see Windows Server 2003, Enterprise Edition: System requirements.

**Review file systems on server partitions**

You must have at least one NTFS partition on domain controllers. It is recommended that you use NTFS on all partitions on all the servers in the domain, because any FAT or FAT32 partition lacks many security features. For example, on FAT or FAT32 partitions, a shared folder can be protected only by the permissions set on the shared folder, not on individual files, and there is no software protection against local access to the partition. For more information, see Reformatting or converting a partition to use NTFS and Convert.

**Service Pack**

Obtain Service Pack 5 or later, because it is necessary for the upgrade.

**Mirror sets or other disk sets**

If you used Windows NT 4.0 to create a volume set, mirror set, stripe set, or stripe set with parity, and you want to run Setup for Windows Server 2003 on that computer, you must prepare the disk set first. For details, see Working with volume, mirror, or stripe sets or stripe sets with parity.

# Planning server roles and the order of server upgrades from Windows NT

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning server roles and the order of server upgrades from Windows NT**

This section provides a list of basic guidelines to use when planning server roles and the order of server upgrades in a domain where all the domain controllers run Windows NT 4.0. For information about how this planning fits into the overall process of the upgrade, see Upgrades in a Windows NT 4.0 Domain.

Use the following guidelines when planning server roles and the order of server upgrades from Windows NT:

- Familiarize yourself with the slight differences between the naming of server roles for servers running Windows NT as compared to servers running Windows Server 2003:

  - In Windows NT, there are two kinds of domain controllers, the primary domain controller, which is limited to one per domain and contains a read-write database, and backup domain controller, which is not limited in number and contains a read-only database.

  - For domain controllers running Windows Server 2003, there is only one kind of domain controller, without a "primary" or "backup" designation. All domain controllers contain matching copies of the user accounts and other Active Directory data in a given domain, and provide read-write access to that data.

  - In addition to domain controllers, the possible roles for either operating system include member server, which belongs to a domain but does not contain a copy of the Active Directory data, and stand-alone server, which belongs to a workgroup instead of a domain.

- Note that in contrast with Windows NT, you can change the role of a server without rerunning Setup. However, as a general practice, it is best to plan the roles ahead of time and change them only as necessary.

- As was true with Windows NT, a domain must have at least one domain controller, but it is recommended that a domain have multiple domain controllers for resilience in the handling of logon requests and directory updates.

- If the features you need most are not Active Directory features, you can focus on upgrading member servers first. For information about reasons for upgrading a particular member server early in the process, see Planning the order of server upgrades in a Windows NT 4.0 domain.

- If the features you need most are Active Directory features, focus on upgrading domain controllers first, and carry out the advance planning that is necessary for this part of the upgrade.

  **Important**

  - When you begin the upgrade of domain controllers running Windows NT 4.0, you must upgrade the primary domain controller first.

  For information about the features that come with Active Directory, see Planning the order of server upgrades in a Windows NT 4.0 domain. For information about the planning necessary before upgrading to Active Directory, see Planning DNS for the upgrade of domain controllers running Windows NT 4.0.

- When you begin upgrading domain controllers, if you have a remote access server that is a *member* server, it is recommended that you upgrade it before the last domain controller is upgraded. Such a server is dependent on domain controllers for user information, and therefore should not be "left behind" when domain controllers are upgraded. As an alternative, you can weaken the security permissions in Active Directory so that the remote access server running Windows NT can read user attributes from the domain controllers running products in the Windows Server 2003 family. For more information, see the *Windows Server 2003 Deployment Kit* at the Microsoft Windows Server System Web site.

For information about obtaining the *Windows Server 2003 Resource Kit*, see Using the Windows Deployment and Resource Kits.

# Planning the order of server upgrades in a Windows NT 4.0 domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning the order of server upgrades in a Windows NT 4.0 domain**
This topic lists the features available with the upgrade of any server running Windows NT 4.0 and also lists the features available with the upgrade of the first domain controller from Windows NT 4.0. Plan the order of server upgrades according to the features you are most interested in.

The following topics provide additional information:

- For information about other features in products in the Windows Server 2003 family, see New Features.
- For more information about upgrading servers in a domain where all domain controllers run Windows NT 4.0, see Upgrades in a Windows NT 4.0 Domain.
- For information about Active Directory, see Active Directory.

**Features available with the upgrade of any server in a Windows NT domain**
If you are interested in the features in the following list, you might want to focus on upgrading the member servers in your Windows NT 4.0 domain. For information about reasons to upgrade the domain controllers, see "Features Available with the First Domain Controller Upgrade in a Windows NT Domain," later in this topic.

- **Management tools**:

  Microsoft Management Console

  Plug and Play

  Device Manager

  Add Hardware Wizard (in Control Panel)

  Support for universal serial bus

  Enhanced Backup utility

- **File system support**: The version of NTFS in the Windows Server 2003 family includes support for disk quotas, the ability to defragment directory structures, and compressed network I/O.
- **Application services**: Win32 Driver Model, DirectX 7.0, and Windows Script Host.
- **Printer protocol support**: Device and protocol support allowing choices from more than 2,500 different printers. Other printing enhancements are included; for example, Internet Printing Protocol support, which allows users to print directly to a URL over an intranet or the Internet.
- **Scalability and availability**: Improved symmetric multiprocessor support as compared to Windows NT.
- **Security**: Encrypting file system.

**Features available with the first domain controller upgrade in a Windows NT domain**
If you are interested in the features in the following list, all of which are features of Active Directory, focus on upgrading the primary domain controller, as opposed to member servers. (The first domain controller that you upgrade must be the primary domain controller. It is recommended that you also upgrade backup domain controllers.) For information about the DNS

planning necessary before beginning the upgrade of domain controllers running Windows NT, see Planning DNS for the upgrade of domain controllers running Windows NT 4.0.

**Note**

- When a domain controller is upgraded, it has all the features of the Windows Server 2003 family operating system, which means it has all the features in the following list, plus the features in "Features available with the upgrade of any server in a Windows NT Domain," earlier in this topic.

When you upgrade a domain controller to run the Windows Server 2003 family operating system instead of Windows NT, many Active Directory features become available in the domain:

- **Simplified management** of network-resource information and user information.
- **Group Policy**, which you can use to set policies that apply across a given site, domain, or organizational unit in Active Directory.
- **Security and authentication features**, including support for Kerberos V5, Secure Sockets Layer (SSL) 3.0, and Transport Layer Security (TLS) using X.509v3 certificates.
- **Directory consolidation**, through which you can organize and simplify the management of users, computers, applications, and devices, and make it easier for users to find the information they want. You can take advantage of synchronization support through interfaces based on LDAP, and work with directory consolidation requirements specific to your applications.
- **Directory-enabled applications and infrastructure**, which make it easier to configure and manage applications and other directory-enabled network components.
- **Scalability without complexity**, a result of Active Directory scaling to millions of objects per domain and using indexing technology and advanced replication techniques to speed performance.
- **Use of Internet standards**, including access through LDAP and a namespace based on the Domain Name System (DNS).
- Active Directory Service Interfaces (ADSI), a powerful development environment.
- **Additional features**, as described in Features for Active Directory.

# Planning DNS for the upgrade of domain controllers running Windows NT 4.0

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning DNS for the upgrade of domain controllers running Windows NT 4.0**
The Domain Name System (DNS) is a hierarchical naming system used for locating computers on the Internet and private TCP/IP networks. This topic describes the planning for DNS that is recommended before you upgrade to Active Directory, that is, before you upgrade the primary domain controller in a Windows NT domain to become a domain controller running Windows Server 2003.

DNS is the name resolution service used by Active Directory. Products in the Windows Server 2003 family can support both the Windows Internet Name Service (WINS) common in Windows NT networks, and DNS. You must deploy DNS to support Active Directory; for information, see Active Directory integration. You can choose to migrate from WINS to DNS or configure your new DNS infrastructure to support your current WINS infrastructure.

For the upgrading of a network with one to five servers and 100 or fewer clients (which is the size of network discussed in these Getting Started topics), only a simple DNS plan is needed, because there will be no parent or child domains. For information about upgrading networks with more than five servers, see the *Windows Server 2003 Deployment Kit* at the Microsoft Windows Server System Web site. For information about obtaining the *Windows Server 2003 Resource Kit*, see Using the Windows Deployment and Resource Kits.

DNS names consist of name labels separated by periods. A simple DNS domain name for a Windows Server 2003 family domain could take the form *domainname.ext* where *ext* is an established "top-level domain" type such as com or org. With a domain name of that form, the form of a computer name in the domain would be:

*computername.domainname.ext*

When you upgrade the primary domain controller in a Windows NT 4.0 domain (the primary domain controller must be upgraded first, before backup domain controllers), you will be offered several different options for the handling of DNS. If there is no DNS server available to work with Active Directory, you will be offered, by default, the option of installing DNS on the domain controller you are upgrading (formerly the primary domain controller). In a network with one to five servers, the simplest approach is to accept this default. Then, after completing the upgrade of the server on which DNS was installed, take note of its static IP address, and configure other computers to send any DNS requests to that IP address.

See the following for more information:

- For additional basic information about DNS and how to configure it during Setup, see Name resolution for TCP/IP.
- For more detailed information about DNS, see DNS.
- For specific information about namespace planning for DNS, see Namespace planning for DNS.

# Performing final preparations and carrying out upgrades of domain controllers running Windows NT 4.0

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Performing final preparations and carrying out upgrades of domain controllers running Windows NT 4.0**
During the final preparations before beginning an upgrade from Windows NT 4.0, you can remove a backup domain controller temporarily from the network to protect your existing network integrity. During the upgrade, there are a few choices to make for the upgrade of a network of two to five servers (a network within the size of those discussed in these Getting Started topics). This topic describes the removal of a backup domain controller from the network, and provides guidelines for making choices during the upgrade.

Before beginning this stage of the upgrade process, be sure to review the important information in other topics listed in Upgrades in a Windows NT 4.0 Domain.

**Removing a backup domain controller from the network**
Before beginning an upgrade of servers in a domain where all the domain controllers run Windows NT 4.0, you can perform an additional precautionary step at the final preparatory stage, when you are confirming that Service Pack 5 or later has been applied, and you are performing server backups.

If you are preparing to upgrade the primary domain controller, to protect your existing network integrity, consider removing a backup domain controller temporarily from the network. To do this, in your existing Windows NT network, choose a backup domain controller, ensure that it has a current copy of the user accounts database, and back it up. Then disconnect its network cable. After the upgrade of your primary domain controller to become a domain controller running Windows Server 2003 (you must upgrade the primary domain controller first), this disconnected system is available for promotion to a Windows NT primary domain controller if necessary. However, in the course of an uneventful upgrade, you would not and could not promote the Windows NT backup domain controller to primary domain controller. Instead, you would continue the upgrade process, eventually reconnecting the disconnected server and upgrading it.

**Running the Active Directory Installation Wizard after upgrading the primary domain controller**
To upgrade the primary domain controller to become a domain controller running Windows Server 2003, you must first run Setup to upgrade the operating system. After the operating system is installed, you can run the Active Directory Installation Wizard (which appears on the screen the first time you log on). This section provides guidelines for making choices in the Active Directory Installation Wizard when upgrading the primary domain controller.

**Important**

- When you begin the upgrade of domain controllers running Windows NT 4.0, you must upgrade the primary domain controller first.

The following table shows the choices that are recommended when running the Active Directory Installation Wizard on the server that was formerly the primary domain controller. These recommendations are for a network of two to five servers (a network within the size of those discussed in these Getting Started topics):

| **Choices on the server that was formerly the primary domain controller** | **Recommended action for a network of two to five servers** |
|---|---|
| **Domain in a new forest** or **Child domain in an existing domain tree** or<br>**Domain tree in an existing forest** | Choose **Domain in a new forest**. |
| Forest functional level of **Windows Server 2003 interim**<br><br>(This choice appears only if **Domain in a new forest** was previously selected.) | Choose **Windows Server 2003 interim**.<br><br>With this domain functional level, you can have domain controllers running Windows NT 4.0 or Windows Server 2003, but no domain controllers running Windows 2000. |

During the upgrade, you can choose the location of three important items: the database containing user accounts and other Active Directory data, the log file, and the system volume (Sysvol) folder. The database and the log file can be on any kind of

partition (FAT, FAT32, or NTFS); the database can expand by as much as a factor of ten from the size it had with Windows NT, so allow plenty of room for it. (Initially, the log file will take up very little space.) The system volume folder must be on an NTFS partition. (For information about NTFS and other file systems, see Choosing a File System for the Installation Partition.)

**Note**

- When you begin upgrading domain controllers, if you have a remote access server that is a *member* server, it is recommended that you upgrade it before the last domain controller is upgraded. For more information, see Planning server roles and the order of server upgrades from Windows NT.

For additional information about upgrades in relation to Active Directory, see Domain and forest functionality and Upgrading from a Windows NT domain. For a general list of topics on Active Directory, see Active Directory.

# Deciding Whether a Computer Will Contain More Than One Operating System

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deciding whether a computer will contain more than one operating system**


- Considerations for computers that contain more than one operating system
- Multiple operating systems and file system compatibility
- Computers that contain multiple Windows Server 2003 family partitions
- Computers that contain Windows 2000 and Windows Server 2003
- Computers that contain Windows NT 4.0 and Windows Server 2003
- Computers that contain Windows 98 or Windows Millennium Edition and Windows Server 2003
- Computers that contain MS-DOS and Windows Server 2003

# Considerations for computers that contain more than one operating system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Considerations for computers that contain more than one operating system**

On a computer with an appropriate disk configuration (outlined in the table later in this topic), you can install more than one operating system, and then choose between the operating systems each time you restart the computer.

**Note**

- On an Itanium architecture-based computer, you can install Windows XP 64-bit Edition (Itanium) and the Itanium-based versions of Windows Server 2003. In this topic and other topics about installing multiple operating systems, information about operating systems that you cannot install does not apply. For example, information about Windows 2000 or Windows NT 4.0 does not apply to an Itanium architecture-based computer.

For example, on an x86-based or x64-based computer, you could set up a server to run Windows Server 2003, Enterprise Edition, most of the time, but allow it to sometimes run Windows NT Server 4.0 Enterprise Edition in order to support an older application. (However, to do this you would need to make specific file system choices and would need Service Pack 5 or later, as described in Multiple operating systems and file system compatibility.) During restarts, a display would appear for a specified number of seconds, allowing you to select between the two operating systems. (You can specify a default operating system that will run if no selection is made during the restart process.)

The following table shows the disk configurations on which you can install more than one operating system.

**Important**

- You must follow the requirements in the following table. For example, on a basic disk, you must install each operating system in a separate partition. This ensures that each operating system does not overwrite crucial files that are needed by another operating system.

| Disk configuration | Requirements for multiple operating systems |
|---|---|
| Basic disk or disks | You can install multiple operating systems, including Windows NT 4.0 and earlier operating systems, on a basic disk. Each operating system must be on a separate partition or logical drive on the disk. A partition or logical drive is a section of the disk that functions as a separate unit. Different partitions often have different drive letters; for example, C: and D:. |
| Single dynamic disk | You can install only one operating system.<br><br>However, if you used Windows 2000 or Windows XP to change a disk with no partitions directly to a dynamic disk, you must revert the disk to basic before you can install an operating system on it. For more information, see the Note after this table. |
| Multiple dynamic disks | Each dynamic disk can contain one installation of Windows 2000, Windows XP, or a product in the Windows Server 2003 family. No other operating systems can start from a dynamic disk.<br><br>However, if you used Windows 2000 or Windows XP to change a disk with no partitions directly to a dynamic disk, you must revert the disk to basic before you can install an operating system on it. For more information, see the Note after this table. |
| Master boot record (MBR) disk on an Itanium architecture-based computer | You cannot start operating systems from an MBR disk on an Itanium architecture-based computer. You must use a GPT disk for this purpose. |

| | |
|---|---|
| GUID partition table (GPT) disk on an Itanium architecture-based computer | You can install one or more operating systems on a GPT disk on an Itanium architecture-based computer. The guidelines in this table for basic and dynamic disks apply to GPT disks on Itanium architecture-based computers.<br><br>On an Itanium architecture-based computer, the operating systems that you can install include Windows XP 64-bit Edition (Itanium) and the Itanium-based versions of Windows Server 2003. On an Itanium architecture-based computer, you cannot install earlier operating systems, such as Windows 2000. |

**Note**

- If you used Windows 2000 or Windows XP to change a disk with no partitions directly to a dynamic disk, you must revert the disk to basic before you can install an operating system on it. All data will be lost in this process, so back it up first. You can then use Windows 2000 or Windows XP to revert the disk, or you can use Windows Server 2003 Setup. To use Windows 2000 or Windows XP, follow the instructions in Help for your operating system. To use Windows Server 2003 Setup, during the partitioning phase, find the dynamic disk among the list of available partitions and then delete it (which erases all data on all volumes). You will be prompted to confirm your action. After you do this, the disk will contain only unpartitioned space, and you can use Setup to create a new (basic disk) partition on it.

**Reasons to install only one operating system**

Setting up a computer so that you can choose between two or more operating systems at startup does have an advantage: it allows you to use applications that run only with a particular operating system. However, there are definite reasons to install only one operating system:

- Each operating system uses valuable disk space.
- Compatibility issues, especially file system compatibility, can be complex. For more information, see Multiple operating systems and file system compatibility.
- On a dynamic disk (a storage type available with products in the Windows Server 2003 family), you can have only one operating system per disk. Dynamic disks also will not work with some operating systems. For more information, see the table earlier in this topic.
- It is no longer necessary to maintain multiple operating systems as a safeguard against problems with starting the computer. With products in the Windows Server 2003 family, you have other options for system recovery. For example, if you have a problem with a newly installed device driver, you can use Safe Mode, in which Windows Server 2003 restarts with default settings and the minimum number of drivers. For more information about Safe Mode and other options, see Understanding Disaster Recovery.

**Specific details to consider**

Before you decide to set up a computer with more than one operating system, review the topics that apply to your situation:

| Operating system that can alternate with Windows Server 2003 | Topics to review |
|---|---|
| Another product in the Windows Server 2003 family | • Computers that contain multiple Windows Server 2003 family partitions |
| Windows 2000 | • Computers that contain Windows 2000 and Windows Server 2003 |
| Windows NT 4.0 | • Multiple operating systems and file system compatibility<br>• Computers that contain Windows NT 4.0 and Windows Server 2003 |
| Windows 98 or Windows Millennium Edition | • Computers that contain Windows 98 or Windows Millennium Edition and Windows Server 2003 |

| | |
|---|---|
| MS-DOS | • Computers that contain MS-DOS and Windows Server 2003 |

# Multiple operating systems and file system compatibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multiple operating systems and file system compatibility**

On computers that contain multiple operating systems, compatibility becomes more complex when you consider file system choices. The file systems to choose from are NTFS, FAT, and FAT32. (For more information, see Choosing a File System for the Installation Partition.)

NTFS is normally the recommended file system because it is more efficient and reliable, and supports important features including Active Directory and domain-based security. With NTFS, however, you need to take file system compatibility into account when considering whether to set up a computer to contain more than one operating system, because with Windows 2000 and the Windows Server 2003 family, NTFS has new features in addition to those in Windows NT. Files that use any new features will be completely usable or readable *only* when the computer is started with Windows 2000 or a product in the Windows Server 2003 family.

For example, a file that uses the new encryption feature will not be readable when the computer is started with Windows NT Server 4.0 or Windows NT 4.0 Server Enterprise Edition, which were released before the encryption feature existed.

For more information about features that affect file accessibility with products in the Windows Server 2003 family, see NTFS.

**Notes**

- If you want to set up a computer with both Windows NT and Windows Server 2003 operating systems, and you want to have an NTFS partition, the only appropriate version of Windows NT is version 4.0 with the latest released Service Pack. Using the latest Service Pack maximizes compatibility between Windows NT 4.0 and Windows Server 2003. (Specifically, you must have Service Pack 5 or later.) Even the latest Service Pack, however, does not provide access to files using the new features in NTFS.

- Using NTFS as the only file system on a computer that contains both Windows Server 2003 and Windows NT is *not* recommended. On these computers, a FAT partition containing the Windows NT 4.0 operating system ensures that when started with Windows NT 4.0, the computer will have access to needed files. In addition, if Windows NT is not installed on the system partition, which is almost always the first partition on the disk, it is recommended that the system partition also be formatted with FAT.

- If you set up a computer so that it starts with Windows NT 3.51 or earlier on a FAT partition, and Windows Server 2003 on an NTFS partition, when that computer starts with Windows NT 3.51, the NTFS partition will not be visible. If you set up a computer this way, and the partition containing Windows NT 3.51 is not the system partition (which is almost always the first partition on the disk), the system partition must also be formatted with FAT.

# Computers that contain multiple Windows Server 2003 family partitions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computers that contain multiple Windows server 2003 family partitions**

You can set up a server so that it contains some combination of Windows XP and products in the Windows Server 2003 family, or contains multiple installations of products in the Windows Server 2003 family. If you plan to do this, note the following:

- Install each operating system on a different partition or, for dynamic disks, on a different disk, and install the applications used with an operating system on the same disk or partition with it. If an application is used with two different operating systems, install it in two places.

- For an x86-based computer, choose any product in the Windows Server 2003 family for installation on a specific partition. For example, you could install Windows Server 2003, Standard Edition, in one location and Windows Server 2003, Enterprise Edition, in another.

  For an x64-based computer, you can choose among Windows XP Professional x64 Edition; the x64-based version of Windows Server 2003, Standard Edition; the x64-based version of Windows Server 2003, Enterprise Edition; and the x64-based version of Windows Server 2003, Datacenter Edition. In addition, you can choose to install any x86-based operating system, such as Microsoft® Windows® XP.

  For an Itanium architecture-based computer, you can choose among Windows XP 64-bit Edition (Itanium); the Itanium-based version of Windows Server 2003, Enterprise Edition; and the Itanium-based version of Windows Server 2003, Datacenter Edition.

- If the computer participates in a domain, use a different computer name for each installation. Because a unique security identifier (SID) is used for each installation of Windows Server 2003 on a domain, the computer name for each installation must be unique, even for multiple installations on the same computer.

- If you set up a server so that it contains some combination of Windows Server 2003 and Windows 2000 or Windows XP, or contains multiple partitions with products in the Windows Server 2003 family, and you want to use the Encrypting File System (EFS) on the computer, you must take certain steps. These steps make encrypted files readable between the different installations.

  - One approach is to ensure that all the installations are in the same domain and that the user of these installations has a roaming profile.

  - Another approach is to export the user's file encryption certificate and associated private key from one installation and import it into the other installations.

  For more information, see Introduction to using Encrypting File System, Using roaming user profiles, and Importing and exporting certificates.

# Computers that contain Windows 2000 and Windows Server 2003

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computers that contain Windows 2000 and Windows Server 2003**

You can set up a server so that it contains both Windows 2000 and Windows Server 2003. If you plan to do this, note the following:

- Install each operating system on a different partition or, for dynamic disks, on a different disk, and install the applications used with an operating system on the same disk or partition with it. If an application is used with two different operating systems, install it in two places.

- Windows Server 2003 must be installed last. Otherwise, important files that are needed for starting Windows Server 2003 will be overwritten.

- If the computer participates in a domain, use a different computer name for each installation. Because a unique security identifier (SID) is used for each installation on a domain, the computer name for each installation must be unique, even for multiple installations on the same computer.

- If you use the Encrypting File System (EFS) on a computer containing Windows 2000 and a product in the Windows Server 2003 family, the encrypted files will not be readable between the different installations unless you take one of the following approaches:

  - One approach is to ensure that all the installations are in the same domain and that the user of these installations has a roaming profile.

  - Another approach is to export the user's file encryption certificate and associated private key from one installation and import it into the other installations.

  For more information, see Introduction to using Encrypting File System, Using roaming user profiles, and Importing and exporting certificates.

# Computers that contain Windows NT 4.0 and Windows Server 2003

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computers that contain Windows NT 4.0 and Windows Server 2003**

If you plan to set up a computer so that it contains Windows NT 4.0 and Windows Server 2003 operating systems, first review the following precautions:

- It is no longer necessary to maintain multiple operating systems as a safeguard against problems with starting the computer. With products in the Windows Server 2003 family, you have other options for system recovery. For example, if you have a problem with a newly installed device driver, you can use Safe Mode, in which Windows Server 2003 restarts with default settings and the minimum number of drivers. For more information about Safe Mode and other options, see Understanding Disaster Recovery.

- Using NTFS as the only file system on a computer that contains both Windows Server 2003 and Windows NT is *not* recommended. For more information, see Multiple operating systems and file system compatibility.

- Make sure that Windows NT 4.0 has been updated with the latest released Service Pack. For details, see Multiple operating systems and file system compatibility.

- Install each operating system on a different partition, and install the applications used with an operating system on the same partition with it. If an application is used with two different operating systems, install it on two partitions.

- Do not install Windows Server 2003 on a compressed drive unless the drive was compressed with the NTFS file system compression feature.

- Windows Server 2003 must be installed last. Otherwise, important files that are needed for starting Windows Server 2003 will be overwritten.

- If the computer participates in a domain, use a different computer name for each installation.

# Computers that contain Windows 98 or Windows Millennium Edition and Windows Server 2003

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computers that contain Windows 98 or Windows Millennium Edition and Windows Server 2003**
If you plan to set up a computer so that it contains Windows 98 or Windows Millennium Edition and Windows Server 2003, note the following:

- Install each operating system on a different partition, and install the applications used with an operating system on the same partition with it. If an application is used with two different operating systems, install it on two partitions.
- Windows 98 or Windows Millennium Edition must be installed on a basic disk on a partition formatted with FAT or FAT32. If either Windows 98 or Windows Millennium Edition is not installed on the system partition, which is almost always the first partition on the disk, the system partition must also be formatted with FAT or FAT32.
- Windows Server 2003 must be installed last. Otherwise, important files that are needed for starting Windows Server 2003 will be overwritten.
- File system compatibility might be an issue. See Multiple operating systems and file system compatibility.

# Computers that contain MS-DOS and Windows Server 2003

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computers that contain MS-DOS and Windows Server 2003**

If you plan to set up a computer so that it contains MS-DOS and Windows Server 2003, note the following:

- Install each operating system on a different partition, and install the applications used with an operating system on the same partition with it. If an application is used with two different operating systems, install it on two partitions.
- MS-DOS must be installed on a basic disk on a partition formatted with FAT. If MS-DOS is not installed on the system partition, which is almost always the first partition on the disk, the system partition must also be formatted with FAT.
- Windows Server 2003 must be installed last. Otherwise, important files that are needed for starting Windows Server 2003 will be overwritten.
- File system compatibility might be an issue. See Multiple operating systems and file system compatibility.

# Choosing a File System for the Installation Partition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing a file system for the installation partition**

- Choosing a file system: NTFS, FAT, or FAT32
- Reformatting or converting a partition to use NTFS
- NTFS compared to FAT and FAT32
- NTFS

# Choosing a file system: NTFS, FAT, or FAT32

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing a file system: NTFS, FAT, or FAT32**
You can choose among three file systems for an installation partition: NTFS, FAT, and FAT32. NTFS is strongly recommended in most situations.

For details about each file system or the process of converting or reformatting a partition, see:

- Reformatting or converting a partition to use NTFS
- NTFS compared to FAT and FAT32
- NTFS

**Note**

- You can use important features such as Active Directory and domain-based security only by choosing NTFS as your file system.

**Important**

- On globally unique identifier (GUID) partition table disks (GPT disks), it is strongly recommended that you use NTFS for the installation partition. However, if you have an Itanium architecture-based computer, and you see that it has a small FAT partition of 100 megabytes (MB) or more, do not delete or reformat this partition. The partition is required for the loading of the operating system. For more information, see Start Setup on an Itanium architecture-based computer.

The following table lists a number of installation or upgrade scenarios for x86-based and x64-based computers (the last two scenarios are fairly uncommon) and provides file system guidelines for each one:

| Installation or upgrade scenario for an x86-based or x64-based computer | File system to use and additional information to read |
|---|---|
| The computer currently uses NTFS only (no FAT or FAT32). | Continue to use NTFS.<br><br>No additional information about file systems needed. |
| The computer is x86-based and has one or more FAT or FAT32 partitions.<br><br>AND<br><br>The computer contains only one operating system, or the operating systems on the computer include Windows 2000, Windows XP, or a product in the Windows Server 2003 family, but no other operating systems.<br><br>If the computer is Itanium architecture-based, see the important note earlier in this topic. | Consider reformatting or converting partitions so all partitions use NTFS.<br><br>For more information, see Reformatting or converting a partition to use NTFS. |
| The computer is x64-based. | Use NTFS. |

| | |
|---|---|
| The computer will contain multiple operating systems, one of which is MS-DOS, Windows 95, Windows 98, or Windows Millennium Edition. | For any partition that must be accessible from MS-DOS, Windows 95, Windows 98, or Windows Millennium Edition, use FAT (or when appropriate, FAT32).<br><br>For more information, see Deciding Whether a Computer Will Contain More Than One Operating System. |
| The computer will contain multiple operating systems, one of which is Windows NT. | Read Multiple operating systems and file system compatibility. |

# Reformatting or converting a partition to use NTFS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Reformatting or converting a partition to use NTFS**

If you have a FAT or FAT32 partition on which you want to install a product in the Windows Server 2003 family, and you want to use NTFS instead, you have two choices:

- You can convert the FAT or FAT32 partition to NTFS. This leaves files intact, although the partition may have somewhat more fragmentation and slower performance than a partition formatted with NTFS. However, it is still advantageous to use NTFS, regardless of whether the partition was formatted with NTFS or converted.

  If you install a product in the Windows Server 2003 family on a FAT or FAT32 partition, you are offered the option to convert the partition to NTFS. You can also convert a FAT or FAT32 partition after Setup by using Convert.exe. For more information about Convert.exe, see Convert.

- You can reformat the partition with NTFS. This erases all files on the partition, but results in less fragmentation and better performance than with a converted partition.

**Important**

- If you have an Itanium architecture-based computer and you see that it has a small FAT partition of 100 MB or more, do not delete or reformat this partition. The partition is required for the loading of the operating system. For more information, see Start Setup on an Itanium architecture-based computer.

If you format a partition during Setup, the file systems choices are listed as NTFS and FAT. The following table provides information about the relationship between partition size and file system choices during Setup.

| State and size of partition | Setup choices and responses (when formatting the partition) |
|---|---|
| Unformatted, less than 2 GB. | Setup offers NTFS or FAT. Setup uses the format chosen. |
| Unformatted, 2 GB or larger, up to a maximum of 32 GB. | Setup offers NTFS or FAT. If FAT is chosen, Setup uses FAT32. |
| Unformatted, larger than 32 GB. | Setup allows only NTFS. |
| Previously formatted with FAT32 and larger than 32 GB. (Partition created with Windows 95, Windows 98, or Windows Millennium Edition.) | No formatting needed, even though an unformatted partition of this size, when formatted during or after Setup for a product in the Windows Server 2003 family, would have to use NTFS. In other words, previously-formatted FAT32 partitions of this size continue to be supported by the Windows Server 2003 family. |

If you format a partition during Setup, you can choose between a quick format and a full format:

**Quick format**

Quick format creates the file system structure on the disk without verifying the integrity of every sector. Choose this method for any disk that has no bad sectors and no history of file-corruption problems that might be related to bad sectors.

**Full format**

A full format identifies and tracks bad sectors so that they are not used for storing data. Choose this method for any disk that has bad sectors or has a history of file-corruption problems that might be related to bad sectors.

# NTFS compared to FAT and FAT32

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NTFS compared to FAT and FAT32**

NTFS has always been a more powerful file system than FAT and FAT32. Windows 2000, Windows XP, and the Windows Server 2003 family include a new version of NTFS, with support for a variety of features including Active Directory, which is needed for domains, user accounts, and other important security features. For more details about features in NTFS, see NTFS.

FAT and FAT32 are similar to each other, except that FAT32 is designed for larger disks than FAT. The file system that works most easily with large disks is NTFS.

The following table describes the compatibility of each file system with various operating systems.

**Note**

- File system choices have no effect on access to files across the network. For example, using NTFS on all partitions on a server does not affect clients connecting across a network to shared folders or shared files on that server, even if those clients run an earlier operating system such as Windows 98 or Windows NT.

| NTFS | FAT | FAT32 |
|---|---|---|
| A computer running Windows 2000, Windows XP, or a product in the Windows Server 2003 family can access files on a local NTFS partition. A computer running Windows NT 4.0 with Service Pack 5 or later might be able to access some files. Other operating systems allow no local access. | Access to files on a local partition is available through MS-DOS, all versions of Windows, and OS/2. | Access to files on a local partition is available only through Windows 95 OSR2, Windows 98, Windows Millennium Edition, Windows 2000, Windows XP, and products in the Windows Server 2003 family. |

The following table compares disk and file sizes possible with each file system.

| NTFS | FAT | FAT32 |
|---|---|---|
| Recommended minimum volume size is approximately 10 MB.<br><br>Maximum volume and partition sizes start at 2 terabytes (TB) and range upward. For example, a dynamic disk formatted with a standard allocation unit size (4 KB) can have partitions of 16 TB minus 4 KB. For more information about maximum volume and partition sizes, see the *Windows Server 2003 Resource Kit, Server Management Guide* at the Microsoft Windows Resource Kits Web site.<br><br>Cannot be used on floppy disks. | Volumes from floppy disk size up to 4 GB.<br><br>Does not support domains. | Volumes from 33 MB to 2 TB can be written to or read using products in the Windows Server 2003 family.<br><br>Volumes up to 32 GB can be formatted as FAT32 using products in the Windows Server 2003 family.<br><br>Does not support domains. |
| Maximum file size is potentially 16 TB minus 64 KB, although files cannot be larger than the volume or partition they are located on. | Maximum file size is 2 GB. | Maximum file size is 4 GB. |

**Notes**

On Itanium-based computers with more than one disk, on x64-based computers, and on x86-based computers running Windows Server 2003 with Service Pack 1 (SP1), you can choose not only the sizes of partitions but the partition style for each disk. The partition style determines the way that information about the partition is stored. There are two partition styles: master boot record (MBR) and GUID partition table (GPT). On GPT disks, unlike MBR disks, data critical to platform operation is located in partitions instead of unpartitioned or hidden sectors. In addition, GPT disks have redundant primary and backup partition tables for improved partition data structure integrity.

On Itanium-based computers, you must install Windows XP 64-bit Edition (Itanium) or the Itanium-based versions of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, on a GPT disk. On x86-based and x64-based computers running Windows Server 2003, you must install the operating system on an MBR disk. Other disks can use either the MBR or GPT partition style. With GPT disks, you can create more partitions and larger volumes and can take advantage of other benefits.

For more information about GPT disks, see GUID partition table. For more information about MBR and GPT partition styles, see Partition styles. For information about obtaining the *Windows Server 2003 Resource Kit*, see Using the Windows Deployment and Resource Kits.

# NTFS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NTFS**

Some of the features you can use when you choose NTFS are:

- Better scalability to large drives. The maximum partition or volume size for NTFS is much greater than that for FAT, and as volume or partition sizes increase, performance with NTFS does not degrade as it does with FAT.
- Active Directory (and domains, which are part of Active Directory). With Active Directory, you can view and control network resources easily. With domains, you can fine-tune security options while keeping administration simple. Domain controllers and Active Directory require NTFS.
- Compression features, including the ability to compress or uncompress a drive, a folder, or a specific file. (However, a file cannot be both compressed and encrypted at the same time.)
- File encryption, which greatly enhances security. (However, a file cannot be both compressed and encrypted at the same time.)
- Permissions that can be set on individual files rather than just folders.
- Remote Storage, which provides an extension to your disk space by making removable media such as tapes more accessible.
- Recovery logging of disk activities, which allows NTFS to restore information quickly in the event of power failure or other system problems.
- Sparse files. These are very large files created by applications in such a way that only limited disk space is needed. That is, NTFS allocates disk space only to the portions of a file that are written to.
- Disk quotas, which you can use to monitor and control the amount of disk space used by individual users.

This is only a partial list of the features in NTFS in the Windows Server 2003 family. For more information about new features, see New Features.

For information about maximum file sizes possible with NTFS, see NTFS compared to FAT and FAT32. For information about converting or reformatting a FAT or FAT32 partition, see Reformatting or converting a partition to use NTFS.

# Planning Disk Partitions or Volumes for New Installations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning disk partitions or volumes for new installations**


- Considerations when planning disk partitions
- Options when partitioning a disk
- Working with dynamic disks and Setup

# Considerations when planning disk partitions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Considerations when planning disk partitions**

You must plan your disk partitions before you run Setup only if both of the following conditions are true:

- You are performing a new installation, not an upgrade.
- The disk on which you are installing is a basic disk, not a dynamic disk. Basic disks are the disk type that existed before Windows 2000; most disks are basic disks. Dynamic disks are disks that once were basic but were changed to dynamic using Windows 2000, Windows XP, or a product in the Windows Server 2003 family. If you plan to install to a dynamic disk, you cannot change the volume or partition sizes on the disk during Setup, and therefore no planning is needed regarding partition sizes. Instead, review the guidelines in Working with dynamic disks and Setup.

Disk partitioning is a way of dividing your physical disk so that each section functions as a separate unit. When you create partitions on a basic disk, you divide the disk into one or more areas that can be formatted for use by a file system, such as FAT or NTFS. Different partitions often have different drive letters (for example, C: and D:). A basic disk can have up to four primary partitions, or three primary partitions and one extended partition. (An extended partition can be subdivided into logical drives, while a primary partition cannot be subdivided.)

**Important**

- If you plan to delete or create partitions on a hard disk, be sure to back up the disk contents beforehand, because these actions will destroy any existing data. As with any major change to disk contents, it is recommended that you back up the entire contents of the hard disk before working with partitions, even if you plan to leave one or more of your partitions alone.

Before you run Setup to perform a new installation, determine the size of the partition on which to install. There is no set formula for figuring a partition size. The basic principle is to allow plenty of room for the operating system, applications, and other files that you plan to put on the installation partition. The files for setting up Windows Server 2003 require approximately 1.25 GB to 3 GB on an x86-based computer and 3 GB to 4 GB on an Itanium architecture-based computer, as described in the topics about system requirements listed in System Requirements and Hardware Compatibility. It is recommended that you allow considerably more disk space than the minimum amount. It is not unreasonable to allow 4 GB to 10 GB on the partition, or more for large installations. This allows space for a variety of items, including optional components, user accounts, Active Directory information, logs, future service packs, the paging file used by the operating system, and other items.

For information about some of the items that require space on the disk, see the following:

- Change virtual memory settings
- Specify what happens when the system stops unexpectedly
- Choosing components to install
- Understanding Active Directory
- Directory data store

When you perform a new installation, you can specify the partition on which to install. If you specify a partition on which another operating system exists, you will be prompted to confirm your choice.

During Setup, create and size only the partition on which you want to install Windows Server 2003. After installation is complete, you can use Disk Management to manage new and existing disks and volumes. This includes creating new partitions from unpartitioned space; deleting, renaming, and reformatting existing partitions; adding and removing hard disks; and changing a basic disk to the dynamic disk storage type, or dynamic to basic. (If you want to have a dynamic disk on a computer that contains more than one operating system, be sure to read Considerations for computers that contain more than one

operating system.)

On Itanium-based computers with more than one disk, on x64-based computers, and on x86-based computers running Windows Server 2003 with Service Pack 1 (SP1), you can plan to choose not only the sizes of partitions but the partition style for each disk. The partition style determines the way that information about the partition is stored. There are two partition styles: master boot record (MBR) and GUID partition table (GPT). On GPT disks, unlike MBR disks, data critical to platform operation is located in partitions instead of unpartitioned or hidden sectors. In addition, GPT disks have redundant primary and backup partition tables for improved partition data structure integrity.

On Itanium-based computers, you must install Windows XP 64-bit Edition (Itanium) or the Itanium-based versions of Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, on a GPT disk. On x86-based and x64-based computers running Windows Server 2003, you must install Windows on an MBR disk. Other disks can use either the MBR or GPT partition style. With GPT disks, you can create more partitions and larger volumes, and can take advantage of other benefits. For more information about GPT disks, see GUID partition table. For more information about MBR and GPT partition styles, see Partition styles. For information about obtaining the *Windows Server 2003 Resource Kit*, see Using the Windows Deployment and Resource Kits.

**Important**

- If you are setting up a computer so that it contains multiple operating systems, you must install Windows Server 2003 on its own partition or logical drive. This ensures that Windows Server 2003 will not overwrite crucial files that are needed by the other operating system. For more information, see Deciding Whether a Computer Will Contain More Than One Operating System.

**Disk partition requirements for Remote Installation Services**

If you plan to use Remote Installation Services on this server so that you can install operating systems onto other computers, a separate partition for use by Remote Installation Services is necessary. Plan on using NTFS on this partition: NTFS is required for the Single Instance Store feature of Remote Installation Services.

If you need to create a new partition for Remote Installation Services, plan on doing it after Setup, and leave enough unpartitioned disk space so that you can create it (at least 4 GB of space is recommended). As an alternative, for the system disk (not cluster disks), you can plan to make the disk a dynamic disk, which allows more flexibility in the use of the disk space than a basic disk. (However, if you want to have a dynamic disk on a computer that contains more than one operating system, see Deciding Whether a Computer Will Contain More Than One Operating System.)

For information about requirements for Remote Installation Services, see Remote Installation Services system requirements. For additional information about disk and partition choices, see Disk Management. For an overview, checklists, and other information about Remote Installation Services, see Remote Installation Services.

# Options when partitioning a disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Options when partitioning a disk**

You can change the partitions on your disk during Setup *only* if you are performing a new installation, not an upgrade. You can modify the partitioning of the disk after Setup by using Disk Management.

If you are performing a new installation, Setup examines the hard disk to determine its existing configuration, and then offers the following options:

- If the hard disk is unpartitioned, you can create and size the partition on which you will install a product in the Windows Server 2003 family.
- If the hard disk is partitioned but has enough unpartitioned disk space, you can create the partition for your Windows Server 2003 family product by using the unpartitioned space.
- If the hard disk has an existing partition that is large enough, you can install a product in the Windows Server 2003 family on that partition, with or without reformatting the partition first. Reformatting a partition erases all data on the partition. If you do not reformat the partition but you do install a Windows Server 2003 family product where there is already an operating system, that operating system will be overwritten, and you must reinstall any applications you want to use with the Windows Server 2003 family product.
- If the hard disk has an existing partition, you can delete it to create more unpartitioned disk space for a partition for a Windows Server 2003 family product. Deleting an existing partition also erases any data on that partition.

# Working with dynamic disks and Setup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Working with dynamic disks and Setup**

A dynamic disk is a disk using the new storage type introduced with Windows 2000. If you changed a disk to dynamic and you want to perform a new installation on the disk, review the following.

- If you used Windows 2000 or Windows XP to change a disk with no partitions directly to a dynamic disk, you must revert the disk to basic before you can install an operating system on it.

  **Caution**

  - All data will be lost in the process of reverting the disk to basic, so back it up first.

  You can use Windows 2000 or Windows XP to revert the disk, or you can use the Setup program for a product in the Windows Server 2003 family. To use Windows 2000 or Windows XP, follow the instructions in Help for your operating system. To use Setup for a product in the Windows Server 2003 family, during the partitioning phase, find the dynamic disk among the list of available partitions and then delete it (which erases all data on all volumes). You will be prompted to confirm your action. After you do this, the disk will contain only unpartitioned space, and you can use Setup to create a new (basic disk) partition on it.

- If you plan to rerun Setup on a computer on which a product in the Windows Server 2003 family was already installed, and the computer contains dynamic disks, be sure to read about the limitations for installing operating systems on disks that are changed to dynamic using a product in the Windows Server 2003 family. For more information, see Dynamic disks and volumes.

# Networks: TCP/IP, IP Addresses, and Name Resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Networks: TCP/IP, IP addresses, and name resolution**


- Overview of networking and TCP/IP
- Internet Protocol (IP) addresses
- Name resolution for TCP/IP

# Overview of networking and TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Overview of networking and TCP/IP**

TCP/IP is the network protocol that provides Internet access. It is the protocol used by most servers, although you can use additional or different network adapters and their associated protocols on your servers.

To use TCP/IP, make sure that each server is provided with an IP address, either a dynamic or automatic address provided through software, or a static address that you obtain and set. Because these addresses are numbers and therefore hard to remember, you will also have to provide users with names that are easier to use. Mapping this type of name to an IP address is called name resolution, and can be accomplished by various methods, primarily the Domain Name System (DNS) and Windows Internet Name Service (WINS). For more information, see Internet Protocol (IP) addresses and Name resolution for TCP/IP.

# Internet Protocol (IP) addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol (IP) addresses**

As outlined in Overview of networking and TCP/IP, using TCP/IP requires that an IP address be provided for each computer. This topic provides a brief description of the options for doing this. For more detailed information, see DHCP and IP Addressing and Routing.

There are two basic approaches for providing an IP address for a server you are installing:

**Dynamic Host Configuration Protocol (DHCP)**

You can provide IP addresses to the computers on your network by configuring one or more DHCP servers, which provide IP addresses dynamically to other computers. A DHCP server must itself be assigned a static IP address.

One server or several servers can provide DHCP along with one or more name resolution services, which are called Domain Name System (DNS) and Windows Internet Name Service (WINS).

If you want to run Setup before you have finalized your decisions about which server to use as your DHCP server and what static IP address to assign to that server, you can choose **Typical settings** in the **Networking Settings** dialog box during Setup and complete the network configuration later. If you do this and there is no DHCP server in the network, Setup will use a limited IP addressing option called Automatic Private IP Addressing (APIPA). During the time that a server is using APIPA, it can communicate only with other computers using APIPA on the same network segment. A server that is using APIPA cannot make connections to the Internet (for browsing or email), and cannot be used with DNS or Active Directory (which depends on DNS).

If you know which server you want to use as your DHCP server, when installing that server, in the **Networking Settings** dialog box in Setup, choose **Custom settings**, and specify a static IP address and related network settings. For more information, see Specify a static IP address and settings needed for DNS and WINS.

**Static IP addressing**

For certain types of servers, you must assign a static IP address and subnet mask during or after Setup. These servers include DHCP servers, DNS servers, WINS servers, and any server providing access to users on the Internet. It is also recommended that you assign a static IP address and subnet mask for each domain controller. If a computer has more than one network adapter, you must assign a separate IP address for each adapter.

If you want to run Setup on a server before you have finalized your decision about the static IP address you want to assign to that server, you can choose **Typical settings** in the **Networking Settings** dialog box during Setup and configure that server later. In this situation, if there is a DHCP server in the network, Setup will obtain an IP address configuration from DHCP. If there is no DHCP server in the network, Setup will use Automatic Private IP Addressing (APIPA). APIPA is described in the previous item in this list.

For more information about static IP addresses, including private IP addresses (which you choose from certain ranges of addresses) and public IP addresses (which you obtain from an Internet service provider), see IP Addressing and Routing.

After you have formed a plan for IP addressing, the next components to consider are those for name resolution, which is the process of mapping a computer name (something that users can recognize and remember) to the appropriate IP address. For more information, see Name resolution for TCP/IP.

# Name resolution for TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Name resolution for TCP/IP**

Name resolution is a process that provides users with easy-to-remember server names, instead of requiring them to use the numerical IP addresses by which servers identify themselves on the TCP/IP network. The name-resolution services are Domain Name System (DNS) and Windows Internet Name Service (WINS).

**DNS**

DNS is a hierarchical naming system used for locating computers on the Internet and private TCP/IP networks. One or more DNS servers are needed in most installations. DNS is required for Internet e-mail, Web browsing, and Active Directory. DNS is often used as a name resolution service in domains with clients running Windows 2000, Windows XP, or products in the Windows Server 2003 family. DNS is installed automatically when you create a domain controller (or when you install Active Directory on an existing member server, which makes it a domain controller), unless Windows Server 2003 software detects that a DNS server already exists for that domain. You can also install DNS by choosing the DNS server role in Manage Your Server or by using Control Panel.

If you plan to install DNS on a server, specify a static IP address on that server and configure that server to use that IP address for its own name resolution.

For more information, see the following:

- To choose from a list of topics about configuring DNS, see DNS.
- For information about installing DNS by choosing the DNS server role in Manage Your Server, see Configuring Roles for Your Server.
- For information about how to install DNS by using Control Panel, see Install a DNS server.
- For information about assigning a static IP address during Setup, see Specify a static IP address and settings needed for DNS and WINS.
- For information about assigning a static IP address after Setup, see Configure TCP/IP for static addressing.
- For information about installing a domain controller (or about installing Active Directory on an existing member server, which makes it a domain controller), see Manage Domain Controllers.

**WINS**

If you provide support for clients running Windows NT or any earlier Microsoft operating system, you might need to install Windows Internet Name Service (WINS) on one or more servers in the domain. You might also need to install WINS if it is required by your applications. You can install WINS by choosing the WINS server role in Manage Your Server or by using Control Panel.

If you plan to install WINS on a server, specify a static IP address on that server.

For more information, see the following:

- To choose from a list of topics about configuring WINS, see WINS.
- For information about installing WINS by choosing the WINS server role in Manage Your Server, see Configuring Roles for Your Server.
- For information about how to install WINS by using Control Panel, see Install a WINS server.
- For information about assigning a static IP address during Setup, see Specify a static IP address and settings needed for DNS and WINS.

- For information about assigning a static IP address after Setup, see Configure TCP/IP for static addressing.

# Deciding Between Workgroups and Domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deciding between workgroups and domains**
This section contains information that will help you decide whether to place your servers in a workgroup or a domain. It also contains other information that, if you create a domain, will help with basic planning of the roles of servers within the domain.

For information about upgrading servers in a domain where all domain controllers run Windows NT 4.0, see Upgrades in a Windows NT 4.0 Domain. For information about upgrading servers in a domain containing one or more Windows 2000 domain controllers, see Upgrades in a domain containing Windows 2000 domain controllers. For information about networks larger than five servers, see the *Windows Server 2003 Deployment Kit* at the Microsoft Windows Server System Web site. For information about obtaining the *Windows Server 2003 Resource Kit*, see Using the Windows Deployment and Resource Kits.

For additional information, see:

- Workgroups compared with domains
- Planning for domain controllers and member servers
- Special aspects of Windows Server 2003 family domain controllers
- Additional information about domains

# Workgroups compared with domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Workgroups compared with domains**

A *domain* is a group of accounts and network resources that share a common directory database and set of security policies, and might have security relationships with other domains. A *workgroup* is a more basic grouping, intended only to help users find objects such as printers and shared folders within that group. Domains are the recommended choice for all networks except very small ones with few users.

In a workgroup, users might have to remember multiple passwords, one for each network resource. (In addition, different users can use different passwords for each resource.) In a domain, passwords and permissions are simpler to keep track of, because a domain has a single, centralized database of user accounts, permissions, and other network details. The information in this database is replicated automatically among domain controllers. You determine which servers are domain controllers and which are simply members of the domain. You can determine these roles not only during Setup but afterward.

Domains, and the Active Directory directory system of which they are a part, provide many options for making resources easily available to users while maintaining good monitoring and security. For more information, see Active Directory.

# Planning for domain controllers and member servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for domain controllers and member servers**

With Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; servers in a domain can have one of two roles: *domain controllers*, which contain matching copies of the user accounts and other Active Directory data in a given domain, and *member servers*, which belong to a domain but do not contain a copy of the Active Directory data. (A server that belongs to a workgroup, not a domain, is called a *stand-alone server*.) It is possible to change the role of a server back and forth from domain controller to member server (or stand-alone server), even after Setup is complete. However, it is recommended that you plan your domain before running Setup and change server roles (and server names) only when necessary.

Multiple domain controllers provide better support for users, compared to a single domain controller. With multiple domain controllers, you have multiple copies of user account data and other Active Directory data; however, it is still important to perform regular backups, including Automated System Recovery backups, and familiarize yourself with the methods for restoring a domain controller. In addition, multiple domain controllers work together to support domain controller functions, such as carrying out logon validations.

For more information, see Active Directory server roles.

# Special aspects of Windows Server 2003 family domain controllers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Special aspects of Windows Server 2003 family domain controllers**

As you manage your Windows Server 2003 family domains, you might want to learn more about operations master roles. This topic describes operations master roles briefly. A variety of domain topics are discussed in more depth in Active Directory.

Operations master roles are special roles that are assigned to one or more domain controllers in an Active Directory domain. The domain controllers that are assigned these roles perform operations that are single-master (not permitted to occur at different places in a network at the same time). For example, the creation of security identifiers for new resources (such as new computers) must be overseen by a single domain controller to ensure that the identifiers are unique.

The first domain controller installed in a domain is automatically assigned all the operations master roles. You can change the assignment of operations master roles after Setup, but in most cases this is not necessary. You will need to be particularly aware of operations masters roles if problems develop on an operations master or if you plan to take one out of service. For more information, see Operations master roles.

# Additional information about domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Additional information about domains**

For information about upgrading servers in a domain where all domain controllers run Windows NT 4.0, see Upgrades in a Windows NT 4.0 Domain. For information about upgrading servers in a domain containing one or more Windows 2000 domain controllers, see Upgrades in a domain containing Windows 2000 domain controllers.

For detailed information about domains and forests, domain functional level and forest functional level, operations master roles, the global catalog, and other domain-related topics, see Active Directory.

For information about setting up multiple domains with structured relationships, see the *Microsoft Windows Server 2003 Deployment Kit* at the Microsoft Windows Server System Web site. For information about obtaining the *Windows Server 2003 Resource Kit*, see Using the Windows Deployment and Resource Kits.

# Preparing Your System Before Running Setup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Preparing your system before running Setup**

- Preparing your system for an upgrade
- Preparing your system for a new installation
- Working with volume, mirror, or stripe sets or stripe sets with parity

# Preparing your system for an upgrade

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Preparing your system for an upgrade**
This topic describes basic steps to take in preparing your server for an upgrade.

If you are upgrading to Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, on server cluster nodes, be sure to read Upgrading to Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition on Cluster Nodes.

**Checking the system log for errors**
Use Event Viewer to review the system log for recent or recurring errors that could cause problems during the upgrade. For information about Event Viewer, see Event Viewer overview, or Help for the operating system that you are running.

**Backing up files**
Before upgrading, it is recommended that you back up your current files, including anything containing configuration information, for example, the System State and the system and boot partitions. You can back up files to a variety of different media, such as a tape drive or the hard disk of another computer on the network.

**Preparing mirror sets and other disk sets for upgrade (Windows NT 4.0 only)**
With the disk management technologies in Windows NT 4.0, you could create volume sets, mirror sets, stripe sets, or stripe sets with parity, each with specific capabilities and limitations. By using dynamic disks, introduced with Windows 2000, you can take advantage of similar technologies, and with the Windows Server 2003 operating systems, you can also extend dynamic volumes without repartitioning or reformatting.

This transition from the technologies used in Windows NT 4.0 means that you must make certain choices before running Setup for Windows Server 2003. For important information about these choices, see Working with volume, mirror, or stripe sets or stripe sets with parity.

**Disconnecting UPS devices**
If you have an uninterruptible power supply (UPS) connected to your target computer, disconnect the connecting serial cable before running Setup. Setup automatically attempts to detect devices connected to serial ports, and UPS equipment can cause problems with the detection process.

**Reviewing hardware and software**
When you start Setup for an upgrade, the first process it carries out is a check for compatible hardware and software on your computer. Setup displays a report before continuing. Use this report, along with information in Relnotes.htm (in the \Docs folder on the Setup CD), to find out whether you need to update your hardware, drivers, or software before upgrading.

### Important

- If you use an Internet service provider (ISP), pay special attention to any recommendations about updating the ISP software necessary for your Internet connectivity.

For the most recent information on compatible hardware and software for products in the Windows Server 2003 family, click the appropriate links in Support resources. If you have hardware devices that do not use Plug and Play, see Taking an inventory of devices that do not use Plug and Play.

# Preparing your system for a new installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Preparing your system for a new installation**
This topic describes basic steps to take in preparing your server for a new installation.

If you are installing Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, on server cluster nodes, be sure to read Installing Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition on Cluster Nodes.

**Checking the system log for errors**
If the computer already has a working operating system, review the system log for recent or recurring errors (especially hardware errors) that could cause problems during the installation. For information about viewing event logs, see Help for the operating system on your computer.

**Backing up files**
Before you perform a new installation, it is recommended that you back up your current files, unless the computer has no files or the current operating system files have been damaged. You can back up files to a variety of different media, such as a tape drive or the hard disk of another computer on the network.

**Uncompressing the drive**
Uncompress any DriveSpace or DoubleSpace volumes before installing. Do not install Windows Server 2003 on a compressed drive unless the drive was compressed with the NTFS file system compression feature.

**Preparing mirror sets and other disk sets (Windows NT 4.0 only)**
With the disk management technologies in Windows NT 4.0, you could create volume sets, mirror sets, stripe sets or stripe sets with parity, each with specific capabilities and limitations. By using dynamic disks, introduced with Windows 2000, you can take advantage of similar technologies, and with Windows Server 2003, you can also extend dynamic volumes without repartitioning or reformatting.

This transition from the technologies used in Windows NT 4.0 means that you must make certain choices before running Setup for Windows Server 2003. For important information about these choices, see Working with volume, mirror, or stripe sets or stripe sets with parity.

**Disconnecting UPS devices**
If you have an uninterruptible power supply (UPS) connected to your target computer, disconnect the connecting serial cable before running Setup. Setup automatically attempts to detect devices connected to serial ports, and UPS equipment can cause problems with the detection process.

# Working with volume, mirror, or stripe sets or stripe sets with parity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Working with volume, mirror, or stripe sets or stripe sets with parity**

With the disk management technologies in Windows NT 4.0, you could create volume sets, mirror sets, stripe sets, or stripe sets with parity, each with specific capabilities and limitations. By using the dynamic disk technology introduced with Windows 2000, you can use similar technologies, with the added flexibility of being able to extend disk volumes without repartitioning or reformatting.

This transition from the technologies used in Windows NT 4.0 means that you must make certain choices before running Setup for Windows Server 2003. Any volume sets, mirror sets, stripe sets, or stripe sets with parity that you created with Windows NT 4.0 are not supported in the Windows Server 2003 operating systems, although they were supported to a limited extent in Windows 2000.

If you used Windows NT 4.0 to create a volume set, mirror set, stripe set, or stripe set with parity, and you want to run Setup for Windows Server 2003 on that computer, you must choose one of the following methods.

For information about other preparations to make before running Setup, see Checklists for installing and upgrading the operating system and Preparing Your System Before Running Setup.

**For a mirror set, break the mirror**

If you are running Windows NT 4.0 on a computer that has a mirror set and you want to run Setup for Windows Server 2003 on that computer, first back up the data (as a safeguard-the data will not be erased) and then break the mirror. Ensure that you have applied Service Pack 5 or later, required before running Setup for Windows Server 2003. Then run Setup.

**For a volume set, stripe set, or stripe set with parity, back up the data and delete the set**

If you are running Windows NT 4.0 on a computer that has a volume set, stripe set, or stripe set with parity, and you want to run Setup for Windows Server 2003 operating systems on that computer, first back up the data. Next, delete the set (which will delete the data). Ensure that you have applied Service Pack 5 or later, required before running Setup for Windows Server 2003. Then run Setup.

After running Setup for Windows Server 2003, you can make the disk dynamic, restore backed-up data as necessary, and make use of the volume options shown in the table later in this section. For more information about dynamic disks, see Dynamic disks and volumes and Using Dynamic Disks.

**If necessary, use the Ftonline support tool**

The preceding methods are the recommended methods for preparing to run Setup for Windows Server 2003 on a computer that contains a volume set, mirror set, stripe set, or stripe set with parity created with Windows NT 4.0. However, if you do not use these methods and you must access one of these sets after running Setup for Windows Server 2003, you can use the Ftonline tool. The Ftonline tool is part of the Windows Server 2003 family Support Tools. For more information, see Install Windows Support Tools.

**Types of multidisk volumes on dynamic disks**

The disk sets described in the previous section have different names in the Windows Server 2003 family than they had in Windows NT 4.0:

| Windows NT 4.0 name | Windows Server 2003 family name |
|---|---|
| Volume set | Spanned volume on a dynamic disk |
| Mirror set | Mirrored volume on a dynamic disk |
| Stripe set | Striped volume on a dynamic disk |
| Stripe set with parity | RAID-5 volume on a dynamic disk |

# Resources for installing and upgrading the operating system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources for installing and upgrading the operating system**

This topic provides sources of additional information about getting a new server ready for use in your enterprise.

- To confirm that your computers, system components, and peripherals are compatible with products in the Windows Server 2003 family, click the appropriate link in Support resources.
- For the most recent information on compatible applications for products in the Windows Server 2003 family, see the software compatibility section of the Microsoft Web site.
- For enhancements and updates to Windows software and other Microsoft software, see Windows Update at the Microsoft Web site. For a description of the different ways you can use Windows Update, see Windows Update.
- For information about the Windows Deployment Resource Kit, Windows Server 2003 Technical Reference, and Windows Server TechCenter, which provide technical background information on products in the Windows Server 2003 family and also on deploying Windows XP, see the following:
    - The *Windows Server 2003 Deployment Kit* at the Microsoft Windows Server System Web site.
    - The *Windows Server 2003 Technical Reference* at the Microsoft Windows Server System Web site.
    - The *Windows Server TechCenter* at the Microsoft Web site.
    - The *Windows XP Professional Resource Kit* at the Microsoft Windows XP Resource Kits Web site.
    - Using the Windows Deployment and Resource Kits
- For information about the differences between Windows Server 2003 and the version of Windows Server 2003 that includes Service Pack 1 (SP1), see Changes to Functionality in Microsoft Windows Server 2003 Service Pack 1.
- Microsoft TechNet is a service that contains resources for Information Technology (IT) professionals. This includes access to white papers and articles written by support professionals at Microsoft, career and industry information, discussion groups, technical feedback, and software. For more information, see TechNet at the Microsoft Web site.
- The Information Technology Infrastructure Library (ITIL) Web site provides guidance for Information Technology professionals. The following is an excerpt from the site: "ITIL consists of a series of books giving guidance on the provision of quality IT services, and on the accommodation and environmental facilities needed to support information technology." (Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

# Troubleshooting for installing and upgrading the operating system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section outlines the steps to troubleshoot Stop messages, which are the result of a system error in the operating system. Following a system error, your computer stops responding, and a message is displayed on a blue or black background.

For information about how to obtain product support, see Product Support at the Microsoft Web site.

This section covers:

- Troubleshooting Stop messages: general strategies
- Troubleshooting specific Stop messages

# Troubleshooting Stop messages: general strategies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting Stop messages: general strategies**
This topic describes a general method for troubleshooting Stop messages or unexpected shutdowns. For more information, see the following:

- For information about specific Stop messages, see Troubleshooting specific Stop messages.
- For articles about specific Stop messages, see Search the Knowledge Base at the Microsoft Web site.
- For information about support options, see Support resources.
- For information about troubleshooting specific technologies or components (for example, disks and file systems), see the *Windows Server 2003 Technical Reference* at the Microsoft Windows Server System Web site.

You can use the System Information tool to gather important information about system settings, problem devices, and hardware resources. For more information, see System Information.

**Important**

- To ensure system stability, install the latest service pack that is appropriate for your system (follow service pack instructions) and use only signed drivers. For more information, see Windows Update at the Microsoft Web site. For a description of the different ways you can use Windows Update, see Windows Update.

Follow these steps for any Stop error that does not display identifying text or specific troubleshooting steps:

1. Use the Windows Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Use the File Signature Verification tool, which identifies unsigned drivers and incompatible system files on your computer. The system files and device driver files that are provided with the Windows Server 2003 family have a Microsoft digital signature, which indicates that the files are original, unaltered system files or that they have been approved by Microsoft for use with Windows. For more information, see Using File Signature Verification.

3. If the Stop message is caused by a driver problem, you can use Device Manager to roll back to a previous version of the driver. For more information, see Roll back to the previous version of a driver. You may need to disable any newly installed drivers.

4. Using a recent version of virus-protection software, check for viruses on your computer. If a virus is found, perform the steps required to eliminate it from your computer. See your antivirus software documentation for these steps.

5. You can search the Knowledge Base for updated information on Stop messages. See Search the Knowledge Base at the Microsoft Web site, and then follow the instructions on the page. When typing keywords, use the Stop message number (for example, **stop 0x0000000A**).

   For more information about articles, troubleshooting wizards, and items that you can download, see Updated technical information.

6. If you can start the operating system, check the System Log in Event Viewer for additional information that might help

determine the device or driver that is causing the problem. Many Stop messages produce specific events that you can search for, such as Event ID 1001 (generated by various Stop messages) and Event ID 6008 (usually caused by an unexpected shutdown or a computer that does not restart correctly). To display a description of an event, in Event Viewer, double-click the event. For more information about Event Viewer, see Event Viewer.

You can search the Knowledge Base for updated information on Event IDs. See Search the Knowledge Base at the Microsoft Web site, and then follow the instructions on the page.

7. If you cannot start your computer, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

8. Remove any newly installed hardware (RAM, adapters, hard disks, modems, and so on).

9. Ensure that you have updated drivers for your hardware devices. Also, ensure that you have the latest system BIOS (for x86-based or x64-based computers) or the correct firmware (for Itanium architecture-based computers). The device or hardware manufacturers can assist you in obtaining these items.

10. Run the system diagnostics supplied by your computer manufacturer, especially the memory check.

11. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

12. If Safe Mode and other startup options do not work, you can consider using the Recovery Console. This method is recommended only if you are an advanced user who can use basic commands to identify and locate problem drivers and files. In addition, you must be an administrator to use the Recovery Console. For more information, see Recovery Console overview.

# Troubleshooting specific Stop messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting specific Stop messages**
This topic discusses how to troubleshoot specific Stop messages. For general information about troubleshooting Stop messages, and for a list of other sources of information, see Troubleshooting Stop messages: general strategies.

You can use the System Information tool to gather important information about system settings, problem devices, and hardware resources. For more information, see System Information.

**What Stop message are you receiving?**

- Stop message 0x0000001E Descriptive text: KMODE_EXCEPTION_NOT_HANDLED
- Stop message 0x000000D1 Descriptive text: DRIVER_IRQL_NOT_LESS_OR_EQUAL
- Stop message 0x000000EA Descriptive text: THREAD_STUCK_IN_DEVICE_DRIVER
- Stop message 0x00000050 Descriptive text: PAGE_FAULT_IN_NONPAGED_AREA
- Stop message 0x000000BE Descriptive text: ATTEMPTED_WRITE_TO_READONLY_MEMORY
- Stop message 0x0000000A on an existing installation Descriptive text: IRQL_NOT_LESS_OR_EQUAL
- Stop message 0x0000000A on a new installation Descriptive text: IRQL_NOT_LESS_OR_EQUAL
- Stop messages 0x00000023 and 0x00000024 Descriptive text: FAT_FILE_SYSTEM or NTFS_FILE_SYSTEM
- Stop message 0x0000007B Descriptive text: INACCESSIBLE_BOOT_DEVICE
- Stop message 0x0000007F Descriptive text: UNEXPECTED_KERNEL_MODE_TRAP
- Stop message 0x0000009F Descriptive text: DRIVER_POWER_STATE_FAILURE
- Stop message 0xC000021A Descriptive text: FATAL_SYSTEM_ERROR
- Stop message 0xC0000221 Descriptive text: BAD_IMAGE_CHECKSUM
- Stop message 0xC0000218 Descriptive text: UNKNOWN_HARD_ERROR

**Stop message 0x0000001E Descriptive text: KMODE_EXCEPTION_NOT_HANDLED**
**Usual cause:** This message has several possible causes.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Check that you have adequate disk space, especially for new installations.

3. Disable any newly installed drivers. You might also want to disable a service associated with the driver or any newly

installed programs.

4. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

5. Ensure that you have the latest system BIOS (for x86-based and x64-based computers) or the correct firmware (for Itanium architecture-based computers). Contact your hardware manufacturer if you need assistance in performing this step.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x0000001E**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0x000000D1 Descriptive text: DRIVER_IRQL_NOT_LESS_OR_EQUAL**
**Usual cause:** An attempt was made to touch pageable memory at a process internal request level (IRQL) that was too high. This is usually caused by drivers using improper addresses. Kernel-mode drivers are forbidden to access pageable memory at a high IRQL.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Use the File Signature Verification tool, which identifies unsigned drivers and incompatible system files on your computer. The system files and device driver files that are provided with the Windows Server 2003 family have a Microsoft digital signature, which indicates that the files are original, unaltered system files or that they have been approved by Microsoft for use with Windows. For more information, see Using File Signature Verification.

3. Use Device Manager to roll back to a previous version of the driver. For more information, see Roll back to the previous version of a driver.

4. Disable the driver identified in the Stop message or any newly installed drivers.

5. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

- When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

6. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, go to the see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x000000D1**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0x000000EA Descriptive text: THREAD_STUCK_IN_DEVICE_DRIVER**
**Usual cause:** The display driver is caught in an infinite loop. This usually indicates a problem with the video hardware or with the display driver.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Disable the driver identified in the Stop message or any newly installed drivers.

3. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

4. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x000000EA**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0x00000050 Descriptive text: PAGE_FAULT_IN_NONPAGED_AREA**
**Usual cause:** Requested data is not found in memory. This might be caused by a printer driver or remote access driver.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Use Device Manager to roll back to a previous version of the driver. For more information, see Roll back to the previous version of a driver.

3. Disable any newly installed drivers.

4. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

5. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x00000050**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0x000000BE Descriptive text: ATTEMPTED_WRITE_TO_READONLY_MEMORY**
**Usual cause:** A device driver attempted to write to read-only memory. This problem can occur during an upgrade, although it is not limited to upgrades.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

3. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

- When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x000000BE**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0x0000000A on an existing installation Descriptive text: IRQL_NOT_LESS_OR_EQUAL**
**Usual cause:** Drivers are using improper memory addresses.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. If you can start the operating system, check the System Log in Event Viewer for additional information that might help determine the device or driver that is causing the problem. To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.

3. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

4. Disable any newly installed drivers and remove any newly added programs.

5. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

6. Remove any newly installed hardware (RAM, adapters, hard disks, modems, and so on).

7. Ensure that you have updated drivers for your hardware devices. Also, ensure that you have the latest system BIOS (for x86-based and x64-based computers) or the correct firmware (for Itanium architecture-based computers). The device manufacturers can assist you in obtaining these items.

8. Run system diagnostics supplied by your hardware vendor, especially a memory check.

9. Disable memory caching in the BIOS (for x86-based and x64-based computers) or EFI (for Itanium architecture-based computers). Contact your hardware manufacturer if you need assistance in performing this step.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x0000000A**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0x0000000A on a new installation Descriptive text: IRQL_NOT_LESS_OR_EQUAL**
**Usual cause:** Drivers are using improper memory addresses.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. During the installation process, at the "Setup is inspecting your computer's hardware configuration" message, press F5. Select the correct computer type when you are prompted. For example, if you have a single-processor computer, select **Standard PC**.

3. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

4. Disable memory caching in the BIOS (for x86-based and x64-based computers) or EFI (for Itanium architecture-based computers). Contact your hardware manufacturer if you need assistance in performing this step.

5. Remove all adapters and disconnect all hardware devices that are not essential.

6. If you are using a SCSI adapter, obtain the latest driver from the adapter vendor, disable sync negotiation, check termination, and check the SCSI IDs of the devices.

7. If you are using Integrated Device Electronics (IDE) devices, define the onboard IDE port as the primary channel. Check the master/subordinate settings for the IDE devices. Remove all IDE devices except for the hard disk. See your hardware manufacturer's documentation for more information.

8. Run the system diagnostics supplied by your hardware vendor, especially the memory check.

9. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x0000000A**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical

information.

**Stop messages 0x00000023 and 0x00000024 Descriptive text: FAT_FILE_SYSTEM or NTFS_FILE_SYSTEM**

**Usual cause:** The hard disk is fragmented or there is excessive file I/O. This Stop error may also be cause by drive-mirroring or some antivirus software.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Disable any antivirus or backup programs, and disable any defragmentation utilities.

3. Check for hard disk corruption by running **chkdsk** and then restarting the computer. If you cannot run **chkdsk**, you might have to move the hard disk to another computer that is running a product in the Windows Server 2003 family and then run the **chkdsk** command. For more information, see chkdsk. (Note that **chkdsk** is available, with a limited set of options, in Recovery Console.)

4. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

**Note**

- You can search the Microsoft Web site for updated information on these Stop messages. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x00000023** or **stop 0x00000024**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0x0000007B Descriptive text: INACCESSIBLE_BOOT_DEVICE**

**Usual cause:** A problem occurred during the initialization of the I/O system (usually the startup device or the file system).

**Solution:**

1. Check for viruses on your computer. This Stop message is frequently displayed when there is a virus in the boot sector.

2. Repair or recover the hard disk. For ways to fix the hard disk, see Repair overview and Recovery Console overview.

3. Remove any newly added hard disks or controllers.

4. If you are using a SCSI adapter, obtain the latest driver from the adapter vendor, disable sync negotiation, check termination, and check the SCSI IDs of the devices.

5. If you are using Integrated Device Electronics (IDE) devices, define the onboard IDE port as the primary channel. Check the master/subordinate settings for the IDE devices. Remove all IDE devices except for the hard disk. See your hardware

manufacturer's documentation for more information.

6. Check for hard disk corruption by running **chkdsk**. If you cannot run chkdsk, you might have to move the hard disk to another computer that is running a product in the Windows Server 2003 family and then run the **chkdsk** command. For more information, see chkdsk. (Note that **chkdsk** is available, with a limited set of options, in Recovery Console.)

7. If the computer will not start normally, try starting it in Last Known Good Configuration. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x0000007B**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0x0000007F Descriptive text: UNEXPECTED_KERNEL_MODE_TRAP**
**Usual cause:** There is a hardware or software problem. This is usually caused by a hardware failure.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Run system diagnostics supplied by your hardware vendor, especially a memory check. This Stop message is often displayed in the case of faulty or mismatched memory.

3. Disable memory caching in the BIOS (for x86-based and x64-based computers) or EFI (for Itanium architecture-based computers). Contact your hardware manufacturer if you need assistance in performing this step.

4. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

5. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x0000007F**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0x0000009F Descriptive text: DRIVER_POWER_STATE_FAILURE**
**Usual cause:** A driver is in an inconsistent or invalid power state.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Disable the driver identified in the Stop message or any newly installed drivers.

3. Update any software that uses filter drivers, such as antivirus or backup programs.

4. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

5. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, go to the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0x0000009F**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0xC000021A Descriptive text: FATAL_SYSTEM_ERROR**
**Usual cause:** The Windows Logon Process system process terminated unexpectedly.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Disable the driver identified in the Stop message or any newly installed drivers.

3. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

4. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0xC000021A**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0xC0000221 Descriptive text: BAD_IMAGE_CHECKSUM**
**Usual cause:** This message might be caused by a corrupted system file or driver or a hardware problem.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Replace the identified file.

3. If the computer will not start normally, try starting it in Last Known Good Configuration or in Safe Mode, and then replace the file. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0xC0000221**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

**Stop message 0xC0000218 Descriptive text: UNKNOWN_HARD_ERROR**
**Usual cause:** A registry file might be corrupted.

**Solution:**

1. Use the Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. If the computer will not start normally, try starting it in Last Known Good Configuration. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

3. If you have an Automated System Recovery (ASR) backup, you might need to use this recovery option. For more information, see Automated System Recovery (ASR) overview.

**Note**

- You can search the Microsoft Web site for updated information on this Stop message. To do this, see Welcome to Microsoft Help and Support at the Microsoft Web site, click a search option, and then follow the instructions on the page. When typing keywords, use **stop 0xC0000218**.

  For more information about articles, troubleshooting wizards, and items you can download, see Updated technical information.

# Installing and Upgrading on Cluster Nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing and upgrading on cluster nodes**
With Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, you can use clustering to ensure that users have constant access to important server-based resources. With clustering, you create several cluster nodes that appear to users as one server. If one of the nodes in the cluster fails, another node begins to provide service (a process known as failover). Mission-critical applications and resources remain continuously available.

The following topics describe the options for upgrading or installing on cluster nodes.

- Before performing an upgrade, see Checklist: Preparation for upgrading a cluster
- Before performing a new installation, see Checklist: Preparation for installing a cluster
- For tips on upgrading and installing on cluster nodes, see Best practices for installing and upgrading cluster nodes
- For help with specific tasks, see How to upgrade cluster nodes
- For general background information, see Concepts for installing and upgrading cluster nodes
- For problem-solving instructions, see Troubleshooting cluster node installations.

# Checklists for installing and upgrading on cluster nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

- Checklist: Preparation for upgrading a cluster
- Checklist: Preparation for installing a cluster

# Checklist: Preparation for upgrading a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Preparation for upgrading a cluster**

| Step | Reference |
|---|---|
| Check that your hardware, including your cluster storage hardware, is compatible with Windows Server 2003, Enterprise Edition.<br><br>Be aware that geographically dispersed (majority node set) server clusters are built on a networking and storage infrastructure with characteristics that are very different than a standard quorum device server cluster, and, as such, are certified under a separate testing program. If you are considering converting your existing server cluster to a geographically dispersed server cluster, you must work with your hardware vendor to obtain a geographically dispersed server cluster configuration that has been certified by Microsoft.<br><br>For more information on majority node set server clusters, see Choosing a Cluster Model. | Hardware compatibility |
| Check with the manufacturer of your cluster storage hardware and be sure you have the drivers you need. | Mass storage drivers and the setup process |
| Review concepts for upgrading cluster nodes, choose the method you will use, and review the instructions for the method you chose.<br><br>**Caution**<br><br>• If you are using any method other than a rolling upgrade (or a rolling upgrade with one or more resources offline), be sure to review and follow the instructions carefully to prevent corruption of the cluster storage. That is, for methods other than a rolling upgrade, be sure to complete the upgrade and start the Cluster service on one node before allowing another node to gain access to the cluster storage. | Concepts for installing and upgrading cluster nodes<br><br>How to upgrade cluster nodes |
| Review the checklist for upgrading the operating system. | Checklist: Performing an upgrade |
| Review procedures for starting and running Setup. | Start Setup on an x86-based computer<br><br>Start Setup on an Itanium architecture-based computer<br><br>Start Setup on an x64-based computer |

| | |
|---|---|
| Carry out the method you chose for upgrading of cluster nodes. | How to upgrade cluster nodes |

# Checklist: Preparation for installing a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Preparation for installing a cluster**

**Important**

- If you are using a shared storage device, make sure that Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, and the Cluster service are installed and running on one server before starting the operating system on another server in the cluster. If the operating system is started on multiple servers before the Cluster service is running on one server, the cluster storage could be corrupted. Once the first server is configured, the remaining servers may be installed simultaneously.

| Step | Reference |
|---|---|
| Check that your hardware, including your cluster storage hardware, is compatible with Windows Server 2003, Enterprise Edition. | Hardware compatibility |
| Check with the manufacturer of your cluster storage hardware and be sure you have the drivers you need. | Mass storage drivers and the setup process |
| Plan basic network details for the cluster, including details about network adapters, IP addressing, and the name and account the cluster will use. Also, plan for the volume needed for cluster configuration data. | Planning and preparing for cluster installation |
| Review the checklist for installation of the operating system and make decisions about such items as whether to install additional components. | Checklist: Performing a new installation |
| Follow the remaining steps for creating a server cluster. | Checklist: Planning and creating a server cluster |

# Best practices for installing and upgrading cluster nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Follow the appropriate checklist when installing or upgrading on cluster nodes**

- Checklist: Preparation for upgrading a cluster
- Checklist: Preparation for installing a cluster

**Use the Typical (full) configuration option when creating a new cluster**

With the Cluster Administrator New Server Cluster Wizard, you can choose between Typical (full) configuration and Advanced (minimum) configuration options. Typical configuration is appropriate for most installations and results in a completely configured cluster. Use the Advanced configuration option only for clusters that have complex storage configurations that the New Server Cluster Wizard cannot validate or for configurations in which you do not want the cluster to manage all of the storage. The following examples describe each situation:

- In some complex storage solutions, such as a fiber channel switched fabric that contains several switches, a particular storage unit might have a different identity on each computer in the cluster. Although this is a valid storage configuration, it violates the storage validation heuristics in the New Server Cluster Wizard. If you have this type of storage solution, you might receive an error when you are trying to create a cluster using the Typical configuration option. If your storage configuration is set up correctly, you can disable the storage validation heuristics and avoid this error by restarting the New Server Cluster Wizard, selecting the Advanced configuration option instead.
- On particular nodes in a cluster, you may want to have some disks that are to be clustered and some disks that are to be kept private. The Typical configuration option configures all disks as clustered disks and creates cluster resources for them all. However, with the Advanced configuration option, you can keep certain disks private because this configuration creates a cluster in which only the quorum disk is managed by the cluster (if you chose to use a physical disk as the quorum resource). After the cluster is created, you must then use Cluster Administrator to add any other disks that you want the cluster to manage.

For more information, see Storage configuration options.

**Important**

- If you are using a shared storage device: Before creating a cluster, when you turn the computer on and start the operating system, it is very important that only one node has access to the cluster disk. Otherwise, the cluster disks can become corrupted. To prevent the corruption of the cluster disks, shut down all but one cluster node, or use other techniques (for example, LUN masking, selective presentation, or zoning) to protect the cluster disks before creating the cluster. Also, before starting the installation of the second and subsequent nodes, ensure that all disks that are to be managed by the cluster have disk resources associated with them. If these disks do not have disk resources associated with them at this time, the disk data will be corrupted because the disks will not be protected and multiple nodes will attempt to connect to them at the same time.
- Do not use Manage Your Server to configure cluster nodes.

**Secure your server cluster**

To keep your server cluster secure from unauthorized users:

- Do not set the Cluster service account to be a member of the domain Administrators group. By giving the minimal possible user rights to the Cluster service account, you avoid potential security issues if that account is compromised.
- Limit and audit access to shared data (for example, files and folders on cluster disks).

- Limit client access to cluster resources.
- Use different accounts for the Cluster service and applications in the cluster.
- Use different Cluster service accounts for multiple clusters.
- If you upgraded from a Windows NT 4.0 or Windows 2000 cluster, remove the read permission on the Windows\Cluster directory and the MSCS directory on the quorum disk (which contains the quorum log) from all nonadministrator user accounts.
- Follow standard best practices for securing servers.

If you want to audit access to shared data, enable auditing on all cluster nodes. For more information, see Securing shared data in a cluster. For more information on the user rights required by the Cluster service account, see Change the account under which the Cluster service runs. For more information on security in a server cluster, see Managing Security in a Cluster. For more information on best practices for securing servers, see Best practices for security.

**Review other server cluster best practices**
For more information, see the Server Cluster Server Cluster Best Practices.

# How to upgrade cluster nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**
These sections explain how to set up clustering with Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. Before beginning the installation or upgrade of a cluster, review Checklist: Preparation for upgrading a cluster or Checklist: Preparation for installing a cluster.

- Upgrade a Cluster from Windows 2000
- Upgrade a Cluster from Windows NT 4.0

# Upgrade a Cluster from Windows 2000

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrade a cluster from Windows 2000**
If you are upgrading from Windows 2000 to Windows Server 2003, Enterprise Edition on cluster nodes, you might be able to perform a *rolling upgrade* of the operating system. In a rolling upgrade, you sequentially upgrade the operating system on each node, making sure that one node is always available to handle client requests. When you upgrade the operating system, the Cluster service is automatically upgraded also. A rolling upgrade maximizes availability of clustered services and minimizes administrative complexity.

The following topics describe the processes and procedures for upgrading from Windows 2000 to Windows Server 2003.

- Perform a rolling upgrade from Windows 2000
- Perform a last node rolling upgrade from Windows 2000
- Upgrade from Windows 2000 on a cluster containing IIS using the provided script
- Upgrade from Windows 2000 on a cluster containing IIS without using the provided script
- Upgrade from Windows 2000 on a cluster containing Message Queuing

# Perform a rolling upgrade from Windows 2000

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To perform a rolling upgrade from Windows 2000**

1. Confirm that your systems are running Windows 2000 and use resources that are supported during a rolling upgrade, as described in "Resource behavior during rolling upgrades" in Related Topics.

2. Confirm that your hardware is designed for or is compatible with Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. For more information, see "Checklist: Preparation for upgrading a cluster" in Related Topics.

3. In Cluster Administrator, click the node that you want to upgrade first.

4. On the **File** menu, click **Pause Node**.

5. In the right details pane, click **Active Groups.**

6. In the details pane, click a group, and then on the **File** menu, click **Move Group**. Repeat this step for each group listed.

   The services will be interrupted during the time they are being moved and restarted on the other node. After the groups are moved, one node is idle, and the other node handles all client requests.

7. Use Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, Setup to upgrade the paused node from Windows 2000.

For information about starting Setup, see Related Topics.

Setup detects the earlier version of clustering on the paused node and automatically installs clustering for Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. The node automatically rejoins the cluster at the end of the upgrade process, but it is still paused and does not handle any cluster-related work.

**Notes**

- As you begin the setup process, you will receive a warning message about using the **tempdrive** option for Setup. Provided you have followed the preceding steps, you can click **OK** and proceed with Setup without using this option.

- Manage Your Server will appear when you initially log on to the newly upgraded node as an Administrator. Close Manage Your Server to continue with the rolling upgrade.

1. To verify that the node that was upgraded is fully functional, perform validation tests such as confirming that the operating system recognizes your network connections and, using Device Manager, ensuring that the operating system recognizes all of your storage devices.

2. In Cluster Administrator, click the node that was paused, and then on the File menu click **Resume Node**.

   **Note**

   - When you use Cluster Administrator for the first time on the upgraded node, Cluster Administrator will prompt you to **Open connection to cluster**. In the **Cluster or server name** field, type a period (**.**). This will connect you to the cluster on the local upgraded node. Although you can connect to the cluster by typing the cluster name, if you do so you will be connected to the cluster on the Windows 2000 node and then will be required to respond to several warning messages.

3. Repeat the preceding steps for the remaining node.

### Important

- For information about what resources are supported during rolling upgrades, see Related Topics.

### Notes

- After you upgrade a cluster node from Windows 2000, you must log on with an account that is a member of both the Domain Admins group and the Administrators group on each cluster node before you can restart the Cluster service. This step is required even if the Domain Admins group is already a member of the Administrators group on each cluster node.

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- If minor problems occur when you are in the midst of a rolling upgrade, we recommend that you complete the upgrade. This might correct the problems.

- tok:changequorumlater1

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Preparation for upgrading a cluster
Start Setup on an x86-based computer
How rolling upgrades work
Restrictions on rolling upgrades
Resource behavior during rolling upgrades
Upgrades that include a Print Spooler resource
Alternatives to rolling upgrades

# Perform a last node rolling upgrade from Windows 2000

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To perform a last node rolling upgrade from Windows 2000**

1. Confirm that your systems are running Windows 2000.

2. Confirm that your hardware is designed for or is compatible with Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. For more information, see "Checklist: Preparation for upgrading a cluster" in Related Topics.

3. Using the information in "Resource behavior during rolling upgrades" in Related Topics, identify the resources in your cluster that are not supported by Windows Server 2003 and remove those resources.

4. Again using the information in "Resource behavior during rolling upgrades" in Related Topics, identify the resources in your cluster that are not supported during rolling upgrades and are not covered by other technology specific upgrade procedures listed in Related Topics.

5. Move all of the resources identified in the previous step to a single node by double clicking the Groups folder in Cluster Administrator's console tree. Click the group that contains the resource and then in the **File** menu, click **Move Group**.

6. Upgrade the operating system on each of the nodes that is not hosting the resources moved in the previous step. After you upgrade each node, close Manage Your Server if it is displayed.

7. Move all of the resources that you moved in the earlier steps from the node that has not been upgraded to one of the newly upgraded nodes.

8. Once all of the resources have been moved, upgrade the final node. After you upgrade, close the **Manage Your Server** dialog box if it is displayed.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- After you upgrade a cluster node from Windows 2000, you must log on with an account that is a member of both the Domain Admins group and the Administrators group on each cluster node before you can restart the Cluster service. This step is required even if the Domain Admins group is already a member of the Administrators group on each cluster node.

- It is not mandatory to move all of the resources to a single node and upgrade that node last. However, following the procedure as described will typically give you the highest availability of the resources that do not support rolling upgrades.

- If minor problems occur when you are performing a rolling upgrade, we recommend that you complete the upgrade. This might correct the problems.

- tok:changequorumlater1

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Alternatives to rolling upgrades
Checklist: Preparation for upgrading a cluster
How rolling upgrades work
Perform a rolling upgrade from Windows 2000
Resource behavior during rolling upgrades
Upgrade from Windows 2000 on a cluster containing IIS without using the provided script

# Upgrade from Windows 2000 on a cluster containing IIS using the provided script

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To upgrade from Windows 2000 on a cluster containing IIS using the provided script**

1. Confirm that your hardware is designed for or is compatible with Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition.

2. Upgrade all nodes except the node containing the IIS instance resource from Windows 2000 to Windows Server 2003. You can perform either a rolling upgrade, or a non-rolling upgrade, whichever is appropriate for your installation. Prior to upgrading, be sure to remove the resources that are not supported by Windows Server 2003. For more information, see Related Topics.

3. If you do not already have a Distributed Transaction Coordinator (DTC) resource on the cluster that you are upgrading, create a DTC resource for this cluster on an upgraded node.

4. On an upgraded node, open Command Prompt and navigate to the %systemroot%\system32\inetsrv directory.

5. Type:

   **iis switch** *cluster name [group name]*

| Value | Description |
|---|---|
| cluster name | The name of a physical node in the cluster that has been upgraded or the name of a virtual server. |
| group name | This is optional. By default the script will upgrade all IIS Web and FTP resources in the cluster. However, if instead you want to upgrade only one particular group, you can specify the name of that resource group. |

**Important**

- The resource group that includes the IIS resource must be on the Windows 2000 node when you run the script.

  The script includes a step that will bring the resource group with the newly created Generic Script resource online. However the script will only bring the group online if all resources in the group were online before the script was run. If the group was only partially online or was offline prior to the script being run, the script will take the entire group offline and will not bring any resources back online after the upgrade.

  Once the script completes, the newly created resource might immediately fail. This is because the W3SVC service is disabled. You can correct this problem by completing the following step.

1. Once the script completes, take the newly created Generic Script resource offline, then start the W3SVC service on this node and set the W3SVC to start automatically.

2. Bring the Generic Script resource online.

3. Using IIS, start the Web site on the upgraded node.

4. Upgrade the final node to Windows Server 2003. Once the upgrade completes, start the W3SVC service on this node and set the service to start automatically.

5. If appropriate, add the final upgraded node to the **Possible Owners** list for the newly created Generic Script resource. For more information, see "Specify which nodes can own a resource" in Related Topics.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- After you upgrade a cluster node from Windows 2000, you must log on with an account that is a member of both the Domain Admins group and the Administrators group on each cluster node before you can restart the Cluster service. This step is required even if the Domain Admins group is already a member of the Administrators group on each cluster node.

- To cluster IIS on Windows Server 2003, you must have a DTC resource on that cluster as well.

- The script will not upgrade IIS resources that are in the cluster group. If you have IIS resources that are in the cluster group, you can instead use the procedure "To upgrade from Windows 2000 without using the provided script on a cluster that includes an IIS resource" in Related Topics.

- Only upgraded nodes can be the owner of the Generic Script Resource created in this procedure because this resource type is not supported in Windows 2000.

- If you prefer to not use the provided script and instead perform each step manually, see "Upgrade from Windows 2000 on a cluster containing IIS without using the provided script" in Related Topics.

- tok:changequorumlater1

- For more information regarding hardware compatibility, see Support resources.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Upgrade from Windows 2000 on a cluster containing IIS without using the provided script
Take a resource offline
Specify which nodes can own a resource
Move a group to another node
Move a resource to another group
Bring a resource online
Modify resource dependencies
Delete a resource
Create a new resource
Checklist: Preparation for upgrading a cluster
Checklist: Creating a clustered IIS Web or FTP service
Start Setup on an x86-based computer
How rolling upgrades work
Restrictions on rolling upgrades
Resource behavior during rolling upgrades
Alternatives to rolling upgrades

# Upgrade from Windows 2000 on a cluster containing IIS without using the provided script

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To upgrade from Windows 2000 on a cluster containing IIS without using the provided script**

1. Confirm that your hardware is designed for or is compatible with Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition.

2. Upgrade all nodes except the node containing the IIS instance resource from Windows 2000 to Windows Server 2003. You can perform either a rolling upgrade, or a nonrolling upgrade, whichever is appropriate for your installation. Prior to upgrading, be sure to remove the resources that are not supported by Windows Server 2003. For more information, see Related Topics.

3. If you do not already have a Distributed Transaction Coordinator (DTC) resource on the cluster that you are upgrading, create a DTC resource for this cluster on an upgraded node.

4. On the Windows 2000 node, note the dependencies of the IIS instance resource. Note resources that depend on the IIS resource and also note what resources IIS itself is dependent on.

5. Take the group containing the IIS instance resource offline For more information on how to do this see "Take a resource offline" in Related Topics.

6. Remove any dependencies on the IIS instance resource.

7. Delete the IIS instance resource.

8. Move the group to a node running Windows Server 2003 and create a Generic Script resource by following the procedure documented in "Checklist: Creating a clustered IIS Web or FTP service" in Related Topics.

   **Notes**

   - As part of this step, recreate the dependencies of the Generic Script resource identical to those of the now deleted IIS resource. Make everything that was dependent on the IIS resource dependent instead on the Generic Script resource. Also make the Generic Script resource dependent on everything that IIS was dependent on.

   - Only upgraded nodes can be the owner of the Generic Script Resource created in this procedure because this resource type is not supported in Windows 2000.

9. Start the W3SVC service on this node and set the service to start automatically.

10. Bring the group containing the new Generic Script resource online.

11. Using IIS, start the Web site on the upgraded node.

12. Upgrade the final node to Windows Server 2003. Once the upgrade is complete, start the W3SVC on this node and set the service to start automatically.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group

might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- After you upgrade a cluster node from Windows 2000, you must log on with an account that is a member of both the Domain Admins group and the Administrators group on each cluster node before you can restart the Cluster service. This step is required even if the Domain Admins group is already a member of the Administrators group on each cluster node.

- To cluster IIS on Windows Server 2003, you must have a DTC resource on that cluster as well.

- Only upgraded nodes can be the owner of the Generic Script Resource created in this procedure because this resource type is not supported in Windows 2000.

- Many of these steps can be automated using a script provided with Windows Server 2003. For more information, see "Upgrade from Windows 2000 on a cluster containing IIS using the provided script" in Related Topics.

- tok:changequorumlater1

- For more information regarding hardware compatibility, see Support resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Upgrade from Windows 2000 on a cluster containing IIS using the provided script
Take a resource offline
Specify which nodes can own a resource
Move a group to another node
Move a resource to another group
Bring a resource online
Modify resource dependencies
Delete a resource
Create a new resource
Checklist: Preparation for upgrading a cluster
Checklist: Creating a clustered IIS Web or FTP service
Start Setup on an x86-based computer
How rolling upgrades work
Restrictions on rolling upgrades
Resource behavior during rolling upgrades
Alternatives to rolling upgrades

# Upgrade from Windows 2000 on a cluster containing Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To upgrade from Windows 2000 on a cluster containing Message Queuing**

1. Confirm that your hardware is designed for or is compatible with Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition.

2. Take all of the Message Queuing cluster resources offline. For more information on how to do this see "Take a resource offline" in Related Topics.

3. For each Message Queuing resource, take the Network Name resource that depends on it offline also.

4. Upgrade all of the nodes from Windows 2000 to Windows Server 2003. You can perform either a rolling upgrade, or a non-rolling upgrade, whichever is appropriate for your installation. Prior to upgrading, be sure to remove the resources that are not supported by Windows Server 2003. For more information, see Related Topics. Manage Your Server will appear when you initially log on to the newly upgraded node as an Administrator. Close Manage Your Server to continue with the upgrade.

5. Bring the Network Name resources back online. Ensure that the Distributed Transaction Coordinator (DTC) resource is online, then bring the Message Queuing resources online.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- After you upgrade a cluster node from Windows 2000, you must log on with an account that is a member of both the Domain Admins group and the Administrators group on each cluster node before you can restart the Cluster service. This step is required even if the Domain Admins group is already a member of the Administrators group on each cluster node.

- tok:changequorumlater1

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Take a resource offline
Specify which nodes can own a resource
Move a group to another node
Move a resource to another group
Bring a resource online
Checklist: Preparation for upgrading a cluster
Start Setup on an x86-based computer
How rolling upgrades work
Restrictions on rolling upgrades
Resource behavior during rolling upgrades

Upgrades that include a Print Spooler resource
Alternatives to rolling upgrades

# Upgrade a Cluster from Windows NT 4.0

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrade a cluster from Windows NT 4.0**
You cannot perform a rolling upgrade directly from Windows NT Server 4.0, Enterprise Edition to Windows Server 2003, Enterprise Edition. You instead have two options:

- You can maintain cluster availability by performing an upgrade to Windows 2000 first, then to Windows Server 2003, Enterprise Edition.

- For some installations, you can perform a nonrolling upgrade directly from Windows NT Server 4.0, Enterprise Edition to Windows Server 2003, Enterprise Edition, which will not allow you to maintain cluster availability. You can do this if, on each node, the folder that contains the clustering software is called **Cluster** (capitalization is unimportant) and is located in the system root folder, which is the folder where the operating system is installed. If the folder is not called **Cluster** or is not located in the system root folder, you cannot upgrade, but must perform a clean installation.

If you upgrade using Windows 2000, perform an upgrade from Windows NT Server 4.0 to Windows 2000 as documented in the Windows 2000 documentation. Then follow the Windows Server 2003, Enterprise Edition upgrade procedures documented in Upgrade a Cluster from Windows 2000. If you have an IIS resource on your Windows NT Server 4.0 cluster, after you complete both upgrades, you must delete the unsupported IIS resource type by opening the Command Prompt and typing the following command: **Cluster restype "IIS Virtual Root" /delete /type**.

To upgrade without any use of the Windows 2000 product, you must interrupt cluster availability. The steps you perform depend on what resources are included in your cluster. For many installations, the following procedure is appropriate:

- Upgrade a cluster from Windows NT

However, if your cluster contains an IIS resource, you must use the following procedure:

- Upgrade from Windows NT on a cluster containing IIS

If your cluster contains an MSMQ or DTC resource, you must perform the procedure documented in article Q315993, "HOW TO: Upgrade a Windows NT 4.0 Cluster with MSDTC and MSMQ Resources" in the Microsoft Knowledge Base.

# Upgrade a cluster from Windows NT

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To upgrade a cluster from Windows NT**

1. Confirm that your hardware is designed for or is compatible with Windows Server 2003, Enterprise Edition. For more information, see "Checklist: Preparation for upgrading a cluster" in Related Topics.

2. On each node, locate the folder that contains the clustering software. Confirm that the folder is called **Cluster** (capitalization is unimportant) and is located in the system root folder, which is the folder where the operating system is installed. For example, if the default was used when the cluster was installed, the folder path will be C:\Winnt\Cluster. If the folder is not called **Cluster** or is not located in the system root folder, you cannot upgrade the cluster to Windows Server 2003, but must perform a clean installation.

3. Confirm that you do not have an IIS, MSMQ, or DTC resource on your cluster. For more information, see Notes and Related Topics.

4. Ensure that Service Pack 5 or later has been applied to all computers that will be upgraded from Windows NT Server 4.0 to Windows Server 2003, Enterprise Edition.

5. As appropriate, notify users that you will be shutting down the applications they use on the cluster.

6. Stop the applications that are made available through the cluster.

7. Remove any resources that are not supported by Windows Server 2003, Enterprise Edition, including NNTP Service Instance, SMTP Service Instance and Time Service resources.

8. Shut down and turn off all nodes but one, or bring them to a shut-down state appropriate to your method of termination.

   **Caution**

   - If you are using a shared storage device, when you upgrade and power on and start the operating system, it is of vital importance that only one node has access to the cluster disk. Otherwise the cluster disks can become corrupted. To prevent the corruption of the cluster disks, shut down all but one cluster node, or use other techniques (for example, LUN masking, selective presentation, or zoning) to protect the cluster disks, before creating the cluster. Once the Cluster service is running properly on one node, the other nodes can be installed and configured simultaneously.

9. Upgrade the operating system on the node that is running. For information about running Setup, see "Start Setup" in Related Topics.

   The cluster software will be upgraded automatically during the operating system upgrade. Note that you cannot make configuration changes such as configuring cluster disks as dynamic disks. For more information, see "Limitations of server clusters" in Related Topics.

   Manage Your Server will appear when you initially log on to the newly upgraded node as an Administrator. Close Manage Your Server to continue with the upgrade.

10. On the upgraded node, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.

11. Check to see that the cluster disks are online in Cluster Administrator.

**Caution**

- Be sure that the cluster disks are online in Cluster Administrator before continuing to the next step. When the disks are online, it means that the Cluster service is working, which means that only one node can access the cluster storage at any given time. Otherwise the cluster storage could be corrupted.

12. Shut down and turn off the upgraded node, or bring it to a shut-down state appropriate to your method of termination.

13. Turn on the other node in the cluster and upgrade the operating system on that node.

14. After the second node is upgraded, turn on the first node.

15. If you want to add additional nodes to the cluster, add them after the first two nodes are upgraded.

**Important**

- Confirm that your entire cluster solution is compatible with products in the Windows Server 2003 family by clicking the appropriate link in Support resources.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- If your Windows NT Server 4.0 cluster includes a DTC or MSMQ resource, you must take additional steps when upgrading directly from Windows NT Server 4.0 to Windows Server 2003, Enterprise Edition. For more information, see article Q315993, "HOW TO: Upgrade a Windows NT 4.0 Cluster with MSDTC and MSMQ Resources" in the Microsoft Knowledge Base.

- If your goal is to have more than two nodes in the cluster, you must use Fibre Channel (not SCSI) for the cluster storage. Before adding additional nodes, ensure that your entire cluster solution is compatible with products in the Windows Server 2003 family.

- When upgrading from Windows NT Server 4.0 to Windows Server 2003, Enterprise Edition, the Cluster service user account requires the additional user right "Act as part of the operating system". If possible, Setup will grant this user right automatically. If Setup cannot grant the user right, you will be prompted to make this change.

  For security reasons, you must grant this user right to the specific user account that is used by the Cluster service. You cannot correct this problem by granting the user right to a security group of which the user account is a member. Typically, you must grant this user right as a local user right; it cannot be a domain-level user right. However, if your node is a domain controller, you can use the domain level user right.

  For more information on setting user rights, refer to "To edit a security setting" in the Windows 2000 help documentation or on Windows NT Server 4.0, open User Manager for Domains, click the **Help** menu in User Manager and refer to "Managing the User Rights Policy."

- tok:changequorumlater1

- Before you can begin an upgrade from Windows NT Server 4.0 directly to Windows Server 2003, Enterprise Edition, you must apply Service Pack 5 or later.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Upgrade from Windows NT on a cluster containing IIS
Checklist: Preparation for upgrading a cluster
Configure how a service is started
Upgrade a Cluster from Windows NT 4.0
Start Setup on an x86-based computer
Start Setup on an Itanium architecture-based computer
Limitations of server clusters
Stop the Cluster service
Delete a resource

# Upgrade from Windows NT on a cluster containing IIS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To upgrade from Windows NT on a cluster containing IIS**

1. Confirm that your hardware is designed for or is compatible with Windows Server 2003, Enterprise Edition.

2. On each node, locate the folder that contains the clustering software. Confirm that the folder is called **Cluster** (capitalization is unimportant) and is located in the system root folder, which is the folder where the operating system is installed. For example, if the default was used when the cluster was installed, the folder path will be C:\Winnt\Cluster. If the folder is not called **Cluster** or is not located in the system root folder, you cannot upgrade the cluster to Windows Server 2003, but must perform a clean installation.

3. As appropriate, notify users that you will be shutting down the applications they use on the cluster.

4. Ensure that Service Pack 5 or later has been applied to all computers that will be upgraded from Windows NT Server 4.0 to Windows Server 2003, Enterprise Edition.

5. Stop the applications that are made available through the cluster.

6. Remove any resources that are not supported by Windows Server 2003, Enterprise Edition, including NNTP Service Instance, SMTP Service Instance and Time Service resources.

7. Set the Cluster service on all nodes to start manually.

8. Shut down and turn off the node that does not contain the IIS resource, or bring it to a shut-down state appropriate to your method of termination.

   **Caution**

   - If you are using a shared storage device, when you upgrade and power on and start the operating system, it is of vital importance that only one node has access to the cluster disk. Otherwise the cluster disks can become corrupted. To prevent the corruption of the cluster disks, shut down all but one cluster node, or use other techniques (for example, LUN masking, selective presentation, or zoning) to protect the cluster disks, before creating the cluster. Once the Cluster service is running properly on one node, the other nodes can be installed and configured simultaneously.

9. On the running node, note the dependencies of the IIS instance resource. Note resources that depend on the IIS resource and also note what resources IIS itself is dependent on.

10. Take the group containing the IIS instance resource offline. For more information on how to do this see "Take a resource offline" in Related Topics.

11. Remove any dependencies on the IIS instance resource.

12. Delete the IIS instance resource and close Cluster Administrator if it is open. For more information "Delete a resource" in Related Topics.

13. Delete the unsupported resource type: Open Command Prompt. Type the following command and press ENTER:

    **Cluster restype "IIS Virtual Root" /delete /type**

14. Stop the Cluster service on the remaining node.

15. Upgrade the operating system on the node that is running. For information about running Setup, see "Start Setup" in Related Topics.

    The cluster software will be upgraded automatically during the operating system upgrade. Note that you cannot make configuration changes such as configuring cluster disks as dynamic disks. For more information, see "Limitations of server clusters" in Related Topics. After you upgrade, close Manage Your Server if it is displayed.

16. Start the Cluster service on the upgraded node.

17. Reconfigure the Cluster service on the upgraded node to start automatically.

18. Shut down and turn off the upgraded node, or bring it to a shut-down state appropriate to your method of termination.

19. Turn on the other node in the cluster and upgrade the operating system on that node. After you upgrade, close Manage Your Server if it is displayed.

    **Caution**

    - If you are using a shared storage device, when you upgrade and power on and start the operating system, it is of vital importance that only one node has access to the cluster disk. Otherwise the cluster disks can become corrupted. To prevent the corruption of the cluster disks, shut down all but one cluster node, or use other techniques (for example, LUN masking, selective presentation, or zoning) to protect the cluster disks, before creating the cluster. Once the Cluster service is running properly on one node, the other nodes can be installed and configured simultaneously.

20. After the second node is upgraded, start the Cluster service on the second upgraded node. The node automatically rejoins the existing cluster.

21. Reconfigure the Cluster service on the second upgraded node to start automatically.

22. Turn on the first node.

23. On one of the upgraded nodes, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.

24. Check to see that the cluster disks are online in Cluster Administrator.

    **Caution**

    - Be sure that the cluster disks are online in Cluster Administrator before continuing to the next step. When the disks are online, it means that the Cluster service is working, which means that only one node can access the cluster storage at any given time. Otherwise the cluster storage could be corrupted.

25. If you do not already have a Distributed Transaction Coordinator (DTC) resource on the cluster that you are upgrading, create a DTC resource on this cluster.

26. On the node that used to contain the IIS resource, create a Generic Script resource by following the procedure documented in "Checklist: Creating a clustered IIS Web or FTP service" in Related Topics.

    **Notes**

    - As part of this step, recreate the dependencies of the Generic Script resource identical to those of the now deleted IIS resource. Make everything that was dependent on the IIS resource dependent instead on the Generic Script

resource. Also make the Generic Script resource dependent on everything that IIS was dependent on.

- Only Windows Server 2003, Enterprise Edition nodes can be the owner of the Generic Script Resource created in this procedure because this resource type is not supported in Windows NT Server 4.0.

27. Start the W3SVC service on all nodes and set the service to start automatically.

28. Bring the group containing the new Generic Script resource online.

29. Using IIS, start the Web site.

30. If you want to add additional nodes to the cluster, add them after the first two nodes are upgraded.

**Important**

- If your cluster uses SCSI, you can have two nodes in the cluster, but no more. If you want to have more than two nodes in the cluster, you must use Fibre Channel for the cluster storage. Also confirm that your entire cluster solution is compatible with products in the Windows Server 2003 family by clicking the appropriate link in Support resources. Your cluster solution must be in the Cluster category of compatible hardware.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To cluster IIS on Windows Server 2003, Enterprise Edition, you must have a DTC resource on that cluster as well.

- This procedure only applies when you upgrade directly from Windows NT Server 4.0 to Windows Server 2003, Enterprise Edition. If you upgrade from Windows NT Server 4.0 to Windows 2000 and then to Windows Server 2003, Enterprise Edition, follow the steps described in the appropriate topics listed in Related Topics.

- When upgrading from Windows NT Server 4.0 to Windows Server 2003, Enterprise Edition, the Cluster service user account requires the additional user right "Act as part of the operating system". If possible, Setup will grant this user right automatically. If Setup cannot grant the user right, you will be prompted to make this change.

  For security reasons, you must grant this user right to the specific user account that is used by the Cluster service. You cannot correct this problem by granting the user right to a security group of which the user account is a member. Typically, you must grant this user right as a local user right; it cannot be a domain-level user right. However, if your node is a domain controller, you can use the domain level user right.

  For more information on setting user rights, refer to "To edit a security setting" in the Windows 2000 help documentation or on Windows NT Server 4.0, open User Manager for Domains, click the **Help** menu in User Manager and refer to "Managing the User Rights Policy."

- Before you can begin an upgrade from Windows NT Server 4.0 directly to Windows Server 2003, Enterprise Edition, you must apply Service Pack 5 or later.

- tok:changequorumlater1

- For more information regarding hardware compatibility, see the Windows Catalog at the Microsoft Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Configure how a service is started
Take a resource offline
Specify which nodes can own a resource
Stop the Cluster service
Move a group to another node
Move a resource to another group
Bring a resource online
Modify resource dependencies
Delete a resource
Checklist: Preparation for upgrading a cluster
Checklist: Creating a clustered IIS Web or FTP service
Upgrade from Windows 2000 on a cluster containing IIS using the provided script
Upgrade from Windows 2000 on a cluster containing IIS without using the provided script
Start Setup on an x86-based computer

# Concepts for installing and upgrading cluster nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts for installing and upgrading cluster nodes**

- Installing Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition on Cluster Nodes
- Upgrading to Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition on Cluster Nodes
- Resources for installing and upgrading cluster nodes

# Installing Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition on Cluster Nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition on cluster nodes**
This section covers:

- Planning and preparing for cluster installation
- Quorum resource options when installing server clusters

# Planning and preparing for cluster installation

Updated: April 10, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning and preparing for cluster installation**
Before carrying out cluster installation, you need to plan hardware and network details.

**Caution**

- If you are using a shared storage device, before creating a cluster, when you turn on the computer and start the operating system, it is very important that only one node has access to the cluster disk. Otherwise, the cluster disks can become corrupted. To prevent the corruption of the cluster disks, shut down all but one cluster node, or use other techniques (for example, LUN masking, selective presentation, or zoning) to protect the cluster disks before creating the cluster. Once the Cluster service is running properly on one node, the other nodes can be installed and configured simultaneously. Each node of your cluster must be running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition.

In your planning, review the following items:

**Cluster hardware and drivers**
Microsoft supports only complete server cluster systems that are compatible with the Windows Server 2003 family. Confirm that your entire cluster solution is compatible with products in the Windows Server 2003 family by clicking the appropriate link in Support resources.

For cluster disks, you must use the NTFS file system and configure the disks as basic disks. You cannot configure cluster disks as dynamic disks, and you cannot use features of dynamic disks such as spanned volumes (volume sets).

Review the manufacturer's instructions carefully before you begin installing cluster hardware. Otherwise, the cluster storage could be corrupted. If your cluster hardware includes a SCSI bus, be sure to carefully review any instructions about termination of the SCSI bus and configuration of SCSI IDs.

To simplify configuration and eliminate potential compatibility problems, consider using identical hardware for all nodes.

**Network adapters on the cluster nodes**
In your planning, decide what kind of communication each network adapter will carry. The following list provides details about the types of communication that an adapter can carry:

- **Only node-to-node communication (private network).** This implies that the server has one or more additional adapters to carry other communication.

  For node-to-node communication, you connect the network adapter to a private network that is used exclusively within the cluster. Note that if the private network uses a single hub or network switch, that piece of equipment becomes a potential point of failure in your cluster.

  The nodes of a cluster must be on the same subnet but you can use virtual LAN (VLAN) switches on the interconnects between two nodes. If you use a VLAN, the point to point, round-trip latency must be less than 1/2 second and the link between two nodes must appear as a single point-to-point connection from the perspective of the Windows operating system running on the nodes. To avoid single points of failure, use independent VLAN hardware for the different paths between the nodes.

  If your nodes use multiple private (node-to-node) networks, it is a best practice for the adapters for those networks to use static IP addresses instead of DHCP.

- **Only client-to-cluster communication (public network).** This implies that the server has one or more additional adapters to carry other communication.

- **Both node-to-node and client-to-cluster communication (mixed network).** When you have multiple network adapters per node, a network adapter that carries both kinds of communication can serve as a backup for other network adapters.

- **Communication unrelated to the cluster.** If a clustered node also provides services unrelated to the cluster, and there are enough adapters in the cluster node, you might want to use one adapter for carrying communication unrelated to the cluster.

The nodes of a cluster must be connected by two or more local area networks (LANs); at least two networks are required to prevent a single point of failure. A server cluster whose nodes are connected by only one network is not a supported configuration. The adapters, cables, hubs, and switches for each network must fail independently. This usually implies that the components of any two networks must be physically independent.

At least two networks must be configured to handle **All communications (mixed network)** or **Internal cluster communications only (private network)**.

The recommended configuration for two adapters is to use one adapter for the private (node-to-node only) communication and the other adapter for mixed communication (node-to-node plus client-to-cluster communication).

Consider choosing a name for each connection that tells what it is intended for. The name makes it easier to identify the connection whenever you are configuring the server.

**Notes**

- If you use fault tolerant network adapters, create multiple private networks instead of a single fault-tolerant network.
- Do not use teaming network adapters on the private network.
- Do not configure a default gateway or DNS or WINS server on the private network adapters. Do not configure private network adapters to use name resolution servers on the public network; otherwise, a name resolution server on the public network might map a name to an IP address on the private network. If a client then received that IP address from the name resolution server, it might fail to reach the address because no route from the client to the private network address exists.
- Configure WINS and/or DNS servers on the public network adapters. If Network Name resources are used on the public networks, set up the DNS servers to support dynamic updates; otherwise the Network Name resources may not fail over correctly. Also, configure a default gateway on the public network adapters. If there are multiple public networks in the cluster, configure a default gateway on only one of these.
- The adapters on a given node must connect to networks in different subnets.
- When you use either the New Server Cluster Wizard or the Add Nodes Wizard to install clustering on a node that contains two network adapters, by default the wizard configures both of the network adapters for mixed network communications. As a best practice, reconfigure one adapter for private network communications only. For more information, see Change how the cluster uses a network.
- Manually configure the communication settings, such as Speed, Duplex Mode, Flow Control and Media Type of each cluster network adapter. Do not use automatic detection. You must configure all of the cluster network adapters to use the same communication settings.
- Confirm that your entire cluster solution is compatible with the products in the Windows Server 2003 family by clicking the appropriate link in Support resources.

**Cluster IP address**
Obtain a static IP address for the cluster itself. You cannot use DHCP for this address.

**IP addressing for cluster nodes**
Determine how to handle the IP addressing for the individual cluster nodes. Each network adapter on each node requires IP addressing. It is a best practice to assign each network adapter a static IP address. As an alternative, you can provide IP addressing through DHCP. If you use static IP addresses, set the addresses for each linked pair of network adapters (linked node-to-node) to be on the same subnet.

Note that if you use DHCP for the individual cluster nodes, it can act as a single point of failure. That is, if you set up your

cluster nodes so that they depend on a DHCP server for their IP addresses, temporary failure of the DHCP server can mean temporary unavailability of the cluster nodes. When deciding whether to use DHCP, evaluate ways to ensure availability of DHCP services, and consider the possibility of using long leases for the cluster nodes. This helps to ensure that they always have a valid IP address.

**Cluster name**
Determine or obtain an appropriate name for the cluster. This is the name administrators will use for connections to the cluster. (The actual applications running on the cluster typically have different network names.) The cluster name must be different from the domain name, from all computer names on the domain, and from other cluster names on the domain.

**Computer accounts and domain assignment for cluster nodes**
Make sure that the cluster nodes all have computer accounts in the same domain. Cluster nodes cannot be in a workgroup.

**Operator user account for installing and configuring the Cluster service**
To install and configure the Cluster service, you must be using an account that is in the local Administrators group on each node. As you install and configure each node, if you are not using an account in the local Administrators group, you will be prompted to provide the logon credentials for such an account.

**Cluster service user account**
Create or obtain the Cluster service user account. This is the name and password under which the Cluster service will run. You need to supply this user name and password during cluster installation.

It is best if the Cluster service user account is an account not used for any other purpose. If you have multiple clusters, set up a unique Cluster service user account for each cluster. The account must be a domain account; it cannot be a local account. However, do not make this account a domain administrator account because it does not need domain administrator user rights.

As part of the cluster setup process, the Cluster service user account is added to the local Administrators group on each node. As well as being a member of the local administrators group, the Cluster service user account requires an additional set of user rights:

- Act as part of the operating system
- Back up files and directories
- Adjust memory quotas for a process
- Increase scheduling priority
- Log on as a service
- Restore files and directories

In addition, by default, the Cluster service account inherits the following user rights as a result of being a member of the local Administrators group:

- Manage auditing and security log
- Debug programs
- Impersonate a client after authentication

If your organization has removed these user rights from the default set of privileges assigned to the local Administrators group, you need to specifically assign these user rights to the Cluster service account.

The preceding user rights are granted to the Cluster service user account as part of the cluster setup process. Be aware that the Cluster service user account will continue to have these user rights even after all nodes are evicted from the cluster. The risk that this presents is mitigated by the fact that these user rights are not granted domain wide, but rather only locally on each former node. However, remove this account from each evicted node if it is no longer needed.

Be sure to keep the password from expiring on the Cluster service user account (follow your organization's policies for password renewal). For more information, see Passwords.

**Volume for important cluster configuration information (checkpoint and log files)**
Plan on setting aside a volume on your cluster storage for holding important cluster configuration information. This information makes up the cluster quorum resource, which is needed when a cluster node stops functioning. The quorum resource provides node-independent storage of crucial data needed by the cluster.

The recommended minimum size for the volume is 500 MB. It is recommended that you do not store user data on any volume in the quorum resource. Do not use Shadow Copies for Shared Folders for the quorum resource. If you plan to put the quorum resource on a disk with multiple NTFS partitions, ensure that all partitions on the disk are assigned drive letters.

### Note

- When planning and carrying out disk configuration for the cluster disks, configure them as basic disks with all partitions formatted as NTFS (they can be either compressed or uncompressed). Partition and format all disks on the cluster storage device before adding the first node to your cluster. Do not configure them as dynamic disks, and do not use spanned volumes (volume sets) or Remote Storage on the cluster disks. Cluster disks on the cluster storage device must be partitioned as master boot record (MBR) and not as GUID partition table (GPT) disks.

The following section describes the physical installation of the cluster storage.

**Beginning the installation of the cluster hardware**
The steps you carry out when first physically connecting and installing the cluster hardware are crucial. Be sure to follow the hardware manufacturer's instructions for these initial steps.

### Important

- Carefully review your network cables after connecting them. Make sure no cables are crossed by mistake (for example, private network connected to public).

**Initial steps to carry out in the BIOS or EFI when using a SCSI shared storage device**
If you are using a SCSI shared storage device, when you first attach your cluster hardware (the shared bus and cluster storage), be sure to work only from the firmware configuration screens on the cluster nodes (a node is a server in a cluster). On a 32-bit computer, use the BIOS configuration screens. On an Itanium architecture-based computer, use the Extensible Firmware Interface (EFI) configuration screens. The instructions from your manufacturer describe whether these configuration screens are displayed automatically or whether you must, after turning on the computer, press specific keys to access them. Follow the manufacturer's instructions for completing the BIOS or EFI configuration process. Remain in the BIOS or EFI configuration screens, and do not allow the operating system to start, during this initial installation phase. Complete the following steps while the cluster nodes are still displaying BIOS or EFI configuration screens, before starting the operating system on the first cluster node.

### Important

- Make sure you understand and follow the manufacturer's instructions for termination of the SCSI bus.

- Make sure that each device on the shared bus (both SCSI controllers and hard disks) has a unique SCSI ID. If the SCSI controllers all have the same default ID (often it is SCSI ID 7), change one controller to a different SCSI ID, such as SCSI ID 6. If there is more than one disk that will be on the shared SCSI bus, each disk must also have a unique SCSI ID. In addition, make sure that the bus is not configured to reset SCSI IDs automatically on startup (otherwise the IDs will change from the settings you specify).

- Ensure that you can scan the bus and see the drives from all cluster nodes (while remaining in the BIOS or EFI configuration screens).

**Initial steps to carry out in the BIOS or EFI when using a fibre channel shared storage device or no shared storage device**

- Turn on a single node. Leave all other nodes turned off.

- During this initial installation phase, remain in the BIOS or Extensible Firmware Interface (EFI) configuration process, and do not allow the operating system to start. While viewing the BIOS or EFI configuration screens, ensure that you can scan the bus and see the drives from the active cluster node. On a 32-bit computer, use the BIOS configuration screens. On an Itanium architecture-based computer, use the EFI configuration screens. Consult the instructions from your manufacturer

to determine whether these configuration screens are displayed automatically or whether you must, after turning on the computer, press specific keys to access them. Follow the manufacturer's instructions for completing the BIOS or EFI configuration process.

**Final steps to complete the installation**

If you have not already installed Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, on the first cluster node, install it before proceeding. After you complete the BIOS or EFI configuration, start the operating system on one cluster node only and complete the configuration of the Cluster service using Cluster Administrator.

With the Cluster Administrator New Server Cluster Wizard, you can choose between Typical (full) configuration and Advanced (minimum) configuration options. Typical configuration is appropriate for most installations and results in a completely configured cluster. Use the Advanced configuration option only for clusters that have complex storage configurations that the New Server Cluster Wizard cannot validate or for configurations in which you do not want the cluster to manage all of the storage. The following examples describe each situation:

- In some complex storage solutions, such as a fiber channel switched fabric that contains several switches, a particular storage unit might have a different identity on each computer in the cluster. Although this is a valid storage configuration, it violates the storage validation heuristics in the New Server Cluster Wizard. If you have this type of storage solution, you might receive an error when you are trying to create a cluster using the Typical configuration option. If your storage configuration is set up correctly, you can disable the storage validation heuristics and avoid this error by restarting the New Server Cluster Wizard, selecting the Advanced configuration option instead.

- On particular nodes in a cluster, you may want to have some disks that are to be clustered and some disks that are to be kept private. The Typical configuration option configures all disks as clustered disks and creates cluster resources for them all. However, with the Advanced configuration option, you can keep certain disks private because this configuration creates a cluster in which only the quorum disk is managed by the cluster (if you chose to use a physical disk as the quorum resource). After the cluster is created, you must then use Cluster Administrator to add any other disks that you want the cluster to manage.

For more information, see Storage configuration options.

**Important**

- If you are using a shared storage device: Before creating a cluster, when you turn the computer on and start the operating system, it is very important that only one node has access to the cluster disk. Otherwise, the cluster disks can become corrupted. To prevent the corruption of the cluster disks, shut down all but one cluster node, or use other techniques (for example, LUN masking, selective presentation, or zoning) to protect the cluster disks before creating the cluster. Also, before starting the installation of the second and subsequent nodes, ensure that all disks that are to be managed by the cluster have disk resources associated with them. If these disks do not have disk resources associated with them at this time, the disk data will be corrupted because the disks will not be protected and multiple nodes will attempt to connect to them at the same time.

- Do not use Manage Your Server or the Configure Your Server Wizard to configure cluster nodes.

For information about further steps for installing a cluster, see Checklist: Validating your clustering system and Checklist: Planning and creating a server cluster.

For information about unattended setup of a cluster, see Create a cluster and "Designing Unattended Installations" at the Microsoft Windows Server System Web site.

# Quorum resource options when installing server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quorum resource options when installing server clusters**

With server clusters on the Windows Server 2003 operating systems, you can now choose between three ways to set up the quorum resource (the resource that maintains the definitive copy of the cluster configuration data and that must always be available for the cluster to run).

The first is a single node server cluster, which has been available in the past and continues to be supported. A single node cluster is often used for development and testing and can be configured with, or without, external cluster storage devices. For single node clusters without an external cluster storage device, the local disk is configured as the cluster quorum device.

The second option is a single quorum device server cluster, which has also been available in earlier Windows versions. This model places the cluster configuration data on a shared cluster storage device that all nodes can access. The general topology is:




This is the most common model and is recommended for most situations. You might choose the single quorum device model if all of your cluster nodes are in the same location and you want to take advantage of the fact such a cluster continues supporting users even if only one node is running.

The third option, which is new for Windows Server 2003, is a "majority node set." A majority node set is a single quorum resource from a Server Cluster perspective; however, the cluster configuration data is actually stored on multiple disks across the cluster. The majority node set resource ensures that the cluster configuration data is kept consistent across the different disks. This allows cluster topologies as follows:




In the majority node set model, every node in the cluster uses a directory on its own local system disk to store the cluster configuration data. If the configuration of the cluster changes, that change is reflected across the different disks. Be aware that it is also possible to have shared storage devices in a majority node set cluster. The exact configuration depends on your installation's requirements.

Only use a majority node set cluster in targeted scenarios, such as:

- Geographically dispersed cluster: A cluster that spans multiple sites.
- Eliminating single points of failure: Although when using a single cluster storage device the quorum disk itself can be made highly available via RAID, the controller port or the Host Bus Adapters (HBA) itself may be a single point of failure.
- Clusters with no shared disks: There are some specialized configurations that need tightly consistent cluster features without having shared disks.
- Clusters that host applications that can fail over, but where there is some other, application-specific way, to replicate or mirror data between nodes.

Do not configure your cluster as a majority node set cluster unless it is part of a cluster solution offered by your Original Equipment Manufacturer (OEM), Independent Software Vendor (ISV), or Independent Hardware Vendor (IHV).

**Cluster model considerations**
Before implementing your cluster, consider what type of quorum resource solution you plan to use. Take into consideration the following differences between single quorum device clusters and majority node set:

### Note

- The following information is presented to help you make basic decisions about the placement and management of your cluster nodes and quorum resource. It does not provide all the details about the requirements for each cluster model, or how each model handles failover situations. If you are not sure which model to use or where you want to place your cluster nodes, refer to the information in Choosing a Cluster Model.

**Node failover behavior**
The failover behavior of the majority node set is significantly different from the behavior of the single quorum device model:

- Using the single quorum device model, you can maintain cluster availability with only a single operational node.

- If you use a majority node set, more than half, or (*Number of nodes configured in the cluster*/2) + 1 nodes must be operational in order to maintain cluster availability. The following table shows the number of node failures that a given majority node set cluster can tolerate yet continue to operate:

| Number of nodes configured in the cluster | Number of node failures allowed before cluster failure | Number of nodes needed to continue cluster operations |
|---|---|---|
| 1 | 0 | 1 |
| 2 | 0 | 2 |
| 3 | 1 | 2 |
| 4 | 1 | 3 |
| 5 | 2 | 3 |
| 6 | 2 | 4 |
| 7 | 3 | 4 |
| 8 | 3 | 5 |

**Geographic Considerations**
You would commonly use a single quorum resource model if all nodes in your cluster will be in the same geographical location. As part of this requirement, your nodes must be connected to the same physical storage device.

A majority node set on the other hand would typically be appropriate if you have geographically dispersed nodes. The cluster configuration data is stored locally on each node on a file share that is shared out to the other nodes on the network. However, those shares must always be accessible or nodes can fail.

There are other specific requirements for geographically dispersed clusters, including the requirement that round-trip latency of the network between cluster nodes be a maximum of 500 milliseconds. For information on cluster solutions that meet all requirements for a geographically dispersed cluster, refer to the hardware compatibility information, which you can find by clicking the appropriate link in Support resources.

**Hardware**
Microsoft supports only complete server cluster systems that are compatible with the Windows Server 2003 family of products. Confirm that your entire cluster solution is compatible with products in the Windows Server 2003 family by clicking the appropriate link in Support resources.

For more information about majority node set clusters, see Choosing a Cluster Model and Create a cluster.

**Important**

- Before implementing a majority node set, it is highly recommended that you read the information in Choosing a Cluster Model. This information will help you understand all the considerations, requirements, and restrictions for each type of quorum resource solution.

# Upgrading to Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition on Cluster Nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrading to Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition on cluster nodes**
If you are upgrading from Windows 2000 to Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition on cluster nodes, you might be able to perform a *rolling upgrade* of the operating system. In a rolling upgrade, you sequentially upgrade the operating system on each node, making sure that one node is always available to handle client requests. When you upgrade the operating system, the Cluster service is automatically upgraded also. A rolling upgrade maximizes availability of clustered services and minimizes administrative complexity.

If you are upgrading from Windows NT 4.0, Enterprise Edition to Windows Server 2003, Enterprise Edition on cluster nodes, you cannot perform a rolling upgrade. For more information about how to perform an upgrade from Windows NT 4.0, Enterprise Edition, see Upgrade a Cluster from Windows NT 4.0.

For more information, see the following topics:

- For a diagram and description of a rolling upgrade, see How rolling upgrades work.
- To determine whether you can perform a rolling upgrade, see Restrictions on rolling upgrades.
- For information about support for various types of resources during a rolling upgrade, see the following:
  - Operation of new resource types in mixed-version clusters
  - Resource behavior during rolling upgrades
  - Upgrades that include a Print Spooler resource
- For information about ways to upgrade your cluster nodes if you cannot perform a rolling upgrade, see Alternatives to rolling upgrades.

# How rolling upgrades work

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How rolling upgrades work**
This topic describes rolling upgrades on server clusters. To decide whether you can perform a rolling upgrade, see Restrictions on rolling upgrades. For information about methods, restrictions, and alternatives to rolling upgrades, see the links at the end of this topic.

There are two major advantages to a rolling upgrade. First, there is a minimal interruption of service to clients. (However, server response time might decrease during the phases in which a fewer number of nodes handle the work of the entire cluster.) Second, you do not have to recreate your cluster configuration. The configuration remains intact during the upgrade process.

In this example, the rolling upgrade starts with two cluster nodes that are running Windows 2000 named *Node 1* and *Node 2*. The following diagram depicts the process:




**Phase 1: preliminary**
Each node runs Windows 2000 Advanced Server with the following software:

- The Cluster service component, one of the optional components of Windows 2000 Advanced Server.
- Applications that support a rolling upgrade. For more information, see the application documentation and Resource behavior during rolling upgrades.

At this point, your cluster is configured so that each node handles client requests (an active/active configuration).

**Phase 2: upgrade node 1**
Node 1 is paused, and Node 2 handles all cluster resource groups while you upgrade the operating system of Node 1 to Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

You are now running a mixed version cluster. A mixed version cluster is a cluster in which the nodes are running different versions of the operating system. Although mixed version clusters are supported by Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, they are not the optimal configuration. It is recommended that you upgrade all nodes to the same operating system by completing the final two phases of the rolling upgrade.

When your cluster is in a mixed version state, it is recommended that you do not make any changes to the cluster configuration, such as adding or removing nodes. It is also recommended that you do not make changes to any of the cluster resources while your cluster is in a mixed version state.

**Phase 3: upgrade node 2**
Node 1 rejoins the cluster. Node 2 is paused and Node 1 handles all cluster resource groups while you upgrade the operating system on Node 2.

**Phase 4: final**

Node 2 rejoins the cluster, and you redistribute the resource groups back to the active/active cluster configuration.

**Notes**

- If your goal is to have more than two nodes in the cluster, you must use Fibre Channel (not SCSI) for the cluster storage. Add the additional nodes after completing the rolling upgrade. Before adding additional nodes, confirm that your entire cluster solution is compatible with products in the Windows Server 2003 family by clicking the appropriate link in Support resources. Your cluster solution must be in the Cluster category of compatible hardware.

- For cluster disks, you must use the NTFS file system and configure the disks as basic disks. You cannot configure cluster disks as dynamic disks, and you cannot use features of dynamic disks such as spanned volumes (volume sets). For more information, see Limitations of server clusters.

For more information about rolling upgrades, see:

- Perform a rolling upgrade from Windows 2000
- Restrictions on rolling upgrades
- Operation of new resource types in mixed-version clusters
- Resource behavior during rolling upgrades
- Upgrades that include a Print Spooler resource
- Alternatives to rolling upgrades

# Restrictions on rolling upgrades

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restrictions on rolling upgrades**

There are several basic restrictions to the rolling-upgrade process. The most basic restrictions are as follows:

- You can perform a rolling upgrade only if you are upgrading from Windows 2000 on the cluster nodes. You cannot perform a rolling upgrade if you are upgrading from Windows NT Server 4.0, Enterprise Edition.

- You cannot mix x86-based and Itanium-based computers within the same cluster. Also, you cannot mix x86-based and x64-based computers within the same cluster. Therefore, you cannot perform a rolling upgrade from an x86-based computer to either an Itanium-based or x64-based computer.

The remaining restrictions involve the beginning of Phase 3, in which you operate a *mixed-version cluster*: a cluster in which the nodes run different versions of the operating system. In order for a mixed-version cluster to work, the different versions of the software running on each node must be prepared to communicate with one another. This requirement leads to several basic restrictions on the rolling-upgrade process.

- For a successful rolling upgrade, every resource that the cluster manages must be capable of a rolling upgrade. For more information, see Resource behavior during rolling upgrades.

- During the mixed-version phase of a rolling upgrade, when the cluster nodes are running different versions of the operating system, do not change the settings of resources (for example, do not change the settings of a printer resource). Also, do not make any changes to the quorum resource.

- During the mixed-version phase of a rolling upgrade, minor problems might occur. We recommend that you complete the upgrade rather than trying to troubleshoot symptoms. Completing the upgrade might correct the problems.

If the preceding restrictions cannot be met, do not perform a rolling upgrade. For more information, see Alternatives to rolling upgrades.

# Operation of new resource types in mixed-version clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Operation of new resource types in mixed-version clusters**

The operation of a mixed-version cluster is complicated if a resource type that you add to the cluster is supported in one version of the operating system but not the other. For example, the Cluster service in Windows Server 2003 supports the Generic Script resource type. However, older versions of the Cluster service do not support it. A mixed-version cluster can run a Generic Script resource on a node running Windows Server 2003 but not on a node running Windows 2000.

The Cluster service transparently sets the possible owners of new resource types to prevent these resources from failing over to a Windows 2000 node of a mixed-version cluster. In other words, when you view the possible owners of a new resource type, a Windows 2000 node will not be in the list, and you will not be able to add this node to the list. If you create such a resource during the mixed-version phase of a rolling upgrade, the resource groups containing those resources will not fail over to a Windows 2000 node.

# Resource behavior during rolling upgrades

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource behavior during rolling upgrades**

Although the Cluster service supports rolling upgrades, not all applications have seamless rolling upgrade behavior. The following table describes the resources that are supported during a rolling upgrade. If you have a resource that is not fully supported during rolling upgrades, see Alternatives to rolling upgrades.

**Note**

- You cannot perform a rolling upgrade on a cluster running Windows NT Server 4.0. Only clusters running Windows 2000 support rolling upgrades to Windows Server 2003, Enterprise Edition.

| Resource | Rolling upgrade notes |
|---|---|
| DHCP | DHCP is not fully supported during rolling upgrades. You must instead perform a "Last node rolling upgrade," as described in Perform a last node rolling upgrade from Windows 2000. |
| File Share | Supported during rolling upgrades. |
| IIS | IIS 6.0 is not supported during rolling upgrades. When upgrading a cluster that contains an IIS resource, you must follow the procedure described in Upgrade from Windows 2000 on a cluster containing IIS without using the provided script or Upgrade from Windows 2000 on a cluster containing IIS using the provided script. |
| IP Address | Supported during rolling upgrades. |
| Local Quorum | Supported during rolling upgrades. |
| MSDTC | Supported during rolling upgrades. However, do not create or delete the MSDTC resource while you are operating as a mixed version cluster. Instead create the MSDTC resource before starting the upgrade on any node, or wait until the upgrade is complete on all nodes. For more information on mixed version clusters, see How rolling upgrades work and Restrictions on rolling upgrades. |
| Message Queuing (also known as MSMQ) | Microsoft Message Queuing is not fully supported during a rolling upgrade. To upgrade a cluster that includes Message Queuing, see Upgrade from Windows 2000 on a cluster containing Message Queuing. |
| Network Name | Supported during rolling upgrades. |
| NNTP Service Instance | NNTP Service Instance is not supported in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. You must remove any NNTP Service Instance resources prior to upgrading. |
| Physical Disk | Supported during rolling upgrades. |
| Print Spooler | The only Print Spooler resources supported during a rolling upgrade are those on LPR ports or standard port monitor ports. See Upgrades that include a Print Spooler resource. |
| SMTP Service Instance | SMTP Service Instance is not supported in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. You must remove any SMTP Service Instance resources prior to upgrading. |

| | |
|---|---|
| Time Service | Time Service is not supported in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. You must remove any Time Service resources prior to upgrading. |
| WINS | WINS is not fully supported during rolling upgrades. You must instead perform a "Last node rolling upgrade," as described in Perform a last node rolling upgrade from Windows 2000. |
| Other resource types | See Relnotes.htm in the \Docs folder on the Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, setup CD. Also see the product documentation that comes with the application or resource. |

When your cluster is in a mixed version state, it is recommended that you do not make any changes to the cluster configuration, such as adding or removing nodes. It is also recommended that you do not make changes to any of the cluster resources while your cluster is in a mixed version state.

**Last node rolling upgrades**

If you have resources on your cluster such as DHCP and WINS that are not fully supported during a rolling upgrade and are not described by other technology specific procedures (such as IIS), you can perform a "last node rolling upgrade."

Last node rolling upgrades are similar to standard rolling upgrades. In a last node rolling upgrade, you upgrade all nodes except the node which contains the unsupported resources. After all other nodes are upgraded, you then upgrade the final node containing the identified resources.

Last node rolling upgrades are a way to accommodate resources that cannot fail over to an older operating system from a node running Windows Server 2003, Enterprise Edition. Be aware that these resources can, however, fail over the other way, that is, from the older operating system to Windows Server 2003, Enterprise Edition. For more information on how to perform a last node rolling upgrade, see Perform a last node rolling upgrade from Windows 2000.

# Upgrades that include a Print Spooler resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrades that include a Print Spooler resource**

If you want to perform a rolling upgrade of a cluster that has a Print Spooler resource, you must consider two issues.

First, the Print Spooler resource only supports upgrades (including rolling upgrades or any other kind of upgrade) of printers on cluster-supported ports (LPR or Standard Port Monitor ports). For information about what to do if your printer is not supported, see Alternatives to rolling upgrades.

Second, when you operate a mixed-version cluster including a Print Spooler resource, note the following:

- Do not change printer settings in a mixed-version cluster with a Print Spooler resource.
- Avoid adding new printers while you are performing a rolling upgrade. If you do add a new printer, when you install the drivers for that printer, be sure to install both the driver for Windows 2000 and the driver for Windows Server 2003 on all nodes.
- If printing preferences or defaults are important, be sure to check them. Printing preferences in Windows Server 2003 will not necessarily correspond to document defaults for the same printer in Windows 2000. This can be affected by differences in the drivers for the two operating systems.

When the rolling upgrade is complete and both cluster nodes are running the updated operating system, you can make any modifications you choose to your printer configuration.

# Alternatives to rolling upgrades

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Alternatives to rolling upgrades**

Certain situations do not allow you to perform a standard rolling upgrade. In these cases, you must perform alternative procedures. These situations are:

- Upgrading from Windows NT Server 4.0, Enterprise Edition. For more information, see Upgrade a Cluster from Windows NT 4.0.

- Upgrading a node that contains resources that are not supported during rolling upgrades. To determine whether your resources are supported during rolling upgrades, see Resource behavior during rolling upgrades. Also be sure to read Relnotes.htm in the \Docs folder of the operating system setup CD, as well as the product documentation that comes with the application or resource. If you have a resource that is not fully supported during rolling upgrades and is not described by a technology specific procedure, such as Internet Information Server (IIS), see Perform a last node rolling upgrade from Windows 2000.

# Resources for installing and upgrading cluster nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This topic provides sources of additional information about installing and upgrading on cluster nodes.

- Microsoft supports only complete server cluster systems that are compatible with the Windows Server 2003 family. Confirm that your entire cluster solution is compatible with products in the Windows Server 2003 family by clicking the appropriate link in Support resources.
- For the most recent information on compatible and certified applications for products in the Windows Server 2003 family, see the software compatibility section of the Microsoft Web site.
- To better ensure high availability, check for enhancements and updates to Windows software and other Microsoft software. See Windows Update at the Microsoft Web site.
- For additional information about deploying clustering, see "Planning Server Deployments" at the Microsoft Windows Server System Web site.
- For additional general information about clustering and the Windows Deployment and Resource Kits, which provide technical background information on products in the Windows Server 2003 family, see the following:
    - The *Windows Server 2003 Deployment Kit* at the Microsoft Windows Server System Web site.
    - The *Windows Server 2003 Technical Reference* at the Microsoft Windows Server System Web site.
    - The *Windows Server TechCenter* at the Microsoft Web site.
    - Using the Windows Deployment and Resource Kits
- Microsoft TechNet is a service that contains resources for Information Technology (IT) professionals. This includes access to white papers and articles about clustering written by support professionals at Microsoft, career and industry information, discussion groups, technical feedback, and software. For more information, see TechNet at the Microsoft Web site.
- The Information Technology Infrastructure Library Web site (ITIL) Web site provides guidance for Information Technology professionals and includes detailed information regarding high availability. The following is an excerpt from the site: "ITIL consists of a series of books giving guidance on the provision of quality IT services, and on the accommodation and environmental facilities needed to support information technology." (Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.)

# Troubleshooting cluster node installations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting cluster node installations**
What problem are you having?

- During cluster creation, the New Server Cluster Wizard did not detect all of the potential cluster disks.
- During cluster creation, the New Server Cluster Wizard configured all of the available storage devices to be managed by the cluster, including storage that you did not want to be managed by the cluster.
- When trying to add a node to a cluster, you receive an error stating that a multinode cluster cannot be created because there is not a quorum-capable resource common to all nodes.
- During cluster creation, either the New Server Cluster Wizard or the Add Nodes Wizard did not detect all of the potential network adapters.
- During cluster creation, you are experiencing problems creating the first node of the cluster.
- During cluster creation, you are experiencing problems creating subsequent nodes after creating the first node of the Cluster.
- During cluster creation, the Physical Disk resources are not automatically created for cluster disks.
- After setting up a new cluster, Cluster Administrator cannot connect to the cluster.

**During cluster creation, the New Server Cluster Wizard did not detect all of the potential cluster disks.**
**Possible Cause:** Your hardware, device driver, or software is not configured correctly.

**Solution:** Check Disk Management to see whether the disks you want to use for shared storage are available.

**See also:** Disk Management

**Possible Cause:** You are using incompatible disk hardware or drivers.

**Solution:** Check your disk cabling to make sure that the cables are connected correctly between the node in question and the shared storage.

Confirm that your hardware, including your cluster disks and other cluster hardware, is compatible with Windows Server 2003, Datacenter Edition, by clicking the appropriate link in Support resources.

Make sure that the host adapter is configured correctly according to the manufacturer's instructions, and that all disks are visible to the host adapter.

Make sure that you have the drivers you need in order to use the disks with the Cluster service. Drivers are available from your hardware manufacturer. Review the manufacturer's instructions carefully before you begin installing cluster hardware.

**Note**

- If you make any changes to your disk hardware or drivers, you must restart the New Server Cluster Wizard for the changes to be detected.

**Possible Cause:** Your cluster storage disks are not formatted as NTFS.

**Solution:** Make sure all partitions on the cluster storage disks are formatted to use the NTFS file system. Any disk with a FAT or FAT32 partition will not be recognized during cluster creation.

**See also:** NTFS

**Possible Cause:** Your cluster storage disks are not configured as basic MBR disks.

**Solution:** Make sure your cluster storage disks are configured as basic MBR disks. Windows Server 2003, Datacenter Edition, offers two types of storage configurations: basic disk and dynamic disk. Basic disks use the same disk structures used by Microsoft Windows NT. Dynamic disks were introduced in Windows 2000.

Cluster servers can recognize only basic master boot record (MBR) disks. If a disk you want to use as shared storage is configured as a dynamic or GUID partition table (GPT) disk, change it back to a basic MBR disk.

**See also:** Converting a basic disk to dynamic; Change a GUID partition table disk into a master boot record disk

**During cluster creation, the New Server Cluster Wizard configured all of the available storage devices to be managed by the cluster, including storage that you did not want to be managed by the cluster.**
**Cause:** With the Cluster Administrator New Server Cluster Wizard, you can choose between Typical (full) configuration and Advanced (minimum) configuration options. The Typical configuration configures all disks as clustered disks and creates cluster resources for them all. The only exception is the system disk, which contains the operating system and is the disk from which the computer boots from. The system disk will not be automatically clustered.

**Solution:** Use the Advanced (minimum) configuration option. With the Advanced configuration option, you can keep certain disks private because the configuration creates a cluster in which only the quorum disk is managed by the cluster (if you chose to use a physical disk as the quorum resource). After the cluster is created, you must then use Cluster Administrator to add any other disks that you want the cluster to manage.

**See also:** Planning and preparing for cluster installation; Storage configuration options

**When trying to add a node to a cluster, you receive an error stating that a multinode cluster cannot be created because there is not a quorum-capable resource common to all nodes.**
**Cause:** The wizard was not able to validate your storage configuration. In some complex storage solutions that contain several switches, such as a fiber channel switched fabric, a particular storage unit might have a different identity on each computer in the cluster. Although this is a valid storage configuration, it violates the storage validation heuristics in the New Server Cluster Wizard and the Add Nodes Wizard.

**Solution:** Verify that your storage configuration is set up correctly. Then, if you are in the Add Nodes Wizard, disable the storage validation heuristics by using the Back button to return to the Select Computer page. On the Select Computer page, click Advanced and select the Advanced (minimum) configuration option. If you are not in the Add Nodes Wizard, restart the wizard and use the Advanced configuration option as previously described.

**See also:** Planning and preparing for cluster installation

**During cluster creation, either the New Server Cluster Wizard or the Add Nodes Wizard did not detect all of the potential network adapters.**
**Possible Cause:** The cluster may not be able to detect a network composed of a crossover cable.

A two-node cluster may use a crossover cable to implement a private network. When the cluster is created on the first node, the second node should be turned on, but stopped in the BIOS or at the startup menu. In this state, the Media Sense feature of Windows may not recognize the network adapter as connected.

When this happens:

- In Network Connections, the icon for the connection that represents the network adapter is shown with a red "X".
- The taskbar displays an icon that, when you pause the mouse on it, displays **Network Cable unplugged**.

**Solution:** If you continue creating the cluster, the crossover cable will be detected when you start the second node. The network will be established in the default mode, which is mixed. You can then change the network mode to private using Cluster Administrator.

**See also:** Manage Networks

**Possible Cause:** Your system's network connections are not running. If the following error message appears in the New Server Cluster Wizard tree, your network connections may not be running:

"Warning: This system is not attached to all defined cluster networks."

**Solution:** Make sure that your network connections are enabled. Otherwise, the Cluster service cannot detect and use them.

To check your network connections:

1. Open Network Connections in Control Panel.
2. Double-click the connection that you want to check.
3. On the **General** tab, check the status of the connection. If the connection is not started, start it.

You must click **Reanalyze** in the New Server Cluster Wizard for the changes to be detected.

**See also:** TCP/IP

**Possible Cause:** Your network hardware is not installed correctly.

**Solution:** Verify that your network cables are connected correctly. Make sure that no cables are crossed by mistake (private-to-public, or public-to-private), and that each network adapter is connected as you planned. With more than two nodes, use a dedicated hub (not a switch) for the private network.

Set all multispeed network adapters on a private or mixed network to the same speed. Do not use the autodetect setting.

If fault-tolerant network adapters are used, create multiple private networks instead of a single fault-tolerant network.

Make sure all network adapters are correctly seated and the correct drivers are installed in accordance with the manufacturer's instructions. Confirm that your hardware is compatible with Windows Server 2003, Datacenter Edition, by clicking the appropriate link in Support resources.

**Possible Cause:** The Cluster networks are not on the correct subnet. All network adapters designated for cluster use on a single node need to be on different subnets. Otherwise, only the first network adapter on a subnet is detected by either the New Server Cluster Wizard or the Add Nodes Wizard, and is used for either cluster or client communication (depending on which role you chose).

However, in order for subsequent nodes to pick up the network roles defined on the first node in the cluster, network adapters in each node must be on the same subnets as the corresponding adapters on the first node.

**Solution:** The Cluster service can detect and use multiple network adapters on the same computer if the following requirements are met:

- The network adapters are bound by TCP/IP.
- You are using hardware that is compatible with Windows Server 2003, Datacenter Edition. To confirm this, click the appropriate link in Support resources.
- Each network adapter is on a different subnet.

Do not use the same IP address for two network adapters, even though they are connected to two different networks.

After the network adapters are configured correctly, use the ping command to check each network adapter for connectivity with the loopback address (127.0.0.1), the adapter's own IP address, and the IP address of another computer.

Make sure that each adapter works correctly and can communicate on each network, and then restart the New Server Cluster Wizard or the Add Nodes Wizard.

**See also:** Configure TCP/IP for static addressing

**Possible Cause:** The cluster does not have network connectivity.

**Solution:** Use the ping command to test connectivity. Ping the node's fully qualified domain name.

If you are unable to ping the node's domain controller by name, consult the Hardware Troubleshooter. You can also learn more about the problem by pinging your domain controller by IP address, and by pinging other network servers by name and IP address, including the node's DNS or WINS server.

**See also:** Manage Networks

**During cluster creation, you are experiencing problems creating the first node of the cluster.**

**Possible Cause:** You are not using a valid IP address for the network interface.

**Solution:** The addresses for each linked pair of network adapters must be on the same subnet.

To check for the IP address, Click **Start**, then **Run** and type **ipconfig /all**, and then press ENTER. This displays your IP configuration, including IP address, subnet mask, default gateway, DNS server address, and DHCP availability. Confirm that these settings are correct. If you are using DHCP, change to static IP addresses.

If no IP address or other IP configuration is displayed, and you see **Media State = Cable Disconnected**, you may not have end-to-end network connectivity from the network adapter to the connecting network device. Make sure your physical connections are secure and valid. If you see an Automatic Private IP Addressing (APIPA) address, change it to a static IP.

Do not use an IP address that is on the same IP range as the public network.

**See also:** TCP/IP

**Possible Cause:** The cluster does not have network connectivity.

**Solution:** Use the ping command to test connectivity. Ping the node's fully qualified domain name.

If you are unable to ping the node's domain controller by name, open Network Connections and check the status of your network adapter. You can also attempt to map a drive to the other node in your cluster to verify communication and authentication. If these solutions do not work, consult the Hardware Troubleshooter. You can also learn more about the problem by pinging your domain controller by IP address, and by pinging other network servers by name and IP address, including the node's DNS or WINS server.

**See also:** Manage Networks

**Possible Cause:** You are not using a unique cluster name.

**Solution:** The cluster name must be different from the domain name, from all computer names on the network, from other cluster names on the network, and from all virtual server names on the network.

Use the **net view** command to determine whether your cluster name is unique.

To use the **net view** command, click **Start**, then **Run** and type the following and then press ENTER:

**net view** *clustername*

*Clustername* is the name of the cluster.

If your cluster name is not unique, assign a unique name. When the New Server Cluster Wizard asks you for the cluster name, type an appropriate new cluster name.

**Possible Cause:** The cluster may not be able to detect a network composed of a crossover cable.

A two-node cluster may use a crossover cable to implement a private network. When the cluster is created on the first node, the second node should be turned on, but stopped in the BIOS or at the startup menu. In this state, the Media Sense feature of Windows may not recognize the network adapter as connected.

When this happens:

- In Network Connections, the icon for the connection that represents the network adapter is shown with a red "X".
- The taskbar displays an icon that, when you pause the mouse on it, displays **Network Cable unplugged**.

**Solution:** If you continue creating the cluster, the crossover cable will be detected when you start the second node. The network will be established in the default mode, which is mixed. You can then change the network mode to private using Cluster Administrator.

**See also:** Manage Networks

**Possible Cause:** The Cluster networks are not on the correct subnet. All network adapters designated for cluster use on a single node need to be on different subnets. Otherwise, only the first network adapter on a subnet is detected by either the New Server Cluster Wizard or the Add Nodes Wizard, and is used for either cluster or client communication (depending on which role you chose).

However, in order for subsequent nodes to pick up the network roles defined on the first node in the cluster, network adapters in each node must be on the same subnets as the corresponding adapters on the first node.

**Solution:** The Cluster service can detect and use multiple network adapters on the same computer if the following requirements are met:

- The network adapters are bound by TCP/IP.
- You are using hardware that is compatible with Windows Server 2003, Datacenter Edition. To confirm this, click the appropriate link in Support resources.
- Each network adapter is on a different subnet.

Do not use the same IP address for two network adapters, even though they are connected to two different networks.

After the network adapters are configured correctly, use the ping command to check each network adapter for connectivity with the loopback address (127.0.0.1), the adapter's own IP address, and the IP address of another computer.

Make sure that each adapter works correctly and can communicate on each network, and then restart the New Server Cluster Wizard or the Add Nodes Wizard.

**See also:** Configure TCP/IP for static addressing

**Possible Cause:** Your network hardware is not installed correctly.

**Solution:** Verify that your network cables are connected correctly. Make sure that no cables are crossed by mistake (private-to-public, or public-to-private), and that each network adapter is connected as you planned.

Set all multispeed network adapters on a private or mixed network to the same speed. Do not use the autodetect setting.

If fault-tolerant network adapters are used, create multiple private networks instead of a single fault-tolerant network.

Make sure all network adapters are correctly seated and the correct drivers are installed in accordance with the manufacturer's instructions. Confirm that your hardware is compatible with Windows Server 2003, Datacenter Edition, by clicking the appropriate link in Support resources.

**Possible Cause:** The WINS client is not configured correctly on the node.

**Solution:** If you use WINS, all cluster resource names that have associated IP addresses are registered with WINS. Use the following instructions to check your WINS configuration.

First, check to see whether the client interconnect has at least one WINS server specified in the IP configuration by clicking **Start**, then **Run** and typing **ipconfig /all**. Confirm that the IP address returned for primary and (if listed) secondary WINS servers are the same IP addresses being used successfully on other computers on the same subnet that are looking for the same resources.

Next determine whether the name is resolved by clicking **Start**, then **Run** and typing **ping *server***, where ***server*** is the network name of the server, and then press ENTER.

If you are using static IP addressing for your cluster nodes, you need to configure WINS on the IP addresses made available on the public network.

If clients have problems connecting to virtual servers, disable WINS on the private cluster network.

**Important**

- Do not disable this for the network adapter used for client communication.

After you configure WINS, if you have already created the cluster on at least one node, ping the cluster name to test for connectivity.

**See also:** Configure TCP/IP to use WINS

**Possible Cause:** TCP/IP is not configured correctly.

**Solution:** If any of the following ping tests fail, there is a TCP/IP configuration problem:

- Pinging the loopback address
- Pinging the local Internet protocol (IP) address (public and private interconnects)
- Pinging the IP address of another node (public and private interconnects)

**Note**

- To test name resolution through the public interconnect, disconnect the private interconnect cable.

If pinging another node by name fails, but you can successfully ping that node's IP address, there is a name-resolution problem. Consult your network administrator for assistance.

TCP/IP errors caused by an incorrect IP address or subnet mask are usually recorded in the system log. Use Event Viewer to check the system log for these errors.

Even if there are no TCP/IP errors in the system log, still check your TCP/IP settings to make sure that they are correct by clicking **Start**, then **Run** and typing **ipconfig /all**. Verify that the listed IP address and subnet mask are the correct values for your computer. If you do not know the correct IP address and subnet mask values for your computer, ask your network administrator.

**See also:** TCP/IP

**Possible Cause:** Your system's network connections are not running. If the following error message appears in the New Server Cluster Wizard tree, your network connections may not be running:

"Warning: This system is not attached to all defined cluster networks."

**Solution:** Make sure that your network connections are enabled. Otherwise, the Cluster service cannot detect and use them.

To check your network connections:

1. Open **Network Connections** in Control Panel.
2. Double-click the connection that you want to check.
3. On the **General** tab, check the status of the connection. If the connection is not started, start it.

You must click **Reanalyze** in the New Server Cluster Wizard for the changes to be detected.

**See also:** TCP/IP

**Possible Cause:** You are not using a valid IP address.

**Solution:** Obtain a static IP address for the cluster itself. You cannot use DHCP to assign this address.

Each network adapter on each cluster node also needs an IP address. Assign a static IP address to each adapter. The addresses for each linked pair of network adapters (linked node-to-node) must be on the same subnet.

If your cluster nodes use DHCP to obtain noncluster IP addresses, at the command prompt, use the **ipconfig** command to verify that you have a valid primary IP address for the networks in question.

Use the **ping** command to confirm that you have IP connectivity to the cluster address and that the IP address is assigned to the correct network. If the ping test fails, you do not have basic network connectivity.

If you made any changes, restart the New Server Cluster Wizard. You must restart this wizard for the changes to be detected.

**See also:** Manage Networks

**Possible Cause:** You do not have the correct domain account and appropriate user rights.

**Solution:** Ensure that you are adhering to the following requirements:

- To create a cluster, you must log on with an account that is a member of the local Administrators group on each node. It

is recommended that you use a domain account. Each node must also be a member of that domain, or must be a member of a domain that is trusted by the domain that contains the account you are logging on to.

- You can administer the Cluster service by using either a domain or a local account; however, other nodes can only be joined to a cluster by using a domain account.

If you make any account changes or changes to user rights, you must click **Reanalyze** in the New Server Cluster Wizard for the changes to be detected. You might also have to log off and back on again for the security changes to take effect.

**See also:** Security

**During cluster creation, you are experiencing problems creating subsequent nodes after creating the first node of the Cluster.**
**Possible Cause:** You are not using a valid IP address for the network interface.

**Solution:** The addresses for each linked pair of network adapters must be on the same subnet.

To check for the IP address, click **Start**, then **Run** and type **ipconfig /all**, and then press ENTER. This displays your IP configuration, including IP address, subnet mask, default gateway, DNS server address, and DHCP availability. Confirm that these settings are correct. If you are using DHCP, change to a static IP address.

If no IP address or other IP configuration is displayed, and you see **Media State = Cable Disconnected**, you might not have end-to-end network connectivity from the network adapter to the connecting network device. Make sure your physical connections are secure and valid.

**See also:** TCP/IP

**Possible Cause:** The cluster does not have network connectivity.

**Solution:** Use the **ping** command to test connectivity. Ping the node's fully qualified domain name.

If you are unable to ping the node's domain controller by name, consult the Hardware Troubleshooter. You can also learn more about the problem by pinging your domain controller by IP address, and by pinging other network servers by name and IP address, including the node's DNS or WINS server.

**See also:** Manage Networks

**Possible Cause:** The cluster might not be able to detect a network composed of a crossover cable.

A two-node cluster might use a crossover cable to implement a private network. When the cluster is created on the first node the second node should be turned on but stopped in the BIOS or at the startup menu. In this state, the Media Sense feature of Windows might not recognize the network adapter as connected.

When this happens:

- In Network Connections, the icon for the connection that represents the network adapter is shown with a red "X".

- The taskbar displays an icon that, when you pause the mouse on it, displays **Network Cable unplugged**.

**Solution:** If you continue creating the cluster, the crossover cable will be detected when you start the second node. The network will be established in the default mode, which is mixed. You can then change the network mode to private using Cluster Administrator.

**See also:** Manage Networks

**Possible Cause:** The Cluster networks are not on the correct subnet.

All network adapters designated for cluster use on a single node need to be on different subnets. Otherwise, only the first network adapter on a subnet is detected by either the New Server Cluster Wizard or the Add Nodes Wizard and is used for either cluster or client communication (depending on which role you chose).

However, in order for subsequent nodes to pick up the network roles defined on the first node in the cluster, network adapters in each node must be on the same subnets as the corresponding adapters on the first node.

**Solution:** The Cluster service can detect and use multiple network adapters on the same computer if the following requirements are met:

- The network adapters are bound by TCP/IP.
- You are using hardware that is compatible with Windows Server 2003, Datacenter Edition. To confirm this, click the appropriate link in Support resources.
- Each network adapter is on a different subnet. Make sure to verify your subnet entry for each node.

Do not use the same IP address for two network adapters, even though they are connected to two different networks.

After the network adapters are configured correctly, use the **ping** command to check each network adapter for connectivity with the loopback address (127.0.0.1), the adapter's own IP address, and the IP address of another computer.

Make sure that each adapter works correctly and can communicate on each network, and then restart either the New Server Cluster Wizard or the Add Nodes Wizard

**See also:** Configure TCP/IP for static addressing

**Possible Cause:** Your network hardware is not installed correctly.

**Solution:** Verify that your network cables are connected correctly. Make sure that no cables are crossed by mistake (private-to-public, or public-to-private), and that each network adapter is connected as you planned.

Check Device Manager for resource conflicts.

Set all multispeed network adapters on a private or mixed network to the same speed. Do not use the autodetect setting.

If fault-tolerant network adapters are used, create multiple private networks instead of a single fault-tolerant network.

Make sure all network adapters are correctly seated and the correct drivers are installed in accordance with the manufacturer's instructions. Confirm that your hardware is compatible with Windows Server 2003, Datacenter Edition, by clicking the appropriate link in Support resources.

**See also:** Device Manager

**Possible Cause:** The WINS client is not configured correctly on the node.

**Solution:** If you use WINS, all cluster resource names that have associated IP addresses are registered with WINS. Use the following instructions to check your WINS configuration.

First, check to see whether the client interconnect has at least one WINS server specified in the IP configuration by clicking **Start**, then **Run** and typing **ipconfig /all**. Confirm that the IP address returned for primary and (if listed) secondary WINS servers are the same IP addresses being used successfully on other computers on the same subnet that are looking for the same resources.

Next determine whether the name is resolved by clicking **Start**, then **Run** and typing **ping *server***, where ***server*** is the network name of the server, and then press ENTER.

If you are using static IP addressing for your cluster nodes, you need to configure WINS on the IP addresses made available on the public network.

If clients have problems connecting to virtual servers, disable NetBios over TCP/IP on the private cluster network.

**Important**

- Do not disable this for the network adapter used for client communication.

After you configure WINS, if you have already created the cluster on at least one node, ping the cluster name to test for connectivity.

**See also:** Configure TCP/IP to use WINS

**Possible Cause:** TCP/IP is not configured correctly.

**Solution:** If any of the following ping tests fail, there is a TCP/IP configuration problem:

- Pinging the loopback address

- Pinging the local Internet protocol (IP) address (public and private interconnects)
- Pinging the IP address of another node (public and private interconnects)

**Note**

- To test name resolution through the public interconnect, disconnect the private interconnect cable.

If pinging another node by name fails, but you can successfully ping that node's IP address, there is a name-resolution problem. Consult your network administrator for assistance.

TCP/IP errors caused by an incorrect IP address or subnet mask are usually recorded in the system log. Use Event Viewer to check the system log for these errors.

Even if there are no TCP/IP errors in the system log, still check your TCP/IP settings to make sure that they are correct by clicking **Start**, then **Run** and typing **ipconfig /all**. Verify that the listed IP address and subnet mask are the correct values for your computer. If you do not know the correct IP address and subnet mask values for your computer, ask your network administrator.

Verify the network adapter's bindings and priority. The public network adapter or adapters should be first, then any private network adapters.

**See also:** TCP/IP

**Possible Cause:** Your system's network connections are not running. If the following error message appears in the New Server Cluster Wizard tree, your network connections may not be running:

"Warning: This system is not attached to all defined cluster networks."

**Solution:** Make sure that your network connections are enabled. Otherwise, the Cluster service cannot detect and use them.

To check your network connections:

1. Open Network Connections in Control Panel.
2. Double-click the connection that you want to check.
3. On the **General** tab, check the status of the connection. If the connection is not started, start it.

You must click **Reanalyze** in the New Server Cluster Wizard for the changes to be detected.

**See also:** TCP/IP

**Possible Cause:** You are not using a valid IP address.

**Solution:** Obtain a static IP address for the cluster itself. You cannot use DHCP to assign this address.

Each network adapter on each cluster node also needs an IP address. Assign a static IP address to each adapter. The addresses for each linked pair of network adapters (linked node-to-node) must be on the same subnet.

If your cluster nodes use DHCP to obtain noncluster IP addresses, at the command prompt, use the **ipconfig** command to verify that you have a valid primary IP address for the networks in question.

Use the **ping** command to confirm that you have IP connectivity to the cluster address and that the IP address is assigned to the correct network. If the ping test fails, you do not have basic network connectivity.

If you made any changes, restart the New Server Cluster Wizard. You must restart this wizard for the changes to be detected.

**See also:** Manage Networks

**Possible Cause:** You do not have the correct domain account and appropriate user rights.

**Solution:** Ensure that you are adhering to the following requirements:

- To create a cluster, you must log on with a domain account that is a member of the local Administrators group on each node, and each node must be a member of that domain.

- You can administer the Cluster service by using either a domain or a local account; however, other nodes can only be joined to a cluster by using a domain account.

If you make any account changes or changes to user rights, you must click **Reanalyze** in the New Server Cluster Wizard for the changes to be detected.

**See also:** Security

**Possible Cause:** Your system does not meet the minimum requirements for a cluster.

**Solution:** Verify that all requirements, including hardware, operating system software configuration, network and disk requirements are met.

**See also:** Planning and preparing for cluster installation

**Possible Cause:** Your cluster nodes are in separate domains or are not in a domain.

**Solution:** All cluster nodes must be members or domain controllers of the same domain, and must use the same account for the Cluster service.

To check computer name and domain name for a node right-click **My Computer**, click **Properties**, and then click the **Computer Name** tab. Make sure each cluster node belongs to the same domain. Use the same account and password for service on each node. This account must have the following user rights:

- Back up files and directories

- Increase scheduling priority

- Log on as a Service

If you made any changes, restart the New Server Cluster Wizard. You must restart this wizard for the changes to be detected.

**Possible Cause:** The cluster name is not online.

**Solution:** To make sure that the cluster name is online:

1. On the first node, open Cluster Administrator by clicking **Start**, then **Run**, then typing **CluAdmin**.

2. Open **Groups**, click **Cluster Group** and check the status of **Cluster Name**.

3. If the status is not **Online**, right-click **Cluster Name**, and then click **Bring Online**.

4. Restart the New Server Cluster Wizard. You must restart this wizard for the changes to be detected.

**Note**

- If the Internet protocol (IP) address is not online, the cluster name will not be online. Make sure that the IP address is online before you try to bring the cluster name online.

**Possible Cause:** The Cluster service is not running on the first node.

**Solution:** Confirm that the Cluster service is running on the first node and that all resources within the cluster group are online before you try to install another node.

To confirm that the Cluster service is started:

1. On the first node, open Computer Management.

2. Double-click **Services and Applications**.

3. Click **Services**.

4. In the details pane, find the Cluster service entry and check its Status column. If the service is not started, right-click it, and then click **Start**.

5. Restart the New Server Cluster Wizard. You must restart this wizard for the changes to be detected.

If the Cluster service was not started, review the records in the event log to see why not.

**During cluster creation, the Physical Disk resources are not automatically created for cluster disks.**
**Cause:** Typically when you create a cluster, the Physical Disk resources are automatically created for cluster disks that use drive letters. However, if your cluster disks use mount points, Physical Disk resources will not be automatically created for those disks.

**Solution:** You must manually create disk resources for cluster disks that use mount points.

For more information, see Checklist: Installing a Physical Disk resource.

**After setting up a new cluster, Cluster Administrator cannot connect to the cluster.**
**Cause:** The new cluster was given a name that corresponds to an existing computer object in Active Directory directory service.

**Solution:** Determine if you want the Network Name resource for the cluster to use Kerberos authentication. If you do, perform the following steps:

1. Using the command line, take the Network Name resource for the cluster offline.

2. Using the command line, enable Kerberos authentication for the Network Name resource.

3. Inspect the permissions on the corresponding computer object in Active Directory and ensure that the Cluster service user account has access to that object.

4. Using the command line, bring the Network Name resource of the cluster online.

As an alternative, if you choose not to use Kerberos authentication for the Network Name resource, you must instead delete the corresponding computer object in Active Directory. Before you do so, first ensure that the computer object is not required elsewhere by your installation.

For more information, see article Q307532, "How to Troubleshoot the Cluster Service Account When It Modifies Objects," and article Q302389, "Description of the Properties of the Cluster Network Name Resource in Windows Server 2003," in the Microsoft Knowledge Base. (http://www.microsoft.com/) Also see Cluster resource and Finding directory information.

# Configuring Roles for Your Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring roles for your server**
The Configure Your Server Wizard provides a central location from which you can install or remove the server roles available on a server running Windows Server 2003. After installing a server role you can use Manage Your Server to manage that role.

Manage Your Server provides a central location from which you can manage the server roles installed through the Configure Your Server Wizard. Manage Your Server starts automatically the first time you log on to your server computer with administrative credentials.

- Server roles
- Typical setup for a first server
- File server role: Configuring a file server
- Print server role: Configuring a print server
- Application server role: Configuring an application server
- Mail server role: Configuring a mail server
- Terminal server role: Configuring a terminal server
- Remote access/VPN server role: Configuring a remote access/VPN server
- Domain controller role: Configuring a domain controller
- DNS server role: Configuring a DNS server
- DHCP server role: Configuring a DHCP server
- Streaming media server role: Configuring a streaming media server
- WINS server role: Configuring a WINS server
- To remove a server role

**Note**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

# Server roles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server roles**
The Windows Server 2003 family provides several server roles. To configure a server role, install the server role by using the Configure Your Server Wizard and manage your server roles by using Manage Your Server. After you finish installing a server role, Manage Your Server starts automatically.

To determine which server role is appropriate for you, review the following information about the server roles that are available with the Windows Server 2003 family:

File server role overview

Print server role overview

Application server role overview

Mail server role overview

Terminal server role overview

Remote access/VPN server role overview

Domain controller role overview

DNS server role overview

DHCP server role overview

Streaming media server role overview

WINS server role overview

**File server role overview**
File servers provide and manage access to files. If you plan to use disk space on this computer to store, manage, and share information such as files and network-accessible applications, configure this computer as a file server.

After configuring the file server role, you can do the following:

- Use disk quotas on volumes formatted with the NTFS file system to monitor and limit the amount of disk space available to individual users. You can also specify whether to log an event when a user exceeds the specified disk space limit or when a user exceeds the specified disk space warning level (that is, the point at which a user is nearing his or her quota limit).
- Use Indexing Service to quickly and securely search for information, either locally or on the network.
- Search in files that are in different formats and languages, either through the **Search** command on the **Start** menu or through HTML pages that users view in a browser.

For more information about implementing this server role, see File server role: Configuring a file server.

**Print server role overview**
Print servers provide and manage access to printers. If you plan to manage printers remotely, manage printers by using Windows Management Instrumentation (WMI), or print from a server or client computer to a print server by using a URL, configure this computer as a print server.

After configuring the print server role, you can do the following:

- Use a browser to manage printers. You can pause, resume, or delete a print job, and view the printer and print job's status.

- Use the new standard port monitor, which simplifies installation of most TCP/IP printers on your network.
- Use Windows Management Instrumentation (WMI), which is the management API created by Microsoft that enables you to monitor and control all system components, either locally or remotely. The WMI Print Provider enables you to manage print servers, print devices, and other printing-related objects from the command line. With WMI Print Provider, you can use Visual Basic (VB) scripts to perform administrative printer functions. For more information, see Windows Management Instrumentation Command-line (WMIC) tool.
- Print from Windows XP clients to print servers running Windows Server 2003 by using a Uniform Resource Locator (URL).
- Connect to printers on your network by using Web point-and-print for single-click installation of a shared printer. You can also install drivers from a Web site.

For more information about implementing this server role, see Print server role: Configuring a print server.

**Application server role overview**
An application server is a core technology that provides key infrastructure and services to applications hosted on a system. Typical application servers include the following services:

- Resource pooling (for example, database connection pooling and object pooling)
- Distributed transaction management
- Asynchronous program communication, typically through message queuing
- A just-in-time object activation model
- Automatic XML Web Service interfaces to access business objects
- Failover and application health detection services
- Integrated security

The Windows Server 2003 family includes an application server that contains all of this functionality and other services for development, deployment, and runtime management of XML Web services, Web applications, and distributed applications.

When you configure this server as an application server you will be installing Internet Information Services (IIS) along with other optional technologies and services such as COM+ and ASP.NET. Together, IIS and the Windows Server 2003 family provide integrated, reliable, scalable, secure, and manageable Web server capabilities over an intranet, the Internet, or through an extranet. IIS is a tool for creating a strong communications platform of dynamic network applications.

For more information about implementing this server role, see Application server role: Configuring an application server.

**Mail server role overview**
To provide e-mail services to users, you can use the Post Office Protocol 3 (POP3) and Simple Mail Transfer Protocol (SMTP) components included with the Windows Server 2003 family. The POP3 service implements the standard POP3 protocol for mail retrieval, and you can pair it with the SMTP service to enable mail transfer. If you plan to have clients connect to this POP3 server and download e-mail to local computers by using a POP3 capable mail client, configure this server as a mail server.

After configuring the mail server role, you can do the following:

- Use the POP3 service to store and manage e-mail accounts on the mail server.
- Enable user access to the mail server so that users can retrieve e-mail from their local computer by using an e-mail client that supports the POP3 protocol (for example, Microsoft Outlook).

For more information about implementing this server role, see Mail server role: Configuring a mail server.

**Terminal server role overview**
With Terminal Server, you can provide a single point of installation that gives multiple users access to any computer that is running a Windows Server 2003 operating system. Users can run programs, save files, and use network resources all from a remote location, as if these resources were installed on their own computer.

After configuring the terminal server role, you can do the following:

- Confirm Internet Explorer Enhanced Security Configuration settings.
- Centralize the deployment of programs on one computer.
- Ensure that all clients use the same version of a program.

### Important

- In addition to configuring a terminal server, you must install Terminal Server Licensing and configure a Terminal Server License Server. Otherwise, your terminal server will stop accepting connections from unlicensed clients when the evaluation period ends 120 days after the first client logon. For more information about Terminal Server Licensing, see Terminal Server Licensing.

For more information about implementing this server role, see Terminal server role: Configuring a terminal server.

**Remote access/VPN server role overview**
Routing and Remote Access provides a full-featured software router and both dial-up and virtual private network (VPN) connectivity for remote computers. It offers routing services for local area network (LAN) and wide area network (WAN) environments. It also enables remote or mobile workers to access corporate networks as if they were directly connected, either through dial-up connection services or over the Internet by using VPN connections. If you plan to connect remote workers to business networks, configure this server as a remote access/VPN server.

Remote access connections enable all of the services that are typically available to a LAN-connected user, including file and print sharing, Web server access, and messaging.

After configuring the remote access/VPN server role, you can do the following:

- Control how and when remote users access your network.
- Provide network address translation (NAT) services for the computers on your network.
- Create custom networking solutions using application programming interfaces (APIs).

For more information about implementing this server role, see Remote access/VPN server role: Configuring a remote access/VPN server.

**Domain controller role overview**
Domain controllers store directory data and manage communication between users and domains, including user logon processes, authentication, and directory searches. If you plan to provide the Active Directory directory service to manage users and computers, configure this server as a domain controller.

### Notes

- You cannot add the domain controller role to a certification authority (CA). If your computer is already a CA, the domain controller role is not available in the Configure Your Server Wizard.
- Computers running Windows Server 2003, Web Edition cannot function as domain controllers. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

After configuring the domain controller role, you can do the following:

- Store directory data and make this data available to network users and administrators. Active Directory stores information about user accounts (for example, names, passwords, phone numbers, and so on), and enables other

authorized users on the same network to access this information.

- Add additional domain controllers to an existing domain to improve the availability and reliability of network services.

- Improve network performance between sites by placing a domain controller in each site. With a domain controller in each site, you can handle client logon processes within the site without using the slower network connection between sites.

For more information about implementing this server role, see Domain controller role: Configuring a domain controller.

**DNS server role overview**
The Domain Name System (DNS) is the TCP/IP name resolution service that is used on the Internet. The DNS service enables client computers on your network to register and resolve user-friendly DNS names. If you plan to make resources in your network available on the Internet, configure this server as a DNS server.

**Important**

- If you plan to include computers on the Internet on your network, use a unique DNS domain name. For more information about DNS namespace planning, see Namespace planning for DNS.

After configuring the DNS server role, you can do the following:

- Host records of a distributed DNS database and use these records to answer DNS queries sent by DNS client computers, such as queries for the names of Web sites or computers in your network or on the Internet.

- Name and locate network resources using user–friendly names.

- Control name resolution for each network segment and replicate changes to either the entire network or globally on the Internet.

- Reduce DNS administration by dynamically updating DNS information.

For more information about implementing this server role, see DNS server role: Configuring a DNS server.

**DHCP server role overview**
Dynamic Host Configuration Protocol (DHCP) is an IP standard designed to reduce the complexity of administering address configurations by using a server computer to centrally manage IP addresses and other related configuration details used on your network. If you plan to perform multicast address allocation, and obtain client IP address and related configuration parameters dynamically, configure this server as a DHCP server.

After configuring the DHCP server role, you can do the following:

- Centrally manage IP addresses and related information.

- Use DHCP to prevent address conflicts by preventing a previously assigned IP address from being used again to configure a new computer on the network.

- Configure your DHCP server to supply a full range of additional configuration values when assigning address leases. This will greatly decrease the time you spend configuring and reconfiguring computers on your network.

- Use the DHCP lease renewal process to ensure that client configurations that need to be updated often (such as users with mobile or portable computers that change locations frequently) can be updated efficiently and automatically by clients communicating directly with DHCP servers.

For more information about implementing this server role, see DHCP server role: Configuring a DHCP server.

**Streaming media server role overview**
Streaming media servers provide Windows Media Services to your organization. Windows Media Services manages, delivers,

and archives Windows Media content, including streaming audio and video, over an intranet or the Internet. If you plan to provide audio and video content streams to clients and to other Windows Media servers, configure this computer as a streaming media server.

After configuring the streaming media server role, you can do the following:

- Stream Windows Media-based content, including entertainment programming and advertising, over a network in a variety of ways to create a Web-based radio or television station.
- Stream internal business communications, such as meetings, speeches, and training, over a corporate intranet, and stream external business communications to customers, vendors, and business partners over the Internet.
- Transmit the information only to users who have been authorized to access the content, by distributing unique user licenses and by limiting access to your server resources.

For more information about implementing this server role, see Streaming media server role: Configuring a streaming media server.

**WINS server role overview**

Windows Internet Name Service (WINS) servers map IP addresses to NetBIOS computer names and NetBIOS computer names back to IP addresses. With WINS servers in your organization, you can search for resources by computer name instead of IP address, which can be easier to remember. If you plan to map NetBIOS names to IP addresses or centrally manage the name-to-address database, configure this server as a WINS server.

After configuring the WINS server role, you can do the following:

- Reduce NetBIOS-based broadcast traffic on subnets by permitting clients to query WINS servers to directly locate remote systems.
- Support earlier Windows and NetBIOS-based clients on your network by permitting these types of clients to browse lists for remote Windows domains without requiring a local domain controller to be present on each subnet.
- Support DNS-based clients by enabling those clients to locate NetBIOS resources when WINS lookup integration is implemented. For more information, see WINS lookup integration.

For more information about implementing this server role, see WINS server role: Configuring a WINS server.

# Typical setup for a first server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Typical setup for a first server**

The following information describes how to install and configure the first server on a network by using the **Typical setup for a first server** option in the Configure Your Server Wizard.

**Important**

- Before you configure the first server on a network, see Checklist: Configuring the typical setup for a first server.

The typical setup for a first server will not run if any of the following conditions are met:

- The computer is running Windows Server 2003, Datacenter Edition.
- The computer is running Windows Server 2003, Web Edition.
- The computer is joined to a domain.
- The computer is already configured as a domain controller.
- The computer is not a domain controller, but the Active Directory Installation Wizard has already been started.
- The computer is a certification authority (CA).
- The computer is already configured as a DNS server.
- The computer is already configured as a DHCP server.
- There are zero IP-enabled network adaptors.
- There is only one IP-enabled network adaptor and the DHCP lease test succeeds.
- The computer is already running Routing and Remote Access.
- The computer does not have at least one NTFS partition.
- The current session is a remote session.

**Typical setup configuration process**

The typical setup configuration process implements the following steps:

- Installs Active Directory and promotes the computer to a domain controller.

  When you promote a server to a domain controller, a domain is automatically created on the network. After you promote a server to a domain controller, you can then promote other servers to domain controllers. For more information, see Domain controllers.

  When you configure your server using the typical setup for a first server, the local administrator's password is automatically set as the Restore Mode Administrator password.

- Sets up an application naming context in Active Directory on this domain controller for use by Telephony API (TAPI) client applications. For more information, see Application directory partitions.

- Installs Domain Name System (DNS) and creates a full domain name for your network.

  DNS is a networking protocol for naming computers and network services that is organized into a hierarchy of domains. TCP/IP networks, such as the Internet, use DNS naming to locate computers and services through user-friendly names. When a user enters a DNS name in an application, DNS services can resolve the name to other information associated with the name, such as an IP address.

  For example, most users prefer a friendly name such as server1.example.microsoft.com to locate a computer such as a mail server or Web server on a network. However, computers use numeric addresses to communicate over a network. To make it easier to use network resources, DNS provides a way to map the user-friendly name for a computer or service to its numeric address. For more information, see DNS domain names.

- Changes the default NetBIOS name. You can change the default NetBIOS name on this page. For example, if your Active Directory DNS domain name is yoursmallbusiness.local, the default NetBIOS domain name is YOURSMALLBUSINE. You may want to change this to CORP or to your own abbreviation for your business name for ease of use. For more information about NetBIOS, see NetBIOS name resolution.

- Assigns a preferred DNS server with the same IP address that you specified for this server. For more information, see How DNS query works.

- Assigns a DNS forwarder. For more information, see Understanding forwarders.

- Installs the DHCP Server service.

  DHCP uses a client/server model. The network administrator establishes one or more DHCP servers that maintain TCP/IP configuration information and provide it to clients. With a DHCP server installed and configured on your network, DHCP-enabled clients can obtain their IP address and related configuration parameters dynamically each time they start and join your network. DHCP servers provide this configuration in the form of an address lease offer to requesting clients.

- Assigns a static IP address, if one is not already assigned to the private network connection. For more information, see Name resolution for TCP/IP.

- Assigns a subnet mask (that is, if none has been configured on this server). By default, the Configure Your Server Wizard assigns a subnet mask of 255.255.255.0. For more information, see Subnet masks.

- Installs the Routing and Remote Access service (that is, if more than one network connection is detected).

**Next steps: Typical setup for a first server**
After you complete the Configure Your Server Wizard, the computer is ready for use as the server on your network. It has been configured as a domain controller, a DNS server, a DHCP server, and possibly as a remote access/VPN server. For more information about each of these roles, see

- Domain controller role: Configuring a domain controller

- DNS server role: Configuring a DNS server

- DHCP server role: Configuring a DHCP server

- Remote access/VPN server role: Configuring a remote access/VPN server

The following table lists some of the additional tasks that you can perform on your server. To manage the new services on this server, do the following:

- To open Administrative Tools, click **Start**, click **Control Panel**, and then double-click **Administrative Tools**.

**Note**

- If you demote this server so that it is no longer a domain controller, you must remove the TAPI application partition by using the Tapicfg.exe utility. For more information, see Remove a TAPI application directory partition.

| Task | Reference |
| --- | --- |
| Make additional configuration changes to the DNS server. | Checklist: Installing a DNS server; Configure a new DNS server |
| Install additional domain controllers, if necessary. | Create an additional domain controller |
| Configure additional options for the DHCP server. | Checklist: Installing a DHCP server |
| (Optional) Monitor and measure various aspects of the DNS server to prevent and troubleshoot performance degradation. | Monitoring DNS server performance |
| (Optional) Monitor and measure various aspects of the DHCP server to prevent and troubleshoot performance degradation. | Monitoring DHCP server performance; DHCP performance monitoring reference |
| (Optional) Install and configure additional domain controllers. | Checklist: Creating an additional domain controller in an existing domain |
| (Optional) Select the appropriate security level for domain administration. | Using the Active Directory Installation Wizard |
| (Optional) Configure ports to allow remote administration. | Windows Firewall Settings |

# File server role: Configuring a file server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File server role: Configuring a file server**
A file server provides a central location on your network where you can store and share files with users across your network. When users require an important file such as a project plan, they can access the file on the file server instead of having to pass the file between their separate computers. If your network users will need access to the same files and network-accessible applications, configure this computer as a file server.

This topic explains how to use both the Manage Your Server and the Configure Your Server Wizard to install and configure a file server. When you have finished setting up a basic file server, you can complete additional tasks by using Manage Your Server. After you complete the Configure Your Server Wizard, you will have a fully functioning file server.

This topic covers:

Before you begin

Configuring your file server

Next steps: Completing additional tasks

**Before you begin**
Before you configure your computer as a file server, verify whether or not:

- The operating system is configured correctly. In the Windows Server 2003 family, file services depend on the appropriate configuration of the operating system and its services. If you have a new installation of a Windows Server 2003 operating system, you can use the default service settings. No further action is necessary. If you upgraded to a Windows Server 2003 operating system or you want to confirm that your services are configured correctly for best performance and security, verify your service settings by comparing them to the table in Default settings for services.

- The computer is joined to an Active Directory domain as a member server. If you want to authenticate clients or publish a shared folder to Active Directory, the file server must be joined to a domain. If you do not need to perform either of these tasks, the file server does not need to be joined to a domain.

- All available disk space is allocated. You can use Disk Management or DiskPart.exe to create a new partition out of unallocated space. For more information see Create a partition or logical drive.

- All existing disk volumes use the NTFS file system. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions.

- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.

  If Windows Firewall is enabled, you must select **File and Printer Sharing** on the **Exceptions** tab in Windows Firewall in order for the File Server Role to function properly.

- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

The following table lists the information that you need to know before you add a file server role.

| Before adding a file server role | Comments |
|---|---|
| Determine whether you want to configure disk quotas. | Use disk quotas to track and control disk space usage for NTFS volumes on a per-volume basis. Quotas prevent users from exceeding the designated disk space by logging an event when a user exceeds a specified disk space limit. |

| | |
|---|---|
| Determine whether you want to use Indexing Service. | Indexing Service creates indexes of the contents and properties of documents located on your local hard drive or on shared network drives. These indexes enable users to perform faster, easier searches. Indexing Service can slow down the server, so use it only if users frequently search the contents of files on this server. |
| Identify the folders that you want to share on the computer, and specify a folder name and description. | Users view the shared resources on this file server based on file name. It is recommended that you create share names that are easy to remember and indicative of the folder contents. For example, suppose users are provided with 2 gigabytes (GB) each for storing their private information on the file server. You might name the top-level folder on your file server Personal Folders, and then name each of the subfolders according to the user's domain name. |
| Determine what type of permissions you want to set on the folders. | Assign the most restrictive permissions that still allow users to perform required tasks. Access control on the NTFS file system provides more security than share permissions alone. |

**Configuring your file server**

To configure a file server, start the Configure Your Server Wizard by doing either of the following:

- From **Manage Your Server**, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **File server**, and then click **Next**.

This section covers:

File Server Disk Quotas

File Server Indexing Service

Summary of Selections

Using the Share a Folder Wizard

Completing the Configure Your Server Wizard

Removing the file server role

**File Server Disk Quotas**

On the **File Server Disk Quotas** page, you can set up disk quotas, which track and control individual user's disk space usage on NTFS volumes on a per-volume basis. The Configure Your Server Wizard automatically applies disk quotas to new users of any NTFS file system, using whatever disk space quota is already applied. You need to change the information on the **File Server Disk Quotas** page only if you want to prevent a server from consuming a certain amount of disk space or if you have a limited amount of disk space. In most cases, you can accept the default system settings.

If you want to log an event when a user exceeds the specified disk space limit or when a user exceeds the specified disk space warning level (that is, the point at which a user is nearing his or her quota limit), you can specify that on this page.

The following table describes the manual configuration options.

| **Settings** | **Comments** |
|---|---|

| | |
|---|---|
| **Set up default disk quotas for new users of this server** | If you want to enable disk quotas, so that you can limit and track disk space usage on this file server, select this check box.<br><br>If you choose to enable disk quotas, you need to set the disk space limit. It is recommended that you set moderately restrictive default limits for all user accounts, and then modify the limits to allow more disk space for users who work with large files. For example, users who work with scanned photographs or artwork might require large amounts of disk space.<br><br>You can also set the warning level so that users are notified when they exceed the specified disk space limit. If you do not want to use a warning level, set this number higher than the disk space limit. |
| **Deny disk space to users exceeding disk space limit** | If you want to limit disk space usage on the file server, configure this setting. If you only want to track disk space usage on a per-user basis, leave this setting blank. |
| **Log an event when the user exceeds any of the following** | If you want to log a system event when the user exceeds the disk space limit or the warning level, configure these settings. You can view system events by using Event Viewer. To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**. |

After you finish, click **Next**.

#### File Server Indexing Service

On the **File Server Indexing Service** page, do one of the following:

- If users will be searching the contents of the files on the server regularly, click **Yes, leave Indexing Service turned on**.
- If you want to conserve CPU and memory resources, click **No, turn Indexing Service off**. Indexing Service can slow server performance.

Indexing Service provides a fast, easy, and secure way for users to search for information locally or on the network. Users can search in files in different formats and languages, either through Search on the Start menu or through HTML pages that they view in a browser.

After you finish, click **Next**.

#### Summary of Selections

On the **Summary of Selections** page, view and confirm the options that you have selected. If you selected File server on the **Server Role** page, the following appears:

- **Install File Server Management**
- **Run the Share a Folder Wizard to add new shared folders or share existing folders**

To apply the selections shown on the **Summary of Selections** page, click **Next**.

#### Using the Share a Folder Wizard

After you click **Next**, the Configure Your Server Wizard starts the Share a Folder Wizard automatically, which you use to configure shared folders. By sharing resources, you make them available for use by other users on your network.

**Important**

- Use caution when sharing folders that contain system files and resources. Verify that the folder or resource that you specify does not contain information that you do not want users to access.

This section describes the following steps in the Share a Folder Wizard:

Folder Path

Name, Description, and Settings

Permissions

## Sharing was Successful

**Folder Path**

On the **Folder Path** page, specify the path to the folder that you want to share. To search for a folder, click **Browse**.

After you finish, click **Next**.

**Name, Description, and Settings**

On the **Name, Description, and Settings** page, specify the following information about the shared folder:

- In **Share name**, type the name you want to use for the shared resource. The share name is required. Choose a name that is short and descriptive, so that it is easy for users to identify.

- In **Description**, type a description of the shared resource. The description is optional. If you are sharing several resources, descriptions can help you organize and identify those resources. The description you type appears in the **Description** column of File Server Management and Shared Folders.

- In **Offline setting**, specify how you want to make the contents of the shared folder available to users when they are not connected to the network. If you want users to have control over which files are available offline, you can accept the default. To change the offline setting, click **Change**. Use the information in the following table to determine what settings you want to use for your offline files.

| Offline setting | Comments |
|---|---|
| **Only the files and programs that users specify will be available offline** | If you want to give users control over which files are available offline, click this option. |
| **All files and programs that users open from the share will be automatically available offline** | If you want to allow all of the files that users open from the shared folder to be automatically available offline, click this option. If you select the **Optimized for performance** check box, all programs will be automatically cached so that users can run them locally. This option is especially useful for file servers that host applications, because it reduces network traffic and improves server scalability. |
| **Files or programs from the share will not be available offline** | If you want to prevent users from storing files offline, click this option. |

After you finish, click **Next**.

**Permissions**

On the **Permissions** page, specify the share permissions for the shared folder. To ensure that only authorized users have access to the information in the folder, you must set permissions on the folder that you created. Share permissions apply only to users who gain access to the resource over the network. They do not apply to users who gain access to the resource from the computer where the resource is stored. Use the following table to determine which share permissions are appropriate.

| Share permission | Comments |
|---|---|
| All users have read-only access | To restrict all access to read-only, click this option. |
| Administrators have full access; other users have read-only access | If you want users to view files and run programs that are located in the shared resource, click this option. Only members of the Administrators group are allowed to change, add, or delete files. Also, only members of the Administrators group are allowed to change the NTFS file permissions on the shared resource. |
| Administrators have full access; other users have read and write access | If you want to restrict access to read and write for all users except members of the Administrators group, click this option. |

| | |
|---|---|
| Use custom share and folder permissions | If you want to grant or deny access to specific users or groups, click this option. You should assign the most restrictive permissions that still allows users to perform necessary functions. |

After you finish, click **Finish**.

**Sharing was Successful**

On the **Sharing was Successful** page, the Share a Folder Wizard displays a status and summary of your selections. If you want to share another folder, click the **When I click Close, run the wizard again to share another folder** check box. When you finish sharing folders, click **Close**.

**Completing the Configure Your Server Wizard**

After you complete the Share a Folder Wizard, the Configure Your Server Wizard displays the **This Server Is Now a File Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at *systemroot*\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

After completing the Configure Your Server Wizard, be sure to visit Windows Update to download any additional updates available. For more information, see Windows Update.

In addition, you should also run the Security Configuration Wizard to secure your file server. For more information see Security Configuration Wizard.

**Removing the file server role**

If you need to reconfigure your server for a different role, you can remove existing server roles. If you remove the file server role, files and folders on this server are no longer shared and network users, programs, or hosts that depend on those shared resources will be unable to connect to them.

To remove the file server role, restart the Configure Your Server Wizard by doing either of the following:

- From **Manage Your Server**, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **File server**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the file server role** check box, and then click **Next**. On the **File Server Role Removed** page, click **Finish**.

**Next steps: Completing additional tasks**

After you complete the Configure Your Server Wizard and create shared resources on the computer, the computer is ready for use as a basic file server that can store, manage, and share information such as files and network-accessible applications. Up to this point, you have completed the following tasks:

- If necessary, established disk space limits by enabling disk quotas.

- If necessary, turned on Indexing Service.

- Created shared folders and set share permissions for each folder.

The Configure Your Server Wizard automatically installs File Server Management, which you use to manage your file server. To open File Server Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **File Server Management**.

The following table lists some of the additional tasks that you can perform on your file server.

| Task | Purpose of task | Reference |
|---|---|---|
| Secure the file server. | To ensure that your file server is secure. | Best practices for security |

| | | |
|---|---|---|
| Implement Encrypting File System (EFS). | To strengthen security of the files and resources on the file server. | Encrypting and decrypting data |
| Configure ports to allow remote administration. | To manage the File Server Role from other computers on the network. | Windows Firewall Settings |
| Set permissions on shared files and folders. | To secure resources on the file server and prevent unauthorized access. Access control on the NTFS file system provides more security than share permissions alone. | Set permissions on a shared resource |
| Make shared resources available offline. | To allow users to store local copies of shared resources, so that they can access these resources when they are not connected to the network. | Configure offline settings for a shared resource |
| Enable shadow copies of shared folders. | To enable shadow copies of shared folders, which provide point-in-time copies of files on network shares. | Enable Shadow Copies for Shared Folders |
| Set up a Distributed File System (DFS). | To make it easier for users to access and manage files that are physically distributed across a network. | Checklist: Creating a distributed file system |
| Ensure that the file server is properly backed up. | To protect data from accidental loss if your system experiences hardware or storage media failure. | Back Up Data |
| Use file compression. | To conserve storage space by compressing files, folders, and programs. | File compression overview |

# Print server role: Configuring a print server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Print server role: Configuring a print server**
If you plan to use this computer to manage and share printers, configure this computer as a print server.

**Note**

- This feature is not included on computers running the Microsoft® Windows Server™ 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

This document explains how to use the Configure Your Server Wizard to quickly meet the most basic requirements of a print server. When you are done setting up a basic print server, you can complete additional configuration tasks, depending on how you want to use this print server.

This topic covers:

- Before you begin
- Configuring your print server
- Next steps: Completing additional tasks

**Before you begin**
Before you configure your server as a print server, verify whether or not:

- The operating system is configured correctly. In the Windows Server 2003 family, print services depend on the appropriate configuration of the operating system and its services. If you have a new installation of a Windows Server 2003 operating system, you can use the default service settings. No further action is necessary. If you upgraded to a Windows Server 2003 operating system or you want to confirm that your services are configured correctly for best performance and security, verify your service settings by comparing them to the table in Default settings for services.
- The computer is joined to an Active Directory domain as a member server. If you want to restrict access to a printer, so that some domain users can print to it and other users cannot, or you want the print server to publish shared printers to Active Directory so that domain users can easily search for those printers, the print server must be joined to a domain. If you do not need to perform either of these tasks, the print server does not need to be joined to a domain.
- All existing disk volumes use the NTFS file system. FAT32 volumes are less secure. For more information about encrypting data stored on NTFS volumes, including spooled print jobs, see Storing Data Securely.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

The following table lists the information that you need to know before you add a print server role.

| Before adding a print server role | Comments |
|---|---|
| Determine the operating system version of the clients that will send jobs to this printer. | You must have this information to select the correct client printer drivers for your client and server computers. After you add this role, the print server can automatically distribute these drivers to the clients. Additionally, the set of client operating systems determines which of these drivers you need to install on the server during the print server role installation. |

| | |
|---|---|
| At the printer, print a configuration or test page that includes manufacturer, model, language, and installed options. | You need this information to choose the correct printer driver. The manufacturer and model are usually enough to uniquely identify the printer and its language. However, some printers support multiple languages, and the configuration printout usually lists them. Also, the configuration printout often lists installed options, such as extra memory, paper trays, envelope feeders, and duplex units. |
| Determine how the print server connects to the printer. | If the printer supports Plug and Play and connects to the print server using infrared technology, a universal serial bus (USB) port, or an IEEE 1394 port, the print server will configure itself automatically. You do not need to follow the remaining steps.<br><br>Otherwise, if the printer is connected to the print server with a cable, note which server port is used. For printers, LPT1 is the most commonly used port.<br><br>If the printer is located away from the print server and uses its own network adapter to receive print jobs, determine the IP address of the network adapter on the printer. |
| (Optional) Determine whether you need a new or updated printer driver. | Most printers are supported by drivers on the installation CD for the Windows Server 2003 operating system. To save time, you can often skip this step because the wizard that you will use to configure your print server provides compatibility information. If the wizard does not list a driver for your printer, you can look for an update from the printer manufacturer or Windows Update. |
| Choose a printer name. | Users running Windows-based client computers choose a printer by using the printer name. The wizard that you will use to configure your print server provides a default name, consisting of the printer manufacturer and model. The printer name is usually less than 31 characters in length. |
| Choose a share name. | A user can connect to a shared printer by typing this name, or by selecting it from a list of share names. The share name is usually less than 8 characters in length for compatibility with MS-DOS and Windows 3.x clients. |
| (Optional) Choose a location description and a comment. | These can help identify the location of the printer and provide additional information. For example, the location could be "Second floor, copy room" and the comment could be "Additional toner cartridges are available in the supply room on floor 1." |

**Configuring your print server**

To set up a print server, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Print server**, and then click **Next**.

This section covers:

- Printers and Printer Drivers

- Summary of Selections

- Using the Add Printer Wizard

- Using the Add Printer Driver Wizard

- Completing the Configure Your Server Wizard

- Removing the print server role

**Printers and Printer Drivers**

On the **Printers and Printer Drivers** page, do one of the following:

- If all of the clients on your network run Windows XP Home Edition, Windows XP Professional, or Windows 2000, click **Windows 2000 and** Windows XP **clients only**.
- If any of the clients run Windows XP 64-bit Edition (Itanium), Windows NT 4.0, Windows Millennium Edition, Windows 98, or Windows 95, click **All Windows clients**.

After you finish, click **Next**.

**Summary of Selections**

On the **Summary of Selections** page, view and confirm the options that you have selected. If you selected **Windows 2000 and** Windows XP **clients only** on the previous page, the following appears:

- **Add printers to this server using the Add Printer Wizard**.

If you selected **All Windows clients** on the previous page, the following appears:

- **Add printers to this server using the Add Printer Wizard.**
- **Add printer drivers to this server using the Add Printer Driver Wizard.**

To apply the selections shown on the **Summary of Selections** page, click **Next**.

**Using the Add Printer Wizard**

After you click **Next**, the Configure Your Server Wizard runs the Add Printer Wizard once for each printer that you want to add. If the wizard finishes and you choose to share at least one printer, your server can be used as a print server. If you cancel the Add Printer Wizard, the Print Spooler service remains installed. If you cancel the Add Printer Wizard and no printers are shared, the server does not add the print server role.

**Important**

- If the printer you want to share supports Plug and Play, do not run the Add Printer Wizard. Plug and Play printers complete the configuration steps in the Add Printer Wizard automatically. If the printer you want to share supports Plug and Play, click **Cancel**.

This section describes the following steps in the Add Printer Wizard:

- Local or Network Printer
    - New Printer Detection
    - Select a Printer Port
    - Specify a Printer
- Install Printer Software
- Use Existing Driver
- Name Your Printer
- Printer Sharing
- Location and Comment
- Print Test Page

- Completing the Add Printer Wizard

**Local or Network Printer**
On the **Local or Network Printer** page of the Add Printer Wizard, choose one of the following options:

- To configure this print sever to send print jobs directly to the printer, click **Local printer attached to this computer**. Typically, print servers send print jobs directly to the printer. A printer with its own network adapter is considered to be a local printer. If you want to send print jobs directly to a printer with its own network adapter, click this option.

- To configure this print server to forward print jobs to a second print server, click **A network printer, or a printer attached to another computer**. For example, you can configure a print server at a branch office to forward print jobs to a print server in the main office. You might do this if regulations require you to create printouts of daily transaction logs and store them at the main office. If you want to do this, click this option.

**Note**

- The **A network printer, or a printer attached to another computer** option is included here because this dialog box is used on all computers running a Windows Server 2003 operating system so that users can connect to a network printer. If you need to print from a computer that is not a print server, click **A network printer, or a printer attached to another computer**.

After you finish, click **Next**.

After you click **Next**, one of the following wizard pages appears:

**New Printer Detection**
If you selected the **Automatically detect and install my Plug and Play printer** check box and the wizard is unable to detect any Plug and Play printers, this page appears. Click **Next**.

To complete the steps on the **Select a Printer Port** page, see Select a Printer Port.

**Select a Printer Port**
If you selected **Local printer attached to this computer**, this page appears.

On the **Select a Printer Port** page, choose one of the following options:

- If a cable connects the printer directly to a port on the print server, under **Use the following port**, click the name of that port. LPT1 is the most commonly used port for this type of printer.

- If the printer has its own network adapter and you want to send print jobs to the printer through the network, click **Create a new port**, and then click the type of port that you want to create. If you do not know what type of port to create, **Standard TCP/IP Port** is recommended.

  If you click **Standard TCP/IP Port**, and then click **Next**, the Add Standard TCP/IP Printer Port Wizard starts. In the Add Standard TCP/IP Printer Port Wizard, click **Next**. On the **Add Port** page, type the name or IP address of the printer. The IP address is usually listed on the printer configuration page. As you type the name or IP address, the wizard completes the **Port Name** field for you. Click **Next**.

  The wizard attempts to connect to the printer. If the wizard is able to connect, the **Completing the Add Standard TCP/IP Printer Port Wizard** page appears, and you can click **Finish**. If the wizard is not able to connect, the **Additional Port Information Required** page appears. If you think that the address or name you entered is not correct, click **Back**, retype the name or address, and then click **Next**.

  If you are sure the address or name is correct, select one of the following device types to identify the printer network adapter:

  - **Standard** is the default. If you click **Standard**, click the manufacturer and model of network adapter from the **Standard** list.

- If the printer network adapter uses nonstandard settings, click **Custom** and then click **Settings**. The **Configure Standard TCP/IP Port Monitor** page appears. Specify the settings that are recommended by the manufacturer of the printer network adapter, and then click **OK**.

After you finish, click **Next**.

**Specify a Printer**
If you selected **A network printer, or a printer attached to another computer**, this page appears.

On the **Specify a Printer** page, choose one of the following options to configure your print server to forward print jobs to another print server:

- If the print server that you want to connect to is available on the network, click **Browse for a printer**, click **Next**, and then, under **Shared printers**, click the server and printer from the list.

- If the print server that you want to connect to is temporarily unavailable on the network, click **Connect to this printer (or to browse for a printer, select this option and click Next)**, and then, in **Name**, type the server and printer names.

- If the print server that you want to connect to belongs to another organization and is available on the Internet, click **Connect to a printer on the Internet or on a home or office network**.

**Important**

- Use the options on this page only if you want your print server to forward print jobs to another print server. If this is not what you want, click **Back**, click **Local printer attached to this computer**, click **Next**, and then follow the steps in Select a Printer Port.

After you finish, click **Next**.

For this configuration path, you can skip some of the following steps in this document. To continue the instructions for this configuration path, see Completing the Add Printer Wizard.

**Install Printer Software**
On the **Install Printer Software** page of the Add Printer Wizard, under **Manufacturer**, click the printer manufacturer, and then, under **Printers**, click the printer model.

**Note**

- Write down the manufacturer and model that you select, because you will need this information later if you use the Add Printer Driver wizard to install printer drivers for other Windows-based clients.

If the manufacturer or model is not listed, try each of the steps outlined in the following table, in sequence, to install the correct printer software.

| Step | Comments |
|---|---|
| Check the configuration printout to confirm the exact spelling of the name of your printer manufacturer and model. | The **Manufacturer** and **Printers** lists show the official product names, which might be different from the names that you normally use. |
| Click **Have Disk**, locate the driver files, and then click **OK**. | If you have printer driver files stored somewhere else, follow these steps. For example, the printer manufacturer might include a CD-ROM containing driver files in the packaging of the printer. |
| Click **Windows Update**. | If you want to look for new or updated drivers that are available from Microsoft as part of Windows Update, click this option. When you click **Windows Update**, the **Manufacturer** and **Printers** lists change to show only the drivers that are available from Windows Update. If the printer is not listed, return to the original list by clicking **Back**, and then clicking **Next**. |

| | |
|---|---|
| Select the manufacturer and model of a compatible printer, and then click **Next**. | To determine which printers are compatible, consult the user guide for your printer. Also, some manufacturers list compatibility information on their Web sites. |

After you finish, click Next.

**Use Existing Driver**

If you add an additional printer that is the same manufacturer and model as one previously installed, the **Use Existing Driver** page appears. Decide whether to keep the same driver or replace it with a new one. If you select **Replace existing driver**, the wizard reinstalls the driver files.

After you finish, click **Next**.

**Name Your Printer**

On the **Name Your Printer** page of the Add Printer Wizard, the default name is the manufacturer and model of the printer. You can change this name so that the printer is easier to use and administer. When using applications, users often select a printer from a list that displays the names of the available printers. To help users decide which printer to select, the application might also list the location or a comment.

Under **Do you want to use this printer as the default printer?**, click **Yes** or **No**. Your response applies only when you print from an application that is running on this print server. Your response does not set this printer as the one that clients use by default.

After you finish, click **Next**.

**Printer Sharing**

**Important**

- You must share at least one printer for this server to act as a print server.

On the **Printer Sharing** page of the Add Printer Wizard, **Share name** is selected by default so that the printer is shared. The default share name is the first 8 letters of the printer manufacturer and model, without spaces. You can change this name so that the printer is easier to use and administer.

For compatibility with clients that run MS-DOS or earlier versions of Windows, type a share name that follows these rules:

- The share name contains only letters, digits, and the period (.).
- The share name contains no more than eight letters and digits, and, optionally, followed by a period, which is followed by no more than 3 letters and digits.

After you finish, click **Next**.

**Location and Comment**

On the **Location and Comment** page of the Add Printer Wizard, in **Location**, type a description of the print server location, and then, in **Comment**, type a comment. This step is optional, but recommended because this information makes it easier to use and administer your print server. Many applications display the comment or the location when the user prints a document, so that the user can choose the most appropriate printer.

After you finish, click **Next**.

**Print Test Page**

On the **Print Test Page** page of the Add Printer Wizard, choose whether to print a test page to confirm that the printer is ready to use.

**Note**

- The test page does not print immediately when you click **Next**. Instead, it prints when you finish the wizard.

After you finish, click **Next**.

**Completing the Add Printer Wizard**

On the **Completing the Add Printer Wizard** page, the **Restart the wizard to add another printer** check box is selected by default. If you leave it selected and click **Finish**, the wizard restarts to add another printer. If you have finished adding all of the

printers that you want to share on this server, clear this check box, and then click **Finish**.

When you click **Finish**, the wizard installs the printer driver files. Then, if you chose to print a test page, the wizard attempts to print that page. If the printer does not receive the test page, you might have selected an incorrect port. However, if the printer receives the test page and prints it incorrectly, you might have selected an incompatible manufacturer and model.

When you started the Configure Your Server Wizard to configure this server as a print server, you selected one of the following options on the **Printers and Printer Drivers** page:

- **Windows 2000 and** Windows XP **clients only**
- **All Windows clients**

If you selected **All Windows clients**, the Add Printer Driver Wizard starts after you click **Finish** in the Add Printer Wizard. You can use the Add Printer Driver Wizard to install client printer drivers onto the print server, which can then automatically distribute them to clients.

**Note**

- The Add Printer Driver Wizard does not communicate with the Add Printer Wizard. Therefore, the Add Printer Driver Wizard does not automatically run once for each printer that you add, and it does not automatically install drivers for the same manufacturer and model of printer. Instead, you must decide how many times to run the Add Printer Driver Wizard, and each time it runs you must decide which manufacturer and model of drivers to install.

#### Using the Add Printer Driver Wizard

If you selected **All Windows clients** on the **Printers and Printer Drivers** page of the Configure Your Server Wizard, the Add Printer Driver Wizard starts after the Add Printer Wizard. If you cancel the Add Printer Driver Wizard, the Print Spooler service remains installed, and any printers you have added remain, but the additional client driver files are not installed on the server, and therefore the server cannot distribute those drivers to clients.

This section describes the following steps in the Add Printer Driver Wizard:

- Printer Driver Selection
- Processor and Operating System Selection
- Completing the Add Printer Driver Wizard

#### Printer Driver Selection

On the **Printer Driver Selection** page of the Add Printer Driver Wizard, select the manufacturer and model of a printer that is shared on this print server, and then click **Next**.

**Important**

- The Add Printer Driver Wizard does not automatically select a manufacturer and model for a printer that you have already added. Instead, it selects the first manufacturer in the list, and the name of the first printer model (in alphabetical order) made by that manufacturer. If possible, select the manufacturer and model of a printer that you have added. If you select a different manufacturer or model, the wizard installs drivers that might not work correctly with your printer.

#### Processor and Operating System Selection

On the **Processor and Operating System Selection** page of the Add Printer Driver Wizard, select the client operating systems and processors.

Drivers for your server operating system are installed automatically when you add a printer. As a result, one of the following is selected automatically and you cannot remove it: **Windows 2000, Windows XP and Windows Server 2003** for x86–based processors, **Windows XP and Windows Server 2003** for Itanium–based processors, or **Windows XP and Windows Server 2003** for x64-based processors.

After you finish, click **Next**.

#### Completing the Add Printer Driver Wizard

On the **Completing the Add Printer Driver Wizard** page, the **Restart the wizard to add another printer driver** check box is selected by default. If you leave it selected and click **Finish**, the wizard restarts to add another printer driver. If you have finished adding all of the printer drivers for all of the printers that you want to share on this server, clear this check box, and then click **Finish**.

**Completing the Configure Your Server Wizard**

After you complete the Add Printer Wizard and, if necessary, the Add Printer Driver Wizard, the Configure Your Server Wizard displays the **This Server is Now a Print Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at systemroot\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

Before you start to use your print server, we recommend the following steps:

- Run Windows Update. For more information, see Windows Update.
- Run the Security Configuration Wizard. For more information, see Security Configuration Wizard Overview.

**Removing the print server role**

If you need to reconfigure your server for a different role, you can remove existing server roles. If you remove the print server role, each client that sent print jobs only to this print server will be unable to print until you reconfigure the client to send print jobs to a different server. Also, each printer managed only by this print server will be unable to receive print jobs until you reconfigure another print server to send print jobs to that printer.

To remove the print server role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Print server**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the print server role** check box, and then click **Next**. On the **Print Server Role Removed** page, click **Finish**.

**Next steps: Completing additional tasks**

After you complete the Configure Your Server Wizard, the server is ready for use as a print server. By following the steps in this document, you have:

- Added one or more printers.
- Shared printers so that clients can send print jobs to the printers.
- If necessary, added client print drivers.

You can use the Add Printer Wizard and Add Printer Driver Wizard to add more printers and client printer drivers. These wizards are available through Manage Your Server.

The following table lists some of the additional tasks that you can perform on your print server.

| Task | Purpose of task | Reference |
|---|---|---|
| Set the configuration to match installed options. | To provide user access to installed printer options, such as an envelope feeder or extra memory, that are available on some printers. If your printer provides additional features, you must update the configuration so that users can use these features. | Set installable options for a printer |

| | | |
|---|---|---|
| Set printing defaults. | To set the default configuration for clients when they connect to the printer. For example, you can set the default layout or paper source. | Set printing defaults |
| Assign printer permissions. | To change the permissions that users have for a printer. | Set or remove permissions for a printer |
| Choose a separator page. | To define a page that appears at the beginning of each printout. | Choose a separator page |
| Configure network clients to use the printer. | To configure clients to connect to the printers that are shared on this print server. | Connect clients to a printer |
| Set advanced printer tasks. | To manage your print server more efficiently and effectively. For example, to schedule alternate printing times, to enable printer location tracking, or to set different priority for different groups. | Use Advanced Options |
| Publish a printer in Active Directory. | To help domain users find printers shared by this print server quickly. For this task, the print server must be a member server. | Publish a printer in Active Directory |
| Configure ports to allow remote administration. | To manage the print server from other computers on the network. | Windows Firewall Settings |

# Application server role: Configuring an application server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application server role: Configuring an application server**
An application server is a core technology that provides key infrastructure and services to applications hosted on a system. Typical application servers include the following services:

- Resource pooling (for example, database connection pooling and object pooling)
- Distributed transaction management
- Asynchronous program communication, typically through message queuing
- A just-in-time object activation model
- Automatic XML Web Service interfaces to access business objects
- Failover and application health detection services
- Integrated security

The Windows Server 2003 family includes all this functionality, in addition to services for development, deployment, and runtime management of XML Web services, Web applications, and distributed applications.

This topic explains the basic steps that you must follow to configure an application server. This process involves using the Configure Your Server Wizard to configure the server as an application server. When you have finished setting up a basic application server, you can complete additional tasks by using Manage Your Server.

This topic covers:

- Before you begin
- Configuring your application server
- Next steps: Completing additional tasks

**Before you begin**
Before you configure your computer as an application server, verify that:

- All existing disk volumes use the NTFS file system. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions. To find out the file system type, in **My Computer** right-click the disk volume, and then click **Properties**. For information about how to convert a FAT partition to an NTFS partition, see "To convert a volume to NTFS from the command prompt" in Windows Help and Support Center.
- Your computer has network connectivity and a static or dynamic IP address.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

The following table lists the information that you need to know before you add an application server role.

| Before adding an application server role | Comments |
| --- | --- |
| Understand the following technologies that are installed automatically while configuring your application server:<br>• Internet Information Services (IIS)<br>• Application Server console<br>• COM+<br>• Distributed Transaction Coordinator (DTC) | • IIS 6.0 is a full-featured Web server that provides the infrastructure for .NET and existing Web applications and Web services.<br>• COM + is an extension to the Component Object Model (COM). COM+ builds on COM's integrated services and features, making it easier for developers to create and use software components in any language, using any tool.<br>• The Application Server console provides a central location from which you can administer your Web applications. To open the Application Server console, in Manage Your Server, click **Manage this application server**.<br>• Distributed Transaction Coordinator (DTC) coordinates COM+ transactions. |
| Determine whether you would like to install FrontPage Server Extensions. | FrontPage Server Extensions enable users on a client computer to publish and administer Web sites on a server remotely over a network. |
| Determine whether you would like to run ASP.NET applications on your server. | ASP.NET is a unified Web development platform that provides the services necessary for developers to build enterprise-class Web applications. You can enable ASP.NET for developing Web applications. |

**Configuring your application server**

To configure an application server, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Application server (IIS, ASP.NET)**, and then click **Next**.

This section covers:

Application Server Options

Summary of Selections

Completing the Configure Your Server Wizard

Removing the application server role

**Application Server Options**

On the **Application Server Options** page, you can choose to install the following optional components with your application server:

- **FrontPage Server Extensions** - FrontPage Server Extensions enable multiple users to administer and publish a Web site from a client computer, remotely. Select this option if you want to enable multiple-users to simultaneously create Web sites, or enable users to create Web applications from their client computers, remotely, over the Internet.

- **Enable ASP.NET** - ASP.NET is a unified Web application platform that provides the services necessary to build and deploy enterprise-class Web applications. ASP.NET offers a new programming model and infrastructure for more secure, scalable, and stable applications that can target any browser or device. If your Web site includes applications that have been developed by using ASP.NET, select this option. If you are not sure that you need to enable ASP.NET, you can enable it later by using IIS Manager. This feature is not available on the Itanium-based versions of the Windows operating systems. By enabling ASP.NET, you can use your application server to host ASP.NET applications. Some of the features of ASP.NET include the following:

- ASP.NET can run side by side with Active Server Pages (ASP) code on Internet Information Services (IIS). If you are already running ASP code you do not need to upgrade your ASP pages, and you can add ASP.NET pages to your applications.

- ASP.NET has enhanced performance.

- ASP.NET supports many languages including Visual Basic .NET, C#, and JScript .NET.

After you finish, click **Next**.

**Summary of Selections**
On the **Summary of Selections** page, you can view and confirm the options that you have selected. If you selected **Application server (IIS, ASP.NET)** on the **Server Role** page, the following appears:

- **Install Internet Information Services (IIS)**

- **Enable COM+ for remote transactions**

- **Enable Microsoft Distributed Transaction Coordinator (DTC) for remote access**

If you selected FrontPage Server Extensions or ASP.NET, the following items also appear:

- **Install FrontPage Server Extensions**

- **Enable ASP.NET**

To apply the selections shown on the **Summary of Selections** page, click **Next**. When you click **Next**, the **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page.

**Completing the Configure Your Server Wizard**
After the components are configured, the Configure Your Server Wizard displays the **This Server is Now an Application Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at *systemroot*\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

If the installation was not successful, the **Cannot Complete** page appears and IIS could not be installed. To troubleshoot the installation, click **Configure Your Server log**.

Run Windows Update. For more information, see Windows Update.

Run the Security Configuration Wizard. For more information, see Security Configuration Wizard Overview.

**Removing the application server role**
If you need to reconfigure your server for a different role, you can remove existing server roles. By removing the application server role, you will uninstall all application server components, such as the IIS. Your server will no longer support serving Web pages, Web applications, or distributed applications.

To remove the application server role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Application server (IIS, ASP.NET)**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the application server role** check box, and then click

**Next**. When you click **Next**, the **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page. On the **Application Server Role Removed** page, click **Finish**.

**Next steps: Completing additional tasks**

After you complete the Configure Your Server Wizard and enable the features that you need to run your applications, the computer is ready for use as a basic application server. Up to this point, you have completed the following tasks:

- Installed Internet Information Services (IIS), ASP.NET, and COM+.
- If necessary, enabled FrontPage Server Extensions.
- If necessary, enabled ASP.NET.

The following table lists some of the additional tasks that you might want to perform on your application server.

| Task | Purpose of task | Reference |
|---|---|---|
| Secure your application server. | To ensure the security of this server, it is recommended that you implement security precautions, such as firewalls and Internet Protocol security (IPSec), before placing it in a production environment. An application server may be targeted by attackers because of its exposure to the Internet and other networks. You can secure your applications by using authentication protocols, access control, Secure Sockets Layer (SSL), and encryption. | Internet Protocol Security (IPSec); Basic Firewall; Security in Microsoft Internet Information Services |
| Secure your files with NTFS. | To secure your Web site, applications, databases, and files use NTFS permissions. This is essential for a secure site. | Set, view, change, or remove permissions on files and folders |
| Configure Web Interface for Remote Administration. | To manage your application server using a Web browser on a remote computer. | Using Web Interface for Remote Administration |
| Create a Web site. | To create a Web site to host your Web applications. | Web Site Setup in Microsoft Internet Information Services |
| Create applications with the latest development tools. | To learn about the latest development tools from Microsoft that can help you develop new applications more quickly and efficiently. | Microsoft Web site |
| Create ASP.NET Web Applications. | To create ASP.NET applications. | Creating ASP.NET Web Application at the Microsoft Web site |
| Secure your ASP.NET Web applications. | To ensure the security of your ASP.NET applications | ASP.NET Web Application Security at the Microsoft Web site |
| Configure ports to allow remote administration. | To manage the application server from other computers on the network. | Windows Firewall Settings |

# Mail server role: Configuring a mail server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Mail server role: Configuring a mail server**

Configure this computer as a mail server to install E-mail Services, which provides e-mail transfer and retrieval services. E-mail Services includes the POP3 service, which provides e-mail retrieval, and the SMTP service, which provides e-mail transfer. Administrators can use the POP3 service to store and manage e-mail accounts on the mail server. After configuring this computer as a mail server, users can connect to the mail server and retrieve e-mail to their local computer using an e-mail client that supports the POP3 protocol, such as Microsoft Outlook.

This topic explains how to use the Configure Your Server Wizard to install and configure E-mail Services. After you have completed the Configure Your Server Wizard, you must perform additional required steps to create mailboxes. After you have completed the Configure Your Server Wizard and created the appropriate mailboxes, you will have a fully-functioning mail server.

You can configure both member servers and stand-alone servers to be a mail server. However, the default authentication method and the available authentication methods will vary. For more information about the default authentication methods and the available authentication methods, see Before you begin and Configure POP3 Service.

This topic covers:


- Before you begin
- Configuring your mail server
- Next steps: Completing additional tasks


**Before you begin**

Before you configure your computer as a mail server, verify that:

- The server on which you intend to install e-mail services has a working Internet connection.
- There is an NTFS partition available. With an NTFS partition, you can take advantage of the increased security provided by disk quotas. For more information about disk quotas, see Configuring disk quotas for the POP3 service.
- You have a registered e-mail domain name. Contact your Internet Service Provider for assistance in registering an e-mail domain name.
- A Mail eXchanger (MX) record for your e-mail domain name exists and matches the name of your server. Contact your Internet service provider (ISP) to create an MX record.
- You have configured your server for static addressing. Contact your Internet Service Provider for the information necessary to configure your server for static addressing. For more information on how to configure your mail server with a static IP address, see Configure TCP/IP for static addressing.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

The following table lists the information that you need to know to before you add a mail server role.

| Before adding a mail server role | Comments |
|---|---|

| | |
|---|---|
| Determine the appropriate level of security for this server. | A server in this role may be targeted by attackers because of its exposure to the Internet and other networks. To ensure the security of this server, it is recommended that you implement security precautions, such as firewalls and Internet Protocol security (IPSec), before placing it in a production environment. For more information, see Internet Protocol Security (IPSec) and Basic Firewall. |
| Determine the appropriate authentication method for your configuration. | You must choose an authentication method before you create any e-mail domains on the mail server. The authentication method can be changed only if there are no existing e-mail domains on the mail server.<br><br>If the computer that you are configuring as a mail server is either a member server or a domain controller, the authentication method setting defaults to Active Directory authentication. Otherwise, the setting defaults to local Windows accounts authentication. |
| Determine that you have a registered e-mail domain name. | The e-mail domain must be a registered domain name and it must match the Mail eXchanger (MX) record created by your ISP. If you do not already have an e-mail domain name, contact your ISP for assistance in registering a domain name.<br><br>**Note**<br><br>• The POP3 service supports top-level and third-level domain names. For example, example.com and mailserver.example.com are both supported.<br><br>• |

**Configuring your mail server**

To configure a mail server, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Mail server (POP3, SMTP)**, and then click **Next**.

This section covers:

- Configure POP3 Service

- Summary of Selections

- Completing the Configure Your Server Wizard

- Creating mailboxes

- Removing the mail server role

**Configure POP3 Service**

On the **Configure POP3 Service** page, under **Authentication method**, click the appropriate method for your deployment. The Windows Server 2003 family supports the authentication methods listed in the following table.

| **Use this authentication method** | **When** |
|---|---|
| Local Windows accounts | Your mail server is not an Active Directory member server, and you want to store user accounts on the server on which the POP3 service is installed |

| | |
|---|---|
| Active Directory-Integrated | Your mail server is a domain controller or a member server |
| Encrypted Password File | Your mail server is not using Active Directory, or you do not want to have user accounts for the POP3 service on the local computer |

The authentication methods that are available to you depend on the configuration of your server:

- If the computer on which the POP3 service is running is a member server in an Active Directory domain, all three authentication methods are available.
- If the computer on which the POP3 service is running is a domain controller, the available authentication methods are Active Directory integrated authentication and encrypted password file authentication.
- Otherwise, the available authentication methods are local Windows accounts authentication and encrypted password file authentication.

Under **E-mail domain name**, type your registered e-mail domain name. You can create additional e-mail domains later by using the POP3 service snap-in or the Winpop command-line tool.

After you finish, click **Next**.

**Summary of Selections**

On the **Summary of Selections** page, you can view and confirm the options that you have selected. If you selected **Mail server (POP3, SMTP)** on the **Server Role** page, the following appears:

- **Install POP3 and Simple Mail Transfer Protocol (SMTP) to enable POP3 mail clients to send and receive mail**

To apply the selections shown on the **Summary of Selections** page, click **Next**. After you click **Next**, the **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page.

**Completing the Configure Your Server Wizard**

After the components are configured, the Configure Your Server Wizard displays the **This Server is Now a Mail Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at *systemroot*\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

At this stage, you have a fully-functioning mail server, but you must also create mailboxes for all of the users in the domain who will be sending or receiving e-mail. Without mailboxes, users cannot send or receive e-mail.

However, before you start to use your mail server, we recommend the following steps:

- Run Windows Update. For more information, see Windows Update.
- Run the Security Configuration Wizard. For more information, see Security Configuration Wizard Overview.
- For the mail server role to work correctly, you may need to manually configure Windows Firewall by adding the ports required by POP3 service and SMTP service. For additional information on the ports required, see Windows Firewall Settings.

**Creating mailboxes**

To send and receive e-mail, each user must have a unique mailbox in the e-mail domain. You can create mailboxes from either the POP3 service MMC snap-in or at the command line. This procedure uses the POP3 service MMC snap-in. For more information about creating mailboxes or administering the POP3 service at the command line, see Winpop. For more information about creating mailboxes, see Create a mailbox.

| Step | Comments |
|---|---|

| | |
|---|---|
| Open the POP3 service MMC snap-in. | To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.<br><br>**Notes**<br><br>• To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.<br><br>• If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure. |
| Create one or more mailboxes. | In the console tree, select the e-mail domain that you specified in the Configure Your Server Wizard (for example, example.com). Right-click the e-mail domain, point to **New**, and then click **Mailbox**. Provide the following information:<br><br>• **Mailbox Name--**the name of the mailbox. The maximum length for a mailbox name is 20 characters for local Windows accounts authentication, and 64 characters for encrypted password file authentication or Active Directory integrated authentication. The minimum length is 1 character.<br><br>• **Password--**the password to access the mailbox.<br><br>• **Confirm Password--**retype the password that was specified in **Password**.<br><br>If you are using Active Directory integrated authentication or local Windows accounts authentication, select the **Create associated user for this mailbox** check box, unless a user account already exists with the same name as the mailbox that you want to create. If the check box is already selected, clear it only if an account already exists with the same name as the mailbox that you want to create. |

**Removing the mail server role**

If you need to reconfigure your server for a different role, you can remove existing server roles. By removing the mail server role, you will uninstall all mail server components, such as the POP3 service and SMTP service. After the mail server components are uninstalled, users will no longer be able to send or receive e-mail using that server. Any e-mail that is stored on the computer will not be affected by removing the mail server role and will remain in the mail store.

To remove the mail server role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Mail server (POP3, SMTP)**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the mail server role** check box, and then click **Next**. After you click **Next**, the **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page. On the **Mail Server Role Removed** page, click **Finish**.

**Next steps: Completing additional tasks**

After you complete the Configure Your Server Wizard and create mailboxes, the computer is ready for use as a mail server. Up to this point, you have completed the following tasks:

- Installed the POP3 service and the SMTP service.

- Configured the POP3 service to use an authentication method.

- Created an e-mail domain.
- Created mailboxes.

The Configure Your Server Wizard automatically installs the POP3 service MMC snap-in, which you use to manage your mail server. To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.

The following table lists additional tasks that you might want to perform on your mail server.

| Task | Purpose of task | Reference |
|---|---|---|
| Provide users with the procedure to configure their e-mail clients to use the mail server. | To connect to the mail server, the user's e-mail client must be configured specifically for the mail server. | Configure E-mail Clients |
| Implement disk quotas. | Disk quotas ensure that the mail store does not use an excessive or unanticipated amount of disk space, which could adversely affect the performance of the server on which the POP3 service is running. You must have an NTFS partition to implement disk quotas. NTFS partitions allow for greater directory and folder security, which better protects e-mail stored on the local hard disk. | Configuring disk quotas for the POP3 service |
| Configure your mail server to require secure e-mail client authentication. | The POP3 service supports Secure Password Authentication (SPA) for Active Directory integrated authentication and local Windows accounts authentication. Secure Password Authentication requires that all e-mail clients transmit both the user name and password using secure authentication. Secure Password Authentication is more secure than the default of plaintext and, therefore, is recommended over plaintext. Secure Password Authentication must be configured on both the server on which e-mail services are running and on every e-mail client that will connect to the mail server. | Configure the mail server to require Secure Password Authentication; Configure Outlook Express for Secure Password Authentication |
| Configure ports to allow remote administration. | To manage the mail server role from other computers on the network. | Windows Firewall Settings |

# Terminal server role: Configuring a terminal server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal server role: Configuring a terminal server**
Configure this computer as a terminal server by installing the Terminal Server component, which provides centralized deployment of applications.

Using a terminal server, users in remote locations can run programs, save files, and use network resources as though those resources were installed on the users' own computers. By installing programs on a terminal server, you can ensure that all users are using the same version of a program. If you plan to use this computer to allow multiple users to access a program at the same time from a single point of installation, configure this computer as a terminal server.

However, if you plan to use this computer for remote administration on Windows Server 2003 operating systems, you do not need to install Terminal Server. Instead, you can use Remote Desktop for Administration (formerly Terminal Services in Remote Administration mode), which is installed by default on computers running one of the Windows Server 2003 operating systems. After you enable remote connections, Remote Desktop for Administration allows you to remotely manage servers from any client over a LAN, WAN, or dial-up connection. Up to two remote sessions, plus the console session, can be accessed at the same time, without requiring Terminal Server Licensing. For more information about Remote Desktop for Administration, see Remote Administration using Terminal Services.

This topic explains how to use the Configure Your Server Wizard to install and configure a terminal server. After you have completed the Configure Your Server Wizard, you must perform the following additional steps in order to have a basic terminal server.

- Confirm Internet Explorer Enhanced Security Configuration settings.
- Configure a Terminal Server License Server. For small deployments, it is acceptable to install both the Terminal Server and Terminal Server Licensing service on the same physical computer. However, for larger deployments, it is recommended that Terminal Server Licensing be installed on a separate server. Install client access licenses (CALs) on the Terminal Server License Server.

  **Important**

- You must configure Terminal Server Licensing correctly in order for your terminal server to continue to accept connections from clients. To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent CAL, or after 120 days, whichever comes first. For more information, see Terminal Server Licensing grace period.
- Install programs on the terminal server.
- Distribute the latest version of Remote Desktop Connection to clients running earlier versions of Remote Desktop Connection for Windows.
- Specify which users have permission to connect to the terminal server.

After you have completed both the Configure Your Server Wizard and these additional required tasks, you will have a basic terminal server.

This topic covers:

Before you begin

Configuring your terminal server

Next steps: Completing additional tasks

**Before you begin**

Before you configure your computer as a terminal server, verify whether or not:

- The operating system is configured correctly. In the Windows Server 2003 family, a terminal server depends on the appropriate configuration of the operating system and its services. If you have a new installation of a Windows Server 2003 operating system, you can use the default service settings. No further action is necessary. If you upgraded to a Windows Server 2003 operating system or you want to confirm that your services are configured correctly for best performance and security, verify your service settings with the table in Default settings for services.

- The computer is a server on a network or in a domain, but is not a domain controller. Installing Terminal Server on a domain controller affects performance because of the additional memory, network traffic, and processor time required to perform the tasks of a domain controller in a domain.

- The computer meets processor and memory requirements for supporting multiple concurrent sessions where different users are logged on. A terminal server requires a minimum of 128 MB RAM, plus additional RAM for each user to support running each user's programs on the server. An additional 10 MB RAM is recommended for each light user, who typically runs one program at a time, and up to 21 MB RAM for each power user, who typically runs three or more programs at the same time. In addition, if you plan to install 16-bit applications on the terminal server, be aware that they consume additional resources when they run in 32-bit environments such as Windows Server 2003 operating systems.

- There are no programs installed on the computer. You should add the Terminal Server role before you install the programs that you want users to access. If there are programs already installed on the computer, you might have to reinstall them to ensure that they work correctly in the Terminal Server environment.

- No users are able to log on remotely to the computer. You should allow users to access the terminal server only after you have installed programs, tested their installation, and performed any tuning necessary for the programs to work in a multisession environment. For information on disabling terminal services connections temporarily, see Disable Terminal Services connections.

- All existing disk volumes use the NTFS file system. FAT32 volumes do not provide either the required level of security for users in a multisession environment or the ability to set file permissions.

- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.

  **Note**

  Later, you will need to configure Windows Firewall to allow an exception, to ensure that clients running Remote Desktop can connect remotely to the terminal server. Additional exceptions might also be needed, depending on your terminal server deployment and network configuration. For more information, see Windows Firewall Settings.

- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

**Configuring your terminal server**

To configure a terminal server, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Terminal server**, and then click **Next**.

This section covers:

Summary of Selections

Completing the Configure Your Server Wizard

Confirming Internet Explorer Enhanced Security Configuration Settings

Configuring a Terminal Server License Server

Installing client access licenses on the Terminal Server License Server

Installing programs on the terminal server

Deploying client software

Giving users permission to access the terminal server

Removing the terminal server role

**Summary of Selections**

On the **Summary of Selections** page, view and confirm the options that you have selected. If you selected **Terminal server** on the **Server Role** page, the following appears:

- **Install Terminal Server**

To apply the selections shown on the **Summary of Selections** page, click **Next**. The following message appears: "During this process, the Configure Your Server Wizard restarts your computer. Before continuing, close any open programs." If you need to close open programs and you want to cancel the configuration of the terminal server role at this time, you must click **Cancel** now. When you click **Cancel**, the **Configure Your Server Wizard** displays the **Cannot Complete** page. To close the Configure Your Server Wizard, click **Finish**. Otherwise, if you click **OK**, the Configure Your Server Wizard begins the configuration process.

Next, the Configure Your Server Wizard displays the message "Installing Terminal Server." The **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page. Then, the Configure Your Server Wizard shuts down the computer and restarts it to accept the configuration changes that make the computer a terminal server.

During the restart process, a dialog box displays progress messages, for example, "Windows is starting up" and "Preparing network connections." Depending on the size of your network, preparing network connections could take some time. When the **Welcome to Windows** dialog box appears, press CTRL+ALT+DEL. In the **Log on to Windows** dialog box, in **Password**, type your password. To complete the process, wait for the Configure Your Server Wizard to appear on the screen.

**Completing the Configure Your Server Wizard**

After your server restarts, the Configure Your Server Wizard displays the **This Server is Now a Terminal Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at systemroot\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

To verify that your server is secure and has the most recent updates, do the following:

1. Run Windows Update. For more information, see Windows Update.

2. Run the Security Configuration Wizard. For more information, see Security Configuration Wizard Overview.

Next, you must complete the following steps so that your server is ready to function as a basic terminal server:

- Confirm Internet Explorer Enhanced Security Configuration settings.

- Configure a Terminal Server License Server.

- Install client access licenses (CALs) on the Terminal Server License Server.

- Install programs on the terminal server.

- Deploy the Remote Desktop Connection .msi file to clients not running Windows XP or Windows Server 2003 operating systems.

- Give users permission to access the terminal server.

A separate window displays checklists that provide information about these additional requirements. The same information is covered in this document.

To run a terminal server, you need to configure Terminal Server Licensing. For small deployments, you can configure Terminal Server Licensing on the same physical computer as the terminal server. For larger deployments, you should install Terminal Server Licensing on a separate server. If a Terminal Server License Server is already installed, you can skip the steps for configuring a Terminal Server License Server and installing CALs, and begin Installing programs on the terminal server. Otherwise, if the Manage Your Server page displays a message indicating that a Terminal Server License Server was not found, you must configure a Terminal Server License Server before you can use your terminal server.

**Confirming Internet Explorer Enhanced Security Configuration settings**
After you complete the Configure Your Server Wizard and install Terminal Server, you can configure Internet Explorer Enhanced Security Configuration settings.

If you activate these settings, Internet Explorer applies the following security settings to a user who logs on as an administrator:

- High security settings to the Internet and Local intranet security zones

- Medium security settings to the Trusted sites zone

By applying high security settings to the Internet and Local intranet security zones, you disable scripts, Microsoft ActiveX® controls, and the Microsoft virtual machine (Microsoft VM) for HTML content in these zones. You also prevent users from downloading files in these zones.

By applying medium security settings to the Trusted sites zone, you set standard browsing functionality. If you use sites for administrative tasks and Web-based applications that an administrator cannot access after you apply these settings, you can add the site addresses to the list of sites in the Trusted sites zone.

To review or change the Internet Explorer Enhanced Security Configuration settings, in Manage Your Server, click **Internet Explorer Enhanced Security Configuration**.

In the Windows Server 2003 family, you can implement enhanced security settings for Internet Explorer for all users and reduce the exposure of your server to Web sites that might pose a security risk. For more information, see Internet Explorer Enhanced Security Configuration.

**Configuring a Terminal Server License Server**
Configure a Terminal Server License Server, either on the same computer for which you have just configured the terminal server role (for small deployments), or on another computer (recommended for larger deployments). A Terminal Server License Server manages licenses for Terminal Services client connections. You are required to activate a Terminal Server License Server only once, after which the Terminal Server License Server becomes the repository for terminal server client licenses. Until the registration process is completed, your Terminal Server License Server can issue temporary licenses for clients.

**Important**

- You must configure Terminal Server Licensing correctly in order for your terminal server to continue to accept connections from clients. To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent CAL, or after 120 days, whichever comes first. For more information, see Terminal Server Licensing grace period.

The easiest and quickest way to activate a Terminal Server License Server is by using the Automatic method. To use this method, the computer running the Terminal Services Licensing service must have a direct connection to the Internet. For information on activation methods for computers that are not connected to the Internet, see Activate a Terminal Server license server by using a Web browser and Activate a Terminal Server license server by using the telephone.

The following table shows the steps you must take to configure and activate a Terminal Server License Server by using the

Automatic method.

| Task | Comments |
| --- | --- |
| Install the Terminal Server Licensing service. | Open **Add or Remove Programs** in Control Panel, and then click **Add/Remove Windows Components**. In the Windows Components Wizard, select the **Terminal Server Licensing** check box, and then click **Next**. If your network includes several domains, or if you are installing the Terminal Server Licensing service on a member server, choose **Your entire enterprise**. If you want to maintain a separate Terminal Server License Server for each domain, or if your network includes workgroups or Windows NT 4.0 domains, choose **Your domain or workgroup**. If you want to change the location of the license server database, specify a new location, and then click **Next**. The **Configuring Components** page displays the progress of configuration changes. On the **Completing the Windows Components Wizard** page, click **Finish**, and then click **Close**.<br><br>To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**. |
| Activate the Terminal Server License Server. | Open **Terminal Server Licensing**, right-click the Terminal Server License Server that you want to activate, and then click **Activate Server**. The Terminal Server License Server Activation Wizard starts. On the **Connection method** page, under **Activation method**, click **Automatic connection**, and then click **Next**. On the **Company Information** page, provide the following required information:<br>• First name<br>• Last name<br>• Company name<br>• Country or region<br><br>Confirm that the information you typed is correct, and then click **Next**. On the next **Company Information** page, you can provide the following optional information:<br>• Email address<br>• Organizational unit<br>• Company address<br>• City<br>• State or province<br>• Postal code<br><br>Confirm that the information you typed is correct, and then click **Next**. On the **Completing the Terminal Server License Server Activation Wizard** page, under **Status**, the following message appears: "Your license server has been successfully activated." If you want to install client licenses now, click **Next**. If you want to postpone the installation of client licenses, clear the **Start Terminal Server Client Licensing Wizard now** check box, and then click **Finish**.<br><br>**Note**<br>• To open Terminal Server Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**. |

**Installing client access licenses on the Terminal Server License Server**

After you activate a Terminal Server License Server, the next step is to install client access licenses (CALs) on the Terminal Server License Server.

**Important**

- Your Terminal Server License Server can issue temporary licenses. Temporary licenses are designed to allow you ample time to deploy a license server, and they allow clients to connect to the terminal server for 90 days. There is no limit to the number of temporary licenses that a license server can issue, but a single client is only issued a temporary license once. After the temporary license expires, the client can only connect to the terminal server if the license server can issue a permanent CAL, or if the terminal server is still within its licensing grace period. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent CAL, or after 120 days, whichever comes first. For more information, see Terminal Server Licensing grace period.

CALs are digitally-signed certificates that each client stores locally. All CALs are installed on a Terminal Server License Server. When a client attempts to log on to a terminal server for the first time, the terminal server recognizes that the client has not been issued a CAL and locates a Terminal Server License Server to issue a new CAL to the client. For information about specific license requirements, see the Microsoft Web Site.

Before you install CALs, you must have your licensing agreement numbers ready, and know which method you used to purchase them.

The easiest and quickest way to install CALs on a Terminal Server License Server is by using the Automatic method. To use this method, the computer running the Terminal Services Licensing service must have a direct connection to the Internet. For information on installing CALs for computers that are not connected to the Internet, see Install client access licenses by using a Web browser and Install client access licenses by using the telephone.

The following table shows the steps you must take to install CALs on a Terminal Server License Server by using the Automatic method.

| Task | Comments |
|---|---|
| Install CALs on the Terminal Server License Server. | On the Terminal Server License Server, open **Terminal Server Licensing**. Verify that the installation method for the Terminal Server License Server is set to Automatic by right-clicking the Terminal Server License Server for which you want to install CALs, and then clicking **Properties**. If necessary, on the **Installation Method** tab, change the installation method to **Automatic connection**, and then click **OK**.<br><br>In the Terminal Server Licensing console tree, right-click the Terminal Server License Server on which you want to install CALs, click **Install Licenses**, and then click Next. The Terminal Server CAL Installation Wizard starts. On the **Licensing program** page, choose the license program under which you purchased your licenses, and then click **Next**. On the **License Code** page, type the license code for each license you have purchased, and then click **Add** after each entry. After you have typed all of the license codes, click **Next**. The **Completing the Terminal Server CAL Installation Wizard** page displays a message that the CALs were successfully installed. To close the wizard, click **Finish**.<br><br>**Note**<br><br>• To open Terminal Server Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Server Licensing**. |

**Installing programs on the terminal server**

At this stage, you have accomplished the following tasks:

- Completed the Configure Your Server Wizard and configured the terminal server role on your server.
- Installed Terminal Server Licensing.
- Activated the Terminal Server License Server.
- Installed CALs on the Terminal Server License Server.

Now you are ready to install programs on the terminal server. **Add or Remove Programs** in Control Panel is the preferred method for program installation, and you should use this method whenever possible. This section describes how to use **Add or Remove Programs** to install programs on a terminal server.

There are other program installation methods, such as the `change user` command, Windows Installer packages (.msi files), and Group Policy Software Installation. For more information about the `change user` command, see Install a program using the change user command. For more information about using Windows Installer, see Assigned and published programs. For more information about Group Policy, see Group Policy.

For improved performance and reduced network traffic, install programs on the local drive of the terminal server instead of on a file server. Ensure that you have enough space to install programs on NTFS file system drives instead of on FAT32 drives. NTFS drives allow you to set file permissions, which you cannot do on FAT32 drives.

If you are installing published programs, you must use another installation method, such as Group Policy Software Installation.

For performance and security reasons, you should use 32-bit programs whenever possible. Most 32-bit programs use the registry to read and write program settings and need to write only to specific registry values. Running 16-bit programs can reduce the number of users a processor supports by 40 percent and increase the memory required for each user by 50 percent. In addition, some 16-bit programs must be able to write to the directory where the program's .ini file is stored.

RAM and CPU requirements increase approximately linearly with the number of sessions running. To reduce RAM and CPU requirements, consider restricting user or group access to certain program types, disabling unnecessary program features, or installing programs on separate terminal servers.

Some programs have known installation issues in a multisession environment. For information about programs that require installation scripts in order to work correctly in a multisession environment, see Optimizing Applications for Windows 2000 Terminal Services and Windows NT Server 4.0, Terminal Server Edition at the Microsoft Web site.

**Application compatibility considerations**

You should install programs from the console session of the terminal server. You can install programs from a remote console session, but this is not the preferred method for installing programs.

Some programs require an application compatibility script to be run after the program is installed. The scripts are stored in the systemroot\Application Compatibility Scripts\Install directory on the terminal server.

You should be aware of the implications of the security mode in which the terminal server operates. There are two security modes:

- *Full security* provides the most secure environment for users connecting to a terminal server. To run in this mode, applications must be written to run in the security context of an ordinary user. For Windows Server 2003 operating systems and Windows 2000, full security is the default.

- *Relaxed security* enables you to run programs that otherwise might not work at all in the more rigorous Full security mode. However, in Relaxed security mode (also known as Windows NT 4.0/Terminal Server Edition permissions compatibility mode), any user on the system can change files and registry settings in many places throughout the system, although others users' data files might not be visible. A malicious user could exploit this situation by replacing a known and trusted program with a program of the same name but some harmful intent. If the operating system on your terminal server was installed using the Upgrade method, the security mode might be set to Relaxed security. When in doubt, you should choose Full security, test your applications in that mode, and change the security mode only if your test results indicate the need to do so.

The following table shows the steps you must take to install programs on a terminal server, using **Add or Remove Programs**.

| Task | Comments |
|---|---|
| Ensure that no users are logged on to the terminal server. | Send a message to all users who are logged on to the terminal server. Program installation often requires restarting the computer, and their sessions will be disconnected. You should not allow users to access the terminal server until programs have been installed and tested. |
| Disable Terminal Services connections temporarily. | Right-click **My Computer**, click **Properties**, click the Remote tab, and then clear the **Allow users to connect remotely to this computer** check box. |

| | |
|---|---|
| Specify **Full Security** as the security mode. | Open **Terminal Services Configuration**. In the console tree, click **Server Settings**, right-click **Permission Compatibility**, and then click **Properties**. In the **Permission Compatibility** dialog box, click **Full Security**, and then click **OK**.<br><br>**Note**<br><br>• To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**. |
| Install programs from a CD or floppy disk. | Ensure that you are logged on as a member of the Administrators group on the terminal server. Open **Add or Remove Programs** in Control Panel, and then click **Add New Programs**. Click **CD or Floppy**. Insert the CD or floppy disk into the appropriate drive, and then click **Next**. Verify that the installation file is specified correctly in the **Open** box on the **Run Installation Program** page, and then click **Finish**. Follow the instructions in the program's installation wizard. After the program is installed, edit and run any applicable scripts to tailor the program for a multisession environment.<br><br>**Note**<br><br>• To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**. |
| Test the installation. | Ensure that event logging is enabled by opening Services in Administrative Tools. Create a temporary user account that mimics the settings of the user or users who will access the program, and use the account to log on to the terminal server. Start the program and step through some basic tasks. Then, use Event Viewer to determine which files or directories need Write access and which registry keys require Read access by the user for correct operation. Note that this process might not find all files, directories, and registry keys for which the application requires access in all user scenarios. The only way to ensure that you have accounted for all access requirements is to perform tasks manually.<br><br>Some programs enable users to start other programs. For example, Microsoft Access has a toolbar that can be used to start other Microsoft Office programs. If you want users to have access only to specified programs when they log on to the terminal server, you should disable toolbar access from within programs that you install on the terminal server.<br><br>**Note**<br><br>• To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**. |
| Tune programs for multisession use. | Use a text editor such as Notepad to modify any scripts, and then run the scripts to tune any programs that require it. To obtain the scripts, see Optimizing Applications for Windows 2000 Terminal Services and Windows NT Server 4.0, Terminal Server Edition at the Microsoft Web site (http://www.microsoft.com/). |
| Run application compatibility scripts. | Navigate to the *systemroot*\Application Compatibility Scripts\Install directory on the terminal server and run scripts for any programs that require them. |
| Enable remote connections on the terminal server. | Right-click **My Computer**, click **Properties**, click the **Remote** tab, and then check the **Allow users to connect remotely to your computer** check box.<br><br>**Note**<br><br>• Depending on your desktop settings, **My Computer** might not appear on your desktop. To show or hide desktop icons, right-click somewhere on the desktop, click **properties**, click the **Desktop** tab, click **Customize Desktop**, and then, under **Desktop icons**, select the check box next to the icon you want to display, or clear the check box next to the icon you want to hide. |

**Deploying Client Software**

Remote Desktop Connection, formerly known as the Terminal Services Client, is installed automatically on computers running Windows XP and Windows Server 2003 operating systems. For performance and security reasons, computers running earlier versions of Microsoft Windows, including Windows 2000 Server, Windows 2000 Professional, Windows NT 4.0, Windows 98, and Windows 95, should have the latest version of Remote Desktop Connection installed.

There are several ways to deploy the client software:

- Share the Msrdpcli.msi file and use Microsoft IntelliMirror to distribute it to workstations running Windows 2000.
- Download Remote Desktop Connection directly from the Microsoft Web site.
- Place the .msi file in a shared folder residing on a server on the network.

This topic describes how to install the client software from a shared folder residing on a server on the network.

Before you deploy the client software, decide whether you want the software to be installed for the use of a single user or for anyone who uses the client computer. You will make this choice during the deployment process.

The following table shows the steps you must take to deploy the latest version of Remote Desktop Connection to clients running earlier versions of either Windows or Remote Desktop Connection.

| Task | Comments |
|---|---|
| Share the client setup folder. | On the computer running a Windows Server 2003 operating system, open **Windows Explorer**. Navigate to the systemroot\System32\Clients\Tsclient\win32 folder, right-click the win32 folder, click **Sharing and Security**. On the **Sharing** tab, click **Share this folder**, and then click **OK**.<br>**Note**<br>• To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**. |
| Install Remote Desktop Connection. | On the client computer, click **Start**, click **Run**, and then, in **Open**, type *ServerName*\win32, where *ServerName* is the name of the computer where the shared folder is located. Double-click the msrdpcli.msi file to start the InstallShield Wizard for Remote Desktop Connection, and then click **Next**. Read the License Agreement, click **I accept the terms in the license agreement**, and then click **Next**. Type your name and organization in the **Customer Information** page, click **Anyone who uses this computer (all users)**, and then click **Next**. On the **Ready to Install the Program** page, either click **Back** to review or change any of your installation settings, or click **Install** to begin the installation. To complete the installation, click **Finish**. |

**Giving users permission to access the terminal server**

By default, on Windows Server 2003 operating systems, members of the Administrators and Remote Desktop Users groups can use Terminal Services connections to connect to a remote computer. The Remote Desktop Users group is not populated by default, so you must decide which users and groups should have permission to log on remotely, and then manually add them to this group.

**Important**

- You must use the Remote Desktop Users group to grant selected users and groups the necessary permission to make Terminal Services connections to remote computers.

Membership in the Remote Desktop Users group does not also put the user into the local Users group. Depending on the contents of your local Users group, you might need to add the user to that group also.

Before you give users permission to access the terminal server, you must:

- Check the membership of the Administrators group to ensure that you know who has access to the terminal server.
- Decide which users should have permission to access the terminal server.
- Determine which users must also be added to the local Users group.

The following table shows the steps you must take to give users permission to access the terminal server.

| Task | Comments |
|---|---|
| Add users to the Remote Desktop Users group. | Open **Computer Management (Local)**, and in the console tree, **click Local Users and Groups**. In the details pane, double-click the **Groups** folder, double-click **Remote Desktop Users**, and then click **Add**. In the **Select Users** dialog box, click **Locations** to specify the search location. To specify the types of objects that you want to search for, click **Object Types**. In this case, you want to search for Users or Groups. Type the name that you want to add in the **Enter the object names to select (examples)** box, and then click **Check Names**. When the name is located, click **OK**.<br>**Note**<br>• To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**. |
| Add users to the local Users group, if they are not already members. | Open **Computer Management (Local)**, and in the console tree, **click Local Users and Groups**. In the details pane, double-click the **Groups** folder, double-click **Users**, and then click **Add**. In **the Select Users** dialog box, click **Locations** to specify the search location. To specify the types of objects that you want to search for, click **Object Types**. In this case, you want to search for Users or Groups. Type the name that you want to add in the **Enter the object names to select (examples)** box, and then click **Check Names**. When the name is located, click **OK**.<br>**Note**<br>• To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**. |

**Removing the terminal server role**

If you need to reconfigure your server for a different role, you can remove existing server roles. If you remove the terminal server role, you will need to reinstall all software, review and update any file or registry permissions for which you changed default values, and review and update any software restriction policies that were used to control programs running on the terminal server.

To remove the terminal server role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Terminal server**, and then click **Next**. On **the Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the terminal server role** check box, and then click **Next**. The following message appears: "During this process, the Configure Your Server Wizard restarts your computer. Before continuing, close any open programs." If you need to close open programs and you want to cancel the removal of the Terminal Server role at this

time, you must click **Cancel** now. When you click **Cancel**, the Configure Your Server Wizard displays the **Cannot Complete** page. To close the Configure Your Server Wizard, click **Finish**. Otherwise, if you click **OK**, the Configure Your Server Wizard begins the removal process.

Next, the Configure Your Server Wizard displays the "Removing Terminal Server" message. The **Configuring Components** page of the Windows Components Wizard appears, displays messages about the configuration changes being made to the computer, and then closes. The Configure Your Server Wizard shuts down the computer and restarts it to accept the configuration changes that remove this role.

During the restart process, a dialog box displays progress messages, for example, "Windows is starting up" and "Preparing network connections." Depending on the size of your network, preparing network connections could take some time. When the **Welcome to Windows** dialog box appears, press CTRL+ALT+DEL. In the **Log on to Windows** dialog box, in **Password**, type your password. To complete the process, wait for the Configure Your Server Wizard to appear on the screen. On the **Terminal Server Role Removed** page, click **Configure Your Server log** to see a record of your changes, and then click **Finish**.

After you remove the terminal server role, you should:

- Reinstall all software.
- Review any file or registry permissions for which you changed default values and, if necessary, make changes.
- Review any software restriction policies used to control programs running on the terminal server and, if necessary, make changes.

**Next steps: Completing additional tasks**

After you complete the Configure Your Server Wizard and associated tasks, the computer is ready for use as a basic terminal server that can accept multiple connections from remote clients. Up to this point, you have completed the following tasks:

- Run the Configure Your Server Wizard.
- Activated a Terminal Server License Server.
- Installed CALs on the Terminal Server License Server.
- Installed applications on the terminal server.
- Deployed the Remote Desktop Connection .msi file to clients not running Windows XP or Windows Server 2003 operating systems.
- Configured user permissions for user access to the terminal server.

The following table lists some additional tasks you might want to perform on your terminal server.

| Task | Purpose of task | Reference |
|---|---|---|
| Manage Terminal Services connections. | To enable, disable, rename, or delete a connection. | Manage Terminal Services Connections |
| Specify connection permissions. | To grant terminal server access only to selected users and groups.<br><br>To identify which users and groups are permitted to perform a given task or tasks on the terminal server. | Managing Terminal Services users; Managing Permissions on Connections |

| | | |
|---|---|---|
| Configure terminal server settings using either Group Policy or Terminal Services Configuration. | To configure settings such as Active Desktop, temporary folders, and session limits for individual users. | Configure Server Settings |
| Deploy Remote Desktop Web Connection. | To allow users to create a Remote Desktop connection within Internet Explorer, even though the Remote Desktop Connection client is not installed on their computers. | About Remote Desktop Web Connection |
| Control programs running in a terminal server session. | To protect terminal servers and users from unknown, or possibly malicious, programs. | Using Software Restriction Policies in Windows XP and the Windows Server 2003 family to Protect Against Unauthorized Software at the Microsoft Web site |
| Configure Session Directory settings. | To ensure that users are transparently reconnected to the original server hosting their disconnected Terminal Server sessions. This task applies to terminal servers that are part of a cluster of terminal servers, and requires that a server running either Windows Server 2003, Enterprise Edition, or a Windows Server 2003, Datacenter Edition, is visible on the network, and has the Session Directory service enabled. This session directory server should not be the server on which the Terminal Server role is configured. | Load balancing and terminal servers |
| Configure ports to allow incoming connections to terminal servers. | • To ensure that clients running Remote Desktop Connection can connect remotely to terminal servers.<br>• Additional port configuration might also be required, if Remote Desktop Web connection is deployed, SSL is enabled on the Web server, and if the terminal server license server and terminal server are on opposite sides of a firewall. | Windows Firewall Settings |

# Remote access/VPN server role: Configuring a remote access/VPN server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access/VPN server role: Configuring a remote access/VPN server**
You can configure a server that allows remote users to access resources on your private network over dial-up or virtual private network (VPN) connections. This type of server is called a remote access/VPN server. Remote access/VPN servers can also provide network address translation (NAT). With NAT, the computers on your private network can share a single connection to the Internet. With VPN and NAT, your VPN clients can determine the IP addresses of the computers on your private network, but other computers on the Internet cannot.

This topic explains the basic steps for configuring a remote access/VPN server using Manage Your Server, the Configure Your Server Wizard, and the Routing and Remote Access Server Setup Wizard. After you finish configuring a basic remote access/VPN server, you can complete additional configuration tasks, depending on how you want to use the remote access/VPN server.

This topic covers:

Before you begin

Configuring your remote access/VPN server

Next steps: Completing additional tasks

**Before you begin**
Before you configure your server as a remote access/VPN server, you should verify whether or not:

- The operating system is configured correctly. In the Windows Server 2003 family, remote access/VPN depend on the appropriate configuration of the operating system and its services. If you have a new installation of a product in the Windows Server 2003 family, you can use the default service settings. No further action is necessary. If you upgraded to a product in the Windows Server 2003 family or you want to confirm that your services are configured correctly for best performance and security, verify your service settings by comparing them to the table in Default settings for services.

- Your server is correctly configured for optimal security for your network needs. Because your remote access/VPN server will connect your private network, the Internet, and your remote clients, you must make sure the server is secure. The security of your private network depends on the security of your remote access/VPN server. For more information, see Security information for remote access.

- This computer has two network interfaces, one that connects to the Internet and one that connects to the private network. The connection to the Internet must be a dedicated connection with enough bandwidth that VPN users can connect to your private network and users on your private network can connect to the Internet. The connection to computers on your private network must be made through a hardware device, such as a network adapter.

- All needed network protocols have been installed for your network interfaces. For more information, see Network interfaces.

- Windows Firewall is disabled on the server that you want to configure for remote access/VPN. You will configure the Basic Firewall feature of Routing and Remote Access during setup, which will serve in place of Windows Firewall.

- Internet Connection Sharing is disabled on the server that you want to configure for remote access/VPN. Internet Connection Sharing is not compatible with Routing and Remote Access. Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

The following table lists the information that you need to know before you configure a remote access/VPN server.

| Before adding a remote access/VPN server role | Comments |
|---|---|
| Determine which network interface connects to the Internet and which network interface connects to your private network. | During configuration, you will be asked to choose which network interface connects to the Internet. If you specify the incorrect interface, your remote access/VPN server will not operate correctly. |
| Determine whether remote clients will receive IP addresses from a Dynamic Host Configuration Protocol (DHCP) server on your private network or from the remote access/VPN server that you are configuring. | If you have a DHCP server on your private network, the remote access/VPN server can lease 10 addresses at a time from the DHCP server and assign those addresses to remote clients. If you do not have a DHCP server on your private network, the remote access/VPN server can automatically generate and assign IP addresses to remote clients. If you want the remote access/VPN server to assign IP addresses from a range that you specify, you must determine what that range should be. |
| Determine whether you want connection requests from VPN clients to be authenticated by a Remote Authentication Dial-In User Service (RADIUS) server or by the remote access/VPN server that you are configuring. | Adding a RADIUS server is useful if you plan to install multiple remote access/VPN servers, wireless access points, or other RADIUS clients to your private network. For more information, see Internet Authentication Service. |
| Determine whether VPN clients can send DHCP messages to the DHCP server on your private network. | If a DHCP server is on the same subnet as your remote access/VPN server, DHCP messages from VPN clients will be able to reach the DHCP server after the VPN connection is established. If a DHCP server is on a different subnet than your remote access/VPN server, make sure that the router between subnets can relay DHCP messages between clients and the server. If your router is running a Windows Server 2003 operating system, you can configure the DHCP Relay Agent service on the router to forward DHCP messages between subnets. |
| Verify that all users have user accounts that are configured for dial-up access. | Before users can connect to the network, they must have user accounts on the remote access/VPN server or in Active Directory. Each user account on a stand-alone server or a domain controller contains properties that determine whether that user can connect. On a stand-alone server, you can set these properties by right-clicking the user account in Local Users and Groups and clicking **Properties**. On a domain controller, you can set these properties by right-clicking the user account in the Active Directory Users and Computers console and clicking **Properties**. For more information, see Dial-in properties of a user account and Windows interface administrative tool reference A-Z: Active Directory Users and Computers. |

**Configuring your remote access/VPN server**

To configure a remote access/VPN server, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Configuration Options** page, click **Custom configuration** and click **Next**. On the **Server Role** page, click **Remote access/VPN server**, and then click **Next**.

This section describes the steps in the Routing and Remote Access Server Setup Wizard for configuring a remote access/VPN server that is not part of an Active Directory domain or part of a network with DNS or DHCP servers. If you follow these steps, you will configure a remote access/VPN server that provides both dial-up and VPN access for remote access clients, provides

NAT for computers on your private network, generates and assigns IP addresses for remote access clients, and locally authenticates connection requests.

This section covers:

Summary of Selections

Using the Routing and Remote Access Server Setup Wizard

Completing the Configure Your Server Wizard

Completing configuration in Routing and Remote Access

Removing the remote access/VPN server role

**Summary of Selections**

On the **Summary of Selections** page, you can view and confirm the options that you have selected. If you clicked **Remote access/VPN server** on the **Server Role** page, the following line appears:

- **Run the Routing and Remote Access Server Setup Wizard to set up routing and VPN**

To apply the selections shown on the **Summary of Selections** page, click **Next**. The Configure Your Server Wizard starts the Routing and Remote Access Server Setup Wizard. If you cancel the Routing and Remote Access Server Setup Wizard, your remote access/VPN server will not be configured, the Routing and Remote Access service will not be started, and the Configure Your Server Wizard will display the **Cannot Complete** page.

When you complete the Routing and Remote Access Server Setup Wizard and the Configure Your Server Wizard, the Routing and Remote Access service is started automatically.

**Using the Routing and Remote Access Server Setup Wizard**

After you choose the remote access/VPN role and confirm your Summary of Selections by clicking **Next** in the Configure Your Server Wizard, the Routing and Remote Access Server Setup Wizard starts.

This section describes the following steps in the Routing and Remote Access Server Setup Wizard:

Configuration

VPN Connection

IP Address Assignment

Name and Address Translation Services

Address Assignment Range

Managing Multiple Remote Access Servers

Completing the Routing and Remote Access Server Setup Wizard

**Configuration**

On the **Configuration** page, click **Virtual Private Network (VPN) access and NAT**, and click **Next**.

**Important**

- This document describes the **Virtual Private Network (VPN) access and NAT** configuration only. If you decide to choose a different configuration, review the documentation for Routing and Remote Access before you complete the Routing and Remote Access Server Setup Wizard. This document will not help you complete any other role than **Virtual Private Network (VPN) access and NAT**. For more information about other configurations, see Common server configurations for remote access servers.

**VPN Connection**

On the **VPN Connection** page, click the network interface that connects this computer to the Internet. The network interface that you choose will be configured to receive connections from VPN clients. Any interface that you do not choose will be configured as a connection to your private network.

In Network Interfaces, the **Enable security on the selected interface by setting up Basic Firewall** check box will already be

selected. Do not clear this check box. This option configures Basic Firewall, a dynamic packet filtering service that helps protect your private network from unsolicited network traffic.

After you finish, click **Next**.

**IP Address Assignment**

On the **IP Address Assignment** page, the **Automatically** option is selected automatically. Do not change the selection. This selection configures your server to generate and assign IP addresses to remote clients.

After you finish, click **Next**.

**Name and Address Translation Services**

On the **Name and Address Translation Services** page, the **Enable basic name and address services** option is selected automatically. Do not change the selection. This selection configures your server to automatically assign IP addresses to any computer on your private network that requests one. The selection also configures your server to forward name resolution requests to a DNS server on the Internet.

After you finish, click **Next**.

**Address Assignment Range**

The **Address Assignment Range** page displays the range of addresses that is defined for assignment to any computer on your network that requests one. This range is generated based on the IP address of the network adapter you chose on the **VPN Connection** page. Review the information presented.

After you finish, click **Next**.

**Managing Multiple Remote Access Servers**

On the **Managing Multiple Remote Access Servers** page, the **No, use Routing and Remote Access to authenticate connection requests** option is selected automatically. Do not change the selection. This selection configures your server to authenticate connection requests locally by using Windows authentication, Windows accounting, and locally stored remote access policies.

After you finish, click **Next**.

**Completing the Routing and Remote Access Server Setup Wizard**

On the **Completing the Routing and Remote Access Server Setup Wizard** page, review the summary information. Verify that:

- The correct network interface is configured to provide VPN access.
- Dial-up and VPN clients are assigned to your private network for addressing.
- Client connections are accepted and authenticated using remote access policies for this remote access/VPN server.
- NAT is configured for the correct network interface.
- Clients will be assigned IP addresses from the correct range.

If any of the summary information is incorrect, click **Back**, and then change the information.

If you click **Finish**, you will not be able to open the Routing and Remote Access Server Setup Wizard again, unless you either remove the remote access/VPN server role from within the Configure Your Server Wizard or disable Routing and Remote Access from the Routing and Remote Access snap-in.

After you have ensured that the summary information is correct, click **Finish**. A message will appear informing you that, to support the relaying of DHCP messages from remote access clients to a DHCP server, you must open Routing and Remote Access on the remote access/VPN server and configure DHCP Relay Agent with the IP address of a DHCP server. Click **OK**. The Routing and Remote Access service will be started automatically, and the Configure Your Server Wizard will reappear.

**Completing the Configure Your Server Wizard**

After you complete the Routing and Remote Access Server Setup Wizard, the Configure Your Server Wizard displays the **This Server is Now a Remote Access/VPN Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at *systemroot*\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click

**Finish**.

After you complete the Configure Your Server Wizard:

- Run Windows Update. For more information, see Windows Update.
- Run the Security Configuration Wizard. For more information, see Security Configuration Wizard Overview.

You are now ready to complete configuration for your remote access/VPN server in Routing and Remote Access.

**Completing configuration in Routing and Remote Access**

To open Routing and Remote Access, click **Manage this remote access/VPN server** from Manage Your Server. You can also open Routing and Remote Access from Administrative Tools. To open Administrative Tools, click **Start**, click **Control Panel**, and then double-click **Administrative Tools**.

In Routing and Remote Access, double-click the server you have just configured, and then click **Remote Access Policies**. The default remote access policy is set to deny all access to everyone. Your users will not be able to connect to your remote access/VPN server until you edit the default policy to allow access or replace the default policy with your own policies. Review the default policy by double-clicking it, and then edit it to specify the access that you want to allow your users.

If this server has been previously configured as a remote access/VPN server or if IAS has been configured on this server, the remote access policy or policies that appear in Routing and Remote Access might be configured differently from the default remote access policy. Carefully review all of your remote access policies to ensure that you have allowed and denied remote access according to your network needs. Be sure that you are not accidentally allowing or denying more remote access to your network than you intend.

For more information, see Add a remote access policy, Introduction to remote access policies, and Re-create the default remote access policy.

After you have configured remote access policies, you have completed all necessary configuration for a remote access/VPN server on a network without a DHCP server. If your network uses a DHCP server, you must also configure the DHCP Relay Agent before configuration is complete. In Routing and Remote Access, double-click **IP Routing**, right-click **DHCP Relay Agent**, and click **Properties**. Type the IP address of the DHCP server for your network in **Server address**, click **Add**, and then click **OK**.

**Removing the remote access/VPN server role**

If you need to reconfigure your server for a different role, you can remove existing server roles. When you remove the remote access/VPN server role, your server will no longer provide dial-up or VPN access for remote access clients. Additionally, your server will no longer provide NAT for computers on your private network. Remote users will not be able to connect to your private network, and the computers on your private network might not be able to connect to the Internet. Basic Firewall will no longer protect the computers on your private network. After you remove the remote access/VPN server role, consider enabling Windows Firewall or adding another firewall to protect the computers on your private network, if your network does not already have one. Test your private network to make sure computers on the private network have the level of access to the Internet required by your business needs. Reconfigure remote access policies in IAS to deny all remote access attempts.

To remove the remote access/VPN server role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Remote access/VPN server**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the remote access/VPN server role** check box, and then click **Next**. In the dialog box that asks you to confirm that you want to disable the router and remove the remote access configuration, click **Yes**. On the **Remote Access/VPN Server Role Removed** page, click **Finish**.

**Next steps: Completing additional tasks**

After you complete the Configure Your Server Wizard and complete configuration in Routing and Remote Access, your server is ready for use as a remote access/VPN server that provides both VPN access and NAT. Up to this point, you have completed

the following:

- Started the Routing and Remote Access service.
- Configured your server to accept VPN and dial-up connections.
- Configured your server to provide NAT for your private network.
- Configured remote access policies to allow users to connect to the server running Routing and Remote Access.
- If the network uses a DHCP server, configured the DHCP Relay Agent to forward DHCP messages from remote access clients to the DHCP server.

If you have completed all of these tasks, you have created a basic remote access/VPN server that will allow remote computers to connect to your server with a dial-up or a VPN connection and provide network address translation (NAT) for your private network.

The following table lists additional tasks that you might want to perform on your remote access/VPN server.

| **Task** | **Purpose of task** | **Reference** |
|---|---|---|
| Configure static packet filters. | To add static packet filters to better protect your network. | Add local host filters; Packet filtering |
| Configure services and ports. | To choose which services on your private network you want to make available for remote access users. | Configure services and ports |
| Adjust logging levels for routing protocols. | To configure the level of event details that you want to log. You can decide what information you want to track your log files. | Log details for a routing protocol |
| Configure the number of VPN ports. | To add or remove VPN ports. | Add PPTP or L2TP ports |
| Create a Connection Manager profile for your users. | To manage the client connection experience for your users and simplify troubleshooting client connections. | Connection Manager Administration Kit |
| Add Certificate Services. | To configure and manage a certification authority (CA) on a server for use in a public key infrastructure (PKI). | Certificate Services; Computer certificates for L2TP/IPSec VPN connections |
| Increase remote access security. | To protect your remote users and your private network by enforcing the use of secure authentication methods, requiring higher levels of data encryption, and more. | Security information for remote access |
| Increase VPN security. | To protect your remote users and your private network by requiring the use of secure routing and tunneling protocols, configuring account lockout, and more. | VPN Security |

# Domain controller role: Configuring a domain controller

Updated: August 26, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain controller role: Configuring a domain controller**
This topic is about configuring a domain controller that runs Windows Server 2003. For more information about installing Active Directory Domain Services in Windows Server 2008, see AD DS Installation and Removal Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkID=139657).

Domain controllers store data and manage user and domain interactions, including user logon processes, authentication, and directory searches. If you plan to use this server to provide the Active Directory directory service to network users and computers, configure this server as a domain controller.

To configure a server as a domain controller, install Active Directory on the server. There are four options available in the Active Directory Installation Wizard. You can create an additional domain controller in an existing domain, a domain controller for a new child domain, a domain controller for a new domain tree, or a domain controller for a new forest. If you are not sure which role you need, read about each role by clicking the role option.

### Notes

- If you have already installed a domain controller role and you want to view next steps, in the list below, click the domain controller configuration that you installed, and then click **Next steps: Completing additional tasks**.

- If you need to reconfigure your server for a different role, you can remove existing server roles. By removing the domain controller role, you will uninstall Active Directory from this server. After Active Directory has been uninstalled, this server will no longer participate in replication of directory objects and domain-based user authentication requests. For more information, see the sections below.

Click the type of domain controller role that you want to create:

Creating an additional domain controller for an existing domain

Creating a domain controller for a new forest

Creating a domain controller for a new child domain

Creating a domain controller for a new domain tree

**Creating an additional domain controller for an existing domain**
Create additional domain controllers when you want to improve the availability and reliability of network services. By adding additional domain controllers, you can provide fault tolerance, balance the load of existing domain controllers, provide additional infrastructure support to sites, and improve performance by making it easier for clients to connect to a domain controller when they log on to the network. For example, as seen in the following illustration, by adding a new domain controller (DC2) to the microsoft.com domain, it helps offset the load on other domain controllers.




This topic explains the basic steps that you must follow to create an additional domain controller in your organization.

This process involves using the **Configure Your Server Wizard** and the **Active Directory Installation Wizard** to install Active Directory on this server. When you have finished setting up your domain controller, you can complete additional configuration tasks.

This topic covers:

Before you begin

Configuring your domain controller

Next steps: Completing additional tasks

**Before you begin**

Before you configure your server as a domain controller, verify whether or not:

- TCP/IP configuration settings for the server are correct, particularly those used for DNS name resolution. For more information, see Configure TCP/IP to use DNS.
- All existing disk volumes use the NTFS file system. Active Directory requires at least one NTFS volume in which to store the SYSVOL folder and its contents. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

The following table lists the information that you need to know before you add a domain controller.

| **Before adding a domain controller role** | **Comments** |
|---|---|
| Determine which sites require a domain controller. | If your network is divided into sites, it is good practice to put at least one domain controller in each site to enhance network performance. When users log on to the network, a domain controller must be contacted as part of the logon process. If clients have to connect to a domain controller located in a different site, the logon process can take a long time. |
| Determine whether to add an additional domain controller over the network or through backup media taken from an existing domain controller. | With the Windows Server 2003 family, you can install Active Directory on member servers using a restored backup taken from a domain controller running Windows Server 2003. You can store this backup on any backup media (for example, tape, CD, or DVD) or a shared network resource. By using restored backup files to create an additional domain controller, you greatly reduce the network bandwidth used when you install Active Directory over a shared network resource. You still need network connectivity to replicate all new objects and recent changes for existing objects to the new domain controller. For more information about creating an additional domain controller from backup media, see Creating an additional domain controller. |
| Determine whether you want your new domain controller to host a global catalog. | A global catalog stores a full copy of all objects in the directory for its host domain and a partial copy of all objects for all other domains in the forest. To optimize network performance in a multiple-site environment, consider adding global catalogs for select sites. In a single-site environment, a single global catalog is usually sufficient to cover common Active Directory queries. However, in a multiple-site environment it is recommended that you use global catalogs in each site. For more information about when to add global catalogs in a multiple site environment, see Global catalogs and sites. |
| Domain controllers running Windows 2000 or Windows Server 2003 are available. | The domain must have at least one other domain controller running Windows 2000 or Windows Server 2003 to add an additional domain controller. Active Directory domain controllers cannot be configured as backup domain controllers (BDCs) for Windows NT domains. |

| | |
|---|---|
| Obtain the administrative credentials necessary to add a domain controller. | To add an additional domain controller to an existing domain, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. |
| Identify the DNS domain name of the Active Directory domain to which you want to add the additional domain controller. | You need to provide the DNS domain name when you use the Active Directory Installation Wizard. |

**Configuring your domain controller**

To configure a domain controller, start the **Configure Your Server Wizard** by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard.**

On the **Server Role** page, click **Domain Controller (Active Directory)**, and then click **Next**.

This section describes each of the steps in this process and outlines the required choices and decisions you will make as you configure your domain controller.

The following sections cover these configuration steps:

Summary of Selections

Using the Active Directory Installation Wizard

Completing the Configure Your Server Wizard

Removing the domain controller role

**Summary of Selections**

On the **Summary of Selections** page of the Configure Your Server Wizard, you can view and confirm the options that you have selected. If you selected **Domain Controller (Active Directory)** on the previous page, the following appears:

- **Run the Active Directory Installation Wizard to configure this server as a domain controller**

To apply the selections shown on the **Summary of Selections** page, click **Next**.

**Using the Active Directory Installation Wizard**

After you click **Next**, the Active Directory Installation Wizard starts automatically. If this is the first time you have installed Active Directory on a server, on the **Welcome** page of the Active Directory Installation Wizard, click **Active Directory Help**for more information about Active Directory.

After you finish reading about Active Directory, click **Next**. You can return to this page from any place in the wizard until you click **Finish** on the last page. On the **Operating System Compatibility** page, read the information and then click **Next**. If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.

This section describes the following steps in the Active Directory Installation Wizard:

Domain Controller Type

Network Credentials

Additional Domain Controller

Database and Log Folders

Shared System Volume

Directory Services Restore Mode Administrator Password

Summary

**Domain Controller Type**
On the **Domain Controller Type** page, click **Domain controller for a new domain**.

After you finish, click **Next**.

**Network Credentials**
On the **Network Credentials** page, type the user name, password, and user domain of the user account that you want to use.

| Option | Comments |
|---|---|
| User name | Type the name of a user account that has the necessary administrative credentials. The user account must be a member of the Domain Admins group (in the forest root domain), or a member of the Enterprise Admins group, or must have been delegated the appropriate authority |
| Password | Type the password for the user account. This should always be a strong password. For more information, see Strong passwords. |
| Domain | Type the full DNS name of the domain in which this user name and password are valid. |

**Additional Domain Controller**
On the **Additional Domain Controller** page, type the full DNS name of the domain to which you want to add this domain controller. A full DNS name is also referred to as a fully qualified domain name (FQDN). Active Directory domains are named with DNS names and follow the same hierarchical structure of DNS.

After you finish, click **Next**.

**Database and Log Folders**
On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location. To avoid any problems with installing or removing Active Directory, it is important to confirm that you have sufficient disk space to host the directory database and log files. The Active Directory Installation Wizard requires 250 megabytes (MB) of disk space for the Active Directory database and 50 MB for the log files. It is recommended that you store these files on an NTFS partition.

After you finish, click **Next**.

**Shared System Volume**
On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location. The Sysvol folder must be stored on an NTFS volume since it contains files that are replicated between domain controllers in a domain or forest. These files include scripts, Windows NT 4.0 and earlier system policies, the NETLOGON and SYSVOL shares, and Group Policy settings.

After you finish, click **Next**.

**Directory Services Restore Mode Administrator Password**
On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password that you want to assign to the restore mode Administrator account for the server. You should use strong passwords for directory restore mode passwords. For more information, see Strong passwords.

**Important**

- You must know this password to restore a backup copy of the System State for this domain controller.

You use this password when the domain controller starts in Directory Services Restore Mode. If this is the first time you have installed Active Directory on a server, click **Active Directory Help** for more information about the restore mode password.

After you finish, click **Next**.

**Summary**
On the **Summary** page, review the information, and then click **Next**.

After you complete the installation, click **Finish**. To restart your computer and implement the changes, click **Restart Now**.

**Completing the Configure Your Server Wizard**
After your server restarts, the Configure Your Server Wizard displays the **This Server is Now a Domain Controller** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at systemroot\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

After you complete the Configure Your Server Wizard, be sure to visit Windows Update to download any additional updates that are available. For more information, see Windows Update.

In addition, you should also run the Security Configuration Wizard to further secure your domain controller. You can run Windows Firewall on a domain controller, but only if you run the Security Configuration Wizard to install and configure it. For more information see Security Configuration Wizard.

**Removing the domain controller role**
If you need to reconfigure your server for a different role, you can remove existing server roles. By removing the domain controller role, you will uninstall Active Directory from this server. After Active Directory has been uninstalled, this server will no longer participate in replication of directory objects and domain-based user authentication requests.

To remove the domain controller role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Domain Controller (Active Directory)**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the domain controller role** check box, click **Next**, and then follow the steps in the Active Directory Installation Wizard.

---

**Next steps: Completing additional tasks**
After you complete the Active Directory Installation Wizard, the server is configured as a domain controller. You can use it to store data, manage objects, and provide information to users, computers and applications. Up to this point, you have created a new domain or additional domain controller for an existing domain.

The following table lists some of the additional tasks that you might want to perform on your domain controller.

| Task | Purpose of task | Reference |
|---|---|---|
| Secure the new domain controller in a locked room. | To ensure that no one can physically access the domain controller. | Domain controllers; Securing Active Directory |
| Use strong encryption techniques. | To secure account password information stored on the new domain controller. | The system key utility |

**Creating a domain controller for a new forest**
Create a domain controller for a new forest when you want to upgrade a Windows NT domain to become the first domain in a new forest, segment your network for administrative autonomy, provide a boundary to administer data, isolate the scope of directory replication, or use a noncontiguous DNS namespace that is different from an existing forest on your network. For example, as seen in the following illustration, the microsoft.com forest is the first Active Directory domain in an organization.




This topic explains the basic steps that you must follow to configure a domain controller for a new forest in your organization.

This process involves using the Configure Your Server Wizard and the Active Directory Installation Wizard to install Active Directory on this server. When you have finished setting up your domain controller, you can complete additional configuration tasks.

This topic covers:

Before you begin

Configuring your domain controller

Next steps: Completing additional tasks

**Before you begin**

Before you configure your server as a domain controller, verify whether or not:

- TCP/IP configuration settings for the server are correct, particularly those used for DNS name resolution. For more information, see Configure TCP/IP to use DNS.
- All existing disk volumes use the NTFS file system. Active Directory requires at least one NTFS volume in which to store the SYSVOL folder and its contents. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

The following table lists the information that you need to know before you add a domain controller for a new forest.

| **Before adding a new domain controller role for a new forest** | **Comments** |
|---|---|
| Verify that DNS is properly configured for your organization. | You need to confirm that DNS is properly configured on your network and that it supports dynamic updates and service (SRV) resource records. If you are setting up Active Directory for the first time in your organization, and you do not currently have a DNS infrastructure configured, the Active Directory Installation Wizard sets up and configures DNS on this server during the Active Directory installation process. Active Directory requires DNS to function and share the same hierarchical domain structure. For example, microsoft.com is a DNS domain and an Active Directory domain. |
| Obtain the administrative credentials necessary to create a forest. | To create a new forest, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. |

**Configuring your domain controller**
To configure a domain controller, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Domain Controller (Active Directory)**, and then click **Next**.

This section describes each of the steps in this process and outlines the required choices and decisions you will make as you configure your domain controller. The following sections cover these configuration steps:

Summary of Selections

Using the Active Directory Installation Wizard

Completing the Configure Your Server Wizard

Removing the domain controller role

**Summary of Selections**
On the **Summary of Selections** page of the Configure Your Server Wizard, you can view and confirm the options that you have selected. If you selected **Domain Controller (Active Directory)** on the previous page, the following appears:

- **Run the Active Directory Installation Wizard to configure this server as a domain controller**

To apply the selections shown on the **Summary of Selections** page, click **Next**.

**Using the Active Directory Installation Wizard**
After you click **Next**, the Active Directory Installation Wizard starts automatically. If this is the first time you have installed Active Directory on a server, on the **Welcome** page of the Active Directory Installation Wizard click **Active Directory Help**for more information about Active Directory.

After you finish reading about Active Directory, click **Next**. You can return to this page from any place in the wizard until you click **Finish** on the last page. On the **Operating System Compatibility** page, read the information and then click **Next**. If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.

This section describes the following steps in the Active Directory Installation Wizard:

Domain Controller Type

Create New Domain

New Domain Name

NetBIOS Domain Name

Database and Log Folders

Shared System Volume

DNS Registration Diagnostics

Permissions

Directory Services Restore Mode Administrator Password

Summary

**Domain Controller Type**
On the **Domain Controller Type** page, click **Domain controller for a new domain**.

After you finish, click **Next**.

**Create New Domain**
On the **Create New Domain** page, click **Domain in a new forest**.

After you finish, click **Next**.

**New Domain Name**
On the **New Domain Name** page, type the full DNS name for the new domain. Provide a full DNS name for the new Active Directory forest that you are about to create (for example, headquarters.example.microsoft.com). A full DNS name is also referred to as a fully qualified domain name (FQDN). Active Directory domains are named with DNS names and follow the same hierarchical structure of DNS. When choosing DNS names to use for your Active Directory forest, start with the registered DNS domain suffix that your organization has reserved for use on the Internet, such as microsoft.com.

Although Dcpromo.exe in Windows Server 2008 and Windows Server 2003 allows you to create a single-label DNS domain name, you should not use a single-label DNS name for a domain for several reasons. In Windows Server 2008 R2, Dcpromo.exe does not allow you to create a single-label DNS name for a domain. For more information, see http://go.microsoft.com/fwlink/?LinkId=92467.

After you finish, click **Next**.

**NetBIOS Domain Name**
On the **NetBIOS Domain Name** page, verify the NetBIOS name. Although Active Directory domains are named according to DNS naming standards, you still need to define a NetBIOS name when you create Active Directory domains. NetBIOS names should match the first label of the DNS domain name whenever possible. When the Active Directory domain has a first label DNS name that is different from its NetBIOS name, the FQDN is constructed using the DNS domain name, not the NetBIOS name. For example, if the first label in the full DNS domain name is "child" (child.microsoft.com is the FQDN) and the NetBIOS domain name is "sales", the FQDN remains "child.microsoft.com".

After you finish, click **Next**.

**Database and Log Folders**
On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location. To avoid any problems with installing or removing Active Directory, it is important to confirm that you have sufficient disk space to host the directory database and log files. The Active Directory Installation Wizard requires 250 megabytes (MB) of disk space for the Active Directory database and 50 MB for the log files. It is recommended that you store these files on an NTFS partition.

After you finish, click **Next**.

**Shared System Volume**
On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location. The Sysvol folder must be stored on an NTFS volume since it contains files that are replicated between domain controllers in a domain or forest. These files include scripts, Windows NT 4.0 and earlier system policies, the NETLOGON and SYSVOL shares, and Group Policy settings.

After you finish, click **Next**.

**DNS Registration Diagnostics**
On the **DNS Registration Diagnostics** page, verify that the DNS settings are correct.

If a diagnostic error appears under **Diagnostic Results**, click **Help** for more information about how to resolve the error.

After you finish, click **Next**.

**Permissions**
On the **Permissions** page, click the level of application compatibility that you want with pre-Windows 2000, Windows 2000, or Windows Server 2003 operating systems.

On servers running Windows NT 4.0 and earlier, read access for user and group information is assigned to anonymous users so that existing applications, including Microsoft BackOffice, SQL Server, and some non-Microsoft applications, function correctly. In Windows 2000 and the Windows Server 2003 family, members of the Anonymous Logon group have read access to this information only when the group is added to the Pre-Windows 2000 Compatible Access group.

| Option | Comments |
|---|---|

| | |
|---|---|
| Permissions compatible with pre-Windows 2000 server operating systems | Click this option if you want the Anonymous Logon group and the Everyone security groups to be added to the Pre-Windows 2000 Compatible Access group. |
| Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems | Click this option to prevent members of the Anonymous Logon group from gaining read access to user and group information. |

After you select one of these options, you can manually switch between the backward compatible and high-security settings on Active Directory objects. To do this, open Active Directory Users and Computers, and then add the Anonymous Logon security group to the pre-Windows 2000 Compatible Access security group.

After you finish, click **Next**.

**Directory Services Restore Mode Administrator Password**
On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password that you want to assign to the restore mode Administrator account for the server. You should use strong passwords for directory restore mode passwords. For more information, see Strong passwords.

**Important**

- You must know this password to restore a backup copy of the System State for this domain controller.

You use this password when the domain controller starts in Directory Services Restore Mode. If this is the first time you have installed Active Directory on a server, click Active Directory Help for more information about the restore mode password.

After you finish, click **Next**.

**Summary**
On the **Summary** page, review the information, and then click **Next**.

After you complete the installation, click **Finish**. To restart your computer and implement the changes, click **Restart Now**.

**Completing the Configure Your Server Wizard**
After your server restarts, the Configure Your Server Wizard displays the **This Server is Now a Domain Controller** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at systemroot\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

After you complete the Configure Your Server Wizard, be sure to visit Windows Update to download any additional updates that are available. For more information, see Windows Update.

In addition, you should also run the Security Configuration Wizard to further secure your domain controller. You can run Windows Firewall on a domain controller, but only if you run the Security Configuration Wizard to install and configure it. For more information see Security Configuration Wizard.

**Removing the domain controller role**
If you need to reconfigure your server for a different role, you can remove existing server roles. By removing the domain controller role, you will uninstall Active Directory from this server. After Active Directory has been uninstalled, this server will no longer participate in replication of directory objects and domain-based user authentication requests.

To remove the domain controller role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Domain Controller (Active Directory)**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the domain controller role** check box, click **Next**, and then follow the steps in the Active Directory Installation Wizard.

**Next steps: Completing additional tasks**
After you complete the Active Directory Installation Wizard, the server is configured as a domain controller. You can use it to store data, manage objects, and provide information to users, computers and applications. Up to this point, you have created a domain controller for a new forest.

The following table lists some of the additional tasks that you might want to perform on your domain controller.

| Task | Purpose of task | Reference |
|---|---|---|
| Secure the new domain controller in a locked room. | To ensure that no one can physically access the domain controller. | Domain controllers; Securing Active Directory |
| Use strong encryption techniques. | To secure account password information stored on the new domain controller. | The system key utility |
| Verify and authenticate the validity of each user. | To enhance forest-wide security by using public key cryptography. | Public Key Infrastructure |
| Require all domain users to use strong passwords. | To prevent unauthorized access to your organization. | Strong passwords |
| Enable audit policy. | To receive notification of actions that could pose a security risk. | Auditing Policy |
| Enforce account lockouts on user accounts. | To decrease the possibility of an attacker compromising your domain through repeated logon attempts. | User and computer accounts |
| Enforce password history on user accounts. | To decrease the possibility of an attacker compromising your domain. | Enforce password history |
| Enforce minimum and maximum password ages on user accounts. | To decrease the possibility of an attacker compromising your domain. | Minimum password age; Maximum password age |
| Implement SID Filtering. | To prevent attacks from malicious users who might try to grant elevated user rights to another user account. | "Using Security Identifier (SID) Filtering to Prevent Elevation of Privilege Attacks" at the Microsoft Web site. (http://www.microsoft.com/) |
| Implement smart cards. | To provide tamper-resistant user authentication and e-mail security. | Smart cards overview |
| Restrict user, group, and computer access to shared resources and filter Group Policy settings. | To secure resources. | Group types |
| Create forest trusts (as appropriate) | To manage the security relationship between two forests, and simplify security administration and authentication across forests. | Forest trusts |
| Assign user rights to new security groups. | To specifically define the administrative role of members in the domain. | Group types |

**Creating a domain controller for a new child domain**
Create a domain controller for a new child domain when you want to create a domain that shares a contiguous namespace with one or more existing domains. This means that the name of the new domain contains the full name of the parent domain. For example, as seen in the following illustration, child.microsoft.com is a child domain of microsoft.com. As a best practice, create new domains as children of the forest root domain.

microsoft.com forest




This topic explains the basic steps that you must follow to configure a domain controller for a new child domain in your organization.

This process involves using the Configure Your Server Wizard and the Active Directory Installation Wizard to install Active Directory on this server. When you have finished setting up your domain controller, you can complete additional configuration tasks.

This topic covers:

Before you begin

Configuring your domain controller

Next steps: Completing additional tasks

**Before you begin**

Before you configure your server as a domain controller, verify whether or not:

- TCP/IP configuration settings for the server are correct, particularly those used for DNS name resolution. For more information, see Configure TCP/IP to use DNS.
- All existing disk volumes use the NTFS file system. Active Directory requires at least one NTFS volume in which to store the SYSVOL folder and its contents. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

The following table lists the information that you need to know before you add a domain controller.

| Before adding a domain controller | Comments |
|---|---|

| | |
|---|---|
| Identify the DNS domain name of the Active Directory domain to which you want to add the domain controller. | You need to provide the DNS domain name for the parent domain of this new child domain. |
| Verify that the network speed is adequate when installing Active Directory. | The server on which you want to install Active Directory and create a child domain should have access to the network over a high-speed connection. |
| Determine which sites require a domain controller. | If your network is divided into sites, it is good practice to put at least one domain controller in each site to enhance network performance. When users log on to the network, a domain controller must be contacted as part of the logon process. If clients have to connect to a domain controller located in a different site, the logon process can take a long time. |
| Determine whether you want your new domain controller to host a global catalog. | A global catalog stores a copy of all Active Directory objects in a forest on a domain controller. The global catalog stores a full copy of all objects in the directory for its host domain and a partial copy of all objects for all other domains in the forest. To optimize network performance in a multiple-site environment, consider adding global catalogs for select sites. In a single-site environment, a single global catalog is usually sufficient to cover common Active Directory queries. However, in a multiple-site environment it is recommended that you use global catalogs in each site. For more information about when to add global catalogs in a multiple site environment, see Global catalogs and sites. |
| Obtain the administrative credentials necessary to create a child domain. | To add a new child domain, you must be a member of the Domain Admins group (in the parent domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. |

**Configuring your domain controller**

To configure a domain controller, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Domain Controller (Active Directory)**, and then click **Next**.

This section describes each of the steps in this process and outlines the required choices and decisions you will make as you configure your domain controller.

The following sections cover these configuration steps:

Summary of Selections

Using the Active Directory Installation Wizard

Completing the Configure Your Server Wizard

Removing the domain controller role

**Summary of Selections**

On the **Summary of Selections** page of the Configure Your Server Wizard, you can view and confirm the options that you have selected. If you selected **Domain Controller (Active Directory)** on the previous page, the following appears:

- **Run the Active Directory Installation Wizard to configure this server as a domain controller**

To apply the selections shown on the **Summary of Selections** page, click **Next**.

**Using the Active Directory Installation Wizard**
After you click **Next**, the Active Directory Installation Wizard starts automatically. If this is the first time you have installed Active Directory on a server, on the **Welcome** page of the Active Directory Installation Wizard, click **Active Directory Help**for more information about Active Directory.

After you finish reading about Active Directory, click **Next**. You can return to this page from any place in the wizard until you click **Finish** on the last page. On the **Operating System Compatibility** page, read the information and then click **Next**. If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.

This section describes the following steps in the Active Directory Installation Wizard:

Domain Controller Type

Create New Domain

Network Credentials

Child Domain Installation

NetBIOS Domain Name

Database and Log Folders

Shared System Volume

DNS Registration Diagnostics

Permissions

Directory Services Restore Mode Administrator Password

Summary

**Domain Controller Type**
On the **Domain Controller Type** page, click **Domain controller for a new domain**.

After you finish, click **Next**.

**Create New Domain**
On the **Create New Domain** page, click **Child domain in an existing domain tree**.

After you finish, click **Next**.

**Network Credentials**
On the **Network Credentials** page, type the user name, password, and domain of the user account that you want to use.

| Option | Comments |
|---|---|
| User name | Type the name of a user account that has the necessary administrative credentials. The user account must be a member of the Domain Admins group (in the parent domain), or a member of the Enterprise Admins group, or must have been delegated the appropriate authority. |
| Password | Type the password for the user account. This should always be a strong password. For more information, see Strong passwords. |
| Domain | Type the full DNS name of the domain in which this user name and password are valid. |

After you finish, click **Next**.

**Child Domain Installation**
On the **Child Domain Installation** page, verify the parent domain, and then the name for the new child domain. Active Directory domains are named with DNS names and follow the same hierarchical structure of DNS. When choosing DNS names to use for your child domain, consider names with either geographical or divisional scope within your organization (for

example, newyork.microsoft.com or sales.microsoft.com).

Following the DNS standards, domain names are fully qualified domain names (FQDN), which consist of the name of the domain appended to the names of the parent domains and root domains using the dot (.) character format. For example, an Active Directory domain named "child" and a parent domain named "microsoft.com" have an FQDN (also referred to as a complete DNS name) of "child.microsoft.com".

After you finish, click **Next**.

**NetBIOS Domain Name**

On the **NetBIOS Domain Name** page, verify the NetBIOS name. Although Active Directory domains are named according to DNS naming standards, you still need to define a NetBIOS name when you create Active Directory domains. NetBIOS names should match the first label of the DNS domain name whenever possible. When the Active Directory domain has a first label DNS name that is different from its NetBIOS name, the FQDN is constructed using the DNS domain name, not the NetBIOS name. For example, if the first label in the full DNS domain name is "child" (child.microsoft.com is the FQDN) and the NetBIOS domain name is "sales", the FQDN remains "child.microsoft.com".

**Database and Log Folders**

On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location, and then click Next. To avoid any problems with installing or removing Active Directory, it is important to confirm that you have sufficient disk space to host the directory database and log files. The Active Directory Installation Wizard requires 250 megabytes (MB) of disk space for the Active Directory database and 50 MB for the log files. It is recommended that you store these files on an NTFS partition.

After you finish, click **Next**.

**Shared System Volume**

On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location. The Sysvol folder must be stored on an NTFS volume since it contains files that are replicated between domain controllers in a domain or forest. These files include scripts, Windows NT 4.0 and earlier system policies, the NETLOGON and SYSVOL shares, and Group Policy settings.

After you finish, click **Next**.

**DNS Registration Diagnostics**

On the **DNS Registration Diagnostics** page, verify that the DNS settings are correct.

If a diagnostic error appears under **Diagnostic Results**, click **Help** for more information about how to resolve the error.

After you finish, click **Next**.

**Permissions**

On the **Permissions** page, click the level of application compatibility that you want with pre-Windows 2000, Windows 2000, or Windows Server 2003 operating systems.

On servers running Windows NT 4.0 and earlier, read access for user and group information is assigned to anonymous users so that existing applications, including Microsoft BackOffice, SQL Server, and some non-Microsoft applications, function correctly. In Windows 2000 and the Windows Server 2003 family, members of the Anonymous Logon group have read access to this information only when the group is added to the Pre-Windows 2000 Compatible Access group.

| Option | Comments |
|---|---|
| Permissions compatible with pre-Windows 2000 server operating systems | Click this option if you want the Anonymous Logon group and the Everyone security groups to be added to the Pre-Windows 2000 Compatible Access group. |
| Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems | Click this option to prevent members of the Anonymous Logon group from gaining read access to user and group information. |

After you select one of these options, you can manually switch between the backward compatible and high-security settings on Active Directory objects. To do this, open Active Directory Users and Computers, and then add the Anonymous Logon security group to the pre-Windows 2000 Compatible Access security group.

After you finish, click **Next**.

**Directory Services Restore Mode Administrator Password**

On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password that you want to assign to the restore mode Administrator account for the server. You should use strong passwords for directory restore mode passwords. For more information, see Strong passwords.

**Important**

- You must know this password to restore a backup copy of the System State for this domain controller.

You use this password when the domain controller starts in Directory Services Restore Mode. If this is the first time you have installed Active Directory on a server, click **Active Directory Help** for more information about the restore mode password.

After you finish, click **Next**.

**Summary**
On the **Summary** page, review the information, and then click **Next**.

After you complete the installation, click **Finish**. To restart your computer and implement the changes, click **Restart Now**.

**Completing the Configure Your Server Wizard**
After your server restarts, the Configure Your Server Wizard displays the **This Server is Now a Domain Controller** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at systemroot\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

After you complete the Configure Your Server Wizard, be sure to visit Windows Update to download any additional updates that are available. For more information, see Windows Update.

In addition, you should also run the Security Configuration Wizard to further secure your domain controller. You can run Windows Firewall on a domain controller, but only if you run the Security Configuration Wizard to install and configure it. For more information see Security Configuration Wizard.

**Removing the domain controller role**
If you need to reconfigure your server for a different role, you can remove existing server roles. By removing the domain controller role, you will uninstall Active Directory from this server. After Active Directory has been uninstalled, this server will no longer participate in replication of directory objects and domain-based user authentication requests.

To remove the domain controller role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Domain Controller (Active Directory)**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the domain controller role** check box, click **Next**, and then follow the steps in the Active Directory Installation Wizard.

**Next steps: Completing additional tasks**
After you complete the Active Directory Installation Wizard, the server is configured as a domain controller. You can use it to store data, manage objects, and provide information to users, computers and applications. Up to this point, you have created a new domain or additional domain controller for an existing domain.

The following table lists some of the additional tasks that you might want to perform on your domain controller.

| Task | Purpose of task | Reference |
|---|---|---|
| Secure the new domain controller in a locked room. | To ensure that no one can physically access the domain controller. | Domain controllers; Securing Active Directory |

| | | |
|---|---|---|
| Use strong encryption techniques. | To secure account password information stored on the new domain controller. | The system key utility |
| Require all domain users to use strong passwords. | To prevent unauthorized access to your organization. | Strong passwords |
| Enable audit policy. | To receive notification of actions that could pose a security risk. | Auditing Policy |
| Enforce account lockouts on user accounts. | To decrease the possibility of an attacker compromising your domain through repeated logon attempts. | User and computer accounts |
| Enforce password history on user accounts. | To decrease the possibility of an attacker compromising your domain. | Enforce password history |
| Enforce minimum and maximum password ages on user accounts. | To decrease the possibility of an attacker compromising your domain. | Minimum password age; Maximum password age |
| Implement SID Filtering. | To prevent attacks from malicious users who might try to grant elevated user rights to another user account. | "Using Security Identifier (SID) Filtering to Prevent Elevation of Privilege Attacks" at the Microsoft Web site.(http://www.microsoft.com/) |
| Implement smart cards. | To provide tamper-resistant user authentication and e-mail security. | Smart cards overview |
| Restrict user, group, and computer access to shared resources and filter Group Policy settings. | To secure resources. | Group types |
| Assign user rights to new security groups. | To specifically define the administrative role of members in the domain. | Group types |

**Creating a domain controller for a new domain tree**

Create a domain controller for a new domain tree when you want to create a domain that has a DNS namespace that is not related to the other domains in the forest. This means that the name of the tree root domain (and all of its children) does not have to contain the full name of the parent domain. A forest can contain one or more domain trees. For example, as seen in the following illustration, msn.com is a new domain tree in the microsoft.com forest.




This topic explains the basic steps that you must follow to configure a domain controller for a new domain tree in your organization.

This process involves using the Configure Your Server Wizard and the Active Directory Installation Wizard to install Active

Directory on this server. When you have finished setting up your domain controller, you can complete additional configuration tasks.

This topic covers:

Before you begin

Configuring your domain controller

Next steps: Completing additional tasks

**Before you begin**

Before you configure your server as a domain controller, verify whether or not:

- TCP/IP configuration settings for the server are correct, particularly those used for DNS name resolution. For more information, see Configure TCP/IP to use DNS.
- All existing disk volumes use the NTFS file system. Active Directory requires at least one NTFS volume in which to store the SYSVOL folder and its contents. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.
- The following table lists the information that you need to know before you add a domain controller.

| **Before adding a domain controller role** | **Comments** |
|---|---|
| Verify that the network speed is adequate when installing Active Directory. | The server on which you want to install Active Directory should have access to the network over a high-speed connection. |
| Determine whether you want your new domain controller to host a global catalog. | A global catalog stores a copy of all Active Directory objects in a forest on a domain controller. The global catalog stores a full copy of all objects in the directory for its host domain and a partial copy of all objects for all other domains in the forest. To optimize network performance in a multiple-site environment, consider adding global catalogs for select sites. In a single-site environment, a single global catalog is usually sufficient to cover common Active Directory queries. However, in a multiple-site environment it is recommended that you use global catalogs in each site. For more information about when to add global catalogs in a multiple site environment, see Global catalogs and sites. |
| Obtain the administrative credentials necessary to add a new domain tree. | To create a new domain tree, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. |

**Configuring your domain controller**

To configure a domain controller, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Domain Controller (Active Directory)**, and then click **Next**.

This section describes each of the steps in this process and outlines the required choices and decisions you will make as you configure your domain controller. The following sections cover these configuration steps:

Summary of Selections

Using the Active Directory Installation Wizard

Completing the Configure Your Server Wizard

Removing the domain controller role

**Summary of Selections**
On the Summary of Selections page of the Configure Your Server Wizard, you can view and confirm the options that you have selected. If you selected **Domain Controller (Active Directory)** on the previous page, the following appears:

- **Run the Active Directory Installation Wizard to configure this server as a domain controller**

To apply the selections shown on the **Summary of Selections** page, click **Next**.

**Using the Active Directory Installation Wizard**
After you click **Next**, the Active Directory Installation Wizard starts automatically. If this is the first time you have installed Active Directory on a server, on the **Welcome** page of the Active Directory Installation Wizard, click **Active Directory Help**for more information about Active Directory.

After you finish reading about Active Directory, click **Next**. You can return to this page from any place in the wizard until you click **Finish** on the last page. On the **Operating System Compatibility** page, read the information and then click **Next**. If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.

This section describes the following steps in the Active Directory Installation Wizard:

Domain Controller Type

Create New Domain

Network Credentials

New Domain Tree

NetBIOS Domain Name

Database and Log Folders

Shared System Volume

DNS Registration Diagnostics

Permissions

Directory Services Restore Mode Administrator Password

Summary

**Domain Controller Type**
On the **Domain Controller Type** page, click Domain controller for a new domain.

After you finish, click **Next**.

**Create New Domain**
On the **Create New Domain** page, click **Domain tree in an existing forest**.

After you finish, click **Next**.

**Network Credentials**
On the **Network Credentials** page, type the user name, password, and user domain of the user account that you want to use.

| Option | Comments |
|---|---|
| User name | Type the name of a user account that has the necessary administrative credentials. The user account must be a member of the Domain Admins group (in the forest root domain), or a member of the Enterprise Admins group, or must have been delegated the appropriate authority. |
| Password | Type the password for the user account. This should always be a strong password. For more information, see Strong passwords. |
| Domain | Type the full DNS name of the domain in which this user name and password are valid. |

**New Domain Name**
On the **New Domain Name** page, type the full DNS name for the new domain tree. Provide a full DNS name for the new domain tree that you want to create. Active Directory domains are named with DNS names and follow the same hierarchical structure of DNS.

Following the DNS standards, domain names are fully qualified domain names (FQDN), which consist of the name of the domain appended to the names of the parent domains and root domains using the dot (.) character format.

Although Dcpromo.exe in Windows Server 2008 and Windows Server 2003 allows you to create a single-label DNS domain name, you should not use a single-label DNS name for a domain for several reasons. In Windows Server 2008 R2, Dcpromo.exe does not allow you to create a single-label DNS name for a domain. For more information, see http://go.microsoft.com/fwlink/?LinkId=92467.

Although Dcpromo.exe in Windows Server 2008 and Windows Server 2003 allows you to create a single-label DNS domain name, you should not use a single-label DNS name for a domain for several reasons. In Windows Server 2008 R2, Dcpromo.exe does not allow you to create a single-label DNS name for a domain. For more information, see http://go.microsoft.com/fwlink/?LinkId=92467.

After you finish, click **Next**.

**NetBIOS Domain Name**
On the **NetBIOS Domain Name** page, verify the NetBIOS name. Although Active Directory domains are named according to DNS naming standards, you still need to define a NetBIOS name when you create Active Directory domains. NetBIOS names should match the first label of the DNS domain name whenever possible. When the Active Directory domain has a first label DNS name that is different from its NetBIOS name, the FQDN is constructed using the DNS domain name, not the NetBIOS name. For example, if the first label in the full DNS domain name is "child" (child.microsoft.com is the FQDN) and the NetBIOS domain name is "sales", the FQDN remains "child.microsoft.com".

After you finish, click **Next**.

**Database and Log Folders**
On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location. To avoid any problems with installing or removing Active Directory, it is important to confirm that you have sufficient disk space to host the directory database and log files. The Active Directory Installation Wizard requires 250 megabytes (MB) of disk space for the Active Directory database and 50 MB for the log files. It is recommended that you store these files on an NTFS partition.

After you finish, click **Next**.

**Shared System Volume**
On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location. The Sysvol folder must be stored on an NTFS volume since it contains files that are replicated between domain controllers in a domain or forest. These files include scripts, Windows NT 4.0 and earlier system policies, the NETLOGON and SYSVOL shares, and Group Policy settings.

After you finish, click **Next**.

**DNS Registration Diagnostics**

On the **DNS Registration Diagnostics** page, verify the DNS settings are correct.

If a diagnostic error appears under **Diagnostic Results**, click **Help** for more information about how to resolve the error.

After you finish, click **Next**.

**Permissions**
On the **Permissions** page, click the level of application compatibility that you want with pre-Windows 2000 and Windows 2000 or Windows Server 2003 operating systems.

On servers running Windows NT 4.0 and earlier, read access for user and group information is assigned to anonymous users so that existing applications, including Microsoft BackOffice, SQL Server, and some non-Microsoft applications, function correctly. In Windows 2000 and the Windows Server 2003 family, members of the Anonymous Logon group have read access to this information only when the group is added to the Pre-Windows 2000 Compatible Access group.

| Option | Comments |
|---|---|
| Permissions compatible with pre-Windows 2000 server operating systems | Click this option if you want the Anonymous Logon group and the Everyone security groups to be added to the Pre-Windows 2000 Compatible Access group. |
| Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems | Click this option to prevent members of the Anonymous Logon group from gaining read access to user and group information. |

After you select one of these options, you can manually switch between the backward compatible and high-security settings on Active Directory objects. To do this, open Active Directory Users and Computers, and then add the Anonymous Logon security group to the pre-Windows 2000 Compatible Access security group.

**Directory Services Restore Mode Administrator Password**
On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password that you want to assign to the restore mode Administrator account for the server. You should use strong passwords for directory restore mode passwords. For more information, see Strong passwords.

**Important**

- You must know this password to restore a backup copy of the System State for this domain controller.

You use this password when the domain controller starts in Directory Services Restore Mode. If this is the first time you have installed Active Directory on a server, click **Active Directory Help** for more information about the restore mode password.

After you finish, click **Next**.

**Summary**
On the **Summary** page, review the information, and then click **Next**.

After you complete the installation, click **Finish**. To restart your computer and implement the changes, click **Restart Now**.

**Completing the Configure Your Server Wizard**
After your server restarts, the Configure Your Server Wizard displays the **This Server is Now a Domain Controller** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at systemroot\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

After you complete the Configure Your Server Wizard, be sure to visit Windows Update to download any additional updates that are available. For more information, see Windows Update.

In addition, you should also run the Security Configuration Wizard to further secure your domain controller. You can run Windows Firewall on a domain controller, but only if you run the Security Configuration Wizard to install and configure it. For more information see Security Configuration Wizard.

**Removing the domain controller role**
If you need to reconfigure your server for a different role, you can remove existing server roles. By removing the domain controller role, you will uninstall Active Directory from this server. After Active Directory has been uninstalled, this server will no longer participate in replication of directory objects and domain-based user authentication requests.

To remove the domain controller role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Domain Controller (Active Directory)**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the domain controller role** check box, click **Next**, and then follow the steps in the Active Directory Installation Wizard.

**Next steps: Completing additional tasks**

After you complete the Active Directory Installation Wizard, the server is configured as a domain controller. You can use it to store data, manage objects, and provide information to users, computers and applications. Up to this point, you have created a new domain or additional domain controller for an existing domain.

The following table lists some of the additional tasks that you might want to perform on your domain controller.

| Task | Purpose of task | Reference |
|---|---|---|
| Secure the new domain controller in a locked room. | To ensure that no one can physically access the domain controller. | Domain controllers; Securing Active Directory |
| Use strong encryption techniques. | To secure account password information stored on the new domain controller. | The system key utility |
| Require all domain users to use strong passwords. | To prevent unauthorized access to your organization. | Strong passwords |
| Enable audit policy. | To receive notification of actions that could pose a security risk. | Auditing Policy |
| Enforce account lockouts on user accounts. | To decrease the possibility of an attacker compromising your domain through repeated logon attempts. | User and computer accounts |
| Enforce password history on user accounts. | To decrease the possibility of an attacker compromising your domain. | Enforce password history |
| Enforce minimum and maximum password ages on user accounts. | To decrease the possibility of an attacker compromising your domain. | Minimum password age; Maximum password age |
| Implement SID Filtering. | To prevent attacks from malicious users who might try to grant elevated user rights to another user account. | "Using Security Identifier (SID) Filtering to Prevent Elevation of Privilege Attacks" at the Microsoft Web site.(http://www.microsoft.com/) |
| Implement smart cards. | To provide tamper-resistant user authentication and e-mail security. | Smart cards overview |
| Restrict user, group, and computer access to shared resources and filter Group Policy settings. | To secure resources. | Group types |
| Assign user rights to new security groups. | To specifically define the administrative role of members in the domain. | Group types |

# DNS server role: Configuring a DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS server role: Configuring a DNS server**
Domain Name System (DNS) servers host records of a distributed DNS database and use the records they host to resolve DNS name queries sent by DNS client computers, such as queries for the names of Web sites or computers in your network or on the Internet. If you plan to use this computer to answer DNS queries for computers in your network, then add the DNS server role.

This topic explains the basic steps that you must follow to configure a DNS server for either a small organization or a branch office in a large organization. For both scenarios, this topic explains the basic steps you must follow to configure a DNS server and configure it with a DNS zone for your network. This topic will also cover how to forward queries for external resources to a DNS server either run by an Internet service provider (ISP), as in the case of a small organization, or run by the central office of a large organization, as in the case of the branch office.

This process involves using the Configure Your Server Wizard, the Configure a New DNS Server Wizard, and Manage Your Server to configure the server as a DNS server. When you have finished setting up a DNS server, you can complete additional configuration tasks, depending on how you want to use the DNS server.

**Note**

- A DNS server is not usually necessary in a small business because the Windows Internet Name Service (WINS) name resolution method is used to locate network resources, and resources on the Internet are located using the DNS servers run by an ISP. However, as more networks are becoming integrated with the Internet, DNS is becoming more common in small networks.

  Using DNS in your network does not necessarily require that you administer a DNS infrastructure. If you have a small network in which information is maintained dependably, then you can choose to have your DNS namespace administered by a different organization that specializes in DNS administration, such as your government or an ISP. In this case, the different organization will host and administer your DNS information for you or integrate your computers with an existing DNS server hosted in its network.

This topic covers:

Before you begin

Configuring your DNS server

Next steps: Completing additional tasks

**Before you begin**
Before you configure your computer as a DNS server, verify whether or not:

- The operating system is configured correctly. In the Windows Server 2003 family, the DNS Server service depends on the appropriate configuration of the operating system and its services, such as TCP/IP. If you have a new installation of a Windows Server 2003 operating system, then you can use the default service settings. No further action is necessary. If you upgraded to a Windows Server 2003 operating system, or if you want to verify that your service settings are configured correctly for the best performance and security, then see Default settings for services.

- All available disk space is allocated. You can use Disk Management or DiskPart.exe to create a new partition from unallocated space. For more information, see Create a partition or logical drive.

- All existing disk volumes use the NTFS file system. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions.

- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.

- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see

Security Configuration Wizard Overview.

The following table lists the information that you need to know before you add a DNS server role.

| **For all organizations** | |
|---|---|
| **Before adding a DNS server role** | **Comments** |
| Determine if you are adding the DNS server role to support Active Directory. | If you are going to deploy Active Directory, then the DNS servers used to support Active Directory will be installed and configured automatically by the Active Directory Installation Wizard. For more information, see Typical setup for a first server and Domain controller role: Configuring a domain controller. |
| Inventory the security policies of your network and company to see how they can be maintained when broadcasting DNS data over the Internet. | DNS was originally designed as an open protocol and is therefore vulnerable to attackers. Windows Server 2003 DNS provides features for a very secure DNS infrastructure. To configure DNS to support these security policies, have your company's security policies available when designing and deploying your DNS server, zones, and resource records. For more information, see Security information for DNS. |
| Review the DNS checklist. | Review the information in Checklist: Installing a DNS server. |
| **For a small organization** | |
| **Before adding a DNS server role** | **Comments** |
| Choose the first DNS domain name for your company. | Choosing the first DNS domain name for your company involves selecting a domain name that is unique within the DNS namespace of the Internet.<br><br>If your organization has a Web site, use your existing Web site name as a starting point for your DNS domain name. If the name of your Web site is www.humongousinsurance.com, create the first domain name as an extension of this name using the subdomain name *corp*, for example, corp.humongousinsurance.com. |
| Check with your ISP to determine that your network Internet Protocol (IP) addresses are registered with an Internet registrar. | In order for your DNS deployment to work over the Internet, both the IP addresses and DNS domain name used by your network must be registered with an authorized Internet registrar. These organizations are responsible for assigning IP addresses and DNS domain names and keeping public records of the assignments.<br><br>If you are connected to the Internet, then your company's network is most likely a subnet of your ISP's network. In this case, the IP addresses of the subnet will have been registered with the Internet registrar. |
| Register your DNS domain name with an Internet registrar. | Register the DNS domain name you will use for your company even if you are deploying DNS in a private network. If you do not register the name and later attempt to use it on the Internet, or connect to a network that is connected to the Internet, you might find that the name is unavailable because some other company has registered the name. You can also have your ISP register your DNS domain name for you.<br><br>When deciding on your DNS domain name, search the Internet to see what domain names are available using one of the Web sites provided by an authorized Internet registration authority. To search the Internet for available domain names, see the Internet Network Information Center (InterNIC) Web site. Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here. |
| Understand that the name of the first DNS zone this DNS server will host is the same as the DNS domain name you registered. | When you configure the DNS server role, you will define the first DNS zone that this DNS server will host using the domain name of the DNS domain of your network, such as corp.humongousinsurance.com. |

| | |
|---|---|
| Obtain the IP address of one or more DNS servers hosted by your ISP to use as a forwarder. | You will configure the DNS server with a forwarder to send queries for names that are not in your network to a DNS server at your ISP. |
| **For a branch office in a large organization** | |
| **Before adding a DNS server role** | **Comments** |
| Obtain the DNS domain name for your network from your central office. | The first DNS domain name for your branch office is a subdomain of a domain used in the network at your central office. For example, if your central office uses the domain name corp.humongousinsurance.com, the DNS domain name for your branch office could be seattle.corp.humongousinsurance.com. Always confirm that your DNS domain name has been properly delegated from the central office. |

**Configuring your DNS server**

To configure a DNS server, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Configure Your Server Wizard**.

On the **Server Role** page, click **DNS server**, and then click **Next**.

This section covers:

Summary of Selections

Setting a static DNS server IP address

Using the Configure a DNS Server Wizard

Completing the Configure Your Server Wizard

Removing the DNS server role

**Summary of Selections**

On the **Summary of Selections** page, view and confirm the options that you have selected. If you selected **DNS Server** on the **Server Role** page, the following appears:

- **Install DNS Server**

- **Run the Configure a DNS Server Wizard to configure DNS**

If the **Summary of Selections** page lists these two items, click **Next**. If the **Summary of Selections** page does not list these two items, click **Back** to return to the **Server Role** page, click **DNS server**, and then click **Next**.

**Setting a static DNS server IP address**

After you click **Next**, the Configure Your Server Wizard installs the DNS Server service. During the installation of the DNS Server service, the Configure Your Server Wizard determines whether the IP address for this server is static or is configured automatically. DNS servers are located by DNS clients by using static IP addresses, and an IP address that is automatically configured can cause problems for DNS clients when the IP address changes.

If this server is currently configured to obtain its IP address automatically, the **Configuring Components** page of the Windows Components Wizard appears, and prompts you to configure this server with a static IP address. In the **Local Area Connection Properties** dialog box, click **Internet Protocol (TCP/IP)**, and then click **Properties**. In the **Internet Protocols (TCP/IP) Properties** dialog box, click **Use the following IP address**, and then type the static IP address, subnet mask, and default gateway for this server. In **Preferred DNS server**, type the IP address of this server. In **Alternate DNS server**, type the

IP address of the DNS server hosted by your ISP or central office. When you finish setting up the static addresses for your DNS server, click **OK**, and then click **Close**.

### Notes

- For a small organization, the static IP address for this server will be used to register the DNS domain name for your company with an authorized Internet registrar. The Internet registrar will map the DNS domain name for your company with the IP address so that computers on the Internet looking for computers in your network will know the IP address of the DNS server for your network.

- For a branch office, the static IP address for this server will be used in the domain name delegation configured on a DNS server at the central office of your organization. Computers in your organization and on the Internet looking for computers in your network will use the IP address of the DNS server for your network. For this reason, it is very important that you do not change the IP address of this server after you have added the DNS server role.

#### Using the Configure a DNS Server Wizard

After you click **Close**, the Configure a DNS Server Wizard starts. If you cancel the Configure a DNS Server Wizard, the DNS Server service remains installed, but it cannot resolve names until you configure DNS settings. If you choose to configure DNS settings later, you can do so using the DNS console.

This section describes the following steps in the Configure a DNS Server Wizard:

Select Configuration Action

Primary Server Location

Zone Name

Dynamic Update

Forwarders

Completing the Configure a DNS Server Wizard

#### Select Configuration Action

On the **Select Configuration Action** page, select **Create a forward lookup zone**, and then click **Next**.

#### Primary Server Location

To specify that this DNS server will host a DNS zone that contains DNS resource records for your network resources, on the **Primary Server Location** page, select **This server maintains the zone**, and then click **Next**.

#### Zone Name

On the **Zone Name** page, in **Zone name**, specify the name of the DNS zone for your network, and then click **Next**. The name of the zone is the same as the name of the DNS domain for your small organization or branch office.

#### Dynamic Update

On the **Dynamic Update** page, click **Allow both nonsecure and secure dynamic updates**, and then click **Next**. This will automate the update of the DNS resource records for the resources in your network.

#### Forwarders

On the **Forwarders** page, click **Yes, it should forward queries to DNS servers with the following IP addresses**, and then click **Next**. By selecting this configuration, you forward all DNS queries for DNS names outside of your network to a DNS server at either your ISP or central office. Type one or more IP addresses used by DNS servers run by either your ISP or central office.

#### Completing the Configure a DNS Server Wizard

On the **Completing the Configure a DNS Server Wizard** page of the Configure a DNS Server Wizard, you can click **Back** to change any of the settings. To apply your selections, click **Finish**.

#### Completing the Configure Your Server Wizard

After you complete the Configure a DNS Server Wizard, the Configure Your Server Wizard displays the **This Server is Now a DNS Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at *systemroot*\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

To verify that your server is secure and has the most recent updates, do the following:

1. Run Windows Update. For more information, see Windows Update.

2. Run the Security Configuration Wizard. For more information, see Security Configuration Wizard Overview.

**Removing the DNS server role**

If you need to reconfigure your server for a different role, you can remove existing server roles. If you remove the DNS server role from this server you should configure the TCP/IP settings of this server, and any clients using this server as their DNS server, with the IP address of a different DNS server.

To remove the DNS server role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **DNS server**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the DNS server role** check box, and then click **Next**. After you click **Next**, the **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page. On the **DNS Server Role Removed** page, click **Finish**.

**Next steps: Completing additional tasks**

After you complete the Configure Your Server Wizard and the Configure a New DNS Server Wizard, your server is ready for use as a DNS server. Up to this point, you have completed the following tasks:

- Set the DNS server to use a static IP address.

- Configured the DNS zone for your network.

- Configured the DNS server to forward all DNS queries for DNS names outside your network to a DNS server at your ISP or central office.

When you complete the Configure Your Server Wizard, it automatically installs the DNS console, which you use to manage your DNS server. To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

The following table lists some of the additional tasks that you might want to perform on your DNS server.

| Tasks | Purpose of task | Reference |
|---|---|---|
| Configure the computers in your network to use this DNS server as their preferred DNS server. | To connect your DNS clients to the DNS server and dynamically update the zone with the DNS resource records required for name resolution. | Configure TCP/IP to use DNS; Dynamic update |
| If this server is a multihomed computer, configure the DNS Server service to respond to queries on the local network IP address only. | To secure your DNS server from external queries. For example, a server acting as proxy server can have two network adapters, one for the intranet and one for the Internet. If that server is also running the Windows Server 2003 DNS Server service, you can configure the service to only use the intranet network adapter. By configuring this server to respond only to queries on the local network IP address, you will secure the server from unwanted Internet queries. | Restrict a DNS server to listen only on selected addresses |

| | | |
|---|---|---|
| Verify the server configuration. | To ensure that the DNS configuration performed using the Configure a New DNS Server Wizard is correct. | Verifying server configuration |
| Verify DNS server responsiveness using the **nslookup** command. | To verify that the DNS server is able to resolve DNS queries for resources in your network. | Verify DNS server responsiveness using the nslookup command |
| Verify that a resource records exist in DNS. | To verify that the computers that use the DNS server can be located on the network. | To verify A resource records exist in DNS |
| Configure ports to allow remote administration. | To manage the DNS server from other computers on the network. | Windows Firewall Settings |

# DHCP server role: Configuring a DHCP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP server role: Configuring a DHCP server**

DHCP servers centrally manage IP addresses and related information and provide it to clients automatically. This allows you to configure client network settings at a server, instead of configuring them on each client computer. If you want this computer to distribute IP addresses to clients, then configure this computer as a DHCP server.

This topic explains the basic steps that you must follow to configure a DHCP server. When you have finished setting up a basic DHCP server, you can complete additional configuration tasks, depending on how you want to use the DHCP server.

This topic covers:

Before you begin

Configuring your DHCP server

Next steps: Completing additional tasks

**Before you begin**

Before you configure your computer as a DHCP server, verify that:

- You are familiar with DHCP concepts such as scopes, leases, and options. For more information, see DHCP terminology.
- The operating system is configured correctly. In the Windows Server 2003 family, DHCP depends on the appropriate configuration of the operating system and its services. If you have a new installation of a Windows Server 2003 operating system, then you can use the default service settings. No further action is necessary. If you upgraded to a Windows Server 2003 operating system, or if you want to confirm that your services are configured correctly for best performance and security, then verify your service settings using the table in Default settings for services.
- This computer has a static IP address. For more information, see Configure TCP/IP for static addressing.
- All existing disk volumes use the NTFS file system. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

When you add the DHCP server role, you create one scope that defines the range of IP addresses that the DHCP server allocates to the clients on one subnet. You need to create one scope for each subnet that has clients that you want to manage using DHCP. The following table lists the information that you need to know before you add the DHCP server role, so that you can create the first scope. You need to collect the same information for each additional scope.

| Before adding a DHCP server role | Comments |
|---|---|
| Review DHCP security issues. | Security issues might affect the way you deploy DHCP servers. For more information about DHCP security best practices, see Security information for DHCP. |
| Identify the range of IP addresses that the DHCP server should allocate to the clients. | Use the entire range of consecutive IP addresses that make up the local IP subnet. In many cases, a private address range is the best choice. For more information and a list of all the IP address ranges approved for use on private networks, see RFC 1918, "Address Allocation for Private Internets", at the Internet Engineering Task Force Web site. Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here. |

| | |
|---|---|
| Determine the correct subnet mask for the clients. | When the DHCP server leases an IP address to a client, the server can specify additional configuration information, including the subnet mask. |
| Identify any IP addresses that the DHCP server should not allocate to clients. | For example, a server or a network-connected printer often has a static IP address, and the DHCP server must not offer this IP address to clients. |
| Decide the duration of the lease of the IP addresses. | The default is eight days. In general, the duration of the lease should be equal to the average time that the clients on this subnet are active. For example, the ideal duration may be longer than eight days if the clients are desktop computers that are rarely turned off, or it may be shorter than eight days if the clients are mobile devices that frequently leave the network or are moved between subnets. |
| (Optional) Identify the IP address of the router (default gateway) that the clients should use to communicate with clients on other subnets. | When the DHCP server leases an IP address to a client, the server can specify additional configuration information, including the IP address of the router. |
| (Optional) Identify the name of the DNS domain of the clients. | When the DHCP server leases an IP address to a client, the server can specify additional configuration information, including the name of the DNS domain to which the clients belong. |
| (Optional) Identify the IP address of the DNS server that the clients should use. | When the DHCP server leases an IP address to a client, the server can specify additional configuration information, including the IP address of the DNS server that the clients should contact to resolve the name of another computer. |
| (Optional) Identify the IP address of the WINS server that the clients should use. | When the DHCP server leases an IP address to a client, the server can specify additional configuration information, including the IP address of the WINS server that the clients should contact to resolve the NetBIOS name of another computer. |

**Configuring your DHCP server**

To configure a DHCP server, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **DHCP server**, and then click **Next**.

This section covers:

Summary of Selections

Using the New Scope Wizard

Completing the Configure Your Server Wizard

Removing the DHCP server role

**Summary of Selections**

On the **Summary of Selections** page, view and confirm the options that you have selected. If you selected **DHCP Server** on the previous page, the following appears:

- **Install DHCP Server**

- **Run the New Scope Wizard to configure a new DHCP scope**

To apply the selections shown on the **Summary of Selections** page, click **Next**. After you click **Next**, the **Configuring**

**Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page.

**Using the New Scope Wizard**

After you click **Next**, the Configure Your Server Wizard installs the DHCP Server service, and then starts the New Scope Wizard to collect information and add the DHCP server role. If you cancel the New Scope Wizard, the DHCP Server service remains installed, but it cannot distribute IP addresses to clients until you create a scope. If you choose to create a scope later, you can do so using the DHCP console.

This section describes the following steps in the New Scope Wizard:

Scope Name

IP Address Range

Add Exclusions

Lease Duration

Configure DHCP Options

Router (Default Gateway)

Domain Name and DNS Servers

WINS Servers

Activate Scope

Completing the New Scope Wizard

**Scope Name**

On the **Scope Name** page, in **Name**, type a name for the scope that you are creating. In **Description**, type a description (this is optional). Most networks have several subnets, and each subnet requires its own scope, so a DHCP server usually manages multiple scopes. Choose a name and description that help you distinguish between the various scopes.

After you finish, click **Next**.

**IP Address Range**

On the **IP Address Range** page, define the range of IP addresses in this scope by typing the IP addresses at the start and the end of that range. The wizard uses the IP addresses that you type to determine the correct subnet mask. The correct subnet mask automatically appears in **Subnet mask**.

In the unusual case where the clients on this subnet need to use a subnet mask other than the one that the wizard provides, you must type it in **Subnet mask**, or type the number of bits of the subnet mask in **Length**.

After you finish, click **Next**.

**Add Exclusions**

On the **Add Exclusions** page, you can define the IP addresses that the DHCP server should not allocate to clients. For example, the DHCP server itself has a static IP address that must not be allocated to clients. The same is true of the default gateway and of various network devices, such as network-connected printers. You must exclude these IP addresses so that the DHCP server does not allocate them to clients.

It is recommended that you exclude more IP addresses than you currently need because it is easier to truncate an exclusion range than it is to expand it. Exclude IP addresses from the beginning or the end of the range of possible IP addresses, not from the middle. For example, if the range of IP addresses on this subnet is from 10.0.0.1 to 10.0.0.255, and you want to exclude ten IP addresses, then define the exclusion range as either of the following:

- 10.0.0.1 to 10.0.0.10
- 10.0.0.245 to 10.0.0.255

For each range of IP addresses that you want to exclude, type the IP address at the beginning of the range in **Start IP address**, type the IP address at the end of the range in **End IP address**, and then click **Add**.

This step eases client administration, but it is optional. If you leave all fields of this page blank and click **Next**, clients will still be

able to obtain IP addresses from the DHCP server.

After you finish, click **Next**.

**Lease Duration**
On the **Lease Duration** page, you can define how long a client can use an IP address from this scope.

The DHCP server leases IP addresses to its clients. Each lease has an expiration date and time. The client must renew the lease if it is going to continue to use that IP address. The default duration of the lease is eight days.

This step eases client administration, but it is optional. If you leave all fields of this page blank and click **Next**, clients will still be able to obtain IP addresses from the DHCP server.

After you finish, click **Next**.

**Configure DHCP Options**
On the **Configure DHCP Options** page, you can specify whether to configure DHCP options.

It is recommended that you accept the default, **Yes, I want to configure these options now**, and click **Next**, even if you do not want to configure the options. If you choose **No, I will configure these options later**, then the following occurs:

- The New Scope Wizard creates the scope but does not activate it, and you must use the DHCP console to activate the scope in order for clients to receive IP addresses from this scope. For information about how to activate a scope using the DHCP console, see Activate a scope.

- You will not see the next four wizard pages described below. Continue reading at Completing the New Scope Wizard.

**Router (Default Gateway)**
On the **Router (Default Gateway)** page, you can specify the routers (also known as default gateways) that clients should use. You can add as many IP addresses as there are routers on this subnet.

This step eases client administration, but it is optional. If you leave all fields of this page blank and click **Next**, clients will still be able to obtain IP addresses from the DHCP server.

After you finish, click **Next**.

**Domain Name and DNS Servers**
On the **Domain Name and DNS Servers** page, you can specify the name of the domain that clients on this subnet should use when they resolve DNS names.

You can also specify the DNS server that clients should use to resolve DNS names. You can type the IP address of that DNS server, or you can type its name and click **Resolve**, and the wizard will determine the IP address for you. You can add several DNS servers.

This step eases client administration, but it is optional. If you leave all fields of this page blank and click **Next**, clients will still be able to obtain IP addresses from the DHCP server.

After you finish, click **Next**.

**WINS Servers**
On the **WINS Servers** page, you can specify the WINS server that clients should communicate with to register and resolve NetBIOS names. You can type the IP address of that WINS server, or you can type its name and click **Resolve**, and the wizard will determine the IP address for you. You can add several WINS servers.

This step eases client administration, but it is optional. If you leave all fields of this page blank and click **Next**, clients will still be able to obtain IP addresses from the DHCP server.

After you finish, click **Next**.

**Activate Scope**
On the **Activate Scope** page, you can activate the scope or choose to activate it later. In most cases, you should accept the default and activate the scope now. If you choose to activate the scope later, you can do so using the DHCP console. You must activate the scope to allow clients on the subnet of this scope to obtain IP addresses from the DHCP server.

After you finish, click **Next**.

**Completing the New Scope Wizard**
On the **Completing the New Scope Wizard** page, you can click **Back** to change any of the settings. To apply your selections, click **Finish**.

**Completing the Configure Your Server Wizard**
After you complete the New Scope Wizard, the Configure Your Server Wizard displays the **This Server is Now a DHCP Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at *systemroot*\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

To verify that your server is secure and has the most recent updates, do the following:

1. Run Windows Update. For more information, see Windows Update.

2. Run the Security Configuration Wizard. For more information, see Security Configuration Wizard Overview.

**Removing the DHCP server role**
If you need to reconfigure your server for a different role, you can remove existing server roles. If you remove the DHCP server role, you must reconfigure clients to obtain IP addresses from another DHCP server. Also, if you want clients to continue using the IP addresses distributed by this server, you must configure another DHCP server to distribute those addresses.

To remove the DHCP server role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **DHCP server**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the DHCP server role** check box, and then click **Next**. After you click **Next**, the **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page. On the **DHCP Server Role Removed** page, click **Finish**.

**Next steps: Completing additional tasks**
After you complete the Configure Your Server Wizard, the computer is ready for use as a basic DHCP server that can manage IP addresses and related information. Up to this point, you have installed the DHCP Server service and created one scope that manages IP addresses and related information for the clients on one subnet. If you want to manage clients on other subnets, you must create additional scopes. If you have not activated the scope, then you must activate the scope to allow clients on the subnet of this scope to obtain IP addresses from the DHCP server.

When you complete the Configure Your Server Wizard, it automatically installs the DHCP console, which you use to manage your DHCP server. To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

The following table lists some of the additional tasks that you might want to perform on your DHCP server.

| Task | Purpose of task | Reference |
|---|---|---|
| Activate existing scopes. | To activate scopes that you created with the New Scope Wizard that are not currently active. | Activate a scope |
| Authorize the DHCP server in Active Directory. | To detect unauthorized DHCP servers and prevent them from starting or running on your network. | Authorizing DHCP servers |
| Add support for a routed network. | To enable the DHCP server to act as a relay agent, which forwards DHCP messages between subnets. | Configure the DHCP Relay Agent |
| Assign a server-based scope option. | To simplify administration by setting default values that all new scopes will inherit. | Assign a server-based option |

| | | |
|---|---|---|
| Change or view scope properties. | To change or view the properties of an existing scope. | Change or view scope properties |
| Create new scopes. | To create scopes that support clients on other subnets. | Create a new scope |
| Add a client reservation. | To define an IP address that the DHCP server should not allocate to a client because another client is permanently using it. | Add a client reservation |
| Configure ports to allow remote administration. | To manage the DHCP server from other computers on the network. | Windows Firewall Settings |

# Streaming media server role: Configuring a streaming media server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Streaming media server role: Configuring a streaming media server**
You can use Windows Media Services to stream audio and video content to clients over the Internet or an intranet. Clients might be computers or devices that play back the content using a player, such as Windows Media Player, or they might be computers running Windows Media Services (called Windows Media servers) that proxy, cache, or redistribute your content. Clients can also be custom applications that have been developed with the Windows Media Software Development Kit (SDK).

If you want this computer to provide audio and video content streams to clients and to other Windows Media servers, then configure this computer as a streaming media server.

### Notes

- This content is not available in this preliminary release.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition operating system. For more information, see Overview of Windows Server 2003, Web Edition.

This topic explains the basic steps that you must follow to configure a streaming media server. When you have finished the basic steps, you can complete additional configuration tasks, depending on how you want to use the streaming media server.

This topic covers:

Before you begin

Configuring your streaming media server

Next steps: Completing additional tasks

**Before you begin**
Before you configure your computer as a streaming media server, verify whether or not:

- The operating system is configured correctly. In the Windows Server 2003 family, Windows Media Services depends on the appropriate configuration of the operating system and its services. If you have a new installation of a Windows Server 2003 operating system, you can use the default service settings. No further action is required. If you upgraded to a Windows Server 2003 operating system, or if you want to confirm that your services are configured correctly for best performance and security, then verify your service settings using the table in Default settings for services.
- All existing disk volumes use the NTFS file system. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

**Configuring your streaming media server**
To configure a streaming media server, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and

then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Streaming media server**, and then click **Next**.

This section covers:

Summary of Selections

Completing the Configure Your Server Wizard

Completing the streaming media server role configuration

Removing the streaming media server role

**Summary of Selections**

On the **Summary of Selections** page, view and confirm the options that you have selected. If you selected **Streaming media server** on the **Server Role** page, the following appears:

- **Install Windows Media Services**

To apply the selections shown on the **Summary of Selections** page, click **Next**. After you click **Next**, the **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page.

After you click **Next** on the **Summary of Selections** page, the Configure Your Server Wizard installs Windows Media Services. Unlike many other services, Windows Media Services installs without any input from the administrator.

**Completing the Configure Your Server Wizard**

After the components are configured, the Configure Your Server Wizard displays the **This Server is Now a Streaming Media Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at *systemroot*\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

Before you start to use your streaming media server, we recommend the following steps:

- Run Windows Update. For more information, see Windows Update.

- Run the Security Configuration Wizard. For more information, see Security Configuration Wizard Overview.

**Completing the streaming media server role configuration**

After you complete the Configure Your Server Wizard, the computer is ready for use as a basic streaming media server that can provide digital media content to clients and to other streaming media servers. Additional configuration is usually required, and the specific steps depend on your requirements. This section explains the basic decisions you must make in order to configure the streaming media server.

**Note**

- If you have installed Windows Media Services on this computer, you can view Windows Media Services Help. To open Windows Media Services Help, click **Start**, click **Run**, and then type **hh wmserver.chm**.

The streaming media server role supports many scenarios, which Windows Media Services Help describes in detail. For more information, see "Scenarios" in Windows Media Services Help. Most scenarios require you to reconfigure an existing publishing point or create a new one. You need to make one decision, or possibly two, and the results determine which of the three main publishing point configurations you should use. The following table shows how the decisions relate to configurations.

| If you | And you want to use | Use this publishing point configuration |
|---|---|---|
| Want clients to control playback | One server connection per client | On-demand, unicast |

| | | |
|---|---|---|
| Do not want clients to control playback | One server connection per client | Broadcast, unicast |
| Do not want clients to control playback | One server connection, shared by all clients | Broadcast, multicast |

**Control of playback**
Control of playback means that the client should be able to start, stop, pause, rewind, and fast-forward digital media content. With on-demand, unicast, the client controls playback, and the user experience is similar to playing a movie from a VCR or a DVD player. This type of playback requires an on-demand publishing point. An on-demand publishing point distributes pre-recorded content, such as audio and video files. When you add the streaming media server role, the wizard creates an on-demand publishing point named **<Default>**. You can distribute your media files from this publishing point, or you can create another publishing point. The configuration steps for an on-demand publishing point are similar to those presented in "Stream Windows Media files on-demand" in Windows Media Services Help. If you choose to use an on-demand publishing point, you must use unicast delivery.

If the client does not control playback, the user experience is similar to viewing a television program. This type of playback requires a broadcast publishing point. This type of publishing point distributes pre-recorded and live content. When you add the streaming media server role, the wizard creates a broadcast publishing point named **Sample_Broadcast** that contains sample content. You should leave this sample broadcast publishing point intact and create a new broadcast publishing point. If you choose to create a broadcast publishing point, see Server connections.

For more information about on-demand and broadcast publishing point types, see "About publishing point types" in Windows Media Services Help.

**Server connections**
With unicast broadcast, the server creates a separate connection to each client. As a result, unicast delivery can consume a large amount of network bandwidth. For example, delivering the same content to 100 clients simultaneously consumes 100 times as much network bandwidth as delivering the content to one client. However, unicast delivery does not require any configuration of network routers and switches. The steps to configure a publishing point this way are similar to those presented in "Use your server to publish live content from Windows Media Encoder" in Windows Media Services Help. For more information about unicast delivery, see "Delivering content as a unicast stream" in Windows Media Services Help.

With multicast broadcast, the server does not create a connection to any client. Instead, the server delivers the content to a Class D Internet Protocol (IP) address on the network, and any client on the network can receive it. This conserves network bandwidth. For example, a multicast delivery to 100 clients consumes only as much bandwidth as delivery to one client. However, many networks by default do not support multicast delivery. To support multicast delivery, the network routers and switches between the server and the clients must be configured to transmit Class D IP addresses and interpret multicast information packets. The steps to configure a publishing point for multicast broadcast are similar to those presented in "Use your server to broadcast a stream published by Windows Media Encoder" in Windows Media Services Help. For more information about multicast delivery, see "Delivering content as a multicast stream" in Windows Media Services Help.

**Using firewalls**
If you plan to stream content from a Windows Media server on a network to clients on the Internet, you may need to open additional ports in your firewall to prevent clients from having problems receiving the content. Also, to make sure that your content is always available to clients that connect to your server using a URL that starts with a Microsoft Media Server (MMS) moniker (mms://), ensure that ports on your firewall are opened for all of the connection protocols that might be used during protocol rollover. For more information, see "Using firewalls" in Windows Media Services Help.

**Removing the streaming media server role**
If you need to reconfigure your server for a different role, you can remove existing server roles. If you remove the streaming media server role, clients will no longer be able to connect to the publishing points of this server, and encoders will no longer be able to send media streams through the server. All content stored or distributed on this server will become unavailable.

To remove the streaming media server role, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **Streaming media server**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the streaming media server role** check box, and then click **Next**. After you click **Next**, the **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page. On the **Streaming Media Server Role Removed** page, click **Finish**.

**Next steps: Completing additional tasks**

Up to this point, you have installed Windows Media Services. The installation added Windows Media Services to your server, installed the Help file, and created two publishing points that contain sample content.

The following table lists some of the additional tasks that you might want to perform on your streaming media server.

| Task | Purpose of task | Reference |
|---|---|---|
| Configure security options. | To control access to the streaming media server and its content. | "Configuring security options" in Windows Media Services Help |
| Take the tour. | To become more familiar with Windows Media Services capabilities. | Click **Start**, click **Run**, and then type **%systemroot%\system32\windows media\server\admin\mmc\hta\tour_.hta** |
| Review streaming media terms and concepts. | To become more familiar with streaming media concepts such as unicast and multicast, on-demand and broadcast, archiving, publishing points, and announcing content. | "Understanding terminology" in Windows Media Services Help |
| Determine how many streaming media servers you need. | To plan ahead for the number of servers you need to install. | "Streaming media system overview" in Server roles, and "Capacity planning" in Windows Media Services Help |
| Identify port conflicts. | To prevent problems when Windows Media Services attempts to use the same TCP port as a Web server. | "Using HTTP streaming and other services on the same computer" in Windows Media Services Help |
| Review instructions for upgrading an earlier version of Windows Media Services. | To ensure that you understand how to upgrade other servers running earlier versions of Windows Media Services. | "Upgrading Windows Media Services" in Windows Media Services Help |
| Enable the Windows Audio service. | The Windows Audio service is disabled by default on a new installation of Windows Server 2003, Datacenter Edition, or Windows Server 2003, Enterprise Edition. This does not prevent the server from streaming audio to clients, but you should enable audio in order to test content playback on the server. | Enable or disable a service for a hardware profile |
| Start the Windows Media Services management interface. | To configure your streaming media server. | Click **Start**, click **Run**, and then type **%systemroot%\system32\windows media\server\admin\mmc\wmsadmin.msc** |
| Manage your streaming media server. | To configure your streaming media server to stream content over an intranet or the Internet. Before you begin streaming content, you must configure settings for your server running Windows Media Services, add and configure publishing points, and set up your content. | "Managing your Windows Media server" in Windows Media Services Help |

| | | |
|---|---|---|
| Log data and events. | To record the activity of clients that connect to your content. | "Logging data and events" in Windows Media Services Help |
| Manage and produce content. | Content management methods and priorities will differ from one project to another based on a variety of factors, such as audience demographics, content type, and available equipment. | "Content management and production" in Windows Media Services Help |
| Decide how to obtain content from the Windows Media Encoder. | To determine whether to configure the encoder to push a stream to the server or to configure the server to pull a stream from the encoder. | "Sourcing from an encoder" in Windows Media Services Help |
| Implement a cache/proxy system. | To store the most recently streamed content for use by other clients seeking the same material. During live broadcasts, cache/proxy servers can perform a task called stream splitting, which allows many unicast clients to receive content while only a single stream is sent from the origin server. | "Implementing a cache/proxy system" in Windows Media Services Help |
| Administer the streaming media server remotely. | To manage the streaming media server from other computers on the network. | "Using Windows Media Services Administrator for the Web" in Windows Media Services Help. |

# WINS server role: Configuring a WINS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WINS server role: Configuring a WINS server**

Windows Internet Name Service (WINS) servers dynamically map IP addresses to computer names (NetBIOS names). This allows users to access resources by computer name instead of by IP address. If you want this computer to keep track of the names and IP addresses of other computers in your network, configure this computer as a WINS server.

This topic explains the basic steps that you must follow to configure a WINS server. When you have finished setting up a basic WINS server, you can complete additional configuration tasks, depending on how you want to use the WINS server.

This topic covers:

Before you begin

Configuring your WINS server

Next steps: Completing additional tasks

**Before you begin**

Before you configure your computer as a WINS server, verify that:

- You are familiar with WINS concepts such as NetBIOS names, WINS server, WINS clients, and replication partners. For more information, see Understanding WINS.
- The operating system is configured correctly. In the Windows Server 2003 family, WINS depends on the appropriate configuration of the operating system and its services. If you have a new installation of a product in the Windows Server 2003 family, you can use the default service settings. No further action is necessary. If you upgraded to a product in the Windows Server 2003 family, or if you want to confirm that your services are configured correctly for best performance and security, then verify your service settings using the table in Default settings for services.
- You know how many WINS servers you need to install and where to locate each server on your network. When you add the WINS server role, you configure this server to maintain a database of computer names and IP addresses. In a large network, you may need to add the WINS server role to additional servers in order to ensure that client computers always have access to at least one WINS server. For more information, see Planning WINS networks.
- This computer has a static IP address. For more information, see Configure TCP/IP for static addressing.
- All existing disk volumes use the NTFS file system. FAT32 volumes are not secure, and they do not support file and folder compression, disk quotas, file encryption, or individual file permissions.
- Windows Firewall is enabled. For more information, see Enable Windows Firewall with no exceptions.
- The Security Configuration Wizard is installed and enabled. For information about the Security Configuration wizard, see Security Configuration Wizard Overview.

**Configuring your WINS server**

To configure a WINS server, start the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **WINS server**, and then click **Next**.

Summary of Selections

Completing the Configure Your Server Wizard

Removing the WINS server role

**Summary of Selections**
On the **Summary of Selections** page, view and confirm the options that you have selected. If you selected **WINS server** on the previous page, the following appears:

- **Install WINS**

To apply the selections shown on the **Summary of Selections** page, click **Next**. After you click **Next**, the **Configuring Components** page of the Windows Components Wizard appears, and then closes automatically. You cannot click **Back** or **Next** on this page. The Configure Your Server Wizard installs the WINS Server service. Unlike many other services, the WINS service installs without any input from the administrator.

If you cancel Configure Your Server, then WINS Server service is not installed. To install it later, restart Manage Your Server and add the WINS role.

**Completing the Configure Your Server Wizard**
After the components are configured, the Configure Your Server Wizard displays the **This Server is Now a WINS Server** page. To review all of the changes made to your server by the Configure Your Server Wizard or to ensure that a new role was installed successfully, click **Configure Your Server log**. The Configure Your Server Wizard log is located at *systemroot*\Debug\Configure Your Server.log. To close the Configure Your Server Wizard, click **Finish**.

To verify that your server is secure and has the most recent updates, do the following:

1. Run Windows Update. For more information, see Windows Update.

2. Run the Security Configuration Wizard. For more information, see Security Configuration Wizard Overview.

**Removing the WINS server role**
If you need to reconfigure your server for a different role, you can remove existing server roles. If you remove the WINS server role, and this server is the only WINS server that clients can use to register and resolve computer names, you need to add the WINS role to another server. Also, if this server is configured to replicate WINS database information with other WINS servers, then you must reconfigure replication on those other WINS servers.

To remove the WINS server role, restart the Configure Your Server Wizard by doing either of the following:

- From Manage Your Server, click **Add or remove a role**. By default, Manage Your Server starts automatically when you log on. To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

On the **Server Role** page, click **WINS server**, and then click **Next**. On the **Role Removal Confirmation** page, review the items listed under **Summary**, select the **Remove the WINS server role** check box, and then click **Next**. On the **WINS Server Role Removed** page, click **Finish**.

**Next steps: Completing additional tasks**
After you complete the Configure Your Server Wizard, the computer is ready for use as a basic WINS server that can keep track of server IP addresses and provide this information when a client requests it. Up to this point, you have installed the WINS Server service on one server. If you want to support WINS clients on a complex network, you may need to install additional WINS servers on other subnets.

The following table lists some of the additional tasks that you might want to perform on your WINS server.

| Task | Purpose of task | Reference |
| --- | --- | --- |
| View WINS name records registered at the server. | Verify that the WINS server is functioning correctly. | Finding and viewing WINS records |
| Modify WINS server defaults. | WINS uses several default server configuration parameters that determine how NetBIOS name records are managed in the WINS server database. These parameters are usually acceptable. You may want to modify them in special circumstances, such as when a host name change must be made, or when you renumber the network so that clients and servers use a different set of IP addresses. | Modifying WINS server defaults |
| Configure replication settings on primary and secondary WINS servers. | WINS servers replicate database changes to each other so that each WINS server has the same information about the servers on the network and the IP address of each. | Configuring WINS replication; Configure Replication |
| Configure ports to allow remote administration. | To manage the WINS server from other computers on the network. | Windows Firewall Settings |

# To remove a server role

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a server role**

1. Open Manage Your Server.
2. Under **Managing Your Server Roles**, click **Add or remove a role**.
3. In the Configure Your Server Wizard, on the **Preliminary Steps** page, click **Next**.
4. On the **Server Role** page, click the role that you want to remove, and then click **Next**.
5. On the **Role Removal Confirmation** page, select the **Remove the *server* role** check box, and then click **Next**.
6. On the ***Server* Role Removed** page, click **Finish**.

**Important**

- When you remove a server role, you might break dependencies that exist between the server role and other programs. Be sure to test removing the server role in a test environment before doing so on your network.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Manage Your Server, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
- For more information about removing specific server roles and any dependencies that might be affected, see the following:

  File server role: Configuring a file server

  Print server role: Configuring a print server

  Application server role: Configuring an application server

  Mail server role: Configuring a mail server

  Terminal server role: Configuring a terminal server

  Remote access/VPN server role: Configuring a remote access/VPN server

  Domain controller role: Configuring a domain controller

  DNS server role: Configuring a DNS server

  DHCP server role: Configuring a DHCP server

  Streaming media server role: Configuring a streaming media server

WINS server role: Configuring a WINS server

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Windows Product Activation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Product Activation**

Due to piracy and other forms of unauthorized use, users cannot always be sure that they have a genuine copy of a Windows Server 2003 operating system. The goal of product activation is to reduce a form of piracy known as casual copying or "softlifting". Casual copying is the sharing and installation of software that is not in compliance with the software's end user license agreement, or EULA, and is estimated to contribute to half of all pirated installations. Microsoft developed Product Activation for Windows Server 2003 operating systems to help ensure that each Windows Server 2003 operating system license is installed in compliance with the EULA and is not installed on more than the limited number (usually one) of computers allowed by the product EULA. Please note, however, that volume licensed versions of Windows Server 2003 operating systems do not require activation.

During software installation, the set up wizard prompts you to enter a product key usually located on the back of the Windows Server 2003 operating system CD-ROM folder. The product key is a 25 character alphanumeric code shown in five groups of five characters each (for example, BCDFG-12345-HJKLM-67890-NPQRS). Keep the product key in a safe location and do not share it with others. This product key forms the basis for your ability to install and use Windows.

The product key also forms the basis for the product ID that is created when members of the Windows Server 2003 family are installed. Each licensed instance of Windows Server 2003 operating systems has a unique Product ID. The Product ID has 20 characters arranged like this: XXXXX-XXX-XXXXXXX-XXXXX. It is listed in the properties for **My Computer**.

For purposes of product activation only, a non-unique hardware hash is also created from general information that is included in the system components. At no time are files on the hard drive scanned, nor is personally-identifiable information of any kind used to create the hardware identifier. None of the information collected during product activation will be used to personally identify you. To ensure your privacy, the hardware identifier is created by what is known as a "one-way hash". To produce a one-way hash, information is processed through an algorithm to create a new alphanumeric string. It is impossible to calculate the original information from the resulting string.

The hardware identifier is used together with the product ID to create a unique installation ID. Whether you choose to activate by using an Internet connection or by speaking with a Microsoft customer service representative, the installation ID is the only piece of information required to activate Windows Server 2003 operating systems.

If you activate via an Internet connection, the installation ID is sent automatically. When you decide to activate over the Internet, the Windows Product Activation Wizard attempts to establish an online connection to Microsoft by way of the Internet. If you do not subscribe to an Internet Service Provider but do have a modem connected to a phone line, the wizard detects the modem and attempts to make a direct connection to Microsoft.

If an online connection cannot be established, you are prompted to contact a customer service representative by telephone. In that message, the installation ID is displayed to you. The customer service representative will ask you to read the installation ID over the telephone. Alternatively, automated telephone response systems will guide you through the activation process.

The installation ID records an association of the product ID to your computer and a confirmation is sent back. The product key can now be used to install Windows on that computer an unlimited number of times. However, if you need to install Windows on a different computer using that product key, you might need to contact a Microsoft customer service representative by telephone.

Until you have activated your copy of your Windows Server 2003 operating system, an Activate icon appears in the Notification Area of the Status Bar of the Desktop. You can click the icon to initiate activation. This icon will not appear in the Notification Area after you have activated the member of the Windows Server 2003 family. Check the Application Event Log to verify activation.

You have normally have a grace period of days--specified in your EULA--in which to activate your Windows Server 2003 operating systems product installation. If the grace period expires and you have not completed activation, you will not be able to complete the login process and Terminal Server sessions will be unavailable. Administrator credentials are required for reactivation. Please note that evaluation CDs and upgrades are not issued an additional grace period. You can logon in Safe Mode to access any available data.

**Notes**

- To view the EULA, click **Start**/**Run**, type "eula.txt".

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- For further information about Windows Product Activation, see How to activate Windows and Windows Product Activation Concepts.

- Windows Product Activation is not available on the Itanium-based versions of the Windows operating systems.

- For the most recently updated information regarding Windows Product Activation, see Microsoft Product Activation Privacy Policy at Microsoft Web site.

# How to activate Windows

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Activate Windows on the Internet
- Activate Windows with your modem when you have no Internet service provider
- Activate Windows by phone
- Activate Windows using Microsoft Text Telephone

# Activate Windows on the Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate Windows on the Internet**

1. Click **Start**, point to **All Programs**, and then click **Activate Windows**.
2. In **Windows Product Activation**, select **Yes, let's activate Windows over the Internet now**. After this selection, you will be given the option to register at the same time.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a system tools item, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click the appropriate icon. Click **Activate Windows**.
- Alternatively, you can open the Windows Product Activation Wizard by clicking **Start**, then **Run**, and then typing "**oobe/msoobe /a**".
- You are strongly encouraged to activate your Windows Server 2003 operating system before using it in a productive environment to prevent loss of time later.
- If you choose not to activate Windows now, reminders will appear periodically during your grace period. After that time has expired, your login process will remain incomplete until activation.
- If you have not yet activated your Windows Server 2003 operating system, you can do so at any time by clicking the Windows Activation icon in the Notification Area on the Status Bar on the Desktop to initiate activation.
- If you have neither an Internet connection or a modem, see "To activate Windows by phone" in Related Topics for instructions on how to contact a customer service representative by phone.
- Activating any of the Windows Server 2003 operating systems is fast and does not require your name or other personal information.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Windows Product Activation
Understanding activation and registration
Microsoft Windows Product Activation privacy statement
Understanding product keys
Activate Windows using Microsoft Text Telephone
Activate Windows with your modem when you have no Internet service provider
Activate Windows by phone
Default local groups
Default groups
Using Run as

# Activate Windows with your modem when you have no Internet service provider

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate Windows with your modem when you have no Internet service provider**

1. Click **Start**, point to **All Programs**, and then click **Activate Windows**.
2. In **Windows Product Activation**, select **Yes, let's activate Windows over the Internet now**. After this selection, you will be given the option to register at the same time.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a system tools item, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click the appropriate icon. Click **Activate Windows**.
- Even without Internet service, you can use your modem to activate, or activate and register, your copy of Windows.
- You are strongly encouraged to activate your Windows Server 2003 operating system before using it in a productive environment to prevent loss of time later.
- When Windows operating systems require activation, a shortcut logo appears in the Notification Area on the Status Bar on the Desktop. You can begin the activation process by clicking on this icon. Once you have activated, this icon disappears. Confirm activation by checking with your Application Event Log.
- Alternatively, you can open the Windows Product Activation wizard by clicking **Start**, then **Run**, and type in "**oobe/msoobe /a**".
- If you choose not to activate now, reminders will appear periodically during the grace period. After that time has expired, you will be required to activate in order to continue to use your Windows Server 2003 operating system.
- Activating any of the Windows Server 2003 operating systems is fast and does not require your name or other personal information.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Windows Product Activation
Understanding activation and registration
Microsoft Windows Product Activation privacy statement
Understanding product keys
Activate Windows using Microsoft Text Telephone

# Activate Windows by phone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate Windows by phone**

1. Click **Start**, click **All Programs** and then click **Activate Windows**.
2. In **Windows Product Activation**, select **Yes, I want to telephone a customer service representative to activate Windows**.
3. After you select the appropriate location, you are provided with a telephone number and an installation ID that you can give to the customer service representative to complete the activation process.

## Notes

- To open a system tools item, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click the appropriate icon. Click **Activate Windows**.
- Alternatively, you can open the Windows Product Activation Wizard by clicking **Start**, then **Run**, and type in "**oobe/msoobe /a**".
- You are strongly encouraged to activate your Windows Server 2003 server operating system before using it in a productive environment to prevent loss of time later.
- If you have not yet activated your Windows Server 2003 operating system, you can initiate activation at any time by clicking the Windows Activation icon in the Notification Area. Once you have activated, this icon disappears. If you choose not to activate now, reminders will appear periodically during the grace period. After that time has expired, you will be required to activate in order to continue to use your Windows Server 2003 operating system interactively.
- Activating any of the Windows Server 2003 operating systems is fast and does not require your name or other personal information.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Accessing information about accessibility
Windows Product Activation
Understanding activation and registration
Microsoft Windows Product Activation privacy statement
Understanding product keys
Activate Windows using Microsoft Text Telephone

# Activate Windows using Microsoft Text Telephone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate Windows using Microsoft Text Telephone**

1. Dial (800) 718-1599 on your TT/TTY modem from the United States or Canada.
2. Once you are connected to Microsoft, you will be guided by on-screen instructions for activation.
3. You are strongly encouraged to activate your Windows Server 2003 operating system before using it in a productive environment to prevent loss of time later.

## Notes

- Microsoft Text Telephone (TT/TTY) services are available for individuals who are deaf or hard-of-hearing.
- From outside the United States, dial United States number (716) 871-6859, using your TT/TTY modem, and then follow the on-screen instructions for activation.
- TT/TTY service is available only in English.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Accessing information about accessibility
Windows Product Activation
Understanding activation and registration

# Windows Product Activation Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts...**

- Understanding activation and registration
- Microsoft Windows Product Activation privacy statement
- Understanding product keys

# Understanding activation and registration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding activation and registration**
There are important distinctions between activating your Windows operating system and registering it.

Product activation and product registration are not the same. Activation is completely anonymous; it requires absolutely no personally identifiable information to complete. Activation is required and ensures that each Windows product is not installed on more than the limited number (usually one) of computers allowed in the software's end user license agreement, or EULA.

During activation, you can also register your copy of Windows. Registration is **not** required; however registering your copy ensures that you receive product support, product update information, and other benefits. The process registration takes only a few minutes to complete. Personal information (for example, "contact" information such as an e-mail address) is required if you decide to register.

You have a short grace period in which to activate your Windows product installation. If the grace period expires and you have not completed activation, login will be incomplete until you activate.

For more information, see the following.

- Activate Windows on the Internet
- Activate Windows with your modem when you have no Internet service provider
- Activate Windows by phone
- Activate Windows using Microsoft Text Telephone

# Microsoft Windows Product Activation privacy statement

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Windows Product Activation privacy statement**
Microsoft asks that you activate Windows so that we may verify that your installation is done with a genuine Microsoft product. When you activate Windows, you are not required to send any personal information to Microsoft. Activation is completely anonymous.

During activation, Windows creates a unique non-hardware identification that represents the configuration of the computer at the time of activation. The hardware identification does not include any personal information, any information about software or data that may reside on your computer, or any information about the specific make or model of your computer. The hardware identification identifies only the computer and only for the sole purpose of activation. Windows can detect and tolerate changes to your computer configuration. Minor upgrades will not require re-activation. If you completely overhaul your computer you may be required to activate Windows again.

During activation you may also provide your name and other contact information if you want to register your product with Microsoft. If you choose to register your product, the information that you provide will be securely stored by Microsoft and will be protected from disclosure to third parties without your consent.

If you believe that Microsoft has not adhered to these privacy principles, please notify us by postal mail at the following address:

**Microsoft**

**Attn: Microsoft Product Activation**

**One Microsoft Way**

**Redmond, Washington 98052-6399 USA**

We will use commercially reasonable efforts to determine and correct the problem promptly.

Here are national and international **Customer Service Centers** telephone numbers:

| Country/Region | Telephone Number |
|---|---|
| Argentina | 54-11-4317-2626 |
| Australia | + 61 2 9870 2131 |
| Brazil | (11) 3328-3396 |
| United Kingdom | + 44 870 2411963 |
| Hong Kong SAR | + 852 2904-2198 |
| India | +91 11 629 2640 |
| Japan | +81-3-6831-3460 |
| New Zealand | +64 800 800 004 |
| United States | 716-871-2929 |

# Understanding product keys

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding product keys**

Product keys are required for all versions of the Windows Server 2003 family, for all languages, and for both 32-bit and 64-bit versions. Product keys are 25-character alphanumeric strings that are formatted in groups of five characters, separated by dashes. You must enter the product key during Setup to complete the installation. (If you acquired your computer from an Original Equipment Manufacturer, the manufacturer may have preconfigured the product key for the Windows Server 2003 operating system software that is preinstalled on the computer).

Windows Server 2003 operating systems require you to keep your product keys always available for your use and in a secure place. You must use only the product key assigned to this specific Windows Server 2003 operating system installation if you ever need to reinstall the operating system. If you attempt to use a product key from another Microsoft product, a beta program, or a different version of Windows Server 2003 operating system, your installation will fail.

Each product key can only be used to install the software on a certain number of unique computers. The actual number of computers depends on your End User License Agreement (EULA). If you have received more than one product key, it is essential that you keep track of which product key you use to install the Windows Server 2003 operating system on each computer.

**Automating the product key entry during Windows Product Activation**

If you perform an unattended Setup, you can automate entry of the product key. In the **[UserData]** section of your unattended answer file, include the value of your product key in the **ProductKey** entry. For more information, see the Microsoft Windows Preinstallation Reference, located in the \Support\Tools\deploy.cab file on the Windows CD.

Products in the Windows Server 2003 family provide two methods to automate activation.

You can script Windows Product Activation through the use of WMI (Windows Management Instrumentation), which provides a number of useful query and activation methods.

# Windows Product Activation Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- I can't find Activate Windows program.
- My Windows operating system has been activated, but I need to read the online Help.
- I didn't activate the product, and now I can't logon.

**I can't find Activate Windows program.**
**Cause**: The key icon for Windows Product Activation is located in the Notification Area of the Desktop and may not be obvious.

**Solution**: Locate the icon on the lower right side of the Desktop. Alternatively, click **Start**, point to **All Programs**, and then click **Activate Windows**.

**My Windows operating system has been activated, but I need to read the online Help.**
**Cause**: Once the operating system has been activated, the icon for activation is removed from the Notification Area on the Desktop.

**Solution**: Open the Help and Support Center from the **Start** menu. Click **Getting Started** and select "Windows Product Activation" from either list. Alternatively, click **Start**, then **Run**. Type "hh wpa.chm", click **OK**.

1. Click **Start**.

   **Where?**

   - **Help and Support**/Getting Started/Windows Product Activation

**See also:** Windows Product Activation.

**I didn't activate the product, and now I can't logon.**
**Cause:** After the grace period, Windows Product Activation does not shut down the computer, but it does prevent interactive logon. For example, if the computer or its service is running on the 30th day, it will continue to run on the 31st day if the computer is not activated and not rebooted. On the next reboot and login, the user will be blocked from logging on.

**Solution:** The user can still boot into the Safe Mode minimum (without networking) to activate, upgrade, or simple retrieve data.

# Accessibility for People with Disabilities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessibility for people with disabilities**
Microsoft is committed to making its products and services easier for everyone to use. This topic provides links to information on features and services that make products in the Windows Server 2003 family more accessible for people with disabilities.

**Note**

- The information in this topic applies only to users who license Microsoft products in the United States. If you obtained this product outside the United States, your package contains a subsidiary information card that lists telephone numbers and addresses for Microsoft support services. You can contact your subsidiary to find out whether the type of products and services described in this topic are available in your area.

The following sources of information can help you learn more about accessibility features and services:

- To learn more about Microsoft products that are available for people with disabilities, see Accessibility at the Microsoft Web site.
- For detailed information about accessibility features and options, see Accessibility Options overview.
- For more information about using keyboard shortcuts in Microsoft Management Console (MMC), see Accessibility for MMC.

In addition, the following topics provide basic information about accessibility features, options, and resources.

- Accessibility Features and Options
- Getting More Accessibility Information

**Note**

- To display Accessibility Options Help, click **Start**, click **Run**, and type **hh access.chm**.

# Accessibility Features and Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessibility features and options**
Products in the Windows Server 2003 family include a number of accessibility features that can improve display, sound, mouse, and keyboard settings for users who are blind, have motion disabilities, or are deaf or hard of hearing. Many accessibility features are also useful to people without disabilities.

For detailed information about accessibility features and options, see Accessibility Options overview.

- Accessibility features
- Setting up accessibility options
- Using the keyboard to move the mouse pointer
- Configuring Windows for people who are blind or have impaired vision
- Configuring Windows for people who have a mobility impairment

**Note**

- To display Accessibility Options Help, click **Start**, click **Run**, and type **hh access.chm**.

For a list of links to other information about accessibility features, see Accessibility for People with Disabilities.

# Accessibility features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessibility features**

Products in the Windows Server 2003 family include the following utilities to help meet your vision, hearing, or mobility needs:

- **Magnifier.** Enlarges a portion of your screen display for easier viewing.
- **Narrator.** Reads information on your screen, including dialog box names, menus, and contents. Narrator provides a minimum level of accessibility so that people who are blind or who have low vision can interact with products in the Windows Server 2003 family during configuration or emergency situations.

  **Notes**

  - Narrator is available only with U.S. English.
  - Narrator requires the Windows Audio service, which by default has a **Startup Type** of **Disabled** on computers running Windows Server 2003, Web Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. To use Narrator, make sure the Windows Audio service is started and set the **Startup Type** to **Automatic**. For more information, see Default settings for services.
- **On-Screen Keyboard.** Allows a pointing device such as a mouse or a switch input device to type on a keyboard displayed on your screen. On-Screen Keyboard is compatible with many, but not all, keyboards and keyboard settings.

Products in the Windows Server 2003 family also offer customizable options for users who are blind, have motion disabilities, or are deaf or hard of hearing:

- **High-contrast desktop color schemes.** Provide more ways to change the colors and font sizes on your screen.
- **Sound schemes.** Give useful audio feedback for important on-screen events, such as the opening and closing of windows.
- **High-visibility mouse pointer schemes.** Give you more options to help you keep track of the pointer visually.

For a list of links to other information about accessibility features, see Accessibility for People with Disabilities.

# Setting up accessibility options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up accessibility options**

You can configure Accessibility options in products in the Windows Server 2003 family in several ways to meet vision, hearing, or mobility needs. Accessibility features are installed on your computer by default.

You can quickly set your accessibility options by using the Accessibility Wizard. The Accessibility Wizard asks you questions about your accessibility needs and automatically configures settings for you. Or, you can open Accessibility Options in Control Panel to directly customize keyboard, display, and mouse functions.

### Note

- If you are using an input device other than a mouse, instead of clicking, use the appropriate action to complete commands or select items.

**To set your accessibility options by using the Accessibility Wizard**

1. Open the **Accessibility Wizard**.
2. Follow the instructions that appear on your screen.

### Note

- To open the Accessibility Wizard, click **Start**, point to **All programs**, point to **Accessories**, point to **Accessibility**, and then click **Accessibility Wizard**.

**To enable or configure accessibility options by using Control Panel**

- Open **Accessibility Options**.

### Note

- To open Accessibility Options, click **Start**, click **Control Panel**, and then double-click **Accessibility Options**.

You can turn some accessibility features on and off by using keyboard combinations, regardless of whether you have the features enabled. When you use one of the following keyboard combinations, your computer responds with both auditory (beeps) and visual (dialog box) cues that indicate that the accessibility function is on or off.

| To | Press |
|---|---|
| Switch FilterKeys on and off | RIGHT SHIFT for eight seconds |
| Switch High Contrast on and off | LEFT ALT+LEFT SHIFT+PRINT SCREEN |
| Switch MouseKeys on and off | LEFT ALT+LEFT SHIFT+NUM LOCK |
| Switch StickyKeys on and off | SHIFT five times |
| Switch ToggleKeys on and off | NUM LOCK for five seconds |

# Using the keyboard to move the mouse pointer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the keyboard to move the mouse pointer**

If you have trouble using a mouse, you can change the MouseKeys settings so that you can use the numeric keypad to move the mouse pointer. You can do this through the Accessibility Wizard, or you can use the following procedure to set this option in Control Panel.

**To turn MouseKeys on**

1. Open Accessibility Options.
2. On the **Mouse** tab, select the **Use MouseKeys** check box.

The following illustration shows which keys to use to replicate mouse actions.




**Note**

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configuring Windows for people who are blind or have impaired vision

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Windows for people who are blind or have impaired vision**

Using Display in Control Panel, people who have impaired vision or are blind can select larger fonts and icons, increase the size of screen elements by using a lower screen resolution, or change the size of windows and window borders. Using Mouse in Control Panel, you can adjust mouse properties to display larger or more visible mouse pointers and control the speed and animation of the mouse pointer.

When using Windows-based programs, you can change font sizes and colors, background color, and window size. In some programs, you can magnify the elements that are displayed in the program's window.

Using the **Properties** or **Default** dialog boxes in MS-DOS programs, you can change font types, sizes, and colors, background color, and the size of the cursor and program window.

The following tools are available with Windows.

**High Contrast**

Using High Contrast, you can specify that Windows uses colors and fonts that are designed for easy reading.

**Magnifier**

Using Magnifier, you can magnify a portion of your screen to make it more easily readable.

**Narrator**

Narrator is a text-to-speech program that reads aloud the text that is displayed on your screen.

**ToggleKeys**

ToggleKeys causes tones to sound when you press CAPS LOCK, NUM LOCK, and SCROLL LOCK.

The accessibility tools that ship with Windows are intended to provide a minimum level of functionality for users with special needs. Most users with disabilities will need utility programs with more advanced functionality for daily use.

**See Also**

**Other Resources**

Windows XP Professional: Accessibility Options overview

# Configuring Windows for people who have a mobility impairment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Windows for people who have a mobility impairment**

People who have motion-related disabilities can use Windows to configure their computer system to provide a more comfortable computing environment and help minimize the impact of mobility impairments on their work.

Keyboard, part of the Printers and Other Hardware category in Control Panel, provides keyboard layouts for users of alternative keyboards. This is especially useful for those who type with one hand.

Mouse, part of the Printers and Other Hardware category in Control Panel, configures the mouse for right-handed or left-handed users; adjusts double-click speed, pointer speed, and acceleration; and causes the mouse pointer to go directly to default buttons in dialog boxes.

**FilterKeys**

Using FilterKeys, located in Accessibility Options in Control Panel, you can tell Windows to ignore brief or repeated keystrokes.

**MouseKeys**

Using MouseKeys, located in Accessibility Options in Control Panel, you can control the pointer with the numeric keypad on your keyboard.

**On-Screen Keyboard**

On-Screen Keyboard users can type using a pointing device.

**SerialKeys**

SerialKeys, located in Accessibility Options in Control Panel, provides alternative access to keyboard and mouse features.

**StickyKeys**

StickyKeys, located in Accessibility Options in Control Panel, enables you to press the SHIFT, CTRL, ALT, or key, and have it remain active until another key is pressed.

The accessibility tools that ship with Windows are intended to provide a minimum level of functionality for users with special needs. Most users with disabilities will need utility programs with more advanced functionality for daily use. For information about accessibility products and aids for Windows operating systems, see the accessibility page (go to http://www.microsoft.com/ and search for "accessibility") on the Microsoft Web site.

**See Also**

**Other Resources**

Windows XP Professional: Accessibility Options overview

# Getting More Accessibility Information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Getting more accessibility information**

For detailed information about accessibility features and options, see Accessibility Options overview.

For a list of links to information about accessibility features, see Accessibility for People with Disabilities.

For additional information, see:

- Documentation in more accessible formats
- Downloading accessibility notes and utilities
- Other accessibility products and services

**Note**

- To display Accessibility Options Help, click **Start**, click **Run**, and type **hh access.chm**.

# Documentation in more accessible formats

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Documentation in more accessible formats**
Documentation for many Microsoft products is available in alternative formats. If you have difficulty reading or handling printed documentation, you can obtain many Microsoft publications. See Accessibility at the Microsoft Web site.

You can also obtain Microsoft publications from Recording for the Blind & Dyslexic, Inc. RFB&D distributes these documents to registered, eligible members of their distribution service, either on audio cassettes or on floppy disks. The RFB&D collection contains more than 80,000 titles, including Microsoft product documentation and books from Microsoft Press.

For more information, contact Recording for the Blind & Dyslexic at:

Recording for the Blind & Dyslexic, Inc.

20 Roszel Road Princeton,

NJ 08540

Phone: (800) 803-7201 (within the United States and Canada) or (609) 452-0606

Recording for the Blind & Dyslexic Web site

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Downloading accessibility notes and utilities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Downloading accessibility notes and utilities**

If you have a modem or another type of network connection, you can view or download accessibility information from the following network services:

- Accessibility information at the Microsoft Web site.

- Microsoft Internet server at ftp.microsoft.com

  Open the /Softlib folder, and then use the readme or index files to learn about the site. In the index, search for words such as "customizing" and "disabilities." Choose a file and locate it in the /Softlib/Mslfiles folder. On your hard disk, create a folder, then run the file you chose from its current location. Follow the instructions to uncompress the file into the folder you created.

# Other accessibility products and services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Other accessibility products and services**

Other products, services, and resources for people with disabilities are available from Microsoft and other organizations.

Microsoft provides a catalog of accessibility aids that you can use with Microsoft operating systems, including the Windows Server 2003 family operating systems. To obtain this catalog, see Accessibility at the Microsoft Web site.

For information on products and services for the Macintosh operating system for people with disabilities, see the Apple Computer Web site.

The Trace R&D Center at the University of Wisconsin-Madison maintains a Web site with a variety of resources for people with disabilities and developers of assistive technology. For information, contact:

Trace R&D Center

University of Wisconsin-Madison

5901 Research Park Boulevard

Madison, WI 53719-1252

Phone: (608) 262-6966

TTY/TDD: (608) 263-5408

Fax: (608) 262-8848

Trace Center Web site

Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Common Administrative Tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Common administrative tasks**

This section provides information about common administrative tasks:

- Backing up and restoring data
- Changing group memberships
- Checking event logs
- Creating administrative scripts
- Creating logon scripts
- Creating user and group accounts
- Deploying and upgrading software
- Installing a DHCP server
- Installing a domain controller
- Managing applications on a local computer
- Managing applications remotely
- Managing directory replication
- Management tasks for disks and volumes
- File and folder management
- Managing network printers
- Managing servers remotely
- Managing services
- Monitoring network traffic
- Monitoring security-related events
- Monitoring server performance
- Resetting user passwords
- Safeguarding your system
- Scheduling tasks
- Setting up DNS

- Setting up TCP/IP
- Setting user and group security

# Backing up and restoring data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backing up and restoring data**

The Backup utility helps you protect data from accidental loss if your system's hardware or storage media fails. For example, you can use Backup to create a duplicate copy of the data on your hard disk and then archive the data on another storage device. The backup storage medium can be a logical drive such as your hard disk, a separate storage device such as a removable disk, or an entire library of disks or tapes organized and controlled by a robotic changer. If the original data on your hard disk is accidentally erased or overwritten, or becomes inaccessible because of a hard disk malfunction, you can easily restore the data from the archived copy.

Two of the most common tasks are backing up files to a file or tape and restoring files from a file or tape. You can also back up data from the Managing Backup from the command line. For more information about other tasks for backing up and restoring data, see Backup How To ....

**To back up files to a file or tape**

1. Open Backup.

   The Backup or Restore wizard starts by default, unless it is disabled. You can use this wizard or go to the next step to work in **Advanced Mode**.

2. Click the Advanced Mode link on the Backup or Restore wizard.

3. Click the **Backup** tab, and then on the **Job** menu, click **New**.

4. Select the files and folders you want to back up by clicking the box to the left of a file or folder in **Click to select the check box for any drive, folder, or file that you want to back up**.

5. In **Backup destination**, do one of the following:

   - Choose **File** if you want to back up files and folders to a file. This is selected by default.
   - Choose a tape device if you want to back up files and folders to a tape.

6. In **Backup media or file name**, do one of the following:

   - If you are backing up files and folders to a file, type a path and file name for the backup (.bkf) file, or click the **Browse** button to find a file.
   - If you are backing up files and folders to a tape, choose the tape you want to use.

7. Select any backup options you want, such as the backup type and the log file type, by clicking the **Tools** menu, and then clicking **Options**. When you have finished selecting backup options, click **OK**.

8. Click **Start Backup**, and then make any changes to the **Backup Job Information** dialog box.

9. If you want to set advanced backup options such as data verification or hardware compression, click **Advanced**. When you have finished setting advanced backup options, click **OK**.

10. Click **Start Backup** to start the backup operation.

**Caution**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

**Notes**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.
- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
- If the Backup or Restore wizard does not start by default, you can still use it to back up files by clicking **Wizard Mode** on the **Welcome** tab.
- To back up and restore Microsoft SQL Server database files, it is recommended that you use the built-in backup and restore utilities that SQL has. For more information, see the Microsoft SQL Server documentation.
- Some tape drives may not support hardware compression.
- System State data contains most elements of a system's configuration, but it might not include all of the information that you require to recover your system from a failure. Therefore, it is recommended that you back up all boot and system volumes, including the System State, when you back up your system.
- You can only back up the System State data on a local computer. You cannot back up the System State data on a remote computer.
- Backup files usually have the extension .bkf. You can use any extension you like, but it is highly recommended that you use .bkf which has file associations that will ensure that the backup file is recognizable.
- Backup operators and administrators can back up and restore encrypted files and folders without decrypting the files or folders.
- If you have Windows Media Services running on your computer, and you want to back up the files associated with these services, see "Running Backup with Windows Media Services" in the Windows Media Services online documentation. You must follow the procedures outlined in the Windows Media Services online documentation before you can back up or restore files associated with Windows Media Services.

**To restore files from a file or tape**

1. Open Backup.

   The Backup or Restore wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore wizard.

3. Click the **Restore and Manage Media** tab, in **Click to select the check box for any drive, folder, or file that you want to restore**, select the files and folders you want to restore by clicking the box to the left of a file or folder.

4. In **Restore files to**, do one of the following:

   - Click **Original location** if you want the backed up files and folders to be restored to the folder or folders they were in when they were backed up. Skip to step 6.

- Click **Alternate location** if you want the backed up files and folders to be restored to a folder that you designate. This option will preserve the folder structure of the backed up data; all folders and subfolders will appear in the alternate folder you designate.

- Click **Single folder** if you want the backed up files and folders to be restored to a folder that you designate. This option will not preserve the folder structure of the backed up data; the files will appear only in the folder that you designate.

5. If you selected **Alternate location** or **Single folder**, type a path for the folder under **Alternate location**, or click the **Browse** button to find the folder.

6. On the **Tools** menu, click **Options**, click the **Restore** tab, and then do one of the following:

   - Click **Do not replace the file on my computer** if you do not want the restore operation to copy over files that are already on your hard disk.

   - Click **Replace the file on disk only if the file on disk is older** if you want the restore operation to replace older files on your disk with newer files from your backup.

   - Click **Always replace the file on my computer** if you want the restore operation to replace files on your disk regardless of whether the backup files are newer or older.

   - Click **OK** to accept the restore options you have set.

7. Click **Start Restore**.

8. If you want to change any of the advanced restore options, such as restoring security settings and junction point data, click **Advanced**. When you are done setting advanced restore options click **OK**.

9. Click **OK** to start the restore operation.

**Caution**

- Choosing **Always replace the file on my computer** could cause you to lose data if a file you are working on is contained in the backup.

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

**Notes**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- You can also use the Restore wizard to restore files by clicking **Restore Wizard** on the **Tools** menu.

- To back up and restore Microsoft SQL Server database files, it is recommended that you use the built-in backup and restore utilities that SQL has. For more information, see the Microsoft SQL Server documentation.

- In order to restore the System State data on a domain controller, you must first start your computer in Directory Services Restore Mode. This will allow you to restore the SYSVOL directory and the Active Directory directory service database.

- You can only restore the System State data on a local computer. You cannot restore the System State data on a remote computer.

- Backup operators and administrators can back up and restore encrypted files and folders without decrypting the files or folders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Changing group memberships

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Changing group memberships**

A group is a collection of users that you can use to simplify the administration of user permissions and rights. In addition, you can use a group to delegate administrative tasks, filter Group Policy settings, and create e-mail distribution lists. Users belonging to a particular group receive all the permissions and rights assigned to that group. By changing group memberships for a user, you can quickly change the resources to which that user has access, as well as the tasks delegated to the user and the Group Policy settings that apply to the user. You can change the membership of Active Directory groups to change users' permissions and rights within a domain or forest. You can also change the membership of local groups to change users' permissions and rights on a specific computer.

Some of the most common tasks are adding or removing members from Active Directory groups and adding or removing members from groups on a local computer. You can also use the command line to change group memberships, either in a Managing Active Directory from the command line or on a Managing local groups from the command line. For more information about other tasks for managing group memberships in a domain, see Manage Groups. For more information about other tasks for managing group memberships on a local computer, see Manage Local Groups.

**To add or remove a member from an Active Directory group**

1. Open Active Directory Users and Computers.
2. In the console tree, double-click the domain node.
3. Click the folder that contains the group to which you want to add or remove a member.
4. In the details pane, right-click the group, and then click **Properties**.
5. Click the **Members** tab, and then do one of the following:
   - To add a member to a group, click **Add**. In **Enter the object names to select**, type the name of the user, group, or computer that you want to add to the group, and then click **OK**.
   - To remove a member from a group, click the member you want to remove, and then click **Remove**.

### Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- In addition to users and computers, groups can include contacts and other groups.
- To add members to a group, you can click the members you want to add, click the **Add to a group** toolbar button, and then type the name of the group to which they will be added. You can also drag a member object to a group, or right-click the object, and then click **Add to a group** from the shortcut menu.
- When you administer a domain, security principals in the parent domain or other trusted domains are not visible on the **Member Of** tab of a domain users properties. The only domain accounts that you can add or view are the present domain groups. Only domain groups in the present domain are shown, even if the member belongs to other trusted domain groups.

**To add or remove a member from a group on a local computer**

1. Open Computer Management.

2. In the console tree, click **Groups**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Groups

3. Right-click the group to which you want to add or remove a member and click **Properties**.

4. Do one of the following:

   - To add a member to a group, click **Add**. In **Enter the object names to select**, type the name of the user, group, or computer that you want to add to the group, and then click **OK**.

   - To remove a member from a group, click the member you want to remove, and then click **Remove**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- A user who belongs to a group has all the rights and permissions granted to that group. If a user is a member of more than one group, then the user has all the rights and permissions granted to every group he or she belongs to. For more information, see Default local groups.

- You should not add a new user to the Administrators group unless the user will perform only administrative tasks. For more information, see Why you should not run your computer as an administrator.

- If the computer participates in a domain, you can add user accounts and global groups from that domain and from trusted domains to a local group.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Checking event logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checking event logs**

Depending on its server role, a computer running a Windows Server 2003 operating system records events in the following types of logs: application, security, directory service, File Replication service, and DNS server. You can use Event Viewer to monitor these logs and gather information about the hardware, software, and system problems on a computer. Event Viewer displays five types of events within each log: error, warning, information, success audit, and failure audit.

Some of the most common tasks are viewing an event log and connecting to another computer. You can also check event logs from the command line. For more information, see Managing event logs from the command line. For more information about other tasks for checking event logs, see Event Viewer How To....

**To view an event log**

1. Open Event Viewer.
2. In the console tree, click the log you want to view.
3. In the details pane, view the list of individual events.
4. If you want to see more details about a specific event, in the details pane, double-click the event.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials. However, you must be a member of the Administrators group on the target computer to open the Security log, or any log when targeting a remote computer.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- To refresh the view, on the **Action** menu, click **Refresh**. When you open a log, Event Viewer displays the current information for the log. While you view the log, the display is not updated unless you refresh the Event Viewer window. If you switch to another log and then return to the first log, the display for the first log is automatically refreshed. When you refresh the view of an event log, only the view is updated, no changes are made to the log. The **Refresh** command is not available for archived logs, because those files can no longer be updated.

**To connect to another computer**

1. Open Event Viewer.
2. In the console tree, right click **Event Viewer (Local)**, and then click **Connect to another computer**.
3. Click **Another computer**, type the name of the computer, and then click **OK**. You can also find another computer on the network by clicking **Browse**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event**

**Viewer**.

- The other computer can be a workstation running Windows XP Home Edition, Windows XP Professional, Windows 2000 Professional, Windows NT Workstation, or a server or domain controller running Windows NT Server, Windows 2000 Server, or a Windows Server 2003 operating system.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Creating administrative scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating administrative scripts**

With scripts, you can automate a wide range of common administrative tasks, such as:

- Administering users, groups, and organizational units.
- Configuring computers.
- Managing event logs.
- Maintaining files, folders, and disks.
- Monitoring performance, processes, and services.

The Windows Server 2003 family supports two scripting environments:

- The command processor runs files that contain batch language commands. Batch language has limited capabilities, but earlier operating systems supported it and you may need to continue using existing batch files. For more information about batch files, see Using batch files.
- Windows Script Host (WSH) runs files that contain Microsoft Visual Basic Scripting Edition (VBScript) or Jscript commands. VBScript and Jscript provide all the capabilities of batch language plus many more. For example, scripts written in VBScript or Jscript can interact with Active Directory Service Interfaces (ADSI) to manage objects stored in Active Directory, and they can interact with Windows Management Instrumentation (WMI) to access system resources.

In addition to the two scripting languages, WSH provides two runtime programs: WScript.exe and CScript.exe. After you create a WSH script containing VBScript or Jscript commands, you use WScript or CScript to run the script. WScript runs the script as a Windows-based process and CScript runs the script as a console-based process.

Administrators often want to manage objects in Active Directory, such as organizational units, groups, and users. The following example script shows how you can use WSH, VBScript, and ADSI to create an Active Directory organizational unit, group, and user. After creating the three Active Directory objects in the current domain, the script assigns a password to the new user account, enables the user account, and adds the user account to the group.

**To create an example script**

1. Open Notepad.
2. Copy and paste, or type, the following:

```
Set oRoot   = GetObject("LDAP://rootDSE")
Set oDomain = GetObject("LDAP://" &
oRoot.Get("defaultNamingContext"))

Set oOU=oDomain.Create("organizationalUnit",
"ou=Demo OU")
oOU.Put "Description", "Demonstration OU"
oOU.SetInfo

Set oGroup = oOU.Create("Group", "cn=Demo
Group")
```

```
oGroup.Put "sAMAccountName", "DemoGroup"
oGroup.Put "Description", "Demonstration
Group"
oGroup.SetInfo

Set oUser = oOU.Create("User", "cn=Demo User")
oUser.Put "sAMAccountName", "DemoUser"
oUser.Put "Description", "Demonstration User"
oUser.SetInfo

oUser.SetPassword "qW5rty"

oUser.AccountDisabled = False
oUser.SetInfo

oGroup.Add oUser.ADSPath
```

3. On the **File** menu, click **Save As**.

4. In **Save as type**, click **All Files**.

5. In **File name**, type a file name, followed by **.vbs**, and then click **Save**. The file is saved as a VBScript.

### Notes

- To open Notepad, click **Start**, point to **All programs**, point to **Accessories**, and then click **Notepad**.

- The root of the current domain cannot contain an organizational unit named Demo OU.

#### To run the example script

1. Open Command Prompt.

2. Change to the directory where the script is located.

3. Type the following:

   **wscript***ScriptName***.vbs**

   where *ScriptName***.vbs** is the name of the script.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This example starts the script from the command prompt. You can also start the script by clicking **Start**, clicking **Run**, typing **wscript** *ScriptName***.vbs** in **Open**, and then clicking **OK**.

- This example runs the script as a Windows-based process. To run it as a console process, type **cscript** instead of **wscript**.

For more information, see Windows Script Host, Windows Script at the Microsoft Web site, and the *System Administration Scripting Guide* at the Microsoft Windows Resource Kits Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Creating logon scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating logon scripts**

You can use logon scripts to assign tasks that will be performed when a user logs on to a particular computer. The scripts can carry out operating system commands, set system environment variables, and call other scripts or executable programs. The Windows Server 2003 family supports two scripting environments: the command processor runs files containing batch language commands, and Windows Script Host (WSH) runs files containing Microsoft Visual Basic Scripting Edition (VBScript) or Jscript commands. You can use a text editor to create logon scripts. Some tasks commonly performed by logon scripts include:

- Mapping network drives.
- Installing and setting a user's default printer.
- Collecting computer system information.
- Updating virus signatures.
- Updating software.

The following example logon script contains VBScript commands that use Active Directory Service Interfaces (ADSI) to perform three common tasks based on a user's group membership:

1. It maps the H: drive to the home directory of the user by calling the WSH Network object's MapNetworkDrive method in combination with the WSH Network object's UserName property.
2. It uses the ADSI IADsADSystemInfo object to obtain the current user's distinguished name, which in turn is used to connect to the corresponding user object in Active Directory. Once the connection is established, the list of groups the user is a member of is retrieved by using the user's memberOf attribute. The multivalued list of group names is joined into a single string by using VBScript's Join function to make it easier to search for target group names.
3. If the current user is a member of one of the three groups defined at the top of the script, then the script maps the user's G: drive to the group shared drive, and sets the user's default printer to be the group printer.

**To create an example logon script**

1. Open Notepad.
2. Copy and paste, or type, the following:

```
Const ENGINEERING_GROUP     = "cn=engineering"
Const FINANCE_GROUP         = "cn=finance"
Const HUMAN_RESOURCES_GROUP = "cn=human resources"

Set wshNetwork = CreateObject("WScript.Network")
wshNetwork.MapNetworkDrive "h:",
"\\FileServer\Users\" & wshNetwork.UserName

Set ADSysInfo = CreateObject("ADSystemInfo")
Set CurrentUser = GetObject("LDAP://" &
ADSysInfo.UserName)
strGroups = LCase(Join(CurrentUser.MemberOf))
```

```
If InStr(strGroups, ENGINEERING_GROUP) Then

    wshNetwork.MapNetworkDrive "g:",
    "\\FileServer\Engineering\"
    wshNetwork.AddWindowsPrinterConnection
    "\\PrintServer\EngLaser"
    wshNetwork.AddWindowsPrinterConnection
    "\\PrintServer\Plotter"
    wshNetWork.SetDefaultPrinter
    "\\PrintServer\EngLaser"

ElseIf InStr(strGroups, FINANCE_GROUP) Then

    wshNetwork.MapNetworkDrive "g:",
    "\\FileServer\Finance\"
    wshNetwork.AddWindowsPrinterConnection
    "\\PrintServer\FinLaser"
    wshNetWork.SetDefaultPrinter
    "\\PrintServer\FinLaser"

ElseIf InStr(strGroups, HUMAN_RESOURCES_GROUP) Then

    wshNetwork.MapNetworkDrive "g:",
    "\\FileServer\Human Resources\"
    wshNetwork.AddWindowsPrinterConnection
    "\\PrintServer\HrLaser"
    wshNetWork.SetDefaultPrinter
    "\\PrintServer\HrLaser"

End If
```

3. On the **File** menu, click **Save As**.

4. In **Save in**, click the directory that corresponds to the domain controller's Netlogon shared folder (usually *SystemRoot*\SYSVOL\Sysvol*DomainName*\Scripts where *DomainName* is the domain's fully qualified domain name).

5. In **Save as type**, click **All Files**.

6. In **File name**, type a file name, followed by **.vbs**, and then click **Save**. WSH uses the .vbs extension to identify files that contain VBScript commands.

### Notes

- To open Notepad, click **Start**, point to **All programs**, point to **Accessories**, and then click **Notepad**.

- To use the example logon script, you need to change the group names, network drive letters, and Universal Naming Convention (UNC) paths to match your system environment.

- To run a logon script, you need to assign the script to a user or a group. For more information, see Assign a logon script to a user or group.

For more information about creating and using logon scripts, see Logon Scripts, Windows Script at the Microsoft Web site, and the Microsoft Windows Resource Kits Web site.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Creating user and group accounts

Updated: March 25, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating user and group accounts**
User accounts are used to authenticate, authorize or deny access to resources for, and audit the activity of individual users on your network. A group account is a collection of user accounts that you can use to assign a set of permissions and rights to multiple users simultaneously. A group can also contain contacts, computers, and other groups. You can create user accounts and group accounts in Active Directory to manage domain users. You can also create user accounts and group accounts on a local computer to manage users specific to that computer.

Some of the most common tasks are creating user accounts in Active Directory, creating group accounts in Active Directory, creating user accounts on a local computer, and creating groups on a local computer. You can also use the command line to create user and group accounts in Managing Active Directory from the command line or on a Managing local groups from the command line. For more information about other tasks for managing Active Directory user accounts and groups, see Manage Users, Groups, and Computers. For information about other tasks for managing user accounts and groups on a local computer, see Local Users and Groups How To....

**To create a user account in Active Directory**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the folder in which you want to add a user account.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder*

3. Point to **New**, and then click **User**.

4. In **First name**, type the user's first name.

5. In **Initials**, type the user's initials.

6. In **Last name**, type the user's last name.

7. Modify **Full name** to add initials or reverse order of first and last names.

8. In **User logon name**, type the user logon name, click the UPN suffix in the drop-down list, and then click **Next**.

   If the user will use a different name to log on to computers running Windows 95, Windows 98, or Windows NT, then you can change the user logon name as it appears in **User logon name (pre-Windows 2000)** to the different name.

9. In **Password** and **Confirm password**, type the user's password, and then select the appropriate password options.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- To add a user, you can also click

  on the toolbar.

- To add a user, you can also copy any previously created user account.

- A new user account with the same name as a previously deleted user account does not automatically assume the permissions and group memberships of the previously deleted account because the security ID (SID) for each account is unique. To duplicate a deleted user account, all permissions and memberships must be manually recreated.

- When a user account is created with the new user wizard from within the details pane, you can quickly edit the user properties by closing the wizard, clicking the new account, and then pressing ENTER. To open the new user wizard from within the details pane, right-click in the details pane, click **New**, and then click **User**.

- For interoperability with other directory services, you can create an InetOrgPerson user object. To create a new inetOrgPerson, in step three, click **InetOrgPerson** instead of **User**.

- When creating a new user, the full name attribute is created in the *FirstNameLastName* format by default. The full name attribute also governs the display name format is shown in the global address list. You can change the display name format by using ADSI Edit. If you do so, this will also change the full name format. For more information, see article Q250455, "How to Change Display Names of Active Directory Users" in the Microsoft Knowledge Base.

**To create a group account in Active Directory**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the folder in which you want to add a new group.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder*

3. Point to **New**, and then click **Group**.

4. Type the name of the new group.

   By default, the name you type is also entered as the pre–Windows 2000 name of the new group.

5. In **Group scope**, click one of the options.

6. In **Group type**, click one of the options.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- To add a group, you can also click the folder in which you want to add the group, and then click

  on the toolbar.

- If the domain in which you are creating the group is set to the domain functional level of Windows 2000 mixed, you can select only security groups with **Domain local** or **Global** scopes. For more information, see Group scope.

- When a group account is created with the new group wizard from within the details pane, you can quickly edit the group account properties by closing the wizard, clicking the new account, and then pressing ENTER. To open the new group wizard from within the details pane, right-click in the details pane, click **New**, and then click **Group**.

- A group name cannot be identical to any other group name in the domain.

- A group name (pre–Windows 2000) (**samAccountName** object attribute) can contain up to 63 uppercase or lowercase characters.

- A group name (pre–Windows 2000) (**samAccountName** object attribute) cannot consist solely of periods or spaces.

- In Active Directory Users and Computers, by default, the name that you type is also entered as the pre–Windows 2000 name of the new group.

- A group name (CN) can contain up to 256 uppercase or lowercase characters, except for the following:

  " / \ [ ] : ; | = , + * ? < >

- A group name (CN) cannot consist solely of spaces.

- When you use the **net group** command to create a new group account (**net group <group_name> /add /domain**, for example, **net group Group1 /add /domain**), if you specify a group name that is longer than 64 characters, the directory service sets the group's CN to the automatically generated **objectSID** of the newly created group account and the **samAccountName** object attribute assumes the name that you specify in the **net group** command. For more information about **net group** name restrictions, see NetGroupAdd (http://go.microsoft.com/fwlink/?LinkId=159751).

**To create a user account on a local computer**

1. Open Computer Management.

2. In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

3. On the **Action** menu, click **New User**.

4. Type the appropriate information in the dialog box.

5. Select or clear the check boxes for:

   - **User must change password at next logon**
   - **User cannot change password**
   - **Password never expires**
   - **Account is disabled**

6. Click **Create**, and then click **Close**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- A user name cannot be identical to any other user or group name on the computer being administered. It can contain up to 20 uppercase or lowercase characters, except for the following:

  " / \ [ ] : ; | = , + * ? < >

  A user name cannot consist solely of periods (.) or spaces.

- In **Password** and **Confirm password**, you can type a password containing up to 127 characters. However, if the network consists of computers running Windows 95 or Windows 98, consider using passwords no longer than 14 characters. If your password is longer, you may not be able to log on to the network from those computers.

- You should not add a new local user to the local Administrators group unless the user will perform only administrative tasks. For more information, see Why you should not run your computer as an administrator.

**To create a group on a local computer**

1. Open Computer Management.

2. In the console tree, click **Groups**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Groups

3. On the **Action** menu, click **New Group**.

4. In **Group name**, type a name for the new group.

5. In **Description**, type a description of the new group.

6. To add one or more users to a new group, click **Add**.

7. In the **Select Users, Computers, or Groups** dialog box, do the following:

   - To add a user or group account to this group, under **Enter the object names to select**, type the name of the user account or group account that you want to add, and then click **OK**.

   - To add a computer account to this group, click **Object Types**, select the **Computers** check box, and then click **OK**. Under **Enter the object names to select**, type the name of the computer account that you want to add, and then click **OK**.

8. In the **New Group** dialog box, click **Create**, and then click **Close**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- A local group name cannot be identical to any other group or user name on the local computer being administered. It can contain up to 256 uppercase or lowercase characters, except for the following:

  " / \ [ ] : ; | = , + * ? < >

  A group name cannot consist solely of periods (.) or spaces.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

#### Change History

| Date | Revision |
|---|---|
| March 25, 2011 | Corrected the description of the maximum length of the **CN** and **sAMAccountName** attributes for a group object. |

# Deploying and upgrading software

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying and upgrading software**
You can deploy and upgrade software to remote computers in managed environments by using Group Policy Software Installation to assign Windows Installer packages. Windows Installer packages are deployed and managed within a Group Policy object, which is in turn associated with a particular Active Directory container--either a site, a domain, or an organizational unit. You can use Add or Remove programs in Control Panel to install, upgrade, or manage an application on a local computer. You can also use Remote Desktop Connection to install or upgrade an application by using Add or Remove programs on a remote computer. For more information, see Add or Remove programs overview and Remote Desktop Connection.

Some of the most common tasks are deploying software to remote computers, upgrading software on remote computers, and installing or upgrading software on a local computer.

**To deploy software to remote computers**

1. Open Group Policy Object Editor.
2. Do one of the following:
   - To assign software applications to computers, in the console tree, double-click **Computer Configuration**.
   - To assign or publish software applications to users, in the console tree, double-click **User Configuration**.
3. Double-click **Software Settings**, and then click **Software Installation**.

   **Where?**

   - *Group Policy object*/Computer Configuration or User Configuration/Software Settings/Software Installation
4. Right-click **Software Installation**, click **New**, and then click **Package**.
5. Click the Windows Installer package you want to assign, and then click **Open**.
6. In **Deploy Software**, click **Assigned**.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- Group Policy Software Installation is available to computers running Windows 2000, Windows XP Professional, and the Windows Server 2003 family in an Active Directory environment.
- To use Group Policy Software Installation, you must create a new Group Policy object or edit an existing Group Policy object for a site, domain, or organizational unit. You can also link a Group Policy object to a site, domain, or organizational unit in Active Directory Users and Computers or Active Directory Sites and Services. For more information about ways to open Group Policy Object Editor, see Group Policy.
- Many software applications come with Windows Installer packages (.msi files). For more information, see your software manufacturer's documentation.

**To upgrade software on remote computers**

1. Open Group Policy Object Editor.

2. Do one of the following:

   - To upgrade software applications on remote computers, in the console tree double-click, **Computer Configuration**.

   - To upgrade software applications for users, in the console tree, double-click **User Configuration**.

3. Double-click **Software Settings**, and then click **Software Installation**.

   **Where?**

   - *Group Policy object*/Computer Configuration or User Configuration/Software Settings/Software Installation

4. Right-click **Software Installation**, click **New**, and then click **Package**.

5. Click the Windows Installer package that will serve as the upgrade package, and then click **Open**.

6. In **Deploy Software**, click **Assigned**.

7. In the details pane, right-click the Windows Installer package that will function as the upgrade (not the package to be upgraded).

8. Click **properties**, and then click the **Upgrades** tab.

9. Click **Add** to create or add to the list of packages that are to be upgraded by the current package.

10. Under **Choose a package from**, click **Current Group Policy object (GPO)** or **A specific GPO** as the source of the package to be upgraded. If you click **A specific GPO**, click **Browse**, and then click the Group Policy object that you want.

11. Review the list of packages under **Package to upgrade**, which lists all of the other packages that are assigned or published within the selected Group Policy object. Depending on the Group Policy object, this list may have zero or more entries.

12. Click the package that you want to upgrade, and then do one of the following:

    - To replace an application with a completely different application, click **Uninstall the existing package, then install the upgrade package**.

    - To install a newer version of the same product while retaining the user's application preferences, click **Package can upgrade over the existing package**.

13. On the **Upgrades** tab, select the **Required upgrade for existing packages** check box if you want the upgrade to be mandatory. If this is an upgrade under Computer Configuration in the Group Policy console tree, the check box is unavailable and selected, because packages can only be assigned to computers, not published.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- Group Policy Software Installation is available to computers running Windows 2000, Windows XP Professional, and the Windows Server 2003 family in an Active Directory environment.

- To use Group Policy Software Installation, you must create a new Group Policy object or edit an existing Group Policy object for a site, domain, or organizational unit. You can also link a Group Policy object to a site, domain, or organizational unit in Active Directory Users and Computers or Active Directory Sites and Services. For more information about ways to open Group Policy Object Editor, see Group Policy.

- Many software applications come with Windows Installer packages (.msi files). For more information, see your software manufacturer's documentation.

**To install or upgrade software on a local computer by using Add or Remove programs**

1. Open Add or Remove programs in Control Panel.

2. Click **Add New programs**, and then click **CD or Floppy**.

3. Follow the instructions on your screen.

### Notes

- To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**.

- When using Add or Remove programs, you can only install programs that were written for Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Installing a DHCP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing a DHCP server**

All computers on a TCP/IP network must have an IP address in order for the network to properly function. You can configure IP addresses manually at each computer, or you can install a DHCP server that automatically assigns IP address leases to each client computer on the network. Most client operating systems seek an IP address lease by default, so no configuration on the client computer is necessary to implement a DHCP-enabled network.

Before a DHCP server can start leasing IP addresses to client computers, you must create and activate a scope. A scope is a range of possible IP addresses for a network. Before you create a scope, make sure that the IP address range you want to use provides enough IP addresses for all of the computers on your network. Also, determine whether any devices on your network, such as DNS servers, WINS servers, or legacy printers, will need to use static IP addresses. If there are devices that need static IP addresses, create an exclusion range of IP addresses at the beginning of the IP address range. An exclusion range is a group of IP addresses that the DHCP server will not lease to client computers. Once the exclusion range has been defined, you can assign all statically configured devices an IP address from the exclusion range.

Some of the most common tasks are installing a DHCP server, creating a scope, and activating a scope. You can also perform these tasks from the Managing DHCP from the command line. For more information about other tasks for managing DHCP, see DHCP How To....

**To install a DHCP server**

1. Open Windows Components wizard.
2. Under **Components**, scroll to and click **Networking Services**.
3. Click **Details**.
4. Under **Subcomponents of Networking Services**, click **Dynamic Host Configuration Protocol (DHCP)**, and then click **OK**.
5. If prompted, type the full path to the Windows Server 2003 distribution files, and click **Continue**.

   Required files are copied to your hard disk.

### Notes

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- DHCP servers must be configured with a static IP address.

**To create a new scope**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **New Scope**.

4. Follow the instructions in the New Scope wizard.

### Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Depending on the starting and ending IP addresses for your scope, the DHCP console suggests a default subnet mask useful for most networks. If you know a different subnet mask is required for your network, you can modify the value as needed.
- When you finish creating a new scope, you might need to complete additional tasks, such as activating the scope for use or assigning scope options.

#### To activate a scope

1. Open DHCP.
2. In the console tree, click the applicable scope.

   **Where?**

   - DHCP/*Applicable DHCP server*/*Applicable superscope* (if used)/*Applicable scope*

3. On the **Action** menu, click **Activate**.

### Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Activating a scope is needed only for new scopes to start lease distribution. You must activate a scope to make it available for DHCP clients.
- Do not activate a scope until you specify the options you want.
- The **Action** menu command changes to **Deactivate** when the selected scope is currently activated. Do not deactivate a scope unless you are retiring it permanently from use on your network.
- If you want to use DHCP to assign IP address leases to clients on a subnet that is separated from your DHCP server by a router, you must have DHCP message forwarding enabled on the router. If you are using a computer running a member of the Windows Server 2003 family as a router, you can use the DHCP Relay Agent service provided with the Routing and Remote Access service. For more information, see DHCP/BOOTP Relay Agents and Configure the DHCP Relay Agent.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Installing a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing a domain controller**

Domain controllers provide network users and computers with the Active Directory directory service, which stores and replicates directory data and manages user interactions with the domain, including user logon processes, authentication, and directory searches. Every domain must contain at least one domain controller. You install a domain controller by installing Active Directory on any member or stand-alone server (except those with restrictive license agreements).

When you install the first domain controller in your organization, you are creating the first domain (also called the root domain) and the first forest. You can add additional domain controllers to an existing domain to provide fault tolerance, improve service availability, and balance the load of existing domain controllers.

You can also install a domain controller to create a new child domain or new domain tree. Create a new child domain when you want a new domain that shares a contiguous namespace with one or more domains. This means that the name of the new domain contains the full name of the parent domain. For example, sales.microsoft.com would be a child domain of microsoft.com. Create a new domain tree only when you need a domain whose Domain Name System (DNS) namespace is not related to the other domains in the forest. This means that the name of the new domain tree's root domain (and all of its children) does not contain the full name of the parent domain. A forest can contain one or more domain trees.

Before installing a new domain controller, you will need to consider pre-Windows 2000 compatible security levels and identify the DNS name of the domain. For more information, see Checklist: Creating an additional domain controller in an existing domain.

The most commonly performed tasks when installing a domain controller are creating a new domain in a new forest, creating a new child domain in an existing domain tree, creating a new domain tree in an existing forest, and adding a domain controller to an existing domain.

For information about important decisions you need to make when installing a domain controller, see Using the Active Directory Installation Wizard.

**To create a new domain in a new forest**

1. Open the Active Directory Installation Wizard.
2. On the **Domain Controller Type** page, click **Domain controller for a new domain**, and then click **Next**.
3. On the **Create New Domain** page, click **Domain in a new forest**, and then click **Next**.
4. On the **New Domain Name** page, type the full DNS name for the new domain, and then click **Next**.
5. On the **NetBIOS Domain Name** page, verify the NetBIOS name, and then click **Next**.
6. On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location, and then click **Next**.
7. On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location, and then click **Next**.
8. On the **DNS Registration Diagnostics** page, verify if an existing DNS server will be authoritative for this forest or, if necessary, choose to install and configure DNS on this server by clicking **Install and configure the DNS server on this computer, and set this computer to use this DNS server as its preferred DNS server**, and then click **Next**.
9. On the **Permissions** page, select one of the following:

    - **Permissions compatible with pre-Windows 2000 Server operating systems**

- **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**

10. Review the **Summary** page, and then click **Next** to begin the installation.

11. Restart the computer.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- The server on which you perform this procedure will be promoted to the first domain controller in the forest root domain.

- The wizard options on the **Permissions** page affect application compatibility with pre-Windows 2000 and Windows Server 2003 operating systems and are not related to domain functionality.

- The Active Directory Installation Wizard allows Active Directory domain names up to 64 characters or up to 155 bytes. Although the limit of 64 characters is usually reached before the limit of 155 bytes, the opposite could be true if the name contains Unicode characters consuming three bytes. These limits do not apply to computer names.

**To create a new child domain in an existing domain tree**

1. Open the Active Directory Installation Wizard.

2. On the **Domain Controller Type** page, click **Domain controller for a new domain**, and then click **Next**.

3. On the **Create New Domain** page, click **Child domain in an existing domain tree**, and then click **Next**.

4. On the **Network Credentials** page, type the user name, password, and user domain of the user account you want to use for this operation, and then click **Next**.

5. On the **Child Domain Installation** page, verify the parent domain and type the new child domain name, and then click **Next**.

6. On the **NetBIOS Domain Name** page, verify the NetBIOS name, and click **Next**.

7. On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location, and then click **Next**.

8. On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location, and then click **Next**.

9. On the **DNS Registration Diagnostics** page, verify the DNS configuration settings are accurate, and then click **Next**.

10. On the **Permissions** page, select one of the following:

    - **Permissions compatible with pre-Windows 2000 Server operating systems**

    - **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**

11. On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password you want to assign to this server's Administrator account that will be used when the computer is started in Directory Services Restore

Mode, and then click **Next**.

12. Review the **Summary** page, and then click **Next** to begin the installation.

13. Restart the computer.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- The server on which you perform this procedure will be promoted to the first domain controller in a new child domain.

- When a child domain is added to an existing tree domain, a two-way, transitive, parent and child trust is established by default.

- The wizard options on the **Permissions** page affect application compatibility with pre-Windows 2000 and Windows Server 2003 operating systems and are not related to domain functionality.

- The Active Directory Installation Wizard allows Active Directory domain names up to 64 characters or up to 155 bytes. Although the limit of 64 characters is usually reached before the limit of 155 bytes, the opposite could be true if the name contains Unicode characters consuming three bytes. These limits do not apply to computer names.

#### To create a new domain tree in an existing forest

1. Open the Active Directory Installation Wizard.

2. On the **Domain Controller Type** page, click **Domain controller for a new domain**, and then click **Next**.

3. On the **Create New Domain** page, click **Domain tree in an existing forest**.

4. On the **Network Credentials** page, type the user name, password, and user domain of the user account you want to use for this operation, and then click **Next**.

5. On the **New Domain Tree** page, type the full DNS name for the new domain, and then click **Next**.

6. On the **NetBIOS Domain Name** page, verify the NetBIOS name, and then click **Next**.

7. On the **Database and Log Folders** page, type the location in which to install the database and log folders, or click **Browse** to choose a location, and then click **Next**.

8. On the **Shared System Volume** page, type the location in which to install the Sysvol folder, or click **Browse** to choose a location, and then click **Next**.

9. On the **DNS Registration Diagnostics** page, verify if an existing DNS server will be authoritative for this forest or, if necessary, choose to install and configure DNS on this server by clicking **Install and configure the DNS server on this computer, and set this computer to use this DNS server as its preferred DNS server**, and then click **Next**.

10. On the **Permissions** page, select one of the following:

    - **Permissions compatible with pre-Windows 2000 Server operating systems**

    - **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**

11. On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password you want to assign this server's Administrator account that will be used when the computer is started in Directory Services Restore Mode, and then click **Next**.

12. Review the **Summary** page, and then click **Next** to begin the installation.

13. Restart the computer.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- The server on which you perform this procedure will be promoted to the first domain controller in a new domain tree.

- When a new domain tree is created in an existing forest, a two-way, transitive, tree root trust is established by default.

- The wizard options on the **Permissions** page affect application compatibility with pre-Windows 2000 and Windows Server 2003 operating systems and are not related to domain functionality.

- The Active Directory Installation Wizard allows Active Directory domain names up to 64 characters or up to 155 bytes. Although the limit of 64 characters is usually reached before the limit of 155 bytes, the opposite could be true if the name contains Unicode characters consuming three bytes. These limits do not apply to computer names.

**To install an additional domain controller in an existing domain**

1. Open the Active Directory Installation Wizard.

2. On the **Domain Controller Type** page, click **Additional domain controller for an existing domain** , and then click **Next** .

3. On the **Network Credentials** page, type the user name, password, and user domain of the user account you want to use for this operation, and then click **Next** . See Note below for more information.

4. On the **Additional Domain Controller** page, enter the full DNS name of the existing domain for which the server will become a domain controller, and then click **Next** .

5. On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location, and then click **Next** .

6. On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location, and then click **Next** .

7. On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password you want to assign to this server's Administrators account that will be used when the computer is started in Directory Services Restore Mode, and then click **Next**.

8. Review the **Summary** page, and then click **Next** to begin the installation.

9. Restart the computer.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using

Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To create an additional domain controller from restored backup files, start the Active Directory Installation Wizard by typing **dcpromo /adv** at a command prompt.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Managing applications on a local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing applications on a local computer**
Managing applications involves adding or removing applications, changing the configuration of an application, or scheduling an application to run at a specific time. You can use Add or Remove Programs in Control Panel to add or remove an application, modify an existing application, repair a damaged application, or add or remove components of Windows. You can use Task Scheduler or the **at** command to schedule an application to run at a specific time.

You can also update your operating system and other Microsoft software by using Windows Update in Help and Support Center to download items such as security fixes, critical updates, the latest Help files, drivers, and Internet products. For more information, see Windows Update.

Some of the most common tasks are adding, removing, or changing an application and scheduling an application to run at a specific time. You can also run an application at a scheduled time from the command line by using the At command. For more information about other tasks for managing applications, see Managing applications remotely.

**To add, remove, or change an application**

1. Open Add or Remove Programs in Control Panel.

2. Do one of the following:

    - To add a new application from a CD or floppy disk, click **Add New Programs**, click **CD or Floppy**, and then follow the instructions on your screen.

    - To add a new application from your network, click **Add New Programs**, click the program you want to add, and then click **Add**. Follow the instructions on your screen.

    - To change or remove an application, click **Change or Remove Programs**, click the program you want to change or remove, and then click **Change/Remove**. Follow the instructions on your screen.

### Caution

- Some programs do not provide Change options and may be removed without prompting you further.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**.

- Only applications written for Windows operating systems can be used with Add or Remove Programs.

- Many software applications provide tools for setting configuration properties or adding or removing components. You will often see these tools listed in an application's program group on the **Start** menu or in Control Panel. For more information, see the software manufacturer's documentation.

**To schedule an application to run at a specific time**

1. Open Task Scheduler.

2. Double-click **Add Scheduled Task**.

3. Follow the instructions in the Scheduled Task Wizard.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or Server Operators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.

- For more information, see Task Scheduler.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Managing applications remotely

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing applications remotely**

You can manage the applications for an entire system from a remote location by using the remote administration features of the Windows Server 2003 family. With the Remote Desktops snap-in, you can access remote computers for installations, deletions, or modifications of applications on those computers. In addition, with Group Policy Software Installation, you can assign or publish applications. You assign an application to users or computers when you want everyone to have the application on his or her computer. When the users log on to their computers, the application installs. You publish an application to users when you want the application to be available to the users who are managed by your Group Policy object. These users then determine when to install the application.

Some of the most common tasks are creating a new connection to a remote computer, assigning an application to users or computers, and publishing an application to users. For more information about other tasks for managing applications remotely, see Remote Desktops snap-in How To....

**To create a new connection to a remote computer**

1. Open Remote Desktops snap-in.
2. In the console tree, right-click **Remote Desktops**.
3. Click **Add new connection**.
4. In the **Add New Connection** dialog box, in **Server name or IP address**, type the name or TCP/IP address of the server. Or, click **Browse** to look for a server.
5. Optionally, in **Connection Name**, type a friendly name for this connection.
6. Clear the **Connect to console** check box if you do not want to connect to the console of the server.
7. Optionally, in **Logon information**, supply the user name, password and domain to automatically log on to the server. Click **OK**.

### Notes

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.
- Depending on the configuration of the Terminal server, you might be prompted to supply your password upon connection even if you specify it in the **Add New Connection** dialog box.

**To assign an application to users or computers**

1. Open Group Policy Software Installation.
2. In the console tree, click **Software Installation**.

   **Where?**

   - *Group Policy object*/Computer Configuration (or User Configuration)/Software Settings/Software Installation

3. In the details pane, right-click, point to **New**, and then click **Package**.

4. In the **Open** dialog box, click the Windows Installer package, and then click **Open**.

5. In the **Deploy Software** dialog box, click **Assigned**, and then click **OK**.

### Notes

- To open Group Policy Software Installation, see Group Policy (pre-GPMC).

- The **Open** dialog box shows the packages that are located at a software distribution point that you specify as the default.

- If the Windows Installer package is located on a different network drive, navigate to the correct software distribution point in the **Open** dialog box, click the package, and then click **Open**.

#### To publish an application to users

1. Open Group Policy Software Installation.

2. In the details pane, right-click, point to **New**, and then click **Package**.

3. In the **Open** dialog box, click the Windows Installer package to be published, and then click **Open**.

4. In the **Deploy Software** dialog box, click **Published**.

### Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Software Installation, see Group Policy (pre-GPMC).

- Packages can be published only to users, not to computers.

- The **Open** dialog box shows the packages that are located at a software distribution point that you specify as the default.

- If the Windows Installer package is located on a different network drive, navigate to the correct software distribution point in the **Open** dialog box, click the package, and then click **Open**.

- The application is available for users to install either by using Add or Remove Programs in Control Panel or by opening a file with an extension that you have associated with the application.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Managing directory replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing directory replication**
Except for very small networks, directory data must reside in more than one place on the network to be equally useful to all users. Through replication, the Active Directory directory service maintains replicas of directory data on multiple domain controllers, ensuring directory availability and performance for all users. Active Directory relies on configuration information that you provide about sites, subnets, and site links to manage and optimize the process of replication.

Some of the most common tasks are creating sites, creating site links, and creating a subnet and associating it with a site. For more information about other tasks for managing directory replication, see Manage Sites.

**To create a site**

1. Open Active Directory Sites and Services.
2. Right-click the **Sites** folder, and then click **New Site**.
3. In **Name**, type the name of the new site.
4. Click a site link object, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- For more information about sites, see Sites overview.

**To create a site link**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click the intersite transport protocol you want the site link to use, and then click **New Site Link**.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*inter-site transport protocol you want the site link to use*

3. In **Name**, type the name to be given to the link.
4. Click two or more sites to connect, and then click **Add**.
5. Configure the site link's cost, schedule, and replication frequency.

### Caution

- If you create a site link that uses SMTP, you must have an enterprise certification authority (Enterprise CA) available and

SMTP must be installed on all domain controllers that will use the site link.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- Creating a site link between two or more sites is a way to influence replication topology. By creating a site link, you provide Active Directory with information about what connections are available, which ones are preferred, and how much bandwidth is available. Active Directory uses this information to choose times and connections for replication that will afford the best performance.
- For more information about configuring the site link's cost, schedule, and replication frequency, see Configure Replication Between Sites and Managing replication.

#### To create a subnet and associate it with a site

1. Open Active Directory Sites and Services.
2. In the console tree, double-click **Sites**.
3. Right-click **Subnets**, and then click **New Subnet**.
4. In **Address**, enter the subnet address.
5. In **Mask**, enter the subnet mask that describes the range of addresses included in this site's subnet.
6. Select a site with which to associate this subnet, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- Any IP address within the subnet range is a valid entry.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Management tasks for disks and volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing disks and volumes**
Managing disks and volumes includes creating and formatting partitions, logical drives, and volumes; setting disk quotas to limit disk usage; defragmenting volumes to improve file-system performance; and checking for file-system errors and bad sectors on a hard disk. The Windows Server 2003 family provides many tools you can use to effectively manage disks and volumes on new or existing systems. These tools include Disk Management, Disk Defragmenter, disk quotas, and error checking.

Some of the most common tasks are creating partitions or logical drives, formatting basic volumes, extending basic volumes, and defragmenting volumes. You can also manage disks and volumes from the Managing disks and volumes from the command line. For information about disk quotas, see Disk Quotas. For information about error checking, see Detecting and Repairing Disk Errors. For more information about other tasks for managing disks and volumes, see Disk Managment How To....

**To create a partition or logical drive**

1. Open Computer Management (Local).

2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.

3. Right-click an unallocated region of a basic disk, and then click **New Partition**, or right-click free space in an extended partition, and then click **New Logical Drive**.

4. In the New Partition wizard, click **Next**, click **Primary partition**, **Extended partition**, or **Logical drive**, and then follow the instructions on your screen.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- You can create primary partitions, extended partitions, and logical drives only on basic disks. You should create basic volumes instead of dynamic volumes if this computer also runs MS-DOS, Windows 95, Windows 98, Windows Millinnium Edition, Windows NT 4.0, or Windows XP Home Edition.

- On a master boot record (MBR) disk, you can create up to four primary partitions, or three primary partitions, one extended partition, and unlimited logical drives.

- On a GUID partition table (GPT) disk, you can create up to 128 primary partitions.

**To format a basic volume**

1. Open Computer Management (Local).

2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.

3. Right-click the partition, logical drive, or basic volume you want to format (or reformat), and then click **Format**.

4. Select the options you want, and then click **OK**.

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- You cannot format system, boot, OEM, or unknown partitions.

- File compression is supported only on NTFS volumes with cluster sizes 4 KB and smaller.

- If you select the **Perform a quick format** check box, the files are removed from the disk but the disk is not scanned for bad sectors. Use this option only if this disk has been previously formatted and you are sure the disk is not damaged.

**To extend a basic volume**

1. Open Command Prompt.

2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the basic volume you want to extend.

4. At the **DISKPART** prompt, type:

   **select volume***n*

   Selects the basic volume, *n*, you want to extend into contiguous, empty space on the same disk.

5. At the **DISKPART** prompt, type:

   **extend** [**size=***n*]

   Extends the selected volume by size=*n* megabytes (MB).

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command**

**prompt**.

- To extend a basic volume, it must either not be formatted with a file system or it must be formatted with the NTFS file system.
- You can only extend a basic volume onto the same disk.
- You can only extend a basic volume if it is followed by contiguous unallocated space.
- For more information about Disk Management and the **diskpart** command, see Disk Management and DiskPart.

**To defragment a volume**

1. Open Disk Defragmenter.
2. Click the volume that you want to defragment, and then click **Defragment**.

   After the defragmentation is complete, Disk Defragmenter displays the results in **Estimated disk usage after defragmentation**.
3. Click **View Report** to view the Defragmentation Report, which displays detailed information about the volume you defragmented.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Disk Defragmenter, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Disk Defragmenter**.
- You should analyze volumes before defragmenting them. This tells you whether you need to take the time to defragment them.
- A volume must have at least 15% free space for Disk Defragmenter to completely and adequately defragment it. Disk Defragmenter uses this space as a sorting area for file fragments. If a volume has less than 15% free space, Disk Defragmenter will only partially defragment it. To increase the free space on a volume, delete unneeded files or move them to another disk.
- You cannot defragment volumes that the file system has marked as dirty, which indicates possible corruption. You must run **chkdsk** on a dirty volume before you can defragment it. You can determine if a volume is dirty by using the **fsutil dirty query** command. For more information about the **chkdsk** command, see chkdsk. For more information about the **fsutil dirty** command, see Fsutil: dirty.
- The time that Disk Defragmenter takes to defragment a volume depends on several factors, including the size of the volume, the number and size of files on the volume, the percentage of fragmentation in the volume, and available system resources.
- You can defragment only local file system volumes, and you can run only one instance of Disk Defragmenter at a time.
- To stop or temporarily interrupt defragmenting a volume, click **Stop** or **Pause**, respectively.
- If you start Disk Defragmenter while performing a backup on the same volume, Disk Defragmenter stops.
- Running the **defrag** command and Disk Defragmenter are mutually exclusive. If you are using Disk Defragmenter to

defragment a volume and you run the **defrag** command at a command prompt, the **defrag** command fails. Conversely, if you run the **defrag** command, and open Disk Defragmenter, the defragmentation options in Disk Defragmenter are unavailable.

- For more information about Disk Defragmenter and the **defrag** command, see Disk Defragmenter and Defrag.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# File and folder management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing files and folders**
Managing files and folders includes storing and securing resources, making those resources available to network users, and managing changes in those resources. The Windows Server 2003 family provides many tools you can use to manage files and folders. These tools include Shared Folders, shadow copies of shared folders, Distributed File System (DFS), Encrypting File System (EFS), and Offline Files. When a folder is shared, users can connect to the folder over the network and gain access to the contents of the shared folder. With shadow copies of shared folders, users can view the contents of network folders as they existed at points of time in the past.

Some of the most common tasks are sharing a folder or drive, enabling shadow copies of shared folders, and changing settings for shadow copies of shared folders. You can also manage files and folders from the Managing shared folders from the command line. For more information about other tasks for managing files and folders, see Shared Folders How To....

**To share a folder or drive**

1. Open Computer Management.

2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. On the **Action** menu, click **New Share**.

4. Follow the steps in the Share a Folder Wizard, and then click **Finish**.

**Important**

- When you create a shared resource, determine how to secure that resource. There are several methods you can use to secure shared resources. For more information, see Share permissions.

**Notes**

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**To enable shadow copies of shared folders**

1. Open Computer Management (Local).

2. In the console tree, right-click **Shared Folders**, point to **All Tasks**, and then click **Configure Shadow Copies**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Select the volume where you want to enable shadow copies of shared folders, and then click **Enable**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can also access shadow copies of shared folders by right-clicking on a fixed NTFS volume in **My Computer** or in **Disk Management**, clicking **Properties**, and then clicking the **Shadow Copies** tab.
- You can also use Disk Management to manage shadow copies of shared folders on remote servers. On the **Action** menu, point to **All Tasks**, and then click **Configure Shadow Copies**.
- Enabling shadow copies of shared folders also creates a default schedule and settings for subsequent shadow copies. To make changes to these defaults, click **Settings**
- Shadow copies are enabled on a per volume basis; you cannot select specific shares on a volume to be copied or not copied.

**To change settings for shadow copies of shared folders**

1. Open Computer Management (Local).
2. In the console tree, right-click **Shared Folders**, point to **All Tasks**, and then click **Configure Shadow Copies**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Select the volume where you want to make changes, and then click **Settings**.
4. In this dialog box, you can change the settings described in the following table.

| Setting | Notes |
|---|---|
| **Storage volume** | This indicates where you would like to store the shadow copies of the selected volume. The default is to use the same volume. It is recommended that you use a separate volume on another disk which provides better performance for heavily used file servers. You can only change the storage volume when there are no shadow copies present. If you need to change the storage volume for a volume that already has been enabled, you must delete all of the shadow copies on that volume, then change the storage volume. |
| **Details** | Clicking **Details** will open a dialog box that lists the shadow copies that are currently stored. This dialog box also gives information about the total space and available storage space on the disk. |
| **Storage limits** | The default size is 10 percent of the size of the volume that holds the source files being shadow copied. If the shadow copies are stored on a separate volume from the source files, you should change this default to reflect the amount of the storage volume that you are willing to dedicate to shadow copies. The storage limit must be at least 100 MB, which will only allow for a single shadow copy to be retained. If you set a restrictive storage limit, you should test to ensure that the number of shadow copies you have scheduled will fit within those restraints. If shadow copies are deleted prematurely due to storage limits, you may be defeating the purpose of enabling shadow copies of shared folders. |

| | |
|---|---|
| **Schedule** | Clicking **Schedule** will open the task scheduler with the information needed to create a task schedule for taking shadow copies of shared folders on a regular basis. Before creating the schedule, take a look at your current users work patterns and design a strategy that will schedule shadow copies at a time of day that will work best for your users. The default schedule is Monday through Friday at 7:00 A.M. and 12:00 P.M. For more information about creating a schedule for shadow copies of shared folders, see Defining a schedule for shadow copies. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Managing network printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing network printers**
To create a secure and effective network printing environment, you must consider variables such as what client operating systems are used on your network, what printing permissions you want to assign to different users, and what procedures you will use to manage the printers once they are connected to the network. The easiest way to connect and manage network printers is through Active Directory. You can also use Group Policy to change the default behavior of the printing environment and to provide computers and users a standard set of preferences.

Some of the most common tasks are publishing a printer in Active Directory, remotely managing printers, setting Group Policy for printers, and setting or removing permissions for a printer. You can also manage network printers from the Managing printing from the command line. For more information about other tasks for managing network printers, see Printing How To....

**To publish a printer in Active Directory**

1. Open Printers and Faxes.
2. Right-click the printer you want to publish, and then click **Sharing**.
3. On the **Sharing** tab, click **Share this printer**, and then type a name for the shared printer.
4. Select the **List in the Directory** check box to publish the printer in Active Directory.

### Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- Only a shared printer can be published. If you stop sharing a printer, the printer is no longer published.
- If you add a printer using the Add Printer Wizard and click **Share this printer**, the printer will be published by default if the Group Policy settings **Automatically publish new printers in Active Directory** and **Allow printers to be published** are enabled, as they are by default.
- You must open Active Directory Users and Groups to publish printers that are managed through print servers that are running Windows NT 4.0 or an earlier operating system.
- Printers you add using the Add Printer Wizard are shared by default if you add them from a computer running a product in the Windows Server 2003 family. However, printers are not shared by default if you add them using the Add Printer Wizard from a computer running Windows XP.
- You must have the Manage Printers permission for the printer you want to share or to publish.
- You can create printer names that contain spaces, special characters, and more than eight characters. However, some clients do not recognize or correctly handle such names. The entire qualified name of a shared printer (for example \\PRINTER2\PSCRIPT) must contain 32 or fewer characters.

**To remotely manage printers**

1. Double-click **My Network Places**, and then locate the print server for the printers you want to manage.
2. Double-click the print server, double-click the **Printer folder** icon on that server, and then click a printer.

   **Important**

- To facilitate stronger network security, remote printer management is not available by default. To enable remote printer management, in Group Policy, you must enable the Allow Print Spooler to Accept Client Connections policy.

3. Change the print server, printer, or printing preference settings as required.

### Notes

- You can change printer settings only if you have Manage Printers permission.
- If you are having trouble finding your printer and you are logged on to a Windows Server 2003 domain, you can find your printer by clicking **Start**, pointing to **Search**, and then clicking **Printers**.
- You can also drag the Printers and Faxes folder of any print server (or specific printers from any Printers and Faxes folder) to the Printers and Faxes folder on your computer. This is an easy way to manage any printer without having to search for it.

#### To set Group Policy for printers

1. Start Group Policy according to the object you want to set printer policy to.
2. After selecting the properties page of the object you want to set printer policy to, select the Group Policy node.
    - If you want to set policies that apply only to computers, expand the **Computer Configuration** node, and then expand **Administrative Templates**.
    - If you want to set policies that apply only to users, expand the **User Configuration** node, expand **Administrative Templates**, and then expand **Control Panel**.
3. Double-click **Printers** to open a listing of policies.
4. Double-click the printer policy you want to set.
5. On the **Policy** tab, enable or disable the policy by selecting or clearing the appropriate radio button. With some policies, you might need to enter additional information.

### Note

- If you do not want to change the current state of the policy setting, leave it as it is (not configured) to save processing time.

#### To set or remove permissions for a printer

1. Open Printers and Faxes.
2. Right-click the printer for which you want to set permissions, click **Properties**, and then click the **Security** tab.
3. Do one of the following:
    - To change or remove permissions from an existing user or group, click the name of the user or group.
    - To set up permissions for a new user or group, click **Add**. In **Select Users, Computers, or Groups**, type the name of the user or group you want to set permissions for, and then click **OK** to close the dialog box.

4. In **Permissions**, click **Allow** or **Deny** for each permission you want to allow or deny, if necessary. Or, to remove the user or group from the permissions list, click **Remove**.

### Notes

- To change device settings, you must have the Manage Printers permission. For information about printing security permissions, see Assigning printer permissions.
- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- To view or change the underlying permissions that make up Print, Manage Printers, and Manage Documents, click the **Advanced** button.
- A printer must be shared in order for the permission settings to affect the users and groups listed.
- You can also view the permissions assigned to you by clicking the group you belong to on the **Security** tab. For information on finding out what group you belong to, see Identify members of a local group.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Managing servers remotely

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing servers remotely**
Windows Server 2003 operating systems provide tools that you can use to manage servers from a remote location. These tools expand your flexibility because you can work as though you are physically present at each computer in your organization. The tools include the Remote Desktops snap-in, Active Directory Users and Computers, Web Interface for Remote Administration, and the Windows Server 2003 Administration Tools Pack. By understanding the advantages and security requirements of each tool, you can choose the most appropriate one for your remote administration and management tasks. For example, if you want to remotely manage a group of servers instead of just one server, Active Directory is the preferred tool to use. Each server that you want to manage remotely must be enabled for remote administration.

Some of the most common tasks are managing servers remotely by using the Remote Desktops snap-in and managing servers remotely by using Active Directory Users and Computers. For more information on installing and using the Windows Server 2003 Administration Tools Pack to manage servers remotely from Windows XP Professional, see Windows Server 2003 Administration Tools Pack. For more information about using Web Interface for Remote Administration, see Using Web Interface for Remote Administration. For more information about other tasks for managing servers remotely, see Remote Desktop for Administration How To....

**To manage servers remotely by using the Remote Desktops snap-in**

1. Open Remote Desktops snap-in.
2. In the console tree, right-click **Remote Desktops**.
3. Click **Add new connection**.
4. In the **Add New Connection** dialog box, in **Server name or IP address**, type the name or TCP/IP address of the server. Or, click **Browse** to look for a server.
5. Optionally, in **Connection Name**, type a friendly name for this connection.
6. Clear the **Connect to console** check box if you do not want to connect to the console of the server.
7. Optionally, in **Logon information**, supply the user name, password and domain to automatically log on to the server. Click **OK**.

### Important

- To complete this procedure, the Remote Desktop option must be enabled on each remote computer.

### Notes

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.
- Depending on the configuration of the Terminal server, you might be prompted to supply your password upon connection even if you specify it in the **Add New Connection** dialog box.

**To manage servers remotely by using Active Directory Users and Computers**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Computers**.

**Where?**

- Active Directory Users and Computers/*domain node*/Computers

Or, click the container that contains the computer you want to manage.

3. In the details pane, right-click the computer, and then click **Manage**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- Clicking **Manage** starts Computer Management from which you can administer remote computers. You must have administrative credentials on the local computer to view certain information or to modify computer properties using Computer Management.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Managing services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing services**
A service is an application type that runs in the background and is similar to a UNIX daemon application. Services typically provide access to key features such as file servers, Web servers, database servers, and other server-based applications to users, both locally and across the network.

Some of the most common tasks associated with managing services are starting and stopping a service, disabling a service for a hardware profile, and changing the startup method for a service. For more information about other tasks for managing services, see Services How To....

To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled. You should not change the startup method of a service unless you are sure you are choosing the appropriate startup method. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services.

**To start, stop, pause, resume, or restart a service**

1. Open Services.

2. In the details panel, do one of the following:

    - Click the service, and then, on the **Action** menu, click **Start**, **Stop**, **Pause**, **Resume**, or **Restart**.

    - Right-click the service, and then click **Start**, **Stop**, **Pause**, **Resume**, or **Restart**.

### Caution

- If you stop, start, or restart a service, any dependent services are also affected. Starting a service does not automatically restart its dependent services. For more information, see View service dependencies.

- Changing the default service settings might prevent key services from running correctly. It is especially important to use caution when changing the **Startup type** and **Log on as** settings of services that are configured to start automatically. In most cases, it is recommended that you not change the **Allow service to interact with desktop** setting.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.

- To start a service with startup parameters, right-click the service, click **Properties**, and then type the parameters in **Start parameters** before you click **Start**. These settings are not persistent; they are are used only once, and then the default settings are restored. (A backslash (\) is treated as an escape character; type two backslashes for each backslash in a parameter.)

**To enable or disable a service for a hardware profile**

1. Open Services.

2. In the details pane, right-click the service that you want to enable or disable, and then click **Properties**.

3. On the **Log On** tab, click the hardware profile that you want to configure.

4. Click **Enable** or **Disable**, and then click **OK**.

### Caution

- Changing the default service settings might prevent key services from running correctly. It is especially important to use caution when changing the **Startup type** and **Log on as** settings of services that are configured to start automatically. In most cases, you should not change the **Allow service to interact with desktop** setting.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.

- You may want to set up a hardware profile to test any newly installed services. For example, you can create two hardware profiles for a newly installed service: one profile in which the service is enabled and a second profile in which the service is disabled. This way, you can troubleshoot any problems that may occur, such as a driver not loading properly.

- Use caution when disabling services by using this method. If you disable a service for a certain hardware profile, the **Hardware Profile** settings override the **Startup type** setting for that service.

- Use the options in the **System Properties** dialog box to create and manage hardware profiles.

- If you receive a **Configuration Manager** error message when you open the **Properties** dialog box, verify that the Remote Registry service is running on the target computer. If the Remote Registry service is not running, or if the target computer is running Windows NT 4.0, you cannot view or modify hardware profile settings, but you can perform any other actions. For more information about troubleshooting services, see Troubleshooting Services Snap-in. Also, see "Managing System Services" at the Microsoft Windows Resource Kits Web site.

#### To start, stop, or change the startup method for a service

1. Open Services.

2. In the details pane, right-click the service that you want to configure, and then click **Properties**.

3. On the **General** tab, in **Startup type**, click **Automatic**, **Manual**, or **Disabled**.

4. To specify the user account that the service can use to log on, click the **Log On** tab, and then do one of the following:

    - To specify that the service uses the LocalSystem account, click **Local System account**.

    - To specify that the service uses the LocalService account, click **This account**, and then type **NT AUTHORITY\LocalService**.

    - To specify that the service uses the NetworkService account, click **This account**, and then type **NT AUTHORITY\NetworkService**.

    - To specify another account, click **This account**, click **Browse**, and then specify a user account in the **Select User** dialog box. When you are finished, click **OK**.

5. Type the password for the user account in the **Password** box and in the **Confirm password** box, and then click **OK**. If

you select the Local Service account or Network Service account, the password must be blank.

### Caution

- Most services are not designed to have the default account changed. Changing the default account of a service could result in the service failing to start.
- Changing the default service settings may prevent key services from running correctly. It is especially important to use caution when changing the **Startup type** and **Log on as** settings of services that are configured to start automatically.
- In most cases, it is recommended that you not change the **Allow service to interact with desktop** setting. If you allow the service to interact with the desktop, any information that the service displays on the desktop will also be displayed on the interactive user's desktop. A malicious user could then take control of the service or attack it from the interactive desktop.

### Notes

- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.
- The Local Service account and Network Service account are configured with a null password. Whatever password information you supply is ignored. For more information about service accounts, see Services permissions.
- It is recommended that user accounts that are used to log on as a service have the **Password never expires** check box selected in their properties dialog box and that they have strong passwords.
- If account lockout policy is enabled and the account is locked out, the service will not start.
- If you enable or disable a service and you encounter a problem starting the computer, you can start the computer in Safe Mode. In Safe Mode, core services that are required to start the operating system are started in a default scheme, regardless of any changes that are made to the service settings. After the computer is in Safe Mode, you can change the service configuration or restore the default configuration.
- If you specify an account that does not have permission to log on as a service, the Services snap-in automatically grants the appropriate permissions to that account on the computer that you are managing.
- For more information about the user accounts that a service uses to log on, see "Managing System Services" at the Microsoft Windows Resource Kits Web site.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Monitoring network traffic

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring network traffic**
As an administrator, you need to monitor and detect problems with traffic on your network. With Network Monitor, you can gather information about the network traffic that flows to and from the network adapter of the computer on which it is installed. Once you capture the information, you can use Network Monitor to analyze the information, diagnose problem traffic patterns, and devise strategies to prevent future network traffic problems.

Some of the most common tasks are installing Network Monitor, specifying data frame patterns to capture, capturing network frames, and viewing a specific frame. You can also monitor network traffic from the Managing Network Monitor from the command line. For more information about other tasks for monitoring network traffic, see Network Monitor How To ....

**To install Network Monitor**

1. Open Windows Components wizard.
2. In the Windows Components wizard, click **Management and Monitoring Tools**, and then click **Details**.
3. In **Subcomponents of Management and Monitoring Tools**, select the **Network Monitor Tools** check box, and then click **OK**.
4. If you are prompted for additional files, insert the installation CD for your operating system, or type a path to the location of the files on the network.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.
- This procedure automatically installs the Network Monitor driver.

**To specify frame data patterns to capture**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Filter**.
4. In the capture filter decision tree, double-click the **AND (Pattern Matches)** line.
5. Complete the Pattern Match dialog box, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.

- You can define up to four data patterns in a capture filter.

- To capture frames that match any of multiple patterns, create a logical branch by selecting one of the patterns in the decision tree and clicking **OR**.

- To exclude frames that match a pattern, select the pattern in the decision tree and then click **NOT**.

- If you specify more than one pattern in the capture filter decision tree, by default, Network Monitor treats each as a logical OR branch.

**To capture network frames**

1. Open Network Monitor.

2. If prompted, select the local network from which you want to capture data by default.

3. On the **Capture** menu, click **Buffer Settings**, and then set the buffer and frame size as appropriate.

4. On the **Capture** menu, click **Start**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.

- To pause, stop, or display the data capture, on the **Capture** menu, click **Pause**, **Stop**, or **Display Captured Data**. You can also stop and display the capture by clicking **Stop and View**.

- Network Monitor displays session statistics of the first 100 unique network sessions that it detects. To reset statistics and view information about the next 100 detected network sessions, click the **Capture** menu, and then click **Clear Statistics**.

- If you want to capture data from multiple local networks at the same time, install an adapter and start an instance of Network Monitor for each network. In each instance, specify the network from which you want to capture data by, on the **Capture** menu, clicking **Networks** and selecting a network.

- If prompted for additional files, insert the installation CD for your operating system, or type a path to the location of the files on the network.

**To view a specific frame**

1. Open Network Monitor.

2. If prompted, select the local network from which you want to capture data by default.

3. Do one of the following:

- On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
- On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.

4. On the **Display** menu, click **Go To Frame**.
5. In the Go To dialog box, type the number of the frame that you want to view, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. For more information, see Set, view, change, or remove permissions on files and folders.
- The frame number is listed in the **Frame** column.
- Frames are numbered in the order that Network Monitor captures them.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Monitoring security-related events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring security-related events**
It is important to monitor security related events so that you can find out if changes are made to security policies or other objects, who made the changes, and when they were made. You can monitor security-related events by setting up an auditing policy. An auditing policy is made up of all of the auditing settings that you configure for individual security event categories.

Some of the most common tasks for monitoring security-related events are defining or modifying auditing policy settings for an event category on your local computer, defining or modifying auditing policy settings for an event category across your organization, and viewing the security log. For more information about other tasks for auditing security-related events, see Auditing Security Events How To....

**To define or modify auditing policy settings for an event category on your local computer**

1. Open Local Security Policy.

2. In the console tree, click **Audit Policy**.

   **Where?**

   - Security Settings/Local Policies/Audit Policy

3. In the details pane, double-click an event category that you want to change the auditing policy settings for.

4. Do one or both of the following, and then click **OK**.

   - To audit successful attempts, select the **Success** check box.

   - To audit unsuccessful attempts, select the **Failure** check box.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

- To audit object accesses, enable auditing of the object access event category by following the steps above. Then, enable auditing on the specific object. For information about how to enable auditing on an object, see Apply or modify auditing policy settings for a local file or folder or Apply or modify auditing policy settings for an object using Group Policy.

**To define or modify auditing policy settings for an event category across your organization**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the domain or organizational unit that you want to set Group Policy for.

3. Click **Properties**, and then click the **Group Policy** tab.

4. Click **Edit** to open the Group Policy object (GPO) that you want to edit. You can also click **New** to create a new GPO, and

then click **Edit**.

5. In the console tree, click **Audit Policy**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Local Policies/Audit Policy

6. In the details pane, double-click an event category that you want to change the auditing policy settings for.

7. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.

8. Do one or both of the following, and then click **OK**.

   - To audit successful attempts, select the **Success** check box.

   - To audit unsuccessful attempts, select the **Failure** check box.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- If you are on a server or workstation joined to the domain, you can open Active Directory Users and Computers if you click **Start**, click **Run**, type **mmc**, and then click **OK**. On the File menu, click **Add/Remove Snap-in**, click **Add**, double-click **Active Directory Users and Computers**, click **Close**, and then click **OK**.

- To audit object accesses, enable auditing of the object access event category by following the steps above. Then, enable auditing on the specific object. For information about how to enable auditing on an object, see Apply or modify auditing policy settings for a local file or folder or Apply or modify auditing policy settings for an object using Group Policy.

#### To view the security log

1. Open Event Viewer.

2. In the console tree, click **Security**. The details pane lists individual security events.

3. In the details pane, double-click an event to view more details.

### Notes

- You must be logged on as a member of the Administrators group or you must have been granted the **Manage auditing and security log** right in Group Policy to perform this procedure.

- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.

- For more information about security events, see "Security Events" on the Microsoft Windows Resource Kits Web site

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Monitoring server performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring server performance**

In today's business environment, administrators need to ensure their computer systems are efficient and reliable. To optimize the performance of your servers, you need the data supplied by performance monitoring. This topic describes the most common tasks associated with monitoring the performance of a simple server configuration, one that has a few client workstations connected to a single facility that contains one or more servers.

To monitor the performance of a simple server configuration, you need to collect three different types of performance data over a period of time:

- General performance data: Information that can help you identify short-term trends such as memory leaks. After a month or two of data collection, you can average the results and save them in a more compact format. This archived data can assist you in capacity planning as your business grows, and later help you to evaluate the effectiveness of your plan.

- Baseline performance data: Information that can help you discover changes that occur slowly, over time. By comparing the current state of your system with historical data, you can troubleshoot and tune your system. Because this information is collected only periodically, there is no need to compress it for storage.

- Data for service level reports: Information that can help you ensure that your system meets a certain service or performance level, and which you will likely present to decision makers who are not performance analysts. How often you collect and maintain this data depends on your specific business needs.

To collect all three types of data, you can use Performance Logs and Alerts to create a counter log. You can also collect this information from the Monitoring performance from the command line. You can then run the log over time, either manually or with automated scheduling. You can customize a counter log by adding objects and adding counters. For more information about monitoring server performance, including how to further modify counter logs, create reports in various formats, and use the collected data, see Performance monitoring How To....

**To create a counter log**

1. Open Performance.

2. Double-click **Performance Logs and Alerts**, and then click **Counter Logs**.

   Any existing counter logs will be listed in the details pane. A green icon indicates that a log is running; a red icon indicates that a log has stopped.

3. Right-click a blank area of the details pane, and click **New Log Settings**.

4. In **Name**, type the name of the counter log, and then click **OK**.

5. On the **General** tab, click **Add Objects** and select the performance objects you want to add, or click **Add Counters** to select the individual counters you want to log.

   - You must use the Run as text box and Password button to collect counter data from a remotely managed computer. These features enable you to specify the logon account name for the log. You can access them from the **General** tab.

6. If you want to change the default file and schedule information, make the changes on the **Log Files** tab and the **Schedule** tab.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

- To save the settings for a counter log, trace log, or alert, right-click the log or alert in the details pane and click **Save Settings As**. You can then specify an .htm file in which to save the settings. To reuse the saved settings for a new log or alert, right-click the details pane, and then click **New Log Settings From** or **New Alert Settings From**. This is an easy way to generate new settings from a log or alert configuration. You can also open the HTML file in Internet Explorer to display a System Monitor graph.

- Although the extension for a log file is .blg, the file format is changed in the Windows server 2003 family and therefore cannot be read on an earlier version of the operating system. For example, if you copy a binary file from Windows Server 2003, Standard Edition, to Windows 2000, you will be unable to read it on the latter system.

**To add objects to a log**

1. Open Performance.

2. Double-click **Performance Logs and Alerts**, and then double-click **Counter Logs** or **Alerts**.

3. In the details pane, double-click the log or alert you want to modify.

4. On the **General** tab, click **Add Objects**. For each object that you want to add to the log, perform the following steps:

    - To log objects from the computer on which the Performance Logs and Alerts service will run, click **Use local computer counter objects**.

      Or, to log objects from a specific computer regardless of where the service is run, click **Select counter objects from computer:** and specify the Universal Naming Convention (UNC) name, such as *MyLogserver*, of the computer you want to monitor.

    - In **Performance object**, click an object to monitor, and then click **Add**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

- When creating counter log settings for export, make sure to select **Use local computer counter objects**. Otherwise, counter logs will obtain data from the computer named in the text box.

- For a description of an object, click the object in **Performance counters** and then click **Explain**.

- When you specify the name of a remote computer using **Select counter objects from computer**, only the name of the local computer and those computer names entered during the current Performance session are listed. Computer names entered in previous sessions will not be listed.

- The Run as feature enables you to run log collections under different accounts. For example, if you need to log data from a remote computer that requires administrative credentials, you can use the feature to specify an account with the necessary permissions.

**To add counters to a log**

1. Open Performance.

2. Double-click **Performance Logs and Alerts**, and then click **Counter Logs**.

3. In the details pane, double-click the log you want to modify.

4. On the **General** tab, click **Add Counters**. For each counter or group of counters that you want to add to the log, perform the following steps:

   - To log counters from the computer on which the Performance Logs and Alerts service will run, click **Use local computer counters**.

     Or, to log counters from a specific computer regardless of where the service is run, click **Select counters from computer** and specify the Universal Naming Convention (UNC) name, such as *MyLogserver*, of the computer you want to monitor.

   - In **Performance object**, select an object to monitor.

   - In **Select counters from list**, click one or more counters to monitor.

   - To monitor all instances of the selected counters, click **All Instances**. (Binary logs can include instances that are not available at log startup but subsequently become available.)

     Or, to monitor particular instances of the selected counters, click **Select Instances From List**, and then click an instance or instances to monitor.

   - Click **Add**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

- When creating a monitoring console for export, make sure to select **Use local computer counters**. Otherwise, counter logs will obtain data from the computer named in the text box, regardless of where the console file is installed.

- For a description of a counter, click the counter in **Performance counters** and then click **Explain**.

- Some object types have several instances. For example, if a system has multiple processors, the Processor object type will have multiple instances. The PhysicalDisk object type has two instances if a system has two disks. Some object types, such as Memory and server, have only a single instance. If an object type has multiple instances, you can add counters to track statistics for each instance, or in many cases, for all instances at once. By default, counters are shown with both the instance name and an instance index. To turn off this feature, right-click the System Monitor graph, click **Properties**, and click the **General** tab.

- When you specify the name of a remote computer using **Select counters from computer**, only the name of the local computer and those computer names entered during the current Performance session are listed. Computer names entered in previous sessions will not be listed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Resetting user passwords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resetting user passwords**
Users occasionally forget their passwords to their local user accounts. You can manually reset a user's password. However, by doing so, the following information is lost:

- E-mail that is encrypted with the user's public key
- Internet passwords that are saved on the computer
- Files that the user has encrypted

To prevent information loss when you reset a local user account password, you can create and use a password reset disk.

Some of the most common tasks when administering local user accounts are resetting a local user account password, resetting a domain user account password, and creating a password reset disk. You can also reset a user password from the Dsmod. For more information about other tasks for using passwords, see Password How To....

**To reset a local user account password**

1. Open Computer Management.
2. In the console tree, double-click **Local Users and Groups**.

   **Where?**

   - Computer Management (Local)/System Tools/Local Users and Groups

3. In the console tree, click **Users**.
4. In the details pane, right-click the user name, and then click **Set Password**.
5. Read the warning message, and then, if you want to continue, click **Proceed**.
6. In **New password** and in **Confirm password**, type a new password, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**To reset a domain user account password**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

- Active Directory Users and Computers/*Domain name*/Users

3. In the details pane, right-click the user name, and then click **Reset Password**.

4. Type a new password in **New Password** and in **Confirm New Password**, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- If you are on a server or workstation joined to the domain, you can open Active Directory Users and Computers if you click **Start**, click **Run**, type **mmc**, and then click **OK**. On the File menu, click **Add/Remove Snap-in**, click **Add**, double-click **Active Directory Users and Computers**, click **Close**, and then click **OK**.

#### To create a password reset disk

1. Press CTRL+ALT+DEL, and then click **Change Password**.

2. In **User name**, type the user name of the account that you want to create a password reset disk for.

3. In **Log on to**, click your *LocalComputerName*.

4. Click **Backup**.

5. Follow the steps in the Forgotten Password Wizard until the procedure is complete. Store the password reset disk in a safe and secure place.

### Notes

- Password reset disks can only be used for local computer accounts, not for domain accounts.

- You cannot change your password and create a password reset disk at the same time. If you type a new password in **New Password** and in **Confirm New Password** before you click **Backup**, the new password information is not saved. When the wizard prompts you for your current user account password, you must type the old password. You can change your password after you create the password reset disk.

- You can change your password any time after you create a password reset disk. You do not have to create a new password reset disk if you change your password or if the password is reset manually.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Safeguarding your system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Safeguarding your system**

To safeguard your system against a serious failure, you can use the Backup tool to create an Automated System Recovery (ASR) set on a regular basis. The Automated System Recovery Wizard creates a two-part backup that you can use to recover your system after all other recovery attempts have failed, or after you have replaced the hard disk. ASR backs up the system state, system services, and all disks associated with the operating system components. It also creates a startup disk that contains information about the backup, the disk configurations (including basic and dynamic volumes), and how to accomplish a restore. You should create a new ASR set after any major change to the system and also on a regular schedule as part of a comprehensive backup plan.

Some of the most common tasks are creating an Automated System Recovery set using Backup and recovering from a system failure using Automated System Recovery. For more information about other system recovery tools, see Repair overview.

**To create an Automated System Recovery set using Backup**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use the Backup or Restore wizard to create an ASR set by answering **All information on this computer** in the **What do you want to backup?** section. Otherwise, you can go to the next step to create an ASR set in **Advanced Mode**.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. On the **Tools** menu, click **ASR Wizard**.

4. Follow the instructions that appear on your screen.

### Important

- You will need a blank 1.44 MB floppy disk to save your system settings and media to contain the backup files. If your computer does not have a floppy drive, see Backup Troubleshooting.

### Notes

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- This procedure will only back up those system files that are necessary for starting your system. You will need to backup your data separately.

- After creating the ASR set, label this floppy disk and backup media carefully, and keep them together. In order to use the backup media, you must have the floppy disk that you created with that set of media. You cannot use a floppy disk created at a different time or with a different set of media. You must also have your Setup CD available at the time you perform Automated System Recovery.

- Keep the ASR set in a secure location. The ASR set contains information about your systems configuration that could be used to do harm to your system.

- If you are backing up a server cluster, run the Automated System Recovery Preparation Wizard on all nodes of the cluster and make sure that the Cluster service is running when you start each Automated System Recovery backup. Make sure

that one of the nodes on which you run the Automated System Recovery Preparation Wizard is listed as the owner of the quorum resource while the wizard is running.

**To recover from a system failure using Automated System Recovery**

1. Make sure you have the following on hand before you begin the recovery procedure:

    - Your previously created Automated System Recovery (ASR) floppy disk.
    - Your previously created backup media.
    - The original operating system installation CD.
    - If you have a mass storage controller and you are aware that the manufacturer has supplied a separate driver file for it (different from driver files available on the Setup CD), obtain the file (on a floppy disk) before you begin this procedure.

2. Insert the original operating system installation CD into your CD drive.
3. Restart your computer. If prompted to press a key in order to start the computer from CD, press the appropriate key.
4. If you have a separate driver file as described in step 1, use the driver as part of Setup by pressing F6 when prompted.
5. Press F2 when prompted at the beginning of the text-only mode section of Setup. You will be prompted to insert the ASR floppy disk you have previously created.
6. Follow the directions on the screen.
7. If you have a separate driver file as described in step 1, press F6 (a second time) when prompted after the system reboots.
8. Follow the directions on the screen.

## Notes

- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.
- ASR will not restore your data files.
- If you are restoring a server cluster in which all nodes failed and the quorum disk cannot be restored from backup, use Automated System Recovery on each node in the original cluster to restore the disk signatures and the partition layout of the cluster disks (quorum and nonquorum).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Scheduling tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scheduling tasks**
By using Task Scheduler, you can schedule commands, programs, or scripts to run at specific times. For example, if you create a script to install updates on employee workstations, you can schedule the script to run in the evening, when most employees are not at work.

Some of the most common tasks are scheduling a new task and scheduling a task on a remote computer. You can also schedule tasks from the Creating and managing scheduled tasks from the Command Line. For more information about other tasks you can perform by using Task Scheduler, see Task Scheduler How To....

**To schedule a new task**

1. Open Task Scheduler.
2. Double-click **Add Scheduled Task**.
3. Follow the instructions in the Scheduled Task Wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators, Backup Operators, or Server Operators, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- When creating a scheduled task, you must enter a user name and password, either in the Add Scheduled Task Wizard or in the **Run as** box on the **Task** tab of the properties dialog box for the scheduled task. When the scheduled task runs, the program runs as if it were started by the user you specified, with that user's security context.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- Confirm that the system date and time on your computer are accurate, as Task Scheduler relies on this information to run scheduled tasks. To verify or change this information, double-click the time indicator on the taskbar.
- If you want to configure advanced settings for the task, select the **Open advanced properties for this task when I click Finish** check box in the final page of the wizard. This opens the properties dialog box for the task. You can then change the program being run on the **Task** tab, fine-tune the schedule on the **Schedule** tab, customize settings on the **Settings** tab, or set user and group permissions on the **Security** tab.

**To schedule a task on a remote computer**

1. Open My Network Places.
2. Click **Search for Computers**.
3. In the **Computer name** box, enter the computer name, and then click **Search**.
4. In the details pane, double-click the remote computer on which you want to view scheduled tasks.
5. Double-click **Scheduled Tasks**.
6. From the **File** menu, point to **New**, and then click **Scheduled Task**.

7. Type a name for the new task, and then press ENTER.

8. Right-click the new task, and then click **Properties**.

9. Do the following:

    - To specify the program being run, in **Run**, type the path for the new program.

    - To set the user account name and password that the scheduled task will run as, edit or enter a new name in the **Run as** box, and then click **Set password**.

    - To set the schedule for the task, click the **Schedule** tab.

    - To customize the settings for the task, click the **Settings** tab.

    - To set user and group permissions for the task, click the **Security** tab.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or have been delegated the appropriate authority, on the remote computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open My Network Places, click **Start**, and then click **My Computer**. Under **Other Places**, click **My Network Places**.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Setting up DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up DNS**

Domain Name System (DNS) is used in Transmission Control Protocol/Internet Protocol (TCP/IP) networks, such as the Internet, to locate computers and services through user-friendly names. When a user enters a DNS name in an application, DNS can resolve the name to other information associated with the name, such as an IP address. If your network is homogeneous (all hosts use the same operating system), and it uses a name resolution method other than DNS, you can continue to use that method without needing DNS. If your network is heterogeneous (the hosts use different operating systems), or it connects to the Internet, you will want to deploy DNS as the Internet consists of protocols that require DNS.

Setting up DNS in your network does not necessarily require that you administer a large DNS infrastructure or become an expert DNS administrator. If you have a very small network in which information is maintained dependably, you may choose to have your DNS namespace administered by a different organization that specializes in DNS administration, such as your government or an Internet Service Provider (ISP). In this case, the other organization will host and administer your DNS zone data for you or integrate your hosts with an existing DNS zone hosted in their network. Once your network expands beyond a small number of hosts, you will want to begin administering DNS yourself in order to better serve the name resolution needs of your expanding network.

Setting up a DNS server typically involves configuring the DNS server with DNS zones to administer the DNS domain names in your network, adding DNS resource records for the hosts in your network to your DNS zones, and delegating administration for these zones by creating a delegation from the parent DNS server previously authoritative for a domain name to the child DNS server that is accepting responsibility for a domain name. Lastly, a DNS server administrator should maintain the integrity of DNS zone data by securing the DNS servers in the network. You can also set up a DNS server from the command line. For more information, see DNS tools. For more information about setting up DNS, see DNS How To....

**To configure a DNS server**

1. Open DNS.
2. If needed, add and connect to the applicable server in the console.
3. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*Applicable DNS server*

4. On the **Action** menu, click **Configure a DNS server**.
5. Follow the instructions in the Configure a DNS Server Wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- If the DNS server is running locally, you do not need to perform step 2.
- The Configure a DNS Server Wizard allows you to specify zones, root hints, and forwarders.
- If you are setting up DNS for your intranet and you want your computers to connect to the Internet also, configure the

root hints with the IP addresses of the DNS servers located at your ISP.

- When you finish configuring the server, you will need to complete additional tasks, such as configuring your network hosts to use this DNS server as their preferred DNS server, enabling dynamic update for the zones on this DNS server, or adding resource records to its zones.

**To add a resource record to a zone**

1. Open DNS.
2. In the console tree, right-click the applicable zone and click **Other New Records**.
3. In **Select a resource record type** list box, select the type of resource record you want to add.
4. Click **Create Record**.
5. In **New Resource Record**, enter the information needed to complete the resource record.
6. After you specify all of the necessary information for the resource record, click **OK** to add the new record to the zone.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Add a resource record for every host in your network that will use DNS. If you have configured your DNS zones to accept dynamic updates, and your network hosts support dynamic update and are configured to use this DNS server as their preferred DNS server, then your network hosts will add their resource records automatically.
- For more information about the different types of resource records you can add, see Resource records reference.

**To create a zone delegation**

1. Open DNS.
2. In the console tree, right-click the applicable subdomain, and then click **New Delegation**.
3. Follow the instructions provided in the New Delegation Wizard to finish creating the new delegated domain.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- All domains (or subdomains) that appear as part of the applicable zone delegation must be created in the current zone prior to performing delegation as described here. As necessary, use the DNS console to first add domains to the zone before completing this procedure. For more information, see Delegating zones.

**To secure a DNS server**

1. Open DNS.

2. In the console tree, right-click the applicable DNS server, and then click **Properties**.

3. On the **Interfaces** tab, ensure that the DNS server is listening to the correct interfaces.

4. On the **Advanced** tab, select the **Secure cache against pollution** check box.

5. If this DNS server will not perform recursive resolution, then select the **Disable recursion** check box.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

- If a DNS zone is stored in Active Directory, you can also secure that DNS zone and its resource records by using the security features of Active Directory.

- By default, a DNS zone is only permitted to allow zone transfers from the DNS servers that are listed in the DNS zone properties.

- By default, the DNS service listens for DNS message communications on all IP addresses that are configured for the server computer.

- Server IP addresses that are added here need to be statically managed. If you later change or remove addresses specified here from the TCP/IP configurations maintained at this server, update this list accordingly.

- The **Secure cache against pollution** check box is selected by default.

- If you disable recursion on this DNS server, you cannot use this DNS server to answer recursive queries from DNS clients or to list this DNS server as a DNS forwarder. Disable recursion on DNS servers that only perform iteration with other DNS servers.

- By disabling recursive resolution on this DNS server, you can prevent a Denial of Service (DoS) attack in which a malicious user attempts to have this DNS server respond to recursive queries for a domain in a DNS zone under the malicious user's control.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Setting up TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up TCP/IP**

By default, computers running Windows Server 2003 operating systems are configured as Dynamic Host Configuration Protocol (DHCP) clients, which means they can accept leases from any available DHCP server. TCP/IP computers (hosts) on a DHCP network automatically obtain TCP/IP configurations from DHCP servers that are located elsewhere on that network. This is the most common scenario for mid-sized and large TCP/IP networks. For more information about DHCP, see DHCP.

In certain cases, you will need to configure TCP/IP for static addressing. Configuring TCP/IP for static addressing is the process of manually assigning unique, permanent addresses to network devices, and providing static address information about other network resources (for example, DHCP servers, routers, and default gateways). If you are setting up a computer to run on a network that does not have a DHCP server, but that does provide Windows Internet Name Service (WINS), you will need to configure TCP/IP to use WINS. For more information about WINS, see WINS.

Some of the most common tasks for setting up TCP/IP are configuring TCP/IP for static addressing, and configuring TCP/IP to use WINS. You can also set up TCP/IP from the Managing TCP/IP from the command line. For more information about other tasks for setting up TCP/IP, see TCP/IP How To....

**To configure TCP/IP for static addressing**

1. Open Network Connections.
2. Right-click the network connection you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Use the following IP address**, and do one of the following:
   - For a local area connection, in **IP address**, **Subnet mask**, and **Default gateway**, type the IP address, subnet mask, and default gateway addresses.
   - For all other connections, in **IP address**, type the IP address.
5. Click **Use the following DNS server addresses**.
6. In **Preferred DNS server** and **Alternate DNS server**, type the primary and secondary DNS server addresses.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If your network does not have a DNS server, you can skip steps 5 and 6.

**To configure TCP/IP to use WINS**

1. Open Network Connections.

2. Right-click the network connection you want to configure, and then click **Properties**.

3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.

4. Click **Advanced**, click the **WINS** tab, and then click **Add**.

5. In **TCP/IP WINS server**, type the IP address of the WINS server, and then click **Add**.

6. Repeat steps 4 and 5 for each WINS server IP address you want to add, and then click **OK**.

7. To enable the use of the Lmhosts file to resolve remote NetBIOS names, select the **Enable LMHOSTS lookup** check box. This option is enabled by default.

8. To specify the location of the file that you want to import into the Lmhosts file, click **Import LMHOSTS**, and then select the file in the **Open** dialog box.

9. To enable or disable NetBIOS over TCP/IP, do the following:

   - To enable the use of NetBIOS over TCP/IP, click **Enable NetBIOS over TCP/IP**.

   - To disable the use of NetBIOS over TCP/IP, click **Disable NetBIOS over TCP/IP**.

   - To have the DHCP server determine whether NetBIOS over TCP/IP is enabled or disabled, click **Default**. For more information, see Notes.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- If you are using a DHCP server to allocate WINS server IP addresses, you do not need to add WINS server addresses.

- You can no longer configure the NetBIOS scope ID on the **WINS** tab (previously available in Windows NT 4.0). To configure the NetBIOS scope ID, set the following registry value to the name of the scope ID that you want to use:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\ScopeID**

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Setting user and group security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting user and group security**

In order to secure a computer and its resources, you must decide what tasks and actions users or groups of users can perform. The tasks and actions that a user or group of users can perform are determined by the user rights that you assign to them. For example, if a trusted member of the Users group needed to monitor the security log, you could grant the user the "Manage auditing and security log" user right instead of adding the user to a more privileged group, such as the Administrators group. Similarly, you can secure an object, such as a file or folder, by assigning permissions.

Some of the most common tasks are assigning user rights on your local computer, assigning user rights throughout your organization, and setting file and folder permissions. For more information about other tasks for setting user and group security, see Access Control How To....

**To assign user rights on your local computer**

1. Open Local Security Settings.
2. In the console tree, click **User Rights Assignment**.

   **Where?**

   - Security Settings/Local Policies/User Rights Assignment
3. In the details pane, double-click the user right that you want to modify.
4. Do one of the following, and then click **OK**.

   - To add a user or group, click **Add a User or Group**. In **Select Users or Groups**, type the user name or group name that you want to add, and then click **OK**.
   - To remove a user or group, select the user or group in the list, and then click **Remove**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**To assign user rights throughout your organization**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the Group Policy object for which you want to edit security settings.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the Group Policy object you want to edit. Or, click **New** to create a new Group Policy object, and then click **Edit**.

5. In the console tree, click **User Rights Assignment**.

   **Where?**

   - *GroupPolicyObject* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/Local Policies/User Rights Assignment

6. In the details pane, double-click the user right that you want to modify.

7. If this security setting has not yet been defined, select the **Define these policy settings** check box.

8. Do one of the following, and then click **OK**.

   - To add a user or group, click **Add a User or Group**. In **Select Users or Groups**, type the user name or group name that you want to add, and then click **OK**.

   - To remove a user or group, select the user or group in the list, and then click **Remove**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- If you are on a server or workstation joined to the domain, you can open Active Directory Users and Computers if you click **Start**, click **Run**, type **mmc**, and then click **OK**. On the File menu, click **Add/Remove Snap-in**, click **Add**, double-click **Active Directory Users and Computers**, click **Close**, and then click **OK**.

- Always test a newly created policy on a test organizational unit before applying it to your network.

- When you change a security setting and click **OK**, that setting will take effect in the next refresh of settings.

- The security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes.

**To set, view, change, or remove permissions on files and folders**

1. Open Windows Explorer.

2. Right-click the file or folder for which you want to set permissions, click **Properties**, and then click the **Security** tab.

3. Do one of the following:

   - To set permissions for a group or user that does not appear in the **Group or user names** box, click **Add**. Type the name of the group or user for which you want to set permissions, and then click **OK**.

   - To change or remove permissions from an existing group or user, click the name of the group or user.

4. Do one of the following:

- To allow or deny a permission, in the **Permissions for *User or Group*** box, select the **Allow** or **Deny** check box.
- To remove the group or user from the **Group or user names** box, click **Remove**.

### Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- In the Windows Server 2003 family, the Everyone group no longer includes Anonymous Logon.
- You can set file and folder permissions only on drives formatted to use NTFS.
- If the check boxes under **Permissions for *User or Group*** are shaded or if the **Remove** button is unavailable, then the file or folder has inherited permissions from the parent folder. For more information, see How inheritance affects file and folder permissions.
- When adding a new user or group, by default, this user or group will have **Read & Execute**, **List Folder Contents**, and **Read** permissions.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Troubleshooting Strategies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting strategies**

This section provides links to information about troubleshooting:

- Troubleshooting A-Z
- Troubleshooting Stop messages: general strategies
- Troubleshooting compatibility issues
- Troubleshooting audio, imaging, and CD-burning peripheral devices
- Server Cluster Troubleshooting
- Using Event Viewer
- Using the Shutdown Event Tracker
- Using online resources
- Troubleshooters

# Troubleshooting A-Z

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting A-Z**
This section contains pointers to troubleshooting information available in Help.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**
Active Directory: Troubleshooting Active Directory

Audio: Troubleshooting audio, imaging, and CD-burning peripheral devices

**B**
Backup Troubleshooting

**C**
Cameras: Troubleshooting audio, imaging, and CD-burning peripheral devices

CD-burning: Troubleshooting audio, imaging, and CD-burning peripheral devices

Certificate Services Troubleshooting

Certificate Templates Troubleshooting

Clustering: Server Cluster Troubleshooting

Compatibility: Troubleshooting compatibility issues

Connection Manager Administration Kit (CMAK): Troubleshooting process

Connection Point Services Troubleshooting

**D**
DHCP Troubleshooting

Disk Defragmenter: Troubleshooting Disk Defragmenter

Disk Management: Troubleshooting Disk Management

Disk Quotas: Troubleshooting Disk Quotas

DNS: Troubleshooting DNS

**E**
Emergency Management Services: Troubleshooting Emergency Management Services

Encrypting File System Troubleshooting

Event Viewer Troubleshooting

**I**
Imaging: Troubleshooting audio, imaging, and CD-burning peripheral devices

IAS Troubleshooting

IPSec Troubleshooting

**L**
Licensing Troubleshooting

Local Users and Groups: Troubleshooting Local Users and Groups

**M**
MMC: Troubleshooting MMC

Modems Troubleshooting

**N**
Network Connections Troubleshooting

Using Network Diagnostics

Network Load Balancing Clusters

Network Monitor: Troubleshooting Network Monitor

Novell NetWare Integration Troubleshooting

**P**
Performance monitoring: Troubleshooting Performance monitoring

Printing Troubleshooting

**R**
Remote Access Troubleshooting

Remote Installation Services: Troubleshooting Remote Installation Services

Remote Desktop for Administration: Troubleshooting Remote Desktop for Administration

Removable storage: Troubleshooting Removable Storage

Routing: Troubleshooting routing

**S**
Scanners: Troubleshooting audio, imaging, and CD-burning peripheral devices

Security Troubleshooting

Server Cluster Troubleshooting

Services for Macintosh Troubleshooting

Shared folders: Troubleshooting Shared Folders

Shutdown Event Tracker: Troubleshooting Shutdown Event Tracker

SNMP: Troubleshooting SNMP

Software Restriction Policies Troubleshooting

Stop messages: Troubleshooting Stop messages: general strategies

Stored User Name and Password Troubleshooting

System Monitor: Troubleshooting Performance monitoring

**T**
Task Scheduler: Troubleshooting Task Scheduler

TCP/IP: Troubleshooting TCP/IP

Telephony Troubleshooting

Windows Time Service Troubleshooting

**V**
VPN Troubleshooting

**W**
Windows ATM services: ATM Troubleshooting

Windows Installer Troubleshooting

WINS Troubleshooting

Windows Server 2003 Wireless Troubleshooting

# Troubleshooting Stop messages: general strategies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting Stop messages: general strategies**
This topic describes a general method for troubleshooting Stop messages or unexpected shutdowns. For more information, see the following:

- For information about specific Stop messages, see Troubleshooting specific Stop messages.
- For articles about specific Stop messages, see Search the Knowledge Base at the Microsoft Web site.
- For information about support options, see Support resources.
- For information about troubleshooting specific technologies or components (for example, disks and file systems), see the *Windows Server 2003 Technical Reference* at the Microsoft Windows Server System Web site.

You can use the System Information tool to gather important information about system settings, problem devices, and hardware resources. For more information, see System Information.

**Important**

- To ensure system stability, install the latest service pack that is appropriate for your system (follow service pack instructions) and use only signed drivers. For more information, see Windows Update at the Microsoft Web site. For a description of the different ways you can use Windows Update, see Windows Update.

Follow these steps for any Stop error that does not display identifying text or specific troubleshooting steps:

1. Use the Windows Online Crash Analysis tool at the Microsoft Web site. You can use this tool to send error reports to Microsoft and track their status by using your Microsoft Passport information. You can access the Online Crash Analysis Web site by using the Error Reporting service or by using your Web browser. When it is enabled, the Error Reporting service monitors your system for kernel and user mode faults that are related to operating system components and applications. With kernel-mode reporting, you can obtain more information about the problem or condition that caused the Stop error. For more information, see System and program error reporting overview.

2. Use the File Signature Verification tool, which identifies unsigned drivers and incompatible system files on your computer. The system files and device driver files that are provided with the Windows Server 2003 family have a Microsoft digital signature, which indicates that the files are original, unaltered system files or that they have been approved by Microsoft for use with Windows. For more information, see Using File Signature Verification.

3. If the Stop message is caused by a driver problem, you can use Device Manager to roll back to a previous version of the driver. For more information, see Roll back to the previous version of a driver. You may need to disable any newly installed drivers.

4. Using a recent version of virus-protection software, check for viruses on your computer. If a virus is found, perform the steps required to eliminate it from your computer. See your antivirus software documentation for these steps.

5. You can search the Knowledge Base for updated information on Stop messages. See Search the Knowledge Base at the Microsoft Web site, and then follow the instructions on the page. When typing keywords, use the Stop message number (for example, **stop 0x0000000A**).

   For more information about articles, troubleshooting wizards, and items that you can download, see Updated technical information.

6. If you can start the operating system, check the System Log in Event Viewer for additional information that might help

determine the device or driver that is causing the problem. Many Stop messages produce specific events that you can search for, such as Event ID 1001 (generated by various Stop messages) and Event ID 6008 (usually caused by an unexpected shutdown or a computer that does not restart correctly). To display a description of an event, in Event Viewer, double-click the event. For more information about Event Viewer, see Event Viewer.

You can search the Knowledge Base for updated information on Event IDs. See Search the Knowledge Base at the Microsoft Web site, and then follow the instructions on the page.

7. If you cannot start your computer, try starting it in Last Known Good Configuration or in Safe Mode, and then remove or disable newly added programs or drivers. For information about how to start your computer in Safe Mode, see Start the computer in Safe Mode. For more information about how to start your computer in Last Known Good Configuration, see Start the computer using the last known good configuration.

   **Important**

   - When you use Last Known Good Configuration, system setting changes made after the last successful startup are lost.

8. Remove any newly installed hardware (RAM, adapters, hard disks, modems, and so on).

9. Ensure that you have updated drivers for your hardware devices. Also, ensure that you have the latest system BIOS (for x86-based or x64-based computers) or the correct firmware (for Itanium architecture-based computers). The device or hardware manufacturers can assist you in obtaining these items.

10. Run the system diagnostics supplied by your computer manufacturer, especially the memory check.

11. Confirm that your hardware is designed for the Windows Server 2003 family by clicking the appropriate link in Support resources.

12. If Safe Mode and other startup options do not work, you can consider using the Recovery Console. This method is recommended only if you are an advanced user who can use basic commands to identify and locate problem drivers and files. In addition, you must be an administrator to use the Recovery Console. For more information, see Recovery Console overview.

# Troubleshooting compatibility issues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting compatibility issues**
What problem are you having?

- An application stops responding, quits, or behaves improperly.
- Hardware does not function properly.

**An application stops responding, quits, or behaves improperly.**
**Cause:** The version of the application is not compatible with this version of Windows.

**Solution:** Start the Program Compatibility Wizard to detect and test compatibility settings. To start the Program Compatibility Wizard, click **Start**, click **Run**, and then type **hcp://system/compatctr/compatmode.htm**.

**See also:** Accessing information about compatibility; Troubleshooting specific Stop messages

**Cause:** The display settings are not correct.

**Solution:** Some older games and applications may require a 256-color quality or 640 x 480 screen resolution to function properly. Use the **Display Properties** dialog box to change the display settings. To open Display, click **Start**, click **Control Panel**, and then double-click **Display**.

**Note**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**Hardware does not function properly.**
**Cause:** The hardware is not compatible with this release, or it needs to be updated to work with this release.

**Solution:** To confirm that your hardware is designed for the Windows Server 2003 family, see the compatibility information in Support resources.

**See also:** Hardware compatibility

**Notes**

- For the latest information about products that are designed for the Windows Server 2003 family, see Hardware compatibility information in the Windows Catalog at the Microsoft Web site.
- For more information about the Designed for Windows logo program, see Designed for Windows at the Microsoft Web site.

# Troubleshooting audio, imaging, and CD-burning peripheral devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting audio, imaging, and CD-burning peripheral devices**
What problem are you having?

- The audio device does not work.
- The imaging device does not work.
- The CD-burning device does not work.

**The audio device does not work.**
**Cause:** If you receive the following message: "The service cannot be started, either because it is disabled or because it has no enabled devices associated with it," the device requires that the Windows Audio service be enabled. By default, the Windows Audio service required by this peripheral device is started only on computers running Windows Server 2003, Standard Edition.

**Solution:** If sound is required, start this service, and then set **Startup type** to **Automatic**.

Most peripheral devices are supported by drivers on the installation CD for the Windows Server 2003 family. If the CD does not list a driver for your device, look for an update from the device manufacturer or from Windows Update. To check if your hardware is compatible, see Support resources.

**See also:** Start, stop, pause, resume, or restart a service; Configure how a service is started

**Note**

- The hardware drivers for this peripheral device cannot be updated unless the Windows Audio service is started.

**The imaging device does not work.**
**Cause:** If you receive the following message: "The service cannot be started, either because it is disabled or because it has no enabled devices associated with it," the device requires that the Windows Image Acquisition (WIA) service be enabled. By default, the Windows Image Acquisition (WIA) service required by this peripheral device is disabled.

**Solution:** If imaging is required, start this service, and then set **Startup type** to **Manual**.

Most peripheral devices are supported by drivers on the installation CD for the Windows Server 2003 family. If the CD does not list a driver for your device, you can look for an update from the device manufacturer or from Windows Update. To check if your hardware is compatible, see Support resources.

**See also:** Start, stop, pause, resume, or restart a service; Configure how a service is started

**Note**

- The hardware drivers for this peripheral device cannot be updated unless the Windows Image Acquisition (WIA) service is started.

**The CD-burning device does not work.**
**Cause:** The IMAPI CD-Burning COM Service required by this peripheral device is disabled.

**Solution:** If CD burning is required, start this service and set **Startup type** to **Manual**, and then restart the server.

Most peripheral devices are supported by drivers on the installation CD for the Windows Server 2003 family. If the CD does not list a driver for your device, you can look for an update from the device manufacturer or from Windows Update. To check if your hardware is compatible, see Support resources.

**See also:** Start, stop, pause, resume, or restart a service; Configure how a service is started

**Note**

- With some third-party software, you may be able to burn CDs without starting the IMAPI CD-Burning COM Service.

# Using Event Viewer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Event Viewer**

You can use Event Viewer to monitor events that are recorded in the Application, Security, and System logs. By using the event logs in Event Viewer, you can gather information about hardware, software, and system problems. You can also monitor security events. For more information, see Event Viewer.

# Using the Shutdown Event Tracker

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Shutdown Event Tracker**

Shutdown Event Tracker provides a simple and standard mechanism that you can use to consistently document the reasons for shutting down or restarting your computer. You can then use this information to analyze the root causes of shutdowns and to develop a more complete understanding of your system environment. For more information, see Shutdown Event Tracker.

# Using online resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using online resources**

The Microsoft Web site contains support information, including the latest downloads and Knowledge Base articles written by support professionals at Microsoft:

- Product Support Services. Search frequently asked questions (FAQs) by product, browse the product support newsgroups, and contact Microsoft Support.
- Search the Knowledge Base. Search the Microsoft Knowledge Base of technical support information and self-help tools for Microsoft products.
- Windows Server Community. The official online community for enthusiasts of the Windows server operating systems.
- Windows Deployment and Resource Kits. Find information about the *Windows Deployment and Resource Kits* that are available for Windows products.
- Windows Resource Kits--Web Resources. Find information about and resources for a variety of Windows topics.
- Microsoft TechNet. Search for troubleshooting information, service packs, patches, and downloads for your system.

# Troubleshooters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooters**
Troubleshooters are tools that walk you through technical problems that might occur with your computer. For more information, see the following:

- Using troubleshooters

# Quick Guide for Finding Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quick guide for finding tools**

This section provides links to topics and tools for carrying out basic administrative tasks:

- Checklists Provided in Windows Server 2003 Help
- Best practices: pointers to recommendations and tips
- New ways to do familiar Windows tasks
- Configuring Roles for Your Server
- Troubleshooting A-Z
- Step-by-step guides for deployment
- Windows interface administrative tool reference A-Z
- Command-line reference A-Z
- Windows Resource Kit Tools Help
- Install Windows Support Tools

**Notes**

- For information about Microsoft® Visual Basic®, Scripting Edition, (VBScript) for system administrators, see "Scripting Solutions for System Administrators" at the Microsoft Windows Resource Kits Web site.
- For information about scripting, see "Scripting for the Enterprise" at the Microsoft Windows Resource Kits Web site.

# Checklists Provided in Windows Server 2003 Help

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

A checklist describes the steps required to set up, configure, or implement a Windows Server 2003, Enterprise Edition, component. Each step refers you to background information or instructions that are needed to complete the step. To perform a task, find the component in the following table and click the checklist.

Windows Server 2003, Enterprise Edition, offers all the features of Windows Server 2003, Standard Edition, as well as enhanced memory support, support for additional processors, and clustering. To learn more about memory support, see Enterprise Memory Architecture. To find checklists for Network Load Balancing clusters and server clusters, see Clustering in the following table.

| Component | Checklist |
|---|---|
| Active Directory | Active Directory Checklists |
| Backing up and restoring data | Checklists for Backing Up and Restoring Data |
| Certificate Services | Checklist: Creating a certification hierarchy with an offline root certification authority<br>Checklist: Configuring certificate autoenrollment |
| Clustering | Checklist: Enabling and configuring Network Load Balancing<br>Checklists: Creating Server Clusters<br>Checklists: Installing Server Cluster Resources<br>Checklist: Preparation for upgrading a cluster<br>Checklist: Preparation for installing a cluster |
| Configuring your server | Checklist: Configuring the typical setup for a first server |
| Connection Point Services (CPS) | Connection Point Services Checklists |
| DHCP | DHCP Checklists |
| Disk Defragmenter | Checklist: Defragmenting volumes |
| Disk Management | Checklist: Adding a new disk |
| Disk quotas | Checklist: Setting up disk quotas |
| Distributed File System (DFS) | Checklist: Creating a distributed file system |
| DNS | DNS Checklists |
| enterprise Memory Architecture | Checklist: Enabling application memory tuning<br>Checklist: Enabling Physical Address Extension (PAE) |
| Installing and upgrading | Checklist: Performing an upgrade<br>Checklist: Performing a new installation<br>Checklist: Preparation for upgrading a cluster<br>Checklist: Preparation for installing a cluster |

| | |
|---|---|
| Internet Authentication Service (IAS) | IAS checklists |
| IP version 6 | Checklist: Installing and configuring IPv6 |
| Licensing | Licensing Checklists |
| Managing security templates | Checklist: Securing computers using Security Configuration Manager |
| Message Queuing | Checklist: Installing a Message Queuing resource |
| Modems | Checklist: Installing and configuring modems |
| Network Connections | Checklist: Configuring network connections<br>Checklist: Setting up a home or small office network |
| Network Load Balancing clusters | Checklist: Enabling and configuring Network Load Balancing |
| Network Monitor | Checklist: Monitoring network traffic on your local computer |
| Printing | Checklist: Setting up a printer |
| Remote Installation Services (RIS) | Checklist: Installing Remote Installation Services |
| Remote Storage | Checklist: Setting up Remote Storage |
| Removable Storage | Checklist: Configure Removable Storage |
| Routing and Remote Access | Checklist: Installing and configuring the remote access server<br>Checklist: Installing and configuring the router |
| Security | Checklist: Securing computers using Security Configuration Manager<br>Checklist: Setting up object access auditing |

| | |
|---|---|
| server clusters | Checklist: Preparation for installing a cluster<br>Checklist: Preparation for upgrading a cluster<br>Checklist: Planning and creating a server cluster<br>Checklist: Validating your clustering system<br>Checklist: Creating a new group<br>Checklist: Creating a new resource<br>Checklist: Creating a server cluster application server<br>Checklist: Creating a server cluster file server<br>Checklist: Creating a server cluster print server<br>Checklist: Creating a server cluster virtual server<br>Checklist: Creating a clustered IIS Web or FTP service<br>Checklist: Installing a DHCP Service resource<br>Checklist: Installing a Distributed File System (DFS) root resource<br>Checklist: Installing a Distributed Transaction Coordinator resource<br>Checklist: Installing a File Share resource<br>Checklist: Installing a Generic Application resource<br>Checklist: Installing a Generic Script resource<br>Checklist: Installing a Generic Service resource<br>Checklist: Installing an IP Address resource<br>Checklist: Installing a Local Quorum resource<br>Checklist: Installing a Majority Node Set resource<br>Checklist: Installing a Message Queuing resource<br>Checklist: Installing a Network Name resource<br>Checklist: Installing a Physical Disk resource<br>Checklist: Installing a Print Spooler resource<br>Checklist: Installing a Volume Shadow Copy Service Task resource<br>Checklist: Installing a WINS Service resource |
| Services for Macintosh | Services for Macintosh Checklists |
| Smart card setup | Checklist: Deploying smart cards for logging on to Windows |
| SNMP service | Checklist: Implementing the SNMP service |
| Software restriction policies | Checklist: Using software restriction policies to help protect your computer against an e-mail virus |
| Stored User Names and Passwords | Checklist: Using Stored User Names and Passwords safely |

| | |
|---|---|
| System recovery | Checklist: Safeguarding your system<br>Checklist: Recovering a system that will not start<br>Checklist: Restoring data |
| TCP/IP | Checklist: Configuring TCP/IP |
| Telephony | Telephony Checklists |
| Terminal server | Checklists: Setting up Terminal Server |
| Virtual private networks (VPNs) | VPN Checklists |
| Windows ATM services | Checklist: Installing ATM services |
| WINS service | Checklist: Installing a WINS server |

# Best practices: pointers to recommendations and tips

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

This section contains pointers to product support recommendations and product tips.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**

Active Directory Best practices

Best practices for Active Directory Domains and Trusts

Best practices for Active Directory Schema

Best practices for Active Directory Sites and Services

Best practices for Active Directory Users and Computers

Best practices for assigning permissions on Active Directory objects

Auditing Security Events Best practices

**B**

Best practices for Backup

**C**

Certificate Services Best practices

Certificate Templates Best practices

Connection Point Services Best Practices

**D**

Best practices for disaster recovery

Best practices for Disk Defragmenter

Best practices for Disk Management

Best practices for Disk Quotas

Best practices for Distributed File System (DFS)

DHCP Best Practices

DNS best practices

**E**

Encrypting File System best practices

Best practices for implementing security measures

Best practices for selecting and configuring hardware

Best practices for using in-band tools

**I**

Best practices for installing and upgrading the operating system

IAS Best Practices

IPSec Best practices

**L**

Best practices for Licensing

Local Users and Groups Best practices

**M**
Modem best practices

Performance monitoring best practices

**N**
Network Connections Best Practices

Network Load Balancing Best practices

Network Monitor Best practices

Best practices for Client Service for NetWare

Best practices for NWLink

**P**
Password Best practices

Performance monitoring best practices

Best practices for permissions and user rights

Best practices for assigning permissions on Active Directory objects

Printing Best practices

**R**
Registry best practices

Best Remote Access practices

Remote Administration Best practices

Remote Desktop for Administration Best practices

Best practices for Remote Installation Services

Best practices for Remote Storage

Best practices for Removable Storage

Routing best practices

**S**
Best practices for security

Best practices for Security Configuration and Analysis

Auditing Security Events Best practices

Best practices for Security Settings

Best practices for Security Templates

Server Cluster Best Practices

Services best practices

Services for Macintosh Best practices

Best practices for Shared Folders

Shutdown Event Tracker Best practices

SNMP Best practices

Software Restriction Policies Best Practices

Stored User Name and Password Best practices

Performance monitoring best practices

**T**

Best practices for TCP/IP

Telephony Best practices

Telnet Server best practices

Terminal Server Best practices

Terminal Services User Properties Best practices

Windows Time Service Best practices

**U**

Best practices for installing and upgrading the operating system

User Profiles best practices

Best practices for permissions and user rights

**V**

VPN best practices

**W**

ATM Best practices

Registry best practices

WINS Best Practices

# New ways to do familiar Windows tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

Many components have changed names or locations since Windows NT 4.0 and Windows 2000. To find the new way to do a task, click the letter that the name of the component starts with, and then click the name of the component.


A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**

New ways to do familiar Active Directory tasks

Active Directory Domains & Trusts: New ways to do familiar tasks

New ways to do familiar Licensing tasks

Active Directory Users & Computers: New ways to do familiar tasks

New ways to do familiar Services for Macintosh tasks

**B**

Backup: New ways to do familiar tasks

Tape Devices: New ways to do familiar tasks

Backup: New ways to do familiar tasks

**C**

CD Player: New ways to do familiar tasks

New ways to do familiar Novell NetWare Integration tasks

New ways to do familiar Server Cluster tasks

New ways to do familiar command-line tasks

**D**

Devices: New ways to do familiar tasks

Devices: New ways to do familiar tasks

New ways to do DHCP tasks

New ways to do DHCP tasks

New ways to do Network Connections tasks

New ways to do familiar tasks for Disk Management

New ways to do familiar tasks for Disk Management

New ways to do familiar DNS tasks

New ways to do familiar DNS tasks

**F**

Find

Find computer

Find files or folders

Find printers

View Options: New ways to do familiar tasks

**H**
New ways to do familiar modem tasks

**I**
Inbox

Internet: New ways to do familiar tasks

Internet: New ways to do familiar tasks

**L**
New ways to do familiar Licensing tasks

New ways to do familiar performance monitoring tasks

**M**
New ways to do familiar Services for Macintosh tasks

New ways to do familiar Active Directory tasks

Media Player

Message Queuing (also known as MSMQ)

New ways to do familiar modem tasks

New ways to do familiar command-line tasks

MSMQ Explorer

Multimedia: New ways to do familiar tasks

My Briefcase: New ways to do familiar tasks

**N**
New ways to do Network Connections tasks

New ways to do familiar Network Load Balancing tasks

New ways to do familiar Network Monitor tasks

New ways to do familiar Novell NetWare Integration tasks

New ways to do familiar Novell NetWare Integration tasks

**O**
Outlook Express

**P**
New ways to do familiar performance monitoring tasks

New ways to do familiar modem tasks

Ports: New ways to do familiar tasks

New ways to do familiar Printing tasks

New ways to do familiar Printing tasks

New ways to do familiar tasks with TCP/IP

New ways to do VPN tasks

**R**
Regional Settings: New ways to do familiar tasks

Regional Settings: New ways to do familiar tasks

New ways to do familiar Remote Access tasks

New ways to do familiar Remote Desktop for Administration tasks

New ways to do familiar Routing tasks

New ways to do familiar Routing tasks

New ways to do familiar Remote Access tasks

**S**

SCSI Adapters: New ways to do familiar tasks

SCSI Adapters: New ways to do familiar tasks

New ways to do familiar Security tasks

Server: New ways to do familiar tasks

New ways to do familiar Server Cluster tasks

Active Directory Users & Computers: New ways to do familiar tasks

New ways to do familiar Services for Macintosh tasks

Server: New ways to do familiar tasks

New ways to do familiar SNMP tasks

Sounds

Multimedia: New ways to do familiar tasks

New ways to do familiar performance monitoring tasks

**T**

Tape Devices: New ways to do familiar tasks

New ways to do familiar Telephony tasks

New ways to do familiar tasks with TCP/IP

New ways to do familiar Telephony tasks

New ways to do familiar Terminal Server tasks

New ways to do familiar Remote Desktop for Administration tasks

**U**

Active Directory Users & Computers: New ways to do familiar tasks

New ways to do User Profiles tasks

**V**

View Options: New ways to do familiar tasks

New ways to do VPN tasks

**W**

CD Player: New ways to do familiar tasks

New ways to do WINS tasks

New ways to do familiar Network Load Balancing tasks

# Configuring Roles for Your Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring roles for your server**
The Configure Your Server Wizard provides a central location from which you can install or remove the server roles available on a server running Windows Server 2003. After installing a server role you can use Manage Your Server to manage that role.

Manage Your Server provides a central location from which you can manage the server roles installed through the Configure Your Server Wizard. Manage Your Server starts automatically the first time you log on to your server computer with administrative credentials.

- Server roles
- Typical setup for a first server
- File server role: Configuring a file server
- Print server role: Configuring a print server
- Application server role: Configuring an application server
- Mail server role: Configuring a mail server
- Terminal server role: Configuring a terminal server
- Remote access/VPN server role: Configuring a remote access/VPN server
- Domain controller role: Configuring a domain controller
- DNS server role: Configuring a DNS server
- DHCP server role: Configuring a DHCP server
- Streaming media server role: Configuring a streaming media server
- WINS server role: Configuring a WINS server
- To remove a server role

**Note**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

# Troubleshooting A-Z

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting A-Z**
This section contains pointers to troubleshooting information available in Help.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**
Active Directory: Troubleshooting Active Directory

Audio: Troubleshooting audio, imaging, and CD-burning peripheral devices

**B**
Backup Troubleshooting

**C**
Cameras: Troubleshooting audio, imaging, and CD-burning peripheral devices

CD-burning: Troubleshooting audio, imaging, and CD-burning peripheral devices

Certificate Services Troubleshooting

Certificate Templates Troubleshooting

Clustering: Server Cluster Troubleshooting

Compatibility: Troubleshooting compatibility issues

Connection Manager Administration Kit (CMAK): Troubleshooting process

Connection Point Services Troubleshooting

**D**
DHCP Troubleshooting

Disk Defragmenter: Troubleshooting Disk Defragmenter

Disk Management: Troubleshooting Disk Management

Disk Quotas: Troubleshooting Disk Quotas

DNS: Troubleshooting DNS

**E**
Emergency Management Services: Troubleshooting Emergency Management Services

Encrypting File System Troubleshooting

Event Viewer Troubleshooting

**I**
Imaging: Troubleshooting audio, imaging, and CD-burning peripheral devices

IAS Troubleshooting

IPSec Troubleshooting

**L**
Licensing Troubleshooting

Local Users and Groups: Troubleshooting Local Users and Groups

**M**
MMC: Troubleshooting MMC

Modems Troubleshooting

**N**
Network Connections Troubleshooting

Using Network Diagnostics

Network Load Balancing Clusters

Network Monitor: Troubleshooting Network Monitor

Novell NetWare Integration Troubleshooting

**P**
Performance monitoring: Troubleshooting Performance monitoring

Printing Troubleshooting

**R**
Remote Access Troubleshooting

Remote Installation Services: Troubleshooting Remote Installation Services

Remote Desktop for Administration: Troubleshooting Remote Desktop for Administration

Removable storage: Troubleshooting Removable Storage

Routing: Troubleshooting routing

**S**
Scanners: Troubleshooting audio, imaging, and CD-burning peripheral devices

Security Troubleshooting

Server Cluster Troubleshooting

Services for Macintosh Troubleshooting

Shared folders: Troubleshooting Shared Folders

Shutdown Event Tracker: Troubleshooting Shutdown Event Tracker

SNMP: Troubleshooting SNMP

Software Restriction Policies Troubleshooting

Stop messages: Troubleshooting Stop messages: general strategies

Stored User Name and Password Troubleshooting

System Monitor: Troubleshooting Performance monitoring

**T**
Task Scheduler: Troubleshooting Task Scheduler

TCP/IP: Troubleshooting TCP/IP

Telephony Troubleshooting

Windows Time Service Troubleshooting

**V**
VPN Troubleshooting

**W**
Windows ATM services: ATM Troubleshooting

Windows Installer Troubleshooting

WINS Troubleshooting

Windows Server 2003 Wireless Troubleshooting

# Step-by-step guides for deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Step-by-step guides for deployment**
The Microsoft Web site provides tutorials that you can use to learn about and deploy the operating system. These technical walkthroughs provide step-by-step instructions and illustrations for installing and configuring key features.

You can access the Step-by-Step Guides at the Microsoft Web site.

The following list includes some of the topics that you can find in these guides:

**Infrastructure**

- Building a common network infrastructure

**Active Directory**

- Setting up additional domain controllers
- Delegating control of objects
- Managing Active Directory
- Managing replication topology
- Using LDAP and VBScript to manage Active Directory
- Extending Active Directory using the Schema and display specifiers

**Management services**

- Understanding Group Policy
- Managing remote installation of the operating system
- Managing application deployments for groups of users and computers
- Managing user data and user settings

**Security services**

**Setting up and configuring certificates**

- Setting up a certification authority
- Configuring a domain to trust an external certification authority
- Managing Certificate Services
- Obtaining and managing certificates
- Configuring automatic certificate requests

**Using certificates**

- Requesting certificates using Web pages
- Mapping certificates to user accounts to be used on a Web site

**Additional security features**

- Configuring local security policy on individual computers
- Managing security in an enterprise network
- Encrypting files
- Using Kerberos V5
- Encrypting IP network traffic using Internet Protocol security (IPSec)

**File and print services**

- Creating logical views of directories and files

**Clustering**

- Installing a typical two-node cluster on Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition

# Windows interface administrative tool reference A-Z

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows interface administrative tool reference A-Z**

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**

Windows interface administrative tool reference A-Z: Access control

Windows interface administrative tool reference A-Z: Accessibility Options

Windows interface administrative tool reference A-Z: Active Directory Domains and Trusts

Windows interface administrative tool reference A-Z: Active Directory Schema

Windows interface administrative tool reference A-Z: Active Directory Sites and Services

Windows interface administrative tool reference A-Z: Active Directory Users and Computers

Windows interface administrative tool reference A-Z: Add Hardware

Windows interface administrative tool reference A-Z: Add or Remove Programs

Windows interface administrative tool reference A-Z: Authorization Manager

**B**

Windows interface administrative tool reference A-Z: Backup

**C**

Windows interface administrative tool reference A-Z: Certificate Templates

Windows interface administrative tool reference A-Z: Certificates

Windows interface administrative tool reference A-Z: Certification Authority

Windows interface administrative tool reference A-Z: Cluster Administrator

Windows interface administrative tool reference A-Z: Computer Management

Windows interface administrative tool reference A-Z: Configure Your Server Wizard

**D**

Windows interface administrative tool reference A-Z: Date and Time

Windows interface administrative tool reference A-Z: Device Manager

Windows interface administrative tool reference A-Z: DHCP

Windows interface administrative tool reference A-Z: Disk Cleanup

Windows interface administrative tool reference A-Z: Disk Defragmenter

Windows interface administrative tool reference A-Z: Disk Management

Windows interface administrative tool reference A-Z: Display

Windows interface administrative tool reference A-Z: Distributed File System

Windows interface administrative tool reference A-Z: DNS

Windows interface administrative tool reference A-Z: Domain Controller Security Policy

Windows interface administrative tool reference A-Z: Domain Security Policy

**E**

Windows interface administrative tool reference A-Z: Event Viewer

**F**

Windows interface administrative tool reference A-Z: Folder Options

Windows interface administrative tool reference A-Z: Fonts

**G**

Windows interface administrative tool reference A-Z: Game Controllers

Windows interface administrative tool reference A-Z: Group Policy

**I**

Windows interface administrative tool reference A-Z: Indexing Service

Windows interface administrative tool reference A-Z: Infrared and Wireless Link

Windows interface administrative tool reference A-Z: Internet Authentication Service

Windows interface administrative tool reference A-Z: Internet Options

**K**

Windows interface administrative tool reference A-Z: Keyboard

**L**

Windows interface administrative tool reference A-Z: Licensing (Administrative Tools)

Windows interface administrative tool reference A-Z: Licensing (Control Panel)

Windows interface administrative tool reference A-Z: Local Security Settings

Windows interface administrative tool reference A-Z: Local Users and Groups

**M**

Windows interface administrative tool reference A-Z: Microsoft Management Console

Windows interface administrative tool reference A-Z: Mouse

Windows interface administrative tool reference A-Z: My Computer Information

**N**

Windows interface administrative tool reference A-Z: Network Connections

Windows interface administrative tool reference A-Z: Network Diagnostics

Windows interface administrative tool reference A-Z: Network Load Balancing Manager

Windows interface administrative tool reference A-Z: Network Monitor

**O**

Windows interface administrative tool reference A-Z: ODBC Data Source Administrator

**P**

Windows interface administrative tool reference A-Z: Performance

Windows interface administrative tool reference A-Z: Phone and Modem Options

Windows interface administrative tool reference A-Z: Power Options

Windows interface administrative tool reference A-Z: Printers and Faxes

**R**

Windows interface administrative tool reference A-Z: Regional and Language Options

Windows interface administrative tool reference A-Z: Remote Assistance

Windows interface administrative tool reference A-Z: Remote Storage

Windows interface administrative tool reference A-Z: Removable Storage

Windows interface administrative tool reference A-Z: Routing and Remote Access

**S**

Windows interface administrative tool reference A-Z: Scanners and Cameras

Windows interface administrative tool reference A-Z: Security Configuration and Analysis

Windows interface administrative tool reference A-Z: Security Templates

Windows interface administrative tool reference A-Z: Services

Windows interface administrative tool reference A-Z: Sounds and Audio Devices

Windows interface administrative tool reference A-Z: Speech

Windows interface administrative tool reference A-Z: Stored User Names and Passwords

Windows interface administrative tool reference A-Z: System

Windows interface administrative tool reference A-Z: System Configuration Utility

Windows interface administrative tool reference A-Z: System Information

**T**

Windows interface administrative tool reference A-Z: Task Manager

Windows interface administrative tool reference A-Z: Task Scheduler

Windows interface administrative tool reference A-Z: Telephony

Windows interface administrative tool reference A-Z: Terminal Server Licensing

Windows interface administrative tool reference A-Z: Terminal Services Configuration

Windows interface administrative tool reference A-Z: Terminal Services Manager

**W**

Windows interface administrative tool reference A-Z: WINS

Windows interface administrative tool reference A-Z: Windows Management Instrumentation Command-line (WMIC)

# Command-line reference A-Z

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Command-line reference A-Z**
To find information about a command, on the A-Z menu, click the letter that the command starts with, and then click the command name.

In addition to the tools installed with the Windows Server 2003 family there are over 40 support tools that you can install directly from the product CD. You can use these tools to diagnose and resolve computer problems. For more information about these support tools and how to install them, see Install Windows Support Tools.

For more information about changes to the functionality of MS-DOS commands, new command-line tools, command shell functionality, configuring the command prompt, and automating command-line tasks, see Command-line reference.

You must run command-line tools at the prompt of the Cmd.exe command interpreter. To open Command Prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**. To view help at the command line, at the command prompt, type:

*CommandName* **/?**

As a security best practice, consider using an account without administrative credentials to perform routine, nonadministrative tasks. To perform only specific tasks with an account with administrative credentials, you can use the Runas command.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**
Adprep

Arp

Assoc

At

Atmadm

Attrib

**B**
Using batch files

Bootcfg

Break

**C**
Cacls

Call

Certreq

Certutil

Terminal Services change commands

Chcp

chdir

chkdsk

Chkntfs

Choice

Cipher

Clip

Cls

Cluadmin

Cluster

Cmd

Cmdkey

Cmstp

Color

Command shell overview

Comp

Compact

Convert

Copy

Cprofile

CScript

CSVDE

**D**

Date

Dcgpofix

Defrag

Del (erase)

Dfscmd

Dir

Directory service command-line tools

Diskcomp

Diskcopy

DiskPart

Doskey

Driverquery

Dsadd

Dsget

Dsmod

Dsmove

Dsquery

Dsrm

**E**

Echo

Endlocal

Eventcreate

Eventquery.vbs

Eventtriggers

Evntcmd

Exit

Expand

**F**

Fc

Using filters

Find

Findstr

Finger

Flattemp

For

Forfiles

Format

Freedisk

Fsutil

Ftp

Ftp subcommands

Ftype

**G**

Getmac

Gettype

Goto

Gpresult

Gpupdate

Graftabl

**H**

Help

Helpctr

Hostname

**I**

If

Iisapp.vbs: IIS application query script

Iisback.vbs: IIS backup management script

iiscnfg.vbs: IIS configuration script

iisext.vbs: IIS Web service extension script

iisftp.vbs: IIS FTP site management script

IISFtpdr.vbs: IIS FTP directory script

iisreset

iisvdir.vbs: IIS virtual directory script

iisweb.vbs: IIS Web site management script

Inuse

Ipconfig

Ipxroute

Irftp

**L**

Label

LDIFDE

Lodctr

Logman

Lpq

Lpr

**M**

Macfile

Md

Mkdir

Mmc

Mode

More

Mountvol

Move

Using the IP multicasting utilities

MS-DOSsubsystem commands

Msiexec (command-line options)

Msinfo32

**N**

Nbtstat

Net services overview

Net services commands

Netsh overview

Netsh commands for AAAA

Netsh commands for DHCP

Netsh diagnostic (diag) commands

Netsh commands for Interface IP

Netsh commands for Interface IPv6

Netsh commands for Interface Portproxy

Netsh commands for Internet Protocol security

Netsh commands for Network Bridge

Netsh commands for remote access (ras)

Netsh commands for Routing

Netsh commands for RPC

Netsh commands for WINS

Netstat

Nlb

Nlbmgr

Nslookup

Nslookup subcommands

Ntbackup

Ntcmdprompt

Ntdsutil

Ntsd

**O**

Openfiles

**P**

Pagefileconfig.vbs

Path

Pathping

Pause

Pbadmin

Pentnt

Perfmon

Ping

Popd

Powercfg

Print

Prncnfg.vbs

Prndrvr.vbs

Prnjobs.vbs

Prnmngr.vbs

Prnport.vbs

Prnqctl.vbs

Prompt

Pushd

**Q**
Terminal Services query commands

**R**
Rasdial

Rcp

Rd

Recover

Using command redirection operators

Reg

Regsvr32

Relog

Rem

Rename (ren)

Replace

Reset session

Rexec

Risetup

Rmdir

Route

Rsh

Rsm

Rss

Runas

**S**
SC

Schtasks

Secedit

Services for Macintosh commands

Set

Setlocal

Setx

Shift

Shutdown

Sort

Start

Subst

Sysocmgr

Systeminfo

System File Checker (sfc)

**T**
Takeown

Tapicfg

Taskkill

Tasklist

Tcmsetup

TCP/IP utilities and services

Telnet Command-line Reference for Telnet Server

Terminal Services commands

Tftp

Time

Timeout

Title

Tracerpt

Tracert

Tree

Tsecimp

Type

Typeperf

**U**
Unlodctr

**V**
Ver

Verify

Vol

Vssadmin

**W**
W32tm

waitfor

Where

Whoami

Winnt

Winnt32

Winpop

Wlbs

Windows Management Instrumentation Command-line (WMIC) tool

WScript

**X**

Xcopy

# Windows Resource Kit Tools Help

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Resource Kit Tools Help**

**Important**

- Some of the links on this page are to English-only content on the Web, regardless of the language of your operating system. In addition, if you install a Microsoft® Windows® Deployment or Resource Kit on a non-English operating system or on an operating system with a Multilingual User Interface Pack (MUI), you will see some English content mixed with non-English content in Help and Support Center. This behavior should be expected if you browse through tools documentation in the Tools By Category listing in the Tools Center in Help and Support Center or if you search anywhere in Help and Support Center by a tool name.

Windows Resource Kit Tools Help includes documentation on tools designed to help system administrators and IT professionals deploy, support, and maintain Windows Server 2003 family operating systems.

You can use Windows Resource Kit Tools to manage the Active Directory® directory service, administer security features, work with the registry, automate recurring jobs, and perform many other important tasks.

Windows Resource Kit Tools are not part of the Windows operating system. You must install them separately.

For information about Windows Resource Kit Tools and Tools Help, or to browse additional Windows Deployment and Resource Kit documentation on the Web, go to the Windows Deployment and Resource Kits Web site.

# Administration and Scripting Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administration and scripting tools**
This section covers:

- Management Strategies and Tools
- Command-line reference
- Configuration and Management Tools
- Remote Administration Tools
- Monitoring and Status Tools
- Scheduling and Automation Tools
- Scripting Tools

# Management Strategies and Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Management Strategies and Tools**
Windows Server 2003 family operating systems contain a rich set of tools and features that you can use to manage your Windows environment.

For information about the technologies, issues, and methods to consider when deciding which tools to use to perform management tasks, see Management strategies and tools overview.

To find the tool or feature you can use to perform a specific task, see Using Management Tools and Features.

# Management strategies and tools overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Management strategies and tools overview**
When deciding to use a tool or feature available in the Microsoft® Windows Server 2003 family operating systems, it is helpful to have an understanding of:

- The technology or feature set a tool is based on, or part of
- Whether you can perform a task using a graphical user interface (GUI)-based tool, command-line tool, or script
- How security and remote administration issues apply to your daily and long-term IT operations goals and objectives.

Adding the scope and limitations of your IT environment and corporate business goals to your analysis of this information, you can make optimal choices about which tools to use.

**Management technologies and feature sets**
Most of the tools and features available in Windows Server 2003 family operating systems are based on, or are a subset of, the following technologies and feature sets. For more information about each technology or feature set, see its corresponding link under **Overview**.

| Technology or feature set | Overview |
|---|---|
| Microsoft® Active Directory directory service | Active Directory Overview |
| Group Policy | Group Policy overview |
| Microsoft® IntelliMirror management technologies | IntelliMirror |
| Microsoft Management Console (MMC) | MMC overview |
| Terminal Services | Terminal Services overview |
| Windows Management Instrumentation (WMI) | Windows Management Instrumentation overview |

**Choosing a management method**
You can perform management tasks using one, some, or all three of the following types of management tools:

- graphical user interface (GUI)-based
- command-line
- script or script-based

You can perform many of the same tasks using any of these methods, but based on your network environment and management goals and objectives, there might be advantages to using one type of tool over another.

**Advantages to using GUI-based tools**

- You can use Microsoft Management Console (MMC) Consoles to group the tools you use most often so that you can view and modify your client, server, and network components and services from a single location.
- You are less likely to type a command and related parameters incorrectly because you configure tasks through a GUI.
- You do not need to remember a lot of syntax and parameters--again, there is less opportunity to make mistakes.

**Advantages to using command-line tools**

- You can more easily automate a variety of management tasks.
- You have more options to customize a task because a number of parameters might be available to use with the command-line tool.

**Advantages to using scripts or script-based tools**

- You can customize a task or set of tasks down to the last detail.
- You can use scripts to fully automate a task or set of tasks.
- You can run a script from either a GUI-based or command-line tool

For more information about creating and using scripts, see the "System Administration Scripting Guide" at the Microsoft Windows Resource Kits Web site.

**Issues important to system administrators**
As more and more business is done over the Internet and from remote locations, security and remote administration will continue to be among the most important considerations in managing your Windows environment. For more information about these issues, click the corresponding links under **Overview**.

| **Issue** | **Overview** |
|---|---|
| Security | Security overview |
| Remote Administration | Remote administration overview |

# Using Management Tools and Features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using management tools and features**

Each of the following topics represents day-to-day activities you perform to manage your Windows environment. Within each activity are common tasks related to the activity, with corresponding links to more information about both graphical user interface (GUI)-based tools and command-line tools available in Windows Server 2003 family operating systems.

- Automating processes and tasks
- Deploying and managing applications
- Managing the Active Directory directory service
- Managing computers
- Managing disks and data
- Managing network services
- Managing printers and faxes
- Managing security
- Managing servers
- Monitoring and controlling services on computers

# Automating processes and tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automating processes and tasks**

- Using the Windows interface
- Using command-line tools

**Using the Windows interface**

The following table lists some of the most common tasks for automating processes and tasks, and the tools you can use to perform them. For more information, click the corresponding link under **Tool or Feature**.

| Task | Tool or feature |
|---|---|
| Automate routine or repetitive tasks. | Using batch files |
| Assign scripts to define settings for a group of users. | Group Policy (pre-GPMC) |
| Assign tasks that will be performed whenever a user logs on to a particular computer, such as mapping network drivers, starting background processes, and initializing user environment variables. | Logon Scripts |
| Run a group of commands in batch mode against a specified computer; save a configuration script in a text file for archival purposes, or to help you configure other computers. | Netsh overview |
| Schedule any script, program or document to run at a scheduled time. | Task Scheduler |
| Write a management script or batch file to automate administrative tasks. | Windows Management Instrumentation Command-line |
| Run scripts by clicking a script file on the Windows desktop or by typing the name of a script file at the command prompt. | Windows Script Host |
| Synchronize time on the computers in your network. | Windows Time Service |

For more information about automating installations, see "Automating and Customizing Installations" at the Microsoft Windows Resource Kits Web site.

For more information about creating and using scripts, see the "System Administration Scripting Guide" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for automating processes and tasks, and the command-line tools you can use to perform them. For more information, click the corresponding link under **Command-line tool**.

| Task | Command-line tool |
|---|---|
| Call one batch program from another without stopping the parent batch program. | Call |
| Choose from among a set of user option keys such as Yes, No and Continue in a batch program. | Choice |
| Redirect command output from the command line to the Clipboard. | Clip |
| Recalls commands, edits command lines, and creates macros. | Doskey |

| | |
|---|---|
| End localization of environment changes in a batch file, restoring environment variables to their values before the matching **setlocal** command. | Endlocal |
| Create a custom event in a specified event log. | Eventcreate |
| List the events and event properties from one or more event logs. | Eventquery.vbs |
| Display and configure event triggers on local or remote computers. | Eventtriggers |
| Exit the current batch script or the Cmd.exe program (that is, the command interpreter) and return to the program that started Cmd.exe or to the Program Manager. | Exit |
| Run a specified command for each file in a set of files. | For |
| Select files in a folder or tree for batch processing. | Forfiles |
| Set the system environment variable %ERRORLEVEL% to the value associated with the specified Windows operating system. | Gettype |
| Within a batch program, directs the operating system to a label; The operating system then processes the commands that begin on the next line. | Goto |
| Perform conditional processing in batch programs. | If |
| Sets the command path in the PATH environment variable, which is the set of directories used to search for executable files. Used without parameters, **path** displays the current command path. | Path |
| Suspend the processing of a batch program and display a message that prompts the user to press any key to continue. | Pause |
| Change the current directory to the directory stored by the **pushd** command. | Popd |
| Store the name of the current directory for use by the **popd** command before changing the current directory to the specified directory. | Pushd |
| Include comments (remarks) in a batch file or in your configuration files. | Rem |
| Communicate with the Service Controller and installed services; retrieve or set control information about services. | SC |
| Schedule commands and programs to run periodically or at a specific time; manage the Task Scheduler service | Schtasks |
| Display, set, or remove environment variables. | Set |
| Start localization of environment variables in a batch file. | Setlocal |
| Set environment variables in the local or system environment, without requiring programming or scripting. | Setx |
| Change the position of batch parameters in a batch file. | Shift |
| Pause the command processor for the specified number of seconds. Timeout is typically used in batch files. | Timeout |
| Diagnose problems occuring with Windows Time. | W32tm |
| Synchronize multiple computers across a network. | waitfor |
| Locate and display all files that match the given parameter; by default, the search is done along the current directory and in the PATH environment variable. | Where |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Deploying and managing applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying and managing applications**

- Using the Windows interface
- Using command-line tools

**Using the Windows interface**

The following table lists some of the most common tasks for deploying and managing applications, and the tools you can use to perform them. For more information, click the corresponding link under **Tool or Feature**.

| Task | Tool or feature |
|---|---|
| Create a specific desktop configuration of installed applications for a group of users; manage desktop configurations for a group of users. | Group Policy (pre-GPMC) |
| Set up and manage an Internet Information Services (IIS) server to host Web applications and services. | Internet Information Services |
| Track and maintain software licenses. | Administering and tracking licenses |
| Create, save, and open administrative tools (in the form of snap-ins) that manage hardware, software and network components of a local or remote computer, or multiple computers. | Microsoft Management Console |
| Set up remote computers with installation images, without requiring the installation compact disc. | Remote Installation Services |
| Install and remove software; modify, repair, and advertise existing applications. | Windows Installer |
| Download and install updates on your computer, such as security fixes, critical updates, the latest Help files, and device drivers. | Windows Update |

For more information about deploying applications to servers and client computers, see "Planning Server Deployments" at the Microsoft Windows Resource Kits Web site.

For more information about automating and customizing installations, see "Automating and Customizing Installations" at the Microsoft Windows Resource Kits Web site.

For more information about software installation tools, see "Part Two: Change and Configuration Management" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for deploying and managing applications, and the command-line tools you can use to perform them. For more information, click the corresponding link under **Command-line tool**.

| Task | Command-line tool |
|---|---|
| Create and manage backup copies of the Internet Information Services (IIS) configuration (metabase and schema) of a remote or local computer. | lisback.vbs: IIS backup management script |
| Import or export all or selected elements of an Internet Information Services (IIS) metabase on a local or remote computer. | liscnfg.vbs: IIS configuration script |
| Create, delete, list, start, stop, or pause FTP sites on local or remote servers running Windows Server 2003 with Internet Information Services (IIS) 6.0. | lisftp.vbs: IIS FTP site management script |

| | |
|---|---|
| Create, delete, or display virtual directories on FTP sites of servers running Windows Server 2003 with Internet Information Services (IIS) 6.0. | IISFtpdr.vbs: IIS FTP directory script |
| Create or delete virtual directories of Web sites on servers running Windows Server 2003 with Internet Information Services (IIS) 6.0. | Iisvdir.vbs: IIS virtual directory script |
| Create, delete, list, start, stop, or pause Web sites on servers running Windows Server 2003 with Internet Information Services (IIS) 6.0. | Iisweb.vbs: IIS Web site management script |
| Install, modify, and perform operations in conjunction with Windows Installer. | Msiexec (command-line options) |
| Manage Remote Installation Services on a server. | Risetup |
| Communicate with the Service Controller and installed services; retrieve or set control information about services. | SC |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Managing the Active Directory directory service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing the Active Directory directory service**

- Using the Windows interface
- Using command-line tools

**Using the Windows interface**

The following table lists some of the most common tasks for managing the Active Directory directory service, and the command-line tools you can use to perform them. For more information, click the corresponding link under **Tool or Feature**.

| Task | Tool or feature |
|---|---|
| Administer Active Directory. | Administering Active Directory |
| Deploy Active Directory. | Deploying Active Directory |
| Create a specific desktop configuration of installed applications for a group of users, and manage desktop configurations for a group of users. | Group Policy (pre-GPMC) |
| Secure and manage user accounts and groups stored locally on your computer. | Local Users and Groups |
| Set up your directory structure. | Introduction to Active Directory |

For more information about deploying Active Directory, see "Part 1: Designing and deploying directory services" at the Microsoft Windows Resource Kits Web site.

For more information about managing Active Directory, see the "Directory services guide" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for managing the Active Directory directory service, and the command-line tools you can use to perform them. For more information, click the corresponding link under **Command-line tool**.

| Task | Command-line tool |
|---|---|
| Import or export data from the Active Directory directory service using files that store data in the comma-separated variable (CSV) format. | CSVDE |
| Restore default Group Policy objects to their original state (after initial installation). | Dcgpofix |
| Manage objects and information about those objects in Active Directory. | Directory service command-line tools |
| Display Group Policy settings and Resultant Set of Policy (RSOP) for a user or computer. | Gpresult |
| Refresh local and Active Directory-based Group Policy settings, including security settings. | Gpupdate |

| | |
|---|---|
| Create, modify, or delete directory objects on computers running Windows Server 2003 family operating systems or Windows XP Professional; extend the schema, export Active Directory user and group information to other applications or services, or populate Active Directory with data from other directory services. | LDIFDE |
| Manage the Active Directory database. | Ntdsutil |
| Administer phone books. | Pbadmin |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Managing computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing computers**

- Using the Windows interface
- Using command-line tools

**Using the Windows interface**

The following table lists some of the most common tasks for managing computers, and the tools you can use to perform them. For more information, click the corresponding link under **Tool or Feature**.

| Task | Tool or feature |
|---|---|
| On a single local or remote computer: create and manage shared resources, view a list of users who are connected to the computer, or view device configurations and add new device drivers. | Computer Management |
| Use a graphical view of the hardware that is installed on your computer; if necessary, update the drivers for hardware devices, modify hardware settings and troubleshoot problems. | Device Manager |
| Monitor security and other events that occur on a local or remote computer, including hardware and software-generated events, and system problems. | Event Viewer |
| Create a specific desktop configuration for a group of users (for example, programs available to users, programs that appear on the user's desktop, Start menu options). | Group Policy (pre-GPMC) |
| Install hardware and manage hardware-related issues. | Hardware |
| Secure and manage user accounts and groups stored locally on a computer. | Local Users and Groups |
| Create, save, and open administrative tools (in the form of snap-ins) that manage hardware, software, and network components of a local or remote computer, or multiple computers. | Microsoft Management Console |
| Manage how a computer connects to the network. | Network Connections Overview |
| Connect computers to printers. | Netsh overview |
| View or edit configuration information on a computer. | The Registry |
| On local or remote computers: display system configuration information, such as hardware configurations, computer components, installed software, and driver information. | System Information |
| View and change system properties on a local or remote computer. | System Properties |
| Manage user profiles. | User profiles overview |
| Manage a single or multiple local or remote computers in a single command. | Windows Management Instrumentation Command-line |
| Synchronize time on the computers in your network. | Windows Time Service |

For more information about deploying computers, see "Part 1: Client Deployments" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for managing computers, and the command-line tools you can use to perform them. For more information, click the corresponding link under **Command-line tool**.

| Task | Command-line tool |
| --- | --- |
| Configure, query, or change Boot.ini file settings. | Bootcfg |
| Manage Terminal Services. | Terminal Services change commands |
| Display or change the identification number of an active console code page. | Chcp |
| Install or removes a Connection Manager service profile. | Cmstp |
| Displays the current system date setting. Used without parameters, **date** displays the current system date setting and prompts you to type a new date. | Date |
| Display properties for all installed device drivers. | Driverquery |
| Create a custom event in a specified event log. | Eventcreate |
| List the events and event properties from one or more event logs. | Eventquery.vbs |
| Display and configure event triggers on local or remote computers. | Eventtriggers |
| Display a comprehensive view of the hardware, system components, and software environment on the local computer. | Msinfo32 |
| Display and configure the virtual memory paging file on a computer. | Pagefileconfig.vbs |
| Detect floating point division error (if present) in the Pentium chip, disable floating point hardware, or start floating point emulation. | Pentnt |
| Control the power settings on a computer. | Powercfg |
| Check Terminal Services status information. | Terminal Services query commands |
| Automate the connection process for any Microsoft client. | Rasdial |
| Add, change, or displays registry subkey information and values in registry entries. | Reg |
| Register .dll files as command components in the registry. | Regsvr32 |
| Reset (delete) a session from a terminal server. | Reset session |
| Run commands on remote computers running the Rexec service (daemon); authenticates a user name on a remote computer before executing a command. | Rexec |
| Run commands on remote computers running the RSH service or daemon. | Rsh |
| Communicate with the Service Controller and installed services; retrieve or set control information about services. | SC |
| Shut down or restart a local or remote computer. | Shutdown |
| Install a limited set of optional components. | Sysocmgr |
| Display detailed configuration information about a computer and its operating system, including operating system configuration, security information, product ID, and hardware properties, such as RAM, disk space, and network cards. | Systeminfo |
| Scan and verify the versions of all protected system files after you restart your computer. | System File Checker (sfc) |

| | |
|---|---|
| Set up or disable the telephony application programming interface (TAPI) client. | Tcmsetup |
| Administer a local or a remote computer running Telnet Server. | Telnet Command-line Reference for Telnet Server |
| Administer Terminal Services. | Terminal Services commands |
| Display or set the system time. | Time |
| Import assignment information from an Extensible Markup Language (XML) file into the TAPI server security file (Tsec.ini); you can also use this command to display the list of TAPI providers and the lines devices associated with each of them, validate the structure of the XML file without importing the contents, and check domain membership. | Tsecimp |
| Displays the Windows XP version number. | Ver |
| Diagnose problems occuring with Windows Time. | W32tm |
| Install or upgrade to a product in the Windows Server 2003 family. | Winnt32 |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Managing disks and data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing disks and data**


- Using the Windows interface
- Using command-line tools


**Using the Windows interface**

The following table lists some of the most common tasks for managing disks and data, and the tools you can use to perform them. For more information, click the corresponding link under **Tool or Feature**.

| Task | Tool or feature |
|---|---|
| Back up and restore data. | Backup overview |
| Set properties for storage devices. | Computer Management |
| Analyze local volumes and consolidate fragmented files and folders on your hard disk. | Disk Defragmenter |
| Manage hard disks and the volumes they contain. | Disk Management |
| Monitor and limit disk space use on NTFS volumes. | Disk Quotas |
| Manage files and folders. | Managing Files and Folders |
| Work with network files and program even when you are not connected to the network. | Offline Files overview |
| Extend disk space on your server without adding more hard disks. | Remote Storage |
| Track a computer's removeable storage media and manage the hardware libraries that contain them. | Removable Storage |
| View the contents of network folders as they existed at points of time in the past. | Shadow Copies for Shared Folders |

For more information about managing disks and data, see "Part One: Managing Data and Devices" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for managing disks and data, and the command-line tools you use to perform them. For more information, click the corresponding link under **Command-line tool**.

| Task | Command-line tool |
|---|---|
| Display or modify file name extension associations. | Assoc |
| Display, set, or remove the read-only, archive, system, and hidden attributes assigned to files or directories. | Attrib |
| Display the name of the current directory or changes the current folder. | Chdir |
| Create and display a status report for a disk based on the file system; correct errors on the disk. | chkdsk |
| Display or specify whether automatic system checking is scheduled to be run on a FAT, FAT32, or NTFS volume when the computer is started. | Chkntfs |
| Display or alter the encryption of folders and files on NTFS volumes. | Cipher |

| | |
|---|---|
| Compare the contents of two files or sets of files byte by byte on the same drive or on different drives, or in the same directory or in different directories. | Comp |
| Display and alter the compression of files or directories on NTFS partitions. | Compact |
| Convert FAT and FAT32 volumes to NTFS. | Convert |
| Copy one or more files from one location to another. | Copy |
| Clean specified profiles of wasted space and, if user-specific file associations are disabled, remove these associations from the registry. | Cprofile |
| Locate and consolidate fragmented boot files, data files, and folders on local volumes. | Defrag |
| Delete specified files. | Del |
| Manage a distributed file system. | Dfscmd |
| Display a list of a directory's files and subdirectories. | Dir |
| Compare the contents of two floppy disks. | Diskcomp |
| Copy the contents of the floppy disk in the source drive to a formatted or unformatted floppy disk in the destination drive. | Diskcopy |
| Manage objects (disks, partitions, or volumes) by using scripts or direct input. | DiskPart |
| Expand one or more compressed files. | Expand |
| Compare two files and display the differences between them. | Fc |
| Used in conjunction with the command redirection pipe character (\|), a command filter is a command within a command that reads the command's input, transforms the input, and then writes the output. Filter commands help you sort, view, and select parts of a command output. Filter commands divide, rearrange, or extract portions of the information that passes through them. | Using filters |
| Searches for a specific string of text in a file or files. After searching the specified file or files, **find** displays any lines of text that contain the specified string. | Find |
| Searches for patterns of text in files using regular expressions. | Findstr |
| Enable or disable flat temporary folders. | Flattemp |
| Format the disk in the specified volume to accept Windows files. | Format |
| Check to see if the specified amount of disk space is available before continuing with an installation process. | Freedisk |
| Perform many FAT and NTFS file system related tasks, such as managing reparse points, managing sparse files, or dismounting a volume. | Fsutil |
| Displays or modifies file types used in file name extension associations. Used without parameters, **ftype** displays the file types that have open command strings defined. | Ftype |
| Replace locked operating system files. | Inuse |
| Create, change, or delete the volume label (that is, the name) of a disk. | Label |
| Manage File Server for Macintosh servers, volumes, directories, and files. | Macfile |
| Creates a directory or subdirectory. | Mkdir |
| Create, delete, or list a volume mount point. **Mountvol** is a way to link volumes without requiring a drive letter. | Mountvol |
| Move one or more files from one directory to a specified directory. | Move |

| | |
|---|---|
| Back up disks and data. | Ntbackup |
| Query, display, or disconnect open files or files opened by network users; enable or disable the system Maintain Objects List global flag. | Openfiles |
| Copy files between a Windows Server 2003 family, or Windows XP computer and a system running **rshd**, the remote shell service (daemon). | Rcp |
| Remove (delete) a directory. | Rd |
| Recover readable information from a bad or defective disk. | Recover |
| Change the name of a file or a set of files. | Rename (ren) |
| Replace files in the destination directory with files in the source directory that have the same name; add unique file names to the destination directory. | Replace |
| Remove (delete) a directory. | Rmdir (Rd) |
| Manage media resources using Removable Storage | Rsm |
| Manage Remote Storage. | Rss |
| Read input, and sort data; write the results to the screen, to a file, or to another device. | Sort |
| Associate a path with a drive letter. | Subst |
| Display the directory structure of a path or of the disk in a drive graphically. | Tree |
| Displays the contents of a text file. Use the **type** command to view a text file without modifying it. | Type |
| Display the disk volume label and serial number, if they exist. | Vol |
| Displays current volume shadow copy backups and all installed shadow copy writers and providers in the command window. The Windows Server 2003 family version of this command contains more parameters than the one that is in Windows XP. | Vssadmin |
| Locate and display all files that match the given parameter. | Where |
| Copy files and directories, including subdirectories. | Xcopy |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Managing network services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing network services**

- Using the Windows interface
- Using command-line tools

**Using the Windows interface**

The following table lists some of the most common tasks for managing network services, and the tools you can use to perform them. For more information, click the corresponding link under **Tool or Feature**.

| **Task** | **Tool or feature** |
|---|---|
| Manage availability and scalability of objects on your network. | Availability and Scalability |
| Set up and configure a DHCP or DNS server. | Configuring Roles for Your Server |
| Configure communication links for network connections, modems, and wireless networking. | Configuring Communication Links |
| Configure communication services for telephony. | Configuring Communication Services |
| Enable dynamic allocation of IP addresses and other related configuration details for clients on your network, by administering the DHCP service. | DHCP |
| Administer the Domain Name System (DNS) service. DNS enables client computer to resolve DNS domain names of network objects. Once domain names are resolved on your computer, you can access resources offered by other computers on your network or other networks, such as the Intranet. | DNS |
| Configure and administer Internet and E-mail services. | Internet and E-mail Services |
| Configure Internet Protocol Security. | Internet Protocol Security (IPSec) |
| Manage bandwidth issues, such as Quality of Service (QoS) and Windows Asynchronous Transfer Mode (ATM) Services. | Managing Bandwidth |
| Choose and manage remote connections and methods through routing and remote access, Connection Manager, Connection Point Services, virtual private networks, and Internet Authentication Service. | Managing Remote Connections |
| Create, save, and open administrative tools (in the form of snap-ins) that manage hardware, software, and network components of a local or remote computer, or multiple computers. | Microsoft Management Console |
| Set up, configure and manage network connections. | Network Connections |

| | |
|---|---|
| Display information about your computer, operating system, network, and network adapters; perform standard tests for network connectivity. | Using Network Diagnostics |
| Gather information about network traffic status. | Network Monitor |
| Obtain status information about network hosts such as servers, workstations, or routers, bridges, and hubs from a centrally-located computer on network using TCP/IP. | SNMP |
| Enable Windows-based computers to connect and share information with other Microsoft and non-Microsoft systems. | TCP/IP Core Networking Service |
| Resolve names for computer and groups on a network using the NetBIOS naming protocol. | WINS |

For more information about deploying network services, see "Deploying Network Services" at the Microsoft Windows Resource Kits Web site.

For more information about managing network services, see the "Networking Guide" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for managing network services, and the command-line tools you can use to perform them. For more information, click the corresponding link under **Command-line tool**.

| **Task** | **Command-line tool** |
|---|---|
| Display and modify entries in the Address Resolution Protocol (ARP) cache | Arp |
| Monitor connections and addresses that are registered by the ATM Call Manager on an asynchronous transfer mode (ATM) network; display statistics for incoming and outgoing calls on ATM adapters. | Atmadm |
| Transfer files to and from a computer running a File Transfer Protocol (FTP) server service such as Internet Information Services. | Ftp |
| Display the media access control (MAC) address and list of network protocols associated with each address for all network cards in each computer, either locally or across a network. | Getmac |
| Display the host name portion of the full computer name of the computer. | Hostname |
| Display all current TCP/IP network configuration values or refresh Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS) settings. | Ipconfig |
| Display or modify information about the routing tables used by the IPX protocol. | Ipxroute |
| Send files over an infrared link. | Irftp |
| Display the configuration of a multicast router to aid in the troubleshooting of multicast forwarding and routing problems | Using the IP multicasting utilities |
| Display NetBIOS over TCP/IP (NetBT) protocol statistics, NetBIOS name tables for both the local computer and remote computers, or the NetBIOS name cache. | Nbtstat |
| Check network status and other configuration information. | Net services overview |
| Display or modify the network configuration of a local or remote computer that is currently running. | Netsh overview |
| Display active TCP connections, ports on which the computer is listening, Ethernet statistics, the IP routing table, IPv4 statistics (for the IP, ICMP, TCP, and UDP protocols), and IPv6 statistics (for the IPv6, ICMPv6, TCP over IPv6, and UDP over IPv6 protocols). | Netstat |

| | |
|---|---|
| Manage network load balancing. | Nlb |
| Display information that you can use to diagnose Domain Name System (DNS) infrastructure. . | Nslookup |
| Obtain information about network latency and network loss at intermediate hops between a source and destination; send multiple Echo Request messages to each router between a source and destination over a period of time and then compute results based on the packets returned from each router. | Pathping |
| Verify IP-level connectivity to another TCP/IP computer by sending Internet Control Message Protocol (ICMP) Echo Request messages. | Ping |
| Display and modify the entries in the local IP routing table. | Route |
| Manage TCP/IP. | TCP/IP utilities and services |
| Transfer files to and from a remote computer, typically a computer running UNIX, that is running the Trivial File Transfer Protocol (TFTP) service or daemon. | Tftp |
| Determine the path taken to a destination by sending Internet Control Message Protocol (ICMP) Echo Request or ICMPv6 messages to the destination with incrementally increasing Time to Live (TTL) field values. | Tracert |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Managing printers and faxes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing printers and faxes**

- Using the Windows interface
- Using command-line tools

**Using the Windows interface**

The following table lists some of the most common tasks for managing printers and faxes, and the tools you can use to perform them. For more information, click the corresponding link under "Tool or feature."

| Task | Tool or feature |
|---|---|
| Enable network users to send and receive faxes through remote fax printer connections. | Fax Service Manager |
| Install, connect to, and manage printers. | Printing |

For more information about printing, see "Network Printing" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for managing printers and faxes, and the command-line tools you can use to perform them. For more information, click the corresponding link under "Command-line tool."

| Task | Command-line tool |
|---|---|
| Start or stop the File Server for Macintosh service or the Print Server for Macintosh service. | Services for Macintosh commands |
| Display the status of a print queue on a computer running Line Printer Daemon (LPD). | Lpq |
| Send a file to a computer running Line Printer Daemon (LPD) in preparation for printing. | Lpr |
| Send a text file to a printer. | Print |
| Configure or display configuration information about a printer. | Prncnfg.vbs |
| Add, delete, or list printer drivers. | Prndrvr.vbs |
| Pause, resume, cancel, or lists print jobs. | Prnjobs.vbs |
| Add, delete, or list printers or printer connections; set or display the default printer. | Prnmngr.vbs |
| Create, delete, or list standard TCP/IP printer ports, in addition to displaying and changing port configuration. | Prnport.vbs |
| Print a test page; pause or resume a printer, or clear a printer queue. | Prnqctl.vbs |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Managing security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing security**

- Using the Windows interface
- Using command-line tools

**Using the Windows interface**

The following table lists some of the most common tasks for managing security, and the tools you can use to perform them. For more information, click the corresponding link under **Tool or Feature**.

| Task | Tool or feature |
|---|---|
| Secure a computer by granting users or groups specific user rights, or assign permissions to an object to enable users or groups to perform specific actions on that object. | Access Control |
| Set up an auditing policy to record user or system activity in specified event categories. | Auditing Security Events |
| Confirm the identity of any user trying to log on to a domain or access network resources. | Authentication |
| Grant access to assigned user roles related to job function. | Authorization Manager |
| Encrypt files and directories that are stored on a disk. | Encrypting File System |
| Verify and authenticate the validity of each party involved in an electronic transaction through the use of public key cryptography. | Public Key Infrastructure |
| Create, apply, and edit the security for your local computer, organizational unit, or domain. | Security Configuration Manager |
| Understand security. | Security overview |
| Identify software and control its ability to run on your local computer, organizational unit, domain, or site. | Software Restriction Policies |

For more information about deploying security, see "Part 2: Deploying Distributed Security Services" at the Microsoft Windows Resource Kits Web site.

For more information about managing security, see "Part One: Security" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for managing security, and the command-line tools you can use to perform them. For more information, click the corresponding link under **Command-line tool**.

| Task | Command-line tool |
|---|---|
| Display or modify discretionary access control list (DACL) files. | Cacls |
| Request certificates from a certification authority (CA). | Certreq |
| Process and display certification authority (CA) configuration information, configure Certificate Services, back up and restore CA components, and verify certificates, key pairs, and certificate chains. | Certutil |

| | |
|---|---|
| Create, list or delete stored user names and passwords or credentials. | Cmdkey |
| Enable a user to run specific tools and programs with different permissions than the user's current logon provides. | Runas |
| Configure and analyze system security by comparing your current configuration to at least one template. | Secedit |
| Recover access to a file that previously was denied by making the administrator owner of the file. | Takeown |
| Display domain or computer name, user name, group names, logon identifier and privileges for the user currently logged on. | Whoami |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Managing servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing servers**

- Using the Windows interface
- Using command-line tools

**Using the Windows interface**

The following table lists some of the most common tasks for managing your server, and the tools you can use to perform them. For more information, click the corresponding link under **Tool or Feature**.

| Task | Tool or feature |
|---|---|
| Add and manage roles for your server. | Configuring Roles for Your Server |
| Install and upgrade the operating system. | Installing and Upgrading the Operating System |

For more information about deploying servers, see "Part 2: Server Deployments" at the Microsoft Windows Resource Kits Web site.

For more information about managing servers, see the "Server Management Guide" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for managing your server, and the command-line tools you can use to perform them. For more information, click the corresponding link under **Command-line tool**.

| Task | Command-line tool |
|---|---|
| Connect to, configure and manage server clusters. | Cluadmin |
| Create a new cluster or administer an existing cluster. | Cluster |
| Create and manage backup copies of the Internet Information Services (IIS) configuration (metabase and schema) of a remote or local computer. | Iisback.vbs: IIS backup management script |
| Import or export all or selected elements of an Internet Information Services (IIS) metabase on a local or remote computer. | Iiscnfg.vbs: IIS configuration script |
| Create, delete, list, start, stop, or pause FTP sites on local or remote servers running Windows Server 2003 with Internet Information Services (IIS) 6.0. | Iisftp.vbs: IIS FTP site management script |
| Create, delete, or display virtual directories on FTP sites of servers running Windows Server 2003 with Internet Information Services (IIS) 6.0. | IISFtpdr.vbs: IIS FTP directory script |
| Create or delete virtual directories of Web sites on servers running Windows Server 2003 with Internet Information Services (IIS) 6.0. | Iisvdir.vbs: IIS virtual directory script |

| | |
|---|---|
| Create, delete, list, start, stop, or pause Web sites on servers running Windows Server 2003 with Internet Information Services (IIS) 6.0. | Iisweb.vbs: IIS Web site management script |
| Start or stop the File Server for Macintosh service or the Print Server for Macintosh service. | File and print server for Macintosh |
| Create, modify, or delete directory objects on computers running Windows Server 2003 operating systems or Windows XP Professional; extend the schema, export Active Directory user and group information to other applications or services, or populate Active Directory with data from other directory services. | LDIFDE |
| Manage a server running Remote Installation Services. | Risetup |
| Communicate with the Service Controller and installed services; retrieve or set control information about services. | SC |
| Display detailed configuration information about a computer and its operating system, including operating system configuration, security information, product ID, and hardware properties, such as RAM, disk space, and network cards. | Systeminfo |
| Install or upgrade to a product in the Windows Server 2003 family. | Winnt32 |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Monitoring and controlling services on computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring and controlling services on computers**

- Using the Windows interface
- Using command-line tools

**Using the Windows interface**

The following table lists some of the most common tasks for monitoring and controlling services, and the tools you can use to perform them. For more information, click the corresponding link under "Tool or feature".

| Task | Tool or feature |
|---|---|
| Collect performance data automatically from local or remote computers, and package data so that it can be viewed by System Monitor or exported to a spreadsheet or database for further analysis and report generation. | Performance Logs and Alerts overview |
| Track how often users shut down or restart their computers. | Shutdown Event Tracker |
| Display system configuration information for local and remote computers, including hardware configurations, computer components, software and drivers. | System Information |
| Collect and view extensive data about the usage of hardware resources and the activity of system services on your computer or other computers on the network. | System Monitor overview |
| View and change system properties on a local or remote computer. | System Properties |
| Monitor key indicators of your computer's performance. | Task Manager overview |
| Synchronize time on the computers in your network. | Windows Time Service |

For more information about monitoring and controlling services on computers, see "Part Four: Operating System Services" at the Microsoft Windows Resource Kits Web site.

**Using command-line tools**

The following table lists some of the most common tasks for monitoring and controlling services, and the command-line tools you can use to perform them. For more information, click the corresponding link under "Command-line tool."

| Task | Command-line tool |
|---|---|
| Configure the translation of events to traps, trap destinations, or both based on information in a configuration file. | Evntcmd |
| Register new Performance counter names and Explain text for a service or device driver; saves or restore counter settings and explain text. | Lodctr |
| Manage and schedule performance counter and event trace log collections on local and remote computers. | Logman |
| Display system status, change system settings, or reconfigure ports or devices. | Mode |
| Display a comprehensive view of the hardware, system components, and software environment on the local computer. | Msinfo32 |

| | |
|---|---|
| Open a Performance console configured with the System Monitor Active-X control and the Performance Logs and Alerts Service snap-in in the Microsoft Management Console. | Perfmon |
| Extract performance counters from performance counter logs into other formats, such as text-TSV (for tab-delimited text), text-CSV (for comma-delimited text), binary-BIN, or SQL. | Relog |
| Communicate with the Service Controller and installed services; retrieve or set control information about services. | SC |
| Display detailed configuration information about a computer and its operating system, including operating system configuration, security information, product ID, and hardware properties, such as RAM, disk space, and network cards. | Systeminfo |
| End one or more tasks or processes. | Taskkill |
| Display a list of applications, services, and the Process ID (PID) currently running on either a local or a remote computer. | Tasklist |
| Process event trace logs or real-time data from instrumented event trace providers, and generate trace analysis reports and CSV (comma-delimited) files for the events generated. | Tracerpt |
| Display performance counter data. | Typeperf |
| Remove performance counter names and explain text for a service or device driver from the system registry. | Unlodctr |
| Diagnose problems occurring with Windows Time. | W32tm |

For more information about the command shell, see Command shell overview. For information about tools used to manage the command-prompt window, see Manage the command-prompt window. For an alphabetical list of all of the command-line tools, see Command-line reference A-Z.

# Command-Line Reference

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Command-line reference**
You can use Command-line reference to familiarize yourself with new and enhanced command-line tools, to configure your command prompt, to learn about the command shell, and to automate command-line tasks by using batch files or scripting tools such as CScript and WMIC.

- For information about functionality that has been changed or is no longer available in the Microsoft® Windows Server 2003 family, see New ways to do familiar command-line tasks.
- For a list of new command-line tools, see New command-line tools.
- For a complete list of command-line tools, see Command-line reference A-Z.
- For general background information, see Command shell overview.
- For help with configuring the command prompt, see Configure the command prompt.

# New ways to do familiar command-line tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The Windows Server 2003 family retains and enhances almost all MS-DOS functionality. The following tables explain:

- Changes to the functionality of MS-DOS commands.
- MS-DOS commands that are no longer available in the Windows Server 2003 family.
- MS-DOS commands that are no longer available on Windows Server 2003, Datacenter Editionor Windows Server 2003, Enterprise Edition.

**Changes toMS-DOScommands**

The following table lists changes to the functionality of MS-DOS commands.

| Command | Changed features |
|---|---|
| Chcp | Changes code pages for all Command Prompt windows. |
| Diskcomp | Command-line options **/1** and **/8** are not supported. |
| **diskperf** | No longer needed to act upon a server computer. You can still use **diskperf** for remote administration. |
| Diskcopy | Command-line option **/1** is not supported. |
| Doskey | Available for all character-based programs that accept buffered input. **Doskey** has been improved by a series of enhancements. |
| Format | Supports 1.44 MB floppy drive only. 20.8 MB floptical (magneto-optical) drive is no longer supported. Command-line options **/b**, **/s**, **/u**, **/1**, **/4**, and **/8** are no longer supported. |
| Label | The caret (^) and ampersand (&) symbols can be used in a volume label. |
| Mode | The **mode lpt***N*[**:**][**cols**=*c*][**lines**=*N*] and **mode lpt***N*[**:**][*c*][**,**[*N*]] syntaxes are no longer supported. |
| Print | Command-line options **/b**, **/c**, **/m**, **/p**, **/q**, **/s**, **/t**, and **/u** are not supported. |
| Prompt | New character combinations enable you to add ampersands ($a), parentheses ($c and $f), and spaces ($s) to your prompt. |
| recover | Recovers files only. |
| rmdir | New **/s** command-line option deletes directories containing files and subdirectories. |
| Sort | Does not require TEMP environment variable. File size is unlimited. |
| Xcopy | New **/g** command-line option creates decrypted destination files. |

**MS-DOScommands not supported in the Windows Server 2003 family**

The following MS-DOS commands are not supported in the Windows Server 2003 family.

| Command | New procedure or reason for obsolescence |
|---|---|
| **assign** | Not supported. |
| **backup** | Not supported. |

| | |
|---|---|
| **choice** | Not supported. |
| **ctty** | Not supported. |
| **dblspace** | Not supported. |
| **deltree** | The **rmdir /s** command deletes directories containing files and subdirectories. For more information about the **rmdir** command, see Rmdir. |
| **dosshell** | Function is now inherent. |
| **drvspace** | Not supported. |
| **emm386** | Function is now inherent. |
| **fasthelp** | Not supported. Use the **help** command at the command prompt. For more information about the **help** command, see Help. For a complete list of command-line tools, see Command-line reference A-Z. |
| **fdisk** | Not supported. Use the **diskpart** command, which enables you to manage disks, partitions, or volumes. For more information about the **diskpart** command, see DiskPart. |
| **include** | Multiple configurations of the MS-DOS subsystem are not supported. |
| **interlnk** | Not supported. |
| **intersrv** | Not supported. |
| **join** | Increased partition size and an improved file system eliminate the need to join drives. |
| **memmaker** | MS-DOS subsystem memory usage is automatically optimized. |
| **menucolor** | Multiple configurations of the MS-DOS subsystem are not supported. |
| **menudefault** | Multiple configurations of the MS-DOS subsystem are not supported. |
| **menuitem** | Multiple configurations of the MS-DOS subsystem are not supported. |
| **mirror** | Not supported. |
| **msav** | Not supported. |
| **msbackup** | Use the **ntbackup** command for computers with tape drives, or the **xcopy** command for computers without tape drives. For more information about the **ntbackup** command, see Ntbackup. For more information about the **xcopy** command, see Xcopy. |
| **mscdex** | It is unnecessary to configure the MS-DOS subsystem to use CD-ROM drives. The Windows Server 2003 family provides access to CD-ROM drives for the MS-DOS subsystem. |
| **msd** | Use the System Information snap-in. To start System Information, click **Start**, click **Run**, and then type **msinfo32**. |
| **numlock** | Not supported. |
| **power** | Not supported. |
| **restore** | Not supported. |
| **scandisk** | Not supported. Use the **chkdsk** command, which enables you to create and display a status report for a disk based on the file system. Chkdsk also lists and corrects errors on the disk. For more information about the **chkdsk** command, see chkdsk. |
| **smartdrv** | Caching for the MS-DOS subsystem is automatic. |
| **submenu** | Multiple configurations of the MS-DOS subsystem are not supported. |

| | |
|---|---|
| **sys** | Windows Server 2003 family operating systems will not fit on a standard 1.2 MB or 1.44 MB floppy disk. |
| **undelete** | Not supported. |
| **unformat** | Not supported. |
| **vsafe** | Not supported. |

**MS-DOScommands no longer available on Windows Server 2003, Datacenter Edition or Windows Server 2003, Enterprise Edition**
The following 16-bit MS-DOS subsystem commands are not available on Windows Server 2003, Datacenter Editionor Windows Server 2003, Enterprise Edition.

| Command | Function |
|---|---|
| **append** | Enables programs to open data files in specified folders as if these files were in the current folder. |
| **debug** | Starts Debug.exe, a program that you can use to test and debug MS-DOS executable files. |
| **edit** | Starts MS-DOS Editor, which creates and changes ASCII text files. |
| **edlin** | Starts Edlin.exe, a line-oriented text editor with which you can create and change ASCII files. |
| **exe2bin** | Converts .exe (executable files) to binary format. |
| **expand** | Expands one or more compressed files. |
| **fastopen** | Windows XP and the Windows Server 2003 family of products do not use this command. It is included only to preserve compatibility with existing MS-DOS files, but it has no effect at the command line because the functionality is automatic. |
| **forcedos** | Starts the specified program in the MS-DOS subsystem. This command is necessary only for those MS-DOS programs that are not recognized as MS-DOS programs by the Windows Server 2003 family. |
| **graphics** | Loads a program into memory that allows you to print on a printer the displayed contents of the screen when you are using a color or graphics adapter. |
| **loadfix** | Ensures that a program is loaded above the first 64KB of conventional memory, and then runs the program. |
| **loadhigh (lh)** | Loads a program into the upper memory area (UMA). Loading a program into the UMA leaves more room in conventional memory for other programs. |
| **mem** | Displays information about allocated memory areas, free memory areas, and programs that are currently loaded into memory in the MS-DOS subsystem. |
| **nlsfunc** | Loads information specific to country/region. Windows XP and the Windows Server 2003 family of products do not use this command. It is included only to preserve compatibility with existing MS-DOS files, but it has no effect at the command line because the functionality is automatic. |
| **setver** | Sets the MS-DOS version number that the MS-DOS subsystem reports to a program. |
| **share** | Windows XP and the Windows Server 2003 family of products do not use this command. It is included only to preserve compatibility with existing MS-DOS files, but it has no effect at the command line because the functionality is automatic. |

# New command-line tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New command-line tools**

The following table lists the new Windows Server 2003 family command-line tools.

| Command | Features |
|---|---|
| Adprep | Prepares Windows 2000 domains and forests for an upgrade to Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. |
| Bootcfg | Configures, queries, or changes Boot.ini file settings. |
| choice | Prompts the user to make a choice in a batch program by displaying a prompt and pausing for the user to choose from a set of user-option keys. |
| Clip | Redirects command output from the command line to the Clipboard. |
| Cmdkey | Creates, lists, and deletes stored user names and passwords or credentials. |
| Defrag | Locates and consolidates fragmented boot files, data files, and folders on local volumes. |
| DiskPart | Manages disks, partitions, or volumes. |
| Driverquery | Queries for a list of drivers and driver properties. |
| Dsadd | Adds a computer, contact, group, organization unit, or user to a directory. |
| Dsget | Displays selected attributes of a computer, contact, group, organizational unit, server or user in a directory. |
| Dsmod | Modifies an existing user, computer, contact, group or organizational units in a directory. |
| Dsmove | Moves any object from its current location in the directory to a new location, as long as the move can be accommodated within a single domain controller, and renames an object without moving it in the directory tree. |
| Dsquery | Queries and finds a list of computers, groups, organizational units, servers, or users in the directory by using specified search criterion. |
| Dsrm | Deletes an object of a specific type or any general object from the directory. |
| Eventcreate | Enables an administrator to create a custom event in a specified event log. |
| Eventquery.vbs | Lists the events and event properties from one or more event logs. |
| Eventtriggers | Displays and configures event triggers on local or remote machines. |
| Expand | Expands one or more compressed files. |
| Forfiles | Selects files in a folder or tree for batch processing. |
| Freedisk | Checks for available disk space before continuing with an installation process. |
| Fsutil | Manages reparse points, managing sparse files, dismounting a volume, or extending a volume. |
| Getmac | Obtains the media access control (MAC) address and list of network protocols. |

| | |
|---|---|
| Gettype | Sets the system environment variable %ERRORLEVEL% to the value associated with the specified Windowsoperating system. |
| Gpresult | Displays Group Policy settings and Resultant Set of Policy (RSoP) for a user or a computer. |
| Helpctr | Starts Help and Support Center. |
| Iisapp.vbs: IIS application query script | Reports the process identifiers (PIDs) of currently running W3pwp.exe processes serving a particular application pool. |
| Iisback.vbs: IIS backup management script | Creates and manages backup copies of the Internet Information Services (IIS) configuration (metabase and schema) of a remote or local computer. |
| Iiscnfg.vbs: IIS configuration script | Imports and exports all or selected parts of the configuration of Internet Information Services (IIS) on a local or remote computer. |
| Iisext.vbs: IIS Web service extension script | Configures and manages Web service extensions, applications, and individual files on servers running Windows Server 2003 with Internet Information Services (IIS) 6.0. |
| Iisftp.vbs: IIS FTP site management script | Creates, deletes, and lists FTP sites on servers that are running Internet Information Services (IIS) 6.0. Also, starts, stops, pauses, and continues FTP sites. |
| IISFtpdr.vbs: IIS FTP directory script | Creates and deletes virtual directories of FTP sites on servers that are running Internet Information Services (IIS) 6.0. |
| Iisvdir.vbs: IIS virtual directory script | Creates and deletes virtual directories of Web sites on servers that are running Internet Information Services (IIS) 6.0. |
| Iisweb.vbs: IIS Web site management script | Creates, deletes, and lists Web sites on servers that are running Internet Information Services (IIS) 6.0. Also, starts, stops, pauses, and continues the Web sites. |
| Inuse | Replaces locked operating system files. |
| Logman | Manages and schedules performance counter and event trace log collections on local and remote systems. |
| Nlb | Replaces Wlbs.exe for managing and controlling network load balancing operations. |
| Nlbmgr | Configures and manages Network Load Balancing clusters and all cluster hosts from a single computer. |
| Openfiles | Queries, displays, or disconnects open files. |
| Pagefileconfig.vbs | Displays and configures the paging file Virtual Memory settings of a system. |
| Perfmon | Enables you to open a Performance console configured with settings files from Windows NT 4.0 version of Performance Monitor. |
| Prncnfg.vbs | Configures or displays configuration information about a printer. |
| Prndrvr.vbs | Adds, deletes, and lists printer drivers from local or remote print servers. |
| Prnjobs.vbs | Pauses, resumes, cancels, and lists print jobs. |
| Prnmngr.vbs | Adds, deletes, and lists printers or printer connections, in addition to setting and displaying the default printer. |
| Prnport.vbs | Creates, deletes, and lists standard TCP/IP printer ports, in addition to displaying and changing port configuration. |
| Prnqctl.vbs | Prints a test page, pauses or resumes a printer, and clears a printer queue. |
| Relog | Extracts performance counters from performance counter logs into other formats, such as text-TSV (for tab-delimited text), text-CSV (for comma-delimited text), binary-BIN, or SQL. |

| | |
|---|---|
| Rss | Enables Remote Storage, which is used for extending server disk space. |
| SC | Retrieves and sets information about services. Tests and debugs service programs. |
| Schtasks | Schedules commands and programs to run periodically or at a specific time. Adds and removes tasks from the schedule, starts and stops tasks on demand, and displays and changes scheduled tasks. |
| setx | Sets environment variables in the local or system environment, without requiring programming or scripting. |
| Shutdown | Shuts down or restarts a local or remote computer. |
| Systeminfo | Queries the system for basic system configuration information. |
| Takeown | Allows an administrator to recover access to a file that previously was denied by making the administrator the owner of the file. |
| Taskkill | Ends one or more tasks or processes. |
| Tasklist | Displays a list of applications, services, and the Process ID (PID) currently running on either a local or a remote computer. |
| Timeout | Pauses the command processor for the specified number of seconds. |
| Tracerpt | Processes event trace logs or real-time data from instrumented event trace providers and allows you to generate trace analysis reports and comma-separated value files (.csv) for the events generated. |
| Tsecimp | Imports assignment information from an Extensible Markup Language (XML) file into the TAPI server security file (Tsec.ini). |
| Typeperf | Writes performance counter data to the command window or to a supported log file format. |
| waitfor | Synchronizes multiple computers across a network by using signals. |
| Where | Locates and displays all files that match the given parameter. |
| Whoami | Returns domain name, computer name, user name, group names, logon identifier, and privileges for the user who is currently logged on. |
| Windows Management Instrumentation Command-line (WMIC) tool | Eases the use of Windows Management Instrumentation (WMI) and systems managed through WMI. |

# Command-line reference A-Z

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Command-line reference A-Z**
To find information about a command, on the A-Z menu, click the letter that the command starts with, and then click the command name.

In addition to the tools installed with the Windows Server 2003 family there are over 40 support tools that you can install directly from the product CD. You can use these tools to diagnose and resolve computer problems. For more information about these support tools and how to install them, see Install Windows Support Tools.

For more information about changes to the functionality of MS-DOS commands, new command-line tools, command shell functionality, configuring the command prompt, and automating command-line tasks, see Command-line reference.

You must run command-line tools at the prompt of the Cmd.exe command interpreter. To open Command Prompt, click **Start**, click **Run**, type **cmd**, and then click **OK**. To view help at the command line, at the command prompt, type:

*CommandName* **/?**

As a security best practice, consider using an account without administrative credentials to perform routine, nonadministrative tasks. To perform only specific tasks with an account with administrative credentials, you can use the Runas command.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**A**
Adprep

Arp

Assoc

At

Atmadm

Attrib

**B**
Using batch files

Bootcfg

Break

**C**
Cacls

Call

Certreq

Certutil

Terminal Services change commands

Chcp

chdir

chkdsk

Chkntfs

Choice

Cipher

Clip

Cls

Cluadmin

Cluster

Cmd

Cmdkey

Cmstp

Color

Command shell overview

Comp

Compact

Convert

Copy

Cprofile

CScript

CSVDE

**D**

Date

Dcgpofix

Defrag

Del (erase)

Dfscmd

Dir

Directory service command-line tools

Diskcomp

Diskcopy

DiskPart

Doskey

Driverquery

Dsadd

Dsget

Dsmod

Dsmove

Dsquery

Dsrm

**E**

Echo

Endlocal

Eventcreate

Eventquery.vbs

Eventtriggers

Evntcmd

Exit

Expand

**F**

Fc

Using filters

Find

Findstr

Finger

Flattemp

For

Forfiles

Format

Freedisk

Fsutil

Ftp

Ftp subcommands

Ftype

**G**

Getmac

Gettype

Goto

Gpresult

Gpupdate

Graftabl

**H**

Help

Helpctr

Hostname

**I**

If

Iisapp.vbs: IIS application query script

Iisback.vbs: IIS backup management script

iiscnfg.vbs: IIS configuration script

iisext.vbs: IIS Web service extension script

iisftp.vbs: IIS FTP site management script

IISFtpdr.vbs: IIS FTP directory script

iisreset

iisvdir.vbs: IIS virtual directory script

iisweb.vbs: IIS Web site management script

Inuse

Ipconfig

Ipxroute

Irftp

**L**

Label

LDIFDE

Lodctr

Logman

Lpq

Lpr

**M**

Macfile

Md

Mkdir

Mmc

Mode

More

Mountvol

Move

Using the IP multicasting utilities

MS-DOSsubsystem commands

Msiexec (command-line options)

Msinfo32

**N**

Nbtstat

Net services overview

Net services commands

Netsh overview

Netsh commands for AAAA

Netsh commands for DHCP

Netsh diagnostic (diag) commands

Netsh commands for Interface IP

Netsh commands for Interface IPv6

Netsh commands for Interface Portproxy

Netsh commands for Internet Protocol security

Netsh commands for Network Bridge

Netsh commands for remote access (ras)

Netsh commands for Routing

Netsh commands for RPC

Netsh commands for WINS

Netstat

Nlb

Nlbmgr

Nslookup

Nslookup subcommands

Ntbackup

Ntcmdprompt

Ntdsutil

Ntsd

**O**

Openfiles

**P**

Pagefileconfig.vbs

Path

Pathping

Pause

Pbadmin

Pentnt

Perfmon

Ping

Popd

Powercfg

Print

Prncnfg.vbs

Prndrvr.vbs

Prnjobs.vbs

Prnmngr.vbs

Prnport.vbs

Prnqctl.vbs

Prompt

Pushd

**Q**
Terminal Services query commands

**R**
Rasdial

Rcp

Rd

Recover

Using command redirection operators

Reg

Regsvr32

Relog

Rem

Rename (ren)

Replace

Reset session

Rexec

Risetup

Rmdir

Route

Rsh

Rsm

Rss

Runas

**S**
SC

Schtasks

Secedit

Services for Macintosh commands

Set

Setlocal

Setx

Shift

Shutdown

Sort

Start

Subst

Sysocmgr

Systeminfo

System File Checker (sfc)

**T**
Takeown

Tapicfg

Taskkill

Tasklist

Tcmsetup

TCP/IP utilities and services

Telnet Command-line Reference for Telnet Server

Terminal Services commands

Tftp

Time

Timeout

Title

Tracerpt

Tracert

Tree

Tsecimp

Type

Typeperf

**U**
Unlodctr

**V**
Ver

Verify

Vol

Vssadmin

**W**
W32tm

waitfor

Where

Whoami

Winnt

Winnt32

Winpop

Wlbs

Windows Management Instrumentation Command-line (WMIC) tool

WScript

**X**

Xcopy

# Command shell overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Command shell overview**

The command shell is a separate software program that provides direct communication between the user and the operating system. The non-graphical command shell user interface provides the environment in which you run character-based applications and utilities. The command shell executes programs and displays their output on the screen by using individual characters similar to the MS-DOS command interpreter Command.com. The Windows server operating system command shell uses the command interpreter Cmd.exe, which loads applications and directs the flow of information between applications, to translate user input into a form that the operating system understands.

You can use the command shell to create and edit batch files (also called scripts) to automate routine tasks. For example, you can use scripts to automate the management of user accounts or nightly backups. You can also use the command-line version of Windows Script Host, CScript, to run more sophisticated scripts in the command shell. For more information, see Run Windows Script Host. You can perform operations more efficiently by using batch files than you can by using the user interface. Batch files accept all commands that are available at the command line. For more information about batch files and scripting, see Using batch files.

You can customize the command prompt window for easier viewing and to increase control over how you run programs. For more information about customizing the command prompt window, see Configure the command prompt.

**Using command syntax**

Syntax appears in the order in which you must type a command and any parameters that follow it. The following example of the **xcopy** command illustrates a variety of syntax text formats:

**xcopy** *Source* [*Destination*] [**/w**] [**/p**] [**/c**] [**/v**] [**/q**] [**/f**] [**/l**] [**/g**] [**/d**[**:***MM-DD-YYYY*]] [**/u**] [**/i**] [**/s** [**/e**]] [**/t**] [**/k**] [**/r**] [**/h**] [{**/a** | **/m**}] [**/n**] [**/o**] [**/x**] [**/exclude:***File1*[**+**[*File2*]][**+**[*File3*]] [{**/y** | **/-y**}] [**/z**]

The following table explains how to interpret the different text formats.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Using multiple commands and conditional processing symbols**

You can run multiple commands from a single command line or script using conditional processing symbols. When you run multiple commands with conditional processing symbols, the commands to the right of the conditional processing symbol act based upon the results of the command to the left of the conditional processing symbol. For example, you might want to run a command only if the previous command fails. Or, you might want to run a command only if the previous command is successful.

You can use the special characters listed in the following table to pass multiple commands.

| Character | Syntax | Definition |
|---|---|---|
| **&** [**...**] | *Command1* **&***Command2* | Use to separate multiple commands on one command line. Cmd.exe runs the first command, and then the second command. |

| | | |
|---|---|---|
| **&&** [...] | *Command1* **&&***Command2* | Use to run the command following **&&** only if the command preceding the symbol is successful. Cmd.exe runs the first command, and then runs the second command only if the first command completed successfully. |
| **\|\|** [...] | *Command1* **\|\|** *Command2* | Use to run the command following **\|\|** only if the command preceding **\|\|** fails. Cmd.exe runs the first command, and then runs the second command only if the first command did not complete successfully (receives an error code greater than zero). |
| **( )** [...] | **(** *Command1* **&** *Command2***)** | Use to group or nest multiple commands. |
| **;** or **,** | *Command1 Parameter1***;***Parameter2* | Use to separate command parameters. |

**Notes**

- The ampersand (&), pipe (|), and parentheses ( ) are special characters that must be preceded by the escape character (^) or quotation marks when you pass them as arguments.
- If a command completes an operation successfully, it returns an exit code of zero (0) or no exit code. For more information about exit codes, see Using the Windows Deployment and Resource Kits.

**Nesting command shells**

You can nest command shells within Cmd.exe by opening a new instance of Cmd.exe at the command prompt. By default, each instance of Cmd.exe inherits the environment of its parent Cmd.exe application. By nesting instances of Cmd.exe, you can make changes to the local environment without affecting the parent application of Cmd.exe. This allows you to preserve the original environment of Cmd.exe and return to it after you terminate the nested command shell. The changes you make in the nested command shell are not saved.

To nest a command shell, at the command prompt, type:

**cmd**

A message similar to the following appears:

```
Microsoft (R) Windows Server 2003, Standard Edition (TM)(C) Copyright 1985-2002 Microsoft Corp.
```

To close the nested command shell, type **exit**.

You can localize changes even further in an instance of Cmd.exe (or in a script) by using the **setlocal** and **endlocal** commands. **Setlocal** creates a local scope and **endlocal** terminates the local scope. Any changes made within the **setlocal** and **endlocal** scope are discarded, thereby leaving the original environment unchanged. You can nest these two commands to a maximum of 32 levels. For more information about the **setlocal** and **endlocal** commands, see Setlocal and Endlocal.

**Using environment variables with Cmd.exe**

The Cmd.exe command-shell environment is defined by variables that determine the behavior of the command shell and the operating system. You can define the behavior of the command-shell environment or the entire operating system environment by using two types of environment variables, system and local. System environment variables define the behavior of the global operating system environment. Local environment variables define the behavior of the environment of the current instance of Cmd.exe.

System environment variables are preset in the operating system and available to all Windows server operating system processes. Only users with administrative credentials can change system variables. These variables are most commonly used in logon scripts.

Local environment variables are only available when the user for whom they were created is logged on to the computer. Local variables set in the **HKEY_CURRENT_USER**hive are valid only for the current user, but define the behavior of the global operating system environment.

The following list describes the various types of variables in descending order of precedence:

1. Built-in system variables
2. System variables found in the **HKEY_LOCAL_MACHINE** hive

3. Local variables found in the **HKEY_CURRENT_USER** hive

4. All environment variables and paths set in the Autoexec.bat file

5. All environment variables and paths set in a logon script (if present)

6. Variables used interactively in a script or batch file

In the command shell, each instance of Cmd.exe inherits the environment of its parent application. Therefore, you can change the variables in the new Cmd.exe environment without affecting the environment of the parent application.

The following table lists the system and local environment variables for the Windows server operating system.

| **Variable** | **Type** | **Description** |
|---|---|---|
| %ALLUSERSPROFILE% | Local | Returns the location of the All Users Profile. |
| %APPDATA% | Local | Returns the location where applications store data by default. |
| %CD% | Local | Returns the current directory string. |
| %CMDCMDLINE% | Local | Returns the exact command line used to start the current Cmd.exe. |
| %CMDEXTVERSION% | System | Returns the version number of the current Command Processor Extensions. |
| %COMPUTERNAME% | System | Returns the name of the computer. |
| %COMSPEC% | System | Returns the exact path to the command shell executable. |
| %DATE% | System | Returns the current date. Uses the same format as the **date /t** command. Generated by Cmd.exe. For more information about the **date** command, see Date. |
| %ERRORLEVEL% | System | Returns the error code of the most recently used command. A non zero value usually indicates an error. |
| %HOMEDRIVE% | System | Returns which local workstation drive letter is connected to the user's home directory. Set based on the value of the home directory. The user's home directory is specified in Local Users and Groups. |
| %HOMEPATH% | System | Returns the full path of the user's home directory. Set based on the value of the home directory. The user's home directory is specified in Local Users and Groups. |
| %HOMESHARE% | System | Returns the network path to the user's shared home directory. Set based on the value of the home directory. The user's home directory is specified in Local Users and Groups. |
| %LOGONSERVER% | Local | Returns the name of the domain controller that validated the current logon session. |
| %NUMBER_OF_PROCESSORS% | System | Specifies the number of processors installed on the computer. |
| %OS% | System | Returns the operating system name. Windows 2000 displays the operating system as Windows NT. |
| %PATH% | System | Specifies the search path for executable files. |
| %PATHEXT% | System | Returns a list of the file extensions that the operating system considers to be executable. |
| %PROCESSOR_ARCHITECTURE% | System | Returns the chip architecture of the processor. Values: x86 or IA64 (Itanium-based). |
| %PROCESSOR_IDENTFIER% | System | Returns a description of the processor. |
| %PROCESSOR_LEVEL% | System | Returns the model number of the processor installed on the computer. |

| | | |
|---|---|---|
| %PROCESSOR_REVISION% | System | Returns the revision number of the processor. |
| %PROMPT% | Local | Returns the command prompt settings for the current interpreter. Generated by Cmd.exe. |
| %RANDOM% | System | Returns a random decimal number between 0 and 32767. Generated by Cmd.exe. |
| %SYSTEMDRIVE% | System | Returns the drive containing the Windows Server operating system root directory (that is, the system root). |
| %SYSTEMROOT% | System | Returns the location of the Windows Server operating system root directory. |
| %TEMP% and %TMP% | System and User | Returns the default temporary directories that are used by applications available to users who are currently logged on. Some applications require TEMP and others require TMP. |
| %TIME% | System | Returns the current time. Uses the same format as the **time /t** command. Generated by Cmd.exe. For more information about the **time** command, see Time. |
| %USERDOMAIN% | Local | Returns the name of the domain that contains the user's account. |
| %USERNAME% | Local | Returns the name of the user who is currently logged on. |
| %USERPROFILE% | Local | Returns the location of the profile for the current user. |
| %WINDIR% | System | Returns the location of the operating system directory. |

**Setting environment variables**

Use the **set** command to create, change, delete, or display environment variables. The **set** command alters variables in the current shell environment only.

To view a variable, at a command prompt, type:

**set** *VariableName*

To add a variable, at a command prompt, type:

**set***VariableName***=***Value*

To delete a variable, at a command prompt, type:

**set** *VariableName* **=**

You can use most characters as variable values, including white space. If you use the special characters <, >, |, &, or ^, you must precede them with the escape character (^) or quotation marks. If you use quotation marks, they are included as part of the value because everything following the equal sign is taken as the value. Consider the following examples:

- To create the variable value **new&name**, type:

  **set varname=new^&name**

- To create the variable value **"new&name"**, type:

  **set varname="new&name"**

- If you type **set varname=new&name** at the command prompt, an error message similar to the following appears:

  ```
  "'name' is not recognized as an internal or external command, operable program or
  batch file."
  ```

Variable names are not case-sensitive. However, **set** displays the variable exactly as you typed it. You can combine uppercase and lowercase letters in your variable names to make your code more readable (for example, UserName).

**Notes**

- The maximum individual environment variable size is 8192bytes.
- The maximum total environment variable size for all variables, which includes variable names and the equal sign, is 65,536KB.

**Substituting environment variable values**

To enable the substitution of variable values at the command line or in scripts, enclose the variable name in percent signs (that is, **%***VariableName***%**). By using percent signs, you ensure that Cmd.exe references the variable values instead of making a literal comparison. After you define variable values for a variable name, enclose the variable name in percent signs. Cmd.exe searches for all instances of the variable name and replaces it with the defined variable value. For example, if you create a script that contains different values (for example, user names) and you want to define the USERNAME environment variable for each user with these values, you can write one script using the variable USERNAME enclosed in percent signs. When you run this script, Cmd.exe replaces %USERNAME% with the variable values, which eliminates the need to perform this task manually for each user. Variable substitution is not recursive. Cmd.exe checks variables once. For more information about variable substitution, see For and Call.

# Configure the command prompt

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the command prompt**

1. Open Command Prompt.
2. Click the upper-left corner of the Command Prompt window, and then click **Properties**.
3. Click the **Options** tab.
4. In **Command History**, type or select **999** in **Buffer Size**, and then type or select **5** in **Number of Buffers**.
5. In **Edit Options**, select the **Quick Edit Mode** and **Insert Mode** check boxes.
6. Click the **Layout** tab.
7. In **Screen Buffer Size**, type or select **2500** in **Height**.
8. Do any of the following optional tasks:
   - In **Screen Buffer Size**, increase **Width**.
   - In **Window Size**, increase **Height**.
   - In **Window Size**, increase **Width**.
   - Clear the **Let system position window** check box, and then, in **Window Position**, change the values in **Left** and **Top**.
9. In the **Apply Properties** dialog box, click **Save properties for future windows with same title**.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To open **Command Prompt Properties** from the keyboard, press ALT+SPACEBAR+P.
- By selecting the **Quick Edit Mode** check box, you enable copy and paste from the Command Prompt window. To copy, select the text in the Command Prompt window with your left mouse button, and then right-click. To paste, either at the command prompt or in a text file, right-click.
- By increasing the screen buffer size to 999, you enable scrolling through the Command Prompt window.
- By increasing the number of buffers to five, you increase the number of lines in the Command Prompt window to 5000.
- To enable or disable file and directory name completion on a computer or user logon session, run Regedit.exe and set the following **REG_DWORD** value:

  **HKEY_LOCAL_MACHINE\Software\Microsoft\Command Processor\CompletionChar\REG_DWORD**

To set the **REG_DWORD** value, use the hexadecimal value of a control character for a particular function (for example, **0×9** is TAB and **0×08** is BACKSPACE). User-specified settings take precedence over computer settings, and command-line options take precedence over registry settings. For more information about file and directory name completion, see Related Topics.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Cmd
Command-line reference A-Z
Command shell overview

# Configuration and Management Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuration and management tools**
This section covers:

- Control Panel
- Group Policy
- Resultant Set of Policy
- Computer Management
- Local Users and Groups
- The Registry
- Windows Management Instrumentation Command-line
- Windows Management Instrumentation Control
- Windows Management Instrumentation Tester
- Microsoft Management Console

# Control Panel

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Control Panel**
Control Panel provides a set of special-purpose administrative tools that you can use to configure the Windows, applications, and services environment. Control Panel includes default items that you can use for common tasks (for example, Display and Add Hardware). User-installed applications and services can also insert icons into Control Panel.

There are two view options for Control Panel: Category View and Classic View. Category View displays Control Panel icons according to the type of task that the user wants to perform. Classic View displays Control Panel icons in a view that is familiar to users of previous versions of Windows.

For more information about setting view options for Control Panel, see:

- To set Control Panel to Classic View
- To set Control Panel to Category View

For more information about Control Panel icons, see:

- Windows interface administrative tool reference A-Z: Accessibility Options
- Windows interface administrative tool reference A-Z: Add Hardware
- Windows interface administrative tool reference A-Z: Add or Remove Programs
- Windows interface administrative tool reference A-Z
- Windows interface administrative tool reference A-Z: Date and Time
- Windows interface administrative tool reference A-Z: Display
- Windows interface administrative tool reference A-Z: Folder Options
- Windows interface administrative tool reference A-Z: Fonts
- Windows interface administrative tool reference A-Z: Game Controllers
- Windows interface administrative tool reference A-Z: Internet Options
- Windows interface administrative tool reference A-Z: Keyboard
- Windows interface administrative tool reference A-Z: Licensing (Control Panel)
- Windows interface administrative tool reference A-Z: Mouse
- Windows interface administrative tool reference A-Z: Network Connections
- Windows interface administrative tool reference A-Z: Phone and Modem Options
- Windows interface administrative tool reference A-Z: Power Options
- Printers and Faxes

- Specify regional and language options and user's name during Setup
- Windows interface administrative tool reference A-Z: Scanners and Cameras
- Windows interface administrative tool reference A-Z: Task Scheduler
- Windows interface administrative tool reference A-Z: Sounds and Audio Devices
- Windows interface administrative tool reference A-Z: Speech
- Stored User Names and Passwords
- Windows interface administrative tool reference A-Z: System

**Note**

- To add Control Panel to My Computer, on the **Tools** menu in Control Panel, click **Folder Options**, click the **View** tab, and then, under **Advanced settings**, select the **Show Control Panel in My Computer** check box.

# Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy**

You can use Group Policy to manage features included with the Microsoft® Windows Server 2003 family, such as Group Policy Software Installation, Administrative Templates, Folder Redirection, Remote Installation Services, Security Settings, Scripts (Startup/Shutdown and Logon/Logoff) and Internet Explorer Maintenance.

To use Group Policy without the Group Policy Management Console (GPMC), see Group Policy (pre-GPMC).

To use Group Policy with the Group Policy Management Console, see Group Policy Management.

Microsoft makes available new procedures, and new and updated tools online. For the latest information about Group Policy and tools that you can use to manage Group Policy, see Group Policy at the Microsoft Web site.

# Group Policy (pre-GPMC)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy**

You can use Group Policy to define policy settings that are applied to computers or users.

- Before you use Group Policy, see Checklist: Implementing Group Policy.
- To find features that have been moved in the Microsoft® Windows Server 2003 family, see New ways to do familiar Group Policy tasks (pre-GPMC).
- For tips about using Group Policy, see Group Policy Best Practices (pre-GPMC).
- For help with specific tasks, see Group Policy How To... (pre-GPMC).
- For general background information, see Group Policy Concepts.
- For problem-solving instructions, see Troubleshooting Group Policy.

# Checklist: Implementing Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Implementing Group Policy**

| | Step | References |
|---|---|---|
| | Review key concepts about Group Policy. | Group Policy overview; Group Policy settings overview |
| | Open Group Policy Object Editor. | Open Group Policy Object Editor; Ways to open Group Policy Object Editor |
| | Create a new Group Policy object or edit an existing Group Policy object. | Create a new Group Policy object; Edit a Group Policy object |
| | Configure Group Policy settings. | Software Settings; Windows Settings; Use Administrative Templates |
| | (Optional) Link a Group Policy object to a site, domain, or organizational unit. | Link a Group Policy object to a site, domain, or organizational unit |
| | (Optional) Set Group Policy Software Installation options. | Group Policy Software Installation overview; Set Group Policy Software Installation defaults; Use Group Policy to work with Applications; Set Application Properties |
| | (Optional) Configure Folder Redirection to redirect special folders. | Folder Redirection; Use Folder Redirection |

# New ways to do familiar Group Policy tasks (pre-GPMC)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

Windows NT 4.0 introduced the System Policy Editor, which you could use to create a system policy to control user actions and the work environment and to enforce system configuration settings for all computers that run Windows NT 4.0. Policy settings define the various components of the desktop environment, including the applications that are available to users, the applications that appear on users' desktops, and the options that are displayed on the **Start** menu.

Group Policy and its extensions, included with Windows XP Professional, Windows XP 64-bit Edition (Itanium), and the Windows Server 2003 family, replace many of the tools that you are familiar with in Windows NT 4.0.

The following table lists common tasks for configuring policy. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Set policies on users and computers in a site | Not applicable | Group Policy, accessed through Active Directory Sites and Services |
| Set policies on users and computers in a domain | System Policy Editor (Poledit.exe) | Group Policy, accessed through Active Directory Users and Computers |
| Set policies on users and computers in an organizational unit | Not applicable | Group Policy, accessed through Active Directory Users and Computers |
| Link a Group Policy object to a site, domain, or organizational unit | Not applicable | Group Policy, accessed through Active Directory Users and Computers |
| Use security groups to filter the scope of policy | Not applicable | Permissions on the **Security** tab in the Group Policy object's properties dialog box. |
| Manage software | For an administrator, Systems Management Server. For a user, Add/Remove programs in Control Panel. | Systems Management Server and the three Group Policy Software Installation and Maintenance tools:<br>• Group Policy Software Installation overview, an extension to the Group Policy Object Editor snap-in<br>• Windows Installer<br>• Add or Remove Programs in Control Panel |
| Create a safe user interface for editing the registry | Windows NT 4.0 Administrative Templates for System Policy Editor | Administrative Templates for Group Policy<br>For an explanation of how the role of .adm files has changed in the Windows Server 2003 family, see The role of Administrative Templates. |

| | | |
|---|---|---|
| Perform general administrative tasks | Administrative wizards, User Manager, and Server Manager | Microsoft Management Console (MMC) snap-ins: Active Directory Users and Computers, Active Directory Sites and Services, and Open Group Policy Object Editor and its extensions. |

The following table lists a common task for configuring Folder Redirection in Group Policy. The user interface for performing this task is different in this version of Windows from the way it was in Windows 2000. For more information about additional changes to Group Policy, see New Group Policy features (pre-GPMC).

| **If you want to** | **In Windows 2000 use** | **In this version of Windows use** |
|---|---|---|
| Redirect a special folder to a shared network directory with a Universal Naming Convention (UNC) path that is personalized to the user | A UNC path that incorporates *UserName*, for example, *ServerName**ShareName**UserName*\My Documents | The special folder's properties dialog box, as described in Redirect special folders to the root directory |

# Group Policy Best Practices (pre-GPMC)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- Best practices for Group Policy objects
- Best practices for Group Policy Software Installation
- Best practices for Folder Redirection

# Best practices for Group Policy objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for Group Policy objects**

**Do not process policy settings that are not configured.**

- If a Group Policy object contains only settings that are set to **Not Configured**, you can avoid processing these settings by disabling User Configuration or Computer Configuration. This expedites the startup and logon processes for those users and computers that are subject to the Group Policy object. For more information, see Disable the User Configuration settings in a Group Policy object, Disable the Computer Configuration settings in a Group Policy object, User Configuration and Computer Configuration.

- To prevent an entire Group Policy object from affecting a site, domain, or organizational unit, see Unlink a Group Policy object from a site, domain, or organizational unit and Disable a Group Policy object link. With these procedures, you can enable or re-link the Group Policy object.

- If you never want to use a certain Group Policy object again, see Delete a Group Policy object.

**Use Block Policy inheritance and No Override sparingly.**

- Routine use of these features makes it difficult to troubleshoot policy. If you must use them, see Block policy inheritance and Prevent a Group Policy object from being overridden.

**Do not use the same name for different Group Policy objects.**

- By using the same name for two different Group Policy objects, you do not cause Group Policy to function incorrectly, but you might cause confusion. For more information, see Group Policy objects.

### Caution

- If you type a name for a Group Policy object that is longer than 255 characters, the name is truncated without warning to 255 characters.

**Minimize the number of WMI Filters used with Group Policy objects.**

- The more WMI Filters that are applied to a Group Policy object, the longer it will take to process the object.

**Filter policy based on security group membership.**

- Users who do not have an access control entry (ACE) directing that a particular Group Policy object be applied to them can avoid the associated logon delay, because the Group Policy object is not processed for those users.

- Filtering can only be done by using membership in security groups. For more information, see Filter the scope of Group Policy according to security group membership.

- To see the ACEs, click the **Security** tab in the **Properties** dialog box for a Group Policy object.

**Override user-based Group Policy with computer-based Group Policy only when necessary.**

- Override user-based Group Policy with computer-based Group Policy only if you want the desktop configuration to be the same regardless of who logs on. The mechanism for doing this is called loopback, an advanced Group Policy setting that is useful in certain closely managed environments, such as laboratories, classrooms, public kiosks, and reception

areas. The **User Group Policy loopback processing mode** policy setting is located in Group Policy Object Editor. For more information, see Order of processing settings.

**Use Group Policy rather than System Policy.**

- Use System Policy only to manage computers that run an operating system that is earlier than Windows 2000 or if you need to manage desktops for multiple users on a stand-alone computer. For more information, see Migration Issues.

**Avoid assigning Group Policy objects across domains.**

- The processing of Group Policy objects slows the startup and logon processes if Group Policy is obtained from another domain.

**Do not set File System policy on a drive or directory, such as Sysvol, that is replicated by the NTFS file replication system (FRS).**

- Settings that are under File System in Group Policy Object Editor can cause excessive replication and can waste network bandwidth. For more information, see File System security settings.

**Do not link a Group Policy object to the same organizational unit more than once.**

- When more than one link for the same organizational unit is applied to a single object, the client-side Group Policy extension may interpret the links differently and produce unexpected resultant set of policy.

  For more information, see Link a Group Policy object to a site, domain, or organizational unit and Unlink a Group Policy object from a site, domain, or organizational unit.

# Best practices for Group Policy Software Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for Group Policy Software Installation**

**Understand the difference between Group Policy Software Installation and Systems Management Server.**

- For simple software installation and deployment scenarios, Group Policy provides Group Policy Software Installation. For more information, see Group Policy Software Installation. For software installation scenarios where scheduling, inventory, reporting, status, and support for installation across a wide area network (WAN) is required, use Systems Management Server (SMS).

**Specify application categories for your organization.**

- Using categories makes it easier for users to find an application in Add or Remove Programs in Control Panel. You can define application categories, such as Sales Applications, Accounting Applications, and so on. For more information, see Specify categories for applications to be managed.

**Verify that Windows Installer packages are correctly transformed before they are published or assigned.**

- Transforms are applied to packages at the time of assignment or publication. Transforms (.mst files) are customizations that are applied to Windows Installer packages. A transform is applied at the time of assignment or publication, not at the time of installation. In practical terms, this means that you should verify that the **Modifications** tab in the package **Properties** dialog box is set up the way that you want it before you click **OK**. If you neglect to do this, and you assign or publish a transformed package before you have completely configured it, you can remove the software and republish or reassign it, or you can upgrade the software with a completely transformed version. For more information, see Remove a managed application, Upgrade an application, and Add or remove modifications for an application package.

**Assign or publish just once per Group Policy object.**

- It is recommended that a Windows Installer package be used to assign an application or to publish an application no more than once in the same Group Policy object. For example, if you assign Microsoft Office to the computers that are affected by a Group Policy object, do not assign or publish it to the users who are affected by the Group Policy object. For more information, see Assign an application and Publish an application.

**Take advantage of authoring tools.**

- Developers who are familiar with the files, registry entries, and other requirements that are necessary for an application to work properly can author native Windows Installer packages by using tools that are available from various software vendors.

**Repackage existing software.**

- You can use commercially available tools to create Windows Installer packages for software that does not include natively authored .msi files. These tools work by comparing a computer's state before and after installation. For best results, install them on a computer that is free of other application software (clean installation).

**Assign or publish applications at a high level in the Active Directory hierarchy.**

- Because Group Policy settings apply by default to child Active Directory containers, it is efficient to assign or publish applications by linking a Group Policy object to a parent organizational unit or domain. Use security descriptors (ACEs) on the Group Policy object for finer control over who receives the software. For more information, see Filter the scope of Group Policy according to security group membership.

**Use required, rather than optional, upgrades.**

- For the correct upgrade procedures, both required and optional, see Upgrade an application.
- The **Upgrades** tab for a package in Group Policy Software Installation has a **Required upgrade for existing packages** check box. It is recommended that you select this check box. If two users use the application on one computer, and one user upgrades the application and the other does not, both versions of the application exist on the computer. Some applications do not support this configuration.
- Authenticated Users need the Read and Apply Group Policy ACE to be able to install applications from the software distribution point.
- Administrators need Full Control to manage software.

**Use Group Policy Software Installation properties for widely scoped control.**

- By using Group Policy Software Installation, you save administrative keystrokes when you assign or publish a large number of packages with similar properties in a single Group Policy object. For example, when all of the software is published and it all comes from the same software distribution point. For specific procedures, see Set Options for Group Policy Software Installation.

**Use Windows Installer package properties for fine control.**

- Use this method for assigning or publishing a single package. For specific procedures, see Use Group Policy to work with Applications.

# Best practices for Folder Redirection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for Folder Redirection**

**Accept the default settings.**

- In general, accept the default Folder Redirection settings.

**Allow the system to create the folders.**

- If you create the folders yourself, they will not have the correct permissions.
- If you must, or already have, created the folders, the information in the "Folder Redirection permissions" section in the Folder Redirection Help topic might not apply.

**Do not redirect My Documents to the home directory unless you have already deployed home directories in your organization.**

- Redirection to the home directory offers less security than standard folder redirection, and it is offered only for backward compatibility. For more information, see Folder Redirection and Redirect My Documents to the home directory.

**If you redirect My Documents to the home directory, and if your users log on to the domain through Terminal Server clients, do not specify a separate Terminal Services Home Directory.**

- The properties dialog box for a user in Active Directory has a **profile** tab, where the **Home Folder** can be set, and also a **Terminal Services profile** tab, where the **Terminal Services Home Directory** can be set. If these network folders are different, logging on to the terminal server causes the profile Home Folder to be copied to the Terminal Services Home Directory. Logging off the terminal server causes copying to occur in the opposite direction. This is probably not the intended behavior, and it might result in delay in the logon and logoff processes for users. For more information, see Redirect My Documents to the home directory.

**Enable client-side caching.**

- This is especially important for users with portable computers. For information about how to configure Folder Redirection and Offline Files to work together, see Folder Redirection.
- Always enable the **Synchronize all offline files before logging off** Group Policy setting in the Group Policy Object Editor under \Computer Configuration (or User Configuration)\Administrative Templates\Network\Offline Files. By enabling this setting, you ensure that offline files are fully synchronized and that all of the files in the user's redirected folder are available when the user is working offline. If this setting is not enabled, the system performs only a cursory synchronization, and as a result only recently used files are cached.

  **Note**

  - On client computers running Windows XP Professional and Windows Server 2003 operating systems, redirected folders are automatically cached to client-side caching. In Windows 2000, you must enable this feature.

**Use fully qualified universal naming convention (UNC) paths.**

- Although you can use paths like C:\FolderName, it is not recommended because the path might not exist on the target computer. Use fully qualified UNC paths.

  **Caution**

- UNC paths that are longer than 260 characters are truncated. If the path is truncated, redirection fails.

**Always place the My Pictures folder in the My Documents folder.**

- My Pictures can be redirected independently of My Documents, or it can follow My Documents (to remain its subfolder whenever My Documents is redirected), as it does by default. The default behavior is recommended unless you have a specific reason (such as server scalability) for separating My Pictures from My Documents.

**Consider the impact of policy removal.**

- Consider the behavior of your Folder Redirection policies at the time of policy removal, as described in the table of policy removal considerations in Folder Redirection. Changing the redirection option to **Not configured** does not redirect the folder to the user's local profile location. It will continue to be redirected to the previous location. If you want to return the folder to the local user profile location, use the **Redirect to the local user profile** policy setting, as described in Redirect special folders to the local profile location.

# Group Policy How To... (pre-GPMC)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Open Group Policy Object Editor
- Work with Group Policy Objects
- Work with Group Policy Object Properties
- Use Administrative Templates
- Use Startup, Shutdown, Logon, and Logoff Scripts
- Use Group Policy Software Installation
- Use Folder Redirection
- Configure Group Policy
- Delegate Control of Group Policy

# Open Group Policy Object Editor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Open Group Policy Object Editor**

- Open Group Policy from Active Directory Users and Computers
- Open Group Policy from Active Directory Sites and Services
- Open Group Policy as an MMC snap-in

# Open Group Policy from Active Directory Users and Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open Group Policy Object Editor from Active Directory Users and Computers**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

   **Where?**

   - Active Directory Users and Computers/*domain/organizational unit/child organizational unit*

3. Click **Properties**, and then click the **Group Policy** tab.
4. Do one of the following:

   - To edit an existing Group Policy object, click the Group Policy object in the list, and then click **Edit**.
   - To create a new Group Policy object, click **New**, type a name for the new Group Policy object, and then click **Edit**.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- You can open Group Policy Object Editor in several ways, depending on the action that you want to perform and the object to which you want to apply Group Policy.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ways to open Group Policy Object Editor

# Open Group Policy from Active Directory Sites and Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open Group Policy Object Editor from Active Directory Sites and Services**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click the site for which you want to set Group Policy.

   **Where?**

   - Active Directory Sites and Services/Sites/*site name*

3. Click **Properties**, and then click the **Group Policy** tab.
4. Do one of the following:

   - To edit an existing Group Policy object, click the Group Policy object in the list, and then click **Edit**.
   - To create a new Group Policy object, click **New**, type a name for the new Group Policy object, and then click **Edit**.

**Notes**

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- You can open Group Policy Object Editor in several ways, depending on the action that you want to perform and the object to which you want to apply Group Policy.

**See Also**
**Concepts**
Ways to open Group Policy Object Editor

# Open Group Policy as an MMC snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open Group Policy Object Editor as an MMC snap-in**

1. Open Microsoft Management Console.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. On the **Standalone** tab, click **Add**.
4. In the **Available Standalone Snap-ins** list, click **Group Policy Object Editor**, and then click **Add**.
5. In the **Select Group Policy Object** dialog box, click **Local Computer** to edit the local Group Policy object, or click **Browse** to find the Group Policy object that you want to edit.
6. Click **Finish**, click **Close**, and then click **OK**. Group Policy Object Editor opens the Group Policy object for you to edit.

## Notes

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- If you want to save a Group Policy Object Editor console and choose which Group Policy object opens in it from the command line, select the **Allow focus of the Group Policy Snap-in to be changed when launching from the command line** check box in the **Select Group Policy Object** dialog box.
- You can open Group Policy Object Editor in several ways, depending on the action that you want to perform and the object to which you want to apply Group Policy.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ways to open Group Policy Object Editor

# Work with Group Policy Objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with Group Policy objects**


- Edit a Group Policy object
- Edit the local Group Policy object
- Create a new Group Policy object
- Delete a Group Policy object
- Link a Group Policy object to a site, domain, or organizational unit
- Block policy inheritance
- Unlink a Group Policy object from a site, domain, or organizational unit
- Disable a Group Policy object link
- Prevent a Group Policy object from being overridden

# Edit a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit a Group Policy object**

1. Open the Group Policy object that you want to edit.
2. In the console tree, double-click the folders to view the policies in the details pane.
3. In the details pane, double-click a policy to open the **Properties** dialog box, and then change the policy settings.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- You must have Read and Write permissions for a Group Policy object to open it.
- Because changes to a Group Policy object take place immediately, you might want to disable the Group Policy object while you are editing it. For information about disabling a Group Policy object, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable a Group Policy object link
Edit the local Group Policy object

# Edit the local Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit the local Group Policy object**

1. Open Group Policy Object Editor.
2. In the console tree, double-click the folders to view the policies in the details pane.
3. In the details pane, double-click a policy to open the **Properties** dialog box, and then change the policy settings.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open the local Group Policy object, click **Start**, click **Run**, type **gpedit.msc**, and then click **OK**.
- To open Group Policy Object Editor, see Related Topics.
- Changes that you make to the local Group Policy object are saved automatically.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Local Group Policy
Group Policy objects

# Create a new Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new Group Policy object**

1. Do one of the following:

   - To create a Group Policy object that is linked to a domain or to an organizational unit, open Active Directory Users and Computers.

   - To create a Group Policy object that is linked to a site, open Active Directory Sites and Services.

2. In the console tree, right-click the site, domain, or organizational unit to which you want the newly created Group Policy object to be linked.

   The Group Policy object will be stored in the current domain, that is, the domain that contains the domain controller that is being used by Active Directory Users and Computers or Active Directory Sites and Services.

3. Click **Properties**, and then click the **Group Policy** tab.

4. Click **New**, type a name for the Group Policy object, and then click **Close**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- The newly created Group Policy object is linked by default to the site, domain, or organizational unit that you select when you create the Group Policy object,and its settings apply to that site, domain, or organizational unit. If you want to unlink the Group Policy object from that site, domain, or organizational unit, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable a Group Policy object link
Delete a Group Policy object

# Delete a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a Group Policy object**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain, or right-click any organizational unit in the domain.
3. Click **Properties**, and then click the **Group Policy** tab.
4. To find all the Group Policy objects that are stored in the domain, click **Add**.
5. In the **Add a Group Policy Object Link** dialog box, click the **All** tab, right-click the Group Policy object that you want to delete, and then click **Delete**.
6. Click **Yes**, click **Cancel**, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- When you delete a Group Policy object, any sites, domains, or organizational units to which it is linked are no longer affected by it. You might want to disable the Group Policy object instead. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Prevent a Group Policy object from being overridden
Disable a Group Policy object link

# Link a Group Policy object to a site, domain, or organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To link a Group Policy object to a site, domain, or organizational unit**

1. Do one of the following:

   - To link to a domain or an organizational unit, open Active Directory Users and Computers.
   - To link to a site, open Active Directory Sites and Services.

2. In the console tree, right-click the site, domain, or organizational unit to which you want the Group Policy object to be linked.
3. Click **Properties**, and then click the **Group Policy** tab.
4. To add the Group Policy object to the **Group Policy Object Links** list, click **Add**. This opens the **Add a Group Policy Object Link** dialog box.
5. Click the **All** tab, click the Group Policy object that you want, and then click **OK**.
6. In the properties dialog box for the site, domain, or organizational unit, click **OK**.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- You can link a Group Policy object to specify that its settings are applied to users and computers in the site, domain, or organizational unit and that its settings are applied to users and computers in Active Directory containers that inherit policy settings from the site, domain, or organizational unit.

Related Topics

Block policy inheritance

Unlink a Group Policy object from a site, domain, or organizational unit

Disable a Group Policy object link

Prevent a Group Policy object from being overridden

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Block policy inheritance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To block policy inheritance**

1. To block policy inheritance in a domain or organizational unit, open Active Directory Users and Computers.
2. In the console tree, right-click the domain, or organizational unit in which you want to block Group Policy inheritance, and then click **Properties**.
3. Click the **Group Policy** tab, select the **Block Policy inheritance** check box, and then click **OK**. This option is recommended, and it is selected by default.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- The **Block Policy inheritance** option blocks Group Policy objects that apply higher in the Active Directory hierarchy of domains, and organizational units. It does not block Group Policy objects if they have **No Override** enabled.
- The **Block Policy inheritance** option is set only on domains, and organizational units, but not on individual Group Policy objects.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Prevent a Group Policy object from being overridden
Disable a Group Policy object link

# Unlink a Group Policy object from a site, domain, or organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To unlink a Group Policy object from a site, domain, or organizational unit**

1. Do one of the following:

   - To unlink a Group Policy object from a domain or organizational unit, open Active Directory Users and Computers.
   - To unlink a Group Policy object from a site, open Active Directory Sites and Services.

2. In the console tree, right-click the site, domain, or organizational unit from which you want to unlink the Group Policy object.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Click the Group Policy object that you want to unlink, and then click **Delete**.
5. In the **Delete** dialog box, click **Remove the link from the list**, and then click **OK**.

### Important

- If you click **Remove the link and delete the Group Policy object permanently** in the **Delete** dialog box, all sites, domains, and organizational units to which the Group Policy object is linked will no longer have those Group Policy settings applied to them, and the Group Policy object will be deleted.

### Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- Unlinking prevents the Group Policy object from affecting that site, domain, or organizational unit.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Link a Group Policy object to a site, domain, or organizational unit
Delete a Group Policy object

# Disable a Group Policy object link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable a Group Policy object link**

1. Do one of the following:

    - To disable a Group Policy object link to a domain or organizational unit, open Active Directory Users and Computers.
    - To disable a Group Policy object link to a site, open Active Directory Sites and Services.

2. In the console tree, right-click the site, domain, or organizational unit to which the Group Policy object is linked, and then click **Properties**.
3. Click the **Group Policy** tab.
4. Right-click the Group Policy object link that you want to disable, click **Disabled**, and then click **Yes**. This switches the Disabled state to Active, and a check appears in the **Disabled** column.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- When you disable a Group Policy object link, the settings in the Group Policy object no longer apply to users or computers in the site, domain, or organizational unit to which the Group Policy object was linked; and they no longer apply to users and computers in child Active Directory containers that inherit those Group Policy settings.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Block policy inheritance
Disable the User Configuration settings in a Group Policy object
Disable the Computer Configuration settings in a Group Policy object

# Prevent a Group Policy object from being overridden

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prevent a Group Policy object from being overridden**

1. Do one of the following:

   - For a Group Policy object that is linked to a domain or organizational unit, open Active Directory Users and Computers.

   - For a Group Policy object that is linked to a site, open Active Directory Sites and Services.

2. In the console tree, right-click the site, domain, or organizational unit to which the Group Policy object is linked.

3. Click **Properties**, and then click the **Group Policy** tab.

4. In the **Group Policy Object Links** list, right-click the Group Policy object link that you want to enforce, click **No Override**, and then click **OK**. A check mark appears in the **No Override** column.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- If you set **No Override** on a Group Policy object link, the Group Policy settings are enforced on all users or computers in the site, domain, or organizational unit and on all users and computers in Active Directory containers that inherit Group Policy from it.

- Group Policy settings that are enforced by using **No Override** cannot be blocked.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Block policy inheritance
Link a Group Policy object to a site, domain, or organizational unit

# Work with Group Policy Object Properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with Group Policy object properties**

- Filter the scope of Group Policy according to security group membership
- Find the sites, domains, and organizational units to which a Group Policy object is linked
- Disable the User Configuration settings in a Group Policy object
- Disable the Computer Configuration settings in a Group Policy object

# Filter the scope of Group Policy according to security group membership

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To filter the scope of Group Policy according to security group membership**

1. Open the Group Policy object whose scope you want to filter.

2. In the console tree, right-click the icon or name of the Group Policy object, and then click **Properties**.

3. Click the **Security** tab, and then click the security group through which you want to filter this Group Policy object. If you want to change the list of security groups through which to filter this Group Policy object, use the **Add** and **Remove** buttons to add or remove security groups.

4. In the **Permissions** box for the selected security group, select or clear the appropriate check boxes to set permissions as shown in the following table, and then click **OK**.

| Your intention | Permissions | Result |
|---|---|---|
| Members of this security group have this Group Policy object applied to them. | Set **Apply Group Policy** to **Allow**.<br><br>Set **Read** to **Allow**. | This Group Policy object applies to members of this security group, unless they are members of at least one other security group that has **Apply Group Policy** set to **Deny**, or **Read** set to **Deny**, or both. |
| Members of this security group are exempt from this Group Policy object. | Set **Apply Group Policy** to **Deny**.<br><br>Set **Read** to **Deny**. | This Group Policy object never applies to members of this security group, regardless of the permissions these members have in other security groups. |
| Membership in this security group is irrelevant to whether the Group Policy object should be applied. | Set **Apply Group Policy** to neither **Allow** nor **Deny**.<br><br>Set **Read** to neither **Allow** nor **Deny**. | This Group Policy object applies to members of this security group if and only if they have both **Apply Group Policy** and **Read** set to **Allow** as members of at least one other security group. They also must not have **Apply Group Policy** or **Read** set to **Deny** as members of any other security group. |

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Object Editor, see Related Topics.

- Group Policy objects are applied only to sites, domains, and organizational units. Group Policy settings affect only the users and computers that they contain. In particular, Group Policy objects are not linked to security groups.

- The location of a security group in Active Directory has no relation to, and no effect on, filtering through that security group as described in this procedure.

- If a user or computer is not contained in a site, domain, or organizational unit that is subject to a Group Policy object, either directly through a link or indirectly through inheritance, there is no combination of permissions on any security group that can cause those Group Policy settings to affect that user or computer.

- Filtering at the Group Policy object level, as described in this procedure, causes the Group Policy object to be processed or not processed as a whole. The Group Policy Software Installation and Folder Redirection extensions use security groups to refine control beyond the Group Policy object level. Except for Folder Redirection and Group Policy Software Installation, security groups are not used to filter individual settings, or subsets, of a Group Policy object. For control over individual settings, edit or create a Group Policy object instead.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object
Edit a Group Policy object
Create a new Group Policy object

# Find the sites, domains, and organizational units to which a Group Policy object is linked

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find the sites, domains, and organizational units to which a Group Policy object is linked**

1. Open the Group Policy object.
2. In the console tree, right-click the icon or name of the Group Policy object, and then click **Properties**.
3. Click the **Links** tab, and then click **Find Now**. The sites, domains, and organizational units to which the Group Policy object is linked are listed in the **Sites, Domains or Organizational Units found** box.

## Notes

- To open Group Policy Object Editor, see Related Topics.
- If the Group Policy object is linked to more than one domain, you can limit your search for organizational units to one domain at a time by using the **Domain** list.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Disable the User Configuration settings in a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable the User Configuration settings in a Group Policy object**

1. Open the Group Policy object that you want to edit.
2. In the console tree, right-click the icon or name of the Group Policy object, and then click **Properties**.
3. On the **General** tab, select the **Disable User Configuration settings** check box, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- When you disable the User Configuration settings in a Group Policy object, they no longer affect any site, domain, or organizational unit to which this Group Policy object is linked.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Disable the Computer Configuration settings in a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable the Computer Configuration settings in a Group Policy object**

1. Open the Group Policy object that you want to edit.
2. In the console tree, right-click the icon or name of the Group Policy object, and then click **Properties**.
3. On the **General** tab, select the **Disable Computer Configuration settings** check box, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- When you disable the Computer Configuration settings in a Group Policy object, they no longer affect any site, domain, or organizational unit to which this Group Policy object is linked.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Use Administrative Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Administrative Templates**

- Add or remove an Administrative Template (.adm file)
- Use the view provided by Administrative Templates
- Filter the view provided by Administrative Templates

# Add or remove an Administrative Template (.adm file)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add or remove an Administrative Template (.adm file)**

1. Open the Group Policy object that you want to edit.
2. In the console tree, right-click **Administrative Templates**.

   **Where?**

   - *Group Policy object*/Computer Configuration (or User Configuration)/Administrative Templates

3. Click **Add/Remove Templates**.
4. Do one of the following:

   - To add a template, click **Add**. In the **Policy Templates** dialog box, click the template that you want to add, and then click **Open**.
   - To remove a template, in the **Current Policy Templates** list, click the template, and then click **Remove**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- Because the role of Administrative Templates in the Windows Server 2003 family is different from Windows NT 4.0, the use of Windows NT 4.0 .adm files on client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems is not recommended.

  If you do use the older .adm files to create namespaces in the Group Policy console, you can control whether the Windows NT 4.0 System Policy settings are visible in the details pane. To do this, in the console tree, click **Administrative Templates**. On the **View** menu, click **Filtering**, and then select or clear the **Only show policy settings that can be fully managed** check box. By selecting this check box, you cause System Policy settings to not be visible in the details pane. You can also use the Group Policy setting **Enforce Show Policies Only** to do this.

  **Where?**

  - *Group Policy object*/User Configuration/Administrative Templates/System/Group Policy

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services

Ways to open Group Policy Object Editor
Edit the local Group Policy object
The role of Administrative Templates
Windows NT 4.0 system policies

# Use the view provided by Administrative Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use the view provided by Administrative Templates**

1. Open Group Policy Object Editor.
2. In the console tree, click the folder under Administrative Templates that contains the policy settings that you want to set.

   **Where?**

   - *Group Policy object*/Computer Configuration (or User Configuration)/Administrative Templates/*policy folder*

3. At the bottom of the details pane, click the **Extended** tab.
4. In the **Setting** column, click the name or icon for a setting to read a description of the setting.
5. To change that setting from its default (not configured) state, double-click the name or icon for the setting.
6. On the **Setting** tab, click one of the buttons described in the following table.

| Button | Result |
|---|---|
| **Not Configured** | The registry is not modified. |
| **Enabled** | The registry reflects that the policy setting is selected. |
| **Disabled** | The registry reflects that the policy setting is not selected. |

7. Select any other available options that you want on the **Setting** tab, and then click **OK**.
8. To view and set other settings in the current folder, click **Previous Setting** or **Next Setting**.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- In Group Policy Object Editor, there are two Administrative Templates branches. Settings for **HKEY_LOCAL_MACHINE** are under Computer Configuration, and settings for **HKEY_CURRENT_USER** are under User Configuration.
- If computer policy conflicts with user policy, computer policy generally takes precedence. Application authors can disregard that convention, however, and if they do, Group Policy cannot enforce it.
- The **Extended** tab feature displays the text that explains the policy setting and indicates which versions of Windows are supported as clients for the setting. If you prefer a view of the policy setting without the explanatory text, click the **Standard** tab.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor

# Filter the view provided by Administrative Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To filter the view provided by Administrative Templates**

1. Open Group Policy Object Editor.
2. In the console tree, click **Administrative Templates**, or click a folder under Administrative Templates that contains the policy settings that you want to filter.

   **Where?**

   - *Group Policy object*/Computer Configuration (or User Configuration)/Administrative Templates/*policy folder*

3. On the **View** menu, click **Filtering**.
4. Select the **Filter by Requirements information** check box.
5. In the list box, select the check boxes for the categories that you want to make visible.
6. To hide settings that are not configured, select the **Only show configured policy settings** check box. If you do this, only enabled or disabled settings will be visible.
7. To hide Windows NT 4.0 system policy settings, select the **Only show policy settings that can be fully managed** check box. This option is recommended, and it is selected by default.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- This filtering procedure is useful if there are too many Administrative Templates settings at one time in Group Policy Object Editor. This filtering only selects which settings are visible in the console. It does not affect whether the settings apply to users or computers, and it is not the same as filtering the scope of Group Policy according to security group membership.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor

# Use Startup, Shutdown, Logon, and Logoff Scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use startup, shutdown, logon, and logoff scripts**

- Set up scripts on the domain controller
- Assign computer startup scripts
- Assign computer shutdown scripts
- Assign user logon scripts
- Assign user logoff scripts

# Set up scripts on the domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up scripts on the domain controller**

- Copy the script and dependent files to the Netlogon shared folder, or another shared folder, on the domain controller from which you want the script to run.

**Note**

- You must be a domain administrator to configure scripts on a domain controller.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Logon script replication
Startup, shutdown, logon, and logoff scripts

# Assign computer startup scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign computer startup scripts**

1. Open Group Policy Object Editor.
2. In the console tree, click **Scripts (Startup/Shutdown)**.

   **Where?**

   - *Group Policy object*/Computer Configuration/Windows Settings/Scripts (Startup/Shutdown)

3. In the details pane, double-click **Startup**.
4. In the **Startup Properties** dialog box, click **Add**.
5. In the **Add a Script** dialog box, do the following:

   - In **Script Name**, type the path to the script, or click **Browse** to search for the script file in the Netlogon shared folder on the domain controller.
   - In **Script Parameters**, type any parameters that you want, the same way as you would type them on the command line. For example, if your script includes parameters called //logo (display banner) and //I (interactive mode), type **//logo //I**.

6. In the **Startup Properties** dialog box, specify the options that you want as follows:

   - **Startup Scripts for *Group Policy object*--**Lists all the scripts that are currently assigned to the selected Group Policy object. If you assign multiple scripts, the scripts are processed in the order that you specify. To move a script up in the list, click it, and then click **Up**. To move a script down in the list, click it, and then click **Down**.
   - **Add--**Opens the **Add a Script** dialog box, where you can specify any additional scripts to use.
   - **Edit--**Opens the **Edit Script** dialog box, where you can modify script information, such as name and parameters.
   - **Remove--**Removes the selected script from the **Startup Scripts** list.
   - **Show Files--**Displays the script files that are stored in the selected Group Policy object.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- Startup scripts are run as Local System, and they have the full rights that are associated with being able to run as Local System.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Assign computer shutdown scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign computer shutdown scripts**

1. Open Group Policy Object Editor.
2. In the console tree, click **Scripts (Startup/Shutdown)**.

   **Where?**

   - *Group Policy object*/Computer Configuration/Windows Settings/Scripts (Startup/Shutdown)

3. In the details pane, double-click **Shutdown**.
4. In the **Shutdown Properties** dialog box, click **Add**.
5. In the **Add a Script** dialog box, do the following:

   - In **Script Name**, type the path to the script, or click **Browse** to search for the script file in the Netlogon shared folder on the domain controller.
   - In **Script Parameters**, type any parameters that you want, the same way as you would type them on the command line. For example, if your script includes parameters called //logo (display banner) and //I (interactive mode), type **//logo //I**.

6. In the **Shutdown Properties** dialog box, specify the options that you want as follows:

   - **Shutdown Scripts for *Group Policy object*--**Lists all the scripts that are currently assigned to the selected Group Policy object. If you assign multiple scripts, the scripts are processed in the order that you specify. To move a script up in the list, click it, and then click **Up**. To move a script down in the list, click it, and then click **Down**.
   - **Add--**Opens the **Add a Script** dialog box, where you can specify any additional scripts to use.
   - **Edit--**Opens the **Edit Script** dialog box, where you can modify script information, such as name and parameters.
   - **Remove--**Removes the selected script from the **Shutdown Scripts** list.
   - **Show Files--**Displays the script files that are stored in the selected Group Policy object.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- Shutdown scripts are run as Local System, and they have the full rights that are associated with being able to run as Local System.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Assign user logon scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign user logon scripts**

1. Open Group Policy Object Editor.
2. In the console tree, click **Scripts (Logon/Logoff)**.

   **Where?**

   - *Group Policy object*/User Configuration/Windows Settings/Scripts (Logon/Logoff)

3. In the details pane, double-click **Logon**.
4. In the **Logon Properties** dialog box, click **Add**.
5. In the **Add a Script** dialog box, do the following:

   - In **Script Name**, type the path to the script, or click **Browse** to search for the script file in the Netlogon shared folder on the domain controller.
   - In **Script Parameters**, type any parameters that you want, the same way as you would type them on the command line. For example, if your script includes parameters called //logo (display banner) and //I (interactive mode), type **//logo //I**.

6. In the **Logon Properties** dialog box, specify the options that you want as follows:

   - **Logon Scripts for *Group Policy object*--**Lists all the scripts that are currently assigned to the selected Group Policy object. If you assign multiple scripts, the scripts are processed in the order that you specify. To move a script up in the list, click it, and then click **Up**. To move a script down in the list, click it, and then click **Down**.
   - **Add--**Opens the **Add a Script** dialog box, where you can specify any additional scripts to use.
   - **Edit--**Opens the **Edit Script** dialog box, where you can modify script information, such as name and parameters.
   - **Remove--**Removes the selected script from the **Logon Scripts** list.
   - **Show Files--**Displays the script files that are stored in the selected Group Policy object.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Assign user logoff scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign user logoff scripts**

1. Open Group Policy Object Editor.
2. In the console tree, click **Scripts (Logon/Logoff)**.

   **Where?**

   - *Group Policy object*/User Configuration/Windows Settings/Scripts (Logon/Logoff)

3. In the details pane, double-click **Logoff**.
4. In the **Logoff Properties** dialog box, click **Add**.
5. In the **Add a Script** dialog box, do the following:

   - In **Script Name**, type the path to the script, or click **Browse** to search for the script file in the Netlogon shared folder on the domain controller.
   - In **Script Parameters**, type any parameters that you want, the same way as you would type them on the command line. For example, if your script includes parameters called //logo (display banner) and //I (interactive mode), type **//logo //I**.

6. In the **Logoff Properties** dialog box, specify the options the you want as follows:

   - **Logoff Scripts for *Group Policy object*--**Lists all the scripts that are currently assigned to the selected Group Policy object. If you assign multiple scripts, the scripts are processed in the order that you specify. To move a script up in the list, click it, and then click **Up**. To move a script down in the list, click it, and then click **Down**.
   - **Add--**Opens the **Add a Script** dialog box, where you can specify any additional scripts to use.
   - **Edit--**Opens the **Edit Script** dialog box, where you can modify script information, such as name and parameters.
   - **Remove--**Removes the selected script from the **Logoff Scripts** list.
   - **Show Files--**Displays the script files that are stored in the selected Group Policy object.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- Logoff scripts are run as User, not Administrator, and their rights are limited accordingly.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Use Group Policy Software Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Group Policy Software Installation**


- Open Group Policy Software Installation
- Set Options for Group Policy Software Installation
- Use Group Policy to work with Applications
- Set Application Properties
- Add or remove modifications for an application package

# Open Group Policy Software Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open Group Policy Software Installation**

1. Open Group Policy Object Editor.
2. To assign software applications to computers, in the console tree double-click **Computer Configuration**. To assign or publish software applications to users, in the console tree double-click **User Configuration**.
3. Double-click **Software Settings**, and then click **Software installation**.

   **Where?**

   - *Group Policy object*/Computer Configuration (or User Configuration)/Software Settings/Software installation

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Set Options for Group Policy Software Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set options for Group Policy Software Installation**

- Set Group Policy Software Installation defaults
- Specify automatic installation options based on file name extension
- Specify categories for applications to be managed

# Set Group Policy Software Installation defaults

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set Group Policy Software Installation defaults**

1. Open Group Policy Software Installation.

2. In the console tree, right-click **Software installation**.

   **Where?**

   - *Group Policy object*/Computer Configuration (or User Configuration)/Software Settings/Software installation

3. Click **Properties**, and on the **General** tab specify the following options:

   - In **Default package location**, specify the default software distribution point.

   - In **New packages**, specify how to add new packages to user settings. By default, packages can either be published or assigned. (To computers, they can only be assigned). If you want to select these options for each package, click **Display the Deploy Software dialog box**. For even more control on a per-package basis, click **Advanced**.

   - In **Installation user interface options**, click **Basic** or **Maximum**, depending on how apparent you want the installation process to be to the users.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Software Installation, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation

# Specify automatic installation options based on file name extension

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify automatic installation options based on file name extension**

1. Open Group Policy Software Installation.

2. In the console tree, right-click **Software installation**.

   **Where?**

   - *Group Policy object*/Computer Configuration (or User Configuration)/Software Settings/Software installation

3. Click **Properties**, and then click the **File Extensions** tab.

4. On the **File Extensions** tab, specify the following options:

   - **Select file extension--**Displays a list of all available file name extensions for the managed software. When an extension is selected, all applications in Active Directory that are associated with the selected file name extension are displayed in the **Application precedence** list box. For example, for an .htm extension, many applications can be associated, such as Microsoft Word, Microsoft Internet Explorer, Netscape Communicator, and so on.

   - **Application precedence--**Lists all the applications that are associated with the selected file name extension. By default, the application with the highest precedence is at the top of the list.

   - **Up--**Moves the selected application higher in the precedence list. The application at the top of the list is automatically installed if a document with the selected file name extension is invoked before the application has been installed. The **Up** button is unavailable when the selected application is at the top of the list.

   - **Down--**Moves the selected application lower in the priority list. The **Down** button is unavailable when the selected application is at the bottom of the list.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Software Installation, see Related Topics.

- When you upgrade an application, turn off the autoinstall option for the older version.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy Software Installation
Open Group Policy Object Editor

Group Policy Software Installation overview
Group Policy Software Installation

# Specify categories for applications to be managed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify categories for applications to be managed**

1. Open Group Policy Software Installation.
2. In the console tree, right-click **Software installation**.

   **Where?**

   - *Group Policy object*\Computer Configuration (or User Configuration)\Software Settings\Software installation

3. Click **Properties**, and then click the **Categories** tab.
4. To create or change the list of categories under which programs appear in Add or Remove Programs in Control Panel, do any of the following:

   - To create a new category, click **Add**, and then type the new category name. If you do not have permission to add categories, this button is unavailable.
   - To modify an existing category, click the category that you want to modify, click **Modify**, and then change the text of the selected category. If you do not have permission to edit categories, this button is unavailable.
   - To remove a category from the **Categories** list, click the category name, and then click **Remove**. If you do not have permission to edit categories, this button is unavailable.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Software Installation, see Related Topics.
- The changes that you make to the application categories apply throughout the domain in which this Group Policy object is stored.
- To perform this procedure, you must be a domain administrator or have equivalent rights.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation

# Use Group Policy to work with Applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with applications**

- Assign an application
- Publish an application
- Upgrade an application
- Remove a managed application

# Assign an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign an application**

1. Open Group Policy Software Installation.
2. In the console tree, right-click **Software installation**, point to **New**, and then click **Package**.

   **Where?**

   - *Group Policy object*\Computer Configuration (or User Configuration)\Software Settings\Software installation

3. In the **Open** dialog box, click the Windows Installer package, and then click **Open**.
4. In the **Deploy Software** dialog box, click **Assigned**, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Software Installation, see Related Topics.
- The **Open** dialog box shows the packages that are located at a software distribution point that you specify as the default. For instructions on how to set the default software distribution point, see Related Topics.
- If the Windows Installer package is located on a different shared network directory, navigate to the correct software distribution point in the **Open** dialog box, click the package, and then click **Open**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set Group Policy Software Installation defaults
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation

# Publish an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To publish an application**

1. Open Group Policy Software Installation.
2. Right-click in the details pane, point to **New**, and then click **Package**.
3. In the **Open** dialog box, click the Windows Installer package to be published, and then click **Open**.
4. In the **Deploy Software** dialog box, click **Published**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Software Installation, see Related Topics.
- Packages can be published only to users, not to computers.
- The **Open** dialog box shows the packages that are located at a software distribution point that you specify as the default. For instructions on how to set the default software distribution point, see Related Topics.
- If the Windows Installer package is located on a different shared network directory, navigate to the correct software distribution point in the **Open** dialog box, click the package, and then click **Open**.
- The application is available for users to install either by using Add or Remove Programs in Control Panel or by opening a file with a file name extension that you have associated with the application. For information about how to associate a file name extension with an application, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set Group Policy Software Installation defaults
Specify automatic installation options based on file name extension
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation

# Upgrade an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To upgrade an application**

1. Open Group Policy Software Installation.
2. In the details pane, right-click the new Windows Installer upgrade package (that is, not the package to be upgraded).
3. Click **Properties**, and then click the **Upgrades** tab.
4. Click **Add** to create or add to the list of packages that you want to upgrade with the new upgrade package.
5. In the **Add Upgrade Package** dialog box, under **Choose a package from**, do one of the following:
   - Click **Current Group Policy object (GPO)**.
   - Click **A specific GPO**, click **Browse**, and then click the Group Policy object that you want to upgrade.
6. Review the list of packages under **Package to upgrade**, which lists all of the other packages that are assigned or published within the selected Group Policy object.
7. Click the package that you want to upgrade, and then do one of the following:
   - To replace an application with a completely different application (perhaps an application from a different vendor), click **Uninstall the existing package, then install the upgrade package**.
   - To install a newer version of the same product while retaining the user's application preferences, document type associations, and so on, click **Package can upgrade over the existing package**.
8. (Optional) To make the upgrade mandatory, on the **Upgrades** tab, select the **Required upgrade for existing packages** check box.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- Before you can perform this procedure, you must first assign or publish the upgrade package.
- Depending on the Group Policy object, the list of packages under **Package to upgrade** might have zero or more entries.
- If this is an upgrade under Computer Configuration in the Group Policy console tree, the **Required upgrade for existing packages** check box is selected and unavailable because packages can only be assigned to computers, not published.
- To open Group Policy Software Installation, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation

# Remove a managed application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a managed application**

1. Open Group Policy Software Installation.
2. In the details pane, right-click the application that you want to remove, point to **All Tasks**, and then click **Remove**.
3. In the **Remove Software** dialog box, click one of the following removal methods:
   - To specify that the application is removed the next time a user logs on or restarts the computer, click **Immediately uninstall the software from users and computers**.
   - To specify that users can continue to use the application if they have already installed it, click **Allow users to continue to use the software, but prevent new installations**. If users have removed the application or if they have never installed it, they will not be able to install it.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Software Installation, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation

# Set Application Properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set application properties**

- View properties for managed applications
- Edit installation options for applications
- Specify application categories for Add or Remove Programs in Control Panel
- Set the autoinstall option for an application
- Set permissions for Group Policy Software Installation

# View properties for managed applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view properties for managed applications**

1. Open Group Policy Software Installation.
2. In the details pane, right-click the application for which you want to view properties, and then click **Properties**.

**Note**

- To open Group Policy Software Installation, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation

# Edit installation options for applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit installation options for applications**

1. To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

2. Open Group Policy Software Installation.

3. In the details pane, right-click the application for which you want to set installation options, and then click **Properties**.

4. In the properties dialog box for the application, click the **Deployment** tab.

5. On the **Deployment** tab, click one of the following options under **Deployment type**:

   - **Published--**Users in the selected site, domain, or organizational unit can install the application by using either Add or Remove Programs in Control Panel or file activation.

   - **Assigned--**Users in the selected site, domain, or organizational unit receive this application the next time that they log on (for assignment to users) or when the computer restarts (for assignment to computers).

6. Under **Deployment options**, select the check boxes for the following options that you want:

   - **Auto-install this application by file extension activation--**Select this option if you want to use the application precedence for the file name extension as it is determined on the **File Extensions** tab in the **Software installation Properties** dialog box. If the application is not already installed, selecting this option causes the application to be installed when the user opens a file with this file association, for example, by double-clicking the file.

   - **Uninstall this application when it falls out of the scope of management--**Select this option if you want the application to be removed when users log on, or when computers start up, if they move to a site, domain, or organizational unit for which the application is not deployed.

   - **Do not display this package in Add or Remove Programs in Control Panel--**This option removes the most obvious method for users to uninstall the application. This may be useful from time to time during the software life cycle.

   - **Install this application at logon--**This option fully installs the application using the package defaults. It does not merely advertise the application to the user by means of a shortcut. Note that the package defaults themselves might call for installation of all features or of only some features.

7. Under **Installation user interface options**, click one of the following options:

   - **Basic--**Only progress bars and errors are displayed.

   - **Maximum--**The entire user interface that is supported by the package is displayed.

8. If you want to set the following advanced options, click **Advanced**:

   - **Ignore language when deploying this package--**Select this check box if you are installing an application whose

language differs from the language of the operating system.

- **Remove existing installs of this product for users, if the product was not installed by Group Policy-based Software Installation--**This option might be useful, for example, if company policy does not allow users to install applications from their own compact discs.

- **Make this 32-bit X86 application available to Win64 machines--**Specifies whether the 32-bit application should be assigned or published to 64-bit computers.

- **Include OLE class and product information--**Specifies whether to deploy information about COM components with the package so that software on the client can install them from Active Directory as needed, in a manner similar to activation by file extension.

### Notes

- To open Group Policy Software Installation, see Related Topics.

- The **Advanced Deployment Options** dialog box, which appears when you click **Advanced**, provides advanced diagnostic information, which can be useful for troubleshooting. The following table lists and describes this information.

| Field | Data |
|---|---|
| Product code | A globally unique identifier (GUID) representing the product |
| Deployment Count | The number of times the package has been deployed in this Group Policy object |
| Script name | The full network path, including GUIDs, to the application assignment script (.aas file) |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation

# Specify application categories for Add or Remove Programs in Control Panel

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify application categories for Add or Remove Programs in Control Panel**

1. Open Group Policy Software Installation.
2. In the details pane, double-click the managed application for which you want to set categories.
3. In the properties dialog box for the application, click the **Categories** tab.
4. On the **Categories** tab, do any of the following:
   - To associate an application category that has been defined for your organization with this application, click a category in the **Available categories** list, and then click **Select**.
   - To remove a category that is associated with this application, click a category in the **Selected categories** list, and then click **Remove**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Software Installation, see Related Topics.
- The categories that you set in this dialog box generally pertain to published applications only, because assigned applications do not appear in Add or Remove Programs in Control Panel. The application appears in the selected categories in Add or Remove Programs, which the user can use to install the application.
- This procedure is not used to create or remove categories, but only to associate applications with existing categories. To manage the categories themselves, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation
Specify categories for applications to be managed

# Set the autoinstall option for an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the autoinstall option for an application**

1. Open Group Policy Software Installation.
2. In the console tree, click **Software installation**.

**Where?**

- *Group Policy object*\Computer Configuration (or User Configuration)\Software Settings\Software installation

1. In the details pane, right-click the application that you want to be installed automatically when the user tries to open a file with the associated extension.
2. Click **Properties**, and then click the **Deployment** tab.
3. Under **Deployment options**, select the **Auto-install this application by file extension activation** check box, and then click **OK**. This option is recommended, and it is selected by default.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Software Installation, see Related Topics.
- When you upgrade an application, turn off the autoinstall option for the older version.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy Software Installation
Open Group Policy Object Editor

# Set permissions for Group Policy Software Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions for Group Policy Software Installation**

1. Open the Group Policy object that you want to edit.
2. In the console tree, right-click the icon or name of the Group Policy object, and then click **Properties**.
3. Click the **Security** tab, and in the **Group or user names** box, click the security group for which you want to set permissions.
4. Do any of the following:
   - To change the set of security groups that have permissions on this Group Policy object, click **Add** or **Remove** to add or remove security groups.
   - If the security group represents administrators (who manage software for users and computers in the organization), in the **Permissions** box for the selected security group, select the **Allow** check box for Full Control, and then click **OK**.
   - If the security group represents users (who use software that is assigned or published by an administrator), in the **Permissions** box for the selected security group, select the **Allow** check box for both Apply Group Policy and Read, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- The permissions that you set in this procedure pertain only to the Group Policy object. In addition, the network must have a shared folder (a software distribution point) containing the packages to be installed. Administrators must have Full Control of the software distribution point, and users must have Read permission. For information about permissions and how to set them, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object
Access Control

# Add or remove modifications for an application package

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add or remove modifications for an application package**

1. Open Group Policy Software Installation.
2. In the console tree, right-click **Software installation**, point to **New**, and then click **Package**.

**Where?**

- *Group Policy object*\Computer Configuration (or User Configuration)\Software Settings\Software installation

1. In the **Open** dialog box, click the Windows Installer package, and then click **Open**.
2. In the **Deploy Software** dialog box, click **Advanced**, and then click **OK**.
3. In the properties dialog box for the package, click the **Modifications** tab.
4. Do any of the following:
   - To add modifications, click **Add**. In the **Open** dialog box, browse to the transform file (.mst), and then click **Open**.
   - To remove modifications, on the **Modifications** tab, click the modification that you want to remove, and then click **Remove**. Repeat this step until you have removed each modification that you do not want.
5. When you are sure that the modifications are configured the way you want them, click **OK**.

   **Important**

   - Do not click **OK** until you have finished configuring the modifications. When you click **OK**, the package is assigned or published immediately. If the modifications are not properly configured, you will have to uninstall the package or upgrade the package with a correctly configured version.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Software Installation, see Related Topics.
- The **Open** dialog box shows the packages that are located at a software distribution point that you specify as the default. For instructions on how to set the default software distribution point, see Related Topics.
- If the Windows Installer package is located on a different shared network directory, navigate to the correct software distribution point in the **Open** dialog box, click the package, and then click **Open**.
- Modifications (.mst files) are applied to Windows Installer packages (which have the .msi extension) in an order that you specify. This order must be determined before the application is assigned or published.
- You can add multiple modifications, which are also called transforms. The transforms are applied in the order that you

specify in the **Modifications** list, from top to bottom. To rearrange the list, click a transform in the list, and then click **Move Up** or **Move Down**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set Group Policy Software Installation defaults
Open Group Policy Software Installation
Open Group Policy Object Editor
Group Policy Software Installation overview
Group Policy Software Installation
Upgrade an application
Remove a managed application

# Use Folder Redirection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Folder Redirection**

- Redirect Special Folders to the Same Location for All Users
- Redirect Special Folders to Different Locations Based on Security Group Membership

# Redirect Special Folders to the Same Location for All Users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Redirect special folders to the same location for all users**

- Redirect My Documents to the home directory
- Redirect special folders to the root directory
- Redirect special folders to a specific path
- Redirect special folders to the local profile location

# Redirect My Documents to the home directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To redirect My Documents to the home directory**

1. Open a Group Policy object that is linked to the site, domain, or organizational unit that contains the users whose My Documents folders you want to redirect.

2. In the console tree, double-click **Folder Redirection** to display My Documents.

   **Where?**

   - *Group Policy object*\User Configuration\Windows Settings\Folder Redirection\My Documents

3. Right-click **My Documents**, and then click **properties**.

4. On the **Target** tab, in **Setting**, click **Basic - Redirect everyone's folder to the same location**.

5. Under **Target folder location**, click **Redirect to the user's home directory**, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Object Editor, see Related Topics.

- This option is intended only for organizations that have already deployed home directories and that want to maintain compatibility with their existing home directory environment.

- If you redirect My Documents to the home directory, domain administrators have full control over the user's My Documents folder. This is the case even if you select the **Grant the user exclusive rights to My Documents** check box on the **Settings** tab in the **My Documents properties** dialog box.

- Redirection to the home directory requires a client computer running Windows XP Professional, Windows XP 64-bit Edition (Itanium), or a Windows Server 2003 operating system. It does not work for Windows NT clients, Windows 2000 clients, or Windows XP Home Edition clients.

- In redirection of the home directory, security is not checked, and permissions are not changed. Folder Redirection assumes that the administrator has set directory security correctly. This relaxed security is why redirection to the home directory is only recommended if it is needed for backward compatibility.

- In home directory redirection, no ownership checks are made. Normally, folder redirection fails if a user is not the owner of the folder to which he or she is being redirected. Because redirection to the home directory is intended for use in a legacy environment, this ownership check is skipped.

- Users must have the home directory property set correctly on their user object in Active Directory. The client computer finds the path for the user's home directory from the user object in Active Directory at logon time. Users who are affected by Folder Redirection policy must have this path set correctly, or folder redirection will fail.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Redirect Special Folders to the Same Location for All Users
Redirect Special Folders to Different Locations Based on Security Group Membership
Windows Management Instrumentation (WMI) Filters
Folder Redirection

# Redirect special folders to the root directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To redirect special folders to the root directory**

1. Open a Group Policy object that is linked to the site, domain, or organizational unit that contains the users whose special folders you want to redirect.

2. In the console tree, double-click **Folder Redirection** to display the special folder that you want to redirect.

   **Where?**

   - *Group Policy object*\User Configuration\Windows Settings\Folder Redirection

3. Right-click the special folder that you want to redirect (for example, **Desktop** or **My Documents**), and then click **Properties**.

4. On the **Target** tab, in **Setting**, click **Basic - Redirect everyone's folder to the same location**.

5. Under **Target folder location**, click **Create a folder for each user under the root path**.

6. In **Root Path**, type a Universal Naming Convention (UNC) path (for example, **\\servername\sharename**), and then click **OK**.

7. In the properties dialog box for the special folder, click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Object Editor, see Related Topics.

- The user name and folder name are appended to the UNC path automatically.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Redirect Special Folders to the Same Location for All Users
Redirect Special Folders to Different Locations Based on Security Group Membership
Windows Management Instrumentation (WMI) Filters
Folder Redirection

# Redirect special folders to a specific path

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To redirect special folders to a specific path**

1. Open a Group Policy object that is linked to the site, domain, or organizational unit that contains the users whose special folders you want to redirect.

2. In the console tree, double-click **Folder Redirection** to display the special folder that you want to redirect.

   **Where?**

   - *Group Policy object name*\User Configuration\Windows Settings\Folder redirection

3. Right-click the special folder that you want to redirect (for example, **Desktop** or **My Documents**), and then click **Properties**.

4. On the **Target** tab, in **Setting**, click **Basic - Redirect everyone's folder to the same location**.

5. Under **Target folder location**, click **Redirect to the following location**.

6. In **Root Path**, type a UNC path (for example, *****ServerName******ShareName*), or type a locally valid path (for example, **C:***FolderName*).

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Object Editor, see Related Topics.

- In most cases, you want to redirect special folders to the root directory rather than to a specific location, as described in this procedure. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Redirect Special Folders to the Same Location for All Users
Redirect Special Folders to Different Locations Based on Security Group Membership
Windows Management Instrumentation (WMI) Filters
Folder Redirection
Redirect special folders to the root directory

# Redirect special folders to the local profile location

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To redirect special folders to the local profile location**

1. Open a Group Policy object that is linked to the site, domain, or organizational unit that contains the users whose special folders you want to redirect.

2. In the console tree, double-click **Folder Redirection** to display the special folder that you want to redirect.

   **Where?**

   - *Group Policy object*\User Configuration\Windows Settings\Folder Redirection

3. Right-click the special folder that you want to redirect (for example, **Desktop** or **My Documents**), and then click **Properties**.

4. On the **Target** tab, in **Setting**, click **Basic - Redirect everyone's folder to the same location**.

5. Under **Target folder location**, click **Redirect to the local user profile location**, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Object Editor, see Related Topics.

- This procedure is useful for returning redirected folders to their original, default location.

- If you change the setting to **No administrative policy specified** on the **Target** tab of the properties dialog box for a special folder, you do not cause the special folder to be redirected back to its default location.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Redirect Special Folders to the Same Location for All Users
Redirect Special Folders to Different Locations Based on Security Group Membership
Windows Management Instrumentation (WMI) Filters
Folder Redirection

# Redirect Special Folders to Different Locations Based on Security Group Membership

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Redirect special folders to different locations based on security group membership**

- Redirect My Documents to the home directory based on security group membership
- Redirect special folders to the root directory based on security group membership
- Redirect special folders to a specific path based on security group membership
- Redirect to the local profile location based on security group membership

# Redirect My Documents to the home directory based on security group membership

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To redirect My Documents to the home directory based on security group membership**

1. Open a Group Policy object that is linked to the site, domain, or organizational unit that contains the users whose My Documents folders you want to redirect.

2. In the console tree, double-click **Folder Redirection** to display My Documents.

   **Where?**

   - *Group Policy object*\User Configuration\Windows Settings\Folder Redirection\My Documents.

3. Right-click **My Documents**, and then click **properties**.

4. On the **Target** tab, in **Setting**, click **Advanced - Specify locations for various user groups**.

5. Click **Add**.

6. In the **Specify Group and Location** dialog box, under **Security Group Membership**, type the name of the security group, or click **Browse** to search for it.

7. Under **Target Folder Location**, click **Redirect to the user's home directory**, and then click **OK**.

8. In the **My Documents properties** dialog box, click **OK**.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Object Editor, see Related Topics.

- This option is intended only for organizations that have already deployed home directories and that want to maintain compatibility with their existing home directory environment.

- If you redirect My Documents to the home directory, domain administrators have full control over the user's My Documents folder. This is the case even if you select the **Grant the user exclusive rights to My Documents** check box on the **Settings** tab in the **My Documents properties** dialog box.

- Redirection to the home directory requires a client computer running Windows XP Professional, Windows XP 64-bit Edition (Itanium), or a Windows Server 2003 operating system. It does not work for Windows 2000 clients, Windows NT clients, or Windows XP Home Edition clients.

- In home directory redirection, security is not checked, and permissions are not changed. Folder Redirection assumes that the administrator has set directory security correctly. This relaxed security is why redirection to the home directory is only recommended if it is needed for backward compatibility.

- In home directory redirection, no ownership checks are made. Normally, folder redirection fails if a user is not the owner of the folder to which he or she is being redirected. Because redirection to the home directory is intended for use in a

legacy environment, this ownership check is skipped.

- Users must have the home directory property set correctly on their user object in Active Directory. The client computer finds out the path for the user's home directory from the user object in Active Directory at logon time. Users who are affected by Folder Redirection policy must have this path set correctly, or folder redirection will fail.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Redirect Special Folders to the Same Location for All Users
Redirect Special Folders to Different Locations Based on Security Group Membership
Windows Management Instrumentation (WMI) Filters
Folder Redirection

# Redirect special folders to the root directory based on security group membership

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To redirect special folders to the root directory based on security group membership**

1. Open a Group Policy object that is linked to the site, domain, or organizational unit that contains the users whose special folders you want to redirect.

2. In the console tree, double-click **Folder Redirection** to display the special folder that you want to redirect.

   **Where?**

   - Group Policy object\User Configuration\Windows Settings\Folder Redirection.

3. Right-click the special folder that you want to redirect (for example, **Desktop** or **My Documents**), and then click **Properties**.

4. On the **Target** tab, in **Setting**, click **Advanced - Specify locations for various user groups**.

5. Click **Add**.

6. In the **Specify Group and Location** dialog box, under **Security Group Membership**, type the name of the security group, or click **Browse** to search for it.

7. Under **Target Folder Location**, click **Create a folder for each user under the root path**.

8. In **Root Path**, type the UNC path to use for redirection (for example, **\\servername\sharename**), and then click **OK**.

9. In the properties dialog box for the special folder, click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Object Editor, see Related Topics.

- The user name and folder name are appended to the UNC path automatically.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Redirect Special Folders to the Same Location for All Users
Redirect Special Folders to Different Locations Based on Security Group Membership
Windows Management Instrumentation (WMI) Filters
Folder Redirection

# Redirect special folders to a specific path based on security group membership

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To redirect special folders to a specific path based on security group membership**

1. Open a Group Policy object that is linked to the site, domain, or organizational unit that contains the users whose special folders you want to redirect.

2. In the console tree, double-click **Folder Redirection** to display the special folder that you want to redirect.

   **Where?**

   - *Group Policy object*\User Configuration\Windows Settings\Folder Redirection

3. Right-click the special folder that you want to redirect (for example, **Desktop** or **My Documents**), and then click **Properties**.

4. On the **Target** tab, in **Setting**, click **Advanced - Specify locations for various user groups**.

5. Click **Add**.

6. In the **Specify Group and Location** dialog box, under **Security Group Membership**, type the name of the security group, or click **Browse** to search for it.

7. Under **Target Folder Location**, click **Redirect to the following location**.

8. In **Root Path**, type the UNC path to use for redirection (for example, **\\servername\sharename**), and then click **OK**.

9. In the properties dialog box for the special folder, click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Object Editor, see Related Topics.

- In most cases, you want to redirect special folders to the root directory rather than to a specific location, as described here. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Redirect Special Folders to the Same Location for All Users
Redirect Special Folders to Different Locations Based on Security Group Membership
Windows Management Instrumentation (WMI) Filters

Folder Redirection
Redirect special folders to the root directory

# Redirect to the local profile location based on security group membership

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To redirect special folders to the local profile location based on security group membership**

1. Open a Group Policy object that is linked to the site, domain, or organizational unit that contains the users whose special folders you want to redirect.

2. In the console tree, double-click **Folder Redirection** to display the special folder that you want to redirect.

   **Where?**

   - *Group Policy object*\Computer Configuration\Windows Settings\Folder Redirection

3. Right-click the special folder that you want to redirect (for example, **Desktop** or **My Documents**), and then click **Properties**.

4. On the **Target** tab, in **Setting**, click **Advanced - Specify locations for various user groups**.

5. Click **Add**.

6. In the **Specify Group and Location** dialog box, under **Security Group Membership**, type the name of the security group, or click **Browse** to search for it.

7. Under **Target Folder Location**, click **Redirect to the local user profile location**, and then click **OK**.

8. In the properties dialog box for the special folder, click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.

- To open Group Policy Object Editor, see Related Topics.

- This procedure is useful for returning redirected folders to their original, default location.

- If you change the setting to **No administrative policy specified** on the **Target** tab of the properties dialog box for a special folder, you do not cause the special folder to be redirected back to its default location.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Redirect Special Folders to the Same Location for All Users
Redirect Special Folders to Different Locations Based on Security Group Membership
Windows Management Instrumentation (WMI) Filters

# Configure Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Group Policy**

- Extend the functionality of Group Policy
- Refresh Group Policy immediately
- Set permissions on a Group Policy object
- Set the Group Policy refresh rate for computers
- Set other policies for Group Policy

# Extend the functionality of Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To extend the functionality of Group Policy**

1. Open Group Policy Object Editor as an MMC snap-in.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. Click the **Extensions** tab, click the snap-ins that you want to extend, click the extensions that you want to add, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ways to open Group Policy Object Editor
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Edit the local Group Policy object
Advanced methods of extending Group Policy

# Refresh Group Policy immediately

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To refresh Group Policy immediately**

1. Click **Start**, and then click **Run**.
2. In the **Open** box, type **gpupdate**, and then click **OK**.

### Note

- To learn more about **gpupdate**, **gpresult**, and other command-line tools, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Gpupdate
Gpresult
Command-line reference

# Set permissions on a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions on a Group Policy object**

1. Open Group Policy Object Editor.
2. In the console tree, right-click the icon or the name of the Group Policy object for which you want to set permissions, and then click **Properties**.
3. Click the **Security** tab, and then specify the following options:
   - **Add--**Opens the **Select Users, Computers, or Groups** dialog box, where you can specify the users and groups for whom you want to assign permissions.
   - **Remove--**In the **Group or user names** box, removes selected users or groups and their associated permissions from this object.
   - **Permissions for Authenticated Users--**Lists the standard permissions that you can allow or deny to users, for example, Full Control, Read, Write, and so on.
   - **Advanced--**Use this option to set special permissions, auditing information, and owner information for the selected object.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- You can only set permissions on a nonlocal Group Policy object.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Access control overview
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object
Filter the scope of Group Policy according to security group membership

# Set the Group Policy refresh rate for computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the Group Policy refresh rate for computers**

1. Open Group Policy Object Editor.
2. In the console tree, click **Group Policy**.

   **Where?**

   - *Group Policy object*\Computer Configuration\Administrative Templates\System\Group Policy

3. In the details pane, double-click **Group Policy refresh interval for computers**.
4. On the **Setting** tab, click **Enabled**.
5. In the **Minutes** boxes, type or click the number of minutes for the refresh interval and random offset, and then click **OK**.
6. In the details pane, double-click **Turn off background refresh of Group Policy**.
7. On the **Setting** tab, click **Disabled**, and then click **OK**.

**Notes**

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- By default, the System.adm file is installed in Computer Configuration\Administrative Templates so that you can view these options. For information about how to add or remove an Administrative Templates file like System.adm, see Related Topics.
- Client computers request policy from a domain controller when any of the following events occur:
  - A computer starts up.
  - A user logs on.
  - An application requests a refresh through the **RefreshPolicy()** application programming interface (API).
  - The user requests an immediate refresh.
  - One of the Group Policy refresh interval policies is enabled, and the interval has transpired.
  - Policy is not imposed by the domain controller; it is requested by the client computer.
- The minimum setting for the Group Policy refresh interval (zero minutes) actually takes about seven seconds between requests for a refresh. Very short refresh interval settings should not be used in a production environment, except briefly for testing.

- Group Policy Software Installation and Folder Redirection are not affected by policy refreshes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object
Add or remove an Administrative Template (.adm file)

# Set other policies for Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set other policies for Group Policy**

1. Open Group Policy Object Editor.
2. Do one of the following:
   - To edit a computer policy, in the console tree click **Group Policy**.
   - **Where?**
   - *Group Policy object*/Computer Configuration/Administrative Templates/System/Group Policy
   - To edit a user policy, in the console tree click **Group Policy**.
   - **Where?**
   - *Group Policy object*/User Configuration/Administrative Templates/System/Group Policy
3. In the details pane, double-click the icon or the name of the policy that you want to set.
4. On the **Setting** tab, select the options that you want, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- To open Group Policy Object Editor, see Related Topics.
- By default, the most common Administrative Templates are displayed. You can add Administrative Templates by right-clicking the Administrative Templates folder, and then clicking **Add/Remove Templates**. For more information, see Add or remove an Administrative Template (.adm file).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object
Add or remove an Administrative Template (.adm file)

# Delegate Control of Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Delegate control of Group Policy**

- Delegate Group Policy object editing
- Delegate creation of Group Policy objects
- Delegate Group Policy object linking

# Delegate Group Policy object editing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate Group Policy object editing**

1. Open Group Policy Object Editor.
2. In the console tree, right-click the icon or name of the Group Policy object, and then click **Properties**.
3. Do any of the following:
   - To add a user or group of users, click the **Security** tab, and then click **Add**. Type the name of the user or group of users, and then click **OK**.
   - To give a user or group Group Policy object editing permissions, in **Permissions for Authenticated Users**, select the check boxes that correspond to the permissions you want to give. When you are finished, click **OK**.
   - If you do not want to apply the policy to a user or group of users, under **Permissions for Authenticated Users**, clear the **Allow** check box for Apply Group Policy. When you are finished, click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Admins group or the Enterprise Admins group.
- To open Group Policy Object Editor, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Delegate creation of Group Policy objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate creation of Group Policy objects**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.
3. In the **Name** column in the details pane, double-click **Group Policy Creator Owners**.
4. In the **Group Policy Creator Owners Properties** dialog box, click the **Members** tab.
5. Click **Add**, type the name of the user or group of users, and then click **OK**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Admins group or the Enterprise Admins group.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- By default, only domain administrators, enterprise administrators, Group Policy Creator Owners, and the operating system can create new Group Policy objects. If the domain administrator wants a nonadministrator or a group to be able to create Group Policy objects, that user or group can be added to the Group Policy Creator Owners security group. When a user who is not an administrator, but who is a member of the Group Policy Creator Owners group, creates a Group Policy object, that user becomes the creator and owner of the Group Policy object; therefore, that user can edit the Group Policy object. Being a member of the Group Policy Creator Owners group gives the user full control of only those Group Policy objects that the user creates or those Group Policy objects that are explicitly delegated to that user. It does not give the nonadministrator user any additional rights over other Group Policy objects for the domain--these users are not granted rights over Group Policy objects that they did not create.
- When an administrator creates a Group Policy object, the Domain Admins group becomes the Creator Owner of the Group Policy object.
- When you delegate this task to nonadministrators, also consider delegating the ability to manage the links for a specific organizational unit. The reason for this is that, by default, nonadministrators cannot manage links, and the inability to manage links prevents them from being able to use Active Directory Users and Computers to create a Group Policy object.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Delegate Group Policy object linking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate Group Policy object linking**

1. Open Active Directory Users and Computers.
2. Right-click the organizational unit to which you want to delegate the right to link Group Policy objects, and then click **Delegate Control**.
3. In the Delegation of Control Wizard, click **Next**, and then click **Add**.
4. In the **Select Users, Computers, or Groups** dialog box, enter the object name to select, and then click **OK**, and then click **Next**.
5. On the Task to Delegate page, in **Delegate the following common tasks**, select the **Manage Group Policy links** check box, and then click **Next**.
6. Click **Finish**.

## Notes

- To complete this procedure, you must be logged on as a member of the Domain Admins group or the Enterprise Admins group.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To open Group Policy Object Editor, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object

# Group Policy Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**
This section provides general background information about Group Policy, and about the change and configuration management technologies that use Group Policy.

- Group Policy overview
- Group Policy settings overview
- Understanding Group Policy
- Using Group Policy
- Group Policy Resources

# Group Policy overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy overview**

To begin using Group Policy immediately, see Ways to open Group Policy Object Editor.

Group Policy settings define the various components of the user's desktop environment that a system administrator needs to manage, for example, the programs that are available to users, the programs that appear on the user's desktop, and options for the **Start** menu. To create a specific desktop configuration for a particular group of users, use Group Policy Object Editor. Group Policy settings that you specify are contained in a Group Policy object, which is in turn associated with selected Active Directory® objects--sites, domains, or organizational units.

Group Policy applies not only to users and client computers, but also to member servers, domain controllers, and any other Microsoft® Windows® 2000 computers within the scope of management. By default, Group Policy that is applied to a domain (that is, applied at the domain level, just above the root of Active Directory Users and Computers) affects all computers and users in the domain. Active Directory Users and Computers also provides a built-in Domain Controllers organizational unit. If you keep your domain controller accounts there, you can use the Group Policy object Default Domain Controllers Policy to manage domain controllers separately from other computers.

Group Policy includes policy settings for User Configuration, which affect users, and for Computer Configuration, which affect computers. For more information, see User Configuration and Computer Configuration.

With Group Policy you can do the following:

- **Manage registry-based policy with Administrative Templates.** Group Policy creates a file that contains registry settings that are written to the User or Local Machine portion of the registry database. User profile settings that are specific to a user who logs on to a given workstation or server are written to the registry under **HKEY_CURRENT_USER** (HKCU), and computer-specific settings are written under **HKEY_LOCAL_MACHINE** (HKLM). For procedural information, see Use Administrative Templates. For technical details, see "Implementing Registry-based Policy" at the Microsoft Web site.

- **Assign scripts.** This includes such scripts as computer startup, shutdown, logon, and logoff. For more information, see Use Startup, Shutdown, Logon, and Logoff Scripts.

- **Redirect folders.** You can redirect folders, such as My Documents and My Pictures, from the Documents and Settings folder on the local computer to network locations. For more information about Folder Redirection, see Use Folder Redirection.

- **Manage applications.** With Group Policy you can assign, publish, update, or repair applications by using Group Policy Software Installation. For more information, see Use Group Policy Software Installation.

- **Specify security options.** To learn about setting security options, see Security Settings extension to Group Policy.

**Group Policy objects that exist by default**

Each computer that runs Windows 2000, Microsoft® Windows® XP Professional , or Windows Server 2003 operating systems has exactly one Group Policy object stored locally. This local Group Policy object contains a subset of the settings that are available in nonlocal Group Policy objects. For more information, see Local Group Policy.

By default, when Active Directory is set up, two nonlocal Group Policy objects are created:

- Default Domain Policy is linked to the domain, and it affects all users and computers in the domain (including computers that are domain controllers) through policy inheritance. For more information, see Policy inheritance.

- Default Domain Controllers Policy is linked to the Domain Controllers organizational unit, and it generally only affects domain controllers, because computer accounts for domain controllers are kept exclusively in the Domain Controllers organizational unit.

### Caution

- If you move a domain controller outside of the Domain Controllers organizational unit, the default Domain Controllers policy is no longer applied. If you need to keep a domain controller in a different organizational unit, link the default Domain Controllers policy to that organization unit.

#### How and when Group Policy is applied

To view a detailed chronology for the application of policy, see Order of events when starting up and logging on.

#### User and computer policy

User policy settings are located under User Configuration in Group Policy, and they are obtained when a user logs on. Computer policy settings are located under Computer Configuration, and they are obtained when a computer starts. For more information, see User Configuration and Computer Configuration.

Users and Computers are the only types of Active Directory objects that receive policy. Specifically, policy does not apply to security groups. Instead, for performance reasons, security groups are used to filter the policy by means of an Apply Group Policy access control entry (ACE), which can be set to **Allow** or **Deny** or left unconfigured. For more information, see Filter the scope of Group Policy according to security group membership.

#### Order of application

Policies are applied in this order:

1. The unique local Group Policy object.
2. Site Group Policy objects, in administratively specified order.
3. Domain Group Policy objects, in administratively specified order.
4. Organizational unit Group Policy objects, from largest to smallest organizational unit (parent to child organizational unit) and in administratively specified order at the level of each organizational unit.

For more information, see Order of processing settings.

By default, policies that are applied later overwrite previously applied policies when the policies are inconsistent. If the policies are not inconsistent, however, earlier and later policies both contribute to the effective policy. For more information, see Group Policy Precedence.

#### Filtering policy by security group membership

A security group ACE on a Group Policy object can be set to **Not configured** (no preference), **Allowed**, or **Denied**. **Denied** takes precedence over **Allowed**. For more information, see Filter the scope of Group Policy according to security group membership.

#### Blocking policy inheritance

Policies that would otherwise be inherited from higher sites, domains, or organizational units can be blocked at the site, domain, or organizational unit level. For more information, see Block policy inheritance.

#### Enforcing policy from above

You can set policies that would otherwise be overwritten by policies in child organizational units to **No Override** at the Group Policy object level. For more information, see Prevent a Group Policy object from being overridden.

### Notes

- Policies set to **No Override** cannot be blocked.
- The **No Override** and **Block** options should be used sparingly. Casual use of these advanced features complicates troubleshooting. For tips about using Group Policy, see Troubleshooting Group Policy and Group Policy Best Practices (pre-GPMC).

# Group Policy settings overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy settings overview**
With Group Policy settings, you can define the configuration of a user's desktop environment. Settings are available for Computer Configuration and User Configuration. In addition to client computers and users, you can apply these settings to member servers and domain controllers within the scope of an Active Directory forest.

**Administrative Templates folder**
The Administrative Templates folders are available under Computer Configuration and User Configuration in Group Policy Object Editor. These folders contain registry settings. For more information, see Administrative Templates.

The following table contains links to Help for all of the registry settings included in the five default Administrative Templates.

| For these settings | Open this Administrative Template | Reference |
|---|---|---|
| System settings that control features such as desktop, network connections, shared folders, and Control Panel. | System (System.adm) | Group Policy settings for Windows 2000, Windows XP, and the Windows Server 2003 family |
| Internet Explorer settings that control features such as security zones, proxy settings, search, and temporary Internet files. | Internet Explorer (Inetres.adm) | Group Policy settings for Internet Explorer |
| Windows Media Player settings that control features such as proxy settings, codec, and network buffering. This tool is not available on the Itanium-based versions of the Windows operating systems. This tool is not available on the x64-based versions of the Windows operating systems. | Windows Media Player (Wmplayer.adm) | Group Policy settings for Windows Media Player |
| NetMeeting settings to control features such as application sharing, audio, video, and chat. This tool is not available on the Itanium-based versions of the Windows operating systems. This tool is not available on the x64-based versions of the Windows operating systems. | NetMeeting (Conf.adm) | Group Policy settings for NetMeeting v3.01 |
| Windows Update settings that control features such as automatic software updates through the Internet. | Windows Update (Wuau.adm) | Group Policy settings for Windows Update |

**Software Settings folder**
The Software Settings folders are available under Computer Configuration and the User Configuration in Group Policy Object Editor. The Computer Configuration\Software Settings folder contains software settings that apply to all users who log on to the computer. This folder contains software installation settings, and it might contain other settings that are placed there by independent software vendors. The User Configuration\Software Settings folder contains software settings that apply to users regardless of which computer they log on to. This folder also contains software installation settings, and it might contain other settings that are placed there by independent software vendors. For more information, see Software Settings.

**Windows Settings folder**
The Windows Settings folders are available under Computer Configuration and User Configuration in Group Policy Object Editor. The Computer Configuration\Windows Settings folder contains Windows settings that apply to all users who log on to the computer. This folder also contains the following items: Security Settings and Scripts. The User Configuration\Windows Settings folder contains Windows settings that apply to users regardless of which computer they log on to. This folder also contains the following items: Folder Redirection, Security Settings, and Scripts. For more information, see Windows Settings.

**Security Settings folder**
The Security Settings folders are available under Computer Configuration and User Configuration in Group Policy Object Editor. Security settings or security policies are rules that you configure on a computer or multiple computers that protect

resources on a computer or network. With security settings you can specify the security policy of an organizational unit, domain, or site. To view a complete list and detailed information about security settings, see Security Setting Descriptions. To specify security settings on a local computer, see Local Security Policy. For more information, see Security Settings extension to Group Policy.

# Understanding Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Group Policy**
This section covers:

- New Group Policy features (pre-GPMC)
- Group Policy objects
- Policy inheritance
- Storage of Group Policy objects
- User Configuration
- Computer Configuration
- Administrative Templates
- Software Settings
- Fast Logon Optimization
- Windows Settings
- Folder Redirection
- Windows Management Instrumentation (WMI) Filters
- Startup, shutdown, logon, and logoff scripts
- Group Policy Software Installation
- Group Policy Precedence

# New Group Policy features (pre-GPMC)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features**

The Windows 2000 change and configuration management technologies have been improved in the Windows Server 2003 family. For information about specific user interface changes in the Windows Server 2003 family, see New ways to do familiar Group Policy tasks (pre-GPMC).

Following is a list of the new features included with Group Policy in the Windows Server 2003 family:

- **Administrative Templates**

  There are over 220 new administrative policies than were previously available on Windows 2000. There is a new supported keyword in .adm files that tells you which versions of Windows are supported as clients for the setting. For more information, see "Implementing Registry-based Policy" at the Microsoft Web site. There is also a new Administrative Templates Web view that uses the supported keyword to show you which operating systems are supported clients for individual settings. For information about how to use the Web view, see Use the view provided by Administrative Templates. For more information about Administrative Templates, see Administrative Templates.

  Administrative Templates online Help is now included with Windows. The following table contains links to the .adm Help files.

| .adm file | Policy settings |
|---|---|
| System.adm | Open Help for System Settings |
| Wmplayer.adm | Open Help for Windows Media Player settings |
| Inetres.adm | Open Help for Internet Explorer settings |
| Conf.adm | Open Help for NetMeeting settings |

- **Command-line tools**

  The new **gpudate** command replaces the Windows 2000 command **secedit /refreshpolicy**. For more information, see Gpupdate.

  The Windows 2000 Resource Kit command **gpresult** has been enhanced, and it is now included with Windows XP Professional, Windows XP 64-bit Edition (Itanium), and the Windows Server 2003 family. For more information, see Gpresult.

- **Folder Redirection**

  It is now easier to use Folder Redirection. It is no longer necessary to type variables such as *UserName* as part of the redirection path. There is a new option for redirecting My Documents to a user's home directory. This is for use in environments where home directories are already deployed. For more information, see Use Folder Redirection.

- **Group Policy Software Installation**

  Group Policy Software Installation now has an advanced deployment option to enable or disable availability of 32-bit programs to 64-bit computers. With another new advanced option, you can enable or disable (default) publication of OLE class information about a software package. As an administrator, you also have an option that forces assigned applications to be installed at deployment time. For more information, see Edit installation options for applications.

- **Resultant Set of Policy**

Resultant Set of Policy (RSoP) is a new tool that you can use to plan for and troubleshoot policy. For more information, see Resultant Set of Policy.

- **Support across forests**

  Group Policy features are now supported in multiple forest environments allowing authenticated access to resources in separate, networked environments. You can manage and enforce policy objects for user and computer objects in remote trusted forests and analyze policies using Resultant Set of Policy. For more information, see Group Policy features that are supported across forests and Accessing resources across forests.

# Group Policy objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy objects**

Policy settings are stored in Group Policy objects. Group Policy Object Editor can be thought of as an application whose document type is the Group Policy object, just as a word processor might use .doc or .txt files.

There are two kinds of Group Policy objects: local and nonlocal. Local Group Policy objects are stored on individual computers. Only one local Group Policy object exists on a computer, and it has a subset of the settings that are available in a nonlocal Group Policy object. Local Group Policy object settings can be overwritten by nonlocal settings if they are in conflict; otherwise, both groups of settings apply. For more information, see Local Group Policy.

Nonlocal Group Policy objects, which are stored on a domain controller, are available only in an Active Directory environment. They apply to users and computers in the site, domain, or organizational unit with which the Group Policy object is associated. For more information, see Storage of Group Policy objects.

# Policy inheritance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policy inheritance**

In general, Group Policy is passed down from parent to child containers within a domain, which you can view by using Active Directory Users and Computers. Group Policy is not inherited from parent to child domains, for example, from wingtiptoys.com to sales.wingtiptoys.com. Active Directory Domains and Trusts, which you can use to manage relationships of this type, is not related to Group Policy.

If you assign a specific Group Policy setting to a high-level parent container, that Group Policy setting applies to all containers beneath the parent container, including the user and computer objects in each container. However, if you explicitly specify a Group Policy setting for a child container, the child container's Group Policy setting overrides the parent container's setting.

If a parent organizational unit has policy settings that are not configured, the child organizational unit does not inherit them. Policy settings that are disabled are inherited as disabled. In addition, if a policy setting is configured (enabled or disabled) for a parent organizational unit and the same policy setting is not configured for a child organizational unit, the child inherits the parent's enabled or disabled policy setting.

If a policy setting that is applied to a parent organizational unit and a policy setting that is applied to a child organizational unit are compatible, the child organizational unit inherits the parent policy setting, and the child's setting is also applied.

If a policy setting that is configured for a parent organizational unit is incompatible with the same policy setting that is configured for a child organizational unit (because the setting is enabled in one case and disabled in the other), the child does not inherit the policy setting from the parent. The policy setting in the child is applied.

**Blocking inheritance**

You can block policy inheritance at the domain or organizational-unit level by opening the properties dialog box for the domain or organizational unit and selecting the **Block Policy inheritance** check box. For more information, see Block policy inheritance.

**Enforcing inheritance**

You can enforce policy inheritance by setting the **No Override** option on a Group Policy object link.

When you select the **No Override** check box, you force all child policy containers to inherit the parent's policy, even if that policy conflicts with the child's policy and even if **Block Inheritance** has been set for the child.

You can set **No Override** on a Group Policy object link by opening the properties dialog box for the site, domain, or organizational unit and making sure that the **No Override** check box is selected. For more information, see Prevent a Group Policy object from being overridden.

### Notes

- Policies that are set to **No Override** cannot be blocked.
- It is recommended that **No Override** and **Block** be used sparingly. Casual use of these advanced features makes troubleshooting more complicated. For more information, see Troubleshooting Group Policy and Group Policy Best Practices (pre-GPMC).

# Storage of Group Policy objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Storage of Group Policy objects**
Each computer that runs Windows XP Professional, Windows XP 64-bit Edition (Itanium), or the Windows Server 2003 operating systems, has exactly one local Group Policy object (GPO). It is stored in *systemroot*\System32\GroupPolicy.

Group Policy objects, other than the local Group Policy object, are virtual objects. The policy setting information of a GPO is actually stored in two locations: the Group Policy container and the Group Policy template. The Group Policy container is an Active Directory container that stores GPO properties, including information on version, GPO status, and a list of components that have settings in the GPO. The Group Policy template is a folder structure within the file system that stores Administrative Template-based policies, security settings, script files, and information regarding applications that are available for Group Policy Software Installation. The Group Policy template is located in the system volume folder (Sysvol) in the \Policies subfolder for its domain. For more information about the local Group Policy object, see Local Group Policy.

**Group Policy container**
The Group Policy container is a directory service object. It includes subcontainers for computer and user Group Policy information. The Group Policy container contains the following data:

- **Version information--**Used to verify that the information is synchronized with Group Policy template information.
- **Status information--**Indicates whether the Group Policy object is enabled or disabled for this site, domain, or organizational unit.
- **List of components--**Specifies which extensions to Group Policy have settings in the Group Policy object.

The Group Policy container stores information for Group Policy Software Installation and for Folder Redirection, which are extensions of the Group Policy Object Editor. For more information, see Group Policy Software Installation overview and Folder Redirection.

**Group Policy template**
The Group Policy template is a folder of domain controllers for the storage domain of the Group Policy object. A typical Group Policy template folder might look like the following example:

*systemroot*\Sysvol\SYSVOL\Streetmarket.com\Policies\

{34975054-fd77-df75-54fe-074936850457}

**Subfolders of the Group Policy template**
The Group Policy template folder contains subfolders, including, but not limited to, the following:

- **Adm--**Contains all the .adm files for this Group Policy template.
- **Scripts--**Contains all the scripts and related files for this Group Policy template.
- **User--**Includes a Registry.pol file that contains the registry settings that are to be applied to users. When a user logs on to a computer, this Registry.pol file is downloaded and applied to the **HKEY_CURRENT_USER** portion of the registry. The User folder contains an Applications subfolder.
- **User\Applications--**Contains the application advertisement script files (.aas) that are used by the operating system-based installation service. These files are applied to users.
- **Machine--**Includes a Registry.pol file that contains the registry settings that are to be applied to computers. When a computer initializes, this Registry.pol file is downloaded and applied to the **HKEY_LOCAL_MACHINE** portion of the registry. The Machine folder contains an Applications subfolder.
- **Machine\Applications--**Contains the .aas files that are used by the operating system-based installation service. These

files are applied to computers.

# User Configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User Configuration**

You can use User Configuration in Group Policy to set policies that apply to users, regardless of which computer they log on to.

User Configuration typically contains subitems for Software Settings, Windows Settings, and Administrative Templates, but because you can add or remove extensions from Group Policy Object Editor, the exact set of subitems that you see might be different. For more information about these subitems, see Software Settings, Windows Settings, and Administrative Templates.

# Computer Configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computer Configuration**
With Computer Configuration in Group Policy, you can set policies that are applied to computers, regardless of who logs on to the computers.

Computer Configuration typically contains Software Settings, Windows Settings, and Administrative Templates. However, because you can add or remove extensions in Group Policy Object Editor, what you see might be slightly different.

For more information about these subitems, see Software Settings, Windows Settings, and Administrative Templates.

# Administrative Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administrative Templates**
Administrative Templates provide policy information for the items that appear under the Administrative Templates folder in the console tree of the Group Policy Object Editor.

**Administrative Template .adm files**
An .adm file consists of a hierarchy of categories and subcategories that together define how the policy settings appear. It also contains the following information:

- Registry locations that correspond to each setting.
- Options or restrictions in values that are associated with each setting.
- For many settings, a default value.
- Explanations of what each setting does.
- The versions of Windows that support each setting.

The Administrative Templates for Group Policy contain all registry-based policy information. User configurations are saved in **HKEY_CURRENT_USER** (HKCU), and computer configurations are saved in **HKEY_LOCAL_MACHINE** (HKLM). HKCU and HKLM place registry information that is specific to Group Policy under \Software\Policies or under \Software\Microsoft\Windows\CurrentVersion\Policies. Therefore, there are five areas of the registry that contain Group Policy registry settings.

Administrative Templates offer these new features:

- A description of each policy setting in Administrative Templates appears in searchable online Help.

| Administrative Template | Description |
|---|---|
| System.adm | System settings |
| Inetres.adm | Internet Explorer settings |
| Wmplayer.adm | Windows Media Player settings.<br>This tool is not available on the Itanium-based versions of the Windows operating systems.<br>This tool is not available on the x64-based versions of the Windows operating systems. |
| Conf.adm | NetMeeting settings.<br>This tool is not available on the Itanium-based versions of the Windows operating systems.<br>This tool is not available on the x64-based versions of the Windows operating systems. |
| Wuau.adm | Windows Update settings. |

- A supported keyword in the .adm files that tells you which versions of Windows are supported as clients for the setting. For information about the supported keyword, see "Implementing Registry-based Policy" at the Microsoft Web site.
- A Web view that shows the explain text of each setting and the new supported information. For information about how to use the Web view, see Use the view provided by Administrative Templates.
- The ability to filter settings from the view to reduce screen clutter when you are using Administrative Templates. For

information about how to use this new feature, see Filter the view provided by Administrative Templates.

If you plan to create entries for Administrative Templates, populate the namespace by using the following naming convention, which is also used in the registry: *CompanyName**product**version* (or *CompanyName**product&version*). For example, the operating system settings for Windows are in \Microsoft\Windows.

For information about adding an .adm file, see Add or remove an Administrative Template (.adm file). For more information about Administrative Templates, see Administrative Templates included with this version of Windows and The role of Administrative Templates.

**Registry.pol files**

The Administrative Templates extension of Group Policy saves information in Registry.pol files. These files contain the customized registry settings that you specify (by using Group Policy) that are to be applied to the computer or user portion of the registry:

- One of the Registry.pol files contains registry settings that are specific to the **HKEY_LOCAL_MACHINE** key; it is stored in the GPT\Machine folder.

- The other Registry.pol file contains registry settings that are specific to the **HKEY_CURRENT_USER** key; it is stored in the GPT\User subdirectory.

# Software Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Software Settings**

Software Settings are available under both User Configuration and Computer Configuration in the console tree of Group Policy Object Editor.

Computer Configuration\Software Settings is for software settings that apply to all users who log on to the computer. This folder contains the Software installation subitem, and it might contain other subitems that are placed there by independent software vendors.

User Configuration\Software Settings is for software settings that apply to users regardless of which computer they log on to. This folder also contains the Software installation subitem, and it might contain other subitems that are placed there by independent software vendors.

For more information about the Computer Configuration and User Configuration folders, see User Configuration and Computer Configuration. For more information about the Software installation subitem, see Group Policy Software Installation overview.

# Fast Logon Optimization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Fast Logon Optimization**

The Fast Logon Optimization feature is set by default for both domain and workgroup members. This results in the asynchronous application of policy when the computer starts up and when the user logs on. This application of policy is similar to a background refresh. It can reduce the length of time it takes for the logon dialog box to appear and the length of time it takes for the desktop to become available to the user.

You can change the default setting by using the Group Policy MMC snap-in. This feature is enabled by default, but you can disable it by using the following policy: Administrative Templates\System\Logon**Always wait for the network at computer startup and logon**.

Policies are processed according to the criteria listed in the following table.

| Default policy process for clients | At startup | At logon | At policy refresh |
|---|---|---|---|
| Windows 2000 | Sync | Sync | Async |
| Windows Server 2003 family | Sync | Sync | Async |
| Windows XP Professional and Windows XP 64-bit Edition (Itanium) | Async | Async | Async |

**Notes**

- Logon Optimization will not be enabled and policies will be processed synchronously when a user is logging on for the first time, the user has a roaming profile, the user has a HomeDir, and the user has a logon script specified in the User object. Folder Redirection and Group Policy Software Installation require a synchronous application of policy. Under these conditions, computer startup can still be asynchronous. However, since logon is synchronous, logon does not exhibit optimization.

- Client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems support Fast Logon Optimization in any domain environment.

# Windows Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Settings**

Windows Settings are available under both User Configuration and Computer Configuration in the console tree of Group Policy Object Editor.

Computer Configuration\Windows Settings is for windows settings that apply to all users who log on to the computer. It includes two subitems: Security Settings and Scripts.

User Configuration\Windows Settings is for windows settings that apply to users regardless of which computer they log on to. It includes three subitems: Folder Redirection, Security Settings, and Scripts.

For more information about the Computer Configuration and User Configuration folders, see User Configuration and Computer Configuration. For more information about the settings in the Windows Settings subitems, see Security Settings extension to Group Policy, Startup, shutdown, logon, and logoff scripts, and Folder Redirection.

# Folder Redirection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Folder Redirection**

In Group Policy Object Editor, you can use Folder Redirection to redirect certain special folders to network locations. Special folders are those folders, such as My Documents and My Pictures, that are located under Documents and Settings. Folder Redirection is located under User Configuration in the console tree of Group Policy Object Editor.

There are several basic options for Folder Redirection. For each basic option, there is an advanced version of that option. The advanced version provides for finer control by allowing redirection that is based on security group membership. For more information about specific procedures, see Use Folder Redirection. The following table describes the different folder redirection options that are available to you.

| Special folder | Notes |
|---|---|
| Application Data | A Group Policy setting controls the behavior of Application Data when client-side caching is enabled. Look in User Configuration\Administrative Templates\Network\Offline Files in the console tree of Group Policy Object Editor. |
| Desktop | Desktop can be redirected independently of all the other special folders. |
| My Documents | My Documents is the default location in the shell for users to save their documents and pictures. |
| My Documents\My Pictures | My Pictures can be redirected independently of My Documents, or it can be made to follow My Documents (to remain its subfolder whenever My Documents is redirected), as it does by default. The default behavior is recommended unless you have a specific reason (such as server scalability) for separating My Pictures from My Documents. If these folders are separated, a shortcut takes the place of the My Pictures folder in My Documents. |
| Start Menu | When Start Menu is redirected, its subfolders always follow. |

**Caution**

- When creating a shared folder redirection directory, limit access to only those users that need access. Redirected folders can contain personal information such as confidential documents and EFS certificates, care should be taken to protect access to the shared folder. Restrict access to the shared folder to only those users that need access. You can also create a security group for users that require permissions for a particular shared folder, and limit access to only those users.

  When creating the shared folder, hide it by putting a **$** after the share name. This hides the shared folder from casual browsers, and it will not be visible in My Network Places.

**Advantages of redirecting My Documents**

Some of the following benefits pertain to redirecting any folder, but redirecting My Documents can be particularly advantageous because this folder tends to become large over time.

- When roaming user profiles are used, only the network path to the My Documents folder is part of the roaming user profile, not the My Documents folder. Therefore, its contents do not have to be copied back and forth between the client computer and the server each time the user logs on or off, and the process of logging on or off can be much faster than it was in Windows NT 4.0.

- Even if a user logs on to various computers on the network, his or her documents are always available.

- Offline File technology gives users access to My Documents even when they are not connected to the network. This is particularly useful for people who use portable computers. For more information, see Make a file or folder available offline.

- Data that is stored in a shared network folder can be backed up as part of routine system administration. This is safer

because it requires no action on the part of the user.

- As an administrator, you can use Group Policy to set disk quotas, limiting the amount of space that is taken up by users' special folders.

- Data that is specific to a user can be redirected to a different hard disk on the user's local computer from the hard disk that holds the operating system files. This makes the user's data safer if the operating system is reinstalled.

For tips about using Folder Redirection, see Best practices for Folder Redirection.

**Granting exclusive rights to special folders**

The **Settings** tab in each folder's properties dialog box contains a check box labeled **Grant the user exclusive rights to My Documents**. If you select this check box, the user and the local system have full control over the folder, and no one else, not even the administrator, has any rights to it. If you clear this check box, no changes are made to the permissions on the folder. Whatever permissions are in effect by default remain in effect.

**Policy removal considerations with regard to Folder Redirection**

The following table summarizes what happens to redirected folders and their contents when the Group Policy object no longer applies.

| Move the contents of the special folder to the new location setting | Policy Removal option | Results when policy is removed |
|---|---|---|
| **Enabled** | **Redirect the folder back to the user profile location when policy is removed** | • The special folder returns to its user profile location.<br>• The contents are copied, not moved, back to the user profile location.<br>• The contents are not deleted from the redirected location.<br>• The user continues to have access to the contents, but only on the local computer. |
| **Disabled** | **Redirect the folder back to the user profile location when policy is removed** | • The special folder returns to its user profile location.<br>• The contents are not copied or moved to the user profile location.<br>**Caution**<br>• If the contents of a folder are not copied to the user profile location, the user cannot see them.<br>• |

| | | |
|---|---|---|
| Either **Enabled** or **Disabled** | **Leave the folder in the new location when policy is removed** | • The special folder remains at its redirected location.<br>• The contents remain at the redirected location.<br>• The user continues to have access to the contents at the redirected folder. |

**Folder Redirection and Offline Files**

The Offline Files technology applies to any mounted or mapped drive that contains documents or data that a user might want to use offline. Offline Files does not depend on Folder Redirection. It is set up and configured on shared network servers separately from the Folder Redirection snap-in. Offline Files enables the user to do useful work even when the user is not connected to the network, for example, on a portable computer or in the event of router failure. For more information, see Offline Files.

If you use redirected folders of any type, it is recommended that you set up Offline Files as described in the following table.

| Special Folder | Offline File configuration |
|---|---|
| My Documents | Autocaching for documents (or manual caching for documents, if you want users to have to manually make files and folders available for offline use) |
| My Pictures | Autocaching for documents (or manual caching for documents, if you want users to have to manually make files and folders available for offline use) |
| Application Data | Autocaching for programs |
| Desktop | Autocaching for programs if the desktop is Read Only |
| Start Menu | Autocaching for programs |

**Folder Redirection permissions**

This is an advanced topic. If you let Folder Redirection create folders for you, which is the recommended procedure, correct permissions are set automatically. Usually, knowledge of these permissions is not necessary. However, there are two reasons the permissions might be of interest:

- Sometimes, even though it is not recommended, administrators create the redirected folders before Folder Redirection creates them. The following table shows what permissions have to be set for Folder Redirection to work.
- Redirection of My Documents to the home directory provides more relaxed security than standard folder redirection. The following table shows what security is in effect in the standard case.

**NTFS permissions required for the root folder**

| User account | Folder Redirection defaults | Minimum permissions needed |
|---|---|---|
| Creator/owner | Full Control, this folder, subfolders, and files | Full Control, this folder, subfolders, and files |
| Administrators | No permissions | No permissions |
| Everyone | No permissions | No permissions |
| Local System | Full Control, this folder, subfolders, and files | Full Control, this folder, subfolders, and files |

| | | |
|---|---|---|
| Security group of users who need to put data on the shared network server | N/A | List Folder/Read Data, Create Folders/Append Data - **This folder only** |

**Share-level (SMB) permissions required for the root folder**

| User Account | Folder Redirection defaults | Minimum permissions needed |
|---|---|---|
| Everyone | Full Control | No permissions (Use security group) |
| Security group of users who need to put data on the shared network server | N/A | Full Control |

**NTFS permissions required for each user's redirected folder**

| User account | Folder Redirection defaults | Minimum permissions needed |
|---|---|---|
| *UserName* | Full Control, owner of folder | Full Control, owner of folder |
| Local System | Full Control | Full Control |
| Administrators | No permissions | No permissions |
| Everyone | No permissions | No permissions |

# Windows Management Instrumentation (WMI) Filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Management Instrumentation (WMI) Filters**

Group Policy provides policy setting information to the WMI Common Information Model (CIM)-compliant object repository (CIMON database). With WMI Filters, you can specify a WMI-based query to filter the effect of a Group Policy object in the CIMON database. WMI Filters can be written and customized using the WMI Query Language (WQL).

For more information, see WMI Overview.

# Startup, shutdown, logon, and logoff scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Startup, shutdown, logon, and logoff scripts**

Group Policy Object Editor includes two extensions for script deployment:

- **Scripts (Startup/Shutdown)--**You can use this extension, which is located under Computer Configuration\Windows Settings in the console tree of Group Policy Object Editor, to specify scripts that are to run when the computer starts up or shuts down. These scripts run as Local System.

- **Scripts (Logon/Logoff)--**You can use this extension, which is located under User Configuration\Windows Settings in the console tree of Group Policy Object Editor, to specify scripts that are to run when the user logs on or logs off the computer. These scripts run as User, not as Administrator.

For more information, see Use Startup, Shutdown, Logon, and Logoff Scripts.

The Windows Server 2003 family includes Windows Script Host, a language-independent scripting host for 32-bit Windows platforms that includes both Visual Basic Scripting Edition (.vbs) and JScript (.js) scripting engines. You can use Windows Script Host to run .vbs and .js scripts directly on the Windows desktop or in the command console, without the need to embed these scripts in an HTML document.

For more information about Windows Script Host, see Windows Script Host overview. To learn about Windows Script Host and the Object Model and to obtain script samples, see the "Windows Script" at the Microsoft Web site.

# Group Policy Software Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Software Installation**
This section covers:

- Group Policy Software Installation overview
- Managing Group Policy Software Installation
- Add or Remove Programs in Control Panel
- File types that you use with Group Policy Software Installation

# Group Policy Software Installation overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Software Installation overview**

With Group Policy Software Installation, a software management feature, you can manage software throughout its life cycle within the organization. Group Policy Software Installation works in conjunction with Group Policy and Active Directory. It is one of three Group Policy Software Installation and Maintenance components, which are described in the following table.

| Component | Role |
|---|---|
| Group Policy Software Installation (an extension of Group Policy Object Editor) | Used by administrators to manage software |
| Windows Installer | Installs software that is packaged in Windows Installer files |
| Add or Remove Programs in Control Panel | Used by users to manage software on their own computers. For more information, see Add or Remove Programs in Control Panel. |

Group Policy Software Installation helps you specify how applications are installed and maintained within your organization. With Group Policy Software Installation, you can manage an application within a Group Policy object, which is in turn associated with a particular Active Directory container--either a site, a domain, or an organizational unit. Applications can be managed in one of two modes: assigned or published.

For tips about using Group Policy Software Installation, see Best practices for Group Policy Software Installation. To begin using Group Policy Software Installation right away, see Open Group Policy Software Installation. For more information, see Group Policy Software Installation.

**Assigning to users**

When you assign an application to a user, the application is advertised to the user the next time the user logs on to a workstation. The application advertisement follows the user, regardless of which physical computer he or she actually uses. The application is installed the first time that the user activates the application on the computer, either by selecting the application on the **Start** menu or by activating a document that is associated with the application. For more information, see Assign an application.

You assign an application when you want everyone to have the application on his or her computer. For example, suppose that you want all users in a marketing department to have Microsoft Excel on their computers. A Group Policy object manages every user in marketing. When you assign Excel within the marketing Group Policy object, Excel is advertised on every marketing user's computer. When an assigned application is advertised, it is not actually installed on the computer. In this case, the application advertisement installs only enough information about Excel to make the Excel shortcuts appear on the **Start** menu and to make the necessary file associations (.xls) appear in the registry.

When these users log on to their computers, Excel appears on their **Start** menu. When they select Excel on the **Start** menu for the first time, Excel installs. A user can also install an advertised application by opening a document that is associated with the application, either by file name extension or by COM-based activation. If a user who has not yet activated Excel from the **Start** menu clicks an Excel spreadsheet to open it, Excel installs and the spreadsheet opens.

A user can delete an assigned application, but the assigned application is advertised again the next time the user logs on. It installs the next time a user selects it on the **Start** menu.

**Assigning to computers**

When you assign an application to the computer, the application is advertised and the installation is performed when it is safe to do so. Typically, this happens when the computer starts up, so that there are no competing processes on the computer. For more information, see Assign an application.

**Publishing to users**

When you publish an application to users, the application does not appear as installed on the users' computers. No shortcuts are visible on the desktop or on the **Start** menu, and no changes are made to the local registry on the users' computers. Instead, published applications store their advertisement attributes in Active Directory. Information, such as the application's name and file associations, is exposed to the users in the Active Directory container. The application is then available for the user to install by using Add or Remove Programs in Control Panel or by clicking a file that is associated with the application,

such as an .xls file for Microsoft Excel. For more information, see Publish an application and Add or Remove Programs in Control Panel.

You publish an application when you want the application to be available to people who are managed by the Group Policy object, in case a user wants the application. With published applications, it is up to each person to decide whether or not to install the published application.

For example, if you publish Microsoft Image Composer to users who are managed by the marketing Group Policy object, and a marketing user wants to install Image Composer, the user can go to Add or Remove Programs in Control Panel and install the application. If users do not install Image Composer by using Add or Remove Programs, and if the .jpg file name extension for the image document is associated with Image Composer, Image Composer can be installed for users when they first open any .jpg document.

**Application assignment scripts**

Application assignment scripts are one of the file types that are often encountered in Group Policy Software Installation. For every published or assigned application in a particular Group Policy object, an application assignment script (.aas file) is generated and stored in that domain's Group Policy object. These script files contain the advertisement information about the application configuration. For more information, see File types that you use with Group Policy Software Installation.

**The user's view of software installation**

At the user's computer, system components--including Winlogon, the shell, object linking and embedding (OLE), the Lightweight Directory Access Protocol (LDAP) client, and the local registry--provide the user's view of software installation. Winlogon is the privileged agent that applies software installation policy. The shell and the OLE are enhanced to be Active Directory-aware and to communicate with Windows Installer to perform setup actions. The LDAP client provides the capability to search and query Active Directory.

By using Add or Remove Programs, users can browse for and install software from Active Directory in a managed environment or from local media (in a nonmanaged environment or, if policy permits, from local media in a managed environment).

# Managing Group Policy Software Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Group Policy Software Installation**
Administrators generally complete the following phases to manage software for their organization: preparation, management, and removal. For tips about using Group Policy Software Installation, see Best practices for Group Policy Software Installation. For more general information, see Group Policy Software Installation overview.

**Preparation**
Before you can use Group Policy Software Installation, you should have a Windows Installer package (.msi) for the program that you want to install. The package is often supplied with the software. If a program does not have a Windows Installer package, you must generate one. Other utilities are available for repackaging a program that you plan to install. You can then use transforms to further customize a package. Transforms, which have the .mst extension, are sometimes called modifications. For more information, see Add or remove modifications for an application package and File types that you use with Group Policy Software Installation.

The next step is to create a shared network directory, called a software distribution point, that contains the packages (.msi files), any transforms (.mst files), and the program files and components. Administration is simpler if packages and program files are kept together, although this is not absolutely necessary. (Packages and transforms do have to be kept together.) As an administrator, you might also benefit from using distributed file systems to help manage these software distribution points.

Finally, verify that users can read from the software distribution point and write to the target of the installation, particularly if the program is written to a network file server.

### Note

- You can still deploy applications that do not have an .msi file by using a .zap setup file. For more information, see File types that you use with Group Policy Software Installation.

**Management**
You can use Group Policy Software Installation to manage programs within your organization. You can assign programs to either computers or users, or you can publish programs to users.

Program assignment and publishing to users have the following software life cycle phases:

- **Evaluation phase--**For evaluation purposes, you might want to have only a few users try out and evaluate the new program. During this phase, the old version of the program is the default version, and the majority of users are still using it. A new user installs the assigned or published version. This is often called a pilot. The pilot is a good time to experiment with Group Policy Software Installation options, such as assigning an application and publishing an application, which is less forceful. You can also try setting the autoinstall option for an application, so that the software is installed when the user attempts to open a file with an extension that you have associated with the application. You can also try the **Basic** or **Maximum** options for the installation user interface, which you can set by using the procedure Set Group Policy Software Installation defaults. These options control the level of visibility of the installation process to the user.

  For more information, see Assign an application, Publish an application, and Set the autoinstall option for an application.

- **Rollout phase--**During rollout, you begin making the new program available to more users. The new version of the program is now the default, and the majority of the users are using the new version. A new user only has the option to install the new program. You can leave certain users unaffected if upgrading them at this time would affect their work.

**Removal**
As an administrator, you must select one of the following options, which are mutually exclusive, when you remove a program:

- **Immediately uninstall--**This option requires that no one currently uses the program. Those users who still have the program on their computers receive an advertisement that removes the program.

- **Just prevent new installations--**This option allows users to continue using the software that they already have.

For more information, see Remove a managed application.

**Deletion from the server**

When you are sure that an old version of a program is no longer needed by the organization, you can delete the Windows Installer package and the program files and components. You might want to archive the package and the program files and components and then delete the files from the software distribution point.

# Add or Remove Programs in Control Panel

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Add or Remove Programs in Control Panel**
By using Add or Remove Programs in Control Panel, you can perform a number of tasks--primarily, the task of installing an application from local media, such as a CD-ROM or floppy disk; a defined Active Directory location, for an assigned and published application on a corporate network; or the Internet. You can also use Add or Remove Programs to remove or modify an existing application or to repair a damaged application.

Add or Remove Programs also offers the following features:

- Provides users with a single user interface for managing and maintaining their programs.
- Supports the installation of programs when a user has the program package (on local media), as well as installation from either a corporate or an Internet environment.
- Exposes the Windows Installer component while maintaining support for other installation technologies.

For more information, see Add or Remove Programs overview.

# File types that you use with Group Policy Software Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File types that you use with Group Policy Software Installation**

You can use the file types listed in the following table with Group Policy Software Installation.

For more information about specific Group Policy Software Installation procedures, see Use Group Policy Software Installation. For general information about Group Policy Software Installation, see Group Policy Software Installation.

| File type | File name extension | Description |
|---|---|---|
| Windows Installer packages | .msi | Your software vendor typically provides these files to facilitate installation of a specific application. Keep these files with any other necessary files at the software distribution point for the software that you are managing. |
| Transforms | .mst | Also called modifications, these files customize the installation of a Windows Installer package at the time of assignment or publication. For example, they might specify a subset of a suite of applications. |
| Patches | .msp | Bug fixes, service packs, and similar files can be distributed in this form.<br><br>Patches should not be used for major changes, and their effects are limited in the following ways:<br><br>• They cannot remove components or features.<br><br>• They cannot change product codes.<br><br>• They cannot remove or change the names of shortcuts, files, or registry keys. |
| .zap files | .zap | These files are created with a text editor, such as Notepad. They can only be published (not assigned), and they specify a Setup executable program that appears in Add or Remove Programs in Control Panel for users with administrative rights on the local computer.<br><br>**Note**<br>The correct use of quotation marks in .zap file syntax can be difficult. The following rules apply:<br><br>• The path and name of the Setup executable program must always be wrapped in quotation marks.<br><br>• If there are no command-line arguments, they must be wrapped in two sets of quotation marks.<br><br>Example syntax using an absolute path:<br><br>**SetupCommand=""***ServerName******ShareName******FolderName***\setup.exe""**<br><br>**SetupCommand="***ServerName******ShareName******FolderName***\setup.exe" /***argument*<br><br>Example syntax using a relative path:<br><br>**SetupCommand=""setup.exe""**<br><br>**SetupCommand="setup.exe"/***argument* |

| | | |
|---|---|---|
| Application assignment scripts | .aas | These files contain instructions that are associated with the assignment or publication of a package. |

# Group Policy Precedence

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy precedence**
You can apply Active Directory-based Group Policy by linking Group Policy objects to Active Directory objects (sites, domains, or organizational units). Group Policy settings under User Configuration apply to users, and Group Policy settings under Computer Configuration apply to computers. The locations of the user and computer accounts in Active Directory determine what Group Policy objects are used to supply policy. Note that although some settings are user interface settings--for example, the background bitmap or the ability to use the **Run** command on the **Start** menu--they can also be applied to computers.

Group Policy is cumulative. Child directory service containers inherit policy from parent containers, and Group Policy processing occurs in the following order: local, site, domain, and then successive organizational units, from highest organizational unit (farthest from the user or computer account) to lowest organizational unit (actually containing the user or computer account). This means that if you link a Group Policy object to a high-level parent container (domain or organizational unit), this Group Policy object applies to all containers beneath the parent container, that is, to the user and computer objects in each container. However, if you explicitly enable or disable some Group Policy settings in another Group Policy object that is linked to a child container, the child container's Group Policy object settings override the parent container's Group Policy object settings.

You can enforce Group Policy on child directory containers by setting the **No Override** option on a Group Policy object link. You can also prevent inheritance of Group Policy from parent directory containers by using the **Block Policy inheritance** option on a domain or organizational unit.

**Important**

- **No Override** and **Block Policy inheritance** are advanced options, and they are not recommended for casual use, because they change the default behavior of policy inheritance, as described in the previous paragraph, and this can complicate troubleshooting.

**Note**

- In the event of a conflict, Group Policy settings take precedence over user profile settings.

This section covers:

- Order of events when starting up and logging on
- Order of processing settings
- The role of applications in Group Policy precedence

# Order of events when starting up and logging on

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Order of events when starting up and logging on**

The following sequence shows the order in which computer policy and user policy are applied when a computer starts and a user logs on:

1. Network starts. Remote procedure Call System Service (RPCSS) and Multiple Universal Naming Convention provider (MUP) start.

2. An ordered list of Group Policy objects is obtained for the computer. The list might depend on these factors:

   - Whether the computer is part of a domain and therefore subject to Group Policy through Active Directory.

   - The location of the computer in Active Directory.

   - Whether the list of Group Policy objects has changed. If the list of Group Policy objects has not changed, no processing is done. You can use a policy setting to change this behavior.

3. Computer policy is applied. These are the settings under Computer Configuration from the gathered list. This is synchronous by default and in the following order: local, site, domain, organizational unit, child organizational unit, and so on. No user interface appears while computer policies are processed.

   For details about the order in which settings are processed when user or computer policy is applied, see Order of processing settings.

4. Startup scripts run. This is hidden and synchronous by default; each script must complete or time out before the next one starts. The default time-out is 600 seconds. You can use several policy settings to modify this behavior.

5. The user presses CTRL-ALT-DEL to log on.

6. After the user is validated, the user profile loads; it is governed by the policy settings that are in effect.

7. An ordered list of Group Policy objects is obtained for the user. The list might depend on these factors:

   - Whether the user is part of a domain and therefore subject to Group Policy through Active Directory.

   - Whether loopback is enabled, and the state (**Merge** or **Replace**) of the loopback policy setting.

   - The location of the user in Active Directory.

   - Whether the list of Group Policy objects has changed. If the list of Group Policy objects has not changed, no processing is done. You can use a policy setting to change this behavior.

8. User policy is applied. These are the settings under User Configuration from the gathered list. This is synchronous by default and in the following order: local, site, domain, organizational unit, child organizational unit, and so on. No user interface appears while user policies are processed.

   For details about the order in which settings are processed when user or computer policy is applied, see Order of processing settings.

9. Logon scripts run. Unlike Windows NT 4.0 scripts, Group Policy-based logon scripts are hidden and asynchronous by

default. The user object script runs last in a normal window.

10. The operating system user interface that is prescribed by Group Policy appears.

**Important**

- Three special cases deserve consideration during migration:

    - If the computer account object is in a Windows NT 4.0 domain and the user account object is in Active Directory, only computer (not user) System Policy is processed when the user logs on. Then, user (not computer) Group Policy is processed.

    - If the computer account object is in Active Directory and the user account object is in a Windows NT 4.0 domain, computer (not user) Group Policy is processed during computer startup. When the user logs on, user (not computer) System Policy is processed.

    - If the computer and user accounts of the computer running Windows XP Professional or a Windows Server 2003 operating system are members of a Windows NT 4.0 domain, only System Policy (not Group Policy) for the computer and user is applied when the user logs on.

**Note**

- Several of these events can be modified. You can set policies to:

    - Reverse the synchronized or asynchronized defaults for running scripts and applying policy.

    - Specify when scripts time out. By default, scripts time out after 600 seconds.

    - Change whether scripts are run hidden, minimized, or in a normal window.

# Order of processing settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Order of processing settings**

This section provides details about the order in which Group Policy settings for users and computers are processed. For information about where the processing of policy settings fits into the framework of computer startup and user logon, see steps 3 and 8 in Order of events when starting up and logging on.

Group Policy settings are processed in the following order:

1. **Local Group Policy object--**Each computer has exactly one Group Policy object that is stored locally.

2. **Site--**Any Group Policy objects that have been linked to the site are processed next. Processing is synchronous and in an order that is specified by the administrator.

3. **Domain--**Processing of multiple domain-linked Group Policy objects is synchronous and in an order specified by the administrator.

4. **Organizational units--**Group Policy objects that are linked to the organizational unit that is highest in the Active Directory hierarchy are processed first, then Group Policy objects that are linked to its child organizational unit, and so on. Finally, the Group Policy objects that are linked to the organizational unit that contains the user or computer are processed.

   At the level of each organizational unit in the Active Directory hierarchy, one, many, or no Group Policy objects can be linked. If several Group Policy objects are linked to an organizational unit, their processing is synchronous and in an order that is specified by the administrator.

This order means that the local Group Policy object is processed first, and Group Policy objects that are linked to the organizational unit of which the computer or user is a direct member are processed last, which overwrites the earlier Group Policy objects.

**Exceptions to the default order**

The default order for processing settings is subject to the following exceptions:

- Any Group Policy object that is linked to a site, domain, or organizational unit (not a local Group Policy object) can be set to **No Override** with respect to that site, domain, or organizational unit, so that none of its policy settings can be overridden. When more than one Group Policy object has been set to **No Override**, the one that is highest in the Active Directory hierarchy (or higher in the hierarchy that is specified by the administrator at each fixed level in Active Directory) takes precedence.

  Note that **No Override** and **Disabled** are settings on Group Policy objects links, not on the Group Policy objects. A Group Policy object can be linked several times to the same organizational unit, and **No Override** and **Disabled** can be configured independently on each of the links. (Although multiple links from one Group Policy object to a single organizational unit are seldom useful, this capability illustrates the flexibility of the Group Policy infrastructure.)

  For information about how to set links as **No Override** and **Disabled**, see Prevent a Group Policy object from being overridden and Disable a Group Policy object link.

- At any site, domain, or organizational unit, you can mark Group Policy inheritance selectively as **Block Policy inheritance**. Group Policy object links that are set to **No Override** are always applied, however, and they cannot be blocked.

  The **Block Policy inheritance** setting is applied directly to the site, domain, or organizational unit. It is not applied to Group Policy objects, nor is it applied to Group Policy object links. **Block Policy inheritance** deflects all Group Policy settings that would reach the site, domain, or organizational unit from above (by way of links to parents in the Active Directory hierarchy), no matter what Group Policy objects those settings originate from. However, **Block Policy**

**inheritance** does not deflect Group Policy settings from Group Policy objects that are linked directly to the site, domain, or organizational unit that has **Block Policy inheritance** enabled.

- A computer that is a member of a workgroup processes only the local Group Policy object.

- Loopback is an advanced Group Policy setting that is useful on computers in certain closely managed environments, such as kiosks, laboratories, classrooms, and reception areas. For a description of loopback, click the **Explain** tab after you double-click **User Group Policy loopback processing mode** in the details pane of the Microsoft Management Console (MMC), which is located under Computer Configuration\Administrative Templates\System\Group Policy.

  Loopback provides alternatives to the default method of obtaining the ordered list of Group Policy objects whose User Configuration settings affect a user. By default, a user's settings come from a Group Policy object list that depends on the user's location in Active Directory. The ordered list goes from site-linked to domain-linked to organizational unit-linked Group Policy objects, with inheritance determined by the location of the user in Active Directory and in an order that is specified by the administrator at each level.

  Loopback can be set to **Not Configured**, **Enabled**, or **Disabled**, as can any other Group Policy setting. In the **Enabled** state, loopback can be set to **Merge** or **Replace**.

  - **Loopback with Replace--**In the case of **Loopback with Replace**, the Group Policy object list for the user is replaced in its entirety by the Group Policy object list that is already obtained for the computer at computer startup (during step 2 in Order of events when starting up and logging on). The User Configuration settings from this list are applied to the user.

  - **Loopback with Merge--**In the case of **Loopback with Merge**, the Group Policy object list is a concatenation. The default step 2 list for computers in Order of events when starting up and logging on is appended to the default step 7 list for users, and the user gets the User Configuration settings in the concatenated list. Note that the Group Policy object list that is obtained for the computer is applied later, and therefore it has precedence if it conflicts with settings in the user's list.

# The role of applications in Group Policy precedence

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The role of applications in Group Policy precedence**

The Group Policy infrastructure cannot force applications to use Group Policy. For example, applications that have always looked for registry entries in a particular place that is outside the approved registry trees that are reserved for Group Policy will continue to look there. Group Policy does not copy registry settings from the Group Policy areas to the Windows NT 4.0 System Policy areas, nor does it copy Windows NT 4.0 System Policy settings into the reserved Group Policy areas. For more information, see Windows NT 4.0 system policies and Migration Issues.

Guidelines for writing applications that take advantage of Group Policy are available in the Group Policy Platform SDK that is used by developers. For more information, see "Group Policy" at the Microsoft Web site. Administrators, who need to anticipate the behavior of applications that they do not author, should be aware of the steps taken by properly written applications to obtain the registry data that the applications need.

**Initial state of the registry**

After the computer starts and the user logs on, the Group Policy registry areas are rewritten to hold the cumulative Group Policy settings that are in effect from the local Group Policy object and from Active Directory. If Windows NT 4.0 System Policy is enabled (which is not recommended for client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems, or Windows 2000 clients), other registry areas might also change when the user logs on.

**Order of events**

1. The user starts the application.

2. A typical user-oriented application looks for registry data in the Group Policy reserved area **HKEY_LOCAL_MACHINE\Software\Policies**. If the application finds what it needs, it looks no further.

   However, an application that changes or replaces features at the level of the operating system (such as the **Run** command), and which therefore affects the behavior of applications other than itself, might look in the following Group Policy reserved area: **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies**. If the application finds what it needs, it looks no further.

3. The application looks for **HKEY_LOCAL_MACHINE** registry data outside the Group Policy reserved area, and if it finds what it needs, it looks no further.

4. The application looks for **HKEY_CURRENT_USER** registry data outside the Group Policy reserved area, and if it finds what it needs, it looks no further.

5. The application uses .ini files (not recommended) or default settings.

Newly authored, native Windows XP applications might perform only the first three steps. Applications that use Group Policy under Windows XP, but which remain compatible with earlier versions of Windows, continue to step 4 when they run on older operating systems. Applications that predate Windows 2000 never perform steps 2 and 3. Instead, they go directly to step 4.

The order of this sequence of events implies the following:

- For properly written applications that run on client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems, Group Policy takes precedence over Windows NT 4.0 System Policy. For more information, see Policy inheritance, Order of events when starting up and logging on, and Order of processing settings.

- Older versions of applications, which are not aware of Active Directory and Group Policy, continue to function in Windows XP Professional, Windows XP 64-bit Edition (Itanium), and the Windows Server 2003 family as they did under Windows NT 4.0.

- **HKEY_LOCAL_MACHINE** settings take precedence over **HKEY_CURRENT_USER** settings, as they did under Windows NT 4.0.

# Using Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Group Policy**
This section covers:

- Ways to open Group Policy Object Editor
- Local Group Policy
- Administrative Templates included with this version of Windows
- Advanced methods of extending Group Policy
- Group Policy Permissions
- Group Policy and Network Infrastructure
- Migration Issues
- Managing Group Policy from the command line

# Ways to open Group Policy Object Editor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ways to open Group Policy Object Editor**

You can open Group Policy Object Editor in several ways, depending on the action that you want to perform and the object to which you want to apply Group Policy. The following table lists the ways in which you can open Group Policy Object Editor.

| To apply Group Policy to | Do this |
|---|---|
| Local computer | In Group Policy Object Editor, edit the local Group Policy object, as described in Edit the local Group Policy object. |
| Another computer | Open the local Group Policy object that is stored on the network computer, as described in Open Group Policy as an MMC snap-in, and then browse to the network computer. You must have administrative rights on the network computer. |
| Site | Open Group Policy Object Editor, as described in Open Group Policy from Active Directory Sites and Services, and then link a Group Policy object to the intended site. |
| Domain | Open Group Policy Object Editor as described in Open Group Policy from Active Directory Users and Computers, and then link a Group Policy object to the intended domain. |
| Organizational unit | Open Group Policy Object Editor as described in Open Group Policy from Active Directory Users and Computers, and then link a Group Policy object to the intended organizational unit. You can also link a Group Policy object to an organizational unit that is higher in the Active Directory hierarchy, so that the organizational unit can inherit Group Policy settings. |
| Existing Group Policy object or set of Group Policy objects | Create and save your own custom Microsoft Management Console (MMC). For more information, see Microsoft Management Console. |

You can also open Group Policy Object Editor by clicking **Start**, clicking **Run**, and then typing a command. Or, you can open Group Policy Object Editor at the command prompt in an MS-DOS window. The following table contains some examples of how you might open Group Policy Object Editor from the command line in different situations.

| To edit | At a command line, use this command-line syntax |
|---|---|
| The local Group Policy object on *ThisComputerName* | **gpedit.msc /gpcomputer:"** *ThisComputerName* **"** |
| The local Group Policy object on *ThisComputerName* (specified in the DNS style) | **gpedit.msc /gpcomputer:"** *ThisComputerName.domain.com* **"** |
| Group Policy objects stored in Active Directory | **gpedit.msc /gpobject:"LDAP://CN={31B2F340-016D-11D2-945F-00C04FB984F9},CN=Policies,CN=System,DC=WingTipToys,DC=com"**<br><br>**Note**<br><br>• The GUID for the Group Policy object above is a fictitious example. |

For information about other Group Policy command-line tools, see Managing Group Policy from the command line.

# Local Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local Group Policy**

Each computer running Windows 2000, Windows XP Professional, Windows XP 64-bit Edition (Itanium), or a Windows Server 2003 operating system has exactly one local Group Policy object. In these objects, Group Policy settings are stored on individual computers, whether or not they are part of an Active Directory environment or a networked environment. For more information, see Security Settings extension to Group Policy, Folder Redirection, and Group Policy Software Installation overview.

Local Group Policy objects contain fewer settings than nonlocal Group Policy objects, particularly under Security Settings. Local Group Policy objects do not support Folder Redirection or Group Policy Software Installation.

Because its settings can be overwritten by Group Policy objects that are associated with sites, domains, and organizational units, the local Group Policy object is the least influential object in an Active Directory environment. In a non-networked environment (or in a networked environment that does not have a domain controller), the local Group Policy object's settings are more important, because they are not overwritten by other Group Policy objects.

To edit the local Group Policy object that is stored on your local computer, use the Group Policy Object Editor snap-in. For more information, see Edit the local Group Policy object.

If you want to edit the local Group Policy object that is stored on another computer on the network, open Group Policy as a stand-alone Microsoft Management Console (MMC) snap-in and browse to the Group Policy object that you want. For more information, see Open Group Policy as an MMC snap-in.

The local Group Policy object resides in *systemroot*\System32\GroupPolicy.

Computers that run Windows NT 4.0 or earlier do not have a local Group Policy object, and they do not recognize nonlocal Group Policy. For more information, see Migration Issues.

# Administrative Templates included with this version of Windows

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administrative Templates included with this version of Windows**

Five Administrative Templates (.adm) files are included with this version of Windows. The following table contains a brief description of each of the .adm files. For more information about .adm files, see Administrative Templates.

| Administrative Template | Description | Help for Administrative Templates policy settings |
|---|---|---|
| System.adm | System settings | System settings |
| Inetres.adm | Internet Explorer settings | Internet Explorer settings |
| Wmplayer.adm | Windows Media Player settings. This tool is not available on the Itanium-based versions of the Windows operating systems. This tool is not available on the x64-based versions of the Windows operating systems. | Windows Media Player settings |
| Conf.adm | NetMeeting settings. This tool is not available on the Itanium-based versions of the Windows operating systems. This tool is not available on the x64-based versions of the Windows operating systems. | NetMeeting settings |
| Wuau.adm | Windows Update settings. | Windows Update settings |

When you add an .adm file to Group Policy, as described in Add or remove an Administrative Template (.adm file), by default only the settings that are contained in the genuine Group Policy trees (the trees that correspond to the reserved Group Policy registry areas) are visible in the console.

If you want to change this default behavior and make not only Group Policy settings but also the older System Policy settings visible, see Filter the view provided by Administrative Templates. Clear the **Only show policy settings that can be fully managed** check box, as described in that procedure. Your ability to clear this check box in the **Filtering** dialog box is subject to the **Enforce Show Policies Only** Group Policy setting.

When you make preference settings appear in Group Policy Object Editor, the icons for these settings are colored red. The icons for true Group Policy settings are colored blue.

**Important**

- It is highly recommended that you do not use System Policy-type .adm files when you manage client computers running Windows 2000 Professional, Windows XP Professional, Windows XP 64-bit Edition (Itanium), or Windows Server 2003 operating systems because this can lead to undesirably persistent registry settings. For more information, see Migration Issues.

# Advanced methods of extending Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Advanced methods of extending Group Policy**
The following table lists advanced methods that you can use to extend Group Policy functionality.

| Method | Advantages | Disadvantages |
|---|---|---|
| For administrators only: Write an .adm file | • Easier than writing an extension to Group Policy.<br>• Everything you need is included in the Windows Server 2003 family.<br>• The .adm files are text files, which you can edit in Notepad.<br>• You can write from scratch or modify existing templates to suit your purpose. | • This option is only suitable for registry-based policies.<br>• The user interface is limited. Finding where to set individual policies might be difficult. |
| For software developers only: Write an extension to Group Policy | • Richer user interface.<br>• The full Group Policy application programming interface (API) is available to the developer.<br>• Not limited to registry-based policies. | • Because writing an extension takes considerable development time, this option is not worthwhile unless it will be used widely or often.<br>• Requires software development kit (SDK) material that is not included with the Windows Server 2003 family. |

**Creating custom .adm files**
You can consider creating custom .adm files if the supplied template, System.adm, is inadequate. However, it is recommended that you use the supplied template.

Administrative Templates propagate registry settings to a large number of computers without requiring you to have detailed knowledge of the registry. In the Windows Server 2003 family, Group Policy is used to set registry-based policies. In Windows NT 4.0, the System Policy Editor (Poledit.exe) was used to set System Policy.

For more information, see "Implementing Registry-based Policy" at the Microsoft Web site.

**Creating Microsoft Management Console extensions**
Developers can create extensions to Microsoft Management Console (MMC) snap-ins to provide program-specific user interfaces for setting Group Policy. For more information, see "Group Policy" at the Microsoft Web site and "MSDN Library" at the Microsoft Web site.

# Group Policy Permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy permissions**

If you want to delegate policy-related tasks to other people, including nonadministrators, see Delegate Control of Group Policy.

This section covers:

- Default permissions
- Policy for Microsoft Management Console

# Default permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Default permissions**

The following table describes the default permissions on Group Policy objects.

| Security group | Default settings |
|---|---|
| Authenticated users | Read, Apply Group Policy (AGP) |
| Local system | Full Control (includes AGP) |
| Domain administrators | Read, Write, Create Child, Delete Child, AGP |
| Administrators | Read, Write, Create Child, AGP |
| Group Policy Creator Owners | Read, Write, AGP |

By default, the Group Policy object Default Domain Policy cannot be deleted by any administrator. The purpose of this restriction is to prevent the accidental deletion of this Group Policy object, which contains important and required settings for the domain. If Default Domain Policy must be deleted for some reason, the Delete permission must be given explicitly to the intended group. This is an advanced access control entry (ACE) on the Group Policy object. For more information, see Set permissions for Group Policy Software Installation.

# Policy for Microsoft Management Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policy for Microsoft Management Console**

There are several Group Policy settings in one of the default Administrative Templates files, System.adm, that govern the use of the Microsoft Management Console (MMC). You can delegate administrative rights by saving previously configured consoles. For more information, see Delegate Control of Group Policy. You can find MMC policy settings in Group Policy Object Editor. You can then double-click each item for which you want to set an administrative Group Policy setting that governs the use of that item in MMC.

For more information, see Microsoft Management Console, Group Policy Permissions, and Configure Group Policy.

# Group Policy and Network Infrastructure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy and network infrastructure**
This section covers:

- Group Policy features that are supported across forests
- Group Policy in multiplatform networks
- Group Policy in replicated environments
- Group Policy over slow links
- Group Policy on sites

# Group Policy features that are supported across forests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy features that are supported across forests**

The Windows Server 2003 family introduces a powerful set of new features that enable you to authenticate and authorize access to resources from separate, networked forests. With full trusts established between forests, you can manage Group Policy throughout Active Directory regardless of the forest, which provides greater flexibility especially in large organizations. For more information about accessing resources across forests in Active Directory, see Accessing resources across forests. For more information on forest trusts, see Forest trusts.

Following is a list of the Group Policy features that are supported across forests:

- **Interactive logon**

  The following table lists the interactive logon tasks that are supported across forests.

| Task | Support across the domain? | Support across domains in the same forest? | Support interactive logon across forests? |
|---|---|---|---|
| Apply policy to a user object | Yes | Yes | Yes |
| Apply policy to a computer object | Yes | Yes | Yes |
| Support Group Policy loopback processing | Yes | Yes | Yes |

  For example, User1 is located in forestA and computer1 is located in forestB. User1 goes to forestB and logs into computer1. In this scenario, computer policy is applied to computer1 from forestB, user policy is applied to user1 on computer1 from forestA. Site policy might be applied depending on the subnet to which computer1 belongs.

  For interactive logon across forests with loopback processing, Group Policy objects with loopback enabled for merge or replace are applied across forests under the same logon model as specified for interactive logon. For more information about loopback, see Order of processing settings.

- **Network logon**

  The following Group Policy extensions are supported for network logon across forests:

  - Software distribution points located in another forest
  - Logon scripts on a shared network directory in another forest.
  - Roaming user profiles stored on a shared network directory in another forest. For more information, see Using roaming user profiles.
  - Redirected folders stored on a file share in another forest. For more information, see Folder Redirection.

- **Delegation**

  Delegation across forests is supported for managing Group Policy links. Other tasks such as creating, deleting, or modifying Group Policy objects across forests is not supported.

# Group Policy in multiplatform networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy in multiplatform networks**
Many networks include computers with different operating systems and versions. To allow administrators the greatest control, Group Policy and its extensions can be applied to computers running Windows 2000 Professional, Windows 2000 Server, Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems.

For client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems that are operating in networks with Windows 2000 Active Directory, the following features and limitations apply:

- **Fast Logon Optimization**

  The Windows Server 2003 family provides a Fast Logon Optimization feature. This results in the asynchronous application of policy when the computer starts up and when the user logs on. This feature is supported for client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems that are operating in any environment. This feature is not available to Windows 2000 clients. For more information, see Fast Logon Optimization.

- **Administrative Templates**

  The Windows Server 2003 family supports all registry policy settings available in Windows 2000. In addition, the Windows Server 2003 family has 212 new registry settings. These new policy settings only apply to client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems. If specified for Windows 2000 clients, the settings are ignored. You can deploy new .adm files provided with Windows XP Professional and Windows Server 2003 operating systems from Windows 2000 Active Directory. For more information, see Administrative Templates.

- **Folder Redirection**

  The Windows Server 2003 family contains a new Folder Redirection with a simplified user interface. You can configure folder redirection for client computers running Windows 2000, Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems. With client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems, you can redirect folders to a HomeDir or redirect a folder back to a local profile. These options cannot be applied to Windows 2000 clients. For more information, see Folder Redirection.

- **Roaming User profiles**

  This feature contains new registry policy settings to prevent Roaming User profile changes from propagating to the server, allow only local user profiles, and add Administrators security group to Roaming User profiles. These features can only be applied to client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems. Users can roam between client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems and Windows 2000 clients, however, only policies that can be applied to each platform will be used. For more information, see Using roaming user profiles.

- **Security**

  With a software restriction policy, you can identify software on client computers and prevent it from running. Software can be identified by file name, path, URL, and MD5 Hash. This policy must be configured in a Group Policy object from the security settings snap-in from a client computer running Windows XP Professional, Windows XP 64-bit Edition (Itanium), or Windows Server 2003 operating systems. This feature is not supported for Windows 2000 clients. For more information, see Security Settings extension to Group Policy.

- **Group Policy Software Installation**

The Windows Server 2003 family contains numerous enhancements to the Group Policy Software Installation user interface. In addition, this version of Windows also support the deployment of 32-bit MSI applications to client computers running Windows XP 64-bit Edition (Itanium), 64-bit versions of Windows Server 2003, Enterprise Edition, and 64-bit versions of Windows Server 2003, Datacenter Edition. You cannot install 64-bit applications on 32-bit clients. Although Group Policy provides a robust solution for software installation and distribution, Microsoft Systems Management Server is the preferred solution for enterprise software planning, distribution, and management. For more information, see Group Policy Software Installation.

- **WMI filtering**

  With WMI Filters, you can specify a WMI-based query to filter the application of a Group Policy object. WMI Filters are written using WMI query language (WQL). WMI Filters only apply to client computers running Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems. WMI Filters associated with a Group Policy object will be ignored by Windows 2000 clients.

# Group Policy in replicated environments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy in replicated environments**

In a domain that contains more than one domain controller, Active Directory information takes time to propagate from one domain controller to another. This topic describes the replication mechanism as it relates to Group Policy.

By default, Group Policy objects are created or edited only on the domain controller that is holding the primary domain controller (PDC) emulator operations master token. This token moves from one domain controller to another over time, as Active Directory information is replicated to keep the domain controllers synchronized.

An organizational unit must be replicated to the domain controller that is holding the token before the Group Policy object that is created there can be linked to it, which allows the Group Policy settings to be applied. If you use Active Directory Users and Computers, you can create an organizational unit on any domain controller. These considerations are more significant if the intradomain links are slow. For more information, see Group Policy over slow links and Group Policy on sites.

**Options governing selection of a domain controller**

When you click the name of the Group Policy object in the console tree of Group Policy Object Editor, the **View** menu contains a command called **DC Options**. This command opens the **Options for domain controller selection** dialog box, where you can specify a domain controller to use for editing Group Policy. The options in this dialog box are as follows:

- **The one with the Operations Master token for the PDC emulator--**This is the default and preferred option, and it is the best option from the standpoint of data safety.

- **The one used by the Active Directory snap-ins--**This option uses the same domain controller as the utility from which the Group Policy Object Editor was invoked, if Group Policy was started this way.

- **Use any available domain controller--**This option allows Group Policy Object Editor to choose any available domain controller. This is the least safe option, because different administrators could theoretically edit a Group Policy object simultaneously, with an indeterminate outcome. On the other hand, when you use this option, it is likely that a domain controller in the local site will be selected. If only one administrator can administer Group Policy on a large domain with several sites, the performance gain might be worthwhile.

**Domain controller selection set through Group Policy**

In addition to the **DC Options** command, there is also a Group Policy setting for domain controller selection. This is part of the System.adm Administrative Template that is loaded into Group Policy Object Editor by default. You can find this Group Policy setting in Group Policy Object Editor under User Configuration\Administrative Templates\System\Group Policy. You can double-click **Group Policy domain controller selection** in the details pane, and then click the appropriate setting in the list.

# Group Policy over slow links

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy over slow links**

If the computer is a member of the domain that the Routing and Remote Access server belongs to or is trusted to, you can apply Group Policy remotely. This is true whether you are logging on through Routing and Remote Access or logging on with cached credentials, and then establishing a Routing and Remote Access connection.

Group Policy is not applied to computers that are members of a workgroup or of a foreign domain (a domain that does not contain the computer and that is not in a trust relationship with the computer's domain). Although the connection may still be made, access to domain resources may be adversely affected because of mismatched Internet Protocol security (IPSec).

By default, registry-based policies are always applied, and they cannot be turned off. Security settings are also applied by default, but all other settings are not applied. For all but the registry settings, the default behavior can be switched to apply or not apply.

A remote access connection is not necessarily a slow link, nor is a local area network (LAN) necessarily a fast link. By default, the fast or slow status of a link is based on a test ping to the server. If it takes less than 2000 milliseconds (two seconds), it is considered a fast link. If it takes more time, it is considered a slow link. You can set this value by using the Group Policy setting **Slow network connection time-out for user profiles**, which is located under Computer Configuration\Administrative Templates\System\User Profiles.

Several other Group Policy settings that are related to slow links appear in this folder as well. Also, check the Group Policy setting **Group Policy slow link detection**, which is located under Computer Configuration\Administrative Templates\System\Group Policy.

For more information, see Group Policy in replicated environments and Group Policy on sites. For a more information about slow link detection, see "Group Policy" at the Microsoft Web site.

# Group Policy on sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy on sites**

Group Policy objects that are applied to Active Directory site objects affect all computers in the site. Directory information is replicated and available among all the domain controllers in the site and to any domain controllers in sites for which a site link has been established. Therefore, any Group Policy object that is linked to a site is applied to all computers in that site, without regard to which domain (in the forest) contains the computers.

This allows multiple domains within a forest to receive the same Group Policy object (and included policies), although the Group Policy object exists only as a stored entity on a single domain and it must be read from that domain when the affected clients read their site-linked Group Policy.

If child domains are set up across wide area network (WAN) boundaries, the site setup should take this into account. If it does not, the computers in a child domain access a site-linked Group Policy object across a WAN link. This increases the processing time for Group Policy.

For more information, see Group Policy in replicated environments and Group Policy over slow links.

**Note**

- Group Policy is not usually applied to sites. It is usually applied to domains and organizational units. For more information, see Ways to open Group Policy Object Editor.

# Migration Issues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Migration issues**
This section covers:

- Client operating systems
- The role of Administrative Templates
- The role of System Policy Editor
- Windows NT 4.0 system policies

# Client operating systems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Client operating systems**

This section describes several operating system issues regarding clients for policy.

- Windows Server 2003 family

  Group Policy fully supports client computers that run Windows Server 2003 operating systems. These servers can either be ordinary member servers of Active Directory or domain controllers. Group Policy fully supports both types of computers.

- Windows 2000 Server

  Computers that run Windows 2000 Server can either be ordinary member servers of Active Directory or domain controllers. Group Policy fully supports both types of computers.

- Windows XP Professional and Windows 2000 Professional

  Group Policy fully supports client computers that run Windows XP Professional, Windows XP 64-bit Edition (Itanium) and Windows 2000 Professional.

- Windows 95, Windows 98, and Windows Millennium Edition

  It is recommended that Windows 95, Windows 98 and Windows Millennium Edition clients run System Policy Editor (Poledit.exe) on the local computer. This ensures that the Config.pol file that is created is compatible with those operating systems. Copy the resultant Config.pol file to the SysVol folder of the domain controller. Group Policy does not apply.

- Windows NT 4.0 Workstation and Windows NT 4.0 Server

  These clients continue to be fully supported by Windows NT 4.0 System Policy. You can use System Policy Editor (Poledit.exe) to write Windows NT 4.0 .adm files. System Policy Editor can be obtained from a Windows NT 4.0 CD. Windows NT 4.0 does not use Active Directory. Because computers that run Windows NT 4.0 do not have local Group Policy objects, Group Policy does not apply. For more information, see Local Group Policy.

- Windows XP Home Edition, Windows NT 3.51, Windows 3.1, and MS-DOS

  Group Policy does not apply to these operating systems.

# The role of Administrative Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The role of Administrative Templates**
In the Windows Server 2003 family, Administrative Templates use the .adm file name extension, as they did in Windows NT 4.0. However, their role is slightly different.

**Role in earlier versions of Windows**
In earlier versions of Windows, Administrative Templates were ANSI-encoded text files. They created a namespace within System Policy Editor for convenient editing of the registry. Administrative Templates provided a friendlier user interface than the Registry Editor (Regedit.exe). They also added a degree of safety by exposing only the registry keys that are explicitly mentioned in the .adm file.

**Role in the Windows Server 2003 family**
The Windows Server 2003 family of products include several .adm files, which are listed in Administrative Templates included with this version of Windows. You can also write additional .adm files. The new version of the .adm language is a superset of the previous version--older templates can create a user interface in Group Policy, but new templates cannot create a user interface in System Policy Editor. For more information on the new .adm language, see "Implementing Registry-based Policy" at the Microsoft Web site.

The Windows Server 2003 family supports Unicode-based .adm files.

For information about using Windows NT 4.0 .adm files in a Windows 2000, Windows XP Professional, Windows XP 64-bit Edition (Itanium), or Windows Server 2003 family environment, see Windows NT 4.0 system policies.

# The role of System Policy Editor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The role of System Policy Editor**

Although System Policy Editor (Poledit.exe) is mostly replaced by Group Policy, it is still useful under some circumstances. If you are upgrading from earlier versions of Windows, see Migration Issues.

- **Management of computers that run Windows 95 or Windows 98--**System Policy Editor must be run locally on computers that run Windows 98 or Windows 95 to create Config.pol files that are compatible with the local operating system.

- **Management of computers that run Windows NT 4.0 Workstation and Windows NT 4.0 Server--**These computers also need their own .pol file: NTconfig.pol.

- **Management of stand-alone computers--**A Windows 2000 computer that is not joined to any domain is not subject to nonlocal Group Policy through Active Directory. The only Group Policy that applies to such a computer is local Group Policy, which contains settings for all users of that computer. For more information on local Group Policy, see Local Group Policy. To provide settings for multiple users, use System Policy Editor to write a Registry.pol file. Only the Windows 2000 version of System Policy Editor is compatible with Windows 2000.

Although earlier versions of System Policy Editor worked only with ASCII-encoded .adm files, the Windows 2000 version also supports Unicode-encoded .adm files.

For client computers running Windows 2000, Windows XP Professional, Windows XP 64-bit Edition (Itanium), and Windows Server 2003 operating systems, use Group Policy instead. For tips about using Group Policy, see Group Policy Best Practices (pre-GPMC).

# Windows NT 4.0 system policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows NT 4.0 system policies**
If a Windows NT 4.0 client computer is upgraded to Windows XP Professional, Windows XP 64-bit Edition (Itanium), or a Windows Server 2003 operating system, the computer receives only Computer Configuration Group Policy, not Windows NT 4.0 computer System Policy. The policy that is received by the user who logs on is User Configuration Group Policy, if the user account is in Active Directory. If the user account is managed by a Windows NT 4.0 domain controller, the user receives Windows NT 4.0 user System Policy. For more information, see Computer Configuration and User Configuration.

Policies that are included with Windows 2000, Windows XP Professional, Windows XP 64-bit Edition (Itanium), and the Windows Server 2003 family set registry keys and values only in one of these four reserved trees:

- **HKEY_LOCAL_MACHINE\Software\Policies**
- **HKEY_CURRENT_USER\Software\Policies**
- **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies**
- **HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Policies**

The first two trees are preferred. All four trees are secure, and they cannot be modified by a user who is not an administrator. When Group Policy changes for any reason, these trees are wiped clean and the new policies are rewritten. Windows NT 4.0 policies do not respect these special trees, and they can write to any part of the registry. After such a policy is applied, it persists until the value is intentionally reversed, either by a counteracting Windows NT 4.0 policy or by editing the registry.

**Caution**

- It is possible, using the procedure to add or remove an Administrative Template (.adm file), for Administrative Templates to be added to Group Policy that set registry values outside the approved Group Policy trees, although the Windows Server 2003 family does not include any such .adm file. Choosing to install Windows NT 4.0 .adm files can lead to undesirable persistent registry settings. For more information, see Administrative Templates and Add or remove an Administrative Template (.adm file).

  If you must use the procedure in this manner, see the information about the **Only show policy settings that can be fully managed** check box in Filter the view provided by Administrative Templates and the information about the **Enforce Show Policies Only** Group Policy setting in Add or remove an Administrative Template (.adm file).

# Managing Group Policy from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Group Policy from the command line**
In addition to using Group Policy Object Editor, you can use command-line tools to manage Group Policy. For example, you can use **gpresult** to see what policy is in effect and to troubleshoot problems. You can use **gpupdate** to refresh policy immediately and to specify certain options at the command line.

For more information about Group Policy command-line tools, see the following:

- Gpresult
- Gpupdate

### Notes

- **Gpupdate** replaces the Windows 2000 command **secedit /refreshpolicy**.
- To start Group Policy Object Editor from the command line, see Ways to open Group Policy Object Editor.

# Group Policy Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

**Important**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

For more information about Group Policy, see the Microsoft Group Policy home page on the Microsoft Web site.

**Related Windows tools and online Help topics**

You can use the following tools with Group Policy:

- **IntelliMirror--**IntelliMirror is a set of powerful features for change and configuration management. IntelliMirror ensures that users' data, software, and personal settings are available when they move from one computer to another, and persist when their computers are connected to the network. IntelliMirror features are a subset of Group Policy, which in turn requires the Active Directory directory service. Not all of the tools associated with IntelliMirror technologies are available on all versions of Windows. Of those tools that are available on computers running Windows Server 2003 operating systems, not all of them are appropriate for managing Windows computers. For more information, see IntelliMirror and Help for specific tools.

- **Resultant Set of Policy (RSoP)--**Resultant Set of Policy (RSoP) is an addition to Group Policy that makes it easier to implement and troubleshoot policy. RSoP is a query engine that polls existing policies--and also planned policies--and then reports the results of those queries. It polls existing policies based on site, domain, domain controller, and organizational unit. RSoP gathers this information from the CIMOM database (commonly referred to as "WMI"). For more information, see Resultant Set of Policy.

- **User profiles--**On computers running Windows Server 2003 operating systems, user profiles automatically create and maintain the desktop settings for each user's work environment on the local computer. A user profile is created for each individual user when logging on to a computer for the first time. You set a profile path from the **Profile** tab of a User object in Active Directory. For more information, see User Profiles.

- **Windows Installer--**Windows Installer is an integral part of IntelliMirror, and a core component of the Windows Group Policy-based change and configuration management technology. By using the IntelliMirror technologies, Group Policy, and change and configuration management, you can approve certain applications, specifying that all configuration operations on those applications (installation, removal, and repair) run as the local system account. For more information, see Windows Installer.

# Troubleshooting Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

An important part of troubleshooting Group Policy problems is to consider dependencies between components. For example, Group Policy Software Installation relies on Group Policy, and Group Policy relies on Active Directory. Active Directory relies on proper configuration of network services. When trying to fix problems that appear in one component, it is generally helpful to check whether components, services, and resources on which the component relies are working correctly. Event logs are useful for tracking problems that are caused by this type of hierarchical dependency.

What problem are you having?

- I need to know what policy settings are in effect.
- I want to refresh policy.
- I want to check basic network connectivity.
- I want to enable logging.
- I need help with the individual Group Policy settings that are included in Administrative Templates.
- I cannot open a Group Policy object in Group Policy Object Editor even though I have Read access to it.
- When I try to edit a Group Policy object, I get the "Failed to open the Group Policy object" error.
- When I try to edit a Group Policy object, I get the "Missing Active Directory Container" error.
- When I try to edit a Group Policy object, I get the "Snap-in failed to initialize" error.
- Group Policy is not being applied to users and computers in a security group that contains those users and computers, even though a Group Policy object is linked to an organizational unit that contains that security group.
- Group Policy is not affecting users and computers in a site, domain, or organizational unit.
- Group Policy is not affecting users and computers in an Active Directory container.
- Group Policy is not taking effect on the local computer.
- Software that you assigned to a computer is not available on the computer.
- Software that you assigned or published to a user is not available on the user's computer.
- Published applications do not appear in Add or Remove Programs in Control Panel.
- Document activation of a published application does not cause the application to install.
- The user receives an error message such as "The feature you are trying to install cannot be found in the source directory."
- The user receives an error message such as "This product is not installed" or "Feature ID not registered."
- After removal of an application, the shortcuts for the application continue to appear on the user's desktop.
- A user receives an error message such as "Another installation is already in progress."

- The user opens an already installed application, and Windows Installer starts.
- A user receives error messages such as "Active Directory will not allow the package to be deployed" or "Cannot prepare package for deployment."
- When I click one of the Group Policy Software Installation icons in the console tree of Group Policy Object Editor, I get the error message "Snap-in failed to initialize."
- I accidentally deleted default Group Policy objects.
- I accidentally deleted the domain controllers Sysvol directory or parts of the Sysvol directory that contain default Group Policy objects.
- I configured policies and I am now being denied access.

**I need to know what policy settings are in effect.**
**Solution:** To see what policy is in effect, use Resultant Set of Policy (RSoP) or **gpresult**.

You can also get a report on policy from Help and Support Center by following these steps:

1. Click the **Start** button, and then click **Help and Support**.
2. Click **Support**.
3. Under **See Also**, click **Advanced System Information**.
4. Under **Advanced System Information**, click **View Group Policy settings applied**.

The report includes the following:

- User name and domain
- Computer name and domain
- When User Settings and Computer Settings were last applied
- Folder redirection details
- Logon, logoff, startup and shutdown scripts
- Installed software
- Administrative Templates
- Security Settings
- Connection and proxy settings for IE Maintenance

**See also:** Resultant Set of Policy; Gpresult

**I want to refresh policy.**
**Solution:** Use **gpupdate** to refresh policy immediately. With **Gpupdate** you can specify certain options at the command line. **Gpupdate** replaces and improves on the Windows 2000 command **secedit /refreshpolicy**. Some policy items, such as computer-assigned software, require a reboot to take effect, however. User-assigned software requires the user to log on and log off.

**See also:** Gpupdate; Refresh Group Policy immediately

**I want to check basic network connectivity.**
**Solution:** Try to ping a domain controller.

Verify that DNS is working properly.

**See also:** Ping; DNS; About Network Connections; Configure a connection to a remote network; Using local area connections; Test a TCP/IP configuration by using the ping command

**I want to enable logging.**
**Solution:** Set a registry key that causes Group Policy diagnostic logging to be written to a file named Userenv.log on the client computer.

You can also set a registry key that enables verbose logging to the event log so that you can see it with Event Viewer.

**See also:** Registry Editor; Event Viewer; "Monitoring Group Policy with Log Files" and "Verbose Logging to Event Log" sections in "Troubleshooting Group Policy" at the Microsoft Web site

**I need help with the individual Group Policy settings that are included in Administrative Templates.**
**Solution:** This information is included in Help.

**See also:** Administrative Templates included with this version of Windows

**I cannot open a Group Policy object in Group Policy Object Editor even though I have Read access to it.**
**Cause:** As an administrator, you must have Full Control of the Group Policy object to open it in Group Policy Object Editor.

**Solution:** Verify that you are a member of a security group with Full Control on the Group Policy object. For example, a domain administrator can manage Active Directory-based Group Policy. An administrator on a computer can edit the local Group Policy object on that computer.

**When I try to edit a Group Policy object, I get the "Failed to open the Group Policy object" error.**
**Cause:** This usually is due to a networking problem, specifically, a problem with the domain name system (DNS) configuration.

**Solution:** Verify that DNS is working properly.

**See also:** DNS

**When I try to edit a Group Policy object, I get the "Missing Active Directory Container" error.**
**Cause:** This is caused by Group Policy attempting to link a Group Policy object to an organizational unit that it cannot find. The organizational unit might be deleted, or it might be created on another domain controller but not replicated to the domain controller that you are using.

**Solution:** Limit the number of administrators who can make structural changes to Active Directory, or who can edit a Group Policy object, at any one time. Allow changes to replicate before making changes that affect the same organizational unit or Group Policy object.

**When I try to edit a Group Policy object, I get the "Snap-in failed to initialize" error.**
**Cause:** This may be happen if Group Policy cannot find framedyn.dll.

**Solution:** If you use installation scripts, verify that your scripts place the *systemroot*\system32\wbem directory in the system path. By default, *systemroot*\system32\wbem is in the system path already; therefore, you are not likely to encounter this issue if you do not use installation scripts.

**Group Policy is not being applied to users and computers in a security group that contains those users and computers, even though a Group Policy object is linked to an organizational unit that contains that security group.**
**Cause:** This is correct behavior. Group Policy affects only the users and computers that are contained in sites, domains, and organizational units. Group Policy objects are not applied to security groups.

**Solution:** Link Group Policy objects to sites, domains, and organizational units only. The location of a security group in Active Directory is unrelated to whether Group Policy applies to the users and computers in that security group.

**See also:** Filter the scope of Group Policy according to security group membership

**Group Policy is not affecting users and computers in a site, domain, or organizational unit.**
**Cause:** Group Policy settings can be prevented, intentionally or inadvertently, from affecting users and computers in several ways. A Group Policy object can be disabled so that it does not affect users, computers, or both. It also must be linked directly to an organizational unit that contains the users and computers, or it must be linked to a parent domain or organizational unit,

so that the Group Policy settings apply through inheritance.

When multiple Group Policy objects apply, they are processed in this order: local, site, domain, organizational unit. By default, settings that are applied later have precedence. In addition, Group Policy can be blocked at the level of any organizational unit, or it can be enforced through a setting of **No Override** that is applied to a particular Group Policy object link.

Finally, the user or computer must belong to one or more security groups that have the appropriate permissions set.

**Solution:** Do the following:

- Verify that the intended policy is not being blocked.
- Verify that no overriding policy that is set at a higher level of Active Directory has been set to **No Override**.
- If **Block** and **No Override** are both used, **No Override** takes precedence.
- Verify that the user or computer is not a member of any security group for which the Apply Group Policy access control entry (ACE) is set to **Deny**.
- Verify that the user or computer is a member of at least one security group for which the Apply Group Policy access control entry (ACE) is set to **Allow**.
- Verify that the user or computer is a member of at least one security group for which the Read access control entry (ACE) is set to **Allow**.

**See also:** Policy inheritance; Filter the scope of Group Policy according to security group membership

**Group Policy is not affecting users and computers in an Active Directory container.**
**Cause:** Group Policy objects cannot be linked to Active Directory containers other than sites, domains, and organizational units.

**Solution:** Link a Group Policy object to an organizational unit that is a parent to the Active Directory container. Then, by default, those settings are applied to the users and computers in the container through inheritance.

**See also:** Filter the scope of Group Policy according to security group membership

**Group Policy is not taking effect on the local computer.**
**Cause:** Local policies are the weakest. Any Active Directory-based policy can overwrite them.

**Solution:** Check to see what Group Policy objects are being applied through Active Directory and if those Group Policy objects have settings that are in conflict with the local settings.

**Software that you assigned to a computer is not available on the computer.**
**Cause:** The client computer has not been restarted.

**Solution:** Restart the client computer.

**Software that you assigned or published to a user is not available on the user's computer.**
**Cause:** The user has not logged off and then logged on.

**Solution:** Have the user log off and then log on.

**Published applications do not appear in Add or Remove Programs in Control Panel.**
**Cause:** Several causes are possible:

- Group Policy is not applied.
- Active Directory cannot be accessed.
- Users do not have any published applications in the Group Policy objects that apply to them.
- The client is running Terminal Server.

**Solution:** Investigate each possibility in turn. Note that Group Policy Software Installation is not supported for Terminal Server clients.

**Document activation of a published application does not cause the application to install.**
**Cause:** As an administrator, you did not set autoinstall.

**Solution:** Set autoinstall.

**See also:** Set the autoinstall option for an application

**The user receives an error message such as "The feature you are trying to install cannot be found in the source directory."**
**Cause:** This could be caused by network or permissions problems.

**Solution:** Verify that the network is working correctly.

**See also:** Set permissions for Group Policy Software Installation

**The user receives an error message such as "This product is not installed" or "Feature ID not registered."**
**Cause:** This could be caused by a user with a roaming user profile logging on to two computers simultaneously, which is an unsupported use of roaming user profiles.

**Solution:** Have the user log off one computer before logging on to another. That way, the application shortcuts will be updated and valid on the second computer, and they will not give error messages.

**After removal of an application, the shortcuts for the application continue to appear on the user's desktop.**
**Cause:** The user has created shortcuts, and Windows Installer does not recognize them.

**Solution:** The user must remove the shortcuts manually.

**A user receives an error message such as "Another installation is already in progress."**
**Cause:** An uninstallation might be taking place in the background, with no user interface presented to the user, or perhaps the user has inadvertently triggered two installations simultaneously, which is not supported.

**Solution:** The user can try again later.

**The user opens an already installed application, and Windows Installer starts.**
**Cause:** An application might be undergoing automatic repair, or a user-required feature is being added.

**Solution:** No action is required.

**A user receives error messages such as "Active Directory will not allow the package to be deployed" or "Cannot prepare package for deployment."**
**Cause:** The package might be corrupted, or there might be a networking problem.

**Solution:** Use an uncorrupted package. Investigate the possibility of network problems and take appropriate action.

**When I click one of the Group Policy Software Installation icons in the console tree of Group Policy Object Editor, I get the error message "Snap-in failed to initialize."**
**Cause:** This may be happen if Group Policy Software Installation cannot find framedyn.dll.

**Solution:** If you use installation scripts, verify that your scripts place the *systemroot*\system32\wbem directory in the system path. By default, *systemroot*\system32\wbem is in the system path already; therefore, you are not likely to encounter this issue if you do not use installation scripts.

**I accidentally deleted default Group Policy objects.**
**Cause:** This can happen when you delete Group Policy objects from Active Directory Users and Computers.

**Solution:** Use **dcgpofix** to restore default Group Policy objects to their original state.

**See also:** Dcgpofix

**I accidentally deleted the domain controllers Sysvol directory or parts of the Sysvol directory that contain default Group Policy objects.**
**Cause:** This can happen when you delete directories from the system root.

**Solution:** Restore the Sysvol directory. Use the Dcgpofix command-line tool to restore default Group Policy objects to their original state. It is very important to backup the Sysvol directory with a regularly scheduled backup procedure.

**See also:** Dcgpofix

**I configured policies and I am now being denied access.**

**Cause:** This can happen in default Group Policy objects when setting File system security policy on the Sysvol directory or setting excessive security for domain administrator policies.

**Solution:** Logon as Administrator and use the Dcgpofix command-line tool to restore default domain and domain controller Group Policy objects to their original state.

**See also:** Dcgpofix

# Group Policy Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Management**
The Group Policy Management Console snap-in provides a single user interface through which to manage Group Policy across an enterprise. The Group Policy Management Console consists of a Microsoft Management Console (MMC) snap-in and a set of scriptable interfaces for managing Group Policy.

- Before you use Group Policy Management Console, see Checklist: Using the Group Policy Management Console.
- To find ways to perform common tasks using Group Policy Management Console, see New ways to do familiar tasks using GPMC.
- For help with specific tasks, see GPMC How To....
- For general background information, see GPMC Concepts.

# Checklist: Using the Group Policy Management Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Using the Group Policy Management Console**

This section shows recommended steps to start using Group Policy Management Console (GPMC).

| Step | Reference |
|---|---|
| Check system requirements for GPMC | You need Windows Server 2003 or Windows XP Professional Service Pack 1 (SP1) to install GPMC. Some features require a Windows Server 2003 domain controller. For details, see Deployment considerations for Group Policy. |
| Install GPMC | Follow the GPMC installation instructions from the Microsoft Web site or installation media where you obtained the installer package (GPMC.msi). Run GPMC.msi, read and accept the End User License Agreement (EULA), and finish the installation. |
| Check for last-minute issues | Consult the release notes that were installed with GPMC, in the **%systemroot%\Program Files\GPMC** folder. |
| Start GPMC | Open GPMC. For step-by-step instructions, see Start Group Policy Management Console. |

# New ways to do familiar tasks using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common Group Policy tasks. Using Group Policy Management Console (GPMC), the procedures for performing these tasks are different from the way you performed these tasks in the past, without the Group Policy Management snap-in. Previously Administrators used Active Directory Users and Computers and Active Directory Sites and Services to manage Group Policy. Now you use GPMC to perform all Group Policy-related tasks.

| If you want to | On a computer without GPMC | With Group Policy Management snap-in |
|---|---|---|
| Create a Group Policy object (GPO) and a link to it | On the **Group Policy** tab for a site, domain, or organizational unit, click **New**. | In GPMC, right-click the relevant site, domain, or organizational unit and then click **Create and link a GPO here**. For more information, see Create or delete a Group Policy object and Link a Group Policy object using GPMC. |
| Create an unlinked GPO | Navigate to the **Group Policy** properties tab for a site, domain, or organizational unit, click **Add**. On the **All** tab, right-click, and then click **New**. | In GPMC, navigate to **Group Policy Objects**, right-click it, and then click **New**. For more information, see Create or delete a Group Policy object. |
| Edit a GPO | Navigate to the **Group Policy** properties tab for a site, domain, or organizational unit, select the GPO, and then click **Edit**. This opens Group Policy Object Editor. | In GPMC, navigate to **Group Policy Objects**, right-click the GPO, and then click **Edit**. For more information, see Edit a Group Policy object from GPMC.<br>**Note**<br>• Group Policy Object Editor MMC snap-in is still used to edit GPOs and is opened when you click **Edit** in GPMC. |
| Link an existing GPO to a site, domain, or organizational unit | Navigate to the **Group Policy** properties tab for a site, domain, or organizational unit, click **Add**, and then select a GPO. | In GPMC, navigate to a site, domain, or organizational unit, right-click it, and then click **Link an existing GPO here**. For more information, see Link a Group Policy object using GPMC. |
| Use security groups to filter the scope of policy | Edit the access control entry (ACE) permission option for **Apply Group Policy** on the **Security** tab in the **Properties** dialog box for the GPO. | In GPMC, click a GPO or GPO Link, select the **Scope** tab, and then use the **Add** and **Remove** buttons to control the groups, users, and computers on which the GPO applies. For more information, see Filter using security groups. |
| Delegate permissions on GPOs | Navigate to the properties of a GPO, click **Permissions for Authenticated Users**, and select the check boxes that correspond to the permissions you want to give. | In GPMC, click the GPO, click the **Delegation** tab, and then click the **Add** or **Remove buttons**. For more information, see Delegate Group Policy tasks. |
| Determine Resultant Set of Policy with Group Policy Results (logging mode) | In Resultant Set of Policy snap-in, right-click **Resultant Set of Policy**, and then click **Generate RSoP Data**. Select **Logging Mode** in the Resultant Set of Policy Wizard. | In GPMC, right-click **Group Policy Results**, and then click **Group Policy Results Wizard**. For more information, see Determine Resultant Set of Policy with Group Policy Results. |
| Simulate Resultant Set of Policy using Group Policy Modeling | In Resultant Set of Policy snap-in, right-click **Resultant Set of Policy**, and then click **Generate RSoP Data**. Select **Planning Mode** in the Resultant Set of Policy Wizard. | In GPMC, right-click **Group Policy Modeling**, and then click **Group Policy Modeling Wizard.** For more information, see Simulate Resultant Set of Policy using Group Policy Modeling. |

| | | |
|---|---|---|
| Add a WMI filter | Navigate to the Properties of a Group Policy object, and click the **WMI Filter** tab. | In GPMC, right-click **WMI Filters** and click **New**. For more information, see Create, import, export, copy, and paste WMI filters. |

# GPMC How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Start Group Policy Management Console
- Create or delete a Group Policy object
- Edit a Group Policy object from GPMC
- Add a forest, site, or domain to the Group Policy Management Console
- Search for a Group Policy object
- Specify a domain controller
- View, print, and save a report of Group Policy object settings
- Determine Resultant Set of Policy with Group Policy Results
- Simulate Resultant Set of Policy using Group Policy Modeling
- Set Group Policy Management Console user interface options
- Back Up, Restore, Copy, and Import
- Control Group Policy Object Scope
- Delegate Group Policy tasks
- Use Script and Command Line Methods

# Start Group Policy Management Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Group Policy Management console**

- To open Group Policy Management, click **Start,** click **Control Panel,** click **Administrative Tools,** and then click **Group Policy Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Create or delete a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create or delete a Group Policy object**

1. Open Group Policy Management.
2. Depending upon whether you want to create or delete, use one of the following procedures:
   - Create
   - Create and link
   - Delete

**Create**

1. In the console tree, right-click **Group Policy Objects** in the forest and domain in which you want to create a Group Policy object (GPO).

   **Where?**

   *Forest name*/Domains/*Domain name*/Group Policy Objects

2. Click **New**.
3. In the **New GPO** dialog box, specify a name for the new GPO, and then click **OK**.

**Create and link**

1. In the console tree, right-click the *domain name* in the forest in which you want to create and link a Group Policy object (GPO).

   **Where?**

   *Forest name*/Domains/*Domain name*

2. Click **Create and Link a GPO Here**.
3. In the **New GPO** dialog box, specify a name for the new GPO, and then click **OK**.

**Delete**

1. In the console tree, double-click **Group Policy Objects** in the forest and domain containing the Group Policy object (GPO) that you want to delete.

   **Where?**

   *Forest name*/Domains/*Domain name*/Group Policy Objects

2. Right-click the GPO, and then click **Delete**.
3. When prompted to confirm the deletion, click **OK**.

**Notes**

- To create a GPO, you must have GPO creation privileges. By default only domain administrators, enterprise administrators, and members of the Group Policy creator owners group can create Group Policy objects. To delegate GPO creation permissions to additional groups and users, go to **Group Policy Objects** in the desired domain and click the Delegation tab.

- To delete a GPO, you must have **Edit Settings**, **Delete**, **Modify Security** permissions for the GPO.

- When you use this procedure to create a GPO, no links are created to the GPO, but you can add links within the same forest by right-clicking any domain, site, or organizational unit, and then clicking **Link Existing GPO**. Alternatively, you can both create and link a GPO by right-clicking any domain or organizational unit and then clicking **Create and Link a GPO Here**.

- When you delete a GPO, Group Policy Management attempts to delete all links to that GPO in the domain of the GPO. However, to delete a link to a GPO, you must have permission to link Group Policy objects for the organizational unit or domain. If you do not have rights to delete a link, the GPO will be deleted, but the link will remain. Links from other domains and sites are not deleted. The link to a deleted GPO appears in Group Policy Management as **Not Found**. To delete **Not Found** links, you must either have permission on the site, domain or organizational unit containing the link, or ask someone with sufficient rights to delete it.

- Group Policy objects are distinguished in the Active Directory by GUID, and it is theoretically possible for more than one GPO to have the same friendly name. The Group Policy Management snap-in prevents the creation of Group Policy objects with duplicate friendly names, but the Group Policy infrastructure does not enforce uniqueness of friendly names. Therefore, it is possible for duplication of friendly names to occur if you use legacy tools to create Group Policy objects, if replication is slow, or if you use a script to perform operations on Group Policy objects.

- You cannot delete the Default Domain Controllers policy or the Default Domain policy.

- Before deleting a GPO, you can check for cross-domain links by navigating to the **Scope** tab of the GPO you want to delete and, in the **Display links in this location** box, selecting **Entire Forest**. You can then select all links, right click the selection, and click **Delete link**. This procedure ensures that cross-domain links are deleted before you delete the GPO.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy Management Console Overview
Start Group Policy Management Console

# Edit a Group Policy object from GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit a Group Policy object**

1. Open Group Policy Management.

2. In the console tree, double-click **Group Policy Objects** in the forest and domain containing the Group Policy object (GPO) that you want to edit.

   **Where?**

   - *Forest name*/Domains/*Domain name*/Group Policy Objects

3. Right-click the GPO, and then click **Edit**.

4. In the **Group Policy Object Editor** console, edit the settings as appropriate.

   **Important**

   - Avoid editing the default domain policy. If you need to apply Group Policy settings to the entire domain, create a new GPO, link it to the domain, and create the settings in that GPO.

   - The default domain policy and default domain controllers policy are vital to the health of any domain. You should not edit the Default Domain Controller policy or the Default Domain policy, except in the following cases:

     It is recommended that account policy be set in the Default Domain Policy.

     If you install applications on domain controllers that require modifications to User Rights or Audit Policies, the modifications must be made in the Default Domain Controllers Policy.

## Notes

- To edit a GPO, you must have **edit** permissions for the GPO that you want to edit.

- To edit IPSec policy settings from within a GPO, you must be a member of the domain administrators group. Otherwise, you can only assign or unassign an existing IPSec policy.

- You can also edit a GPO by clicking any link to the GPO in the console tree and following the above steps. When using Group Policy Object Editor to edit the settings in the GPO, any changes that you make are global to the GPO, so they will impact all locations where the GPO is linked.

- To edit the local Group Policy object, you must open it by clicking **Start**, clicking **Run**, typing **gpedit.msc**, and then clicking **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Group Policy Management Console Overview
Start Group Policy Management Console

# Add a forest, site, or domain to the Group Policy Management Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Add a forest, site, or domain to the Group Policy Management Console**

1. Open Group Policy Management.
2. Depending upon whether you want to add a forest, site, or domain, use one of the following procedures:
   - Forest
   - Domain
   - Site

**Forest**

1. In the console tree, right-click **Group Policy Management**, and then click **Add Forest**.
2. In the **Add Forest** dialog box, type the DNS or NetBIOS name of any domain in the forest, and then click **OK**.

   You can specify either the DNS name or the NetBIOS name of any domain in the forest. If you specify a NetBIOS name, you must confirm that the NetBIOS name corresponds to the DNS name of the domain.

   The forest is added to Group Policy Management Console, along with the domain that you specified.

**Domain**

1. In the console tree, double-click the forest in which you want to add a domain, right-click **Domains**, and then click **Show Domains**.
2. In the **Show Domains** dialog box, select the check box next to each domain to add in Group Policy Management Console and then click **OK**.

**Site**

1. In the console tree, double-click the forest in which you want to add a site, right-click **Sites**, and then click **Show Sites**.
2. In the **Show Sites** dialog box, select the check box next to each site to add to Group Policy Management Console and then click **OK**.

## Notes

- Adding a forest to Group Policy Management does not create a new forest in Active Directory. This function simply allows you to view and manage an existing forest using Group Policy Management Console. Likewise, adding a domain or site to Group Policy Management does not actually create a new domain or site in Active Directory.
- When you add a forest, the forest displayed in Group Policy Management is named after the forest root domain for that forest. By default, Group Policy Management shows a single domain in the forest, which is the domain that you specified in the **Add forest** dialog box. To show other domains, you must add them.

- By default, sites do not appear in the console. Sites are shown only if you explicitly add them.

- Because Group Policy is not inherited across domains, domains are always shown as peers of one another, regardless of their actual relationship.

- You can use the **Add Forest** function to add externally trusted domains, even if you do not have forest trust with the entire forest. You must, however, have trust between the domain you want to add and the domain of your user object. In the **Add Forest** dialog box, enter the name of the externally trusted domain that you want to add.

- When you close Group Policy Management Console, it preserves the current list of forests, domains, and sites shown.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy Management Console Overview
Start Group Policy Management Console

# Search for a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To search for a Group Policy object**

1. Open Group Policy Management.

2. In the console tree, double-click the forest containing the domain that you want to search for a Group Policy object (GPO), double-click **Domains**, right-click the domain, and then click **Search**.

3. In the **Search for Group Policy Objects** dialog box, in the **Search for GPOs in this domain** box, select a domain or **All domains shown in this forest**.

4. In the **Search item** box, select the type of object on which you want to base your search.

   If you select **Security Group**, the **Select User, Computer, or Group** dialog box appears. Specify the appropriate object type, location of the object, and object name, and then click **OK**.

   You can choose **GPO-links** on the **Search item** dropdown menu to find unlinked GPOs and GPOs linked across domains.

5. In the **Condition** box, select the condition that you want to use in the search.

6. In the **Value** box, select or specify the value that you want to use to filter the search, and then click **Add**.

7. Repeat steps 4 through 6 until you complete the definition of all search criteria, and then click **Search**.

   The search is based on the intersection of the criteria specified, so a GPO must meet all criteria specified for it to be returned in the results.

8. When search results are returned, do one of the following:

   - To save the search results, click **Save results** and then, in the **Save GPO Search Results** dialog box, specify the file name for the saved results, and then click **Save**.

   - To navigate to a GPO found in the search, double-click the GPO in the search results list.

   - To clear the search results, click **Clear**.

9. Repeat steps 3 through 8 until you have completed all required searches, and then click **Close**.

### Notes

- You can also open the search dialog box by by right-clicking a forest and then clicking Search. In this case, the Search for GPOs in this domain box defaults to **All domains shown in this forest**.

- If a setting is enabled, and then all the settings in that extension are removed, there can be false-positive search for certain types of settings. This happens because the GPO has the extension listed as active. The extensions with this behavior are Security Settings, Software Installation, Folder Redirection, Internet Explorer Maintenance, and Encrypting File System (EFS).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy Management Console Overview
Start Group Policy Management Console

# Specify a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify a domain controller**

1. Open Group Policy Management.
2. In the console tree, take one of the following actions, depending on whether you want to change the domain controller for a domain or a site:
   - If you want to change the domain controller for a domain, double-click the forest that contains the domain for which you want to change the domain controller, double-click **Domains**, right-click the domain, and then click **Change Domain Controller**.
   - If you want to change the domain controller for a site, double-click the forest that contains the sites for which you want to change the domain controller, right-click **Sites**, and then click **Change Domain Controller**. In the **Look in this domain** list of the **Change Domain Controller** dialog box, select the domain in which the domain controller resides.
3. To change the domain controller for either a site or a domain, do one of the following:
   - To choose the primary domain controller, select **The domain controller with the Operations Master token for the PDC emulator**, and then click **OK**.
   - To choose any available domain controller, select **Any available domain controller**, and then click **OK**.
   - To choose any available domain controller running Windows Server 2003 or later, select **Any available domain controller running Windows Server 2003 or later**, and then click **OK**.
   - To choose a specific domain controller, select **This domain controller**, select a domain controller from the list, and then click **OK**.

## Notes

- All operations performed in the domain, including operations on Group Policy objects (GPOs), security groups, and organizational units in this domain use the domain controller you specify using the **Change Domain Controller** function.
- All operations performed on links to sites use the domain controller that you specify for the sites, but all operations on Group Policy objects (GPOs) use the domain controller you specify for the domain in which the GPO resides.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy Management Console Overview
Start Group Policy Management Console

# View, print, and save a report of Group Policy object settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view, print, and save a report of Group Policy object settings**

1. Open Group Policy Management.
2. In the console tree, double-click the forest containing the domain or organizational unit which contains the Group Policy object (GPO) you want to work with, and then select the GPO.
3. Click the **Settings** tab to view reports in GPMC.
4. If you want to print or save the report, right-click the settings report in the results pane and do one of the following:
   - Select **Print** to print the report.
   - Select **Save Report** to save the report.

## Notes

- To customize the information displayed in a report, click **show** or **hide** to display only the data you want to view or print.
- On computers with Internet Explorer Enhanced Security Configuration enabled a prompt appears when you attempt to view reports in the Group Policy Management Console (GPMC) snap-in. This happens because the reports contain a script that allows you to expand and collapse sections of the reports using **Show** and **Hide**. To use the show and hide functionality in these reports, you must add the about:security_MMC.exe site to the Trusted sites zone in Internet Explorer. This site represents all Web pages that are hosted inside MMC. You can do this by clicking **Add** on the **Internet Explorer** prompt. This opens the **Trusted sites** dialog box with the correct entry (about:security_mmc.exe) for the page being called by GPMC. Select **Add** and then **Close** on the **Trusted sites** dialog box to add this site to the Trusted sites zone. If the about:security_mmc.exe site is not added to the Trusted sites zone, the reports appear fully expanded and you cannot collapse the report.
- To view a description of a setting in Administrative Templates, click the setting name in the report. This feature is only available for settings in the Administrative Templates section of the report.
- From within Group Policy Management, it is not possible to use the keyboard to navigate inside an HTML report. To navigate inside the HTML save the report to a file and then use Internet Explorer to view the report. The reports have been designed to support keyboard navigation when viewed directly in Internet Explorer.
- To view a saved report directly in a Web browser, you must use Internet Explorer 6 or Netscape 7. Netscape 7 does not support functionality that enables you to show or hide data in reports.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy Management Console Overview
Start Group Policy Management Console

# Determine Resultant Set of Policy with Group Policy Results

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Determine Resultant Set of Policy with Group Policy Results**

1. Open Group Policy Management.
2. In the console tree, double-click the forest in which you want to create a Group Policy Results query, right-click **Group Policy Results**, and then click **Group Policy Results Wizard**.
3. In the **Group Policy Results Wizard**, click **Next** and enter the appropriate information.
4. After completing the wizard, click **Finish**.
5. If you want to print or save the report, right-click the settings report in the results pane and do one of the following:
   - Select **Print** to print the report.
   - Select **Save Report** to save the report.

### Notes

- To access Group Policy Results data for a user or computer, you must have **Read Group Policy Results data** permission on the domain or organizational unit that contains the user or computer, or you must be a member of a local Administrator's group on the targeted computer.
- To customize the information displayed in a report, click **show** or **hide** to display only the data you want to view or print.
- On computers with Internet Explorer Enhanced Security Configuration enabled a prompt appears when you attempt to view reports in the Group Policy Management Console (GPMC) snap-in. This happens because the reports contain a script that allows you to expand and collapse sections of the reports using Show and Hide. To enable the Show and Hide functionality in these reports, you must add the about:security_MMC.exe site to the Trusted sites zone in Internet Explorer. This site represents all Web pages that are hosted inside MMC. You can do this by clicking **Add** on the **Internet Explorer** prompt. This opens the **Trusted sites** dialog box with the correct entry (about:security_mmc.exe) for the page being called by GPMC. Select **Add** and then **Close** on the **Trusted sites** dialog box to add this site to the Trusted sites zone. If the about:security_mmc.exe site is not added to the Trusted sites zone, the reports appear fully expanded and you cannot collapse the report.
- From within Group Policy Management, it is not possible to use the keyboard to navigate inside an HTML report. To navigate inside the HTML save the report to a file and then use Internet Explorer to view the report. The reports have been designed to support keyboard navigation when viewed directly in Internet Explorer.
- If you open a previously saved console, and want to save a Group Policy Modeling or Group Policy Results report in XML, the report must be regenerated using the **Rerun query** option.
- To view a saved report directly in a Web browser, you must use Internet Explorer 6 or Netscape 7. Netscape 7 does not support functionality that enables you to show or hide data in reports.
- The functionality provided by Group Policy Results was previously referred to as Resultant Set of Policy - logging mode in Windows XP.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy Management Console Overview
Start Group Policy Management Console

# Simulate Resultant Set of Policy using Group Policy Modeling

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To simulate Resultant Set of Policy using Group Policy Modeling**

1. Open Group Policy Management.
2. In the console tree, double-click the forest in which you want to create a Group Policy Modeling query, right-click **Group Policy Modeling**, and then click **Group Policy Modeling Wizard**.
3. In the **Group Policy Modeling Wizard**, click **Next** and enter the appropriate information.
4. When finished, click **Finish**.
5. If you want to print or save the report, right-click the settings report in the results pane and do one of the following:
   - Select **Print** to print the report.
   - Select **Save Report** to save the report.

**Notes**

- To create a Group Policy Modeling query, you must have the **Perform Group Policy Modeling analyses** permission on the domain or organizational unit that contains the objects on which you want to run the query.
- **Group Policy Modeling** is only available if at least one domain controller in the forest is running Windows Server 2003.
- To customize the information displayed in a report, click **show** or **hide** to display only the data you want to view or print.
- On computers with Internet Explorer Enhanced Security Configuration enabled a prompt appears when you attempt to view reports in the Group Policy Management Console (GPMC) snap-in. This happens because the reports contain a script that allows you to expand and collapse sections of the reports using Show and Hide. To enable this Show and Hide functionality in these reports, you must add the about:security_MMC.exe site to the Trusted sites zone in Internet Explorer. This site represents all Web pages that are hosted inside MMC. You can do this by clicking **Add** on the **Internet Explorer** prompt. This opens the **Trusted sites** dialog box with the correct entry (about:security_mmc.exe) for the page being called by GPMC. Select **Add** and then **Close** on the **Trusted sites** dialog box to add this site to the Trusted sites zone. If the about:security_mmc.exe site is not added to the Trusted sites zone, the reports appear fully expanded and you cannot collapse the report.
- In Windows Server 2003 without Group Policy Management installed, Group Policy Modeling was referred to as Resultant Set of Policy (RSoP) planning mode.
- From within Group Policy Management, it is not possible to use the keyboard to navigate inside an HTML report. To navigate inside the HTML save the report to a file and then use Internet Explorer to view the report. The reports have been designed to support keyboard navigation when viewed directly in Internet Explorer.
- If you open a previously saved console, and want to save a Group Policy Modeling or Group Policy Results report in XML, the report must be regenerated using the **Rerun query** option.
- To view a saved report directly in a Web browser, you must use Internet Explorer 6 or Netscape 7. Netscape 7 does not support functionality that enables you to show or hide data in reports.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy Management Console Overview
Start Group Policy Management Console

# Set Group Policy Management Console user interface options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set Group Policy Management Console user interface options**

1. Open Group Policy Management.
2. On the **View** menu, click **Options**.
3. Choose a type of view option from the list below:
   - Columns
   - Reporting
   - General

**Columns**

1. In the **Options** dialog box, click the **Columns** tab.
2. Select a table from the **Table location** drop down box.
3. In **Columns displayed in this order**, add a check next to each column that should be displayed.
4. Select a column and use the **Move Up** and **Move Down** buttons to set the order of columns.
5. Repeat steps 2 through 4 for each table in the **Table location** drop down box you want to change, and then click **OK**.
6. You can optionally check **Save customized column order and sizes** to make column order and size selections persistent. This setting is checked by default.

**Reporting**

1. In the **Options** dialog box, click the **Reporting** tab.
2. Select one of the following options.
   - To specify that Group Policy Managemetn Console (GPMC) should use the local adm files,select **Default**. If an .adm file is not found locally, and then GPMC will look in the Group Policy objects (GPO) folder in the sysvol folder.
   - To specify a custom location that takes precedence over the default locations, select **Custom**, and then enter the location where .adm files are stored.

     **Notes**

   - Group Policy Management Console (GPMC) uses files with an .adm extension to display the friendly names of policy settings when generating HTML reports for GPOs, Group Policy Modeling, and Group Policy Results. These options let you control the location from which GPMC reads these .adm files.
   - When searching for a given .adm file, GPMC uses the .adm file that is first found based on the search order described above, regardless of time stamp.

- If you add or modify an .adm file in an existing location, you need to restart GPMC in order for GPMC to pick up the change in the .adm file when displaying HTML reports. (You do not need to restart GPMC if you switch between the **Default** and **Custom** option buttons, or change the custom location used for .adm files, and click **OK**.)

- GPMC never copies the .adm file to the sysvol folder, unlike the Group Policy Object Editor.

**General**

1. In the **Options** dialog box, click the **General** tab.

2. Select from the following options.

   - To display only forests or domains with two-way trusts to the user's domain, select **Enable trust checking**. This setting is checked by default.

   - To show domain name controller names in parentheses after each domain in GPMC, select **Show domain controllers after domain names**.

   - To display a confirmation dialog box to distinguish between GPOs and GPO links each time a GPO or GPO link is selected in GPMC, select **Show confirmation dialog box to distinguish between GPOs and GPO links**. This setting is checked by default.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy Management Console Overview
Start Group Policy Management Console
Administrative templates overview for GPMC

# Back Up, Restore, Copy, and Import

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Back Up, Restore, Copy, and Import**

- Back up a Group Policy object using GPMC
- Restore a backed-up Group Policy object using GPMC
- Copy a Group Policy object using GPMC
- Import a Group Policy object using GPMC
- Create a migration table

# Back up a Group Policy object using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up a Group Policy object**

1. Open Group Policy Management.

2. In the console tree, double-click **Group Policy Objects** in the forest and domain containing the Group Policy object (GPO) that you want to back up.

   **Where?**

   - *Forest name*/Domains/*Domain name*/Group Policy Objects

3. To backup a single GPO, right-click the GPO, and then click **Back Up**. To backup all GPOs in the domain, right-click **Group Policy Objects** and click **Back Up All**.

4. In the **Backup Group Policy Object** dialog box, in the **Location** box, enter the path to the location at which you want to store the GPO backup(s), or click **Browse**, locate the folder in which you want to store the GPO backup(s), and then click **OK**.

5. In the **Description** box, type a description for the GPO(s) that you want to back up, and then click **Backup**. If you are backing up multiple GPOs, the description will apply to all GPOs you back up.

6. After the operation completes, click **OK**.

   **Important**

   - To secure backed up GPOs, ensure that only authorized administrators have permission to access the folder to which you are exporting the GPO.

## Notes

- You must have **Read** permissions on the GPO and **Write** permissions on the folder containing the GPO backup to perform this procedure.

- You can also perform this procedure or a related task by using one or more of the sample scripts included with Group Policy Management. For more information, see "Scripting Group Policy tasks" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a migration table
Restore a backed-up Group Policy object using GPMC
Copy a Group Policy object using GPMC
Import a Group Policy object using GPMC
Backup using GPMC
Copy using GPMC

Restore using GPMC
Import using GPMC
Migration tables
The migration table editor
Scripting Group Policy tasks using GPMC
Start Group Policy Management Console

# Restore a backed-up Group Policy object using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore a backed-up Group Policy object**

1. Open Group Policy Management.
2. In the console tree, in the forest and domain containing the Group Policy object (GPO) that you want to restore, locate **Group Policy Objects**.

   **Where?**

   - *Forest name*\Domains*Domain name*\Group Policy Objects.

3. Do one of the following:

   - To restore a previous version of an existing GPO
   - To restore a deleted GPO

**To restore a previous version of an existing GPO**

1. Double-click **Group Policy Objects**, right-click the GPO which you want to restore, and then click **Restore from Backup**.
2. When the **Restore Group Policy Object Wizard** appears, follow the instructions and provide the appropriate information about the GPO that you want to restore, and then click **Finish**.
3. After the **Restore Group Policy Object Wizard** completes the restore operation, click **OK**.

**To restore a deleted GPO**

1. Right-click **Group Policy Objects** and then choose **Manage Backups**.
2. In the **Manage Backups** dialog box, in the **Backup location** box, type the path to the backup folder. You can also use **Browse** to locate the backup folder.
3. In the **Backed up GPOs** box, select the GPO that you want to restore from the list of GPO backups shown, and then click **Restore**.
4. When prompted to confirm the restore operation, click **OK**.

## Notes

- You must have **Edit, delete, and modify security** permissions on the GPO and **Read** permissions on the folder containing the GPO backup to restore an existing GPO.
- You must have privileges to **create** GPOs in the domain and **Read** permissions on the file system location of the backed up GPO to restore a GPO that has been deleted.
- You can also restore an existing or deleted GPO using the **Manage backups** function by right-clicking **Domains** or **Group Policy Objects**.

- The **Manage Backups** dialog box can be used to restore either an existing or deleted GPO. The **Manage Backups** dialog box can be opened either by right-clicking **Domains** or **Group Policy Objects** in a given domain. When **Manage Backups** is opened by right clicking **Group Policy Objects**, only GPO backups from that domain are shown. In contrast, when **Manage Backups** is opened by right clicking **Domains**, all GPO backups are shown, regardless of domain.

- You can also perform this procedure or a related task by using one or more of the sample scripts included with Group Policy Management. For more information, see "Scripting Group Policy tasks" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a migration table
Back up a Group Policy object using GPMC
Copy a Group Policy object using GPMC
Import a Group Policy object using GPMC
Backup using GPMC
Copy using GPMC
Restore using GPMC
Import using GPMC
Migration tables
The migration table editor
Scripting Group Policy tasks using GPMC
Start Group Policy Management Console

# Copy a Group Policy object using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy a Group Policy object**

1. Open Group Policy Management.
2. In the console tree, right-click the Group Policy object (GPO) that you want to copy, and then click **Copy**.

   **Where?**

   - *Forest name*/Domains/*Domain name*/Group Policy Objects/*Group Policy object name*

3. Do one of the following:

   - To create a copy of the GPO in the same domain as the source GPO, right-click **Group Policy Objects**, click **Paste**, specify permissions for the new GPO in the **Copy GPO** box, and then click **OK**.
   - To create a copy of the GPO in a different domain, double-click the destination domain, right-click **Group Policy Objects**, and then click **Paste**. Answer all the questions in the cross-domain copying wizard that appears, and then click **Finish**.

## Notes

- You must have privileges to create GPOs in the destination domain to complete this procedure.
- For copy operations to another domain or forest, you may need to specify a migration table to update references to security principals and UNC paths in the source GPO to new values in the destination GPO. For more information about migration tables see "Migration Tables," "The migration table editor," and "Create a migration table" in Related Topics.
- You can also perform this procedure or a related task by using one or more of the sample scripts included with Group Policy Management. For more information, see "Scripting Group Policy tasks" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a migration table
Back up a Group Policy object using GPMC
Import a Group Policy object using GPMC
Backup using GPMC
Copy using GPMC
Restore using GPMC
Import using GPMC
Migration tables
The migration table editor
Restore a backed-up Group Policy object using GPMC
Scripting Group Policy tasks using GPMC
Start Group Policy Management Console

# Import a Group Policy object using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import a Group Policy object**

1. Open Group Policy Management.
2. In the console tree, double-click **Group Policy Objects** in the forest and domain containing the Group Policy object (GPO) into which you want to import settings.

   Where?

   - *Forest name*
   - Domains
   - *Domain name*
   - Group Policy Objects
3. Right-click the GPO, and then click **Import Settings**.
4. Follow the instructions in the **Import Settings Wizard**, and then click **Finish**.

## Notes

- You must have **Edit** permissions on the GPO into which you want to import settings to complete this procedure.
- When importing settings from a backed up GPO in another domain or forest, if the GPO settings contain references to UNC paths or security principals, you may need to update these references to new values in the destination GPO. You can specify a migration table during the import operation to update the references to security principals and UNC paths in the destination GPO. For more information about migration tables see "Migration Tables," "The migration table editor," and "Create a migration table" in Related Topics.
- You can also perform this procedure or a related task by using one or more of the sample scripts included with Group Policy Management. For more information, see "Scripting Group Policy tasks" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a migration table
Back up a Group Policy object using GPMC
Copy a Group Policy object using GPMC
Backup using GPMC
Copy using GPMC
Restore using GPMC
Import using GPMC
Migration tables

The migration table editor
Restore a backed-up Group Policy object using GPMC
Scripting Group Policy tasks using GPMC
Start Group Policy Management Console

# Create a migration table

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a migration table**

1. Open Group Policy Management.

2. In the console tree, double click the forest from which you want to create a migration table.

3. Right-click **Domains** and then click **Open Migration Table Editor**

   - or -

   Right-click **Group Policy Objects**, and then click **Open Migration Table Editor**.

4. Specify a source you want to include in the migration table by doing the following:

   a. Under **Source Name**, type the source name or right-click the cell and click **Browse** to find and then specify the object type, location and name.

      If you enter the source name manually, you must type the exact name of the User, Computer, Group or UNC Path referenced in the source GPO. For example, **TestDomain\PolicyAdmins** or **\\server1\publicshare.** The type of the source name must match the source type specified in the migration table.

      Security principals can be specified using any of the following formats:

      UPN. For example, someone@testdomain.com.

      SAM. For example, testdomain\someone.

      DNS. For example, testdomain.com\someone.

      Free text. For example, someone. You must specify the type as Free Text or SID in this case.

      SID. For example, S-1-11-111111111-111111111-1111111111-1112. You must specify the type as Free Text or SID in this case.

   b. Under **Source Type** specify the type of source in the cell adjacent to the source name cell by using the drop down list and selecting the appropriate type. If you have used **Browse** to find the source name, the source type is entered by default.

   c. Under **Destination Name**, type the destination name or right-click the cell adjacent to the source name and source type cells, and then do one of the following:

      Click **Browse** to find a user, computer or group

      - or -

      Point to **Set Destination**, and then click the appropriate selection. You cannot use the **No Destination** option for UNC paths.

5. Repeat step 4 for each source you want to include in the migration table.

6. On the menu, click **File**, and then click **Save**.

7. Specify the location where you want to save the migration table, and then click **Save**.

**Notes**

- You can use the migration table editor to check that destination entries can be resolved. This is important to know before you copy or import a GPO, because if there are unresolvable destination entries in the migration table, then the copy or import operation might fail. To validate your migration table, from the **Tools** menu, click **Validate**.

- You can auto-populate a migration table by scanning one or more GPOs or backup GPOs to extract all references to security principals and UNC paths, and automatically enter these items into the table as source name entries.

  **To auto-populate a migration table from a GPO:**

  1. From the **Tools** menu, click **Auto-populate from GPO**.

  2. In the **Look in this domain** box, select the domain that contains the GPO.

  3. In the **Group Policy objects** box, click the GPO or GPOs from which you want to populate the migration table, check **Include security principals from the DACL on the GPO during scan** if appropriate, and then click **OK**.

  **To auto-populate a migration table from a backup GPO:**

  1. From the **Tools** menu, click **Auto-populate from Backup**.

  2. In the **Backup location** box, type the path to the folder that contains the backup GPO, or click **Browse** to locate the folder.

  3. In the **Backed up GPOs** box, select one or more backed up GPOs, check **Show only the latest version of each GPO** or **Include security principals from the DACL on the GPO during scan** if appropriate, and then click **OK**.

  To complete the migration table for either scenario, you only need to adjust the destination values. By default, the destination name for each entry will be set to **Same as source** when you use either of the Populate options.

- Performing this task does not require you to have administrative credentials.

- You can also perform this procedure or a related task by using one or more of the sample scripts included with Group Policy Management. For more information, see "Scripting Group Policy tasks" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Back up a Group Policy object using GPMC
Restore a backed-up Group Policy object using GPMC
Copy a Group Policy object using GPMC
Import a Group Policy object using GPMC
Migration tables
The migration table editor
Scripting Group Policy tasks using GPMC
Start Group Policy Management Console

# Control Group Policy Object Scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Control Group Policy Object Scope**

- Link a Group Policy object using GPMC
- Link a WMI filter to a Group Policy object
- Filter using security groups
- Create, import, export, copy, and paste WMI filters.
- Enforce a Group Policy object link using GPMC
- Block inheritance using GPMC
- Get user settings from the machine account (loopback processing)
- Disable a Group Policy object link using GPMC
- Disable the User or Computer settings in a Group Policy object using GPMC

# Link a Group Policy object using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To link a Group Policy object**

1. Open Group Policy Management.
2. In the console tree, locate the site, domain or organizational unit to which you want to link a Group Policy object (GPO)

   Where?

   - *Forest name*
   - Domains or *Sites*
   - *Site name*, *Domain name* or *organizational unit name*
3. Do one of the following:

   - To link an existing GPO, right-click the domain or organizational unit within the domain, and then click **Link an Existing GPO**. In the **Select GPO** dialog box, click the GPO which you want to link, and then click **OK**.
   - To link a new GPO, right-click the domain or organizational unit within a domain, and then click **Create and Link a GPO Here**. In the **Name** box, type a name for the new GPO, and then click **OK**.

## Notes

- To link an existing GPO to a site, domain, or organization unit, you must have **Link GPOs** permission on that site, domain, or organizational unit. By default, only Domain Administrators and Enterprise Administrators have this privilege for domains and organizational units, and only Enterprise Administrators and Domain Administrators of the forest root domain have this privilege for sites.
- To create and link a GPO, you must have **Link GPOs** permissions on the desired domain or organizational unit, and you must have permission to create GPOs in that domain. By default, only Domain Administrators, Enterprise Administrators, and Group Policy Creator owners have permission to create GPOs.
- The **Create and Link a GPO Here** option is not available for sites, because it is unclear in which domain to create the GPO. The user must first create a GPO in any domain in the forest, and then use the **Link an Existing GPO** option to link it to the site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy objects overview for GPMC
Local Group Policy objects overview for GPMC
Start Group Policy Management Console

# Filter using security groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To filter using security groups**

1. Open Group Policy Management.
2. In the console tree, double-click **Group Policy Objects**, which contains the Group Policy object (GPO) to which you want to apply security filtering (console tree location is: *Forest name*/Domains/*Domain name*/Group Policy Objects).
3. Click the GPO.
4. In the results pane, on the **Scope** tab, click **Add**.
5. In the **Enter the object name to select** box, type the name of the group, user, or computer that you want to add to the security filter. Click **OK**.

   If **Authenticated Users** appears in the **Security Filtering** section of the **Scope** tab, select this group and click **Remove**. This will ensure that only members of the group or groups you added in step 5 can receive the settings in this GPO.

## Notes

- You must have **Edit settings, delete, and modify security** permissions on the GPO to perform these procedures.
- The settings in a GPO will apply only to users and computers that are contained in the domain, organizational unit, or organizational units where the GPO is linked, and that are specified in, or are members of a group that are specified in **Security Filtering**. You can specify multiple groups, users or computers in the security filter for a single GPO.
- You can also perform this procedure or a related task by using one or more of the sample scripts included with Group Policy Management. For more information, see "Scripting Group Policy tasks" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Security filtering using GPMC
Create, import, export, copy, and paste WMI filters
WMI filtering using GPMC
Scripting Group Policy tasks using GPMC
Start Group Policy Management Console

# Link a WMI filter to a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To link a WMI filter to a Group Policy object**

1. Open Group Policy Management.
2. In the console tree, double-click **Group Policy Objects** in the forest and domain containing the Group Policy object (GPO) to which you want to link a WMI filter.

   **Where?**

   *Forest name*/Domains/*Domain name*/Group Policy Objects

3. Click the GPO.
4. In the results pane, on the **Scope** tab, under **WMI Filtering**, select a WMI filter from the drop-down list box.
5. When prompted to confirm the selection, click **Yes**.

**Notes**

- You can also link a WMI filter to a GPO using one of the following methods:

  1. In the console tree, find the WMI filter you want to link to a GPO.

     **Where?**

     *Forest name*/Domains/*Domain name*/WMI Filters/*WMI filter name*

  2. Do one of the following:

     Click the WMI filter.

     In the details pane, on the **General** tab, under **GPOs that use this WMI filter**, right-click, and then click **Add**.

     Under **Group Policy objects**, click the GPO for which you want to link the WMI filter, and then click **OK**.

     -or-

     In the console tree, drag and drop the WMI filter to the GPO you want to link it to.

- To link a WMI filter to a GPO, you must have Edit permissions on a GPO.
- Only one WMI filter can be linked to a GPO. If you try to link a WMI filter to a GPO that already has another WMI filter linked to it, you will be asked whether you want to replace the old filter with the new one.
- You can only link a WMI filter to a GPO in the same domain.
- **WMI Filters** (as well as **WMI Filtering** on the **Scope** tab for GPOs) is only available if at least one domain controller in the domain is running Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Create, import, export, copy, and paste WMI filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create, import, export, copy, and paste WMI filters**


- To create a WMI filter
- To import a WMI filter
- To export a WMI filter
- To copy and paste a WMI filter


**To create a WMI filter**

1. Open Group Policy Management.
2. In the console tree, right-click **WMI Filters** in the forest and domain in which you want to create a WMI filter.

   **Where?**

   *Forest name*/Domains/*Domain name*/WMI Filters

3. Click **New**.
4. In the **New WMI Filter** dialog box, type a name for the new WMI Filter in the **Name** box, type a description of the filter in the **Description** box.
5. Click **Add**.
6. In the **WMI Query** dialog box, either leave the default namespace (root\CIMv2) or enter another namespace by doing one of the following:

   - In the **Namespace** box, type the name of the namespace that you want to use for the WMI query.
   - Click **Browse**, select a namespace from the list, and then click **OK**.

7. Type a WMI query in the **Query** box, and then click **OK**.
8. To add more queries, repeat steps 5 through 7 to add each query.
9. After adding all queries, click **Save**.

**Notes**

- You must have privileges to create WMI filters in the domain in which you want to create the filter. By default the Domain Administrators, Enterprise Administrators, and Group Policy Creator Owners groups have this permission.
- **WMI Filters** is only available if at least one domain controller in the domain is running Windows Server 2003. The same is true for **WMI Filtering** on the **Scope** tab for Group Policy objects (GPOs).
- To delete a WMI filter, in the console tree, right-click the WMI filter and then click **Delete**. When asked to confirm

whether you want to delete the WMI filter, click **Yes**.

**Where?**

*Forest name*/Domains/*Domain name*/WMI Filters/*WMI filter name*

#### To import a WMI filter

1. Open Group Policy Management.

2. In the console tree, right-click **WMI Filters** in the forest and domain into which you want to import a WMI filter.

   **Where?**

   *Forest name*/Domains/*Domain name*/WMI Filters

3. Click **Import**.

4. In the **Import WMI Filter** dialog box, enter the .mof file that contains the WMI filter you want to import, and then click **Open**.

5. In the **Import WMI Filter** dialog box, specify the appropriate information in the **Name** and **Description** boxes, and then do one of the following:

   - To import the WMI filter without changing it, click **Import**.

   - To change the WMI filter and then import it, click **Edit**, make the appropriate changes to the namespace or query, click **OK**, and then click **Import**.

### Notes

- You must have privileges to create WMI filters in the domain in which you want to create the filter to complete this procedure. By default the Domain Administrators, Enterprise Administrators, and Group Policy Creator Owners groups have this permission.

- **WMI Filters** is only available if at least one domain controller in the domain is running WindowsServer 2003. The same is true for **WMI Filtering** on the **Scope** tab for Group Policy objects (GPOs).

- Importing a WMI filter creates a new filter instead of modifying an existing filter.

#### To export a WMI filter

1. Open Group Policy Management.

2. In the console tree, find the WMI filter that you want to export.

   **Where?**

   *Forest name*/Domains/*Domain name*/WMI Filters/*WMI filter name*

3. Right-click the WMI filter that you want to export, and then click **Export**.

4. In the **Export WMI filter** dialog box, select a secure file system location from the **Save in** drop-down list box, enter a name for the filter that you want to export, and then click **Save**.

   **Important**

- To secure exported WMI filters, ensure that only authorized administrators have permission to access the folder to which you are exporting the filter.

- To maximize security, you should store WMI filters in a secure file system location on servers running Windows 2000 Server or Windows Server 2003.

### Notes

- **WMI Filters** is only available if at least one domain controller in the domain is running Windows Server 2003. The same is true for **WMI Filtering** on the **Scope** tab for Group Policy objects (GPOs).

- You can also export a WMI filter by clicking **WMI Filters** and, in the results pane, clicking the **Contents** tab, right-clicking the filter name, and then clicking **Export**.

- Exported WMI filters are stored in the file system as .mof files, which use a general-purpose format for storing various types of WMI data. Group Policy Management can import only .mof files that contain a single WMI filter.

#### To copy and paste a WMI filter

1. Open Group Policy Management.

2. In the console tree, click the WMI filter in the forest and domain containing the WMI filter that you want to copy.

   **Where?**

   *Forest name*/Domains/*Domain name*/WMI Filters/*WMI filter name*

3. Right-click the WMI filter you want to copy, and then click **Copy** to copy the filter to the clipboard.

4. In the destination domain, right-click **WMI Filters**, and then click **Paste**.

### Notes

- You must have privileges to create WMI filters in the domain in which you want to create the filter to complete this procedure. By default the Domain Administrators, Enterprise Administrators, and Group Policy Creator Owners groups have this permission.

- **WMI Filters** is only available if at least one domain controller in the domain is running Windows Server 2003. The same is true for **WMI Filtering** on the **Scope** tab for Group Policy objects (GPOs).

**See Also**
**Concepts**
WMI filtering using GPMC
Delegate creation of WMI filters using GPMC
Delegate an individual WMI filter using GPMC
Filter using security groups
Start Group Policy Management Console

# Enforce a Group Policy object link using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enforce a Group Policy object link**

1. Open Group Policy Management.
2. In the console tree, double-click the forest containing the domain, site, or organizational unit containing the link for which you want to change enforcement of a Group Policy object (GPO) link, and then do one of the following:
   - To enforce a GPO link to a domain, double-click **Domains**, and then double-click the domain containing the GPO link.
   - To enforce a GPO link to an organizational unit, double-click **Domains**, double-click the domain containing the organizational unit, and then double-click the organizational unit (including any parent or child organizational unit) containing the GPO link.
   - To enforce a GPO link to a site, double-click **Sites**, and double-click the site containing the GPO link.
3. Right-click the GPO link, and then click **Enforced** to enable or disable enforcement. A check mark next to **Enforced** indicates that the link is enforced.

## Notes

- You must have **Link GPOs** permission for the domain, site, or organizational unit containing the GPO link to complete this procedure.
- To determine whether a GPO link is enforced, you can also click the GPO link and look at the information in the **Links** section of the **Scope** tab.
- You can also enforce a GPO link by clicking on the GPO link in the console tree, right-clicking the link under **Links** on the **Scope** tab in the results pane, clicking **Enforced** to enable or disable enforcement for the link, and then clicking **OK**.
- **Enforced** is referred to as **No Override** on computers where Group Policy Management Console has not been installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Link a Group Policy object using GPMC
Block inheritance using GPMC
Disable a Group Policy object link using GPMC
Disable the User or Computer settings in a Group Policy object using GPMC
Managing inheritance of Group Policy
Start Group Policy Management Console

# Block inheritance using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To block inheritance**

1. Open Group Policy Management.
2. In the console tree, double-click the forest containing the domain or organizational unit for which you want to block inheritance for Group Policy object (GPO) links, and then do one of the following:
   - To block inheritance of the GPO links of an entire domain, double-click **Domains**, and then right-click the domain.
   - To block inheritance for an organizational unit, double-click **Domains**, double-click the domain containing the organizational unit, and then right-click the organizational unit.
3. Click **Block Inheritance**.

## Notes

- You must have **Link GPOs** permission for the domain or organizational unit to complete this procedure.
- If a domain or organizational unit is set to block inheritance, it will appear with a blue exclamation mark in the console tree.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Link a Group Policy object using GPMC
Enforce a Group Policy object link using GPMC
Disable a Group Policy object link using GPMC
Disable the User or Computer settings in a Group Policy object using GPMC
Managing inheritance of Group Policy
Start Group Policy Management Console

# Get user settings from the machine account (loopback processing)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To get user settings from the machine account (loopback processing)**

1. Open Group Policy Management.
2. In the console tree, double-click the forest containing the domain or organizational unit which contains the computer accounts.
3. Double-click the domain or organization unit.
4. Right-click a Group Policy object (GPO) that is linked to the domain or organizational unit, and then click **Edit**. This opens Group Policy Object Editor. If there is no existing GPO, use the **Create and Link a GPO Here** option.
5. In the results pane of the Group Policy Object Editor, double-click **User Group Policy loopback processing mode**.

   Where?

   - *GPO name*
   - Computer Configuration
   - Administrative Templates
   - System
   - Group Policy
   - User Group Policy loopback processing mode
6. Click **Enabled**, and then in the **Mode** box, select either **Merge** or **Replace**. Click **OK**.

### Notes

- You must have **Edit** permissions on the GPO to perform this procedure.
- The Loopback option applies to all GPOs applied to a given computer. It cannot be specified on a per GPO basis.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Loopback processing with merge or replace
Group Policy object editor overview for GPMC
Start Group Policy Management Console

# Disable a Group Policy object link using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable a Group Policy object link**

1. Open Group Policy Management.
2. In the console tree, double-click the forest containing the domain, site, or organizational unit containing the Group Policy object (GPO) link that you want to disable, and then do one of the following:
   - To disable a GPO link to a domain, double-click **Domains**, and then double-click the domain containing the GPO link.
   - To disable a GPO link to an organizational unit, double-click **Domains**, double-click the domain containing the organizational unit, and then double-click the organizational unit (including any parent or child organizational unit) containing the GPO link.
   - To disable a GPO link to a site, double-click **Sites**, and double-click the site containing the GPO link.
3. Right-click the GPO link. A check mark next to **Link Enabled** indicates that the link is currently enabled.
4. Click **Link Enabled** to remove the check mark and disable the link.

## Notes

- You must have **Link GPOs** permission for the domain, site, or organizational unit containing the GPO link to complete this procedure. If you do not have **Link GPOs** permission, **Link Enabled** is unavailable.
- To determine whether a GPO link is enabled, you can also click the GPO link and look at the information in the **Links** section of the **Scope** tab.
- You can also disable a GPO link by clicking on the GPO link in the console tree, right-clicking the link under **Links** on the **Scope** tab in the results pane, clicking **Link Enabled** to enable or disable the link, and then clicking **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Link a Group Policy object using GPMC
Enforce a Group Policy object link using GPMC
Disable the User or Computer settings in a Group Policy object using GPMC
Managing inheritance of Group Policy
Start Group Policy Management Console

# Disable the User or Computer settings in a Group Policy object using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable the user or computer settings for a Group Policy object**

1. Open Group Policy Management.
2. In the console tree, double-click the forest containing the domain or organizational unit which contains the Group Policy object (GPO).
3. Double-click the domain or organizational unit.
4. Right-click the GPO that contains the user or computer settings you want to disable, point to **GPO Status**, and then do one of the following:
   - Click **User settings disabled** to disable user settings for the GPO. A check mark next to **User settings disabled** indicates that user settings are disabled.
   - Click **Computer settings disabled** to disable computer settings for the GPO. A check mark next to **Computer settings disabled** indicates that computer settings are disabled.

## Notes

- You must have **Edit** permission on the GPO to complete this procedure.
- You can also modify GPO Status from the **Details** tab of a GPO.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable a Group Policy object link using GPMC
Link a Group Policy object using GPMC
Managing inheritance of Group Policy
Scripting Group Policy tasks using GPMC
Start Group Policy Management Console

# Delegate Group Policy tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Delegate Group Policy tasks**

- Delegate creation of Group Policy objects using GPMC
- Delegate an individual Group Policy object using GPMC
- Delegate policy-related permissions on a domain, OU, or site using GPMC
- Delegate creation of WMI filters using GPMC
- Delegate an individual WMI filter using GPMC

# Delegate creation of Group Policy objects using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate creation of Group Policy objects**

1. Open Group Policy Management.
2. In the console tree, click **Group Policy Objects** in the forest and domain for which you want to delegate creation rights for Group Policy objects (GPOs).

   **Where?**

   - *Forest name*/Domains/*Domain name*/Group Policy Objects

3. In the results pane, click the **Delegation** tab, and then do one of the following.

   - To add a new group or user, using the Group Policy Creator Owners group
   - To add a new group or user to have the same permissions as the Group Policy Creator Owners Group
   - To remove creation rights from a group or user

**To add a new group or user, using the Group Policy Creator Owners group**

1. In the **Groups and Users** list box, double-click **Group Policy Creator Owners**.
2. In the **Group Policy Creator Owners** properties dialog box, select the **Members** tab, and click **Add**.
3. In the **Select User, Computer, or Group** dialog box, click **Object Types**, select the types of objects for which you want to delegate creation rights, and then click **OK**.
4. Click **Locations**, select either **Entire Directory** or the domain or organizational unit containing the object for which you want to delegate creation rights, and then click **OK**.
5. In the **Enter the object name to select** box, enter the name of the object to which you want to delegate creation rights by doing one of the following:

   - If you know the name, type it, and then click **OK**.
   - To search for the name, click **Advanced**, enter the search criteria, click **Find Now**, select the name in the list box, click **OK**, and then click **OK**.

**To add a new group or user to have the same permissions as the Group Policy Creator Owners group**

1. Click **Add**.
2. In the **Select User, Computers, or Groups** dialog box, click **Object Types**, select the types of objects to which you want to delegate creation rights for GPOs, and then click **OK**.
3. Click **Locations**, select either **Entire Directory** or the domain or organizational unit containing the object to which you want to delegate creation rights, and then click **OK**.

4. In the **Enter the object name to select** box, enter name of the object to which you want to delegate creation rights by doing one of the following:

    - If you know the name, type it, and then click **OK**.

    - To search for the name, click **Advanced**, enter the search criteria, click **Find Now**, select the name in the list box, click **OK**, and then click **OK**.

**To remove creation rights from a group or user**

1. Do one of the following:

    - To remove a name from the **Groups and users** list box, delete the name of the group or user, and then click **Remove**.

    - To remove a member of Group Policy Creator Owners group, double-click **Group Policy Creator Owners**, click the **Members** property tab, select the name to be removed, and then click **Remove**.

2. When prompted to confirm the removal of the delegation privilege, click **Yes**, and then click **OK**.

## Notes

- You must be a member of the Domain Administrators or Enterprise Administrators to complete this procedure.

- You can also remove a group or user from the permissions list by right-clicking the name of the group or user in the **Groups and users** list box on the **Delegation** tab, and then clicking **Remove**.

- If you want to delegate permissions to groups and users in the same domain as the GPOs, it is recommended that you add them to the Group Policy Creator Owners group since this is the original method that was available in Windows 2000. However, you cannot add groups or users from another domain to the Group Policy Creator Owners group.

- To delegate GPO creation permissions to groups or users in another domain, you must explicitly grant that group GPO creation permissions (without using Group Policy Creator Owners group). Alternatively you can create a domain local group in the domain where you want to delegate GPO creation permission and grant that group permission to create GPOs. Then add members to this domain local group from other domains as needed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Scripting Group Policy tasks using GPMC
Delegate policy-related permissions on a domain, OU, or site using GPMC
Delegation and policy-related permissions
Start Group Policy Management Console
Delegate an individual Group Policy object using GPMC

# Delegate an individual Group Policy object using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate an individual Group Policy object**

1. Open Group Policy Management.

2. In the console tree, double-click **Group Policy Objects** in the forest and domain containing the Group Policy object (GPO) for which you want to add or remove permissions.

   **Where?**

   *Forest name*/Domains/*Domain name*/Group Policy Objects

3. Click the GPO.

4. In the results pane, click the **Delegation** tab, and then do one of the following:

   - To delegate permissions for a group or user
   - To change permissions for a group or user
   - To remove permissions for a group or user

**To delegate permissions for a group or user**

1. Click **Add**.

2. In the **Select User, Computer, or Group** dialog box, click **Object Types**, select the types of objects for which you want to add GPO permissions, and then click **OK**.

3. Click **Locations**, select either **Entire Directory** or the domain or organizational unit containing the object for which you want to add GPO permissions, and then click **OK**.

4. In the **Enter the object name to select** box, enter the name of the object for which you want to add GPO permissions by performing one of the actions in the following table.

   - If you know the name, type it, and then click **OK**.
   - To search for the name, click **Advanced**, enter the search criteria, click **Find Now**, select the name in the list box, click **OK**, and then click **OK**.
   - In the **Permissions** box of the **Add Group or User** dialog box, select the appropriate permissions from the drop down list, and then click **OK**.

**To change permissions for a group or user**

1. In the **Groups and users** list box, right-click the name of the group or user for which you want to change permissions, and then click the appropriate permission level: **Read only**, **Edit Settings**, or **Edit, delete, modify security**.

2. When prompted to confirm the change in delegation permissions, click **OK**.

**To remove permissions for a group or user**

1. In the **Groups and users** list box, right-click the name of the group or user for which you want to remove permissions, and then click **Remove**.

2. When prompted to confirm the removal of the delegation privilege, click **OK**.

## Notes

- You must have **Edit settings, delete, and modify security** permissions on the GPO to perform these procedures.

- Groups and users that have **Custom** in the **Allowed Permissions** column in the **Groups and users** list box on the **Delegation** tab have permissions that do not match one of the three standard levels of permissions. To view the permissions for groups with custom permissions or to set custom permissions, click **Advanced**. To change permissions to one of the standard levels, follow the steps above for *change permissions for a group or user*.

- At least one group or user must have **Edit settings, delete, and modify security** permissions on each GPO. You cannot remove a group or user if it is the only group or user on a GPO that has this permission set.

- Permissions inherited from parent containers cannot be removed.

- You can also perform this procedure or a related task by using one or more of the sample scripts included with Group Policy Management. For more information, see "Scripting Group Policy tasks" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Scripting Group Policy tasks using GPMC
Delegate creation of Group Policy objects using GPMC
Delegate policy-related permissions on a domain, OU, or site using GPMC
Delegation and policy-related permissions
Start Group Policy Management Console

# Delegate policy-related permissions on a domain, OU, or site using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Delegate policy-related permissions on a domain, organizational unit, or site**

To delegate permissions to link Group Policy objects

1. Open Group Policy Management.
2. In the console tree, do one of the following:
   - To delegate permission to link Group Policy objects (GPOs) to either the domain or an organizational unit, click the domain or the organizational unit (console tree location is: *Forest name*/Domains/*Domain name*/Organizational unit).
   - To delegate permission to link GPOs to a site, click the site (console tree location is: *Forest name*/Sites/*Site name*).
3. In the results pane, click the **Delegation** tab.
4. In the **Permission** drop down list box, select **Link GPOs**, and then specify permissions to link GPOs for a group or user by doing one of the following:
   - To add a new group or user to the permissions list for domains and organizational units
   - To add a new group or user to the permissionslist for sites
   - To change the inheritance of permissions on a domain or organizational unit for a group or user
   - To remove a group or user from the permissions list
   - To delegate permissions for generating Group Policy Modeling data
   - To delegate permissions to read Group Policy Results data

**To add a new group or user to the permissions list for domains and organizational units**

1. On the **Delegation** tab, click **Add**.
2. In the **Select User, Computer, or Group** dialog box, click **Object Types**, select the types of objects to which you want to delegate permissions for the domain, site, or organizational unit, and then click **OK**.
3. Click **Locations**, select either **Entire Directory** or the domain or organizational unit containing the object to which you want to delegate permissions, and then click **OK**.
4. In the **Enter the object name to select** box, enter name of the object to which you want to delegate permissions by doing one of the following:
   - If you know the name, type it, and then click **OK**.
   - To search for the name, click **Advanced**, enter the search criteria, click **Find Now**, select the name in the list box,

click **OK**, and then click **OK**.

5. In the **Add Group or User** dialog box, in the **Permissions** drop down list box, select the level to which you want permissions to apply for this group or user, and then click **OK**.

**To add a new group or user to the permissions list for sites**

1. On the **Delegation** tab, click **Add**.

2. In the **Select User, Computer, or Group** dialog box, click **Object Types**, select the types of objects to which you want to delegate permissions for the domain, site, or organizational unit, and then click **OK**.

3. Click **Locations**, select either **Entire Directory** or the domain or organizational unit containing the object to which you want to delegate permissions, and then click **OK**.

4. In the **Enter the object name to select** box, enter name of the object to which you want to delegate permissions by doing one of the following:

   - If you know the name, type it, and then click **OK**.

   - To search for the name, click **Advanced**, enter the search criteria, click **Find Now**, select the name in the list box, click **OK**, and then click **OK**.

**To change the inheritance of permissions on a domain or organizational unit for a group or user**

- On the **Delegation** tab of the domain or organizational user, in the **Groups and users** list box, right-click the name of the group or user, and then click **This container only** or **This container and children** to specify the level to which you want permissions to apply for this group or user.

**To remove a group or user from the permissions list**

1. On the **Delegation** tab, in the **Groups and users** list box, select the name of the group or user to be removed, and then click **Remove**.

2. When prompted to confirm the removal, click **OK**.

## Notes

- To delegate permissions to Link GPOs to a site, domain or organizational unit, you must have **Modify Permissions** on that site, domain or organizational unit. By default, only Domain Administrators, Enterprise Administrators, have this permission.

- Users and groups with permission to link GPOs to a specific site, domain, or organizational unit can link GPOs, change link order, and set block inheritance on that site, domain, or organizational unit.

- You cannot remove groups and users that inherit permissions from a parent container.

- Some entries in the **Groups and users** list box, such as **System**, do not have an associated property dialog box, so **Properties** is unavailable for these entries.

**To delegate permissions for generating Group Policy Modeling data**

1. Open Group Policy Management.

2. In the console tree, click the domain or organizational unit for which you want to delegate Group Policy Modeling

permissions.

**Where?**

- *Forest name*/Domains/*Domain name*/Organizational unit (if applicable)

3. In the results pane, click the **Delegation** tab.

4. In the **Permission** box, select **Perform Group Policy Modeling analyses**, and then do one of the following:

| | |
|---|---|
| **To add a new group or user to the permissions list** | a. On the **Delegation** tab, click **Add**.<br><br>b. In the **Select User, Computer, or Group** dialog box, click **Object Types**, select the types of objects to which you want to delegate permissions for the domain, site, or organizational unit, and then click **OK**.<br><br>c. Click **Locations**, select either **Entire Directory** or the domain or organizational unit containing the object to which you want to delegate permissions, and then click **OK**.<br><br>d. In the **Enter the object name to select** box, enter name of the object to which you want to delegate permissions by doing one of the following:<br><br>• If you know the name, type it, and then click **OK**.<br><br>• To search for the name, click **Advanced**, enter the search criteria, click **Find Now**, select the name in the list box, click **OK**, and then click **OK**.<br><br>• In the **Add Group or User** dialog box, in the **Permissions** drop down list box, select the level to which you want permissions to apply for this group or user, and then click **OK**. |
| **To change the inheritance for a group or user in the permissions list** | On the **Delegation** tab, in the **Groups and users** list box, right-click the name of the group or user, and then click **This container only** or **This container and children** to specify the level to which you want permissions to apply for this group or user. |
| **To remove a group or user from the permissions list** | a. On the **Delegation** tab, in the **Groups and users** list box, select the name of the group or user to be removed, and then click **Remove**.<br><br>b. When prompted to confirm the removal, click **OK**. |

**Notes**

- To delegate permissions to perform Group Policy Modeling analyses for objects in a domain or organizational unit, you must have **Modify Permissions** on that domain or organizational unit. By default, only Domain Administrators and Enterprise Administrators have this permission.

- You cannot delegate permission to perform Group Policy Modeling analyses for sites.

- In earlier versions of Group Policy, Group Policy Modeling was referred to as Resultant Set of Policy (RSoP) planning mode.

**To delegate permissions to read Group Policy Results data**

1. Open Group Policy Management.

2. In the console tree, click the domain or organizational unit for which you want to delegate permission to read Group Policy Results data.

   **Where?**

   - Group Policy Management Console/*Forest name*/Domains/*Domain name*/Organizational unit (if applicable)

3. In the results pane, click the **Delegation** tab.

4. In the **Permissions** dropdown list, select **Read Group Policy Results data**, and then do one of the following:

| | |
|---|---|
| **To add a new group or user to the permissions list** | a. On the **Delegation** tab, click **Add**.<br>b. In the **Select User, Computer, or Group** dialog box, click **Object Types**, select the types of objects to which you want to delegate permissions for the domain, site, or organizational unit, and then click **OK**.<br>c. Select the user or group to who permission should be delegated.<br>d. In the **Add Group or User** dialog box, in the **Permissions** drop down list box, select the level to which you want permissions to apply for this group or user, and then click **OK**. |
| **To change inheritance for a group or user in the permissions list** | On the **Delegation** tab, in the **Groups and users** list box, right-click the name of the group or user, and then click **This container only** or **This container and children** to specify the level to which you want permissions to apply for this group or user. |
| **To remove a group or user from the permissions list** | a. On the **Delegation** tab, in the **Groups and users** list box, select the name of the group or user to be removed, and then click **Remove**.<br>b. When prompted to confirm the removal, click **OK**. |

**Notes**

- To delegate permissions to read Group Policy Results data for objects in a domain or organizational unit, you must have **Modify Permissions** on that domain or organizational unit. By default, only Domain Administrators and Enterprise Administrators have this permission.

- You cannot delegate permission to read Group Policy Results data for sites.

- In earlier versions of Group Policy, Group Policy Results was referred to as Resultant Set of Policy (RSoP) logging mode.

**See Also**
**Concepts**
Link a Group Policy object using GPMC
Delegation and policy-related permissions
Scripting Group Policy tasks using GPMC
Security settings overview for GPMC
Security filtering using GPMC
Filter using security groups
Simulate Resultant Set of Policy using Group Policy Modeling
Determine Resultant Set of Policy with Group Policy Results
Start Group Policy Management Console

# Delegate creation of WMI filters using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate creation of WMI filters**

1. Open Group Policy Management.

2. In the console tree, click **WMI Filters** in the forest and domain in which you want to delegate management permissions for all WMI filters.

   **Where?**

   *Forest name*/Domains/*Domain name*/WMI Filters

3. In the results pane, click the **Delegation** tab, and then do one of the following:

   - To add a new group or user with management permissions on all WMI filters
   - To change the permission levels on all WMI filters for a group or user
   - To remove a group or user from the permissions list for all WMI filters

**To add a new group or user with management permissions on all WMI filters**

1. Click **Add**.

2. In the **Select Users, Computers, or Groups** dialog box, click **Object Types**, select the types of objects to which you want to delegate permissions on all WMI filters, and then click **OK**.

3. Click **Locations**, select either **Entire Directory** or the domain or organizational unit containing the object to which you want to delegate permissions, and then click **OK**.

4. In the **Enter the object name to select** box, enter name of the object to which you want to delegate permissions by doing one of the following:

   a. If you know the name, type it, and then click **OK**.

   b. To search for the name, click **Advanced**, enter the search criteria, click **Find Now**, select the name in the list box, click **OK**, and then click **OK**.

5. In the **Add Group or User** dialog box, select the permissions level you want to assign to the group or user, and then click **OK**.

**To change the permission levels on all WMI filters for a group or user**

- In the list box, right-click the name of the group or user, and then click **Creator Owner** or **Full Control** to specify the permissions level you want to assign to the group or user.

**To remove a group or user from the permissions list for all WMI filters**

- In the **Groups and users** list box, select the name of the group or user to have permissions on all WMI filters removed, and then click **Remove**.

- or -

1. If the group or user is a member of another group with these permissions, double-click the group it belongs to. In the property dialog box that appears, click the group or user you want to remove, and then click **Remove**.

2. When prompted to confirm the removal of the delegation privilege, click **OK**.

### Notes

- You must be a domain administrator or Enterprise Administrator to delegate permissions on all WMI filters in a domain.

- Users with Full control permissions can create and control all WMI filters in the domain, including WMI filters created by others. Users with Creator owner permissions can create WMI filters, but can only control WMI filters that they create.

- If you remove Group Policy Creator Owners from the permissions list, users who create Group Policy objects (GPOs) can no longer create WMI filters unless they are explicitly given that permission through membership in another group.

- All users must have Read access to all WMI filters. Otherwise, Group Policy stops processing when it encounters a WMI filter that cannot be read. You cannot use Group Policy Management to remove Read permissions from WMI filters.

- **WMI Filters** is only available if at least one domain controller in the domain is running Windows Server 2003. The same is true for **WMI Filtering** on the **Scope** tab for GPOs.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delegation and policy-related permissions
WMI filtering using GPMC
Delegate an individual WMI filter using GPMC
Delegate creation of Group Policy objects using GPMC
Delegate an individual Group Policy object using GPMC
Delegate policy-related permissions on a domain, OU, or site using GPMC
Start Group Policy Management Console

# Delegate an individual WMI filter using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate an individual WMI filter**

1. Open Group Policy Management.
2. In the console tree, click the WMI filter you want to delegate permissions for.

   **Where?**

   *Forest name*/Domains/*Domain name*/WMI Filters/*WMI filter name*

3. In the results pane, click the **Delegation** tab, and then do one of the following:

   - To add a new group or user with management permissions on the WMI filter
   - To change the permissions levels on the WMI filter for a group or user in the permission list
   - To remove a group or user from the permissions list for the WMI filter

**To add a new group or user with management permissions on the WMI filter**

1. Click **Add**.
2. In the **Select User, Computer, or Group** dialog box, click **Object Types**, select the types of objects to which you want to delegate permissions on the WMI filter, and then click **OK**.
3. Click **Locations**, select either **Entire Directory** or the domain or organizational unit containing the object to which you want to delegate permissions, and then click **OK**.
4. In the **Enter the object name to select** box, enter name of the object to which you want to delegate permissions by doing one of the following:

   - If you know the name, type it, and then click **OK**.
   - To search for the name, click **Advanced**, enter the search criteria, click **Find Now**, select the name in the list box, click **OK**, and then click **OK**.

5. In the **Add Group or User** dialog box, in the **Permissions** box, select the permissions level you want to assign to the group or user, and then click **OK**.

**To change the permissions levels on the WMI filter for a group or user in the permission list**

1. In the list box, right-click the name of the group or user, and then click **Edit** or **Full Control** to specify the permissions level you want to assign to that group or user.
2. When prompted to confirm the change, click **OK**.

**To remove a group or user from the permissions list for the WMI filter**

1. In the list box, select the name of the group or user for which you want to have permissions removed, and then click

**Remove**.

2. When prompted to confirm the removal of the delegation privilege, click **OK**.

### Notes

- You must have Full Control permissions on a WMI filter to change its permissions.
- You cannot remove or change inherited permissions for WMI filters.
- All users must have Read access to all WMI filters. Otherwise, Group Policy stops processing when it encounters a WMI filter that cannot be read.
- You cannot use Group Policy Management to remove Read permissions from WMI filters.
- To open Group Policy Management, click **Start,** click **Programs,** click **Administrative Tools,** and then click **Group Policy Management**.
- **WMI Filters** is only available if at least one domain controller in the domain is running Windows Server 2003.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delegation and policy-related permissions
WMI filtering using GPMC
Delegate creation of WMI filters using GPMC
Delegate creation of Group Policy objects using GPMC
Delegate an individual Group Policy object using GPMC
Delegate policy-related permissions on a domain, OU, or site using GPMC
Start Group Policy Management Console

# Use Script and Command Line Methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Script and Command Line Methods**


- Run a script for a Group Policy-related task
- Refresh Group Policy settings with GPUpdate.exe
- Determine Resultant Set of Policy with GPResult.exe

# Run a script for a Group Policy-related task

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run a script for a Group Policy related task**

- From the command line, enter the path and name of the script with the appropriate parameters. If cscript is not your default scripting host, be sure to precede this script name with "cscript" and enclose the path name in quotes if it contains spaces.

## Notes

- Group Policy Management Console (GPMC) provides over 30 sample scripts that you can use to perform various administrative tasks. All of the scripts are located in **%programfiles%\gpmc\scripts** directory on any computer where you installed GPMC.
- All of these scripts execute from the command line. Executing a script with the /? switch displays the usage for that script.
- The sample scripts echo output to the command window and must be executed using cscript.exe. If cscript.exe is not your default scripting host, you will need to explicitly specify cscript.exe on the command line. For example, to run the ListAllGPOs.wsf script, type the following at the command line: "cscript ListAllGPOs.wsf".
- You can run "cscript //H:cscript" from a command line to make cscript the default scripting host.
- These sample scripts are based on the COM interfaces that are provided by Group Policy Management Console. The interfaces are documented in the Group Policy Management Console SDK, which is located at **%program%files\gpmc\scripts\gpmc.chm** on any computer where you installed GPMC. The GPMC SDK is also available in the Platform SDK. The GPMC SDK is only available in English.
- You can modify and distribute the sample scripts in accordance with the terms of the End User License Agreement.
- For more information, see Scripting Group Policy tasks using GPMC.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Refresh Group Policy settings with GPUpdate.exe

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To refresh Group Policy settings with GPUpdate**

The **gpupdate** command refreshes local and Active Directory–based Group Policy settings, including security settings on the computer from where it is run. You can use **gpupdate** locally on Windows XP and higher computers to refresh policy immediately. On computers running Windows 2000, this functionality is provided by the using the **secedit** command with the refreshpolicy option.

**Syntax**

**Gpupdate** [**/target:**{**computer**|**user**}] [**/force**] [**/wait:***value*] [**/logoff**] [**/boot**]

**Parameters**

**/target:{computer|user}**

Processes only the *computer* settings or the current *user* settings. By default, both the computer settings and the user settings are processed.

**/force**

Ignores all processing optimizations and reapplies all settings. The Group Policy engine on the client tracks versions of the GPOs that are applied to the user and computer. By default, if none of the GPO versions change and the list of GPOs remains the same, the Group Policy engine will not reprocess policy. This option overrides this optimization and forces the Group Policy engine to reprocess all policy information.

**/wait: *value***

Number of seconds that policy processing waits to finish. The default is 600 seconds. *0* means "no wait"; *-1* means "wait indefinitely."

**/logoff**

Logs off after the refresh has completed. This is required for those Group Policy client-side extensions that do not process on a background refresh cycle but that do process when the user logs on, such as user Software Installation and Folder Redirection. This option has no effect if there are no extensions called that require the user to log off.

**/boot**

Restarts the computer after the refresh has completed. This is required for those Group Policy client-side extensions that do not process on a background refresh cycle but that do process when the computer starts up, such as computer Software Installation. This option has no effect if there are no extensions called that require the computer to be restarted.

**/?**

Displays help at the command prompt.

**Remarks**

- If you make a syntax error, a syntax summary appears that looks similar to this Help topic.

**Examples**

The following examples show how you can use the **gpupdate** command:

- **gpupdate**
- **gpupdate /target:computer**
- **gpupdate /force /wait:**100
- **gpupdate /boot**

**Formatting legend**

| Format | Meaning |
|---|---|

| | |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Determine Resultant Set of Policy with GPResult.exe

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To determine resultant set of policy with GPResult**

Gpresult is a command line tool and must be run from the command prompt. The **gpresult** command displays Resultant Set of Policy (RSOP) for a user or a computer. You can use gpresult to see what policy is in effect and to troubleshoot problems on Windows XP and higher computers.

**Important**

The gpresult.exe that ships with Windows XP and Windows Server 2003 family is completely different than the original version of gpresult.exe that shipped in the Windows 2000 Resource Kit. The newer version cannot be used to view policy information for computers running Windows 2000.

**Syntax**

**Gpresult** [**/s** *computer* [**/u** *domain\user* **/p** *password*]] [**/user** *TargetUserName*] [**/scope** {**user**|**computer**}] [**/v**] [**/z**]

**Parameters**

***/s computer***

Specifies the name or IP address of a remote computer. (Do not use backslashes.) The default is the local computer.

**/u *domain\user***

Runs the command with the account permissions of the user that is specified by *user* or *domain\user*. The default is the permissions of the current logged-on user on the computer that issues the command.

**/p *password***

Specifies the password of the user account that is specified in the **/u** parameter.

**/user *TargetUserName***

Specifies the user name of the user whose RSOP data is to be displayed.

**/scope {user|computer}**

Displays either **user** or **computer** results. Valid values for the **/scope** parameter are **user** or **computer**. If you omit the **/scope** parameter, **gpresult** displays both **user** and **computer** settings.

**/v**

Specifies that the output display verbose policy information.

**/z**

Specifies that the output display all available information about Group Policy. Because this parameter produces more information than the **/v** parameter, redirect output to a text file when you use this parameter (for example, **gpresult /z >policy.txt**).

**/?**

Displays help at the command prompt.

**Remarks**

- Because you can apply overlapping levels of policies to any computer or user, the Group Policy feature generates a resulting set of policies when the user logs on to the network. **Gpresult** displays the resulting set of policies that were enforced on the computer when the specified user logged on.

**Examples**

The following examples show how you can use the **gpresult** command:

**gpresult /user targetusername /scope computer**

**gpresult /s srvmain /u maindom\hiropln /p p@ssW23 /user targetusername /scope USER**

**gpresult /s srvmain /u maindom\hiropln /p p@ssW23 /user targetusername /z >policy.txt**

**gpresult /s srvmain /u maindom\hiropln /p p@ssW23**

**Formatting legend**

| Format | Meaning |
| --- | --- |
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# GPMC Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section contains the following information:

- Group Policy Management Console Overview
- Group Policy objects overview for GPMC
- Local Group Policy objects overview for GPMC
- Group Policy object editor overview for GPMC
- Group Policy Object Editor Extensions
- Group Policy replication and domain controller selection
- Group Policy processing and precedence
- Controlling the Scope of Group Policy Objects using GPMC
- Scripting Group Policy tasks using GPMC
- Deployment considerations for Group Policy
- Group Policy Management Console icons reference
- GPMC Resources

# Group Policy Management Console Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Management Console Overview**
This section contains the following information:

- Domains in Group Policy Management Console
- Forests in Group Policy Management Console
- Backup, Restore, Import, Copy, and Migration Tables
- Delegation and policy-related permissions

# Domains in Group Policy Management Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domains in Group Policy Management Console**
Group Policy Management Console (GPMC) supports management of multiple domains at the same time, with each domain grouped by forest in the console.

For information about GPMC in a multi-forest environment, see Forests in Group Policy Management Console.

**Adding and removing domains**
By default, only a single domain is shown in GPMC. When you first start GPMC using either the pre-configured snap-in (gpmc.msc) or a custom MMC console, GPMC displays the domain that contains the user account you used to start GPMC. When you add a new forest, GPMC displays the domain that was entered in the **Add Forest** dialog box.

You can specify domains in each forest that you want to manage using GPMC by adding and removing the domains shown in the console.

- Adding domains
- Removing domains

**Adding domains**
Under each forest, right-click the **Domains** node and click **Show Domains**. In the **Show Domains** dialog box, select each domain that you want to display in GPMC.

To add domains in another forest, you must first add the new forest that contains that domain.

You can add externally trusted domains, even if you do not have forest trust with the entire forest. By default, you must have two-way trust between the domain you want to add and the domain of your user object. You can also add domains across a one-way trust by disabling the trust detection feature of GPMC, using the **Options** dialog box on the **View** menu. To add externally trusted domains, you must first use the **Add Forest** dialog box to add one domain from a forest that contains the externally trusted domains. Once this forest is added, you can add any domains in that forest that are trusted by right-clicking the Domains node of the forest and clicking **Show Domains**.

For step-by-step instructions, see Add a forest, site, or domain to the Group Policy Management Console.

**Removing domains**
Remove domains from the console by using the **Show Domains** dialog box, available by right-clicking the **Domains** node. In the **Show Domains** dialog box, clear check boxes if you want to remove them from GPMC. You can also right-click an existing domain in the console and click **Remove**.

**See Also**
**Concepts**
Forests in Group Policy Management Console
Add a forest, site, or domain to the Group Policy Management Console

# Forests in Group Policy Management Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Forests in Group Policy Management Console**
Group Policy Management Console (GPMC) supports management of multiple forests from within the console, when there is trust between the target forest and the forest of your user object.

### Notes

- GPOs cannot be linked outside a forest.
- GPMC also supports operations such as backup and import, to allow migrating GPOs between untrusted domains in separate forests.

**Adding and removing forests**
When you first open Group Policy Management, it displays only the forest containing the domain of your user account (you, the GPMC user). You can add additional existing forests, as long as there is trust between the forest of your user account and the other forest. Forest trust is a new feature of Windows Server 2003. You can also add externally trusted domains even if there is no forest trust, provided there is an explicit external trust.

You can enable support for adding forests and domains with one way trusts, by turning off the trust detection feature of GPMC, using the Options dialog box on the View menu. By default, GPMC requires two-way trust to the target forest, either via a two-way external trust between domains, or a two way forest trust between the forests.

It is not possible to add a forest to which you have no trust.

In GPMC, each forest shows the following child nodes: **Domains**, **Sites**, and **Group Policy Results**. Forests with the Windows Server 2003 schema will also show **Group Policy Modeling**.

- Adding forests
- Removing forests

**Adding forests**
Add forests one at a time, by right-clicking the Group Policy Management root node and clicking **Add Forest**. You specify either the NetBIOS name or the DNS name of a domain in the forest.

When adding a forest, the forest displayed in Group Policy Management is named after the forest root domain for that forest and includes a single domain, which is the domain that you specified in the **Add forest** dialog box. To show other domains, you can add them by right-clicking the **Domains** node and clicking **Show Domains**.

For step-by-step instructions to add a forest, see Add a forest, site, or domain to the Group Policy Management Console.

**Removing forests**
You remove forests from GPMC one at a time by right-clicking each forest node and clicking **Remove**.

**See Also**
**Concepts**
Domains in Group Policy Management Console
Add a forest, site, or domain to the Group Policy Management Console

# Backup, Restore, Import, Copy, and Migration Tables

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backup, Restore, Import, Copy, and Migration Tables**
With Group Policy Management Console (GPMC) you can back up, restore, import, or copy Group Policy objects (GPOs). When copying or importing GPOs across domains, you can use migration tables to facilitate the copy or import operation.

This section contains the following information:

- Backup using GPMC
- Restore using GPMC
- Import using GPMC
- Copy using GPMC
- Migration tables
- The migration table editor

# Backup using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backup**
Backing up a Group Policy object (GPO) copies the data in the GPO to the file system. The backup function also serves as the export capability for GPOs. A GPO backup can be used to restore the GPO to the backed-up state, or to import the settings in the backup to another GPO. For step-by-step instructions, see Back up a Group Policy object using GPMC.

**What is saved in a backup**
Backing up a GPO saves all information that is stored inside the GPO to the file system. This includes the following information:

- The GPO globally unique identifier (GUID) and domain
- GPO Settings
- The discretionary access control list (DACL) on the GPO
- The WMI filter link, if there is one, but not the filter itself
- Links to IP Security Policies, if any
- XML report of the GPO settings, which can be viewed as HTML from within GPMC
- Date and time stamp of when the backup was taken
- User-supplied description of the backup

**What is not saved in a backup**
Backing up a GPO only saves data that is stored inside the GPO. Data that is stored outside the GPO includes the following information:

- Links to a site, domain or organizational unit
- WMI filter
- IP Security policy

This data is not available when the backup is restored to the original GPO or imported into a new one.

**Managing Backups**
Multiple backups of the same or different GPO can be stored in the same file system location. Each backup is identified by a unique backup ID. The collection of backups in a given file system location can be managed using the **Manage Backups** dialog box in GPMC or through the scriptable interfaces. The **Manage Backups** dialog box is available by right-clicking either the **Domains** node or the **Group Policy Objects** node in a given domain. When you open **Manage Backups** from the Group Policy Objects node, the view is automatically filtered to show only backups of GPOs from that domain. When opened from the **Domains** node, the **Manage Backups** dialog box shows all backups, regardless of which domain they are from.

**Securing GPO backups**
To prevent unauthorized access and tampering, you should store GPO backups in a secure location with strong file system security (using discretionary access control lists, also known as DACLs) so that only authorized administrators have access to the backups.

For additional security, you can store these files on computers running Windows 2000 or later operating systems. This will prevent tampering with the data while in transit over the network. This is because all SMB traffic between computers running Windows 2000 and above is digitally signed by default. Furthermore, mutual authentication is used for communications between computers running Windows 2000 or later, which ensures that the data comes from a trusted source.

**See Also**

**Concepts**

Create a migration table
Restore a backed-up Group Policy object using GPMC
Copy a Group Policy object using GPMC
Import a Group Policy object using GPMC
Backup using GPMC
Copy using GPMC
Restore using GPMC
Import using GPMC
Migration tables
The migration table editor
Scripting Group Policy tasks using GPMC

# Restore using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restore**

Restoring a Group Policy object (GPO) re-creates the GPO from the data in the backup. A restore operation can be used in both of the following cases: the GPO was backed up but has since been deleted, or the GPO is live and you want to roll back to a known previous state. The end effect is the same in either case, except as noted below in some cases for GPOs that contain software installation settings. A restore operation replaces the following components of a GPO:

- GPO settings
- The Discretionary Access Control List (DACL) on the GPO
- WMI filter links (but not the filters themselves)

The restore operation does not restore objects that are not part of the GPO. This includes:

- Links to a site, domain, or organizational unit. Links are an attribute of the site, domain, or organizational unit, not the GPO. Any existing links in these containers will continue to be used, for example, when restoring an existing GPO to a previous state. However, if you have deleted a GPO and all links to the GPO, you must add these links back after restoring the GPO.
- WMI filters. A GPO only contains a link to the WMI filter. If the WMI filter does not exist at the time of restore, the link will be dropped, otherwise the link will be restored.
- IPSec Policies. A GPO only contains a link to the IPSec policy. If the IPSec policy does not exist at the time of restore, the link will be dropped, otherwise the link will be restored.

GPMC identifies the GPO by its domain and globally unique identifier (GUID). The purpose of a restore operation is to return the GPO to its original state, so the restore operation retains the original GUID even if it is recreating a deleted GPO. This is a key difference between the restore operation and the import or copy operations. You cannot use restore to transfer GPOs to different domains or forests. That capability is provided by import and copy.

GPMC identifies the GPO by its domain and globally unique identifier (GUID), a restore operation retains the original GUID even if it is recreating a deleted GPO. This is a key difference between the restore operation and the import or copy operations. You cannot use restore to transfer GPOs to different domains or forests. That capability is provided by import and copy.

**Restoring a GPO from backup: permissions required**

The permissions necessary to perform a restore of a GPO vary, depending on whether you are restoring an existing GPO or if you are restoring a GPO that has been deleted since it was backed up. Version numbers for the GPO are handled differently as well. The following table summarizes the situation for existing and deleted GPOs:

| GPO state | Permission needed to restore GPO from backup | GPO version number |
|---|---|---|
| Existing GPO | The user must have **Edit settings, delete, and modify** permissions on the GPO, as well as read access to the file system location where the backup is stored. This does NOT require GPO creation rights. | Incremented by 1, which will trigger client refreshes of settings from the GPO. |
| Deleted GPO | The user must have the right to create GPOs in the domain, as well as read access to the file system location where the backup is stored. | Retained unchanged from the backed-up GPO. |

**Restoring GPOs with software installation settings**

When restoring a *deleted* GPO that contains Software Installation settings, some side effects are possible depending on the circumstances under which the GPO is restored.

When restoring a GPO that contains software installation settings, it is possible that:

- Cross-GPO upgrade relationships that upgrade applications in the GPO being restored, if any, are not preserved after restore. A cross-GPO upgrade is one in which the administrator has specified that an application should upgrade another application, and the two applications are not in the same GPO. Note that upgrade relationships are preserved when applications—in the GPO being restored—upgrade applications in other GPOs.

- If the client computer has not yet detected that the GPO has been deleted (either because the user has not logged on again or rebooted since the deletion of the GPO), and the application was deployed with the option to **Uninstall this application when it falls out of scope of management** then the next time the user logs on:

    - Published applications that the user has previously installed will be removed.

    - Assigned applications will be uninstalled before re-installation.

This issue can be avoided if all of the following conditions are met:

- You perform the restore on a Windows Server 2003 domain controller instead of a Windows 2000 domain controller.

- The user performing the restore has permissions to re-animate tombstone objects in the domain.

- The time elapsed between deletion and restoration of the GPO does not exceed the tombstone lifetime specified in Active Directory.

Tombstone re-animation is a new feature of Windows Server 2003 Active Directory. By default, only Domain Admins and Enterprise Admins have this permission but you can delegate this right to additional users at the domain level using the ACL editor.

As a general rule, if you deploy software using Group Policy, it is recommended that you perform the restoration of GPOs that contain application deployments using a domain controller running Windows Server 2003 and that you grant the tombstone re-animation right to the users who will be performing restoration of those GPOs.

For more details on this issue, see the Group Policy Administration with Group Policy Management Console (GPMC) (http://www.microsoft.com/).

Finally, when restoring a GPO that contains software installation settings, if you are using categories to classify applications, the application in the restored GPO will appear in its original category only if the category exists at the time of restoration. Note that the category definition is not part of the GPO.

**See Also**
**Concepts**
Backup using GPMC
Restore a backed-up Group Policy object using GPMC
Backup, Restore, Import, Copy, and Migration Tables

# Import using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Import**
The Import operation transfers settings into an existing GPO in Active Directory using a backed-up GPO in the file system location as its source. Import operations can be used to transfer settings from one GPO to another GPO within the same domain, to a GPO in another domain in the same forest, or to a GPO in a domain in a different forest. The import operation always places the backed-up settings into an existing GPO. It erases any pre-existing settings in the destination GPO. Import does not require trust between the source domain and destination domain. Therefore it is useful for transferring settings across forests and domains that don't have trust. Importing settings into a GPO does not affect its discretionary access control list (DACL), links on sites domains or organizational units to that GPO, or a link to a WMI filter.

For step-by-step instructions, see Import a Group Policy object using GPMC.

If you have separate test and production forests, with or without trust, it is a good idea to test GPOs in the test forest before importing them into the production forest.

When using import to transfer GPO settings to a GPO in a different domain or different forest, you may want to use a migration table in conjunction with the import operation. A migration table allows you to facilitate the transfer of references to security groups, users, computers, and UNC paths in the source GPO to new values in the destination GPO. For information about migration tables, see Migration tables. For step-by-step instructions on creating a migration table, see Create a migration table.

**See Also**
**Concepts**
Copy a Group Policy object using GPMC
Import a Group Policy object using GPMC
Import using GPMC
Migration tables
The migration table editor
Scripting Group Policy tasks using GPMC

# Copy using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Copy**
A copy operation allows you to transfer settings from an existing Group Policy object (GPO) in Active Directory directly into a new GPO. The new GPO created during the copy operation is given a new globally unique identifier (GUID) and is unlinked. You can use a copy operation to transfer settings to a new GPO in the same domain, another domain in the same forest, or a domain in another forest. Because a copy operation uses an existing GPO in Active Directory as its source, trust is required between the source and destination domains. Copy operations are suited for moving Group Policy between production environments, and for migrating Group Policy that has been tested in a test domain or forest to a production environment, as long as there is trust between the source and destination domains.

For step-by-step instructions, see Copy a Group Policy object using GPMC.

**Copying** is similar to a backing up followed by importing, but there is no intermediate file system step, and a new GPO is created as part of the copy operation. For information about backup, see Backup using GPMC.

The import operation, by contrast with the copy operation, does not require trust. For information about the import operation, see Import using GPMC.

**Specifying the discretionary access control list (DACL) on the new GPO**
You have two options for specifying the DACL to use on the new GPO:

- Use the default permissions that are used when creating new GPOs.
- Preserve the DACL from the source GPO. For this option, you can specify a migration table, used to facilitate the transfer of references to security groups, users, computers, and UNC paths in the source GPO to new values in the destination GPO. If you specify a migration table for the copy operation, and you choose the option to preserve the DACL from the source GPO, the migration table will apply to any security principals in the DACL. For information on migration tables, see Migration tables.

**Copying within a domain compared with copying to another domain**
Copying a GPO to another domain is slightly different from copying it to the same domain:

When copying a GPO within the same domain, you have a simple choice of two options, just described, for choosing the DACL. However, for copy operations to another domain, GPMC presents a wizard to facilitate the operation. The wizard guides you through the following choices:

- Choice of DACL for the new GPO, as described earlier.
- Specification of migration table, if applicable. A migration table allows you to facilitate the transfer of references to security groups, users, computers, and UNC paths in the source GPO to new values in the destination GPO. For more details, see Migration tables.

When copying a GPO within the same domain, any link to a WMI filter is preserved. However, when copying a GPO to a new domain, the link is dropped because WMI filters can only be linked to GPOs within the same domain.

**See Also**
**Concepts**
Copy a Group Policy object using GPMC
Migration tables
Backup, Restore, Import, Copy, and Migration Tables

# Migration tables

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Migration tables**

When you copy or import a Group Policy object (GPO) from one domain to another, you can use a migration table to tell Group Policy Management Console (GPMC) how domain-specific data should be treated.

The key challenge when migrating GPOs from one domain to another is that some information in the GPO is actually specific to the domain where the GPO is defined. When transferring the GPO to a new domain, it may not always be desirable, or even possible, to use the exact same settings. Items that can be specific to a domain include references to security principals such as users, groups, and computers as well as UNC paths. The solution is to modify these references in the GPO that are domain-specific, during the import or copy operation, so that the settings in the destination GPO are written with the appropriate information for the destination domain. GPMC supports this capability using migration tables.

A migration table is a file that maps references to users, groups, computers, and UNC paths in the source GPO to new values in the destination GPO. A migration table consists of one or more mapping entries. Each mapping entry consists of a type, source reference, and destination reference. If you specify a migration table when performing an import or copy, each reference to the source entry will be replaced with the destination entry when writing the settings into the destination GPO.

The migration table will apply to any references in the settings within a GPO, whether you are performing an import or copy operation. In addition, during a copy operation, if you choose the option to preserve the discretionary access control list (DACL) on the GPO, the migration table will also apply to both the DACL on the GPO and the DACLs on any software installation settings in the GPO.

Migration tables store the mapping information as XML, and have their own file name extension, .migtable. You can create migration tables using the Migration Table Editor (MTE). The MTE is a convenient tool for viewing and editing migration tables without having to work in, or be familiar with, XML. The MTE is associated with the .migtable extension so that when you double click a migration table, it opens in the MTE. The MTE is installed with GPMC. For information about MTE, see The migration table editor. You can also create and edit migration tables using any XML editor or using the GPMC scripting interfaces.

For more information about using migration tables for importing GPOs from one forest to another, see the white paper, Migrating GPOs Across Domains with GPMC.

**Settings impacted by migration tables**

There are two types of settings that may not transfer well across domains:

- Users, groups, and computers (security principals) that are referenced in the following places: the settings of the GPO, in the ACL for the GPO itself, and the ACL for any software installation settings in the GPO. Security principals do not transfer well for several reasons. For instance, domain local groups are invalid in external domains but other groups are valid if there is trust in place. Even with a trust between domains, it may not always be appropriate to use the exact same group in the new domain. If there is no trust, then none of the security principals in the source domain will be available in the destination domain.

- UNC paths. When you are migrating GPOs across domains that do not have trust, such as from test to production environments, users in the destination domain may not have access to paths in the source domain.

The following items can contain security principals and can be modified using a migration table:

- Security policy settings of the following types:
    - User rights assignment
    - Restricted groups
    - Services
    - File system

    - Registry

- Advanced folder redirection policies.

- The GPO Discretionary Access Control List (DACL), if you choose to preserve it during a copy operation.

- The DACL on software installation objects. This is only preserved if the option to copy the GPO DACL is specified. Otherwise the default DACL is used.

  **Note**

    - Security Principals referenced in Administrative Templates settings cannot be migrated using a migration table.

The following items can contain UNC paths and can be modified using a migration table:

- Folder redirection policies.

- Software installation policies (for software distribution points).

- Pointers to scripts deployed through Group Policy (such as startup and shutdown scripts) that are stored outside the GPO. Note that the script itself is not copied as part of the GPO copy operation, unless the script is stored inside the source GPO.

**Options for specifying migration tables**
Migration tables are specified when performing import and copy operations. There are three options for using migration tables with import and copy:

- Do not use a migration table. This option copies the GPO exactly as it is. All references to security principals and UNC paths are copied identically.

- Use a migration table. This option maps any references in the GPO that are in the migration table. References that are not contained in the migration table are copied as is.

- Use a migration table exclusively. This option requires that all references to security principals and UNC paths that are referenced in the source GPO be specified in the migration table. If a security principal or UNC path is referenced in the source GPO and is not included in the migration table, the import or copy operation fails.

In addition, cross-domain copy operations will apply the migration table to the DACL on the GPO (and any software installation settings) if you choose the option to **Preserve or migrate the existing permissions.**

When performing a copy or import, the wizard scans the source GPO to determine if there are any references to security principals or UNC paths in the GPO. If there are, you have the opportunity to specify a migration table. During cross-domain copy operations, if you specify the option to **Preserve or migrate the permissions on the GPO**, the wizard will always present the opportunity to specify a migration table, because a DACL by definition contains security principals.

**Contents of Migration tables**
A migration table consists of one or more mapping entries. Each mapping entry has three pieces of information: Type, source and destination. The following table describes these categories in more detail.

| | **Description** |
|---|---|

| | |
|---|---|
| **Source Type** | The type of domain-specific information in the domain for the source GPO. Migration tables support the following types:<br>• User<br>• Computer<br>• Domain Local Group<br>• Domain Global Group<br>• Universal Group<br>• UNC Path<br>• Free Text or security identifier (SID). Note: This category is only for use with security principals that are specified as free text and raw SIDs.<br><br>**Note**<br>• Built-in groups such as "Administrators" and "Account Operators" have the same SID in all domains. If references to built-in groups are stored in the GPO using their resolved format (based on the underlying SID), they cannot be mapped using migration tables. However, if references to built-in groups are stored as free text, you can map them using a migration table, and in this case, you must specify source type="Free Text or SID."<br><br>When viewed in an XML editor, the types are represented as follows:<br>• `<Type>User</Type>`<br>• `<Type>Computer</Type>`<br>• `<Type>LocalGroup</Type>`<br>• `<Type>GlobalGroup</Type>`<br>• `<Type>UniversalGroup</Type>`<br>• `<Type>UNCPath</Type>`<br>• `<Type>Unknown</Type>` |

| | |
|---|---|
| **Source Name** | The exact name of the User, Computer, Group or UNC Path referenced in the source GPO. For example, **TestDomain\PolicyAdmins** or **\\server1\publicshare.** The type of the source name must match the source type specified in the migration table.<br><br>Security principals can be specified using any of the following formats:<br><br>● UPN. For example, "someone@example.com".<br><br>● SAM. For example, "example\someone".<br><br>● DNS. For example, "example.com\someone".<br><br>● Free text. For example, "someone". You must specify the type as Unknown Group in this case.<br><br>● SID. For example, S-1-11-111111111-111111111-1111111111-1112. You must specify the type as Free Text or SID in this case. |

| | |
|---|---|
| **Destination Name** | The destination name specifies how the name of the User, Computer, Group or UNC Path in the source GPO should be treated upon transfer to the destination GPO.<br><br>There are four choices:<br><br>1. **Same as source**. Copy without changes. Equivalent to not putting the source value in the migration table at all.<br>  - MTE value: <Same As Source><br>  - XML tag:<br>  `<DestinationSameAsSource />`<br>2. **None**. Removes the User, Computer or Group from the GPO. This option cannot be used with UNC paths.<br>  - MTE value: <None><br>  - XML tag: <DestinationNone /><br>3. **Map by relative name**. For example, map **SourceDomain\Group1** to **TargetDomain\Group1**. This option cannot be used with UNC paths.<br>  - MTE value: <Map by Relative name><br>  - XML tag: `<DestinationByRelativeName />`<br>4. **Explicitly specify value.** In the destination GPO, replace the source value with the exact literal value you specify.<br>  - MTE value:<br>  `<exact name to use in target>`<br>  - XML tag: `<Destination>`*Exact name to use in target domain*`</Destination>`<br><br>Security principals can be specified using any of the following formats:<br><br>- UPN. For example, "someone@example.com".<br>- SAM. For example, "example\someone".<br>- DNS. For example, "example.com\someone".<br>- Free text. For example, "someone".<br><br>Raw SIDs cannot be used for the destination name. |

A sample migration table can be found at **%programfiles%\GPMC\scripts\SampleMigrationTable.migtable**.

**Creating Migration Tables Manually**

Valid migtable files must contain the following entries, which identifies the namespace used by .migtables. If you are creating migtables by hand, you need to add this. Otherwise MTE does this for you.

```
<?xml version="1.0" encoding="utf-16"? >

<MigrationTable xmlns="http://www.microsoft.com/GroupPolicy/GPOOperations/MigrationTable">
```

Each mapping entry appears in the following format:

```
<Mapping>

<Type> Type </Type>

<Source> Source </Source>

<Destination> Destination </Destination>

</Mapping>
```

Migration Tables can also be created using the GPMC scripting interfaces. For example, see the CreateMigrationTable.wsf sample script that is included with GPMC.

**See Also**
**Concepts**
The migration table editor
Create a migration table
Scripts overview for GPMC

# The migration table editor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The Migration Table Editor**

The Migration Table Editor (MTE) is provided with Group Policy Management Console (GPMC) to facilitate the editing of migration tables. Migration tables are used for copying or importing Group Policy objects (GPOs) from one domain to another, in cases where the GPOs include domain-specific information that must be updated during copy or import. For information about migration tables, see Migration tables.

MTE displays three columns of information:

| Source Name | Source Type | Destination Name |
|---|---|---|
| Name of the user, computer, group, or UNC path, referenced in the source GPO. | The type of the reference named in the Source Name column. For example, Domain Global Group. | The new value, after copy or import to the new GPO; or the method used to calculate the new value. |

Note that you do not specify Destination Type. The Destination Type is determined during the import or copy operation by checking the actual reference identified in the Destination Name.

**Migration Table Editor Features**

MTE offers these features:

- You can edit text fields using Cut, Copy, and Paste options, for example when you right-click an item in the **Source Name** column.
- You can add computers, users, and groups, for source and destination names, by using a **Browse** dialog box.
- You can fill in **Source Type** fields by using a drop-down menu.
- MTE provides automatic syntax checking to make sure the XML file is properly populated and also ensures that only compatible entries are entered in the table. For example, \\server1\share1 is not a valid security group name name, and server1 is not a valid UNC path so MTE prompts the user to correct those entries.
- You can validate the overall migration table using MTE, in addition to basic syntax checking using the **Validate Table** option, which is available on the **Tools** menu. The **Validate Table** option checks to ensure:
    - The existence of destination security principals and UNC paths. This is important to know before you copy or import a GPO, because if there are unresolvable entries in the migration table, then the copy or import operation might fail.
    - Source entries with UNC paths do not have destinations of MapByRelative name or NoDestination.
    - Type of the destination entry in the table matches the real type in the destination source.
- You can auto-populate a migration table by scanning one or more GPOs or backups to extract all references to security principals and UNC paths, and then enter these items into the table as source name entries. This capability is provided by the **Populate from GPO** and **Populate from Backup** options available on the **Tools** menu. To complete the migration table, you only need to adjust the destination values. By default, the destination name for each entry will be set to **Same as source** when you use either of the Populate options.

  **Note**

    - With either auto-populate option, when the source GPO or backup contains references to built-in groups such as **Administrators** or **Backup Operators**, these references will not be entered into the migration table if the references are stored in the source GPO or backup in their resolved format [based on security identifier (SID)].

This is because built-in groups have the same SID in all domains and therefore cannot be mapped to new values using a migration table. However, if references to built-in groups are stored as free text, they will be captured by the auto-populate options. In this case, you can map them as long as you specify the source type as "Free Text or SID."

**Accessing Migration Table Editor**

MTE can be opened in the following ways:

- In GPMC, right-click either the **Domains** node in a given forest or the **Group Policy Objects** container in a given domain, and then click **Open Migration Table Editor**.
- In GPMC, when you use wizards to import or copy GPOs that have references to users, groups, computers, or UNC paths, you can open MTE to edit, examine, or create a migration table.
- When you double-click a migration table (.migtable), it opens in MTE.
- You can open MTE from the command line. The executable, mtedit.exe, is located in **Program Files\GPMC**. Example: mtedit table1.migtable /domain:domain1.wingtiptoys.com /DC:DC1.

**Command-line support for Migration Table Editor**

### Syntax

```
Mtedit.exe [/domain:value][/DC:value] [/forest:value][/DisableTrustChecking][filename]
[/?]
```

### Parameters

`/Domain:value`

Specifies the domains to be used with the "Populate from GPO" option. The domain must be specified in DNS format. When used with the /DC parameter, this parameter can be specified multiple times in the order to include multiple domains.

`/DC:value`

Specifies the domain controller to use for the previously specified domain. If no domain controller is specified, any will be used. You can specify either DNS or NetBIOS format for the name of the domain controller.

`/Forest:value`

If specified, the Migration Table Editor will use all domains in the specified forest to populate the Domain list in the Select GPO dialog box used for the "Populate from GPO" option. This parameter can be specified multiple times in the order to include domains from the more than one forest. The forest name must be specified as the DNS name of the forest root domain in the forest.

`/DisableTrustChecking`

If this parameter is not specified, you can only access domains to which you have two-way trust.

`Filename`

Specifies the name of a migration table you want to open in the Migration Table Editor.

`/?`

Displays the usage.

**See Also**

**Concepts**

Migration tables

Create a migration table

# Delegation and policy-related permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Delegation and policy-related permissions**

With Group Policy Management Console (GPMC), the following tasks can be delegated:

- Create Group Policy objects (GPOs) in a domain.
- Set permissions on a GPO.
- Set policy-related permissions on site, domain or organizational unit.
    - Link GPOs to a given site, domain or organizational unit.
    - Perform Group Policy Modeling analyses on a given domain or organizational unit (but not on a site).
    - Read Group Policy Results data for objects in a given domain or organizational unit (but not on a site).
- Create WMI filters in a domain.
- Set permissions on a WMI filter.

GPMC simplifies delegation by managing the various access control entries (ACEs) required for a task as a single bundle of permission for the task. If you want to see the access control list (ACL) in detail, you can click the **Advanced** button on the **Delegation** tab.

**Delegating creation of GPOs**

Creating GPOs is a user right of the Group Policy Creator Owners (GPCO) group by default but can be delegated to any group or user. There are two methods to grant a group or user this right:

- Add the user or group to membership of the Group Policy Creator Owners (GPCO) group. This was the only method available prior to GPMC.
- Explicitly grant the group or user permission to create GPOs. This method is newly available with GPMC.

You can manage this permission using the **Delegation** tab on the **Group Policy Objects** container for a given domain in GPMC. This tab shows the groups that have permission to create GPOs in the domain, including the GPCO group. From this tab, you can modify the membership of existing groups with this permission, or add new groups.

The ability to grant users permissions to create GPOs without using GPCO was added to facilitate the delegation of GPO creation to users outside the domain. Because the Group Policy Creator Owners group is a domain global group, it cannot contain members from outside the domain. Thus, prior to GPMC, this task could not be delegated to members outside the domain.

It is recommended that for users and groups within the domain, you continue to use the GPCO group to grant them GPO creation rights. If you require that users outside the domain have the ability to create GPOs, then create a new domain local group in the domain ("GPCO - External"), grant that group GPO creation rights in the domain, and then add external domain users to that group.

Adding a user to the membership of GPCO, or granting the user GPO creation permissions directly using the new method available in GPMC, is identical in terms of permissions. Users have the ability to create GPOs in the domain, but do not have permissions on GPOs created by other users. For example, granting a user the ability to create GPOs in the domain does not give the user the ability to edit or delete existing GPOs, or the ability to link the GPO to a site, domain or organizational unit.

For step-by-step instructions, see Delegate creation of Group Policy objects using GPMC.

**Delegating an individual GPO**

There are five permission options on GPOs in GPMC user interface. Each corresponds to a set of individual NT permissions. The correspondence is summarized in the following table.

| Option in GPMC user interface | Corresponding NT permission in ACL Editor |
|---|---|
| Read | Allow **Read** Access on the GPO |
| Edit settings | Allow **Read**, **Write**, **Create Child Objects**, and **Delete Child Objects** |
| Edit, delete, and modify security | Allow **Read**, **Write**, **Create Child Objects**, **Delete Child Objects**, **Delete**, **Modify Permissions**, and **Modify Owner**. This essentially grants full control on the GPO, except that the **Apply Group Policy** permission is not set. |
| Read (from Security Filtering) | This setting cannot be set directly, but appears on the delegation tab if the user has **Read** and **Apply Group Policy** permissions to the GPO. |
| Custom | Any other combination of permissions, including the use of **Deny** will show up as **Custom** in the display. GPMC can only set custom permission sets by clicking the Advanced button and opening the ACL editor. |

Permissions on a GPO are managed from the Delegation tab of that GPO. For step-by-step instructions, see Delegate an individual Group Policy object using GPMC.

**Delegating Group Policy-related permissions on sites, domains, and organizational units**

The ability to link GPOs to a site, domain or organizational unit is a permission that is specific to that site, domain or organizational unit. In GPMC, this permission can be managed using the **Delegation** tab on the site, domain or organizational unit when you click the **Link GPOs** option in the permission drop-down list box. At the individual permission level in Active Directory, this allows **Read** and **Write** access to the gPLink and gPOptions attributes on the site, domain, or organizational unit. By default, only Domain Administrators and Enterprise Administrators have this permission.

Group Policy Modeling allows the user to simulate the resultant set of policy for objects in a domain or organizational unit. In GPMC, this permission can be managed using the **Delegation** tab on the domain or organizational unit when you click the **Perform Group Policy Modeling Analyses** option in the **permission** drop-down list box. This feature is only available to Domain Administrators and Enterprise Administrators by default but can be delegated to other users or groups. At the individual permission level in Active Directory, this delegation is equivalent to granting the user or group the **Generate Resultant Set of Policy (Planning)** permission on the domain or organizational unit. This permission is only available in forests that have the Windows Server 2003 schema or later.

Group Policy Results allows the user to read Resultant Set of Policy logging data for objects in a domain or organizational unit. In GPMC, this permission can be managed using the **Delegation** tab on the domain or organizational unit when you click the **Read Group Policy Results data** option in the **permission** drop-down list box. By default, only users with local administrator rights on the target computer can remotely access Group Policy Results data. However, this right can be delegated to other users or groups. Delegation is performed on a domain or organizational unit. Users with this permission can read Group Policy Results data for any object in that container. This requires the **Generate Resultant Set of Policy (Logging)** permission on the domain or organizational unit. This permission is only available in forests with the Windows Server 2003 schema.

For step-by-step instructions for delegation of linking of GPOs, delegation of Group Policy Modeling analyses, and delegation of Group Policy Results, see Delegate policy-related permissions on a domain, OU, or site using GPMC.

**Delegating creation of WMI filters**

The ability to create WMI filters is a per-domain right. In GPMC, there are two levels of permission for creating WMI filters. **Creator Owner** allows the user to create new WMI Filters in the domain, but does not grant them permissions on WMI filters created by other users. **Full Control** allows the user to create WMI filters, and grants them full control on all WMI Filters in the domain, including new filters that are created after they are granted this right.

When a new WMI filter is created it is stored in the WMIPolicy container in the domain's System container in Active Directory Users and Computers. It is the permissions on the WMIPolicy container that govern the rights a user has to create, edit, and delete WMI Filters.

In GPMC, you can manage these permissions from the **Delegation** tab of the **WMI Filters** container in a given domain.

For step-by-step instructions, see Delegate creation of WMI filters using GPMC.

**Delegating an individual WMI Filter**

GPMC provides the ability to delegate rights on an individual WMI Filter. There are two levels of delegation of an individual WMI filters. **Edit** allows the user or group to edit the WMI Filter. **Full Control** allows the user or group to edit, delete, and modify security on the WMI Filter.

All users have **Read** access to all WMI filters through the Authenticated Users group, because it is required to allow policy processing. GPMC does not allow you to remove read permission.

These permissions are managed using the **Delegation** tab of a WMI filter. For step-by-step instructions, see Delegate an individual WMI filter using GPMC.

**See Also**

**Concepts**

Delegate creation of Group Policy objects using GPMC
Delegate policy-related permissions on a domain, OU, or site using GPMC
Delegate an individual Group Policy object using GPMC

# Group Policy objects overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy objects**

Policy settings are stored in Group Policy objects (GPOs). Settings for each GPO are edited using the Group Policy Object Editor. After installation of the Group Policy Management Console (GPMC), Group Policy Object Editor is usually opened from GPMC. For information about Group Policy Object Editor, see Group Policy object editor overview for GPMC. For information about GPMC, see Group Policy Management Console Overview.

There are two kinds of GPOs:

- **Active Directory-based GPOs**. These are stored in a domain and replicate to all the domain controllers for the domain. They are available only in an Active Directory environment. They apply to users and computers in a site, domain, or organizational unit to which the Group Policy object is linked. This is the primary mechanism through which Group Policy is used in an Active Directory environment.

- **Local GPOs**. There is just one local GPO stored on each computer. Local GPOs are the least influential GPOs in an Active Directory environment, and local GPOs have only a subset of the settings found in Active Directory-based GPOs. For information about local GPOs, see Local Group Policy objects overview for GPMC.

**Active Directory-based GPOs**

Active Directory-based GPOs can be linked to a domain, site, or organizational unit to apply their settings.

A GPO can be linked to more than one site, domain or organizational unit.

A site, domain or organizational unit can have more than one GPO linked to it. In that case there are rules that determine which settings prevail in case of conflict. See Group Policy processing and precedence for information about these rules. In the case where multiple GPOs are linked to a particular site, domain, or organizational unit, you can prioritize the order and therefore the precedence with which these GPOs are applied. By default, the last applied configured setting is used. Settings are applied in this order: Local, Site, Domain, and organizational unit. For example, suppose **Add Logoff to the Start Menu** were configured in GPOs as follows:

| The local GPO | Site-linked GPO | Domain-linked GPO | Organizational unit-linked GPOs |
|---|---|---|---|
| Disabled | Not Configured | Disabled | Enabled for the first applied (link order 2) GPO, and Not Configured for the second applied GPO (link order 1) |

In that case Enabled wins, and Logoff will appear on the user's start menu.

GPMC introduces support for copy, import, backup, and restore of GPOs. See Back Up, Restore, Copy, and Import and Backup, Restore, Import, Copy, and Migration Tables for information about GPMC support for these tasks.

For information on modifying the scope of a GPO, see Controlling the Scope of Group Policy Objects using GPMC.

**User settings and computer settings**

GPO settings are divided between **User Configuration**, which holds settings that are applied to users when they log on, and **Computer Configuration**, which holds settings that are applied to computers when they start up (boot). Most settings are found in only one section, but a few, like **Run logon scripts synchronously**, are found in both. If they are found in both, and there is a disagreement, the computer setting is used.

User Configuration and Computer Configuration are further subdivided into a customizable set of MMC extensions to Group Policy. To learn about the default extensions, see Group Policy Object Editor Extensions.

**Changing the status of a GPO**

The status of a GPO is **Enabled** by default. It can be changed to **User settings disabled**, which disables the User Configuration of the GPO, or **Computer settings disabled**, which disables the Computer Configuration of the GPO, or **All settings disabled**, which disables the entire GPO. When a client computer processes a GPO, disabled portions of the GPO are not evaluated.

The status of a GPO is **Enabled** by default. It can be changed to **User settings disabled**, which affects all settings under **User Configuration**, or **Computer settings disabled**, which affects all settings under **Computer Configuration**, or **All settings disabled**, which disabled the entire GPO.

When you change the status of a GPO, all sites, domains and organizational units that get policy from the GPO are affected. Thus, disabling a GPO is more far-reaching than disabling one of it links.

**Notes**

- **Enforce** (previously known as "no override") on a GPO link takes precedence over **Block Inheritance** on a domain or organizational unit.
- If you turn on **Enforced** and turn off **Link Enabled** for a GPO link, then the GPO does not apply.
- **Block Inheritance** does not deflect Group Policy settings from GPOs that are linked directly to the domain or organizational unit that has **Block Inheritance** enabled.

**Default permissions on GPOs**

The default permissions on Group Policy objects are shown in the following table.

| Security group | Default permissions as shown in GPMC |
|---|---|
| Authenticated Users | Security Filtering on GPO Scope tab; "Read (from Security Filtering)" on Delegation tab |
| SYSTEM | Edit settings, delete, modify security |
| Domain Admins | Edit settings, delete, modify security |
| Enterprise Admins | Edit settings, delete, modify security |
| ENTERPRISE DOMAIN CONTROLLERS | Read |

**Note**

- The Enterprise domain controllers group is not granted any permissions on GPOs in a pure Windows 2000 forest that does not have the Windows Server 2003 schema.

The special Group Policy objects **Default Domain Policy** and **Default Domain Controllers Policy** cannot be deleted. The purpose of this restriction is to prevent the accidental deletion of these Group Policy objects, which contain important and required settings for the domain.

**See Also**

**Concepts**

Create or delete a Group Policy object
Copy a Group Policy object using GPMC
Control Group Policy Object Scope

# Local Group Policy objects overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local Group Policy objects**
The local Group Policy object (local GPO) is stored on each individual computer, in the hidden **%systemroot%\System32\GroupPolicy** directory. Each computer running a version of Windows Server 2003, Windows 2000, or Windows XP Professional has exactly one local GPO, regardless of whether the computers are part of an Active Directory environment.

**Limitations of the local Group Policy object**
Local GPOs do not support certain extensions, such as Folder Redirection or Group Policy Software Installation. Local GPOs do support many security settings, but the Security Settings extension of the Group Policy Object Editor does not support remote management of local GPOs. Thus, for example, the command line **gpedit.msc /gpcomputer:"Computer1"** will allow you to edit the local GPO on Computer1, but **Security Settings** will not appear.

Local GPOs are always processed, but are the least influential GPOs in an Active Directory environment, because Active Directory-based GPOs have precedence.

For more information about the Group Policy extensions that differ in the local GPO as compared with Active Directory-based GPOs, see Folder redirection overview for GPMC, Software installation overview for GPMC, and Security settings overview for GPMC.

**See Also**
**Concepts**
Group Policy objects overview for GPMC
Group Policy Management Console Overview

# Group Policy object editor overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy object editor**
Group Policy Object Editor is an MMC snap-in used to edit the policy settings in Group Policy objects (GPOs).

**Opening Group Policy Object Editor**
In Group Policy Management Console (GPMC), open Group Policy Object Editor by right-clicking a node for a GPO or a GPO link and then clicking **Edit**.

**Note**

- You must have **edit** rights on the GPO in order to open it in Group Policy Object Editor.

**Configuring Group Policy settings**
Group Policy Object Editor consists of two main sections: **User Configuration**, which holds settings that are applied to users (at logon and periodic background refresh) and **Computer Configuration**, which holds settings that are applied to computers (at startup and periodic background refresh). These sections are further divided into the different types of policies that can be set, such as Administrative Templates, Security, or Folder Redirection. For more information about these categories of policy settings, see Group Policy Object Editor Extensions.

You configure Group Policy settings by navigating to the appropriate location in each section. For example, if you want to set an Administrative Templates policy setting in a GPO to prevent users from seeing the Run command, you would need to enable the policy setting **Remove Run Menu from Start Menu**. This is located under **User Configuration\Administrative Templates\Start Menu and Task Bar**. You edit most policy settings by double-clicking the title of the policy setting, which opens a dialog box that provides specific options. In Administrative Templates policy settings, for example, you can choose to enable or disable the policy setting or leave it as not configured. In other areas, such as Security Settings, you can select a check box to define a policy setting and then set available parameters.

The Group Policy Object Editor provides different ways of learning about the function or definition of specific policy settings. In most cases, when you can double click the title of a policy setting, the dialog box contains any relevant defining information about the policy setting. For Administrative Templates policy settings, the Group Policy Object Editor provides explain text directly in the Web view of the console. You also can find this explain text by double-clicking the policy setting and then clicking the **Explain** text tab. In either case, this text shows operating system requirements, defines the policy setting, and includes any specific details about the effect of enabling or disabling the policy setting.

**Managing Multiple Platforms in Group Policy Object Editor**
Because new Administrative Template policy settings have been added that only work on specific versions of the operating system such as Windows XP Professional or Windows Server 2003, the Group Policy Object Editor enables you to view only the Administrative Template policy settings that might be applied in your users' work environment. For example, you may want to edit only policy settings that could be applied on client computers running Windows 2000 Service Pack 3. You can specify these options in the **Filtering** dialog box, available by clicking a node in the Administrative Templates section, clicking the **View** menu and then clicking **Filtering**.

For more information about editing Group Policy objects see:

- Group Policy objects overview for GPMC
- Edit a Group Policy object from GPMC
- Group Policy Object Editor Extensions

To learn how Group Policy tasks differ after you install GPMC, see New ways to do familiar tasks using GPMC.

**Opening Group Policy Object Editor from the command line**
It is sometimes useful to open Group Policy Object Editor from the command line. By specifying the appropriate command line switches, you can automatically connect to a specific machine or GPO.

The following table illustrates the available command-line switches used to open Group Policy Object Editor under different circumstances.

| Command-line parameters | Effect |
|---|---|
| **gpedit.msc** | Edit the local GPO on the local computer. |
| **gpedit.msc /gpcomputer:"Computer1"**<br><br>**Notes**<br><br>• The quotation marks are mandatory.<br><br>• There is no space after **/gpcomputer:** | You can edit the local GPO on Computer1. |
| **gpedit.msc /gpcomputer:"ThisComputer.TailspinToys.com"**<br><br>**Notes**<br><br>• The quotation marks are mandatory.<br><br>• There is no space after **/gpcomputer:** | You can edit the local GPO on Computer1, with the computer specified in the DNS style rather than the NetBIOS style. |
| **gpedit.msc /gpobject:**"LDAP:"LDAP_PATH_TO_GPO"<br><br>**Notes**<br><br>• LDAP path to the GPO is based on the GUID of the GPO, for example: "LDAP://CN={31B2F340-016D-11D2-945F-00C04FB984F9},CN=Policies,CN=System,DC=WingTipToys,DC=com".<br><br>• You can optionally specify the domain controller to use in the LDAP path as the first item in the LDAP path. You can use either a NETBIOS or DNS format for the name of the domain controller. For example, this path would specify that GPEdit could use domain controller named DC13.TailspinToys.com: "LDAP://DC13.TailspinToys.com/CN={31B2F340-016D-11D2-945F-00C04FB984F9},CN=Policies,CN=System,DC=WingTipToys,DC=com". | You can edit an Active Directory–based GPO by passing in the LDAP path to the GPO. |

**See Also**
**Concepts**
Group Policy objects overview for GPMC
Control Group Policy Object Scope
Group Policy Object Editor Extensions
Edit a Group Policy object from GPMC

# Group Policy Object Editor Extensions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Object Editor Extensions**
Group Policy Object Editor is an MMC snap-in, and like all snap-ins, its functionality can be customized or extended by means of MMC snap-in extensions. When you use Group Policy Object Editor, various extensions are included by default, and are visible in the tree view with icons of their own.

This section contains information about the following Group Policy Object Editor extensions:

- Administrative templates overview for GPMC
- Security settings overview for GPMC
- Software installation overview for GPMC
- Folder redirection overview for GPMC
- Scripts overview for GPMC
- Internet Explorer Maintenance overview for GPMC
- Remote Installation Services overview for GPMC

# Administrative templates overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administrative Templates**

Administrative templates, (or .adm files), enable administrators to control registry settings using Group Policy. These settings appear under the **Administrative Templates** folder for both user configuration and computer configuration in the console tree of the Group Policy Object Editor, and in HTML reports produced by GPMC.

It is important to understand that .adm files are not the actual settings that are deployed to client operating systems. The .adm file is simply a template file (implemented as text file with an .adm extension) that provides the friendly name for the setting and an explanation. This template file is used to populate the user interface. The settings that are deployed to clients are contained in the registry.pol file inside the GPO. On Windows XP and Windows Server 2003, each registry setting contains a "Supported on" tag that indicates which operating system versions support that policy setting. If a setting is specified and deployed to a client operating system that does not support that setting, the settings are ignored. These .adm files are stored in two locations by default: inside GPOs, and in the %windir%\inf folder on the local computer.

Windows includes a predefined set of Administrative template files that define the registry settings that can be configured in a Group Policy object (GPO). The .adm files can be added or removed from the Group Policy Object Editor by right-clicking **Administrative Templates** and clicking **Add/Remove Templates**. Adding or removing .adm files does not affect which policies are processed by the Group Policy engine. It only affects whether a specific Administrative Template policy setting is displayed in the Group Policy Object Editor. For example, if you removed all the .adm files from the GPO via the **Add/Remove Templates** dialog box, no Administrative Template policy settings would be displayed under the Administrative Templates node. This will not affect the policies already stored in the Registry.pol file.

**Contents of .adm files**

An .adm file consists of a hierarchy of categories and subcategories that together define how the policy settings are to appear in the Group Policy Object Editor user interface (UI), and also contains information about which registry locations control the settings.

The following information is included in .adm files:

- Registry locations that correspond to each setting in the Administrative Templates section of the Group Policy Object Editor.
- Options or restrictions in values that are associated with each setting. These are only restrictions for the user interface. There is no checking of value ranges during the actual policy processing.
- Check box, editbox, and other methods of parameter input.
- For many settings, a default value to display.
- Explanations of what each setting does, and about settings that affect and are affected by it, is included in Help text embedded in the .adm file. This information is displayed on an **Explain** tab in Group Policy Object Editor.
- The versions of Windows that support each setting are indicated by use of the **Suppported** keyword.
- Registry location for the setting, hive name, key name, and value name. Type of registry key, whether DWORD, REG_SZ, or other type. Binary data is not supported.

**Standard .adm files included with Windows**

The predefined set of Administrative template files that define the registry settings for controlling various aspects of the operating system are grouped into five .adm files listed below.

A description of each policy setting in Administrative Templates appears in searchable online Help.

| Administrative Template | Description |
|---|---|
| System.adm | System settings |
| Inetres.adm | Internet Explorer settings |
| Wmplayer.adm | Windows Media Player settings. This tool is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. |
| Conf.adm | NetMeeting settings. This tool is not available on Windows XP Professional, 64-Bit Edition and the 64-bit versions of the Windows Server family. |
| Wuau.adm | Windows Update settings. |

**True Policies or Preferences**
The standard Administrative Templates shipped with Windows only contain true policies, as distinct from preferences. The registry information for true policies is located under \Software\Policies or under \Software\Microsoft\Windows\CurrentVersion\Policies in the HKEY_CURRENT_USER or HKEY_LOCAL_MACHINE hives.

Policy settings that are stored in these specific locations of the registry are known as *true policies*. Storing settings here has the following advantages:

- These trees are secure and cannot be modified by a non-administrator.
- When Group Policy changes, for any reason, these trees are cleaned, and the new policies are then rewritten.

This prevents the behavior that was often present in Windows NT 4.0, whereby System Policies resulted in persistent settings in the user and computer registry. The policy remained in effect until the value was reversed, either by a counteracting policy or by editing the registry. These settings are stored outside the approved registry locations above and are known as *preferences*.

All the policy settings in the standard Administrative Templates shipped with Windows use registry settings in the Policies trees of the registry. This means that they will *not*cause persistent settings in the registry when the GPO that applies to them is no longer in effect.

The Group Policy Object Editor does not show preferences. If you add an .adm file that contains preferences rather than policy settings, you will have to clear the **Only show policy settings that can be fully managed** check box on the **Filtering** dialog box under the **View** menu of Group Policy Object Editor, otherwise the settings will not be visible in the console.

If you plan to create entries for Administrative Templates, populate the namespace by using the following naming convention, which is also used in the registry: *CompanyName**product**version* (or *CompanyName**product&version*). For example, the operating system settings for Windows are in \Microsoft\Windows.

For information about creating custom policy settings in .adm files, see "Implementing Registry-based Policy" at the Microsoft Web site. (http://msdn.microsoft.com/)

**Registry.pol files**
Group Policy Object Editor, when setting Administrative Templates policy settings, saves information in Registry.pol files. These Unicode files contain the customized registry settings that you specify (by using Group Policy) that are to be applied to the computer or user portion of the registry. The Registry.pol files contain the actual Group Policy settings used by the Group Policy engine during processing. They contain instructions to add or delete registry keys, corresponding to the settings specified by the administrator in the Group Policy Object Editor.

- One of the Registry.pol files contains registry settings that are specific to the **HKEY_LOCAL_MACHINE** key. It is stored in the GPT\Machine folder.
- The other Registry.pol file contains registry settings that are specific to the **HKEY_CURRENT_USER** key. It is stored in the GPT\User subdirectory.

**See Also**
**Concepts**

Group Policy Object Editor Extensions
Deployment considerations for Group Policy

# Security settings overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security settings**
The Group Policy object editor allows you to edit security settings for users or computers.

Where?

- *Group Policy object name*/Computer (or User) Configuration/Windows Settings/Security Settings

With **Security Settings**, you can require minimum password complexity, control the ability to log on to computers remotely, enable software restriction policies, set IP security, and much more.

**Note**

- Security settings under **Account Policies** for domain user accounts work when the Group Policy object is linked to a domain but are ignored when the Group Policy object is linked to an organizational unit. However, local accounts created on computers affected by the policy will have to follow restrictions defined in the GPO, even when the GPO is linked to an organizational unit.

Where?

- *Group Policy object name*/Computer Configuration/Windows Settings/Security Settings/Account Policies

For information about security settings, see Security Settings extension to Group Policy and Security Setting Descriptions.

For information about specific extensions of Security Settings, see these topics:

- Wireless Networking
- Software Restriction Policies
- Internet Protocol Security (IPSec)
- Encrypting File System
- Public Key Policies

# Software installation overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Software installation**
You can use the Software Installation extension of Group Policy to centrally manage software distribution in your organization. You can assign and publish software for groups of users and computers using this extension.

- Assigning Applications
- Publishing Applications

**Assigning Applications**
When you assign applications to users or computers, the applications are automatically installed on their computers at logon (for user-assigned applications) or startup (for computer-assigned applications.)

When assigning applications to users, the default behavior is that the application will be advertised to the computer the next time the user logs on. This means that the application shortcut appears on the **Start** menu, and the registry is updated with information about the application, including the location of the application package and the location of the source files for the installation. With this advertisement information on the user's computer, the application is installed the first time the user tries to use the application. In addition to this default behavior, Windows XP Professional and Windows Server 2003 clients support an option to fully install the package at logon, as an alternative to installation upon first use. Note that if this option is set, it is ignored by computers running Windows 2000, which will always advertise user-assigned applications.

When assigning applications to computers, the application is installed the next time the computer boots up. Applications assigned to computers are not advertised, but are installed with the default set of features configured for the package. Assigning applications through Group Policy requires that the application setup is authored as a Windows Installer (.msi) package.

**Publishing Applications**
You can also publish applications to users, making the application available for users to install. To install a published application, users can use **Add or Remove Programs** in Control Panel, which includes a list of all published applications that are available for them to install. Alternatively, if the administrator has selected the **Auto-install this application by file extension activation** feature, users can open a document file associated with a published application. For example, double clicking an .xls file will trigger the installation of Microsoft Excel, if it is not already installed. Publishing applications only applies to user policy; you cannot publish applications to computers.

To take advantage of all of the features of Group Policy Software Installation, it is best to use applications that include a Windows Installer (.msi) package. For example, published MSI packages support installation for users who do not have administrative credentials.

However, you can also publish legacy setup programs using a .zap file. These applications will be displayed in **Add or Remove Programs** like any other published application, but typically can only be installed by users with administrative credentials.

A .zap file is a simple text file that describes the path to the setup program, as well as any arguments to be passed on the command line.

A simple example illustrating the syntax of a .zap file is shown below:

```
[Application]
FriendlyName = Microsoft Works 4.5a
SetupCommand = ""\\DeploymentServer\Apps\Works 4.5a\Standard\Setup.exe""
```

**Note**

- When using quotes in zap files, the following rules apply:
    - The path and name of the setup executable must always be quoted.
    - If there are no command-line arguments, they must be quoted twice.

**Additional Information**

Group Policy Software Installation supports a variety of advanced features, such as patching applications and applying transforms at deployment time, as well as filtering based on language, processor architecture, and so on.

For more information about Software Installation, see Software Installation and Maintenance(http://www.microsoft.com/) and Step-by-Step Guide to Software Installation and Maintenance(http://www.microsoft.com/) at the Microsoft Web site.

**See Also**

**Concepts**

Group Policy Object Editor Extensions

Group Policy Management Console Overview

# Folder redirection overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Folder redirection**

User settings and user files are normally stored in the local user profile, under the **Documents and Settings** folder. The files in local user profiles are only accessible from the current computer, which makes it difficult for users who use more then one computer to work with their data and synchronize settings between multiple computers. Two different technologies exist to address this problem: Roaming Profiles and Folder Redirection. Both these technologies have their advantages and they can be used separately or together to create a seamless user experience from one computer to another. They also provide additional options for administrators managing user data.

Folder Redirection allows administrators to redirect the path of a folder to a new location. The location can be a folder on the local computer or a directory on a network file share. Users have the ability to work with documents on a server as if the documents were based on a local drive. The documents in the folder are available to the user from any computer on the network. Folder Redirection is located under **Windows Settings** in the console tree of the Group Policy Object Editor.

**Where?**

- *Group Policy object name*/User Configuration/Windows Settings/Folder redirection

There are five folders that can be redirected:

- **My Documents**
- **My Pictures** (Redirection of My Pictures is provided for backward compatibility reasons, and not as a primary feature. You can only redirect the My Pictures folder using the Windows XP or Windows Server 2003 operating system if you are editing GPO that was originally created on Windows 2000 computer and that redirected My Pictures.)
- **Application Data**
- **Desktop**
- **Start Menu** (Only for Terminal Server users.)

**Advantages of Folder Redirection**

- When roaming user profiles are used, only the network path to the folder is part of the roaming user profile, not the folder. Therefore, its contents do not have to be copied back and forth between the client computer and the server each time the user logs on or off, and the process of logging on or off can be much faster.
- Even if a user logs on to various computers on the network, their documents are always available.
- Offline File technology gives users access to the folder even when they are not connected to the network. This is particularly useful for people who use portable computers.
- Data that is stored in a shared network folder can be backed up as part of routine system administration. This is safer because it requires no action on the part of the user.
- As an administrator, you can use Group Policy to set disk quotas, limiting the amount of space that is taken up by users' special folders.
- Data that is specific to a user can be redirected to a different hard disk on the user's local computer from the hard disk that holds the operating system files. This makes the user's data safer in case the operating system has to be reinstalled.
- The Volume Shadowcopy Service feature of Windows Server 2003 can be used by users to recover older copies of the

files in their redirected folders without intervention of technical support personnel.

- Folder Redirection is a secure way to keep user data on a network, provided that the option **Grant the user exclusive rights** is turned on.

#### Selecting a Folder Redirection target

The **Target** tab of the folder's Properties dialog box allows the administrator to select the location of the redirected folder on a network or in local user profile. You can choose between **Basic** and **Advanced** settings. Basic redirection will be applied to all users affected by the GPO. The **Advanced** setting allows you to provide different redirection locations for users based on the users' security group membership. Using either advanced or basic redirection you can set different targets for folder redirection:

- **Create a folder for each user under the root path**. This option will create a folder in the form **\\server\share*User Account Name**Folder Name***. Each user will get a unique path to their redirected folder. This option is not available for the Start Menu folder.

- **Redirect to the following location**. This option will use an explicit path to the redirection location. This can cause multiple users to share the same path to the redirected folder.

- **Redirect to the local user profile location**. This option will move the location of the folder to the local user profile under **Documents and Settings**.

- The My Documents folder has one more option: **Redirect to the user's home directory**. To use this option, the administrator needs to use Active Directory Users and Computers to set the **home directory** attribute on the user objects for each user affected by the policy.

  **Note**

  - The **Redirect to the user's home directory** option causes the value of the **Grant the user exclusive rights** option on the **settings** tab to be ignored.

#### Granting exclusive rights to special folders

The **Settings** tab in the Properties dialog box contains a check box labeled **Grant the user exclusive rights to My Documents**. This option controls permissions on the newly-created redirected folder. If this option is selected, and the target directory doesn't yet exist, Folder Redirection will create the directory and set permissions on this directory to allow only the user and the local system full control over the folder. If this option is selected, and the target directory already exists, then Folder Redirection will verify ownership of the directory. If the directory is owned by someone else, Folder Redirection will fail. Whatever permissions are in effect by default remain in effect. If you clear this check box, no changes are made to the permissions on the folder. Whatever permissions are in effect by default remain in effect. The Start Menu folder can only be redirected with this option turned off.

It is recommended to secure the contents of all folders except the Start Menu by selecting the **Grant the user exclusive rights** option on the **Settings** tab. The Start Menu folder is special in this respect, because redirection of the Start Menu is recommended to be used only for Terminal Server users, with all users sharing the same Start Menu folder. Therefore, the Start Menu folder doesn't have an option to select **Grant the user exclusive rights**.

#### Policy removal considerations with regard to Folder Redirection

The following table summarizes what happens to redirected folders and their contents when the Group Policy object no longer applies.

| Move the contents of the special folder to the new location setting | Policy Removal option | Results when policy is removed |
|---|---|---|

| | | |
|---|---|---|
| **Enabled** | **Redirect the folder back to the user profile location when policy is removed** | • The special folder returns to its user profile location.<br>• The contents are copied, not moved, back to the user profile location.<br>• The contents are not deleted from the redirected location.<br>• The user continues to have access to the contents, but only on the local computer. |
| **Disabled** | **Redirect the folder back to the user profile location when policy is removed** | • The special folder returns to its user profile location.<br>• The contents are not copied or moved to the user profile location.<br>**Caution**<br>• If the contents of a folder are not copied to the user profile location, the user cannot see them. |
| Either **Enabled** or **Disabled** | **Leave the folder in the new location when policy is removed** | • The special folder remains at its redirected location.<br>• The contents remain at the redirected location.<br>• The user continues to have access to the contents at the redirected folder. |

**Important**

- Changing the redirection option to **Not Configured** does not redirect the folder to the local profile, this option means that Folder Redirection no longer controls the location of the folder-if a folder was previously redirected it will continue to be redirected to the previous location. If an administrator wanted to return the folder to the local user profile they should use the **Redirect to the local user profile setting.**

**Folder Redirection and Offline Files**

The Offline Files technology applies to any mounted or mapped drive that contains documents or data that a user might want to use offline. Offline Files does not depend on Folder Redirection. It is set up and configured on shared network servers separately from the Folder Redirection snap-in. Offline Files enables the user to do useful work even when not connected to the network, for example, on a portable computer or in the event of router failure.

If you use redirected folders of any type, it is recommended that you set up Offline Files as described in the following table.

| Special Folder | Offline File configuration |
|---|---|
| My Documents | Autocaching for documents (or manual caching for documents, if you want users to have to manually make files and folders available for offline use) |
| My Pictures | Autocaching for documents (or manual caching for documents, if you want users to have to manually make files and folders available for offline use) |

| | |
|---|---|
| Application Data | Autocaching for programs |
| Desktop | Autocaching for programs if the desktop is Read Only |
| Start Menu | Autocaching for programs |

**Folder Redirection permissions**

It is recommended to let Folder Redirection create folders for you. If you let Folder Redirection create folders for you, correct permissions are set automatically. Usually, knowledge of these permissions is not necessary. However, there are three reasons the permissions might be of interest:

- An administrator must always secure the network file share used for keeping redirected folder data by setting correct share level (SMB) permissions and NTFS file system permissions.

- Sometimes administrators create the redirected folders before Folder Redirection creates them. The directories created by administrators cannot be used together with the **Grant the user exclusive rights** option because Folder Redirection verifies ownership of the directories before allowing the folder to be redirected. The following table shows what permissions have to be set for Folder Redirection to work with the **Grant the user exclusive rights** option turned off.

- Redirection of My Documents to the home directory provides more relaxed security than standard folder redirection. The following table shows what security is in effect in the standard case.

**NTFS permissions required for the root folder**

| **User account** | **Folder Redirection defaults** | **Minimum permissions needed** |
|---|---|---|
| Creator Owner | Full Control, subfolders and files only | Full Control, subfolders and files only |
| Administrators | No permissions | No permissions |
| Everyone | No permissions | No permissions |
| Local System | Full Control, this folder, subfolders, and files | Full Control, this folder, subfolders, and files |
| Security group of users who need to put data on the shared network server | N/A | List Folder/Read Data, Create Folders/Append Data - **This folder only** |

**Share-level (SMB) permissions required for the root folder**

| **User Account** | **Folder Redirection defaults** | **Minimum permissions needed** |
|---|---|---|
| Everyone | Full Control | No permissions (Use security group) |
| Security group of users who need to put data on the shared network server | N/A | Full Control |

**NTFS permissions required for each user's redirected folder**

| **User account** | **Folder Redirection defaults** | **Minimum permissions needed** |
|---|---|---|
| *UserName* | Full Control, owner of folder | Full Control, owner of folder |
| Local System | Full Control | Full Control |
| Administrators | No permissions | No permissions |

| | | |
|---|---|---|
| Everyone | No permissions | No permissions |

**See Also**
**Concepts**
Group Policy Object Editor Extensions
Group Policy Management Console Overview

# Scripts overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scripts**
Group Policy Object Editor includes two extensions for script deployment:

- **Scripts (Startup/Shutdown)—**You can use this extension, which is located under **Computer Configuration\Windows Settings** in the console tree of Group Policy Object Editor, to specify scripts that are to run when the computer starts up or shuts down. These scripts run as Local System.

- **Scripts (Logon/Logoff)—**You can use this extension, which is located under **User Configuration\Windows Settings** in the console tree of Group Policy Object Editor, to specify scripts that are to run when the user logs on or logs off the computer. These scripts run under the account of the user who is logging on or off.

There are five script types:

- Group Policy logon scripts.

- Group Policy logoff scripts.

- Group Policy startup scripts.

- Group Policy shutdown scripts.

- Legacy logon scripts of the style that predates Group Policy. These are specified on the User object.

**Controlling script behavior**
The main administrative template file, system.adm, is present by default, and includes several settings that control how Group Policy processes logon, logoff, startup and shutdown scripts. To see these settings, in the Group Policy Object Editor console tree, click **Scripts**.

**Where?**

- *Group Policy object name*/**Computer** (or **User**) **Configuration**/**Administrative templates**/**System**/**Scripts**

**Notes**

- By default, scripts run invisibly to the user. However, there are policy settings that make logon or logoff scripts run visibly.

- Scripts that run synchronously in order from the top to the bottom of the list shown in the Properties dialog box for scripts of the given type. Logon scripts run asynchronously by default, and there is a policy setting to make them run synchronously. Startup scripts run synchronously by default, and there is a policy setting to make them run asynchronously.

- Whenever possible, it is recommended to use Group Policy scripts rather than legacy-style scripts attached to the user object.

**See Also**
**Concepts**
Group Policy Object Editor Extensions
Scripting Group Policy tasks using GPMC

# Internet Explorer Maintenance overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Explorer Maintenance**
Group Policy supports the administration and customization of Microsoft Internet Explorer on computers running Windows 2000, Windows XP, and Windows Server 2003.

Settings for these features can be customized using Group Policy:

**Browser User Interface**
Settings for the user interface and the appearance of the browser are available for the following areas:

- Browser Title
- Custom Logo and Animated Bitmaps
- Toolbar Customization (toolbar background and toolbar buttons)

**Connection Settings**
Settings for the presetting and control of the user's connection settings, such as dial-up and local area network (LAN) connections, are available for the following areas:

- Connection Settings
- Automatic Browser Configuration
- Proxy Settings
- User Agent String

**Custom URLs**
Settings for specifying the URLs that appear on users' browsers, such as the Home page, Search page, and customized Favorites and Links, are available for the following options:

- Favorites and Links
- Important URLs

**Security**
Settings for specifying presetting and controlling user security settings, such as security zones, privacy settings, and content ratings, are available for the following areas:

- Security Zones and Content Ratings
- Authenticode Settings

**Program Associations**
Settings for specifying which Internet programs are used by default for common Internet tasks, such as reading e-mail or viewing newsgroups, are available in the following area:

- Programs

For more information on Internet Explorer Maintenance, see the Internet Explorer Maintenance Help, which is accessible in Windows Server 2003, either from Group Policy Object Editor or, through this link: Internet Explorer Maintenance overview for

GPMC.

**Notes**

- Windows Server 2003 includes a new default security configuration for Internet Explorer, called Internet Explorer Enhanced Security Configuration (ESC). ESC impacts the Security Zones and Privacy settings within the Internet Explorer Maintenance settings of a GPO. The Security Zones and Privacy settings can either be ESC enabled or not. For more details, see Deployment considerations for Group Policy

- If you are running Group Policy Management Console on a computer running Windows XP Professional, Internet Explorer Maintenance Help might not be available. To access the Help, install Windows Help from the Windows Server 2003 installation CD onto a computer running Windows XP Professional.

To customize Internet Explorer before deployment or to administer Internet Explorer on other operating systems, you can use the Profile Manager. The Profile Manager is part of the Internet Explorer Administration Kit (IEAK), which is available in the Windows 2000 and Windows XP Resource Kit and on the Microsoft Windows Web site.

**See Also**
**Concepts**
Group Policy Object Editor Extensions
Security settings overview for GPMC

# Remote Installation Services overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services**
Remote Installation Services (RIS) is an optional component of the Windows Server 2003 operating system. It works with other Windows technologies to enable administrators or users to remotely install a copy of the Windows Server 2003, Windows XP Professional, or Windows 2000 operating systems. Administrators use the Remote Installation Services extension of Group Policy only to specify which options are presented to users by the Client Installation Wizard.

Client computers that install an operating system through RIS must be equipped with network cards that either support Pre-boot Execution Environment (PXE), or are supported by the RIS remote boot floppy disk. The client computer starts (boots), connects to the network, and accesses the RIS server to install the operating system. The RIS server queries Active Directory to determine remote installation options defined for the user via policy setting. Based on that result, RIS determines which screens to send to the pre-boot RIS client code for display to the user.

The RIS settings that control what the user sees are visible in the Group Policy Object Editor console tree under the **Remote Installation Services** node.

**Where?**

- *Group Policy object name*/User Configuration/Windows Settings/Remote Installation Services

You can customize these settings using Group Policy Object Editor:

**Automatic Setup**
This option supports the predefinition of the computer name and a location within Active Directory for the client computer accounts.

**Custom Setup**
This option supports user definition of a unique name for a computer and specification of where the computer account is created within Active Directory.

**Restart Setup**
This option restarts an operating system installation attempt if it fails prior to completion.

**Tools**
This option supports user access to tools from the Client Installation Wizard.

For each of the Client Installation Wizard options, the following choices are available:

**Enable**
Users to whom this policy applies are offered the specific option.

**Not Configured**
The policy settings of the parent container apply to the specific option. For example, if you choose **Not Configured** and the administrator for the entire domain has set Group Policy specific to Remote Installation Services, the policy that is set on the domain is applied to all users affected by that policy.

**Disabled**
Users affected by this policy cannot access this installation option.

For more information on Remote Installation Services and the options offered by the Client Installation Wizard, see the Remote Installation Services Help, which is accessible in Windows Server 2003 family, either from Group Policy Object Editor or the **%systemroot%\Windows\Help** folder (RISconcepts.chm).

**Notes**

- Windows XP Professional does not show RIS settings in the Group Policy Object Editor (GPOE) unless the Admin Tools Pack is installed.

- If you are running Group Policy Management Console on a computer running Windows XP Professional, Remote

Installation Services Help is not available by default. To access the Help, install Windows Help from the Windows Server 2003 installation CD onto a computer running Windows XP Professional.

For more information about RIS, see Remote Operating System Installation(http://www.microsoft.com/) on the Microsoft Web site.

**See Also**
**Concepts**
Group Policy Object Editor Extensions
Security settings overview for GPMC

# Group Policy replication and domain controller selection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy replication and domain controller selection**

In a domain that contains more than one domain controller, Group Policy information takes time to propagate, or replicate, from one domain controller to another. Low bandwidth network connections between domain controllers slow replication. The Group Policy infrastructure has mechanisms to manage these issues.

Each Group Policy object (GPO) is stored partly in the Sysvol on the domain controller and partly in Active Directory. The Group Policy Management Console (GPMC) and Group Policy Object Editor present and manage the GPO as a single unit. For example, when you set permissions on a GPO in GPMC, GPMC is actually setting permissions on objects in both Active Directory and the Sysvol. It is not recommended that you manipulate these separate objects independently outside of GPMC and the Group Policy Object Editor. It is important to understand that these two separate components of a GPO rely on different replication mechanisms. The file system portion is replicated through the File Replication Service (FRS), independently of the replication handled by Active Directory.

The tools used to manage Active Directory and Group Policy, such as GPMC, the Group Policy Object Editor, and Active Directory Users and Computers all communicate with domain controllers. If there are several domain controllers available, changes made to objects like users, computers, organizational units, and GPOs may take time to appear on other domain controllers. The administrator may see different data depending on the last domain controller on which changes were made and which domain controller they are currently viewing the data from.

For example, if you create a GPO on one domain controller and immediately attempt to link it on another domain controller, the operation could fail. In each domain, GPMC uses the same domain controller for all operations in that domain, in order to avoid any synchronization issues. This includes all operations on GPOs, organizational units, and security groups in that domain. In addition, when the Group Policy Object Editor is opened from GPMC, it will also use the same domain controller in use by GPMC. Finally, GPMC uses the same domain controller for all operations on sites within a given forest. This domain controller for sites is used to read and write information about the links to GPOs that exist on any given site; information regarding the GPO itself is obtained from the domain controller of the domain hosting the GPO. This domain controller is used to read and write information about the links to GPOs that exist on any given site; information regarding the GPO itself is obtained from the domain controller of the domain hosting the GPO.

By default, when you add a new domain to the console, GPMC uses the PDC emulator in that domain to help ensure that all administrators are using the same domain controller. For managing sites, GPMC uses the PDC emulator in the user's domain by default. You can change the default choice of domain controller using the **Change Domain Controller** dialog box in GPMC. If you are located at a remote site with a slow connection to the default domain controller, you may want to do this.

It is important for administrators to consider the choice of domain controller in order to avoid replication conflicts particularly because both Active Directory and FRS use multi-master replication. This is especially important to consider because GPO data resides in both Active Directory and on Sysvol, and two independent replication mechanisms must be used to replicate GPO data to the various domain controllers in the domain. If two administrators are simultaneously editing the same GPO on different domain controllers, it is possible for the changes written by one administrator to be overwritten by another administrator, depending on replication latency.

**Important**

- If multiple administrators manage a common GPO, it is recommended that all administrators use the same domain controller when editing a particular GPO, to avoid collisions in FRS.

**Options governing selection of a domain controller for GPMC**

In GPMC, when you right-click a domain or the sites container and click **Change Domain Controller**, you see a **Change Domain Controller** dialog box. The domain controller options for GPMC are:

- **The one with the Operations Master token for the PDC emulator.** This is the default and preferred option.
- **Use any available domain controller.** This is the least safe option.
- **Use any available domain controller that is running Windows Server 2003 or later.** This option is useful if you are restoring deleted GPOs that contain software installation settings. If possible, it is recommended to perform restoration

of GPOs containing software installation settings on domain controllers running Windows Server 2003. For more information, see Restore using GPMC.

- **This domain controller.** This option allows you to choose a specific domain controller from a list of domain controllers in the domain.

If you are changing the domain controller for a site, you can also choose any available trusted domain from the **Look in this domain** drop-down list box in the **Change Domain Controller** dialog box.

When you open the Group Policy Object Editor from GPMC it always uses the same domain controller that is targeted in GPMC for the domain where that GPO is located.

For step-by-step instructions for choosing a domain controller, see Specify a domain controller.

**See Also**
**Concepts**
Deployment considerations for Group Policy
Group Policy Management Console Overview

# Group Policy processing and precedence

Updated: April 3, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy processing and precedence**
The Group Policy objects (GPOs) that apply to a user (or computer) do not all have the same precedence. Settings that are applied later can override settings that are applied earlier.

**Order of processing settings**
This section provides details about the order in which Group Policy settings for users and computers are processed. For information about where the processing of policy settings fits into the framework of computer startup and user logon, see steps 3 and 8 in Startup and logon, in this topic.

Group Policy settings are processed in the following order:

1. **Local Group Policy object—**Each computer has exactly one Group Policy object that is stored locally. This processes for both computer and user Group Policy processing.

2. **Site—**Any GPOs that have been linked to the site that the computer belongs to are processed next. Processing is in the order that is specified by the administrator, on the **Linked Group Policy Objects** tab for the site in Group Policy Management Console (GPMC). The GPO with the lowest **link order** is processed last, and therefore has the highest precedence.

3. **Domain—**Processing of multiple domain-linked GPOs is in the order specified by the administrator, on the **Linked Group Policy Objects** tab for the domain in GPMC. The GPO with the lowest **link order** is processed last, and therefore has the highest precedence.

4. **Organizational units—**GPOs that are linked to the organizational unit that is highest in the Active Directory hierarchy are processed first, then GPOs that are linked to its child organizational unit, and so on. Finally, the GPOs that are linked to the organizational unit that contains the user or computer are processed.

   At the level of each organizational unit in the Active Directory hierarchy, one, many, or no GPOs can be linked. If several GPOs are linked to an organizational unit, their processing is in the order that is specified by the administrator, on the **Linked Group Policy Objects** tab for the organizational unit in GPMC. The GPO with the lowest **link order** is processed last, and therefore has the highest precedence.

This order means that the local GPO is processed first, and GPOs that are linked to the organizational unit of which the computer or user is a direct member are processed last, which overwrites settings in the earlier GPOs if there are conflicts. (If there are no conflicts, then the earlier and later settings are merely aggregated.)

**Exceptions to the default order of processing settings**
The default order for processing settings is subject to the following exceptions:

- A GPO link may be **enforced**, or **disabled**, or both. By default, a GPO link is neither enforced nor disabled.

- A GPO may have its user settings disabled, its computer settings disabled, or all settings disabled. By default, neither user settings nor computer settings are disabled on a GPO.

- An organizational unit or a domain may have **Block Inheritance** set. By default, **Block Inheritance** is not set.

For information about the above modifications to default behavior, see Managing inheritance of Group Policy.

- A computer that is a member of a workgroup processes only the local Group Policy object.

- Loopback may be enabled. For information about loopback, see Loopback processing with merge or replace.

**Startup and logon**

The following sequence shows the order in which computer policy and user policy are applied when a computer starts and a user logs on:

1. Network starts. Remote Procedure Call System Service (RPCSS) and Multiple Universal Naming Convention Provider (MUP) start. (Windows XP Professional is an exception to this rule. By default, on Windows XP Professional, Group Policy processing does not wait for the network to start. This default behavior can be changed by a policy setting.)

2. An ordered list of GPOs is obtained for the computer. The list might depend on these factors:

   - Whether the computer is part of a domain and therefore subject to Group Policy through Active Directory.

   - The location of the computer in Active Directory.

   - Whether the list of Group Policy objects has changed. If the list of GPOs has not changed, no processing is done. You can use a policy setting to change this behavior.

   - Any enforcements/blocks set by the administrator. An administrator can set a policy at a higher level to always apply, a process called enforcement, previously known as no override. Alternatively, an administrator can set a container to block any policies from higher levels from applying. Note: A policy set by the administrator at a higher level to be enforced will continue to be applied even with a block from above set.

3. Computer policy is applied. These are the settings under Computer Configuration from the gathered list. GPOs are applied in the following order: local, site, domain, organizational unit, child organizational unit, and so on. No notification appears while computer policies are processed. Verbose logging to show notification of computer policies processed can be turned on via policy.

   For details about the order in which settings are processed when user or computer policy is applied, see Order of processing settings, in this topic.

4. Startup scripts run. This is hidden and synchronous by default; each script must complete or time out before the next one starts. The default time-out is 600 seconds. You can use several policy settings to modify this behavior.

**Note**

- Any version of Windows XP Professional provides a fast logon optimization feature. By default, computers with these operating systems do not wait for the network to start when they boot up. After logon, policy is processed in the background once the network is available. This means that at logon and startup, the computer will continue to use the earlier policy settings. Therefore, for settings that can only be applied at boot or logon (such as software installation and folder redirection) more than one logon can be required by the user after the initial change is made to the GPO. This policy is controlled by the setting in **Computer Configuration\Administrative Templates\System\Logon\Always wait for the network at computer startup and logon.** This feature is not available on versions of Windows 2000 or Windows 2003 Server.

1. The user presses CTRL-ALT-DEL to log on.

2. After the user is validated, the user profile loads; it is governed by the policy settings that are in effect.

3. An ordered list of GPOs is obtained for the user. The list might depend on these factors:

   - Whether the user is part of a domain and therefore subject to Group Policy through Active Directory.

   - Whether loopback is enabled, and the state (**Merge** or **Replace**) of the loopback policy setting.

   - The location of the user in Active Directory.

- Any enforcements/blocks set by the administrator. An administrator can set a policy at a higher level to always apply, a process called enforcement, previously known as no override. Alternatively, an administrator can set on a container to block any policies from higher levels from applying. Note: A policy set by the administrator at a higher level to be enforced will continue to be applied even with a block from above set.

4. User policy is applied. These are the settings under User Configuration from the gathered list. GPOs are processed in the following order: local, site, domain, organizational unit, child organizational unit, and so on. No notification appears while user policies are processed. (Notification can be turned on through policy.)

   For details about the order in which settings are processed when user or computer policy is applied, see Order of processing settings, in this topic.

5. Logon scripts run. Unlike Windows NT 4.0 scripts, Group Policy–based logon scripts are hidden and asynchronous by default. The user object script runs last in a normal window. There is also a time-out on logon scripts.

6. The operating system user interface that is prescribed by Group Policy appears.

**Important**

- Three special cases deserve consideration during migration:

  - If the computer account object is in a Windows NT 4.0 domain and the user account object is in Active Directory, System Policy for the computer only (not user) is processed when the user logs on. Then, user (not computer) Group Policy is processed.

  - If the computer account object is in Active Directory and the user account object is in a Windows NT 4.0 domain, computer (not user) Group Policy is processed during computer startup. When the user logs on, System Policy for the user (not computer) is processed.

  - If the computer and user accounts of the computer are members of a Windows NT 4.0 domain, only System Policy (not Group Policy) for the computer and user is applied when the user logs on.

**Note**

- Several of these events can be modified. You can set policies to:

  - Reverse the synchronized or asynchronized defaults for running script.

  - Specify when scripts time out. By default, scripts time out after 600 seconds.

  - Change whether scripts are run hidden, minimized, or in a normal window.

**See Also**
**Concepts**
Deployment considerations for Group Policy
Controlling the Scope of Group Policy Objects using GPMC

# Controlling the Scope of Group Policy Objects using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Controlling the Scope of Group Policy Objects**

The process of targeting Group Policy objects (GPOs) to specific users and computers is referred to as scoping the GPO. In order to apply GPOs to users or computers, you must first link them to sites, domains, or organizational units. In addition, you can change the default order in which GPOs are processed by changing the link order, blocking policy inheritance, enforcing a GPO link (previously known as no override), and disabling a GPO link. After linking GPOs, you have two additional options for applying greater precision: security filtering and WMI filtering on the GPO. Loopback processing allows you to ensure that a consistent set of policies is applied to any user logging on to a particular computer.

This section contains the following information:

- Managing inheritance of Group Policy
- Security filtering using GPMC
- WMI filtering using GPMC
- Loopback processing with merge or replace

# Managing inheritance of Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing inheritance of Group Policy**

To apply the settings of a Group Policy object (GPO) to the users and computers of a domain, site, or organizational unit, you can link that domain site or organizational unit to that GPO. You can add one or more GPO links to each domain, site, and organizational unit in Group Policy Management Console. The settings deployed by GPOs linked to higher containers (parent container) in Active Directory are inherited by default to child containers and combine with any settings deployed in GPOs linked to child containers. If multiple GPOs attempt to set a setting to conflicting values, the GPO with the highest precedence sets the setting. GPO processing is based on a last writer wins model, and GPOs that are processed later have precedence over GPOs that are processed sooner. Group Policy objects are processed according to the following order:

1. The local Group Policy object (LPGO) is applied.
2. GPOs linked to sites.
3. GPOs linked to domains
4. GPOs linked to organizational units. In the case of nested organizational units, GPOs associated with parent organizational units are processed prior to GPOs associated with child organizational units.

Links to a specific site, domain, or organizational unit are applied in reverse sequence based on link order. For example, a GPO with Link Order 1 has highest precedence over other GPOs linked to that container.

You can view the precedence order of GPOs for a given site, domain or organizational unit by navigating to the **Group Policy Inheritance** tab for any site, domain, or organizational unit. Note that when looking on the **Group Policy Inheritance** tab for a domain or organizational unit, GPOs linked to sites are not shown. This is the specific site that a computer is in is not known ahead of time. Also, when viewing a site, the only difference between the **Group Policy Inheritance** tab and the **Linked Group Policy Objects** tab is that the former takes into account the enforcement (described below) attribute.

For more background information about GPO link processing and precedence, including the default order for processing, see Group Policy processing and precedence.

You can further control precedence and how GPO links are applied to specific domains, sites, or organizational units by doing the following:

- Changing the link order
- Blocking Group Policy inheritance
- Enforcing a GPO link
- Disabling a GPO link

**Changing the link order**

Within each domain, site, and organizational unit, the link order controls when links are applied. To change the precedence of a link, you can change the link order, moving each link up or down in the list to the appropriate location. The link with the higher order (with 1 being the highest order) has the higher precedence for a given site, domain, or organizational unit. For example, if you add six GPO links and later decide that you want the last one that you added to have highest precedence, you can move the GPO link to the top of the list.

**Blocking Group Policy inheritance**

You can block policy inheritance for a domain or organizational unit. Using block inheritance prevents GPOs linked to higher sites, domains, or organizational units from being automatically inherited by the child-level. By default, children inherit all GPOs from the parent, but it is sometimes useful to block inheritance. For example, if you want to apply a single set of policies to an entire domain except for one organizational unit, you can link the required GPOs at the domain level (from which all organizational units inherit policies by default), and then block inheritance only on the organizational unit to which the policies

should not be applied.

**Enforcing a GPO link**
You can specify that the settings in a GPO link should take precedence over the settings of any child object by setting that link to Enforced. GPO-links that are enforced cannot be blocked from the parent container. Without enforcement from above, the settings of the GPO links at the higher level (parent) are overwritten by settings in GPOs linked to child organizational units, if the GPOs contain conflicting settings. With enforcement, the parent GPO link always has precedence. By default, GPO links are not enforced. In tools prior to GPMC, "enforced" was known as "No override."

**Disabling a GPO link**
By default, processing is enabled for all GPO links. You can completely block the application of a GPO for a given site, domain, or organizational unit by disabling the GPO link for that domain, site, or organizational unit. Note that this does not disable the GPO itself, and if the GPO is linked to other sites, domains or organizational units, they will continue to process the GPO, if their links are enabled.

For more information about these tasks, see Control Group Policy Object Scope.

### Important

- GPO links set to enforce (no override) cannot be blocked.
- The enforce and block inheritance options should be used sparingly. Casual use of these advanced features complicates troubleshooting.

In addition to using GPO links to apply policies, you can also control how GPOs are applied by using security filters or WMI filters.

**See Also**
**Concepts**
Security filtering using GPMC
WMI filtering using GPMC
Controlling the Scope of Group Policy Objects using GPMC

# Security filtering using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security filtering**
Security filtering is a way of refining which users and computers will receive and apply the settings in a Group Policy object (GPO). Using security filtering, you can specify that only certain security principals within a container where the GPO is linked apply the GPO. Security group filtering determines whether the GPO as a whole applies to groups, users, or computers; it cannot be used selectively on different settings within a GPO.

In order for the GPO to apply to a given user or computer, that user or computer must have both **Read** and **Apply Group Policy (AGP)** permissions on the GPO, either explicitly, or effectively though group membership.

By default, all GPOs have **Read** and **AGP** both **Allowed** for the Authenticated Users group. The Authenticated Users group includes both users and computers. This is how all authenticated users receive the settings of a new GPO when it is applied to an organizational unit, domain or site. However, you can change these permissions to limit the scope to a specific set of users, groups, or computers within the organizational unit, domain, or site.

Group Policy Management manages these permissions as a single unit, and displays the security filtering for the GPO on the GPO Scope tab. Using Group Policy Management Console (GPMC), you can add and remove groups, users, and computers to be used as security filters for each GPO. In addition, security principals used for security filtering also appear on **Delegation** tab for a GPO as having **Read (from Security Filtering)**, since they have read access to the GPO.

To modify security filtering, you add or remove groups in the Security Filtering section on the Scope tab of a GPO. In practice, you don't have to set the two access control entries (ACEs), because GPMC sets both for you when you set security filtering. For the step-by-step procedure, see Filter using security groups.

In addition, The **Read** and **AGP** permissions are visible separately, and able to be set independently of one another, through the access control list (ACL) editor. In GPMC, the **Security Filtering** section of the **Scope** tab of a GPO shows only whether the GPO will apply. If you want to see the permissions separately, you can open the ACL editor by clicking the **Advanced** button on the **Delegation** tab for the GPO.

**Notes**

- GPOs cannot be linked directly to users, computers, or security groups. They can only be linked to sites, domains and organizational units. However, by using security filtering, you can narrow the scope of a GPO so that it applies only to a single group, user, or computer.

- Granting **Read** and **AGP** is not sufficient to ensure that the GPO is processed for a user or computer. The GPO also has to be linked to a site, domain or organizational unit containing the user or computer, directly or through inheritance.

- A GPO with security filtering set to **Read** and **AGP** doesn't necessarily apply to all security principals that have security filtering. It only applies to them if those user or computer objects are in the container or child container that is linked to the GPO.

- The location of a security group in Active Directory is irrelevant to security group filtering and, more generally, irrelevant to Group Policy processing.

**See Also**
**Concepts**
WMI filtering using GPMC
Managing inheritance of Group Policy

# WMI filtering using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WMI filtering**

Windows Management Instrumentation (WMI) filters allow you to dynamically determine the scope of Group Policy objects (GPOs) based on attributes of the target computer.

When a GPO that is linked to a WMI filter is applied on the target computer, the filter is evaluated on the target computer. If the WMI filter evaluates to false, the GPO is not applied (except if the client computer is running Windows 2000, in which case the filter is ignored and the GPO is always applied). If the WMI filter evaluates to true, the GPO is applied.

WMI makes data about a target computer available for administrative use. Such data can include hardware and software inventory, settings, and configuration information. For example, WMI exposes hardware configuration data such as CPU, memory, disk space, and manufacturer, as well as software configuration data from the registry, drivers, file system, Active Directory, the Windows Installer service, networking configuration, and application data.

A WMI filter consists of one or more queries based on this data. If all queries are true, the GPO linked to the filter will be applied. The queries are written using the WMI Query Language (WQL), a SQL-like language. Queries can be combined with AND and OR logical operators to achieve whatever effect the administrator wants. Each query is executed against a particular WMI namespace. When you create a query, you must specify the namespace. The default is root\CIMv2, which is appropriate for most WMI queries.

The WMI filter is a separate object from the GPO in the directory. To apply a WMI filter to a GPO, you link the filter to the GPO. This is shown in the WMI filtering section on the Scope tab of a GPO. Each GPO can have only one WMI filter, however the same WMI filter can be linked to multiple GPOs.

WMI filters, like GPOs, are stored on a per-domain basis. A WMI filter and the GPO it is linked to must be in the same domain.

### Notes

- Client support for WMI filters exists only on Windows XP, Windows Server 2003, and later operating systems. Windows 2000 clients will ignore any WMI filter and the GPO is always applied, regardless of the WMI filter.

- WMI filters are only available in domains that have at least one Windows Server 2003 domain controller. In an environment consisting only of Windows 2000 domains, the WMI filter node in Group Policy Management Console (GPMC) is not shown.

**Using WMI filters with GPMC**

Using GPMC, you can create and delete WMI filters, link and unlink WMI filters, copy and paste WMI filters, import and export WMI filters, and view and edit attributes of WMI filters.

**Sample code**

The following table shows sample code for several WMI filters.

| Criterion | Administrator's intent | WMI filter |
|---|---|---|
| Configuration | Avoid turning on netmon on computers that can have multicasting turned on. | `Select * from Win32_NetworkProtocol where SupportsMulticasting = true` |
| Time zone | Apply policy on all servers located on the East Coast of the United States. | `Root\cimv2 ; Select * from win32_timezone where bias =-300` |
| Hotfix | Apply a policy on computers that have a specific hotfix. | `Root\cimv2 ; Select * from Win32_QuickFixEngineering where HotFixID = 'q147222'` |

| | | |
|---|---|---|
| Software inventory | Assign software only on computers already having either of two software packages. | `Root\cimv2;Select * from Win32_Product where name = "MSIPackage1" OR name = "MSIPackage2"` |
| Operating system | Only target computers running Windows XP Professional. | `Root\CimV2; Select * from Win32_OperatingSystem where Caption = "Microsoft Windows XP Professional"` |
| Resources | Target only machines that have at least 600 megabytes (MB) available. | `Root\CimV2; Select * from Win32_LogicalDisk where FreeSpace > 629145600 AND Description <> "Network Connection"` |
| Make or model | Target Toshiba Tecra models 800 and 810. | `Root\CimV2; Select * from Win32_ComputerSystem where manufacturer = "Toshiba" and Model = "Tecra 800" OR Model = "Tecra 810"` |

**See Also**
**Concepts**
Managing inheritance of Group Policy
Security filtering using GPMC
Create, import, export, copy, and paste WMI filters

# Loopback processing with merge or replace

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Loopback processing with merge or replace**
Loopback is an advanced Group Policy setting that is useful on computers in certain closely managed environments, such as servers, kiosks, laboratories, classrooms, and reception areas. Loopback only works when both the user account and the computer account are in a Windows 2000 or later domain. Loopback does not work for computers joined to a workgroup. Setting loopback causes the User Configuration settings in GPOs that apply to the computer to be applied to every user logging on to that computer, *instead of* (in replace mode) or *in addition to* (in merge mode) the User Configuration settings of the user. This allows you to ensure that a consistent set of policies is applied to any user logging on to a particular computer, regardless of their location in Active Directory. Loopback is controlled by the following setting, **User Group Policy loopback processing mode**, which is located under Computer Configuration\Administrative Templates\System\Group Policy in Group Policy Object Editor (GPMC).

By default, a user's policy settings come from the set of GPOs that are applied to the user object in Active Directory. During Group Policy processing on the client, the Group Policy engine assembles an ordered list of GPOs from the site, domain, and all organizational units for that user object.

Loopback can be set to **Not Configured**, **Enabled**, or **Disabled**. In the **Enabled** state, loopback can be set to **Merge** or **Replace**. In either case the user only receives user-related policy settings.

- **Loopback with Replace—**In the case of **Loopback with Replace**, the GPO list for the user is replaced in its entirety by the GPO list that is already obtained for the computer at computer startup (during step 2 in Group Policy processing and precedence). The User Configuration settings from this list are applied to the user.

- **Loopback with Merge—**In the case of **Loopback with Merge**, the Group Policy object list is a concatenation. The default list of GPOs for the user object is obtained, as normal, but then the list of GPOs for the computer (obtained during computer startup) is appended to this list. Because the computer's GPOs are processed after the user's GPOs, they have precedence if any of the settings conflict.

**See Also**
**Concepts**
Get user settings from the machine account (loopback processing)
Managing inheritance of Group Policy

# Resultant set of policy overview for GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resultant Set of Policy**
Resultant set of policy (RSoP) is a feature of Group Policy that makes implementation, troubleshooting, and planning of Group Policy easier. When multiple Group Policy objects (GPOs) apply to a given user or computer, they can contain conflicting policy settings. For most policy settings, the final value of the policy setting is set only by the highest precedent GPO that contains that setting. RSoP helps you understand and identify the final set of policy that is applied as well as settings that did not apply as a result of policy inheritance.

Specifically, Resultant Set of Policy helps you determine the following:

- The final value of the setting that is applied as a result of all the GPOs.
- The final GPO that set the value of this setting (also known as the winning GPO).
- Precedence details that show any other GPOs that attempted to set this setting and the value that each GPO attempted to set for that policy setting.

**Example showing precedence**
Assume that there are three GPOs, GPO-A, GPO-B, and GPO-C, that all apply to the same user. GPO-A has precedence over GPO-B, which has precedence over GPO-C. Each GPO has been configured to set a desktop background. GPO-A sets the value to "Red Moon Desert." GPO-B sets it to "Bliss." GPO-C sets it to "Windows XP." The resultant set of policy data in this example shows that when these three GPOs are applied to a user, the following results occur:

- The background will be set to "Red Moon Desert."
- This background value was set by GPO A because it had highest precedence
- GPO-B attempted to set the background to "Bliss" and GPO-C attempted to set it to "Windows XP". The precedence was GPO-A, followed by GPO-B, followed by GPO-C.

**RSoP in Group Policy Management Console**
The Group Policy Management Console (GPMC) exposes two different resultant set of policy capabilities that are provided by Windows:

- Group Policy Results
- Group Policy Modeling

**Group Policy Results**
This represents the actual policy data that is applied to a given computer and user. It is obtained by querying the target computer and retrieving the RSoP data that was applied to that machine. The Group Policy Results capability is provided by the client operating system and requires Windows XP, Windows Server 2003 or later. Outside of GPMC, Group Policy Results is referred to as RSoP - logging mode.

**Group Policy Modeling**
This is a simulation of what would happen under circumstances specified by an administrator. Group Policy Modeling requires that you have at least one domain controller running Windows Server 2003 because this simulation is performed by a service running on a domain controller that is running Windows Server 2003. With Group Policy Modeling, you can either simulate the RSOP data that would be applied for an existing configuration, or you can perform "what-if" analyses by simulating hypothetical changes to your directory environment and then calculating the RSOP for that hypothetical configuration. For example, you can simulate changes to security group membership, or changes to the location of the user or computer object in Active Directory. Outside of GPMC, Group Policy Modeling is referred to as RSoP - planning mode.

In GPMC, resultant set of policy data is obtained using Group Policy Modeling or Group Policy Results wizards. GPMC provides

an HTML report of the RSoP data. This report shows the final value of the winning settings and the winning GPO that set that value. When you create a Group Policy Modeling or Group Policy Results report, the report is shown in GPMC under the appropriate node. Right-clicking this report and choosing Advanced View opens the RSoP snap-in, which provides additional information, enabling you to verify precedence for a policy setting. The Precedence tab, available from the dialog box for a policy setting, shows all GPOs that attempted to set a particular setting and the value for each GPO.

**RSoP and WMI**

RSoP relies on Windows Management Instrumentation (WMI) for storing and processing information about the policy settings that apply. When Group Policy is processed on a Windows XP or Windows Server 2003 computer, the relevant data is stored in the local WMI repository.

In logging mode, the Group Policy Results feature queries the WMI repository on the target computer, obtains information about the policy settings, and displays it in either the HTML presentation or in RSoP MMC snap-in.

In planning mode, the Group Policy Modeling feature calls the Resultant Set of Policy Provider Service on a domain controller. Administrators use the Group Policy Modeling wizard to specify information about the simulated policy environment such as the target user and computer, membership in security groups, and so forth. The RSOP provider service then simulates the resultant set of policy for this configuration. The results of this simulation are stored to a local WMI database on the domain controller before the information is passed back to GPMC, where it can be easily viewed by the administrator.

**See Also**

**Concepts**

Determine Resultant Set of Policy with Group Policy Results

Simulate Resultant Set of Policy using Group Policy Modeling

# Scripting Group Policy tasks using GPMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scripting Group Policy tasks**

The Group Policy Management Console (GPMC) provides a comprehensive set of COM interfaces for scripting many Group Policy-related operations. The interfaces are documented in the Group Policy Management Console SDK, which is located at **%programfiles%\gpmc\scripts\gpmc.chm** on any computer where you installed GPMC. (The Group Policy Management Console SDK is only available in English.)

When you install GPMC, a set of sample scripts illustrating the use of these interfaces are installed to the **%programfiles%\gpmc\scripts** directory.

The sample scripts address real-world administrative problems and scenarios. You can perform various tasks such as finding all Group Policy objects (GPOs) in a domain that have duplicate names or generating a list of all GPOs in a domain whose settings are disabled or partially disabled.

**Note**

- Scripted control of individual settings inside a GPO is not provided.

The scripts include examples written in VBScript and JScript. They all have the .wsf extension and are executed through Windows Script Host (WSH), which is included with Windows Server 2003 and other recent versions of Windows. All of the scripts should be executed from the command line. Executing a script with the **/?** switch displays the usage for that script.

The sample scripts echo output to the command window and must be executed using cscript.exe. If cscript.exe is not your default scripting host, you will need to explicitly specify cscript.exe on the command line. For example: "d:\program files\gpmc\scripts>cscript ListAllGPOs.wsf".

**Note**

- You can run "cscript //H:cscript" from a command line to make cscript the default scripting host.

You can modify and distribute the sample scripts in accordance with the terms of the End User License Agreement.

Many of the sample scripts rely on a library of common helper functions contained in the file Lib_CommonGPMCFunctions.js. If you copy these scripts to another location, you must also copy this library file to that location for the script samples to work.

**Scripting samples for Group Policy administrative tasks**

The following table shows the administrative tasks you can perform with sample scripts included with Group Policy Management.

| Administrative task | Script name | Description |
|---|---|---|
| Back up all GPOs in a domain | BackupAllGPOs.wsf | Backs up all GPOs in a domain to the specified folder. |
| Back up a GPO | BackupGPO.wsf | Given a GPO name or a globally unique identifier (GUID), backs up the GPO to the specified folder. |
| Copy a GPO | CopyGPO.wsf | Creates a new GPO and copies the settings from the source GPO into the new destination GPO, given a source GPO name or GUID and a new destination GPO name. |
| Create a policy environment using an XML representation | CreateEnvironmentFromXML.wsf | Reads an XML file that specifies a policy environment; for example, organizational units, GPOs, links, and security groups. The script can either create the environment in a domain by creating the objects, or delete the environment by deleting objects specified in the XML file. |

| | | |
|---|---|---|
| Create a GPO with default options | CreateGPO.wsf | Creates a GPO with the specified name, in the current domain, using the default options. |
| Create a migration table | CreateMigrationTable.wsf | Creates migration tables that can be edited and used to map paths and security principals to new values when importing and copying GPOs across domains. |
| Create an XML representation of a policy environment | CreateXMLFromEnvironment.wsf | Reads an existing policy environment and creates an XML file representing that environment. The XML file captures information about organizational units, GPOs, and GPO links, and security on GPOs. You can use this script in conjunction with the CreateEnvironmentFromXML.wsf script to create a replica of a domain for staging purposes. |
| Delete a GPO | DeleteGPO.wsf | Deletes the specified GPO when given a GPO name or GUID. By default the script deletes links to that GPO within the same domain. |
| Grant Permissions for all GPOs in a Domain | GrantPermissionOnAllGPOs.wsf | Grants a user or group the specified level of permission for all GPOs in the specified domain. |
| Import settings into a GPO | ImportGPO.wsf | Imports the settings from the specified backup to an existing destination GPO in the specified domain. |
| Import multiple GPOs into a domain | ImportAllGPOs.wsf | Creates a new GPO and imports settings into that GPO for each backed-up GPO stored at a specific file system location. |
| Restore a GPO | RestoreGPO.wsf | Restores a backed-up GPO. |
| Restore all GPOs | RestoreAllGPOs.wsf | Restores all GPOs that are stored at a given file system location |
| Grant permissions for GPOs linked to a domain, organizational unit, or site | SetGPOPermissionsBySOM.wsf | Grants a user or group the specified permission type for all GPOs that are linked to a specified domain, organizational unit, or site. You can specify **Read, Apply, Edit, FullEdit,** or **None** for the permission type. |
| Set GPO permissions | SetGPOPermissions.wsf | Sets the permission level for a security principal on a given GPO. You can specify **Read, Apply, Edit, FullEdit,** or **None** for the permission type. |
| Set permissions to create GPOs | SetGPOCreationPermissions.wsf | Grants or removes the ability to create GPOs in a domain for a given security principal. |
| Set policy-related permissions on a given site, domain, or organizational unit | SetSOMPermissions.wsf | Sets policy-related permissions on a given site, domain, or organizational unit. |
| List disabled GPOs | FindDisabledGPOs.wsf | Prints all GPOs in the specified domain that are disabled or partially disabled. |

| | | |
|---|---|---|
| List GPO information | DumpGPOInfo.wsf | Prints the information for a specific GPO, including creation time, modification time, owner, status, version number, security groups that filter the GPO, security groups that have full control, edit, read, or custom permissions, and links. |
| List scope of management information | DumpSOMInfo.wsf | Prints information for a specific site, domain, or organizational unit, including GPO links and policy-related permissions. |
| List GPO by policy extension | FindGPOsByPolicyExtension.wsf | Prints all GPOs in the specified domain for which a specific policy extension is configured; for example, find all GPOs that contain the Software Installation or Folder Redirection policy settings. |
| List GPOs by security group | FindGPOsBySecurityGroup.wsf | Prints all GPOs on which a given security principal has the specified permission or effective permission. You can specify **Read, Apply, Edit,** or **Fulledit** for the permission type. |
| List GPOs with duplicate names | FindDuplicateNamedGPOs.wsf | Prints all GPOs in the specified domain that have duplicate names. |
| List GPOs without Apply permission | FindGPOsWithNoSecurityFiltering.wsf | Prints all GPOs in the specified domain that do not apply to anyone because **Apply** permission is not set on the GPO. |
| Listing GPOs Orphaned in SYSVOL | FindOrphanedGPOsInSYSVOL.wsf | Finds and prints all GPOs in SYSVOL with no corresponding component in Active Directory. |
| List domains, organizational units, and sites with external GPO links | FindSOMsWithExternalGPOLinks.wsf | Prints all domains, organizational units, and sites in the specified domain that link to a GPO in a different domain. |
| List unlinked GPOs in a domain | FindUnlinkedGPOs.wsf | Prints all GPOs in the specified domain that have no links. Links outside the domain, including site links, are not checked. |
| Get reports for all GPOs | GetReportsForAllGPOs.wsf | Takes a domain name, and gets reports for all GPOs in that domain. |
| Get reports for GPO | GetReportsForGPO.wsf | Generates XML and HTML reports for a given GPO. |
| List all GPOs in a domain | ListAllGPOs.wsf | Prints all GPOs in the specified domain. |
| Print the scope of management policy tree | ListSOMPolicyTree.wsf | Prints a list of all organizational units in the specified domain with the list of GPOs that are linked to the domain and each organizational unit. |
| List GPO backups in a given file system location | QueryBackupLocation.wsf | Prints information about all backed up GPOs at the file system location specified by the user. |

**See Also**
**Concepts**
Group Policy Object Editor Extensions
Scripts overview for GPMC

# Deployment considerations for Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deployment Considerations for Group Policy**
This topic describes deployment issues related to managing Group Policy in enterprise environments. This includes interoperability issues such as handling a mix of client operating system versions, a mix of domain versions, as well as new features included with Windows Server 2003.

To view an item, click the appropriate link:

- Applying Group Policy on different client operating systems
- Using Administrative template files in mixed environments
- Administering a mix of Windows 2000 and Windows Server 2003 domains
- Upgrading Windows 2000 domains to Windows Server 2003 domains and interaction with Group Policy Modeling
- Using Group Policy features across forests
- Group Policy for sites
- Using Group Policy and Internet Explorer enhanced security configuration

**Applying Group Policy on different client operating systems**
Group Policy can apply to client computers running any of the following operating systems or later:

- Any version of Windows 2000.
- Any version of Windows XP Professional.
- Any version of Windows Server 2003.

Except for Local Group Policy objects (LGPOs), Group Policy requires Active Directory. In order for computer configuration settings to apply, the computer account must be in an Active Directory domain. In order for user configuration settings to apply, the user account must be in an Active Directory domain.

Group Policy is not applied to computers running any of the following operating systems:

- Windows 95
- Windows 98
- Windows Millennium
- Windows NT

**Using Administrative template files in mixed environments**
Administrative templates, (or .adm files), enable administrators to control registry settings using Group Policy. Windows comes with a predefined set of Administrative template files, which are implemented as text files (with an .adm extension), that define the registry settings that can be configured in a Group Policy object (GPO). These .adm files are stored in two locations by default: inside GPOs and in the %windir%\inf folder on the local computer.

As new versions of Windows are released, new policy settings are added. In addition to supporting these new settings, each successive version of Windows supports all registry policy settings that were available in earlier versions of Windows. For

example, the Windows Server 2003 family supports all registry policy settings available in Windows 2000 and Windows XP.

It is important to understand that .adm files are not the actual settings that are deployed to client operating systems. The adm file is simply a template file that provides the friendly name for the setting and an explanation. This template file is used to populate the user interface. The settings that are deployed to clients are contained in the registry.pol file inside the GPO. On Windows XP and Windows Server 2003, each registry setting contains a "Supported on" tag that indicates which operating system versions support that policy setting. If a setting is specified and deployed to a client operating system that does not support that setting, the settings are ignored.

Because all successive iterations of .adm files include settings from earlier versions, and because there is no harm if a new setting is applied inadvertently to a computer running an earlier operating system that does not support that setting, it is recommended to always create and edit GPOs from a computer that has the latest .adm files available. Note that the behavior for handling .adm files in GPMC and the Group Policy Object Editor differs.

**Handling .adm files in Group Policy Object Editor**

- The Group Policy Object Editor uses .adm files to display available policy settings in the Administrative Templates section of a GPO.
- By default it attempts to read .adm files from the GPO (from the Sysvol on the domain controller). Alternatively, the .adm file can be read from the local workstation computer. This behavior can be controlled by a policy setting.
- By default, if the version of the .adm file found on the local computer is newer (based on the time stamp of the file) than the version on the Sysvol, the local version is copied to the Sysvol and is then used to display the settings. This behavior can be controlled by a policy setting.
- If the GPO contains registry settings for which there is no corresponding .adm file, these settings cannot be seen in the Group Policy Object Editor. However, the policy settings are still active and will be applied to users or computers targeted by the GPO.

**Handling .adm files in GPMC**

- GPMC uses .adm files to display the friendly names of policy settings when generating HTML reports for GPOs, Group Policy Modeling, and Group Policy Results.
- By default, GPMC uses the local .adm file, regardless of time stamp. If the file is not found, then GPMC will look in the GPO's folder on Sysvol.
- The user can specify an alternate path for where to find .adm files. If specified, this takes precedence over the previous locations.
- GPMC never copies the .adm file to the Sysvol.

**Administering a mix of Windows 2000 and Windows Server 2003 domains**

GPMC exposes features that are available in the underlying operating system. Because new features have been added to Group Policy since Windows 2000, certain features will only be available in GPMC depending on the operating system that has been deployed on the domain controllers. This section describes these dependencies. In general, there are three key issues that determine whether a feature is available in GPMC:

- Whether the forest supports the Windows Server 2003 schema for Active Directory. Certain features are only available once the schema is upgraded. This is the first step that must be taken before any Windows Server 2003 domain controller can be deployed in an existing Windows 2000 forest. The schema is a forest-wide configuration and is upgraded by running **ADPrep /ForestPrep**. ADPrep is a utility included on the Windows Server 2003 CD. Note that it is possible to have the Windows Server 2003 schema in a forest with all Windows 2000 domain controllers.
- Whether there is at least one domain controller in the forest that is running Windows Server 2003. Group Policy Modeling must be performed on a domain controller running Windows Server 2003.
- Whether a domain contains the Windows Server 2003 domain configuration. This is implemented once **ADPrep**

**/DomainPrep** is run in that domain. This is the first step that must be taken before any Windows Server 2003 domain controller can be deployed in an existing Windows 2000 domain.

Note that there is no dependency from the Group Policy perspective on whether a domain is in native mode or mixed mode.

**Delegation of Group Policy Results and Group Policy Modeling**
In order to delegate either Group Policy Modeling or Group Policy Results, the Active Directory schema in the forest must be the Windows Server 2003 schema. Note that you can use Group Policy Results even without this schema, but only users with local administrative privileges on the target computer can remotely access Group Policy Results data. Thus, if the forest does not have the Windows Server 2003 schema, the delegation pages in GPMC for organizational units and domains will not show these permissions.

**Group Policy Modeling**
Group Policy Modeling is a simulation that is performed by a service that can only run on a domain controller running Windows Server 2003 or later. As long as there is at least one domain controller running Windows Server 2003 in the forest, you can use Group Policy Modeling. GPMC will only show the Group Policy Modeling node in the user interface if the Windows Server 2003 schema is present.

**WMI filtering**
WMI filters are only available in domains that have the Windows Server 2003 configuration. Although none of the domain controllers need to be running Windows Server 2003, you must have run **ADPrep /DomainPrep** in this domain. Also note that WMI filters are only evaluated by clients running Windows XP, Windows Server 2003, or later. WMI filters associated with a Group Policy object will be ignored by Windows 2000 clients and the GPO will always be applied on Windows 2000. For more information, see WMI filtering using GPMC.

If **ADPrep /DomainPrep** has not been run in a given domain, the **WMI Filters** node will not be present, and the GPO scope tab will not have a WMI filters section.

**Upgrading Windows 2000 domains to Windows Server 2003 domains and interaction with Group Policy Modeling**
Group Policy Modeling is a new feature of Windows Server 2003 that simulates the resultant set of policy for a given configuration. The simulation is performed by a service that runs on Windows Server 2003 domain controllers. In order to perform the simulation in cross-domain scenarios, the service must have read access to all GPOs in the forest.

In a Windows Server 2003 domain (whether it is upgraded from Windows 2000 or installed as new), the Enterprise Domain Controllers group is automatically given read access to all newly created GPOs. This ensures that the service can read all GPOs in the forest.

However, if the domain was upgraded from Windows 2000, any existing GPOs that were created before the upgrade do not have read access for the Enterprise Domain Controllers group. When you click a GPO, GPMC detects this situation and notifies the user that Enterprise Domain Controllers do not have read access to all GPOs in this domain. To solve this problem, you can use one of the sample scripts provided with GPMC, **GrantPermissionOnAllGPOs.wsf**. This script can update the permissions for all GPOs in the domain. To use this script:

- Ensure that the person running this script is either a Domain Admin or has permissions to modify security on all GPOs in the domain.
- Open a command prompt and navigate to the **%programfiles%\gpmc\scripts** folder by typing:

  ```
  CD /D %programfiles%\gpmc\scripts
  ```

- Type the following:

  ```
  Cscript GrantPermissionOnAllGPOs.wsf "Enterprise Domain Controllers"
  /Permission:Read /Domain:value
  ```

The value of domain parameter is the DNS name of the domain.

**Using Group Policy features across forests**
The Windows Server 2003 family introduces a new feature called Forest Trust that enables you to authenticate and authorize access to resources from separate, networked forests. With trusts established between forests, you can manage Group Policy throughout your enterprise, which provides greater flexibility especially in large organizations. For more information on forest trusts, see Forests in Group Policy Management Console.

This section describes Group Policy behavior in an environment with forest trust enabled:

- It is not possible to link a GPO to a domain in another forest.

- With Forest trust, it is possible that a user in forest B could log onto a computer in forest A. In this case, when the computer starts up, it will process policy for the computer configuration from Forest A, as usual. When a user from Forest B logs on, where they receive their policy settings from depends on the value of the **Allow Cross-Forest User Policy and Roaming Profiles** policy setting.

    - When this setting is **Not Configured**, no user-based policy settings are applied from the user's forest. Instead, loopback Group Policy processing will be applied, using the Group Policy objects scoped to the computer. Users will receive a local profile instead of their roaming profile.

    - When this setting is **Enabled**, the behavior is exactly the same as Windows 2000 Server, User policy is applied from the user's forest and a roaming user profile is allowed from the trusted forest.

    - When this setting is **Disabled**, the behavior is the same as **Not Configured**.

  This setting is available on Windows Server 2003 located at: **Computer Configuration\Administrative Templates\System\Group Policy\Allow Cross-Forest User Policy and Roaming Profiles.**

- It is possible to deploy Group Policy settings to users and computers in the same forest, but have those settings reference servers in other trusted forests. For example, the shares that host software distribution points, redirected folders, logon scripts, and roaming user profiles could be in another trusted forest.

- Group Policy Modeling requires that both the user and the computer be in the same forest. If you want to simulate a user from Forest A logging on to a computer in Forest B, you must perform two separate Group Policy Modeling simulations: one for the user configuration and the other for the computer configuration.

- Delegation across forests is supported for managing Group Policy. For example, you can delegate to someone in Forest B the ability to perform Group Policy Modeling simulations on objects in Forest A.

**Group Policy for sites**

GPOs that are linked to Active Directory site objects affect all computers in the site. Therefore, any Group Policy object that is linked to a site is applied to all computers in that site, without regard to which domain (in the forest) contains the computers.

This means that computers in multiple domains within a forest can potentially receive the same Group Policy object when you use site-linked GPOs. It also means that a given client computer could receive GPOs from multiple domains in the forest. Note that the GPO exists as a stored entity only in a single domain and it must be read from that domain when the affected clients read their site-linked Group Policy.

If multiple domains are set up across wide area network (WAN) boundaries, the site configuration should take this into account. If it does not, the computers in another domain access a site-linked GPO across a WAN link. This increases the processing time for Group Policy.

In general, it is recommended that you link GPOs to domains and organizational units rather than sites.

**Using Group Policy and Internet Explorer enhanced security configuration**

Windows Server 2003 includes a new default security configuration for Internet Explorer, called Internet Explorer Enhanced Security Configuration (ESC). ESC impacts the Security Zones and Privacy settings within the Internet Explorer Maintenance settings of a GPO. The Security Zones and Privacy settings can either be ESC enabled or not.

- When you edit settings for Security Zones and Privacy settings in a GPO from a computer where ESC is enabled, that GPO will contain ESC-enabled settings. When you look at the HTML report for that GPO, the Security Zones and Privacy header will be appended with the text **(Enhanced Security Configuration enabled)**.

- When you edit settings for Security Zones and Privacy settings in a GPO from a computer where ESC is not enabled , that GPO will contain ESC disabled settings. ESC is not enabled on any computer running Windows 2000 or Windows XP, nor on computers running Windows Server 2003 where ESC has been explicitly disabled.

ESC settings deployed through Group Policy will only be processed on and applied by computers where ESC is enabled. ESC settings will be ignored on computers where ESC is not enabled (all computers running Windows 2000 and Windows XP, and Windows Server 2003 computers where ESC has been explicitly disabled). The converse is also true: A GPO that contains non-ESC settings will only be processed on and applied by computers where ESC is not enabled.

**See Also**
**Concepts**
Managing inheritance of Group Policy
Internet Explorer Maintenance overview for GPMC
Security settings overview for GPMC

# Group Policy Management Console icons reference

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy Management Console icons reference**
Group Policy Management Console (GPMC) provides context-specific icons to dynamically represent the state of console objects. There are different icons for Group Policy objects (GPOs), links, WMI filters, sites, forests, and other objects and containers.

**General usage of icons**
Some icons exist in only one version. For example, the icon for GPMC always looks the same:

Other icons have variations in appearance. The variations indicate the state of the object. For example, when the icon for a domain or organizational unit is shown with a blue circle and exclamation point superimposed, such as:

and ,

this means that inheritance is blocked on that domain or organizational unit.

**Glossary of icons**
The following is a list of icons used in GPMC.

GPO icons

GPO link icons

Icons showing whether inheritance is blocked

WMI filtering icons

Resultant Set of Policy icons

Other GPMC icons

**GPO icons**

| Icon | Description |
|---|---|
| | GPO is either fully enabled or half enabled. Either all settings are enabled, or either user configuration or computer configuration settings are enabled. |
| | GPO that is completely disabled. Both user configuration and computer configuration settings are both disabled. |

**GPO link icons**

| Icon | Description |
|---|---|
| | Enabled, non-enforced GPO link. |
| | Enabled, enforced GPO link. |
| | Enabled link to inaccessible GPO. The GPO exists but is not readable, perhaps because you do not have Read access to it. |
| | Enabled link to nonexistent GPO. Either the GPO has been deleted but the link somehow remains, or the GPO was created on another domain controller and it has not yet replicated to the domain controller that GPMC is using. |
| | Unenforced link to completely disabled GPO. |
| | Enforced link to completely disabled GPO. |
| | Disabled link to enabled or half-enabled GPO. |
| | Enforced and disabled link to enabled or half-enabled GPO. |
| | Disabled link to inaccessible GPO. The GPO exists but is not readable, perhaps because you do not have Read access to it. |

| | |
|---|---|
| | Disabled link to nonexistent GPO. Either the GPO has been deleted but the link somehow remains, or the GPO was created on another domain controller, and it has not yet replicated to the domain controller that GPMC is using. |
| | Disabled link to completely disabled GPO. |
| | Enforced and disabled link to completely disabled GPO. |

**Icons showing whether inheritance is blocked**

| Icon | Description |
|---|---|
| | Domain with inheritance allowed. |
| | Domain with inheritance blocked. |
| | Organizational unit with inheritance allowed. |
| | Organizational unit with inheritance blocked. |

**WMI filtering icons**

| Icon | Description |
|---|---|
| | WMI filters container. |
| | WMI filter. |

**Resultant Set of Policy icons**

| Icon | Description |
|---|---|
| | Group Policy Modeling container. |
| | Group Policy Modeling query. |
| | Group Policy Results container. |
| | Group Policy Results query. |

**Other GPMC icons**

| Icon | Description |
|---|---|
| | Forest. |
| | Domains container. |
| | Security principal that could not be resolved. |
| | Security group. |
| | User object. |
| | Computer object. |
| | GPO container |
| | Sites container. |
| | Site. |

For more information about GPMC, see Group Policy Management Console Overview.

# GPMC Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This section contains links to more information about Group Policy and related technologies.

- Microsoft.com Group Policy Home Page. Provides an entry point for Group Policy documentation on the Web. Includes links to documentation, knowledge base articles, support information, and newsgroups.

- Windows Server 2003 Deployment Kit, Designing a Managed Environment Book. Describes the technologies in Windows Server 2003 associated with deployment of a managed environment. Has significant coverage of Group Policy and related IntelliMirror technologies. Includes planning, designing and implementation guidance.

- Troubleshooting Windows Server 2003 Group Policy White paper. Provides a structured guide to troubleshooting Group Policy operations. Covers Group Policy processing itself, dependent technologies and relevant troubleshooting tools.

- Group Policy Administration using the Group Policy Management Console White paper. Provides technical details of functionality in GPMC.

- Migrating GPOs Across Domains with GPMC. Explains how to migrate GPOs from one domain to another using GPMC.

- Windows Server 2003 Group Policy Infrastructure White paper. Describes a range of topics related to Group Policy at both the server and the client level. Includes detailed Group Policy processing as well as many best practices useful to the Group Policy administrator.

- **Group Policy Management Console Software Development Kit (SDK).** Provides information about how to use the COM interfaces of Group Policy Management, which support scripting many of the operations supported by Group Policy Management Console.

  **Note**

  - This is available when you install GPMC. See "gpmc.chm" in the **%programfiles%\gpmc\scripts** directory. The SDK is only available in English.

# Resultant Set of Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resultant Set of Policy**
You can use Resultant Set of Policy (RSoP) to simulate and test policy settings that are applied to computers or users using Group Policy.

- For help with specific tasks, see Resultant Set of Policy How To....
- For general background information, see RSoP Concepts.
- For problem-solving instructions, see Troubleshooting RSoP.

# Resultant Set of Policy How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Run RSoP
- Use the RSoP Snap-in
- View Policy Settings with RSoP
- Delegate control of RSoP

# Run RSoP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Run RSoP**

- Run an RSoP query on a computer account
- Run an RSoP query on a user account
- Run an RSoP query on a domain
- Run an RSoP query on an organizational unit
- Run an RSoP query on a site
- Run an RSoP query on a local computer

# Run an RSoP query on a computer account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run an RSoP query on a computer account**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the computer account on which you want to run RSoP.

**Where?**

- Active Directory Users and Computers*domain node*\Computers*computer account*

1. Point to **All Tasks**, and then click **Resultant Set of Policy (Logging)** or **Resultant Set of Policy (Planning)**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- You can also run an RSoP query by adding the RSoP snap-in to a Microsoft Management Console (MMC). For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run an RSoP query on a domain
Run an RSoP query on a site
Run an RSoP query on a user account
Run an RSoP query on a local computer
Run an RSoP query on an organizational unit
Open RSoP as an MMC snap-in

# Run an RSoP query on a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run an RSoP query on a user account**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the user account on which you want to run RSoP.

**Where?**

- Active Directory Users and Computers*domain node*\Users*user account*

1. Point to **All Tasks**, and then click **Resultant Set of Policy (Logging)** or **Resultant Set of Policy (Planning)**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- You can also run an RSoP query by adding the RSoP snap-in to a Microsoft Management Console (MMC). For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run an RSoP query on a domain
Run an RSoP query on a site
Run an RSoP query on a computer account
Run an RSoP query on a local computer
Run an RSoP query on an organizational unit
Open RSoP as an MMC snap-in

# Run an RSoP query on a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run an RSoP query on a domain**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the domain on which you want to run RSoP.

**Where?**

- Active Directory Users and Computers*domain node*

1. Point to **All Tasks**, and then click **Resultant Set of Policy (Planning)**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- You can also run an RSoP query by adding the RSoP snap-in to a Microsoft Management Console (MMC). For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run an RSoP query on a user account
Run an RSoP query on a site
Run an RSoP query on a computer account
Run an RSoP query on a local computer
Run an RSoP query on an organizational unit
Open RSoP as an MMC snap-in

# Run an RSoP query on an organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run an RSoP query on an organizational unit**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the organizational unit on which you want to run RSoP.

**Where?**

- Active Directory Users and Computers*domain node**organizational unit**child organizational unit*

1. Point to **All Tasks**, and then click **Resultant Set of Policy (Planning)**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- You can also run an RSoP query by adding the RSoP snap-in to a Microsoft Management Console (MMC). For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run an RSoP query on a domain
Run an RSoP query on a site
Run an RSoP query on a computer account
Run an RSoP query on a local computer
Run an RSoP query on a user account
Open RSoP as an MMC snap-in

# Run an RSoP query on a site

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run an RSoP query on a site**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click the site for which you want to run RSoP.

**Where?**

- Active Directory Sites and Services/*domain node*/Sites/*site*

1. Point to **All Tasks**, and then click **Resultant Set of Policy (Planning)**.

**Notes**

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- If you use this method to open RSoP, you cannot change the site name in the RSoP query.
- You can also run an RSoP query by adding the RSoP snap-in to a Microsoft Management Console (MMC). For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run an RSoP query on a domain
Run an RSoP query on a user account
Run an RSoP query on a computer account
Run an RSoP query on a local computer
Run an RSoP query on an organizational unit
Open RSoP as an MMC snap-in

# Run an RSoP query on a local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run an RSoP query on a local computer**

1. Open Microsoft Management Console.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. On the **Standalone** tab, click **Add**.
4. In the **Available Standalone Snap-in** box, click **Resultant Set of Policy**, and then click **Add**.
5. In Microsoft Management Console, right click **Resultant Set of Policy**, and then click **Generate RSoP Data**.
6. In the Resultant Set of Policy Wizard, click **Next**.
7. On the **Mode Selection** page, click **Logging mode**, and then click **Next**.
8. On the **Computer Selection** page, click **This computer**, and then click **Next** to continue the Resultant Set of Policy Wizard.

## Notes

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- Logging mode is the only mode that is available for an RSoP query on a local computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run an RSoP query on a domain
Run an RSoP query on a site
Run an RSoP query on a computer account
Run an RSoP query on a user account
Run an RSoP query on an organizational unit

# Use the RSoP Snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use the RSoP snap-in**

The following topics describe how to use the Resultant Set of Policy (RSoP) snap-in.

RSoP does not display policy settings introduced in Windows Vista® with Service Pack 1 and later releases. Therefore, to view all policy settings that apply to computers and users on computers running Windows Vista SP1 and later, use the Group Policy Results (GPresult.exe) command-line tool rather than RSoP. For information about Group Policy results, see GPResult (http://go.microsoft.com/fwlink/?LinkId=140844).

- Open RSoP as an MMC snap-in
- Access RSoP data for an existing computer and user (logging mode)
- Simulate RSoP data for a computer and user (planning mode)
- Save data from an RSoP query
- Change an existing RSoP query
- Refresh an RSoP query

# Open RSoP as an MMC snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To open RSoP as an MMC snap-in

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Microsoft Management Console.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. On the **Standalone** tab, click **Add**.
4. In the **Available Standalone Snap-in** box, click **Resultant Set of Policy**, and then click **Add**.

### Notes

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- You must add a Resultant Set of Policy snap-in to Microsoft Management Console (MMC) for each new RSoP query.

#### Using a command line

1. Open Command Prompt.
2. Do one of the following:

    - To open RSoP as an MMC snap-in and display RSoP logging mode for the currently logged on user and computer, type:

      **rsop.msc**

    - To open RSoP as an MMC snap-in and display RSoP logging mode for the specified namespace and target computer, type:

      **rsop.msc /RsopNamespace:***NameSpace***/RsopTargetComp:***TargetComputer*

      **Note**

This syntax is intended for use with a WMI script.

### Note

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
Access RSoP data for an existing computer and user (logging mode)

# Access RSoP data for an existing computer and user (logging mode)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To access RSoP data for an existing computer and user (logging mode)**

1. Open Microsoft Management Console.
2. Add a Resultant Set of Policy snap-in to a Microsoft Management Console (MMC).
3. In the console tree, right-click **Resultant Set of Policy**, and then click **Generate RSoP Data**.
4. In the Resultant Set of Policy Wizard, click **Next**.
5. On the **Mode Selection** page, click **Logging mode**, and then click **Next**.
6. On the **Computer Selection** page, specify the computer on which you want to run RSoP, and then click **Next**.
7. On the **User Selection** page, specify the user for which you want to collect RSoP data, and then click **Next**.
8. On the **Summary of Selections** page, click **Next**, and then wait for RSoP to finish processing the data.
9. On the **Completing the Resultant Set of Policy Wizard** page, click **Finish**.
10. In the console tree, click the newly created RSoP query to view the data.

**Notes**

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- To open RSoP as a Microsoft Management Console (MMC) snap-in, see Related Topics.
- The Resultant Set of Policy MMC snap-in runs the Resultant Set of Policy Wizard to generate Resultant Set of Policy (RSoP) data. On the last page of the Resultant Set of Policy Wizard, there is a **Gather extended error information** check box, which is selected by default. When this check box is selected, the Resultant Set of Policy MMC snap-in enumerates all of the events in the application event log of the target computer. This option is useful when you are using Logging mode because it allows you to collect error information on the client computers, where errors are more likely to occur. You can also use this option in Planning mode, which targets a domain controller, to collect error information in case the RSoP simulation on the domain controller experiences problems.
- Selecting the **Gather extended error information** check box, slows down the processing speed. It is recommended that you use this option if errors appear in the RSoP snap-in after the first attempt. Errors that appear in Logging mode indicate problems that occurred on the target computer the last time that the application of Group Policy was attempted. In Planning mode, the errors might indicate an issue with incorrect permissions or other configuration issue with the domain controller targeted by the simulation.
- Only users and computers that are currently logged on to the domain are visible.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Run RSoP

Open RSoP as an MMC snap-in

Use the RSoP Snap-in

View Policy Settings with RSoP

Save data from an RSoP query

# Simulate RSoP data for a computer and user (planning mode)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To simulate RSoP data for a computer and user (planning mode)**

- 1. Add the Resultant Set of Policy snap-in to a Microsoft Management Console (MMC).
- 2. Start the Resultant Set of Policy Wizard in planning mode, and then click Next.
- 3. On the User and Computer Selection page, specify the user and computer, and then click Next.
- 4. On the Advanced Simulation Options page, specify the simulation options that you want to apply, and then click Next.
- 5. On the User Security Groups page, verify that the information is complete, and then click Next.
- 6. On the Computer Security Groups page, verify that the information is complete, and then click Next.
- 7. On the WMI Filters for Users page, verify the list of WMI filters, and then click Next.
- 8. On the WMI Filters for Computers page, verify the list of WMI filters, and then click Next.
- 9. On the Summary of Selections page, verify the domain controller, and then click Next.
- 10. On the Completing the Resultant Set of Policy Wizard page, click Finish.

**1. Add the Resultant Set of Policy snap-in to a Microsoft Management Console (MMC).**

- For more information, see Related Topics.

**2. Start the Resultant Set of Policy Wizard in planning mode, and then click Next.**

- In the console tree, right-click **Resultant Set of Policy**, and then click **Generate RSoP Data**.
- In the Resultant Set of Policy Wizard, click **Next**.
- On the **Mode Selection** page, click **Planning mode**, and then click **Next**.

**3. On the User and Computer Selection page, specify the user and computer, and then click Next.**

- On the **User and Computer Selection** page, under **User information** and **Computer information**, specify the name of the user and computer, and then click **Next**.

**4. On the Advanced Simulation Options page, specify the simulation options that you want to apply, and then click Next.**

- If you want to simulate a slow network connection, select the **Slow network connection** check box.
- If you want to simulate loopback processing, select the **Loopback processing** check box. If you select the **Loopback processing** check box, click either **Replace** to simulate only applied Group Policy objects (GPOs) based on the computer, or **Merge** to simulate applied GPOs based on the computer and the user.

**5. On the User Security Groups page, verify that the information is complete, and then click Next.**

- The security groups that are currently assigned appear under **Security groups**.

- To search for an alternate user security group, click **Add**, type the name of the security group, click **Check Names**, click a security group, and then click **OK**.

- To restore security groups to those currently assigned to the user, click **Restore Defaults**.

**6. On the Computer Security Groups page, verify that the information is complete, and then click Next.**

- The security groups that are currently assigned appear under **Security groups**.

- To search for an alternate computer security group, click **Add**, type the name of the security group, click **Check Names**, click a security group, and then click **OK**.

- To restore security groups to those currently assigned to the computer, click **Restore Defaults**.

**7. On the WMI Filters for Users page, verify the list of WMI filters, and then click Next.**

- To automatically evaluate all WMI filters that are associated with any Group Policy object (GPO) as true, click **All linked filters**.

- To specify which WMI filters associated with any GPO that you want to automatically evaluate as true, click **Only these filters**. If the filter you want is not in the list, click **List Filters**.

- To remove a filter from the list, click the filter name, and then click **Remove**. Remove any filters that would be considered a false condition for the designated computer.

**8. On the WMI Filters for Computers page, verify the list of WMI filters, and then click Next.**

- To automatically evaluate all WMI filters that are associated with any Group Policy object (GPO) as true, click **All linked filters**.

- To specify which WMI filters associated with any GPO that you want to automatically evaluate as true, click **Only these filters**. If the filter you want is not in the list, click **List Filters**.

- To remove a filter from the list, click the filter name, and then click **Remove**. Remove any filters that would be considered a false condition for the designated computer.

**9. On the Summary of Selections page, verify the domain controller, and then click Next.**

- The wizard might take several minutes to run the simulation.

**10. On the Completing the Resultant Set of Policy Wizard page, click Finish.**

- In the console tree, click the RSoP query to view the data.

### Notes

- To open RSoP as a Microsoft Management Console (MMC) snap-in, see Related Topics.

- The Resultant Set of Policy MMC snap-in runs the Resultant Set of Policy Wizard to generate Resultant Set of Policy (RSoP) data. On the last page of the Resultant Set of Policy Wizard, there is a **Gather extended error information** check box, which is selected by default. When this check box is selected, the Resultant Set of Policy MMC snap-in enumerates all of the events in the application event log of the target computer. This option is useful when you are using Logging mode because it allows you to collect error information on the client computers, where errors are more likely to occur. You can also use this option in Planning mode, which targets a domain controller, to collect error information in case the RSoP simulation on the domain controller experiences problems.

- Selecting the **Gather extended error information** check box, slows down the processing speed. It is recommended that you use this option if errors appear in the RSoP snap-in after the first attempt. Errors that appear in Logging mode indicate problems that occurred on the target computer the last time that the application of Group Policy was attempted. In Planning mode, the errors might indicate an issue with incorrect permissions or other configuration issue with the domain controller targeted by the simulation.

- If the information is complete after your specify your settings in a particular wizard page and you want to run RSoP immediately, select the **Skip to the final page of this wizard without collecting additional data** check box, and then click **Next**.

- You can use this procedure to generate RSoP data based on containers, computers, and users.

- You must be a domain administrator or you must have delegated RSoP (planning) authority over the organizational unit to run RSoP.

- You must be an enterprise administrator to cross domain boundaries in the same forest.

- If you only have authority over one of the destinations, RSoP returns data for that destination only.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
Use the RSoP Snap-in
View Policy Settings with RSoP
Save data from an RSoP query

# Save data from an RSoP query

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save data from an RSoP query**

1. From within an existing RSoP query in Microsoft Management Console (MMC), in the console tree, click **Console Root**, and then in the details pane double-click ***user account* on *computer account* - RSoP**.
2. On the **View** menu, click **Archive data in console file**.
3. On the **File** menu, click **Save**.
4. In **File name**, type a name for the file, and then click **Save**.

## Notes

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- To open RSoP as a Microsoft Management Console (MMC) snap-in, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
Change an existing RSoP query
Refresh an RSoP query

# Change an existing RSoP query

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change an existing RSoP query**

1. From within an existing RSoP query in Microsoft Management Console (MMC), in the console tree, right-click the RSoP query that you want to change, and then click **Change Query**.

2. In the Resultant Set of Policy Wizard, select the appropriate step for the query option that you want to change by clicking **Next**.

3. To close the wizard, click **Finish**.

## Notes

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.

- To open RSoP as a Microsoft Management Console (MMC) snap-in, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
Save data from an RSoP query
Refresh an RSoP query

# Refresh an RSoP query

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To refresh an RSoP query**

- From within an existing RSoP query in Microsoft Management Console (MMC), in the console tree, right-click the RSoP query that you want to refresh, and then click **Refresh Query**.

## Notes

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- To open RSoP as a Microsoft Management Console (MMC) snap-in, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
Change an existing RSoP query
Save data from an RSoP query

# View Policy Settings with RSoP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**View policy settings with RSoP**

- View individual policy settings with RSoP
- View a list of all Group Policy objects associated with an RSoP query
- View the scope of management associated with an RSoP query
- View Group Policy object revision information in RSoP
- View an RSoP report in HTML
- View an RSoP report at the command line

# View individual policy settings with RSoP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view individual policy settings with RSoP**

1. Open an RSoP query.
2. In the console tree, double-click ***user account* on *computer account* - RSoP**, and then click the folders under it to view individual policy settings.

### Note

- To open RSoP, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
Open RSoP as an MMC snap-in

# View a list of all Group Policy objects associated with an RSoP query

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view a list of all Group Policy objects associated with an RSoP query**

1. Open an RSoP query.
2. In the console tree, double-click ***user account* on *computer account* - RSoP**.
3. In the console tree, right-click **Computer Configuration** or **User Configuration**, and then click **Properties**.
4. On the **General** tab, select the **Display all GPOs and filtering status** check box.
5. In the details pane, the list of Group Policy objects appears in the **Group Policy Object** column, and the filtering status appears in the **Filtering** column as either **Applied** or **Not Applied (Empty)**.

## Notes

- To open RSoP, see Related Topics.
- The list includes Group Policy objects that are not applied to the object that is queried.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
View the scope of management associated with an RSoP query
View Group Policy object revision information in RSoP
View Policy Settings with RSoP

# View the scope of management associated with an RSoP query

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the scope of management associated with an RSoP query**

1. Open an RSoP query.
2. In the console tree, double-click ***user account* on *computer account* - RSoP**.
3. Right-click **Computer Configuration** or **User Configuration**, and then click **Properties**.
4. On the **General** tab, select the **Display scope of management** check box.
5. In the details pane, the CNAME of each Group Policy object in Active Directory appears in the **Scope of Management** column.

### Note

- To open RSoP, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
View Group Policy object revision information in RSoP
View Policy Settings with RSoP

# View Group Policy object revision information in RSoP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view Group Policy object revision information in RSoP**

1. Open an RSoP query.
2. In the console tree, double-click ***user account* on *computer account* - RSoP**.
3. Right-click **Computer Configuration** or **User Configuration**, and then click **Properties**.
4. On the **General** tab, select the **Display revision information** check box.
5. In the details pane, the location of the revision information appears in the **Revision** column.

### Note

- To open RSoP, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
View the scope of management associated with an RSoP query
View a list of all Group Policy objects associated with an RSoP query
View Policy Settings with RSoP

# View an RSoP report in HTML

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view an RSoP report in HTML**

1. Click **Start**, and then click **Help and Support**.
2. Under **Support Tasks**, click **Tools**.
3. Under **Tools**, click **Help and Support Center Tools**.
4. Under **Help and Support Center Tools**, click **Advanced System Information**.
5. Under **Advanced System Information**, click **View Group Policy settings applied**.
6. Scroll to the results that you want to view.

### Note

- To hide details, click the arrow in the upper-right corner of a category.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View Group Policy settings in HTML
View Policy Settings with RSoP
Gpresult

# View an RSoP report at the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view an RSoP report at the command line**

1. Open Command Prompt.
2. At the command prompt, redirect the output to a text file by typing:

   **gpresult /z >gp.txt**

3. At the command prompt, open the file by typing:

   **notepad gp.txt**

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view all of the **gpresult** command-line parameters, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Gpresult
View Group Policy settings in HTML

# Delegate control of RSoP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate control of RSoP**

1. Open Active Directory Users and Computers.
2. In the **Name** column, right-click the organizational unit that you want to delegate control of Resultant Set of Policy (RSoP), and then click **Delegate Control**.
3. In the Delegation of Control Wizard, click **Next**.
4. On the **Users or Groups** page, click **Add**.
5. In the **Name** box, type the name of the user or group to which you want to delegate control, and then click **OK**.
6. Verify that the user or group that appears in the **Selected users and groups** box is correct, and then click **Next**.
7. On the **Tasks to Delegate** page, click **Delegate the following common tasks**.
8. Select the **Generate Resultant Set of Policy (logging)** check box or the **Generate Resultant Set of Policy (planning)** check box (you can select both), and then click **Next**.
9. Verify the information that appears, and then click **Finish**.
10. Click the folders in the console tree to view the data in the details pane.

**Notes**

- To open Active Directory Users and Computers, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Delegated control is inherited by all child containers below the container to which control is delegated.
- You must be an enterprise administrator to delegate RSoP control at the domain and site level.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
Run an RSoP query on a domain
Use the RSoP Snap-in

# RSoP Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- RSoP overview
- Using RSoP

# RSoP overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RSoP overview**
Resultant Set of Policy (RSoP) is an addition to Group Policy that makes policy implementation and troubleshooting easier. RSoP is a query engine that polls existing policies and planned policies, and then reports the results of those queries. It polls existing policies based on site, domain, domain controller, and organizational unit. RSoP gathers this information from the Common Information Management Object Model (CIMOM) database (otherwise known as CIM-compliant object repository) through Windows Management Instrumentation (WMI).

RSoP provides details about all policy settings that are configured by an Administrator, including Administrative Templates, Folder Redirection, Internet Explorer Maintenance, Security Settings, Scripts, and Group Policy Software Installation.

When policies are applied on multiple levels (for example, site, domain, domain controller, and organizational unit), the results can conflict. RSoP can help you determine a set of applied policies and their precedence (the order in which policies are applied).

RSoP consists of two modes: planning mode and logging mode. With planning mode, you can simulate the effect of policy settings that you want to apply to a computer and user. Logging mode reports the existing policy settings for a computer and user that is currently logged on.

The Resultant Set of Policy Wizard helps you create an RSoP query. You can open the wizard from Microsoft Management Console (MMC), Active Directory Users and Computers, or Active Directory Sites and Services. You must run the wizard at least once to create an RSoP query. When complete, the wizard displays the query results in the RSoP snap-in in MMC. From here, you can save, change, and refresh your queries. You can create many RSoP queries by adding multiple Resultant Set of Policy snap-ins to MMC, one RSoP snap-in per query.

**RSoP and the CIMOM database**
RSoP uses the CIMOM database through WMI. When a computer logs on to a network, information such as the computer hardware, Group Policy Software Installation settings, Internet Explorer Maintenance settings, Scripts, Folder Redirection settings, and Security Settings, is written to the CIMOM database. When you start RSoP in logging mode, RSoP reports policy settings that have been applied from information provided in the CIMOM database.

**RSoP, Active Directory, and policy precedence**
Unlike the CIMOM database, Active Directory® directory services stores objects regardless of the state of a computer or user. Group Policy uses Group Policy objects (GPOs) in Active Directory to store policy settings. With Group Policy, administrators can:

- Deploy registry keys (Administrative Templates)
- Deploy software (Group Policy Software Installation)
- Deploy security: Safer, IPSec (Security Settings)
- Deploy scripts: logon, logoff, startup, shutdown (Scripts)
- Configure the browser (Internet Explorer Maintenance)

After you define a policy setting for an object, it is applied the next time that object logs on. When an object logs on to a network, the policy settings are applied in the following order:

- Local policy
- Site-level policy
- Domain-level policy
- Domain controller policy (if the domain controller is left in the domain controller container)

- Organizational unit policy

When a Group Policy object overwrites the settings of a different GPO that was applied previously, the new GPO has precedence over the GPO that it has overwritten. When a Group Policy object has a **no overwrite** attribute, it has precedence over all of the policies that are applied subsequently. RSoP can simulate and test the application of policy settings and precedence to Group Policy objects in Active Directory.

**RSoP and Group Policy Software Installation**
A significant part of Group Policy are the software settings extensions, which monitor Group Policy Software Installation. In an RSoP report, RSoP displays which applications are available for any given user or computer, as well as any software setting changes that are advertised or applied. By identifying all of the software that is available for a given user, as well as updates and configuration changes, RSoP makes deployment scenario planning and implementation easier.

**RSoP and security issues**
RSoP provides the following features that you can use to determine which comprehensive security policy meets your needs:

- Provides security templates for creating and assigning security settings for one or more computers. A security template is a file representation of a security setting configuration. It can be applied to a local computer or it can be imported to a Group Policy object in Active Directory. When you import a security template to a Group Policy object, Group Policy processes the security template and makes the corresponding changes to the members of that Group Policy object, which can be users or computers. RSoP verifies those changes. By polling the system and displaying the resultant policy, RSoP indicates a misapplied or overwritten policy setting and the policy setting's precedence, which enables you to fix a security breach.

- RSoP reports the scope of a Group Policy object according to security group membership. RSoP does this through Group Policy filtering.

- Processes and displays the resulting policy for any computer or user. Through individual security settings, Administrators can define a security policy in Active Directory that contains specific security settings for nearly all security areas. Security settings in a local Group Policy object can also establish a security policy on a local computer. When there are conflicts, security settings that are defined in Active Directory always override any security settings that are defined locally.

# Using RSoP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using RSoP**

RSoP provides the following tools that you can use to obtain the information necessary to best configure Group Policy for your organization:

- **The RSoP snap-in--**You can use the RSoP snap-in to create detailed reports about applied policy settings. The RSoP snap-in has two modes: logging mode and planning mode. Logging mode displays policy settings that are applied to computers or users that have logged on. For more information, see RSoP logging mode. Planning mode simulates policy settings that will be applied to a computer or user, or you can use planning mode to check assigned policy settings for a computer that is not currently available or a user that is not currently logged on. For more information, see RSoP planning mode. For information about the RSoP snap-in, see Use the RSoP Snap-in.

- **Advanced System Information - Policy--**In Help and Support Center, you can use **Advanced System Information - Policy** to view an HTML report of Group Policy results. Similar to RSoP logging mode, this tool provides a detailed assessment of policy settings applied to the computer and user that are currently logged on. You can print this report. For more information, see View an RSoP report in HTML.

- **Gpresult--**You can use the **gpresult** command at the command line to get similar detailed information provided by RSoP logging mode. You can export query results to a text file. For more information, see Gpresult.

This section covers:

- Using the RSoP Snap-in

- View Group Policy settings in HTML

# Using the RSoP Snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the RSoP snap-in**

This section covers:

- RSoP logging mode
- RSoP planning mode
- RSoP planning mode options
- Viewing RSoP query results

# RSoP logging mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RSoP logging mode**
Resultant Set of Policy (RSoP) logging mode helps administrators review existing policy settings that have been applied to computers and users.

Logging mode is most helpful in the following situations:

- You want to discover which policy settings are applied to a computer or user.
- You want to discover failed or overwritten policy settings.
- You want to see how security groups affect policy settings.

For more information, see Access RSoP data for an existing computer and user (logging mode) and View individual policy settings with RSoP.

To run RSoP logging mode on a remote computer, you must be logged on as a member of the Domain Admins security group, Enterprise Admins security group, or be delegated **Generate Resultant Set of Policy (logging)** rights. For more information about security groups, see Default groups. For more information about delegating control of RSoP, see Delegate control of RSoP. Any local user account can run a single RSoP logging mode query when logged onto a local computer.

# RSoP planning mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RSoP planning mode**
Resultant Set of Policy (RSoP) planning mode helps administrators plan for growth and reorganization. With RSoP planning mode, you can poll existing Group Policy objects (GPOs) for all of the policy settings that can be applied. Then, you can use the results to construct a "What if?" scenario to predict the effect of changes in policy settings.

Planning mode is most helpful in the following situations:

- You want to simulate the affect of specific policy settings on a computer or user, domain, organizational unit, or site.
- The user is in Active Directory only (for example, a new account).
- You want to test policy precedence in the following situations:
    - The user and the computer are in different security groups.
    - The user and the computer are in different organizational units.
    - The user or the computer is moving to a new location.
- You want to create a slow network simulation.
- You want to create a loopback simulation.

After you decide that a change is necessary, you can run a series of RSoP planning simulations to see what will happen to a user or a group of users if they are moved to another location or security group, or even to another computer. This includes what policy settings will be applied and which files will be loaded automatically after you implement the change. For more information, see RSoP planning mode options.

To run RSoP planning mode on a remote computer, you must be logged on as a member of the Domain Admins security group, Enterprise Admins security group, or be delegated **Generate Resultant Set of Policy (planning)** rights. For more information about security groups, see Default groups. For more information about delegating control of RSoP, see Delegate control of RSoP.

# RSoP planning mode options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RSoP planning mode options**
Resultant Set of Policy (RSoP) provides several planning mode options. You can configure the Resultant Set of Policy Wizard to run each option separately or in conjunction with other options. For more information about running RSoP in planning mode, see RSoP planning mode and Simulate RSoP data for a computer and user (planning mode).

**Slow network connection**
This option simulates a slow network connection in the reporting of Group Policy settings. If the rate at which data is transferred (that is, from the domain controller that provides a policy update to the computers in this group) is slower than the rate that is specified by this policy, the system considers the connection to be slow. The system's response to a slow network connection varies among policy settings. During slow connections, some policy settings might not be applied. RSoP simulates a slow network connection and displays the expected results.

**Loopback processing**
Loopback processing simulates the application of alternate user policy settings when a user logs on to a computer that is affected by the alternate user policy settings. The loopback processing simulation directs the system to apply the set of Group Policy objects (GPOs) for the computer to any user who logs on to a computer that is affected by the alternate user policy settings. It is intended for special-use computers, such as those in public places, laboratories, and classrooms, where you must modify the user policy based on the computer that is being used.

By default, the user's Group Policy objects determine which user policy settings apply. When you apply loopback processing, you need to select one of the following options:

- **Replace--**Indicates that the user policies that are defined in the computer's Group Policy objects replace the user policies that are normally applied to the user.

- **Merge--**Indicates that the user policies that are defined in the computer's Group Policy objects and the user policy settings that are normally applied to the user are combined. If the policy settings conflict, the user policy settings in the computer's Group Policy objects take precedence over the user's existing policy settings.

When you run the query, the computer's Group Policy objects determine which set of Group Policy objects applies.

**Site name**
This option specifies which site and physical location of the subnet that RSoP will use. A site is required. A domain controller, is not required. Use a site to test policy settings that are based on startup or logon being on a subnet other than the one on which the query is currently being run. Use a domain controller to test policy for computers and users that could be operating in location other than the one in which the query is being run, such as a field office.

**Alternate security groups**
This option simulates the application of alternate security groups to both computer and user configurations. With this option you can run a query that simulates what happens when a user or computer belongs to alternate security groups and how this affects the application of the Group Policy object. Administrators can use security groups to further refine which computers and users a Group Policy object influences. For any Group Policy object, administrators can filter the Group Policy object's effect on computers or users that are members of specified security groups. You can set this filtering by using the standard access control list (ACL) editor. To use the ACL editor, click the **Security** tab in a Group Policy object's properties dialog box. Administrators can also use the ACL editor to delegate who can modify the Group Policy object.

**WMI filters**
This option enables RSoP to include configured Windows Management Instrumentation (WMI) filters. With a WMI filter, an administrator can specify a WMI-based query to filter the effect of a Group Policy object based on certain information, such as installed hardware and software. You can find WMI filters in the properties dialog box for a Group Policy object. In addition, RSoP also displays existing WMI filters that are currently enabled.

# Viewing RSoP query results

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing RSoP query results**

After you run the Resultant Set of Policy Wizard, you can view the results from the RSoP query for computer configuration and user configuration in the RSoP snap-in. In the RSoP query in Microsoft Management Console (MMC), under Computer Configuration and User Configuration, the query results for Software Settings, Windows Settings, and Extra Registry Settings appear in the MMC details pane.

**Software Settings**

You can view the Resultant Set of Policy (RSoP) results for Software Settings in the details pane, which displays the five default columns described in the following table.

| Column | Description |
|---|---|
| Name | Lists the name of the deployed package |
| Version | Lists the software version of the deployed package |
| Deployment state | Indicates whether the package is assigned or published |
| Source | Specifies the source location of the deployed package |
| Origin | Lists the name of the Group Policy object that deployed the package |

**Windows Settings**

You can view the Resultant Set of Policy (RSoP) results for Windows Settings in the details pane, which contains results for the following:

- **Scripts**--If a policy setting applies scripts, four default columns appear in the details pane, which are described in the following table.

  | Column | Description |
  |---|---|
  | Name | Lists the name of the script |
  | Parameters | Lists any parameters that are assigned to the script |
  | Last Executed | Lists the date that the script was last run |
  | GPO Name | Lists the name of the Group Policy object that assigned the script |

- **Security Settings**--The Security Settings results in Windows Settings appear the same way that they do for Group Policy. However, when you double-click a policy setting, the **Security Policy Settings** tab is not available and there is a new **Precedence** tab. The **Precedence** tab that indicates which Group Policy objects affect this policy setting from newest to oldest.

- **Internet Explorer Maintenance**--The Internet Explorer Maintenance results appear the same way that they do for Group Policy. However, for each policy setting there is an additional **Precedence** tab that indicates which Group Policy objects affect this policy setting in order from newest to oldest.

- **Administrative Templates**--You can view the Resultant Set of Policy (RSoP) results for Administrative Templates in the details pane. The details pane appears the same way that it does for Group Policy, except that only those policy settings that have been changed from the default setting appear and there is an additional column called **GPO Name**, which displays the Group Policy object name.

  To view more information about an Administrative Template setting, in the details pane double-click the setting. A dialog box appears with three tabs. The following table lists and describes these tabs.

| Tab | Description |
| --- | --- |
| Setting | Indicates the state of the policy setting (unavailable) |
| Explain | Describes the policy setting |
| Precedence | Indicates Group Policy object precedence, from newest to oldest |

**Extra Registry Settings**

You can view the Resultant Set of Policy (RSoP) results for Extra Registry Settings in the details pane. Extra Registry Settings is a collection of policy-based registry settings that do not have an Administrative Template associated with them. Because there is no association with an Administrative Template, there is no explanation of the setting except for the registry key and the Group Policy object that set the key. If no Extra Registry Settings are applied, Extra Registry Settings does not appear in the RSoP query results.

### Note

- To view an HTML report of Group Policy policy settings, you can use **Advanced System Information - Policy** in Help and Support Center. You can use this feature to quickly determine which policy settings are currently being applied, without running RSoP logging mode. For more information, see View Group Policy settings in HTML.

# View Group Policy settings in HTML

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**View Group Policy settings in HTML**

In Help and Support Center, the **Advanced System Information - Policy** tool provides an HTML report of Group Policy settings and Resultant Set of Policy (RSoP) as applied to the local computer and currently logged on user. This tool uses the same functions as RSoP logging mode to determine policy settings specified in the system logging provider (CIMOM database through WMI). Although this tool does not provide the level of detail provided by the RSoP snap-in, it can help you quickly determine the policy settings that are currently applied. For more information, see View an RSoP report in HTML.

Use **Advanced System Information - Policy** to view the following:

- Group Policy results for local computer, which includes the following:
    - Computer Information: Computer name, domain, site, and last time Group Policy was applied
    - Applied Group Policy objects
    - Security group membership when Group Policy was applied
    - Startup scripts
    - Shutdown scripts
    - Security settings - Restricted groups
    - Security settings - File system
    - Security settings - Registry
    - Programs Installed (by using Group Policy Software Installation)
    - Programs listed in Add or Remove Programs (determined by last use of current user)
    - Registry settings
- Group Policy results for current user, which includes the following:
    - User Information: User name, domain, and last time Group Policy was applied
    - Applied Group Policy objects
    - Security group membership when Group Policy was applied
    - Internet Explorer automatic browser configuration
    - Internet Explorer proxy server settings
    - Logon scripts
    - Logoff scripts
    - Redirected folders

- Programs installed (by using Software Installation)
- Programs listed in Add or Remove Programs (determined by last use of current user)
- Registry settings

**Note**

- The RSoP snap-in provides more detailed information about applied policy settings. For more information, see Using the RSoP Snap-in.

# Troubleshooting RSoP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The Active Directory query for the specified user, user container, computer, or computer container failed.
- I received the error message "The container path entered is actually a user object. Please select a container which holds user objects instead."
- I received the error message "The container path entered is actually a computer object. Please select a container which holds computer objects instead."
- I received the error message "Failed to generate RSoP data due to the error listed below. (The planning mode provider failed.)"
- I received the error message "The DomainName domain does not contain a domain controller with the RSoP planning mode functionality available."
- I received the error message "The domain controller you have selected does not have the RSoP planning mode functionality available. Please select a different domain controller."
- I received the error message "The selected computer currently does not have RSoP logging enabled. Do you want to continue?"
- I received the error message "The RSoP snap-in was unable to enumerate the users on the selected computer due to the error listed below."
- I received the error message "The selected computer does not support RSoP logging. RSoP logging support is available in operating system releases after Windows 2000."

**The Active Directory query for the specified user, user container, computer, or computer container failed.**
**Cause:** The domain controller that holds this object is unavailable, or the domain controller that holds this object is not responding to requests for information.

**Solution:** Do the following:

- Verify that the domain controller is available.
- Check the network connections, responses, and throughput.

**Cause:** The client does not have access to the Read properties of the object that is queried.

**Solution:** Verify that the client has access to the Read properties of the object that is queried.

**I received the error message "The container path entered is actually a user object. Please select a container which holds user objects instead."**
**Cause:** The indicated path points to a user account instead of to a container.

**Solution:** Verify that the path points to a container object.

**I received the error message "The container path entered is actually a computer object. Please select a container which holds computer objects instead."**
**Cause:** The indicated path points to a computer account instead of to a container.

**Solution:** Verify that the path points to a container object.

**I received the error message "Failed to generate RSoP data due to the error listed below. (The planning mode provider failed.)"**
**Cause:** The planning mode provider failed.

**Solution:** See the event log file for more details.

**I received the error message "The DomainName domain does not contain a domain controller with the RSoP planning mode functionality available."**
**Cause:** This message appears when the client browses for a domain controller. The client polls a Windows 2000 domain controller, but Windows 2000 does not have planning mode functionality. The user is informed that no domain controller is available.

**Solution:** Connect to a different domain, or upgrade the domain to Windows Server 2003, Standard Edition or Windows Server 2003, Enterprise Edition.

**I received the error message "The domain controller you have selected does not have the RSoP planning mode functionality available. Please select a different domain controller."**
**Cause:** This message is generated when the client browses for a domain controller. The client is running in a Windows 2000 domain, but Windows 2000 does not have planning mode functionality. The user is informed that no domain controller is available.

**Solution:** Connect to a different domain controller, or upgrade the domain controller and domain to Windows Server 2003, Standard Edition or Windows Server 2003, Enterprise Edition.

**I received the error message "The selected computer currently does not have RSoP logging enabled. Do you want to continue?"**
**Cause:** The client is trying to poll a computer that does not have the **Turn off Resultant Set of Policy logging** policy setting disabled.

**Solution:** Disable the **Turn off Resultant Set of Policy logging** policy setting.

**I received the error message "The RSoP snap-in was unable to enumerate the users on the selected computer due to the error listed below."**
**Cause:** The RSoP snap-in cannot enumerate a list of users on a specific domain controller, usually because the client user does not have Read permissions for that domain controller.

**Solution:** Try a different domain controller, or obtain Read permissions for the user.

**I received the error message "The selected computer does not support RSoP logging. RSoP logging support is available in operating system releases after Windows 2000."**
**Cause:** The client is trying to poll a computer with Windows 2000, and Windows 2000 does not support RSoP logging mode.

**Solution:** Poll a different computer, or update the operating system to Windows XP Professional or a product in the Windows Server 2003 family.

# Computer Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computer Management**

You can use Computer Management to manage local or remote computers.

- For help with specific tasks, see Computer Management How To....
- For general background information, see Computer Management Concepts.
- For problem-solving instructions, see Troubleshooting Computer Management.

# Computer Management How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Connect to another computer
- Export information from Computer Management

# Connect to another computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to another computer**

1. Open Computer Management.
2. In the console tree, right-click **Computer Management**, and then click **Connect to another computer**.
3. In the **Select Computer** dialog box, type the name of the computer that you want to manage, or click **Browse** to find it, and then click **OK**.

## Notes

- You must have the appropriate user rights or you must have been delegated the appropriate authority on the target computer to connect to that computer.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- The name of the computer that you are managing appears in parentheses next to Computer Management in the console tree.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections
Change virtual memory settings
**Other Resources**
Send a console message to all connected computers
Send a console message to a computer that you are managing

# Export information from Computer Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export information from Computer Management**

1. Open Computer Management.
2. In the console tree, select the item that you want to export information for.
3. On the **Action** menu, click **Export List**. This opens the **Export List** dialog box.
4. In **Save in**, click or type the file location where you want to save the exported file.
5. In **File name**, type the name of the file.
6. In **Save as type**, click the file type, and then click **Save**.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- The name of the computer that you are managing appears in parentheses next to Computer Management in the console tree.
- You can export information to a tab-delimited text (.txt) file, which is the default, or to a comma-delimited (.csv) file.
- This export functionality may not be available for some items in the console tree.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to another computer
**Other Resources**
Send a console message to all connected computers
Send a console message to a computer that you are managing

# Computer Management Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Computer Management:

- Computer Management Overview
- Understanding Computer Management

# Computer Management Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computer Management overview**
Computer Management is a collection of administrative tools that you can use to manage a single local or remote computer. It combines several administrative utilities into a console tree, and it provides easy access to administrative properties and tools.

You can use Computer Management to:

- Monitor system events, such as logon times and application errors.
- Create and manage shared resources.
- View a list of users who are connected to a local or remote computer.
- Start and stop system services, such as Task Scheduler and Indexing Service.
- Set properties for storage devices.
- View device configurations and add new device drivers.
- Manage applications and services.

Computer Management contains three items: System Tools, Storage, and Services and Applications. For more information, see:

- System Tools overview
- Storage overview
- Services and Applications overview

**Note**

- If you are not a member of the Administrators group, you may not have the administrative credentials required to view or modify some properties or to perform some tasks that Computer Management hosts.

# System Tools overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Tools overview**

System Tools is the first item in the Computer Management console tree. You can use the default tools--Event Viewer, Shared Folders, Local Users and Groups, Performance Logs and Alerts, and Device Manager--to manage system events and performance on the specified computer.

- For more information about Event Viewer, see Event Viewer.
- For more information about Shared Folders, see Shared Folders.
- For more information about Local Users and Groups, see Local Users and Groups.
- For more information about Performance Logs and Alerts, see Performance Logs and Alerts overview.
- For more information about Device Manager, see Device Manager.

# Storage overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Storage overview**
Storage is the second item in the Computer Management console tree. It displays storage devices that are installed on the specified computer. You can use the default tools--Removable Storage, Disk Defragmenter, and Disk Management--to manage the properties of storage devices.

- For more information about Removable Storage, see Removable Storage.
- For more information about Disk Defragmenter, see Disk Defragmenter.
- For more information about Disk Management, see Disk Management.

# Services and Applications overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services and Applications overview**
Services and Applications is the third item in the Computer Management console tree. You can use the default tools--Services, WMI Control, Telephony, and Indexing Service--to help you manage services and applications on the specified computer. For example, you can use Services to view and manage the properties of the Net Logon service.

- For more information about Services, see Services Snap-in.
- For more information about WMI Control, see WMI Overview.
- For more information about Telephony, see Telephony.
- For more information about Indexing Service, see Indexing Service.

**Note**

- The actual set of tools and applications that you see may vary, depending on which services are installed on the specified computer. You may not see the installed services and applications if the specified computer is a remote computer and either of the following conditions is true: you are not a member of the Administrators group on the specified computer or the Remote Registry service is not running on the specified computer.

# Understanding Computer Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Computer Management**
Computer Management has a two-pane view that is similar to Windows Explorer. The console tree is in the left pane, and the details pane is in the right pane. When you select an item in the console tree, information about that item appears in the details pane. This information might include details about the contents, data, or subtools that are available for the selected item. For example, if you click **Sessions** under Shared Folders in System Tools, you can view information about the current user sessions, if any, on the target computer.

The console tree contains three items: System Tools, Storage, and Services and Applications. For more information about these items, see:

- System Tools overview
- Storage overview
- Services and Applications overview

For more information about working with a console tree in Microsoft Management Console (MMC), see Microsoft Management Console.

**Note**

- If you are not a member of the Administrators group, you may not have the administrative credentials required to view or modify some properties or to perform some tasks that Computer Management hosts.

# Troubleshooting Computer Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- I cannot access the tools that are hosted in Computer Management on a remote computer.

**I cannot access the tools that are hosted in Computer Management on a remote computer.**
**Cause:** The Remote Registry service is not running on the remote computer.

**Solution:** Ensure that the Remote Registry service is started on the remote computer. For more information, see Start, stop, pause, resume, or restart a service. To start the service, you need the appropriate administrative credentials on the remote computer.

**See also:** Troubleshooting Performance monitoring for System Monitor and Performance Logs and Alerts

**Note**

- You may also receive this error message if the remote computer is running Windows 95. Computer Management does not support remote access to computers running Windows 95.

# Local Users and Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local Users and Groups**

Use Local Users and Groups to create and manage users and groups that are stored locally on a computer.

- For tips about using Local Users and Groups, see Local Users and Groups Best practices.
- For help with specific tasks, see Local Users and Groups How To....
- For general background information, see Local Users and Groups Concepts.
- For problem-solving instructions, see Troubleshooting Local Users and Groups.

# Local Users and Groups Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **As a security best practice, it is recommended that you do not log on to your computer with administrative credentials.**

  When you are logged on to your computer without administrative credentials, you can use Run as to accomplish administrative tasks.

  For more information, see Using Run as.

- **To further secure your local computer, it is recommended that you implement the following security guidelines:**

  - Limit the number of users in the Administrators group since members of the Administrators group on a local computer have Full Control permissions on that computer.

    For more information, see Why you should not run your computer as an administrator.

  - Rename or disable the Administrator account.

    For more information, see Local user accounts.

  - Leave the Guest account disabled. The Guest account is used by people who do not have an actual account on the computer. The Guest account does not require a password, so it is a security risk. The Guest account is disabled by default, and it is recommended that it stay disabled.

    For more information, see Local user accounts.

  - Some default user rights assigned to specific default local groups may allow members of those groups to gain additional rights on your computer, including administrative rights. Therefore, your organization must equally trust all personnel that are members of the Administrators, Power Users, Print Operators, and Backup Operators groups.

    For more information about these groups, see Default local groups.

  - Review important security considerations about local users and groups.

    For more information, see Default security settings for groups.

- **Use passwords no longer than 14 characters if you are on a network with computers running Windows 95 and Windows 98.**

  You can create a password containing up to 127 characters. However, computers running Windows 95 and Windows 98 support passwords up to only 14 characters. If your password is longer, you may not be able to log on to the network from those computers.

  For more information, see Create a local user account.

# Local Users and Groups How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Manage Local Users
- Manage Local Groups
- Start programs as an administrator

# Manage Local Users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage local users**

- Create a local user account
- Reset the password for a local user account
- Disable or activate a local user account
- Delete a local user account
- Rename a local user account
- Assign a logon script to a local user account
- Assign a home folder to a local user account

# Create a local user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a local user account**

1. Open Computer Management.
2. In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

3. On the **Action** menu, click **New User**.
4. Type the appropriate information in the dialog box.
5. Select or clear the check boxes for:

   - **User must change password at next logon**
   - **User cannot change password**
   - **Password never expires**
   - **Account is disabled**

6. Click **Create**, and then click **Close**.

**Notes**

- To perform this procedure, you must be a member of the Power Users group, or the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- A user name cannot be identical to any other user or group name on the computer being administered. It can contain up to 20 uppercase or lowercase characters, except for the following:

  " / \ [ ] : ; | = , + * ? < >

  A user name cannot consist solely of periods (.) or spaces.

- In **Password** and **Confirm password**, you can type a password containing up to 127 characters. However, if the network consists of computers running Windows 95 or Windows 98, consider using passwords no longer than 14 characters. If your password is longer, you may not be able to log on to the network from those computers.
- The use of strong passwords and appropriate password policies can help protect your computer from attack. For more

information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Local user accounts
Why you should not run your computer as an administrator
Assign a logon script to a local user account
Assign a home folder to a local user account
Strong passwords

# Reset the password for a local user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reset the password for a local user account**

1. Open Computer Management.
2. In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

3. Right-click the user account for which you want to reset the password, and then click **Set Password**.
4. Read the warning message, and if you want to continue, click **Proceed**.
5. In **New password** and in **Confirm password**, type a new password and then click **OK**.

## Notes

- To perform this procedure, you must be the member of the Power Users group who created the account or a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Strong passwords improve computer and network security. For more information, see Related Topics.
- Resetting a local user account password may result in some data loss for that user. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Strong passwords
Preventing data loss from password resets
Best practices for security

# Disable or activate a local user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable or activate a local user account**

1. Open Computer Management.
2. In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

3. Right-click the user account you want to change, and then click **Properties**.
4. Do one of the following:

   - To disable the selected user account, select the **Account is disabled** check box.
   - To activate the selected user account, clear the **Account is disabled** check box.

## Notes

- To perform this procedure, you must be a member of the Power Users group, or the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- When a user account is disabled, the user is not permitted to log on. The account appears in the details pane with an X on the icon.
- Before activating a disabled account, ensure that the account was not locked out for security reasons. For more information, see Related Topics.
- When a user account is activated, the user is permitted to log on normally.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Local Users and Groups Best practices
Local user accounts
Account lockout policy overview

# Delete a local user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a local user account**

1. Open Computer Management.
2. In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

3. Right-click the user account that you want to delete, and then click **Delete**.

## Notes

- To perform this procedure, you must be the member of the Power Users group who created the account or a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- When you need to delete a user account, it is recommended that you first disable the account. When you are certain that disabling the account has not caused a problem, you can safely delete it.
- A deleted user account cannot be recovered.
- The Administrator and Guest account cannot be deleted.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Local user accounts
Delete a local user account
Disable or activate a local user account

# Rename a local user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a local user account**

1. Open Computer Management.
2. In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

3. Right-click the user account you want to rename, and then click **Rename**.
4. Type the new user name, and then press ENTER.

## Notes

- To perform this procedure, you must be the member of the Power Users group who created the account or a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Because it retains its security identifier (SID), a renamed user account retains all its other properties, such as its description, password, group memberships, user profile, account information, and any assigned permissions and user rights.
- A user name cannot be identical to any other user or group name of the computer being administered. It can contain up to 20 uppercase or lowercase characters, except for the following:

  " / \ [ ] : ; | = , + * ? < >

  A user name cannot consist solely of periods (.) or spaces.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Local user accounts
Assign a logon script to a local user account
Assign a home folder to a local user account

# Assign a logon script to a local user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign a logon script to a local user account**

1. Open Computer Management.
2. In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

3. Right-click the user account you want, and then click **Properties**.
4. On the **Profile** tab, in **Logon script**, type the file name and relative path of the script.

## Notes

- To perform this procedure, you must be a member of the Power Users group, or the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Logon scripts can contain malicious commands. It is recommended that you be familiar with the contents of the logon script before assigning it to a user.
- Logon scripts that are stored on a local computer apply only to users logging on to that local computer.
- Local logon scripts must be stored in a shared folder, or subfolders of the shared folder, named **Netlogon**. If this folder does not exist by default, you must create it. To specify a logon script that is stored in a subfolder of the Netlogon folder, precede the file name with the relative path to that folder. For example, to assign the Startup.bat logon script stored in *ComputerName*\Netlogon*FolderName* to a local user, in **Logon script**, type: *FolderName*\Startup.bat

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Logon script assignment
Assign a home folder to a local user account
Add a member to a local group

# Assign a home folder to a local user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign a home folder to a local user account**

1. Open Computer Management.
2. In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

3. Right-click the user account for which you want to specify a home folder, and then click **Properties**.
4. On the **Profile** tab, do one of the following:

   - To specify a local home folder, click **Local path**, and then type the path. For example, c:\users\nicolette
   - To specify a home folder on a shared resource, click **Connect**, click the proper drive letter, and then type the network path. For example, \\airedale\users\nathan

## Notes

- To perform this procedure, you must be a member of the Power Users group, or the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- The My Documents folder provides a convenient alternative to home folders, but does not replace them. Each user has a My Documents folder on the boot volume. To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.
- If no home folder is assigned, the system assigns a default local home folder to the user account (on the root directory where the operating system files are installed as the initial version).
- To specify a network path for the home folder, you must first create the shared resource and set permissions that allow the user access.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Managing files for local user accounts
Shared Folders overview
Assign a logon script to a local user account
Add a member to a local group

# Manage Local Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage local groups**

- Create a local group
- Add a member to a local group
- Identify members of a local group
- Delete a local group
- Allow members of the Anonymous Logon group to be members of the Everyone group on a local computer

# Create a local group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a local group**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management.
2. In the console tree, click **Groups**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Groups

3. On the **Action** menu, click **New Group**.
4. In **Group name**, type a name for the new group.
5. In **Description**, type a description of the new group.
6. To add one or more members to the new group, click **Add**.
7. In the **Select Users, Computers, or Groups** dialog box, do the following:

   - To add a user or group account to this group, under **Enter the object names to select**, type the name of the user account or group account that you want to add, and then click **OK**.
   - To add a computer account to this group, click **Object Types**, select the **Computers** check box, and then click **OK**. Under **Enter the object names to select**, type the name of the computer account that you want to add, and then click **OK**.

8. In the **New Group** dialog box, click **Create**, and then click **Close**.

## Notes

- To perform this procedure, you must be a member of the Power Users group, or the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- A local group name cannot be identical to any other group or user name on the local computer being administered. It can contain up to 256 uppercase or lowercase characters, except for the following:

" / \ [ ] : ; | = , + * ? < >

A group name cannot consist solely of periods (.) or spaces.

#### Using a command line

1. Open Command Prompt.

2. To create a group, type:

   **net localgroup***GroupName***/add**

### Notes

- To perform this procedure, you must be a member of the Power Users group, or the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- A local group name cannot be identical to any other group or user name on the local computer being administered. It can contain up to 256 uppercase or lowercase characters, except for the following:

  " / \ [ ] : ; | = , + * ? < >

  A group name cannot consist solely of periods (.) or spaces.

- For more information about the **net localgroup** command, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Default local groups
Add a member to a local group
**Other Resources**
Net localgroup

# Add a member to a local group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a member to a local group**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management.
2. In the console tree, click **Groups**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Groups

3. Right-click the group in which you want to add a member, click **Add to Group**, and then click **Add**.
4. In the **Select Users, Computers, or Groups** dialog box, do the following:

   - To add a user account or group account to this group, under **Enter the object names to select**, type the name of the user account or group account that you want to add to the group, and then click **OK**.
   - To add a computer account to this group, click **Object Types**, select the **Computers** check box, and then click **OK**. Under **Enter the object names to select**, type the name of the computer account that you want to add to the group, and then click **OK**.

## Notes

- To perform this procedure, you must be the member of the Power Users group who created the account or a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To remove a member from a local group, select the user account, computer account, or group account in **Members**, and then click **Remove**.
- All the rights and permissions assigned to a group are assigned to all members of that group. For more information, see Related Topics.
- Limit the number of users in the Administrators group since members of the Administrators group on a local computer have Full Control permissions on that computer.
- If the computer is joined to a domain, you can add user accounts, computer accounts, and group accounts from that domain and from trusted domains to a local group.

**Using a command line**

1. Open Command Prompt.

2. To add a member to the Power Users group, type:

   **net localgroup "***PowerUsersGroupName***""***MemberName***"/add**

   You must include the quotation marks.

## Notes

- To perform this procedure, you must be the member of the Power Users group who created the account or a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- All the rights and permissions assigned to a group are assigned to all members of that group. For more information, see Related Topics.

- Limit the number of users in the Administrators group since members of the Administrators group on a local computer have Full Control permissions on that computer.

- If the computer is joined to a domain, you can add user accounts, computer accounts, and group accounts from that domain and from trusted domains to a local group.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Default local groups
Why you should not run your computer as an administrator
**Other Resources**
Net localgroup

# Identify members of a local group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To identify members of a local group

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Computer Management.
2. In the console tree, click **Groups**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Groups

3. Right-click the group in which you want to identify members, and click **Properties**.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

#### Using a command line

1. Open Command Prompt.
2. To list members of the Power Users group, type:

   **net localgroup"***PowerUsersGroupName***"**

   You must include the quotation marks.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- For more information about the **net localgroup** command, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Default local groups
**Other Resources**
Net localgroup

# Delete a local group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a local group**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management.
2. In the console tree, click **Groups**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Groups

3. Right-click the group you want to delete, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Power Users group, or the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- The following default groups cannot be deleted:
  - Administrators
  - Backup Operators
  - Guests
  - Network Configuration Operators
  - Performance Log Users
  - Performance Monitor Users
  - Power Users
  - Print Operators
  - Remote Desktop Users
  - Replicator

    - Users

- A deleted group cannot be recovered.

- Deleting a local group removes only the group; it does not delete the user accounts, computer accounts, or group accounts that were members of that group.

- If you delete a group and then create another group with the same group name, you must set new permissions for the new group; it will not inherit the permissions that were assigned to the old group.

#### Using a command line

1. Open Command Prompt.

2. To delete a group, type:

   **net localgroup***GroupName***/delete**

### Notes

- To perform this procedure, you must be a member of the Power Users group, or the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The following default groups cannot be deleted:

    - Administrators

    - Backup Operators

    - Guests

    - Network Configuration Operators

    - Performance Log Users

    - Performance Monitor Users

    - Power Users

    - Print Operators

    - Remote Desktop Users

    - Replicator

    - Users

- A deleted group cannot be recovered.

- Deleting a local group removes only the group; it does not delete the user accounts, computer accounts, or group

accounts that were members of that group.

- If you delete a group and then create another group with the same group name, you must set new permissions for the new group; it will not inherit the permissions that were granted to the old group.

- For more information about the **net localgroup** command, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Default local groups
Delete a local user account
**Other Resources**
Net localgroup

# Allow members of the Anonymous Logon group to be members of the Everyone group on a local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow members of the Anonymous Logon group to be members of the Everyone group on a local computer**

1. Open Local Security Settings.
2. In the console tree, double-click **Local Policies**, and then click **Security Options**.
3. In the details pane, right-click **Network access: Let Everyone permissions apply to anonymous users**, and then click **Properties**.
4. On the **Local Security Settings** tab, click **Enabled**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Settings, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Local Security Policy**.
- By default, the inclusion of the Everyone security group in the anonymous user access token is disabled.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network access: Let Everyone permissions apply to anonymous users
Network access: Shares that can be accessed anonymously
Default security settings for groups

# Start programs as an administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Start programs as an administrator**

- Run a program with administrative credentials
- Create a shortcut using the runas command

# Run a program with administrative credentials

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run a program with administrative credentials**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. In Windows Explorer, or on the **Start** menu, right-click the program executable file (or icon) that you want to open, and then click **Run as**.
2. Click **The following user**.
3. In **User name** and **Password**, type the administrator account name and password.

## Notes

- The use of Run as is not limited to administrator accounts.
- You can use Run as to start a program in the context of an administrator account. The administrator context is only used for that specific program and is only available until that program is closed.
- You can also use Run as to start two separate instances of an administrative application, such as Active Directory Domains and Trusts. Each instance can be directed at a different domain or forest to more efficiently perform interdomain or interforest administration tasks.
- You can open most Control Panel items using Run as by holding down the SHIFT key, and then right-clicking the item.
- Some programs do not support Run as.
- If you try to start a program, such as Microsoft Management Console (MMC) or a Control Panel item, from a network location using Run as, it could fail if the credentials used to connect to the shared resource are different from the credentials used to start the program. The credentials used to run the program might not be able to gain access to the same shared resource. For more information about troubleshooting the **runas** command, see Related Topics.
- If Run as fails, the Secondary Logon service may not be running. For more information about how to start a service, see Related Topics.

**Using a command line**

1. Open Command Prompt.
2. To run a program, type:

   **runas/user:***UserName* **"***ProgramNamePathToProgramFile***"**

| Value | Description |
|---|---|

| | |
|---|---|
| *UserName* | Specifies the user credentials to use when running the specified program. You can use the format *user*@*domain* or *domain**user*. |
| " *ProgramName PathToProgramFile* " | Specifies the program name and path to the program file name that you want to open with the specified user credentials.<br><br>For example, to open Active Directory Users and Computers, use the following format:<br><br>**"mmc %windir%\system32\dsa.msc"**<br><br>To open a specific file using Notepad, use the following format:<br><br>**"notepad c:\myfile.txt"** |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- The use of the **runas** command is not limited to administrator accounts.
- You can use the **runas** command to start a program in the context of an administrator account. The administrator context is only used for that specific program and is only available until that program is closed.
- You can also use the **runas** command to start two separate instances of an administrative application, such as Active Directory Domains and Trusts. Each instance can be directed at a different domain or forest to more efficiently perform interdomain or interforest administration tasks.
- Some programs do not support the **runas** command.
- If you try to start a program, such as Microsoft Management Console (MMC) or a Control Panel item, from a network location using the runas command, it could fail if the credentials used to connect to the shared resource are different from the credentials used to start the program. The credentials used to run the program might not be able to gain access to the same shared resource. For more information about troubleshooting the **runas** command, see Related Topics.
- If the **runas** command fails, the Secondary Logon service may not be running. For more information about how to start a service, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting Local Users and Groups
Why you should not run your computer as an administrator
Runas
Using Run as
Why you should not run your computer as an administrator
Differences in default security settings
Default security settings for groups
Start, stop, pause, resume, or restart a service
Create a shortcut using the runas command

# Create a shortcut using the runas command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a shortcut using the runas command**

1. Right-click the desktop, point to **New**, and then click **Shortcut**.
2. In **Type the location of the item**, type **runas** and the command parameters you want to use. In the following table, see examples for creating shortcuts.
3. Click **Next**, type a name for the shortcut, and then click **Finish**.

**Examples**

| To create a shortcut to | Type |
|---|---|
| A command prompt with administrative credentials | **runas /user:** *ComputerName* **\administrator cmd** |
| Computer Management with administrative credentials | **runas /user:** *ComputerName* **\administrator "mmc %windir%\system32\compmgmt.msc"** |
| Active Directory Users and Computers with domain administrative credentials | **runas /user:** *DomainName* **\administrator "mmc %windir%\system32\dsa.msc"** |
| Active Directory Users and Computers in another forest. | **runas /netonly /user:** *DomainName* **** *UserName* **"mmc %windir%\system32\dsa.msc"** |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- Use this procedure if you want to perform administrative tasks when you are logged on as a member of another group. The use of the **runas** command is not limited to administrator accounts.
- Microsoft Management Consoles (MMC) will not display the credentials under which it is running. Be aware of this when you are using the **runas** command.
- Some programs do not support the **runas** command.
- If you try to start a program, such as Microsoft Management Console (MMC) or a Control Panel item, from a network location using the **runas** command, it could fail if the credentials used to connect to the shared resource are different from the credentials used to start the program. The credentials used to run the program might not be able to gain access to the same shared resource. For more information about troubleshooting the **runas** command, see Related Topics.
- If the **runas** command fails, the Secondary Logon service may not be running. For more information about how to start a service, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Troubleshooting Local Users and Groups
Why you should not run your computer as an administrator
Why you should not run your computer as an administrator
Differences in default security settings
Default security settings for groups
Runas
Using Run as
Start, stop, pause, resume, or restart a service

# Local Users and Groups Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Local Users and Groups:

- Understanding Local Users and Groups
- Administering Local Users and Groups

# Understanding Local Users and Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Local Users and Groups**
This section covers:

- Local Users and Groups overview
- Local user accounts
- Default local groups

# Local Users and Groups overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local Users and Groups overview**

Local Users and Groups is located in Computer Management, a collection of administrative tools that you can use to manage a single local or remote computer. You can use Local Users and Groups to secure and manage user accounts and groups stored locally on your computer. A local user or group account can be assigned permissions and rights on a particular computer and that computer only. Local Users and Groups is available on the following client and server operating systems:

- Client computers running Microsoft® Windows® 2000 Professional or Windows XP Professional
- Member servers running a product in the Microsoft Windows 2000 Server family or the Windows Server 2003 family
- Stand-alone servers running a product in the Microsoft Windows 2000 Server family or the Windows Server 2003 family

Using Local Users and Groups you can limit the ability of users and groups to perform certain actions by assigning them rights and permissions. A right authorizes a user to perform certain actions on a computer, such as backing up files and folders or shutting down a computer. A permission is a rule associated with an object (usually a file, folder, or printer), and it regulates which users can have access to the object and in what manner.

You cannot use Local Users and Groups to view local user and group accounts once a member server has been promoted to a domain controller. However, you can use Local Users and Groups on a domain controller to target remote computers (that are not domain controllers) on the network. Use Active Directory Users and Computers to manage users and groups in Active Directory.

For more information, see:

- Local user accounts
- Default local groups
- Default security settings for groups
- Active Directory Users and Computers (Windows Server 2003)

# Local user accounts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local user accounts**

The Users folder located in the Local Users and Groups Microsoft Management Console (MMC) displays the default user accounts as well as the user accounts you create. These default user accounts are created automatically when you install a stand-alone server or member server running Windows Server 2003. The following table describes each default user account on servers running Windows Server 2003.

| Default user account | Description |
|---|---|
| Administrator account | The Administrator account has full control of the server and can assign user rights and access control permissions to users as necessary. This account must be used only for tasks that require administrative credentials. It is highly recommended that you set up this account to use a strong password. For more information, see Strong passwords. For additional security considerations for accounts with administrative credentials, see Local Users and Groups Best practices.<br><br>The Administrator account is a member of the Administrators group on the server. The Administrator account can never be deleted or removed from the Administrators group, but it can be renamed or disabled. Because the Administrator account is known to exist on many versions of Windows, renaming or disabling this account will make it more difficult for malicious users to try and gain access to it. For more information about how to rename or disable a user account, see Rename a local user account and Disable or activate a local user account.<br><br>The Administrator account is the account you use when you first set up the server. You use this account before you create an account for yourself.<br><br>**Important**<br><br>• Even when the Administrator account has been disabled, it can still be used to gain access to a computer using Safe Mode. |
| Guest account | The Guest account is used by people who do not have an actual account on the computer. A user whose account is disabled, but not deleted, can also use the Guest account. The Guest account does not require a password. The Guest account is disabled by default, but you can enable it.<br><br>You can set rights and permissions for the Guest account just like any user account. By default, the Guest account is a member of the default Guests group, which allows a user to log on to a server. Additional rights, as well as any permissions, must be granted to the Guests group by a member of the Administrators group. The Guest account is disabled by default, and it is recommended that it stay disabled. |
| HelpAssistant account (installed with a Remote Assistance session) | The primary account used to establish a Remote Assistance session. This account is created automatically when you request a Remote Assistance session and has limited access to the computer. The HelpAssistant account is managed by the Remote Desktop Help Session Manager service and will be automatically deleted if no Remote Assistance requests are pending. For more information about Remote Assistance, see Administering Remote Assistance. |

For more information about creating and managing local user accounts, see Manage Local Users.

# Default local groups

Updated: May 24, 2012

Applies To: Windows 7, Windows 8, Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Vista, Windows XP

**Default local groups**

The Groups folder located in the Local Users and Groups Microsoft Management Console (MMC) displays the default local groups as well as the local groups that you create. The default local groups are automatically created when you install a stand-alone server or a member server running Windows Server 2003. Belonging to a local group gives a user the rights and abilities to perform various tasks on the local computer. For more information about domain-based groups, see Default groups.

You can add local user accounts, domain user accounts, computer accounts, and group accounts to local groups. However, you cannot add local user accounts and local group accounts to domain group accounts. For more information about adding members to local groups, see Add a member to a local group.

**Note**

- To learn what group you need to be a member of to perform a particular procedure, many procedural topics under **How To** in Help and Support Center provide a note that identifies this information.

The following table provides descriptions of the default groups located in the Groups folder and lists the assigned user rights for each group. These rights are assigned within the local security policy. For complete descriptions of the user rights listed in the table, see User Rights Assignment. For information about editing these rights, see Assign user rights for your local computer.

| Group | Description | Default user rights |
|---|---|---|
| Administrators | Members of this group have full control of the server and can assign user rights and access control permissions to users as necessary. The Administrator account is also a default member. When this server is joined to a domain, the Domain Admins group is automatically added to this group. Because this group has full control of the server, add users with caution. For more information, see Default local groups and Default groups. | Access this computer from the network; Adjust memory quotas for a process; Allow log on locally; Allow log on through Terminal Services; Back up files and directories; Bypass traverse checking; Change the system time; Create a pagefile; Debug programs; Force shutdown from a remote system; Increase scheduling priority; Load and unload device drivers; Manage auditing and security log; Modify firmware environment variables; Perform volume maintenance tasks; Profile single process; Profile system performance; Remove computer from docking station; Restore files and directories; Shut down the system; Take ownership of files or other objects. |
| Backup Operators | Members of this group can back up and restore files on the server, regardless of any permissions that protect those files. This is because the right to perform a backup takes precedence over all file permissions. They cannot change security settings. | Access this computer from the network; Allow log on locally; Back up files and directories; Bypass traverse checking; Restore files and directories; Shut down the system. |
| DHCP Administrators (installed with the DHCP Server service) | Members of this group have administrative access to the Dynamic Host Configuration Protocol (DHCP) Server service. This group provides a way to assign limited administrative access to the DHCP server only, while not providing full access to the server. Members of this group can administer DHCP on a server using the DHCP console or the Netsh command, but are not able to perform other administrative actions on the server. | No default user rights. |

| | | |
|---|---|---|
| DHCP Users (installed with the DHCP Server service) | Members of this group have read-only access to the DHCP Server service. This allows members to view information and properties stored at a specified DHCP server. This information is useful to support staff when they need to obtain DHCP status reports. | No default user rights. |
| Guests | In a computer joined to the domain, members of this group have a temporary profile created at log on, and when the member logs off, the profile is deleted. Profiles in workgroup environments are not deleted. The Guest account (which is disabled by default) is also a default member of this group.<br><br>Members of this group will have a temporary profile created at log on, and when the member logs off, the profile will be deleted. The Guest account (which is disabled by default) is also a default member of this group. | No default user rights. |
| HelpServicesGroup | This group allows administrators to set rights common to all support applications. By default, the only group member is the account associated with Microsoft support applications, such as Remote Assistance. Do not add users to this group. | No default user rights. |
| Network Configuration Operators | Members of this group can make changes to TCP/IP settings and renew and release TCP/IP addresses. This group has no default members. | No default user rights. |
| Performance Monitor Users | Members of this group can monitor performance counters on the server locally and from remote clients without being a member of the Administrators or Performance Log Users groups. | No default user rights. |
| Performance Log Users | Members of this group can manage performance counters, logs and alerts on the server locally and from remote clients without being a member of the Administrators group. | No default user rights. |
| Power Users | Members of this group can create user accounts and then modify and delete the accounts they have created. They can create local groups and then add or remove users from the local groups they have created. They can also add or remove users from the Power Users, Users, and Guests groups. Members can create shared resources and administer the shared resources they have created. They cannot take ownership of files, back up or restore directories, load or unload device drivers, or manage security and auditing logs. | Access this computer from the network; Allow log on locally; Bypass traverse checking; Change the system time; Profile single process; Remove computer from docking station; Shut down the system. |
| Print Operators | Members of this group can manage printers and print queues. | No default user rights. |

| | | |
|---|---|---|
| Remote Desktop Users | Members of this group can remotely log on to a server.<br><br>For more information, see Enabling users to connect remotely to the server. | Allow log on through Terminal Services. |
| Replicator | The Replicator group supports replication functions. The only member of the Replicator group should be a domain user account used to log on the Replicator services of a domain controller. Do not add user accounts of actual users to this group. | No default user rights. |
| Terminal Server Users | This group contains any users who are currently logged on to the system using Terminal Server. Any program that a user can run with Windows NT 4.0 will run for a member of the Terminal Server User group. The default permissions assigned to this group enable its members to run most earlier programs. | No default user rights |
| Users | Members of this group can perform common tasks, such as running applications, using local and network printers, and locking the server. Users cannot share directories or create local printers. By default, the Domain Users, Authenticated Users, and Interactive groups are members of this group. Therefore, any user account created in the domain becomes a member of this group. | Access this computer from the network; Allow log on locally; Bypass traverse checking. |
| WINS Users (installed with WINS service) | Members of this group are permitted read-only access to Windows Internet Name Service (WINS). This allows members to view information and properties stored at a specified WINS server. This information is useful to support staff when they need to obtain WINS status reports. | No default user rights. |

For more information about the most common default groups, see Default security settings for groups.

# Administering Local Users and Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Local Users and Groups**
This section covers:

- Managing files for local user accounts
- Managing local groups from the command line
- Why you should not run your computer as an administrator

# Managing files for local user accounts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing files for local user accounts**
Administrators can use home folders and My Documents to collect user files in one location. User files located in one place simplifies the backup process and makes access control management easier to maintain.

Home folders can be a local folder or a folder located on a shared resource and can be assigned to a single user or many users. When a home folder is assigned to a user, it becomes the user's default folder for the **Open** and **Save As** dialog boxes, for command prompt sessions, and for all programs without a defined working folder.

The My Documents folder is an alternative for home folders, but it does not replace them. When a user tries to save or open a file, most programs determine whether to use the home folder or My Documents in one of two ways:

- Some programs first look in the home folder for files that match the type of file that is to be opened or saved (for example, *.doc or *.txt). If a file with that extension is found, the program opens the home folder and ignores My Documents. If a file of that type is not found, the program opens My Documents.

- In other programs, the home folder is ignored, regardless of whether the home folder contains any files.

For more information about how to assign home folders, see Assign a home folder to a local user account.

When upgrading servers from Windows NT 4.0 or earlier, programs that have stored documents in the home folder will continue to open and save documents in the home folder. However, if a program is installed after you have upgraded, the program uses My Documents to open and save files.

# Managing local groups from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing local groups from the command line**

The following command-line tool can be used to manage local groups.

| Name | Description |
| --- | --- |
| Net localgroup | This command adds, displays, or modifies local groups. |

Additional command-line tools are available in the Support tools folder of the product compact disc and from the Resource Kits.

For information about other command-line tools, see Command-line reference A-Z.

# Why you should not run your computer as an administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Why you should not run your computer as an administrator**

Running your computer as a member of the Administrators group makes the system vulnerable to Trojan horses and other security risks. The simple act of visiting an Internet site or opening an e-mail attachment can be damaging to the system. An unfamiliar Internet site or e-mail attachment may have Trojan horse code that can be downloaded to the system and executed.

If you are logged on as an administrator of a local computer, a Trojan horse could reformat your hard drive, delete your files, and create a new user account with administrative access. If you are logged on as a member of the Domain Admins group, Enterprise Admins group, or Schema Admins group in Active Directory, a trojan horse could create a new domain user account with administrative access and put schema, configuration, or domain data at risk.

On a local computer, it is recommended that you add your domain user account *only* to the Users group (and not to the Administrators group) to perform routine tasks, including running programs and visiting Internet sites. When it becomes necessary to perform administrative tasks on the local computer or in Active Directory, use Run as to start a program using administrative credentials.

Run as allows you to accomplish administrative tasks without exposing your computer or data stored in Active Directory to unnecessary risk. For more information, see Using Run as. For more information about how to use Run as, see Run a program with administrative credentials.

If you need to perform administrative tasks, such as upgrading the operating system or configuring system parameters, then log off and log back on as an administrator.

# Troubleshooting Local Users and Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Received the "The local policy of this system does not permit you to logon interactively" error message, or unable to log on locally.
- Runas command failed.
- Cannot log on to a network computer running Windows 95 or Windows 98.
- Cannot access the Computer Management extension snap-ins on a remote computer.

**Received the "The local policy of this system does not permit you to logon interactively" error message, or unable to log on locally.**
**Cause:** The ability to log on locally to a stand-alone server or member server running Windows Server 2003 is controlled by local security policy.

**Solution:** Add the user account to the Users local group or use the local security policy to assign the particular user or group the right to Allow log on locally.

**See also:** Allow log on locally;Assign user rights for your local computer; Local Security Policy overview

**Runas command failed.**
**Cause:** The Secondary Logon service is not running.

**Solution:** Start the Secondary Logon service.

**See also:** Start, stop, pause, resume, or restart a service

**Cause:** The Secondary Logon service is running; however, you do not have permissions on the shared resource that is running the Secondary Logon service.

**Solution:** Verify that your local user account in Computer Management has the permissions needed to access the shared resource on which the Secondary Logon service is installed. To view permissions in Computer Management, in Control Panel, double-click **Administrative Tools**, and then double-click **Computer Management**.

**See also:** Runas

**Cause:** The Secondary Logon service is running; however, you cannot start or install a program from a shared resource using the **runas** command.

**Solution:** Using the **runas** command, map a network drive to the shared resource using administrative credentials from your local computer, and then start the program again. To do this, open a command prompt by typing the following command in the **Run** dialog box:

**runas /user:administrator cmd**

When prompted, type the administrative password and when the command prompt appears, type:

**net use z: ** *Server* **** *Share* **/user:** *DomainName* **** *UserName* *

This will connect you to the shared resource while still under the security context of the local administrator account.

**See also:** Runas

**Cannot log on to a network computer running Windows 95 or Windows 98.**
**Cause:** Your password is longer than 14 characters.

**Solution:** Create a new user account or change your password to the 14 character maximum for Windows 95 and Windows 98.

**See also:** Reset the password for a local user account

**Cannot access the Computer Management extension snap-ins on a remote computer.**
**Cause:** The Remote Registry service is not running on the remote computer.

**Solution:** Ensure that the Remote Registry service is started on the remote computer. You need appropriate permissions on the remote computer to start the service.

**See also:** Start, stop, pause, resume, or restart a service

**Note**

- You might also receive this error message if the remote computer is running Windows 95. Computer Management does not support remote access to computers that are running Windows 95.

# The Registry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The registry**
The Microsoft® Windows Server 2003 family registry is a database that contains system configuration information for the Windows Server 2003 family operating system.

- For tips about using the registry and Registry Editor, see Registry best practices.
- For help with specific tasks, see Registry How To....
- For general background information, see Registry concepts.

# Registry best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

In the Windows Server 2003 family, system configuration information is centrally located in the registry. While this simplifies the administration of a computer or network, one incorrect edit to the registry can disable the operating system. The following list provides some best practices for using the registry and Registry Editor safely:

- **Before making changes to the registry, make a backup copy.**

  You can back up the registry by using a program such as Backup. After you make changes to the registry, update your Automated System Recovery (ASR) disk. For troubleshooting purposes, keep a list of the changes you make to the registry. For more information, see the "Windows Server 2003 family Registry Reference" at the Microsoft Windows Resource Kits Web site.

- **Do not replace the** Windows Server 2003 **family registry with the registry of another version of the Windows or** Windows NT **operating systems.**

- **Edit the registry carefully.**

  Incorrectly editing the registry may severely damage your system.

- **Limit the number of people who have access to the registry.**

  For example, because members of the Administrators group have full access to the registry, add only users who need such access to the Administrators group. Alternately, you can use Registry Editor to set permissions for specific keys and subtrees, or simply remove Registry Editor from the computers of users whom you do not want to alter the registry.

- **Never leave Registry Editor running unattended.**

- **Do not run Registry Editor as Administrator, except when necessary.**

  A malicious user who has access to Registry Editor running as Administrator can cause severe damage to the operating system and software. Only run Registry Editor as Administrator if you need to see or change keys that you cannot otherwise access.

- **Use Run as to run Registry Editor as Administrator or a different user.**

  To maintain registry security when making changes to a key that requires administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the Regedit icon, clicking Run as, and choosing an account in the local Administrators group. The Regedit icon does not appear by default from the Start menu. To access the icon, open the Windows or WINNT folder on your computer.

- **Do not access HKEY_CURRENT_USER or HKEY_CLASSES_ROOT using services that impersonate more than one user.**

  HKEY_CURRENT_USER and HKEY_CLASSES_ROOT are aliases for other keys in the registry. For example, HKEY_CURRENT_USER is an alias for the key in HKEY_USERS for the user who is currently logged in. If a service accesses one of these keys while running as one user, and again while running as another user, the operating system might still have the alias data cached from the first user. If this happens, an error will result because the second user will not have access rights for keys belonging to the first user.

# Registry How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Open Registry Editor
- Change Keys and Values
- Import or Export Registry Keys
- Load or Unload Registry Hives
- Maintain Registry Security
- Print all or part of the registry

# Open Registry Editor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open Registry Editor**

- Click **Start**, click **Run**, type **regedit**, and then click **OK**.

## Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. Whenever possible, use tools and programs other than Registry Editor to edit the registry.

## Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- Some previous versions of Windows had two versions of Registry Editor: regedit.exe and regedt32.exe. The version of Registry Editor that ships with the Windows Server 2003 family replaces both versions from previous operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Registry best practices
Load or Unload Registry Hives

# Change Keys and Values

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Change keys and values**


- Find a string, value, or key
- Add a registry key to Favorites
- Add a registry key
- Add a value to a registry key entry
- Change a value
- Delete a registry key or value
- Rename a registry key or value
- Connect to a registry over the network
- Disconnect from a network registry
- Copy a registry key name
- Restore the registry

# Find a string, value, or key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a string, value, or key**

1. Open Registry Editor.
2. On the **Edit** menu, click **Find**.
3. In **Find what**, type the string, value, or key you want to find.
4. Select the **Keys**, **Values**, **Data**, and **Match whole string only** check boxes to match the type of search you want, and then click **Find Next**.

## Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

## Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- You can speed up the search by looking through only one type of information; for example, if you know that you are looking for a entry, you can clear the **Keys** and **Data** check boxes, so that you will not search through those objects.
- To repeat the search, press F3.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore the registry
Add a registry key to Favorites
Add a registry key
Copy a registry key name
Add a value to a registry key entry
Change a value

# Add a registry key to Favorites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a registry key to Favorites**

1. Open Registry Editor.
2. Select the registry key you want to add to **Favorites**.
3. On the **Favorites** menu, click **Add to Favorites**.
4. In the **Add to Favorites** dialog box, accept the default registry key name or type a new one.

   The registry key is added to the **Favorites** list. You can then return to this list by simply selecting it from the **Favorites** menu.

## Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

## Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- You can create a list of favorite/frequently visited registry keys.
- To remove a registry key from the **Favorites** list, on the **Favorites** menu, click **Remove Favorite**, and you can select one or more registry keys to remove from the **Favorites** list.
- Renaming a favorite does not rename the corresponding registry key.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore the registry
Add a registry key
Find a string, value, or key
Change a value

# Add a registry key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a registry key**

1. Open Registry Editor.
2. In the registry tree (on the left), click the registry key under which you would like to add a new key.
3. On the **Edit** menu, point to **New**, and then click **Key**.
4. Type a name for the new key, and then press ENTER.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- If you make a mistake that results in your computer not starting properly, you can restore the registry. For instructions, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore the registry
Add a registry key to Favorites
Copy a registry key name
Change a value

# Add a value to a registry key entry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a value to a registry key entry**

1. Open Registry Editor.
2. Click the key or entry where you want to add the new value.
3. On the **Edit** menu, point to **New**, and then click the type of value you want to add: **String Value**, **Binary Value**, **DWORD Value**, **Multi-String Value**, or **Expandable String Value**.
4. Type a name for the new value, and then press ENTER.

## Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

## Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- If you make a mistake that results in your computer not starting properly, you can restore the registry. For instructions on how to restore the registry, see Related Topics.
- When you add a value, the value name and data type are saved with a default value. For instructions on how to modify the default value, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore the registry
Change a value

# Change a value

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a value**

1. Open Registry Editor.
2. Select the entry you want to change.
3. On the **Edit** menu, click **Modify**.
4. In **Value data**, type the new data for the value, and then click **OK**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- If you make a mistake that results in your computer not starting properly, you can restore the registry. For instructions, see Related Topics.
- To make changes to a registry key, you must have the appropriate permissions. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore the registry
Add a value to a registry key entry
Maintain Registry Security

# Delete a registry key or value

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a registry key or value**

1. Open Registry Editor.
2. Click the key or entry you want to delete.
3. On the **Edit** menu, click **Delete**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- You can delete keys and values from your registry. However, you cannot delete a predefined key (such as HKEY_CURRENT_USER) or change the name of a predefined key.
- If you make a mistake that results in your computer not starting properly, you can restore the registry. For instructions, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a registry key
Add a value to a registry key entry
Restore the registry

# Rename a registry key or value

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a registry key or value**

1. Open Registry Editor.
2. Click the key or entry you want to rename.
3. On the **Edit** menu, click **Rename**.
4. Type the new name, and then press ENTER.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- If you make a mistake that results in your computer not starting properly, you can restore the registry. For information, see Related Topics.
- You cannot rename the root keys or the default value of a key.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore the registry
Add a registry key to Favorites
Copy a registry key name
Change a value

# Connect to a registry over the network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a registry over a network**

1. Open Registry Editor.
2. On the **File** menu, click **Connect Network Registry**.
3. In the **Select Computer** dialog box, type the name of the computer to whose registry you want to connect.

## Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

## Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must be logged on as an administrator or a member of the Administrators group on both your computer and on the remote computer before you can make any changes to the remote computer's registry. Network policy settings may also prevent you from completing this procedure.
- If you make a mistake that results in your computer not starting properly, you can restore the registry. For instructions on how to restore the registry, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore the registry
Disconnect from a network registry

# Disconnect from a network registry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect from a network registry**

1. Open Registry Editor.
2. On the **File** menu, click **Disconnect Network Registry**.
3. In the **Disconnect Network Registry** dialog box, click the name of the computer from whose registry you want to disconnect.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- If you make a mistake that results in your computer not starting properly, you can restore the registry. For instructions on how to restore the registry, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore the registry
Connect to a registry over the network

# Copy a registry key name

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy a registry key name**

1. Open Registry Editor.
2. In the registry tree (on the left), click a registry key.
3. On the **Edit** menu, click **Copy Key Name**.
4. Paste the name of the registry key into another program or document.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- If you make a mistake that results in your computer not starting properly, you can restore the registry. For instructions on how to restore the registry, see Related Topics.
- For information on how to find a registry key, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore the registry
Find a string, value, or key
Add a registry key to Favorites
Copy a registry key name
Rename a registry key or value
Change a value

# Restore the registry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore the registry**

1. Click **Options**, and then click **Print** to print these instructions. (If you are using the Help and Support Center, click **Print** above the topic area.) They will not be available after you shut down your computer in step 3.

2. Open Registry Editor.

3. Click **Start**, and then click **Shut Down**.

4. In the list, click **Restart**, and then click **OK**.

5. When you see the message **Please select the operating system to start**, press F8.

6. Use the arrow keys to highlight **Last Known Good Configuration**, and then press ENTER.

   NUM LOCK must be off before the arrow keys on the numeric keypad will function.

7. Use the arrow keys to highlight an operating system, and then press ENTER.

## Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

## Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.

- Choosing **Last Known Good Configuration** provides a way to recover from problems such as a newly added driver that may be incorrect for your hardware. It does not solve problems caused by corrupted or missing drivers or files.

- When you choose **Last Known Good Configuration**, Windows restores information in registry key HKLM\System\CurrentControlSet only. Any changes you have made in other registry keys remain.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Startup options
Start the computer using the last known good configuration

# Import or Export Registry Keys

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Import or export registry keys**

- Export all or part of the registry to a text file
- Import some or all of the registry
- Export a registry key to a hive file
- Import a registry key from a hive file

# Export all or part of the registry to a text file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export all or part of the registry to a text file**

1. Open Registry Editor.
2. On the **File** menu, click **Export**.
3. In **File name**, enter a name for the registry file.
4. Under **Export range**, do one of the following:
   - To back up the entire registry, click **All**.
   - To back up only a particular branch of the registry tree, click **Selected branch** and enter the name of the branch you want to export.
5. Click **Save**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- You can use a text editor like Notepad to work with the registry files you create by exporting.
- You can save registry files in the Windows format, in the format used in Windows 95, Windows 98, and Windows NT 4.0, as binary hive files, or as text files. Registry files are saved with .reg extensions, and text files are saved with .txt extensions.
- In Windows Explorer, double-clicking a file with the .reg extension imports the file into the computer's registry.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Import some or all of the registry
Export a registry key to a hive file

# Import some or all of the registry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import some or all of the registry**

1. Open Registry Editor.
2. On the **File** menu, click **Import Registry File**.
3. Find the file you want to import, click the file to select it, then click **Open**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Note

- In Windows Explorer, double-clicking a file with the .reg extension imports the file into the computer's registry.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Export all or part of the registry to a text file
Import a registry key from a hive file

# Export a registry key to a hive file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export a registry key to a hive file**

1. Open Registry Editor.
2. Select the key that you want to save as a file.
3. On the **File** menu, click **Export**.
4. In the **Export Registry File** dialog box, in **Save in**, click the drive, folder, or network computer and folder where you want to save the hive.
5. In **File name**, enter a name for the hive.
6. In **Save as type**, click **Registry Hive Files**.
7. Click **Save**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- Registry Editor provides a number of commands that are designed primarily for the maintenance of your system. For example, **Load Hive** and **Unload Hive** allow a part of your system to be temporarily downloaded onto another computer for maintenance. Before a hive can be loaded or restored, it must be saved as a key, either to a floppy disk or to your hard disk.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Import a registry key from a hive file
Load a hive into the registry
Unload a hive from the registry

# Import a registry key from a hive file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import a registry key from a hive file**

1. Open Registry Editor.
2. Select the keys in which you want to restore the hive.
3. On the **File** menu, click **Import**.
4. In **Look in**, select the drive, folder, or network computer and folder in which the hive is located.
5. In **Files of type**, click **Registry Hive Files**.
6. Select the correct file name for the hive.
7. Click **Open**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- A restored hive overwrites an existing registry key and becomes a permanent part of your configuration. For example, to perform maintenance on part of your system, you can use **Export** to save a hive to a disk. When you are ready, you can then use **Import** on the **File** menu to restore the saved key to your system.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Export a registry key to a hive file
Load a hive into the registry
Unload a hive from the registry

# Load or Unload Registry Hives

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Load or unload registry hives**

- Load a hive into the registry
- Unload a hive from the registry

# Load a hive into the registry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To load a hive into the registry**

1. Open Registry Editor.
2. In the registry tree (on the left), click either the HKEY_USERS or HKEY_LOCAL_MACHINE keys.
3. On the **File** menu, click **Load Hive**.
4. In **Look in**, click the drive, folder, or network computer and folder that contains the hive you want to load.
5. Click **Open**.
6. In **Key Name**, type the name that you want to assign to the hive, and then click **OK**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- **Load Hive** and **Unload Hive** affect only the HKEY_USERS and HKEY_LOCAL_MACHINE keys and are active only when these predefined keys are selected. When you load a hive into the registry, the hive becomes a subkey of one of these keys.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Unload a hive from the registry
Import a registry key from a hive file
Export a registry key to a hive file

# Unload a hive from the registry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To unload a hive from the registry**

1. Open Registry Editor.
2. Select a hive that you have previously loaded onto your system.
3. On the **File** menu, click **Unload Hive**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- **Load Hive** and **Unload Hive** affect only the HKEY_USERS and HKEY_LOCAL_MACHINE keys and are active only when these predefined keys are selected. When you load a hive into the registry, the hive becomes a subkey of one of these keys.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Load a hive into the registry
Import a registry key from a hive file
Export a registry key to a hive file

# Maintain Registry Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maintain registry security**

- Assign permissions to a registry key
- Assign special access to a registry key
- Add users or groups to the Permissions list
- Remove a user or group from the Permissions list
- Grant full control of a registry key
- Audit activity on a registry key
- Add users or groups to the Audit list
- Remove a user or group from the Audit list
- Take ownership of a registry key

# Assign permissions to a registry key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign permissions to a registry key**

1. Open Registry Editor.
2. Click the key to which you want to assign permissions.
3. On the **Edit** menu, click **Permissions**.
4. Assign an access level to the selected key as follows:
   - To grant the user permission to read the key contents, but not save any changes made to the file, under Permissions for *name*, for **Read**, select the **Allow** check box.
   - To grant the user permission to open, edit, and take ownership of the selected key, under Permissions for *name*, for **Full Control**, select the **Allow** check box.
   - To grant the user special permission in the selected key, click **Advanced**.
5. If you are assigning permissions to a subkey and you want the inheritable permissions assigned to the parent key to apply to the subkey also, click **Advanced** and select the **Inherit from parents the permission entries that apply to child objects. Include these with entries explicitly defined here** check box.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- If you own a registry key, you can specify the users and groups that can open that key. To determine who can open your registry keys, you need to assign permissions to them. You can add or remove users or groups from those authorized to access your registry keys at any time.
- The **Special Permissions** check boxes indicate whether custom permissions have been set for this key, but you cannot set special permissions by clicking these check boxes. Click **Advanced** to set special permissions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Audit activity on a registry key
Grant full control of a registry key
Add users or groups to the Permissions list
Remove a user or group from the Permissions list
Assign special access to a registry key

# Assign special access to a registry key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign special access to a registry key**

1. Open Registry Editor.
2. Click the key to which you want to assign special access.
3. On the **Edit** menu, click **Permissions**.
4. Click **Advanced**, and then double-click the user or group to whom you want to assign special access.
5. Under **Permissions**, select the **Allow** or **Deny** check box for each permission you want to allow or deny.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- Assigning special access is useful for situations in which you need to assign permissions that are not defined by either Read or Full Control.
- If you are assigning permissions to a subkey and you want the inheritable permissions assigned to the parent key to apply to the subkey also, select the **Inherit from parents the permission entries that apply to child objects. Include these with entries explicitly defined here** check box.
- If you are assigning permissions to a parent key and want any existing permissions of its subkeys replaced with the parent key's inheritable permissions, select the **Replace permission entries on all child objects with entries shown here that apply to child objects** check box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign permissions to a registry key
Add users or groups to the Permissions list

# Add users or groups to the Permissions list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add users or groups to the Permissions list**

1. Open Registry Editor.
2. Click the key whose Permissions list you want to change.
3. On the **Edit** menu, click **Permissions**, and then click **Add**.
4. In the **Select Users, Computers, or Groups** dialog box, click **Locations**, and then click the computer or domain of the users and groups you want to view.
5. Type the name or names of the users or groups you would like to add, separating each name with a semicolon. Click **Examples** to see some examples of the syntax you can use to enter names. Click **Check Names** to validate names with the directory. When you are finished entering names, click **OK**.
6. In the **Permissions** dialog box, under **Permissions** for *name*, assign a type of access to the selected user or group as follows:
   - To grant the user permission to read the key contents but not to save any changes made to it, select the **Allow** check box for **Read**.
   - To grant the user permission to open, edit, and take ownership of the selected key, select the **Allow** check box for **Full Control**.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

**Notes**

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- If the check boxes under **Permissions** are shaded, the key has inherited permissions from the parent object key.
- To allow permissions assigned to a parent key to apply to its subkeys also, click **Advanced**, and select the **Inherit from parent the permission entries that apply to child objects. Include these with entries explicitly defined here** check box.
- In the **Select Users, Computers, or Groups** dialog box, if you type the name, rather than selecting it, click **Check Names** before clicking **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign permissions to a registry key
Remove a user or group from the Permissions list
Assign special access to a registry key

# Remove a user or group from the Permissions list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a user or group from the Permissions list**

1. Open Registry Editor.
2. Click the key whose Permissions list you want to change.
3. On the **Edit** menu, click **Permissions**.
4. Under **Group or user names:** , Click the name of the user or group that you want to remove from the Permissions list.
5. Click **Remove**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add users or groups to the Permissions list
Assign permissions to a registry key
Assign special access to a registry key

# Grant full control of a registry key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To grant full control of a registry key**

1. Open Registry Editor.
2. Click the key to which you want to grant full control.
3. On the **Edit** menu, click **Permissions**.
4. Under **Group or user names**, click the user to whom you want to grant full control of your registry key.
5. Under **Permissions for** *name*, where *name* represents the name of the user to whom you are granting full control of the key, select the **Allow** check box for **Full Control**.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

**Notes**

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- You can permit another user to take ownership of a registry key only if you are the current owner of the key.
- To permit a user to to take ownership of a registry key, you must first assign the user Full Control of the key.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign permissions to a registry key
Take ownership of a registry key

# Audit activity on a registry key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To audit activity on a registry key**

1. Open Registry Editor.
2. Click the key you want to audit.
3. On the **Edit** menu, click **Permissions**.
4. Click **Advanced**, and then click the **Auditing** tab.
5. Double-click the name of a group or user.
6. Under **Access**, select or clear the **Successful** and **Failed** check boxes for the activities that you want to audit or to stop auditing:

| Select | To audit |
|---|---|
| **Query Value** | Any attempts to read a entry from a registry key |
| **Set Value** | Any attempts to set entries in a registry key |
| **Create Subkey** | Any attempts to create subkeys on a selected registry key |
| **Enumerate Subkeys** | Any attempts to identify the subkeys of a registry key |
| **Notify** | Any notification events from a key in the registry |
| **CreateLink** | Any attempts to create a symbolic link in a particular key |
| **Delete** | Any attempts to delete a registry object |
| **Write DAC** | Any attempts to write a discretionary access control list on the key |
| **Write Owner** | Any attempts to change the owner of the selected key |
| **Read Control** | Any attempts to open the discretionary access control list on a key |

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

**Notes**

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- If your computer is connected to a network, network policy settings might prevent you from completing this procedure.

- You must first add users and groups before specifying the events to audit.
- Auditing activity can slow the computer down significantly. Consider auditing only failures, and not successes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add users or groups to the Audit list
Remove a user or group from the Audit list

# Add users or groups to the Audit list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add users or groups to the Audit list**

1. Open Registry Editor.
2. Click the key you want to audit.
3. On the **Edit** menu, click **Permissions**.
4. Click **Advanced**, click the **Auditing** tab, and then click **Add**.
5. Click **Object Types**, select the type or types of users or groups you want to find, and then click **OK**.
6. Click **Locations**, select the computer or domain of the users or groups you want to view, and then click **OK**.
7. Type the name of the user or group you want to add and then click **OK** to open the **Auditing Entry** dialog box, or click **Advanced** to search for a user, computer, or group based on parameters you set.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- You must first add a user or group before specifying which events to audit.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Remove a user or group from the Audit list
Audit activity on a registry key

# Remove a user or group from the Audit list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a user or group from the audit list**

1. Open Registry Editor.
2. Click the key whose Audit list you want to change.
3. On the **Edit** menu, click **Permissions**.
4. Click **Advanced**, and then click the **Auditing** tab.
5. Click the user or group that you want to remove, and then click **Remove**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add users or groups to the Audit list
Audit activity on a registry key

# Take ownership of a registry key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To take ownership of a registry key**

1. Open Registry Editor.
2. Click the key you want to take ownership of.
3. On the **Edit** menu, click **Permissions**.
4. Click **Advanced**, and then click the **Owner** tab.
5. Under **Change owner to**, click the new owner, and then click **OK**.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- You can take ownership of a registry key if you are logged on as an administrator or if you have been specifically assigned the permission to take ownership of the registry key by the current owner.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign permissions to a registry key
Audit activity on a registry key

# Print all or part of the registry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print all or part of the registry**

1. Open Registry Editor.
2. Click the computer or top-level key of the registry area you want to print.
3. On the **File** menu, click **Print**.
4. You can print the entire registry by clicking **All** or only part of the registry by clicking **Selected Branch** and typing a particular branch in the text box, and then click **OK**.

## Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

## Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- The average registry can take hundreds of pages to print. Consider printing only those branches you need.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Registry Editor
**Other Resources**
Registry Editor overview

# Registry concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Registry overview
- Registry structure
- Securing the registry

# Registry overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Registry overview**

In the Microsoft® Windows Server 2003 family, system configuration information is stored centrally in a hierarchical database called the registry. The registry replaces many of the .ini, .sys, and .com configuration files used in Windows for MS-DOS® and Microsoft LAN Manager.

For more information about managing the registry, see "Windows Server 2003 Family Registry Reference" at the Microsoft Windows Resource Kits Web site. For information about Windows Server 2003 support, see Support resources.

# Registry structure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Registry structure**

The registry is organized in a hierarchical structure of subtrees and their keys, subkeys, and entries.

The contents of the registry for one computer may vary widely from that of another, depending on the devices, services, and programs installed on each computer.

Keys can have subkeys and subkeys can, in turn, have subkeys. While most information in the registry is stored on disk and is considered permanent, some information, stored in *volatile keys*, is overwritten each time the operating system starts.

**Registry subtrees**

Windows Server 2003 family operating systems have two registry subtrees: **HKEY_LOCAL_MACHINE** and **HKEY_USERS**. However, to make the information in the registry easier to find, the Registry Editor utilities display five subtrees, three of which are aliases of other parts of the registry. These five subtrees are listed and described in the following table.

| Root key name | Description |
|---|---|
| **HKEY_LOCAL_MACHINE** | Contains information about the local computer system, including hardware and operating system data such as bus type, system memory, device drivers, and startup control data. |
| **HKEY_CLASSES_ROOT** | Contains information used by various OLE technologies and file-class association data. A particular key or value exists in **HKEY_CLASSES_ROOT** if a corresponding key or value exists in either **HKEY_LOCAL_MACHINE\SOFTWARE\Classes** or **HKEY_CURRENT_USER\SOFTWARE\Classes**. If a key or value exists in both places, the **HKEY_CURRENT_USER** version is the one that appears in **HKEY_CLASSES_ROOT**. |
| **HKEY_CURRENT_USER** | Contains the user profile for the user who is currently logged on interactively (as opposed to remotely), including environment variables, desktop settings, network connections, printers, and program preferences. This subtree is an alias of the **HKEY_USERS** subtree and points to **HKEY_USERS***security ID of current user*. |
| **HKEY_USERS** | Contains information about actively loaded user profiles and the default profile. This includes information that also appears in **HKEY_CURRENT_USER**. Users who are accessing a server remotely do not have profiles under this key on the server; their profiles are loaded into the registry of their own computers. |
| **HKEY_CURRENT_CONFIG** | Contains information about the hardware profile used by the local computer system at startup. This information is used to configure settings such as the device drivers to load and the display resolution to use. This subtree is part of the **HKEY_LOCAL_MACHINE** subtree and points to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Hardware Profiles\Current**. |

Each root key name begins with **HKEY_** to indicate to software developers that this is a *handle* that a program can use. A handle is a value used to identify a resource so that a program can access it.

**Note**

- The behavior of **HKEY_CLASSES_ROOT** in Windows 2000 and Windows Server 2003 family operating systems is different from that of the same key in previous versions of Windows. In Windows 2000 and Windows Server 2003 family operating systems, a feature called per-user class registration provides a view of **HKEY_CLASSES_ROOT** that contains both user-specific class registrations and computer-specific class registration. With this feature, different users can have different sets of class registrations (including those for file associations, COM components, and MIME types). Previously, all users shared registrations, so that one user could alter registrations for other users. Per-user class registration isolates each user's registrations and allows the administrator to increase security on **HKEY_LOCAL_MACHINE\Software\Classes**.

**Registry hives and files**

The term hive describes a body of keys, subkeys, and values that is rooted at the top of the registry hierarchy. A hive is backed

by a single file and a .log file that are in the *systemroot*\System32\Config or the *systemdrive*\Documents and Settings*username* folders. The string *systemroot* is a placeholder for the directory name that contains the Windows Server 2003 family system files; the string *systemdrive* is a placeholder for the partition that contains the Windows Server 2003 family system files. For example, if the Windows Server 2003 family system files for a user named Annie are in a directory called Windows in the C: partition, the hive files are stored in the C:\Windows\System32\Config and C:\Documents and Settings\Annie folders. Hives may also be called registry files or registry log files.

By default, most hive files (**DEFAULT**, **SAM**, **SECURITY**, **SOFTWARE**, and **SYSTEM**) are stored in the *systemroot*\System32\Config folder. In the Windows Server 2003 family of operating systems, the location of user profile information for each user of a computer, including the Ntuser.dat and Ntuser.dat.log, may depend on whether the installation of the operating system was a fresh installation or whether it was installed as an upgrade from Windows NT or Windows 2000. In fresh installations, the Ntuser.dat and Ntuser.dat.log files are stored in the *systemdrive*\Documents and Settings*username* folder. In installations that are upgrades from Windows NT or Windows 2000, the Ntuser.dat and Ntuser.dat.log files are stored in the *systemroot*\Profiles*username* folder.

Each hive in the Windows Server 2003 family registry is associated with a set of standard files. The following table shows the standard hives and files for a computer running an operating system in the Windows Server 2003 family.

| Registry hive | File names |
|---|---|
| **HKEY_LOCAL_MACHINE\SAM** | Sam and Sam.log |
| **HKEY_LOCAL_MACHINE\SECURITY** | Security and Security.log |
| **HKEY_LOCAL_MACHINE\SOFTWARE** | Software and Software.log |
| **HKEY_LOCAL_MACHINE\SYSTEM** | System and System.log |
| **HKEY_CURRENT_CONFIG** | System and System.log |
| **HKEY_CURRENT_USER** | Ntuser.dat and Ntuser.dat.log |
| **HKEY_USERS\.DEFAULT** | Default and Default.log |

**Entries in the registry keys**

Each registry key or subkey can contain data called entries. Some entries store information that is specific to each user, while others store information that applies to all users of a computer. An entry has three parts: the name of the value, the data type of the value, and the value itself. The three parts of entries always appear in the following order.

```
        Name                Data type            Value
DependOnService:  REG_MULTI_SZ:  Tcpip Nbtsys Streams
```

Data types describe the format of the data. Data types from 0 through 0x7FFFFFFF are reserved for definition by the system. Programs are encouraged to use these data types, but data types from 0x80000000 through 0xFFFFFFFF are also reserved for use by programs.

The following table lists the data types currently defined and used by the system.

| Data type | Description |
|---|---|
| **REG_BINARY** | Raw binary data. Most hardware component information is stored as binary data and is displayed in Registry Editor in hexadecimal format. |
| **REG_DWORD** | Data represented by a number that is 4 bytes long. Many parameters for device drivers and services are this type and are displayed in Registry Editor in binary, hexadecimal, or decimal format. |
| **REG_EXPAND_SZ** | A variable-length data string. This data type includes variables that are resolved when a program or service uses the data. |
| **REG_MULTI_SZ** | A multiple string. Values that contain lists or multiple values in a form that people can read are usually this type. Entries are separated by spaces, commas, or other marks. |
| **REG_SZ** | A fixed-length text string. |

| | |
|---|---|
| **REG_FULL_RESOURCE_DESCRIPTOR** | A series of nested arrays designed to store a resource list for a hardware component or driver. |
| | |

For more information about managing the registry, see "Windows Server 2003 Family Registry Reference" at the Microsoft Windows Resource Kits Web site.

# Securing the registry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing the registry**

The Registry contains sensitive data about the computer and its applications and files. A malicious user can use the registry to cause severe damage to the computer. It is important that you maintain a high level of registry security.

The registry is given a high level of security by default. Administrators are given full access to the entire registry, while other users generally have full access to the keys related to their own user accounts (which include HKEY_CURRENT_USER) and read-only access to the keys related to the computer and its software. Users have no access to the keys related to other users' accounts. Users with the appropriate permissions for a given key can make changes to the permissions for that key and for any keys it contains.

For help securing the registry, see Maintain Registry Security. For help using permissions, see Access Control.

# Windows Management Instrumentation Command-line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Management Instrumentation Command-line**
Windows Management Instrumentation command-line (WMIC) provides you with a simple command-line interface to Windows Management Instrumentation (WMI), so you can take advantage of WMI to manage computers running the Windows Server 2003 family of operating systems. WMIC interoperates with existing shells and utility commands, and can be easily extended by scripts or other administration-oriented applications.

- For tips about using WMIC, see WMIC Best practices.
- For help with specific tasks, see WMIC How To....
- For general background information , see WMIC Concepts.
- For problem-solving instructions, see Troubleshooting WMIC.

# WMIC Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Use the WHERE () clause for specific instances to be returned for an alias.**

- The WHERE() clause must include the WHERE clause for a valid WQL query. The parentheses are only required for complex criteria. See the WMI SDK for WQL syntax. The WHERE() clause must immediately follows the alias name. For example: `process where (processID>300) list brief`

- When the WHERE clause is WHERE NAME='<value>', it is sufficient to give only the value. For example: `SERVICE CLIPSRV`

  Instead of:

  ```
  SERVICE WHERE NAME='CLIPSRV'
  ```

**Use /OUTPUT, /APPEND, and /RECORD commands to redirect output.**

- Commands that output in HTML format (such as CLASS) typically have their output directed to a file. You can then view that file for results of the command. You can redirect the output can by using the /OUTPUT, /APPEND, and /RECORD commands.

**Enclose values in quotation marks if the values contain special characters such as dashes, slashes, or spaces.**

**Use these methods of query for large domains or databases**

- ALIAS

  ```
  WMIC ALIAS PROCESS LIST BRIEF
  ```

- USERACCOUNT

  ```
  WMIC USERACCOUNT WHERE "Name='PutUserNameHere' and Domain='PutDomainNameHere'"
  ```

- FSDIR

  ```
  WMIC FSDIR WHERE Name='c:\\WINDOWS'
  ```

- DATAFILE

  ```
  WMIC DATAFILE WHERE Name="c:\\boot.ini"
  WMIC DATAFILE WHERE "PATH='\\windows\\' and Extension='exe' and FileSize>'108032'"
  GET LastAccessed, LastModified, Name, FileSize
  ```

- NTEVENT

  ```
  WMIC NTEVENT WHERE "LogFile='system' and Type>'0'" GET Message, TimeGenerated
  WMIC NTEVENT WHERE "LogFile='system' and Type>'0'" GET Message, TimeGenerated
  /FORMAT:htable >c:\MySystemEvents.htm
  ```

- Querying:

Use quotes "" to separate WHERE expressions instead of parentheses () :

```
WMIC NTEVENT WHERE "LogFile='system' and Type>'4'"
```

[ Compare to:

```
WMIC NTEVENT WHERE (LogFile='system' and TYPE>4)
```

In this case, Command Shell will misinterpret

```
">4)"
```

and will try to redirect output to file named

```
"4)"
```

]

For information about validation of users in WMI, see Managing WMI Security.

For information about WMIC security, authentication, and authorization, see Security with WMIC.

For more information, see Windows Management Instrumentation Command-line.

# WMIC How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Run WMIC in interactive mode
- Run WMIC in non-interactive mode
- Create a format in WMIC
- Edit a format in WMIC

# Run WMIC in interactive mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run WMIC in interactive mode**

1. Click **Start**, **Run**.
2. Type:

   **WMIC**, and then press ENTER.
3. Type the alias, command, or global switch you want WMIC to perform.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- When you finish using WMIC in interactive mode, you can use the **Exit** or **Quit** command to exit WMIC.
- For more information about the WMIC commands, type **/?** at the command prompt.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Running Windows Management Instrumentation Command-line
WMIC verbs
Security with WMIC

# Run WMIC in non-interactive mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run WMIC in non-interactive mode**

1. Open Command prompt.
2. At the command prompt, type:

   **WMIC**

   and the command, aliases, or global switches, you want WMIC to perform.
3. Press ENTER.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the command prompt window, click Start, point to All Programs, point to Accessories, and then click Command Prompt.
- In non-interactive mode, WMIC returns you to the command prompt after executing the command.
- For more information about the WMIC commands, type **WMIC /?** at the command prompt.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Running Windows Management Instrumentation Command-line
WMIC verbs
Security with WMIC

# Create translation table in WMIC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a translation table in WMIC**

1. Locate the Managed Object Format (MOF) file, **CliEgAliases.mof**, located in the %WINDIR%\system32\wbem directory.
2. Open the MOF file in Notepad.
3. Create a new translation table as an instance of **MSFT_CliTranslateTable**.
4. Define the **Name** property of the translation table being created.
5. Define the **FromValue** and **ToValue** properties of Instance of **MSFT_CliTranslateTableEntry**.
6. Save the file, using UNICODE encoding.
7. From the command prompt, type:

```
Mofcomp CliEgAliases.mof
```

### Note

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a format in WMIC
Running Windows Management Instrumentation Command-line
WMIC verbs
Security with WMIC

# Create a format in WMIC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a format in WMIC**

1. Create an XSLT-formatted file using the XSL language (for example, *MyNewFormat.xsl*).
2. Copy the file into the **%WINDIR%\system32\wbem** directory.
3. Locate the **xsl-mappings.xml** file in the **%WINDIR%\system32\wbem** directory.
4. Add an entry in the **xsl-mappings.xml** file which will map the format to Keyword:

```
<XSLFORMAT KEYWORD="MyNewKeyword">MyNewFormat.xsl</XSLFORMAT>
```

### Note

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create translation table in WMIC
Running Windows Management Instrumentation Command-line
WMIC verbs
Security with WMIC

# Edit a format in WMIC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit a format in WMIC**

1. Locate an existing XSLT format (an .xsl file) in the **%WINDIR%\system32\wbem** directory.
2. Make the required changes and save.
3. Locate the **xsl-mappings.xml** file in the **%WINDIR%\system32\wbem** directory.
4. Modify the related entry in the xsl-mappings.xml file:

```
<XSLFORMAT KEYWORD="MyModifiedKeyword">MyModifiedFormat.xsl</XSLFORMAT>
```

### Note

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a format in WMIC
Running Windows Management Instrumentation Command-line
WMIC verbs
Security with WMIC

# WMIC Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**
This section provides general background information about the Windows Management Intrumentation Command-line (WMIC).

- Windows Management Instrumentation overview
- Understanding Windows Management Instrumentation Command-line
- Running Windows Management Instrumentation Command-line
- Using Windows Management Instrumentation Command-line

# Windows Management Instrumentation overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Management Instrumentation overview**
Windows Management Instrumentation (WMI) is the Microsoft implementation of Web-Based Enterprise Management Initiative (WBEM), an industry initiative to establish standards for accessing and sharing management information over an enterprise network. WMI provides integrated support for the Common Information Model (CIM)--the data model that describes the objects that exist in a management environment.

WMI includes an object repository, which is the database of object definitions, and the WMI Object Manager, which handles the collection and manipulation of objects in the repository and gathers information from WMI providers. WMI providers act as intermediaries between WMI and components of the operating system, applications, and other systems. For example, the registry provider draws information from the registry, while the SNMP provider provides data and events from SNMP devices. Providers provide information about their components and might provide methods to manipulate the components, properties that can be set or events that can alert you to changes in the components.

WMI can be used by computer management tools, such as Microsoft Systems Management Server, to help you manage your computers. WMI is also used by other Microsoft technologies and tools, such as Microsoft Health Monitor and Microsoft Operations Manager, and by other vendors of computer management systems. You can use WMI with programming or scripting systems (such as Windows Script Host) to retrieve configuration details about most aspects of your computer systems, including server applications, or to affect changes to your systems. For more details, see the WMI SDK.

Several management tools are WMI-enabled, including System Properties, System Information, and the Dependencies component of Services. Brief descriptions of these components follow:

- System Properties lets you view and change system properties on a local or remote computer. You can restart a remote computer to apply settings changes or to detect new hardware, view the computer name and domain information for other computers on your network, or change the settings for the virtual memory paging file on a computer that might run programs requiring a lot of memory. For more information about System Properties, see System Properties.

- System Information collects and displays configuration information about your system. This is is especially useful when a support technician troubleshoots your system. For more information, see System Information.

- Services helps you manage the services on your computer. Services dependencies identify the services on which the current service depends on and the services that depend on it. For more information about Services, see Services Concepts.

For technical information about developing the WMI system, see the Windows Management Instrumentation documentation at the Microsoft Platform Software Development Kit (SDK) Web site.

For information about using the WMI Tester to view and modify WMI objects, see Windows Management Instrumentation Tester overview.

For information about using the command-line interface to WMI, see Windows Management Instrumentation Command-line.

For information about using the WMI Control to configure WMI settings and security on a remote or local computer, see WMI Control Overview.

# Understanding Windows Management Instrumentation Command-line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Windows Management Instrumentation Command-line**
This section covers:

- Windows Management Instrumentation Command-line overview
- WMIC architecture

# Windows Management Instrumentation Command-line overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Management Instrumentation Command-line overview**
The Window Management Instrumentation command-line (WMIC) provides you with a simple command-line interface to Windows Management Instrumentation (WMI), so you can take advantage of WMI to manage computers running Microsoft Windows. WMIC interoperates with existing shells and utility commands, and can be easily extended by scripts or other administration-oriented applications.

WMI Providers enable you to use WMI to manage a wide variety of hardware components, operating system subsystems, and application systems. WMIC can be used with all the schemas implemented by those WMI Providers.

WMIC can be used to remotely manage any computer with WMI that is a Windows domain member. WMIC does not have to be available on the remotely managed computer in order for WMIC to manage it.

You can use WMIC to accomplish tasks in the following typical scenarios:

- Local management of a computer--you are at the computer and use the WMIC command to manage it.
- Remote management of a computer--you are at one computer and use WMIC to manage another computer.
- Remote management of multiple computers--you are at on one computer and use WMIC to manage multiple computers with a single command.
- Remote management of a computer (using a remote session)--you use a remote session technology (such as Telnet or Terminal Services) to connect to a remote computer and manage it with WMIC.
- Automated management using administrative scripting--you use WMIC to write a simple management script (batch file) to automate the management of a computer (local, remote, or multiple computers--serially or simultaneously).

**Notes**

- For general background information on WMI, see Windows Management Instrumentation overview.
- For information about how WMI validates users, see Managing WMI security.

# WMIC architecture

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding WMIC architecture**

This topic discusses the relationship of WMIC to WMI, to other computers, and to scripting languages. It also discusses the relationship between aliases, formats, translation tables, and other WMIC elements.

WMIC has an engine to parse, interpret, and execute alias operations called from the command line. The engine is an executable file, called WMIC.exe, which is run on the computer that the user is connected to. Any extension to the WMI command-line interpreter must be done through the addition of alias file definitions or registrations.

Aliases, or friendly names, are defined in Management Object Format (MOF), a language for defining classes and instances in the CIM, and are related to classes in WMI schema.

The WMIC.exe engine receives a command from its stdin interface. The command is parsed by the tool and if it involves an alias, the alias definition is obtained from the current WMI operational namespace (by default, root\cli). The tool interprets the original command with the alias command and current environmental settings, such as the current target namespace. The command is executed as a series of WMI API calls against the target namespace. Any namespace might be used as a target if the namespace is consistent with the alias being processed.

The information returned for the API calls is retrieved in an XML format. The retrieved data is modified by the XSL style sheet, which is associated with the alias. It is applied to the data primarily for presentation formatting. The resulting text is presented back to the user using the stdout interface. The target system and management system need not be on the same computer. However, the original computer must have sufficient credentials to access the target computer.

The style sheet applied to the output defaults is the standard FULL style sheet if no other is specified.

For more information about using WMIC, see Using Windows Management Instrumentation Command-line.

For information about WMI security, see Managing WMI security.

For information about WMIC security, see Security with WMIC.

For more information, see Windows Management Instrumentation Command-line.

# Security with WMIC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**General information about the security model**
Only users in the local Administrators group can start WMIC. The WMI security access model is enforced on any call from WMIC to WMI. The operating system enforces operating system-level security for any operations executed through WMI. For example, you cannot access files through WMI that you cannot access from Windows Explorer.

WMIC is a client for WMI, so all security checks are performed within WMI components. Before connecting to a remote computer, the component uses **PING** functionality to verify validity of the input (remote computer status).

When WMIC is used in a Telnet, Terminal Services, or similar session, all commands are carried out in the context of the user issuing the command.

WMIC allows extension of Aliases and XSLs. Security is addressed by NTFS which enforces administrator-only credential requirements for accessing WMIC's files.

Except for method parameters, WMIC allows input data only through the **/NODE** switch. A list of computer nodes can be provided as a text file using the **@FileList** option. The file and its location should be secured by the Administrator, and the input file should preferably be read from the WBEM directory.

**User permissions**
No special user permissions are required to run WMIC. Implicitly, to be able to use WMIC, a user needs to have full write access to both the WMIC namespaces and a registry key (HKLM\Software\Microsoft\WBEM).

**Authentication and authorization**
Authentication and authorization are provided by the system when these resources are opened. They are stored in a secured location and the Discretionary Access Control List (DACL) is used to verify that client has access to them. They are wholly dependent on Distributed Component Object Model (DCOM) security as well as Windows Management (WINMGMT) for remote access in providing it with user name, password, and requested levels. The **/AUTHLEVEL** switch is used for securing WMIC network traffic. In the Windows Server 2003 family operating systems, the default level is set to **PKTPRIVACY** providing packet privacy encryption level. You can modify the authentication level using the **/AUTHLEVEL** switch.

**Example**:

```
WMIC /AUTHLEVEL:Pktprivacy /NODE:"testcomputer.microsoft.com" BIOS
```

**Configuration data**
The configuration information is stored in the registry and in the WMI repository. The registry configuration holds timestamps and validity of Managed Object Format(MOF)-compiled status for schemas. At the start of each session of WMIC, this is evaluated and WMIC will update the WMI repository when schemas are not valid or were updated by a user.

**Miscellaneous**
WMIC is a powerful command-line tool that allows you to quickly do any task for which you have permissions. If you are not familiar with the WMI environment, you can use the **/INTERACTIVE** switch to set a warning for delete operations. For example:

```
WMI /INTERACTIVE:ON
```

For information about how WMI validates users, see Managing WMI security.

For more information, see Windows Management Instrumentation Command-line.

# Using Windows Management Instrumentation Command-line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Windows Management Instrumentation Command-line**

This section covers:

- Running Windows Management Instrumentation Command-line
- Using WMIC help
- WMIC aliases
- WMIC verbs
- WMIC switches
- WMIC commands
- Security with WMIC
- Creating and Customizing WMIC Components

# Running Windows Management Instrumentation Command-line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Running Windows Management Instrumentation Command-line**
When you run Windows Management Instrumentation Command-line (WMIC) for the first time, it compiles Alias definitions, or Managed Object Format (MOF) files. Whenever MOF files are changed, WMIC recompiles the files and updates the repository.

You can choose to use WMIC in interactive mode or non-interactive mode. The following gives you some information about the characteristics and functions of the two modes:

**Interactive mode**
WMIC in interactive mode returns you to the WMIC command prompt after executing a command. **Root\cli** is the default WMIC role, and is there when you starts WMIC in interactive mode. The **Exit** or **Quit** command is used to exit WMIC.

Interactive mode is useful if you are entering a series of WMIC commands.

To start WMIC in interactive mode see Run WMIC in interactive mode.

The following examples include aliases, commands, and global switches that you can enter at the WMIC prompt when in interactive WMIC mode:

- **CONTEXT** - lists all environment variables (Global switches).
- **OS** - displays data about operating system in a default TABLE format.
- **OS list full** - displays **OS** alias data in **LIST** format.
- **/?** - lists descriptions of available environment variables and alias.
    - **OS /?** - lists options available for operating systems.
    - **OS list /?** - specifies formats for "list" operations.
    - **OS list full /?** - lists switches available.
    - **OS list full /format /?** - lists keywords for available formats.
    - **OS call /?** - lists methods (verbs) available.
    - **OS get /?** - lists properties available.

**Non-Interactive mode**
WMIC in non-interactive mode exits WMIC after performing the command you entered. You are returned to the WMIC command prompt when the command is completed. WMIC aliases, global switches, or commands can also be used in non-interactive mode.

Non-interactive mode is useful if you use WMIC for a batch procedure, or if you only need to execute one WMIC command.

To start WMIC in non-interactive mode see Run WMIC in non-interactive mode.

The following examples show how you can run commands in non-interactive mode:

- <PROMPT>`wmic os get /format:hform>MyOS.htm`

The output is directed to a html file, and you are returned to the command prompt where you started.

- <PROMPT>`wmic /?`

  The Help is listed, and you are returned to the command prompt where you started

For more information, see Windows Management Instrumentation Command-line.

# Using WMIC help

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using WMIC help**
WMIC includes on-line Help at the command line. At any level, you can type **/?** for additional details. By itself, **/?** gives you the global switches and the aliases available in the current role. When used after an alias, **/?** gives you the verbs and switches available for that alias. Following a verb, **/?** gives you the details for that verb.

For example:

For a list of alias and syntax, at the command prompt, type `wmic` to start the WMIC shell and then type:

```
/?
```

To display options available for the Process alias, from within the WMIC shell type:

```
process /?
```

**/?:FULL** gives all available details. This can be especially useful for the GET and CALL verbs.

**Notes**

- You can also use a hyphen (**-?**) instead of a forward slash (**/?**).
- For the PATH and CLASS command, progressive help is available for a singular NODE only (on both local and remote computers).

For more information, see Windows Management Instrumentation Command-line.

# WMIC aliases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WMIC aliases**
Aliases are intermediary facilitators that help you access the WMI infrastructure when you are running WMIC. Aliases are friendly names used to capture the features of a WMI class that are relevant to some specific task such as disk or network administration.

Aliases can be used to provide:

- Clearer names for WMI classes, properties, and method.
- Useful output formats of arranged properties.

  The output formats can include specific property values or can be formatted in a manner appropriate to some specific presentation strategy or function. For example, an alias might have a "BRIEF" format that will list only property values essential for the identification of the objects visible through the alias. Management data is retrieved in XML format and processed by built-in or custom XSL output formats. For more information, see Creating and editing formats in WMIC.

WMIC provides aliases for the most commonly used classes of management information. You can add aliases for other classes. Thus, aliases in WMIC may be of the following types:

- Default aliases in the root\cli namespace.
  - For a list of aliases included with WMIC, type `/?` at the WMIC command prompt.
  - For more information about a specific alias, type: `alias_name /?`
  - For more information about alias schema, type: `Alias alias_name list brief`
- Aliases added by you in the root\cli namespace.
- Aliases defined by you in other namespaces. If you define aliases in other namespaces, you must use the **/ROLE** switch to specify the namespace when using it with WMIC. For more information, see Creating and editing aliases.

For information about WMI security, see Managing WMI security.

For information about WMIC security, see Security with WMIC.

For more information, see Windows Management Instrumentation Command-line.

# WMIC verbs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WMIC Verbs**

Most aliases support the following verbs. Aliases might also support other verbs. To use verbs in WMIC, enter the alias name followed by the verb. For more information about **<alias> /?**, see "Aliases" in Help. If an alias does not support a verb, you receive the message "provider is not capable of the attempted operation."

| Verb | Action | Parameters or Verb-specific switches | Example |
|---|---|---|---|
| ASSOC | Returns the result of the query: `Associators of {<wmi object>}` Where <wmi object> is the path of objects returned by the **PATH** or **CLASS** commands. The results are instances associated with the object. When ASSOC is used with an alias, the classes associated with the class underlying the alias are returned. By default, the output for class is in TABLE format. See the following table of switches for /ASSOC. | Optionally, an output file format, such as LIST, MOF, or other. | OS ASSOC |
| CALL | Executes methods. | Method and parameter list if appropriate. Parameter lists are comma delimited. Use **SERVICE CALL /?** to get a list of available methods and their parameters for the current alias. | SERVICE WHERE CAPTION='TELNET' CALL STARTSERVICE |
| CREATE | Creates a new instance and sets the property values for the new instance. This cannot be used to create a new class. | Properties equated to values, delimited with commas. Use **CREATE /?** for a list of property names for the alias. | ENVIRONMENT CREATE Name="WMIC_test",VariableValue="WMIC_test_value",UserName="*SYSTEM*" |
| DELETE | Deletes the current instance or set of instances. This can be used to delete a class. | /INTERACTIVE (prompt to confirm) or /NOINTERACTIVE (do not prompt to confirm). | PROCESS WHERE NAME="CALC.EXE" DELETE |

| | | | |
|---|---|---|---|
| GET | Get specific properties. | Property name or switch. See the table of switches for /GET below. Also use **GET /?** for a list of property names and switches for the alias. | PROCESS GET NAME |
| LIST | Show data. LIST is the default verb. | See the following tables of adverbs and switches for LIST. | PROCESS LIST BRIEF |
| SET | Property set operations. | Properties equated to values, delimited with commas. Use **SET /?** for a list of property names for the alias. | ENVIRONMENT WHERE Name="WMIC_test" SET VariableValue="WMIC_test_value1" |

The **List** verb has the following adverbs. To use adverbs in WMIC, enter the alias name followed by a verb and adverb. For more information about **<alias> /?**, see "Aliases" in Help.

| **Adverb** | **Results** |
|---|---|
| BRIEF | A core set of the properties. |
| FULL | The full set of properties. This is the default set of LIST properties. |
| INSTANCE | The instance paths only. |
| STATUS | The status and related properties of the object. |
| SYSTEM | System properties. |
| Alias-specific or user format | Alias-specific or user defined formats might be defined by providing distinct lists of properties and a format to be used in displaying them. |
| WRITEABLE | The writeable properties of the objects. |

The **List** verb has the following switches. To use verb-specific switches in WMIC, enter the alias name followed by a switch (verbs and adverbs might also be used). For more information about **<alias> /?**, see "Aliases" in Help.

| **Switch** | **Effect** |
|---|---|
| /TRANSLATE: <translation table> | Translate the output using the translation table named by the command. BasicXml and NoComma are translation tables included with WMIC. |
| /EVERY:<interval> | Repeat the command every X seconds; X is determined by the interval. |
| /FORMAT:<format specifier> | Specify a keyword or XSL file name to format the data, as explained in the following note. |

The **Get** verb has the following switches.

| **Switch** | **Effect** |
|---|---|
| /VALUE | The output is formatted with each value listed on a separate line and with the name of the property. |
| /ALL | The output is formatted as a table. The default output format is /ALL. |

| | |
|---|---|
| /TRANSLATE: <translation table> | Translate the output using the translation table named by the command. BasicXml and NoComma are translation tables included with WMIC. |
| /EVERY:<interval> | Return values every X seconds, X is the interval. |
| /FORMAT:<format specifier> | Specify a keyword or an XSL file name to format the data, as explained in the following note. |

The **Assoc** verb has the following switches.

| Switch | Effect |
|---|---|
| /RESULTCLASS: <classname> | The returned endpoints associated with the source object must belong to or be derived from the specified class. |
| /RESULTROLE: <rolename> | The returned endpoints must play a particular role in their association with the source object. |
| /ASSOCCLASS: <assocclass> | The returned endpoints must be associated with the source through the specified class or one of its derived classes. |

**Notes**

- For both the LIST and GET verbs, the format specifier is used to designate either the name of an XSL file or the standard formats TABLE or LIST. **/FORMAT:TABLE** is equivalent to **/ALL**, which is the default output format.

- MOF-formatted output can be captured to a file (using the **/RECORD** switch) and the file can be compiled using Mofcomp.exe on another computer to import it into WMI on the other system. However, make sure that computer-specific properties are edited as appropriate for the other computer.

- The order of the /FORMAT and /TRANSLATE switches is important. If /TRANSLATE follows /FORMAT, the output is formatted and then translated. If /FORMAT follows /TRANSLATE, the output is translated and then formatted.

**Examples**

The following examples demonstrate the effective use of the verbs and their adverbs and switches.

- PROCESS GET NAME /FORMAT:XML.XSL /TRANSLATE:BASICXML

  The process names are output in XML format. Without the **/TRANSLATE** switch, the characters around the values are displayed as their HTML equivalents.

- PROCESS WHERE Handle="0" ASSOC

  The instances associated with a specific object (unlike the classes associated with an alias) are returned.

For information about WMI security, see Managing WMI security.

For information about WMIC security, authentication, and authorization, see Security with WMIC.

For more information, see Windows Management Instrumentation Command-line.

# WMIC switches

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WMIC switches**

Global switches are used to set defaults for the WMIC environment. You can view the current value of the conditions that the global switches set by entering the **CONTEXT** command.

| Switch | Condition | Possible values | Default |
|---|---|---|---|
| /NAMESPACE | The namespace that the aliases typically use | Any namespace | root\cimv2 |
| /ROLE | The namespace that WMIC typically looks in for aliases and other WMIC information. When the role changes, the WMIC interactive command prompt changes to match the role. WMIC only has one role by default (\\root\cli), so you only have switch roles if you have a management product that has defined other roles, or if other roles have been created at your organization. | Any namespace that contains aliases. | root\cli |
| /NODE | Computer names, comma delimited. All commands are synchronously executed against all computers listed in this value. File names must be prefixed with '@'. Computer names within a file must be comma-delimited, or put on separate lines, or both. | Any computer name, a list of computer names, or a file with computer names in it. | Local computer name |
| /IMPLEVEL | Impersonation level | Anonymous, Identify, Impersonate, Delegate | Impersonate |
| /AUTHLEVEL | Authentication level | Default, None, Connect, Call, Pkt, Pktintegrity, Pktprivacy | Pktprivacy |
| /LOCALE | Locale | MS_409 (English), MS_411 (Japanese), MS_40B (Finnish), and so on. | The default language on the computer when WMIC is installed. |
| /PRIVILEGES | Enable all privileges. | ENABLE or DISABLE | Enabled |
| /TRACE | The success or failure of all functions used to execute WMIC commands is displayed. | ON or OFF | Off |
| /RECORD | Records all output to an XML file. Output is also displayed at the command prompt. | File name | Not set and no default file name exists. |
| /INTERACTIVE | Typically, delete commands are confirmed. | ON or OFF | OFF in NON-INTERACTIVE mode; ON in INTERACTIVE mode |

| | | | |
|---|---|---|---|
| /FAILFAST | Whether or not the **/NODE** computers are checked before trying to execute the WMIC commands against them. When FAILFAST is ON, WMIC pings the computers in the **/NODE** switch before sending WMIC commands to them. If they do not respond to the ping, the WMIC commands are not executed for them. | ON or OFF | OFF |
| /USER | A user name to be used by WMIC when accessing the /NODE computers or computers specified in aliases. You are prompted for the password. A user name cannot be used with a local computer. | Any user name. | Not set |
| /PASSWORD | A password to be used by WMIC when accessing the /NODE computers (possibly including the local computer). The password is visible at the command line. | Any password | Not set |
| /OUTPUT | Specifies a mode for output redirection. All output is directed to the destination given only. Output does not appear at the command line. The destination is cleared before the output begins. | STDOUT, CLIPBOARD, or a file name. STDOUT is the command line. Clipboard is the Windows clipboard. The output can then be pasted to any program that accepts data in the format produced. For more details, see the following note. | STDOUT |
| /APPEND | Specifies a mode for output redirection. All output is directed to the destination given only. Output does not appear at the command line. The destination is not cleared before the output begins. The new output is appended to the current contents of the destination. | STDOUT, CLIPBOARD, or a file name. STDOUT is the command line. Clipboard is the Windows clipboard. The output can then be pasted to any program that accepts data in the format produced. See the following note for more details. | STDOUT |
| /AGGREGATE | Used with the **LIST** and **GET/EVERY** switch. If AGGREGATE is ON, **LIST** and **GET** display their results when all computers in the NODE property have either responded or timed out. If AGGREGATE is OFF, **LIST** and **GET** display their results as soon as they are received. | ON or OFF | ON |
| /AUTHORITY | Specifies the authority type for the connection. | Needed if the value assigned to the IMPLEVEL switch is Delegate (for example, /IMPELEVEL:Delegate). It contains the authority definition string: "kerberos:*TargetDomainName**TargetComputerName*". For this setting to succeed, computers need to have trust for delegation enabled on the Domain Controller. | Not set |

| **Switch** | **Example** |
|---|---|
| /NAMESPACE | /NAMESPACE:\\root or /NAMESPACE:SMS |
| /ROLE | /ROLE:\\root |
| /NODE | /NODE:"*TESTSERVER1*","*TESTSERVER2*",@"*C:\COMPUTERLIST.TXT*" |
| /IMPLEVEL | /IMPLEVEL:Impersonate |

| | |
|---|---|
| /AUTHLEVEL | /AUTHLEVEL:default |
| /LOCALE | /LOCALE:MS_411 |
| /PRIVILEGES | /PRIVILEGES:ENABLE |
| /TRACE | /TRACE:ON |
| /RECORD | /RECORD:output.xml |
| /INTERACTIVE | /INTERACTIVE:ON |
| /FAILFAST | /FAILFAST:ON |
| /USER | /USER:JSMITH or /USER:"" (to reset the user name and password to nothing). |
| /PASSWORD | /PASSWORD:PASSWORD |
| /OUTPUT | /OUTPUT:CLIPBOARD |
| /APPEND | /APPEND:CLIPBOARD |
| /AGGREGATE | /AGGREGATE:OFF |
| /AUTHORITY | /AUTHORITY:"kerberos:*TargetDomainName**TargetComputerName*" |

**Notes**

- To browse for possible ROLE and NAMESPACE namespaces, use the command `WMIC PATH __NAMESPACE`. To see available namespaces in other namespaces, use the **/NAMESPACE** switch to change a target namespace. For example, use `/NAMESPACE:\\root` to change the ROOT namespace.

- The **GET**, **LIST**, and **ASSOC** verbs have WMIC verbs.

- The **/APPEND** and **/OUTPUT** switches can contradict each other (indicating that the results should be sent to different places). When they do contradict each other, **/OUTPUT** predominates in all cases except where **/APPEND** is directing the results to a file. In that case, the results go to both the **/OUTPUT** and **/APPEND** destinations.

- Switches can be used on the same line as commands. For example, you can use `/NODE:"SERVER1" SERVICE` to get a list of all services on SERVER1. The **/NODE** switch works as described above, and all subsequent commands are also applied to SERVER1. However, if **/OUTPUT** and **/APPEND** are used on the same line as a command, the results redirection specified by **/OUTPUT** and **/APPEND** only apply to that command. The results of subsequent commands are directed to the previous value of the OUTPUT and APPEND switches.

- To reset a switch, set its value to "". For example: `/NODE:""` sets NODE back to the current computer name.

For more information, see Windows Management Instrumentation Command-line.

# WMIC commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WMIC commands**

The following commands are available at all times:

| Command | Purpose |
|---|---|
| CLASS | Escape from the default alias mode of WMIC to access classes in the WMI schema directly. |
| PATH | Escape from the default alias mode of WMIC to access instances in the WMI schema directly. |
| CONTEXT | Display the current values of all global switches. |
| QUIT | Exit from WMIC. |
| EXIT | Exit from WMIC. |

For more information about the global switches displayed by CONTEXT, see WMIC switches.

For information about how WMI validates users, see Managing WMI security.

For information about WMIC security, authentication, and authorization, see Security with WMIC.

# Creating and Customizing WMIC Components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating and customizing WMIC components**
This section covers:

- Alias namespaces and classes.
- Creating and editing aliases.
- Creating and editing formats in WMIC.
- Creating and editing translation tables.

# Alias namespaces and classes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Alias namespaces and classes**

Examples of WMI alias classes are listed in the following table:

| Alias class | Description |
|---|---|
| MSFT_CliTranslateTable | Verb-specific switches. WMIC includes translation tables called BasicXml and NoComma. |
| MSFT_CliTranslateTableEntry | Used in the translation table. |
| MSFT_CliSeeAlso | The MSFT_CliSeeAlso association provides a relationship between an alias and any aliases with some kind of functional relationship to the alias. You can extend the alias schema to allow aliases to be categorized and navigated in various ways. Use the **SeeAlso** verb of the **Alias** alias to see this information. |
| MSFT_CliAlias | MSFT_CliAlias class is the principal class in the alias schema. Instances of this class are aliases. |

An **MSFT_CliAlias** instance has the following properties:

| Property | Description |
|---|---|
| FriendlyName | The name of the alias. It must be unique. |
| Description | A description of the alias. This is the descriptive text when **/?** is entered at the WMIC command line. |
| Formats | A list, each of which has a name and a list of properties (objects of the class MSFT_CliProperty) to be displayed for that format. All formats are objects of the class MSFT_CliFormat. |
| Verbs | Verbs: A list, each of which are the various behaviors available through this alias. The behaviors come in two forms:<br>• Standard verbs, which are directly supported by the utility.<br>• User-defined verbs, which must map to some method defined for the target of the alias.<br>All verbs are objects of the class MSFT_CliVerb. |
| Qualifiers | A list, similar to WMI qualifiers. All qualifiers are objects of the class MSFT_Qualifier. |
| Target | A list, similar to WMI qualifiers. All qualifiers are objects of the class MSFT_Qualifier. |
| PWhere clause | Optional WHERE clause that limits the Target. It have substitution values which are the parameters of the alias. The substitution values are marked with **#**. If multiple parameters are needed, they are matched with the # markers in sequence. |
| Connection | Details on which computers to connect to, the security details to be used, and so on. If a connection is not specified, the computers to be accessed is the value of **/NODE**, and the namespace is the value of **/NAMESPACE**. If a user name and password are not provided, then the value of **/USER** and **/PASSWORD**, if available, is used (otherwise the current account is used). Connection is an object of the class MSFT_CliConnection. |
| View an alias schema | Use a metaalias Alias to view an alias schema; example: `ALIAS OS` |

**MSFT_CliAlias** properties contain the following classes to define objects that are used in the properties:

| Class | Used in | Description |
|---|---|---|
| MSFT_CliFormat | aliases | The MSFT_CliFormat class defines a named format that determines the set of properties to be displayed by the utility on a **LIST** command. Each MSFT_CliFormat instance has a name, a list of properties and a Format property that contains the name of a style sheet (XSL file) that must be used to format the property display. |
| MSFT_CliProperty | formats | Each MSFT_CliFormat instance contains a list of MSFT_CliProperty instances. Instances of the MSFT_CliProperty class define the set of properties to be displayed for a given format. Each MSFT_CliProperty instance has a derivation that might be either a constant, the name of a property or a WMI command line **GET** command. Each MSFT_CliProperty also has a localized description.<br><br>Formats are related to a specific XSL by the NAME property of the format object. FULL, BRIEF, and INSTANCE correspond to particular XSLs, and the NAME can specify the name of a custom XSL file. |
| MSFT_CliVerb | aliases | Instances of the MSFT_CliVerb class represent verbs defined for the alias. Verbs are either standard verbs (**GET**, **LIST**, **CALL** and **SET** ) or user defined. Standard verbs have descriptions and usages but no parameters. |
| MSFT_CliParam | verbs | |
| MSFT_CliQualifier | aliases and verbs | WMIC qualifiers serve the same purpose as WMI qualifiers. |
| MSFT_CliConnection | aliases | Instances of this class provides connection details. |

**Notes**

- When viewing the role namespaces using WMI browsing tools (such as CIM Studio), only the MSFT_CliAlias and MSFT_CliTranslateTable classes have instances. Their properties might include objects of the other classes, but the other classes will not have instances of their own.

- For this reason, when creating aliases (see Creating and editing aliases), you create instances primarily in the MSFT_CliAlias class. You create instances of the other classes as properties of the MSFT_CliAlias class or as properties of its properties.

For more information, see Windows Management Instrumentation Command-line.

# Creating and editing aliases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating and editing aliases**

You can use the WMI tools to create aliases as Managed Object Format (MOF) files. You can use WMI Common Information Model (CIM) Studio to generate a MOF file for a selected class. This file should be edited to create an instance of MSFT_CliAlias. The MOF files can then be used to register the alias. Aliases are registered with WMI by compiling the MOF form of the alias.

When viewing the role namespaces using WMI browsing tools (such as WMI CIM Studio), only the MSFT_CliAlias and MSFT_CliTranslateTable classes have instances. Their properties might include objects of the other classes, but the other classes do not have instances of their own. For this reason, when creating aliases, you primarily create instances in the MSFT_CliAlias class. You create instances of the other classes as properties of the MSFT_CliAlias class or as properties of its properties.

Localization is supported by using localized text arrays defined in the schema. Localization of aliases descriptions is done through the MFL file (CliEgAliases.mfl).

**Adding roles** to any WMI namespace can be done, for example, by copying WMIC's CliEgAliases.mof file, and then changing the namespace reference to the namespace that is to be used as a role. The new file can then be compiled with MOFCOMP.exe. Aliases can then be added as described above.

Any WMI namespace can serve as a role; but without the additional steps above, such a role would not have WMIC's interactive help or the base classes required for aliases.

For more information, see Windows Management Instrumentation Command-line.

# Creating and editing formats in WMIC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating and editing formats in WMIC**

Windows Management Instrumentation Command-line (WMIC) uses two types of formats to display properties: **LIST** format and **OUTPUT** format.

**LIST** formats allow displaying and recording of the properties of an alias in a specific way, such as:

| List Format | Properties |
|---|---|
| `Process list brief` | displays a short list of properties. |
| `Process list full` | displays a full list of properties (this is similar to `process get /all`). |
| `Process list memory` | lists properties related to a process memory. |

To see what LIST formats are available for an alias, type:

```
<alias name> list /?.
```

For example:

```
os list /?
```

Each LIST format has a default output format name specified as part of alias definition. For example, FULL is usually associated with TABLE, while BRIEF, with LIST (value list).

As an optional step, add the LIST formats group to the alias class. This makes the formats discoverable and easier to use. You can also specify the default stylesheet to be used for the format output.

**OUTPUT** formats are provided by WMIC using the `/FORMAT` switch to display and record output in a desired way. WMIC uses Extensible Style sheet Language (XSL) to transform and format the Extensible Markup Language (XML) output. The provided formats are stored in the WBEM directory as XSL files. Formats are used with **GET**, or **LIST**. Specific formats can be specified using **GET /FORMAT:**, **LIST /FORMAT:**, **GET /ALL**, or **GET /VALUE**.

Use an XSL editor. You might want to first capture some XML data to use for experimenting while using the XSL editor.

**Custom formats** may be created to display WMIC queries. If you create custom formats, you should also define them in the **Formats** property of the **MSFT_CliAlias** class. Each format must have a name that is unique within the alias.

**Other Output formats** are included in WMIC. These formats can be set as defaults or as specified using keywords:

| Keywords | Output Formats |
|---|---|
| mof.xsl | Converts XML data to Managed Object Format (MOF) format. If aliases or other WMI objects are defined in XML format, this XSL can be used to convert them to MOFs so that they can be MOFCOMPed. - /FORMAT:MOF format. |
| texttable.xsl | /ALL or /FORMAT:TABLE. |
| texttable.xsl | /VALUE format—one line per value. |

For more information on mapping, see **XSL-MAPPINGS.XML** in **%windir%\system32\wbem**.

**Additional XSL files** that are included in WMI and can be useful when you use WMIC are:

| XSL Files | |
|---|---|
| xml.xsl | Similar to rawxml.xsl, except that the output is in HTML format. |
| hform.xsl | Html output—the data is displayed in a vertical table (one row per property or instance name). The table can be very long. |

| | |
|---|---|
| htable.xsl | Html output—the data is displayed in a horizontal HTML table (one row per instance). The table can be very wide. you can specify the column that the data is typically sorted by. Enclose this file name in quotation marks. The column and its datatype are specified by using the following syntax: `/FORMAT:"htable.xsl":"datatype=<datatype>":"sortby=<property>"`. For example, `PROCESS GET /FORMAT:"htable.xsl":"datatype=number":"sortby=ProcessId"`. |
| rawxml.xsl | To output XML data in XML format. All WMIC data is extracted in XML format, so this format gives you data the way that WMIC sees it. This can be useful for creating your own XSL files. Outputting data in XML format is also useful if you want to provide the data to systems that accept XML data. |
| csv.xsl | Comma-delimited—the data is displayed with commas between each value. This format is suitable for importing the data into Microsoft Excel or a similar program. |

**Caution**

- When editing default views, be sure to save your changes to a format outside WMI. If you update WMI or the operating system, the new version of WMIC might return the aliases or output formats to the defaults, overwriting your changes.

For step-by-step procedures on how to create and edit formats, see Create a format in WMIC and Edit a format in WMIC.

For more information, see Windows Management Instrumentation Command-line.

# Creating and editing translation tables

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating and editing translation tables**

Like formats, translation tables allow you to change the output of WMIC commands to suit your needs. Translation tables are not as flexible or powerful as formats, but they are much simpler to create and edit.

Translation tables are used when you specify **/TRANSLATE** on the **LIST** or **GET** verbs. WMIC includes translation tables called BasicXML and NoComma. BasicXML converts the HTML symbols for greater-than and less-than characters into the actual characters.

Translation tables are defined in the MSFT_CliTranslateTable WMI class.

For more information, see Windows Management Instrumentation Command-line.

# Troubleshooting WMIC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Parameter not valid
- Value map does not contain the input value for this property
- RPC error on the unsuccessful connection to a computer or node
- Query fails or is slow
- Output is not properly aligned in rows and columns
- Localized version of WMIC is not available (in some Multilingual User Interface (MUI) or Language Interface Packs
- XSL files on Windows XP Professional are not supporting arrays
- Input parameter is of OBJECT type

**Parameter not valid**
**Cause:** A parameter name is misspelled or does not apply to the object in question. A required parameter might also be missing, which can include key values or writeable values.

**Solution:** You must specify the name, user name, and variable value when you create an environment variable.

**Value map does not contain the input value for this property**
**Cause:** Could be that the needed parameter is not used in the query.

**Solution:** You can interpret the value for an input parameter by referring to the schema information, the method's description qualifier in the corresponding Managed Object Format (MOF) file, or to the MSDN description of the method.

Some schema information can be retrieved using WMIC, for example:

- `wmic:root\cli>alias OS list full /format:table`
- `wmic:root\cli>process call /?:full`

**RPC error on the unsuccessful connection to a computer or node.**
**Cause:** The computer or node has a firewall set to block incoming calls.

**Solution:**

- Ping computer's IP for the availability of the computer.
- Disable firewall on computer.

**Cause:** The computer has UNIX, Linux, Windows 95, Windows 98, Windows Millennium Edition, or Windows NT 4 operating system that has no network-enabled WMI or WBEM support.

**Solution:**

- Enable WMI or WBEM support in the operating system.

- Deploy operating systems with network-enabled WMI or WBEM support.

**Query fails or is slow**
**Cause:** The query does not have enough parameters.

**Solution:** Narrow down query using **Named** parameters or **WHERE** clause. For example:

- `wmic:root\cli>Alias OS list brief`
- `wmic:root\cli>nteventlog application get /format:list`
- `wmic:root\cli>UserAccount where "Domain='DomainName' and Name='UserName'" get /all`

**Output is not properly aligned in rows and columns.**
**Cause:** Display space is too small.

**Solution:**

- Increase Screen Buffer Size for Command (CMD) window.
- Use WordPad.exe instead of Notepad to open large text files.
- Use more criteria on query to display only necessary information.
- Use a custom XSL file to display custom formatting.

**Localized version of WMIC is not available (in some Multilingual User Interface (MUI) or Language Interface Packs)**
**Cause:** Installed localized version of operating system.

**Solution:** Use the /LOCALE:ms_409 switch to set up English Locale as the default.

**XSL files on Windows XP Professional are not supporting arrays**
**Cause:** Operating system does not support XSL file arrays.

**Solution:** Use XSL from the Windows Server 2003 family to query other computers, including Windows XP Professional

**Input parameter is of OBJECT type**
**Cause:** Support is not fully implemented yet for input parameters which are of the OBJECT type. The default value (an empty string: "") is the only value currently supported.

**Solution:** Use the following command to set sharing of the c:\test directory for three users:

```
WMIC SHARE CALL Create "","test","3","TestShareName","","c:\test",0
```

# Windows Management Instrumentation Control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Management Instrumentation Control**

- For help on using WMI Control, see Windows Management Instrumentation Control How To....
- To understand more about WMI Control, see Windows Management Instrumentation Control Concepts.

# Windows Management Instrumentation Control How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Turn WMI error logging on or off
- Backup or restore the WMI repository
- Change the WMI scripting settings
- Authorize WMI users and set permissions
- Modify permissions or delete authorized users

# Turn WMI error logging on or off

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To turn WMI error logging on or off**

1. Open WMI Control.
2. In the console tree, right-click **WMI Control**, and then click **Properties**.
3. Click the **Logging** tab.
4. To turn logging off, click **Disabled** under **Logging level**.
5. To turn logging on click either **Errors only** or **Verbose** under **Logging Level**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the WMI Control console, click **Start**, click **Run**, type **wmimgmt.msc**, and then click **OK**.
- **Errors only** reports errors that occur within WMI. **Verbose** reports all actions and errors.
- Verbose logging can negatively impact system performance, so select **Verbose** only when you need more extensive information about the events leading to errors.
- Log files are maintained for the WMI core and framework, as well as for each WMI-enabled provider.
- You can set the maximum size for log files under **Maximum size**. When a log file has reached the maximum size, it is renamed with the .lo_ extension, and a new log file is created with the .log extension. When that file reaches the defined maximum size, the first file is deleted and the process is repeated.
- You can set logging levels on a remote computer or a local computer. To access a remote computer, right-click **WMI Control**, click **Connect to another computer**, click **Another computer**, and then type the name of the computer to which you want to connect. If you are using WMI Control from the Computer Management console, right-click the **Computer Management** node to connect to the other computer.
- It is recommended that you include log files when reporting issues associated with WMI services.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
WMI Control Overview
Working with MMC console files

# Backup or restore the WMI repository

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To backup or restore the WMI repository**

1. Open WMI Control.
2. In the console tree, right-click **WMI Control**, and then click **Properties**.
3. Click the **Backup/Restore** tab.
4. Click the **Back Up Now** button or **Restore Now** button.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the WMI Control console, click **Start**, click **Run**, type **wmimgmt.msc**, and then click **OK**.
- When you click the **Back Up Now** button to back up the repository manually, you can back up to a specified file name, and later restore from that saved file. This can be useful if you want to save a specific configuration of classes.
- You can backup or restore a remote computer or a local computer. To access a remote computer, right-click **WMI Control**, click **Connect to another computer**, click **Another computer**, and then type the name of the computer to which you want to connect. If you are using WMI Control from the Computer Management console, right-click the **Computer Management** node to connect to the other computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
WMI Control Overview
Working with MMC console files

# Change the WMI scripting settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the WMI scripting settings**

1. Open WMI Control.
2. In the console tree, right-click **WMI Control**, and then click **Properties**.
3. Click the **Advanced** tab.
4. Under **Default namespace for scripting**, click **Change**.
5. In the **Browser for Namespace** dialog box, select a different default namespace as a target for your scripts.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the WMI Control console, click **Start**, click **Run**, type **wmimgmt.msc**, and then click **OK**.
- Set the default targeted namespace to one you use most often in your WMI scripts. If no namespace is identified within a script, the default is used.
- If the computer you are manage runs Microsoft Windows NT 4.0, you can define whether WMI scripts can be run from within Active Server Pages (ASP). If you enable the option, your ASP connects to Windows Management services using the process level identity of Internet Information Services, which might compromise security.
- You can change the scripting namespace on a remote computer or a local computer. To access a remote computer, right-click **WMI Control**, click **Connect to another computer**, click **Another computer**, and then type the name of the computer to which you want to connect. If you are using WMI Control from the Computer Management console, right-click the **Computer Management** node to connect to the other computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
WMI Control Overview
Working with MMC console files
Managing WMI security

# Authorize WMI users and set permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To authorize WMI users and set permissions**

1. Open WMI Control.
2. In the console tree, right-click **WMI Control**, and then click **Properties**.
3. Click the **Security** tab.
4. Select the namespace for which you want to give a user or group access, and then click **Security**.
5. In the **Security** dialog box, click **Add**.
6. In the **Select Users, Computers, or Groups** dialog box, enter the name of the object (user or group) that you want to add. Click the **Check Names** button to verify your entry and then click **OK**. You might have to change the location or use the **Advanced** button to query for objects. See the dialog box help for more details.
7. In the **Security** dialog box, under **Permissions**, select the permissions to allow or deny the new user or group.

| Level | Description |
|---|---|
| Execute Methods | Allows methods exported from the WMI classes or instances to be run. |
| Full Write | Allows full read, write, and delete access to all WMI objects, classes, and instances. |
| Partial Write | Allows write access to static WMI objects. |
| Provider Write | Allows write access to objects that are provided by providers. |
| Enable Account | Allows read access to WMI objects. |
| Remote Enable | Allows remote access to the namespace. |
| Read Security | Allows read-only access to WMI security information. |
| Edit Security | Allows read and write access to WMI security information. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the WMI Control console, click **Start**, click **Run**, type **wmimgmt.msc**, and then click **OK**.
- You can set permissions on a remote computer or a local computer. To access a remote computer, right-click **WMI Control**, click **Connect to another computer**, click **Another computer**, and then type the name of the computer to which you want to connect. If you are using WMI Control from the Computer Management console, right-click the **Computer Management** node to connect to the other computer.
- On computers running Windows 95, Windows 98, or Windows ME, all users have full control locally. Security settings are only relevant for remote connection to a computer running Windows 95, Windows 98, or Windows ME.
- You can delete a user's or group's authorization to access WMI services by selecting that user or group and clicking

**Remove**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
WMI Control Overview
Modify permissions or delete authorized users
Managing WMI security
Working with MMC console files
Permissions for files and folders

# Modify permissions or delete authorized users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify permissions or delete authorized users**

1. Open WMI Control.
2. In the console tree, right-click **WMI Control**, and then click **Properties**.
3. Click the **Security** tab.
4. Select the namespace for which you want to change a user's or group's permissions or delete a user or group, and then click **Security**.
5. In the **Security** dialog box, under **Name**, select the user or group you want to remove or whose permissions you want to modify.
6. Perform one of the following steps:
   - To remove the user or group, click **Remove**.
   - To change security settings, under **Permissions**, modify the permissions which are allowed or denied to the user or group by clearing or selecting the appropriate check boxes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the WMI Control console, click **Start**, click **Run**, type **wmimgmt.msc**, and then click **OK**.
- WMI security is verified only when a user logs on to Windows Management services. Thus, any changes made to a user's permissions while the user is connected do not take effect until the next time the user logs on to WMI. For example, if a user's access is revoked, the changes do not take effect until the user logs off Windows Management services and attempts to log on again.
- You can modify user or group permissions on a remote computer or a local computer. To access a remote computer, right-click **WMI Control**, click **Connect to another computer**, click **Another computer**, and then type the name of the computer to which you want to connect. If you use WMI Control from the Computer Management console, right-click the **Computer Management** node to connect to the other computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
WMI Control Overview
Authorize WMI users and set permissions
Managing WMI security
Working with MMC console files

# Windows Management Instrumentation Control Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- WMI Control Overview
- WMI Overview
- Managing WMI security

# WMI Control Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WMI Control overview**
Windows Management Instrumentation (WMI) Control is a tool that enables you to configure WMI settings on a remote computer or local computer. Using the WMI Control, you can manage the following tasks remotely.

**Authorize users or groups and set permission levels**
You can enable an individual user, group, or namespace to access network objects and perform WMI tasks and services. For example, you can enable a group to manage WMI's Common Information Model (CIM) objects on their local computers.

**Configure error logging**
You can turn error logging on or off and, if turned on, set it to report errors only (the default) or all actions (verbose). Error logging can help you troubleshoot WMI problems. You can also define a maximum size for log files and their folder location.

**Back up the repository**
You can configure the WMI Control to back up your repository on a regular schedule, or you can do it manually at any time. The repository is the database of objects that you can access through WMI. You can also restore a previous version of the repository.

**Change the default namespace for scripting**
You can change the default namespace that is targeted in WMI scripts.

For more information, see WMI Overview

For information about using the command-line interface to WMI, see Windows Management Instrumentation Command-line.

# WMI Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WMI Overview**

Windows Management Instrumentation (WMI) is the Microsoft implementation of WBEM, an industry initiative to establish standards for accessing and sharing management information over an enterprise network. WMI is WBEM-compliant and provides integrated support for the Common Information Model (CIM), the data model that describes the objects that exist in a management environment.

WMI includes a CIM-compliant object repository, which is the database of object definitions, and the CIM Object Manager, which handles the collection and manipulation of objects in the repository and gathers information from WMI providers. WMI providers act as intermediaries between WMI and components of the operating system, applications, and other systems. For example, the registry provider draws information from the registry, while the SNMP provider provides data and events from SNMP devices. Providers provide information about their components, and might provide methods to manipulate the components, properties that can be set, or events that can alert you to changes in the components.

WMI can be used by computer management tools, such as Microsoft Systems Management Server, to help you manage your computers. WMI is also used by other Microsoft technologies and tools, such as Microsoft Health Monitor and Microsoft Operations Manager, and by other vendors of computer management systems. You can use WMI with programming or scripting systems (such as Windows Script Host) to retrieve configuration details about most aspects of your computer systems, including server applications, or to affect changes to your systems. For more information see the WMI SDK.

Several management tools are WMI-enabled, including System Properties, System Information, and the Dependencies component of Services. Brief descriptions of these components follow:

- System Properties lets you view and change system properties on a local or remote computer. You can restart a remote computer to apply settings changes or to detect new hardware, view the computer name and domain information for other computers on your network, or change the settings for the virtual memory paging file on a computer that might run programs requiring a lot of memory. For more information about System Properties, see System Properties.

- System Information collects and displays configuration information about your system. This is is especially useful when a support technician is troubleshooting your system. For more information, see System Information.

- Services helps you manage the services on your computer. Services dependencies identify the services on which the current service depends and the services that depend on it. For more information about Services, see Services Concepts.

For technical information about developing the WMI system, see the Windows Management Instrumentation documentation at the Microsoft Platform Software Development Kit (SDK) Web site

For information about using the WMI Tester to view and modify CIM objects, see Windows Management Instrumentation Tester overview.

For information about using the command-line interface to WMI, see Windows Management Instrumentation Command-line.

For information about using the WMI Control to configure WMI settings and security on a remote or local computer, see WMI Control Overview.

# Managing WMI security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing WMI security**

Windows Management Instrumentation (WMI) supports a limited form of security that validates each user before the user is allowed to connect to WMI, on either the local computer or a remote computer. This security is layered on top of the operating system security. WMI does not override or circumvent security provided by the operating system.

By default, all members of the Administrators group have full control of the WMI services on the managed computer. All others have read/write/execute permissions on their local computer only.

Permissions can be changed by adding a user to the Administrators group on the managed computer or by authorizing users or groups in WMI and setting their permission level. Access is based on WMI namespaces.

On computers running Microsoft Windows 95, Windows 98, or Windows Millennium Edition (ME), all users have full control locally. However, permissions can be set for users managing a Windows 95, Windows 98, or Windows ME computer remotely.

Security is checked only when a user connects to the Windows Management service. Thus, any changes made to a user's permissions while the user is connected do not take effect until the next time the user starts a WMI service. For example, if a user's access is revoked, the changes do not take effect until the user exits from WMI and attempts to connect to the service again.

For more information, see Modify permissions or delete authorized users.

# Windows Management Instrumentation Tester

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Management Instrumentation Tester**
Using Windows Management Instrumentation (WMI) Tester, also called WBEMTest, you can view and change Common Information Model (CIM) classes, instances, and methods. You can also use WBEMTest to troubleshoot Windows Management Instrumentation (WMI) and programs that depends on WMI.

- For help with specific tasks, see Windows Management Instrumentation Tester How to....
- For general background information, see Windows Management Instrumentation Tester Concepts.

# Windows Management Instrumentation Tester How to...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**


- Display WMI tester dialog box
- Connect to WMI
- Enumerate Classes
- Create a class
- Delete a class
- Open an instance
- Delete an instance
- Execute a notification query
- Execute a notification query

# Display WMI tester dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To display WMI tester dialog box**

1. Click **Start**, click **Run**, and then type **wbemtest.exe**.

## Notes

- Performing this task on the local computer does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- Performing this task on a remote computer requires that you are a member of the Administrators group on the remote computer, or that you have been delegated the appropriate authority.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to WMI
Using WBEMTest user interface
Connect dialog box
Managing WMI security

# Connect to WMI

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to WMI**

1. Click **Start**, click **Run**, and then type **wbemtest.exe**.
2. In the **Windows Management Instrumentation (WMI) Tester** dialog box, click **Connect**.
3. In the **Namespace** box, type the path to the namespace that you want to connect to.

   **Example?**

   - To connect to the namespace where the Win32 provider classes reside, type:

     **root\cimv2**

4. Click **Connect**.

## Notes

- Performing this task on the local computer does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- Performing this task on a remote computer requires that you are a member of the Administrators group on the remote computer, or that you have been delegated the appropriate authority.
- The online help procedures use the **root\cimv2** namespace. To step through the procedures, you must be connected to this namespace instead of the default namespace.
- When you start WBEMTest, the **Windows Management Instrumentation Tester** dialog box appears. Before you can use WBEMTest functions, you must log on to WMI.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect dialog box
Managing WMI security

# Enumerate Classes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To enumerate classes

1. Open the **Windows Management Instrumentation Tester** dialog box.
2. In the **Windows Management Instrumentation (WMI) Tester** dialog box, click **Connect**.
3. In the **Namespace** box, type the path to the namespace that you want to connect to.

   **Example?**

   - To connect to the namespace where the Win32 provider classes reside, type:

     **root\cimv2**

4. Click **Connect**.
5. Click **Enum Classes**.
6. In the **Superclass Info** dialog box, type the name of the superclass to enumerate the child classes of a specific class.

   - Select **Immediate only** option if you want to see only the immediate subclasses of the superclass you specify.
   - Select **Recursive** option if you want to see all the classes of the superclass, subclasses of the superclass, and all the subclasses of those subclasses.

7. To see all the top level classes for the current namespace, leave the **Superclass Info** dialog box empty and select **Immediate only** option.
8. To see all the classes in a namespace, subclasses of the superclass, and all the subclasses of those subclasses, leave the **Superclass Info** dialog box empty and select **Recursive** option.

### Notes

- Performing this task on the local computer does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- Performing this task on a remote computer requires that you are a member of the Administrators group on the remote computer, or that you have been delegated the appropriate authority.
- To open the **Windows Management Instrumentation Tester** dialog box, click **Start**, click **Run**, and then type **WBEMTest.exe**.
- You can display the child classes of a superclass by clicking **Enum Classes** in the **Windows Management Instrumentation Tester** dialog box.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Superclass Info dialog box
Query Result dialog box

# Create a class

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a class**

1. Open the **Windows Management Instrumentation Tester** dialog box.
2. In the **Windows Management Instrumentation (WMI) Tester** dialog box, click **Connect**.
3. In the **Namespace** box, type the path to the namespace that you want to connect to.

   **Example?**

   - To connect to the namespace where the Win32 provider classes reside, type:

     **root\cimv2**

4. Click **Connect**.
5. Click the **Create Class** button.

   - If you are creating a new base class, leave the **Optional Superclass** dialog box empty.
   - If you are specifying an existing class definition, the new class is based on the current namespace.
   - For example, if you specify **Win32_LogicalDisk** as the superclass, the new child class inherits the properties of the **Win32_LogicalDisk** class.

6. Click **Continue**.
7. In the **Object Editor** dialog box, select the **__Class** property.
8. Select **Edit Property**.
9. In the **Property Editor** dialog box, click **Not Null** and then type the name for the new class in the text box.
10. Click **Save Property**.
11. In the **Object Editor** dialog box, define the properties, qualifiers, and methods for the new class.
12. Click **Save Object** to create the new class.

**Notes**

- Performing this task on the local computer does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- Performing this task on a remote computer requires that you are a member of the Administrators group on the remote computer, or that you have been delegated the appropriate authority.
- To open the **Windows Management Instrumentation Tester** dialog box, click **Start**, click **Run**, and then type

**WBEMTest.exe**.

- To use the **Win32_LogicalDisk** example above, you must be logged in to the **root\cimv2** namespace.
- You can create a new class definition in the current namespace by clicking **Create Class** in the **Windows Management Instrumentation Tester** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Object Editor dialog box
Property Editor dialog box

# Delete a class

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a class**

You can delete a class in the current namespace with **Delete Class** on the **Windows Management Instrumentation Tester** dialog box.

1. Open the **Windows Management Instrumentation Tester** dialog box.
2. In the **Windows Management Instrumentation (WMI) Tester** dialog box, click **Connect**.
3. In the **Namespace** box, type the path to the namespace that you want to connect to.

   **Example?**

   - To connect to the namespace where the Win32 provider classes reside, type:

     **root\cimv2**

4. Click **Connect**.
5. Click **Delete Class**.
6. In the **Get Class Name** dialog box, type the name of the class you want to delete, and then click **OK**.

   A dialog box is displayed to warn you that the class might have derived classes or instances. This dialog box is displayed even if there are no derived classes or instances.

7. If you are sure you want to delete the class, click **Yes** to delete the class.

## Notes

- Performing this task on the local computer does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- Performing this task on a remote computer requires that you are a member of the Administrators group on the remote computer, or that you have been delegated the appropriate authority.
- To open the **Windows Management Instrumentation Tester** dialog box, click **Start**, click **Run**, and then type **WBEMTest.exe**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Get Class Name dialog box

Create a class

# Delete an instance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an instance**

1. Open the **Windows Management Instrumentation Tester** dialog box.
2. In the **Windows Management Instrumentation (WMI) Tester** dialog box, click **Connect**.
3. In the **Namespace** box, type the path to the namespace that you want to connect to.

   **Example?**

   - To connect to the namespace where the Win32 provider classes reside, type:

     **root\cimv2**

4. Click **Connect**.
5. Click **Delete Instance**.

   The **Get Object Path** dialog box appears.

6. Type the relative object path for the instance you want to delete.

   **Example?**

   - To delete an instance of MyClass where the key property is F122, type:

     **MyClass.KeyProperty="F122"**

## Notes

- Performing this task on the local computer does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- Performing this task on a remote computer requires that you are a member of the Administrators group on the remote computer, or that you have been delegated the appropriate authority.
- To open the **Windows Management Instrumentation Tester** dialog box, click **Start**, click **Run**, and then type **WBEMTest.exe**.
- You cannot delete dynamic instances of static classes unless the provider supports this capability. For example, if you attempt to delete an instance of **Win32_Process**, a "Provider Not Capable" message appears because the provider does not support the deletion of dynamic instances.
- You can delete an instance of a class by clicking **Delete Instance** in the **Windows Management Instrumentation Tester** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Get Object Path dialog box

# Execute a notification query

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To execute a notification query**

1. Open the **Windows Management Instrumentation Tester** dialog box.
2. In the **Windows Management Instrumentation (WMI) Tester** dialog box, click **Connect**.
3. In the **Namespace** box, type the path to the namespace that you want to connect to.

   **Example?**

   - To connect to the namespace where the Win32 provider classes reside, type:

     **root\cimv2**

4. Click **Connect**.
5. Click **Notification Query**.
6. In the **Enter Query** box, type an event notification query.

   **Example?**

   - An event query that notifies you when a log event occurs:

     SELECT * FROM __instancecreationevent WHERE TargetInstance ISA 'Win32_NTLogEvent' AND (TargetInstance.LogFile="Application" OR TargetInstance.LogFile="System")

7. Click **Apply**.

   - The text **Operation in Progress** appear in the lower-right corner of the **Query Result** dialog box, indicating that WMI is processing your query. Because event notification queries are never complete, the **Operation in Progress** status message is displayed on the **Query Result** dialog box until you stop the query.
   - For event notification queries that are running semisynchronously, you must specify batch count and time-out values. These values determine how events are extracted from the enumerator.
   - If your event notification query is running asynchronously, events are displayed in the **Query Result** dialog box as the events are received.
   - If you want to view an event notification, double-click the event in the **Query Result** dialog box. The **Instance Of** dialog box appears.
   - To stop the query, click **Close** .

**Notes**

- Performing this task on the local computer does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- Performing this task on a remote computer requires that you are a member of the Administrators group on the remote computer, or that you have been delegated the appropriate authority.

- To open the **Windows Management Instrumentation Tester** dialog box, click **Start**, click **Run**, and then type **WBEMTest.exe**.

- If you want your event notification query to run asynchronously, select the **Execute asynchronously** check box on the **Windows Management Instrumentation Tester** dialog box.

- You can query for event notifications with **Notification Query** on the **Windows Management Instrumentation Tester** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Query dialog box
Query Result dialog box

# Open an instance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open an instance**

1. Open the **Windows Management Instrumentation Tester** dialog box.
2. In the **Windows Management Instrumentation (WMI) Tester** dialog box, click **Connect**.
3. In the **Namespace** box, type the path to the namespace that you want to connect to.

   **Example?**

   - To connect to the namespace where the Win32 provider classes reside, type:

     **root\cimv2**

4. Click **Connect**.
5. Click **Open Instance**.
6. In the **Get Object Path** dialog box, type the relative object path for the instance you want to retrieve. And, then click **OK** to use the **Object Editor** dialog box to view or edit the class instance you have opened.

   **Example?**

   - To open an instance of the **Win32_LogicalDisk** class, type:

     **Win32_LogicalDisk.DeviceID="C:"**

## Notes

- Performing this task on the local computer does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- Performing this task on a remote computer requires that you are a member of the Administrators group on the remote computer, or that you have been delegated the appropriate authority.
- To open the **Windows Management Instrumentation Tester** dialog box, click **Start**, click **Run**, and then type **WBEMTest.exe**.
- You can open an instance of a class with **Open Instance** on the **Windows Management Instrumentation Tester** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Get Object Path dialog box
Object Editor dialog box

# Windows Management Instrumentation Tester Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Windows Management Instrumentation Tester overview
- Using WBEMTest user interface

# Windows Management Instrumentation Tester overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Management Instrumentation Tester overview**
Windows Management Instrumentation (WMI) Tester, also called WBEMTest, is a general-purpose tool for viewing and modifying WMI classes, instances, and methods during the development of Windows Management Instrumentation (WMI) providers and WMI applications. You can also use WBEMTest to troubleshoot WMI and programs that depend on WMI.

Use WBEMTest to perform the following tasks:

- Create, enumerate, and delete WMI classes and instances.
- Execute methods.
- Run queries.
- Display the Managed Object Format (MOF) code for classes and instances.

Because WBEMTest acts as an interface to the COM API for WMI, you can also use WBEMTest to check queries, classes, and instances when developing WMI providers and WMI client applications.

The executable file for WBEMTest, Wbemtest.exe, is installed in the WBEM directory of your Windows system directory. On Windows XP systems, this is the %WINDIR%\system32\Wbem folder.

For information about using the command-line interface to WMI, see Windows Management Instrumentation Command-line.

For information about using the WMI Control to configure WMI settings and security on a remote or local computer, see WMI Control Overview.

For information about WMI security, see Managing WMI security.

# Using WBEMTest user interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using WBEMTest user interface**
Two dialog boxes make up the primary interface to Windows Management Instrumentation (WMI) Tester, also called WBEMTest, functions.

- The **Windows Management Instrumentation Tester** dialog box provides access to all WBEMTest functions. For information on how to display this dialog box, see Display WMI tester dialog box.
- The **Connect** dialog box enables you to log on to WMI and connect to a namespace.

For more information about other WBEMTest dialog boxes, see:

- Connect dialog box
- Object Editor dialog box
    - Object Editor - Qualifier, Property, and Method functions
    - Object Editor dialog box - other functions
- Get Object Path dialog box
- Get Class Name dialog box
- Superclass Info dialog box
- Query dialog box
- Query Result dialog box
- Class Info dialog box
- Execute Method dialog box
- Property Editor dialog box
- Qualifier Editor dialog box
- Method Editor dialog box
- Query Result dialog box for refresher
- Get Object For Refresher dialog box
- MOF dialog box

For information about WMI security, see Managing WMI security.

# Connect dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connect dialog box**

Before you can use the functions of WBEMTest, you must log on to WMI. Click **Connect** on the **Windows Management Instrumentation Tester** dialog box to log on to WMI. The **Connect** dialog box appears.

Typically, you can accept the default settings displayed in the **Connect** dialog box, described as follows.

- The **Namespace** boxes displays the default namespace for your computer. You can specify a different namespace on your computer or on another computer. For example, for access to the classes of the Win32 provider, you can specify **root\cimv2**. To access the cimv2 namespace on a remote computer named **myserver**, you can specify **\\myserver\root\cimv2**. To access remote computers, you must grant permission to read or write from the remote computer using the WMI Control application, Wbemcntl.exe.

- The **Connection:** boxes specify properties of the connection to WMI. The **Using:** box specifies that **IWbemLocator** is used for obtaining the initial namespace pointer to Windows Management on a particular host computer. The **Returning:** box specifies that a **IWbemServices** pointer is used for accessing Windows Management information. The **Completion** box specifies that processing is synchronous.

- The **Credentials** boxes enable you to specify a user name, password, and authority. You can log on to WMI using credentials different from those of the user who is currently logged on. If you leave these boxes blank, WMI uses the settings of the current Windows logged on user.

- The **Locale** box enables you to specify a different locale by entering a localization code. To use the current locale, leave this box blank. Otherwise, it it must contain the location from which you want information. For Microsoft locale identifiers, the format of the string is "MS_xxx", where *xxx* is a string in hexadecimal form that indicates the Locale Identifier (LCID). For example, American English would appear as "MS_409".

- The **How to interpret empty password** option buttons control what happens if you leave the password box blank. WMI processes a blank password differently from no password at all. The default value for this box is NULL and causes WBEMTest to use the current Windows logged on user password to connect to WMI. To log on to WMI using an empty password, select the **Blank** option button.

- The **Impersonation level** option buttons enable you to set the Distributed COM (DCOM) impersonation level that WBEMTest typically uses. In most cases, you can use the default value, **Impersonate**,for this option button. However, some providers might require specific impersonation settings. Consult the documentation for the provider you are using if you are not sure which value to use. Additionally, the impersonation level that you specify must be applicable to the operating system you use. For example, if you specify an impersonation level of **Delegate** and WBEMTest is not running on a computer that uses the Windows XP Home Edition or Windows XP Professional operating system, this setting is ignored.

  For more information about impersonation, see Impersonation Level Enumeration.

- The **Authentication level** option buttons enable you to specify the Distributed COM (DCOM) authentication level you want to use. These authentication levels control how information is authenticated by DCOM. Typically, you can use **Packet**, the default value for this option button.

For more information on WBEMTest, see COM API for WMI. Also see the Windows Management Instrumentation documentation at the Microsoft Platform SDK Web site

# Object Editor dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Object Editor dialog box**
The **Object Editor** dialog box appears when you double-click a class or instance from the **Query Result** dialog box. The **Object Editor** dialog box also appears when you click **Open Class** or **Enum Class** on the **Windows Management Instrumentation Tester** dialog box.

You can use the **Object Editor** dialog box to:

- Edit the qualifiers, properties, and methods for the selected class or instance.
- View the MOF code that creates the class or instance.
- View instances of the class, the derived classes, and the superclass of the class, if you are editing a class definition.
- Refresh the view of the object after you make changes.

For example, when you click **Open Class** in the **Windows Management Instrumentation Tester** dialog box, and then type **Win32_LogicalDisk** in the **Get Class Name** dialog box, the **Object Editor for Win32_LogicalDisk** dialog box appears.

If you double-click any item in a list box, the appropriate editor opens for that object. For example, if you double-click the **reset** method in the methods list box, the method editor for the **reset** method appears.

For more information on WBEMTest, see COM API for WMI. Also see the Windows Management Instrumentation documentation at the Microsoft Platform SDK Web site

# Object Editor - Qualifier, Property, and Method functions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Object Editor - Qualifier, Property, and Method functions**

In the **Object editor**, you can do the following:

- Open the **Qualifier Editor** dialog box by clicking **Add Qualifier**. You can use this feature to add qualifiers that apply to the class or instance as a whole as opposed to a specific property of the class or instance. Type in the values for the new qualifier and click **Save Qualifier** to create the new qualifier.
- Open the **Qualifier Editor** dialog box for the qualifier that is selected, by clicking **Edit Qualifier**. You can use this to change the value or flavors assigned to a qualifier.
- Delete the currently selected qualifier, by clicking **Delete Qualifier**.
- Not display system properties in the **Object Editor** dialog box, by clicking **Hide System Properties**.
- Define new properties and assign values to classes or instances, by clicking **Add Property**. When you click **Add Property**, the **Property Editor** dialog box appears.
- Change the definition for the currently selected property, by clicking **Edit Property**.
- Delete the currently selected property, by clicking **Delete Property**.
- Define the methods for a class or instance, by clicking **Add Method**. When you click **Add Method**, the **Method Editor** dialog box appears.
- Edit the methods for a class or instance, by clicking **Edit method**. When you click **Edit Method**, the **Method Editor** dialog box appears.
- Delete the currently selected method, by clicking **Delete Method**.

For more information on WBEMTest, see COM API for WMI. Also see the Windows Management Instrumentation documentation at the Microsoft Platform SDK Web site

# Object Editor dialog box - other functions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Object Editor dialog box - other functions**

In the **Object Editor** dialog box, you can do the following:

- Close the **Object Editor** dialog box, by clicking **Close**.
- Save the changes to the current class or instance, by clicking **Save Object**.
- Open the **MOF** dialog box, which contains the managed object format (MOF) code that you can use to construct the current object, by clicking **Show MOF**. It is important to remember that the MOF code shown in this dialog box might not be identical to the actual MOF code that created the class or instance. The original MOF file might have comments and different formatting from text shown, for more information. see Managed Object Format
- Run a schema query that shows the class from which the current class or instance is derived, by clicking **Superclass**.
- Run a schema query that shows the child classes of the current class or instance, by clicking **Derived**.
- Show instances of the current class, by clicking **Instances**.
- Control how the class are updated, by using the **Update type** option. The following tables describe the two groups of option buttons that control the update process.

| Update Type | Meaning |
|---|---|
| **Create only** | Used for creation only. The save operation fails if the class or instance already exists. |
| **Update only** | Causes this call to update. |
| **Either** | Causes the class or instance to be created if it does not exist, or overwritten if it exists already. |

| Update Type | Meaning |
|---|---|
| **Compatible** | Allows a class to be updated are no derived classes and instances for that class exist. It also allows updates in all cases if qualifiers are changed, such as the description qualifier. This is the default value for this option and is used for compatibility with previous versions of WMI. If the class has instances, the update fails. |
| **Safe** | Allows updates of classes even when there are child classes, providing the change does not cause any conflicts with child classes. An example of an update to this value is adding to the base class a new property that was not previously mentioned in any of the child classes. If the class has instances, the update fails. |
| **Force** | Forces updates of classes when conflicting child classes exist. For example, this value forces an update if a class qualifier was defined in a child class, and the base class tries to add the same qualifier in conflict with the existing one. In force mode, this conflict is resolved by deleting the conflicting qualifier in the child class. If the class has instances, the update fails. |

For more information on WBEMTest, see COM API for WMI. Also see the Windows Management Instrumentation documentation at the Microsoft Platform SDK Web site

# Get Object Path dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Get Object Path dialog box**

Use the **Get Object Path** dialog box to specify the name of a class or the relative path to a class instance.

Object paths have the format: [Class].[KeyProperty] = Value

For example, an instance of a class called Myclass that has a key property called Keyproperty with a value of "A" the object path is: MyClass.KeyProperty="A".

For more information, see WMI Object Path Requirements.

# Get Class Name dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Get Class Name dialog box**

Use the **Get Class Name** dialog box to specify class names for the WBEMTest Open Class function.

# Superclass Info dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Superclass Info dialog box**

The **Enum Classes** function uses the **Superclass Info** dialog box. If you select the **Recursive** option button, WBEMTest also enumerates all subclasses of the class you specify, instead of only the immediate subclasses of that class.

# Query dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Query dialog box**
Use the **Query** dialog box to compose a query, and then execute it. To display the **Query** dialog box, click **Query** on the **Windows Management Instrumentation Tester** dialog box.

In the **Enter Query** text box, type the WBEMTest query .

In the **Query Type** list box, you can select a query language. Currently, only Windows Management Instrumentation (WMI) Query Language (WQL) is supported.

You can use the **Retrieve class prototype** check box to indicate that you do not want instances from this query, and then to the class that the instances might belong to if returned.

For more information, see Querying WMI.

# Query Result dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Query Result dialog box**

The **Query Result** dialog box is displayed when you perform a function that enumerates classes or instances. This dialog box also appears for event notification queries to display events. For asynchronous event notification queries, this dialog box displays the events as received. New events are delivered when you run a semisynchronous event notification query.

If the query results are instances, the relative path for the instances is displayed. You can double-click any of the objects in the **Query Result** dialog box to display the **Object Editor** dialog box for that object.

For more information, see Object Editor dialog box.

# Class Info dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Class Info dialog box**

Use the **Class Info** dialog box to specify the name of the class for which you want to see instances. This dialog box appears when you click **Enum Instances** on the **Windows Management Instrumentation Tester** dialog box. If you select **Immediate Only**, the default value, only instances of the class you specify are displayed. If you select **Recursive**, instances of the subclasses of the class you specify are displayed.

For more information on WBEMTest, see COM API for WMI. Also see the Windows Management Instrumentation documentation at the Microsoft Platform SDK Web site

# Execute Method dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Execute Method dialog box**

In the **Execute Method** dialog box, you can do the following:

- Edit the object path, by clicking **Edit Object Path** and set the object path to refer to a specific instance if necessary.
- Show the current object path in the **ObjectPath** text box.
- Close this dialog box without executing the method, by clicking **Dismiss**.
- Display the methods available for this class in the **Method** drop-down list .
- Set the various parameters the method might require, by clicking **Edit In Parameters**.
- Reset all the input parameters to the default settings, by clicking **Clear In Parameters**.
- View the output parameters for the method, by clicking **Edit Out Parameters**. The output parameters contain data returned from a method.

**Notes**

- To display the **Execute Method** dialog box, click **Execute Method** on the **Windows Management Instrumentation Tester** dialog box.
- You can use the **Execute Method** dialog box to run a method of a class or class instance.

# Property Editor dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Property Editor dialog box**

Use the Property Editor dialog box to create or edit properties of classes. In some cases, for example, when you edit an existing class, you might not be able to set all the fields on this dialog box.

For more information, see the Windows Management Instrumentation documentation at the Microsoft Platform SDK Web site

# Qualifier Editor dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Qualifier Editor dialog box**

You can define the name, data type, option, and value of a qualifier in the **Qualifier Editor** dialog box.

You can assign five common qualifier options to your qualifier by selecting the check boxes in the option section. The following options are available:

- Propagates to instances
- Propagates to derived classes
- Allows overrides
- Propagated
- Amended

For more information, see WMI-Specific Qualifiers.

# Method Editor dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Method Editor dialog box**

In the **Method Editor** dialog box, you can do the following:

- Enable the **Edit Input Arguments** button by clicking the **Enable Input Args** option button.

  To open the **Object Editor** dialog box for an instance of the __Parameters class, click the **Edit Input Arguments** button. You can use the object editor to edit or define new input parameters for a method.

- Enable the **Edit Output Arguments** button by clicking the **Enable Output Args** option button.

  To display the **Object Editor** dialog box for an instance of the __Parameters class, click the **Edit Output Arguments** button. You can use the object editor to edit or define new output parameters for a method.

- Add, delete, and edit the qualifiers for this method in the **Qualifiers** section.

- Save new methods and update existing methods by clicking **Save Method**.

**Notes**

- To display the **Method Editor** dialog box, in the **Windows Management Instrumentation Tester** dialog box, click the appropriate **IWbemServices** class button. In the next dialog box, select a Method, and then click the appropriate **Method** button.

- You can use the **Method Editor** dialog box to add methods to classes and class instances, and to edit methods.

# Query Result dialog box for refresher

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Query Result dialog box for refresher**

In the **Query Result** dialog box for refresher, you can do any of the following:

- Display the **Get Object For Refresher** dialog box where you can add objects or enumerators to a refresher, by clicking **Add**. Type the name of a class that supports refreshers in the **Object Path** text box to add a refresher to the query.
- Delete objects or enumerators from a refresher by clicking **Delete**.
- Populate the instances with data, clicking **Refresh**.

**Notes**

- To display the **Query Result** dialog box for refresher, on the **Windows Management Instrumentation Tester** dialog box, click **Create Refresher**.
- The refreshers you create are empty until you click **Refresh**. Click **Refresh** as often as necessary to update the data in your refreshers.

For more information on WBEMTest, see COM API for WMI. Also see the Windows Management Instrumentation documentation at the Microsoft Platform SDK Web site

# Get Object For Refresher dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Get Object For Refresher dialog box**
In the **Get Object For Refresher** dialog box, you can do the following:

- Refresh an object, by typing the relative path to an instance and then clicking **Add Object**.
- Refresh an enumerator, by typing the class name of your refresher and then clicking **Add Enum**.

**Notes**

- To display the **Get Object For Refresher** dialog box, click **Add** in the **Query Result - REFRESHER** dialog box.
- To specify the object path to an instance you are adding to your refresher or to specify the Class name of the enumerator you want to add to your refresher, use the **Get Object For Refresher** dialog box.

For more information on WBEMTest, see COM API for WMI. Also see the Windows Management Instrumentation documentation at the Microsoft Platform SDK Web site.

# MOF dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MOF dialog box**

The **MOF** dialog box appears when you click **Show MOF** on the **Object Editor** dialog box. You can use the **MOF** dialog box to display Managed Object Format (MOF) code when editing a class or class instance in the **Object Editor** dialog box.

For information, see Object Editor dialog box.

# Microsoft Management Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Management Console**
Microsoft Management Console (MMC) hosts administrative tools that you can use to administer networks, computers, services, and other system components.

- For help with specific tasks, see MMC How To....
- For MMC background information, see MMC Concepts.
- For information about accessibility features, see Accessibility for MMC.
- For problem-solving instructions, see Troubleshooting MMC.
- For Help information for Windows Server 2003 family of operating systems, see Access Windows Help.
- For security suggestions, see MMC Best Practices and Runas.

# MMC Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Log on as a normal user**

- When you log on with administrative rights, you open the systems over which you have control to attack. For more information, see Why you should not run your computer as an administrator.
- Use the **Runas** command to open specific tools when you need to use them. For more information, see Runas.
- For tools that you use frequently, create shortcuts that use the **Runas** command.These shortcuts will help you work efficiently without compromising security. For more information, see Create a shortcut using the runas command.
- ***When you complete your task, close the tool***. This will eliminate an opportunity for possible security breach of system or computer access by someone using your administrative credentials.

**Follow the security guidelines for the tools that you use.**

- Check the Help that accompanies each tool for security guidelines.

For more security information, see Best practices for security, Security overview, and Local Users and Groups Best practices.

# MMC How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Access Windows Help
- Open MMC and saved console files
- Author an MMC console file
- Modify a saved MMC console file
- Work with MMC and saved console files
- Set policy for MMC and snap-ins

# Access Windows Help

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To access Windows Help**

1. Click **Start**, then click **Help and Support** to access the **Help and Support Center**.

## Notes

- Once Help and Support are open, you can navigate to all Help topics.
- Also see Viewing classic Help in Related Topics. Classic Help is Help in the format that shipped with Windows 2000. It is the same content that appears in **Help and Support Center** with the only difference in the format.

  Related Topics

  Viewing classic Help

  Windows Server 2003 Product Help

  Getting Started

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Open MMC and saved console files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Open MMC and saved console files**

- Open MMC
- Open a saved MMC console for local and remote computers
- Create a desktop shortcut to open MMC

# Open MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open MMC**

- Do one of the following:
  - Click **Start**, click **Run**, type **mmc**, and then click **OK**.
  - At the command prompt, type **mmc**, and then press ENTER.

The complete command-line syntax for MMC is:

mmc*path******filename.msc* [/a] [/64] [/32]

*path* **** *filename.msc*

This command starts MMC and opens a saved console. You need to specify the complete path and file name for the saved console file. If you do not specify a console file, MMC opens a new console.

You can use environment variables to create command lines or shortcuts that do not depend on the explicit location of console files. For instance, if the path to a console file is in the system folder (for example, **mmc***c:\windows\system32\console_name.msc*), you can use the expandable data string %systemroot% to specify the location (**mmc***%systemroot%\system32\console_name.msc*). This may be useful if you are delegating tasks to people in your organization who are working on different computers.

**/a**

This switch opens a saved console in author mode and is used to make changes to saved consoles. When consoles are opened with this option, they are opened in author mode, regardless of their default mode. This does not permanently change the default mode setting for files. When you omit this option, MMC opens console files according to their default mode settings.

**/64**

Opens the 64-bit version of MMC (MMC64). Use this option only if you are running a Microsoft 64-bit operating system. For more information, see Running 32-bit and 64-bit snap-ins in Windows Server 2003 and Windows XP Professional operating systems.

**/32**

Opens the 32-bit version of MMC (MMC32). When running a Microsoft 64-bit operating system, you can run 32-bit snap-ins by opening MMC with this command-line option. For more information, see Running 32-bit and 64-bit snap-ins in Windows Server 2003 and Windows XP Professional operating systems.

**Notes**

- After you open MMC or a console file in author mode, you can open any existing console by clicking **Open** on the **File** menu.
- You can use the command line to create shortcuts for opening MMC and saved consoles. A command-line command is used in the **Run** command on the **Start** menu, in any command-prompt window, in shortcuts, or in any batch file or program that calls the command.
- 32-bit is the default.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Add a snap-in to a new MMC console for a local computer
Open a saved MMC console for local and remote computers
Install a program managed by an MMC snap-in
Create a desktop shortcut to open MMC
Using files created with earlier versions of MMC with MMC 2.0

# Open a saved MMC console for local and remote computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open a saved MMC console for local and remote computers**
There are several ways to open a saved MMC console:

- Open the folder where your console is located, and then click the console.
- Click **Start**, click **Run**, type the name of the console file, and then click **OK**.
- You can also open console files by using Run as. With this feature, you can log on to your computer with user rights to perform routine tasks, but open console files to perform administrative tasks with Administrator rights as necessary. You can use this feature by typing **runas** at the command prompt or by right-clicking an .msc file, and then clicking **Run as**.
- If you are using Windows XP Professional or Windows XP Home Edition, click **Start**, click **Control Panel**, click**Administrative Tools**, and then click the console you want to open.
- If you are using Windows 2000, click **Start**, click **Settings**, click **Control Panel**, click**Administrative Tools**, and then click the console you want to open.

### Note

- By default, consoles are saved in the following folder: %systemdrive%\Documents and Settings*username*\Start Menu\programs\Administrative Tools.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save an MMC console file
Add a snap-in to a new MMC console for a local computer
Using files created with earlier versions of MMC with MMC 2.0
Using the runas command with MMC

# Create a desktop shortcut to open MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a desktop shortcut to open MMC**

1. Right-click an open area on the desktop, point to **New**, and then click **Shortcut**.

2. Follow the instructions on the screen. A shortcut will appear on the desktop.

## Notes

- You can drag a shortcut that you create on the desktop to other folders in the operating system, such as the folder where you keep your .msc files or the folder for the **Start** menu.

- You can also create a shortcut to MMC in folders by using Windows Explorer. To do so, click the folder in which you want the shortcut, and on the **File** menu, point to **New**, click **Shortcut**, and then follow the instructions on the screen.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open MMC
Using files created with earlier versions of MMC with MMC 2.0

# Author an MMC console file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Author an MMC console file**


- Add a snap-in to a new MMC console for a local computer
- Add a published snap-in to a new MMC console
- Add an extension snap-in to an MMC console
- Add a published extension snap-in to an MMC console
- Create a taskpad view in a saved MMC console
- Create tasks for a taskpad view in saved MMC console
- Edit a taskpad view in a saved MMC console
- Edit tasks in a taskpad view in a saved MMC console
- Add an item to the Favorites list in an MMC console
- Organize the Favorites list in an MMC console
- Install a program managed by an MMC snap-in
- Set MMC console options
- Save an MMC console file

# Add a snap-in to a new MMC console for a local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a snap-in to a new MMC console for a local computer**

1. Open MMC.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Snap-in**, double-click the item you want to add, and, if prompted, do one of the following:
   - Click **Local computer: (the computer this console is running on)**, and then click **Finish**.
   - If a wizard appears, follow the instructions.
4. To add another item to the console, repeat step 3.

**Notes**

- To open MMC, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- If you add a snap-in to a console, and the snap-in becomes damaged, try adding another instance of the same snap-in to the console. Configure it as needed, and then remove the first instance of the snap-in.
- If a console is saved to the per-user Administrative Tools folder (located at *systemdrive*\Documents and Settings*user*\Start Menu\Programs\Administrative Tools), it becomes available in the Administrative Tools folder on the **Programs** menu for that user.
- You can only access snap-ins published in Active Directory.
- A remote computer is entirely snap-in specific.
- If a snap-in does not appear in the list, you must first install the program, device, or service administered by the snap-in. You can access snap-ins not locally installed from the **Add/Remove Snap-in** dialog box in the Active Directory directory service. You must be using one of the following operating systems for this process to work:
  - Is running Windows 98 with Windows Installer installed;

    –or–
  - Is running Windows NT 4.0 with Windows Installer installed;

    or
  - Is running Windows Millennium Edition;

    –and–
  - Is part of a Windows domain.
  - For more information, see Help.

- To make an item other than the console root subordinate to an item in the console tree, click the appropriate item in **Snap-ins added to** before you click **Add** in step 2.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save an MMC console file
Install a program managed by an MMC snap-in
Add a published snap-in to a new MMC console
Using files created with earlier versions of MMC with MMC 2.0

# Add a published snap-in to a new MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a published snap-in to a new MMC console**

1. Open MMC.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click the published snap-in you want to add, and then click **Add**.
4. If a wizard for Windows Installer appears, follow the instructions, and then do one of the following:
   - Click **Local computer: (the computer this console is running on)**, and then click **Finish**.
   - If a wizard for the snap-in appears, follow the instructions.
5. To add another item to the console, repeat step 3.

## Notes

- To open MMC, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- For snap-ins available from the Active Directory directory service, **Not Installed** displays in the **Vendor** column.
- You can only access extension snap-ins published in Active Directory if the computer you are using is running Windows 98 with Windows Installer, Windows NT 4.0 with Windows Installer, or Windows Millinnium Edition. In addition, the computer must be part of a Windows domain. For more information about publishing applications and Windows Installer, see Help.
- If a console is saved to the per-user Administrative Tools folder (located at *systemroot*\Documents and Settings*User*\Start Menu\Programs\Administrative Tools), it is then available in the Administrative Tools folder on the **Programs** menu for that user.
- To make an item subordinate to an item in the console tree other than the console root, click the appropriate item in **Snap-ins added to** before you click **Add** in step 2.
- For more information on Windows Installer, go to MSDN Online Downloads.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save an MMC console file
Add a published extension snap-in to an MMC console

# Add an extension snap-in to an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an extension snap-in to an MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc path***filename***.msc/a**, and then click **OK**.

2. On the **File** menu, click **Add/Remove Snap-in**, and then click the item that you want to extend.

3. On the **Extensions** tab, in **Available extensions**, select the check box next to the extension you want to add, and then click **OK**.

## Notes

- To add every extension available for a given item, on the **Extensions** tab, select the **Add all extensions** check box. To add only certain extensions, clear this check box.
- You can also use this procedure to remove extensions from a console by clearing the check box for an extension that was previously selected.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a published extension snap-in to an MMC console
Using files created with earlier versions of MMC with MMC 2.0
Open a saved MMC console for local and remote computers

# Add a published extension snap-in to an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a published extension snap-in to an MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. On the **File** menu, click **Add/Remove Snap-in** .
3. On the **Extensions** tab, in **Snap-ins that can be extended**, click the snap-in that you want to extend.
4. In **Available extensions**, select the check box next to the published extension you want to add, and then click **Download**.
5. If one or more wizards appear, follow the instructions on your screen, and then click **OK**.

### Important

- If you select the **Add all extensions** check box on the **Extensions** tab, only extension snap-ins that are locally installed and registered are added to the console. You must specifically download extension snap-ins from the directory service to make them available in a console. Thus, to add all published extensions for a given item, do not select the **Add all extensions** check box, but instead select the check box next to each published extension.

### Notes

- For extension snap-ins that are available from the Active Directory directory service, the snap-in names are followed by the phrase **(not installed)**.
- You can only access extension snap-ins published in Active Directory if the computer you are using is running Windows 98 with Windows Installer, Windows NT 4.0 with Windows Installer, or Windows Millennium Edition. In addition, the computer must be part of a Windows domain. For more information about publishing applications and Windows Installer, see Help.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a published snap-in to a new MMC console
Viewing extensions to snap-ins
Open a saved MMC console for local and remote computers

# Create a taskpad view in a saved MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a taskpad view in a saved MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc path***filename***.msc/a**, and then click **OK**.

2. In the console tree, click a snap-in item.
3. On the Action menu, click **New Taskpad View**.
4. Follow the instructions in the New Taskpad View wizard.
5. If you want to create tasks immediately after you create the taskpad view, select the **Start New Task wizard** check box in the final screen of the wizard.

**Note**

- You can also create a taskpad view in a new console, but you must first add a snap-in to the console.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Taskpad views and tasks
Create tasks for a taskpad view in saved MMC console
Add a snap-in to a new MMC console for a local computer
Open a saved MMC console for local and remote computers

# Create tasks for a taskpad view in saved MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create tasks for a taskpad view in a saved MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. In the console tree, click an item associated with the taskpad view; and, on the **Action** menu, click **Edit Taskpad View**.
3. On the **Tasks** tab, click **New**.
4. Follow the instructions in the New Task wizard.

## Notes

- If you are creating a new taskpad view and want to add a task, you can start the New Task wizard by selecting the **Start New Task wizard** check box in the final screen of the New Taskpad View wizard.
- On the **Shortcut Menu Command** page of the New Task wizard, in **Command source**, if you selected **Tree item** task and cannot view all items in the console tree, you must expand the console tree manually before starting the New Task wizard.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Taskpad views and tasks
Create a taskpad view in a saved MMC console
Edit tasks in a taskpad view in a saved MMC console
Open a saved MMC console for local and remote computers

# Edit a taskpad view in a saved MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit a taskpad view in a saved MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. In the console tree, click an item associated with the taskpad view. On the **Action** menu, click **Edit Taskpad View**.

3. On the **General** tab, do the following:

   - To change the name for the taskpad view, in **Name**, type a new name. The name appears in the upper left corner of the taskpad view and on the tab on the bottom edge of the taskpad view.
   - To change the description for a taskpad view, in **Description**, type a new description. The description appears under the name in the taskpad view.
   - To change the list format, or to configure the taskpad view to contain tasks only, under **Style for the details pane,** click a new list format, or click **No list**.
   - To remove the **Standard** tab from view, check **Hide Standard tab**.
   - To configure task descriptions, under **Style for task descriptions**, click **Text** or **InfoTip**. Text descriptions are displayed to the right of a task. InfoTip descriptions are displayed as pop-up windows when a user points to tasks with the mouse.
   - To change the size of the list, in **List size**, click a new size for the width (Vertical list) or height (Horizontal list) of the list.

4. Click **Options**, and then do one of the following:

   - To associate the taskpad view that shows only the currently selected item in the console tree, click **Selected tree item**.
   - To associate the taskpad view that shows all items in the console tree of the same type as the currently selected item, click **All tree items that are the same type as the selected tree item**.

   **Note**

   - If you click this option and want the taskpad view to appear by default, select the **Change default display to this taskpad view for these tree items** check box, instead of the view on the **Standard** tab.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Taskpad views and tasks
Create a taskpad view in a saved MMC console
Create tasks for a taskpad view in saved MMC console
Open a saved MMC console for local and remote computers

# Edit tasks in a taskpad view in a saved MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit tasks in a taskpad view in a saved MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. In the console tree, click an item associated with the taskpad view, and on the **Action** menu, click **Edit Taskpad View**.

3. On the **Tasks** tab, under **Display these tasks**, click a task, and then do any of the following:

   - To move the selected task up the display list in the taskpad view, click **Move Up**.
   - To move the selected task down the display list in the taskpad view, click **Move Down**.
   - To delete the task from the display list in the taskpad view, click **Remove**.

4. On the **Tasks** tab, under **Display these tasks**, click **Modify**, and then do any of the following:

   - To change the name of a task, on the **General** tab, in **Task name**, type a new name.
   - To change the description for a task, on the **General** tab, in **Description**, type a new description. Depending on how you configured the taskpad view, the result appears either to the right of the task or on the screen when a user points to the task with the mouse.
   - To change the icon for the task, on the **Task Icon** tab, click a new icon, and then click **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Taskpad views and tasks
Create a taskpad view in a saved MMC console
Create tasks for a taskpad view in saved MMC console
Open a saved MMC console for local and remote computers

# Add an item to the Favorites list in an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an item to the Favorites list in an MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. In the console tree, click the item that you want to add to the Favorites list.
3. In the details pane, click either the **Standard** tab or the tab for the taskpad view you want to add.
4. On the **Favorites** menu, click **Add to Favorites**.
5. In **Create in**, do one of the following:

   - To add an item to an existing folder, click the folder in which you want the item to reside, and then click **OK**.
   - To add an item to a new folder, click the folder you want as the parent, and then click **New Folder**. In **Folder name**, type a name, click **OK**, and then click **OK** again.

**Note**

- The **Standard** tab and tabs for taskpad views are only available in the details pane of a console if you click an item in the console tree that is configured to have a taskpad view. Otherwise, no tabs are visible in the details pane.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Organize the Favorites list in an MMC console
Open a saved MMC console for local and remote computers

# Organize the Favorites list in an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To organize the Favorites list in an MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. On the **Favorites** menu, click **Organize Favorites**, and then do the following:

   - To add a folder, click **Create Folder**, type a name in **Folder name**, and then click **OK**.
   - To move an item to a folder, click the item under the Favorites folder, click **Move to Folder**, click the folder where you want the item to appear, and then click **OK**.
   - To change the name of an item, click the item under the Favorites folder, click **Rename**, type a new name, and then press ENTER.
   - To remove an item, click the item under the Favorites folder, and then click **Delete**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add an item to the Favorites list in an MMC console

# Install a program managed by an MMC snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a program managed by an MMC snap-in**

1. Click **Start**.
2. Click **Control Panel**.
3. Click **Add or Remove programs**.
4. Follow the instructions on the screen.

If you are using Windows 2000, do the following:

1. Click **Start**.
2. Click **Settings**.
3. Click **Control Panel**.
4. Click **Add or Remove programs**.
5. Follow instructions on the screen.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Set MMC console options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set MMC console options**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. On the **File** menu, click **Options**.
3. To change the title for the console, type a new title in the box.
4. To change the icon for the console, click **Change Icon**. In **File Name**, type the path to a file containing icons.
5. Under **Current icon** click an icon, and then click **OK**.
6. To change the default mode for the console, in **Console mode**, click the mode that you want the console to open in.
7. If the default console mode for the console is one of the user modes, do the following:

   - To prevent users from editing the console, select the **Do not save changes to this console** check box.
   - To enable users to access the **Customize View** dialog box, select the **Allow the user to customize views** check box.

## Notes

- If you change the title of a console, the title bar does not display the path of items clicked in the console tree. However, if you do not change the title, or if you delete the title, the paths are displayed.
- Many files contain icons that can be used for custom consoles. For example, you can find more icons in Shell32.dll, located in the *%systemroot%*\System32 folder.
- To stop receiving the message **Save console settings to** *console name***?**, which appears each time you close a console opened in author mode, set the default mode for the console to one of the user modes and then clear the **Do not save changes to this console** check box. Any changes to the console are saved by default.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Open a saved MMC console for local and remote computers
Install a program managed by an MMC snap-in
Using files created with earlier versions of MMC with MMC 2.0

# Save an MMC console file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save an MMC console file**

- In an MMC console opened in author mode, on the **File** menu, click **Save**.

## Notes

- To open a saved console in author mode, right-click the .msc file and then click **Author** or click **Start**, then click **Run**, and type **mmc path**_filename_**.msc /a**.

- If the console is not opened in author mode, the **Save** option on the **File** menu is not available. In this case, saving is determined by whether the Do not save changes to this console check box was selected when the console was configured. This check box is available by clicking **Options** on the **File** menu. If this check box was not selected, changes to the console are automatically saved when you close MMC. If it was selected, changes to the console are discarded when you close MMC.

- If a console is saved to the per-user Administrative Tools folder (located at %systemdrive%\Documents and Settings\user\Start Menu\Programs\Administrative Tools), it is then available in the Administrative Tools folder on the **Programs** menu for that user.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Install a program managed by an MMC snap-in
Open MMC
Using files created with earlier versions of MMC with MMC 2.0
Saving consoles

# Modify a saved MMC console file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modify a saved MMC console file**

- Rename an item on the console tree of an MMC console
- Remove a snap-in or other item from an MMC console
- Make a new console window from an MMC console
- Make a new console window from part of an MMC console tree
- Open a saved MMC console for local and remote computers

# Rename an item on the console tree of an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename an item on the console tree of an MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. In the console tree, click the item that you want to rename.

3. On the **Action** menu, click **Rename**, and then type the new name.

### Note

- Not all items in the console tree can be renamed. If the **Rename** command does not appear on the **Action** menu, you cannot rename the item.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Open a saved MMC console for local and remote computers
Install a program managed by an MMC snap-in

# Remove a snap-in or other item from an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a snap-in or other item from an MMC console**

1. Open a saved console in author mode by doing one of the following:

    - Right-click the .msc file, and then click **Author**.
    - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. On the **File** menu, click **Add/Remove Snap-in**.

3. On the **Standalone** tab, click the item you want to remove, and then click **Remove**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Open a saved MMC console for local and remote computers
Install a program managed by an MMC snap-in
Using files created with earlier versions of MMC with MMC 2.0

# Make a new console window from an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a new console window from an MMC console**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. On the **Window** menu, click **New Window**.

### Note

- The new console window is rooted at the console root of the saved console.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Install a program managed by an MMC snap-in
Open a saved MMC console for local and remote computers
Make a new console window from part of an MMC console tree
Using files created with earlier versions of MMC with MMC 2.0

# Make a new console window from part of an MMC console tree

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a new console window from part of an MMC console tree**

1. Open a saved console in author mode by doing one of the following:

   - Right-click the .msc file, and then click **Author**.
   - Click **Start**, click **Run**, type **mmc** *path**filename.msc* **/a**, and then click **OK**.

2. In the console tree, click the item that you want to become the console root of the new console.

3. On the **Action** menu, click **New Window from Here**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Install a program managed by an MMC snap-in
Open a saved MMC console for local and remote computers
Make a new console window from an MMC console

# Work with MMC and saved console files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with MMC and saved console files**


- Viewing extensions to snap-ins
- Using disk cleanup tab
- Hide or display features of a saved MMC console
- Hide or display menus and toolbars for snap-ins
- Reorder columns in an MMC console
- Hide or display columns in an MMC console
- Filter rows in an MMC console
- Export columns in an MMC console to a text file
- View the associated snap-in and item for a property page

# Viewing extensions to snap-ins

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view extensions to snap-ins**

1. Click **Start**, go to **Additional Programs** or **All Programs**, go to **Administrative Tools**, and then click the console file name.
2. In the **Add/Remove Snap-in** dialog box, go to **Snap-ins added to:** and select the snap-in you want to use.
3. Click the **Extensions** tab.
4. Extensions will be displayed in the **Available extensions** box.
5. Extensions from other snap-ins can be displayed from selections in the **Snap-ins that can be extended:** item list.

## Notes

- You must be in author mode to view the **Add/Remove snap-ins** menu.
- By default, user added snap-ins are saved in Administrative Tools (**%SystemDrive%/Documents and Settings/%UserName%/Start/Programs/Administrative Tools**). They may be saved elsewhere at the user's discretion.
- To open the console file in author mode, right-click the console file name and select **Author**.
- In the opened MMC console, click **File**. From the drop-down menu, select **Add/Remove Snap-in**.
- The **Add/Remove snap-in** menu is not available in any of the user modules.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a published extension snap-in to an MMC console
Add a published snap-in to a new MMC console
View the associated snap-in and item for a property page

# Using disk cleanup tab

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To use the disk cleanup tab

1. Open MMC.
2. Click **File** on the menu bar.
3. Click **Options** on the drop-down menu. The **Options** dialog box opens, displaying two tabs.
4. Click the **Disk Cleanup** tab.
5. To delete the console files, click **Delete Files**.
6. Close the dialog box by clicking **OK**.
7. To cancel, click **Cancel** at the bottom of the dialog box.

### Notes

- To open MMC, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- The text on the **Disk Cleanup** tab states that changes to the view of a console file are saved in your profile. The amount of disk space used on the computer for these console changes is indicated on the **Disk Cleanup** tab.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Hide or display features of a saved MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To hide or display features of a saved MMC console**

1. Open a saved console.
2. On the **View** menu, click **Customize**.
3. Under **MMC**, do the following:
   - To display or hide the console tree, select or clear the **Console tree** check box.
   - To display or hide the **Action** and **View** menus, select or clear the **Standard menus (Action and View)** check box.
   - To display or hide the console toolbar, select or clear the **Standard toolbar** check box.
   - To display or hide the status bar at the bottom of the console window, select or clear the **Status bar** check box.
   - To display or hide the description bar along the top of the details pane, select or clear **Description bar**.
   - To display or hide the tabs along the bottom of the details pane, select or clear the **Taskpad navigation tabs** check box.

### Important

- If you clear the **Standard menus (Action and View)** check box and close the **Customize View** dialog box, you cannot access commands on the **Action** and **View** menus, including the **Customize** command. In this situation, to access the **Customize View** dialog box, click the system menu (the icon in the upper-left corner of the console), and then click **Customize View**.

### Notes

- To open a saved console in author mode, right-click the .msc file and then click **Author**.
- Click **Start**, click **Run**, and then type **mmc** *path**filename.msc* **/a**.
- As you select or clear check boxes, view the console to confirm that the changes you make are what you expected.
- If you hide the **Console tree**, the **Favorites** menu, if applicable, will still be visible on the menu bar.
- Displaying or hiding the tabs along the bottom of the details pane can also be done from any pre-existing taskpad view. Select **Action** from the menu, click **Edit taskpad view**, select or deselect the **Hide Standard tab** check box, and then click **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Add a snap-in to a new MMC console for a local computer
Open a saved MMC console for local and remote computers
Install a program managed by an MMC snap-in
Create a desktop shortcut to open MMC

# Hide or display menus and toolbars for snap-ins

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To hide or display menus and toolbars for snap-ins**

1. Open a saved console.
2. On the **View** menu, click **Customize**.
3. Under **Snap-in**, do one or more of the following:
   - To display or hide menus specific to snap-ins, select or clear the **Menus** check box.
   - To display or hide toolbars specific to snap-ins, select or clear the **Toolbars** check box.

## Notes

- To open a saved console in author mode, right-click the .msc file and then click **Author**. Click **Start**, then click **Run**, and type **mmc** *path**filename.msc* **/a**.
- When you select or clear the **Menus** and **Toolbars** check boxes, the menus and toolbars are displayed or hidden for all snap-ins in the console, not just the currently selected snap-in. If toggling the check boxes does not change the view, the currently selected snap-in does not have custom menus or toolbars.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Open a saved MMC console for local and remote computers
Install a program managed by an MMC snap-in
Create a desktop shortcut to open MMC

# Reorder columns in an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reorder columns in an MMC console**

1. In an open MMC console, click an item in the console tree that displays columns in the details pane.
2. On the **View** menu, click **Add/Remove Columns**.
3. In the **Add/Remove Columns** dialog box, under **Displayed columns**, click a column name, and then click **Move Up** or **Move Down** to change the position of the column.

## Notes

- You can also reorder columns in the details pane by using a mouse to drag a column heading to the left or right of its original position. As you drag a column, highlighting between the column headings indicates the new position of the column.
- You can resize columns by using the mouse to drag column headings.
- This feature is not enabled for all items. If you do not see the **Add/Remove Columns** option on the **View** menu, you cannot use this feature for the selected item.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open a saved MMC console for local and remote computers
Hide or display columns in an MMC console
Export columns in an MMC console to a text file
Filter rows in an MMC console

# Hide or display columns in an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To hide or display columns in an MMC console**

1. In an open MMC console, click an item in the console tree that displays columns in the details pane.
2. On the **View** menu, click **Add/Remove Columns**.
3. In the **Add/Remove Columns** dialog box, do the following:
   - To hide a column, in **Displayed columns**, click the column you want to hide, and then click **Remove**.
   - To display a column, in **Available columns**, click the column you want to display, and then click **Add**.

## Notes

- You might not be able to hide some columns in the details pane.
- The hide or display columns feature is not enabled for all items. If you do not see the **Choose Columns** option on the **View** menu, you cannot use this feature for the selected item.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open a saved MMC console for local and remote computers
Reorder columns in an MMC console
Export columns in an MMC console to a text file
Filter rows in an MMC console

# Filter rows in an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To filter rows in an MMC console**

1. In an open MMC console, click an item in the console tree that displays columns in the details pane.

2. On the **View** menu, click **Filter**. Modify the options on the **Filter** tab of the **Application Properties** dialog box as desired.

   Filters appear below the column headings in the details pane.

### Note

- This feature is not enabled for all items. If you do not see the **Filter** option on the **View** menu, you cannot use this feature for the selected item.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open a saved MMC console for local and remote computers

# Export columns in an MMC console to a text file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export columns in an MMC console to a text file**

1. In an open MMC console, click an item in the console tree that displays columns in the details pane.
2. On the **Action** menu, click **Export List**.
3. In the **Export List** dialog box, enter the following information:
   - In **Save in**, click a location to save the file.
   - In **File name**, type a name for the file.
   - In **Save as type**, click a file format in the list.
   - To export only certain rows, select the **Save Only Selected Rows** check box.
4. Click **Save**.

## Notes

- To export only certain rows in the details pane, select the rows.
- If you want to export the entire contents of the details pane, skip the preceding step.
- Some snap-ins do not provide the ability to select more than one row in the details pane. For these snap-ins, you can export all the rows in the details pane, but not selected rows only.
- Unicode file formats are not available in **Save as type** for operating systems earlier than Windows NT 4.0.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open a saved MMC console for local and remote computers
Hide or display columns in an MMC console
Reorder columns in an MMC console

# View the associated snap-in and item for a property page

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the associated snap-in and item for a property page**

1. In the console tree or details pane of an open MMC console, right-click the item for which you want to view properties and then click **Properties**.

2. In the properties dialog box, do the following:

    - To view which snap-in provided a particular property page in the dialog box, hold down the CTRL key and point to the label of one of the tabs.

    - To view the location in the console tree for the item associated with the properties dialog box, hold down the CTRL key and point to the title bar of the properties dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Set policy for MMC and snap-ins

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set policy for MMC and snap-ins**

- Restrict access to author mode in MMC
- Restrict access to author mode in MMC for a domain
- Restrict access to a permitted list of snap-ins
- Restrict access to a permitted list of snap-ins for a domain
- Permit or restrict access to a snap-in
- Permit or restrict access to a snap-in for a domain

# Restrict access to author mode in MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restrict access to author mode in MMC**

1. Open MMC.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Snap-in**, click **Group Policy Object Editor**, and then click **Add**.
4. In the **Select Group Policy** wizard, do one of the following:
   - To edit the local Group Policy Object, leave the **Group Policy Object** box set to **Local Computer**, and then click **Finish**.
   - To specify a different Group Policy Object, click **Browse**, select the object, click **OK**, and then click **Finish**.
5. Click **Finish**, click **Close**, and then click **OK**.

   The Group Policy Object Editor snap-in opens the specified Group Policy Object Editor.
6. In the console tree, click **Microsoft Management Console**.

   **Where?**

   - *PolicyName* Policy/User Configuration/Administrative Templates/Windows Components/Microsoft Management Console
7. In the details pane, double-click **Restrict the user from entering author mode**.
8. On the **Setting** tab, do one of the following:
   - To allow the user to use author mode in MMC, click **Not Configured** or **Disabled**.
   - To restrict the user from using author mode in MMC, click **Enabled**.

**Notes**

- To open MMC, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- Once **Restrict the user from entering author mode** has been enabled, any attempt to reaccess the console results in an error message indicating blocked access to the console in author mode. Users with administrator credentials can regain access to author mode by clicking **Start/Run/gpedit.msc** which opens the Group Policy Object Editor for local computer policy. Next, click **User Configuration/Administrative Templates/Windows Components/Microsoft Management Console.** Then double-click **Restrict the user from entering author mode** in the details pane. Unclick the radio button **Enabled** so that either **Not Configured or Disabled** has been selected.
- This feature applies only in a network where Group Policy Object Editor has been configured; a system administrator must enable Group Policy Object Editor. Not available in Windows XP Home Edition.

- For more information, click the **Explain** tab in the **Restrict the user from entering author mode properties** dialog box, and see Help.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Setting Group Policy in MMC
Group Policy (pre-GPMC)

# Restrict access to author mode in MMC for a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restrict access to author mode in MMC for a domain**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the organizational unit for which you want to configure policy, and then click **Properties**.
3. On the **Group Policy** tab, click **Edit** on the property page.
4. In the console tree, click **Microsoft Management Console**.

   **Where?**

   - *PolicyName* Policy/User Configuration/Administrative Templates/Windows Components/Microsoft Management Console

5. In the details pane, double-click **Restrict the user from entering author mode**.
6. On the **Policy** tab, do one of the following:

   - To allow the user to use author mode in MMC, click **Not Configured** or **Disabled**.
   - To restrict the user from using author mode in MMC, click **Enabled**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Active Directory applies only in a network where Group Policy Object Editor has been configured. You must be a domain Administrator, or have equivalent rights, and use a computer configured as a domain controller to configure Group Policy Object Editor for a domain.
- For more information, click the **Explain** tab in the **Restrict the user from entering author mode properties** dialog box, and see Help.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Setting Group Policy in MMC
Group Policy (pre-GPMC)

# Restrict access to a permitted list of snap-ins

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restrict access to a permitted list of snap-ins**

1. Open MMC.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Snap-in**, click **Group Policy Object Editor**, and then click **Add**.
4. In the **Select Group Policy** wizard, do one of the following:
   - To edit the local Group Policy Object, leave the **Group Policy Object** box set to **Local Computer**, and then click **Finish**.
   - To specify a different Group Policy Object, click **Browse**, select the object, click **OK**, and then click **Finish**.
5. Click **Finish**, click **Close**, and then click **OK**.

   The Group Policy Object Editor snap-in provides access to the specified Group Policy object.
6. In the console tree, click **Microsoft Management Console**.

   **Where?**

   - *PolicyName* Policy/User Configuration/Administrative Templates/Windows Components/Microsoft Management Console
7. In the details pane, double-click **Restrict users to the explicitly permitted list of snap-ins**.
8. On the **Policy** tab, do one of the following:
   - To permit the user to access snap-ins that are not explicitly restricted, click **Not Configured** or **Disabled**.
   - To restrict the user from accessing any snap-in that is not explicitly permitted, click **Enabled**.

**Notes**

- To open MMC, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- This feature is not available in Windows XP Home Edition.
- This feature applies only in a network where Group Policy object has been configured; a system administrator must enable Group Policy object.
- If you enable this policy, only permitted snap-ins appear in the list of available snap-ins in the **Add Standalone Snap-in** dialog box in MMC.
- For more information, click the **Explain** tab in the **Restrict users to the explicitly permitted list of snap-ins**

**properties** dialog box, and see Help.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer
Setting Group Policy in MMC
Permit or restrict access to a snap-in
Group Policy (pre-GPMC)

# Restrict access to a permitted list of snap-ins for a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restrict access to a permitted list of snap-ins for a domain**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the organizational unit for which you want to configure policy, and then click **Properties**.
3. On the **Group Policy** tab, click **Edit**.

   The Group Policy Object Editor appears.
4. In the console tree, click **Microsoft Management Console**.

   **Where?**

   - *PolicyName* Policy/User Configuration/Administrative Templates/Windows Components/Microsoft Management Console
5. In the details pane, double-click **Restrict users to the explicitly permitted list of snap-ins**.
6. On the **Policy** tab, do one of the following:

   - To permit the user to access snap-ins that are not explicitly restricted, click **Not Configured** or **Disabled**.
   - To restrict the user from accessing any snap-in that is not explicitly permitted, click **Enabled**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Active Directory applies only in a network where Group Policy Object Editor has been configured. You must be a domain Administrator, or have administrative credentials, and use a computer configured as a domain controller to configure Group Policy Object Editor for a domain.
- If you enable this policy, only permitted snap-ins appear in the list of available snap-ins in the **Add Standalone Snap-in** dialog box in MMC.
- For more information, click the **Explain** tab in the **Restrict users to the explicitly permitted list of snap-ins properties** dialog box, and see Help.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a snap-in to a new MMC console for a local computer

Setting Group Policy in MMC
Permit or restrict access to a snap-in for a domain
Group Policy (pre-GPMC)

# Permit or restrict access to a snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To permit or restrict access to a snap-in**

1. Open MMC.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Snap-in**, click **Group Policy Object Editor**, and then click **Add**.
4. In the **Select Group Policy** wizard, do one of the following:
   - To edit the local Group Policy Object, leave the **Group Policy Object** box set to **Local Computer**, and then click **Finish**.
   - To specify a different Group Policy Object, click **Browse**, select the object, click **OK**, and then click **Finish**.
5. Click **Finish**, click **Close**, and then click **OK**.
6. In the console tree, click **Restricted/Permitted snap-ins**.

   **Where?**

   - *PolicyName* Policy/User Configuration/Administrative Templates/Microsoft Management Console/Restricted/Permitted snap-ins
7. In the details pane, double-click the snap-in that you want to permit or restrict, and then do one of the following:
   - To enable the user to access this snap-in (unless the user is restricted by the **Restrict users to the explicitly permitted list of snap-ins** policy), click **Not Configured**.
   - To permit the user to access this snap-in, click **Enabled**.
   - To restrict the user from accessing this snap-in, click **Disabled**.

## Notes

- To open MMC, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- This feature applies only in a network where Group Policy object has been configured; a system administrator must enable Group Policy object.
- For more information, click the **Explain** tab in the *snap-in* **Properties** dialog box, and see Help.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Add a snap-in to a new MMC console for a local computer
Setting Group Policy in MMC
Restrict access to a permitted list of snap-ins
Permit or restrict access to a snap-in for a domain
Group Policy (pre-GPMC)

# Permit or restrict access to a snap-in for a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To permit or restrict access to a snap-in for a domain**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the organizational unit for which you want to configure policy, and then click **properties**.
3. On the **Group Policy** tab, click **Edit**.

   The Group Policy Object Editor appears.
4. In the console tree, click **Restricted/Permitted snap-ins**.

   **Where?**

   - *PolicyName* Policy/User Configuration/Administrative Templates/Windows Components/Microsoft Management Console/Restricted/Permitted snap-ins
5. In the details pane, double-click the snap-in that you want to permit or restrict, and then do one of the following:

   - To enable the user to access this snap-in (unless the user is restricted by the **Restrict users to the explicitly permitted list of snap-ins** policy), click **Not Configured**.
   - To permit the user to access this snap-in, click **Enabled**.
   - To restrict the user from accessing this snap-in, click **Disabled**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, point to **programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Active Directory is not available in Windows XP Professional or Windows XP Home Edition.
- This feature applies only in a network where Group Policy Object Editor has been configured. You must be a domain Administrator, or have administrative credentials, and use a computer configured as a domain controller to configure Group Policy Object Editor for a domain.
- When you restrict or explicitly permit access to a snap-in, the snap-in is added to a list of restricted or permitted snap-ins. The restricted list takes precedence over the permitted list, so that if the same snap-in exists on both lists, access to the snap-in is restricted.
- For more information, click the **Explain** tab in the *snap-in* **properties** dialog box, and see Help.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Add a snap-in to a new MMC console for a local computer
Setting Group Policy in MMC
Restrict access to a permitted list of snap-ins for a domain
Group Policy (pre-GPMC)

# MMC Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

- MMC overview
- MMC security
- Understanding MMC
- Using MMC
- Running 32-bit and 64-bit snap-ins in Windows Server 2003 and Windows XP Professional operating systems
- MMC command-line options
- MMC Resources

# MMC overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MMC overview**

You can use Microsoft Management Console (MMC) to create, save, and open administrative tools (called MMC consoles) that manage the hardware, software, and network components of your Windows system. MMC can be run on the various Windows 9x and Windows NT operating systems, as well as Windows XP Home Edition, Windows XP Professional, and the Windows Server 2003 family of operating systems.

MMC does not perform administrative functions, but hosts tools that do. The primary type of tool you can add to a console is called a snap-in. Other items that you can add include ActiveX controls, links to Web pages, folders, taskpad views, and tasks.

There are two ways that you can use MMC: in user mode, working with existing MMC consoles to administer a system, or in author mode, creating new consoles or modifying existing MMC consoles. For more information about the differences between user and author mode, see Console access options.

Keep security in mind at all times. It is good practice for administrators to use an account with restrictive permissions to perform routine, non-administrative tasks, and to use an account with administrative credentials only when performing specific administrative tasks. To accomplish this without logging off and back on, log on with a regular user account, and then use the **Runas** command to run the tools that require the administrative credentials. For detailed information, see Runas.

For detailed information about all security options, see Best practices for security, Security overview, and Local Users and Groups Best practices.

For Help with other aspects of the Windows Server 2003 family of operating systems, click **Start**, then **Help and Support** to access the **Help and Support Center**.

# MMC security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MMC security**
Keep security fully in mind at all times.

For detailed information about security options, see:

- Why you should not run your computer as an administrator.
- Run a program with administrative credentials.
- Create a shortcut using the runas command.
- Using the runas command with MMC.
- Best practices for security, Security overview, and Local Users and Groups Best practices for the broadest discussion and links to related security information.

# Understanding MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding MMC**
This section covers:

- MMC consoles
- Group Policy and MMC
- MMC in author mode

# MMC consoles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MMC consoles**

A new MMC console consists of a window divided into two panes. The left pane displays the console tree. The console tree shows the items that are available in the console. The right pane contains the details pane. The details pane shows information about, and functions pertaining to, these items. As you click different items in the console tree, the information in the details pane changes. The details pane can display many types of information including Web pages, graphics, charts, tables, and columns.

Each console has its own menus and toolbar, separate from those of the main MMC window, that help a user perform tasks. For more information about the main MMC window, see The MMC window.

The operating system you are using may already have preconfigured and saved console files available on the **Programs** menu or in the Administrative Tools folder in Control Panel.

Consoles which have been created with older versions of MMC (versions 1.1 and 1.2) can be read by MMC version 2.0. However, when files are saved by MMC version 2.0, they are automatically converted and cannot be restored to earlier versions. For further information, see Using files created with earlier versions of MMC with MMC 2.0.

# Group Policy and MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy and MMC**

Only administrators can use Group Policy to restrict or allow access to specific snap-ins or restrict the ability of a user or group to use author mode in MMC. This applies in a network where Group Policy has been configured.

To set policies for users of a particular computer, you must be an administrator for that computer or have equivalent rights. To set policies for an organizational unit in a domain, you must be an administrator for that domain or have equivalent rights. You cannot use MMC to create computer policies; you can *only* create user policies.

For more information about Group Policy, see Help.

For more information about applying Group Policy to MMC, see Setting Group Policy in MMC.

# MMC in author mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MMC in author mode**

You can use MMC in author mode to create new consoles or modify existing consoles. To create an administrative tool, you add snap-ins and other items to a console. You can create additional console windows that provide views of the various items that make up a console. You can also create taskpad views that contain shortcut links to run menu commands from different locations and tools in the console, as well as run command-line functions. After a console is saved, you can distribute it to users.

This section covers:

- The MMC window
- Snap-ins
- Taskpad views and tasks
- The console tree and console root
- The Favorites list
- Console access options

# The MMC window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The MMC window**

The components of an MMC console are contained in the MMC window. This window has several menus and a toolbar that provide commands to open, create, and save MMC consoles. The menu and toolbar on the MMC window are called the main menu bar and the main toolbar, respectively. In addition, there is a status bar at the bottom of the window and a description bar along the top of the details pane. To view the MMC window, open MMC by clicking **Start**, then **Run**, typing "mmc", and then clicking **OK**.

When you open a new MMC console, a console window appears in the workspace in the MMC window. In the console window, you can assemble and configure a new console and then work with the tools in the console. After you add items to a console, you can hide the main menu bar, main toolbar, description bar, and status bar to prevent users from making changes to the console. For more information, see Hide or display features of a saved MMC console and Hide or display menus and toolbars for snap-ins.

# Snap-ins

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Snap-ins**

A snap-in is the basic component of an MMC console. Snap-ins always reside in a console; they are not run outside of MMC.

When you install a component that has a snap-in associated with it, the snap-in is available to anyone creating a console on that computer (unless restricted by a user policy). For information about user policies, see Group Policy and MMC.

**Stand-alone and extension snap-ins**

MMC supports two types of snap-ins: stand-alone snap-ins and extension snap-ins. You can add a stand-alone snap-in, typically called a snap-in, to a console tree without adding another item first. An extension snap-in, typically called an extension, is always added to a stand-alone or extension snap-in that is already on the console tree. When extensions are enabled for a snap-in, they operate on the objects controlled by the snap-in, such as a computer, printer, modem, or other device.

When you add a snap-in or extension to a console, it might appear as a new item in the console tree, or it might add shortcut menu items, additional toolbars, additional property pages, or wizards to a snap-in already installed in the console.

**Adding snap-ins to a console**

You can add a single snap-in or multiple snap-ins and other items to a console. In addition, you can add multiple instances of a particular snap-in to the same console to administer different computers or to repair a damaged console. Each time you add a new instance of a snap-in to a console, any variables for the snap-in are set at default values until you configure the snap-in. For instance, if you configure a particular snap-in to manage a remote computer and then add a second instance of the snap-in, the second instance will not automatically be configured to manage the remote computer.

Typically, you can only add snap-ins that are installed on the computer you are using to author a console. However, if your computer is part of a domain, you can use MMC to download any snap-ins that are not locally installed, but that are available in the Active Directory directory service. For more information about adding published snap-ins and extensions to a console, see Creating consoles. For more information about distributing software using Active Directory, see Delegating administration.

# Taskpad views and tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Taskpad views and tasks**

Taskpad views are pages to which you can add views of the details pane of a console, as well as shortcuts to functions both inside and outside a given console. You can use these shortcuts to run tasks such as starting wizards, opening property pages, performing menu commands, running command lines, and opening Web pages. You can configure a taskpad view so that it contains all the tasks a given user might need. In addition, you can create multiple taskpad views in a console, so that you can group tasks by function or user.

A taskpad view may make it easier for new users to perform their jobs. For instance, you can add applicable tasks to a taskpad view and then hide the console tree, so that users can begin using tools before they are familiar with the location of particular items in the console tree or operating system.

You might also use taskpad views to make complex tasks easier. For instance, if you must frequently perform a task involving multiple snap-ins and other tools, you can present tasks to open or run the necessary dialog boxes, property pages, command lines, and scripts in a single location.

For more information about creating taskpad views and tasks, see Creating consoles.

# The console tree and console root

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The console tree and console root**

The console tree is a hierarchical structure in the left pane of an MMC console. The console tree shows the items that are available in a console. These items can include folders, snap-ins, controls, Web pages, and other tools.

To view the console tree of a new console, open MMC by clicking **Start**, then **Run**, type "mmc", and then click **OK**.

In a new console, the only item on the console tree is a folder labeled Console Root. You can create the administrative functions of the console by adding items to the console. For instructions about adding items to a console, see Creating consoles.

You do not need to put a console window on the console root. You can locate a console window on any item in the console tree. This hides items in the console tree above the old root and focuses the window on the new console root and any administrative tools under it.

For instructions about changing where a console is rooted, see Make a new console window from part of an MMC console tree.

**Containers on the console tree**

A container is any item on the console tree to which objects are added. Double-click ⊞ to expand a container and view its contents, and the minus sign ⊟ to collapse the container. The default container is a folder, which you can use to group related items in a console tree. However, you can also use MMC to add tools to items other than folders, which then also become containers.

The console root is a container. Its only function is to contain the console tree. Typically, the top item in a snap-in is a folder or another container.

Another type of item that can contain other items is a viewable item. When you click a viewable item, it displays a list, text, or graphic information in the details pane instead of additional items. Web pages are viewable items because they require a browser, which can open only in the details pane. Do not add items to viewable items because user might hide the console tree. If so, they might not know that a viewable item contains other items. Instead, if you have a collection of viewable items, you can add a folder to the console tree and group the viewable items in the folder. Then a user can select from all the items when the console tree is hidden.

For instructions about hiding the console tree, see Hide or display features of a saved MMC console.

One type of item that cannot contain other items is a leaf. When you click a leaf, it lists items in the details pane. These are typically items that are individually contained on the console tree. However, when items number in the hundreds or thousands, a snap-in uses a leaf instead of a container.

**Adding new functionality to a console**

You might find that none of the snap-ins provided by the **Add/Remove Snap-in** command provide the exact capabilities you require. You can try the following possible solutions:

- Check for additional snap-ins or extensions that might have been posted on the Web by Microsoft or other vendors.
- Author your own extension to an existing snap-in, or author a new stand-alone snap-in.

For more information about these options, see MMC Resources.

# The Favorites list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The Favorites list**

The Favorites list is located on the **Favorites** menu. You can use the Favorites list to create shortcuts to items in the console tree. For your convenience, you can add shortcuts to tools that you use often or tools that are several levels down in the console tree. You can also use the Favorites list to make it easier for new users to complete tasks. For instance, you can create a console that has shortcuts in the Favorites list to only the items in the console tree that a given user needs to perform a task. You can then hide the console tree to further simplify the view.

You can also use the Favorites list to organize taskpad views. For example, a complex console might have multiple taskpad views distributed among the multiple items in the console tree. However, you can add all these taskpad views to the Favorites list, which enables users to access them from a single location.

For more information, see Creating consoles.

# Console access options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Console access options**

When building a custom console, you can assign the console one of two access options: author mode or user mode. There are, in turn, three levels of user mode, so that there are four options for default access to a console:

| Mode | Description |
|---|---|
| Author mode | Enables full customization of the MMC console, including the ability to add or remove snap-ins, create new windows, create Favorites and taskpads, and access to all the options of the **Customize View** and **Options** dialog boxes. Users creating a custom console file for themselves or others typically use this mode. The resulting console is usually saved in one of the following user modes. |
| User mode-full access | The same as author mode, except that users cannot add or remove snap-ins, change console options, create Favorites, or create taskpads. |
| User mode-limited access, multiple windows | Provides access only to those parts of the console tree that were visible when the console file was saved. Users can create new windows but cannot close any existing windows. |
| User mode-limited access, single window | Provides access only to those parts of the console tree that were visible when the console file was saved. Users cannot create new windows. |

You can configure these options in the **Options** dialog box in MMC. For instructions about opening this dialog box, see Set MMC console options.

Changes made to consoles in author mode and consoles in user mode are saved differently. If you work with a console in author mode, when you close the console you are prompted to save your changes. However, if you work with a console in user mode and have cleared the **Do not save changes to this console** check box, available by clicking **Options** on the **File** menu, changes will be saved when you close the console.

If any of the following conditions apply, the default mode for a console is ignored and a console is opened in author mode:

For more information about console modes and their command-line syntax for MMC, see Open MMC.

Author mode access for MMC is unnecessary for users who do not need to create or change MMC consoles. A system administrator can configure user profile settings to prevent users from opening MMC in author mode, by disabling the **/a** option or the shortcut menu option. In addition, an administrator can use Group Policy settings to prevent users from opening MMC and saved consoles in author mode. For more information, see Group Policy and MMC.

# Using MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using MMC**

This section covers:

- Runas
- Using files created with earlier versions of MMC with MMC 2.0
- Using MMC consoles
- Using property pages in saved consoles
- Setting Group Policy in MMC
- Using the MMC help system
- Authoring MMC consoles
- Running 32-bit and 64-bit snap-ins in Windows Server 2003 and Windows XP Professional operating systems

# Using the runas command with MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the runas command with MMC**

It is good practice for administrators to use an account with restrictive permissions to perform routine, nonadministrative tasks, and to use an account with broader permissions only when performing specific administrative tasks. To accomplish this without logging off and back on, log on with a regular user account, and then use the **runas** command to run the tools that require the broader permissions.

With the **runas** command, you can run programs (*.exe), saved MMC consoles (*.msc), shortcuts to programs and saved MMC consoles, and Control Panel items. You can run them as an administrator while you are logged on to your computer as a member of another group, such as the Users or Power Users group.

You can use the **runas** command start any program, MMC console, or Control Panel item. As long as you provide the appropriate user account and password information, the user account has the ability to log on to the computer, and the program, MMC console, or Control Panel item is available on the system and to the user account.

With the **runas** command, you can administer a server in another forest (the computer from which you run a tool and the server you administer are in different domains).

If you try to start a program, MMC console, or Control Panel item from a network location using **runas**, it might fail because the credentials used to connect to the network file share are different from the credentials used to start the program. The latter credentials may not be able to gain access to the same network file share.

Some items, such as the Printers folder and desktop items, are opened indirectly and cannot be started with the **runas** command.

If the **runas** command fails, the **Secondary Logon** service might not be running or the user account you are using might not be valid. To check the status of the **Secondary Logon** service, in Computer Management, click **Services and Applications**, and then click **Services**. To test the user account, try logging on to the appropriate domain using the account.

For details about the **runas** command, see Runas.

# Using files created with earlier versions of MMC with MMC 2.0

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using files created with earlier versions of MMC with MMC 2.0**
MMC 2.0 reads consoles created with older versions of MMC.

Methods of saving files created in older versions of MMC include the following:

- If you work in MMC 2.0 authoring mode with a file created in an earlier version of MMC, you can save the file only in MMC 2.0.
- You can work with an MMC 1.2 file in MMC 2.0 in any of the user modes.
- In user modes, MMC 2.0 saves console changes from an earlier version in the new format into the user's profile directory. That converted file has all of the user's changes, such as window resizing. Additionally, the original console (from the earlier version of MMC) remains intact in its original format. However, when converted files are saved, the earlier format cannot be restored.

When you open a console from an earlier MMC version, MMC 2.0 opens it with the selected modified settings. However, if you copy this modified console to another computer, or delete the settings through the MMC menu **File/Options**, you will lose these modifications.

For more information about MMC, see:

- Add a snap-in to a new MMC console for a local computer
- Create a desktop shortcut to open MMC
- Install a program managed by an MMC snap-in
- Open a saved MMC console for local and remote computers

# Using MMC consoles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using MMC consoles**

You might need to use preconfigured MMC consoles (commonly called "saved consoles") that are part of your operating system or an application and never need to build custom consoles of your own. Typically, preconfigured consoles are available in the Administrative Tools folder in Control Panel or from the **Start** menu. If you save a custom console to a per-user Administrative Tools folder (located at *systemdrive*\Documents and Settings*user*\Start Menu\Programs\Administrative Tools), it is then available in the Administrative Tools folder on the **Programs** menu for that user.

It is likely that any preconfigured consoles that came as part of your operating system are configured to open in one of the three user modes. The default mode is **User mode - limited access, single window**. For more information about default console mode settings, see Console access options.

**Customizing the view for MMC and saved consoles**

You can use the **Customize View** dialog box on the **Customize** command on the **View** menu to hide or display elements for a console. One of the elements that you can hide is the **View** menu itself. If you do this, but then want to reconfigure the view, you can also access the **Customize View** dialog box from the **Customize View** command on the **System** menu. For more information, see Hide or display features of a saved MMC console and Hide or display menus and toolbars for snap-ins.

**Working with columns in saved consoles**

In saved consoles that display columns in the details pane, you can customize how the columns and rows appear. For instance, you can reorder or hide columns. You can also reorder rows alphabetically or chronologically by clicking the column heading. In addition, with certain snap-ins, you can filter columns based on additional attributes. If this feature is enabled, a row of drop-down list boxes that contain options for filtering is displayed beneath the column headings. For more information, see Reorder columns in an MMC console, Hide or display columns in an MMC console, and Filter rows in an MMC console.

You can also export the contents of columns to a text file. For more information, see Export columns in an MMC console to a text file.

If you customize the columns in a console, your settings are saved from session to session.

**Using consoles from older versions of MMC**

Consoles which have been created with older versions of MMC (versions 1.1 and 1.2) can be read by MMC version 2.0. However, saving these files in MMC 2.0 automatically converts them, and the older format cannot be restored. For further information, see Using files created with earlier versions of MMC with MMC 2.0.

# Using property pages in saved consoles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using property pages in saved consoles**

Much of the work done using consoles is performed from property pages. In fact, certain snap-ins or extensions may be accessed in a console only from a property page. For some tabs, you may be able to determine which snap-in provided a particular property page if you hold down the CTRL key while you point with the mouse to the label of the tab for the property page. Also, if you hold down the CTRL key while you point to the title bar of a particular properties dialog box, you can determine the path in the console tree to the item associated with the property page. For more information, see View the associated snap-in and item for a property page.

The properties dialog boxes for consoles often become hidden behind consoles and other tools. If you cannot find a property page that you were working on, simply reopen the properties dialog box or minimize the other tools on your screen.

# Setting Group Policy in MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting Group Policy in MMC**

This feature applies only in a network where Group Policy has been configured; a system administrator must enable Group Policy. For detailed information, see Group Policy overview.

| Policy | Action | Result |
|---|---|---|
| Enable **Restrict the user from entering author mode**. | Open a console file that is configured to open in author mode. | Users with administrator credentials can access author mode. Restrict access to author mode in MMC.<br><br>An error message appears indicating denied access to console in author mode. |
| Disable or do not configure **Restrict the user from entering author mode**. | Open a console file that is configured to open in author mode. | Use of author mode is unrestricted. |
| Enable **Restrict users to the explicitly permitted list of snap-ins**. | Open a console that contains a snap-in that is not on the permitted list, or a snap-in that is extended by a snap-in that is not on the permitted list. | An error message appears.<br><br>**Note**<br><br>• If you enable this policy, only explicitly permitted snap-ins appear in the list of available snap-ins in the **Add Standalone Snap-in** dialog box in MMC--all other snap-ins are restricted. |
| Disable or do not configure **Restrict users to the explicitly permitted list of snap-ins**. | Open a console that contains a snap-in that is not on the permitted list, or a snap-in that is not extended by a snap-in on the permitted list. | |
| Explicitly permit a snap-in. | Open a console that contains the snap-in, or a snap-in that is extended by the permitted snap-in. | Use of the console that contains the snap-in is unrestricted. |
| Explicitly restrict a snap-in. | Open a console file that contains the snap-in, or contains a snap-in that is extended by the restricted snap-in. | An error message appears. |
| Do not configure the policy to permit or restrict a snap-in. | Open a console file that contains the snap-in, or contains a snap-in that is extended by this snap-in. | Use of the console that contains the snap-in is unrestricted unless the user, group, or organizational unit is restricted to a list of explicitly permitted snap-ins. |

To open MMC or open a console in author mode:

- Type **mmc** at the command prompt.
- Type **mmc***path**filename*.**msc/a** at the command prompt.
- Right-click an .msc file, and then click **Author**.
- Open a console file that is configured to open in author mode.

To open a console file that contains the snap-in, or contains a snap-in that is extended by this snap-in:

- Type **mmc***path******filename*.**msc** at the command prompt.
- Double-click the .msc file.
- Right-click the .msc file, and then click **Open**.

# Using the MMC help system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the MMC help system**

You can access Help for MMC from the **Help** menu on the main toolbar, from the button on the toolbar of a console window, on the system menu of the MMC window or a console window, by right-clicking an item in the console tree or details pane, or by pressing the F1 key.

Help for MMC displays a combined table of contents that includes general Help for MMC and Help specific to each snap-in included in the console. Snap-in Help is provided only for existing Help files that were associated to snap-ins by their authors. If you are authoring a console, you may want to provide additional documentation for tasks associated with the console or provide information specifically related to your network.

# Authoring MMC consoles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authoring MMC consoles**
This section covers:

- Creating consoles
- Adding items to the console tree
- Adding taskpad views and tasks
- Viewing the console tree in a console window
- Rooting console windows on items in the console tree
- Adding items to the Favorites list
- Adding items to the Favorites list
- Saving consoles
- Using custom consoles

For more information about author mode, see MMC in author mode.

# Creating consoles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating consoles**

Before you author a console, identify the tasks the console will perform, the components to be administered, and the snap-ins and controls that are needed to perform the tasks. You should also consider whether you need to create a taskpad view and tasks. After you make these decisions, you can open a new console and start adding items to the console tree.

To create or edit a console, open the console in author mode (use the command-line option **/a**). To view the complete command-line syntax, see Open MMC.

# Adding items to the console tree

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding items to the console tree**

You can help users locate the components they need in the console by arranging items hierarchically or in groups on the console tree. To add items to the console tree, use the **Add/Remove Snap-in** command on the **File** menu of the main toolbar of MMC.

In the **Add/Remove Snap-in** dialog box, **Snap-ins added to** determines the item on the console tree under which new items are added. The default value is **Console Root**. Click an item in **Snap-ins added to** to locate an object elsewhere on the console tree.

The **Add Standalone Snap-in** dialog box displays a list of available snap-ins. For computers that are members of a domain, this list includes both locally-installed snap-ins and snap-ins published in the Active Directory directory service. For snap-ins that are available in Active Directory, **Not Installed** appears in the **Vendor** column.

If you have already installed a snap-in, and you want to enable one of its extension snap-ins, use the **Extensions** tab in the **Add/Remove Snap-in** dialog box. On this tab, select any item in the console tree that can be extended and view the extension snap-ins that you can enable or disable. When you enable an extension snap-in, it is automatically inserted in the console tree under the selected item. If there is more than one instance of a snap-in on the console tree, all instances of the snap-in are extended.

After an extension snap-in is enabled, you might notice that it appears in **Snap-ins that can be extended** along with the stand-alone snap-ins. This means that the extension snap-in also has extensions that you can enable.

Like stand-alone snap-ins, an administrator can also publish extension snap-ins in Active Directory. Extension snap-ins available to a user from Active Directory appear in the list of available extension snap-ins with any locally installed extension snap-ins, except that they are followed by the phrase **(not installed)**. You must specifically download extension snap-ins from the directory service to make them available in a console. Thus, to download all extensions for a given item, do not select the **Add all extensions** check box on the **Extensions** tab of the **Add/Remove Snap-in** dialog box. Instead, clear this check box and select the check boxes next to each extension that you want to download.

To add multiple instances of the same snap-in to a console either to manage multiple remote computers from the same console or to repair a damaged console, add a new instance of the snap-in to the console, configure it as needed, and then remove the old instance of the snap-in from the console. You might need to do this if you noticed that a tool is behaving differently from expected, that its configuration is outdated, or that it is timing out because resources have been removed.

For information about adding items to the console tree, see Add a snap-in to a new MMC console for a local computer, Add a published snap-in to a new MMC console, Add an extension snap-in to an MMC console, and Add a published extension snap-in to an MMC console.

# Adding taskpad views and tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding taskpad views and tasks**

Before you add taskpad views and tasks to a console, determine how many taskpad views you need. If you need multiple taskpad views, you also need to determine how the tasks are divided among those views. In addition, you should decide what kind of taskpad view you want to use--one that displays a list and tasks, or one that displays tasks only.

To create a taskpad view for a console, the console must contain at least one snap-in. Use the New Taskpad View wizard to configure the titles, headings, and lists that appear in the taskpad view, and to define whether a taskpad view is associated with a single item or multiple items in the console tree.

After you complete the New Taskpad View wizard, use the New Task wizard to add tasks to the taskpad view. Tasks can include menu commands for the items in the console, as well as commands that are run from a command prompt. Create commands to act on part of the console tree or details pane, or to open another component on your computer. However, if you create a task from a menu command for an item in the console tree, and then you remove that item, the task is disabled.

For information about creating and editing taskpad views and tasks, see Create a taskpad view in a saved MMC console, Create tasks for a taskpad view in saved MMC console, Edit a taskpad view in a saved MMC console, and Edit tasks in a taskpad view in a saved MMC console.

# Viewing the console tree in a console window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing the console tree in a console window**

When the console tree is visible, it appears in the left pane of a console window. Optionally, hide the console tree. For information about hiding the console tree, see Hide or display features of a saved MMC console.

What you view in a console window is determined by where the console window is rooted on the console tree. The initial window in a new console is rooted at Console Root, as are any windows you create with the **New Window** command on the **Window** menu on the main toolbar.

Close console windows to hide the console root or entire portions of the console tree. Then configure the console to prevent users from viewing hidden portions of the tree. However, the entire console tree is always saved when you save a console.

# Rooting console windows on items in the console tree

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Rooting console windows on items in the console tree**

While you can use MMC to make consoles that meet every administrative need for a group, you can also design simple consoles for less experienced users. For instance, after you add items to a console tree, open additional console windows rooted on any item in the console tree. Use this method to create console windows that display specific administrative components of the console tree. Then close windows that show portions of the console tree that users do not need.

Closing a window does not remove any items from the console tree; when you close a window, you remove only that view of the console tree. The entire tree is still saved when you save the console, and you can always view the console root.

For information about changing the console root, see Make a new console window from part of an MMC console tree.

For information about viewing a console root that is hidden, see Make a new console window from an MMC console.

# Adding items to the Favorites list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding items to the Favorites list**
Use the Favorites list to create shortcuts to items in the console tree and to taskpad views. Then access these items from the **Favorites** menu.

Note: The **Favorites** tab in the Console scope pane has been removed in MMC 2.0 but is present in MMC 1.2.

For more information about creating the Favorites list, see Add an item to the Favorites list in an MMC console and Organize the Favorites list in an MMC console.

**Configuring console options**
Use the **Options** dialog box, available from the **File** menu, to choose an icon for, change the title of, and choose the default mode for a console.

Icons are available from many sources including the Shell32.dll file, located in the *systemroot*\system32 folder.

When changing the title of a console, keep in mind that, by default, as you click items in the console tree, the title bar displays the path to the selected item. But if you change the title of a console, the title bar will not display this path. However, if you delete the title, the default behavior is restored.

When considering which mode to set as the default for a console, remember that if you use one of the user modes, you can open the console at any time by using the **Author** or **Run As** commands--you do not need to use author mode for most management tasks.

For more information about console options, see Set MMC console options and Console access options.

# Configuring console options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring console options**

Use the **Options** dialog box, available from the **File** menu, to choose an icon for, change the title of, and choose the default mode for a console.

Icons are available from many sources including the Shell32.dll file, located in the *systemroot*\system32 folder.

When changing the title of a console, keep in mind that, by default, as you click items in the console tree, the title bar displays the path to the selected item. If you change the title of a console, the title bar will not display this path. However, if you delete the title, the default behavior is restored.

When considering which mode to set as the default for a console, remember that if you use one of the user modes, you can open the console at any time by using the **Author** or **Run As** commands--you do not need to use author mode for most management tasks.

For more information about console options, see Set MMC console options and Console access options.

# Saving consoles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Saving consoles**

Saving an MMC console to a file preserves the list of loaded snap-ins for the console, the arrangement and contents of console windows in the main MMC window, the default mode, and information about permissions. All the configuration settings for the tools and controls are saved with the console and restored when the console file is opened. You can open a console file on different computers or even different networks and restore the saved settings for all the tools.

Console files have an .msc (management saved console) file name extension. The operating system you are using may already have preconfigured console files that are available from the Administrative Tools folder in Control Panel. Many applications also include preconfigured consoles, which are typically available from the **Start** menu. If you create console files and save them in your per-user Administrative Tools folder (located at *%systemdrive%*\Documents and Settings*user*\Start Menu\Programs\Administrative Tools), they are available from the Administrative Tools folder on the **Programs** menu for you only--the console files are not available for other users of the computer. After you save a console, you can distribute it by using a floppy disk, e-mail, or your network.

If you are working with a console in author mode, you can save changes to the console by using the **Save** command on the **File** menu. If you are working with a console in one of the user modes, saving changes to the console is determined by whether the **Do not save changes to this console** check box (available by clicking **Options** on the **File** menu) was selected when the console was configured. If this check box was not selected, changes to the console are automatically saved when you close MMC; if it was selected, changes to the console are discarded when you close MMC. For more information about saving consoles, see Save an MMC console file. For more information about console modes, see Console access options.

# Using custom consoles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using custom consoles**
After you create and save an MMC console, you can use it on your local computer, send it to other users in e-mail, post it on your network or the Web, or copy it to a floppy disk and install it on other computers. You can also use MMC and the Active Directory directory service to publish consoles or assign consoles to users. For more information about distributing software, see Active Directory.

**Requirements to use a console**
To use a console, you must have access to the services and administrative tools included in the console, either installed on the local computer or available on the network. You must also have administrative permissions for the components on the system that is administered by the console.

**Opening a saved console**
If you save a console to the per-user Administrative Tools folder (located at *systemdrive*\Documents and Settings*user*\Start Menu\programs\Administrative Tools), it is then available to that user in the Administrative Tools folder on the **programs** menu.

If you know the name and location of a console file, you can open it as you would any other document:

- By double-clicking the .msc file
- By right-clicking the .msc file, and then clicking **Open**
- From a command prompt

On computers running Windows 2000, Windows XP Professional, or an operating system in the Windows Server 2003 family, you can log on to your computer with user rights and perform routine tasks. However, you must use the **Run As** command to open console files and perform administrative tasks that require Administrator rights. You can use this command by typing **runas** at the command prompt or by right-clicking an .msc file, and then clicking **Run As**.

For more information about opening consoles, see Open a saved MMC console for local and remote computers, and Open MMC.

# Running 32-bit and 64-bit snap-ins in Windows Server 2003 and Windows XP Professional operating systems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Running 32-bit and 64-bit snap-ins in Windows Server 2003 and Windows XP Professional operating systems**

The 64-bit versions of the Microsoft® Windows°Server™°2003 and Microsoft® Windows®°XP Professional operating systems can run both the 32-bit Microsoft Management Console (MMC32) and the 64-bit MMC (MMC64). MMC32 can run only 32-bit snap-ins, and MMC64 can run only 64-bit snap-ins. However, some snap-ins are available in both 32-bit and 64-bit form. A snap-in is considered available in 64-bit form if it defined by an **HKEY_CLASSES_ROOT\CLSID***snap-in clsid***\InprocServer32** key in the registry, where *snap-in clsid* is the registered CLSID for the snap-in.

This topic describes how to determine and control which version of MMC is run in 64-bit versions of Windows operating systems.

**Start-up considerations if /32 or /64 is specified**

MMC can be run from the command line by using a command such as the following that includes a console file name:

**%systemdrive%\Documents and Settings** *username* **\Start Menu\programs\Administrative Tools\myconsole.msc**

On the 64-bit version of a Windows operating system, you can include one of the following command-line parameters:

**/32**
Specifies MMC32.

**/64**
Specifies MMC64. This is the default on the 64-bit version of a Windows operating system.

If both of these command-line parameters are used, the last one on the command line takes effect. MMC64 yields transparently to MMC32 if certain conditions are met, as described in the following section.

**Start-up considerations if a console file is specified**

If the command for running MMC does not contain a **/32** or **/64** parameter, but does contain a console file name, MMC64 examines the console file before determining which version to run. The decision sequence is as follows:

1. If all snap-ins in the console are available in 64-bit form, MMC64 is used.

2. If all snap-ins in the console are available in 32-bit form and some are not available in 64-bit form, MMC32 is used. (MMC64 starts MMC32, and then MMC64 terminates.)

3. If some of the snap-ins are available only in 64-bit form and others are available only in 32-bit form, the version of MMC that can run the majority of the snap-ins is used. The remaining snap-ins are not loaded.

   For example, if a console contains three snap-ins that are available only in 32-bit form and two snap-ins that are available only in 64-bit form, MMC32 is used and only the three 32-bit snap-ins appear when the console is opened. The two snap-ins that are available only in 64-bit form would not be loaded.

4. If some of the snap-ins are available only in 64-bit form and an equal number of snap-ins are available only in 32-bit form, you will be prompted to specify which version of MMC to run. If you choose MMC64, only the 64-bit snap-ins are loaded. If you choose MMC32, only the 32-bit snap-ins are loaded.

**When to specify the path to the console file**

In Windows XP and Windows Server 2003 operating systems, the system files are stored in the %windir%\system32 folder. In the 64-bit versions of these operating systems, the \system32 folder contains 64-bit files, including 64-bit consoles such as eventvwr.msc. If there is also a 32-bit version of the snap-in, it resides in the %windir%\syswow64 folder. In this example, the 32-bit version of Event Viewer would be %windir%\syswow64\eventvwr.msc.

If there is both a 64-bit version of the console in the \system32 folder and a 32-bit version in the \syswow64 folder and you want to explicitly run the 32-bit version, you must bypass the MMC decision-making process by specifying both the path and

the **/32** command-line parameter. For example:

**MMC /32 %windir%\syswow64\eventvwr.msc**

**Note**

- Double-clicking a shortcut to a file in the \syswow64 folder is like double-clicking the file in Windows Explorer: it will run MMC without the **/32** command-line parameter. To open the 32-bit consoles in the \syswow64 folder, click **Start**, click **Run**, and then type the command specifying both the path and the **/32** command-line parameter, instead of using a shortcut.

# MMC command-line options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MMC command-line options**

Using MMC command-line options, you can open a specific MMC console, open MMC in author mode, or specify that the 32-bit or 64-bit version of MMC is opened.

**Syntax**

**mmc** *path* **** *filename.msc*[**/a**] [**/64**] [**/32**]

**Parameters**

***path* \ *filename.msc***

Starts MMC and opens a saved console. You need to specify the complete path and file name for the saved console file. If you do not specify a console file, MMC opens a new console.

**/a**

Opens a saved console in author mode. Used to make changes to saved consoles.

**/64**

Opens the 64-bit version of MMC (MMC64). Use this option only if you are running a Microsoft 64-bit operating system and want to use a 64-bit snap-in.

**/32**

Opens the 32-bit version of MMC (MMC32). When running a Microsoft 64-bit operating system, you can run 32-bit snap-ins by opening MMC with this command-line option when you have 32-bit only snap-ins.

**Remarks**

- Using the *path******filename.msc* command-line option

  You can use environment variables to create command lines or shortcuts that do not depend on the explicit location of console files. For instance, if the path to a console file is in the system folder (for example, **mmc** *c:\winnt\system32\console_name.msc*), you can use the expandable data string %systemroot% to specify the location (**mmc**%systemroot%\system32\console_name.msc). This may be useful if you are delegating tasks to people in your organization who are working on different computers.

- Using the **/a** command-line option

  When consoles are opened with this option, they are opened in author mode, regardless of their default mode. This does not permanently change the default mode setting for files; when you omit this option, MMC opens console files according to their default mode settings.

- After you open MMC or a console file in author mode, you can open any existing console by clicking **Open** on the **Console** menu.

- You can use the command line to create shortcuts for opening MMC and saved consoles. A command-line command works with the **Run** command on the **Start** menu, in any command-prompt window, in shortcuts, or in any batch file or program that calls the command.

**Formatting Legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |

| | |
|---|---|
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Program output |

**See Also**

**Concepts**

Command-line reference A-Z
Set
Dir
Path
Running 32-bit and 64-bit snap-ins in Windows Server 2003 and Windows XP Professional operating systems

# MMC Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

For more information about MMC, see the following resources:

- The Personal Support Center Web site at the Microsoft Web site contains general support information and articles written by support professionals at Microsoft.

- For further information, see Windows Server 2003, Standard Edition Resource Kit at the Microsoft Web site.

# Accessibility for MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessibility for MMC**
In addition to Microsoft Windows accessibility products and services, MMC provides keyboard shortcuts for selecting commands and for navigating in and between console windows:

- Navigation in MMC
- MMC main window keyboard shortcuts
- MMC console window keyboard shortcuts

Snap-ins and other items that are added to an MMC console may also provide accessibility features. For more information, refer to the Help or other documentation for the console item.

# Navigation in MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Navigation in MMC**

You can browse an MMC console tree by expanding and collapsing branches by clicking the plus sign ⊞ and minus sign ⊟ icons, or by double-clicking with the mouse. The following table lists the keystrokes you can use to move in and between console windows.

| Keystroke action | Result |
|---|---|
| TAB or F6 | Moves forward between panes in the active console window. |
| SHIFT+TAB or SHIFT+F6 | Moves backward between panes in the active console window. |
| CTRL+TAB or CTRL+F6 | Moves forward between console windows. |
| CTRL+SHIFT+TAB or CTRL+SHIFT+F6 | Moves backward between console windows. |
| PLUS SIGN (+) on the numeric keypad | Expands the selected item. |
| MINUS SIGN (-) on the numeric keypad | Collapses the selected item. |
| Asterisk (*) on the numeric keypad | Expands the entire console tree below the root item in the active console window. |
| UP ARROW | Moves the selection up one item in a pane. |
| DOWN ARROW | Moves the selection down one item in a pane. |
| PAGE UP | Moves the selection to the top item visible in a pane. |
| PAGE DOWN | Moves the selection to the bottom item visible in a pane. |
| HOME | Moves the selection to the first item in a pane. |
| END | Moves the selection to the last item in a pane. |
| RIGHT ARROW | Expands the selected item. If the selected item does not contain hidden items, behaves like DOWN ARROW. |
| LEFT ARROW | Collapses the selected item. If the selected item doesn't contain exposed items, behaves like UP ARROW. |
| ALT+RIGHT ARROW | Moves the selection to the next item. Performs the same function as the Forward arrow on the toolbar. |
| ALT+LEFT ARROW | Moves the selection to the previous item. Performs the same function as the Back arrow on the toolbar. |

The following table lists the mouse actions you can use to move through an MMC console tree.

| Mouse action | Result |
|---|---|
| Click | Selects an item. |
| Double-click | Displays or hides items contained by the selected item. Displays properties for or opens an item. |

| | |
|---|---|
| Right-click | Displays the **Action** shortcut menu for the selected item. |

**Note**

- If the mouse buttons do not behave as expected, the right and left buttons might be switched. The computer you are using might be configured so that the default actions result from the left and right mouse buttons being reversed.

# MMC main window keyboard shortcuts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MMC main window keyboard shortcuts**

The following table lists keyboard shortcuts for the menu commands that act on the entire console or the main window of a console.

| Action | Result |
|---|---|
| CTRL+O | Opens a saved console. |
| CTRL+N | Opens a new console. |
| CTRL+S | Saves the open console. |
| CTRL+M | Adds or removes a console item. |
| CTRL+W | Opens a new window. |
| F5 | Refreshes the content of all console windows. |
| ALT+SPACEBAR | Displays the MMC window menu. |
| ALT+F4 | Closes the console. |
| ALT+A | Displays the **Action** menu. |
| ALT+V | Displays the **View** menu. |
| ALT+F | Displays the **File** menu. |
| ALT+O | Displays the **Favorites** menu. |

**Note**

- The CTRL+W keyboard shortcut is available for consoles opened in author mode or user mode—full access. Other CTRL+ shortcuts are only available for consoles opened in author mode.

# MMC console window keyboard shortcuts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MMC console window keyboard shortcuts**

The following table lists keyboard shortcuts for the menu commands that act on the active console window in a console or on the contents of a console window.

| Action | Result |
|---|---|
| CTRL+P | Prints the current page or active pane, only if the snap-in implements the printing function. |
| ALT+MINUS SIGN | Displays the window menu for the active console window. |
| SHIFT+F10 | Displays the **Action** shortcut menu for the selected item. |
| F1 | Opens the Help topic, if any, for the selected item. |
| F5 | Refreshes the content of all console windows. |
| CTRL+F10 | Maximizes the active console window. |
| CTRL+F5 | Restores the active console window. |
| ALT+ENTER | Displays the properties dialog box, if any, for the selected item. |
| F2 | Renames the selected item. |
| CTRL+F4 | Closes the active console window. When a console has only one console window, this closes the console. |

# Troubleshooting MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Help for a snap-in cannot be found.
- A snap-in is not listed in the Add Standalone Snap-in dialog box.
- When using a console, one or more of the following occurs: the console is not behaving as expected, the console is timing out, or error messages are displayed that you have not seen before.

**Help for a snap-in cannot be found.**
**Cause:** The snap-in may not provide a Help file to merge into the MMC Help table of contents. Instead, the snap-in may use a different method to provide Help.

**Solution:** Open a console in author mode that contains the snap-in, click the item in the console tree for which you want Help, and then click the **Help** menu. A separate command for the snap-in Help may be displayed on the **Help** menu above the **Help Topics** command. If you do not see a separate command, view the contents of the toolbars and details pane for a Help button or additional icon. If you do not see a separate command, button, or icon, then no Help is available from MMC for the snap-in.

**A snap-in is not listed in the Add Standalone Snap-in dialog box.**
**Cause**: The corresponding service or software may not be installed on your computer.

**Solution**: You must install the component to view its snap-in.

**See also**: Install a program managed by an MMC snap-in.

**Cause**: If you are using a computer running Windows 2000, Windows XP Professional, or an operating system in the Windows Server 2003 family, an administrator may have set a policy that restricts you from accessing the snap-in.

**Solution**: An administrator can verify whether a policy setting is restricting access to the snap-in, and change the policy if necessary.

**See also**: Setting Group Policy in MMC

**When using a console, one or more of the following occurs: the console is not behaving as expected, the console is timing out, or error messages are displayed that you have not seen before.**
**Cause**: One of the following might be the cause. The console file may have become damaged. Or, because of changes to your network, the configuration of the console may have become outdated.

**Solution**: Determine which snap-in is causing the problem in the console and then refer to the troubleshooting information in the documentation for that snap-in to find out how to repair or reconfigure the snap-in. To determine which snap-in is causing the problem, right-click the item in the console tree or the details pane, and then click **properties**. Hold down the CTRL key and point to the labels of the tabs in the dialog box.

If you cannot directly repair or reconfigure the console, try adding another instance of the snap-in you are using to the console, configure it as needed, and then remove the old version of the snap-in.

**See also**: Add a snap-in to a new MMC console for a local computer, Snap-ins, and View the associated snap-in and item for a property page.

# Remote Administration Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote administration tools**
This section covers:

- Remote Administration
- Using Web Interface for Remote Administration
- Administering Remote Assistance
- Remote Administration using Terminal Services
- Windows Server 2003 Administration Tools Pack Overview
- Emergency Management Services

# Remote Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Administration**
Using tools included with Microsoft® Windows Server 2003 family operating systems CDs, you can remotely manage your system servers running Microsoft® Windows 2000 and Windows Server 2003 family operating systems. You can also remotely manage Windows Server 2003 family operating systems from a computer using Microsoft® Windows® XP Professional.

- For general information, see Remote administration overview.
- For tips about managing your server remotely, see Remote Administration Best practices.
- For other sources of information, see Remote Administration Resources.

# Remote administration overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote administration overview**

The following table presents remote management tasks and links to information about the tools. Unless indicated otherwise, these tools are used on Microsoft® Windows Server™ 2003 operating systems and, in some cases, Microsoft® Windows® XP Professional.

| Task | Tools | Comments |
|---|---|---|
| Manage several servers by performing similar tasks. | Microsoft Management Console<br><br>Author an MMC console file<br><br>Work with MMC and saved console files | You can use appropriate saved MMC consoles, if available, or you can create custom MMC consoles for tasks you frequently either delegate or perform. |
| Log onto one computer remotely and manage it as if you were there. | Remote Desktop Connection | In addition to Windows Server 2003 operating systems, this feature can be used with Microsoft® Windows® 2000 Server. The Remote Desktop feature must be enabled on the remote computer. |
| Switch among several computers that you want to log on to remotely. | Remote Desktops snap-in | In addition to Windows Server 2003 operating systems, this feature can be used with Windows 2000 Server. The Remote Desktop feature must be enabled on the remote computer. |
| Administer a server running Windows Server 2003 from any computer on your network. | Remote Desktop for Administration | In addition to Windows Server 2003 operating systems, this feature can be used with Windows 2000 Server. The Remote Desktop feature must be enabled on the remote computer. |
| Manage servers from Windows XP Professional. | Windows Server 2003 Administration Tools Pack Overview<br><br>Remote Desktop Connection | Windows Server 2003 Administration Tools Pack must be installed on Windows XP Professional to permit server management using these tools. The Remote Desktop feature must be enabled on servers to permit use of the Remote Desktop Connection. |
| Manage your Web server using a Web browser on a remote computer. | Using Web Interface for Remote Administration | Use Internet Explorer, version 6 or later, for remote administration using Web Interface. Web Interface is not installed by default for Web servers. |
| Assist in real time an operator on a remote computer running Windows Server 2003 family operating system, Windows XP Professional or Microsoft® Windows® XP Home Edition. | Administering Remote Assistance (www.microsoft.com) | The Remote Assistance connection requires explicit permission from the connected user, and both computers must be running Windows XP Professional or Windows Server 2003 operating systems. |
| Connect to a remote computer when that computer cannot access a network or is not in a functional state because of hardware or software failure. | Telnet | All commands must be entered through the command line. Capabilities are limited, and security is minimal. |

| | | |
|---|---|---|
| Manage disk space on servers by copying files to a library of magnetic tapes or magneto-optical disks. | Remote Storage | Some Windows components require configuration before they can be used and will be displayed in the **Add or Remove Windows Components** in **Add or Remove programs** in **Control Panel**. Remote storage requires installation. |

**Notes**

- Depending on the Group Policy settings on your network, you might not be able to perform some of these tasks.
- Some features require that specific ports be opened on the server in order for the role to be administered remotely. For information about these port requirements, see Service overview and network port requirements for the Windows Server system (http://go.microsoft.com/fwlink/?LinkID=61531) and Windows Firewall Settings (http://go.microsoft.com/fwlink/?LinkID=38651).
- You can use command-line tools to automate many processes and tasks. For more information, see Automating processes and tasks.
- You must physically access each computer to restart, change hardware, and enable the Remote Desktop feature. Once the Remote Desktop feature has been enabled, it will remain enabled.
- For an explanation of the various Remote Desktop terminologies, see Remote Administration using Terminal Services.
- For more information about remote administration, see the following:
    - Common Administrative Tasks
    - New Features
    - Quick Guide for Finding Tools
    - Remote Administration Resources

# Remote Administration Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**To perform many tasks on a single, remote server, use Remote Desktop for Administration.**

- To remotely administer a server running a Windows Server 2003 family operating system, you do not need to install Terminal Server. Instead, use Remote Desktop for Administration. On these operating systems, Remote Desktop for Administration is on by default but is not enabled. After you enable remote connections, Remote Desktop for Administration allows you to administer your server from any computer on your network running Windows Server 2003 family, Windows XP Professional, and Windows 2000. Windows Server 2003 , Web Edition, does not support remote administration.

- You can use Remote Desktop with Windows 2000 Server. You need to install Terminal Services for Remote Administration as described in the Windows 2000 Help.

  For more information, see Enable or disable Remote Desktop and Remote Desktop for Administration.

**To do similar tasks on many servers, use MMC snap-ins.**

- Most of the administration tools provided with Windows Server 2003 family operating systems are Microsoft Management Console (MMC) snap-ins that can be used to administer remote servers as well as your local computer. The advantages of using MMC include the following:

  - You can create a console with the tools you use for the tasks you do most often.

  - You can set the focus for a tool to any of the servers you administer, and switch between servers and tools within a single MMC console.

  For more information, see Microsoft Management Console.

**To manage servers running Windows Server 2003 family operating systems from a computer using Windows XP Professional , use Windows Server 2003 Administration Tools Pack**

- Windows Server 2003 Administration Tools Pack when installed on Windows XP Professional (ordinarily a client operating system) provides many remote management tools for managing your servers.

- Windows Server 2003 Administration Tools Pack comes as adminpak.msi on the Windows Server 2003 family operating system and is easily installed from the server CD onto Windows XP Professional by installing the .msi from the server to the Windows XP Professional operating system.

- The Administration Tools Packs in the Windows Server 2003 family and Windows 2000 are not mutually compatible. You must use the Windows 2000 Administrative Tools to administer Windows 2000, and the Windows Server 2003 Administration Tools Pack to administer the Windows Server 2003 family operating system.

  For more information, see Windows Server 2003 Administration Tools Pack Overview.

# Remote Administration Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**
The following resources provide additional information about remote administration.

**Windows Server 2003 Resource Kit**
The Microsoft Windows Server 2003 Resource Kit contains additional tools that you can use for remote administration.

**Microsoft Platform Software Development Kits (SDK)**
For technical information about Active Directory Scripting Interfaces (ADSI), Microsoft Management Console (MMC), Windows Management Instrumentation (WMI) Software Development Kits (SDK), see the Microsoft Web site.

**TechNet**
TechNet is a service that contains resources for information technology (IT) professionals. This includes access to articles written by support professionals at Microsoft, career and industry information, discussion groups, technical feedback, and software.

For more information, see Microsoft TechNet.

# Using Web Interface for Remote Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Web Interface for Remote Administration**
You can use Web Interface for Remote Administration to manage your application server by using a Web browser on a remote computer.

**To install Web Interface for Remote Administration of an application server**

1. Open the Windows Components Wizard.
2. In the Windows Components Wizard, select the **Application Server** check box, and then click **Details**.
3. Select the **Internet Information Services (IIS)** check box, click **Details**, select the **World Wide Web Service** check box, click **Details**, and then select the **Remote Administration (HTML)** check box.

### Notes

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Web Interface for Remote Administration of an application server is installed by default on Windows Server 2003, Web Edition only. It is not installed by default on other Windows Server 2003 operating systems.
- Internet Explorer version 6.0 or later is recommended for Remote Administration.
- Installation of Remote Administration is not supported on a domain controller.

**To open Web Interface for Remote Administration**

1. Click **Start**, and then click **Run**.
2. Type:

   **https://*ComputerName*:8098**

   Where *ComputerName* is the name of the computer on which you want to run Remote Administration.

### Note

- To open Web Interface for Remote Administration Help, open Web Interface for Remote Administration, and then click the **Help** tab.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Remote Administration using Terminal Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Administration using Terminal Services**
In Microsoft® Windows Server 2003 family operating systems, Terminal Services technology is the basis for several features that enable you to connect to remote computers and perform administrative tasks.

- Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode) provides remote server management capabilities for Windows Server 2003 family operating systems. Using this feature, you can administer a server from virtually any computer on your network. No license is required for up to two simultaneous remote connections in addition to the server console session. A corresponding desktop version of Remote Desktop for Administration is available on Microsoft® Windows® XP Professional, and is called Remote Desktop.

- The Remote Desktops MMC snap-in allows you to create remote connections to the console session of multiple terminal servers, as well as computers running Windows 2000 or Windows Server 2003 family operating systems.

- Remote Desktop Connection, available on Windows Server 2003 family operating systems as well as on Microsoft® Windows® XP operating systems, enables you to log on to a remote computer and perform administrative tasks, even from a client computer that is running an earlier version of Windows.

This section provides information about using Remote Desktop for Administration and its related tools:

- Remote Desktop for Administration

- Remote Desktops snap-in

- Remote Desktop Connection

# Remote Desktop for Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Desktop for Administration**
Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode) provides remote access to the desktop of any computer running a Microsoft® Windows Server 2003 family operating system, allowing you to administer your server from virtually any computer on your network.

- To find features that have been moved, see New ways to do familiar Remote Desktop for Administration tasks.
- For tips about using Remote Desktop for Administration, see Remote Desktop for Administration Best practices.
- For help with specific tasks, see Remote Desktop for Administration How To....
- For general background information, see Remote Desktop for Administration Concepts.
- For problem-solving instructions, see Troubleshooting Remote Desktop for Administration.

# New ways to do familiar Remote Desktop for Administration tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**
This table lists common tasks for Remote Desktop for Administration (formerly Terminal Services in Remote Administration mode). The user interface for performing these tasks is different in Windows Server 2003 family operating systems than it was in Microsoft Windows NT 4.0 or Microsoft Windows 2000.

Please choose your upgrade path:

- Windows NT 4.0 to one of the Windows Server 2003 family operating systems
- Windows 2000 to one of the Windows Server 2003 family operating systems

**Windows NT 4.0 to one of the Windows Server 2003 family operating systems**

| If you want to | In Windows NT 4.0 use | In Windows Server 2003 family operating systems do this |
|---|---|---|
| Create a Terminal Services connection | Terminal Server Client | Open Remote Desktop Connection. For more information, see Remote Desktop Connection. |
| Remotely administer your server with Terminal Services | Not available | Use Remote Desktop for Administration. For more information, see Remote Desktop for Administration overview. |
| Create a Terminal Services connection using the World Wide Web | Not available | Use Remote Desktop Web Connection. For more information, see Remote Desktop Web Connection. |

**Windows 2000 to one of the Windows Server 2003 family operating systems**

| If you want to | In Windows 2000 use | In Windows Server 2003 family operating systems do this |
|---|---|---|
| Create a Terminal Services connection | Client Connection Manager and Terminal Services Client | Open Remote Desktop Connection. For more information, see Remote Desktop Connection. |
| Remotely administer your server with Terminal Services | Use Terminal Services in Remote Administration mode, enabled with **Add or Remove Programs** | Remote Desktop for Administration. For more information, see Remote Desktop for Administration overview. |
| Create a Terminal Services connection using the World Wide Web | Terminal Services Advanced Client | Use Remote Desktop Web Connection. For more information, see Remote Desktop Web Connection. |

# Remote Desktop for Administration Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### Best practices

#### Security

- **Add yourself to the Remote Desktop Users group for your computer.**

  When you do this, you do not have to log on as an administrator to access your computer remotely. You should not add yourself to the Administrators group, and you should avoid running your computer while you are logged on as an administrator unless you are doing tasks that are restricted to administrators. For most computer activity, log on as a member of the Users or Power Users group. If you need to perform an administrator-only task, log on as an administrator, perform the task, and then log off. For more information, see Why you should not run your computer as an administrator.

- **Require all members of the Remote Desktop Users group to log on with a strong password.**

  This is especially important if your computer is connected directly to the Internet. For more information, see Strong passwords.

- **Install the latest version of Remote Desktop Connection.**

  If you are using the Windows 2000 32-bit Terminal Services Client, you should upgrade to the Windows Server 2003 Remote Desktop Connection.

  For more information, see Install Remote Desktop Connection (32-bit computers).

- **Review general security considerations.**

  For more information, see Best practices for security.

#### Remote Administration

- **Coordinate Remote Administration Tasks with Other Administrators**

  The Remote Desktop feature of Windows Server 2003 family operating systems is not meant to provide a managed multiuser experience. Using the two remote connections plus the console can implement a collaborative operation, but it is not designed to support general access by multiple simultaneous administrators. In particular, ensure that administrators do not run potentially destructive applications at the same time. For example, two administrators trying to reconfigure the disk subsystem can undermine each other's work, or even worse, destroy data. Use the Terminal Services Manager tool or the Query user command line utility to check for the presence of other administrators.

  To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.

- **Remote Administration Is Not Application Serving**

  Many general office applications require special installation, install scripts, or environment management to perform well in a remote session. Terminal Services provides these when you install Terminal Server, but they are not available for Remote Desktop for Administration. For general desktop and application remote access requirements, use a dedicated server with Terminal Server installed.

  For more information, see Terminal Server.

- **Configure the Remote Desktop Session to Disconnect When Connection is Broken**

  This is the default setting, and is especially important if you perform system updates over unreliable network connections. If a session is interrupted due to a network problem, the session goes into a disconnected state and continues executing the processes that were running before the interruption occurred. If the session is configured to reset when the connection breaks, all processes running in that session will stop, which is similar to stopping an application by using End Task.

- **Configure Disconnect and Reset Timeouts**

  Because it is not possible to have more than two remote sessions, remote administrators might be locked out of a server if two remote sessions (using different user accounts) are in either an active or disconnected state. When configuring disconnect timeouts, it is critical that sessions that are disconnected do not get reset prematurely. For this reason, it may be useful to perform remote administration tasks that should not be accidentally reset using a shared administrator account, such as a local machine account. You can use the account **Properties** tab to configure this account not to reset after it is disconnected.

  **Notes**

  - Group Policy settings may override settings in the user account **Properties** tab.
  - Using a shared administrator account will not allow you to track individual administrators using the computer.

  For more information, see Configuring Remote Desktop for Administration.

- **Avoid Tasks that Require Reboots**

  Some tasks, for example system upgrades and domain controller promotion, require reboots at their completion. These tasks work correctly from within a Remote Desktop session, but you should be aware that something as simple as a floppy disk in the drive or a bad boot sector on the disk could prevent the server from restarting. Therefore, it is advisable not to remotely reboot mission-critical servers unless you have the ability to physically intervene at the server if a problem occurs.

- **Avoid relying on Server Console Messages**

  It is not possible to see server console messages when logging on using Terminal Services unless you are connected to the console session. Therefore, it is good practice to check the Server event logs, rather than relying on a system pop-up.

- **Administrator Collaboration**

  Using the Terminal Services Manager, you can control another Terminal Services session remotely. However, you cannot interact with the console from the Terminal Services Manager.

  For more information, see Remotely control a session.

# Remote Desktop for Administration How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**


- Enable or disable Remote Desktop
- Connect to Remote Desktop for Administration
- Add users to the Remote Desktop Users group
- Share the client setup folder
- Install Remote Desktop Connection (32-bit computers)
- Install Remote Desktop Web Connection
- Connect to another computer using Remote Desktop Web Connection
- Connect to the console session of a server using Remote Desktop for Administration
- Change the icon associated with the shortcut to a connection file

# Enable or disable Remote Desktop

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable Remote Desktop**
**Using Group Policies (best practice)**

1. Open Group Policy.

2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Allows users to connect remotely using Terminal Services** setting.

3. Do one of the following:

   - To enable Remote Desktop, click **Enabled**.

   - To disable Remote Desktop, click **Disabled**.

     If you disable Remote Desktop while users are connected to the target computers, the computers maintain their current connections, but will not accept any new incoming connections.

**Important**

When you enable Remote Desktop on a computer, you enable the capability for other users and groups to log on remotely to the computer. However, you must also decide which users and groups should be able to log on remotely, and then manually add them to the Remote Desktop Users group. For more information, see Enabling users to connect remotely to the server and Add users to the Remote Desktop Users group.

You should thoroughly test any changes you make to group Policy settings before applying them to users or computers. Fore more information about testing policy settings, see Resultant Set of Policy

**Note**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an administrator. Then, you must start Group Policy by using the **Active Directory Users and Computers** snap-in.

- If the **Allows users to connect remotely using Terminal Services** Group Policy setting is set to **Not Configured**, the **Enable Remote Desktop on this computer** setting (on the **Remote** tab of the **System Properties** dialog box) on the target computers takes precedence. Otherwise, the **Allows users to connect remotely using Terminal Services** Group Policy setting takes precedence.

- Be aware of the security implications of remote logons. Users who log on remotely can perform tasks as though they were sitting at the console. For this reason, you should ensure that the server is behind a firewall. For more information, see VPN servers and firewall configuration and Security information for IPSec.

- You should require all users who make remote connections to use a strong password. For more information, see Strong passwords.

- Remote Desktop is disabled by default in Windows Server 2003 operating systems.

**Using System Properties**

1. Open **System** in Control Panel.

2. On the **Remote** tab, select or clear the **Enable Remote Desktop on this computer** check box, and then click **OK**.

### Important

When you enable Remote Desktop on a computer, you enable the capability for other users and groups to log on remotely to the computer. However, you must also decide which users and groups should be able to log on remotely, and then manually add them to the Remote Desktop Users group. For more information, see Enabling users to connect remotely to the server and Add users to the Remote Desktop Users group.

### Note

- You must be logged on as a member of the Administrators group to enable or disable Remote Desktop.

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

- Any configuration set with Group Policy overrides the configuration set by using **System properties**, as described in this procedure.

- Be aware of the security implications of remote logons. Users who log on remotely can perform tasks as though they were sitting at the console. For this reason, you should ensure that the server is behind a firewall. For more information, see VPN servers and firewall configuration and Security information for IPSec.

- You should require all users who make remote connections to use a strong password. For more information, see Strong passwords.

- Remote Desktop is disabled by default in Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Getting the client software
Remote Desktop Connection
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Connect to Remote Desktop for Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to Remote Desktop for Administration**

- At the server
- At the remote computer

**At the server**

Be sure you have the appropriate permission to use Remote Desktop for Administration. You must be an Administrator, or you must be a member of the Remote Desktop Users group. To learn more about joining the Remote Desktop Users group, see Add users to the Remote Desktop Users group

1. Be sure **Remote Desktop for Administration** is enabled.
2. You must have the network computer name or IP address of the server.
3. You must not turn off the server.
4. Print this page and take these instructions to the remote computer.

### Notes

- To determine your network computer name, open **System** in Control Panel. On the **Computer Name** tab, locate the **Full computer name** item. Your computer name is the part of the text string to the left of the first period.
- To print this page, right-click anywhere on the text, and choose **Print** from the context menu.

**At the remote computer**

- From Windows Server 2003 family operating systems
- From an operating system other than a Windows Server 2003 family operating system

**From Windows Server 2003 family operating systems**

1. Connect to your company network by means of Local Area Network (LAN), modem, or Virtual Private Network (VPN) connection.
2. To open Remote Desktop Connection, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
3. Type in the name or IP address of the server you want to connect to, and then click **Connect.**

**From an operating system other than a Windows Server 2003 family operating system**

1. Log on to your Internet Service Provider (ISP).
2. Install the necessary client software from the Remote Desktop Connection Software Download page.
3. To open Remote Desktop Connection, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

4. Type in the name or IP address of the server you want to connect to, and then click **Connect**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Getting the client software
Install Remote Desktop Connection (32-bit computers)

# Add users to the Remote Desktop Users group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add users to the Remote Desktop Users group**

1. Open Computer Management.
2. In the console tree, click the **Local Users and Groups** node.
3. In the details pane, double-click the **Groups** folder.
4. Double-click **Remote Desktop Users**, and then click **Add...**.
5. On the **Select Users** dialog box, click **Locations...** to specify the search location.
6. Click **Object Types...** to specify the types of objects you want to search for.
7. Type the name you want to add in the **Enter the object names to select (examples):** box.
8. Click **Check Names**.
9. When the name is located, click **OK**.

## Notes

- By default, the Remote Desktop Users group is not populated. You must decide which users and groups should have permission to log on remotely, and then manually add them to the group.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Change the permissions a user or group has to a connection
Remove a user or group from permissions lists

# Share the client setup folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To share the client setup folder**

1. On your computer that is running one of the Windows Server 2003 family operating systems, open Windows Explorer.
2. Browse to the %Systemroot%\System32\Clients\Tsclient folder, where *systemroot* is the folder containing your Windows installation files (for example, C:\Windows\System32\Clients\Tsclient).
3. Right-click the \Tsclient folder, and then click **Sharing and Security...**.
4. Click **Share this folder.** The share name is displayed as **tsclient.** You can change the user limit by clicking **Allow this number of users** and choosing a number from the scroll box.

## Notes

- To share folders and drives, you must be logged on as a member of the Administrators, Server Operators, or Power Users group.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Remote Desktop Connection (32-bit computers)

# Install Remote Desktop Connection (32-bit computers)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Remote Desktop Connection (32-bit computers)**

1. On a computer running one of the Windows Server 2003 family operating systems, share the client setup folder.
2. From the computer running Windows 95, Windows 98, Windows NT 4.0 or Windows 2000, connect to the Local Area Network that contains the computer running one of the Windows Server 2003 family operating systems.
3. Click **Start**, and then click **Run**.
4. In **Open**, type the following:

   ```
   \\computername\Tsclient\Win32\Setup.exe
   ```

   Where *computername* is the network computer name of the computer running one of the Windows Server 2003 family operating systems. Click **OK**.
5. Follow the directions that appear on your screen.

### Notes

- To share folders and drives, you must be logged on as a member of the Administrators, Server Operators, or Power Users group.
- You can also install Remote Desktop Connection directly from the Microsoft Web Site.
- To determine your network computer name, open **System** in Control Panel. On the **Computer Name** tab, locate the **Full computer name** item. Your computer name is the part of the text string to the left of the first period.
- To use Microsoft IntelliMirror to deploy Remote Desktop Connection, share the Msrdpcli.msi file in the \Tsclient\Win32 folder.
- To use Group Policy to deploy Remote Desktop Connection, assign the program on a per-computer basis to a terminal server. For more information, see article Q324750, "HOW TO: Assign Software to a Specific Group By Using a Group Policy in Windows Server 2003" in the Microsoft Knowledge Base.
- When you use the Msrdpcli.msi file to install Remote Desktop Connection on computers that are not running one of the Windows Server 2003 family operating systems, the existing shortcut on the **Communications** menu is updated. In addition, the installation process adds a shortcut for Remote Desktop Connection to the **All Programs** menu. This also applies if you install the .msi file on a computer running a Windows XP operating system, which already has a version of Remote Desktop Connection installed by default.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Share the client setup folder
Enable or disable Remote Desktop

# Install Remote Desktop Web Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Remote Desktop Web Connection**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. Click **Application Server**, and then click **Details**.
4. Select **Internet Information Services**, and then click **Details**.
5. In the **Subcomponents of Internet Information Services** list, select **World Wide Web Service**, and then click **Details**.
6. In the **Subcomponents of World Wide Web Service** list, click the **Remote Desktop Web Connection** check box. Click **OK**. Click **OK** again, and then click **OK** a third time.
7. In the Windows Components wizard, click **Next**.
8. Verify that the RDP-Tcp connection is checked in the **Connection types:** box, and then click **Next** to begin the installation of the Remote Desktop Web Connection.

## Notes

- You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings may also prevent you from completing this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to another computer using Remote Desktop Web Connection

# Connect to another computer using Remote Desktop Web Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to another computer using Remote Desktop Web Connection**

1. Ensure that Remote Desktop Web Connection is installed and running on the Web server.

2. Ensure that your client computer has an active network connection and that the DNS server service (or other name resolution method) is functioning.

3. On your client computer, start Microsoft Internet Explorer.

4. In the **Address** box, type the Uniform Resource Locator (URL) for the home directory of the Web server hosting Remote Desktop Web Connection.

   The URL is "http://" followed by the Windows Networking name of your server, followed by the path of the directory containing the Remote Desktop Web Connection files (default = /Tsweb/). (Note the forward slash marks.) For example, if your Web server is registered with the DNS server as "Admin1", in the **Address** box you type: **http://admin1/tsweb/**, and then press **ENTER**. The Remote Desktop Web Connection page appears on the screen.

5. In **Server**, type the name or IP address of the remote computer to which you want to connect.

6. Optionally, specify the screen size and logon information for your connection.

7. Optionally, select the **Send logon information for this connection** checkbox and type your user name and domain.

8. Click **Connect**.

## Notes

- Remote Desktop Web Connection requires Internet Explorer 5 or higher.

- You may need to specify the fully qualified domain name of the computer you are connecting to if that computer is not part of your local area network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Setting up Remote Desktop Web Connection
Enable or disable Remote Desktop

# Connect to the console session of a server using Remote Desktop for Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to the console session of a server**

Using the Remote Desktops MMC Snap-in

Using the command line

**Using the Remote Desktops MMC Snap-in**

1. Open Remote Desktops snap-in.
2. If you have not already done so, create the connection to the terminal server or computer to which you want to connect.
3. In the console tree, right-click the connection.
4. In the context menu, click **Connect**.

### Note

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.

**Using the command line**

1. Open Command Prompt.
2. Type:

   **mstsc /console**.
3. Remote Desktop Connection will start. Type the computer name or IP address of the computer you want to connect to in the **Computer** box.
4. Configure any other desired options, and then click **Connect**.

### Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at a command prompt, type:

  **mstsc /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Enable or disable Remote Desktop
Terminal Services commands
Create a new connection with the Remote Desktops snap-in

# Change the icon associated with the shortcut to a connection file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the icon associated with the shortcut to a connection file**

1. Create a connection file (.rdp file), or use an existing .rdp file.
2. Open Windows Explorer and create a shortcut to the .rdp file.
3. Right-click the shortcut, click **Properties,** and then click **Change Icon.**
4. Select an icon in the list on the **Change Icon** dialog box, and then click **OK.**
5. Click **OK** to close the **Shortcut Properties** dialog box.
6. Deploy the .rdp file and the shortcut to the appropriate users.

## Notes

- You might also have to deploy the icon (.ico) file, if it is not available on the users' computers.
- The default connection file, Default.rdp, is stored as a hidden file in **My Documents**. Connection files that you create are also stored in **My Documents**, but they are not hidden.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a connection file
Save your connection settings to a file

# Remote Desktop for Administration Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Remote Desktop for Administration:

- Remote Desktop for Administration overview
- Using Remote Desktop for Administration for remote server administration
- Configuring Remote Desktop for Administration
- Deploying Remote Desktop for Administration

# Remote Desktop for Administration overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Desktop for Administration overview**
Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode) provides remote access to the desktop of computers running any Microsoft® Windows Server 2003 family operating system, allowing you to administer your server from virtually any computer on your network. For detailed information about using Remote Desktop for Administration for remote server administration, see Using Remote Desktop for Administration for remote server administration.

**Important**

- Be aware of the security implications of remote logons. Users who log on remotely can perform tasks as though they were sitting at the console. For this reason, you should ensure that the server is behind a firewall. For more information, see VPN servers and firewall configuration and Security information for IPSec.

- You should require all users who make remote connections to use Strong passwords.

To use Remote Desktop for Administration, you must enable remote connections at the server you plan to connect to. For more information, see Enable or disable Remote Desktop.

To connect to Remote Desktop for Administration from a remote computer, use Remote Desktop Connection. For more information, see Remote Desktop Connection.

**Notes**

- Remote Desktop for Administration is disabled by default in Windows Server 2003 family operating systems.

- By default, the Remote Desktop Users group is not populated. You must decide which users and groups you want to have permission to log on remotely, and then manually add them to the group. For more information, see Add users to the Remote Desktop Users group.

- Remote administration of servers with Remote Desktop for Administration is available on any computer running a Windows Server 2003 family operating system. A simpler version of Remote Desktop is also available on Microsoft Microsoft® Windows® XP Professional. For information on installing and configuring Terminal Server, see Terminal Server.

# Using Remote Desktop for Administration for remote server administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Remote Desktop for Administration for remote server administration**

Remote Desktop for Administration can greatly reduce the overhead associated with remote administration. Enabled by Terminal Services technology, Remote Desktop for Administration is specifically designed for server management. Therefore, it does not install the application-sharing and multiuser capabilities or the process scheduling of the full Terminal Server component (formerly called Terminal Services in Application Server mode). As a result, Remote Desktop for Administration can be used on an already busy server without noticeably affecting CPU performance, which makes it a convenient and efficient service for remote management.

Remote Desktop for Administration does not require you to purchase special licenses for client computers that access the server. It is not necessary to install Terminal Server Licensing when using Remote Desktop for Administration.

Administrators can also fully administer computers running Windows Server 2003 family operating systems from computers running earlier versions of Windows by installing Remote Desktop Connection.

**Notes**

- Remote Desktop for Administration is disabled by default in Windows Server 2003 family operating systems. For more information, see Enable or disable Remote Desktop.
- To connect to Remote Desktop for Administration from a remote computer, use Remote Desktop Connection. For more information, see Remote Desktop Connection.

**How Remote Desktop for Administration works**

You can use Remote Desktop for Administration to manage a network remotely using a configuration similar to the one shown in the following illustration.




Remote Desktop for Administration provides remote access to the server desktop by using the Terminal Services Remote Desktop Protocol (RDP) on port 3389. RDP transmits the user interface to the client session, and also transmits keyboard and mouse clicks from the client to the server. You can create up to two simultaneous remote connections. Each session you log on to is independent of other client sessions as well as the server console session. In essence, you can use Remote Desktop for Administration to log on to the server remotely as though you were logged on locally.

If you need to connect to the server console session remotely (for example, to access applications that direct only their user interface to the console), either use the Remote Desktops snap-in management tool or use Remote Desktop Connection from the command line. When you attempt to connect to the console session, whether remotely or locally, you will be notified if there is already another user connected to the console session. The notification message will be shown after your logon credentials are validated, and will include information about the user who is logged on to the console session, including user name, location of logon (local or remote), and the state of the session (in use, locked, or idle).

**Important**

- Be aware of the security implications of remote logons. Users who log on remotely can perform tasks as though they were sitting at the console. For this reason, you should ensure that the server is behind a firewall. For more information, see VPN servers and firewall configuration and Security information for IPSec.

- You should require all users who make remote connections to use a strong password. For more information, see Strong passwords.

The connection to Remote Desktop for Administration uses TCP/IP, either over an existing network connection or by remote access. A remote access server running one of the Windows Server 2003 family operating systems provides two different types of remote access connectivity:

- Network and Dial-up Connections
- Virtual private networking

The following illustration shows how you can connect to a computer running one of the Windows Server 2003 family operating systems from a remote location using remote access.




For more information, see Remote Access.

**Administering Windows Server 2003 family operating systems remotely**

After you are connected to a computer running a Windows Server 2003 family operating system, you can use Remote Desktop for Administration to remotely administer the server and its local computers. Remote Desktop for Administration gives you access to a variety of administrative tools you can use to configure and manage computers. Through a Terminal Services session, you can access Microsoft Management Console (MMC), Active Directory, Systems Management Server, network configuration tools, and most other administrative tools.

Remote Desktop for Administration is extremely useful because it provides remote access to most configuration settings, including Control Panel, which usually cannot be configured remotely. Also, using Remote Desktop for Administration can be particularly convenient for diagnosing a problem and testing multiple solutions quickly.

You can access the servers from anywhere in the world by using a wide-area network (WAN), a virtual private network (VPN), or a dial-up connection. You can start time-consuming batch administrative jobs (for example, tape backups), disconnect, and later reconnect to the corporate network to check progress.

Server application and operating system upgrades can be completed remotely as well as tasks that are not usually possible unless you are sitting at the console, such as domain controller promotion/demotion and disk defragmentation. Server file system tasks such as copying large files and virus scanning are much more efficient when performed within a Remote Desktop for Administration session, rather than using utilities that are executed from a client computer.

Administrative tasks are quicker and more intuitive than using command line utilities, although it is still possible to open a command shell.

**Note**

- For some third-party applications, pop-up messages cannot be seen from a Terminal Services session. This is because there is a different security context or desktop for the connected session that does not display the application's pop-up messages. The pop-up messages in these instances will go directly to the console. If you need to see these messages, connect to the console session using Remote Desktop Connection from the command line or the Remote Desktops MMC snap-in.

  For more information, see Connect to the console session of a server using Remote Desktop for Administration.

# Configuring Remote Desktop for Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Remote Desktop for Administration**
You can use the Terminal Services administrative tools to configure computers and users for Remote Desktop for Administration connections.

In addition to the standard administrative tools, Windows Server 2003 family operating systems include the following Terminal Services tools, Windows Server 2003 tool extensions, and commands for administering Terminal Services connections, computers, and users. While many of these tools are primarily intended for configuring Terminal Server for multiple users, they can also be used to configure Remote Desktop for Administration connections.

**Terminal Services Group Policies**
Use Terminal Services Group Policies to configure individual computers or groups of computers, or to assign policies for Remote Desktop for Administration users or user groups.

To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Terminal Services Manager**
Use Terminal Services Manager to manage and monitor users, sessions, and processes on any terminal server on the network.

- Display information about servers, sessions, users, and processes.
- Connect to and disconnect from sessions.
- Monitor sessions.
- Reset sessions.
- Send messages to users.
- Log off users.
- Terminate processes.

For more information on Terminal Services Manager, see Terminal Services Manager Overview.

To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Terminal Services Configuration**
During the installation of any Windows Server 2003 family operating system, a connection is configured for the Remote Desktop Connection (RDP) protocol. The connection provides the link clients use to log on to a session on the server for either Remote Desktop for Administration or Terminal Server connections. After installation, you can use Terminal Services Configuration (Tscc.msc) to change the properties of this connection on the local computer or to add new connections.

**Note**

- Many of these settings can also be set with Terminal Services Group Policy or Windows Management Instrumentation (WMI), in which case Terminal Services Configuration settings are overridden.

Use Terminal Services Configuration to:

- Name a connection.
- Specify a connection type.

- Specify a connection transport and transport properties.
- Set the maximum number of sessions allowed.
- Enable or disable logons through the connection.
- Set connection time-outs.
- Set the level of encryption.
- Set whether to disconnect broken connections.
- Enable or disable session remote control.
- Enable or disable automatic logons.
- Specify a program to run automatically when a user logs on.
- Override user profile settings for wallpaper.
- Set permissions on the connection.
- Set client device mapping and connection parameters.

For more information on Terminal Services Configuration, see Terminal Services Configuration Overview.

To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Active Directory Users and Computers and Local Users and Groups extensions**
Terminal Services extends Active Directory Users and Computers and Local Users and Groups to allow you to control Terminal Services features for each user. You can also use Terminal Services Group Policies to control these features, in which case settings you have specified might be overridden.

Use the Terminal Services extension to:

- Set the path to the Terminal Services user profile of each user.
- Enable or disable logons.
- Set time limits on session.
- Set whether to disconnect or reset a broken connection.
- Enable or disable remote control.
- Specify a program to run automatically when a user logs on.
- Automatically connect client drives and printers at logon.

For more information on the Terminal Services extensions, see Terminal Services Configuration Overview.

To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

Open Computer Management.

**System Monitor counters**
Terminal Services extends System Monitor by adding User and Session objects and their counters. You can use these objects and counters to monitor the resources that a user or session consumes. Terminal Services also adds counters to the Process and System objects.

For more information on Terminal Services performance monitoring, see Performance monitoring.

**Task Manager additional fields**
Terminal Services provides two additional fields, **ID** (session ID) and **User**, to Task Manager for monitoring and ending processes for all sessions. For more information, see Task Manager overview.

**Command line**
Use Terminal Services command-line functions to configure Terminal Services from the command prompt.

| Command | Function |
|---|---|
| change logon | Temporarily disables logons to a Terminal server. |
| change port | Changes COM port mappings for MS-DOS program compatibility. |
| change user | Changes the .ini file mapping for the current user. |
| Cprofile | Removes user-specific file associations from a user's profile. |
| dbgtrace | Enables and disables debug tracing. |
| flattemp | Enables or disables flat temporary directories. |
| logoff | Ends a client's session. |
| msg | Sends a message to one or more clients. |
| query process | Displays information about processes. |
| query session | Displays information about Terminal Services sessions. |
| query termserver | Lists Terminal Services servers on the network. |
| query user | Displays information about users logged on to the system. |
| register | Registers a program so that it has special execution characteristics. |
| reset session | Allows you to reset (delete) a session from the Terminal server. |
| shadow | Allows you to monitor or remotely control an active session of another user. |
| tscon | Connects to another existing Terminal Services session. |
| tsdiscon | Disconnects a client from a Terminal Services session. |
| tskill | Terminates a process. |
| tsprof | Copies user configuration and changes profile path. |
| tsshutdn | Shuts down a Terminal Services server. |

For more information on Terminal Services command-line functions, see Terminal Services commands.

**Terminal Services WMI Provider**
With the Terminal Services WMI provider, you can create customized scripts for configuring, managing, and querying terminal servers. For more information, see Configuring Terminal Services with WMI.

# Deploying Remote Desktop for Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying Remote Desktop for Administration**
This section covers:

- Getting the client software
- Enabling users to connect remotely to the server
- Deploying Remote Desktop Web Connection

# Getting the client software

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Getting the client software**

To set up your client computer to connect to Remote Desktop for Administration, you need the following:

- Access to the remote computer (the computer running a Windows Server 2003 family operating system) by means of Local Area Network (LAN), modem, or Virtual Private Network (VPN) connection.

- Remote Desktop Connection or Terminal Services client installed on the client computer. See the following table for instructions on installing the client software for your operating system.

  **Important**

  - You should install the latest version of Remote Desktop Connection. If you are using the Windows 2000 32-bit Terminal Services Client, you should upgrade to the version of Remote Desktop Connection that is provided with all Windows Server 2003 family operating systems.

| Operating System | Client version | How to access |
|---|---|---|
| Windows Server 2003 family operating systems (all versions) | **Remote Desktop Connection** installed by default | **Start/Programs/Accessories/Communications/Remote Desktop Connection** |
| Windows XP (all versions) | **Remote Desktop Connection** installed by default | **Start/Programs/Accessories/Communications/Remote Desktop Connection** |
| Microsoft Windows 2000 Professional | Not installed | See Install Remote Desktop Connection (32-bit computers) |
| Microsoft Windows 2000 Server | Not installed | See Install Remote Desktop Connection (32-bit computers) |
| Microsoft Windows 95 & Microsoft Windows 98 | Not installed | See Install Remote Desktop Connection (32-bit computers) |
| Windows NT 4.0 | Not installed | See Install Remote Desktop Connection (32-bit computers) |
| Microsoft Windows CE | Not installed | Go to the Microsoft Web Site and follow the directions to install the Windows CE client. |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Enabling users to connect remotely to the server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling users to connect remotely to the server**
To enable users to connect remotely to a terminal server, you must ensure that:

- Remote Desktop is enabled on the server. For instructions on how to enable Remote Desktop, see Enable or disable Remote Desktop.
- Users have the appropriate rights and permissions to log on remotely to the server.

To perform these two tasks, you must be logged on as a member of the Administrators group.

**Using the Remote Desktop Users group to grant access to a terminal server**
You can easily manage permissions and rights for a terminal server on a per-computer basis, by using the Remote Desktop Users group. The Remote Desktop Users group is one of the built-in users groups available when you install one of the Windows Server 2003 operating systems. Members of this group are able to log on remotely to a terminal server on which Remote Desktop is enabled.

By default, the Remote Desktop Users group is not populated. Therefore, you must decide which users and groups should have access to log on remotely to a terminal server, and then add them to this group. For instructions on how to add users to the Remote Desktop Users group, see Add users to the Remote Desktop Users group.

It is highly recommended that you use the Remote Desktop Users group to grant individuals access to terminal servers, rather than assigning the required permissions manually

**Caution**

If you alter the default permissions on the Remote Desktop Users group or remove this group, members of this group might lose the ability to log on remotely to terminal servers.

**Assigning Permissions Manually**
In some cases, it might be necessary to manage access to a terminal server on a per-connection basis and to manually customize rights and permissions. Keep in mind that if you do not use the Remote Desktop Users group to grant users access to log on to a terminal server, you must assign those users the same rights and permissions manually. The required rights and permissions are as follows:

| Right or permission | Comments |
|---|---|
| Allow log on through Terminal Services | This right determines which users or groups have permission to log on as a Terminal Services client.<br><br>Configure by navigating to **Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment**.<br><br>For configuration instructions, see Edit security settings on a Group Policy object. |

| | |
|---|---|
| User Access | • This permission type grants the following special permissions: Query Information, Logon, and Connect. These special permissions allow a user to:<br>• Log on to a session on the terminal server.<br>• Query information about a session.<br>• Send messages to other user sessions.<br>• Connect to another session.<br><br>Configure by using Terminal Services Configuration.<br><br>For configuration instructions, see Change the permissions a user or group has to a connection. |

You can also explicitly deny an individual user access to a terminal server by modifying that user's terminal server profile. Doing this is useful if you have granted a large group access to a terminal server, but want to make individual exceptions as needed. For instructions, see Deny a user permissions to log on to terminal servers.

**See Also**
**Concepts**
Managing Terminal Services users
Configuring Terminal Services with TSCC

# Deploying Remote Desktop Web Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying Remote Desktop Web Connection**
This section provides information about deploying Remote Desktop Web Connection. Remote Desktop Web Connection is a Web application consisting of an ActiveX component and sample connection pages. When you deploy Remote Desktop Web Connection on a Web server, you can provide client connectivity to terminal servers and other computers via Internet Explorer and TCP/IP.

This section provides general background information about deploying Remote Desktop Web Connection:

- About Remote Desktop Web Connection.
- Setting up Remote Desktop Web Connection.
- Providing for Client Security
- About the Sample Web Page
- Customizing Remote Desktop Web Connection

# About Remote Desktop Web Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**About Remote Desktop Web Connection**

Remote Desktop Web Connection is a Web application that consists of an ActiveX control, sample web pages, and other files. When deployed on a Web server, Remote Desktop Web Connection allows users to create a Terminal Services connection within Internet Explorer, even if a Terminal Services client program is not installed on the user's computer.

Remote Desktop Web Connection offers the following benefits:

- **Simple, straightforward deployment of a Terminal Services client.** With Remote Desktop Web Connection, deploying a Terminal Services connection can be as easy as sending a URL.
- **Support for roaming users.** Users who are away from their computers can use Remote Desktop Web Connection to gain secure access to their primary workstation or favorite terminal server from any computer running Windows and Microsoft Internet Explorer 5 or later.
- **A lowest common denominator cross-platform system.** Remote Desktop Web Connection may meet the needs of organizations that want identical clients on different platforms.
- **Delivery of extranet applications.** Corporations that want to deploy Terminal Services connections to vendors, suppliers, or customers can use Remote Desktop Web Connection to distribute them easily, inexpensively, and efficiently over the Internet.

**How Remote Desktop Web Connection works**

When a Web user accesses a Web page that contains the embedded Remote Desktop Web Connection ActiveX Client Control, the following sequence of events occurs:

- If necessary, Internet Explorer (IE), version 5 or later, downloads the .cab file pointed to by the CODEBASE parameter of the OBJECT tag. This is necessary in the following circumstances:
    - The computer that accessed the Web page does not have an installation of the full Remote Desktop Connection program or the Terminal Services Advanced Client with Web support.
    - The installed version of the .cab file on the client computer is older than the version on the Web page.
- After downloading, IE decompresses the .cab file and installs the ActiveX Client Control in IE's default location for downloaded controls on the target computer.
- The default connection page appears, asking the user for server, logon, and connection information.
- The Remote Desktop opens in the Web page. Depending on the parameters passed and the settings of the remote computer to which the user is trying to connect, the Windows logon screen appears.

For information on using Remote Desktop Web Connection from the client computer, see Connect to another computer using Remote Desktop Web Connection.

# Setting up Remote Desktop Web Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up Remote Desktop Web Connection**
Remote Desktop Web Connection is an optional World Wide Web Service component of Internet Information Services, which is included in Windows XP Professional and Windows Server 2003 family operating systems. Remote Desktop Web Connection is not installed by default, but must be installed using **Add or Remove programs**.

When you install Remote Desktop Web Connection, the files are copied by default to the **%systemroot%\Web\Tsweb** directory of your web server. The included sample default.htm page can be used as is, or you can modify it to meet the needs of your application.

On the client side, Remote Desktop Web Connection requires that the computer have a TCP/IP connection to the Internet or network, and be running Microsoft Internet Explorer 5 or later.

**Important**

- If you are setting up a Web site for yourself or for some non-Administrator users of a server running one of the Windows Server 2003 family operating systems, you must disable the Internet Explorer Enhanced Security Configuration setting for Internet Explorer. Otherwise, all server users, including Administrators, will be prevented from downloading files or ActiveX controls such as Remote Desktop Web Connection from the Web site. To change this setting, use **Add/Remove Windows Components** in **Add or Remove programs**. For more information, see Internet Explorer Enhanced Security Configuration.

For more information, see:

- Install Remote Desktop Web Connection.
- Connect to another computer using Remote Desktop Web Connection.
- About the Sample Web Page.

For information on customizing Remote Desktop Web Connection, see the Microsoft Platform SDK.

# Providing for Client Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providing for client security**

To allow clients to protect themselves from potentially untrustworthy servers, some properties of the Remote Desktop Web Connection ActiveX component object are restricted to certain Internet Explorer security zones. This means that when a Web user accesses the page from a computer in a less secure zone, these properties are disabled.

Some of these restricted properties are:

**StartProgram**
Specifies a program to start upon connection.

**WorkDir**
Specifies the working directory for the program specified in StartProgram.

**FullScreen**
Specifies whether the connection will be displayed as full screen or windowed mode.

Servers using Remote Desktop Web Connection can display these restricted properties in the following Internet Explorer security zones:

- My Computer
- Local Intranet
- Trusted Sites

Servers using Remote Desktop Web Connection cannot display restricted these properties in the following zones:

- Internet
- Restricted Sites

For more information about Remote Desktop ActiveX Control interfaces, see Remote Desktop Control ActiveX Interface (http://go.microsoft.com/fwlink/?LinkId=102079).

# About the Sample Web Page

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**About the Sample Web page**

When you install Remote Desktop Web Connection, a sample Web page, Default.htm, is copied to the installation directory on your computer. This page can be run as is, or it can be modified.

Default.htm is a default logon page that is designed to collect connection information from the user. This page contains a required field for the terminal server or remote computer to which the user wants to connect and optional fields for screen size and user logon information (username and domain). After this data is collected, it is passed to the ActiveX Client Control (Msrdp.ocx) to initiate a remote desktop session. The entire process works as follows:

1. If necessary, IE downloads the .cab file containing the ActiveX component and installs it on the user's computer.
2. The user types in the fully qualified domain name or IP address of the remote computer. Optionally, the user specifies a screen size for the connection. If no screen size is specified, the session opens in Full Screen mode (if the user is in a trusted zone. See Providing for Client Security for more information).
3. Optionally, the user supplies automatic logon information (username, domain) for the remote session.
4. When the user clicks the **Connect** button, Default.htm passes the parameters to the ActiveX Client Control.
5. The remote desktop opens in the browser window (or full screen).

When Default.htm opens for the first time in Internet Explorer, a dialog box appears and asks if the user wants to download the Msrdp.ocx ActiveX component. This happens only under the following conditions:

- The computer that accessed the Web page does not have an installation of the full Terminal Services Client with Web support.
- The installed version of the .cab file on the client computer is older than the version on the Default.htm page.

For security reasons, some options that are available in Remote Desktop Web Connection are restricted to certain Internet Explorer security zones when called by the Remote Desktop Web Connection ActiveX Client Control. See Providing for Client Security for more information.

The screen size specified by the user in this form is the size of the remote desktop upon connection. The remote desktop appears at the specified size in the browser window or, if the session opens in the default full-screen mode, the browser controls and borders disappear and the remote desktop expands to the size of the user's screen. As with the full Remote Desktop Connection program, the user can toggle between full-screen and windowed mode by pressing CTRL+ALT+BREAK.

The size of the remote desktop session is specified in Default.htm. To change the screen size, the user must return to Default.htm (after disconnecting the session) and edit the connection. If no screen size is specified, the remote desktop opens in the browser window in the default full-screen mode. At this resolution, the IE tool bars and window frames disappear, and the remote desktop expands to match the full dimensions of the user's screen. As with the full Remote Desktop Connection client, the user can toggle between full-screen and windowed mode by pressing CTRL+ALT+BREAK.

See Providing for Client Security for information on specifying a program to run upon connection.

# Customizing Remote Desktop Web Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Customizing Remote Desktop Web Connection**

The sample Default.htm page is installed with Remote Desktop Web Connection and can be used as is, or can be customized to fit the needs of your organization. Web authors can embed the Remote Desktop Web Connection ActiveX Control in a Web page with the HTML <OBJECT> tag, and can use the scriptable interfaces, for example, to specify a specific program to run upon connection, to allow seamless logon, and to specify advanced configuration options for the client connection.

If you have enabled a Virtual Channels application in your Terminal Services deployment, you can also make this application available to client computers that access the terminal server by means of the Remote Desktop Web Connection ActiveX component.

For a complete documentation of the full set of Remote Desktop Web Connection ActiveX Control parameters, see the Microsoft Platform SDK.

# Troubleshooting Remote Desktop for Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Client fails to connect with "This initial program cannot be started" message.
- Cannot disconnect active RAS connections with Terminal Services.
- The connection does not automatically log on to the server. I am prompted for my user name and password on connection.
- The Remote Desktop connection is very slow.
- The screen saver on my Remote Desktop is blank.

**Client fails to connect with "This initial program cannot be started" message.**
**Cause:** The client connection might have been configured to start a program upon connection, but the program path and filename are incorrect.

**Solution:** Check the properties for the client connection to ensure that the program path and filename are correct, and that the program exists in this location on the server.

**Cause:** The server administrator might have specified a program to start upon connection, using Group Policy, Terminal Services Configuration, or other Terminal Services administration tools.

**Solution:** Check the server configuration to ensure that the program path and filename specified for this user are correct, and that the program exists in this location on the server.

**Cannot disconnect active RAS connections with Terminal Services.**
**Cause:** When a RAS connection is established in a Terminal Services session, the current ID of the session is added to the port. When disconnecting the RAS connection, the current session ID must match the session ID that was added to the port, or it cannot be disconnected.

**Solution:** A RAS connection can only be disconnected from inside the session in which it was established.

**The connection does not automatically log on to the server. I am prompted for my user name and password on connection.**
**Cause:** Due to encryption enhancements for Windows Server 2003 family operating systems, the Windows NT 4.0 version of Terminal Services Client does not recognize the user name and password in the automatic logon portion of the connection file.

**Solution:** From the Windows NT 4.0 version of Client Connection Manager, right-click the connection and select **Properties**. On the **General** tab, select **Automatic logon**, and then type the user name, password, and domain.

**Solution:** Upgrade the client software to Remote Desktop Connection.

**Cause:** The server administrator has configured the server to deny automatic logon.

**Solution:** Type your user name and password when prompted.

**Cause:** User name and password information has not been added to Remote Desktop Connection.

**Solution:** In Remote Desktop Connection, on the **General** tab, type your user name and password, and then select the **Save my password** checkbox.

**The Remote Desktop connection is very slow.**
**Cause:** A low-bandwidth connection can slow Remote Desktop performance.

**Solution:** On the remote computer, set the background selection to **None.**

**Solution:** Do not run memory-intensive applications, for example, video games, over Remote Desktop.

**The screen saver on my Remote Desktop is blank.**

**Cause:** When a screen saver is activated on Remote Desktop, it is blank. This is true regardless of whether you have previously selected a different screen saver.

**Solution:** Continue to use the blank screen saver.

# Remote Desktops snap-in

Updated: May 9, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Desktops snap-in**

The Remote Desktops snap-in enables you to manage Remote Desktop connections to terminal servers and computers that are running Windows 2000 Server or Windows Server 2003 operating systems. The following topics provide information about this snap-in:

- For general background information, see Remote Desktops snap-in overview.
- For help with specific tasks, see Remote Desktops snap-in How To....

# Remote Desktops snap-in How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Add the Remote Desktops snap-in
- Create a new connection with the Remote Desktops snap-in
- Connect to a session
- Disconnect from a session
- Change an existing connection
- Delete a connection

# Add the Remote Desktops snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add the Remote Desktops snap-in**

1. Run MMC by doing one of the following:

   - Click **Start**, click **Run**, type `mmc`, and then click **OK.**
   - At a command prompt, type `mmc`, and then press ENTER.

2. On the **Console** menu, click **File,** and then click **Add/Remove Snap-in.**
3. In the dialog box that opens, click **Add.** In the **Available Standalone Snap-ins** list, click **Remote Desktops**.
4. Click **Add,** and then click **Close.**
5. Click **OK** to close the **Add/Remove Snap-in** dialog box.

**Note**

- You can also open the Remote Desktops snap-in from the Administrative Tools folder in Control Panel. To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Enable or disable Remote Desktop

# Create a new connection with the Remote Desktops snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new connection**

1. Open Remote Desktops snap-in.
2. In the console tree, right-click **Remote Desktops**.
3. Click **Add new connection...**.
4. In the **Add New Connection** dialog box, in **Server name or IP address**, type the name or TCP/IP address of the server. Click **Browse** to look for a server.
5. Optionally, in **Connection Name**, type a friendly name for this connection.
6. Uncheck **Connect to console** if you do not want to connect to the console session on the server.

   **Important**

   - A connection to the console session always provides a full desktop session, just as you would see if you were logged on to the server locally. If you want to start a specific program on connection to the server, you must uncheck **Connect to console** ands specify a program path and file name on the **Other** tab.

7. Optionally, In **Logon information**, supply the User name, Password and Domain for automatic logon to the server. Click **OK**.

**Notes**

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.
- Depending on the configuration of the Terminal server, you might be prompted to supply your password upon connection even if you specify it in the **Add New Connection** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a session
Disconnect from a session

# Connect to a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a session**

1. Open Remote Desktops snap-in.
2. If you have not already done so, create the connection to the terminal server or computer to which you want to connect.
3. In the console tree, right-click the connection.
4. On the context menu, click **Connect**.

### Note

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a new connection with the Remote Desktops snap-in
Disconnect from a session

# Disconnect from a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect from a session**

1. Open Remote Desktops snap-in.
2. In the console tree, right-click the connection.
3. On the context menu, click **Disconnect**.

## Notes

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.
- Your session remains active on the server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a new connection with the Remote Desktops snap-in
Connect to a session

# Change an existing connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Change an Existing Connection**

- Change the server to which to connect
- Change logon information
- Change client screen area
- Specify a program to start upon connection using Remote Desktops snap-in

# Change the server to which to connect

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the server to which to connect**

1. Open Remote Desktops snap-in.
2. In the console tree, right-click the connection you want to change.
3. In the context menu, click **Properties**.
4. On the **General** tab, in **Server name or IP address**, type the name or TCP/IP address of the server, or click **Browse** to look for a server.
5. Optionally, in **Connection name**, type a friendly name for this connection.
6. Optionally, In **Logon information**, check the checkbox for automatic logon, and then provide the User name, Domain, and Password. Click **OK**.

## Notes

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.
- Changes to the connection take effect the next time you connect to the server.
- Depending on the configuration of the computer to which you are connecting, you might be prompted to supply your password upon connection even if you specify it in the **Logon information** area.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a session
Disconnect from a session

# Change logon information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change logon information**

1. Open Remote Desktops snap-in.
2. In the console tree, right-click the connection you want to change.
3. In the context menu, click **Properties**.
4. On the **General** tab, under **Logon information**, type a new user name, password, and domain name, and then click **OK**.

### Note

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a session
Disconnect from a session

# Change client screen area

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change client screen area**

1. Open Remote Desktops snap-in.
2. If you have not already done so, create the connection.
3. In the console tree, right-click the connection you want to change.
4. In the context menu, click **Properties**.
5. On the **Screen Options** tab, choose the desktop size you want or enter a custom desktop size. The default is **Expand to fill MMC Result Pane.**

## Notes

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.
- Changes to the connection take effect the next time you connect to the server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a session
Disconnect from a session

# Specify a program to start upon connection using Remote Desktops snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify the program to start upon connection**

1. Open Remote Desktops snap-in.
2. If you have not already done so Create a new connection with the Remote Desktops snap-in.
3. In the console tree, right-click the connection you want to change.
4. In the context menu, click **Properties**.
5. On the **General** tab, uncheck **Connect to console**.

   **Important**

   - You must uncheck this option in order to have a specific program start on connection to the server.

6. On the **Other** tab, click **Start the following program on connection** and type the program path and file name. If you do not specify the program to start, the connection starts at the Windows desktop.
7. Check the **Redirect local drives when logged on to the remote computer** checkbox to allow the remote computer to access drives on the local computer.

## Notes

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.
- Changes to the connection take effect the next time you connect to the server.
- If you choose to redirect local drives when logged on to the remote computer, be aware of the security implications of redirection. For more information, see Making local resources available to a terminal server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a session
Disconnect from a session

# Delete a connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a connection**

1. Open Remote Desktops snap-in.
2. In the console tree, right-click the connection you want to delete.
3. In the context menu, click **Delete**.

### Note

- To open Remote Desktops, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Desktops**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a new connection with the Remote Desktops snap-in
Change an existing connection

# Remote Desktops snap-in overview

Updated: May 9, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Desktops snap-in overview**

The Remote Desktops snap-in is a tool that administrators can use to administer multiple computers from a single remote location. By using the Remote Desktops snap-in, you can create Remote Desktop connections to multiple servers (including terminal servers) that are running Windows 2000 Server or Windows Server 2003 operating systems. You can also create connections to computers that are running the Windows XP Professional operating system. A navigable tree display provides easy switching between connections.

The Remote Desktops snap-in is included with Windows Server 2003. It is also included with the Windows Server 2003 Administration Tools Pack (Adminpak.msi). The Administration Tools Pack provides server management tools that enable administrators to remotely manage Windows 2000 Server-based or Windows Server 2003-based servers from a Windows XP-based or Windows Server 2003-based computer. You can download the Windows Server 2003 Service Pack 2 (SP2) version of the Administration Tools Pack from the following Web sites:

- Windows Server 2003 SP2 Administration Tools Pack for x86 editions (http://go.microsoft.com/fwlink/?LinkId=114415)
- Windows Server 2003 SP2 Administration Tools Pack for x64 editions (http://go.microsoft.com/fwlink/?LinkId=114416)

By default, when you connect to a Windows Server 2003-based computer, the Remote Desktops snap-in connects you to the console session of the computer that you specify in the connection. To run a specific program upon connection, create a new connection that specifies the program name, and ensure that the default behavior is not selected. For more information, see Specify a program to start upon connection using Remote Desktops snap-in.

Most of the procedures that you can perform with the Remote Desktops snap-in require that you be logged on as an administrator or as a member of the Administrators group. If your computer is connected to a network, network policy settings may prevent you from completing a given procedure.

# Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Desktop Connection**
With Remote Desktop Connection, you can easily connect to a terminal server or to any computer running Remote Desktop. All you need is network access and permissions to connect to the other computer. Optionally, you can specify special settings for your connection and then save the settings for the next time you connect.

For help with specific tasks, see Remote Desktop Connection How To....

For general background information, see Remote Desktop Connection Concepts.

For troubleshooting information, see Troubleshooting Remote Desktop Connection.

**Notes**

- To enable your computer to accept remote connections, your computer must be running Windows NT4 Terminal Server Edition, Windows 2000 Server, Windows XP Professional, or a Windows Server 2003 operating system. Check Help for your operating platform for more information.

- To open **Remote Desktop Connection**, click **Start**, point to **programs** or **All programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

- To view a list of available computers on the domain, click the **Computer** arrow and select **<Browse for more...>**.

# Remote Desktop Connection How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Create a new connection with Remote Desktop Connection
- Reestablish a prior connection
- Save your connection settings to a file
- Open a saved connection
- Connect to the console session of a server using Remote Desktop Connection
- Disconnect without ending a session
- Log off and end the session
- Copy and paste a file from a local computer to a remote computer
- Copy and paste a file from a remote computer to a local computer
- Change Connection Settings

# Create a new connection with Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new connection**

1. Open **Remote Desktop Connection**.

2. In **Computer**, type a computer name or IP address. The computer can be a terminal server, or it can be a computer running Windows XP Professional or a Windows Server 2003 operating system that has Remote Desktop enabled and for which you have Remote Desktop permissions.

   To connect to the console session of the remote computer, in **Computer**, type:

   ***computername or IP address* /console**

3. Click **Connect**.

   The **Log On to Windows** dialog box appears.

4. In the **Log On to Windows** dialog box, type your user name, password, and domain (if required), and then click **OK**.

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **programs** or **All programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

- To change your connection settings, (such as screen size, automatic logon information, and performance options), click **Options** before you connect. For more information, see Change Connection Settings.

Related Topics

Reestablish a prior connection

Save your connection settings to a file

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Reestablish a prior connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reestablish a prior connection**

1. In the **Remote Desktop Connection** window, click the **Computer** arrow and select the name of the computer to which you want to connect.

2. Click **Connect**.

   The **Log On to Windows** dialog box appears.

3. In the **Log On to Windows** dialog box, type your user name, password, and domain (if required), and then click **OK**.

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

- To view a list of available computers on the domain, click the **Computer** arrow and select **<Browse for more...>**.

Related Topics

Create a new connection with Remote Desktop Connection

Save your connection settings to a file

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Save your connection settings to a file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save your connection settings to a file**

1. Open **Remote Desktop Connection**.
2. Click **Options**.
3. Specify the connection settings that you want for this connection, (such as screen size, automatic logon information, and performance options).
4. On the **General** tab, click **Save As**.
5. Enter a file name for the saved connection file, and then click **Save**.

## Notes

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- Connections are saved as Remote Desktop protocol (.rdp) files. An .rdp file contains all of the information for a connection to a terminal server, including the **Options** settings that were configured when the file was saved. You can customize any number of .rdp files, including files for connecting to the same server with different settings. For example, you can save a file that connects to *MyServer* in full screen mode and another file that connects to the same computer in 800×600 screen size. The default connection file, Default.rdp, is stored as a hidden file in **My Documents.** Connection files that you create are also stored in **My Documents,** but they are not hidden. To edit an .rdp file and change the connections settings it contains, right-click the file and then click **Edit**.
- For more information about specifying connection settings, see Change Connection Settings.

Related Topics

Create a new connection with Remote Desktop Connection

Open a saved connection

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Open a saved connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open a saved connection**

1. In Windows Explorer, open the **My Documents** folder.
2. Double-click the .rdp file for the connection you want to open.

**Note**

- A Remote Desktop protocol (.rdp) file contains all of the information for a connection to a remote computer, including the **Options** settings that were configured when the file was saved. You can customize any number of .rdp files, including files for connecting to the same server with different settings. For example, you can save a file that connects to *MyServer* in full screen mode and another file that connects to the same computer in 800×600 screen size. The default connection file, Default.rdp, is stored as a hidden file in **My Documents.** Connection files that you create are also stored in **My Documents,** but they are not hidden. To edit an .rdp file and change the connections settings it contains, right-click the file and then click **Edit**.

Related Topics

Create a new connection with Remote Desktop Connection

Save your connection settings to a file

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Connect to the console session of a server using Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to the console session of a server**

1. Open Command prompt.

2. Type:

   **mstsc/v:***server***/console**

| Value | Description |
|---|---|
| Server | The DNS name or IP address of the server to which you want to connect |

**Notes**

- To open command prompt, Click **Start**, point to **programs** or **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can also connect to the console session of a server by specifying "/console" when you type a server name or IP address in the **Computer** field on the **General** tab on the Remote Desktop Connection **Options** dialog box.

- Each computer has only one console session. When you connect to the console session remotely, other users may be unable to log on to the computer locally.

- Some client option settings are not applied when a connection is made to the console session of a remote computer running Windows XP. For example, if the client has specified a program to start upon connection to a remote computer, that program does not start if the connection is made to the console session of the remote computer. Instead, the default desktop appears. If the connection is made to a session other than the console session, the specified program starts as expected. The console session is the only session available for Remote Desktop on Windows XP Professional, so the default desktop always appears when a connection is made, even if the client has specified a different program.

- To view the complete syntax for this command, at a command prompt, type:

  **mstsc /?**

Related Topics

Create a new connection with Remote Desktop Connection

Reestablish a prior connection

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Disconnect without ending a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect without ending a session**

1. In the Remote Desktop Connection window, click **Start**, and then click **Shut Down**.

   The **Shut Down Windows** dialog box appears.

2. Click **Disconnect**, and then click **OK**.

## Notes

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- You can also disconnect from a session by closing the Remote Desktop Connection window.
- Remote Desktop Connection automatically reconnects to this session the next time you connect to this computer (if the remote computer is configured to allow reconnection of disconnected sessions).

Related Topics

Log off and end the session

Change Connection Settings

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Log off and end the session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To log off and end the session**

1. In the **Remote Desktop Connection** window, click **Start**, and then click **Shut Down**.

   The **Shut Down Windows** dialog box appears.

2. Click **Log Off <*username*>**, and then click **OK**.

**Note**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

Related Topics

Disconnect without ending a session

Remote Desktop Connection Concepts

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Copy and paste a file from a local computer to a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy and paste a file from a local computer to a remote computer**

1. Use Remote Desktop Connection to establish a connection to a remote computer.

2. Click **Start** on the task bar of the remote computer, and then click **My Computer**, or double-click the **My Computer** icon on the desktop of the remote computer.

   When Windows Explorer opens on the remote computer, it displays the drives on the remote computer. If you have chosen to make your local drives available in a session, Windows Explorer on the remote computer also displays the drives on your local computer. The following example shows how the drives on your local computer are named when they appear in Windows Explorer on the remote computer.

   ```
   <driveletter> on <computername>
   ```

   *tsclient* is the name that Remote Desktop Connection assigns to your local computer.

3. Right-click the file residing on your local computer that you want to transfer to the remote computer, and then click **Copy.**

4. Navigate in the same Windows Explorer window to the location on the remote computer where you want to paste the file.

5. Right-click the folder icon and then click **Paste.**

**Note**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

Related Topics

Copy and paste a file from a remote computer to a local computer

Cutting and pasting to and from remote sessions

Accessing local files and folders in a remote session

Make your local disk drives available in a session

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Copy and paste a file from a remote computer to a local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy and paste a file from a remote computer to a local computer**

1. Use Remote Desktop Connection to establish a connection to a remote computer.

2. Click **Start** on the task bar of the remote computer, and then click **My Computer**, or double-click the **My Computer** icon on the desktop of the remote computer.

   When Windows Explorer opens on the remote computer, it displays the drives on the remote computer. If you have chosen to make your local drives available in a session, Windows Explorer on the remote computer also displays the drives on your local computer. The following example shows how the drives on your local computer are named when they appear in Windows Explorer on the remote computer.

   ```
   <driveletter> on tsclient
   ```

   *tsclient* is the name that Remote Desktop Connection assigns to your local computer.

3. Right-click the file residing on the remote computer that you want to transfer to the local computer, and then click **Copy.**

4. Navigate in the same Windows Explorer window to the location on the local computer where you want to paste the file.

5. Right-click the folder icon and then click **Paste.**

**Note**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

Related Topics

Copy and paste a file from a local computer to a remote computer

Cutting and pasting to and from remote sessions

Accessing local files and folders in a remote session

Make your local disk drives available in a session

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Change Connection Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change connection settings**

- Provide information for automatic logon
- Reconnect if connection is dropped
- Change the screen size for connections
- Change the color settings for connections
- Improve the performance of connections
- Display the connection bar in full-screen mode
- Specify a program to start upon connection Remote Desktop Connection
- Configure connections to play audio
- Make your local disk drives available in a session
- Make your local printer available in a session
- Make your local serial port available in a session
- Make a smart card available in a session
- Configure Windows shortcut keys in a session
- Configure bitmap caching

# Provide information for automatic logon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To provide information for automatic logon**

1. In the **Remote Desktop Connection** window, click **Options**.
2. On the **General** tab, under **Logon settings**, in the **User name** box, verify that your user name is correct.
3. If required, in the **Domain** box, type the domain to automatically log on to.
4. Optionally, clear the **Save my password** check box, and then enter your password in the **Password** box.
5. Click **Connect**.

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- Depending on the policies of your network or the configuration of the terminal server to which you want to connect, you might be prompted to supply your password upon connection, even if you supply it for automatic logon.
- When you click **Connect**, the information you have entered in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Create a new connection with Remote Desktop Connection

Change Connection Settings

Save your connection settings to a file

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Reconnect if connection is dropped

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reconnect if connection is dropped**

1. Open **Remote Desktop Connection**.
2. Click **Options**
3. On the **Experience** tab, select the **Reconnect if connection is dropped** check box.

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- To change your connection settings, (such as screen size, automatic logon information, and performance options), click **Options** before you connect.

Related Topics

Reestablish a prior connection

Save your connection settings to a file

Change Connection Settings

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Change the screen size for connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the screen size for connections**

1. In the **Remote Desktop Connection** window, click **Options**.
2. On the **Display** tab, under **Remote desktop size**, drag the slider to choose the size of your remote desktop. Drag the slider all the way to the right for full screen.
3. Click **Connect**.

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- When you click **Connect**, the information you have entered in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Save your connection settings to a file

Using Terminal Services shortcut keys

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Change the color settings for connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the color settings for connections**

1. In the **Remote Desktop Connection** window, click **Options**.
2. On the **Display** tab, in the **Colors** list, click the color resolution you want.
3. Click **Connect**.

## Notes

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- Depending on the policies of your network or the configuration of the terminal server to which you want to connect, the maximum color resolution for your remote connection might be limited. If the color resolution you select is higher than the allowable setting, Remote Desktop Connection creates the connection at the maximum allowable color resolution.
- When you click **Connect**, the information you have entered in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Save your connection settings to a file

Change Connection Settings

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Improve the performance of connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To improve the performance of connections**

1. In the **Remote Desktop Connection** dialog box, click **Options**, and then click the **Experience** tab.

2. In the **Performance** box, choose the speed of your connection.

   To improve the performance of the connection, you can allow certain characteristics of the remote Windows session to appear if they are enabled on the remote computer. The characteristics that can be changed depend on the speed of your connection.

   If you choose **Custom**, you can check any combination of boxes. Your choices are saved in a connection (.rdp) file.

## Notes

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

- The default connection speed is **Modem (56 Kbps).** This speed offers very good performance for most networks. Use the faster speed settings to enable richer graphical features, such as desktop wallpaper or menu sliding and fading. To select a combination of individual effects, use the **Custom** setting.

- Depending on the speed of your connection, you can specify the following characteristics:

  - Desktop background

  - Show window contents while dragging

  - Menu fading and sliding

  - Themes

  - Bitmap caching

Related Topics

Save your connection settings to a file

Configure bitmap caching

Change Connection Settings

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Display the connection bar in full-screen mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To display the connection bar in full-screen mode**

1. In the **Remote Desktop Connection** window, click **Options.**
2. On the **Display** tab, check the **Display the connection bar when in full screen mode** box.
3. Click **Connect.**

## Notes

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- For the first five seconds of a remote connection, the connection bar is displayed so you can verify that you are connected to the correct remote computer. This occurs even though you have chosen not to display the connection bar.
- When you click **Connect**, the information you have entered in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp. To create connections with different or connection-specific options, see Save your connection settings to a file.

Related Topics

Create a new connection with Remote Desktop Connection

Change Connection Settings

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Specify a program to start upon connection Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify a program to start upon connection**

1. In the **Remote Desktop Connection** window, click **Options**.
2. On the **programs** tab, under **Start a program**, click **Start the following program on connection**.
3. In the **program path and file name** box, type the path and file name of the program that you want to run after the connection is created.
4. Optionally, in the **Start in the following folder:** box, type the path to the working directory of the program.
5. Click **Connect**.

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **programs** or **All programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- This option works only if you are connecting to a terminal server.
- When the connection is made, only the specified program runs on the remote computer.
- Some client option settings are not applied when a connection is made to the console session of a remote computer running Windows XP. For example, if the client has specified a program to start upon connection to a remote computer, that program does not start if the connection is made to the console session of the remote computer. Instead, the default desktop appears. If the connection is made to a session other than the console session, the specified program starts as expected. The console session is the only session available for Remote Desktop on Windows XP Professional, so the default desktop always appears when a connection is made, even if the client has specified a different program.
- There are two ways to end a session.
    - To end the session and close the program, use the program's **Exit** command.
    - To disconnect the session and leave the program running on the remote computer, close the Remote Desktop Connection window.
- To specify a path, type the drive letter followed by a colon (:) and a backslash (\), type the names of the folders and subfolders that contain the file, separating each name with a backslash, and then type the name of the file or folder you want to open.
- When you click **Connect**, the information you have entered in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Save your connection settings to a file

Change Connection Settings

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure connections to play audio

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure connections to play audio**

1. In the **Remote Desktop Connection** window, click **Options**.
2. On the **Local Resources** tab, in the **Remote computer sound** list, click the option you want.
   - To run sound files in your Remote Desktop session and hear them on your local computer, choose **Bring to this computer**.
   - To run sound files in your Remote Desktop session and play them only on the remote computer, choose **Leave at remote computer**.
   - To disable all sound in Remote Desktop sessions, choose **Do not play**.
3. Click **Connect**.

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- Depending on the policies of your network, audio redirection might be disabled for some or all remote connections. Additionally, audio redirection does not work when you connect to a server running Windows 2000 or earlier.
- When you click **Connect**, the information you have entered in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Save your connection settings to a file

Change Connection Settings

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Make your local disk drives available in a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make your local disk drives available in a session**

1. In the **Remote Desktop Connection** window, click **Options.**
2. On the **Local Resources** tab, Under **Local devices**, click **Disk drives.**
3. Click **Connect.**

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- Mapping local drives is not available on computers running Windows 2000 or earlier versions of Windows.
- Depending on the policies of your network, local drive mapping might be disabled for some or all remote connections.
- You might see a Remote Desktop Connection Security Warning, which asks whether you want to connect your local disk drives, local ports, or smart card to the remote computer. To enable one or more of these connections, check the appropriate check boxes. If you choose not to enable these connections, a Remote Desktop Connection is still established, but these resources are not available to the remote computer.
- When you click **Connect**, the information you enter in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Save your connection settings to a file

Change Connection Settings

Making local resources available to a terminal server

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Make your local printer available in a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make your local printer available in a session**

1. In the **Remote Desktop Connection** window, click **Options**.
2. On the **Local Resources** tab, under **Local devices**, click **Printers**.
3. Click **Connect**.

## Notes

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- Depending on the policies of your network, local printer mapping might be disabled for some or all remote connections. However, you can print to network printers.
- When you click **Connect**, the information you have entered in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Save your connection settings to a file

Change Connection Settings

Making local resources available to a terminal server

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Make your local serial port available in a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make your local serial port available in a session**

1. In the **Remote Desktop Connection** window, click **Options.**
2. On the **Local Resources** tab, under **Local devices**, click **Serial ports.**
3. Click **Connect.**

## Notes

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- Depending on the policies of your network, local serial port mapping might be disabled for some or all remote connections.
- You might see a Remote Desktop Connection Security Warning, which asks whether you want to connect your local disk drives, local ports, or smart card to the remote computer. To enable one or more of these connections, check the appropriate check boxes. If you choose not to enable these connections, a Remote Desktop Connection is still established, but these resources are not available to the remote computer.
- When you click **Connect**, the information you have entered in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Save your connection settings to a file

Change Connection Settings

Making local resources available to a terminal server

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Make a smart card available in a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a smart card available in a session**

1. In the **Remote Desktop Connection** window, click **Options.**
2. On the **Local Resources** tab, Under **Local devices**, click **Smart card.**

   **Note**

   - The option to enable smart card redirection is not displayed unless a smart card is present and the service is enabled.

3. Click **Connect.**

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection.**
- Depending on the policies of your network, smart card redirection might be disabled for some or all remote connections.
- You might see a Remote Desktop Connection Security Warning, which asks whether you want to connect your local disk drives, local ports, or smart card to the remote computer. To enable one or more of these connections, check the appropriate check boxes. If you choose not to enable these connections, a Remote Desktop Connection is still established, but these resources are not available to the remote computer.
- When you click **Connect**, the information you enter in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Save your connection settings to a file

Change Connection Settings

Using a smart card to log on to a terminal server

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure Windows shortcut keys in a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Windows shortcut keys in a session**

1. In the **Remote Desktop Connection** window, click **Options**.
2. On the **Local Resources** tab, under **Keyboard**, click the **Apply Windows key combinations** list, and then click the option you want. This option affects the behavior of Windows shortcut key combinations (such as **ALT+TAB**) when you are connected to a remote computer.
   - To configure your connection so that Windows shortcut keys always apply to your local desktop, choose **On the local computer**.
   - To configure your connection so that all Windows shortcut keys apply to the desktop of the remote computer, choose **On the remote computer**.
   - To configure your connection so that Windows shortcut keys apply to the remote computer only when the connection is in full screen mode, choose **In full screen mode only**.
3. Click **Connect**.

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- When you click **Connect**, the information you have entered in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.
- This procedure does not apply if you are using Remote Desktop Connection on a computer running Windows 95, Windows 98, or Windows Millennium Edition.
- CTRL+ALT+DEL always goes to the local desktop. If you need to send CTRL+ALT+DEL to the remote computer, you must use CTRL+ALT+END.
- To toggle between full screen and non-full screen mode during a Terminal Services session, use CTRL+ALT+BREAK.

Related Topics

Using Terminal Services shortcut keys

Save your connection settings to a file

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure bitmap caching

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure bitmap caching**

1. In the **Remote Desktop Connection** window, click **Options**.
2. On the **Experience** tab, verify that the **Bitmap caching** check box is selected. Or, to disable bitmap caching, clear the **Bitmap caching** check box.
3. Click **Connect**.

**Notes**

- To open **Remote Desktop Connection**, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- Bitmap caching can speed up your connection by storing frequently used images on your local hard disk. By default, bitmap caching is enabled for Remote Desktop connections. It is recommended that you use bitmap caching unless instructed otherwise by your network administrator.
- When you click **Connect**, the information you enter in **Options** becomes the default setting for all Remote Desktop connections and is saved in a file called Default.rdp.

Related Topics

Change Connection Settings

Save your connection settings to a file

Improve the performance of connections

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Remote Desktop Connection Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

- Cutting and pasting to and from remote sessions
- Making local resources available to a terminal server
- Using a smart card to log on to a terminal server
- Accessing local files and folders in a remote session
- Printing to your local printer from a remote session
- Using Terminal Services shortcut keys

# Cutting and pasting to and from remote sessions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cutting and pasting to and from remote sessions**
Many Remote Desktop and Terminal Server connections provide clipboard sharing, allowing you to cut and paste from programs running in your Remote Desktop session to programs running on your local computer.

When you cut or copy information from a program, it is moved to the Clipboard and remains there until you clear the Clipboard or until you cut or copy another piece of information. The Clipboard window in ClipBook Viewer (Clipbrd.exe) shows the contents of the Clipboard. You can paste the information from the Clipboard into any document as often as you like. However, the information is only stored on the Clipboard temporarily.

The shared clipboard synchronizes its contents with the local clipboard. You can copy and paste text or graphics from a document within the Remote Desktop Connection window, and paste it into a document on your local computer. For information about copying, pasting, and using local files and folders in a Remote Desktop session, see Accessing local files and folders in a remote session.

**Note**

- Depending on the policies of your network and the version of Windows that is running on the remote computer, clipboard sharing might not work in some instances.

# Making local resources available to a terminal server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Making local resources available to a terminal server**

Depending on the Group Policy settings on the network, you can choose whether to allow a terminal server to have access to the disk drives, serial ports, printers, or smart card on your local computer. This is called resource redirection. Unless Group Policy settings prohibit resource redirection, you can redirect local resources to the terminal server by checking the appropriate checkbox in the Local devices section on the Local Resources tab in Remote Desktop Connection.

Making these resources available to a terminal server means that they can be used by that terminal server for the duration of the session. For example, suppose you choose to make your local disk drive available to the terminal server. While this allows you to easily copy files to or from the terminal server, it also means that the terminal server has access to the contents of your local disk drive. If there are times when this would not be appropriate, you can uncheck the appropriate checkbox to keep your local disk drive, or any other local resource, from being redirected to the terminal server.

By default, resource redirection for disk drives, serial ports, and smart cards is turned off. To turn on resource redirection, see Make your local serial port available in a session, Make your local disk drives available in a session, and Make a smart card available in a session.

# Using a smart card to log on to a terminal server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using a smart card to log on to a terminal server**

Smart cards are the most secure way to log on to a computer that is connected to a network. Because smart cards are tamper-resistant, they offer much stronger security than password-based logons can provide.

Many people choose passwords they can remember easily, which makes password-based logons weak and open to attack. If an unauthorized individual obtains your password, that individual can use your identity to connect to a network. However, to log on to a network by using your smart card, an unauthorized individual would have to obtain both the smart card and your personal identification number (PIN), which is far less likely.

Using a smart card to log on to a terminal server verifies your identity to the terminal server and allows you to access the programs you need to do your work. You can also make the smart card available to the terminal server for the duration of the session, in the same way that you can make your local disk drives available to the terminal server. However, the terminal server cannot access any private information stored on the smart card without your consent.

For information on making a smart card available to a terminal server, see Make a smart card available in a session. The option to enable smart card redirection is not displayed unless a smart card is present and the service is enabled.

# Accessing local files and folders in a remote session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing local files and folders in a remote session**

You can configure Remote Desktop Connection to make your local drives available when you connect to a remote computer.

Within a Remote Desktop Connection session, your local drives appear in Windows Explorer with the following designation:

```
<driveletter> on <computername>
```

To access local drives from a command line or a run line in a remote session, type:

```
\\tsclient\<driveletter>
```

where *driveletter* is the drive letter (without a colon) of the local drive.

**Notes**

- Depending on your network policies, automatic drive mapping might be disabled for some or all computers on your network. To ensure that Remote Desktop Connection is configured to map local drives, see Make your local disk drives available in a session.

- Automatic drive mapping is not available when you are connecting to a terminal server running Windows 2000 or earlier.

# Printing to your local printer from a remote session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing to your local printer from a remote session**
Printer redirection routes printing jobs from the terminal server or Remote Desktop computer to a printer attached to your local computer (also called the "client computer"). There are two ways to provide access to local printers: automatic and manual printer redirection. Use manual redirection when your local printer requires a driver that is not available on the version of Windows that is running on the remote computer.

**Automatic printer redirection**
Printer redirection is automatic when the local printer uses a driver that is installed on the server. When you log on to a session on a terminal server, or to a computer running Windows Professional or Server and Remote Desktop, any printers that are installed on the client (local) computer, as well as network printers, are automatically detected and a local queue is created on the server. The client computer printer settings for the default printer and some properties (such as printing on both sides of the page) are used by the server.

When you disconnect or end the session, the printer queue is deleted, and any incomplete or pending print jobs are lost. Information about the client's local printers and settings are saved on the client computer. On subsequent logons, the printer queue is created using the information stored on the client computer.

If a printer driver is not found on the server, an event is logged and the client printer is not created. To make the printer available, the driver must be manually installed on the server.

**Manual printer redirection**
Printers attached to LPT and COM ports on the client (local) computer can be manually redirected, although manual redirection of printers connected through USB ports is not supported.

To manually redirect a client printer, contact your administrator and provide the name of your computer (or IP address for a Windows-based Terminal). The client must be connected to the remote computer during manual redirection.

After the initial manual redirection, printers are automatically redirected during subsequent logons.

### Notes

- Redirected printers are available for use with applications running on the server. Redirected printers appear in the Printers and Faxes folder in Control Panel. For more information on redirecting a local printer, see Make your local printer available in a session.

- When you disconnect or log off from a session, the printer queue is deleted, and incomplete or pending print jobs are lost.

# Using Terminal Services shortcut keys

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Terminal Services shortcut keys**

If you choose not to apply Windows key combinations to your Terminal Services sessions, you can use these shortcut keys to perform many of the same functions. For more information about applying Windows key combinations, see Configure Windows shortcut keys in a session.

| Shortcut key | Description |
|---|---|
| ALT+PAGE UP | Switches between programs from left to right. |
| ALT+PAGE DOWN | Switches between programs from right to left. |
| ALT+INSERT | Cycles through the programs in the order they were started. |
| ALT+HOME | Displays the **Start** menu. |
| CTRL+ALT+BREAK | Switches the client between a window and full screen. |
| CTRL+ALT+END | Brings up the **Windows Security** dialog box. |
| ALT+DELETE | Displays the **Windows** menu. |
| CTRL+ALT+Minus (-) symbol on the numeric keypad | Places a shadow copy of the active window within the client, on the Terminal server clipboard, which provides the same functionality as pressing PrintScrn on a local computer. |
| CTRL+ALT+Plus (+) symbol on the numeric keypad | Places a shadow copy of the entire client window area on the Terminal server clipboard, which provides the same functionality as pressing ALT+PrintScrn on a local computer. |

**Notes**

- On an NEC98, these shortcut keys are different in two cases:
  - CTRL+ALT+BREAK is replaced by CTRL+GRPH+VF2.
  - CTRL+ALT+END is replaced by CTRL+GRPH+VF1.
- CTRL+ALT+BREAK and CTRL+ALT+END are available in all Remote Desktop Connection sessions, even if you choose to apply Windows key combinations on the remote computer.

# Troubleshooting Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

- Troubleshooting General Remote Desktop Error Messages
- Troubleshooting Remote Desktop Licensing Error Messages
- Troubleshooting Specific Remote Desktop Problems

# Troubleshooting General Remote Desktop Error Messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting General Error Messages**
Which message are you getting?

- Because of an error in data encryption, this session will end. Please try connecting to the remote computer again.
- The remote connection has timed out. Please try connecting to the remote computer again.
- The remote session was disconnected because the total logon time limit was reached. This limit is set by the server administrator or by network policies.
- The specified computer name contains invalid characters. Please verify the name and try again.
- The specified remote computer could not be found. Verify that you have typed the correct computer name or IP address, and then try connecting again.
- The remote session to the remote computer was ended by means of an administration tool. Your administrator might have ended your connection.
- The remote session was disconnected because your session was logged off at the remote computer. Your administrator or another user might have ended your session.
- The remote session was disconnected because another user has connected to the session.
- The remote session was disconnected because the total logon time limit was reached. This limit is set by the server administrator or by network policies.
- The local computer is low on memory. Close some programs, and then connect to the remote computer again.
- The connection was ended because of a network error. Please try connecting to the remote computer again.
- The client could not connect. You cannot connect to the console from a console session of the same computer.
- The client could not connect to the remote computer. Remote connections might not be enabled or the computer might be too busy to accept new connections. It is also possible that network problems are preventing your connection. Please try your connection again later. If the problem continues to occur, contact your administrator.
- Error: Out of memory. The remote session will be disconnected. Close some programs on the local computer, and then try connecting to the remote computer again.
- An internal state error has occurred. The remote session will be disconnected. Your local computer might be low on memory. Close some programs, and then try connecting to the remote computer again.
- Because of a protocol error, this session will be disconnected. Please try connecting to the remote computer again.
- Because of a protocol error (error_code), the remote session will be disconnected. Please try connecting to the remote computer again.
- Because of a security error, the client could not connect to the remote computer. Verify that you are logged on to the network, and then try connecting again.

- The remote session was disconnected because the remote computer is low on memory.
- Because of an unknown error, the remote session will be disconnected. The local computer might be low on memory. Close some programs, and then try connecting to the remote computer again.
- An unrecoverable error occurred while the client was attempting to connect. Close some programs, and then try connecting to the remote computer again.
- An error occurred when the client attempted to connect to the remote computer. Check your system memory and then try the connection again.
- Unable to allocate window resources. The remote session will be disconnected. Close some programs on the local computer, and then try connecting to the remote computer again.

**Because of an error in data encryption, this session will end. Please try connecting to the remote computer again.**
**Cause:** Data encryption provides security for data traveling over network connections. A data encryption error can end a session for security reasons.

**Solution:** Try connecting to the remote computer again. Encryption policies can vary between systems and domains. If you continue to see this error message, contact the server administrator.

**The remote connection has timed out. Please try connecting to the remote computer again.**
**Cause:** The terminal server disconnected your session because you did not send a response within the time limit set for responses.

**Solution:** Try connecting to the terminal server again. If you receive this error message, try to reconnect at a later time. If you continue to receive this error message, contact the server administrator.

**Cause:** A slow network connection, caused by excessive network traffic, has delayed your response to the terminal server.

**Solution:** Try connecting to the terminal server again. If you receive this error message, try to reconnect at a later time. If you continue to receive this error message, contact the server administrator.

**The remote session was disconnected because the total logon time limit was reached. This limit is set by the server administrator or by network policies.**
**Cause:** Your connection to the terminal server exceeded the time limit.

**Solution:** Try connecting to the terminal server again. If you receive this error message, try to reconnect at a later time. If you continue to receive this error message, contact the server administrator.

**The specified computer name contains invalid characters. Please verify the name and try again.**
**Cause:** The name of the remote computer is incorrect. This might be a typing error.

**Solution:** Try typing the name of the remote computer again. If you receive the same message, contact the server administrator to be sure you are using the correct name for the remote computer.

**The specified remote computer could not be found. Verify that you have typed the correct computer name or IP address, and then try connecting again.**
**Cause:** The name or IP address of the remote computer is incorrect. This might be a typing error.

**Solution:** Try typing the name or IP address of the remote computer again. If you receive the same message, contact the server administrator to be sure you are using the correct name or IP address for the remote computer.

**The remote session to the remote computer was ended by means of an administration tool. Your administrator might have ended your connection.**
**Cause:** The most common cause of this error is that the server administrator needed to perform maintenance tasks on the terminal server. These tasks can be performed only when there are no remote connections to the terminal server.

**Solution:** Try to reconnect at a later time. If you receive the same message, contact the server administrator.

**The remote session was disconnected because your session was logged off at the remote computer. Your administrator or another user might have ended your session.**
**Cause:** The most common cause of this error is that the server administrator needed to perform maintenance tasks on the terminal server. These tasks can be performed only when there are no remote connections to the terminal server.

**Solution:** Try to reconnect at a later time. If you receive the same message, contact the server administrator.

**The remote session was disconnected because another user has connected to the session.**
**Cause:** The most common cause of this error is that the server administrator needed to perform maintenance tasks on the terminal server. These tasks can be performed only when there are no remote connections to the terminal server.

**Solution:** Try to reconnect at a later time. If you receive the same message, contact the server administrator.

**The remote session was disconnected because the total logon time limit was reached. This limit is set by the server administrator or by network policies.**
**Cause:** Your connection to the terminal server exceeded the time limit set.

**Solution:** Try connecting to the remote computer again. If you receive the same message, try to reconnect at a later time. If you still have problems connecting, contact the server administrator.

**The local computer is low on memory. Close some programs, and then connect to the remote computer again.**
**Cause:** Not enough RAM available on your computer. If your computer has too little RAM available, it cannot free enough processing capacity to start new functions, such as applications or connections.

**Solution:** Close all unnecessary programs, and try your connection again.

**Cause:** Not enough disk space available on your hard drive. If the hard drive on your computer is full, there might not be enough space available to allow the swap file to operate correctly. The swap file allows your computer to perform some RAM functions even when there is not enough RAM space available.

**Solution:** Free some space on your hard drive by:

- Emptying your recycle bin.
- Deleting temporary Internet files.
- Moving documents and other data files to a different drive.
- Removing programs that you rarely use.
- Deleting old temporary files from your hard disk. Temporary files have a .tmp file name extension, or names that begin with a tilde (~), and are usually found in the C:\Temp or C:\Windows\Temp folder.
- Emptying your recycle bin after you delete files.

**The connection was ended because of a network error. Please try connecting to the remote computer again.**
**Cause:** A network error prevented your computer from communicating with the terminal server.

**Solution:** Try connecting to the remote computer again. If you receive the same message, be sure that your computer is connected to the network. If you are still unable to connect to the remote computer, contact the server administrator.

**The client could not connect. You cannot connect to the console from a console session of the same computer.**
**Cause:** You are logged on to a computer and you are trying to connect to that same computer.

**Solution:** If you are not on the same computer, contact the server administrator.

**The client could not connect to the remote computer. Remote connections might not be enabled or the computer might be too busy to accept new connections. It is also possible that network problems are preventing your connection. Please try your connection again later. If the problem continues to occur, contact your administrator.**
**Cause:** The remote computer might not be set up to accept remote connections.

**Solution:** Try connecting to the remote computer at a later time. If you receive the same message, contact the server administrator.

**Cause:** The remote computer might already have the maximum number of remote connections it is set up to accept.

**Solution:** Try connecting to the remote computer at a later time. If you receive the same message, contact the server administrator.

**Cause:** Network problems might be preventing the connection.

**Solution:** Try connecting to the remote computer at a later time. If you receive the same message, contact the server administrator.

**Error: Out of memory. The remote session will be disconnected. Close some programs on the local computer, and then try connecting to the remote computer again.**
**Cause:** Not enough RAM available on your computer. If your computer has too little RAM available, it cannot free enough processing capacity to start new functions, such as applications or connections.

**Solution:** Close all unnecessary programs, and try your connection again.

**Cause:** Not enough disk space available on your hard drive. If the hard drive on your computer is full, there might not be enough space available to allow the swap file to operate correctly. The swap file allows your computer to perform some RAM functions even when there is not enough RAM space available.

**Solution:** Free some space on your hard drive by:

- Emptying your recycle bin.
- Deleting temporary Internet files.
- Moving documents and other data files to a different drive.
- Removing programs that you rarely use.
- Deleting old temporary files from your hard disk. Temporary files have a .tmp file name extension, or names that begin with a tilde (~), and are usually found in the C:\Temp or C:\Windows\Temp folder.
- Emptying your recycle bin after you delete files.

**An internal state error has occurred. The remote session will be disconnected. Your local computer might be low on memory. Close some programs, and then try connecting to the remote computer again.**
**Cause:** Not enough RAM available on your computer. If your computer has too little RAM available, it cannot free enough processing capacity to start new functions, such as applications or connections.

**Solution:** Close all unnecessary programs, and try your connection again.

**Cause:** Not enough disk space available on your hard drive. If the hard drive on your computer is full, there might not be enough space available to allow the swap file to operate correctly. The swap file allows your computer to perform some RAM functions even when there is not enough RAM space available.

**Solution:** Free some space on your hard drive by:

- Emptying your recycle bin.
- Deleting temporary Internet files.
- Moving documents and other data files to a different drive.
- Removing programs that you rarely use.
- Deleting old temporary files from your hard disk. Temporary files have a .tmp file name extension, or names that begin with a tilde (~), and are usually found in the C:\Temp or C:\Windows\Temp folder.
- Emptying your recycle bin after you delete files.

**Because of a protocol error, this session will be disconnected. Please try connecting to the remote computer again.**
**Cause:** A protocol specifies the way a computer communicates with other computers. A protocol error prevents your computer from communicating with the terminal server to which you were connected.

**Solution:** Try connecting to the remote computer again. If you receive the same message, contact the server administrator.

**Because of a protocol error (error_code), the remote session will be disconnected. Please try connecting to the remote computer again.**
**Cause:** A protocol specifies the way a computer communicates with other computers. A protocol error prevents your computer from communicating with the terminal server to which you were connected.

**Solution:** Try connecting to the remote computer again. If you receive the same message, contact the server administrator, and provide the error code that was displayed in the message.

**Because of a security error, the client could not connect to the remote computer. Verify that you are logged on to the network, and then try connecting again.**
**Cause:** The most common causes are network errors or slowdowns.

**Solution:** Be sure your computer is connected to the network. Try to reconnect at a later time. If you cannot connect, contact the server administrator.

**The remote session was disconnected because the remote computer is low on memory.**
**Cause:** Memory errors have a variety of causes, but the result is the same. The remote computer does not have enough memory to start and maintain the connection you are attempting to make.

**Solution:** Contact the server administrator.

**Because of an unknown error, the remote session will be disconnected. The local computer might be low on memory. Close some programs, and then try connecting to the remote computer again.**
**Cause:** Not enough RAM available on your computer. If your computer has too little RAM available, it cannot free enough processing capacity to start new functions, such as applications or connections.

**Solution:** Close all unnecessary programs, and try your connection again.

**Cause:** Not enough disk space available on your hard drive. If the hard drive on your computer is full, there might not be enough space available to allow the swap file to operate correctly. The swap file allows your computer to perform some RAM functions even when there is not enough RAM space available.

**Solution:** Free some space on your hard drive by:

- Emptying your recycle bin.
- Deleting temporary Internet files.
- Moving documents and other data files to a different drive.
- Removing programs that you rarely use.
- Deleting old temporary files from your hard disk. Temporary files have a .tmp file name extension, or names that begin with a tilde (~), and are usually found in the C:\Temp or C:\Windows\Temp folder.
- Emptying your recycle bin after you delete files.

**An unrecoverable error occurred while the client was attempting to connect. Close some programs, and then try connecting to the remote computer again.**
**Cause:** Not enough RAM available on your computer. If your computer has too little RAM available, it cannot free enough processing capacity to start new functions, such as applications or connections.

**Solution:** Close all unnecessary programs, and try your connection again.

**Cause:** Not enough disk space available on your hard drive. If the hard drive on your computer is full, there might not be enough space available to allow the swap file to operate correctly. The swap file allows your computer to perform some RAM functions even when there is not enough RAM space available.

**Solution:** Free some space on your hard drive by:

- Emptying your recycle bin.
- Deleting temporary Internet files.

- Moving documents and other data files to a different drive.
- Removing programs that you rarely use.
- Deleting old temporary files from your hard disk. Temporary files have a .tmp file name extension, or names that begin with a tilde (~), and are usually found in the C:\Temp or C:\Windows\Temp folder.
- Emptying your recycle bin after you delete files.

**An error occurred when the client attempted to connect to the remote computer. Check your system memory and then try the connection again.**
**Cause:** Not enough RAM available on your computer. If your computer has too little RAM available, it cannot free enough processing capacity to start new functions, such as applications or connections.

**Solution:** Close all unnecessary programs, and try your connection again.

**Cause:** Not enough disk space available on your hard drive. If the hard drive on your computer is full, there might not be enough space available to allow the swap file to operate correctly. The swap file allows your computer to perform some RAM functions even when there is not enough RAM space available.

**Solution:** Free some space on your hard drive by:

- Emptying your recycle bin.
- Deleting temporary Internet files.
- Moving documents and other data files to a different drive.
- Removing programs that you rarely use.
- Deleting old temporary files from your hard disk. Temporary files have a .tmp file name extension, or names that begin with a tilde (~), and are usually found in the C:\Temp or C:\Windows\Temp folder.
- Emptying your recycle bin after you delete files.

**Unable to allocate window resources. The remote session will be disconnected. Close some programs on the local computer, and then try connecting to the remote computer again.**
**Cause:** Not enough RAM available on your computer. If your computer has too little RAM available, it cannot free enough processing capacity to start new functions, such as applications or connections.

**Solution:** Close all unnecessary programs, and try your connection again.

**Cause:** Not enough disk space available on your hard drive. If the hard drive on your computer is full, there might not be enough space available to allow the swap file to operate correctly. The swap file allows your computer to perform some RAM functions even when there is not enough RAM space available.

**Solution:** Free some space on your hard drive by:

- Emptying your recycle bin.
- Deleting temporary Internet files.
- Moving documents and other data files to a different drive.
- Removing programs that you rarely use.
- Deleting old temporary files from your hard disk. Temporary files have a .tmp file name extension, or names that begin with a tilde (~), and are usually found in the C:\Temp or C:\Windows\Temp folder.
- Emptying your recycle bin after you delete files.

# Troubleshooting Remote Desktop Licensing Error Messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting Licensing Error Messages**
This topic lists licensing error messages that can appear on client computers, and it describes the causes of and solutions for these errors. Although these error messages appear on clients, they are frequently caused by problems with the Windows Server 2003 Terminal Server license server or the terminal server. Therefore, when you troubleshoot Terminal Server Licensing issues, it is useful to first determine whether there are server configuration issues or problems with network connectivity.

**Important**

The solutions in this topic are intended to be implemented by administrators. If you are not an administrator, contact your server administrator for assistance in resolving these error messages. If you are an administrator, to help prevent Terminal Server Licensing issues and to more efficiently diagnose issues, it is highly recommended that you see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627), and Windows Server 2003 Terminal Server Licensing Issues and Requirements for Deployment (http://go.microsoft.com/fwlink/?LinkID=23444).

Which message are you getting?

- The remote session was disconnected because the local computer's client access license could not be upgraded or renewed. Please contact the server administrator.
- The remote computer disconnected the session because of an error in the licensing protocol. Please try connecting to the remote computer again or contact your server administrator.
- The remote session was disconnected because there were network problems during the licensing protocol. Please try connecting to the remote computer again.
- A licensing error occurred while the client was attempting to connect. (Licensing timed out.) Please try connecting to the remote computer again.
- The remote session was disconnected because there are no Terminal Server client access licenses available for this computer. Please contact the server administrator.
- The remote session was disconnected because there are no Terminal Server License Servers available to provide a license. Please contact the server administrator.
- Because of a security error, the client could not connect to the terminal server. After making sure that you are logged on to the network, try connecting to the server again.

**The remote session was disconnected because the local computer's client access license could not be upgraded or renewed. Please contact the server administrator.**
**Cause:** The Terminal Server licensing mode on the terminal server might be set to Per Device, while the license server might be configured to issue only Per User client access licenses (CALs). If this is the case, then the license server issues only temporary licenses that cannot be upgraded. When the temporary licenses are within several days of expiring, "Event ID 26, Source: Application Pop-up" appears in the application event log on the client. The event message indicates the number of days remaining before the temporary license expires. Similarly, "Event ID 1011, Source TermService" appears in the application event log on the terminal server.

**Solution:** Change the Terminal Server licensing mode from Per Device to Per User.

**Important**

Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. This does not remove administrators from End User License Agreement (EULA) requirements to have a valid terminal server CAL for each user. Failure to have a Per User CAL for each user, if Per Device CALs are not being used, is a violation of the EULA.

The following text is from the EULA for Windows Server 2003:

"Two different TS CALs are available to you: 'Device' and 'User.' Each TS Device CAL permits one Device (used by any User) to conduct Windows Sessions on any of your Servers. Each TS User CAL permits one User (using any Device) to conduct Windows Sessions on any of your Servers. You may use a mix of TS Device CALs and TS User CALs simultaneously with the Server Software in your environment. You can have a Terminal Server request Per User licenses or Per Device (default) but not both simultaneously."

**See also:** Configure the Terminal Server Licensing mode

**The remote computer disconnected the session because of an error in the licensing protocol. Please try connecting to the remote computer again or contact your server administrator.**
**Cause:** The terminal server might not be able to locate the license server.

**Solution:** Perform the following steps:

1. Verify that the license server is correctly installed.

2. Verify that the Terminal Server Licensing service is running on the license server.

3. Verify that the client, the terminal server, and the license server can communicate by ensuring that Domain Name System (DNS) is configured correctly on each computer. To do this, run the **ping** command from each computer to each computer using the IP address, FQDN, and the NetBIOS name. If any of the **ping** commands fail, verify the DNS configuration on the network.

4. On the terminal server, set a preferred licensing server to connect to. You can do this by using the registry, a Windows Management Instrumentation (WMI) script, or, in Windows Server 2003 with Service Pack 1, you can use Group Policy settings or Terminal Server Configuration.

**See also:** Install Terminal Server Licensing; Set preferred Terminal Server license servers; Terminal Server license server roles

**Solution:** If the previous solution does not resolve this problem, create a backup of the **MSLicensing** registry key and its subkeys on the client, and then remove the original key and subkeys by doing the following:

1. On the client, navigate to the following registry subkey:
   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSLicensing**.

2. Click **MSLicensing.**

3. On the **Registry** menu, click **Export Registry File**.

4. In the **File name** box, type **mslicensingbackup**, and then click **Save**.

5. If you need to restore this registry key in the future, double-click **mslicensingbackup.reg**.

6. On the **Edit** menu, click **Delete**, and then click **Yes** to confirm the deletion of the **MSLicensing** registry subkey.

7. Close Registry Editor, and then restart the computer.

When the client is restarted, the missing registry key is rebuilt.

To resolve this problem automatically, click the Fix this problem link. Then, click Run in the File Download dialog box.

Fix this problem

| **Note** |
| --- |
| This wizard may be in English only; however, the automatic fix also works for other language versions of Windows. |

| Note |
| --- |
| If you are not on the computer that has the problem, save the automatic fix to a flash drive or to a CD so that you can run it on the computer that has the problem. If you do not want to use the automatic fix, you can do it yourself by performing the steps in the procedure above. |

**Caution**

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**See also:** Removing Terminal Server Licenses From an RDP Client (http://go.microsoft.com/fwlink/?LinkId=38560)

**The remote session was disconnected because there were network problems during the licensing protocol. Please try connecting to the remote computer again.**
**Cause:** The terminal server might not be able to locate the license server.

**Solution:** Perform the following steps:

1. Verify that the license server is correctly installed.

2. Verify that the Terminal Server Licensing service is running on the license server.

3. Verify that the client, the terminal server, and the license server can communicate by ensuring that Domain Name System (DNS) is configured correctly on each computer. To do this, run the **ping** command from each computer to each computer using the IP address, FQDN, and the NetBIOS name. If any of the **ping** commands fail, verify the DNS configuration on the network.

4. On the terminal server, set a preferred licensing server to connect to. You can do this by using the registry, a Windows Management Instrumentation (WMI) script, or, in Windows Server 2003 with Service Pack 1, you can use Group Policy settings or Terminal Server Configuration.

**See also:** Install Terminal Server Licensing; Set preferred Terminal Server license servers; Terminal Server license server roles

**Cause:** The license server might be running Windows 2000, and it might be configured to prohibit anonymous connections, except by resources that have been explicitly granted access to the server.

**Solution:** On the Windows 2000 license server, do one of the following:

**Caution**

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

- In the registry, navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\RestrictAnonymous**, and set the **RestrictAnonymous** registry key to a value of **1** or **0**.

- In Local Security Policies, open the appropriate policy and, in the console tree, navigate to Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options. Then, enable either **Do not allow enumeration of SAM accounts and shares** (equivalent to a **RestrictAnonymous** value of **1** or **None**) or **Rely on default permissions** (equivalent to a **RestrictAnonymous** value of **0**).

**Notes**

- If the license server is a member of an Active Directory domain and a conflicting security setting is configured for the license server in Group Policy, the Group Policy setting overrides the local security setting. In this case, to ensure that the security setting that you want to apply takes effect, configure the setting in Group Policy.

- In Windows Server 2003, you cannot set **RestrictAnonymous** to a value of **2** to prohibit anonymous connections. If you need to prohibit anonymous users from being granted the same access that is granted to members of the Everyone

group, you must use the new **Everyone Network access: Let Everyone permissions apply to anonymous users** setting in Local Security Policies.

**Solution:** If the previous solution does not resolve this problem, create a backup of the **MSLicensing** registry key and its subkeys on the client, and then remove the original key and subkeys by doing the following:

1. On the client, navigate to the following registry subkey: **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSLicensing**.
2. Click **MSLicensing.**
3. On the **Registry** menu, click **Export Registry File**.
4. In the **File name** box, type **mslicensingbackup**, and then click **Save**.
5. If you need to restore this registry key in the future, double-click **mslicensingbackup.reg**.
6. On the **Edit** menu, click **Delete**, and then click **Yes** to confirm the deletion of the **MSLicensing** registry subkey.
7. Close Registry Editor, and then restart the computer.

When the client is restarted, the missing registry key is rebuilt.

To resolve this problem automatically, click the Fix this problem link. Then, click Run in the File Download dialog box.

Fix this problem

| **Note** |
| --- |
| This wizard may be in English only; however, the automatic fix also works for other language versions of Windows. |

| **Note** |
| --- |
| If you are not on the computer that has the problem, save the automatic fix to a flash drive or to a CD so that you can run it on the computer that has the problem. If you do not want to use the automatic fix, you can do it yourself by performing the steps in the procedure above. |

**Caution**

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**See Also:** How To Use the RestrictAnonymous Registry Value in Windows 2000 (http://go.microsoft.com/fwlink/?LinkId=38561)

**A licensing error occurred while the client was attempting to connect. (Licensing timed out.) Please try connecting to the remote computer again.**
**Cause:** If you are using Internet Protocol security (IPsec) to help protect traffic over TCP between clients and terminal servers, then packet fragmentation might occur. As a result, some packets might not reach their destination, and client connections to terminal servers might fail.

**Solution:** Configure IPsec to help protect traffic over UDP rather than over TCP.

**See Also:** Define IPSec Policies

**The remote session was disconnected because there are no Terminal Server client access licenses available for this computer. Please contact the server administrator.**
**Cause:** The license server might not have any remaining Per Device CALs to issue.

**Solution:** Use Terminal Server Licensing to determine the number of CALs remaining on the license server. To open Terminal Server Licensing, click **Start**, **Programs**, **Administrative Tools**, and then point to **Terminal Server Licensing**. If the license server does not have any remaining CALs to issue to clients, purchase and install additional CALs as required.

**See also:** Purchase client access licenses; Install Client Access Licenses; Purchasing and installing client access licenses on a Terminal Server license server

**Cause:** The Terminal Server licensing mode on the terminal server might be set to Per Device, while the license server might have only Per User CALs. If this is the case, the license server issues only temporary licenses that cannot be upgraded. When the temporary licenses are within several days of expiring, "Event ID 26, Source: Application Pop-up" appears in the application event log on the client. The event message indicates the number of days remaining before the temporary license expires. Similarly, "Event ID 1011, Source TermService" appears in the application event log on the terminal server.

**Solution:** Change the Terminal Server Licensing mode from Per Device to Per User.

**Important**

Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. This does not remove administrators from End User License Agreement (EULA) requirements to have a valid terminal server CAL for each user. Failure to have a Per User CAL for each user, if Per Device CALs are not being used, is a violation of the EULA.

The following text is from the EULA for Windows Server 2003:

"Two different TS CALs are available to you: 'Device' and 'User.' Each TS Device CAL permits one Device (used by any User) to conduct Windows Sessions on any of your Servers. Each TS User CAL permits one User (using any Device) to conduct Windows Sessions on any of your Servers. You may use a mix of TS Device CALs and TS User CALs simultaneously with the Server Software in your environment. You can have a Terminal Server request Per User licenses or Per Device (default) but not both simultaneously."

**See also:** Configure the Terminal Server Licensing mode

**Cause:** The terminal server might not be able to locate the license server.

**Solution:** Perform the following steps:

1. Verify that the license server is correctly installed.

2. Verify that the Terminal Server Licensing service is running on the license server.

3. Verify that the client, the terminal server, and the license server can communicate by ensuring that Domain Name System (DNS) is configured correctly on each computer. To do this, run the **ping** command from each computer to each computer using the IP address, FQDN, and the NetBIOS name. If any of the **ping** commands fail, verify the DNS configuration on the network.

4. On the terminal server, set a preferred licensing server to connect to. You can do this by using the registry, a Windows Management Instrumentation (WMI) script, or, in Windows Server 2003 with Service Pack 1, you can use Group Policy settings or Terminal Server Configuration.

**See also:** Install Terminal Server Licensing; Set preferred Terminal Server license servers; Terminal Server license server roles

**Solution:** If the previous solution does not resolve this problem, create a backup of the **MSLicensing** registry key and its subkeys on the client, and then remove the original key and subkeys by doing the following:

1. On the client, navigate to the following registry subkey:
   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSLicensing**.

2. Click **MSLicensing.**

3. On the **Registry** menu, click **Export Registry File**.

4. In the **File name** box, type **mslicensingbackup**, and then click **Save**.

If you need to restore this registry key in the future, double-click **mslicensingbackup.reg**.

5. On the **Edit** menu, click **Delete**, and then click **Yes** to confirm the deletion of the **MSLicensing** registry subkey.

6. Close Registry Editor, and then restart the computer.

When the client is restarted, the missing registry key is rebuilt.

To resolve this problem automatically, click the Fix this problem link. Then, click Run in the File Download dialog box.

Fix this problem

| **Note** |
| --- |
| This wizard may be in English only; however, the automatic fix also works for other language versions of Windows. |

| **Note** |
| --- |
| If you are not on the computer that has the problem, save the automatic fix to a flash drive or to a CD so that you can run it on the computer that has the problem. If you do not want to use the automatic fix, you can do it yourself by performing the steps in the procedure above. |

**Caution**

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**See Also:** How To Use the RestrictAnonymous Registry Value in Windows 2000 (http://go.microsoft.com/fwlink/?LinkId=38561)

**The remote session was disconnected because there are no Terminal Server license servers available to provide a license. Please contact the server administrator.**
**Cause:** The terminal server might not be able to locate the license server.

**Solution:** Perform the following steps:

1. Verify that the license server is correctly installed.

2. Verify that the Terminal Server Licensing service is running on the license server.

3. Verify that the client, the terminal server, and the license server can communicate by ensuring that Domain Name System (DNS) is configured correctly on each computer. To do this, run the **ping** command from each computer to each computer using the IP address, FQDN, and the NetBIOS name. If any of the **ping** commands fail, verify the DNS configuration on the network.

4. On the terminal server, set a preferred licensing server to connect to. You can do this by using the registry, a Windows Management Instrumentation (WMI) script, or, in Windows Server 2003 with Service Pack 1, you can use Group Policy settings or Terminal Server Configuration.

**See also:** Install Terminal Server Licensing; Set preferred Terminal Server license servers; Terminal Server license server roles

**Cause:** The client might have exceeded its Terminal Server Licensing grace period, and a license server has not yet been installed or activated to issue the client a CAL.

**Solution:** Install Terminal Server Licensing, activate the license server, and then install and purchase a sufficient number of CALs to support the clients in your organization.

**See also:** Terminal Server Licensing grace period; Install Terminal Server Licensing; Activate a Terminal Server License Server;

Purchase client access licenses; Install Client Access Licenses; Purchasing and installing client access licenses on a Terminal Server license server; Activating a Terminal Server license server

**Because of a security error, the client could not connect to the terminal server. After making sure that you are logged on to the network, try connecting to the server again.**

**Cause:** If you upgraded a Windows NT domain to Windows 2000 or Windows Server 2003, then the certificate on the terminal server might be corrupt. As a result, Windows 2000 Terminal Services clients might be repeatedly denied access to the terminal server.

**Solution:** On each terminal server and client, perform the following steps:

1. On each terminal server, create a backup of the registry.
2. Navigate to the following registry subkey: **HKLM\SYSTEM\CurrentControlSet\Services\TermServices\Parameters**.
3. On the **Registry** menu, click **Export Registry File**.
4. In the **File name** box, type **exported-parameters**, and then click **Save**.

   If you need to restore this registry subkey in the future, double-click **exported-parameters.reg**.
5. Under the **Parameters** registry subkey, right-click each of the following values:
   - **Certificate**
   - **X509 Certificate**
   - **X509 Certificate ID**
6. Click **Delete**, and then click **Yes** to confirm the deletion.
7. Close Registry Editor, and then restart each terminal server.
8. On the client, create a backup of the **MSLicensing** registry key and its subkeys, and then remove the original key and subkeys by doing the following:
   - Navigate to the following registry subkey: **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSLicensing**.
   - Click **MSLicensing**.
   - On the **Registry** menu, click **Export Registry File**.
   - In the **File name** box, type **mslicensingbackup**, and then click **Save**.
   - If you need to restore this registry key in the future, double-click **mslicensingbackup.reg**.
   - On the **Edit** menu, click **Delete**, and then click **Yes** to confirm the deletion of the **MSLicensing** registry subkey.
   - Close Registry Editor, and then restart the client computer.

**Caution**

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Solution:** If the client still cannot connect to the terminal server, perform the following variation of this procedure:

1. Deactivate the license server.

2. Reactivate the license server by using the Telephone connection method in the Terminal Server License Server Wizard.

   When you activate Terminal Server Licensing by using the Telephone option, Terminal Server Licensing uses a different certificate.

3. On each terminal server, create a backup of the registry, and then delete the **Certificate**, **X509 Certificate**, and **X509 Certificate ID** registry keys, as described in the preceding procedure.

4. Close Registry Editor, and then restart each terminal server.

5. On the client, create a backup of the **MSLicensing** registry key and its subkeys on the client, and then remove the original key and subkeys as described in the preceding procedure.

6. Close Registry Editor, and then restart the computer.

When the client is restarted, the missing registry key is rebuilt.

**See also:** Deactivate a Terminal Server license server; Reactivate a Terminal Server license server

**Cause:** Windows XP-based clients might be attempting to connect to a Windows 2000-based Terminal Services server in a low-bandwidth network environment, in which client sessions are encrypted. In this case, IP packet fragmentation can cause encrypted frames that are sent by a client to be decrypted incorrectly.

**Solution:** Obtain the latest service pack for Windows 2000.

**Cause:** The Remote Desktop Protocol (RDP) encryption settings on the terminal server computer and the client might not be compatible. For example, the terminal server might be running 128-bit encryption with an encryption level set to **High**. When this occurs, "Event ID 50, Source: TermDD" appears in the system event log on the terminal server.

**Solution:** Change the RDP encryption level on the terminal server to **Medium** or **Low**.

**See also:** Because of a Security Error, the Client Could Not Connect to the Terminal Server (http://support.microsoft.com/default.aspx?scid=kb;en-us;329896); Terminal Services Sessions Are Disconnected Because of a Decryption Error (http://go.microsoft.com/fwlink/?LinkId=38563); Terminal Services Client Cannot Connect to a Server Running 128-bit Encryption (http://go.microsoft.com/fwlink/?LinkId=38559); Configure authentication and encryption on the server

# Troubleshooting Specific Remote Desktop Problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting Specific problems**
What problem are you having?

- I cannot copy text from an application on the terminal server to another application that is on my local computer.
- I minimized the Remote Desktop Connection dialog box, and then my screen went blank.
- The screen saver on my Remote Desktop is blank.
- I'm supposed to be able to connect to a specific application when I log on to the remote computer, but it doesn't start.
- Remote Desktop Connection is very slow.

**I cannot copy text from an application on the terminal server to another application that is on my local computer.**
**Cause:** It is possible that Remote Desktop Connection was not installed correctly.

**Solution:** Uninstall and then reinstall Remote Desktop Connection. If you need assistance, contact the server administrator.

**I minimized the Remote Desktop Connection dialog box, and then my screen went blank.**
**Cause:** Using a password-protected screen saver can cause this problem. Your applications are still running on the remote computer, but you can no longer use them.

**Solution:** Disconnect Remote Desktop Connection, or ask the server administrator to disconnect you.

**Solution:** Do not minimize the Remote Desktop Connection dialog box if your computer uses a password-protected screen saver.

**Solution:** Do not use a password-protected screen saver on your local computer.

**The screen saver on my Remote Desktop is blank.**
**Cause:** By default, when a screen saver is activated on Remote Desktop, it is blank. This is true regardless of whether you have previously selected a different screen saver.

**Solution:** Continue to use the blank screen saver.

**I'm supposed to be able to connect to a specific application when I log on to the remote computer, but it doesn't start.**
**Solution:** If the remote computer is running Windows XP Professional, the complete contents of the Desktop are always displayed. Start the program you want to use by double-clicking its icon on the Desktop or using the Start menu.

**Cause:** The client connection might have been configured to start a program upon connection, but the program path and file name are incorrect.

**Solution:** Check the properties for the client connection to ensure that the program path and file name are correct. Also, confirm that the program exists in this location on the server.

**Cause:** The server administrator might have specified a different program to start upon connection, using Group Policy, Terminal Services Configuration, or other Terminal Services administration tools.

**Solution:** Contact the server administrator.

**Remote Desktop Connection is very slow.**
**Cause:** A low-bandwidth connection can slow Remote Desktop performance.

**Solution:** On the remote computer, set the background selection to **None.**

**Solution:** Do not run memory-intensive applications over Remote Desktop Connection.

**Solution:** Consider upgrading your modem or purchasing a high-speed service.

# Windows Server 2003 Administration Tools Pack

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003 Administration Tools Pack**
To make the remote management of your servers easier, Microsoft has included Windows Server 2003 Administration Tools Pack on the Microsoft® Windows Server 2003 family CDs Using these CDs, you can install Windows Server 2003 Administration Tools Pack onto computers running the following operating systems.

- Windows Server 2003 family
- Windows XP Professional with Service Pack 1 or later
- Windows XP Professional with QFE Q329357

For tips about administering servers from Windows XP Professional and Windows 2000 Professional, see Windows Server 2003 Administration Tools Pack Best practices.

For information about installing, reinstalling, deploying, and uninstalling Windows Server 2003 Administration Tools Pack, see Windows Server 2003 Administration Tools Pack How to....

For an overview of the tools, see Windows Server 2003 Administration Tools Pack Overview.

For help with specific tasks and links to detailed information about these tools, see Managing remotely with the Windows Server 2003 Administration Tools Pack tools.

For help with common problems, see Troubleshooting Windows Server 2003 Administration Tools Pack.

Windows Server 2003 Administration Tools Pack tools will not work on computers running Microsoft® Windows 2000. For additional information, see Windows 2000 Server online Help.

# Windows Server 2003 Administration Tools Pack Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Check security requirements for each component**

- For a broader security discussion and links to additional security information, see the following.
    - Best practices for security
    - Security overview
    - Local Users and Groups Best practices

**Always install help contents from the Windows Server 2003 operating system**

- The help content on a Windows Server 2003 operating system differs from Windows XP Professional Help. After you install the help content from a Windows Server 2003 operating system onto your computer running Windows XP Professional, you can switch between the different help content as often as you like, but you cannot use both simultaneously. To switch between different help systems, in **Help and Support Center**, click the **Options** menu. Under **Options**, click **Install and share Windows Help**, click **Switch from one operating system's Help content to another**, and then follow the instructions on the screen.

- If you are administering servers running Windows Server 2003 from a computer running Windows XP Professional with Service Pack 1 or Windows XP Professional with hotfix Q329357, you should install the help content from the Windows Server 2003 operating system onto the computer running Windows XP Professional.

  For instructions, see Install Help content from a Windows Server 2003 server operating system onto a computer running Windows XP Professional.

**Make sure that the firewall settings on the server you want to administer permit remote administration**

- With Windows Firewall enabled, you might need to make specific configuration settings in order to use or remotely administer technologies that ship with Windows Server 2003. For information about these settings, see Windows Firewall Settings.

**Administering computers running Windows 2000 Server**

- Windows Server 2003 Administration Tools Pack has been designed solely for use on servers running Windows Server 2003, Windows XP Professional with Service Pack 1 or later, and Windows XP Professional with QFE Q329357 installed. Windows Server 2003 Administration Tools Pack will not run on computers running Windows 2000 operating systems. Windows 2000 operating systems can use only the Windows 2000 version of Administrative Tools.
    - Install Windows 2000 Administrative Tools onto a computer running Windows 2000 professional. For additional information, see Windows 2000 Help.
    - Use the Remote Desktop Connection from Windows XP Professional or Windows 2000 professional to the server running Windows 2000 and manage that server from the remote session. For more information, see Remote Desktop Connection.

# Windows Server 2003 Administration Tools Pack How to...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install or reinstall Windows Server 2003 Administration Tools Pack
- Deploy Windows Server 2003 Administration Tools Pack
- Uninstall Windows Server 2003 Administration Tools Pack
- Install Help content from a Windows Server 2003 server operating system onto a computer running Windows XP Professional

# Install or reinstall Windows Server 2003 Administration Tools Pack

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install or reinstall Windows Server 2003 Administration Tools Pack**

- Using the Windows interface
- Using a command line

**Using the Windows interface**
**If you are using a Windows Server 2003 family CD:**

1. Put your Windows Server 2003 CD into the CD tray of the computer on which you want to install the Administration Tools Pack.
2. When the Welcome screen appears, click **Perform additional tasks**.
3. Click **Browse this CD**.
4. Open the server cd i386 folder.
5. Right-click adminpak.msi, then click **Install**.

**If you are using a Windows Server 2003 family CD image or net share:**

1. Put your Windows Server 2003 CD into the CD tray of the computer on which you want to install the Administration Tools Pack and open the i386 folder on the Windows Server 2003 operating system image or net share.
2. Right-click adminpak.msi, then click **Install** to begin installation of Windows Server 2003 Administration Tools Pack onto the Windows XP Professional operating system.

**Note**

- If the **Welcome to the Microsoft Server 2003 Family** screen does not appear, click **Start**, then **Run**. Click **Browse** and select the drive holding the server cd. In the **Browse** dialog box, set **Files of type** to **All files**, then open the server cd i386 folder. Right-click adminpak.msi, then click **Install**.

**Using a command line**

1. With the Windows Server 2003 family CD in the CD tray of your computer running Windows XP Professional, open Command Prompt.
2. Type the path to adminpak.msi. For example, if the CD drive is drive E, "E:\i386\adminpak.msi".
3. Press **Enter**, and adminpak.msi will install Windows Server 2003 Administration Tools Pack onto the computer running Windows XP Professional with Service Pack 1 operating system.

   **Note**

   - To open the Command Prompt, click **Start**, then click **Run**, and type cmd.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Administrative Tools, click **Start**, click **Control Panel**, and then double-click **Administrative Tools**.

- Installed Windows Server 2003 Administration Tools Pack MMC snap-ins are listed in **Administrative Tools** on the Windows XP Professional **Start** menu. They are also listed in **Administrative Tools** in **Control Panel**.

- You can reinstall Windows Server 2003 Administration Tools Pack at any time by using any of the installation methods. Reinstallation will overwrite any prior installation.

- Command-line tools are not available on the **Start** menu. Command-line tools can be accessed from the command line.

- After upgrading from Windows 2000 Server, you must install Windows Server 2003 Administration Tools Pack. It is not installed automatically during upgrades.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Windows Server 2003 Administration Tools Pack Overview
Deploy Windows Server 2003 Administration Tools Pack
Uninstall Windows Server 2003 Administration Tools Pack
Install Help content from a Windows Server 2003 server operating system onto a computer running Windows XP Professional

# Deploy Windows Server 2003 Administration Tools Pack

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To deploy Windows Server 2003 Administration Tools Pack**

1. If you do not already have a user group for the administrators who will receive adminpak.msi, create one. For more information, see Group Policy in Related Topics.

2. Copy adminpak.msi from the Windows Server 2003 family CD to your software distribution share and publish or assign it as you would any other program. For more information about assigning and publishing software, see Group Policy in Related Topics.

## Notes

- You assign an application when you want the targeted user group to have the application assigned to his or her computer for installation. With Windows Server 2003 Administration Tools Pack installed on Windows XP Professional with Service Pack 1 or Windows XP Professional with QFE Q329357, you click each snap-in on the Administrative Tools menu to install only that snap-in and any related snap-ins.

- You publish an application when you want the targeted user group to be able to install the application at users' discretion.

- You must open all internal snap-ins from the **Start** menu before you open that console or the contained snap-ins do not load.

- If you publish Windows Server 2003 Administration Tools Pack in Active Directory, targeted users can use **Add or Remove programs** in **Control Panel** on the remote computer to install it when needed.

- After Windows Server 2003 Administration Tools Pack has been installed to a Windows XP Professional computer, you can access the tools from the **Start** menu. Click **Start**, point to **programs**, and then point to **Administrative Tools**. Click the name of the tool that you want to use.

- Command line tools are not available on the **Start** menu. Command line tools can be accessed from the command line.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy (pre-GPMC)
Windows Server 2003 Administration Tools Pack Overview
Install or reinstall Windows Server 2003 Administration Tools Pack
Uninstall Windows Server 2003 Administration Tools Pack

# Uninstall Windows Server 2003 Administration Tools Pack

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To uninstall Windows Server 2003 Administration Tools Pack**

1. Open Add or Remove Programs in Control Panel.
2. Select **Change or remove programs**.
3. In Windows Server 2003 Administration Tools Pack, click **Remove**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Windows Server 2003 Administration Tools Pack Overview
Install or reinstall Windows Server 2003 Administration Tools Pack
Deploy Windows Server 2003 Administration Tools Pack

# Install Help content from a Windows Server 2003 server operating system onto a computer running Windows XP Professional

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install help content from a Windows Server 2003 operating system onto a computer running Windows XP professional**

1. On the computer running Windows XP Professional, click **Start**, then **Help and Support**.
2. In Help and Support Center, click the **Options** menu.
3. Under **Options**, click **Install and share Windows Help**, click one of the following:
   - **Install Help content from another Windows computer**
   - **Install Help content from a CD or disk image**
4. Follow the instructions on the screen.

## Notes

- The help content on a Windows Server 2003 operating system differs from Windows XP Professional Help. After you install the help content from a Windows Server 2003 operating system onto your Windows XP Professional computer, you can switch between the different help content as often as you like, but you cannot use both simultaneously. To switch between different help systems, in **Help and Support Center**, click the **Options** menu. Under **Options**, click **Install and share Windows Help**, click **Switch from one operating system's Help content to another**, and then follow the instructions on the screen.
- If you are administering computers running Windows Server 2003 operating systems from a computer running Windows XP Professional, you should install the help content from the Windows Server 2003 operating system onto the computer running Windows XP Professional.
- You can only install and share Help content on computers that are running Windows XP Professional or a product in the Windows Server 2003 family. For example, you cannot install Help content from Windows XP or a product in the Windows Server 2003 family on a computer running Windows 2000, or vice versa.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Viewing Help for the Windows Server 2003 family from another computer
Windows Server 2003 Administration Tools Pack Overview
Deploy Windows Server 2003 Administration Tools Pack
Uninstall Windows Server 2003 Administration Tools Pack

# Windows Server 2003 Administration Tools Pack Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Server 2003 Administration Tools Pack Overview**
Microsoft® Windows Server 2003 Administration Tools Pack provides administrators with tools for managing their servers remotely from platforms:

- The Windows Server 2003 family
- Windows XP Professional with Service Pack 1 (SP1)
- Windows XP Professional with QFE Q329357

These tools are packaged as adminpak.msi on the Microsoft® Windows Server 2003 family CDs. These tools support 32-bit and 64-bit server operating systems.

These tools can be classified into four categories.

- System administration
- Network administration
- Storage management
- Directory services administration

Windows Server 2003 Administration Tools Pack also provides Convenience Consoles which combine certain of these tools into packages for managing or delegating frequent tasks.

For Windows Server 2003 Administration Tools Pack to be installed on Windows XP Professional, you must either apply Service Pack 1 to the operating system or install the hotfix included with Knowledge Base Article Q329357. See article Q329357, "Windows Server 2003 adminpak.msi may not run on Windows XP 2600 Machines," in the Microsoft Knowledge Base.

For more information, see Windows Server 2003 Administration Tools Pack Best practices.

You must have administrative permissions on the local computer to install or run Windows Server 2003 Administration Tools Pack. For security reasons, Windows Server 2003 Administration Tools Pack should be removed if the computer running Windows XP Professional is to be used by a non-administrator. In addition, you should be aware of specific security requirements for each tool.

For more information, see:

- Managing remotely with the Windows Server 2003 Administration Tools Pack tools
- Convenience Consoles in Windows Server 2003 Administration Tools Pack

**Finding Windows Server 2003 Administration Tools Packs**
You can find and download Windows Server 2003 Administration Tools Packs on the following Web sites:

- Windows Server 2003 Administration Tools Pack (http://go.microsoft.com/fwlink/?LinkID=91722)

  Allows administrators to remotely manage Windows 2000 Server and Windows Server 2003.

- Windows Server 2003 Service Pack 1 Administration Tools Pack (http://go.microsoft.com/fwlink/?LinkID=54895)

Allows administrators to remotely manage computers running Windows 2000 Server and Windows Server 2003 family operating systems. This site provides the Service Pack 1 version (build 5.2.3790.1830) of the adminpak.msi file.

- Windows Server 2003 R2 Administration Tools Pack (x86) (http://go.microsoft.com/fwlink/?LinkID=73895)

  Allows administrators to remotely manage computers running 32-bit versions of the Windows Server 2003 R2 family of operating systems. The management tools downloadable from this site allow management of features specific to Windows Server 2003 R2.

- Windows Server 2003 R2 Administration Tools Pack (x64) (http://go.microsoft.com/fwlink/?LinkId=110584)

  Allows administrators to remotely manage computers running 64-bit versions of the Windows Server 2003 R2 family of operating systems. The management tools contained in these packages manage features specific to Windows Server 2003 R2.

# Managing remotely with the Windows Server 2003 Administration Tools Pack tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing remotely with the tools**
**System Administration**

| Task | Tool | In a Convenience Console? |
|---|---|---|
| Support on-demand, image-based, operating system installations over a network connection from a RIS server to a client computer to simplify the management of operating systems and applications, and to improve failure recovery. | Remote Installation Services | No |
| Provide application developers a flexible framework for integrating role-based access control into their applications and administrators who use those applications access in a natural, intuitive way: through assigned user roles relating to job function. | Authorization Manager | No |
| Issue certificates to individuals, computers, or organizations that affirm the identity and other attributes of the certificate subject to other entities | Certification Authorities | Public Key Management |
| Track the licenses that have been issued to many terminal servers concurrently. | Terminal Server Licensing | No |
| Manage and monitor users, sessions, and processes on any server running Terminal Services on the network. | Terminal Services Manager | No |
| Manage Web Services directly through standards-based XML Web service for development tools and business applications. | UDDI Services overview | No |

**Network Administration**

| Task | Tool | In a Convenience Console? |
|---|---|---|
| Define features for remote access clients so that customers can connect to your service. | Connection Manager Administration Kit | No |
| Manage centrally IP addresses and other related configuration details used on your network. | DHCP | IP Address Management |
| Manage files distributed across multiple servers to appear to users as if they reside in one place on the network. | Distributed File System (DFS) | No |
| Manage TCP/IP networks, such as the Internet, to locate computers and services through user-friendly names. | DNS | Active Directory Management and IP Address Management |
| Manage exchange with NNTP snap-in. | Choosing components to install<br><br>Microsoft NNTP Service (Requires installing the NNTP service first). | No |

| | | |
|---|---|---|
| Manage telecommunications with the computer to communicate over direct connections to local computers, phone lines, LANs, WANs, and the Internet. | Telephony | No |
| Map NetBIOS names to IP addresses , particularly for NetBIOS name resolution in routed networks that use NetBIOS over TCP/IP. | WINS | IP Address Management |

**Storage Management**

| Task | Tool | In a Convenience Console? |
|---|---|---|
| Manage a group of independent computer systems working together as a single system to ensure that mission-critical applications and resources remain available to clients. | Windows Clustering | No |
| Manage all hosts of a Network Load Balancing cluster from a single computer. | Network Load Balancing Clusters | No |
| Manage server disk volumes and automatically copy infrequently-used files to a tape or disk library. | Remote Storage | No |

**Directory services administration**

| Task | Tool | In a Convenience Console? |
|---|---|---|
| Manage the trust relationships between domains. | Active Directory<br>Manage Domains<br>Manage Trusts<br>Manage Forests | Active Directory Management |
| Manage users, computers, security groups and other objects in the Active Directory. | Active Directory<br>Manage Users, Groups, and Computers<br>Group Policy (pre-GPMC) | Active Directory Management |
| Replicate Active Directory information. | Active Directory<br>Manage Sites | Active Directory Management |
| Manage stored object information in Active Directory schemas from the command line (schmmgmt.msc). | Active Directory<br>Manage the Schema | No |
| Delegate control to select users or groups, organizational units and objects. | Active Directory<br>Delegating administration | No |
| Manage Active Directory information via the command line. | Directory service command-line tools | No |
| Manage the Active Directory database from the command line. | Ntdsutil | No |
| Manage Group Policy from the command line (gpedit.msc) | Group Policy (pre-GPMC) | No |

# Convenience Consoles in Windows Server 2003 Administration Tools Pack

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Convenience Consoles in Windows Server 2003 Administration Tools Pack**

Convenience consoles combine related features into Microsoft Management Consoles (MMC). For example, to administer and troubleshoot an Active Directory deployment, administrators generally need to use Active Directory Users and Computers, Active Directory Domains and Trusts, Active Directory Sites and Services, and DNS. Windows Server 2003 Administration Tools Pack includes a saved console in MMC called **Active Directory Management** that contains these tools.

The Windows Server 2003 family contains three Convenience Consoles:

| Convenience Console | Contents | Snap-ins |
|---|---|---|
| Active Directory Management | Manage Users, Groups, and Computers<br>Manage Domains<br>Manage Sites<br>DNS | ADMgmt.msc |
| Public Key Management | Certification Authorities<br>Certificate Templates<br>Certificates for Current User<br>Certificates for Local Computer | PKMgmt.msc |
| IP Address Management | DHCP<br>DNS<br>WINS | IPAddrMgmt.msc |

Convenience Consoles are displayed in **Administrative Tools** on computers running Windows XP Professional with Windows installed. These consoles contain snap-ins included separately in Windows.

For further information, see Windows Server 2003 Administration Tools Pack.

# Troubleshooting Windows Server 2003 Administration Tools Pack

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Broken links in Help files
- Trouble installing Windows Server 2003 Administration Tools Pack
- Windows Server 2003 Administration Tools Pack will not install on Windows XP Professional

**Broken links in Help files**
**Cause:** When Windows Server 2003 Administration Tools Pack is installed on Windows XP Professional, some Help links will be broken. Server and Windows XP Professional Help systems differ.

**Solution:** Install the Server Help on Windows XP Professional. Switch to the other installed Help if links in one Help system are broken. In most cases, this will resolve the broken link problem

**See:** Install Help content from a Windows Server 2003 server operating system onto a computer running Windows XP Professional

**Trouble installing Windows Server 2003 Administration Tools Pack**
**Cause**: You need local administrator credentials to install Windows Server 2003 Administration Tools Pack.

**Solution:** Use the Runas command. For more information, see Runas.

**Windows Server 2003 Administration Tools Pack will not install on Windows XP professional**
**Cause:** Windows Server 2003 Administration Tools Pack cannot be installed on computers running Windows XP Professional without first applying Service Pack 1 to the operating system or installing the hotfix contained in Q329357.

**Solution**: Install the QFE or apply Windows XP Professional with Service Pack 1.

**See:** Article Q329357, "Windows Server 2003 adminpak.msi may not run on Windows XP 2600 Machines," Microsoft Knowledge Base.

# Emergency Management Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Emergency Management Services**
The Microsoft® Windows Server 2003 family provides native support for server operation and management that can be performed without a local keyboard, mouse, or monitor. Emergency Management Services is one element of an out-of-band management solution that you can use to manage servers remotely when the operating system is not running properly. Emergency Management Services is available with or without a video card.

- For recommendations on setting up and using Emergency Management Services most effectively, see Emergency Management Services Best Practices.
- For help with using Emergency Management Services to solve problems with your system, see Emergency Management Services How To....
- For general background information on Emergency Management Services, see Emergency Management Services Concepts.
- For help solving problems that might occur in Emergency Management Services, see Troubleshooting Emergency Management Services.

# Emergency Management Services Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- Best practices for using in-band tools
- Best practices for selecting and configuring hardware
- Best practices for implementing security measures

# Best practices for using in-band tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for using in-band tools**

**Try using in-band tools before using Emergency Management Services.**

- If your server becomes unstable, try to use in-band tools before using Emergency Management Services to troubleshoot your operating system.

  This is important because Emergency Management Services is designed to be used only when in-band means of accessing and managing the server are unavailable. For example, you would use Emergency Management Services to force your system to reset only if you are unable to reset it using the standard, in-band tools. In addition, be aware that Emergency Management Services is not a complete management solution; your ability to access and manage the server at all times depends upon your hardware configuration.

  For more information, see In-band management.

**Enable and configure in-band tools during Setup.**

- Because in-band tools can be difficult or impossible to configure remotely after Setup has completed, it is important to enable and configure these tools during Setup. Examples of in-band tools include Terminal Services Remote Desktop for Administration and Telnet.

  For more information, see In-band management.

# Best practices for selecting and configuring hardware

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for selecting and configuring hardware**

**Ensure that the BIOS supports console redirection and the SPCR table.**

- If there is no Serial Port Console Redirection (SPCR) setting, the basic input/output system (BIOS) either enables it automatically or does not support it. It is preferable to have a BIOS that supports the SPCR table. If this is not the case, you must enable and configure Emergency Management Services manually. You can do this during Setup or by using the Bootcfg.exe tools in the operating system. Configure Emergency Management Services to use the same settings as the BIOS console redirection. Be sure to make a note of the baud rate and the serial port that you set.

- You might find it necessary to enable console redirection in the BIOS settings. COM 1 is the preferred COM port. Make a note of the baud rate and the COM port that you set. You must use this information to attach a serial cable properly to the correct port. You also must use this information to configure your terminal emulator to the correct baud rate.

- If the SPCR table is supported by the BIOS, it can be enabled automatically when BIOS redirection is enabled. If the BIOS has a setting for SPCR or SPCR table, be sure to configure it.

- Configure your server to reboot automatically if a Stop error occurs. This allows the server to reboot without requiring you to access !Special Administration Console (!SAC) or the power switch to reboot it manually. This provides maximum availability if a Stop error occurs.

  For more information, see Firmware console redirection.

**Choose the highest baud rate that your hardware and Emergency Management Services allows.**

- This gives you the best performance.

- This also prevents a slow text-mode Setup process.

- For the preferred Remote Installation Services (RIS) implementation, use a redirecting BIOS that prompts you to boot from the network and a version of Startrom.com that does not redirect.

  For more information, see Configuring Serial Ports.

**When using Emergency Management Services with a serial port, use a null modem cable with carrier detect.**

**Select service processors that support Emergency Management Services.**

- A service processor can provide additional management functionality that complements that of Emergency Management Services. This gives you a more complete remote-management solution.

- Use an original equipment manufacturer (OEM) service processor that provides, at a minimum, reset and BIOS console redirect capabilities.

- You must select a service processor that supports Emergency Management Services. In addition, you might find it necessary to enable the service processor so that it provides that support.

- It might be necessary to update your firmware to enable it to support Emergency Management Services. Service processors can require their own firmware updates in addition to system firmware updates. With the proper firmware updates, your current service processors might support Emergency Management Services.

For more information, see Service processors.

**Configure your server to reboot automatically if a Stop error occurs.**

- This makes it unnecessary for you to access !SAC or the power switch to reboot the server manually. It also provides maximum availability if a Stop error occurs.

**Select hardware and firmware that supports VT-UTF8.**

- If your hardware and firmware do not use the same terminal type that Emergency Management Services uses, it can be extremely difficult to send the proper command escape sequences. This is because the command escape sequences change according to whether the hardware, firmware, or Emergency Management Services is controlling the port.

- Hardware and firmware that support VT-UTF8 or VT100+ provide the best compatibility. To prevent incompatibility, you must at least select hardware and firmware that support VT100.

- It might be necessary to update your system and service processor firmware to enable VT-UTF8 or VT100+ output.

  For more information, see VT-UTF8, VT100+, and VT100 conventions.

# Best practices for implementing security measures

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for implementing security measures**

**Limit the ability of users to access your out-of-band infrastructure.**

- When used incorrectly or maliciously, Emergency Management Services and out-of-band management hardware can be made to disrupt the use of a server. To prevent this problem, limit physical access to the server. This will also keep unwanted users from accessing Emergency Management Services through a serial port. To provide this level of physical security, be sure to protect sensitive areas such as racks, computer labs, server farms, and data centers by locking them.

- Use a terminal concentrator with security features. If you access your servers out-of-band through a terminal concentrator, choose one that provides security when you connect through the network.

- Select service processors with well-designed security implementations. Service processors can provide improved access mechanisms, such as network interfaces. With these improvements, however, come the added security risks of network access. You must ensure that any service processor you enable provides a secure access mechanism.

- Consider setting up a separate network dedicated exclusively to management traffic, including that of Emergency Management Services. A separate network for management traffic adds an additional layer of security. Only known, secure management workstations should be granted access to this network and its out-of-band server connections. The management network must not allow any connections to the Internet, and it must authenticate only trusted users.

  For more information, see Remotely administered servers, Service processors, and Terminal concentrators.

# Emergency Management Services How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Verify that Emergency Management Services is enabled
- Start the computer
- View the POST status and errors
- View and manage the system while the loader and kernel initialize
- View and manage the system during text-mode Setup
- Determine the cause of a pause or failure during GUI-mode Setup
- Resolve system resource problems using SAC
- Respond to a Stop error message
- Respond if SAC and !SAC become unavailable
- Change channels

# Verify that Emergency Management Services is enabled

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify that Emergency Management Services is enabled**
You can use Bootcfg to quickly verify that Emergency Management Services is active in an x86-based or Itanium-based system without connecting devices to the out-of-band port.

- At a command prompt, type:

  **bootcfg/query**

### Note

- You can also check the registry to verify Emergency Management Services installation. For more information on working with the registry, see "Resources" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enabling Emergency Management Services
Emergency Management Services Resources
Troubleshooting Emergency Management Services

# Start the computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the computer**

Depending on the remote computer's hardware configuration, you can use one or more of the following approaches to start your computer. For more information, see your hardware manufacturer's documentation.

- Use the interface supplied by a service processor, if one is present, to turn on the remote computer.
- Use the interface supplied by a power switch or uninterruptible power supply (UPS) with remote computer management capabilities, if one is present, to turn on the computer.
- Use wake-on-LAN functionality by sending a network request that "wakes" the remote computer from sleep mode. This feature must be supported by both the system motherboard and firmware. Follow the instructions provided by your computer and network adapter motherboard manufacturers to enable this functionality.
- If none of these approaches is appropriate, you can go to the physical location of the target computer and start it by pressing the power button.

### Note

- For general background information on wake-on-LAN functionality, see "Routing and Remote Access" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

System-state diagram
VT-UTF8, VT100+, and VT100 conventions
Emergency Management Services Resources
Troubleshooting Emergency Management Services
Routing and Remote Access

# View the POST status and errors

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the POST status and errors**

Depending on the remote computer's hardware configuration, one or more of the following approaches can help the computer complete the power-on self test (POST).

For all of these options, you must configure the remote computer's firmware or service processor, if one is present, to redirect the POST. You will then be able to view the progress of the POST. For more information, see to your hardware manufacturer's documentation.

- Connect to the remote computer's service processor, if one is present, and restart the computer.
- If a service processor is present, and the out-of-band tools are available, perform diagnostic procedures using the service processor's out-of-band tools.
- Use a power switch or uninterruptible power supply (UPS) with remote computer management capabilities, if one is present, to restart the remote computer.

### Note

- Windows Server 2003 operating systems cannot start if the remote computer stops responding before completing the POST.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

System-state diagram
Emergency Management Services Resources
Troubleshooting Emergency Management Services

# View and manage the system while the loader and kernel initialize

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view and manage the system while the loader and kernel initialize**

- Ensure that Emergency Management Services is enabled on the remote computer.
- If the remote computer's operating system fails to initialize, note whether an error message appears. If you do not receive an error message, use the **D** command in either Special Administration Console (SAC) or !Special Administration Console (!SAC) to view the log.
- Restart the remote computer using SAC or !SAC and correct the problem.

## Notes

- You can use both firmware console redirection and Emergency Management Services console redirection to monitor the remote system's startup as if you were at a local console. It is preferred that the remote computer's firmware be able to pass the configuration settings of the out-of-band port to the operating system using the Serial Port Console Redirection (SPCR) table. If your system uses Extensible Firmware Interface (EFI), Windows reads the boot settings from the SPCR table. If your system does not use EFI, Windows reads the boot settings from the boot.ini file.
- For more information, see "Verifying that Emergency Management Services is configured" in Related Topics.
- For more information about how to restart the remote computer, see "Startup options" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System-state diagram
Special Administration Console (SAC) and SAC commands
Emergency Management Services Resources
Startup options

# View and manage the system during text-mode Setup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view and manage the system during text-mode Setup**

Depending on the remote computer's hardware configuration, one or more of the following approaches can help you view and manage the system during text-mode Setup. If the remote computer has a serial port console redirection (SPCR) table, Emergency Management Services will automatically display the pertinent information. For more information, see your hardware manufacturer's documentation.

- For x86-based systems, use Remote Installation Services (RIS) to install Windows Server 2003 operating systems.
- If you are not using unattended Setup files, you must choose Express Setup when prompted.
- If Setup fails, note any error messages that appear and attempt to correct any problems you discover.

## Notes

- When you are using GUI-mode Setup, you must perform an unattended Setup. This allows GUI-mode Setup to proceed without prompting you for input. If you do not perform an unattended Setup, GUI-mode Setup will fail.
- You can use RIS-based installations to perform unattended Setups. When you RIS-based installations for this purpose, it is not necessary to use the operating system CD or a floppy disk containing a Winnt.sif file. For more information on RIS and the Winnt.sif file, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

System-state diagram
Special Administration Console (SAC) and SAC commands
RIS and PXE requirements
Emergency Management Services Resources

# Determine the cause of a pause or failure during GUI-mode Setup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To determine the cause of a pause or failure during GUI-mode Setup**

- To determine the reason for a pause or failure during a GUI-mode Setup on a remote computer, use the **ch** command in the Special Administration Console (SAC), or press ESC+TAB. You can use either of these methods to view one of the three setup logging channels.

- If GUI-mode Setup pauses, use the **cmd** command in SAC to open a command-prompt channel. You can then use this command-line environment to navigate the system, and, if necessary, repair the $winnt$.sif file located in the *systemroot*\System32 folder.

## Notes

- When you are using GUI-mode Setup, you must perform an unattended Setup. You can then respond from your local computer when the remote computer prompts you for input during setup. If you do not perform an unattended Setup, GUI-mode Setup will fail.

- If setup fails, but the remote computer is still responding, there might be a problem with the setup files.

- If you examine the setup log channels and are still unable to determine the reason for the pause or failure of GUI-mode Setup, you might need to monitor a local display to determine the cause.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System-state diagram
Emergency Management Services components
Special Administration Console (SAC) and SAC commands
Emergency Management Services Resources

# Resolve system resource problems using SAC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To resolve system resource problems using SAC**

- Use the **T** command in Special Administration Console (SAC) to view the processes that are running on the remote computer, the amount of CPU time they have consumed, and the amount of memory they are using. See "Task list headings" in Related Topics for more information on analyzing the output from the **T** command.
- Use the **R** and **L** SAC commands to modify the priority of any processes that might be preventing the use of in-band tools on the remote computer.
- Use the **K** SAC command to end a process that is causing a problem on the remote computer.
- If the problem persists, or if you do not know which process is causing the problem, use the **restart** SAC command to restart the remote computer.

## Notes

- Before you use SAC to resolve a system resource problem, first try using in-band tools. Use SAC only if using in-band tools fails to resolve the problem.
- You might be able to determine which processes are consuming the most resources by comparing two consecutive task lists.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System-state diagram
Special Administration Console (SAC) and SAC commands
Task list headings
Emergency Management Services Resources

# Respond to a Stop error message

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To respond to a Stop error message**

- View the Stop error message information remotely, and use !Special Administration Console (!SAC) to restart the remote computer. Viewing the kernel log with the **D** command can help to determine the cause of the error, especially if the Stop error message appeared during the startup process.

## Notes

- To display !SAC, press ENTER when viewing the Stop screen.
- You can manually generate a Stop error message and force a memory dump file to be created by using the SAC **crashdump** command. To do so, however, you must have crash dump enabled.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System-state diagram
Special Administration Console (SAC) and SAC commands
!Special Administration Console (!SAC) and !SAC commands
Troubleshooting Emergency Management Services

# Respond if SAC and !SAC become unavailable

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To respond if SAC and !SAC become unavailable**

Depending on the remote computer's hardware configuration, one or more of the following approaches can be helpful if Special Administration Console (SAC) and !SAC become unavailable. For more information, see your hardware manufacturer's documentation.

- Perform diagnostic procedures using the service processor's out-of-band tools if a service processor is present, and if the tools are available.
- Restart the remote computer by using the interface supplied by a service processor, if one is present.
- Restart the remote computer by using the interface supplied by a power switch or an uninterruptible power supply (UPS) with remote computer management capabilities, if one is present.
- If none of these approaches is appropriate, start the remote computer by pressing the power button.

**Note**

- If you restart the remote computer using a power switch or UPS, its firmware must be configured to turn on when power is restored.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

System-state diagram
Emergency Management Services Resources
Troubleshooting Emergency Management Services

# Change channels

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change channels**

- Press ESC+TAB.

## Notes

- To return to the Special Administration Console (SAC) channel, press ESC+TAB, and then press **0**.
- To learn more about managing your system using Emergency Management Services channels, type:

  **ch -?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Emergency Management Services components
Special Administration Console (SAC) and SAC commands
Emergency Management Services Resources
Troubleshooting Emergency Management Services

# Emergency Management Services Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Emergency Management Services overview
- System-state diagram
- Understanding Emergency Management Services
- Setting up Emergency Management Services
- Using Emergency Management Services
- Emergency Management Services Resources

# Emergency Management Services overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Emergency Management Services overview**

Emergency Management Services is a new feature that is included with the Microsoft® Windows Server 2003 family. You can use this feature to perform remote-management and system recovery tasks when the server is unavailable through the network or other standard remote-administration tools and mechanisms. Emergency Management Services is available for x86-based and Itanium-based systems.

When a server is operating normally, you can manage it remotely by using a variety of administrative tools that are provided in your Windows Server 2003 operating system. Using these tools to access and manage the server is referred to as in-band management. It usually requires a network that is fully available. For more information on this topic, see In-band management.

When the network is unavailable, you can use Emergency Management Services to access and manage the server. You do so through an out-of-band connection. This method does not require operating-system network drivers. In addition, you can use an out-of-band connection to troubleshoot the server when it is not fully initialized and functioning. Typically, you can do this without being physically present at the computer. When the appropriate out-of-band hardware is in place, the server requires physical access only when hardware must be installed or replaced. Emergency Management Services also eliminates the need for a keyboard, mouse, monitor, or video adapter on the remotely administered server. Emergency Management Services is available with or without a video card. For more information about this topic, see Out-of-band management.

With Emergency Management Services, you use a straightforward but powerful terminal text mode instead of a graphical user interface (GUI). As a result, you can use Emergency Management Services with a wide range of communication media (including serial connections) and current out-of-band infrastructure (including terminal concentrators and terminal emulators). It also provides interoperability with other platforms, such as UNIX.

Emergency Management Services features are fully or partially available when the operating system loads, when it is active (setting up and starting up), and when it is in distress (responding improperly or experiencing a Stop error). For a visual overview of these situations, see the System-state diagram.

For additional information on how to integrate Emergency Management Services into your infrastructure, see "Deploying Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

# System-state diagram

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System-state diagram**




**System-state diagram conventions**

- The shaded boxes represent situations in which Emergency Management Services features are fully or partially available.
- The dashed lines represent transitions into distress or failure states. The operating system might not be able to recover unless assisted. You can use Emergency Management Services to attempt remote recovery if the system is at least partially functioning.
- The ability to manage the system remotely, which is represented by the non-shaded boxes, varies by system. Consult your system documentation to determine whether your system includes features, such as firmware redirection or a service processor, that can affect your ability to manage the system remotely.
- The solid lines represent normal or expected behavior. In these situations, you would expect the computer to turn on, initialize system firmware, and start the operating system without problems.
- The numbers are a sequential representation of the system states.
- For information about system-recovery procedures that correspond to situations represented in this diagram, see Emergency Management Services How To....

# Understanding Emergency Management Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Emergency Management Services**
This section covers:

- Remotely administered servers
- In-band management
- Out-of-band management
- Emergency Management Services components
- RIS and PXE requirements
- Firmware console redirection
- Emergency Management Services console redirection
- Service processor console redirection

# Remotely administered servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remotely administered servers**
For servers equipped with the proper firmware, Emergency Management Services provides functionality that you can use to administer a server remotely. You can usually do so without a local keyboard, mouse, or display device. Except for hardware maintenance and replacement, all administrative functions that you can accomplish locally should be available remotely. This includes starting your system and performing system-recovery tasks.

Remote administration does present security issues, however. For more information on securing a system that is managed with Emergency Management Services, see Best practices for implementing security measures.

You should be able to perform the following administrative tasks on a remotely administered server:

- Start up a server.
- Change basic input/output system (BIOS) settings.
- View power-on self test (POST) results.
- Install Windows using a standard mechanism such as Remote Installation Services (RIS).
- Select which operating system to start.
- Manage Windows using standard management tools such as Windows Management Instrumentation (WMI), Terminal Services Remote Desktop for Administration, Microsoft Management Console (MMC), Telnet, Windows Script Host, and non-Microsoft tools.
- Manage a Windows Server 2003 operating system when it is unavailable over the network.
- View system Stop errors.
- Reset a server.

To support operations performed without a keyboard, mouse, or display device attached to the system, Windows Server 2003 operating systems are able to boot without most video card support. For Advanced Configuration and Power Interface (ACPI) systems, the Windows Server 2003 family supports operating without a legacy keyboard controller. Universal serial bus (USB) systems running these operating systems can also support optional keyboard and mouse attachments. Consult your hardware documentation to determine whether your system supports this functionality.

Windows Server 2003 operating systems use the out-of-band management port to support Remote Installation Services (RIS). This makes it easier to install and repair servers remotely. For more information about RIS, see RIS and PXE requirements.

Managing a server that is running normally is referred to as in-band management. It typically involves connecting through the network directly to the server. A server that is functioning normally can be administered remotely with tools provided in the Windows Server 2003 family and listed earlier in this topic.

Troubleshooting a server that is not functioning normally often requires out-of-band management because you must access the server without relying on the network. To perform out-of-band management, you must establish a secure connection through a phone line or serial port, or through an additional network connection (possibly on a secondary network). The goal of out-of-band management is always to return the server to a fully functioning state where in-band connection is available.

Although you can use Emergency Management Services to manage the server remotely when the system is not functioning normally, Emergency Management Services must be set up and configured with the appropriate hardware and firmware. This is necessary so you can access the full range of its out-of-band management capabilities. When you have the proper hardware and Emergency Management Services configuration, out-of-band console input/output (I/O) support is provided for all kernel components, the loader, Setup, Recovery Console, and Stop errors.

When the server is functioning properly, support is also provided for a text-mode management console called Special Administration Console (SAC). For more information, see Emergency Management Services components.

For more information about hardware requirements for Emergency Management Services, see Selecting Hardware for Emergency Management Services.

# In-band management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**In-band management**

Windows Server 2003 operating systems provide a variety of in-band tools you can use to manage computers remotely.

Unlike out-of-band connection, in-band connection is available only when the server is initialized and functioning properly. In-band connection relies on operating-system network drivers to establish computer connections.

The most common in-band remote-management hardware device is the network adapter. Analog modems and Integrated Services Digital Network (ISDN) adapters are also classified as network adapters.

You can use remote-access connections to manage servers. The remote-access connection is made through the in-band port that resides on the target server.

You can use the following administrative tools to perform remote-management tasks through an in-band connection:

- Microsoft Management Console (MMC)
- Systems Management Server
- Telnet
- Terminal Services Remote Desktop for Administration
- Windows Management Instrumentation (WMI)
- Remote Command Service (Rcmd.exe and Rcmdsvc.exe). For more information on this tool, see Windows Resource Kit Tools Help.

You can also use appropriate non-Microsoft management tools to perform remote management tasks through an in-band connection.

If you lose the ability to manage your server through an in-band connection, you can manage the server using an out-of-band connection. For more information on this topic, see Out-of-band management.

# Out-of-band management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Out-of-band management**
When you are unable to manage a server over an in-band connection, you can access the server through an alternate but reliable connection called an out-of-band connection. An out-of-band connection does not depend upon operating-system network drivers, and is available even when your server has not loaded or is not functioning properly.

The following list describes situations in which you cannot use in-band management, and where it might therefore be appropriate to use out-of band management:

- The system is turned off. Because Emergency Management Services is unavailable, special hardware or firmware is required to manage systems in this situation.
- The basic input/output system (BIOS) is conducting the power-on self test (POST). Because Emergency Management Services is unavailable, special hardware or firmware is required to manage systems in this situation.
- The server is not functioning properly due to a Stop message event.
- The server is very low on resources, which causes the network driver to be very slow or makes it unable to respond to requests.
- The network stack is malfunctioning or has failed.
- An operating-system component is running (such as the loader or Recovery Console) that does not support in-band communication.
- The server is not yet fully initialized

The serial port is the hardware interface most commonly used by Emergency Management Services. There are several advantages to this. First, the serial port is simple and reliable enough to allow flexibility and availability. Second, the communication standards that govern serial ports are well defined. This makes interoperability with non-Microsoft systems such as UNIX possible. The third advantage is that most x86-based servers have at least one serial port.

For more information about in-band connections and in-band administrative tools that are available in Windows Server 2003 operating systems, see In-band management.

# Emergency Management Services components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Emergency Management Services components**
Emergency Management Services consists of components that are standard features of Windows Server 2003 operating systems, and to which console redirection functionality has been added. Emergency Management Services also includes two remote-management consoles that are unique to it.

**Components that support Emergency Management Services console redirection**
The components that support Emergency Management Services console redirection include Setup loader, text-mode Setup, the Recovery Console, Remote Installation Services (RIS), loader, and Stop error messages. When the operating system is configured to run with Emergency Management Services, these components will redirect their output to the out-of-band management port and to the video card, if one is attached. Emergency Management Services is available, however, with or without a video card. All Emergency Management Services output is accessible using a terminal emulator. For more information about console redirection, see Emergency Management Services console redirection.

There are also Emergency Management Services-specific versions of the Startrom.com file that is used in Remote Installation Services (RIS) installations. For more information about this, see RIS and PXE requirements.

**Remote-management consoles: SAC and!SAC**
The two remote-management consoles designed for Emergency Management Services are Special Administration Console (SAC) and !Special Administration Console (!SAC). You access these consoles through an out-of-band remote management port using terminal software that supports VT-UTF8, VT100+, or VT100.

**Special Administration Console (SAC)**
Special Administration Console (SAC) is the primary Emergency Management Services command-line environment hosted by Windows Server 2003 operating systems. It is separate from the command-line environment and provides different functionality.

Because SAC is available early in the boot process, you can use it to manage the server during normal system operation and initiation of most Windows Server 2003 family components. You can also use it when the system is in Safe Mode and during GUI-mode Setup. When Emergency Management Services is enabled, SAC remains active as long as the kernel is running.

SAC provides a set of commands you can use to perform a number of management tasks that help return your system to a normally functioning state. These tasks include:

- Restarting or shutting down the server.
- Viewing a list of processes that are currently active.
- Ending processes.
- Setting or viewing the Internet Protocol (IP) address of the server.
- Generating a Stop error to create a memory dump file.
- Starting and accessing command prompts.

For more information about using SAC commands, see Special Administration Console (SAC) and SAC commands.

**!Special Administration Console (!SAC)**
!Special Administration Console (!SAC) is an auxiliary Emergency Management Services command-line environment that is hosted by Windows Server 2003 operating systems. It also accepts input and sends output through the out-of-band port. !SAC is a separate entity from both SAC and Windows Server 2003 family command-line environments. After a specific failure point is reached, Emergency Management Services components determine when the shift should be made from SAC to !SAC. !SAC becomes available automatically if SAC fails to load or is not functioning.

!SAC provides a subset of SAC commands you can use to restore system functionality during a failure. Its two primary

functions are:

- To redirect Stop error message text.
- To restart the computer if SAC becomes unavailable.

For more information about using !SAC commands, see !Special Administration Console (!SAC) and !SAC commands.

**Windows command-prompt channels and Emergency Management Services**

The Special Administration Console Helper service (Sacsvr) is a user-mode service you can use to create Windows command-prompt channels when you type **cmd** at the SAC prompt. With the command-prompt channels, you can use in-band tools on a remotely administered server when an in-band connection is not available.

For example, when the network is down, you can still run text-based utilities as you would through Telnet. These tools include command-shell utilities such as **dir**, and text-based console programs such as the boot configuration utility (Bootcfg.exe). You cannot run programs with a graphical user interface (GUI) such as Notepad.exe, however.

Although access to in-band tools through the command-prompt channels provides extensibility and convenience, it does not guarantee that the tools will perform as reliably as they would when run over an in-band connection; the out-of-band functionality and robustness of Emergency Management Services applies only to SAC. In addition, Windows command-prompt channels require system resources, which might cause SAC to be available but the command prompt to be unavailable.

To use a command-prompt channel, you must provide valid logon credentials to a local or domain account, and log on to each command-prompt instance. Press ESC+TAB to switch back and forth between the command prompt channels and SAC. If a command prompt channel becomes unresponsive, you can close it and then open another.

For more information about the SAC commands for managing channels, see Special Administration Console (SAC) and SAC commands.

**Setup log access**

SAC provides access to the setup logs during GUI-mode Setup. You can press ESC+TAB to switch between the setup logs and SAC. When accessing the setup logs from Emergency Management Services, you can see which portions of Setup have completed and whether any errors have occurred. This is a very useful way to check the progress of your setup and to diagnose setup failures.

The three setup log channels are as follows:

**setuplog.txt**
Monitors setup progress.

**setupact.log**
Displays any warnings during setup.

**setuperr.log**
Displays any errors that might occur during setup.

# RIS and PXE requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RIS and PXE requirements**

Remote Installation Services (RIS) supports remote installations of Windows Server 2003 operating systems. This means that you can initiate RIS-based installations remotely on a computer without being physically present at the computer. RIS eliminates the need to use either the operating system CD or a floppy disk containing a Winnt.sif file. For more information, see Enabling Emergency Management Services during a new installation.

- RIS support is available on both x86 and Itanium-based systems.

For remote installations to be possible, the client computers must have a network adapter that is Pre-Boot eXecution Environment (PXE) capable. You can then perform installations using a network connection on systems without CD-ROM drives and on systems that have non-formatted or non-partitioned hard disks.

The firmware should recognize the PXE-capable network adapter as a valid boot device and as part of the boot order selection. In addition, the firmware should provide the option of using the F12 key to initiate a remote boot using a keystroke, thus circumventing the boot order. You can use this ability to change the boot order to implement emergency recovery procedures if necessary.

**Startrom.com**

The Startrom.com file is the first file that an x86-based RIS client downloads to begin the boot process. There are multiple versions of this file. Based on the PXE capabilities of your firmware, you can use Emergency Management Services-specific versions of Startrom.com to make use of console redirection for x86-based RIS clients. The Boot Information Negotiation Layer service (BINLSVC) creates an answer file that allows the rest of Setup to run without user input. Itanium-based RIS clients do not use Startrom.com, however. Instead, you can configure the Extensible Firmware Interface (EFI) to redirect output, and then use the PXE-capable network adapter to boot the computer.

Although you can use the Startrom.com files with systems that are not designed to work with Emergency Management Services, it is better to rely on the firmware to redirect its console output to the out-of-band management port and to use a standard Startrom.com file. This allows Startrom.com to be more flexible in its use of baud rates.

- **Hdlscom1.com and Hdlscom2.com.** For systems that do not support firmware console redirection, these files display the **Press F12 for network boot** prompt, using console redirection to serial port 1 or 2. Users can either press F12 to proceed with the boot process, or they can exit the boot process by not pressing F12. If a user presses the F12 key when prompted, Hdlscom1.com and Hdlscom2.com cause the PXE client to boot into the Client Installation Wizard only when the basic input/output system (BIOS) attempts a network boot. These files are installed by default with RIS.

- **Hdlscom1.n12 and Hdlscom2.n12.** For systems that support firmware console redirection, these files do not display the **Press F12 for network boot** prompt and do not wait for user input. When you use one of these files, the PXE client will always boot into the Client Installation Wizard when the BIOS attempts a network boot.

All Emergency Management Services-specific versions of Startrom.com rely on the BIOS to redirect their output to the out-of-band management port and to the video card, if one is attached. Emergency Management Services is available with or without a video card.

To use a Hdlscom*x*.com or Hdlscom*x*n12 file, you must first rename it to Startrom.com, and then copy it to the BINLSVC server. By default, these files are located in the *ServerName*\reminst\oschooser\i386 folder, where *ServerName* is the name of the RIS server running the BINLSVC.

**Client Installation Wizard**

The Client Installation Wizard automatically directs its entire user interface through the out-of-band port and a local console if one is attached to the computer. The Client Installation Wizard determines which out-of-band port to use by either reading the Serial Port Console Redirection (SPCR) table (if available) or by using the same port as Startrom.com. For more information, see Enabling the SPCR table.

For more information about RIS, see Remote Installation Services.

# Firmware console redirection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Firmware console redirection**

Firmware can provide redirection before the operating system starts. See the System-state diagram for a visual overview of how firmware and Emergency Management Services console redirection work together.

Types of firmware include both basic input/output system (BIOS) for x86-based computers and Extensible Firmware Interface (EFI) for Itanium-based computers. You can perform the following remote-management tasks using firmware console redirection:

- **View the power-on self-test (POST) status remotely.** Typically, computers with firmware that supports console redirection allow a system to complete the POST without a keyboard, mouse, or display device attached.

- **View and edit firmware settings remotely.** You can use the built-in firmware configuration program to change computer hardware settings.

- **Start a Remote Installation Services (RIS)-based setup.** If the firmware supports Pre-Boot execution Environment (PXE), you can use firmware console redirection to initiate Remote Installation Services (RIS)-based setups. In order to support PXE, the firmware must recognize the network adapter as a boot device. If necessary, the firmware must also allow you to change the settings of the BIOS boot order. For more information, see RIS and PXE requirements.

- **Respond to the prompt remotely when starting your computer. Press any Key to Boot from CD** You can do this by using the operating system CD. Before you begin a CD-based Windows Server 2003 operating system setup, enable firmware console redirection so you can view and respond to this prompt remotely. Your computer's BIOS must support the bootable CD (no emulation mode) format. Check with your computer manufacturer if you do not know whether your BIOS includes this feature. The BIOS specifies the boot order. The CD-ROM should be listed as the first boot device. Emergency Management Services is self-configured during a bootable CD installation by reading the Serial Port Console Redirection (SPCR) table.

**Notes**

- The firmware must be able to release control of the serial port to the operating system when the operating system starts.

- Firmware console redirection is typically available in text mode only.

Without firmware console redirection, you cannot perform out-of-band remote management during the period of time between system restart and the initial load phase of the Windows Server 2003 operating system. When firmware redirection is coupled with Emergency Management Services, however, you can manage and monitor a server remotely as soon as the POST process begins and until the operating system components are initialized. It is recommended that the firmware support the SPCR table. For more information, see Enabling the SPCR table.

# Emergency Management Services console redirection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Emergency Management Services console redirection**
Windows Server 2003 family components such as the Windows loader (Ntldr) can provide Emergency Management Services console redirection. Console redirection makes video output appear in text format when the output is sent to the out-of-band port and local monitor, if one is attached. The computer continues to receive keyboard input from the out-of-band port and the local keyboard, if attached. While the operating system is booting, the console takes over redirection from the firmware. See the System-state diagram for a visual overview of this process.

**Important**

- Your computer must have firmware console redirection for you to be able to view information from the remotely administered server. For more information, see Firmware console redirection.

Emergency Management Services does not provide console redirection for the following:

- The **Press any key to boot from CD** prompt that appears when you start your system by using the operating system CD.
- The Windows Server 2003 family GUI-mode Setup.
- Disk related error messages caused by problems such as a damaged master boot record (MBR).
- Startrom.com at baud rates greater than 9600 baud.
- For Itanium-based systems, Emergency Management Services relies on the firmware to redirect the operating system boot menu.

With Emergency Management Services console redirection, you can manage a server remotely through an out-of-band connection to do the following:

- Verify the initial load of Windows Server 2003 family components when administrative tools for in-band management are not yet available.
- For x86-based systems, select an operating system to load on multiple boot systems.
- View the progress of the Windows Server 2003 family text-mode Setup.
- For systems without firmware console redirection, respond to the **F12** prompt remotely when using Remote Installation Services (RIS) to initiate a Setup program.
- Respond to Windows Server 2003 family text-mode Setup prompts.
- View Stop error messages remotely when a system problem has occurred.
- Use the Recovery Console to troubleshoot startup problems remotely.
- Start and access command prompts.

# Service processor console redirection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service processor console redirection**

If your system includes a service processor, enhanced out-of-band management features, including console redirection, might be available to you. When selecting service processor hardware to use with Emergency Management Services, ensure that the hardware supports the following functions:

- Turn the computer on.
- Turn the computer off.
- Reset.
- Console redirection.
- Field-replaceable components (for example, service processors that are installed in sockets rather than soldered permanently in place).

For more information about selecting a service processor to complement Emergency Management Services, see Service processors. If your computer is currently equipped with a service processor, check your system documentation for more information about the out-of-band management features it provides.

# Setting up Emergency Management Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up Emergency Management Services**

This section covers:

- Selecting Hardware for Emergency Management Services
- Configuring Hardware for Emergency Management Services
- Selecting client terminal software for Emergency Management Services
- Enabling Emergency Management Services

# Selecting Hardware for Emergency Management Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Selecting hardware for Emergency Management Services**
Emergency Management Services must be integrated with properly designed hardware in order to manage Windows Server 2003 operating systems completely and successfully.

The three key areas you must address to fully integrate your platform with Emergency Management Services are as follows:

- **Ensure that you have appropriate hardware** for an out-of-band management port. Windows Server 2003 operating systems support a legacy serial port and hardware that uses the Universal Asynchronous Receiver/Transmitter (UART) interface provided by a service processor.

- **The user interface** for the system firmware and out-of-band management service processor should use the same communication medium as Emergency Management Services. If the firmware is using the out-of-band port, the firmware should deliver the medium over to Emergency Management Services cleanly after the power-on self test (POST) process has completed. The firmware should also use the same terminal conventions as Emergency Management Services, using the same keystrokes to represent keys such as F12. These criteria are recommended, not required. For more information, see VT-UTF8, VT100+, and VT100 conventions.

- **The firmware** must be able to pass the configuration settings of the out-of-band port to the operating system using the Serial Port Console Redirection (SPCR) table. The firmware stores configuration settings in the SPCR table, and the operating system reads them if the appropriate settings are in the Boot.ini file or the Extensible Firmware Interface (EFI). For more information, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

For additional information about selecting hardware, see Support resources.

This section covers:

- Serial ports

- Service processors

- Terminal concentrators

# Serial ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Serial ports**
The most common out-of-band hardware interface is the serial port, also called the COM port. This serial port must be a 16550 Universal Asynchronous Receiver/Transmitter (UART). If your system firmware is compatible with Emergency Management Services, it is recommended that you configure it to use the same serial port settings. Emergency Management Services and the Windows debugger cannot share the same COM port. The system can have only a single out-of-band management port. Windows Server 2003 operating systems do not support a configuration in which one out-of-band port sends outgoing communication and a different out-of-band port (on a second device) receives incoming data.

For more information about serial ports, see:

- Enabling firmware console redirection
- Configuring Serial Ports

For additional information about selecting hardware for Emergency Management Services, see:

- "Emergency Management Services" at the Microsoft Windows Resource Kits Web site
- Support resources

# Service processors

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service processors**

To use Emergency Management Services console redirection, Special Administration Console (SAC), or !Special Administration Console (!SAC), the operating system loader or kernel must be at least partially functioning. Therefore, when severe system problems cause a computer to completely stop responding, you must have a remote-management medium that can function entirely without the operating system. Service processors integrated onto system motherboards and Peripheral Component Interconnect (PCI) adapters provide such a medium.

Service processors can function this way because they operate independently of the processors in a computer. In addition, they use their own custom firmware and can even use a separate power supply to enhance reliability. You can communicate directly with a service processor through an out-of-band connection.

The management tools, features, and client interfaces provided by service processors differ from one manufacturer to another. They can range in complexity from a simple Telnet console to a rich Web-browser environment. You should evaluate the management features and troubleshooting tools provided by a service processor to ensure that they meet the needs of your organization.

Service processors can also provide support through media that are more complex than serial ports, such as Transmission Control Protocol/Internet Protocol (TCP/IP) on Ethernet. Service processors are typically integrated into system motherboards or add-in PCI adapters and function separately from the system's CPU.

Some service processors are also called Application Specific Integrated Circuits (ASICs). Depending on their original equipment manufacturer (OEM), service processors or ASICs can be implemented with many different types of out-of-band connection ports, including serial ports, Universal Serial Buses (USBs), modem connections, and RJ-45 Ethernet ports. The type of out-of-band connection you use determines which additional components are required to allow out-of-band access to your servers.

When selecting service processor hardware for use with Emergency Management Services, verify that the following functions are supported:

- Turn on the computer
- Turn off the computer
- Reset
- Access to Emergency Management Services
- Console redirection
- Field replaceable units (FRUs); for example, a hard disk

**Internal hardware interface requirements**

To provide complementary functionality, Emergency Management Services and service processors or ASICs must be available for remote management. The service processor must provide a pass-through connection between the operating system and the remote management client. With this functionality, you can use both Emergency Management Services and the service processor through a single connection medium.

The service processor must include an appropriate 16550 or 16450 Universal Asynchronous Receiver/Transmitter (UART) in its system input/output (I/O) as an internal hardware interface for Emergency Management Services. This UART hardware must be available to Emergency Management Services at all times, including when the operating system loader is running, when the operating system is fully loaded, and when a Stop error has occurred.

The UART interface must be fully described in the Serial Port Console Redirection (SPCR) table, or, in Itanium-based systems, through the Extensible Firmware Interface (EFI) console device path. For more information, see Enabling the SPCR table.

**User interfaces for service processors**

Original equipment manufacturers (OEMs) can define a range of user interfaces for their service processors. Serial ports and RJ-45 Ethernet interfaces are examples of two such user interfaces. Some interfaces might support more than one communication protocol; for example, Ethernet supports Telnet and Hypertext Transfer Protocol (HTTP), as well as others. Other interfaces use a single communication channel, such as a serial port.

With an interface that uses a serial port as its communication channel, you can use only one out-of-band management tool (Emergency Management Services or the service processor) at a time. Emergency Management Services and the service processor must share the interface reliably. Integration between Emergency Management Services and the service processor is especially important during the operating system's boot process because you are interacting with both Emergency Management Services and the service processor using the same serial interface. In addition, you are relying on both Emergency Management Services and the service processor to follow the same I/O conventions.

After a system restart, the service processor must release control of the serial port to Emergency Management Services. For this to occur, internal UART hardware must be available so that the service processor and Emergency Management Services can communicate with each other. If the service processor and Emergency Management Services both use VT100 conventions, communication will be successful, and the service processor can be invoked and exited. For more information about this, see VT-UTF8, VT100+, and VT100 conventions.

When the user interface is a multiple-channel interface, such as an RJ-45 Ethernet, the requirements for providing access to Emergency Management Services change. Access to Emergency Management Services must be available simultaneously with the out-of-band functionality provided by the service processor. The terminal interface that you use to manage the server should support VT100+ conventions.

**Power switches or uninterruptible power supplies (UPS) with remote computer management capabilities**
A power switch or uninterruptible power supply (UPS) provides control of the power that is supplied to the server. It can also provides basic remote-management capabilities, such as the ability to reset the system or turn it off. A power switch or UPS can also allow the serial communication between a management system and a server to pass through its port. That is, the port on the computer used for management can have a serial connection to an external serial port on the UPS, which can then have a serial connection to the server.

If you use a power switch or a UPS, it must provide unrestricted communication between the management computer and the Windows Server 2003 family out-of-band tools. A power switch or UPS can interrupt the out-of-band tools by using an escape sequence that is sent by the management computer. To use a power switch or UPS, the server must be configured to boot automatically when power is applied.

The following requirements must be met before you use a power switch or UPS with Emergency Management Services:

- The UPS must be able to reset the system.
- The UPS user interface must support the escape sequences and VT100+ conventions.

It might be preferable for the power switch or UPS to provide terminal concentrator functionality. For more information, see Terminal concentrators.

The following resources provide additional information about selecting hardware for Emergency Management Services:

- "Emergency Management Services" at the Microsoft Windows Resource Kits Web site
- Support resources

# Terminal concentrators

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal concentrators**

A terminal concentrator is a hardware device that consolidates serial access to multiple servers into a single networked device. You can use this device to monitor a large number of servers simultaneously from one location.

Terminal concentrators include many serial ports connected to multiple servers using null modem cables. Typically, you access terminal concentrators over the network through the Telnet protocol. Terminal concentrators provide an interface through which you can remotely view data on multiple servers that use serial ports as their out-of-band connection.

Terminal concentrators can improve your management of servers because they can establish in-band connections to the servers and then perform out-of-band management tasks. In addition, terminal concentrators make it easier to manage servers for the following reasons:

- You can use terminal concentrators to manage multiple servers without needing to be within a serial cable's distance to the computer.
- Several administrators can simultaneously view the output of different servers.
- Using an out-of-band connection, you can use terminal concentrators to monitor servers methodically. You can also manage multiple servers from one location.

Several companies manufacture terminal concentrators; their setup, features, and configuration details vary. When assessing the appropriateness of a particular terminal concentrator, consider the following:

- The number of serial ports available.
- Built-in Telnet security features, such as passwords.
- Remote-access capabilities.
- The number of Ethernet ports available.

Telnet security features are not standard across terminal concentrators. If your device does not include security features, consider using a secondary private management network accessible through a direct-dial remote access connection or a virtual private network (VPN) connection. You can also use a router to secure network traffic going to the terminal concentrator.

For more information, see Checklist: Installing and configuring the router and Checklist: Installing and configuring the remote access server.

Using Telnet over Internet Protocol security (IPSec) can enhance security. See Internet Protocol Security (IPSec) for additional information.

For more information on terminal concentrators, see Configuring terminal concentrators.

For additional information about selecting hardware for Emergency Management Services, see:

- "Emergency Management Services" at the Microsoft Windows Resource Kits Web site
- Support resources

# Configuring Hardware for Emergency Management Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring hardware for Emergency Management Services**
This section covers:

- Configuring Serial Ports
- Configuring terminal concentrators

# Configuring Serial Ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring serial ports**
This section covers:

- Configuring serial port hardware settings
- Enabling firmware console redirection
- Enabling the SPCR table
- Selecting the terminal definition setting

# Configuring serial port hardware settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring serial port hardware settings**

The serial port, also known as the COM port, is the most common out-of-band hardware interface. For information about selecting serial ports that complement Emergency Management Services, see Serial ports.

Emergency Management Services uses the first serial port (COM 1 at 3F8) by default. If a device already occupies that serial port, you can either move the device to another serial port or use another serial port for Emergency Management Services.

To prepare serial ports for out-of-band communication, you must first verify that motherboard serial ports are enabled. This is important if you have not already attached serial devices to determine whether the ports are working. To conduct this test, follow the manufacturer's instructions on how to access the x86-based basic input/output system (BIOS) or Itanium-based Extensible Firmware Interface (EFI) setup utility. Then navigate to the menu items that pertain to device configuration.

Serial ports allow the configuration of specific hardware resources such as input/output (I/O) addresses and hardware interrupt requests (IRQs). Depending on your system firmware, serial port resources might be user configurable, or they might have **Enabled** or **Disabled** as the only choices available. You should configure the appropriate port so it is enabled and at the default settings for a given COM port.

For more detailed procedures and best practices related to configuring serial ports, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site

# Enabling firmware console redirection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling firmware console redirection**

With firmware console redirection, your system can redirect text output to the out-of-band port.

Typically, you must enable firmware console redirection in the system firmware settings. Consult the original equipment manufacturer (OEM) or the documentation that came with your computer for information about how to change the firmware settings.

It is preferable for a system to indicate its firmware redirection settings through the Serial Port Console Redirection (SPCR) table. For more information about this topic, see Enabling the SPCR table.

For more information about firmware console redirection, see Firmware console redirection.

For more information about possible firmware settings and setup procedures, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

# Enabling the SPCR table

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling the SPCR table**

The Serial Port Console Redirection (SPCR) table provides information on how the system firmware and the service processor use the out-of-band management port. For example, if the out-of-band management port is a serial port, the SPCR table contains information such as serial port number, baud rate, terminal type, and other settings used for out-of-band communication. If the out-of-band management port is a non-legacy Universal Asynchronous Receiver/Transmitter (UART) interface on a service processor, the SPCR table provides information about the location and type of UART interface.

The SPCR table ensures that there is a smooth transition between the system firmware and Emergency Management Services.

The SPCR table is usually enabled automatically if you have a service processor that supports Emergency Management Services, or you have enabled firmware console redirection. See your original equipment manufacturer (OEM) documentation to determine whether your computer supports the SPCR table and whether you must adhere to special requirements to enable it.

Windows Server 2003 operating systems attempt to configure Emergency Management Services hardware resource settings to match SPCR table settings without user intervention. If your system does not have an SPCR table, you must configure the Emergency Management Services settings to match the firmware settings, providing the operating system with the Emergency Management Services out-of-band port parameters.

For more information about required or recommended settings, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

# Selecting the terminal definition setting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Selecting the terminal definition setting**

Terminal definition standards ensure that communications between the remotely administered server and the computer that manages it are standardized. For background information on these standards, see VT-UTF8, VT100+, and VT100 conventions.

Be sure to select a terminal definition setting for your firmware or service processor that is compatible with the Windows Server 2003 family and client terminal software. The best setting is VT-UTF8. If VT-UTF8 is not available, the next best choices, given in order of preference, are VT100+ and VT100. For a list of required or recommended settings, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

# Configuring terminal concentrators

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring terminal concentrators**

A terminal concentrator consolidates access to multiple servers. You can use a terminal concentrator to establish remote out-of-band connections through the Telnet protocol over an in-band network connection. For information about selecting a terminal concentrator to complement Emergency Management Services, see Terminal concentrators.

Telnet security features are not standard across terminal concentrators. If your device does not include security features, consider using a secondary private management network that is accessible through a direct-dial remote access or virtual private network (VPN) connection. You can also use a router to secure network traffic to the terminal concentrator.

For more information, see Checklist: Installing and configuring the router and Checklist: Installing and configuring the remote access server.

You can also use Telnet over Internet Protocol security (IPSec) for increased security. See Internet Protocol Security (IPSec) for additional help.

If you choose not to use a terminal concentrator, you can use null modem cables to connect a single networked computer running a Windows Server 2003 operating system to multiple servers running Emergency Management Services.

For more information on selecting and configuring terminal concentrators, see "Emergency Management Services" in *Windows Server TechCenter* at the Microsoft Web site.

# Selecting client terminal software for Emergency Management Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Selecting client terminal software for Emergency Management Services**
Before you can access Emergency Management Services remotely, you must ensure that you are using compatible client software.

The out-of-band management hardware installed on your server determines which client software you should use. HyperTerminal or other VT-UTF8 terminal emulation software might be an option. For more information, see VT-UTF8, VT100+, and VT100 conventions and the HyperTerminal Help.

If your system has a service processor that requires a specialized client, both the service processor and the client must be capable of displaying Emergency Management Services output. For communication with service processors using RJ-45 Ethernet, Transmission Control Protocol/Internet Protocol (TCP/IP) might be an appropriate choice.

**Viewing non-English output from Emergency Management Services**
Emergency Management Services supports localized output on non-English versions of Windows Server 2003 operating systems. To view non-English characters, you must have an out-of-band management application that supports the VT-UTF8 terminal definition. HyperTerminal supports the VT-UTF8 terminal type.

For some terminals that support the VT-UTF8 terminal type, you must configure the operating system in order to be able to view output from a server that is a non-English language that is different from the language on your local computer.

For example, if the server runs a Chinese version of the Windows Server 2003 operating system, but the computer you use to manage the server runs an English version of the operating system, it might be necessary to configure the local computer's operating system. This is typically done in Regional and Language Options. The procedure can vary, however, according to the terminal application you use. It might also be necessary to install additional fonts.

For step-by-step information on how to view non-English output from Emergency Management Services, see View East Asian languages with the VT-UTF8 terminal emulator in Hyperterminal.

# Enabling Emergency Management Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling Emergency Management Services**

This section covers:

- Enabling Emergency Management Services during a new installation
- Enabling Emergency Management Services during an upgrade (x86-based systems)
- Enabling Emergency Management Services after Setup (x86-based systems)
- Enabling Emergency Management Services after Setup (Itanium-based systems)

# Enabling Emergency Management Services during a new installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling Emergency Management Services during a new installation**

You can initiate an Emergency Management Services embedded installation when you start the computer from the Windows Server 2003 operating system CD if your system supports firmware console redirection and the Serial Port Console Redirection (SPCR) table.

Before you begin a CD-based Windows Server 2003 operating system Setup, enable firmware console redirection. You can then remotely view and respond to the prompt **Press any key to boot from CD-ROM**, which appears when you start the system using the operating system CD.

Emergency Management Services configures itself during a bootable CD installation by reading the SPCR table. If Emergency Management Services is enabled, you are prompted at the end of text-mode Setup to allow Setup to automatically configure your system without user input. You must choose this option; otherwise, the next portion of Setup, known as GUI-mode Setup, completes only if you provide input through a local monitor and keyboard. After Setup is complete, you can reconfigure the setting appropriately.

If your firmware does not support the SPCR table, you must fully automate your setup. If your firmware supports the SPCR table, and you want to disable Emergency Management Services, you can configure Setup using Unattend.txt and Winnt.sif files.

**Unattend.txt and Winnt.sif files**

You must have the Unattend.txt and Winnt.sif files if you want to fully automate the remote installation of a Windows Server 2003 operating system.

A sample Unattend.txt file is provided on the operating system CD. You can either use the default settings or customize your installations by modifying or adding parameters. When you edit the Unattend.txt files, insert the parameters in the **[Data]** section, as shown in the following table:

| [Data] Parameter | Possible values |
|---|---|
| **EMSPort=**{**com1**\|**com2**\|**usebiossettings**} | • **com1** or **com2** (where 1 or 2 specifies serial port 1 or 2). This option is valid for x86-based systems only.<br>• **usebiossettings**<br>This is the default value. This parameter instructs the operating system to detect and use SPCR settings. If you use this parameter and an SPCR table is not present, Emergency Management Services is not enabled. |
| | |
| **EMSBaudRate=** *value* | The default value is 9600 baud, with the values of 19200, 57600, and 115200 possible, depending on the capabilities of the serial port. This must be used with **EMSPort=** or the parameter is ignored. |

You can create a Winnt.sif file by copying an Unattend.txt file to a floppy disk and renaming it to Winnt.sif. Because CD media is read-only, CD-based Windows Server 2003 operating system installations must read custom setup information from the floppy drive. The table later in this topic lists unattended Setup parameters required for Winnt.sif. When editing a Winnt.sif file, insert the parameters listed in both tables into the **[Data]** section.

For remote Setup to complete successfully without a local keyboard or monitor, you must configure your setup scripts so they are completely unattended. For more information about unattended Setups, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

| [Data] Parameter | Possible values |
|---|---|

| | |
|---|---|
| **Autopartition**={**0**\|**1**} | Allows unattended Setup to choose a partition in which to install. The value must be set to **1** for unattended Setups. If you do not set the value, text-mode Setup stops and waits for your input. |
| **MsDosInitiated**={**0**\|**1**} | Informs the Setup Loader that an unattended installation is occurring directly from the CD. The value must be set to **0** for unattended Setups. |
| **UnattendedInstall**={**No**\|**Yes**} | The value must always be set to **Yes** for CD-based unattended installations. |

**Important**

- To use Winnt.sif files, you must set the CD-ROM as the primary boot device and leave the floppy disk containing a Winnt.sif file in the floppy drive.
- Unattended Setups for Windows Server 2003 operating systems do not pause for user input. With console redirection, you can view text-mode Setup remotely for monitoring purposes.

For specific examples and additional information about creating unattended setups, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

**Remote Installation Services (RIS)**

With Remote Installation Services (RIS), you can install a Windows Server 2003 operating system on computers without being physically present at those computers. If the firmware on the target computers supports Pre-Boot eXecution Environment (PXE), console redirection, and the Serial Port Console Redirection (SPCR) table, you can use Emergency Management Services to initiate and complete Setup from remote setup media.

If your firmware does not support the SPCR table, you can configure your Emergency Management Services settings in the RISTNDRD.txt file just as you would configure them in the normal setup scripts, Unattend.txt and Winnt.sif. If your firmware does not support console redirection, you can use special versions of Startrom.com that output to a generic serial port, com1 or com2, at 9600 baud.

For more information about remote installation, see RIS and PXE requirements.

# Enabling Emergency Management Services during an upgrade (x86-based systems)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling Emergency Management Services during an upgrade (x86-based systems)**
During an upgrade, you can enable Setup on an x86-based system by using Winnt32.exe. After inserting the Windows Server 2003 operating system CD into your computer, map a network drive connection to the path containing the installation files. At the command prompt type:

*Drive* **\i386\winnt32.exe**

*Drive* represents the drive letter assigned to the CD-ROM drive or shared folder. Follow this with one or any combination of the command-line parameters listed in the following table:

| [Data] Parameter | Possible values |
|---|---|
| **EMSPort=**{**com1**\|**com2**\|**usebiossettings**} | • **com1** or **com2** (where the numeral 1 or 2 specifies serial port 1 or 2).<br><br>This option is valid for x86-based systems only.<br><br>• **usebiossettings**<br><br>This is the default value. This parameter instructs the operating system to detect and use Serial Port Console Redirection (SPCR) settings. If you use this parameter and an SPCR table is not present, Emergency Management Services is not enabled. |
| **EMSBaudRate=** *value* | The default value is 9600 baud, with other values of 19200, 57600, and 115200 possible, depending on the capabilities of the serial port. This must be used with **EMSPort=** or the parameter is ignored. |

Winnt32.exe copies setup files to your system and then prompts you to confirm a restart. After restarting the system, text-mode Setup begins, followed by GUI-mode Setup. During text-mode Setup, you can view setup messages through Emergency Management Services console redirection. During GUI-mode Setup, you can access Special Administration Console (SAC), but you cannot view Setup prompts.

**Important**

- Winnt32.exe is a GUI-based tool and does not support Emergency Management Services; you cannot start or monitor Winnt32.exe by using the out-of-band port. For this purpose, you must use a GUI-based tool, such as Terminal Services Remote Desktop for Administration.

For more information about using Winnt32.exe to enable Emergency Management Services, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

# Enabling Emergency Management Services after Setup (x86-based systems)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling Emergency Management Services after setup (x86-based systems)**
To enable Emergency Management Services after setting up a Windows Server 2003 operating system, you must edit the Boot.ini file to enable Windows loader console redirection and Special Administration Console (SAC). The Boot.ini file controls startup; it is located on the system partition root.

Use Bootcfg (Bootcfg.exe) to edit the Boot.ini file.

To view examples of how to edit the Boot.ini file, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

You can enable Emergency Management Services console redirection in the Recovery Console. You do this by editing the Winnt.sif file, located in the Cmdcons folder on the system partition root. Note that the Winnt.sif for the Recovery Console is separate from the floppy disk Winnt.sif file that is referenced in Enabling Emergency Management Services during a new installation.

For specific information about setting Recovery Console Bootcfg parameters for Emergency Management Services, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

# Enabling Emergency Management Services after Setup (Itanium-based systems)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling Emergency Management Services after setup (Itanium-based systems)**

To enable Emergency Management Services on Itanium-based installations of a Windows Server 2003 operating system, you must first enable firmware console redirection.

To do so, follow the instructions in your hardware manufacturer's documentation. Then, edit the Extensible Firmware Interface (EFI) boot manager settings using Bootcfg. The **Bootcfg** command enables Emergency Management Services console redirection and Special Administration Console (SAC).

For more information about using Bootcfg to enable Emergency Management Services, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

# Using Emergency Management Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Emergency Management Services**

This section covers:

- Special Administration Console (SAC) and SAC commands
- !Special Administration Console (!SAC) and !SAC commands
- VT-UTF8, VT100+, and VT100 conventions
- Task list headings

# Special Administration Console (SAC) and SAC commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Special Administration Console (SAC) and SAC commands**
You perform most remote-management and recovery tasks using an out-of-band management port with Special Administration Console (SAC). The server must be at least partially functioning in order to use SAC.

After Emergency Management Services is enabled, SAC is always running unless Emergency Management Services components are prevented from loading properly. If Emergency Management Services components are prevented from loading properly, the server starts !Special Administration Console (!SAC). In addition, a set of SAC commands is available for performing management tasks.

The following prompt appears when SAC is active:

**SAC>**

For more information, see Emergency Management Services components and !Special Administration Console (!SAC) and !SAC commands.

**SAC commands**
The following table lists SAC commands that are available for remote management tasks.

| Command | Description |
|---|---|
| **ch** | Lists all channels. |
| **cmd** | Creates Windows command-prompt channels. To use a command-prompt channel, you must provide valid logon credentials. You must log on to each command-prompt instance. |
| **crashdump** | Manually generates a Stop error message and forces a memory dump file to be created. |
| **D** | Dumps the current kernel log. |
| **F** | Toggles the information output by the t-list command, which shows processes only, or shows processes and threads. |
| **I** | If no parameters are passed, this command lists Internet Protocol (IP) information. You can configure IP parameters to display or set the IP address, subnet mask, and gateway of a given network interface device by providing the network number, IP address, and subnet information. To do so, use the following format: *<network#> <IPaddress> <subnet>* |
| **Id** | Displays identification information about the server. |
| **K***<PID>* | Ends the given process. *PID* is the process identification number you specify. |
| **L***<PID>* | Lowers the priority of a process (and any associated child processes) to the lowest possible level. |
| **lock** | Restricts access to Emergency Management Services command-prompt channels. You must provide valid logon credentials to unlock a channel. |
| **M***<PID> <MB-***allow***>* | Limits the memory usage of a process (and any associated child processes) to a specified number of megabytes. *MB* is the number of megabytes you specify. |
| **P** | Causes t-list command output to pause after displaying one full screen of information. |
| **R***<PID>* | Raises the priority of a process and any associated child processes up by one level. |
| **restart** | Restarts the computer. |

| | |
|---|---|
| **S** | If no parameters are passed, this command displays the current date using the 24-hour clock format. You can set the system time by providing the date and, optionally, the time in this format: *mm***/***dd***/***yyyyhh***:***mm*. |
| **shutdown** | Shuts down the computer. Do not use this command unless you can be physically present at the computer when you are ready to restart it. |
| **T** | Lists the processes and threads that are currently running. |
| **? or help** | Lists the available commands. |

**Channel-management commands**

The following table lists the subset of SAC commands available for managing out-of-band port channels.

| Command | Description |
|---|---|
| **ch** | Lists all channels. |
| **ch -si** <*n*> | Changes to the channel with the number, *n*, which you specify. |
| **ch -sn** <*name*> | Changes to the channel with the *name* you specify. |
| **ch -ci** <*n*> | Closes the channel with the number, *n*, which you specify. |
| **ch -cn** <*name*> | Closes the channel with the *name* you specify. |
| <**Esc**><**Tab**> | Enables you to change channels. |
| <**Esc**><**Tab**>**0** | Returns you to the SAC channel. |
| **ch -?** | Displays the channel-management commands Help. |

# !Special Administration Console (!SAC) and !SAC commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**!Special Administration Console (!SAC) and !SAC commands**
You can use !Special Administration Console (!SAC) to perform remote management tasks if the Windows Server 2003 operating system stops functioning due to a Stop error message.

The main functions of !SAC are as follows:

- Redirect Stop error message explanatory text.
- Restart the system.
- Obtain computer identification information.

When !SAC becomes available, the following prompt appears:

**!SAC>**

!SAC includes a subset of SAC commands, which are listed in the following table:

| Command | Description |
|---|---|
| **restart** | Restarts the server immediately. |
| **D** | Displays all log entries (screen pauses at each page of information). |
| **Id** | Displays computer identification. |
| **Q** | Quits !SAC and resumes normal out-of-band port operation. |
| **? or Help** | Lists available commands. |

For additional information about topics related to !SAC, see Emergency Management Services components and Special Administration Console (SAC) and SAC commands.

# VT-UTF8, VT100+, and VT100 conventions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VT-UTF8, VT100+, and VT100 conventions**

Industry terminal definition conventions provide a smooth way for you to send commands from the remote management console--usually Special Administration Console (SAC)-- to the server. These terminal definition conventions ensure a consistent mechanism for managing systems in all states of operation and failure. Windows Server 2003 operating systems use the commands, which are described later in this topic. The firmware, service processors and client terminal software should also use these commands.

The VT100 Terminal Definition is the standard convention used to configure and conduct emergency management tasks with UNIX-based servers. VT100 does not support all keys on the standard PC 101-key layout, however. The VT-UTF8 convention makes available additional keys that are not provided by VT100 for the PC 101 keyboard. With VT-UTF8, you can also use colors and some standardized escape sequences that are commonly used for management tasks. With VT-UTF8, output sent to the serial port can be localized while UNIX interoperability for English is sustained.

For more information about viewing non-English output from Emergency Management Services, see Selecting client terminal software for Emergency Management Services.

VT100+ is a subset of VT-UTF8 that handles only VT100-compatible characters--that is, American National Standards Institute (ANSI) characters 0 through 127. VT-UTF8 is the preferred terminal type for viewing Emergency Management Services output using a terminal emulator, such as HyperTerminal, and for firmware or service processor output. In either case, if VT-UTF8 is not available, your next best choices are to try VT100+ and then VT100, in that order.

**VT-UTF8 and VT100+ keyboard-to-escape-sequence mappings**

If you plan to use a VT100, VT100+, or VT-UTF8 terminal emulator to communicate with Emergency Management Services or some other management console, the key mappings provided later in this topic should occur automatically when you press the associated key.

For example, to send an **F2** keystroke to a remotely administered server from a terminal using the VT100+ or VT-UTF8 terminal definition, press **F2**. If you use VT100 emulation software or a terminal emulator that does not support all the key mappings or management escape sequences, however, you can type the escape sequences directly into the terminal. For example, to send an **F2** from a VT100 terminal, type **<ESC>**, and then type **2**.

The following table defines these VT100+ and VT-UTF8 key mappings.

| VT100+ key | Escape sequence sent |
|---|---|
| **HOME** | <ESC> h |
| **END** | <ESC> k |
| **INSERT** | <ESC> + |
| **DELETE** | <ESC> - |
| **PAGE UP** | <ESC> ? |
| **PAGE DOWN** | <ESC> / |
| **ALT** | <ESC>^A |
| **CTRL** | <ESC>^C |
| **F1** | <ESC> 1 |
| **F2** | <ESC> 2 |
| **F3** | <ESC> 3 |
| **F4** | <ESC> 4 |

| | |
|---|---|
| **F5** | <ESC> 5 |
| **F6** | <ESC> 6 |
| **F7** | <ESC> 7 |
| **F8** | <ESC> 8 |
| **F9** | <ESC> 9 |
| **F10** | <ESC> 0 |
| **F11** | <ESC> ! |
| **F12** | <ESC> @ |

**Note**

- PC function keys are not equivalent to the VT100 terminal function keys.

The following table lists VT-UTF8 escape sequences used to send special commands:

| Action | Escape sequence |
|---|---|
| **Start firmware setup utility** | ^F |
| **Reset**<br><br>The computer will reset without shutting down Windows. This command requires either service processor support, or a power switch or uninterruptible power supply (UPS) with remote computer management capabilities. | <ESC>R<ESC>r<ESC>R |
| **Invoke service processor**<br><br>The service processor monitors the serial traffic for this escape sequence. When it detects the sequence, the service processor assumes control of the out-of-band port and immediately sends an acknowledgement sequence (defined later in this table). You cannot use Emergency Management Services until you exit the service processor. This command requires service processor support. | <ESC>( |
| **Invoke power switch or UPS**<br><br>The power switch or UPS monitors the serial traffic for this escape sequence. When it detects the sequence, the power switch assumes control of the out-of-band port and immediately sends an acknowledgement sequence (defined later in this table). You cannot use Emergency Management Services until you exit the power switch. This command requires a power switch or UPS with remote computer management capabilities. | <ESC>) |
| **Power switch, UPS, or service processor acknowledgement**<br><br>After taking control of the out-of-band port, the power switch, UPS, or service processor sends this escape sequence. This indicates that the power switch, UPS, or service processor has assumed control of the port and is ready for interaction. This command requires either service processor support, or a power switch or UPS with remote computer management capabilities. | <ESC>* |
| **Exit power switch, UPS, or service processor; display a user interface (UI)**<br><br>The power switch, UPS, or service processor releases control of the out-of-band port. Before it does so, it might display a user interface message. For example, it might display a warning message such as: "Do you want to save setting and exit? [y/n]." This command requires either service processor support, or a power switch or UPS with remote computer management capabilities. | Q |

| | |
|---|---|
| **Exit power switch, UPS, or service processor; do not display a UI**<br><br>The power switch, UPS, or service processor releases control of the out-of-band port without displaying a user interface (UI). | <ESC>Q |

For more information, see Selecting the terminal definition setting.

# Task list headings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Task list headings**

The following table describes the output from the Special Administration Console (SAC) **T** command, which lists the processes and threads currently running.

| Task list heading | Description |
|---|---|
| User Time | The amount of user-mode time charged to a process since the last time the computer was restarted. |
| Kernel Time | The amount of kernel-mode time charged to a process since the last time the computer was restarted. |
| ws | Working set, or the amount of physical memory assigned to a process by the operating system. |
| Faults | The page faults that have occurred since the last time the computer was restarted. A page fault is an interrupt that occurs when software attempts to read from or write to a virtual memory location that is marked *not present*. |
| Commit | The amount of memory that a process is using. |
| Pri | The default thread priority. |
| Hnd | The number of handles charged to a process. |
| Thd | The thread count. |
| Pid | The process identifier (PID) or numerical identifier that uniquely distinguishes a process while it runs. |
| Name | The name of the process. |

For more information on using the **T** command, see Resolve system resource problems using SAC.

# Emergency Management Services Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

For information on how to integrate Emergency Management Services into your existing environment, see "Deploying Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

For more detailed information on how to select and configure hardware and enable Emergency Management Services, see "Emergency Management Services" at the Microsoft Windows Resource Kits Web site.

For information about compatible hardware, product support services, and tools, see Support resources.

# Troubleshooting Emergency Management Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting Emergency Management Services**

**Important**

- You can use the approaches provided in this topic to troubleshoot Emergency Management Services components, not to troubleshoot your system. For information about using Emergency Management Services to troubleshoot your system, see Emergency Management Services How To....

What problem are you having?

- The out-of-band port is briefly unavailable; no status information appears.
- Status information appears, but Emergency Management Services does not respond to terminal input.
- The kernel debugger disables !Special Administration Console (!SAC). Garbled characters might appear.

**The out-of-band port is briefly unavailable; no status information appears.**

**Cause:** The Windows Server 2003 family components that provide Emergency Management Services load and unload as part of the startup process, which causes components to become briefly disabled. This is called a "black-out period."

**Solution:** If Special Administration Console (SAC) or !Special Administration Console (!SAC) do not respond, wait for operating system components to initialize before restarting or stopping the system using whatever means are available, such as a service processor or uninterruptible power supply (UPS). If a service processor is present, consult your documentation to determine other diagnostic tools that it provides.

**See also:** Using Emergency Management Services.

**Status information appears, but Emergency Management Services does not respond to terminal input.**

**Cause:** The Windows Server 2003 family components that provide Emergency Management Services load and unload as part of the startup process, which causes components to provide limited functionality. This is called a "gray-out period."

**Solution:** Gray-out periods might become interactive if the component provides !Special Administration Console (!SAC) functionality. If !SAC does not respond, however, wait for operating system components to initialize before restarting or stopping the system using other means, such as a service processor or uninterruptible power supply (UPS). If a service processor is present, consult your documentation to determine other diagnostic tools that it provides.

**See also:** !Special Administration Console (!SAC) and !SAC commands.

**The kernel debugger disables !Special Administration Console (!SAC). Garbled characters might appear.**

**Cause:** The kernel debugger and !SAC cannot use the same serial port simultaneously. In addition, when a kernel debugger is running, !SAC becomes unavailable because the debugger takes control of the system if a Stop error occurs.

**Solution:** Use different serial ports for the kernel debugger and !SAC. It is important to specify the correct COMx parameters in the Boot.ini file. You should also disable kernel debuggers on servers during normal operations.

# Monitoring and Status Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring and status tools**
This section covers:

- Task Manager overview
- Using Network Diagnostics
- System Information
- System Properties
- Services Snap-in
- Event Viewer
- Network Monitor
- Monitoring Performance
- Shutdown Event Tracker
- SNMP
- Device Manager
- Windows Management Instrumentation Control

# Task Manager overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Task Manager overview**
Task Manager provides information about programs and processes running on your computer. It also displays the most commonly used performance measures for processes.

You can use Task Manager to monitor key indicators of your computer's performance. You can see the status of the programs that are running and end programs that have stopped responding. You can also assess the activity of running processes using as many as fifteen parameters, and see graphs and data on CPU and memory usage.

In addition, if you are connected to a network, you can view network status and see how your network is functioning.

If you have more than one user connected to your computer, you can see who is connected, what they are working on, and you can send them a message.

**Programs that are running**
The **Applications** tab shows the status of the programs running on your computer.

On this tab, you can end, switch to, or start a program.

**Processes that are running**
The Processes tab shows information about the processes running on your computer.

For example, you can display information on CPU and memory usage, page faults, handle count, and a number of other parameters.

**Performance measures**
The **Performance** tab displays a dynamic overview of your computer's performance, including:

- Graphs for CPU and memory usage.
- Totals for the number of handles, threads, and processes running on your computer.
- Totals, in kilobytes, for physical, kernel, and commit memory.

**Viewing Network performance**
The **Networking** tab displays a graphical representation of network performance. It provides a simple, qualitative indicator that shows the status of the network(s) that are running on your computer. The Networking tab is displayed only if a network card is present.

On this tab, you can view the quality and availability of your network connection, whether you are connected to one or more than one network.

**Monitoring Sessions**
The **Users** tab displays users who can access this computer, and session status and names. **Client Name** specifies the name of the client computer using the session, if applicable. **Session** provides a name for you to use to perform such tasks as sending another user a message or connecting to another user’s session.

The **Users** tab is displayed only if the computer you are working on has Fast User Switching enabled, and is a member of a workgroup or is a standalone computer. The **Users** tab is unavailable on computers that are members of a network domain.

# Using Network Diagnostics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Network Diagnostics**

With Network Diagnostics, you can display information about your network environment quickly and easily. This tool provides information about your computer, operating system, network, and network adapters. Using Network Diagnostics, you can perform standard tests for network connectivity, such as running the ping command to test the network path to a server or trying to connect to a server. Network Diagnostics reports both successful and unsuccessful attempts to reach network resources.

To open Network Diagnostics, click **Start**, click **Help and Support**, click **Tools**, click **Help and Support Center Tools**, and then click **Network Diagnostics**.

**Note**

- Network tests are run only on your computer to ensure the privacy and integrity of your data. No information is sent to Microsoft.

# System Information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Information**
System Information collects and displays information about the configuration of your system. You can view information for a local computer or for a computer to which you are connected.

- For help with specific tasks, see System Information How To....
- For general background information, see System Information Concepts.

# System Information How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Find system data
- View system history
- Export system data to a text file
- Print system data

# Find system data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find system data**

1. Open System Information.
2. In the **Find what** box at the bottom of the **System Information** window, type the word or words that correspond to the system information for which you are searching.
3. Select the appropriate search option:
   - To search only a portion of the categories in the console tree, select the **Search selected category only** check box. This starts the search at the top of the currently selected category and searches all of its subcategories. To start the search at the root, clear this check box.
   - To search only the categories in the console tree, ignoring any matches in the details pane, select the **Search category names only** check box. To search both the console tree and the details pane, clear this check box.
   - To search all categories in both panes, clear both check boxes.
4. Click **Find**.
5. To search for the item again, click **Find Next**.

## Notes

- To access System Information on a remote computer, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open System Information, click **Start**, and then click **Help and Support**. Under **Support Tasks**, click **Tools**, and then click **Help and Support Center Tools**. Click **Advanced System Information**, and then click **View detailed system information (Msinfo32.exe)**.
- If the **Find** options do not appear at the bottom of the System Information window, on the **Edit** menu, click **Hide Find** to clear the check mark and restore the options.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Information
Print system data
Export system data to a text file

# View system history

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view system history**

1. Open System Information.
2. On the **View** menu, click **System History**.

## Notes

- To access System Information on a remote computer, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open System Information, click **Start**, and then click **Help and Support**. Under **Support Tasks**, click **Tools**, and then click **Help and Support Center Tools**. Click **Advanced System Information**, and then click **View detailed system information (Msinfo32.exe)**.
- System history is collected when changes are made to certain system devices and settings. This system history data is compared periodically to the previous system settings, and if there are any changes, those changes are saved in the System History view. Only the most recent changes are compared against previous settings; therefore, it is possible to make multiple changes to a certain device or driver and have only one of the changes appear in the System History view.
- History data is not collected for all System Information categories.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Find system data
Export system data to a text file
Print system data
System Information

# Export system data to a text file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export system data to a text file**

1. Open System Information.
2. On the **File** menu, click **Export**.
3. In **Save in**, click a file location.
4. In **File name**, type the file name.
5. Click **Save**.

### Important

- The information that you export from System Information may contain sensitive system data. It is important to safeguard this information and prevent unauthorized access to it.

### Notes

- To access System Information on a remote computer, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open System Information, click **Start**, and then click **Help and Support**. Under **Support Tasks**, click **Tools**, and then click **Help and Support Center Tools**. Click **Advanced System Information**, and then click **View detailed system information (Msinfo32.exe)**.
- System information is automatically refreshed before a file is saved.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Information
Find system data
Print system data

# Print system data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print system data**

1. Open System Information.
2. On the **File** menu, click **Print**.
3. In the **Print** dialog box, under **Page Range**, select one of the following options, and then click **Print**:
   - To print all of the system data, click **All**.
   - To print the currently selected category and its subcategories, click **Selection**.

## Notes

- To access System Information on a remote computer, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open System Information, click **Start**, and then click **Help and Support**. Under **Support Tasks**, click **Tools**, and then click **Help and Support Center Tools**. Click **Advanced System Information**, and then click **View detailed system information (Msinfo32.exe)**.
- System information is automatically refreshed before each print job.
- Print output from System Information, especially when you are exporting multiple categories to a print job, can be very large. You can use the **msinfo32** command to limit the categories for which you export data. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Information
Find system data
Export system data to a text file
Msinfo32

# System Information Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about System Information:

- System Information overview
- Managing System Information from the command line
- Understanding System Information

# System Information overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Information overview**

System Information collects and displays system configuration information for local and remote computers. This includes information about hardware configurations, computer components, and software, including signed drivers and unsigned drivers. You can use System Information to quickly find the information that you need to resolve a system problem.

System Information saves data files in a native (.nfo) format. By default, system data is automatically saved in an .nfo file format that is compatible with the current version of System Information. You can also save data as a Version 5.0/6.0 System Information (.nfo) file, which is compatible with older versions of System Information.

You can use System Information to view system history. System history is collected when changes are made to certain system devices and settings. This system history data is compared periodically to the previous system settings, and if there are any changes, those changes are saved in the System History view. Only the most recent changes are compared against previous settings; therefore, it is possible to make multiple changes to a certain device or driver and have only one of the changes appear in the System History view. For more information, see View system history.

You can also open .cab files and .xml files in System Information. When you open .cab files, you can view the contents of those files by using the **Tools** menu.

# Managing System Information from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing System Information from the command line**
In addition to the System Information support tool in Help and Support Center, you can use the **msinfo32** command to view information about a computer that you are managing. With this command, you can:

- Use System Information from a batch file.
- Create .nfo files or .txt files that contain information from specified categories.
- Open System Information, displaying only specified categories.
- Save a System Information file without opening the System Information tool.
- Start System Information for a remote computer. (You must have the appropriate permissions.)
- Create a shortcut that opens System Information in a commonly used configuration.

For more information about the **msinfo32** command, see Msinfo32. For information about other command-line utilities, see Command-line reference A-Z. For information about manageability, see Management Strategies and Tools.

# Understanding System Information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding System Information**
System Information is a support tool that you can use to quickly locate information about your computer and its operating system. The category tree in the left pane contains items in a folder view that is similar to Windows Explorer. The details pane on the right side of the window displays information about the items that you select in the category tree.

This section covers:

- System Summary
- Hardware Resources
- Components
- Software Environment
- Internet Settings

# System Summary

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Summary**
System Summary is the first category in the System Information category tree. The details pane displays general information about your computer and its operating system. You can view information about the operating system name, version, manufacturer, and directory location. You can verify the version of the BIOS or EFI, the processor type, and memory information.

You can use System Summary to:

- Check the BIOS or EFI version and date.
- Determine the directory in which the operating system is installed.
- Verify that new memory is installed correctly by checking Total Physical Memory and Available Physical Memory.
- Check the Page File Space value if the computer is having memory problems. Page File Space is the physical space on the hard disk that the operating system uses to increase the apparent size of RAM.
- Check your activation status. This information appears only if you have not activated the product. For more information, see Windows Product Activation.

For more information about updating your system BIOS or EFI, or about installing new memory, refer to your hardware manufacturer's documentation.

# Hardware Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware Resources**
The Hardware Resources category in System Information displays information about resource assignments and possible sharing conflicts among DMA, Forced Hardware, I/O, IRQs, and Memory resources.

For more information about these items, see the following:

- Conflicts/Sharing
- DMA
- Forced Hardware
- I/O
- IRQs
- Memory

# Conflicts/Sharing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Conflicts/Sharing**

Conflicts/Sharing is located in the Hardware Resources category in System Information. It displays devices that share the same resources, such as direct memory access (DMA), IRQs, and memory.

In general, shared resources are not problematic, but investigating potential conflicts can help resolve issues. If you are having a problem with a device, check Conflicts/Sharing to verify that there are no sharing conflicts. Locate the device in the **Device** column in the details pane, and then verify which resource the device uses in the **Resource** column. Devices that share the same resource have duplicate entries in the **Resource** column. For example, you may have two adapters that share the same IRQ channel, such as IRQ 5. If the problem device shares a resource with another device, there may be a conflict. You can use Device Manager to troubleshoot the problem. For more information about Device Manager, see Device Manager.

For more information about Hardware Resources, see Hardware Resources.

# DMA

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DMA**

DMA is located in the Hardware Resources category in System Information. It displays information about your system's direct memory access (DMA). DMA transfers data between system memory and hardware devices without passing it through the CPU. The **Resource** column in the details pane displays the DMA channel that is being used by the device, which is listed in the **Device** column. The **Status** column displays the status of the device.

If the **Status** column indicates an error, such as Error, Degraded, Unknown, Pred Fail, Starting, Stopping, or Service, you can try the following troubleshooting solutions:

- Use Device Manager to locate and troubleshoot the problem. For more information about Device Manager, see Device Manager.
- Contact the device manufacturer to obtain an updated device driver.
- Reconfigure the settings in the system BIOS or EFI.

For more information about Hardware Resources, see Hardware Resources.

# Forced Hardware

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Forced Hardware**

Forced Hardware is located in the Hardware Resources category in System Information. The **Device** column in the details pane identifies the type of device. The **PNP Device ID** column lists the device ID.

Forced hardware is any device that you have to configure manually or that has user-specified resources, as opposed to system-specified resources. The category Forced Hardware also applies to devices that are not Plug and Play compatible, such as legacy Industry Standard Architecture (ISA) devices.

If there is a conflict with a manually configured hardware device, use Device Manager to locate and troubleshoot the problem. For more information about Device Manager, see Device Manager.

For more information about Hardware Resources, see Hardware Resources.

# I/O

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**I/O**

I/O is located in the Hardware Resources category in System Information. I/O is the communication channel among hardware devices in your computer. The **Resource** column in the details pane displays the resource that is used by an I/O device, which appears in the **Device** column. The **Status** column displays the status of the device.

If there is a sharing conflict with a device that uses an I/O channel, locate the device in Conflicts/Sharing or in I/O, and then use Device Manager to locate and troubleshoot the problem. For more information about Device Manager, see Device Manager.

For more information about Hardware Resources, see Hardware Resources.

# IRQs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IRQs**

IRQs is located in the Hardware Resources category in System Information. It displays information about interrupt request (IRQ) channels on your system. The **Resource** column and the **Device** column indicate which devices are assigned to each channel, and the **Status** column displays the status of the device. The operating system allows you to use only certain IRQs. Other IRQs are reserved for a standard set of devices.

If there is a sharing conflict with an IRQ, such as a mouse not responding, use Device Manager to locate and troubleshoot the problem. For more information about Device Manager, see Device Manager.

For more information about Hardware Resources, see Hardware Resources.

# Memory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Memory**
Memory is located in the Hardware Resources category in System Information. The **Resource** column in the details pane displays the memory address of the device that is listed in the **Device** column.

Memory address ranges are used for communication between devices and the operating system. Each device requires its own range of memory. If you are experiencing problems, and two devices are configured to use the same memory address range, you can try the following troubleshooting solutions:

- Use Device Manager to locate and troubleshoot the problem. For more information about Device Manager, see Device Manager.
- Contact the device manufacturer to obtain an updated device driver.
- Adjust the memory address ranges for COM ports in the system BIOS or EFI. Contact your hardware manufacturer for more information.

For more information about Hardware Resources, see Hardware Resources.

# Components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Components**
The Components category in System Information contains information about the various components and devices on your system. Most of the subcategories contain an **Item** column, which identifies the component, and a **Value** column, which provides additional information about the component, such as drive location or version number.

For more information about these items, see the following:

- Multimedia
- CD-ROM
- Sound Device
- Display
- Infrared
- Input
- Modem
- Network
- Ports
- Storage
- Printing component
- Problem Devices
- USB

# Multimedia

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multimedia**

Multimedia is located in the Components category in System Information. It contains two subcategories: Audio Codecs and Video Codecs. These subcategories contain information about audio and video codecs.

If you experience problems when playing back multimedia files, or if there are errors in multimedia applications, such as Windows Media Player, you can check this Multimedia category to verify their settings.

For more information about Components, see Components.

# CD-ROM

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CD-ROM**

CD-ROM is located in the Components category in System Information. It contains information about the CD-ROM drive, including the drive letter assigned to it and its description, manufacturer, and status.

If you are experiencing problems when you play a CD, try cleaning the CD. If cleaning the CD does not fix the problem, you can use Device Manager to troubleshoot the problem. For more information about Device Manager, see Device Manager.

Occasionally, the drive letter that is assigned to the CD-ROM device is reassigned. You can use Disk Manager to verify the drive letter assignments. For more information about disk management, see Disk Management.

For more information about Components, see Components.

# Sound Device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sound Device**

Sound Device is located in the Components category in System Information. It contains information about the sound devices on your system, including the name, manufacturer, and status of the sound device and driver details.

When you are troubleshooting problems with a sound device, verify that the hardware is properly connected and that the sound control is set correctly. You can also use Device Manager to verify that the sound device driver is signed. For more information about Device Manager, see Device Manager.

For more information about Components, see Components.

# Display

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Display**

Display is located in the Components category in System Information. It lists information about the display devices on your system. You can use this information to determine the manufacturer of a device and then contact the manufacturer for device updates. You can also verify the display resolution, which helps troubleshoot display problems.

To ensure that your hardware meets the minimum compatibility requirements, see Hardware compatibility information in the Windows Catalog (search page) at the Microsoft Web Site. You can also use Device Manager to verify that your system is using signed drivers. For more information about Device Manager, see Device Manager.

For more information about Components, see Components.

# Infrared

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Infrared**

Infrared is located in the Components category in System Information. It contains information about infrared (IR) devices and ports on your system. These devices might not exist on all computers.

For more information about Components, see Components.

# Input

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Input**

Input is located in the Components category in System Information. It contains two subcategories: Keyboard and Pointing Device. These subcategories contain information about your keyboard, pointing device (mouse), and other input devices.

You can troubleshoot keyboard problems by double-clicking **Keyboard** in Control Panel. You can troubleshoot pointing device problems by double-clicking **Mouse** in Control Panel.

For more information about Components, see Components.

**Important**

- If you are troubleshooting an input device that is not a universal serial bus (USB) device, you must shut down the computer before you unplug or reconfigure the device. Refer to your hardware manufacturer's documentation for more information.

# Modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modem**

Modem is located in the Components category in System Information. It displays information about any telephone modems that are installed on your system.

For more information about modems, see Installing and configuring Modems. For more information about Components, see Components.

# Network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network**

Network is located in the Components category in System Information. It contains three subcategories: Adapter, Protocol, and WinSock. These subcategories display information about the network adapters, protocols, and WinSock settings on your system. You can use this information to verify network configurations.

You can also use the Network Diagnostics tool to access information about your network settings. For more information about this tool, see Windows interface administrative tool reference A-Z: Network Diagnostics.

For more information about Components, see Components.

# Ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ports**

Ports is located in the Components category in System Information. It contains two subcategories: Serial and Parallel. For both serial ports and parallel ports, the details pane displays information such as name, status, I/O port (the communication channel among hardware devices that are installed on your computer), IRQ channel, and driver.

For more information about Components, see Components.

# Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Storage**

Storage is located in the Components category in System Information. It contains four subcategories: Drives, Disks, SCSI, and IDE. Storage displays information about the drives, disks, SCSI controllers, and IDE controllers on your system.

Drives displays information for all the drives on your system, including information about the drive letter, file system, size, free space, volume name, and volume serial number. Disks displays information about your disk drive, including the manufacturer, model, size, number of partitions, and partition size. SCSI and IDE display information about each of the device controllers, including the name, manufacturer, status, IRQ channel, and driver.

For more information about Components, see Components.

# Printing component

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing**

Printing is located in the Components category in System Information. It displays information about printers that are installed on your system. The **Name** column in the details pane displays the name of the printer. The **Port Name** column displays the port to which the printer is attached. The **Server Name** column displays the name of the print server, if the printer is on a network.

If you have a printer problem, such as a lost printer connection, use Printers and Faxes to locate and troubleshoot the problem. For more information about Printers and Faxes, see Printing. For more information about the print jobs that are currently printing on your system, see Print Jobs.

For more information about Components, see Components.

# Problem Devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Problem Devices**

Problem Devices is located in the Components category in System Information. It displays information, including the device ID and the error code, about devices that may have a problem. For more information about Components, see Components.

You can use Device Manager to locate and troubleshoot any device problems. For more information about Device Manager, see Device Manager.

The following table contains a list of error codes that may appear in the Problem Devices category, as well as descriptions of the errors. Only error codes that are valid for Windows 2000 and the Windows Server 2003 family of operating systems are included. For more information about troubleshooting these errors, see Article Q245386, "Windows 2000 Configuration Manager Error Code Meanings and Troubleshooting," in the Microsoft Knowledge Base.

| Error Code | Description |
|---|---|
| Code 1 | This device is not configured correctly. |
| Code 2 | Windows cannot load the driver for this device. |
| Code 3 | The driver for this device might be corrupted, or your system may be running low on memory or other resources. |
| Code 4 | This device is not working properly. One of its drivers or your registry might be corrupted. |
| Code 5 | The driver for this device needs a resource that Windows cannot manage. |
| Code 6 | The boot configuration for this device conflicts with other devices. Cannot filter. |
| Code 7 | Cannot filter. |
| Code 8 | The driver loader for the device is missing. |
| Code 9 | This device is not working properly because the controlling firmware is reporting the resources for the device incorrectly. |
| Code 10 | This device cannot start. |
| Code 11 | This device failed. |
| Code 12 | This device cannot find enough free resources that it can use. |
| Code 13 | Windows cannot verify this device's resources. |
| Code 14 | This device cannot work properly until you restart your computer. |
| Code 15 | This device is not working properly because there is probably a renumeration problem. |
| Code 16 | Windows cannot identify all the resources this device uses. |
| Code 17 | This device is asking for an unknown resource type. |
| Code 18 | Reinstall the drivers for this device. |
| Code 19 | Your registry might be corrupted. |
| Code 21 | Windows is removing this device. |
| Code 22 | This device is disabled. |

| | |
|---|---|
| Code 24 | This device is not present, is not working properly, or does not have all its drivers installed. |
| Code 25 | Windows is still setting up this device. (Same text as Code 26.) |
| Code 26 | Windows is still setting up this device. (Same text as Code 25.) |
| Code 27 | This device does not have valid log configuration. |
| Code 28 | The drivers for this device are not installed. |
| Code 29 | This device is disabled because the firmware of the device did not give it the required resources. |
| Code 31 | This device is not working properly because Windows cannot load the drivers required for this device. |
| Code 32 | A driver for this device was not required and has been disabled. |
| Code 33 | Windows cannot determine which resources are required for this device. |
| Code 34 | Windows cannot determine the settings for this device. |
| Code 35 | Your computer's system firmware does not include enough information to properly configure and use this device. |
| Code 36 | This device is requesting a PCI interrupt but is configured for an ISA interrupt (or vice versa). |

# USB

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**USB**

USB is located in the Components category in System Information. It displays information about universal serial bus (USB) devices and controllers on your computer. The **Device** column in the details pane lists each USB device, and the **PNP Device ID** column lists the ID for the device.

If there is a problem with a USB device, first try unplugging and plugging in the device. If that does not fix the problem, you can use Device Manager to locate and troubleshoot the problem. For more information about Device Manager, see Device Manager.

For more information about Components, see Components.

# Software Environment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Software Environment**

The Software Environment category in System Information contains information about your system configuration, including details about system drivers, environment variables, and current print jobs. In Software Environment, you can view information about the following items:

- System Drivers
- Signed Drivers
- Environment Variables
- Print Jobs
- Network Connections System Information
- Running Tasks
- Loaded Modules
- Services
- Program Groups
- Startup Programs
- OLE Registration
- Windows Error Reporting

# System Drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Drivers**

System Drivers is located in the Software Environment category in System Information. It displays information about the system drivers that are installed on your system, including their name, description, type, state, status, and other information. For example, you can check the **State** column in the details pane to see which drivers are running and which are stopped.

For more information about Software Environment, see Software Environment.

# Signed Drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Signed Drivers**
Signed Drivers is located in the Software Environment category in System Information. It displays information about the signed drivers that are installed on your system, including their device name, device class, driver version, manufacturer, device ID, and other information. Signed drivers are drivers that have been certified by Microsoft.

To confirm that your hardware is designed for the Windows Server 2003 family of operating systems, see the compatibility information in Support resources. In addition, you can use Device Manager to verify that your system is using signed drivers. For more information about Device Manager, see Device Manager.

For more information about Software Environment, see Software Environment.

For more information about driver signing, see the following:

- Driver signing for Windows
- Designed for Microsoft Windows logo
- To set file signature verification options

# Environment Variables

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Environment Variables**

Environment Variables is located in the Software Environment category in System Information. It displays information about your computer environment, including the number of processors, the operating system path, and the location of temporary folders.

For more information about Software Environment, see Software Environment.

# Print Jobs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Print Jobs**

Print Jobs is located in the Software Environment category in System Information. It contains information about the current print jobs on your system, including document name, size, owner, time submitted, pages printed, and other information. Only documents that are currently printing appear here. For more information about the printers that are installed on your system, see Printing component.

For more information about Software Environment, see Software Environment.

# Network Connections System Information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Connections**

Network Connections is located in the Software Environment category in System Information. It displays information, including local name, remote name, type, status, and user name, about mapped drives, network connections, and other shared resources to which your computer is connected. For more information about Network Connections, see Network Connections.

For more information about Software Environment, see Software Environment.

# Running Tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Running Tasks**

Running Tasks is located in the Software Environment category in System Information. It displays information about processes that are currently running on your system, including the name, path, ID, and version of each process and the priority that the system assigns to each process.

If there is a problem with a running task, such as a program not responding, use Task Manager to locate and troubleshoot the problem. You can use Task Manager to close a task, if necessary.

For more information about Software Environment, see Software Environment.

# Loaded Modules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Loaded Modules**

Loaded Modules is located in the Software Environment category in System Information. It displays information about programs that are running on your system, including their associated DLLs. You can use Loaded Modules to verify the version of a DLL on your system. You can also check the DLL's size, file date, and path.

For more information about Software Environment, see Software Environment.

# Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services**

Services is located in the Software Environment category in System Information. It displays information about the services that are running on your system, including their display name; short name; state; start mode (auto, manual, or disabled); service type; and path.

You can use the Services snap-in to manage the services on your system. For more information about the Services snap-in, see Services Snap-in.

For more information about Software Environment, see Software Environment.

# Program Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Program Groups**

Program Groups is located in the Software Environment category in System Information. It displays information about groups of programs on your system, including the group name, system path and name, and user name that is assigned to the group. These programs appear on the **Start** menu.

You can manage program groups with **Taskbar and Start Menu** in Control Panel.

For more information about Software Environment, see Software Environment.

# Startup Programs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Startup Programs**

Startup Programs is located in the Software Environment category in System Information. It displays information, including the user name and location, about programs that start automatically on your system, even if they are hidden.

You can manage program groups with **Taskbar and Start Menu** in Control Panel.

For more information about Software Environment, see Software Environment.

# OLE Registration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**OLE Registration**

OLE Registration is located in the Software Environment category in System Information. It contains information, including the object and local server, about components that are registered as OLE programs.

For more information about Software Environment, see Software Environment.

# Windows Error Reporting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Error Reporting**
Windows Error Reporting is located in the Software Environment category in System Information. It contains information, including the time, type, and details, about faults that are generated in the Event Log, which you can see by using Event Viewer. For more information about diagnostic tools such as Event Viewer, see Monitoring and Status Tools.

For more information about Software Environment, see Software Environment.

For more information about error reporting, see the following topics:

- System and program error reporting overview
- Exclude programs from error reporting
- Enable error reporting
- Send an error report to Microsoft
- Disable error reporting

# Internet Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Settings**

The Internet Settings category in System Information contains information about Internet settings on the computer. The information in this category is provided by the browser, and it includes information about the browser version, network connections, and cache summary, as well as certificate settings and security settings.

For more information about Internet settings and connections, see Network Connections.

# System Properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Properties**

Using System Properties, you can view and change system properties on a local or remote computer.

- For help with specific tasks, see System Properties How To....
- For general background information, see System Properties Concepts.

# System Properties How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- View system properties
- Set Performance Options
- Add or change environment variables
- Set Startup and Recovery Options

# View system properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view system properties**

1. Open Computer Management.
2. In the console tree, right-click **Computer Management (Local)**, and then click **Properties**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- The **General** tab shows which version of Windows is installed, the name of the person registered to use this copy of Windows, the type of processor, and the total physical memory of the computer. If the operating system was installed by an original equipment manufacturer (OEM), this tab also displays information about technical support.
- You can view your computer's network identification by clicking the **Computer Name** tab.
- To perform this procedure on a remote computer, right-click **Computer Management (Local)**, click **Connect to another computer**, select **Another computer**, and then type in the name of the remote computer. You can then follow the steps in this procedure, starting at step 2, and substituting **Computer Management (*remote computername*)** for **Computer Management (Local)**. You must be a member of the Administrators group, or you must have been delegated the appropriate authority, on the computer that you specify for *remote computername*.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Properties overview
Perform a local or remote shutdown
Windows Management Instrumentation overview

# Set Performance Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set performance options**

- Change processor scheduling options
- Specify memory usage
- Change virtual memory settings

# Change processor scheduling options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change processor scheduling options**

1. Open Computer Management.
2. In the console tree, right-click **Computer Management (Local)**, and then select **Properties**.
3. On the **Advanced** tab, under **Performance**, click **Settings**.
4. Under **Processor scheduling**, select **Programs** or **Background services**. Click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To perform this procedure on a remote computer, right-click **Computer Management (Local)**, click **Connect to another computer**, select **Another computer**, and then type in the name of the remote computer. You can then follow the steps in this procedure, starting at step 2, and substituting **Computer Management (*remote computername*)** for **Computer Management (Local)**. You must be a member of the Administrators group, or you must have been delegated the appropriate authority, on the computer that you specify for *remote computername*.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Properties overview
View system properties
Perform a local or remote shutdown
Windows Management Instrumentation overview

# Specify memory usage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify memory usage**

1. Open Computer Management.
2. In the console tree, right-click **Computer Management (Local)**, and then click **Properties**.
3. On the **Advanced** tab, under **Performance**, click **Settings**.
4. Under **Memory usage**, select **Programs** or **System cache**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To perform this procedure on a remote computer, right-click **Computer Management (Local)**, click **Connect to another computer**, select **Another computer**, and then type in the name of the remote computer. You can then follow the steps in this procedure, starting at step 2, and substituting **Computer Management (*remote computername*)** for **Computer Management (Local)**. You must be a member of the Administrators group, or you must have been delegated the appropriate authority, on the computer that you specify for *remote computername*.
- On all Windows Server 2003 family operating systems, the **System cache** option is selected by default.
- You must restart your computer before your changes take effect.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change processor scheduling options
Change virtual memory settings

# Change virtual memory settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change virtual memory settings**

1. Open Computer Management.
2. In the console tree, right-click **Computer Management (Local)**, and then select **Properties**.
3. On the **Advanced** tab, under **Performance**, click **Settings**.
4. Under **Virtual memory**, click **Change**.
5. In the **Drive** list, click the drive that contains the paging file you want to change.
6. Under **Paging file size for selected drive**, click **Custom size**, type new paging file sizes in the **Initial Size (MB)** and **Maximum Size (MB)** boxes, and then click **Set**. Click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To perform this procedure on a remote computer, right-click **Computer Management (Local)**, click **Connect to another computer**, select **Another computer**, and then type in the name of the remote computer. You can then follow the steps in this procedure, starting at step 2, and substituting **Computer Management (*remote computername*)** for **Computer Management (Local)**. You must be a member of the Administrators group, or you must have been delegated the appropriate authority, on the computer that you specify for *remote computername*.
- For best performance, set the initial size to not less than the recommended size under **Total paging file size for all drives**. The recommended size is equivalent to 1.5 times the amount of RAM on your system.
- Usually, you should leave the paging file at its recommended size, although you might increase its size if you routinely use programs that require a lot of memory.
- If you decrease the size of the paging file, you must restart your computer to see the effects of those changes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Properties overview
View system properties
Perform a local or remote shutdown
Windows Management Instrumentation overview

# Add or change environment variables

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add or change environment variables**

1. Open Computer Management.
2. In the console tree, right-click **Computer Management (Local)**, and then click **Properties**.
3. On the **Advanced** tab, under **Environment Variables**, click **Settings**.
4. Select a user in the **User variables for** list.
5. Click the name of the user or system variable you want to change or delete.
   - Click **New** to add a new variable name and value.
   - Click **Edit** to change a current variable name and value.
   - Click **Delete** to remove a variable name and value.

## Notes

- You are not required to have administrative credentials to create, edit or delete variables under **User variables for**. Therefore, as a security best practice, consider performing this task as a user without administrative credentials. You must be a member of the Administrators group on the local computer in order to create, edit or delete the variables listed under **System variables**. For more information, see Default local groups.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To perform this procedure on a remote computer, right-click **Computer Management (Local)**, click **Connect to another computer**, select **Another computer**, and then type in the name of the remote computer. You can then follow the steps in this procedure, starting at step 2, and substituting **Computer Management (*remote computername*)** for **Computer Management (Local)**. You must be a member of the Administrators group, or you must have been delegated the appropriate authority, on the computer that you specify for *remote computername*.
- Changes in the registry are saved so they are available automatically the next time you start your computer.
- Changes do not affect programs that are running and may not take effect until you restart your computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Properties overview
View system properties
Windows Management Instrumentation overview

# Set Startup and Recovery Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set startup and recovery options**

- Specify the default operating system for startup
- Specify what happens when the system stops unexpectedly
- Perform a local or remote shutdown

# Specify the default operating system for startup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify the default operating system for startup**

1. Open Computer Management.
2. In the console tree, right-click **Computer Management (Local)**, and then select **Properties**.
3. On the **Advanced** tab, under **Startup and Recovery**, click **Settings**.
4. Under **System startup**, in the **Default operating system** list, click the operating system you want to start if you do not select one from the list displayed during startup.
5. Select **Time to display list of operating systems** box, and then enter the number of seconds the list should be displayed before the default operating system starts automatically.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To perform this procedure on a remote computer, right-click **Computer Management (Local)**, click **Connect to another computer**, select **Another computer**, and then type in the name of the remote computer. You can then follow the steps in this procedure, starting at step 2, and substituting **Computer Management (*remote computername*)** for **Computer Management (Local)**. You must be a member of the Administrators group, or you must have been delegated the appropriate authority, on the computer that you specify for *remote computername*.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Properties overview
View system properties
Windows Management Instrumentation overview

# Specify what happens when the system stops unexpectedly

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify what happens when the system stops unexpectedly**

1. Open Computer Management.
2. In the console tree, right-click **Computer Management (Local)**, and then click **Properties**.
3. On the **Advanced** tab, under **Startup and recovery**, click **Settings**.
4. Under **System startup**, select the **Time to display list of operating systems** check box, and then type the number of seconds you want the list of recovery options to display before the default recovery option is automatically selected.
5. Under **System failure**, select the check boxes that correspond to the actions you want Windows to perform if a Stop error occurs:

   - **Write an event to the system log** specifies that event information will be recorded in the system log.

     On computers running a Windows Server 2003 family operating system, you cannot disable this feature. Windows always writes event information to the system log.

   - **Send an administrative alert** specifies that your system administrator will be notified.

   - **Automatically restart** specifies that Windows will automatically restart your computer.

6. Under **Write debugging information**, choose the type of information you want Windows to record when the system stops unexpectedly:

   **Small memory dump**

   Records the smallest amount of information that will help identify the problem. This option requires a paging file of at least 2 MB on the boot volume of your computer. If you choose this option, Windows creates a new file (64 KB in size) each time the system stops unexpectedly. A history of these files is stored in the directory listed under **Small dump directory**.

   **Kernel memory dump**

   Records only kernel memory, which stores more information than small memory dump but takes less time to complete than the complete memory dump when the system stops unexpectedly. The file is stored in the directory listed under **Dump file**. If you choose this option you must have a sufficiently large paging file on the boot volume. The required size depends on the amount of RAM in your computer (although the maximum amount of space that must be available for a kernel memory dump is 2,060 MB). The following table provides guidelines:

| RAM size | Paging file should be no smaller than |
|---|---|
| 256 MB-1,373 MB | 1.5 times the RAM size |
| 1,374 MB or greater | 2,060 MB (maximum amount of data in a kernel memory dump) |

   **Complete memory dump**

   Not available on computers with 2 or more GB of RAM. Records the entire contents of system memory when the system stops unexpectedly. If you choose this option you must have a paging file on the boot volume large enough to hold all of

the physical RAM plus one megabyte (MB). The file is stored in the directory listed under **Dump file**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- To perform this procedure on a remote computer, right-click **Computer Management (Local)**, click **Connect to another computer**, select **Another computer**, and then type in the name of the remote computer. You can then follow the steps in this procedure, starting at step 2, and substituting **Computer Management (*remote computername*)** for **Computer Management (Local)**. You must be a member of the Administrators group, or you must have been delegated the appropriate authority, on the computer that you specify for *remote computername*.

- You must have at least a 2-MB paging file on the computer's boot volume if you select **Write an event to the system log** or **Send an administrative alert**.

- If you choose **Kernel memory dump**, be sure not to set the paging file smaller than the sizes listed in the procedure, or you might not capture the information you need.

- If you choose either **Kernel memory dump** or **Complete memory dump** and select the **Overwrite any existing file** check box, Windows always writes to the same file name. To save individual dump files, clear the **Overwrite any existing file** check box and change the file name after each Stop error.

- When setting memory dump options, check your settings for the size of the paging file on the boot volume. Set the file size according to the preceding guidelines.

  It is recommended that you use the same value for both the minimum size and the maximum size of the paging file. This prevents the file from constantly growing and shrinking, which in turn prevents unnecessary disk fragmentation.

- Memory dumps might contain sensitive information, so you should store them in a secure location and limit the number of users who can access them. When you transfer a memory dump across the network or send a memory dump over the Internet, always use a secure file-transfer mechanism to guard against the file being compromised by a malicious user.

- In computers with 8 or more processors that have the maximum amount of RAM installed, you might be able to gain some incremental performance improvement by splitting the paging file into multiple paging files. Each paging file must be on a separate physical disk and for reliability, each disk must be part of a hardware RAID-5 volume. For more information about changing the size of the virtual memory paging file, see Related Topics.

- You can save memory by clearing the **Write an event to the system log** and **Send an administrative alert** check boxes. The memory saved depends on the computer, but these features typically require about 60 KB to 70 KB.

- If you contact Microsoft Product Support Services about a Stop error, they might ask for the memory dump file created when the Stop error occurred.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify the default operating system for startup

Change the size of the virtual memory paging file
Perform a local or remote shutdown

# Perform a local or remote shutdown

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To perform a local or remote shutdown**

1. Open Computer Management.
2. In the console tree, right-click **Computer Management (Local)** , and then click **Properties**.
3. On the **Advanced** tab, click **Startup and Recovery**.
4. Click **Shut Down** to open the **Shut Down** dialog box.
5. Under **Action**, select the actions you want to perform on the computer to which you are connected.
6. Under **Force Apps Closed**, select the circumstances under which you want to force applications to close when you shut down or restart the computer, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To perform this procedure on a remote computer, right-click **Computer Management (Local)**, click **Connect to another computer**, select **Another computer**, and then type in the name of the remote computer. You can then follow the steps in this procedure, starting at step 2, and substituting **Computer Management (*remote computername*)** for **Computer Management (Local)**. You must be a member of the Administrators group, or you must have been delegated the appropriate authority, on the computer that you specify for *remote computername*.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Properties overview
View system properties
Windows Management Instrumentation overview

# System Properties Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

- System Properties overview

# System Properties overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Properties overview**

System Properties is a Windows Management Instrumentation (WMI) tool you access through **Computer Management**. Using System Properties, you can view and change system properties on a local computer or remote computer.

For more information about WMI, see Windows Management Instrumentation overview.

You can also view and change system properties on a local computer by accessing **System** in Control Panel. This is a separate tool that enables you to perform additional tasks on your local computer, but does not enable you to view or change system properties on a remote computer. For more information see System Properties Overview - Control Panel.

# Services Snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services**

You can use the Services snap-in to manage services on local or remote computers.

- For tips about using Services, see Services best practices.
- For help with specific tasks, see Services How To....
- For general background information, see Services Concepts.
- For problem-solving instructions, see Troubleshooting Services Snap-in.

**Important**

- To improve performance and security in the Microsoft® Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

# Services best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services best practices**

**Save the default state of services.**

- It is helpful to keep a list of the default service settings in the event that you need to restore these settings. For more information, see Export services information from the Services snap-in. You can also use Security Templates to define a default security template, including system service settings, for your computer or for your network. For more information, see Security Templates overview. For more information about how restore default security settings, see Reapply default security settings.

**Perform adequate testing and backup procedures before changing default settings.**

- Changing the default settings for services might prevent key services from running correctly. It is especially important to use caution when changing the **Startup type** and **Log on as** settings of services that are configured to start automatically.

- In most cases, it is recommended that you not change the **Allow service to interact with desktop** setting. If you allow the service to interact with the desktop, any information that the service displays on the desktop will also be displayed on an interactive user's desktop. A malicious user could then take control of the service or attack it from the interactive desktop.

**Verify service dependencies before starting or stopping services.**

- Ensure that a service you are starting does not depend on a service whose startup type is set to **Disabled** or **Manual**. Also, make sure that a service you are stopping is not required for a dependent service to function properly. For more information, see View service dependencies.

**Notify connected users before stopping or pausing services.**

- Stopping or pausing certain services affects connected users. For example, if you stop the Server service, all file, print, and named pipe sharing over the network is disabled. Also, the affected computers can no longer be administered remotely, so you must start the Server service locally.

  Stopping or pausing a service does not change the startup type of that service. For example, if you stop a service that is configured to start automatically (**Startup type** is set to **Automatic**), that service starts automatically when the computer is restarted.

**When you start disabled services, be sure to select the appropriate startup type.**

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. If you want to enable any of these services, you must select the appropriate startup type. For a table that lists the default settings and provides information about the appropriate startup type for a disabled service, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

# Services How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services how to...**

- Configure how a service is started
- Start, stop, pause, resume, or restart a service
- Enable or disable a service for a hardware profile
- Access the Services snap-in on another computer
- Set up recovery actions to take place when a service fails
- Export services information from the Services snap-in
- View service dependencies

# Configure how a service is started

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure how a service is started**

- Using Services
- Using the command line

**Using Services**

1. Open Services.
2. In the details pane, right-click the service that you want to configure, and then click **Properties**.
3. On the **General** tab, in **Startup type**, click **Automatic**, **Manual**, or **Disabled**.
4. To specify the user account that the service can use to log on, click the **Log On** tab, and then do one of the following:
   - To specify that the service uses the Local System account, click **Local System account**.
   - To specify that the service uses the Local Service account, click **This account**, and then type **NT AUTHORITY\LocalService**.
   - To specify that the service uses the Network Service account, click **This account**, and then type **NT AUTHORITY\NetworkService**.
   - To specify another account, click **This account**, click **Browse**, and then specify a user account in the **Select User** dialog box. When you are finished, click **OK**.
5. Type the password for the user account in **Password** and in **Confirm password**, and then click **OK**. If you select the Local Service account or Network Service account, the password must be blank.

### Note

- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.

**Using the command line**

1. Open Command Prompt.
2. Type:

   **sc config***service name***start=** {**boot**|**system**|**auto**|**demand**|**disabled**}

| Value | Description |
|---|---|
| **sc config** | Modifies the value of a service's entries in the registry and in the Service Control Manager's database. |
| *service name* | Specifies the short name of the service. |

| | |
|---|---|
| **start=** | Specifies the start type for the service. |
| **boot** | A device driver that is loaded by the boot loader. |
| **system** | A device driver that is started during kernel initialization. |
| **auto** | A service that automatically starts each time the computer is restarted and runs even if no one logs on to the computer. |
| **demand** | A service that must be manually started. This is the default value if **start=** is not specified. |
| **disabled** | A service that cannot be started. To start a disabled service, change the start type to one of the other values. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- To view the complete syntax for this command, at the command prompt, type:

  **sc config help**

- For more information about the **sc** command, see Related topics.

**Caution**

- Most services are not designed to have the default account changed. Changing the default account of a service could result in the service failing to start.
- Changing the default service settings may prevent key services from running correctly. It is especially important to use caution when changing the **Startup type** and **Log on as** settings of services that are configured to start automatically.
- In most cases, it is recommended that you not change the **Allow service to interact with desktop** setting. If you allow the service to interact with the desktop, any information that the service displays on the desktop will also be displayed on an interactive user's desktop. A malicious user could then take control of the service or attack it from the interactive desktop.

**Important**

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.
- The Local Service account and Network Service account are configured with a null password. Whatever password information you supply is ignored. For more information about service accounts, see Services permissions.
- It is recommended that user accounts that are used to log on as a service have the **Password never expires** check box selected in their properties dialog box and that they have strong passwords. For more information about creating strong passwords, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- If account lockout policy is enabled and the account is locked out, the service will not start. For more information about account lockout policy, see Related Topics.

- If you enable or disable a service and you encounter a problem starting the computer, you can start the computer in Safe Mode. In Safe Mode, core services that are required to start the operating system are started in a default scheme, regardless of any changes that are made to the service settings. After the computer is in Safe Mode, you can change the service configuration or restore the default configuration. For more information about startup options, see Related Topics.

- If you specify an account that does not have permission to log on as a service, the Services snap-in automatically grants the appropriate permissions to that account on the computer that you are managing.

- For more information about the user accounts that a service uses to log on, see "Services permissions" in Related Topics. Also, see "Managing System Services" at the Microsoft Windows Resource Kits Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start, stop, pause, resume, or restart a service
Default settings for services
Services permissions
SC
Account Lockout Policy
Startup options
Strong passwords

# Start, stop, pause, resume, or restart a service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start, stop, pause, resume, or restart a service**

- Using Services
- Using the command line

**Using Services**

1. Open Services.
2. In the details panel, do one of the following:
    - Click the service, and then, on the **Action** menu, click **Start**, **Stop**, **Pause**, **Resume**, or **Restart**.
    - Right-click the service, and then click **Start**, **Stop**, **Pause**, **Resume**, or **Restart**.

### Notes

- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.
- To start a service with startup parameters, right-click the service, click **Properties**, type the parameters in **Start parameters**, and then click **Start**. These settings are not persistent; they are used only once, and then the default settings are restored. (A backslash (\) is treated as an escape character; type two backslashes for each backslash in a parameter.)

**Using the command line**

1. Open Command Prompt.
2. Type one of the following:
    - To start a service, type:

      **net start***service*

    - To stop a service, type:

      **net stop***service*

    - To pause a service, type:

      **net pause***service*

    - To resume a service, type:

      **net continue***service*

| Value | Description |
|---|---|

| | |
|---|---|
| **net start** | Starts a service. |
| **net stop** | Stops a service. |
| **net pause** | Pauses a service. |
| **net continue** | Continues a service that has been paused. |
| *service* | Specifies the name of service. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- To view the complete syntax for these commands, at the command prompt, type:

  **net help***command*

- For more information about these commands, see Related Topics.

**Caution**

- If you stop, start, or restart a service, any dependent services are also affected. Starting a service does not automatically restart its dependent services. For more information, see View service dependencies.
- Changing the default service settings may prevent key services from running correctly. It is especially important to use caution when changing the **Startup type** and **Log on as** settings of services that are configured to start automatically.
- In most cases, it is recommended that you not change the **Allow service to interact with desktop** setting. If you allow the service to interact with the desktop, any information that the service displays on the desktop will also be displayed on an interactive user's desktop. A malicious user could then take control of the service or attack it from the interactive desktop.

**Important**

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- Some services, such as Remote Procedure Call (RPC), Event Log, and Plug and Play, cannot be stopped by using the Service snap-in or the **net stop** command. These services are required for the operating system to function properly.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure how a service is started

Net start
Net stop
**Other Resources**
Net pause
Net continue

# Enable or disable a service for a hardware profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable a service for a hardware profile**

1. Open Services.
2. In the details pane, right-click the service that you want to enable or disable, and then click **Properties**.
3. On the **Log On** tab, click the hardware profile that you want to configure.
4. Click **Enable** or **Disable**, and then click **OK**.

### Caution

- Use caution when you disable services by using this procedure. If you disable a service for a certain hardware profile, the **Hardware Profile** settings override the **Startup type** setting for that service.
- In most cases, it is recommended that you not change the **Allow service to interact with desktop** setting. If you allow the service to interact with the desktop, any information that the service displays on the desktop will also be displayed on an interactive user's desktop. A malicious user could then take control of the service or attack it from the interactive desktop.

### Important

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.
- You can use the options on the **Hardware** tab in System in Control Panel to create and manage hardware profiles. For more information about creating hardware profiles, see Related Topics.
- You may want to set up a hardware profile to test any newly installed services. For example, you can create two hardware profiles for a newly installed service: one profile in which the service is enabled and a second profile in which the service is disabled. This way, you can troubleshoot any problems that might occur, such as a driver not loading properly.
- If you receive a Connection Manager error message when you open the properties dialog box for a service, verify that the Remote Registry service is running on the specified computer. If the Remote Registry service is not running, or if the specified computer is running Windows NT 4.0, you cannot view or modify hardware profile settings, but you can perform any other actions. For more information, see Troubleshooting Services Snap-in. Also, see "Managing System Services" at the Microsoft Windows Resource Kits Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure how a service is started
Start, stop, pause, resume, or restart a service
Create a hardware profile
SC

# Access the Services snap-in on another computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To access the Services snap-in on another computer**

1. Open Services.
2. In the console tree, right-click **Services**, and then click **Connect to another computer**.
3. In the **Select Computer** dialog box, type the name of the computer that you want to manage, or click **Browse** to find it, and then click **OK**.

### Caution

- Changing the default service settings may prevent key services from running correctly. It is especially important to use caution when changing the **Startup type** and **Log on as** settings of services that are configured to start automatically.
- In most cases, it is recommended that you not change the **Allow service to interact with desktop** setting. If you allow the service to interact with the desktop, any information that the service displays on the desktop will also be displayed on an interactive user's desktop. A malicious user could then take control of the service or attack it from the interactive desktop.

### Important

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.
- If you enable or disable a service and you encounter a problem starting the computer, you can start the computer in Safe Mode. Then, you can change the service configuration or restore the default configuration. In Safe Mode, core services that are required to start the operating system are started in a default scheme, regardless of any changes that are made to the service settings. For more information about Safe Mode and other startup options, see Related Topics.

### Notes

- You must be a domain member on the specified computer to access the Services snap-in.
- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start, stop, pause, resume, or restart a service
Default settings for services
SC
Startup options

# Set up recovery actions to take place when a service fails

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up recovery actions to take place when a service fails**

1. Open Services.
2. In the details pane, right-click the service for which you want to set up recovery actions, and then click **Properties**.
3. On the **Recovery** tab, click the action that you want in **First failure**, **Second failure**, and **Subsequent failures**, and then click **OK**:
   - If you click **Run a Program**, under **Run program**, type the absolute path for the specified computer. Universal Naming Convention (UNC) names are not supported. For example, type **C:\scripts\handlefailure.cmd**, rather than *****computername***\c$\scripts\handlefailure.cmd**.
   - If you click **Restart the Computer**, click **Restart Computer Options** to specify how long the computer waits before restarting. You can also create a message to send to remote users before the computer restarts. This message is sent automatically, using the same mechanism as the one used by the **net send** command and Send Console Message.

### Important

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.
- Recovery actions are available only on computers running Windows 2000, Windows XP, or products in the Windows Server 2003 family.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure how a service is started
Start, stop, pause, resume, or restart a service
SC
**Other Resources**
Net send
Send Console Message Overview

# Export services information from the Services snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export services information from the Services snap-in**

1. Open Services.
2. On the **Action** menu, click **Export List**. This opens the **Export List** dialog box.
3. In **Save in**, click or type the file location where you want to save the exported file.
4. In **File name**, type the name of the file.
5. In **Save as type**, click the file type, and then click **Save**.

### Important

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

### Note

- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start, stop, pause, resume, or restart a service
Default settings for services
Services permissions
SC

# View service dependencies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view service dependencies**

1. Open Services.
2. In the details pane, right-click the service that you want to view dependencies for, and then click **Properties**.
3. Click the **Dependencies** tab.
4. To view services that are associated dependencies of the selected service, in the list on the **Dependencies** tab, click the plus sign next to the service.

### Important

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

### Notes

- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.
- The top list on the **Dependencies** tab identifies the other services that must be running for the selected service to function. The bottom list on the **Dependencies** tab identifies the services that require the selected service to be running so that the dependent services can function properly.
- For more information about service dependencies, see "Managing System Services" at the Microsoft Windows Resource Kits Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start, stop, pause, resume, or restart a service
SC

# Services Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services concepts**

This section provides general background information about Services:

- Services overview
- Services permissions
- Default settings for services

**Important**

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

# Services overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services overview**

A service is an application type that runs in the system background and is similar to a UNIX daemon application. Services provide core operating system features, such as Web serving, event logging, file serving, help and support, printing, cryptography, and error reporting. With the Services snap-in, you can manage services on local or remote computers. The operating system services that are delivered with the Microsoft® Windows Server 2003 family are designed so that only the key services that are required for common server roles are started. For more information on server roles, see Configuring Roles for Your Server.

You can use the Services snap-in to:

- Start, stop, pause, resume, or disable services on remote and local computers.
- Manage services on local and remote computers (on remote computers running Microsoft Windows NT® 4.0, Windows 2000, Windows XP, or products in the Windows Server 2003 family only).
- Set up recovery actions to take place if a service fails, for example, restarting the service automatically or restarting the computer (on computers running Windows NT 4.0, Windows 2000, Windows XP, or products in the Windows Server 2003 family only).
- Enable or disable services for a particular hardware profile.
- Export service information to a .txt or .csv file for system administration purposes.
- View the status and description of each service.

For more information about managing operating system services, see "Managing System Services" at the Microsoft Windows Resource Kits Web site.

**Important**

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

# Services permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services permissions**

There are two types of permissions that apply to services: service account permissions and service permissions. Service account permissions refer to the user rights and credentials that are granted to the service through the logon account. Service permissions refer to the permissions that are required to configure a service.

**Service account permissions**

A service must log on to an account to access resources and objects on the operating system. Most services are not designed to have their default logon account changed. Changing the default account will probably cause the service to fail.

If you select an account that does not have permission to log on as a service, the Services snap-in automatically grants that account the user rights that are required to log on as a service on the computer that you are managing. However, this does not guarantee that the service will start.

It is recommended that the user accounts that are used to log on as a service have the **Password never expires** check box selected in their properties dialog box and that they have strong passwords. For more information, see Strong passwords.

If account lockout policy is enabled and the account is locked out, the service will malfunction. For more information, see Account Lockout Policy.

The following table describes the service logon accounts and how they are used.

| Logon account | Description |
|---|---|
| Local System account | The Local System account is a powerful account that has full access to the system, including the directory service on domain controllers. If a service logs on to the Local System account on a domain controller, that service has access to the entire domain. Some services are configured by default to log on to the Local System account. Do not change the default service setting. |
| Local Service account | The Local Service account is a special, built-in account that is similar to an authenticated user account. The Local Service account has the same level of access to resources and objects as members of the Users group. This limited access helps safeguard your system if individual services or processes are compromised. Services that run as the Local Service account access network resources as a null session with no credentials. |
| Network Service account | The Network Service account is a special, built-in account that is similar to an authenticated user account. The Network Service account has the same level of access to resources and objects as members of the Users group. This limited access helps safeguard your system if individual services or processes are compromised. Services that run as the Network Service account access network resources using the credentials of the computer account. |

**Caution**

- Changing the default service settings may prevent key services from running correctly. It is especially important to use caution when changing the **Startup type** and **Log on as** settings of services that are configured to start automatically.
- In most cases, it is recommended that you not change the **Allow service to interact with desktop** setting. If you allow the service to interact with the desktop, any information that the service displays on the desktop will also be displayed on an interactive user's desktop. A malicious user could then take control of the service or attack it from the interactive desktop.

**Service permissions**

Each service has special permissions that you can grant or deny for each user or group. You can set permissions for individual services by using Security Templates or the **sc** command. For more information, see Security Templates and SC.

The following table lists the individual service permissions that you can apply.

| Permission | Allows you to |
| --- | --- |
| Full Control | Perform all functions. This permission automatically grants all service permissions to the user. |
| Query Template | Determine the configuration parameters that are associated with a service object. |
| Change Template | Change the configuration of a service. This permission is required so that the user can change the startup type. |
| Query Status | Access information about the status of the service. |
| Enumerate Dependents | Determine all the other services that depend on the specified service. |
| Start | Start a service. |
| Stop | Stop a service. |
| Pause and Continue | Pause and continue the service. |
| Interrogate | Report the current status information for the service. |
| User-Defined Control | Send a user-defined control request--or a request that is specific to the service--to the service. |
| Delete | Delete a service. |
| Read Permissions | Read the security permissions that are assigned to the service. |
| Change Permissions | Change the security permissions that are assigned to the service. |
| Take Ownership | Change a security key or change permissions on a service that is not owned by the user. |

**Important**

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. For a table that lists the default settings and provides information about how to enable these services, see Default settings for services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

# Default settings for services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Default settings for services**
For a typical installation, services are configured as automatic, manual, or disabled:

- An automatic service starts automatically when the operating system starts.
- You can start a service manually by using the Services snap-in or by using the **net start** command. A manual service can also be started by a related operating system service, a system device driver, or an action in the user interface that depends on the manual service. For example, both the Telephony service and Remote Access Connection Manager service are configured as **Manual** by default, but both of these services are started during the Windows Product Activation process.
- A disabled service cannot be started automatically or manually. To start a disabled service, you must select an appropriate startup type, which is identified in the tables below. For information on enabling the services required by audio, imaging, and CD-burning peripheral devices, see Troubleshooting audio, imaging, and CD-burning peripheral devices.

For more information about how to configure service startup, see Configure how a service is started. For more information about starting a service from the command line, see Net start.

You can use the **sc** command at the command prompt to configure services and obtain information about them. For example, you can use **sc getkeyname** to obtain the key name, or short name, of the service. For more information about the **sc** command, see SC.

**Important**

- To improve performance and security in the Windows Server 2003 family, several services have been disabled by default that were previously enabled on Windows 2000. See the tables below for information about how to enable these services. Note that these settings apply only to new installations, not upgrades; all previous service configurations are preserved during upgrades to the Windows Server 2003 family.

The following tables list all the default services that may be installed as part of the Windows Server 2003 family. Your list of services may be different, depending on which Windows product is installed, which components and options are selected during installation, and whether any additional programs are installed.

**Alerter**
Short name: Alerter

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Automatic**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local Service |
| Windows Server 2003, Enterprise Edition | Disabled | Local Service |
| Windows Server 2003, Datacenter Edition | Disabled | Local Service |
| Windows Server 2003, Web Edition | Disabled | Local Service |

**Application Layer Gateway Service**
Short name: ALG

| Product | Startup type | Log on as |
|---|---|---|

| | | |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local Service |
| Windows Server 2003, Enterprise Edition | Manual | Local Service |
| Windows Server 2003, Datacenter Edition | Manual | Local Service |
| Windows Server 2003, Web Edition | Manual | Local Service |

**Application Management**
Short name: AppMgmt

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Automatic Updates**
Short name: wuauserv

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Background Intelligent Transfer Service**
Short name: BITS

For information about troubleshooting problems associated with this service, see Background Intelligent Transfer Service (BITS) inactive job timeout.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**ClipBook**
Short name: ClipSrv

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Manual**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |

| | | |
|---|---|---|
| Windows Server 2003, Web Edition | Disabled | Local System |

**COM+ Event System**
Short name: EventSystem

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual/Started | Local System |
| Windows Server 2003, Enterprise Edition | Manual/Started | Local System |
| Windows Server 2003, Datacenter Edition | Manual/Started | Local System |
| Windows Server 2003, Web Edition | Manual/Started | Local System |

**COM+ System Application**
Short name: COMSysApp

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Computer Browser**
Short name: Browser

This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Cryptographic Services**
Short name: CryptSvc

This service is started in Safe Mode, Safe Mode with Networking, and Safe Mode with Command Prompt.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**DCOM Server Process Launcher**
Short name: DcomLaunch

This service is started in Safe Mode, Safe Mode with Networking, and Safe Mode with Command Prompt.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**DHCP Client**
Short name: Dhcp

This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Network Service |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Network Service |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Network Service |
| Windows Server 2003, Web Edition | Automatic/Started | Network Service |

**Distributed File System**
Short name: Dfs

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Distributed Link Tracking Client**
Short name: TrkWks

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Distributed Link Tracking Server**
Short name: TrkSvr

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Manual**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |

| | | |
|---|---|---|
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Distributed Transaction Coordinator**
Short name: MSDTC

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Network Service |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Network Service |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Network Service |
| Windows Server 2003, Web Edition | Automatic/Started | Network Service |

**DNS Client**
Short name: Dnscache

This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Network Service |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Network Service |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Network Service |
| Windows Server 2003, Web Edition | Automatic/Started | Network Service |

**Error Reporting Service**
Short name: ERSvc

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Event Log**
Short name: Eventlog

This service is started in Safe Mode, Safe Mode with Networking, and Safe Mode with Command Prompt.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**File Replication**
Short name: ntfrs

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Help and Support**
Short name: helpsvc

This service is started in Safe Mode, Safe Mode with Networking, and Safe Mode with Command Prompt.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**HTTP SSL**
Short name: W3SSL

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual/Started | Local System |

**Human Interface Device Access**
Short name: HidServ

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**IAS Jet Database Access**
Short name: IASJet

This service is available only on the Itanium-based versions of the Windows Server 2003 family.

| Product | Startup type | Log on as |
|---|---|---|
| Itanium-based version of Windows Server 2003, Enterprise Edition | Manual | Local System |
| Itanium-based version of Windows Server 2003, Datacenter Edition | Manual | Local System |

**IIS Admin**
Short name: IISADMIN

This service is installed by default only on computers running the Windows Server 2003, Web Edition operating system. For other editions, you can this service by using the Configure Your Server Wizard or by using Add or Remove Programs in Control Panel.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**IMAPI CD-Burning COM Service**
Short name: ImapiService

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Manual**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Indexing Service**
Short name: cisvc

For more information about this service, see Windows interface administrative tool reference A-Z: Indexing Service.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Internet Connection Firewall (ICF)/Internet Connection Sharing (ICS)**
Short name: SharedAccess

If you want to enable this service, set **Startup type** to **Manual**.

Internet Connection Firewall is included only in the original release of Windows Server 2003, Standard Edition, and the 32-bit version of Windows Server 2003, Enterprise Edition.

Instead of using Internet Connection Firewall and the Internet Connection Firewall (ICF)/Internet Connection Sharing (ICS) service, computers running Windows Server 2003 with Service Pack 1 (SP1) use Windows Firewall and the Windows Firewall/Internet Connection Sharing (ICS) service.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |

**Intersite Messaging**
Short name: IsmServ

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**IPSec Services**
Short name: PolicyAgent

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Kerberos Key Distribution Center**
Short name: kdc

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**License Logging**
Short name: LicenseService

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Network Service |
| Windows Server 2003, Enterprise Edition | Disabled | Network Service |
| Windows Server 2003, Datacenter Edition | Disabled | Network Service |
| Windows Server 2003, Web Edition | Disabled | Network Service |

**Logical Disk Manager**
Short name: dmserver

This service is started in Safe Mode, Safe Mode with Networking, and Safe Mode with Command Prompt.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Logical Disk Manager Administrative Service**
Short name: dmadmin

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Messenger**
Short name: messenger

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Automatic**. This service must be running for alert notifications and console messages to be received. For more information, see Send Console Message overview. This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Microsoft Software Shadow Copy Provider**
Short name: SWPRV

This service is installed by default only on computers that are running the Windows Server 2003, Web Edition operating system. For other editions, you can install it by using Add or Remove Programs in Control Panel or by using Manage Your Server.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Network Service |
| Windows Server 2003, Enterprise Edition | Manual | Network Service |
| Windows Server 2003, Datacenter Edition | Manual | Network Service |
| Windows Server 2003, Web Edition | Manual | Network Service |

**MSSQL$UDDI**
Short name: MSSQLSERVER

By default, this service is not installed with a typical installation. You can install it by using Add or Remove Programs in Control Panel. When this service installed, the Startup Type for it is Manual.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |

| | | |
|---|---|---|
| Windows Server 2003, Web Edition | Manual | Local System |

**MSSQLserverADHelper**
Short name: MSSQLserverADHelper

By default, this service is not installed with a typical installation. You can install it by using Add or Remove Programs in Control Panel. When this service installed, the Startup Type for it is Manual.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**.NET Framework Support Service**
Short name: CORRTSvc

By default, this service is not installed with a typical installation. You can install it by using Add or Remove Programs in Control Panel. When this service installed, the Startup Type for it is Manual.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Net Logon**
Short name: Netlogon

Until a user joins a Domain, Net Logon is Manual. If a user joins a domain during installation, the service starts and is configured as Automatic.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**NetMeeting Remote Desktop Sharing**
Short name: mnmsrvc

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Manual**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |

| | | |
|---|---|---|
| Windows Server 2003, Web Edition | Disabled | Local System |

**Network Connections**
Short name: Netman

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual/Started | Local System |
| Windows Server 2003, Enterprise Edition | Manual/Started | Local System |
| Windows Server 2003, Datacenter Edition | Manual/Started | Local System |
| Windows Server 2003, Web Edition | Manual/Started | Local System |

**Network DDE**
Short name: NetDDE

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Manual**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Network DDE DSDM**
Short name: NetDDEDSDM

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Manual**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Network Location Awareness (NLA)**
Short name: NLA

This service starts automatically when **Startup type** is set to **Manual**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual/Started | Local System |
| Windows Server 2003, Enterprise Edition | Manual/Started | Local System |
| Windows Server 2003, Datacenter Edition | Manual/Started | Local System |
| Windows Server 2003, Web Edition | Manual/Started | Local System |

**NT LM Security Support Provider**
Short name: NtLmSsp

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Performance Logs and Alerts**
Short name: SysmonLog

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Network Service |
| Windows Server 2003, Enterprise Edition | Manual | Network Service |
| Windows Server 2003, Datacenter Edition | Manual | Network Service |
| Windows Server 2003, Web Edition | Manual | Network Service |

**Plug and Play**
Short name: PlugPlay

This service is started in Safe Mode, Safe Mode with Networking, and Safe Mode with Command Prompt.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Portable Media Serial Number**
Short name: WmdmPmSp

This service is not available on the Itanium-based versions of the Windows Server 2003 family.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Print Spooler**
Short name: Spooler

If you are having trouble with a printer not responding, stop and then restart this service. This action cancels all pending print jobs in the printer queue.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |

| | | |
|---|---|---|
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Protected Storage**
Short name: ProtectedStorage

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Remote Access Auto Connection Manager**
Short name: RasAuto

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Remote Access Connection Manager**
Short name: RasMan

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Remote Administration Service**
Short name: SRVCSURG

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Remote Desktop Help Session Manager**
Short name: RDSessMgr

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |

| | | |
|---|---|---|
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Remote Procedure Call (RPC)**
Short name: RpcSs

This service is started in Safe Mode, Safe Mode with Networking, and Safe Mode with Command Prompt.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Network Service |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Network Service |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Network Service |
| Windows Server 2003, Web Edition | Automatic/Started | Network Service |

**Remote Procedure Call (RPC) Locator**
Short name: RpcLocator

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Network Service |
| Windows Server 2003, Enterprise Edition | Manual | Network Service |
| Windows Server 2003, Datacenter Edition | Manual | Network Service |
| Windows Server 2003, Web Edition | Manual | Network Service |

**Remote Registry**
Short name: RemoteRegistry

To troubleshoot performance problems related to this service, see Troubleshooting Performance monitoring.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local Service |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local Service |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local Service |
| Windows Server 2003, Web Edition | Automatic/Started | Local Service |

**Remote server Manager**
Short name: APPMGR

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Removable Storage**
Short name: NtmsSvc

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |

| | | |
|---|---|---|
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Resultant Set of Policy Provider**
Short name: RSoPProv

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Routing and Remote Access**
Short name: RemoteAccess

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Manual**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Secondary Logon**
Short name: seclogon

You can configure the Secondary Logon service by using the **runas** command. For more information, see Runas.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Security Accounts Manager**
Short name: SamSs

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Server**
Short name: lanmanserver

When you stop the server service, all users who are connected to the computer over the network are disconnected. Also, the affected computer cannot be administered remotely; therefore, you must start the server service locally. When you pause the server service, only members of the Administrators group and Server Operators group can make connections to the server. This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Shell Hardware Detection**
Short name: ShellHWDetection

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Simple Mail Transfer Protocol (SMTP)**
Short name: SMTPSVC

This service in only installed with a typical installation on Windows Server 2003, Web Edition. For other editions, you can install the service by using the Configure Your Server Wizard or by using Add or Remove Programs in Control Panel.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Smart Card**
Short name: SCardSvr

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local Service |
| Windows Server 2003, Enterprise Edition | Manual | Local Service |
| Windows Server 2003, Datacenter Edition | Manual | Local Service |
| Windows Server 2003, Web Edition | Manual | Local Service |

**Special Administration Console Helper**
Short name: Sacsvr

This service runs only when needed, and then it stops automatically.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |

| | | |
|---|---|---|
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**SQLAGENT$UDDI**
Short name: SQLSERVERAGENT

By default, this service is not installed with a typical installation. You can install it by using the Configure Your Server Wizard.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Not installed | Local System |
| Windows Server 2003, Enterprise Edition | Not installed | Local System |
| Windows Server 2003, Datacenter Edition | Not installed | Local System |
| Windows Server 2003, Web Edition | Not installed | Local System |

**System Event Notification**
Short name: SENS

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Task Scheduler**
Short name: Schedule

When you pause the Task Scheduler service, only members of the Administrators group on the local computer can schedule or modify tasks, and only those tasks that are scheduled by members of the Administrators group or the system account will run.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**TCP/IP NetBIOS Helper**
Short name: LmHosts

This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local Service |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local Service |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local Service |

| | | |
|---|---|---|
| Windows Server 2003, Web Edition | Automatic/Started | Local Service |

**Telephony**
Short name: TapiSrv

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Telnet**
Short name: TlntSvr

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Manual**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Terminal Services**
Short name: TermService

This service, although it is set to Manual, will start if Remote Desktop for Administration is enabled. For more information, see Remote Desktop for Administration.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual/Started | Local System |
| Windows Server 2003, Enterprise Edition | Manual/Started | Local System |
| Windows Server 2003, Datacenter Edition | Manual/Started | Local System |
| Windows Server 2003, Web Edition | Manual/Started | Local System |

**Terminal Services Session Directory**
Short name: Tssdis

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Manual**. This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Themes**
Short name: Themes

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Automatic**. This service is not available on the Itanium-based editions of the Windows Server 2003 family.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Uninterruptible Power Supply**
Short name: UPS

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local Service |
| Windows Server 2003, Enterprise Edition | Manual | Local Service |
| Windows Server 2003, Datacenter Edition | Manual | Local Service |
| Windows Server 2003, Web Edition | Manual | Local Service |

**Upload Manager**
Short name: uploadmgr

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Automatic**. This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Virtual Disk Service**
Short name: VDS

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Volume Shadow Copy**
Short name: VSS

| Product | Startup type | Log on as |
|---|---|---|

| | | |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**WebClient**
Short name: WebClient

This service is now disabled by default. If you want to enable this service, set **Startup type** to **Automatic**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local Service |
| Windows Server 2003, Enterprise Edition | Disabled | Local Service |
| Windows Server 2003, Datacenter Edition | Disabled | Local Service |
| Windows Server 2003, Web Edition | Disabled | Local Service |

**Web Element Manager**
Short name: ElementMgr

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Windows Audio**
Short name: AudioSrv

By default, this service is started only on computers that are running Windows Server 2003, Standard Edition. If sound is required, start this service and set **Startup type** to **Automatic**.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Windows Firewall/Internet Connection Sharing (ICS)**
Short name: SharedAccess

Windows Firewall and the Windows Firewall/Internet Connection Sharing (ICS) service were not included in the original release of Windows Server 2003, but are included with Windows Server 2003 Service Pack 1 (SP1).

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition with SP1 | Disabled | Local System |
| Windows Server 2003, Enterprise Edition with SP1 | Disabled | Local System |
| Windows Server 2003, Datacenter Edition with SP1 | Disabled | Local System |
| Windows Server 2003, Web Edition with SP1 | Disabled | Local System |

**Windows Image Acquisition (WIA)**
Short name: stisvc

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Disabled | Local System |
| Windows Server 2003, Enterprise Edition | Disabled | Local System |
| Windows Server 2003, Datacenter Edition | Disabled | Local System |
| Windows Server 2003, Web Edition | Disabled | Local System |

**Windows Installer**
Short name: MSIserver

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual/Started | Local System |
| Windows Server 2003, Enterprise Edition | Manual/Started | Local System |
| Windows Server 2003, Datacenter Edition | Manual/Started | Local System |
| Windows Server 2003, Web Edition | Manual/Started | Local System |

**Windows Management Instrumentation**
Short name: winmgmt

This service is started in Safe Mode, Safe Mode with Networking, and Safe Mode with Command Prompt.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**Windows Management Instrumentation Driver Extensions**
Short name: wmi

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Windows Media Services**
Short name: WMserver

You can install this service by using the Configure Your Server Wizard or by using Add or Remove Programs in Control Panel.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Not installed | Network Service |

| | | |
|---|---|---|
| Windows Server 2003, Enterprise Edition | Not installed | Network Service |
| Windows Server 2003, Datacenter Edition | Not installed | Network Service |
| Windows Server 2003, Web Edition | Not installed | Network Service |

**Windows Time**
Short name: W32time

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic | Local System |
| Windows Server 2003, Enterprise Edition | Automatic | Local System |
| Windows Server 2003, Datacenter Edition | Automatic | Local System |
| Windows Server 2003, Web Edition | Automatic | Local System |

**WinHTTP Web Proxy Auto-Discovery Service**
Short name: WinHttpAutoProxySvc

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Wireless Configuration**
Short name: WZCSVC

This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**WMI Performance Adapter**
Short name: WmiApSrv

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Manual | Local System |
| Windows Server 2003, Enterprise Edition | Manual | Local System |
| Windows Server 2003, Datacenter Edition | Manual | Local System |
| Windows Server 2003, Web Edition | Manual | Local System |

**Workstation**

Short name: lanmanworkstation

This service is started in Safe Mode with Networking.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Standard Edition | Automatic/Started | Local System |
| Windows Server 2003, Enterprise Edition | Automatic/Started | Local System |
| Windows Server 2003, Datacenter Edition | Automatic/Started | Local System |
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

**World Wide Web Publishing**
Short name: W3SVC

You can install the service by using the Configure Your Server Wizard or by using Add or Remove Programs in Control Panel.

| Product | Startup type | Log on as |
|---|---|---|
| Windows Server 2003, Web Edition | Automatic/Started | Local System |

# Troubleshooting Services Snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The service will not start.
- The computer will not start.
- When managing services on a remote computer, I receive a Connection Manager error message when I try to access the properties dialog box for a service.

**The service will not start.**
**Cause:** The account under which the service runs (as indicated on the **Log On** tab of the properties dialog box for the service) has changed.

**Solution:** Examine the System, Application, and Security logs in Event Viewer for "Service failed to start" error messages to identify which account the service uses. For more information about how to use Event Viewer to do this, see Search for specific events. Then, restore the default settings. For a table that lists the default settings, see Default settings for services. If you cannot access the Services snap-in to restore the settings, you can also use the **sc** command. For more information, see SC.

**Cause:** The system configuration has changed.

**Solution:** Use the System Configuration Utility (Msconfig.exe) to restart the computer using a diagnostic startup method. With the System Configuration Utility, you can select which devices and services are loaded at startup, which can help identify problem devices. For more information, see Windows interface administrative tool reference A-Z: System Configuration Utility.

**Cause:** Account lockout policy is enabled, or the password is expired.

**Solution:** If the service account is locked out or the service account password is expired, the service will not start. For more information about how to configure account lockout policy settings, see Apply or modify account lockout policy. For more information about how to apply password policy, see Apply or modify password policy.

**See also:** Account Lockout Policy.

**The computer will not start.**
**Cause:** The account under which the service runs (as indicated on the **Log On** tab of the properties dialog box for the service) has changed.

**Solution:** Start the computer in Safe Mode. In Safe Mode, you can change the service configuration or restore the default configuration. In Safe Mode, core services that are required to start the operating system are started in a default scheme, regardless of any changes that are made to the service settings. For more information, see Startup options.

**When managing services on a remote computer, I receive a Connection Manager error message when I try to access the properties dialog box for a service.**
**Cause:** The Remote Registry service is not running on the remote computer.

**Solution:** Ensure that the Remote Registry service is running on the remote computer. This is necessary only if you want to configure hardware profile settings for the service. For more information, see Start, stop, pause, resume, or restart a service. To start the service, you need the appropriate permissions on the remote computer.

**See also:** Troubleshooting Performance monitoring for System Monitor and Performance Logs and Alerts.

**Note**

- You may also receive this error message if the remote computer is running Windows 95 or Windows NT 4.0.

For information on enabling the services required by audio, imaging, and CD-burning peripheral devices, see Troubleshooting audio, imaging, and CD-burning peripheral devices.

# Event Viewer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Event Viewer**
With Event Viewer, you can monitor events recorded in event logs. Typically a computer stores the Application, Security, and System logs. It could also contain other logs, depending on the computer's role and the applications installed.

- For tips about using Event Viewer, see Event Viewer Best practices.
- For help with specific tasks, see Event Viewer How To....
- For general background information, see Event Viewer Concepts.
- For problem-solving instructions, see Event Viewer Troubleshooting.

# Event Viewer Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Enable security logging for security-critical objects and actions, and then periodically review the security logs on computers in your network.**

- By setting up an auditing policy and reviewing your security logs, you can detect unauthorized activity or intrusion attempts by malicious users or attackers on computers in your network.
- For more information about setting auditing policies, see Auditing Security Events.
- For more information about viewing event logs, see View Event Logs.

# Event Viewer How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- View Event Logs
- Manage Event Logs
- Customize Event Logs
- Use the Security Log

# View Event Logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**View event logs**

- View an event log
- Add another view of an event log
- Search for specific events
- Connect to another computer in Event Viewer
- Specify a sort order in an event log
- Filter events

# View an event log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view an event log**

1. Open Event Viewer.
2. In the console tree, click the log you want to view.
3. In the details pane, view the list of individual events.
4. If you want to see more details about a specific event, in the details pane, double-click the event.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.. However, you must be a member of the Administrators group on the target computer to open the Security log, or any log when targeting a remote computer.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- To refresh the view, on the **Action** menu, click **Refresh**. When you open a log, Event Viewer displays the current information for the log. While you view the log, the display is not updated unless you refresh the Event Viewer window. If you switch to another log and then return to the first log, the display for the first log is automatically refreshed. When you refresh the view of an event log, only the view is updated, no changes are made to the log. The **Refresh** command is not available for archived logs, because those files can no longer be updated.
- For more information about the event header in the details pane, click "The event header" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
The event header
Add another view of an event log
Search for specific events
Connect to another computer in Event Viewer
Event Viewer overview

# Add another view of an event log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add another view of an event log**

1. Open Event Viewer.
2. In the console tree, click the log of which you want to add another view.
3. On the **Action** menu, click **New Log View**. A copy of the log appears in the console tree with a default name. For example, if you selected the **Application** log, the default name of the copy would be **Application (2)**.
4. If you want to change the name of the new log view, on the **Action** menu, click **Rename**.
5. Type the name as you want it to appear in the console tree, and then press ENTER.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.. However, you must be a member of the Administrators group on the target computer to open the Security log, or any log when targeting a remote computer.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- By creating a new log view, you can customize how Event Viewer displays the information of an existing log and then apply a new name to the new log view, while keeping the defaults or the options you set for the original log intact. Log views added to the console tree can be managed and customized in the same way as the default logs.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Customize Event Logs
Event Viewer overview

# Search for specific events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To search for specific events**

1. Open Event Viewer.
2. In the console tree, click the log you want to search.
3. On the **View** menu, click **Find**.
4. Under **Event types**, select the types of events you want to find.
5. In **Event source**, **Category**, **Event ID**, **User**, **Computer**, or **Description**, specify additional information about the event or events you want to find.
6. Click **Find Next**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials. However, you must be a member of the Administrators group on the target computer to open the Security log, or any log when targeting a remote computer.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- In **Description**, you can type any text that matches a portion of an event record description.
- To restore the default search criteria, click **Restore Defaults** before clicking **Find Next**.
- Your search parameters remain in **Find** throughout the current session. The default settings are restored the next time you start Event Viewer.
- If you are looking for groups of events instead of a small number of individual events, you can also filter the log. For more information, see "Filter events" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Filter events
Specify a sort order in an event log

# Connect to another computer in Event Viewer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to another computer in Event Viewer**

1. Open Event Viewer.
2. In the console tree, right click **Event Viewer (Local)**, and then click **Connect to another computer**
3. Click **Another computer**, type the name of the computer, and then click **OK**. You can also find another computer on the network by clicking **Browse**.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or have been delegated the appropriate authority, on the target computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- The other computer can be a workstation running Windows XP Home Edition, Windows XP Professional, Windows 2000 Professional, or Windows NT Workstation, or a server or domain controller running Windows 2000 Server, Windows NT Server, or a Windows Server 2003 family operating system.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View an event log
Event Viewer overview
Managing event logs from the command line
Managing event logs from the command line

# Specify a sort order in an event log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify a sort order in an event log**

1. Open Event Viewer.
2. In the console tree, click the log you want to sort.
3. Click the column heading you want to sort by.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.. However, you must be a member of the Administrators group on the target computer to open the Security log, or any log when targeting a remote computer.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- To reverse the sort order, click the column heading a second time.
- To sort chronologically, on the **View** menu, click **Newest First** or **Oldest First**. The default is **Newest First**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Reset an event log to default settings
Filter events
Search for specific events

# Filter events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To filter events**

1. Open Event Viewer.
2. In the console tree, select the log you want to filter.
3. On the **View** menu, click **Filter**.
4. On the **Filter** tab, specify the characteristics you want to use in the filter.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.. However, you must be a member of the Administrators group on the target computer to open the Security log, or any log when targeting a remote computer.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- Filtering changes only the view, and has no effect on the actual contents of the log. All events are logged continuously, whether the filter is active or not. If you archive a log from a filtered view, all records are saved, even if you select a text format or comma-delimited text format file.
- To return to the default criteria, click **Restore Defaults**.
- To turn off event filtering, on the **View** menu, click **All Records**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Search for specific events
Reset an event log to default settings
Specify a sort order in an event log

# Manage Event Logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage event logs**

- Clear an event log
- Archive an event log
- Open an archived event log

# Clear an event log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To clear an event log**

1. Open Event Viewer.
2. In the console tree, click the log you want to clear.
3. On the **Action** menu, click **Clear all Events**.
4. If you want to save the log before clearing it, click **Yes**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- You cannot clear archived logs; instead, delete the archived log file.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding event logging options
Archive an event log
Open an archived event log
Event Viewer overview
Set event logging options

# Archive an event log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To archive an event log**

1. Open Event Viewer.
2. In the console tree, right-click the log you want to archive, and then click **Save Log File As**.
3. In **File name**, enter a name for the archived log file.
4. In **Save as type**, select a file format, and then click **Save**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- If you archive a log in log-file format (.evt), you can reopen it in Event Viewer. If you archive a log in text (.txt) or comma-delimited (.csv) format, you can reopen the log in other programs, such as word processing or spreadsheet programs, but not with Event Viewer.
- You can only save event logs in log-file format on the computer on which the events occur; for example, when using Event Viewer to view event logs on a remote computer.
- To save a local event log in log-file format on another computer, in **File name**, type *remotecomputername**share*, and then type the path to and name for your archive file.
- When you archive a log file, the entire log is saved, regardless of filtering options. The sort order is not retained when logs are saved. However, you can filter and sort the archived log again when you open it in Event Viewer.
- Archiving has no effect on the current contents of the active log.
- To remove an archived log file from your system, delete the file in Windows Explorer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
File formats for archived logs
Clear an event log
Open an archived event log
Event Viewer overview

# Open an archived event log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open an archived event log**

1. Open Event Viewer.
2. On the **Action** menu, click **Open Log File**.
3. Click the file you want to open. You may need to search for the drive or folder that contains the document.
4. In **Log type**, select the type of log to be opened.
5. In **Display name**, enter the name as you want it to appear in the console tree, and then click **Open**.

## Notes

- Performing this task does not require you to have administrative credentials. However, you must be a member of the Administrators group on the target computer to open the Security log, or any log when targeting a remote computer.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- You can use Event Viewer to open an archived log in log-file format (.evt). If an archived log was saved in text (.txt) or comma-delimited (.csv) format, you can open the log using other programs, such as word processing or spreadsheet programs, but not Event Viewer.
- To remove an archived log file from your system, delete the file in Windows Explorer.
- For more information about event log settings, see "Understanding event logging options" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Archive an event log
Event Viewer overview
Clear an event log

# Customize Event Logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Customize event logs**

- Set event logging options
- Reset an event log to default settings
- Change the maximum event log size

# Set event logging options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set event logging options**

1. Open Event Viewer.
2. In the console tree, click the log you want to set options for.
3. On the **Action** menu, click **Properties**.
4. On the **General** tab, specify the options you want.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding event logging options
Reset an event log to default settings
Change the maximum event log size
Filter events
Specify a sort order in an event log
Event Viewer overview
Settings for Event Logs

# Reset an event log to default settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reset an event log to default settings**

1. Open Event Viewer.
2. In the console tree, click the log you want to reset.
3. On the **Action** menu, click **Properties**.
4. On the **General** tab, click **Restore Defaults**.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or have been delegated the appropriate authority, on the target computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- For more information about event log settings, see *Understanding event logging options* in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding event logging options
Change the maximum event log size
Set event logging options
Filter events
Specify a sort order in an event log
Event Viewer overview

# Change the maximum event log size

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the maximum event log size**

1. Open Event Viewer.
2. In the console tree, click the log you want to change.
3. On the **Action** menu, click **Properties**.
4. On the **General** tab, in **Maximum log size**, specify the new log size in kilobytes.
5. To put the new setting in effect, click **Clear Log**.

   If you want to retain the information currently in the log, click **Yes** when a message appears asking if you want to save the original log before clearing it, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- For more information about event log settings, see *Understanding event logging options* in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding event logging options
Reset an event log to default settings
Set event logging options
Filter events
Specify a sort order in an event log
Event Viewer overview

# Use the Security Log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use the security log**

- View the security log
- Define or modify auditing policy settings for an event category
- Apply or modify auditing policy settings for a local file or folder
- Apply or modify auditing policy settings for an object using Group Policy

# View the security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the security log**

1. Open Event Viewer.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. If you want to see more details about a specific event, in the details pane, double-click the event.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- For more information about security events, see "Security Events" at the Microsoft Windows Resource Kits Web site

**See Also**
**Concepts**
Auditing Security Events
Auditing Policy
Settings for Event Logs
Understanding event logging options

# Define or modify auditing policy settings for an event category

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define or modify auditing policy settings for an event category**

- For your local computer
- For only domain controllers, when you are on a domain controller or on a workstation that hasWindows Server 2003Administration Tools Pack installed
- For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed
- For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain

**For your local computer**

1. Open Local Security Settings.
2. In the console tree, click **Audit Policy**.

   **Where?**

   - Security Settings/Local Policies/Audit Policy

3. In the details pane, double-click an event category that you want to change the auditing policy settings for.
4. Do one or both of the following, and then click **OK**.

   - To audit successful attempts, select the **Success** check box.
   - To audit unsuccessful attempts, select the **Failure** check box.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**For only domain controllers, when you are on a domain controller or on a workstation that has Windows Server 2003 Administration Tools Pack installed**

1. Open Domain Controller Security Policy.
2. In the console tree, click **Audit Policy**.

   **Where?**

- Computer Configuration/Windows Settings/Security Settings/Local Policies/Audit Policy

3. In the details pane, double-click an event category that you want to change the auditing policy settings for.

4. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.

5. Do one or both of the following, and then click **OK**.

   - To audit successful attempts, select the **Success** check box.

   - To audit unsuccessful attempts, select the **Failure** check box.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Domain Controller Security Policy, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Domain Controller Security Policy**.

**For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

3. Click **Properties**, and then click the **Group Policy** tab.

4. Click **Edit** to open the Group Policy object (GPO) that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.

5. In the console tree, click **Audit Policy**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Local Policies/Audit Policy

6. In the details pane, double-click an event category that you want to change the auditing policy settings for.

7. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.

8. Do one or both of the following, and then click **OK**.

   - To audit successful attempts, select the **Success** check box.

   - To audit unsuccessful attempts, select the **Failure** check box.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain**

1. Open Microsoft Management Console (MMC).

2. In the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.

3. Click **Group Policy Object Editor**, and then click **Add**.

4. On the **Select Group Policy Object** page in the Group Policy Wizard, click **Browse**.

5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit--or create a new one, click **OK**, and then click **Finish**.

6. Click **Close**, and then click **OK**.

7. In the console tree, click **Audit Policy**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Local Policies/Audit Policy

8. In the details pane, double-click an event category that you want to change the auditing policy settings for.

9. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.

10. Do one or both of the following, and then click **OK**.

    - To audit successful attempts, select the **Success** check box.

    - To audit unsuccessful attempts, select the **Failure** check box.

    **Notes**

    - To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

    - To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.

**Notes**

- To audit object accesses, enable auditing of the object access event category by following the steps above. Then, enable auditing on the specific object. For information about how to enable auditing on an object, see "Apply or modify auditing

policy settings for a local file or folder" or "Apply or modify auditing policy settings for an object using Group Policy" in Related Topics.

- After your audit policy is configured, events will be recorded in the security log. Open the security log to view these events. For information about the security log, see "Use the security log" in Related Topics.

- The default auditing policy setting for domain controllers is **No Auditing**. This means that even if auditing is enabled in the domain, the domain controllers do not inherit auditing policy locally. If you want domain auditing policy to apply to domain controllers, you must modify this policy setting.

**See Also**
**Concepts**
Auditing Policy
Auditing Security Events
Auditing overview
Apply or modify auditing policy settings for a local file or folder
Apply or modify auditing policy settings for an object using Group Policy
Use the Security Log

# Apply or modify auditing policy settings for a local file or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To apply or modify auditing policy settings for a local file or folder**

1. Open Windows Explorer.
2. Right-click the file or folder that you want to audit, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, and then click the **Auditing** tab.
4. Do one of the following:
   - To set up auditing for a new user or group, click **Add**. In **Enter the object name to select**, type the name of the user or group that you want, and then click **OK**.
   - To remove auditing for an existing group or user, click the group or user name, click **Remove**, click **OK**, and then skip the rest of this procedure.
   - To view or change auditing for an existing group or user, click its name, and then click **Edit**.
5. In the **Apply onto** box, click the location where you want auditing to take place.
6. In the **Access** box, indicate what actions you want to audit by selecting the appropriate check boxes:
   - To audit successful events, select the **Successful** check box.
   - To stop auditing successful events, clear the **Successful** check box.
   - To audit unsuccessful events, select the **Failed** check box.
   - To stop auditing unsuccessful events, clear the **Failed** check box.
   - To stop auditing all events, click **Clear All**.
7. If you want to prevent subsequent files and subfolders of the original object from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.

**Important**

- Before setting up auditing for files and folders, you must enable object access auditing by defining auditing policy settings for the object access event category. If you do not enable object access auditing, you will receive an error message when you set up auditing for files and folders, and no files or folders will be audited. For more information about how to enable object access auditing, see "Define or modify auditing policy settings for an event category" in Related Topics.

**Notes**

- You must be logged on as a member of the Administrators group or you must have been granted the **Manage auditing**

**and security log** right in Group Policy to perform this procedure.

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- For information about how to audit local registry keys, see "Audit activity on a registry key" in Related Topics.
- After object access auditing is enabled, view the security log in Event Viewer to review the results of your changes.
- You can set up file and folder auditing only on NTFS drives.
- If you see the following:
    - In the **Auditing Entry for *File or Folder*** dialog box, in the **Access** box, the check boxes are unavailable ...
    - In the **Advanced Security Settings for *File or Folder*** dialog box, the **Remove** button is unavailable ...

  auditing has been inherited from the parent folder.
- Because the security log is limited in size, select the files and folders to be audited carefully. Also, consider the amount of disk space that you want to devote to the security log. The maximum size for the security log is defined in Event Viewer.

**See Also**
**Concepts**
How inheritance affects file and folder auditing
Selecting where to apply auditing entries
Auditing Security Events
Apply or modify auditing policy settings for an object using Group Policy
Define or modify auditing policy settings for an event category
Audit activity on a registry key

# Apply or modify auditing policy settings for an object using Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To apply or modify auditing policy settings for an object using Group Policy**

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. On the **Select Group Policy Object** page in the Group Policy Wizard, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit--or create a new one, click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. Do one or more of the following:

| To audit | Do this |
|---|---|
| System services | In the console tree, click **System Services**.<br>Where?<br>• Computer Configuration/Windows Settings/Security Settings/System Services<br>In the details pane, right-click the service that you want to apply or modify auditing policy settings for, and then click **Properties**.<br>If it is not already selected, select the **Define this policy setting** check box, and then select the appropriate setting.<br>Click **Edit security**. |
| Registry keys | In the console tree, click **Registry**.<br>Where?<br>• Computer Configuration/Windows Settings/Security Settings/Registry<br>If you want to add a registry key to this GPO to audit, right-click **Registry**, and then click **Add Key**. Browse to the key that you want, and then click **OK**.<br>If you want to apply or modify auditing settings on a registry key that has already been added to this GPO, in the details pane, right-click the registry key, click **Properties**, and then click **Edit Security**. |

| | |
|---|---|
| Files or folders | In the console tree, click **File System**.<br><br>Where?<br><br>• Computer Configuration/Windows Settings/Security Settings/File System<br><br>If you want to add a file or folder to this GPO to audit, right-click **File System**, and then click **Add File**. Browse to the file that you want, or make a new folder, and then click **OK**.<br><br>If you want to apply or modify auditing settings on a file or folder that has already been added to this GPO, in the details pane, right-click the file or folder, click **Properties**, and then click **Edit Security**. |

8. Click **Advanced**, and then click the **Auditing** tab.

9. Do one of the following:

   - To set up auditing for a new user or group, click **Add**. In **Name**, type the name of the user or group that you want, and then click **OK**.
   - To view or change auditing for an existing group or user, click the name that you want, and then click **Edit**.
   - To remove auditing for an existing group or user, click the name that you want, click **Remove**, click **OK**, and then skip the rest of this procedure.

10. Select the appropriate entry in the **Apply onto** list.

11. In the **Access** box, indicate what actions you want to audit by selecting the appropriate check boxes:

    - To audit successful events, select the **Successful** check box.
    - To stop auditing successful events, clear the **Successful** check box.
    - To audit unsuccessful events, select the **Failed** check box.
    - To stop auditing unsuccessful events, clear the **Failed** check box.
    - To stop auditing all events, click **Clear All**.

12. If you want to prevent files and subfolders in the tree from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.

**Important**

- Before setting up auditing for files and folders, you must enable object access auditing by defining auditing policy settings for the object access event category. If you do not enable object access auditing, you will receive an error message when you set up auditing for files and folders, and no files or folders will be audited. For more information about how to enable object access auditing, see "Define or modify auditing policy settings for an event category" in Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- For more information on selecting where to apply auditing entries, see Related Topics.
- You can set up file and folder auditing only on NTFS drives.
- If you see the following:
    - In the **Auditing Entry for *File or Folder*** dialog box, in the **Access** box, the check boxes are unavailable ...
    - In the **Advanced Security Settings for *File or Folder*** dialog box, the **Remove** button is unavailable ...

  auditing is inherited from the parent folder.
- After object access auditing is enabled, view the security log in Event Viewer to review the results of your changes.
- Because the security log is limited in size, select the files and folders to be audited carefully. Also, consider the amount of disk space that you want to devote to the security log. The maximum size for the security log is defined in Event Viewer.

**See Also**
**Concepts**
Auditing Security Events
View the security log
Define or modify auditing policy settings for an event category
Apply or modify auditing policy settings for a local file or folder
Apply or modify auditing policy settings for an object using Group Policy
Selecting where to apply auditing entries

# Event Viewer Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Event Viewer:

- Event Viewer overview
- Auditing Security Events
- Event log overview
- Using the Event Viewer snap-in
- Managing event logs from the command line

# Event Viewer overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Event Viewer overview**
Using Event Viewer, you can view and set logging options for event logs in order to gather information about hardware, software, and system problems.

- For information about event logs, see Event log overview.
- For information about events, see The event header or Types of events.
- For information about event logging options you can set using Event Viewer, see Understanding event logging options.
- For information about setting up auditing policies, see Auditing Security Events.

# Auditing Security Events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Auditing security events**
You can set up audit policy so that user or system activity in specified event categories is recorded. You can monitor security-related activity, such as who accesses an object, if a user logs on to or logs off from a computer, or if changes are made to an auditing policy setting.

- Before defining auditing policy settings in the object access event category, see Checklist: Setting up object access auditing.
- For tips about auditing, see Auditing Security Events Best practices.
- For help with specific tasks, see Auditing Security Events How To....
- For general background information, see Auditing Security Events Concepts.

# Event log overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Event log overview**
By default, a computer running a Windows Server 2003 family operating system records events in three kinds of logs:

**Application log**
The application log contains events logged by applications or programs. For example, a database program might record a file error in the application log. Application developers decide which events to log.

**Security log**
The security log records events such as valid and invalid logon attempts, as well as events related to resource use such as creating, opening, or deleting files or other objects. For example, if logon auditing is enabled, attempts to log on to the system are recorded in the security log.

**System log**
The system log contains events logged by Windows system components. For example, the failure of a driver or other system component to load during startup is recorded in the system log. The event types logged by system components are predetermined by the server.

A computer running a Windows Server 2003 family operating system which is configured as a domain controller records events in two additional logs:

**Directory service log**
The directory service log contains events logged by the Windows Active Directory service. For example, connection problems between the server and the global catalog are recorded in the directory service log.

**File Replication service log**
The File Replication service log contains events logged by the Windows File Replication service. For example, file replication failures and events that occur while domain controllers are being updated with information about system volume changes are recorded in the file replication log.

A computer running Windows configured as a Domain Name System (DNS) server records events in an additional log:

**DNS server log**
The DNS server log contains events logged by the Windows DNS service.

Other types of events and event logs might be available on a computer, depending on what services are installed.

The Event Log service starts automatically when you start Windows.

If you are a member of the Administrators group on the local computer, you can set access permissions on event logs using Group Policy. For more information about security groups, see Default groups. For more information, see Settings for Event Logs.

For more information about viewing events, see Types of events and Interpreting an Event.

# Using the Event Viewer snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Event Viewer**

This section covers:

- Interpreting an Event
- Viewing and Archiving Log Files

# Interpreting an Event

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interpreting an event**

This section covers:

- The event header
- Types of events
- Linking to Microsoft for Help and Support

# The event header

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The event header**

The event header contains the following information:

| Information | Meaning |
|---|---|
| Date | The date the event occurred. The date and time of an event are stored in Universal Time Coordinate (UTC) but always display in the viewer's locale. |
| Time | The time the event occurred. The date and time of an event are stored in UTC but always display in the viewer's locale. |
| User | The name of the user on whose behalf the event occurred. This name is the client ID if the event was actually caused by a server process, or the primary ID if impersonation is not taking place. Where applicable, a security log entry contains both the primary and impersonation IDs. Impersonation occurs when the server allows one process to take on the security attributes of another. |
| Computer | The name of the computer where the event occurred. This is usually the name of your own computer, unless you are viewing an event log on another computer. |
| Source | The software that logged the event, which can be either a program name such as SQL Server, or a component of the system (such as a driver name) or of a large program. For example, "Elnkii" indicates an EtherLink II driver. The Source always remains in its original language. |
| Event | A number identifying the particular event type for this source. The first line of the description usually contains the name of the event type. For example, 6005 is the ID of the event that occurs when the Event log service is started. The first line of the description of such an event is "The Event log service was started." Using the values of Source and Event together, product support representatives can troubleshoot system problems. |
| Type | A classification of the event severity: Error, Information, or Warning in the system and application logs; Success Audit or Failure Audit in the security log. In the Event Viewer normal list view, these are represented by a symbol. |
| Category | A classification of the event by the event source. This information is primarily used in the security log. For example, for security audits, this corresponds to one of the event types for which success or failure auditing can be enabled in Group Policy by a member of the Administrators group. |

For more information, see

- Event log overview
- View an event log
- Types of events
- Understanding event logging options

# Types of events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Types of events**

Event Viewer displays five types of events:

| Event type | Description |
|---|---|
| Error | A significant problem, such as loss of data or loss of functionality. For example, if a service fails to load during startup, an Error will be logged. |
| Warning | An event that is not necessarily significant, but might indicate a possible future problem. For example, when disk space is low, a Warning might be logged. |
| Information | An event that describes the successful operation of an application, driver, or service. For example, when a network driver loads successfully, an Information event will be logged. |
| Success Audit | Any audited security event that succeeds. For example, a user's successful attempt to log on the system will be logged as a Success Audit event. |
| Failure Audit | Any audited security event that fails. For example, if a user tries to access a network drive and fails, the attempt will be logged as a Failure Audit event. |

For more information, see:

- Event log overview
- View an event log
- The event header
- Understanding event logging options

# Linking to Microsoft for Help and Support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Linking to Microsoft for Help and Support**
For more details about an event, you can double-click it in the details pane to open the **Event Properties** dialog box. The **Event Properties** dialog box contains a Uniform Resource Locator (URL) link in the **Description** dialog box.

If you click the URL in the **Event Properties** dialog box of an event generated by a Microsoft application or service, information related to the event display is sent over the Internet to Microsoft in the form of a query, which locates more detailed information from the Microsoft Web site. The collected data is limited to only what is needed to retrieve more information about the event. User names and e-mail addresses, names of files unrelated to the logged event, machine addresses, and any other forms of personally identifiable information are not collected.

### Notes

- Events not generated by Microsoft applications or services might also contain a URL, and might take you to non-Microsoft web sites.

- Administrators can use a Group Policy setting to prevent the URL link in the **Description** dialog box from appearing.

# Viewing and Archiving Log Files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing and archiving log files**
This section covers:

- Understanding event logging options
- File formats for archived logs

# Understanding event logging options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding event logging options**

Using Event Viewer, you can define logging parameters for each kind of event log. To define parameters, right-click a log in the console tree, and then click **Properties**. On the **General** tab, you can set the maximum size of the log and specify whether the events are overwritten or stored for a certain period of time.

The default logging policy is that if a log is full, the oldest events are deleted to make room for new events, provided events are at least seven days old. You can customize this policy, or event log wrapping options, for different logs.

Event log wrapping options include the following.

| Use | To |
|---|---|
| **Overwrite events as needed** | Have new events continue to be written when the log is full. Each new event replaces the oldest event in the log. This option is a good choice for low-maintenance systems. |
| **Overwrite events older than** [*x*] **days** | Retain the log for the number of days you specify before overwriting events. The default is seven days. This option is the best choice if you want to archive log files weekly. This strategy minimizes the chance of losing important log entries and at the same time keeps log sizes reasonable. |
| **Do not overwrite events** | Clear or archive the log manually rather than automatically. Select this option only if you cannot afford to miss an event (for example, for the security log at a site where security is extremely important). |

You can also use Group Policy to set the maximum log sizes and log wrapping options, as well as set access permissions on event logs. For more information, see Settings for Event Logs.

Application and system logging start automatically when you start the computer. For information about configuring security logging, see Auditing Security Events.

**Notes**

- To perform these procedures, you must be a member of the Administrators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform these procedures. For more information, see Default local groups, Default groups, and Using Run as.

- For information about how to set event logging options using Event Viewer, see Set event logging options.

# File formats for archived logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File formats for archived logs**

When you archive an event log, you save it in one of three file formats:

- **Log-file format** (.evt), which enables you to view the archived log again in Event Viewer.
- **Text-file format** (.txt), which enables you to use the information in a program such as a word processing program. You cannot use Event Viewer to open a log file saved in .txt format.
- **Comma-delimited text-file format** (.csv), which enables you to use the information in a program such as a spreadsheet or a flat-file database. You cannot use Event Viewer to open a log file saved in .csv format.

The event description is saved in all archived logs. The data within each individual event record is recorded in the following order:

1. Date
2. Time
3. Source
4. Type
5. Category
6. Event
7. User
8. Computer
9. Description

For more information, see:

- Archive an event log
- Open an archived event log
- Event log overview
- Types of events

# Managing event logs from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing event logs from the command line**

**Opening Event Viewer from the command line**
You can launch Event Viewer from the command-line using `eventvwr.exe` or `mmc.exe`. Either command opens Event Viewer.

**Syntax for eventvwr.exe**

```
eventvwr.exe
```

**Parameters**

```
computername
```

Specifies the computer you want to target, where `computername` is the name of the target computer. If you do not specify a computer, Event Viewer displays the logs on the local computer. You can specify a remote computer for `computername`, if you are a member of the Administrators group, or have been delegated the appropriate authority, on the remote computer.

**or**

**Syntax for mmc.exe**

```
mmc.exe X:\WINDOWS\system32\eventvwr.msc
```

Where `X:\WINDOWS\System32` is the default path to `eventvwr.msc`, and `X` is the folder where Windows is installed.

**Parameters**

```
/computer=computername
```

Specifies the computer you want to target, where `computername` is the name of the target computer. If you do not specify a computer, Event Viewer displays the logs on the local computer. You can specify a remote computer for `computername`, if you are a member of the Administrators group, or have been delegated the appropriate authority, on the remote computer.

```
/auxsource=computername
```

Enables Event Viewer to obtain the information about events from a computer on which the event log service is installed. This can be the computer that generated the log or another computer on the network. You can use the `/auxsource=computername` parameter on computers running a Windows XP or later operating system. To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**Using command-line tools to manage events and event logs**
You can also use command-line utilities to create and query event logs and associate programs with particular logged events. For example, you can use Eventcreate to customize an event entry to a specified event log. Eventquery.vbs is used to list the events and event properties from one or more event logs. Eventtriggers enables you to create event triggers that will run programs upon the occurrence of specific events.

For information about other command-line utilities, see Command-line reference A-Z.

# Event Viewer Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- When viewing an event in a saved event log on a remote computer, you are unable to view the Description or Category fields in the properties dialog box. Instead of a description in the Description field, you see an error message, or instead of a category in the Category field, you see a number.
- When viewing an event in the details pane, you see a globally unique identifier (GUID) or security ID (SID) in the User field. You should see a user name.
- You are unable to view the Security log, but can view other logs.
- You are able to view the Security log, but are unable to view other logs.
- You cannot access the Computer Management extension snap-ins on a remote computer.

**When viewing an event in a saved event log on a remote computer, you are unable to view the Description or Category fields in the properties dialog box. Instead of a description in the Description field, you see an error message, or instead of a category in the Category field, you see a number.**
**Cause:** Event Viewer obtains event description and category information from the computer that contains the log. The event log might have been copied from the computer on which it was generated, to another computer, on which the service that logged these events is not installed.

**Solution:** Open Event Viewer from the command-line using **mmc.exe**, and use the **/auxsource=***computername* parameter. The **/auxsource=***computername* parameter enables Event Viewer to obtain the information about events from a computer on which the event log service is installed. This can be the computer which generated the log or another computer on the network. You can use the **/auxsource=***computername* parameter on computers running a Windows XP or later operating system. To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**Syntax**

**mmc.exe X:\WINDOWS\system32\eventvwr.msc /auxsource=** *computername*

Where X:\WINDOWS\System32 is the default path to eventvwr.msc, and X is the drive where Windows is installed.

**Note**

- To run a command-line utility, click **Start**, and then click **Run**. In the **Open** box, enter the command-line utility and parameters, and then click **OK**.
- If category information is not available, you will also see a number instead of text in the **Category** field in the Event Viewer details pane for an event listing.

**See also:** Mmc; Managing event logs from the command line

**Cause:** You are attempting to view a log on a remote computer and are not a member of the Administrators group, or have not been delegated the appropriate authority on the remote computer.

**Solution:** To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

You might also be able to obtain the same information from a different computer on your network. To do this, run Event

Viewer at a command prompt (Eventvwr.msc) with the **auxsource=***ComputerName* parameter. You can use this parameter to enable Event Viewer to obtain information from a computer on which the service is installed. This can be either the computer that generated the log or another computer on the network.

**Cause:** You are attempting to view a log on a remote computer and the Remote Registry service on the remote computer is stopped. If the remote computer is running a Windows XP or later operating system, the Remote Registry service must be started for you to be able to see the **Description** or **Category** fields in the property page for an event log.

**Solution:** On the remote computer, start the Remote Registry Service. You can manage the Remote Registry service through the **Services and Applications** snap-in in **Computer Management**.

To start the Remote Registry service

1. In **Computer Management**, verify that you are in the console for the remote computer.

2. In the console tree, under **Services and Applications**, click **Services**.

3. In the details pane, double-click **Remote Registry**.

4. In the **General** tab, under **Service status**, click **Start**. Click **OK**.

You might also be able to obtain the same information from a different computer on your network. To do this, run Event Viewer at a command prompt (Eventvwr.msc) with the **auxsource=***ComputerName* parameter. You can use this parameter to enable Event Viewer to obtain information from a computer on which the service is installed. This can be either the computer that generated the log or another computer on the network.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- This parameter is available on computers running Windows XP or a product in the Windows Server 2003 family.

**See also:** Remote Administration;Computer Management; Microsoft Management Console

**When viewing an event in the details pane, you see a globally unique identifier (GUID) or security ID (SID) in the User field. You should see a user name.**
**Cause:** The computer on which you are running Event Viewer is unable to translate the GUID or SID into the user name.

**Solution:** Run Event Viewer on the computer on which the event log was generated.

**Cause:** The user account referenced by the event has been deleted.

**Solution:** You will not be able to retrieve this information through Event Viewer.

**You are unable to view the Security log, but can view other logs.**
**Cause:** You are not a member of the Administrators group, or have not been delegated the appropriate authority on the target computer.

**Solution:** To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**See also:** Use the Security Log

**You are able to view the Security log, but are unable to view other logs.**
**Cause:** Your account has been added to the Guests group on the target computer, and is also a member of the Administrators group.

**Solution:** Remove your account from the **Guests** group on the target computer:

1. Open **Computer Management**. If the target computer is a remote computer, right-click **Computer Management (Local)**, and then click **Connect to another computer**.

2. Under **Local Users and Groups**, click **Groups**.

3. In the details pane, double-click **Guests**.

4. Select the account, and then click **Remove**.

5. Click **OK**.

### Notes

- This problem typically develops when an administrator adds a group containing a broad category of users (such as the **Everyone**, **INTERACTIVE**, OR **Authenticated users** group) to the **Guests** group.

- By default, members of the **Guests** group are explicitly denied access to resources such as event logs. Therefore, avoid adding groups containing users with administrative privileges to the **Guests** group.

- You might need to log off, and then log back on for these changes to take effect.

**See also:** Connect to another computer in Event Viewer.

**You cannot access the Computer Management extension snap-ins on a remote computer.**
**Cause:** The Remote Registry service is not running on the remote computer.

**Solution:** Ensure that the Remote Registry service is started on the remote computer. For more information, see Start, stop, pause, resume, or restart a service. You need appropriate permissions on the remote computer to start the service.

### Note

- You might also receive this error message if the remote computer is running Windows 95. Computer Management does not support remote access to computers that are running Windows 95.

# Network Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Monitor**
Network administrators use Network Monitor to view and detect problems on local area networks (LANs).

- Before using Network Monitor, see Checklist: Monitoring network traffic on your local computer.
- To find features that have been moved in Microsoft® Windows Server 2003 family operating systems, see New ways to do familiar Network Monitor tasks.
- For tips about using Network Monitor, see Network Monitor Best practices.
- For help with specific tasks, see Network Monitor How To ....
- For general background information, see Network Monitor Concepts.
- For problem-solving instructions, see Troubleshooting Network Monitor.

# Checklist: Monitoring network traffic on your local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Monitoring network traffic on your local computer**

| Step | Reference |
|---|---|
| Review key concepts. | Network Monitor overview;<br><br>Network Monitor Best practices |
| Confirm that you are a member of the Administrators group. | Default local groups;<br><br>Edit local security settings |
| Install Network Monitor tools. | Install Network Monitor |
| (Optional) If you want to make sure your captures do not exceed available memory on your computer, then modify your buffer settings. | Modify capture buffer settings |
| (Optional) If you want to exclude traffic from a certain network address from being captured, then create a capture filter. | Capture filters;<br><br>Design a capture filter |
| (Optional) If you want to trigger a capture when a specific string appears in a frame, then create a capture trigger. | Set a capture trigger |
| (Optional) If you want to specify certain protocols or addresses that you want to save to file, then create a display filter. | Design a display filter;Load a display filter |

# New ways to do familiar Network Monitor tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring Network Monitor. The user interface for performing these tasks is almost identical in this version of Windows to the way it was in Windows NT version 4.0.

### Note

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Install Network Monitor | Network in Control Panel. | Open Windows Components wizard.<br><br>For more information, see Install Network Monitor. |
| Install the Network Monitor driver | Network in Control Panel. | Open Network Connections; go to **Local Area Connection Properties**.<br><br>For more information, see Install the Network Monitor driver. |

# Network Monitor Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Running Network Monitor during low-usage times**

  Run Network Monitor at low-usage times or for short periods of time. This decreases the effect that Network Monitor might have on the performance of your computer.

- **Capturing minimum amount of network statistics**

  Use capture filters and capture only as many statistics as you need for evaluation. This prevents you from capturing too much information to make a reasonably quick diagnosis of the problem.

- **Viewing Network Monitor captures**

  Use display filters when viewing captured frames to omit extraneous captured frames.

- **Capturing frames on token ring networks**

  Token ring adapters reject frames larger than a size specified in their configuration settings. If you capture frames on token ring networks, set your network adapter to accept the largest possible token-ring frame size, which is 17 KB. For more information, see the documentation for your network adapter.

# Network Monitor How To ...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install and Configure Network Monitor
- Capture Network Data
- Display Network Data
- Print Network Data

# Install and Configure Network Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and configure Network Monitor**

- Install Network Monitor
- Modify capture buffer settings
- Install the Network Monitor driver

# Install Network Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Network Monitor**

1. Open Windows Components wizard.
2. In the Windows Components wizard, click **Management and Monitoring Tools**, and then click **Details**.
3. In **Subcomponents of Management and Monitoring Tools**, select the **Network Monitor Tools** check box, and then click **OK**.
4. If you are prompted for additional files, insert the installation CD for your operating system, or type a path to the location of the files on the network.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.
- This procedure automatically installs the Network Monitor driver.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the Network Monitor driver
Capture network frames

# Modify capture buffer settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify capture buffer settings**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Buffer Settings**.
4. Specify the maximum size of the capture buffer. If necessary, also adjust the maximum size of frames to be captured, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- If your buffer setting exceeds the available memory on your computer, frames might be dropped.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
The capture process
Design a capture filter
Save captured frames to a file

# Install the Network Monitor driver

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install the Network Monitor driver**

1. Open Network Connections.
2. In **Network Connections**, click **Local Area Connection**, click the **File** menu, and then click **Properties**.
3. In the **Local Area Connection Properties** dialog box, click **Install**.
4. In the **Select Network Component Type** dialog box, click **Protocol**, and then click **Add**.
5. In the **Select Network Protocol** dialog box, click **Network Monitor Driver**, and then click **OK**.
6. If you are prompted for additional files, insert the installation CD for your operating system, or type a path to the location of the files on the network.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- **Network Monitor Driver** does not appear in the **Select Network Protocol** dialog box if the Network Monitor driver is already installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Network Monitor
Network Monitor overview
Network Monitor How To ...

# Capture Network Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture network data**

- Capture network frames
- Save captured frames to a file
- Set a capture trigger
- Capture data in dedicated capture mode
- Identify the largest broadcaster on the network
- Use a Capture Filter

# Capture network frames

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To capture network frames**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Buffer Settings**, and then set the buffer and frame size as appropriate.
4. On the **Capture** menu, click **Start**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- If you add a network interface, you must restart Network Monitor before you can capture data on that interface.
- To pause, stop, or display the data capture, on the **Capture** menu, click **Pause**, **Stop**, or **Display Captured Data**. You can also stop and display the capture by clicking **Stop and View**.
- Network Monitor displays session statistics of the first 100 unique network sessions that it detects. To reset statistics and view information about the next 100 detected network sessions, click the **Capture** menu, and then click **Clear Statistics**.
- If you want to capture data from multiple local networks at the same time, install an adapter and start an instance of Network Monitor for each network. In each instance, specify the network from which you want to capture data by, on the **Capture** menu, clicking **Networks** and selecting a network.
- If prompted for additional files, insert the installation CD for your operating system, or type a path to the location of the files on the network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Network Monitor
Network Monitor Best practices

# Save captured frames to a file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save captured frames to a file**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Buffer Settings**, and then set the buffer and frame size as appropriate.
4. On the **Capture** menu, click **Start**.
5. To halt the data capture temporarily, on the **Capture** menu, click **Pause**.
6. To stop and view the data capture, on the **Capture** menu, click **Stop and View**.
7. On the **File** menu, click **Save As**.
8. Open the folder in which you want to store the file.
9. In the **File Name** box, specify a file name.
10. If necessary, do one of the following:
    - To save a range of frames, in the **From** box, type the beginning frame number, and in the **To** box, type the ending frame number.
    - To save only the frames that appear when the current display filter is in use, select the **Filtered** check box.
11. Click **Save**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- If you add a network interface, you must restart Network Monitor before you can capture data on that interface.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- To change the default network, on the **Capture** menu, click **Networks**, and then choose the network on which you are trying to capture data.
- If you want to use Network Monitor to load the file you saved, use the .cap file extension. Network Monitor automatically uses this file extension if you do not specify one.

- Network Monitor cannot open capture files larger than 1 GB (1,024 MB).

- When you save a range of frames or apply a display filter and then save the filtered frames, Network Monitor renumbers the saved frames, starting with 1. However, you can create a file with the original frame numbers preserved. To do this, on the **File** menu, click **Print**, click the **General** tab, and then select the **Print to file** check box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Capture network frames
Set, view, change, or remove permissions on files and folders
Design a capture filter

# Set a capture trigger

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set a capture trigger**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Trigger**.
4. Do one of the following to specify trigger criteria:
   - To initiate a trigger action when a specific ASCII or hexadecimal string appears in a frame, click **Pattern match**. In the **Pattern** box, type the string you want Network Monitor to detect, and then specify whether the pattern is in hexadecimal or ASCII. If you want, specify where Network Monitor should search for the pattern.
   - To initiate a trigger action when a specific percentage of the capture buffer is full, click **Buffer space**, and then specify the percentage needed.
   - To initiate a trigger action when Network Monitor detects a specific pattern in a frame after a specific percentage of the capture buffer becomes full, click **Buffer space then pattern match**, and specify the percentage and pattern needed.
   - To initiate a trigger action when a specific percentage of the capture buffer becomes full after Network Monitor detects a specific pattern in a frame, click **Pattern match then buffer space**, and specify the pattern and percentage needed.
   - To clear any capture triggers that have been set, click **Nothing**.
5. Do one of the following to specify a trigger action:
   - To have the computer beep, click **Audible Signal Only**.
   - To stop the capture, click **Stop Capture**.
   - To run a command or a program, click **Execute Command Line** and specify the command or program that runs when the conditions for the trigger are met. To specify a program, type the path and the name of the program file, or click **Browse** and navigate to the program file. To use an MS-DOS command, such as **copy**, type **CMD /K**, and then type the command.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- If you click **Pattern match then buffer space** and set **Buffer Space** to **100%**, Network Monitor starts monitoring the capture buffer size with the frame after the frame that contained the pattern match. This means that when the buffer is

100% full, the frame containing the pattern match has been overwritten. If you use the **Pattern match then buffer space** option and need to view the frame that contains the pattern match in order to analyze your network traffic issues, set **Buffer Space** to 75% or smaller.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
The capture process
Capture triggers

# Capture data in dedicated capture mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To capture data in dedicated capture mode**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Dedicated Capture Mode**.
4. On the **Capture** menu, click **Start**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- If you add a network interface, you must restart Network Monitor before you can capture data on that interface.
- During the capture, select options in the Dedicated Mode dialog box.
- In dedicated capture mode, Network Monitor does not display or refresh capture statistics when frames are copied to the temporary capture file. This frees more resources for capturing data.
- Use dedicated capture mode if Network Monitor drops frames due to a lack of resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Capture network frames
Save captured frames to a file
Design a capture filter

# Identify the largest broadcaster on the network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To identify the largest broadcaster on the network**

1. Open Network Monitor.
2. On the **Capture** menu, click **Start**.
3. When you finish capturing data, on the **Capture** menu, click **Stop** (or **Pause**).
4. In the Station Statistics pane (the lowest pane on the screen), double-click the **Broadcasts Sent** column header.

   Or, right-click in the **Broadcasts Sent** column and then click **Sort Column**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- If the Graph pane indicates a high percentage of network use and if network operations are unusually slow, an abrupt increase in broadcasts may be the source of the problem.
- The row containing the highest number of broadcasts that were sent appears at the top of the list. The network address in this row represents the highest broadcaster on your network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Use a Capture Filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use a capture filter**

- Design a capture filter
- Load a capture filter
- Specify frame data patterns to capture
- Specify capture filter protocols
- Save a capture filter
- Filter Data by Address

# Design a capture filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To design a capture filter**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Filter**.
4. In the Capture Filter dialog box, do the following:
   - Specify capture filter protocols.
   - Specify address pairs in a capture filter.
   - Specify frame data patterns to capture.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Specify capture criteria to identify the frames that you want to capture on the network.
- Build a complete capture filter expression by specifying the protocols, address pairs, and data patterns of the frames that you want to include or exclude from the capture.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save a capture filter
Save a display filter
Load a capture filter

# Load a capture filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To load a capture filter**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Filter** and then click **Load**.
4. Select the capture file you want and then click **Open**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- The decision tree that appears reflects the contents of the filter file that you have loaded. Network Monitor applies this filter when the data capture begins.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Design a capture filter
Set, view, change, or remove permissions on files and folders
Capture Window
Network Monitor overview

# Specify frame data patterns to capture

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify frame data patterns to capture**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Filter**.
4. In the capture filter decision tree, double-click the **AND (Pattern Matches)** line.
5. Complete the Pattern Match dialog box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- You can define up to four data patterns in a capture filter.
- To capture frames that match any of multiple patterns, create a logical branch by selecting one of the patterns in the decision tree and clicking **OR**.
- To exclude frames that match a pattern, select the pattern in the decision tree and then click **NOT**.
- If you specify more than one pattern in the capture filter decision tree, by default, Network Monitor treats each as a logical OR branch.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify capture filter protocols
Design a capture filter

# Specify capture filter protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify capture filter protocols**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Filter**.
4. Double-click the **SAP/ETYPE=** line.
5. Enable and disable protocols in the Capture Filter SAPs and ETYPES dialog box and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- All the protocols that you specified in your filter appear on the **SAP/ETYPE=** line in the capture filter decision tree.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify address pairs in a capture filter
Design a capture filter
Capture Filter SAPs and ETYPEs dialog box
Design a capture filter

# Save a capture filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save a capture filter**

1. Open Network Monitor.
2. On the **Capture** menu, click **Filter**.
3. When you finish editing the filter, click **Save**.
4. In the **File name** text box, type a name for the filter.
5. Click **Save**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access.
- Use a .cf file name extension for a capture filter or use a .df file name extension for a display filter.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Filtering data
Set, view, change, or remove permissions on files and folders
Design a display filter
Load a display filter

# Filter Data by Address

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Filter data by address**

- Specify address pairs in a capture filter
- Create an address database
- Edit an address in an address database
- Save the address database to a file
- Load an address file

# Specify address pairs in a capture filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify address pairs in a capture filter**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Filter**.
4. Double-click the **AND (Address Pairs)** line in the decision tree.

   Or, double-click the address pair you want to edit.
5. In the **Address Expression** dialog box, specify address pair properties, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- You can define up to three address pairs in a capture filter.
- In a capture filter, EXCLUDE statements take logical precedence over INCLUDE statements. Regardless of the sequence in which statements appear in the **Capture Filter** dialog box, Network Monitor evaluates EXCLUDE statements first. Therefore, in a filter containing both an EXCLUDE and an INCLUDE statement, if a frame meets the criteria specified in an EXCLUDE statement, Network Monitor discards the frame. Network Monitor does not test that frame against the INCLUDE statement to see if the frame also meets the criteria of the INCLUDE statement.
- When you close the capture filter, Network Monitor optimizes the logic you used in the **(Pattern Matches)** branch of the capture filter decision tree. Therefore, when you re-open the capture filter, it might look different from the way it did when you created it. For example, if you define multiple pattern match statements, Network Monitor automatically processes them with an implicit OR operator and removes any explicit OR operators between such statements.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify frame data patterns to capture
Design a capture filter

# Create an address database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an address database**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Start**.
4. On the **Capture** menu, click **Stop and View**.
5. On the **Display** menu, click **Find All Names**.
6. In the **Find All Names** dialog box, click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- After the frames are processed, a message appears telling you how many unique addresses were added to the database. To display a list of the addresses in the database you just created, on the **Display** menu, click **Addresses**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Capturing Data
How Network Monitor works
Save the address database to a file

# Edit an address in an address database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit an address in an address database**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Addresses**.
4. Click the address that you want to edit.
5. Click **Edit**.
6. Edit the address in the **Address Information** dialog box and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access.
- When you include the edited address in a capture or display filter, your changes are reflected in the filter.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save the address database to a file
Set, view, change, or remove permissions on files and folders
Load an address file

# Save the address database to a file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save the address database to a file**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Addresses**.
4. In the **Address Database** dialog box, click **Save**.
5. Specify the file name and location.
6. Click **Save**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- The next time that you design a capture or display filter that includes address pairs, you can use this database. If you assign the database the Default.adr file name, it becomes your default address database.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create an address database
Set, view, change, or remove permissions on files and folders
Load an address file

# Load an address file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To load an address file**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Capture** menu, click **Addresses**.
4. Click **Load**.
5. If a message box that asks you whether you want to cancel your changes and load another address database appears, click **Yes**.
6. Specify the file name of the address file (.adr) that you want to load.
7. Click **Open**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access.
- If you load an address database while another address database is open, you will lose unsaved changes made to the open database.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create an address database
Set, view, change, or remove permissions on files and folders
Save the address database to a file

# Display Network Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Display network data**

- Open an existing capture file
- Change the Frame Viewer window label
- Display network adapter vendor names
- Display address names
- Add a comment frame to a capture
- Specify the protocol header to display
- Select a protocol color scheme
- Display Frames
- Use a Display Filter

# Open an existing capture file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open an existing capture file**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **File** menu, click **Open**.
4. Select the capture file that you want to open, and then click **Open**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- You can open several capture files at one time; Network Monitor opens a separate Frame Viewer window for each capture file.
- Network Monitor can also open capture files created by some other programs. Some of the other file formats Network Monitor can read are:
    - Ethernet (files with an .ENC extension)
    - Token Ring (files with a .TRC extension)
    - FDDI (files with an .FDC extension)

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save captured frames to a file
Set, view, change, or remove permissions on files and folders
Design a capture filter
Network Monitor overview

# Change the Frame Viewer window label

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the Frame Viewer window label**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Window** menu, click **Label**.
5. In the Label Window dialog box, type the text that you want to appear in the window title bar, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open an existing capture file
Set, view, change, or remove permissions on files and folders

# Display network adapter vendor names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To display network adapter vendor names**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. On the **Options** menu, click **Show Vendor Names**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- When **Show Vendor Names** is active, Network Monitor replaces computer addresses with a twelve-byte code. The first six bytes of the code display the name of the vendor that produced the network adapter on the remote computer. The last six bytes are a host ID.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Display address names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To display address names**

1. Open Network Monitor.
2. On the **Options** menu, click **Show Address Names**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- When **Show Address Names** is active, a check mark appears next to the command name.
- User-assigned address names are displayed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Display network adapter vendor names

# Add a comment frame to a capture

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a comment frame to a capture**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Tools** menu, click **Insert Comment Frame**.
5. Complete the settings in the Insert Comment Frame dialog box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- The comment frame contains the Trail protocol, which includes statistics such as the amount of frames and bandwidth consumed. These statistics are based on a block of frames bounded by the current comment frame and the next comment frame of the same type (or the end of the capture).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Insert Comment Frame dialog box
Set, view, change, or remove permissions on files and folders
Frame Viewer Window

# Specify the protocol header to display

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify the protocol header to display**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Display Captured Data** when you are done capturing frames.
4. On the **Display** menu, click **Options**.
5. In the Display Options dialog box, click **Auto**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify the frame time stamp
Set, view, change, or remove permissions on files and folders
Specify capture filter protocols

# Select a protocol color scheme

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select a protocol color scheme**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Colors**.
5. In the Protocol Colors dialog box, click the protocol names and the color combination you want, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- If you want to specify more than one color combination, repeat this procedure for each protocol color scheme that you want.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open an existing capture file
Set, view, change, or remove permissions on files and folders

# Display Frames

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Display frames**

- View a specific frame
- Search for a frame by its properties
- Expand and collapse frame details
- Specify the frame time stamp
- Set the frame data display font

# View a specific frame

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view a specific frame**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Go To Frame**.
5. In the Go To dialog box, type the number of the frame that you want to view, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- The frame number is listed in the **Frame** column.
- Frames are numbered in the order that Network Monitor captures them.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Search for a frame by its properties
Set, view, change, or remove permissions on files and folders
Frame Viewer Window

# Search for a frame by its properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To search for a frame by its properties**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Find Next Frame**.
5. In the **Find Frame Expression** dialog box, do one of the following:
   - To search for a frame on the basis of its source or destination address, open and complete the Expression dialog box: Address tab.
   - To search for a frame on the basis of the protocols that were used to send it, open and complete the Expression dialog box: Protocol tab.
   - To search for a frame on the basis of protocol properties, open and complete the Expression dialog box: Property tab.
6. To begin searching for the expression displayed near the top of the dialog box, click **OK**.
7. To find the next frame that has the specified properties, on the **Display** menu, click **Repeat Find Next Frame**.
8. To find the previous frame that has the specified properties, on the **Display** menu, click **Repeat Find Previous Frame**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- If the property is found, the frame automatically expands to display the property that you specified in your search.
- The search remains active until you close the active window or begin a new search.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View a specific frame
Set, view, change, or remove permissions on files and folders
Print captured frames

# Expand and collapse frame details

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To expand and collapse frame details**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. Double-click the frame you want to view.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- To expand or collapse a list of frame data, in the Detail pane, double-click the line you want to expand or collapse.
- By default, frame data displayed in the Frame Viewer window Detail pane are collapsed.
- All other lines using the expanded protocol automatically expand as they appear.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Frame Viewer window: Detail pane
Set, view, change, or remove permissions on files and folders
Frame Viewer Window

# Specify the frame time stamp

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify the frame time stamp**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Options**.
5. In the Display Options dialog box, specify one of the available time options, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- The time setting appears under **Time** in the Frame Viewer window Summary pane.
- The **Seconds From Previous Frame** option provides very useful data for troubleshooting performance issues, because it helps you pinpoint gaps in time that can occur in a network transaction.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify the protocol header to display
Set, view, change, or remove permissions on files and folders
Search for a frame by its properties
View a specific frame

# Set the frame data display font

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the frame data display font**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Font**.
5. Specify the font settings in the Font dialog box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open an existing capture file
Set, view, change, or remove permissions on files and folders

# Use a Display Filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use a display filter**

- Design a display filter
- Load a display filter
- Modify the display filter decision tree structure
- Specify address pairs in a display filter
- Specify display filter protocol properties
- Specify display filter protocols
- Save a display filter

# Design a display filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To design a display filter**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Filter**.
5. In the Display Filter dialog box, do the following:
   - Specify display filter protocols.
   - Specify display filter protocol properties.
   - Specify address pairs in a display filter.
6. If needed, you can Modify the display filter decision tree structure.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- A Network Monitor display filter functions like a database query. You can use it to specify the types of captured data you want to examine. You also can use display filters to specify the types of displayed data, such as protocols and computer addresses, that you want to save to a file.
- You can add only one expression at a time. If you specify an expression and then click another tab in the Expression dialog box, that expression is lost. To save the specified expression and add it to the decision tree, you must click **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Filtering data

Set, view, change, or remove permissions on files and folders

Load a display filter

Save a capture filter

Save a display filter

# Load a display filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To load a display filter**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Filter**.
5. Click **Load**.
6. Specify the display filter file (.df), and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- The new display filter decision tree appears in the Display Filter dialog box.
- If you load a display filter, you will overwrite any unsaved changes to the current display filter.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save a capture filter
Save a display filter
Set, view, change, or remove permissions on files and folders
Design a display filter
Filtering data

# Modify the display filter decision tree structure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify the display filter decision tree structure**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Filter**.
5. To add a logical operator to the decision tree in the Display Filter dialog box, click the line below the row where you want to add the operator. Then:
   - To find frames that satisfy all expressions in the branch, click **AND**.
   - To find frames that satisfy one or more of the expressions in the branch, click **OR**.
   - To exclude frames that satisfy the expressions in the branch, click **NOT**.
6. To change a logical operator, click it and then click **Change Operator** until the operator you want appears.
7. To modify an expression, click it and then click **Edit Expression**.
8. To add an expression to a branch, click the line beneath the row where you want to insert the new expression. Then, in the **Add** group, click **Expression**. You can:
   - Specify display filter protocols.
   - Specify display filter protocol properties.
   - Specify address pairs in a display filter.
9. To delete a single expression or logical operator, click the line, and then click **Line** in the **Delete** group.
10. To delete a branch, click the top branch that you want to delete, click **Branch** in the **Delete** group, and click **Yes** to confirm that you want to delete the branch and all the branches underneath it.
11. To clear the decision tree and start over, click **All** in the **Delete** group. Click **Yes** to confirm that you want to start over and clear the decision tree contents.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have

been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.

- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.

- When you add an operator, a new decision tree branch appears.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save a capture filter
Save a display filter
Set, view, change, or remove permissions on files and folders
Design a display filter
Filtering data

# Specify address pairs in a display filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify address pairs in a display filter**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Filter**.
5. In the **Add** group box, click **Expression**.
6. Click the Expression dialog box: Address tab, specify the address pair, and click **OK**.
7. To make the address pair one of many display criteria, click it and then click **OR**.
8. To exclude the address pair from the display, click it and then click **NOT**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify display filter protocols
Set, view, change, or remove permissions on files and folders
Design a display filter
Filtering data

# Specify display filter protocol properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify display filter protocol properties**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Filter**.
5. In the display filter decision tree, click one row below where you want to add the expression.
6. In the **Add** group box, click **Expression**.
7. Click the Expression dialog box: Property tab, specify the protocol property that you want to add to the display filter, and click **OK**.
8. To make the protocol property one of many display criteria, click it and then click **OR**.
9. To exclude the protocol property from the display, click it and then click **NOT**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify address pairs in a display filter
Set, view, change, or remove permissions on files and folders
Design a display filter
Filtering data

# Specify display filter protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify display filter protocols**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **Display** menu, click **Filter**.
5. In the **Add** group box, click **Expression**.
6. Click the Expression dialog box: Protocol tabspecify which protocols in the display filter you want to enable and disable, and click **OK**.
7. To make the protocol one of many display criteria, click it and then click **OR**.
8. To exclude the protocol from the display, click it and then click **NOT**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify display filter protocol properties
Set, view, change, or remove permissions on files and folders
Design a display filter
Filtering data

# Save a display filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save a display filter**

1. Open Network Monitor.
2. On the **Display** menu, click **Filter**.
3. When you finish editing the filter, click **Save**.
4. In the **File name** text box, type a name for the filter.
5. Click **Save**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- Use a .cf file name extension for a capture filter or use a .df file name extension for a display filter.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Filtering data
Set, view, change, or remove permissions on files and folders
Design a display filter
Load a display filter

# Print Network Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Print network data**

- Print captured frames
- Print all data in a capture file
- Set the print filter

# Print captured frames

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print captured frames**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **File** menu, click **Print**.
5. In the Print dialog box: Netmon tab, specify the print options, and then click **Print**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.
- When you print frames, columns that do not fit on the page horizontally are truncated. To avoid losing important information, resize the columns before printing so they display only essential information. On the **Layout** tab, select **Landscape**. If you resize a column to its smallest possible size, it does not print at all.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Capture network frames
Set, view, change, or remove permissions on files and folders
Open an existing capture file
Set the print filter
Print all data in a capture file

# Print all data in a capture file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print all data in a capture file**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Display Captured Data** when you are done capturing frames.
4. On the **File** menu, click **Print**.
5. In the Print dialog box: Netmon tab, click **All** as the print range.
6. Click the **Netmon** tab.
7. Clear the **Apply Display Filter** check box, and then click **Print**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Print captured frames
Set, view, change, or remove permissions on files and folders
Set the print filter

# Set the print filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the print filter**

1. Open Network Monitor.
2. If prompted, select the local network from which you want to capture data by default.
3. Do one of the following:
   - On the **File** menu, click **Open**, and then double-click a saved capture file to open it.
   - On the **Capture** menu, click **Start**, and then click **Stop and View** when you are done capturing frames.
4. On the **File** menu, click **Print**.
5. In the Print dialog box: Netmon tab, click the **Netmon** tab.
6. Select the **Apply Display Filter** check box.
7. Click **Edit Filter**.
8. In the Edit Display Filter dialog box, edit the filter, and then click **OK**.
9. To print the filtered frames, click **Print**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Network Monitor, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Network Monitor**.
- Network Monitor runs with reduced access in which administrative privileges have been removed. If you receive an Access Denied message when you follow this procedure, add your user name to the permissions list of the file or folder that you want to access. See Related Topics for more information about setting permissions on files and folders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Print captured frames
Set, view, change, or remove permissions on files and folders
Print all data in a capture file
Design a display filter
Filtering data

# Network Monitor Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Network Monitor overview
- Understanding Network Monitor
- Using Network Monitor
- Network Monitor Resources

# Network Monitor overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Monitor overview**

By using Network Monitor, you can gather information that helps you keep your network running smoothly, from identifying patterns to preventing or solving problems. Network Monitor provides information about the network traffic that flows to and from the network adapter of the computer on which it is installed. By capturing and analyzing that information, you can prevent, diagnose, and solve many types of network problems. For more information about troubleshooting networks, see Network Monitor Resources.

You can configure Network Monitor to provide specific types of information that are most relevant to you. For example, you can set up triggers so that Network Monitor starts or stops capturing information when a circumstance or set of circumstances arises. You can also set up filters to control what types of information Network Monitor captures or displays. To make analyzing the information easier, you can modify how information appears on the screen, and you can save or print the information for review at a later time. For more information about customizing Network Monitor, see Capturing Data and Filtering data.

The Network Monitor component that ships with Microsoft® Windows Server 2003 family operating systems can capture frames that are sent to or from the computer on which Network Monitor is installed. If you want to capture frames that are sent to or from a remote computer, you must use the Network Monitor component that ships with Microsoft Systems Management Server, which can capture frames sent to or from any computer on which the Network Monitor driver is installed. For more information on Systems Management Server, go to the Microsoft Web site.

The information that Network Monitor provides comes from the network traffic itself, which is divided into frames. These frames contain information such as the address of the computer that sent the frame, the address of the computer to which the frame was sent, and the protocols that exist within the frame. For more information about frames, see How Network Monitor works.

**See Also**

**Concepts**

Specify the protocol header to display
Specify the frame time stamp
Select a protocol color scheme
Set the frame data display font
Print captured frames
Add a comment frame to a capture

# Understanding Network Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Network Monitor**
This section covers:

- How Network Monitor works
- Capturing Data
- Frame Viewer Window

# How Network Monitor works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Network Monitor works**

Data sent over a network is divided into frames. Each frame contains the following information:

- **Source address** The address of the network adapter from which the frame originated.
- **Destination address** The address of the network adapter that is meant to receive the frame. This address can also specify a group of network adapters.
- **Header information** Information specific to each protocol used to send the frame.
- **Data** The information (or a portion of it) being sent.

Every computer on a network segment receives frames transmitted on that segment. The network adapter in each computer retains and processes only those frames that are addressed to that adapter. The rest of the frames are dropped and no longer processed. The network adapter also retains broadcast (and potentially multicast) frames.

After installing Network Monitor, users can capture to a file all the frames sent to, or retained by, the network adapter of the computer on which it is installed. These captured frames can then be viewed or saved for later analysis. Users can design a capture filter so that only certain frames are captured. This filter can be configured to capture frames based on criteria such as source address, destination address, or protocol. Network Monitor also makes it possible for a user to design a capture trigger to initiate a specified action when Network Monitor detects a particular set of conditions on the network. This action can include starting a capture, ending a capture, or starting a program.

By default, the size of the capture buffer is 1 MB. You can reduce the amount of data you capture by shrinking the capture buffer.

**See Also**

**Concepts**

Capture Window

Frame Viewer Window

# Capturing Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capturing data**
This section covers:

- The capture process
- Capture filters
- Capture triggers

# The capture process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The capture process**

The process by which Network Monitor copies frames is referred to as capturing. You can capture all network traffic to and from the local network adapter, or you can set a capture filter and capture a subset of frames. You can also specify a set of conditions that trigger an event. If you create triggers, Network Monitor can respond to events on your network. For example, you can make the operating system start an executable file when Network Monitor detects a particular set of conditions on the network. After you have captured data, you can view it. Network Monitor translates the raw capture data into its logical frame structure.

While Network Monitor captures frames from the network, statistics about the frames appear in the Capture window, which has four panes:

| Pane | Displays |
|---|---|
| Capture window: Graph pane | A graphical representation of frames sent to or from the local computer. |
| Capture window: Session Statistics pane | Statistics about current individual sessions. |
| Capture window: Station Statistics pane | Statistics about frames sent to or from the computer running Network Monitor. |
| Capture window: Total Statistics pane | Summary statistics about frames sent to or from the local computer since the capture process began. |

Network Monitor uses the Network Driver Interface Specification (NDIS) facility to copy all frames it detects to its capture buffer.

For detailed instructions on capturing network data, see Capture network frames.

# Capture filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture filters**
A capture filter functions like a database query that you can use to specify the types of network information you want to monitor. For example, to see only a specific subset of computers or protocols, you can create an address database, use the database to add addresses to your filter, and then save the filter to a file. By filtering frames, you save both buffer resources and time. Later, if necessary, you can load the capture filter file and use the filter again.

**Designing a capture filter**
To design a capture filter, specify decision statements in the **Capture Filter** dialog box. This dialog box displays the filter's decision tree, which is a graphical representation of a filter's logic. When you include or exclude information from your capture specifications, the decision tree reflects these specifications.

**Filtering by protocol**
To capture frames sent using a specific protocol, specify the protocol on the **SAP/ETYPE=** line of the capture filter. For example, to capture only IP frames, disable all protocols and then enable **IP ETYPE 0x800** and **IP SAP 0x6**. By default, all of the protocols that Network Monitor supports are enabled. You can only specify protocols with ETYPE or SAP.

**Filtering by address**
To capture frames sent from a specific computer on your network to your computer or sent from your computer to a specific computer on your network, specify one or more address pairs in a capture filter. You can monitor up to four address pairs simultaneously.

An address pair consists of:

- The addresses of the two computers you want to monitor traffic between.
- Arrows that specify the traffic direction you want to monitor.
- The INCLUDE or EXCLUDE keyword, indicating how Network Monitor should respond to a frame that meets a filter's specifications.

Regardless of the sequence in which statements appear in the **Capture Filter** dialog box, EXCLUDE statements are evaluated first. Therefore, if a frame meets the criteria specified in an EXCLUDE statement in a filter containing both an EXCLUDE and INCLUDE statement, that frame is discarded. Network Monitor does not test that frame by INCLUDE statements to see if it meets that criterion also.

For example, to capture all the traffic from Joe's computer *except* the traffic from Joe to Anne, use the following capture filter address section:

**Addresses**

**include Joe <----> Any**

**exclude Joe <----> Anne**

If there are no INCLUDE lines, *YourComputer* <----> Any is used implicitly.

**Filtering by data pattern**
By specifying a pattern match in a capture filter, you can:

- Limit a capture to only those frames containing a specific pattern of ASCII or hexadecimal data.
- Specify how many bytes (offsets) of the frame must be ignored before the search begins.

When you filter based on a pattern match, you must specify where, in the frame, the search for the pattern should begin. This setting specifies, in bytes, the distance from the beginning of the frame or the end of the topology header to the point at which the pattern might occur. If your network medium has a variable size in the media access control protocol, such as Ethernet or

token ring, specify to count from the end of the topology header.

# Capture triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture triggers**
If you create capture triggers, Network Monitor can respond to events on your network. By default, no trigger is set.

**Types of triggers**
Network Monitor can detect how full your capture buffer is, and whether a specific pattern appears in a captured frame. You can create capture triggers that depend on either or both of these criteria.

If you specify a trigger that depends on a specific pattern appearing in a captured frame, Network Monitor performs the action you specify when it detects a frame that contains the pattern you specify. The pattern can be an ASCII string or a hexadecimal string. You can also specify a trigger which depends on a specific pattern appearing in a captured frame and a certain percentage of the capture buffer being filled. You can also specify whether Network Monitor should start its search at the beginning of each frame, after the end of the header of each frame, or some number of bytes after either of these locations. By default, Network Monitor searches the entire frame for the pattern.

**Trigger actions**
You can choose one of the following actions to occur when the trigger criteria are met:

- The computer beeps.
- Network Monitor stops capturing frames.
- A command that you specify runs.

  To specify a command that starts a program, type the path and the name of the program file, or click **Browse** and navigate to the program file. To use an MS-DOS command, such as **copy**, type **CMD /K**, and then type the command.

# Displaying Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Displaying data**
This section covers:

- The Frame Viewer window
- Filtering data
- Protocol parsers included with Network Monitor

# The Frame Viewer window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The Frame Viewer window**

Network Monitor simplifies data analysis by interpreting raw data collected during the capture and displaying it in the Frame Viewer window.

The Frame Viewer window includes the following panes:

| Pane | Displays |
|---|---|
| Frame Viewer window: Summary pane | General information about the captured frames in the order that they were captured in. |
| Frame Viewer window: Detail pane | The contents of the frame, including the protocols that were used to send it. |
| Frame Viewer window: Hexadecimal pane | Hexadecimal and ASCII representations of the captured data. |

# Filtering data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Filtering data**
Like a capture filter, a display filter helps you to isolate specific types of information. Unlike a capture filter, however, a display filter operates on data that has already been captured. You can use a display filter to configure how much captured data appears in the Frame Viewer Window, or you can use it to save data to a capture file.

Use display filters to determine which frame to display. You can filter a frame by:

- Its source or destination address.
- The protocols it contains.
- The properties and values that it contains. The properties of a protocol collectively indicate the purpose of the protocol.

The structure of the display-filter decision tree is flexible. You can define a simple, rather flat structure, or you can make it complex, as your needs dictate.

**Protocols Used**
When you display captured data, all available information about the captured frames appears in the Frame Viewer window. You can display only the frames which contain a specific protocol by editing the **Protocol ==** line in the Display Filter dialog box.

Network Monitor processes the filter you have designed and applies it to the contents of the Frame Viewer window.

**Protocol Properties**
You can use a display filter to isolate frames that contain specific protocol properties. Protocol properties consist of the elements of information that define the purpose of a protocol. Because the purposes of protocols vary, properties differ from one protocol to another.

As an example, you might capture a large number of frames with SMB protocol but want to examine only those frames in which the SMB protocol was used to create a directory on a remote computer. In this situation, you could isolate frames that include the **Make Directory** SMB command property.

Network Monitor identifies the protocols used to send a frame on the network by using a protocol parser. Each protocol that Network Monitor supports has a corresponding parser.

**Computer Addresses**
When you display captured data, by default all addresses that you capture information from appear in the Frame Viewer window. You can display only those frames that originate or are sent to a specific computer by editing the **ANY <--> ANY** line in the **Edit Display Filter** dialog box.

**Filtering and address databases**
Often, you need to capture only those frames that originate with or are sent to specific computers. To do this, you must know the addresses of the computers on your network.

You can use the **ping** command to find the IP address of a computer if you know its computer name. For more information, see the section "Testing connections by using ping" in Command-line utilities.

**See Also**
**Concepts**
Save a capture filter
Save a display filter
Load a display filter
Design a display filter

# Protocol parsers included with Network Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Protocol parsers included with Network Monitor**

Each protocol that Network Monitor supports has a corresponding parser. The following is a list of the protocols whose parsers are included with Network Monitor.

| Protocol | Description |
|---|---|
| AARP | AppleTalk Address Resolution Protocol (Atalk.dll) |
| ADSP | AppleTalk Data Stream Protocol (Atalk.dll) |
| AFP | AppleTalk File Protocol (Atalk.dll) |
| AH | IP Authentication Header (Tcpip.dll) |
| ARP_RARP | Internet Address Resolution Protocol/Reverse Address Resolution Protocol (Tcpip.dll) |
| ASP | AppleTalk Session Protocol (Atalk.dll) |
| ATMARP | ATM Address Resolution Protocol (Atmarp.dll) |
| ATP | AppleTalk Transaction Protocol (Atalk.dll) |
| BONE | Bloodhound-Oriented Network Entity Protocol (Bone.dll) |
| BOOKMARK | Network Monitor BOOKMARK Protocol (Trail.dll) |
| BPDU | Bridge Protocol Data Unit (Llc.dll) |
| BROWSER | Microsoft Browser (Browser.dll) |
| CBCP | Callback Control Protocol (Ppp.dll) |
| CCP | Compression Control Protocol (Ppp.dll) |
| COMMENT | Network Monitor COMMENT Protocol (Trail.dll) |
| DDP | AppleTalk Datagram Delivery Protocol (Atalk.dll) |
| DHCP | Dynamic Host Configuration Protocol (Tcpip.dll) |
| DNS | Domain Name System (Tcpip.dll) |
| EAP | PPP Extensible Authentication Protocol (Ppp.dll) |
| ESP | IP Encapsulating Security Payload (Tcpip.dll) |
| ETHERNET | Ethernet 802.3 topology (Mac.dll) |
| FDDI | FDDI topology (Mac.dll) |
| FINGER | Internet Finger protocol (Tcpip.dll) |
| FRAME | Base frame properties (Frame.dll) |
| FTP | File Transfer Protocol (Tcpip.dll) |
| GENERIC | Network Monitor GENERIC Protocol (Trail.dll) |

| | |
|---|---|
| GRE | Generic Routing Encapsulation Protocol (Ppp.dll) |
| ICMP | Internet Control Message Protocol (Tcpip.dll) |
| IGMP | Internet Group Management Protocol (Tcpip.dll) |
| IP | Internet Protocol (Tcpip.dll) |
| IPCP | Internet Protocol Control Protocol (Ppp.dll) |
| IPX | NetWare Internetwork Packet eXchange protocol (Ipx.dll) |
| IPXCP | NetWare Internetwork Packet eXchange Control Protocol (Ppp.dll) |
| IPXWAN | NetWare Internetwork Packet eXchange Protocol for Wide Area Networks (Ppp.dll) |
| ISAKMP | Internet Security Association and Key Management Protocol (Ppp.dll) |
| L2TP | Level 2 Tunneling Protocol (L2tp.dll) |
| LAP | AppleTalk Link Access Protocol (Atalk.dll) |
| LCP | Link Control Protocol (Ppp.dll) |
| LLC | Logical Link Control 802.2 Protocol (LLC.dll) |
| LPR | BSD Printer (Ppp.dll) |
| MESSAGE | Network Monitor MESSAGE Protocol (Trail.dll) |
| MSRPC | Microsoft Remote Procedure Call Protocol (Msrpc.dll) |
| NBFCP | NetBIOS Frames Control Protocol (Ppp.dll) |
| NBIPX | NetBIOS on IPX (Ipx.dll) |
| NBP | AppleTalk Name Binding Protocol (Atalk.dll) |
| NBT | Internet NetBIOS Over TCP/IP (Tcpip.dll) |
| NCP | NetWare Core Protocol (Ncp.dll) |
| NDR | NetWare Diagnostic Redirector (Ipx.dll) |
| NETBIOS | Network Basic Input/Output System Protocol (Netbios.dll) |
| NETLOGON | Microsoft Netlogon Broadcasts (Netlogon.dll) |
| NFS | Network File System (Tcpip.dll) |
| NMPI | Microsoft Name Management Protocol on IPX (Ipx.dll) |
| NSP | NetWare Serialization Protocol (Ipx.dll) |
| NWDP | NetWare WatchDog Protocol (Ipx.dll) |
| ODBC | Network Monitor ODBC Protocol (Trail.dll) |
| OSPF | Open Shortest Path First (Tcpip.dll) |
| PAP | AppleTalk Printer Access Protocol (Atalk.dll) |
| PPP | Point-to-Point Protocol (Ppp.dll) |
| PPPCHAP | PPP Challenge Handshake Authentication Protocol (Ppp.dll) |

| | |
|---|---|
| PPPML | Point-to-Point Multilink Protocol (Ppp.dll) |
| PPPPAP | PPP Password Authentication Protocol (Ppp.dll) |
| PPTP | Point-to-Point Tunneling Protocol (Ppp.dll) |
| R_LOGON | Generated RPC for interface logon (Logon.dll) |
| R_LSARPC | Generated RPC for Interface Lsarpc (Lsarpc.dll) |
| R_WINSPOOL | Generated RPC for Interface Winspool (Winspl.dll) |
| RADIUS | Remote Authentication Dial-In User Service Protocol (Ppp.dll) |
| RIP | Internet Routing Information Protocol (Tcpip.dll) |
| RIPX | NetWare Routing Information Protocol (Ipx.dll) |
| RPC | Remote Procedure Call (Tcpip.dll) |
| RPL | Remote Program Load (Llc.dll) |
| RSVP | Resource Reservation Protocol (Rsvp.dll) |
| RTMP | AppleTalk Routing Table Maintenance Protocol (Atalk.dll) |
| SAP | NetWare Service Advertising Protocol (Ipx.dll) |
| SMB | Server Message Block Protocol (Smb.dll) |
| SMT | FDDI MAC Station Management (Mac.dll) |
| SNAP | Sub-Network Access Protocol (Llc.dll) |
| SNMP | Simple Network Management Protocol (Snmp.dll) |
| SPX | NetWare Sequenced Packet eXchange Protocol (Ipx.dll) |
| SSP | Security Support Provider Protocol (Msrpc.dll) |
| STATS | Network Monitor Capture Statistics Protocol (Trail.dll) |
| TCP | Transmission Control Protocol (Tcpip.dll) |
| TMAC | Token Ring MAC layer (Mac.dll) |
| TOKENRING | Token Ring 802.5 Topology (Mac.dll) |
| TPCTL | Test Protocol Control Language (Tpctl.dll) |
| TRAIL | Network Monitor TRAIL Protocol (Trail.dll) |
| UDP | User Datagram Protocol (Tcpip.dll) |
| VINES_FRAG | Banyan Vines Fragmentation Protocol (Vines.dll) |
| VINES_IP | Banyan Vines Internet Protocol (Vines.dll) |
| VINES_TL | Banyan Vines Transport Layer Protocols (Vines.dll) |
| XNS | Xerox Network System (Xns.dll) |
| ZIP | AppleTalk Zone Information Protocol (Atalk.dll) |

**Note**

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows

Server 2003 family.

# Using Network Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Network Monitor**

This section covers:

- Managing Network Monitor from the command line
- Capture Window
- Frame Viewer Window

# Managing Network Monitor from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Network Monitor from the command line**

By using the **start netmon** command, you can start the Network Monitor user interface from the command line. You can also specify parameters for capturing and filtering frames from network traffic going to and coming from the computer on which you have Network Monitor installed.

**Syntax**

**start netmon** [**/net** *Number*] [**/capturefilter** *Path*] [**/displayfilter** *Path*] [**/buffersize** *Number*] [**/quickfilter** {**ATM** | **ETHERNET** | **IP** | **IPX** | **VINES** | **TOKENRING** | **FDDI**}**,** *Address*] [**/quickfiltername** *FilterName*] [**/autostart**] [**/autostop**]

**Parameters**

**/net *Number***

Specifies the number of the network from which you want to capture frames. You can specify any of the network segments to which your computer is connected.

**/capturefilter *Path***

Specifies a capture filter to be used when Network Monitor starts. You can abbreviate **/capturefilter** as **/cf**.

**/displayfilter *Path***

Specifies a display filter to be used when Network Monitor starts. You can abbreviate **/displayfilter** as **/df**.

**/buffersize: *Number***

Specifies the size (in megabytes) of the capture buffer.

**/quickfilter{ ATM| ETHERNET| IP| IPX| VINES| TOKENRING| FDDI} , *Address***

Specifies that Network Monitor should capture frames only of a particular type and from a particular address. You can specify up to three types at a time. You can abbreviate **/quickfilter** as **/qf**. The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

**/quickfiltername *FilterName***

Specifies the name of the filter that you describe with the **/quickfilter** command. This parameter is ignored if you do not specify any **/quickfilter** options. You can abbreviate **/quickfiltername** as **/qfn**.

**/autostart**

Specifies that Network Monitor should begin to capture data at interface startup.

**/autostop**

Specifies that Network Monitor should stop capturing data when the capture buffer is full.

**/?**

Displays help at the command prompt.

**Remarks**

- Network Monitor is not installed by default. To use this command, you must have installed Network Monitor on your computer, and you must run the command from the directory where Network Monitor is installed.

**Examples**

To set the size of your capture buffer to 4 megabytes, type:

**start netmon /buffersize:4**

To load a specific capture filter when you start Network Monitor, type:

**start netmon /capturefilter d:\captures\http.cf**

**Formatting legend**

| Format | Meaning |
|---|---|

| | |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Install and Configure Network Monitor
Command-line reference A-Z
Command shell overview

# Capture Window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture window**
This section covers:

- Capture Window Panes
- Capture Window Dialog Boxes

# Capture Window Panes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Panes**

This section covers:

- Capture window: Graph pane
- Capture window: Session Statistics pane
- Capture window: Station Statistics pane
- Capture window: Total Statistics pane

# Capture window: Graph pane

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture window: Graph pane**

The Graph pane graphically represents the total capture statistics of the current capture data. This pane appears in the upper-left corner of the Capture Window, and it is on by default.

**% Network Utilization**

The percentage of your network adapter's capacity that the current capture uses. This percentage is calculated by dividing the rate at which your adapter is sending and receiving frames by the maximum rate at which your adapter can process those frames.

**Frames Per Second**

The number of frames that the adapter is capturing every second.

**Bytes Per Second**

The number of bytes that the adapter is capturing every second.

**Broadcasts Per Second**

The number of broadcasts that the adapter is capturing every second.

**Multicasts Per Second**

The number of multicasts that the adapter is capturing every second.

Other panes in this window: Capture window: Session Statistics pane, Capture window: Station Statistics pane, Capture window: Total Statistics pane

**See Also**

**Concepts**

Capture Window

How Network Monitor works

# Capture window: Session Statistics pane

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture window: Session Statistics pane**
A session designates the data that is sent to or from the local computer. The Session Statistics pane displays statistics on a per-session basis. The pane also identifies both participants in a session and displays how much information passed in either direction between them.

This pane appears in the left-center portion of the Capture Window, and it is on by default.

**Network Address 1**
The network address of the first participant in a network session.

**1 --> 2**
The number of frames sent from the address listed in the **Network Address 1** column to the address listed in the **Network Address 2** column.

**1 <-- 2**
The number of frames sent from the address listed in the **Network Address 2** column to the address listed in the **Network Address 1** column.

**Network Address 2**
The network address of the second participant in a network session.

**Note**

- Network Monitor reflects session statistics of only the first 100 unique addresses that it detects. To gather statistics on a specific workstation, design a capture filter. To reset statistics and view information about the next 100 detected network sessions, click the **Capture** menu, and then click **Clear Statistics**.

Other panes in this window: Capture window: Graph pane, Capture window: Station Statistics pane, Capture window: Total Statistics pane

**See Also**
**Concepts**
Capture Window
How Network Monitor works

# Capture window: Station Statistics pane

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture window: Station Statistics pane**

The Station Statistics pane displays statistics that describe the network activity of your workstation. This pane appears at the bottom of the Capture Window, and it is on by default.

**Network Address**

The network address of the computer on which the frames were captured.

**Frames Sent**

The number of frames sent from the address listed in the **Network Address** column.

**Frames Rcvd**

The number of frames received by the address listed in the **Network Address** column.

**Bytes Sent**

The number of bytes sent by the address listed in the **Network Address** column.

**Bytes Rcvd**

The number of bytes received by the address listed in the **Network Address** column.

**Directed Frames Sent**

The number of non-broadcast, non-multicast frames transferred over the network from the associated address.

**Multicasts Sent**

The number of times the address listed in the **Network Address** column has sent frames to a subset of computers on the network, by sending "FFFFFFFFFFFF."

**Broadcasts Sent**

The number of times that the address listed in the **Network Address** column has sent frames to all computers on the network.

**Note**

- Network Monitor reflects station statistics of only the first 128 unique addresses that it detects. To gather statistics on a specific workstation, design a capture filter.

Other panes in this window: Capture window: Graph pane, Capture window: Session Statistics pane, Capture window: Total Statistics pane

**See Also**

**Concepts**

Capture Window

Capture filters

# Capture window: Total Statistics pane

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture window: Total Statistics pane**

The Total Statistics pane provides an overall view of network traffic sent to or from the local computer. This pane appears in the upper-right corner of the Capture Window, and it is on by default.

**Network Statistics**

The total amount of traffic that has been sent to or from the local computer since the current capture began. These statistics include:

- The total number of frames sent to or from the local computer.
- The total number of broadcasts sent to or from the local computer.
- The total number of multicasts sent to or from the local computer.
- The total number of bytes sent to or from the local computer.
- The total number of frames dropped.
- The network status. On an Ethernet network, this entry will always be **Normal**. On a token ring network, this entry reflects whether the token is present locally.

**Captured Statistics**

Total statistics that describe the current capture, including:

- The total number of captured frames.
- The total number of frames in the temporary capture file.
- The number of frames dropped when the buffer was exceeded.
- The total number of captured bytes.
- The total number of bytes in the temporary capture file.
- The percentage of allotted buffer space that is in use.
- The number of frames dropped by Network Monitor.

**Per Second Statistics**

Statistical averages of current activity and continual updates of this average to reflect current per-second activity. The statistics in this panel include all frames (even frames that were excluded by a capture filter). Per-second statistics include:

- The average number of frames per second detected since the capture began.
- The average number of bytes per second detected since the capture began.
- The average number of broadcast messages per second detected since the capture began.
- The average number of multicast messages per second detected since the capture began.
- The percentage of network utilization. This statistic shows the percentage of your network adapter's capacity that the

current capture uses. This percentage is calculated by dividing the rate at which your adapter is sending and receiving frames by the maximum rate at which your adapter can process those frames.

**Network Card (MAC) Statistics**

Statistics that reflect average activity detected by your network adapter since the current capture began. The statistics in this pane reflect all the network activity that your network adapter can receive. These statistics include:

- Total frames detected by the network adapter.
- Total broadcast frames detected by the network adapter.
- Total multicast frames detected by the network adapter.
- Total bytes detected by the network adapter.

**Network Card (MAC) Error Statistics**

Statistics that reflect adapter-related errors that have occurred since the capture began. These statistics include:

- Number of errors that occurred because the cyclical redundancy check (CRC) did not match the actual bytes received.
- Number of frames that the network adapter detected but that were dropped because Network Monitor lacked sufficient buffer space.
- Number of frames that the network adapter detected but that were dropped because of hardware constraints.

**Note**

- Not all network adapters support these statistics. If your network adapter does not support a statistic, the value is replaced with **Unsupported**.

Other panes in this window: Capture window: Graph pane, Capture window: Session Statistics pane, Capture window: Station Statistics pane

**See Also**
**Concepts**
Capture Window
How Network Monitor works

# Capture Window Dialog Boxes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dialog boxes**
This section covers:

- Address Database dialog box
- Address Expression dialog box
- Address Information dialog box
- Capture Buffer Settings dialog box
- Capture Filter dialog box
- Capture Filter SAPs and ETYPEs dialog box
- Capture Trigger dialog box
- Dedicated Mode dialog box
- Open File dialog box
- Pattern Match dialog box
- Print dialog box: Netmon tab
- Protocol Parsers dialog box
- Select a Network dialog box

# Address Database dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Address Database dialog box**

Use this dialog box to add, delete, or edit address database information. You can also use this dialog box to save the current address database or load a different one.

**Name**

The user-defined name of the computer.

**Address**

The 4-byte (IPv4) or 16-byte (IPv6) address of the network adapter.

**Type**

The network adapter address type (MAC or IP, for example).

**Comment**

Comments or additional address information.

**Vendor**

Name of the vendor that created the network adapter.

**Add**

Click to open the Address Information dialog box, which you can use to add a new computer address to the address database.

**Edit**

Click to open the Address Information dialog box, which you can use to change the currently selected computer address.

**Delete**

Click to remove the currently selected computer address from the address database.

**Load**

Click to open the **Open Address file** dialog box, which you can use to load a different address database instead of the one currently displayed.

**Save**

Click to open the **Save Addresses as** dialog box, which you can use to save the currently displayed address database to a file. Unless you save any changes that you made, you lose them when you load a new address database or close Network Monitor.

**See Also**

**Concepts**

Create an address database

How Network Monitor works

# Address Expression dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Address Expression dialog box**

Use this dialog box to create a new address expression that is added to the **Address Pairs** section of the current capture filter.

**Include**

Includes capture data that travels between the members of the current address pair. By default, this option is selected.

**Exclude**

Excludes capture data that travels between the members of the address pair.

**Station 1**

The list of first addresses in the address pair. Depending on the arrow that you select in **Direction**, **Station 1** can be the originating or the destination address.

**Direction**

Specify an arrow that indicates the direction of the traffic between the two addresses.

**<-->**

Click to specify frames that travel in either direction between the computers that you specified in **Station 1** and **Station 2**. By default, this arrow is selected.

**-->**

Click to specify frames that travel from **Station 1** to **Station 2**.

**<--**

Click to specify frames that travel from **Station 2** to **Station 1**.

**Station 2**

The list of second addresses in the address pair. Depending on the arrow that you select in **Direction**, **Station 2** can be the originating or the destination address.

**Edit Addresses**

Click to open the Address Database dialog box, which you can use to modify the current address database.

**Note**

- When you specify **Direction**, you cannot construct a statement in which the source address is either Broadcast or Multicast because Broadcast and Multicast are always destination addresses.

**See Also**

**Concepts**

Specify address pairs in a capture filter
Capture Window
Address Database dialog box

# Address Information dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Address Information dialog box**

Use this dialog box to edit individual address entries or add a new entry.

**Name**

The friendly name (computer name, for example) for the **Address**.

**Permanent Name**

Uses the text in the **Name** text box as the friendly name of the computer in the local address database (.adr). By default, this check box is cleared.

**Address**

The address that identifies the computer to the network.

**Type**

The address type that associates with **Address**. By default, the detected address type appears in this list.

For example, suppose that the network associates a name with an Ethernet address and you want to associate the name with an IP address. You select **IP** and then type the IP address in the **Address** text box. Associating a name with a new address type removes its association with the previous address type.

**Comment**

Any additional information that you want to include in the address description. For example, you might enter the type of CPU or network adapter or mention some problem the computer has exhibited.

**See Also**

**Concepts**

Create an address database

How Network Monitor works

# Capture Buffer Settings dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture Buffer Settings dialog box**

Use this dialog box to adjust the size of the data frame and the total amount of frames you want to capture.

**Buffer Size (MB)**

The size of your capture buffer. By default, the buffer size is set to 1.0 MB.

**Frame Size (bytes)**

The number of bytes that you want Network Monitor to capture from each frame. By default, the frame size is **Full** (65,535).

The drop-down list contains numbers in increments of 64, up to 65,472. You can select one of these numbers, or you can type a specific number between 32 and 65,535, inclusive.

**See Also**

**Concepts**

Modify capture buffer settings
Capture network frames
Save captured frames to a file

# Capture Filter dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture Filter dialog box**
This dialog box displays the capture filter decision tree and helps you to specify the frames that you want to capture. You can also load filters that were previously designed and saved to a file. By doing this, you avoid creating filters every time you want to capture filtered data.

You can use the following controls to make changes to the decision tree.

- **SAP/ETYPE =**

  Double-click this line to specify protocols that you want to capture in frames.

- **AND (Address Pairs)**

  Double-click this line to filter on computer address pairs.

  You can specify as many as three address pairs.

- **AND (Pattern Matches)**

  Double-click this line to capture only frames that contain a specified pattern at a specified offset.

  You can specify as many as four pattern matches.

- **Add**

  Use the buttons in this group box to add criteria to the capture filter.

  - **Address**

    Click to open the Address Expression dialog box, which you can use to specify the address pairs of the computers between which you want to capture data.

  - **Pattern**

    Click to open the Pattern Match dialog box, which you can use to capture only those frames containing a specific pattern of ASCII or hexadecimal data.

- **Insert**

  Use the buttons in this group box to add operators to the capture filter.

  - **OR**

    Click to add an OR branch to the capture filter decision tree. This option is available only for pattern matches.

  - **NOT**

    Click to add a NOT branch to the decision tree or to remove a branch of the tree. This option is available only for pattern matches.

- **Edit**

Click to edit the line that is currently selected in the capture filter.

- **Delete**

  Click to delete the line that is currently selected in the capture filter.

- **Load**

  Click to open the **Load Capture Filter** dialog box, which you can use to replace the current capture filter with one that has been saved to a file.

- **Save**

  Click to open the **Save Capture Filter** dialog box, which you can use to save the current capture filter as a file.

## Notes

- In the address portion of a capture filter, EXCLUDE statements take logical precedence over INCLUDE statements. Regardless of the sequence in which statements appear in the **Capture Filter** dialog box, Network Monitor evaluates EXCLUDE statements first. Therefore, in a filter containing both an EXCLUDE and an INCLUDE statement, if a frame meets the criteria specified in an EXCLUDE statement, Network Monitor discards the frame without testing it against the INCLUDE statement to see whether the frame also meets the criteria of the INCLUDE statement.

- When you close the capture filter, Network Monitor optimizes the logic you used in the **Pattern Matches** branch of the capture filter decision tree. Therefore, when you re-open the capture filter, it might look different from when you created it. For example, if you define multiple pattern match statements, Network Monitor automatically processes them with an implicit OR operator and removes any explicit OR operators between such statements.

**See Also**
**Concepts**
Save a capture filter
Save a display filter
Load a capture filter
Design a capture filter
How Network Monitor works

# Capture Filter SAPs and ETYPEs dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture Filter SAPs and ETYPEs dialog box**
Use this dialog box to specify the protocols that Network Monitor captures while your capture filter runs. The choices that you make in this dialog box are reflected in the line that begins with **SAP/ETYPE=** in the decision tree in the Capture Filter dialog box. Only protocols that use an SAP or an ETYPE can be filtered in the capture filter.

- **Enabled Protocols**

  The protocols that Network Monitor uses to capture frames from the network. By default, all protocols are enabled.

  Double-click any protocol in this box to move it to the **Disabled Protocols** table.

  - **Protocol**

    The names of all currently enabled protocols.

  - **Comment**

    A brief description of each enabled protocol.

  - **Type**

    The enabled protocol ETYPE or SAP.

  - **Values**

    Possible values of enabled protocols.

- **Disable**

  Click to move the protocols that you select in the **Enabled Protocols** table to the **Disabled Protocols** table.

- **Disable All**

  Click to disable all available protocols.

- **Enable**

  Click to move protocols that you selected in the **Disabled Protocols** table to the **Enabled Protocols** table.

- **Enable All**

  Click to enable all available protocols.

- **Disabled Protocols**

  The currently disabled protocols, which will not be used to capture network data.

  - **Protocol**

    The names of all currently disabled protocols.

- **Comment**

  A brief description of each disabled protocol.

- **Type**

  The disabled protocol type (ETYPE or SAP).

- **Values**

  Possible values of disabled protocols.

**See Also**
**Concepts**
Design a capture filter
Capture network frames

# Capture Trigger dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capture Trigger dialog box**

Use this dialog box to specify the event that acts as a trigger and the action that results from the event.

- **Trigger On**

  Specify the conditions that induce the trigger action. By default, no trigger is set.

  - **Nothing**

    Turns off trigger actions during data capture.

  - **Pattern match**

    Engages the trigger when a specified hexadecimal or ASCII pattern match occurs.

  - **Buffer space**

    Engages the trigger when a capture fills a specific percentage of the capture buffer.

  - **Pattern match then buffer space**

    Requires Network Monitor to monitor the buffer space after the specific matched pattern occurs and to perform the trigger action when both specified conditions exist.

  - **Buffer space then pattern match**

    Requires Network Monitor to detect when a capture has filled a specified percentage of the capture buffer and then to perform the trigger action when both conditions exist.

- **Buffer Space**

  The percentage of capture buffer that is filled before the trigger action occurs. For example, if you want to trigger action to occur when the capture buffer is three-quarters full, select **75%**. The only accepted values are multiples of 25%.

- **Pattern**

  The pattern you select for the pattern match.

  - **Offset (Hex)**

    The number of hexadecimal bytes from the beginning of the frame to where Network Monitor begins looking for a pattern match. By default, the pattern triggers if the pattern occurs anywhere in the frame.

    The pattern offset can be measured from the beginning of the frame or from the end of the topology header.

  - **From Start of Frame**

    Starts the search for the pattern at the beginning of the frame.

  - **From End of Topology Header**

Starts the pattern search beyond the topology header.

- **Pattern**

  The pattern that you select to engage the trigger.

- **Hex**

  A hexadecimal pattern.

- **ASCII**

  An ASCII pattern.

- **Trigger Action**

  Specify the action that you want to occur when the trigger criteria (pattern match, buffer space, or both) are met.

  - **Audible Signal Only**

    The computer beeps when the trigger condition is met and the capture process continues without interruption. By default, this option is selected.

  - **Stop Capture**

    The capture process stops when the trigger condition is met.

  - **Execute Command Line**

    Runs a command or opens a file when the trigger specifications are met. You can use as many as 259 characters for the command line. Click **Browse** to specify a file you want to run.

### Note

- If you click the **Pattern match then buffer space** option and set **Buffer Space** to **100%**, Network Monitor will overwrite the frame that has the pattern match because it does not start counting the buffer space until after the pattern match is found.

**See Also**
**Concepts**
Set a capture trigger
How Network Monitor works

# Dedicated Mode dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dedicated Mode dialog box**

This dialog box appears when you click **Dedicated Capture Mode** on the **Capture** menu and start capturing data.

**Total Frames Captured**

The number of frames that Network Monitor captures.

**Stop**

Click to stop the capture and return to Network Monitor.

**Stop and View**

Click to stop the capture and to view the frames in the Network Monitor Frame Viewer window.

**Pause (Continue)**

Click to stop (or resume) capturing frames temporarily.

**Normal Mode**

Click to turn off **Dedicated Capture Mode** and return to the Network Monitor Capture window.

Frame capture continues until you explicitly stop the capture process.

**See Also**

**Concepts**

Capture data in dedicated capture mode

How Network Monitor works

# Open File dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Open File dialog box**
Use this dialog box to:

- Open and view a capture file (.cap) in the Frame Viewer Window.
- Load a capture filter (.cf) in the Capture Filter dialog box.
- Load a display filter (.df) in the Display Filter dialog box.
- Load an address file (.adr) in the Address Database dialog box.
- Specify an executable or command file in the Capture Trigger dialog box to run when capture trigger specifications are met.

**Look in**
The drive and directory that contains the file you want to open, load, or specify.

**File name**
The name of the file you want to open, load, or specify.

**Files of type**
The type of files that you want to see in the list of files.

**See Also**
**Concepts**
Open an existing capture file
Load a capture filter
Load a display filter
Load an address file
How Network Monitor works

# Pattern Match dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Pattern Match dialog box**
Use this dialog box to specify that you want to capture only those frames that contain a specific pattern of hexadecimal or ASCII data. You can specify that the pattern can start at the beginning of the frame or that it must be a specified number of bytes (offset) away from the beginning of the frame or the end of the topology header. Specify the offset by using the following controls.

- **Pattern**

  The hexadecimal or ASCII data pattern that you want captured frames to match.

  - **Hex**

    The specified pattern is read in hexadecimal format. By default, this option is selected.

  - **ASCII**

    The specified pattern is read in ASCII format.

- **Offset (in hex)**

  The hexadecimal number that specifies where the search should begin. The way Network Monitor interprets this number depends upon whether you click the **From Start Of Frame** or the **From End Of Topology Header** option.

  - **From Start of Frame**

    The offset you specify is measured from the beginning of the frame. By default, this option is selected.

  - **From End of Topology Header**

    The offset you specify is measured from the end of the topology header.

**See Also**
**Concepts**
Specify frame data patterns to capture
Capture Filter dialog box
How Network Monitor works

# Print dialog box: Netmon tab

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Print dialog box: Netmon tab**
Use this tab to specify printing options specific to your Network Monitor capture file.

- **Output Detail**

  Specify the amount of detailed information that you want to print for each frame.

  - **Print Frame Summary Lines**

    Prints information that appears in the Frame Viewer window Summary pane. By default, this check box is selected.

    **Print Protocol Details**

    Prints information that appears in the Frame Viewer window Details pane. By default, this check box is selected.

| | |
|---|---|
| **Indent Levels** | Indents each successive protocol level when frames are printed. By default, this check box is selected. |
| **Indent Spaces per Level** | The number of spaces that you want to indent each level of detail. By default, this value is set to 0. |
| **Print Same as View** | The printed document reflects exactly the structure of the Frame Viewer window. |
| **Expand All Levels** | The printed document contains an expanded view of all of the protocols listed in the Frame Viewer window Detail pane. By default, this option is selected. |

  - **Print Hex Data**

    Prints the data that appears in the Hexadecimal pane of the Frame Viewer window. By default, this check box is selected.

| | |
|---|---|
| **Lines of Hex Data to Print** | The number of lines of hexadecimal data that you want to print. By default, this value is set to "All."<br><br>If you want to print all the data in a frame, set this option to "All." (This is useful if you are printing only a few frames.) |

- **Filtering**

  Specify whether to print only those frames that match the display filter, which is described in Filtering data.

  - **Apply Display Filter**

    Applies a display filter to the data that you want to print. By default, this check box is cleared.

  - **Edit Filter**

    Click to open the Display Filter dialog box, which you can use to edit the display filter before applying it to the data that you want to print.

- **Add Page Breaks**

  Adds page breaks to the printed material. By default, this check box is cleared.

On the **General** tab, the following options are specific to Network Monitor:

- **Print Range**

  Specify which frames that you want to print.

  Network Monitor preserves the original frame numbers. Saving the filtered frames to a capture (.cap) file causes the original frame numbering to be lost; in a capture file, frame numbers always begin with 1.

  - **All**

    All displayed frames. By default, this option is selected.

  - **Frame**

    A portion of the displayed frames. Specify the range of frames that you want to print in the **From** and **To** boxes.

**See Also**
**Concepts**
Print captured frames
Frame Viewer Window
Design a display filter

# Protocol Parsers dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Protocol Parsers dialog box**
Use this dialog box to enable or disable the parsers that read and interpret data sent to or from your computer with various protocols. When you disable a parser, Network Monitor does not attach the properties of the disabled parser to the captured data in the associated protocol. Any protocol that follows the disabled protocol in the frame will not be parsed.

- **Enabled Protocol Parsers**

  The names and descriptions of all currently enabled parsers. By default, all Network Monitor protocols are enabled.

  - **Name**

    The names of enabled protocols.

  - **Description**

    The descriptions of enabled protocols.

- **Disable**

  Click to move the protocol selected in the **Enabled Protocol Parsers** list to the **Disabled Protocol Parsers** list.

- **Disable All**

  Click to move all protocols to the **Disabled Protocol Parsers** list.

- **Enable**

  Click to move the protocol selected in the **Disabled Protocol Parsers** list to the **Enabled Protocol Parsers** list.

- **Enable All**

  Click to move all protocols to the **Enabled Protocol Parsers** list.

- **Disabled Protocol Parsers**

  All of the parsers that are loaded into the \Parsers subdirectory but that are not enabled. As long as a parser appears in this list, Network Monitor cannot interpret frames using the associated protocol.

  - **Name**

    The names of disabled protocols.

  - **Description**

    The descriptions of disabled protocols.

- **Save as Default**

  Saves the parsers that appear in the **Enabled Protocol Parsers** list as your default configuration.

  By default, this check box is cleared (Network Monitor enables all parsers in the \Parsers subdirectory).

**See Also**

**Concepts**

How Network Monitor works

Specify capture filter protocols

Specify display filter protocols

# Select a Network dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Select a Network dialog box**
This dialog box displays information about the current network configuration. Use this dialog box to:

- Find information about your local configuration.
- Use a different network adapter on the local computer to capture network data.

**Local Computer**
A tree control that displays information about one or more network adapters on the local computer. Select the network adapter for which you want to display information.

**Properties**
This control describes the following information about the network adapter you selected:

- whether the network configuration supports delayed capture.
- whether the network configuration supports Expert Stats (ESP interface).
- whether the network configuration supports real time capture.
- whether the network configuration supports network statistics.
- the type of network connection (for example, a dial-up connection).
- the manufacturer name and model information for the network adapter.
- the MAC address of your network adapter. This is the address assigned by the manufacturer of the adapter.
- the current address of your network adapter. Unless the permanent address of the adapter has been overridden, this listing is identical to the MAC address.
- the Link speed, in bits per second (bps). This is the rate at which you can transfer bits over the network. This information is supplied by the manufacturer of the network adapter.
- the type of network (for example, Ethernet) to which the network adapter can connect. This information is supplied by the manufacturer of the network adapter.
- the maximum number of bytes per frame that the selected network adapter can handle. This information is supplied by the manufacturer of the network adapter.
- whether the network adapter supports promiscuous mode.

**See Also**
**Concepts**
No Networks Found dialog box
Network Monitor overview
Network Monitor How To ...

# Frame Viewer Window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Frame Viewer window**

This section covers:

- Frame Viewer Panes
- Frame Viewer Dialog Boxes

# Frame Viewer Panes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Panes**

This section covers:

- Frame Viewer window: Summary pane
- Frame Viewer window: Detail pane
- Frame Viewer window: Hexadecimal pane

# Frame Viewer window: Summary pane

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Frame Viewer window: Summary pane**

The Frame Viewer Window, Summary pane, displays a list of the frames in the order in which they were captured.

**Frame**

The frame number, which is listed in the order that the frame was captured. Frame numbers begin with 1.

**Time**

The time that the frame was captured. The way that the capture time is expressed depends on the display options that you specify. The time is expressed as one of the following:

- The absolute system time that the frame was captured.
- The time that the frame was captured relative to the beginning of the capture. This value is expressed in microseconds.
- The time that the frame was captured relative to the capture completion time of the frame that precedes it. This value is specified in microseconds.

**Src MAC Addr**

The network address of the computer that sent the frame.

**Dst MAC Addr**

The network address of the computer that the frame was sent to.

**Protocol**

The primary protocol that was used to send the frame.

**Description**

A brief description of the contents of the frame.

**Src Other Addr**

The IP or IPX/XNS address associated with the address displayed in the **Src MAC Addr** column.

**Dst Other Addr**

The IP or IPX/XNS address associated with the address displayed in the **Dst MAC Addr** column.

**Type Other Addr**

The types of the addresses displayed in the **Src Other Addr** and **Dst Other Addr** columns.

**Note**

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

Other panes in this window: Frame Viewer window: Detail pane, Frame Viewer window: Hexadecimal pane

**See Also**

**Concepts**

Frame Viewer Window
Network Monitor overview
Frame Viewer window: Detail pane

# Frame Viewer window: Detail pane

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Frame Viewer window: Detail pane**

The Frame Viewer window Detail pane displays information about the frame that is currently selected in the Frame Viewer window: Summary pane. This information includes details about the selected frame and about the protocols that were used to send it.

You can expand or collapse the levels of detail that appear in this pane by double-clicking the lines that appear next to plus (+) or minus (-) signs.

**Frame**

The base properties of the selected frame. These properties include:

- The date and time that the frame was captured.
- The time (measured in milliseconds) that elapsed between the capture of the current frame and the capture of the frame that precedes it.
- The frame number in the sequence of total frames captured.
- The length of the frame (in bytes).
- The amount of the frame (in bytes) that was captured.
- The number of bytes in the frame not parsed by the base parser.

**Protocol**

The protocol section of the Detail pane consists of all information that appears below the end of the **Frame** section. The length of the section varies, depending on the number of protocols used to send the frame. Network Monitor displays the protocols from the lowest-level protocol to the highest-level protocol.

When you click a line in this pane, Network Monitor selects the raw representation of this line (which appears in the Hexadecimal pane) automatically.

Other panes in this window: Frame Viewer window: Summary pane, Frame Viewer window: Hexadecimal pane

**See Also**

**Concepts**

Frame Viewer Window
Display Filter dialog box
Network Monitor overview

# Frame Viewer window: Hexadecimal pane

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Frame Viewer window: Hexadecimal pane**
The Frame Viewer Window Hexadecimal pane consists of two sections.

- The hexadecimal section displays the contents of the frame in hexadecimal form.
- The ASCII section displays an alphabetical representation of the ASCII contents of the frame.

When you select a line in the Detail pane, the section of the frame that the line represents is selected automatically in both sections of the Hexadecimal pane.

Other panes in this window: Frame Viewer window: Summary pane, Frame Viewer window: Detail pane

**See Also**
**Concepts**
Frame Viewer Window
Network Monitor overview
Frame Viewer window: Summary pane
Frame Viewer window: Detail pane

# Frame Viewer Dialog Boxes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dialog boxes**
This section covers:

- Display Options dialog box
- Expression dialog box
- Expression dialog box: Address tab
- Expression dialog box: Protocol tab
- Expression dialog box: Property tab
- Font dialog box
- Go To dialog box
- Insert Comment Frame dialog box
- Label Window dialog box
- No Networks Found dialog box
- Protocol Colors dialog box
- Save dialog box
- Display Filter dialog box

# Display Options dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Display Options dialog box**
Use this dialog box to specify the time data and the protocols that you want to appear in the Frame Viewer window: Summary pane of the Frame Viewer window.

- **Time**

  Specify how you want the time to display in the **Time** column of the Summary pane.

  - **Time of day**

    Display the system time at which the frame was captured (for example, 10:35:48.358).

  - **Seconds from beginning of capture**

    Display the time (in microseconds) that elapsed between the beginning of the capture and the capture of the frame.

    By default, this option is selected.

  - **Seconds from previous frame**

    Display the time (in microseconds) that elapsed between the capture of a frame and the capture of the frame that precedes it.

- **Summary Protocol**

  Specify the protocol summary information that appears in the **Protocol** column of the Summary pane.

  - **Last protocol in frame**

    Display information about the final protocol that appears in the captured frame. By default, this option is selected.

  - **Auto (Based on protocols in display filter)**

    Display the protocol information that Network Monitor selects for you. The selection is based on the current display filter settings.

**See Also**
**Concepts**
Specify the frame time stamp
Specify the protocol header to display

# Expression dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Expression dialog box**
Use this dialog box to:

- Write or edit an expression in a display filter decision tree. For more information about display filters, see Filtering data.
- Search for a frame in the Frame Viewer Window with criteria that you provide.

The main categories that you can specify options for are represented by tabs.

- Expression dialog box: Address tab
- Expression dialog box: Protocol tab
- Expression dialog box: Property tab

**Note**

- Only one expression string at a time can be added to a display filter or be used as a search criterion.

**See Also**
**Concepts**
Specify address pairs in a display filter
Specify display filter protocols
Specify display filter protocol properties
Search for a frame by its properties

# Expression dialog box: Address tab

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Expression dialog box: Address tab**

Use this tab to specify the address you want to find or to add or edit an address expression in a display filter decision tree. The **Station 1** and **Station 2** boxes display all of the addresses contained in the currently loaded address database. For each address, the friendly name, address, type of address (protocol), and associated comments are included.

**Expression**

The address expression that you want to locate or to add to a display filter decision tree. The display changes as you select items in other areas of the dialog box. When you click **OK**, the displayed expression is added to the display filter decision tree or used as the search criteria for frames. Any other changes you make are lost.

The **Expression** text box appears regardless of which tab you select. You cannot edit the expression directly.

**Station 1 (and Station 2)**

The computer addresses between which you want to monitor traffic or that you want to locate. If you do not specify the local computer as either Station 1 or Station 2, no data will appear when you apply the filter.

**Direction**

An arrow that indicates the traffic direction that you want to monitor or locate. For example, to identify traffic that travels from Computer 1 to Computer 2, select the Computer 1 name, the **-->** symbol, and then the Computer 2 name.

**Edit Addresses**

Click to open the Address Database dialog box, which you can use to edit the selected address.

Other tabs in this dialog box: Expression dialog box: Protocol tab, Expression dialog box: Property tab

**Note**

- You should specify BROADCAST, MULTICAST, and FUNCTIONAL only as destination addresses. If you construct a statement in which the source address is either broadcast or multicast, you filter out all frames. Broadcast and multicast are always destination addresses.

**See Also**

**Concepts**

Expression dialog box

Display Filter dialog box

Open an existing capture file

# Expression dialog box: Protocol tab

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Expression dialog box: Protocol tab**

Use this tab to specify the protocols that you want to find in the Frame Viewer window or to specify protocols in the display filter decision tree.

By default, all protocols are enabled, and the **Expression** box shows **Protocol == Any**. The protocols listed in this tab are those whose parsers are currently available to Network Monitor.

**Expression**

The address expression to search for in frames or to add to a display filter decision tree. As you select items in other areas of the dialog box, the display changes. When you click **OK**, the displayed expression is added to the display filter decision tree or is used as the frame search criterion. Any other changes you make are lost.

The **Expression** box appears regardless of which tab you select. You cannot edit the expression directly.

**Enabled Protocols**

The protocols to include in the display filter or to use to search for a frame.

**Disabled Protocols**

The protocols excluded from the display filter, which will not be used to search for a frame. If there are protocols in the frame in addition to the disabled protocols, then the frame will still be found. For example, if you disable the SMB protocol, SMB frames will be found if they contain enabled protocols.

**Disable**

Click to move selected protocols to the **Disabled Protocols** box.

**Enable**

Click to move the selected protocols to the **Enabled Protocols** box.

**Disable All**

Click to move all protocols to the **Disabled Protocols** box.

**Enable All**

Click to move all protocols to the **Enabled Protocols** box.

Other tabs in this dialog box: Expression dialog box: Address tab, Expression dialog box: Property tab

**Note**

- To exclude all frames that contain the SMB protocol, create a line that states **Protocol == SMB** and put a "**Not**" in front of it.

**See Also**

**Concepts**

Expression dialog box

Display Filter dialog box

Open an existing capture file

# Expression dialog box: Property tab

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Expression dialog box: Property tab**

Use this tab to specify which protocol properties you want to find or to add or edit a protocol property expression in the display filter decision tree.

**Expression**

The address expression that you want to locate or to add to a display filter decision tree. As you select items in other areas of the dialog box, the display changes. When you click **OK**, the displayed expression is added to the display filter decision tree or is used as the frame search criterion. Any other changes that you make are lost.

The **Expression** box appears regardless of which tab you have selected. You cannot edit the expression directly.

**Protocol:Property**

The names and properties of protocols available to Network Monitor. If a plus (+) sign appears next to a protocol name, you double-click the sign to expand the protocol and select a property.

**Relation**

A relational operator used to specify the connection between the protocol property and its possible values. The relational value can be one of the following:

**==**
Equal to

**<>**
Not equal to

**>**
Greater than

**<**
Less than

**>=**
Greater than or equal to

**<=**
Less than or equal to

**contains**
Contains (value)

**exists**
Exists (value)

**Value**

The possible values that the highlighted property is compared to.

When the **Hex** and the **Decimal** or **Ascii** options are visible, you can specify whether the number in the **Value** box is displayed in hexadecimal, decimal, or ASCII form.

Other tabs in this dialog box: Expression dialog box: Address tab, Expression dialog box: Protocol tab

**See Also**

**Concepts**

Expression dialog box
Specify capture filter protocols
Open an existing capture file

# Font dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Font dialog box**

Use this dialog box to change the fonts used to display or print captured frame data.

**Font**

The fonts you can use to display or print captured frame data. You might prefer a monospace font such as Courier or Fixedsys rather than a proportional font such as MS Sans Serif. Depending on your printer configuration, some fonts might not be available for printing. By default, the font is Courier New.

**Font Style**

The font style, such as bold or italic, that you use to display or print captured frame data. By default, the font style is Regular.

**Size**

The type size that gives you the best combination of readability and quantity of data per screen. By default, the type size is 8 points. (Other ways to fit data onto the screen include maximizing the window and rearranging columns.)

**See Also**

**Concepts**

Set the frame data display font

Capture network frames

Network Monitor overview

# Go To dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Go To dialog box**
Use this dialog box to open a specific frame in the currently displayed capture file.

### Frame Number

The number of the frame that you want to highlight.

### Note

- The second field in the Frame Viewer window status bar lists the number of the frame that is currently highlighted, followed by a slash mark (**/**) and the total number of frames in the current capture file.

**See Also**
**Concepts**
View a specific frame
Display Filter dialog box
How Network Monitor works

# Insert Comment Frame dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Insert Comment Frame dialog box**

Use this dialog box to add a frame that contains comments and other information into a capture file. The comment frame contains the Trail protocol, which includes statistics such as the number of frames and bandwidth consumed. These statistics are calculated for a block of frames bounded by the current comment frame and the next comment frame of the same type (or by the end of the capture).

**Frame Number**

The frame position in the capture where Network Monitor inserts the comment frame. By default, the frame number is the current location. Network Monitor places the inserted frame in front of the frame number specified and increases the other frame numbers by one.

The maximum frame number ranges from one more than the number of the final frame in the capture.

**Type of Frame to Insert**

The parser that processes the comment frame after the Trail protocol runs.

If you modify the %netmon%\Parsers\Trail.ini file and map a parser to an integer, you can type that number in this field.

**No Statistics**

Disables the generating of statistics for the comment frame. When **No Statistics** is selected, Network Monitor does not use the comment frame to define a block of any other inserted comment frames for statistics. By default, this check box is selected.

**Apply Current Filter to Statistics**

Uses the current display filter to calculate statistics. By default, this check box is selected.

If you clear this check box, the output shows data for all frames in the capture, including those that have been filtered from the display.

**Enter in a comment for this new frame**

The text that you attach to the frame by using the protocol specified in **Type of Frame to Insert**.

**See Also**

**Concepts**

Add a comment frame to a capture
Design a display filter
Load a capture filter

# Label Window dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Label Window dialog box**

Use this dialog box to change the text that appears in the Frame Viewer Window title bar.

### Description

The new window label.

**See Also**

**Concepts**

Change the Frame Viewer window label

Design a display filter

Network Monitor overview

# No Networks Found dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**No Networks Found dialog box**
This dialog box appears when Network Monitor cannot find installed network drivers on the local computer or if the Network Monitor driver is not installed.

**Do not notify me of this in the future**
Hides this message box the next time you open Network Monitor. Select this check box if you plan to work only with existing Network Monitor capture files. By default, this check box is cleared.

If you receive this message, Install the Network Monitor driver.

**See Also**
**Concepts**
Select a Network dialog box
Install the Network Monitor driver
Network Monitor overview

# Protocol Colors dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Protocol Colors dialog box**
Use this dialog box to specify colors that appear in the panes in the Frame Viewer window. You can use the colors to highlight frames, properties, and hexadecimal data pertaining to that protocol.

- **X**

  Selected protocols apply the current color selection. Click anywhere in the protocol line to select or clear it.

- **Name**

  The names of currently supported protocols.

- **Description**

  A description of each supported protocol.

- **Colors**

  Specifies the colors of the selected protocols.

  - **Foreground**

    The color to apply to the protocol text.

  - **Background**

    The color to apply to the text background.

- **Select All**

  Click to select all listed protocols.

- **Clear All**

  Click to clear all listed protocols.

  This button does not reset the background and foreground colors.

- **Save as Default**

  Automatically uses the specified color combination every time you display data in the Frame Viewer window. By default, this check box is cleared.

**See Also**
**Concepts**
Select a protocol color scheme
Change the Frame Viewer window label
Specify capture filter protocols
Frame Viewer Window

# Save dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Save dialog box**

Use this dialog box to:

- Save captured data in a file (.cap) that you can view in Network Monitor at a later time.
- Save a capture or display filter that you have designed or edited.
- Save an address database to a file.
- Save a text file to which you can print frame data.

**Save in**

The drive and directory that you want to save the file to.

**File name**

The name to save the file as.

**Save as type**

The file format of the information that you want to save. Network Monitor uses these default file types and file name extensions:

| File type | File name extension |
|---|---|
| Capture files | .cap |
| Capture filters | .cf |
| Display filters | .df |
| Address files | .adr |
| Print files | .txt |

**Range**

In the **From** and **To** boxes, specify the numerical range of the frames that you want to save in the capture file.

When you save a range of frames to a capture file, the numbers associated with the frames are changed. In a capture file, frame numbers always begin with 1, regardless of the number associated with the original frame. Similarly, if you apply a display filter and then save the filtered frames, the frame numbers in the capture file begin with 1. If, on the other hand, you select **Print To File** in the **Print** dialog box, the original frame numbers associated with the frames are preserved.

**Filtered**

Applies the current display filter to the capture data when the data is being saved. When you select this check box, the saved frames are renumbered so that no gaps appear in the sequence. By default, this check box is cleared.

**Note**

- The **From**, **To**, and **Filtered** options are available only when you are saving a capture file from the Frame Viewer window.

**See Also**

**Concepts**

Save captured frames to a file
Save a capture filter
Save a display filter
Save the address database to a file

# Display Filter dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Display Filter dialog box**

Use this dialog box to limit your display to frames that have the properties that you specify. For example, you can specify the protocols, computer addresses, and protocol properties that you want to display in the Frame Viewer Window. The following controls help you to make changes to the display filter decision tree.

- **Protocol ==**

  The decision tree branch that lists the protocols that you want to display. By default, this branch is set to **Protocol == Any**. To modify the protocol list, double-click the branch to open the Expression dialog box: Protocol tab.

- *address* **<->** *address*

  The decision tree branch that lists the computer address pairs that you want to display. By default, this branch is set to **Any <-> Any**. To modify the address pairs list, double-click the branch to open the Expression dialog box: Address tab. You must specify the local computer as the data source or destination.

- **Add**

  Decision tree criteria.

  - **Expression**

    Click to open the Expression dialog box, which enables you to specify the computer address pairs, protocols, and protocol properties that you want to display.

- **Insert**

  Logical operators for the decision tree. You can use as many as 4,000 decision tree operators.

  - **AND**

    Click to add an AND branch to the currently selected branch of the decision tree.

  - **OR**

    Click to add an OR branch to the decision tree.

  - **NOT**

    Click to add a NOT branch to the decision tree or to negate a branch of the tree.

- **Edit Expression (Change Operator)**

  If you select an expression in the decision tree, the label on the button in this group is **Edit Expression**. Click this button to open the Expression dialog box, which you can use to edit the selected data.

  If you select a logical operator in the decision tree, the label on the button in this group is **Change Operator**. Click this button to cycle through the logical operators for the selected line.

- **Delete**

  Decision criteria that you want to delete from the tree.

- **Line**

  Click to delete the selected line in the decision tree.

- **Branch**

  Click to delete the currently selected branch in the decision tree.

- **All**

  Click to delete the entire decision tree.

- **Load**

  Click to open the **Open** dialog box, which you can use to replace the current display filter with a display filter that has previously been saved to a file.

- **Save**

  Click to open the **Save Display Filter** dialog box, which you can use to save the current display filter as a file (.df). To open this display filter again later, you can click the **Load** button to open the file.

**See Also**
**Concepts**
Design a display filter
Save a capture filter
Save a display filter
Load a display filter
Filtering data

# Network Monitor Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

For more information about Network Monitor, consult the following resources:

- Windows Server 2003 Resource Kit at the Microsoft Windows Resource Kits Web site.
- Microsoft TechNet at the Microsoft Web site.
- Microsoft Product Support Services at the Microsoft Web site.

# Troubleshooting Network Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Network Monitor will not start.
- Not enough disk space to create a capture buffer.
- Access denied to default capture directory on computers with lockdown policies in place.
- "No NPPs found" error message.

**Network Monitor will not start.**
**Cause: Authenticated Users** has been manually removed from the **Users** group.

**Solution:** Add the Network Monitor user to the **Users** group following the steps described in Add a member to a local group. After the user has been added, they must log-off the computer, and then log back on.

**Cause:** The Discretionary Access Control List (DACL) of the system directory is changed to disallow normal user's access.

**Solution:** Add the Network Monitor user to the DACL of the system directory and assign permissions to allow normal user access by following the steps described in Assign, change, or remove permissions on Active Directory objects or attributes.

**Not enough disk space to create a capture buffer.**
**Cause:** When you set a capture buffer, Network Monitor reserves disk space equal to the size of the capture buffer. If not enough disk space is free, you will get an error.

**Solution:** Make sure that the temporary capture directory has enough free disk space (at least 1 MB free + buffer size).

**Access denied to default capture directory on computers with lockdown policies in place.**
**Cause:** Lockdown computers will not allow you to write to files in the \Windows folder and subfolders, but the default folder for Network Monitor capture files is a subfolder of the \Windows folder.

**Solution:** After you create a capture, save the capture into a folder to which you have read and write access. This folder becomes the new default location for saving capture files.

**"No NPPs found" error message.**
**Cause:** Network Monitor, the Network Monitor driver, or neither is installed properly.

**Solution:** Make sure that the Network Monitor components are installed properly.

1. If the Network Monitor driver is not installed, install it using the steps described in Install the Network Monitor driver.
2. If the Network Monitor driver is installed, delete it and then reinstall it.
3. If reinstalling the Network Monitor driver does not solve the problem, remove Network Monitor and then reinstall it using the steps described in Install Network Monitor.

# Monitoring Performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring performance**

System Monitor and Performance Logs and Alerts support detailed monitoring of the utilization of operating system resources.

- To find features that have been changed in the Windows Server 2003 family, see New ways to do familiar performance monitoring tasks.
- For tips about monitoring performance, see Performance monitoring best practices.
- For help with specific tasks, see Performance monitoring How To....
- For general background information, see Performance monitoring Concepts.
- For problem solving instructions, see Troubleshooting Performance monitoring.

# New ways to do familiar performance monitoring tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for System Monitor and Performance Logs and Alerts (hosted in the Performance console, named Perfmon.msc). The user interface for performing these tasks is different in the Windows Server 2003 family and Windows 2000 from the way it was in Windows NT version 4.0. Many of the toolbar buttons remain the same among the three operating systems and can be used instead of the steps listed in this table.

| **If you want to** | **In Windows NT 4.0 use** | **In Windows 2000 or the Windows Server 2003 family use** |
|---|---|---|
| Create a chart of performance data: See Work with Monitoring Views. | Performance Monitor on the **Administrative Tools** menu. | Open System Monitor. |
| Create a report of performance data: See Change from graph view to histogram or report views. | Performance Monitor on the **Administrative Tools** menu. | Open System Monitor. |
| Create a log: See Create a counter log or Create a trace log. | Performance Monitor on the **Administrative Tools** menu. | Open Performance. |
| Create an alert: See Create an alert. | Performance Monitor on the **Administrative Tools** menu. | Open Performance. |

# Performance monitoring best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Set up a monitoring configuration.** Configure Performance Logs and Alerts to report data for the recommended counters at regular intervals, such as every 10 to 15 minutes. Retain logs over extended periods of time, store data in a database, and query the data to report on and analyze the data as needed for overall performance assessment, trend analysis, and capacity planning.

  For best results, do the following before starting System Monitor or Performance Logs and Alerts on the computer you want to monitor for diagnostic purposes:

  - If the server in question has halted or is not responding, run System Monitor from another computer.

- **Keep monitoring overhead low.** In general, the performance tools are designed for minimal overhead. However, you might find that the overhead increases under the following conditions:

  - You are running System Monitor in graph view.
  - You have selected an option other than the default (current value) for either the System Monitor graph or report views.
  - You are sampling at very frequent intervals (less than three seconds apart).
  - You have selected many different objects and counters.

- Other aspects of performance tool operation that affect performance include file size and disk space taken up by log files. To reduce file size and related disk space usage, extend the update interval. Also, log to a disk other than the one you are monitoring. Frequent logging also adds demand on disk input and output (I/O).

  If monitoring overhead is a concern, run only the Performance Logs and Alerts service and do not monitor using a System Monitor graph.

  During local logging of remote counters, frequent updating can slow performance due to network transport. In this case, it is recommended that you log continuously on remote computers but upload logs infrequently -- for example, once a day.

- **Analyze performance results, and establish a performance baseline.** Review logged data by graphing it using the System Monitor display or exporting it for printing. Compare the values against the counter thresholds shown in Analyzing performance data to verify that resource usage or other activity is within acceptable limits. Set your baseline according to the level of performance that you consider satisfactory for your typical workload.

- **Set alerts.** Set alerts according to the counter values you consider to be unacceptable, as defined by baseline evaluation.

- **Tune performance.** Tune system settings and workload to improve performance and repeat monitoring to examine tuning results. For more information on ways to improve performance, see Solving performance problems.

- **Plan ahead.** Monitor trends for capacity planning and add or upgrade components as needed. Maintain logged data in a database and observe changes to identify changes in resource requirements. After you observe changes in activity or resource demand, you can identify areas that might require additional resources.

# Performance monitoring How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Work with Counters
- Work with Monitoring Views
- Work with Data
- Work with Settings
- Create and Configure Counter Logs
- Create and Configure Trace Logs
- Create and Configure Alerts
- Work with Logs and Alerts

# Work with Counters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with counters**

- Add counters to System Monitor
- Delete counters from System Monitor
- Get counter descriptions

# Add counters to System Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add counters to System Monitor**

1. Open Performance.
2. Right-click the System Monitor details pane and click **Add Counters**.
3. To monitor any computer on which the monitoring console is run, click **Use local computer counters**.

   Or, to monitor a specific computer, regardless of where the monitoring console is run, click **Select counters from computer** and specify a computer name or IP address.
4. In **Performance object**, click an object to monitor.
5. To monitor all counters, click **All counters**.

   Or, to monitor only selected counters, click **Select counters from list** and select the counters you want to monitor.
6. To monitor all instances of the selected counters, click **All instances**.

   Or, to monitor only selected instances, click **Select instances from list** and select the instances you want to monitor.
7. Click **Add**, and then click **Close**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, the Performance Log Users group, or the Performance Monitor Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- By default, the following three counters are displayed in System Monitor -- Memory\Pages/sec, PhysicalDisk(_Total)\Avg. Disk Queue Length, and Processor(_Total)\%Processor Time.
- If you select an object on a remote computer, you might notice a short delay as System Monitor refreshes the list to reflect objects present on that computer.
- Unless you are configuring data from a log, you can select only active instances in the list box.
- Some object types have several instances. For example, if a system has multiple processors, the Processor object type will have multiple instances. The Physical Disk object type has two instances if a system has two disks. Some object types such as Memory and Server have only a single instance. If an object type has multiple instances, you can add counters to track statistics for each instance, or in many cases, for all instances at once.
- To monitor the sum of the values all of a particular counter's instances, select the _Total instance.
- By default, process and thread instances have a #n appended (where n is an integer). To append the process ID instead, create a DWORD value called ProcessNameFormat under the

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PerfProc\Performance subkey and set its value to 2. Likewise, to append the thread ID instead, create a DWORD value called ThreadNameFormat under the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PerfProc\Performance subkey and set its value to 2. To turn off any of these options and return to the default, set the value of the appropriate DWORD to 1.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

- By default, counters are shown with both the instance name and an instance index. To turn off this feature, right-click the details pane, click **Properties**, select the **General** tab, and clear the **Allow duplicate counter instances** check box.

- When you select two instance indexes of the same counter, for example, if you are monitoring multiple threads of a process, note that the instance index number assigned to that instance changes over time as the instance, such as a thread, starts and stops. The index number does not necessarily correspond to the same thread over the life of the process.

- When creating a monitoring console for export, make sure to select **Use local computer counters**. Otherwise, System Monitor will obtain data from the computer named in the text box, regardless of where the console file is installed.

- For a description of a particular counter, click the name of the counter in **Add counters**, and then click **Explain**.

- You can change the default display characteristics assigned to a counter such as the color, line style and width, and graph scale by using the **Data** tab in the **System Monitor Properties** dialog box. To access the **Data** tab, right-click the details pane, and click **Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor a different computer
Delete counters from System Monitor
Quick reference to System Monitor settings

# Delete counters from System Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete counters from System Monitor**

1. Open Performance.
2. In the System Monitor details pane, click the name of the counter in the legend.
3. Press DELETE.

## Notes

To perform this procedure, you must be a member of the Administrators group, the Performance Log Users group, or the Performance Monitor Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- You can also delete counters by using the **Remove** button on the **Data** tab in the **System Monitor Properties** dialog box. The **Remove** button is helpful for deleting a group of counters at once.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add counters to System Monitor
Quick reference to System Monitor settings
Change counters and counter properties

# Get counter descriptions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To get counter descriptions

1. Open Performance.
2. Right click the System Monitor details pane and click **Add Counters**.
3. In **Performance object**, click an object.
4. In **Select counters from list**, click the counter for which you want information.
5. Click **Explain**, and then click the Close button.

### Notes

- To perform this procedure, you must be a member of the Administrators group, the Performance Log Users group, or the Performance Monitor Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Work with Monitoring Views

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with monitoring views**

- Use logged data
- Change counters and counter properties
- Change sampling options
- Monitor a different computer
- Highlight counter data
- Change from graph view to histogram or report views
- Add titles, grids, and other attributes to a graph
- Change background, chart, text, and font properties

# Use logged data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use logged data**

1. Open Performance.
2. Right-click the System Monitor details pane and click **Properties**.
3. On the **Source** tab under **Data Source**, click **Log files**, and then add the file(s) you want to use as a data source.

   To specify a SQL server data source, click **Database**.
4. Click **Time Range**.
5. Click the **Data** tab and click **Add** to open the **Add Counters** dialog box. The counters you selected during log configuration appear.

## Notes

- You do not need administrative credentials to view counters from log files on the local computer.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- Unless you specifically want to monitor start-up events, you should exclude times that include such events from your time window because these temporary high values tend to skew overall performance results.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set parameters for a log
Performance Logs and Alerts interface

# Change counters and counter properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change counters and counter properties**

1. Open Performance.
2. Right-click the System Monitor details pane and click **Properties**.
3. On the **Data** tab, specify the options you want to use:
   - **Add.** Opens the **Add Counters** dialog box, where you can select other counters to add.
   - **Remove.** Removes the counter selected in the counter list.
   - **Color.** Changes the color of the selected counter.
   - **Scale.** Changes the displayed scale of a selected counter in the graph or histogram view. Counter values can be scaled linearly from .0000001 to 1000000.0. You can adjust the counter scale settings to enhance the visibility of counter data in the graph. Changing the scale does not affect the statistics displayed in the value bar.
   - **Width.** Changes the line width of a selected counter. Note that defining a line width can determine the line styles that are available.
   - **Style.** Changes the line style for a selected counter. You can select a style only when you are using the default line width.
4. Click **Apply** and then click **OK**.

## Notes

- If you want to view counters in real-time mode, or set up a monitoring configuration, you must be a member of the Administrators group, the Performance Log Users group, or the Performance Monitor Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- If you want to simply work with data from log files, you do not need administrative credentials on the local computer.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- By default, System Monitor displays Pages/sec, Avg. Disk Queue Length, and % Processor Time counter data for the local computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Add counters to System Monitor
Delete counters from System Monitor
Quick reference to System Monitor settings
Change background, chart, text, and font properties

# Change sampling options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change sampling options**

1. Open Performance.
2. Right-click the System Monitor details pane and click **Properties**.
3. On the **General** tab, specify the sampling option you want:
   - To sample automatically at regular intervals, select the **Sample automatically every** check box and in **seconds** type the length of the interval in seconds. The default interval is one second.
   - To sample manually, clear the **Sample automatically every** check box.

     When manual sampling is selected, click the Update Data button to collect a sample.
4. Click **Apply**, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group, the Performance Log Users group, or the Performance Monitor Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Quick reference to System Monitor settings
Add counters to System Monitor
Delete counters from System Monitor

# Monitor a different computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To monitor a different computer**

1. Open Performance.
2. Right-click the System Monitor details pane and click **Add Counters**.
3. Click **Select counters from computer**, and in the text box, type the name of the computer that you want to monitor (*computer_name*), or enter the IP address of the computer you want to monitor.
4. In **Performance object**, click an object to monitor.
5. In **Select counters from list**, click one or more counters, or click **All counters** to use all counters for that object.
6. To use all counter instances available when System Monitor is in use, click **All instances**.

   Or, to use only selected counter instances, click **Select instances from list**, and select the instances you want to use.
7. Click **Close**.

## Notes

- You do not need administrative credentials to simply view counters on a remotely managed computer unless you use Perfmon with the /WMI option. To use the Windows Management Instrumentation (WMI) interface, you must be a member of the Administrators group on the remotely managed computer.
- To monitor (rather than simply view) counters on remotely managed computers, you must be a member of the Administrators group, the Performance Log Users group, or the Performance Monitor Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- If you select an object on a remote computer, you might notice a short delay as System Monitor refreshes the list to reflect objects present on that computer.
- If you do not see a counter that you want to monitor, it could be because the service or feature that provides the counter object is not installed or enabled on that computer and you might need to add it before you can monitor related data.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Monitor interface
System Monitor buttons and keyboard shortcut reference

# Highlight counter data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To highlight counter data**

1. Open Performance.
2. In the legend, click the counter you want to highlight.
3. Click the Highlight button, or press CTRL+H.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- When you highlight a counter, a thick line replaces the colored chart line. For white or light-colored backgrounds (defined by the Graph Background property), this line is black; for other backgrounds, this line is white.
- When using the System Monitor control in an application that supports Active-X controls, note that the CTRL+H combination used for System Monitor highlighting may conflict with the default key settings in the application. You may need to change the default key settings to support highlighting when you use the System Monitor control (Sysmon.ocx in the systemroot\System32 folder) in other applications.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Monitor interface
System Monitor buttons and keyboard shortcut reference

# Change from graph view to histogram or report views

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change from graph view to histogram or report views**

1. Open Performance.
2. On the toolbar, click the View Histogram button to view a histogram or bar graph.

   Click the View Report button to view a report.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- You can also change the display by clicking the View Log Data button to open the **System Monitor Properties** dialog box. Click the **General** tab, and use the **View** setting to specify your display preference.
- Graph view displays counter data over a given time period in line graph format. The histogram view displays counter data in bar chart format, showing only a single value per counter instance. Report view displays counter data in table format, also showing only a single value per counter instance. In report view, counter names and data values appear in rows beneath their associated performance objects, and a separate column displays each instance and its data.
- For both the histogram and report views, choosing average, minimum, or maximum statistics causes the statistic to be calculated at each sample interval. This incurs some additional performance overhead for realtime data only.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add counters to System Monitor
Delete counters from System Monitor
Quick reference to System Monitor settings

# Add titles, grids, and other attributes to a graph

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add titles, grids, and other attributes to a graph**

1. Open Performance.
2. Right-click the System Monitor details pane and click **Properties**.
3. On the **Graph** tab, specify the graph options you want to use:
   - To change the grid color, use the **Appearance** tab.
   - To add or remove the legend, statistics, or the toolbar, use the **General** tab.
   - To change border style or appearance, use the **General** tab.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change background, chart, text, and font properties
Quick reference to System Monitor settings

# Change background, chart, text, and font properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change background, chart, text, and font properties**

1. Open Performance.
2. Right-click the System Monitor details pane and click **Properties**.
3. Click the **Appearance** tab.
4. To change the color elements, in **Color**, select the graph element for which you want to change the color. To change the color used for the selected graph element, click the **Change** button.

   Or, to change the font used for the text or numbers on your graph, in **Font**, click the **Change** button. In **Font**, set any font options you want. The options include **Font**, **Font style**, **Size**, and **Script**.
5. The available elements in **Color** are:

   - **Graph Background**. The background color of the window area where counter data is charted.
   - **Control Background**. The background color that surrounds the window area where data is charted.
   - **Text**. The text color.
   - **Grid**. The color used for vertical and horizontal grid lines.
   - **Time Bar**. The color used for the timer bar.
6. Click **Apply**, and then click **OK**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Quick reference to System Monitor settings
System Monitor interface

# Work with Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with data**

- Collect performance data using Microsoft SQL Server
- Collect performance data using WMI
- Copy counter data between instances of System Monitor
- Print performance data
- Save graph settings as an HTML page
- Insert the System Monitor control in a Web page
- Save data to a log file

# Collect performance data using Microsoft SQL Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Collect performance data using Microsoft SQL Server**

- Set up a user account on a SQL server
- Set up the System DSN on each test computer

# Set up a user account on a SQL server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up a user account on a SQL server**

1. Open Enterprise Manager.
2. In the console tree, double-click **Microsoft SQL Servers**, click **SQL Server Group**, and then **SQLComputerName (Windows NT)**.
3. Double-click **Databases**, and then double-click your database of choice.
4. Right-click **Users**, and then click the **New Database User** menu.
5. Next to **Login name**, type:

   *domain**username*
6. Next to **User name**, type your user name.
7. Select the **Public** check box, and then select all of the following objects:

   - **db_owner**
   - **db_accessadmin**
   - **db_securityadmin**
   - **db_ddladmin**
   - **db_datareader**
   - **db_datawriter**
   - **db_backupoperator**

   **Important**

   - Do not select **db_denydatareader** or **db_denydatawriter**. These objects allow members to deny Read or Write permissions to data in the database.

**Notes**

- To open Enterprise Manager, click **Start**, point to **Programs**, click **Microsoft SQL Server**, and then click **Enterprise Manager**.
- Before you begin, confirm that a domain and your administrative credentials are set up:
  - You have permissions to read from and write to the SQL Server database.
  - Your user name has administrative credentials on each of the above computers; you will use it to log on at all

times.

- This is a quick start procedure only. Consult your Microsoft SQL Server documentation for detailed instructions on setting up and configuring a SQL database.

- Windows Server 2003 family, not Microsoft SQL Server, authentication is used when you connect to a SQL server for performance monitoring purposes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up the System DSN on each test computer
Create and Configure Counter Logs

# Set up the System DSN on each test computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up the System Data Source Name (DSN) on each test computer**

1. Open ODBC Data Source Administrator.
2. Click the **SystemDSN tab**.
3. Click **Add**.
4. In the **Create New Data Source** wizard, click **SQL Server**, and then **Finish**.
5. In the **Create a New Data Source to SQL Server** dialog box, fill in the **Name** and **Description** text boxes. In the **Server** box, type or select the name of the computer.

   **Important**

   - If you previously configured an **ODBC data source** for the same SQL server, a dialog box may appear which reads, "A new data source will be created with the following configuration." Click **Finish** if you do not want to change databases. Otherwise, follow steps 6 through 9.

6. Click **Next** to accept the following default settings:

   - **With Windows NT authentication using the network login ID.**
   - **Connect to SQL Server to obtain default settings for the additional configuration options.**

7. Click **Next** a second time to accept the following default settings:

   - **Use ANSI quoted identifiers.**
   - **Use ANSI nulls, paddings and warnings.**

   **Important**

The **Change the default database to** check box should be cleared. If, when setting up a user account, you selected a database other than **master**, you should now select the **Change the default database to** check box, and then click that database in the list.

1. Click **Finish** to accept the default setting, **Perform translation for character data.**
2. In the **ODBC Microsoft SQL Server Setup** dialog box, click **Test Data Source**, and verify that the "Tests completed successfully!" message appears.

**Notes**

- Before you begin, confirm that your computer is connected to the domain, and that you are logged into the domain as a user. See Related Topics for more information.
- To open Data Sources (ODBC), click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-

click **Data Sources (ODBC)**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up a user account on a SQL server
Create and Configure Counter Logs

# Collect performance data using WMI

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To collect performance data using WMI**

- Open Command Prompt and type:
- **perfmon.exe /wmi**

## Notes

- To perform this procedure, you must be a member of the Administrators group, the Performance Log Users group, or the Performance Monitor Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- When this option is specified, the program obtains performance data by means of Windows Management Instrumentation (WMI) instead of the registry. WMI Hi-Perf providers can then expose performance counters, which in turn appear in System Monitor or are logged by the Performance Logs and Alerts Service. For Windows XP and the Windows Server 2003 family, all WMI counters now appear through the use of a reverse adapter regardless of whether you use the **/wmi** option.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Use a Windows NT Server 4.0 Performance Monitor settings file
Perfmon

# Copy counter data between instances of System Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy counter data between instances of System Monitor**

1. Open Performance.
2. In the System Monitor details pane, click the Copy Properties button on the toolbar.
3. Open another instance of System Monitor and click the Paste Counter List button on the toolbar to load the counter path data from the Clipboard into the current window.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System Monitor interface
System Monitor buttons and keyboard shortcut reference

# Print performance data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print performance data**

1. Open Performance.
2. Press ALT+PRINT SCREEN to copy a view of the active window.
3. To open **Paint**, click **Start**, point to **All programs**, point to **Accessories**, and then click **Paint**.
4. On the **Edit** menu, click **Paste**.

   If prompted with the alert message "The image in the Clipboard is larger than the bitmap. Would you like the bitmap enlarged?", click **Yes**.
5. On the **File** menu, click **Print**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- For the best quality prints, choose one of the following options:
  - Add the System Monitor control to a Microsoft Office application such as Microsoft Word or Microsoft Excel and print from that program.
  - Save the System Monitor control data as an HTML file and print it from Internet Explorer or another program using **Save As**.
  - Import a log file in comma-separated (.csv) or tab-separated (.tsv) format into a Microsoft Excel spreadsheet and print from that application.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Insert the System Monitor control in a Web page
Save graph settings as an HTML page

# Save graph settings as an HTML page

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save graph settings as an HTML page**

1. Open Performance.
2. Add counters to a graph.
3. Right-click the displayed graph, and click **Save As**.
4. In the **Save As** dialog box, select a path and type the name of the new HTML file.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- The System Monitor configuration is saved to the HTML file, including View type (graph, histogram, report), graph title, and so on.
- You can apply the counter settings that are saved as an HTML file by dragging the HTML file from Windows Explorer onto a System Monitor graph.
- You can print a static System Monitor graph by opening the HTML file in a browser. The graph will display the data at the time the HTML file was saved. If you click on the graph, it will become live.
- You can use the Copy Properties button in the toolbar to transfer all System Monitor settings to the Windows Clipboard as a System Monitor HTML object. You can paste this HTML object into any HTML page.
- Use the Paste Counter List button on the toolbar (or drag and drop) to transfer the counter settings from the saved System Monitor object configuration in the Clipboard or HTML file to the System Monitor.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add counters to System Monitor
Copy counter data between instances of System Monitor

# Insert the System Monitor control in a Web page

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To insert the System Monitor control in a Web page**

1. Open Performance.
2. In the System Monitor toolbar, click the Copy Properties button.
3. Open your Web page in Notepad or a different text editor.
4. Insert the cursor at the point you want the control to appear, right-click, and select **Paste**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- When you open your Web page in the browser, the System Monitor graph will appear static. It represents the performance data as it appeared the moment you copied the graph. To monitor performance in real time, go to the toolbar and click the Freeze Display button.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add titles, grids, and other attributes to a graph
Change background, chart, text, and font properties

# Save data to a log file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save data to a log file**

1. Open Performance.
2. Right-click the System Monitor details pane and click **Save Data As**.
3. Type the name of the saved log file and select the format.
4. Enter the data filter (value *n*) for the saved log file in **Reduce log file size by only writing one out of every*n*original log records**.

   By default, *n* is set to 1, and the saved log file contains the same number of data points as the original log file.
5. Click **Save**.

## Notes

- To perform this task, you do not need administrative credentials on the local computer.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- To select a log file or a SQL server as the data source, see Related Topics.
- The **Save Data As** feature only works on log files and not on real-time data.
- Only the counters that are currently loaded from the log and are displayed in the current view will be relogged, or resampled.
- Although the extension for a log file is .blg, the file format is changed in the Windows Server 2003 family and therefore cannot be read on an earlier version of the operating system. For example, if you copy a binary file from Windows Server 2003, Standard Edition, to Windows 2000, you will be unable to read it on the latter system.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add counters to System Monitor
Use logged data
Copy counter data between instances of System Monitor
Print performance data
Save graph settings as an HTML page

# Work with Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with settings**

- Save a monitoring console
- Save log or alert settings as an HTML page
- Create a custom monitoring console
- Use a Windows NT Server 4.0 Performance Monitor settings file

# Save a monitoring console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save a monitoring console**

1. Open Performance.
2. Using System Monitor or a tool under Performance Logs and Alerts, configure the tool with the settings you want.
3. On the **File** menu, click **Save As**, and then type the name you want for this console. The default directory for storing this file is System32 on the root volume.
4. Click **Save**.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- Only modifications to the System Monitor are saved.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add counters to System Monitor
Create a counter log
Create an alert

# Save log or alert settings as an HTML page

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save log or alert settings as an HTML page**

1. Open Performance.
2. Create a log or an alert.
3. In the details pane, right-click the name of the log or the alert and click **Save Settings As**. By default, the settings are saved in the user's personal default directory; this is typically My Documents.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- As a security precaution, the RunAs user name and password are not copied to the HTML file when you save a log or alert configuration to it.
- A second security feature is that you can use this HTML settings file to create a new log or alert, not to modify the settings of an existing log or alert. In the existing log or alert file, right-click the details pane, and click **New Log Settings From** or **New Alert Settings From**. This is an easy way to generate log settings from an alert configuration.
- For counter logs and alerts, you can open the HTML page in a browser, displaying a System Monitor object. The System Monitor object will specify the current log file as the data source.

  Or, you can drag the HTML file and drop it on a System Monitor object. In this case, the counters specified in the HTML file will be added to the System Monitor display.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add counters to System Monitor
Create a counter log
Create a trace log
Create an alert

# Create a custom monitoring console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a custom monitoring console**

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove Snap-in**.
3. On the **Standalone** tab, click **Add**.
4. In the snap-in list, click **ActiveX Control**, and then click **Add**.
5. In the **Insert ActiveX Control** wizard, click **Next**.

   There is a short delay before the next dialog box appears; this is normal.
6. In **Control category**, click **All Categories**.
7. In **Control type**, click **System Monitor Control**.
8. Click **Next**.

   The System Monitor control provides the functionality of System Monitor.
9. In **Select a name for the ActiveX control**, type a name for the control, and then click **Finish**.

## Notes

- To open MMC, click **Start**, click **Run**, and then type **mmc**.
- Creating a custom console is useful if you want to have System Monitor hosted in the same console with Event Viewer or another tool.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save a monitoring console
Event Viewer overview
Event Viewer
Microsoft Management Console

# Use a Windows NT Server 4.0 Performance Monitor settings file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use a Windows NT Server 4.0 Performance Monitor settings file**

- At a command prompt, type

  **perfmon.exe***file_name*

  where *file_name* is the name of the settings file you want to use.

## Notes

- This procedure will open a Performance console configured with the settings from the Windows NT Server 4.0 version of Performance Monitor settings file.

- This procedure works for the following types of Windows NT Server 4.0 version of Performance Monitor settings files: chart (.pmc), report (.pmr), alert (.pma), and log (.pml). Workspace (.pmw) files that were created with Windows NT Server 4.0 version of Performance Monitor are not fully compatible with System Monitor. Only one of the views previously saved in the workspace will be available in System Monitor.

- To display the Windows NT Server 4.0 settings file in System Monitor, the system temporarily converts the file for use with System Monitor but discards the converted version after the console starts. If you want to save the settings file for permanent use with System Monitor, type the following command:

  **perfmon.exe /HTMLFILE:***converted_filesettings_file*

  where *converted_file* is the name you are giving to the converted file and *settings_file* is the name of the original Windows NT Server 4.0 settings file.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Collect performance data using WMI
Perfmon

# Create and Configure Counter Logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and configure counter logs**

- Create a counter log
- Add objects to a log
- Add counters to a log
- Get details about objects and counters
- Set parameters for a log

# Create a counter log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a counter log**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Counter Logs**.

   Any existing counter logs will be listed in the details pane. A green icon indicates that a log is running; a red icon indicates that a log has stopped.
3. Right-click a blank area of the details pane, and click **New Log Settings**.
4. In **Name**, type the name of the counter log, and then click **OK**.
5. On the **General** tab, click **Add Objects** and select the performance objects you want to add, or click **Add Counters** to select the individual counters you want to log.

   - You must use the Run as text box and Password button to collect counter data from a remotely managed computer. These features enable you to specify the logon account name for the log. You can access them from the **General** tab.
6. If you want to change the default file and schedule information, make the changes on the **Log Files** tab and the **Schedule** tab.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- To save the settings for a counter log, trace log, or alert, right-click the log or alert in the details pane and click **Save Settings As**. You can then specify an .htm file in which to save the settings. To reuse the saved settings for a new log or alert, right-click the details pane, and then click **New Log Settings From** or **New Alert Settings From**. This is an easy way to generate new settings from a log or alert configuration. You can also open the HTML file in Internet Explorer to display a System Monitor graph.
- Although the extension for a log file is .blg, the file format is changed in the Windows server 2003 family and therefore cannot be read on an earlier version of the operating system. For example, if you copy a binary file from Windows Server 2003, Standard Edition, to Windows 2000, you will be unable to read it on the latter system.

#### Using a command line

1. Open Command Prompt

2. Type:

   **logman create counter***collection_name***-c** {*path* [*path*...]} | {**-cf***filename*}

| Value | Description |
|---|---|
| **create counter** *collection_name* | Creates a counter collection query named *collection_name*. |
| **-c** {*path* [*path*...]} \| **-cf** {*filename*} | Specifies the performance counter path to log (*path*), or specifies the path name of the log file that lists these counters (**-cf***filename*). To list multiple counter paths, separate the command line by a space, or use the **-cf** option to list counter paths in an input file, one per line. Use this format [*computer*]*object*[*instance*]*counter*] for each counter path where:<br><br>*Computer* is the optional computer name or the IP address of a remote computer. If a computer name is not supplied, the local computer will be used by default.<br><br>*Object* is the performance counter object (process or memory).<br><br>*Instance* (if any) is the performance object (process).<br><br>*Counter* is the performance counter for the object (processor time). |

For example, to create a collection query named perf_log to collect Total % Processor Time counter data, type:

**logman create counter perf_log -c "\Processor(_Total)\% Processor Time"**

### Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add counters to a log
Set parameters for a log
Define start or stop parameters for a log or alert
Logman

# Add objects to a log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add objects to a log**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then double-click **Counter Logs** or **Alerts**.
3. In the details pane, double-click the log or alert you want to modify.
4. On the **General** tab, click **Add Objects**. For each object that you want to add to the log, perform the following steps:

   - To log objects from the computer on which the Performance Logs and Alerts service will run, click **Use local computer counter objects**.

     Or, to log objects from a specific computer regardless of where the service is run, click **Select counter objects from computer:** and specify the Universal Naming Convention (UNC) name, such as *MyLogServer*, of the computer you want to monitor.

   - In **Performance object**, click an object to monitor, and then click **Add**. Objects are added to the **Counters** list on the **General** tab.
   - When you are finished adding objects to monitor, click **Close**.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- When creating counter log settings for export, make sure to select **Use local computer counter objects**. Otherwise, counter logs will obtain data from the computer named in the text box.
- For a description of an object, click the object in **Performance counters** and then click **Explain**.
- When you specify the name of a remote computer using **Select counter objects from computer**, only the name of the local computer and those computer names entered during the current Performance session are listed. Computer names entered in previous sessions will not be listed.
- The Run as feature enables you to run log collections under different accounts. For example, if you need to log data from a remote computer that requires administrative credentials, you can use the feature to specify an account with the

necessary permissions.

#### Using a command line

1. Open Command Prompt

2. Type:

   **logman update** *collection_name* **-c** {*path* [*path*...]} | {**-cf***filename*}

| Value | Description |
|---|---|
| **update** *collection_name* | Updates the collection query named *collection_name*. For counter collections, modifications to the query will stop, and then restart the collections. |
| **-c** {*path* [*path*...]} \| **-cf** {*filename*} | Specifies the performance object path to log (*path*), or specifies the path name of the log file that lists these objects (**-cf***filename*). To list multiple object paths, separate the command line by a space, or use the **-cf** option to list object paths in an input file, one per line. Use this format [*computer**object*] for each object path where: *Computer* is the optional computer name or the IP address of a remote computer. If a computer name is not supplied, the local computer will be used by default.<br><br>*Object* is the performance counter object (process or memory). |

For example, to add the PhysicalDisk object to the collection query named perf_log on computer \\COMP1, type:

**logman update perf_log -c "\\COMP1\PhysicalDisk(*)*"**

### Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Define start or stop parameters for a log or alert
Set parameters for a log
Logman

# Add counters to a log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add counters to a log**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Counter Logs**.
3. In the details pane, double-click the log you want to modify.
4. On the **General** tab, click **Add Counters**. For each counter or group of counters that you want to add to the log, perform the following steps:

   - To log counters from the computer on which the Performance Logs and Alerts service will run, click **Use local computer counters**.

     Or, to log counters from a specific computer regardless of where the service is run, click **Select counters from computer** and specify the Universal Naming Convention (UNC) name, such as *MyLogServer*, of the computer you want to monitor.

   - In **Performance object**, select an object to monitor.
   - In **Select counters from list**, click one or more counters to monitor.
   - To monitor all instances of the selected counters, click **All Instances**. (Binary logs can include instances that are not available at log startup but subsequently become available.)

     Or, to monitor particular instances of the selected counters, click **Select Instances From List**, and then click an instance or instances to monitor.

   - Click **Add**.
   - When you are finished adding counters, click **Close**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- When creating a monitoring console for export, make sure to select **Use local computer counters**. Otherwise, counter

logs will obtain data from the computer named in the text box, regardless of where the console file is installed.

- For a description of a counter, click the counter in **Performance counters** and then click **Explain**.

- Some object types have several instances. For example, if a system has multiple processors, the Processor object type will have multiple instances. The PhysicalDisk object type has two instances if a system has two disks. Some object types, such as Memory and Server, have only a single instance. If an object type has multiple instances, you can add counters to track statistics for each instance, or in many cases, for all instances at once. By default, counters are shown with both the instance name and an instance index. To turn off this feature, right-click the System Monitor graph, click **Properties**, and click the **General** tab.

- When you specify the name of a remote computer using **Select counters from computer**, only the name of the local computer and those computer names entered during the current Performance session are listed. Computer names entered in previous sessions will not be listed.

**Using a command line**

1. Open Command Prompt

2. Type:

   **logman update** *collection_name* **-c** {*path* [*path*...]} | {**-cf***filename*}

| Value | Description |
|---|---|
| **update** *collection_name* | Updates the collection query named *collection_name*. For counter collections, modifications to the query will stop, and then restart the collections. |
| **-c** {*path* [*path*...]} \| **-cf** {*filename*} | Specifies the performance counter path to log (*path*), or specifies the path name of the log file that lists these counters (**-cf***filename*). To list multiple counter paths, separate the command line by a space, or use the **-cf** option to list counter paths in an input file, one per line. Use this format [*computer*]*object*[*instance*]*counter*] for each counter path where:<br><br>*Computer* is the optional computer name or the IP address of a remote computer. If a computer name is not supplied, the local computer will be used by default.<br><br>*Object* is the performance counter object (process or memory).<br><br>*Instance* (if any) is the performance object (process).<br><br>*Counter* is the performance counter for the object (processor time).<br><br>**Note**<br><br>• An asterisk (*) can be used as a wild card when specifying the counter path. For example, to specify all instances of the % Processor Time, use "\Processor(*)\% Processor Time". |

For example, to add all instances of the % Interrupt Time to the collection query named perf_log on computer \\COMP1, type:

**logman update perf_log -c "\\COMP1\Processor(*)\% Interrupt Time"**

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command**

**prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Define start or stop parameters for a log or alert
Set parameters for a log
Remove counters from a log or alert
Logman

# Get details about objects and counters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To get details about objects and counters**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Counter Logs**.
3. In the details pane, double-click the counter log.
4. Click the **Add Objects** button to open the **Add Objects** dialog box or click the **Add Counters** button to open the **Add Counters** dialog box.
5. In the **Add Objects** dialog box, select **Performance objects**, click an object, then click **Explain**.
6. In the **Add Counters** dialog box, select **Performance object**, click an object, and in **Performance counters**, click a counter and then click **Explain**.

## Notes

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- Clicking **Explain** opens a pop-up window that contains information about the selected object or counter. For more information on counters, including detailed descriptions of how each counter is calculated and guidelines on using counters to diagnose performance problems, see the Performance Counter Help Reference at the Microsoft Windows Resource Kits Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Performance Logs and Alerts interface
Add counters to a log
Add objects to a log

# Set parameters for a log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set parameters for a log**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**.
3. To set parameters for a counter log, double-click **Counter Logs**.

   Or, to set parameters for a trace log, double-click **Trace Logs**.
4. In the details pane, double-click the log.
5. On the **Log Files** tab, complete the following options:

   - **Log file type.** In the list, select the format you want for this log file, complete the options, and then click the **Configure** button. See "Quick reference to log file formats" under Related Topics for more information.
   - **End file names with.** Select this and choose the suffix style you want from the list. Use **End file names with** to distinguish between individual log files with the same log file name that are in a group of logs that have been automatically generated.
   - **Start numbering at.** Set this to the start number for automatic file numbering, when you select **nnnnnn** as the **End file names with**.
   - **Comment.** If appropriate, type a comment or description for the log file.
   - **Overwrite existing log file.** Select this check box if the log file should overwrite and replace the existing log file.
   - **Configure.** Select the configuration parameters using the following options for either **Configure Log Files** or **Configure SQL Logs, based on the log file type selected from the drop-down menu.**

| Option | Description | Applies to |
|---|---|---|
| **Location** | Type the name of the folder where you want the log file created, or click **Browse** to search for the folder. | Configure Log Files |
| **File name** | Type a partial or base name for the log file. You can use **File name** in conjunction with **End file names with** if appropriate. | Configure Log Files |
| **Repository name** | Select the **System DSN** (Data Source Name) from the drop-down list, and then type the **Log set name.** The **Log set name** will be stored in the database within the **System DSN.** | Configure SQL Logs |

6. **Log file size.** Decide, using the following options, whether or not you want to limit the size of the log file.

7. Or, under **Log set size.** Decide, using the following options, whether or not you want to limit the size of the log set.

    - **Maximum limit.** When you select this option, data is continuously collected in a log file until it reaches limits set by disk quotas or the operating system. (For SQL logs, data is collected in a database until it reaches its limits set by the number of records written.)

    - **Limit of.** To define a size limit for the log file, specify the size. For counter and trace logs, specify the maximum size in megabytes. For SQL logs, specify the maximum size in records.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

- If you choose to allow the log file to grow to its maximum size, consider the available space on your disk and any quotas that may be in effect. An error might occur if you run out of disk space due to logging. The **Log file size** option works in conjunction with the **When the log file is full** option on the **Schedule** tab.

- When you select a log file limit or log set size limit, the stop option **When the log file is full** reflects the size limit typed here.

- The **File name** can identify a group of log files with similar parameters.

- The **End file names with** suffix can be a numeric sequence or a character string indicating the date and time the log was started. For example, you might define the **File name** "DailyLog" and the Performance Logs and Alerts service would append the date as a suffix for each file so that you might have a series of files named as follows: DailyLog_19990715, DailyLog_19990716, for each day the log is run.

- The comment text and the file name appear in the details pane in the Performance Logs and Alerts console window.

- Sequential counter or trace log files can grow to a size larger than the maximum file size specified by the user. This can occur for different reasons, depending on the type of log configured. For counter logs, the log service postpones checking the log's file size until after it has collected the last data sample, at which point the file size might already have exceeded the defined limit. For trace logs, the log service writes additional buffers to the log when the log service stops, which can cause the final log file size to be slightly larger than the defined limit.

**Using a command line**

1. Open Command Prompt

2. Type the appropriate command below:

| Command | Action |
|---|---|
| **logman update** *collection_name* **-f** {*bin* \| *bincirc* \| *csv* \| *tsv* \| *SQL*} | Defines the log file format. |

| | |
|---|---|
| **logman update** *collection_name* **-v** {*nnnnn* \| *mmddhhmm*} | Specify the file suffix. |
| **logman update** *collection_name* **-v** *nnnnn* | Set the automatic file numbering start number. |
| **logman update** *collection_name* **-o** *path* | Specify the location of the log file. |
| **logman update** *collection_name* **-o** *DSN!counter_log* | Specify the location of the SQL database and log set where output data will be written. |

| Value | Description |
|---|---|
| **update** *collection_name* | Updates the collection query named *collection_name*. For counter collections, modifications to the query will stop, and then restart the collections. |
| **-f** *bin* \| *bincirc* \| *csv* \| *tsv* \| *SQL* | Specifies the file format used for collecting performance counter and trace data. You can use binary (bin), circular binary (circbin), comma (csv) and tab (tsv) separated, or SQL database (SQL) formats when collecting performance counters.<br><br>**Note**<br><br>• For SQL database formats, you must use the **-o** option in the command line with the DNS!counter_log option, the System Database System Name (System DSN) must be predefined, and you must have Write permissions to the database. The log set *counter_log* is created in the database, and is specified by the DSN.<br><br>If a file format is not specified, it defaults to binary. |
| **-v** *nnnnn* \| *mmddhhmm* | Attaches the version control information to the end of the output file and path name. Use numeric *nnnnn* format, or date format *mmddhhmm* (month, day, 24-hour, minute) for version control. |
| **-o** *path* | Specifies the path name of the output file that collects performance counter or trace data. By default, the collection log file name is the collection query name suffixed by either .blg for performance counters, or .etl for trace data. |
| **-o** *DSN!counter_log* | Specifies the location of the SQL database and log set. |

For example, to specify that the output for counter collection perf_log be written to a SQL log set named daily_log in the perfdb database, type:

**Logman update perf_log -f sql -o perfdb!daily_log**

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Add counters to a log
Select trace log providers and events
Define start or stop parameters for a log or alert
Define trace log buffers
Quick reference to log file formats
Logman

# Create and Configure Trace Logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and configure trace logs**

- Create a trace log
- Select trace log providers and events
- Define trace log buffers

# Create a trace log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a trace log**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Trace Logs**.

   Any existing trace logs will be listed in the details pane. A green icon indicates that a log is running; a red icon indicates that a log has stopped.
3. Right-click a blank area of the details pane, and click **New Log Settings**.
4. In **Name**, type the name of the trace log, and then click **OK**.

   By default, the log file is created in the PerfLogs folder in the root directory on your system drive, a sequence number is appended to the file name you entered, and the sequential trace file type with the .etl extension. Use the **Log Files** and **Advanced** tabs to modify these parameters or define other parameters for your log. To define providers and events to log, use the **General** tab. To specify when you want logging to occur, use the **Schedule** tab.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- To save the settings for a counter log, trace log, or alert, right-click the log or alert in the details pane and click **Save Settings As**. You can then specify an .htm file in which to save the settings. To reuse the saved settings for a new log or alert, right-click the details pane, and then click **New Log Settings From** or **New Alert Settings From**. This is an easy way to generate new settings from a log or alert configuration. You can also open the HTML file in Internet Explorer to display a System Monitor graph.

**Using a command line**

1. Open Command Prompt.
2. To create a trace log type:

   **Logmancreate***tracecollection_name*

| Value | Description |
| --- | --- |
| **create** *trace* | Creates a trace log collection query. |
| *collection_name* | Specifies the name of the collection query. |

For example, to create a trace log collection query named perf_log, type:

**logman create trace perf_log**

Once the collection query is created, providers then need to be enabled.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **Logman /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set parameters for a log
Select trace log providers and events
View or change properties of a log or alert
Logman

# Select trace log providers and events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select trace log providers and events**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Trace Logs**.
3. In the details pane, double-click the log.
4. For a list of the installed providers and their status (enabled or not), click **Provider Status**.

   Only one instance of each trace provider can be enabled at any given time.

   By default, the **Nonsystem providers** option is selected to keep trade logging overhead to a minimum.
5. If you click **Events logged by system provider**, a default provider (the Windows kernel trace provider) is used to monitor processes, threads, and other activity. To define events for logging, click the check boxes as appropriate.
6. If you click **Nonsystem providers**, you can select the data providers you want--for example, if you have written your own providers. Use the **Add** or **Remove** buttons as needed.

## Notes

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- Trace logging of file details and page faults can generate an extremely large amount of data. It is recommended that you limit trace logging using the file details and page fault options to a maximum of two hours.
- The Windows kernel trace provider and some other providers require you to run the collection under the authority of an account with administrative credentials. By default, the service runs under the NT Authority\NetworkService account. Use the Run as text box to specify an administrator account to do this.

**Using a command line**

1. Open Command Prompt.
2. Type the appropriate command below:

| Action | Command |
|---|---|
| To view each installed provider and its status. | **logman query providers** |

| | |
|---|---|
| To specify that data collected by the system trace provider is logged. | **logman update** *collection_name* **-P** "Windows kernel trace" (*process*, *thread*, *disk*, *net*, *page*, *file*) |
| To specify that data collected by a provider other than the Windows kernel trace provider is logged. | **logman update** *collection_name* **-P** *provider* |

| Value | Description |
|---|---|
| **query providers** | Queries the providers installed on the local system. |
| **update** *collection_name* | Updates the collection query named *collection_name*. |
| **-P** "Windows kernel trace" (*process*, *thread*, *disk*, *net*, *page*, *file*) | Specifies "Windows kernEl trace" as the provider that collects data for the trace log. *Process*, *thread*, *disk*, *net*, *pf*, *hf*, *registry*, *image*, and *file* are optional events to include in the trace log.<br><br>Use *process* to include data on the creating and ending processes.<br><br>Use *thread* to include data on the creating and ending threads.<br><br>Use *disk* to include data on disk input/output operations.<br><br>Use *net* to include data on TCP/IP send or receive requests.<br><br>Use *pf* to enable soft page fault tracing.<br><br>Use *hf* to enable hard page fault tracing.<br><br>Use *registry* to include data on registry operations.<br><br>Use *image* to include data on the program name for the loaded process.<br><br>Use *file* to include file I/O data. |
| **-P** *provider* | Specifies the non-system provider that collects data for the trace log. |

For example, to specify that a provider named Nonsystem01 is used to collect data for a trace log named perf_log, type:

**logman update perf_log -P Nonsystem01**

**Notes**

- The Windows kernel trace provider and some other providers require you to run the collection under the authority of an account with administrative credentials. By default, the service runs under the NT Authority\NetworkService account. Use the **-u** option to specify an account with administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set parameters for a log
Create a trace log
View or change properties of a log or alert
Define trace log buffers
Logman

# Define trace log buffers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define trace log buffers**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and click **Trace Logs**.
3. In the details pane, double-click the log, and click the **Advanced** tab on the dialog box that appears.
4. In **Buffer size**, specify the size of the buffer you want to use in kilobytes for trace data collection.
5. In **Minimum**, specify the smallest number of buffers you want to use for trace data.
6. In **Maximum**, specify the largest number of buffers you want to use for trace data.
7. To have the trace provider periodically flush the buffers, select the **Transfer data from buffers to log file at least every** check box and specify the transfer interval in seconds.

**Note**

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **logmanupdate***collection_name***-bs** {*value*} **-nb** {*minmax*} **-ft***[[hh:]mm:]ss*

| Value | Description |
|---|---|
| **update** *collection_name* | Updates the collection query named *collection_name*. |
| **-bs** {*value*} | Specifies the buffer size in *n* kilobytes for the trace log. |
| **-nb** {*minmax*} | Specifies the minimum and maximum number of buffers for the trace log. |
| **-ft** *[[hh:]mm:]ss* | Specifies whether the buffer will be flushed and the interval at which the buffer will be flushed. |

For example, to specify a buffer size of 128 kilobytes, a minimum of 10 buffers, and a maximum of 25 buffers for a trace log named perf_log, type:

**logman update perf_log -bs 128 -nb 10 25**

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a trace log
Select trace log providers and events
Set parameters for a log
Logman

# Create and Configure Alerts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and configure alerts**

- Create an alert
- Define counters and thresholds for an alert
- Define actions for an alert

# Create an alert

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an alert**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Alerts**.

   Any existing alerts will be listed in the details pane. A green icon indicates that an alert is running; a red icon indicates that an alert has stopped.
3. Right-click a blank area of the details pane and click **New Alert Settings**.
4. In **Name**, type the name of the alert, and then click **OK**.
5. To define a comment for your alert, along with counters, alert thresholds, and a sample interval, use the **General** tab.

   To define actions that should occur when counter data triggers an alert, use the **Action** tab. To define when the service should begin scanning for alerts, use the **Schedule** tab.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- To save the settings for a counter log, trace log, or alert, right-click the log or alert in the details pane and click **Save Settings As**. You can then specify an .htm file in which to save the settings. To reuse the saved settings for a new log or alert, right-click the details pane, and then click **New Log Settings From** or **New Alert Settings From**. This is an easy way to generate new settings from a log or alert configuration. You can also open the HTML file in Internet Explorer to display a System Monitor graph.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save graph settings as an HTML page
Define actions for an alert
Define counters and thresholds for an alert

# Define counters and thresholds for an alert

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define counters and thresholds for an alert**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Alerts**.
3. In the details pane, double-click the alert.
4. In **Comment**, type a comment to describe the alert as needed, and click **Add**.
5. For each counter or group of counters that you want to add to the log, perform the following steps:
   a. To monitor counters from the computer on which the Performance Logs and Alerts service will run, click **Use local computer counters**.

      Or, to monitor counters from a specific computer regardless of where the service is run, click **Select counters from computer** and type the Universal Naming Convention (UNC) name, such as *MyLogServer*, of the computer you want to monitor in the text box.

   b. In **Performance object**, click an object to monitor.
   c. In **Select counters from list**, click one or more counters to monitor.
   d. To monitor all instances of the selected counters, click **All Instances**. (Binary logs can include instances that are not available at log startup but subsequently become available.)

      Or, to monitor particular instances of the selected counters, click **Select Instances From List**, and then click an instance or instances to monitor.

   e. Click **Add**. When you are through adding counters, click **Close**.
6. In **Alert when the value is**, specify **Under** or **Over**, and in **Limit**, specify the value that triggers the alert. Perform this step for each counter or group of counters that you added to the log.
7. In **Sample data every**, specify the amount and the unit of measure for the update interval.
8. Complete the alert configuration using the **Action** and **Schedule** tabs.
9. Click **Apply**, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

- You can add a counter to an alert multiple times, and specify different thresholds for each instance.

- To save the settings for a counter log, trace log, or alert, right-click the log or alert in the details pane and click **Save Settings As**. You can then specify an .htm file in which to save the settings. To reuse the saved settings for a new log or alert, right-click the details pane, and then click **New Log Settings From** or **New Alert Settings From**. This is an easy way to generate new settings from a log or alert configuration. You can also open the HTML file in Internet Explorer to display a System Monitor graph.

- Some object types have several instances. For example, if a system has multiple processors, the Processor object type will have multiple instances. The PhysicalDisk object type has two instances if a system has two disks. Some object types, such as Memory and Server, have only a single instance. If an object type has multiple instances, you can add counters to track statistics for each instance, or in many cases, for all instances at once.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Setting up a monitoring configuration
Analyzing performance data
Performance Logs and Alerts interface

# Define actions for an alert

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define actions for an alert**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Alerts**.
3. In the details pane, double-click the alert.
4. Click the **Action** tab.
5. To have Performance Logs and Alerts create an entry visible in Event Viewer, select **Log an entry in the application event log**. This option is selected by default.
6. To have Performance Logs and Alerts trigger the Messenger service to send a message, select **Send a network message to** and type the IP address or name of the computer on which the alert message should be displayed.
7. To run a counter log when an alert occurs, select **Start performance data log** and specify the counter log you want to run.
8. To have a program run when an alert occurs, select **Run this program** and type the file path and name, or click **Browse** to locate the file. Click **Command Line Arguments** and select the appropriate check boxes for arguments to include when the program is run.

## Notes

- To perform this procedure, you must be a member of the Administrators group or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

  **Important**

  - Note the security implications running an interactive program: When an alert stops, it will run the program on the desktop. If you schedule a collection to run on a system that other users have access to under an account with administrative credentials, then the program will appear on the desktop under the authentication of the administrator account, possibly when other users are logged onto the system.

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- When you select **Run this program**, by default, the program runs in the background.
- When an alert occurs, the service creates a process and runs the specified executable or command file. The service also copies any command line arguments that you define to run the file to the command line.
- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create an alert
Start or stop a counter log, trace log, or alert manually

# Work with Logs and Alerts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with logs and alerts**

- Start or stop a counter log, trace log, or alert manually
- Remove counters from a log or alert
- View or change properties of a log or alert
- Define start or stop parameters for a log or alert
- Delete a log or alert

# Start or stop a counter log, trace log, or alert manually

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start or stop a counter log, trace log, or alert manually**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and click **Counter Logs**, **Trace Logs**, or **Alerts**.
3. In the details pane, right-click the name of the log or alert you want to start or stop, and click **Start** to begin the logging or alert activity you defined, or click **Stop** to terminate the activity.

   Or, click the **Start** or **Stop** buttons in the toolbar.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- There might be a slight delay before the log or alert starts or stops, indicated when the icon changes color (from green for started to red for stopped, and vice versa). This delay might reflect the startup time for a service or the fact that you have added many counters. If a log or alert does not start, you can view the Application Log for errors.

**Using a command line**

1. Open Command Prompt
2. Type:

   **logman start***collection_name*

   or

   **logman stop***collection_name*

| Value | Description |
|---|---|
| **start** *collection_name* | Starts the data collection query *collection_name*. Use this option to change from scheduled collections to manual ones. Use the **update** parameter in the command line with begin-time (**-b**), end-time (**-e**), or repeat-time (**-rt**) to reschedule collections. |

| | |
|---|---|
| **stop** *collection_name* | Stops the data collection query collection_name. Use this option to change from scheduled collections to manual ones. Use the **update** parameter in the command line with begin-time (**-b**), end-time (**-e**), or repeat-time (**-rt**) to reschedule collections. |

For example, to stop the collection query named perf_log, type:

**logman stop perf_log**

## Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Performance Logs and Alerts interface
Logman
Event Viewer overview

# Remove counters from a log or alert

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove counters from a log or alert**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Counter Logs** or **Alerts**.
3. In the details pane, double-click the name of the log or alert.
4. Under **Counters**, click the counter you want to remove, and then click **Remove**.

## Notes

- To perform this procedure, you must be a member of the Administrators group or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set parameters for a log
Define start or stop parameters for a log or alert
Add counters to a log

# View or change properties of a log or alert

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view or change properties of a log or alert**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**.
3. Click **Counter Logs**, **Trace Logs**, or **Alerts**.
4. In the details pane, double-click the name of the log or alert.
5. View or change the log properties as needed.

### Notes

To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- To start or stop a log or alert, right-click its name in the details pane, and click **Start** or **Stop** as appropriate. This overrides any scheduled start and stop settings, resetting the start and stop mode to manual.

**Using a command line**

1. Open Command Prompt.
2. Type the appropriate command below:

| Action | Command |
|---|---|
| To update counter log counters and the update interval. | **logman update** *collection_name* **-c** {*path* [*path*...]} \| **-cf** {*filename*} **-si***[[hh:]mm:]ss* |
| To update a counter log collection, file type, file size limits, path and name, and automatic naming parameters. | **logman update** *collection_name***-f** {*bin* \| *bincirc* \| *csv* \| *tsv* \| *SQL*} **-max***value***-o***path***-v** {*nnnnn* \| *mmddhhmm*} |
| To update manual to automated start and stop methods and schedule. | **logman update** *collection_name***-b** *M/d/I h:mm:ss* [*AM* \| *PM*] **-e** *M/d/yyyy h:mm:ss* [*AM* \| *PM*] **-r** |
| To update trace log providers and events to log. | **logman update** *collection_name***-p** *provider* |

| | |
|---|---|
| To update a trace log collection, file type, file size limits, path name, and automatic naming parameters. | **logman update** *collection_name* **-f** {*bin* \| *bincirc*} **-max***value* **-o***path* **-v** {*nnnnn* \| *mmddhhmm*} |
| To update a trace log's buffer size, limits, and transfer interval. | **logman update** *collection_name* **-bs** {*value*} **-nb** {*min max*} **-ft** *[[hh:]mm:]ss* |

| Value | Description |
|---|---|
| **update** *collection_name* | Updates the collection query named *collection_name*. For counter collections, modifications to the query will stop, and then restart the collections. |
| **-c** {*path* [*path*...]} \| **-cf** {*filename*} | Specifies the performance counter path to log (*path*), or specifies the path name of the log file that lists these counters (**-cf***filename*). To list multiple counter paths, separate the command line by a space, or use the **-cf** option to list counter paths in an input file, one per line. Use this format [*computer*]*object*[*instance*]*counter*] for each counter path where:<br><br>*Computer* is the optional computer name or the IP address of a remote computer. If a computer name is not supplied, the local computer will be used by default.<br><br>*Object* is the performance counter object (process or memory).<br><br>*Instance* (if any) is the performance object (process).<br><br>*Counter* is the performance counter for the object (processor time). |
| **-si** *[[hh:]mm:]ss* | Specifies the update interval for counter collection in hours, minutes, and seconds. |
| **-f** *bin* \| *bincirc* \| *csv* \| *tsv* \| *SQL* | Specifies the file format used for collecting performance counter and trace data. You can use binary (bin), circular binary (circbin), comma (csv) and tab (tsv) separated, or SQL database (SQL) formats when collecting performance counters. If a file format is not specified, it defaults to binary. Binary (bin) and circular binary (circbin) are the only types supported for trace collections. |
| **-max** {*value*} | Specifies the maximum size of the collected log file in megabytes. If the log file exceeds the maximum size, the collection will stop. For a SQL database, the maximum size is the number of records to be written. |
| **-o** *path* | Specifies the path name of the output file that collects performance counter or trace data. By default, the collection log file name is the collection query name suffixed by either .blg for performance counters, or .etl for trace data. |
| **-v** {*nnnnn* \| *mmddhhmm*} | Attaches the version control information to the end of the output file and path name. Use numeric *nnnnn* format, or date format *mmddhhmm* (month, day, 24-hour, minute) for version control. |
| **-b** *M/d/yyyy h:mm:ss*[*AM* \| *PM*] | Specifies the begin-time for collections in a 24-hour format. You can also specify begin-time for collections in a 12-hour format by adding AM or PM in the command line. By default, the current day and time is used unless otherwise specified. |
| **-e** *M/d/yyyy h:mm:ss*[*AM* \| *PM*] | Specifies the end-time for collections in a 24-hour format. You can also specify end-time for collections in a 12-hour format by adding AM or PM in the command line. By default, the current day and time is used unless otherwise specified. |
| **-r -[-]r** | Repeats the collection every day at the time periods specified by the **-b** and **-e** options. This is valid only for begin- and end-times specified on the same day, month, and year.<br><br>**-[-]r** turns off the repeat option. |
| **-p** {*provider* [flags [level]]} | Specifies the name of the provider. |
| **-bs** {*value*} | Specifies the buffer size in *n* kilobytes for the trace log. |
| **-nb** {*minmax*} | Specifies the minimum and maximum number of buffers for the trace log. |

| | |
|---|---|
| **-ft** *[[hh:]mm:]ss* | Specifies whether the buffer will be flushed, and the interval at which the buffer will be flushed. |

**Notes**

To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set parameters for a log
Define start or stop parameters for a log or alert
Define actions for an alert
Define counters and thresholds for an alert
Logman

# Define start or stop parameters for a log or alert

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define start or stop parameters for a log or alert**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Performance.
2. Double-click **Performance Logs and Alerts**, and then click **Counter Logs**, **Trace Logs**, or **Alerts**.
3. In the details pane, double-click the name of the log or alert.
4. Click the **Schedule** tab.
5. Under **Start log**, click one of the following options:
   - To start the log or alert manually, click **Manually**. When this option is selected, to start the log or alert, right-click the log name in the details pane, and click **Start**.
   - To start the log or alert at a specific time and date, click **At**, and then specify the time and date.
6. Under **Stop log**, select one of the following options:
   - To stop the log or alert manually, click **Manually**. When this option is selected, to stop the log or alert, right-click the log or alert name in the details pane, and click **Stop**.
   - To stop the log or alert after a specified duration, click **After**, and then specify the number of intervals and the type of interval (days, hours, and so on).
   - To stop the log or alert at a specific time and date, click **At**, and then specify the time and date. (The year box accepts four characters; the others accept two characters.)
   - To stop a log when the log file becomes full, select options as follows: For counter logs and trace logs, click **When the log file is full**. The file will continue to accumulate data according to the file-size limit you set on the **Log Files** tab. When setting this option, take into consideration your available disk space and any disk quotas that are in place. An error might occur if your disk runs out of disk space due to logging.
7. Complete the properties as appropriate for logs or alerts:
8. For logs, under **When a log file closes**, select the appropriate option:
   - If you want to configure a continuous (automated) counter or trace logging, select **Start a new log file**.
   - If you want to run a program after the log file stops (for example, a copy command for transferring completed logs to an archive site), select **Run this command**. Also, type the path and file name of the command to run, or click **Browse** to locate the command.

- For alerts, under **When an alert scan finishes**, select **Start a new scan** if you want to configure continuous alert scanning.

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the group on the local computer, or you must have been delegated the appropriate Performance Log Users authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

- When you schedule a log to close at a specific time and date or close the log manually, the **Start a new log file** option is unavailable.

- If the log is set to start at a scheduled time using **Start At**, the log will restart if the system reboots during the scheduled time. The scheduled time is based on the **Stop** settings.

- When you select **Run this command**, by default, the command runs in the background.

- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.

**Important**

- Note the security implications running an interactive program: When log stops, it will run the program on the desktop. If you schedule a collection to run on a system that other users have access to under an account with administrative credentials, then the program will appear on the desktop under the authentication of the administrator account, possibly when other users are logged onto the system.

- To run an interactive program, open **Services** and double-click **Performance Logs and Alerts**. Click the **Log On** tab, and fill the Local System account and **Allow service to interact with desktop** check boxes.

- When creating a monitoring console for export, make sure to select **Use local computer counters**. Otherwise, System Monitor will obtain data from the computer named in the text box, regardless of where the console file is installed.

- For a description of a particular counter, click the name of the counter in **Add counters**, and then click **Explain**.

- To monitor the sum of the values all of a particular counter's instances, select the _Total instance.

- Unless you are configuring data from a log, you can select only active instances in the list box.

- By default, counters are shown with both the instance name and an instance index. To turn off this feature, right-click the details pane, click **Properties**, select the **General** tab, and clear the **Allow duplicate counter instances** check box.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **logmanupdate***collection_name***-mstartstop**

| Value | Description |
| --- | --- |
| **update** *collection_name* | Updates the collection query named *collection_name*. |
| **-m start stop** | Specifies that the logging will start and stop manually. **Start** specifies the time you want logging to start. **Stop** specifies the time you want logging to stop. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set parameters for a log
Logman

# Delete a log or alert

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To delete a log or alert

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Performance.
2. Double-click **Performance Logs and Alerts**.
3. Click **Counter Logs**, **Trace Logs**, or **Alerts**.
4. In the details pane, right-click the name of the log or alert, and click **Delete**.

### Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **logmandelete***collection_name*

| Value | Description |
|---|---|
| **delete** *collection_name* | Deletes the collection query named *collection_name*. If *collection_name* does not exist, you will receive an error. |

For example, to delete a log named perf_log, type:

**logman delete perf_log**

### Notes

- To perform this procedure, you must be a member of the Administrators group, or the Performance Log Users group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **logman /?**

- This command is used only for deleting logs. You cannot use this command to delete an alert.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Performance Logs and Alerts interface
Logman

# Performance monitoring Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information on performance-monitoring tools in the Windows Server 2003 family:

- Performance console overview
- Understanding the Performance console
- Using the Performance console
- Performance monitoring resources

# Performance console overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Performance console overview**
The Windows Server 2003 family provides the following tools as part of the Performance console:

- System Monitor
- Performance Logs and Alerts
- Task Manager

Monitoring system performance is an important part of maintaining and administering your operating system. You use performance data to do the following:

- Understand your workload and the corresponding effect on your system's resources.
- Observe changes and trends in workloads and resource usage so you can plan for future upgrades.
- Test configuration changes or other tuning efforts .
- Diagnose problems and target components or processes for optimization.

System Monitor and Performance Logs and Alerts provide detailed data about the resources that are used by specific components of the operating system and by programs that have been designed to collect performance data. Graphs provide a display for performance-monitoring data. Logs provide recording capabilities for the data. Alerts send notification to users when a counter value reaches, rises above, or falls below a defined threshold.

Task Manager is another tool that provides performance information about systems running the Windows Server 2003 family. Task Manager provides information about programs and processes that are running on your computer, plus a summary of its processor and memory usage. For information about Task Manager, see Task Manager overview.

Microsoft technical support often uses the results of performance monitoring in problem diagnosis. Therefore, Microsoft recommends that you monitor system performance as part of your administrative routine.

For information on server reliability, availability, and performance and tuning features in the Windows Server 2003 family, see Features for Management, Security, and Best Performance.

For overviews of the performance-monitoring tools, see the following topics:

- System Monitor overview
- Performance Logs and Alerts overview

# System Monitor overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Monitor overview**

With System Monitor, you can measure the performance of your own computer or other computers on a network in the following ways:

- Collect and view the real-time performance data of a local computer or several remote computers: Usability has been modified in the Windows Server 2003 family. For example, you can delete multiple counters at once and display the data properties page for a counter directly from the list window. You can save selected data from a performance log file or an SQL database to a new file for analysis at a later time.

- Also new in the Windows Server 2003 family are two new security groups that help you to ensure that only trusted users can access and manipulate sensitive performance data. These are the Performance Log Users group and the Performance Monitor Users group.

- View data collected either currently or previously in a counter log. With the Windows Server 2003 family, you can now simultaneously view data from multiple log files.

- Present data in a printable graph, histogram, or report view.

- Incorporate System Monitor functionality into applications that support ActiveX controls, for example, Web pages, and Microsoft Word or other applications in Microsoft Office

- Create HTML pages from performance views. Views stored in HTML format can be displayed by a browser.

- Create reusable monitoring configurations that can be installed on other computers using Microsoft Management Console (MMC).

With System Monitor, you can collect and view extensive data about the usage of hardware resources and the activity of system services on computers you administer. You can define the data you want System Monitor to collect in the following ways:

- **Type of data.** To select the data to be collected, you specify performance objects, performance counters, and performance object instances.

  Some objects provide data on system resources (such as memory); others provide data on the operation of applications (for example, system services).

- **Source of data.** System Monitor can collect data from your local computer or from other computers on the network for which you have administrative credentials. By default, administrative credentials are required. In addition, you can include real-time data or data collected previously using counter logs. With the Windows Server 2003 family, you can now view performance data that was previously collected and stored in an SQL database by the Performance Logs and Alerts service.

- **Sampling parameters.** System Monitor supports manual, on-demand sampling or automatic sampling based on a time interval you specify; this functionality applies to real-time data only. When viewing logged data, you can also choose starting and stopping times so that you can view data spanning a specific time range.

In addition to options for defining data content, you have considerable flexibility in designing the appearance of your System Monitor views:

- **Type of display.** System Monitor supports graph, histogram, and report views. The graph view is the default view; it offers the widest variety of optional settings.

- **Display characteristics.** For any of the three views, you can define the colors and fonts for the display. In graph and histogram views, you can select from many different options when you view performance data:
    - Provide a title for your graph or histogram and label the vertical axis.
    - Set the range of values depicted in your graph or histogram.
    - Adjust the characteristics of lines or bars plotted to indicate counter values, using color, width, style, and other graphical features.

For more information about the performance-monitoring process and interface, see Performance console overview and System Monitor interface.

# Performance Logs and Alerts overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Performance Logs and Alerts overview**

With Performance Logs and Alerts you can collect performance data automatically from local or remote computers. You can view logged counter data using System Monitor or export the data to spreadsheet programs or databases for analysis and report generation. The following list explains the capabilities of Performance Logs and Alerts:

- New in the Microsoft® Windows Server 2003 family is the ability to run log collections under different accounts. For example, if you need to log data from a remote computer that requires administrative credentials, you can specify an account with the necessary credentials.

- Also new in the Windows Server 2003 family are two new security groups that help you to ensure that only trusted users can access and manipulate sensitive performance data. These are the Performance Log Users group and the Performance Monitor Users group.

- The Windows Server 2003 family supports log files greater than 1 GB in size, and with its new log-file format, you can append performance data to an existing log file.

- Performance Logs and Alerts collects data in a comma-separated or tab-separated format for easy import to spreadsheet programs. A binary log-file format is also provided for circular logging or for logging instances such as threads or processes that may begin after the log starts collecting data. (Circular logging is the process of continuously logging data to a single file, overwriting previous data with new data.)

- You can also collect data in an SQL database format. This option defines the name of an existing SQL database and log set within the database where the performance data will be read or written. This file format is useful when collecting and analyzing performance data at an enterprise level rather than on a per-computer basis. Logging data directly to a SQL database is supported through open database connectivity (ODBC).

- Counter data collected by Performance Logs and Alerts can be viewed during collection as well as after collection has stopped.

- Because logging runs as a service, data collection occurs regardless of whether any user is logged on to the computer being monitored.

- You can define start and stop times, file names, file sizes, and other parameters for automatic log generation.

- You can manage multiple logging sessions from a single console window.

- You can set an alert on a counter, thereby defining that a message be sent, a program be run, an entry made to the application event log, or a log be started when the selected counter's value exceeds or falls below a specified setting.

Similar to System Monitor, Performance Logs and Alerts supports defining performance objects, performance counters, and performance object instances. It also supports setting sampling intervals for monitoring data about hardware resources and system services. Performance Logs and Alerts also offers other options related to recording performance data:

- **Start and stop logging** either manually on demand or automatically based on a user-defined schedule.

- **Configure additional settings** for automatic logging, such as automatic file renaming, and set parameters for stopping and starting a log file based on the elapsed time or the file size.

- **Create trace logs**. Using the default Windows Server 2003 family data provider or another application provider, trace logs record detailed system application events when certain activities, such as a disk I/O operation or a page fault, occurs. When the event occurs, your operating system logs the system data to a file specified by the Performance Logs and

Alerts service. This differs from the operation of counter logs; when counter logs are in use, the service obtains data from the system when the update interval has elapsed, rather than waiting for a specific event. A parsing tool is required to interpret the trace log output. Developers can create such a tool using application programming interfaces (APIs) provided in the MSDN Library on the Microsoft Web site.

- You can also produce trace analysis reports from trace log output files using the Tracerpt tool. Use this tool to process kernel, Active Directory, and other transactional based trace event logs, and to generate trace analysis reports and .csv files from binary logs.

- Define a program that runs when a log is stopped.

- If you want to export log data to Microsoft Excel, the Performance Logs and Alerts service must be stopped because Microsoft Excel requires exclusive access to the log file. Most other programs are not known to require this exclusive access; therefore, in general you can work with data from a log file while the service is collecting data to that file.

For more information about the Performance Logs and Alerts user interface, see Performance Logs and Alerts interface.

# Understanding the Performance console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding the Performance console**

This section describes the design and operation of System Monitor and Performance Logs and Alerts:

- Performance objects and counters
- Performance tool architecture

# Performance objects and counters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Performance objects and counters**

The Windows Server 2003 family of operating systems obtains performance data from components in your computer as those components are utilized. That data is described as a performance object and is typically named for the component generating the data. For example, the Processor object is a collection of performance data about processors on your system.

Performance objects are built into the operating system, typically corresponding to the major hardware components such as memory, processors, and so on. Other programs might install their own performance objects. For example, services such as Windows Internet Name Service (WINS) or server programs such as Microsoft Exchange provide performance objects, and performance graphs and logs can monitor these objects.

Each performance object provides performance counters that represent data on specific aspects of a system or service. For example, the Pages/sec counter provided by the Memory object tracks the rate of memory paging.

Although your system might typically make available many more objects, following are the most frequently used objects:

- Cache
- Memory
- Objects
- Paging File
- PhysicalDisk
- Process
- Processor
- Server
- System
- Thread

The following table describes the services or features of your operating system that you might be using in your configuration and their corresponding performance objects:

| Feature or service to monitor | Performance object available |
|---|---|
| TCP/IP | ICMP, IP, NBT, TCP, and UDP objects |
| Browser, Workstation, and Server services | Browser, Redirector, and Server objects |
| Windows Internet Name Service (WINS) | WINS object |
| Connection Point Services | PBServer Monitor object |
| Indexing Service | Indexing Service, Indexing Service Filter, and HTTP Indexing Service objects |
| Directory service | NTDS object |
| Print server activity | Print Queue object |

For a description of the data provided by a particular counter that is associated with a performance object, click **Explain** in the **Add Counters** dialog box. For information about adding counters, see Add counters to System Monitor.

Although some objects (such as Memory and Server) have only a single performance object instance, some performance objects can have multiple instances. If an object has multiple instances, you can add counters to track statistics for each instance or for all instances at once.

Depending on how the counter was defined, its value may be:

- The most recent measurement of an aspect of resource utilization. These are called instantaneous counters.

  An example is Process\Thread Count, which shows the number of threads for a particular process as of the last time this was measured.

- The average of the last two measurements over the period between samples.

  An example is Memory\Pages/sec, a rate per second based on the average number of memory pages during the last two samples.

Other types of counters can be defined as described in the Platform Software Development Kit (SDK).

The combination of computer name, object, counter, instance, and instance index is known as the *counter path*. The counter path is typically shown in the tools as follows:

*Computer_name**Object_name*(*Instance_name*#*Index_Number*)*Counter_name*

The *computer_name* is optional. If you do not supply a name, your operating system will use the name of the local computer by default.

# Performance tool architecture

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Performance tool architecture**

Various architectural levels of the Windows Server 2003 family support the operation of System Monitor and Performance Logs and Alerts. For example, these performance tools obtain data by default by means of the registry. You can optionally collect performance data by means of the Windows Management Instrumentation (WMI) interface for hardware resources or applications installed on the system. The Performance Data Helper dynamic-link library (DLL) acts as an intermediary between the WMI and the performance tools to format the data and make calculations, where necessary, to convert raw data for display or reporting. For information about configuring performance tools to collect data using WMI, see Collect performance data using WMI.

DLLs supplied by the operating system provide counters for monitoring the behavior of resources such as memory, processors, disks, and network adapters and protocols. In addition, many services in the Windows Server 2003 family provide their own DLLs that install counters that you can use to monitor their operation.

System Monitor and Performance Logs and Alerts use the Remote Registry service process that manages user logon and logoff operations for remote monitoring connections. The Messenger service in the Windows Server 2003 family sends users alert notification. This service must be running for alert notifications to be received.

# Using the Performance console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the Performance console**

This section covers the following topics:

- Setting up a monitoring configuration
- Quick reference to System Monitor settings
- System Monitor buttons and keyboard shortcut reference
- System Monitor interface
- Performance Logs and Alerts interface
- Quick reference to log file formats
- Analyzing performance data
- Solving performance problems
- Monitoring performance from the command line

# Setting up a monitoring configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up a monitoring configuration**
Setting up a monitoring configuration is the first step in evaluating your system's performance. To monitor your Windows Server 2003 family operating system, you can choose to view data in a graph or collect the data in log files for use in other applications.

**Note**

- For best results, read the following sections for setting up your monitoring configuration. To get started quickly, however, you can use the predefined log settings under Counter Logs. These settings, named System Overview, are configured to create a binary log that, after manual startup, updates every 15 seconds and logs continuously until it achieves a maximum size. If you start logging with these settings, data is saved to the Perflogs folder on the root directory and includes the counters: Memory\Pages/sec, PhysicalDisk(_Total)\Avg. Disk Queue Length, and Processor(_Total)\% Processor Time.

**Choosing the monitoring method**
Graphs are useful for short-term, real-time monitoring of a local or remote computer--for example, when you want to observe a system event while it is happening. Choose the update interval to best capture the type of activity you are interested in. Logs are useful for record-keeping and for extended monitoring, especially of a remote computer; logged data can be exported for report generation and presented as graphs or histograms using System Monitor. Logging is the most practical way to monitor multiple computers.

**Choosing how often to monitor**
For routine monitoring, start by logging activity over 15-minute intervals. If you are monitoring for a specific problem, you might want to vary the interval. If you are monitoring the activity of a specific process at a specific time, set a frequent update interval; however, if you are monitoring a problem that manifests itself slowly, such as a memory leak, use a longer interval.

Also consider the overall length of time you want to monitor when choosing this interval. Updating every 15 seconds is reasonable if you plan to monitor for no more than four hours. If you plan to monitor a system for eight hours or more, set an interval longer than 300 seconds (five minutes). Setting the update interval to a frequent rate (low value) can cause the system to generate a large amount of data. This can be difficult to work with and can increase the overhead of running Performance Logs and Alerts.

Monitoring a large number of objects and counters can also generate large amounts of data and consume disk space. Try to strike a balance between the number of objects you monitor and the sampling frequency to keep log file size within manageable limits.

**Choosing the data to monitor**
Start by monitoring the activity of the following components in order:

1. Memory
2. Processors
3. Disks
4. Network

The following table shows the types of counters recommended for server monitoring. When examining specific resources, you should include other counters for the associated performance object.

| Component | Performance aspect being monitored | Counters to monitor |
|---|---|---|

| | | |
|---|---|---|
| Disk | Usage | Physical Disk\Disk Reads/sec Physical Disk\Disk Writes/sec LogicalDisk\% Free Space Interpret the % Disk Time counter carefully. Because the _Total instance of this counter may not accurately reflect utilization on multiple-disk systems, it is important to use the % Idle Time counter as well. Note that these counters cannot display a value exceeding 100%. |
| Disk | Bottlenecks | Physical Disk\Avg. Disk Queue Length (all instances) |
| Memory | Usage | Memory\Available Bytes Memory\Cache Bytes |
| Memory | Bottlenecks or leaks | Memory\Pages/sec Memory\Page Reads/sec Memory\Transition Faults/sec Memory\Pool Paged Bytes Memory\Pool Nonpaged Bytes Although not specifically Memory object counters, the following are also useful for memory analysis: Paging File\% Usage object (all instances) Cache\Data Map Hits % Server\Pool Paged Bytes and Server\Pool Nonpaged Bytes |
| Network | Throughput | Protocol transmission counters (varies with networking protocol); for TCP/IP: Network Interface\Bytes total/sec Network Interface\Packets/sec Server\Bytes Total/sec or Server\Bytes Transmitted/sec and Server\Bytes Received/sec |
| Processor | Usage | Processor\% Processor Time (all instances) |
| Processor | Bottlenecks | System\Processor Queue Length (all instances) Processor\Interrupts/sec System\Context switches/sec |

If some of the counters listed in the preceding table are not available on your computer, verify that you have installed the necessary services or activated the counters.

**Choosing the computer to use for monitoring**

When monitoring computers remotely, you have some options for collecting data. For example, you can run performance logging on the administrator's computer, drawing data continuously from each remote computer. In another case, you can have each computer running the service to collect data and, at regular intervals, run a batch program to transfer the data to the administrator's computer for analysis and archiving. Note that Performance Logs and Alerts allows you to specify a batch file that will run whenever a given log file is closed for a specified log.

Centralized data collection (that is, collection on a local computer from remote computers that you are monitoring) is simple to implement because only one logging service is running. You can collect data from multiple systems into a single log file. It causes additional network traffic, however, and might be restricted by available memory on the administrator's computer. To do centralized data collection, use the Add Counters dialog box to select a remote computer while running Performance Logs and Alerts on your local computer.

Distributed data collection (that is, data collection that occurs on the remote computers that you are monitoring) does not incur the memory and network traffic problems of centralized collection. It does result in delayed availability of the data, however, requiring that the collected data be transferred to the administrator's computer for review. To do distributed data collection, use the Performance Logs and Alerts node within Computer Management on a local computer to select a remote computer on which to collect data.

**Remote monitoring**

If you plan to monitor remote computers, you must have the appropriate administrative credentials to gain access to them. In order to monitor remote systems from your computer, you must specify a Run as account that has access to the remote computer. The Run as option sets the account for the specified log or alert, not for the entire Performance Logs and Alerts service. By default, the service is started under the local computer's NTAuthority\NetworkService account, which generally has permission only to access services and resources on the local computer. To start the service under a different account, use the Run as option when configuring data collection, or use Services under Computer Management, and update the properties of the Performance Logs and Alerts service.

**Keeping records of performance data**

Performance Logs and Alerts provides you with the ability to log performance data to an SQL database. When you retain your logged data in a database, you can query the information and include it in reports. Using database analysis tools, you can query results and examine the results in detail using a variety of parameters. You can also use logs for trend analysis and capacity planning.

**Exporting your monitoring configuration**

After configuring the performance tools with the counters, update intervals, and other settings you want, you can save those settings on your local computer or for export to another computer. System Monitor settings are saved to the .msc file, while Performance Logs and Alerts settings are stored in the registry. If you save the settings under the name Perfmon.msc, note that

you are permanently changing the configuration of the tools on the computer. Therefore, it is recommended that you save the file under another name. You can also save the settings for a single log or alert to an HTML file, which you can later use as a starting point to create new logs or alerts.

In addition to saving monitoring settings, you can add other tools or add multiple copies of System Monitor; for example, if you want to monitor graph and report views simultaneously. For information about adding tools to a console, see Create a custom monitoring console.

# Quick reference to System Monitor settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quick reference to System Monitor settings**

For all of the following settings, start by right-clicking the System Monitor details pane and clicking **Properties**.

| To add or change this | Use this tab |
|---|---|
| View type: graph, histogram, or report | **General** |
| Objects, counters, and instances (add or delete) | **Data**<br><br>**Note**<br><br>• You can also add counters by right-clicking in the details pane of the System Monitor (as opposed to the **System Monitor Properties** dialog box) and clicking **Add Counters**, or by clicking the Add button on the toolbar. You can also delete counters by pressing the **DELETE** key. This works only if you are using the System Monitor (as opposed to the **System Monitor Properties** dialog box). To remove all counters and specify new ones, click the New Counter Set button on the toolbar. |
| Source of data displayed: current data input to the graph, current data input from a log, archived data input from a log | **Source**<br><br>**Note**<br><br>• You can also use the View Log File Data button on the toolbar to use a log file as the data source or the View Current Activity to view real-time data. |
| Time range for a log and view time range | **Source** |
| Update frequency | **General**<br><br>**Note**<br><br>• You can also use the Freeze Display and Update Data buttons on the toolbar to take samples manually. |
| Report or histogram value type (minimum, maximum, average) | **General** |
| Display of counter legend | **General** |
| Display of last, minimum, and maximum values for a selected counter (the value bar); allow multiple counter instances | **General** |
| Color, width, style, or graph line | **Data** |

| | |
|---|---|
| Scale of counter data values | **Data** |
| Background color of results pane surrounding the graph area, color of graph data-display area | **Appearance** |
| Title of graph | **Graph** |
| Label on value axis, vertical or horizontal grid lines, and upper and lower limits of graph axes | **Graph** |
| Font color in graph display or legend | **Appearance** |
| Font type, size, and style | **Appearance** |
| Grid color, timer bar color | **Appearance** |
| Border style, appearance of graph window | **General** |
| Toolbar | **General** |

# System Monitor buttons and keyboard shortcut reference

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Monitor buttons and keyboard shortcut reference**

Many commonly used functions can be easily accessed by icons found on the System Monitor toolbar, or by using keystroke combinations. The following table displays these functions and their corresponding buttons or keystroke combinations.

| Function | Button | Keyboard Shortcut |
|---|---|---|
| New Counter Set | | Ctrl+E |
| Clear Display | | Ctrl+D |
| View Current Activity | | Ctrl+T |
| View Log File Data | | Ctrl+L |
| View Graph | | Ctrl+G |
| View Histogram | | Ctrl+B |
| View Report | | Ctrl+R |
| Add | | Ctrl+I |
| Delete | | Delete Key |
| Paste Counter List | | Ctrl+V |
| Copy Properties | | Ctrl+C |
| Properties | | Ctrl+Q |
| Freeze Display | | Ctrl+F |
| Update Data | | Ctrl+U |
| Highlight | | Ctrl+H |
| Help | | F1 |

# System Monitor interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Monitor interface**
When you open Performance, a toolbar appears and the graph view area displays the same set of counters as the system overview query in Performance Logs and Alerts. System Monitor begins charting counter values in this graph area immediately. By default, data is updated automatically. Switch to manual update by using the **Freeze Display** button in the toolbar. Use the **Update Data** button to collect each sample on demand. One collection is made per click. To remove all data from the display, click the **Clear Display** button.

The names and associated information for the counters you select are shown in the legend beneath the graph, as follows:

- Computer on which System Monitor is running
- Performance object
- Performance counter
- Performance object instance. Note that by default counter instances are listed by name and numerical index. This index appears after the instance name, represented by a number sign (#) and a number. This index makes it easier to monitor multiple instances, for example, when you are monitoring threads of a process.

  To turn off the index display, right-click the graph, click **Properties**, and clear the **Allow duplicate counter instances** check box. By default, process and thread instances have a #n appended (where n is an integer). To append the process ID instead, create a DWORD value called ProcessNameFormat under the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PerfProc\Performance subkey and set its value to 2. Likewise, to append the thread ID instead, create a DWORD value called ThreadNameFormat under the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PerfProc\Performance subkey and set its value to 2. To turn off any of these options and return to the default, set the value of the appropriate DWORD to 1.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. For more information, see Registry Editor overview.

- Graphical properties of each counter.

**Note**

- To match a line in a graph with the counter for which it is charting values, select a counter in the legend and double-click a position in the line. When highlighting is in effect, the bar or line representing data for the selected counter changes color to black for most background colors. If chart lines are close together, try to find a point in the graph where they diverge. Otherwise, System Monitor may not be able to pinpoint the value you are interested in.

Above the legend columns is the value bar, where you see the last, minimum, maximum, and average values for the counter that is currently selected. These values are calculated over the time period and number of samples displayed in the graph (reflected by the **Duration** value), not over the time that has elapsed since monitoring was started. The **Duration** value in the value bar indicates the total elapsed time displayed in the graph (based on the update interval). The movement of the timer bar across the graph indicates the passing of each update interval. Regardless of the update interval, the view will show up to 100 samples. System Monitor compresses log data as necessary to fit it in the display.

You can define the following attributes of the graph:

- Type of display, with options for graph, histogram, or report.

- Background color of the details pane and of the data-display area.
- Size, type, and style of font used to show text in the display.
- Color, width, and style of line used to chart data.

**Note**

- Default key settings in Microsoft Word may conflict with the CTRL+H combination used for System Monitor highlighting. You may need to change these to support highlighting when the System Monitor control (Sysmon.ocx in the *system_root*\System32 folder) is used in Microsoft Word.

Clicking **Object**, **Counter**, **Instance**, or **Computer** in the counter legend sorts entries in ascending or descending order for that category. For example, to sort all counters by name, click **Counter**.

For more information, see Quick reference to System Monitor settings.

# Performance Logs and Alerts interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Performance Logs and Alerts interface**

In Performance Logs and Alerts you define settings for counter logs, trace logs, and alerts. The details pane of the console window shows logs and alerts that you have created. You can define multiple logs or alerts to run simultaneously. Each log or alert is a saved configuration that you define. If you have configured the log for automatic starting and stopping, a single log can generate many individual log data files. For example, if you were generating a log file for each day's activity, one file would close at 11:59 P.M. today, and a new file would open at 12:00 A.M. tomorrow. The following table explains the query summary information provided by the columns in the details pane.

| Column | Description |
|---|---|
| **Name** | This is the name of the log or alert. It describes the type of data you are collecting or the condition you are monitoring. One log can generate multiple log files.<br><br>**Note**<br><br>• Sample settings, named System Overview, have been predefined for counter logging. You can start logging using these settings or define your own settings, as appropriate. |
| **Comment** | This can be any descriptive information about the log or alert. |
| **Log File Type** | This is the log-file format you define. For alert, the type will always be **alerts**. For trace logs, the type can be **sequential** or **circular**. For logs, this can be **binary**, **binary circular**, **text file (comma delimited)**, **text file (tab delimited)** or **SQL**. |
| **Log File Name** | This is the path and base file name you defined for the files generated by this log. The base file name is used for automatically naming new files. |

To see the parameters defined for each log, double-click the list entry for the log. In the dialog box that appears, you can choose how to name your log files, when logging is scheduled to occur, and what performance objects and counters you want to monitor in your log.

If a log is currently running and collecting data (based on the schedule you defined for the log or alert), a green data icon appears next to the log or alert. If a red icon appears, the log or alert has been defined but is not currently running.

**Note**

- You can configure more than one type of log to run at a time. One log can generate multiple log files if the restart option is selected or if the user starts and stops the log multiple times. However, you will not see these individual log files listed in the console window. Use Windows Explorer to view a listing of these files.

# Quick reference to log file formats

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quick reference to log file formats**

The following table provides a description of each log file format. For step-by-step guidance on selecting a format, see Set parameters for a log.

| Log file format | Description | Applies to |
|---|---|---|
| Text File (Comma delimited) | This option defines a comma-delimited log file (with a .csv extension). Use this format, for example, to export the log data into a spreadsheet program. | Counter logs |
| Text File (Tab delimited) | This option defines a tab-delimited log file (with a .tsv extension). Use this format, for example, to export the log data into a spreadsheet program. | Counter logs |
| Binary File | This option defines a sequential, binary-format log file (with a .blg extension). Use this file format if you want to be able to record data instances that are intermittent--that is, stopping and resuming after the log has begun running. Only binary file formats can accommodate instances that are not persistent throughout the duration of the log. | Counter logs |
| Binary Circular File | This option defines a circular, binary-format log file (with a .blg extension). Use this file format to record data continuously to the same log file, overwriting previous records with new data when the file reaches its maximum size. | Counter logs |
| Circular Trace File | This option defines a circular trace log file (with an .etl extension). Use this file format to record data continuously to the same log file, overwriting previous records with new data when the file reaches its maximum size. | Trace logs |
| Sequential Trace File | This option defines a sequential trace log file (with an .etl extension) that collects data until it reaches a user-defined limit and then closes and starts a new file. | Trace logs |
| SQL Database | This option defines the name of an existing SQL database and log set within the database where the performance data will be read or written. Use this file format to collect performance data at an enterprise level rather than a per computer basis. | Counter logs |

# Analyzing performance data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Analyzing performance data**

Analyzing your monitoring data consists of examining counter values that are reported while your system is performing various operations. You should determine which processes are most active and which programs or threads, if any, are monopolizing a resource. Using this type of performance data analysis, you can understand how your system is responding to workload demands.

As a result of this analysis, you may find that your system performs satisfactorily at some times and unsatisfactorily at others. Depending on the causes of these variations and the degree of difference, you may choose to take corrective action or to accept these variations and delay tuning or upgrading resources for a later time.

The level of system performance that you consider acceptable when your system is handling a typical workload and running all required services is its baseline. The baseline performance is a subjective standard that the administrator determines based on the work environment. It may correspond to a range of counter values, including some that are temporarily unacceptable but which generally indicate the best possible performance under the administrator's specific conditions. The baseline can be the measure used for setting performance standards for your users and can be included in any service agreements you put in place.

**Determining acceptable values for counters**

In general, deciding whether or not performance is acceptable is a subjective judgment that varies significantly with variations in user environments. The following table containing threshold values for specific counters can, however, help you determine whether values reported by your system indicate a problem. If System Monitor consistently reports these values, it is likely that bottlenecks exist on your system and you should take action to tune or upgrade the affected resource. Counters that show usage ratios over an interval are a more informative measurement than averages of instantaneous counter values. In performance data measurements, one or two data points that fall outside normal operating conditions for a brief period of time, for example, during the startup period of a service, can skew the average. This gives a false impression of the overall operating performance during the collection period.

| Resource | Object\Counter | Suggested threshold | Comments |
|---|---|---|---|
| Disk | Physical Disk\% Free Space<br>Logical Disk\% Free Space | 15% | |
| Disk | Physical Disk\% Disk Time<br>Logical Disk\% Disk Time | 90% | |
| Disk | Physical Disk\Disk Reads/sec, Physical Disk\Disk Writes/sec | Depends on manufacturer's specifications | Check the specified transfer rate for your disks to verify that this rate does not exceed the specifications. In general, Ultra Wide SCSI disks can handle 50 to 70 I/O operations per second. Note that whether or not the I/O is sequential or random can have a strong impact on values for disk reads/sec and disk writes/sec. |
| Disk | Physical Disk\Current Disk Queue Length | Number of spindles plus 2 | This is an instantaneous counter; observe its value over several intervals. For an average over time, use Physical Disk\Avg. Disk Queue Length. |
| Memory | Memory\Available Bytes | For larger memory computers, greater than 4 MB | Research memory usage and add memory if needed. |
| Memory | Memory\Pages/sec | *n* pages/sec per pagefile | Research paging activity. Note the amount of I/O going to the disks with pagefiles. |
| Paging file | Paging File\% Usage | Above 70% | Review this value in conjunction with Available Bytes and Pages/sec to understand paging activity on your computer. |
| Processor | Processor\% Processor Time | 85% | Find the process that is using a high percentage of processor time. Upgrade to a faster processor or install an additional processor. |

| | | | |
|---|---|---|---|
| Processor | Processor\Interrupts/sec | Depends on processor; 1,000 interrupts per second is a good starting point | A dramatic increase in this counter value without a corresponding increase in system activity indicates a hardware problem. Identify the network adapter, disk, or other hardware causing the interrupts. |
| Server | Server\Bytes Total/sec | | If the sum of Bytes Total/sec for all servers is roughly equal to the maximum transfer rates of your network, you may need to segment the network. |
| Server | Server\Work Item Shortages | 3 | If the value reaches this threshold, consider adding the DWORD entries **InitWorkItems** (the number of work items allocated to a processor during startup) or **MaxWorkItems** (the maximum number of receive buffers that a server can allocate) to the registry (under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters). The entry **InitWorkItems** can range from 1 to 512 while **MaxWorkItems** can range from 1 to 65,535. Start with any value for **InitWorkItems** and a value of 4,096 for **MaxWorkItems** and keep doubling these values until the Server\Work Item Shortages threshold stays below 3. For information about modifying the registry, see Registry Editor overview.<br>**Caution**<br>• Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |
| Server | Server\Pool Paged Peak | Amount of physical RAM | This value is an indicator of the maximum paging file size and the amount of physical memory. |
| Server | Server Work Queues\Queue Length | 4 | If the value reaches this threshold, there may be a processor bottleneck. This is an instantaneous counter; observe its value over several intervals. |
| Multiple processors | System\Processor Queue Length | 2 | This is an instantaneous counter; observe its value over several intervals. |

For tuning and upgrade suggestions, see Solving performance problems.

For a more detailed discussion of performance-data analysis, see "Part One: Operating System Performance" in the System Performance and Troubleshooting Guide, available at the Microsoft Windows Resource Kits Web site.

# Solving performance problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Solving performance problems**
Analysis of your monitoring data can reveal problems such as excessive demand on certain resources resulting in bottlenecks. This section discusses common causes for bottlenecks and a recommended strategy for tuning and testing.

**Causes of bottlenecks**
Demand may become extreme enough to cause resource bottlenecks for the following reasons:

- Resources are insufficient and additional or upgraded components are required.
- Resources are not sharing workloads evenly and need to be balanced.
- A resource is malfunctioning and needs to be replaced.
- A program is monopolizing a particular resource; this might require substituting another program, having a developer rewrite the program, adding or upgrading resources, or running the program during periods of low demand.
- A resource is incorrectly configured and configuration settings need to be changed.

**Strategy for tuning and testing**
Insufficient memory is the most common cause of serious performance problems in computer systems. If you suspect other problems, check memory counters to rule out a memory shortage. Poor response time on a workstation is most likely to result from memory and processor problems; servers are more susceptible to disk and network problems.

Before you start tuning, consider the following recommendations:

- Make one change at a time. In some cases, a problem that appears to relate to a single component might be the result of bottlenecks involving multiple components. For this reason, it is important to address problems individually.

  Making multiple changes simultaneously may make it impossible to assess the impact of each individual change.

- Repeat monitoring after every change. This is important for understanding the effect of the change and to determine whether additional changes are required. Proceed methodically, making one change to the identified resource at a time and then testing the effects of the changes on performance. Because tuning changes can affect other resources, it is important to keep records of the changes you make and to review after you make a change.
- In addition to monitoring, review event logs, because some performance problems generate output you can display in Event Viewer. For more information about using Event Viewer, see Event Viewer.
- To see whether network components are playing a part in performance problems, compare the performance of programs that run over the network with locally run programs.

**Tuning and upgrade tips by component**
This section lists tuning tips for resources you are monitoring:

**Memory:**

- Increase physical memory above the minimum required.
- Create multiple paging files when you use multiple disks.
- Determine the correct size for the paging files. The recommended initial size of the paging file is between 1 and 1.5 times the amount of RAM available, based on your system; the size needed depends on the workload.

- Ensure that memory settings are properly configured.
- Run memory-intensive programs on your highest performing computers or when system workload is light.

**Disk:**

- Upgrade to a higher speed disk, or add disks. If possible, upgrade the disk controller and the bus.
- On servers, use Disk Management to create striped volumes on multiple physical disks. This solution increases throughput because I/O commands can be issued concurrently.
- Distribute programs among servers. Distributed File System (DFS) can be used to balance workload.
- Isolate tasks that heavily utilize disk I/O on separate physical disks or disk controllers.
- Use Disk Defragmenter to consolidate files optimizing data access and disk space.
- If you want to improve the efficiency of disk access, consider installing the latest driver software for your host adapters or motherboard drivers. Contact your adapter or motherboard manufacturer for information.

**Processor:**

- Add a processor (especially for multithreaded programs). Upgrade to a faster processor, or one with a larger cache.
- On multiprocessor computers, manage the processor affinity with respect to process threads and interrupts.

**Network:**

- Configure your network so that systems shared by the same group of people are on the same subnet.
- Unbind infrequently used network adapters. See Modify the protocol bindings order for more information.
- If you are using more than one protocol, you can set the order in which the workstation and NetBIOS software bind to each protocol. See Network and Dial-up Connections Help for more information. Some reasons for changing the list order are as follows:
    - If the protocol you use most frequently is first in the binding list, average connection time decreases.
    - Some protocols are faster than others for certain network topologies. If you are optimizing a client computer, putting the faster protocol first in the bindings list improves performance. Because the server accepts incoming connections using the protocol chosen by the client computer, there is no reason to reorder server computer bindings.
- Use multiple network adapters. The Windows Server 2003 family of operating systems supports multiple adapters for a given protocol and multiple protocols for a given adapter. Although this configuration can create distinct networks that cannot communicate with one another, it is a way to increase file sharing throughput.

For information on server reliability, availability, and performance and tuning features in the Windows Server 2003 family, see Features for Management, Security, and Best Performance.

For detailed guidance on solving performance problems, see "Part One: Operating System Performance" in the System Performance and Troubleshooting Guide, available at the Microsoft Windows Resource Kits Web site.

# Monitoring performance from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring performance from the command line**

In addition to using the Performance graphical interface, you can use the following command line utilities:

- Logman manages and schedules performance counter and event trace log collections on local and remote systems.
- Perfmon allows you to open a Performance console configured with the System Monitor ActiveX control and Performance Logs and Alerts Service.
- Relog extracts performance counters from performance counter logs into other formats, such as **text file (tab delimited)**, **text file (comma delimited)**, **binary-BIN**, or **SQL**.
- Tracerpt processes event trace logs or real-time data from instrumented event trace providers and allows you to generate trace analysis reports and **CSV** (comma-delimited) files for the events generated.
- Typeperf writes performance counter data to the command window, or to a supported log file format.
- Lodctr registers new Performance counter names and Explain text for a service or device driver, and saves and restores counter settings and Explain text.
- Unlodctr removes Performance counter names and Explain text for a service or device driver from the system registry.

For more information about incorporating these Performance commands into your Windows Management Instrumentation (WMI), see "Part One: Operating System Performance" in the System Performance and Troubleshooting Guide, available at the Microsoft Windows Resource Kits Web site.

For information about other command-line utilities, see Command-line reference. For more information about manageability, see Management Strategies and Tools.

# Performance monitoring resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

- For a more detailed discussion of performance data and its analysis, and on troubleshooting your computer using performance-monitoring tools, see "Part One: Operating System Performance" in the System Performance and Troubleshooting Guide, available at the Microsoft Windows Resource Kits Web site.

- For information about optimizing programs or writing performance counters for programs, and about tuning network and service parameters in the registry to optimize performance, see the MSDN library at the Microsoft Web site.

# Troubleshooting Performance monitoring

Updated: April 25, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Counter values consistently equal zero.
- System Monitor shows gaps in its line graphs.
- Values recorded in a log do not appear in the graph view.
- Objects, counters, or instances seem to be missing or invalid.
- I sometimes see an extremely high value for one instance and not another while I am monitoring processes of the same name.
- While monitoring several threads, if one thread stops, its data appears to be reported for another.
- An error message is displayed if I try to export log data to Microsoft Excel while the Performance Logs and Alerts service is actively collecting data to that log.
- I lost my connection to a remote computer from which I was logging data and I cannot resume logging.
- Help for System Monitor does not appear in MMC Help; only Performance Logs and Alerts is displayed.
- I cannot log performance counters from a remote computer.
- My counters used to work, but since I installed a new application they do not work. Is there a way to restore the old counter settings?
- A remote performance counter query might return incorrect zero value for process elapsed time.

**Counter values consistently equal zero.**
**Cause:** The process being monitored has stopped so there is no data in the performance tools.

**Solution:** If you stopped the process manually, restart it to see the process in System Monitor. Otherwise, check Event Viewer for concurrent entries. You may find an error associated with this process.

**Cause:** The counter DLL was disabled after you selected the corresponding counters in a log or display. The performance tools will not detect that the counter was removed or disabled but will report the counter data as zeros.

**Solution:** Enable the corresponding counter DLL.

**Cause:** You lack the appropriate administrative credentials on the computer being monitored.

**Solution:** An administrator must ensure that your user account has the appropriate administrative credentials to use the performance tools.

**System Monitor shows gaps in its line graphs.**
**Cause:** This could be because data collection was subordinated to a processing activity with a higher priority on a system with a heavy load. When the system has adequate resources to continue with data collection, the graphing will resume as usual. A message appears describing this.

**Solution:** Reduce the performance overhead of system monitoring.

**See also:** Performance monitoring best practices

**Values recorded in a log do not appear in the graph view.**
**Cause:** The graph is limited to 100 samples.

**Solution:** Reduce the selected time range. For more information, see Use logged data.

**Cause:** You did not add counters to the graph view.

**See also:** For more information, see Add counters to System Monitor.

**Objects, counters, or instances seem to be missing or invalid.**
**Cause:** Test routines that run when you open the Performance console have detected a problem with installed counters and have disabled the counters automatically to prevent them from slowing the system. Disabled objects and counters do not appear in the **Add Counters** dialog box.

**Solution #1:** You can enable the disabled counters. Using Registry Editor, change the value under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services***Service_name***\Performance\Disable Performance Counters** from **1** for disable to **0** for enable. Note that counters that have been disabled after initial testing are likely to contain errors and may cause system problems. For information about debugging problems with counter DLLs, see the Microsoft Web site.

**Solution #2:** You can turn off the automatic testing feature for the counter DLL or adjust the level of testing. Using Registry Editor, locate the **Configuration Flags** subkey under **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Perflib** and change the value of the **REG_DWORD** entry for this subkey:

| REG_DWORD value | Performance counter DLL test level |
|---|---|
| 0x00 | Test for buffer alignment errors and report any critical error detected. |
| 0x01 | Do not test for data buffer alignment errors. |
| 0x02 | Do not disable DLLs when errors are detected.<br>**Note**<br>• Your operating system will still write messages to the Event Log if this option is used. |
| 0x04 | Disable all performance counter function testing. |
| 0x08 | Enable the disabling of performance objects if they exceed a time-out value.<br>**Note**<br>• Normally your operating system will merely write a message to the Event Log. |
| 0x10 | Enable the collection of performance and error statistics for each performance object. This is the default setting.<br>**Caution**<br>• Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**Cause:** The process that starts the object counters is not started or installed.

**Solution:** Use Task Manager to verify that the process is running. If not, use Exctrlst.exe.

**See also:** Windows Resource Kit Tools Help.

**Cause:** The counters have not been enabled.

**Solution:** Make sure that the service or feature that provides the counter has been installed or configured.

**Cause:** The DLL that installs the counters is generating errors. An example is possible problems with localization functions by the counter.

**Solution:** Check Event Viewer to see whether the counter DLL or the Performance Data Helper reported any errors. If necessary, you can disable counter DLLs that are causing errors by using Exctrlst.exe.

**See also:** Windows Resource Kit Tools Help.

**Cause:** You are trying to monitor a 16-bit or MS-DOS application. Only 32-bit processes appear in the instances list. Active 16-bit processes appear as threads running in a Windows Virtual DOS Machine (NTVDM) process.

**Solution:** Monitor the application through the NTVDM process.

**Cause:** The instance you want to monitor is not currently active. If you are configuring System Monitor to collect real-time data, you can only select active instances for data collection. (If you are viewing logged data, you can select inactive instances for which the log contains data.) If you select the process and it stops after you've selected it, it will continue to appear in the list box but the reported data will be zeros.

**Solution:** Make sure the instance is active.

**I sometimes see an extremely high value for one instance and not another while I am monitoring processes of the same name.**
**Cause:** The performance tools sometimes misrepresent data for separate instances of processes with the same name by reporting the combined values of the instances as the value of a single instance.

**Solution:** Use the instance index and track the Process\ID Process and Process\Creating Process ID counters.

**While monitoring several threads, if one thread stops, its data appears to be reported for another.**
**Cause:** This is caused by the way threads are numbered. For example, assume you are monitoring three threads, numbered 0, 1, and 2. If thread 0 stops, the remaining threads get renumbered. The original thread 0 is now gone and the original thread 1 is renamed 0. As a result, data for the stopped thread 0 could be reported along with data for the running thread 1 because the previous thread 1 is now the current thread 0.

**Solution:** Use the instance index with Thread ID to track these threads.

**An error message is displayed if I try to export log data to Microsoft Excel while the Performance Logs and Alerts service is actively collecting data to that log.**
**Cause:** Microsoft Excel requires exclusive access to the log file. Most other programs do not require this exclusive access; therefore, in general, you can work with data from a log file while the service is collecting data to that file.

**Solution:** Stop the log collection using Performance Logs and Alerts before trying to use the log data with Microsoft Excel.

**I lost my connection to a remote computer from which I was logging data and I cannot resume logging.**
**Cause:** Logging data from a remote computer requires the use of Remote Registry Service. If the service stops due to failure, by default the system restarts it automatically only once.

**Solution:** If Remote Registry service stops more than once, you must restart the service manually on the second and any subsequent failures. To change this default behavior, use Computer Management or Services (on the **Administrative Tools** menu) to modify the properties for Remote Registry service.

**Help for System Monitor does not appear in MMC Help; only Performance Logs and Alerts is displayed.**
**Cause:** Because System Monitor is designed as an ActiveX control, its behavior differs from other Microsoft Management Console (MMC) snap-ins. For example, System Monitor Help is not available when you right-click System Monitor and then click **Help** on the shortcut menu, or when you click **Help Topics** on the **Help** menu.

**Solution:** Click the **Help** button on the System Monitor toolbar.

**I cannot log performance counters from a remote computer.**
**Cause:** You may not have administrative credentials to view the counters on the remote computer.

**Solution:** Ask the administrator of the remote computer to give you read access to the remote computer.

**Cause:** You may not have the correct administrative credentials granted through Group Policy.

**Solution:** Ask the administrator of the remote computer to grant you access through the following User Rights policies: **Profile System Performance** and **Profile a Single Process**.

**See also:** Group Policy (pre-GPMC).

**Cause:** The Performance Logs and Alerts service may not have permission to log on to the remote computer or to create and update the log file.

**Solution:** Use the Run as text box and **Password** button to specify the logon account name with the necessary permissions for the Performance Logs and Alerts service. You can access these features from the **General** tab.

**Cause:** The Remote Registry service is not running on the remote computer.

**Solution:** Ensure that the Remote Registry service is started on the remote computer. You need appropriate permissions on the remote computer to start the service.

**Note**

- You might also receive this error message if the remote computer is running Windows 95. Computer Management does not support remote access to computers that are running Windows 95.

**See also:** Start, stop, pause, resume, or restart a service.

**Cause:** The size limit of the log file may not be large enough to collect your requested data.

**Solution:** Increase the limit of the log file size or set the file size to the maximum limit allowed.

**See also:** Set parameters for a log.

**Cause:** The log update interval may be too short to allow the server to respond to each periodic data request.

**Solution:** Increase the log update interval.

**See also:** View or change properties of a log or alert.

**My counters used to work, but since I installed a new application they do not work. Is there a way to restore the old counter settings?**

**Cause:** Whenever you install a new application that installs performance counters on your operating system, the operating system creates a backup copy of the performance counter registry entries taken prior to the addition of the new application's performance counters file. This backup is saved as a file named PerfStringBackup.ini. If your operating system successfully adds the new application's performance counters file, it then replaces the previous version of the PerfStringBackup.ini file with the file that it just created.

**Solution:** Restore the performance counter registry to the state described in the uncorrupted backup file by using the following command: **LODCTR /R:***filename* where *filename* is the name of the backup file of the performance counter registry entries.

**Note**

- To make a backup copy at any time of the performance counter registry settings, use the following command:

  **LODCTR /S:***filename*

  where *filename* is any name. You can save these backup files anywhere on your computer.

**A remote performance counter query might return incorrect zero value for process elapsed time.**

**Cause:** If an application queries performance counters through remote registry, some performance counters might appear with zero values.This is caused by a remote registry design limitation. The issue will occur if your application calls WNetAddConnection2W with a CONNECT_UPDATE_PROFILE flag prior to opening a remote registry handle.

**Solution:** This is a known issue that is resolved in Windows Server 2008.

# Shutdown Event Tracker

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shutdown Event Tracker**
Shutdown Event Tracker is a routine part of the shutdown process, providing a way for IT professionals to uniformly document why users restart or shut down their computers. The data recorded by Shutdown Event Tracker help to create a comprehensive picture of an organization's system environment. Shutdown Event Tracker is enabled by default and supported on all Microsoft® Windows Server 2003 family operating systems. It is disabled by default and unsupported on Microsoft® Windows® XP Professional.

- For tips on using Shutdown Event Tracker, see Shutdown Event Tracker Best practices.
- For help with specific tasks, see Shutdown Event Tracker How To....
- For general background information, see Shutdown Event Tracker Concepts.
- For problem-solving instructions, see Troubleshooting Shutdown Event Tracker.

# Shutdown Event Tracker Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Plan to use Shutdown Event Tracker as a best practice.**

- You can use the data provided by Shutdown Event Tracker to create a more complete picture of your system environment. For more information, see Shutdown Event Tracker overview.

- To identify patterns of system restarts and shutdowns, review the event logs in Event Viewer on a regular basis. This information serves as a starting point for isolating system problems. For more information, see Event Viewer.

- To create a record of which system resources reached or exceeded their limits just before an unplanned shutdown, use the System State Data feature. When you decide to restart a server running a Windows Server 2003 family operating system, and you select an unplanned reason for doing so, this feature records information about each process, page file, and disk running on the system. It also records data on overall system usage. You can opt to send this information in a System State Data file to Microsoft using the error reporting functionality in the Windows Server 2003 family of operating systems. For more information, see Configure the System State Data feature on your local computer, Configure the System State Data feature on a domain or organizational unit, and Enable notification and reporting of unplanned shutdowns on your local computer. To access Microsoft's privacy policy on error reporting, see Shutdown Event Tracker Resources.

**Consider security when assigning shutdown user rights.**

- In addition to physically securing your server, use Active Directory to uniformly assign shutdown user rights to trusted users. For guidance, see the Access control overview and Best practices for security for security.

**Use standard reasons instead of custom reasons.**

- Using the standard restart and shutdown reasons provided by Microsoft enables you to remotely browse event logs, and to use automated tools that will be available. in the future. In addition, using predefined reasons is a simpler approach than defining your own reasons. For guidance on how to define custom reasons, as well as a comprehensive list of the standard reasons, see "Tools for Troubleshooting" at the Microsoft Windows Resource Kits Web site.

**Establish guidelines across your organization for choosing reasons and making comments.**

- Users should consistently use the restart and shutdown reasons displayed in the graphical user interface (GUI) to ensure the validity of the data collected. As a system administrator, you should establish and communicate specific guidelines for using each reason. You should also provide standards for information that users type in the **Comment** box, including when a comment is required, what information it should contain, and what form that information should take.

  For example, you may want to establish guidelines such as these:

  - If installing new software involves restarting a computer, use the reason **Application: Installation (Planned)**. Always supply the name of the new software in the **Comment** box.

  - If repairing hardware forces you to immediately restart or shut down a computer, use the reason **Hardware: Maintenance (Unplanned)**. Always supply the name of the hardware in the **Comment** box.

# Shutdown Event Tracker How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Configure Shutdown Event Tracker
- Work with Shutdown Event Tracker

# Configure Shutdown Event Tracker

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Shutdown Event Tracker**

- Configure Shutdown Event Tracker on your local computer
- Configure Shutdown Event Tracker on a domain or organizational unit
- Configure the System State Data feature on your local computer
- Configure the System State Data feature on a domain or organizational unit

# Configure Shutdown Event Tracker on your local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Shutdown Event Tracker on your local computer**

1. Open Group Policy.
2. Navigate to the **Display Shutdown Event Tracker** policy, and double-click it.

   **Where?**

   - Local Computer Policy/Computer Configuration/Administrative Templates/System

3. On the **Setting** tab, click **Not Configured**, **Enabled**, or **Disabled**.
4. Click **Apply**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Group Policy, click **Start**, click **Run**, type **gpedit.msc**, and then click **OK**.
- The **Server Only** and **Workstation Only** settings that appear in the user interface (UI) have an impact on the local computer, but they are intended primarily for use by domain administrators.
- To apply the policy setting to a remote computer, see "Group Policy" under Related Topics, below.
- You cannot use Group Policy to manage computers running Windows NT 4.0, Windows 95, Windows 98, Windows Millennium Edition, and Windows XP Home Edition.
- Shutdown Event Tracker is enabled by default on all Windows Server 2003 family operating systems. It is disabled by default on Windows XP Professional. It is not a functional component of Windows XP Home Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy (pre-GPMC)
Enabling or disabling Shutdown Event Tracker
Configure Shutdown Event Tracker on a domain or organizational unit
Shutdown Event Tracker Resources

# Configure Shutdown Event Tracker on a domain or organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Shutdown Event Tracker on a domain or organizational unit**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

   **Where?**

   - Active Directory Users and Computers [*domain_controller_name.domain_name*]/*domain organizational unit/child organizational unit*

3. Click **properties**, and then click the **Group Policy** tab.
4. Click an existing Group Policy object (GPO) in the list box, and then click **Edit**.

   or

   Click **New** to create a new Group Policy object, and then click **Edit**.

5. Click **Display Shutdown Event Tracker**

   **Where?**

   - *Group policy object name*/Computer Configuration/Administrative Templates/System

6. On the **Setting** tab, click **Not Configured**, **Enabled**, or **Disabled**.
7. Click **Apply**, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- If you choose to enable the Group Policy for **Server Only**, Shutdown Event Tracker appears when you shut down a computer running a Windows Server 2003 family operating system.
- If you choose to enable the Group Policy for **Workstation Only**, the Shutdown Event Tracker appears when you shut down a computer running Windows XP Professional.
- To apply the Group Policy setting to a site, see "Group Policy" under Related Topics, below.
- You cannot use Group Policy to manage computers running Windows NT 4.0, Windows 95, Windows 98,

Windows Millennium Edition, and Windows XP Home Edition.

- Shutdown Event Tracker is enabled by default but not configured on all Windows Server 2003 family operating systems. It is disabled by default on Windows XP Professional. It is not a functional component of Windows XP Home Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy (pre-GPMC)
Enabling or disabling Shutdown Event Tracker
Configure Shutdown Event Tracker on your local computer
Shutdown Event Tracker Resources

# Configure the System State Data feature on your local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the System State Data feature on your local computer**

1. Open Group Policy.
2. Navigate to the **Activate Shutdown Event Tracker System State Data feature** policy, and double-click it.

   **Where?**

   - Local Computer Policy/Computer Configuration/Administrative Templates/System

3. On the **Setting** tab, click **Not Configured**, **Enabled**, or **Disabled**.
4. Click **Apply**, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Group Policy, click **Start**, click **Run**, type **gpedit.msc**, and then click **OK**.
- When this policy is enabled, Shutdown Event Tracker writes out a system state data file when a user with shutdown user rights designates a restart or shutdown as unplanned. This file is stored in the *%windir%\system32\LogFiles\Shutdown* directory, and only members of the Administrators group may view it or optionally send it to Microsoft.
- You can configure the number of days a system state data file will be saved to the *%windir%\system32\LogFiles\Shutdown* directory using the registry. See "Resources" in Related Topics for more information.
- To view Microsoft's privacy policy on error reporting, see "Resources" in Related Topics.
- When this policy is disabled, Shutdown Event Tracker will not create a system state data file in the event of an unplanned restart or shutdown.
- If you do not configure the System State Data feature, the default behavior occurs: By default, this feature is always enabled on the Windows Server 2003 family operating systems.
- To apply the policy setting to a remote computer, see "Group Policy" under Related Topics, below.
- You cannot use Group Policy to manage computers running Windows Server 4.0, Windows 95, Windows 98, Windows Millennium Edition, and Windows XP Home Edition.
- Administrative Templates provide a friendlier user interface (UI) for editing the registry than does Registry Editor (Regedit.exe), as well as an added degree of security: They expose only the registry keys that Microsoft has determined to be frequently useful.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable error reporting
Group Policy (pre-GPMC)
Shutdown Event Tracker overview
Configure the System State Data feature on a domain or organizational unit
Shutdown Event Tracker Resources

# Configure the System State Data feature on a domain or organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the System State Data feature on a domain or organizational unit**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

   **Where?**

   - Active Directory Users and Computers [*domain_controller_name.domain_name*]/*domain/organizational unit/child organizational unit*

3. Click **properties**, and then click the **Group Policy** tab.

4. Click an existing Group Policy object in the list box, and then click **Edit**.

   or

   Click **New** to create a new Group Policy object, and then click **Edit**.

5. Click **Activate Shutdown Event Tracker System State Data feature**

   **Where?**

   - *Group policy object name*/Computer Configuration/Administrative Templates/System

6. On the **Setting** tab, click **Not Configured**, **Enabled**, or **Disabled**.

7. Click **Apply**, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- When this policy is enabled, Shutdown Event Tracker writes out a system state data file when a user with shutdown rights designates a restart or shutdown as unplanned. This file is stored in the *%windir%\system32\LogFiles\Shutdown* directory, and only members of the Administrators group may view it or, optionally, send it to Microsoft.

- You can configure the number of days a system state data file will be saved to the *%windir%\system32\LogFiles\Shutdown* directory using the registry. See "Resources" in Related Topics for more information.

- To view Microsoft's privacy policy on error reporting, see "Resources" in Related Topics.

- When this policy is disabled, Shutdown Event Tracker does not create a system state data file.
- This policy is not configured by default on Windows XP Professional and Windows Server 2003 family operating systems.
- To apply the policy setting to a site, see "Group Policy" under Related Topics, below.
- You cannot use Group Policy to manage computers running Windows Server 4.0, Windows 95, Windows 98, Windows Millennium Edition, and Windows XP Home Edition.
- Administrative Templates provide a friendlier user interface (UI) for editing the registry than does Registry Editor (Regedit.exe), as well as an added degree of security: They expose only the registry keys that Microsoft has determined to be frequently useful.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Active Directory Users and Computers (Windows Server 2003)
Enable error reporting
Group Policy (pre-GPMC)
Shutdown Event Tracker overview
Configure the System State Data feature on your local computer
Shutdown Event Tracker Resources

# Work with Shutdown Event Tracker

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with Shutdown Event Tracker**

- Restart or shut down locally and document the reason
- Restart or shut down remotely and document the reason
- Cancel a restart or shutdown
- Document the reason for an unexpected restart or shutdown of your local computer
- Document the reason for unexpected restarts or shutdowns of remote computers
- Enable notification and reporting of unplanned shutdowns on your local computer

# Restart or shut down locally and document the reason

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To restart or shut down locally and document the reason

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Click **Start**, and then click **Shut Down** or **Restart**.
2. Under **What do you want the computer to do**, click **Restart** or **Shut down**.
3. If you did not anticipate having to restart or shut down the computer at this time, clear the **Planned** check box.
4. Select the appropriate reason from the list.
5. If a comment is required, the **OK** button will not function until you type at least one character in the text box.

### Important

- Do not turn off your computer until a message appears telling you that it is safe to do so. Windows Server 2003 family operating systems store important data in memory while the system is running and need to write the data to the hard disk before you turn off the computer. After the data is saved, the operating system notifies you that it is okay to turn off the computer.

### Notes

- To perform this procedure, you must have the user rights to shut down the local system. For a list of security groups that have this permission by default on computers running Windows Server 2003 family operating systems, see "Privileges" under Related Topics.
- An alternative way to open the expected shutdown dialog box is to press CTRL+ALT+DELETE and then click **Shut Down**.
- Shutdown Event Tracker records data only for the **Restart** and **Shutdown** options.
- If the System State Data feature is enabled and you specify an unplanned reason for the restart or shutdown, a file containing relevant system data will be created in the *%windir%\system32\LogFiles\Shutdown* directory. The System State Data feature is enabled by default on all Windows Server 2003 family operating systems.
- For information on defining custom reasons, see "Best practices" under Related Topics.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **shutdown /s /d 1:1** to initiate an unplanned shut down after a default wait period of 30 seconds, with the major reason, **Hardware**, and the minor reason, **Maintenance**.

Or,

Type:

**shutdown /r /t 60 /d p:4:2** to initiate a planned startup after waiting one minute, with the major reason, **Application**, and the minor reason, **Installation**.

Or,

Type:

**shutdown /p /d p:1:2** to turn off the power to your computer with no time-out period or warning and to indicate that the action was planned with the major reason **Hardware** and the minor reason **Installation**.

| Value | Description |
|---|---|
| **/s** | Shuts down the computer. |
| **/r** | Restarts the computer after shutdown. |
| **/p** | Turns off the local computer with no time-out period or warning. You can use **/p** only with **/d**. If your computer does not support power off functionality, it will shut down when you use **/p**, but the power to the computer will remain on. |
| **/t** *nnn* | Sets the wait period before a restart or shutdown to *xxx* seconds, causing a warning to display on the local console. You can specify from 0 through 600 seconds. If you omit **/t**, the wait period defaults to 30 seconds. |
| **/d**[**p:**]*xx*:*yy* | Lists the reason for the system restart, shutdown, or power off. The following rows describe the parameter values. |
| **p:** | Indicates that the restart or shutdown is planned. If you do not use the **p:** option, Shutdown Event Tracker assumes that the restart or shutdown is unplanned. |
| *xx* | Specifies the major reason number (0-255). |
| *yy* | Specifies the minor reason number (0-65535). |

**Notes**

- To perform this procedure, you must have the user rights to shut down the local system. For a list of security groups that have this permission by default on computers running Windows Server 2003 family operating systems, see "Privileges" under Related Topics.
- For the complete command line syntax, see "Shutdown" under Related Topics, below.
- If the System State Data feature is enabled and you specify an unplanned reason for the restart or shutdown, a file containing relevant system data will be created in the *%windir%\system32\LogFiles\Shutdown* directory. The System State Data feature is enabled by default on all Windows Server 2003 family operating systems.
- For information on defining custom reasons, see "Best practices" under Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restart or shut down remotely and document the reason
Shutdown

Shutdown Event Tracker Resources
Shutdown Event Tracker Best practices
Privileges

# Restart or shut down remotely and document the reason

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restart or shut down remotely and document the reason**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Command Prompt.
2. Type:

   **shutdown /i** to display the Remote Shutdown dialog box (Shutdown.exe).
3. Under **Computers**, click **Add** to enter computer names, or click **Browse** to open the **Find Computers** dialog box.
4. Under **What do you want these computers to do**, click **Restart** or **Shut down**.
5. If you did not anticipate restarting or shutting down the target computers at this time, clear the **Planned** check box.
6. Select the appropriate reason from the list.
7. If a comment is required, the **OK** button will not function until you type at least one character one in the text box.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To perform this procedure, you must have the user rights to force a shutdown from a remote system. For a list of security groups that have this permission by default on computers running Windows Server 2003 family operating systems, see "Privileges" under Related Topics.
- The Remote Shutdown dialog box can be displayed only by typing **shutdown /i** at the command prompt.
- If the remote computer is functioning properly, but Shutdown.exe fails to restart it or shut it down, wait for two minutes before making a second attempt: Either the log on dialog box or the unexpected shutdown dialog box is open on the remote computer, causing an error.
- Shutdown Event Tracker records data only for the **Restart** and **Shutdown** options.
- If the System State Data feature is enabled and you specify an unplanned reason for the restart or shutdown, a file containing relevant system data will be created in the *%windir%\system32\LogFiles\Shutdown* directory. The System State Data feature is enabled by default on all Windows Server 2003 family operating systems.
- For information on defining custom reasons, see "Best practices" under Related Topics.
- The **Warn users of the action** check box is selected by default.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **shutdown /s /m ***ComputerName***/d 1:1** to initiate an unplanned shutdown after a default wait period of 30 seconds, with the major reason, **Hardware**, and the minor reason, **Maintenance**.

   Or,

   Type:

   **shutdown /r /m ***ComputerName***/t 60 /d p:4:2** to initiate a planned startup after waiting for one minute, with the major reason, **Application**, and the minor reason, **Installation**.

| Value | Description |
|---|---|
| **/s** | Shuts down the computer. |
| **/r** | Restarts the computer after shutdown. |
| **/m ** *ComputerName* | Specifies the target computer. |
| **/t** *nnn* | Sets the wait period before a restart or shutdown to *nnn* seconds, causing a warning to display on the local console. You can specify from 0 through 600 seconds. If you omit **/t**, the wait period defaults to 30 seconds. |
| **/d**[**p:**]*xx***:***yy* | Lists the reason for the system restart, shutdown, or power off. The following rows describe the parameter values. |
| **p:** | Indicates that the restart or shutdown is planned. If you do not use the **p:** option, Shutdown Event Tracker assumes that the restart or shutdown is unplanned. |
| *xx* | Specifies the major reason number (from 0 through 255). |
| *yy* | Specifies the minor reason number (from 0 through 65535). |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To perform this procedure, you must have the user rights to force a shutdown from a remote system. For a list of security groups that have this permission by default on computers running Windows Server 2003 family operating systems, see "Privileges" under Related Topics.

- For the complete command line syntax, see "Shutdown" under Related Topics, below.

- If the remote computer is functioning properly, but Shutdown.exe fails to restart it or shut it down, wait for 2 minutes before making a second attempt: Either the log on dialog box or the unexpected shutdown dialog box is open on the remote computer, causing an error.

- If the System State Data feature is enabled and you specify an unplanned reason for the restart or shutdown, a file containing relevant system data will be created in the *%windir%\system32\LogFiles\Shutdown* directory. The System State Data feature is enabled by default on all Windows Server 2003 family operating systems.

- For information on defining custom reasons, see "Best practices" under Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restart or shut down locally and document the reason
Shutdown
Shutdown Event Tracker Resources
Shutdown Event Tracker Best practices
Configure the System State Data feature on your local computer
Configure the System State Data feature on a domain or organizational unit
Privileges

# Cancel a restart or shutdown

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To cancel a restart or shutdown**

1. Open Command Prompt.
2. Type:

   **shutdown/a** within the time-out period.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To perform this procedure on a local computer, you must have user rights to shut down the computer. To perform this procedure on a remotely managed computer, you must have the user rights to force a shutdown from a remote system.
    - You can set the time-out period before a restart or shutdown to *nnn* seconds using the **/t***nnn* option. The valid range is from 0 through –600 seconds, with a default of 30 seconds.
    - You cannot cancel a restart or shutdown using the Shutdown Event Tracker user interface (UI). It is possible to perform this task only from the command line.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restart or shut down locally and document the reason
Restart or shut down remotely and document the reason
Shutdown
Shutdown Event Tracker Resources

# Document the reason for an unexpected restart or shutdown of your local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To document the reason for an unexpected restart or shutdown of your local computer**

1. Under **Why did the computer shut down unexpectedly,** specify the appropriate reason from the list.
2. If your organization has a trouble-ticketing system, you can enter a **Problem Id**.
3. If a comment is required, the **OK** button will not function until you type at least one character in the text box.

### Note

- To perform this procedure, you must have the user rights to shut down the local system. For a list of security groups that have this right by default on computers running the Windows Server 2003 family operating systems, see "Privileges" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Document the reason for unexpected restarts or shutdowns of remote computers
Shutdown Event Tracker overview
Shutdown Event Tracker Resources
Privileges

# Document the reason for unexpected restarts or shutdowns of remote computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To document the reason for unexpected restarts or shutdowns of remote computers**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Command Prompt.
2. Type:

   **shutdown /i** to display the Remote Shutdown dialog box (Shutdown.exe).
3. Under **Computers**, click **Add** to enter computer names, or click **Browse** to open the **Find Computers** dialog box.
4. Under **What do you want these computers to do**, click **Annotate Unexpected Shutdown**.
5. Select the appropriate reason for the unexpected shutdown from the list.
6. If a comment is required, the **OK** button will not function until you type at least one character one in the text box.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- The Remote Shutdown dialog box (Shutdown.exe) can only be displayed by typing **shutdown /i** at the command prompt.
- For information on defining custom reasons, see "Best practices" in Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type: **shutdown/e/m**_ComputerName_**/c "**_Comment_**" /d** xx**:**yy

| Value | Description |
|---|---|
| **/e** | Enables you to document the reason for the unexpected shutdown on the remote computer. |

| | |
|---|---|
| **/m** *ComputerName* | Specifies the target computer. |
| **/c** | Enables you to comment in detail about the reason for the shutdown. You can use a maximum of 511 characters. Comments must be enclosed in quotation marks. |
| **/d** *xx* **:** *yy* | Lists the reason for the system restart or shutdown, where *xx* specifies a major reason number from 0 through 255 and *yy* specifies a minor reason number from 0 through 65,535. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- For the complete command line syntax, see "Shutdown" under Related Topics, below.
- For information on defining custom reasons, see "Best practices" under Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shutdown Event Tracker Best practices
Document the reason for an unexpected restart or shutdown of your local computer
Shutdown Event Tracker overview
Shutdown Event Tracker Resources
Shutdown

# Enable notification and reporting of unplanned shutdowns on your local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable notification and reporting of unplanned shutdowns on your local computer**

1. Enable Shutdown Event Tracker. See "Configure Shutdown Event Tracker on your local computer" in Related Topics for instructions.

2. Open Group Policy.

3. Navigate to **Activate Shutdown Event Tracker System State Data feature**, and double-click it.

   **Where?**

   - Local Computer Policy/Computer Configuration/Administrative Templates/System

4. On the **Setting** tab, click **Enabled**.

5. Click **Apply**, and then click **OK**.

6. Navigate to the **Configure Error Reporting** policy, and double-click it.

   **Where?**

   - Local Computer Policy/Computer Configuration/Administrative Templates/System/Error Reporting

7. On the **Setting** tab, click **Enabled**.

8. Click **Apply**, and then click **OK**.

9. Navigate to the **Report unplanned shutdown events** policy, and double-click it.

   **Where?**

   - Local Computer Policy/Computer Configuration/Administrative Templates/System/Error Reporting/Advanced Error Reporting settings

10. On the **Setting** tab, click **Enabled**.

11. Click **Apply**, and then click **OK**.

12. Navigate to the **Display Error Notification** policy, and double-click it.

    **Where?**

    - Local Computer Policy/Computer Configuration/Administrative Templates/System/Error Reporting

13. On the **Setting** tab, click **Enabled**.

14. Click **Apply**, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- The above procedure causes a system state data file to be created in the *%windir%\system32\LogFiles\Shutdown* directory prior to an unplanned restart or shutdown. The first user with administrative credentials who subsequently logs on to the computer will see a dialog box notification reading "The system has restarted after an unplanned shutdown. A log of this event has been created." From this dialog box, the administrator can navigate to a more detailed description of the system state data file, its contents, and Microsoft's privacy policy for data collection on the Web. The administrator can choose whether or not to send the system state data file to Microsoft by clicking either the **Send Error Report** or **Don't Send** button.

- To enable the creation of the system state data file and notification, but disable error reporting, set all of the policies to **Enabled** except for the **Configure Error Reporting** policy, which should be set to **Disabled**.

- To bypass the administrator and send the error report automatically, set all of the policies to **Enabled** except for the **Display Error Notification** policy, which should be set to **Disabled**.

- You have the option of enabling Corporate Error Reporting by setting a path for uploading data to a central location on your network. You can also customize the error reporting dialogs by replacing instances of the word "Microsoft" with your company's name. For more information, double-click the **Configure Error Reporting** policy and click the **Explain** tab.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restart or shut down locally and document the reason
Restart or shut down remotely and document the reason
Configure Shutdown Event Tracker on your local computer
Shutdown Event Tracker Resources

# Shutdown Event Tracker Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Shutdown Event Tracker and the System State Data feature:

- Shutdown Event Tracker overview
- Using Shutdown Event Tracker
- Shutdown Event Tracker Resources

# Shutdown Event Tracker overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shutdown Event Tracker overview**
Shutdown Event Tracker provides a way for IT professionals to consistently track why users restart or shut down their computers. It does not document why users choose other options, such as **Log off** and **Hibernate**. It gathers the reasons users give for restarts and shutdowns to help create a comprehensive picture of an organization's system environment. Shutdown Event Tracker is enabled by default and supported on all Microsoft® Windows Server 2003 family of operating systems. It is disabled by default and unsupported on Microsoft® Windows® XP Professional.

**"Expected" and "unexpected" startups and shutdowns**
When Shutdown Event Tracker is enabled, the "expected" shutdown dialog box appears when users click **Start** and then **Shutdown**, or when they press CTRL+ALT+DELETE and then click **Shutdown**. This dialog box is different from the standard Windows shutdown dialog box in that it prompts users to supply a reason and a comment to explain the action. An expected startup or shutdown provides the operating system time to complete its usual shutdown routine. By contrast, the computer cannot anticipate an "unexpected" startup or shutdown. If Shutdown Event Tracker is enabled, the unexpected shutdown dialog box appears to the first person with shutdown user rights who logs on to the computer after the startup or shutdown. Like the expected shutdown dialog box, it prompts this user to supply a reason and a comment.

**Planned and unplanned startups and shutdowns**
From the user's point of view, an expected system startup or shutdown can be planned or unplanned. When users have control over the timing of a startup or shutdown, the task is planned. For example, the IT department may reserve a specific time at which to install new applications. By contrast, an unplanned restart or shutdown forces users to immediately perform the task. For example, unresponsive applications might suddenly force users to restart their computers.

Unexpected restarts or shutdowns can also be planned or unplanned. For example, users sometimes choose to shut down their computers by pressing the power button instead of clicking **Start** and then **Shutdown**, or by pressing CTRL+ALT+DELETE and then clicking **Shutdown**. In this case, the shutdown is unexpected by the computer and planned by the user. However, if the computer's power cord is accidentally disconnected, the shutdown is both unexpected and unplanned. In each case, the unexpected dialog box appears to the first person with shutdown user rights to log on to the computer after the event.

For step-by-step guidance and specific examples, see Restart or shut down locally and document the reason and Restart or shut down remotely and document the reason.

**Local and remote startups and shutdowns**
Shutdown Event Tracker and the Shutdown.exe tool enable users to restart or shut down a local computer and one or more remote computers by using either the graphical user interface (GUI) or the Shutdown command. Additionally, IT professionals can perform remote bulk annotations of unexpected shutdowns, an alternative to the time-consuming task of logging on to each computer to record a reason for an unexpected shutdown.

**Root-cause analysis of startups and shutdowns**
Shutdown Event Tracker records the reason for each shutdown or startup through the Event Log service. You can use Event Viewer to open the system log and look for a pattern of events to find the cause of frequent system restarts and shutdowns. For more information, see Event Viewer.

The System State Data feature gathers information for root-cause analysis of unplanned shutdowns. System state data is recorded in a log file when a user who has shutdown user rights specifies an unplanned reason for shutting down a computer. Regardless of who gathers the data, only users with administrative credentials can view it or, optionally, send it to Microsoft. The first person with administrative credentials to log on to the computer after the unplanned shutdown encounters a Windows error reporting dialog box with a link to the system state data file. This file is stored in the *%windir%\system32\LogFiles\Shutdown* directory.

The System State Data feature is available on all Windows Server 2003 family of operating systems. It is not available on Windows XP Professional. For step-by-step guidance, see Configure the System State Data feature on your local computer and Configure the System State Data feature on a domain or organizational unit.

# Using Shutdown Event Tracker

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Shutdown Event Tracker**

This section covers:

- Enabling or disabling Shutdown Event Tracker
- Using Shutdown Event Tracker with non-English Windows Server 2003 family products
- Managing Shutdown Event Tracker from the command line

# Enabling or disabling Shutdown Event Tracker

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling or disabling Shutdown Event Tracker**

You can use either Group Policy to enable or disable Shutdown Event Tracker on your system. For details on how to configure Shutdown Event Tracker by setting the **Display Shutdown Event Tracker** policy, see Configure Shutdown Event Tracker on your local computer or Configure Shutdown Event Tracker on a domain or organizational unit.

When Shutdown Event Tracker is enabled, the expected shutdown dialog box appears when you click **Start** and then click **Shutdown**, or when you press CTRL+ALT+DELETE and then click **Shutdown**. Should an unexpected restart or shutdown occur, the unexpected shutdown dialog box will appear to the first person with shutdown user rights who logs on to the computer.

When Shutdown Event Tracker is disabled, the standard Windows shutdown dialog box appears. It provides a way to restart or shut down your computer; however, it does not provide a way to record whether or not the event is planned or unplanned, or a way to record a reason for it. Should an unexpected restart or shutdown occur, the unexpected shutdown dialog box will *not* appear to the first person with shutdown rights who logs on to the computer. For definitions of planned and unplanned, and expected and unexpected startups and shutdowns, see Shutdown Event Tracker overview.

# Using Shutdown Event Tracker with non-English Windows Server 2003 family products

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Shutdown Event Tracker with non-English Windows Server 2003 family products**

You can configure Shutdown Event Tracker to display standard and custom reasons in languages other than English by modifying the registry. For example, to display the standard shutdown reasons supplied by Microsoft in a language such as Romanian, your operating system's Locale must also be in Romanian, and the Windows Server 2003 family operating system on your computer must also be Romanian. However, custom reasons that you store under a designated subkey will appear in whatever language they are stored. For example, if you define reasons under this subkey in Chinese characters, Shutdown Event Tracker will attempt to display the Chinese characters on the Romanian Windows Server 2003 family operating system. If you have installed the appropriate code page, the Chinese characters will appear. If you have not installed the code page, question marks (?) will appear instead.

For information on how to edit the registry to enable multilanguage support, see "Tools for Troubleshooting" at the Microsoft Windows Resource Kits Web site.

# Managing Shutdown Event Tracker from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Shutdown Event Tracker from the command line**
In addition to using the graphical user interface (GUI) to record restart and shutdown reasons, you can use the command line. For details, see the following:

- Restart or shut down locally and document the reason
- Restart or shut down remotely and document the reason
- Document the reason for unexpected restarts or shutdowns of remote computers

See Shutdown for specific usage information.

# Shutdown Event Tracker Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

See "Tools for Troubleshooting" at the Microsoft Windows Resource Kits Web site for the following information:

- How to manage Shutdown Event Tracker using the registry
- How to turn off the standard restart and shutdown reasons provided by Microsoft
- How to define your own reasons

To view Microsoft's privacy policy in relation to the System State Data feature, see Microsoft Error Reporting.

For information on error and event log messages, technical support, and tools, see Support resources.

# Troubleshooting Shutdown Event Tracker

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Shutdown Event Tracker does not appear when you attempt to shut down your computer. (The standard Windows shutdown dialog box appears. It does not allow you to supply a reason or a comment.)
- A System State Data file is not created before an unplanned shutdown: Upon checking the Event Log for the corresponding 1074 event, you see that the System State Data file name does not appear in the event description.
- You initiated the shutdown sequence using the command line, but the shutdown reason text in the event description is incorrect.
- When you log on to a computer, the error reporting dialog box appears even though the previous shutdown was planned.
- The OK button is disabled after you choose a shutdown reason.
- After you upgrade a server from Windows NT to a Windows Server 2003 family operating system, users who were previously able to shut down the computer can no longer do so.
- When you attempt to shut down a remote computer, you are denied access.
- The Shutdown tool (Shutdown.exe) fails to restart or shut down a remote computer, although the target computer is functioning properly.

**Shutdown Event Tracker does not appear when you attempt to shut down your computer. (The standard Windows shutdown dialog box appears. It does not allow you to supply a reason or a comment.)**
**Cause:** Shutdown Event Tracker is disabled.

**Solution:** If you have not yet enabled Shutdown Event Tracker using Group Policy, do so.

**See also:** Configure Shutdown Event Tracker on your local computer and Configure Shutdown Event Tracker on a domain or organizational unit.

**Cause:** You do not have the rights to shut down the computer.

**Solution:** Assign yourself the rights to shut down each computer that you manage or have your system administrator assign you the appropriate user rights.

**See also:** Privileges and Security overview

**Cause:** An out-of-date shutdown tool from an earlier Resource Kit is installed.

**Solution:** Uninstall any out-of-date shutdown tool.

**Cause:** The Shutdown Event Tracker Group Policy setting has not yet taken effect.

**Solution:** Allow more time for the Group Policy setting to update itself. If this strategy does not yield results, use the gpupdate.exe tool.

**See also:** Group Policy (pre-GPMC) and Gpupdate.

**Cause:** A higher-level Group Policy configuration takes precedence.

**Solution:** If the Group Policy settings are at a higher level and you do not have the proper administrative credentials to change it, contact your system administrator.

**See also:** Windows Server 2003 Technical Reference at the Microsoft Web site, Group Policy (pre-GPMC), and Configure Shutdown Event Tracker on a domain or organizational unit.

**A System State Data file is not created before an unplanned shutdown: Upon checking the Event Log for the corresponding 1074 event, you see that the System State Data file name does not appear in the event description.**
**Cause:** The System State Data feature is disabled.

**Solution:** If you have not yet enabled the System State Data feature using Group Policy, do so. If the problem persists, consult the Microsoft Resource Kit chapter, below.

**See also:** Configure the System State Data feature on your local computer, Configure the System State Data feature on a domain or organizational unit, and "Tools for Troubleshooting" at the Microsoft Windows Resource Kits Web site.

**Cause:** You initiated the shutdown on a computer running Windows XP Professional.

**Solution:** Windows XP Professional does not include the System State Data feature.

**You initiated the shutdown sequence using the command line, but the shutdown reason text in the event description is incorrect.**
**Cause:** The appropriate reason code was not defined on the computer you shut down.

**Solution:** Add the same set of reasons to all the computers in your organization.

**See also:** "Tools for Troubleshooting" at the Microsoft Windows Resource Kits Web site.

**When you log on to a computer, the error reporting dialog box appears even though the previous shutdown was planned.**
**Cause:** The person who performed the previous shutdown specified it as unplanned. This is likely if the shutdown was initiated from the command line, because omitting **p:** causes Shutdown Event Tracker to record an unplanned event.

**Solution:** You cannot use Shutdown Event Tracker to change or annotate the details of an unplanned shutdown that was expected by the computer, after the fact. In other words, the user who initiated the shutdown did so through the Shutdown Event Tracker user interface (UI) or the command line instead of pulling the power cord or pushing the power button. If the unplanned shutdown was initiated through a different tool, contact that vendor to see if an updated version exists, and whether it supports planned and unplanned shutdown flagging.

**See also:** Shutdown Event Tracker overview, Restart or shut down locally and document the reason, and Restart or shut down remotely and document the reason.

**The OK button is disabled after you choose a shutdown reason.**
**Cause:** A comment is required for this reason.

**Solution:** Type explanatory text in the **Comment** box.

**After you upgrade a server from Windows NT to a Windows Server 2003 family operating system, users who were previously able to shut down the computer can no longer do so.**
**Cause:** The shutdown user rights assigned to users on Windows NT systems will be automatically revoked during an upgrade to a Windows Server 2003 family operating system. This is done for security purposes, to ensure that user rights are uniformly assigned.

**Solution:** Use group accounts in Active Directory to ensure that shutdown user rights are uniformly assigned to trusted users.

**See also:** Privileges, Security overview, and Restart or shut down remotely and document the reason.

**When you attempt to shut down a remote computer, you are denied access.**
**Cause:** You must have the appropriate user rights on the remote computer before you attempt to shut it down.

**Solution:** Establish the appropriate user rights on all computers that you must remotely manage. Being able to shut down a computer locally does not mean that you have the right to shut it down remotely.

**See also:** Privileges, Security overview, and Restart or shut down remotely and document the reason.

**The Shutdown tool (Shutdown.exe) fails to restart or shut down a remote computer, although the target computer is functioning properly.**
**Cause:** There are four possible causes:

- The target computer has just restarted and is applying policies.
- A shut down of the target computer is in progress or has already occurred.

- The log on dialog box is open on the target computer.
- An expected or unexpected shutdown dialog box is open on the target computer.

**Solution:** In all of these cases, the operating system is functioning properly. For the last two causes, wait for a few minutes before making another attempt to restart or shut down the target computer.

**See also:** Restart or shut down remotely and document the reason.

# SNMP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SNMP**
Simple network management protocol (SNMP) provides the ability to monitor and communicate status information between a variety of hosts.

- Before installing SNMP, see Checklist: Implementing the SNMP service.
- To find features that have been moved in the Microsoft® Windows Server 2003 family, see New ways to do familiar SNMP tasks.
- For tips about using SNMP, see SNMP Best practices.
- For help with specific tasks, see SNMP How To...
- For general background information, see SNMP Concepts.
- For problem-solving instructions, see Troubleshooting SNMP.

# Checklist: Implementing the SNMP service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Implementing the SNMP service**

| Step | Reference |
|---|---|
| Obtain, install, and configure SNMP management software. | The SNMP service |
| Gather information required to implement SNMP on your network. Requirements include contact persons (administrator), physical computer location, configured SNMP community names, and IP or IPX addresses, or computer names of SNMP management systems on your network. | Using SNMP |
| Install SNMP on this computer. | Install the SNMP service |
| Configure SNMP agent properties. | Configure agent properties |
| If you have an additional TCP/IP service in your network, such as a bridge or router, you should consult RFC 1213 for additional configuration information before proceeding. | "Simple Network Management Protocol" at the Microsoft Windows Resource Kits Web site |
| Configure trap destinations. | Configure traps |
| Configure SNMP security properties. | Configure SNMP security properties |

**Note**

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# New ways to do familiar SNMP tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring Simple Network Management Protocol (SNMP). The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Install the SNMP service. | The SNMP Service option in Windows NT TCP/IP Installation Options. | Windows Components wizard. For more information, see Install the SNMP service. |
| Configure the SNMP agent. | Control Panel, Network, **Services** tab. | **Services** in the Computer Management console. For more information, see Configure agent properties. |
| Configure SNMP traps. | Control Panel, Network, **Services** tab. | **Services** in the Computer Management console. For more information, see Configure traps. |
| Configure SNMP security. | Control Panel, Network, **Services** tab. | **Services** in the Computer Management console. For more information, see Configure SNMP security properties. |

# SNMP Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Update settings regularly.**

- Make sure that you monitor and update the host name settings on an ongoing basis. This will help ensure timely detection of any unauthorized access.

**Allow hosts to accept packets only from specific hosts.**

- When you configure security options, do not click **Accept packets from any host**.

**Configure authentication traps.**

- Take advantage of SNMP security checking by configuring authentication traps on all SNMP agents.

**Verify proper functioning of service-specific components.**

- If you will be monitoring service-specific components, such as Dynamic Host Configuration Protocol (DHCP) or Windows Internet Name Service (WINS), verify that these services have been successfully installed and configured.

**Use IPSec to secure SNMP messages.**

- Reduce the risk of internal attacks by configuring the SNMP agent systems to reject request messages from unauthorized management systems. For security purposes, use Internet Protocol security (IPSec) to secure SNMP messages by creating filter specifications in the appropriate IP filter list between SNMP management systems and agents.

- Remember that SNMP is an insecure protocol.

- If you decide to use SNMP to manage networks, remember that SNMP is an insecure protocol whose effectiveness depends solely on how you implement it.

# SNMP How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install the SNMP service
- Configure agent properties
- Configure traps
- Configure SNMP security properties
- Start or stop the SNMP service
- Configure event to trap translator

# Install the SNMP service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install the SNMP service**

1. Open Windows Components wizard.
2. In **Components**, click **Management and Monitoring Tools** (but do not select or clear its check box), and then click **Details**.
3. Select the **Simple Network Management Protocol** check box, and click **OK**.
4. Click **Next**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.
- SNMP starts automatically after installation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configure agent properties
Configure traps
Configure SNMP security properties
Start or stop the SNMP service

# Configure agent properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure agent properties**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Services and Applications/Services

3. In the details pane, click **SNMP Service**.
4. On the **Action** menu, click **Properties**.
5. On the **Agent** tab, in **Contact**, type the name of the user or administrator for this computer.
6. In **Location**, type the physical location of the computer or the contact.
7. Under **Service**, select the appropriate check boxes for this computer, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- If you change **Contact** or **Location** SNMP settings, your changes take effect within a few minutes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Install the SNMP service
Configure traps
Configure SNMP security properties
Start or stop the SNMP service

# Configure traps

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure traps**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Services and Applications/Services

3. In the details pane, click **SNMP Service**.
4. On the **Action** menu, click **Properties**.
5. On the **Traps** tab, under **Community name**, type the case-sensitive community name to which this computer will send trap messages, and then click **Add to list**.
6. In **Trap destinations**, click **Add**.
7. In **Host name, IP or IPX address**, type information for the host, and click **Add**.
8. Repeat steps 5 through 7 until you have added all the communities and trap destinations you want.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- If you change existing SNMP settings, your changes take effect immediately. The SNMP service does not need to be restarted for your settings to take effect.
- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Install the SNMP service
Configure agent properties
Configure SNMP security properties

Start or stop the SNMP service

# Configure SNMP security properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure SNMP security properties**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Services and Applications/Services

3. In the details pane, click **SNMP Service**.
4. On the **Action** menu, click **Properties**.
5. On the **Security** tab, select **Send authentication trap** if you want a trap message sent whenever authentication fails.
6. Under **Accepted community names**, click **Add**.
7. Under **Community Rights**, select a permission level for this host to process SNMP requests from the selected community.
8. In **Community Name**, type a case-sensitive community name, and then click **Add**.
9. Specify whether or not to accept SNMP packets from a host:

   - To accept SNMP requests from any host on the network, regardless of identity, click **Accept SNMP packets from any host**.
   - To limit acceptance of SNMP packets, click **Accept SNMP packets from these hosts**, click **Add**, type the appropriate host name, IP or IPX address, and then click **Add** again.

**Important**

- By default, SNMP does not respond to any community names presented.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can add additional community and host names as necessary.
- You can make changes to an entry by clicking the entry, and then clicking **Edit**. You can delete a selected entry by clicking **Remove**.

- If you change existing SNMP settings, your changes take effect immediately. The SNMP service does not need to be restarted for your settings to take effect.

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Install the SNMP service
Configure agent properties
Configure traps
Start or stop the SNMP service

# Start or stop the SNMP service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start or stop the SNMP service**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Services and Applications/Services

3. In the details pane, click **SNMP Service**.
4. On the **Action** menu, click **Start**, **Stop**, or **Restart**.

## Notes

- To perform this procedure, you must be a member of the Power Users group, or the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- If you add new extensions to SNMP, you do not need to restart the SNMP service for your changes to take effect.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Install the SNMP service
Configure agent properties
Configure traps
Configure SNMP security properties

# Configure event to trap translator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure event to trap translator**

1. Open Command Prompt.
2. At Command Prompt, type **evntcmd /?** for syntax and options for using the **evntcmd** command.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Trap properties
**Other Resources**
Evntcmd

# SNMP Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- SNMP Overview
- Understanding SNMP
- Using SNMP

# SNMP Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SNMP overview**


- SNMP defined
- The SNMP service

# SNMP defined

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SNMP defined**

The Simple Network Management Protocol (SNMP) is a network management standard widely used in TCP/IP networks.

SNMP provides a method of managing network hosts such as workstation or server computers, routers, bridges, and hubs from a centrally-located computer running network management software. SNMP performs management services by using a distributed architecture of management systems and agents.

Because network management is critical for both auditing and resource management, SNMP can be used to:

- **Configure remote devices.** Configuration information can be sent to each networked host from the management system.
- **Monitor network performance.** You can track the speed of processing and network throughput, and collect information about the success of data transmissions.
- **Detect network faults or inappropriate access.** You can configure trigger alarms on network devices when certain events occur. When an alarm is triggered, the device forwards an event message to the management system. Common types of alarms include a device being shut down and restarted, a link failure being detected on a router, and inappropriate access.
- **Audit network usage.** You can monitor both overall network usage to identify user or group access, and types of usage for network devices and services.

# The SNMP service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The SNMP service**
The simple network management protocol (SNMP) service supports computers running TCP/IP and IPX protocols. It is an optional service that can be installed after the TCP/IP protocol has been successfully configured.

The SNMP service provides an SNMP agent that allows remote, centralized management of computers running:

- Microsoft® Windows® XP
- Microsoft Windows 2000
- Microsoft Windows Server 2003 family

The SNMP agent also allows management of the following services:

- Windows XP or the Windows Server 2003 family and Microsoft Windows 2000-based WINS
- Windows XP or the Windows Server 2003 family and Windows 2000-based DHCP
- Windows XP or the Windows Server 2003 family and Windows 2000-based Internet Information Services
- Microsoft LAN Manager

To access the information that the SNMP agent service provides, you need at least one SNMP management system software application. The SNMP service supports but does not currently include SNMP management software. SNMP management software must be running on the host which acts as the management system.

**Note**

- The IPX/SPX protocol is not available on the Itanium-based versions of the Windows operating systems.

# Understanding SNMP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding SNMP**


- SNMP management systems and agents
- SNMP messages

# SNMP management systems and agents

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SNMP management systems and agents**

Using SNMP requires two components:

- **An SNMP management system.**

  The management system, also called *management console*, sends information and update requests to an SNMP agent. Any computer running SNMP management software is an SNMP management system. The management software application does not need to run on the same host as the SNMP agent.

  The SNMP management system requests information from a managed computer, called an SNMP agent, such as the amount of hard disk space available or the number of active sessions. If the management system has been granted write access to an agent, that management system can also initiate a change to the agent's configuration.

- **An SNMP agent.**

  The SNMP agent responds to management system requests for information. Any computer running SNMP agent software is an SNMP agent. The SNMP service, which is agent software, responds to information requests from one or multiple management systems. The SNMP service can be configured to determine which statistics are tracked and which management systems are authorized to request information.

  In general, agents do not originate messages, but only respond to them. A trap message is the only agent-initiated SNMP communication, and it enhances security. A trap is an alarm-triggering event, such as a system reboot or illegal access, on an agent.

  Management hosts and agents belong to an SNMP community, which is a collection of hosts grouped together for administrative purposes. Defining communities provides security by allowing only management systems and agents within the same community to communicate.

# SNMP messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SNMP messages**

When simple network management protocol (SNMP) management programs send requests to a network device, the agent software on that device receives the requests and retrieves information from MIBs. The agent then sends the requested information back to the initiating SNMP management program. To perform these tasks, the agent uses the following message types:

| SNMP Message | Description |
|---|---|
| Get | The basic SNMP request message. Sent by an SNMP management system, it requests information about a single MIB entry on an SNMP agent. For example, the amount of free drive space. |
| Get-next | An extended type of request message that can be used to browse the entire tree of management objects. When processing a Get-next request for a particular object, the agent returns the identity and value of the object which logically follows the object from the request. The Get-next request is useful for dynamic tables, such as an internal IP route table. |
| Set | If write access is permitted, this message can be used to send and assign an updated MIB value to the agent. |
| Getbulk | Requests that the data transferred by the host agent be as large as possible within given restraints of message size. This minimizes the number of protocol exchanges required to retrieve a large amount of management information. The maximum message size should not be larger than the path maximum transmission unit (MTU), the largest frame size allowed for a single frame on your network, or fragmentation can occur. |
| Trap | An unsolicited message sent by an SNMP agent to an SNMP management system when the agent detects that a certain type of event has occurred locally on the managed host. The SNMP management console that receives a trap message is known as a trap destination. For example, a trap message might be sent on a system restart event. |

Four of these message types are simple request-reply protocols where SNMP uses User Datagram Protocol (UDP). This means that a request from the management system might not arrive at the agent, and the agent's reply might not make it back to the management system. Because SNMP is a connectionless network protocol, SNMP messages are not guaranteed to arrive at their destinations.

You can use Internet Protocol security (IPSec) to secure traffic between SNMP management systems and agents. For more information, see Securing SNMP messages with IPSec.

# Using SNMP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using SNMP**


- Defining communities
- Agent properties
- Trap properties
- SNMP security properties
- Securing SNMP messages with IPSec
- SNMP service properties
- Managing SNMP from the command line

# Defining communities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Defining communities**

You can assign groups of hosts to Simple Network Management Protocol (SNMP) communities for limited security checking of agents and management systems or for administrative purposes. Communities are identified by community names that you assign. A host can belong to multiple communities at the same time, but an agent does not accept a request from a management system outside its list of acceptable community names.

For more information, see SNMP Best practices.

Define communities logically to take advantage of the basic authentication service provided by SNMP. In the following example, there are two communities: Public 1 and Public 2.




- Agent 1 can send traps to Manager 2 and respond to requests from Manager 2 because they are both members of the Public 2 community.
- Agents 2-4 can send traps to Manager 1 and respond to requests from Manager 1 because they are all members, by default, of the Public 1 community.

### Important

- Community names are sent across the network as plaintext. Because attackers can read plaintext with network analysis software, sending SNMP community names across the network represents a potential security risk. You can help protect SNMP messages by configuring Internet Protocol security (IPSec). For more information about configuring SNMP for IPSec, see Securing SNMP messages with IPSec.

### Notes

- There is no relationship between community names and domain or workgroup names. Community names represent a shared password for groups of network hosts, and should be selected and changed as you would change any password.
- Use community names primarily as an element for organization, not security.
- You should not create a community that is named Public and grant it read access. You should also specify hosts from which packets can be accepted, rather than clicking **Accept packets from any host**.

# Agent properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Agent properties**

The simple network management protocol (SNMP) agent provides the related management system with information on activities that occur at the Internet Protocol (IP) network layer.

The SNMP service sends MIB information in response to an SNMP request or in an SNMP trap message.

You can configure the following agent properties by using the **Agent** tab on the **SNMP Service Properties** dialog box:

| **Agent Service** | **Select if this computer:** |
|---|---|
| Physical | Manages physical devices, such as a hard disk partition. |
| Applications | Uses any applications that send data using the TCP/IP protocol suite. This service should always be enabled. |
| Datalink and Subnet | Manages a bridge. |
| Internet | Is an IP gateway (router). |
| End-to-end | Is an IP host. This service should always be enabled. |

You can also configure agent properties such as:

- The location of the contact person.
- The name of the person to contact, such as the network administrator.

For more information on how to configure SNMP agent properties, see Configure agent properties.

# Trap properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Trap properties**

The SNMP service, when configured for an agent, generates trap messages if any specific events occur. These messages are sent to a trap destination. For example, an agent can be configured to initiate an authentication trap if a request for information is sent by an unrecognized management system.

Trap destinations consist of the computer name or the IP or IPX address of the management system. The trap destination must be a network-enabled host that is running SNMP management software. Trap destinations can be configured by a user, but the events (such as a system reboot) that can generate a trap message are internally defined by the SNMP agent.

The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

You can configure trap destinations by using the **Traps** tab in **SNMP Service Properties**.

For more information on how to configure SNMP traps, see Configure traps.

# SNMP security properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SNMP security properties**

Simple network management protocol (SNMP) provides security through the use of authentication traps and community names. In addition, you can restrict SNMP communications for the agent, allowing it to communicate with only a specified list of SNMP management systems. Because the SNMP protocol provides minimal security, you should use SNMP only on trusted networks.

You can configure certain features of SNMP security on the **Security** tab in the **SNMP Service Properties** dialog box. These features provide a minimal level of security to your network, and you should use them in conjunction with other stronger methods of secure network management.

The following options can be configured to enable SNMP security:

- **Accepted community names**. The service has no default community name. By default, SNMP will not respond to any community names presented. You can add, delete, or change multiple community names. If an SNMP request is received from a community that is not on this list, it will generate an authentication trap.

  You can define community rights for a community name. A permission level can be selected in the **SNMP Service Configuration** dialog box, that determines how the SNMP agent processes requests from a selected community. For example, you can configure the permission level to block the SNMP agent from processing any requests from a specific community.

- **Accept SNMP packets from any host**. In this context, no SNMP packets are rejected on the basis of the name or address of the source host or the list of acceptable hosts.

- **Accept SNMP packets from these hosts**. In this context, the list of acceptable hosts means the acceptable SNMP management systems. This option is selected by default with *localhost* as the only hostname. Only SNMP packets received from *localhost* are accepted. SNMP messages from all other hosts are rejected and an authentication trap is sent. This option provides a higher level of security than using a community name, which might contain many hosts.

- **Send authentication trap**. Authentication is the process of verifying that a host name or address is valid. When the SNMP agent receives a request that does not contain the correct community name or is not sent from a member of the acceptable host list, the agent sends an authentication trap message to one or more trap destinations (management systems), indicating the failure of authentication. This option is selected by default.

For more information on how to configure SNMP security, see Configure SNMP security properties.

**Note**

- The default Security tab settings are no communities listed under **Accepted community names** and one host name, *localhost*, under **Accept SNMP packets from these hosts**. The term *localhost* refers to the loopback interface on the local computer. You must add a community name of your choice to **Accepted community names**. For security reasons, naming a community Public is not recommended. You should monitor and update the settings on an ongoing basis to ensure timely detection of any unauthorized access.

# Securing SNMP messages with IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing SNMP messages with IPSec**

By configuring IPSec policies on all SNMP agents and managers, you can help prevent malicious users and attackers from intercepting SNMP messages. If you cannot configure IPSec policies on all SNMP hosts but you want to ensure that all hosts can communicate with each other, you must configure IPSec policies that allow plaintext communication. However, allowing plaintext communication is not recommended.

IPSec does not automatically encrypt the SNMP traffic. You must create filter specifications in the appropriate IP filter list for traffic between the SNMP managers and agents.

To make SNMP messages more secure, you must add two sets of filter specifications to a new or existing IPSec policy on the SNMP-enabled host. The first set of filter specifications regulates typical SNMP traffic or SNMP messages between the SNMP managers and the SNMP agents. This set of filter specifications typically consists of one filter specification for inbound traffic and one for outbound traffic. You must create these specifications in the **IP Filter Properties** dialog box. For more information about creating filter specifications, see Filter list, Set an active IP filter list for a rule, and Add, edit, or remove IP filter lists.

On the **Addresses** tab:

- In **Source address**, click **A specific IP address** and type the address of the SNMP manager.
- In **Destination address**, click **My IP Address**, which refers to the IP address of the SNMP agent to which the policy is assigned.
- Select the **Mirrored** check box to automatically create the filter specification for outbound traffic.

On the **Protocol** tab:

- Under **Select a protocol type**, click **TCP** or **UDP**. If you require both these protocols, create an additional filter specification.
- Click **From any port**, click **To this port**, and type **161**.

The second set of filter specifications regulates trap messages, and it consists of one filter specification for inbound traffic and one for outbound traffic:

On the **Addresses** tab:

- Under **Source address**, click **A specific IP address** and type the IP address of the SNMP agent.
- Under **Destination address**, click **My IP Address**, which refers to the IP address of the SNMP manager to which the policy is assigned.
- Select the **Mirrored** check box to automatically create the outbound filter specification.

On the **Protocol** tab:

- Under **Select a protocol type**, click **TCP** or **UDP**. If you require both these protocols, create an additional filter specification.
- Click **From any port**, click **To this port**, and type **162**.

For more information on SNMP messages, see SNMP messages.

To send secure SNMP messages, you must configure IPSec policies on SNMP managers as well as SNMP agents. The Windows Server 2003 family supports but does not include SNMP management software.

# SNMP service properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SNMP service properties**

You can use the **General**, **Log On**, and **Recovery** tabs in **SNMP Service Properties** to configure how the SNMP service starts, logs on to the system, and recovers from an abnormal program termination of the service or operating system.

For information on changing these properties, see Services Snap-in.

# Managing SNMP from the command line

Published: January 21, 2005

Updated: November 28, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing SNMP from the command line**

You can use the **evntcmd** command to configure SNMP traps based on events recorded in system logs. You can also use this command to specify where trap messages are sent within an SNMP community. For information on how to use the **evntcmd** command to configure traps and trap destinations, see evntcmd at http://www.microsoft.com/resources/documentation/windows/xp/all/proddocs/en-us/evntcmd.mspx. For information on identifying the events from which you can configure traps, see Event Viewer.

You can use the **evntwin** command to start Event to Trap Translator, which is a graphical utility you can use to configure the translation of events to traps. However, you cannot use Event to Trap Translator to configure trap destinations. For information on configuring trap destinations using the SNMP console, see Configure traps.

You must install SNMP before either the **evntcmd** command or the **evntwin** command will work.

For information about other command-line utilities, see Command-line reference. For more information about manageability, see Management Strategies and Tools.

# Troubleshooting SNMP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- My SNMP agent is not functioning.
- My SNMP management system keeps timing out while querying the WINS server.

**My SNMP agent is not functioning.**
**Cause:** An unknown network or SNMP problem.

**Solution:** SNMP error handling is integrated with Event Viewer. Use Event Viewer if you suspect there is a problem with the SNMP service. You can configure a **View** filter to display only SNMP messages.

**See also:** Event Viewer.

**My SNMP management system keeps timing out while querying the WINS server.**
**Cause:** The SNMP time-out period is too short or the WINS server is off-line or unreachable.

**Solution:** Increase the SNMP time-out period on the SNMP management system. If that does not work contact your system administrator to check the WINS server.

# Device Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Device Manager**
You can use Device Manager to update the drivers (or software) for hardware devices, modify hardware settings, and troubleshoot problems.

- For help with specific tasks, see Device Manager How To....
- For general background information, see Device Manager Concepts.

**Caution**

- Incorrectly altering your hardware configuration can damage your system. Be sure to read the appropriate instructions before making changes to your hardware configuration.

# Device Manager How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- View Device Information
- Configure Device Properties and Settings
- Uninstall and Reinstall Devices
- Update and View Device Drivers

# View Device Information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**View device information**

- View devices and resource settings
- View hidden devices
- View the status of a device
- View power allocations for USB hubs
- View bandwidth allocations for a USB host controller
- Print information about a specific device

# View devices and resource settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view devices and resource settings**

1. Open Device Manager.
2. On the **View** menu, click one of the options listed in the following table:

| Option | Description |
|---|---|
| **Devices by type** | Displays devices by the type of device installed, such as monitor or mouse. This is the default view. After you double-click the device type, the connection name is listed below the type. |
| **Devices by connection** | Displays devices by how they are connected in your computer. Each device is listed under the hardware to which it is connected. For example, if a small computer system interface (SCSI) card is listed, the devices attached to the SCSI card are listed beneath it. |
| **Resources by type** | Displays all allocated resources by the type of device using these resources. The resources are as follows:<br>• Direct memory access (DMA) channels<br>• Input/output ports (I/O ports)<br>• Interrupt request (IRQ<br>• Memory addresses |
| **Resources by connection** | Displays all allocated resources by connection type. The resources are: DMA channels, I/O ports, IRQ, and memory addresses. |

**Notes**

- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- For information on how to view hidden devices, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View information about a device driver
View the status of a device
Change resource settings for a device
View hidden devices
Working with MMC console files
Using the runas command with MMC

# View hidden devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view hidden devices**

1. Open Device Manager.
2. On the **View** menu, click **Show hidden devices**.

## Notes

- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- If you no longer want to view hidden devices, click **Show hidden devices** to remove the check mark.
- Hidden devices include non-Plug and Play devices (devices with earlier Windows device drivers).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View devices and resource settings
View information about a device driver
View the status of a device
Change resource settings for a device
Working with MMC console files
Using the runas command with MMC

# View the status of a device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the status of a device**

1. Open Device Manager.
2. Double-click the type of device you want to view.
3. Right-click the specific device you want, and then click **Properties**.
4. On the **General** tab, under **Device status** is a description of the device status.

## Notes

- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- If the device is experiencing a problem, the type of problem is displayed. You might also see a problem code and number and a suggested solution. If you call a support line, this problem code and number can be useful for determining and diagnosing the problem.
- For information on how to resolve hardware device problems, click **Troubleshoot** to start the Hardware Troubleshooter.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View information about a device driver
View devices and resource settings
View hidden devices
Working with MMC console files

# View power allocations for USB hubs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view power allocations for USB hubs**

1. Open Device Manager.
2. Double-click **Universal serial bus controller**.
3. Right-click **USB Root Hub**, and then click **Properties**.
4. On the **Power** tab, view the power required by each device in the **Attached devices** list.

## Notes

- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- Hubs for universal serial bus (USB) devices are self-powered or bus-powered. A self-powered hub (a hub that is plugged into an electrical outlet) provides maximum power to the device. A bus-powered hub (a hub that is plugged into another USB port) provides minimum power. Devices requiring a lot of power, such as cameras, should be plugged into self-powered hubs.
- You can check the power allocations for a USB hub to detect situations in which too many devices are using the hub.
- **Universal serial bus controller** only appears if you have a USB port on your computer.
- The **Power** tab only appears for USB hubs.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View bandwidth allocations for a USB host controller
Working with MMC console files

# View bandwidth allocations for a USB host controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view bandwidth allocations for a USB host controller**

1. Open Device Manager.
2. Double-click **Universal serial bus controller**.
3. Right-click a **USB Universal Host Controller** for your system, and then click **Properties**.
4. On the **Advanced** tab, in the **Bandwidth-consuming devices** list, view the bandwidth used by each device.

## Notes

- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- You can view bandwidth only for a universal serial bus (USB) controller. Some USB devices (for example, modems) might not appear because they do not report bandwidth requirements to the operating system.
- Each USB controller has a fixed amount of bandwidth that all attached devices must share. If you suspect a bandwidth shortage for a USB controller, you can confirm this by viewing the bandwidth allocations currently being used by each device.
- **Universal serial bus controller** appears only if you have a USB port on your computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View power allocations for USB hubs
Working with MMC console files

# Print information about a specific device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print information about a specific device**

1. Open Device Manager.
2. Double-click the type of device you want to print information about, and then click the specific device to select it.
3. On the **Action** menu, click **Print**.
4. In the **Print** dialog box, under **Report type**, click **Selected class or device**.

   You can also print other types of reports from this dialog box. The following table describes the types of reports you can print:

| Option | Description |
|---|---|
| **System summary** | Prints a summary of the devices that are installed on your computer or on the remote computer you are managing. The summary includes the product version number, bus type, registered owner, and resource information such as interrupt request (IRQ) usage or input/output (I/O) port usage. |
| **Selected class or device** | Prints a report for the device or the hardware type you have selected. The report includes device name, device driver location, version and manufacturer information, and resources allocated to the device, such as I/O ports and memory addresses. |
| **All devices and system summary** | Prints a report for each device or hardware type that is installed on your computer and a summary of your system. If you have many devices installed on your system, this could be a large document. |

**Notes**

- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- You can print to a .prn file by selecting the **Print to file** check box.
- **Selected class or device** is the default setting for the type of report to be printed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
View devices and resource settings
Working with MMC console files

# Configure Device Properties and Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure device properties and settings**

- Enable a device
- Disable a device
- Change resource settings for a device
- Change network adapter settings
- Change port settings
- Change PS/2 mouse settings
- Change IDE controller settings
- Change region settings for DVD drives
- Change ISDN adapter settings
- Change power management settings for a device
- Enable or disable digital CD audio
- Configure multimedia devices
- Upgrade from a uniprocessor to a multiprocessor computer

# Enable a device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable a device**

1. Open Device Manager.
2. Double-click the type of device that you want to enable.
3. Right-click the specific device you want, and then click **Enable**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- If you are prompted to restart the computer, the device will not be enabled until the computer has been restarted.
- **Enable** will only be listed if the device is disabled.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Uninstalling devices
Disable a device
View devices and resource settings
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Disable a device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable a device**

1. Open Device Manager.
2. Double-click the type of device that you want to disable.
3. Right-click the specific device you want, and then click **Disable**.

### Important

- If you disable a device that is required to start the system, such as your hard disk drive, these devices will not work, and your system might not start.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- When you disable a device, the physical device stays connected to your computer, but the device driver is disabled. The drivers are available again when you enable the device. It can be useful to disable devices if you want to have more than one hardware configuration for your computer, or if you have a portable computer that you use at a docking station.
- If you are prompted to restart the computer, the device will not be disabled and will continue to function until the computer is restarted.
- After you disable a device and restart your computer (if necessary), the resources allocated to the device are free and can be allocated to another device.
- Some devices cannot be disabled, such as devices for disk drives and processors.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Uninstalling devices
Enable a device
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC
Start and stop non-Plug and Play drivers

# Change resource settings for a device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change resource settings for a device**

1. Open Device Manager.
2. Double-click the type of device you want to change.
3. Double-click the specific device you want to change. If the device has resources settings that you can change, the **Resources** tab will appear in the dialog box.
4. Click the **Resources** tab, and then clear the **Use automatic settings** check box. This check box might be unavailable if you cannot assign a specific resource to the device.
5. In **Setting based on**, click the hardware configuration you want to change.
6. In **Resource settings**, under **Resource type**, click the resource type you want to change. Your choices are as follows:
   - Direct memory access (DMA)
   - Interrupt request (IRQ)
   - Input/output (I/O) port
   - Memory address
7. Click **Change Setting**, and then type a new value for the resource type.
8. Repeat as necessary.

**Caution**

- Changing resource settings improperly can disable your hardware and cause your computer to malfunction or become inoperable. Resource settings should be changed only by users who have expert knowledge of computer hardware and hardware configuration settings.

**Important**

- You must be certain that the new values you apply do not conflict with any other devices. If you have any device conflicts, they appear under **Conflicting device list**.

  **Use automatic settings** is unavailable for devices that have no other settings to configure. It is also unavailable on devices that are controlled by Plug and Play resources and should not be changed by the user.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer**

**Management**, and then click **Device Manager**.

- You should need to assign resources manually only for some older Industry Standard Architecture (ISA) and Plug and Play basic input/output system (BIOS) devices. It is not usually necessary for you to change resource settings because resources are automatically allocated by the system when the hardware is installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View devices and resource settings
Working with MMC console files
Using the runas command with MMC
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as

# Change network adapter settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change network adapter settings**

1. Open Device Manager.
2. Double-click **Network adapters**.
3. Right-click the network adapter for which you want to change settings, and then click **Properties**.
4. On the **Advanced** tab, make the appropriate changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- Device properties for network adapters are device dependent. Read the device manufacturer's instruction manual to determine which settings you must change.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Change port settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change port settings**

1. Open Device Manager.
2. Double-click **Ports (COM & LPT)**.
3. Right-click the port for which you want to change settings, and then click **Properties**.
4. On the **Port Settings** tab, make the appropriate changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- Most hardware devices that use ports set the port speed automatically. The defaults for modem port speed are specified in **Phone and Modem Options** in **Control Panel**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
Configure an installed modem
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Change PS/2 mouse settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change PS/2 mouse settings**

1. Open Device Manager.
2. Double-click **Mice and other pointing devices**.
3. Right-click **PS/2 Compatible Mouse**, and then click **Properties**.
4. On the **Advanced Settings** tab, make the appropriate changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- Some of the advanced settings for a PS/2 compatible mice are device dependent. Read the device manufacturer's instruction manual to determine which settings you can change.
- For information on changing other mouse settings, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
View devices and resource settings
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Change IDE controller settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change IDE controller settings**

1. Open Device Manager.
2. Double-click **IDE ATA/ATAPI controllers**, right-click the Integrated Device Electronics (IDE) device for which you want to change advanced settings, and then click **Properties**.
3. On the **Advanced Settings** tab, make the appropriate changes, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- The **Advanced Settings** tab might not appear for all IDE controller devices.
- Some of the advanced settings for IDE controllers are device dependent. Read the device manufacturer's instruction manual to determine which settings you can change.
- You will not be able to change the setting in **Device Type** if the system has already detected an IDE device on an IDE channel. This prevents you from selecting **None**, which would disable the IDE device. If you want to disable the device, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
Disable a device
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Change region settings for DVD drives

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change region settings for DVD drives**

1. Open Device Manager.
2. Double-click **DVD/CD-ROM drives**, right-click the DVD drive for which you want to change region settings, and then click **Properties**.
3. On the **DVD Region** tab, make the appropriate changes.

### Important

- Because many software decoders require Direct Memory Access (DMA) in order to function, you should check the settings of your Integrated Device Electronics (IDE) ATA/ATAPI controllers to ensure that **DMA if available** is selected in **Transfer Mode**. If applicable, you should also check this setting for the Secondary IDE Channel.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- The **DVD Region** tab might not be available for all DVD drives.
- To configure your DVD drive to play discs from other regions, select a geographic area from the list of regions. Typically, you would select the region in which you reside. It is not usually necessary to change the setting again.
- The region settings can be changed only a limited number of times. When **Changes remaining** equals 1, and you select a new region, an alert will be displayed warning you that after this, you will only be able to play DVDs from that region. Be aware that you might not be able to change the region setting for some DVD hardware and software decoders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
View devices and resource settings
Change IDE controller settings
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Change ISDN adapter settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change ISDN adapter settings**

1. Open Device Manager.
2. Double-click **Network adapters**, right-click the Integrated Services Digital Network (ISDN) adapter card for which you want to change settings, and then click **Properties**.
3. On the **ISDN** tab, in the **Select the switch type or D-channel protocol** list, click a switch type or a D-channel protocol, and then click **Configure**.
4. Make the appropriate changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- Your telephone company or hardware manufacturer should provide the information you need to configure ISDN adapter settings.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
View devices and resource settings
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Change power management settings for a device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change power management settings for a device**

1. Open Device Manager.
2. Double-click the type of device for which you want to change power management settings.
3. Right-click the specific device you want, and then click **Properties**.
4. On the **Power Management** tab, change the power management settings.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- You can use the power management settings to specify that the operating system return from a low-power state to service an event. One example of the effect of doing so is that an incoming modem call will alert the system to service the call even if the system is in hibernation. You can also specify whether the operating system should turn off a particular device if it is not doing something, such as servicing a call.
- The **Power Management** tab appears only for device drivers that support the power management standards, such as network adapters.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
View devices and resource settings
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Enable or disable digital CD audio

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable digital CD audio**

1. Open Device Manager.
2. Double-click **DVD/CD-ROM drives**, right-click the CD-ROM drive, and then click **Properties**.
3. On the **Properties** tab, select or clear the **Enable digital CD audio for this CD-ROM device** check box.

   This check box enables some CD-ROM drives to read the audio data from an audio CD digitally and send it to either a sound card or a pair of universal serial bus (USB) speakers.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- It might be necessary to restart your computer after changing this setting.
- Enabling this feature can improve audio quality because Windows reads the CD directly, bypassing the CD-ROM drive's digital-to-analog converters. Directly accessing the CD improves quality because it eliminates extraneous electrical interference that can leak through the shielding on the cable that connects the CD-ROM drive to the sound card.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
View devices and resource settings
Working with MMC console files
Using the runas command with MMC

# Configure multimedia devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure multimedia devices**

1. Open Device Manager.
2. Double-click **Sound, video and game controllers**.
3. Right-click the device that you want to configure, and then click **Properties**.
4. On the **Properties** tab, locate the specific device, and then double-click it.
5. Change the configurations for the device.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- The **Properties** button is available only if you have drivers other than Windows Driver Model (WDM) drivers on your computer. If The **Properties** button appears dimmed, use the **Resources** tab to configure the device.
- Multimedia devices can include audio or video codecs, game ports, and MIDI devices.
- You might want to reconfigure multimedia settings to adjust compression settings and allow a device driver to be used by a program

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Update or change a device driver
View devices and resource settings
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Upgrade from a uniprocessor to a multiprocessor computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To upgrade from a uniprocessor to a multiprocessor computer**

1. Open Device Manager.
2. Double-click **Computer**.
3. Right-click the computer model, and then click **Properties**.
4. On the **Driver** tab, click **Update Driver**. This starts the Hardware Update Wizard.
5. Click **Install from a list or specific location**, and then click **Next**.
6. From the **Please choose your search and installation options** page, click **Don't search. I will choose the driver to install**, and then click **Next**.
7. From the **Select a Device Driver** page, click the appropriate computer type from the **Models** list.
8. Click **Next**, click **Next** again, and then click **Finish**.

### Caution

- Upgrading from a uniprocessor to a multiprocessor computer can improperly disable your hardware and cause your computer to malfunction or become inoperable. Upgrading should be done only by users who have expert knowledge of computer hardware and hardware configuration settings.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View information about a device driver
Update or change a device driver
Device Manager overview
Using the runas command with MMC

# Uninstall and Reinstall Devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Uninstall and reinstall devices**

- Uninstall a device
- Reinstall a Plug and Play device
- Reinstall a non-Plug and Play device
- Roll back to the previous version of a driver
- Undock a portable computer
- Unplug or eject devices

# Uninstall a device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To uninstall a device**

1. Open Device Manager.
2. Double-click the type of device that you want to uninstall.
3. Right-click the specific device you want to uninstall, and then click **Uninstall**.
4. In the **Confirm Device Removal** dialog box, click **OK**.
5. When you are done uninstalling the device, remove it from the computer.

## Notes

- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- If you are prompted to restart the computer, the device will not be removed and will continue to function until the computer is restarted.
- You do not usually need to use Device Manager to uninstall a Plug and Play device. Simply disconnect the Plug and Play device from your computer. You might have to restart your computer, however, depending on the type of device. Check the device manufacturer's instructions for more information.
- Uninstalling a device does not remove the device drivers from your computer's hard disk.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Uninstalling devices
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as

# Reinstall a Plug and Play device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reinstall a Plug and Play device**

1. Open Device Manager.
2. Double-click the type of device that you want to reinstall.
3. Right-click the specific device you want, and then click **Uninstall**.
4. In the **Confirm Device Removal** dialog box, click **OK**. If you are prompted to restart the computer, the device will be reinstalled when the computer starts again. If you are not prompted to restart the computer, continue to the next step.
5. To reinstall the device, on the **Action** menu, click **Scan for hardware changes**.
6. Follow the instructions on the screen.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- You should reinstall a Plug and Play device only if it is working improperly or it has stopped working altogether. Before you reinstall a device, you should try to restart your computer and check the device to determine whether it is functioning properly. If it is not, try reinstalling the device.
- You do not usually need to reinstall or uninstall a Plug and Play device. With Plug and Play devices, you can simply remove the device from your computer. To remove a device properly, you should follow the device manufacturer's installation/removal instructions.
- **Scan for hardware changes** will not reinstall a Plug and Play device if it detects the device and its driver. If the device has not been uninstalled, **Scan for hardware changes** will not detect it as a change. You must uninstall a Plug and Play device before scanning for hardware changes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Uninstalling devices
Reinstall a non-Plug and Play device
Device Manager overview
View devices and resource settings
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas

Using Run as
Using the runas command with MMC

# Reinstall a non-Plug and Play device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reinstall a non-Plug and Play device**

1. In **Control Panel**, open Add Hardware.
2. In the **Add Hardware Wizard**, click **Next**, and then follow the instructions on your screen.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Reinstall a Plug and Play device
Device Manager overview
View devices and resource settings
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC

# Roll back to the previous version of a driver

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To roll back to the previous version of a driver**

1. Open Device Manager.
2. Right-click the device for which you want to reinstall the previous version of the driver, then click **Properties**.
3. Click the **Drivers** tab.
4. Click **Roll Back Driver**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Uninstalling devices
Enable a device
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Using the runas command with MMC
Start and stop non-Plug and Play drivers

# Undock a portable computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To undock a portable computer**

1. Click **Start**, and then click **Eject PC**. A message appears saying that it is safe to undock your computer.

2. Undock your portable computer, or, if you have a motorized docking station, your portable computer will undock automatically.

### Warning

- Removing a portable computer without using **Eject PC** is not supported and can result in data loss or even hardware damage. To avoid loss or damage, use **Eject PC** to undock your computer, or shut down the system and undock the portable computer while your system is off. Undocking a portable computer while it is in either standby mode or hibernate mode is also not supported as this is likely to result in data loss and system instability.

### Notes

- **Eject PC** appears on the **Start** menu only if a computer is Advanced Configuration and Power Interface (ACPI)-enabled. If **Eject PC** does not appear on the **Start** menu, your computer is not ACPI-enabled, might not support undocking, or you might lack the security permissions required to undock your portable computer. See your administrator for these permissions.

- Some computers have an eject button on the docking station. Pressing the eject button button instead of following this procedure also undocks the laptop. Read the manufacturer's instructions for more information.

- To work on a file or folder that is located on the network, ensure that you make the item available on your undocked laptop. For more information on using offline folders, see **Offline Files overview** in Related Topics.

- For optimal laptop performance, you should change your power scheme using **Power Options** in **Control Panel**.

- Programs located on a network will not be available on your laptop when you undock.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Uninstalling devices
Unplug or eject devices
Synchronizing offline items when you log on to or log off from a computer
Make a file or folder available offline
Scheduling offline items to synchronize
Synchronizing offline items
Why you should not run your computer as an administrator
Run a program with administrative credentials
Offline Files overview

# Unplug or eject devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To unplug or eject devices**

1. In the notification area next to your task bar, double-click the **Safe Removal** icon. Safe Removal displays a list of Plug and Play devices that support safe removal and that are currently attached to the system.

   If you do not see the **Safe Removal** icon, your device does not support safe removal, and you cannot unplug or eject your device using Safe Removal.

2. In **Safely Remove Hardware**, in the list of devices, select the device that you want to unplug or eject, and then click **Stop**. This tells the system that you will be unplugging or ejecting the device.

3. In **Stop a Hardware device**, click **OK**.

   A notification appears stating that it is now safe to unplug or eject the device.

### Caution

- Unplugging or ejecting a device that supports safe removal without first using the Safe Removal application to warn the system can cause data to be lost or your system to become unstable. For example, if a device is unplugged during a data transfer, data loss is likely. If you use Safe Removal, however, you can warn the system before you unplug or eject a device, preventing possible loss of data.

### Note

- For removable storage devices that can safely be removed while the system is on, the system disables write caching by default. It does this so the devices can be removed without loss of data. When write caching is disabled, however, the system will likely experience slower performance, so be sure to enable write caching when you have finished removing the device.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Undock a portable computer

# Update and View Device Drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Update and view device drivers**

- Update or change a device driver
- View information about a device driver
- Start and stop non-Plug and Play drivers
- Change the startup type for non-Plug and Play drivers

# Update or change a device driver

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To update or change a device driver**

1. Open Device Manager.
2. Double-click the type of device you want to update or change.
3. Right-click the specific device driver you want to update or change, and then click **Update Driver** to open the Hardware Update Wizard. Follow the instructions in the wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- When you are upgrading a driver or installing new hardware, and your computer has access to the Internet, your computer will automatically check Windows Update for an appropriate device driver. For more information on Windows Update, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View information about a device driver
Device Manager overview
Working with MMC console files
Run a program with administrative credentials
Why you should not run your computer as an administrator
Runas
Using Run as
Start and stop non-Plug and Play drivers
Using the runas command with MMC

# View information about a device driver

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view information about a device driver**

1. Open Device Manager to view information about a device driver.
2. Double-click the type of device you want to view.
3. Right-click the specific device you want, and then click **Properties**.
4. On the **Driver** tab, click **Driver Details**.

## Notes

- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- This information can also help you determine the file version of the device driver.
- An icon appears next to device drivers that are digitally signed. For more information on digitally signed drivers, see "Driver Signing for Windows" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Device Manager overview
Update or change a device driver
View devices and resource settings
View hidden devices
Working with MMC console files
Start and stop non-Plug and Play drivers
Driver signing for Windows

# Start and stop non-Plug and Play drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start and stop non-Plug and Play drivers**

1. Open Device Manager.
2. Click the **View** menu, and then click **Show hidden devices**.
3. In the list of devices, double-click **Non-Plug and Play Drivers**, right-click a device, and then click **Properties**.
4. On the **Driver** tab, click **Start** or **Stop** to start or stop the device driver, and then click **OK**.

### Caution

- Starting or stopping a non-system device incorrectly can cause it to stop working.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- If the **Start** button on the **Driver** tab is unavailable, the driver is already loaded.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Update or change a device driver
Change the startup type for non-Plug and Play drivers
View devices and resource settings
View information about a device driver
Using the runas command with MMC

# Change the startup type for non-Plug and Play drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the startup type for non-Plug and Play drivers**

1. Open Device Manager.
2. Click the **View** menu, and then click **Show hidden devices**.
3. Double-click **Non-Plug and Play Drivers**, right-click a device, and then click **Properties**.
4. On the **Driver** tab, in **Startup** select a **Type**, and then click **OK**.

### Caution

- Improperly changing the startup type of a boot device or system device can cause the system to become unstable or to no longer boot, or cause the device to stop working altogether. You should fully understand the instructions before changing this setting.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.
- The **Startup Type** defaults to the type required by the driver when the driver is first installed. You can use the options described in the following table to change the default. Use caution when changing the startup type; be sure that you fully understand the instructions before you change this setting.

| Type | Description |
|---|---|
| **Automatic** | Starts every time the system starts, after the boot and system devices start. Click this option for devices that are not critical to basic system operation. |
| **Boot** | Starts every time the system starts, before any other devices start. Click this option for critical devices that are essential to system operation. |
| **Demand** | Starts when the device is detected or it is needed for a specific event. Click this option for devices that are not critical to basic system operation or required for device detection. |
| **System** | Starts every time the system starts, after the boot devices start. Click this option for critical devices that are essential to system operation. |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Update or change a device driver

View devices and resource settings
View information about a device driver
Using the runas command with MMC

# Device Manager Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides background about Device Manager. It provides information on th following topics:

- Device Manager overview
- Configuring devices
- Uninstalling devices
- Hardware profiles overview

# Device Manager overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Device Manager overview**
Device Manager provides a graphical view of the hardware that is installed on your computer. You can use Device Manager to update the drivers (or software) for hardware devices, modify hardware settings, and troubleshoot problems.

**Warning**

- Incorrectly altering your hardware configuration can damage your system. Be sure to read the appropriate instructions before making changes to your hardware configuration.

You can use Device Manager to do the following:

- Determine whether the hardware on your computer is working properly.
- Change hardware configuration settings.
- Identify the device drivers that are loaded for each device and obtain information about each device driver.
- Change advanced settings and properties for devices.
- Install updated device drivers.
- Disable, enable, and uninstall devices.
- Roll back to the previous version of a driver.
- Print a summary of the devices that are installed on your computer.

Device Manager is usually used to check hardware status and update device drivers on the computer. Advanced users with thorough understanding of computer hardware might also use Device Manager's diagnostic features to resolve device conflicts and change resource settings, but this should be done with extreme caution.

**Caution**

- Changing resource settings improperly can disable your hardware and cause your computer to malfunction or become inoperable. Only users who have expert knowledge of computer hardware and hardware configurations should change resource settings.

You will not usually use Device Manager to change resource settings because resources are allocated automatically by the system during hardware setup.

You must be a member of the Administrators group to manage device drivers.

You can use Device Manager to manage devices only on a local computer. On a remote computer, Device Manager will work only in read-only mode.

**Notes**

- If you cannot access Device Manager from the Computer Management extension snap-ins on a remote computer, ensure that the Remote Registry service is started on the remote computer. For more information, see Start, stop, pause, resume, or restart a service. You must have appropriate permissions on the remote computer to start the service.
- You might also receive this error message if the remote computer is running Microsoft® Windows 95. Computer Management does not support remote access to computers that are running Windows 95.

For more information, see the "Device Management" chapter that is provided on the Microsoft Resource Kits Web site.

# Configuring devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring devices**
When you install a Plug and Play device, Windows automatically configures the device so it will work properly with the other devices that are installed on your computer. As part of that configuration process, Windows assigns a unique set of system resources to the device you are installing. These resources can include one or more of the following:

- Interrupt request (IRQ) line numbers.
- Direct memory access (DMA) channels.
- Input/output (I/O) port addresses.
- Memory address ranges.

Each resource that is assigned to your device must be unique or the device will not function properly. For Plug and Play devices, Windows automatically ensures that these resources are configured properly.

Occasionally, two devices require the same resources, resulting in a device conflict. If this occurs, you can manually change the resource settings to be sure each setting is unique. However some resources, such as interrupts on PCI devices, can be shared, depending on the drivers and computer.

When you install a non-Plug and Play device, the resource settings for the device are not automatically configured. Depending on the type of device you are installing, you may have to manually configure these settings. Instructions for doing so should be supplied in the manual that came with your device.

Generally, you should not change resource settings manually, because when you do, the settings become fixed, and Windows will then have less flexibility when allocating resources to other devices. If too many resources become fixed, Windows may not be able to install new Plug and Play devices.

You can manually configure devices using Device Manager.

**Warning**

- Changing resource settings improperly can disable your hardware and cause your computer to malfunction or become inoperable. Resource settings should only be changed if you are certain the new settings do not conflict with other hardware, or if a hardware manufacturer has provided you with specific resource settings for a device.

**See Also**
**Concepts**
Device Manager overview
Hardware overview

# Uninstalling devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Uninstalling devices**

Although you can usually uninstall a plug and play device by disconnecting or removing the device, some devices might require that you turn off the computer first. To ensure that you do this properly, consult the device manufacturer's instructions for device installation and removal.

Uninstalling a non-Plug and Play device typically involves two steps:

- Using Device Manager to uninstall the device.
- Removing the device from your computer.

You can use Device Manager to notify the system that you want to uninstall a non-Plug and Play device. After you notify the system that you are uninstalling a device, you must physically disconnect or remove the device from your computer. For example, if the device is connected to a port on the outside of your computer, you shut down your computer, disconnect the device from the port, and then unplug the power cord for the device.

You need not uninstall a device that you want to disable but that you want to remain attached to the computer--a modem, for example. Instead of uninstalling it, you can simply disable the Plug and Play device. When you disable such a device, it stays physically connected to your computer, but Windows updates the system registry so that the device drivers for the device you have disabled are no longer loaded when you start your computer. The drivers are available again, however, when you enable the device. This is useful if you need to switch between two hardware devices, such as a networking card and a modem, or if you need to troubleshoot a hardware problem.

For instructions on how to uninstall, reinstall, and roll back devices, see Uninstall and Reinstall Devices.

For background information on hardware and instructions for installing devices, see Hardware profiles overview and Configuring devices.

# Hardware profiles overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware profiles overview**

A hardware profile is a set of instructions that tells Windows which devices to start when you start your computer or which settings to use for each device. When you first install Windows, a hardware profile called Profile 1 (for laptops, the profiles are Docked Profile or Undocked Profile) is created. By default, every device that is installed on your computer at the time you install Windows is enabled in the Profile 1 hardware profile.

Hardware profiles are especially useful if you have a portable computer. Most portable computers are used in a variety of locations, and hardware profiles will let you change which devices your computer uses when you move it from location to location. For example, you may have one profile named Docking Station Configuration for using your portable computer at a docking station with hardware components such as a CD-ROM drive and a network adapter. And you may have a second profile named Undocked Configuration for using your portable computer in a hotel or on an airplane, when you are not using a network adapter or a CD-ROM but you are using a modem and a portable printer.

If there is more than one hardware profile on your computer, you can designate a default profile that will be used every time you start your computer. You can also have Windows ask you which profile to use every time you start your computer. Once you create a hardware profile, you can use Device Manager to disable and enable devices that are in the profile. When you disable a device in a hardware profile, the device drivers for the device are not loaded when you start your computer.

**See Also**

**Concepts**

View properties of, copy, rename, or delete a hardware profile
Create a hardware profile
Device Manager overview
Disable a device
Enable a device
Hardware overview
System Properties overview

**Other Resources**

Installing devices overview

# Windows Management Instrumentation Control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Management Instrumentation Control**

- For help on using WMI Control, see Windows Management Instrumentation Control How To....
- To understand more about WMI Control, see Windows Management Instrumentation Control Concepts.

# Scheduling and Automation Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scheduling and automation tools**
This section covers:

- Task Scheduler
- Logon Scripts
- Windows Script Host

# Task Scheduler

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Task Scheduler**

You can use Task Scheduler to schedule commands, programs, or scripts to run at specific times.

- For tips about using Task Scheduler, see Task Scheduler Best practices.
- For help with specific tasks, see Task Scheduler How To....
- For general background information, see Task Scheduler Concepts.
- For problem-solving instructions, see Troubleshooting Task Scheduler.

# Task Scheduler Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Set appropriate permissions on the file your scheduled task will run.**

- To avoid a malicious user from corrupting a program set to run by a scheduled task, restrict access for the program file to only the specific users who need access.

- For more information about setting permissions on a file, see Set, view, change, or remove permissions on files and folders.

**Make sure you have correct password information for your users.**

- Before creating a scheduled task, check that the account and password information you have for the user that you will enter in the **Add Scheduled Task** wizard or **Run as** box in the scheduled task's property dialog box is accurate. If an account password expires any scheduled task for which that user has been specified will not run. For more information, see Task Scheduler and security.

**Encourage your users to use strong passwords.**

- Scheduled tasks always run under the context of a specific account, and the set of permissions for which the group it belongs to provides. To prevent a malicious user from modifying a scheduled task using another user's account name and password, encourage the use of strong passwords. For more information about strong passwords, see Strong passwords.

**Choose carefully when deciding which tasks to schedule under the System account.**

- Tasks scheduled under the System (or AT Service account) will run regardless of which user is logged on to a computer. Often these tasks will run in the background, unnoticed by the user.

  For more information about the implications of running a task under the System account, see Act as part of the operating system.

**Before stopping a scheduled task that is running, check for other scheduled tasks that are dependent on the task.**

- If you stop a scheduled task that is running without checking for dependent tasks, other tasks might not run or run incorrectly.

**Only create scheduled tasks on domain controllers if absolutely necessary.**

- If you are a member of the Administrators group, and the local computer on which you create a scheduled task is a domain controller, the scheduled task could affect objects in the entire domain.

**Avoid entering a member of the Domain group in the Add Scheduled Task wizard or Run as box when creating a scheduled task.**

- It is possible that a malicious user who is a member of the Administrators group could obtain those credentials and access the entire domain.

# Task Scheduler How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Schedule Tasks
- Manage Task Scheduler

# Schedule Tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Schedule tasks**

- Schedule a new task
- Modify a scheduled task
- Remove a scheduled task
- Stop a scheduled task that is running
- Run a scheduled task immediately
- Schedule a task on a remote computer

# Schedule a new task

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To schedule a new task**

1. Open Task Scheduler.
2. Double-click **Add Scheduled Task**.
3. Follow the instructions in the Scheduled Task wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators, Backup Operators, or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- When creating a scheduled task, you must enter a user name and password, either in the **Add Scheduled Task Wizard** or in the **Run as** box in the **Task** tab of the scheduled task's property dialog box. When the scheduled task runs, the program you've scheduled runs as if it were started by the user you specified, with that user's security context.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- Confirm that the system date and time on your computer are accurate, as Task Scheduler relies on this information to run scheduled tasks. To verify or change this information, double-click the time indicator on the taskbar.
- If you want to configure advanced settings for the task, select the **Open advanced properties for this task when I click Finish** check box in the final page of the wizard. This opens the properties dialog box for the task. You can then change the program being run in the **Task** tab, fine-tune the schedule in the **Schedule** tab, customize settings in the **Settings** tab, or set user and group permissions in the **Security** tab.
- If the task program requires command-line options, type them in **Run**, after the task path.
- If the path to the task program includes spaces, type double quotation marks ("") around the entire task path. For example:

  **"C:\Program Files\Windows Media Player\Mplayer2.exe"**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Task Scheduler Best practices
Task Scheduler overview
Modify a scheduled task
Creating and managing scheduled tasks from the Command Line
Strong passwords

# Modify a scheduled task

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify a scheduled task**

1. Open Task Scheduler.
2. Right-click the task you want to modify, and then click **Properties**.
3. Do one or more of the following:
   - To change the program being run, in **Run**, type the path for the new program.
   - To change the user account name and password that the scheduled task will run as, edit or enter a new name in the **Run as** box, and then click **Set password**.
   - To change the schedule for the task, click the **Schedule** tab.
   - To customize the settings for the task, click the **Settings** tab.
   - To set user and group permissions for the task, click the **Security** tab.

## Notes

- To perform this procedure, you must be a member of the Administrators, Backup Operators, or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- Confirm that the System date and time on your computer are accurate, as Task Scheduler relies on this information to run scheduled tasks. To verify or change this information, double-click the time indicator on the taskbar.
- If the task program requires command-line options, type them in **Run**, after the task path.
- If the path to the task program includes spaces, type double quotation marks ("") around the entire task path. For example:

  **"C:\Program Files\Windows Media Player\Mplayer2.exe"**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Schedule a new task
Creating and managing scheduled tasks from the Command Line

# Remove a scheduled task

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a scheduled task**

1. Open Task Scheduler.
2. Right-click the task that you want to remove, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Administrators, Backup Operators, or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- You can also remove a scheduled task by selecting a task in the details pane, and then clicking **Delete** on the **File** menu.
- Removing a scheduled task removes only the .job file from the schedule. The program file the task runs is not removed from the hard disk.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modify a scheduled task
Creating and managing scheduled tasks from the Command Line

# Stop a scheduled task that is running

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop a scheduled task that is running**

1. Open Task Scheduler.
2. Click the task that you want to stop.
3. From the **File** menu, click **End Scheduled Task**.

### Notes

- To perform this procedure, you must be a member of the Administrators, Backup Operators, or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- If you stop a scheduled task that is currently running, you might experience a delay (up to 3 minutes) before the task shuts down.
- To restart a stopped task, right-click the task, and then click **Run**.

### Note

- Before stopping a scheduled task that is running, make sure you check for other scheduled tasks that are dependent on the task. If you stop a scheduled task that is running without checking what other tasks are dependent on it, other tasks might not run or run incorrectly.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Pause the Task Scheduler service
Stop the Task Scheduler service

# Run a scheduled task immediately

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run a scheduled task immediately**

1. Open Task Scheduler.
2. Right-click the task that you want to run, and then click **Run**.

## Notes

- To perform this procedure, you must be a member of the Administrators, Backup Operators, or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- You can also run a scheduled task immediately by selecting a task in the details pane, and then clicking **Run** on the **File** menu.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Stop a scheduled task that is running
Modify a scheduled task

# Schedule a task on a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To schedule a task on a remote computer**

1. Open My Network Places.
2. Click **Search for Computers**.
3. In the **Computer name** box, enter the computer name, and then click **Search**.
4. In the details pane, double-click the remote computer on which you want to view scheduled tasks.
5. Double-click **Scheduled Tasks**.
6. On the **File** menu, point to **New**, and then click **Scheduled Task**.
7. Type a name for the new task, and then press ENTER.
8. Right-click the new task, and then click **Properties**.
9. Do the following:
    - To specify the program being run, in **Run**, type the path for the new program.
    - To set the user account name and password that the scheduled task will run as, type a new name in **Run as**, and then click **Set password**.
    - To set the schedule for the task, click the **Schedule** tab.
    - To customize the settings for the task, click the **Settings** tab.
    - To set user and group permissions for the task, click the **Security** tab.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or have been delegated the appropriate authority, on the remote computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Network Places, click **Start**, and then click **My Computer**. Under **Other Places**, click **My Network Places**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Creating and managing scheduled tasks from the Command Line
Schedule a new task
Modify a scheduled task

# Manage Task Scheduler

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage Task Scheduler**

- Stop the Task Scheduler service
- Restart the Task Scheduler service
- Pause the Task Scheduler service
- Receive notification of missed tasks
- Set the AT Service Account
- View a log of past scheduled tasks
- Set power management options for a scheduled task
- Configure Task Scheduler using Computer Management

# Stop the Task Scheduler service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop the Task Scheduler service**

1. Open Task Scheduler.
2. On the **Advanced** menu, click **Stop Using Task Scheduler**.

## Notes

- To perform this procedure, you must be a member of the Administrators or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- Tasks will not run at their scheduled times until you restart the Task Scheduler service.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restart the Task Scheduler service
Pause the Task Scheduler service
Administering Task Scheduler using Computer Management
Configure Task Scheduler using Computer Management

# Restart the Task Scheduler service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restart the Task Scheduler service**

1. Open Task Scheduler.
2. On the **Advanced** menu, click **Start Using Task Scheduler**.

## Notes

- To perform this procedure, you must be a member of the Administrators or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Stop the Task Scheduler service
Pause the Task Scheduler service
Administering Task Scheduler using Computer Management
Configure Task Scheduler using Computer Management

# Pause the Task Scheduler service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To pause the Task Scheduler service**

1. Open Task Scheduler.
2. On the **Advanced** menu, click **Pause Task Scheduler**.

## Notes

- To perform this procedure, you must be a member of the Administrators or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- The **Pause Task Scheduler** command is useful if you do not want scheduled tasks to run at the same time you are installing software or running another program (such as a game).
- To resume the schedules for all tasks, on the **Advanced** menu, click **Continue Task Scheduler**. Once the service is restarted, tasks will run at their previously scheduled times.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Stop the Task Scheduler service
Restart the Task Scheduler service
Stop a scheduled task that is running
Administering Task Scheduler using Computer Management
Configure Task Scheduler using Computer Management

# Receive notification of missed tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To receive notification of missed tasks**

1. Open Task Scheduler.
2. On the **Advanced** menu, click **Notify Me of Missed Tasks**.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- If you enable the **Notify Me of Missed Tasks** setting, the missed tasks dialog box displays at the foreground of the computer display, in front of the logon screen, when there are missed tasks to report. Anyone who sits at the computer can view the list of missed tasks.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- A check mark next to **Notify Me of Missed Tasks** indicates that the feature is turned on. You can click **Notify Me of Missed Tasks** again to clear the check mark and turn off the feature.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View a log of past scheduled tasks
Schedule a task on a remote computer
Administering Task Scheduler using Computer Management
Configure Task Scheduler using Computer Management

# Set the AT Service Account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the AT Service Account**

1. Open Task Scheduler.
2. On the **Advanced** menu, click **AT Service Account**.
3. Click **This Account**, and then type the user account that will run tasks that are scheduled by using the **at** command.
4. In the **Password** and **Confirm Password** boxes, type the password for the account.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- The AT service account is used by tasks that have been scheduled by using the **at** command.
- By default, tasks scheduled using the **at** command stop after 72 hours. For more information, see **At** in Related Topics.
- Unless you click **This Account** and type the user account name and password, tasks are scheduled using **System Account**, which is the default selection in the **AT Service Account Configuration** dialog box. If you run the **at** command for tasks by using the System Account, the scheduled tasks run using the default System account.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Task Scheduler Best practices
Logon rights
Creating and managing scheduled tasks from the Command Line
Act as part of the operating system
At

# View a log of past scheduled tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view a log of past scheduled tasks**

1. Open Task Scheduler.

2. On the **Advanced** menu, click **View Log**.

## Notes

- To perform this procedure, you must be a member of the Administrators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.

- The Task Scheduler log file lists all past scheduled tasks under all users, including members of the Administrators group and the System account. The file is located in the \Windows\Tasks folder, and is named SchedLgU.txt. In the SchedLgU.txt file, to see the most recent task that has run, find the following text:

  **[ ***** Most recent entry is above this line ***** ]**. The most recent task that has run is the entry directly above this text.

- The default log file size is 32 kilobytes.

- When the log file reaches its maximum size, it automatically starts recording new information at the beginning of the log file and writes over old log file information.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set, view, change, or remove permissions on files and folders
Receive notification of missed tasks

# Set power management options for a scheduled task

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set power management options for a scheduled task**

1. Open Task Scheduler.
2. Right-click the task for which you want to set power management options, and then click **Properties**.
3. On the **Settings** tab, under **Power Management**, do one or more of the following:
   - To conserve the battery power of your computer, select the **Don't start the task if the computer is running on batteries** check box.
   - To stop the task if the computer starts running on battery power, select the **Stop the task if battery mode begins** check box.
   - To wake your computer from hibernation or standby to run a task, select the **Wake the computer to run this task** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.
- By default, power management is set to refrain from starting scheduled tasks if the computer is running on battery power. If you are using a portable computer that runs on batteries and you want a scheduled task to run, clear the default check boxes.
- If the power is interrupted or if the battery fails, the scheduled task that is running might not complete.
- The **Wake the computer to run this task** check box is only displayed on computers that support it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Creating and managing scheduled tasks from the Command Line
**Other Resources**
Power Options overview

# Configure Task Scheduler using Computer Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Task Scheduler using Computer Management**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management (Local)/Services and Applications

3. In the details pane, double-click **Task Scheduler**.
4. Edit the appropriate settings in the **General**, **Log On**, **Recovery**, or **Dependencies** tab.

### Note

- To perform this procedure, you must be a member of the Administrators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Task Scheduler using Computer Management
Creating and managing scheduled tasks from the Command Line

# Task Scheduler Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Task Scheduler:

- Task Scheduler overview
- Task Scheduler and security
- Administering Task Scheduler using Computer Management
- Creating and managing scheduled tasks from the Command Line

# Task Scheduler overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Task Scheduler overview**

Using Task Scheduler, you can schedule any script, program, or document to run at a time that is most convenient for you.

When you first install your Windows Server 2003 family operating system, the Task Scheduler service is enabled. By default, you must be a member of the Administrators, Backup Operators, or Server Operators group on the local computer, to view, add, edit, or delete scheduled tasks, or to disable, pause, or restart the Task Scheduler service. For more information, see Task Scheduler and Security.

With Task Scheduler, you can:

- Schedule a task to run daily, weekly, monthly, or at certain times (such as system startup).
- Change the schedule for a task.
- Stop a scheduled task.
- Customize how a task runs at a scheduled time.

Each scheduled task you create is stored as a .job file in the \Windows\Tasks folder. The .job file contains the properties and configuration information for the task. You can create a scheduled task on your computer and then drag the .job object over to a remote computer. After you drag a task from one computer to another, you must update the account information for the task before it will run. Removing a scheduled task removes only the .job file from the schedule. The program file the task runs is not removed from the hard disk.

For tips on using Task Scheduler, see Task Scheduler Best practices.

# Task Scheduler and security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Task Scheduler and security**

When creating a scheduled task, you must enter a user name and password, either in the **Add Scheduled Task Wizard** or in the **Run as** box in the **Task** tab of the scheduled task's property dialog box. When the scheduled task runs, the program you've scheduled runs as if it were started by the user you specified, with that user's security context. For example, if the user specified for a scheduled task is a member of the Backup Operators group on the local computer, the program specified in the scheduled task file runs as if a member of the Backup Operators group is logged onto the local computer. If another user is logged on to the computer at the time a scheduled task specified for a different user runs, the task runs but is not visible to the current user.

By default, to schedule a task, you must be a member of the Administrators, Backup Operators, or Server Operators group on the local computer. By default, when creating a scheduled task, you cannot enter a user who belongs to a group that has more rights than the group you belong to. For example, if you are a member of the Backup Operators group on the local computer, you cannot specify a member of the Administrators group when creating a scheduled task. However, a member of the Administrators group can enable a member of any group to create or modify scheduled tasks, by using the **cacls** command to modify the discretionary access control list (DACL) of the **Tasks** folder. By default, the **Tasks** folder is located in the **\Windows** folder on the hard drive of the local computer, for example **C:\Windows\Tasks**. For more information about using **cacls**, see Cacls. For more information about groups, see Default local groups and Default groups.

**Note**

- A user for whom you assign permissions to the **Tasks** folder using **cacls** will be able to access scheduled tasks for all users. Choose which users to give access to the **Tasks** folder judiciously.

In Windows Server 2003 family operating systems, passwords have an expiration date. If you are scheduling recurring tasks that run indefinitely, you need to be aware of the expiration date on your passwords. If the password changes for an account on a domain, then you must update the tasks scheduled to run under that account.

When a job is created through Task Scheduler, account information for all tasks that use the same **run as** account is stored only once. Task Scheduler validates the **run as** account when a job is created, which prevents incorrect **run as** account information from being saved and, as a result, prevents other existing jobs that use the same account from being affected. If the password validation fails, a job file is created, but it will not run.

If the **run as** account password has expired or the account has been deleted and recreated, the password must be updated for jobs to run. Updating the password for one job automatically updates it for all jobs that use the same **run as** account.

When you create a scheduled task, the user credentials you specify for that specific task are securely stored on the local computer.

For security-related tips on using Task Scheduler, see Task Scheduler Best practices.

**Notes**

- To view the user and group permissions for a scheduled task, right-click the task, click **Properties**, and then click the **Security** tab.
- To view or change advanced user and group permissions, on the **Security** tab, click **Advanced**, click the user for whom you want to view or change permissions, and then click **View/Edit**.
- For information about viewing past scheduled tasks, see View a log of past scheduled tasks.
- For general information about Task Scheduler, see Task Scheduler overview.

# Administering Task Scheduler using Computer Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Task Scheduler using Computer Management**

Using Computer Management, you can do the following to a local or remote computer console:

- Stop, start, or pause the Task Scheduler service
- Specify start parameters for the Task Scheduler service
- Specify logon settings for the Task Scheduler service
- Select the computer's response if the Task Scheduler service fails
- Set or verify dependencies for the Task Scheduler service

For more information, see Configure Task Scheduler using Computer Management.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- The **Task Scheduler Properties (*computername*)** dialog box is in the **Services** snap-in. To make the dialog box available in MMC, you must add the **Services** snap-in to your Computer Management console.

# Creating and managing scheduled tasks from the Command Line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating and managing scheduled tasks from the Command Line**

In addition to using Task Scheduler, you can use the **at** and **schtasks** command-line utilities to create and manage scheduled tasks.

For more information about the **at** command, including all available parameters, see At. For more information about the **schtasks** command, including all available parameters, see Schtasks. For information about other command-line utilities, see Command-line reference A-Z.

# Troubleshooting Task Scheduler

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Task does not run at the scheduled time.

**Task does not run at the scheduled time.**
**Cause:** The password for the user you specified when creating the scheduled task might have changed.

**Solution:** Open the scheduled task, and then update the password for the user.

**Note**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**See also:** Modify a scheduled task; Schedule a new task; Task Scheduler and security; Task Scheduler Best practices

or

**Cause:** The password for the user you specified when creating the scheduled task might be blank.

**Solution:** Assign a password for the user account for which you are scheduling the task. For more information, see User and Group Accounts. After the user has a password, modify the scheduled task to include the password.

**Note**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**See also:** Modify a scheduled task; Schedule a new task; Task Scheduler and security; Task Scheduler Best practices

or

**Cause:** Task Scheduler cannot find the executable for the program that is supposed to run at the scheduled time because the path to it is set incorrectly.

**Solution:**

1. Open Task Scheduler.
2. Double-click the scheduled task.
3. In the **Task** tab, verify that the paths to the program executable are correct in the **Run** and **Start in** boxes.

**Note**

- To perform this procedure, you must be a member of the Administrators, Backup Operators, or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

**See also:** Task Scheduler Best practices; Schedule a new task; Creating and managing scheduled tasks from the Command Line; Configure Task Scheduler using Computer Management

or

**Cause:** The task was scheduled to run under another user account than the user currently logged on.

**Solution:**

1. Open Task Scheduler.

2. In the **Task** tab, in the **Run as** box, and then type the correct user name.

**Note**

- To perform this procedure, you must be a member of the Administrators, Backup Operators, or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

or

**Cause:** The Task Scheduler service might be stopped.

**Solution:**

1. Open Task Scheduler.

2. In the **Advance** menu, click **Start Using Task Scheduler**.

**Note**

- To perform this procedure, you must be a member of the Administrators or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

**See also:** Administering Task Scheduler using Computer Management

or

**Cause:** The Task Scheduler service might be disabled.

**Solution:**

1. Open Computer Management.

2. In the console tree, click **Services**.

   **Where?**

   - Computer Management (Local)/Services and Applications

3. In the details pane, double-click **Task Scheduler**.

4. In the **General** tab, in **Startup type** list box, select **Automatic** or **Manual**. Click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators or Server Operators group, or have been delegated the appropriate authority, on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

You can also stop and start the Task Scheduler service in the Microsoft Management Console: open the **Task Scheduler Properties (*computername*)** dialog box, and then under the **General** tab, under **Service status**, click **Start** or **Stop**.

**See also:** Administering Task Scheduler using Computer Management

# Logon Scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logon Scripts**
Logon scripts are groups of commands that run when a user logs on to a computer. Administrators design logon scripts to automate the process of configuring a user environment. Administrators use the tools described below to manage logon scripts for individual user accounts, especially those defined locally.

- For help with specific tasks, see Logon Scripts How To....
- For general background information, see Logon Scripts Concepts.

Administrators use a different set of tools to manage scripts for groups of users and computers defined in Active Directory®. For more information about Group Policy startup, shutdown, logon, and logoff scripts, see Group Policy (pre-GPMC).

# Logon Scripts How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Assign a logon script to a user or group
- Assign a logon script to a user in the Active Directory

# Assign a logon script to a user or group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign a logon script to a user or group**

1. Open Computer Management.
2. In the console tree, click **Users**.

   **Where?**

   - System Tools/Local Users and Groups/Users

3. Double-click the user to which you want to assign a logon script.
4. Click the **Profile** tab.
5. In the **Logon script** field, enter the path and name of the logon script you want to assign to that user, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can store local logon script files in any local or shared folder that the user can access during logon.
- For most members of the Windows server 2003 family, the default location for local logon scripts is the *%systemroot%*\System32\Repl\Imports\Scripts folder, which is shared using the share name **netlogon**. However, Windows Server 2003, Standard Edition, does not create this folder.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy (pre-GPMC)

# Assign a logon script to a user in the Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign a logon script to a user in the Active Directory**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.
3. Double-click the user to which you want to assign the logon script.
4. Click the **Profile** tab.
5. In the **Logon script** field, type the path and name of the logon script you want to assign to that user, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to Active Directory
Group Policy overview

# Logon Scripts Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about logon scripts:

- Logon scripts overview
- Understanding logon scripts
- Logon script assignment
- Logon script replication
- Logon Scripts Environment variables

# Logon scripts overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logon scripts overview**

A system administrator can use logon scripts to assign tasks that will be automatically performed whenever a user logs on to a particular computer system. These scripts can use system environment variables, and can also call other scripts or executable programs. Logon scripts are frequently used to map network drives, start background processes, and initialize user environment variables.

# Understanding logon scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding logon scripts**

A logon script runs automatically whenever a user logs on to a computer running a member of the Windows Server 2003 family of operating systems. The script can contain operating system commands, such as those that make network connections or start programs. Logon scripts can also set environment variables to specify information such as the computer search path and the directory for temporary files. A logon script is usually a batch file (.bat or .cmd file name extension), but any executable program can be used.

Logon scripts are optional. You can use them to configure user working environments by creating network connections and starting programs. Logon scripts are useful when you want to affect the user work environment without managing all aspects of it.

Script files are text files that contain script commands. The Windows Server 2003 family of operating systems supports these types of scripts:

- Batch file commands are stored in text files with the .bat or .cmd file name extension. Batch files automate simple series of tasks that would otherwise be run from a command line. Scripts written using batch file commands are run by the command shell. For more information about the command shell, see Command shell overview.

- Visual Basic Scripting Edition (VBScript) commands are stored in text files with the .vbs file name extension, and JScript commands are stored in text files with the .js file name extension. VBScript and JScript allow the administrator to construct sophisticated scripts. The Windows Script Host can run these scripts from the desktop of the computer, or from a command line. For more information about the Windows Script Host, see Windows Script Host.

After you create a logon script, you can assign it to one or more local users, sites, domains, or organizational units (OUs).

For more information about specific tasks related to assigning scripts, see Logon Scripts How To....

# Logon script assignment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logon script assignment**

On computers running operating systems in the Windows Server 2003 family, you can assign a logon script to a user account. When a user logs on and a path to a logon script is present in the user account, the file is located and run.

You can also assign logon and logoff scripts, and computer startup and shutdown scripts by using the Group Policy snap-in. These scripts apply to the entire scope of users and computers for which a particular Group Policy object applies.

In **Computer Management**, you can use the **User Property** dialog box to assign logon scripts to user accounts by typing the file name (for example, Clerks.bat) in the **Logon script** text box. At logon, the server authenticating the logon locates an assigned logon script. It looks for the specified file following the local logon script path on the server (usually %systemroot%\SYSVOL\sysvol\domain_name\scripts). If a relative path is provided before the file name (for example, Admins\User1.bat), the server looks for the logon script in that subdirectory of the logon script path.

The entry in the **Logon Script** text box specifies only the file name (and, optionally, the relative path) and does not create the actual logon script. You create a logon script with the specified name and place it in the appropriate directory on the appropriate replication export server.

You can place a logon script in a local directory on a user's computer. You would typically use this location, however, when you are administering user accounts that exist on a single computer rather than in a domain. This logon script runs only when a user logs on locally to the computer and does not run when the user logs on to the domain. You must place the logon script using the computer's logon script path or in a subdirectory of that logon script path. The default location for local logon scripts is the %systemroot%\System32\Repl\Imports\Scripts folder. This folder is not created on an new installation of Windows XP. A folder must be created and shared with the name **netlogon**; for step-by-step instructions, see Share a folder or drive. The NTFS permissions of this folder should allow users and server operators only read and execute permissions, and should allow administrators full control. The folder is created and shared automatically on domain controllers, so you should not attempt to create a **netlogon** folder on a domain controller manually.

For more information, see: Logon Scripts How To..., Group Policy overview, and Privileges.

# Logon script replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logon script replication**

A logon script is always downloaded from the server that validates a user's logon request. For users with accounts on Windows Server 2003 family domains that have one or more domain controllers, any one of the domain controllers can authorize a user's logon attempt. In order for logon scripts to always work for users, logon scripts must exist for all user accounts in a domain, on every domain controller in that domain.

Windows Server 2003, Standard Edition, domain controllers and servers use File Replication Service (FRS) to replicate system policy and logon scripts to every domain controller in a domain. When you make a change to a logon script stored in the SYSVOL share of a domain controller, FRS replicates this change to the SYSVOL shares of all other domain controllers in the domain. If the administrator manages scripts for a domain using Active Directory Users and Computers, then the logon scripts are typically located in the %systemroot%\SYSVOL\sysvol*domain_name*\Scripts folder. The system will look for scripts in the **netlogon** share of the domain controller that authenticates the user.

If the administrator uses Group Policy to manage logon scripts on a standalone system, then the user logon scripts are typically located in the %systemroot%\System32\GroupPolicy\User\Scripts\Logon directory.

For specific information on how to configure the logon script path for a user in a domain, see Assign a logon script to a user in the Active Directory.

For background information about replication, see Understanding Sites and Replication and Replication overview.

For instructions about configuring replication, see Configure Replication Between Sites.

# Logon Scripts Environment variables

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Environment variables**
When managing multiple user and group accounts, you may need to make the same change to many accounts. You can use environment variables to replace static names or labels with a dynamic variable. The dynamic variable is replaced by the appropriate environment data when called.

**Changing the system environment variables**
Servers running Windows Server 2003 require certain information to find programs, to allocate memory space for some programs to run, and to control various programs. This information, called the system and user environment variables, can be viewed using the **System Properties** in **Control Panel** on the **Advanced** tab. These environment variables are similar to those that can be set in the MS-DOS operating system, such as PATH and TEMP.

The *system environment variables* are the same regardless who is logged on to the computer. If you are logged on as a member of the Administrators group, you can add new variables or change the values.

The *user environment variables* can be different for each user of a particular computer. They include any environment variables you want to define, or variables defined by your applications, such as the path where application files are located.

After you change any environment variables, the new values are saved in the registry, making them automatically available the next time you start your computer.

If any conflict exists between environment variables, variables are set in this way:

- System environment variables are set first.
- Variables defined in the Autoexec.bat file (except for Path variables) are set next, and override system variables.
- User environment variables defined in **System Properties** are set next, and override both the system and Autoexec.bat variables.
- Path variables defined in the Autoexec.bat file are set last.

Path settings, unlike other environment variables, are cumulative. The full path (what you see when you type **path** at the command prompt) is created by appending the path contained in Autoexec.bat to the paths defined in **System Properties** in **Control Panel**.

**Using system environment variables in user profile paths, home directory paths, and logon scripts**
Any system environment variable on a Windows-based client computer can be used in a user account profile path, logon script path, home directory path, and within a logon script itself. For a list of system environment variables, see Command shell overview

This feature helps to ensure that logon scripts and user profiles run most efficiently on domains that span wide area network (WAN) links, especially if users work at more than one site.

For example, suppose you have two physical sites, one located in Paris, and the other located in London. On every computer running an operating system in the Windows Server 2003 family at the London site, you can set the *servername* system environment variable to the computer name of a backup domain controller in London. On the computers in Paris, you can set *servername* in a similar way, using the computer name of the backup domain controller in Paris. For domain user accounts, you can use %servername% in the logon script paths (for example, \\%servername%\scripts). When a user logs on, the logon script is always loaded from a server at the local site.

For more information, see the Resource Kit.

For more information about constructing scripts, see System Administration Scripting Guide at the Microsoft Windows Resource Kits Web site.

# Windows Script Host

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Script Host**
Windows Script Host creates an environment in which script engines run the commands contained in script files. The Windows Script Host supports several script engines, including JScript and VBScript. Windows Script Host is ideal for non-interactive scripting needs, and it allows you to run scripts from the Windows desktop or from the command prompt.

- For help with specific tasks, see Windows Script Host How To....
- For general background information, see Windows Script Host Concepts.

# Windows Script Host How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Run Windows Script Host
- Run Windows Script Host samples
- Set global properties for all scripts that are run with the Windows-based script host

# Run Windows Script Host

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run Windows Script Host**

Click the task that you want to perform.

- To start a script so that it runs in a Windows environment
- To start a script so that it runs in a command-line environment

**To start a script so that it runs in a Windows environment**

1. Open Command Prompt.
2. Type: **wscript**[*ScriptName*] [**//B**] [**//D**] [**//E:***Engine*] [**//H:CScript** | **//H:WScript**] [**//I**] [**//Job:***Identifier*] [**//Logo** | **//NoLogo**] [**//S**] [**//T:***Number*] [**//X**] [**//?**] [*ScriptArguments*]

**Parameters**

***ScriptName***

Specifies the path and file name of the script file.

**//B**

Specifies batch mode, which does not display alerts, scripting errors, or input prompts.

**//D**

Starts the debugger.

**//E: *Engine***

Specifies the engine that is used to run the script.

**//H:cscript**

Registers Cscript.exe as the default script host for running scripts.

**//H:wscript**

Registers Wscript.exe as the default script host for running scripts. This is the default.

**//I**

Specifies interactive mode, which displays alerts, scripting errors, and input prompts. This is the default and the opposite of **//B**.

**//Job: *Identifier***

Runs the job identified by *Identifier* in a .wsf script file.

**//Logo**

Specifies that the Windows Script Host banner is displayed in the console before the script runs. This is the default and the opposite of **//Nologo**.

**//Nologo**

Specifies that the Windows Script Host banner is not displayed before the script runs.

**//S**

Saves the current command-prompt options for the current user.

**//T: *Number***

Specifies the maximum time the script can run (in seconds). You can specify up to 32,767 seconds. The default is no time limit.

**//X**

Starts the script in the debugger.

**//?**
Displays available command parameters and provides help for using them.

***ScriptArguments***
Specifies the arguments passed to the script. Each script argument must be preceded by a slash (/).

**Remarks**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Each parameter is optional; however, you cannot specify script arguments without specifying a script. If you do not specify a script or any script arguments, Wscript.exe displays the **Windows Script Host Settings** dialog box, which you can use to set global scripting properties for all scripts that Wscript.exe runs on the local computer.
- The **//T** parameter prevents excessive running of scripts by setting a timer. When the time exceeds the specified value, WScript interrupts the script engine and ends the process.
- Windows script files usually have one of the following file name extensions: .wsf, .vbs, .js.
- If you double-click a script file with an extension that has no association, the **Open With** dialog box appears. Select Wscript or Cscript, and then select **Always use this program to open this file type**. This registers Wscript.exe or Cscript as the default script host for files of this file type.
- You can set properties for individual scripts. See Related Topics for more information.
- Windows Script Host can use .wsf script files. Each .wsf file can use multiple scripting engines and perform multiple jobs.

**To start a script so that it runs in a command-line environment**

1. Open Command Prompt.
2. Type: **cscript**[*ScriptName*] [**//B**] [**//D**] [**//E:***Engine*] [**//H:CScript** | **//H:WScript**] [**//I**] [**//Job:***Identifier*] [**//Logo** | **//NoLogo**] [**//S**] [**//T:***Number*] [**//X**] [**//U**] [**//?**] [*ScriptArguments*]

**Parameters**
***ScriptName***
Specifies the path and file name of the script file.

**//B**
Specifies batch mode, which does not display alerts, scripting errors, or input prompts.

**//D**
Starts the debugger.

**//E: *Engine***
Specifies the engine that is used to run the script.

**//H:cscript**
Registers Cscript.exe as the default script host for running scripts.

**//H:wscript**
Registers Wscript.exe as the default script host for running scripts. This is the default.

**//I**
Specifies interactive mode, which displays alerts, scripting errors, and input prompts. This is the default and the opposite of **//B**.

**//Job:** ***Identifier***
Runs the job identified by *Identifier* in a .wsf script file.

**//Logo**
Specifies that the Windows Script Host banner is displayed in the console before the script runs. This is the default and the opposite of **//Nologo**.

**//Nologo**
Specifies that the Windows Script Host banner is not displayed before the script runs.

**//S**
Saves the current command-prompt options for the current user.

**//T:** ***Number***
Specifies the maximum time the script can run (in seconds). You can specify up to 32,767 seconds. The default is no time limit.

**//U**
Specifies Unicode for input and output that is redirected from the console.

**//X**
Starts the script in the debugger.

**//?**
Displays available command parameters and provides help for using them. This is the same as typing Cscript.exe with no parameters and no script.

***ScriptArguments***
Specifies the arguments passed to the script. Each script argument must be preceded by a slash (/).

**Remarks**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Each parameter is optional; however, you cannot specify script arguments without specifying a script. If you do not specify a script or any script arguments, Cscript.exe displays the Cscript.exe syntax and the valid host options.
- The **//T** parameter prevents excessive running of scripts by setting a timer. When the run time exceeds the specified value, CScript interrupts the script engine and ends the process.
- Windows script files usually have one of the following file name extensions: .wsf, .vbs, .js.
- You can set properties for individual scripts. See Related Topics for more information.
- Windows Script Host can use .wsf script files. Each .wsf file can use multiple scripting engines and perform multiple jobs.
- If you double-click a script file with an extension that has no association, the **Open With** dialog box appears. Select Wscript or Cscript, and then select **Always use this program to open this file type**. This registers Wscript.exe or Cscript as the default script host for files of this file type.
- You can set properties for individual scripts. See Related Topics for more information.
- Windows Script Host can use .wsf script files. Each .wsf file can use multiple scripting engines and perform multiple jobs.

**Formatting legend**

| Format | Meaning |
|---|---|

| | |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Windows Script Host overview
Command-line reference

# Run Windows Script Host samples

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run Windows Script Host samples**

1. Open the Command Prompt.

2. Type the following:

   **cscript***samples folder******script name*

   Where:

   *samples folder* is the path to the folder into which you downloaded the samples.

   *script name* is the name of the sample script you want to run, including the file name extension (for example, .js or .vbs).

   For example, if you saved the samples to C:\Samples, and you wanted to run the Network.vbs sample, you would type:

   **cscript c:\sample\network.vbs**

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You must download the Windows Script Host sample files from the **Microsoft Windows Script Technologies** home page at the Microsoft Web site (http://www.microsoft.com/) before you can run them.

- You can also double-click the sample script file to run a script using the Wscript.exe script host.

**See Also**
**Concepts**
Windows Script Host overview
Run Windows Script Host

# Set global properties for all scripts that are run with the Windows-based script host

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set global properties for all scripts that are run with the Windows-based script host**

1. Click **Start**, and then click **Run**.
2. In the **Open** box, type *wscript*, and then click **OK**. The **Windows Script Host Settings** dialog box appears.
3. Select the **Stop script after specified number of seconds** check box.
4. In the **seconds** box, enter the time limit you want to place on all scripts.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- With one setting, you can set properties for all scripts that are run with the Windows-based script host. Every script that is run with this host automatically uses these settings.

**See Also**
**Concepts**
Windows Script Host overview

# Set properties of individual scripts that are run with the Windows-based script host

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set properties for individual scripts that are run with the Windows-based script host**

1. Open Windows Explorer. Browse to the folder with the script whose properties you want to set, and right-click it.
2. Select **Properties**, and then click the **Script** tab.
3. Select the **Stop script after specified number of seconds** check box.
4. In the **seconds** box, enter the time limit you want to place on this script.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- These settings override global settings.
- After clicking **OK**, a command file is created that binds the property settings to the script. It has the file name extension .wsh. You can use this command file as an alternate script file to run the script.

For more information, see the Windows Script Technologies Web site at the Microsoft Web site (http://www.microsoft.com/).

**See Also**
**Concepts**
Windows Script Host overview
Set global properties for all scripts that are run with the Windows-based script host
Run Windows Script Host

# Windows Script Host Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about the Windows Script Host:

- Windows Script Host overview
- Windows Script Host object model

# Windows Script Host overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Script Host overview**

You can use Windows Script Host to run scripts by clicking a script file on the Windows desktop or by typing the name of a script file at the command prompt. Similar to Microsoft Internet Explorer, Windows Script Host serves as a controller of ActiveX scripting engines. Unlike Internet Explorer, Windows Script Host has very low memory requirements and is ideal for both interactive and non-interactive scripting, such as logon scripting and administrative scripting.

There are two versions of the Windows Script Host: a Windows-based version (Wscript.exe), which provides a property sheet for setting script properties, and a command-prompt-based version (Cscript.exe), which provides command-line switches for setting script properties. You can run either of these by typing Wscript.exe or Cscript.exe at the command prompt.

Previously, the only native scripting language supported by the Windows operating system was that of batch files. Although batch language is fast and small, its features are limited compared to VBScript and JScript. For example, the ability to control program flow was not built into the batch language. Today, the Windows Script Host architecture allows you to take advantage of these powerful scripting languages, although support is still provided for batch files.

This version of WSH adds several areas of new functionality:

- **Improved argument handling.** Handling and documenting command-line arguments is now simpler. For more information, see **What's New in WSH 5.6** at the Microsoft Web site (http://www.microsoft.com/)

- **Remote scripting.** You can load a script onto several remote computer systems which you can then run at the same time.

- **A running process is an object.** You can sink events, access standard I/O streams, and share environment variables with other processes.

- **Flexible working directory.** You can determine and modify the path of the current working directory of the active process.

- **Improved script security.** The new script security model makes it safer to distribute and run scripts. For more information, see Windows Script (http://www.microsoft.com/), Microsoft Windows Resource Kits Web site (http://www.microsoft.com/)., and Security and Windows Script Host (http://www.microsoft.com/).

# Windows Script Host object model

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Script Host object model**

The Windows Script Host (WSH) environment includes a built-in library, WScript, and four COM (Component Object Model)-based libraries: WshShell, WshNetwork, WshController and the Scripting Runtime Library. Together, WScript, WshShell, WshNetwork and WshController are referred to as the WSH Objects. Your scripts access the functionality of these libraries through the WSH Object Model.

Scripts use object models to access features of COM-based libraries. In COM-based libraries an object model documents the features of the library. Scripts call an object's methods and manipulate an object's properties to utilize the features required to accomplish a task.

For more information about the Windows Script Host object model, see the Windows Script Host Object Model page at the Microsoft Web site.

For general information about Windows scripting, see the Windows Script Technologies Web site at the Microsoft Web site, and Microsoft Windows Resource Kits Web site.

# Scripting Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scripting tools**
This section covers:

- Using batch files
- Netsh
- Logon Scripts
- Windows Management Instrumentation Command-line

For information about Microsoft® Visual Basic® Scripting Edition (VBScript), see "VBScript Primer" at the Microsoft Windows Resource Kits Web site.

For more information about creating scripts, see "Creating Enterprise Scripts" at the Microsoft Windows Resource Kits Web site. For more information about managing scripts, see "Managing Enterprise Scripts" at the Microsoft Windows Resource Kits Web site.

# Using batch files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using batch files**

With batch files, which are also called batch programs or scripts, you can simplify routine or repetitive tasks. A batch file is an unformatted text file that contains one or more commands and has a .bat or .cmd file name extension. When you type the file name at the command prompt, Cmd.exe runs the commands sequentially as they appear in the file.

You can include any command in a batch file. Certain commands, such as **for**, **goto**, and **if**, enable you to do conditional processing of the commands in the batch file. For example, the **if** command carries out a command based on the results of a condition. Other commands allow you to control input and output and call other batch files.

The standard error codes that most applications return are 0 if no error occurred and 1 (or higher value) if an error occurred. To determine the meaning of specific error codes, refer to your application help documentation.

For more information about batch file operations, see the following topics:

- Using batch parameters
- Using filters
- Using command redirection operators

For more information about commands that you can use in batch files, click a command:

- Call
- Choice
- Echo
- Endlocal
- For
- Goto
- If
- Pause
- Rem
- Setlocal
- Shift

# Netsh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh**

This section provides information about Netsh commands and their use:

- Netsh overview
- Security information for Netsh commands
- Enter a netsh context
- Netsh DHCP example
- Netsh WINS example
- Netsh Commands

# Netsh overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh overview**

**Netsh** is a command-line scripting utility that allows you to, either locally or remotely, display or modify the network configuration of a computer that is currently running. **Netsh** also provides a scripting feature that allows you to run a group of commands in batch mode against a specified computer. **Netsh** can also save a configuration script in a text file for archival purposes or to help you configure other servers.

**Netsh contexts**

**Netsh** interacts with other operating system components using dynamic-link library (DLL) files. Each Netsh helper DLL provides an extensive set of features called a context, which is a group of commands specific to a networking component. These contexts extend the functionality of **netsh** by providing configuration and monitoring support for one or more services, utilities, or protocols. For example, Dhcpmon.dll provides **netsh** the context and set of commands necessary to configure and manage DHCP servers.

To run a **netsh** command, you must start **netsh** from the Cmd.exe prompt and change to the context that contains the command you want to use. The contexts that are available to you depend on which networking components you have installed. For example, if you type **dhcp** at the Netsh command prompt, you change to the DHCP context, but if you do not have DHCP installed the following message appears:

```
The following command was not found: dhcp.
```

For more information about Netsh contexts, see the following topics:

- Netsh commands for AAAA
- Netsh commands for DHCP
- Netsh diagnostic (diag) commands
- Netsh commands for Interface IP
- Netsh commands for Interface IPv6
- Netsh commands for Interface Portproxy
- Netsh commands for Internet Protocol security
- Netsh commands for Network Bridge
- Netsh commands for remote access (ras)
- Netsh commands for Routing
- Netsh commands for RPC
- Netsh commands for WINS

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Using multiple contexts**
A context can exist within a context. For example, within the Routing context, you can change to the IP and IPX subcontexts.

To display a list of commands and subcontexts that you can use within a context, at the **netsh** prompt, type the context name, and then type either **/?** or **help**. For example, to display a list of subcontexts and commands that you can use in the Routing context, at the **netsh** prompt (that is, netsh>), type either of the following:

**routing /?**

**routing help**

To perform tasks in another context without changing from your current context, type the context path of the command you want to use at the **netsh** prompt. For example, to add the Local Area Connection interface in the IGMP context without changing to the IGMP context, at the **netsh** prompt, type:

**routing ip igmp add interface "Local Area Connection" startupqueryinterval=21**

**Running Netsh commands from the Cmd.exe command prompt**
When you run **netsh** from the Cmd.exe command prompt, it uses the following syntax. To run these commands on a remote Windows 2000 Server, you must first use Remote Desktop Connection to connect to a Windows 2000 Server that is running Terminal Services. There might be functional differences between Netsh context commands on Windows 2000 and the Windows Server 2003 family.

The following describes the netsh command:

**netsh**
Netsh is a command-line scripting utility that allows you to, either locally or remotely, display or modify the network configuration of a currently running computer. Used without parameters, **netsh** opens the Netsh.exe command prompt (that is, netsh>).

**Syntax**
**netsh**[**-a***AliasFile*] [**-c***Context*] [**-r***RemoteComputer*] [{*NetshCommand* | **-f***ScriptFile*}]

**Parameters**
**-a**
Returns you to the **netsh** prompt after running *AliasFile*.

***AliasFile***
Specifies the name of the text file that contains one or more **netsh** commands.

**-c**
Changes to the specified **netsh** context.

***Context***
Specifies the **netsh** context. The following table lists the available **netsh** contexts.

| Context | Description |
|---|---|
| Netsh commands for AAAA | Shows and sets the configuration of the authentication, authorization, accounting, and auditing (AAAA) database used by the Internet Authentication Service (IAS) and the Routing and Remote Access service. |
| Netsh commands for DHCP | Administers DHCP servers and provides an equivalent alternative to console-based management. |
| Netsh diagnostic (diag) commands | Administers and troubleshoots operating system and network service parameters. |
| Netsh commands for Interface IP | Configures the TCP/IP protocol (including addresses, default gateways, DNS servers, and WINS servers) and displays configuration and statistical information. |

| | |
|---|---|
| Netsh commands for Interface IPv6 | Queries and configures IPv6 interfaces, address, caches, and routes. |
| Netsh commands for Interface Portproxy | Administers servers that act as proxies between IPv4 and IPv6 networks and applications. |
| Netsh commands for Interface IPv6 | Provides an equivalent alternative to the console-based management and diagnostic capabilities provided by the IP Security Policy Management and IP Security Monitor snap-ins available in the Microsoft Management Console (MMC). By using the Netsh commands for IPSec, you can configure and view static or dynamic IPSec main mode settings, quick mode settings, rules, and configuration parameters. |
| Netsh commands for Network Bridge | Enables or disables Layer 3 compatibility mode and shows configuration information for the Network Bridge adapters. |
| Netsh commands for remote access (ras) | Administers remote access servers. |
| Netsh commands for Routing | Administers Routing servers. |
| Netsh commands for RPC | Changes, resets or displays selective system binding settings |
| Netsh commands for WINS | Administers WINS servers. |

**-r**
Configures a remote computer.

***RemoteComputer***
Specifies the remote computer to configure.

***NetshCommand***
Specifies the **netsh** command that you want to run.

**-f**
Exits Netsh.exe after running the script.

***ScriptFile***
Specifies the script that you want to run.

**/?**
Displays help at the command prompt.

**Remarks**

- If you specify **-r** followed by another command, **netsh** executes the command on the remote computer and then returns

to the Cmd.exe command prompt. If you specify **-r** without another command, **netsh** opens in remote mode. The process is similar to using **set machine** at the Netsh command prompt. When you use **-r**, you set the target computer for the current instance of **netsh** only. After you exit and reenter **netsh**, the target computer is reset as the local computer. You can run **netsh** commands on a remote computer by specifying a computer name stored in WINS, a UNC name, an Internet name to be resolved by the DNS server, or an IP address.

**Running Netsh commands from the Netsh.exe command prompt**

**Netsh** uses the following standard commands in all contexts that you can run from a Netsh.exe command prompt (that is, netsh>). To run these Netsh commands on a remote Windows 2000 Server, you must first use Remote Desktop Connection to connect to a Windows 2000 Server that is running Terminal server. There might be functional differences between Netsh context commands on Windows 2000 and members of the Windows Server 2003 family.

To view the command syntax, click a command:

- ..
- abort
- add helper
- alias
- bye
- commit
- delete helper
- dump
- exec
- exit
- help
- offline
- online
- popd
- pushd
- quit
- set file
- set machine
- set mode
- show
- unalias

**..**
Moves to the context that is one level up.

**Syntax**
**..**

**Parameters**
**/?**
Displays help at the command prompt.

**abort**
Discards any changes made in offline mode. **Abort** has no effect in online mode.

**Syntax**
**abort**

**Parameters**
**/?**
Displays help at the command prompt.

**add helper**
Installs the helper DLL in **netsh**.

**Syntax**
**add helper** *DLLName*

**Parameters**
***DLLName***
Required. Specifies the name of the helper DLL that you want to install.

**/?**
Displays help at the command prompt.

**alias**
Adds an alias that consists of a user-defined character string, which **netsh** treats as equivalent to another character string. Used without parameters, **alias** displays all available aliases.

**Syntax**
**alias**[*AliasName*] [*String1* [*String2* ...]]

**Parameters**
**alias[ *AliasName*]**
Displays the specified alias.

**alias[ *AliasName*][ *String1*[ *String2*...]]**
Sets *AliasName* to the specified strings.

**/?**
Displays help at the command prompt.

**Examples**
The following **netsh** sample script sets two **netsh** aliases, Shaddr and Shp, and then leaves the Netsh command prompt in the Interface IP context:

```
alias shaddr show interface ip addr
alias shp show helpers
interface ip
```

If you type **shaddr** at the Netsh command prompt, Netsh.exe interprets this as the command **show interface ip addr**. If you type **shp** at the Netsh command prompt, Netsh.exe interprets this as the command **show helpers**.

**bye**
Exits Netsh.exe.

**Syntax**
**bye**

**Parameters**
**/?**

Displays help at the command prompt.

**commit**
Commits any changes made in the offline mode to the router. **Commit** has no effect in online mode.

**Syntax**
**commit**

**Parameters**
**/?**
Displays help at the command prompt.

**delete helper**
Removes the helper DLL from **netsh**.

**Syntax**
**delete helper** *DLLName*

**Parameters**
***DLLName***
Required. Specifies the name of the helper DLL that you want to uninstall.

**/?**
Displays help at the command prompt.

**dump**
Creates a script that contains the current configuration. If you save this script to a file, you can use the file to restore configuration settings that have been changed. Used without parameters, **dump** displays all of the **netsh** context configurations.

**Syntax**
**dump** [*FileName*]

**Parameters**
**[ *FileName*]**
Specifies the name of the file to which you want to redirect output.

**/?**
Displays help at the command prompt.

**exec**
Loads a script file and runs commands from it.

**Syntax**
**exec** *ScriptFile*

**Parameters**
***ScriptFile***
Required. Specifies the name of the script that you want to load and run.

**Remarks**

- The *ScriptFile* can run on one or more computers.

**exit**
Exits Netsh.exe.

**Syntax**
**exit**

**Parameters**
**/?**
Displays help at the command prompt.

**help**
Displays help.

**Syntax**

{**/?** | **?** | **help** | **h**}

**Parameters**
**none**

**offline**
Sets the current mode to offline.

**Syntax**
**offline**

**Parameters**
**/?**
Displays help at the command prompt.

**Remarks**

- Changes that you make in this mode are saved, but you need to run the **commit** or **online** command to set the changes in the router.
- When you switch from offline mode to online mode, changes that you made in offline mode are reflected in the configuration that is currently running.
- Changes that you make in online mode are immediately reflected in the configuration that is currently running.

**online**
Sets the current mode to online.

**Syntax**
**online**

**Parameters**
**/?**
Displays help at the command prompt.

**Remarks**

- Changes that you make in online mode are immediately reflected in the configuration that is currently running.
- When you switch from offline mode to online mode, changes that you made in offline mode are reflected in the configuration that is currently running.

**popd**
Restores a context from the stack.

**Syntax**
**popd**

**Parameters**
**/?**
Displays help at the command prompt.

**Remarks**

- Used in conjunction with **pushd**, **popd** enables you to change the context, run the command in the new context, and then resume the prior context.

**Examples**
The following sample script changes a context from the root context to the **interface ip** context, adds a static IP route, and then returns to the root context:

`netsh>` **pushd**

`netsh>` **interface ip**

```
netsh interface ip> set address local static 10.0.0.9 255.0.0.0 10.0.0.1 1

netsh interface ip> popd

netsh>
```

**pushd**
Saves the current context on a first-in-last-out (FILO) stack.

**Syntax**
**pushd**

**Parameters**
**/?**
Displays help at the command prompt.

**Remarks**

- Used in conjunction with **popd**, **pushd** enables you to change the context, run the command in the new context, and then resume the prior context.

**quit**
Exits Netsh.exe.

**Syntax**
**quit**

**Parameters**
**/?**
Displays help at the command prompt.

**set file**
Copies the Command Prompt window output to a file.

**Syntax**
**set file** {**open***FileName* | **append***FileName* | **close**}

**Parameters**
**open *FileName***
Sends the Command Prompt window output to the specified file.

**append *FileName***
Appends the Command Prompt window output to the specified existing file.

**close**
Stops sending output and closes a file.

**/?**
Displays help at the command prompt.

**Remarks**

- If the specified *FileName* does not currently exist, **netsh** creates a new file with that name. If the specified *FileName* does currently exist, **netsh** overwrites the existing data.

**Examples**
To create a new log file called Session.log and copy all succeeding **netsh** input and output to Session.log, type:

**set file open c:\session.log**

**set machine**
Sets the current computer on which to perform configuration tasks. Used without parameters, **set machine** sets the local computer.

**Syntax**
**set machine** [[*ComputerName*=]*String*]

**Parameters**

***ComputerName***
Specifies the name of the computer on which to perform configuration tasks.

***String***
Specifies the location of the remote computer.

**/?**
Displays help at the command prompt.

**Remarks**

- You can run commands on multiple computers from a single script. You can use **set machine** in a script to specify a destination computer (for example, Computer A), and then run the commands that follow **set machine** on this computer (that is, Computer A). You can then use **set machine** to specify another destination computer (for example, Computer B), and then run commands on this computer (that is, Computer B).

**set mode**
Sets the current mode to online or offline.

**Syntax**
**set mode** {**online** | **offline**}

**Parameters**
**online**
Sets the current mode to online.

**offline**
Sets the current mode to offline.

**/?**
Displays help at the command prompt.

**show**
Displays alias, helper, and mode information.

**Syntax**
**show** {**alias** | **helper** | **mode**}

**Parameters**
**alias**
Lists all defined aliases.

**helper**
Lists all top-level helpers.

**mode**
Displays the current mode.

**/?**
Displays help at the command prompt.

**unalias**
Deletes the specified alias.

**Syntax**
**unalias** *AliasName*

**Parameters**
***AliasName***
Required. Specifies the name of the alias.

**/?**
Displays help at the command prompt.

**Formatting legend**

| Format | Meaning |
|---|---|

| | |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

# Security information for Netsh commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for Netsh commands**

Netsh commands are available for some of the networking technologies, and the security issues for each technology are distinct. You can use the links provided to review security issues, considerations, and best practices for each . In addition, you can review the recommendations in Best practices for security.

| Technology | Context name | Reference |
|---|---|---|
| Dynamic Host Configuration Protocol (DHCP) | **netsh dhcp** | Security information for DHCP |
| Internet Authentication Service (IAS) | **netsh aaaa** | Security information for IAS |
| Internet Protocol security (IPSec) | **netsh ipsec** | Security issues with IP |
| Internet Protocol version 6 (IPv6) | **netsh interface ipv6** and **netsh interface portproxy** | Security information for IPv6 |
| Remote Access | **netsh ras** | Security information for remote access; Security issues for VPN |
| Routing | **netsh routing** | Security information for routing; Security issues for VPN |
| Transmission Control Protocol/Internet Protocol (TCP/IP) | **netsh interface ip** and **netsh bridge** | Internet Protocol Security (IPSec); IPSec Best practices |
| Windows Internet Name Service (WINS) | **netsh wins** | Security information for WINS |

# Enter a netsh context

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enter a netsh context**

1. Open Command Prompt.
2. At the command prompt, type:

   **netsh**

3. Type one of the following:

| To enter this context | Type |
|---|---|
| AAAA | **aaaa** |
| DHCP | **dhcp** |
| Network Diagnostics (diag) | **diag** |
| Interface | **interface** |
| Interface IP | **interface ip** |
| Interface IPv6 | **interface ipv6** |
| Interface Portproxy | **interface portproxy** |
| Internet Protocol security | **ipsec** |
| Network Bridge | **bridge** |
| Remote Access | **ras** |
| Routing | **routing** |
| RPC Helper | **rpc** |
| WINS | **wins** |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To enter a context, you can type only enough letters in the context name to allow **netsh** to identify the context. For example, to enter the WINS context from the **netsh** prompt (that is, netsh>), you can type **w**, and then press ENTER.
- Most of the contexts listed above have one or more subcontexts. For more information, see Related Topics.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Netsh commands for DHCP
Netsh diagnostic (diag) commands
Netsh commands for Interface IP
Netsh commands for Interface IPv6
Netsh commands for Interface Portproxy
Netsh commands for remote access (ras)
Netsh commands for Routing
Netsh commands for WINS

# Netsh DHCP example

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh DHCP example**

You can use Netsh commands for Dynamic Host Configuration Protocol (DHCP) in batch files and other scripts to automate tasks. The following example batch file demonstrates how to use Netsh commands for DHCP to perform a variety of related tasks.

In the circumstance of this example procedure, DHCP-01 is a DHCP server with the IP address 192.168.0.2. The procedure adds a new scope to DHCP-01 with the name MyScope, IP address 192.168.10.0, subnet mask 255.255.255.0, and comment NewScope. It then configures the scope with an address range (192.168.10.1 through 192.168.10.254), an exclusion range (192.168.10.1 through 192.168.10.25), and router IP addresses (DHCP option 003). The scope is then set to an active state.

For more information, see Setting Up Scopes, Setting Up Options, and DHCP options.

For more information and a complete list of Netsh commands for DHCP, see Netsh commands for DHCP.

In the following example procedure, lines that contain comments are preceded by "rem," for remark. Netsh ignores comments.

```
rem one DHCP server:
rem (DHCP-01) 192.168.0.2

rem 1. Connect to (DHCP-01), and add the scope MyScope with IP address 192.168.10.0,
rem 1.1  subnet mask 255.255.255.0, and the comment NewScope.
netsh dhcp server 192.168.0.2 add scope 192.168.10.0 255.255.255.0 MyScope NewScope

rem 2. Connect to (DHCP-01 MyScope), and add IP address range 192.168.10.1 to
192.168.10.254 for distribution
rem 2.1  and the default ClientType of DHCP.
netsh dhcp server 192.168.0.2 scope 192.168.10.0 add iprange 192.168.10.1
192.168.10.254

rem 3. Connect to (DHCP-01 MyScope), and add IP exclusion range 192.168.10.1 to
192.168.10.25
rem 3.1  and the default ClientType of DHCP.
netsh dhcp server 192.168.0.2 scope 192.168.10.0 add excluderange 192.168.10.1
192.168.10.25

rem 4. Connect to (DHCP-01 MyScope), and set the value of option code 003
rem 4.1   to list two router IP addresses (10.1.1.1, 10.1.1.2).
netsh dhcp server 192.168.0.2 scope 192.168.10.0 set optionvalue 003 IPADDRESS 10.1.1.1
10.1.1.2

rem 5.  Connect to (DHCP-01 MyScope), and set the scope state to active.
netsh dhcp server 192.168.0.2 scope 192.168.10.0 set state 1

rem 6. End example batch file.
```

The following table lists the **netsh dhcp** commands that are used in this example procedure.

| Command | Description |
|---|---|
| **server** | Shifts the current DHCP command-line context to the server that is specified by either its name or IP address. |
| **add scope** | Adds a new scope to the specified DHCP server. |
| **scope** | Switches the command context to the DHCP scope that is specified by its IP address. |
| **add iprange** | Adds a range of IP addresses to the current scope. |

| | |
|---|---|
| **add excluderange** | Adds a range of excluded addresses to the current scope. |
| **set optionvalue** | Sets an option value for the current scope. |
| **set state** | Sets or resets the state of the current scope to either an active or inactive state. |

# Netsh WINS example

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh WINS example**

You can use Netsh commands for Windows Internet Name Service (WINS) in batch files and other scripts to automate tasks. The following example batch file demonstrates how to use Netsh commands for WINS to perform a variety of related tasks.

In this example procedure, WINS-A is a WINS server with the IP address 192.168.125.30 and WINS-B is a WINS server with the IP address 192.168.0.189. The procedure adds a dynamic name record with IP address 192.168.0.205, MY_RECORD [04h], to WINS-A, sets WINS-B as a push/pull replication partner of WINS-A, connects to WINS-B, and then sets WINS-A as a push/pull replication partner of WINS-B. Next the procedure initiates a push replication from WINS-A to WINS-B, and then connects to WINS-B to verify that the new record, MY_RECORD, was replicated successfully.

For more information, see WINS Replication, NetBIOS names reference, and Configure Replication.

For a complete list of Netsh commands for WINS, see Netsh commands for WINS.

In the following example procedure, lines that contain comments are preceded by "rem," for remark. Netsh ignores comments.

```
rem two WINS servers:
rem (WINS-A) 192.168.125.30
rem (WINS-B) 192.168.0.189

rem 1. Connect to (WINS-A), and add the dynamic name MY_RECORD [04h] to the (WINS-A)
database.
netsh wins server 192.168.125.30 add name Name=MY_RECORD EndChar=04 IP={192.168.0.205}

rem 2. Connect to (WINS-A), and set (WINS-B) as a push/pull replication partner of
(WINS-A).
netsh wins server 192.168.125.30 add partner Server=192.168.0.189 Type=2

rem 3. Connect to (WINS-B), and set (WINS-A) as a push/pull replication partner of
(WINS-B).
netsh wins server 192.168.0.189 add partner Server=192.168.125.30 Type=2

rem 4. Connect back to (WINS-A), and initiate a push replication to (WINS-B).
netsh wins server 192.168.125.30 init push Server=192.168.0.189 PropReq=0

rem 5. Connect to (WINS-B), and check that the record MY_RECORD [04h] was replicated
successfully.
netsh wins server 192.168.0.189 show name Name=MY_RECORD EndChar=04

rem 6. End example batch file.
```

The following table lists the **netsh wins** commands that are used in this example procedure.

| Command | Description |
|---|---|
| **server** | Shifts the current WINS command-line context to the server specified by either its name or IP address. |
| **add name** | Registers a name to the server. |
| **add partner** | Adds a replication partner to the server. |
| **init push** | Initiates and sends a push trigger to a WINS server. |
| **show name** | Displays detailed information for a particular record in the WINS server database. |

# Netsh Commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands**

This section provides a comprehensive netsh command reference, including syntax, parameters, and examples.

- Netsh commands for AAAA
- Netsh commands for DHCP
- Netsh diagnostic (diag) commands
- Netsh commands for Interface IP
- Netsh commands for Interface IPv6
- Netsh commands for Interface Portproxy
- Netsh commands for Internet Protocol security
- Netsh commands for Network Bridge
- Netsh commands for remote access (ras)
- Netsh commands for Routing
- Netsh commands for RPC
- Netsh commands for WINS

**Notes**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Netsh commands for AAAA

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for AAAA**

You can use commands in the Netsh AAAA context to show and set the configuration of the authentication, authorization, accounting, and auditing (AAAA) database used by the Internet Authentication Service (IAS) and the Routing and Remote Access service. The AAAA database is also known as the IAS database (Ias.mdb). The primary use of commands in the Netsh AAAA context is to:

- Export the configuration of one IAS server, including registry keys and the IAS database (ias.mdb), as a Netsh script using either the **dump** command or one of the **show** commands.
- Import the configuration to another IAS server using the **netsh exec** command, and a Netsh script that contains the **set config** command.

You can run these commands from the Windows Server 2003 family command prompt or from the command prompt for the Netsh AAAA context. For these commands to work at the Windows Server 2003 family command prompt, you must type **netsh aaaa** before typing commands and parameters as they appear in the syntax below. There might be functional differences between Netsh context commands on Windows 2000 and the Windows Server 2003 family.

For more information on Netsh, see Netsh overview and Enter a netsh context.

**Netsh AAAA command reference**

- dump
- set config
- show clients
- show config
- show connection_request_policies
- show logging
- show remote_access_policies
- show server_settings
- show version

**dump**

Displays the configuration of the IAS database file (Ias.mdb) as a Netsh command script.

**Syntax**

**dump**

**Remarks**

- The **dump** command displays the Netsh command script that you can use to duplicate the configuration of the server running IAS or the Routing and Remote Access service on which the command is executed. The Netsh command script contains the configuration of the IAS server, including the registry keys and database file (Ias.mdb), in a compressed text format as a large data block. This large data block is used by the **set config** command within the script to import the configuration of a saved data block into an existing IAS database on the same or another computer using the **netsh exec** command. To save the Netsh command script to a file, type: **netsh aaaa show config >***Path\File.txt*
- The **dump** command is not supported on computers running Windows 2000 server.

**set config**

Configures the IAS server and IAS database (Ias.mdb) with the registry keys and database imported from the specified data block.

**Syntax**

**set config** [**type**={**server_settings** | **clients** | **connection_request_policies** | **logging** | **remote_access_policies**] **blob=***DataBlock*}

**Parameters**

**blob= *DataBlock***

Required. Specifies the file that contains the configuration of the IAS server, including registry keys and IAS database (Ias.mdb), in a compressed text format as output by the **dump** or **show** commands.

**type=**
Specifies the kind of data that you want to import from the *DataBlock* to the local server configuration.

**server_settings**
Specifies that the server configuration settings are imported from the *DataBlock* to the local server. The imported settings include server Description, system event log settings, registry keys, and ports that are used specifically for the IAS service.

**clients**
Specifies that the settings for **RADIUS clients**, as viewed in the IAS console, are imported to the local server from *DataBlock*.

**connection_request_policies**
Specifies that the IAS settings for **Connection Request Processing**, including **Connection Request Policies** and **Remote RADIUS server Groups**, are imported to the local server from *DataBlock*.

**logging**
Specifies that the IAS settings for **Remote Access Logging** are imported to the local server from *DataBlock*.

**remote_access_policies**
Specifies that the IAS settings for **Remote Access Policies** are imported to the local server from *DataBlock*.

**Remarks**

- Running the **set config** command manually is not supported. This command is used only within a Netsh command script that is created by using the **dump** or **show** commands.
- To view the version of the IAS database on which the Netsh script is being run, use the **show version** command. Scripts older than those created with Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition are supported.
- The configuration of IAS servers running on products in the Windows 2000 server family can be imported into products in the Windows Server 2003 family with **set config**. The reverse, however, is not possible.

**show clients**
Displays the RADIUS client list for the local IAS server.

**Syntax**
**show clients**

**Remarks**

- This command dumps the RADIUS client list of the IAS server on which the command is executed. In the IAS console, this client list is displayed in **RADIUS Clients**.
- The Netsh command script contains the local server settings, including registry keys and the IAS database (Ias.mdb), as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh exec** command. To save the Netsh command script to a file, type **netsh aaaa show clients >***Path\File.txt* at a command prompt.
- This command is not available on computers running Windows 2000 Server.

**show config**
Displays the configuration of the IAS database file (Ias.mdb) as a Netsh command script.

**Syntax**
**show config**

**Remarks**

- The **show config** command is equivalent to the **dump** command.
- The **show config** command displays the Netsh command script that you can use to duplicate the configuration of the server running IAS or the Routing and Remote Access service on which the command is run. The Netsh command script contains the configuration of the IAS server, including registry keys and the IAS database (Ias.mdb), as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh exec** command. To save the Netsh command script to a file, type **netsh aaaa show config >***Path\File.txt* at a command prompt.
- This command is not available on computers running Windows 2000 Server.

**show connection_request_policies**
Displays the Connection Request Processing policies configuration of the IAS server in script format.

**Syntax**
**show connection_request_policies**

**Remarks**

- This command displays Connection Request Policies for the IAS server on which the command is run. In the IAS console, these policies are displayed in **Connection Request Processing** and include **Connection Request Policies** and **Remote RADIUS server Groups**.

- The Netsh command script contains the local server settings as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh exec** command. To save the Netsh command script to a file, type **netsh aaaa show connection_request_policies >***Path**File.txt* at a command prompt.

- This command is not available on computers running Windows 2000 Server.

**show logging**
Displays the logging configuration for the local IAS server.

**Syntax**
**show logging**

**Remarks**

- This command displays the Remote Access Logging configuration for the IAS server on which the command is run. In the IAS console, this information is displayed in **Remote Access Logging**.

- The Netsh command script contains the local server settings as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh exec** command. To save the Netsh command script to a file, type **netsh aaaa show logging >***Path**File.txt* at a command prompt.

- This command is not available on computers running Windows 2000 Server.

**show remote_access_policies**
Displays all objects within the remote access policy.

**Syntax**
**show remote_access_policies**

**Remarks**

- This command displays the Remote Access Policies for the IAS server on which the command is run. In the IAS console, this information is displayed in **Remote Access Policies**.

- The Netsh command script contains the local server settings as a large data block in a compressed text format. To import the configuration into an existing IAS server on either the same or another computer, you can use the **netsh exec** command. To save the Netsh command script to a file, type **netsh aaaa show remote_access_policies >***Path**File.txt* at a command prompt.

- This command is not available on computers running Windows 2000 Server.

- The following registry keys and values are displayed:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn\DefaultDomain\REG_SZ**

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy\Allow LM Authentication\REG_DWORD**

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy\Default User Identity\REG_SZ**

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy\User Identity Attribute\REG_DWORD**

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy\Override User-Name\REG_DWORD**

**show server_settings**
Displays the configuration of the local IAS server as a Netsh command script.

**Syntax**
**show server_settings**

**Remarks**

- This command displays the configuration of server settings for the IAS server on which the command is run. These settings include:

  - Server description

  - Settings for accounting and authorization events in the system event log

  - Ports used by the service

- Registry keys and their values

In the IAS console, this information is displayed in server **Properties**.

- The Netsh command script contains the local server settings as a large data block in a compressed text format. This large data block is used within the script by the **netsh exec** command to import the configuration of a saved data block into an existing IAS server on the same or another computer. To save the Netsh command script to a file, type **netsh aaaa show server_settings >***Path\File.txt* at a command prompt.

- This command is not available on computers running Windows 2000 Server.

- The following registry keys and values are displayed:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IAS\Parameters\Allow SNMP Set\REG_DWORD**

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout\MaxDenials\REG_DWORD**

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout\ResetTime\REG_DWORD**

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IAS\Parameters\Ping User-Name\REG_SZ**

**show version**
Displays the version of the IAS database.

**Syntax**
**show version**

**Examples**
To display the Netsh command script that shows the IAS database in its current state, type:

**show config**

To display the version of the IAS database running on a server, type:

**show version**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Netsh commands for DHCP
Netsh diagnostic (diag) commands
Netsh commands for Interface IP
Netsh commands for remote access (ras)
Netsh commands for Routing
Netsh commands for WINS
Command-line reference A-Z
Command shell overview

# Netsh commands for DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for DHCP**
The Netsh commands for DHCP offer a command-line tool that helps with the administration of DHCP servers and provides an equivalent alternative to console-based management. This can be useful in the following situations:

- When managing DHCP servers in wide area networks (WANs), commands can be used in interactive mode at the Netsh command prompt to better manage across slow-speed network links.
- When managing a large number of DHCP servers, commands can be used in batch mode at the Netsh command prompt to help script and automate recurring administrative tasks that need to be performed for all DHCP servers.

You can run these commands from the Windows Server 2003 family command prompt or from the command prompt for the Netsh DHCP context. For these commands to work at the Windows Server 2003 family command prompt, you must type **netsh dhcp** before typing commands and parameters as they appear in the syntax below. There might be functional differences between Netsh context commands on Windows 2000 and the Windows Server 2003 family.

For more information about how to use the Netsh commands that are provided for DHCP, see Use DHCP Command-line Tools.

For an example of how to use Netsh commands for DHCP in batch files and scripts, see Netsh DHCP example.

For more information about Netsh commands, see The Netsh Command-Line Utility.

For more information about **netsh**, see Netsh overview and Enter a netsh context.

- Netsh DHCP
- Netsh DHCP server
- Netsh DHCP server scope
- Netsh DHCP server mscope

**Netsh DHCP**
The following commands are available at the **dhcp>** prompt, which is rooted within the **netsh** environment.

To view the command syntax, click a command:

- add server
- delete server
- server
- show server

**add server**
Adds a DHCP server to the list of authorized servers in Active Directory.

**Syntax**
**add server** *ServerDNS ServerIP*

**Parameters**
***ServerDNS***
Required. Specifies the DHCP server to add. Identifies the server by DNS domain name.

***ServerIP***
Required. Specifies the DHCP server to add. Identifies the server by IP address.

**Examples**
In the following example, this command adds a server with the DNS name dhcpsrv1.example.microsoft.com and IP address 10.2.2.2 to the list of authorized servers in Active Directory.

**add server dhcpsrv1.example.microsoft.com 10.2.2.2**

**delete server**
Deletes a DHCP server from the list of authorized servers in Active Directory.

**Syntax**
**delete server** *ServerDNS ServerIP*

**Parameters**
***ServerDNS***
Required. Specifies the DHCP server to delete. Identifies the server by DNS domain name.

***ServerIP***
Required. Specifies the DHCP server to delete. Identifies the server by IP address.

**Examples**
In the following example, this command deletes the server with the DNS name dhcpsrv1.example.microsoft.com and IP address 10.2.2.2 from the list of authorized servers in Active Directory.

**delete server dhcpsrv1.example.microsoft.com 10.2.2.2**

**server**
Shifts the current Netsh DHCP command-line context to a different DHCP server. Used without parameters, **server** shifts the current command-line context to the local computer.

**Syntax**
**server** [{*****ServerNetBIOSName* | *****ServerFQDN* | *ServerIP*}]

**Parameters**
**{ *ServerNetBIOSName*| *ServerFQDN*| *ServerIP*}**
Specifies the DHCP server to which you want to shift the current command-line context. Identifies the server by NetBIOS name, fully qualified domain name (FQDN), or IP address.

**Examples**
In the first example, this command changes the Netsh DHCP command-line context to the DHCP server with the NetBIOS name \\DHCP-SRV1.

In the second example, this command changes the Netsh DHCP command-line context to the DHCP server with IP address 10.0.0.1.

In the third example, this command changes the Netsh DHCP command-line context to the DHCP server with the FQDN \\example.microsoft.com.

**server \\DHCP-SRV1**

**server 10.0.0.1**

**server \\example.microsoft.com**

**show server**
Displays a list of authorized servers in Active Directory.

**Syntax**
**show server**

**Parameters**
none

**Netsh DHCP server**
The following commands are available at the **dhcp server>** prompt, which is rooted within the **netsh** environment.

To view the command syntax, click a command:

- add class
- add mscope
- add optiondef
- add scope
- delete class
- delete dnscredentials
- delete mscope
- delete optiondef
- delete optionvalue
- delete scope
- delete superscope
- dump
- export
- import
- initiate auth
- mscope
- scope
- set auditlog
- set databasebackupinterval
- set databasebackuppath
- set databasecleanupinterval
- set databaseloggingflag
- set databasename
- set databasepath
- set databaserestoreflag
- set detectconflictretry
- set dnsconfig
- set dnscredentials

- set optionvalue
- set server
- set userclass
- set vendorclass
- show all
- show auditlog
- show bindings
- show class
- show detectconflictretry
- show dnsconfig
- show dnscredentials
- show mibinfo
- show mscope
- show optiondef
- show optionvalue
- show scope
- show server
- show dbproperties
- show serverstatus
- show userclass
- show vendorclass
- show version

**add class**
Adds a class to the specified DHCP server.

**Syntax**
add class *ClassName* [*ClassComment*] [*Data*] [[**IsVendor=**]{**0** | **1**}] [[**IsBinary=**]{**B** | **b**}]

**Parameters**
***ClassName***
Required. Specifies the name of the class to create.

***ClassComment***
Specifies the comment to associate with the class.

***Data***
Specifies class data in ASCII or binary format depending on the **IsBinary** parameter. If unspecified, it is assumed that the format of the data is ASCII.

**[ IsVendor=]{0 | 1}**
Indicates whether the class to add is a user class or a vendor class: **0**-User class (default), **1**-Vendor class.

**[ IsBinary=]{B | b}**
Indicates if the class data is in binary format. If unspecified, it is assumed that the format of the data is ASCII.

**Examples**
In the first example, this command creates a vendor class named MyClass, with the comment TestClass, and the class data TestDataForMyClass in ASCII format.

In the second example, this command creates a user class named MyUserClass with binary data of 4345335532 (ASCII-CE3U2).

The third example is an alternate command syntax for the second example.

**add class MyClass TestClass TestDataForMyClass 1**

**add class MyUserClass TestClass1 4345335532 B**

**add class MyUserClass TestClass1 4345335532 0 b**

**add mscope**
Adds a multicast scope to the specified DHCP server.

**Syntax**
**add mscope** *MscopeName* [*MscopeComment*] [*TTL*]

**Parameters**
***MscopeName***
Required. Specifies the name of the multicast scope to add.

***MscopeComment***
Specifies a description of the current multicast scope.

***TTL***
Specifies the Time-To-Live value for the multicast scope. The valid range for this field is 1 to 255, with a default of 32.

**Examples**
In the following example, this command creates a multicast scope with the name My MultiCast Scope, a comment of MyTestMulticastScope, and a multicast TTL value of 40.

**add mscope "My MultiCast Scope" MyTestMulticastScope 40**

**add optiondef**
Adds a new option type definition to the specified DHCP server.

**Syntax**
**add optiondef** *OptCode OptName*{**BYTE** | **WORD** | **DWORD** | **STRING** | **IPADDRESS**} [[**IsArray=**]{**0** | **1**}] [**vendor=***VendorClass*] [**comment=***OptComment*] [*DefValue*]

**Parameters**
***OptCode***
Required. Specifies the option type to define. Identifies the option type with a unique option code. For the value of the code, choose a number between 0 and 255.

***OptName***
Required. Specifies the name of the option type.

**{BYTE | WORD | DWORD | STRING| IPADDRESS}**
Required. Indicates the data type of the option.

**[IsArray=]{0 | 1}**
Indicates whether the data type is arrayed or non-arrayed: **0**-Unary or non-arrayed (default), **1**-Arrayed.

**vendor= *VendorClass***

Specifies the vendor class for the option. If unspecified, the default vendor class, the DHCP standard options class, is used.

**comment=** ***OptComment***
Specifies a comment to add to the new option type definition. If unspecified, the default value is NULL.

***DefValue***
Specifies a default value to use for this option type, if one is not specified otherwise.

**Remarks**

- **IsArray** must be **0** for STRING type options.

**Examples**
In the first example, this command defines a new option type named ExtensionsPath with code 18, a STRING data type, and joins it to the DHCP standard options class, assigning it a default value of c:\Temp.

In the second and third examples, this command defines a new option type named TestOption with code 231 and an IPADDRESS data type. In the second example, the option is defined to provide only a single IP address. IsArray and *DefValue* are not necessary because defaults are assumed. In the third example, *DefValue*, and IsArray are required to create an array containing the IP address values 10.1.1.1, 10.1.1.2, and 10.1.1.3.

**add optiondef 18 ExtensionsPath STRING 0 c:\Temp**

**add optiondef 231 TestOption IPADDRESS**

**add optiondef 231 TestOption IPADDRESS 1 vendor=MyVendorClass 10.1.1.1 10.1.1.2 10.1.1.3**

**add scope**
Adds a scope to the specified DHCP server.

**Syntax**
**add scope** *ScopeAddress SubnetMask ScopeName*[*ScopeComment*]

**Parameters**
***ScopeAddress***
Required. Specifies the scope to add. Identifies the scope by IP network number.

***SubnetMask***
Required. Specifies the subnet mask for the scope to add.

***ScopeName***
Required. Specifies an identifying name for the scope.

***ScopeComment***
Specifies a comment for the scope.

**Examples**
In the following example, this command adds a scope for the DHCP server with scope address 10.2.2.0, subnet mask 255.255.255.0, a scope name of MyScope, and a scope comment of MyComment.

**add scope 10.2.2.0 255.255.255.0 MyScope MyComment**

**delete class**
Deletes a class from the specified DHCP server.

**Syntax**
delete class *Name*

**Parameters**
***Name***
Required. Specifies the user or vendor class to delete.

**Examples**
In the following example, this command deletes the class named MyClass.

**delete class MyClass**

**delete dnscredentials**

Deletes the credentials used for all DNS dynamic updates.

**Syntax**

**delete dnscredentials dhcpforceflag**

**Parameters**

none

**delete mscope**

Deletes a multicast scope from the specified DHCP server.

**Syntax**

**delete mscope** *MscopeName*

**Parameters**

***MscopeName***

Required. Specifies the multicast scope to delete. *MscopeName* is case-sensitive.

**Examples**

In the following example, this command deletes the multicast scope MyMulticastScope from the currently specified DHCP server.

**delete mscope MyMulticastScope**

**delete optiondef**

Deletes a defined option type from the specified DHCP server.

**Syntax**

**delete optiondef** *OptCode* [[**vendor=**]*VendorName*]

**Parameters**

***OptCode***

Required. Specifies the unique identifier for the option type to delete.

**vendor= *VendorName***

Specifies the vendor class name with which the option type is associated. If unspecified, the option type definition is deleted from the non-vendor specific DHCP standard options class.

**Examples**

In the following example, this command deletes an option type with code 18 from the DHCP standard options class.

**delete optiondef 18**

**delete optionvalue**

Deletes the current value of the specified option type for the specified DHCP server.

**Syntax**

**delete optionvalue** *OptCode*[[**vendor=**]*VendorClass*] [[**user=**]*UserClass*]

**Parameters**

***OptCode***

Required. Specifies the unique identifier of the global server option with a value to delete.

**vendor= *VendorClass***

Specifies the vendor class name with which the option is associated. If the tag is not included, the global *VendorName* that was set by **set vendorclass** is assumed. If the tag is used, but no value is specified, then no vendor class is used. This parameter is available only for computers running Windows 2000 Server.

**user= *UserClass***

Specifies the user class name with which the option is associated. If the tag is not included, the global *UserName* that was set by **set userclass** is assumed. If the tag is used, but no value is specified, then no user class is used. This parameter is available only for computers running Windows 2000 Server.

**Examples**

In the following example, this command deletes the current value of the option type with code 18.

**delete optionvalue 18**

**delete scope**

Deletes a scope from the specified DHCP server.

**Syntax**
**delete scope** *ScopeAddress* {**DHCPNOFORCE** | **DHCPFULLFORCE**}

**Parameters**
***ScopeAddress***
Required. Specifies the address of the scope to delete.

**{DHCPNOFORCE | DHCPFULLFORCE}**
Required. Indicates the force with which the scope will be deleted: **DHCPNOFORCE**-Deletes the scope if there are no active clients in the scope, **DHCPFULLFORCE**-Deletes the scope even if there are active clients in the scope.

**Examples**
In the following example, this command forces the deletion of the scope 10.2.2.0 from the DHCP server and ignores all warnings.

**delete scope 10.2.2.0 dhcpfullforce**

**delete superscope**
Deletes a superscope from the specified DHCP server.

**Syntax**
**delete superscope** *SuperscopeName*

**Parameters**
***SuperscopeName***
Required. Specifies the name of the superscope to delete.

**Remarks**

- You can set the superscope from the scope context with the **set superscope** command.
- The superscope name is case-sensitive.

**Examples**
In the following example, this command deletes the superscope named My Superscope from the specified DHCP server.

**delete superscope "My Superscope"**

**dump**
Dumps the configuration of the local DHCP server to the command prompt window when run within the netsh environment.

**Syntax**
**dump**

When run at the command prompt or in a batch file, output can be saved in a text file. Used without parameters, this command dumps the configuration of the local server to the command prompt window.

**Syntax**
**netsh dhcp server [ *IPAddress* ] dump >**[*PathAndFileName*]

**Parameters**
***IPAddress***
Specifies the IP address of the DHCP server for which the configuration is output.

***PathAndFileName***
Specifies both the location where the file is saved, and the name of the destination file to which to dump DHCP server configuration. If unspecified, the DHCP server configuration is dumped to the command prompt window.

**Remarks**

- This command dumps the DHCP configuration as text output within the command prompt window. You can copy or redirect this output to a file to view it in a text editor. After file output is obtained, you can use the **exec** command to configure another DHCP server with the same configuration.
- Before performing the **dump** command at the source server, reconcile all scopes and fix any reported inconsistencies.

- The following command dumps the current configuration for the local DHCP server (the source) to a text file named Dhcpcfg.dmp. This file can be copied to another (destination) DHCP server and used to recreate the source configuration on it. You can recreate the configuration with the following command:

  **dump > dhcpcfg.dmp**

- Before you use the **exec** command to execute a dump file at the destination server, complete the following steps:

  1. Close the DHCP console at the destination server. Do not reopen the console until after you have run the **exec** command.

  2. Delete all default option definitions and any user and vendor classes at the destination server.

- These examples demonstrate how to use the **delete optiondef** and **del class** Netsh commands to delete all default option definitions and any user and vendor classes. You can run these commands at the **dhcp server>** prompt, from within the **netsh** environment.

  - **delete optiondef 76**

  - **del class "Default BOOTP Class"**

  - **del class "Microsoft Windows 98 Options"**

- On the destination server, the **exec** command is used to load and execute the saved configuration:

  **netsh exec dhcpcfg.dmp**

- After you use the **exec** command to load the dump file, you must reconcile all scopes. Use **net stop dhcpserver** to stop the DHCP Server service and **net start dhcpserver** to restart it. Once the service is restarted, DHCP database changes take effect.

**Examples**

The first command, which is run within the netsh environment, dumps the configuration of the local DHCP server to the command prompt window.

The second command, which is run at the command prompt, dumps the configuration of a DHCP server with IP address 192.168.0.1 to a text file named Dhcpcfg.dmp at the location C:\Dhcp\ on the local computer.

The third command, which is run at the command prompt, dumps the configuration of a DHCP server with IP address 192.168.0.1 to a text file named Dhcpcfg.dmp on a shared network folder named \\Backup\Dhcp\.

**dump**

**netsh dhcp server 192.168.0.1 dump > C:\Dhcp\Dhcpcfg.dmp**

**netsh dhcp server 192.168.0.1 dump > \\Backup\Dhcp\Dhcpcfg.dmp**

**export**

Exports the DHCP service configuration to a file.

**Syntax**

**export** [*Path*]*FileName* {**all** | *ScopeList*}

**Parameters**

**[*Path*] *FileName***

Required. Specifies, by name, the file where the DHCP configuration will be stored. If the path, the file name, or both contain spaces, quotation marks must be used.

**{all | *ScopeList*}**

Required. Specifies which scopes you want to export. The parameter **all** exports all scopes. The parameter *ScopeList* exports the scopes that correspond to the IP addresses you list. Each IP address in the list must be separated by spaces.

**Remarks**

- This command works only on the local server.
- While the **export** command runs, the DHCP service is stopped and does not respond to DHCP clients seeking new leases or lease renewals.
- If the DHCP service has a large number of scopes or a large number of client address leases, this command can take a long time to run.

**Examples**
In the first example, this command exports the complete DHCP service configuration to the file c:\Temp\Dhcpdb.

In the second example, this command exports the DHCP configuration for scopes 10.0.0.0 and 192.168.0.0 to the file c:\Temp\Dhcpdb.

In the third example, this command exports the complete DHCP service configuration to the file c:\My Folder\Dhcp Configuration. Note that both the path and file name contain spaces, so quotation marks are used.

**export c:\Temp\Dhcpdb all**

**export c:\Temp\Dhcpdb 10.0.0.0 192.168.0.0**

**export "c:\My Folder\Dhcp Configuration" all**

**import**
Imports a DHCP service configuration from a file to the local service.

**Syntax**
**import** [*Path*]*FileName* {**all** | *ScopeList*]

**Parameters**
**[*Path*] *FileName***
Required. Specifies, by name, the file from which the DHCP configuration will be imported. If the path, the file name, or both contain spaces, quotation marks must be used.

**{all | *ScopeList*}**
Required. Specifies which scopes you want to import. The parameter **all** imports all scopes represented in the file you specify. The parameter *ScopeList* imports the scopes that correspond to the IP addresses you list. Each IP address in the list must be separated by spaces.

**Remarks**

- This command works only on the local server.
- While the **import** command runs, the DHCP service is stopped and does not respond to DHCP clients seeking new leases or lease renewals.
- If the DHCP service has a large number of scopes or a large number of client address leases, this command can take a long time to run.

**Examples**
In the first example, this command imports the complete DHCP service configuration from the file c:\Temp\Dhcpdb.

In the second example, this command imports the DHCP configuration for scopes 10.0.0.0 and 192.168.0.0 from the file c:\Temp\Dhcpdb

In the third example, this command imports the complete DHCP service configuration from the file c:\My Folder\Dhcp Configuration. Note that both the path and file name contain spaces, so quotation marks are used.

**import c:\Temp\Dhcpdb all**

**import c:\Temp\Dhcpdb 10.0.0.0 192.168.0.0**

**import "c:\My Folder\Dhcp Configuration" all**

**initiate auth**
Initiates authorization of the specified DHCP server in Active Directory.

**Syntax**
**initiate auth**

**Parameters**
none

**mscope**
Switches the command context to a DHCP multicast scope.

**Syntax**
**mscope** *MscopeName*

**Parameters**
***MscopeName***
Required. Specifies the name of the multicast scope to which to switch the command context.

**Remarks**

- This command enables you to use any of the available DHCP multicast scope-level commands on the specified DHCP multicast scope. Subsequent operations will be performed on this multicast scope at the specified server.
- The name of the multicast scope is case-sensitive.

**Examples**
In the following example, this command switches the command context to a multicast scope named My Multicast Scope.

**mscope "My Multicast Scope"**

**scope**
Switches the command context to a DHCP scope.

**Syntax**
**scope** *ScopeIPAddress*

**Parameters**
***ScopeIPAddress***
Required. Specifies the scope to which to switch the command context. Identifies the scope by network or subnet IP address.

**Remarks**

- This command enables you to use any of the available DHCP scope-level commands on the specified DHCP scope. Subsequent operations will be performed on this scope at the specified server.

**Examples**
In the following example, this command switches the command context to the scope with IP address 10.0.0.0.

**scope 10.0.0.0**

**set auditlog**
Sets the audit log path parameters for the specified DHCP server.

**Syntax**
**set auditlog** *NewAuditPath*

**Parameters**
***NewAuditPath***
Required. Specifies a local folder for the DHCP server to use when storing audit log files. A valid folder path is required, and the folder must be local to the DHCP server where the path is defined.

**Remarks**

- In order for this change to take effect, you need to stop and restart the DHCP service after you run this command.

**Examples**
In the following example, this command sets the DHCP server to store audit log files in the directory c:\logpath.

**set auditlog c:\logpath**

**set databasebackupinterval**
Sets the backup interval for the specified DHCP server.

**Syntax**
**set databasebackupinterval** *NewInterval*

**Parameters**
***NewInterval***
Required. Specifies the backup interval in minutes.

**Examples**
In the following example, this command sets the database backup interval for the specified DHCP server to 1,440 minutes (24 hours).

**set databasebackupinterval 1440**

**set databasebackuppath**
Sets the database backup path for the specified DHCP server.

**Syntax**
**set databasebackuppath** *NewBackupPath*

**Parameters**
***NewBackupPath***
Required. Specifies a local folder for the DHCP server to use when storing audit log files. A valid folder path is required, and the folder must be local to the DHCP server where the path is defined.

**Examples**
In the following example, this command sets the database backup path to c:\windows\dhcp\backup.

**set databasebackuppath c:\windows\dhcp\backup**

**set databasecleanupinterval**
Sets the database cleanup interval of the specified DHCP server.

**Syntax**
**set databasecleanupinterval** *NewInterval*

**Parameters**
***NewInterval***
Required. Specifies the database cleanup interval, in minutes.

**Examples**
In the following example, this command sets the database cleanup interval to 10,080 minutes (every seven days).

**set databasecleanupinterval 10080**

**set databaseloggingflag**
Sets or resets the database logging flag for the specified DHCP server. Used without parameters, this command resets the database logging flag.

**Syntax**
**set databaseloggingflag** [{**0** | **1**}]

**Parameters**
**{0 | 1}**
Indicates the database logging flag setting: **0**-Resets the flag (default), **1**-Sets the flag.

**Examples**
In the following example, this command resets the database logging flag for the currently specified DHCP server.

**set databaseloggingflag 0**

**set databasename**
Sets the name of the DHCP server database file for the specified DHCP server.

**Syntax**
**set databasename** *NewFileName*

**Parameters**
***NewFileName***
Required. Specifies a new name for the DHCP server database file.

**Remarks**

- In order for this change to take effect, you need to stop and restart the DHCP Server service.

**Examples**
In the following example, this command sets the DHCP server database file name to newdatabase.mdb.

**set databasename newdatabase.mdb**

**set databasepath**
Sets the path of the DHCP server database file for the specified DHCP server.

**Syntax**
**set databasepath** *NewPath*

**Parameters**
***NewPath***
Required. Specifies the path of the DHCP database file.

**Remarks**

- In order for this change to take effect, you need to stop and restart the DHCP Server service.

**Examples**
In the following example, this command sets the path of the DHCP server database file to c:\windows\dhcp.

**set databasepath c:\windows\dhcp**

**set databaserestoreflag**
Sets or resets the database restore flag for the specified DHCP server. Used without parameters, this command resets the database restore flag.

**Syntax**
**set databaserestoreflag** [{**0** | **1**}]

**Parameters**
**{0 | 1}**
Indicates the setting of the database restore flag: **0**-Resets the flag (default), **1**-Sets the flag.

**Examples**
In the following example, this command sets the database restore flag for the currently specified DHCP server.

**set databaserestoreflag 1**

**set detectconflictretry**
Sets the number of conflict detection attempts for the specified DHCP server.

**Syntax**
**set detectconflictretry** *RetryNumber*

**Parameters**
***RetryNumber***
Required. Specifies the number of conflict detection attempts that the DHCP server will make before it leases a scope IP address to a client.

**Remarks**

- The maximum value for *RetryNumber* is five. All values greater than five will default down to five.

**Examples**
In the following example, this command sets the DHCP server to make four attempts to detect if a there is a conflict for a scope IP address before the DHCP server uses this scope IP address in a new address lease.

**set detectconflictretry 4**

**set dnsconfig**
Sets the DNS dynamic update configuration for the specified DHCP server.

**Syntax**
**set dnsconfig** [**Enable=**]{**0** | **1**} [[**Update=**]{**0** | **1**}] [[**Lookup=**]{**0** | **1**}] [[**NonDyn=**]{**0** | **1**}]

**Parameters**
**[Enable=]{0 | 1}**
Required. Indicates whether to enable or disable the dynamic updates of DNS client information: **0**-Disables dynamic updates of DNS client information, **1**-Enables dynamic updates of DNS client information.

**[Update=]{0 | 1}**
Indicates the type of update lookups: **0**-Updates according to client request, **1**-Updates forward and reverse name lookups.

**[Lookup=]{0 | 1}**
Indicates whether to enable or disable forward lookups for clients when leases expire: **0**-Disables forward lookups, **1**-Enables forward lookups.

**[NonDyn=]{0 | 1}**
Indicates whether to update DNS information for clients that do not support dynamic updates: **0**-Disables updates for clients that do not support dynamic updates, **1**-Enables these updates.

**Remarks**

- This command is only supported for DHCP servers running Windows 2000 Server and members of the Windows Server™ 2003 family.
- The configured Preferred DNS server, in the TCP/IP properties of the DHCP server, will be updated with DHCP client records when dynamic updates are enabled. If the Preferred DNS server is not available on the network, the Alternate DNS server will be updated.

**Examples**
In the following example, this command enables the dynamic update of DNS information by the DHCP server and configures updates to always be made for both forward and reverse lookups. This example also disables the option to perform forward lookups when leases expire and enables updates for non-dynamic clients.

**set dnsconfig 1 1 0 1**

**set dnscredentials**
Sets DNS dynamic update credentials for the specified DHCP server.

**Syntax**
**set dnscredentials** *UserName Domain Password*

**Parameters**
***UserName***
Required. The user name associated with a valid domain account.

***Domain***
Required. The domain in which the user account is located.

***Password***
Required. The password for the user account.

**Examples**
In the first example, this command sets the DNS dynamic update credentials to the values specified.

In the second example, this command sets the value of *UserName* to User1 and the value of *Domain* to Domain1, but the password is not specified. Instead an asterisk (*) is used. When an asterisk (*) is used in place of a valid password, the user is

prompted for a password. Without the correct password for the specified *UserName*, the credentials cannot be changed.

**set dnscredentials User1 Domain1 Password1**

**set dnscredentials User1 Domain1 ***

#### set optionvalue

Sets a DHCP server option value that will be applied for all scopes defined at the specified DHCP server.

**Syntax**

**set optionvalue** *OptCode* {**BYTE** | **WORD** | **DWORD** | **STRING** | **IPADDRESS**} [[**user=**]*UserName*] [[**vendor=**]*VendorName*] [*OptionValue*]

**Parameters**

***OptCode***
Required. Specifies the unique identifier for the option type with the value to set.

**{BYTE | WORD | DWORD | STRING | IPADDRESS}**
Required. Specifies the data type for the option type with the value to set.

**[user=]*UserName***
Sets the applicable user class name to which this command will apply. If unspecified, the default user class is assumed.

**[vendor=]*VendorName***
Sets the applicable vendor class name to which this command will apply. If unspecified, the DHCP standard options class is assumed.

**[*OptionValue*]**
Specifies the new server default value for the option type identified by *OptCode*. The value must be of the associated data type.

**Remarks**

- Vendor and user classes are supported only for DHCP servers running Windows 2000 Server.
- To modify the current defaults for an unspecified class, use either **set userclass** or **set vendorclass** commands.

**Examples**

In the first example, this command sets the value of option code 003 to list two router IP addresses, 10.1.1.1 and 10.1.1.2.

In the second example, this command further qualifies the first example and sets the specified value to apply only for those scope clients that identify themselves as members of the vendor class Vendor1. In order for this example to work, this class must already be defined at the server, with the specified option type defined for its use.

**set optionvalue 003 IPADDRESS 10.1.1.1 10.1.1.2**

**set optionvalue 003 IPADDRESS vendor=Vendor1 10.1.1.1 10.1.1.2**

#### set server

Changes the Netsh command-line context to the specified server. Used without parameters, the local server is assumed.

**Syntax**

**set server**[{*ServerIP* | *****ServerDNS*}]

**Parameters**

**{*ServerIP* | *ServerDNS*}**
Specifies the IP address or the DNS name of the specified DHCP server.

**Remarks**

- When this command is used, and the Netsh command-line context is changed to another server, all subsequent command operations are performed on the specified server. To perform command operations on another server, this command must be executed again, with the IP address or DNS name of the new server provided as the parameter.

**Examples**

In the following example, this command switches the DHCP server command-line context to the server at IP address 10.1.1.1

**set server 10.1.1.1**

**set userclass**
Sets the name for the current user class. Used without parameters, the current user class is reset to the default user class

**Syntax**
**set userclass** [*UserClass*]

**Parameters**
***UserClass***
Specifies the name of the user class for the currently specified DHCP server.

**Remarks**

- This command is available only for use with DHCP servers running Windows 2000 Server.

**Examples**
In the first example, this command sets the user class name to MyUserClass.

In the second example, this command resets the current user class to the default user class.

**set userclass MyUserClass**

**set userclass**

**set vendorclass**
Sets the name for the current vendor class. Used without parameters, the current vendor class is reset to the default class, DHCP standard options.

**Syntax**
**set vendorclass** [*VendorClass*]

**Parameters**
***VendorClass***
Specifies the name of the vendor class for the current server.

**Remarks**

- This command is available only for use with DHCP servers running Windows 2000 Server.

**Examples**
In the first example, this command sets the vendor class name to MyVendorClass.

In the second example, this command resets the current vendor class to the default class, DHCP standard options.

**set vendorclass MyVendorClass**

**set vendorclass**

**show all**
Displays all status and configuration information for the specified DHCP server.

**Syntax**
**show all**

**Parameters**
**none**

**Remarks**

- This command lists all current server status and configuration details, including the stored management information base (MIB) information, for the currently specified DHCP server.

**show auditlog**
Displays all audit log information for the specified DHCP server.

**Syntax**
**show auditlog**

**Parameters**
**none**

**show bindings**
Displays bindings information for the specified DHCP server.

**Syntax**
**show bindings**

**Parameters**
none

**Remarks**

- This command displays bindings information for the current DHCP server, indicating which active network connections that are configured with static IP addresses are enabled or disabled for use in servicing DHCP clients on your network.
- The DHCP service automatically disables network connections that obtain their IP address configuration dynamically from service bindings. This happens by default.

**show class**
Enumerates and displays all class information for the specified DHCP server. This includes both user and vendor class information.

**Syntax**
**show class**

**Parameters**
**none**

**show detectconflictretry**
Displays the configured number of conflict detection attempts for the specified DHCP server.

**Syntax**
**show detectconflictretry**

**Parameters**
**none**

**Remarks**

- This command displays the current number of ping retries that are used to attempt the detection of address conflicts for the scope IP addresses that are distributed by the server.

**show dnsconfig**
Displays the DNS dynamic update configuration for the specified DHCP server.

**Syntax**
**show dnsconfig**

**Parameters**
**none**

**Remarks**

- If DNS configuration has not been enabled or set, you can use the DNS console or the **set dnsconfig** command to configure these settings.

**show dnscredentials**
Displays the current DNS dynamic update credentials.

**Syntax**
**show dnscredentials**

**Parameters**
**none**

**show mibinfo**
Displays management information base (MIB) information for the specified DHCP server.

**Syntax**
**show mibinfo**

**Parameters**
none

**show mscope**
Displays all information about multicast scopes for the specified DHCP server.

**Syntax**
**show mscope**

**Parameters**
**none**

**show optiondef**
Displays all defined and available options types for use at the specified DHCP server. Used without parameters, all options are displayed.

**Syntax**
**show optiondef** [**vendor=***VendorName*]

**Parameters**
**vendor= *VendorName***
Indicates which defined and available option types to display. If unspecified, the default for *VendorName* is the default vendor class that is currently set for the DHCP server.

**Remarks**

- The default vendor class can be reset with the **set vendorclass** command.
- *VendorName* is applicable only for DHCP servers running Windows 2000 Server.

**Examples**
In the first example, this command displays all defined option types for the current vendor class.

In the second example, this command displays the defined option types that are available for use with the specified class Vendor1.

**show optiondef**

**show optiondef vendor=Vendor1**

**show optionvalue**
Displays all of the available option values that are currently set for the specified DHCP server. Used without parameters, all option values are displayed, including those that are used for both user and vendor classes.

**Syntax**
**show optionvalue**[{**user=***UserName* | **vendor=***VendorName*}]

**Parameters**
**user= *UserName***
Specifies the default values that are set for the option types that are available for use with the currently set user class.

**vendor=** ***VendorName***

Specifies the default values that are set for the option types that are available for use with the currently set vendor class.

**Remarks**

- **User=** *UserName* and **vendor=***VendorName* are only available for use with DHCP servers running Windows 2000 Server.

- If you provide values for *UserName* or *VendorName*, the respective tags (**user=** or **vendor=**) are required. If the tags are not included, the default that is assumed is the current user class, previously set with the **set userclass** command.

**Examples**

In the following example, this command displays all of the available option types that have values set at the currently specified DHCP server.

**show optionvalue**

**show scope**

Displays information about the scopes for the specified DHCP server.

**Syntax**

**show scope**

**Parameters**

**none**

**show server**

Displays information about the specified DHCP server, including the fully qualified domain name and IP address of the server.

**Syntax**

**show server**

**Parameters**

**none**

**show dbproperties**

Displays information about server database configuration for the specified DHCP server.

**Syntax**

**show dbproperties**

**Parameters**

**none**

**Remarks**

- This command displays the following database configuration information:

    - DatabaseName

    - DatabasePath

    - DatabaseBackupPath

    - DatabaseBackupInterval

    - DatabaseLoggingFlag

    - DatabaseRestoreFlag

    - DatabaseCleanupInterval

**show serverstatus**
Displays status information for the specified DHCP server.

**Syntax**
**show serverstatus**

**Parameters**
**none**

**Remarks**

- This is an example of the type of status information that this command displays:

```
Server Attrib - Server Servicing Clients :TRUE

Server Attrib - Dynamic BootP Support Enabled :TRUE

Server Attrib - DHCP Server Part Of DS :TRUE

Server Attrib - DHCP Server Bindings Aware :TRUE

Server Attrib - Administrative Rights :TRUE
```

**show userclass**
Displays the current user class setting at the specified DHCP server.

**Syntax**
**show userclass**

**Parameters**
**none**

**Remarks**

- This is an example of the type of setting information that this command displays:

```
Current Class Name set for the Server dhcpsrv1.example.microsoft.com is None.
```

**show vendorclass**
Displays the current vendor class setting at the specified DHCP server.

**Syntax**
**show vendorclass**

**Parameters**
**none**

**Remarks**

- This is an example of the type of setting information that this command displays:

```
Current Vendor Name set for the Server dhcpsrv1.example.microsoft.com is None.
```

**show version**
Displays current version information for the specified DHCP server.

**Syntax**
**show version**

**Parameters**

**none**

**Remarks**

- This is an example of the type of version information that this command displays:

```
The version of the DHCP Server 192.168.0.100 is 5.6.
```

**Netsh DHCP server scope**

The following commands are available at the **dhcp server scope>** prompt, which is rooted within the **netsh** environment.

To view the command syntax, click a command:

- add excluderange
- add iprange
- add reservedip
- delete excluderange
- delete iprange
- delete lease
- delete optionvalue
- delete reservedip
- delete reservedoptionvalue
- dump
- initiate reconcile
- set comment
- set name
- set optionvalue
- set reservedoptionvalue
- set scope
- set state
- set superscope
- show clients
- show clientsv5
- show excluderange

- show iprange
- show optionvalue
- show reservedip
- show reservedoptionvalue
- show scope
- show state

**add excluderange**
Adds a range of addresses to exclude from distribution in the current scope.

**Syntax**
**add excluderange** *StartIP EndIP*

**Parameters**
***StartIP***
Required. Specifies the IP address that starts the exclusion range.

***EndIP***
Required. Specifies the IP address that ends the exclusion range.

**Examples**
In the following example, this command excludes the IP addresses in the range 10.2.2.10 to 10.2.2.20 from distribution in the scope.

**add excluderange 10.2.2.10 10.2.2.20**

**add iprange**
Adds a range of IP addresses to the current scope.

**Syntax**
**add iprange** *StartIP EndIP* [{**DHCP** | **BOOTP** | **BOTH**}] [*MaxBootP*]

**Parameters**
***StartIP***
Required. Specifies the IP address that starts the range.

***EndIP***
Required. Specifies the IP address that ends the range.

**{DHCP | BOOTP | BOTH}**
Specifies the type of client to be serviced by this scope and IP range. **DHCP** is the default.

***MaxBootP***
Specifies the maximum number of BOOTP clients.

**Remarks**

- *MaxBootP* is useful only if *ClientType* is set to support BOOTP-type clients for the scope. If an optional command-line option is used, all command-line options (mandatory as well as optional) that come before the optional command that is used are required and must appear in proper sequence.

**Examples**
In the following example, this command adds DHCP clients that are in the IP address range 10.2.2.10 to 10.2.2.20 to the distribution range in the scope.

**add iprange 10.2.2.10 10.2.2.20**

**add reservedip**
Reserves an IP address for use by a specified media access control (MAC) address in the current scope.

**Syntax**

**add reservedip** *ReservedIP MACAddress* [*ClientName*] [*ClientComment*] [{**DHCP** | **BOOTP** | **BOTH**}]

**Parameters**

***ReservedIP***

Required. Specifies the IP address to reserve.

***MACAddress***

Required. Specifies the physical hardware or MAC address to associate with the reserved IP address.

***ClientName***

Specifies a client name to associate with this reserved client entry. If unspecified, a client name is not associated with this reserved client entry.

***ClientComment***

Specifies a client comment to associate with this reserved client entry. If unspecified, a client comment is not associated with this reserved client entry.

**{DHCP | BOOTP | BOTH}**

Specifies the type of clients to associate with this reserved client entry. **DHCP** is the default.

**Remarks**

- If an optional command-line option is used, all command-line options (mandatory as well as optional) that come before the optional command-line option that is used are required and must appear in proper sequence.

**Examples**

In the following example, this command reserves the IP address 10.2.2.32 for use by the DHCP client that identifies its MAC address as 08-00-2b-30-36-9b when it obtains a lease in the current scope.

**add reservedip 10.2.2.32 08002b30369B**

**delete excluderange**

Deletes a range of previously excluded IP addresses from the current scope.

**Syntax**

**delete excluderange** *StartIP EndIP*

**Parameters**

***StartIP***

Required. Specifies the IP address that starts the exclusion range.

***EndIP***

Required. Specifies the IP address that ends the exclusion range.

**Examples**

In the following example, this command deletes the exclusion range that starts with an IP address of 10.2.1.5 and ends with an IP address of 10.2.1.10 from the current scope.

**delete excluderange 10.2.1.5 10.2.1.10**

**delete iprange**

Deletes a range of IP addresses from the current scope.

**Syntax**

**delete iprange** *StartIP EndIP*

**Parameters**

***StartIP***

Required. Specifies the IP address that starts the range to delete.

***EndIP***

Required. Specifies the IP address that ends the range to delete.

**Examples**

In the following example, this command deletes IP addresses in the range 10.2.1.5 to 10.2.1.10 from the current scope.

**delete iprange 10.2.1.5 10.2.1.10**

**delete lease**
Deletes an IP address lease from the current scope.

**Syntax**
**delete lease** [{*LeaseIP* | *****HostName* | **AllBadAddresses** | **AllRasServerAddresses**}]

**Parameters**
***{LeaseIP* | *HostName* | AllBadAddresses | AllRasServerAddresses}**
Required. Specifies the IP address or the host name associated with the leased address. If **AllBadAddresses** is specified, all records marked as BAD_ADDRESS are removed. If **AllRasServerAddresses** is specified, all leases obtained by all Routing and Remote Access Service servers are removed.

**Examples**
In the first example, this command deletes the IP address lease 10.2.1.32 from the current scope.

In the second example, this command deletes all IP address leases marked as BAD_ADDRESS from the current scope.

**delete lease 10.2.1.32**

**delete lease AllBadAddresses**

**delete optionvalue**
Removes or clears the currently set scope option value from the current scope.

**Syntax**
**delete optionvalue** *OptCode*[**user=***UserName*] [**vendor=***VendorName*]

**Parameters**
***OptCode***
Required. Specifies the unique identifier of the option whose value is to delete.

**user= *UserName***
Specifies the user class from which to delete the currently set option value. If the tag is provided, but the value unspecified, the current default user class is assumed.

**vendor= *VendorName***
Specifies the vendor class from which to delete the currently set option value. If the tag is provided, but the value unspecified, the current default vendor class is assumed.

**Remarks**

- Vendor and user classes are only supported for DHCP servers running Windows 2000 Server.
- If you specify a user class, a vendor class, or both, this command deletes the set option value from only the scope clients that are identified as members of the specified class or classes.
- To modify the current defaults for an unspecified class that is used with this command, use either the **set userclass** or **set vendorclass** commands.

**Examples**
In the first example, this command deletes the currently set option value for option code 18 from the current scope.

In the second example, this command deletes the currently set option value for option code 18 from the vendor class MyVendorClass.

**delete optionvalue 18**

**delete optionvalue 18 vendor="MyVendorClass"**

**delete reservedip**
Deletes a reservation for an IP address in the current scope.

**Syntax**
**delete reservedip** *ReservedIP MACAddress*

**Parameters**
***ReservedIP***

Required. Specifies the IP address reservation to delete from the current scope.

***MACAddress***
Required. Specifies the media access control (MAC) or physical hardware address string for which the IP address was reserved.

**Examples**
In the following example, this command removes the IP address reservation 10.2.1.32 for the MAC address 08002B30369B from the current scope.

**delete reservedip 10.2.1.32 08002B30369B**

**delete reservedoptionvalue**
Deletes an option value that is currently assigned to a reserved client in the current scope.

**Syntax**
**delete reservedoptionvalue** *ReservedIP OptCode* [**User=***UserName*] [**vendor=***VendorName*]

**Parameters**
***ReservedIP***
Required. Specifies the IP address of the reserved client.

***OptCode***
Required. Specifies the unique code for the option type that is currently assigned to the reserved client.

**User= *UserName***
Specifies the user class from which to delete the option value. If the tag is provided, but no value is specified, the current default user class is assumed.

**vendor= *VendorName***
Specifies the vendor class from which to delete the option value. If the tag is provided, but no value is specified, the current vendor class is assumed.

**Remarks**

- Vendor and user classes are only supported for DHCP servers running Windows 2000 Server.
- To modify the current defaults for an unspecified class used with this command, use either **set userclass** or **set vendorclass**.

**Examples**
In the following example, this command deletes the currently set option value for the option identified by code 18 for reserved IP address of 10.2.2.32 in the current scope.

**delete reservedoptionvalue 10.2.2.32 18**

**dump**
Dumps the configuration of the current scope to the command prompt window when run within the netsh environment.

**Syntax**
**dump**

When run at the command prompt or in a batch file, output can be saved in a text file. Used without parameters, this command dumps the configuration of the current scope to the command prompt window.

**Syntax**
**netsh dhcp server { *ServerName* | *IPAddress* } scope *ScopeID* dump >**[*PathAndFileName*]

**Parameters**
***ServerName* | *IPAddress***
Specifies the server name or IP address of the DHCP server for which the scope configuration is output.

***ScopeID***
Required. Specifies the IP address of the scope for which the configuration is output.

***PathAndFileName***
Specifies both the location where the file is saved, and the name of the destination file to which to dump the scope

configuration. If unspecified, the scope configuration is dumped to the command prompt window.

**Examples**
The first command, which is run within the netsh environment, dumps the configuration of the current scope to the command prompt window.

The second command, which is run from the command prompt, dumps the configuration of the local server scope 192.168.1.0 to a file named Scopecfg.dmp on the shared network folder \\Backup\Dhcp\.

**dump**

**netsh dhcp server scope 192.168.1.0 dump >** *\\Backup\Dhcp\Scopecfg.dmp*

**initiate reconcile**
Checks and reconciles the current scope. Used without parameters, this command verifies the scopes and check for inconsistencies but does not fix any inconsistencies that it finds in the database.

**Syntax**
**initiate reconcile** [**fix**]

**Parameters**
**[fix]**
Indicates that the command will fix, if possible, any inconsistencies that it finds in the database.

**Examples**
In the first example, this command verifies the scopes and checks for inconsistencies.

In the second example, this command verifies the scopes, checks for inconsistencies, and fixes any inconsistencies that are found.

**initiate reconcile**

**initiate reconcile fix**

**set comment**
Sets the comment for the current scope. Used without parameters, **set comment** deletes the current comment.

**Syntax**
**set comment** [*NewComment*]

**Parameters**
**[*NewComment*]**
Specifies a new or modified comment string for the scope.

**Examples**
In the first example, this command sets a comment string that has no spaces.

In the second example, this command sets a comment string that includes spaces.

**set comment NewCommentNoSpaces**

**set comment "New Comment With Spaces"**

**set name**
Sets the name of the current scope.

**Syntax**
**set name** *NewName*

**Parameters**
***NewName***
Required. Specifies the new name of the scope.

**Examples**
In the first example, this command sets a name that has no spaces.

In the second example, this command sets a name that includes spaces.

**set name NewNameNoSpaces**

**set name "New Name With Spaces"**

#### set optionvalue

Sets an option value for the current scope.

#### Syntax

**set optionvalue** *OptCode*{**BYTE** | **WORD** | **DWORD** | **STRING** | **IPADDRESS**} [**user=***UserName*] [**vendor=***VendorName*] *OptionValue*

#### Parameters

***OptCode***

Required. Specifies the code for the option type whose value is to be set.

**{BYTE | WORD | DWORD | STRING | IPADDRESS}**

Required. Specifies the data type for the option type whose value is to be set.

**user= *UserName***

Specifies the user class. If the tag is provided, but the value is unspecified, the current default user class is assumed.

**vendor= *VendorName***

Specifies the vendor class. If the tag is provided, but the value is unspecified, the current default vendor class is assumed.

***OptionValue***

Required. Specifies the assigned value for the option type that is specified in *OptCode*. If the option type supports an array that contains more than a single numeric or IP address value, provide the additional values, in the order that you prefer them, at the end of the command, with each value separated by a space.

#### Remarks

- Vendor and user classes are only supported for DHCP servers running Windows 2000 Server.
- To modify the current defaults for an unspecified class used with this command, use either **set userclass** or **set vendorclass**.

#### Examples

In the first example, this command sets the value of option code 003 to list two router IP addresses (10.1.1.1, 10.1.1.2).

In the second example, this command sets the value that is specified in the first example to apply only to those scope clients that identify themselves as members of the vendor class Vendor1, a class that was previously defined at the server with this specified option type defined for its use.

**set optionvalue 003 IPADDRESS 10.1.1.1 10.1.1.2**

**set optionvalue 003 IPADDRESS vendor=Vendor1 10.1.1.1 10.1.1.2**

#### set reservedoptionvalue

Sets the value of an option for a reservation IP address in the current scope.

#### Syntax

**set reservedoptionvalue** *ReservedIP OptCode* {**BYTE** | **WORD** | **DWORD** | **STRING** | **IPADDRESS**} [**user=***UserName*] [**vendor=***VendorName*] *OptValue*

#### Parameters

***ReservedIP***

Required. Specifies the reserved IP address for which the option value is to be set.

***OptCode***

Required. Specifies the code for the option type whose value is to be set.

**{BYTE | WORD | DWORD | STRING | IPADDRESS}**

Required. Specifies the data type for the option type whose value is to be set.

**user= *UserName***

Specifies either the current default user class or the class specified as *UserName*. If the tag is provided, but no value is specified, the current default user class is assumed.

**vendor= *VendorName***

Specifies either the current default vendor class or the class specified as *VendorName*. If the tag is provided, but no value is specified, the current default vendor class is assumed.

***OptValue***
Required. Specifies the assigned value for the option type specified in *OptCode*. If the option type supports an array that contains more than a single numeric or IP address value, provide the additional values, in the order that you prefer them, at the end of the command with each value separated by a space.

**Remarks**

- Vendor and user classes are only supported for DHCP servers running Windows 2000 Server.
- To modify the current defaults for an unspecified class used with this command, use either **set userclass** or **set vendorclass**.

**Examples**
In the following example, this command sets the value of the router option (code 003) for the reserved client IP address of 10.1.1.50 in the current scope to set IP addresses of 10.1.1.1 and 10.1.1.2 for its configured routers (default gateways).

**set reservedoptionvalue 10.1.1.50 003 IPADDRESS 10.1.1.1 10.1.1.2**

**set scope**
Sets the scope for use in subsequent operations.

**Syntax**
**set scope** *ScopeAddress*

**Parameters**
***ScopeAddress***
Required. Specifies the IP address of the scope to use in subsequent command operations.

**Examples**
In the following example, this command sets the current scope to 10.2.2.0 for subsequent operations.

**set scope 10.2.2.0**

**set state**
Sets or resets the state of the current scope to either the active or inactive state. Used without parameters, this command activates the scope.

**Syntax**
**set state** [{**0** | **1** | **2** | **3**}]

**Parameters**
**{0 | 1 | 2 | 3}**
Indicates the state of the scope: **0**-Deactivates the scope, **1**-Activates the scope (default), **2**-Deactivates the scope and marks the scope as "Switched," **3**-Activates the scope and marks the scope as "Switched."

**Remarks**

- **2** and **3** are typically used for switched networks or networks where multiple logical networks are hosted on a single physical network.

**Examples**
In the first example, this command activates a scope.

In the second example, this command deactivates a scope.

**set state 1**

**set state 0**

**set superscope**
Sets the superscope to use in subsequent operations.

**Syntax**

**set superscope** *SuperscopeName* {**0** | **1**}

**Parameters**

***SuperscopeName***

Required. Specifies the name of the superscope to include the current scope. The *SuperscopeName* is case-sensitive.

**{0 | 1}**

Required. Indicates whether to set the state of the superscope to active or inactive: **0**-Deactivates the superscope, **1**-Activates the superscope.

**Examples**

In the following example, this command adds the current scope to the superscope MySuperScope and activates the superscope.

**set superscope MySuperScope 1**

**show clients**

Displays all of the available version 4 clients for the current scope. Used without parameters, this command displays the following information for each client: IP address, subnet mask, unique ID, lease expiration, and type.

**Syntax**

**show clients**[{**0** | **1**}]

**Parameters**

**{0 | 1}**

Indicates the detail level of the output: **0**-Shows the following information for each client: IP address, subnet mask, unique ID, lease expiration, and data type (default), **1**-Shows all of the information that **0** provides and also displays the fully qualified domain name of each client.

**Remarks**

- To view the output of this command effectively, increase the width of the command prompt window to at least 95 characters.

**show clientsv5**

Displays all of the available version 5 clients for the current scope. Used without parameters, this command displays the following information for each client: IP address, subnet mask, unique ID, lease expiration, and type.

**Syntax**

**show clientsv5**[{**0** | **1**}]

**Parameters**

**{0 | 1}**

Indicates the detail level of the output: **0**-Shows the following information for each client: IP address, subnet mask, unique ID, lease expires, type (default), **1**-Shows all of the information that **0** provides and also displays the fully qualified domain name of each client.

**Remarks**

- To view the output of this command effectively, increase the width of the command prompt window to at least 95 characters.

**show excluderange**

Displays all of the currently set exclusion ranges of IP addresses for the current scope.

**Syntax**

**show excluderange**

**Parameters**

**none**

**Remarks**

- Use other scope-level commands to add and delete exclusion ranges.

**show iprange**
Displays all of the address ranges that are available for the current scope.

**Syntax**
**show iprange**

**Parameters**
**none**

**show optionvalue**
Displays all of the option values that are set for the current scope. Used without parameters, this command assumes the current default user and vendor classes.

**Syntax**
**show optionvalue**[**user=***UserName*] [**vendor=***VendorName*]

**Parameters**
**user= *UserName***
Specifies that the options that are set for the specified user class will display. If the tag is provided, but no value is specified, the current default user class is assumed.

**vendor= *VendorName***
Specifies that the options that are set for the specified vendor class will display. If the tag is provided, but no value is specified, the current default vendor class is assumed.

**Remarks**

- Vendor and user classes are only supported for DHCP servers running Windows 2000 Server.
- To modify the current defaults for an unspecified class used with this command, use either **set userclass** or **set vendorclass**.

**Examples**
In the following example, this command displays all options and values set for the current scope for the user defined class My User Class.

**show optionvalue user="My User Class"**

**show reservedip**
Displays all of the IP addresses that are currently reserved for the current scope.

**Syntax**
**show reservedip**

**Parameters**
**none**

**show reservedoptionvalue**
Displays all currently set option values for a reserved client IP address in the current scope.

**Syntax**
**show reservedoptionvalue** *ReservedIP*[**user=***UserName*] [**vendor=***VendorName*]

**Parameters**
***ReservedIP***
Required. Specifies the IP address reservation for which currently assigned options are to display.

**user= *UserName***
Specifies that the options that are set for the specified user class will display. If the tag is provided, but no value is specified, the current default user class is assumed.

**vendor= *VendorName***
Specifies that the options that are set for the specified vendor class will display. If the tag is provided, but no value is specified, the current default vendor class is assumed.

**Remarks**

- Vendor and user classes are only supported for DHCP servers running Windows 2000 Server.
- To modify the current defaults for an unspecified class used with this command, use either **set userclass** or **set vendorclass**.

**Examples**
In the following example, this command displays the option values set for the reserved IP address 10.2.2.100 in the current scope.

**show reservedoptionvalue 10.2.2.100**

**show scope**
Displays information for the current scope.

**Syntax**
**show scope**

**Parameters**
**none**

**show state**
Displays the state of the current scope, indicating whether it is active or inactive.

**Syntax**
**show state**

**Parameters**
**none**

**Netsh DHCP server mscope**
The following commands can run from the **dhcp server mscope>** prompt, which is rooted within the **netsh** environment.

To view the command syntax, click a command:

- add excluderange
- add iprange
- delete excluderange
- delete iprange
- dump
- initiate reconcile
- set comment
- set lease
- set mscope
- set name
- set state
- set ttl

- show clients
- show excluderange
- show iprange
- show lease
- show mibinfo
- show mscope
- show state
- show ttl

**add excluderange**
Adds a range of excluded addresses to the current multicast scope.

**Syntax**
**add excluderange** *StartIP EndIP*

**Parameters**
***StartIP***
Required. Specifies the IP address that starts the exclusion range.

***EndIP***
Required. Specifies the IP address that ends the exclusion range.

**Remarks**

- The exclusion range must be a subset of the overall scope address range.

**Examples**
In the following example, this command adds an exclusion range that starts with 224.2.2.10 and ends with 224.2.2.20 to the distribution range of the current multicast scope.

**add excluderange 224.2.2.10 224.2.2.20**

**add iprange**
Adds a range of IP addresses to the current multicast scope.

**Syntax**
**add iprange** *StartIP EndIP*

**Parameters**
***StartIP***
Required. Specifies the IP address that starts the range.

***EndIP***
Required. Specifies the IP address that ends the range.

**Remarks**

- The range must be within the valid range of multicast IP addresses (from 224.0.0.0 to 239.255.255.255).

**Examples**
In the following example, this command adds the IP address range 224.2.2.10 to 224.2.2.20 to the distribution range in the multicast scope.

**add iprange 224.2.2.10 224.2.2.20**

**delete excluderange**

Deletes an exclusion range of previously excluded IP addresses in the current multicast scope.

**Syntax**
**delete excluderange** *StartIP EndIP*

**Parameters**
***StartIP***
Required. Specifies the IP address that starts the exclusion range.

***EndIP***
Required. Specifies the IP address that ends the exclusion range.

**Examples**
In the following example, this command removes a multicast scope exclusion with a range of IP addresses that starts at 224.2.2.10 and ends with 224.2.2.20 for the multicast scope.

**delete excluderange 224.2.2.10 224.2.2.20**

**delete iprange**
Deletes a range of IP addresses from the current multicast scope.

**Syntax**
**delete iprange** *StartIP EndIP*

**Parameters**
***StartIP***
Required. Specifies the IP address that starts the range to delete.

***EndIP***
Required. Specifies the IP address that ends the range to delete.

**Examples**
In the following example, this command deletes the IP address range 224.2.2.10 to 224.2.2.20 from the overall range of the full multicast scope.

**delete iprange 224.2.2.10 224.2.2.20**

**dump**
Dumps the configuration of the current multicast scope to the command prompt window when run within the netsh environment.

**Syntax**
**dump**

When run at the command prompt or in a batch file, output can be saved in a text file. Used without parameters, this command dumps the configuration of the current multicast scope to the command prompt window.

**Syntax**
**netsh dhcp server [** ***ServerName*** **|** ***IPAddress*** **] mscope** ***MscopeID*** **dump >**[*PathAndFileName*]

**Parameters**
***ServerName | IPAddress***
Specifies the server name or IP address of the DHCP server for which the scope configuration is output.

***MscopeID***
Required. Specifies the name of the multicast scope for which the configuration is output.

***PathAndFileName***
Specifies both the location where the file is saved, and the name of the destination file to which to dump the multicast scope configuration. If unspecified, scope configuration is dumped to the command prompt window.

**Examples**
The first command, which is run within the netsh environment, dumps the configuration of the current multicast scope to the command prompt window.

The second command, which is run from the command prompt, dumps the configuration of the local server multicast scope TestMscope to the file Mscopecfg.dmp on the shared network folder \\Backup\Dhcp\.

**dump**

**netsh dhcp server mscope TestMscope dump >** *\\Backup\Dhcp\Mscopecfg.dmp*

#### initiate reconcile

Checks and reconciles the current multicast scope.

**Syntax**

**initiate reconcile**

**Parameters**

**none**

**Remarks**

- This command checks the integrity of the current multicast scope by comparing the current contents of the server database with a mirrored copy of the same information in the Windows registry. If inconsistencies are detected in the database, they are repaired based on the information that is duplicated in the registry. Repair is always attempted for any inconsistency that is found.

#### set comment

Sets the comment for the current multicast scope.

**Syntax**

**set comment** *NewComment*

**Parameters**

***NewComment***

Required. Specifies the new or revised comment for the multicast scope

**Examples**

In the first example, this command modifies the multicast scope comment with a new comment that contains no spaces.

In the second example, this command modifies the multicast scope comment with a new comment that contains spaces.

**set comment NewCommentNoSpaces**

**set comment "New Comment With Spaces"**

#### set lease

Sets the lease duration for the multicast scope IP addresses.

**Syntax**

**set lease** *Time*

**Parameters**

***Time***

Required. Specifies the lease duration for clients of the multicast scope. Specifying **-1** sets the duration of the IP address lease to an unlimited or infinite time.

**Examples**

In the following example, this command sets the lease duration for the clients of the current multicast scope to 691200 seconds (eight days).

**set lease 691200**

#### set mscope

Sets the multicast scope to use in subsequent operations.

**Syntax**

**set mscope** *NewName*

**Parameters**

***NewName***

Required. Specifies the name of the multicast scope to which the command context is changed. *NewName* is case-sensitive.

**Remarks**

- This command changes the context of the **netsh dhcp server mscope>** prompt from one multicast scope to another.

- This command does not rename the current multicast scope. To change the name of the current multicast scope, use the command **set name**.

**Examples**
In the first example, this command changes the command context to a multicast scope named MyMulticastScope. Note that the multicast scope name contains no spaces.

In the second example, this command changes the command context to a multicast scope named My Multicast Scope. Note that the multicast scope name contains spaces.

**set mscope MyMulticastScope**

**set mscope "My Multicast Scope"**

**set name**
Changes the name of the current multicast scope.

**Syntax**
**set name** *NewName*

**Parameters**
***NewName***
Required. Specifies a new name for the current multicast scope.

**Examples**
In the first example, this command sets a new name that contains no spaces for the current multicast scope.

In the second example, this command sets a new name that contains spaces for the current multicast scope.

**set name NewNameNoSpaces**

**set name "New Name With Spaces"**

**set state**
Sets or resets the state of the current multicast scope to either an active or inactive state.

**Syntax**
**set state** {**0** | **1**}

**Parameters**
**{0 | 1}**
Required. Sets the state of the current multicast scope: **0**-Deactivates the current multicast scope, **1**-Activates the current multicast scope .

**Examples**
In the first example, this command activates the current multicast scope.

In the second example, this command deactivates the current multicast scope.

**set state 1**

**set state 0**

**set ttl**
Sets the Time-To-Live (TTL) value for the current multicast scope.

**Syntax**
**set ttl** *TTL*

**Parameters**
***TTL***
Required. Specifies the Time-to-Live (TTL) value. The valid range for this value is a number from 1 to 255.

**Examples**
In the following example, this command sets the TTL value for the current multicast scope to 32.

**set ttl 32**

**show clients**
Displays all available clients for the current multicast scope.

**Syntax**
**show clients**

**Parameters**
**none**

**Remarks**

- To view the output of this command effectively, increase the width of the command prompt window to at least 95 characters.

**show excluderange**
Displays all currently excluded ranges of IP addresses for the current multicast scope.

**Syntax**
**show excluderange**

**Parameters**
**none**

**Remarks**

- If no exclusion ranges have been previously defined for the scope, this command outputs an empty list.

**show iprange**
Displays all available IP address ranges for the current multicast scope.

**Syntax**
**show iprange**

**Parameters**
**none**

**show lease**
Displays the current lease duration settings for the current multicast scope.

**Syntax**
**show lease**

**Parameters**
**none**

**show mibinfo**
Displays management information base (MIB) information for the current multicast scope.

**Syntax**
**show mibinfo**

**Parameters**
**none**

**show mscope**
Displays information for the current multicast scope.

**Syntax**
**show mscope**

**Parameters**
**none**

**show state**
Displays the state of the current multicast scope.

**Syntax**
**show state**

**Parameters**
**none**

**show ttl**
Displays the Time-To-Live (TTL) value for the current multicast scope.

**Syntax**
**show ttl**

**Parameters**
**none**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Netsh commands for AAAA
Netsh diagnostic (diag) commands
Netsh commands for Interface IP
Netsh commands for remote access (ras)
Netsh commands for WINS
Command-line reference A-Z
Command shell overview

# Netsh diagnostic (diag) commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh diagnostic (diag) commands**
You can use the Netsh Network Diagnostics commands, or diag context, to administer and troubleshoot operating system and network service parameters from the command line. The command prompt for the Netsh diagnostic context is **netsh diag>**.

The following commands are available at the command prompt for the context (as distinguished from the Windows Server 2003 family command prompt). The Netsh Diag context is new to the Windows Server 2003 family. You cannot run Netsh Diag commands on a Windows 2000 Server.

You can run these commands from the Windows Server 2003 family command prompt or from the command prompt for the Netsh diagnostic context. For these commands to work at the Windows Server 2003 family command prompt, you must type **netsh diag** before typing commands and parameters as they appear in the syntax below. There might be functional differences between Netsh context commands on Windows 2000 and the Windows Server 2003 family.

For more information about **netsh**, see Netsh overview and Enter a netsh context.

To view the command syntax, click a command:

- connect ieproxy
- connect iphost
- connect mail
- connect news
- dump
- gui
- ping
- ping adapter
- ping dhcp
- ping dns
- ping gateway
- ping ieproxy
- ping ip
- ping iphost
- ping loopback
- ping mail
- ping news
- ping wins

- show adapter
- show all
- show client
- show computer
- show dhcp
- show dns
- show gateway
- show ieproxy
- show ip
- show mail
- show modem
- show news
- show os
- show test
- show version
- show wins

**connect ieproxy**
Establishes, verifies, and then drops a TCP/IP connection with the proxy server specified in Internet Explorer.

**Syntax**
**connect ieproxy**

**Parameters**
none

**Remarks**

- If no proxy server is specified in Internet Explorer, the command generates output similar to the following:

```
Internet Explorer Proxy (Not Configured)
    Not Configured
```

- If you have specified in Internet Explorer a proxy server named corpproxy with a port of 80, the command generates output similar to the following:

```
Internet Explorer Proxy (corpproxy)
    IEProxyPort = 80
    IEProxy = corpproxy
    Server appears to be running on port(s) [80]
```

**connect iphost**
Establishes, verifies, and then drops a TCP/IP connection to a specified host through a specified port.

**Syntax**
**connect iphost** *Host Port*

**Parameters**
***Host***
Specifies the IP address, NetBIOS name, or fully qualified domain name of the host to which you want to connect.

***Port***
Specifies the port through which you want to connect.

**Examples**
To connect to the host named microsoft.com through port 28, type:

**connect iphost microsoft.com 28**

**connect mail**
Establishes, verifies, and then drops a TCP/IP connection with the node specified as the mail server in the default Outlook Express profile of the current user.

**Syntax**
**connect mail**

**Parameters**
none

**connect news**
Establishes, verifies, and then drops a TCP/IP connection with the node specified as the news server in the default Outlook Express profile of the current user.

**Syntax**
**connect news**

**Parameters**
none

**dump**
Creates a script that contains the current configuration.

**Syntax**
**dump**

**Parameters**
none

**Remarks**

- If you save this script to a file, you can use it to restore altered configuration settings.

**gui**
Starts the network diagnostics tool in Help and Support Center.

**Syntax**
**gui**

**Parameters**
none

**ping**
Verifies connectivity with a specified remote host.

**Syntax**
**ping** *Host*

**Parameters**
***Host***
Specifies the IP address, NetBIOS name, or fully qualified domain name of the host with which you want to verify

connectivity.

**ping adapter**
Verifies connectivity with computers and other devices (for example, routers) identified in the TCP/IP properties of the specified network adapter or adapters. Used without parameters, the **ping adapter** command verifies connectivity through all installed adapters.

**Syntax**
**ping adapter** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}]

**Parameters**
***{ IndexNumber | FullAdapterName| PartialAdapterName}***
Specifies the adapter or adapters for which you want to verify connectivity. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**Remarks**

- This command verifies connectivity with all computers and other devices configured in the TCP/IP properties for each network adapter you specify. For example, if you have a network adapter whose TCP/IP properties are configured with IP addresses for a preferred DNS server, a default gateway, and a DHCP server, **ping adapter** verifies connectivity with each of these devices.

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.

- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**
To verify connectivity with all computers and other devices configured in the TCP/IP properties for each network adapter, type:

**ping adapter**

To verify connectivity with all computers and other devices configured in the TCP/IP properties for each network adapter with an index number of 2, type:

**ping adapter 2**

To verify connectivity with all computers and other devices configured in the TCP/IP properties for each network adapter with "net" at the beginning of its name, type:

**ping adapter net***

**ping dhcp**
Verifies connectivity with the Dynamic Host Configuration Protocol (DHCP) servers configured in the TCP/IP properties for the specified adapter or adapters. Used without parameters, the **ping dhcp** command verifies connectivity with DHCP servers through all installed adapters.

**Syntax**
**ping dhcp** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}]

**Parameters**
***{IndexNumber| FullAdapterName| PartialAdapterName}***
Specifies the adapter or adapters for which you want to verify connectivity. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**
To verify connectivity with the DHCP server configured in the TCP/IP properties for each network adapter, type:

**ping dhcp**

To verify connectivity with the DHCP server configured in the TCP/IP properties for the network adapter with an index number of 2, type:

**ping dhcp 2**

To verify connectivity with the DHCP server configured in the TCP/IP properties for each network adapter with "net" at the beginning of its name, type:

**ping dhcp net***

**ping dns**
Verifies connectivity with the Domain Name System (DNS) servers configured in the TCP/IP properties for the specified adapter or adapters. Used without parameters, the **ping dns** command verifies connectivity with DNS servers through all installed adapters.

**Syntax**
**ping dns** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}]

**Parameters**
**{ *IndexNumber*| *FullAdapterName*| *PartialAdapterName*}**
Specifies the adapter or adapters for which you want to verify connectivity. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**
To verify connectivity with the DNS server configured in the TCP/IP properties for each network adapter, type:

**ping dns**

To verify connectivity with the DNS server configured in the TCP/IP properties for the network adapter with an index number of 2, type:

**ping dns 2**

To verify connectivity with the DNS server configured in the TCP/IP properties for each network adapter with "net" at the beginning of its name, type:

**ping dns net***

**ping gateway**
Verifies connectivity with the default gateway servers configured in the TCP/IP properties for the specified adapter or adapters. Used without parameters, the **ping gateway** command verifies connectivity through all installed adapters.

**Syntax**

**ping gateway** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}]

**Parameters**

***{ IndexNumber| FullAdapterName| PartialAdapterName}***

Specifies the adapter or adapters for which you want to verify connectivity. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**

To verify connectivity with the gateway server configured in the TCP/IP properties for each network adapter, type:

**ping gateway**

To verify connectivity with the gateway server configured in the TCP/IP properties for the network adapter with an index number of 2, type:

**ping gateway 2**

To verify connectivity with the gateway server configured in the TCP/IP properties for each network adapter with "net" at the beginning of its name:

**ping gateway net***

**ping ieproxy**

Verifies connectivity with the Web server proxy for Internet Explorer.

**Syntax**

**ping ieproxy**

**Parameters**

none

**ping ip**

Verifies connectivity with all IP addresses for the specified adapter or adapters. Used without parameters, the **ping ip** command verifies connectivity through all installed adapters.

**Syntax**

**ping ip** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}]

**Parameters**

***{ IndexNumber| FullAdapterName| PartialAdapterName}***

Specifies the adapter or adapters for which you want to verify connectivity. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.

- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**
To verify connectivity with each network adapter, type:

**ping ip**

To verify connectivity with the network adapter with an index number of 2, type:

**ping ip 2**

To verify connectivity with each network adapter with "net" at the beginning of its name, type:

**ping ip net***

**ping iphost**
Verifies connectivity with a remote or local host.

**Syntax**
**ping iphost** *Node*

**Parameters**
***Node***
Specifies a host by its IP address, its NetBIOS name, or its fully qualified domain name.

**Examples**
To verify connectivity with the local host (at the loopback address of 127.0.0.1), type:

**ping iphost 127.0.0.1**

To verify connectivity with the host microsoft.com, type:

**ping iphost microsoft.com**

**ping loopback**
Verifies connectivity with the local host loopback address (127.0.0.1).

**Syntax**
**ping loopback**

**Parameters**
none

**ping mail**
Verifies connectivity with the mail server configured in the default Outlook Express profile of the current user.

**Syntax**
**ping mail**

**Parameters**
none

**ping news**
Verifies connectivity with the mail server configured in the default Outlook Express profile of the current user.

**Syntax**
**ping news**

**Parameters**
none

**ping wins**
Verifies connectivity with the preferred and alternate Windows Internet Name Service (WINS) servers for the specified network adapter or adapters. Used without parameters, **ping wins** verifies connectivity through all installed adapters.

**Syntax**
**ping wins** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}]

**Parameters**

**{ *IndexNumber*| *FullAdapterName*| *PartialAdapterName*}**
Specifies the adapter or adapters for which you want to verify connectivity. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**
To verify connectivity with the WINS servers whose IP addresses are configured in the TCP/IP properties of each installed adapter, type:

**ping wins**

To verify connectivity with the WINS servers whose IP addresses are configured in the TCP/IP properties of each installed adapter with an index number number of 2, type:

**ping wins 2**

**show adapter**
Shows information about the specified logical adapter or adapters. Used without parameters, the **show adapter** command shows information about all adapters defined for this server.

**Syntax**
**show adapter** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}] [{**/v** | **/p**}]

**Parameters**
**{ *IndexNumber*| *FullAdapterName*| *PartialAdapterName*}**
Specifies the adapter or adapters for which you want to show information. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- This command lists each adapter by index number in square brackets followed by the name of the adapter as it is registered. Typical output from this command follows:

  Network Adapters

  1. [00000001] 3Com 3C920 Integrated Fast Ethernet Controller (3C905C-TX Compatible)
  2. [00000002] RAS Async Adapter
  3. [00000003] WAN Miniport (L2TP)
  4. [00000004] WAN Miniport (PPTP)

5. [00000005] WAN Miniport (PPPOE)

6. [00000006] Direct Parallel

7. [00000007] WAN Miniport (IP)

8. [00000008] WAN Miniport (Network Monitor)

**Examples**
To display the logical adapters defined on the computer, type:

**show adapter**

**show all**
Lists all network objects defined for the local server. Used without parameters, the **show all** command lists overview configuration information for the following hardware and software components: Default Outlook Express Mail, Default Outlook Express News, Internet Explorer Web Proxy, Loopback, Computer System, Operating System, Version, Modems, Network Adapters, and Network Clients.

**Syntax**
**show all** [{**/v** | **/p**}]

**Parameters**
**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- The output from this command can be lengthy and includes configuration information for mail, news, and proxy servers. It lists information for the TCP/IP loopback address and details about the computer, the operating system, and all network adapters and clients.

**show client**
Lists all network clients defined for the specified adapter or adapters. Used without parameters, **show client** lists clients for all installed adapters.

**Syntax**
**show client** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}] [{**/v** | **/p**}]

**Parameters**
**{ *IndexNumber*| *FullAdapterName*| *PartialAdapterName*}**
Specifies the adapter or adapters for which you want to list clients. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.

- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**

To list all network clients defined for all network adapters in the host, type:

**show client**

To list all fields with non-null values for all network clients defined for adapter number 2, type:

**show client 2 /p**

**show computer**
Lists Windows Management Interface (WMI) details for the server. Used without parameters, the **show computer** command lists the NetBIOS name of the local computer.

**Syntax**
**show computer** [{**/v** | **/p**}]

**Parameters**
**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Examples**
To list all non-null WMI interface parameters for the server:

**show computer /p**

**show dhcp**
Lists all DHCP servers defined for the specified adapter. Used without parameters, the **show dhcp** command lists the servers defined for all installed adapters.

**Syntax**
**show dhcp** {*IndexNumber* | *FullAdapterName* | *PartialAdapterName*} [{**/v** | **/p**}]

**Parameters**
***{ IndexNumber| FullAdapterName| PartialAdapterName}***
Specifies the adapter or adapters for which you want to list servers. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**
To list details of all DHCP servers defined for all network adapters, type:

**show dhcp /v**

To list details of any DHCP servers defined for network adapter 2, type:

**show dhcp 2**

**show dns**
Lists all DNS servers defined for a specified adapter. Used without parameters, the **show dns** command lists servers defined for all installed adapters.

**Syntax**

**show dns** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}] [{**/v** | **/p**}]

**Parameters**

**{ *IndexNumber*| *FullAdapterName*| *PartialAdapterName*}**

Specifies the adapter or adapters for which you want to list servers. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**{ /v| /p}**

Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**

To list all non-null details of all DNS servers defined for all network adapters, type:

**show dns /p**

To list any DNS servers defined for network adapter 2, type:

**show dns 2**

**show gateway**

Lists all Internet gateways defined for the specified adapter or adapters. Used without parameters, the **show gateway** command lists gateways defined for all installed adapters.

**Syntax**

**show gateway** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}] [{**/v** | **/p**}]

**Parameters**

**{ *IndexNumber*| *FullAdapterName*| *PartialAdapterName*}**

Specifies the adapter or adapters for which you want to list gateways. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**{ /v| /p}**

Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**

To list all available details about all Internet gateways defined for all adapters, type:

**show gateway /v**

To list all non-null values available for any gateways defined for network adapter 3, type:

**show gateway 3 /p**

**show ieproxy**
Lists Internet Explorer proxy servers defined for the specified adapter or adapters. Used without parameters, **show ieproxy** lists the servers for all installed adapters.

**Syntax**
**show ieproxy** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}] [{**/v** | **/p**}]

**Parameters**
**{ *IndexNumber*| *FullAdapterName*| *PartialAdapterName*}**
Specifies the adapter or adapters for which you want to list servers. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, you must type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, you must type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, you must type an asterisk at the beginning of the parameter and at the end of the parameter.

**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**
To list all Internet Explorer proxy servers defined for all network adapters in the server, type:

**show ieproxy**

**show ip**
Shows information about the specified adapter or adapters for which TCP/IP is enabled. Used without parameters, the **show ip** command shows information for all installed adapters.

**Syntax**
**show ip** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}] [{**/v** | **/p**}]

**Parameters**
**{ *IndexNumber*| *FullAdapterName*| *PartialAdapterName*}**
Specifies the adapter or adapters for which you want to show information. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.

**Examples**
To list the index number, name, and TCP/IP address for all adapters for which TCP/IP is enabled, type:

**show ip**

**show mail**
Shows the Outlook Express mail server specified for the local computer. Used without parameters, the **show mail** command lists overview configuration information for Default Outlook Express Mail.

**Syntax**
**show mail** [{**/v** | **/p**}]

**Parameters**
**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Examples**
To display the mail server configured for the local computer, type:

**show mail**

**show modem**
Lists details about the specified modem device or devices configured on the local computer. Used without parameters, **show modem** lists details about all modem devices.

**Syntax**
**show modem** [{*IndexNumber* | *FullModemName* | *PartialModemName*}] [{**/v** | **/p**}]

**Parameters**
***{ IndexNumber| FullModemName| PartialModemName}***
Specifies the modem device or devices for which you want to list details. The *IndexNumber* parameter specifies the device by its index number. The *FullModemName* parameter specifies the device by its name as registered in the operating system. The *PartialModemName* specifies the device or devices whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Examples**
To show all details, even those with null values, about all modem devices configured on the local computer, type:

**show modem /v**

To show the name and index number of all modem devices configured on the local computer with "dialup" at the beginning of their names, type:

**show modem dialup***

**show news**
Displays configuration information for the news server defined for the local computer. Used without parameters, the **show news** command lists overview configuration information for Default Outlook Express News.

**Syntax**
**show news** [{**/v** | **/p**}]

**Parameters**
**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Examples**
To display the name and the TCP/IP address of the news server, type:

**show news**

**show os**
Displays information about the operating system running on the local computer.

**Syntax**
**show os** [**/v**]

**Parameters**
**/v**
Specifies that fields with a null value should appear.

**Remarks**

- Information shown by this command is collected by the WMI provider win32_operatingsystem.

**show test**
Shows all network objects defined for the local computer and verifies connectivity with each object using the **ping** command. Used without parameters, the **show test** command lists overview configuration information for the following hardware and software components: Default Outlook Express Mail, Default Outlook Express News, Internet Explorer Web Proxy, Loopback, Computer System, Operating System, Version, Modems, Network Adapters, and Network Clients. In addition, connectivity is verified with the Loopback adapter and all IP addresses configured as a client option in the TCP/IP properties of each network adapter.

**Syntax**
**show test** [{**/v** | **/p**}]

**Parameters**
**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- The output of this command can be lengthy and might require a few minutes to complete. The output includes configuration information for mail, news, and proxy servers as well as the results of attempts to verify connectivity to those servers using the **ping** command. This command lists information for and verifies connectivity with the TCP/IP loopback address. This command lists details about the computer and the operating system as well as all network adapters and clients. Connectivity is verified with network adapters and clients using the **ping** command, and the results are displayed.

**Examples**
To show all information available to the Netsh diagnostic context, type:

**show test /v**

**show version**
Shows the version of the operating system and, optionally, the version of Windows Management Infrastructure (WMI). Used without parameters, the **show version** command lists the version number of the operating system installed on the local computer.

**Syntax**
**show version** [{**/v** | **/p**}]

**Parameters**
**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- The first line of output for this command shows the version of the operating system. The second line shows the version of WMI. Typical output for this command follows:

```
Version (5.1.2416)
    Version = 5.1.2416
    BuildVersion = 2416.0000
```

**Examples**
To show the version of the operating system and WMI, including fields with a null value, type:

**show version /v**

**show wins**
Lists WINS servers for the specified adapter or adapters. Used without parameters, the **show wins** command lists servers for all installed adapters.

**Syntax**
**show wins** [{*IndexNumber* | *FullAdapterName* | *PartialAdapterName*}] [{**/v** | **/p**}]

**Parameters**
**{ *IndexNumber*| *FullAdapterName*| *PartialAdapterName*}**
Specifies the adapter or adapters for which you want to list servers. The *IndexNumber* parameter specifies the adapter by its index number. The *FullAdapterName* parameter specifies the adapter by its name as registered in the operating system. The *PartialAdapterName* parameter specifies the adapter or adapters whose names contain the sequence of characters specified in the parameter. If the sequence of characters must appear at the beginning of the name, type an asterisk at the end of the parameter. If the sequence of characters must appear at the end of the name, type an asterisk at the beginning of the parameter. If the sequence of characters can appear anywhere in the name, type an asterisk at the beginning of the parameter and at the end of the parameter.

**{ /v| /p}**
Specifies whether fields with a null value should appear. The **/v** parameter specifies that the fields should appear. The **/p** parameter specifies that the fields should not appear.

**Remarks**

- You can use the **show adapter** command to display the index numbers and the full names of all logical adapters defined on the computer.
- If you specify an adapter by its index number, you do not need to type the leading zeroes.
- Typical output for this command is as follows:

```
WINS Servers
    2. [00000002] 3Com 3C920 Integrated Fast Ethernet Controller (3C905C-TX
Compatible)
    WINSPrimaryServer = 172.16.44.17
    WINSSecondaryServer = 172.16.32.22
```

**Examples**
To list the WINS servers defined for all adapters on the server, type:

**show wins**

To list the WINS servers specified for the adapter with index number 2, type:

**show wins 2**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |

| | |
|---|---|
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Netsh commands for AAAA
Netsh commands for DHCP
Netsh commands for Interface IP
Netsh commands for remote access (ras)
Netsh commands for WINS
Command-line reference A-Z
Command shell overview

# Netsh commands for Interface IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for Interface IP**
You can use commands in the Netsh Interface IP context to configure the TCP/IP protocol (including addresses, default gateways, DNS servers, and WINS servers) and to display configuration and statistical information.

You can run these commands from the Windows Server 2003 family command prompt or from the command prompt for the Netsh interface IP context. For these commands to work at the Windows Server 2003 family command prompt, you must type **netsh interface ip** before typing commands and parameters as they appear in the syntax below. There might be functional differences between Netsh context commands on Windows 2000 and the Windows Server 2003 family.

For more information on Netsh, see Netsh overview and Enter a netsh context.

To view the command syntax, click a command:

- set address
- add address
- delete address
- show address
- set dns
- add dns
- delete dns
- show dns
- set wins
- add wins
- delete wins
- show wins
- show icmp
- show interface
- show ipaddress
- show ipnet
- show ipstats
- show joins
- show tcpconn

- show tcpstats
- show udpconn
- show udpstats
- show config
- show offload
- delete arpcache
- dump

**set address**
Configures an IP address and a default gateway on a specified interface.

**Syntax**
**set address** [**name=**]*InterfaceName* [**source=**]{**dhcp** | **static** [**addr=**]*IPAddress*[**mask=**]*SubnetMask* [**gateway=**]{**none** | *DefaultGateway* [[**gwmetric=**]*GatewayMetric*]}}

**Parameters**
**[ name =] *InterfaceName***
Required. Specifies the name of the interface for which you want to configure address and gateway information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ source=]{ dhcp| static[ addr=] *IPAddress*[ mask=] *SubnetMask*[ gateway=]{ none| *DefaultGateway*[[ gwmetric=] *GatewayMetric*]}}**
Required. Specifies whether the IP address to configure originates from a Dynamic Host Configuration Protocol (DHCP) server or is static. If the address is static, *IPAddress* specifies the address to configure, and *SubnetMask* specifies the subnet mask for the IP address being configured. If the address is static, you must also specify whether you want to leave the current default gateway (if any) in place or configure one for the address. If you configure a default gateway, *DefaultGateway* specifies the IP address of the default gateway to be configured, and *GatewayMetric* specifies the metric for the default gateway to be configured.

**/?**
Displays help at the command prompt.

**Remarks**

- You can use this command to switch an interface between DHCP configuration and static configuration. If the interface is configured statically, the DNS and WINS server addresses for this interface must also be statically configured.
- Omitting the *DefaultGateway* parameter does not clear gateway configuration but preserves whatever was previously configured.

**add address**
Adds an IP address and a default gateway on a specified interface configured with a static IP address.

**Syntax**
**add address** [**name=**]*InterfaceName* [**addr=**]*IPAddress* [**mask=**]*SubnetMask*[[**gateway=**] *DefaultGateway* [**gwmetric=**]*GatewayMetric*]

**Parameters**
**[ name=] *InterfaceName***
Required. Specifies the name of the interface for which you want to add address and gateway information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ addr=] *IPAddress*[ mask=] *SubnetMask***
Required. Specifies the IP address to add and the subnet mask for that IP address.

**[ gateway=] *DefaultGateway*[ gwmetric=] *GatewayMetric***
Specifies the IP address of the default gateway to add and the metric for that default gateway.

**/?**
Displays help at the command prompt.

**delete address**
Deletes an IP address or a default gateway on a statically configured interface.

**Syntax**
**delete address** [**name=**]*InterfaceName* [**addr=**] *IPAddress* [[**gateway=**]{*DefaultGateway* | **all**}]

**Parameters**
**[ name=] *InterfaceName***
Required. Specifies the name of the interface for which you want to delete address and gateway information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ addr=] *IPAddress***
Required. Specifies the IP address to delete.

**[ gateway=]{ *DefaultGateway*| all}**
Specifies whether to delete one default gateway or all default gateways. If only one default gateway should be deleted, *DefaultGateway* specifies the IP address of the default gateway to be deleted.

**/?**
Displays help at the command prompt.

**Remarks**

- This command will not delete the last IP address on the interface.

**show address**
Displays information about static IP addresses and default gateways on a specified interface. Used without parameters, **show address** displays address information for all interfaces.

**Syntax**
**show address** [[**name=**]*InterfaceName*]

**Parameters**
**[ name=] *InterfaceName***
Specifies the name of the interface for which you want to display address information. The *InterfaceName* must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**/?**
Displays help at the command prompt.

**Remarks**

- For interfaces configured using DHCP, use the **ipconfig** command to display the assigned IP addresses, lease parameters, and other DHCP information.

**set dns**
Configures a DNS server address for a specified interface.

**Syntax**
**set dns** [**name=**]*InterfaceName* [**source=**]{**dhcp** | **static** }[**addr=**]{*IP Address* | **none**} [**register=**]{**none** | **primary** | **both**}

**Parameters**
**[ name=] *InterfaceName***
Required. Specifies the name of the interface for which you want to set DNS information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ source=]{ dhcp| static }**

Required. Specifies whether the IP address of the DNS server is configured by DHCP or is static.

**[ addr=]{ *IP Address*| none }**
If the IP address is static, *IP Address* specifies the IP address of the DNS server to configure, and **none** specifies that the DNS configuration should be removed.

**[ register=]{ none| primary| both }**
**None** specifies whether to disable Dynamic DNS resolution. **Primary** registers the computer name under the primary DNS suffix only. **Both** registers the computer name under both the primary DNS suffix as well as under the connection-specific suffix.

**/?**
Displays help at the command prompt.

**Remarks**

- If the interface is already statically configured, the **static** parameter replaces the existing list of DNS server addresses with the one specified in the command.

- Only if the [source] is **static** is the [addr] option available for configuring a static list of DNS server IP addresses for the specified interface.

**Examples**
**set dns name="Local Area Connection" source=dhcp**

**set dns "Local Area Connection" static 10.0.0.1 primary**

**add dns**
Adds a DNS server to a list of DNS servers for a specified interface.

**Syntax**
**add dns** [**name=**]*InterfaceName* [**addr=**] *DNSAddress* [[**index=**]*DNSIndex*]

**Parameters**
**[ name=] *InterfaceName***
Required. Specifies the name of the interface for which you want to add DNS information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ addr=] *DNSAddress***
Required. Specifies the IP address of the DNS server to add.

**[ index=] *DNSIndex***
Specifies the position of the added DNS server in the list of DNS servers for the interface.

**/?**
Displays help at the command prompt.

**Remarks**

- A lower number for the *DNSIndex* parameter corresponds to a higher preference. If no index is specified, then the server is added with the lowest preference.

**delete dns**
Deletes a DNS server or all DNS servers from a list of DNS servers for a specified interface or for all interfaces.

**Syntax**
**delete dns** [**name=**]*InterfaceName* [**addr=**]{*DNSAddress* | **all**}

**Parameters**
**[ name=] *InterfaceName***
Required. Specifies the name of the interface for which you want to delete DNS information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ addr=]{ *DNSAddress*| all}**
Required. Specifies whether to delete the address of one DNS server or all servers for all interfaces. If only one DNS server should be deleted, *DNSAddress* specifies the IP address of the DNS server to delete.

**/?**
Displays help at the command prompt.

**show dns**
Displays the DNS configuration of a specified interface. Used without parameters, **show dns** displays the DNS configurations of all interfaces.

**Syntax**
**show dns** [[**name=**]*InterfaceName*]

**Parameters**
**[ name=] *InterfaceName***
Specifies the name of the interface whose DNS configuration you want to display. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**/?**
Displays help at the command prompt.

**set wins**
Sets WINS server configuration to either DHCP or static mode for a specified interface.

**Syntax**
**set wins** [**name=**]*InterfaceName* [**source=**]{**dhcp** | **static** [**addr=**]{*WINSAddress* | **none** }}

**Parameters**
**[ name=] *InterfaceName***
Required. Specifies the name of the interface for which you want to set WINS information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ source=]{ dhcp| static[ addr=]{ *WINSAddress*| none}}**
Required. Specifies whether the IP address of the WINS server to configure should be assigned by DHCP or is static. If the IP address is static, *WINSAddress* specifies the IP address of the WINS server to configure, and **none** specifies that the WINS configuration should be removed.

**/?**
Displays help at the command prompt.

**Remarks**

- If the interface is already statically configured, the **static** parameter replaces the existing WINS server address list with the one specified in the **set wins** command.

**add wins**
Adds a WINS server to a list of WINS servers for a specified interface.

**Syntax**
**add wins** [**name=**]*InterfaceName* [**addr=**] *WINSAddress* [[**index=**]*WINSIndex*]

**Parameters**
**[ name=] *InterfaceName***
Required. Specifies the name of the interface for which you want to add WINS information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ addr=] *WINSAddress***
Required. Specifies the IP address of the WINS server to add.

**[ index=] *WINSIndex***
Specifies the position of the added WINS server in the WINS server list for that interface.

**/?**

Displays help at the command prompt.

**Remarks**

- A lower number for the *WINSIndex* parameter corresponds to a higher preference. If no index is specified, then the server is added with the lowest preference.

**delete wins**

Deletes a WINS server or servers from a list of WINS servers for a specified interface or all interfaces.

**Syntax**

**delete wins** [**name=**]*InterfaceName* [**addr=**]{*WINSAddress* | **all**}

**Parameters**

**[ name=] *InterfaceName***

Required. Specifies the name of the interface for which you want to delete a WINS server or servers. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ addr=]{ *WINSAddress*| all}**

Required. Specifies whether to delete only one server for an interface or all servers for all interfaces. If only one server should be deleted, *WINSAddress* specifies the IP address of the WINS server to delete.

**/?**

Displays help at the command prompt.

**show wins**

Displays the WINS configuration for a specified interface. Used without parameters, **show wins** displays the WINS configuration for all interfaces.

**Syntax**

**show wins** [[**name=**]*InterfaceName*]

**Parameters**

**[ name=] *InterfaceName***

Specifies the name of the interface whose WINS information you want to display. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**/?**

Displays help at the command prompt.

**show icmp**

Displays ICMP statistics. Used without parameters, **show icmp** displays the statistics only once.

**Syntax**

**show icmp** [[**rr=**]*RefreshRate*]

**Parameters**

**[ rr=] *RefreshRate***

Specifies the refresh rate (the number of seconds between refreshing the display of the statistics).

**/?**

Displays help at the command prompt.

**show interface**

Displays statistics for a specified interface. Used without parameters, **show interface** displays statistics for all interfaces only once.

**Syntax**

**show interface** [[**index=**]*InterfaceIndex*] [[**rr=**]*RefreshRate*]

**Parameters**

**[ index=] *InterfaceIndex***

Specifies the interface index, an integer that identifies the interface.

**[ rr=] *RefreshRate***

Specifies the refresh rate (the number of seconds between refreshing the display of the statistics).

**/?**
Displays help at the command prompt.

**Remarks**

- To obtain the interface index for an interface, review the output of the **show interface** command used without parameters.

**show ipaddress**
Displays information for a specified IP address. Used without parameters, **show ipaddress** displays information for all IP addresses on all interfaces once.

**Syntax**
**show ipaddress** [[**index=**]*IPAddress*] [[**rr=**]*RefreshRate*]

**Parameters**
**[ index=] *IPAddress***
Specifies an IP address of an interface.

**[ rr=] *RefreshRate***
Specifies the refresh rate (the number of seconds between refreshing the display of the statistics).

**/?**
Displays help at the command prompt.

**Remarks**

- To obtain the IP addresses of all interfaces, review the display of the **show ipaddress** command used without parameters.

**show ipnet**
Displays the contents of the Address Resolution Protocol (ARP) cache, which contains the hardware addresses of resolved next-hop IP addresses. Used without parameters, **show ipnet** displays the information once.

**Syntax**
**show ipnet** [[**rr=**]*RefreshRate*]

**Parameters**
**[ rr=] *RefreshRate***
Specifies the refresh rate (the number of seconds between refreshing the display of the statistics).

**/?**
Displays help at the command prompt.

**show ipstats**
Displays IP statistics. Used without parameters, **show ipstats** displays the statistics once.

**Syntax**
**show ipstats** [[**rr=**]*RefreshRate*]

**Parameters**
**[ rr=] *RefreshRate***
Specifies the refresh rate (the number of seconds between refreshing the display of the statistics).

**/?**
Displays help at the command prompt.

**show joins**
Displays IP multicast groups that have been joined for the specified IP address. Used without parameters, **show joins** displays information for all IP addresses.

**Syntax**
**show joins** [[**index=**]*IPAddress*]

**Parameters**

**[ index=] *IPAddress***
Specifies an IP address of an interface.

**/?**
Displays help at the command prompt.

**Remarks**

- To obtain the IP addresses for all interfaces, review the display of the **show ipaddress** command used without parameters.

**show tcpconn**
Displays information on a specified TCP connection. Used without parameters, **show tcpconn** displays information for all TCP connections once.

**Syntax**
**show tcpconn** [[**index=**]{*LocalIPAddress* | *LocalPort* | *RemoteIPAddress* | *RemotePort*}] [[**rr=**]*RefreshRate*]

**Parameters**
**[ index=]{ *LocalIPAddress*| *LocalPort*| *RemoteIPAddress*| *RemotePort*}**
Specifies the connection about which to display information. The *LocalIPAddress* parameter specifies an IP address of an interface. The *LocalPort* parameter specifies a TCP port for a local process. The *RemoteIPAddress* parameter specifies an IP address of a remote host. The *RemotePort* parameter specifies a TCP port for a remote process.

**[ rr=] *RefreshRate***
Specifies the refresh rate (the number of seconds between refreshing the display of the information).

**/?**
Displays help at the command prompt.

**show tcpstats**
Displays TCP statistics. Used without parameters, **show tcpstats** displays the statistics once.

**Syntax**
**show tcpstats** [[**rr=**]*RefreshRate*]

**Parameters**
**[ rr=] *RefreshRate***
Specifies the refresh rate (the number of seconds between refreshing the display of the statistics).

**/?**
Displays help at the command prompt.

**show udpconn**
Displays information about the UDP ports used for each IP address. Used without parameters, **show udpconn** displays UDP port information for all IP addresses once.

**Syntax**
**show udpconn** [[**index=**]{*LocalIPAddress* | *LocalPort*}] [[**rr=**]*RefreshRate*]

**Parameters**
**[ index=]{ *LocalIPAddress*| *LocalPort*}**
Specifies the connection about which to display information. The *LocalIPAddress* parameter specifies an IP address of an interface. The *LocalPort* parameter specifies a UDP port for a local process.

**[ rr=] *RefreshRate***
Specifies the refresh rate (the number of seconds between refreshing the display of the statistics).

**/?**
Displays help at the command prompt.

**show udpstats**
Displays UDP statistics. Used without parameters, **show udpstats** displays the statistics once.

**Syntax**
**show udpstats** [[**rr=**]*RefreshRate*]

**Parameters**

**[ rr=] *RefreshRate***
Specifies the refresh rate (the number of seconds between refreshing the display of the statistics).

**/?**
Displays help at the command prompt.

**show config**
Displays IP address and other configuration information for a specified interface. Used without parameters, **show config** displays configuration information for all interfaces.

**Syntax**
**show config** [[**name=**]*InterfaceName*]

**Parameters**
**[ name=] *InterfaceName***
Specifies the name of the interface for which you want to display configuration information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**/?**
Displays help at the command prompt.

**show offload**
Displays the tasks that can be performed by the network adapter for the specified interface corresponding to installed network hardware. Used without parameters, **show offload** displays offload information for all interfaces corresponding to installed network hardware.

**Syntax**
**show offload** [[**name=**]*InterfaceName* ]

**Parameters**
**[ name=] *InterfaceName***
Specifies the name of the interface for which you want to display offload information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**/?**
Displays help at the command prompt.

**delete arpcache**
Removes the entries in the Address Resolution Protocol (ARP) cache for a specified interface. Used without parameters, **delete arpcache** removes the entries in the ARP caches of all interfaces.

**Syntax**
**delete arpcache** [[**name=**]*InterfaceName*]

**Parameters**
**[ name=] *InterfaceName***
Specifies the name of the interface for which you want to remove the ARP cache entries. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**/?**
Displays help at the command prompt.

**dump**
Displays the current configuration as a series of Netsh Interface IP commands.

**Syntax**
**dump**

**Parameters**
none

**Remarks**

- To create a configuration script file, redirect the output of this command to a file.

**Netsh Interface IP example**

The following command configures the interface named **Local Area Connection** with the static IP address 10.0.5.99, the subnet mask of 255.255.255.0, and a default gateway of 10.0.5.1:

```
set address name="Local Area Connection" source=static addr=10.0.5.99 mask=255.255.255.0 gateway=10.0.5.1
gwmetric=1
```

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Netsh commands for AAAA
Netsh commands for DHCP
Netsh diagnostic (diag) commands
Netsh commands for remote access (ras)
Netsh commands for WINS
Command-line reference A-Z
Command shell overview

# Netsh commands for Interface IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for Interface IPv6**
The Netsh commands for Interface IPv6 provide a command-line tool that you can use to query and configure IPv6 interfaces, address, caches, and routes.

In addition, the Interface IPv6 context of **netsh** has a subcontext for 6to4. You can use the commands in the **netsh interface IPv6 6to4** context to configure or display the configuration of the 6to4 service on either a 6to4 host or a 6to4 router.

You can run these commands from the Windows Server™ 2003 family command prompt or from the command prompt for the **netsh interface IPv6** context. For these commands to work at the Windows Server™ 2003 family command prompt, you must type **netsh interface ipv6** before typing commands and parameters as they appear in the syntax below. There might be functional differences between **netsh** context commands on Windows 2000 and the Windows Server™ 2003 family.

To view help for a command at the command prompt, type *CommandName***/?**, where *CommandName* is the name of the command.

For more information about **netsh**, see Netsh overview and Enter a netsh context.

To view the command syntax, click a command:

- 6to4
- add 6over4tunnel
- add address
- add dns
- add prefixpolicy
- add route
- add v6v4tunnel
- delete address
- delete destinationcache
- delete dns
- delete interface
- delete neighbors
- delete prefixpolicy
- delete route
- dump
- install
- isatap

- renew
- reset
- set address
- set global
- set interface
- set mobility
- set prefixpolicy
- set privacy
- set route
- set state
- show address
- show bindingcacheentries
- show destinationcache
- show dns
- show global
- show interface
- show joins
- show mobility
- show neighbors
- show prefixpolicy
- show privacy
- show routes
- show siteprefixes
- uninstall

**6to4**
Specifies that the **6to4** context of **netsh interface IPv6 6to4** is used.

**Syntax**
**6to4**

**add 6over4tunnel**
Creates a 6over4 interface using the specified IPv4 address.

**Syntax**

**add 6over4tunnel** [[**interface=**]*String*] [**localaddress=**]*IPv4Address* [[**store=**]{**active** | **persistent**}]

**Parameters**

**[[ interface=] *String*]**
Specifies an interface name or index.

**[ localaddress=] *IPv4Address***
Required. Specifies the IPv4 address that is encapsulated.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**

This example command creates a 6over4 interface using the IPv4 address 10.1.1.1 on the interface named "Private."

**add 6over4tunnel "Private" 10.1.1.1**

**add address**

Adds an IPv6 address to a specified interface. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s). For example, 2d represents two days.

**Syntax**

**add address** [[**interface=**]*String*] [**address=**]*IPv6Address* [[**type=**]{**unicast** | **anycast**}] [[**validlifetime=**]{*Integer* | **infinite**}] [[**preferredlifetime=**]{*Integer* | **infinite**}] [[**store=**]{**active** | **persistent**}]

**Parameters**

**[[ interface=] *String*]**
Specifies an interface name or index.

**[ address=] *IPv6Address***
Required. Specifies the IPv6 address to add.

**[[ type=]{ unicast| anycast}]**
Specifies whether a unicast address (**unicast**) or an anycast address (**anycast**) is added. The default selection is **unicast**.

**[[ validlifetime=]{ *Integer*| infinite}]**
Specifies the lifetime over which the address is valid. The default value is **infinite**.

**[[ preferredlifetime=]{ *Integer*| infinite}]**
Specifies the lifetime over which the address is preferred. The default value is **infinite**.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**

This example command adds the IPv6 address FE80::2 to the interface named "Private."

**add address "Private" FE80::2**

**add dns**

Adds a new DNS server IP address to the statically-configured list of DNS servers for the specified interface.

**Syntax**

**add dns** [**interface=**]*String* [**address=**]*IPAddress* [[**index=**]*Integer*]

**Parameters**

**[ interface=] *String***
Required. Specifies, by name, which interface will have a DNS server IP address added to its list of DNS server IP addresses.

**[ address=] *IPAddress***
Required. Specifies the IPv6 address of the DNS server to add to the list.

**[[ index=] *Integer*]**
Specifies the position on the statically-configured list in which to place the DNS server IP address specified in **address**. By default, the DNS server IP address is added to the end of the list.

**Remarks**
If an index is specified, the DNS server is placed in that position in the list.

**Examples**
In the first example command, a DNS server with the IPv6 address FEC0:0:0:FFFF::1 is added to the list of DNS server IP addresses for the interface named "Local Area Connection." In the second example, a DNS server with the IPv6 address FEC0:0:0:FFFF::2 is added at index 2 as the second server on the list of servers for the interface named "Local Area Connection".

**add dns "Local Area Connection" FEC0:0:0:FFFF::1**

**add dns "Local Area Connection" FEC0:0:0:FFFF::2 index=2**

**add prefixpolicy**
Adds a source and destination address selection policy for a specified prefix.

**Syntax**
**add prefixpolicy** [**prefix=**]*IPv6Address***/***Integer* [**precedence=**]*Integer* [**label=**]*Integer* [[**store=**]{**active** | **persistent**}]

**Parameters**
**[ prefix=] *IPv6Address/Integer***
Required. Specifies the prefix for which to add a policy in the policy table. *Integer* specifies the prefix length.

**[ precedence=] *Integer***
Required. Specifies the precedence value used for sorting destination addresses in the policy table.

**[ label=] *Integer***
Required. Specifies the label value that allows for policies that require a specific source address prefix for use with a destination address prefix.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
This example command adds a prefix policy for prefix ::/96, with a precedence value of 3 and a label value of 4.

**add prefixpolicy ::/96 3 4**

**add route**
Adds a route for a specified prefix. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s). For example, 2d represents two days.

**Syntax**
**add route** [**prefix=**]*IPv6Address***/***Integer* [[**interface=**]*String*] [[**nexthop=**]*IPv6Address*] [[**siteprefixlength=**]*Integer*] [[**metric=**]*Integer*] [[**publish=**]{**no** | **yes** | **immortal**}] [[**validlifetime=**]{*Integer* | **infinite**}] [[**preferredlifetime=**]{*Integer* | **infinite**}] [[**store=**]{**active** | **persistent**}]

**Parameters**
**[ prefix=] *IPv6Address/Integer***
Required. Specifies the prefix for which to add a route. *Integer* specifies the prefix length.

**[[ interface=] *String*]**
Specifies an interface name or index.

**[[ nexthop=] *IPv6Address*]**
Specifies the gateway address, if the prefix is not on-link.

**[[ siteprefixlength=] *Integer*]**
Specifies the prefix length for the entire site, if the prefix is not on-link.

**[[ metric=] *Integer*]**
Specifies the route metric.

**[[ publish=]{ no| yes| immortal}]**
Specifies whether routes are advertised (**yes**), advertised with an infinite lifetime (**immortal**), or not advertised (**no**) in Route Advertisements. The default selection is **no**.

**[[ validlifetime=]{ *Integer*| infinite}]**

Specifies the lifetime over which the route is valid. The default value is **infinite**.

**[[ preferredlifetime=]{ *Integer*| infinite}]**
Specifies the lifetime over which the route is preferred. The default value is **infinite**.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
This example command adds a route on the interface named "Internet" with a prefix of 3FFE:: and a prefix length of 16 bits (3FFE::/16). The **nexthop** value is FE80::1.

**add route 3FFE::/16 "Internet" FE80::1**

**add v6v4tunnel**
Creates an IPv6-in-IPv4 tunnel.

**Syntax**
**add v6v4tunnel** [[**interface=**]*String*] [**localaddress=**]*IPv4Address* [**remoteaddress=**]*IPv4Address* [[**neighbordiscovery=**]{**enabled** | **disabled**}] [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ interface=] *String*]**
Specifies an interface name or index.

**[ localaddress=] *IPv4Address***
Required. Specifies the IPv4 address of the local tunnel endpoint.

**[ remoteaddress=] *IPv4Address***
Required. Specifies the IPv4 address of the remote tunnel endpoint.

**[[ neighbordiscovery=]{ enabled| disabled}]**
Specifies whether Neighbor Discovery is enabled (**enabled**) or disabled (**disabled**) on the interface. The default selection is **disabled**.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
This example command creates an IPv6-in-IPv4 tunnel between the local address 10.0.0.1 and the remote address 192.168.1.1 on the interface "Private."

**add v6v4tunnel "Private" 10.0.0.1 192.168.1.1**

**delete address**
**Syntax**
**delete address** [[**interface=**]*String*] [**address=**]*IPv6Address* [[**store=**]{**active** | **persistent**}]

Modifies an IPv6 address on a specified interface.

**Parameters**
**[[ interface=] *String*]**
Specifies an interface name or index.

**[ address=] *IPv6Address***
Required. Specifies the IPv6 address to delete.

**[[ store=]{ active| persistent}]**
Specifies whether the deletion lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
This example command deletes the address FE80::2 from the interface named "Private".

**delete address "Private" FE80::2**

**delete destinationcache**

Clears the destination cache. If an interface is specified, clears the cache only on that interface. If an address is also specified, deletes only that destination cache entry.

**Syntax**

**delete destinationcache** [[**interface=**]*String*] [[**address=**]*IPv6Address*]

**Parameters**

**[[ interface=] *String*]**

Specifies an interface name or index.

**[[ address=] *IPv6Address*]**

Specifies the address of the destination.

**Remarks**

When no parameters are specified, all entries in the destination caches for all interfaces are deleted.

**Examples**

This example command deletes the destination cache for the interface named "Private."

**delete destinationcache "Private"**

**delete dns**

Deletes statically configured DNS server IPv6 addresses for a specific interface.

**Syntax**

**delete dns** [**interface=**]*String* [[**address=**]{*IPv6Address* | **all**}]

**Parameters**

**[ interface=] *String***

Required. Specifies the interface, by name, for which you want to remove a DNS server from the list of DNS servers.

**[[ address=]{ *IPv6Address*| all}]**

Specifies the DNS server IPv6 address to delete. If **all** is specified, all DNS server IPv6 addresses on the list for the interface are deleted.

**Examples**

In the first example command, the DNS server IPv6 address FEC0:0:0:FFFF::1 is deleted from the list of addresses for the connection named "Local Area Connection." In the second example command, all DNS server IPv6 addresses are deleted for the connection named "Local Area Connection."

**delete dns "Local Area Connection" FEC0:0:0:FFFF::1**

**delete dns "Local Area Connection" all**

**delete interface**

Deletes a specified interface from the IPv6 stack.

**Syntax**

**delete interface** [[**interface=**]*String*] [[**store=**]{**active** | **persistent**}]

**Parameters**

**[[ interface=] *String*]**

Specifies an interface name or index.

**[[ store=]{ active| persistent}]**

Specifies whether the deletion lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**

This example command deletes the interface named "Private."

**delete interface "Private"**

**delete neighbors**

Specifies that all entries in the neighbor cache are deleted. If an interface is specified, clears the cache only on that interface. If an address is also specified, deletes only that neighbor cache entry.

**Syntax**

**delete neighbors** [[**interface=**]*String*] [[**address=**]*IPv6Address*]

**Parameters**

**[[ interface=] *String*]**
Specifies an interface name or index.

**[[ address=] *IPv6Address*]**
Specifies the address of the neighbor.

**Examples**

This example command removes all entries from the neighbor cache on the interface named "Private."

**delete neighbors "Private"**

**delete prefixpolicy**

Deletes the source and destination address selection policy for a specified prefix.

**Syntax**

**delete prefixpolicy** [**prefix=**]*IPv6Address***/***Integer* [[**store=**]{**active** | **persistent**}]

**Parameters**

**[ prefix=] *IPv6Address/Integer***
Required. Specifies the prefix (*IPv6Address*) and prefix length (*Integer*) to delete from the policy table.

**[[ store=]{ active| persistent}]**
Specifies whether the deletion lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**

This example command deletes the prefix ::/96 from the policy table.

**delete prefixpolicy ::/96**

**delete route**

Deletes an IPv6 route.

**Syntax**

**delete route** [**prefix=**]*IPv6Address***/***Integer* [[**interface=**]*String*] [[**nexthop=**]*IPv6Address*] [[**store=**]{**active** | **persistent**}]

**Parameters**

**[ prefix=] *IPv6Address/Integer***
Required. Specifies the prefix of the route to delete.

**[[ interface=] *String*]**
Specifies an interface name or index.

**[[ nexthop=] *IPv6Address*]**
Specifies the gateway address, if the prefix is not on-link.

**[[ store=]{ active| persistent}]**
Specifies whether the deletion lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**

This example command deletes the route with the prefix 3FFE::/16 and the gateway FE80::1 from the interface named "Internet."

**delete route 3FFE::/16 "Internet" FE80::1**

**dump**

Dumps the network adapter IPv6 configuration to the command prompt window when run within the netsh context. When used in a batch file or script, output can be saved in a text file.

**Syntax**

**netsh interface ipv6 dump > [ *PathAndFileName* ]**

**Parameters**

**[ *PathAndFileName*]**
Specifies both the location where to the file is saved and the name of the destination file to which the configuration is saved.

**Remarks**

- After file output is obtained, you can use the **netsh exec** command to either configure another computer with the same IPv6 configuration or to restore the original configuration on the same computer.

- All IPv6 configuration information is saved with the dump command. For example, if an ISATAP or a 6to4 configuration is defined on an interface, the **dump** command saves these settings in the text file.

**Examples**
In the first example, the command is run manually at the **netsh interface ipv6** context of a command prompt. The IPv6 configuration is displayed in the command prompt window, and can be copied and pasted into a text file. In the second example, the **dump** command is run in a batch file, and the configuration is saved to a text file named Ipv6_conf.txt at the location C:\Temp.

**dump**

**netsh interface ipv6 dump > C:\temp\ipv6_conf.txt**

**install**
Installs IPv6.

**Syntax**
**install**

**isatap**
Specifies that the **isatap** context of **netsh interface IPv6 isatap** is used.

**Syntax**
**isatap**

**Remarks**

- Intrasite Automatic Tunnel Addressing Protocol (ISATAP) is used for communication between IPv6 and IPv4 nodes within an IPv4 site. It is described in the Internet draft titled "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)" (draft-ietf-ngtrans-isatap-*00*.txt). For more information, see IPv6 traffic between nodes on different subnets of an IPv4 internetwork.

**renew**
Restarts IPv6 interfaces.

**Syntax**
**renew** [[**interface=**]*String*]

**Parameters**
**[[ interface=] *String*]**
Specifies an interface name or index.

**Examples**
**renew "Private"**

**reset**
Resets the IPv6 configuration state.

**Syntax**
**reset**

**set address**
Modifies an IPv6 address on a specified interface. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s). For example, 2d represents two days.

**Syntax**
**set address** [[**interface=**]*String*] [**address=**]*IPv6Address* [[**type=**]{**unicast** | **anycast**}] [[**validlifetime=**]{*Integer* | **infinite**}] [[**preferredlifetime=**]{*Integer* | **infinite**}] [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ interface=] *String*]**
Specifies an interface name or index.

**[ address=] *IPv6Address***
Required. Specifies the IPv6 address to modify.

**[[ type=]{ unicast| anycast}]**
Specifies whether the address is marked as a unicast address (**unicast**) or as an anycast address (**anycast**). The default selection is **unicast**.

**[[ validlifetime=]{ *Integer*| infinite}]**
Specifies the lifetime over which the address is valid. The default value is **infinite**.

**[[ preferredlifetime=]{ *Integer*| infinite}]**
Specifies the lifetime over which the address is preferred. The default value is **infinite**.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
This example command sets the address FE80::2 on the interface named "Private" as an anycast address.

**set address "Private" FE80::2 anycast**

**set global**
Modifies global configuration parameters.

**Syntax**
**set global** [[**defaultcurhoplimit=**]*Integer*] [**neighborcachelimit=**]*Integer* [[**routecachelimit=**]*Integer*] [[**reassemblylimit=**]*Integer*] [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ defaultcurhoplimit=] *Integer*]**
Specifies the default hop limit of packets sent.

**[ neighborcachelimit=] *Integer***
Required. Specifies the maximum number of neighbor cache entries.

**[[ routecachelimit=] *Integer*]**
Specifies the maximum number of route cache entries.

**[[ reassemblylimit=] *Integer*]**
Specifies the maximum size of the reassembly buffer.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
This example command sets global parameters for all IPv6-enabled interfaces on the computer. The default hop limit is set to 32, the maximum number of neighbor cache entries is set to 100, and the maximum number of route cache entries is 100,000.

**set global 32 100 100000**

**set interface**
Modifies interface configuration parameters.

**Syntax**
**set interface** [[**interface=**]*String*] [[**forwarding=**]{**enabled** | **disabled**}] [[**advertise=**]{**enabled** | **disabled**}] [[**mtu=**]*Integer*] [[**siteid=**]*Integer*] [[**metric=**]*Integer*] [[**firewall=**]{**enabled** | **disabled**}] [[**siteprefixlength=**]*Integer*] [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ interface=] *String*]**
Specifies an interface name or index.

**[[ forwarding=]{ enabled| disabled}]**
Specifies whether packets arriving on this interface can be forwarded to other interfaces. The default selection is **disabled**.

**[[ advertise=]{ enabled| disabled}]**

Specifies whether Router Advertisements are sent on this interface. The default selection is **disabled**.

**[[ mtu=]** ***Integer*****]**
Specifies the MTU of this interface. The default MTU is the natural MTU of the link.

**[[ siteid=]** ***Integer*****]**
Specifies the site scope zone identifier.

**[[ metric=]** ***Integer*****]**
Specifies the interface metric, which is added to route metrics for all routes over the interface.

**[[ firewall=]{ enabled| disabled}]**
Specifies whether to operate in firewall mode.

**[[ siteprefixlength=]** ***Integer*****]**
Specifies the default length of the global prefix for the entire site.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
This example command sets the interface with the name "Private," with a siteid of 2 and a metric of 2. All other parameter values are left at the default values.

**set interface "Private" siteid=2 metric=2**

**set mobility**
Modifies mobility configuration parameters.

**Syntax**
**set mobility** [[**security=**]{**enabled** | **disabled**}] [[**bindingcachelimit=**]*Integer*] [[**correspondentnode=**]**enabled** | **disabled**] [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ security=]{ enabled| disabled}]**
Specifies whether binding updates must be secured.

**[[ bindingcachelimit=]** ***Integer*****]**
Specifies the maximum number of binding cache entries.

**[[ correspondentnode=] enabled| disabled]**
Specifies whether Correspondent Node functionality is set to **enabled** or the default of **disabled**.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
**set mobility security=disabled bindingcachelimit=1000 corr=enabled**

**set prefixpolicy**
Modifies a source and destination address selection policy for a specified prefix.

**Syntax**
**set prefixpolicy** [**prefix=**]*IPv6Address***/***Integer* [**precedence=**]*Integer* [**label=**]*Integer* [[**store=**]{**active** | **persistent**}]

**Parameters**
**[ prefix=]** ***IPv6Address/Integer***
Required. Specifies the prefix for which to add a policy in the policy table. *Integer* specifies the prefix length.

**[ precedence=]** ***Integer***
Required. Specifies the precedence value used for sorting destination addresses in the policy table.

**[ label=]** ***Integer***
Required. Specifies the label value that allows for policies that require a specific source address prefix for use with a destination address prefix.

**[[ store=]{ active| persistent}]**

Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

#### Examples

This example command sets a policy in the policy table for the prefix ::/96, with a precedence value of 3 and a label value of 4.

**set prefixpolicy ::/96 3 4**

### set privacy

Modifies parameters related to temporary address generation. If **randomtime=** is specified, **maxrandomtime=** is not used. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s). For example, 2d represents two days.

#### Syntax

**set privacy** [[**state=**]{**enabled** | **disabled**}] [[**maxdadattempts=**]*Integer*] [[**maxvalidlifetime=**]*Integer*] [[**maxpreferredlifetime=**]*Integer*] [[**regeneratetime=**]*Integer*] [[**maxrandomtime=**]*Integer*] [[**randomtime=**]*Integer*] [[**store=**]{**active** | **persistent**}]

#### Parameters

**[[ state=]{ enabled| disabled}]**
Specifies whether temporary addresses are enabled.

**[[ maxdadattempts=] *Integer*]**
Specifies the number of duplicate address detection attempts made. The default value is 5.

**[[ maxvalidlifetime=] *Integer*]**
Specifies the maximum lifetime over which a temporary address is valid. The default value is 7d (seven days).

**[[ maxpreferredlifetime=] *Integer*]**
Specifies the maximum lifetime over which an anonymous is preferred. The default value is 1d (one day).

**[[ regeneratetime=] *Integer*]**
Specifies the duration of time that elapses when a new address is generated prior to deprecating a temporary address. The default value is 5s (five seconds).

**[[ maxrandomtime=] *Integer*]**
Specifies the upper limit to use when computing a random delay at boot. The default value is 10m (ten minutes).

**[[ randomtime=] *Integer*]**
Specifies a time value to use, instead of a value generated at boot.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

### set route

Modifies route parameters. Time values can be expressed in days (d), hours (h), minutes (m), and seconds (s). For example, 2d represents two days.

#### Syntax

**set route** [**prefix=**]*IPv6Address***/***Integer* [[**interface=**]*String*] [[**nexthop=**]*IPv6Address*] [[**siteprefixlength=**]*Integer*] [[**metric=**]*Integer*] [**publish=**]{**no** | **yes** | **immortal**}] [[**validlifetime=**]{*Integer* | **infinite**}] [[**preferredlifetime=**]{*Integer* | **infinite**}] [[**store=**]{**active** | **persistent**}]

#### Parameters

**[ prefix=] *IPv6Address/Integer***
Required. Specifies the prefix (*IPv6Address*) and prefix length (*Integer*) of the route to modify.

**[[ interface=] *String*]**
Specifies an interface name or index.

**[[ nexthop=] *IPv6Address*]**
Specifies the gateway address, if the prefix is not on-link.

**[[ siteprefixlength=] *Integer*]**
Specifies the prefix length for the entire site, if the prefix is not on-link.

**[[ metric=] *Integer*]**
Specifies the route metric.

**[[ publish=]{ no| yes| immortal}]**
Specifies whether routes are advertised (**yes**), advertised with an infinite lifetime (**immortal**), or not advertised (**no**) in Route Advertisements. The default selection is **no**.

**[[ validlifetime=]{ *Integer*| infinite}]**
Specifies the lifetime over which the route is valid. The default value is **infinite**.

**[[ preferredlifetime=]{ *Integer*| infinite}]**
Specifies the lifetime over which the route is preferred. The default value is **infinite**.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Examples**
This example command sets a route on the interface named "Internet." The route prefix is 3FFE::, and has a length of 16 bits. The gateway address, defined by the **nexthop=** parameter, is FE80::1.

**set route 3FFE::/16 "Internet" FE80::1**

**set state**
Enables or disables IPv4 compatibility. The default value for all parameters is disabled.

**Syntax**
**set state** [[**6over4=**]{**enabled** | **disabled** | **default**}] [[**v4compat=**]{**enabled** | **disabled** | **default**}]

**Parameters**
**[[6over4=]{enabled| disabled| default}]**
Specifies whether 6over4 interfaces are created. To both disable and delete 6over4 compatible interfaces, specify **default**. To disable 6over4 compatible interfaces without deleting them, specify **disabled**.

**[[ v4compat=]{ enabled| disabled| default}]**
Specifies whether IPv4 compatible interfaces are created. To both disable and delete IPv4 compatible interfaces, specify **default**. To disable IPv4 compatible interfaces without deleting them, specify **disabled**.

**Examples**
In the first example command, IPv4-compatible addresses are disabled, and any previously existing interfaces are deleted. In the second example command, IPv4-compatible addresses are enabled

**set state default**

**set state 6over4=disabled v4compat=enabled**

**show address**
Displays all IPv6 addresses, or all addresses on a specified interface.

**Syntax**
**show address** [[**interface=**]*String*] [[**level=**]{**normal** | **verbose**}] [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ interface=] *String*]**
Specifies an interface name or index.

**[[ level=]{ normal| verbose}]**
Specifies whether one line per interface is displayed (**normal**) or additional information is displayed for each interface (**verbose**). When no interface is specified, the default selection is **normal**. When an interface is specified, the default selection is **verbose**.

**[[ store=]{ active| persistent}]**
Specifies whether active (**active**) or persistent (**persistent**) addresses are displayed. The default selection is **active**.

**show bindingcacheentries**
Displays all binding cache entries.

**Syntax**
**show bindingcacheentries**

**show destinationcache**
Displays destination cache entries. If an interface is specified, displays the cache only on that interface. If an address is also

specified, displays only that destination cache entry.

**Syntax**

**show destinationcache** [[**interface=**]*String*] [[**address=**]*IPv6Address*]

**Parameters**

**[[ interface=] *String*]**

Specifies an interface name or index.

**[[ address=] *IPv6Address*]**

Specifies the address of the destination.

**show dns**

Displays the DNS server configuration for a specific interface or interfaces.

**Syntax**

**show dns** [[**interface=**]*String*]

**Parameters**

**[[ interface=] *String*]**

Specifies the interface, by name, for which you want to display configured DNS server IPv6 addresses. If no interface is specified, servers for all interfaces are displayed.

**Examples**

In this example command, DNS server IPv6 addresses configured on the "Local Area Connection" interface are displayed.

**show dns "Local Area Connection"**

**show global**

Displays global configuration parameters.

**Syntax**

**show global** [[**store=**]{**active** | **persistent**}]

**Parameters**

**[[ store=]{ active| persistent}]**

Specifies whether active (**active**) or persistent (**persistent**) information is displayed. The default selection is **active**.

**show interface**

Displays information about all interfaces, or about a specified interface.

**Syntax**

**show interface** [[**interface=**]*String*] [[**level=**]{**normal** | **verbose**}] [[**store=**]{**active** | **persistent**}]

**Parameters**

**[[ interface=] *String*]**

Specifies an interface name or index.

**[[ level=]{ normal| verbose}]**

Specifies whether one line per interface is displayed (**normal**) or additional information is displayed for each interface (**verbose**). When no interface is specified, the default selection is **normal**. When an interface is specified, the default selection is **verbose**.

**[[ store=]{ active| persistent}]**

Specifies whether active (**active**) or persistent (**persistent**) interfaces are displayed. The default selection is **active**.

**show joins**

Displays all IPv6 multicast addresses, or all multicast addresses on a specified interface.

**Syntax**

**show joins** [[**interface=**]*String*] [[**level=**]{**normal** | **verbose**}]

**Parameters**

**[[ interface=] *String*]**

Specifies an interface name or index.

**[[ level=]{ normal| verbose}]**

Specifies whether one line per interface is displayed (**normal**) or additional information is displayed for each interface (**verbose**). When no interface is specified, the default selection is **normal**. When an interface is specified, the default

selection is **verbose**.

**show mobility**
Displays mobility configuration parameters.

**Syntax**
**show mobility** [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ store=]{ active| persistent}]**
Specifies whether active (**active**) or persistent (**persistent**) information is displayed. The default selection is **active**.

**show neighbors**
Displays neighbor cache entries. If an interface is specified, displays only the cache on that interface. If an address is also specified, displays only that neighbor cache entry.

**Syntax**
show neighbors [[**interface=**]*String*] [[**address=**]*IPv6Address*]

**Parameters**
**[[ interface=] *String*]**
Specifies an interface name or index.

**[[ address=] *IPv6Address*]**
Specifies the address of the neighbor.

**show prefixpolicy**
Displays prefix policy table entries used in source and destination address selection.

**Syntax**
**show prefixpolicy** [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ store=]{ active| persistent}]**
Specifies whether active (**active**) or persistent (**persistent**) information is displayed. The default selection is **active**.

**show privacy**
Displays privacy configuration parameters.

**Syntax**
**show privacy** [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ store=]{ active| persistent}]**
Specifies whether active (**active**) or persistent (**persistent**) information is displayed. The default selection is **active**.

**show routes**
Displays route table entries.

**Syntax**
**show routes** [[**level=**]{**normal** | **verbose**}] [[**store=**]{**active** | **persistent**}]

**Parameters**
**[[ level=]{ normal| verbose}]**
Specifies whether only normal routes (**normal**) or routes used for loopback (**verbose**) are displayed. The default selection is **normal**.

**[[ store=]{ active| persistent}]**
Specifies whether active (**active**) or persistent (**persistent**) routes are displayed. The default selection is **active**.

**show siteprefixes**
Displays the site prefix table.

**Syntax**
**show siteprefixes**

**uninstall**
Uninstalls IPv6.

**Syntax**

**uninstall**

**Netsh interface IPv6 6to4**

You can use the following commands in the **netsh interface IPv6 6to4** context to display the configuration of or configure the 6to4 service on either a 6to4 host or a 6to4 router. For more information about using the 6to4 service, see IPv6 traffic between nodes in different sites across the Internet (6to4).

For more information about Netsh, see Netsh overview.

- set interface
- show interface
- set relay
- show relay
- set routing
- show routing
- set state
- show state
- reset

**set interface**

Configures the 6to4 service on an interface.

**Syntax**

**set interface** [**name=**] *InterfaceName* [[**routing=**] {**enabled** | **disabled** | **default**}]

**Parameters**

**[ name=] *InterfaceName***

Required. Specifies the name of the interface for which you want to set 6to4 service configuration. *InterfaceName* must match the name of the interface specified in Network Connections. If *InterfaceName* contains any spaces, it must be enclosed in quotes.

**[[ routing=] { enabled| disabled| default}]**

Specifies whether the forwarding of 6to4 packets received on the interface is enabled, disabled, or set to its default value.

**Remarks**

- This command enables, disables, or sets to default the routing behavior of the 6to4 service on a specified interface.

The default setting for the **routing=** parameter is **enabled**, which enables routing on private interfaces if Internet Connection Sharing (ICS) is used.

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**show interface**

Displays the 6to4 service routing configuration on all interfaces, or on a specified interface.

**Syntax**

**show interface** [[**name=**] *InterfaceName*]

**Parameters**

**[[ name=] *InterfaceName*]**

Specifies the name of the interface for which you want to display the 6to4 service configuration. *InterfaceName* must match the name of the interface specified in Network Connections. If *InterfaceName* contains any spaces, it must be enclosed in quotes.

**Remarks**

- If an interface name is not specified, the 6to4 configuration for all interfaces is displayed.

**set relay**

Configures the name of the 6to4 relay router for the 6to4 service. Additionally, specifies how often the name is resolved and the state of the relay component for the 6to4 service.

**Syntax**

**set relay** [[**name=**] {*RelayDNSName* | **default**}] [[**state=**] {**enabled** | **disabled** | **automatic** | **default**}] [[**interval=**] {*ResInterval* | **default**}]

**Parameters**

**[[ name=] { *RelayDNSName*| default}]**

Specifies either the fully qualified domain name (FQDN) of a 6to4 relay router on the IPv4 Internet (*RelayDNSName*), or sets the relay name to its default value of 6to4.ipv6.microsoft.com (**default**).

**[[ state=] { enabled| disabled| automatic| default}]**

Specifies whether the state of the relay component for the 6to4 service is enabled, disabled, automatically enabled if a public IPv4 address is configured, or set to its default value.

**[[ interval=] { *ResInterval*| default}]**

Specifies how often the name of the relay router is resolved in minutes (*ResInterval*) or sets the resolution interval to its default value of 1440 minutes (**default**).

**Remarks**

- The 6to4 relay router is a router that provides an access point between the IPv4 Internet and the 6bone (the native IPv6 portion of the Internet). In order to access 6bone resources from a 6to4 router, the 6to4 router encapsulates 6to4 traffic with an IPv4 header and sends it to the IPv4 address of the 6to4 relay router. The 6to4 relay router removes the IPv4 header and forwards the traffic to the 6bone. For return traffic, the 6to4 relay router encapsulates IPv6 traffic and sends it to the 6to4 router at the 6to4 host's site.

- The default name of the 6to4 relay router is 6to4.ipv6.microsoft.com.

- The default state is **automatic**, which enables the forwarding of native IPv6 traffic to a relay router when a public IPv4 address is assigned to any interface.

- The default resolution interval is 1440 minutes (once each day).

**show relay**

Displays the relay router configuration for the 6to4 service.

**Syntax**

**show relay**

**set routing**

Sets both the state of routing and the inclusion of site-local address prefixes in Router Advertisements that are sent by the 6to4 router.

**Syntax**

**set routing** [[**routing=**] {**enabled** | **disabled** | **automatic** | **default**}] [[**sitelocals=**] {**enabled** | **disabled** | **default**}]

**Parameters**

**[[ routing=] { enabled| disabled| automatic| default}]**

Specifies whether the state of routing on a 6to4 router is enabled, disabled, automatically enabled if Internet Connection Sharing (ICS) is enabled, or set to its default value.

**[[ sitelocals=] { enabled| disabled| default}]**
Specifies whether the advertising of site-local address prefixes, in addition to 6to4 address prefixes, is enabled, disabled, or set to its default value.

**Remarks**

- The default setting for the **routing=** parameter is **automatic**, which enables routing on private interfaces when ICS is used.
- The default setting for the **sitelocals=** parameter is **enabled**, which enables the advertising of site-local prefixes when site-local addresses are configured on private interfaces.

### Note

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**show routing**
Displays the routing configuration of the 6to4 service.

**Syntax**
**show routing**

**set state**
Configures the state of the 6to4 service.

**Syntax**
**set state** [[**state=**] {**enabled** | **disabled** | **default**}] [[**undoonstop=**] {**enabled** | **disabled** | **default**}] [[**6over4=**] {**enabled** | **disabled** | **default**}]

**Parameters**
**[[ state=] { enabled| disabled| default}]**
Specifies whether the state of the 6to4 service is enabled, disabled, or set to its default value.

**[[ undoonstop=] { enabled| disabled| default}]**
Specifies whether the reversal of all automatic configuration that has been performed by the 6to4 service occurs when the service stops is enabled, disabled, or set to its default value.

**Remarks**

- The default setting for the **state=** parameter is **enabled**, which enables the 6to4 service.
- The default setting for the **undoonstop=** parameter is **enabled**, which reverses all automatic configuration performed by the 6to4 service when the service is stopped.

**show state**
Displays the state of the 6to4 service.

**Syntax**
**show state**

**reset**
Resets the 6to4 service.

**Syntax**
**reset**

**Netsh interface ipv6 isatap**
Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) is an address assignment and tunneling mechanism for communication between IPv6/IPv4 nodes within an IPv4 site. It is described in the Internet draft titled "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)" (draft-ietf-ngtrans-isatap-*00*.txt). You can use the following commands to configure the ISATAP router.

- set router
- show router

**set router**
Specifies the Intra-Site Automatic Tunneling Address Protocol (ISATAP) router information, including router name, state, and resolution interval.

**Syntax**
**set router** [[**name=**]{*String* | **default**}] [[**state=**]{**Enabled** | **Disabled** | **Default**}] [[**interval**]=*Integer*]

**Parameters**
**[[ name=]{ *String*| default}]**
Specifies whether the router is named with a string. If **default** is specified, the system reverts to using the default name.

**[[ state=]{ Enabled| Disabled| Default}]**
Specifies whether the ISATAP router relays packets between subnets.

**[[ interval]= *Integer*]**
Specifies the router resolution interval, in minutes. The default interval is 1440 (24 hours).

**Examples**
The following example command sets the router name to isatap, enables the router, and sets the resolution interval to 120 minutes:

**set router isatap enabled 120**

**show router**
Displays configuration information for the ISATAP router.

**Syntax**
**show router**

**Remarks**
This command displays the router name, the relay state, and the resolution interval.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Netsh commands for Interface Portproxy
Netsh overview
Command shell overview
IP Version 6

# Netsh commands for Interface Portproxy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for Interface Portproxy**
The Netsh Interface Portproxy commands provide a command-line tool for use in administering servers that act as proxies between IPv4 and IPv6 networks and applications. You can use these commands to establish proxy service in the following ways:

- IPv4-configured computer and application messages sent to other IPv4-configured computers and applications.
- IPv4-configured computer and application messages sent to IPv6-configured computers and applications.
- IPv6-configured computer and application messages sent to IPv4-configured computers and applications.
- IPv6-configured computer and application messages sent to other IPv6-configured computers and applications.

When writing batch files or scripts using these commands, each command must be preceded by **netsh interface portproxy**. For example, when using the **delete v4tov6** command to specify that the portproxy server delete an IPv4 port and address from the list of IPv4 addresses for which the server listens, the batch file or script must use the following syntax:

**netsh interface portproxy delete v4tov6 listenport=** {*Integer* | *ServiceName*} [[**listenaddress=**] {*IPv4Address*| *HostName*}] [[**protocol=**]**tcp**]

You can run these commands from the Windows Server 2003 family command prompt or from the command prompt for the Netsh interface portproxy context. For these commands to work at the Windows Server 2003 family command prompt, you must type **netsh interface portproxy** before typing commands and parameters as they appear in the syntax below. There might be functional differences between Netsh context commands on Windows 2000 and the Windows Server 2003 family.

For more information about **Netsh**, see Netsh overview and Enter a netsh context.

To view the command syntax, click a command:

- add v4tov4
- add v4tov6
- add v6tov4
- add v6tov6
- delete v4tov4
- delete v4tov6
- delete v6tov4
- delete v6tov6
- reset
- set v4tov4
- set v4tov6
- set v6tov4

- set v6tov6
- show all
- show v4tov4
- show v4tov6
- show v6tov4
- show v6tov6

#### add v4tov4

Specifies that the portproxy server listen for messages sent to a specific port and IPv4 address, and maps a port and IPv4 address to which to send the messages received after establishing a separate TCP connection.

#### Syntax

**add v4tov4 listenport=** {*Integer* | *ServiceName*} [[**connectaddress=**] {*IPv4Address* | *HostName*}] [[**connectport=**] {*Integer* | *ServiceName*}] [[**listenaddress=**] {*IPv4Address*| *HostName*}] [[**protocol=**]**tcp**]

#### Parameters

**listenport**

Required. Specifies the IPv4 port, by port number or service name, on which to listen.

**connectaddress**

Specifies the IPv4 address to which to connect. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**connectport**

Specifies the IPv4 port, by port number or service name, to which to connect. If **connectport** is not specified, the default is the value of **listenport** on the local computer.

**listenaddress**

Specifies the IPv4 address for which to listen. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**protocol**

Specifies the protocol to use. Currently, only Transmission Control Protocol (TCP) is supported.

**/?**

Displays help at the command prompt.

#### add v4tov6

Specifies that the portproxy server listen for messages sent to a specific port and IPv4 address, and maps a port and IPv6 address to which to send the messages received after establishing a separate TCP connection.

#### Syntax

**add v4tov6 listenport=** {*Integer* | *ServiceName*} [[**connectaddress=**] {*IPv6Address* | *HostName*}] [[**connectport=**] {*Integer* | *ServiceName*}] [[**listenaddress=**] {*IPv4Address*| *HostName*}] [[**protocol=**]**tcp**]

#### Parameters

**listenport**

Required. Specifies the IPv4 port, by port number or service name, on which to listen.

**connectaddress**

Specifies the IPv6 address to which to connect. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**connectport**

Specifies the IPv6 port, by port number or service name, to which to connect. If **connectport** is not specified, the default is the value of **listenport** on the local computer.

**listenaddress**

Specifies the IPv4 address on which to listen. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**protocol**
Specifies the protocol to use. Currently only TCP is supported.

**/?**
Displays help at the command prompt.

#### add v6tov4

Specifies that the portproxy server listen for messages sent to a specific port and IPv6 address, and maps a port and IPv4 address to which to send the messages received after establishing a separate TCP connection.

##### Syntax

**add v6tov4 listenport=** {*Integer* | *ServiceName*} [[**connectaddress=**] {*IPv4Address* | *HostName*}] [[**connectport=**] {*Integer* | *ServiceName*}] [[**listenaddress=**] {*IPv6Address*| *HostName*}] [[**protocol=**]**tcp**]

##### Parameters

**listenport**
Required. Specifies the IPv6 port, by port number or service name, on which to listen.

**connectaddress**
Specifies the IPv4 address to which to connect. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**connectport**
Specifies the IPv4 port, by port number or service name, to which to connect. If **connectport** is not specified, the default is the value of **listenport** on the local computer.

**listenaddress**
Specifies the IPv6 address on which to listen. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**protocol**
Specifies the protocol to use. Currently only TCP is supported.

**/?**
Displays help at the command prompt.

#### add v6tov6

Specifies that the portproxy server listen for messages sent to a specific port and IPv6 address, and maps a port and IPv6 address to which to send the messages received after establishing a separate TCP connection.

##### Syntax

**add v6tov6 listenport=** {*Integer* | *ServiceName*} [[**connectaddress=**] {*IPv6Address* | *HostName*}] [[**connectport=**] {*Integer* | *ServiceName*}] [[**listenaddress=**] {*IPv6Address*| *HostName*}] [[**protocol=**]**tcp**]

##### Parameters

**listenport**
Required. Specifies the IPv6 port, by port number or service name, on which to listen.

**connectaddress**
Specifies the IPv6 address to which to connect. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**connectport**
Specifies the IPv6 port, by port number or service name, to which to connect. If **connectport** is not specified, the default is the value of **listenport** on the local computer.

**listenaddress**
Specifies the IPv6 address on which to listen. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**protocol**
Specifies the protocol to use. Currently only TCP is supported.

**/?**
Displays help at the command prompt.

#### delete v4tov4

Specifies that the portproxy server delete an IPv4 address from the list of IPv4 ports and addresses for which the server listens.

##### Syntax

**delete v4tov4 listenport=** {*Integer* | *ServiceName*} [[**listenaddress=**] {*IPv4Address*| *HostName*}] [[**protocol=**]**tcp**]

##### Parameters

**listenport**

Required. Specifies the IPv4 port to delete.

**listenaddress**

Specifies the IPv4 address to delete. If an address is not specified, the default is the local computer.

**protocol**

Specifies the protocol to use. Currently only TCP is supported.

**/?**

Displays help at the command prompt.

#### delete v4tov6

Specifies that the portproxy server delete an IPv4 port and address from the list of IPv4 addresses for which the server listens.

##### Syntax

**delete v4tov6 listenport=** {*Integer* | *ServiceName*} [[**listenaddress=**] {*IPv4Address*| *HostName*}] [[**protocol=**]**tcp**]

##### Parameters

**listenport**

Required. Specifies the IPv4 port to delete.

**listenaddress**

Specifies the IPv4 address to delete. If an address is not specified, the default is the local computer.

**protocol**

Specifies the protocol to use. Currently only TCP is supported.

**/?**

Displays help at the command prompt.

#### delete v6tov4

Specifies that the portproxy server delete an IPv6 port and address from the list of IPv6 addresses for which the server listens.

##### Syntax

**delete v6tov4 listenport=** {*Integer* | *ServiceName*} [[**listenaddress=**] {*IPv6Address*| *HostName*}] [[**protocol=**]**tcp**]

##### Parameters

**listenport**

Required. Specifies the IPv6 port to delete.

**listenaddress**

Specifies the IPv6 address to delete. If an address is not specified, the default is the local computer.

**protocol**

Specifies the protocol to use. Currently only TCP is supported.

**/?**

Displays help at the command prompt.

#### delete v6tov6

Specifies that the portproxy server delete an IPv6 address from the list of IPv6 addresses for which the server listens.

##### Syntax

**delete v6tov6 listenport=** {*Integer* | *ServiceName*} [[**listenaddress=**] {*IPv6Address*| *HostName*}] [[**protocol=**]**tcp**]

##### Parameters

**listenport**

Required. Specifies the IPv6 port to delete.

**listenaddress**

Specifies the IPv6 address to delete. If an address is not specified, the default is the local computer.

**protocol**
Specifies the protocol to use. Currently only TCP is supported.

**/?**
Displays help at the command prompt.

#### reset
Resets the IPv6 configuration state.

#### Syntax
**reset**

#### set v4tov4
Modifies the parameter values of an existing entry on the portproxy server created with the **add v4tov4** command, or adds a new entry to the list that maps port/address pairs.

#### Syntax
**set v4tov4 listenport=** {*Integer* | *ServiceName*} [[**connectaddress=**] {*IPv4Address* | *HostName*}] [[**connectport=**] {*Integer* | *ServiceName*}] [[**listenaddress=**] {*IPv4Address*| *HostName*}] [[**protocol=**]**tcp**]

#### Parameters
**listenport**
Required. Specifies the IPv4 port, by port number or service name, on which to listen.

**connectaddress**
Specifies the IPv4 address to which to connect. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**connectport**
Specifies the IPv4 port, by port number or service name, to which to connect. If **connectport** is not specified, the default is the value of **listenport** on the local computer.

**listenaddress**
Specifies the IPv4 address for which to listen. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**protocol**
Specifies the protocol to use. Currently, only Transmission Control Protocol (TCP) is supported.

**/?**
Displays help at the command prompt.

#### Remarks

- You can change any parameter values for an existing portproxy server entry. If values are not supplied, no changes are made.

#### set v4tov6
Modifies the parameter values of an existing entry on the portproxy server created with the **add v4tov6** command, or adds a new entry to the list that maps port/address pairs.

#### Syntax
**set v4tov6 listenport=** {*Integer* | *ServiceName*} [[**connectaddress=**] {*IPv6Address* | *HostName*}] [[**connectport=**] {*Integer* | *ServiceName*}] [[**listenaddress=**] {*IPv4Address*| *HostName*}] [[**protocol=**]**tcp**]

#### Parameters
**listenport**
Required. Specifies the IPv4 port, by port number or service name, on which to listen.

**connectaddress**
Specifies the IPv6 address to which to connect. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**connectport**
Specifies the IPv6 port, by port number or service name, to which to connect. If **connectport** is not specified, the default is the value of **listenport** on the local computer.

**listenaddress**
Specifies the IPv4 address on which to listen. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**protocol**
Specifies the protocol to use. Currently only TCP is supported.

**/?**
Displays help at the command prompt.

**Remarks**

- You can change any parameter values for an existing portproxy server entry. If values are not supplied, no changes are made.

**set v6tov4**
Modifies the parameter values of an existing entry on the portproxy server created with the **add v6tov4** command, or adds a new entry to the list that maps port/address pairs.

**Syntax**
**set v6tov4 listenport=** {*Integer* | *ServiceName*} [[**connectaddress=**] {*IPv4Address* | *HostName*}] [[**connectport=**] {*Integer* | *ServiceName*}] [[**listenaddress=**] {*IPv6Address*| *HostName*}] [[**protocol=**]**tcp**]

**Parameters**
**listenport**
Required. Specifies the IPv6 port, by port number or service name, on which to listen.

**connectaddress**
Specifies the IPv4 address to which to connect. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**connectport**
Specifies the IPv4 port, by port number or service name, to which to connect. If **connectport** is not specified, the default is the value of **listenport** on the local computer.

**listenaddress**
Specifies the IPv6 address on which to listen. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**protocol**
Specifies the protocol to use. Currently only TCP is supported.

**/?**
Displays help at the command prompt.

**Remarks**

- You can change any parameter values for an existing portproxy server entry. If values are not supplied, no changes are made.

**set v6tov6**
Modifies the parameter values of an existing entry on the portproxy server created with the **add v6tov6** command, or adds a new entry to the list that maps port/address pairs.

**Syntax**
**set v6tov6 listenport=** {*Integer* | *ServiceName*} [[**connectaddress=**] {*IPv6Address* | *HostName*}] [[**connectport=**] {*Integer* | *ServiceName*}] [[**listenaddress=**] {*IPv6Address*| *HostName*}] [[**protocol=**]**tcp**]

**Parameters**
**listenport**
Required. Specifies the IPv6 port, by port number or service name, on which to listen.

**connectaddress**
Specifies the IPv6 address to which to connect. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If an address is not specified, the default is the local computer.

**connectport**
Specifies the IPv6 port, by port number or service name, to which to connect. If **connectport** is not specified, the default is the value of **listenport** on the local computer.

**listenaddress**
Specifies the IPv6 address on which to listen. Acceptable values are IP address, computer NetBIOS name, or computer DNS name. If you do not specify an address, the default is the local computer.

**protocol**
Specifies the protocol to use. Currently only TCP is supported.

**/?**
Displays help at the command prompt.

**Remarks**

- You can change any parameter values for an existing portproxy server entry. If values are not supplied, no changes are made.

**show all**
Displays all portproxy parameters, including port/address pairs for v4tov4, v4tov6, v6tov4, and v6tov6.

**Syntax**
**show all**

**show v4tov4**
Displays v4tov4 portproxy parameters.

**Syntax**
**show v4tov4**

**show v4tov6**
Displays v4tov6 portproxy parameters.

**Syntax**
**show v4tov6**

**show v6tov4**
Displays v6tov4 portproxy parameters.

**Syntax**
**show v6tov4**

**show v6tov6**
Displays v6tov6 portproxy parameters.

**Syntax**
**show v6tov6**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Checklist: Installing and configuring IPv6

Understanding IPv6

Command shell overview

# Netsh commands for Internet Protocol security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for Internet Protocol security (IPSec)**
The Netsh commands for Internet Protocol security (IPSec) provide an equivalent alternative to the console-based management and diagnostic capabilities provided by the IP Security Policy Management and IP Security Monitor snap-ins available in the Microsoft Management Console (MMC). By using the Netsh commands for IPSec, you can configure and view static or dynamic IPSec main mode settings, quick mode settings, rules, and configuration parameters.

Administering IPSec from the command line is especially useful when you want to:

- Script IPSec configuration.
- Extend the security and manageability of IPSec by configuring the following features, which are not available in the IP Security Policy Management snap-in: IPSec diagnostics, default traffic exemptions, strong certificate revocation list (CRL) checking, IKE (Oakley) logging, logging intervals, computer startup security, and computer startup traffic exemptions.

You can run these commands from the Windows Server™ 2003 family command prompt or from the command prompt for the **netsh ipsec** context. For these commands to work at the Windows Server™ 2003 family command prompt, you must type **netsh ipsec** before typing commands and parameters as they appear in the syntax below.

**Netsh ipsec static mode commands**
You can use the **netsh ipsec static** commands to perform the same management and monitoring tasks that you can perform by using the IP Security Policy Management and IP Security Monitor consoles. By using these commands, you can create, modify, and assign IPSec policies without immediately affecting the configuration of the active IPSec policy.

**Netsh ipsec dynamic mode commands**
You can use the **netsh ipsec dynamic** commands to display the active state of IPSec and to immediately affect the configuration of the active IPsec policy. These commands directly configure the security policy database (SPD). Changes that you make to an IPSec policy while using these commands take effect only while the IPSec service is running. If the IPSec service is stopped, the dynamic policy settings are discarded. Although most of these commands take effect immediately, several configuration commands still require you to restart the IPSec service or restart the computer before they take effect. For more information about these commands, see the syntax descriptions for the **netsh ipsec dynamic set config** commands.

### Caution

- Because the IPSec Policy Agent does not interpret the **netsh ipsec dynamic** commands, you must be knowledgeable in the application of IKE main mode and quick mode policies to use these commands effectively. Exercise caution when using these commands, because it is possible to create invalid IPSec policy configurations without warning.

### Notes

- The Netsh commands for IPSec can only be used to configure IPSec policies on computers running members of the Windows Server™ 2003 family.

  To use the command line to configure IPSec policies on computers running Windows XP, use Ipseccmd.exe, which is provided on the Windows XPCD, in the \Support\Tools folder. To use the command line to configure IPSec policies on computers running Windows 2000, use Ipsecpol.exe, which is provided with the *Windows 2000 Server Resource Kit*.
- For more information about **netsh**, see Netsh overview and Enter a netsh context.
- For more information about Netsh commands, see The Netsh Command-Line Utility.

**Netsh ipsec**
The following commands are available at the **ipsec**> prompt, which is rooted within the netsh environment.

**Note**

- Although the **dump** command is available at the **ipsec**> prompt, it is not functional.

To view the command syntax, click a command:

- static
- dynamic

**static**
Switches to the static context.

**Syntax**
**static**

**Parameters**
none

**dynamic**
Switches to the dynamic context.

**Syntax**
**dynamic**

**Parameters**
none

**Netsh ipsec static**
The following commands are available at the **ipsec static**> prompt, which is rooted within the netsh environment.

To view the command syntax, click a command:

- add filter
- add filteraction
- add filterlist
- add policy
- add rule
- delete all
- delete filter
- delete filteraction
- delete filterlist
- delete policy
- delete rule
- exportpolicy
- importpolicy
- restorepolicyexamples

- set defaultrule
- set filteraction
- set filterlist
- set policy
- set rule
- set store
- show all
- show filteraction
- show filterlist
- show gpoassignedpolicy
- show policy
- show rule

**add filter**
Adds a filter to the specified filter list.

**Syntax**
**add filter filterlist= srcaddr= dstaddr=** [**description=**][**protocol=**][**mirrored=**] [**srcmask=**][**dstmask=**][**srcport=**] [**dstport=**]

**Parameters**
**filterlist=** ***String***
Required. Specifies the name of the filter list to which the filter is added. Each filter defines a set of inbound or outbound network traffic to be secured.

**srcaddr={ Me| Any|** ***IPAddress*****|** ***DNSName*****|** ***ServerType*****}**
Required. Specifies the source IP address, DNS name, or server type for the IP traffic. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**dstaddr={ Me| Any|** ***IPAddress*****|** ***DNSName*****|** ***ServerType*****}**
Required. Specifies the destination IP address, DNS name, or server type for the IP traffic. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ description=*String*]**
Provides information about the IP filter.

**[ protocol={ ANY| ICMP| TCP| UDP| RAW|** ***Integer*** **}]**
Specifies the IP protocol if, in addition to addressing information, you want to filter a specific IP protocol. The default value is **ANY**, meaning all protocols are used for the filter.

**[ mirrored={ yes| no}]**
Specifies whether to create a mirrored filter. Use **yes** to create two filters based on the filter settings--one for traffic to the destination and one for traffic from the destination. The default value is **yes**.

**[ srcmask={** ***Mask*****|** ***Prefix*****}]**
Specifies the source address subnet mask or the prefix of the packets to be filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ dstmask={** ***Mask*****|** ***Prefix*****}**

Specifies the destination address subnet mask or the prefix value of the packets to be filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ srcport=*Port*]**
Specifies the source port number of the packets to be filtered. This option only applies if you are filtering TCP or UDP packets. If 0 is specified, packets sent from any port are filtered. The default is any.

**[ dstport=*Port*]**
Specifies the destination port number of the packets to be filtered. This option only applies if you are filtering TCP or UDP packets. If 0 is specified, packets sent to any port are filtered. The default is any.

**Remarks**

- If a filter list does not exist, it is created.
- Do not create a filter list with the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec filter lists (for example, **delete filterlist all**).
- To filter any packets sent from or to the computer, you can use **srcaddr=Me** or **dstaddr=Me**.
- To filter packets sent from or to any computer, you can use **srcaddr=Any** or **dstaddr=Any**.
- All string parameters are case-sensitive.

**add filteraction**
Creates a filter action with the specified quick mode security methods.

**Syntax**
**add filteraction name=** [**description=**][**qmpfs=**][**inpass=**] [**soft=**][**action=**][**qmsecmethods=**]

**Parameters**
**name= *String***
Required. Specifies the name of the filter action to be created.

**[ description=*String*]**
Provides information about the filter action.

**[ qmpfs={ yes| no}]**
Specifies whether to enable session key perfect forward secrecy (PFS). If **yes** is specified, new master key material is renegotiated each time a new session key is required. The default value is **no**.

**[ inpass={ yes| no}]**
Specifies whether to allow an incoming packet that matches the configured filter list to be unsecured, but require IPSec-secured communication when replying. The default value is **no**.

**[ soft={ yes| no}]**
Specifies whether to fall back to unsecured communication with other computers that do not support IPSec, or when IPSec negotiations with an IPSec-capable computer fail. The default value is **no**.

**[ action={ permit| block| negotiate}]**
Specifies whether to permit traffic without negotiating IP security. If **permit** is specified, traffic is transmitted or received without negotiating or applying IP security. If **block** is specified, traffic is blocked. If **negotiate** is specified, IP security is used with the specified list of security methods. The default value is **negotiate**.

**[ qmsecmethods="*Neg1Neg2*"]**
Specifies one or more security methods, separated by spaces and defined by the following format:

{**ESP** [*ConfAlg*,*AuthAlg*]**:***k/s* | **AH** [*HashAlg*]**:***k/s* | **AH** [*HashAlg*+**ESP***ConfAlg*,*AuthAlg*]**:***k/s*}]

Where:

***ConfAlg***
Specifies the encryption algorithm. *ConfigAlg* can be **DES** (Data Encryption Standard), **3DES**, or **none**.

***AuthAlg***

Specifies the integrity algorithm. *AuthAlg* can be **MD5** (Message Digest 5), **SHA1** (Secure Hash Algorithm 1), or **none**.

***HashAlg***
Specifies the hash function. *HashAlg* can be **MD5** (Message Digest 5) or **SHA1**.

***k***
Specifies the session key lifetime in kilobytes. After the specified number of kilobytes of data is transferred, a new session key for the quick mode SA is generated. The default value is 100000 kilobytes.

***s***
Specifies the session key lifetime in seconds. The default value is 3600 seconds.

**Remarks**

- Do not create a filter action with the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec filter actions (for example, **delete filteraction all**).
- If **action=permit** or **action=block** are specified, and **qmsecmethods** parameters are specified, the **qmsecmethods** parameters are not used. In addition, if **qmpfs=yes**, **inpass=yes**, or **soft=yes** are specified, those parameters are also not used.
- Session key regeneration will start based on whichever interval, seconds or kilobytes, is reached first. If you do not configure new intervals, the default intervals are used.
- If you do not specify **qmsecmethods=** (quick mode security methods), the following default values are used:
    - ESP [3DES, SHA1]:100000k/3600s
    - ESP [3DES, MD5]:100000k/3600s
- The preference order of each quick mode security method is determined by the order in which it was specified in the command.
- All string parameters are case-sensitive.

**add filterlist**
Creates an empty filter list with the specified name.

**Syntax**
**add filterlist name=** [**description=**]

**Parameters**
**name=** ***String***
Required. Specifies the name of the filter list to be created.

**[ description=*String*]**
Provides information about the filter list.

**Remarks**

- Do not create a filter list with the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec filter lists (for example, **delete filterlist all**).
- All string parameters are case-sensitive.

**add policy**
Creates an IPSec policy with the specified name.

**Syntax**
**add policy name=** [**description=**][**mmpfs=**][**qmpermm=**] [**mmlifetime=**][**activatedefaultrule=**][**pollinginterval=**] [**assign=**][**mmsecmethods=**]

**Parameters**

**name= *String***
Required. Specifies the name of the IPSec policy to be created.

**[ description=*String*]**
Provides information about the IPSec policy.

**[ mmpfs={ yes| no}]**
Specifies whether to enable master key perfect forward secrecy (PFS). If **yes** is specified, main mode security SAs are reauthenticated and new master key keying material is negotiated each time session key material for a quick mode SA is required. The default value is **no**.

**[ qmpermm=*Integer*]**
Specifies the number of times that master keying material can be used to derive the session key. The default value is 0, meaning an unlimited number of quick mode SAs can be derived from the main mode SA.

**[ mmlifetime=*Integer*]**
Specifies the number of minutes after which a new master key will be generated. The default value is 480 minutes.

**[ activatedefaultrule={ yes| no}]**
Specifies whether to activate the default response rule for this IPSec policy. The default value is **yes**.

**[ pollinginterval=*Integer*]**
Specifies how often IPSec polls for changes to this policy. The default value is 180 minutes.

**[ assign={ yes| no}]**
Specifies whether to assign this IPSec policy (only one IPSec policy can be assigned) The default value is **no**.

**[ mmsecmethods="*SecMeth1SecMeth2*"]**
Specifies one or more key exchange security methods, separated by spaces and defined by the following format: *ConfAlg-HashAlg-GroupNumb*, where:

***ConfAlg***
Specifies the encryption algorithm. *ConfAlg* can be **DES** (Data Encryption Standard) or **3DES**.

***HashAlg***
Specifies the hash function. *HashAlg* can be **MD5** (Message Digest 5) or **SHA1** (Secure Hash Algorithm 1).

***GroupNum***
Specifies the Diffie-Hellman group to be used for the base keying material. *GroupNumb* can be: **1** (low, protects with 768 bits of keying material), **2** (medium, protects with 1024 bits), and **3** (high, protects with 2048 bits).

**Remarks**

- Do not create a policy with the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec policies (for example, **delete policy all**).
- Because only one IPSec policy can be assigned, if a policy is currently assigned and you assign a new policy, the currently assigned policy is automatically unassigned.
- If **set store=domain** is specified (when the IPSec policy is stored in Active Directory), **assign** has no effect. To assign a policy to a Group Policy object, you must first create a policy by using the **add policy** command, and then use the **set store** command.
- If **mmpfs=yes** is specified (master key PFS is enabled), by default **qmperm** is set to 1 and not configurable, because each new session will cause the master key keying material to be renegotiated.
- If you do not specify **mmsecmethods=** (key exchange security methods), the following default values are used:
    - 3DES-SHA1-2
    - 3DES-MD5-2
    - 3DES-SHA1-3

- IPSec peers must have at least one common key exchange security method (one that uses the same settings) or negotiations will fail.

- If the number of quick mode negotiations will exceed the value set for the number of quick mode negotiations per main mode negotiation during the main mode lifetime, a new main mode negotiation occurs.

- All string parameters are case-sensitive.

**add rule**

Creates a rule that links the specified IPSec policy, filter list, and filter action with the specified authentication methods.

**Syntax**

**add rule name= policy= filterlist= filteraction=**[**tunnel=**][**conntype=**] [**activate=**][**description=**][**kerberos=**][**psk=**] [**rootca=**]

**Parameters**

**name=** ***String***

Required. Specifies the name of the IPSec rule to be created.

**policy=** ***String***

Required. Specifies the name of the IPSec policy that contains this rule.

**filterlist=** ***String***

Required. Specifies the name of the IP filter list for this rule.

**filteraction=** ***String***

Required. Specifies the name of the filter action for this rule.

**[ tunnel={** ***IPAddress*****|** ***DNSName*****}]**

Specifies the IP address or DNS name of the tunnel endpoint for tunnel mode. By default, this option is not specified and transport mode is used.

**[ conntype={ lan| dialup| all}]**

Specifies whether the rule applies only to remote access or dial-up connections or to local area network (LAN) connections, or to all connections. The default value is **all**.

**[ activate={ yes | no}]**

Specifies whether to activate this rule for the specified IPSec policy. The default value is **yes**.

**[ description=*String*]**

Provides information about the rule.

**[ kerberos={ yes | no}]**

Specifies whether to use the Kerberos V5 protocol as an authentication method.

**[ psk=*String*]**

Specifies the string of characters to use for the preshared key, if a preshared key is used as an authentication method.

**[ rootca="*String*certmap:{ yes| no} excludecaname:{ yes| no} "]**

Specifies certificate authentication options, where:

***String***

Specifies the distinguished name of the certificate, if a certificate is used as an authentication method.

**certmap:{ yes| no}**

Specifies whether to enable certificate-to-account mapping. You can enable certificate-to-account mapping to verify that the certificate is being used by a trusted computer.

**excludecaname:{ yes| no}**

Specifies whether to exclude from the certificate request the list of trusted root CA names from which a certificate is accepted.

**Remarks**

- Do not create a rule with the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec rules (for example, **delete rule all**).

- You can only use Kerberos V5 authentication or certificate-to-account mapping for computers that are members of an Active Directory domain.

- Although you can use only one preshared key for authentication, you can use multiple certificates by specifying the **rootca** parameter once, for each certificate you want to use.

- All certificate authentication parameters must be contained within quotation marks. Embedded quotation marks must be replaced by a backslash followed by an apostrophe (\').

- All string parameters are case-sensitive.

- The preference order of each authentication method is determined by the order in which it was specified in the command.

- If no authentication methods are specified, dynamic defaults are used. By default, IPSec policies use Kerberos V5 authentication. If the computer has a computer certificate, any root CA to which the computer certificate chains is also used for authentication.

- If **excludecaname:yes** is specified, the list of trusted root CAs is not sent as part of the certificate request, which prevents the potential disclosure of sensitive information about the trust relationships of a computer. To enhance security for computers connected to the Internet, specify this option.

- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. In addition, preshared keys are stored in plaintext.

- IPSec peers must have at least one common authentication method or communication will fail.

**delete all**
Deletes all IPSec policies, filter lists, and filter actions.

**Syntax**
**delete all**

**Parameters**
None.

**delete filter**
Deletes a filter from a filter list that matches the specified parameters.

**Syntax**
**delete filter filterlist= srcaddr= dstaddr=**[**protocol=**] [**srcmask=**][**dstmask=**][**srcport=**] [**dstport=**][**mirrored=**]

**Parameters**
**filterlist=** ***String***
Required. Specifies the name of the filter list to which the filter was added.

**srcaddr={ Me| Any|** ***IPAddress*****|** ***DNSName*****|** ***ServerType*****}**
Required. Specifies the source IP address, DNS name, or server type for the IP traffic being matched. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**dstaddr={ Me| Any|** ***IPAddress*****|** ***DNSName*****|** ***ServerType*****}**
Required. Specifies the destination IP address, DNS name, or server type for the IP traffic being matched. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ protocol={ ANY| ICMP| TCP| UDP| RAW|** ***Integer*** **}]**
Specifies the IP protocol if, in addition to addressing information, a specific IP protocol is filtered. A value of **ANY** matches filters with a protocol setting of **any**.

**[ srcmask={** ***Mask*****|** ***Prefix*****}]**

Specifies the source address subnet mask or the prefix of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ dstmask={ *Mask*| *Prefix*}]**
Specifies the destination address subnet mask or the prefix value of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ srcport=*Port*]**
Specifies the source port number of the packets being filtered. This option only applies if you are filtering TCP or UDP packets. A value of **0** matches filters with a source port setting of **any**.

**[ dstport=*Port*]**
Specifies the destination port number of the packets being filtered. This option only applies if you are filtering TCP or UDP packets. A value of **0** matches filters with a destination port setting of **any**.

**[ mirrored={ yes| no}]**
Specifies whether a mirrored filter was created.

**Remarks**

- This command only deletes a filter that matches the exact parameters specified.
- If no optional parameters are specified, all filters that match the specified (required) parameters are deleted.
- To delete a filter that filtered any packets sent from or to the computer, you can use **srcaddr=Me** or **dstaddr=Me**.
- To delete a filter that filtered packets sent from or to any computer, you can use **srcaddr=Any** or **dstaddr=Any**.
- All string parameters are case-sensitive.

**delete filteraction**
Deletes the specified filter action, or all filter actions.

**Syntax**
**delete filteraction name=** | **all**

**Parameters**
**name= *String*| all**
Required. Specifies the name of the filter action to delete. Or, if **all** is specified, all filter actions are deleted.

**Remarks**

- All string parameters are case-sensitive.

**delete filterlist**
Deletes the specified filter list, or all filter lists.

**Syntax**
**delete filterlist name=** | **all**

**Parameters**
**name= *String*| all**
Required. Specifies the name of the filter list to delete. Or, if **all** is specified, all filter lists are deleted.

**Remarks**

- All string parameters are case-sensitive.

**delete policy**
Deletes the specified IPSec policy and all associated rules, or all IPSec policies.

**Syntax**
**delete policy name=** | **all**

**Parameters**

**name= *String*| all**

Required. Specifies the name of the IPSec policy to delete. Or, if **all** is specified, all IPSec policies are deleted.

**Remarks**

- All string parameters are case-sensitive.

**delete rule**

Deletes a specified rule, or all rules from the specified IPSec policy.

**Syntax**

**delete rule name= | ID= | allpolicy=**

**Parameters**

**name= *String*| ID=*Integer*| all**

Required. Specifies the rule to delete. If either the rule name or the rule ID (the number identifying the position of the rule in the policy rule list) is specified, the corresponding rule is deleted. If **all** is specified, all rules are deleted.

**policy= *String***

Required. Specifies the name of the policy from which one or more rules are deleted.

**Remarks**

- The default response rule cannot be deleted.
- After a rule is deleted, all IDs for the remaining rules change accordingly.
- All string parameters are case-sensitive.

**exportpolicy**

Exports all IPSec policy information to the specified file.

**Syntax**

**exportpolicy file=**

**Parameters**

**file= *String***

Required. Specifies the name of the file into which the IPSec policy information is exported.

**Remarks**

- By default, when an IPSec policy is imported into a file, the .ipsec extension added to the file name.
- To enhance interoperablity in a mixed environment with computers running Windows 2000, limit the name of the file to which you want the policy information saved to 60 characters.
- All string parameters are case-sensitive.

**importpolicy**

Imports all IPSec policy information from the specified IPSec file.

**Syntax**

**importpolicy file=**

**Parameters**

**file= *String***

Required. Specifies the name of the file from which the IPSec policy information is imported.

**Remarks**

- All string parameters are case-sensitive.

**restorepolicyexamples**
Restores the default IPSec policies.

**Syntax**
**restorepolicyexamples release=**

**Parameters**
**release={ win2K| Win2003}**
Required. Specifies the version of the default IPSec policies being restored. If **win2K** is specified, the default IPSec policies that were provided with Windows 2000 are restored. If **Win2003** is specified, the default IPSec policies that were provided with the Windows Server™ 2003 family are restored.

**Remarks**

- Restoring default IPSec policies will overwrite any changes to the original default policies, filter lists, and filter actions, even when the names of these configuration items have been changed. If you have modified these items and you do not want the modifications to be overwritten, do not restore the default policies.

- You can only restore default IPSec policies for computer-based IPSec policies. You cannot restore the default IPSec policies for IPSec policies in Active Directory.

**set defaultrule**
Modifies the default response rule for the specified policy.

**Syntax**
**set defaultrule policy=**[**qmpfs=**][**activate=**] [**qmsecmethods=**][**kerberos=**][**psk=**][**rootca=**]

**Parameters**
**policy=** ***String***
Required. Specifies the name of the IPSec policy for which the default response rule is to be modified.

**[ qmpfs={ yes| no}]**
Specifies whether to enable session key perfect forward secrecy (PFS). If **yes** is specified, new master key material is renegotiated each time a new session key is required. The default value is **no**.

**[ activate={ yes | no}]**
Specifies whether to activate this rule for the specified IPSec policy. The default value is **yes**.

**[ qmsecmethods="*Neg1Neg2*"]**
Specifies one or more security methods, separated by spaces and defined by the following format:

{**ESP** [*ConfAlg*,*AuthAlg*]:*k*/*s* | **AH** [*HashAlg*]:*k*/*s* | **AH** [*HashAlg*+**ESP***ConfAlg*,*AuthAlg*]:*k*/*s*}]

Where:

***ConfAlg***
Specifies the encryption algorithm. *ConfigAlg* can be **DES** (Data Encryption Standard), **3DES**, or **none**.

***AuthAlg***
Specifies the integrity algorithm. *AuthAlg* can be **MD5** (Message Digest 5), **SHA1** (Secure Hash Algorithm 1), or **none**.

***HashAlg***
Specifies the hash function. *HashAlg* can be **MD5** (Message Digest 5) or **SHA1**.

***k***
Specifies the session key lifetime in kilobytes. After the specified number of kilobytes of data is transferred, a new session key for the quick mode SA is generated. The default value is 100,000 kilobytes.

***s***
Specifies the session key lifetime in seconds. The default value is 3600 seconds.

**[ kerberos={ yes | no}]**
Specifies whether to use the Kerberos V5 protocol as an authentication method.

**[ psk=*String*]**
Specifies the string of characters to use for the preshared key, if a preshared key is used as an authentication method.

**[ rootca="*String*certmap:{ yes| no} excludecaname:{ yes| no} "]**
Specifies certificate authentication options, where:

***String***
Specifies the distinguished name of the certificate, if a certificate is used as an authentication method.

**certmap:{ yes| no}**
Specifies whether to enable certificate-to-account mapping. You can enable certificate-to-account mapping to verify that the certificate is being used by a trusted computer.

**excludecaname:{ yes| no}**
Specifies whether to exclude from the certificate request the list of trusted root CA names from which a certificate is accepted.

**Remarks**

- You can only use Kerberos V5 authentication or certificate-to-account mapping for computers that are members of an Active Directory domain.
- Although you can use only one preshared key for authentication, you can use multiple certificates by specifying the **rootca** parameter once, for each certificate you want to use.
- All certificate authentication parameters must be contained within quotation marks. Embedded quotation marks must be replaced by a backslash followed by an apostrophe (\').
- All string parameters are case-sensitive.
- The preference order of each authentication method is determined by the order in which it was specified in the command.
- If no authentication methods are specified, dynamic defaults are used. By default, IPSec policies use Kerberos V5 authentication. If the computer has a computer certificate, any root CA to which the computer certificate chains is also used for authentication.
- If **excludecaname:yes** is specified, the list of trusted root CAs is not sent as part of the certificate request, which prevents the potential disclosure of sensitive information about the trust relationships of a computer. To enhance security for computers connected to the Internet, specify this option.
- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. In addition, preshared keys are stored in plaintext.
- IPSec peers must have at least one common authentication method or communication will fail.
- Modifying authentication methods will overwrite all previous authentication methods, even if the previous authentication methods were different. For example, if **kerberos=yes** and **psk=yes** were previously specified, and then you specify **kerberos=no**, the **psk=yes** parameter will also be overwritten, and preshared key authentication will no longer be used.
- *ConfAlg* and *AuthAlg* cannot both be set to **none**.

**set filteraction**
Modifies a filter action.

**Syntax**
**set filteraction name=** | **guid=**[**newname=**] [**description=**][**qmpfs=**][**inpass=**] [**soft=**][**action=**][**qmsecmethods=**]

**Parameters**
**name= *String*| guid= *guid***
Required. Specifies the name or global unique identifier (GUID) of the filter action to modify.

**[ newname=*String*]**
Specifies the new name of the filter action.

**[ description=*String*]**
Provides information about the filter action.

**[ qmpfs={ yes| no}]**
Specifies whether to enable session key perfect forward secrecy (PFS). If **yes** is specified, new master key material is renegotiated each time a new session key is required.

**[ inpass={ yes| no}]**
Specifies whether to allow an incoming packet that matches the configured filter list to be unsecured, but require IPSec-secured communication when replying.

**[ soft={ yes| no}]**
Specifies whether to fall back to unsecured communications with other computers that do not support IPSec, or when IPSec negotiations with an IPSec-capable computer fail.

**[ action={ permit| block| negotiate}]**
Specifies whether to permit traffic without negotiating IP security. If **permit** is specified, traffic is transmitted or received without negotiating or applying IP security. If **block** is specified, traffic is blocked. If **negotiate** is specified, IP security is used, with the specified list of security methods.

**[ qmsecmethods="*Neg1Neg2*"]**
Specifies one or more security methods, separated by spaces and defined by the following format:

{**ESP** [*ConfAlg*,*AuthAlg*]**:***k/s* | **AH** [*HashAlg*]**:***k/s* | **AH** [*HashAlg*+**ESP***ConfAlg*,*AuthAlg*]**:***k/s*}]

Where:

***ConfAlg***
Specifies the encryption algorithm. *ConfigAlg* can be **DES** (Data Encryption Standard), **3DES**, or **none**.

***AuthAlg***
Specifies the integrity algorithm. *AuthAlg* can be **MD5** (Message Digest 5), **SHA1** (Secure Hash Algorithm 1), or **none**.

***HashAlg***
Specifies the hash function. *HashAlg* can be **MD5** (Message Digest 5) or **SHA1** (Secure Hash Algorithm 1).

***k***
Specifies the session key lifetime in kilobytes. After the specified number of kilobytes of data is transferred, a new session key for the quick mode SA is generated. The default value is 100000 kilobytes.

***s***
Specifies the session key lifetime in seconds. The default value is 3600 seconds.

**Remarks**

- If you specify a new name for the filter action, do not use the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec filter actions (for example, **delete filteraction all**).
- If **action=permit** or **action=block** are specified, do not set **qmpfs=yes**, **inpass=yes**, or **soft=yes**.
- Session key regeneration will start based on whichever interval, seconds or kilobytes, is reached first. If you do not configure new intervals, the default intervals are used.
- If **qmsecmethods=** (quick mode security methods) were not previously specified for this filter action, the following default values are used:
    - ESP [3DES, SHA1]:100000s/3600k
    - ESP [3DES, MD5]:100000s/3600k
- The preference order of each quick mode security method is determined by the order in which it was specified in the command.
- All string parameters are case-sensitive.

**set filterlist**
Modifies a filter list.

**Syntax**
**set filterlist name=** [**newname=**] [**description=**]

**Parameters**
**name= *String***
Required. Specifies the name of the filter list to modify.

**[ newname=*String*]**
Specifies the new name of the filter list.

**[ description=*String*]**
Provides information about the filter list.

**Remarks**

- If you specify a new name for the filter list, do not use the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec filter lists (for example, **delete filterlist all**).

- All string parameters are case-sensitive.

**set policy**
Modifies an IPSec policy.

**Syntax**
**set policy name= newname=** [**description=**][**mmpfs=**][**qmpermm=**] [**mmlifetime=**][**activatedefaultrule=**] [**pollinginterval=**][**assign=**][**gponame=**][**mmsecmethods=**]

**Parameters**
**name= *String*| guid=*guid***
Required. Specifies the name or GUID of the IPSec policy to modify.

**newname= *String***
Required. Specifies the new name of the IPSec policy.

**[ description=*String*]**
Provides information about the IPSec policy.

**[ mmpfs={ yes| no}]**
Specifies whether to enable master key perfect forward secrecy (PFS). If **yes** is specified, main mode security SAs are reauthenticated and new master key keying material is negotiated each time session key material for a quick mode SA is required.

**[ qmpermm=*Integer*]**
Specifies the number of times that master keying material can be used to derive the session key.

**[ mmlifetime=*Integer*]**
Specifies the number of minutes after which a new master key will be generated.

**[ activatedefaultrule={ yes| no}]**
Specifies whether to activate the default response rule for this IPSec policy.

**[ pollinginterval=*Integer*]**
Specifies how often IPSec polls for changes to this policy. The default value is 180 minutes.

**[ assign={ yes| no}]**
Specifies whether to assign this IPSec policy.

**[ gponame=*String*]**
Specifies the name of the Group Policy object to which the IPSec policy is assigned. This parameter is only applicable if you are configuring policy for a computer that is an Active Directory domain member.

**[ mmsecmethods="*SecMeth1SecMeth2*"]**

Specifies one or more key exchange security methods, separated by spaces and defined by the following format: *ConfAlg-HashAlg-GroupNumb*, where:

***ConfAlg***
Specifies the encryption algorithm. *ConfAlg* can be **DES** (Data Encryption Standard) or **3DES**.

***HashAlg***
Specifies the hash function. *HashAlg* can be **MD5** (Message Digest 5) or **SHA1** (Secure Hash Algorithm 1).

***GroupNum***
Specifies the Diffie-Hellman group to be used for the base keying material. *GroupNumb* can be: **1** (low, protects with 768 bits of keying material), **2** (medium, protects with 1024 bits), and **3** (high, protects with 2048 bits).

**Remarks**

- If you specify a new name for the policy, do not use the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec policies (for example, **delete policy all**).
- If **set store=domain** is specified (when the IPSec policy is stored in Active Directory), **assign** will have no effect.
- If **mmpfs=yes** is specified (master key PFS is enabled), by default **qmperm** is set to 1 and not configurable, because each new session will cause the master key keying material to be renegotiated.
- IPSec peers must have at least one common key exchange security method (one that uses the same settings) or negotiations will fail.
- You can only specify a Group Policy object name if **set store=domain**.
- All string parameters are case-sensitive.

**set rule**
Modifies a rule in an IPSec policy.

**Syntax**
**set rule name=** | **ID=policy=** [**newname=**][**description=**][**filterlist=**] [**filteraction=**] [**tunnel=**][**conntype=**] [**activate=**] [**kerberos=**][**psk=**][**rootca=**]

**Parameters**
**name=** ***String*** | **ID=*Integer***
Required. Specifies the name or ID (the number identifying the position of the rule in the policy rule list) of the rule to modify.

**policy=** ***String***
Required. Specifies the name of the IPSec policy that contains this rule.

**[ newname=*String*]**
Specifies the new name of the rule.

**[ description=*String*]**
Provides information about the rule.

**[ filterlist=*String*]**
Specifies the name of the IP filter list for this rule.

**[ filteraction=*String*]**
Specifies the name of the filter action for this rule.

**[ tunnel={ *IPAddress*| *DNSName*}]**
Specifies the IP address or DNS name of the tunnel endpoint for tunnel mode.

**[ conntype={ lan| dialup| all}]**
Specifies whether the rule applies only to remote access or dial-up connections or to local area network (LAN) connections, or to all connections.

**[ activate={ yes | no}]**
Specifies whether to activate this rule for the specified IPSec policy.

**[ kerberos={ yes | no}]**
Specifies whether to use the Kerberos V5 protocol as an authentication method.

**[ psk=*String*]**
Specifies the string of characters to use for the preshared key, if a preshared key is used as an authentication method.

**[ rootca="*String*certmap:{ yes| no} excludecaname:{ yes| no} "]**
Specifies certificate authentication options, where:

***String***
Specifies the distinguished name of the certificate, if a certificate is used as an authentication method.

**certmap:{ yes| no}**
Specifies whether to enable certificate-to-account mapping. You can enable certificate-to-account mapping to verify that the certificate is being used by a trusted computer.

**excludecaname:{ yes| no}**
Specifies whether to exclude from the certificate request the list of trusted root CA names from which a certificate is accepted.

**Remarks**

- You can only use Kerberos V5 authentication or certificate-to-account mapping for computers that are members of an Active Directory domain.
- Although you can use only one preshared key for authentication, you can use multiple certificates by specifying the **rootca** parameter once, for each certificate you want to use.
- All certificate authentication parameters must be contained within quotation marks. Embedded quotation marks must be replaced by a backslash followed by an apostrophe (\').
- All string parameters are case-sensitive.
- The preference order of each authentication method is determined by the order in which it was specified in the command.
- If no authentication methods are specified, dynamic defaults are used. By default, IPSec policies use Kerberos V5 authentication. If the computer has a computer certificate, any root CA to which the computer certificate chains is also used for authentication.
- If **excludecaname:yes** is specified, the list of trusted root CAs is not sent as part of the certificate request, which prevents the potential disclosure of sensitive information about the trust relationships of a computer. To enhance security for computers connected to the Internet, specify this option.
- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. In addition, preshared keys are stored in plaintext.
- IPSec peers must have at least one common authentication method or communication will fail.
- Modifying authentication methods will overwrite all previous authentication methods, even if the previous authentication methods were different. For example, if **kerberos=yes** and **psk=yes** were previously specified, and then you specify **kerberos=no**, the **psk=yes** parameter will also be overwritten, and preshared key authentication will no longer be used.

**set store**
Sets the current IPSec policy storage location.

**Syntax**
**set store location=** [**domain=**]

**Parameters**

**location={ local| persistent| domain}**

Required. Specifies the storage location for the IPSec policy.

**[ domain=]**

Specifies the name of the domain where the IPSec policy is stored, if the policy is stored in Active Directory (when **location=domain** is specified).

**Remarks**

- The **set store** command only works from within the netsh environment, that is:
    - If you run this command from the command prompt for the **netsh ipsec** context).
    - If you run a batch file by using the **netsh.exe** command.
- The persistent store contains IPSec policies that can be assigned to secure this computer at start up, before the local policy or domain-based policy is applied. A persistent IPSec policy provides security in the event of a failure because it remains in effect whether the local policy or domain-based policy is applied or not (for example, an IPSec policy might not be applied if it is corrupted). For enhanced security, it is recommended that you create and assign a persistent policy.
- The local store contains IPSec policies that can be assigned to secure this computer. If a domain policy is available, the domain policy is applied instead of the local policy.
- The domain store contains IPSec policies that can be assigned to secure groups of computers in a domain.
- It is recommended that the persistent policy be the most restrictive of all policies. Domain policies and local policies should complement persistent policies.
- Use the **set machine** command to configure a remote computer.
- All string parameters are case-sensitive.

**show all**

Displays configuration information for all IPSec policies, rules, filter lists, and filter actions.

**Syntax**

**show all** [**format=**] [**wide=**]

**Parameters**

**[ format={ list| table}]**

Specifies whether to display IPSec configuration information in screen or tab-delimited format. The default value is **list**, meaning that output is displayed in screen format.

**[ wide={ yes | no}]**

Specifies whether to allow the display of IPSec configuration information to exceed the screen width of 80 characters. The default value is **no**, meaning that the display of configuration information is limited to the screen width.

**Remarks**

- Because the **show all** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show all** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

2. Then type:

   **ipsec static show all**

3. To stop sending output and close the file, type:

   **set file close**

**If you are not in the netsh environment**

- At the command prompt, type:

  **netsh ipsec static show all >***FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.

  - Use Task Manager to end the **Netsh** program.

**show filteraction**
Displays configuration information for one or more filter actions.

**Syntax**
**show filteraction name=** | **rule=** | **all** [**level=**][**format=**] [**wide=**]

**Parameters**
**name=** ***String*| rule=*String*| all**
Required. Specifies one or more filter actions for which configuration information is to be displayed. If **name** is specified, the filter action with the specified name is displayed. If **rule** is specified, all filter actions associated with the specified rule are displayed. If **all** is specified, all filter actions are displayed.

**[ level={ verbose| normal}]**
Specifies the level of information to display. If **verbose** is specified, information about the security methods, policy storage location, and whether session key perfect forward secrecy (PFS) is enabled is displayed, in addition to basic filter action information. The default value is **normal**.

**[ format={ list| table}]**
Specifies whether to display IPSec configuration information in screen or tab-delimited format. The default value is **list**, meaning that output is displayed in screen format.

**[ wide={ yes | no}]**
Specifies whether to allow the display of IPSec configuration information to exceed the screen width of 80 characters. The default value is **no**, meaning that the display of configuration information is limited to the screen width.

**Remarks**

- All string parameters are case-sensitive.

- Because the **show filteraction** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show filteraction** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

**set file open** *FileName*.txt

2. Then type:

   **ipsec static show filteraction***Name* | *Rule* | **all** [**level=verbose** | **normal**]

3. To stop sending output and close the file, type:

   **set file close**

**If you are not in the netsh environment**

- At the command prompt, type:

  **netsh ipsec static show filteraction***Name* | *Rule* | **all** [**level=verbose** | **normal**] >*FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.

  - Use Task Manager to end the **Netsh** program.

**show filterlist**

Displays configuration information for one or more filter lists.

**Syntax**

**show filterlist name=** | **rule=** | **all** [**level=**][**format=**][**resolvedns=**] [**wide=**]

**Parameters**

**name=** ***String*| rule=*String*| all**

Required. Specifies one or more filter lists to display. If **name** is specified, the filter list with the specified name is displayed. If **rule** is specified, all filter lists associated with the specified rule are displayed. If **all** is specified, all filter lists are displayed.

**[ level={ verbose| normal}]**

Specifies the level of information to display. If **verbose** is specified, the source, destination, and type of IP traffic defined by each filter are displayed, in addition to basic filter list information. The default value is **normal**.

**[ format={ list| table}]**

Specifies whether to display IPSec configuration information in screen or tab-delimited format. The default value is **list**, meaning that output is displayed in screen format.

**[ resolvedns={ yes | no}]**

Specifies whether to resolve the Domain Name System (DNS) or NETBIOS computer name associated with an IP address when displaying sources or destinations. If **yes** is specified, **level** must also be set to **verbose**, or the DNS names are not displayed. The default value is **no**.

**[ wide={ yes | no}]**

Specifies whether to allow the display of IPSec configuration information to exceed the screen width of 80 characters. The default value is **no**, meaning that the display of configuration information is limited to the screen width.

**Remarks**

- All string parameters are case-sensitive.

- Because the **show filterlist** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show all** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

1. At the netsh prompt, type:

   **set file open** *FileName*.txt

2. Then type:

   **ipsec static show filterlist***Name* | *rule* | **all** [**level=verbose** | **normal**][**resolvedns=yes** | **no**

3. To stop sending output and close the file, type:

   **set file close**

**If you are not in the netsh environment**

- At the command prompt, type:

  **netsh ipsec static show filterlist***Name* | *rule* | **all** [**level=verbose** | **normal**][**resolvedns=yes** | **no** >*FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.

  - Use Task Manager to end the **Netsh** program.

**show gpoassignedpolicy**
Displays configuration information for the active IPSec policy assigned to the specified Group Policy object.

**Syntax**
**show gpoassignedpolicy** [**name=** ]

**Parameters**
**[ name=*String*]**
Specifies the name of the Group Policy object to which the active IPSec policy is assigned. If no name is specified, the local IPSec policy is displayed.

**Remarks**

- You can only specify a Group Policy object name if **set store=domain**.

- All string parameters are case-sensitive.

  To save output in a text file for the **show gpossignedpolicy** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec static show gpoassignedpolicy** [*Name*]

  3. To stop sending output and close the file, type:

     **set file close**

**If you are not in the netsh environment**

- At the command prompt, type:

  **netsh ipsec static show gpoassignedpolicy** [*Name*] >*FileName*.txt

**show policy**

Displays configuration information for the specified IPSec policy, or for all IPSec policies.

**Syntax**

**show policy name=** | **all** [**level=**] [**format=**] [**wide=**]

**Parameters**

**name=** ***String*****| all**

Required. Specifies the name of the IPSec policy to display or, if **all** is specified, that all IPSec policies are displayed.

**[ level={ verbose| normal}]**

Specifies the level of information to display. If **verbose** is specified, the security methods and authentication method are displayed, in addition to information about filter actions and rules. The default value is **normal**.

**[ format={ list| table}]**

Specifies whether to display IPSec configuration information in screen or tab-delimited format. The default value is **list**, meaning that output is displayed in screen format.

**[ wide={ yes | no}]**

Specifies whether to allow the display of IPSec configuration information to exceed the screen width of 80 characters. The default value is **no**, meaning that the display of configuration information is limited to the screen width.

**Remarks**

- All string parameters are case-sensitive.

- Because the **show policy** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show policy** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec static show policy***Name* | **all** >*FileName*.txt

  3. To stop sending output and close the file, type:

     **set file close**

  **If you are not in the netsh environment**

  - At the command prompt, type:

    **netsh ipsec static show policy***Name* | **all** >*FileName***.txt**

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

- Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.

- Use Task Manager to end the **Netsh** program.

**show rule**

Displays configuration information for a rule for a specified policy, or for all rules for a specified policy.

**Syntax**

**show rule name=** | **ID=** | **all** | **defaultpolicy=** [**type=**][**level=**][**format=**] [**wide=**]

**Parameters**

**name={ *String*| ID=*Integer*| all| default}**

Required. Specifies one or more rules to display. If either the rule name or the rule ID (the number identifying the position of the rule in the policy rule list) is specified, the corresponding rule is displayed. If **all** is specified, all rules for the specified policy are displayed. If **default** is specified, the default response rule is displayed.

**policy= *String***

Required. Specifies the name of the policy for which the specified rule, or all rules, are displayed.

**[ type={ transport| tunnel}]**

Specifies whether to display all transport rules or all tunnel rules. The default value is to display all rules.

**[ level={ verbose| normal}]**

Specifies the level of information to display. If **verbose** is specified, information about associated filter actions is displayed, in addition to basic information about the rule. The default value is **normal**.

**[ format={ list| table}]**

Specifies whether to display IPSec configuration information in screen or tab-delimited format. The default value is **list**, meaning that output is displayed in screen format.

**[ wide={ yes | no}]**

Specifies whether to allow the display of IPSec configuration information to exceed the screen width of 80 characters. The default value is **no**, meaning that the display of configuration information is limited to the screen width.

**Remarks**

- If you use the **type** parameter, you must also use the **all** parameter (you must specify **show rule all**).

- All string parameters are case-sensitive.

- Because the **show rule** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show rule** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec static show rule***Name* | *ID* | **all** | **default***Policy* [**level=verbose** | **normal**]

  3. To stop sending output and close the file, type:

     **set file close**

**If you are not in the netsh environment**

- At the command prompt, type:

  **netsh ipsec static show rule***Name* | *ID* | **all** | **default***Policy* [**level=verbose** | **normal**] >*FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.

  - Use Task Manager to end the **Netsh** program.

**Netsh ipsec dynamic**

The following commands are available at the **ipsec dynamic** > prompt, which is rooted within the netsh environment.

To view the command syntax, click a command:

- add mmpolicy
- add qmpolicy
- add rule
- delete all
- delete mmpolicy
- delete qmpolicy
- delete rule
- set config
- set mmpolicy
- set qmpolicy
- set rule
- show all
- show config
- show mmfilter
- show mmpolicy
- show mmsas
- show qmfilter
- show qmpolicy
- show qmsas
- show rule

- show stats

**add mmpolicy**
Creates an IPSec main mode policy with the specified name and adds it to the security policy database (SPD).

**Syntax**
**add mmpolicy name=** [**qmpermm=**] [**mmlifetime=**][**softsaexpirationtime=**][**mmsecmethods=**]

**Parameters**
**name=** ***String***
Required. Specifies the name of the IPSec policy to be created.

**[ qmpermm=*Integer*]**
Specifies the number of times that master keying material can be used to derive the session key. The default value is 0, The default value is 0, meaning an unlimited number of quick mode SAs can be derived from the main mode SA.

**[ mmlifetime=*Integer*]**
Specifies the number of minutes after which a new master key is generated. The default value is 480 minutes.

**[ softsaexpirationtime=*Integer*]**
Specifies the number of minutes after which an unprotected security association expires. The default value is 480 minutes.

**[ mmsecmethods="*SecMeth1SecMeth2*"]**
Specifies one or more key exchange security methods, separated by spaces and defined by the following format: *ConfAlg-HashAlg-GroupNumb*, where:

***ConfAlg***
Specifies the encryption algorithm. *ConfAlg* can be **DES** or **3DES**.

***HashAlg***
Specifies the hash function. *HashAlg* can be **MD5** or **SHA1**.

***GroupNum***
Specifies the Diffie-Hellman group to be used for the base keying material. *GroupNumb* can be: **1** (low, protects with 768 bits of keying material), **2** (medium, protects with 1024 bits), and **3** (high, protects with 2048 bits).

**Remarks**

- Do not create a main mode policy with the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec main mode policies (for example, **delete mmpolicy all**).
- If the number of quick mode negotiations will exceed the value set for the number of quick mode negotiations per main mode negotiation during the main mode lifetime, a new main mode negotiation occurs.
- If you do not specify **mmsecmethods=** (key exchange security methods), the following default values are used:
    - 3DES-SHA1-2
    - 3DES-MD5-2
    - 3DES-SHA1-3
- IPSec peers must have at least one common key exchange security method (one that uses the same settings) or negotiations will fail.
- All string parameters are case-sensitive.

**add qmpolicy**
Creates an IPSec quick mode policy with the specified name and adds it to the SPD.

**Syntax**

**add qmpolicy name=** [**soft=**][**pfsgroup=**][**qmsecmethods=**]

**Parameters**

**name=** ***String***

Required. Specifies the name of the IPSec quick mode policy to be created.

**[ soft={ yes| no}]**

Specifies whether to fall back to unsecured communications with other computers that do not support IPSec, or when IPSec negotiations with an IPSec-capable computer fail. The default value is **no**.

**[ pfsgroup={ grp1| grp2| grp3| grpmm| nopfs}]**

Specifies the Diffie-Hellman group to use for session key PFS. If **grp1** is specified, Group 1 (low) is used. If **grp2** is specified, Group 2 (medium) is used. If **grp3** is specified, Group 2048 (high) is used. If **grpmm** is specified, the group value is taken from the current main mode settings. The default value is **nopfs**, meaning session key PFS is disabled.

**[ qmsecmethods="*Neg1Neg2*"]**

Specifies one or more security methods, separated by spaces and defined by the following format:

{**ESP** [*ConfAlg*,*AuthAlg*]**:***k/s* | **AH** [*HashAlg*]**:***k/s* | **AH** [*HashAlg*+**ESP***ConfAlg*,*AuthAlg*]**:***k/s*}]

Where:

***ConfAlg***

Specifies the encryption algorithm. *ConfigAlg* can be **DES** (Data Encryption Standard), **3DES**, or **none**.

***AuthAlg***

Specifies the integrity algorithm. *AuthAlg* can be **MD5** (Message Digest 5), **SHA1** (Secure Hash Algorithm 1), or **none**.

***HashAlg***

Specifies the hash function. *HashAlg* can be **MD5** (Message Digest 5) or **SHA1** (Secure Hash Algorithm 1).

***k***

Specifies the session key lifetime in kilobytes. After the specified number of kilobytes of data is transferred, a new session key for the quick mode SA is generated. The default value is 100000 kilobytes.

***s***

Specifies the session key lifetime in seconds. The default value is 3600 seconds.

**Remarks**

- Do not create a quick mode policy with the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec quick mode policies (for example, **delete qmpolicy all**).
- If you do not specify **qmsecmethods=** (quick mode security methods), the following default values are used:
    - ESP [3DES, SHA1]100000k/3600s
    - ESP [3DES, MD5]100000k/3600s
- *ConfAlg* and *AuthAlg* cannot both be set to **none**.
- IPSec peers must have the same **pfsgroup** enabled (that is, both peers must use the same Diffie-Hellman group for session key PFS), or communication will fail.
- For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.
- All string parameters are case-sensitive.

**add rule**

Creates an IPSec rule with the specified main mode policy and quick mode policy and adds it to the SPD.

**Syntax**

**add rule srcaddr= dstaddr= mmpolicy=**[**qmpolicy=**][**protocol=**][**srcport=**][**dstport=**][**mirrored=**][**conntype=**] [**actioninbound=**][**actionoutbound=**][**srcmask=**][**dstmask=**][**tunneldstaddress=**][**kerberos=**][**psk=**][**rootca=**]

**Parameters**

**srcaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}**
Required. Specifies the source IP address, DNS name, or server type for the IP traffic. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**dstaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}**
Required. Specifies the destination IP address, DNS name, or server type for the IP traffic. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**mmpolicy= *String***
Required. Specifies the name of the main mode policy.

**[ qmpolicy=] *String***
Specifies the name of the quick mode policy. Required if **actioninbound=negotiate** or **actionoutbound=negotiate** are specified.

**[ protocol={ ANY| ICMP| TCP| UDP| RAW| *Integer* }]**
Specifies the IP protocol if, in addition to address information, you want to filter a specific IP protocol. The default value is **ANY**, meaning all protocols are used for the filter.

**[ srcport=*Port*]**
Specifies the source port number of the packets to be filtered. This option only applies if you are filtering TCP or UDP packets. If 0 is specified, packets sent from any port are filtered. The default is any.

**[ dstport=*Port*]**
Specifies the destination port number of the packets to be filtered. This option only applies if you are filtering TCP or UDP packets. If 0 is specified, packets sent to any port are filtered. The default is any.

**[ mirrored={ yes| no}]**
Specifies whether to create a mirrored filter. Use **yes** to create two filters based on the filter settings, one for traffic to the destination and one for traffic from the destination. The default value is **yes**.

**[ conntype={ lan| dialup| all}]**
Specifies whether the rule applies only to remote access or dial-up connections or to local area network (LAN) connections, or to all connections. The default value is **all**.

**[ actioninbound={ permit| block| negotiate}]**
Specifies the action that IPSec is required to take for inbound traffic. If **permit** is specified, traffic is received without negotiating or applying IP security. If **block** is specified, traffic is blocked. If **negotiate** is specified, IP security is used, with the list of security methods specified in the main mode and quick mode policies. The default value is **negotiate**.

**[ actionoutbound={ permit| block| negotiate}]**
Specifies the action that IPSec is required to take for outbound traffic. If **permit** is specified, traffic is sent without negotiating or applying IP security. If **block** is specified, traffic is blocked. If **negotiate** is specified, IP security is used, with the list of security methods specified in the main mode and quick mode policies. The default value is **negotiate**.

**[ srcmask={ *Mask*| *Prefix*}]**
Specifies the source address subnet mask or the prefix of the packets to be filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ dstmask={ *Mask*| *Prefix*}]**
Specifies the destination address subnet mask or the prefix value of the packets to be filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ tunneldstaddress={ *IPAddress*| *DNSName*}]**
Specifies whether the traffic is tunneled and, if it is, the IP address or DNS name of the tunnel destination (the computer or gateway on the other side of the tunnel).

**[ kerberos={ yes | no}]**
Specifies whether to use the Kerberos V5 protocol as an authentication method.

**[ psk=*String*]**
Specifies the string of characters to use for the preshared key, if a preshared key is used as an authentication method.

**[ rootca="*String*certmap:{ yes| no} excludecaname:{ yes| no} "]**

Specifies certificate authentication options, where:

***String***

Specifies the distinguished name of the certificate, if a certificate is used as an authentication method.

**certmap:{ yes| no}**

Specifies whether to enable certificate-to-account mapping. You can enable certificate-to-account mapping to verify that the certificate is being used by a trusted computer.

**excludecaname:{ yes| no}**

Specifies whether to exclude from the certificate request the list of trusted root CA names from which a certificate is accepted.

**Remarks**

- Do not create a rule with the name **all**. Doing this creates a conflict with the **netsh ipsec** option to select all IPSec rules (for example, **delete rule all**).
- If the filter action for both inbound and outbound traffic (**actioninbound** and **actionoutbound**) are set to **Permit** or **Block**, a quick mode filter is not required.
- If a tunnel rule is specified, **mirror** should be set to **no** (by default, **mirror** is set to **yes**). For IPSec tunnels, you must create two rules--one rule describes the traffic to be sent through the tunnel (outbound traffic) and the other describes the traffic to be received through the tunnel (inbound). Next, create two rules that use the inbound and outbound filter lists in your policy.
- To filter any packets sent from or to the computer, you can use **srcaddr=Me** or **dstaddr=Me**.
- To filter packets sent from or to any computer, you can use **srcaddr=Any** or **dstaddr=Any**.
- You can only use Kerberos V5 authentication or certificate-to-account mapping for computers that are members of an Active Directory domain.
- Although you can use only one preshared key for authentication, you can use multiple certificates by specifying the **rootca** parameter once, for each certificate you want to use.
- All certificate authentication parameters must be contained within quotation marks. Embedded quotation marks must be replaced by a backslash followed by an apostrophe (\').
- All string parameters are case-sensitive.
- The preference order of each authentication method is determined by the order in which it was specified in the command.
- If no authentication methods are specified, dynamic defaults are used. By default, IPSec policies use Kerberos V5 authentication. If the computer has a computer certificate, any root CA to which the computer certificate chains is also used for authentication.
- If **excludecaname:yes** is specified, the list of trusted root CAs is not sent as part of the certificate request, which prevents the potential disclosure of sensitive information about the trust relationships of a computer. To enhance security for computers connected to the Internet, specify this option.
- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. In addition, preshared keys are stored in plaintext.
- IPSec peers must have at least one common authentication method or communication will fail.

**delete all**

Deletes all IPSec policies, filters, and authentication methods, if possible, from the SPD.

**Syntax**

**delete all**

**Parameters**

None.

**delete mmpolicy**

Deletes the specified IPSec main mode policy, or all IPSec main mode policies, from the SPD.

**Syntax**

**delete mmpolicy name=** | **all**

**Parameters**

**name=** ***String*****| all**

Required. Specifies the name of the IPSec main mode policy to delete. Or, if **all** is specified, all IPSec main mode policies are deleted.

**Remarks**

- If a rule is associated with the main mode policy, you must delete the rule before you can delete the policy.
- All string parameters are case-sensitive.

**delete qmpolicy**

Deletes the specified IPSec quick mode policy, or all IPSec quick mode policies, from the SPD.

**Syntax**

**delete qmpolicy** [**name=**] | [**all**]

**Parameters**

**name=** ***String*****| all**

Required. Specifies the name of the IPSec quick mode policy to delete. Or, if **all** is specified, all IPSec quick mode policies are deleted.

**Remarks**

- If a rule is associated with the quick mode policy, you must delete the rule before you can delete the policy.
- All string parameters are case-sensitive.

**delete rule**

Deletes an IPSec rule from the SPD.

**Syntax**

**delete rule srcaddr= dstaddr= protocol= srcport= dstport= mirrored= conntype=**[**srcmask=**][**dstmask=**] [**tunneldstaddress=**]

**Parameters**

**srcaddr={ Me| Any|** ***IPAddress*****|** ***DNSName*****|** ***ServerType*****}**

Required. Specifies the source IP address, DNS name, or server type for the IP traffic. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**dstaddr={ Me| Any|** ***IPAddress*****|** ***DNSName*****|** ***ServerType*****}**

Required. Specifies the destination IP address, DNS name, or server type for the IP traffic. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**protocol={ ANY| ICMP| TCP| UDP| RAW|** ***Integer*** **}**

Required. Specifies the IP protocol used for the filter.

**srcport=** ***Port***

Required. Specifies the source port number of the packets being filtered. This option only applies if you are filtering TCP or UDP packets. A value of **0** matches filters set to a source port of **0** or **any**.

**dstport=** ***Port***

Required. Specifies the destination port number of the packets being filtered. This option only applies if you are filtering TCP or UDP packets. A value of **0** matches filters set to a destination port of **0** or **any**.

**mirrored={ yes| no}**
Required. Specifies whether the rule was created with mirrored filters.

**conntype={ lan| dialup| all}**
Required. Specifies whether the rule to be deleted applies only to remote access or dial-up connections or to local area network (LAN) connections, or to all connections.

**[ srcmask={ *Mask*| *Prefix*}]**
Specifies the source address subnet mask or the prefix of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ dstmask={ *Mask*| *Prefix*}]**
Specifies the destination address subnet mask or the prefix value of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ tunneldstaddress={ *IPAddress*| *DNSName*}]**
Specifies whether the traffic is tunneled and, if it is, the IP address or DNS name of the tunnel destination (the computer or gateway on the other side of the tunnel).

**Remarks**

- To filter any packets sent from or to the computer, you can use **srcaddr=Me** or **dstaddr=Me**.
- To filter packets sent from or to any computer, you can use **srcaddr=Any** or **dstaddr=Any**.

**set config**
Creates or modifies the following IPSec settings: IPSec diagnostics, default traffic exemptions, strong certificate revocation list (CRL) checking, IKE (Oakley) logging, logging intervals, computer startup security, and computer startup traffic exemptions.

**Syntax**
**set config** [**property=**] [**value=**]

**Parameters**
**[ property=]{ ipsecdiagnostics value=| ipsecexempt value=| ipseclogginterval value=| ikelogging value=| strongcrlcheck value=| bootmode value=| bootexemptions value=}**
Required. Specifies the name of the IPSec setting to be created or modified and a value for the setting,

where:

**ipsecdiagnostics value={ 0| 1| 2| 3| 4| 5| 6| 7}**
Specifies whether to enable IPSec diagnostic logging and, if so, which level of logging to provide. The default value is **0**, meaning that logging is disabled. If you change the value for this setting, you must restart the computer for the new value to take effect.

You can specify other values as follows, to enable different levels of logging:

- When **1** is specified, bad SPI packets (the total number of packets for which the Security Parameters Index or SPI was incorrect), IKE negotiation failures, IPSec processing failures, packets received with invalid packet syntax, and other errors are recorded in the System log. Unauthenticated hashes (with the exception of the "Clear text received when should have been secured" event) are logged as well.
- When **2** is specified, inbound per-packet drop events are recorded in the System log.
- When **3** is specified, level 1 and level 2 logging are performed. In addition, unexpected clear text events (packets that are sent or received in plaintext) are also recorded.
- When **4** is specified, outbound per-packet drop events are recorded in the System log.
- When **5** is specified, level 1 and level 4 logging are performed.

- When **6** is specified, level 2 and level 4 logging are performed.
- When **7** specified, all levels of logging are performed.

**ipsecexempt value={ 0| 1| 2| 3}**
Specifies whether to modify the default IPSec traffic exemption (traffic that is not matched against IPSec filters but is still permitted). The default value is **3**, meaning that only IKE traffic is exempted from IPSec filtering. If you change the value for this setting, you must restart the computer for the new value to take effect.

You can specify other values as follows:

- If **0** is specified, multicast, broadcast, RSVP, Kerberos, and IKE traffic is exempted from IPSec filtering.
- If **1** is specified, Kerberos and RSVP traffic is not exempted from IPSec filtering (multicast, broadcast, and IKE traffic is exempted).
- If **2** is specified, multicast and broadcast traffic is not exempted from IPSec filtering (RSVP, Kerberos, and IKE traffic is exempted).

**ipseclogginterval value={ *Integer*}**
Specifies the interval, in seconds, after which IPSec event logs are sent to the System log. For *Integer*, valid values range from **60** through **86400**. The default value is **3600**. If you change the value for this setting, you must restart the computer for the new value to take effect.

**ikelogging value={ 0| 1}**
Specifies whether to enable IKE (Oakley) logging, to generate details about the SA establishment process. The default value is **0**, meaning that IKE logging is disabled.

**strongcrlcheck value={ 0| 1| 2}**
Specifies the level of CRL checking to use. If **0** is specified, CRL checking is disabled. If **1** is specified, standard CRL checking is used, and certificate validation fails only if the certificate is determined to be revoked. If **2** is specified, strong CRL checking is used, and certificate validation fails if any CRL check error occurs. The default value is **1**.

**bootmode value={ stateful| block| permit}**
Specifies the action that IPSec is required to take when the computer starts. If **stateful** is specified, only the following traffic is permitted during computer startup: outbound traffic initiated by the computer during startup, inbound traffic that is sent in response to the outbound traffic, and DHCP traffic. If **block** is specified, all inbound and outbound traffic is blocked until a local IPSec policy or a domain-based IPSec policy is applied. If **permit** is specified, all traffic is transmitted and received. The default value is **stateful**. If you use either stateful filtering or if you specify that traffic be blocked during computer startup, you can also use the **bootexemptions** parameter to specify traffic types that you want to exempt from IPSec filtering during computer startup.

If you change the value for this setting, you must restart the computer for the new value to take effect.

**bootexemptions value= *Exemption1Exemption2***
Specifies one or more IPSec traffic exemptions from startup security, separated by spaces and defined by the following format for TCP and UDP traffic: *protocol:srcport:dstport:direction* and the following format for non-TCP/UDP traffic: *protocol:direction*, where:

***protocol* ={ ICMP| TCP| UDP| RAW| *Integer* }**
Specifies the IP protocol type to exempt from IPSec filtering during computer startup.

***srcport* = *Port***
Specifies the source port number of the packets to exempt from IPSec filtering during computer startup. A value of 0 means that any source port is exempted.

***dstport* = *Port***
Specifies the destination port number of the packets to exempt from IPSec filtering during computer startup. A value of 0 means that any destination port is exempted.

***direction* ={ inbound | outbound}**
Specifies the direction of the traffic to exempt from IPSec filtering during computer startup.

**Remarks**

- Use strong CRL checking (set **property=strongcrlcheck value=2**) if the CRL distribution point must be reachable on the network and certificates can only be validated if no CRL check error occurs.

- IPSec can only negotiate security associations for Kerberos traffic if your IPSec policy does not use Kerberos as the authentication method. If Kerberos is required for authentication, you must exempt Kerberos traffic by using the **ipsecexempt** parameter.

- In Windows 2000 and Windows XP, by default, all broadcast, multicast, Internet Key Exchange (IKE), Kerberos, and Resource Reservation Protocol (RSVP) traffic was exempted from IPSec filtering. In the Windows Server™ 2003 family, only IKE traffic is exempted from IPSec filtering by default. All other traffic types are now matched against IPSec filters, and you can configure block or permit filter actions specifically for multicast and broadcast traffic (IPSec does not negotiate security associations for multicast and broadcast traffic).

  As a result of this change in default behavior for IPSec, you should verify the behavior of IPSec policies designed for Windows 2000 or Windows XPand determine whether to configure explicit permit filters to permit specific traffic types. To restore the default behavior of Windows 2000 and Windows XPfor IPSec policies, edit the following registry key: **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IPSec**. Add a new DWORD value named **NoDefaultExempt** and assign to it a value of **0**. For more information about adding values to registry keys, see Add a value to a registry key entry.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

- Modifying IPSec traffic exemptions from startup security (that is, modifying the **bootexemptions=** parameter) will overwrite all previous exemptions from startup security.

**set mmpolicy**

Modifies an IPSec main mode policy and writes the changes to the SPD.

**Syntax**

**set mmpolicy name=** [**qmperm=**][**mmlifetime=**][**softsaexpirationtime=**][**mmsecmethods=**]

**Parameters**

**name=** ***String***

Required. Specifies the name of the IPSec main mode policy to modify.

**[ qmpermm=*Integer*]**

Specifies the number of times that master keying material is used to derive the session key. A value of 0 means that an unlimited number of quick mode SAs can be derived from the main mode SA.

**[ mmlifetime=*Integer*]**

Specifies the number of minutes after which a new master key is generated.

**[ softsaexpirationtime=*Integer*]**

Specifies the number of minutes after which an unprotected security association expires.

**[ mmsecmethods="*SecMeth1SecMeth2*"]**

Specifies one or more key exchange security methods, separated by spaces and defined by the format *ConfAlg-HashAlg-GroupNumb*, where:

***ConfAlg***

Specifies the encryption algorithm. *ConfAlg* can be **DES** or **3DES**.

***HashAlg***

Specifies the hash function. *HashAlg* can be **MD5** or **SHA1**.

***GroupNum***

Specifies the Diffie-Hellman group to be used for the base keying material. *GroupNumb* can be: **1** (low, protects with 768

bits of keying material), **2** (medium, protects with 1024 bits), and **3** (high, protects with 2048 bits).

**Remarks**

- IPSec peers must have at least one common key exchange security method (one that uses the same settings) or negotiations will fail.
- If the number of quick mode negotiations will exceed the value set for the number of quick mode negotiations per main mode negotiation during the main mode lifetime, a new main mode negotiation occurs.
- All string parameters are case-sensitive.

**set qmpolicy**

Modifies an IPSec quick mode policy and writes the changes to the SPD.

**Syntax**

**set qmpolicy name=** [**soft=**][**pfsgroup=**][**qmsecmethods=**]

**Parameters**

**name=** ***String***

Required. Specifies the name of the IPSec quick mode policy to modify.

**[ soft={ yes| no}]**

Specifies whether to fall back to unsecured communications with other computers that do not support IPSec, or when IPSec negotiations with an IPSec-capable computer fail.

**[ pfsgroup={ grp1| grp2| grp3| grpmm| nopfs}]**

Specifies the Diffie-Hellman group to use for session key PFS. If **grp1** is specified, Group 1 (low) is used. If **grp2** is specified, Group 2 (medium) is used. If **grp3** is specified, Group 2048 (high) is used. If **grpmm** is specified, the group value is taken from the current main mode settings. A value of **nopfs** means that session key PFS is disabled.

**[ qmsecmethods="*Neg1Neg2*"]**

Specifies one or more security methods, separated by spaces and defined by the following format:{**ESP** [*ConfAlg*,*AuthAlg*]**:***k/s* | **AH** [*HashAlg*]**:***k/s* | **AH** [*HashAlg*+**ESP***ConfAlg*,*AuthAlg*]**:***k/s*}]

Where:

***ConfAlg***

Specifies the encryption algorithm. *ConfigAlg* can be **DES** (Data Encryption Standard), **3DES**, or **none**.

***AuthAlg***

Specifies the integrity algorithm. *AuthAlg* can be **MD5** (Message Digest 5), **SHA1** (Secure Hash Algorithm 1), or **none**.

***HashAlg***

Specifies the hash function. *HashAlg* can be **MD5** or **SHA1**.

***k***

Specifies the session key lifetime in kilobytes. After the specified number of kilobytes of data is transferred, a new session key for the quick mode SA is generated. The default value is 100000 kilobytes.

***s***

Specifies the session key lifetime in seconds. The default value is 3600 seconds.

**Remarks**

- The preference order of each quick mode security method is determined by the order in which it was specified in the command.
- IPSec peers must have the same **pfsgroup** enabled (that is, both peers must use the same Diffie-Hellman group for session key PFS), or communication will fail.
- For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.

- All string parameters are case-sensitive.

**set rule**

Modifies an IPSec rule that defines a set of filters and writes the changes to the SPD.

**Syntax**

**set rule srcaddr= dstaddr= protocol= srcport= dstport= mirrored= conntype=** [**srcmask=**][**dstmask=**] [**tunneldstaddress=**][**mmpolicy=**][**qmpolicy=**][**actioninbound=**][**actionoutbound=**][**kerberos=**][**psk=**][**rootca=**]

**Parameters**

**srcaddr={ Me| Any|** ***IPAddress*****|** ***DNSName*****|** ***ServerType*****}**

Required. Specifies the source IP address, DNS name, or server type for the IP traffic. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**dstaddr={ Me| Any|** ***IPAddress*****|** ***DNSName*****|** ***ServerType*****}**

Required. Specifies the destination IP address, DNS name, or server type for the IP traffic. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**protocol={ ANY| ICMP| TCP| UDP| RAW|** ***Integer*** **}**

Required. Specifies the IP protocol used for the filter.

**srcport=** ***Port***

Required. Specifies the source port number of the packets being filtered. This option only applies if you are filtering TCP or UDP packets. A value of **0** matches filters set to a source port of **0** or **any**.

**dstport=** ***Port***

Required. Specifies the destination port number of the packets being filtered. This option only applies if you are filtering TCP or UDP packets. A value of **0** matches filters set to a destination port of **0** or **any**.

**mirrored={ yes| no}**

Required. Specifies whether the rule was created with mirrored filters.

**conntype={ lan| dialup| all}**

Required. Specifies whether the rule applies only to remote access or dial-up connections or to local area network (LAN) connections, or to all connections.

**[ srcmask={** ***Mask*****|** ***Prefix*****}]**

Specifies the source address subnet mask or the prefix of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ dstmask={** ***Mask*****|** ***Prefix*****}]**

Specifies the destination address subnet mask or the prefix value of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ tunneldstaddress={** ***IPAddress*****|** ***DNSName*****}]**

Specifies whether the traffic is tunneled and, if it is, the IP address or DNS name of the tunnel destination (the computer or gateway on the other side of the tunnel).

**[ mmpolicy=*****String*****]**

Specifies the name of the main mode policy.

**[ qmpolicy=*****String*****]**

Specifies the name of the quick mode policy.

**[ actioninbound={ permit| block| negotiate}]**

Specifies the action that IPSec is required to take for inbound traffic. If **permit** is specified, traffic is received without negotiating or applying IP security. If **block** is specified, traffic is blocked. If **negotiate** is specified, IP security is used, with the list of security methods specified in the main mode and quick mode policies.

**[ actionoutbound={ permit| block| negotiate}]**

Specifies the action that IPSec is required to take for outbound traffic. If **permit** is specified, traffic is sent without negotiating or applying IP security. If **block** is specified, traffic is blocked. If **negotiate** is specified, IP security is used, with the list of security methods specified in the main mode and quick mode policies.

**[ kerberos={ yes | no}]**

Specifies whether to use the Kerberos V5 protocol as an authentication method.

**[ psk=*String*]**
Specifies the string of characters to use for the preshared key, if a preshared key is used as an authentication method.

**[ rootca="*String*certmap:{ yes| no} excludecaname:{ yes| no} "]**
Specifies certificate authentication options, where:

***String***
Specifies the distinguished name of the certificate, if a certificate is used as an authentication method.

**certmap:{ yes| no}**
Specifies whether to enable certificate-to-account mapping. You can enable certificate-to-account mapping to verify that the certificate is being used by a trusted computer.

**excludecaname:{ yes| no}**
Specifies whether to exclude from the certificate request the list of trusted root CA names from which a certificate is accepted.

**Remarks**

- You can modify the following parameters: **mmpolicy=**, **qmpolicy=**, **actioninbound=**, and **actionoutbound=**. All other parameters are used to identify the rule that you want to modify, and therefore they cannot be modified.
- If the filter action for both inbound and outbound traffic (**actioninbound** and **actionoutbound**) are set to **Permit** or **Block**, a quick mode filter is not required.
- If a tunnel rule is specified, **mirror** should be set to **no** (by default, **mirror** is set to **yes**). For IPSec tunnels, you must create two rules: one rule describes the traffic to be sent through the tunnel (outbound traffic) and the other describes the traffic to be received through the tunnel (inbound). Next, create two rules that use the inbound and outbound filter lists in your policy.
- To filter any packets sent from or to the computer, you can use **srcaddr=Me** or **dstaddr=Me**.
- To filter packets sent from or to any computer, you can use **srcaddr=Any** or **dstaddr=Any**.
- You can only use Kerberos V5 authentication or certificate-to-account mapping for computers that are members of an Active Directory domain.
- Although you can use only one preshared key for authentication, you can use multiple certificates by specifying the **rootca** parameter once, for each certificate you want to use.
- All certificate authentication parameters must be contained within quotation marks. Embedded quotation marks must be replaced by a backslash followed by an apostrophe (\').
- All string parameters are case-sensitive.
- The preference order of each authentication method is determined by the order in which it was specified in the command.
- If no authentication methods are specified, dynamic defaults are used. By default, IPSec policies use Kerberos V5 authentication. If the computer has a computer certificate, any root CA to which the computer certificate chains is also used for authentication.
- If **excludecaname:yes** is specified, the list of trusted root CAs is not sent as part of the certificate request, which prevents the potential disclosure of sensitive information about the trust relationships of a computer. To enhance security for computers connected to the Internet, specify this option.
- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. In addition, preshared keys are stored in plaintext.
- IPSec peers must have at least one common authentication method or communication will fail.

- Modifying authentication methods will overwrite all previous authentication methods, even if the previous authentication methods were different. For example, if **kerberos=yes** and **psk=yes** were previously specified, and then you specify **kerberos=no**, the **psk=yes** parameter will also be overwritten, and preshared key authentication will no longer be used.

**show all**
Displays configuration information for all IPSec policies, filters, statistics, and security associations in the SPD.

**Syntax**
**show all** [**resolvedns=**]

**Parameters**
**[ resolvedns={ yes | no}]**
Specifies whether to resolve the Domain Name System (DNS) or NETBIOS computer name associated with an IP address when displaying sources or destinations.

**Remarks**

- Because the **show all** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show all** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec dynamic show all**

  3. To stop sending output and close the file, type:

     **set file close**

  **If you are not in the netsh environment**

  - At the command prompt, type:

    **netsh ipsec dynamic show all >***FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.

  - Use Task Manager to end the **Netsh** program.

**show config**
Displays values for the following IPSec settings: IPSec diagnostics, default traffic exemptions, strong certificate revocation list (CRL) checking, IKE (Oakley) logging, logging intervals, computer startup security, and computer startup traffic exemptions.

**Syntax**
**show config**

**Parameters**
None.

**Remarks**

- To save output in a text file for the **show config** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec dynamic show config**

  3. To stop sending output and close the file, type:

     **set file close**

  **If you are not in the netsh environment**

  - At the command prompt, type:

    **netsh ipsec dynamic show config >***FileName*.txt

**show mmfilter**

Displays configuration information for the specified IPSec main mode filter, or for all IPSec main mode filters, in the SPD.

**Syntax**

**show mmfilter name=** | **all** [**type=**] **srcaddr=dstadd=** [**srcmask=**][**dstmask=**] [**resolvedns=**]

**Parameters**

**name= *String*| all**
Required. Specifies the name of the IPSec main mode filter to display. Or, if **all** is specified, all IPSec main mode filters are displayed.

**type={ generic| specific}**
Specifies whether to display generic or specific main mode filters. The default value is **generic**.

**[ srcaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}]**
Specifies the source IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ dstaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}]**
Specifies the destination IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ srcmask={ *Mask*| *Prefix*}]**
Specifies the source address subnet mask or the prefix of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ dstmask={ *Mask*| *Prefix*}]**
Specifies the destination address subnet mask or the prefix value of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ resolvedns={ yes | no}]**
Specifies whether to resolve the Domain Name System (DNS) or NETBIOS computer name associated with an IP address when displaying sources or destinations. The default value is **no**.

**Remarks**

- All string parameters are case-sensitive.

- Because the **show mmfilter** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show mmfilter** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec dynamic show mmfilter***Name* | **all**

  3. To stop sending output and close the file, type:

     **set file close**

  **If you are not in the netsh environment**

  - At the command prompt, type:

    **ipsec dynamic show mmfilter***Name* | **all >=***FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.

  - Use Task Manager to end the **Netsh** program.

**show mmpolicy**

Displays configuration information for the specified IPSec main mode policy, or for all IPSec main mode policies, in the SPD.

**Syntax**

**show mmpolicy name=** | **all**

**Parameters**

**name=** ***String*| all**

Required. Specifies the name of the IPSec main mode policy to display. Or, if **all** is specified, all IPSec main mode policies are displayed.

**Remarks**

- All string parameters are case-sensitive.

- Because the **show mmpolicy** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show mmpolicy** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

2. Then type:

   **ipsec dynamic show mmpolicy***Name* | **all**

3. To stop sending output and close the file, type:

   **set file close**

**If you are not in the netsh environment**

- At the command prompt, type:

  **ipsec dynamic show mmpolicy***Name* | **all >***FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.

  - Use Task Manager to end the **Netsh** program.

**show mmsas**

Displays the IPSec main mode security associations for the specified source and destination addresses, or all IPSec main mode security associations, in the SPD.

**Syntax**

**show mmsas** [**all**] [**srcaddr=**][**dstaddr=**][**format=**] [**resolvedns=**]

**Parameters**

**[ all]**
Specifies that all main mode security associations are displayed.

**[ srcaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}]**
Specifies the source IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ dstaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}]**
Specifies the destination IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ format={ list| table}]**
Specifies whether to display IPSec configuration information in screen or tab-delimited format. The default value is **list**, meaning that output is displayed in screen format.

**[ resolvedns={ yes | no}]**
Specifies whether to resolve the Domain Name System (DNS) or NETBIOS computer name associated with an IP address when displaying sources or destinations. The default value is **no**.

**Remarks**

- If no parameters are specified, all main mode security associations are displayed.

- All string parameters are case-sensitive.

- Because the **show mmsas** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show mmsas** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

1. At the netsh prompt, type:

   **set file open** *FileName*.txt

2. Then type:

   **ipsec dynamic show mmsas**

3. To stop sending output and close the file, type:

   **set file close**

**If you are not in the netsh environment**

- At the command prompt, type:

  **netsh ipsec dynamic show mmsas >***FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:
  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.
  - Use Task Manager to end the **Netsh** program.

**show qmfilter**

Displays configuration information for the specified quick mode filter, or for all quick mode filters, in the SPD.

**Syntax**

**show qmfilter name=** | **all** [**type=**] [**srcaddr=**][**dstaddr=**][**srcmask=**][**dstmask=**][**protocol=**][**srcport=**][**dstport=**] [**actioninbound=**][**actionoutbound=**][**resolvedns=**]

**Parameters**

**name=** ***String*| all**
Required. Specifies the name of the IPSec quick mode filter to display, or, if **all** is specified, that all IPSec quick mode filters are displayed.

**[ type={ generic| specific}]**
Specifies whether to display generic or specific quick mode filters. The default value is **generic**.

**[ srcaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}]**
Specifies the source IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ dstaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}]**
Specifies the destination IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ srcmask={ *Mask*| *Prefix*}]**
Specifies the source address subnet mask or the prefix of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ dstmask={ *Mask*| *Prefix*}]**
Specifies the destination address subnet mask or the prefix value of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ protocol={ ANY| ICMP| TCP| UDP| RAW| *Integer* }]**
Specifies the IP protocol if, in addition to addressing information, a specific IP protocol is filtered. The default value is **ANY**, meaning all protocols are used for the filter.

**[ srcport=*Port*]**
Specifies the source port number of the packets being filtered. This option only applies if you are filtering TCP or UDP

packets. If 0 is specified, packets sent from any port are filtered. The default is any.

**[ dstport=*Port*]**
Specifies the destination port number of the packets being filtered. This option only applies if you are filtering TCP or UDP packets. If 0 is specified, packets sent to any port are filtered. The default is any.

**[ actioninbound={ permit| block| negotiate}]**
Specifies the action that IPSec is required to take for inbound traffic. The default value is **negotiate**.

**[ actionoutbound={ permit| block| negotiate}]**
Specifies the action that IPSec is required to take for outbound traffic. The default value is **negotiate**

**[ resolvedns={ yes | no}]**
Specifies whether to resolve the Domain Name System (DNS) or NETBIOS computer name associated with an IP address when displaying sources or destinations. The default value is **no**.

**Remarks**

- All string parameters are case-sensitive.
- Because the **show qmfilter** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show qmfilter** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec dynamic show qmfilter** *Name* | **all**

  3. To stop sending output and close the file, type:

     **set file close**

  **If you are not in the netsh environment**

  - At the command prompt, type:

    **netsh ipsec dynamic show qmfilter** *Name* | **all** >*FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:
  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.
  - Use Task Manager to end the **Netsh** program.

**show qmpolicy**
Displays configuration information for the specified IPSec quick mode policy, or for all IPSec quick mode policies, in the SPD.

**Syntax**
**show qmpolicy name=** | **all**

**Parameters**
**name=** ***String*| all**
Required. Specifies the name of the IPSec quick mode policy to display. Or, if **all** is specified, all IPSec quick mode policies are

displayed.

**Remarks**

- All string parameters are case-sensitive.
- Because the **show qmpolicy** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show qmpolicy** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec dynamic show qmpolicy***Name* | **all**

  3. To stop sending output and close the file, type:

     **set file close**

  **If you are not in the netsh environment**

  - At the command prompt, type:

    **netsh ipsec dynamic show qmpolicy***Name* | **all** >*FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:
  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.
  - Use Task Manager to end the **Netsh** program.

**show qmsas**
Displays the IPSec quick mode security associations for the specified source and destination addresses, or all IPSec quick mode security associations, in the SPD.

**Syntax**
**show qmsas** [**all**][**srcaddr=**][**dstaddr=**][**protocol=**][**format=**][**resolvedns=**]

**Parameters**
**[ all]**
Specifies that all IPSec quick mode security associations are displayed.

**[ srcaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}]**
Specifies the source IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ dstaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*]**
Specifies the destination IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ protocol={ ANY| ICMP| TCP| UDP| RAW| *Integer* }]**
Specifies the IP protocol if, in addition to addressing information, a specific IP protocol is being used for the security association. The default value is **ANY**, meaning all protocols are used for the security association.

**[ format={ list| table}]**
Specifies whether to display IPSec configuration information in screen or tab-delimited format. The default value is **list**, meaning that output is displayed in screen format.

**[ resolvedns={ yes | no}]**
Specifies whether to resolve the Domain Name System (DNS) or NETBIOS computer name associated with an IP address when displaying sources or destinations. The default value is **no**.

**Remarks**

- Because the **show qmsas** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show qmsas** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec dynamic show qmsas all**

  3. To stop sending output and close the file, type:

     **set file close**

  **If you are not in the netsh environment**

  - At the command prompt, type:

    **netsh ipsec dynamic show qmsas all** >*FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.

  - Use Task Manager to end the **Netsh** program.

**show rule**
Displays configuration information for one or more IPSec rules in the SPD.

**Syntax**
**show rule** [**type=**][**srcaddr=**] [**dstaddr=**][**srcmask=**][**dstmask=**] [**protocol=**][**srcport=**][**dstport=**][**actioninbound=**] [**actionoutbound=**][**resolvedns=**]

**Parameters**
**[ type={ transport| tunnel}]**
Specifies whether to display a transport rule or a tunnel rule. The default value is to display all rules.

**[ srcaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}]**
Specifies the source IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ dstaddr={ Me| Any| *IPAddress*| *DNSName*| *ServerType*}]**
Specifies the destination IP address, DNS name, or server type for the IP traffic being filtered. You can use **WINS**, **DNS**, **DHCP**, or **gateway** for *ServerType*.

**[ srcmask={ *Mask*| *Prefix*}]**
Specifies the source address subnet mask or the prefix of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ dstmask={ *Mask*| *Prefix*}]**
Specifies the destination address subnet mask or the prefix value of the packets being filtered. You can specify a prefix value in the range of 1 through 32. The default value is the mask of 255.255.255.255.

**[ protocol={ ANY| ICMP| TCP| UDP| RAW| *Integer* }]**
Specifies the IP protocol if, in addition to addressing information, a specific IP protocol is used for the rule. The default value is **ANY**, meaning all protocols are used for the rule.

**[ srcport=*Port*]**
Specifies the source port number of the packets being filtered. This option only applies if you are filtering TCP or UDP packets. If 0 is specified, packets sent from any port are filtered. The default is any.

**[ dstport=*Port*]**
Specifies the destination port number of the packets being filtered. This option only applies if you are filtering TCP or UDP packets. If 0 is specified, packets sent to any port are filtered. The default is any.

**[ actioninbound={ permit| block| negotiate}]**
Specifies the action that IPSec is required to take for inbound traffic. The default value is **negotiate**.

**[ actionoutbound={ permit| block| negotiate}]**
Specifies the action that IPSec is required to take for outbound traffic. The default value is **negotiate**.

**[ resolvedns={ yes | no}]**
Specifies whether to resolve the Domain Name System (DNS) or NETBIOS computer name associated with an IP address when displaying sources or destinations. The default value is **no**.

**Remarks**

- All string parameters are case-sensitive.

- Because the **show rule** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show rule** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec dynamic show rule**

  3. To stop sending output and close the file, type:

     **set file close**

  **If you are not in the netsh environment**

  - At the command prompt, type:

    **netsh ipsec dynamic show rule >***FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:

- Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.
- Use Task Manager to end the **Netsh** program.

**show stats**
Displays main mode and quick mode statistics for IPSec.

**Syntax**
**show stats** [**type=**]

**Parameters**
**[ type=all| ike| ipsec]**
Specifies the IPSec statistics to display. If **all** is specified, IPSec main mode and quick mode statistics are displayed. If **ike** is specified, only IPSec main mode statistics are displayed. If **ipsec** is specified, only IPSec quick mode statistics are displayed.

**Remarks**

- Because the **show stats** command can result in lengthy, rapidly scrolling output, consider saving the output in a text file, unless you only need to view limited portions.

  To save output in a text file for the **show stats** command, do either of the following:

  **If you are in the netsh environment (netsh>)**

  1. At the netsh prompt, type:

     **set file open** *FileName*.txt

  2. Then type:

     **ipsec dynamic show stats**

  3. To stop sending output and close the file, type:

     **set file close**

  **If you are not in the netsh environment**

  - At the command prompt, type:

    **netsh ipsec dynamic show stats >***FileName*.txt

- To stop the output of IPSec configuration information, you must exit **Netsh** by doing one of the following:
  - Close the **Netsh** window by clicking the **X** icon in the upper-right corner of the window.
  - Use Task Manager to end the **Netsh** program.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |

| | |
|---|---|
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

# Netsh commands for Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for Network Bridge**
For adapters that are part of the Network Bridge, you can use commands in the **netsh bridge** context to enable or disable Layer 3 compatibility mode and show configuration information for these adapters.

You can run these commands from the Windows Server 2003 family command prompt or from the **netsh bridge** context. To successfully run these commands at the Windows Server 2003 family command prompt, you must type **netsh bridge** before typing the commands and parameters as they appear in this topic.

For more information about **netsh**, see Netsh overview and Enter a netsh context.

To view the command syntax, click a command:

- show adapter
- set adapter

### Note

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**show adapter**
Displays adapter identification, adapter names, and the state of the Layer 3 compatibility mode of adapters that are part of Network Bridge.

**Syntax**
**show adapter** [[**id=**]*AdapterID*]

**Parameters**
**[[ id=] *AdapterID*]**
Optional. Specifies the information to display for a specific adapter.

***/?***
Displays help at the command prompt.

**Examples**
In the first example, this command lists all of the adapters that are part of Network Bridge, and provides adapter ID numbers, friendly names, and the state of the Layer 3 compatibility mode that corresponds to each adapter.

In the second example, this command lists the adapter ID, friendly name, and the state of the Layer 3 compatibility mode information for adapter 2.

**show adapter**

**show adapter 2**

**set adapter**
This command modifies the configuration of a specified adapter that is part of Network Bridge by setting the state of the adapter to either enable or disable network layer (Layer 3) compatibility mode.

**Syntax**
**set adapter** [**id=**]*AdapterID* [**forcecompatmode=**]{**enable** | **disable**}

**Parameters**
**[ id=] *AdapterID***
Required. Specifies which adapter that is part of Network Bridge to configure.

**[ forcecompatmode=]{ enable| disable}**
Required. Specifies whether Layer 3 mode compatibility is enabled or disabled for the specified adapter.

**/?**
Displays help at the command prompt.

**Remarks**

- This command should only be run if an enabled Network Bridge does not forward network traffic to all network segments. This condition can exist if one or more network adapters (typically wireless) included in the bridge do not support promiscuous mode.

- For more information about determining adapter ID values, see show adapter.

**Examples**
In this example, the command is used to force adapter 2 to run in Layer 3 compatibility mode.

**set adapter 2 forcecompatmode=enable**

## Notes

- The **install** and **uninstall** commands are not supported in the **netsh bridge** context. Network Bridge can only be set up or removed through Network Connections. For more information, see Set up Network Bridge or Remove Network Bridge.

- The **dump** command is not supported in the **netsh bridge** context.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Network Bridge
Remove Network Bridge
Set up Network Bridge
Command-line reference A-Z
Command shell overview

# Netsh commands for remote access (ras)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for remote access (ras)**
You can administer remote access servers by typing commands at the command prompt for the **Netsh ras** context. By using the **Netsh ras** command prompt, you can administer servers more quickly over slow network connections, and you can create scripts that automate the administration of many servers.

For more information on Netsh, see Netsh overview and Enter a netsh context.

To view the command syntax, click a command:

- help
- show activeservers
- show client
- dump
- show tracing
- set tracing
- show authmode
- set authmode
- add authtype
- delete authtype
- show authtype
- add link
- delete link
- show link
- add multilink
- delete multilink
- show multilink
- add registeredserver
- delete registeredserver
- show registeredserver
- show user

- set user
- Netsh ras diagnostics context commands
- Netsh ras IP context commands
- Netsh ras AppleTalk context commands
- Netsh ras AAAA context commands

**help**
Displays command-line help for commands you can use in the Netsh ras context.

**Syntax**
**help**

**Parameters**
none

**Remarks**

- You can also display command-line help by typing any of the following at the command prompt for the Netsh ras context: **?**, **/?**, or **-?**.

**show_activeservers**
Displays a list of remote access server advertisements.

**Syntax**
**show activeservers**

**Parameters**
none

**show_client**
Lists remote access clients connected to this server.

**Syntax**
**show client**

**Parameters**
none

**dump**
Displays the configuration of the remote access server in script form.

**Syntax**
**dump**

**Parameters**
none

**show_tracing**
Shows whether tracing is enabled for the specified component. Used without parameters, **show tracing** lists all installed components and whether tracing is enabled for each.

**Syntax**
**show tracing** [*Component*]

**Parameters**
***Component***
Specifies the component for which to display information.

**set_tracing**
Enables or disables tracing for the specified component.

**Syntax**

**set tracing** *Component* {**enabled** | **disabled**}

**Parameters**

***Component***

Required. Specifies the component for which you want to enable or disable tracing. Use ***** to specify all components.

**{ enabled| disabled}**

Required. Specifies whether to enable or disable tracing for the specified component.

**Remarks**

- To see a list of all installed components, use the **show tracing** command without parameters.

**Examples**

To set tracing for the PPP component, type:

**set tracing ppp enabled**

**show_authmode**

Shows whether dial-up clients using certain types of devices should be authenticated.

**Syntax**

**show authmode**

**Parameters**

none

**set_authmode**

Specifies whether dial-up clients using certain types of devices should be authenticated.

**Syntax**

**set authmode** {**standard** | **nodcc** | **bypass**}

**Parameters**

**{ standard| nodcc| bypass}**

Required. Specifies whether dial-up clients using certain types of devices should be authenticated. The parameter **standard** specifies that clients using any type of device should be authenticated. The parameter **nodcc** specifies that clients using any type of device except a direct-connect device should be authenticated. The parameter **bypass** specifies that no clients should be authenticated.

**add_authtype**

Adds an authentication type to the list of types through which the remote access server should attempt to negotiate authentication.

**Syntax**

**add authtype** {**pap** | **spap** | **md5chap** | **mschap** | **mschapv2** | **eap**}

**Parameters**

**{ pap| spap| md5chap| mschap| mschapv2| eap}**

Required. Specifies which authentication type to add to the list of types through which the remote access server should attempt to negotiate authentication. The **pap** parameter specifies that the remote access server should use the Password Authentication Protocol (plaintext). The **spap** parameter specifies that the remote access server should use the Shiva Password Authentication Protocol. The **md5chap** parameter specifies that the remote access server should use the Challenge Handshake Authentication Protocol (using the Message Digest 5 hashing scheme to encrypt the response). The **mschap** parameter specifies that the remote access server should use the Microsoft Challenge-Handshake Authentication Protocol. The **mschapv2** parameter specifies that the remote access server should use Version 2 of MSCHAP. The **eap** parameter specifies that the remote access server should use Extensible Authentication Protocol.

**Remarks**

- The remote access server will attempt to negotiate authentication by using protocols in order from the most secure to the least secure. After both the client and the server have agreed on an authentication type, PPP negotiation proceeds according to the appropriate RFCs.

**delete_authtype**

Deletes an authentication type from the list of types through which the remote access server should attempt to negotiate

authentication.

**Syntax**

**delete authtype**{**pap** | **spap** | **md5chap** | **mschap** | **mschapv2** | **eap**}

**Parameters**

**{ pap| spap| md5chap| mschap| mschapv2| eap}**

Required. Specifies which authentication type to delete from the list of types through which the remote access server should attempt to negotiate authentication. The **pap** parameter specifies that the remote access server should not use the Password Authentication Protocol (plaintext). The **spap** parameter specifies that the remote access server should not use the Shiva Password Authentication Protocol. The **md5chap** parameter specifies that the remote access server should not use the Challenge Handshake Authentication Protocol (using the Message Digest 5 hashing scheme to encrypt the response). The **mschap** parameter specifies that the remote access server should not use the Microsoft Challenge-Handshake Authentication Protocol. The **mschapv2** parameter specifies that the remote access server should not use Version 2 of MSCHAP. The **eap** parameter specifies that the remote access server should not use Extensible Authentication Protocol.

**show_authtype**

Lists the authentication type (or types) that the remote access server uses to attempt to negotiate authentication.

**Syntax**

**show authtype**

**Parameters**

none

**add_link**

Adds a link property to the list of link properties PPP should negotiate.

**Syntax**

**add link** {**swc** | **lcp**}

**Parameters**

**{ swc| lcp}**

Required. Specifies which link property to add to the list of link properties PPP should negotiate. The parameter **swc** specifies that software compression (MPPC) should be added. The parameter **lcp** specifies that Link Control Protocol extensions from the PPP suite of protocols should be added.

**delete_link**

Deletes a link property from the list of link properties PPP should negotiate.

**Syntax**

**delete link** {**swc** | **lcp**}

**Parameters**

**{ swc| lcp}**

Required. Specifies which link property to delete from the list of link properties PPP should negotiate. The parameter **swc** specifies that software compression (MPPC) should be deleted. The parameter **lcp** specifies that Link Control Protocol extensions from the PPP suite of protocols should be deleted.

**show link**

Displays the link properties PPP should negotiate.

**Syntax**

**show link**

**Parameters**

none

**add multilink**

Adds a multilink type to the list of multilink types PPP should negotiate.

**Syntax**

**add multilink** {**multi** | **bacp**}

**Parameters**

**{ multi| bacp}**

Required. Specifies which multilink type to add to the list of multilink types PPP should negotiate. The parameter **multi** specifies that multilink PPP sessions should be added. The parameter **bacp** specifies that Bandwidth Allocation Control

Protocol should be added.

**delete multilink**
Deletes a multilink type from the list of multilink types PPP should negotiate.

**Syntax**
**delete multilink** {**multi** | **bacp**}

**Parameters**
**{ multi| bacp}**
Required. Specifies which multilink type to delete from the list of multilink types PPP should negotiate. The parameter **multi** specifies that multilink PPP sessions should be deleted. The parameter **bacp** specifies that Bandwidth Allocation Control Protocol should be deleted.

**show multilink**
Shows the multilink types PPP should negotiate.

**Syntax**
**show multilink**

**Parameters**
none

**add registeredserver**
Registers the specified server as a remote access server in the specified Active Directory domain. Used without parameters, **add registeredserver** registers the computer from which you type the command in its primary domain.

**Syntax**
**add registeredserver** [[**domain=**]*DomainName*] [[**server=**]*ServerName*]

**Parameters**
**[ domain=] *DomainName***
Specifies, by domain name, the domain in which to register the server. If you do not specify a domain, the server is registered in its primary domain.

**[ server=] *ServerName***
Specifies, by DNS name or IP address, the server to register. If you do not specify a server, the computer from which you type the command is registered.

**delete registeredserver**
Deletes the registration of the specified server as a remote access server from the specified Active Directory domain. Used without parameters, **delete registeredserver** deletes the registration of the computer from which you type the command from its primary domain.

**Syntax**
**delete registeredserver** [[**domain=**]*DomainName*] [[**server=**]*ServerName*]

**Parameters**
**[ domain=] *DomainName***
Specifies, by domain name, the domain from which to remove the registration. If you do not specify a domain, the registration is removed from the primary domain of the computer from which you type the command.

**[ server=] *ServerName***
Specifies, by IP address or DNS name, the server whose registration you want to remove. If you do not specify a server, the registration is removed for the computer from which you type the command.

**show registeredserver**
Displays status information about the specified server registered as a remote access server in the specified Active Directory domain. Used without parameters, **show registeredserver** displays information about the computer from which you type the command in its primary domain.

**Syntax**
**show registeredserver** [[**domain=**]*DomainName*] [[**server=**]*ServerName*]

**Parameters**
**[ domain=] *DomainName***
Specifies, by domain name, the domain in which the server about which you want to display information is registered. If you do not specify a domain, information is displayed about the server as it is registered in the primary domain of the computer

from which you type the command.

**[ server=] *ServerName***
Specifies, by IP address or DNS name, the server about which you want to display information. If you do not specify a server, information about the computer from which you typed the command is displayed.

**show user**
Displays the properties of a specified remote access user or users. Used without parameters, **show user** displays the properties of all remote access users.

**Syntax**
**show user** [**name=***UserName*] [[**mode=**] {**permit** | **report**}]

**Parameters**
**name= *UserName***
Specifies, by logon name, the user whose properties you want to display. If you do not specify a user, the properties of all users are displayed.

**mode={ permit| report}**
Specifies whether to show properties for all users or only those whose dial-up permission is set to permit. The **permit** parameter specifies that properties should be displayed only for users whose dial-up permission is permit. The **report** parameter specifies that properties should be displayed for all users.

**setuser**
Sets the properties of the specified remote access user.

**Syntax**
**set user** [**name=**]*UserName* [**dialin**] {**permit** | **deny** | **policy**} [**cbpolicy**] {**none** | **caller** | **admin** [**cbnumber=**]*CallbackNumber*}

**Parameters**
**name= *UserName***
Required. Specifies, by logon name, the user for which you want to set properties.

**[ dialin] { permit| deny| policy}**
Required. Specifies under what circumstances the user should be allowed to connect. The **permit** parameter specifies that the user should always be allowed to connect. The **deny** parameter specifies that the user should never be allowed to connect. The **policy** parameter specifies that remote access policies should determine whether the user is allowed to connect.

**[ cbpolicy] { none| caller| admin[ cbnumber=] *CallbackNumber*}**
Required. Specifies the callback policy for the user. The callback feature saves the user the cost of the phone call used to connect to a remote access server. The **none** parameter specifies that the user should not be called back. The **caller** parameter specifies that the user should be called back at a number specified by the user at connection time. The **admin** parameter specifies that the user should be called back at the number specified by the *CallbackNumber* parameter.

**Remarks**

- For users in a mixed-mode domain, the **policy** parameter and the **deny** parameter are equivalent.

**Examples**
To allow GuestUser to connect and be called back at 4255550110, type:

**set user guestuser permit admin 4255550110**

**Netsh ras diagnostics context commands**
The following commands are specific to the **ras diagnostics** context within the Netsh environment.

To view the command syntax, click a command:

- dump
- show installation
- show logs

- show configuration
- show all
- show cmtracing
- set cmtracing
- show modemtracing
- set modemtracing
- show rastracing
- set rastracing
- show securityeventlog
- set securityeventlog
- show tracefacilities
- set tracefacilities

**dump**
Displays the configuration of Remote Access Diagnostics in script form.

**Syntax**
**dump**

**Parameters**
none

**show installation**
Creates a Remote Access Diagnostic Report that includes only diagnostics results for Information Files, Installation Check, Installed Networking Components, and Registry Check and delivers the report to a location you specify.

**Syntax**
**show installation** [**type=**] {**file** | **email**} [**destination=**] {*FileLocation* | *E-mailAddress*} [[**compression=**] {**enabled** | **disabled**}] [[**hours=**] *NumberOfHours*] [[**verbose=**] {**enabled** | **disabled**}]

**Parameters**
**[ type=] { file| email}**
Required. Specifies whether you want to save the report to a file or you want to send it to an e-mail address.

**[ destination=] { *FileLocation*| *E-mailAddress*}**
Required. Specifies the full path and file name to which the report should be saved or the full e-mail address to which the report should be sent.

**[[ compression=] { enabled| disabled}]**
Specifies whether to compress the report into a .cab file. If you do not specify this parameter, the report is compressed if you send it to an e-mail address but not if you save it to a file.

**[[ hours=] *NumberOfHours*]**
Specifies the number of past hours for which to show activity in the report. This parameter must be an integer between 1 and 24. If you do not specify this parameter, all past information is included.

**[[ verbose=] { enabled| disabled}]**
Specifies the amount of data to include in the report. If you do not specify this parameter, only minimal data is included.

**Remarks**

- The Remote Access Diagnostic Report is saved as an HTML file (.htm).

#### show logs

Creates a Remote Access Diagnostic Report that contains only diagnostics results for Tracing Logs, Modem Logs, Connection Manager Logs, IP Security Log, Remote Access Event Logs, and Security Event Logs and delivers the report to a location you specify.

**Syntax**

**show logs** [**type=**] {**file** | **email**} [**destination=**] {*FileLocation* | *E-mailAddress*} [[**compression=**] {**enabled** | **disabled**}] [[**hours=**] *NumberOfHours*] [[**verbose=**] {**enabled** | **disabled**}]

**Parameters**

**[ type=] { file| email}**
Required. Specifies whether you want to save the report to a file or you want to sent it to an e-mail address.

**[ destination=] { *FileLocation*| *E-mailAddress*}**
Required. Specifies the full path and file name to which the report should be saved or the full e-mail address to which the report should be sent.

**[[ compression=] { enabled| disabled}]**
Specifies whether to compress the report into a .cab file. If you do not specify this parameter, the report is compressed if you send it to an e-mail address but not if you save it to a file.

**[[ hours=] *NumberOfHours*]**
Specifies the number of past hours for which to show activity in the report. This parameter must be an integer between 1 and 24. If you do not specify this parameter, all past information will be included in the report.

**[[ verbose=] { enabled| disabled}]**
Specifies the amount of data to include in the report. If you do not specify this parameter, minimal data is included.

**Remarks**

- The Remote Access Diagnostic Report is saved as an HTML file (.htm).

#### show configuration

Creates a Remote Access Diagnostic Report that includes only diagnostics results for Installed Devices, Process Information, Command-Line Utilities, and Phone Book Files and delivers the report to a location you specify.

**Syntax**

**show configuration** [**type=**] {**file** | **email**} [**destination=**] {*FileLocation* | *E-mailAddress*} [[**compression=**] {**enabled** | **disabled**}] [[**hours=**] *NumberOfHours*] [[**verbose=**] {**enabled** | **disabled**}]

**Parameters**

**[ type=] { file| email}**
Required. Specifies whether the report should be saved to a file or sent to an e-mail address.

**[ destination=] { *FileLocation*| *E-mailAddress*}**
Required. Specifies the full path and file name to which the report should be saved or the full e-mail address to which the report should be sent.

**[[ compression=] { enabled| disabled}]**
Specifies whether to compress the report into a .cab file. If you do not specify this parameter, the report is compressed if you send it to an e-mail address but not if you save it to a file.

**[[ hours=] *NumberOfHours*]**
Specifies the number of past hours for which to show activity in the report. This parameter must be an integer between 1 and 24. If you do not specify this parameter, all past information is included.

**[[ verbose=] { enabled| disabled}]**
Specifies the amount of data to include in the report. If you do not specify this parameter, minimal data is included.

**Remarks**

- The Remote Access Diagnostic Report is saved as an HTML file (.htm).

**show all**
Creates a Remote Access Diagnostic Report for all remote access logs and delivers the report to a location you specify.

**Syntax**
**show all** [**type=**] {**file** | **email**} [**destination=**] {*FileLocation* | *E-mailAddress*} [[**compression=**] {**enabled** | **disabled**}] [[**hours=**] *NumberOfHours*] [[**verbose=**] {**enabled** | **disabled**}]

**Parameters**
**[ type=] { file| email}**
Required. Specifies whether you want to save the report to a file or sent it to an e-mail address.

**[ destination=] { *FileLocation*| *E-mailAddress*}**
Required. Specifies the full path and file name to which the report should be saved or the full e-mail address to which the report should be sent.

**[[ compression=] { enabled| disabled}]**
Specifies whether to compress the report into a .cab file. If you do not specify this parameter, the report is compressed if you send it to an e-mail address but not if you save it to a file.

**[[ hours=] *NumberOfHours*]**
Specifies the number of past hours for which to show activity in the report. This parameter must be an integer between 1 and 24. If you do not specify this parameter, all past information is included.

**[[ verbose=] { enabled| disabled}]**
Specifies the amount of data to include in the report. If you do not specify this parameter, minimal data is included.

**Remarks**

- The Remote Access Diagnostic Report is saved as an HTML file (.htm).
- The Remote Access Diagnostic Report that you create by using the **show all** command contains all diagnostic information available. You can use other commands, such as **show configuration**, to create reports that contain a narrower scope of information.

**show cmtracing**
Shows whether information about Connection Manager connections is being logged.

**Syntax**
**show cmtracing**

**Parameters**
none

**set cmtracing**
Enables or disables logging of information about all Connection Manager connections.

**Syntax**
**set cmtracing** {**enabled** | **disabled**}

**Parameters**
**{ enabled| disabled}**
Required. Specifies whether you want information about Connection Manager connections to be logged. The **enabled** parameter specifies that you want information to be logged. The **disabled** parameter specifies that you do not want information to be logged.

**Remarks**

- Each log file for a Connection Manager connection contains information that is specific to the history of that connection. Each Connection Manager connection generates a unique log file. Depending on the configuration of the connection, the log file might be saved in any of various locations. For more information, see Connection Manager Administration Kit and Troubleshooting process.
- This command will be be overridden for a particular connection if the user enables or disables logging for that connection or if the user installs and uses a Connection Manager profile for which an administrator has enabled or disabled logging.

**show modemtracing**
Shows whether modem tracing is enabled or disabled.

**Syntax**
**show modemtracing**

**Parameters**
none

**set modemtracing**
Enables or disables modem tracing for all modems installed for the local computer.

**Syntax**
**set modemtracing** {**enabled** | **disabled**}

**Parameters**
**{ enabled| disabled}**
Required. Specifies whether you want modem activity for each modem to be logged. The **enabled** parameter specifies that you want activity to be logged. The **disabled** parameter specifies that you do not want activity to be logged.

**Remarks**

- This command enables or disables the logging of modem activity for all modems. If you want to enable or disable modem logging for a specific modem, you should configure the modem in Device Manager.

**show rastracing**
Shows whether all activity for remote access components is traced.

**Syntax**
**show rastracing**

**Parameters**
none

**set rastracing**
Enables or disables tracing and logging of all activity for all remote access components or for a specific remote access component.

**Syntax**
**set rastracing** [**component=**] {*Component* | *} {**enabled** | **disabled**}

**Parameters**
**[ component=] { *Component*| *}**
Required. Specifies whether you want to enable or disable tracing and logging for a component that you specify or for all components. The *Component* parameter specifies the component for which you want to enable or disable tracing and logging.

**{ enabled| disabled}**
Required. Specifies whether you want activity to be traced and logged. The **enabled** parameter specifies that you want activity to be traced and logged. The **disabled** parameter specifies that you do not want activity to be traced and logged.

**Remarks**

- If you enable tracing and logging, a log file named *Component*.txt is created in the %windir%\tracing directory for each component being traced.

**show securityeventlog**
Shows whether security events are being logged.

**Syntax**
**show securityeventlog**

**Parameters**
none

**set securityeventlog**
Enables or disables logging of all security events.

**Syntax**
**set securityeventlog** {**enabled** | **disabled**}

**Parameters**
**{ enabled| disabled}**
Required. Specifies whether you want security events to be logged. The **enabled** parameter specifies that you want security events to be logged. The **disabled** parameter specifies that you do not want security events to be logged.

**show tracefacilities**
Shows whether all activity for all remote access components or for a remote access component that you specify is being traced and logged.

**Syntax**
**show tracefacilities**

**Parameters**
none

**set tracefacilities**
Enables or disables tracing and logging of all activity for all remote access components that are configured on the local computer.

**Syntax**
**set tracefacilities** [**state=**] {**enabled** | **disabled** | **clear**}

**Parameters**
**[ state=] { enabled| disabled| clear}**
Required. Specifies whether you want to enable tracing for all remote access components, to disable tracing, or to clear all logs generated by tracefacilities. The **enabled** parameter specifies that you want to enable tracing. The **disabled** parameter specifies that you want to disable tracing. The **clear** parameter specifies that you want to clear all logs.

**Netsh ras IP context commands**
The following commands are specific to the **ras IP** context within the Netsh environment.

To view the command syntax, click a command:

- dump
- show config
- set negotiation
- set access
- set addassign
- set addrreq
- set broadcastnameresolution
- show broadcastnameresolution
- add range
- delete range
- delete pool

**dump**
Creates a script that contains the IP configuration of a remote access server. If you save this script to a file, you can use it to

restore IP configuration settings.

**Syntax**
**dump**

**Parameters**
none

**show config**
Displays the current IP configuration of the remote access server.

**Syntax**
**show config**

**Parameters**
none

**set negotiation**
Specifies whether the remote access server should allow IP to be configured for any client connections the server accepts.

**Syntax**
**set negotiation** {**allow** | **deny**}

**Parameters**
**{ allow| deny}**
Required. Specifies whether to permit IP over client connections. The **allow** parameter allows IP over client connections. The **deny** parameter prevents IP over client connections.

**set access**
Specifies whether IP network traffic from any client should be forwarded to the network or networks to which the remote access server is connected.

**Syntax**
**set access** {**all** | **serveronly**}

**Parameters**
**{ all| serveronly}**
Required. Specifies whether clients should be able to reach the remote access server and any networks to which it is connected. The **all** parameter allows clients to reach networks through the server. The **serveronly** parameter allows clients to reach only the server.

**set addrassign**
Sets the method by which the remote access server should assign IP addresses to its clients.

**Syntax**
**set addrassign** {**auto** | **pool**}

**Parameters**
**{ auto| pool}**
Required. Specifies whether IP addresses should be assigned by using DHCP or from a pool of addresses held by the remote access server. The **auto** parameter specifies that addresses should be assigned by using DHCP. If no DHCP server is available, a random, private address is assigned. The **pool** parameter specifies that addresses should be assigned from a pool.

**set addrreq**
Specifies whether dial-up clients should be able to request their own IP addresses.

**Syntax**
**set addrreq** {**allow** | **deny**}

**Parameters**
**{ allow| deny}**
Required. Specifies whether clients should be able to request their own IP addresses. The **allow** parameter allows clients to request addresses. The **deny** parameter prevents clients from requesting addresses.

**set broadcastnameresolution**
Enables or disables broadcast name resolution using NetBIOS over TCP/IP.

**Syntax**

**set broadcastnameresolution** {**enabled** | **disabled**}

**Parameters**

**{ enabled| disabled}**

Required. Specifies whether to enable or disable broadcast name resolution using NetBIOS over TCP/IP.

**show broadcastnameresolution**

Displays whether broadcast name resolution using NetBIOS over TCP/IP has been enabled or disabled for the remote access server.

**Syntax**

**show broadcastnameresolution**

**Parameters**

none

**add range**

Adds a range of addresses to the pool of static IP addresses that the remote access server can assign to clients.

**Syntax**

**add range** [**from**=]*StartingIPAddress* [**to**=]*EndingIPAddress*

**Parameters**

**[ from=] *StartingIPAddress*[ to=] *EndingIPAddress***

Required. Specifies the range of IP addresses to add. The *StartingIPAddress* parameter specifies the first IP address in the range. The *EndingIPAddress* parameter specifies the last IP address in the range.

**Examples**

To add the range of IP addresses 10.2.2.10 to 10.2.2.20 to the static pool of IP addresses that the remote access server can assign, type:

**add range from=10.2.2.10 to=10.2.2.20**

**delete range**

Deletes a range of addresses from the pool of static IP addresses that the remote access server can assign to clients.

**Syntax**

**delete range** [**from**=]*StartingIPAddress* [**to**=]*EndingIPAddress*

**Parameters**

**[ from=] *StartingIPAddress*[ to=] *EndingIPAddress***

Required. Specifies the range of IP addresses to delete. The *StartingIPAddress* parameter specifies the first IP address in the range. The *EndingIPAddress* parameter specifies the last IP address in the range.

**Examples**

To delete the range of IP addresses 10.2.2.10 to 10.2.2.20 from the pool of static IP addresses that the remote access server can assign, type:

**delete range from=10.2.2.10 to=10.2.2.20**

**delete pool**

Deletes all addresses from the pool of static IP addresses that the remote access server can assign to clients.

**Syntax**

**delete pool**

**Parameters**

none

**Netsh ras AppleTalk context commands**

The following commands are specific to the **ras AppleTalk** context within the Netsh environment.

This feature is not available on the Itanium-based versions of the Windows operating systems. This content is not available in this preliminary release.

To view the command syntax, click a command:

- dump

- show config
- set negotiation

**dump**
Creates a script that contains the AppleTalk configuration of the remote access server. If you save this script to a file, you can use it to restore AppleTalk configuration settings.

**Syntax**
**dump**

**Parameters**
none

**show config**
Displays the current AppleTalk configuration of the remote access server.

**Syntax**
**show config**

**Parameters**
none

**set negotiation**
Specifies whether the remote access server should allow AppleTalk to be configured for any client connections the server accepts.

**Syntax**
**set negotiation** {**allow** | **deny**}

**Parameters**
**{ allow| deny}**
Required. Specifies whether to allow AppleTalk configuration. The **allow** parameter allows configuration. The **deny** parameter prevents configuration.

**Netsh ras AAAA context commands**
The following commands are specific to the **ras AAAA** context within the Netsh environment.

To view the command syntax, click a command:

- dump
- add acctserv
- delete acctserv
- set acctserv
- show authserv
- add authserv
- delete authserv
- set authserv
- show authserv
- set acco
- show acco

- set authe
- show authe

**dump**
Displays the AAAA configuration of a remote access server in script form.

**Syntax**
**dump**

**Parameters**
none

**add acctserv**
Specifies the IP address or the DNS name of a RADIUS server to use for accounting and specifies accounting options.

**Syntax**
**add acctserv** [**name**=]*ServerID*[[**secret**=]*SharedSecret*] [[**init-score**=]*ServerPriority*] [[**port**=]*Port*] [[**timeout**=]*Seconds*] [[**messages**] {**enabled** | **disabled**}]

**Parameters**
**[ name=] *ServerID***
Required. Specifies, by IP address or DNS name, the RADIUS server.

**[ secret=] *SharedSecret***
Specifies the shared secret.

**[ init-score=] *ServerPriority***
Specifies the initial score (server priority).

**[ port=] *Port***
Specifies the port to which accounting requests should be sent.

**[ timeout=] *Seconds***
Specifies the timeout period, in seconds, during which the RADIUS server can be idle before it should be marked unavailable.

**[ messages] { enabled| disabled}**
Specifies whether to send accounting on/off messages. The **enabled** parameter specifies that messages should be sent. The **disabled** parameter specifies that messages should not be sent.

**delete acctserv**
Deletes a RADIUS accounting server.

**Syntax**
**delete acctserv** [**name**=]*ServerID*

**Parameters**
**[ name=] *ServerID***
Required. Specifies, by DNS name or IP address, which server to delete.

**set acctserv**
Provides the IP address or the DNS name of a RADIUS server to use for accounting.

**Syntax**
**add acctserv** [**name=**]*ServerID* [[**secret=**]*SharedSecret*] [[**init-score=**]*ServerPriority*] [[**port=**]*Port*] [[**timeout=**]*Seconds*] [[**messages**] {**enabled** | **disabled**}]

**Parameters**
**[ name=] *ServerID***
Required. Specifies, by IP address or DNS name, the RADIUS server.

**[ secret=] *SharedSecret***
Specifies the shared secret.

**[ init-score=] *ServerPriority***
Specifies the initial score (server priority).

**[ port=] *Port***
Specifies the port on which to send the authentication requests.

**[ timeout=] *Seconds***
Specifies, in seconds, the amount of time that should elapse before the RADIUS server is marked unavailable.

**[ messages=] { enabled| disabled}**
Specifies whether accounting on/off messages should be sent.

#### show acctserv

Displays detailed information about an accounting server. Used without parameters, **show acctserv** displays information about all configured accounting servers.

#### Syntax

**show acctserv** [[**name**=]*ServerID*]

#### Parameters

**[ name=] *ServerID***
Specifies, by DNS name or IP address, the RADIUS server about which to display information.

#### add authserv

Provides the IP address or the DNS name of a RADIUS server to which authentication requests should be passed.

#### Syntax

**add authserv** [**name**=]*ServerID*[[**secret**=]*SharedSecret*] [[**init-score**=]*ServerPriority*] [[**port**=]*Port*] [[**timeout**=]*Seconds*] [[**signature**] {**enabled** | **disabled**}]

#### Parameters

**[ name=] *ServerID***
Required. Specifies, by IP address or DNS name, the RADIUS server.

**[ secret=] *SharedSecret***
Specifies the shared secret.

**[ init-score=] *ServerPriority***
Specifies the initial score (server priority).

**[ port=] *Port***
Specifies the port to which authentication requests should be sent.

**[ timeout=] *Seconds***
Specifies the timeout period, in seconds, during which the RADIUS server can be idle before it should be marked unavailable.

**[ signature] { enabled| disabled}**
Specifies whether to use digital signatures. The **enabled** parameter specifies that digital signatures should be used. The **disabled** parameter specifies that digital signatures should not be used.

#### delete authserv

Deletes a RADIUS authentication server.

#### Syntax

**delete authserv** [**name**=]*ServerID*

#### Parameters

**[ name=] *ServerID***
Required. Specifies, by DNS name or IP address, which server to delete.

#### set authserv

Provides the IP address or the DNS name of a RADIUS server to which authentication requests should be passed.

#### Syntax

**set authserv** [**name=**]*ServerID* [[**secret=**]*SharedSecret*] [[**init-score=**]*ServerPriority*] [[**port=**]*Port*] [[**timeout=**]*Seconds*] [[**signature**] {**enabled** | **disabled**}]

#### Parameters

**[ name=] *ServerID***
Required. Specifies, by IP address or DNS name, the RADIUS server.

**[ secret=] *SharedSecret***

Specifies the shared secret.

**[ init-score=]** ***ServerPriority***
Specifies the initial score (server priority).

**[ port=]** ***Port***
Specifies the port on which to send the authentication requests.

**[ timeout=]** ***Seconds***
Specifies the amount of time, in seconds, that should elapse before the RADIUS server is marked unavailable.

**[ signature=] { enabled| disabled}**
Specifies whether digital signatures should be used.

**show authserv**
Displays detailed information about an authentication server. Used without parameters, **show authserv** displays information about all configured authentication servers.

**Syntax**
**show authserv** [[**name**=]*ServerID*]

**Parameters**
**[ name=]** ***ServerID***
Specifies, by DNS name or IP address, the RADIUS server about which to display information.

**set acco**
Specifies the accounting provider.

**Syntax**
**set acco** {**windows** | **radius** | **none**}

**Parameters**
**{ windows| radius| none}**
Required. Specifies whether accounting should be performed and by which server. The **windows** parameter specifies that Windows security should perform accounting. The **radius** parameter specifies that a RADIUS server should perform accounting. The **none** parameter specifies that no accounting should be performed.

**show acco**
Displays the accounting provider.

**Syntax**
**show acco**

**Parameters**
none

**set authe**
Specifies the authentication provider.

**Syntax**
**set authe** {**windows** | **radius**}

**Parameters**
**{ windows| radius}**
Required. Specifies which technology should perform authentication. The **windows** parameter specifies that Windows security should perform authentication. The **radius** parameter specifies that a RADIUS server should perform authentication.

**show authe**
Displays the authentication provider.

**Syntax**
**show authe**

**Parameters**
none

**Formatting legend**

| Format | Meaning |
|---|---|

| | |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**See Also**

**Concepts**

Netsh commands for AAAA
Netsh commands for DHCP
Netsh diagnostic (diag) commands
Netsh commands for Interface IP
Netsh commands for WINS
Command-line reference A-Z
Command shell overview

# Netsh commands for Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for Routing**
You can use the Netsh commands for Routing to administer routing servers from the command line, rather than through the Routing and Remote Access console. Administering servers from the command line is most useful in the following situations:

- You can administer routing servers more quickly from the command line than you can using a console. The ability is especially important if you are using a slow-speed network link, for example, a connection over a wide area network (WAN).
- You can administer routing servers more consistently by incorporating these commands into a script that you can run on many servers.

You can use the following commands at the command prompt for the appropriate Netsh context to administer routing servers. To run these Netsh commands on a remote Windows 2000 Server, you must first use Remote Desktop Connection to connect to a Windows 2000 Server that is running Terminal server. There might be functional differences between Netsh context commands on Windows 2000 and the Windows Server 2003 family.

For more information on Netsh, see Netsh overview and Enter a netsh context.

**Netsh Routing IP IGMP commands**
The following commands are available for use in the Netsh Routing IP IGMP context.

To view the command syntax, click a command:

- add interface
- install
- delete interface
- uninstall
- set interface
- set global
- show interface
- show global
- show ifstats
- show iftable
- show grouptable
- show rasgrouptable
- show proxygrouptable

**add interface**
Configures IGMP on the specified interface.

**Syntax**

**add interface** [**InterfaceName=**][**InterfaceName=**]*InterfaceName* [[**IgmpPrototype=**]{**igmprtrv1** | **igmprtrv2** | **igmprtrv3** | **igmpproxy**}] [[**IfEnabled=**]{**enable** | **disable**}] [[**RobustVar**=]*Integer*] [[**GenQueryInterval=**]*Integer*] [[**GenQueryRespTime=**]*Integer*] [[**StartUpQueryCount=**]*Integer*] [[**StartUpQueryInterval=**]*Integer*] [[**LastMemQueryCount=**]*Integer*] [[**LastMemQueryInterval=***Integer*] [[**AccNonRtrAlertPkts=**]{**yes** | **no**}]

**Parameters**

**[ InterfaceName=][ InterfaceName=] *InterfaceName***

Required. Specifies the name of the interface on which you want to configure IGMP. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ IgmpPrototype=]{ igmprtrv1| igmprtrv2| igmprtrv3| igmpproxy}**

Specifies the type of protocol to be configured on the interface. The default type is igmprtrv2.

**[ IfEnabled=]{ enable| disable}**

Specifies whether the protocol you are configuring should be enabled or disabled. The default is enabled.

**[ RobustVar=]= *Integer***

Specifies the robustness variable. The value should be greater than zero (0). The default value is 2.

**[ GenQueryInterval=] *Integer***

Specifies the interval, in seconds, at which general queries should be sent on the interface. The default value is 125 seconds.

**[ GenQueryRespTime=] *Integer***

Specifies the maximum response time, in seconds, by which hosts should respond to a general query. The default value is 10 seconds.

**[ StartUpQueryCount=] *Integer***

Specifies the number of general queries that will be sent during startup. The default value is 2.

**[ StartUpQueryInterval=] *Integer***

Specifies the interval, in seconds, between general queries sent during startup. The default value is 31 seconds.

**[ LastMemQueryCount=] *Integer***

Specifies the number of group-specific queries sent when notice of a host leaving the group is received. The default value is 2.

**[ LastMemQueryInterval=] *Integer***

Specifies the interval, in milliseconds, between group-specific queries. The default value is 1000 milliseconds.

**[ AccNonRtrAlertPkts=]{ yes| no}**

Specifies whether to accept IGMP packets that might not have the router alert option set. The default is **yes**.

**Examples**

To modify the default startup query interval to 21 seconds for use with IGMP configuration of the interface named Local Area Connection, type:

**add interface "Local Area Connection" startupqueryinterval=21**

**install**

Installs an IGMP router/proxy and sets global logging. Used without parameters, **install** installs the IGMP router/proxy and specifies that errors related to IGMP are logged but no other IGMP-related events are logged.

**Syntax**

**install** [[**LogLevel**=]{**none** | **error** | **warn** | **info**}]

**Parameters**

**[ LogLevel=]{ none| error| warn| info}**

Specifies which events should be logged. The **none** parameter specifies no events related to IGMP should be logged. The **error** parameter specifies that only errors related to IGMP should be logged. The **warn** parameter specifies that only warnings related to IGMP should be logged. The **info** parameter specifies that all events related to IGMP should be logged.

**delete interface**

Removes an IGMP router or proxy from the specified interface.

**Syntax**

**delete interface** [**InterfaceName=**][**InterfaceName=**]*InterfaceName*

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, an interface from which you want to remove an IGMP router or proxy. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**Examples**

**delete interface "Local Area Connection"**

**uninstall**

Removes IGMP router/proxy from all interfaces.

**Syntax**

**uninstall**

**Parameters**

none

**set interface**

Changes interface configuration parameters.

**Syntax**

**set interface** [**InterfaceName=**]*InterfaceName* [[**IgmpPrototype=**]{**igmprtrv1** | **igmprtrv2** | **igmprtrv3** | **igmpproxy**}] [[**IfEnabled=**]{**enable** | **disable**}] [[**RobustVar=**]*Integer*] [[**GenQueryInterval=**]*Integer*] [[**GenQueryRespTime=**]*Integer*] [[**StartUpQueryCount=**]*Integer*] [[**StartUpQueryInterval=**]*Integer*] [[**LastMemQueryCount=**]*Integer*] [[**LastMemQueryInterval=**]*Integer*] [[**AccNonRtrAlertPkts=**]{**yes** | **no**}]

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies the name of the interface on which you want to configure IGMP. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ IgmpPrototype=]{ igmprtrv1| igmprtrv2| igmprtrv3| igmpproxy}**

Specifies the type of protocol to be configured on the interface. The default type is igmprtrv2.

**[ IfEnabled=]{ enable| disable}**

Specifies whether the protocol you are configuring should be enabled or disabled. The default is enabled.

**[ RobustVar=] *Integer***

Specifies the robustness variable. The value should be greater than zero (0). The default value is 2.

**[ GenQueryInterval=] *Integer***

Specifies the interval, in seconds, at which general queries should be sent on the interface. The default value is 125 seconds.

**[ GenQueryRespTime=] *Integer***

Specifies the maximum response time, in seconds, by which hosts should respond to a general query. The default value is 10 seconds.

**[ StartUpQueryCount=] *Integer***

Specifies the number of general queries that will be sent during startup. The default value is 2.

**[ StartUpQueryInterval=] *Integer***

Specifies the interval, in seconds, between general queries sent during startup. The default value is 31 seconds.

**[ LastMemQueryCount=] *Integer***

Specifies the number of group-specific queries sent when notice of a host leaving the group is received. The default value is 2.

**[ LastMemQueryInterval=] *Integer***

Specifies the interval, in milliseconds, between group-specific queries. The default value is 1,000 milliseconds.

**[ AccNonRtrAlertPkts=]{ yes| no}**

Specifies whether to accept IGMP packets that do not have the router alert option set. The default is to accept those packets.

**Examples**

To modify the default startup query interval to 21 seconds for use with IGMP configuration of an interface named Local Area Connection, type:

**set interface "Local Area Connection" startupqueryinterval=21**

**set global**
Sets IGMP global parameters. Used without parameters, **set global** specifies that warnings are logged but that no other events related to IGMP are logged.

**Syntax**
**set global** [**LogLevel**=]{**none** | **error** | **warn** | **info**}

**Parameters**
**[ LogLevel=]{ none| error| warn| info}**
Required. Specifies which events should be logged. The **none** parameter specifies that no events related to IGMP should be logged. The **error** parameter specifies that only errors related to IGMP should be logged. The **warn** parameter specifies that only warnings related to IGMP should be logged. The **info** parameter specifies that all events related to IGMP should be logged.

**show interface**
Displays the IGMP configuration of the specified interface. Used without parameters, **show interface** displays IGMP configuration information for all interfaces.

**Syntax**
**show interface** [**InterfaceName=**]*InterfaceName*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface for you want to display information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**show global**
Displays the global IGMP parameters.

**Syntax**
**show global**

**Parameters**
none

**show ifstats**
Displays IGMP statistics for the specified interface. Used without parameters, **show ifstats** shows IGMP statistics once for all interfaces.

**Syntax**
**show ifstats** [[**Index**=]*Integer*] [[**Rr**=]*Integer*]

**Parameters**
**[ Index=] *Integer***
Specifies a number identifying the interface for which you want to display information.

**[ Rr=] *Integer***
Specifies, in seconds, how often you want the information to be refreshed.

**Examples**
To show IGMP statistics every 5 seconds for an interface with an index value of 1001, type:

**show ifstats index=1001 rr=5**

**show iftable**
Shows IGMP host groups for a specified interface. Used without parameters, **show iftable** shows IGMP host groups for all interfaces once.

**Syntax**
**show iftable** [[**Index**=]*Integer*] [[**Rr**=]*Integer*]

**Parameters**
**[ Index=] *Integer***
Specifies, by index number, the interface for which you want to display information.

**[ Rr=] *Integer***

Specifies, in seconds, how often you want the information to be refreshed.

**Examples**
To show IGMP host groups every 5 seconds for an interface with an index value of 1001, type:

**show iftable index=1001 rr=5**

**show grouptable**
Shows the IGMP hosts group table for a multicast group. Used without parameters, **show grouptable** shows the IGMP hosts group tables once for all multicast groups on all interfaces associated with the router.

**Syntax**
**show grouptable** [[**Index**=]*Integer*] [[**Rr**=]*Integer*]

**Parameters**
**[ Index=] *Integer***
Specifies, by IP address, the multicast group for which you want to display information.

**[ Rr=] *Integer***
Specifies, in seconds, how often you want the information to be refreshed.

**Examples**
To display the IGMP hosts group table every 5 seconds for an IGMP multicast group identified by the IP address of 192.168.100.239, type:

**show grouptable index=192.168.100.239 rr=5**

**show rasgrouptable**
Displays the hosts group table for a remote access client interface. Used without parameters, **show rasgrouptable** displays the hosts group tables once for all remote access client interfaces.

**Syntax**
**show rasgrouptable** [[**Index**=]*IPAddress*] [[**Rr**=]*Integer*]

**Parameters**
**[ Index=] *IPAddress***
Specifies, by IP address, the remote access client interface for which you want to display information.

**[ Rr=] *Integer***
Specifies, in seconds, how often you want the information to be refreshed.

**Examples**
To display the hosts group table for a remote access client interface with an IP address of 10.10.1.100 and to refresh the information every five seconds, type:

**show rasgrouptable index=10.10.1.100 rr=5**

**show proxygrouptable**
Displays the IGMP hosts group table for an IGMP proxy interface. Used without parameters, **show proxygrouptable** displays information for all interfaces once.

**Syntax**
**show proxygrouptable** [[**InterfaceName=**]*InterfaceName*] [[**Rr**=]*Integer*]

**Parameters**
**[ InterfaceName=] *InterfaceName***
Specifies, by name, the interface for which you want to display information.

**[ Rr=] *Integer***
Specifies, in seconds, how often you want the information to be refreshed.

**Examples**
To display IGMP hosts group tables for all IGMP proxy interfaces and to refresh the information every five seconds, type:

**show proxygrouptable rr=5**

**Netsh Routing IP Routerdiscovery commands**
The following commands are available for use in the Netsh Routing IP Routerdiscovery context.

To view the command syntax, click a command:

- add interface
- set interface
- delete interface
- show interface

**add interface**
Configures router discovery for the specified interface.

**Syntax**
**add interface** [**InterfaceName=**]*InterfaceName* [[**Disc=**]{**enable** | **disable**}] [[**MinInt=**]*Integer*] [[**MaxInt=**]*Integer*] [[**Life=**]*Integer*] [[**Level=**]*Integer*]

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface for which you want to configure router discovery. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ Disc=]{ enable| disable}**
Specifies whether router discovery should occur on this interface. The **enable** parameter specifies that router discovery should occur, and the **disable** parameter specifies that router discovery should not occur.

**[ MinInt=] *Integer***
Specifies a minimum interval, in minutes, between when router discovery packets are sent.

**[ MaxInt=] *Integer***
Specifies a maximum interval, in minutes, between when router discovery packets are sent.

**[ Life=] *Integer***
Specifies the number of minutes for which a router discovery packet sent from this router or server is valid.

**[ Level=] *Integer***
Specifies a number indicating preference for using this interface as the default gateway. Where multiple interfaces are enabled for router discovery, a higher number used here indicates a higher level of preference.

**Remarks**

- When added, router discovery advertisements are sent at random times bounded by the set minimum and maximum intervals.

**Examples**
To enable router discovery for the Local Area Connection interface, specifying a minimum interval of 7 minutes, a maximum interval of 10 minutes, a life of 30 minutes for router discovery packets, and a zero (0) or no preference level, type:

**add interface "Local Area Connection" enable 7 10 30 0**

**set interface**
Configures router-discovery for an interface.

**Syntax**
**set interface** [**InterfaceName=**]*InterfaceName* [[**Disc=**]{**enable** | **disable**}] [[**MinInt=**]*Integer*] [[**MaxInt=**]*Integer*] [[**Life=**]*Integer*] [[**Level=**]*Integer*]

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface for which you want to configure router discovery. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ Disc=]{ enable| disable}**
Specifies whether router discovery should occur on this interface. The **enable** parameter specifies that router discovery should occur, and the **disable** parameter specifies that router discovery should not occur.

**[ MinInt=]** ***Integer***
Specifies, in minutes, a minimum interval between when router discovery packets are sent.

**[ MaxInt=]** ***Integer***
Specifies, in minutes, a maximum interval between when router discovery packets are sent.

**[ Life=]** ***Integer***
Specifies, in minutes, how long a router discovery packet sent from this router or server is valid.

**[ Level=]** ***Integer***
Specifies a number indicating preference for using this interface as the default gateway. A higher number used here specifies a higher level of preference.

**Examples**
To enable router discovery for the Local Area Connection interface, specifying a minimum interval of 7 minutes, a maximum interval of 10 minutes, a life of 30 minutes for router discovery packets, and a zero (0) or no preference level, type:

**set interface "Local Area Connection" enable 7 10 30 0**

**delete interface**
Deletes router-discovery configuration.

**Syntax**
**delete interface** [**InterfaceName=**]*InterfaceName*

**Parameters**
**[ InterfaceName=]** ***InterfaceName***
Required. Specifies, by name, the interface for which you want to delete router-discovery configuration. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**Examples**
To delete the router-discovery configuration for an interface named Local Area Connection, type:

**delete interface name="Local Area Connection"**

**show interface**
Displays router-discovery information. Used without parameters, **show interface** displays information for all interfaces.

**Syntax**
**show interface** [[**InterfaceName=**]*InterfaceName*]

**Parameters**
**[ InterfaceName=]** ***InterfaceName***
Specifies, by name, the interface for which you want to display information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**Examples**
To display router-discovery information for an interface named Local Area Connection, type:

**show interface "Local Area Connection"**

**Netsh Routing IP RIP commands**
The following commands are available for use in the Netsh Routing IP RIP context.

To view the command syntax, click a command:

- add peerfilter
- add acceptfilter
- add announcefilter

- add neighbor
- add interface
- add helper
- delete peerfilter
- delete acceptfilter
- delete announcefilter
- delete neighbor
- delete interface
- delete helper
- set interface
- set flags
- set global
- show interface
- show flags
- show global
- show ifstats
- show ifbinding

**add peerfilter**
Adds a filter for servers that can be accepted as peers.

**Syntax**
**add peerfilter** [**Server**=]*IPAddress*

**Parameters**
**[ Server=] *IPAddress***
Required. Specifies, by IP address, the peer server whose traffic you want to filter.

**add acceptfilter**
Adds an acceptance filter for routes received on an interface.

**Syntax**
**add acceptfilter** [**InterfaceName=**]*InterfaceName* [**Addr=**]*IPAddress1IPAddress2*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, an interface where RIP is available for use. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ Addr=] *IPAddress1IPAddress2***
Required. Specifies an IP address pairing. The first address is for the low IP address, and the second address is for the high IP

address to use in building the filter.

**Examples**

To add an acceptance filter for an interface named Local Area Connection for an IP range of 10.0.0.2 to 10.11.21.154, type:

**add acceptfilter "Local Area Connection" 10.0.0.2 10.11.21.154**

**add announcefilter**

Adds a filter for routes announced on an interface.

**Syntax**

**add announcefilter** [**InterfaceName=**]*InterfaceName* [**Addr=**]*IPAddress1IPAddress2*

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies the name of an interface where RIP is available for use. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ Addr=] *IPAddress1IPAddress2***

Required. Specifies an IP address pairing. The first address is for the low IP address, and the second address is for the high IP address to use in building the filter.

**Examples**

To add a filter for routes announced on an interface named Local Area Connection for a range of IP addresses from 10.0.0.2 to 10.11.21.154, type:

**add announcefilter Local Area Connection 10.0.0.2 10.11.21.154**

**add neighbor**

Adds a RIP neighbor on an interface.

**Syntax**

**add neighbor** [**InterfaceName=**]*InterfaceName* [**addr=**]*IPAddress*

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, an interface where RIP is available for use. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ addr=] *IPAddress***

Required. Specifies the IP address of the RIP neighbor you want to add.

**Examples**

To add a RIP neighbor located at an IP address of 10.0.0.2 for an interface named Local Area Connection, type:

**add neighbor "Local Area Connection" 10.0.0.2**

**add interface**

Adds RIP to a specified interface.

**Syntax**

**add interface** [**InterfaceName=**]*InterfaceName* [[**Metric=**]*Integer*] [[**UpdateMode=**]{**demand** | **periodic**}] [[**Announce=**]{**none** | **rip1** | **rip1compat** | **rip2**}] [[**Accept=**]{**none** | **rip1** | **rip1compat** | **rip2**}] [[**Expire=**]*Integer*] [[**Remove=**]*Integer*] [[**Update=**]*Integer*] [[**AuthMode=**]{**authnone** | **authsimplepassword**}] [[**Tag=**]*Integer*] [[**Unicast=**]{**also** | **only** | **disable**}] [[**AccFiltMode=**]{**include** | **exclude** | **disable**}] [[**AnnFiltMode=**]{**include** | **exclude** | **disable**}] [[**Password=**]*String*]

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface on which to add RIP. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ Metric=] *Integer***

Specifies a metric value for routes based on this interface.

**[ UpdateMode=]{ demand| periodic}**

Specifies when updates will occur. The **demand** parameter specifies that updates should occur on demand. The **periodic** parameter specifies that updates should occur at regular intervals.

**[ Announce=]{ none| rip1| rip1compat| rip2}**
Specifies which routes should be announced. The **none** parameter specifies that all routes should be announced. The **rip1** parameter specifies that only RIP version 1 routes should be announced. The **rip1compat** parameter specifies that announcements should be made in RIP version-compatible mode. The **rip2** parameter specifies that only RIP version 2 routes should be announced.

**[ Accept=]{ none| rip1| rip1compat| rip2}**
Specifies which routes should be accepted. The **none** parameter specifies that all routes should be accepted. The **rip1** parameter specifies that only RIP version 1 routes should be accepted. The **rip1compat** parameter specifies that routes should be accepted in RIP version-compatible mode. The **rip2** parameter specifies that only RIP version 2 routes should be accepted.

**[ Expire=]** ***Integer***
Specifies, in seconds, the route expiration interval.

**[ Remove=]** ***Integer***
Specifies, in seconds, the route removal interval.

**[ Update=]** ***Integer***
Specifies, in seconds, the full update interval.

**[ AuthMode=]{ authnone| authsimplepassword}**
Specifies whether authorization should be required. The **authnone** parameter specifies that no authorization should be required. The **authsimplepassword** parameter specifies that password authorization should be required.

**[ Tag=]** ***Integer***
Specifies the route tag.

**[ Unicast=]{ also| only| disable}**
Specifies the unicast mode. The **also** parameter specifies use also unicast mode. The **only** parameter specifies use only unicast mode. The **disable** parameter specifies that unicast mode should be disabled.

**[ AccFiltMode=]{ include| exclude| disable}**
Specifies whether acceptance filters should be included or disabled. The **include** parameter specifies that acceptance filters should be included. The **exclude** parameter specifies that acceptance filters should not be included. The **disable** parameter specifies that acceptance filters should be disabled.

**[ AnnFiltMode=]{ include| exclude| disable}**
Specifies whether announcement filters should be included or disabled. The **include** parameter specifies that announcement filters should be included. The **exclude** parameter specifies that announcement filters should not be included. The **disable** parameter specifies that announcement filters should be disabled.

**[ Password=]** ***String***
Specifies a plaintext password no longer than 16 characters.

#### add helper
Installs a helper DLL under RIP.

#### Syntax
**add helper** *DLLName*

#### Parameters
***DLLName***
Required. Specifies a DLL that supports this context.

#### delete peerfilter
Deletes a filter for an accepted peer server.

#### Syntax
**delete peerfilter** [**Server**=]*IPAddress*

#### Parameters
**[ Server=]** ***IPAddress***
Required. Specifies, by IP address, the peer server whose traffic you do not want to filter.

**delete acceptfilter**
Deletes an acceptance filter for routes received on a specified interface.

**Syntax**
**delete acceptfilter** [**InterfaceName=**]*InterfaceName* [**Addr=**]*IPAddress1IPAddress2*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, an interface where RIP is available for use. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ Addr=] *IPAddress1IPAddress2***
Required. Specifies an IP address pairing. The first address is for the low IP address, and the second address is for the high IP address used in building the filter.

**Examples**
To delete an acceptance filter for the IP ranges of 10.0.0.2 to 10.11.21.154 on an interface named Local Area Connection, type:

**delete acceptfilter "Local Area Connection" 10.0.0.2 10.11.21.154**

**delete announcefilter**
Deletes an announcement filter set for an interface.

**Syntax**
**delete announcefilter** [**InterfaceName=**]*InterfaceName* [**Addr=**]*IPAddress1IPAddress2*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies the name of an interface where RIP is available for use. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ Addr=] *IPAddress1IPAddress2***
Required. Specifies an IP address pairing. The first address specifies the low IP address, and the second address specifies the high IP address used in building the filter.

**Examples**
To delete an announcement filter for the IP range 10.0.0.2 to 10.11.21.154 for an interface named Local Area Connection, type:

**delete announcefilter "Local Area Connection" 10.0.0.2 10.11.21.154**

**delete neighbor**
Deletes a RIP neighbor from an interface.

**Syntax**
**delete neighbor** [**InterfaceName=**]*InterfaceNameIPAddress*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, an interface where RIP is available for use. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

***IPAddress***
Required. Specifies, by IP address, the RIP neighbor to delete.

**Examples**
To delete a RIP neighbor located at 10.0.0.2 for an interface named Local Area Connection, type:

**delete neighbor "Local Area Connection" 10.0.0.2**

**delete interface**
Removes RIP from the specified interface.

**Syntax**
**delete interface** [**InterfaceName=**]*InterfaceName*

**Parameters**

**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, an interface from which you want to remove RIP. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**delete helper**
Removes a specified helper DLL from Netsh.

**Syntax**
**delete helper** *HelperName*

**Parameters**
***HelperName***
Required. Specifies, by file name, the helper DLL to remove.

**set interface**
Configures RIP on a specified interface.

**Syntax**
**set interface** [**InterfaceName=**]*InterfaceName* [[**Metric=**]*Integer*] [[**UpdateMode=**]{**demand** | **periodic**}] [[**Announce=**]{**none** | **rip1** | **rip1compat** | **rip2**}] [[**Accept=**]{**none** | **rip1** | **rip1compat** | **rip2**}] [[**Expire=**]*Integer*] [[**Remove=**]*Integer*] [[**Update=**]*Integer*] [[**AuthMode=**]{**authnone** | **authsimplepassword**}] [[**Tag=**]*Integer*] [[**Unicast=**]{**also** | **only** | **disable**}] [[**AccFiltMode=**]{**include** | **exclude** | **disable**}] [[**AnnFiltMode=**]{**include** | **exclude** | **disable**}] [[**Password=**]*String*]

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface on which to configure RIP. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ Metric=] *Integer***
Specifies a metric value for routes based on this interface.

**[ UpdateMode=]{ demand| periodic}**
Specifies when updates should occur. The **demand** parameter specifies that updates should occur on demand. The **periodic** parameter specifies that updates should occur at regular intervals.

**[ Announce=]{ none| rip1| rip1compat| rip2}**
Specifies which routes should be announced. The **none** parameter specifies that all routes should be announced. The **rip1** parameter specifies that only RIP version 1 routes should be announced. The **rip1compat** parameter specifies that announcements should be made in RIP version-compatible mode. The **rip2** parameter specifies that only RIP version 2 routes should be announced.

**[ Accept=]{ none| rip1| rip1compat| rip2}**
Specifies which routes should be accepted. The **none** parameter specifies that all routes should be accepted. The **rip1** parameter specifies that only RIP version 1 routes should be accepted. The **rip1compat** parameter specifies that acceptance should occur using RIP version-compatible mode. The **rip2** parameter specifies that only RIP version 2 routes should be accepted.

**[ Expire=] *Integer***
Specifies, in seconds, how much time should elapse before the route expires.

**[ Remove=] *Integer***
Specifies, in seconds, how much time should elapse before the route is removed.

**[ Update=] *Integer***
Specifies, in seconds, how often full updates occur.

**[ AuthMode=]{ authnone| authsimplepassword}**
Specifies whether authorization should be required. The **authnone** parameter specifies that no authorization should be required. The **authsimplepassword** parameter specifies that password authorization should be required.

**[ Tag=] *Integer***
Specifies the route tag.

**[ Unicast=]{ also| only| disable}**
Specifies the unicast mode. The **also** parameter specifies use also unicast mode. The **only** parameter specifies use only

unicast mode. The **disable** parameter specifies that unicast mode should be disabled.

**[ AccFiltMode=]{ include| exclude| disable}**
Specifies whether acceptance filters should be included. The **include** parameter specifies that acceptance filters should be included. The **exclude** parameter specifies that acceptance filters should not be included. The **disable** parameter specifies that acceptance filters should be disabled.

**[ AnnFiltMode=]{ include| exclude| disable}**
Specifies whether announcement filters should be included or disabled. The **include** parameter specifies that announcement filters should be included. The **exclude** parameter specifies that announcement filters should not be included. The **disable** parameter specifies that announcement filters should be disabled.

**[ Password=]** ***String***
Specifies a plaintext password no longer than 16 characters long.

**Examples**
To modify the RIP configuration on the interface named Local Area Connection, specifying a metric value of 2 and that updates should occur on demand, type:

**set interface "Local Area Connection" metric=2 updatemode=periodic**

**set flags**
Sets RIP-related flags for a specified interface.

**Syntax**
**set flags** [**InterfaceName=**]*InterfaceName* [**clear,**][**splithorizon,**][**poisonreverse,**][**triggeredupdates,**][**cleanupupdates**] [**accepthostroutes,**][**acceptdefaultroutes,**][**senddefaultroutes,**][**nosubnetsummary**]

**Parameters**
**[ InterfaceName=]** ***InterfaceName***
Required. Specifies, by name, the interface on which to set flags. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ clear,][ splithorizon,][ poisonreverse,][ triggeredupdates,][ cleanupupdates,][ accepthostroutes,][ acceptdefaultroutes,][ senddefaultroutes,][ nosubnetsummary]**
Specifies policies for RIP-related flags. The **clear** parameter specifies that all set RIP flags should be cleared. The **splithorizon** parameter specifies that RIP for split horizon should be enabled. The **poisonreverse** parameter specifies that RIP for poison reverse should be enabled. The **triggeredupdates** parameter specifies that RIP triggered updates should occur. The **cleanupupdates** parameter specifies that cleanup updates for RIP should occur. The **accepthostroutes** parameter specifies that host routes should be accepted. The **acceptdefaultroutes** parameter specifies that default routes should be accepted. The **senddefaultroutes** parameter specifies that default routes should be sent. The **nosubnetsummary** parameter specifies that subnet summarization should not occur.

**Remarks**

- The last policy parameter should not have a comma after it.

**set global**
Sets global RIP parameters.

**Syntax**
**set global** {[[**loglevel**=]{**none** | **error** | **warn** | **info**}] | [[**mintrig**=]*Integer*] | [[**peermode**=]{**include** | **exclude** | **disable**}]}

**Parameters**
**[ loglevel=]{ none| error| warn| info}**
Specifies which RIP events should be logged. The **none** parameter specifies that no RIP events should be logged. The **error** parameter specifies that only errors related to RIP should be logged. The **warn** parameter specifies that only warnings related to RIP should be logged. The **info** parameter specifies that all events related to RIP should be logged.

**[ mintrig=]** ***Integer***
Specifies, in minutes, the minimum amount of time that should elapse between triggers.

**[ peermode=]{ include| exclude| disable}**
Specifies whether peers should be included and whether peer mode should be disabled. The **include** parameter specifies that peers should be included.The **exclude** parameter specifies that peers should not be included.The **disable** parameter

specifies that peer mode should be disabled.

**show interface**
Shows RIP configuration for the specified interface. Used without parameters, **show interface** displays configuration information for all interfaces.

**Syntax**
**show interface** [[**InterfaceName=**]*InterfaceName*]

**Parameters**
**[ InterfaceName=] *InterfaceName***
Specifies, by name, the interface whose information you want to display. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**show flags**
Displays RIP flags set for a specified interface. Used without parameters, **show flags** displays information for all interfaces.

**Syntax**
**show flags** [[**InterfaceName=**]*InterfaceName*]

**Parameters**
**[ InterfaceName=] *InterfaceName***
Specifies, by name, the interface for which you want to display information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**show global**
Shows RIP global parameters.

**Syntax**
**show global**

**Parameters**
none

**show ifstats**
Displays RIP statistics for a specified interface. Used without parameters, **show ifstats** displays information once for all interfaces.

**Syntax**
**show ifstats** [[**Index**=]*Integer*] [[**Rr**=]*Integer*]

**Parameters**
**[ Index=] *Integer***
Specifies, by index number, the interface for which you want to display information.

**[ Rr=] *Integer***
Specifies, in seconds, how often you want the information to be refreshed.

**Remarks**

- To list the index value for each interface, use the **show interface** command in the Netsh Interface IP context.

**Examples**
To display RIP statistics every 5 seconds for an interface with an index value of 1, type:

**show ifstats index=1 rr=5**

**show ifbinding**
Displays RIP interface IP address binding. Used without parameters, **show ifbinding** displays information once for all interfaces.

**Syntax**
**show ifbinding** [[**Index**=]*Integer*] [[**Rr**=]*Integer*]

**Parameters**
**[ Index=] *Integer***

Specifies, by index number, the interface for which you want to display information.

**[ Rr=] *Integer***
Specifies, in seconds, how often you want the information to be refreshed.

**Remarks**

- To list the index value for each interface, use the **show interface** command in the Netsh Interface IP context.

**Examples**
To display the bindings every 5 seconds for an interface with an index value of 1, type:

**show ifbinding index=1 rr=5**

**Netsh Routing IP DHCP Relay Agent commands**
The following commands are available for use in the Netsh Routing IP DHCP Relay Agent context.

To view the command syntax, click a command:

- add dhcpserver
- add interface
- delete interface
- delete dhcpserver
- set global
- set interface
- show interface
- show global
- show ifconfig
- show ifbinding
- show ifstats

**add dhcpserver**
Adds a DHCP server to the list of DHCP servers whose messages will be forwarded by DHCP Relay Agent.

**Syntax**
**add dhcpserver** *IPAddress*

**Parameters**
***IPAddress***
Required. Specifies, by IP address, the DHCP server you want to add.

**add interface**
Enables DHCP Relay Agent on the specified interface.

**Syntax**
**add interface** [**InterfaceName=**]*InterfaceName*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface on which you want to enable DHCP Relay Agent.

**delete interface**

Disables DHCP Relay Agent on the specified interface.

**Syntax**

**delete interface** [**InterfaceName=**]*InterfaceName*

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface on which you want to disable DHCP Relay Agent. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**delete dhcpserver**

Removes a DHCP server from the list of DHCP servers whose messages will be forwarded by DHCP Relay Agent.

**Syntax**

**delete dhcpserver** *IPAddress*

**Parameters**

***IPAddress***

Required. Specifies, by IP address, the DHCP server to be deleted.

**set global**

Sets global parameters for DHCP Relay Agent configuration.

**Syntax**

**set global** [**LogLevel=**]{**none** | **error** | **warn** | **info**}

**Parameters**

**[ LogLevel=]{ none| error| warn| info}**

Required. Specifies which DHCP Relay Agent events should be logged. The **none** parameter specifies that no events related to DHCP Relay Agent should be logged. The **error** parameter specifies that only errors related to DHCP Relay Agent should be logged. The **warn** parameter specifies that only warnings related to DHCP Relay Agent should be logged. The **info** parameter specifies that all events related to DHCP Relay Agent should be logged.

**set interface**

Updates the DHCP Relay Agent configuration on an interface.

**Syntax**

**set interface** [**InterfaceName=**]*InterfaceName* [**relaymode=**]{**enable** | **disable**} [[**maxhop=**]*Integer*] [[**minsecs=**]*Integer*]

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface on which you want to update the configuration of DHCP Relay Agent. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ relaymode=]{ enable| disable}**

Required. Specifies whether the DHCP Relay Agent is enabled or disabled. The **enable** parameter specifies that DHCP Relay Agent should be enabled for this interface. The **disable** parameter disables DHCP Relay Agent for this interface.

**[ maxhop=] *Integer***

Specifies the number of hops a DHCP packet can make before it should be dropped and should no longer be relayed.

**[ minsecs=] *Integer***

Specifies the minimum number of seconds from boot time that must appear in a DHCP packet before it should be sent to a DHCP server in the list of servers whose messages will be forwarded by DHCP Relay Agent.

**Examples**

To update the DHCP Relay Agent configuration on the interface named Local Area Connection with a maximum hop count of 16 and a value of 4 seconds for the minimum seconds from boot time, type:

**set interface Local Area Connection enable 16 4**

**show interface**

Displays configuration information for DHCP Relay Agent on the specified interface.

**Syntax**

**show interface** [**InterfaceName=**]*InterfaceName*

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface for which you want to display information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**show global**

Displays DHCP Relay Agent global configuration.

**Syntax**

**show global** [[**Rr=**]*Integer*]

**Parameters**

**[ Rr=] *Integer***

Specifies, in seconds, how often you want the information to be refreshed.

**show ifconfig**

Displays DHCP Relay Agent configuration information for the specified interface. Used without parameters, **show ifconfig** displays information once for all interfaces where DHCP Relay Agent is enabled.

**Syntax**

**show ifconfig** [[**Index**=]*Integer*] [[**Rr=**]*Integer*]

**Parameters**

**[ Index=] *Integer***

Specifies, by index value, the interface for which you want to display information.

**[ Rr=] *Integer***

Specifies, in seconds, how often you want the information to be refreshed.

**show ifbinding**

Displays IP address bindings for the specified interface. Used without parameters, **show ifbinding** displays information once for all interfaces where DHCP relay is enabled for use.

**Syntax**

**show ifbinding** [[**Index**=]*Integer*] [[**Rr=**]*Integer*]

**Parameters**

**[ Index=] *Integer***

Specifies, by index value, the interface for which you want to display information.

**[ Rr=] *Integer***

Specifies, in seconds, how often you want the information to be refreshed.

**show ifstats**

Displays DHCP Relay Agent statistics for the specified interface. Used without parameters, **show ifstats** displays information once for all interfaces where DHCP relay is enabled for use.

**Syntax**

**show ifstats** [[**Index**=]*Integer*] [[**Rr=**]*Integer*]

**Parameters**

**[ Index=] *Integer***

Specifies, by index value, the interface for which you want to display information.

**[ Rr=] *Integer***

Specifies, in seconds, how often you want the information to be refreshed.

**Netsh Routing IP OSPF commands**

The following commands are available for use in the Netsh Routing IP OSPF context.

## Notes

- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

To view the command syntax, click a command:

- add range
- add area
- add virtif
- add neighbor
- add interface
- add routefilter
- add protofilter
- install
- delete range
- delete area
- delete virtif
- delete neighbor
- delete interface
- delete routefilter
- delete protofilter
- uninstall
- set area
- set virtif
- set interface
- set global
- set routefilter
- set protofilter
- show global
- show area
- show virtif
- show interface
- show routefilter

- show protofilter
- show area
- show virtif

**add range**
Adds the IP address and mask pair that define a range of addresses belonging to this area.

**Syntax**
**add range** [**Areaid=**]*IPAddress* [**Range=**]*IPAddressSubnetMask*

**Parameters**
**[ Areaid=] *IPAddress***
Required. Specifies the IP address of the area to which the range is added.

**[ Range=] *IPAddressSubnetMask***
Required. Specifies the IP address and the mask of the range to add.

**Remarks**

- OSPF ranges are used to summarize the routes within the OSPF area.

**add area**
Adds an ID for the area.

**Syntax**
**add area** [**Areaid=**]*IPAddress*

**Parameters**
**[ Areaid=] *IPAddress***
Required. Specifies the IP address of the area to which you want to add the range.

**Remarks**

- The **Areaid** parameter is a 32-bit number expressed in dotted decimal notation that identifies the OSPF area. The area ID does not need to coincide with an IP address or an IP network ID. The area ID of 0.0.0.0 is reserved for the backbone. If the area represents a subnetted network, you can use the IP network number of the subnetted network for the area ID.

**add virtif**
Adds the virtual interface for a router.

**Syntax**
**add virtif** [**transareaid=**]*TransitAreaID* [**virtnbrid=**]*VirtualNeighborID*

**Parameters**
**[ transareaid=] *TransitAreaID***
Required. Specifies the IP address of the transit area used to exchange routing information between the backbone area border router and the area border router of the noncontiguous area.

**[ virtnbrid=] *VirtualNeighborID***
Required. Specifies, by IP address, the router ID of the virtual neighbor to add.

**add neighbor**
Adds an OSPF neighbor to the specified interface.

**Syntax**
**add neighbor** [**InterfaceName=**]*InterfaceName* [**addr**=]*IPAddress* [**nbraddr**=]*IPAddress* [**nbrprio=**]*Priority*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface to which you want to add a neighbor. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ addr=] *IPAddress***
Required. Specifies the OSPF interface, by IP address.

**[ nbraddr=] *IPAddress***
Required. Specifies the IP address of the neighbor you want to add.

**[ nbrprio=] *Priority***
Required. Specifies the OSPF router priority of the neighbor you want to add.

**add interface**
Configures OSPF on the specified interface.

**Syntax**
**add interface** [**InterfaceName=**]*InterfaceName* [**areaid**=]*IPAddress* [**addr**=]*IPAddress* [**mask**=]*SubnetMask*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface on which you want to add OSPF. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ areaid=] *IPAddress***
Required. Specifies the IP address of the area to which the interface belongs.

**[ addr=] *IPAddress***
Required. Specifies the IP address of the interface to add.

**[ mask=] *SubnetMask***
Required. Specifies the subnet mask of the interface to add.

**add routefilter**
Adds an OSPF route filter.

**Syntax**
**add routefilter** [**Filter=**]*IPAddressSubnetMask*

**Parameters**
**[ Filter=] *IPAddressSubnetMask***
Required. Specifies, by IP address and subnet mask, the route you want to filter.

**add protofilter**
Adds an OSPF protocol filter.

**Syntax**
**add protofilter** [**filter=**]{**autostatic** | **local** | **rip** | **snmp** | **nondod** | **static**}

**Parameters**
**[ filter=]{ autostatic| local| rip| snmp| nondod| static}**
Required. Specifies the protocol filter you want to add. The **autostatic** parameter specifies a RIP autostatic route. The **local** parameter specifies a route added by the network to which the computer is connected. The **rip** parameter specifies that the filter is added using RIP. The **snmp** parameter specifies that the route is added using SNMP. The **nondodm** parameter specifies an RTM route that is present only when a DOD interface is connected. The **static** parameter specifies an RTM route.

**Remarks**

- The **add protofilter** command is used on an OSPF router acting as an Autonomous System Boundary Router. The filters you add filter other route sources by protocol source, such as static routes and RIP, so that they are not redistributed into the OSPF Autonomous System.

**install**
Installs OSPF.

**Syntax**
**install**

**Parameters**
none

**delete range**
Deletes the specified range for the area ID.

**Syntax**
**delete range** [**Areaid=**]*IPAddress* [**Range=**]*IPAddressSubnetMask*

**Parameters**
**[ Areaid=] *IPAddress***
Required. Specifies, by IP address, the area you want to delete.

**[ Range=] *IPAddressSubnetMask***
Required. Specifies the IP address and the subnet mask of the area range you want to delete.

**delete area**
Deletes the specified area.

**Syntax**
**delete area** [**Areaid=**]*IPAddress*

**Parameters**
**[ Areaid=] *IPAddress***
Required. Specifies, by IP address, the area you want to delete.

**Remarks**
The area ID is a 32-bit number expressed in dotted decimal notation that identifies the OSPF area. The area ID does not need to coincide with an IP address or an IP network ID. The area ID of 0.0.0.0 is reserved for the backbone. If the area represents a subnetted network, you can use the IP network number of the subnetted network for the area ID.

**delete virtif**
Deletes the specified virtual interface.

**Syntax**
**delete virtif** [**transareaid=**]*TransitAreaID* [**virtnbrid=**]*VirtualNeighborID*

**Parameters**
**[ transareaid=] *TransitAreaID***
Required. Specifies the IP address of the transit area used to exchange routing information between the backbone area border router and the area border router of the noncontiguous area.

**[ virtnbrid=] *VirtualNeighborID***
Required. Specifies the router ID of the virtual neighbor to delete.

**delete neighbor**
Deletes the specified OSPF neighbor.

**Syntax**
**delete neighbor** [**InterfaceName=**]*InterfaceName* [**addr**=]*IPAddress* [**nbraddr**=]*IPAddress* [**nbrprio=**]*Priority*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface from which you want to delete a neighbor. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ addr=] *IPAddress***
Required. Specifies, by IP address, the interface from which you want to delete a neighbor.

**[ nbraddr=] *IPAddress***
Required. Specifies the IP address of the neighbor to delete.

**[ nbrprio=] *Priority***
Required. Specifies the OSPF router priority of the neighbor to delete.

**delete interface**
Deletes OSPF from the specified interface.

**Syntax**
**delete interface** [**InterfaceName=**]*InterfaceName* [**addr**=]*IPAddress* [**mask**=]*SubnetMask*

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface from which you want to delete OSPF. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ addr=] *IPAddress***

Required. Specifies, by IP address, the interface to delete.

**[ mask=] *SubnetMask***

Required. Specifies, by subnet mask, the interface to delete.

**delete routefilter**

Deletes the specified OSPF route filter.

**Syntax**

**delete routefilter** [**Filter=**]*IPAddressSubnetMask*

**Parameters**

**[ Filter=] *IPAddressSubnetMask***

Required. Specifies the IP address and the subnet mask of the route filter you want to delete.

**delete protofilter**

Deletes the specified OSPF protocol filter.

**Syntax**

**delete protofilter** [**filter=**]{**autostatic** | **local** | **rip** | **snmp** | **nondod** | **static**}

**Parameters**

**[ filter=]{ autostatic| local| rip| snmp| nondod| static}**

Required. Specifies the protocol filter you want to delete. The **autostatic** parameter specifies a RIP autostatic route. The **local** parameter specifies a route added by the network to which the computer is connected. The **rip** parameter specifies that the filter is added using RIP. The **snmp** parameter specifies that the route is added using SNMP. The **nondodm** parameter specifies an RTM route that is present only when a DOD interface is connected. The **static** parameter specifies an RTM route.

**Remarks**

- The **delete protofilter** command is used on an OSPF router acting as an Autonomous System Boundary Router. The filters you add filter other route sources by protocol source, such as static routes and the Routing Information Protocol (RIP), so that they are not redistributed into the OSPF Autonomous System.

**uninstall**

Deletes OSPF.

**Parameters**

none

**set area**

Changes parameters for the specified area.

**Syntax**

**set area** [**areaid=**]*AreaID* [[**auth=**]{**none** | **password**}] [[**stubarea**=]{**yes** | **no**}] [[**metric**=]*StubMetric*] [[**sumadv**=]{**yes** | **no**}]

**Parameters**

**[ areaid=] *AreaID***

Required. Specifies, by IP address, the OSPF area for which you want to change parameters.

**[ auth=]{ none| password}**

Specifies whether passwords are required for the area. If passwords are required for an area, all interfaces that are in the same area and on the same network segment must use the same password, and interfaces that are in the same area but on different network segments can have different passwords. By default, passwords are required and the password is 12345678. Passwords are transmitted in plaintext format, so this option is for identification, not security.

**[ stubarea=]{ yes| no}**

Specifies whether this area is configured as a stub area, which is an OSPF area that does not enumerate external routes. Routes from outside the OSPF Autonomous System (AS) are not flooded into or through stub areas. Routing to AS external destinations in these areas is based only on a summary default route. This reduces the overhead for the routers in the stub

area.

**[ metric=] *StubMetric***
Specifies the cost of the summary default route that the router advertises into the stub area.

**[ sumadv=]{ yes| no}**
Specifies whether to import summary advertisements.

**set virtif**
Configures parameters for the specified virtual interface.

**Syntax**
**set virtif** [**transareaid=**]*TransitAreaID* [**virtnbrid =**] *VirtualNeighborID* [[**transdelay=**]*TransitDelay*] [[**retrans=**]*ReTransmitInterval*] [[**hello=**]*HelloInterval*] [[**dead=**]*DeadInterval*] [[**password=**]*Password*]

**Parameters**
**[ transareaid=] *TransitAreaID***
Required. Specifies the IP address of the transit area.

**[ virtnbrid=] *VirtualNeighborID***
Required. Specifies the router ID of the virtual neighbor.

**[ transdelay=] *TransitDelay***
Specifies the estimated number of seconds it takes to transmit a link state update packet over this interface. If you do not specify this parameter, the estimated delay is set to 1 second.

**[ retrans=] *ReTransmitInterval***
Specifies the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. A typical value for a local area network is 5 seconds.

**[ hello=] *HelloInterval***
Specifies, in seconds, the interval between transmissions of hello packets by the router on the interface. This setting must be the same for all routers that are connected to the same network. A typical value for a local area network is 10 seconds.

**[ dead=] *DeadInterval***
Specifies, in seconds, how long a router can fail to respond before a neighboring router should consider this router is malfunctioning. This setting should be an integral multiple of the hello interval (commonly 4). This value must be the same for all OSPF router interfaces that are connected to the same network segment.

**[ password=] *Password***
Specifies the password for this interface, if passwords are required for the area. All interfaces that are in the same area and on the same network must use the same password. By default, passwords are required and the password is 12345678.

**set interface**
Configures OSPF parameters for an interface.

**Syntax**
**set interface** [**interfacename=**]*InterfaceName* [**addr=**]*IPAddress* [**mask=**]*SubnetMask* [[**state=**]{**enable** | **disable**}] [[**areaid=**]*AreaID*] [[**type=**]{**nbma** | **p2p** | **broadcast**}] [[**prio=**]*Priority*] [[**transdelay=**]*TransitDelay*] [[**retrans=**]*ReTransmitInterval*] [[**hello=**]*HelloInterval*] [[**dead=**]*DeadInterval*] [[**poll=**]*PollInterval*] [[**metric=**]*Metric*] [[**password=**]*Password*] [[**mtu=**]*MTUSize*]

**Parameters**
**[ interfacename=] *InterfaceName***
Required. Specifies, by name, the interface for which you want to configure parameters. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ addr=] *IPAddress***
Required. Specifies the IP address of the interface.

**[ mask=] *SubnetMask***
Required. Specifies the subnet mask of the interface.

**[ state=]{ enable| disable}**
Specifies whether the interface is enabled. The **enable** parameter enables the interface. The **disable** parameter disables the interface.

**[ areaid=]** ***AreaID***
Specifies the IP address of the interface area.

**[ type=]{ nbma| p2p| broadcast}**
Specifies the type of OSPF interface for which you want to configure parameters. The **nbma** parameter specifies a non-broadcast multiple access (NBMA) OSPF interface. The **p2p** parameter specifies a point-to-point OSPF interface. The **broadcast** parameter specifies a broadcast OSPF interface.

**[ prio=]** ***Priority***
Specifies the interface priority.

**[ transdelay=]** ***TransitDelay***
Specifies the estimated number of seconds it takes to transmit a link state update packet over this interface. If you do not specify this parameter, the estimated delay is set to 1 second.

**[ retrans=]** ***ReTransmitInterval***
Specifies the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. A typical value for a local area network is 5 seconds.

**[ hello=]** ***HelloInterval***
Specifies, in seconds, the interval between transmissions of hello packets by the router on the interface. This setting must be the same for all routers that are connected to the same network. A typical value for a local area network is 10 seconds.

**[ dead=]** ***DeadInterval***
Specifies, in seconds, how long a router can fail to respond before a neighboring router should consider this router to be malfunctional. This setting should be an integral multiple of the hello interval (commonly 4). This value must be the same for all OSPF router interfaces connected to the same network segment.

**[ poll=]** ***PollInterval***
Specifies the number of seconds between OSPF network polls for non-broadcast multiple access (NBMA) interfaces only. You should set the poll interval to be at least twice as long as the dead interval. A typical value for an X.25 network is two minutes.

**[ metric=]** ***Metric***
Specifies the cost of sending a packet out of this interface. The value set here is advertised as the link cost for this interface in the link state advertisement of the router. Faster interfaces usually have lower costs. The maximum setting is 32,767.

**[ password=]** ***Password***
Specifies the password for this interface, if passwords are required for the area. All interfaces that are in the same area and on the same network must use the same password. By default, passwords are required and the password is 12345678.

**[ mtu=]** ***MTUSize***
Specifies, in bytes, the maximum size of IP packets carrying OSPF information that can be sent without fragmentation. The default IP MTU for an Ethernet network is 1,500 bytes.

#### set global
Configures global OSPF parameters.

#### Syntax
**set global** [**routerid=**]*IPAddress* [**asborder=**]{**yes** | **no**} [[**loglevel=**]{**none** | **error** | **warn** | **info**}]

#### Parameters
**[ routerid=]** ***IPAddress***
Required. Specifies, by IP address, the router for which you want to configure parameters.

**[ asborder=]{ yes| no}**
Required. Specifies whether the router acts as an Autonomous System (AS) boundary router.

**[ loglevel=]{ none| error| warn| info}**
Specifies which events should be logged. The **none** parameter specifies that no events related to OSPF should be logged. The **error** parameter specifies that only errors related to OSPF should be logged. The **warn** parameter specifies that only warnings related to OSPF should be logged. The **info** parameter specifies that all events related to OSPF should be logged.

#### set routefilter
Configures the OSPF route filter action.

#### Syntax

**set routefilter** [**action=**]{**drop** | **accept**}

**Parameters**

**[ action=]{ drop| accept}**

Required. Specifies whether the route filter drops or accepts packets. The **drop** parameter specifies that packets should be dropped. The **accept** parameter specifies that packets should be accepted.

**set protofilter**

Configures OSPF protocol filter action.

**Syntax**

**set protofilter** [**action=**]{**drop** | **accept**}

**Parameters**

**[ action=]{ drop| accept}**

Required. Specifies whether the protocol filter accepts or drops packets. The **drop** parameter specifies that packets should be dropped. The **accept** parameter specifies that packets should be accepted.

**show global**

Displays global OSPF parameters.

**Syntax**

**show global**

**Parameters**

none

**show area**

Displays area parameters.

**Syntax**

**show area**

**Parameters**

none

**show virtif**

Displays parameters for all virtual interfaces.

**Syntax**

**show virtif**

**Parameters**

none

**show interface**

Displays OSPF configuration information for the specified interface.

**Syntax**

**show interface** [**InterfaceName=**]*InterfaceName*

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface for which you want to display information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**show routefilter**

Displays information about the OSPF route filter.

**Syntax**

**show routefilter**

**Parameters**

None

**show protofilter**

Displays information about the OSPF protocol filter.

**Syntax**

**show protofilter**

**Parameters**
none

**show area**
Displays information about OSPF areas.

**Syntax**
**show area**

**Parameters**
none

**show virtif**
Displays information about OSPF virtual interfaces.

**Syntax**
**show virtif**

**Parameters**
none

**Netsh Routing IP Autodhcp commands**
The following commands are available for use in the Netsh Routing IP Autodhcp context.

To view the command syntax, click a command:

- add exclusion
- delete exclusion
- install
- set global
- set interface
- show global
- show interface
- uninstall

**add exclusion**
Adds an exclusion to the DHCP allocator scope.

**Syntax**
**add exclusion** *IPAddress*

**Parameters**
***IPAddress***
Required. Specifies an IP address to exclude from the DHCP allocator scope.

**delete exclusion**
Deletes an exclusion from the DHCP allocator scope.

**Syntax**
**delete exclusion** *IPAddress*

**Parameters**
***IPAddress***
Required. Specifies an IP address already excluded from the DHCP allocator scope.

**install**

Installs the routing protocol under IP.

**Syntax**
**install**

**Parameters**
none

**set global**
Configures global parameters for DHCP allocation.

**Syntax**
**set global** {[[**scopenetwork=**]*IPAddress*] | [[**scopemask=**]*SubnetMask*] | [[**leasetime=**]*Minutes*] | [[**loglevel=**]{**none** | **error** | **warn** | **info**}]}

**Parameters**
**[ scopenetwork=] *IPAddress***
Specifies the IP address for the DHCP allocator scope.

**[ scopemask=] *SubnetMask***
Specifies the subnet mask associated with the IP address for the scope.

**[ leasetime=] *Minutes***
Specifies, in minutes, the lease duration time.

**[ LogLevel=]{ none| error| warn| info}**
Specifies which events should be logged. The **none** parameter specifies that no events related to DHCP should be logged. The **error** parameter specifies that only errors related to DHCP should be logged. The **warn** parameter specifies that only warnings related to DHCP should be logged. The **info** parameter specifies that all events related to DHCP should be logged.

**Examples**
To set the IP network for DHCP allocation to 10.10.10.0, with an associated scope subnet mask of 255.255.255.0, a lease time of 11520 minutes (8 days), and logging of information for all related events, type:

**set global 10.10.10.0 255.255.255.0 11520 info**

To adjust only the lease time globally for all scopes on all interfaces to a value of 4320 minutes (3 days), type:

**set global leasetime=4320**

**set interface**
Configures DHCP allocator parameters for the specified interface.

**Syntax**
**set interface** [**InterfaceName=**]*InterfaceName* [**mode=**]{**enable** | **disable**}

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface on which to configure parameters. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ mode=]{ enable| disable}**
Required. Specifies whether DHCP allocation is enabled or disabled for the interface.

**show global**
Displays the DHCP allocator global configuration.

**Syntax**
**show global**

**Parameters**
none

**show interface**
Displays the DHCP allocator configuration for the specified interface.

**Syntax**
**show interface** [**InterfaceName=**]*InterfaceName*

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface for which you want to display information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**uninstall**

Removes the routing protocol under IP.

**Syntax**

**uninstall**

**Parameters**

none

**Netsh Routing IP Dnsproxy commands**

The following commands are available for use in the Netsh Routing IP Dnsproxy context.

To view the command syntax, click a command:

- set global
- install
- set interface
- show global
- show interface
- uninstall

**set global**

Sets global DNS proxy parameters.

**Syntax**

**set global** [**querytimeout=**]*Integer* [[**dnsmode=**]{**enable** | **disable**}] [[**loglevel=**]{**none** | **error** | **warn** | **info**}]

**Parameters**

**[ querytimeout=] *Integer***

Required. Specifies a time-out, in seconds, for proxied DNS queries.

**[ dnsmode=]{ enable| disable}**

Specifies whether DNS proxy operation is enabled or disabled.

**[ loglevel=]{ none| error| warn| info}**

Specifies which events should be logged. The **none** parameter specifies that no events related to DNS proxy should be logged. The **error** parameter specifies that only errors related to DNS proxy should be logged. The **warn** parameter specifies that only warnings related to DNS proxy should be logged. The **info** parameter specifies that all events related to DNS proxy should be logged.

**Examples**

To specify that queries should time out after 10 seconds, to enable DNS proxy operation, and to prevent logging of DNS proxy events, type:

**set global 10 enable none**

**install**

Installs the routing protocol under IP.

**Syntax**

**install**

**Parameters**

none

**set interface**
Configures DNS proxy parameters for an interface.

**Syntax**
**set interface** [**InterfaceName=**]*InterfaceName* [[**mode=**]{**enable** | **disable** | **default**}]

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface for which to configure parameters. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ mode=]{ enable| disable| default}**
Specifies whether DNS proxy is enabled, disabled, or set to use the default for the specified interface. The **enable** parameter specifies that DNS proxy is enabled. The **disable** parameter specifies that DNS proxy is disabled. The **default** parameter specifies that DNS proxy is set according to the default for the interface.

**show global**
Displays the DNS proxy global configuration.

**Syntax**
**show global**

**Parameters**
none

**show interface**
Displays the DNS proxy configuration for the specified interface.

**Syntax**
**show interface** [**InterfaceName=**]*InterfaceName*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface for which you want to display information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**uninstall**
Removes the routing protocol under IP.

**Syntax**
**uninstall**

**Parameters**
none

**Netsh Routing IP NAT commands**
The following commands are available for use in the Netsh Routing IP NAT context.

To view the command syntax, click a command:

- add addressmapping
- add addressrange
- add interface
- add portmapping
- delete addressmapping
- delete addressrange
- delete interface

- delete portmapping
- set global
- set interface
- show global
- show interface

**add addressmapping**
Adds an IP address mapping to the network address translation (NAT) address pool for the specified interface.

**Syntax**
**add addressmapping** [**InterfaceName=**]*InterfaceName* [**public=**]*IPAddress* [**private=**]*IPAddress* [**inboundsessions=**] {**enable** | **disable**}

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface on which you want to use network address translation. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ public=] *IPAddress***
Required. Specifies the IP address used on a public network.

**[ private=] *IPAddress***
Required. Specifies the IP address in the pooled range of addresses available for use on the private network.

**[ inboundsessions=]{ enable| disable}**
Required. Specifies whether inbound sessions are enabled or disabled. The **enable** parameter specifies that inbound sessions are enabled. The **disable** parameter specifies that inbound sessions are disabled.

**Examples**
To map the public IP address 10.11.11.1 to the private IP address 10.10.10.1 and to disable inbound sessions for this NAT table mapping, type:

**add addressmapping Local Area Connection 10.11.11.1 10.10.10.1 disable**

**add addressrange**
Adds an address range to the NAT address pool for the specified interface.

**Syntax**
**add addressrange** [**InterfaceName=**]*InterfaceName* [**start=**]*IPAddress* [**end=**]*IPAddress* [**mask=**]*SubnetMask*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface to whose address pool you want to add an address range.

**[ start=] *IPAddress***
Required. Specifies the starting IP address for the address range.

**[ end=] *IPAddress***
Required. Specifies the ending IP address for the address range.

**[ mask=] *SubnetMask***
Required. Specifies the IP subnet mask associated with the network range bounded by the start and end IP addresses.

**add interface**
Configures NAT on the specified interface.

**Syntax**
**add interface** [**InterfaceName=**]*InterfaceName* [[**mode=**]{**full** | **addressonly** | **private**}]

**Parameters**
**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface on which you want to configure NAT. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ mode=]{ full| addressonly| private}**
Specifies whether the interface uses full, addressonly, or private mode.

**add portmapping**
Adds a protocol port mapping on the NAT interface.

**Syntax**
**add portmapping** [**InterfaceName=**]*InterfaceName* [**proto=**]{**tcp** | **udp**} [**publicip=**]{*IPAddress* | **0.0.0.0**} [**publicport=**]*Integer* [**privateip=**]*IPAddress* [**privateport=**]*Integer*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface for which you want to add a port mapping. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ proto=]{ tcp| udp}**
Required. Specifies whether the protocol type is set to TCP or UDP.

**[ publicip=]{ *IPAddress*| 0.0.0.0}**
Required. Specifies an external IP address on the public network or 0.0.0.0 to indicate any IP address not specified within the private network address range.

**[ publicport=] *Integer***
Required. Specifies the public protocol port by using a number from 0 to 9999.

**[ privateip=] *IPAddress***
Required. Specifies an IP address within the private network range.

**[ privateport=] *Integer***
Required. Specifies the private protocol port by using a number from 0 to 9999.

**delete addressmapping**
Deletes an address mapping from the NAT address pool for the specified interface.

**Syntax**
**delete addressmapping** [**InterfaceName=**]*InterfaceName* [[**public=**]*IPAddress*]

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface for which you want to delete an address mapping. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ public=] *IPAddress***
Specifies the IP address to be deleted. If you do not specify an address, all address mappings from the NAT address pool are deleted from the interface.

**delete addressrange**
Deletes an address range from the NAT address pool for the specified interface.

**Syntax**
**delete addressrange** [**InterfaceName=**]*InterfaceName* [**start=**]*IPAddress*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface from which you want to delete an address range. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ start=] *IPAddress***
Required. Specifies the starting IP address of the range you want to delete.

**delete interface**

Removes network address translation (NAT) from the specified interface.

**Syntax**

**delete interface** [**InterfaceName=**]*InterfaceName*

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface for which you want to remove the use of NAT. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**delete portmapping**

Deletes a protocol port mapping from the specified NAT-enabled interface.

**Syntax**

**delete portmapping** [**InterfaceName=**]*InterfaceName* [[**proto=**]{**tcp** | **udp**}] [**publicip=**]{*IPAddress* | **0.0.0.0**}] [[**publicport=**]*Integer*] [[**privateip=**]*IPAddress*] [**privateport=**]*Integer*]

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface for which you want to delete a port mapping. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, **"***Interface Name***"**).

**[ proto=]{ tcp| udp}**

Specifies the protocol type. The **tcp** parameter specifies that the protocol type is Transmission Control Protocol. The **udp** parameter specifies that the protocol type is User Datagram Protocol.

**[ publicip=]{ *IPAddress*| 0.0.0.0}**

Specifies an external IP address on the public network or 0.0.0.0 to indicate any IP address not specified within the private network address range.

**[ publicport=] *Integer***

Specifies the public protocol port by using a number from 0 to 9999.

**[ privateip=] *IPAddress***

Specifies an IP address within the private network range.

**[ privateport=] *Integer***

Specifies the private protocol port by using a number from 0 to 9999.

**set global**

Sets global parameters for NAT.

**Syntax**

**set global** {[[**tcptimeoutmins=**]*Integer*] | [[**udptimeoutmins=**]*Integer*] | [[**LogLevel**=]{**none** | **error** | **warn** | **info**}]}

**Parameters**

**[ tcptimeoutmins=] *Integer***

Specifies, in minutes, the time-out value for TCP mappings.

**[ udptimeoutmins=] *Integer***

Specifies, in minutes, the time-out value for UDP mappings.

**[ LogLevel=]{ none| error| warn| info}**

Specifies which events should be logged. The **none** parameter specifies that no events related to NAT should be logged. The **error** parameter specifies that only errors related to NAT should be logged. The **warn** parameter specifies that only warnings related to NAT should be logged. The **info** parameter specifies that all events related to NAT should be logged.

**set interface**

Configure NAT parameters for the specified interface.

**Syntax**

**set interface** [**InterfaceName=**]*InterfaceName* [**mode=**]{*Full* | *AddressOnly* | *Private*}

**Parameters**

**[ InterfaceName=] *InterfaceName***

Required. Specifies, by name, the interface for which you want to configure parameters. The *InterfaceName* parameter must

match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**[ mode=]{ *Full*| *AddressOnly*| *Private*}**
Required. Specifies whether the interface should be enabled for full, addressonly, or private mode.

**show global**
Displays network address translation (NAT) global configuration.

**Syntax**
**show global**

**Parameters**
none

**show interface**
Displays network address translation (NAT) configuration for the specified interface.

**Syntax**
**show interface** [**InterfaceName=**]*InterfaceName*

**Parameters**
**[ InterfaceName=] *InterfaceName***
Required. Specifies, by name, the interface for which you want to display information. The *InterfaceName* parameter must match the name of the interface as specified in Network Connections. If *InterfaceName* contains spaces, use quotation marks around the text (for example, "*Interface Name*").

**Netsh Routing example**
To enable router discovery for the interface named Local Area Connection, specifying a minimum interval of 7 minutes, a maximum interval of 10 minutes, a life of 30 minutes for router discovery packets, and a zero (0) or no preference level, type the following from the netsh routing prompt:

**add interface "Local Area Connection" enable 7 10 30 0**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Netsh commands for AAAA
Netsh commands for DHCP
Netsh diagnostic (diag) commands
Netsh commands for Interface IP
Netsh commands for remote access (ras)
Netsh commands for WINS
Command-line reference A-Z
Command shell overview

# Netsh commands for RPC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for RPC**
The Netsh commands for RPC provide a command-line tool that you can use to change, reset, or display either specified or all subnet and interface settings.

*Interface* refers to an IPv4 address that is active on the computer and has a network adapter associated with it. *Subnet* refers to a set of zero or more interfaces that resolve to the same subnet address when applied with the appropriate mask.

You can run the Netsh RPC commands from the command prompt for the **netsh rpc** context. For these commands to work at the Windows Server 2003 family command prompt, you must type **netsh rpc** before typing commands and parameters as they appear in the syntax below.

For more information about **netsh**, see Netsh overview and Enter a netsh context.

To view the command syntax, click a command:

- add
- delete
- show interfaces
- show settings
- reset

**add**
This command adds specified subnets. Each subnet must be in dotted-decimal format (for example, 192.168.255.255).

**Syntax**
**add** *Subnet1*[*Subnet2*]

**Parameters**
***Subnet1***
Required. Specifies the first subnet in the list of added subnets.

**[ *Subnet2*]**
Specifies the second subnet in the list of added subnets.

**/?**
Displays help at the command prompt.

**Examples**
The following example adds two interfaces with subnet addresses of 192.168.255.255 and 172.16.1.0.

**add 192.168.255.255 172.16.1.0**

**delete**
This command deletes specified subnets. Each subnet must be in dotted-decimal format (for example, 192.168.255.255).

**Syntax**
**delete** *Subnet1*[*Subnet2*]

**Parameters**
***Subnet1***
Required. Specifies the first subnet in the list of deleted subnets.

**[ *Subnet2*]**
Specifies the second subnet in the list of deleted subnets.

**/?**
Displays help at the command prompt.

**Examples**
The following example deletes two interfaces with subnet addresses of 192.168.255.255 and 172.16.1.0.

**delete 192.168.255.255 172.16.1.0**

**show interfaces**
This command displays all interfaces and the state of each (enabled or disabled) on the computer, grouped by subnet.

**Syntax**
**show interfaces**

**Parameters**
**/?**
Displays help at the command prompt.

**show settings**
This command displays the selective binding settings for each subnet. Setting values are **none**, **add list**, or **delete list**. If the selective binding settings are either **add list** or **delete list**, then the associated subnets are also displayed.

**Syntax**
**show settings**

**Parameters**
**/?**
Displays help at the command prompt.

**reset**
This command sets the selective binding settings to **none**, which specifies that RPC listens on all interfaces.

**Syntax**
**reset**

**Parameters**
**/?**
Displays help at the command prompt.

## Note

- To run all of these commands (with the exception of **show settings**), you must either have administrative credentials or be a user with permission to write to the **HKEY_LOCAL_MACHINE** registry key.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Netsh overview
Command shell overview

# Netsh commands for WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for WINS**

The Netsh commands for WINS offer a command-line tool for administering WINS servers that can be used as an equivalent alternative to console-based management. This is useful in the following situations:

- When managing WINS servers in wide area networks (WANs), commands can be used in interactive mode at the Netsh command prompt to better manage slow-speed network links.
- When managing a large number of WINS servers, commands can be used in batch mode to script and automate recurring administrative tasks that need to be performed for all WINS servers.

You can run these commands from the Windows Server 2003 family command prompt or from the command prompt for the Netsh wins context. For these commands to work at the Windows Server 2003 family command prompt, you must type **netsh wins** before typing commands and parameters as they appear in the syntax below. There might be functional differences between Netsh context commands on Windows 2000 and the Windows Server 2003 family.

**Notes**

- Unless otherwise specified, field tags (such as **Name=** and **EndChar=**) are optional. If tags are used, they must be applied to all of the parameters in the command. If tags are omitted, a value must be provided for each of the fields that are supported in this command.
- The terms Persona Non Grata and Persona Grata refer to the Windows Server 2003 family features **Block records for these owners** and **Accept records only for these owners**, respectively.
- For more information about Netsh, see Netsh overview and Enter a netsh context.
- For an example of how to use Netsh commands for WINS in batch files and scripts, see Netsh WINS example.
- For more information about how to use the Netsh commands for WINS, see Use WINS Command-line Tools.

**Netsh WINS**

The following commands are available at the **wins>** prompt, which is rooted within the netsh environment.

To view the command syntax, click a command:

- server
- add
- add name
- add partner
- add pgserver
- add pngserver
- check
- check database

- check name
- check version
- delete
- delete name
- delete owners
- delete partner
- delete pgserver
- delete pngserver
- delete records
- init
- init backup
- init import
- init pull
- init pullrange
- init push
- init replicate
- init restore
- init scavenge
- init search
- reset
- reset statistics
- set
- set autopartnerconfig
- set backuppath
- set burstparam
- set defaultparam
- set logparam
- set migrateflag

- set namerecord
- set periodicdbchecking
- set pgmode
- set pullparam
- set pullpartnerconfig
- set pushparam
- set pushpartnerconfig
- set replicateflag
- set startversion
- show
- show browser
- show database
- show info
- show name
- show partner
- show partnerproperties
- show pullpartnerconfig
- show pushpartnerconfig
- show recbyversion
- show reccount
- show server
- show serverstatistics
- show version
- show versionmap

**server**
Switches to the specified server context. Used without parameters, the local WINS server is assumed as the default.

**Syntax**
**server** [{****_ServerName_ | _IPAddress_}]

**Parameters**

**{ *ServerName*| *IPAddress*}**
Specifies the remote (non-local) WINS server to which to switch the Netsh WINS command-line context. Identifies the server by server name or IP address.

**Examples**
In the first example, this command changes the Netsh context to the local WINS server.

In the second example, this command changes the Netsh context to a remote server named \\WINS-SRV1.

In the third example, this command changes the Netsh context to a remote server with IP address 10.0.0.1.

**server**

**server \\WINS-SRV1**

**server 10.0.0.1**

**Netsh WINS server**
The following commands are available at the **wins server>** prompt, which is rooted within the netsh environment.

To view the command syntax, click a command:

**add**
**add name**
Adds a name record to the database on the specified WINS server.

**Syntax**
**add name** [**Name=**]*ComputerName* [[**EndChar=**]*16thCharInHex*] [[**Scope=**]*ScopeName*] [[**RecType=**]{**0** | **1**}] [[**Group=**]{**0** | **1** | **2** | **3** | **4**}] [[**Node=**]{**0** | **1** | **3**}] [**IP=**]**{***IPAddress1*[**,***IPAddress2*,*IPAddress3*... ]**}**

**Parameters**
**[ Name=] *ComputerName***
Required. Specifies the name to add and register in the WINS server database.

**[ EndChar=] *16thCharInHex***
Indicates the 16th character of the name in hexadecimal. If unspecified, record types are added based on the value of **Group=**. For more information, see Remarks.

**[ Scope=] *ScopeName***
Specifies the name of the NetBIOS scope.

**[ RecType=]{ 0| 1}**
Indicates the record type: **0**-Static (default), **1**-Dynamic.

**[ Group=]{ 0| 1| 2| 3| 4}**
Indicates the group type: **0**-Unique (default), **1**-Group, **2**-Internet **3**-Multihomed, **4**-Domain Name.

**[ Node=]{ 0| 1| 3}**
Indicates the NetBIOS node type: **0**-B Node, **1**-P Node (default), **3**-H Node.

**[ IP=] {*IPAddress1*[ ,*IPAddress2*,*IPAddress3*...] }**
Required. Specifies the IP address of the name record to add. If the computer is multihomed, more than one IP address can be specified. Use braces ({}) to enclose one or more IP addresses, and separate multiple IP addresses with commas.

**Remarks**

- If **EndChar=** and **Group=** are both used, **EndChar=** overrides **Group=**.
- When **EndChar=** is not specified, the following record types are added to the WINS database based on the value of **Group=**:

| Group parameter value | Record type added |
|---|---|
| Unique group (0) | 03h, 20h, 00h |
| Normal group (1) | 1eh |

| | |
|---|---|
| Special group/Internet (2) | 20h |
| Multihomed group (3) | 03h, 20h, 00h |
| Special group/domain (4) | 1ch |

**Examples**
In the first example, this command adds a default NetBIOS record set of the 00h, 03h, and 20h records for a computer named HOSTA with a mapped IP address of 10.0.0.1.

In the second example, this command adds a 20h record for HOSTB, a multihomed computer with IP addresses 10.0.0.2 and 10.0.0.3, on SUBNETA, with a static record type and P-Node configuration.

**add name Name=HOSTA IP={10.0.0.1}**

**add name HOSTB 20 SUBNETA 0 3 1{10.0.0.2,10.0.0.3}**

**add partner**
Adds a replication partner to the list of replication partners on the specified WINS server.

**Syntax**
**add partner** [**Server=**]*IPAddress* [[**NetBIOS=**]*ServerNetBIOSName*] [[**Type=**]{**0** | **1** | **2**}]

**Parameters**
**[ Server=] *IPAddress***
Required. Specifies the replication partner to add to the WINS server. Identifies the partner server by IP address.

**[ NetBIOS=] *ServerNetBIOSName***
Specifies the replication partner to add to the WINS server. Identifies the partner server by NetBIOS name.

**[ Type=]{ 0| 1| 2}**
Indicates the type of partner to add: **0**-Pull, **1**-Push, **2**-Both (default).

**Remarks**

- If the IP address does not resolve, the NetBIOS name parameter is used to create the partner record. Otherwise, the NetBIOS parameter is not used.

**Examples**
In the following example, this command adds a push/pull replication partner with an IP address of 10.2.2.2 to the current WINS server.

**add partner 10.2.2.2**

**add pgserver**
Adds a list of Persona Grata servers to the specified WINS server.

**Syntax**
**add pgserver** [**Servers=**]**{***ListOfServerIPAddresses***}**

**Parameters**
**[ Servers=] {*ListOfServerIPAddresses*}**
Required. Specifies the Persona Grata server or servers to add. Identifies each server by its IP address, enclosed by braces ({}). If you provide a list, IP addresses must be separated by commas and enclosed by braces ({}).

**Remarks**

- This command adds the new server or servers after checking for duplicates. The Persona Grata (PG) list is used to specify owners (other WINS servers) whose records are to be accepted during replication.

**Examples**
In the following example, this command adds the WINS servers 10.2.2.2 and 10.2.2.3 to the Persona Grata list for the current WINS server.

**add pgserver {10.2.2.2,10.2.2.3}**

**add pngserver**
Adds a list of Persona Non Grata servers to the specified WINS server.

**Syntax**
**add pngserver** [**Servers=**]**{***ListOfServerIPAddresses***}**

**Parameters**
**[ Servers=] {*ListOfServerIPAddresses*}**
Required. Specifies the Persona Non Grata server or servers to add. Identifies each server by its IP address, enclosed by braces ({}). If a you provide a list, IP addresses must be separated by commas and enclosed by braces ({}).

**Remarks**

- This command adds the new server or servers after checking for duplicates. The Persona Non Grata (PNG) list is used to specify owners (other WINS servers) whose records are to be blocked during replication.

**Examples**
In the following example, this command adds the WINS servers 10.2.2.2 and 10.2.2.3 to the Persona Non Grata list for the current WINS server.

**add pngserver {10.2.2.2,10.2.2.3}**

**check**
**check database**
Checks the consistency of the WINS database. Used without parameters, a consistency check is performed on all replicas whose verification interval has expired. The consistency check is not be performed immediately if overload conditions exist, but occurs after the configured verification interval period has passed.

**Syntax**
**check database** [[**All=**]{**1** | **0**}] [[**Force=**]{**1** | **0**}]

**Parameters**
**[ All=]{ 1| 0}**
Indicates whether to perform a consistency check on all replicas or to perform a consistency check on the replicas whose verification interval has expired: **1**-All, **0**-Otherwise (default).

**[ Force=]{ 1| 0}**
Indicates whether to override the WINS consistency check under overload conditions: **1**-Override, **0**-Otherwise (default).

**Remarks**

- It is a network and resource intensive operation to check all replicas for consistency.

- Overload condition occurs when the consistency check command is repeated within a duration of one hour.

**Examples**
In the following example, this command does a consistency check for all replicas. If overload conditions exist, the consistency check is performed after the configured verification interval has passed.

**check database 1**

**check name**
Checks a list of name records against a specified set of WINS servers.

**Syntax**
**check name** {**NameFile=***FileName* | **NameList={***NamesList***}**} {**ServerFile=***FileName* | **ServerList={***ListOfServerIPAddresses***}**} [**InclPartner=**{**Y** | **y**}]

**Parameters**
**{ NameFile=*FileName*| NameList={*NamesList*}}**
Required. Specifies the file that contains the list of names to check or specifies a list of names to check. If you provide a list, names must be separated by commas and enclosed by braces ({}).

**{ ServerFile=*FileName*| ServerList={*ListOfServerIPAddresses*}}**
Required. Specifies the file with the list of server IP addresses or specifies a list of the server IP addresses to check. If you provide a list, IP addresses must be separated by commas and enclosed by braces ({}).

**InclPartner={ Y| y}**
Specifies that all partners will be included in the list of servers. Partners are not included by default.

**Remarks**

- Tags are required for this command. Whenever inconsistencies are found, repair is attempted.
- The format for *Name* in both *NameFile* and *NameList* is *Name*16thchar*, where *Name* is the record name and 16thchar is the 16th character for the record.
- In the file, name entries are separated by commas or carriage returns/line feeds.

**Examples**
In the first example, this command checks WINS data against the Names.txt and Servers.txt files, assuming those files to be located in the specified path location (the root of drive D).

In the second example, this command checks the names Name1 with 16th character 03, 00 and 20h respectively against the servers with IP addresses 10.2.2.2 and 10.2.2.3.

**check name NameFile=D:\Names.txt ServerFile=D:\Servers.txt**

**check name NameList={Name1*03,Name1*00,Name1*20} ServerList={10.2.2.2,10.2.2.3}**

**check version**
Checks the consistency of version ID numbers for WINS record owners in the WINS database.

**Syntax**
**check version** [**Server=**]*IPAddress* [[**File=**]*FileName*]

**Parameters**
**[ Server=] *IPAddress***
Required. Specifies the IP address of the start WINS server.

**[ File=] *FileName***
Specifies the file in which to store the index to IP address table. If omitted, **check version** creates the file Wins.rec in the current directory and stores the index to IP address table there.

**Remarks**

- If you provide an invalid file name, *FileName* is set to Wins.rec in the current directory.
- You might need to increase the width of the command window to make it wide enough to view the table correctly.

**Examples**
In the following example, this command verifies the version number consistency of the records for a WINS server with the IP address 10.2.2.2 at its replication partners.

**check version 10.2.2.2**

**delete**
**delete name**
Deletes a registered name from the WINS server database.

**Syntax**
**delete name** [**Name=**]*RegisteredName* [**EndChar=**]*16thCharInHex* [[**Scope=**]*ScopeName*]

**Parameters**
**[ Name=] *RegisteredName***
Required. Specifies the registered name to be deleted.

**[ EndChar=] *16thCharInHex***
Required. Indicates the 16th character of the name of the record to be deleted in hexadecimal.

**[ Scope=] *ScopeName***
Specifies the NetBIOS scope identifier for the record to be deleted.

**Examples**
In the following example, this command deletes the messenger (03h) records with the record name HOSTA from the WINS database.

**delete name Name=HOSTA EndChar=03**

**delete owners**
Deletes or tombstones a list of owners and their records from the database on the specified WINS server.

**Syntax**
**delete owners** [**Servers=**]**{***ListOfServerIPAddresses***}**[[**Op=**]{**0** | **1**}]

**Parameters**
**[ Servers=] {*ListOfServerIPAddresses*}**
Required. Specifies a list of the owner servers with records to delete. Identifies the servers by IP addresses that must be separated by commas and enclosed by braces ({}).

**[ Op=]{ 0| 1}**
Indicates your option: **0**-Tombstone (default), **1**-Delete.

**Examples**
In the following example, this command deletes all records owned by the WINS server with the IP address 10.0.0.1.

**delete owners Servers={10.0.0.1} Op=1**

**delete partner**
Deletes either all replication partners or a specified replication partner from the list of replication partners on the specified WINS server. When used without parameters, all replication partners are deleted from both the pull and the push partner lists without a confirmation prompt.

**Syntax**
**delete partner** [[**Server=**]*IPAddress*] [[**Type=**]{**0** | **1** | **2**}] [[**Confirm=**]{**Y** | **y**}]

**Parameters**
**[ Server=] *IPAddress***
Specifies the partner to delete. Identifies the partner by IP address.

**[ Type=]{ 0| 1| 2}**
Indicates the type of partner list to delete: **0**-Pull, **1**-Push, **2**-Both (default).

**[ Confirm=]{ Y| y}**
Initiates and completes the operation without prompting for confirmation.

**Examples**
In the following example, this command deletes a push/pull type partner with an IP address of 10.0.0.1 from the current WINS server without prompting for confirmation.

**delete partner Server=10.0.0.1 Type=2 Confirm=Y**

**delete pgserver**
Deletes all Persona Grata servers or specific Persona Grata servers from the list of Persona Grata servers on the specified WINS server.

**Syntax**
**delete pgserver** [**Servers=**]**{***ListOfServerIPAddresses***}**

**Parameters**
**[ Servers=] {*ListOfServerIPAddresses*}**
Required. Specifies a list of Persona Grata servers to delete. Identifies Persona Grata servers by IP addresses separated by commas and enclosed by braces ({}).

**Remarks**

- The Persona Grata (PG) list specifies owners (other WINS servers) whose records are to be accepted during replication. If you use an empty set of braces ({}), **delete pgserver** deletes all of the servers on the PG server list.

**Examples**
In the following example, this command deletes the WINS servers 10.2.2.2 and 10.2.2.3 from the Persona Grata list for this

WINS server.

**delete pgserver Servers={10.2.2.2,10.2.2.3}**

**delete pngserver**
Deletes all Persona Non Grata servers or specific Persona Non Grata servers from the list of Persona Non Grata servers on the specified WINS server.

**Syntax**
**delete pngserver** [**Servers=**]**{***ListOfServerIPAddresses***}**

**Parameters**
**[ Servers=] {*ListOfServerIPAddresses*}**
Required. Specifies a list of Persona Non Grata servers to delete. Identifies Persona Non Grata servers by IP addresses separated by commas and enclosed by braces ({}).

**Remarks**

- The Persona Non Grata (PNG) list specifies owners (other WINS servers) whose records are to be blocked during replication. If you use an empty set of braces ({}), **delete pngserver** deletes all of the servers on the PNG server list.

**Examples**
In the following example, this command deletes the WINS servers 10.2.2.2 and 10.2.2.3 from the Persona Non Grata list for this WINS server.

**delete pngserver Servers={10.2.2.2,10.2.2.3}**

**delete records**
Deletes all records or a set of records from the current WINS server or marks records as tombstoned.

**Syntax**
**delete records** [**MinVer=**]**{***High***,***Low***}** [**MaxVer=**]**{***High***,***Low***}** [[**Op=**]{**0** | **1**}]

**Parameters**
**[ MinVer=] {*High*,*Low*}**
Required. Specifies the minimum high and the minimum low record versions from which to begin deletion. High and low values are separated by a comma and enclosed by braces ({}).

**[ MaxVer=] {*High*,*Low*}**
Required. Specifies the maximum high and maximum low record versions at which to end deletion. High and low values are separated by a comma and enclosed by braces ({}).

**[ Op=]{ 0| 1}**
Indicates your option: **0**-Delete (default), **1**-Tombstone.

**Remarks**

- If you set both **MinVer=** and **MaxVer=** to {0,0}, **delete records** deletes all records.

**Examples**
In the first example, this command marks all records on the WINS server as tombstoned.

In the second example, this command deletes all records between and including 2de2d and 2de3f.

**delete records MinVer={0,0} MaxVer={0,0} Op=1**

**delete records MinVer={0,2de2d} MaxVer={0,2de3f}**

**init**
**init backup**
Initiates backup of the WINS database to a specified directory. Used without parameters, this command provides a full backup to the default backup path that is currently set.

**Syntax**
**init backup** [[**Dir=**]*BackupDir*] [[**Type=**]{**0** | **1**}]

**Parameters**

**[ Dir=] *BackupDir***
Specifies the directory in which backups are stored. If unspecified, the default backup path that is currently set is used.

**[ Type=]{ 0| 1}**
Indicates the type of backup: **0**-Full backup (default), **1**-Incremental backup.

**Remarks**

- WINS backup can only be used for a local backup on the same server. Backup of remote WINS servers is not supported.
- Backup files will automatically be created under the wins_bak subdirectory of the **Dir=** directory.

**Examples**
In the following example, this command initiates a full backup of the WINS database on the local WINS server to the folder C:\WINSfiles\wins_bak. WINS will create the wins_bak subdirectory.

**init backup Dir=C:\WINSfiles Type=0**

**init import**
Initiates the import of static mappings from an Lmhosts file.

**Syntax**
**init import** [**File=**]*ImportFile*

**Parameters**
**[ File=] *ImportFile***
Required. Specifies the Lmhosts-formatted text file used for importing static mappings into the WINS database.

**Examples**
In the following example, this command imports the static mappings of the Fixedip.sam file, located at the C:\WINSfiles path, as new static mappings in the local WINS database.

**init import File=C:\WINSfiles\Fixedip.sam**

**init pull**
Initiates and sends a pull trigger to another WINS server.

**Syntax**
**init pull** [**Server=**]{*PullServerName* | *IPAddress*}

**Parameters**
**[ Server=]{ *PullServerName*| *IPAddress*}**
Required. Specifies the WINS server to which to send a pull trigger. Identifies the server by server name or IP address.

**Remarks**

- If you provide a server name that cannot be resolved to an IP address, the operation fails.

**Examples**
In the following example, this command initiates and sends a pull trigger to WINS server 10.2.2.2.

**init pull Server=10.2.2.2**

**init pullrange**
Initiates and pulls a range of records from another WINS server.

**Syntax**
**init pullrange** [**Owner=**]*OwnerServerIP* [**Server=**]*PullServerIP* [**MaxVer=**]**{***High*,*Low***}** [**MinVer=**]**{***High*,*Low***}**

**Parameters**
**[ Owner=] *OwnerServerIP***
Required. Specifies the WINS server with records to be pulled. Identifies the server by IP address.

**[ Server=] *PullServerIP***
Required. Specifies the WINS server that pulls the range of records. Identifies the server by IP address.

**[ MaxVer=] {*High*,*Low*}**

Required. Specifies the maximum high and the maximum low version number of the records to be pulled. High and low values are separated by a comma and enclosed by braces ({}).

**[ MinVer=] {*High*,*Low*}**
Required. Specifies the minimum high and the minimum low version number of the records to be pulled. High and low values are separated by a comma and enclosed by braces ({}).

**Remarks**

- If you set both **MaxVer=** and **MinVer=** to {0,0}, **init pullrange** pulls the entire range of records.

**Examples**
In the following example, this command initiates and sends a pull trigger to WINS server 10.2.2.3 for all records owned by WINS server 10.2.2.2.

**init pullrange Owner=10.2.2.2 Server=10.2.2.3 MaxVer={0,0} MinVer={0,0}**

**init push**
Initiates and sends a push trigger to another WINS server.

**Syntax**
**init push** [**Server=**]{****_PushServerName_ | _IPAddress_} [[**PropReq=**]{**0** | **1**}]

**Parameters**
**[ Server=]{ *PushServerName*| *IPAddress*}**
Required. Specifies the WINS server to which to send a push trigger. Identifies the server by name or IP address.

**[ PropReq=]{ 0| 1}**
Indicates your propagation request: **0**-Push to the specified server only (default), **1**-Have server propagate the push trigger to all of its partners.

**Remarks**

- If you provide a server name that cannot be resolved to an IP address, the operation fails.

**Examples**
In the following example, this command initiates and sends a push trigger to WINS server 10.2.2.2 and has it propagate the trigger to all of its replication partners.

**init push Server=10.2.2.2 PropReq=1**

**init replicate**
Initiates and forces an immediate replication of the database with replication partners.

**Syntax**
**init replicate**

**Parameters**
none

**Remarks**

- This command initiates push/pull replication with all replication partners for this server.

**init restore**
Initiates the restoration of the WINS database from a directory and file at the specified WINS server.

**Syntax**
**init restore** [**Dir=**]*RestoreDir*

**Parameters**
**[ Dir=] *RestoreDir***
Required. Specifies the directory from which to restore files.

**Remarks**

- WINS backups can only be restored locally on the same server. You cannot restore the WINS database from a remote

WINS backups can only be restored locally on the same server. You cannot restore the WINS database from a remote computer.

- **Dir=** should contain a subdirectory wins_bak, which will contain the database file, but this subdirectory should not be included in the **Dir=** parameter.

**Examples**

In the following example, this command restores the WINS database from the backup database file location of C:\WINSfiles.

**init restore Dir=C:\WINSfiles**

**init scavenge**

Initiates the scavenging of the WINS database for the specified WINS server.

**Syntax**

**init scavenge**

**Parameters**

none

**Remarks**

- This command removes outdated information from the WINS database.

**init search**

Initiates a search for the specified record name in the WINS database.

**Syntax**

**init search** [**Name=**]*RecordName* [[**EndChar=**]*EndCharacter*] [[**Case=**]{**0** | **1**}] [[**File=**]*FileName*]

**Parameters**

**[ Name=] *RecordName***

Required. Specifies the name of the record for which to search.

**[ EndChar=] *EndCharacter***

Indicates the 16th character of the NetBIOS name record. If unspecified, the search returns all records that match the string that you entered as a value for **Name=**. For more information, see Remarks.

**[ Case=]{ 0| 1}**

Indicates name case preference: **0**-Ignore name case (default), **1**-Use name case.

**[ File=] *FileName***

Specifies the file in which to store ASCII plain text output. If unspecified, output is displayed in the command prompt window.

**Remarks**

- If you do not supply a value for **EndChar=**, the search returns all records that match the string that is entered as a value for **Name=**. For example, if the database contains a record named MYPC-1, a search that uses MYPC as the value of the name parameter returns all records starting with this name pattern, regardless of record type:

| Record Name | Record Type |
|---|---|
| MYPC-1 | [00h]Workstation |
| MYPC-1 | [01h]Other |
| MYPC-1 | [03h]Messenger |

If you specify a value for **EndChar=**, only records that match the record name and record type are returned. For example, if the value for **Name=** is MYPC and the value for EndChar is 00, the following search results return:

| Record Name | Record Type |
|---|---|
| MYPC-1 | [00h]Workstation |

**Examples**
In the following example, this command searches for records named dc1 with an end character of 1c, ignores name case, and outputs the information to the names.txt file located in C:\WINSfiles.

**init search Name=dc1 EndChar=1c Case=0 File=C:\WINSfiles \names.txt**

**reset**
**reset statistics**
Resets the statistics for the local WINS server.

**Syntax**
**reset statistics**

**Parameters**
none

**set**
**set autopartnerconfig**
Sets the automatic replication partner configuration information for the server. Used without parameters, automatic partner configuration is reset (disabled).

**Syntax**
**set autopartnerconfig** [[**State=**]{**0** | **1**}] [[**Interval=**]*Value*] [[**TTL=**]*Value*]

**Parameters**
**[ State=]{ 0| 1}**
Indicates the automatic replication partner configuration information setting: **0**-Reset (default), **1**-Set.

**[ Interval=] *Value***
Specifies the multicast interval value, in seconds.

**[ TTL=] *Value***
Specifies the Time-To-Live value for multicast datagrams between partner WINS servers. The TTL value must be within the range of 1 to 32.

**Examples**
In the following example, this command sets the automatic configuration information for this WINS server for a multicast interval of 3,600 seconds and a Time to Live (TTL) value of 20.

**set autopartnerconfig State=0 Interval=3600 TTL=20**

**set backuppath**
Sets the backup parameters for the specified WINS server.

**Syntax**
**set backuppath** [**Dir=**]*NewPath* [[**Shutdown=**]{**0** | **1**}]

**Parameters**
**[ Dir=] *NewPath***
Required. Specifies the backup path for the server database. If unspecified, the default path is C:\.

**[ Shutdown=]{ 0| 1}**
Indicates the setting for backup during shutdown: **0**-Disable backup of the WINS database during shutdown (default), **1**-Enable backup of the WINS database during shutdown.

**Examples**
In the following example, this command sets the WINS server to automatically back up the WINS database to C:\WINSfiles each time that the server is shut down.

**set backuppath Dir=C:\WINSfiles Shutdown=1**

**set burstparam**
Sets burst handling parameters for the specified WINS server.

**Syntax**
**set burstparam** [**State=**]{**0** | **1**} [[**Value=**]*QueueSize*]

**Parameters**
**[ State=]{ 0| 1}**

Required. Indicates the setting for burst handling: **0**-Reset or disable burst handling for the WINS server, **1**-Set or enable burst handling for the WINS server.

**[ Value=] *QueueSize***
Specifies the value of the burst queue size. The queue size must be within the range of 50 to 5,000.

**Examples**
In the following example, this command sets the size of the burst queue to permit 1,000 client registration and renewal requests before the WINS server uses burst handling.

**set burstparam State=1 Value=1000**

**set defaultparam**
Sets the default parameters for WINS server configuration.

**Syntax**
**set defaultparam**

**Parameters**
none

**Remarks**

- This command sets all the configuration parameters for the WINS server to their default values. It is recommended that you run this command after WINS service installation to configure the server with default parameter settings.

**set logparam**
Sets or resets database and detailed event logging parameters. Used without parameters, the current values are maintained.

**Syntax**
**set logparam** [[**DBChange=**]{**0** | **1**}] [[**Event=**]{**0** | **1**}]

**Parameters**
**[ DBChange=]{ 0| 1}**
Indicates whether you want to log database changes: **1**-Enable logging of database changes to JET log files (default), **0**-Reset this option.

**[ Event=]{ 0| 1}**
Indicates the detail of events: **1**-Log detailed events to the system event log , **0**-Reset this option (default).

**Remarks**

- The WINS database engine can log the database changes to separate files. Because this option improves the performance of database operations, it is turned on by default.

- The WINS server can log detailed event information in the system event log. Because this option is resource-intensive, it is turned off by default.

**Examples**
In the following example, this command sets logging parameters for this WINS server to enable the logging of database changes to JET files and to log detailed events to the system event log.

**set logparam Event=1**

**set migrateflag**
Sets the migration flag for the specified WINS server.

**Syntax**
**set migrateflag** [**State=**]{**0** | **1**}

**Parameters**
**[ State=]{ 0| 1}**
Required. Indicates the setting of the migrate flag: **0**-Disable (migrate off), **1**-Enable (migrate on)

**Remarks**

- The migrate flag determines if unique static mappings can be overwritten by dynamic records.

**Examples**
In the following example, this command sets the migrate flag to permit the overwriting of unique static mappings by dynamic records for the specified WINS server.

**set migrateflag State=1**

**set namerecord**
Sets the intervals that are used to determine when WINS name records are renewed, deleted, and verified at the specified WINS server. Used without parameters, current values are maintained.

**Syntax**
**set namerecord** [[**Renew=**]*Value*] [[**Extinction=**]*Value*] [[**ExTimeout=**]*Value*] [[**Verification=**]*Value*]

**Parameters**
**[ Renew=] *Value***
Specifies the interval renewal rate, in seconds.

**[ Extinction=] *Value***
Specifies the extinction interval, in seconds.

**[ ExTimeout=] *Value***
Specifies the extinction time-out, in seconds.

**[ Verification=] *Value***
Specifies the verification interval, in seconds.

**Remarks**

- All parameters are optional. For parameters not used, previously set or default values are maintained.

**Examples**
In the following example, this command sets each of the values for name records on this WINS server.

**set namerecord Renew=3600 Extinction=7200 ExTimeout=86400 Verification=2073600**

**set periodicdbchecking**
Sets the periodic database consistency checking parameters for the specified WINS server. Used without parameters, current values are maintained.

**Syntax**
**set periodicdbchecking** [[**State=**]{**0** | **1**}] [[**MaxRec=**]*Value*] [[**CheckAgainst=**]{**0** | **1**}] [[**CheckEvery=**]*Value*] [[**Start=**]*Value*]

**Parameters**
**[ State=]{ 0| 1}**
Required. Indicates whether to check the state of the WINS database: **0**-Disable, **1**-Enable(default). If you set the value of this parameter to 0, all checks are disabled regardless of the values of other parameters.

**[ MaxRec=] *Value***
Specifies the maximum number of records that are checked each period. If unspecified, the default is 30,000.

**[ CheckAgainst=]{ 0| 1}**
Indicates the server to check against: **0**-Owner server (default), **1**-Randomly selected server.

**[ CheckEvery=] *Value***
Specifies how often the WINS database should be checked for consistency, in hours. If unspecified, the default is 24 hours.

**[ Start=] *Value***
Specifies the time, in seconds, at which to start the consistency check.

**Remarks**

- In order to set periodic consistency check intervals, it is recommended that you specify values for the **CheckEvery** and **Start** parameters. In order to adjust existing periodic consistency check settings, specify values only for **State** and the parameter or parameters that you want to change. Parameters for which you do not supply a value remain unchanged.

**Examples**
In the following example, this command checks the WINS database against a random replication partner every hour starting at 3,600 seconds.

**set periodicdbchecking State=1 MaxRec=1000 CheckAgainst=1 CheckEvery=1 Start=3600**

**set pgmode**
Sets the Persona Grata mode or the Persona Non Grata mode.

**Syntax**
**set pgmode** [**Mode=**]{**0** | **1**}

**Parameters**
**[ Mode=]{ 0| 1}**
Required. Specifies which mode to set: **0**-Sets the Persona Non Grata (PNG) mode, **1**-Sets the Persona Grata (PG) mode.

**Remarks**

- Use the **add pgserver** and the **delete pgserver** commands to update the Persona Grata list.

**Examples**
In the following example, this command sets the Persona Grata mode. During replication, WINS will accept only the records that are owned by the servers that are specified in the Persona Grata list.

**set pgmode Mode=1**

**set pullparam**
Sets the default pull partner parameters for the specified WINS server. Used without parameters, previously set or default values are maintained.

**Syntax**
**set pullparam** [[**State=**]{**0** | **1**}] [[**Strtup=**]{**0** | **1**}] [[**Start=**]*Value*] [[**Interval=**]*Value*] [[**Retry=**]*Value*]

**Parameters**
**[ State=]{ 0| 1}**
Required. Indicates whether you want persistent connections to be used with pull partners: **0**-Do not use persistent connections with pull partners, **1**-Use persistent connections with pull partners.

**[ Strtup=]{ 0| 1}**
Indicates whether pull replication should start at startup: **0**-Do not start pull replication at startup (default), **1**-Start pull replication at startup.

**[ Start=] *Value***
Specifies the start time for the pull trigger, in seconds.

**[ Interval=] *Value***
Specifies the replication interval, in seconds.

**[ Retry=] *Value***
Specifies the number of times to retry replication.

**Remarks**

- If you do not specify a value for **Interval=**, only one replication occurs.

- If you do not specify a value for **Start=**, execution of the command starts the pull trigger.

**Examples**
In the following example, this command sets the default pull partner parameters for this WINS server to use a persistent connection with its pull partners, begin replication at startup, start the pull trigger in 23,600 seconds at an interval of 259,200 seconds, and retry replication seven times.

**set pullparam State=1 Strtup=1 Start=23600 Interval=259200 Retry=7**

**set pullpartnerconfig**
Sets the configuration parameters for the specified pull partner.

**Syntax**
**set pullpartnerconfig** [[**State=**]{**0** | **1**}] [**Server=**]*ServerName* [[**Start=**]*Value*] [[**Interval=**]*Value*]

**Parameters**
**[ State=]{ 0| 1}**
Indicates the state of the persistent connection flag: **0**-Disable, **1**-Enable (default).

**[ Server=] *ServerName***
Required. Specifies the pull replication partner by server name or IP address.

**[ Start=] *Value***
Specifies the start time for replication, in seconds.

**[ Interval=] *Value***
Specifies the interval between pull replications, in seconds.

**Remarks**

- If you provide a server name for **Server=** that cannot be resolved to an IP address, the function fails.

**Examples**
In the following example, this command sets the pull partner configuration for the current WINS server for the partner WINS server 10.2.2.2. Pull replication starts in 3,600 seconds and runs every 4,320 seconds.

**set pullpartnerconfig State=1 Server=10.2.2.2 Start=3600 Interval=4320**

**set pushparam**
Sets the default push partner parameters for the specified WINS server.

**Syntax**
**set pushparam** [**State=**]{**0** | **1**} [[**Strtup=**]{**0** | **1**}] [[**AddChange=**]{**0** | **1**}] [[**Update=**]*Value*]

**Parameters**
**[ State=]{ 0| 1}**
Required. Indicates whether you want persistent connections to be used with push partners: **0**-Do not use persistent connections with push partners, **1**-Use persistent connections with push partners.

**[ Strtup=]{ 0| 1}**
Indicates whether you want push replication to start at startup: **0**-Do not start push replication at startup (default), **1**-Start push replication at startup.

**[ AddChange=]{ 0| 1}**
Indicates whether you want replication to be triggered upon address change: **0**-Do not trigger replication upon address change (default), **1**-Trigger replication upon address change.

**[ Update=] *Value***
Specifies the number of changes in version ID that will occur before replication.

**Examples**
In the following example, this command sets the default push partner parameters for this WINS server to use persistent connections with its push partners, start push replication at startup and when addresses change, and push to partners after the version ID has been updated twice for records owned at this WINS server.

**set pushparam State=1 Strtup=1 AddChange=1 Update=2**

**set pushpartnerconfig**
Sets the configuration parameters for the specified push partner.

**Syntax**
**set pushpartnerconfig** [[**State=**]{**0** | **1**}] [**Server=**]*****ServerName* [**Update=**]*Value*

**Parameters**
**[ State=]{ 0| 1}**
Indicates the state of the persistent connection flag: **0**-Disable, **1**-Enable (default).

**[ Server=] *ServerName***
Required. Specifies the push partner server for which you want to set configuration parameters. Identifies the server by server name or IP address.

**[ Update=] *Value***
Required. Specifies the threshold for updates.

**Remarks**

- If you provide a name in the server parameter cannot be resolved to an IP address, the function fails.

**Examples**
In the following example, this command sets the push partner configuration for the current WINS server for a partner WINS server with IP address 10.2.2.2 and sets two as the limit of changes in version ID allowed before push partners are informed about them.

**set pushpartnerconfig State=1 Server=10.2.2.2 Update=2**

**set replicateflag**
Sets the replication flag for the specified WINS server.

**Syntax**
**set replicateflag** [**State=**]{**0** | **1**}

**Parameters**
**[ State=]{ 0| 1}**
Required. Indicates the setting of the replication flag: **0**-Disable (replication flag off), **1**-Enable (replication flag on).

**Examples**
In the following example, this command enables the replication flag for this server.

**set replicateflag State=1**

**set startversion**
Sets the value of the beginning (start) version ID number for the database.

**Syntax**
**set startversion** [**Version=**]**{***High***,***Low***}**

**Parameters**
**[ Version=] {*High*,*Low*}**
Required. Specifies the start version ID number for the WINS database. High and low values are separated by a comma and enclosed by braces ({}).

**Examples**
In the following example, this command sets the beginning version ID number for the WINS database to 0.

**set startversion Version={0,0}**

**show**
**show browser**
Displays all active domain master browser [1Bh] records for the specified WINS server.

**Syntax**
**show browser**

**Parameters**
none

**show database**
Displays the database and records for a partial or complete list of specified owner servers.

**Syntax**
**show database** [**Servers=**]**{***IPAddresses***}** [[**RecType=**]{**0** | **1** | **2**}] [[**Count=**]*Count*] [[**Start=**]{**0** | **1**}] [[**EndChar=**]*16thCharInHex*] [[**File=**]*FileName*]

**Parameters**
**[ Servers=] {*IPAddresses*}**

Required. Specifies a list of the owner servers whose records are to be displayed. Identifies owner servers by IP addresses, which must be separated by commas and enclosed in braces ({}).

**[ RecType=]{ 0| 1| 2}**
Indicates the type of record: **0**-Both (default), **1**-Static, **2**-Dynamic.

**[ Count=]** ***Count***
Specifies the number of records that you want to display. If unspecified, all records that match the query parameters are be returned.

**[ Start=]{ 0| 1}**
Indicates where to start sorting the display: **0**-Start from the beginning of the database (default), **1**-Start from the end of the database.

**[ EndChar=]** ***16thCharInHex***
Indicates the 16th character of the records, in hexadecimal. This number filters the records that are retrieved from the database. If unspecified, all record types are returned.

**[ File=]** ***FileName***
Specifies the file in which formatted output is stored. If unspecified, output is displayed in the command prompt window.

**Remarks**

- If you specify an empty set ({}) for **Servers=**, **show database** records are displayed for all owner servers.

**Examples**
In the following example, this command outputs the first 1000 dynamic records in the WINS database for the server 10.2.2.2 to the text file output.txt.

**show database Servers={10.2.2.2} RecType=2 Count=1000 Start=0 Endchar=0 File=C:\WINSfiles\output.txt**

**show info**
Displays configuration information for the specified WINS server.

**Syntax**
**show info**

**Parameters**
none

**Remarks**

- If set, this command displays the following WINS configuration information for the current server:
    - Backup path
    - Name record setting
    - Periodic database consistency checking
    - Burst handling
    - Logging parameters
    - Start version count

**show name**
Queries and displays detailed information for a specified record in the current WINS server database.

**Syntax**
**show name** [**Name=**]*RecordName* [**EndChar=**]*16thCharInHex* [[**Scope=**]*ScopeName*]

**Parameters**
**[ Name=]** ***RecordName***

Required. Specifies the record for which to query.

**[ EndChar=]** ***16thCharInHex***
Required for all queries except unique type records. Indicates the 16th character in hexadecimal.

**[ Scope=]** ***ScopeName*****]**
Specifies the NetBIOS scope, if used.

**Examples**
In the following example, this command shows currently stored information for name record HOSTA, of File Server type, [20h] that exists in the WINS database on the specified WINS server.

**show name Name=HOSTA EndChar=20**

**show partner**
Displays all pull partners, push partners, or pull and push partners for the specified WINS server. When used without parameters, this command displays all push partners, pull partners, and push/pull partners for the specified WINS server.

**Syntax**
**show partner** [[**Type=**]{**0** | **1** | **2** | **3**}]

**Parameters**
**[ Type=]{ 0| 1| 2| 3}**
Indicates the type of partner to display: **0**-All (default), **1**-Pull, **2**-Push, **3**-Both push and pull.

**Examples**
In the following example, this command displays the server name, IP address, and replication type of each replication partner for this WINS server.

**show partner**

**show partnerproperties**
Displays default partner configuration information for the specified WINS server.

**Syntax**
**show partnerproperties**

**Parameters**
none

**Remarks**

- This command displays the auto partner configuration and default push/pull parameters for the current WINS server.

**show pullpartnerconfig**
Displays configuration information for a pull partner.

**Syntax**
**show pullpartnerconfig** [**Server=**]{*Name* | *IPAddress*}

**Parameters**
**[ Server=]{ *Name*| *IPAddress*}**
Required. Specifies the pull partner for which you want to display configuration information. Identifies the partner by name or IP address.

**Remarks**

- If you provide a name that cannot be resolved to an IP address, the function fails.

**Examples**
In the following example, this command displays the pull partner configuration information for WINS server 10.2.2.2.

**show pullpartnerconfig Server=10.2.2.2**

**show pushpartnerconfig**
Displays configuration information for a push partner.

**Syntax**

**show pushpartnerconfig** [**Server=**]{*Name* | *IPAddress*}

**Parameters**

**[ Server=]{ *Name*| *IPAddress}***

Required. Specifies the push partner for which you want to display configuration information. Identifies the push partner by name or IP address.

**Remarks**

- If you provide a name that cannot be resolved to an IP address, the function fails.

**Examples**

In the following example, this command displays the push partner configuration information for WINS server 10.2.2.2.

**show pushpartnerconfig Server=10.2.2.2**

**show recbyversion**

Displays a set of records for an owner, specified as a range of version IDs, and sorts the records by version.

**Syntax**

**show recbyversion** [**Server=**]*IPAddress* [**MaxVer=**]**{***High*,*Low***}** [**MinVer=**]**{***High*,*Low***}** [[**Name=**]*RecordName*] [[**EndChar=**]*16thChar*] [**Case=**]{**0** | **1**}

**Parameters**

**[ Server=] *IPAddress***

Required. Specifies the owner server for which you want to display a set of records. Identifies the server by IP address.

**[ MaxVer=] {*High*,*Low*}**

Required. Specifies the range of records to return by the maximum high and the maximum low version number. High and low values are separated by a comma and enclosed by braces ({}).

**[ MinVer=] {*High*,*Low*}**

Required. Specifies the range of records to return by minimum high and the minimum low version number. High and low values are separated by a comma and enclosed by braces ({}).

**[ Name=] *RecordName***

Specifies the records for which to look.

**[ EndChar=] *16thChar***

Indicates the character for which to look, in hexadecimal.

**[ Case=]{ 0| 1}**

Indicates whether to create a case-sensitive comparison: **0**-Ignore case (default), **1**-Create a case-sensitive comparison.

**Remarks**

- If **Name=** is not specified, **EndChar=** and **Case=** are ignored.
- If you set both **MinVer=** and **MaxVer=** to {0,0}, **show recbyversion** displays the entire range of records.

**Examples**

In the following example, this command displays the total number of records (by record version) owned by WINS server 10.2.2.2.

**show recbyversion Server=10.2.2.2 MaxVer={0,0} MinVer={0,0} Name=dc1 EndChar=1c Case=0**

**show reccount**

Displays the number of records owned by a specific WINS server.

**Syntax**

**show reccount** [**Server=**]*IPAddress* [**MaxVer=**]**{***High*,*Low***}** [**MinVer=**]**{***High*,*Low***}**

**Parameters**

**[ Server=] *IPAddress***

Required. Specifies the owner server for which you want to display the number of records. Identifies the server by IP address.

**[ MaxVer=] {*High*,*Low*}**
Required. Specifies the maximum high and the maximum low version of records to show. High and low values are separated by a comma and enclosed by braces ({}).

**[ MinVer=] {*High*,*Low*}**
Required. Specifies the minimum high and the minimum low version of records to show. High and low values are separated by a comma and enclosed by braces ({}).

**Remarks**

- If you set both **MinVer=** and **MaxVer=** to {0,0}, **show reccount** counts the entire range of records.

**Examples**
In the following example, this command shows the total record count for WINS server 10.2.2.2.

**show reccount Server=10.2.2.2 MaxVer={0,0} MinVer={0,0}**

**show server**
Displays information for the specified WINS server.

**Syntax**
**show server**

**Parameters**
none

**Remarks**

- This command displays the name and IP address of the specified WINS server.

**show serverstatistics**
Displays statistics for the specified WINS server.

**Syntax**
**show statistics**

**Parameters**
none

**show version**
Displays the current maximum version counter value for the specified WINS server.

**Syntax**
**show version**

**Parameters**
none

**show versionmap**
Displays the owner ID to maximum version number mappings table for the specified WINS server.

**Syntax**
**show versionmap**

**Parameters**
none

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |

| | |
|---|---|
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Netsh commands for AAAA
Netsh commands for DHCP
Netsh diagnostic (diag) commands
Netsh commands for Interface IP
Netsh commands for remote access (ras)
Command-line reference A-Z
Command shell overview

# Software Deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Software deployment**
This section covers:

- Group Policy
- Windows Update
- Automatic Updates
- Licensing
- Windows Installer
- Remote Installation Services
- Remote Administration
- Terminal Services

# Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy**

You can use Group Policy to manage features included with the Microsoft® Windows Server 2003 family, such as Group Policy Software Installation, Administrative Templates, Folder Redirection, Remote Installation Services, Security Settings, Scripts (Startup/Shutdown and Logon/Logoff) and Internet Explorer Maintenance.

To use Group Policy without the Group Policy Management Console (GPMC), see Group Policy (pre-GPMC).

To use Group Policy with the Group Policy Management Console, see Group Policy Management.

Microsoft makes available new procedures, and new and updated tools online. For the latest information about Group Policy and tools that you can use to manage Group Policy, see Group Policy at the Microsoft Web site.

# Windows Update

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Update**
Windows Update is the online extension of Windows. Use Windows Update to download items such as security fixes, critical updates, the latest Help files, drivers, and Internet products. New content is added to the site regularly, so you can always get the most recent updates and fixes to protect your computer and keep it running smoothly. For more information, see Windows Update at the Microsoft Web site. Also, see Using Windows Update.

With Windows Update Catalog, you can select updates that you plan to deploy later, which means you can control how and when the updates are deployed. For more information, see Windows Update Catalog at the Microsoft Web site.

With Automatic Updates, which is built into Windows XP and products in the Windows Server 2003 family, the operating system can automatically download and install critical updates from the Windows Update site. For more information, see Automatic Updates.

With Microsoft Windows Update Services, you can deploy a limited version of the Windows Update site inside your firewall. You can control how and when updates are deployed, and deliver updates to users through your intranet rather than having updates occur directly from the Internet. For more information about Windows Update Services, see Evaluate Windows Update Services Now at the Microsoft Web site.

**Note**

- You must be logged on as an administrator or as a member of the Administrators group to update system files, including Help files.

# Automatic Updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Automatic Updates**
Windows can keep your computer up-to-date automatically with the latest updates and enhancements. You do not need to search for critical updates and information because Windows delivers them directly to your computer. Windows recognizes when you are online and uses your Internet connection to search for downloads from the Windows Update Web site. An icon appears in the notification area each time new updates are available.

**Difference between Windows Update and Automatic Updates**
Windows Update and Automatic Updates are two separate components designed to work together to keep Windows secure. Windows Update is Microsoft-owned Web site from which Windows users can download critical and non-critical software. Automatic Updates enables you to automatically interact with the Windows Update Web site to obtain the critical software updates. These items include security fixes, critical updates, and critical drivers. You have full control over the level of this interaction. For example, you could choose to automatically download and install updates as they become available, or simply be notified when new updates become available.

- For tips about using Automatic Updates, see Best practices for Automatic Updates.
- For help with specific tasks, see Automatic Updates How to....
- For general background information, see Automatic Updates Concepts.

# Best practices for Automatic Updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for Automatic Updates**

**Use a schedule to keep Windows automatically updated**
Scheduled updates keep your operating system up-to-date, and minimize the time you need to spend logged onto your computer as a member of the Administrators group. By minimizing the time you spend logged onto your computer as a member of the Administrators group, you keep your computer more secure.

**See Also**
**Concepts**
Schedule automatic updates

# Automatic Updates How to...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automatic Updates How to...**


Turn on Automatic Updates

Schedule automatic updates

Change how updates are downloaded and installed

Download updates that apply to your computer

Review and install downloaded updates

View a list of updates installed on your computer

Check if there are new updates for your computer

Stop Windows from notifying you about updates that you have already decided not to download or install (hide updates)

Restore hidden updates

Uninstall an update

# Turn on Automatic Updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Turn on Automatic Updates**

You must be logged on as a computer administrator to complete this procedure.

1. Open **System**, and then click the **Automatic Updates** tab.

   – or –

   If you are running Windows 2000, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Automatic Updates**.

2. Click **Automatic (recommended)**.

3. Under **Automatically download recommended updates for my computer and install them**, select the day and time you want Windows to install updates.

## Notes

- To open System, click **Start**, click **Control Panel**, click **Performance and Maintenance**, and then click **System**.

- Automatic Updates provides high-priority updates, which include security and other critical updates that can help protect your computer. It's a good idea to visit the Windows Update Web site (http://www.microsoft.com/) on a regular basis to get optional updates, such as recommended software and hardware updates, that can help improve your computer's performance.

**See Also**

**Concepts**

Options for setting up Automatic Updates on your computer
Why use scheduled updates to help protect your computer?
Windows Update privacy statement

# Schedule automatic updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Schedule automatic updates**

You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

1. Open **System**, and then click the **Automatic Updates** tab.

   – or –

   If you are running Windows 2000, click **Start**, point to **Settings**, click Control Panel, and then double-click **Automatic Updates**.

2. Click **Automatic (recommended)**.

3. Under **Automatically download recommended updates for my computer and install them**, set the time and day you want Windows to install updates, and then click **Apply**.

**Notes**

- To open System, click **Start**, click **Control Panel**, click **Performance and Maintenance**, and then click **System**.

- Scheduled updates are automatically installed at the scheduled time, regardless of who is logged on to the computer. If you choose to be notified before updates can be installed, you must be logged on as an administrator to get notifications and to install updates.

- Some updates require you to restart your computer before they can take effect. Windows will notify any user who is logged on to the computer before restarting; however, only members of the local Administrators group can stop the computer from restarting. Make sure you save your work before restarting your computer and remind other users to save their work, especially before scheduled installation times.

# Change how updates are downloaded and installed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Change how updates are downloaded and installed**
You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

1. Open System, and then click the **Automatic Updates** tab.

   – or–

   If you are running Windows 2000, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Automatic Updates**.

2. Select a setting, and then click **Apply**.

You can select from the following options:

- Automatic (recommended)
- Download updates for me, but let me choose when to install them
- Notify me but don't automatically download or install updates
- Turn off Automatic Updates

**Automatic (recommended)**
When you are connected to the Internet, Windows finds and downloads updates in the background—you are not notified or interrupted during this process and the updates do not interfere with other downloads. If you do not change the default schedule, updates that have been downloaded to your computer will be installed at 3 A.M.

If your computer is turned off during a scheduled update, Windows will install the updates the next time you start your computer. If you need to help complete the installation process, Windows will notify you. For example, you might need to accept an End User License Agreement (EULA) before some updates can be installed. If you need to restart your computer for an update to take effect, Windows will notify you but will also restart your computer at the scheduled time.

**Download updates for me, but let me choose when to install them**
With this option, updates will be downloaded to your computer and Windows will alert you that they are ready to be installed. To receive alerts, you must be a member of the Administrators group for your computer. When you are connected to the Internet, Windows finds and downloads updates in the background—you are not notified or interrupted during this process and the updates do not interfere with other downloads.

After downloading is complete, the **Windows Update** icon appears in the notification area and an alert pops up, letting you know that the updates are ready to be installed. To review and install available updates, click the icon or the alert. You can install all or some of the available updates.

**Notify me but don't automatically download or install updates**
To manually download and install updates, you must be a member of the Administrators group for your computer. Windows checks for important updates and notifies you if any are available but they are not delivered or installed on your computer until you choose to do so. When Windows finds updates for your computer, the **Windows Update** icon appears in the notification area and an alert pops up, letting you know that updates are ready to be downloaded. After you click the icon or the alert, you can select some or all of the updates to download.

Windows then downloads the updates in the background—you aren't notified or interrupted during this process and the updates do not interfere with other downloads. When downloading is complete, the **Windows Update** icon appears in the notification area again, this time to let you know that the updates are ready to be installed. You can choose to install all or some of the available updates.

**Turn off Automatic Updates**
You will never be notified when important updates are available for your computer, and you will never be asked to download or install them. This means that your computer can be vulnerable to security threats and harmful viruses that can damage your computer or your files. Viruses can also spread over the Internet to other people with whom you exchange e-mail, share files, or work with on a network.

New viruses and security threats are continually developed by attackers, so helping protect your computer is an ongoing process. If you don't turn on Automatic Updates, we recommend that you regularly install updates from the Windows Update Web site (http://www.microsoft.com/).

### Notes

- To open System, click **Start**, click **Control Panel**, click **Performance and Maintenance**, and then click **System**.

- Only users with administrator privileges may add or remove programs, including Windows updates. It is strongly recommended that you log out of the computer administrator account when you are not performing tasks that require administrator privileges. If you are logged on as an administrator when your computer is the target of a virus or malicious user, the attack can cause extensive damage. For example, it might be able to reformat your hard drive, delete all your files, or create a new administrator account so the attacker can take over your computer. For more information about user accounts and why you should not run your computer as an administrator, see Help and Support.

# Download updates that apply to your computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Download updates that apply to your computer**
Based on your Automatic Updates settings, Windows can notify you when updates that apply to your computer are ready to download. You can review each update and choose whether to download it. To help protect your computer, we recommend that you download and install all updates that apply to your computer as soon as they are available.

1. When you are notified that new updates are ready to be downloaded, click the **Windows Update** icon that appears in the notification area.

2. By default, all updates are selected to be downloaded. If you do not want Windows to download an update, clear the check box beside it.

3. Click **Download**.

### Notes

- If you are using the **Automatic (recommended)** setting, Windows automatically downloads important updates as soon as they are available. You are not notified or interrupted during the process.

- Windows Update software might automatically be updated prior to applying any other updates.

**See Also**
**Concepts**
Review and install downloaded updates
Stop Windows from notifying you about updates that you have already decided not to download or install (hide updates)
Turn on Automatic Updates

# Review and install downloaded updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Review and install downloaded updates**

1. When the **Windows Update** icon appears in the notification area, or when you are notified that updates are ready to install, click the icon.

2. To install all available updates, click **Express install (Recommended)**, or, to choose which downloads to install, click **Custom install (Advanced)**. By default, all updates are selected to be installed. If you do not want Windows to download an update, clear the check box next to it.

Windows notifies you when the updates are installed. Windows will also notify you if you need to accept an End User License Agreement (EULA) or restart your computer to complete the installation process.

### Notes

- Windows might prompt you to restart your computer after installing certain updates. Please restart your computer when prompted; otherwise, your computer might not work properly.

- If you are using the **Automatic (recommended)** setting, Windows automatically installs important updates on the schedule that you selected. You are not notified or interrupted during the process unless you need to accept an End User License Agreement (EULA) or restart your computer.

- Windows Update software might automatically be updated prior to applying any other updates.

**See Also**
**Concepts**
Options for setting up Automatic Updates on your computer
Review and install downloaded updates
Turn on Automatic Updates

# View a list of updates installed on your computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**View a list of updates installed on your computer**

1. Open **Add or Remove Programs**.
2. On the top right, select the **Show updates** check box. All installed updates appear at the top of your program list.

**Notes**

- To open Add or Remove Programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove Programs**.
- You can also view your installation history at the Windows Update Web site (http://www.microsoft.com/).

# Check if there are new updates for your computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Check if there are new updates for your computer**
If you turned on Automatic Updates, Windows automatically checks for important updates for your computer. These updates can include critical updates, security updates, and service packs. Windows notifies you about (or automatically installs) important updates, based on your Automatic Updates settings.

It is also a good idea to routinely check the Windows Update Web site to see if other types of updates, such as driver updates, are available. You must be logged on as a computer administrator to complete this procedure. You must also be connected to the Internet.

1. Go to the Windows Update Web site (http://www.microsoft.com/).

2. Click **Express Install**.

If updates are found, we recommend that you first install all high-priority updates, and then install all other updates such as software and hardware updates.

### Notes

- Some updates require you to accept an End User License Agreement (EULA) before you can install them. If you decline, these updates are not installed. If you change your mind, you can install them later by visiting the Windows Update Web site (http://www.microsoft.com/).

- Some updates require you to restart your computer before the updates can take effect. Windows will notify any user who is logged on to the computer before restarting; however, only members of the local Administrators group can delay the restart. Make sure you save your work before restarting your computer and remind other users to save their work, especially before scheduled installation times.

**See Also**
**Concepts**
Options for setting up Automatic Updates on your computer
Turn on Automatic Updates

# Stop Windows from notifying you about updates that you have already decided not to download or install (hide updates)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Stop Windows from notifying you about updates that you have already decided not to download or install (hide updates)**
You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure. You must also be connected to the Internet.

1. When the **Windows Update** icon appears in the notification area, or when you are notified that updates are ready to download, click the icon.

2. Review the recommended updates. If you do not want to download an item, clear the check box next to it.

3. Click **Download**. (If you cleared all the check boxes, click **Close**.)

4. In the confirmation dialog box, select the **Don't notify me about this update again** check box, and then click **OK**.

### Note

- Hiding an update means that Windows will no longer try to download it when checking for updates that apply to your computer. If you change your mind and want to download a hidden update later, go to the Windows Update Web site (http://www.microsoft.com/), click **Custom Install**, and then click **Restore hidden updates**.

**See Also**
**Concepts**
Automatic Updates: Frequently Asked Questions
Restore hidden updates

# Restore hidden updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restore hidden updates**

You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

1. Open System, and then click the **Automatic Updates** tab.

   – or –

   If you are running Windows 2000, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Automatic Updates**.

2. Click **Restore hidden updates**.

The next time you connect to the Internet, Windows will check for any high-priority updates for your computer, including updates that you previously declined. Windows will either download the updates automatically or notify you that they are available, depending on your Automatic Updates setting.

Alternately, you can restore hidden updates by visiting the Windows Update Web site.

1. Go to the Windows Update Web site (http://www.microsoft.com/).

2. Click **Custom Install**.

3. In the left pane, click **Restore Hidden Updates**, and then select the updates you want to install.

4. Click **Install**.

**Notes**

- To open System, click **Start**, click **Control Panel**, click **Performance and Maintenanc**e, and then click **System**.

- When you restore hidden updates, Windows determines which updates still apply to your computer. If you cannot find a specific hidden update, you might have installed a more recent update that replaced the original.

- Some updates require you to accept an End User License Agreement (EULA) before you can install them. If you decline the agreement, the updates are not installed. If you change your mind, you can install the updates later by visiting the Windows Update Web site (http://www.microsoft.com/).

- Some updates require you to restart your computer before they take effect. Windows will notify any user who is logged on to the computer before restarting; however, only members of the local Administrators group can stop the computer from restarting. Make sure you save your work before restarting your computer and remind other users to save their work, especially before scheduled installation times.

# Uninstall an update

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Uninstall an update**

You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

1. Open Add or Remove Programs.
2. On the top right, select the **Show updates** check box. All installed updates appear at the top of your program list.
3. Click the update you want to uninstall, and then click **Remove**.

**Note**

- To open Add or Remove Programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove Programs**.

# Automatic Updates Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Automatic Updates.

- Automatic Updates overview
- How does Automatic Updates work?
- Why use scheduled updates to help protect your computer?
- Automatic Updates: Frequently Asked Questions
- Options for setting up Automatic Updates on your computer
- Windows Update privacy statement

# Automatic Updates overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automatic Updates overview**

With Automatic Updates, Windows routinely checks for updates that can help protect your computer against the latest viruses and other security threats. These high-priority updates are available through the Windows Update Web site and include security updates, critical updates, or service packs.

When you turn on Automatic Updates, you don't have to search for updates online or worry that critical fixes might be missing. Windows automatically downloads and installs them for you, using a schedule that you determine. If you prefer to download and install updates yourself, you can also set up Automatic Updates to notify you whenever any high-priority updates become available.

**See Also**

**Concepts**

Automatic Updates: Frequently Asked Questions

Options for setting up Automatic Updates on your computer

# How does Automatic Updates work?

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How does Automatic Updates work?**
When you turn on Automatic Updates, Windows routinely checks the Windows Update Web site for high-priority updates that can help protect your computer from the latest viruses and other security threats. These updates can include security updates, critical updates, and service packs. Depending on the setting you choose, Windows automatically downloads and installs any high-priority updates that your computer needs, or notifies you as these updates become available.

How are updates downloaded?

How are updates installed?

**How are updates downloaded?**
When you connect to the Internet, Windows sends data to the Windows Update Web site about the way your computer is set up so the service can determine which updates your computer needs. Windows does not use your name, address, e-mail address, or any information that can be used to identify you or contact you. For details, see the Windows Update privacy statement.

It makes no difference whether you use a dial-up or broadband connection; important updates are downloaded behind the scenes and do not interfere with other downloads. If you disconnect from the Internet before updates are fully downloaded, nothing is lost. The download process will continue the next time you connect to the Internet.

**How are updates installed?**
You do not have to be connected to the Internet for Windows to install new updates. But updates must be installed, not just downloaded, before they can help protect your computer. If you use the **Automatic (recommended)** setting, new updates are installed at 3 A.M. However, you can change the time or frequency of scheduled updates to best suit your needs.

If your computer is turned off during a scheduled update, updates are installed the next time you start your computer. You can also choose to have Windows notify you when new updates are available, and then install them yourself.

You might be asked to accept an End User License Agreement (EULA) before certain updates can be installed. Other updates might require you to restart your computer before the installation process can be completed. If you are an administrator for your computer, you can delay the restart. Otherwise, Windows notifies you and then restarts your computer for you so the updates can begin to help protect your computer. It is always a good idea to save your work frequently, and to remind other users of your computer to save their work, especially before scheduled installation times.

### Note

- Installing updates before you shut down your computer is another way to keep your computer up to date and more secure. This option is available only in Microsoft Windows XP Service Pack 2 (SP2) and only if important updates have been downloaded but not yet installed. Do not turn off or unplug your computer while updates are installing. Windows will automatically turn off your computer after the updates are installed.

**See Also**
**Concepts**
Automatic Updates: Frequently Asked Questions
Options for setting up Automatic Updates on your computer
Turn on Automatic Updates

# Why use scheduled updates to help protect your computer?

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Why use scheduled updates to help protect your computer?**

If you have not turned on Automatic Updates, your computer is more vulnerable to viruses and other security threats. When you turn on Automatic Updates, Windows routinely checks the Windows Update Web site for high-priority updates that can help protect your computer against attacks. High-priority updates include security updates, critical updates, and service packs.

Scheduled updates offer:

**Convenience**

You don't have to remember to visit the Windows Update Web site or other online security sites to get updates. And you don't have to choose which updates are most important for protecting your computer—Automatic Updates looks for security updates, critical updates, and service packs, and installs them on the schedule that you set.

**Reliability**

Updates are downloaded behind the scenes whenever you're connected to the Internet. The downloading process doesn't interfere with other downloads or interrupt you while you work. If you disconnect from the Internet before updates are fully downloaded, the download process will continue the next time you connect to the Internet. Don't forget, however, that updates must be installed to take effect. When you schedule updates, Windows automatically installs them for you, unless you choose to be notified and then install them yourself.

**Up-to-date software**

You set the schedule for Windows to install new updates. This means that Windows is routinely checking for and installing any important updates that your computer needs. Even if you choose to be notified and install updates yourself, you still learn about important updates as they are released. By scheduling updates, you can quickly and easily keep your computer up to date—without thinking about it.

**See Also**

**Concepts**

Automatic Updates: Frequently Asked Questions
Automatic Updates How to...
Turn on Automatic Updates

# Automatic Updates: Frequently Asked Questions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automatic Updates: Frequently Asked Questions**


Why are updates important?

How does Automatic Updates work?

What if ...?


**Why are updates important?**


- What are updates?
- How can I keep my computer updated?
- Why should I use Automatic Updates?


**What are updates?**

Microsoft offers important updates, which include security and other critical updates, to help protect your computer against new viruses and other security threats that can spread over the Internet or a network. Other updates contain enhancements, such as upgrades and tools that can help your computer run more smoothly. What are all of the different types of Microsoft software updates?

**How can I keep my computer updated?**

**Turn on Automatic Updates**

Get high-priority updates, which include security and other critical updates, automatically delivered to your computer and installed on a schedule that you set. Turn on Automatic Updates and choose a setting by opening System and then clicking the **Automatic Updates** tab. (If you are running Windows 2000, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Automatic Updates**.)

**Visit the Windows Update Web site (http://www.microsoft.com/)**

Go here on a regular basis to get high priority and optional updates for your computer.

**Why should I use Automatic Updates?**

Turning on Automatic Updates is the easiest way to help protect your computer against new and ongoing attacks. When you turn on Automatic Updates, Windows regularly checks the Windows Update Web site for important updates that your computer might need. These updates can include security updates, critical updates, and service packs. Windows downloads updates behind the scenes when you are online—the updates do not interfere with other downloads and you aren't interrupted while you're working.

Windows automatically installs important updates for you, though in some cases you may be asked to accept an End User License Agreement (EULA) or restart your computer to finish the process. Updates are installed on a schedule that you set, so you don't have to search the Windows Update Web site for new updates or worry that something important might be missing. You can also set up Automatic Updates to notify you when updates are available so you can download and install updates yourself.

**How does Automatic Updates work?**


- How are updates downloaded?
- How are important updates installed?
- Are Automatic Updates the same updates as those that are available from the Windows Update Web site?
- How long will it take to download or install an update?
- Will I need to restart my computer?

- Will I need to accept an End User License Agreement (EULA)?
- How do I know which updates are already installed?
- How can I be sure that updates are safe to install?
- How much does it cost to use Automatic Updates?
- What type of information is sent to Microsoft during the update process?

**How are updates downloaded?**
When you connect to the Internet, Windows sends data about the way your computer is set up to the Windows Update Web site so the service can determine which updates your computer might need. Windows does not send your name, address, e-mail address, or any information that can be used to identify you or contact you. If updates are available, they are downloaded behind the scenes when you are online—they do not interfere with other downloads and you are not interrupted.

If you use a broadband connection, updates can be downloaded at any time, even if you are not actively using the Internet. If you use a dial-up connection, updates can be downloaded only when you are connected to the Internet. If you disconnect from the Internet before updates are fully downloaded, the download process will continue the next time you connect to the Internet.

Remember that updates must be installed before they can help protect your computer. Remember, too, that if you do not use the **Automatic (recommended)** setting, Windows can only remind you about updates if you are logged on as a member of the Administrators group. (Only administrators can add or remove programs, including updates.) To notify you, the **Windows Update** icon appears in the notification area and an alert appears, stating that new updates are ready to be downloaded.

**How are important updates installed?**
How updates are installed depends on the Automatic Update settings you choose. If you use the **Automatic (recommended)** setting, updates are installed at 3 A.M. However, you can change the time or frequency of scheduled updates to best suit your needs. If your computer is turned off during a scheduled update, updates are installed the next time you start your computer.

In some cases, an update will require you to accept an End User License Agreement (EULA). If so, Windows notifies you. If an update requires you to restart your computer before it can be installed, Windows notifies you to restart your computer for security maintenance. Make sure you save your work frequently and remind other users to save their work, especially before scheduled installation times.

Remember that if you do not use the **Automatic (recommended)** setting, Windows can only notify you about updates if you are logged on as a member of the Administrators group. (Only administrators can add or remove programs, including updates.) To notify you, the **Windows Update** icon appears in the notification area and an alert appears, stating that new updates are ready to be installed.

You do not have to be connected to the Internet for updates to be installed.

**Are Automatic Updates the same updates as those that are available from the Windows Update Web site?**
Yes. However, Automatic Updates delivers only the latest important updates that can help protect your computer against viruses and other security threats. These updates can include security updates, critical updates and service packs and should be installed as soon as they are released. You can download and install other types of updates, such as upgrades and tools that help improve the performance of your computer, by visiting the Windows Update Web site (http://www.microsoft.com/) on a regular basis.

**How long will it take to download or install an update?**
It depends on the type of Internet connection you have and the size and number of updates that are being downloaded or installed. You can view the progress of downloads by hovering over the **Windows Update** icon which will appear in the notification area. Remember that updates are downloaded behind the scenes so you are not interrupted. You can view more details about the installation progress by clicking the **Windows Update** icon while updates are being installed. If there are a large number of updates and the download process is expected to take a long time (due to infrequent or slow Internet connections), Windows will also notify you at least once every other day about updates that have successfully finished downloading.

**Will I need to restart my computer?**
It depends. Some updates cannot be installed or will not take effect until you restart your computer. For those, Windows notifies you to restart your computer for security maintenance. If you are an administrator for your computer, you can delay the restart; otherwise, Windows warns you and then restarts your computer for you. Make sure you save your work and remind other users to save their work, especially before scheduled installation times.

**Will I need to accept an End User License Agreement (EULA)?**
Sometimes. Some updates require you to accept a EULA before the updates can be installed. These updates can not be automatically installed so an alert and the Windows Update icon will appear in the notification area to let you know that updates are ready to be installed. Click the icon to view and accept the EULA for these updates and then install them.

**How do I know which updates are already installed?**
To see a list of updates already installed on your computer, go to the Windows Update Web site (http://www.microsoft.com/). In the left pane, under **Other Options**, click **View installation history**. In Microsoft Windows XP Service Pack 2 (SP2), you can also open **Add or Remove Programs** in Control Panel and select the **Show Updates** check box.

**How can I be sure that updates are safe to install?**
Windows uses Secure Socket Layer (SSL) to encrypt the transfer of system information and updates between your computer and the Windows Update Web site. Each file that you download using Automatic Updates has a digital signature from Microsoft. Digital signatures are designed to ensure the authenticity and integrity of signed files. Automatic Updates will not install files that do not contain the correct digital signature.

**How much does it cost to use Automatic Updates?**
There is no charge to you for this service. Updates from Microsoft are provided at no cost, as part of maintenance and support service for Windows users. Microsoft is committed to helping protect your computing environment, no matter what tasks you perform at your computer. Depending on your Internet connection, standard local and long distance phone charges and Internet service charges might apply.

**What type of information is sent to Microsoft during the update process?**
To provide you with the best possible service, Windows Update collects some system configuration data in order to provide you with updates that apply to your computer. This information includes:

- Computer make and model
- Windows version number
- Internet Explorer version number
- Version numbers of other Microsoft software for which updates might be available
- Plug and Play ID numbers of hardware devices
- Region and Language setting
- Globally Unique Identifier (GUID)
- Product ID and Product Key
- BIOS name, revision number, and revision date

Microsoft is committed to helping protect your privacy. Windows Update does not collect your name, address, e-mail address, or any other form of personally identifiable information. Your Internet Protocol (IP) address is logged when you connect to the Windows Update site, but this address will only be used to generate aggregate statistics. For more details on how the information is protected, see the Windows Update privacy statement. You can always find the privacy statement at the Windows Update Web site (http://www.microsoft.com/).

**What if ...?**

- I'm not an administrator for my computer?
- My Internet connection disconnects before an update is downloaded?
- My Automatic Updates options are grayed out?
- I hid an update but want to install it later?

- I can't hide an update?
- I can't uninstall an update?
- I want to uninstall some of the updates from a security package?
- I've restored my hidden updates but cannot find the one I want?
- I share my computer with other users—do we get the same updates?
- I have a home network—do all computers get the same updates?
- My computer is in hibernate or standby mode during a scheduled update?
- My computer is turned off during a scheduled update?
- I never turn off my computer?

**I'm not an administrator for my computer?**
To help protect your computer, use the **Automatic (recommended)** setting, which checks for new updates and automatically downloads and installs them—regardless of whether you or other users are administrators for your computer.

Only users with administrator privileges may add or remove programs, including Windows updates. If you are not an administrator for your computer, you need to contact someone who is so they can change your Automatic Updates settings to the recommended setting.

It is strongly recommended that you log out of the computer administrator account when you are not performing tasks that require administrator privileges. If you are logged on as an administrator when your computer is the target of a virus or malicious user, the attack can cause extensive damage. For example, it might be able to reformat your hard drive, delete all your files, or create a new administrator account so the attacker can take over your computer. For more information about user accounts and why you should not run your computer as an administrator, see Help and Support.

**My Internet connection disconnects before an update is downloaded?**
If you disconnect from the Internet before updates are fully downloaded, the download process will continue from where it was interrupted the next time you connect to the Internet. If you use a dial-up connection, we recommend that you try to stay connected long enough for Windows to download all applicable updates. You do not need to be connected to the Internet to install updates, just to download them.

**My Automatic Updates options are grayed out?**
If your computer is part of a network, your network administrator might have configured your Automatic Updates options for you. Network administrators often need to implement systemwide changes to help make sure that the network operates smoothly and more securely whenever changes, such as updates, are applied. Your options might also be grayed out if you are not an administrator for your computer. Only administrators can add or remove programs, including updates.

**I hid an update but want to install it later?**
You can install all hidden updates by opening **System**, and, on the **Automatic Updates** tab, **clicking Restore hidden updates**. (If you are running Windows 2000, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Automatic Updates**.) The next time you connect to the Internet, Windows will check for all applicable updates for your computer, old or new.

To choose which hidden updates you want to install, go to the Windows Update Web site (http://www.microsoft.com/) and click **Custom Install**. In the left pane, under **Other Options**, click **Restore hidden update**s and then select the updates you want to download and install.

**I can't hide an update?**
If your computer is part of a network, a network administrator has determined that the update must be installed for security maintenance. The update will be installed on the schedule that the network administrator sets. In rare cases, Microsoft will issue an update deemed critical to how the Windows Update Web site or Automatic Updates functions. These updates cannot be hidden, so you will always be notified that they are available to be downloaded or installed so you can continue using Automatic Updates.

**I can't uninstall an update?**

Only administrators can add or remove programs, including updates. If you are not an administrator for your computer, you must contact an administrator to manage updates. If your computer is part of a network and you uninstall an update, your network administrator may re-install it as part of systemwide security or maintenance policy. In rare cases, some updates can not be uninstalled.

**I want to uninstall some of the updates from a security package?**
Some security updates can contain multiple updates, which are downloaded and installed as a package. You cannot uninstall some of the updates within a security package; you need to uninstall the entire package.

**I've restored my hidden updates but cannot find the one I want?**
If you visit the Windows Update Web site (http://www.microsoft.com/) but cannot find the update in the **Restore hidden updates** list, the update no longer applies to your computer. Updates no longer apply when another update, such as a critical update, security update, service pack, or other type of update, includes the same solution for your computer and that later update has already been installed.

**I share my computer with other users—do we get the same updates?**
Yes. Updates apply to the computer, not to each user account. If you install an update, the update applies for all users of the computer. Similarly, if you decline or uninstall an update, it will not be available for other users.

**I have a home network—do all computers get the same updates?**
Not necessarily. If you use more than one computer, including laptops, in a home network, each computer must have access to the Internet and you must turn on Automatic Updates for each computer separately. Also, if any computers run different versions of Windows, different updates will apply.

It's a good idea to use the **Automatic (recommended)** setting for all your computers so you can keep Windows up to date and help protect your home network, no matter who uses the computers. Unless you use the **Automatic (recommended)** setting for all computers on your network, or set policies using a domain controller, you need to be sure that updates are downloaded and installed for each computer separately.

**My computer is in hibernate or standby mode during a scheduled update?**
Updates will not be installed at the scheduled time. Instead, updates will be installed the next time you start your computer. This means you need to shut down your computer, not let it hibernate or stand by, and then restart it before the updates can be installed.

If you are an administrator for your computer and you usually hibernate or suspend it, you can install available updates as soon as they become available. The **Windows Update** icon appears in the notification area whenever updates are fully downloaded and ready to install. If you like, you can click the icon and install new updates before the scheduled time.

**My computer is turned off during a scheduled update?**
Any updates that have been downloaded are not installed. Instead, updates are installed the next time you start your computer. If you are an administrator for your computer, the **Windows Update** icon appears in the notification area whenever updates are ready to install. If you like, you can click the icon and install new updates before the scheduled time.

**I never turn off my computer?**
If you use the **Automatic (recommended)** setting, or scheduled updates, Windows will download new updates whenever you are connected to the Internet. Updates are installed at the scheduled time, as long as your computer is turned on. If any updates require you to accept an End User License Agreement (EULA), Windows will notify you and the updates will not be installed until you view and accept the agreement.

Some updates cannot be installed unless your computer is restarted. For these, Windows notifies you that your computer needs to be restarted for security maintenance. If you are an administrator for your computer, you can delay the restart. Otherwise, Windows will warn you and then restart your computer so the updates can begin to help protect your computer.

# Options for setting up Automatic Updates on your computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Options for setting up Automatic Updates on your computer**
To choose when and how updates will be delivered to your computer, you have four options:

- Automatic (recommended)
- Download updates for me, but let me choose when to install them
- Notify me but don't automatically download or install updates
- Turn off Automatic Updates

**Automatic (recommended)**
When you are connected to the Internet, Windows finds and downloads updates in the background—you are not notified or interrupted during this process and the updates do not interfere with other downloads. If you do not change the default schedule, updates that have been downloaded to your computer will be installed at 3 A.M.

If your computer is turned off during a scheduled update, Windows will install the updates the next time you start your computer. If you need to help complete the installation process, Windows will notify you. For example, you might need to accept an End User License Agreement (EULA) before some updates can be installed. If you need to restart your computer for an update to take effect, Windows will notify you and will restart your computer at the scheduled time.

**Download updates for me, but let me choose when to install them**
To receive alerts, you must be a member of the Administrators group for your computer. When you are connected to the Internet, Windows finds and downloads updates in the background—you are not notified or interrupted during this process and the updates do not interfere with other downloads.

After downloading is complete, the **Windows Update** icon appears in the notification area and an alert pops up, letting you know that the updates are ready to be installed. To review and install available updates, click the icon or the alert. You can install all or some of the available updates.

**Notify me but don't automatically download or install updates**
To manually download and install updates, you must be a member of the Administrators group for your computer. Windows checks for important updates and notifies you if any are available; updates are not delivered or installed on your computer until you choose to do so. When Windows finds updates for your computer, the **Windows Update** icon appears in the notification area and an alert pops up, letting you know that updates are ready to be downloaded. After you click the icon or the alert, you can select some or all of the updates to download.

Windows downloads the updates in the background—you are not notified or interrupted during this process and the updates do not interfere with other downloads. When downloading is complete, the **Windows Update** icon appears in the notification area again, this time to let you know that the updates are ready to be installed. You can choose to install all or some of the available updates.

**Turn off Automatic Updates**
You will never be notified when important updates are available for your computer, and you will never be asked to download or install them. This means that your computer can be vulnerable to security threats and harmful viruses that can damage your computer or your files. Viruses can also spread over the Internet to other people with whom you exchange e-mail, share files, or work with on a network.

New viruses and security threats are continually developed by attackers, so helping protect your computer is an ongoing process. If you do not turn on Automatic Updates, we recommend that you regularly install updates from the Windows Update Web site (http://www.microsoft.com/).

### Notes

- To open System, click **Start**, click **Control Panel**, click **Performance and Maintenance**, and then click System.

- Installing updates before you shut down your computer is another way to keep your computer up to date and more secure. This option is available only in Microsoft Windows XP Service Pack 2 (SP2) and only if important updates have been downloaded but not yet installed. Do not turn off or unplug your computer while updates are installing. Windows will automatically turn off your computer after the updates are installed.

- Only users with administrator privileges may add or remove programs, including Windows updates. It is strongly recommended that you log out of the computer administrator account when you are not performing tasks that require administrator privileges. If you are logged on as an administrator when your computer is the target of a virus or malicious user, the attack can cause extensive damage. For example, it might be able to reformat your hard drive, delete all your files, or create a new administrator account so the attacker can take over your computer. For more information about user accounts and why you should not run your computer as an administrator, see Help and Support.

**See Also**
**Concepts**
Automatic Updates: Frequently Asked Questions
How does Automatic Updates work?
Automatic Updates How to...

# Windows Update privacy statement

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Update privacy statement**
Last updated: 4/15/2004

View the latest version of this privacy statement at the Windows Update Web site (http://www.microsoft.com/).

Microsoft is committed to protecting your privacy.

**What data is collected—and why?**
Windows Update collects general system information from your computer with each visit, so that you receive the updates that work best with your computer. The information is also used to generate aggregate statistics about how the Windows Update Web site is used and which systems need support, so that we can improve our service. This information includes:

- Computer make and model
- Version information for the operating system, browser, and any other Microsoft software for which updates might be available
- Plug and Play ID numbers of hardware devices
- Region and language setting
- Globally Unique Identifier (GUID)
- Product ID and Product Key
- BIOS name, revision number, and revision date

Your Internet Protocol (IP) address is logged when you connect to the Windows Update site, but this address will only be used to generate aggregate statistics.

**How is this data used?**
Windows Update collects the Product ID and Product Key to confirm that you are running a validly licensed copy of Windows. The Product ID and Product Key collected are not retained after you are finished using Windows Update, unless the Product ID is not valid.

To generate accurate statistics, Windows Update evaluates a Globally Unique Identifier (GUID) that is stored on your computer to uniquely identify it. The GUID does not contain any information that can be used to identify you. A GUID is assigned and tracked in the following cases:

To provide customers with the best possible service, Windows Update tracks and records the number of individual computers that visit the site and whether the download and installation of specific updates succeeded or failed. Windows Update records the GUID of the computer that attempted the download, the ID of the item that was requested, whether updates were required, and the configuration information listed above.

Windows Update logs an additional GUID if you provide responses about whether help and troubleshooting articles were useful in resolving your problem. This allows Windows Update to provide you with increasingly helpful and relevant information.

Microsoft collects information about the pages our customers visit within microsoft.com, including Windows Update. This information might include: your IP address, browser type, operating system, domain name, the time at which you accessed the site, and referring Web site addresses. This site visitation data is identified only by a unique ID number used solely for this purpose.

**About surveys**
Occasionally you might be invited to participate in a survey about the way you use the Windows Update Web site. Each survey includes a privacy statement that details the terms and use of any information submitted with that survey.

# Licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Licensing**

This section provides information about the licensing program for the Microsoft® Windows Server 2003 family.

- Before you begin setting up a licensing model, see Licensing Checklists.
- To find features that have been changed in the Windows Server 2003 family, see New ways to do familiar Licensing tasks.
- For tips on setting up licensing, see Best practices for Licensing.
- For help with specific tasks, see Licensing How To ....
- For general information about licensing, see Licensing Concepts.
- For problem-solving instructions, see Licensing Troubleshooting.
- For information about required licenses for your Windows Server 2003 operating system deployment, see the Microsoft Web Site.

**Note**

- This section does not contain information about Terminal Server licensing. For information about Terminal Server licensing, see Terminal Server Licensing.

# Licensing Checklists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

- Checklist: Configure licensing on a server
- Checklist: Set up licensing for the enterprise

# Checklist: Configure licensing on a server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configure licensing on a server**

This checklist shows the steps that are required to configure licensing on a local server. It also lists several optional steps that you might need to take after licensing is configured. To configure licensing on a remote server, see Checklist: Set up licensing for the enterprise.

| | Step | Reference |
|---|---|---|
| | Review key concepts. | Administering and tracking licenses |
| | Add Per Server licenses by using the **Licensing** option in Control Panel. | Add Per Server licenses on the local server |
| | Configure replication on a server. | Licensing replication; Configuring licensing replication; Configure licensing replication on the local server |
| | (Optional) Delete Per Server licenses by using the **Licensing** option in Control Panel. | Delete Per Server licenses on the local server |
| | (Optional) Change the licensing mode from Per Server to Per Device or Per User on a server. | Change the local server to Per Device or Per User licensing |

# Checklist: Set up licensing for the enterprise

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Set up licensing for the enterprise**

When you create a Windows Server 2003 operating system site, license logging is also set up for the site. However, you might need to change the licensing setup for the site, or for individual servers in the site. To configure licensing locally on a server, using **Licensing** in Control Panel, see Checklist: Configure licensing on a server.

This checklist shows the steps that are required to configure licensing on a remote server.

| | Step | Reference |
|---|---|---|
| | Review key concepts. | Administering licensing for the enterprise |
| | View license information for the enterprise. | Find the site license server for a site; View licensing for a site |
| | Select the licensing mode for servers in a site or change the licensing mode for a server to Per Device or Per User. | Change the licensing mode for a server to Per Device or Per User licensing |
| | Configure replication for servers in a domain. | Tracking Your Organization's Licenses; Configure replication for a server |
| | Create license groups when the number of users is not the same as the number of devices. | License groups; Create a new license group |
| | Change the site license server. | Change the site license server for a site |

# New ways to do familiar Licensing tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The user interface for Licensing in the Windows Server 2003 family is similar to that in Windows 2000 or Microsoft Windows NT 4.0. The following table lists the tasks that are different.

| **If you want to** | **In Windows NT 4.0** | **In the Windows Server 2003 family and Windows 2000** |
|---|---|---|
| Change the master server | Use **Licensing** in Control Panel or **Licensing** in **Administrative Tools**. | Not defined for Windows Server 2003 operating systems or Microsoft Windows 2000 domain. For more information, see Licensing replication or Configure replication for a server. |
| Change the site license server | Not defined for Windows NT 4.0 domain. | Open Active Directory Sites and Services. For more information, see Licensing replication, Change the site license server for a site, or Configure replication for a server. |

# Best practices for Licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

Consider the following factors when configuring and tracking licensing for your organization:

- **Per Device or Per User licensing mode is the best option if your clients frequently use multiple servers on the network.**

  The Per Device or Per User licensing mode enables all network devices or users to access all the servers on a network. The number of simultaneous connections to any server is unlimited. Per Device or Per User is the normal licensing mode for a server product that is installed on multiple servers in a network.

  For more information, see the Microsoft Web Site .

- **Per Server is the best licensing option when a server product is installed on only one server accessed at any time by no more than a subset of your users.**

  Per Server connections to a particular server are allocated on a first-come, first-served basis, and are limited to the number of Client Access Licenses allocated to the server. This might be more economical for a single-server network or for a server product that is accessed on one server by only one department or group in an organization. As your network grows, you are allowed a one-time conversion to Per Device or Per User licensing.

  For more information, see the Microsoft Web Site .

- **If you install a server product on a single server and you are uncertain which licensing mode is appropriate, choose Per Server.**

  As your network grows, or when you install the product on additional servers, you can make a one-time switch from Per Server to Per Device or Per User licensing. The licensing agreement does not permit the reverse switch from Per Device or Per User to Per Server licensing.

- **When the number of users is different from the number of devices, create a license group to allocate licenses for Per Device or Per User licensing.**

  Per Device or Per User licensing requires one Client Access License for each device. However, the License Logging service assigns and tracks licenses by user name. You should create license groups to correctly allocate licenses and track usage when multiple users share one or more devices or when one user has multiple devices.

  For more information, see License groups.

- **To replicate license data from a server running Windows NT 4.0 to the licensing database for a** Windows Server 2003 **operating system site, specify the site license server as the enterprise server for replication on the Windows NT 4.0 domain controller.**

  By default, Windows NT 4.0 servers replicate to the domain controller or to an enterprise licensing server. On all servers running Windows Server 2003 , the License Logging service replicates automatically to the site license server. In the replication options for a Windows NT 4.0 domain controller, always choose to replicate to the enterprise server and always specify the site license server as the enterprise server.

  For more information, see:

  - Configuring licensing replication
  - Configuring Windows NT 4.0 licensing replication

# Licensing How To ...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Manage Licensing on the Local Server
- Manage Licensing for an Enterprise
- Manage Products and Product Licenses
- Manage User Licensing

# Manage Licensing on the Local Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage licensing on the local server**

- Add Per Server licenses on the local server
- Delete Per Server licenses on the local server
- Change the local server to Per Device or Per User licensing
- Configure licensing replication on the local server

# Add Per Server licenses on the local server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add Per Server licenses on the local server**

1. On the server you want to configure, open Choose Licensing Mode.
2. In **Product**, select the product for which you want to add licenses.
3. Click **Add Licenses**.
4. In **Product**, select the product.
5. In **Quantity**, type the number of new licenses purchased by your organization, and then click **OK**.
6. If you agree to the licensing terms, select the **I agree that** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon. In this case, click **License.**
- You can set the number of Per Server connections to 0 for a Windows Server 2003 operating system if it is used only as an operating system for an application and does not provide basic network services (such as file sharing, print sharing, remote access, and authenticated access).
- This procedure is only for registering the purchase of Client Access Licenses on a server running in Per Server mode.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete Per Server licenses on the local server
Using Licensing

# Delete Per Server licenses on the local server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete Per Server licenses on the local server**

1. On the server you want to configure, open Choose Licensing Mode.
2. In **Product**, click the server product from which you want to remove licenses.
3. Click **Remove Licenses**.
4. Under **Serial Number**, click the certificate you want to change.
5. In **Number of Licenses to remove**, type the number of licenses you want to remove, and then click **Remove**.
6. To close the dialog box, click **Cancel**.

   Clicking **Cancel** closes the dialog box but does not cancel any changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon. In this case, click **Licensing**.
- After removing licenses for a product, you can allocate them to a different server.
- This procedure can also be performed remotely with the Licensing administrative tool. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Per Server client access licenses
Manage Licensing on the Local Server

# Change the local server to Per Device or Per User licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the local server to Per Device or Per User licensing**

1. On the server you want to configure, open Choose Licensing Mode.
2. In **Product**, click the product for which you want to change the licensing mode.
3. Click **Per Device or Per User**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon. In this case, click **Licensing**.
- The Per Device or Per User licensing mode requires a Client Access License for every device or user that connects to the selected server product. After a device or user is licensed, it can access any server running the product. Typically, this is more economical on networks where clients connect to more than one server.
- If you are uncertain which mode to use, leave the server product in **Per Server** mode. The license agreement allows a one-time, one-way change from Per Server to Per Device or Per User licensing. You cannot switch from Per Device or Per User to Per Server.
- If you are performing a one-time change from Per Server to Per Device or Per User licensing mode, the Client Access Licenses allocated to the server in Per Server mode can be added to the pool of Per Device or Per User licenses available for client devices or users accessing the server product on any server in the site. Use Licensing in Administrative Tools to add Per Device or Per User licenses for the product.
- This procedure can also be performed remotely with the Licensing administrative tool. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Licensing

# Configure licensing replication on the local server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure licensing replication on the local server**

1. On the server you want to configure, open Choose Licensing Mode.
2. Click **Replication**.
3. Under **Replication Frequency**, specify the desired interval for replicating licensing information to the site license server:

   To manually specify a time for daily replication on this server, click **Start At**, and then enter a time.

   Click **Start Every**, and then enter the interval at which the system should stagger the start times if multiple servers are set to this option.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon. In this case, click **Licensing**.
- If you select **Start At** under **Replication Frequency** and multiple servers replicate to the site license server, you must vary the starting times on each server to balance the load. Replication might not start until 15 minutes after the specified time.
- The default is **Start Every** with a 24-hour interval. You can change this to an interval from one to 72 hours.
- This procedure can also be performed remotely with the Licensing administrative tool. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Licensing replication
Configuring licensing replication

# Manage Licensing for an Enterprise

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage licensing for an enterprise**

- Find the site license server for a site
- View licensing for a site
- View licensing for a Windows NT 4.0 domain
- View licensing for a server
- Change the licensing mode for a server to Per Device or Per User licensing
- Configure replication for a server
- Change the site license server for a site

# Find the site license server for a site

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find the site license server for a site**

1. Open Active Directory Sites and Services.
2. In the console tree, click the site you want to configure.

   **Where?**

   - Active Directory Sites and Services/Sites/*Site*

3. In the details pane, click **Licensing Site Settings**.
4. On the **Action** menu, click **Properties**.
5. Under **Licensing Computer**, read the **Computer** and **Domain** for the site license server.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- The default site license server is the first domain controller created in the site, but the site license server does not have to be a domain controller. For optimal performance, however, it is recommended that the site license server and domain controller be in the same site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the site license server for a site
Licensing replication

# View licensing for a site

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view licensing for a site**

1. Open Licensing.
2. On the **License** menu, click **Select Domain**.
3. Enter the name of the site license server for the domain.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- To find the name of the site license server for a site, see "Find the site license server for a site" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Find the site license server for a site
Administering licensing for the enterprise

# View licensing for a Windows NT 4.0 domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view licensing for a Windows NT 4.0 domain**

1. Open Licensing.
2. On the **License** menu, click **Select Domain**.
3. In **Domain**, enter the name of the domain to administer.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- After you select a domain, Licensing displays the data pertaining to the server to which licensing information for the domain is replicated. For a Windows NT 4.0 domain, this is the enterprise server specified for licensing replication on the primary domain controller. If no enterprise server is specified, this is the primary domain controller for the domain.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Windows NT 4.0 licensing replication
Licensing overview

# View licensing for a server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view licensing for a server**

1. Open Licensing.
2. On the **Server Browser** tab, expand the domain for the server.
3. Right-click the server.
4. Click **Properties**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- To find out which licensing modes are available for a server product, click the server product on the **Products View** tab in **Properties**, and then click **Edit**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage Licensing for an Enterprise
Manage Licensing on the Local Server

# Change the licensing mode for a server to Per Device or Per User licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the licensing mode for a server to Per Device or Per User licensing**

1. Open Licensing.
2. On the **Server Browser** tab, expand the domain name.
3. Right-click the site license server in the domain, and then click **Properties**.
4. On the **Products View** tab, click the product, and then click **Edit**.
5. In **Choose Licensing Mode**, click **Per Device or Per User**, and then click **OK**.
6. After reading the information in **Per Device or Per User Licensing**, select the check box to agree to be bound by the license agreements for this product, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**. In this case, click **License.**
- The Per Device or Per User licensing mode requires a Client Access License for every device or user that connects to the selected server product. Licensed devices or users can access any server running the product. Typically, this is more economical on networks where client devices or users connect to more than one server.
- If you are uncertain which mode to use, leave the server product in Per Server mode. The license agreement allows a one-time, one-way change from Per Server to Per Device or Per User licensing. You cannot switch from Per Device or Per User to Per Server.
- If you are performing the one-time change from Per Server to Per Device or Per User licensing mode, the Client Access Licenses allocated to the server in Per Server mode can be added to the pool of Per Device or Per User licenses available for client devices or users accessing the server product on any server in the site.
- To perform this procedure locally on a server, click **Licensing** in Control Panel.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete Per Device or Per User client access licenses
Add Per Device or Per User client access licenses

# Configure replication for a server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure replication for a server**

1. Open Licensing.
2. On the **Server Browser** tab, expand the domain.
3. Right-click the server to manage, and then click **Properties**.
4. Click the **Replication** tab.
5. Under **Replication Frequency**, specify the interval at which licensing information is replicated to the site license server:
   - To manually specify a time for daily replication on this server, click **Start At** and enter a time.
   - To manually set an interval of time between licensing information replication, click **Start Every** and enter the desired interval.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- If you select **Start At** under **Replication Frequency** and multiple servers replicate to the site license server, you must vary the starting times on each server to balance the load. Replication might not start until 15 minutes after the specified time.
- Windows Server 2003 operating systems replicate automatically to a domain controller, and domain controllers replicate to the site license server. If you have a standalone server in a Windows NT 4.0 domain, that server replicates to the primary domain controller. The option to specify a master server applies only to servers running Windows NT 4.0 and is unavailable when you are managing a server running Windows 2000 or Windows Server 2003.
- When you are managing a server running Windows NT 4.0, you can use the **Master Server** options to specify where the server replicates. For more information, see "Licensing replication" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Licensing replication
Configuring Windows NT 4.0 licensing replication

# Change the site license server for a site

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the site license server for a site**

1. Open Active Directory Sites and Services.
2. In the console tree, click the site you want to configure.

   **Where?**

   - Active Directory Sites and Services/Sites/*Site*

3. In the details pane, click **Licensing Site Settings**.
4. On the **Action** menu, click **Properties**.
5. Under **Licensing Computer**, to select a different server to be the site license server, click **Change**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- To retain the licensing history for your enterprise, stop the License Logging service on the new site license server, and then copy the following files from the old site license server to the same location on the new site license server:

  - *Systemroot*\System32\Cpl.cfg, which contains the purchase history for your organization.
  - *Systemroot*\Lls\Llsuser.lls, which contains user information about the number of connections.
  - *Systemroot*\Lls\Llsmap.lls, which contains license group information.

  Restart the License Logging service after you copy these files. This transfer must be completed immediately after you change to the new site license server, or you might lose licensing data replicated from other servers.

- For optimal performance, install the site license server in the site whose licenses you want to manage. The default site license server is the first domain controller created in the site, but the site license server does not have to be a domain controller.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start, stop, pause, resume, or restart a service

Find the site license server for a site

# Manage Products and Product Licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage products and product licenses**

- View product license information
- Add Per Device or Per User client access licenses
- Delete Per Device or Per User client access licenses
- Add Per Server client access licenses
- Delete Per Server client access licenses

# View product license information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view product license information**

1. Open Licensing.
2. On the **Products View** tab, right-click the server product name.
3. Click **Properties**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- The three tabs in the **Properties** dialog box for a server product show the following information.

| Tab | Shows |
|---|---|
| **Clients** | Users have used the server product in Per Device or Per User licensing mode. |
| **Purchase History** | The date, quantity, and administrator for license purchases. You can add or delete licenses. |
| **Server Browser** | The license mode for servers in the site and the number of licenses for Per Server mode. You can add or remove Per Server licenses, change the licensing mode, or configure replication of licensing information. |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering licensing for the enterprise
Tracking Your Organization's Licenses

# Add Per Device or Per User client access licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add Per Device or Per User client access licenses**

1. Open Licensing.
2. On the **License** menu, click **New License**.
3. In **Product**, click the server product you want to manage.
4. In **Quantity**, type the number of licenses purchased.
5. In **Comment**, you can type a note to aid in identifying the purchase.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- Client access licenses added in Per Device or Per User licensing mode are added to the pool of licenses for Windows Server 2003 operating systems. As devices or users access the product anywhere in the site, they are allocated licenses from the pool, one license for each device or user. After the pool of available licenses is depleted, license violations occur when additional devices or users access the product.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete Per Device or Per User client access licenses

# Delete Per Device or Per User client access licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete Per Device or Per User client access licenses**

1. Open Licensing.
2. On the **Products View** tab, click the server product from which you want to delete licenses.
3. On the **License** menu, click **Delete**.
4. In **Number of Licenses to remove**, type the number of licenses to be deleted.
5. Click **Remove**.
6. Click **Cancel** to close the dialog box. Clicking **Cancel** does not cancel your changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Per Device or Per User client access licenses

# Add Per Server client access licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add Per Server Client Access Licenses**

1. Open Licensing.
2. On the **Server Browser** tab, expand **Enterprise**, the domain name, and then the server you want to manage. The list of server products installed on the server is displayed.
3. Right-click the server product for which you want to add licenses, and then click **Properties**.
4. Click **Add Licenses**.
5. In **Quantity**, type the number of licenses to add.
6. After reading the information in **Per Server Licensing**, select the check box to agree to be bound by the license agreements for this product, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**. In this case, click **License.**
- Client Access Licenses added in Per Server mode allow for connections to a server product on a particular server. When a user connects to the server product on that server, the connection consumes a license. When the user disconnects from the server product, the License Logging service makes the license available to another user.
- You can set the number of Per Server connections to 0 for a Windows Server 2003 operating system if it is used only as an operating system for an application and does not provide basic network services (such as file sharing, remote access, and authenticated access).
- This procedure can also be performed locally on a server with the Licensing option in Control Panel. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete Per Server client access licenses
Administering and tracking licenses

# Delete Per Server client access licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete Per Server Client Access Licenses**

1. Open Licensing.
2. On the **Server Browser** tab, expand **Enterprise**, the domain name, and then the server you want to manage. The list of server products installed on the server is displayed.
3. Right-click the server product for which you want to remove licenses, and then click **Properties**.
4. Click **Remove Licenses**.
5. In **Number of licenses to remove**, type the number of licenses.
6. Click **Remove**, and click **Yes** to confirm the change.
7. Click **Cancel** to close the dialog box, and then click **OK** to save the changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- This procedure can also be performed locally on a server with the Licensing option in Control Panel. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Per Server client access licenses
Administering and tracking licenses

# Manage User Licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage users**

- View user licensing information
- Revoke a user's Client Access License for a product
- Delete a user from a license
- Create a new license group
- Edit or view a license group

# View user licensing information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view user licensing information**

1. Open Licensing.
2. On the **Clients** tab, right-click the user's name.
3. Click **Properties**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- The user properties show the number of times a user has accessed server products. You can also revoke a user's product license.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Revoke a user's Client Access License for a product
Delete a user from a license
Manage User Licensing

# Revoke a user's Client Access License for a product

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To revoke a user's Client Access License for a product**

1. Open Licensing.
2. Click the **Products View** tab.
3. Right-click the product, and then click **Properties**.
4. On the **Clients** tab, click the user.
5. Click **Revoke**, and then click **Yes.**

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- Revoking a license returns it to the pool of available Per Device or Per User licenses for the server product.
- You might want to revoke a user's license if the usage statistics show that the user does not access the product often enough to warrant retaining the Client Access License.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete Per Device or Per User client access licenses
Add Per Device or Per User client access licenses
Delete a user from a license
Tracking user licenses for multiple sites

# Delete a user from a license

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a user from a license**

1. Open Licensing.
2. On the **Clients** tab, select the name of the user.
3. On the **License** menu, click **Delete**, and then click **Yes**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.
- Deleting a user releases licenses allocated to the user to the respective product license pools for Microsoft server products. Perform this procedure when a user has been removed from the network or no longer accesses the server products regularly.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View user licensing information
Manage User Licensing

# Create a new license group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new license group**

1. Open Licensing.
2. On the **Options** menu, point to **Advanced**, and then click **New License Group**.
3. In **Group Name**, type a short descriptive name.
4. In **Description**, type additional details to help identify the group.
5. In **Licenses**, enter the number of licenses needed in Per Device or Per User licensing mode.
6. Click **Add**.
7. In **List Users from**, click the appropriate domain.
8. Under **Users**, select the users you want in the group, and then click **Add**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
License groups
Manage User Licensing

# Edit or view a license group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit or view a license group**

1. Open Licensing.
2. On the **Options** menu, point to **Advanced**, and then click **Edit License Groups**.
3. Perform any of the following tasks.

| To | Do this |
|---|---|
| Add a new license group | Click **Add**. |
| Delete a license group | Click the group, and then click **Delete**. |
| Change the description, number of licenses, or users | Click the group, and then click **Edit**. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Licensing, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Licensing**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
License groups

# Licensing Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

- Licensing overview
- Using Licensing
- Licensing Resources

# Licensing overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Licensing overview**

A software license grants you the legal right to use a piece of software. For each software program you use, you need a license, which is granted to you and documented in the license agreement. Your use of a software product is governed by the terms of the license agreement, as well as by copyright law. The unauthorized duplication of software constitutes copyright infringement, whether it is done for sale, for free distribution, or for personal use.

The license agreement that governs the use of Microsoft products is called the End User License Agreement (EULA). It is a contract between you and Microsoft. By acquiring the software and opening it, you accept the terms of the license agreement and you are granted the right to use the software. It is important to keep the EULA in your files. It is proof that you are the legal, licensed user of the software program.

You must have server licenses for a Microsoft® Windows Server 2003 operating system and for other server products. You must also have a Client Access License for each device or user that connects to that server.

The following topics provide more detailed information about licensing.

- For information about how to track licensing, see Using Licensing.
- For information about where to find additional Microsoft licensing resources, including licensing for specific Windows Server 2003 operating systems, and information about enterprise licensing, see Licensing Resources.
- For information about required licenses for your Windows Server 2003 operating system deployment, see the Microsoft Web Site.

# Using Licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Licensing**

- Administering and tracking licenses
- Tracking Your Organization's Licenses

# Administering and tracking licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering and tracking licenses**

Tracking licenses manually within single domain or across multiple domains can be time consuming and expensive. Windows Server 2003 operating systems provide two utilities that automate this task and help you minimize the time and expense of license tracking. You can use these utilities to view data and manage licensing for server products and to ensure that you are complying with the Windows Server 2003 operating system licensing requirements.

**Important**

- You must be an administrator or a member of the Administrators group in order to use either the **Licensing** option in **Control Panel** or the **Licensing** tool in **Administrative Tools**.

**Licensing in Control Panel**

The **Licensing** option in Control Panel manages licensing requirements for a single computer running a Windows Server 2003 operating system. Use it locally on the computer you are managing to:

- Add or remove Client Access Licenses for a server product licensed in Per Server mode.
- Change the licensing mode from Per Server to Per Device or Per User (one time only).
- Configure the time or frequency of replication of the server's licensing information to the centralized database on the site license server.

For more information, see Manage Licensing on the Local Server.

**Licensing in Administrative Tools**

Use **Licensing** in **Administrative Tools** to manage licensing for the enterprise. It enables you to administer and track licensing data for a site in the centralized licensing database on the site license server. You can:

- View Per Server and Per Device or Per User licenses for the site or a group in the site.
- Manage the purchase or deletion of licenses for server products on network servers.
- View usage statistics for each user.
- Remotely manage licensing on servers for server products licensed in Per Server mode.
- Manage license replication on the network.

For more information, see Administering licensing for the enterprise.

# Tracking Your Organization's Licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tracking your organization's licensing**


- Administering licensing for the enterprise
- Licensing replication
- Configuring licensing replication
- Configuring Windows NT 4.0 licensing replication
- License groups
- Tracking user licenses for multiple sites

# Administering licensing for the enterprise

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering licensing for the enterprise**

To administer Windows Server 2003 operating system licensing for sites or the enterprise, use **Licensing** in **Administrative Tools**. You can track licensing data replicated to the site license server for sites or replicated to the primary domain controller for Windows NT 4.0 domains. For more information, see View licensing for a site and View licensing for a Windows NT 4.0 domain.

**Licensing** in **Administrative Tools** divides tasks into Purchase History, Products View, Clients, and Server Browser.

**Purchase History: To view license purchases**

The **Purchase History** tab in **Licensing** provides a historical overview of the licenses purchased for a site. It also shows the quantity of licenses purchased and deleted, the date of addition or deletion, and the name of the administrator who performed the action.

**Products View: To view and manage product licensing**

The **Products View** tab of **Licensing** provides licensing information for server products in a site. The information displayed for a server product depends on the licensing mode for the product.

For a product in Per Device or Per User mode, **Product View** shows how many licenses have been purchased for the product and how many licenses have been allocated to users or devices for all the products.

For Per Server mode, **Product View** shows the total number of licenses that have been purchased for each product on all the servers in the site. It also shows the maximum number of Per Server connections reached to the current date on all servers with the product in the site.

You can use both the numerical and graphical information shown in **Products View** to determine the following information about your licenses:

| Symbol | Licensing status |
|---|---|
| | The product is in compliance with legal licensing requirements. |
| | The product is not in compliance with legal license requirements. |
| | The product has reached the legal limit, and you might want to purchase additional licenses. |

In **Products View**, you can perform the following licensing tasks:

- Add new Client Access Licenses for a product
- Delete Client Access Licenses for a product
- View client accesses, purchase history, and server information for a product.

**Clients: To view and manage client licensing**

The **Clients** tab lists devices and users that have used the products in a domain or organization. For each device or user, it shows licensed and unlicensed usage of Windows Server 2003 operating systems. You can perform the following licensing tasks on this tab:

- View product usage by devices or users.
- Add new Client Access Licenses.
- Create, view, and edit license groups.

  **Important**

- Each device or user in a license group requires a client access license.

- Revoke a license allocated to a device or user.

You can use both the numerical and graphical information shown in **Clients** to determine the following licensing status:

| Symbol | Licensing status |
|---|---|
| | Users in compliance with legal licensing requirements (licensed usage). |
| | Users in violation of legal licensing requirements (unlicensed usage). |

**Server Browser: To view and manage server licensing**

The **Server Browser** tab lists all the sites, domains, and servers in your organization. It shows the Windows Server 2003 operating systems installed on each server. You can use this tab to view and configure licensing for server products on servers in any site or domain for which you have administrative authority.

Use the **Server Browser** tab to view the licensing configuration of any server in the enterprise. You can perform any of the following administrative tasks at the server level:

- View Windows Server 2003 operating systems installed on a server.

- View the licensing mode for each product on a server.

- For products in Per Device or Per User licensing mode, view the number of Client Access Licenses purchased.

- For any server product in Per Device or Per User licensing mode, add or remove Client Access Licenses for a product.

- Open and manage a selected item's properties at the server and product levels.

- For any Windows server, view and change the time or frequency of licensing replication.

- For a computer running Windows NT 4.0 Server, view and change the location where licensing replication information is stored.

The properties for each server list all the Windows Server 2003 operating systems installed on the server. For each product, you can add or remove Per Server licenses or change the licensing mode. The properties for each product show the licensing mode and, if the product is in Per Server mode, the number of licenses.

# Licensing replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Licensing replication**

An administrator can ensure an organization's legal compliance with the license agreements for Windows Server 2003 operating systems by monitoring license purchases, deletions, and usage. This licensing information is collected on a server by the License Logging service. The License Logging service is disabled by default.

The License Logging service on each server in a site replicates this licensing information to a centralized database on a server, called the site license server, for the site. A site administrator or administrator for the site license server can then use **Licensing** in **Administrative Tools** to view the licensing history for the entire site stored on the site license server.

Typically, the site license server is the first domain controller created for the site. In a large organization with multiple sites, the site license server in each site collects licensing information for that site.

To find which server is the site license server, see Find the site license server for a site. To change the site license server to a different server, see Change the site license server for a site. To view licensing information for any site on your network, see View licensing for a site.

# Configuring licensing replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring licensing replication**

On a computer running Windows 2000 Server or a Windows Server 2003 server operating system, you can configure how often licensing data is replicated to the domain controller or site license server. On a computer running Windows NT 4.0 Server, you can also configure where the data is replicated. These settings can be configured either locally on a server or remotely.

To configure licensing replication locally on a server, use the **Licensing** option in Control Panel. For more information, see Configure licensing replication on the local server. To configure replication remotely for a server in your network, use Licensing in Administrative Tools, as described in Configure replication for a server.

**Replication options**

The following table describes the options for configuring replication.

| Option | Use to | Use when |
|---|---|---|
| Start At | Specify a time each day when licensing information is replicated to the site license server. | You want to manually set up a starting time for replication on this computer.<br><br>You must vary the starting times on servers replicating to a particular master or enterprise server. |
| Start Every | Set the frequency for licensing replication. | The system staggers the times licensing data from each server is replicated. |

**Balancing the license replication load**

The **Replication Frequency** settings (**Start At** and **Start Every** in the table) affect how long licensing changes on individual servers take to propagate to the enterprise server. The default is to start every 24 hours, and the system automatically staggers replication to balance the load. You can lengthen or shorten the interval to change how quickly data is propagated.

If you want to control the time at which replication occurs, select **Start At** on each replicating server and specify a time. To balance the load, you must manually vary the starting time of each server replicating to a particular server. The actual replication might start up to 15 minutes before or after the specified time.

**Additional replication options for earlier versions of Windows servers**

Servers running Windows 2000 or Windows Server 2003 automatically replicate to the domain controller, and domain controllers replicate to the site license server. For Microsoft Windows NT 3.1 or Windows NT 4.0 servers on your network, you must configure where they replicate their licensing data. You do this with the **Master Server** options in the **Replication Configuration** dialog box. These options appear only for Windows NT 3.1 or Windows NT 4.0 Servers.

| Option | Use to | Use when |
|---|---|---|
| Domain Controller | Specify that this computer replicates to the primary domain controller. | This computer is a member server or a backup domain controller in a Windows NT 4.0 domain.<br><br>This computer is a standalone server or a primary domain controller, and you do not want to replicate further. |
| Enterprise Server | Specify that the master server is an enterprise server, and enter the name of the enterprise server. | This computer is a standalone server and part of an enterprise, and you want to replicate its licensing information directly to an enterprise-licensing database.<br><br>The computer is the primary domain controller for its domain, and you want to replicate to an enterprise-licensing database. |

For more information, see Configuring Windows NT 4.0 licensing replication and View licensing for a Windows NT 4.0 domain. To configure licensing locally on a Windows NT 3.1 or Windows NT 4.0 server, see online Help on the server.

# Configuring Windows NT 4.0 licensing replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Windows NT 4.0 licensing replication**

A server running Windows NT Server 4.0 replicates licensing data to a designated master server. By default, a server running Windows NT Server 4.0 replicates licensing data to the domain controller in its domain. An administrator can configure the domain controller to either retain the licensing data or replicate it to a designated *enterprise server*, which collects licensing data from all domains in an enterprise. This is different from a server running Windows 2000 or Windows Server 2003 , which always replicates to the site license server.

If you are administering licensing for a Windows 2000 or Windows Server 2003 operating system site, you can still configure replication for any server running Windows NT Server 4.0 in your network. For a Windows NT 4.0 domain controller, you can select **Domain Controller** (the default setting) to retain the licensing information for the domain on the domain controller. Alternatively, you can specify that the licensing information for the domain replicate to a specified enterprise server. Preferably, you want to consolidate licensing information from a Windows NT 4.0 domain controller with other information for your enterprise. To do this, choose to replicate licensing information to an enterprise server and specify the site license server as the enterprise server.

If a server running Windows NT Server 4.0 replicates to a domain controller that then replicates to an enterprise licensing database, data can take longer to reach the site license server than with a server running Windows 2000 or Windows Server 2003 replicating directly to the site license server. For example, with a replication interval of 24 hours, information can take up to 24 hours to replicate to a domain controller and then up to 24 additional hours to replicate to a site license server. When the option to replicate to the enterprise server is available for a server that is not a domain controller, you can speed the replication process by specifying that replication should go directly to the site license server.

**Important**

- Do not do this if you want to be able to view licensing data separately for a Windows NT 4.0 domain. For more information, see View licensing for a Windows NT 4.0 domain.

For more information about replication in Windows 2000 or Windows Server 2003 operating systems, see Licensing replication. For more information about master server replication options for servers running Windows NT 4.0, see " Additional replication options for earlier versions of Windows servers" in Configuring licensing replication. To locate the site license server for a site, see Find the site license server for a site. To view replication settings for a Windows NT 4.0 Server, follow the procedure Configure replication for a server.

# License groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**License groups**

In Per Device or Per User licensing mode, a license is required for each client device or user that accesses a server product. License Logging service, rather than allocating licenses to computers, allocates them to users when they access a server product for the first time. This fulfills the legal requirement of one Client Access License for each client device only when the number of client devices is the same as the number of users on a network--most simply, when each user has one device.

Having a number of users that is different from the number of devices might provide incorrect licensing information. To obtain correct licensing information, you should organize users in *license groups*. A license group pools the licenses for devices used by users in a group. When users in a group access the server product, they use the group license and do not consume additional licenses. Thus, you can maintain the one-to-one correspondence between licenses and devices.

Create license groups to obtain correct licensing information under the following conditions:

- **Many people use one device.** This can occur when people share jobs or when employees who are working different shifts use the same device.
- **Many users use many devices, but there are more users than devices.** This situation might exist in a university's computer lab or in a retail store.
- **One user uses more than one device.** This often occurs when software developers create programs on one device and test their applications on several other devices.

To create a license group, specify the users in the group and allocate one license for each client device used to access the server. The number of licenses allocated to a group should correspond to the number of devices used by members of the group. This can be either fewer or more than the number of users.

- If 10 people are using one device, you need only one Client Access License. The group would include all of the users, but have only one license.
- If 100 users are using 10 devices, you need 10 Client Access Licenses, one for each of the 10 devices. The group would list the 100 users and have 10 Client Access Licenses.
- If one user uses 5 devices, you need a Client Access License for each device. The group would list the single user and have 5 licenses.

To create a new license group, see Create a new license group. To add, delete, or change a license group, see Edit or view a license group.

# Tracking user licenses for multiple sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tracking user licenses for multiple sites**

Users accessing a Windows Server 2003 operating system in more than one site might be allocated multiple licenses for the product. The License Logging service for Windows Server 2003 operating systems collects site-licensing information on the site license server for a single site. This enables a geographical or organizational unit in an organization to purchase and manage their software licenses separately. However, because each site allocates licenses to users independently of other sites, a user connecting to a product in more than one site might be allocated more than one Client Access License for the same product.

If your organization has multiple sites and you notice unexpected consumption of software licenses, you should check for multiple licenses allocated to users in different sites. To search for multiple Client Access Licenses allocated to the same user in more than one site, you must compare the product information for different sites. To compare the users listed in **Clients** for different sites, use **Licensing** in **Administrative Tools**. For more information, see View product license information.

If a user from a different site is allocated a license for a server product in both sites, you can either delete the user or revoke the license. However, do not do this if the user uses a different device for each site--for example, if the user regularly travels between corporate offices and has a different device in each location. The license allocation would then be correct.

Multiple sites for an organization normally reflect groups of users that are distinct geographically or organizationally, working on distinct subnets owned by the organization. If you find that users are constantly accessing servers in different sites, you might want to reconsider how you structure the sites or how you distribute resources among sites.

# Licensing Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

**Licensing for the enterprise**
For more information on licensing for an enterprise, see Software Licensing at the Microsoft Web site.

# Licensing Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

This section provides troubleshooting information about Licensing.

- License Logging service messages
- Troubleshooting licensing

# License Logging service messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**License Logging service messages**

You might see one or more of the following informational or error messages from the License Logging service.

**License Logging service messages**

| |
|---|
| **3 Server** *ServerName* **has requested that license database replication to it be delayed.**<br><br>Indicates that the specified Site License Server or domain controller is unable to accept replicated data at this time. License Logging service will reattempt replication in 15 minutes. |
| **4 License database replication to server** *ServerName* **has started.**<br><br>Indicates that replication to the specified Site License Server or domain controller has started successfully. |
| **5 License database replication to server** *ServerName* **has completed successfully.**<br><br>Indicates that replication to the specified Site License Server or domain controller has completed successfully. |

**Warning and error messages**

| |
|---|
| **201 No license was available for user** *UserName* **using product** *ProductName* **. Use Licensing from the Administrative Tools folder to ensure that you have sufficient licenses.**<br><br>Indicates a licensing violation. All the licenses for the specified product have been assigned to users. Use Licensing in Administrative Tools to analyze usage. You can find out how many licenses should be purchased or revoke licenses from users who are not using them. |
| **202 The product** *ProductName* **is out of licenses. Use Licensing from the Administrative Tools folder for more information on which users are out of compliance and how many licenses should be purchased.**<br><br>Indicates a licensing violation. All the licenses for the specified product have been assigned to users. Use Licensing in Administrative Tools to find out which users are out of compliance. You can find out how many licenses should be purchased or revoke licenses from users who are not using them. |
| **203 The user data could not be saved. The data is the error code.**<br><br>Indicates that changes to the user information could not be saved. The License Logging service may be able to successfully save the data later. If not, the data will be lost. Data that is replicated from other servers will be received again during the next replication. Data entered on this server will have to be re-entered. |
| **204 The license group mapping data could not be saved to LlsMap.LLS. The data is the error code.**<br><br>Indicates that changes to the license groups could not be saved. The License Logging service may be able to successfully save the data later. If not, the data will be lost. Data replicated from other servers will be received again during the next replication. Data entered on this server will have to be re-entered. |
| **205 The purchased license data could not be saved to CPL.CFG. The data is the error code.**<br><br>Indicates that changes to license purchase information could not be saved. The License Logging service may be able to successfully save the data later. If not, the data will be lost. Data replicated from other servers will be received again during the next replication. Data entered on this server will have to be re-entered. |
| **208 The saved user data could not be restored.**<br><br>Indicates that user data could not be read from a saved file. The License Logging service will overwrite the file the next time it saves information. Data that was received from other servers will be retrieved during the next replication, but data entered for this server will have to be re-entered. |

**209 The saved license group data could not be restored.**

Indicates that user data could not be read from a saved file. The License Logging service will overwrite the file the next time it saves information. Data that is received from other servers will be retrieved during the next replication, but data entered for this server will have to be re-entered.

**210 The saved purchased license data could not be restored.**

Indicates that the license groups data could not be read from a saved file. The License Logging service will overwrite the file the next time it saves information. Data that is received from other servers will be retrieved during the next replication, but data entered for this server will have to be re-entered.

**213 Replication of license information failed because the License Logging service on server** *ServerName* **could not be contacted.**

Indicates that either the indicated Site License Server or domain controller is unavailable on the network or the License Logging service on that computer is stopped. Verify that the computer is on the network and is running. If the computer is available on the network, use Services in the Computer Management System Tools to verify that the License Logging service is running on that computer.

**214 The License Logging service encountered an error while initiating replication to server** *ServerName* **.**

Indicates that the License Logging service on the indicated server was contacted, but an error occurred before replication could begin. This is probably caused by communication problems on the network. The License Logging service will retry replication every 15 minutes until it succeeds.

**215 License database replication to server** *ServerName* **was unsuccessful.**

Indicates that replication to the named server began successfully but failed to complete. This is probably caused by communication problems on the network. The License Logging service will retry replication every 15 minutes until it succeeds.

**218 The certificate database could not be saved.**

Indicates that License Logging service could not save the certificate database. Check whether the server has enough disk space.

**219 The certificate database could not be restored.**

Indicates that the database file may be corrupted. No action is required, because the database will be automatically regenerated.

**220 License database replication cannot be performed to server** *ServerName* **because the version of Windows NT installed there does not support the License Logging service.**

Indicates that the specified server is available on the network, but License Logging service is not installed, because it is running a version of Windows NT earlier than 3.51.

**221 License usage for product** *ProductName* **is nearing the maximum number of Per Server licenses purchased. Consult Licensing from the Administrative Tools folder for more information.**

Indicates that the specified product is licensed in Per Device or Per User mode. Nearly all the Client Access Licenses are assigned, and you may need to purchase additional licenses to provide for the clients connecting to this product on your network. You can also revoke licenses from clients that are no longer using the product. Use Licensing in Administrative Tools to analyze usage.

**222 License usage for product** *ProductName* **is at the maximum number of Per Server licenses purchased. New license requests cannot be granted. Consult Licensing in Administrative Tools for more information.**

Indicates that the specified product is licensed in Per Device or Per User mode. All the Client Access Licenses are being used, and you need to purchase additional licenses to provide for the client devices connecting to this product on your network. You can also revoke licenses from client devices that are no longer using the product. Use Licensing in Administrative Tools to analyze usage.

# Troubleshooting licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting licensing**
What problem are you having?

- No more connections can be made to a server product because the number of connections has reached the maximum that the server can accept.
- The number of Per Device or Per User Client Access Licenses allocated is more than the number of client devices accessing a Windows Server 2003 operating system.
- The number of Client Access Licenses allocated is fewer than the number of client devices used to access a product.
- Client Access Licenses for a product are allocated to users who belong to another site.

**No more connections can be made to a server product because the number of connections has reached the maximum that the server can accept.**
**Cause:** The server product might be configured with Per Server licensing and the number of licenses might be exhausted.

**Solution:** Check license usage for the product on the server. The user can wait until others stop accessing the product. To eliminate the problem, you can purchase more licenses for the product.

**See also:** View product license information

**The number of Per Device or Per User Client Access Licenses allocated is more than the number of client devices accessing a Windows Server 2003 operating system.**
**Cause:** A Client Access License is required for each client device or user that accesses a Windows Server 2003 operating system. Any of the following situations can consume extra licenses:

- A user who accessed a server product has been removed from the network.
- A user is still on the network, but no longer uses the server product.
- Several users access a product from a single client device.

**Solution:** To release allocated licenses, do the following:

- When you remove a user from the network, use **Licensing** in **Administrative Tools** to delete the user.
- Use **Licensing** in **Administrative Tools** to check how recently and often a user accesses a particular server product. Revoke a Client Access License if a user no longer uses the product.
- When several users share one or more client devices, create a license group to allocate the same number of Client Access Licenses among the users as there are client devices.

**See also:** View user licensing information, Delete a user from a license, Revoke a user's Client Access License for a product, or Create a new license group

**The number of Client Access Licenses allocated is fewer than the number of client devices used to access a product.**
**Cause:** A user accesses a server product from more than one device.

**Solution:** Whenever one or more users use several different devices, create a license group to allocate the pooled Client Access Licenses for the client devices to the users. This can cause more Client Access Licenses to be allocated than there are client devices accessing a product.

**See also:** Create a new license group

**Client Access Licenses for a product are allocated to users who belong to another site.**

**Cause:** License Manager tracks usage for a single site, and does not distinguish users who access a particular server product in more than one site from a single client device.

**Solution:** Use **Licensing** in **Administrative Tools** to revoke licenses or delete users from another site. Do not revoke a Client Access License from a user who uses different devices that belong to the different sites. The number of licenses then corresponds correctly to the number of client devices. If a user in a site shares a device with a user from another site (for example, a visitor from another division of the company), then create a license group to correctly allocate the corresponding Client Access License to both users.

**See also:** Tracking user licenses for multiple sites, Delete a user from a license, Revoke a user's Client Access License for a product, or Create a new license group

# Windows Installer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer**
Windows Installer makes it easy for you to manage the software installation and removal process, and modify, repair, and advertise existing applications.

- For help with specific tasks, see Windows Installer How To....
- For general background information, see Windows Installer Concepts.
- For problem-solving instructions, see Windows Installer Troubleshooting.

# Windows Installer How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Use Windows Installer packages with Windows Explorer
- Use Windows Installer packages with the Add or Remove Programs utility
- Use Windows Installer from the command line
- Configure Group Policy for Windows Installer

# Use Windows Installer packages with Windows Explorer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use Windows Installer packages with Windows Explorer**

1. Open Windows Explorer.
2. Right-click the Windows Installer package (.msi file), and then click **Install**, **Repair**, or **Uninstall**.

### Note

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Use Windows Installer packages with the Add or Remove Programs utility
Use Windows Installer from the command line

# Use Windows Installer packages with the Add or Remove Programs utility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use Windows Installer packages with Add or Remove Programs**

1. Open Add or Remove Programs in Control Panel.
2. Click **Change or Remove Programs**, click the program that you want to change or remove, and then do one of the following:
   - To change a program, click **Change**.
   - To remove a program, click **Remove**.

### Note

- To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Use Windows Installer packages with Windows Explorer
Use Windows Installer from the command line

# Use Windows Installer from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Windows Installer from the command line**


- Install a transform using the command line
- Install and advertise a product using a transform at the command line
- Install a package using the command line
- Remove a package using the command line
- Repair a package using the command line
- Advertise a package using the command line
- Create an administrative installation package from the command line
- Create an installation log file using the command line

# Install a transform using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a transform using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /i***package***TRANSFORMS=***TransformList*

| Value | Description |
|---|---|
| **/i** | Installs a product. |
| *package* | Specifies the name of the Windows Installer package file. |
| **TRANSFORMS=***TransformList* | Specifies what transform (.mst) files should be applied to the package. *TransformList* is a list of paths separated by semicolons. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Install a package using the command line

# Install and advertise a product using a transform at the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install and advertise a product using a transform at the command line**

1. Open Command Prompt.

2. Type:

   **msiexec***package* **/j**{**u** | **m**} **/t***TransformList*

| Value | Description |
|---|---|
| *package* | Specifies the name of the Windows Installer package file. |
| **/j** | Advertises a product. This command-line option ignores any property values typed at the command line. |
| **u** | Advertises to the current user. |
| **m** | Advertises to all users of this computer. |
| **/t** *TransformList* | Applies transform to advertised package and specifies what transform (.mst) files should be applied to the package. *TransformList* is a list of paths separated by semicolons. |

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Install a package using the command line

# Install a package using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a package using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /i** {*package* | *ProductCode*}

| Value | Description |
|---|---|
| **/i** | Installs a product. |
| *package* | Specifies the name of the Windows Installer package file. |
| *ProductCode* | Specifies the globally unique identifier (GUID) of the Windows Installer package. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Create an administrative installation package from the command line

# Remove a package using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a package using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /x** {*package* | *ProductCode*}

| Value | Description |
|---|---|
| **/x** | Uninstalls a product. |
| *package* | Specifies the name of the Windows Installer package file. |
| *ProductCode* | Specifies the globally unique identifier (GUID) of the Windows Installer package. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Repair a package using the command line

# Repair a package using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To repair a package using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec/f**[**p**][**o**][**e**][**d**][**c**][**a**][**u**][**m**][**s**][**v**] {*package* | *ProductCode*}

| Value | Description |
|---|---|
| **/fp** | Reinstalls only if the file is missing. |
| **/fo** | Reinstalls if the file is missing or if an older version is installed. |
| **/fe** | Reinstalls if the file is missing or an equal or older version is installed. |
| **/fd** | Reinstalls if the file is missing or a different version is installed. |
| **/fc** | Reinstalls if the file is missing or the stored checksum does not match the calculated value. |
| **/fa** | Forces all of the files to be reinstalled. |
| **/fu** | Rewrites all of the required user-specific registry entries. |
| **/fm** | Rewrites all of the required computer-specific registry entries. |
| **/fs** | Overwrites all of the existing shortcuts. |
| **/fv** | Runs from source and recaches the local package. |
| *package* | Specifies the name of the Windows Installer package file. |
| *ProductCode* | Specifies the globally unique identifier (GUID) of the Windows Installer package. |

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The default parameter for this command-line option is **/fpecms**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Remove a package using the command line

# Advertise a package using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To advertise a package using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /j**{**u** | **m**} *package*

| Value | Description |
|---|---|
| **/j** | Advertises a product. This command-line option ignores any property values entered on the command line. |
| **u** | Advertises to the current user. |
| **m** | Advertises to all users of this computer. |
| *package* | Specifies the name of the Windows Installer package file. |

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you want to install the application with elevated privileges, use the **/jm** command-line option.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Use Windows Installer from the command line

# Create an administrative installation package from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an administrative installation package from the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /a***package*

| Value | Description |
|---|---|
| **/a** | Applies the administrative installation option. |
| *package* | Specifies the name of the Windows Installer package file. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Install a package using the command line

# Create an installation log file using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an installation log file using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /i** *package***/L** [**i**][**w**][**e**][**a**][**r**][**u**][**c**][**m**][**p**][**v**][**+**][**!**][*****] *LogFile***.txt**

| Value | Description |
|---|---|
| **/i** | Installs a product. |
| *package* | Specifies the name of the Windows Installer package. |
| **/L** | Indicates to log information. |
| **/Li** | Logs status messages. |
| **/Lw** | Logs nonfatal warnings. |
| **/Le** | Logs all error messages. |
| **/La** | Logs start up of actions. |
| **/Lr** | Logs action-specific records. |
| **/Lu** | Logs user requests. |
| **/Lc** | Logs initial user interface parameters. |
| **/Lm** | Out-of-memory. |
| **/Lp** | Logs Terminal properties. |
| **/Lv** | Specifies verbose output. |
| **/L+** | Appends log to the existing file. |
| **/L!** | Flushes each line to the log. |
| **/L*** | Wildcard. Logs all information except for the **v** command-line option. |
| *LogFile***.txt** | Specifies the name and path of the text log file. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

- To include the **v** option in a log file using the wildcard flag, type **/L*v** at the command prompt.

- The Windows Installer log file options can also be used with the uninstall and repair processes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Advertise a package using the command line

# Configure Group Policy for Windows Installer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Group Policy for Windows Installer**

- Configure Windows Installer Group Policy computer options
- Configure Windows Installer Group Policy user options

# Configure Windows Installer Group Policy computer options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Windows Installer Group Policy computer options**

1. Open Group Policy.
2. In the console tree, click the **Windows Installer** node.

   **Where?**

   - *applicable policy name*/Computer Configuration/Administrative Templates/Windows Components/Windows Installer

3. In the details window, double-click one of the following policy settings:

   - **Disable Windows Installer**
   - **Always install with elevated privileges**
   - **Prohibit rollback**
   - **Remove browse dialog box for new source**
   - **Prohibit patching**
   - **Disable IE security prompt for Windows Installer scripts**
   - **Enable user control over installs**
   - **Enable user to browse for source while elevated**
   - **Enable user to use media source while elevated**
   - **Enable user to patch elevated products**
   - **Allow admin to install from Terminal Services session**
   - **Cache transforms in secure location on workstation**
   - **Logging**
   - **Prohibit User Installs**
   - **Turn off creation of System Restore Checkpoints**
   - **Prohibit removal of updates**
   - **Enforce upgrade components rule**
   - **Prohibit non-administrators from applying vendor signed updates**

- **Baseline file cache maximum size**

4. In the **Group Policy Property** dialog box, set the options you want to use, and then click **OK**.

### Notes

- To open Group Policy to edit the local Group Policy object, click **Start**, click **Run**, type **gpedit.msc**, and then press ENTER.
- For more information about Windows Installer Group Policy options, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Managing options for computers through Group Policy
Using Windows Installer

# Configure Windows Installer Group Policy user options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Windows Installer Group Policy user options**

1. Open Group Policy.
2. In the console tree, click the **Windows Installer** node.

   **Where?**

   - *applicable policy name*/User Configuration/Administrative Templates/Windows Components/Windows Installer

3. In the details window, double-click one of the following policy settings:

   - **Always install with elevated privileges**
   - **Search order**
   - **Prohibit rollback**
   - **Prevent removable media source for any install**

4. In the **Group Policy Property** dialog box, set the options you want to use, and then click **OK**.

## Notes

- To open Group Policy to edit the local Group Policy object, click **Start**, click **Run**, type **gpedit.msc**, and then press ENTER.
- For more information about Windows Installer Group Policy options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Managing options for users through Group Policy
Using Windows Installer

# Windows Installer Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Windows Installer:

- Windows Installer overview
- Understanding Windows Installer
- Using Windows Installer

# Windows Installer overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer overview**

Windows Installer is a component of the Windows operating system that simplifies the application installation process.

Windows Installer manages the installation and removal of applications by applying a set of centrally defined setup rules during the installation process. You can also use this service to modify, repair, or remove an existing application. Windows Installer technology consists of the Windows Installer service for the Windows operating systems and the package (.msi) file format that contains information about application set up and installation.

Windows Installer is also an extensible software management system. It manages the installation, addition, and deletion of software components, monitors file resiliency, and maintains basic failure recovery with rollbacks. Windows Installer supports installing and running software from multiple sources, and developers can customize Windows Installer to install custom applications.

# Understanding Windows Installer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Windows Installer**
This section covers:

- Windows Installer features
- Windows Installer technologies

# Windows Installer features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer features**

With Windows Installer and the .msi package file format, software installation and removal are more reliable and resilient. Windows Installer also provides a larger set of installation options. Windows Installer performs the following tasks:

- **Restores original computer state upon installation failure--**Windows Installer keeps track of all changes made to the system during the application installation process. If the installation fails, Windows Installer can restore, or roll back, the system to its initial state.

- **Helps prevent interapplication conflicts--**Windows Installer enforces installation rules that help to prevent conflicts with shared resources between existing applications. Such conflicts can be caused when an install operation makes updates to a dynamic link library (.dll) shared by an existing application or when an operation deletes a .dll shared by another application.

- **Reliably removes existing programs--**Windows Installer can reliably uninstall any program it previously installed. It removes all of the associated registry entries and application files, except for those shared by other installed software. You can uninstall an application at any time after a successful installation. (Removal should not be confused with rollback, which restores a computer to its initial state when an installation failure has occurred.)

- **Diagnoses and repairs corrupted applications--**An application can query Windows Installer to determine whether an installed application has missing or corrupted files. If any are detected, Windows Installer repairs the application by recopying only those files found to be missing or corrupted.

- **Supports on-demand installation of application features--**Windows Installer can be configured to initially install a minimal subset of an application. Later, additional components can be automatically installed the first time the user accesses features that require those components. This is known as advertising. For example, Windows Installer could install Microsoft Word with a minimal set of features. The first time the user tried to access a mail merge function (not included with the original installation), Windows Installer would automatically install the mail merge component. Similarly, Windows Installer can also purge components that go unused in an application. For example, Windows Installer could be configured to remove the mail merge component if it goes unused for 60 days.

- **Supports unattended application installation--**Installation packages can be configured to require no installation process interaction from the user. During the installation process, Windows Installer can query the computer for desktop attributes, including determining whether any applications were previously installed by Windows Installer.

- **Supports 32-bit and 64-bit applications--**32-bit applications can be installed on 64-bit machines.

- **Supports Microsoft .NET Framework--**The Microsoft .NET Framework is a new platform for building integrated, service-oriented applications that gather information from and interact with a wide variety of sources, regardless of the platforms or languages in use. The .NET Framework and the common language runtime can deliver write-once, compile-once, run-anywhere application development. Specifically, the .NET Framework delivers code reuse, code specialization, resource management, multilanguage development, security, deployment, and administration.

- **Integrated with side-by-side components--**This feature eliminates .dll version conflicts by permitting an application to be bound to the version of the component it was designed and tested with, regardless of the computer that hosts the application. Side-by-side components support the simultaneous execution of multiple versions of each component.

- **Integrated with software restriction policies--**This feature provides virus protection support, including protection from Trojan horse viruses and worms propagated through e-mail and the Web. Software restriction policies make the simple, point-and-click, active context user experience safe. Group Policy implements them as part of the list of trusted applications, and Windows Installer operates with applications permitted by these software restriction policies.

# Windows Installer technologies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer technologies**

Windows Installer technologies are divided into two parts that work in combination: a client-side installer service (Msiexec.exe) and a package file (.msi file). Windows Installer uses the information contained within a package file to install the application. Following are the Windows Installer technologies:

- **Installer service--**Windows Installer is an operating system service that allows the operating system to manage the installation process.

- **Installer program--**The Msiexec.exe program is a component of Windows Installer. This program uses a dynamic link library, Msi.dll, to read the package files (.msi), apply transforms (.mst), and incorporate command-line options. The installer performs all installation-related tasks: copying files onto the hard disk, making registry modifications, creating shortcuts on the desktop, and displaying dialog boxes to query user installation preferences when necessary.

  When Windows Installer is installed on a computer, the file association capabilities of the operating system are modified to recognize the .msi file type. When a user double-clicks a file with the .msi extension, the operating system associates the .msi file with Windows Installer and runs the Msiexec.exe application.

- **Installation package file--**Each package (.msi) file contains a relational type database that stores all the instructions and data required to install (and uninstall) the program across many installation scenarios. For example, a package file could contain instructions for installing an application when a prior version of the application is already installed. The package file could also contain instructions for installing the software on a computer where that application has never been present.

- **Transforms--**The installation process can be manipulated by applying transforms (.mst) to the installation database. A transform makes changes to elements of the database. For example, Windows Installer can use a transform file to change the language in the user interface of an application. The Windows Installer transform files modify the installation package file at installation time, and can therefore dynamically affect the installation behavior.

  If you are using Group Policy Software Installation, you should note that a transform is applied at the time of assignment or publication, in addition to the time of installation. In practical terms, this means that you should ensure the **Modifications** tab of the package properties dialog box is set up the way you want it to be before clicking **OK**. If you neglect to do this and deploy an incorrectly transformed package, then you either have to remove the software and redeploy it, or upgrade the software with a correctly transformed version.

  Customization transforms, much like patches, remain cached on the computer. These transforms are applied to the base package file whenever Windows Installer needs to perform a configuration change to the installation package. Transforms are applied only at initial installation; they cannot be applied to an application that has already been installed.

# Using Windows Installer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Windows Installer**
This section covers:

- Windows Installer in managed environments
- Msiexec (command-line options)
- Managing options for computers through Group Policy
- Managing options for users through Group Policy

# Windows Installer in managed environments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer in managed environments**

Windows Installer is an integral part of IntelliMirror, and a core component of the Group Policy-based change and configuration management technology. By using the IntelliMirror technologies, Group Policy, and change and configuration management, administrators can approve certain applications, specifying that all configuration operations on those applications (installation, removal, and repair) run as the local system account. The administrator has control and management of the file system and registry, and Windows Installer performs user-initiated software installations. Only applications that are approved by the administrator run with elevated privileges. For more information about Windows Installer, see Using Windows Installer.

Administrators use Active Directory, IntelliMirror, and Group Policy to assign and publish applications to groups of users or computers within the enterprise. Active Directory is a secure, distributed, partitioned, and replicated directory service that provides management services. These services include a standardized method of locating resources within the enterprise, and applying Group Policy to objects managed by Active Directory. For more information about Active Directory, see Active Directory.

# Msiexec (command-line options)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Msiexec**
Provides the means to install, modify, and perform operations on Windows Installer from the command line.

- To install or configure a product
- To use the administrative installation option
- To repair a product
- To uninstall a product
- To advertise a product
- To set logging level
- To apply an update
- To install a transform using the command line
- To install and advertise a product using a transform with the command line
- To set the user interface level
- To display copyright information for the Windows Installer
- To call the system API DllRegisterServer to self-register modules passed on the command line
- To call the system API DllUnRegisterServer to unregister modules passed on the command line

**To install or configure a product**
**Syntax**
**msiexec /i** {*Package* | *ProductCode*}

**Parameters**
**/i**
Installs or configures a product.

***Package***
Specifies the name of the Windows Installer package file.

***ProductCode***
Specifies the globally unique identifier (GUID) of the Windows Installer package.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To install a product from A:\Example.msi, type:

**msiexec /i A:\Example.msi**

**To use the administrative installation option**
**Syntax**
**msiexec /a** *Package*

**Parameters**
**/a**
Applies the administrative installation option.

***Package***
The name of the Windows Installer package file.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**To repair a product**
**Syntax**
**msiexec /f** [**p**][**o**][**e**][**d**][**c**][**a**][**u**][**m**][**s**][**v**]{*Package* | *ProductCode*}

**Parameters**
**/f**
Enables one or more of the command-line options listed in the following table.

| Command | Description |
|---|---|
| **p** | Reinstalls only if file is missing. |
| **o** | Reinstalls if file is missing or if an older version is installed. |
| **e** | Reinstalls if file is missing or an equal or older version is installed. |
| **d** | Reinstalls if file is missing or a different version is installed. |
| **c** | Reinstalls if file is missing or the stored checksum does not match the calculated value. |
| **a** | Forces all files to be reinstalled. |
| **u** | Rewrite all required user-specific registry entries. |
| **m** | Rewrites all required computer-specific registry entries. |
| **s** | Overwrites all existing shortcuts. |
| **v** | Runs from source and re-caches the local package. |

***Package***
Name of the Windows Installer package file.

***ProductCode***
Globally unique identifier (GUID) of the Windows Installer package.

**Remarks**

- This option ignores any property values that you type at the command line.
- The default parameter for this command-line option is **/fpecms**.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To repair the installation package, type:

**msiexec /fpecms Example.msi**

**To uninstall a product**
**Syntax**
**msiexec /x** {*Package* | *ProductCode*}

**Parameters**
**/x**
uninstalls a product.

***Package***
Name of the Windows Installer package file.

***ProductCode***
Globally unique identifier (GUID) of the Windows Installer package.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To remove or uninstall a package, type:

**msiexec /x Example.msi**

**To advertise a product**
**Syntax**
**msiexec /j** [{**u** | **m**}] *Package*

**msiexec** {**u** | **m**} *Package* **/t** *TransformList*

**msiexec** {**u** | **m**} *Package* **/g** *LanguageID*

**Parameters**
**/j**
Advertises a product.

**u**
Advertises to the current user.

**m**
Advertises to all users of the computer.

***Package***
Specifies the Windows Installer package file.

**/g *LanguageID***
Identifies the language.

**/t *TransformList***
Applies transform to advertised package.

**Remarks**

- This option ignores any property values that you type at the command line.
- If you want to install the application with elevated privileges, use **/jm**.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To advertise a package to all users of this computer, type:

**msiexec /jm Example.msi**

**To set logging level**
**Syntax**
**msiexec /L** [**i**][**w**][**e**][**a**][**r**][**u**][**c**][**m**][**p**][**v**][+][**!**]*LogFile***.txt**

**Parameters**
**/L**
Specifies the path to the log file.

**i**
Logs status messages.

**w**
Logs nonfatal warnings.

**e**
Logs all error messages.

**a**
Logs startup of actions.

**r**
Logs action-specific records.

**u**
Logs user requests.

**c**
Logs initial user interface parameters.

**m**
Logs out-of-memory.

**p**
Logs terminal properties.

**v**
Logs verbose output. To use **v**, specify **/L*v**.

**+**
Appends to existing file.

**!**
Flushes each line to the log.

*****
Logs all information except for the **v** option. This is a wildcard.

***LogFile* .txt**
Name and path of the text log file.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- To include the **v** option in a log file using the wildcard flag, type **/L*v** at the command prompt.
- The Windows Installer log file options can also be used with the uninstall and repair processes.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To install a package and create a log file that contains the information related to the status, out-of-memory, and error messages, type:

**msiexec /i Example.msi /Lime logfile.txt**

**To apply an update**
**Syntax**
**msiexec /p** *UpdatePackage*

**Parameters**
**/p**
Applies an update.

***UpdatePackage***
Specific update.

**Remarks**

- To apply an update to an administrative installation package, use the following syntax:

  **msiexec/p***UpdatePackage***/a***Example.msi*

- Windows Installer command-line options are not case-sensitive.

- For more information about the Windows Installer command-line options, see Related Topics.

**To install a transform using the command line**
**Syntax**
**msiexec /i** *Package* **TRANSFORMS**=*TransformList*

**Parameters**
**/i**
Installs or configures a product.

***Package***
Specifies the Windows Installer package file.

**TRANSFORMS=**
Property that is used to specify what transform (.mst) files should be applied to the package.

***TransformList***
List of paths separated by semicolons.

**Remarks**

- Windows Installer command-line options are not case-sensitive.

- For more information about the Windows Installer command-line options, see Related Topics.

**To install and advertise a product using a transform with the command line**
**Syntax**
**msiexec /i** *Package* **/j**[**u**][**m**] **/t***TransformList*

**Parameters**
**/i**
Installs or configures a product.

***Package***
Name of the Windows Installer package file.

**/j**
Advertises a product. This option ignores any property values entered on the command line.

**u**

Advertises to the current user.

**m**
Advertises to all users of this computer.

**/t**
Applies transform to advertised package.

***TransformList***
List of paths separated by semicolons.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**To set the user interface level**
**Syntax**
**msiexec /q**{**n** | **b** | **r** | **f** | **n+** | **b+** | **b-**}

**Parameters**
**/qn**
Displays no user interface.

**/qb**
Displays a basic user interface.

**/qr**
Displays a reduced user interface with a modal dialog box displayed at the end of the installation.

**/qf**
Displays the full user interface with a modal dialog box displayed at the end.

**/qn+**
Displays no user interface, except for a modal dialog box displayed at the end.

**/qb+**
Displays a basic user interface with a modal dialog box displayed at the end.

**/qb-**
Displays a basic user interface with no modal dialog boxes.

**Remarks**

- **/qb+-** is not a supported user interface level. The modal box is not displayed if the user cancels the installation.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To display the basic user interface options during the package installation of Example.msi, type:

**msiexec /qb Example.msi**

**To display copyright information for the Windows Installer**
**Syntax**
**msiexec** {**/?** | **/h**}

**Parameters**
**{ /?| /h}**
Displays the Windows Installer version and copyright information.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**To call the system API DllRegisterServer to self-register modules passed on the command line**
**Syntax**
**msiexec /y** *Module*

**Parameters**
**/y**
Calls the system API DllRegisterServer to self-register modules passed on the command line.

***Module***
Specifies the file name of module.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Remarks**

- This option is used only for registry information that cannot be added using the registry tables of the .msi file.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
The following example shows how you can use the **msiexec /y** command:

**msiexec /y my_file.dll**

**To call the system API DllUnRegisterServer to unregister modules passed on the command line**
**Syntax**
**msiexec /z** *Module*

**Parameters**
**/z**
Calls the system API DllUnRegisterServer to unregister modules passed on the command line.

***Module***
File name of module.

**Remarks**

- This option is used only for registry information that cannot be added using the registry tables of the .msi file.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
The following example shows how you can use the **msiexec /z** command:

**msiexec /z my_file.dll**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |

| | |
|---|---|
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Install a package using the command line
Install a transform using the command line
Install and advertise a product using a transform at the command line
Create an administrative installation package from the command line
Repair a package using the command line
Remove a package using the command line
Advertise a package using the command line
Create an installation log file using the command line
Command-line reference A-Z
Command shell overview

# Managing options for computers through Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing options for computers through Group Policy**
Windows Installer can be configured by using Group Policy and Active Directory to manage the computer installation options.

The following table shows policies concerned with managing the Group Policy computer options for Windows Installer. To configure these policies, start Group Policy, then in the console tree, click the **Windows Installer** node.

**Where?**

- *applicable policy name*/Computer Configuration/Administrative Templates/Windows Installer

| Policy | Details |
|---|---|
| **Disable Windows Installer** | Disables or restricts the use of Windows Installer.<br><br>This policy can prevent users from installing software on their system, or permit users to install only those programs offered by an administrator.<br><br>If you enable this policy, you can use the options in the **Disable Windows Installer** dialog box to establish an installation policy.<br><br>• **Never--**Windows Installer is fully enabled. Users can install and upgrade software. Windows Installer is enabled by default in Windows.<br><br>• **For non-managed apps only--**Users can install only those programs that an administrator assigns (offers on the desktop) or publishes (adds to Add or Remove programs).<br><br>• **Always--**Windows Installer is disabled.<br><br>**Note**<br><br>• This policy affects Windows Installer only. It does not prevent users from using other methods to install and upgrade programs. |
| **Always install with elevated privileges** | Directs Windows Installer to use system permissions when it installs any program on the system.<br><br>This policy extends to all programs the elevated privileges usually reserved for programs that have been assigned to the user (offered on the desktop) or the computer (installed automatically), or made available in Add or Remove programs in Control Panel. This policy lets users install programs that require access to directories that the user might not have permission to view or change, including directories on highly restricted computers.<br><br>If you disable this policy, or do not configure it, the system applies the current user's permissions when it installs programs that are not distributed or offered by an administrator.<br><br>**Note**<br><br>• This policy appears in both the Computer Configuration and User Configuration folders. To make this policy effective, you must enable the policy in both folders. |

| | |
|---|---|
| **Prohibit rollback** | Prohibits Windows Installer from generating and saving the files it needs to reverse an interrupted or unsuccessful installation.<br><br>**Prohibit rollback** also prevents Windows Installer from recording the original state of the system and frequency of changes it makes during installation, and prevents Windows Installer from retaining files it intends to delete later. As a result, Windows Installer cannot restore the computer to its original state if the installation does not complete.<br><br>This policy is designed to reduce the amount of temporary disk space required to install programs. **Prohibit rollback** also prevents malicious users from interrupting an installation to gather data about the internal state of the computer or to search secure system files. However, because an incomplete installation can render the system or a program inoperable, do not use this policy unless essential.<br><br>This policy appears in both the Computer Configuration and User Configuration folders. To make this policy effective, you must enable the policy in both folders. |
| **Remove browse dialog box for new source** | Prevents users from searching for installation files when adding features or components to an installed program.<br><br>This policy disables the **Browse** button next to the **Use feature from** list in the Windows Installer dialog box. Users must select an installation file source from the **Use features from** list configured by the administrator.<br><br>This policy applies even when the installation is running in the user's security context.<br><br>If you disable this policy or do not configure it, the **Browse** button is enabled when an installation is running in the user's security context, but only administrators can browse when an installation is running with elevated system privileges, such as installations offered on the desktop or in Add or Remove programs.<br><br>This policy affects Windows Installer only. It does not prevent users from using other browsers, such as Windows Explorer or My Network Places, to search for installation files. |
| **Prohibit patching** | Prevents users from installing patches to upgrade their programs.<br><br>Patches are updates or upgrades that replace only those program files that have changed. Because patches can be easy vehicles for malicious programs, some installations prohibit their use.<br><br>This policy applies only to installations that run in the user's security context. By default, users who are not administrators cannot apply patches to installations that run with elevated system privileges, such as those offered on the desktop or in Add or Remove programs. |
| **Disable IE security prompt for Windows Installer scripts** | Allows Web-based programs to install software on the computer without notifying the user.<br><br>By default, when a script hosted by an Internet browser attempts to install a program on the system, the system warns the user and prompts the user to either allow or refuse the installation. This policy suppresses the warning and allows the installation to proceed without user intervention.<br><br>This policy is designed for enterprises that use Web-based tools to distribute programs to their employees. However, because it might pose a security risk, it should be applied cautiously. |
| **Enable user control over installs** | Permits users to change installation options that typically are available only to administrators.<br><br>This policy bypasses some of the security features of Windows Installer. It permits installations to complete that otherwise would be halted due to a security violation.<br><br>The security features of Windows Installer prevent users from changing installation options that are typically reserved for administrators, such as specifying the directory to which files are installed. If Windows Installer detects that an installation package has permitted the user to change a protected option, it stops the installation and displays an error message. These security features operate only when the installation program is running in a privileged security context in which the installation program has access to directories that are denied to the user.<br><br>This policy is designed for less restrictive environments. It can be used to circumvent errors in an installation program that prevent software from being installed. |

| | |
|---|---|
| **Enable user to browse for source while elevated** | Permits users to search for installation files during privileged installations.<br><br>This policy enables the **Browse** button on the **Use feature from** dialog box. Users can search for installation files, even when the installation program is running, with elevated system privileges. By default, only system administrators can browse during installations with elevated privileges, such as installations offered on the desktop or displayed in Add or Remove programs.<br><br>Because the installation is running with elevated system privileges, users can browse through directories that their own permissions would not allow.<br><br>This policy does not affect installations that run in the user's security context. For more information, see the **Remove browse dialog box for new source** policy. |
| **Enable user to use media source while elevated** | Permits users to install programs from removable media, such as floppy disks and CD-ROMs, during privileged installation.<br><br>This policy permits all users to install programs from removable media, even when the installation program is running with elevated system privileges. By default, users can install programs from removable media only when the installation runs in the user's security context. During privileged installations, such as those offered on the desktop or displayed in Add or Remove programs, only system administrators can install from removable media.<br><br>This policy does not affect installations that run in the user's security context. By default, users can install from removable media when the installation runs in their own security context.<br><br>For more information, see the **Prevent removable media source for any install** policy in Managing options for users through Group Policy. |
| **Enable user to patch elevated products** | Permits users to upgrade programs during privileged installations.<br><br>This policy permits all users to install patches, even when the installation program is running with elevated privileges. Patches are updates or upgrades that replace only those program files that have changed. Because patches can easily be vehicles for malicious programs, some installations prohibit their use.<br><br>By default, only system administrators can apply patches during installations with elevated privileges, such as installations offered on the desktop or displayed in Add or Remove programs.<br><br>This policy does not affect installations that run in the user's security context. By default, users can install patches to programs that run in their own security context. For more information about prohibiting patch installations, see the **prohibit patching** policy. |
| **Allow admin to install from Terminal Services session** | Permits Terminal Services administrators to install and configure programs remotely.<br><br>By default, system administrators can install programs only when they are logged on to the computer on which the program is being installed. This policy creates a special exception for computers running Terminal Services.<br><br>This policy affects system administrators only. Other users cannot install programs remotely. |
| **Cache transforms in secure location on workstation** | Copies of transform files are saved in a secure location on the local computer.<br><br>Transform files consist of instructions to modify or customize a program during installation. By default, Windows Installer stores transform files in the Application Data directory in the user's profile. When a user reinstalls, removes, or repairs an installation, the transform file is available, even if the user is on a different computer, or isn't connected to the network.<br><br>If you enable this policy, the transform file is saved in a secure location on the user's computer instead of in the user profile. Because Windows Installer requires the transform file to repeat an installation in which the transform file was used, the user must be using the same computer or be connected to the original or identical media to reinstall, remove, or repair the installation.<br><br>This policy is designed for enterprises that must take special precautions to prevent unauthorized or malicious editing of transform files. |

| | |
|---|---|
| **prohibit User Installs** | Allows you to configure user installs. To configure this setting, set it to enabled and use the drop-down list to select the behavior you want.<br><br>If this setting is not configured, or if the setting is enabled and **Allow User Installs** is selected, the installer allows and makes use of products that are installed per user, and products that are installed per computer. If the installer finds a per-user install of an application, this hides a per-computer installation of that same product.<br><br>If this setting is enabled and **Hide User Installs** is selected, the installer ignores per-user applications. This causes a per-computer installed application to be visible to users, even if those users have a per-user install of the product registered in their user profile.<br><br>If this setting is enabled and **prohibit User Installs** is selected, the installer prevents applications from being installed per user, and it ignores previously installed per-user applications. An attempt to perform a per-user installation causes the installer to display an error message and stop the installation. This setting is useful in environments where the administrator only wants per-computer applications installed, such as on a kiosk or a Windows Terminal Server. |
| **Turn off creation of System Restore Checkpoints** | Stops Windows Installer from generating System Restore checkpoints when installing applications.<br><br>System Restore enables users, in the event of a problem, to restore their computers to a previous state without losing personal data files. By default, Windows Installer automatically creates a System Restore checkpoint each time an application is installed, so that users can restore their computer to the state it was in before installing the application.<br><br>If you disable this setting or do not configure it, Windows Installer automatically creates a System Restore checkpoint each time an application is installed.<br><br>**Note**<br><br>• This setting only applies to Windows XP Professional. |
| **Logging** | Specifies the types of events that Windows Installer records in its transaction log for each installation. The log, msi.log, appears in the *systemroot*\Temp directory. The Windows Installer logging options are listed in the following table. To indicate that an event type is recorded, enter the letter representing the event type. You can enter the letters in any order, and list as many or as few event types as you want. To disable logging, delete all of the letters from the box. If you disable this policy, or do not configure it, Windows Installer logs the default event types, represented by the letters **iweap**. You can include all possible events in the installation log file by entering **iwearucmpvo**. However, doing so will add considerable time to the installation process. |

| Mode | Log entry |
|---|---|
| **i** | Status messages |
| **w** | Non-fatal warnings |
| **e** | All error messages |
| **a** | Startup of actions |
| **r** | Action-specific records |
| **u** | User requests |
| **c** | Initial user interface parameters |
| **m** | Out-of-memory |
| **p** | Terminal properties |
| **v** | Verbose output |
| **o** | Out of disk space messages |

# Managing options for users through Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing options for users through Group Policy**

Windows Installer can be configured by using Group Policy and Active Directory to manage the user installation options.

This table shows policies concerned with managing the Group Policy user options for Windows Installer. To configure these policies, start Group Policy, then in the console tree, click the **Windows Installer** node.

**Where?**

- *applicable policy name*/User Configuration/Administrative Templates/Windows Installer

| Policy | Details |
|---|---|
| **Always install with elevated privileges** | Directs Windows Installer to use system permissions when it installs any program on the system.<br><br>This policy extends to all programs the elevated privileges usually reserved for programs that have been assigned to the user (offered on the desktop) or the computer (installed automatically), or made available in Add or Remove Programs in Control Panel. This policy lets users install programs that require access to directories that the user might not have permission to view or change, including directories on highly restricted computers.<br><br>If you disable this policy, or do not configure it, the system applies the current user's permissions when it installs programs that are not distributed or offered by an administrator.<br><br>**Note**<br><br>• This policy appears both in the Computer Configuration and User Configuration folders. To make this policy effective, you must enable the policy in both folders. |
| **Search order** | Specifies the order in which Windows Installer searches for installation files.<br><br>By default, Windows Installer searches the network first, then removable media (floppy drive, CD-ROM, or DVD), and finally, the Internet (URL).<br><br>To change the search order, enable the policy, and then type the following letters for each file source in the order that you want Windows Installer to search. For example, if you want to search the network, the Internet, and then removable media, type **num**. To exclude a file source, omit or delete the letter representing that source type.<br><br>• **n** - Adds the network in the source file search.<br><br>• **m** - Adds the removable media in the source file search.<br><br>• **u** - Adds the Internet in the source file search. |

| | |
|---|---|
| **Prohibit rollback** | Prohibits Windows Installer from generating and saving the files it needs to reverse an interrupted or unsuccessful installation.<br><br>This policy prevents Windows Installer from recording the original state of the system and sequence of changes it makes during installation. It also prevents Windows Installer from retaining files it intends to delete later. As a result, Windows Installer cannot restore the computer to its original state if the installation does not complete.<br><br>This policy is designed to reduce the amount of temporary disk space required to install applications. It also prevents malicious users from interrupting an installation to gather data about internal state of the computer or to search secure system files. However, because an incomplete installation can render the system or an application inoperable, this policy should not be used unless it is essential.<br><br>**Note**<br><br>• This policy appears in both the Computer Configuration and User Configuration folders. To make this policy effective, you must enable the policy in both folders. |
| **Prevent removable media source for any install** | Prevents users from installing programs from removable media.<br><br>If a user tries to install a program from removable media, such as CD-ROMs, floppy disks, or DVDs, a message informs the user that the feature cannot be found.<br><br>This policy applies even when the installation is running in the user's security context.<br><br>If you disable or do not configure this policy, users can install from removable media when the installation is running in their own security context, but only system administrators can use removable media when an installation is running with elevated system privileges, such as installations offered on the desktop or in Add or Remove Programs. |

# Windows Installer Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The installation package does not install correctly.
- The installation process stops before completing the installation.

**The installation package does not install correctly.**
**Cause:** The installation packaged has been corrupted.

**Solution:** Repair the installation package by using the Windows Installer repair option.

**See also:** Repair a package using the command line

**The installation process stops before completing the installation.**
**Cause:** Windows Installer was not able to read the installation package, or conditions on your computer prevented the package from being installed successfully.

**Solution:** Windows Installer creates entries in the Windows event log. To view the log, do the following:

1. Open Event Viewer.
2. Review the application log for Windows Installer entries.

**Notes**

- To open Event Viewer, click **Start**, point to **Settings**, and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Event Viewer**.
- You can configure Windows Installer to create a custom installation log file.

**See also:** Create an installation log file using the command line

# Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services**
Remote Installation Services (RIS) is an optional component that is included in the Microsoft® Windows Server 2003™ family operating systems.

You can use RIS to create installation images of operating systems or of complete computer configurations, including desktop settings and applications. You can then make these installation images available to users at client computers. You can also specify which RIS server will provide installations to a given client computer, or you can allow any RIS server to provide the installation.

When you use RIS, a CD-ROM is not required for installation of an operating system on a client computer. The client computers must support remote booting with the Pre-Boot eXecution Environment (PXE) ROM, or they must be started with a remote-startup floppy disk and contain a network adapter that is supported by that disk.

You can use RIS to create automated installation images of Windows Server 2003 family operating systems, Windows XP, and Windows 2000.

- Before you begin using RIS, see Checklist: Installing Remote Installation Services.
- For tips about using RIS, see Best practices for Remote Installation Services.
- For help with specific tasks, see Remote Installation Services How To....
- For general background information, see Remote Installation Services Concepts.
- For problem-solving instructions, see Troubleshooting Remote Installation Services.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Checklist: Installing Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing Remote Installation Services**

| Step | Reference |
|---|---|
| Review key concepts. | Remote Installation Services Concepts |
| Ensure that both your server and client hardware meet the Remote Installation Services (RIS) hardware requirements. | Remote Installation Services system requirements |
| Ensure that your network is based on TCP/IP, and that a Domain Name System (DNS) server exists on the network. You do not need to use the Microsoft version of DNS. | Install a DNS server; Configure a new DNS server |
| Ensure that an authorized Dynamic Host Configuration Protocol (DHCP) server exists on the network. You do not need to use the Microsoft version of DHCP. A non-Microsoft DHCP server does not need to be authorized. | Install a DHCP server; Checklist: Installing a DHCP server |
| Ensure that Active Directory exists on the network. | Active Directory Overview; Using the Active Directory Installation Wizard |
| Run the Remote Installation Services Setup Wizard. | Install Remote Installation Services |

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Best practices for Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Follow best practices for network security in any network that includes PXE-enabled client computers.**

- For more information, see PXE architecture, RIS, and security considerations.

**Use the appropriate number of Remote Installation Services (RIS) servers on your network.**

- In a small local area network (LAN)--for example, one physical subnet without a router--a single RIS server can serve all Pre-Boot eXecution Environment (PXE) remote boot-enabled client computers up to the network bandwidth or server resource limitations.

  In a routed environment, use the settings that are available when you prestage clients to direct each client to be serviced by the RIS server that is in closest proximity to it. Also, set the Dynamic Host Configuration Protocol (DHCP) forwarding option to allow your routers to forward client requests to the RIS servers. In branch offices, if the only links to the branch site are low-speed links, use a RIS server at the branch site. Do not use RIS over low-speed links on a wide area network (WAN).

**Install RIS on a physical disk separate from the one on which the operating system is installed.**

- Installing RIS on a disk separate from the one on which the operating system is installed will help ensure optimal performance.

**Learn how to control the number of installation options and operating system choices a user can access within the Client Installation Wizard.**

- Restricting installation options increases the number of successful operating system installations that can be completed without assistance from technical support or administrative staff. By default, RIS presents one installation option and operating system option to the user. For more information about other installation options, see Installation options available through RIS.

  To maintain appropriate security and simplify the client installation process for users, review information in the *Windows Server 2003 Deployment Kit* regarding settings in .osc and .sif files. For details, see "Designing RIS Installations" at the Microsoft Windows Resource Kits Web site.

  For information about the Setup Manager tool and .sif files, see deploy.chm in \Support\Tools\deploy.cab on the Windows CD.

**Use the Remote Installation Preparation Wizard (RIPrep) image format to deploy a standard corporate desktop configuration across different types of client hardware throughout your organization.**

- With the Remote Installation Preparation image format, administrators can replicate the installation image of an existing client computer, including locally installed applications and operating system configuration changes, to an available RIS server on the network. After the installation image is replicated, it can be installed remotely by any supported client computer. When you create new images using RIPrep, the image is available only to clients that have the same type of Hardware Abstraction Layer (HAL). Images that were created using RIPrep are available to all clients that are authorized to view them. For more information, see Creating an installation image with RIPrep.

**Use RIS with computers that contain the PXE-based remote boot ROM.**

- For computers that do not contain the PXE-based remote boot ROM, use the remote boot floppy generator (Rbfg.exe) to

create a floppy disk. You can then use the RIS feature for these computers. For more information on the PXE-based remote boot ROM, see PXE architecture.

**When using the Client Installation Wizard, use standard ASCII characters for user name, password, and domain name information.**

- The Client Installation Wizard does not support extended ASCII character sets (such as those containing ü é and other non-standard characters). You should therefore limit the characters used for user name, password, and domain name information to standard ASCII characters (OEM characters 32-126).

### Note

- This topic does not apply to Windows Server 2003, Web Edition.

# Remote Installation Services How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install Remote Installation Services
- Manage Remote Installation Services
- Add RIS Client Installation Images
- Manage RIS Client Computers
- Manage Security for Remote Installation Services

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Install Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Remote Installation Services**

1. If you have not already installed Remote Installation Services through Control Panel or through an unattended Setup answer file, open Add or Remove Programs and use **Add/Remove Windows Components** to install the RIS component.

2. Open the Remote Installation Services Setup Wizard.

3. In the wizard, click **Next**.

You are prompted for the following information:

**Remote Installation Services Drive and Directory**
Enter the disk drive and directory where you want to install Remote Installation Services (RIS). The disk should be dedicated to the RIS server and contain enough space to store as many client images as you plan to host with this server. The recommended minimum amount of space is four gigabytes.

**Initial Settings**
Select **Respond to client computers requesting service** if you want the RIS server to begin responding to client computers as soon as you finish the wizard. Select **Do not respond to unknown client computers** if you want the RIS server to respond only to prestaged client computers in Active Directory. You can change these settings later.

**Source Path**
Enter the location of the client images. This can be either the compact disc or a shared folder that contains the installation files.

**Friendly Description and Help Text**
Enter the friendly description for the client computer installation. This description will be displayed to users or clients of this server. The Help text describes the operating system installation choices to users or clients of RIS.

**Notes**

- To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**.

- This topic does not apply to Windows Server 2003, Web Edition.

- To start the Remote Installation Services Setup Wizard, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Remote Installation Services Setup**.

- The Installation Wizard might prompt you to restart the computer if RIS was added after Setup.

- For information about the technologies that RIS requires to identify network computers (for example, DHCP and other technologies), see Related Topics.

- For information about RIS system requirements, including important information about requirements for the partition on which you install RIS, see Related Topics.

- To manage a RIS server from a computer running Windows XP, install the Windows Server 2003 Administration Tools Pack on the computer running Windows XP. To manage a RIS server from any remote computer, use Remote Desktop

for Administration. For more information, see Related Topics.

- Do not use Network Load Balancing in front of the RIS server. If you do, RIS will not work. This is because the computer will appear to Active Directory as if it has two Internet Protocol (IP) addresses and is therefore a multihomed computer. Similarly, a computer with more than one network adapter is also multihomed, and cannot be used as a RIS server.

- Installation images that are stored in locations other than *Server_name*\REMINST\Setup*Language*\Images

  and referenced through junction points cannot be used with the Single Instance Store (SIS) because the SIS groveler agent does not work with junction points.

- RIS does not currently support the Encrypting File System (EFS). The Distributed File System (DFS) is not supported as a target location. You can, however, run RIS on a computer that is also running DFS.

- To ensure that all of the services have started successfully, open Event Viewer and review the entries for System Log and Application Log. For more information about using Event Viewer, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configure Remote Installation Services
View an event log
Using the Event Viewer snap-in
Event Viewer overview
Running the Remote Installation Services Setup Wizard
Technologies used by Remote Installation Services to identify network computers
Remote Installation Services system requirements
Windows Server 2003 Administration Tools Pack
Using Remote Desktop for Administration for remote server administration

# Manage Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage Remote Installation Services**

- Configure Remote Installation Services
- Configure settings for new clients for RIS
- Manage client installation images
- Manage troubleshooting tools
- Verify the Remote Installation Services configuration

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Configure Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Remote Installation Services**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the applicable Remote Installation Services (RIS) server.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable RIS server*

3. Click **Properties**, and then in the **Properties** dialog box, click the **Remote Install** tab.
4. In the **Remote Install** dialog box, select or clear check boxes as needed:

   - **Respond to client computers requesting service**

     Select this option to enable the RIS server to begin responding to client computers. Clear it to prevent the RIS server from responding.

   - **Do not respond to unknown client computers**

     Select this option if you want this RIS server to respond only to prestaged client computers in Active Directory. Clear it if you want this RIS server to respond to any client computer requesting remote installation.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- This topic does not apply to Windows Server 2003, Web Edition.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- For information about prestaging clients, see Related Topics.
- To manage a RIS server from a computer running Windows XP, install the Windows Server 2003 Administration Tools Pack on the computer running Windows XP. To manage a RIS server from any remote computer, use Remote Desktop for Administration. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files
Install Remote Installation Services
Configure settings for new clients for RIS
Remote Installation Services administration overview
Prestage client computers
Choosing appropriate group memberships for RIS administrators
Windows Server 2003 Administration Tools Pack
Using Remote Desktop for Administration for remote server administration
Create a shortcut using the runas command

# Configure settings for new clients for RIS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure settings for new clients for RIS**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the applicable Remote Installation Services (RIS) server.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable RIS server*

3. Click **Properties**. In the **Properties** dialog box, click the **Remote Install** tab, and then click **Advanced Settings**.
4. In the **Advanced Settings** dialog box, click the **New Clients** tab.
5. In **Generate client computer names using**, either select the client computer naming format you want to use, or click **Customize** to create a client computer naming format.
6. To specify where to create the client computer account, click one of the following options:

   - **Default directory service location**
   - **Same location as that of the user setting up the client computer**
   - **The following directory service location**

7. If you chose the last option, click **Browse** and specify where to create the computer accounts.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- This topic does not apply to Windows Server 2003, Web Edition.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- This procedure applies only to client computer accounts that are not prestaged.
- After performing this procedure, be sure to assign the appropriate permissions to users who install from RIS on the client computer. This is necessary so that users can create computer accounts within the chosen domain or organizational unit. For more information, see Related Topics.
- To manage a RIS server from a computer running Windows XP, install the Windows Server 2003 Administration Tools Pack on the computer running Windows XP. To manage a RIS server from any remote computer, use Remote Desktop for Administration. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Allow a user or group to create a RIS managed computer account in the domain
Windows Server 2003 Administration Tools Pack
Using Remote Desktop for Administration for remote server administration
Create a shortcut using the runas command

# Manage client installation images

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage client installation images**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the applicable Remote Installation Services (RIS) server.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable RIS server*

3. Click **Properties**.
4. In the **Properties** dialog box, click the **Remote Install** tab, and then click **Advanced Settings**.
5. In the **Remote Installation Services Properties** dialog box, click the **Images** tab.
6. In the **Images** dialog box, click the installation image or unattended setup answer file.
7. Click one of the following options:

   - **Add**
   - **Remove**
   - **Properties**
   - **Refresh**

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- This topic does not apply to Windows Server 2003, Web Edition.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- For information about adding a new client installation image, see Related Topics.
- To manage a RIS server from a computer running Windows XP, install the Windows Server 2003 Administration Tools Pack on the computer running Windows XP. To manage a RIS server from any remote computer, use Remote Desktop for Administration. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add RIS Client Installation Images
Advanced settings for installation images
Manage Security for Remote Installation Services
Windows Server 2003 Administration Tools Pack
Using Remote Desktop for Administration for remote server administration
Create a shortcut using the runas command

# Manage troubleshooting tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To manage troubleshooting tools

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the applicable Remote Installation Services (RIS) server.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable RIS server*

3. Click **Properties**, and then in the **Properties** dialog box, click the **Remote Install** tab.
4. Click **Advanced Settings**, and then click the **Tools** tab.
5. In the **Tools** dialog box, click **Tool**.
6. Click one of the following options:

   - **Remove**
   - **Properties**
   - **Refresh**

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- This topic does not apply to Windows Server 2003, Web Edition.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- You can use this procedure to view the list of tools or remove a tool from the list. For more information about maintenance and troubleshooting tools and options for RIS, see Related Topics.
- To manage a RIS server from a computer running Windows XP, install the Windows Server 2003 Administration Tools Pack on the computer running Windows XP. To manage a RIS server from any remote computer, use Remote Desktop for Administration. For more information, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files
Manage client installation images
Maintenance and troubleshooting tools and options
Choosing appropriate group memberships for RIS administrators
Windows Server 2003 Administration Tools Pack
Using Remote Desktop for Administration for remote server administration
Create a shortcut using the runas command

# Verify the Remote Installation Services configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify the Remote Installation Services configuration**

1. Make sure you have the installation CDs available. You might be prompted for them during the process of verifying the server.

2. Go to the Remote Installation Services (RIS) server you want to verify, or use Remote Desktop for Administration to connect to the RIS server you want to verify.

3. Open Active Directory Users and Computers.

4. In the console tree, right-click the RIS server whose configuration you want to verify.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable RIS server*

5. Click **Properties**, and then in the **Properties** dialog box, click the **Remote Install** tab.

6. Click **Verify Server**. The Check Server Wizard opens.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- This topic does not apply to Windows Server 2003, Web Edition.

- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.

- Use Verify Server if the server seems to be failing or is behaving inconsistently, or if you need to restore a RIS volume from backup.

- If you need to restore a RIS volume from backup, verify the server configuration before you restore the volume.

- You can only use Verify Server locally on a RIS server, or by connecting to that server through Remote Desktop for Administration. You cannot use Verify Server if you are performing RIS administration tasks from a remote server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Install Remote Installation Services

Configure settings for new clients for RIS
Technologies and components used by Remote Installation Services
Remote Desktop for Administration
Create a shortcut using the runas command

# Add RIS Client Installation Images

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Add RIS client installation images**

- Add a new client operating system installation image
- Associate an unattended setup answer file with a RIS installation image
- Create a Remote Installation Preparation Wizard image

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Add a new client operating system installation image

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a new client operating system installation image**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the applicable Remote Installation Services (RIS) server.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable RIS server*

3. Click **Properties**, and then, in the **Properties** dialog box, click the **Remote Install** tab.
4. Click **Advanced Settings**, and then click the **Images** tab.
5. Click **Add**.
6. Click **Add a new installation image**, and then click **Next** to start the Remote Installation Services Setup Wizard.
7. Click **Next**, and then type the location of the installation files. The location can be either a compact disc or a shared folder.
8. Type the name of the folder for the new installation image.
9. Enter the friendly description and Help text, and then click **Next**.
10. Specify whether you want to use new .osc screens and back up the previous .osc screens, use new .osc screens and overwrite the previous screens, or leave the previous .osc screens in place.
11. Review the installation summary, and then click **Finish**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- This topic does not apply to Windows Server 2003, Web Edition.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- After completing this procedure, you can associate an unattended setup answer file with the installation image. For more information, see Related Topics.
- For information about .osc screens and how they relate to installation options made available through RIS, see Related Topics.
- You cannot perform this procedure remotely. It must be performed locally on the RIS server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add RIS Client Installation Images
Set Group Policy for client installation options
Installation options available through RIS
Associate an unattended setup answer file with a RIS installation image
Create a shortcut using the runas command

# Associate an unattended setup answer file with a RIS installation image

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To associate an unattended setup answer file with a RIS installation image**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the applicable Remote Installation Services (RIS) server.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable RIS server*

3. Click **Properties**, and then in the **Properties** dialog box, click the **Remote Install** tab.
4. Click **Advanced Settings**, and then click the **Images** tab.
5. Click **Add**.
6. Click **Associate a new answer file to an existing image**, and then click **Next**.
7. Click one of the following options:

   - **Windows image sample files**
   - **Another remote installation server**
   - **An alternate location**

8. Select the installation image to which the answer file will be associated, and then click **Next**.
9. Select the answer file that you want to associate with the installation image, and then click **Next**.
10. Enter the friendly description and Help text, and then click **Next**.
11. Review the settings summary, and then click **Finish** to complete the process.

**Notes**

- To perform this procedure, you must have permission to write to the folder that contains the RIS folders and files (the default name for the folder is RemoteInstall). For more information about the permissions and group membership needed to perform this task and other tasks related to RIS administration, see Related Topics.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- This topic does not apply to Windows Server 2003, Web Edition.

- The installation images that appear in the list in the Add Wizard include flat images only. Flat images are images that were created with the Remote Installation Services Setup Wizard. They do not include images created with the Remote Installation Preparation (RIPrep) Wizard.

- For more information about installation options available through RIS, and the ways these options can be controlled through unattended setup answer files (.sif files), see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add a new client operating system installation image
Set permissions for administrators who manage client installation images for RIS
Manage Security for Remote Installation Services
Choosing appropriate group memberships for RIS administrators
Installation options available through RIS
Create a shortcut using the runas command

# Create a Remote Installation Preparation Wizard image

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a Remote Installation Preparation Wizard image**

1. Install the operating system on the client computer that you will use to create the installation image. You can install it using either Remote Installation Services (RIS) or the product CD. For more information, see Related Topics.

2. Install any additional applications and modify the local configuration settings of the source client computer.

   **Important**

   - Before starting the RIPrep utility, be sure to remove all confidential information from the source client computer.

3. Click **Start**, and then click **Run**.

4. Type the Universal Naming Convention (UNC) path of the RIPrep utility, and then click **OK**.

   Example:

   *Server_name*\REMINST\Admin\I386\RIPrep.exe

5. Read the information provided in the Remote Installation Preparation Wizard Welcome screen, and then click **Next**.

6. Type the server name to which this image will be copied, and then click **Next**. By default this is the current RIS server from which you are running the wizard.

7. Type the name of the folder on the RIS server to which this installation image will be copied, and then click **Next**.

8. Type the friendly description and the Help text, and then click **Next**. This information is displayed by the Client Installation Wizard when RIS clients request network services.

9. After you type the friendly description and the Help text, a screen might appear showing that multiple user profiles were detected, or that services are running that RIPrep determines should be stopped, but which it cannot stop. The screen lists the services that are running and should be stopped.

   If this screen appears, do one of the following:

   - It is recommended that you stop all programs or services on the source computer before proceeding. Review the list of programs or services that are currently running on the source computer, close any running applications, and then click **Next**.

   - If you choose not to stop the programs, skip the screen; to do so, click **Next**.

10. Review the settings summary, and then click **Next**.

11. Review the information from **Completing the Remote Installation Preparation Wizard**, and then, to replicate the source computer installation image onto the RIS server, click **Next**.

**Important**

- A 64-bit version of the Remote Installation Preparation (RIPrep) Wizard is included with the 64-bit version of Windows Server 2003, Enterprise Edition. For more information about using the RIPrep Wizard to create an installation image of a 64-bit operating system, see article 891128, "Updated Remote Installation Service (RIS) functionality in Windows Server 2003 Service Pack 1," in the Microsoft Knowledge Base.

- You cannot use RIPrep to create an image of a Windows 2000 operating system if Internet Information Services (IIS) is installed.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the computer where you create the configuration that RIPrep copies, or you must have been delegated the appropriate authority. Members of the Domain Admins group might also be able to perform this procedure. You must also have permission to write to the RIS folders on the RIS server (the default name for the folder that contains all the RIS folders is RemoteInstall). One way of ensuring that you have permission to write to the RIS folders on the RIS server is to be explicitly assigned the permission. Another way is to have membership in the local Administrators group on the RIS server. For more information, see Default local groups and Default groups.

- This topic does not apply to Windows Server 2003, Web Edition.

- The source computer shuts down when the image replication process is complete.

- The abbreviated Setup program runs automatically when you restart the source computer. Complete the setup process if you want to use this client computer to create another installation image.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Creating an installation image with RIPrep
Perform a network installation using RIS
Working with MMC console files
Associate an unattended setup answer file with a RIS installation image
Add a new client operating system installation image
Set permissions for administrators who manage client installation images for RIS
Choosing appropriate group memberships for RIS administrators
Operating systems supported by Remote Installation Services

# Manage RIS Client Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage RIS client computers**

- Prestage client computers
- Remove existing client computers
- Find RIS client computers
- Locate the GUID for client computers
- Perform a network installation using RIS
- Create a remote boot disk

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Prestage client computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prestage client computers**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the applicable organizational unit that will contain the new client computer.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain/Applicable organizational unit*

3. Click **New**, and then click **Computer**.
4. Type the client computer name, click **Next**, and then click **This is a managed computer**.
5. In the text entry field, type the client computer's globally unique identifier (GUID), and then click **Next**.
6. Click one of the following options to specify which server or servers will support this client computer:

   - **Any available Remote Installation Services (RIS) server**

     Selecting this option specifies that this client computer can be serviced by any RIS server.

   - **The following RIS server**

     Selecting this option designates a specific server to service this client computer.

**Notes**

- To select the sequence **New** and then **Computer**, you must have Create Computer Objects permissions for the Computers container in Active Directory. Membership in the Account Operators group provides these permissions. As a security best practice, instead of logging on as a member of Account Operators or an administrative group, consider using Run as to perform this procedure. For more information about permissions and group memberships required for tasks related to RIS administration, see Related Topics. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- This topic does not apply to Windows Server 2003, Web Edition.
- Make sure you have the client computer GUID before you start this procedure. For information about how to locate the GUID, see Related Topics.
- You can predetermine how a specific client computer network account is identified, by prestaging the client computer account object within Active Directory. This information is used to identify and route the client computer during a RIS installation.
- You can use prestaging as one part of a security strategy for RIS because with prestaging, you can prevent unknown client computers (computers that are not prestaged) from being installed through RIS. For information about other

security considerations, related to PXE architecture and RIS, see Related Topics.

- You must also set the permissions required by the users of the prestaged client computer. For more information on setting these permissions, see Related Topics.

- When a client computer is being prestaged into a domain with multiple domain controllers, the replication delay of the client computer account information can cause a client computer to be serviced by another RIS server. For more information on controlling information replication, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configure settings for new clients for RIS
Find RIS client computers
Locate the GUID for client computers
Set permissions required by RIS users who use prestaged client computers
Domain controllers
Force replication over a connection
Configure site link replication frequency
PXE architecture, RIS, and security considerations
Choosing appropriate group memberships for RIS administrators
Create a shortcut using the runas command

# Remove existing client computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove existing client computers**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the applicable client computer account.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable client computer account*

3. Click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- This topic does not apply to Windows Server 2003, Web Edition.

- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Prestage client computers
Create a shortcut using the runas command

# Find RIS client computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find RIS client computers**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the applicable Remote Installation Services (RIS) server.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable RIS server*

3. Click **Properties**, and then in the **Properties** dialog box, click the **Remote Install** tab.
4. In the **Remote Install** dialog box, click **Show Clients**.
5. Scan the list for the appropriate client computer. If you know the globally unique identifier (GUID) of the client computer, you can also type the GUID and then click **Find Now**.

## Notes

- To perform this task, the only access required is read access to the domain. Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- This topic does not apply to Windows Server 2003, Web Edition.
- The Show Clients feature searches for all client computers that are prestaged for this RIS server.
- When using the Show Clients feature in multiple RIS server environments, the search results might include client computers from multiple servers that have similar names. For example, if you have multiple RIS servers with computer names such as RISsvr1, RISsvr10, and RISsvr100, and you search for RISsvr1, the list will display the client computers that install from all three servers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Locate the GUID for client computers
RIS client computer properties
Create a shortcut using the runas command

# Locate the GUID for client computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To locate the GUID for client computers**

- The computer's globally unique identifier (GUID) (sometimes called the universally unique identifier [UUID]) appears in the following areas:
    - A label on the side of the computer case.
    - A label within the computer case.
    - The basic input/output system (BIOS) of the client computer. You can collect this information through the Windows Management Instrumentation (WMI) interface by using the universally unique identifier (UUID) value of the Win32_ComputerSystemProduct class. For more information about this, see the MSDN Library at the Microsoft Web site.
    - Network traffic from the client computer. You can use a network utility to sniff the network traffic and locate the DHCPDiscover packet; within that field will be the 128-bit, 16-byte GUID/UUID.

## Notes

- The computer's GUID is supplied by the manufacturer. For more information on the computer's GUID, see "RIS client computer properties" in Related Topics.
- This topic does not apply to Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find RIS client computers
RIS client computer properties

# Perform a network installation using RIS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To perform a network installation using RIS**

1. After the Client Installation Wizard starts, press ENTER to continue the process.

2. Enter the user name, password, and domain name into the appropriate fields, pressing TAB to move from one field to the next.

3. Several options are available through the Client Installation Wizard; Automatic Setup is the default display option. The other three options are available only if you have changed the Default Domain Policy in Group Policy to allow them.

   - **Automatic Setup**

     This option predefines the computer name and a location in a container within Active Directory for the client computer accounts.

   - **Custom Setup**

     With this option selected, users can define a unique name for this computer and specify where the computer account will be created within Active Directory.

   - **Restart Setup**

     This option restarts an operating system installation attempt if it fails prior to completion.

   - **Maintenance and Troubleshooting**

     With this option selected, users can access tools from the Client Installation Wizard.

4. Follow any additional instructions that appear on the screen.

**Important**

- The RIS administrator can control some of the choices and screens seen by users of RIS client computers. For more information about installation options available through RIS, see Related Topics.

**Notes**

- A user who performs this procedure must first have a valid user name, password, and domain name. The user must also have appropriate permissions. For example, if the computer account has not been prestaged, the user must have permission to create the computer account within the appropriate domain or organizational unit. If the computer account has been prestaged, other permissions must be set. For more information, see Related Topics.

- This topic does not apply to Windows Server 2003, Web Edition.

- It is recommended that a user logging on to a RIS client computer type the user name, password, and domain name in the respective fields. Users should not combine user name and domain name with the @ symbol in between.

- A network installation using RIS can be initiated by either the user of a Pre-Boot eXecution Environment (PXE) remote boot-enabled client computer or by a special remote boot disk provided for pre-Net PC/PC98 computers.

- To invoke a network service boot (network installation using RIS) on client computers that are not remote-boot enabled, insert the boot disk into the disk drive and start the client computer. The client computer must have a Peripheral Component Interconnect (PCI)-based network adapter from the list of specific adapters supported. To determine which network adapters RIS supports, use the boot disk–generating utility. For more information, see the topic about creating a remote boot disk in Related Topics.

- The client computer's basic input/output system (BIOS) must be configured to use the network adapter as its primary boot device.

- When running the Client Installation Wizard, use only standard ASCII characters (OEM characters 32-126) for user name, password, or domain name information. The Client Installation Wizard does not support extended ASCII character sets (such as those containing ü, é, and other extended characters).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Set Group Policy for client installation options
Controlling client installations through Group Policy
Set permissions required by RIS users who use prestaged client computers
Allow a user or group to create a RIS managed computer account in the domain
Create a remote boot disk
Installation options available through RIS

# Create a remote boot disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a remote boot disk**

1. Click **Start**, and then click **Run**.
2. Type the Universal Naming Convention (UNC) path of the Rbfg.exe utility, and then click **OK**.

   Example:

   *Server_name*\REMINST\Admin\I386\RBFG.exe
3. Insert a formatted disk into the disk drive.
4. Click **Create Disk**.
5. When the disk is ready, click **Close**, and then remove the disk from the disk drive.

### Important

- Use the remote boot disk on client computers that do not have a remote boot-enabled ROM. A computer on which you use the remote boot disk must have a Peripheral Component Interconnect (PCI)-based network adapter supported by the Rbfg.exe utility. PC Card adapters, CardBus adapters, and adapters that are not PCI-based are not supported. You can view the list of supported network adapters by starting the Rbfg.exe utility (as described in this procedure) and clicking **Adapter List**.
- Some network adapters supported by the Rbfg.exe utility might not be supported by the operating system you plan to install on the client computer. For more details, see the documentation for that operating system. To determine which network adapters are supported by the Windows Server 2003 family operating systems, see Support Resources in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

### Note

- This topic does not apply to Windows Server 2003, Web Edition.

**See Also**
**Concepts**
Working with MMC console files
Perform a network installation using RIS
Remote Installation Services boot disk

# Manage Security for Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage security for Remote Installation Services**

- Set permissions required by RIS users who use prestaged client computers
- Allow a user or group to create a RIS managed computer account in the domain
- Set Group Policy for client installation options
- Allow or prevent the installing of a RIS image by a user or group
- Allow or prevent the viewing and installing of a RIS image by a user or group
- Set permissions for administrators who manage client installation images for RIS

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Set permissions required by RIS users who use prestaged client computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions required by RIS users who use prestaged client computers**

1. Open Active Directory Users and Computers.
2. On the **View** menu, ensure that **Advanced Features** is checked.
3. In the console tree, right-click the applicable client computer account.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable client computer account*

4. Click **Properties**.
5. In the **Properties** dialog box, click the **Security** tab, and then click **Add**.
6. In the **Select Users, Computers or Groups** dialog box, type the name of the user or group that will use Remote Installation Services (RIS) to install the prestaged client computer, and then click **OK**.
7. Click the user or group you have added.
8. Under **Allow**, ensure that **Read** and **Reset password** are selected, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- This topic does not apply to Windows Server 2003, Web Edition.
- If you set these permissions for a group, remember to add users to that group.
- For client computer accounts that are prestaged in another Active Directory folder location, expand the Active Directory Users and Computer console and select the appropriate client computer accounts.
- This procedure pertains to prestaged client computers. For client computers that are not prestaged, the user must be able to create a computer account in the domain. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files
Allow a user or group to create a RIS managed computer account in the domain
Managing Active Directory from MMC
Prestage client computers
Choosing appropriate group memberships for RIS administrators
Create a shortcut using the runas command

**Other Resources**

Managing the Active Directory schema from MMC

# Allow a user or group to create a RIS managed computer account in the domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow a user or group to create a RIS managed computer account in the domain**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the applicable domain.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*

3. Click **Delegate control**.
4. In the Delegation of Control Wizard, click **Next**.
5. In the **Users or Groups** dialog box, click **Add**.
6. In the **Select Users, Computers, or Groups** dialog box, type the name of the **User** account or the **Security Group** (preferred) that contains the users who create Remote Installation Services (RIS) managed computer accounts in the domain (that is, who install from RIS on client computers that are not prestaged).
7. Click **OK**, and then click **Next**.
8. In **Tasks to Delegate**, click **Join a computer to the domain**, and then click **Next**.
9. Review the delegation of control summary information, and then click **Finish**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- This topic does not apply to Windows Server 2003, Web Edition.
- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.
- This procedure pertains to RIS client computers that are not prestaged. For information about setting permissions that pertain to prestaged client computers, see Related Topics.
- If you performed this procedure for a group, remember to add users to that group.
- It is recommended that you limit the user to a specific container in which the RIS server is configured to create clients. Otherwise, the user will be able to create computer objects in any container in the domain.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Set permissions required by RIS users who use prestaged client computers
Allow or prevent the installing of a RIS image by a user or group
Manage Security for Remote Installation Services
Choosing appropriate group memberships for RIS administrators
Create a shortcut using the runas command

# Set Group Policy for client installation options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set Group Policy for client installation options**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or the applicable organizational unit.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)

3. Click **Properties**, and then click the **Group Policy** tab.
4. In the **Group Policy** dialog box, click the Group Policy object, and then click **Edit** to start Group Policy.
5. Click **User Configuration**, click **Windows Settings**, and then click **Remote Installation Services**.
6. Double-click the **Choice Options** object, and then select from among the following Client Installation Wizard choices:

   - **Automatic Setup**

     Use this option if you want to predefine the computer name and a location within Active Directory for the client computer accounts. (Another alternative is prestaging client computer accounts.) For more information about computer names and the options for predefining them, or for information about prestaging, see Related Topics.

   - **Custom Setup**

     Use this option to allow users to define a unique name for this computer and specify where the computer account will be created within Active Directory.

   - **Restart Setup**

     This option restarts an operating system installation attempt if it fails prior to completion.

   - **Tools**

     This option gives users the ability to access tools, including the Recovery Console, from the Client Installation Wizard. For more information about maintenance and troubleshooting tools and options for RIS, see Related Topics.

7. Click one of the following Group Policy options for each of the Client Installation Wizard choices selected in the previous step:

   - **Enabled**

     Use this option to offer the installation option to users to whom this policy applies.

   - **Disabled**

If this option is set for a specific installation option, users affected by this policy cannot access that installation option within the Client Installation Wizard.

- **Not Configured**

  Use this policy to accept the policy settings of the parent container. For example, if the administrator for the entire domain has set Group Policy that is specific to Remote Installation Services, and the administrator of this container has chosen the Not Configured option, the policy that is set on the domain is applied to all users who are affected by that policy.

8. In the **Choice Options Properties** dialog box, click **OK** to apply your changes.

9. Close Group Policy, and then, in the **Group Policy** dialog box, click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- This topic does not apply to Windows Server 2003, Web Edition.

- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.

- The **Not Configured** installation choice is the default setting.

- **Automatic Setup** is the default display option for users of client computers during the Client Installation Wizard process.

- The security settings that have been set on a domain controller through Group Policy within the Active Directory Users and Computers console might not take effect for the defined user or group accounts for up to eight hours.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Set permissions required by RIS users who use prestaged client computers
Controlling client installations through Group Policy
Maintenance and troubleshooting tools and options
Computer names
Group Policy (pre-GPMC)
Group Policy overview
Create a shortcut using the runas command

# Allow or prevent the installing of a RIS image by a user or group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow or prevent the installing of a RIS image by a user or group**

1. Open Active Directory Users and Computers.
2. On the **View** menu, ensure that **Advanced Features** is checked.
3. In the console tree, right-click the applicable Remote Installation Services (RIS) server.

   **Where?**

   - Active Directory Users and Computers/*Applicable domain*/*Applicable organizational unit* (such as Computers)/*Applicable RIS server*

4. Click **Properties**.
5. In the **Properties** dialog box, click the **Remote Install** tab, and then click **Advanced Settings**.
6. In the **Remote Installation Services Properties** dialog box, click the **Images** tab.
7. In the **Images** dialog box, click the installation image that you want to add or remove from the choices available to a user or group, and then click **Properties**.
8. Click **Permissions**.
9. Click the **Security** tab and do one of the following:

   - To make this installation image available to a group or user that does not appear in the **Group or user names** box, click **Add**. In **Select Users, Computers, or Groups**, type the name of the group or user, click **OK**, and then click the name of the group you just added.
   - To check or change the permissions of this installation image for an existing group or user, click the name of the group or user.

   **Important**

   - The **Group or user names** box always includes **SYSTEM**. Do not remove this entry, and do not change the permissions on it. This entry is necessary for RIS.

10. Do one of the following:

    - To prevent the selected group or user from installing this image, click **Remove**.
    - To allow the selected group or user to install this image, under **Allow**, ensure that **Read & Execute** and **Read** are selected.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For information about assigning appropriate permissions so that a RIS administrator who is not a member of these groups can perform this procedure, see Related Topics.

- This topic does not apply to Windows Server 2003, Web Edition.

- To open Active Directory Users and Computers, click **Start**, click **Run**, type **dsa.msc**, and then click **OK**. For information about creating a shortcut so you can easily open Active Directory Users and Computers with **runas**, see Related Topics.

- For information about creating new installation images, see Related Topics.

- **Authenticated Users** appears by default in **Group or user names**. This means that all authenticated users can install a RIS image. To allow a more limited group (instead of all authenticated users) to install a RIS image, add the appropriate group and allow the permissions described in this procedure. Click **Authenticated Users**, and then click **Remove**.

- With products in the Windows Server 2003 family, **Authenticated Users** does not by default contain anonymous users or guests. The **Everyone** group contains guests. As described in this procedure, it is recommended that you exclude **Everyone** from the list of groups for which permissions are assigned on the RIS image, and use **Authenticated Users** or a more restricted group for assigning these permissions. For more information about differences in default security settings for groups such as **Authenticated Users** (as compared to the settings in earlier operating systems), see Related Topics.

- If you leave **SYSTEM** but remove all other names from **Group or user names**, no one will be able to perform remote installations with this image.

- If you previously restricted permissions on the folder and subfolders containing the image, you must adjust those permissions, not just the permissions described in this procedure, to allow intended users to install the image. For instructions on how to do this, see the topic about allowing the viewing and installing of a RIS image in Related Topics.

- Setting permissions as described in this procedure affects the user's ability only to install a particular RIS image. For information about how to exert greater control and affect not only the user's ability to install a particular RIS image, but also the user's ability to see the files in that image, see Related Topics. To prevent confusion, once you have chosen a particular approach, it is recommended that you continue to use that approach whenever modifying permissions for images.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add RIS Client Installation Images
Create a Remote Installation Preparation Wizard image
Choosing appropriate group memberships for RIS administrators
Manage Security for Remote Installation Services
Allow or prevent the viewing and installing of a RIS image by a user or group
Differences in default security settings
Create a shortcut using the runas command

# Allow or prevent the viewing and installing of a RIS image by a user or group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow or prevent the viewing and installing of a RIS image by a user or group**

1. On the Remote Installation Services (RIS) server, open Windows Explorer.
2. Click the folder that contains the RIS folders and files. (The default name for the folder is RemoteInstall.)
3. Click the **Setup** folder.
4. Click the folder for the language of the installation image you want to confirm or change access to, and then click the **Images** folder.
5. Right-click the folder that contains the image that you want to confirm or change access to, and click **Properties**.
6. Click the **Security** tab.
7. For best security, if several names appear in **Group or user names**, and one of those groups is **Everyone**, click that group and then click **Remove**

   **Important**

   - The **Group or user names** box always includes **SYSTEM**. Do not remove this entry, and do not change the permissions on it. This entry is necessary for RIS.

8. Do one of the following:

   - In **Group or user names**, if you see the name of a group or user for whom you want to confirm or change access to a RIS image, click the name of the group or user.
   - In **Group or user names**, if you do not see the name of a group or user for whom you want to confirm or change access to a RIS image, click **Add**. Type the name of the group or user, click **OK**, and then click the name of the group or user.

9. Do one of the following:

   - To prevent the selected group or user from viewing and installing this image, click **Remove**.
   - To allow the selected group or user to view and install this image, under **Allow**, ensure that **Read & Execute**, **List Folder Contents**, and **Read** are selected.

10. When the list of group and user names and their permissions match the access you want to allow or prevent, click **Advanced**.
11. Select the check box marked **Replace permission entries on all child objects with entries shown here that apply to child objects**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

- This topic does not apply to Windows Server 2003, Web Edition.

- **Authenticated Users** appears by default in **Group or user names**. This means that all authenticated users can view and install a RIS image. To allow a more limited group (instead of all authenticated users) to view and install a RIS image, add the appropriate group and allow the permissions described in this procedure. Then click **Authenticated Users** and click the **Remove** button.

- With products in the Windows Server 2003 family, **Authenticated Users** does not, by default, contain anonymous users or guests. The **Everyone** group contains guests. As described in this procedure, it is recommended that you exclude **Everyone** from the list of groups for which permissions are assigned on the RIS image; then use **Authenticated Users** or a more restricted group for assigning these permissions. For more information about differences in default security settings for groups such as **Authenticated Users** (as compared to the settings in earlier operating systems), see Related Topics.

- If you leave **SYSTEM** but remove all other names from **Group or user names**, no one will be able to perform remote installations with this image.

- Using this procedure, you can control the user's ability to install a particular RIS image, as well as the user's ability to see the files in that image. You might want to exert less control, however. You can do this by changing permissions that, unless overridden by this procedure, affect the user's ability only to install (not see) a particular RIS image. For more information about this, see Related Topics. Once you have chosen a particular approach, it is recommended that you continue to use that approach whenever modifying permissions for images; this will help prevent confusion.

- If you prevent a user from viewing or reading the files in an image, you also prevent the user from installing the image.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add RIS Client Installation Images
Create a Remote Installation Preparation Wizard image
Choosing appropriate group memberships for RIS administrators
Allow or prevent the installing of a RIS image by a user or group
Differences in default security settings

# Set permissions for administrators who manage client installation images for RIS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions for administrators who manage client installation images for RIS**

1. Create a list of people who will manage user access to client installation images for Remote Installation Services (RIS). The list might include the same people who will manage accounts for prestaging for RIS.

2. Create a group account (or determine the name of the group account if it already exists), and make the people on the list members of the group.

3. On a computer from which you can set folder permissions for the RIS server containing the installation images that will be managed by the group, open Windows Explorer.

4. Right-click the folder that contains the RIS folders and files (the default name for the folder is RemoteInstall).

5. Click the **Setup** folder.

6. Click the folder for the language of the installation image that you want to set permissions for.

7. Click the **Images** folder.

8. Do one of the following:

   - If you want the group that you created to manage all images in the **Images** folder, right-click the **Images** folder and then click **Properties**.

   - If you want the group that you created to manage only certain images in the **Images** folder, right-click a folder that contains an image that you want the group to manage, and then click **Properties**.

9. Click the **Security** tab.

10. Do one of the following:

    - If the group you created for those who manage client installation images does not appear in the **Group or user names** box, click **Add**. Type the name of the group, and then click **OK**.

    - If the group you created for those who manage client installation images appears in the **Group or user names** box, click the name of the group.

11. In the **Permissions for User or Group** box, under **Allow**, make sure the **Full Control** check box is selected.

12. When the list of group and user names and their permissions match the access you want to allow or prevent, click **Advanced**.

13. Select the check box marked **Replace permission entries on all child objects with entries shown here that apply to child objects**.

14. If you have another folder for which you want to set permissions, navigate to the folder and repeat the process for

setting the permissions.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- This topic does not apply to Windows Server 2003, Web Edition.
- Consider using this procedure if there are people in your organization who manage accounts and permissions but do not configure other aspects of RIS. You can use this procedure to place those who manage accounts and permissions in a group such as Account Operators instead of putting them in a group such as Domain Admins or Enterprise Admins. For more information setting appropriate permissions for RIS administrators, see Related Topics.
- Managing user access to client installation images includes the following:
    - Associating an unattended setup answer file with an installation image
    - Allowing or preventing the installing of a RIS image by a user or group
    - Allowing or preventing the viewing and installing of a RIS image by a user or group
- For security reasons, set these permissions on the folder, as described in this procedure, not on the shared folder. Permissions on a folder are controlled by the NTFS file system, which provides greater security than permissions set on a shared folder.
- For information about creating a group or adding a member to a group, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Allow or prevent the viewing and installing of a RIS image by a user or group
Allow or prevent the installing of a RIS image by a user or group
Create a new group
Add a member to a group
Choosing appropriate group memberships for RIS administrators

# Remote Installation Services Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Remote Installation Services (RIS). It covers the following topics:

- Remote Installation Services overview
- Understanding Remote Installation Services
- Using Remote Installation Services
- Remote Installation Services Resources

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Remote Installation Services overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services overview**
Remote Installation Services (RIS) is a change-and-configuration-management feature that is included in the Microsoft® Windows Server 2003 family operating systems.

You can use RIS to provide on-demand, image-based, operating system installations. This occurs over a network connection from a RIS server to a client computer. RIS can simplify the management of operating systems and applications and improve failure recovery.

You can use RIS to do the following:

- Provide an operating system to users on demand. When a user starts a client computer, even if that computer contains no operating system, the RIS server can respond by installing an operating system over the network; no CD is required. To support this, client computers must use the Pre-Boot eXecution Environment (PXE), a remote-boot technology that allows the client computer to begin a boot sequence from the network adapter. For more information about the process through which RIS installs an operating system, see PXE architecture and Technologies and components used by Remote Installation Services.

- Provide operating system images that include specific settings and applications--for example, an image that complies with a corporate desktop standard. A particular group of users can be offered the image or images that you designate for that group. For more information, see Creating an installation image with RIPrep.

- Create automated installation images of products in the Windows Server 2003 family, as well as images of Windows XP and Windows 2000. For more information on the operating systems RIS supports, see Operating systems supported by Remote Installation Services.

In addition to using RIS in the ways already described, you can combine RIS with IntelliMirror features, such as user documents and settings, Software Installation, and Group Policy. This can improve the efficiency of computer management in your organization and reduce the number of technical support service calls. For more information about combining these features, see IntelliMirror.

For information on another alternative for starting a remote computer, see Emergency Management Services.

**Notes**

- RIS is not included in the Windows Server 2003, Web Edition operating system. For more information, see Overview of Windows Server 2003, Web Edition.

- This topic does not apply to Windows Server 2003, Web Edition.

# Understanding Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Remote Installation Services**
This section covers:

- Technologies used by Remote Installation Services to identify network computers
- Remote Installation Services Architecture

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Technologies used by Remote Installation Services to identify network computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Technologies used by Remote Installation Services to identify network computers**
Remote Installation Services (RIS) requires several supporting technologies that identify computers on the network. The following three technologies are included as components in the Windows Server 2003 family:

- Dynamic Host Configuration Protocol (DHCP)
- Domain Name System (DNS)
- Active Directory

DHCP and DNS need not be Microsoft versions to work with RIS. You must use the Microsoft version of Active Directory when using RIS, however. DHCP, DNS, and Active Directory can be installed (along with RIS) on one server or on several servers. (You are installing Active Directory when you make a server a domain controller.) Before adding RIS to a server that provides DHCP, DNS, or Active Directory, it is recommended that you consider server load and server performance.

RIS also depends on a unique identifier that computer manufacturers assign to each computer. This identifier, which is called the globally unique identifier (GUID), is also known as the universally unique identifier (UUID). As its name suggests, it gives each computer a unique identity. Each RIS client computer connects to the RIS server through the network and identifies itself by its GUID/UUID. To find a computer's unique identifier, see Locate the GUID for client computers.

When RIS detects a client computer's request for service, it responds in one of the following two ways, depending on how you have set up your network:

- A RIS server that you have designated for that client computer recognizes the client computer from its unique identifier and responds with the available installation choices.
- Any RIS server answers all network service requests from client computers.

For information on specifying how RIS responds to requests, see Remote Installation Services administration overview.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Remote Installation Services Architecture

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services architecture**
This section covers:

- Technologies and components used by Remote Installation Services
- PXE architecture
- Remote Installation Services directory service planning
- Remote Installation Services system requirements
- Remote Installation Services boot disk
- Network service boot

For information about the sequence of steps through which you create an installation image by using the Remote Installation Preparation (RIPrep) Wizard, see Creating an installation image with RIPrep.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Technologies and components used by Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Technologies and components used by Remote Installation Services**

Remote Installation Services (RIS) uses or depends upon a variety of technologies and components.

**Network infrastructure**

RIS depends on a network infrastructure that can identify computers on the network. RIS therefore requires that DHCP, DNS, and Active Directory be available on the network.

**Remote-boot technology**

RIS uses remote-boot technology to allow a computer with no operating system to begin a boot sequence from the network adapter, after which an operating system can be installed across the network. This remote-boot technology is called the Pre-Boot eXecution Environment (PXE). For more information, see Network service boot and PXE architecture. For client computers that are not PXE-enabled--that is, pre-Net PC/PC98 computers--you can use a special remote boot disk if the client computer has a supported network adapter. For more information, see Remote Installation Services boot disk and Create a remote boot disk.

**Technology to create an installation image**

RIS includes technology that you can use to create an image of an operating system that will be installed on client computers. You can create the image in one of two formats: a flat image or a Remote Installation Preparation Wizard (RIPrep) image format. Use the flat image option to create an image directly from a set of operating system files-- for example, the files on a CD. Use the RIPrep image format to create an image that includes an operating system with specific settings and applications-- for example, an image that complies with a corporate desktop standard. For more information, see Add a new client operating system installation image and Creating an installation image with RIPrep.

**Technologies for communication between RIS client and server**

RIS includes technologies that establish communication between a client that has begun a PXE boot sequence and the RIS server that contains installation images available to that client. For this process, RIS uses DHCP to provide the client with an Internet Protocol (IP) address. RIS then downloads the Client Installation Wizard, which prompts the user to log on and provides a menu of installation options customized for that user. You can control these images through Group Policy. For more information about the Client Installation Wizard, see Client Installation Wizard and Installation options available through RIS.

**Services added by RIS**

When RIS is enabled, the following services are automatically added to the server:

- **Remote Installation (BINLSVC)**

  This service listens for client network service requests and provides overall management of the RIS environment. BINLSVC ensures that the correct files are passed to the client computer. It also ensures that when the settings for a prestaged client specify that it is to be answered by a particular RIS server, the response comes from the specified server. If the client computer has not been prestaged, BINLSVC creates the client computer account object within Active Directory.

- **Trivial File Transfer Protocol Daemon (TFTPD)**

  RIS server uses the Trivial File Transfer Protocol Daemon (TFTPD) to download the initial files required for the remote installation process to begin. The most common file downloaded to the client using TFTPD is Startrom.com, which is responsible for bootstrapping the client computer. The default version of Startrom.com then prompts the user to press F12 to initiate a network installation. If the user responds to the prompt, the Client Installation Wizard is downloaded to begin the remote installation process.

- **Single Instance Store (SIS)**

  Single Instance Store (SIS) reduces the overall storage required on the RIS volume. The SIS drivers contain a feature,

called a groveler agent, that scans the RIS volume for duplicate files. If the SIS groveler finds duplicate files, it copies the original file into the SIS and leaves a link file in its place. The link file contains information about the original file, such as its current location, size, and attributes. If an image contains duplicate files, those duplicates are copied into the store; as a result, less disk space is required on the RIS server.

SIS has two limitations: it cannot act upon any files that are referenced through junction points, and it cannot be used with any file system except NTFS file system, which is the file system to use on RIS servers.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# PXE architecture

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PXE architecture**

Remote Installation Services (RIS) uses the Dynamic Host Configuration Protocol (DHCP) that follows the Pre-Boot eXecution Environment (PXE) architecture to bootstrap a client computer.

When a new PXE remote boot-enabled client computer starts for the first time, it uses the DHCP to request an Internet Protocol (IP) address and the IP address of an active RIS server. As part of the initial request, the client computer sends out its globally unique identifier (GUID), which identifies the client computer within Active Directory.

After the client computer receives an IP address from the DHCP server, it then requests service from a RIS server. The request to the RIS server is a broadcast, which means all of the RIS servers on the network will receive the request. If the RIS server is allowed to answer, it queries its preferred domain controller for the client computer's GUID. If the server fails to find the GUID, it queries the global catalog for the forest. If the GUID is found in either location, the client computer is considered "known." If the GUID is not found, it is considered "unknown." In this case --that is, the query by the RIS server found no computer account object in Active Directory with this computer's GUID --only RIS servers that are configured to answer unknown clients will do so.

If the client is known, all the available RIS servers query the domain to determine whether the client computer account has been prestaged with settings that specify that the client computer is to be answered by a particular server. If so, only the designated server answers the service request. The other servers respond by telling the client which server will answer it. This is called a referral. If the client computer account settings do not require the client computer to be answered by a particular server, all RIS servers answer it if they are aware of this client and configured to answer it. Each RIS server then offers its own files to download.

For information on configuring how RIS servers respond to client computers, see Configure Remote Installation Services. For information on prestaging clients computers and configuring the setting that specifies whether the client computer is to be answered by a particular server, see Prestage client computers.

After the client computer receives a reply to its service request, it initiates a Trivial File Transfer Protocol (TFTP) download of the bootstrap program. The Remote Installation Services implementation of this program is called Startrom.com (located at *RIS_server_name*\REMINST\oschooser\i386\startrom.com).

After Startrom.com has been downloaded, the client computer runs it. The default version of Startrom.com then prompts the user to press F12 to initiate a network installation. If the user fails to press F12 within three seconds, the network boot is stopped, and the client computer attempts to start up from the next available boot device. When the user presses F12, the client computer uses TFTP to download the Client Installation Wizard. The wizard then requests the user to log on to the network. For more information about user choices available during remote installation, see Perform a network installation using RIS and Installation options available through RIS.

For information about security limitations related to the PXE standard, see PXE architecture, RIS, and security considerations.

**Considerations for 64-bit client computers**

For 64-bit client computers, Startrom.com is not required and the F12 prompt is not provided to the user. The 64-bit client computer simply downloads the Extensible Firmware Interface (EFI) version of the Client Installation Wizard (located at *RIS_server_name*\REMINST\oschooser\ia64\oschoice.efi) when the user selects a remote installation option during computer startup.

The 64-bit client computers can request service from RIS servers even if the servers are configured not to answer clients. These client computers can request service if they have received a referral from one of the other RIS servers in the same forest.

### Note

- This topic does not apply to Windows Server 2003, Web Edition.

# Remote Installation Services directory service planning

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services directory service planning**
The Remote Installation Services (RIS) environment relies on a well-designed and well-planned Active Directory architecture. Normally, determining the physical location of a server can be a challenge. This can become even more difficult when you must find and configure several servers that are located in different buildings, different offices, or on different floors. Using Active Directory, however, can simplify this process.

In a given domain, Active Directory provides organizational units or containers that you can use to organize users and resources into logical administrative groups. This makes it easier to locate and configure servers at multiple locations. For more information on planning Active Directory architecture, see Organizational units.

You can verify a RIS server using the **Verify Server** command, which is available only locally on a RIS server or by connecting to that server through Remote Desktop for Administration. The computer account object is located in the Computers container when the RIS server is a member server of a domain. For performance reasons, it is recommended that you install RIS on a member server, not a domain controller. For more information, see Verify the Remote Installation Services configuration.

To manage a RIS server from a computer running Windows XP, install the Windows Server 2003 Administration Tools Pack on the computer running Windows XP. For more information, see Windows Server 2003 Administration Tools Pack.

To manage a RIS server from any remote computer, use Remote Desktop for Administration. For more information, see Using Remote Desktop for Administration for remote server administration.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Remote Installation Services system requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services system requirements**
This section describes the system requirements for Remote Installation Services (RIS) servers, client computers, and network adapters.

**Server hardware requirements**

- The server hardware must meet minimum requirements for the product version of the Windows Server 2003 family that you are installing. For example, if you are installing Windows Server 2003, Enterprise Edition, your computer must meet the minimum requirements for this product. For more information, see Windows Server 2003, Enterprise Edition: System requirements.
- 4-gigabyte (GB) or larger disk drive for RIS server folder tree.
- 10 or 100 Mb/sec Windows-compatible network adapter that supports TCP/IP (100 Mb/sec recommended). A RIS server cannot be a multihomed computer; that is, it can contain only one network adapter.

**Client computer hardware requirements**

- The client computer must meet the minimum requirements for the operating system that is to be installed.
- Pre-Boot eXecution Environment (PXE) DHCP-based boot ROM version 1.00 or greater.

**Supported client network adapters**
To obtain a list of network adapters that RIS supports, run the Rbfg.exe utility. For additional information, see Create a remote boot disk.

**Additional considerations**

- RIS requires a significant amount of disk space. You should dedicate an entire partition to the RIS folder tree.
- RIS cannot be installed on the same partition as the system volume or boot volume.
- The volume on which you install RIS must be formatted with the NTFS file system.
- RIS does not currently support the Encrypting File System (EFS). The Distributed File System (DFS) is not supported as a target location; however, you can run RIS on a computer that is also running DFS.
- Windows Catalog contains a compilation of systems and hardware that have been extensively tested for stability and compatibility with products in the Windows Server 2003 family. The information in Windows Catalog is the guide Microsoft Product Support Services uses to determine whether a system is supported for use with the Windows Server 2003 family. For more information about hardware compatibility, see Support resources.

### Note

- This topic does not apply to Windows Server 2003, Web Edition.

# Remote Installation Services boot disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services boot disk**
The Remote Installation Services (RIS) boot disk is used with client computers that do not have a remote boot-enabled ROM. The boot disk simulates the Pre-Boot eXecution Environment (PXE) boot process for computers that lack a formal remote boot ROM. For more information about the PXE boot process, see PXE architecture.

The boot disk is analogous to a boot ROM, which uses the floppy drive to install the operating system from the RIS server. The remote boot disk-generating utility (Rbfg.exe) is located in the following folder:

*Server_name*\REMINST\Admin\I386\Rbfg.exe

The boot disk must be placed before the hard disk in the boot sequence in order for the boot disk to work.

To use the remote boot disk on a client computer that does not have a remote boot-enabled ROM, insert the disk into the disk drive and start the client computer. The client computer must have a Peripheral Component Interconnect (PCI)-based network adapter from the list of specific adapters supported. To determine which network adapters RIS supports, use the boot disk-generating utility. For more information, see Create a remote boot disk.

For more information about performing a client installation, see Perform a network installation using RIS.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Network service boot

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network service boot**
A network service boot can be initiated by either the user of a Pre-Boot eXecution Environment (PXE) remote boot-enabled client computer or by a special remote boot disk provided for pre-Net PC/PC98 computers.

To invoke the network service boot request, the user presses a particular key (for example, F12) when prompted during the client computer startup sequence. The client computer receives an Internet Protocol (IP) address from the DHCP server, requests a network service boot, and then downloads the Client Installation Wizard.

For more information about how a remote service boot is carried out through PXE, see PXE architecture.

For more information about user choices and actions during a network service boot, see Perform a network installation using RIS.

The user is prompted for user name, password, and domain name during the initial phase of the Client Installation Wizard process. Remote Installation Services uses Group Policy within Active Directory to determine which installation options are appropriate for the user. The user is then given a list of operating system installation choices that are determined by the user's credentials or security group membership.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Using Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Remote Installation Services**
This section covers:

- Remote Installation Services server authorization
- Running the Remote Installation Services Setup Wizard
- Using Risetup from the command line
- Remote Installation Services administration overview
- Creating an installation image with RIPrep
- RIS client computer properties
- Security Information for Remote Installation Services
- Advanced Configuration Settings
- The Remote Installation Process

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Remote Installation Services server authorization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services server authorization**
The Remote Installation Services (RIS) server authorization feature prevents unauthorized RIS servers from being added to a network where Active Directory is in use.

You can use RIS to specify which RIS servers can accept and process requests, and which RIS servers can only service clients on the network. Before a RIS server can accept requests, it must be authorized to run in Active Directory. To authorize a RIS server in Active Directory, you must first log on with an account that is a member of the **Enterprise Admins** group in the domain where you want the RIS server to be authorized. You should run **RISETUP -Check** to authorize the server.

When a RIS server attempts to start on the network, Active Directory is queried and the Internet Protocol (IP) address is compared to the list of authorized RIS servers. If a match is found, the RIS server is authorized and can start on the network. If a match is not found, the server is not authorized. In this case, RIS will not answer.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Running the Remote Installation Services Setup Wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Running the Remote Installation Services Setup Wizard**

Before running the Remote Installation Services Setup Wizard, you must ensure that you have an account in the Administrators group on the server on which you want to run the wizard. You should also be sure to complete all prerequisite steps outlined in Checklist: Installing Remote Installation Services.

You can select Remote Installation Services (RIS) as an optional component after completing Setup. To do this, use Add or Remove programs in Control Panel. You can also install RIS during Setup if Setup is run using an unattended answer file. In this case, after you have selected RIS, the Setup Wizard guides you through the necessary steps of installation and configuration. For complete instructions on installing RIS, see Install Remote Installation Services.

The Setup Wizard presents the following options:

- The first time the wizard is run, you will be prompted for a location for the \RemoteInstall directory tree. This location must be on an NTFS file system volume and cannot be on the system volume or boot volume.

- **Respond to client computers requesting service**

  Select this option to enable the RIS server to begin responding to client computers. Clear this option to prevent the RIS server from responding.

- **Do not respond to unknown client computers**

  Select this option if you want this RIS server to respond only to known (prestaged) client computers. Clear this option if you want this RIS server to respond to any client computer requesting remote installation.

  Prestaging can be one element of your security strategy for RIS. This is because you can use prestaging to prevent unknown client computers from being installed through RIS. For information about security considerations related to PXE architecture and RIS, see PXE architecture, RIS, and security considerations.

- The Setup Wizard prompts you for the location of the installation files. This can be either the compact disc or a network location that contains the installation files.

- The Setup Wizard prompts you for a directory name to which the installation files should be copied on the RIS server. You should name the directory something that represents the operating system you are copying--for example, WindowsXP.pro.

- The Setup Wizard prompts you for a friendly description and Help text describing this operating system installation image. The friendly description and Help text is displayed to users during the client installation process on a remote boot-enabled client computer. Enter a description of the operating system installation choice. One example of a friendly description would be "Windows XP Professional for Sales Staff." The Help text is displayed when the user selects the friendly description within the Client Installation Wizard.

- The Setup Wizard displays the summary information of the setup process. To change any of this information, click **Back**.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Using Risetup from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Risetup from the command line**
With Risetup, you can manage a server running Remote Installation Services (RIS) from the command line. You can use Risetup to add new installation images or to verify a server and ensure that it is running properly.

To run Risetup, at a command prompt, type:

**risetup**

For complete information on using Risetup, see Risetup.

For information about other command-line utilities, see Command-line reference A-Z.

For background information about RIS, see Remote Installation Services Concepts.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Risetup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Risetup**

Manages a server running Remote Installation Services (RIS). Adds new installation images or verifies a server to ensure that it is running properly.

**Syntax**

risetup [-add] [-check] [/auto *FileName.inf*]

**Parameters**

| Parameter | Explanation |
|---|---|
| -add | Adds a new installation image to the Remote Installation Services (RIS) server. |
| -check | Verifies that the server is functioning properly. Also, runs only the server portion of Risetup, which resolves most problems related to a RIS server that is not running properly. |
| /? | Displays help at the command prompt. |
| /auto *FileName.inf* | Configures the server in an unattended manner as defined by *FileName.inf*. The path to the folder location of the .inf file must be under 255 characters.<br><br>The following table lists the lines of an example .inf file and explains each line: |

| Lines in the example .inf file | Requirement and explanation |
|---|---|
| `[Version]Signature = "$Windows NT$"` | Required. Heading and version line. |
| `[Risetup]` | Required. Heading. |
| `RootDir = "f:\RemoteInstall"` | Required. Drive and folder that contain the Remote Installation files and folders. You must include this line every time RISetup is run. If RIS is already installed, the path must remain unchanged from the current path for the RIS installation. The maximum number of bytes is 127. |
| `Source ="e:\"` | Required. Path to the source CD or shared folder from which to create an installation image. The maximum number of bytes allowed for this line is 260. For information about the supported operating systems for this type of image (a flat image), see Related Topics. |
| `Directory = "Windows_XP_Pro"` | Required. Directory in which to place the new installation image. The maximum number of bytes is 39. |
| `Description = "Windows XP"` | Required. Short description of the installation image. The maximum number of bytes is 65. |
| `HelpText = "For English-speaking users."` | Required. Long description of the installation image. The maximum number of bytes is 260. |
| `Screens = "overwrite"` | Required. Action to perform on existing .osc files, which contain text that displays in screens seen by users of the Client Installation Wizard. Specify only one of the following choices:<br><br>"overwrite"<br><br>"backup"<br><br>"leavealone" |

| | |
|---|---|
| `Architecture = "x86"` | Required. Can be one of:<br>• **x86** for an image that will be installed on x86 computers<br>• **ia64** for an image that will be installed on Itanium-based computers.<br>• **amd64** for an image that will be installed on x64-based computers. |
| `Language = "English"` | Required. Language of the image. |

**Examples**

To reauthorize the RIS server in Active Directory, type:

**RISETUP -Check**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Command-line reference A-Z
Running the Remote Installation Services Setup Wizard
Install Remote Installation Services
Operating systems supported by Remote Installation Services

# Remote Installation Services administration overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Installation Services administration overview**
To administer Remote Installation Services (RIS), you use the property pages that reside on specific objects within Active Directory.

To access the property pages, right-click the server object you want to administer, and then click **Properties**. For more information on configuring RIS, see Configure Remote Installation Services.

You can set the following options from the **Remote Install** property page:

**Client servicing**

You can specify either that servers respond to all clients requesting service, or that they respond only to known (prestaged) clients. Use the following options to specify how the RIS server responds to requests for service from client computers:

- **Respond to clients requesting service.** If you select this option, RIS is enabled and will respond to client computers requesting service.

- **Do not respond to unknown client computers.** If you select this option, the RIS server responds only to prestaged client computers in Active Directory. Otherwise, the RIS server responds to any client computer requesting remote installation.

For more information about prestaging, see Prestage client computers.

**Verify Server**

Select this option if the server seems to be failing or is behaving inconsistently, or if you need to restore a RIS volume from backup. This option brings up a wizard that checks to determine whether all of the settings, services, and configuration options are correctly set and functioning.

Before you select this option, make sure you have the installation compact discs available. You might be prompted for them during the Verify Server procedure.

You can only use Verify Server locally on a RIS server, or by connecting to that server through Remote Desktop for Administration. You cannot use Verify Server if you are performing RIS administration tasks from a remote server.

For more information, see Verify the Remote Installation Services configuration.

**Important**

- If you are verifying the server configuration because you need to restore a RIS volume from backup, you must verify the server configuration before you restore the volume.

**Show Clients**

If you select this option, you can search for RIS clients within Active Directory.

Selecting this option displays a list of client computers sorted by their globally unique identifier (GUID). The list also includes client computers that have been prestaged.

For more information, see RIS client computer properties.

**Advanced Settings**

If you select this option, you can control how a client computer is installed. The advanced setting options are as follows:

- **Automatic client computer account naming format**

- **Active Directory location of client computer accounts**

- **Manage the installation images and pre-operating system maintenance tools installed on this RIS server**

For more information, see Advanced Configuration Settings.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Creating an installation image with RIPrep

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating an installation image with RIPrep**

With the Remote Installation Preparation (RIPrep) Wizard, you can create an image of an existing installation and replicate that image to an available Remote Installation Services (RIS) server on the network. To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. You must also have permission to write to the RIS folders on the RIS server (the default name for the folder that contains all the RIS folders is RemoteInstall).

One way to ensure that you have permission to write to the RIS folders on the RIS server is to be explicitly assigned the permission. Another way is to have membership in the local Administrators group on the RIS server. For more information, see Default local groups and Default groups.

1. It is recommended that you start by using RIS to install the operating system that you have chosen for the installation image. RIPrep can be used to create images only for the particular operating systems it supports, however. For complete information on which operating systems RIPrep supports, see Operating systems supported by Remote Installation Services.

   It is not necessary to install the operating system using RIS; you can install it from the CD instead.

2. After you have installed the operating system, you should install the client computer applications that you want to deploy with clients.

   **Important**

   - If you want to install applications that do not comply with the Windows Installer technology, this is the step at which you should do so. You can also install applications that do comply with the Windows Installer technology.

3. After you have installed the appropriate client computer applications, you should configure the source computer to conform to any required company desktop standards. For example, you might want to define specific screen colors, set the background bitmap to a company-based logo, remove any games installed by the base operating system, or configure Internet Explorer proxy settings. To save your desktop configuration, copy to the Default User profile the profile of the user who was logged in while the configuration changes were made. For instructions on how to copy user profiles, see Copy a user profile.

   **Important**

   - Before starting the Remote Installation Preparation Wizard, be sure to remove all confidential information from the source client computer.

4. You should then close all applications and run the Remote Installation Preparation Wizard. For more information, see Create a Remote Installation Preparation Wizard image.

5. The wizard configures the source computer to a generic state, removing anything that is unique to the client installation, such as the computer's unique security identifier (SID), computer name, and any registry settings unique to the client source computer.

6. The wizard then prompts you for the installation information required by the image-creation process. This information includes the location where the client installation image should be replicated, the name of the directory it should be copied to on the server, and a friendly description and associated Help text that will describe the installation image to users who run the Client Installation Wizard.

Because of memory limitations during the remote boot process, there are limits to the total text you can provide for a set of installation image choices. With a relatively large number of installation image choices, the friendly descriptions and Help text must be limited in length; with fewer choices, the friendly descriptions and Help text can be longer. It is not unusual to encounter these limits to total text when the number of installation image choices for client computers reaches approximately 12 to 16.

7. After the replication is complete, the installation image is automatically added to the list of available operating system installation options and is available to client computers that use the remote boot technology. When you create new images using RIPrep, the image is available only to client computers that have the same type of Hardware Abstraction Layer (HAL). Images that were created using RIPrep are available to all clients that are authorized to view them.

   **Note**

   - A remote boot-enabled client computer need not contain hardware that is identical (other than the HAL) to that of the source computer used to create the installation image. The Remote Installation Preparation Wizard uses the Plug and Play feature to detect any differences between the source and destination computers' hardware during the image installation process. This ensures that, for any Plug and Play device, the correct device driver is installed.

8. You can control the installation choices available to clients by setting access control permissions on the installation image .sif files or, for greater control, on the folders that contain the installation image. For more information, see Allow or prevent the installing of a RIS image by a user or group and Allow or prevent the viewing and installing of a RIS image by a user or group.

**Factors to consider when planning for RIPrep**

The following sections describe factors to consider when planning to create an installation image with RIPrep. These include:

- Requirements for the RIPrep replication process
- Installation images, the product identification number, and licensing
- Disk characteristics on the source client computer and the destination computers

**Requirements for the RIPrep replication process**

When planning to run RIPrep, it is important to understand requirements that must be met by the source client computer, the image, and the medium onto which the image is replicated.

- The Remote Installation Preparation Wizard currently supports the replication of a single disk and a single boot partition of an installation to a single RIS server. This means that the operating system and all of the applications that make up the standard installation image must reside on a single partition of the source client computer.
- The files on the source client computer must not be encrypted. Replication of encrypted files is not supported.
- Modifications to replicated installation images are not supported.
- The wizard allows source image replication only to available RIS servers. Currently, source replication to alternate drives or media types is not supported.

**Installation images, the product identification number, and licensing**

- When you use the Remote Installation Preparation Wizard (RIPrep) to create an installation image of a client computer that was originally installed using a retail version of the operating system, you must modify the RIS unattended setup answer file (sif) to include the product identification number (PID). The PID is the unique identification number specific to each copy of the operating system.

  If the PID is not entered in the .sif file, the installation process will stop and prompt the user for the product identifier

information during the installation of that RIPrep image. To prevent prompting your users for the PID, add the product identifier to the [UserData] section of the .sif file associated with this installation image.

For example, you would type the following (including the dashes and quotation marks) into the [UserData] section of the .sif file:

**ProductKey = "xxxxx-xxxxx-xxxxx-xxxxx-xxxxx"**

For more information about product activation, see Windows Product Activation.

- The source client computer will shut down when the image replication process is complete. The abbreviated Setup program runs automatically when the source computer restarts. You must complete the setup process if you want to use this client computer to create another installation image. You can create additional installation images only three times if you are using a Retail PID. A volume licensing PID permits you to create installation images more than three times, however.

- All copies of Microsoft software made or installed using RIS must be properly licensed. All copies of non-Microsoft software made or installed using RIS must also be properly licensed; it is the licensee's obligation to ensure that non-Microsoft software is licensed to make any such copies.

**Disk characteristics on the source client computer and the destination computers**

When planning to run RIPrep, it is important to study the disk configuration that exists on the source client computer and then compare it to the disk configurations on the client computers on which the image will be installed. The following requirements apply:

- The destination computer's disk capacity must be equal to or greater than that of the source computer.

- The RIPrep installation image maintains the volume and partition characteristics of the disk on the source computer.

  The following example illustrates this, but describes the use of the file allocation table (FAT) and FAT32 file systems, which are not recommended. The recommended file system for RIS is the NTFS file system. The example is as follows. If you create an installation image from a source computer with a 2-gigabyte (GB) FAT volume, and you then install the image on a client computer with a 4-GB drive, the resulting installation will format the drive as a 4-GB volume using FAT32. The change in the destination computer file system type is dictated by the limits of the FAT file system. This type of change in the destination computer file system type (from FAT to another file system) does not work if the destination volume is larger than 32 GB. For destination volumes larger than 32 GB, the source computer must use the NTFS file system. For more information about file system compatibility and limitations, see Choosing a file system: NTFS, FAT, or FAT32.

- Disk preparation of the destination computer must be identical to that of the source computer; any remaining disk capacity on the destination computer will be formatted unless you follow the instructions in the next item in this list. For example, if the source computer's disk capacity is 1 GB and the disk capacity of the destination computer is 2 GB, the entire 2 GB will be formatted on the destination computer.

- You can change the information in the .sif file associated with an installation image to restrict the disk reformatting so that it will be the same as that which the source computer used to create the installation image. For example, open the Riprep.sif file located in the *Server_name*\REMINST\Setup*Applicable language*\Images*Applicable image name*\I386\Templates\Riprep.sif folder and modify the **UseWholeDisk** parameter to equal **NO**. When a client computer installs this image, the disk will be formatted to match the capacity of the source computer, and the remainder of the destination computer's disk will be unformatted.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# RIS client computer properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RIS client computer properties**

Remote Installation Services (RIS) uses property pages to contain information about RIS client computers. For information about viewing a list of client computers, see Find RIS client computers. To view or change information about RIS client computers, you can right-click a client listing, and then click **Properties**.

**The computer's unique identifier (GUID)**

One important property that appears on the property pages is the computer's globally unique identifier (GUID), also known as the universally unique identifier (UUID). The RIS server uses this identifier to locate a client computer. The following list provides more details about the GUID:

- The GUID is supplied by the client computer manufacturer. For information about how to find the GUID on a computer, see Locate the GUID for client computers.

- When you enter a GUID, you usually enter it in the display format, {*dddddddd-dddd-dddd-dddd-dddddddddddd*}, where *d* is a hexadecimal text digit. There are 8 digits before the first hyphen, 4 digits in each of the next three groups of digits, and 12 digits after the last hyphen. This is illustrated in the following example:

  {67452301-AB89-EFCD-0123-456789ABCDEF}

- The GUID uses only hexadecimal text digits and hyphens; that is, the characters in the following list (not case sensitive):

```
0 1 2 3 4 5 6 7 8 9
a b c d e f - A B C D E F
```

- Do not type spaces in a GUID.

- The GUID can also be entered in wire format. In a GUID in wire format, the initial 16 digits are transposed as compared to a GUID in display format. The following example compares wire format to display format:

  GUID in wire format: 0123456789ABCDEF0123456789ABCDEF

  Same GUID in display format: {67452301-AB89-EFCD-0123-456789ABCDEF}

**Finding client computers**

Each client computer is identified by its GUID. In Active Directory, you can use the Show Clients feature to display a list of client computers sorted by their GUIDs. The list includes client computers that have been prestaged. The search can be conducted over the entire Active Directory structure or limited to a specific domain. (For more information on GUIDs, see Locate the GUID for client computers and Prestage client computers.)

The Show Clients search process appends a wildcard search attribute to the current RIS server computer name. For example, if the RIS server is named RISsvr1, the Show Clients feature will use RISsvr1* for the server name.

When you use the Show Clients feature in multiple RIS server environments, the search results might include client computers from multiple servers. For example, if you have multiple RIS servers with computer names such as RISsvr1, RISsvr10, and RISsvr100, and you search for RISsvr1, the list will display the client computers that install from all three servers.

For more information about searching Active Directory for RIS client computer accounts, see Find RIS client computers.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Security Information for Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for Remote Installation Services**
This section provides background information about the following security topics:

- PXE architecture, RIS, and security considerations
- Choosing appropriate group memberships for RIS administrators
- Setting the LAN Manager Authentication Level on a network that includes RIS

For specific information on how to set permissions and user rights for Remote Installation Services (RIS), see Manage Security for Remote Installation Services.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# PXE architecture, RIS, and security considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PXE architecture, RIS, and security considerations**
The Pre-Boot eXecution Environment (PXE) architecture used by Remote Installation Services (RIS) presents some security risks that are inherent in PXE due to its design. PXE is an industry standard, so any remote-installation method that uses PXE carries the same risks.

These security risks include the following:

- PXE does not provide a way to prevent an unknown server from performing remote installations on PXE-enabled client computers. If a server can establish a connection with the clients, it can perform remote installations on them.

- PXE does not provide a way to fully prevent packet spoofing. This means that packets sent by an attacker could be received by a client computer and incorporated into that client computer's installation.

- PXE does not provide a way to prevent an unknown PXE-enabled computer from installing from a server if the PXE-enabled computer can connect to the network. RIS provides some security not inherent in PXE, however, because RIS performs remote installation only after the user has logged on. A user who lacks a valid user name and password cannot use RIS to perform an installation.

  In addition, you can achieve a somewhat greater degree of security with RIS if you prestage your client computers and configure your RIS servers to respond only to known (prestaged) clients. Then, if an intruder succeeds in connecting an unknown, PXE-enabled client computer to your RIS server, no installation files will be sent to that client computer. The intruder will not gain information about the configuration you use on your RIS client computers. For more information about prestaging, see Prestage client computers.

Because of these security issues, it is recommended that you place protections around any network that contains PXE-enabled client computers.

You can take steps, such as those described in the following list, to ensure that computers controlled by people outside of your organization cannot make a connection to PXE-enabled clients:

- Use a firewall, and configure it appropriately.

- Use appropriate auditing and monitoring to detect intrusions into the network.

- Restrict physical access to the network.

- Use strong passwords throughout your organization.

- Follow other best practices for secure networks.

For more information about the preceding recommendations, see Best practices for security for security and also the Best practices for installing and upgrading the operating system topic in Getting Started.

For information about the sequence of events that occurs when a client computer is started with PXE, see PXE architecture.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Choosing appropriate group memberships for RIS administrators

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing appropriate group membership for RIS administrators**

If there are people in your organization whose responsibilities include configuring Remote Installation Services (RIS) servers and creating installation images, make those people members of an administrative group such as Domain Admins or Enterprise Admins. This will allow them to carry out all RIS configuration tasks. For the single task of completing installation of a RIS server in a way that authorizes the RIS server in Active Directory, membership in Enterprise Admins is required. For more information about authorization, see Remote Installation Services server authorization.

Administrators in your organization should also be provided with user accounts that are not in any administrative group. With such accounts, administrators can follow the security best practice of logging on as a user and then using Run as to perform administrative tasks. For more information, see Using Run as and Create a shortcut using the runas command.

Whenever possible, choose group memberships and assign permissions in a way that provides only necessary access to domain accounts and resources. For example, perhaps there are people in your organization whose responsibilities include managing accounts and permissions, but do not include configuring RIS servers or creating client installation images. Instead of giving these people membership in a group such as Domain Admins or Enterprise Admins, give then membership in a group such as Account Operators and grant them permissions for folders on the RIS server. For more information, see Default local groups and Default groups.

The following table illustrates the group membership and permissions required for people who manage only the tasks of RIS related to permissions, computer accounts, and user accounts:

| **Tasks** | **Permission or group membership required for the administrator who performs those tasks** |
|---|---|
| Managing client installation images, including:<br>• Associating an unattended setup answer file with an installation image.<br>• Allowing or preventing the installing of a RIS image by a user or group.<br>• Allowing or preventing the viewing and installing of a RIS image by a user or group.<br><br>For more information about these tasks, see:<br>• Manage Security for Remote Installation Services | Full Control on the **Images** folder, or on one or more subfolders within the **Images** folder.<br><br>For more information, about setting this permission, see:<br>• Set permissions for administrators who manage client installation images for RIS |

| | |
|---|---|
| Managing accounts in domains where prestaging is used, including:<br>• Prestaging client computers.<br>• Setting permissions required by RIS users who use prestaged client computers.<br>• Removing an account for a computer that is no longer being prestaged.<br><br>For more information about these tasks, see:<br>• Prestage client computers<br>• Set permissions required by RIS users who use prestaged client computers | Membership in Account Operators.<br><br>For more information, see:<br>• Default groups<br>• Add a member to a group |

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Setting the LAN Manager Authentication Level on a network that includes RIS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting the LAN Manager Authentication Level on a network that includes RIS**

You can use the LAN Manager Authentication Level to set a level that specifies whether certain authentication protocols are available for network communication. When a user logs on at a RIS client, Remote Installation Services (RIS) uses either the first version of the NTLM authentication protocol or the second version, NTLMv2. Of the two protocols, NTLMv2 is significantly more secure because of the way it handles encryption keys.

For information about choosing the most appropriate LAN Manager Authentication Level in a network that includes RIS, see the table that appears later in this topic. For information about the different levels and how to view or change the setting, see Network security: LAN Manager authentication level.

When you set the LAN Manager Authentication Level, you should consider both your requirements for network security and your requirements for using a variety of operating systems. If you select the highest level (**Send NTLMv2 response only\refuse LM & NTLM**), only NTLMv2, the more secure protocol, is used; all computers involved in authentication must run software that supports NTLMv2. If you select a lower level (**Send NTLM response only**), NTLMv2 is used whenever possible, and NTLM is used only when required--that is, when one or more computers involved in authentication do not support NTLMv2.

The LAN Manager Authentication Level that is most appropriate depends upon which operating system is installed on the domain controllers, on the RIS servers, and on the clients you want to install. The following table provides criteria you can use to determine the appropriate LAN Manager Authentication Level.

| Criteria for "Send NTLM response only" | Criteria for "Send NTLMv2 response only\refuse LM & NTLM" |
|---|---|
| **Domain controllers: any**<br><br>Domain controllers can include those running a Windows 2000 Server operating system without a specific service pack. They do not have to be limited as described in the other column of this table. | **Domain controllers: run either a Windows 2000 Server operating system with a specific service pack, or they run a Windows Server 2003 operating system.**<br><br>In order to work with this authentication level, domain controllers must run either a Windows 2000 Server operating system with a specific service pack, or they must run a Windows Server 2003 operating system. For information about the service pack required in order to run a Windows 2000 Server operating system with this authentication level, search the Knowledge Base at the Microsoft Web site.<br><br>In addition, RIS servers and clients must also meet the requirements in this table. |
| **RIS servers: any**<br><br>RIS servers can include those that run Windows 2000. They do not have to be limited to those running Windows Server 2003. | **RIS servers: run Windows Server 2003 only**<br><br>In order to work with this authentication level, all RIS servers must run Windows Server 2003. In addition, domain controllers and clients must also meet the requirements in this table. |
| **Clients: any**<br><br>Clients can include those that run Windows 2000 or Windows XP without a service pack. They do not have to be limited to those running Windows Server 2003 or those running Windows XP with Service Pack 1 (SP1) or later. | **Clients: run Windows XP with SP1 or later, or run Windows Server 2003**<br><br>In order to work with this authentication level, clients must run either Windows XP with SP1 or later, or run Windows Server 2003. In addition, domain controllers and RIS servers must also meet the requirements in this table. |

| Overall effect of using "Send NTLM response only" | Overall effect of using "Send NTLMv2 response only\refuse LM & NTLM" |
| --- | --- |
| Security is not as strong, but you can choose from a wider range of operating systems. | Security is stronger, but you must choose from a narrower range of operating systems. |

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Advanced Configuration Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Advanced configuration settings**
The advanced configuration settings of Remote Installation Services (RIS) determine how RIS creates client computer names, manages operating system images, and manages diagnostic tool applications.

This section covers:

- Advanced settings for new clients
- Advanced settings for installation images
- Maintenance and troubleshooting tools and options
- Computer names

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Advanced settings for new clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Advanced settings for new clients**

You can use the advanced settings of Remote Installation Services (RIS) to configure settings for new clients. With advanced settings, you can define a policy that provides each remote installation client computer with a name that is unique on the network. For information on how to do this, see Computer names.

You can use advanced settings to determine the default directory service container in which the client computer machine account objects will be stored. You can also group client computers within a specific Active Directory domain or organizational unit.

A user setting up the client computer and using advanced settings must have the appropriate user rights to create computer accounts within Active Directory. For more information, see Allow a user or group to create a RIS managed computer account in the domain.

**Computer account location**

This option is available in advanced settings. Use it if you want to define the default Active Directory container for all remote installation client computer accounts prior to installation. To specify a location, select one of the following:

**Default directory service location**

Specifies that the client computer account object is created in the Computers container by default during the domain join operation. The client computer becomes a member of the same domain as the RIS server installing the client.

**Same location as that of the user setting up the client computer**

Specifies that the client computer account object is created within the same Active Directory container as the user setting up the computer. For example, if a user logs on whose user account currently resides within the Users container, the client computer account is also created within the Users container.

**The following directory location**

With this option, you can predetermine where client computer account objects are created in Active Directory. This option affects all client computers installed from this server.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Advanced settings for installation images

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Advanced settings for installation images**
You can use Remote Installation Services (RIS) to manage the operating system installation options for RIS users. During the base installation of RIS, you are prompted to add a default operating system installation image to the remote installation-enabled server. This image installation process copies the compact disc contents and directory structure to the RIS server.

The unattended installation process uses text files--which are called setup information (.sif) files --to store the configuration settings for the installation image. The information contained in the .sif file is processed each time a client computer uses the Client Installation Wizard to choose the installation image. By modifying the .sif file for the installation image, you can predetermine the configuration settings for each client installation image. For example, if all of your client computers have the same type of display devices, you can set the screen resolution prior to installation by modifying the [Display] information of the .sif file. For more information on unattended Setup, see the Microsoft Windows Preinstallation Reference, located in \Support\Tools\deploy.cab on the CD for products in the Windows Server 2003 family.

**Important**

- The information you put in the [ImageType] field of the .sif file is limited to 64 bytes. If you put in more than 64 bytes, the installation will fail.

For additional information about ways you can use .sif files, see Installation options available through RIS.

If you are creating an installation image using the Remote Installation Preparation Wizard (RIPrep), you can also use a .sif file to change how the client computer disk drive is partitioned. While the Client Installation Wizard is running, the installation image that the user selects is copied to the client computer. If the default disk preparation of the destination computer is not identical to that of the source computer, any remaining disk capacity on the destination computer is partitioned. For example, if the source computer disk capacity is 1 gigabyte (GB) and the disk capacity of the destination computer is 2 GB, the entire 2 GB are formatted on the destination computer.

By changing the information in the .sif file associated with an installation image, you can restrict the disk reformatting to be the same as the source computer used to create the installation image. For example, suppose you open the Riprep.sif file and modify the UseWholeDisk parameter to **NO**. When a client computer installs this image, the disk is formatted to match the capacity of the source computer. The disk and source computer then become identical. The rest of the destination computer's disk remains unformatted.

For additional background information about how RIPrep works, see Creating an installation image with RIPrep. For instructions on running RIPrep, see Create a Remote Installation Preparation Wizard image.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Maintenance and troubleshooting tools and options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maintenance and troubleshooting tools and options**
You can control whether Remote Installation Services (RIS) offers certain types of maintenance and troubleshooting tools at the beginning of a remote installation. Although the client computer's hard disk might be empty prior to installing the operating system, the computer can still be managed with these maintenance and troubleshooting tools. The tools can therefore be useful not only to users, but also to administrators and technical support staff.

The maintenance and troubleshooting tools also give administrators an easy way to update a client computer system--for example, by the flashing of the basic input/output system (BIOS)--before or after the operating system installation.

If tools are provided to you by an independent software vendor (ISV) or original equipment manufacturer (OEM), you can control whether users are able to use these tools. You do this with the Tools tab of Advanced settings in RIS. (ISVs and OEMs make the tools available through external setup programs.) The tools appear in the list on the Tools tab and become available to administrators and users when they run the Client Installation Wizard. For information about viewing the Tools tab of Advanced settings in RIS, see Manage troubleshooting tools.

You can also make the Recovery Console available to a user of a RIS client computer if the client computer is configured to support it. When the Recovery Console is available to a user, it is displayed in the Client Installation Wizard under "Maintenance and Troubleshooting." For information on using RIS to start the Recovery Console, see Remote Installation Services Resources.

The Remote Installation Services (RIS) tools feature includes the following options:

- **Remove**

  Removes the template setup information file (.sif) associated with the currently selected tool. The .sif file contains the friendly description text, which explains the tool's purpose to the users of a RIS client computer. If you choose the Remove option, the template file is deleted, and, although the tool remains on the disk, it is not offered to RIS clients when they run the Client Installation Wizard.

  **Important**

  - If you want to make a tool temporarily unavailable, or you want to make it available only to certain users, do not remove it; doing so deletes the .sif file associated with the tool. Instead, use the **Properties** option and restrict permissions on the .sif file associated with the tool (**Read** and **Read & Execute** permissions are required by tool users).

- **Properties**

  Provides access to property information for the selected tool. Use this option to change the friendly description text and to retrieve specific details about the selected tool. You can also use this option to change permissions for the tool.

- **Refresh**

  Refreshes the contents of the display window.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Computer names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computer names**

You can use the advanced settings of Remote Installation Services (RIS) to define the naming policy that gives each client computer a unique name. This unique name, which is used to identify the client computer on the network, is similar to the network basic input/output system (NetBIOS) name used in previous versions of Windows and Windows NT.

When you define your naming policy, you can either choose from a list of automatic naming options (for example, **First initial, last name**), or you can click the **Customized** button and create a customized naming policy. If you want to ensure that each client computer is given a unique computer name, choose from the list of automatic naming options. All of the automatic naming options that are based on the user's name (for example, **First initial, last name**) add an incremental number at the end to ensure that the computer name is unique.

To create a custom naming policy, click **Customize**. You can then select one of the options in the following list to create the customized computer name automatically.

**%First**
The user's first name is used as the computer name.

**%Username**
The user name is used as the computer name.

**%[ 0][ *n*] #**
The computer name includes an incremental number containing *n* digits. To include zero-padding in the incremental number, type a zero as shown. For example, if you use %03#, a three-digit number in the range from 001 through 999 is used.

**%Last**
The user's last name is used as the computer name.

**%MAC**
The network card media access control address is used as the computer name.

**Important**

- Do not create a naming convention that could result in duplicate computer names. One way to avoid duplication in computer names is to use the incremental number option described in the previous list.

When you are defining your computer naming policy, you can limit the length of computer names by adding a numerical value to the text string. For example, if you want to limit the computer name to three characters of the user's last name with a two-digit incremental number appended, you would use the string: **%3Last%02#** Thus, for the fifth computer user whose last name is Smith, the corresponding computer name would be Smi05.

You can use customized naming options separately, or you can combine them. For example, if you want the computer name to consist of the first three letters of the user's first name, followed by the first three letters of the user's last name, followed by a three-digit incremental number, use the following string:**%3First%3Last%03#**

For instructions on how to configure the computer naming format, see Configure settings for new clients for RIS.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# The Remote Installation Process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The remote installation process**
This section covers:

- Client Installation Wizard
- Installation options available through RIS
- Operating systems supported by Remote Installation Services
- Controlling client installations through Group Policy

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Client Installation Wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Client Installation Wizard**

Users of a remote boot-enabled client computer can use the Client Installation Wizard to select installation options, operating systems, and maintenance and troubleshooting tools.

The wizard prompts the user for a user name, password, domain name, and other information as required. After the user's credentials have been validated, the wizard displays the installation options that are available to the user. After the user selects an option, the selected operating system installation image is copied to the client computer's local hard disk. You can predetermine some installation options using Group Policy.

Users of the Client Installation Wizard must type only standard ASCII characters (OEM characters 32-126) for user name, password, or domain name information. The Client Installation Wizard does not support extended ASCII character sets (such as those containing ü, é, or other extended characters).

For more information, see the following:

- For more information about how a user logs on through the Client Installation Wizard and performs a remote installation, see Perform a network installation using RIS.
- For information about using Group Policy to control some installation options, see Controlling client installations through Group Policy.

As an administrator, you can edit .osc files to determine which screens the Client Installation Wizard presents to the user. For more information about .osc files, see see "Designing RIS Installations" at the Microsoft Windows Resource Kits Web site. For basic information about other ways you can control some of the choices and screens seen by users of RIS client computers, see Installation options available through RIS.

**Important**

- By default, the client computer's hard disk is reformatted during the operating system installation. Therefore, all locally stored files are removed during installation. To prevent the loss of important data, you should back up your files before installation. You can also choose not to reformat the computer's hard disk during installation.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Installation options available through RIS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installation options available through RIS**
The Client Installation Wizard is a text-based setup wizard that guides the user through the remote installation process. You can use Group Policy to control which installation options are offered to users. For more specific information about this, see Controlling client installations through Group Policy.

You can specify that the user be presented with automatic setup or custom setup, or that a previous setup attempt be restarted. You can also specify that the user be given maintenance and troubleshooting options.

**Automatic Setup**
Use this option if you want to predefine the computer name and a location within Active Directory for the client computer accounts. For more information about computer names and the options for predefining them, see Computer names. You can also prestage the client computer; for more information about this, see Prestage client computers.

**Custom Setup**
Use this option to allow users to define a unique name for this computer and specify where the computer account will be created within Active Directory. Typically, you use custom setup for users outside of your own organization who are setting up computers.

**Restart setup ("Restart a Previous Setup Attempt" in the Client Installation Wizard)**
This option restarts a failed operating system installation attempt if the installation procedure fails during text-mode setup. This option does not copy files from where the previous installation attempt failed. However, the user is not required to answer any questions that were already answered during the previous setup attempt.

**Tools ("Maintenance and Troubleshooting" in the Client Installation Wizard)**
This option gives users access to tools that are used to maintain and troubleshoot client computers. These tools include memory virus scanners, updating the flashing for system basic input/output system (BIOS), and computer diagnostic utilities. In addition, if Recovery Console has been made available to a client computer, and if the client computer is configured to support the Recovery Console, it is displayed in the Client Installation Wizard under "Maintenance and Troubleshooting." For more information about maintenance and troubleshooting tools, see Maintenance and troubleshooting tools and options.

**Additional options and other considerations**

- Remote Installation Services (RIS) does not fully support unattended installations on computers that contain Microsoft Internet Security and Acceleration (ISA) devices or devices that are non-Plug and Play.

- Users who are setting up a remote boot-enabled client computer are not offered installation choices when automatic setup is their only installation option. For more information, see Controlling client installations through Group Policy.

- The unattended installation setup answer files (.sif) are text files that describe the installation process for an installation image. An unattended installation template, which you can customize, is provided by RIS. The .sif file must be associated with one of the available installation images. For information about associating a .sif file with an installation image, see Associate an unattended setup answer file with a RIS installation image. For information about creating unattended installation files and automating installations, see "Automating and Customizing Installations" at the Microsoft Windows Resource Kits Web site.

  One of the things you can control through a .sif file is the enabling of the Administrator account (which by default is disabled when the client joins the domain). If you do this, be sure to use a tool called Setup Manager to ensure that the Administrator password is encrypted in the .sif file before being transmitted over the network. For information about this, see the following:

  - For specific information about working with .osc and .sif files for RIS, see "Designing RIS Installations" at the Microsoft Windows Resource Kits Web site.

  - For information about Setup Manager, see Deploy.chm in \Support\Tools\deploy.cab on the Windows CD.

- By default, the Client Installation Wizard uses the Welcome.osc file located in \RemoteInstall\OSChooser to manage client installation image choices. For multiple language installation image options, you must remove the Welcome.osc file and rename the MultiIng.osc file to Welcome.osc. The Client Installation Wizard then provides the user with a menu of multiple language choices. You can also edit Welcome.osc to create custom language options. For more information about .osc files, see "Designing RIS Installations" at the Microsoft Windows Resource Kits Web site.

- For most client computers, the user must press a particular key (for example, F12) to initiate the remote installation process. Some client computers have settings in the basic input/output system (BIOS) that cause a prompt to be displayed when it is time for the user to press the key. Not all client computers display a prompt. By default, RIS causes a prompt to be displayed on the client computer, telling the user to press F12. To change this so that RIS does not cause a prompt to be displayed, in *RIS_server_name*\REMINST\oschooser\i386, rename Startrom.com to Startrom.old, and then rename Startrom.n12 to Startrom.com.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Operating systems supported by Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Operating systems supported by Remote Installation Services**
You can use Remote Installation Services (RIS) to enable automated installations of an operating system on new or existing client computers.

Two types of operating system images are supported by the Client Installation Wizard: flat image and the Remote Installation Preparation Wizard (RIPrep) image formats.

Using a flat image is similar to setting up a client directly from the installation CD, except that the operating system image source files reside on an available RIS server. For information about creating a flat image, see Add a new client operating system installation image.

When you create an RIPrep image, you can install the necessary client computer applications and configure the client computer settings, such as monitor settings and network printers. For background information about RIPrep, see Creating an installation image with RIPrep. For instructions on how to run RIPrep, see Create a Remote Installation Preparation Wizard image.

The Client Installation Wizard displays both the flat image and RIPrep image formats in the list of available installation options. The installation process begins when the user of the client computer selects one of the displayed operating system images.

The table below lists the operating systems supported by Remote Installation Services. It also shows the image type that is supported for each operating system.

| Operating system | Flat image | RIPrep image |
|---|---|---|
| Windows XP Professional | X | X |
| Windows Server 2003, Standard Edition | X | X |
| Windows Server 2003, Web Edition | X | X |
| Windows Server 2003, Enterprise Edition | X | X |
| 64-bit version of Windows Server 2003, Enterprise Edition | X | X |
| Windows 2000 Professional | X | X |
| Windows 2000 Server | X | X |
| Windows 2000 Advanced Server | X | X |

**Important**

- A 64-bit version of the Remote Installation Preparation (RIPrep) Wizard is included with the 64-bit version of Windows Server 2003, Enterprise Edition. For more information about using the RIPrep Wizard to create an installation image of a 64-bit operating system, see article 891128, "Updated Remote Installation Services (RIS) functionality in Windows Server 2003 Service Pack 1," in the Microsoft Knowledge Base.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Controlling client installations through Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Controlling client installations through Group Policy**

You can use the Group Policy snap-in to control whether the user of a RIS client computer must supply a computer name during installation, whether the remote installation attempt should be restarted if it fails, and whether users have access to tools before the operating system is installed. Group Policy is applied to sites, domains, and organizational units, and can be inherited within Active Directory. For more information on Group Policy, see Group Policy overview.

The Group Policy options for the Client Installation Wizard are as follows:

**Automatic Setup**

Use this option if you want to predefine the computer name and a location within Active Directory for the client computer accounts. For more information about computer names and the options for predefining them, see Computer names. You can also prestage client computers; for more information, see Prestage client computers.

**Custom Setup**

Use this option to allow users to define a unique name for this computer and specify where the computer account will be created within Active Directory.

**Restart Setup**

This option restarts an operating system installation attempt if it fails prior to completion.

**Tools**

Use this option to allow users to access tools from the Client Installation Wizard. For more information about how tools are made available, see Maintenance and troubleshooting tools and options.

For each of the Client Installation Wizard choices, the following options are available:

**Enable**

Users to whom this policy applies are offered this installation option.

**Not Configured**

The policy settings of the parent container apply to this option. For example, if you choose **Not Configured** and the administrator for the entire domain has set Group Policy specific to Remote Installation Services, the policy that is set on the domain is applied to all users who are affected by that policy.

**Disabled**

Users affected by this policy cannot access this installation option.

For procedural information on controlling access through Group Policy, see Set Group Policy for client installation options.

For information on the options offered by the Client Installation Wizard, see Installation options available through RIS.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Remote Installation Services Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

- For more information about deploying Remote Installation Services (RIS), see "Designing RIS Installations" at the Microsoft Windows Resource Kits Web site. For information about the Resource Kits, see Using the Windows Deployment and Resource Kits.
- For information about unattended installation, see deploy.chm in \Support\Tools\deploy.cab on the Windows CD.
- For information about compatible hardware, Windows Update, and product support, see Support resources.
- For information about starting the Recovery Console from RIS, see article Q222478, "Creating a Template to Run Recovery Console Using a Remote Install Server," in the Microsoft Knowledge Base.
- For information about using File Replication service (FRS) to replicate RIS images, see article Q273594, "How to: Replicate RIS Images Using DFS and NTFRS," in the Microsoft Knowledge Base.
- For information about configuring a router to forward DHCP packets to a client computer, see article Q257579, "PXE Clients Do Not Receive an IP Address from a DHCP Server Across a Router," in the Microsoft Knowledge Base.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Troubleshooting Remote Installation Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting Remote Installation Services**

What problem are you having?

- I am not sure I have the correct Pre-Boot eXecution Environment (PXE) ROM version.
- I am not sure whether the client computer has received an Internet protocol (IP) address and has contacted the Remote Installation Services (RIS) server.
- I am trying to locate the globally unique/universally unique identifier (GUID/UUID) on a client computer so I can prestage the client computer within Active Directory for use with RIS.
- Command settings are not being processed during the unattended installation.
- Language choice options do not appear during the Client Installation Wizard session.
- The client computer is prestaged to a RIS server but is being serviced by a different server.
- I have client computers that use the older remote boot protocol, Remote program Load (RPL)-based ROMs.
- I am investigating questions of security related to the Pre-Boot portion of the PXE-based remote boot ROM.
- I am evaluating methods for replicating all of the operating system installation images currently located on one RIS server to other RIS servers on the network to achieve consistency across all client installations.
- I am evaluating whether to place a RIS server and another vendor's remote boot server on the same network.
- I want to manage the RIS servers remotely from computers running Windows XP Professional on my network.
- I want to add more network adapters to the RIS boot disk.
- I want to know whether I can use the Active Directory object attributes to create a naming format that can be used with the RIS automatic computer naming feature.
- After the restoration of a backup of a RIS volume, RIS does not function properly.
- I created an image on my RIS server, but the installation image does not appear in the list of available operating system installation options when I attempt to install the image on a client computer.
- I created an image using the Remote Installation Preparation (RIPrep) Wizard, started the Add Wizard, and clicked Associate a new answer file to an existing image, but I do not see the RIPrep image in the list in the Add Wizard.
- I am evaluating the use of RIPrep in a situation where I want to preserve the file attributes and security settings defined on the source computer.
- I want to use RIPrep in a situation where there are hardware differences between the source computer that is used to create the RIPrep based installation image and the destination computer that will install the image.
- I am evaluating whether I can use RIPrep to support an installation on multiple disks or multiple partitions on the client computer.

- I am evaluating the use of RIPrep in a situation where the disk sizes differ between the source computer used to create the image and the destination computer that will receive the image.

**I am not sure I have the correct Pre-Boot eXecution Environment (PXE) ROM version.**
**Solution:** When the Net PC or client computer containing a remote boot ROM starts, the version of the PXE ROM appears on the screen. Remote Installation Services (RIS) supports .99n or greater PXE ROMs, except in a few situations that require the .99L version. You might be required to obtain a newer version of the PXE-based ROM code from your original equipment manufacturer (OEM) if you have problems with the existing ROM version installed on a client computer.

**I am not sure whether the client computer has received an Internet protocol (IP) address and has contacted the Remote Installation Services (RIS) server.**
**Solution:** When the client computer boots, you will see the PXE boot ROM begin to load and initialize. The following sequence occurs with most PC98 and Net PCs, PXE ROM-based computers, and the computers using the RIS boot disk.

**Remote Boot ROM Load Sequence:**

**Step 1:**

The client computer displays the message "DHCP." This indicates that the client is requesting an Internet protocol (IP) address from the DHCP server. This can also mean that the client has obtained an IP address from the DHCP server and is awaiting a response from the RIS server. To verify that the client is receiving an IP address, check the IP leases that have been granted on your DHCP server.

**Troubleshooting**

If the client does not receive the message, an IP address might not have been received or the Remote Installation (BINLSVC) server might not be responding. Consider the following:

- Is the DHCP server available and has the service started? The DHCP server and remote installation servers must be authorized in Active Directory before they can provide IP addresses. Check to ensure that the service has started and that other clients that are not remote boot-enabled are receiving IP addresses on this segment.

- Does the DHCP server have a defined IP address scope, and has it been activated?

- Is there a router between the client and the DHCP server that is not allowing DHCP packets through?

- Are there any error messages in the event log under the system log for DHCP?

**Step 2:**

When the client receives an IP address from the DHCP server, the message might change to "BINL." This indicates that the client successfully leased an IP address and is now waiting to contact the RIS server. The client will eventually time out and post the error message "No Bootfile received from DHCP, BINL, or Bootp."

**Troubleshooting**

If the client does not receive the BINL message, this indicates the client is not receiving a response from the RIS server. Consider the following:

- Is the server available and has Remote Installation Services started? RIS servers must be authorized before they can respond to client computers. Run **Risetup -Check** to authorize a RIS server. For more information, see Remote Installation Services server authorization.

- Are other remote boot-enabled clients receiving the Client Installation Wizard? If so, this client computer either is not supported or is having remote boot ROM-related problems. Check the version of the PXE ROM on the client computer. Also, check Active Directory to see whether the administrator has prestaged this client computer to a RIS server that is offline or unavailable to the client computer.

- Has an additional network card been added to the RIS server? If so, the server must be re-authorized in Active Directory. Run **RISETUP -Check** to re-authorize the server.

- Is a router between the client and the RIS server not allowing the DHCP-based requests or responses through? The RIS

server communicates by way of the DHCP packet type during the initial service request/response sequence. You might need to configure the router to forward the DHCP packets. For more information, see article Q257579, "PXE Clients Do Not Receive an IP Address from a DHCP Server Across a Router," in the Microsoft Knowledge Base.

- Are there any error messages in the event log under the system or application logs specific to Remote Installation (BINLSVC), Domain Name System (DNS), or Active Directory?

- Does the client computer have an account in Active Directory? If not, the RIS server might be configured to respond only to known client computers.

- Is the RIS server also a DHCP server, and does it have multiple network adapters? If so, then there must be a scope for the subnet of each network adapter. For more information, see Configuring scopes.

**Step 3:**

The client then changes to Trivial File Transfer protocol (TFTP) or prompts the user to press a particular key (for example, F12). This indicates that the client has contacted the RIS server and is waiting to receive the first image file-CIW. You might not see the BINL and TFTP message on some computers because this sequence can occur very rapidly.

**Troubleshooting**

If the client computer does not get a response from the RIS server, the client will time out and send an error message saying that it did not receive a file from DHCP, BINL, or TFTP. In this case, the RIS server did not answer the client computer. Do the following:

- Stop and restart Remote Installation (BINLSVC). To do so, click **Start**, and then click **Run**. In the **Run** dialog box, type **cmd** in the text field, and click **OK**. In the console window, type the following:

  **Net Stop BINLSVC**

  **Net Start BINLSVC**

- Check the RIS server properties to ensure that the **Respond to client computers requesting service** option is selected and that **Do not respond to unknown client computers** is cleared unless you have prestaged the client computer in Active Directory prior to starting the client computer.

- If the client computer does not receive an answer after it has attempted to stop and restart the service, and you have checked the remote installation server object properties to ensure that the correct settings have been set, check the event log on the RIS server for any errors relating to DHCP, DNS, and RIS (BINLSVC).

- Check the TFTP Daemon to ensure that it has actually started, and check the event log to see if there are any TFTP errors.

**Step 4:**

At this point, the client should have downloaded and displayed the Client Installation Wizard application; a Welcome screen greeting the user should appear.

**I am trying to locate the globally unique/universally unique identifier (GUID/UUID) on a client computer so I can prestage the client computer within Active Directory for use with RIS.**
**Solution:** The GUID/UUID for most client computers that are PC98 or Net PC-compliant is provided in the system basic input/output system (BIOS) of the computer. Computer manufacturers are encouraged to ship either a floppy disk containing a comma-separated file or spreadsheet that includes a mapping of serial numbers to GUIDs/UUIDs. You can use this information to script the prestaging of client computers within Active Directory.

The GUID/UUID also appears on the outside of the computer. If it does not appear in the expected locations, one method for discovering it is to use a network utility to sniff the network traffic of the client computer and locate the DHCPDiscover packet; within that field will be the 128-bit, 16-byte GUID/UUID. Another method is to use the Windows Management Instrumentation (WMI) interface-- specifically, the universally unique identifier (UUID) value of the Win32_ComputerSystemproduct class. For more information about this, see the MSDN Library at the Microsoft Web site.

**Command settings are not being processed during the unattended installation.**

**Cause:** Incorrect directory information is being used. Correct directory information is required when the "OemPreinstall = yes" setting in a .sif file is being used.

**Solution:** Change the directory information to \\RemoteInstall\Setup*applicable_language*\Images*applicable_name*\$oem$.

**Language choice options do not appear during the Client Installation Wizard session.**
**Cause:** By default, RIS uses the Welcome.osc file to manage client installation image choices. To make multiple language installation image options available, you must replace the default Welcome.osc file with the Multilng.osc file.

**Solution:** The Client Installation Wizard uses the Welcome.osc file located in the \\RemoteInstall\OSChooser folder to manage client installation image choices. When you remove the Welcome.osc file and rename the Multilng.osc file to Welcome.osc, the Client Installation Wizard will also offer a menu of language choices to the user. You can edit the Welcome.osc file to create custom language options.

**The client computer is prestaged to a RIS server but is being serviced by a different server.**
**Cause:** When you prestage a client computer into a domain with multiple domain controllers, the replication delay of the client computer account information can cause a client computer to be serviced by another RIS server.

**Solution:** You can either wait for the computer account information to be propagated during the next scheduled replication session, or you can modify the replication frequency between your domain controllers. For more information on configuring replication, see Replication overview.

**I have client computers that use the older remote boot protocol, Remote program Load (RPL)-based ROMs.**
**Solution:** There is no support for the older RPL-based remote boot, so computers using this cannot use RIS. RIS uses the PXE DHCP-based remote boot ROMs.

**I am investigating questions of security related to the Pre-Boot portion of the PXE-based remote boot ROM.**
**Solution:** There are inherent security limitations with PXE. The entire boot ROM sequence and operating system installation or replication process is not secure for packet type encryption, client or server spoofing, or wire sniffer-based mechanisms. You should exercise caution when using RIS on your corporate network. Allow only authorized RIS servers on your network. Also, control the number of administrators who are allowed to install or configure RIS servers. For more information, see Security Information for Remote Installation Services.

**I am evaluating methods for replicating all of the operating system installation images currently located on one RIS server to other RIS servers on the network to achieve consistency across all client installations.**
**Solution:** RIS does not currently provide a mechanism for replication of operating system images from one RIS server to another. There are several alternative mechanisms you can use, however. One solution is to use the strong replication features of the Systems Management Server product. This product provides for scheduled replication, compression, and low-speed link features. You can also use other vendor solutions for operating system image replication. If you do this, however, ensure that the replication mechanism you choose supports maintaining the file attributes and security settings of the source images.

**I am evaluating whether to place a RIS server and another vendor's remote boot server on the same network.**
**Solution:** You can place multiple-vendor remote boot/installation (RB/RI) servers on one physical network. It is important to understand, however, that the remote boot PXE ROM code is not currently able to detect the difference between vendors' RB/RI servers. As a result, when a remote boot-enabled client computer starts and requests the Internet protocol (IP) address of an RB/RI server, all of the available servers will respond to that client; thus, the client has no way to ensure that it is being serviced by a specific RB/RI server.

With RIS, however, an administrator can prestage client computers into Active Directory and determine which RIS server will service a client computer. By configuring the RIS server to answer only known (prestaged) client computers, the administrator is assured that the client will be serviced by the correct RIS server.

Not all of the other RB/RI vendors are able to ignore service requests. You might need to isolate the specific vendors' servers on the network so that clients are not answered by these vendors' RB/RI servers.

**I want to manage the RIS servers remotely from computers running Windows XP Professional on my network.**
**Solution:** If you are an administrator in the domain, and you have installed the Windows Server 2003 Administration Tools Pack on the computer running Windows XP Professional, you can administer most of the RIS configuration settings. There are some items that you cannot manage, however. For example, you cannot remotely add more operating system installation images to the RIS servers from a computer running Windows XP Professional.

**I want to add more network adapters to the RIS boot disk.**
**Solution:** For this release of RIS, the Rbfg.exe utility cannot be modified to accommodate additional network adapters. Updates to Rbfg.exe can be made available through the same distribution channels that are used for other software updates and feature packs, however. You can obtain information about possible updates to Rbfg.exe (and the network adapter support in those updates) in the same way that you obtain information about other software updates.

**I want to know whether I can use the Active Directory object attributes to create a naming format that can be used with the RIS automatic computer naming feature.**
**Solution:** This is not supported. Although the existing attributes supported with the automatic computer naming feature do use Active Directory, not all Active Directory object attributes are currently supported.

**After the restoration of a backup of a RIS volume, RIS does not function properly.**
**Cause:** Backup restored the volume without a Single Instance Store (SIS) directory.

**Solution:** Verify the configuration of the RIS volume, then restore the volume again.

For information about how to verify the configuration of a RIS volume, see Verify the Remote Installation Services configuration.

For information about how to back up and restore SIS directories, see "Backup and Restore" at the Microsoft Windows Resource Kits Web site.

**I created an image on my RIS server, but the installation image does not appear in the list of available operating system installation options when I attempt to install the image on a client computer.**
**Cause:** The client computer has a different type of hardware abstraction layer (HAL) than the image.

**Solution:** When you use RIPrep to create new images, RIPrep ensures that the images viewed by the client computers are only those that are compatible with the computer. A client computer must have the same type of hardware abstraction layer (HAL) as the image. If it does not, the image will not appear in the list of available operating system installation options.

**Cause:** Permissions have been set that prevent users from seeing a particular installation image.

**Solution:** For information about setting permissions that permit or deny access to installation images, see Allow or prevent the installing of a RIS image by a user or group and Allow or prevent the viewing and installing of a RIS image by a user or group.

**I created an image using the Remote Installation Preparation (RIPrep) Wizard, started the Add Wizard, and clicked Associate a new answer file to an existing image, but I do not see the RIPrep image in the list in the Add Wizard.**
**Solution:** The installation images that appear in the list in the Add Wizard include only flat images; that is, those images created with the Remote Installation Services Setup Wizard. The list does not include images created with the Remote Installation Preparation (RIPrep) Wizard.

**I am evaluating the use of RIPrep in a situation where I want to preserve the file attributes and security settings defined on the source computer.**
**Solution:** RIPrep supports this. The file attributes and security settings that are defined on the source computer are preserved on the destination computer that installs that image. For more information about RIPrep, including information about requirements for the RIPrep replication process, see Creating an installation image with RIPrep.

**I want to use RIPrep in a situation where there are hardware differences between the source computer that is used to create the RIPrep based installation image and the destination computer that will install the image.**
**Solution:** You can use RIPrep in a situation where hardware on the source computer and the destination computer is different. The one limitation is that the computers must use the same type of Hardware Abstraction Layer (HAL). For example, if the source computer is an Advanced Configuration Power Interface (ACPI)-based computer, it uses a HAL that includes ACPI. If you attempt to install this RIPrep image on a computer that is not based or enabled on ACPI, the operating system installation process will fail.

**I am evaluating whether I can use RIPrep to support an installation on multiple disks or multiple partitions on the client computer.**
**Solution:** In this release of Remote Installation Services, the RIPrep utility supports only a single disk with a single partition. You can use RIPrep to make an image of a computer that contains more than one partition or disk, but RIPrep will only image the system partition.

**I am evaluating the use of RIPrep in a situation where the disk sizes differ between the source computer used to create the image and the destination computer that will receive the image.**
**Solution:** For details about working with different disk characteristics on the source client computer and the destination computer, see Creating an installation image with RIPrep.

### Note

- This topic does not apply to Windows Server 2003, Web Edition.

# Remote Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Administration**
Using tools included with Microsoft® Windows Server 2003 family operating systems CDs, you can remotely manage your system servers running Microsoft® Windows 2000 and Windows Server 2003 family operating systems. You can also remotely manage Windows Server 2003 family operating systems from a computer using Microsoft® Windows® XP Professional.

- For general information, see Remote administration overview.
- For tips about managing your server remotely, see Remote Administration Best practices.
- For other sources of information, see Remote Administration Resources.

# Remote administration overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote administration overview**

The following table presents remote management tasks and links to information about the tools. Unless indicated otherwise, these tools are used on Microsoft® Windows Server™ 2003 operating systems and, in some cases, Microsoft® Windows® XP Professional.

| Task | Tools | Comments |
|---|---|---|
| Manage several servers by performing similar tasks. | Microsoft Management Console<br><br>Author an MMC console file<br><br>Work with MMC and saved console files | You can use appropriate saved MMC consoles, if available, or you can create custom MMC consoles for tasks you frequently either delegate or perform. |
| Log onto one computer remotely and manage it as if you were there. | Remote Desktop Connection | In addition to Windows Server 2003 operating systems, this feature can be used with Microsoft® Windows® 2000 Server. The Remote Desktop feature must be enabled on the remote computer. |
| Switch among several computers that you want to log on to remotely. | Remote Desktops snap-in | In addition to Windows Server 2003 operating systems, this feature can be used with Windows 2000 Server. The Remote Desktop feature must be enabled on the remote computer. |
| Administer a server running Windows Server 2003 from any computer on your network. | Remote Desktop for Administration | In addition to Windows Server 2003 operating systems, this feature can be used with Windows 2000 Server. The Remote Desktop feature must be enabled on the remote computer. |
| Manage servers from Windows XP Professional. | Windows Server 2003 Administration Tools Pack Overview<br><br>Remote Desktop Connection | Windows Server 2003 Administration Tools Pack must be installed on Windows XP Professional to permit server management using these tools. The Remote Desktop feature must be enabled on servers to permit use of the Remote Desktop Connection. |
| Manage your Web server using a Web browser on a remote computer. | Using Web Interface for Remote Administration | Use Internet Explorer, version 6 or later, for remote administration using Web Interface. Web Interface is not installed by default for Web servers. |
| Assist in real time an operator on a remote computer running Windows Server 2003 family operating system, Windows XP Professional or Microsoft® Windows® XP Home Edition. | Administering Remote Assistance (www.microsoft.com) | The Remote Assistance connection requires explicit permission from the connected user, and both computers must be running Windows XP Professional or Windows Server 2003 operating systems. |
| Connect to a remote computer when that computer cannot access a network or is not in a functional state because of hardware or software failure. | Telnet | All commands must be entered through the command line. Capabilities are limited, and security is minimal. |

| | | |
|---|---|---|
| Manage disk space on servers by copying files to a library of magnetic tapes or magneto-optical disks. | Remote Storage | Some Windows components require configuration before they can be used and will be displayed in the **Add or Remove Windows Components** in **Add or Remove programs** in **Control Panel**. Remote storage requires installation. |

**Notes**

- Depending on the Group Policy settings on your network, you might not be able to perform some of these tasks.
- Some features require that specific ports be opened on the server in order for the role to be administered remotely. For information about these port requirements, see Service overview and network port requirements for the Windows Server system (http://go.microsoft.com/fwlink/?LinkID=61531) and Windows Firewall Settings (http://go.microsoft.com/fwlink/?LinkID=38651).
- You can use command-line tools to automate many processes and tasks. For more information, see Automating processes and tasks.
- You must physically access each computer to restart, change hardware, and enable the Remote Desktop feature. Once the Remote Desktop feature has been enabled, it will remain enabled.
- For an explanation of the various Remote Desktop terminologies, see Remote Administration using Terminal Services.
- For more information about remote administration, see the following:
    - Common Administrative Tasks
    - New Features
    - Quick Guide for Finding Tools
    - Remote Administration Resources

# Remote Administration Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**To perform many tasks on a single, remote server, use Remote Desktop for Administration.**

- To remotely administer a server running a Windows Server 2003 family operating system, you do not need to install Terminal Server. Instead, use Remote Desktop for Administration. On these operating systems, Remote Desktop for Administration is on by default but is not enabled. After you enable remote connections, Remote Desktop for Administration allows you to administer your server from any computer on your network running Windows Server 2003 family, Windows XP Professional, and Windows 2000. Windows Server 2003 , Web Edition, does not support remote administration.

- You can use Remote Desktop with Windows 2000 Server. You need to install Terminal Services for Remote Administration as described in the Windows 2000 Help.

  For more information, see Enable or disable Remote Desktop and Remote Desktop for Administration.

**To do similar tasks on many servers, use MMC snap-ins.**

- Most of the administration tools provided with Windows Server 2003 family operating systems are Microsoft Management Console (MMC) snap-ins that can be used to administer remote servers as well as your local computer. The advantages of using MMC include the following:

  - You can create a console with the tools you use for the tasks you do most often.

  - You can set the focus for a tool to any of the servers you administer, and switch between servers and tools within a single MMC console.

  For more information, see Microsoft Management Console.

**To manage servers running Windows Server 2003 family operating systems from a computer using Windows XP Professional , use Windows Server 2003 Administration Tools Pack**

- Windows Server 2003 Administration Tools Pack when installed on Windows XP Professional (ordinarily a client operating system) provides many remote management tools for managing your servers.

- Windows Server 2003 Administration Tools Pack comes as adminpak.msi on the Windows Server 2003 family operating system and is easily installed from the server CD onto Windows XP Professional by installing the .msi from the server to the Windows XP Professional operating system.

- The Administration Tools Packs in the Windows Server 2003 family and Windows 2000 are not mutually compatible. You must use the Windows 2000 Administrative Tools to administer Windows 2000, and the Windows Server 2003 Administration Tools Pack to administer the Windows Server 2003 family operating system.

  For more information, see Windows Server 2003 Administration Tools Pack Overview.

# Remote Administration Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**
The following resources provide additional information about remote administration.

**Windows Server 2003 Resource Kit**
The Microsoft Windows Server 2003 Resource Kit contains additional tools that you can use for remote administration.

**Microsoft Platform Software Development Kits (SDK)**
For technical information about Active Directory Scripting Interfaces (ADSI), Microsoft Management Console (MMC), Windows Management Instrumentation (WMI) Software Development Kits (SDK), see the Microsoft Web site.

**TechNet**
TechNet is a service that contains resources for information technology (IT) professionals. This includes access to articles written by support professionals at Microsoft, career and industry information, discussion groups, technical feedback, and software.

For more information, see Microsoft TechNet.

# Terminal Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services**
Terminal Services provides remote access to a Microsoft® Windows® desktop through "thin client" software, allowing the client computer to serve as a terminal emulator. Terminal Services transmits only the user interface of the program to the client. The client then returns keyboard and mouse clicks to be processed by the server. Each user logs on and sees only their individual session, which is managed transparently by the server operating system and is independent of any other client session. To connect to a computer via Terminal Services, use Remote Desktop Connection (the new Terminal Services client). For more information, see Remote Desktop Connection.

Terminal Server uses Terminal Services technology to provide an effective and reliable way to distribute Windows-based programs with a network server. With Terminal Server, a single point of installation allows multiple users access to the desktop of a computer running one of the Microsoft Windows Server™ 2003 family operating systems, where they can run programs, save files, and use network resources. Terminal Server is not installed by default. For more information about Terminal Server, see Terminal Server. (Terminal Server is not available on Windows Server 2003, Web Edition.)

**Important**

- It is not necessary to install Terminal Server for remote administration of your server. Remote Desktop for Administration is installed by default. To use Remote Desktop for Administration, you must first enable remote connections. For more information, see Enable or disable Remote Desktop. For more information about Remote Desktop for Administration, see Remote Desktop for Administration.

Terminal Server has its own method for licensing clients that log on to terminal servers, separate from the licensing method for clients running one of the Windows Server 2003 family operating systems. Clients must receive a valid license issued by a license server before they are allowed to log on to a Terminal Server. For more information about Terminal Server Licensing, see Terminal Server Licensing.

A Terminal Services connection provides the link clients use to log on to a session on the server. A TCP/IP connection is configured on port 3389 when Terminal Services is installed. Using Terminal Services Configuration, you can change the default properties of the connection, as well as other settings that apply to Terminal Services sessions. For more information about Terminal Services configuration, see Terminal Services Configuration.

User accounts can be set up and managed for a specific computer using Local Users and Groups or for a domain of computers using Active Directory Users and Computers. For more information about Terminal Services user properties, see Terminal Services User Properties.

Terminal Services Manager can be used to view information about terminal servers that reside in trusted domains. Use this tool to monitor users, sessions, and applications on each terminal server, and to carry out assorted actions to manage the server. For more information about Terminal Services Manager, see Terminal Services Manager.

# Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Server**
On Microsoft® Windows Server™ 2003 family operating systems, the Terminal Server feature gives users at client computers throughout your network access to Windows-based programs installed on terminal servers. With Terminal Server, you can provide a single point of installation that allows multiple users access to the Windows Server 2003 family operating system desktop, where they can run programs, save files, and use network resources, all from a remote location, as if these resources were installed on their own computers.

For computers running Microsoft® Windows® XP or Windows Server 2003 family operating systems, the Terminal Services client program (Remote Desktop Connection) is already installed. Remote Desktop Connection can also be installed on other 32-bit Windows-based operating systems.

Terminal Server is not included in Microsoft® Windows Server™ 2003, Web Edition.

- Before you begin using Terminal Server, see Checklists: Setting up Terminal Server.
- To find features that have been moved in this release, see New ways to do familiar Terminal Server tasks.
- For tips about using Terminal Server, see Terminal Server Best practices.
- For help with specific tasks, see Terminal Server How To....
- For general background information, see Terminal Server Concepts.
- For problem-solving instructions, see Troubleshooting Task Scheduler.

**Important**

- It is not necessary to install Terminal Server for remote administration of your server. Remote Desktop for Administration is installed by default. To use Remote Desktop for Administration, you must first enable remote connections. For more information, see Enable or disable Remote Desktop. For more information about Remote Desktop for Administration, see Remote Desktop for Administration.

**Note**

- For additional information about deploying Terminal Server, see "Hosting Applications with Terminal Server" at the Microsoft Windows Resource Kit Web site.

# Checklists: Setting up Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists: Setting up Terminal Server**

- Checklist: Plan your Terminal Server deployment
- Checklist: Configure Terminal Server
- Checklist: Configure Terminal Server Licensing
- Checklist: Install applications
- Checklist: Choose client software deployment options

# Checklist: Plan your Terminal Server deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Plan your Terminal Server deployment**

| Step | Reference |
|---|---|
| Read the "Guidelines for Deploying Terminal Server" white paper and the Terminal Services overview. | Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627)<br><br>Terminal Services overview |
| Review information about hosting applications with Terminal Server. | Using Terminal Server to host applications centrally<br><br>Hosting Applications with Terminal Server (http://go.microsoft.com/fwlink/?LinkID=4746) |
| Plan your hardware configuration. | System Planning |
| Follow pre-installation recommendations. | Before Installing Terminal Server |
| Determine whether you need to change the Internet Explorer Enhanced Security Configuration setting for Internet Explorer.<br><br>**Note**<br><br>• This applies to servers running one of the Windows Server 2003 operating systems where the installation of Terminal Server was performed using the Configure Your Server Wizard. If you install Terminal Server using **Add or Remove Programs**, you can change this setting during the installation process. | Terminal server role: Configuring a terminal server; Internet Explorer Enhanced Security Configuration |
| Understand how to use Terminal Services administrative tools. | Terminal Services administrative tools |
| Read the "Windows Server 2003 Terminal Server Licensing" white paper and the Terminal Server Licensing overview.<br><br>**Important**<br><br>• You must configure Terminal Server Licensing correctly in order for your terminal server to continue to accept connections from clients. To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent client access license (CAL), or after 120 days, whichever comes first. | Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220); Terminal Server Licensing overview |
| Review program considerations. | Program Considerations |

Checklist: Configure Terminal Server

# Checklist: Configure Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configure Terminal Server**

| Step | Reference |
|---|---|
| Configure Terminal Server by using Terminal Services Configuration. | Configure Terminal Services Connections |
| Read about using Network Load Balancing with terminal servers. | Load balancing and terminal servers |
| Review the steps for enabling and configuring Network Load Balancing. | Checklist: Enabling and configuring Network Load Balancing |
| Configure Session Directory settings. | Configure Session Directory settings |
| Configure Terminal Server Licensing. This step is required.<br>**Caution**<br>• You must configure Terminal Server Licensing correctly in order for your terminal server to continue to accept connections from clients. To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent client access license (CAL) or after 120 days, whichever comes first. | Checklist: Configure Terminal Server Licensing |
| Install applications. | Checklist: Install applications |
| Choose client software deployment options | Checklist: Choose client software deployment options |

Checklist: Plan your Terminal Server deployment

Checklist: Configure Terminal Server Licensing

# Checklist: Configure Terminal Server Licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configure Terminal Server Licensing**

| Step | Reference |
|---|---|
| • Read the "Windows Server 2003 Terminal Server Licensing" white paper and the Terminal Server Licensing overview. | Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220)<br><br>Terminal Server Licensing overview |
| Choose a Terminal Server license server role. | Terminal Server license server roles |
| Install Terminal Server Licensing. | Install Terminal Server Licensing |
| Choose a license server activation method. | Activating a Terminal Server license server |
| Activate the license server. | Activate a Terminal Server License Server |
| Review information about purchasing and installing client access licenses (CALs). | Purchasing and installing client access licenses on a Terminal Server license server |
| Decide what type of CAL to purchase. | Choosing the Licensing Model (http://go.microsoft.com/fwlink/?LinkId=100439) |
| Purchase CALs. | Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) |
| Install CALs.<br><br>**Caution**<br>You must configure Terminal Server Licensing correctly in order for your terminal server to continue to accept connections from clients. To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent client access license (CAL), or after 120 days, whichever comes first. | Install Client Access Licenses |
| Ensure that the terminal server can detect the license server.<br><br>**Note**<br>Although Terminal Server attempts to detect a license server automatically, it is recommended that you explicitly specify a preferred license server that a terminal server connects to. | Set preferred Terminal Server license servers |

| | |
|---|---|
| Ensure that the Terminal Server mode matches the type of CALs purchased. | Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) |

Checklist: Configure Terminal Server

Checklist: Install applications

# Checklist: Install applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Install applications**

| | Step | Reference |
|---|---|---|
| | Review program compatibility information. | Program Compatibility |
| | Prepare for installation. | Getting ready to install programs |
| | Install programs. | Starting the installation |
| | Test the application installation. | Testing the installation |

Checklist: Configure Terminal Server Licensing

Checklist: Choose client software deployment options

# Checklist: Choose client software deployment options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Choose client software deployment options**

| | Step | Reference |
|---|---|---|
| | Determine Remote Desktop Connection deployment method. | About the Terminal Services Client MSI Setup Package |
| | Install Remote Desktop Connection. | Install Remote Desktop Connection (32-bit computers) |
| | Review Remote Desktop Web Connection concepts. | Deploying Remote Desktop Web Connection |
| | Add users to the Remote Desktop Users group. | Add users to the Remote Desktop Users group |

Checklist: Install applications

# New ways to do familiar Terminal Server tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for Terminal Services. The user interface for performing these tasks is different in Windows Server 2003 family operating systems from the way it was in Microsoft Windows NT® 4.0 or Microsoft Windows® 2000.

Please choose your upgrade path:

- Windows NT 4.0 to one of the Windows Server 2003 family operating systems
- Windows 2000 to one of the Windows Server 2003 family operating systems

**Windows NT 4.0 to one of the Windows Server 2003 family operating systems**

| If you want to | In Windows NT 4.0 use | In Windows Server 2003 family operating systems use |
|---|---|---|
| Connect to a terminal server | **Terminal Services Client** in **Program Manager** | Open Remote Desktop Connection.<br><br>For more information, see Remote Desktop Connection |
| Configure a terminal server | **Terminal Server Connection Configuration** in **Administrative Tools** | To configure groups of servers or individual servers, open Group Policy and use Terminal Services Group Policies.<br><br>To configure a terminal server locally, use Terminal Services Configuration.<br><br>For more information, see Terminal Services Configuration Overview |
| Use Terminal Services for Remote Administration | Not available | **Remote Desktop** (installed by default).<br><br>For more information, see Remote Desktop for Administration. |
| Manage the terminal server and view all sessions, users, and processes | **Terminal Server Administration** in **Administrative Tools** | Terminal Services Manager.<br><br>For more information, see Terminal Services Manager. |
| Manage Terminal Server licenses | **Terminal Server License Manager** in **Administrative Tools** | Terminal Server Licensing.<br><br>For more information, see Terminal Server Licensing. |
| Manage Terminal Server user accounts | **User Manager** in **Administrative Tools** | Terminal Services Group Policies.<br><br>For more information, see Managing Terminal Services users with Group Policy |

**Windows 2000 to one of the Windows Server 2003 family operating systems**

| If you want to | In Windows 2000 use | In Windows Server 2003 family operating systems use |
|---|---|---|
| Connect to a terminal server | **Terminal Services Client** and **Client Connection Manager** | Open Remote Desktop Connection.<br><br>For more information, see Remote Desktop Connection |

| | | |
|---|---|---|
| Configure a terminal server | **Terminal Services Configuration** in **Administrative Tools** | To configure groups of servers or individual servers, open Group Policy and use Terminal Services Group Policies.<br><br>To configure a terminal server locally, use Terminal Services Configuration.<br><br>For more information, see Terminal Services Configuration Overview |
| Use Terminal Services for Remote Administration | Terminal Services in Remote Administration mode (must be installed using **Add or Remove Programs**) | **Remote Desktop** (installed by default).<br><br>For more information, see Remote Desktop for Administration. |
| Manage the terminal server and view all sessions, users, and processes | **Terminal Services Manager** in **Administrative Tools** | Terminal Services Manager.<br><br>For more information, see Terminal Services Manager. |
| Manage Terminal Server licenses | **Terminal Services Licensing** in **Administrative Tools** | Terminal Server Licensing.<br><br>For more information, see Terminal Server Licensing. |
| Manage Terminal Server user accounts | For a domain user account, use **Active Directory Users and Computers**. For a local user account, use **Computer Management**. | Terminal Services Group Policies.<br><br>For more information, see Managing Terminal Services users with Group Policy |

# Terminal Server Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### Best practices

- Install Terminal Server on a standalone server and not on a domain controller.

  Installing Terminal Server on a domain controller can affect the performance of the server because of the additional memory, network traffic, and processor time required to perform the tasks of a domain controller in a domain.

- Install Terminal Server on an NTFS file system partition.

  NTFS provides greater security for users in a multisession environment who access the same data structures.

- When shutting down a terminal server, use the **tsshutdn** command instead of the **Shut Down** option on the **Start** menu. This will shut down the server in a controlled manner. The **Shut Down** option on the **Start** menu does not notify users before ending user sessions and is not recommended. Ending a user's session without warning can result in loss of data at the client. For more information about using **tsshutdn**, see Tsshutdn.

- Back up your license server regularly.

  Backing up your license server regularly protects data from accidental loss due to hardware or storage failure. Create a duplicate copy of the data on your hard disk and then archive the data on another storage device such as a removable disk or tape. In the event that the original data on your hard disk is accidentally erased or overwritten, or becomes inaccessible because of a hard disk failure, you can easily restore the data from the archived copy.

#### Program installation

- Use **Add or Remove Programs** in Control Panel to install applications on the terminal server. Using this method, you can install programs for multisession use. For more information about installing applications on a terminal server, see Starting the installation.

- After you install Terminal Server, do not use **Add or Remove Programs** to switch Terminal Server on and off.

  Terminal Server installs programs for use in a multisession environment. Programs that were installed while Terminal Server was installed might not work correctly when Terminal Server is uninstalled. Instead, reinstall all programs for use without Terminal Server if you decide to remove the component.

- Check for application compatibility scripts before installing programs for use with Terminal Server.

  Many commonly used programs have been tested for compatibility. Some programs require minor changes to the installation. Scripts are available for these programs and must be run after the program installation is complete. Scripts are located in the systemroot in \Application Compatibility Scripts\Install.

#### Configuring Terminal Services

- **Use Terminal Services Group Policy to configure one or more terminal servers, or to manage Terminal Server user settings.** Terminal Server Group Policies can be applied on individual computers or on groups of computers belonging to a single organizational unit.

#### Managing users

- **Use Terminal Services-specific groups**

Create User Groups that are specifically for Terminal Services users. Maintaining users through groups is much easier and less time consuming than managing users individually.

Windows Server 2003 family server operating systems contain a default User group called Remote Desktop Users, which has been specifically created to manage Terminal Server users. This group is not populated by default. You must add users to the Remote Desktop Users group if you want them to be able to establish remote connections.

- **Use Terminal Services-specific profiles**

  Assign a separate profile for logging onto Terminal Services. Many of the common options that are stored in profiles, such as screen savers and animated menu affects, are not desirable when using Terminal Services. Assigning a specific profile allows users to get the most out of the system they are using without expending additional server resources. For information about assigning a Terminal Services specific profile, see Change a user's Terminal Services profile path.

- **Use mandatory profiles**

  Use a mandatory Terminal Services profile that is created to suit the needs of all of the different types of clients and that provides the best server performance. Be aware that 16-bit computers and Windows-based terminals might not support some screen resolutions.

- **Set time limits**

  Setting limits on the duration of client connections can improve server performance. You can set the limits on how long a session lasts, how long a disconnected session is allowed to remain active on the server, and the time allowed for a session to remain connected, yet idle. For information about setting session limits, see Configuring session limits.

- **Use the Starting program option**

  If you have users who need access to only one application on the terminal server, use the **Starting program** option to restrict users to that application. For more information, see Specify a program to start on session connection.

- **Create preconfigured connection files for users or groups of users**

  To make connecting to Terminal Services easier, you can supply users with preconfigured connection files. Collections of connection files can also be made either for different departments within your organization or for different job titles. preconfigured connection files are created using Remote Desktop Connection. For more information, see Managing Terminal Services connection files.

# Terminal Server How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install Terminal Server
- Disable Terminal Services connections
- Uninstall Terminal Server
- Deploy Applications
- Deploy Client Software
- Use Client Software

# Install Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Terminal Server**

1. Open Configure Your Server Wizard.
2. Click **Next** and follow the instructions on the **Preliminary Steps** screen.
3. Click **Next**, and on the **Server Role** screen, select **Terminal Server**.
4. Click **Next** and follow the instructions in the wizard.
5. Some programs might not work correctly when Terminal Server is installed. If so, you must reinstall these programs for multisession access using **Add or Remove Programs** after Terminal Server is installed. For more information, see Program Compatibility.

**Warning**

- Terminal Server Licensing is a required component that licenses clients on a terminal server. You must install Terminal Server Licensing or your terminal server will stop accepting connections from unlicensed clients 120 days from the date of the first client logon.

  You should not install Terminal Server Licensing on the same computer with Terminal Server.

  For more information on Terminal Server Licensing, see Terminal Server Licensing.

**Important**

- When you install Terminal Server using the Configure Your Server Wizard, as a security precaution the Internet Explorer Enhanced Security Configuration is set by default to prevent Administrators from downloading files or running ActiveX controls from the Web. If you want to change this setting to either permit all users to download files or run ActiveX controls or prevent all users from doing so, use **Add/Remove Windows Components** in **Add or Remove Programs**. This setting applies to servers running one of the Windows Server 2003 family operating systems. For more information, see Terminal server role: Configuring a terminal server and Internet Explorer Enhanced Security Configuration.

**Important**

- You must use the Remote Desktop Users group to control remote connection permissions for Terminal Server. For more information, see Managing Terminal Services users and Add users to the Remote Desktop Users group.

**Important**

- It is not necessary to install Terminal Server for remote administration of your server. Remote Desktop for Administration is installed by default. To use Remote Desktop for Administration, you must first enable remote connections. For more information on enabling remote connections, see Enable or disable Remote Desktop and Enabling users to connect remotely to the server.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For

more information, see Default local groups, Default groups, and Using Run as.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.

- You can also use **Add or Remove Programs** to install Terminal Server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Terminal Server Licensing
Activate a Terminal Server License Server
Install Client Access Licenses
Enable or disable Remote Desktop

# Disable Terminal Services connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable Terminal Services connections**

1. Click the **System** in Control Panel.

2. On the **Remote** tab, clear the **Allow users to connect remotely to your computer** check box, and then click **OK**.

   The terminal server no longer accepts new connections, but existing connections are maintained until closed by the users. This allows administrators to shut down a terminal server gradually.

## Notes

- You must be logged on as Administrator or a member of the Administrators group to disable Terminal Services connections.

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change logon
Switching Terminal Services on and off

# Uninstall Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To uninstall Terminal Server**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add or Remove Windows Components** to open the Windows Components wizard.
3. In the Windows Components wizard, deselect the **Terminal Server** check box, and then click **Next**.
4. Follow the instructions in the wizard.

### Important

- Use this procedure only if you want to completely uninstall Terminal Server. To disable Terminal Services connections temporarily (for example, to install an application on a terminal server), you can do so using the **Remote** tab in System Properties and clearing the **Allow users to connect remotely to your computer** checkbox.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Switching Terminal Services on and off
Disable Terminal Services connections
Disable a Terminal Services connection
Change logon

# Deploy Applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploy applications**

- Install a program using Add or Remove Programs
- Install a program using the change user command

# Install a program using Add or Remove Programs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a program using Add or Remove Programs**

1. Log on to the terminal server as an administrator and close all programs.
2. Open Add or Remove Programs in Control Panel.
3. In **Add or Remove Programs**, click **Add New Programs**.
4. Select the method to install the program and follow the instructions in the wizard.
5. The installation wizard will guide you through the rest of the installation process. For security purposes, install the program on an NTFS file system partition.
6. If applicable, edit and run the program installation script.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- If you are prompted to restart the computer to finish the installation, click **Finish** in **Add or Remove Programs** before accepting the prompt to restart the computer. It might be necessary to log on as the same user after the server has been rebooted in order to finish the application installation properly. Do not allow other users to log on before this step has been completed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Program Considerations
Deploying Applications

# Install a program using the change user command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a program by using the change user command**

1. Log on to the terminal server as an administrator and close all applications.
2. To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
3. Type **change user /install** and press ENTER.
4. Install the program on a local NTFS file system formatted drive as directed by the installation program.
5. At the command prompt, type **change user /execute** when installation is complete.
6. If applicable, edit and run the program installation script.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- If prompted by a program to install an additional program, you can leave your initial program running.
- You can type **change user /query** at the **Command Prompt** to see whether execute mode is enabled and .ini file mapping is on.
- You should use **Add or Remove Programs** to install programs whenever possible.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing Terminal Services from the command line

# Deploy Client Software

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploy client software**

- Share the client setup folder at the server
- Create a connection file
- Install the 32-bit Terminal Services Client (Remote Desktop Connection)

# Share the client setup folder at the server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To share the client setup folder**

1. On a computer that is running one of the Windows Server 2003 family operating systems, open Windows Explorer.
2. Browse to the systemroot\System32\Clients\Tsclient\win32 folder, where *systemroot* is the folder containing your Windows installation files (for example, C:\Windows\System32\Clients\Tsclient\win32).
3. Right-click the \win32 folder, click **Sharing and security**, check **Share this folder**, and then click **OK**.

### Note

- To share folders and drives, you must be logged on as a member of the Administrators, Server Operators, or Power Users group.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the 32-bit Terminal Services Client (Remote Desktop Connection)
Install Remote Desktop Connection (32-bit computers)

# Create a connection file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a connection file**

1. Open **Remote Desktop Connection**.
2. In **Computer**, type a computer name or IP address.
3. Click **Options**.
4. Specify the connection settings that you want for this connection, (such as screen size, automatic logon information, and performance options).
5. On the **General** tab, click **Save As**.
6. Enter a file name for the saved connection file, and then click **Save**.

**Note**

- To open Remote Desktop Connection, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing Terminal Services connection files
Change Connection Settings
Connect to a terminal server using Remote Desktop Connection

# Install the 32-bit Terminal Services Client (Remote Desktop Connection)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install the 32-bit Terminal Services Client (Remote Desktop Connection)**

1. At a command prompt, type *drive*:setup (where drive is the letter of the drive containing the Terminal Services Client disk or path to shared installation files), and press ENTER.

2. Follow the instructions in the Wizard.

## Notes

- Remote Desktop Connection is already installed on Windows XP and Windows Server 2003 family operating systems.

- To use Group Policy to deploy Remote Desktop Connection, assign the program on a per-computer basis to a terminal server. For more information, see article Q324750, "HOW TO: Assign Software to a Specific Group By Using a Group Policy in Windows Server 2003" in the Microsoft Knowledge Base.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Share the client setup folder at the server
Remove the Terminal Services Client
Terminal Services Client Overview

# Remove the Terminal Services Client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove the Terminal Services Client**

1. In the **Terminal Services Client group**, click **Uninstall**.
2. In the **Terminal Services Client Setup** dialog box, click **Remove All** to remove all Terminal Services Client files from your client computer.
3. Click **Yes** to confirm the removal.
4. In the **Terminal Services Client - Restart Windows** dialog box, click **Restart Windows** to complete.

### Note

- If you installed the client from a floppy disk, you must place the installation disk in the drive during this procedure. Placing the disk in the drive during uninstall reactivates the disk so that it can be used to install the client for another user.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the 32-bit Terminal Services Client (Remote Desktop Connection)
Remote Desktop Connection overview

# Use Client Software

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use client software**

- Connect to a terminal server using Remote Desktop Connection
- Use Terminal Services shortcut keys
- Check the version number of Remote Desktop Connection
- Disconnect without ending a session using Remote Desktop Connection
- Log off and end the session using Remote Desktop Connection

# Connect to a terminal server using Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a terminal server using Remote Desktop Connection**

1. To open Remote Desktop Connection, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

   The **Remote Desktop Connection** dialog box appears.

2. In **Computer**, type a computer name or IP address. In the **Computer** list, you can also choose **Browse for more...** to select a computer from your domain or network.

3. Click **Options**>> to configure the connection settings in the available tabs.

4. Click **Save As...** to save the current connection settings or click **Open** to choose a previously saved connection file.

5. Click **Connect**.

   The **Logon Information** dialog box appears.

6. Type your user name, password, and domain (if required), and then click **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a connection file
Terminal Services Client Overview

# Use Terminal Services shortcut keys

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Terminal Services shortcut keys**
If you choose not to apply Windows key combinations to your terminal server sessions, you can use these Terminal Services shortcut keys to perform many of the same functions.

| Shortcut key | Description |
|---|---|
| ALT+PAGE UP | Switches between programs from left to right. |
| ALT+PAGE DOWN | Switches between programs from right to left. |
| ALT+INSERT | Cycles through the programs in the order they were started. |
| ALT+HOME | Displays the **Start** menu. |
| CTRL+ALT+BREAK | Switches the client between a window and full screen. |
| CTRL+ALT+END | Displays the **Windows Security** dialog box. |
| ALT+DELETE | Displays the **Windows** menu. |
| CTRL+ALT+Minus (-) symbol on the numeric keypad | Places a shadow copy of the active window, within the client, on the terminal server clipboard (provides the same functionality as pressing PrintScrn on a local computer.) |
| CTRL+ALT+Plus (+) symbol on the numeric keypad | Places a shadow copy of the entire client window area on the terminal server clipboard (provides the same functionality as pressing ALT+PrintScrn on a local computer.) |

**Note**

- On an NEC98, these shortcut keys are different in two cases:
  - CTRL+ALT+BREAK is replaced by CTRL+GRPH+VF2.
  - CTRL+ALT+END is replaced by CTRL+GRPH+VF1.

**See Also**
**Concepts**
Terminal Services Client Overview
Configure Windows shortcut keys in a session

# Check the version number of Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check the version number of Remote Desktop Connection**

1. To open Remote Desktop Connection, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

   The **Remote Desktop Connection** dialog box appears.

2. Click the computer icon in the upper-left corner of the title bar, and then click **About**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the 32-bit Terminal Services Client (Remote Desktop Connection)
Install Remote Desktop Connection (32-bit computers)

# Disconnect without ending a session using Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect without ending a session**

1. In the remote desktop window, click **Start**, and then click **Shut Down**.

   The **Shut Down Window** dialog box appears.

2. Select **Disconnect** from the list at the top of this dialog, and then click **OK**.

### Note

- Remote Desktop Connection reconnects to this session the next time you connect to this terminal server (if the connection is configured for reconnection of disconnected sessions).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Log off and end the session using Remote Desktop Connection
Log off or disconnect

# Log off and end the session using Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To log off and end the session**

1. In the remote desktop window, click **Start**, and then click **Shut Down**.

   The **Shut Down Windows** dialog box appears.

2. Click **Log Off**, and then click **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disconnect without ending a session using Remote Desktop Connection
Log off or disconnect

# Terminal Server Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Terminal Server:

- Getting Started with Terminal Server
- Deploying Terminal Server
- Deploying Applications
- Deploying Client Software
- Terminal Services Client Overview

# Getting Started with Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Getting started with Terminal Server**

This topic covers information about getting started with Terminal Server. The following topics are included:

- Terminal Services overview
- Using Terminal Server to host applications centrally
- Terminal Services administrative tools

# Terminal Services overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services overview**

Terminal Services provides remote access to a Microsoft® Windows® desktop through "thin client" software, allowing the client computer to serve as a terminal emulator. Terminal Services transmits only the user interface of the program to the client. The client then returns keyboard and mouse clicks to be processed by the server. Each user logs on and sees only their individual session, which is managed transparently by the server operating system and is independent of any other client session. Client software can run on a number of client hardware devices, including computers and Windows-based terminals. Other devices, such as Macintosh computers or UNIX-based workstations, can use additional third-party software to connect to a server running Terminal Server.

Terminal Services is the underlying technology for several features and components of Microsoft Windows Server™ 2003 family operating systems. Among these are Terminal Server and Remote Desktop for Administration.

Terminal Server provides an effective and reliable way to distribute Windows-based programs with a network server. With Terminal Server, a single point of installation allows multiple users to access the desktop on a server running one of the Windows Server 2003 family operating systems. Users can run programs, save files, and use network resources as if they were sitting at that computer. For more information on using Terminal Server for multisession application hosting, see Using Terminal Server to host applications centrally.

Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode) provides remote access to the desktop of any computer running one of the Windows Server 2003 family operating systems, allowing you to administer your server--even a Microsoft Windows 2000 server--from virtually any computer on your network. For more information on using Remote Desktop for Administration for remote server administration, see Remote Desktop for Administration.

**Terminal Services benefits:**

- **Brings** Windows Server 2003 **family operating systems to desktops faster.** Terminal Services helps bridge the gap while older desktops are migrated to Microsoft Microsoft® Windows® XP Professional, providing a virtual desktop experience of any Windows Server 2003 family operating system to computers that are running earlier versions of Windows. Terminal Services clients are available for many different desktop platforms including Microsoft MS-DOS®, Windows-based terminals, Macintosh, and UNIX. Additionally, a Web-based version of the Terminal Services client (Remote Desktop Web Connection) provides Terminal Services connectivity to computers with Web access and an Internet Explorer browser. (Connectivity for MS-DOS, Macintosh, and UNIX-based computers requires additional software.)

- **Takes full advantage of existing hardware.** Terminal Services extends the model of distributed computing by allowing computers to operate as both thin clients and full-featured personal computers simultaneously. Computers can continue to be used as they have been within existing networks while also functioning as thin clients capable of emulating the Windows XP Professional desktop.

**Terminal Server benefits:**

- **Centralized deployment of programs.** With Terminal Server, all program execution, data processing, and data storage occur on the server, centralizing the deployment of programs. Terminal Server ensures that all clients can access the same version of a program. Software is installed only once on the server, rather than on every desktop throughout the organization, reducing the costs associated with updating individual computers. For more information on centralized deployment of programs, see Using Terminal Server to host applications centrally.

**Remote Desktop for Administration benefits:**

- **Remote administration.** Terminal Services Remote Desktop for Administration provides remote administration for Windows Server 2003 family operating systems, giving system administrators a method of remotely managing their server from any client over a LAN, WAN or dial-up connection. Up to two remote sessions, plus the console session, can be accessed simultaneously. Terminal Server licensing is not required to use this feature.

# Using Terminal Server to host applications centrally

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Terminal Server to host applications centrally**

Terminal Server provides an effective and reliable way to distribute Windows-based programs by making them available on a network server. With Terminal Server, a single point of installation allows multiple users to access the desktop of a server running one of the Windows Server 2003 family operating systems. Users can run programs, save files, and use network resources as though they were sitting at that computer. Terminal Server delivers the Windows desktop and Windows-based applications to computers that might not typically be able to run Windows operating systems.

**Note**

- Terminal Server is not included in Windows Server 2003, Web Edition.

**Ease of use**

| Feature | Description | For more information, see: |
|---|---|---|
| **Improved client: Remote Desktop Connection** | Renamed **Remote Desktop Connection**, the Terminal Services client features a more streamlined user interface that allows you to connect to terminal servers, change connection settings, save connection information, and connect to a remote computer. The old Terminal Services Client Connection Manager tool is no longer needed, although you can still use any connections that were created with that tool. With Remote Desktop Connection, users can create connection files, which are predefined connections to servers for either a single program or full desktop access, and save them in **My Documents**. You can ensure that the same version of a single program is being used across the computing environment by creating a connection file and distribute it along with the Remote Desktop Connection client software. | Remote Desktop Connection overview |
| **Automated local printer support** | Terminal Services can both add and automatically reconnect printers attached to Terminal Services clients. | Providing client access to local printers |
| **Drive and file system redirection** | Users can see their local drives while logged on to a terminal server. In Microsoft Windows Explorer, users' local drives appear in the session folder tree as <driveletter> on <client>. | Providing client access to the local file system |
| **Audio redirection** | Users of a terminal server can direct audio output from the terminal server to the audio device on their local computer. | Providing client access to server audio |
| **Clipboard redirection** | Users can cut and paste between the programs running on the local computer and the terminal server. | Shared clipboard |
| **Performance enhancements** | Enhancements to caching, including persistent caching, packet utilization, and frame size, provide significant performance improvements over earlier versions of Terminal Services. | Bitmap caching |
| **Roaming disconnect support** | This feature enables users to disconnect from a session without logging off. A session can remain active while disconnected, enabling the user to reconnect to the existing session from another computer or at a later time. Logon is required for reconnection, keeping each session secure at all times. | Log off or disconnect |
| **Multiple logon support** | Users can log on to multiple sessions simultaneously from one or more clients, to multiple terminal servers, or to a single server multiple times. As a result, the user can do several tasks concurrently or run multiple unique desktop sessions. | Terminal Services User Properties |

**Manageability**

| Feature | Description | For more information, see: |
|---|---|---|
| **Terminal Services Group Policies** | With the addition of Terminal Services Group Policies, administrators can use one tool (Group Policy) to manage Terminal Services-related settings throughout an organization. Terminal Services nodes have been added to the Computer and User section of the Administrative Group Policy template (System.adm). | Configuring Terminal Services with Group Policy |
| **Terminal Services WMI provider** | The Windows Management Instrumentation (WMI) provider for Terminal Services allows you to use scripts to manage, configure, and troubleshoot terminal servers remotely, bypassing the Terminal Services tools, command-line options, and other management methods. | Configuring Terminal Services with WMI |
| **Session remote control** | Support staff can either view or control another Terminal Services session. Keyboard input, mouse movements, and display graphics are shared between two Terminal Services sessions, giving support staff the ability to diagnose and resolve configuration problems, as well as provide user training from a remote location. This feature is particularly useful for organizations with branch offices. | Using remote control |
| **Load balancing** | Terminal Services Connection Management (TSCM), combined with load balancing technology, allows Terminal Services clients to connect to the least busy member of a group of terminal servers. Session Directory services ensure that disconnected sessions are reconnected to the same server where the original connection was established. | Load balancing and terminal servers |
| **Terminal Services Web Client: Remote Desktop Web Connection** | Remote Desktop Web Connection is an ActiveX control that provides virtually the same functionality as the full Remote Desktop Connection client, but is designed to deliver this functionality over the Web. When embedded in a Web page, Remote Desktop Web Connection can host a Terminal Services client session even if the full client is not installed on the user's computer. | About Remote Desktop Web Connection |
| **Microsoft Windows Installer (MSI) package for client deployment** | Remote Desktop Connection can be installed with Windows Installer. Windows Installer provides a quick and efficient means for deploying client software to target computers, using either a network-based share point or Microsoft IntelliMirror. | About the Terminal Services Client MSI Setup Package |
| **Windows-based terminals** | Windows-based terminals are available from a variety of manufacturers. They use a custom implementation of the Microsoft Windows CE operating system and the Remote Desktop Protocol (RDP). | |
| **Terminal Server Licensing** | Terminal Server Licensing helps system administrators and purchasing offices track clients and their associated licenses. | Terminal Server Licensing |
| **Distributed File System (DFS) support** | Support for Distributed File System (DFS) enables users to connect to a DFS share, and administrators to host a DFS share from a terminal server. | Distributed File System overview |
| **Terminal Services Manager** | Administrators can use Terminal Services Manager to query and manage Terminal Services sessions, users, and processes. | Terminal Services Manager |
| **Terminal Services Configuration** | Terminal Services Configuration allows you to create, modify, configure, and delete Remote Desktop Protocol (RDP) connections. You can also use Terminal Services Configuration to view and set terminal server settings. | Terminal Services Configuration |
| **Integration with Local Users and Groups and Active Directory Users and Computers** | Administrators can create accounts for Terminal Services users the same way they create accounts for users on any Windows Server 2003 family operating system. Extra fields exist for specifying information specific to Terminal Services, such as the Terminal Services Profile Path and Home Directory. | Terminal Services users |

| | | |
|---|---|---|
| **Integration with System Monitor** | Integration with System Monitor enables administrators to monitor Terminal Services system performance, including tracking processor use, memory allocation, and paged memory usage and swapping per-user session. | Performance monitoring |
| **Messaging support** | Administrators can alert users to important information, such as system shutdowns, upgrades, or new programs. | Send a message to a user |
| **Remote administration with Remote Desktop for Administration** | Any member of the Administrators group with access to the Terminal Services administrative utilities can remotely manage all aspects of a server. Remote Desktop for Administration is installed automatically on all Windows Server 2003 family operating systems. | Remote Desktop for Administration |
| **Configurable session time-out** | Administrators can reduce server resource usage by configuring session time-outs. Administrators can specify the length of an active session and how long a session can remain idle on the server. | Configuring session limits |

**Security**

| Feature | Description | For more information, see: |
|---|---|---|
| **Security modes** | Terminal Server can run in one of two security modes, depending on the applications you plan to run:<br><br>**Full Security**<br>Select this option to take advantage of new security features in Windows Server 2003 family operating systems and provide the most secure environment for your terminal server.<br><br>**Relaxed Security**<br>This option allows access to the system registry, enabling most legacy applications to run as they did under Windows 2000, Terminal Server Edition. | Choosing the security mode for a terminal server |

| | | |
|---|---|---|
| **Encryption** | Terminal Services has four levels of encryption.<br><br>• By default, Terminal Services connections between any of the Windows Server 2003 family operating systems and client computers running the Remote Desktop Connection client are encrypted with **High** (128-bit) encryption.<br><br>• The **FIPS Compliant** encryption level encrypts and decrypts data sent from client to server and from server to client with the Federal Information Processing Standard (FIPS) encryption algorithms using the Microsoft cryptographic modules. If FIPS compliance has already been enabled by the **System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing** Group Policy, administrators cannot change the encryption level for Terminal Services connections either by changing the Terminal Services **Set client connection encryption level** Group Policy setting or by using Terminal Services Configuration.<br><br>• For legacy clients that support only low levels of encryption, administrators can set the encryption level to **Client Compatible.**<br><br>• The **Low** level encrypts data sent from the client to the server using 56-bit encryption.<br><br>**Important**<br><br>• Data sent from the server to the client is not encrypted.<br><br>In addition, the Terminal Services logon process includes change password, unlock desktop, and unlock screen saver features. The logon process is encrypted, ensuring secure transfer of user name and password. | Configuring authentication and encryption System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing |
| **Limit logon attempts and connection time** | Administrators can limit the number of user logon attempts to prevent unauthorized access to a server. Additionally, the connection time of an individual user or a group of users can be limited. | Terminal Services Configuration |

# Terminal Services administrative tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services administrative tools**
In addition to the standard administrative tools, the following Terminal Services tools, tool extensions, and commands for administering Terminal Services connections, computers, and users are included in all of the Windows Server 2003 family operating systems.

- Terminal Services Group Policies.
- Terminal Services Manager.
- Terminal Services Configuration.
- Remote Desktops MMC snap-in.
- Active Directory Users and Computers and Local Users and Groups extensions.
- Terminal Server Licensing.
- System Monitor counters.
- Task Manager additional fields.
- Multiuser support in Add or Remove Programs.
- Commands.

**Terminal Services Group Policies**
Use Terminal Services Group Policies to configure individual terminal servers or groups of terminal servers, or to assign policies for Terminal Server users or user groups. For more information, see Configuring Terminal Services with Group Policy.

**Terminal Services Manager**
Use Terminal Services Manager to manage and monitor users, sessions, and processes on any server running Terminal Services on the network. Use it to:

- Display information about servers, sessions, users, and processes.
- Connect to and disconnect from sessions.
- Monitor sessions.
- Reset sessions.
- Send messages to users.
- Log off users.
- Terminate processes.

For more information about Terminal Services Manager, see Terminal Services Manager.

**Terminal Services Configuration**
During the installation of a Windows Server 2003 family operating system, a connection is configured for the Remote Desktop

Protocol (RDP). This connection provides the link clients use to log on to a session on the server for either Remote Desktop for Administration or Terminal Server connections. After installation, you can use Terminal Services Configuration (Tscc.msc) to change the properties of this connection on the local computer or to add new connections.

**Note**

- Many of these settings can also be set with Terminal Services Group Policy or Windows Management Instrumentation (WMI), in which case Terminal Services Configuration settings are overridden.

Use Terminal Services Configuration to:

- Name a connection.
- Specify a connection type.
- Specify a connection transport and transport properties.
- Set the maximum number of sessions allowed.
- Enable or disable logons through the connection.
- Set connection time-outs.
- Set the level of encryption.
- Set whether to disconnect broken connections.
- Enable or disable session remote control.
- Enable or disable automatic logons.
- Specify a program to run when a user logs on.
- Override user profile settings for wallpaper.
- Set permissions on the connection.
- Set client device mapping and connection parameters.

For more information about Terminal Services Configuration, see Terminal Services Configuration.

**Remote Desktops MMC snap-in**
This Microsoft Management Console (MMC) Remote Desktops snap-in allows you to create connections to multiple terminal servers, configure them to run specific programs upon connecting, and switch easily between Terminal Services connections by clicking the connection in the console tree. A special MMC version of the Terminal Services client displays the desktop of the selected computer in the MMC results pane.

For more information about the Remote Desktops snap-in, see Remote Desktops snap-in.

**Active Directory Users and Computers and Local Users and Groups extensions**
Terminal Services extends Active Directory Users and Computers and Local Users and Groups to allow you to control Terminal Services features for each user. You can also use Terminal Services Group Policies to control these features, although settings you have specified might be overridden.

Use the Terminal Services extension to:

- Set the path to the Terminal Services user profile for each user.
- Enable or disable logons.

- Set time limits on session.
- Set whether to disconnect or reset a broken connection.
- Enable or disable remote control.
- Specify a program to run when a user logs on.
- Make client drives and printers connect at logon.

**Terminal Server Licensing**
Use Terminal Server Licensing to register and track licenses for Terminal Server clients. For more information, see Terminal Server Licensing.

### Important

- You must install Terminal Server Licensing or your terminal server will stop accepting connections from unlicensed clients 120 days from the date of the first client logon.

**System Monitor counters**
Terminal Services extends System Monitor by adding User and Session objects and their counters. You can use these objects and counters to monitor the resources that a user or session consumes. Terminal Services also adds counters to the Process and System objects.

For more information on System Monitor, see System Monitor overview.

**Task Manager additional fields**
Terminal Services provides two additional fields, **ID** (session ID) and **User**, to Task Manager for monitoring and ending processes for all sessions.

**Multiuser support in Add or Remove Programs**
Add or Remove programs ensures that applications are installed for use in a multisession environment.

**Commands**
The following table shows the commands you can use to manage Terminal Services at the command prompt.

| Command | Function |
|---|---|
| Change logon | Temporarily disable logons to a terminal server. |
| Change port | Changes COM port mappings for MS-DOS program compatibility. |
| Change user | Changes the .ini file mapping for the current user. |
| Cprofile | Removes user-specific file associations from a user's profile. |
| Flattemp | Enables or disables flat temporary directories. |
| Logoff | Ends a client session. |
| Msg | Sends a message to one or more clients. |
| Mstsc | Runs Remote Desktop Connection. |
| Query process | Displays information about processes. |
| Query session | Displays information about Terminal Services sessions. |
| Query termserver | Lists Terminal Services servers on the network. |

| | |
|---|---|
| Query user | Displays information about users logged on to the system. |
| Register | Registers a program so that it has special execution characteristics. |
| Reset session | Allows you to reset (delete) a session from the terminal server. |
| Shadow | Allows you to monitor or remotely control an active session of another user. |
| Tscon | Connects to another existing Terminal Services session. |
| Tsdiscon | Disconnects a client from a Terminal Services session. |
| Tskill | Terminates a process. |
| Tsprof | Copies user configuration and changes profile path. |
| Tsshutdn | Shuts down a Terminal Services server. |

# Deploying Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying Terminal Server**

This section provides information about deploying Terminal Server:

- Installing Terminal Server
- System Planning

# Installing Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Terminal Server**

This section covers information about installing Terminal Server:

- Before Installing Terminal Server
- Choosing the security mode for a terminal server
- Switching Terminal Services on and off
- Managing Terminal Services from the command line

When you install one of the Windows Server 2003 family operating systems, Remote Desktop for Administration is installed automatically. With Remote Desktop for Administration, you can create up to two concurrent remote connections to your server for remote administration and other tasks. To use Remote Desktop for Administration, you must first enable remote connections. For more information, see Enable or disable Remote Desktop.

To use Terminal Services for multiuser application hosting, you must install the Terminal Server component using the **Configure Your Server Wizard**. (Terminal Server is not included in Windows Server 2003, Web Edition.)

For information on using Remote Desktop for Administration, see Remote Desktop for Administration.

# Before Installing Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Before installing Terminal Server**

Review the following information before you install Terminal Server:

- Read System Planning for system sizing and performance considerations.
- Read "Hosting Applications with Terminal Server" at the Microsoft Windows Resource Kit Web site.
- Plan to install Terminal Server on a standalone server, not on a domain controller. When installed on a domain controller, Terminal Server performance is affected because of the additional memory, network traffic, and processor time required to perform the tasks of a domain controller in a domain.
- Install Terminal Server on a server installed on an NTFS file system partition. NTFS provides greater security for users in a multisession environment.
- Understand the effect of the Internet Explorer Enhanced Security Configuration setting on users of a terminal server. When you install Terminal Server using the Configure Your Server Wizard, this setting prohibits Administrators from downloading files or ActiveX controls from the Web by default as a security precaution. However, all non-Administrator users can download files or ActiveX controls. After you install Terminal Server, you can change this setting to allow both Administrators and non-Administrator users to perform these tasks, or to prevent all users from performing these tasks. Use **Add/Remove Windows Components** in **Add or Remove Programs** to change this setting. This setting applies to servers running one of the Windows Server 2003 family operating systems. For more information, see Terminal server role: Configuring a terminal server and Internet Explorer Enhanced Security Configuration.
- Install the programs you want to use with Terminal Server *after* Terminal Server is installed. This allows Terminal Server to register the programs for multisession access, making the programs easier to manage. For more information about installing programs, see Installing Programs.

For instructions on installing Terminal Server, see Install Terminal Server

# Installing Terminal Services in remote administration mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Terminal Services in remote administration mode**
It is no longer necessary to install Terminal Services in Remote Administration mode (now called Remote Desktop for Administration). When you install one of the Windows Server 2003 family operating systems, Remote Desktop for Administration is installed automatically. To use Remote Desktop for Administration, you must first enable remote connections. For more information, see Enable or disable Remote Desktop.

For information on using Remote Desktop for Administration, see:

- Remote Desktop for Administration

# Choosing the security mode for a terminal server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing the security mode for a terminal server**
In Windows 2000 and Windows Server 2003 family operating systems, you can choose one of two security modes:

- Relaxed security: Windows NT 4.0/Terminal Server Edition permissions compatibility mode
- Full security: Windows 2000/Windows Server 2003 family operating systems permissions mode

In both modes, the TERMINAL SERVER USER security descriptor is applied to many keys in the registry and directories in the file system. The difference between the two modes is based on the way the security descriptor is applied.

- In Relaxed Security mode, the security descriptor is added to each member of the Users group. On Windows NT 4.0, when a user's permissions are evaluated to determine access to a file or registry entry, the presence of the security descriptor authorizes access.

  This application of the security descriptor allows members of the Users group to run programs that otherwise might not work at all in the more rigorous Full Security permissions mode used on Windows 2000 and Windows Server 2003 family operating systems. However, when you choose Relaxed Security mode, any user on the system can change files and registry settings in many places throughout the system, although others users' data files may not be visible. A malicious user could exploit this situation by replacing a known and trusted program with a program of the same name but some bad intent.

- By contrast, Full Security mode does not apply the descriptor to each user. Instead, applications must be written to run in the context of an ordinary user. When in doubt, you should choose Full Security mode and test your applications in that mode.

  **Important**

  - Use Relaxed Security mode only if your test results indicate that it is necessary.

For information on changing the security mode, see Switch default permissions for application compatibility.

# Switching Terminal Services on and off

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Switching Terminal Services on and off**
To disable Terminal Services connections temporarily (for example, to install an application on a terminal server), go to the **Remote** tab in System Properties and clear the **Allow users to connect remotely to your computer** checkbox. (This checkbox is cleared by default when you install any of the Windows Server 2003 family operating systems.) When you disable remote connections, Terminal Services remains active on the computer, but no new remote connections are accepted.

If Terminal Server is installed on your computer, you can also uninstall it by using **Add or Remove programs**. When you do so, the computer is returned to its default state, Remote Desktop for Administration.

**Caution**

- Many older applications were not designed for multiple users and need to be modified before they will run properly on a terminal server. If you decide to uninstall Terminal Server using **Add or Remove programs**, the applications you installed and configured to run on a terminal server might not run the next time you install Terminal Server. You must configure or install the application again. For this reason, you should not use **Add or Remove programs** to temporarily disable Terminal Server.

**Note**

- For more information on disabling Terminal Services connections using the **Remote** tab in System Properties, see Disable Terminal Services connections

# Managing Terminal Services from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Terminal Services from the command line**

You can manage users, sessions, processes, and terminal servers by using command-line utilities. The following table lists the Terminal Services administrative utilities that you can invoke from the command line.

| Command | Function |
|---|---|
| Change logon | Disables, temporarily, logons to a terminal server. |
| Change port | Changes COM port mappings for MS-DOS program compatibility. |
| Change user | Changes the .ini file mapping for the current user. |
| Cprofile | Removes user-specific file associations from a user profile. |
| Flattemp | Enables or disables flat temporary directories. |
| Logoff | Logs off a user from a session and deletes the session from the server. |
| Msg | Sends a message to a user or group of users. |
| Mstsc | Displays the Remote Desktop Connection to establish a connection with a Terminal Server. |
| Query process | Displays information about processes running on a terminal server. |
| Query session | Displays information about sessions on a terminal server. |
| Query termserver | Displays a list of all terminal servers on the network. |
| Query user | Displays information about user sessions on a terminal server. |
| Register | Registers applications to execute in a global context on the system. |
| Reset session | Resets a session to known initial values. |
| Shadow | Monitors another user's session. |
| Tscon | Connects to another existing terminal server session. |
| Tsdiscon | Disconnects a client from a terminal server session. |
| Tskill | Ends a process. |
| Tsprof | Copies user configuration and changes profile path. |
| Tsshutdn | Shuts down a terminal server. |

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

# System Planning

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System planning**
This section presents information on system sizing and system components that can help you decide on a hardware configuration for a terminal server that gives your users optimal performance.

- System sizing
- Peripherals that affect performance
- Application performance
- Performance monitoring
- Load balancing and terminal servers

**Note**

- For additional information about deploying Terminal Server, see "Hosting Applications with Terminal Server" at the Microsoft Windows Resource Kit Web site.

# System sizing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System sizing**

Most computer users can be placed into one of three categories:

- **Light users.** A light user typically runs one program at a time--for example, a word processing program or a line-of-business program.
- **Structured task users.** A structured task user simultaneously runs two or more programs but does not frequently switch between programs. This user runs programs whose data-processing demands on the system are not heavy--for example, productivity programs such as word processing and an Internet browser.
- **Power users.** A power user is a more sophisticated user who runs three or more programs, and often has several active at the same time. A power user runs programs whose data processing demands on the system are heavy, such as detailed queries on large databases.

When determining how powerful your servers need to be, consider the types of users who will be using them. Light users place the lightest load on computational resources while power users consume the most computational resources. Also consider the types of programs your server runs. Keep in mind that a terminal server shares executable resources between individual users, just as any of the Windows Server 2003 family operating systems shares executable resources (.dll, .exe, and so on) between individual programs. As a result, the memory requirements for additional users running the same program are typically less than the requirements for the first user to load the application.

For adequate performance, a terminal server requires a Pentium or higher processor. A terminal server requires a minimum of 128 MB RAM, plus additional RAM for each user to support running each user's programs on the server. An additional 10 MB RAM is recommended for each light user and up to 21 MB RAM for each power user.

In general, processor and memory requirements scale linearly: You can support double the number of users on a multiprocessor-capable Pentium system by doubling the number of processors and doubling the amount of memory. Purchasing a system that supports multiple processors, even if you initially purchase only one processor, allows you to add capacity easily as your requirements grow.

Performance can vary depending on the programs you run on a terminal server. To judge the processing power and RAM required for each user, test your system by running the programs you use.

The bus architecture in single and multiprocessor systems also affects system performance. The ISA (AT bus) architecture is low bandwidth and is not recommended for a terminal server. Use a higher-performance bus, such as EISA, MCA, PCI, PCI-X, or InfiniBand for best performance. All of these buses support the sustained high data transfer rates typically required when running Terminal Server.

Peripherals can affect the performance of a terminal server. For more information, see Peripherals that affect performance.

# Peripherals that affect performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Peripherals that affect performance**

Hard-disk throughput also affects the performance of a terminal server. SCSI disk drives and adapters, especially devices compatible with Fast SCSI and SCSI-2, have significantly better throughput than ST-506, IDE, or ESDI disk drives and adapters.

For highest disk performance, consider using a SCSI RAID controller. RAID (Redundant Array of Independent Disks) controllers place data on multiple disks to increase disk performance and improve data reliability.

Although the Remote Desktop Protocol (RDP) used with a terminal server causes minimal network load, a high-performance network adapter is recommended. This is particularly important if many users require either access to data stored on network servers or run client/server programs. Consider using two network adapters, dedicating one to RDP traffic only.

If a multiport asynchronous communication adapter is installed for supporting dial-up users, be sure to use an intelligent (microprocessor-based) adapter to both reduce interrupt overhead and increase throughput.

Applications can also affect performance. For more information, see Application performance.

# Application performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application performance**

Windows Server 2003 family operating systems are 32-bit environments, whereas Microsoft Windows 3.11 is a 16-bit environment. The Windows Server 2003 family operating systems run 16-bit programs through a process called Windows on Windows (WOW), translating 16-bit programs in enhanced mode. This process causes 16-bit programs to consume additional system resources. Running 16-bit programs on a terminal server can reduce the number of users a processor can support by 40 percent and increase the memory required for each user by 50 percent. For this reason, use 32-bit versions of programs whenever possible.

Here are additional considerations for installing programs on a terminal server:

- Some video active applications can decrease performance. Turn off the video effects of an application for better performance.
- MS-DOS applications are not recommended because they can consume all of the available CPU.
- Application features that are always on and run in the background, such as automatic spelling and grammar checking in Microsoft Word, require more system resources.

For more information about 16-bit and 32-bit programs, see Comparing 16-bit and 32-bit programs. For more information about installing programs on a terminal server, see Installing Programs.

# Performance monitoring

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Performance monitoring**
Use the performance monitoring tools supplied with the Windows Server 2003 family operating systems to monitor system performance and the effects of configuration changes on system throughput. Some of the most important measurements for performance monitoring include:

- Processor utilization
- Hard-disk I/O rates
- Memory utilization
- Pagefile activity

A good way to estimate how many users a server can support is to measure system performance with two to five users on the system, and then scale the results.

Terminal Services adds the Terminal Services User and Terminal Services Session objects and their counters to the set of objects you can observe using System Monitor.

For more information about performance monitoring, see "WSRK System Performance and Troubleshooting Guide Part One: Operating System Performance" at the Microsoft Windows Resource Kit Web site.

# Load balancing and terminal servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Load balancing and terminal servers**
Load Balancing pools the processing resources of several servers using the TCP/IP networking protocol. You can use this service with a cluster of terminal servers to scale the performance of a single terminal server by distributing sessions across multiple servers. Session Directory (included in Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition) keeps track of disconnected sessions on the cluster, and ensures that users are reconnected to those sessions.

**Terminal Services Connection Management**
Terminal Services Connection Management works with your load balancing service to ensure that users are transparently reconnected to the original server hosting their disconnected Terminal Server sessions. The components of Terminal Services Connection Management are:

- A network load-balancing solution (Network Load Balancing, Domain Name System (DNS) round-robin, or a third-party solution).
- Two or more terminal servers, logically grouped into a Terminal Server cluster.
- A session directory server. This computer can be any running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, that is visible on the network and running the Terminal Services Session Directory service. As a Best Practice, make the Session Directory server a highly available network server that is not itself running Terminal Server.

The Terminal Server Connection Management process is as follows:

1. When the user logs on to the terminal server cluster, the terminal server receiving the initial client logon request sends a query to the session directory server.
2. The session directory server checks the username against its database and sends the result to the requesting server.
    - If the user has no disconnected sessions, logon continues at the server hosting the initial connection.
    - If the user has a disconnected session on another server, the client session is passed to that server and logon continues.
3. When the user logs on to the disconnected session, the session directory is updated.

**Terminal Services Session Directory**
The Terminal Services Session Directory service is a database that keeps track of sessions on terminal servers in a cluster and provides the information used at connection time to connect users to existing sessions.

When the Session Directory service starts up, it creates the "Session Directory Computers" local group. By default, this group is not populated. You must choose the individual computers or groups you want to participate in the Session Directory service, and then manually Changing group memberships to the Session Directory Computers group.

For more information about configuring Session Directory settings, see Configure Session Directory settings and "Hosting Applications with Terminal Server" at the Microsoft Windows Resource Kit Web site. For more information about adding members to a security group, see Changing group memberships.

**To start Session Directory service**

1. Open Services.

2. Double-click **Terminal Services Session Directory Properties**.

3. To start the service, click **Start**. To ensure that the service runs each time the server boots, in **Startup Type**, click **Automatic**.

### Note

- To open the Services snap-in manually, click **Start**, point to **Administrative Tools**, and then click **Services**.

#### To create a Terminal Server Cluster

1. Open the Group Policy object for the organizational unit that contains the terminal servers you want to participate in the cluster.

2. In Administrative Templates, in the Windows Components/Terminal Services/Session Directory folder, enable the Session Directory Active setting.

3. In the Session Directory Server setting, click Enable, type the name of the server running the Session Directory service, and then click **OK**.

4. In the Session Directory Cluster name setting, click Enable, specify a name for this Terminal Server cluster, and then click **OK**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Deploying Applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying applications**

To deploy an application for use by Terminal Server users, you install it on the terminal server.

For information about program considerations, see Program Considerations

For information about installing programs for use on a terminal server, see Installing Programs

For information about compatibility of programs with Terminal Server, see Program Compatibility

# Program Considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Program considerations**
Consider the following before selecting and installing programs on Terminal Server.

- Assigned and published programs
- Comparing 16-bit and 32-bit programs
- System resource considerations
- Programs starting programs

**Note**

- For additional information about deploying Terminal Server, see "Hosting Applications with Terminal Server" at the Microsoft Windows Resource Kit Web site.

# Assigned and published programs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Assigned and published programs**

Windows Server 2003 family operating systems provide Windows Installer for deploying programs within your organization. You can assign a program to either computers or users, or you can publish programs to users. For more information on using Windows Installer to install programs, see Managing options for computers through Group Policy.

Because Terminal Server is designed to provide multisession access to programs, additional consideration is needed before deploying applications in environments that contain terminal servers.

Applications are installed on the terminal server, and are available to any user with remote access to the terminal server. Terminal Server cannot accept published programs because they are published on a per-user basis. Additionally, assigned programs must be assigned on a per-computer basis, rather than on a per-user basis.

When installing assigned programs, members of the Administrators group on the terminal server can install the program from either the console session or from a remote session.

**Caution**

- Program installation often requires restarting the computer. Always ensure that no users are logged on to the terminal server before starting any program installation.
- If the program is installed from a remote session, the session is disconnected or reset if the program installation requires rebooting the server. When the server reboots, your session is reset and you need to again log on to the terminal server to ensure that the program is correctly installed.
- Whenever possible, install programs from the console session.

**Deploying programs on terminal servers by using Windows Installer**

Microsoft Windows Installer installs programs at different levels, depending on how the .msi package was created. Many .msi packages require minimal installation where only a small number of program components are installed. When a user first invokes the program, there are enough components installed for Windows Installer to know where to find the remainder of the components and install the program. For Terminal Server, members of the Users security group cannot invoke a Windows Installer installation. Therefore, it is important to install these programs locally. To do this, a transform file (.mst) can be used. Transform files appear as modifications to .msi packages and tell Windows Installer which bits to install locally. Since transform files might not be available for all programs, you can either ensure that all necessary components are fully installed during installation or create a transform file. For more information on transform files, see Windows Installer technologies.

Legacy programs can also be repackaged as Windows Installer programs, using tools that capture the changes made to the system when the program is installed. These changes are then used to create an .msi package for the legacy program. This method allows you to continue using programs developed specifically for your organization.

# Comparing 16-bit and 32-bit programs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Comparing 16-bit and 32-bit programs**

Terminal Server supports 16-bit and 32-bit Windows-based programs. Running 16-bit programs can reduce the number of users a processor supports by 40 percent and increase the memory required for each user by 50 percent. In addition, some 16-bit programs must be able to write to the directory where the program .ini file is stored. Most 32-bit programs use the registry to read and write program settings and need only to write to specific registry values. For both performance and security reasons, 32-bit programs are recommended whenever possible.

For more information about application performance, see Application performance.

# System resource considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System resource considerations**

Because Terminal Server is designed for multisession use, you should limit the number of installed programs that require large amounts of system resources. Consider restricting user or group access to certain program types, disabling unnecessary features that require the most resources, or installing these programs on separate servers.

Check system requirements for each program carefully and remember that RAM and CPU requirements increase with the number of sessions running. While there is not a true one-to-one correlation between these, this gives you a basis for projecting program performance.

For more information, see System Planning.

# Programs starting programs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Programs starting programs**

Some programs enable users to start other programs. For example, Microsoft Access has a toolbar that can be used to start other Microsoft Office programs. In this situation, you should disable toolbar access to other programs from the program that is installed on the terminal server.

Many Microsoft products include the Microsoft System Information (Msinfo.exe) program. There is the potential that users can breach security or gather system information, such as server hardware configuration, with this program. Set file permissions on the Microsoft System Information program so that only administrators can use it to gather system configuration information, aid in problem resolution, and start programs by the appropriate users.

For more information about setting file permissions, see Set, view, change, or remove permissions on files and folders.

# Installing Programs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing programs**
This section covers:

- Getting ready to install programs
- Starting the installation
- Testing the installation

**Note**

- For additional information about deploying Terminal Server, see "Hosting Applications with Terminal Server" at the Microsoft Windows Resource Kit Web site.

# Getting ready to install programs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Getting ready to install programs**

You should install programs on a terminal server before giving clients access to the terminal server. Doing so ensures that you can test each program before clients access it, and can reduce the time it takes to perform program tuning.

If you are installing a program after clients have been granted access, you must ensure that there are no users logged on to the terminal server during installation. You can send clients a message detailing the time and duration of the installation and then disable all connections before starting.

For more information about currently connected sessions, see Terminal Services Manager Overview.

For more information about sending messages to users, see Send a message to a user.

For more information about disabling a connection, see Disable a Terminal Services connection.

**Note**

- When installing programs, log on using an administrator account.

# Starting the installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Starting the installation**

Use either of the following methods to install programs for multisession Terminal Server access:

- Use **Add or Remove Programs** in Control Panel (recommended).
- Use the **change user** command at the command prompt before and after installing the program.

One of the primary functions of the **change user** command is to ensure that program files are installed to the systemroot rather than the windows subdirectory of the user's home directory (%homepath%\windows). This makes the programs available for multisession access.

Before the program is installed, type **change user /install** at the command prompt to place the system in install mode and turn off .ini file mapping. The system then records how the setup APIs initially install the program.

After the program is installed, type **change user /execute** at the command prompt to return the system to execute mode, restore .ini file mapping, and redirect user-specific data to the user's home directory.

When the user opens the program, user-specific registry setting files (.ini, .dll, .ocx, and so on) are propagated as needed to the user's home directory.

**Add or Remove Programs**, which runs the **change user** command, is the preferred method of installing programs. Enter **change user** at the command prompt only when you install a program by another method and want to ensure multisession access. For example, when Internet Explorer prompts you to install an add-on program, use **change user** at the command prompt to ensure that the program is installed for multisession access.

For more information, see Install a program using Add or Remove Programs.

**Important**

- Microsoft provides scripts for some programs that need tuning in order to work correctly in a multisession environment. Before installing a program, see the Microsoft Web Site for a list of programs that require scripts and information about how to use them.

**Notes**

- Do not copy program files from one directory to another or manually edit the registry during installation.
- Install programs on the terminal server local drive rather than on a file server. This improves program performance by reducing network traffic.
- Install programs on NTFS file system drives rather than on FAT drives. You can only set file permissions on NTFS drives.

# Testing the installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Testing the installation**

Some programs installed using **change user** might not work at first because of the locked-down state of the system. Some 16-bit programs must be able to write to the directory where the program .ini file is stored. Some 32-bit programs create registry entries that are needed for proper operation.

Enable event logging and then log onto the server using a temporary user account created to mimic the user or users that will be accessing the program. Start the program and go through the basic procedures for using it. Then, use Event Viewer to determine which files or directories need Write access and which registry keys require Read access by the user for correct operation.

**Note**

- This process might not find all files, directories, and registry keys for which access is required by the application in all user scenarios. The only way to ensure that you have accounted for all access requirements is to perform tasks manually.

Consider creating two user accounts and testing with each account before proceeding. This allows you to make changes to program preferences and templates, while ensuring that those changes do not affect other users of the same group.

Keep track of all files, and any issues that arise. For information about how to tune programs for multisession use, see Program Compatibility.

# Program Compatibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Program compatibility**
Many programs, particularly older ones, were not designed for multisession use. As a result, some programs need to be tuned before they can run on a terminal server.

Microsoft provides installation scripts for some programs that have known installation issues. Run the scripts after the application installation process is complete and before running the application the first time.

**Application compatibility scripts**
Many commonly used applications have been tested for compatibility with Terminal Server. For maximum performance on a terminal server, some applications require minor changes after installation. Scripts are available for these applications and must be run after the application installation is complete. The scripts are located in systemroot in \Application Compatibility Scripts\Install.

Many newer applications, such as Microsoft Office XP, do not require modification to work properly with Terminal Server. For information about specific applications requiring modification for compatibility with Terminal Server, see the Microsoft Web Site.

**Note**

- For information about specific script capabilities and how to modify them for custom installations, refer to the notes within the script for the application you are installing. You can modify scripts by using a text editor such as **Notepad**.

In addition to running scripts, you might need to perform some additional tuning. The following topics describe some of the changes that might be required:

- Registry settings
- File and directory issues
- Program compatibility flags

# Registry settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Registry settings**
Programs store configuration information in a system database called the registry. For configuration purposes, the two most important sections of the registry are labeled HKEY_LOCAL_MACHINE and HKEY_CURRENT_USER.

You should store information that pertains to the entire system in HKEY_LOCAL_MACHINE. For example, a program might store the following information in this section:

- The components of a program that were installed.
- The path used to load program components.
- The path to a shared database. For example, a point-of-sale program where all users reference the same database.

Programs store user-specific information in the HKEY_CURRENT_USER section, which is part of the user profile. When a user logs on, a profile is loaded into the system registry where it becomes HKEY_CURRENT_USER. When the user logs off, any changes to HKEY_CURRENT_USER are written back to the user profile. A program might store the following types of information in this section:

- Per-user paths are particularly important for multiuser operation. For example, paths to custom dictionaries, mailboxes, configuration files, and temporary directories.
- Per-user preference settings. For example, some users might want to enable background spell checking while others might want to disable it.

**Per-user data in HKEY_LOCAL_MACHINE**
Problems arise if programs store per-user data in HKEY_LOCAL_MACHINE instead of HKEY_CURRENT_USER. If a program stores personal preferences, such as color and font size, in HKEY_LOCAL_MACHINE, users are not able to specify custom preferences.

**Settings created when a program is run**
A program might create HKEY_CURRENT_USER registry settings the first time a user runs the program, rather than during installation. If you do not run the program while install mode is still active, the HKEY_CURRENT_USER settings are not copied to HKEY_LOCAL_MACHINE. The first time each user runs the program, HKEY_CURRENT_USER will be loaded with the default settings. Depending on the program, this might or might not be appropriate. If the default settings are not sufficient, they need to be corrected for each user individually.

To avoid this problem, run the program once before clicking **Finish** in **Add\Remove Programs** (or before running **change user /execute**). This will cause the program to create the necessary registry settings.

If you have already left install mode and have not run the program, re-enter install mode with **change user /install** and run the program. If you have already run the program while in execute mode, create another administrator account, log on with this new account, run **change user /install**, run the program, and then run **change user /execute** to leave install mode. For more information about the change user command, see Change user.

# File and directory issues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File and directory issues**
The files and directories a program creates and uses might present problems if the program was not originally designed for a multiuser environment.

**Files that store per-user data**
A program might use a single path for a file that contains per-user data. In some cases, this can mean that if one user updates this file, it affects all users. More severe problems include preventing multiple users from running the program simultaneously and corruption of the file, which can cause the program to fail for all users.

**Note**

- Most files do not contain per-user data and are safe to share. Some files can be classified either way, depending on the system administrator's preference. For instance, document template files can be made read-only and safely shared. Or, users can have their own set of template files that they can update as desired.

If you have difficulty finding a shared file, you can use file auditing to track successful write operations.

Another path issue arises when programs store files in public directories such as systemroot. Some public directories have read-only permission for nonadministrative users, which prevents accidental modification of common files. However, if a program attempts to write a file into one of these shared directories, the program fails.

A frequent symptom of this problem is a program that can be successfully executed by an administrative user, but fails for a nonadministrative user. To isolate this problem, use file auditing to log all write operations that fail.

Some programs use temporary directories to store information that relates to the current user. The program might not create a separate directory for each user, thus causing temporary data for multiple users to be stored in a single location. If the program is not intelligent about handling this data, one user can delete data that is still needed by another user.

**Updating the path**
When you identify a file that stores per-user data, update the path so that a copy of the file resides in each user's home directory.

Programs can store a file path in several places, including:

- HKEY_LOCAL_MACHINE
- HKEY_CURRENT_USERS
- A configuration file
- Hard-coded within the program itself

Some per-user files are generated by the program. Other per-user files, such as document templates, must be provided for the user by copying them into the user's home directory. If you specified the path as a subdirectory of the user's home directory, you might need to create the subdirectory as well.

Per-user tasks can be performed in a script that is run during logon. The script checks the user's home directory to determine whether the files or directories in question already exist. If not, it can copy them from a master. Using a script allows you both to update current users with a new program and to set up new users.

**Permissions and accessibility**
Another path-related issue arises when a program cannot locate files in or below systemroot.

This problem occurs with some programs that use the **GetWindowsDirectory()** function to determine the name of the directory where Terminal Server was installed. Programs that use this function benefit from an automatic remapping. The function returns the Windows subdirectory of the user's home directory.

This automatic remapping can be disabled with a program compatibility flag. This is appropriate if the files under systemroot are being used only for read access. For more information about compatibility flags, see Program compatibility flags.

# Program compatibility flags

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Program compatibility flags**

Program compatibility flags are registry entries that provide the system with information about compatibility issues for particular programs. Some of these entries are created by default. Administrators can create new entries or change existing entries to alleviate many program compatibility issues.

Following is a list of compatibility flags that can be stored under

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\ Windows NT\CurrentVersion\Terminal Server\ Compatibility\Applications*app_name*

The keys are created using the program's executable name and the flag is stored as a **DWORD** value.

| Value | Description |
|---|---|
| 0x00000001 | MS-DOS-based program |
| 0x00000002 | OS/2-based program |
| 0x00000004 | Windows-based 16-bit program |
| 0x00000008 | Windows-based 32-bit program |
| 0x0000000C | Windows-based 16-bit and 32-bit program |
| 0x0000000F | Any version of a program |
| 0x00000010 | Return user name instead of computer name for **GetComputerName**. |
| 0x00000020 | Return Terminal Server build number instead of Windows 2000 build number for **GetVersion**. |
| 0x00000040 | Synchronize user .ini file to system version.* |
| 0x00000080 | Do not substitute user \Windows directory.** |
| 0x00000100 | Disable registry mapping for program or registry key. |
| 0x00000200 | Per-object user/system global mapping |
| 0x00000400 | Return system \Windows directory instead of user \Windows directory for **GetWindowsDir**. |
| 0x00000800 | Limit the reported physical memory for **GlobalMemoryStatus**. |
| 0x00001000 | Log object creation to file. |
| 0x20000000 | Do not put program to sleep on unsuccessful keyboard polling (Windows-based 16-bit programs only). |

* This setting adds new entries from the system version of the .ini file to the user's version, but does not delete any existing entries in the user's .ini file. By default, the user's .ini file is deleted when it is out of date with the system version.

** This setting requires that the user's \Windows directory not be substituted for the system root directory for any paths inside the .ini file when the system version of the .ini file is copied to the user's \Windows directory. It is recommended that you make this substitution so that the program uses the user's \Windows directory for all of the modifications it makes to data or .ini files.

For more information about registry settings, see Registry settings.

# Deploying Client Software

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying client software**
This section provides information about setting up and configuring Remote Desktop Connection (formerly known as the Terminal Services client) on client computers.

**Important**

- You should install the latest version of Remote Desktop Connection. If you are using the Windows 2000 32-bit Terminal Services Client, you should upgrade to the version of Remote Desktop Connection that is provided with each of the Windows Server 2003 family operating systems.

- About the Terminal Services Client MSI Setup Package

- Remote Desktop Web Connection

# About the Terminal Services Client MSI Setup Package

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**About the Terminal Services Client MSI Setup Package**
The Terminal Services Client MSI Setup package is a tool for deploying Remote Desktop Connection (formerly known as the Terminal Services client) and associated files on computers running earlier versions of Windows. These versions include Microsoft Windows 2000 Server, Microsoft Windows 2000 Professional, Microsoft Windows NT 4.0, Microsoft Windows 98, and Microsoft Windows 95. The MSI Setup package can also be placed on a share for distribution to workstations running Windows 2000 by using Microsoft IntelliMirror.

For procedural information on installing the client with the MSI setup package, see:

- Share the client setup folder at the server
- Install the 32-bit Terminal Services Client (Remote Desktop Connection)

For more information on using Group Policy, see Group Policy overview.

**Note**

- Remote Desktop Connection is installed automatically with Windows XP and Windows Server 2003 family operating systems.

# Remote Desktop Web Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Desktop Web Connection**
Remote Desktop Web Connection is an ActiveX control that provides virtually the same functionality as the full Terminal Services client, but it is designed to deliver this functionality over the Web. When embedded in a Web page, Remote Desktop Web Connection can host a client session with a terminal server, even if the full Terminal Server client is not installed on a user's computer.

For more information, see Deploying Remote Desktop Web Connection.

# Terminal Services Client Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services Client overview**
The Terminal Services Client is a thin client that delivers the 32-bit functionality of Microsoft® Windows Server™ 2003 family operating systems to a wide range of desktop hardware devices. The client appears as a window within the local desktop environment. It contains only the minimum amount of software necessary to establish a connection to the server and present the user interface. All other operating system functions, as well as applications, run on the server. The client can be run on a range of desktop hardware:

- Computers running Microsoft Windows 95, Microsoft Windows 98, Microsoft Windows ME, Microsoft Windows NT 4.0, Microsoft Windows 2000 or Microsoft Windows Server 2003 operating systems.
- Windows-based Terminal devices (embedded).

For more information about Remote Desktop Connection, the new Terminal Services client, see Remote Desktop Connection overview.

For information about specific client features, see:

- Bitmap caching
- Shared clipboard
- Providing client access to local printers
- Log off or disconnect
- Managing Terminal Services connection files
- Using a dial-up connection to access a terminal server

**Note**

- The Terminal Services Client was known as the Terminal Server Client in Windows NT 4.0.

# Remote Desktop Connection overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Desktop Connection overview**
Remote Desktop Connection (formerly known as the Terminal Services client) is installed by default on all Windows Server 2003 family operating systems. You can use Remote Desktop Connection to connect to terminal servers, or to the desktop of a computer running one of the Windows Server 2003 family operating systems or Windows XP for remote administration.

Remote Desktop Connection allows you to create and configure your connection, save your connection settings to a file, and open and edit your saved connection files, all in the same program. For more information on Remote Desktop Connection, see Remote Desktop Connection.

### Notes

- When you install a Windows XP or one of the Windows Server 2003 family operating systems, Remote Desktop Connection is installed, but not enabled. For more information on enabling remote connections, see Enable or disable Remote Desktop.

- To use Group Policy to deploy Remote Desktop Connection, assign the program on a per-computer basis to a terminal server. For more information, see article Q324750, "HOW TO: Assign Software to a Specific Group By Using a Group Policy in Windows Server 2003" in the Microsoft Knowledge Base.

# Bitmap caching

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Bitmap caching**

Bitmap caching places commonly used bitmaps, such as the one that displays the **Start** button, in a file on the client computer. This improves performance by minimizing the amount of display information that must be passed over a connection.

Bitmap caching is enabled on a per-session basis from the client software and requires 10MB of free hard-drive space on the client computer.

For more information, see Configure bitmap caching.

# Shared clipboard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shared clipboard**

On Windows Server 2003 family operating systems, Terminal Services provides seamless clipboard sharing. Clipboard contents can be available to applications locally on a user's computer and to programs within a Terminal Services session.

The shared clipboard synchronizes its contents with the local clipboard and can be viewed using the Windows ClipBook Viewer (clipbrd.exe). For example, you can copy text from a document within a session, and paste it into a document on your local computer. You cannot, however, copy and paste files or folders. The shared clipboard is intended only for data such as text and graphics.

Clipboard sharing can be disabled using Terminal Services Group Policies or Terminal Services Configuration. For more information, see Managing drive, printer, and device mappings for clients.

For more information about the ClipBook Viewer, see ClipBook Viewer overview.

# Providing client access to local printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providing client access to local printers**
Terminal Services provides printer redirection, which routes printing jobs from the terminal server to a printer attached to the client computer, or to a network printer available to the client computer.

Print redirection is enabled by default, but it can be disabled using Terminal Services Group Policies (see Configuring Terminal Services with Group Policy) or on a per-connection basis using Terminal Services Configuration (see Connect client drives and printers at logon).

There are two ways to provide client access to local printers:

- Automatic printer redirection--for Remote Desktop Connection.
- Manual printer redirection--for 16-bit clients, Windows-based Terminals, and local printers that require drivers other than those shipped with Windows Server 2003 family operating systems.

### Notes

- Bidirectional printing is not supported.
- Redirected printers are available for use with applications running on the server. Redirected printers appear in the Printers folder in Control Panel.
- Clients see only their own redirected printers in the **Printers** dialog box.

**Automatic printer redirection**
For Remote Desktop Connection, printer redirection is automatic when the local printer uses a driver that is installed on the server.

When a client logs on to Terminal Services, local printers attached to LPT, COM, and USB ports on the client computer are automatically detected and a local queue is created on the server. The client computer printer settings for the default printer and some properties, such as printing on both sides of the page, are used by the server.

Information about the client's local printers and settings are saved on the client computer. On subsequent logons, the printer queue is created using the information stored on the client computer.

### Warning

- When a client disconnects or ends the session, the printer queue is deleted and any incomplete or pending print jobs are lost.

If a printer driver is not found on the server, an event is logged and the client printer is not created. To make the printer available, the driver must be manually installed on the server.

**Manual printer redirection**
Administrators can manually redirect printers that are attached to LPT and COM ports on the client's local computer so they can be used with Remote Desktop Connection.

### Note

- Manual redirection of printers connected through USB ports is not supported.

To manually redirect a client printer, contact your administrator and provide the name of your computer (or IP address for a Windows-based Terminal). The client must be connected to the terminal server during manual redirection.

The administrator uses **Add Printers** in **Printers** in Control Panel to manually add the printer. The administrator uses the client computer name to select the printer port from the list of available ports in **Add Printer**. Ports for all clients currently connected to the server appear in the **Port** list. Terminal Services client ports are listed as follows:

| Port | Description |
|---|---|
| TS | computername:*PortX* (where *Port* is the type of port and *X* is the port number) |

When a client disconnects or ends the session, the printer queue is deleted, and any incomplete or pending print jobs are lost. Information about client local printers and settings are saved on the client computer. On subsequent logons, the printer queue is created using the information stored on the client computer.

**Notes**

- When automatic redirection for 32-bit client printers fails, an event is logged. Typically, this event in the log contains information about the reason for redirection failure and, if applicable, which server-side driver is missing.
- When the client disconnects or logs off, the printer queue is deleted, and incomplete or pending print jobs are lost.
- After the initial manual redirection, printers are automatically redirected during subsequent logons.

# Providing client access to the local file system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providing client access to the local file system**

On Windows Server 2003 family operating systems, Terminal Services provides client drive redirection, making the local file system available on the remote desktop within a Terminal Services session. If drive redirection is specified, client drives appear in Windows Explorer with the designation "tsclient <driveletter>".

When you install Terminal Server, client drive mapping is disabled by default. To enable this feature, use Terminal Services Group Policies or the Terminal Services Configuration tool.

# Providing client access to server audio

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providing client access to server audio**

On Windows Server 2003 family operating systems, Terminal Server administrators can enable users to play server audio on their local computers. This feature is disabled by default when you install Terminal Server. To enable it, configure the **Allow audio redirection** policy in Terminal Services Group Policies. You can also enable audio redirection locally using the Terminal Services Configuration tool. For more information, see Connect client drives and printers at logon.

# Log off or disconnect

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Log off or disconnect**

Terminal Services provides two options when closing a client session: you can either log off or disconnect.

Disconnecting leaves the session running on the server. You can reconnect to the server and resume the session. If you are running a task on the server, such as running a time consuming query on a database, you can start the task and disconnect from the session. Later, you can log back onto the server, reenter the session, and check the results. This is especially helpful when operating over a remote access connection on a long-distance toll line. For information about how to disconnect from a Terminal Services session, see Disconnect without ending a session.

Network administrators can limit the amount of time a disconnected session is allowed to remain on the server. For more information, see Set time-out settings for disconnected, active, and idle sessions.

Logging off ends the session running on the server. Any applications running within the session will be closed and unsaved changes made to open files will be lost. The next time you log onto the server, a new session is created. For information about how to log off from a Terminal Services session, see Log off and end the session.

If printer redirection is enabled for the session, when the session is disconnected or logged off, the printer queue is deleted and any incomplete or pending print jobs are lost.

# Managing Terminal Services connection files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Terminal Services connection files**

Remote Desktop Connection allows administrators and users to set up predefined connection files to servers for single program or full desktop access. By using Remote Desktop Connection, you can save connection information and settings to a connection file that is saved in the My Documents/Remote Desktops folder with the .rdp extension. This way, for example, an administrator who wants to provide a single line-of-business program across the computing environment can create the connection file, and distribute that file to all desktops.

For information on how to use the Remote Desktop Connection to save connection settings to a file, see Create a connection file.

# Using a dial-up connection to access a terminal server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using a dial-up connection to access a terminal server**

If you have remote access to a network that contains terminal servers, you can access those terminal servers using a dial-up connection from a remote location.

To use remote access to connect to a terminal server, you must first configure a dial-up connection on your computer. Then you can connect to the network where the terminal server resides and use Remote Desktop Connection to create a connection to the terminal server itself. You can then log on to the terminal server as usual.

Administrators can grant users the ability to reconnect to a session if the connection is broken, for example, if a phone connection is interrupted. For more information, see Set time-out settings for disconnected, active, and idle sessions to allow reconnection on a per-connection basis or Change settings for when a session limit is reached or a connection is broken on a per-user basis.

# Troubleshooting Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The connection does not automatically log on to the server. I am prompted for my User name and Password on connection.
- Installation of an .msi package fails from a remote Terminal Services session.
- Installed program does not work correctly.
- Client fails to connect with "This initial program cannot be started" message.

**The connection does not automatically log on to the server. I am prompted for my User name and Password on connection.**
**Cause:** Encryption enhancements for Windows 2000 and Windows Server 2003 family operating systems cause the Windows NT 4.0 version of the Terminal Services client not to recognize the user name and password in the automatic logon portion of the connection file.

**Solution:** From the Windows NT 4.0 version of Client Connection Manager, right-click the connection and select **Properties**. On the **General** tab, check **Automatic logon** and type the User name, Password, and Domain.

**Installation of an .msi package fails from a remote Terminal Services session.**
**Cause:** You must be logged in as a member of the Administrators group to install .msi packages for Terminal Services from a remote session.

**Solution:** Log on as a member of the Administrators group or use Run as to install the .msi package.

**See also:** Install a program using Add or Remove Programs

**Installed program does not work correctly.**
**Cause:** The program was installed before Terminal Server was installed.

**Solution:** Uninstall the program, and reinstall using either **Add or Remove Programs** or the **change user** command.

**See also:** Starting the installation

**Cause:** The program was installed using **change user /execute** rather than **change user /install**.

**Solution:** Uninstall the program, and reinstall using either **Add or Remove Programs** or the **change user /install** command.

**See also:** Starting the installation

**Cause:** Program compatibility scripts were not run properly.

**Solution:** Run the program scripts again.

**See also:** Program Compatibility

**Client fails to connect with "This initial program cannot be started" message.**
**Cause:** The client connection might have been configured to start a program upon connection, but the program path and filen ame are incorrect.

**Solution:** Check the properties for the client connection to ensure that the program path and file name are correct, and that the program exists in this location on the server.

**Cause:** The server administrator might have specified a program to start upon connection, using Group Policy, Terminal Services Configuration, or other Terminal Services administration tools.

**Solution:** Check the server configuration to ensure that the program path and file name specified for this user are correct,

and that the program exists in this location on the server.

# Terminal Server Licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Server Licensing**
A valid license, issued by a Terminal Server license server, is required before a client is permitted to log on to a terminal server.

**Important**

- The Terminal Server Licensing method is separate from the licensing method used for Microsoft® Windows Server 2003 operating system clients. For more information on licensing for Windows Server 2003 operating systems, see Licensing.

A license server stores all client licenses. A terminal server must be able to connect to an activated license server before clients can be issued permanent licenses. When you activate a license server, Microsoft provides the server with a digital certificate that validates server ownership and identity. By using this certificate, a license server can make transactions with Microsoft and receive client licenses for your terminal servers. If you install a license server but do not activate it, that server only issues temporary licenses.

For small deployments, it is acceptable to install both the terminal server and Terminal Server Licensing on the same physical computer. However, for larger deployments consider installing Terminal Server Licensing on a separate server. Note that one license server can serve many terminal servers concurrently.

**Important**

- You must configure Terminal Server Licensing correctly in order for your terminal server to continue to accept connections from clients. To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent client access license (CAL), or after 120 days, whichever comes first.

Before installing the license server, determine which of the two license server roles you require—a Domain license server or an Enterprise license server. By default, a license server is installed as an Enterprise license server. For more information, see Terminal Server license server roles.

After installing the license server, you must activate it and install CALs by using the Terminal Server License Server Activation Wizard. For information about activating a license server, see Activating a Terminal Server license server.

The following resources provide additional information about Terminal Server and Terminal Server Licensing.

- For information about deploying Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627)

- For detailed information about Terminal Server Licensing, see Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

- For help with specific tasks, see Terminal Server Licensing How To....

- For general conceptual information, see Terminal Server Licensing Concepts.

**Note**

- You can use any computer running one of the Windows Server 2003 operating systems to administer servers remotely with Remote Desktop for Administration (formerly Terminal Services in Remote Administration mode). You do not need to install Terminal Server. A Terminal Server license server and the Terminal Server Licensing administrative tool are not required for Remote Desktop for Administration. For more information, see Remote Desktop for Administration.

# Terminal Server Licensing How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install Terminal Server Licensing
- Configure a Terminal Server License Server
- Configure a Terminal Server to Receive Licenses

# Install Terminal Server Licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Terminal Server Licensing**

1. Open Add or Remove Programs in Control Panel.

2. Click **Add/Remove Windows Components** to start the Windows Components wizard.

3. In **Windows Components**, select **Terminal Server Licensing**, and then click **Next**.

4. In **Terminal Server Licensing Setup**, do one of the following:

   - Click **Your entire enterprise** to choose an Enterprise license server. You can only install an Enterprise license server on a domain controller or a member server in a domain, not on a stand-alone server. This role is published to Active Directory.

   - Click **Your domain or workgroup** to choose a Domain license server. If you select this option, the computer on which you install the license server does not have to be a domain controller. However, if you want the license server to be automatically discovered by terminal servers that communicate with it, the computer on which you install the license server must be a domain controller. If the computer on which you plan to install the license server is in a workgroup, it can only be automatically discovered by terminal servers on the same subnet.

   For more information about Terminal Server license server roles and the license server discovery process, see Terminal Server license server roles.

5. Provide the file system location where the license server database should be installed on the Terminal Server license server, click **Next**, and then click **Finish**.

   The default location for the license server database is systemroot\System32\LServer.

**Important**

- Before you install Terminal Server Licensing, verify that the operating system of the license server is compatible with the operating system of the terminal server. Note the following:

  - A license server running Windows Server 2000 cannot issue client access licenses (CALs) to terminal servers running Windows Server 2003 operating systems.

  - A license server running a Windows Server 2003 operating system can issue CALs to terminal servers running Windows Server 2000 or Windows Server 2003, or to terminal servers running Windows Server 2000 and terminal servers running Windows Server 2003 in a mixed environment.

**Notes**

- If you are installing a Domain license server, you must be a member of the Administrators group on the local computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

  If you are installing an Enterprise License Server, you must meet one of the following conditions:

  - You must be a member of the Enterprise Admins group in Active Directory.

- You must have the right to create the TS-Enterprise-License-Server object in the same site in which the license server will reside in Active Directory.

As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

- Terminal Serviceslicensingopen

- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://www.microsoft.com/) and Windows Server 2003 Terminal Server Licensing (http://www.microsoft.com/).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Server Licensing overview
Activate a Terminal Server License Server
Install Client Access Licenses
Terminal Server license server roles

# Configure a Terminal Server License Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure a Terminal Server license server**

- Activate a Terminal Server License Server
- Purchase client access licenses
- Install Client Access Licenses
- Deactivate a Terminal Server license server
- Reactivate a Terminal Server license server
- Change Terminal Server license server properties

# Activate a Terminal Server License Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Activate a Terminal Server license server**

The Microsoft Clearinghouse is the facility that Microsoft maintains to activate license servers, issue client access licenses (CALs) to license servers, and to reactivate license servers. You can activate a Terminal Server license server by using one of three methods: Automatic, Web Browser, or Telephone. For instructions on activating a license server and locating a telephone number to contact the Clearinghouse, see the following procedures:

- Locate the Microsoft Clearinghouse telephone number for your country or region
- Activate a Terminal Server license server by using the telephone
- Activate a Terminal Server license server automatically
- Activate a Terminal Server license server by using a Web browser

# Locate the Microsoft Clearinghouse telephone number for your country or region

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To locate the Microsoft Clearinghouse telephone number for your country or region**

- When the Terminal Server license server is running
- When the Terminal Server license server is not running

**When the Terminal Server license server is running**

1. Open **Terminal Server Licensing**.
2. In the details pane, right-click the Terminal Server license server, and then click **Properties**.
3. On the **Installation Method** tab, in the **Installation Method** list, click **Telephone**.
4. In **Select Country or Region**, click your country or region, and then click **OK**.
5. In the details pane, right-click the license server, click **Install Licenses**, and then click **Next**.
6. Note the phone number that is displayed, and then click **Cancel**.
7. To change the server installation method from Telephone to another value, in the details pane, right-click the license server, click **Properties**, change the installation method as needed on the **Installation Method** tab, and then click **OK**.

**When the Terminal Server license server is not running**

1. Open **Terminal Server Licensing.**
2. In the details pane, right-click the Terminal Server license server, and then click **Activate Server** to start the Terminal Server License Server Activation Wizard.
3. On the **Welcome** page, click **Next**.
4. On the **Connection Method** page, in **Activation method**, click **Telephone**, and then click **Next**.
5. On the **Country or Region** page, click your country or region, and then click **Next**.

## Notes

- The Microsoft Clearinghouse is the facility that Microsoft maintains to activate license servers, issue CALs to license servers, recover CALs, and reactivate license servers. It is not the facility for purchasing CALs. For instructions on purchasing CALs, see Related Topics.
- After you purchase CALs, make sure to retain the License Purchase Agreement information. Having this information with you will facilitate communication with the Microsoft Clearinghouse, should you need assistance with recovering CALs or other Terminal Server Licensing tasks that are supported by the Clearinghouse.
- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing

(http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Purchase client access licenses
Purchasing and installing client access licenses on a Terminal Server license server
Terminal Server Licensing overview

# Activate a Terminal Server license server by using the telephone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate a Terminal Server license server by using the telephone**

1. Open Terminal Server Licensing.
2. In the console tree, right-click the Terminal Server license server you want to activate, and then click **Activate Server** to start the Terminal Server License Server Activation Wizard.
3. In the **Activation method** list, select **Telephone**, and then click **Next**.
4. Follow the instructions in the wizard.

### Important

- To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent client access license (CAL), or after 120 days, whichever comes first. For more information about the licensing grace period, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- Terminal Serviceslicensingopen
- Licensing wizard properties, such as installation method and company information, set during the activation process can be changed later.
- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install client access licenses by using the telephone
Terminal Server Licensing grace period
Connect to a Terminal Server license server
Activating a Terminal Server license server
Change Terminal Server license server properties

# Activate a Terminal Server license server automatically

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate a Terminal Server license server automatically**

1. Open Terminal Server Licensing.

2. In the console tree, right-click the Terminal Server license server you want to activate, and then click **Activate Server** to start the Terminal Server license server Activation Wizard.

3. In the **Activation method** list, select **Automatic connection (recommended)**, and then click **Next**.

4. Follow the instructions in the wizard.

   If necessary, at a later time you can change any license server properties, such as activation method and company information, that you set during the activation process.

**Important**

- If the Internet Explorer Enhanced Security Configuration setting is in effect on your Terminal Server license server, you must add the Microsoft Clearinghouse Web site (https://activate.microsoft.com/) to the Trusted sites zone in Internet Explorer before you can activate your Terminal Server license server. If you try to connect to the Microsoft Clearinghouse Web site and it is not yet a trusted site, you are prompted to specify whether to add it to the Trusted sites zone.

  For more information, see Internet Explorer Enhanced Security Configuration.

**Important**

- To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent client access license (CAL), or after 120 days, whichever comes first. For more information about the licensing grace period, see Terminal Server Licensing grace period.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Terminal Serviceslicensingopen

- Information that you provide over the Internet or World Wide Web to the Microsoft Clearinghouse is encrypted for your security and is securely stored by the Microsoft Clearinghouse in a dedicated database. This information is kept separate from information you might have provided to Microsoft in other contexts, and is not disclosed to third parties. For more information about Terminal Server license management and your privacy, see Related Topics.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install client access licenses automatically
Connect to a Terminal Server license server
Activating a Terminal Server license server
Change Terminal Server license server properties
Terminal Server license management and your privacy

# Activate a Terminal Server license server by using a Web browser

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate a Terminal Server license server by using a Web browser**

1. Open Terminal Server Licensing.

2. In the console tree, right-click the Terminal Server license server you want to activate, and then click **Activate Server** to start the Terminal Server License Server Activation Wizard.

3. In the **Activation method** list, select **Web Browser**, and then click **Next**.

4. Follow the instructions in the wizard.

   If necessary, at a later time you can change any license server properties, such as activation method and company information, that you set during the activation process.

### Important

- If the Internet Explorer Enhanced Security Configuration setting is in effect on your Terminal Server license server, you must add the Microsoft Clearinghouse Web site (https://activate.microsoft.com/) to the Trusted sites zone in Internet Explorer before you can activate your Terminal Server license server. If you try to connect to the Microsoft Clearinghouse Web site and it is not yet a trusted site, you are prompted to specify whether to add it to the Trusted sites zone. For more information, see Internet Explorer Enhanced Security Configuration.

### Important

- To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent client access license (CAL), or after 120 days, whichever comes first. For more information about the licensing grace period, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Terminal Serviceslicensingopen

- Information that you provide over the Internet or World Wide Web to the Microsoft Clearinghouse is encrypted for your security and is securely stored by the Microsoft Clearinghouse in a dedicated database. This information is kept separate from information you might have provided to Microsoft in other contexts, and is not disclosed to third parties. For more information about Terminal Server license management and your privacy, see Related Topics.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install client access licenses by using a Web browser
Connect to a Terminal Server license server
Terminal Server Licensing grace period
Activating a Terminal Server license server
Change Terminal Server license server properties
Terminal Server license management and your privacy

# Purchase client access licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To purchase client access licenses**

- In the United States, call (800) 426-9400 or visit the Licensing Program Reseller Web site (http://go.microsoft.com/fwlink/?LinkID=26222).
- In Canada, call the Microsoft Resource Centre at (877) 568-2495.
- If you are outside the United States or Canada and have Internet access, please review the Worldwide Microsoft Licensing (http://go.microsoft.com/fwlink/?LinkID=26223) Web sites or contact your local Microsoft subsidiary on the Microsoft Worldwide Home Web site (http://go.microsoft.com/fwlink/?LinkID=26226).

**Note**

- You must purchase CALs and install them on the computer running Terminal Server Licensing. The Terminal Server license server then issues them to terminal servers for distribution to clients. For more information about purchasing and installing CALs on a license server, see Related Topics.
- Verify that you have purchased enough CALs for the Terminal Server license server and that the Terminal Server Licensing mode matches the type of CAL available on the license server. Terminal Server licensing modes determine the type of CAL that a license server issues to a client. In Windows Server 2003, two types of CALs are available: Per User and Per Device. A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices. In this case, only one CAL is needed instead of a CAL for each device. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used.
- After you purchase CALs, make sure to retain the License Purchase Agreement information. Having this information with you will facilitate communication with the Microsoft Clearinghouse, should you need assistance with recovering CALs or other Terminal Server Licensing tasks that are supported by the Clearinghouse. The Microsoft Clearinghouse is the facility that Microsoft maintains to activate license servers, issue CALs to license servers, recover CALs, and reactivate license servers. For information about how to locate the Microsoft Clearinghouse telephone number for your region, see Related Topics.
- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Locate the Microsoft Clearinghouse telephone number for your country or region
Required licenses
Install Client Access Licenses
Purchasing and installing client access licenses on a Terminal Server license server

# Install Client Access Licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install client access licenses**

You can install client access licenses (CALs) by using one of three methods: the Internet, a Web browser, or telephone. Installing CALs by using the Internet is the easiest method. For instructions on installing CALs, see the following procedures:

- Install client access licenses automatically
- Install client access licenses by using a Web browser
- Install client access licenses by using the telephone
- Repeat the installation of a client access license

# Install client access licenses automatically

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install client access licenses automatically**

1. Open Terminal Server Licensing.

2. Verify that the installation method for the Terminal Server license server is set to **Automatic connection (recommended)** by right-clicking the Terminal Server license server for which you want to install client access licenses (CALs), and then clicking **Properties**. On the **Installation Method** tab, change the installation method if necessary.

3. In the console tree, right-click the Terminal Server license server for which you want to install CALs, click **Install Licenses** to start the Terminal Server CAL Installation Wizard, and then click **Next**.

   The previous steps are not necessary if the Terminal Server CAL Installation Wizard is already started.

4. In **Program and Client License Information**, provide the required information for your licensing program to receive your CALs, and then click **Next**.

   The Microsoft Clearinghouse processes your request, and installs the encrypted CALs on your Terminal Server license server.

5. Click **Finish** to complete the process.

   The Terminal Server license server can now issue licenses to clients that connect to a Terminal server.

## Notes

- Terminal Serviceslicensingopen

- You cannot install CALs until your Terminal Server license server is activated.

- Contact the Microsoft Clearinghouse if you receive an error during the CAL installation process. To find the nearest phone number, you must first go to **Licensing Wizard Properties**, select **Telephone** as the installation method, and then click **OK.** Then, right-click the Terminal Server license server and click **Install licenses** to display the appropriate number.

- On computers running one of the Windows Server 2003 operating systems, you can install Per Device and Per User CALs. Per User CALs, however, are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. On computers running the Windows 2000 Server operating system, you can install Per Device and Internet Connector CALs.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Activate a Terminal Server license server automatically
Connect to a Terminal Server license server
Recovering client access licenses on a Terminal Server license server
Purchasing and installing client access licenses on a Terminal Server license server

# Install client access licenses by using a Web browser

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install client access licenses by using a Web browser**

1. Open Terminal Server Licensing.

2. Verify that the installation method for the Terminal Server license server is set to **Web Browser** by right-clicking the Terminal Server license server for which you want to install client access licenses (CALs), and then clicking **Properties**. On the **Installation Method** tab, change the installation method if necessary.

3. In the console tree, right-click the Terminal Server license server for which you want to install client access licenses, click **Install Licenses** to start the Terminal Server CAL Installation Wizard, and then click **Next**.

   The previous steps are not necessary if the Terminal Server CAL Installation Wizard is already started.

4. Connect to the Terminal Server Licensing web site (https://activate.microsoft.com). You can connect to this site on any computer that has Internet access.

   On the **Windows Terminal Services** Web page, click **Install Client Access License Tokens** and then click **Next**.

5. Provide the following required information:

   - License server ID, a 35-digit number, in groups of 5 numerals, which is displayed in the Terminal Server CAL Installation Wizard.
   - Purchase method (License Program)
   - Last name or surname
   - First name or given name
   - Company name
   - Country or region

   You can also provide the optional information requested, such as company address, e-mail address, and phone number. In the organizational unit field, you can describe the unit within your organization that this server will serve.

   Click **Next**.

6. Type the CAL numbers in the **License Codes:** boxes. If you have more than 5 numbers, type the first 5 and then click **Add**. Continue typing additional numbers and clicking **Add** until all numbers are entered, and then click **Next**.

7. On the next page, verify that all of the information you have entered is correct, and then click **Next** to submit your request to the Microsoft Clearinghouse. The Web page then displays a CAL ID generated by the Microsoft Clearinghouse.

8. Type the CAL ID in the **Enter the license key pack ID in the boxes below. This ID is obtained from the Terminal Server Licensing Web site** text box, and then click **Next**.

   The CALs are installed on your Terminal Server license server.

9. Click **Finish** to complete the process.

The Terminal Server license server can now issue licenses to clients that connect to a terminal server.

## Notes

- Terminal Serviceslicensingopen
- You cannot install CALs until your Terminal Server license server is activated.
- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Activate a Terminal Server license server by using a Web browser
Connect to a Terminal Server license server
Recovering client access licenses on a Terminal Server license server
Purchasing and installing client access licenses on a Terminal Server license server

# Install client access licenses by using the telephone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install client access licenses by using the telephone**

1. Open Terminal Server Licensing.

2. Verify that the installation method for the Terminal Server license server is set to **Telephone** by right-clicking the Terminal Server license server for which you want to install client access licenses (CALs), and then clicking **Properties**. On the **Installation Method** tab, change the installation method if necessary.

3. Right-click the Terminal Server license server for which you want to install client access licenses, click **Install Licenses** to start the Terminal Server CAL Installation Wizard, and then click **Next**.

   The previous steps are not necessary if the Terminal Server CAL Installation Wizard is already started.

4. Call the Customer Support Center at the number that is displayed on the **Obtain client license key pack** page of the wizard, and give the representative your Terminal Server license server ID and the required information for the licensing program under which you purchased client licenses. The representative then processes your request to install CALs, and gives you a unique ID for the CAL.

5. In **Type the client access license ID in the boxes below**, type the CAL ID provided by the representative, and then click **Next**.

   The client access licenses are installed on your Terminal Server license server.

6. Click **Finish** to complete the process.

   The Terminal Server license server can now issue licenses to clients that connect to a terminal server.

## Notes

- Terminal Serviceslicensingopen

- You cannot install CALs until your Terminal Server license server is activated.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Activate a Terminal Server license server by using the telephone
Connect to a Terminal Server license server
Recovering client access licenses on a Terminal Server license server
Purchasing and installing client access licenses on a Terminal Server license server
Locate the Microsoft Clearinghouse telephone number for your country or region

# Repeat the installation of a client access license

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To repeat the installation of a client access license**

1. Open Terminal Server Licensing.

2. In the console tree, right-click the Terminal Server license server for which you want to reload the client access license (CAL), point to **Advanced**, and then click **Repeat Last Installation**. The Licensing Wizard starts.

3. In **Information Needed**, confirm that the required information is correct, and then click **Next**.

   Your request is sent to the Microsoft Clearinghouse for processing.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Terminal Serviceslicensingopen

- You cannot repeat the installation of a CAL by using the Web Browser installation method.

- You might need to reinstall Terminal Server CALs in situations such as the following:

  - A catastrophic failure or any other event that causes the loss of or corruption to the licensing database, when a known good backup is not available. In this case, you must reinstall Terminal Server Licensing, reactivate the license server, and then reinstall the previously issued CALs.

  - If you upgrade or migrate a Windows 2000 license server to Windows Server 2003. In this case, you must reinstall the Windows 2000 Terminal Server CALs on the Windows Server 2003 license server.

    In each of these cases, you must call the Microsoft Clearinghouse to reinstall the appropriate CALs. For information about how to call the Clearinghouse, see Locate the Microsoft Clearinghouse telephone number for your country or region.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Client Access Licenses
Connect to a Terminal Server license server
Change Terminal Server license server properties
Recovering client access licenses on a Terminal Server license server
Purchasing and installing client access licenses on a Terminal Server license server

# Deactivate a Terminal Server license server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To deactivate a Terminal Server license server**

- Using the Automatic method
- Using the telephone method

**Using the Automatic method**

1. Open Terminal Server Licensing.
2. In the console tree, right-click the Terminal Server license server you want to deactivate, point to **Advanced**, and then click **Deactivate Server**. The Terminal Server License Server Wizard starts.
3. In **Required Information**, confirm that your name, phone number, company, and country or region are correct, and then click **Next**.

   Your request to deactivate the Terminal Server license server is sent to the Microsoft Clearinghouse for processing.
4. Click **Finish**.

**Using the telephone method**

1. Open Terminal Server Licensing.
2. In the console tree, right-click the Terminal Server license server you want to deactivate, point to **Advanced**, and then click **Deactivate Server**. The Terminal Server License Server Wizard starts.
3. Select **Telephone**, and then click **Next.**
4. Specify your location, and then click **Next.**
5. Call the telephone number displayed in the wizard, and give the Customer Support representative the Product ID that is displayed below the telephone number.
6. In the boxes in the wizard, type the 35-digit confirmation code provided by the Customer Support representative, and then click **Next.**
7. Click **Finish**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- Terminal Serviceslicensingopen
- For information about the Microsoft privacy statement, see Terminal Server license management and your privacy.
- You cannot deactivate a Terminal Server license server by using the Web Browser connection method.

- You might need to deactivate a Terminal Server license server when the certificate of the server has expired or becomes corrupted, or when the server is being redeployed. However, you will be prompted to reactivate the Terminal Server license server when its registration has expired.

- When you deactivate a Terminal Server license server, you will not be able to license additional clients from this server until it is reactivated.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Reactivate a Terminal Server license server
Connect to a Terminal Server license server
Change Terminal Server license server properties

# Reactivate a Terminal Server license server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reactivate a Terminal Server license server**

1. Open Terminal Server Licensing.

2. In the console tree, right-click the Terminal Server license server you want to reactivate, point to **Advanced**, and then click **Reactivate Server**. The Terminal Server License Server Wizard starts.

3. In **Information Needed**, confirm that your name, phone number, and e-mail address are correct.

   In **Reason**, select the appropriate reason for reactivating the license server, and click **Finish**.

   Your request to reactivate the license server is sent to the Microsoft Clearinghouse for processing.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Terminal Serviceslicensingopen

- For information about the Microsoft privacy statement, see Related Topics.

- When reactivating a license server by using the telephone, you will receive a telephone call confirming that the license server has been reactivated.

- When a license server is reactivated, a record of your licenses is retained. Licenses that were already issued remain valid, and any un-issued licenses are also valid, but must be reissued through the Microsoft Clearinghouse.

- You must reactivate a license server when its registration has expired.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Deactivate a Terminal Server license server
Terminal Server license management and your privacy
Connect to a Terminal Server license server
Change Terminal Server license server properties

# Change Terminal Server license server properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change Terminal Server license server properties**

1. Open Terminal Server Licensing.

2. In the console tree, right-click the Terminal Server license server for which you want to change properties, and then click **Properties**.

   The **Licensing Wizard Properties** dialog box opens.

3. Click any of the following tabs to modify properties:

   - Click **Installation Method** to select a different method for communicating with the Microsoft Clearinghouse. You can select **Automatic connection (recommended), Web Browser,** or **Telephone.**

   - Click **Required Information** to make changes to your name, company, and country or region.

   - Click **Optional Information** to make changes to your e-mail address and company address or to your organizational unit.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Terminal Serviceslicensingopen

- For information about the Microsoft privacy statement, see Terminal Server license management and your privacy.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Server Licensing overview
Terminal Server Licensing administrative tool

# Configure a Terminal Server to Receive Licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure a Terminal Server to receive licenses**

- Configure the Terminal Server Licensing mode
- Set preferred Terminal Server license servers
- Connect to a Terminal Server license server

# Configure the Terminal Server Licensing mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the Terminal Server Licensing mode**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click **Terminal Server Licensing Mode**.
3. Click **Enabled**.
4. Click either **Per User** or **Per Device**, and then click **OK**.

   A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices.

### Important

- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- This Group Policy setting is only available on Windows Server 2003 with Service Pack 1.
- Use this procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an administrator, and then open Group Policy by using the **Active Directory Users and Computers** snap-in.
- Group Policy overrides the configuration set with the Terminal Services Configuration tool.
- By default, when the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.
- Users can create more than one session by specifying a different program to start on connection for each session.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.

2. In the console tree, click **Server Settings**.

3. In the details pane, right-click **Licensing Mode**, and then click **Properties**.

4. Select either **Per User** or **Per Device**, and then click **OK**.

   A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices.

## Notes

- Terminal Servicesccopen

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

- By default, when the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.

- Users can create more than one session by specifying a different program to start on connection for each session.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Server Licensing
Terminal Server Licensing overview
Terminal Server license server roles

# Set preferred Terminal Server license servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set preferred Terminal Server license servers**
At a stand-alone server:

- Using Group Policies (best practice)
- Using Terminal Services Configuration
- Using the registry

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click **Use the specified Terminal Server license servers**.
3. Click **Enabled**.
4. Enter the names of the Terminal Server license servers you want this Terminal Server to contact first when attempting to obtain client access licenses (CALs), and then click **OK**.

### Important

- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- This Group Policy setting is only available on Windows Server 2003 with Service Pack 1, or with the Service Pack 1 versions of System.adm Group Policy template files installed in your Terminal Server environment.
- Use this procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an administrator. Then, you must open Group Policy through the **Active Directory Users and Computers** snap-in.
- You can also set preferred license servers by using a Windows Management Instrumentation (WMI) script. For three scripts that you can use to set preferred license servers, delete preferred license servers, or query preferred license servers, see "Preferred License Server WMI Scripts" in Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220). .

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Server Settings**.

3. In the details pane, right-click **License server discovery mode**, and then click **Properties**.

4. Click **Use these license servers**, and enter the names or IP addresses of the license servers in the text box, separating them with commas.

5. Click **Check names**, make any necessary corrections, and click **OK**.

### Notes

- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

- This feature is only available on Windows Server 2003 with Service Pack 1.

- You can also set preferred license servers by using a Windows Management Instrumentation (WMI) script. For three scripts that you can use to set preferred license servers, delete preferred license servers, or query preferred license servers, see "Preferred License Server WMI Scripts" in Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

#### Using the registry

1. Click **Start**, click **Run**, type **regedit**, and then click **OK**.

2. Locate, and then click, the following key in the registry:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\TermService\Parameters**

3. On the **Edit** menu, point to **New**, click **Key**, and then type **LicenseServers** to name the new key.

4. Locate, and then click, the following key in the registry:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\TermService\Parameters\LicenseServers**

5. On the **Edit** menu, point to **New**, click **Key**, and then type ***ServerName*** where *ServerName* is the NetBIOS name of the license server that you want to use, and then press **ENTER**.

   The new key name can be any of the following designations that represent the license server:

   - The NetBIOS name of the server

   - The fully-qualified domain name (FQDN) of the server

   - The IP address of the server

6. Restart your computer

### Notes

- You can also set preferred license servers by using a Windows Management Instrumentation (WMI) script. For three scripts that you can use to set preferred license servers, delete preferred license servers, or query preferred license servers, see "Preferred License Server WMI Scripts" in Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

- For more information about Terminal Server, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627)

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change Terminal Server license server properties
Resultant Set of Policy
Terminal Server license server roles

# Connect to a Terminal Server license server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a Terminal Server license server**

1. Open Terminal Server Licensing.
2. On the **Action** menu, click **Connect**.
3. In **Server**, type the name of the Terminal Server license server to which you want to connect, and then click **OK**.

## Notes

- Terminal Serviceslicensingopen
- When Terminal Server Licensing opens, it locates all license servers in your workgroup or domain.
- To connect to a license server, a terminal server must have the **Access this computer from the network** logon right for the target license server. Additionally, if the License Server Security Group setting in Group Policy is enabled, the terminal server computer account must be a member of this group. Enabling this Group Policy setting and applying it to your license server creates a local group called Terminal Services Computers. The license server issues licenses only to the terminal servers in this group. You must add both the terminal servers for which you need to provide licenses and any license servers that might need to acquire licenses to this group for each license server. For more information, see Related Topics.
- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Activate a Terminal Server License Server
Install Client Access Licenses
Controlling the issuance of Terminal Server licenses

# Terminal Server Licensing Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

- Terminal Server Licensing overview
- Terminal Server Licensing grace period
- Controlling the issuance of Terminal Server licenses
- Understanding Terminal Server Licensing
- Using Terminal Server Licensing

# Terminal Server Licensing overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Server Licensing overview**
Terminal Server has its own method for licensing clients that log on to terminal servers, separate from the licensing method for Microsoft® Windows Server 2003 family operating system clients. When a client connects to a terminal server, the terminal server determines if a client access license (CAL) is needed. The terminal server then requests a CAL from the Terminal Server license server, and then the CAL is issued to the client by the terminal server. Although Terminal Server provides a licensing grace period during which no license server is required, after the grace period ends, clients must receive a valid CAL issued by a license server before they can log on to a terminal server.

**Important**

- If you use Remote Desktop for Administration, two concurrent connections can log on. You do not need a Terminal Server license server for these connections. For more information on Remote Desktop for Administration, see Remote Desktop for Administration overview.

**Prerequisites**
To use Terminal Server, you must also have at least one license server. For small deployments, it is acceptable to install both the Terminal Server and Terminal Server Licensing service on the same physical computer; however, for larger deployments, it is recommended that Terminal Server Licensing be installed on a separate server.

Before installing Terminal Server Licensing, verify that the operating system of the license server will be compatible with the operating system of the terminal server. Note the following:

- A Windows 2000 license server cannot issue CALs to terminal servers running Windows Server 2003 operating systems.

- A license server running a Windows Server 2003 operating system can issue CALs to terminal servers running Windows Server 2000 or Windows Server 2003, or to terminal servers running Windows Server 2000 and terminal servers running Windows Server 2003 in a mixed environment.

After you install Terminal Server Licensing, you must activate the license server so that the license server can issue permanent CALs for Terminal Server clients. After the license server is activated and you purchase CALs, it provides a secure way to install, and issue permanent CALs to Terminal Server clients. If you do not activate a license server, it issues temporary CALs, which allow clients to connect to the terminal server for 90 days.

**Terminal Server Licensing Components**
Following are descriptions of the primary Terminal Server Licensing components:

**Microsoft Clearinghouse**
The Microsoft Clearinghouse is the facility that Microsoft maintains to activate license servers, issue CALs to license servers, recover CALs, and reactivate license servers. The Clearinghouse stores information about all activated license servers and CALs that have been issued. This information helps you track client usage of the terminal servers within your organization to ensure that you have purchased a sufficient number of CALs. The Clearinghouse is accessed directly when you choose the Internet (Automatic) activation method in the Terminal Server Licensing Activation Wizard. If you do not have Internet connectivity, you can also contact the Clearinghouse by phone. For more information, see Locate the Microsoft Clearinghouse telephone number for your country or region.

You cannot purchase CALs through the Microsoft Clearinghouse. For information about purchasing CALs, see Purchase client access licenses. After you purchase CALs, make sure to retain the License Purchase Agreement information. Having this information with you will facilitate communication with the Microsoft Clearinghouse, should you need assistance with recovering CALs or other Terminal Server Licensing tasks that are supported by the Clearinghouse.

**License server**
A license server stores all client licenses. A terminal server must be able to connect to an activated license server before clients can be issued permanent CALs.

Before installing the license server, determine which of the two license server roles you require—a Domain license server or an

Enterprise license server. During Terminal Server Licensing setup, you select one of these roles. By default, a license server is installed as an Enterprise license server. For more information, see Terminal Server license server roles.

**Important**

- You must configure Terminal Server Licensing correctly in order for your terminal server to continue to accept connections from clients. To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent CAL, or after 120 days, whichever comes first. For more information, see Terminal Server Licensing grace period.

To activate a license server, use the Terminal Server License Server Activation Wizard. When you activate a license server, Microsoft provides the server with a limited-use digital certificate that validates server ownership and identity. Microsoft uses the X.509 industry standard certificate for this purpose. The license server then uses this certificate to obtain and issue CALs to clients. For information about activating a license server, see Activating a Terminal Server license server.

**Terminal server**

A terminal server is a server on which the Terminal Server component of the Windows Server 2003 operating system is installed. A terminal server gives clients access to Windows-based applications that run entirely on the server, and supports multiple client sessions on the server. If a client does not have a license, the terminal server requests one from the license server. For instructions on installing Terminal Server, see Installing Terminal Server.

**Note**

- You can use any computer running Windows Server 2003 family operating systems to administer servers remotely with Remote Desktop for Administration (formerly Terminal Services in Remote Administration mode). You do not need to install Terminal Server. A license server and the Terminal Server Licensing administrative tool are not required for Remote Desktop for Administration. For more information about Remote Desktop for Administration, see Remote Desktop for Administration.

**Client licenses**

After the Terminal Server Licensing grace period ends, you must purchase and install the appropriate number of CALs for each device or user that needs connect to a terminal server. In addition, you must verify that the Terminal Server Licensing mode that you specify matches the type of CAL available on the license server. Terminal Server Licensing modes determine the type of CAL that a license server issues to a client. In Windows Server 2003, two types of CALs are available: Per User and Per Device. A Per Device CAL gives each client computer or device the right to access a terminal server that is running a Windows Server 2003 operating system. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices. In this case, only one CAL is needed instead of a CAL for each device.

When the Per Device licensing mode is used, and a client logs on to a terminal server for the first time, the terminal server issues the client a temporary license by default. When a client logs on to a terminal server for the second time, if the license server is activated and enough Per Device CALs are available, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.

For instructions on configuring the Terminal Server Licensing mode, see Configure the Terminal Server Licensing mode. For information about the types of licenses supported by a license server that runs Windows Server 2003 family operating systems, see Required licenses.

**Resources**

For more information about Terminal Server and Terminal Server Licensing, see:

- Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627)
- Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).
- Terminal Server license management and your privacy

# Terminal Server Licensing grace period

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Server Licensing grace period**

To allow ample time for you to deploy a Terminal Server license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent client access license (CAL), or after 120 days, whichever comes first. For information about the client licenses that are supported by a license server that runs Windows Server 2003 family operating systems, see Required licenses.

The license server issues CALs to the terminal server, which in turn issues them to clients. In order for a license server to issue permanent CALs, you must activate the license server and then purchase and install the appropriate number of permanent CALs. For information about how to activate a license server, see Activate a Terminal Server License Server. For information about how to purchase CALs, see Purchase client access licenses. If a license server is not activated, it issues temporary licenses. These temporary licenses allow clients to connect to the terminal server for 90 days. There is no limit to the number of temporary licenses that a license server can issue, but a single client is only issued a temporary license once. After the temporary license expires, the client can only connect to the terminal server if the license server can issue a permanent CAL, or if the terminal server is still within its licensing grace period.

Terminal Server Licensing modes determine the type of CAL that a license server issues to a client. In Windows Server 2003, two types of CALs are available: Per User and Per Device. A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices. In this case, only one CAL is needed instead of a CAL for each device.

By default, when the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.

For each permanent Per Device CAL that is issued, an expiration period is applied. This expiration period is a random number of days between 52 to 89 days after the license is issued. The terminal server always attempts to renew these CALs seven days before they expire. This functionality facilitates the automatic recovery of Per Device CALs that are lost due to hardware failure, operating system reinstallation, and other, similar events. For more information about the expiration period for Per Device CALs and about recovering CALs, see Recovering client access licenses on a Terminal Server license server.

For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

# Controlling the issuance of Terminal Server licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Controlling the issuance of Terminal Server licenses**

By default, a Terminal Server license server issues a license to any computer that requests one. For security reasons, you might want to restrict the issuance of licenses to specific computers or groups. You can use the **license server Security Group** Group Policy setting to control which servers are eligible to receive licenses.

**Terminal Services Computers**, a local security group, is created by default the first time Terminal Server Licensing is started. When the target computer is a domain controller, this group is a domain local group. For more information about enabling this policy setting, see Control the issuance of Terminal Server licenses.

By default, the **Terminal Services Computers** security group is not populated. You must choose the computers or groups you want to receive licenses and manually add them to this group. Any changes you make to the membership of this group take effect only after you stop and restart the Terminal Server Licensing service. For more information about adding members to a security group, see Changing group memberships.

The most efficient way to manage Terminal Server computer accounts is to create a global group that contains the accounts of all terminal servers that must be licensed. Then, add this global group to the local (or domain local) **Terminal Services Computers** group. This method allows a domain administrator to manage a single list of computer accounts.

For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

# Understanding Terminal Server Licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Terminal Server Licensing**


- Required licenses
- Terminal Server Licensing administrative tool
- Terminal Server license management and your privacy

# Required licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Required licenses**

You must obtain the necessary licenses as part of your deployment of Terminal Server on your network. A license server that runs a Windows Server 2003 operating system supports the following types of licenses:

- **Windows Server 2003 Terminal Server Device CALs**. These CALs are purchased for devices that connect to a terminal server running a Windows Server 2003 operating system.

- **Windows Server 2003 Terminal Server User CALs**. These CALs are purchased for users who connect to a terminal server running a Windows Server 2003 operating system.

- **Windows Server 2003 Terminal Server External Connector Licenses**. These licenses are purchased to allow unlimited by external users (for example, business partners) to a terminal server running a Windows Server 2003 operating system. Terminal Server External Connector licenses are not managed; therefore, to use these licenses, you must set your terminal server to Per User mode. For instructions on configuring the Terminal Server Licensing mode, see Configure the Terminal Server Licensing mode.

- **Windows 2000 Terminal Services CALs**. These licenses are purchased for devices that connect to a Windows 2000 terminal server.

- **Windows 2000 Terminal Services Internet Connector Licenses**. These licenses are purchased to allow up to 200 simultaneous anonymous connections to a Windows 2000 terminal server by non-employees across the Internet.

- **Windows 2000 Built-in Licenses**. Clients that are running Windows 2000 Professional or Windows XP are issued a license from the built-in pool of licenses when connecting to a Windows 2000 terminal server.

- **Temporary Licenses**. When a Windows Server 2003 family terminal server requests a Windows Server 2003 Per Device CAL, or when a Windows 2000 terminal server requests a Windows 2000 CAL, and the license server has no licenses, it issues a temporary license to the connecting client (if the client device has no license). The license server tracks the issuance and expiration of these licenses. Temporary licenses are designed to allow ample time for you to deploy a license server, and they allow unlicensed clients to connect to the terminal server for 90 days. There is no limit to the number of temporary licenses that a license server can issue, but a single client is only issued a temporary license once. After the temporary license expires, the client can only connect to the terminal server if the license server can issue a permanent CAL, or if the terminal server is still within its licensing grace period.

**Important**

- For a license server to function correctly, the operating system of the license server must be compatible with the operating system of the terminal server. Note the following:

    - A Windows 2000 license server cannot issue CALs to terminal servers running Windows Server 2003 operating systems.

    - A license server running a Windows Server 2003 operating system can issue CALs to terminal servers running Windows Server 2000 or Windows Server 2003, or to terminal servers running Windows 2000 and terminal servers running Windows Server 2003 in a mixed environment

**Notes**

- Verify that you have purchased enough CALs for the Terminal Server license server and that the Terminal Server Licensing mode matches the type of CAL available on the license server. Terminal Server Licensing modes determine the type of CAL that a license server issues to a client. In Windows Server 2003, two types of CALs are available: Per User and

Per Device. A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices. In this case, only one CAL is needed instead of a CAL for each device. Per User CALs, however, are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used.

- By default, when the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

# Terminal Server Licensing administrative tool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Server Licensing administrative tool**
The Terminal Server Licensing administrative tool can help you manage Terminal Server licenses. Using Terminal Server Licensing, you can perform the following tasks after connecting to a Terminal Server license server:

- Activate a Terminal Server license server.
- Install client access licenses (CALs).
- Reactivate a license server.
- Deactivate a license server.
- Repeat the installation of CALs.
- Change the properties of the license server.

To perform these tasks, you must have either Internet access or access to a telephone.

You can also use Terminal Server Licensing to connect to any Terminal Server license server on your network and view information about the licenses on that server. You can view the following Terminal Server licensing information:

- The list of installed client license key packs.
- The total number of licenses in each client license key pack, and the number of available and issued licenses in each key pack.
- The date a license was issued to each computer, and the name of the computer to which each license is issued.
- The issue and expiration dates of each temporary license and the name of the computer to which the license was issued.

When the number of clients requesting licenses from a license server exceeds the number of un-issued licenses stored on the server, you receive a reminder to purchase new licenses and register them. The reminder appears as an event in the system log of the Event Viewer.

For step-by-step procedures on using Terminal Server Licensing, see Terminal Server Licensing How To....

You can also remotely manage a license server from Windows XP computers that have the Windows Server 2003 Administration Tools Pack installed. You can install the Windows Server 2003 Administration Tools Pack, which includes the Terminal Server Licensing tool, on computers that run any of the following operating systems:

- Windows XP Professional with hotfix Q329357
- Windows XP Professional with Service Pack 1 or later
- Windows Server 2003 family

The Windows Server 2003 Administration Tools Pack is included on the Windows Server 2003 family CD. For more information, see Windows Server 2003 Administration Tools Pack Overview.

**Note**

- You cannot manage license servers running Windows Server 2003 from computers that run Windows 2000.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

# Terminal Server license management and your privacy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Server license management and your privacy**

Microsoft asks you to activate your Terminal Server license server and client access licenses (CALs) to verify that you are a licensed user and to provide customized benefits and services. When you provide complete and accurate information, Microsoft customer service representatives can successfully re-create the licensing information on your license server if technical problems occur. This prevents the possible loss of valuable licensing information. The information you provide might also be used internally at Microsoft to perform aggregate quality testing of the Terminal Server Licensing program.

You do not need to provide all of the requested information. However, some fields are required, and if you do not provide complete information, the Microsoft Clearinghouse might be unable to activate your Terminal Server license server and install CALs.

Information provided over the Internet or World Wide Web to the Microsoft Clearinghouse is encrypted for your security and is securely stored by the Microsoft Clearinghouse in a dedicated database. Microsoft uses this information only for the reasons detailed above. This information is kept separate from information you might have provided to Microsoft in other contexts, and is not disclosed to third parties.

To access and make changes to the data provided to the Microsoft Clearinghouse in this license activation process, call your customer service representative. To find the appropriate phone number to call, select the telephone activation method in the Terminal Server License Server Activation Wizard. For information about activating a Terminal Server license server by using the telephone, see Activate a Terminal Server license server by using the telephone.

If you believe that Microsoft has not complied with these privacy principles, please notify us by postal mail at the following address:

Microsoft Corporation

Attn: Terminal Services License Management Program

One Microsoft Way

Redmond, WA 98052-6399

Commercially reasonable efforts are made promptly to determine and correct the problem.

# Using Terminal Server Licensing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Terminal Server Licensing**


- Terminal Server license server roles
- Activating a Terminal Server license server
- Purchasing and installing client access licenses on a Terminal Server license server
- Recovering client access licenses on a Terminal Server license server
- Backing up a Terminal Server license server

# Terminal Server license server roles

Updated: January 19, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Server license server roles**
A Terminal Server license server stores all client licenses. A terminal server must be able to connect to an activated license server before clients can be issued permanent client access licenses (CALs). When you activate a license server, Microsoft provides the server with a digital certificate that validates server ownership and identity. By using this certificate, a license server can make transactions with Microsoft and receive client licenses for your terminal servers. If you install a license server but do not activate it, that server only issues temporary licenses.

For small deployments, it is acceptable to install both the terminal server and Terminal Server Licensing on the same physical computer; however, for larger deployments, consider installing Terminal Server Licensing on a separate server. Note that one Terminal Server license server can serve many terminal servers concurrently.

### Important

- You must configure Terminal Server Licensing correctly in order for your terminal server to continue to accept connections from clients. To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent CAL, or after 120 days, whichever comes first. For more information about the licensing grace period, see Terminal Server Licensing grace period.

Before installing a license server, determine which of the two license server roles you require—a Domain license server or an Enterprise license server. During Terminal Server Licensing setup, you select one of these roles. By default, the Enterprise license server role is selected.

If you are installing a license server in an Active Directory domain environment, you can select either the Enterprise license server or the Domain license server role. If you are installing a license server in a workgroup or a non-Active Directory domain environment, you must select the Domain license server role. To activate the license server quickly, install Terminal Server Licensing on a server that has Internet access.

**Enterprise license server**
You can only install an Enterprise license server on a domain controller or a member server in a domain, not on a stand-alone server. This role is published to Active Directory.

An Enterprise license server is an appropriate choice if your network includes several domains, and if you want to maintain one license server that can issue licenses to terminal servers across the different domains. An Enterprise license server can serve terminal servers in any domain, but the domain must be a Windows Server 2003 domain or Windows 2000 Active Directory domain. An Enterprise license server is automatically discovered by any terminal server that is in the same site as the license server. By default, an Enterprise license server only serves terminal servers in the same site.

**Domain license server**
A Domain license server is an appropriate choice if you want to maintain a separate license server for each domain. Terminal servers can access domain license servers only if they are in the same domain as the license server.

**License server discovery process**
Terminal servers use a discovery process to locate license servers. This process begins when the Terminal Server service starts. Although Terminal Server attempts to detect a license server automatically, you can explicitly specify a preferred license server to which the terminal server connects. For instructions on configuring a preferred license server through Group Policy or by modifying the registry, see Set preferred Terminal Server license servers. You can also set preferred license servers by using a Windows Management Instrumentation (WMI) script. For three scripts that you can use to set preferred license servers, delete preferred license servers, or query preferred license servers, see "Preferred License Server WMI Scripts" in Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

During the discovery process, terminal servers attempt to contact license servers in the following order:

1. Any preferred license servers that are specified in Group Policy.

2. Any Enterprise license servers or Domain license servers that are specified in the registry.

   In Windows Server 2003, the relevant registry key is **LicenseServers** , and it is located under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\TermService\Parameters**.

   In Windows 2000, the relevant registry key is **DefaultLicenseServer**, and it is located under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\TermService\Parameters**.

   In Windows Server 2003, you can specify multiple license servers in the registry. In Windows 2000, you can only specify one license server.

3. Enterprise license servers that are specified in Active Directory.

4. Domain license servers.

**Enterprise license server discovery**
As mentioned in this topic, an Enterprise license server is automatically discovered by any terminal server that is in the same site as the license server and by default, only serves terminal servers in the same site. You can, however, publish the license server to another site in Active Directory. A terminal server polls Active Directory every 60 minutes to locate an Enterprise license server, and continues to poll every 60 minutes, even when an Enterprise license server is located.

**Domain license server discovery**
If you are installing a Domain license server, the computer on which you install the license server must be a domain controller for it to be automatically discovered by terminal servers that communicate with it. If you install a Domain license server on a member server, you must specify this server as a preferred license server on all terminal servers that need to communicate with it. If the computer on which you plan to install the Domain license server is in a workgroup, it can only be automatically discovered by terminal servers on the same subnet.

If you install an Enterprise license server on a member server, you do not need to specify it as a preferred license server, because Enterprise license servers are automatically discovered.

When a terminal server attempts to locate a Domain license server, the terminal server queries all domain controllers within the same site. If the terminal server cannot find an available license server in the same site, the terminal server will not attempt to query any other domain controllers (for example, other domain controllers in the same domain). A terminal server attempts to locate a Domain license server every 15 minutes. After a Domain license server is located, the terminal server continues to search for a Domain license server every two hours. If the terminal server cannot locate a Domain license server, it resumes searching every 15 minutes.

## Notes

- To improve security, you should consider removing workgroups and replacing them with domains.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

# Activating a Terminal Server license server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Activating a Terminal Server license server**

You must activate a Terminal Server license server and then purchase and install the appropriate number of client access licenses (CALs) before the license server can issue permanent CALs to Terminal Server clients. A license server that has been installed but not activated only issues temporary licenses. These temporary licenses allow clients to connect to the terminal server for 90 days. For information about purchasing and installing CALs, see Purchasing and installing client access licenses on a Terminal Server license server.

To activate a license server, use Terminal Server Licensing. When you activate a license server, the Microsoft Clearinghouse provides the server with an X.509 digital certificate that validates server ownership and identity. Using this certificate, a license server can make subsequent transactions with the Microsoft Clearinghouse and receive permanent CALs for your terminal servers.

Information that you provide over the Internet or World Wide Web to the Microsoft Clearinghouse is encrypted for your security and is securely stored by the Microsoft Clearinghouse in a dedicated database. This information is kept separate from information you might have provided to Microsoft in other contexts, and is not disclosed to third parties. For more information, see Terminal Server license management and your privacy.

**Important**

- If the Internet Explorer Enhanced Security Configuration setting is in effect on your Terminal Server license server, you must add the Microsoft Clearinghouse Web site (https://activate.microsoft.com/) to the Trusted sites zone in Internet Explorer before you can activate your Terminal Server license server. If you try to connect to the Microsoft Clearinghouse Web site and it is not yet a trusted site, you are prompted to specify whether to add it to the Trusted sites zone. For more information, see Internet Explorer Enhanced Security Configuration.

The following table shows the three ways to activate a Terminal Server license server.

| Automatic | Web Browser | Telephone |
|---|---|---|
| This is the fastest and easiest method. To use this method, the server running Terminal Server Licensing must have an Internet connection. | Use this method to connect directly to the secure Microsoft Terminal Services Web site. This method also requires an Internet connection, but the connection can be made from any computer, not just the computer that is running Terminal Server Licensing. | If you do not have an Internet connection, use this method to contact the Microsoft Clearinghouse and receive an ID number to activate the license server by phone. For instructions on how to locate the appropriate local number to call, see Locate the Microsoft Clearinghouse telephone number for your country or region. |

For instructions on activating a license server by using any of these methods, see Activate a Terminal Server License Server.

You are required to activate a license server only once, after which the license server becomes the repository for terminal server client licenses.

**Notes**

- If you upgrade a Terminal Server license server from Windows 2000, you must reactivate it.

- The Microsoft Clearinghouse is the facility that Microsoft maintains to activate license servers, issue CALs to license servers, recover CALs, and reactivate license servers. You cannot purchase CALs from the Clearinghouse. For instructions on how to purchase CALs, see Purchase client access licenses. After you purchase CALs, make sure to retain the License Purchase Agreement. If you contact the Clearinghouse, the information in the Agreement will help them assist you.

- To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. For more information about the licensing grace period, see Terminal Server Licensing grace period.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

# Purchasing and installing client access licenses on a Terminal Server license server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Purchasing and installing client access licenses on a Terminal Server license server**
After a Terminal Server license server is activated, you can purchase and then install client access licenses (CALs) on that server so that it can issue those licenses to terminal servers upon request. The CALs are stored and tracked on the license server.

For instructions on purchasing CALs, see Purchase client access licenses. After you purchase CALs, make sure to retain the License Purchase Agreement information. Having this information with you will facilitate communication with the Microsoft Clearinghouse, should you need assistance with recovering CALs or other Terminal Server Licensing tasks that are supported by the Clearinghouse. The Clearinghouse is the facility that Microsoft maintains to activate license servers, issue CALs to license servers, recover CALs, and reactivate license servers.

The Clearinghouse processes your request to install and issue CALs, in the quantity requested, to your license server. Then, you install the CALs, using one of the three methods you used to activate a license server—Automatic (recommended), Web Browser, and Telephone. For instructions on installing CALs by using any of these methods, see Install Client Access Licenses. For instructions on locating the Clearinghouse telephone number if you do not have access to the Internet, see Locate the Microsoft Clearinghouse telephone number for your country or region.

Client licenses are digitally signed certificates that each client stores locally. After client licenses are installed on a license server, the server can issue licenses. When a client computer attempts to log on to a terminal server for the first time, the terminal server recognizes that the client has not been issued a license and locates a license server to issue a new license to the client. Terminal Server Licensing modes determine the type of CAL that a license server issues to a client. In Windows Server 2003, two types of CALs are available: Per User and Per Device. A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices. In this case, only one CAL is needed instead of a CAL for each device.

By default, when the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.

If you want to transfer a license from one computer to another, you must call the Microsoft Clearinghouse. According to the terms of your End User License Agreement (EULA) for Terminal Server Client Access Licenses, you can make this transfer only once.

### Notes

- You must configure Terminal Server Licensing correctly in order for your terminal server to continue to accept connections from clients. To allow ample time for you to deploy a license server, Terminal Server provides a licensing grace period, during which no license server is required. During this grace period, a terminal server can accept connections from unlicensed clients without contacting a license server. The grace period begins the first time the terminal server accepts a client connection. It ends after you deploy a license server and that license server issues its first permanent CAL, or after 120 days, whichever comes first. For more information about the licensing grace period, see Terminal Server Licensing grace period.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

# Recovering client access licenses on a Terminal Server license server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovering client access licenses on a Terminal Server license server**

After a Terminal Server license server is activated, you can purchase and then install client access licenses (CALs) on that server so that it can issue those licenses to terminal servers upon request. The terminal servers, in turn, issue licenses to clients. Terminal Server licensing modes determine the type of CAL that a license server issues. In Windows Server 2003, two types of CALs are available: Per User and Per Device. A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices. In this case, only one CAL is needed instead of a CAL for each device.

To configure the Terminal Server licensing mode, you must set the licensing mode on the terminal server to match the type of CAL that you purchase and install on the license server. To set the licensing mode on the terminal server, you can use Terminal Services Configuration. For more information, see Configure the Terminal Server Licensing mode.

When the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license by default. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.

By default, for terminal servers and license servers that are running Windows Server 2003 or Windows 2000 Service Pack 3 (or that are running Windows 2000 Service Pack 2 and earlier with the hotfix described in Microsoft Security Bulletin MS01-052 applied), an expiration period is applied for each Per Device CAL that is issued. This expiration period, which is not configurable, is a random number of days between 52 to 89 days after the license is issued. When a client connects to a terminal server, the expiration date for the CAL is checked. If the expiration is set to occur within seven days, the terminal server connects to the license server and renews the CAL. During this time, the CAL is given another expiration period of 52 to 89 days. If the license server is not available, the terminal server attempts to replace the CAL each time the client logs on to a terminal server session. After the CAL expires, the license server returns any CAL that was not renewed to the group of available CALs.

This functionality facilitates the automatic recovery of Per Device CALs that are lost due to hardware failure, operating system reinstallation, and other similar events. This functionality is also available for terminal servers and license servers running earlier versions of Windows 2000 that have the hotfix described in Microsoft Security Bulletin MS01-052 applied. For the hotfix to work correctly, you must apply it to all terminal servers and license servers in your organization.

**Important**

- CALs cannot be recovered automatically in the following cases:

    - CALs that are issued by license servers and terminal servers that are running Windows 2000 Service Pack 2 and earlier and that do not have the hotfix described in Microsoft Security Bulletin MS01-052 applied. In such cases, these CALs remain assigned to the devices to which they are issued. If the CALs are lost, you must contact the Microsoft Clearinghouse to recover them. If you do need to contact the Clearinghouse to recover CALs, make sure that you have the License Purchase Agreement information available. Having this information with you will facilitate the CAL recovery process. For information about how to contact the Clearinghouse, see Locate the Microsoft Clearinghouse telephone number for your country or region.

    - Catastrophic failure that results in the loss of the licensing database. In the event of a failure that results in the loss of the licensing database when a known good backup is not available, you must reinstall Terminal Server Licensing and reactivate the license server, and then contact the Clearinghouse to re-issue any previously issued CALs. The original CALs that were issued are based on the identification number of the license server at the time of issuance. A license server ID number is a unique 35-character alphanumeric representation of the certificate of a license server. This certificate is provided during license server activation, and it is used to obtain CALs. If the license server ID number changes, CALs that are based on the old license server ID number cannot be re-installed.

A re-issued set of CALs is displayed in the Terminal Services Licensing administrative tool as another CAL installation. Other than being a separate CAL installation, re-issued CALs are not distinguishable from an original installation.

The following problems might result from a Terminal Server licensing mode mismatch:

- If a terminal server is set to Per Device mode and the license server with which it communicates only has Per User CALs, the license server only issues temporary licenses. When the temporary licenses are close to expiration, the following message appears on the terminal server: "Your terminal services temporary client license will expire in N days. Please contact your system administrator to get a permanent license." In this case, Event ID 26, Source: Application Pop-up appears in the application event log.

- Similarly, Terminal Server clients that run Windows XP or Windows 2000 might receive the following error message, if the client cannot connect to the terminal server for this reason: "The remote session was disconnected because the local computer client access license could not be upgraded or renewed." In this case, Event ID 1011, Source TermService appears in the application event log on the terminal server.

To correct these problems, switch the Terminal Server Licensing mode from Per Device to Per User. For information about switching modes, see Configure the Terminal Server Licensing mode.

Keep in mind that even though Per User CALs are not monitored by Terminal Server, this does not remove administrators from End User License Agreement (EULA) requirements to have a valid terminal server (TS) CAL for each user. Failure to have a valid Per User CAL for each user, if Per Device CALs are not being used, is a violation of the EULA.

The following text is from the EULA for Windows Server 2003.

Two different TS CALs are available to you: "Device" and "User." Each TS Device CAL permits one Device (used by any User) to conduct Windows Sessions on any of your Servers. Each TS User CAL permits one User (using any Device) to conduct Windows Sessions on any of your Servers. You may use a mix of TS Device CALs and TS User CALs simultaneously with the Server Software in your environment. You can have a Terminal Server request Per User licenses or Per Device (default) but not both simultaneously.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

# Backing up a Terminal Server license server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backing up a Terminal Server license server**

You should back up your Terminal Server license server regularly by using the Backup tool. This will help protect your licensing data from accidental loss if your system experiences hardware or storage failure.

When backing up a Terminal Server license server, back up the System State and Repair data, and the %systemdrive%\system32\lserver directory. Doing this ensures that data in both the registry and the license server database is backed up.

If you restore the System State data and the database to the same license server computer, any un-issued licenses are restored correctly as long as you have not replaced the operating system on the computer or removed and reinstalled Terminal Server Licensing. If you do not restore to the same computer, any un-issued licenses that are detected are not restored and an event will appear in the system log of Event Viewer. You can still restore the un-issued licenses by contacting the Microsoft Clearinghouse by phone and requesting that the licenses be reissued. When you contact the Clearinghouse, make sure that you have the License Purchase Agreement available. The information in the Agreement will facilitate the process of reissuing the licenses. For information about how to contact the Clearinghouse, see Locate the Microsoft Clearinghouse telephone number for your country or region.

For more information about using Backup, see Backup.

For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

# Terminal Services Configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services Configuration**

You can configure new connections for Terminal Services, modify the settings of existing connections, and delete connections by using the Terminal Services Configuration tool (TSCC.msc) or Group Policy (gpedit.msc).

- For help with specific tasks, see Terminal Services Configuration How To....
- For general background information, see Terminal Services Configuration Concepts.

# Terminal Services Configuration How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install Terminal Server
- Configure Terminal Services Connections
- Manage Terminal Services Connections
- Configure Server Settings
- Configure Group Policy Settings

# Install Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Terminal Server**

1. Open Configure Your Server Wizard.
2. Click **Next** and follow the instructions on the **Preliminary Steps** screen.
3. Click **Next**, and on the **Server Role** screen, select **Terminal Server**.
4. Click **Next** and follow the instructions in the wizard.
5. Some programs might not work correctly when Terminal Server is installed. If so, you must reinstall these programs for multisession access using **Add or Remove Programs** after Terminal Server is installed. For more information, see Program Compatibility.

**Warning**

- Terminal Server Licensing is a required component that licenses clients on a terminal server. You must install Terminal Server Licensing or your terminal server will stop accepting connections from unlicensed clients 120 days from the date of the first client logon.

  You should not install Terminal Server Licensing on the same computer with Terminal Server.

  For more information on Terminal Server Licensing, see Terminal Server Licensing.

**Important**

- When you install Terminal Server using the Configure Your Server Wizard, as a security precaution the Internet Explorer Enhanced Security Configuration is set by default to prevent Administrators from downloading files or running ActiveX controls from the Web. If you want to change this setting to either permit all users to download files or run ActiveX controls or prevent all users from doing so, use **Add/Remove Windows Components** in **Add or Remove Programs**. This setting applies to servers running one of the Windows Server 2003 family operating systems. For more information, see Terminal server role: Configuring a terminal server and Internet Explorer Enhanced Security Configuration.

**Important**

- You must use the Remote Desktop Users group to control remote connection permissions for Terminal Server. For more information, see Managing Terminal Services users and Add users to the Remote Desktop Users group.

**Important**

- It is not necessary to install Terminal Server for remote administration of your server. Remote Desktop for Administration is installed by default. To use Remote Desktop for Administration, you must first enable remote connections. For more information on enabling remote connections, see Enable or disable Remote Desktop and Enabling users to connect remotely to the server.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For

more information, see Default local groups, Default groups, and Using Run as.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.

- You can also use **Add or Remove Programs** to install Terminal Server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Terminal Server Licensing
Activate a Terminal Server License Server
Install Client Access Licenses
Enable or disable Remote Desktop

# Configure Terminal Services Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Terminal Services connections**

- Make a new connection
- Specify a maximum number of sessions that can connect to the server
- Configure Authentication and Encryption
- Specify a program to start automatically when a user logs on
- Allow a user to log on to the server automatically
- Manage Remote Control
- Configure Time-out and Reconnection Settings
- Manage Client Settings
- Manage Connection Permissions

# Make a new connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a new connection**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. On the **Action** menu, click **Create New Connection**.
4. The Terminal Services Connection wizard appears. Follow the instructions on your screen to configure the connection.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
- Only one Remote Desktop Protocol (RDP) connection can be configured for each network adapter in a terminal server. Typically, the RDP connection that is configured when Terminal Services is installed is the only connection needed. You can configure additional connections using RDP if you install an additional network adapter for each connection.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring Terminal Services with TSCC

# Specify a maximum number of sessions that can connect to the server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify a maximum number of sessions that can connect to the server**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Limit number of connections** setting, and then click **Enabled**.
3. In the **TS Maximum Connections allowed** box, type the maximum number of connections you want to allow, and then click **OK**.

### Important

- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

### Note

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy by using the **Active Directory Users and Computers** snap-in.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection for which you want to specify a maximum number of sessions, and then click **Properties**.
4. On the **Network Adapter** tab, click **Maximum connections**, type the maximum number of sessions that can connect to the server, and then click **Apply**.

### Notes

- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.
- By default, the connection is configured to allow an unlimited number of sessions to connect to the server.
- Restricting the number of sessions improves performance because fewer sessions are demanding system resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring Terminal Services with TSCC
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Configure Authentication and Encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Authentication and Encryption**

Request a computer certificate for server authentication

Configure authentication and encryption on the server

Request a certification authority certificate for the client

Configure authentication on the client by using Remote Desktop Connection

Configure authentication on the client by using an .rdp file

# Request a computer certificate for server authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To request a computer certificate for server authentication**

Submit an advanced certificate request for a computer certificate by using the Web

Submit a computer certificate request by using the Certificate Request Wizard

**Submit an advanced certificate request for a computer certificate by using the Web**

1. Open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the Web server that is running Windows Server 2003 and where the certification authority (CA) that you want to access is located.
3. Click **Request a certificate**.
4. Click **Advanced certificate request**.
5. Click **Create and submit a certificate request to this CA**.
6. Provide identifying information as required.
7. In the **Name** box, type the name of the terminal server for which you are requesting the certificate.
8. In **Type of Certificate needed**, click **Server Authentication Certificate**.
9. In **Key Options**, click **Create a new key set**, and then specify additional options as follows:
   - In **Cryptographic service provider (CSP)**, click **Microsoft RSA SChannel Cryptographic Provider**. This CSP supports key derivation for the SSL2, PCT1, SSL3, and TLS1 protocols.
   - In **Key usage**, click **Exchange**. Exchange means that the private key can be used to enable the exchange of sensitive information.
   - Select the **Mark Keys as exportable** check box. Doing this saves the public and private key to a PKCS #12 file. This is useful if you want to copy a certificate for use on another computer.
10. Select the **Store certificate in the local computer certificate store** check box.

    **Important**

    - This check box must be selected for TLS authentication to function.
11. Click **Submit**.
12. If you see the **Certificate Issued** Web page, click **Install this certificate**. For instructions on how to retrieve a CA certificate from a Windows Server 2003 CA, see Related Topics.
13. If you are finished using the Certificate Services Web pages, close Internet Explorer.

**Notes**

- To request a certificate, a subject must have Read and Enroll permissions on the certificate template.
- To open Internet Explorer, click **Start**, point to **All Programs**, and then click **Internet Explorer**.
- This procedure is for obtaining certificates from a Windows Server 2003 CA. For instructions on how to obtain certificates by using the Windows 2000 Certificate Services Web pages, see Related Topics. Alternatively, you can purchase a certificate from a non-Microsoft vendor and install the certificate manually.
- Use this procedure if you need to request the following certificates:
    - Certificates that are issued from a stand-alone CA.
    - Certificates that are based on a template configured to obtain the subject name from the subject.
    - Certificate types that require approval before the certificates are issued.
- You can use the second procedure in this topic, "Submit a computer certificate request by using the Certificate Request Wizard," if the certificates meet the following requirements:
    - The certificates are issued from an enterprise CA.
    - The certificates are based on templates where the subject name is generated by Windows.
    - The certificate type does not require approval before the certificates are issued.
- For a client to make a connection to the terminal server, the client must trust the root of the server's certificate. The client computers must have the certificate of the root CA that issued the server certificate in their Trusted Root Certification Authorities store, which you can view in the Certificates snap-in.

**Submit a computer certificate request by using the Certificate Request Wizard**

1. If you are submitting a computer certificate request for the local terminal server, open Certificates (Local Computer) If you are submitting a computer certificate request for a remote terminal server, open Certificates (*Servername)*.
2. In the console tree, click one of the following:
    - **Certificates (Local Computer)** (for the local terminal server)
    - **Certificates (*Servername*)** (for a remote terminal server).
3. If you are not in Certificate Purpose view mode, on the **View** menu, click **Options**, click **Certificate Purpose, and then click OK**.
4. In the details pane, under **Intended Purposes**, double-click **Server Authentication**.
5. On the **Action** menu, point to **All Tasks**, and then click **Request New Certificate** to start the Certificate Request Wizard.
6. On the **Certificate Types** page, click **Server Authentication**, and then select the **Advanced** check box.
7. On the **Cryptographic Service Provider** page, click **Microsoft RSA SChannel Cryptographic Provider**. This CSP supports key derivation for the SSL2, PCT1, SSL3, and TLS1 protocols. '
8. Select the key length. It is recommended that you use a key length of 1024 bits (the default setting) or higher.

9. Select the **Mark Keys as exportable** check box. Doing this enables you to save the public and private key to a PKCS #12 file. This is useful if you want to copy a certificate for use on another computer.

10. **Specify whether to enable strong private key protection**.

11. On the **Certification authority** page, if more than one CA is available, click **Browse**, and then click the name of the CA that will issue the certificate.

12. On the **Certificate Friendly Name and Description** page, type a friendly name for the new certificate.

13. On the wizard completion page, click **Finish**.

### Notes

- To request a certificate, a subject must have Read and Enroll permissions on the certificate template.

- To open Certificates for a computer account, click **Start**, click **Run**, type **mmc**, and then click **OK**. On the **File** menu, click **Open**, navigate to the console file that you want to open, and then click **Open**. In the console tree, click the **Certificates snap-in that contains the computer account you want**.

- If you have not already created an MMC console that contains Certificates, see Related Topics for instructions on managing certificates for a computer.

- This procedure is for obtaining certificates from a Windows Server 2003 CA. However, you can also use the Windows Certificate Request Wizard to obtain certificates from a Windows 2000 CA, or you can purchase a certificate from a non-Microsoft vendor and install the certificate manually.

- To use this procedure to request certificates, the certificates must meet the following requirements:

    - The certificates are issued from an enterprise CA.

    - The certificates are based on templates where the subject name is generated by Windows.

    - The certificate type does not require approval before the certificates are issued.

- Use the first procedure in the topic, "Submit an advanced certificate request for a computer certificate via the Web," if you need to request the following certificates:

    - Certificates that are issued from a stand-alone CA.

    - Certificates that are based on a template configured to obtain the subject name from the subject.

    - The certificate type requires approval before the certificates are issued.

- For a client to make a connection to the terminal server, the client must trust the root of the server's certificate. The client computers must have the certificate of the root CA that issued the server certificate in their Trusted Root Certification Authorities store, which you can view in the Certificates snap-in.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

#### See Also

**Concepts**

Working with MMC console files
Manage certificates for a computer
Requesting certificates
Retrieve a certification authority certificate from a Windows Server 2003 CA
Use Windows 2000 Certificate Services Web Pages
Configuring Terminal Services with Group Policy
Configuring Terminal Services with TSCC

# Configure authentication and encryption on the server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure authentication and encryption on the server**

- Configure the encryption level by using Group Policies
- Using Terminal Services Configuration

**Configure the encryption level by using Group Policies**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Encryption and Security**, double-click the **Set client connection encryption level** setting, and then click **Enabled**.
3. To set the encryption level, do one of the following:
   - To set the encryption level to **Client Compatible**, **High Level**, or **Low Level**, in the **Encryption Level** list, click the level that you want, and then click **OK**. For information about these encryption levels, see Note, at this end of this topic.
   - To enable FIPS compliant encryption, click **OK** to close the **Set client connection encryption level Properties** dialog box, and then navigate to **Computer Configuration**, **Windows Settings**, **Security Settings**, **Local Policies**, **Security Options**. Double-click **System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing**, and then click **Enabled**.
4. If you set the encryption level to **High Level** or if you enabled FIPS compliant encryption and you want to use Transport Level Security (TLS) 1.0 to authenticate the server, you must enable TLS by using the Terminal Services Configuration tool and meet additional configuration requirements.

   You cannot use Group Policy to enable TLS authentication. For more information, see Using Terminal Services Configuration, later in this topic.

**Important**

- Any encryption level settings that you configure in Group Policy override the configuration that you set by using the Terminal Services Configuration tool. Also, if you enable the System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing Group Policy setting, this setting overrides the **Set client connection encryption level** Group Policy setting.
- In order for clients to be able to connect to a terminal server that uses FIPS compliant encryption, you must upgrade these clients to use the RDP 5.2 (Windows Server 2003) client. You can install this client from Windows Server 2003 terminal servers. For more information, see Remote Desktop Connection for Windows Server 2003 [5.2.3790] (http://go.microsoft.com/fwlink?/LinkID=41068).
- Use this procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an administrator. Then, you must open Group Policy by using the **Active Directory Users and Computers** snap-in.
- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.

2. In the console tree, click **Connections**.

3. In the details pane, right-click the connection you want to modify, and then click **Properties**.

4. On the **General** tab, in **Security layer**, select a security method. The security method that you select determines whether the terminal server is authenticated to the client, and the level of encryption that you can use. You can select from these security methods.

   - The **Negotiate** method uses TLS 1.0 to authenticate the server, if TLS is supported. If TLS is not supported, the server is not authenticated.

   - The **RDP Security Layer** method uses native Remote Desktop Protocol encryption to secure communications between the client and server. If you select this setting, the server is not authenticated.

   - The **SSL** method requires the use of TLS 1.0 to authenticate the server. If TLS is not supported, the connection fails. This method is only available if you select a valid certificate, as described in Step 6.

   If you select **Negotiate** or **SSL**, for TLS to function correctly, you must also set the encryption level to **High**, or you must enable FIPS compliant encryption by using Group Policy or Terminal Server Configuration. Additional server and client configuration requirements must also be met. For more information about requirements and tasks for configuring Terminal Server to support TLS authentication, see Configuring authentication and encryption.

5. In **Encryption level**, click the level that you want. You can select **Low**, **Client Compatible**, **High**, **or FIPS Compliant**. Fo**r more information about these levels, see Notes, at the end of this topic.**

6. **To use TLS 1.0 to authenticate the server, in Certificate, click Browse, click Select Certificate, and then click the certificate that you want to use. The certificate must be an X.509 certificate with a corresponding private key. For instructions on how to verify whether the certificate has a corresponding private key, see Notes.**

7. To specify that clients log on to the terminal server by typing their credentials in the default Windows logon dialog box, select the **Use standard Windows logon interface** check box.

**Note**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

- Any encryption level settings that you configure in Group Policy override the configuration that you set by using the Terminal Services Configuration tool. Also, if you enable the System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing Group Policy setting, this setting overrides the **Set client connection encryption level** Group Policy setting.

- When you change the encryption level, the new encryption level takes effect the next time a user logs on. If you require

multiple levels of encryption on one server, install multiple network adapters and configure each adapter separately.

- To verify that certificate has a corresponding private key, in Terminal Services Configuration, right-click the connection for which you want to view the certificate, click the **General** tab, click **Edit**, click the certificate that you want to view, and then click **View Certificate**. At the bottom of the **General** tab, the statement, "You have a private key that corresponds to this certificate" should appear. You can also view this information by using the Certificates snap-in.

- The **FIPS compliant** setting (the **System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing** setting in Group Policy or the **FIPS Compliant** setting in Terminal Server Configuration) encrypts and decrypts data sent from the client to the server and from the server to the client, with the Federal Information Processing Standard (FIPS) 140-1 encryption algorithms, using Microsoft cryptographic modules. For more information, see FIPS 140 Evaluation (http://go.microsoft.com/fwlink/?LinkID=34627).

- The **High** setting encrypts data sent from the client to the server and from the server to the client by using strong 128-bit encryption.

- The **Client Compatible** setting encrypts data sent between the client and the server at the maximum key strength supported by the client.

- The **Low** setting encrypts data sent from the client to the server using 56-bit encryption.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring authentication and encryption
Request a computer certificate for server authentication
Request a certification authority certificate for the client
Configure authentication on the client by using Remote Desktop Connection
Configuring Terminal Services with Group Policy

# Request a certification authority certificate for the client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To request a certification authority certificate for the client**

1. Open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the Web server running Windows Server 2003 where the certification authority (CA) you want to access is located.
3. Click **Download a CA certificate**, **certificate chain**, or **CRL**.
4. If you want to trust all the certificates issued by this CA, click **Install this CA certificate chain**. A certificate chain contains the chain of certificates that can be traced back to the root CA.
5. If you are finished using the Certificate Services Web pages, close Internet Explorer.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Internet Explorer, click **Start**, point to **All Programs**, and then click **Internet Explorer**.
- Use this procedure to ensure that clients trust the root of the server's certificate (that is, clients have the certificate of the root CA that issued the server certificate in their Trusted Root Certification Authorities store). This requirement must be met for Transport Level Security (TLS) to be used for server authentication when clients connect to a terminal server. For more information about client and server requirements and tasks for configuring Terminal Server to support TLS authentication, see Configuring authentication and encryption.
- Choosing **Install this CA certification path** is useful when you want to establish trust in a subordinate CA, but you do not currently have the certificate of the root CA in your certificate store.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring authentication and encryption
Configure authentication on the client by using Remote Desktop Connection
Request a computer certificate for server authentication
Configure authentication and encryption on the server
Requesting certificates

# Configure authentication on the client by using Remote Desktop Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure authentication on the client by using Remote Desktop Connection**

1. Open Remote Desktop Connection.
2. Click **Options**.
3. On the **Security** tab, in the **Authentication** list, click the authentication option that you want. You can select from the following options:

- **No authentication**. If you select this option, the server is not authenticated.
- **Attempt authentication.** If you select this option, Transport Level Security (TLS) 1.0 is used to authenticate the terminal server, if TLS 1.0 is supported and correctly configured.
- **Require authentication.** If you select this option, TLS is required to authenticate the server. If TLS is not supported or correctly configured, the connection fails. This option is only available for clients that are connecting to terminal servers running Windows Server 2003 Service Pack 1 (SP1).

**Important**

- If you plan to use TLS authentication, you must configure clients and servers correctly to support this option. Clients must:
    - Run Windows 2000 or Windows XP.
    - Use the RDP 5.2 (Windows Server 2003) client. You can install this client from Windows Server 2003 terminal servers. For more information, see Remote Desktop Connection for Windows Server 2003 [5.2.3790] (http://go.microsoft.com/fwlink?/LinkID=41068).
    - Trust the root of the server's certificate.
- For more information about client and server requirements and tasks for configuring Terminal Server to support TLS authentication, see Configuring authentication and encryption.

**Note**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Remote Desktop Connection, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.
- If you select **Attempt authentication** and the server certificate has expired or is not issued by a trusted root CA, or if the name of the certificate does not match the name of the client computer, you can choose to continue the connection without TLS server authentication. Other errors, however, will cause the connection to fail.
- Support for TLS 1.0 authentication can also be enabled on client computers through configuration and distribution of an .rdp file. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring authentication and encryption
Request a certification authority certificate for the client
Configure authentication on the client by using an .rdp file
Request a computer certificate for server authentication
Configure authentication and encryption on the server
Requesting certificates

# Configure authentication on the client by using an .rdp file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure authentication on the client by using an .rdp file**

1. Open Remote Desktop Connection.
2. Click **Options**.
3. On the **Security** tab, in the **Authentication** list, click the authentication option that you want. You can select from the following options:
   - **No authentication**. If you select this option, the server is not authenticated.
   - **Attempt authentication.** If you select this option, Transport Level Security (TLS) 1.0 is used to authenticate the terminal server, if TLS 1.0 is supported and correctly configured.
   - **Require authentication.** If you select this option, TLS is required to authenticate the server. If TLS is not supported or correctly configured, the connection fails. This option is only available for clients that are connecting to terminal servers running Windows Server 2003 Service Pack 1 (SP1).
4. On the **General** tab, click **Save As**.
5. Enter a file name for the saved connection file, and then click **Save**. Connections are saved as Remote Desktop protocol (.rdp) files.
6. Distribute the .rdp file to client computers.

**Important**

- If you plan to use TLS authentication, you must configure clients and servers correctly to support this option. Clients must:
  - Run Windows 2000 or Windows XP.
  - Use the RDP 5.2 (Windows Server 2003) client. You can install this client from Windows Server 2003 terminal servers. For more information, see Remote Desktop Connection for Windows Server 2003 [5.2.3790] (http://go.microsoft.com/fwlink?/LinkID=41068).
  - Trust the root of the server's certificate.
- For more information about client and server requirements and tasks for configuring Terminal Server to support TLS authentication, see Configuring authentication and encryption.

**Note**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Remote Desktop Connection, click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, point to **Communications**, and then click **Remote Desktop Connection**.

- You can also edit the .rdp file by using Microsoft Notepad, where the parameter is **authentication level**. You can specify a value of 0, 1, or 2 where:

  - 0 = No authentication
  - 1 = Require authentication
  - 2 = Attempt authentication

  For example, you can enter the following:

  **authentication level:i:0**

  Where**i**specifies that the value is an integer, and **0** specifies the authentication level (in this example, no authentication).

- If you select **Attempt authentication** and the server certificate has expired or is not issued by a trusted root CA, or if the name of the certificate does not match the name of the client computer, you can choose to continue the connection without TLS server authentication. Other errors, however, will cause the connection to fail.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring authentication and encryption
Request a certification authority certificate for the client
Configure authentication on the client by using Remote Desktop Connection
Request a computer certificate for server authentication
Configure authentication and encryption on the server
Requesting certificates

# Specify a program to start automatically when a user logs on

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify a program to start when a user logs on**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration, Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Start a program on connection** setting, and then click **Enabled**.
   - If the status is to **Not Configured**, then Terminal Services cannot start a specified initial program automatically when a client connects to a terminal server, unless an initial program is specified in the default user profile and in Remote Desktop Connection or the Terminal Services client. If an initial program is not specified, then the user must start programs by using the default desktop that is displayed during the Terminal Services session.
   - If the status is set to **Disabled**, then Terminal Services cannot start a specified initial program automatically when a client connects to a terminal server. Instead, the user must start programs by using the default desktop that is displayed during the Terminal Services session.
3. If you selected **Enabled**, in the **Program path and file name** box, type the path and file name of the program you want to start when a user logs on.
4. In the **Working Directory** box, type the working directory, and then click **OK**.

### Important

- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

### Notes

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an administrator. Then, you must start Group Policy by using the **Active Directory Users and Computers** snap-in.
- These policies affect every client that connects to the terminal server. To specify a program to start on a per-user basis, use the corresponding policy under **User Configuration**.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections.**
3. In the details pane, right-click the connection for which you want to specify an initial program, and then click **Properties**.

4. On the **Environment** tab, under **Initial program**, select **Start the following program when the user logs on**. This option allows you to configure an initial program for the connection.

    - If you select **Do not allow an initial program to be launched. Always show desktop**, then Terminal Services cannot start a specified initial program automatically when a client connects to a terminal server. Instead, the user must start programs by using the default desktop that is displayed during the Terminal Services session.

    - If you select **Run initial program specified by user profile and Remote Desktop Connection or Terminal Services client**, then the program that is specified in the default user profile and in Remote Desktop Connection or the Terminal Services client will run when the client connects to the terminal server.

5. If you selected **Start the following program when the user logs on**, in **Program path and file name**, type the path and file name of the program that you want to start when the user logs on to the terminal server.

6. In **Start in**, type the working directory path for the program, and then click **OK**.

### Notes

- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

- Some client option settings are not applied when a connection is made to the console session of a remote computer running Windows XP. For example, if the client has specified a program to start upon connection to a remote computer, that program does not start if the connection is made to the console session of the remote computer. Instead, the default desktop appears. If the connection is made to a session other than the console session, the specified program starts as expected. The console session is the only session available for Remote Desktop on Windows XP Professional, so the default desktop always appears when a connection is made, even if the client has specified a different program.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring Terminal Services with TSCC
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Allow a user to log on to the server automatically

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow a user to log on to the server automatically**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection through which you want to allow a user to log on automatically, and then click **Properties**.
4. On the **Logon Settings** tab, click **Always use the following logon information**.

   This setting enables you to configure logon information for a user. However, you should be aware that when all users are logged in under the same account, it is more difficult to track down a user who might be causing trouble.
5. In **User name**, type the name of the user you want to allow to log on to the server automatically.
6. In **Domain**, type the name of the domain to which the user's computer belongs.
7. In both **Password** and **Confirm password**, type the user's password.
8. Click **Always prompt for password** to specify that the user is prompted for a password before logging on to the server.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files

# Manage Remote Control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage remote control**

- Configure remote control settings
- Disable remote control

# Configure remote control settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure remote control settings**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Set rules for remote control of Terminal Services user sessions** setting, and then click **Enabled**.
3. In the **Options** box, click the desired remote control permissions, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.
- These policies affect every client that connects to the terminal server. To configure Remote Control on a per-user basis, use the corresponding policy under **User Configuration**.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection for which you want to configure remote control, and then click **Properties**.
4. On the **Remote Control** tab, click Use remote control with the following settings to configure remote control for the connection.
5. To display a message on the client, asking permission to view or take part in the session, select the **Require user's permission** check box.
6. Under **Level of control**, click **View** the session to specify that the user's session can be viewed only, or click **Interact with the session** to specify that the user's session can be actively controlled with your keyboard and mouse.

   **Note**

   - To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and

then double-click **Terminal Services Configuration**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

- Configuring remote control on a per-connection basis affects all sessions that use the connection. You can configure remote control on a per-user basis using the Terminal Services User Configuration Group Policies, or by using the Terminal Services Extension to Local Users and Groups and Active Directory Users and Computers.

- Changes to remote control permissions are not applied to sessions that are connected when the change is made.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Using remote control
Disable remote control
Remotely control a session
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Disable remote control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable remote control**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Set rules for remote control of Terminal Services user sessions** setting, and then click **Enabled**.
3. In the **Options** box, click "No remote control allowed", and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

**Notes**

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.
- These policies affect every client that connects to the terminal server. To configure Remote Control on a per-user basis, use the corresponding policy under **User Configuration**.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection for which you want to disable remote control, and then click **Properties**.
4. On the **Remote Control** tab, click **Do not allow remote control**.

   **Note**

   - To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For

more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Using remote control
Configure remote control settings
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Configure Time-out and Reconnection Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure time-out and reconnection settings**

- Set time-out settings for disconnected, active, and idle sessions
- Change settings for when a session limit is reached or connection is broken
- Change settings for reconnecting disconnected Citrix ICA sessions

# Set time-out settings for disconnected, active, and idle sessions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set time-out settings for disconnected, active, and idle sessions**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Sessions**, enable the settings you want to configure. See the policy explain text for more information.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.
- These policies affect every client that connects to the terminal server. To define Session settings on a per-user basis, use the Sessions policies under **User Configuration**.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection for which you want to modify time-out settings, and then click **Properties**.
4. On the **Sessions** tab, above **End a disconnected session**, select the **Override user settings** check box. This allows you to configure time-out settings for the connection.
5. Configure the following time-out settings as appropriate:

   - In **End a disconnected session**, select the maximum amount of time that a disconnected session remains on the server. When the time limit is reached, the disconnected session ends. When a session ends, it is permanently deleted from the server. Select **Never** to allow disconnected sessions to remain on the server indefinitely.
   - In **Active session limit**, select the maximum amount of time that a user's session can remain active on the server. When the time limit is reached, either the user is disconnected from the session or the session ends. When a session ends, it is permanently deleted from the server. Select **Never** to allow the session to continue indefinitely.
   - In **Idle session limit**, select the maximum amount of time that an idle session (a session without client activity)

remains on the server. When the time limit is reached, either the user is disconnected from the session or the session ends. When a session ends, it is permanently deleted from the server. Select **Never** to allow idle sessions to remain on the server indefinitely.

### Notes

- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
- Group Policy overrides the configuration set with the Terminal Services Configuration tool.
- These settings affect every client that uses the connection to connect to the terminal server. To define Session settings on a per-user basis, use the User Configuration Group Policies or the Terminal Services Extension to Local Users and Groups and Active Directory Users and Computers.t

### Note

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring session limits
Change settings for when a session limit is reached or connection is broken
Change settings for reconnecting disconnected Citrix ICA sessions
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Change settings for when a session limit is reached or connection is broken

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change settings for when a session limit is reached or connection is broken**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Sessions**, double-click the **Terminate session when time limits are reached setting**, and do one of the following:

   - To ensure that users are always disconnected from the session, allowing the session to be reconnected (default behavior), click **Disabled**, and then click **OK**.
   - To configure Terminal Server to terminate sessions rather than disconnect them, click **Enabled**, and then click **OK**. See the policy explain text for more information.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

**Notes**

- This policy affects every client that connects to the terminal server. To define Session settings on a per-user basis, use the Sessions policies under **User Configuration**.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection you want to modify, and then click **Properties**.
4. On the **Sessions** tab, above **When session limit is reached or connection is broken**, select the **Override user settings** check box.

   This allows you to configure these settings for the connection.
5. Under **When session limit is reached or connection is broken**, do one of the following:

- Click **Disconnect from session** to disconnect the user from the session, allowing the session to be reconnected.

- Click **End session** to end the session. When a session ends, it is permanently deleted from the server. Any running applications are forced to shut down, which can result in loss of data at the client.

### Notes

- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

- A connection to the server can be broken for any reason, including by request, connection error, idle time-out, or connection time-out.

- These settings affect every client that uses the connection to connect to a terminal server. To define settings on a per-user basis, use Terminal Services Group Policies or the Terminal Services Extension to Local Users and Groups.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring session limits
Set time-out settings for disconnected, active, and idle sessions
Change settings for reconnecting disconnected Citrix ICA sessions
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Change settings for reconnecting disconnected Citrix ICA sessions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change settings for reconnecting disconnected Citrix ICA sessions**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. n **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Sessions**, double-click the **Allow reconnection from original client only** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

**Notes**

- This policy affects every client that connects to the terminal server. To define Session settings on a per-user basis, use the Sessions policies under **User Configuration**.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection you want to modify, and then click **Properties**.
4. On the **Sessions** tab, above **Allow reconnection**, select the **Override user settings** check box.

   This allows you to configure reconnection settings for the connection.
5. Under **Allow reconnection**, do one of the following:

   - Click **From any client** to allow the user to reconnect to a disconnected session from any computer.
   - Click **From previous client** to allow the user to reconnect to a disconnected session from only the client computer where the session originated.

**Note**

- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

- This setting is supported only for Citrix ICA-based clients.

- By default, you can reconnect to a disconnected session from any computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring session limits
Set time-out settings for disconnected, active, and idle sessions
Change settings for when a session limit is reached or connection is broken
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Manage Client Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage client settings**

- Connect client drives and printers at logon
- Configure settings for mapping client devices
- Limit the maximum color depth

# Connect client drives and printers at logon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect client drives and printers at logon**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Client/Server data redirection** double-click the specific setting for the feature you want to configure, and then click **Enabled**.

   - To disable the mapping of client printers in Terminal Services sessions, enable **Do not allow client printer redirection**.
   - To disable the mapping of client drives in Terminal Services sessions, enable **Do not allow client drive redirection**.
   - To disable the setting of the default client printer as the default printer in Terminal Services sessions, enable **Do not set default client printer to be default printer in a session**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

**Note**

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection to which you want to connect client drives and printers, and then click **Properties**.
4. On the **Client Settings** tab, under **Connection**, clear the **Use connection settings from user settings** check box.

   This allows you to configure these settings for the connection.
5. Choose from the following options:

- Select the **Connect client drives at logon** check box to reconnect to all mapped client drives while logging on. (This option is supported for clients running any of the Windows Server 2003 family operating systems.)

- Select the **Connect client printers at logon** check box to reconnect to all mapped local client printers while logging on.

- Select the **Default to main client printer** check box to print to the default printer of the client. Otherwise, the default printer of the server is used as the default printer for all client sessions.

**Notes**

- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

- These settings affect all clients that use the connection to log on to a terminal server. If you want to define settings on a per-user basis, use Terminal Services Group Policies or the Terminal Services Extension to Local Users and Groups.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure settings for mapping client devices
Managing drive, printer, and device mappings for clients
Make your local disk drives available in a session
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Configure settings for mapping client devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure settings for mapping client devices**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Client/Server data redirection** double-click the specific setting for the feature you want to configure, and then click **Enabled**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.
   - To disable the redirection of clipboard data between server and client in Terminal Services sessions, enable **Do not allow clipboard redirection**.
   - To enable the redirection of audio data between the server and the client in Terminal Services sessions, enable **Allow audio redirection**. (Note that audio redirection is disabled by default in Terminal Server).
   - To disable the mapping of client COM ports in Terminal Services sessions, enable **Do not allow COM port redirection**.
   - To disable the mapping of client printers in Terminal Services sessions, enable **Do not allow client printer redirection**.
   - To disable the mapping of client LPT ports in Terminal Services sessions, enable **Do not allow LPT port redirection**.
   - To disable the mapping of client drives in Terminal Services sessions, enable **Do not allow drive redirection**.
   - To disable specifying of the default client printer as the default printer in a Terminal Services session, enable **Do not set default client printer to be default printer in a session**. The client printer is designated as the default printer for a Terminal Services session.

   **Note**

   - Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Note**

- These policies affect every client that connects to the terminal server. To define client device settings on a per-user basis, use the corresponding policies under **User Configuration**.

#### Using Terminal Services Configuration

1. Open Terminal Services Configuration.

2. In the console tree, click **Connections**.

3. In the details pane, right-click the connection for which you want to configure client device mappings, and then click **Properties**.

4. On the **Client Settings** tab, under **Disable the following**, select the appropriate check boxes:

   - **Drive mapping** to disable client-drive mapping. By default, **Drive mapping** is enabled (cleared).

   - **Windows printer mapping** to disable client Windows printer mapping. By default, **Windows printer mapping** is enabled (cleared). When enabled, clients can map Windows printers, and all client printer queues reconnect at logon. However, when both LPT and COM port mappings are disabled, you will be unable to manually add printers. When **Windows printer mapping** is disabled (checked), clients are unable to map Windows printers and client printer queues are not reconnected at logon. However, you will be able to manually reconnect printers if **LPT port mapping** or **COM port mapping** is enabled (cleared).

   - **LPT port mapping** to disable client LPT port mapping. By default, **LPT port mapping** is enabled (cleared). When enabled, client LPT ports are mapped for printing and are available in the port list of the Add Printer wizard. You need to manually create the printer to the LPT port using the Add Printer wizard. When **LPT port mapping** is disabled (checked), client LPT ports are not automatically mapped. You will be unable to manually create printers using LPT ports.

   - **COM port mapping** to disable client COM port mapping. By default, **COM port mapping** is disabled (cleared). When enabled, client COM ports are mapped for printing and are available in the port list of the Add Printer wizard. You will need to manually create the printer to the COM port using the Add Printer wizard. When **COM port mapping** is disabled (checked), client COM ports are not automatically mapped. You will be unable to manually create printers using COM ports.

   - **Clipboard mapping** to disable client clipboard mapping. By default, **Clipboard mapping** is enabled (cleared).

   - **Audio mapping** to disable client audio mapping. By default, **Audio mapping** is disabled (checked).

   **Note**

   - To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Connect client drives and printers at logon
Managing drive, printer, and device mappings for clients
Make your local disk drives available in a session
Make your local printer available in a session
Make your local serial port available in a session
Configure connections to play audio
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Limit the maximum color depth

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To limit the maximum color depth

- Using Group Policies (best practice)
- Using Terminal Services Configuration

#### Using Group Policies (best practice)

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Limit maximum color depth** setting, and then click **Enabled**.
3. In the **Color Depth** box, click the desired maximum color depth, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

### Note

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

#### Using Terminal Services Configuration

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection on which you want to limit the maximum client color depth, and then click **Properties**.
4. On the **Client Settings** tab, select the **Limit Maximum Color Depth** check box.

   This allows you to configure settings for the connection.
5. Select the desired maximum color depth from the list, and then click **OK**

   **Note**

   - To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

### Note

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have

been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Viewing client information
Group Policy (pre-GPMC)

# Manage Connection Permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage connection permissions**
You can customize permissions on a per-connection basis using the Terminal Services Configuration tool.

- Add users and groups to permissions lists
- Change the permissions a user or group has to a connection
- Remove a user or group from permissions lists

**Note**

- As a best practice, set Terminal Services permissions on a per-computer basis.

# Add users and groups to permissions lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add users and groups to permissions lists**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection for which you want to configure permissions, and then click **Properties**.
4. On the **Permissions** tab, click **Add...**.
5. On the **Select Users or Groups** dialog box, click **Locations...** to specify the search location.
6. Click **Object Types...** to specify the types of objects you want to search for.
7. Click **Check Names**.
8. When the name is located, click **OK**. The name now appears in the **Group or user names:** list on the **Permissions** tab.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
- After you add a user or group, default permissions are applied. To change permissions, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Change the permissions a user or group has to a connection
Remove a user or group from permissions lists

# Change the permissions a user or group has to a connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the permissions a user or group has to a connection**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection for which you want to change permissions, and then click **Properties**.
4. On the **Permissions** tab, click **Advanced** to open the **Advanced Security Settings** dialog box.
5. In **Permission Entries**, select the user or group for which you want to change permissions, and then click **Edit...** to open the **Permission Entry** dialog box.
6. In **Permissions**, select or clear, as appropriate, the **Allow** or **Deny** check boxes next to the permissions you want to set for the group.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
- You must use the Remote Desktop Users group to control remote access to Terminal Server and Remote Desktop for Administration.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add users and groups to permissions lists
Remove a user or group from permissions lists

# Remove a user or group from permissions lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a user or group from permission lists**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection for which you want to change permissions, and then click **Properties**.
4. On the **Permissions** tab, in **Group or user names:**, select the user or group you want to remove, and then click **Remove**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
- You must use the Remote Desktop Users group to control remote access to Terminal Server and Remote Desktop for Administration.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Change the permissions a user or group has to a connection
Add users and groups to permissions lists

# Manage Terminal Services Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage Terminal Services connections**

- Disable a Terminal Services connection
- Enable a Terminal Services connection
- Rename a Terminal Services connection
- Delete a Terminal Services connection

# Disable a Terminal Services connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable a connection**

1. Open Terminal Services Configuration.

2. In the console tree, click **Connections**.

3. In the details pane, right-click the connection you want to disable, point to **All Tasks**, and then click **Disable Connection**.

   Click **Yes** to confirm that you want to disable the connection. If there are active sessions on the server, you are warned that these sessions will end if you disable the connection. Ending an active session without warning can result in loss of data at the user's session.

   The connection is disabled and no users can log on to the server using the connection until it is re-enabled.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

- You should disable all connections while you install new applications. This prevents users from logging on to the server during installation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable Terminal Services connections
Terminal Services Configuration Overview

# Enable a Terminal Services connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable a connection**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection you want to enable, point to **All Tasks**, and then click **Enable Connection**.

   The connection is enabled and users can log on to the terminal server using the connection.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable a Terminal Services connection
Terminal Services Configuration Overview
Enable or disable Remote Desktop

# Rename a Terminal Services connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a connection**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection you want to rename, point to **All Tasks**, and then click **Rename Connection**.
4. In **New name**, type a new name for the connection, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services Configuration Overview
Configuring Terminal Services with TSCC

# Delete a Terminal Services connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a connection**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection you want to delete, and then click **Delete**.

   Click **Yes** to confirm that you want to delete the connection.

   **Caution**

   - If there are active sessions on the server, you are warned that these sessions end if you delete the connection. Ending an active session without warning can result in loss of data at the user's session.

   The connection is permanently deleted from the server.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services Configuration Overview
Configuring Terminal Services with TSCC

# Configure Server Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure server settings**

- Delete or retain temporary folders when exiting
- Use separate temporary folders for each session
- Disable Active Desktop
- Switch default permissions for application compatibility
- Configure Session Directory settings
- Set the terminal server licensing mode
- Limit users to one remote session

# Delete or retain temporary folders when exiting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete or retain temporary folders when exiting**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Temporary folders**, double-click the **Do not delete temp folder upon exit** setting.
3. To retain users' temporary folders when a user logs off from a session, click **Enabled**. To delete users' temporary folders when a user logs off, click **Disabled**.
4. Click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

**Note**

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Server Settings**.
3. In the details pane, right-click **Delete temporary folders on exit**, and then click **Properties**.
4. Click **Yes** to have temporary folders deleted when a user logs off from a session, or click **No** to retain temporary folders, and then click **OK**.

   **Note**

   - To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group

might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring Terminal Services with Group Policy

# Use separate temporary folders for each session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use separate temporary folders for each session**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Temporary folders**, double-click the **Do not use temp folders per session** setting.
3. To disable the creation of a separate temporary folder for each terminal server session, click **Enabled**. To enable the creation of a separate temporary folder for each terminal server session, click **Disabled**.
4. Click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Server Settings**.
3. In the details pane, right-click **Use temporary folders per session**, and then click **Properties**.
4. Click **Yes** to create a separate temporary folder for each new session or **No** if you do not want to create temporary folders for each new session.

   **Note**

   - To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring Terminal Services with Group Policy

# Disable Active Desktop

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable Active Desktop**

1. Open Terminal Services Configuration.
2. In the console tree, click **Server Settings**.
3. In the details pane, right-click **Active Desktop**, and then click **Properties**.
4. Select the **Disable Active Desktop** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services Configuration Overview
Configuring Terminal Services with TSCC

# Switch default permissions for application compatibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To switch default permissions for application compatibility**

1. Open Terminal Services Configuration.
2. In the console tree, click **Server Settings**.
3. In the details pane, right-click **Permission Compatibility**, and then click **Properties**.
4. Select **Full Security** to provide the most secure environment or select **Relaxed Security** to provide an environment that is compatible with most legacy applications, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
- You can configure this option only when Terminal Server is installed on a server. This option does not apply to servers that are used for remote administration.
- If you select **Full Security**, Terminal Server Users will have the same permissions as members of the Users group. You should use Full Security whenever possible.
- If you select **Relaxed Security**, all Users will have full access to critical registry and file system locations. This may be necessary in order to run legacy applications.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Terminal Server to host applications centrally
Terminal Services Configuration Overview
Choosing the security mode for a terminal server

# Configure Session Directory settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Session Directory settings**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Session Directory**, enable the **Join Session Directory** setting.
3. Enable the **Session Directory Server** setting and, in **Session Directory Server**, type the name of the server where the Terminal Services Session Directory service is running.
4. Enable the **Session Directory Cluster Name** setting and, in **Session Directory Cluster Name**, type the name of the cluster to which the Terminal Server belongs. See the policy explain text for more information.
5. Optionally, enable the **Terminal Server IP Address Redirection** setting. See the policy explain text for more information.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Server Settings**.
3. In the details pane, right-click **Session Directory**, and then click **Properties**.
4. Select the **Join session directory** check box.
5. In **Cluster name:**, type the name of the Terminal Server cluster for this terminal server.

6. In **Session directory server name:**, type the DNS name or IP address of the domain server where the Terminal Services Session Directory service is running.

   The server name must be a valid server name, and it cannot be left blank.

7. Select an IP address and network adapter from the **Network adapter and IP address session directory should redirect users to:** list.

8. Optionally, uncheck the **IP address redirection (uncheck for routing token redirection)** check box to have client devices reconnect to disconnected sessions by using the virtual IP address of the terminal server cluster. You should uncheck this check box if clients have visibility only to the virtual IP address of the cluster and cannot connect to the IP address of an individual terminal server. This option is checked by default, which enables clients to reconnect by using the individual IP addresses of the terminal servers in the Session Directory.

   **Note**

   - To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Load balancing and terminal servers

# Set the terminal server licensing mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the terminal server licensing mode**

1. Open Terminal Services Configuration.
2. In the console tree, click **Server Settings**.
3. In the details pane, right-click **Licensing**, and then click **Properties**.
4. Select either **Per User** or **Per Device**, and then click **OK**.

   A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
- Group Policy overrides the configuration set with the Terminal Services Configuration tool. For information about how to configure the Terminal Server Licensing mode option by using Group Policy, see Configure the Terminal Server Licensing mode.
- By default, when the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.
- Users can create more than one session by specifying a different program to start on connection for each session.
- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Server Licensing overview

Install Terminal Server Licensing
Activate a Terminal Server License Server

# Limit users to one remote session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To limit users to one remote session**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

| **Warning** |
| --- |
| If users have been limited to one remote session, connecting to the session by using a different protocol is not supported. To access a remote session by using a different protocol, the user must log off from the current session. |

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Restrict Terminal Services users to a single remote session** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.
2. In the console tree, click **Server Settings**.
3. In the details pane, right-click **Restrict each user to one session**, and then click **Properties**.
4. Check the **Restrict each user to one session** check box, and then click **OK**.

   **Note**

   - To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and

then double-click **Terminal Services Configuration**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

- Users will always be able to create more than one session by specifying a different program to start on connection for each session.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify a program to start upon connection Remote Desktop Connection
Configuring Terminal Services with Group Policy
Configuring Terminal Services with TSCC

# Configure Group Policy Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Group Policy settings**

- Ensure the session state is consistent with the client state
- Prevent automatic license upgrades
- Allow users to connect remotely using Terminal Services
- Remove the Disconnect item from the Shut Down dialog
- Remove the Windows Security item from the Start menu
- Secure Terminal Server RPC traffic
- Control the issuance of Terminal Server licenses
- Use the specified Terminal Server license servers
- Configure the Terminal Server Licensing mode
- Always prompt client for password upon connection
- Enforce removal of Remote Desktop wallpaper
- Disallow smart card device redirection

# Ensure the session state is consistent with the client state

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To ensure the session state is consistent with the client state**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Keep-Alive Messages** setting, and then click **Enabled**. select the desired **Keep-Alive interval**.
3. In the **Keep-Alive interval** box, type the desired keep-alive interval, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Viewing session states
Group Policy (pre-GPMC)

# Prevent automatic license upgrades

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prevent automatic license upgrades**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Licensing**, double-click the **Prevent License Upgrade** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Terminal Server Licensing overview
Group Policy (pre-GPMC)

# Allow users to connect remotely using Terminal Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow users to connect remotely using Terminal Services**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Allow users to connect remotely using Terminal Services** setting.
3. Do one of the following:
   - To enable Remote Desktop, click **Enabled.**
   - To disable Remote Desktop, click **Disabled**.

   If you disable Remote Desktop while users are connected to the target computers, the computers maintain their current connections, but will not accept any new incoming connections.

### Important

- When you enable Remote Desktop on a computer, you enable the capability for other users and groups to log on remotely to the computer. However, you must also decide which users and groups should be able to log on remotely, and then manually add them to the Remote Desktop Users group. For more information, see Enabling users to connect remotely to the server and Add users to the Remote Desktop Users group.
- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use this procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an administrator. Then, you must start Group Policy by using the **Active Directory Users and Computers** snap-in.
- Alternatively, you can enable Remote Desktop by selecting the **Enable Remote Desktop on this computer** check box (on the **Remote** tab of the **System Properties** dialog box) on the target computers. For more information, see Enable or disable Remote Desktop.

  If the **Allow users to connect remotely using Terminal Services** Group Policy setting described in this procedure is set to **Not Configured**, then the **Enable Remote Desktop on this computer** setting on the **System Properties** dialog box of the target computers takes precedence. Otherwise, the Group Policy setting takes precedence.
- Be aware of the security implications of remote logons. Users who log on remotely can perform tasks as though they were sitting at the console. For this reason, you should ensure that the server is behind a firewall. For more information, see VPN servers and firewall configuration and Security information for IPSec.
- You should require all users who make remote connections to use a strong password. For more information, see Strong

passwords.

- Remote Desktop is disabled by default in Windows Server 2003 family operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# To restrict local administrators from customizing permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To restrict local administrators from customizing permissions

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Do not allow local administrators to customize permissions** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

### Note

- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in. For information on applying Group Policy to groups of users or computers, see Configuring Terminal Services with Group Policy. For information on configuring the local Group Policy object, see Group Policy overview.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Remove the Disconnect item from the Shut Down dialog

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove the Disconnect item from the Shut Down dialog**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Remove Disconnect item from Shut Down dialog** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Log off or disconnect
Group Policy (pre-GPMC)

# Remove the Windows Security item from the Start menu

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove the Windows Security item from the Start menu**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Remove Windows Security item from Start menu** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)

# Secure Terminal Server RPC traffic

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To secure Terminal Server RPC traffic**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Encryption**, **RPC Security Policy**, double-click the **Secure Server (Require Security)** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The remote procedure call (RPC) interface is used for administering and configuring Terminal Services.
- If the status of this setting is Enabled, the terminal server accepts requests from RPC clients that support secure requests, and does not allow unsecured communication with untrusted clients.
- If the status of this setting is Disabled, the terminal server always requests security for all RPC traffic. However, unsecured communication is allowed for RPC clients that do not respond to the request.
- If the status of this setting is Not Configured, unsecured communication is allowed. This is the default.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Group Policy (pre-GPMC)
IPv6 protocol features
Default settings for services

# Control the issuance of Terminal Server licenses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To control the issuance of Terminal Server licenses**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **License Server Security Group** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- If the status of this setting is Enabled, only the computers you add to the Terminal Services Computers security group will receive licenses. By default, this group is not populated. You must add computers or groups to this security group if you want them to receive licenses. Any changes you make to the membership of this group take effect only after you stop and restart the Terminal Server Licensing service. For more information, see Related Topics.
- If the status of this setting is either Disabled or Not Configured, the Terminal Server License Server issues a license to any computer that requests one. The Terminal Services Computers group is not deleted or changed in any way. This is the default behavior.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Controlling the issuance of Terminal Server licenses
Group Policy (pre-GPMC)
Changing group memberships

# Use the specified Terminal Server license servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Use the specified Terminal Server license servers** setting.
3. Click **Enabled**.
4. Enter the names of the Terminal Server license servers you want this Terminal Server to contact first when attempting to obtain client access licenses (CALs). Then click **OK**.

### Important

- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- You can also set preferred license servers by using Terminal Services Configuration (TSCC), the registry, or a Windows Management Instrumentation (WMI) script. For more information about setting preferred license servers by using TSCC or the registry, see Related Topics. For three WMI scripts that you can use to set preferred license servers, delete preferred license servers, or query preferred license servers, see "Preferred License Server WMI Scripts" in Windows Server 2003 Terminal Server Licensing (http://www.microsoft.com/).
- This Group Policy setting is only available on Windows Server 2003 with Service Pack 1, or with the Service Pack 1 versions of System.adm Group Policy template files installed in your Terminal Server environment.
- Use this procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an administrator. Then, you must open Group Policy through the **Active Directory Users and Computers** snap-in.
- For more information about Terminal Server, see Guidelines for Deploying Terminal Server (http://www.microsoft.com/).

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change Terminal Server license server properties
Terminal Server license server roles
Set preferred Terminal Server license servers

# Configure the Terminal Server Licensing mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the Terminal Server Licensing mode**

- Using Group Policies (best practice)
- Using Terminal Services Configuration

**Using Group Policies (best practice)**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click **Terminal Server Licensing Mode**.
3. Click **Enabled**.
4. Click either **Per User** or **Per Device**, and then click **OK**.

   A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices.

### Important

- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- This Group Policy setting is only available on Windows Server 2003 with Service Pack 1.
- Use this procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an administrator, and then open Group Policy by using the **Active Directory Users and Computers** snap-in.
- Group Policy overrides the configuration set with the Terminal Services Configuration tool.
- By default, when the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.
- Users can create more than one session by specifying a different program to start on connection for each session.

**Using Terminal Services Configuration**

1. Open Terminal Services Configuration.

2. In the console tree, click **Server Settings**.

3. In the details pane, right-click **Licensing Mode**, and then click **Properties**.

4. Select either **Per User** or **Per Device**, and then click **OK**.

   A Per Device CAL gives each client computer or device the right to access a terminal server that is running Windows Server 2003. A Per User CAL gives one user the right to access a terminal server from an unlimited number of devices.

## Notes

- Terminal Servicesccopen

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Group Policy overrides the configuration set with the Terminal Services Configuration tool.

- By default, when the Per Device licensing mode is used, after a client logs and authenticates to a terminal server for the first time, and after the terminal server locates a license server, the terminal server issues the client a temporary license. After a client logs on and authenticates to a terminal server for the second time, if the license server is activated and at least one Per Device CAL is installed on the license server and not yet issued, the terminal server issues the client a permanent, Per Device CAL. Per User CALs are not monitored by Terminal Server. This means that even though there is a Per User CAL in the license server database, the Per User CAL is not decremented when it is used. Additionally, if you use the Per User licensing mode, when a client logs on to a terminal server for the second time, the temporary license is not upgraded to a permanent CAL.

- Users can create more than one session by specifying a different program to start on connection for each session.

- For more information about Terminal Server and Terminal Server Licensing, see Guidelines for Deploying Terminal Server (http://go.microsoft.com/fwlink/?LinkID=34627) and Windows Server 2003 Terminal Server Licensing (http://go.microsoft.com/fwlink/?LinkID=26220).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Server Licensing
Terminal Server Licensing overview
Terminal Server license server roles

# Always prompt client for password upon connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To always prompt client for password upon connection**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, **Encryption and Security**, double-click the **Always prompt client for password upon connection** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Provide information for automatic logon
Group Policy (pre-GPMC)

# Enforce removal of Remote Desktop wallpaper

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enforce removal of Remote Desktop wallpaper**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Enforce Removal of Remote Desktop Wallpaper** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Improve the performance of connections
Group Policy (pre-GPMC)

# Disallow smart card device redirection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disallow smart card device redirection**

1. Open Group Policy.
2. In **Computer Configuration**, **Administrative Templates**, **Windows Components**, **Terminal Services**, double-click the **Do not allow smart card device redirection** setting.
3. Click **Enabled**, and then click **OK**.

   **Important**

   - You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Use the above procedure to configure the local Group Policy object. To change a policy for a domain or an organizational unit, you must log on to the primary domain controller as an Administrator. Then, you must invoke Group Policy through the **Active Directory Users and Computers** snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring Terminal Services with Group Policy
Authentication and smart card support

# Terminal Services Configuration Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides information about configuring Terminal Services.

- Terminal Services Configuration Overview
- Configuring Terminal Services with Group Policy
- Configuring Terminal Services with WMI
- Configuring Terminal Services with TSCC
- Managing Sessions
- Managing Permissions on Connections
- Managing drive, printer, and device mappings for clients
- Using temporary folders with terminal servers

# Terminal Services Configuration Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services Configuration overview**
Several tools exist for configuring terminal servers, Terminal Services user settings, Terminal Services connections, and Terminal Services sessions. These tools include Terminal Services Group Policies, the Terminal Services WMI provider, the Terminal Services Configuration Tool, the Terminal Services extension to Local Users and Groups and Active Directory Users and Computers. Additionally, connection settings that are not configured at the group, computer or user level can be set in the client software itself on a per-session basis.

**Which tool should I use?**
Typically, you use the tool that gives you the most control with the least administrative overhead, given the constraints of your usage profile. For an environment containing only terminal servers running one of the Windows Server 2003 family operating systems, using Terminal Services Group Policies will allow you to configure all settings that will apply across an organizational unit or group of Terminal Services users. The Active Directory Service Interfaces (ADSI) further extend the Terminal Services extensions to Local Users and Groups and Active Directory Users and Computers by allowing you to create scripts tailored to your Terminal Services configuration requirements. You can also configure individual Windows Server 2003 family operating systems by using Group Policies on the local Group Policy object.

However, there are situations where Group Policy is not the preferred tool. In mixed environments where several different Windows platforms are present, it might be necessary to use a combination of tools--for example, configuring Windows Server 2003 family operating systems with Group Policies, while using the Terminal Services Configuration tool to configure servers that are running previous versions of Windows.

If two or more connections are present on the same computer, and you want to configure each connection differently, Group Policies cannot be used. Instead, use the Terminal Services Configuration tool, which allows you to configure Terminal Services settings on a per-connection basis.

**Configuration order of precedence**
Configuration settings take precedence in the following order. Settings higher on the list take precedence over those below:

- Computer level Group Policies (if set)
- User level Group Policies (if set)
- Local computer configuration set with Terminal Services Configuration tool
- User level policies set with Local Users and Groups
- Local client settings

**Note**

- Configuration settings applied with Terminal Services WMI operate at the same order of precedence as the same settings when set with the corresponding configuration tool. In general, Group Policy settings always override settings applied with WMI.

For more information about configuring Terminal Services with Group Policy, see Configuring Terminal Services with Group Policy.

For more information about configuring Terminal Services with Terminal Services Configuration, see Terminal Services Configuration.

For more information about configuring Terminal Services user policies with Local Users and Groups, see Managing Users with the Terminal Services Extension to Local Users and Groups.

For more information about using ADSI, see Programming interfaces and the Active Directory Programmer's Guide at the Microsoft Web site .Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Configuring Terminal Services with Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Terminal Services with Group Policy**
You can use Group Policy to configure Terminal Services connection settings, set user policies, configure terminal server clusters, and manage Terminal Services sessions. You can enable Group Policy for users of a computer, for individual computers, or for groups of computers belonging to an organizational unit of a domain. To set policies for users of a particular computer, you must be an administrator for that computer. To set policies for an organizational unit in a domain, you must be an administrator for that domain.

**Enabling Group Policy on an individual computer**
To set Terminal Services policies setting for a particular computer or for users of that computer, open the Group Policy Object Editor snap-in to edit the local Group Policy. For more information, see Group Policy overview.

When you open the Group Policy Object Editor snap-in, click to expand Administrative Templates, and then click the Terminal Services folder containing the policies you want want to set. The settings for those policies appear in the details pane.

**Where?**

- *policy_name* Policy/Computer Configuration (or User Configuration)/Administrative Templates/Windows Components/*Terminal Services*

To determine the Enabled, Disabled, and Not Configured behavior of each policy, see the explain text for that policy in the Group Policy Object Editor snap-in.

For more information, see Group Policy overview.

**Enabling Group Policy for an organizational unit in a domain**
To set Terminal Services policies for a domain, you must use a computer that is configured as a domain controller and you must be an Administrator for that domain.

After you configure your computer as a domain controller, the Active Directory Users and Computers console appears in the Administrative Tools folder on the **Programs** menu. You can use this tool to set policies for organizational units in the domain. Before you begin, you need to confirm that **Group Policy** is selected (on the **Extensions** tab in the **Add/Remove Snap-in** dialog box). Then you need to confirm that the **Administrative Templates (Users)** extension is selected for Group Policy.

**Note**

- Terminal Services Group Policies can be used to administer only computers running one of the Windows Server 2003 family operating systems.

# Configuring Terminal Services with WMI

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Terminal Services with WMI**
In Microsoft® Windows Server 2003 family operating systems, you can use the Terminal Services WMI provider to create scripts for remote administration of terminal servers.

**About the Terminal Services WMI Provider**
The Terminal Services Windows Management Instrumentation (WMI) provider allows administrators to create customized scripts for configuring, managing, and querying terminal servers. It contains properties and methods that can perform the same tasks as the Terminal Services configuration tools and command-line utilities provided with Windows Server 2003 family operating systems, but remotely and through scripted applications.

**How the Provider Works**
Windows Management Instrumentation (WMI) defines a unified architecture for describing, accessing, and instrumenting objects. Part of this architecture is a large database of WMI classes used to carry out remote management tasks on specific objects.

A provider extends the WMI schema of classes to allow WMI to work with new types of objects. The Terminal Services Provider defines classes for querying and configuring Terminal Services.

The Terminal Services Provider is defined in %systemroot%\System32\Wbem\tscfgwmi.mof.

A general description of the classes associated with the provider can be seen in the following table. For complete list of properties and methods associated with each class, please see the embedded comments in the Terminal Services provider MOF file (tscfgwmi.mof). The contents of this file can be viewed with the Windows Management Instrumentation Tester (Wbemtest.exe), which is located in the %systemroot%\System32\Wbem\. To view the .mof file in Wbemtest, use the procedure below.

**To view the Terminal Services Provider classes in Wbemtest.exe**

1. Open Wbemtest.exe
2. Click **Connect**.
3. In **Namespace**, type **root\Cimv2**
4. Click **Connect**.
5. Click **Open Instance**.
6. In **Object Path**, type the name of the Terminal Services Provider class (see the table below for the names of the classes.) Click **OK**.
7. In **Properties**, click the name of the property you want to view, and then click **Show MOF**.

**Terminal Services WMI Classes**
The Terminal Services Provider has the following WMI classes.

| WMI Class | Description |
|---|---|
| Win32_TerminalService | The Win32_TerminalService class is a subclass of Win32_Service class and inherits all its properties and methods. In addition, Win32_TerminalService represents the Element property of the Win32_TerminalServiceToSetting association. |
| Win32_TSSessionDirectory | Defines the configuration for Win32_TSSessionDirectorySetting. This includes properties such as Session Directory store, Cluster Name and Additional parameters. |

| | |
|---|---|
| Win32_TerminalServiceSetting | Defines the configuration for TerminalServerSetting. This includes properties such as Terminal Server Mode, Licensing, Active Desktop, Permissions Capability, Deletion of Temporary folders and per-session Temporary folders. |
| Win32_Terminal | Associates a TerminalSetting and its several configuration setting groups such as: General, Logon, Session, Environment, Remote Control, Client, Network Adapter and Permission. |
| Win32_TSGeneralSetting | Defines the configuration for properties such as Protocol, Transport, Comment, Windows authentication and Encryption Level |
| Win32_TSLogonSetting | Defines the configuration for properties such as ClientLogonInfoPolicy, UserName, Domain, and Password. |
| Win32_TSSessionSetting | Defines the configuration for properties such as time limit policy for connections and session limits for active, idle, and disconnected sessions. |
| Win32_TSEnvironmentSetting | Defines the configuration for properties such as Initial program policy, Initial program path, Start in, and Client wallpaper. |
| Win32_TSRemoteControlSetting | Defines the configuration for properties such as Remote Control policy and level of control. |
| Win32_TSClientSetting | Defines the configuration for properties such as Connection policy, Windows printer mapping, Comport mapping, etc. |
| Win32_TSNetworkAdapterSetting | Defines the configuration for properties such as LAN Adapter and Maximum number of connections |
| Win32_TSNetworkAdapterListSetting | Enumerates the list of network adapters that are configurable for Win32_Terminal based on the Terminal Protocol and Transport. The properties include: TerminalProtocol, Transport, NetworkAdapterID, and Description. |
| Win32_TSPermissionsSetting | Represents an association between a Win32_Terminal and its configuration settings (Win32_TSAccount). Contains properties for adding accounts to defined permission sets and restoring default permissions. |
| Win32_TSAccount | Defines the configuration for properties such as SID, Allowed and denied permissions. |

**Important**

- This documentation describes only the classes, methods, and associations specific to Terminal Services. Top-level WMI classes (those with the CIM_ prefix) that are defined in the MOF file are not specific to Terminal Services and are described in the Platform SDK documentation for Windows Management Instrumentation (WMI).

For an overview of WMI, see WMI Overview.

For a complete description and definition of WMI, as well as technical information about developing for the WMI system, see the WMI Software Development Kit (SDK). The WMI SDK is released by MSDN as part of the Microsoft Platform SDK.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configuring Terminal Services with TSCC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Terminal Services with TSCC**
A Terminal Services connection provides the link clients use to log on to a session on the server. A TCP/IP connection is configured when Terminal Services is installed. Using Terminal Services Configuration, you can change the default properties of the connection or add new connections.

**Terminal Services Configuration**
When you open Terminal Services Configuration you will see that a connection has already been configured. This is called the RDP-TCP connection. Typically, this is the only connection that needs to be configured for clients to connect to the server for Remote Desktop for Administration or application sharing with Terminal Server. Only one RDP (Remote Desktop Protocol) connection can be configured for each network adapter. If you want to configure additional RDP connections, you must install additional network adapters.

With Terminal Services Configuration, you can reconfigure the properties of the RDP-TCP connection, which includes limiting the amount of time client sessions can remain active on the server, setting protection levels for encryption, and selecting which permissions you want users and groups to have. Some connection properties can also be configured on a per-user basis using Terminal Services Group Policies or the Terminal Services extension to Local Users and Groups. For example, you can set different session time limits for each user when you use the Terminal Services extension to Local Users and Groups. Using Terminal Services Configuration, you can only set session time limits on a per-connection basis, which means the same time limit applies to all users who log on to the server using the connection.

In addition to configuring connections, you can use Terminal Services Configuration to configure settings that apply to the terminal server. These include settings for temporary folders, default connection security, and licensing. For more information, see Configure Server Settings

**Citrix ICA-based client**
Terminal Services Configuration is also used to configure connections for Citrix ICA-based clients. These connections can be configured using TCP/IP, asynchronous, IPX/SPX, or NetBIOS transport protocols when the Citrix ICA protocol and Citrix ICA-based client software are added to the system.

### Important

- As a best practice, use Terminal Services Group Policies to configure connection settings wherever possible. For more information, see Configuring Terminal Services with Group Policy.

# Managing Sessions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing sessions**

This section provides general background information about managing Terminal Services sessions:

- Configuring authentication and encryption
- Using remote control
- Configuring session limits

# Configuring authentication and encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring authentication and encryption**

**TLS authentication overview**
Remote Desktop Protocol (RDP) provides data encryption, but it does not provide authentication to verify the identity of a terminal server. In Windows Server 2003 Service Pack 1 (SP1), you can enhance the security of Terminal Server by configuring Terminal Services connections to use Transport Layer Security (TLS) 1.0 for server authentication, and to encrypt terminal server communications. TLS is a standard protocol that is used to provide secure Web communications on the Internet or intranets. It enables clients to authenticate servers or, optionally, servers to authenticate clients. It also provides a secure channel by encrypting communications. TLS is the latest and most secure version of the Secure Sockets Layer (SSL) protocol. One important difference is that TLS 1.0 applies a Keyed-Hashing for Message Authentication Code (HMAC) algorithm, whereas SSL 3.0 applies the Message Authentication Code (MAC) algorithm. The HMAC produces an integrity check value as the MAC does, but with a hash function construction that makes the hash much harder to break.

The TLS and SSL protocols comprise two layers: the Handshake Protocol Layer and the Record Protocol Layer. For authentication purposes, the Handshake Protocol uses an X.509 certificate to provide strong evidence to a second party that helps prove the identity of the party that holds the certificate and the corresponding private key. A certificate is a digital form of identification that is usually issued by a certification authority (CA) and contains identification information, a validity period, a public key, a serial number, and the digital signature of the issuer. A CA is a mutually trusted third party that confirms the identity of a certificate requestor (usually a user or computer), and then issues the requestor a certificate. The certificate binds the requestor's identity to a public key. CAs also renew and revoke certificates as necessary. For example, if a client is presented with a server's certificate, the client computer might try to match the server's CA against the client's list of trusted CAs. If the issuing CA is trusted, the client will verify that the certificate is authentic and has not been tampered with. Finally, the client will accept the certificate as proof of identity of the server.

For more information about TLS, see SSL/TLS in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkID=19646) and RFC 2246, The TLS Protocol Version 1.0 (http://go.microsoft.com/fwlink/?LinkID=40979).

**Prerequisites for Configuring Server Authentication**
By default, Terminal Server uses native RDP encryption and does not authenticate the server. For TLS to be used for server authentication and encryption of terminal server communications, the server and client must be correctly configured.

**Server prerequisites**
For TLS authentication to function correctly, terminal servers must meet the following prerequisites:

- Terminal servers must run Windows Server 2003 SP1.
- You must obtain a certificate for the terminal server. You can do this by doing any of the following:
    - Use Windows Server 2003 Certificate Services Web Pages or Use Windows 2000 Certificate Services Web Pages.
    - Use the Windows Server 2003 Certificate Request Wizard or Windows Server 2000 Certificate Request Wizard. For more information, see Requesting certificates.
    - Purchase a certificate from a non-Microsoft vendor and install the certificate manually.

**Note**

- If you plan to obtain a certificate by using the Certificate Web pages or Certificate Request Wizard, a public key infrastructure (PKI) must be configured correctly to issue SSL-compatible X.509 certificates to the terminal server. For information about PKI configuration, see Public Key Infrastructure and Deploying a Public Key Infrastructure.Each certificate must be configured as follows:
    - The certificate is a computer certificate.

- The intended purpose of the certificate is server authentication.
- The certificate has a corresponding private key.
- The certificate has a cryptographic service provider (CSP) that can be used for the TLS protocols (for example, Microsoft RSA/SChannel Cryptographic Provider). For more information, see Microsoft Cryptographic Service Providers (http://go.microsoft.com/fwlink/?LinkID=40983).
- The certificate is stored in the terminal server's Personal store. You can view this store by using the Certificates snap-in. For more information, see Certificate stores.

**Client prerequisites**

For TLS authentication to function correctly, clients must meet the following prerequisites:

- Clients must run Windows 2000 or Windows XP.
- Clients must be upgraded to use the RDP 5.2 (Windows Server 2003) client. You can install this client-side Remote Desktop Connection package by using the %systemdrive\system32\clients\tsclient\win32\msrdpcli.msi file. The msrdpcli.msi file is located on Windows Server 2003 terminal servers. Installing this file from the terminal server installs the 5.2 version of Remote Desktop Connection to the %systemdrive\Program files\Remote Desktop directory on the destination computer. For more information, see Remote Desktop Connection for Windows Server 2003 [5.2.3790] (http://go.microsoft.com/fwlink?/LinkID=41068).
- Clients must trust the root of the server's certificate. That is, clients must have the certificate of the CA that issued the server certificate in their Trusted Root Certification Authorities store. You can view this store by using the Certificates snap-in.

**Tasks for Configuring Server Authentication**

For TLS authentication to function correctly, you must complete the following configuration tasks on terminal servers and clients. The procedure links contain step-by-step instructions for completing these tasks.

**Server Configuration Tasks**

1. Request a computer certificate for server authentication, if you do not already have a certificate that meets the requirements described in "Server prerequisites," earlier in this topic.

   If you have been granted the appropriate permissions, there are two primary ways to explicitly request certificates:

   - You can explicitly request certificates by using the Windows Server 2003 or Windows 2000 Certificate Services Web pages.
   - You can use the Windows Server 2003 or Windows 2000 Certificate Request Wizard.

   Alternatively, you can buy a certificate from a non-Microsoft vendor.

2. Configure authentication and encryption on the server. For TLS to function correctly, in Terminal Server Configuration, on the **General** tab of the **RDP-tcpProperties** dialog box, you must do the following:

   - Select a certificate that meets the requirements described in "Server prerequisites" earlier in this topic.
   - Set the **Security layer** to **Negotiate** or **SSL**.
   - Set the **Encryption level** to **High,** or enable Federal Information Processing Standard (FIPS) compliant encryption. You can enable FIPS-compliant encryption by using Group Policy as well.
   - You cannot enable TLS by using Group Policy.

**Note**

If you enable TLS authentication in a session directory farm, in order for TLS authentication to remain enabled when a terminal server connection is redirected, you must configure all of the servers that are members of the session directory farm to use one of the following settings:

- Set the **Security layer** to **SSL**.
- Set the **Security layer** to **Negotiate**. If you use this configuration, TLS authentication is only enabled if the client supports it.

**Authentication and Encryption Settings**

The following table summarizes the combinations of **Security layer** and **Encryption level** settings that are valid and invalid in Windows Server 2003 Service Pack 1, whether TLS authentication and encryption can be used, and the key strength provided by each encryption level.

| ***Security Layers*** | | **Encryption** | **Levels** | |
|---|---|---|---|---|
| | **Client Compatible** | **Low**<br><br>**RC4--56-bit** | **High**<br><br>**RC4--128-bit** | **FIPS-compliant** |
| ***Negotiate*** | Valid. The Maximum security layer and encryption level supported by the client are used. If TLS is supported for server authentication, TLS is used rather than RDP, and High or FIPS-compliant encryption level is used. | Valid | Valid | Valid |
| ***SSL (TLS 1.0)*** | Valid. TLS is required for server authentication. High or FIPS-compliant encryption level is used. If TLS is not supported, connections fail. | Not Valid. The encryption level defaults instead to Client Compatible. High or FIPS-compliant encryption level is used, if supported by client. | Valid | Valid |
| ***RDP*** | Valid. However, TLS cannot be used for server authentication. RDP encryption is used and the maximum encryption level supported by the client is used. | Valid | Valid | Valid |

For more information about encryption options, see "Determining the level of encryption," later in this topic.

**Client configuration tasks**

1. Request a certification authority certificate for the client.

   Use this procedure to request a certificate chain from a Windows Server 2003 CA. A certificate chain contains the chain of certificates that can be traced back to the root CA. The certificate that you retrieve by using this procedure is stored in the Trusted Root Certification Authorities store.

2. Configure authentication on the client by using Remote Desktop Connection or Configure authentication on the client by using an .rdp file.

   Keep in mind that clients must run Windows 2000 or Windows XP and that they must use the RDP 5.2 (Windows Server 2003) client.

**Determining the level of encryption**

For Terminal Services connections, data encryption can protect your data by encrypting it on the communications link between the client and the server. Encryption protects against the risk of unauthorized transmission interception on the link between server and client.

By default, Terminal Services connections are encrypted at the highest level of security available (128-bit). However, some older versions of the Terminal Services client do not support this high level of encryption. If your network contains such legacy clients, you can set the encryption level of the connection to send and receive data at the highest encryption level supported by the client. In order for clients to be able to connect to a terminal server that uses FIPS-compliant encryption, you must upgrade these clients to use the RDP 5.2 (Windows Server 2003) client. For more information, see Remote Desktop Connection for Windows Server 2003 [5.2.3790] (http://go.microsoft.com/fwlink?/LinkID=41068).

There are four levels of encryption available:

| **Level of encryption** | **Description** |
|---|---|
| FIPS-compliant | This level encrypts and decrypts data sent from client to server and from server to client with the Federal Information Processing Standard (FIPS) encryption algorithms using the Microsoft cryptographic modules. You can enable FIPS-compliant encryption by using Group Policy or Terminal Services Configuration. For more information, see FIPS 140 Evaluation (http://go.microsoft.com/fwlink?/LinkID=36329).<br><br>**Important**<br><br>• Any encryption level settings that you configure in Group Policy override the configuration that you set by using the Terminal Services Configuration tool. Also, if you enable the System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing Group Policy setting, this setting overrides the **Set client connection encryption level** Group Policy setting.<br><br>• You must enable FIPS-compliant encryption or set the encryption level to **High** if you plan to use TLS for server authentication. |
| High | This level encrypts data sent from client to server and from server to client by using strong 128-bit encryption. Use this level when the terminal server is running in an environment containing 128-bit clients only (such as Remote Desktop Connection clients). Clients that do not support this level of encryption will not be able to connect.<br><br>**Important**<br><br>• You must enable FIPS-compliant encryption or set the encryption level to **High** if you plan to use TLS for server authentication. |
| Client Compatible | This level encrypts data sent between the client and the server at the maximum key strength supported by the client. Use this level when the terminal server is running in an environment containing mixed or legacy clients. |
| Low | This level encrypts data sent from the client to the server using 56-bit encryption.<br><br>**Important**<br><br>• Data sent from the server to the client is not encrypted. |

**See Also**
**Concepts**
Working with MMC console files
Request a computer certificate for server authentication
Configure authentication and encryption on the server
Request a certification authority certificate for the client
Configure authentication on the client by using Remote Desktop Connection
Configuring Terminal Services with Group Policy

# Using remote control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using remote control**

You can monitor the actions of a client logged on to a terminal server by remotely controlling the user's session from another session. Remote control allows you to either observe or actively control another session. If you choose to actively control a session, you will be able to input keyboard and mouse actions to the session. A message can be displayed on the client session asking permission to view or take part in the session before the session is remotely controlled. You can use Terminal Services Group Policies or Terminal Services Configuration to configure remote control settings for a connection and Terminal Services Manager to initiate remote control on a client session. For instructions on configuring remote control, see Configure remote control settings. For instructions on using remote control, see Remotely control a session.

Windows Server 2003 family operating systems also support Remote Assistance, which allows greater versatility for controlling another user's session. Remote Assistance also provides the ability to chat with the other user. For more information, see Windows interface administrative tool reference A-Z: Remote Assistance.

### Notes

- Remote control can also be configured on a per-user basis using Group Policies or the Terminal Services Extension to Local Users and Groups and Active Directory Users and Computers.
- The console session cannot remotely control another session, and a client session cannot remotely control the console session.
- The computer used for your session must be capable of supporting the video resolution used at the remotely controlled client session. Otherwise, the operation fails.

# Configuring session limits

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring session limits**

For a Terminal Services connection, you can limit the amount of time that active, disconnected, and idle (without client activity) sessions remain on the server. This is useful because sessions that remain running indefinitely on the server consume valuable system resources. When a session limit is reached for active or idle sessions, you can select to either disconnect the user from the session or end the session. A user who is disconnected from a session can reconnect to the same session later. When a session ends, it is permanently deleted from the server and any running applications are forced to shut down, which can result in loss of data at the client. When a session limit is reached for a disconnected session, the session ends, which permanently deletes it from the server. Sessions can also be allowed to continue indefinitely.

Session limits configured in Terminal Services Configuration apply to all sessions using a given connection. Session limits can also be configured on a per-user basis using the Terminal Services Extension to Local Users and Groups and Active Directory Users and Computers.

For information on configuring time-out and reconnection settings, see Configure Time-out and Reconnection Settings.

**Disconnected sessions**

A session is disconnected when any of the following occurs: by request, connection error, or the active or idle session limit is reached.

**Reconnecting to a disconnected session**

By default, a terminal server allows you to reconnect to a disconnected session from any computer. However, you can restrict users and allow them to only reconnect from the computer where the session originated. This option is supported only for clients that provide a serial number when connecting, such as Citrix ICA-based clients. For more information, see Change settings for reconnecting disconnected Citrix ICA sessions.

### Note

- As a best practice, use Terminal Services Group Policies to configure these settings wherever possible.

  For more information, see Configuring Terminal Services with Group Policy.

# Managing Permissions on Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing permissions on connections**
Permissions define how users and groups can access a terminal server, allowing you to secure the server. Terminal Services permissions can be handled easily on a per-computer basis, using the **Remote Desktop Users** user group and the **RemoteInteractiveLogon** right. In some cases, however, it might be necessary to manage permissions on a per-connection basis.

**Managing Permissions on a per-connection basis**
Terminal Services Configuration allows you to manage permissions on an individual connection by means of the **Permissions** tab in the connection **Properties** page.

The TCP/IP connection installed with Terminal Services comes with a set of default permissions. You can modify these default permissions by setting different permissions for different users or groups, adjusting them to fit the requirements of your organization. You must be logged on as a member of the Administrators group to manage connection permissions.

The default permissions on Terminal Services objects are as follows:

| Group | Permission |
|---|---|
| Administrators | Full control |
| LOCAL SERVICE | Service permissions |
| NETWORK SERVICE | Service permissions |
| Remote Desktop Users | User access |
| SYSTEM | Full control |
| Guest | Guest access |

For information about managing permissions for users and groups, see Adding users and groups.

For information about the types of permissions used to control access to Terminal Services, see Controlling connection access.

# Controlling connection access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Controlling connection access**

Use the permissions provided for Terminal Services to control how users and groups access the server. You can alter the default permissions to restrict individual users and groups from performing certain tasks, such as logging off a user from a session or ending sessions. You manage permissions from Terminal Services Configuration. You must be logged in as a member of the Administrators group to set permissions.

By default, there are three types of permissions: Full control, User access, and Guest access.

You can configure these by setting permissions for which users or groups can use a specific task. You can set the following permissions.

| Permission | Allows you to |
|---|---|
| Query Information | Query sessions and servers for information. |
| Set Information | Configure connection properties. |
| Remote Control | View or actively control another user's session. |
| Logon | Log on to a session on the server. |
| Logoff | Log off a user from a session. Be aware that logging off a user without warning can result in loss of data at the client. |
| Message | Send a message to another user's sessions. |
| Connect | Connect to another session. |
| Disconnect | Disconnect a session. |
| Virtual Channels | Use virtual channels. Be aware that turning off virtual channels disables some Terminal Services features such as clipboard and printer redirection. |

**Note**

- You must use the Remote Desktop Users group to control remote access to Terminal Server and Remote Desktop for Administration. For more information, see Enabling users to connect remotely to the server.

# Adding users and groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding users and groups**

You control how users and groups access a terminal server by setting permissions. You should assign permissions that reflect the level of access you want to permit. Before adding users and groups to permissions lists, consider what types of access are required for each individual user or group.

You can add any user, group, or computer in your network to permissions lists. If your server belongs to a domain, you can also add user accounts and global groups from both the local domain and trusted domains. Local groups in some domains can also be added.

You must use the Remote Desktop Users group to control remote connection permissions for Remote Desktop for Administration and Terminal Server. The Remote Desktop Users group allows the same access as the Users group, with the additional ability to connect remotely. By using the Remote Desktop Users group, you save administrative resources by not having to set up these rights for each user individually. For information about how to add users to the Remote Desktop Users group, see Add users to the Remote Desktop Users group.

# Full control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Full control**
Full control allows users to:

- Query information about a session.
- Modify connection parameters.
- Reset (end) a session.
- Remotely control another user's session.
- Log on to a session on the server.
- Log off a user from a session.
- Send a message to another user's session.
- Connect to another session.
- Disconnect a session.
- Use virtual channels, which provides access from a server program to client devices.

For information on controlling connection access, see Controlling connection access.

# User access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User access**

User access allows users to:

- Log on to a session on the server.
- Query information about a session.
- Send messages to other user sessions.
- Connect to another session.

For information on controlling connection access, see Controlling connection access.

# Guest access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Guest access**

Guest access allows users to:

- Log on to a session on the server.

For information on controlling connection access, see Controlling connection access.

# Service permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service permissions**

Local and Network service permissions allow:

- Query information about a session.
- Send a message to another session.

For information on controlling connection access, see Controlling connection access.

# Managing drive, printer, and device mappings for clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing drive, printer, and device mappings for clients**
Because client sessions can establish multiple data channels between client and server, users can map to local devices, such as drives and printers. By default, drive and printer mappings that a user sets in a client session are temporary and are not available the next time the user logs on to the server. However, using Terminal Services Configuration, you can specify that client mappings be restored when the user logs on. In addition, you can disable specific client devices so that a user cannot map the device. Users can map the following devices:

- Drives (this option is supported for Remote Desktop Connection clients logging on to one of the Windows Server 2003 family operating systems and for Citrix ICA connections)
- Windows printers
- LPT ports
- COM ports
- Smart cards
- Clipboard
- Audio (this option is supported for Remote Desktop Connection clients logging on to one of the Windows Server 2003 family operating systems and for Citrix ICA connections)

**Notes**

- As a best practice, use Terminal Services Group Policies to configure these settings wherever possible.

  For more information, see Configuring Terminal Services with Group Policy.

- Settings for client drive, printer, and device mappings configured in Terminal Services Configuration apply to all client sessions on the server. If you want to specify client settings on a per-user basis, use Terminal Services Group Policies or the Terminal Services extensions to Local Users and Groups.
- For clients running Remote Desktop Connection, these settings must also be configured in the client software. For more information, see Make your local disk drives available in a session, Make your local printer available in a session, Make your local serial port available in a session, and Make a smart card available in a session.

# Using temporary folders with terminal servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using temporary folders with terminal servers**

By default, a terminal server creates a separate temporary folder for each new session on the server, enabling each user to store individual temporary files. To reclaim disk space, these temporary folders are deleted when the user logs off from a session. If you want to change these default settings for a terminal server, use Group Policy (best practice) or Terminal Services Configuration to modify them. For more information, see Delete or retain temporary folders when exiting or Use separate temporary folders for each session.

# Terminal Services User Properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services user properties**

This topic covers managing the information contained in the Terminal Services Extension to Local Users and Groups and Active Directory Users and Computers. The following topics are included:

- For information about best practices, see Terminal Services User Properties Best practices.
- For help with specific tasks, see Terminal Services User Properties How To....
- For general background information, see Terminal Services User Properties Concepts.

# Terminal Services User Properties Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Use Terminal Services-specific groups**

  Create User Groups that are specifically for Terminal Services users. Windows Server 2003 family operating systems contain a default User group called Remote Desktop Users, which has been specifically created to manage Terminal Server users. Maintaining users through groups is much easier and less time consuming than managing users individually.

- **Use Terminal Services-specific profiles**

  Assign a separate profile for logging onto Terminal Services. Many common options, such as screen savers and animated menu affects, which are stored in profiles, are not wanted when you use Terminal Services. Assigning a specific profile allows users to get the most out of the system they are working with without expending additional server resources. For information about assigning a Terminal Services specific profile, see Change a user's Terminal Services profile path.

- **Use mandatory profiles**

  Use a mandatory Terminal Services profile that is created to suit the needs of all of the different types of clients and that provides the best server performance. Be aware that 16-bit computers and Windows-based Terminals might not support some screen resolutions.

- **Set time limits**

  Setting limits on the duration of client connections can improve server performance. You can set the limits on how long a session lasts, how long a disconnected session is allowed to remain active on the server, and the time allowed for a session to remain connected, yet idle. For information about setting session limits, see Sessions.

- **Use Starting program**

  If you have users who need access to only one application on the terminal server, use the Starting program option. This can be done for all users by using Terminal Services Configuration or on a per-user basis by using either the Terminal Services Extension to Local Users and Groups or Active Directory Users and Computers. For information about using Starting program, see Specify a program to start on session connection.

- **Create preconfigured connection files for users or groups of users**

  To make connecting to Terminal Services easier, you can supply users with preconfigured connection files. Collections of connection files can also be made either for different departments within your organization or for different job titles. preconfigured connection files are created using Remote Desktop Connection. For more information, see Managing Terminal Services connection files.

# Terminal Services User Properties How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**


- Add the Terminal Services Extension to Local Users and Groups
- Add the Terminal Services Extension to Active Directory Users and Computers
- Change a user's Terminal Services profile path
- Change the path to a user's home directory
- Deny a user permissions to log on to terminal servers
- Set the maximum duration of an active session
- Set the maximum time a disconnected session remains active
- Set the maximum idle time for a session
- Change settings for when a session limit is reached or a connection is broken
- Change settings for reconnecting disconnected sessions
- Specify a program to start on session connection
- Connect Citrix ICA client drives at logon
- Connect client printers at logon
- Default to the client's main printer
- Enable remote control
- Require user's permission to observe or control sessions
- Set the level of control over a user's sessions

# Add the Terminal Services Extension to Local Users and Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add the Terminal Services Extension to Local Users and Groups**

1. Open Microsoft Management Console.
2. On the **File** menu, click **Add/Remove Snap-in**. The **Add/Remove Snap-in** dialog box appears.
3. Click **Add**. The **Add Standalone Snap-in** dialog box appears.
4. Select **Local Users and Groups**, and then click **Add**. The **Choose Target Machine** dialog box appears.
5. Select the computer to which to add this console, and then click **Finish**.
6. In the **Add Standalone Snap-in** dialog box, click **Close**.
7. In the **Add/Remove Snap-in** dialog box, click the **Extensions** tab.
8. Click **Terminal Services - Extension**, and then click **OK**.

## Notes

- You must be logged on as a computer administrator to complete this procedure.
- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- The Terminal Services extension is already included in Local Users and Groups if Terminal Server is installed on the server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Managing Terminal Services users
Managing Terminal Services users with Group Policy

# Add the Terminal Services Extension to Active Directory Users and Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add the Terminal Services Extension to Active Directory Users and Computers**

1. Open Microsoft Management Console.
2. On the **File** menu, click **Add/Remove Snap-in**. The **Add/Remove Snap-in** dialog box appears.
3. Click **Add**. The **Add Standalone Snap-in** dialog box appears.
4. Select **Active Directory Users and Computers**, click **Add**, and then click **Close**.
5. In the **Add/Remove Snap-in** dialog box, click the **Extensions** tab.
6. Click **Terminal Services - Extension**, and then click **OK**.

## Notes

- You must be logged on as a computer administrator to complete this procedure.
- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- To complete this procedure, the computer must be joined to an Active Directory domain, and you must have permission to edit the Group Policy object.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Managing Terminal Services users
Managing Terminal Services users with Group Policy

# Change a user's Terminal Services profile path

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a user's Terminal Services profile path**

- Using Group Policies (best practice)
- Using Terminal Services Extension to Local Users and Groups

**Using Group Policies (best practice)**

1. Choose from the following:

   - For a domain user account, from a Domain Controller, open Active Directory Users and Computers.

     In the console tree, right click the domain node, and then click **Properties**. In the Group Policy tab, double-click the Group Policy object, or click **Edit**.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
   - For a local user account, open Group Policy.

2. In the Group Policy console tree, click **Terminal Services** under **Computer Configuration**.

   **Where?**

   - Local Computer Policy/Computer Configuration/Administrative Templates/Windows Components/Terminal Services

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

3. Double-click the **Set path for TS Roaming Profiles** setting, and then click **Enabled**.

4. In the **Profile path** box, type the path for Terminal Services roaming profiles, and then click **OK**.

**Important**

- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

**Using Terminal Services Extension to Local Users and Groups**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which users are located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user whose profile path you want to change.

3. On the **Terminal Services Profile** tab, under **Profile Path:**, type the new path to the user's profile, and then click **OK**.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Terminal Server Best practices
Terminal Services Profile

# Change the path to a user's home directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the path to a user's home directory**

- Using Group Policies (best practice)
- Using Terminal Services Extension to Local Users and Groups

**Using Group Policies (best practice)**

1. Choose from the following:

   - For a domain user account, from a Domain Controller, open Active Directory Users and Computers.

     In the console tree, right click the domain node, and then click **properties**. In the Group Policy tab, double-click the Group Policy object, or click **Edit**.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
   - For a local user account, open Group Policy.

2. In the Group Policy console tree, click **Terminal Services** under **Computer Configuratio**n.

   **Where?**

   - Local Computer Policy/Computer Configuration/Administrative Templates/Windows Components/Terminal Services

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

3. Double-click the **TS User Home Directory** setting, and then click **Enabled**.

4. In the **Location** box, click the location for the Terminal Services home directory.

5. In the **Home Dir Root Path** box, type the root path for the Terminal Services home directory.

6. If the root path is on the network, click the drive letter for the mapped drive in the **Drive Letter** box. Click **OK**.

   **Important**

- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

**Using Terminal Services Extension to Local Users and Groups**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user whose home directory you want to change.

3. On the **Terminal Services Profile** tab, do one of the following:

   - If the Terminal Services home folder is on the local server, click **Local path**, and then type the path to the home directory.

   - If the Terminal Services home directory is on a network share, click **Connect**, select a drive to connect, and then type the network path.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Terminal Services Profile

# Deny a user permissions to log on to terminal servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To deny a user permissions to log on to terminal servers**

1. Choose from the following:

   - For a domain user account, open **Active Directory Users and Computers**.

     In the console tree, expand the domain node, and then click the folder in which the user profile is located.

   **Note**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open **Computer Management (Local)**.

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Note**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

2. Double-click the name of the user whose settings you want to change.

3. On the **Terminal Services Profile** tab, select the **Deny this user permissions to log on to any terminal server** check box, and then click **OK**.

**Important**

Clearing the **Deny this user permissions to log on to any terminal server**check box does not automatically allow the user permissions to log on to a terminal server. In addition to clearing this check box, you must also ensure that the user has been added to the Remote Desktop Users group on the appropriate terminal server, and that Remote Desktop has been enabled on the terminal server. For information about adding users to the Remote Desktop Users group, enabling Remote Desktop, and enabling users to connect to a terminal server, see Related Topics.

1. **Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Terminal Services Profile
Add users to the Remote Desktop Users group
Enable or disable Remote Desktop
Enabling users to connect remotely to the server

# Set the maximum duration of an active session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the maximum duration of an active session**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account that you want to change.

3. On the **Sessions** tab, in **Active session limit**, select the amount of time to wait before the connection times out, and then click **OK**.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Terminal Server Best practices
Sessions

# Set the maximum time a disconnected session remains active

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the maximum time a disconnected session remains active**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users** (console tree location is: Computer Management\System Tools\Local Users and Groups\Users).

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account for which you want to change settings.

3. On the **Sessions** tab, in **End a disconnected session**, select the amount of time you want to elapse before the disconnected session is reset, and then click **OK**.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Terminal Services users
Terminal Server Best practices
Sessions

# Set the maximum idle time for a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the maximum idle time for a session**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users** (console tree location is: Computer Management/System Tools/Local Users and Groups/Users).

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

2. Double-click the user account for which you want to change settings.

3. On the **Sessions** tab, in **Idle session limit**, select the amount of time you want to elapse before the idle session times out, and then click **b**.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Terminal Services users
Terminal Server Best practices
Sessions

# Change settings for when a session limit is reached or a connection is broken

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change settings for when a session limit is reached or a connection is broken**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account for which you want to change settings.

3. On the **Sessions** tab, under **When a session limit is reached or connection is broken**, click **Disconnect from session** to allow the session to be reconnected.

   Or, click **End session** to end the session. Click **OK**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Terminal Server Best practices
Sessions

# Change settings for reconnecting disconnected sessions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change settings for reconnecting disconnected sessions**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account for which you want to change settings.

3. On the **Sessions** tab, under **Allow reconnection**, click **From any client** to allow connection to that session from any client computer.

   Or, clic**k From originating client only** to allow only the client computer where the session was created to reconnect to that session. This option is supported only for Citrix ICA clients. Click **OK**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Terminal Server Best practices
Sessions

# Specify a program to start on session connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify a program to start on session connection**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account for which you want to change the Starting program.

3. On the **Environment** tab, under **Starting program**, select the **Start the following program at logon** check box.

4. In **Program file name**, type the name of the program to start when the user logs on. You can also specify a working directory by typing the path in **Start in**. Click **OK**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- Some client option settings are not applied when a connection is made to the console session of a remote computer running Windows XP. For example, if the client has specified a program to start upon connection to a remote computer, that program does not start if the connection is made to the console session of the remote computer. Instead, the default desktop appears. If the connection is made to a session other than the console session, the specified program starts as

expected. The console session is the only session available for Remote Desktop on Windows XP Professional, so the default desktop always appears when a connection is made, even if the client has specified a different program.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Environment

# Connect Citrix ICA client drives at logon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect Citrix ICA client drives at logon**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account for which you want to connect client drives.

3. On the **Environment** tab, under **Client devices**, select the **Connect client drives at logon** check box, and then click **OK**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- This procedure is for Citrix ICA clients only. For clients that use the Remote Desktop Protocol (RDP), client drive redirection is supported if the server is running one of the Windows Server 2003 family operating systems and the client is running Remote Desktop Connection.

- In Windows Explorer, client drives appear in the session folder tree in the form `\\TSClient\<driveletter>$` if drive redirection has been specified as a Remote Desktop Connection option on the client.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Environment
Making local resources available to a terminal server

# Connect client printers at logon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To connect client printers at logon

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account for which you want to connect client printers.

3. On the **Environment** tab, under **Client devices**, select the **Connect client printers at logon** check box, and then click **OK**.

### Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Providing client access to local printers
Environment

# Default to the client's main printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To default to the client's main printer

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account for which you want to change the printer default.

3. On the **Environment** tab, under **Client devices**, select the **Default to main client printer**, and then click **OK**.

### Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Terminal Services users

Providing client access to local printers

Environment

# Enable remote control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To enable remote control

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account for which you want to enable remote control.

3. On the **Remote control** tab, select the **Enable remote control** check box, and then click **OK**.

### Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
  This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Terminal Services users

Remote control

# Require user's permission to observe or control sessions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To require user's permission to observe or control sessions**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user is located.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user whose permission you want to require.

3. On the **Remote control** tab, select the **Require user's permission** check box, and then click **OK**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Remote control

# Set the level of control over a user's sessions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the level of control over a user's sessions**

1. Choose from the following:

   - For a domain user account, open Active Directory Users and Computers.

     In the console tree, expand the domain node, and then click the folder in which the user account is located.

   **Important**

   - To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

   - For a local user account, open Computer Management (Local).

     In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

   **Important**

   - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

2. Double-click the user account for which you want to set the level of control.

3. On the **Remote control** tab, under **Level of control**, click **View the user's session** to allow only session observation.

   Or, click **Interact with the session** to allow full remote control of the session. Click **OK**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**. This procedure works only if the server has been promoted to a domain controller.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Terminal Services users
Remote control

# Terminal Services User Properties Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Managing Terminal Services users
- Managing Terminal Services users with Group Policy
- Terminal Services users

# Managing Terminal Services users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Terminal Services Users**
Each user who logs on to a Terminal Services session must have a user account either on the server or in a domain on the network that the server is on. The Terminal Services user account contains additional information about the user that determines when users log on, under what conditions, and how specific desktop settings are stored. Windows Server 2003 family operating systems contain a built-in User group called Remote Desktop Users, which is used to manage Terminal Services users.

**About the Remote Desktop Users group**
When you install one of the Windows Server 2003 family operating systems, the Remote Desktop Users group is one of the built-in user groups on your computer. Members of this group have the same access as members of the Users group, but they have the additional ability to log on remotely to the computer.

By default, this group is not populated when you install Terminal Server on your computer. You must choose the users and groups that you want to have permission to log on remotely to the terminal server, and manually add them to the Remote Desktop Users group. This increases the security of remote connections, and also allows you to install any required programs before users start connecting to the terminal server.

The **Select Remote Users** button on the **Remote** tab of the **System Properties** dialog box allows you to add users to the Remote Desktop Users group. However, if a server running a Windows Server 2003 family operating system is being used as a domain controller in a Windows 2000 domain, this button is disabled. This domain controller can have domain groups, but not local groups, so in this situation you cannot use this button to add users to the Remote Desktop Users group.

For information on how to add users to the Remote Desktop Users group, see Add users to the Remote Desktop Users group.

# Managing Terminal Services users with Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Terminal Services Users with Group Policy**
You can use Group Policy to configure Terminal Services user profile information. Group Policy configures these options on a per-computer basis, for easy configuration of all users who connect to a specific computer. For environments where individual user properties are not available to local administrators, using Group Policy provides an alternative to the Terminal Services Extension to Local Users and Groups.

**Important**

- You should thoroughly test any changes you make to Group Policy settings before applying them to users or computers. For more information on testing policy settings, see Resultant Set of Policy.

Use the **Set path for TS Roaming Profiles** setting to configure a network path for Terminal Services user profiles. For more information, see Change a user's Terminal Services profile path.

Use the **TS User Home Directory** setting to configure a local or network path as the root of the home directory for Terminal Services users. For more information, see Change the path to a user's home directory.

For more information about Terminal Services profiles, see Terminal Services Profile.

For general information about using Terminal Server Group Policy, see Configuring Terminal Services with Group Policy.

# Managing Users with the Terminal Services Extension to Local Users and Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Users with the Terminal Services Extension to Local Users and Groups**
This topic covers information about managing Terminal Services users with the Terminal Services Extension to Local Users and Groups. The following topics are included:

- For help with specific tasks, see Terminal Services User Properties How To....
- For general background information, see Terminal Services User Properties Concepts.

# Terminal Services users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services users**
Each user who logs on to a terminal server must have a user account either on the server or in a domain on the network that the terminal server is on. The Terminal Services user account contains additional information about the user that determines when users log on, under what conditions, and how specific desktop settings are stored. Windows Server 2003 family operating systems contain a built-in User group called Remote Desktop Users, which has been specifically created to manage Terminal Server users.

For more information about the Remote Desktop Users group and for managing users with Group Policies, see Managing Terminal Services users.

In addition to the Terminal Services Group Policies, User accounts can be set up and managed for a specific computer using Local Users and Groups or for a domain of computers using Active Directory Users and Computers.

**Note**

- User settings set with Local Users and Groups and with Active Directory Users and Computers are overridden by Group Policy settings at both the computer and user level. For more information, see Terminal Services Configuration Overview.

The following topics describe the Terminal Services Extension tabs that appear in the **Properties** dialog box for a user:

- Environment
- Sessions
- Remote control
- Terminal Services Profile

# Environment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Environment**
The Terminal Services Extension **Environment** tab on the **Properties** dialog box contains settings for creating the client working environment.

**Starting program**
If a starting program is specified in the user's settings, that program opens whenever the user connects to a terminal server. It is the only application that can be used by the user. When the user closes the application, the connection to the terminal server closes, as though the user has logged off of the terminal server.

You can also use Terminal Services Configuration to specify a starting program that is the only application available to all users who connect to the terminal server. For more information about setting the starting program on a connection basis, see Specify a program to start automatically when a user logs on.

**Client devices**
You can configure Terminal Services to connect local client printers at logon. When the client logs on to the server, the local printer is detected and the appropriate printer driver is installed on the terminal server. If multiple printers are connected, you can also default all print jobs to the main client printer. For more information, see Connect client printers at logon or Providing client access to local printers.

# Sessions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sessions**

The **Sessions** tab of the Terminal Services Extension provides settings for limiting the time length of sessions based on their current state (active, idle, or disconnected). You can also specify what action to take when a session has reached a time limit.

| Time-out setting | Description |
|---|---|
| End a disconnected session | Specifies the maximum duration that a disconnected session is retained. The session is reset and can no longer be reconnected after the time limit has expired. For more information, see Set the maximum time a disconnected session remains active. |
| Active session limit | Specifies the maximum connection duration. When the time limit is reached, the session will either be disconnected, leaving the session active on the server, or reset. Set the maximum duration of an active session. |
| Idle session limit | Specifies the maximum idle time (time without connection activity) allowed before the session is disconnected or reset. The session is disconnected or reset when the interval elapses without any activity at the connection. Set the maximum idle time for a session. |

**Note**

- If **Never** is selected, the timer is disabled.

When the connection to the server is broken for any reason, including a request, a connection error, or an idle or active limit is reached, the client session is either placed into the disconnected state or it is reset. A session that is reset can no longer be activated by a user and is closed. A disconnected session is left active on the server and is available for reconnection. For more information, see Change settings for when a session limit is reached or a connection is broken.

**Disconnected sessions**

By default, a terminal server allows you to reconnect to a disconnected session from any computer. However, to allow only the original client computer to reconnect the session, you can specify that a user can only reconnect to a disconnected session from the computer where the session originated. This option is supported only for Citrix ICA clients that provide a serial number when connecting.

# Remote control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote control**
You can monitor the actions of a client logged on to a terminal server by using remote control from another session. Remote control allows you to either observe or actively control a client session. If you choose to actively control a client session, you will be able to input keyboard and mouse actions to the session. You can warn a client that you want to remotely control their session by choosing to display a message on the client that asks permission to view or take part in the session. You can use either Local Users and Groups (for local users) or Active Directory Users and Computers (for domain users) to enable Remote control for a user account. For more information about enabling remote control, see Enable remote control.

**Notes**

- You should use care when assigning permissions on connections. For more information on permissions, see Controlling connection access.
- The system console cannot remotely control a client session.
- The computer you are logged on to must be capable of supporting the video resolution used at the remotely controlled client session. Otherwise, the operation fails.
- To end remote control, press `CTRL+ASTERISK (*)`. You can define another hot key to perform this function.

# Terminal Services Profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services profile**
A Terminal Services profile to assign each user a profile is specifically created for terminal server sessions. Administrators can then create user profiles tailored to the Terminal Services environment.

The Terminal Services profile can be used to restrict access to applications by removing them from the user's **Start** menu. Administrators can also create and store network connections to printers and other resources for use during user sessions.

For more information about profiles, see User Profiles.

**Terminal Services home folder**
You can specify a path to a home folder to be used for terminal server sessions. This can be either a local folder or a network share.

**Allow logon to terminal server**
You can grant access to a specific terminal server on a per-user basis. If this option is left disabled, the user is not allowed to log onto any terminal servers.

# Terminal Services Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services Manager**
Use Terminal Services Manager to view information about terminal servers in trusted domains, including all sessions, users, and processes for each terminal server. You can also use this tool to perform various actions to manage the server.

- For help with specific tasks, see Terminal Services Manager How To....
- For general background information, see Terminal Services Manager Concepts.

# Terminal Services Manager How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install Terminal Server
- Manage Users, Sessions, and Processes
- Work with Terminal Servers

# Install Terminal Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Terminal Server**

1. Open Configure Your Server Wizard.
2. Click **Next** and follow the instructions on the **Preliminary Steps** screen.
3. Click **Next**, and on the **Server Role** screen, select **Terminal Server**.
4. Click **Next** and follow the instructions in the wizard.
5. Some programs might not work correctly when Terminal Server is installed. If so, you must reinstall these programs for multisession access using **Add or Remove Programs** after Terminal Server is installed. For more information, see Program Compatibility.

**Warning**

- Terminal Server Licensing is a required component that licenses clients on a terminal server. You must install Terminal Server Licensing or your terminal server will stop accepting connections from unlicensed clients 120 days from the date of the first client logon.

  You should not install Terminal Server Licensing on the same computer with Terminal Server.

  For more information on Terminal Server Licensing, see Terminal Server Licensing.

**Important**

- When you install Terminal Server using the Configure Your Server Wizard, as a security precaution the Internet Explorer Enhanced Security Configuration is set by default to prevent Administrators from downloading files or running ActiveX controls from the Web. If you want to change this setting to either permit all users to download files or run ActiveX controls or prevent all users from doing so, use **Add/Remove Windows Components** in **Add or Remove Programs**. This setting applies to servers running one of the Windows Server 2003 family operating systems. For more information, see Terminal server role: Configuring a terminal server and Internet Explorer Enhanced Security Configuration.

**Important**

- You must use the Remote Desktop Users group to control remote connection permissions for Terminal Server. For more information, see Managing Terminal Services users and Add users to the Remote Desktop Users group.

**Important**

- It is not necessary to install Terminal Server for remote administration of your server. Remote Desktop for Administration is installed by default. To use Remote Desktop for Administration, you must first enable remote connections. For more information on enabling remote connections, see Enable or disable Remote Desktop and Enabling users to connect remotely to the server.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For

more information, see Default local groups, Default groups, and Using Run as.

- To open the Configure Your Server Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Configure Your Server Wizard**.

- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.

- You can also use **Add or Remove Programs** to install Terminal Server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Terminal Server Licensing
Activate a Terminal Server License Server
Install Client Access Licenses
Enable or disable Remote Desktop

# Manage Users, Sessions, and Processes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage users, sessions, and processes**

- Connect to another session
- Disconnect a Terminal Server session
- Log off a user from a session
- Send a message to a user
- Reset a session
- View status information for a session
- Remotely control a session
- End a process

# Connect to another session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to another session**

1. Open Terminal Services Manager.
2. Right-click the session to which you want to connect, and then click **Connect**.

   You are connected to the new session and disconnected from your previous session.

## Notes

- You can always connect to a session in which you are logged on with the same user account. To connect to another user's session, you must have either Full Control or User Access permission.
- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- You can connect to another session only from within an existing session. You must open Terminal Services Manager, or use the **tscon** command, from within a session to be able to connect to another session.
- You can connect to a session only if it is in either an active or disconnected state.
- You cannot connect to another session from the console session.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disconnect a Terminal Server session
Viewing the sessions associated with a server or domain
Tscon
Managing Permissions on Connections

# Disconnect a Terminal Server session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect from a session**

1. Open Terminal Services Manager.
2. Right-click the session you want to disconnect, and then click **Disconnect**.

## Notes

- You can always disconnect from your own sessions, but you must have Full Control permission to disconnect a session belonging to another user.
- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- When you disconnect a session, the session remains attached to the terminal server in the disconnected state and any applications that are running continue to run. The applications remain running when the session is reconnected, with no loss of data.
- You can also use the **tsdiscon** command to disconnect a session.
- You can reset a session that is in the disconnected state, deleting the session from the server. You must have Full Control permission to do this.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to another session
Reset a session
Tsdiscon
Managing Permissions on Connections

# Log off a user from a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To log off a user from a session**

1. Open Terminal Services Manager.
2. On the **Users** tab, right-click the name of the user you want to log off from a session, and then click **Log Off**.

   **Caution**

   - Logging off a user from a session without warning can result in loss of data at the user's session. Instead, send a message to the user using **Send Message** on the **Actions** menu to warn the user before taking this action.

## Notes

- You can always log off from your own sessions, but you must have Full Control permission to log off another user from a session.
- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- When you log off a user, all processes end and the session is deleted from the server.
- You can also use the **logoff** command to log off a user from a session.
- You cannot log off a user from the console session.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Send a message to a user
Logoff
Managing Permissions on Connections

# Send a message to a user

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send a message to a user**

1. Open Terminal Services Manager.
2. Right-click the session or user to which you want to send a message, and then click **Send Message**.

   The **Send Message** dialog box appears.
3. In **Message title**, type the title of the message.
4. In **Message**, type the information you want to send to the user. To start a new paragraph, press CTRL+ENTER.
5. To send the message, click **OK**.

## Notes

- You must have Message permission in order to send messages. For more information on connection permissions, see Managing Permissions on Connections.
- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- Use this feature to notify users of impending disconnection, server status, or other system information.
- You can send messages only to users whose sessions are in the active or connected state, including the console session.
- You can also use the **msg** command to send messages to users.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msg
Managing Permissions on Connections

# Reset a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reset a session**

1. Open Terminal Services Manager.
2. Right-click the session you want to reset, and then click **Reset**.

   The session is immediately deleted from the server.

## Notes

- You can always reset your own sessions, but you must have Full Control permission to reset a session belonging to another user.
- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- Be aware that resetting a user's session without warning can result in loss of data on the session.
- Reset a session only when it malfunctions or appears to have stopped responding.
- If you reset a listener session, all sessions that use the connection are reset.
- You can also use the **reset session** command to reset a session.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disconnect a Terminal Server session
Log off a user from a session
Reset session
Managing Permissions on Connections

# View status information for a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view status information for a session**

1. Open Terminal Services Manager.
2. Right-click the session for which you want status information, and then click **Status**.

   The session's **Status** information is displayed.
3. To refresh this information, click **Refresh Now**.

   To reset counters, click **Reset Counters**.

## Notes

- You can always view status information for your own sessions. To view status information for a session belonging to another user, you must have either Full Control or User Access permission, or be specifically assigned the Query Information advanced permission. For more information on connection permissions, see Managing Permissions on Connections.
- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- You cannot view status information for the console or listener sessions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to another session
Remotely control a session
Managing Permissions on Connections

# Remotely control a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remotely control a session**

1. Open Terminal Services Manager.
2. Right-click the session you want to monitor, and then click **Remote Control**.

   The **Remote Control** dialog box appears.
3. In **Hot key**, select the keys you want to use to end a remote control session, and then click **OK**. The default hot key is CTRL+* (using * from the numeric keypad only).

   When you want to end remote control, press CTRL+* (or whatever hot key you have defined).

## Notes

- You must have Full Control permission to remotely control another session.
- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- To configure remote control settings for a connection, use Terminal Services Configuration. Remote control can also be configured on a per-user basis by using the Terminal Services Extension to Local Users and Groups and Active Directory Users and Computers.
- Before monitoring begins, the server warns the user that the session is about to be remotely controlled, unless this warning is disabled. Your session might appear to be frozen for a few seconds while it waits for a response from the user.
- When you enter the remote control session, your current session shares every input and output with the session you are monitoring.
- Your session must be capable of supporting the video resolution used at the session you are remotely controlling or the operation fails.
- The console session can neither remotely control another session nor can it be remotely controlled by another session.
- You can also use the **shadow** command to remotely control another session.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using remote control
Disconnect a Terminal Server session
End a process
Shadow
Managing Permissions on Connections

# End a process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To end a process**

1. Open Terminal Services Manager.
2. On the **Processes** tab, under the **User** column, right-click the process you want to end, and then click **End Process**.

## Notes

- You must have Full Control permission to end a process.
- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- Be aware that ending a process without warning can result in loss of data at the user's session.
- You might need to end a process because the application is not responding.
- You can also use the **tskill** command to end a process.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Tskill
Terminal Services Manager How To...

# Work with Terminal Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with terminal servers**

- Connect to terminal servers
- Disconnect from terminal servers
- Find all terminal servers in a domain
- Add a terminal server to the Favorites list
- Remove a terminal server from the Favorites list

# Connect to terminal servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to terminal servers**

1. Open Terminal Services Manager.
2. Expand the domain containing the servers to which you want to connect.
3. Do one of the following:
   - To connect to a specific terminal server, right-click the server, and then click **Connect**.
   - To connect to a computer that is not displayed in the navigation pane of the Remote Desktops snap-in, right-click **All Listed Servers**, click **Connect to Computer**, specify a computer name, and then click **OK**.

## Notes

- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- You should connect to only one terminal server at a time. When a server is connected, Terminal Services Manager queries the server for information about its sessions and processes. Connecting to more than one server at a time can overload system resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disconnect from terminal servers
Find all terminal servers in a domain

# Disconnect from terminal servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect from terminal servers**

1. Open Terminal Services Manager.
2. Expand the domain containing the servers from which you want to disconnect.
3. Do one of the following:
   - To disconnect from a specific terminal server, right-click the server, and then click **Disconnect**.
   - To disconnect from all terminal servers in the domain, right-click the domain, and then click **Disconnect from All Servers in Domain**.

## Notes

- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- To disconnect from all servers in all domains in your network, right-click **All Listed Servers**, and then click **Disconnect from All Servers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to terminal servers
Find all terminal servers in a domain

# Find all terminal servers in a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find all terminal servers in a domain**

1. Open Terminal Services Manager.
2. Right-click the domain whose terminal servers you want to find, and then click **Refresh Servers in Domain**.

## Notes

- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.
- Depending on the number of servers on your network, this may take an extended period of time.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disconnect from terminal servers
Connect to terminal servers

# Add a terminal server to the Favorites list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a terminal server to the Favorites list**

1. Open Terminal Services Manager.
2. Expand the domain containing the server you want to add to the Favorites list.
3. Right click the server, and then click **Add To Favorites**.

### Note

- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disconnect from terminal servers
Find all terminal servers in a domain

# Remove a terminal server from the Favorites list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a terminal server from the Favorites list**

1. Open Terminal Services Manager.
2. Expand the **Favorite servers** node.
3. Right click the server, and then click **Remove From Favorites**.

### Note

- To open Terminal Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Manager**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disconnect from terminal servers
Find all terminal servers in a domain

# Terminal Services Manager Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides background information about administering servers that have Terminal Server installed.

- Terminal Services Manager Overview
- Using Terminal Services Manager
- The Terminal Services Manager Window

# Terminal Services Manager Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal Services Manager overview**

Use Terminal Services Manager to view information about terminal servers that reside in trusted domains. Use this tool to monitor users, sessions, and applications on each terminal server, and to carry out assorted actions to manage the server. See individual procedures in Terminal Services Manager How To... for the permission levels required to perform each task.

When a user creates a session by connecting to a terminal server from a client computer, the session appears in the **Session** list in Terminal Services Manager. In addition, the name of the user who logs on by using the session appears in the **Users** list. Any applications run in the user's session can be monitored on the **Processes** list. Therefore, you can oversee all users, sessions, and processes on a terminal server from one location.

Terminal Services Manager also allows you to manage a terminal server by using the commands on the **Actions** menu. For more information, see Using Terminal Services Manager.

**Console session**

In Terminal Services Manager, the system console session is identified as session 0 in the **Session** list when you connect to a terminal server. A console session is defined as the session you connect to at the physical console of the remote computer, as though you were logging on locally instead of remotely. You can send a message to the console session, but you cannot perform any of the other administrative actions on it.

**Listener sessions**

Listener sessions are different from regular sessions. These sessions listen for and accept new Remote Desktop Protocol (RDP) client connections, thereby creating new sessions for the client requests. If you have configured more than one connection in Terminal Services Configuration, several listener sessions are available.

You have the option to reset a listener session. However, this is not recommended, because doing so resets all sessions that use the same Terminal Services connection. Resetting a user's session without warning can result in loss of data at the client.

**Idle sessions**

To optimize the performance of a terminal server, idle sessions are initialized by the server before client connections are made. These sessions are available to clients for connection. Two idle sessions are created by default. User sessions can also be in idle state. For information on configuring session limits, see Configuring session limits.

# Using Terminal Services Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Terminal Services Manager**

Terminal Services Manager is used to monitor the sessions, users, and processes of terminal servers in trusted domains. In addition to monitoring, you can manage a terminal server by using the actions provided by Terminal Services Manager. Actions are found on the **Actions** menu and are described below. Most require you to be logged in as a member of the Administrators group to perform them.

| **Action** | **Description** | **Permission Level Required** |
|---|---|---|
| Connect | Enables a user to connect to another session. Be aware that connecting to a session that is currently being used by another user might be very disruptive to that user. When you connect to another session, you are disconnected from your previous session. If you have created more than one session on a server, you can use this option to switch between these sessions. You cannot connect to another session from the system console. For more information, see Connect to another session. | User Access or Full Control |
| Disconnect | Disconnects a user from a session. The session remains attached to the terminal server in the *disconnected* state and currently running applications continue to run. When you attempt to reconnect to the server, you are reconnected to the same session from which you disconnected, even if you are reconnecting from a different computer. Applications that were left open when you disconnected remain running when you reconnect to the session, with no loss of data. This is very useful if you change locations often (for example, from work to home). For more information, see Disconnect a Terminal Server session. | Full Control |
| Send Message | Enables a user to send a message to another user's session. For example, an administrator might want to send a message to a user before disconnecting or logging off the user from a session. For more information, see Send a message to a user. | User Access or Full Control |
| Remote Control | Enables a user to observe or remotely control another user's session, allowing you to monitor the activities at the session and interact with it as needed. A user whose session you want to monitor can be warned before monitoring begins. Remote control is configured either with Terminal Services Configuration or, on a per-user basis, with the Terminal Services Extension to Local Users and Groups and Active Directory Users and Computers. You cannot remotely control another session from the console session. For more background information, see Using remote control. For information about using remote control, see Remotely control a session. | Full Control |
| Reset | Enables you to delete a session instantly. Be aware that resetting a user's session without warning can result in loss of data at that session. Reset a session only when it malfunctions or appears to have stopped responding. Resetting a listener session resets all sessions that use the connection. For more information, see Reset a session. | Full Control |
| Status | Enables you to monitor session-related counters, such as incoming and outgoing bytes and frames. Status information for console and listener sessions cannot be displayed. For more information, see View status information for a session. | User Access or Full Control |
| Log Off | Enables you to log off a user from a session on the server. Be aware that logging off a user without warning can result in loss of data at the user's session. When you log off a user, all processes end and the session is deleted from the server. For more information, see Log off a user from a session. | Full Control |
| End Process | Enables you to end a process that is running at a user's session. This is useful when an application has stopped responding. Be aware that ending a process without warning can result in loss of data at the user's session. For more information, see End a process. | Full Control |

# The Terminal Services Manager Window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the Terminal Services Manager window**

The Terminal Services Manager window has a two-pane view. The navigation pane displays the domains, the terminal servers, and the sessions on the network. The tabs in the details pane display the results of the item selected in the navigation pane. For example, when you click a server in the navigation pane, a series of tabs appear in the details pane displaying lists of users, sessions, and processes associated with that server.

Typically, you perform tasks by selecting an item in the navigation pane and then taking action on that item in the details pane by using shortcut menus and toolbars.

For information about monitoring terminal servers, click the following:

- Viewing the users on a server or domain
- Viewing the sessions associated with a server or domain
- Viewing the processes associated with a server or domain
- Viewing the processes associated with a session
- Viewing client information
- Viewing session states

# Viewing the users on a server or domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing the users on a server or domain**

Using Terminal Services Manager, you can view information about the user's session, the state of the session, and other properties for each user on a terminal server. To view user information on a terminal server, click the server or domain in the navigation pane, and then click the **Users** tab.

Each field on the **Users** tab is described in the following table.

| Field | Description |
|---|---|
| Server | Displays the servers to which users are logged on. This field is shown only when you click a domain in the navigation pane. |
| User | Displays all users associated with the server or domain. |
| Session | Displays the sessions associated with a server or all servers in the domain. All sessions are displayed regardless of status. |
| ID | Displays the numeric ID that identifies the session to the server. |
| State | Displays the status of a session. Session states include *Active*, *Connected*, *ConnectQuery*, *RemoteControl*, *Listen*, *Disconnected*, *Idle*, *Down*, and *Init*. For more information, see Viewing session states. |
| Idle time | Displays the number of minutes that have elapsed since the last keyboard or mouse input from a session. |
| Logon time | Displays the time at which the user logged on, if applicable. |
| Comment | Provides additional information about the session, such as its location. This field is optional. |

# Viewing the sessions associated with a server or domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing the sessions associated with a server or domain**

Using Terminal Services Manager, you can view information about the user, state, and other properties of each session on a terminal server. To view a session on a terminal server, click the server or domain in the navigation pane, and then click the **Sessions** tab.

Each field on the **Sessions** tab is described in the following table.

| Field | Description |
|---|---|
| Server | Displays the servers with which the sessions are associated. This field is shown only when you click a domain in the navigation pane. |
| Session | Displays the sessions associated with a server. All sessions are displayed regardless of status. |
| User | Displays the users associated with the server or domain. |
| ID | Displays the numeric ID that identifies the session to the server. |
| State | Displays the status of a session. Session states include *active*, *conn*, *connq*, *shadow*, *listen*, *disc*, *idle*, *down*, and *init*. For more information, see Viewing session states. |
| Type | Specifies the type of client using the session or if the session is the system console. |
| Client name | Specifies the name of the client computer using the session, if applicable. |
| Idle time | Displays the number of minutes that have elapsed since the last keyboard or mouse input from a session. |
| Logon time | Displays the time at which the user logged on, if applicable. |
| Comment | Provides additional information about the session, such as its location. This field is optional. |

# Viewing the processes associated with a server or domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing the processes associated with a server or domain**

Using Terminal Services Manager, you can view current information about all processes on a terminal server. To view the processes, click the server or domain in the navigation pane, and then click the **Processes** tab.

Each field on the **Processes** tab is described in the following table.

| Field | Description |
|---|---|
| Server | Displays the servers with which the processes are associated. This field is shown only when you click a domain in the navigation pane. |
| User | Displays the users associated with the processes. |
| Session | Displays the sessions associated with the processes. All sessions are displayed regardless of status. |
| ID | Displays the numeric ID that identifies the session to the server. |
| PID | Displays the numeric ID that identifies the process to the server. |
| Image | Specifies the executable program that created the process. |

For information on how to end a process, see End a process.

# Viewing the processes associated with a session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing the processes associated with a session**

Using Terminal Services Manager, you can view current information about all processes associated with a particular session. To view the processes, click the session in the navigation pane, and then click the **Processes** tab.

Each field on the **Processes** tab is described in the following table.

| Field | Description |
|---|---|
| ID | Displays the numeric ID that identifies the session to the server. |
| PID | Displays the numeric ID that identifies the process to the server. |
| Image | Specifies the executable program that created the process. |

For information on how to end a process, see End a process.

# Viewing client information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing client information**

Using Terminal Services Manager, you can view client information about a particular session. To view client information, click the session in the navigation pane, and then click the **Information** tab.

Each field on the **Information** tab is described in the following table.

| Field | Description |
|---|---|
| User name | The name of the person who is logged in to the session on the server. |
| Client name | The name of the client computer. |
| Client build number | The version of the software installed on the client computer |
| Client directory | The directory in which the client is installed. |
| Client product ID | Used by Citrix Metaframe. |
| Client Hardware ID | Used by Citrix Metaframe. |
| Client address | The IP address of the client. |
| Server buffers | The size of the server buffers. |
| Client buffers | The size of the client buffers. |
| Client Color Depth | The number of colors in the color palette used for the remote session. |
| Client Resolution | The video resolution of the remote session. |
| Modem Name | Used by Citrix Metaframe. |
| Encryption Level | The encryption level being used for this session. |
| Client License | Used by Citrix Metaframe. |

# Viewing session states

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing session states**

The following table describes each possible state of a session.

| Session state | Description |
|---|---|
| Active | The session is connected, and a user is logged on to the server. |
| Connected | The session is connected, but there is no user logged on to the server. |
| ConnectQuery | The session is in the process of connecting. If this state continues, it indicates a problem with the connection. |
| RemoteControl | The session is in the process of remotely controlling another session. |
| Listen | The session is ready to accept a client connection. |
| Disconnected | The user is disconnected from the session, but the session is still attached to the server and can be reconnected at any time. |
| Idle | The session is initialized and ready to accept a connection. To optimize the performance of a server, two default (idle) sessions are initialized before any client connections are made. |
| Down | The session failed to initialize correctly or could not be terminated, and is not available. If this state continues, it indicates a problem with the connection of the session. |
| Init | The session is in the process of initializing. |

# Client Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Client computers**
This section covers:

- Windows Installer
- Remote Administration
- Offline Files
- User Profiles
- Network Connections
- Connect clients to a printer

# Windows Installer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer**
Windows Installer makes it easy for you to manage the software installation and removal process, and modify, repair, and advertise existing applications.

- For help with specific tasks, see Windows Installer How To....
- For general background information, see Windows Installer Concepts.
- For problem-solving instructions, see Windows Installer Troubleshooting.

# Windows Installer How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Use Windows Installer packages with Windows Explorer
- Use Windows Installer packages with the Add or Remove Programs utility
- Use Windows Installer from the command line
- Configure Group Policy for Windows Installer

# Use Windows Installer packages with Windows Explorer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use Windows Installer packages with Windows Explorer**

1. Open Windows Explorer.
2. Right-click the Windows Installer package (.msi file), and then click **Install**, **Repair**, or **Uninstall**.

### Note

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Use Windows Installer packages with the Add or Remove Programs utility
Use Windows Installer from the command line

# Use Windows Installer packages with the Add or Remove Programs utility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use Windows Installer packages with Add or Remove Programs**

1. Open Add or Remove Programs in Control Panel.
2. Click **Change or Remove Programs**, click the program that you want to change or remove, and then do one of the following:
   - To change a program, click **Change**.
   - To remove a program, click **Remove**.

### Note

- To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Use Windows Installer packages with Windows Explorer
Use Windows Installer from the command line

# Use Windows Installer from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Windows Installer from the command line**


- Install a transform using the command line
- Install and advertise a product using a transform at the command line
- Install a package using the command line
- Remove a package using the command line
- Repair a package using the command line
- Advertise a package using the command line
- Create an administrative installation package from the command line
- Create an installation log file using the command line

# Install a transform using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a transform using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /i***package***TRANSFORMS=***TransformList*

| Value | Description |
|---|---|
| **/i** | Installs a product. |
| *package* | Specifies the name of the Windows Installer package file. |
| **TRANSFORMS=***TransformList* | Specifies what transform (.mst) files should be applied to the package. *TransformList* is a list of paths separated by semicolons. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Install a package using the command line

# Install and advertise a product using a transform at the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install and advertise a product using a transform at the command line**

1. Open Command Prompt.

2. Type:

   **msiexec***package* **/j**{**u** | **m**} **/t***TransformList*

| Value | Description |
|---|---|
| *package* | Specifies the name of the Windows Installer package file. |
| **/j** | Advertises a product. This command-line option ignores any property values typed at the command line. |
| **u** | Advertises to the current user. |
| **m** | Advertises to all users of this computer. |
| **/t** *TransformList* | Applies transform to advertised package and specifies what transform (.mst) files should be applied to the package. *TransformList* is a list of paths separated by semicolons. |

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Install a package using the command line

# Install a package using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a package using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /i** {*package* | *ProductCode*}

| Value | Description |
|---|---|
| **/i** | Installs a product. |
| *package* | Specifies the name of the Windows Installer package file. |
| *ProductCode* | Specifies the globally unique identifier (GUID) of the Windows Installer package. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Create an administrative installation package from the command line

# Remove a package using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a package using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /x** {*package* | *ProductCode*}

| Value | Description |
|---|---|
| **/x** | Uninstalls a product. |
| *package* | Specifies the name of the Windows Installer package file. |
| *ProductCode* | Specifies the globally unique identifier (GUID) of the Windows Installer package. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Repair a package using the command line

# Repair a package using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To repair a package using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec/f**[**p**][**o**][**e**][**d**][**c**][**a**][**u**][**m**][**s**][**v**] {*package* | *ProductCode*}

| Value | Description |
|---|---|
| **/fp** | Reinstalls only if the file is missing. |
| **/fo** | Reinstalls if the file is missing or if an older version is installed. |
| **/fe** | Reinstalls if the file is missing or an equal or older version is installed. |
| **/fd** | Reinstalls if the file is missing or a different version is installed. |
| **/fc** | Reinstalls if the file is missing or the stored checksum does not match the calculated value. |
| **/fa** | Forces all of the files to be reinstalled. |
| **/fu** | Rewrites all of the required user-specific registry entries. |
| **/fm** | Rewrites all of the required computer-specific registry entries. |
| **/fs** | Overwrites all of the existing shortcuts. |
| **/fv** | Runs from source and recaches the local package. |
| *package* | Specifies the name of the Windows Installer package file. |
| *ProductCode* | Specifies the globally unique identifier (GUID) of the Windows Installer package. |

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The default parameter for this command-line option is **/fpecms**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Remove a package using the command line

# Advertise a package using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To advertise a package using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /j**{**u** | **m**} *package*

| Value | Description |
|---|---|
| **/j** | Advertises a product. This command-line option ignores any property values entered on the command line. |
| **u** | Advertises to the current user. |
| **m** | Advertises to all users of this computer. |
| *package* | Specifies the name of the Windows Installer package file. |

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you want to install the application with elevated privileges, use the **/jm** command-line option.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Use Windows Installer from the command line

# Create an administrative installation package from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an administrative installation package from the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /a***package*

| Value | Description |
|---|---|
| **/a** | Applies the administrative installation option. |
| *package* | Specifies the name of the Windows Installer package file. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Install a package using the command line

# Create an installation log file using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an installation log file using the command line**

1. Open Command Prompt.

2. Type:

   **msiexec /i** *package***/L** [**i**][**w**][**e**][**a**][**r**][**u**][**c**][**m**][**p**][**v**][**+**][**!**][*****] *LogFile***.txt**

| Value | Description |
|---|---|
| **/i** | Installs a product. |
| *package* | Specifies the name of the Windows Installer package. |
| **/L** | Indicates to log information. |
| **/Li** | Logs status messages. |
| **/Lw** | Logs nonfatal warnings. |
| **/Le** | Logs all error messages. |
| **/La** | Logs start up of actions. |
| **/Lr** | Logs action-specific records. |
| **/Lu** | Logs user requests. |
| **/Lc** | Logs initial user interface parameters. |
| **/Lm** | Out-of-memory. |
| **/Lp** | Logs Terminal properties. |
| **/Lv** | Specifies verbose output. |
| **/L+** | Appends log to the existing file. |
| **/L!** | Flushes each line to the log. |
| **/L*** | Wildcard. Logs all information except for the **v** command-line option. |
| *LogFile***.txt** | Specifies the name and path of the text log file. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Windows Installer command-line options are not case-sensitive. For more information about **msiexec** command-line options, see Related Topics.

- To include the **v** option in a log file using the wildcard flag, type **/L*v** at the command prompt.

- The Windows Installer log file options can also be used with the uninstall and repair processes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Msiexec (command-line options)
Advertise a package using the command line

# Configure Group Policy for Windows Installer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Group Policy for Windows Installer**

- Configure Windows Installer Group Policy computer options
- Configure Windows Installer Group Policy user options

# Configure Windows Installer Group Policy computer options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Windows Installer Group Policy computer options**

1. Open Group Policy.
2. In the console tree, click the **Windows Installer** node.

   **Where?**

   - *applicable policy name*/Computer Configuration/Administrative Templates/Windows Components/Windows Installer

3. In the details window, double-click one of the following policy settings:

   - **Disable Windows Installer**
   - **Always install with elevated privileges**
   - **Prohibit rollback**
   - **Remove browse dialog box for new source**
   - **Prohibit patching**
   - **Disable IE security prompt for Windows Installer scripts**
   - **Enable user control over installs**
   - **Enable user to browse for source while elevated**
   - **Enable user to use media source while elevated**
   - **Enable user to patch elevated products**
   - **Allow admin to install from Terminal Services session**
   - **Cache transforms in secure location on workstation**
   - **Logging**
   - **Prohibit User Installs**
   - **Turn off creation of System Restore Checkpoints**
   - **Prohibit removal of updates**
   - **Enforce upgrade components rule**
   - **Prohibit non-administrators from applying vendor signed updates**

- **Baseline file cache maximum size**

4. In the **Group Policy Property** dialog box, set the options you want to use, and then click **OK**.

### Notes

- To open Group Policy to edit the local Group Policy object, click **Start**, click **Run**, type **gpedit.msc**, and then press ENTER.
- For more information about Windows Installer Group Policy options, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Managing options for computers through Group Policy
Using Windows Installer

# Configure Windows Installer Group Policy user options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Windows Installer Group Policy user options**

1. Open Group Policy.
2. In the console tree, click the **Windows Installer** node.

   **Where?**

   - *applicable policy name*/User Configuration/Administrative Templates/Windows Components/Windows Installer

3. In the details window, double-click one of the following policy settings:

   - **Always install with elevated privileges**
   - **Search order**
   - **Prohibit rollback**
   - **Prevent removable media source for any install**

4. In the **Group Policy Property** dialog box, set the options you want to use, and then click **OK**.

## Notes

- To open Group Policy to edit the local Group Policy object, click **Start**, click **Run**, type **gpedit.msc**, and then press ENTER.
- For more information about Windows Installer Group Policy options, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Managing options for users through Group Policy
Using Windows Installer

# Windows Installer Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Windows Installer:

- Windows Installer overview
- Understanding Windows Installer
- Using Windows Installer

# Windows Installer overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer overview**

Windows Installer is a component of the Windows operating system that simplifies the application installation process.

Windows Installer manages the installation and removal of applications by applying a set of centrally defined setup rules during the installation process. You can also use this service to modify, repair, or remove an existing application. Windows Installer technology consists of the Windows Installer service for the Windows operating systems and the package (.msi) file format that contains information about application set up and installation.

Windows Installer is also an extensible software management system. It manages the installation, addition, and deletion of software components, monitors file resiliency, and maintains basic failure recovery with rollbacks. Windows Installer supports installing and running software from multiple sources, and developers can customize Windows Installer to install custom applications.

# Understanding Windows Installer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Windows Installer**
This section covers:

- Windows Installer features
- Windows Installer technologies

# Windows Installer features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer features**

With Windows Installer and the .msi package file format, software installation and removal are more reliable and resilient. Windows Installer also provides a larger set of installation options. Windows Installer performs the following tasks:

- **Restores original computer state upon installation failure--**Windows Installer keeps track of all changes made to the system during the application installation process. If the installation fails, Windows Installer can restore, or roll back, the system to its initial state.

- **Helps prevent interapplication conflicts--**Windows Installer enforces installation rules that help to prevent conflicts with shared resources between existing applications. Such conflicts can be caused when an install operation makes updates to a dynamic link library (.dll) shared by an existing application or when an operation deletes a .dll shared by another application.

- **Reliably removes existing programs--**Windows Installer can reliably uninstall any program it previously installed. It removes all of the associated registry entries and application files, except for those shared by other installed software. You can uninstall an application at any time after a successful installation. (Removal should not be confused with rollback, which restores a computer to its initial state when an installation failure has occurred.)

- **Diagnoses and repairs corrupted applications--**An application can query Windows Installer to determine whether an installed application has missing or corrupted files. If any are detected, Windows Installer repairs the application by recopying only those files found to be missing or corrupted.

- **Supports on-demand installation of application features--**Windows Installer can be configured to initially install a minimal subset of an application. Later, additional components can be automatically installed the first time the user accesses features that require those components. This is known as advertising. For example, Windows Installer could install Microsoft Word with a minimal set of features. The first time the user tried to access a mail merge function (not included with the original installation), Windows Installer would automatically install the mail merge component. Similarly, Windows Installer can also purge components that go unused in an application. For example, Windows Installer could be configured to remove the mail merge component if it goes unused for 60 days.

- **Supports unattended application installation--**Installation packages can be configured to require no installation process interaction from the user. During the installation process, Windows Installer can query the computer for desktop attributes, including determining whether any applications were previously installed by Windows Installer.

- **Supports 32-bit and 64-bit applications--**32-bit applications can be installed on 64-bit machines.

- **Supports Microsoft .NET Framework--**The Microsoft .NET Framework is a new platform for building integrated, service-oriented applications that gather information from and interact with a wide variety of sources, regardless of the platforms or languages in use. The .NET Framework and the common language runtime can deliver write-once, compile-once, run-anywhere application development. Specifically, the .NET Framework delivers code reuse, code specialization, resource management, multilanguage development, security, deployment, and administration.

- **Integrated with side-by-side components--**This feature eliminates .dll version conflicts by permitting an application to be bound to the version of the component it was designed and tested with, regardless of the computer that hosts the application. Side-by-side components support the simultaneous execution of multiple versions of each component.

- **Integrated with software restriction policies--**This feature provides virus protection support, including protection from Trojan horse viruses and worms propagated through e-mail and the Web. Software restriction policies make the simple, point-and-click, active context user experience safe. Group Policy implements them as part of the list of trusted applications, and Windows Installer operates with applications permitted by these software restriction policies.

# Windows Installer technologies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer technologies**

Windows Installer technologies are divided into two parts that work in combination: a client-side installer service (Msiexec.exe) and a package file (.msi file). Windows Installer uses the information contained within a package file to install the application. Following are the Windows Installer technologies:

- **Installer service--**Windows Installer is an operating system service that allows the operating system to manage the installation process.

- **Installer program--**The Msiexec.exe program is a component of Windows Installer. This program uses a dynamic link library, Msi.dll, to read the package files (.msi), apply transforms (.mst), and incorporate command-line options. The installer performs all installation-related tasks: copying files onto the hard disk, making registry modifications, creating shortcuts on the desktop, and displaying dialog boxes to query user installation preferences when necessary.

  When Windows Installer is installed on a computer, the file association capabilities of the operating system are modified to recognize the .msi file type. When a user double-clicks a file with the .msi extension, the operating system associates the .msi file with Windows Installer and runs the Msiexec.exe application.

- **Installation package file--**Each package (.msi) file contains a relational type database that stores all the instructions and data required to install (and uninstall) the program across many installation scenarios. For example, a package file could contain instructions for installing an application when a prior version of the application is already installed. The package file could also contain instructions for installing the software on a computer where that application has never been present.

- **Transforms--**The installation process can be manipulated by applying transforms (.mst) to the installation database. A transform makes changes to elements of the database. For example, Windows Installer can use a transform file to change the language in the user interface of an application. The Windows Installer transform files modify the installation package file at installation time, and can therefore dynamically affect the installation behavior.

  If you are using Group Policy Software Installation, you should note that a transform is applied at the time of assignment or publication, in addition to the time of installation. In practical terms, this means that you should ensure the **Modifications** tab of the package properties dialog box is set up the way you want it to be before clicking **OK**. If you neglect to do this and deploy an incorrectly transformed package, then you either have to remove the software and redeploy it, or upgrade the software with a correctly transformed version.

  Customization transforms, much like patches, remain cached on the computer. These transforms are applied to the base package file whenever Windows Installer needs to perform a configuration change to the installation package. Transforms are applied only at initial installation; they cannot be applied to an application that has already been installed.

# Using Windows Installer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Windows Installer**
This section covers:

- Windows Installer in managed environments
- Msiexec (command-line options)
- Managing options for computers through Group Policy
- Managing options for users through Group Policy

# Windows Installer in managed environments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Installer in managed environments**

Windows Installer is an integral part of IntelliMirror, and a core component of the Group Policy-based change and configuration management technology. By using the IntelliMirror technologies, Group Policy, and change and configuration management, administrators can approve certain applications, specifying that all configuration operations on those applications (installation, removal, and repair) run as the local system account. The administrator has control and management of the file system and registry, and Windows Installer performs user-initiated software installations. Only applications that are approved by the administrator run with elevated privileges. For more information about Windows Installer, see Using Windows Installer.

Administrators use Active Directory, IntelliMirror, and Group Policy to assign and publish applications to groups of users or computers within the enterprise. Active Directory is a secure, distributed, partitioned, and replicated directory service that provides management services. These services include a standardized method of locating resources within the enterprise, and applying Group Policy to objects managed by Active Directory. For more information about Active Directory, see Active Directory.

# Msiexec (command-line options)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Msiexec**
Provides the means to install, modify, and perform operations on Windows Installer from the command line.

- To install or configure a product
- To use the administrative installation option
- To repair a product
- To uninstall a product
- To advertise a product
- To set logging level
- To apply an update
- To install a transform using the command line
- To install and advertise a product using a transform with the command line
- To set the user interface level
- To display copyright information for the Windows Installer
- To call the system API DllRegisterServer to self-register modules passed on the command line
- To call the system API DllUnRegisterServer to unregister modules passed on the command line

**To install or configure a product**
**Syntax**
**msiexec /i** {*Package* | *ProductCode*}

**Parameters**
**/i**
Installs or configures a product.

***Package***
Specifies the name of the Windows Installer package file.

***ProductCode***
Specifies the globally unique identifier (GUID) of the Windows Installer package.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To install a product from A:\Example.msi, type:

**msiexec /i A:\Example.msi**

**To use the administrative installation option**
**Syntax**
**msiexec /a** *Package*

**Parameters**
**/a**
Applies the administrative installation option.

***Package***
The name of the Windows Installer package file.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**To repair a product**
**Syntax**
**msiexec /f** [**p**][**o**][**e**][**d**][**c**][**a**][**u**][**m**][**s**][**v**]{*Package* | *ProductCode*}

**Parameters**
**/f**
Enables one or more of the command-line options listed in the following table.

| Command | Description |
|---|---|
| **p** | Reinstalls only if file is missing. |
| **o** | Reinstalls if file is missing or if an older version is installed. |
| **e** | Reinstalls if file is missing or an equal or older version is installed. |
| **d** | Reinstalls if file is missing or a different version is installed. |
| **c** | Reinstalls if file is missing or the stored checksum does not match the calculated value. |
| **a** | Forces all files to be reinstalled. |
| **u** | Rewrite all required user-specific registry entries. |
| **m** | Rewrites all required computer-specific registry entries. |
| **s** | Overwrites all existing shortcuts. |
| **v** | Runs from source and re-caches the local package. |

***Package***
Name of the Windows Installer package file.

***ProductCode***
Globally unique identifier (GUID) of the Windows Installer package.

**Remarks**

- This option ignores any property values that you type at the command line.
- The default parameter for this command-line option is **/fpecms**.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To repair the installation package, type:

**msiexec /fpecms Example.msi**

**To uninstall a product**
**Syntax**
**msiexec /x** {*Package* | *ProductCode*}

**Parameters**
**/x**
uninstalls a product.

***Package***
Name of the Windows Installer package file.

***ProductCode***
Globally unique identifier (GUID) of the Windows Installer package.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To remove or uninstall a package, type:

**msiexec /x Example.msi**

**To advertise a product**
**Syntax**
**msiexec /j** [{**u** | **m**}] *Package*

**msiexec** {**u** | **m**} *Package* **/t** *TransformList*

**msiexec** {**u** | **m**} *Package* **/g** *LanguageID*

**Parameters**
**/j**
Advertises a product.

**u**
Advertises to the current user.

**m**
Advertises to all users of the computer.

***Package***
Specifies the Windows Installer package file.

**/g *LanguageID***
Identifies the language.

**/t *TransformList***
Applies transform to advertised package.

**Remarks**

- This option ignores any property values that you type at the command line.
- If you want to install the application with elevated privileges, use **/jm**.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To advertise a package to all users of this computer, type:

**msiexec /jm Example.msi**

**To set logging level**
**Syntax**
**msiexec /L** [**i**][**w**][**e**][**a**][**r**][**u**][**c**][**m**][**p**][**v**][+][**!**]*LogFile***.txt**

**Parameters**
**/L**
Specifies the path to the log file.

**i**
Logs status messages.

**w**
Logs nonfatal warnings.

**e**
Logs all error messages.

**a**
Logs startup of actions.

**r**
Logs action-specific records.

**u**
Logs user requests.

**c**
Logs initial user interface parameters.

**m**
Logs out-of-memory.

**p**
Logs terminal properties.

**v**
Logs verbose output. To use **v**, specify **/L*v**.

**+**
Appends to existing file.

**!**
Flushes each line to the log.

*****
Logs all information except for the **v** option. This is a wildcard.

***LogFile* .txt**
Name and path of the text log file.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- To include the **v** option in a log file using the wildcard flag, type **/L*v** at the command prompt.
- The Windows Installer log file options can also be used with the uninstall and repair processes.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To install a package and create a log file that contains the information related to the status, out-of-memory, and error messages, type:

**msiexec /i Example.msi /Lime logfile.txt**

**To apply an update**
**Syntax**
**msiexec /p** *UpdatePackage*

**Parameters**
**/p**
Applies an update.

***UpdatePackage***
Specific update.

**Remarks**

- To apply an update to an administrative installation package, use the following syntax:

  **msiexec/p***UpdatePackage***/a***Example.msi*

- Windows Installer command-line options are not case-sensitive.

- For more information about the Windows Installer command-line options, see Related Topics.

**To install a transform using the command line**
**Syntax**
**msiexec /i** *Package* **TRANSFORMS**=*TransformList*

**Parameters**
**/i**
Installs or configures a product.

***Package***
Specifies the Windows Installer package file.

**TRANSFORMS=**
Property that is used to specify what transform (.mst) files should be applied to the package.

***TransformList***
List of paths separated by semicolons.

**Remarks**

- Windows Installer command-line options are not case-sensitive.

- For more information about the Windows Installer command-line options, see Related Topics.

**To install and advertise a product using a transform with the command line**
**Syntax**
**msiexec /i** *Package* **/j**[**u**][**m**] **/t***TransformList*

**Parameters**
**/i**
Installs or configures a product.

***Package***
Name of the Windows Installer package file.

**/j**
Advertises a product. This option ignores any property values entered on the command line.

**u**

Advertises to the current user.

**m**
Advertises to all users of this computer.

**/t**
Applies transform to advertised package.

***TransformList***
List of paths separated by semicolons.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**To set the user interface level**
**Syntax**
**msiexec /q**{**n** | **b** | **r** | **f** | **n+** | **b+** | **b-**}

**Parameters**
**/qn**
Displays no user interface.

**/qb**
Displays a basic user interface.

**/qr**
Displays a reduced user interface with a modal dialog box displayed at the end of the installation.

**/qf**
Displays the full user interface with a modal dialog box displayed at the end.

**/qn+**
Displays no user interface, except for a modal dialog box displayed at the end.

**/qb+**
Displays a basic user interface with a modal dialog box displayed at the end.

**/qb-**
Displays a basic user interface with no modal dialog boxes.

**Remarks**

- **/qb+-** is not a supported user interface level. The modal box is not displayed if the user cancels the installation.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
To display the basic user interface options during the package installation of Example.msi, type:

**msiexec /qb Example.msi**

**To display copyright information for the Windows Installer**
**Syntax**
**msiexec** {**/?** | **/h**}

**Parameters**
**{ /?| /h}**
Displays the Windows Installer version and copyright information.

**Remarks**

- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**To call the system API DllRegisterServer to self-register modules passed on the command line**
**Syntax**
**msiexec /y** *Module*

**Parameters**
**/y**
Calls the system API DllRegisterServer to self-register modules passed on the command line.

***Module***
Specifies the file name of module.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Remarks**

- This option is used only for registry information that cannot be added using the registry tables of the .msi file.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
The following example shows how you can use the **msiexec /y** command:

**msiexec /y my_file.dll**

**To call the system API DllUnRegisterServer to unregister modules passed on the command line**
**Syntax**
**msiexec /z** *Module*

**Parameters**
**/z**
Calls the system API DllUnRegisterServer to unregister modules passed on the command line.

***Module***
File name of module.

**Remarks**

- This option is used only for registry information that cannot be added using the registry tables of the .msi file.
- Windows Installer command-line options are not case-sensitive.
- For more information about the Windows Installer command-line options, see Related Topics.

**Examples**
The following example shows how you can use the **msiexec /z** command:

**msiexec /z my_file.dll**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |

| | |
|---|---|
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Install a package using the command line
Install a transform using the command line
Install and advertise a product using a transform at the command line
Create an administrative installation package from the command line
Repair a package using the command line
Remove a package using the command line
Advertise a package using the command line
Create an installation log file using the command line
Command-line reference A-Z
Command shell overview

# Managing options for computers through Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing options for computers through Group Policy**
Windows Installer can be configured by using Group Policy and Active Directory to manage the computer installation options.

The following table shows policies concerned with managing the Group Policy computer options for Windows Installer. To configure these policies, start Group Policy, then in the console tree, click the **Windows Installer** node.

**Where?**

- *applicable policy name*/Computer Configuration/Administrative Templates/Windows Installer

| Policy | Details |
|---|---|
| **Disable Windows Installer** | Disables or restricts the use of Windows Installer.<br><br>This policy can prevent users from installing software on their system, or permit users to install only those programs offered by an administrator.<br><br>If you enable this policy, you can use the options in the **Disable Windows Installer** dialog box to establish an installation policy.<br><br>• **Never--**Windows Installer is fully enabled. Users can install and upgrade software. Windows Installer is enabled by default in Windows.<br><br>• **For non-managed apps only--**Users can install only those programs that an administrator assigns (offers on the desktop) or publishes (adds to Add or Remove programs).<br><br>• **Always--**Windows Installer is disabled.<br><br>**Note**<br><br>• This policy affects Windows Installer only. It does not prevent users from using other methods to install and upgrade programs. |
| **Always install with elevated privileges** | Directs Windows Installer to use system permissions when it installs any program on the system.<br><br>This policy extends to all programs the elevated privileges usually reserved for programs that have been assigned to the user (offered on the desktop) or the computer (installed automatically), or made available in Add or Remove programs in Control Panel. This policy lets users install programs that require access to directories that the user might not have permission to view or change, including directories on highly restricted computers.<br><br>If you disable this policy, or do not configure it, the system applies the current user's permissions when it installs programs that are not distributed or offered by an administrator.<br><br>**Note**<br><br>• This policy appears in both the Computer Configuration and User Configuration folders. To make this policy effective, you must enable the policy in both folders. |

| | |
|---|---|
| **Prohibit rollback** | Prohibits Windows Installer from generating and saving the files it needs to reverse an interrupted or unsuccessful installation.<br><br>**Prohibit rollback** also prevents Windows Installer from recording the original state of the system and frequency of changes it makes during installation, and prevents Windows Installer from retaining files it intends to delete later. As a result, Windows Installer cannot restore the computer to its original state if the installation does not complete.<br><br>This policy is designed to reduce the amount of temporary disk space required to install programs. **Prohibit rollback** also prevents malicious users from interrupting an installation to gather data about the internal state of the computer or to search secure system files. However, because an incomplete installation can render the system or a program inoperable, do not use this policy unless essential.<br><br>This policy appears in both the Computer Configuration and User Configuration folders. To make this policy effective, you must enable the policy in both folders. |
| **Remove browse dialog box for new source** | Prevents users from searching for installation files when adding features or components to an installed program.<br><br>This policy disables the **Browse** button next to the **Use feature from** list in the Windows Installer dialog box. Users must select an installation file source from the **Use features from** list configured by the administrator.<br><br>This policy applies even when the installation is running in the user's security context.<br><br>If you disable this policy or do not configure it, the **Browse** button is enabled when an installation is running in the user's security context, but only administrators can browse when an installation is running with elevated system privileges, such as installations offered on the desktop or in Add or Remove programs.<br><br>This policy affects Windows Installer only. It does not prevent users from using other browsers, such as Windows Explorer or My Network Places, to search for installation files. |
| **Prohibit patching** | Prevents users from installing patches to upgrade their programs.<br><br>Patches are updates or upgrades that replace only those program files that have changed. Because patches can be easy vehicles for malicious programs, some installations prohibit their use.<br><br>This policy applies only to installations that run in the user's security context. By default, users who are not administrators cannot apply patches to installations that run with elevated system privileges, such as those offered on the desktop or in Add or Remove programs. |
| **Disable IE security prompt for Windows Installer scripts** | Allows Web-based programs to install software on the computer without notifying the user.<br><br>By default, when a script hosted by an Internet browser attempts to install a program on the system, the system warns the user and prompts the user to either allow or refuse the installation. This policy suppresses the warning and allows the installation to proceed without user intervention.<br><br>This policy is designed for enterprises that use Web-based tools to distribute programs to their employees. However, because it might pose a security risk, it should be applied cautiously. |
| **Enable user control over installs** | Permits users to change installation options that typically are available only to administrators.<br><br>This policy bypasses some of the security features of Windows Installer. It permits installations to complete that otherwise would be halted due to a security violation.<br><br>The security features of Windows Installer prevent users from changing installation options that are typically reserved for administrators, such as specifying the directory to which files are installed. If Windows Installer detects that an installation package has permitted the user to change a protected option, it stops the installation and displays an error message. These security features operate only when the installation program is running in a privileged security context in which the installation program has access to directories that are denied to the user.<br><br>This policy is designed for less restrictive environments. It can be used to circumvent errors in an installation program that prevent software from being installed. |

| | |
|---|---|
| **Enable user to browse for source while elevated** | Permits users to search for installation files during privileged installations.<br><br>This policy enables the **Browse** button on the **Use feature from** dialog box. Users can search for installation files, even when the installation program is running, with elevated system privileges. By default, only system administrators can browse during installations with elevated privileges, such as installations offered on the desktop or displayed in Add or Remove programs.<br><br>Because the installation is running with elevated system privileges, users can browse through directories that their own permissions would not allow.<br><br>This policy does not affect installations that run in the user's security context. For more information, see the **Remove browse dialog box for new source** policy. |
| **Enable user to use media source while elevated** | Permits users to install programs from removable media, such as floppy disks and CD-ROMs, during privileged installation.<br><br>This policy permits all users to install programs from removable media, even when the installation program is running with elevated system privileges. By default, users can install programs from removable media only when the installation runs in the user's security context. During privileged installations, such as those offered on the desktop or displayed in Add or Remove programs, only system administrators can install from removable media.<br><br>This policy does not affect installations that run in the user's security context. By default, users can install from removable media when the installation runs in their own security context.<br><br>For more information, see the **Prevent removable media source for any install** policy in Managing options for users through Group Policy. |
| **Enable user to patch elevated products** | Permits users to upgrade programs during privileged installations.<br><br>This policy permits all users to install patches, even when the installation program is running with elevated privileges. Patches are updates or upgrades that replace only those program files that have changed. Because patches can easily be vehicles for malicious programs, some installations prohibit their use.<br><br>By default, only system administrators can apply patches during installations with elevated privileges, such as installations offered on the desktop or displayed in Add or Remove programs.<br><br>This policy does not affect installations that run in the user's security context. By default, users can install patches to programs that run in their own security context. For more information about prohibiting patch installations, see the **prohibit patching** policy. |
| **Allow admin to install from Terminal Services session** | Permits Terminal Services administrators to install and configure programs remotely.<br><br>By default, system administrators can install programs only when they are logged on to the computer on which the program is being installed. This policy creates a special exception for computers running Terminal Services.<br><br>This policy affects system administrators only. Other users cannot install programs remotely. |
| **Cache transforms in secure location on workstation** | Copies of transform files are saved in a secure location on the local computer.<br><br>Transform files consist of instructions to modify or customize a program during installation. By default, Windows Installer stores transform files in the Application Data directory in the user's profile. When a user reinstalls, removes, or repairs an installation, the transform file is available, even if the user is on a different computer, or isn't connected to the network.<br><br>If you enable this policy, the transform file is saved in a secure location on the user's computer instead of in the user profile. Because Windows Installer requires the transform file to repeat an installation in which the transform file was used, the user must be using the same computer or be connected to the original or identical media to reinstall, remove, or repair the installation.<br><br>This policy is designed for enterprises that must take special precautions to prevent unauthorized or malicious editing of transform files. |

| | |
|---|---|
| **prohibit User Installs** | Allows you to configure user installs. To configure this setting, set it to enabled and use the drop-down list to select the behavior you want.<br><br>If this setting is not configured, or if the setting is enabled and **Allow User Installs** is selected, the installer allows and makes use of products that are installed per user, and products that are installed per computer. If the installer finds a per-user install of an application, this hides a per-computer installation of that same product.<br><br>If this setting is enabled and **Hide User Installs** is selected, the installer ignores per-user applications. This causes a per-computer installed application to be visible to users, even if those users have a per-user install of the product registered in their user profile.<br><br>If this setting is enabled and **prohibit User Installs** is selected, the installer prevents applications from being installed per user, and it ignores previously installed per-user applications. An attempt to perform a per-user installation causes the installer to display an error message and stop the installation. This setting is useful in environments where the administrator only wants per-computer applications installed, such as on a kiosk or a Windows Terminal Server. |
| **Turn off creation of System Restore Checkpoints** | Stops Windows Installer from generating System Restore checkpoints when installing applications.<br><br>System Restore enables users, in the event of a problem, to restore their computers to a previous state without losing personal data files. By default, Windows Installer automatically creates a System Restore checkpoint each time an application is installed, so that users can restore their computer to the state it was in before installing the application.<br><br>If you disable this setting or do not configure it, Windows Installer automatically creates a System Restore checkpoint each time an application is installed.<br><br>**Note**<br><br>• This setting only applies to Windows XP Professional. |
| **Logging** | Specifies the types of events that Windows Installer records in its transaction log for each installation. The log, msi.log, appears in the *systemroot*\Temp directory. The Windows Installer logging options are listed in the following table. To indicate that an event type is recorded, enter the letter representing the event type. You can enter the letters in any order, and list as many or as few event types as you want. To disable logging, delete all of the letters from the box. If you disable this policy, or do not configure it, Windows Installer logs the default event types, represented by the letters **iweap**. You can include all possible events in the installation log file by entering **iwearucmpvo**. However, doing so will add considerable time to the installation process. |

| Mode | Log entry |
|---|---|
| **i** | Status messages |
| **w** | Non-fatal warnings |
| **e** | All error messages |
| **a** | Startup of actions |
| **r** | Action-specific records |
| **u** | User requests |
| **c** | Initial user interface parameters |
| **m** | Out-of-memory |
| **p** | Terminal properties |
| **v** | Verbose output |
| **o** | Out of disk space messages |

# Managing options for users through Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing options for users through Group Policy**

Windows Installer can be configured by using Group Policy and Active Directory to manage the user installation options.

This table shows policies concerned with managing the Group Policy user options for Windows Installer. To configure these policies, start Group Policy, then in the console tree, click the **Windows Installer** node.

**Where?**

- *applicable policy name*/User Configuration/Administrative Templates/Windows Installer

| Policy | Details |
|---|---|
| **Always install with elevated privileges** | Directs Windows Installer to use system permissions when it installs any program on the system.<br><br>This policy extends to all programs the elevated privileges usually reserved for programs that have been assigned to the user (offered on the desktop) or the computer (installed automatically), or made available in Add or Remove Programs in Control Panel. This policy lets users install programs that require access to directories that the user might not have permission to view or change, including directories on highly restricted computers.<br><br>If you disable this policy, or do not configure it, the system applies the current user's permissions when it installs programs that are not distributed or offered by an administrator.<br><br>**Note**<br><br>• This policy appears both in the Computer Configuration and User Configuration folders. To make this policy effective, you must enable the policy in both folders. |
| **Search order** | Specifies the order in which Windows Installer searches for installation files.<br><br>By default, Windows Installer searches the network first, then removable media (floppy drive, CD-ROM, or DVD), and finally, the Internet (URL).<br><br>To change the search order, enable the policy, and then type the following letters for each file source in the order that you want Windows Installer to search. For example, if you want to search the network, the Internet, and then removable media, type **num**. To exclude a file source, omit or delete the letter representing that source type.<br><br>• **n** - Adds the network in the source file search.<br><br>• **m** - Adds the removable media in the source file search.<br><br>• **u** - Adds the Internet in the source file search. |

| | |
|---|---|
| **Prohibit rollback** | Prohibits Windows Installer from generating and saving the files it needs to reverse an interrupted or unsuccessful installation.<br><br>This policy prevents Windows Installer from recording the original state of the system and sequence of changes it makes during installation. It also prevents Windows Installer from retaining files it intends to delete later. As a result, Windows Installer cannot restore the computer to its original state if the installation does not complete.<br><br>This policy is designed to reduce the amount of temporary disk space required to install applications. It also prevents malicious users from interrupting an installation to gather data about internal state of the computer or to search secure system files. However, because an incomplete installation can render the system or an application inoperable, this policy should not be used unless it is essential.<br><br>**Note**<br><br>• This policy appears in both the Computer Configuration and User Configuration folders. To make this policy effective, you must enable the policy in both folders. |
| **Prevent removable media source for any install** | Prevents users from installing programs from removable media.<br><br>If a user tries to install a program from removable media, such as CD-ROMs, floppy disks, or DVDs, a message informs the user that the feature cannot be found.<br><br>This policy applies even when the installation is running in the user's security context.<br><br>If you disable or do not configure this policy, users can install from removable media when the installation is running in their own security context, but only system administrators can use removable media when an installation is running with elevated system privileges, such as installations offered on the desktop or in Add or Remove Programs. |

# Windows Installer Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The installation package does not install correctly.
- The installation process stops before completing the installation.

**The installation package does not install correctly.**
**Cause:** The installation packaged has been corrupted.

**Solution:** Repair the installation package by using the Windows Installer repair option.

**See also:** Repair a package using the command line

**The installation process stops before completing the installation.**
**Cause:** Windows Installer was not able to read the installation package, or conditions on your computer prevented the package from being installed successfully.

**Solution:** Windows Installer creates entries in the Windows event log. To view the log, do the following:

1. Open Event Viewer.
2. Review the application log for Windows Installer entries.

**Notes**

- To open Event Viewer, click **Start**, point to **Settings**, and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Event Viewer**.
- You can configure Windows Installer to create a custom installation log file.

**See also:** Create an installation log file using the command line

# Remote Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Administration**
Using tools included with Microsoft® Windows Server 2003 family operating systems CDs, you can remotely manage your system servers running Microsoft® Windows 2000 and Windows Server 2003 family operating systems. You can also remotely manage Windows Server 2003 family operating systems from a computer using Microsoft® Windows® XP Professional.

- For general information, see Remote administration overview.
- For tips about managing your server remotely, see Remote Administration Best practices.
- For other sources of information, see Remote Administration Resources.

# Remote administration overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote administration overview**

The following table presents remote management tasks and links to information about the tools. Unless indicated otherwise, these tools are used on Microsoft® Windows Server™ 2003 operating systems and, in some cases, Microsoft® Windows® XP Professional.

| Task | Tools | Comments |
|---|---|---|
| Manage several servers by performing similar tasks. | Microsoft Management Console<br><br>Author an MMC console file<br><br>Work with MMC and saved console files | You can use appropriate saved MMC consoles, if available, or you can create custom MMC consoles for tasks you frequently either delegate or perform. |
| Log onto one computer remotely and manage it as if you were there. | Remote Desktop Connection | In addition to Windows Server 2003 operating systems, this feature can be used with Microsoft® Windows® 2000 Server. The Remote Desktop feature must be enabled on the remote computer. |
| Switch among several computers that you want to log on to remotely. | Remote Desktops snap-in | In addition to Windows Server 2003 operating systems, this feature can be used with Windows 2000 Server. The Remote Desktop feature must be enabled on the remote computer. |
| Administer a server running Windows Server 2003 from any computer on your network. | Remote Desktop for Administration | In addition to Windows Server 2003 operating systems, this feature can be used with Windows 2000 Server. The Remote Desktop feature must be enabled on the remote computer. |
| Manage servers from Windows XP Professional. | Windows Server 2003 Administration Tools Pack Overview<br><br>Remote Desktop Connection | Windows Server 2003 Administration Tools Pack must be installed on Windows XP Professional to permit server management using these tools. The Remote Desktop feature must be enabled on servers to permit use of the Remote Desktop Connection. |
| Manage your Web server using a Web browser on a remote computer. | Using Web Interface for Remote Administration | Use Internet Explorer, version 6 or later, for remote administration using Web Interface. Web Interface is not installed by default for Web servers. |
| Assist in real time an operator on a remote computer running Windows Server 2003 family operating system, Windows XP Professional or Microsoft® Windows® XP Home Edition. | Administering Remote Assistance (www.microsoft.com) | The Remote Assistance connection requires explicit permission from the connected user, and both computers must be running Windows XP Professional or Windows Server 2003 operating systems. |
| Connect to a remote computer when that computer cannot access a network or is not in a functional state because of hardware or software failure. | Telnet | All commands must be entered through the command line. Capabilities are limited, and security is minimal. |

| | | |
|---|---|---|
| Manage disk space on servers by copying files to a library of magnetic tapes or magneto-optical disks. | Remote Storage | Some Windows components require configuration before they can be used and will be displayed in the **Add or Remove Windows Components** in **Add or Remove programs** in **Control Panel**. Remote storage requires installation. |

**Notes**

- Depending on the Group Policy settings on your network, you might not be able to perform some of these tasks.
- Some features require that specific ports be opened on the server in order for the role to be administered remotely. For information about these port requirements, see Service overview and network port requirements for the Windows Server system (http://go.microsoft.com/fwlink/?LinkID=61531) and Windows Firewall Settings (http://go.microsoft.com/fwlink/?LinkID=38651).
- You can use command-line tools to automate many processes and tasks. For more information, see Automating processes and tasks.
- You must physically access each computer to restart, change hardware, and enable the Remote Desktop feature. Once the Remote Desktop feature has been enabled, it will remain enabled.
- For an explanation of the various Remote Desktop terminologies, see Remote Administration using Terminal Services.
- For more information about remote administration, see the following:
    - Common Administrative Tasks
    - New Features
    - Quick Guide for Finding Tools
    - Remote Administration Resources

# Remote Administration Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**To perform many tasks on a single, remote server, use Remote Desktop for Administration.**

- To remotely administer a server running a Windows Server 2003 family operating system, you do not need to install Terminal Server. Instead, use Remote Desktop for Administration. On these operating systems, Remote Desktop for Administration is on by default but is not enabled. After you enable remote connections, Remote Desktop for Administration allows you to administer your server from any computer on your network running Windows Server 2003 family, Windows XP Professional, and Windows 2000. Windows Server 2003 , Web Edition, does not support remote administration.

- You can use Remote Desktop with Windows 2000 Server. You need to install Terminal Services for Remote Administration as described in the Windows 2000 Help.

  For more information, see Enable or disable Remote Desktop and Remote Desktop for Administration.

**To do similar tasks on many servers, use MMC snap-ins.**

- Most of the administration tools provided with Windows Server 2003 family operating systems are Microsoft Management Console (MMC) snap-ins that can be used to administer remote servers as well as your local computer. The advantages of using MMC include the following:

  - You can create a console with the tools you use for the tasks you do most often.

  - You can set the focus for a tool to any of the servers you administer, and switch between servers and tools within a single MMC console.

  For more information, see Microsoft Management Console.

**To manage servers running Windows Server 2003 family operating systems from a computer using Windows XP Professional , use Windows Server 2003 Administration Tools Pack**

- Windows Server 2003 Administration Tools Pack when installed on Windows XP Professional (ordinarily a client operating system) provides many remote management tools for managing your servers.

- Windows Server 2003 Administration Tools Pack comes as adminpak.msi on the Windows Server 2003 family operating system and is easily installed from the server CD onto Windows XP Professional by installing the .msi from the server to the Windows XP Professional operating system.

- The Administration Tools Packs in the Windows Server 2003 family and Windows 2000 are not mutually compatible. You must use the Windows 2000 Administrative Tools to administer Windows 2000, and the Windows Server 2003 Administration Tools Pack to administer the Windows Server 2003 family operating system.

  For more information, see Windows Server 2003 Administration Tools Pack Overview.

# Remote Administration Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**
The following resources provide additional information about remote administration.

**Windows Server 2003 Resource Kit**
The Microsoft Windows Server 2003 Resource Kit contains additional tools that you can use for remote administration.

**Microsoft Platform Software Development Kits (SDK)**
For technical information about Active Directory Scripting Interfaces (ADSI), Microsoft Management Console (MMC), Windows Management Instrumentation (WMI) Software Development Kits (SDK), see the Microsoft Web site.

**TechNet**
TechNet is a service that contains resources for information technology (IT) professionals. This includes access to articles written by support professionals at Microsoft, career and industry information, discussion groups, technical feedback, and software.

For more information, see Microsoft TechNet.

# Offline Files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Offline Files**
Using Offline Files, you can continue to work with network files and programs even when you are not connected to the network.

- For help with specific tasks, see Offline Files How To....
- For general background information, see Offline Files Concepts.

# Offline Files How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Set up your computer to use Offline Files
- Make a file or folder available offline
- Change how Offline Files responds to network disconnection
- Encrypt offline files
- Make Web pages available offline

# Set up your computer to use Offline Files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up your computer to use Offline Files**

1. Open My Computer.
2. On the **Tools** menu, click **Folder Options**.
3. Click the **Offline Files** tab, select the **Enable Offline Files** check box, and then do one of the following:
   - To do a full synchronization, select the **Synchronize all offline files before logging off** check box.
   - To do a quick synchronization, clear the **Synchronize all offline files before logging off** check box.

## Notes

- To open My Computer, click **Start**, and then click **My Computer**.
- A full synchronization ensures that you have the most current versions of every network file that has been made available offline. A quick synchronization ensures that you have complete versions of all of your offline files, but they might not be the most current versions.
- **Enable Offline Files** is controlled through Group Policy, as is **Amount of disk space to use for temporary offline files**. To adjust these settings, you must have the appropriate administrative credentials.
- By default, Offline Files is enabled in Windows XP Professional. However, by default, Offline Files is not enabled in the Windows Server 2003 family because Terminal Services is not compatible with Offline Files. To enable Offline Files on a server running Windows Server 2003 , right-click **My Computer**, and then click **properties**. In the **System properties** dialog box, click the **Remote** tab. Under **Remote Desktop**, clear the **Allow users to connect remotely to this computer** check box, click **OK**, and then follow the steps in this procedure.
- Even if your computer is configured to use Offline Files, you must still select the network files and folders that you want to make available offline.
- To view a list of all of the network files that are available offline, in the **Folder Options** dialog box, click the **Offline Files** tab, click **View Files**, and view the files in the Offline Files Folder.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working offline
Synchronizing offline items when you log on to or log off from a computer
Make a file or folder available offline
Encrypt offline files
Make Web pages available offline
Working with Web pages offline
Change how Offline Files responds to network disconnection

# Make a file or folder available offline

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a file or folder available offline**

1. Open My Computer.
2. In My Computer, click the shared network file or folder that you want to make available offline.
3. On the **File** menu, click **Make Available Offline**.
4. Follow the steps in the Offline Files Wizard.

### Important

- **Make Available Offline** appears on the **File** menu only if your computer is set up to use Offline Files. For more information about setting up your computer to use Offline Files, see Related Topics.

### Notes

- To open My Computer, click **Start**, and then click **My Computer**.
- You can make individual files from a shared network folder, or an entire shared network folder, available offline. If an entire folder is made available offline, any files placed in that shared folder will automatically be made available offline the next time the computers are synchronized.
- If a shortcut to a file is made available offline, that file is made available offline. If a shortcut to a folder is made available offline, the contents of that folder are not made available offline.
- When the status of your network connection changes, an Offline Files icon appears in the notification area, and a reminder balloon appears over the icon to notify you of the change. If the reminder balloon notifies you that you are offline, you can continue to work with your files as you normally do, or you can click the Offline Files icon in the notification area for more information about the status of your connection.
- When you open files on a shared folder marked for automatic caching, those files are automatically cached and become available to you when you are offline.
- To stop making a file or folder available offline, right-click the item, and then click **Make Available Offline** to remove the check mark.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up your computer to use Offline Files
Make Web pages available offline
Synchronizing offline items
Working offline
Working with Web pages offline
Change how Offline Files responds to network disconnection

# Change how Offline Files responds to network disconnection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change how Offline Files responds to network disconnection**

1. Open My Computer.
2. On the **Tools** menu, click **Folder Options**.
3. Click the **Offline Files** tab, and then click **Advanced**.
4. Under **When a network connection is lost**, do one of the following:
   - To receive notification that the network connection has been lost and to continue working with offline files, click **Notify me and begin working offline**.
   - To make offline files and folders unavailable to you when you lose a network connection, click **Never allow my computer to go offline**.

**Notes**

- To open My Computer, click **Start**, and then click **My Computer**.
- To set up individual computers to respond differently from the default settings, in the **Offline Files - Advanced Settings** dialog box, click **Add**. Under **Computer**, type the path to the computer, and then, under **When a network connection is lost**, click the setting that you want for this particular computer. For example, if you want to continue working offline by default and you do not want to work offline on any files from a particular computer on the network, you can add that computer's name under **Computer**, and then click **Never allow my computer to go offline** as the behavior for that specific computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up your computer to use Offline Files
Working offline
Working with Web pages offline

# Encrypt offline files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To encrypt offline files**

1. Open My Computer.
2. On the **Tools** menu, click **Folder Options**.
3. Click the **Offline Files** tab, and then select the **Encrypt offline files to secure data** check box.

## Notes

- To open My Computer, click **Start**, and then click **My Computer**.
- The **Encrypt offline files to secure data** check box is disabled under the following conditions:
  - You are not an administrator on the computer.
  - The local drive is not NTFS or does not support encryption.
  - Your system administrator has implemented an encryption policy for Offline Files.
  - Offline Files is not available in Windows XP Home Edition.
  - In the Windows Server 2003 family, Offline Files is not enabled by default. To enable Offline Files, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Encrypting File System overview
Encrypting offline files
Set up your computer to use Offline Files

# Make Web pages available offline

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make Web pages available offline**

1. Open Internet Explorer.
2. Access the Web page or Web site you want to make available offline.
3. On the **Favorites** menu, click **Add to Favorites**.
4. In the **Add Favorite** dialog box, select the **Make available offline** check box.
5. To specify a schedule for updating that page and how much content to download, click **Customize**, and then follow the instructions in the Offline Favorite Wizard.

### Note

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working offline
Make a file or folder available offline
Change how Offline Files responds to network disconnection
Working with Web pages offline

# Offline Files Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

- Offline Files overview
- Using Offline Files

# Offline Files overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Offline Files overview**
Using Offline Files, you can continue to work with network files and programs even when you are not connected to the network.

If you lose your connection to the network or undock your portable computer, your view of shared network resources that have been made available offline remains the same as when you were connected. You can continue to work with them as you normally do. You have the same access permissions to those files and folders as you would have if you were connected to the network. When the status of your connection changes, an Offline Files icon appears in the notification area and a reminder balloon appears over the notification area to notify you of the change.

When your network connection is restored or you dock your portable computer, any changes that you made while working offline are updated to the network by default. When you and someone else on the network have made changes to the same file, you have the option of saving your version of the file to the network, keeping the other version, or saving both.

For more information about synchronizing files with Synchronization Manager, see Synchronization Manager overview.

**Notes**

- You can initiate synchronization manually or you can set Synchronization Manager to control when offline files are synchronized with the network. Offline Files controls whether a full or quick synchronization is performed. A full synchronization ensures that you have the most current version of every network file that has been made available offline. A quick synchronization is much faster than a full synchronization, but may not provide the most current version of every network file that has been made available offline. However, a quick synchronization does ensure that a complete version of every file is available so that you can continue working.

- Any shared files or folders on a Microsoft network can be made available offline. You can make files available offline from any computer that supports Server Message Block-based (SMB) File and Printer Sharing, including Microsoft® Windows® 95, Windows® 98, Windows NT® version 4.0, Microsoft® Windows® XP, and the Windows Server 2003 family. Offline Files is not available on Novell NetWare networks.

- The system administrator can set up shared folders so that every network file you open in that shared folder is automatically made available offline, or so that only those network files you choose are made available offline.

- If a shortcut to a file is made available offline, that file is made available offline. If a shortcut to a folder is made available offline, the contents of that folder will not be made available offline.

- In the Windows Server 2003 family, Offline Files is not enabled by default. To enable Offline Files, see Set up your computer to use Offline Files.

# Using Offline Files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Offline Files**


- Working offline
- Sharing a drive or folder
- Handling file conflicts
- Encrypting offline files
- Working with Web pages offline
- Offline Files for mobile users

# Working offline

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Working offline**

Offline Files notifies you if the status of your network connection changes. When the status of your network connection changes, an Offline Files icon appears in the notification area and a reminder balloon appears over the notification area to notify you of the change. If the reminder balloon notifies you that you are offline, you can continue to work with your files as you normally do, or you can click the Offline Files icon in the notification area for more information about the status of your connection.

If you are working offline (either because you are disconnected from the network or because you undocked your portable computer), you can still browse network drives and shared folders in My Computer or My Network Places. A red X appears over any disconnected network drives. You will be able to see only those files that you made available offline and any files that you created after the network connection was lost. For more information about making a file or folder available offline, see Make a file or folder available offline.

Your permissions on the network files and folders remain the same whether you are connected to the network or working offline. For example, a read-only document on a mapped network drive would remain read-only if you were disconnected from the network.

When you are disconnected from the network, you can print to local printers, but you cannot print to shared printers on the network.

After you reconnect to the network, Synchronization Manager updates the network files with changes that you made while working offline. For more information about synchronizing files with Synchronization Manager, see Synchronizing offline items.

# Sharing a drive or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sharing a drive or folder**

You can share a folder on your computer with other users on the network. When you share a folder, you can decide the permissions that other users will be allowed or denied when they access the files and programs in that folder over the network. You can also specify whether other users will be able to make the shared folder available offline. For more information, see Share a folder or drive.

To make shared network files and programs available offline, Offline Files stores a version of them in a cache. The computer can access this cache while disconnected from the network. There are three caching options to choose from when sharing files:

- **Only the files and programs that users specify will be available offline**

  This caching option provides offline access to only those files and programs that someone using your shared folder specifically identifies. This caching option is ideal for a shared network folder containing files and programs that are to be accessed and modified by several people. This is the default option when you set up a shared folder to be used offline.

- **All files and programs that users open from the shared folder will automatically be available offline**

  This caching option makes every file and program opened from your shared folder automatically available offline. Files and programs that are not opened are not available offline. This caching option is ideal for making files and programs available offline that are read, referenced, or run, but not changed in the process.

  Automatic caching for files and programs provides offline access to shared folders containing files and programs that are not to be changed. Automatic caching for programs reduces network traffic because offline files are opened directly, without accessing the network versions in any way, and generally start and run faster than the network versions. To provide automatic caching for files only, clear the **Optimized for Performance** check box.

  When you use automatic caching for files and programs, be sure to restrict permissions on the files and programs contained in the shared folder to read-only access.

- **Files or programs from the shared folder will not be available offline**

  All files or programs in the selected shared folder will not be available offline.

**Notes**

- Sharing folders and drives requires special security consideration. For more information, see Securing shared resources.

- The default cache size is set to 10 percent of available drive space. You can modify this setting in the **Offline Files** tab in **Folder Options**. Cache size settings affect only the autocached files. Manually cached files and programs using the **Only the files and programs that users specify will be available offline** option are restricted only by actual disk space available to the user. For more information, see Related Topics.

- By default, Offline Files stores shared network files and programs in the *systemroot*\CSC directory on the user's hard drive. For example, if your system directory is C:\WINDOWS, your cache is C:\WINDOWS\CSC.

# Handling file conflicts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Handling file conflicts**

When you synchronize files, the files that you opened or updated while you were disconnected from the network are compared to the files that are saved on the network. As long as the same files you changed have not been changed by someone else while you were offline, your changes are copied to the network. Following are common file conflict scenarios:

- If someone else made changes to the same network file that you updated offline, you can choose to keep your version, keep the version on the network, or keep both versions. If you choose to save both versions of the file, you need to provide a different file name for your version of the file.
- If you delete a network file on your computer while working offline and, in the meantime, someone else on the network is working in that file, the file is deleted from your computer, but it is not deleted from the network.
- If you change a network file while working offline and someone else on the network deletes that file, you can choose to save your version to the network or delete it from your computer.
- If you are disconnected from the network when a new file is added to a shared network folder that you have made available offline, that new file will be added to your computer when you reconnect and synchronize.

For more information about synchronizing files with Synchronization Manager, see Synchronizing offline items and Synchronization Manager.

# Encrypting offline files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encrypting offline files**

With Offline Files, you can specify that offline files are encrypted. If you choose to encrypt offline files, only the files on your local computer are encrypted. The files on the network are not affected. Use encryption to keep offline files and folders safe from intruders who might gain unauthorized physical access to your computer.

You do not need to decrypt an encrypted file or folder before using it. However, an intruder who tries to access encrypted files or folders will be prevented from doing so.

For more information about working with encrypted files, see Encrypting File System overview and Encrypt offline files.

# Working with Web pages offline

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Working with Web pages offline**

Using Microsoft Internet Explorer, you can make Web pages or entire Web sites available offline for viewing, including related links. A wizard prompts you to specify how much content to make available offline. You can also set up a schedule to synchronize content automatically.

Making Web pages available offline is a feature of Internet Explorer and is configured separately from the Offline Files feature in Windows. The folder options you specify for Offline Files, such as encryption and synchronization, do not apply to Web pages that you make available offline. Instead, use the options that are available for offline Web pages through Internet Explorer.

For more information, see Make Web pages available offline.

# Offline Files for mobile users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Offline Files for mobile users**

If you travel frequently and rely on a portable computer to accomplish your work, Offline Files provides a way to ensure that you are working with the most current versions of the network files and that your work will be synchronized with the network when you reconnect. Using Offline Files, you can navigate through shared folders and mapped network drives regardless of whether you are connected to a network.

Follow these steps when working with Offline Files on a portable computer:

- **Make a file or folder available offline**

  Before you disconnect your portable computer from the network, you must first select the shared network files or folders that you want to use when you are disconnected and make them available offline. For more information, see Make a file or folder available offline.

- **Synchronize files when you log on to or log off from a computer**

  Set up Synchronization Manager to reflect the way you want your portable computer to synchronize your files when you log on to and log off from the network. For more information, see Synchronizing offline items when you log on to or log off from a computer.

- **Disconnect a portable computer from the network**

  After you have completed the previous tasks, you can disconnect your portable computer and begin working offline. In Windows XP Professional and the Windows Server 2003 family, the computer synchronizes automatically whenever the user suspends the computer. This policy is not controllable through the user interface.

- **Work offline**

  While working offline, you can still browse network drives and shared folders in My Computer or My Network Places. You will be able to see only those files that you made available offline and any files that you created after losing the network connection. A red X appears over any disconnected network drives. Your permissions on the network files and folders remain the same whether you are connected to the network or working offline. For more information, see Working offline.

- **Handle file conflicts**

  When you are finished working offline and you reconnect your portable computer to the network, you must synchronize your files. By synchronizing your files, you update the network files with changes that you made while working offline. By having Synchronization Manager automatically synchronize your files when you log on, you ensure that the latest file changes are reflected on the network and that any file conflicts are handled properly. For more information, see Handling file conflicts.

# User Profiles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User profiles**

User profiles automatically create and maintain the desktop settings for each user's work environment on the local computer.

- To find features that have been moved in the Microsoft® Windows Server 2003 family, see New ways to do User Profiles tasks.
- For tips about using user profiles, see User Profiles best practices.
- For help with specific tasks, see User Profiles How To....
- For general background information, see User Profiles Concepts.

# New ways to do User Profiles tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring user profiles. The user interface for performing these tasks is different in this version of Windows from the way was in Windows NT 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Add a path to a user profile | **User Manager** in **Administrative Tools** | Active Directory Users and Computers, and then click **Users**. |
| View the contents of a user profile | The profile is stored in Profiles typically in C:\Windows or C:\Winnt*UserName* | Windows Explorer. The profile is stored in the \Documents and Settings*UserName* folder. |
| Copy a user profile | **System** utility in Control Panel | System Properties, click the **Advanced** tab, and then, under **User Profiles**, click **Settings**. |
| Add a home directory to a path | **User Manager** in **Administrative Tools** | Active Directory Users and Computers, and then click **Users**. |
| Add a logon script to a user profile | **User Manager** in **Administrative Tools** | Group Policy accessed through Active Directory Users and Computers |

# User Profiles best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Allow for different hardware configurations.**

- Because user profiles can be used on various types of client computers, you should keep in mind that these client computers can have different hardware configurations, particularly different video cards and display monitors.
- Because a user profile determines screen placement and window size, a workstation's display hardware affects how well the user profile works. For example, the window setup in a user profile created for a computer with a Super VGA monitor might not look correct when loaded on a computer with a regular VGA monitor.

**Use the same type of video hardware when you create or edit a user profile for a single user.**

- When you create or edit a user profile for a single user, use a computer with the same type of video hardware as the computer that the user typically uses.

**Create a single mandatory user profile for a group of users only if they all use computers with the same type of video hardware.**

- When you create a mandatory user profile for several users, create a single user profile for the whole group of users only if they all use computers with the same type of video hardware.

**Do not use Offline Folder caching on roaming user profile shared directories.**

- It is important to turn off Offline Folder caching for shared directories where roaming user profiles are stored. If you do not turn off Offline Folder caching for a user's profile, you might experience synchronization problems when both Offline Folders and roaming user profiles try to synchronize the files in a user's profile. This does not affect your ability to use Offline Folders with redirected folders such as My Documents.

**Do not use Encrypted File System (EFS) on files in a roaming user profile.**

- The Encrypted File System (EFS) is not compatible with roaming user profiles. If you encrypt profile folders or files in the user profile using EFS, the user's profile will not roam.

**Do not set disk quotas too low for users with roaming user profiles.**

- If a user's disk quotas are set too low, roaming user profile synchronization might fail. Make sure enough disk space is allocated to allow the system to create a temporary duplicate copy of a user's profile. The temporary profile is created in the user's context as part of the synchronization process, so it debits the user's quota.

**When creating a roaming profile shared directory, limit access to only those users that need access.**

- Because a users roaming profile can contain personal information such as confidential documents and EFS certificates, care should be taken to protect access to the shared directory. Restrict access to the shared directory to only those users that need access. You can also create a security group for users that have profiles on a particular shared directory, and limit access to only those users.
- Only give users the minimum amount of permissions needed.
- When creating the shared directory, hide it by putting a **$** after the share name. This hides the shared directory from casual browsers, and it will not be visible in My Network Places.

**Use at least Windows 2000 servers to host user roaming profile shared directories.**

- Because a users roaming profile contains personal information which is copied to and from a client computer, and the server hosting the roaming profile, it is important to ensure that data is protected as it travels over the network. Potential threats to the privacy and integrity of a user's data come from intercepting the data as it passes over the network, tampering with the data as it passes over the network, and spoofing the server hosting the user's data. Features such as Kerberos, IPSec, and Server Message Block (SMB) signing included in Windows 2000 and the Windows Server 2003 family can help to secure a user's data.

**Always use the NTFS file system for volumes holding user's data.**

- Configure servers hosting roaming profiles to use the NTFS File System. Unlike FAT, NTFS supports Discretionary access control lists (DACLs) and system access control lists (SACLs), which control who can perform operations on a file and what events will trigger logging of actions performed on a file.

# User Profiles How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Create a roaming user profile
- Copy a user profile
- Delete a user profile
- Create a preconfigured user profile
- Create a mandatory user profile
- Add a home folder to a profile
- Switch between a roaming and local user profile

# Create a roaming user profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a roaming user profile**

1. Create a folder on the server where user profiles will be stored. This will be the top-level folder that contains all the individual user profiles.

2. Configure the folder as a shared folder and give all users **Full Control** permissions.

3. Open Active Directory Users and Computers and navigate to the individual's User object.

4. Right-click the user's name and click **Properties**.

5. Click the **Profile** tab. For the Profile Path, type the path to the shared folder where the user profile is to be stored. For example, for a user whose network name is pwest, the following path, **\\NetworkShare\Profiles\%username%**, would create a directory called pwest in the Profiles folder on the server used to store user profiles.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- Use a full path in each user account:

  *****Server******ShareName******UserName*

- The user profile path location can be on any server. It does not have to be a domain controller. When the user logs on, Windows checks the user's account to see if a user profile path exits. If it does, the user profile is located by the system and copied to the local computer.

- To provide better security, user profiles and home folders should be created on an NTFS volume.

- The user account can be created in any organizational unit or container within the domain. The default location for the user account is the Users folder.

- The Windows Server 2003 family does not support the use of encrypted files within the roaming user profile.

- Roaming user profiles that are used with Terminal Services clients are not replicated to the server until the interactive user logs off and the interactive session is closed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Using roaming user profiles
Working with MMC console files
Add a home folder to a profile

# Copy a user profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy a user profile**

1. Open System in Control Panel.
2. On the **Advanced** tab, under **User Profiles**, click **Settings**.
3. Under **Profiles stored on this computer**, click the user profile you want to copy, and then click **Copy To**.
4. Do one or more of the following:

   **To specify where the new profile will be saved**

   - In **Copy profile to**, type the location for the new profile, or click **Browse** to select the path.

   **To specify who is permitted to use the copied profile**

   a. In **Permitted to use**, click **Change**.

   b. In the **Select User or Group** dialog box, in **Enter the object name to select**, add the user, group, or built-in security principle or click **Object Types** to select an object type.

      To specify a domain to search, in the **Select User or Group** dialog box, click **Locations**, and then select the domain.

      To further narrow your search, in the **Select User or Group** dialog box, click **Advanced**.

   c. Click **OK**

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open System, click **Start**, click **Control Panel**, and then double-click **System**.
- To open System from a command line as an administrator, type:

  **runas /user:** *computername* **\Administrator "rundll32.exe shell32.dll,Control_RunDLL sysdm.cpl"**

- You cannot copy or delete a user profile that belongs to the currently logged on user or any user whose profile is in use.
- If you copy the profile to a new location, you must update the **User Profile Path** entry for the user's account to refer to this new location as well.
- You cannot use Windows Explorer or any other file management utility to copy user profiles.
- You can issue a network default user profile for all new users in the domain. For more information about creating a

preconfigured user profile, see **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
User profiles overview
Delete a user profile
Switch between a roaming and local user profile
System Properties Overview - Control Panel
Add a home folder to a profile
Create a preconfigured user profile

# Delete a user profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a user profile**

1. Open System in Control Panel.
2. On the **Advanced** tab, under **User Profiles**, click **Settings**.
3. Under **Profiles stored on this computer**, click the user profile you want to delete, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open System, click **Start**, click **Control Panel**, and then double-click **System**.
- To open System from a command line as an administrator, type:

  **runas /user:** *computername* **\Administrator "rundll32.exe shell32.dll,Control_RunDLL sysdm.cpl"**

- You cannot copy or delete a user profile that belongs to the currently logged on user or any user whose profile is in use.
- You can delete a roaming user profile by taking ownership of the profile's folder on the domain controller and then deleting it. This does not remove roaming profile information that has been downloaded to any of the computers in that domain. Downloaded profile information must be deleted individually from each computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
User profiles overview
Copy a user profile
Switch between a roaming and local user profile
System Properties Overview - Control Panel

# Create a preconfigured user profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a preconfigured user profile**

1. Create a new user account that will be used as a template for the preconfigured user profile.

   For more information, see Create a local user account.

2. Log on as the new user, then customize the desktop and install applications to configure this user's profile for the user profile template.

3. Log off, and then log on as the administrator.

4. Open System in Control Panel.

5. On the **Advanced** tab, under **User Profiles**, click **Settings**.

6. Under **Profiles stored on this computer**, select the user that you created in step 1, and click **Copy To**.

   - If you want a domain-wide default profile, enter the path to NETLOGON\Default User on the domain controller. This creates the default user profile for the domain.

   - If you want to change the default profile for the local computer only, copy the profile to the *systemroot*\Documents and Settings\Default User folder.

7. In the **Copy To** dialog box, under **Permitted to use**, click **Change**.

8. In the **Select User or Group** dialog box, in **Enter the object name to select**, type **Everyone**.

   This sets the profile as the default for everyone in this domain.

### Caution

- If you are using a roaming profile and install a program on one computer while simultaneously logged on to another computer, you might overwrite crucial program-related registry settings stored in your roaming profile, thus preventing you from running those programs.

  For example: You are logged on to computer A and computer B. You install a program on computer B and then log off computer B. Computer B stores the shortcuts for the application, and the registry is saved to your roaming profile. Computer A does not get updated profile information until you log off and log on again.

  When you log off from computer A, however, the computer writes to the registry stored in the roaming profile (which now includes the Microsoft Windows Installer (MSI) registration for the program you installed on computer B) with the stale registry information from computer A. The program shortcuts remain in your roaming profile but the Windows Installer data stored in the registry settings is lost, preventing you from running the programs.

  You can repair your roaming profile by repairing or reinstalling the program on computer B or by installing the program on computer A.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have

been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open System, click **Start**, click **Control Panel**, and then double-click **System**.

- To open System from a command line as an administrator, type:

  **runas /user:** *computername* **\Administrator "rundll32.exe shell32.dll,Control_RunDLL sysdm.cpl"**

- You cannot copy or delete a user profile that belongs to the currently logged on user or any user whose profile is in use.

- The first time a user logs on, a copy of the preconfigured user profile is returned from the server instead of a copy of the default profile on the local computer. Thereafter, the user profile functions the same as a standard roaming user profile does. Each time the user logs off, the user profile is saved locally and is also copied to the server.

- The Windows Server 2003 family of operating systems does not support the use of encrypted files within the roaming user profiles.

- Roaming user profiles used with Terminal Services clients are not replicated to the server until the interactive user logs off and the interactive session is closed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a user profile
Switch between a roaming and local user profile
System Properties Overview - Control Panel
Add a home folder to a profile
Create a mandatory user profile
Create a roaming user profile
User profiles overview
Using roaming user profiles

# Create a mandatory user profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a mandatory user profile**

1. Open Active Directory Users and Computers.
2. In the details section, right-click the applicable user account, and then click **Properties**.

   **Where?**

   - Active Directory Users and Computers/applicabledomain/applicable container (such as Users)/applicable user account

3. Click the **Profile** tab.
4. In **Profile path**, type the path information ending with the .man file name extension.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- A mandatory user profile is a preconfigured user profile. The user can still modify the desktop, but the changes are not saved when the user logs off. The next time the user logs on, the mandatory user profile is downloaded again. User profiles become mandatory when you rename the NTuser.dat file on the server to NTuser.man. This extension makes the user profile read-only.
- Mandatory user profiles do not allow changes to be applied to the user profile stored on the server.
- Profile management should be done preferentially by policy. Mandatory profile use, although permitted, is less manageable and more prone to create administration problems, thus it is not recommended.
- Use a full path in each user account:

  *****Server******ShareName******UserName*

  For *ShareName*, create a Profiles folder if it does not already exist, and share the folder with authenticated users allowing read-only permissions. The shared folder must be created before the user profile can be used.

- To provide better security, user profiles and home folders should be created on an NTFS volume.
- When creating a mandatory profile, make sure you set the appropriate access permissions for the user or groups of users who will use this profile.
- You can also create a mandatory user profile by using Windows Explorer to rename the NTuser.dat file to NTuser.man.
- The administrator can assign the same mandatory user profile to as many users as needed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
User profiles overview
Using roaming user profiles

# Add a home folder to a profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a home folder to a profile**

1. Open Active Directory Users and Computers.
2. In the details section, right-click the applicable user account, and then click **Properties**.

   **Where?**

   - Active Directory Users and Computers/*applicable domain*/*applicable container* (such as Users)/*applicable user account*

3. Click the **Profile** tab.
4. In **Home Folder**, type the directory folder path information.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- The Windows Server 2003 family includes a folder called My Documents, which offers an alternative to home folders, but does not replace them. All users have a My Documents folder in their user profile.
- To provide better security, user profiles and home folders should be created on an NTFS volume.
- Users can change the target folder location of their My Documents folders by clicking **Start**, right-clicking **My Documents**, clicking **Properties**, and then specifying a new location in **Target**. This capability can be enabled or disabled by the system administrator.
- You can use Group Policy to change the target folder location of any user's My Documents folder, disable a user's ability to change the target folder location, remove the My Documents icon from the desktop, and choose whether programs use My Documents or a different folder as the default location for opening and saving files.
- If you do not assign a home folder, the system uses the local user profile folder as the home folder (by default, *systemdrive*\Documents and Settings*UserName*). The home directory can use the same location as the My Documents folder.
- When you use Terminal Services on computers running Windows Server 2003 operating systems, the user profile is the default home directory.
- When creating the user's home folder, the folder corresponding to the user's account name is immediately created if it does not exist. This is in contrast to roaming user profiles, where the profile folder for the user is created after the user logs on.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing files for local user accounts
Assign a home folder to a local user account
Working with MMC console files
Create a roaming user profile

# Switch between a roaming and local user profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To switch between a roaming and local user profile**

1. Open System in Control Panel.
2. On the **Advanced** tab, under **User Profiles**, click **Settings**.
3. Under **Profiles stored on this computer**, click the user profile you want to change, and then click **Change Type**.
4. Select **Roaming profile** or **Local profile**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open System, click **Start**, click **Control Panel**, and then double-click **System**.
- To open System from a command line as an administrator, type:

  **runas /user:** *computername* **\Administrator "rundll32.exe shell32.dll,Control_RunDLL sysdm.cpl"**

- If **Roaming profile** is unavailable, the profile is a local user profile. For more information about how to create roaming user profiles, see **Related Topics**.
- If you use a roaming user profile on more than one computer simultaneously, the profile preserves the registry settings from the computer you log off from last.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Copy a user profile
Delete a user profile
System Properties Overview - Control Panel
User profiles overview
Create a roaming user profile

# User Profiles Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about user profiles.

- New features and changes to user profiles
- User profiles overview
- Understanding User Profiles
- Using User Profiles

# New features and changes to user profiles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features and changes to user profiles**

The following new features and changes in the Windows Server 2003 family affect user profiles:

**New location for User Profile configuration**

User profiles is located on the **Advanced** tab of the **System Properties** dialog box.

**Error Message changes**

All events have unique identifications and details.

**Roaming profiles are saved to the server at logoff with user registry keys open**

Previously, applications that kept registry keys open at logoff time prevented the system from saving any changes to roaming profiles until those registry keys were closed. With the Windows Server 2003 family, roaming profiles that are changed during a session will save these changes even if registry keys are open.

**Delete and Copy To button availability**

The **Delete** and **Copy To** buttons are unavailable for the profile belonging to the currently logged on user or for any user whose profile is in use.

**Group Policy changes**

The following changes to Group Policy in the Windows Server 2003 family affect user profiles:

- User profiles policies now have their own node in the Group Policy.

  **Where?**

  - *Group Policy object*/User Configuration or Computer Configuration/Administrative Templates/System/User Profiles

- **Prevent Roaming Profile changes from propagating to the server--**This policy determines if the changes a user makes to their roaming profile are merged with the server copy of their profile. If this policy is set at user logon, the user will receive their roaming profile, but any changes a user makes to their profile will not be merged to their roaming profile at logoff.

- **Add the Administrator security group to roaming user profiles--**This policy allows an administrator to choose the same behavior as Windows NT version 4.0, and permit the administrators group to have full control of the user's profile directories. In Windows 2000, the default file permissions for newly generated roaming profiles are full control, or read and write access for the user, and no file access for the Administrators group.

- **Only allow local user profiles--**This setting determines whether roaming user profiles are available on a particular computer. By default, when roaming profile users log on to a computer, their roaming profile is copied to the local computer. By using this setting, an administrator can prevent users configured to use roaming profiles from receiving their roaming profile on a specific computer.

# User profiles overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User profiles overview**
On computers running Windows Server 2003 operating systems, user profiles automatically create and maintain the desktop settings for each user's work environment on the local computer. A user profile is created for each individual user when the user logs on to a computer for the first time.

**Advantages of using user profiles**
User profiles provide several advantages:

- More than one user can use the same computer. When users log on to their individual workstations, they receive the desktop settings as they existed when they logged off.
- Customization of the desktop environment made by one user does not affect another user's settings.
- User profiles can be stored on a server so that they can follow users to any computer running a Microsoft® Windows NT® or later operating system on the network. These are called roaming user profiles.

As an administrative tool, user profiles provide these options:

- You can create a default user profile that is appropriate for the user's tasks.
- You can set up a mandatory user profile that does not save changes made by the user to the desktop settings. Users can modify the desktop settings of the computer while they are logged on, but none of these changes are saved when they log off. The mandatory profile settings are downloaded to the local computer each time the user logs on. For more information on mandatory profiles, see Create a mandatory user profile.
- You can specify the default user settings that will be included in all of the individual user profiles.

**User profile types**
A user profile defines customized desktop environments, which include individual display settings, network and printer connections, and other specified settings. You or your system administrator can define your desktop environment.

Types of user profiles include:

- **Local user profile--**A local user profile is created the first time you log on to a computer and is stored on a computer's local hard disk. Any changes made to your local user profile are specific to the computer on which you made the changes.
- **Roaming user profile--**A roaming user profile is created by the system administrator and is stored on a server. This profile is available every time you log on to any computer on the network. Changes made to your roaming user profile are updated on the server.
- **Mandatory user profile--**A mandatory user profile is a roaming profile that can be used to specify particular settings for individuals or an entire group of users. Only system administrators can make changes to mandatory user profiles.
- **Temporary user profile--**A temporary profile is issued any time that an error condition prevents the users profile from being loaded. Temporary profiles are deleted at the end of each session. Changes made by the user to their desktop settings and files are lost when the user logs off.

# Understanding User Profiles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding user profiles**
This section covers:

- Contents of a user profile
- Settings saved in a user profile

# Contents of a user profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Contents of a user profile**

Every user profile begins as a copy of Default User, which is a default user profile stored on each computer running a Windows Server 2003 family operating system. The NTuser.dat file within Default User contains Windows Server 2003 family configuration settings. Every user profile also uses the common program groups contained in the All Users folder.

The user profile folders contain various items including the desktop and Start menu. The following table lists and describes the contents of each user profile folder.

| User profile folder | Contents |
|---|---|
| Application Data | Program-specific data (for example, a custom dictionary). Program vendors decide what data to store in this user profile folder. |
| Cookies | User information and preferences. |
| Desktop | Desktop items, including files, shortcuts, and folders. |
| Favorites | Shortcuts to favorite locations on the Internet. |
| Local Settings | Application data, history, and temporary files. Application data roams with the user by way of roaming user profiles. |
| My Documents | User documents and subfolders. |
| My Recent Documents | Shortcuts to the most recently used documents and accessed folders. |
| NetHood | Shortcuts to My Network Places items. |
| PrintHood | Shortcuts to printer folder items. |
| SendTo | Shortcuts to document-handling utilities. |
| Start Menu | Shortcuts to program items. |
| Templates | User template items. |

**NTuser.dat file**

The NTuser.dat file is the registry portion of the user profile. When a user logs off of the computer, the system unloads the user-specific section of the registry (that is, **HKEY_CURRENT_USER**) into NTuser.dat and updates it. For more information about the registry, see Registry structure.

**All Users folder**

Although they are not copied to user profile folders, the settings in the All Users folder are used to create the individual user profiles. The Windows Server 2003 family supports two program group types:

- Common program groups are always available on a computer, no matter who is logged on.
- Personal program groups are private to the user who creates them.

Common program groups are stored in the All Users folder under the Documents and Settings folder. The All Users folder also contains per-computer settings for the Desktop and the **Start** menu.

## Notes

- The My Documents, My Pictures, Favorites, Start Menu, and Desktop folders are the only folders displayed in Windows Explorer by default. The NetHood, PrintHood, Local Settings, Recent, and Templates folders are hidden and do not appear in Windows Explorer. To view these folders and their contents in Windows Explorer, on the **Tools** menu, point to **Folder options**, click the **View** tab, and then click **Show hidden files and folders**.

- On computers running Windows Server 2003 operating systems with the NTFS file system, only members of the Administrators group can create, delete, or modify the common program groups.

# Settings saved in a user profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Settings saved in a user profile**

A user profile contains configuration preferences and options for each user.

The following table is a sample of the settings that are contained in a user profile.

| Source | Parameters saved |
|---|---|
| Windows Explorer | All user-definable settings for Windows Explorer. |
| My Documents | User-stored documents. |
| My Pictures | User-stored picture items. |
| Favorites | Shortcuts to favorite locations on the Internet. |
| Mapped network drive | Any user-created mapped network drives. |
| My Network Places | Links to other computers on the network. |
| Desktop contents | Items stored on the Desktop and shortcuts. |
| Screen colors and fonts | All user-definable computer screen colors and display text settings. |
| Application data and registry hive | Application data and user-defined configuration settings. |
| Printer settings | Network printer connections. |
| Control Panel | All user-defined settings made in Control Panel. |
| Accessories | All user-specific program settings affecting the user's Windows environment, including Calculator, Clock, Notepad, and Paint. |
| Windows Server 2003 family-based programs | Any program written specifically for Windows Server 2003 family can be designed so that it tracks program settings on a per-user basis. If this information exists, it is saved in the user profile. |
| Online user education bookmarks | Any bookmarks placed in the Windows Server 2003 family Help system. |

**Note**

- The My Documents, My Pictures, Favorites, Start Menu, and Desktop folders are the only folders displayed in Windows Explorer by default. The NetHood, PrintHood, Recent, and Templates folders are hidden and do not appear in Windows Explorer. To view these folders and their contents in Windows Explorer, from the **Tools** menu, point to **Folder options**, click the **View** tab, and then click **Show hidden files and folders**.

# Using User Profiles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using user profiles**
This section covers:

- Using roaming user profiles
- Local user profiles

# Using roaming user profiles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using roaming user profiles**

Roaming user profiles enable users to log on to computers in a domain while preserving their user profile settings. User profiles are stored at an administrator specified server location. When a user logs on and has been authenticated within the directory service, the user profile, including user settings and documents, are copied to the local computer. User profile changes made on the local computer are then captured. Changes will then be copied to the user profile stored on the server and be applied the next time the user logs on. Mandatory user profiles are similar, but do not allow changes to be copied to the user profile stored on the server. For more information about implementing roaming user profiles, see Create a roaming user profile and Create a mandatory user profile.

From Active Directory, you can assign a server location for user profiles. If you type a user profile path into a user's domain account, a copy of the user's local user profile is saved both locally and in the user profile path location when the user logs off. The next time that user logs on, the profile stored in the user profile path location is compared to the copy in the local user profile folder, and the most recent copy of the profile is opened. The local user profile becomes a roaming user profile because of the centralized domain location. The user's settings and documents are available wherever the user logs on.

If the server is not available, the local cached copy of the roaming user profile is used. If the user has not logged on to the computer before, a new local user profile is created. In either case, if the centrally stored user profile is not available when the user logs on, it is not updated when the user logs off. If the user profile is not downloaded because of server problems, it is not loaded when the user logs off.

# Local user profiles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local user profiles**

The location of a user profile depends on whether the user is starting from a clean installation of a Windows Server 2003 family operating system or an upgrade from a previous version of Windows. The following table lists the operating system and the location of user profiles for that scenario.

| Operating system | User profile location |
|---|---|
| Windows Server 2003 family clean installation (no previous operating system) | *systemdrive*\Documents and Settings (for example, C:\Documents and Settings) |
| Windows Server 2003 family upgrade from Windows 2000 | *systemdrive*\Documents and Settings (for example, C:\Documents and Settings) |
| Windows Server 2003 family upgrade from Windows NT 4.0 | *systemroot*\Profiles (for example, C:\Windows or C:\WINNT\Profiles) |
| Windows Server 2003 family upgrade from Windows NT 3.51 | *systemdrive*\Documents and Settings (for example, C:\Documents and Settings) |
| Windows Server 2003 family upgrade from Windows 95 or Windows 98 | *systemdrive*\Documents and Settings (for example, C:\Documents and Settings) |

When no preconfigured server-based roaming user profile exists for a user, a user profile folder is created for that user when the user logs on to the computer for the first time. The contents of the Default User folder are then copied to the new user profile folder. The user profile, along with the common program group settings in the All Users folder, creates the user's desktop. When the user logs off, any changes made to the default settings during the session are saved to the new user profile folder. The user profile in the Default User folder remains unchanged.

If the user has a user account on the local computer in addition to a domain user account, or more than one domain user account, the local user profile is different for each account because different user profiles are generated for each user who logs on. When the user logs off, changed settings are saved to only one user profile, depending on which account the user used to log on to the computer.

# Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Connections**
Network Connections provides connectivity between your computer and the Internet, a network, or another computer. With Network Connections, you can configure settings to reach local or remote network resources or functions.

Network Connections combines Microsoft® Windows® NT 4.0 Dial-Up Networking with features that were formerly located in the Network Control Panel, such as network protocol and service configuration. Each connection in the Network Connections folder contains a set of features that creates a link between your computer and another computer or network. By using Network Connections, performing a task, such as modifying a network protocol, is as easy as right-clicking a connection and then clicking **Properties**.

- Before using Network Connections, see Checklist: Configuring network connections.
- Before setting up a home or small office network, see Checklist: Setting up a home or small office network.
- To find features that have been moved, see New ways to do Network Connections tasks.
- For tips about using Network Connections, see Network Connections Best Practices.
- For help with specific tasks, see Network Connections How To....
- For general background information, see Network Connections Concepts.
- For problem-solving instructions, see Network Connections Troubleshooting.

# Checklist: Configuring network connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring network connections**

| | Step | Reference |
|---|---|---|
| | Review key concepts. | Network Connections Overview |
| | Connect to a local area network. | Make a local area connection |
| | Connect to a network by using a phone line. | Make a dial-up connection by using a phone line;<br>Connect to a remote network |
| | Connect to an Internet service provider (ISP). | Make an Internet connection;<br>Connect to a remote network |
| | Connect to a private network through the Internet (virtual private network). | Make a virtual private network (VPN) connection;<br>Connect to a remote network |
| | Create a direct connection to another computer. | Make a direct network connection;<br>Connect to a remote network |
| | Let other computers connect to this computer. | Accept incoming network connections |
| | Monitor a connection. | Monitor a connection |

# Checklist: Setting up a home or small office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Setting up a home or small office network**

| Step | Reference |
|---|---|
| Decide whether to use home and small office networking or Routing and Remote Access for your home or small office network. | Home and small office networking compared to Routing and Remote Access |
| If you decide to use home and small office networking, review home and small office networking concepts. | Home and Small Office Networking; Using Home and Small Office Networking Features |
| Decide how you want to save your Internet connection credentials (you will be prompted to select credential settings when you fist attempt to activate your Internet connection or when you run the New Connection Wizard). | Dial-up credentials |
| Familiarize yourself with the icons that indicate that Internet Connection Firewall (ICF), Windows Firewall, Network Bridge, or Internet Connection Sharing (ICS) have been enabled on a network connection. | Network connection icons |
| Enable ICF or Windows Firewall to protect your home or small office network. | Enable or disable ICF |
| Set up Network Bridge to connect network segments. | Set up Network Bridge |
| Enable ICS to provide a shared Internet connection. | Enable ICS |

**Notes**

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

# New ways to do Network Connections tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

This table lists common tasks for Network Connections. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Connect to a dial-up server | **Dial-Up Networking** in My Computer | Network Connections. For more information, see Connect to a remote network. |
| Connect to a network by using a modem or ISDN line | **Dial-Up Networking** in My Computer (click **New**) | Network Connections. For more information, see Make a dial-up connection by using a phone line. |
| Connect to a network by using a virtual private network (VPN) | **Dial-Up Networking** in My Computer (click **New**) | Network Connections. For more information, see Make a virtual private network (VPN) connection. |
| Configure a computer for incoming connections | **Network** in Control Panel | Network Connections. For more information, see Accept incoming network connections. |
| Configure the TCP/IP protocol | **Network** in Control Panel | Network Connections. For more information, see Configure TCP/IP settings. |
| Configure a connection | **Dial-Up Networking** in My Computer (click **More**) | Network Connections. For more information, see Configure a connection to a remote network. |
| Add a client or service | **Network** in Control Panel | Network Connections. For more information, see Add a network component. |
| Add a protocol | **Network** in Control Panel | Network Connections. For more information, see Add a network component. |
| Add optional networking components | **Network** in Control Panel | Network Connections. For more information, see Add optional networking components. |
| Monitor incoming or outgoing connections | **Dial-Up Monitor** in Control Panel | Network Connections. For more information, see Monitor a connection. |

# Network Connections Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

The following list provides best practices for implementing Network Connections:

- **When using multiple network adapters, rename each local area connection.**

  Your operating system detects network adapters and automatically creates a local area connection in the Network Connections folder for each network adapter. If more than one network adapter is installed, you can eliminate possible confusion by immediately renaming each local area connection to reflect the network that it connects to.

  You need to add or enable the network clients, services, and protocols that are required for each connection. When you do so, the client, service, or protocol is added or enabled in all other network and dial-up connections.

- **Verify required connection settings for your network adapter.**

  If your network administrator or Internet service provider (ISP) requires static settings, you may need one or more of the following:

  - A specific IP address.
  - DNS addresses.
  - DNS domain name.
  - Default gateway address.
  - WINS addresses.

  DHCP is enabled by default. This gets your IP address and other configuration information automatically. Automated IP settings are used for all connections, and they eliminate the need to configure settings such as DNS, WINS, and so on.

  For more information about enabling DHCP, see Configure TCP/IP settings.

- **Create dial-up, VPN, or direct connections by using the New Connection wizard.**

  After you create a connection, you can copy the connections, rename them and modify the connection settings. By doing so, you can easily create different connections to accommodate multiple modems, ISPs, dialing profiles, and so on.

  For more information, see Copy a network or dial-up connection.

- **Specify the order in which network providers and protocols are accessed.**

  By changing the order of protocols bound to those providers, you can improve performance. For example, if your LAN connection is enabled to access NetWare and Microsoft Windows networks, which use IPX and TCP/IP, but your primary connection is to a Microsoft Windows network that uses TCP/IP, you can move **Microsoft Windows Network** to the top of the **Network providers** list on the **Provider Order** tab, and move **Internet Protocol (TCP/IP)** to the top of the **File and Printer Sharing for Microsoft Networks** binding on the **Adapters and Bindings** tab.

  The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

  For more information about modifying the order of network providers, see Modify the network provider order. For more information about modifying the order of protocol bindings, see Modify the protocol bindings order.

- **Only install and enable the network protocols that you need.**

  Limiting the number of protocols on your computer enhances its performance and reduces network traffic.

  If your computer encounters a problem with a network or dial-up connection, it attempts to establish connectivity by using every network protocol that is installed and enabled. By only installing and enabling the protocols that your system can use, the operating system does not attempt to connect with protocols it cannot use, and returns status information to you more efficiently.

# Network Connections How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Connect to a remote network
- Disconnect from a remote network
- Add optional networking components
- Make Network and Dial-up Connections
- Configure Network and Dial-up Connections
- Monitor Network Connections
- Set up and Manage Home and Small Office Networking

# Connect to a remote network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a remote network**

1. Open Network Connections.

2. Double-click the connection you want to connect to the network.

3. If you are prompted, in the **Connect***connection type* dialog box, type your user name, password, and logon domain.

   If **Logon domain** does not appear in this dialog box, and you want to log on to a domain, type your user name and the domain name in one of two ways:

   - Your user principal name prefix (your user name) and your user principal name suffix (your domain name), joined by the "at" sign (@). For example, user@sales.westcoast.microsoft.com.

   - Your domain name and your user name, separated by the backslash (\) character. For example, sales\user.

     Note that the suffix in the first example is a fully-qualified DNS domain name. Your administrator might have created an alternative suffix to simplify the logon process. For example, creating a user principal name suffix of "microsoft" allows the same user to log on by using the much simpler user@microsoft.com. For more information, see your system administrator.

   Instead of having to type your domain name, you can also enable Windows to include your logon domain to the **Connect***connection type* dialog box.

   a. Click **Properties** if the button is enabled.

   b. On the **Options** tab, select the **Include Windows logon domain** check box.

   Once you are connected to the network, you can minimize your connection window and use e-mail, Windows Explorer, and so on.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- To connect to a remote network you can also right-click the connection that you want to connect to the network and click **Connect**.

- Authentication happens at two levels: logging on to the local computer domain and logging on as a valid user.

- If you are on a local area network, your computer logs on to the local computer domain automatically even before you see the logon prompt. When you type your user credentials at the logon prompt, the computer checks your user credentials, opens a user context, and makes your personal data available to you. You are now logged on at both levels.

- If you are not on a local area network, the attempt to log on to the local computer domain fails silently. In this case, user logon can proceed in two ways.

  - When you log on at the logon prompt, the computer checks your cached credentials from the last time you were actually online with the domain controller. If they match, a user context associated with your name is opened. You

are now logged on to the computer as a valid user but you are not logged on to the network. To access your personal files you can now create a remote access connection by using your remote access credentials or by using your cached credentials if your computer is a member of the domain.

- Select the **Log on using dial-up networking** check box if your computer is a member of the domain. In this case, the computer will attempt to make a remote access connection before processing your credentials. Once connected it uses the user logon credentials as though you were on a local area network.

- The logon domain name you type should name the domain that the server is in. This is not the DNS domain name given by some PPP/SLIP providers.

- If you are using IP connectivity on your local network connection and on your remote connection, you may not be able to see all computers on your local network. This is because after you connect, your remote connection becomes your default path for network routing. Consequently, you see computers on the remote network, and you see other computers on the same LAN segment to which your computer is connected. But you cannot communicate with computers on networks that were previously reached through a router on your local LAN.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Automating a Network or Dial-up Connection
Authentication process
Make a virtual private network (VPN) connection

# Disconnect from a remote network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect from a remote network**

1. Open Network Connections.
2. Right-click the connection you want to disconnect, and then click **Disconnect**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you disable a local area connection, it disables the network adapter driver until the connection is reactivated.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a remote network
Make a virtual private network (VPN) connection

# Add optional networking components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add optional networking components**

1. Open Network Connections.
2. On the **Advanced** menu, click **Optional Networking Components**.
3. Do one of the following:
   - To install all optional networking components, select the **Management and Monitoring Tools**, **Networking Services**, and **Other Network File and Print Services** check boxes, and then click **Next**.
   - To install specific optional networking components, select its corresponding check box, and then click **Details**.

     In **Subcomponents of***component*, select the check boxes next to the components you want to install, and then click **OK**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Examples of optional networking components include Simple Network Management Protocol (SNMP), Simple TCP/IP Services, and Print Services for UNIX.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Optional networking components
New ways to do Network Connections tasks

# Make Network and Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Make network and dial-up connections**

- Make a virtual private network (VPN) connection
- Make a local area connection
- Manage Network Connections
- Set up ISDN Connections
- Set up Dial-up Connections
- Configure Dial-up Options
- Make Internet Connections
- Make Direct Network Connections
- Set up Incoming Connections

# Make a virtual private network (VPN) connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a virtual private network (VPN) connection**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Connect to the network at my workplace**, and then click **Next**.
4. Click **Virtual Private Network connection**, and click **Next**.
5. Type a name for this connection in **Company Name**, and click **Next**.
6. If you have already established a dial-up connection, do one of the following:
   - If you need to establish a connection with your ISP or some other network before tunneling to your destination computer or network, click **Automatically dial this initial connection**, click a connection in the list, and then click **Next**.
   - If you do not want to automatically dial an initial connection, click **Do not dial the initial connection**, and then click **Next**.
7. Type the host name or IP address of the VPN server to which you are connecting, and then click **Next**.
8. Do one of the following:
   - If you want this connection to be made available to all users on this computer, click **Anyone's use**, and then click **Next**.
   - If you want to reserve the connection for your own use, click **My use only**, and then click **Next**.
9. If you want to create a shortcut for this connection on your desktop, select the **Add a shortcut to this connection to my desktop** check box.
10. Click **Finish**.

### Notes

- To make the connection available to all users, you must be logged on as Administrator or as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You can create multiple VPN connections by copying them in the Network Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple hosts, security options, and so on.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Virtual Private Network (VPN) Connections
Troubleshooting network and dial-up connections
PPP operation and protocols
Outsourcing dial-up networks
Point-to-Point Protocol (PPP)

# Make a local area connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a local area connection**

- Typically, most computers running a member of the Windows Server 2003 family are connected to a local area network. When you start your computer, your network adapter is detected and the local area connection automatically starts. Unlike other types of connections, the local area connection is created automatically, and you do not have to click the local area connection in order to start it.

Once the local area connection is running, you can choose to disable it. To do so, do the following.

1. Open Network Connections.
2. Right-click the **LAN or High Speed Internet** connection that you want to disable.
3. Click **Disable**.

If you need to enable a disabled **LAN or High Speed Internet** connection, do the following.

1. Right-click the connection that you want to enable.
2. Click **Enable**.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- A local area connection is automatically created for each network adapter that is detected.
- If more than one network adapter is installed, you can eliminate possible confusion by immediately renaming each local area connection to reflect the network that it connects to.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using local area connections
View the status of a local area connection
Set up Network Components
Network connection icons

# Manage Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage network connections**

- Copy a network or dial-up connection
- Show a network connection icon in the notification area
- Rename a network connection icon
- Repair a LAN or high-speed Internet connection

# Copy a network or dial-up connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy a network or dial-up connection**

1. Open Network Connections.
2. Right-click your dial-up, VPN, or direct connection, and then click **Create Copy**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You cannot copy local area and incoming connections.
- You can create multiple dial-up, virtual private network, or direct connections by copying them in the Network Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple modems, ISPs, dialing profiles, and so on.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Network connection icons

# Show a network connection icon in the notification area

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To show a network connection icon in the notification area**

1. Open Network Connections.
2. Right-click the icon for the network connection that you want to display in the notification area, and then click **Properties**.
3. Select the **Show icon in notification area when connected** check box, and then click **OK**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If a network connection icon is displayed in the notification area, you can access basic connection and statistical information by pointing to the icon.
- If the **Show icon in notification area when connected** check box is selected for the shared Internet connection, an icon for the shared Internet connection appears in the notification area each time that connection is enabled. When the shared Internet connection is disabled, the icon disappears from the notification area.
- It is a good practice to select this check box when you are using wireless connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Home and small office networking features
Rename a network connection icon

# Rename a network connection icon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a network connection icon**

1. Open Network Connections.
2. Right-click the icon for the network connection that you want to rename, and click **Rename**.
3. Type a new name for the selected network connection icon.

## Notes

- You must be logged on as a member of the Administrators group to rename a network connection icon. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You should rename each network connection icon with a name that is easy to remember.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a network connection icon in the notification area
Home and small office networking features

# Repair a LAN or high-speed Internet connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To repair a LAN or high-speed Internet connection**

1. Open Network Connections.
2. Before using the **Repair** option, try the following steps:
   - Check the status indicator lights on the DSL or cable modem to verify that the connection to your Internet service provider (ISP) is working.
   - Turn off the DSL or cable modem, and then turn it on again.
   - If your Internet connection is not working after these steps, contact your ISP.
3. If media is connected properly, under **LAN or High-Speed Internet**, right-click the connection that you want to repair, and then click **Repair**.

## Notes

- You must be a member of the Administrators group, the Power Users group, or the Network Configuration Operators group to complete this procedure. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you have more than one network adapter installed in your computer, it is recommended that you rename your Internet connection to easily identify it. For example, you could call your Internet connection "Internet connection" or use the name of your ISP.
- The **Repair** option is not available for media disconnected connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet service providers
Copy a network or dial-up connection
Network connection icons

# Set up ISDN Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up ISDN connections**

- Make a dial-up connection by using an ISDN line
- Install an ISDN adapter
- Configure ISDN settings

# Make a dial-up connection by using an ISDN line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a dial-up connection by using an ISDN line**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. In **Network Connection Type**, do one of the following:
   - Click **Connect to the Internet**. In **Internet Connection**, click **Connect using a dial-up modem**, and then click **Next**.
   - Click **Connect to the network at my workplace**. In **Network Connection**, click **Dial-up connection**, and then click **Next**.
4. Follow the instructions in the New Connection wizard.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To enable automatic dialing for all users, create a connection, save your user name and password for all users, and then enable Internet Connection Sharing for the new connection.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using ISDN lines
Configuring multiple device dialing
Install an ISDN adapter
Configure ISDN settings
Troubleshooting network and dial-up connections

# Install an ISDN adapter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install an ISDN adapter**

1. Install an Integrated Services Digital Network (ISDN) adapter and start your computer. If you install more than one ISDN adapter, you must restart your computer after you have installed the ISDN drivers. Otherwise, all available ISDN ports may not show up when you configure your network connection for ISDN.

2. You can use Device Manager to configure your system so that your ISDN adapter knows what type of telephone switch it is connected to. The switch type simply refers to the brand of equipment and software revision level that the telephone company uses to provide you with ISDN service. There are only a few types of switches in the world and usually just one in countries/regions other than the United States.

## Notes

- To use Device Manager, you must be logged on as an administrator or a member of the Administrators group.

- To open Device Manager, click **Start**, right-click **My Computer**, click **Manage**, and then click **Device Manager**.

- If your ISDN adapter is internal to your computer, it appears in **Network adapters**. If the ISDN adapter is external to your computer, it appears in **Modems**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure ISDN settings
ISDN access
Using ISDN lines
Set up ISDN Connections

# Configure ISDN settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure ISDN settings**

1. Open Network Connections.
2. Right-click the dial-up connection that uses ISDN, and then click **Properties**.
3. On the **General** tab, in **Connect using**, click the ISDN device, and then click **Configure**.
4. In the **ISDN Configure** dialog box, do one or both of the following:
   - In **Line type**, click the line type you want to use. The line types are listed from highest to lowest quality.
   - If you want to start negotiation with the line type you selected, and then negotiate with a lower-quality line type depending on the line's condition, select the **Negotiate line type** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Depending on the type of ISDN adapter that you are using, when you click **Configure** in step 3, you may be presented with **Modem Configuration** settings.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using ISDN lines
Troubleshooting network and dial-up connections

# Set up Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up dial-up connections**

- Make a dial-up connection by using a phone line
- Configure a modem for a dial-up connection
- Make a dial-up connection by using X.25
- Start Remote Access Auto Connection Manager

# Make a dial-up connection by using a phone line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a dial-up connection by using a phone line**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. In **Network Connection Type**, do one of the following:
   - Click **Connect to the Internet**. In **Internet Connection**, click **Connect using a dial-up modem**, and then click **Next**.
   - Click **Connect to the network at my workplace**. In **Network Connection**, click **Dial-up connection**, and then click **Next**.
4. Follow the instructions in the New Connection wizard.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To enable automatic dialing for all users, create a connection, save your user name and password for all users, and then enable Internet Connection Sharing for the new connection.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using phone lines and modems
Using an intermediary device
Install a modem
Configure a modem for a dial-up connection
Install a modem

# Configure a modem for a dial-up connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a modem for a dial-up connection**

1. Open Network Connections.
2. Right-click the dial-up connection you want to configure, and then click **Properties**.
3. On the **General** tab, under **Connect using**, click the modem you want to configure, and then click **Configure**.
4. Under **Hardware Features**, select the check box options you want to enable.
5. If you want to enable the modem speaker, select the **Enable modem speaker** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To ensure compatibility, you need to use the same kind of modem as the one connected to the remote access server, select the same initial speed, and enable the same features. If you do not select the same model, at least select a modem with the same ITU-T standard as the modem on the server.
- Selecting a feature that is not supported by your modem has no effect on its performance.
- Network Connections automatically configures connections according to the devices that are available. For example, you might use a laptop and docking station when you are in the office, with your connection configured to use the docking station modem. When you are on the road, and the docking station modem is not available, the connection is automatically configured to use the laptop's PCMCIA modem. When you return to the office and re-dock your laptop into its docking station, Network Connections detects that the docking station modem is available again, and automatically reconfigures the connection to use the docking station modem.
- Modem changes made as described in this procedure are unique to this connection. To make global changes, you will need to change each modem.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install a modem
Information about specific modems

# Make a dial-up connection by using X.25

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a dial-up connection by using X.25**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. In **Network Connection Type**, click **Connect to the Internet** or **Connect to the network at my workplace**, and then click **Next**.
4. In **Network Connection**, click **Connect using a dial-up modem** or **Dial-up connection**, and then click **Next**.
5. Follow the instructions in the New Connection wizard.
6. When the wizard is finished, right-click the new connection, and then click **Properties**.
7. On the **Options** tab, click **X.25**.
8. In **Network**, click your X.25 network provider.
9. In **X.121 address**, type the X.121 address (the X.25 equivalent of a phone number) for the server you want to call.
10. If required by the X.25 host computer, in **User Data**, type additional connection information.
11. In **Facilities**, type any additional facility options you want to request from your X.25 provider. (For example, some providers support the **/r** option to specify reverse charging.)

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To enable automatic dialing for all users, create a connection, save your user name and password for all users, and then enable Internet Connection Sharing for the new connection.
- For details on facility options, consult your X.25 documentation or provider.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Using X.25
Troubleshooting network and dial-up connections

# Start Remote Access Auto Connection Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Remote Access Auto Connection Manager**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. In the details pane, right-click **Remote Access Auto Connection Manager**, and then click **Start**.

   Under the **Status** column, **Started** appears.

### Note

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure redial options
Configuring redial options
Working with MMC console files

# Configure Dial-up Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure dial-up options**

- Configure client callback options
- Configure autodial
- Enable logging
- Enable multiple device dialing
- Configure phone number modifiers
- Configure redial options
- Assign multiple phone numbers to a connection
- Enable operator-assisted calls or manual dialing

# Configure client callback options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure client callback options**

1. Open Network Connections.
2. On the **Advanced** menu, click **Remote Access Preferences**.
3. On the **Callback** tab, do one of the following:
   - If you do not want to use callback, click **No callback**.
   - If you want to decide whether to use callback at the time you connect, click **Ask me during dialing when the server offers**.
   - If you want to consistently use callback, click **Always call me back at the number(s) below**, and select the modem or device you want to be called back.

     If **Phone number** is blank for the device you have selected, click **Edit**, and then type the number.

     If you want to remove a modem or device from the list of possible callback devices, click the modem or device, and then click **Delete**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Your system administrator can override any callback settings that you configure. Through server settings, the administrator can deny the use of callback, allow you to set callback options, or require callback to a particular number. Contact your system administrator to determine your callback options.
- To use callback, Link Control Protocol (LCP) extensions must be enabled.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Callback
Start requesting LCP extensions in PPP
PPP operation and protocols

# Configure autodial

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure autodial**

1. Open Network Connections.
2. On the **Advanced** menu, click **Remote Access Preferences**.
3. On the **Autodial** tab, select the check box next to the location where you want to enable autodial, and then click **OK**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Autodial is enabled by default in Windows XP. In support of this feature, the Remote Access Auto Connection Manager is also started by default.
- Autodial maps the resources that are required to initially make a connection, but it does not map resources after the connection is made. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Autodial
Automating a Network or Dial-up Connection
Start Remote Access Auto Connection Manager

# Enable logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable logging**

1. Open Network Connections.
2. On the **Advanced** menu, click **Remote Access Preferences**.
3. On the **Diagnostics** tab, do one or more of the following:
   - To enable logging for all remote access connections, select the **Enable logging** check box.
   - To clear logs created during a previous connection, click **Clear Log**.
   - To generate a report, click **Generate Report**.
4. If you clicked **Generate Report** in the previous step, do one of the following:
   - To view the diagnostic report, click **View report**.
   - To save the report to a file, click **Save report to this file**, and specify the location and the name of the file you want to create. Specify the location by typing it or by clicking **Browse** and browsing to the location you want.
   - To send the report to an e-mail address, click **Send report to this e-mail address**, and type the e-mail address to which you want to send the report.
   - To quickly generate a more concise report, select the **Generate summary report** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Creating a diagnostic report in HTML can take several minutes. If you need only configuration and installation information, you can create a summary report.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure autodial
Configure callback for incoming connections
Logging remote access connections

# Enable multiple device dialing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable multiple device dialing**

1. Open Network Connections.
2. Right-click the connection on which you want to enable the dialing of multiple devices, and then click **Properties**.
3. On the **Options** tab, in **Multiple devices**, do one of the following:
   - If you want to dial only the first available device, click **Dial only first available device**.
   - If you want to use all of your devices, click **Dial all devices**.
   - If you want to dynamically dial and hang up devices as needed, click **Dial devices only as needed**, and then click **Configure**.

     In **Automatic dialing**, click the **Activity at least** percentage and **Duration at least** time you want to set. Another line is dialed when connection activity reaches this level for the amount of time that you specify.

     In **Automatic hangup**, click the **Activity no more than** percentage and **Duration at least** time you want to set. A device is disconnected when connection activity decreases to this level for at least the amount of time that you specify.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you selected **Dial devices only as needed**, the last multilinked device ignores the **Automatic hangup** setting, and a twenty-minute time-out is used for the last device.
- If you use multiple devices to dial a server that requires callback, then only one of your multilinked devices is called back. This is because only one number is stored in a user account. Therefore, only one device connects and all other devices fail to complete the connection, and your connection loses Multilink functionality.

  You can avoid this problem if the multilinked phonebook entry is an ISDN with two channels that have the same phone number.
- Multiple device dialing is available only if multiple adapters are available on the computer.
- If you select **Dial all devices**, dropped links in the multilinked bundle are not automatically reinitialized. You can force links to reinitialize by selecting **Dial devices only as needed**, then **Configure**, and then setting easily achieved **Automatic dialing** conditions which cause another line to be dialed. For example, set **Activity at least** to 1% and **Duration at least** to 3 seconds.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Configuring multiple device dialing
ISDN access
Incoming connections
Configuring dial-up options
Using ISDN lines

# Configure phone number modifiers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure phone number modifiers**

1. Open Network Connections.
2. Right-click the dial-up connection you want to configure, and then click **Properties**.
3. On the **General** tab, modify the phone number by using one or more characters from the following list.

| Character | Function |
|---|---|
| , (comma) | Pauses briefly (two seconds for most modems) before continuing. |
| P | Switches from touch-tone to pulse/rotary dialing. |
| T | Switches from pulse/rotary to touch-tone dialing. |
| $ | Waits for a calling card prompt tone. |

For example, when you dial in from a hotel room, you may have to add 9 to the number you are calling to get an outside line. Your number may look similar to the following:

9,555-0100

The comma after the 9 gives a pause long enough to get an outside line before continuing to dial the full number.

**Notes**

- To configure phone number modifiers, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- By selecting the **Use dialing rules** check box, and then clicking **Dialing Rules**, you can create a **Location** for which specific dialing rules automatically apply. For example, if you frequently connect from home and need to disable call waiting, you can create a location for use from home that automatically dials the code to disable call waiting before the connection call is placed.
- If you see unfamiliar characters being dialed before your number, check to make sure the **Use dialing rules** check box is not selected. If it is selected, a location may be enforcing a dialing rule that you do not need. For example, a location may be selected that always dials 1 to accommodate long distance dialing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Specify area code rules
Use a calling card for long distance dialing
Change dialing properties for a location
Create a new location
Specify numbers to dial an outside line

Modify an existing calling card
Create a new calling card

# Configure redial options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure redial options**

1. Open Network Connections.
2. Right-click the dial-up connection you want to configure, and then click **Properties**.
3. On the **Options** tab, in **Redialing options**, click the arrows to set the number of times the dial-up connection is automatically redialed if the first attempt to connect fails.
4. In **Time between redial attempts**, click the down arrow to set the amount of time to pause between the end of one attempt and the beginning of another.
5. In **Idle time before hanging up**, click the down arrow to set the amount of idle time before a client hangs up a connection.
6. If you want the connection to automatically redial if the line is dropped, select the **Redial if line is dropped** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The **Time between redial attempts** pause lets the device reset itself before redialing. The default is one minute. If that is not enough time, increase this setting. You can also experiment with shorter times, but if you make the pause too short, the device does not have time to reset itself.
- Keep in mind that the remote access server also has a disconnect timer. If **Idle time before hanging up** is set to a high time-out number, there is no guarantee that the connection will not disconnect before the time passes. This is because the server may choose to disconnect you before the time set here.
- If you are using Internet Explorer, the browser's own settings for dial-up connections may conflict with these redial options.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring redial options
Start Remote Access Auto Connection Manager

# Assign multiple phone numbers to a connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign multiple phone numbers to a connection**

1. Open Network Connections.
2. Right-click the dial-up connection to which you want to assign multiple phone numbers, and then click **Properties**.
3. On the **General** tab, click **Alternates**.
4. If you want the connection to attempt a different phone number if the first number in the list is not successful, select the **If number fails, try next number** check box.
5. If you want the connection to move the first successful phone number to the top of the list under **Phone numbers** and use it on the next connection attempt, click **Move successful number to top of list**.
6. To add a new phone number to the list, click **Add**, and in **Phone number**, type the phone number.
7. If you want the connection to use established area code and dialing rules, select the **Use dialing rules** check box.
8. When you are finished adding the new phone number, click **OK**.
9. If you want to change the order of phone numbers under **Phone numbers**, click a phone number, and then click the **Up** button or **Down** button until the phone number is in the correct order.
10. If you want to remove a phone number under **Phone numbers**, click the phone number, and then click **Delete**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you have alternate phone numbers defined for a connection, you can specify a particular phone number at the time you initiate your connection.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Configure Dial-up Options
Configuring dial-up options

# Enable operator-assisted calls or manual dialing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable operator-assisted calls or manual dialing**

1. Open Network Connections.
2. On the **Advanced** menu, click **Operator-Assisted Dialing**.
3. Double-click the connection you want to dial.
4. Pick up the telephone handset, and then dial the number or ask the operator to dial it for you.

   The number assigned to the entry is displayed in the dialog box for easy reference.
5. Immediately after you have finished dialing, click **Dial**.
6. Hang up the handset only after the modem takes control of the line, which is typically signaled with a click followed by silence.

   It is always safe to replace the handset once Network Connections begins verifying your user name and password. The status message will remind you of this.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If operator-assisted dialing is enabled, a check mark appears next to **Operator-Assisted Dialing** on the **Advanced** menu.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Make Internet Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Make Internet connections**

- Make an Internet connection
- Make a broadband connection using Point-to-Point Protocol over Ethernet (PPPoE)

# Make an Internet connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make an Internet connection**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Connect to the Internet**, and then click **Next**.
4. Click **Connect using a dial-up modem**, and click **Next**.
5. If prompted, select the check box for the modem you want to use for this connection, and click **Next**.
6. Type the name of your ISP in **ISP Name**, and click **Next**.
7. Type the phone number of your ISP, including the area code, in **Phone number**, and click **Next**.
8. Make the connection available to any user or only to yourself by selecting **Anyone's use** or **My use only**.
9. In **User name**, type the account name your ISP assigned to you. In **Password** and **Confirm password**, type the password for your account name, and click **Next**.
10. If you want to create a shortcut for this connection on your desktop, select the **Add a shortcut to this connection to my desktop** check box.
11. Click **Finish.**

## Notes

- To make a connection available to any user, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To enable automatic dialing for all users, create a connection, save your user name and password for all users, and then enable Internet Connection Sharing for the new connection.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Internet Connections
Troubleshooting network and dial-up connections
Make a virtual private network (VPN) connection

# Make a broadband connection using Point-to-Point Protocol over Ethernet (PPPoE)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a broadband connection using Point-to-Point Protocol over Ethernet (PPPoE)**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Connect to the Internet**, and then click **Next**.
4. To create a high-speed connection that requires you to sign in, click **Connect using a broadband connection that requires a user name and password**, and click **Next**.
5. Type the name of your ISP in **ISP Name**, and click **Next**.
6. Do one of the following:
   - If you want this connection to be made available to all users on your network, click **Anyone's use**, and then click **Next**.
   - If you want to reserve the connection for your own use, click **My use only**, and then click **Next**.
7. In **User name**, type the account name your ISP assigned to you. In **Password** and **Confirm password**, type the password for your account name, and click **Next**.
8. To create a shortcut for this connection on your desktop, select the **Add a shortcut to this connection to my desktop** check box.
9. Click **Finish**.

## Notes

- To make a connection available to any user, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Point-to-Point Protocol over Ethernet (PPPoE) Protocols
Point-to-Point Protocol (PPP)
Make Internet Connections

# Make Direct Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Make direct network connections**

- Make a direct network connection
- Make a direct network connection by using a DirectParallel cable

# Make a direct network connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a direct network connection**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Set up an advanced connection**, and then click **Next**.
4. Click **Connect directly to another computer**, click **Next**, and then follow the instructions in the New Connection Wizard.

## Notes

- To create a direct network connection that acts as a host, you must be logged on as Administrator or be a member of the Administrators group. Guest direct network connections do not require administrator credentials.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you specify your connection as a host when you create it, the connection appears as **Incoming Connections** in the Network Connections folder.
- You can create multiple direct connections by copying them in the Network Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple ports, host computers, and so on.
- Direct connections can bypass authentication requirements. This is useful for devices such as handheld PCs. You must configure this setting in the host incoming connection. For more information, see Related Topics.
- If you create a direct connection by using a serial (RS-232C) cable, the port that you select in the New Connection Wizard is enabled for connections that use a null modem.
- If you are logged on to your computer as Administrator or a member of the Administrators group when you create a direct connection, you are presented with a list of connection devices to choose from that includes all of the parallel ports for the computer, infrared ports that are installed and enabled, and COM ports. If you need to use a COM port for a direct connection, ask your system administrator to configure one of the COM ports on your computer with a null modem by using Phone and Modem Options in Control Panel.
- Users do not need to use direct connections to allow access to shared resources, such as files and printers, over a local area network. In order to enable shared access to resources on the local computer, you must enable file and print sharing, share the resources, and then set up the appropriate permissions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using direct connections
Grant incoming connection access rights to your computer

Share a folder or drive

# Make a direct network connection by using a DirectParallel cable

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a direct network connection by using a DirectParallel cable**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Set up an advanced connection**, and then click **Next**.
4. Click **Connect directly to another computer**, click **Next**, and then do one of the following:
   - If you want your computer to act as the host, click **Host**, and click **Next**. Click **Direct Parallel (LPT1)** in the **Device for this connection**, and click **Next**.
   - If you want your computer to act as the guest, click **Guest**, and click **Next**. Type the name of the other computer in **Computer Name**, and click **Next**.
5. Follow the remaining instructions in the New Connection Wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- For information about contacting Parallel Technologies or their local dealers to order DirectParallel cables, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DirectParallel cable access
To contact Parallel Technologies

# Set up Incoming Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up incoming connections**

- Accept incoming network connections
- Grant incoming connection access rights to your computer
- Configure an incoming connection to use TCP/IP
- Configure callback for incoming connections

# Accept incoming network connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To accept incoming network connections**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Set up an advanced connection**, and then click **Next**.
4. Click **Accept incoming connections**, click **Next**, and then follow the instructions in the New Connection Wizard.

## Notes

- To create an incoming network connection, you must be a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you make another incoming network connection, and you use the New Connection Wizard again, the existing incoming network connection is reconfigured.
- For large numbers of incoming connections on a server that operates as part of a distributed network or as a domain controller, use Routing and Remote Access to create a remote access server.
- If your incoming connection and Fax Service have problems working together (for example, you cannot receive incoming connection calls on a device enabled for Fax receive), the modem may not support adaptive answer. Check your modem documentation to verify that you need to disable Fax receive for that device to accept incoming connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Incoming connections
Incoming connection clients
Configure an incoming connection to use TCP/IP
Grant incoming connection access rights to your computer

# Grant incoming connection access rights to your computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To grant incoming connection access rights to your computer**

1. Open Network Connections.
2. Right-click **Incoming Connections**, and then click **Properties**.
3. On the **General** tab, do one or more of the following:
   - In **Devices**, click the check boxes next to the devices through which incoming connections can connect.
   - If you have more than one device selected, and you want to enable Multilink dialing, select the **Enable multilink** check box.
   - If you want to enable others to make private connections to your computer by tunneling through the Internet or other network, select the **Allow others to make private connections to my computer by tunneling through the Internet or other network** check box.
   - If you want to show icons on the taskbar when incoming connections are connected, select the **Show icon in notification area when connected** check box.
4. On the **Users** tab, do one or more of the following:
   - If you want to allow a user to connect, select the check box next to the user's name.
   - If you want to revoke a user's right to connect, clear the check box next to a user's name.
   - To create a new user and grant them access to connect, click **New**.
   - To remove a user, click **Delete**.
   - To modify a user's name, password, and callback options, click **Properties**.
   - If you want directly connected devices to connect without providing a password, select the **Always allow directly connected devices such as palmtop computers to connect without providing a password** check box.
5. On the **Networking** tab, do one or more of the following:
   - If you want to enable a network component, in **Network components**, select the check box next to the component name.
   - If you want to disable a network component, in **Network components**, clear the check box next to the component name.
   - To add a network component, click **Install**.
   - To remove a network component, click **Uninstall**.

- To configure a network component, click **Properties**.

**Notes**

- To modify an incoming connection, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you enable or disable TCP/IP for incoming connections, the settings affect only incoming connections. Other protocols, such as AppleTalk, are enabled or disabled for all connection types.

Related Topics

Configure callback for incoming connections

Incoming connections

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure an incoming connection to use TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure an incoming connection to use TCP/IP**

1. Open Network Connections.
2. Right-click **Incoming Connections**, and then click **Properties**.
3. On the **Networking** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. If you want incoming users to be able to access the local area network on which the dial-up computer resides, select the **Allow callers to access my local area network** check box.
5. Do one of the following:

   - If you want to automatically assign TCP/IP addresses, select the **Assign TCP/IP addresses automatically using DHCP** check box.
   - If you want to specify TCP/IP addresses, click **Specify TCP/IP addresses**, and do the following:

     In **From**, type the starting IP address.

     In **To**, type the ending IP address.

     Based on the addresses you type in **From** and **To**, the number of addresses that are allocated is displayed in **Total**. The starting and ending IP addresses that you choose depend on the originating address space and the number of addresses needed. For more information on defining IP addresses for incoming connections, see Related Topics.

6. If you want incoming connections to be able to use a specific TCP/IP address, select the **Allow calling computer to specify its own IP address** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Most TCP/IP networks use subnets in order to effectively manage routed IP addresses. The closest matching subnet is calculated for the range that you specified in **From** and **To**. The range of addresses in the closest matching subnet may exceed the range that you specified. Unless the addresses specified in **From** and **To** are subnet boundaries, the range based on the calculated subnet will be larger than the range that you specified. To avoid this, specify a range that falls on subnet boundaries. For example, if you are using the 10.0.0.0 private network ID for your intranet, a range that falls on subnet boundaries is 10.0.1.168 to 10.0.1.175. Or, if you are using the 192.168.0.0 private network ID for your intranet, a range that falls on subnet boundaries is 192.168.1.0 to 192.168.1.255.
- If you select **Allow callers to access my local area network** or **Allow calling computer to specify its own IP address**, these changes take effect immediately.
- If you select **Assign TCP/IP addresses automatically using DHCP**, but there is no DHCP server available, random addresses from the range 169.254.0.1 to 169.254.255.254 are assigned.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Incoming connections and IP addresses
To configure a default gateway

# Configure callback for incoming connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure callback for incoming connections**

1. Open Network Connections.
2. Right-click **Incoming Connections**, and then click **Properties**.
3. On the **Users** tab, click the user for which you want to configure callback, and then click **Properties**.
4. On the **Callback** tab, do one of the following:
   - If you do not want to use callback for incoming connections, click **Do not allow callback**.
   - If you want to give incoming connections the option of setting a specific callback number, click **Allow the caller to set the callback number**.
   - If you want to consistently use the same callback number, click **Always use the following callback number**, and type the number.

**Note**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Grant incoming connection access rights to your computer
Callback

# Configure Network and Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure network and dial-up connections**

- Configure a connection to a remote network
- Set up Network Components
- Secure Network or Dial-up Connections

# Configure a connection to a remote network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a connection to a remote network**

1. Open Network Connections.
2. Right-click the dial-up connection that you want to configure, and then click **Properties**.
3. Do one or more of the following:
   - To configure dialing devices, phone numbers, host address, country/region codes, or dialing rules, click the **General** tab.
   - To configure dialing and redialing options, multilink configuration, or X.25 parameters, click the **Options** tab.
   - To configure identity authentication, data encryption, or terminal window and scripting options, click the **Security** tab.
   - To configure the remote access server and protocols used for this connection, click the **Networking** tab.
   - To enable or disable Internet Connection Firewall, Windows Firewall, Internet Connection Sharing, and on-demand dialing, click the **Advanced** tab. For more information, see Related Topics.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- Depending on the type of connection you are configuring, different options appear in the properties dialog box.
- For more information about a specific item on a tab, right-click the item, and then click **What's This?**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring a network or dial-up connection
Set up Network Components
Secure Network or Dial-up Connections
Configuring Authentication and Data Encryption

Enable ICS
Enable or disable ICF
Enable ICS on-demand dialing

# Set up Network Components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up network components**

- Configure TCP/IP settings
- Add a network component
- Remove a network component
- Enable a network component
- Disable a network component
- Modify the protocol bindings order
- Modify the network provider order
- Configure Client for Microsoft Networks
- Configure File and Printer Sharing for Microsoft Networks
- PPP Settings

# Configure TCP/IP settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure TCP/IP settings**

1. Open Network Connections.
2. Right-click the connection you want to configure, and then click **Properties**.
3. Do one of the following:
   - If the connection is a local area connection, on the **General** tab, in **This connection uses the following items**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
   - If the connection is a dial-up or VPN connection, on the **Networking** tab, in **This connection uses the following items**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
   - If the connection is an incoming connection, on the **Networking** tab, in **Network components**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Do one of the following:
   - If you want IP settings to be assigned automatically, click **Obtain an IP address automatically**, and then click **OK**.
   - If you want to specify an IP address or a DNS server address, do the following:

     Click **Use the following IP address**, and in **IP address**, type the IP address.

     Click **Use the following DNS server addresses**, and in **Preferred DNS server** and **Alternate DNS server**, type the addresses of the primary and secondary DNS servers.
5. To configure DNS, WINS, and IP settings, click **Advanced**.
6. In a local area connection, selecting the **Obtain an IP address automatically** option enables the **Alternate Configuration** tab. Use this to enter alternate IP settings if your computer is used on more than one network. To configure DNS, WINS, and IP settings, click **User configured** on the **Alternate Configuration** tab.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You should use automated IP settings (DHCP) whenever possible for all connections because they eliminate the need to configure settings such as IP address, DNS server address, and WINS server address.
- The **Alternate Configuration** settings specify a second set of IP settings that are used when a DHCP server is not available. This is very useful for laptop users who often switch between two different network environments such as DHCP and static IP network environments.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
TCP/IP
New ways to do Network Connections tasks

# Add a network component

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a network component**

1. Open Network Connections.
2. Right-click the connection to which you want to add a network component, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, click **Install**.
   - If this is a dial-up, VPN, or incoming connection, on the **Networking** tab, click **Install**.
4. In the **Select Network Component Type** dialog box, click either **Client**, **Service**, or **Protocol**, and then click **Add**.
5. Do one of the following:
   - If you do not have an installation disk for the component, click the appropriate client, service, or protocol, and then click **OK**.
   - If you have an installation disk for the component, click the appropriate client, service, or protocol, click **Have Disk**, insert the installation disk into the selected drive, and then click **OK**.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You should only install the network components that you need, for the following reasons:
  - Network performance is enhanced and network traffic is reduced when only the required protocols and clients are installed.
  - If a problem with a network or dial-up connection occurs, your computer attempts to establish connectivity by using every network protocol that is installed and enabled. By only installing and enabling the protocols that your system can use, the operating system does not attempt to connect with protocols it cannot use, and returns status information to you more efficiently.
  - Excessive services can hinder performance on your local computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Remove a network component
Disable a network component
Enable a network component

## Add optional networking components

# Remove a network component

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a network component**

1. Open Network Connections.
2. Right-click the connection on which you want to remove a network component, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, click the client, service, or protocol you want to remove, and then click **Uninstall**.
   - If this is a dial-up or VPN connection, on the **Networking** tab, in **This connection uses the following items**, click the client, service, or protocol you want to remove, and then click **Uninstall**.
   - If this is an incoming connection, on the **Networking** tab, in **Network components**, click the client, service, or protocol you want to remove, and then click **Uninstall**.
4. In the **Uninstall***component name* dialog box, click **Yes**.

**Note**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Add a network component
Enable a network component
Disable a network component

# Enable a network component

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable a network component**

1. Open Network Connections.
2. Right-click a connection, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, select the check box next to the client, service, or protocol you want to enable.
   - If this is a dial-up or VPN connection, on the **Networking** tab, in **This connection uses the following items**, select the check box next to the client, service, or protocol you want to enable.
   - If this is an incoming connection, on the **Networking** tab, in **Network components**, select the check box next to the client, service, or protocol you want to enable.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You should only enable the network components that you need, for the following reasons:
  - Network performance is enhanced and network traffic is reduced when only the required protocols and clients are enabled.
  - If your computer encounters a problem with a network or dial-up connection, it attempts to establish connectivity by using every network protocol that is installed and enabled. By only enabling the protocols that your system can use, the operating system does not attempt to connect with protocols it cannot use, and returns status information to you more efficiently.
  - Excessive services can hinder performance on your local computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Disable a network component

# Disable a network component

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable a network component**

1. Open Network Connections.
2. Right-click the connection on which you want to disable a network component, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, clear the check box next to the client, service, or protocol you want to disable.
   - If this is a dial-up, VPN, or incoming connection, on the **Networking** tab, in **Use these components with this connection**, clear the check box next to the client, service, or protocol you want to disable.

**Note**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Enable a network component

# Modify the protocol bindings order

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify the protocol bindings order**

1. Open Network Connections.
2. On the **Advanced** menu, click **Advanced Settings**.
3. In **Connections**, click the connection that you want to modify.
4. Click the protocol that you want to move up or down in the list, and then click the up or down arrow button.

## Notes

- To modify the protocol bindings order, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Configuring a network or dial-up connection
Set up Network Components

# Modify the network provider order

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify the network provider order**

1. Open Network Connections.
2. Click the connection you want to modify, and on the **Advanced** menu, click **Advanced Settings**.
3. On the **Provider Order** tab, in **Network Providers**, click the provider you want to move up or down in the list, and then click the **Up** or **Down** button.

## Notes

- To modify the network provider order, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Network providers are accessed in the order they appear in this list.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Set up Network Components
Configuring a network or dial-up connection

# Configure Client for Microsoft Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Client for Microsoft Networks**

1. Open Network Connections.
2. Right-click a connection, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, click **Client for Microsoft Networks**, and then click **Properties**.
   - If this is a dial-up, VPN, or incoming connection, on the **Networking** tab, in **Use these components with this connection**, click **Client for Microsoft Networks**, and then click **Properties**.
4. In **Name service provider**, click an RPC name service provider.
5. If you selected **DCE Cell Directory Service**, in **Network address** type the network address of your provider.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Windows Locator is the default name service provider.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Client for Microsoft Networks
Set up Network Components

# Configure File and Printer Sharing for Microsoft Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure File and Printer Sharing for Microsoft Networks**

1. Open Network Connections.
2. Right-click a connection, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, click **File and Printer Sharing for Microsoft Networks**, and then click **Properties**.
   - If the connection is a dial-up or VPN connection, on the **Networking** tab, in **This connection uses the following items**, click **File and Printer Sharing for Microsoft Networks**, and then click **Properties**.
   - If the connection is a incoming connection, on the **Networking** tab, in **Network components**, click **File and Printer Sharing for Microsoft Networks**, and then click **Properties**.
4. Do one of the following:
   - To optimize the server for a small number of clients, click **Minimize memory used**.
   - To optimize the server for a mixed usage of file and printer sharing in addition to other services, such as serving as an interactive workstation, click **Balance**.
   - To dedicate as many resources as possible to file and print server services, click **Maximize data throughput for file sharing**.
   - To optimize server memory for distributed applications that perform their own memory caching, such as Microsoft SQL Server, click **Maximize data throughput for network applications**.
5. To specify that LAN Manager 2.*x* clients on your network can browse for shared resources on this computer, select the **Make browser broadcasts to LAN Manager 2.x clients** check box.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You can only configure File and Printer Sharing for Microsoft Networks on a server.
- To share local folders, you must enable File and Printer Sharing for Microsoft Networks. For more information, see Related Topics.
- The File and Printer Sharing for Microsoft Networks component is the equivalent of the Server service in Windows NT 4.0.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
File and Printer Sharing for Microsoft Networks
Share a folder or drive

# PPP Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPP settings**

- Stop requesting LCP extensions in PPP
- Start requesting LCP extensions in PPP
- Turn off IP header compression in PPP
- Disable software compression

# Stop requesting LCP extensions in PPP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop requesting LCP extensions in PPP**

1. Disconnect from the remote server.
2. Open Network Connections.
3. Right-click the dial-up connection you want to configure, and then click **Properties**.
4. On the **Networking** tab, click **Settings**.
5. In the **PPP Settings** box, clear the **Enable LCP extensions** check box.

### Note

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting network and dial-up connections
Start requesting LCP extensions in PPP
PPP operation and protocols

# Start requesting LCP extensions in PPP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start requesting LCP extensions in PPP**

1. Disconnect from the remote server.
2. Open Network Connections.
3. Right-click the dial-up connection you want to configure, and then click **Properties**.
4. On the **Networking** tab, click **Settings**, and then select the **Enable LCP extensions** check box.

### Note

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting network and dial-up connections
Stop requesting LCP extensions in PPP
PPP operation and protocols

# Turn off IP header compression in PPP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To turn off IP header compression in PPP**

1. Disconnect from the remote server.
2. Open Network Connections.
3. Right-click the dial-up connection you want to configure, and then click **Properties**.
4. On the **Networking** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
5. Click **Advanced**, and then clear the **Use IP header compression** check box.

## Notes

- To turn off IP header compression, you must be logged on as an administrator or a member of the Administrators group. Some connections do not display the **IP header compression** check box.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting network and dial-up connections
ISP access methods
PPP Settings

# Disable software compression

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable software compression**

1. Open Network Connections.
2. Right-click the dial-up, VPN, or incoming connection you want to configure, and then click **Properties**.
3. On the **Networking** tab, click **Settings**, and clear the **Enable software compression** check box.

### Note

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting network and dial-up connections
Start requesting LCP extensions in PPP
Stop requesting LCP extensions in PPP
Using phone lines and modems

# Secure Network or Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Secure network or dial-up connections**

- Configure identity authentication and data encryption settings
- Enable smart card or other certificate authentication
- Terminal and Scripting Options

For information on security considerations for routing, remote access and VPN, see Security information for routing, Security information for remote access and Security issues for VPN.

# Configure identity authentication and data encryption settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure identity authentication and data encryption settings**

1. Open Network Connections.
2. Right-click the dial-up connection you want to configure, and then click **Properties**.
3. On the **Security** tab, do one of the following:

   - To select preconfigured combinations of identity authentication methods and data encryption requirements, click **Typical (recommended settings)**, and in **Validate my identity as follows**, click a method to use for validation.

     Depending on your **Validate my identity as follows** selection, you can select (enable) or clear (disable) the **Automatically use my Windows logon name and password (and domain if any)** or **Require data encryption (disconnect if none)** check boxes. The following table shows which selections are available.

     If **Require data encryption (disconnect if none)** is not enabled, encryption is optional. To prevent encryption, click **Advanced (custom settings)**, and then click **Settings**. In **Data encryption**, click **No encryption allowed (server will disconnect if it requires encryption)**.

     For dial-up connections:

     | Validate my identity as follows | Automatically use my Windows logon name and password (and domain if any) | Require data encryption (disconnect if none) |
     |---|---|---|
     | Allow unsecured password | Unavailable | Unavailable |
     | Require secured password | Available | Available |
     | Use smart card | Unavailable | Available |

     For VPN connections:

     | Validate my identity as follows | Automatically use my Windows logon name and password (and domain if any) | Require data encryption (disconnect if none) |
     |---|---|---|
     | Require secured password | Available | Available (enabled by default) |
     | Use smart card | Unavailable | Available (enabled by default) |

   - To individually enable, configure, and disable authentication methods and encryption requirements, click **Advanced (custom settings)**, and then click **Settings**.

**Important**

- Modifying **Advanced (custom settings)** requires a knowledge of security protocols.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- Among the advanced custom settings is the **Allow older MS-CHAP version for Windows 95 servers** check box. You must select this check box to configure a connection to a server running Windows 95.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dial-up connection authentication and data encryption
VPN connection authentication and data encryption

# Enable smart card or other certificate authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable smart card or other certificate authentication**

1. Open Network Connections.
2. Right-click the dial-up, VPN, or incoming connection on which you want to use smart card or other certificate authentication, and then click **Properties**.
3. If you are using typical settings for your smart card, on the **Security** tab, click **Typical (recommended settings)**, and in the **Validate my identity as follows** list, click **Use smart card**, and then click **OK**.
4. If you are individually enabling, configuring, and disabling authentication methods and encryption requirements, on the **Security** tab, click **Advanced (custom settings)**, and click **Settings**. In **Logon security**, click **Use Extensible Authentication Protocol (EAP)**, click **Smart card or other certificate (encryption enabled)**, click **Properties**, and then do the following:
   - If you want to use the certificate that resides on your smart card, click **Use my smart card**.
   - If you want to use the certificate that resides in the certificate store on your computer, click **Use a certificate on this computer**.
   - If you want to verify that the server certificate presented to your computer has not expired, has the correct signature, and has a trusted root certification authority, select the **Validate server certificate** check box.
   - If you only want to connect to servers within a particular domain, select the **Connect only if server name ends with** check box, and then type the name of the domain.
   - To specify that the root certification authority for your server certificate must be in a particular root certification authority, in **Trusted root certification authority**, click the appropriate certification authority.
   - To use a different user name when the user name in the smart card or certificate is not the same as the user name in the domain that you are logging on to, select the **Use a different user name for the connection** check box.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If, for example, you only want to connect to servers that reside in Microsoft.com, you type **Microsoft.com** in **Connect only if server name ends in**.
- If, for example, you are working for a consulting company where you need to log on to the domain of the company to which you are assigned, but your smart card contains a user name specific to your home company, select the **Use a different user name for the connection** check box.
- If you select the **Use a different user name for the connection** check box, your certificate is exported without private keys and submitted to your system administrator to be explicitly mapped to your domain user account. For more information, see Related Topics.
- If you select the **Connect only if server name ends with** check box, and do not type a domain name, at the time you connect you are prompted to use the domain name in the server certificate.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Smart card and other certificate authentication
Mobile users and certificates
Map a certificate to a user account

# Terminal and Scripting Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal and scripting options**

- Use the Terminal feature to log on to a remote computer
- Activate a logon script
- Log and view modem commands

# Use the Terminal feature to log on to a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use the Terminal feature to log on to a remote computer**

1. Open Network Connections.
2. Right-click the dial-up connection on which you want to use Terminal, and then click **Properties**.
3. On the **Security** tab, select the **Show terminal window** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- After you connect, the After Dial Terminal window appears and prompts you from the remote computer display. You then use this window to log on to the remote computer. After you have completed all interactions with the remote computer, click **Done**. At this point, authentication on the remote access server begins.
- If the logon sequence does not vary, you can write a script that automatically passes information to the remote computer during the logon sequence, which enables completely automatic connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using the Terminal feature for the remote logon process
Automating the Dial-up Logon Process by Using Switch.inf Scripts
Activate a logon script
Log and view modem commands

# Activate a logon script

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate a logon script**

1. Open Network Connections.
2. Right-click the dial-up connection on which you want to run a script, and then click **Properties**.
3. On the **Security** tab, select the **Run script** check box.
4. To use an existing script file, click **Browse**.

   Or, to use Switch.inf, click **Edit**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- In addition to Switch.inf, the Windows 95 scripting language is supported for computers running the Windows Server 2003 family that connect to computers running operating systems other than Windows. You may find the Windows 95 scripting language easier to use than Switch.inf. For more information, see Related Topics.
- If you click **Edit** and load Switch.inf into Notepad, immediately resave the file with a different name to preserve the default Switch.inf file.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Creating Switch.inf Scripts for Network and Dial-up Connections
Log and view modem commands
Troubleshooting logon scripts

# Log and view modem commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To log and view modem commands**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem whose commands you want to log.
3. Click **Properties**, then click the **Diagnostics** tab.
4. In **Logging**, select the **Append to Log** check box if you are working on a computer running Windows XP or **Record a Log** check box if you are working on a computer running a member of the Windows Server 2003 family.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Click **View Log** to display the log file. This file identifies the modem and displays information about how it is functioning.
- Commands sent to the modem are captured in the file:

  *systemroot*\System32\ModemLog_*Model*.txt.

  The path is usually C:\Windows\System32\. *Model* is the name of the modem as it appears in the list of installed modems on the **Modems** tab of **Phone and Modem Options**.
- In Windows XP, logging is always turned on, and the log is overwritten at the beginning of every session unless you select the **Append to Log** check box. In the Windows Server 2003 family, logging is turned off unless you select the **Record a Log** check box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modem diagnostics
Test a modem
Log and view modem commands
Troubleshooting modems

# Monitor Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor network connections**

- Monitor a connection
- View the status of a local area connection

# Monitor a connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To monitor a connection**

1. Open Network Connections.
2. Do one of the following:
   - To monitor activity for the current session, right-click the connection you want to monitor, and then click **Status**.
   - To automatically enable the status monitor each time the connection is active, right-click the connection, click **Properties**, and then select the **Show icon in taskbar when connected** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If your computer is set up to accept incoming connections, a connection with a user name assigned to it appears in the Network Connections folder as each user connects. You can view the progress of incoming connections by right-clicking a named connection, and then clicking **Status**.
- You can also monitor a connection by clicking **Status** on the **File** menu or by clicking the status icon in the taskbar. You can also right-click the taskbar icon and click **Status**.
- If you want to view the general status (**Disconnected**, **Connected**) of all of your connections, on the **View** menu, click **Details**.
- You also can disconnect an active connection by clicking **Disconnect** in the **Status** dialog box.
- By default, the status monitor is disabled for local area connections, but enabled for all other types of connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitoring network and dial-up connections
Troubleshooting network and dial-up connections
Log and view modem commands
Configure a connection to a remote network

# View the status of a local area connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the status of a local area connection**

1. Open Network Connections.
2. Do one of the following:
   - To monitor activity for the current session, right-click the local area connection, and then click **Status**.
   - To automatically enable the status monitor each time the connection is active, right-click the local area connection, click **Properties**, and then select the **Show icon in taskbar when connected** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You can view and modify the network components of a local area connection, such as network protocol settings, by clicking **Properties**.
- By default, the status monitor is disabled for local area connections, but enabled for all other types of connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using local area connections
Make a local area connection

# Set up and Manage Home and Small Office Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up and manage home and small office networking**

- Set up and Manage Internet Connection Firewall (ICF)
- Set up and Manage Network Bridge
- Set up and Manage Internet Connection Sharing (ICS)
- Manage service definitions for ICF or ICS

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Set up and Manage Internet Connection Firewall (ICF)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up and manage Internet Connection Firewall (ICF)**

- Enable or disable ICF
- Enable or disable Internet Control Message Protocol requests for ICF
- Enable or disable ICF security log options
- View the ICF security log
- Change the file size of the ICF security log
- Change the path and file name of the ICF security log
- Restore the default settings to the ICF security log

# Enable or disable ICF

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable ICF**

1. Open Network Connections.
2. Under **Dial-up** or **LAN or High-Speed Internet**, right-click the connection that you want to protect, and then click **Properties**.
3. On the **Advanced** tab, under **Internet Connection Firewall**, do one of the following:
   - To enable Internet Connection Firewall (ICF), select the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box.
   - To disable ICF, clear the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box.

**Caution**

- If you clear the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box, the firewall on this connection is disabled. When the firewall is disabled, your computer is exposed to possible external attacks through the unprotected connection.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You need to enable ICF only on connections that connect to the Internet. You should not enable ICF on virtual private network (VPN) connections or on network client computers because ICF interferes with network activities such as file and printer sharing.
- If ICF is enabled on two or more connections on a single computer, the ICF settings are global. If you select a setting or change a setting in **Logging Options** or **Services** on any one of the connections on which ICF is enabled, that setting is applied to all of the ICF firewalls on that computer.
- ICF cannot be enabled on the private connections of the Internet Connection Sharing host computer, on Network Bridge, or on the adapters that are included in Network Bridge.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- This topic applies only to product features available in the original release of Windows Server 2003.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Home and small office networking features
Internet Connection Firewall
Internet Connection Firewall security log
Service definitions for Internet Connection Firewall and Internet Connection Sharing

# Enable or disable Internet Control Message Protocol requests for ICF

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable Internet Control Message Protocol requests for ICF**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **ICMP** tab, do one of the following:
   - To enable Internet Control Message Protocol (ICMP) options, select the check box next to each type of request for information to which you would like your computer to respond.
   - To disable ICMP options, clear any or all of the ICMP check boxes.

### Caution

- All ICMP options are disabled by default. If you enable ICMP options on a computer running ICF, that computer may become visible to the Internet and vulnerable to attacks.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable ICF
Home and small office networking features
Internet Connection Firewall

# Enable or disable ICF security log options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable ICF security log options**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, under **Logging Options**, select from the following:

   To enable ICF security log options:

   - To enable the logging of all dropped packets that originate from either the home or small office network or the Internet, select the **Log dropped packets** check box.
   - To enable the logging of all successful connections that originate from either the home or small office network or the Internet, select the **Log successful connections** check box.

   Or, to disable ICF security log options:

   - To disable the logging of all dropped packets that originate from either the home or small office network or the Internet, clear the **Log dropped packets** check box.
   - To disable the logging of all successful connections that originate from either the home or small office network or the Internet, clear the **Log successful connections** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Security logging is not enabled by default when ICF is enabled.
- ICF functions whether security logging is enabled or disabled.
- If ICF is enabled on two or more connections on a single computer, the ICF settings are global. If you select a setting or change a setting in **Logging Options** or **Services** on any one of the connections on which ICF is enabled, that setting is applied to all of the ICF firewalls on that computer.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall security log
Internet Connection Firewall
Enable or disable ICF

# View the ICF security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the ICF security log**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, under **Log file options**, in **Name**, click **Browse**.
5. Scroll to the ICF security log, right-click the file, and then click **Open**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The default name for the ICF security log is pfirewall.log, and the default location is in the Windows directory.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall
Internet Connection Firewall security log
Enable or disable ICF
Enable or disable ICF security log options

# Change the file size of the ICF security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the file size of the ICF security log**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, under **Log file options**, in **Size limit**, use the arrow buttons to adjust the size limit.

### Caution

- If the maximum allowable size for pfirewall.log is exceeded, the information that pfirewall.log contains is saved as pfirewall.log.old, and new information is saved as pfirewall.log. If the maximum allowable size for pfirewall.log is exceeded again, the information that pfirewall.log contains is saved as pfirewall.log.old, and the information that had been in pfirewall.log.old is overwritten. New information is always saved in pfirewall.log.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If ICF is enabled on two or more connections on a single computer, the ICF settings are global. If you select a setting or change a setting in **Logging Options** or **Services** on any one of the connections on which ICF is enabled, that setting is applied to all of the ICF firewalls on that computer.
- The default size limit of the ICF security log is 4,096 kilobytes.
- You must enable ICF before you can adjust the size limit of the ICF security log.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall security log
Enable or disable ICF
Internet Connection Firewall

# Change the path and file name of the ICF security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the path and file name of the ICF security log**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, under **Log file options**, click **Browse**, and then do the following:
   - To change the path of the ICF security log, navigate to the location where you want the ICF security log to be saved. (If you do not want to change the path, do not navigate to a new location.)
   - To change the file name of the ICF security log, in **File name**, type a new name, and then click **Open**. (If you do not want to change the default file name, in **File name**, type **pfirewall**, and then click **Open**.)

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The default name for the ICF security log is pfirewall.log, and the default location is in the Windows directory.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall security log
Enable or disable ICF
Enable or disable ICF security log options

# Restore the default settings to the ICF security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore the default settings to the ICF security log**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, click **Restore Defaults**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The default size limit of the log is 4,096 kilobytes.
- The default name for the ICF security log is pfirewall.log, and the default location is in the Windows directory.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall security log
Home and small office networking features
Internet Connection Firewall
Enable or disable ICF

# Set up and Manage Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up and manage Network Bridge**

- Set up Network Bridge
- Add a connection to or remove a connection from Network Bridge
- Enable or disable Network Bridge
- Remove Network Bridge

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Set up Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up Network Bridge**

1. Open Network Connections.
2. Under **LAN or High-Speed Internet**, hold down the CTRL key, and then click each of the adapters that you want to include in the bridge.
3. Right-click one of the connections that you have selected, and then click **Bridge Connections**.

**Caution**

- If Internet Connection Firewall (ICF), Windows Firewall, or Internet Connection Sharing (ICS) are not enabled on your network, do not set up Network Bridge between the public Internet connection and the private network connection. Setting up Network Bridge between the public Internet connection and the private network connection creates an unprotected link between your network and the Internet, leaving your network vulnerable to external attacks. When ICF, Windows Firewall, or ICS are enabled, this risk is mitigated.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Connections that run ICF, Windows Firewall, or adapters that provide shared Internet connections cannot be a part of the Network Bridge and do not appear under **LAN or High-Speed Internet**. Similarly, the **Bridge Connections** menu command is available only for adapters that you can add as connections to Network Bridge.
- You must select at least two private network connections in order for **Bridge Connections** to create a bridge.
- You can create only one Network Bridge on a computer, but the bridge can accommodate numerous network connections.
- Network adapters that previously appeared under **LAN or High-Speed Internet**, appear under **Network Bridge** when you add them as connections to the bridge. Network adapters that provide Internet connectivity, such as DSL and cable modems, cannot be bridged and always appear under **LAN or High-Speed Internet**.
- Only Ethernet, IEEE-1394 adapters, or Ethernet-compatible adapters, such as wireless and Home Phoneline Network adapters, can be a part of Network Bridge.
- Network Bridge with a wireless or IEEE-1394 connection supports traffic using Internet Protocol version 4 (IPv4) only.
- You cannot create Network Bridge on a computer running Windows 2000 or earlier versions of Windows.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a connection to or remove a connection from Network Bridge
Enable or disable Network Bridge
Remove Network Bridge
Home and small office networking features
Network Bridge

# Add a connection to or remove a connection from Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a connection to or remove a connection from Network Bridge**

1. Open Network Connections.
2. Under **Network Bridge**, right-click **Network Bridge**, and then click **Properties**.
3. On the **General** tab, under **Adapters**, do one of the following.
   - To add a connection to Network Bridge, select the check box next to each adapter you want to add to the bridge, and then click **OK**.
   - To remove a connection from Network Bridge, clear the check box next to each adapter you want to remove from the bridge, and then click **OK**.

**Caution**

- If Internet Connection Firewall (ICF), Windows Firewall, or Internet Connection Sharing (ICS) are not enabled on your network, do not set up Network Bridge between the public Internet connection and the private network connection. Setting up Network Bridge between the public Internet connection and the private network connection creates an unprotected link between your network and the Internet, leaving your network vulnerable to external attacks. When ICF, Windows Firewall, or ICS are enabled, this risk is mitigated.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you remove network adapters from Network Bridge and less than two adapters remain, Network Bridge does not perform its intended function but continues to use system resources.
- You must set up Network Bridge before you can add connections to it. For more information, see Related Topics.
- Connections that run ICF, Windows Firewall, or adapters that provide shared Internet connections cannot be a part of Network Bridge and do not appear under **Adapters**. Similarly, the **Add to Bridge** menu command is available only for adapters that you can add as connections to Network Bridge.
- Network adapters that previously appeared under **LAN or High-Speed Internet**, appear under **Network Bridge** when you add them as connections to the bridge. Network adapters that provide Internet connectivity, such as DSL and cable modems, cannot be bridged and always appear under **LAN or High-Speed Internet**.
- Only Ethernet, IEEE-1394 adapters, or Ethernet-compatible adapters, such as wireless and Home Phoneline Network adapters, can be a part of Network Bridge.
- Network Bridge with a wireless or IEEE-1394 connection supports traffic using Internet Protocol version 4 (IPv4) only.

- You can also add a connection to or remove a connection from Network Bridge by opening Network Connections, right-clicking the adapter that you want to add, and then clicking **Add to Bridge**.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up Network Bridge
Enable or disable Network Bridge
Remove Network Bridge
Home and small office networking features
Network Bridge

# Enable or disable Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable Network Bridge**

1. Open Network Connections.
2. Under **Network Bridge**, right-click **Network Bridge**, and do one of the following:
   - To enable Network Bridge, click **Enable**.
   - To disable Network Bridge, click **Disable**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you remove or disable Network Bridge, computers on your network are unable to communicate with computers that are on different network segments. Some of the computers on your network may lose Internet connectivity.
- You must set up Network Bridge before you can enable or disable it. For more information, see Related Topics.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up Network Bridge
Add a connection to or remove a connection from Network Bridge
Remove Network Bridge
Home and small office networking features
Network Bridge

# Remove Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove Network Bridge**

1. Open Network Connections.
2. Under **Network Bridge**, right-click **Network Bridge**, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you remove or disable Network Bridge, computers on your network are unable to communicate with computers that are on different network segments. Some of the computers on your network may lose Internet connectivity.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up Network Bridge
Add a connection to or remove a connection from Network Bridge
Enable or disable Network Bridge
Home and small office networking features
Network Bridge

# Set up and Manage Internet Connection Sharing (ICS)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up and manage Internet Connection Sharing (ICS)**

- Enable ICS
- Enable ICS on-demand dialing
- Configure client computers for ICS

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Enable ICS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable ICS**

1. Open Network Connections.
2. Under **Dial-up** or **LAN or High-Speed Internet**, right-click the connection on which you want to enable (ICS), and then click **Properties**.
3. On the **Advanced** tab, under **Internet Connection Sharing**, select the **Allow other network users to connect through this computer's Internet connection** check box.
4. If present, under **Home networking connection**, in the **Select a private network connection** drop-down box, select the connection that connects directly to your network. For more information, see Notes.

**Important**

- You can enable ICS only when two or more network adapters are installed.
- If your computer has multiple private connections that connect to private LAN segments, you should bridge the private connections before you enable ICS. For information about how to set up Network Bridge, see Related Topics.
- When ICS is enabled, it automatically provides an IP address for the private connection of the ICS host.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- **Home networking connection** and the **Select a private network connection** drop-down box are present only when two or more adapters that can connect to the private network are installed.
- For information about the protocols, services, interfaces, and routes that are automatically configured by ICS when you connect to the Internet in a home or small office network, see Related Topics.
- If you are having difficulty connecting the ICS host to the Internet, you can use **Repair** to obtain refreshed IP address configurations from your Internet service provider (ISP).
- If a computer on your network is unable to communicate with other network computers, you can use **Repair** to force the connection to refresh its IP address configuration from the ICS host.
- Before using **Repair**, ensure that the client computers are plugged into the network and that the ICS host is running, configured, and connected to the network and the Internet.
- For information about how to configure client computers for ICS, see Related Topics.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connecting to the Internet in a home or small office network
Manage service definitions for ICF or ICS
Network Bridge
Configure client computers for ICS
Connecting to the Internet in a home or small office network

# Enable ICS on-demand dialing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable ICS on-demand dialing**

1. Open Network Connections.
2. Right-click the dial-up connection on which you want to enable ICS on-demand dialing, and then click **Properties**.
3. On the **Advanced** tab, under **Internet Connection Sharing**, do one of the following:
   - To enable ICS on-demand dialing, select the **Establish a dial-up connection whenever a computer on my network attempts to access the Internet** check box.
   - To disable ICS on-demand dialing, clear the **Establish a dial-up connection whenever a computer on my network attempts to access the Internet** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- In order for ICS on-demand dialing to work correctly, the shared Internet connection must be available for all ICS clients, and the dial-up credentials must be saved for all users. For more information, see Related Topics.
- **Establish a dial-up connection whenever a computer on my network attempts to access the Internet** is selected by default. Unless you clear the check box, clients can activate the shared Internet connection from their computers by starting their Internet browsers or mail programs.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting demand-dial routing
Dial-up credentials

# Configure client computers for ICS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure client computers for ICS**

1. Insert the Windows Server 2003, Standard Edition or Windows Server 2003, Enterprise Edition CD into the CD-ROM drive of your computer running Windows 98, Windows 98 Second Edition, or Windows Millennium Edition.

2. When the Welcome screen appears, click **Perform additional tasks**, and then click **Browse this CD**.

3. Double-click **Support**, and then double-click **Netsetup**.

4. Follow the instructions that appear on your screen.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- Enable Internet Connection Sharing (ICS) on your ICS host computer before you attempt to configure client computers for your home or small office network. For more information, see Related Topics.

- You can run Netsetup.exe only on computers running Windows 98, Windows 98 Second Edition, or Windows Millennium Edition.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connecting to the Internet in a home or small office network
Enable ICS
Home and Small Office Networking Overview

# Manage service definitions for ICF or ICS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage service definitions for ICF or ICS**

1. Open Network Connections.
2. Right-click the shared Internet connection or the Internet connection on which Internet Connection Firewall (ICF) or Internet Connection Sharing (ICS) is enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Services** tab, do one of the following:
   - If you want to enable a service, under **Services** select the check box next to service that you want to enable, and then enter the required information.
   - If you want to add a service definition, click **Add**, and then enter the required information.
   - If you want to edit a service definition, click the service that you want to edit, click **Edit**, and then change the appropriate information.
   - If you want to delete a service definition, click the service that you want to delete, and then click **Delete**.

**Caution**

- This feature is intended for advanced users. When you add or edit a service definition, you create a port mapping that can expose your network to the Internet. For more information about service definitions, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Before you can add, edit, or delete service definitions, you must enable ICF, ICS, or both.
- You can obtain the correct configuration values for the TCP or UDP port number from the service's documentation or Web site.
- No two service definitions can use the same port numbers simultaneously.
- For service definitions that are preconfigured, you can edit only the field that is labeled **Name or IP address of the computer hosting this service on your network**.
- For service definitions that you have added, you can edit only the fields that are labeled **Name or IP address of the computer hosting this service on your network**, **External Port number for this service**, and **Internal Port number for this service**.
- You can delete only the service definitions that you have added. You cannot delete service definitions that are preconfigured.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Service definitions for Internet Connection Firewall and Internet Connection Sharing
Internet Connection Firewall
Enable or disable ICF

# Network Connections Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Network Connections:

- Network Connections Overview
- Understanding Network Connections
- Using Network Connections
- Administering Network Connections

# Network Connections Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Connections overview**
This section covers:

- About Network Connections
- Hardware requirements for network and dial-up connections
- Network connection types
- Dial-up connections
- Network connection icons
- Protocols and Access Methods
- Home and Small Office Networking Overview

# About Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**About Network Connections**

Network Connections provides connectivity between your computer and the Internet, a network, or another computer. With Network Connections, you can gain access to network resources and functionality, whether you are physically located at the location of the network or in a remote location. Connections are created, configured, stored, and monitored from within the Network Connections folder. For information about creating connections, see Make Network and Dial-up Connections. For a list of network connection icons, see Network connection icons.

Each connection in the Network Connections folder contains a set of features that you can use to create a link between your computer and another computer or network. Outgoing connections contact a remote access server by using a configured access method (LAN, modem, ISDN line, cable modem, DSL line, X.25 line, and so on) to establish a connection with a remote network. Conversely, an incoming connection enables a computer running Microsoft® Windows® XP or a stand-alone computer running a member of the Windows Server 2003 family to be contacted by other computers. This means your computer can operate as a remote access server. Whether you are connected locally (LAN), remotely (dial-up, VPN, and so on), or both, you can configure any connection so that it can perform any needed network function. For example, you can print to network printers, access network drives and files, browse other networks, or access the Internet.

Because all services and communication methods are configured within the connection, you do not need to use external management tools to configure connection settings. For example, the settings for a dial-up connection include features to be used before, during, and after connecting. These include the modem you use to dial, the type of password encryption you want to use upon connecting, and the network protocols you use after you connect. Connection status, which includes the duration and speed of a connection, is viewed from the connection itself; you do not need to use an external status tool. For information about configuring a connection, see Configure a connection to a remote network.

Logon and domain security, support for security hosts, data encryption, authentication, and callback provide secure network access for network and dial-up connections. For more information about security features, see Security Features of Network Connections.

# Hardware requirements for network and dial-up connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware requirements for network and dial-up connections**
Depending on your configuration, you may need some or all of the following hardware:

- One or more network adapters with a Network Driver Interface Specification (NDIS) driver for LAN connectivity
- One or more compatible modems and an available COM port
- ISDN adapter (if you are using an ISDN line)
- DSL adapter
- X.25 adapter or PAD (if you are using X.25)
- Analog telephone line, ISDN line, X.25 line, or DSL line
- Smart card reader
- Wireless adapter

To verify the compatibility of all hardware, see the compatibility information at Support resources.

For information about configuring a connection, see Configure a connection to a remote network. For a list of network connection icons, see Network connection icons.

# Network connection types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network connection types**

You can use Network Connections to create several types of connections between computers. The following table shows each type of connection you can create, the communication methods you can use to create each type of connection, and an example of each connection type. For more information, click the appropriate connection type in the table.

| Connection type | Communication method | Example |
|---|---|---|
| Dial-up connections | Modem, ISDN, X.25 | Connect to a private network or the Internet by using remote access |
| Virtual Private Network (VPN) Connections | VPN connections to organization networks over the Internet using PPTP or L2TP | Connect securely to a private network over the Internet |
| Connection Manager Administration Kit | See dial-up or VPN connections | Connect to a private network using a collection of dial-up or VPN settings provided by your network administrator in a self-installing profile |
| Using local area connections | Ethernet, token ring, cable modem, DSL, FDDI, IP over ATM, IrDA, wireless, WAN technologies (T1, Frame Relay), PPPoE | Connect directly to a local area network, to your cable modem, or to your DSL modem through an Ethernet adapter or a similar device |
| Using direct connections | Serial cabling, infrared link, DirectParallel cable | Connect a handheld computer running Microsoft® Windows® CE to a desktop computer so you can synchronize information |
| Incoming connections | See dial-up, VPN, or direct connections | Accept dial-up, VPN, or direct connections from other computers |

For a list of icons that appear in the Network Connections folder, see Network connection icons.

# Dial-up connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up connections**

A dial-up connection connects you to a network or the Internet by using a device that utilizes the telephone network. This device can be a modem that uses a standard phone line, an ISDN adapter with a high-speed ISDN line, DSL, broadband, a cable modem, or an X.25 network.

If you are a typical user, you may have one or two dial-up connections, perhaps to the Internet and to your corporate network.

You can create multiple dial-up connections by copying them in the Network Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple modems, dialing profiles, and different destinations.

For more information about creating dial-up connections, see Set up Dial-up Connections.

# Network connection icons

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network connection icons**

The following table displays the different icons that appear in the Network Connections folder.

| Icon | Description | Location |
|---|---|---|
| | The local area connection is active | Network Connections folder |
| | The local area connection is disabled | Network Connections folder |
| | The default remote access connection | Network Connections folder |
| | An incoming connection | Network Connections folder |
| | The media is disconnected | Network Connections folder |
| | The media is disconnected | Taskbar |
| | The driver is disabled | Network Connections folder |
| | The adapter in this connection has an invalid IP address | Network Connections folder |
| | The adapter in this connection has an invalid IP address | Taskbar |
| | A connection on which Internet Connection Sharing (ICS) has been enabled. | Network Connections folder |
| | A connection on which Internet Connection Firewall (ICF) has been enabled. | Network Connections folder |
| | A connection on which Network Bridge has been enabled. | Network Connections folder |

**Notes**

- This topic applies only to product features available in the original release of Windows Server 2003.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Protocols and Access Methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Protocols and access methods**
This section covers:

- Network communications
- TCP/IP
- IPX
- AppleTalk
- ISDN access
- Serial (RS-232) cable access
- DirectParallel cable access
- SLIP
- Point-to-Point Protocol (PPP)

**Note**

- The IPX/SPX protocol is not available on the Itanium-based versions of the Windows operating systems.

# Network communications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network communications**
Network protocols, access methods, and server protocols help create a communications link between your computer and a network.

Whether information is transported from your computer to a server by using a direct serial cable, or across a secure VPN from your Internet service provider (ISP) to the corporate network, different combinations of access methods and protocols are used to send information. These access methods and protocols are listed as follows:

**Network protocols**

- TCP/IP
- IPX
- AppleTalk

**Access methods**

- Using phone lines and modems
- ISDN access
- Serial (RS-232) cable access
- DirectParallel cable access

**Server protocols**

- Point-to-Point Protocol (PPP)
- SLIP
- Layer Two Tunneling Protocol (L2TP)
- Point-to-Point Tunneling Protocol (PPTP)

The IPX/SPX protocol is not available on the Itanium-based versions of the Windows operating systems.

# TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP**
Transmission Control Protocol/Internet Protocol (TCP/IP) is the most popular protocol, and the basis for the Internet. Its routing capabilities provide maximum flexibility in an enterprise-wide network.

On a TCP/IP network, you must provide IP addresses to clients. Clients may also require a naming service or a method for name resolution. This section explains IP addressing and name resolution for Network Connections on TCP/IP networks. It also describes the FTP and Telnet tools that are provided by the Microsoft® Windows Server 2003 family.

**Assigning IP addresses to remote access connections**
Each remote computer that connects to a remote access server on a Windows TCP/IP network is provided an IP address by the server. This IP address is either automatically provided through the server by DHCP on the target network or selected from the static range that is assigned to the remote access server. For more information see Incoming connections and IP addresses.

If you use a specific IP address, the remote access server must be configured to permit users to request a specific address.

**Name resolution for network and dial-up connections**
In addition to requiring an IP address, network and dial-up connections for a NetBT gateway on a TCP/IP network may require a mechanism to resolve computer names to IP addresses. You can use any one of the following mechanisms for name resolution: DNS, WINS, broadcast name resolution, and the Hosts and Lmhosts files.

Incoming connections clients are assigned the same WINS and DNS servers that are assigned to the incoming connections computer. The TCP/IP settings of a network or dial-up connection can override these automatic assignments.

Incoming connections that do not have a DNS or WINS server rely completely on broadcast name resolution, which occurs by using broadcasts to locate network resources that match a computer name.

**Internet tools**
The Windows Server 2003 family provides File Transfer Protocol (FTP) and Telnet. FTP is a character-based utility that permits you to connect to FTP servers and transfer files. Telnet is a command-line application that lets you log on to remote computers and issue commands as if you were at the keyboard of the remote computer. Multiple variations of FTP, Telnet, and other services based on earlier Internet standards are also available on the Internet or commercially.

For information about configuring a connection, see Configure a connection to a remote network. For information about configuring TCP/IP, see Configure TCP/IP settings.

### Note

- FTP and Telnet tools run on top of TCP/IP networks.

# IPX

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPX**
Internetwork Packet Exchange (IPX) is the protocol that is used on many Novell NetWare networks. This section explains how to integrate Network Connections into an IPX network.

**Support for NetWare and IPX**
Network Connections supports NetWare and IPX networks by including:

- NWLink IPX/SPX/NetBIOS Compatible Transport Protocol, also known as NWLink, which is the Microsoft® Windows Server 2003 family implementation of the IPX protocol.
- Client Service for NetWare client component, which provides access to NetWare shared file and printer resources.
- Point-to-Point Protocol (PPP) IPX Configuration Protocol (IPXCP), which supports IPX over PPP for dial-up, virtual private network (VPN), and incoming connections. PPP connections use IPXCP, as defined in RFC 1552, "The PPP Internetwork Packet Exchange Control Protocol (IPXCP)," to configure PPP connections for IPX.

For information about configuring a connection, see Configure a connection to a remote network. For information about configuring NWLink IPX/SPX/NetBIOS Compatible Transport Protocol, see Configure NWLink.

**Notes**

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.
- IPX routing is not supported in the Windows Server 2003 family.
- IPX cannot be configured for incoming or dial-up connections in the Windows Server 2003 family.

# AppleTalk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AppleTalk**

Networking on an Apple Macintosh computer is implemented through the AppleTalk protocol. Applications and processes can communicate across a single AppleTalk network or an AppleTalk internet, which is a number of interconnected AppleTalk networks. By using AppleTalk, applications and processes can transfer and exchange data and share resources, such as printers and file servers. AppleTalk remote access is supported by the AppleTalk Control Protocol (ATCP).

With ATCP, Macintosh clients can run the AppleTalk network protocol over the Point-to-Point Protocol (PPP). With ATCP support, a remote user can access a Web server over TCP/IP, print a document to an AppleTalk printer, and connect to a Macintosh file server (over TCP/IP or AppleTalk), all with the same dial-up connection over PPP.

For more information, see AppleTalk over PPP.

For information about configuring a connection, see Configure a connection to a remote network.

# ISDN access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ISDN access**

To enhance your network connection speed, you can use an Integrated Services Digital Network (ISDN) line. Whereas standard phone lines typically transmit at up to 56 kilobits per second (Kbps), ISDN lines can transmit at speeds of 64 or 128 Kbps.

An ISDN line must be installed by the phone company at both the server and at the remote site. ISDN also requires that you install an ISDN adapter on the server and your computer. Costs for ISDN equipment and lines may be higher than standard modems and phone lines. However, the speed of communication reduces the duration of the connection, possibly saving toll charges.

An ISDN line comes with two B channels that transmit data at 64 Kbps, and one D channel for signaling that transmits data at 16 Kbps. You can configure each B channel to operate as a port. With some ISDN drivers, you can aggregate the channels. This means you can statically assign a higher bandwidth by configuring both B channels to act as a single port. With this configuration, line speed increases to 128 Kbps. For information about configuring a connection to use ISDN, see Make a dial-up connection by using an ISDN line or Configure ISDN settings.

The Multilink feature performs channel aggregation for ISDN. Multilink combines multiple physical links into a logical bundle. This aggregate link increases the bandwidth of a connection. In addition, you can allocate multiple links dynamically, which means ISDN lines are used only as they are required. This eliminates excess bandwidth, representing a significant efficiency advantage to users. You can configure the conditions under which extra lines are dialed, and underused lines are disconnected, through Network Connections settings. For more information, see Configuring multiple device dialing.

# Serial (RS-232) cable access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Serial (RS-232) cable access**

Suppose you have two networks in the same location that are not physically connected. To use resources on both networks from one computer, you can use RS-232 cables and COM ports to create null modems. You connect an RS-232 cable from a COM port to a COM port on each remote access server. To the servers, the COM ports then appear to be modems. You can use server remote access, a part of the integrated Routing and Remote Access service, to create network access.

For information about creating a direct connection by using serial access, see Make a direct network connection.

# DirectParallel cable access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DirectParallel cable access**

The DirectParallel feature supports computer-to-computer connections by using standard and enhanced (ECP) parallel ports, which use Parallel Technologies' Basic or Fast parallel cables. For information about configuring a direct network connection to use a DirectParallel cable, see Make a direct network connection by using a DirectParallel cable. Direct cable connection cables are available from several sources. You can obtain DirectParallel cables from Parallel Technologies or from their local dealers. For more information, see To contact Parallel Technologies.

# SLIP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SLIP**
Serial Line Internet Protocol (SLIP) is an older remote access standard typically used by UNIX remote access servers. Network Connections supports SLIP for dial-up connections.

When you connect to a SLIP server, a **Windows Terminal** dialog box appears for an interactive logon session with the UNIX SLIP server. The UNIX logon overrides and prevents the remote access logon session from appearing. After a connection is established, remote network access becomes transparent to the user.

The RFCs supported in remote access are:

- RFC 1055, "A Nonstandard for Transmission of IP Datagrams Over Serial Lines: SLIP"
- RFC 1144, "Compressing TCP/IP Headers for Low-Speed Serial Links"

For information about configuring a network or dial-up connection, see Configure a connection to a remote network.

**Notes**

- Microsoft® Windows® XP does not support SLIP for incoming connections.
- You must use the TCP/IP protocol and a serial COM port to connect to a SLIP server.

# Point-to-Point Protocol (PPP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Point-to-Point Protocol (PPP)**
The Point-to-Point Protocol (PPP) is a set of standard protocols that allow remote access software from different vendors to interoperate. A PPP-enabled connection can dial into remote networks through any industry-standard PPP server. PPP also permits a remote access server to receive calls from, and provide network access to, other vendors' remote access software that complies with the PPP standards.

The PPP standards also permit advanced features that are not available with older standards such as SLIP. PPP supports several authentication methods, as well as data compression and encryption. With most PPP implementations, you can automate the entire logon sequence.

PPP also supports multiple LAN protocols. You can use TCP/IP or IPX as your network protocol.

The IPX/SPX protocol is not available on the Itanium-based versions of the Windows operating systems.

PPP is the basis for the Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) protocols, which are used in secure virtual private network (VPN) connections.

PPP is the primary standard for most remote access computing.

This section covers:

- PPP operation and protocols
- Point-to-Point Protocol over Ethernet (PPPoE) Protocols

# PPP operation and protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPP operation and protocols**

Connections made using Point-to-Point Protocol (PPP) must adhere to standards established in PPP RFCs. For information on PPP RFCs, see "Introduction to TCP/IP" at the Microsoft Windows Resource Kits Web site. This topic gives an overview of PPP operations and the protocols used in a PPP connection. For information about how to configure PPP settings, see PPP Settings.

**PPP connection process**

After a physical or logical connection is made to a PPP-based remote access server, the following negotiations establish a PPP connection:

1. Negotiate the use of the link

   PPP uses Link Control Protocol (LCP) to negotiate link parameters such as the maximum PPP frame size, the use of Multilink, and the use of a specific PPP authentication protocol.

2. Authenticate the remote access computer

   The remote access client and the remote access server exchange messages according to the negotiated authentication protocol. If EAP is being used, the client and the server negotiate a specific EAP method, known as an EAP type, and then exchange the messages of that EAP type.

3. Use callback

   If callback for the dial-up connection is configured, the physical connection is terminated, and the remote access server calls the remote access client back.

4. Negotiate the use of network protocols

   This process involves using a series of network control protocols (NCPs) to configure the network protocols the remote access client uses. For example, the NCP for TCP/IP is Internet Protocol Control Protocol (IPCP). If the remote access client is configured to use TCP/IP, IPCP is used to assign an IP address to the client and to configure the client with the IP addresses of DNS and WINS servers.

The resulting connection remains active until the line is disconnected for any of the following reasons:

- The user or the administrator explicitly hangs up the line.
- The line is disconnected due to idle time-out.
- An unrecoverable link error occurs.

**Link Control Protocols**

Link Control Protocol (LCP) establishes and configures PPP framing. PPP framing defines how data is encapsulated before it is transmitted on the wide area network. The PPP standard framing format ensures that remote access software from any vendor can communicate and recognize data packets from any remote access software that adheres to the PPP standards.

PPP and the Microsoft® Windows Server 2003 family use variants of High-level Data Link Control (HDLC) framing for serial or ISDN connections.

**Authentication protocols**

Authentication protocols are negotiated immediately after link quality is determined and before the network layer is negotiated. For more information about available authentication protocols, see Authentication Protocols and Methods.

**Network control protocols**

Network control protocols establish and configure different network protocol settings for TCP/IP, IPX, and AppleTalk. The following table describes network control protocols used with a PPP connection.

The IPX/SPX protocol is not available on the Itanium-based versions of the Windows operating systems.

| Network control protocol | Description |
|---|---|
| Internet Protocol Control Protocol (IPCP) | IPCP is used to configure TCP/IP on the remote access client. Configuration parameters include an IP address and the IP addresses of DNS and WINS servers. |
| Internetwork Packet Exchange Control Protocol (IPXCP) | IPXCP is used to configure IPX on the remote access client. IPXCP is widely implemented by PPP vendors. Configuration parameters include the IPX network and node numbers. |
| AppleTalk Control Protocol (ATCP) | ATCP is used to configure AppleTalk on the remote access client. Configuration parameters include an AppleTalk address. |

**Note**

- The **PPP settings** dialog box includes the option **Negotiate multi-link for single-link connections**. This option enables separation of high and low-priority channels over a single link connection. If your remote access server supports this feature, you might notice improved audio quality. However, because this feature is incompatible with many remote access servers, you should leave it disabled unless instructed otherwise. For more information on PPP settings, see PPP Settings.

# Point-to-Point Protocol over Ethernet (PPPoE) Protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Point-to-Point Protocol over Ethernet (PPPoE) Protocols**
Using Point-to-Point Protocol over Ethernet (PPPoE) and a broadband modem, LAN clients can gain individual authenticated access to high-speed data networks. By combining Ethernet and Point-to-Point Protocol (PPP), PPPoE provides an efficient way to create a separate connection to a remote server for each user. Access, billing, and choice of service are managed on a per-user basis, rather than a per-site basis. Clients can create both dial-up and dial-on-demand connections using PPPoE.

**PPPoE RFCs**
Requests for Comments (RFCs) are an evolving series of technical reports, proposals for protocols, and protocol standards used by the Internet community. PPPoE is defined in RFCs published by the Internet Engineering Task Force (IETF) and other working groups. The following is an example of a PPPoE RFC:

**RFC number** - 2516A

**Title** - Method for Transmitting PPP Over Ethernet (PPPoE)

**Obtaining RFCs**
You can obtain RFCs from the RFC Editor Web site. This Web site is currently maintained by members of the Information Sciences Institute (ISI) who publish a classified listing of all RFCs. RFCs are classified as one of the following: approved Internet standards, proposed Internet standards (circulated in draft form for review), Internet best practices, or For Your Information (FYI) documents.

### Notes

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.
- PPPoE is supported on clients running Microsoft® Windows® XP. All members of the Windows Server 2003 family can connect to and route PPPoE traffic. However, they cannot accept an incoming PPPoE connection.

For more information, see Make a broadband connection using Point-to-Point Protocol over Ethernet (PPPoE).

# Home and Small Office Networking Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Home and small office networking overview**
With home and small office networking, even if you do not have the resources, time, or experience of a network administrator, you can build a small network in your home or small office. Home and small office networking is a group of features that limits hardware requirements and automates configurations for routing and private IP addressing. This reduces the work that is typically required to connect and configure a network. It also reduces the associated costs.

Home and small office networking provides many of the benefits of traditional networks and offers benefits that are uniquely suited for a home or small office network. With home and small office networking you can:

- Share one Internet connection with all of the computers on your network.
- Provide local private network services, such as name resolution and addressing, automatically.
- Share printers or other hardware devices, such as cameras and scanners, from any computer on your network.
- Share folders and files with other computers on your network.
- Use one computer to secure your entire network and protect your Internet connection.
- Easily establish network connections across two or more types of network media, such as IEEE 1394, Ethernet, wireless, and Home Phoneline Network adapters.

This section covers:

- Home and small office networking features
- Home and small office networking compared to Routing and Remote Access

# Home and small office networking features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Home and small office networking features**
Home and small office networking uses the following features to simplify the implementation and administration of small networks.

**Internet Connection Firewall (ICF) and Windows Firewall**
Home and small office networking offers computers that are directly connected to the Internet the added security of firewall protection. The firewall checks all communications that cross the connection between your network and the Internet and is selective about which responses from the Internet it allows. ICF and Windows Firewall protect only the computer on which they are enabled. However, if ICF or Windows Firewall are enabled on the Internet Connection Sharing (ICS) host computer, ICS clients that use the shared Internet connection for Internet connectivity are protected because they cannot be seen from outside your network. For this reason, you should always enable ICF or Windows Firewall on the ICS host computer. In addition, if there are clients on your network with direct Internet connections, or if you have a stand-alone computer that is connected to the Internet, then it is recommended that you enable ICF or Windows Firewall on those Internet connections as well.

For more information about ICF, see Internet Connection Firewall. For more information about Windows Firewall, see Help: Windows Firewall.

**Network Bridge**
Network Bridge removes the need for routing and bridging hardware in a home or small office network that consists of multiple LAN segments. With Network Bridge, multiple LAN segments become a single IP subnet, even if the LAN segments are of mixed network media types. Network Bridge automates the configuration and management of the address allocation, routing, and name resolution that is typically required in a network that consists of multiple LAN segments.

For more information see Network Bridge.

**Internet Connection Sharing (ICS)**
ICS allows you to provide Internet access for a home or small office network by using one common connection as the Internet gateway. The ICS host is the only computer that is directly attached to the Internet. Multiple ICS clients simultaneously use the common Internet connection and benefit from Internet services as if the clients were directly connected to the Internet service provider (ISP) (for more information, see Dial-up credentials). Security is enhanced when ICS is enabled because only the ICS host computer is visible to the Internet. The addresses of ICS clients are hidden from the Internet rendering ICS clients invisible to the Internet. In addition, ICS simplifies the configuration of small networks by providing local private network services, such as name resolution and addressing.

For more information about ICS, see Connecting to the Internet in a home or small office network.

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

# Home and small office networking compared to Routing and Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Home and small office networking compared to Routing and Remote Access**
Home and small office networking and the Routing and Remote Access service are both included in Microsoft® Windows® XP and the Windows Server 2003 family to provide the following networking capabilities:

- Network address translation (NAT)
- Packet filtering
- Dial-up access
- Virtual private network (VPN) access
- Automated IP addressing for network clients
- Name resolution
- Internet connections

Because home and small office networking features and Routing and Remote Access share common drivers, they cannot coexist, and you need to decide which to implement in order to meet the needs of your networking environment.

**Home and small office networking**
You can use home and small office networking features--Internet Connection Firewall (ICF), Windows Firewall, Network Bridge, and Internet Connection Sharing (ICS)--to simplify the configuration of your home or small office network. These features are intended for networks that consist of two to 10 computers when you want to:

- Protect computers that are connected to the Internet using an easily configured firewall. For more information, see Internet Connection Firewall.
- Connect LAN segments simply, by using the **Bridge Connections** menu command. For more information, see Network Bridge.
- Provide simultaneous Internet access for up to 10 computers using a single dial-up or high-speed Internet connection. For more information, see Connecting to the Internet in a home or small office network.
- Support up to 10 simultaneous dial-up or VPN connections to provide remote computers with access to network resources. For more information, see Incoming connections.

**Routing and Remote Access**
You can use Routing and Remote Access in combination with Internet Security and Acceleration (ISA) server to meet the needs of your small business with high security needs, your medium-sized private business, or your enterprise network that spans multiple subnets and supports up to 1,000 computers running Windows Server 2003, Standard Edition or up to 5,000 computers running Windows Server 2003, Enterprise Edition. Internet connectivity for branch offices is routed through the corporate routing and firewall infrastructure. Use Routing and Remote Access when you want to:

- Provide local and branch office computers with high-security Internet access.
- Connect branch offices with corporate intranets, and share resources as if all computers are connected to the same LAN.

- Protect network interfaces with static packet filters or dynamic packet filters.
- Support up to 1,000 simultaneous dial-up or VPN connections to provide remote computers with access to corporate network resources.

**Comparisons**

The following table summarizes the differences between the way that home and small office networking features (and incoming connections) and Routing and Remote Access (with ISA server) implement basic networking services.

| Service | Description | Home and small office networking features (and incoming connections) | Routing and Remote Access (with ISA server) |
|---|---|---|---|
| **NAT** | Hides internally managed IP addresses from external networks by translating private internal addresses to public external addresses. This reduces IP address registration costs by letting you use unregistered IP addresses internally, with translation to a small number of registered IP addresses externally. It also hides the internal network structure, reducing the risk of attacks against internal systems. | ICS | Routing and Remote Access NAT |
| **Dynamic packet filtering** | Provides protection from unsolicited traffic for the private network. Permits only traffic that is sent in response to an internal request. | ICF, Windows Firewall | Basic Firewall |
| **VPN and dial-up access** | Allows clients on a remote computer to connect to a private network and to access network resources as if the computer was physically attached to the network. | Incoming connections | Routing and Remote Access |
| **Address assignment** | Automates the assignment of client IP addresses on the private network in order to configure clients to allow client-access to network resources. | ICS DHCP allocator | Routing and Remote Access NAT, with a DHCP allocator or a DHCP server |
| **DNS name resolution** | Converts the names of computers and other network devices, such as printers, to IP addresses. | ICS DNS Proxy | NAT, with a DNS proxy or a DNS server |
| **Internet connections** | Provides a high-speed or dial-up connection for the computer that is connected to the Internet to use to publish its services to the private network. | Configured through Network Connections | Configured through Routing and Remote Access |

The following table summarizes which components, features, and connection types are best suited to provide networking services in a given networking scenario.

| Service | Scenario: Small office, non-domain network | Scenario: Small office, domain network | Scenario: Medium office network with NAT traversal | Scenario: Medium office network without NAT traversal | Scenario: Enterprise network | Scenario: Branch office network |
|---|---|---|---|---|---|---|
| **NAT** | ICS | ICS | Routing and Remote Access | Routing and Remote Access | Routing and Remote Access | Routing and Remote Access with ISA (on the corporate network) |

| | | | | | | |
|---|---|---|---|---|---|---|
| **Packet filtering** | ICF, Windows Firewall | ICF, Windows Firewall | Basic Firewall or ISA | Basic Firewall or ISA | ISA | ISA (on the corporate network) |
| **VPN/remote access** | Incoming connections | Incoming connections | Routing and Remote Access | Routing and Remote Access | Routing and Remote Access | Routing and Remote Access with ISA (on the corporate network) |
| **Address assignment** | ICS DHCP allocator | DHCP server | DHCP server | DHCP server | DHCP server | DHCP server |
| **Internet connection** | Dial-up, ISDN, broadband, DSL, or LAN | Dial-up, ISDN, broadband, DSL, or LAN (with ISA) | Any combination of T1 or T3, dial-up, ISDN, broadband, DSL, and LAN | Any combination of dial on demand (DoD) routing, T1 or T3, dial-up, ISDN, broadband, DSL, or LAN | Any combination of dial on demand (DoD) routing, T1 or T3, dial-up, ISDN, broadband, DSL, or LAN | Any combination of dial on demand (DoD) routing, T1 or T3, dial-up, ISDN, broadband, DSL, or LAN |
| **NAT traversal** | Not available | Not available | ISA | Does not apply | ISA | ISA (on the corporate network) |
| **DNS** | DNS proxy | DNS server | DNS server | DNS server | DNS server | DNS server |

**Notes**

- In the case of ICS with a DNS or DHCP server, ICS discovers the DHCP service or DNS service and, if present, disables the DNS proxy, DHCP allocator, or both.
- Because ICS, ICF, and Windows Firewall are not compatible with Routing and Remote Access, you cannot enable ICS or ICF if Routing and Remote Access is configured. Likewise, if ICS or ICF are enabled, they must first be disabled in order to configure Routing and Remote Access.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Understanding Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Network Connections**
This section covers:

- Security for Network Connections
- Tunneling Protocols
- Home and Small Office Networking

# Security for Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security for network connections**

This section covers:

- Security Features of Network Connections
- Intermediary Devices

# Security Features of Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security features of Network Connections**

You can configure your dial-up, virtual private network (VPN), broadband, and direct connections to enforce various levels of password authentication and data encryption. Authentication methods range from unencrypted to custom, using Extensible Authentication Protocol (EAP). EAP provides an infrastructure to support a wide range of authentication methods, including smart cards, certificates, one-time passwords, and public keys. You can also specify the type of data encryption, depending on the level of password authentication (MS-CHAP, MS-CHAP v2, or EAP-TLS) that you select. Finally, you can configure callback options to increase dial-up security.

User accounts and domains provide security with encrypted authentication and encrypted data. For more information about security, see Security.

This section covers:

- Callback
- How security works at connection
- Group Policy for Network Connections
- Group Policy settings that prohibit home and small office networking on your domain
- Dial-up credentials
- Data Encryption and Encryption Methods
- Network Authentication

# Callback

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Callback**
The callback feature instructs the remote access server to disconnect after authentication and then to call you back. Callback provides cost advantages to you and security advantages to your network. By immediately hanging up and then dialing you back, callback reduces phone charges. Required callback enhances network security by ensuring that only users from specific locations can access the server. By dropping the call, and then calling back a moment later to the preassigned callback number, most impersonators can be thwarted.

If callback has been required by your administrator, the following steps occur when your call reaches the remote access server:

1. The server determines whether your user name and password are correct.

2. If they are correct, the server disconnects and then returns the call.

If your administrator has configured callback for your connection through the **Remote Access Preferences** menu option, the following steps occur when your call reaches the remote access server:

1. The server determines whether your user name and password are correct.

2. If they are correct, the **Callback** dialog box appears on your computer.

3. You then type the current callback number in the dialog box and wait for the server to disconnect and return the call.

   Optionally, you can press ESC to cancel the callback process and remain connected.

For more information about configuring your callback options, see Configure client callback options.

**Notes**

- Your callback options are determined by your system administrator. The administrator may deny the use of callback, allow user-defined callback, or require callback to a particular number. You can contact your system administrator to determine your callback options.

- If your computer is configured to accept incoming connections, you can enforce callback options on that computer. For more information, see Configure callback for incoming connections.

# How security works at connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How security works at connection**

The following steps describe what happens during a call to a remote access server:

1. The remote access client computer calls a remote access server.

2. Depending on the authentication methods that have been negotiated, one of the following happens:

   - If you are using PAP or SPAP:

     The client sends its user name and password to the server.

     The server checks the account credentials against the user database.

   - If you are using CHAP or MS-CHAP:

     The server sends a challenge to the client.

     The client sends an encrypted response to the server.

     The server checks the response against the user database.

   - If you are using MS-CHAP v2:

     The server sends a challenge to the client.

     The client sends an encrypted response and a challenge to the server.

     The server checks the encrypted response of the client against the user database and sends the client its own encrypted response.

     The client verifies the encrypted response of the server.

   - If you are using certificate-based authentication:

     The server sends its computer certificate to the client.

     If the client is configured for certificate-based authentication, the client validates the server certificate.

     The client sends its user certificate to the server.

     The server verifies that the user certificate is valid and has not been revoked.

3. If the connection attempt is both authenticated (the user credentials are valid, the user account is enabled and not locked out, and the connection is occurring within the allowed logon hours) and authorized (for incoming connections, the user account has remote access permission), the server accepts the remote access connection.

   If callback is enabled, the server calls your computer back and repeats steps 2 through 3.

**Notes**

- For the Routing and Remote Access service, authorization of the connection is determined by the dial-in properties of the

user account and remote access policies. Remote access policies are a set of conditions and connection settings that give network administrators more flexibility in granting remote access permissions and usage. If the settings of the connection do not match at least one of the remote access policies that apply to your connection, the connection attempt is rejected, regardless of the dial-in properties of the user account.

- Network Connections authentication methods control access to your network, not to resources on the network. Once you are connected to a network, resource access control is managed through access control lists (ACLs) in various ways: Local Users and Groups, Active Directory, Group Policy, file and printer sharing, and so on.

- Authentication methods are designed for remote computing environments where the computer that is dialing in does not have access to the network until credentials, such as user name and password, are validated.

- For information about configuring security options for your connections, such as identity authentication protocols and data encryption settings, see Secure Network or Dial-up Connections.

# Group Policy for Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy for Network Connections**
By using the Group Policy snap-in, an administrator can configure the desktop settings of many computers at once. Some Group Policy settings apply specifically to Network Connections, such as user access to connections and the ability to change the properties of connections. These settings can be found in the Administrative Templates of the Group Policy snap-in.

### Where?

- *Group_Policy_object_name*/Computer Configuration (or User Configuration)/Administrative Templates/Network/Network Connections

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- If you do not have write access to a Group Policy object, you cannot see the settings contained in that object.

For an overview of Group Policy, Group Policy objects, and administrative templates, see Group Policy overview. For more information on how to use Group Policy to define policy settings that are applied to computers or users, see Group Policy (pre-GPMC).

# Group Policy settings that prohibit home and small office networking on your domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy settings that prohibit home and small office networking on your domain**

Group Policy settings define the various components of the user's environment that a system administrator can manage. The following Group Policy settings are available to prohibit home and small office networking on your domain:

- **prohibit use of Internet Connection Sharing on your DNS domain network** determines whether computer users with administrator accounts can enable and configure Internet Connection Sharing (ICS) on network connections on your domain.

- **prohibit use of Internet Connection Firewall on your DNS domain network** determines whether computer users with administrator accounts can enable and configure Internet Connection Firewall (ICF) on network connections on your domain.

- **prohibit installation and configuration of Network Bridge on your DNS domain network** determines whether computer users with administrator accounts can enable Network Bridge on your domain.

**Group Policy for ICS**

ICS allows a user with an administrator account to provide Internet access for a home or small office network, using one common connection as the Internet gateway, and to provide local private network services, such as name resolution and addressing. The ICS host is the only computer that is directly connected to the Internet. Multiple ICS clients simultaneously use the common (shared) Internet connection. ICS is available only on computers that have two or more network connections. Users with administrator accounts enable ICS by accessing the **Advanced** tab on network connections. For more information about ICS, see Connecting to the Internet in a home or small office network.

If you do not want users on your domain to be able to enable or configure ICS, apply the **prohibit use of Internet Connection Sharing on your DNS domain network** Group Policy setting. When you apply the **prohibit use of Internet Connection Sharing on your DNS domain network** Group Policy setting, users with administrator accounts cannot enable or configure ICS because the **Advanced** tab on network connections becomes unavailable, and the related wizard options on computers running Windows XP Professional are blocked.

Computer users with user rights are prohibited from configuring ICS, regardless of this Group Policy setting.

**Group Policy for ICF**

ICF allows a user with an administrator account to activate a firewall to protect the public connection of a small network or a single computer that is connected to the Internet. ICF is enabled by default only on dial-up connections. Users with administrator accounts can enable ICF on other types of connections, such as LAN and wireless connections, by accessing the **Advanced** tab on network connections. For more information about ICF, see Internet Connection Firewall.

If you do not want users to be able to enable or configure ICF, apply the **prohibit use of Internet Connection Firewall on your DNS domain network** Group Policy setting. When you apply the **prohibit use of Internet Connection Firewall on your DNS domain network** Group Policy setting, users with administrator accounts cannot enable or configure ICF because the **Advanced** tab becomes unavailable, and the related wizard options on computers running Windows XP Professional are blocked.

Computer users with user rights are prohibited from configuring ICF, regardless of this Group Policy setting.

**Group Policy for Network Bridge**

Network Bridge automates the configuration that is required to route traffic between multi-segment networks comprised of a single type of media or mixed media. With Network Bridge, no configuration is required, and you do not need purchase additional hardware such as routers or bridges. The Network Bridge menu command **Bridge Connections** is available only when two or more network adapters are present. By default, Network Bridge is disabled, but users with administrator accounts can use **Bridge Connections** to enable Network Bridge. For more information about Network Bridge, see Network Bridge.

If you do not want users on your domain to be able to enable or configure Network Bridge, apply the **prohibit installation and configuration of Network Bridge on your DNS domain network** Group Policy setting. When you apply the **prohibit**

**installation and configuration of Network Bridge on your DNS domain network** Group Policy setting, users with administrator accounts cannot enable or configure Network Bridge because the **Bridge Connections** command is removed from the menu on network connection icons, and the related wizard options on computers running Windows XP Professional are blocked.

Computer users with user rights are prohibited from configuring Network Bridge, regardless of this Group Policy setting.

**Important**

The Group Policy settings for ICS and ICF are location-aware. A location-aware Group Policy setting applies only when a computer is connected to the domain that it was connected to when the setting was last refreshed. If a computer, such as a laptop computer, is connected to your domain, but the setting on the computer was last refreshed on another domain, then the Group Policy setting on your network do not apply to that computer.

For example, if ICS is enabled on a computer that is on your network, and then you apply the Group Policy setting for ICS, ICS continues to run as normal until you restart the computer or log off of the network and then log on again. The same is true if the Group Policy setting for ICS is set on your network, and then a computer on which ICS is enabled joins your network. In this case also, ICS continues to run as normal until you restart the new computer or log off of the network and then log on again. This example applies to ICF as well.

The Group Policy setting for Network Bridge works somewhat differently. If Network Bridge is enabled on a computer that is on your network, and then you apply the Group Policy setting for Network Bridge, or if the Group Policy setting for Network Bridge is set on your network, and then a computer on which Network Bridge is enabled joins your network, Network Bridge continues to exist in the Network Connections folder and to use system resources, but the functionality of Network Bridge changes. In these cases, Network Bridge continues to receive and send data over the network connections that are included in the bridge but, for security reasons, no longer forwards traffic from one network connection to another. This change is immediate; it does not require that you restart the computer or log off of the network and then log on again.

For an overview of Group Policy, see Group Policy (pre-GPMC).

**Note**

- This topic applies only to product features available in the original release of Windows Server 2003.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Dial-up credentials

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up credentials**

Credentials are required to successfully connect to the Internet using dial-up connections. Your credentials are the combination of your ISP account name and password. Your computer user profile identifies your account as either an administrator account or a user account. If you log on using user credentials, you will have two options for how your credentials are saved when you create a dial-up connection. If you log on with administrative credentials, you will have three options for how your credentials are saved when you create a dial-up connection. If you have account profiles with both administrator and user credentials, only the credentials for the account that you have logged on with will apply.

If you log on with user account credentials, you can create only dial-up connections for **My use only**. These connections will be successful only when you are logged on to the computer, and they will not be accessible to any other users. When creating user account dial-up connections, you can either save your credentials as **Me only** (for your use only) or not save your credentials. If you save your credentials as **Me only**, and you are logged on to the computer using the account you used when you created the connection, your credentials will be automatically supplied when using the connection. If you do not save your credentials, a dialog box will be provided for you to enter your credentials each time you are logged on to the computer using the account you used when creating the connection and you attempt to connect to the Internet.

If you log on with administrative credentials, you can create connections for either **Anyone's use** or **My use only**. If you create a connection for **My use only**, the connection options are the same as those created by anyone with a user account. Connections created for **Anyone's use** can be seen by any user who logs on to that computer, but connection success to the Internet is affected by how credentials were saved. When creating a connection for **Anyone's use**, you can save your credentials so that they will be supplied automatically for **Anyone who logs on to this computer** or **Me only**. You can also choose not to save your credentials.

If you save your credentials for **Anyone who logs on to this computer**, your saved credentials will automatically be supplied for any user who has logged on to the computer and attempts to connect to the Internet using that connection or on-demand dialing. If you save your credentials but do not make them available for **Anyone who logs on to this computer**, then your saved credentials will be automatically supplied only if you are logged on to the computer. Other users will see a dialog box that prompts them for your user name and password; they will be unable to connect to the Internet using this connection unless they can provide your credentials. If you do not save your credentials, a dialog box will appear and prompt you or any other user to enter your credentials before the connection to the Internet can be completed.

The following summarizes the experience of a user logging on with user credentials and experience of a user logging on with administrative credentials:

**Logging on with user credentials**

| Connection type created | Credential save options | Local computer access |
|---|---|---|
| • **My use only** | • **Me only**<br>• Not saved | This connection will work only when you are logged on with the account that you used when you created the connection. If you did not save your credentials, you will be prompted to supply them every time you use the connection. You cannot configure Internet Connection Sharing (ICS), Internet Connection Firewall (ICF), or Windows Firewall on this connection. If you want to enable ICS or configure ICF or Windows Firewall, you should log on with administrative credentials before you create the dial-up connection because only administrators can access the **Advanced** tab. If you want to enable ICS, ICF, or Windows Firewall on a dial-up connection that was created with user credentials, the user who created the connection must be promoted to an administrator. |

| ICS client Internet access through My use only ICS connections |
|---|

| |
|---|
| Users of ICS client computers can connect to the Internet through connections on the ICS host computer that are saved for **My use only** if:<br>• The user who created the connection saved credentials and is logged on to the ICS host computer.<br>• The user who created the connection did not save credentials but that user or another user can supply those credentials when prompted by the ICS host computer to do so. |

**Logging on with administrative credentials**

| Connection type created | Credential save options | Local computer access |
|---|---|---|
| • **My use only** | • **Me only**<br>• Not saved | This connection will work only when you are logged on with the account that you used when you created the connection. If you did not save your credentials, you will be prompted to supply them every time you use the connection. You cannot configure Internet Connection Sharing (ICS), Internet Connection Firewall (ICF), or Windows Firewall on this connection. If you want to enable ICS or configure ICF or Windows Firewall, you should log on with administrative credentials before you create the dial-up connection because only administrators can access the **Advanced** tab. If you want to enable ICS, ICF, or Windows Firewall on a dial-up connection that was created with user credentials, the user who created the connection must be promoted to an administrator. |
| • **Anyone's use** | • **Anyone who uses this computer**<br>• **Me only**<br>• Not saved | All users can see the Internet connection. If you saved your credentials for all users, any user who is logged on to the computer can connect to the Internet. If you saved your credentials as **Me only**, the connection is automatically completed only if you log on to the computer as the administrator who saved the credentials when you created the connection. If you did not save your credentials, when you or another user attempts to connect to the Internet, your credentials must be supplied when prompted in order to complete the connection. You can access the **Advanced** tab to enable ICS or configure ICF or Windows Firewall because you logged on as an administrator when you created the dial-up connection. |

| **ICS client Internet access through Anyone's use connections** |
|---|
| Users of ICS client computers can connect to the Internet through connections on the ICS host computer that are saved for **Anyone's use** only if:<br>• The administrator who created the connection saved credentials for use by anyone who uses the computer. (This attempt will succeed whether a user is logged on to the ICS host computer or the ICS client computer.)<br>• The administrator who created the connection saved credentials as **User only** and is now logged on to the ICS host computer with the same credentials.<br>• The administrator who created the connection did not save credentials, but a user can supply those credentials when prompted by the ICS host computer. |

You must enable on-demand dialing in order for programs to automatically initiate dial-up connections to your Internet service provider. For more information, see Enable ICS on-demand dialing.

You must enable ICS in order for computers on your network to share a common Internet connection. For more information, see Connecting to the Internet in a home or small office network.

**Notes**

- As a security measure, you should log on using user credentials rather than administrative credentials.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

# Data Encryption and Encryption Methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data encryption and encryption methods**

This section covers:

- Data encryption
- Microsoft Point-to-Point Encryption (MPPE)
- Internet Protocol Security (IPSec) Encryption

# Data encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data encryption**

You can think of encryption as locking something valuable into a strong box with a key. Sensitive data is encrypted by using an encryption algorithm and a key, which renders it unreadable without the knowledge of the key. Data encryption keys are determined at connect time between the connecting computers. The use of data encryption can be initiated by your computer or by the server you are connecting to.

Network Connections supports two types of encryption:

- Microsoft MPPE, which uses RSA RC4 encryption.
- An implementation of Internet Protocol security (IPSec) that uses Data Encryption Standard (DES) encryption.

Both MPPE and IPSec support multiple levels of encryption, as shown in the following table.

| Encryption type | Level of encryption supported |
|---|---|
| MPPE Standard | 40-bit, 56-bit |
| MPPE Strong | 128-bit |
| IPSec DES | 56-bit |
| IPSec Triple DES | 3DES |

For more information, see the following:

- Microsoft Point-to-Point Encryption (MPPE)
- Internet Protocol Security (IPSec) Encryption

**Notes**

- Server controls are flexible and can be set to deny the use of encryption, require a specific encryption strength, or allow your computer to select an encryption strength. By default, most servers allow encryption and allow clients to select their encryption methods. This works for most computers. A system administrator sets encryption requirements. To determine your options, contact your system administrator.
- To enable MPPE-based data encryption for dial-up or virtual private network (VPN) connections, you must select the MS-CHAP, MS-CHAP v2, or EAP-TLS authentication methods. These authentication methods generate the keys used in the encryption process.
- Virtual Private Networks (VPNs) use encryption depending on the type of server they are connecting to. If the VPN connection is configured to connect to a PPTP server, then MPPE encryption is used. If the VPN is configured to connect to an L2TP server, then IPSec encryption methods are used. If the VPN is configured for an Automatic server type, which is the default selection, then L2TP, and its associated IPSec encryption, will be attempted first, then PPTP, and its associated MPPE encryption.

# Microsoft Point-to-Point Encryption (MPPE)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Point-to-Point Encryption (MPPE)**
Microsoft Point-to-Point Encryption (MPPE) encrypts data in Point-to-Point Protocol (PPP)-based dial-up connections or Point-to-Point Tunneling Protocol (PPTP) virtual private network (VPN) connections. 128-bit key (strong), 56-bit key, and 40-bit key (standard) MPPE encryption schemes are supported. MPPE provides data security for the PPTP connection that is between the VPN client and the VPN server.

For more information about configuring connections, see Configure a connection to a remote network.

**Note**

- MPPE requires encryption keys as generated by Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2), or Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication process.

# Internet Protocol Security (IPSec) Encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol security (IPSec)**
The long-term direction for secure networking, IPSec is a suite of cryptography-based protection services and security protocols. Because it requires no changes to applications or protocols, you can easily deploy IPSec for existing networks.

IPSec provides computer-level authentication, as well as data encryption, for VPN connections that use the L2TP protocol. IPSec is negotiated between your computer and an L2TP-based VPN server before an L2TP connection is established. This negotiation secures both passwords and data.

L2TP uses standard PPP-based authentication protocols, such as EAP, MS-CHAP, MS-CHAP v2, CHAP, SPAP, and PAP with IPSec.

Encryption is determined by the IPSec Security Association, or SA. A security association is a combination of a destination address, a security protocol, and a unique identification value, called a Security Parameters Index (SPI). The available encryptions include:

- Data Encryption Standard (DES), which uses a 56-bit key.
- Triple DES (3DES), which uses three 56-bit keys and is designed for high-security environments.

For more information about IPSec, see Internet Protocol Security (IPSec).

For more information about configuring connections, see Configure a connection to a remote network.

# Network Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication**
This section covers:

- Authentication process
- Authentication Protocols and Methods

# Authentication process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication process**
Authentication in the Windows Server 2003 family consists of two parts: an interactive logon process and a network authentication process. Successful user authentication depends on both of these processes.

**Interactive logon process**
The interactive logon process confirms the user's identification to either a domain account or a local computer. Depending on the type of user account, the interactive logon process is different:

- With a domain account, a user logs on to the network with a password or smart card by using single sign-on credentials stored in Active Directory. By logging on with a domain account, an authorized user can access resources in the domain and any trusting domains. Kerberos V5 is used for authentication if a password is used to log on to a domain account. Kerberos V5 authentication with certificates is used if a smart card is used instead.

- With a local computer account, a user logs on to a local computer by using credentials stored in the Security Accounts Manager (SAM), which is the local security account database. Any workstation or member server can store local user accounts.

**Network authentication process**
Network authentication confirms the user's identification to any network service that the user is attempting to access. To provide this type of authentication, many mechanisms are supported, including Kerberos V5, Secure Socket Layer/Transport Layer Security (SSL/TLS), and, for compatibility with Windows NT 4.0, LAN Manager.

Users who use a domain account do not see network authentication. Users who use a local computer account must provide credentials (such as a user name and password) every time they access a network resource. By using the domain account, the user has credentials that are automatically used for single sign-on.

For information about authentication methods and protocols used by Network Connections, see Authentication Protocols and Methods.

# Authentication Protocols and Methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication protocols and methods**

With Network Connections, you can use the following authentication protocols and methods:

- Password Authentication Protocol (PAP)
- Challenge Handshake Authentication Protocol (CHAP)
- Shiva Password Authentication Protocol (SPAP)
- Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)
- Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2)
- Extensible Authentication Protocol (EAP)
- Smart card and other certificate authentication

# Password Authentication Protocol (PAP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Password Authentication Protocol (PAP)**
Password Authentication Protocol (PAP) is a simple authentication protocol in which the user name and password is sent to the remote access server in a plaintext (unencrypted) form. Using PAP is strongly discouraged because your passwords are easily readable from the Point-to-Point Protocol (PPP) packets exchanged during the authentication process. PAP is typically used only when connecting to older UNIX-based remote access servers that do not support more secure authentication protocols.

To configure a connection for PAP, see Configure identity authentication and data encryption settings.

### Notes

- You cannot use Microsoft Point-to-Point Encryption (MPPE) if PAP is used to authenticate the connection.
- If the connection is configured to require a secured password and you attempt to connect to a server that is configured only for PAP, the Windows XP remote access client terminates the connection.

# Challenge Handshake Authentication Protocol (CHAP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Challenge Handshake Authentication Protocol (CHAP)**
Challenge Handshake Authentication Protocol (CHAP) is a widely supported authentication method in which a representation of the user's password, rather than the password itself, is sent during the authentication process. With CHAP, the remote access server sends a challenge to the remote access client. The remote access client uses a hash algorithm (also known as a hash function) to compute a Message Digest-5 (MD5) hash result based on the challenge and a hash result computed from the user's password. The remote access client sends the MD5 hash result to the remote access server. The remote access server, which also has access to the hash result of the user's password, performs the same calculation using the hash algorithm and compares the result to the one sent by the client. If the results match, the credentials of the remote access client are considered authentic. A hash algorithm provides one-way encryption, which means that calculating the hash result for a data block is easy, but determining the original data block from the hash result is mathematically infeasible.

To configure a connection for CHAP, see Configure identity authentication and data encryption settings.

### Notes

- You cannot use Microsoft Point-to-Point Encryption (MPPE) if CHAP is used to authenticate the connection.

- If you are using CHAP to authenticate a connection to a remote access server running Windows 2000 and Routing and Remote Access, the user account of the connecting client must be configured to allow the storage of the password in a reversibly encrypted form. For more information on configuring user accounts for storing a reversibly encrypted form of a password, see Enable reversibly encrypted passwords in a domain.

# Shiva Password Authentication Protocol (SPAP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shiva Password Authentication Protocol (SPAP)**
Shiva Password Authentication Protocol (SPAP) is a simple encrypted password authentication protocol supported by Shiva remote access servers. With SPAP, the remote access client sends an encrypted password to the remote access server. SPAP uses a two-way encryption algorithm. The remote access server decrypts the password and uses the plaintext form to authenticate the remote access client.

To configure a connection for SPAP, see Configure identity authentication and data encryption settings.

**Note**

- You cannot use Microsoft Point-to-Point Encryption (MPPE) if SPAP is used to authenticate the connection.

# Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)**
Microsoft created MS-CHAP to authenticate remote Windows-based workstations, integrating the functionality to which LAN-based users are accustomed with the hashing algorithms used on Windows networks. Like CHAP, MS-CHAP uses a challenge-response mechanism to authenticate connections without sending any passwords.

MS-CHAP uses the Message Digest 4 (MD4) hashing algorithm and the Data Encryption Standard (DES) encryption algorithm to generate the challenge and the response. MS-CHAP also provides mechanisms for reporting connection errors and for changing the user's password. The response packet is in a format designed to work with networking products in Windows 95, Windows 98, Windows Millennium Edition, Windows NT, Windows 2000, Windows XP, and the Windows Server 2003 family.

To configure a connection for MS-CHAP, see Configure identity authentication and data encryption settings.

**Notes**

- If you are configuring a connection to a server running Windows 95, you must use a specific version of MS-CHAP. To enable this older version of MS-CHAP, see Configure identity authentication and data encryption settings.
- Unlike CHAP, MS-CHAP does not require that the user's password be stored in a reversibly encrypted form.
- During the MS-CHAP authentication process, shared secret encryption keys for Microsoft Point-to-Point Encryption (MPPE) are generated.

# Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2)**
The Windows Server 2003 family supports MS-CHAP v2, which provides mutual authentication, the generation of stronger initial data encryption keys for Microsoft Point-to-Point Encryption (MPPE), and different encryption keys for sent and received data. To minimize the risk of password compromise during a password change, support for older methods of the MS-CHAP password change are not supported.

Because MS-CHAP v2 is more secure than MS-CHAP, it is offered before MS-CHAP (if enabled) for all connections.

MS-CHAP v2 is supported by computers running Windows XP, Windows 2000, Windows 98, Windows Millennium Edition, and Windows NT version 4.0. Computers running Windows 95 support MS-CHAP v2 only for VPN connections, not for dial-up connections.

To configure a connection for MS-CHAP v2, see Configure identity authentication and data encryption settings.

**Note**

- MS-CHAP v2 is a mutual authentication protocol, which means that both the client and the server prove that they have knowledge of the user's password. First, the remote access server asks the remote access client for proof by sending a challenge to the client. Then the remote access client asks the remote access server for proof by sending a challenge back to the server. If the server cannot prove that it has knowledge of the user's password by correctly answering the challenge from the client, the client terminates the connection. Without mutual authentication, a remote access client cannot establish a connection to an unauthorized remote access server.

# Extensible Authentication Protocol (EAP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Extensible Authentication Protocol (EAP)**
Extensible Authentication Protocol (EAP) extends Point-to-Point Protocol (PPP) by allowing arbitrary authentication methods that use credential and information exchanges of arbitrary lengths. EAP was developed in response to an increasing demand for authentication methods that use security devices, such as smart cards, token cards, and crypto calculators. EAP provides an industry-standard architecture for supporting additional authentication methods within PPP.

By using EAP, you can support additional authentication schemes, known as EAP types. These schemes include token cards, one-time passwords, public key authentication using smart cards, and certificates. EAP, in conjunction with strong EAP types, is a critical technology component for secure virtual private network (VPN) connections. Strong EAP types, such as those based on certificates, offer better security against brute-force or dictionary attacks and password guessing than password-based authentication protocols, such as CHAP or MS-CHAP.

To find out if an EAP type is being used in your organization, contact your network administrator.

To configure a connection for EAP, see Configure identity authentication and data encryption settings.

The Windows Server 2003 family supports two EAP types:

- EAP-MD5 CHAP (equivalent to the CHAP authentication protocol)
- EAP-TLS (used for user certificate-based authentication).

EAP-TLS is a mutual authentication method, which means that both the client and the server prove their identities to each other. During the authentication process, the remote access client sends its user certificate and the remote access server sends its computer certificate. If either certificate is not sent or is invalid, the connection is terminated.

**Note**

- During the EAP-TLS authentication process, shared secret encryption keys for Microsoft Point-to-Point Encryption (MPPE) are generated.

# Smart card and other certificate authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Smart card and other certificate authentication**

If a user certificate is installed, either in the certificate store on your computer or on a smart card, Extensible Authentication Protocol (EAP) is enabled, and **Smart Card or other Certificate EAP type (EAP-TLS)** is selected, you can use certificate-based authentication in a single network logon process, which provides tamper-resistant storage of authentication information.

A certificate is an encrypted set of authentication credentials. A certificate includes a digital signature from the certification authority that issued the certificate. In the EAP-TLS certificate authentication process, your computer presents its user certificate to the remote access server, and the remote access server presents its computer certificate to your computer, providing mutual authentication. Certificates are authenticated by using a public key to verify the included digital signature. The digital signature is contained in a trusted root certification authority certificate that is stored on your computer. These trusted root certificates are the basis for certificate verification. The Windows Server 2003 family provides many trusted root certificates. You should add or remove trusted root certificates only if your system administrator advises.

Certificates can reside either in the certificate store on your computer or on a smart card. A smart card is a credit-card-sized device that is inserted into a smart card reader, which is either installed internally in your computer or connected externally to your computer.

If you configure the advanced security settings for a connection, you can choose to use a smart card or other certificate, and you can specify particular certificate requirements. For example, you can specify that the computer certificate for the server must be validated, that the server name must end in a particular value, and that the computer certificate must be issued by a specific trusted root certification authority for the server.

When you double-click **New Connection Wizard** in the Network Connections folder and a smart card reader is installed, Windows detects it and prompts you to use it as the authentication method for the connection. If you decide not to use the smart card at the time you create a connection, you can modify the connection to use the smart card or other certificate at a later time. For more information, see Enable smart card or other certificate authentication.

If you are a member of an Active Directory domain, and need to request a certificate, see Request a certificate. If you are not a member of an Active Directory domain, or you need to request a certificate from the Internet, see Submit a user certificate request via the Web. For information about mobile users and certificates, see Mobile users and certificates.

# Intermediary Devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Intermediary devices**
This section covers:

- Using an intermediary device
- Security hosts
- Writing scripts for intermediary devices

# Using an intermediary device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using an intermediary device**

Network Connections supports various kinds of intermediary devices (security hosts and switches) between the remote access client and the remote access server. These devices include:

- Modem-pool switches
- Security hosts
- X.25 networks

Before you connect to the remote access server, you can have one of two possible dialog boxes (user input and computer response screens) with each intermediary device:

- Static (a dialog box that never changes and requires no input from you)
- Interactive (a dialog box that always changes and requires input from you)

You must configure your connection to work with each intermediary device.

If you require both static and interactive dialog boxes, you must take the following steps:

- Write a script for the static dialog box. For more information, see Writing scripts for intermediary devices.
- Activate Terminal mode for the interactive dialog boxes. For more information, see Use the Terminal feature to log on to a remote computer.

If you require only one kind of dialog box, take only one of the preceding steps. For example:

- If your clients connect through only one intermediary device with a static dialog box (such as an X.25 network), skip step 2.
- If your clients connect through a security host with an interactive dialog box, skip step 1.

# Security hosts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security hosts**
A security host is an authentication device that verifies whether a connection is authorized to connect to a remote access server. This verification supplements security already supplied by Network Connections and by the Windows Server 2003 family. There are two kinds of security hosts:

- Those which perform authentication checks during a modem call

  This type of security host sits between you and the remote access server and performs an authentication check before remote access server authentication. This type of security host generally provides an extra layer of security by requiring a hardware key of some sort in order to provide authentication. Verification that you are in physical possession of the key takes place before access to the remote access server is granted. With this open architecture, your system administrator can select from a variety of security hosts to augment the security in Network Connections, but the authentication check is restricted to modem calls.

  For example, one such security system consists of two hardware devices: the security host and the security card. The security host is installed between the remote access server and its modem. The security card is a small unit the size of a credit card that resembles a pocket calculator without keys. The security card displays a different access number every minute. This number is synchronized with a similar number calculated in the security host every minute. When connecting, you send a PIN number and the number on the security card to the host. If these match the number calculated on the host, the security host connects you with the remote access server.

  Another security host of the same type prompts you to type in a user name (which may or may not be the same as the remote access user name) and a password (which differs from the remote access password). The security host must be configured to allow the remote access server to initialize the modem before the security functions take effect. The remote access server must also be able to directly initialize the modem connected to the security host without security checks from the security host.

- Those which are called during the authentication process of the connection

  This type of security host provides custom authentication during the remote access authentication process. This authentication may augment or replace the standard verification of your network credentials by Routing and Remote Access in the Windows Server 2003 family. RADIUS servers are examples of this type of security host in that they perform user authentication on behalf of Routing and Remote Access. With the introduction of the Extensible Authentication Protocol (EAP), other vendors can create interfaces between remote access authentication and their own proprietary servers. Servers of this type are used to verify smart cards and other forms of extended authentication.

**Note**

- To make other security devices work with your connection, you must activate Terminal mode to interact with the security host. For more information, see Use the Terminal feature to log on to a remote computer.

# Writing scripts for intermediary devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Writing scripts for intermediary devices**

Each script for a static dialog box is a section of the Switch.inf file. You add a script for each intermediary device to the Switch.inf file. For an example of X.25 network scripts, see the Pad.inf file, located in the *systemroot*\System32\Ras folder.

When you connect by using modem pools or security hosts, you must have a dialog with that device. The dialog occurs either before dialing or after dialing. Although preconnect and postconnect dialogs can be either static or interactive, the preconnect dialog box is normally static and the postconnect is normally interactive.

The included Switch.inf file provides a generic script that will probably work with little or no modification when you are connecting to PPP servers. You should try to connect by using the generic Switch.inf script. If that does not work, you can copy--and then modify--the generic script to match the logon sequence of the remote computer to which you want to connect.

For more information about Switch.inf scripts, see Automating the Dial-up Logon Process by Using Switch.inf Scripts. For information about configuring a connection to use a Switch.inf script, see Activate a logon script.

# Tunneling Protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tunneling protocols**
This section covers:

- Point-to-Point Tunneling Protocol (PPTP)
- Layer Two Tunneling Protocol (L2TP)

# Point-to-Point Tunneling Protocol (PPTP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Point-to-Point Tunneling Protocol (PPTP)**

You can access a private network through the Internet or other public network by using a virtual private network (VPN) connection with the Point-to-Point Tunneling Protocol (PPTP).

PPTP enables the secure transfer of data from a remote computer to a private server by creating a VPN connection across IP-based data networks. PPTP supports on-demand, multiprotocol, virtual private networking over public networks, such as the Internet.

Developed as an extension of the Point-to-Point Protocol (PPP), PPTP adds a new level of enhanced security and multiprotocol communications over the Internet. By using the new Extensible Authentication Protocol (EAP) with strong authentication methods such as certificates, data transfer through a PPTP-enabled VPN connection is as secure as within a single LAN at a corporate site. For more information on EAP, see Extensible Authentication Protocol (EAP).

PPTP encapsulates IP or IPX protocols inside of PPP datagrams. This means that you can remotely run applications that are dependent upon particular network protocols. The tunnel server performs all security checks and validations, and enables data encryption, which makes it much safer to send information over nonsecure networks. You can also use PPTP in private LAN-to-LAN networking.

The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.




PPTP requires IP connectivity between your computer and the server. If you are directly attached to an IP LAN and can reach a server, then you can establish a PPTP tunnel across the LAN. If you are creating a tunnel over the Internet, and your normal Internet access is a dial-up connection to an ISP, you must dial up your Internet connection before you can establish the tunnel.

For more information about VPNs, see Virtual Private Network (VPN) Connections. For information about creating a VPN, see Make a virtual private network (VPN) connection.

# Layer Two Tunneling Protocol (L2TP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Layer Two Tunneling Protocol (L2TP)**

You can access a private network through the Internet or other public network by using a virtual private network (VPN) connection with the Layer Two Tunneling Protocol (L2TP). L2TP is an industry-standard Internet tunneling protocol with roughly the same functionality as the Point-to-Point Tunneling Protocol (PPTP). The Windows Server 2003 family implementation of L2TP is designed to run natively over IP networks. This implementation of L2TP does not support native tunneling over X.25 or Frame Relay networks.

Based on the Layer Two Forwarding (L2F) and Point-to-Point Tunneling Protocol (PPTP) specifications, you can use L2TP to set up tunnels across intervening networks. Like PPTP, L2TP encapsulates Point-to-Point Protocol (PPP) frames, which in turn encapsulate IP or IPX protocols, thereby allowing users to remotely run applications that are dependent upon specific network protocols.

The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.




With L2TP, the computer running a member of the Windows Server 2003 family that you are logging on to performs all security checks and validations, and enables data encryption, which makes it much safer to send information over nonsecure networks. By using Internet Protocol security (IPSec) authentication and encryption, data transfer through an L2TP-enabled VPN is as secure as within a single LAN at a corporate site. For more information about IPSec, see Internet Protocol Security (IPSec) Encryption.

For more information about VPNs, see Virtual Private Network (VPN) Connections. For information about creating a VPN connection, see Make a virtual private network (VPN) connection.

# Home and Small Office Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Home and small office networking**
This section covers:

- Internet Connection Firewall
- Internet Connection Firewall security log
- Network Bridge
- Service definitions for Internet Connection Firewall and Internet Connection Sharing

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Internet Connection Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Connection Firewall**

A firewall is a security system that acts as a protective boundary between a network and the outside world. Internet Connection Firewall (ICF) is firewall software that is used to set restrictions on what traffic is allowed to enter your network from the Internet. ICF protects your network against external threats by allowing safe network traffic to pass through the firewall into your network, while denying the entrance of unsafe traffic.

If your network uses Internet Connection Sharing (ICS) to provide shared Internet access to multiple computers, ICF should be enabled on the shared Internet connection.




Additionally, you should enable ICF on the Internet connection of any computer that is connected directly to the Internet. ICF can provide protection to a single computer connected to the Internet with a cable modem, a DSL adapter or modem, or a dial-up modem. You should not enable ICF on VPN connections because it will interfere with the operation of file sharing and other VPN functions.

When ICF has been enabled on a network connection, the following network connection icon appears in Network Connections: . To check to see whether ICF is enabled, or to enable ICF, see Enable or disable ICF. For information about using ICF, see Protecting your home or small office network using Internet Connection Firewall.

**How ICF works**

ICF is considered a "stateful" firewall. A stateful firewall is one that monitors all aspects of the communications that cross its path and inspects the source and destination address of each message that it handles. To prevent unsolicited traffic from the public side of the connection from entering the private side, ICF keeps a table of all communications that have originated from the ICF computer. When used in conjunction with ICS, ICF tracks all traffic that has originated from the ICF/ICS computer and all traffic that has originated from private network computers. All inbound traffic from the Internet is compared against entries in the table. Inbound Internet traffic is allowed to reach the computers in your network only when there is a matching entry in the table that shows that the communication exchange originated from your computer or private network.

To thwart common hacking attempts, such as port scanning, communications that originate from the Internet are dropped by the firewall. Rather than sending you notifications about firewall activity, ICF silently discards unsolicited communications, because such notifications could be sent frequently enough to become a distraction. Instead, ICF creates a security log to track this activity. For more information, see Internet Connection Firewall security log.

Services can be configured to allow unsolicited traffic from the Internet to be forwarded by the ICF computer to the private network. For example, if you are hosting an HTTP Web server service, and the HTTP service is enabled on your ICF computer, unsolicited HTTP traffic will be forwarded by the ICF computer. A set of operational information, known as a service definition, is required by ICF to allow the unsolicited Internet traffic to be forwarded to the Web server on your private network. Service definitions for ICF work on a per connection basis. If your network has multiple firewall connections, you should configure service definitions on all firewall connections. For information about service definitions, see Service definitions for Internet Connection Firewall and Internet Connection Sharing. For information about adding and editing service definitions, see Manage service definitions for ICF or ICS.

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and

Windows Server 2003, Enterprise Edition.

- This topic applies only to product features available in the original release of Windows Server 2003.

# Internet Connection Firewall security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Connection Firewall security log**

The Internet Connection Firewall (ICF) security log allows advanced users to choose which information to log. With ICF security logging you can:

- **Log dropped packets**. Log all dropped packets that originate from either the home or small office network or the Internet.
- **Log successful connections**. Log all successful connections that originate from either the home or small office network or the Internet.

When you select the **Log dropped packets** check box, information is collected each time ICF detects and denies traffic attempts. For example, if your Internet Control Message Protocol (ICMP) settings are not set to allow incoming echo requests, such as those sent by the **ping** and **tracert** commands, and an echo request from outside your network is received, the echo request is dropped, and an entry is made in the log.

When you select the **Log successful connections** check box, information is collected each time ICF detects and permits traffic attempts. For example, when a computer on your network successfully connects to a Web site, an entry is produced in the log.

The ICF security log is produced using the W3C Extended Log File Format, a format similar to the format that is used in common log analysis tools. For information about how to view an ICF security log, see View the ICF security log. To save the ICF security log using your choice of name and location, see Change the path and file name of the ICF security log.

ICF security logging is considered an advanced option and is not enabled by default. In order for the ICF security log to accept new data, you must select one or both of the logging options. For more information, see Enable or disable ICF security log options.

The ICF security log has two sections:

- The *header* provides information about the version of the ICF security log and the fields that are available for data entry. The header information is presented as a static list.
- The *body* contains compiled data that is entered as a result of traffic attempting to cross the firewall. The fields in the body of the ICF security log are entered from left to right across the page. The body of the ICF security log is a dynamic list--new data entries are entered at the bottom of the log.

The following tables define the information that the ICF security log contains:

**Header items**

| Item | Description | Example |
|---|---|---|
| #Version | Displays which edition of the ICF security log is installed. | 1.0 |
| #Software | Provides the name of the ICF security log. | Microsoft Internet Connection Firewall |
| #Time | Indicates that all of the timestamps in the log are in local time. | Local |
| #Fields | Displays a static list of the fields that are available for ICF security log entries, if data is available. Fields include: date, time, action, protocol, src-ip, dst-ip, src-port, dst-port, size, tcpflags, tcpsyn, tcpack, tcpwin, icmptype, icmpcode, and info. | Src-port, dst-port, size |

**Body data**

| Field | Description | Example |
|---|---|---|
| Date | Specifies the year, month, and day on which the recorded transactions occurred. Dates are recorded in the format:<br><br>`YY-MM-DD`<br><br>where YY is the year, MM is the month, and DD is the day. | 2001-01-27 |
| Time | Specifies the hour, minute, and seconds at which the recorded transaction occurred. Times are recorded in the format:<br><br>`HH:MM:SS`<br><br>where HH is the hour in 24 hour format, MM is minutes, and SS is seconds. | 21:36:59 |
| Action | Specifies which operation was observed by the firewall. The options available to the firewall are OPEN, CLOSE, DROP, and INFO-EVENTS-LOST. An INFO-EVENTS-LOST action indicates the number of events that happened but were not recorded in the log. | DROP |
| Protocol | Specifies which protocol was used for the communication. A protocol entry can be TCP, UDP, ICMP, or, if the protocol that was used was not TCP, UDP, or ICMP, a number. | ICMP |
| Src-ip | Specifies the source IP address (the IP address of the computer attempting to establish communications). The source IP is recorded in the dotted decimal format:<br><br>`(number).(number).(number).(number)` | 192.168.0.0 |
| Dst-ip | Specifies the destination IP address (the IP address of the destination for a communication attempt). The destination IP is recorded in the dotted decimal format:<br><br>`(number).(number).(number).(number)` | 192.168.0.1 |
| Src-port | Specifies the port number of the source (sending) computer. A src-port entry is recorded as a whole number, ranging from 1 to 65,535. Only TCP and UDP return a valid src-port entry. All other protocols are invalid for src-port and result in an entry of **-**. | 4039 |
| Dst-port | Specifies the port number of the destination computer. A dst-port entry is recorded as a whole number, ranging from 1 to 65,535. Only TCP and UDP return a valid dst-port entry. All other protocols are invalid for dst-port and result in an entry of **-**. | 53 |
| Size | Specifies the packet size in bytes. | 60 |
| Tcpflags | Specifies the TCP control flags found in the TCP header of an IP packet:<br><br>• **A**ck-Acknowledgment field significant<br><br>• **F**in-No more data from sender<br><br>• **P**sh-Push function<br><br>• **R**st-Reset the connection<br><br>• **S**yn-Synchronize sequence numbers<br><br>• **U**rg-Urgent pointer field significant<br><br>Flags are written as uppercase letters. The entry information for tcpflags is provided for users with an in-depth knowledge of Transmission Control Protocol (TCP). Additional information about TCP can be found in RFC 793. | AFP |
| Tcpsyn | Specifies the TCP sequence number in the packet. The entry information for tcpsyn is provided for users with an in-depth knowledge of TCP. | 1315819770 |

| | | |
|---|---|---|
| Tcpack | Specifies the TCP acknowledgement number in the packet. The entry information for tcpack is provided for users with an in-depth knowledge of TCP. | 0 |
| Tcpwin | Specifies the TCP window size in the packet. The size is specified in bytes. The entry information for tcpack is provided for users with an in-depth knowledge of TCP. | 64240 |
| Icmptype | Specifies a number that represents the type field of the ICMP message. | 8 |
| Icmpcode | Specifies a number that represents the code field of the ICMP message. | 0 |
| Info | Specifies an information entry that depends on the type of action that occurred. For example, an INFO-EVENTS-LOST action displays the number of events that occurred but were not placed in the log since the last occurrence of this event type. | 23 |

The character (**-**) is entered when no information is available for a field.

If ICF is enabled on two or more connections on a single computer, the ICF settings are global. If you select a setting or change a setting in **Logging Options** or **Services** on any one of the connections on which ICF is enabled, that setting is applied to all of the ICF firewalls on that computer.

**Notes**

- If the maximum allowable size for pfirewall.log is exceeded, the information that pfirewall.log contains is saved as pfirewall.log.old, and new information is saved as pfirewall.log. If the maximum allowable size for pfirewall.log is exceeded again, the information that pfirewall.log contains is saved as pfirewall.log.old, and the information that had been in pfirewall.log.old is overwritten. New information is always saved in pfirewall.log.

- You can obtain RFCs from the RFC Editor Web site. This Web site is currently maintained by members of the Information Sciences Institute (ISI) who publish a classified listing of all RFCs. RFCs are classified as one of the following: approved Internet standards, proposed Internet standards (circulated in draft form for review), Internet best practices, or For Your Information (FYI) documents.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

- This topic applies only to product features available in the original release of Windows Server 2003.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

# Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Bridge**
Network Bridge makes it inexpensive and easy to connect LAN segments. A LAN segment is a section of network media that connects computers. Frequently, a network has more than one LAN segment. Prior to Windows XP; Windows Server 2003, Standard Edition; and Windows Server 2003, Enterprise Edition, if you wanted to have a network with more than one LAN segment, then you had two options: IP routing and hardware bridging. IP routing required you to buy hardware routers or configure computers to act like routers, configure IP addressing for each computer on each network segment, and configure each network segment as a separate subnet. Hardware bridging did not require difficult configurations, but it did require you to purchase hardware bridges. Additionally, if you were using different types of network media, you needed to create a separate subnet for each type of media.

In contrast, the Network Bridge feature that is available with Windows XP; Windows Server 2003, Standard Edition; and Windows Server 2003, Enterprise Edition allows you to connect LAN segments simply by clicking the **Bridge Connections** menu command. No configuration is required, and you do not need to purchase additional hardware, such as routers or bridges. Network Bridge automates the configuration that is required to route traffic between multi-segment networks that consist of a single type of media or mixed media.




**Network Bridge example**
Suppose you have a small office network with four computers (PC1, PC2, PC3, and PC4) and one Ethernet hub. The four computers are running Windows XP; Windows Server 2003, Standard Edition; or Windows Server 2003, Enterprise Edition and have the following hardware installed:

- PC1 has an adapter connecting it to the Internet, an Ethernet network adapter, an HPNA network adapter, and a wireless adapter.
- PC2 has an Ethernet network adapter.
- PC3 has an HPNA network adapter.
- PC4 has a wireless network adapter.

The Ethernet adapters on PC1 and PC2 are connected to a common Ethernet hub to form the first LAN segment. PC1 is connected to PC3 with the HPNA adapter to form a second LAN segment, and PC1 is connected to PC4 with the wireless adapter to form a third LAN segment.

You can use Network Bridge to connect the Ethernet network adapter, the HPNA network adapter, and the wireless network adapter on PC1. Network Bridge can forward traffic from one LAN segment to another and enable all of your computers to communicate with each other.




Without Network Bridge (or additional routing configurations or bridging hardware), only PC1 can communicate with each of the other computers because PC1 is the only computer that has connections to all three LAN segments. Because PC2, PC3, and PC4 use different types of network media, they are on different LAN segments, and they are unable to communicate with any computer other than PC1 (without Network Bridge or additional routing configurations or bridging hardware).

**Network Bridge considerations**

Only one bridge can exist on a computer, but it can be used to bridge as many different network connections as the computer can physically accommodate. For information about creating Network Bridge, see Set up Network Bridge.

In order to include the ICS private connection as part of a Network Bridge, you must temporarily disable ICS. Be careful not to include the ICS public connection as part of a Network Bridge. To do so creates an unprotected link between your network and the Internet, leaving your network vulnerable to intrusions, and prevents you from enabling ICF or ICS. For more information about using ICS, see Connecting to the Internet in a home or small office network.

Only Ethernet, IEEE-1394 adapters, or Ethernet-compatible adapters, such as wireless and Home Phoneline Network adapters, can be a part of Network Bridge. Network Bridge with a wireless or IEEE-1394 connection supports traffic using Internet Protocol version 4 (IPv4) only.

**Spanning tree algorithm**

Network Bridge uses the IEEE spanning tree algorithm (STA) to establish a loop-free forwarding topology. When there are multiple paths in a bridged network, loops can form, and the simple forwarding rules of a bridge can cause forwarding storms--a condition in which the same frame is relayed endlessly from one network segment, across the bridge, to another segment. STA provides an automated mechanism to selectively disable bridge forwarding on individual ports as is necessary, in order to ensure that the network topology is loop-free. There is no configuration needed to configure Network Bridge for STA.

### Note

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Service definitions for Internet Connection Firewall and Internet Connection Sharing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service definitions for Internet Connection Firewall and Internet Connection Sharing**

Services run transparently to support other programs. An example of a service is an HTTP Web server that supports hosting Web pages from your home or small office network. If the HTTP Web server service is enabled for your network, you can use that service to supply your own Web pages to the Internet. In addition to having a Web server on your network, you can have an FTP service that permits users to upload or download files through the Internet, an Internet mail server for storing and forwarding mail, and a variety of other services.

In order to permit traffic to flow from the Internet to the computer hosting the service, you must use a preconfigured service or add the service to the **Services** list by entering information about the operational settings of the services on the **Services** tab of the Internet Connection Firewall (ICF) or Internet Connection Sharing (ICS) host computer. The operational settings are known as a "service definition," because they define the environment that the service requires in order to operate. Service definitions for ICF and ICS provide the port mappings that are required in order for ICF to allow unsolicited traffic to travel from the Internet to the network.

Some service definitions are preconfigured for ICF and ICS. FTP server, Internet Mail server, Secure Web server, Telnet server, and Web server are examples of preconfigured (built-in) services. These services can run on ICS client computers or on your ICF or ICS host computer. You can add a new service to your network by installing the service software on one of your network computers and then adding the service definition so that ICF and ICS, if one or both are enabled, will allow the service to be accessed from the Internet.

Services for ICF and ICS can run on an ICS client computer or an ICS host computer. When you add a service definition to an ICS host computer, it is made available to ICS clients. For information about how to add or edit a service definition, see Manage service definitions for ICF or ICS.

**Notes**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

# Using Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Network Connections**

Network Connections is automatically installed. This section covers:

- Monitoring network and dial-up connections
- Making Network Connections
- Using Home and Small Office Networking Features

# Monitoring network and dial-up connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring network and dial-up connections**
You can obtain information about an active connection by using the **Status** menu command on the **File** menu.

By using **Status**, you can view:

- The duration of a connection.
- The speed at which you initially connected.

  For a single-link connection and for individual links in a multilink connection, this speed is negotiated (and fixed) at the time the connection or link is established. For multilink connections, this speed is equal to the sum of the speeds of the individual links. For multilink connections, this speed varies as links are added or deleted.
- For local area connections, the number of bytes transmitted and received during a connection. For other types of connections, the number of bytes transmitted and received during a connection, and the associated compression and error statistics.
- The diagnostic tools that you can use for a connection, if any. For example, the Windows Network Troubleshooter, TCP/IP, Autoping, and TCP/IP Windows IP Configuration. These tools are explicitly registered in the Windows Server 2003 family by their providers, whether the provider is Microsoft or another organization.

You can also generate diagnostic reports for remote access connections. For more information on remote access preferences and generating diagnostic reports, see Logging remote access connections and Enable logging.

For more information on monitoring connections, see Monitor a connection.

**Note**

- If your computer is set up to accept incoming connections, a connection icon with a user name assigned to it appears in the Network Connections folder as each user connects. You can view the progress of incoming connections by right-clicking a named connection, and then clicking **Status**.

# Making Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Making network connections**

This section covers:

- Using local area connections
- Using direct connections
- Using Remote Access Methods
- Using Incoming Connections
- Using Internet Connections
- Virtual Private Network (VPN) Connections

# Using local area connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using local area connections**

Typically, computers running Windows are connected to a local area network (LAN). When you install Windows, your network adapter is detected, and a local area connection is created. It appears, like all other connection types, in the Network Connections folder. By default, a local area connection is always activated. A local area connection is the only type of connection that is automatically created and activated.

If you disable your local area connection, the connection is no longer automatically activated. Because your hardware profile remembers this, it accommodates your location-based needs as a mobile user. For example, if you travel to a remote sales office and use a separate hardware profile for that location that does not enable your local area connection, you do not waste time waiting for your network adapter to time out. The adapter does not even try to connect.

If your computer has more than one network adapter, a local area connection icon for each adapter is displayed in the Network Connections folder.

Examples of LAN connections include Ethernet, token ring, cable modems, DSL, FDDI, IP over ATM, IrDA (Infrared), wireless, and ATM-emulated LANs. Emulated LANs are based on virtual adapter drivers such as the LAN Emulation Protocol.

If changes are made to your network, you can modify the settings of an existing local area connection to reflect those changes. The **General** tab of the *Local Area Connection***Status** dialog box allows you to view connection information such as connection status, duration, speed, signal strength, amounts of data transmitted and received, and any diagnostic tools available for a particular connection. The **Support** tab contains information on:

- The address type which indicates how the address was assigned. For example the TCP/IP address is assigned by DHCP.
- The IP address currently assigned for the session.
- The IP subnet mask for the IP address currently assigned for the session.
- The default gateway address of the IP device that allows access to other protocols.

The **Support** tab also has a **Details** button that displays detailed information about the properties of the network connection. This includes the addresses of dependent external devices. For information about using the **Status** menu option, see View the status of a local area connection.

If you install a new LAN adapter in your computer, the next time you start your computer, a new local area connection icon appears in the Network Connections folder. Plug and Play functionality finds the network adapter and creates a local area connection for it. If you are using a laptop computer, you can add a PC card while the computer is on. Plug and Play will identify the new card without you having to restart your laptop computer. The local area connection icon is immediately added to the folder. You cannot manually add local area connections to the Network Connections folder.

You can configure multiple LAN adapters through the **Advanced Settings** menu option. You can modify the order of adapters that are used by a connection, and the associated clients, services, and protocols for the adapter. You can modify the provider order in which this connection gains access to information on the network, such as networks and printers.

You configure the device a connection uses, and all of the associated clients, services, and protocols for the connection, through the **Properties** menu option. Clients define the access of the connection to computers and files on your network. Services provide features such as file and printer sharing. Protocols, such as TCP/IP, define the language your computer uses to communicate with other computers.

Depending on the status of your local area connection, the icon changes appearance in the Network Connections folder, or a separate icon appears in the taskbar. If a LAN adapter is not detected by your computer, a local area connection icon does not appear in the Network Connections folder. For a list of connection icons, see Network connection icons.

# Using direct connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using direct connections**

With Network Connections, you can create a physical connection to another computer by using a serial cable, a DirectParallel cable, a modem, infrared devices or other methods. For example, you might want to connect a handheld PC running Windows CE to your network. Or, you may have two or more networks in the same location that are not physically connected. If you want to use resources on both networks from one computer, you can use a serial cable connection with an RS-232C null modem cable. You connect an RS-232C cable from a COM port on the computer to a COM port on your remote access server, which is used to create network access. You can also use an RS-232C null modem as a substitute for a network adapter if your computer is located physically near (less than 50 feet of cable) a remote access server.

The DirectParallel driver supports computer-to-computer connections by using standard and enhanced (ECP) parallel ports, which use Parallel Technologies' Basic or Fast parallel cables. For more information about using a DirectParallel cable for a direct connection, see DirectParallel cable access.

Remote access authentication methods control access to your network, not resources on the network. Having a Windows CE device directly attached to a computer on your network is analogous to having a computer attached to your network, but not logged on. In that instance, you attached and accessed the network without using authentication of any kind.

Once you are present on a network, resource control is managed in various ways: Local Users and Groups, Active Directory, Group Policy, file and printer sharing, and so on.

Therefore, if you have set up a direct network connection using, for example, a Windows CE device, you need to ensure that resource control is in place that enforces access to restricted folders, printers, and so on.

Authentication methods are designed for remote computing, where the computer dialing in does not have access to the network until credentials such as user name and password are exchanged.

For information about creating a direct network connection, see Make a direct network connection.

**Note**

- To create or modify a direct network connection, you must be a member of the Administrators group.

# Using Remote Access Methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using remote access methods**

This section covers:

- Using phone lines and modems
- Using ISDN lines
- Using X.25

# Using phone lines and modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using phone lines and modems**

The most common dial-up connection is made by using DSL, a cable modem, or a standard analog telephone line and a modem. Standard analog phone lines are available worldwide and meet most requirements of a mobile user. Standard analog phone lines are also called PSTN (public switched telephone network) or POTS (plain old telephone service).

Hundreds of modems are compatible with Windows XP and all members of the Windows Server 2003 family. Most modems that comply with industry standards should interoperate. However, many difficult-to-detect problems are the result of incompatible modems. To verify the compatibility of a modem, see the compatibility information at Support resources.

Modems are automatically detected. This is especially useful for users who are not sure what modem is installed (for example, an internal modem). You can use shared modem pools for dial-up connections. Modem data compression and error control are available on most modems. However, built-in software compression offers enhanced performance over modem data compression.

For information about connecting to a network by using a phone line, see Make a dial-up connection by using a phone line.

# Using ISDN lines

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using ISDN lines**

To enhance connection speeds, you can use an Integrated Services Digital Network (ISDN) line. Whereas standard phone lines typically transmit from 28.8 to 56 kilobits per second (Kbps), typical ISDN devices are capable of 64 or 128 Kbps. This is slower than local area networks supported by high-speed data communications technology, but faster than analog telephone lines. ISDN provides end-to-end digital connectivity between your computer and a remote computer or network by using a local telephone line for network connection.

An ISDN line must be installed by the phone company at your site and at the remote access server. Also, an ISDN adapter must be installed in place of a modem in both your computer and the remote access server. For information about installing an ISDN adapter, see Install an ISDN adapter. Costs for ISDN equipment and lines may be higher than standard modems and phone lines. However, the speed of communication reduces the duration of the connection time, thereby possibly saving toll charges.

An ISDN line comes with two B channels that transmit data at 64 Kbps, and one D channel for signaling that transmits data at 16 Kbps. You can configure each B channel to operate as a port. With this configuration, line speed increases to 128 Kbps. Channel aggregation schemes are specific to each driver, so you must use the same driver and equipment on the client and server.

Multilink dialing over ISDN combines multiple physical links into a logical bundle. This aggregate link increases your bandwidth. To use Multilink, both your connection and the remote access server must have Multilink enabled. For more information about the Multilink feature, see Configuring multiple device dialing.

For information about creating an ISDN network connection, see Make a dial-up connection by using an ISDN line. For information about configuring ISDN network connection settings, see Configure ISDN settings.

# Using X.25

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using X.25**

An X.25 network transmits data with a packet-switching protocol, bypassing noisy telephone lines. This protocol relies on an elaborate worldwide network of packet-forwarding nodes that can participate in delivering an X.25 packet to its designated address.

Network Connections supports X.25 by using packet assemblers/disassemblers (PADs) and X.25 cards. You can also use a modem and special dial-up X.25 carriers in place of a PAD or X.25 smart card on your computer.

Remote access clients running Windows XP or a member of the Windows Server 2003 family can either use an X.25 card or dial in to an X.25 PAD to create connections. To accept incoming connections on a computer using X.25 running Windows XP or a member of the Windows Server 2003 family, you must use an X.25 card.

For information about creating an X.25 connection, see Make a dial-up connection by using X.25.

# Using Incoming Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using incoming connections**

This section covers:

- Incoming connections
- Incoming connections and IP addresses
- Incoming connection clients

# Incoming connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incoming connections**
With the Windows Server 2003 family, you can create incoming connections and allow even computers that are members of a domain to run incoming connections.

You can configure an incoming connection to accept the following connection types: Dial-up connections (modem, ISDN, X.25), Virtual Private Network (VPN) Connections (PPTP, L2TP), or Using direct connections (serial, infrared, DirectParallel). On a computer running Windows 2000 professional or Windows XP Professional, an incoming connection can accept up to three incoming calls, up to one of each of these types. On a computer running Windows 2000 Server, the number of inbound calls is only limited by the computer and its hardware configuration. For information about Windows NT Server (and other operating systems) that can accept an incoming connection, see Incoming connection clients.

You can use the Multilink feature to configure multiple modems or ISDN adapters for use with incoming dial-up connections. For more information on Multilink, see Configuring multiple device dialing.

When you create a connection, the users that can connect to your incoming connection and their network protocols are selected. Each user that connects to an incoming connection must have established a local user account. For more information about user settings such as name, password, and callback rights, see Grant incoming connection access rights to your computer. For more information about configuring your connection for the TCP/IP network protocol, see Configure an incoming connection to use TCP/IP.

For information about creating an incoming connection, see Accept incoming network connections.

**Notes**

- For large numbers of incoming connections on a computer that is running Windows 2000 Server or a member of the Windows Server 2003 family and that operates as part of a distributed network or as a domain controller, you can use Routing and Remote Access to create a remote access server.

- To create an incoming connection, you must be a member of the Administrators group.

- Users do not need to use incoming connections to allow access to shared resources, such as files and printers, over a local area network. In order to enable shared access to resources on the local computer, you must enable file and print sharing, share the resources, and then set the appropriate permissions. For information about sharing a folder, see Share a folder or drive.

- The Guest account is disabled by default for both incoming connections and in the Computer Management snap-in. Guests can connect only after you enable the Guest account in the Computer Management snap-in and in the **Users** tab of the properties dialog box of the connection.

# Incoming connections and IP addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incoming connections and IP addresses**
The incoming connections computer will assign IP addresses to the remote computers that connect to it. These addresses can be allocated in one of two ways:

- They are automatically allocated by DHCP.
- They are assigned based on a range of addresses that is defined in the TCP/IP properties for the incoming connection.

If you want to specify a range of IP addresses, you must provide a starting IP address and an ending IP address for the range. Network Connections calculates the number of addresses that can be used within the starting and ending addresses of the range. For more information about specifying the range, and the effects of a subnet on the address range, see Configure an incoming connection to use TCP/IP.

If the incoming connections computer is attached to another network, and the specified address range consists of a subset of IP addresses for the network to which the computer is attached, you must ensure that the addresses in the range are not already assigned to other TCP/IP nodes within the attached network. This is recommended because using a subset of the larger network's addresses requires less manual configuration of the intranet nodes.

If you must use a range of addresses that is not a subset of the network to which the incoming connections computer is attached, then steps must be taken to ensure that computers on the network attached to the incoming connections computer can reach the incoming connection clients. To do this, on each of the computers on the network you can create a route to the incoming client computers with the IP address of the incoming connections computer as the associated gateway.

For more information, see To configure a default gateway.

**Notes**

- If Internet Connection Sharing (ICS) is already enabled on the dial-up, virtual private network (VPN), or local area connection of an incoming connections computer, you do not need to configure the default gateway setting. ICS automatically makes the ICS host computer the default gateway for all of the computers on the intranet.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Incoming connection clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incoming connection clients**
The Windows Server 2003 family supports incoming connections from clients running Windows XP, Windows 2000, and the following operating systems:

- Windows NT version 4.0 clients

  Windows NT version 4.0 clients can take full advantage of all incoming connection features except dynamic allocation of multilinked lines, Extensible Authentication Protocol (EAP), IPSec client to server data security, and L2TP for tunneling across intervening networks.

- Windows NT version 3.1, 3.5, and 3.51 clients

  Windows NT version 3.5*x* clients can take full advantage of incoming connection features, except for Multilink functionality. Windows NT 3.5*x* clients negotiate the logon process and authentication with the incoming connection. You can also configure these clients to use scripts that completely automate the logon process.

  Windows NT version 3.1 clients do not support the implementation of Point-to-Point Protocol (PPP) that was introduced in Windows NT 3.5. Only Windows NT 3.5*x* or other PPP clients provide the support necessary to run TCP/IP or IPX applications on clients that directly communicate with servers on the LAN by using TCP/IP or IPX.

- Windows 98 clients

  Windows 98 clients can take full advantage of incoming connection features, except dynamic allocation of multilinked lines and the Extensible Authentication Protocol. Windows 98 clients negotiate the logon process and authentication with the server. You can also configure these clients to use scripts that completely automate the logon process.

- Windows 95 clients

  Windows 95 clients can take full advantage of incoming connection features, except the use of Multilink, dynamic allocation of multilinked lines, and Extensible Authentication Protocol. Windows 95 clients negotiate the logon process and authentication with the server. You can also configure these clients to use scripts that completely automate the logon process.

- Macintosh clients

  Macintosh users can connect to a computer running Services for Macintosh and access Macintosh file volumes and AppleTalk printers by using the AppleTalk Control Protocol (ATCP).

- PPP clients

  PPP clients that other companies offer and that use TCP/IP can access an incoming connection. The incoming connection automatically negotiates authentication with PPP clients; special configuration for these PPP clients is not required.

  For more information about your PPP client, see the software documentation for your PPP client.

For information about creating and configuring incoming connections, see Set up Incoming Connections.

**Notes**

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

- IPX cannot be configured for incoming or dial-up connections in the Windows Server 2003 family.

# Using Internet Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Internet connections**
With Network Connections, connecting to the Internet is easy. You can use the following components to gain access to the Internet:

- The TCP/IP protocol that is enabled for your network connection.
- If you are part of a corporate domain, a user account with remote access permissions.
- A modem or other connection to an Internet service provider (ISP).
- An account with an ISP.

This section covers:

- ISP access methods
- Internet service providers

# ISP access methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ISP access methods**
You can create a connection to an Internet service provider (ISP) through a modem and telephone line, a Point-to-Point Protocol over Ethernet (PPPoE) broadband connection, an ISDN adapter and ISDN line, or an X.25 network.

If you are using a modem, you will want the fastest possible speed to reduce downloading time while on the Internet. A modem of 28.8 Kbps or above is recommended. For more information about modems, see Installing and configuring Modems.

Your Internet connection dials the ISP and logs you on. The logon sequence varies based on the requirements of the ISP that you are calling. Point-to-Point Protocol (PPP) connections are often entirely automatic. Serial Line Internet Protocol (SLIP) connections may require that you log on using the terminal window, which may or may not permit you to automate the logon process with Switch.inf file scripts. You can set the terminal window to appear during the logon process by selecting **Show terminal window** on the **Security tab** of the **Dial-up Connection Properties** dialog box.

For more information, see:

- TCP/IP
- Internet service providers
- Make an Internet connection

**Note**

- Before you create an Internet connection, check with your ISP to verify the required connection settings. A connection to your ISP may require one or more of the following settings:
    - A specific IP address.
    - IP header compression (for PPP).
    - DNS addresses and domain names.
    - Optional settings, such as Internet Protocol security (IPSec).

For information about creating an Internet connection, see Make an Internet connection

# Internet service providers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet service providers**

There are Internet service providers (ISPs) around the world. Similar to other online services or bulletin boards, you dial the service number and log on to the remote system. Once connected, you have access to the Internet and any other services, such as electronic mail, that are provided by the ISP. Fees usually apply for commercial ISPs.

For information about creating an Internet connection, see Make an Internet connection.

# Virtual Private Network (VPN) Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private network (VPN) connections**

A virtual private network (VPN) extends the capabilities of a private network to encompass links across shared or public networks such as the Internet. With a VPN, you can send data between two computers across a shared or public network in a manner that emulates a point-to-point link on a private network. Virtual private networking is the act of creating and using a virtual private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information, which allows the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data is encrypted for confidentiality. Packets that are intercepted on the shared or public network cannot be read without the encryption keys. The link in which the private data is encapsulated and encrypted is a VPN connection.

The following table describes the advantages of using VPN connections. For more information, click an advantage.

| Advantage | Example |
|---|---|
| Cost advantages | The Internet is used as a connection instead of a long-distance telephone number or 1-800 service. Because an Internet service provider (ISP) maintains communications hardware, such as modems and ISDN adapters, to access the Internet, your network requires less hardware to purchase and manage. |
| Outsourcing dial-up networks | You can make a local call to the telephone company or ISP, which then connects you to a remote access server and your corporate network. The telephone company or ISP manages the modems and telephone lines required for dial-up access. Because the ISP supports complex configurations of communication hardware, a network administrator is free to centrally manage user accounts at the remote access server. |
| Enhanced security through VPN | The connection over the Internet is encrypted and secure. Authentication and encryption are enforced by the VPN server. Sensitive data is hidden from unauthorized users, but it is accessible to users authorized through the connection. |
| Network protocol support | You can remotely run any application that depends on the most common network protocols, including TCP/IP and IPX. The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family. |
| IP address security | Because information sent over a VPN is encrypted, the addresses you specify are protected, and the Internet sees only the external IP address. For organizations with private addresses, this advantage is substantial, because no administrative costs are incurred from having to change IP addresses for remote access over the Internet. |

You can create a VPN connection by dialing an ISP or by connecting directly to the Internet, as shown in the following examples.

**Example 1**

- In the first example, the VPN user makes a call to an ISP. After a connection is established with the ISP, the ISP server then makes a call to the remote access server that establishes the Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP) tunnel. These protocols are automatically installed on your computer. After the VPN connection is made, you can access the corporate network, as shown in the following illustration.




**Example 2**

- In the second example, a user who is already connected to the Internet uses a VPN connection to connect to the VPN server. Examples of this type of user include a person whose computer is connected to a local area network, a cable

modem user, or a subscriber of a service such as ADSL, where IP connectivity is established immediately after the user's computer is turned on. The VPN client makes a tunnel through the Internet and connects to the PPTP-enabled or L2TP-enabled VPN server. After the VPN connection is made, the user can access the corporate network, achieving the same functionality as the user in the preceding example.




For information about creating a VPN connection, see Make a virtual private network (VPN) connection. For information about the ways Active Directory can tighten security measures for VPNs, see VPN administration.

**Note**

- Connecting directly to the Internet means direct IP access without going through an ISP. (For example, some hotels allow you to use an Ethernet cable to connect to the Internet.)

# Cost advantages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cost advantages**

By using virtual private networks (VPNs), mobile professionals and telecommuters can access their corporate local area network (LAN) over the Internet for a fraction of the cost of traditional remote access solutions. Just as it is more cost effective to make use of the phone company's massive switching system, rather than to string up your own telephone lines and invest in your own switch, a connection to the IP address or DNS name of your Point-to-Point Tunneling Protocol (PPTP)-enabled VPN server makes use of the physical base of routers that are available on the Internet, asynchronous transfer mode (ATM) switches, and digital and analog lines, without sacrificing security.

You can use a network adapter to dial in to a virtual private network, as you would dial a modem to connect to a traditional remote server. VPNs offer corporations a release from the cost and maintenance obligations of components such as modem banks and dedicated analog (1-800) phone lines. The modems and their associated infrastructure remain centralized at the Internet service provider's locations, without sacrificing security or the ability to control remote connections. At the same time, the benefits of secure access to private data are ensured by additional VPN-required authenticated access, encryption, and user data compression.

For information about creating a VPN connection, see Make a virtual private network (VPN) connection.

# Outsourcing dial-up networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Outsourcing dial-up networks**

The communications hardware that is available for supporting dial-up needs can be complicated and not well-integrated. For a large organization, putting together a remote access server requires modems, serial controllers, and many cables. Furthermore, many solutions do not provide a single integrated way to efficiently support V.34 and ISDN dial-up lines.

Many corporations like to outsource dial-up access to their corporate backbone networks in a manner that is cost-effective, hassle-free, protocol-independent, secure, and that requires no changes to the existing network addressing. Virtual WAN support by using Virtual Private Network (VPN) Connections is one way an Internet service provider (ISP) can meet the needs of corporations. In this way, the outsourcing company maintains and manages remote access modems and lines, leaving a system administrator free to manage users and their authentication needs.

This type of solution takes advantage of proven PPP authentication, encryption, and compression technologies. For information about creating a VPN connection, see Make a virtual private network (VPN) connection.

The following illustration shows an example of outsourcing.




The connection does not need to have a PPTP driver; the client simply makes a PPP connection to the modem pool or communications server. The communication server or modem pool must implement PPTP for communication with the remote access server.

# Enhanced security through VPN

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enhanced security through VPN**

Virtual private network (VPN) connections that are enabled with Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) are authenticated by using Point-to-Point Protocol (PPP) user-level authentication methods. For PPTP connections, you can use only the Microsoft Challenge Authentication Protocol (MS-CHAP), MS-CHAP version 2 (MS-CHAP v2) or Extensible Authentication Protocol-Transport Level Security (EAP-TLS). For L2TP connections, you can use any PPP authentication protocol because the PPP authentication message exchange is encrypted with Internet Protocol Security (IPSec); however, MS-CHAP, MS-CHAP v2, or EAP-TLS is recommended.

For the most secure PPTP VPN connections, do the following:

- Use EAP-TLS with user certificates or smart cards.
- Require the use of the strongest possible encryption.
- Use a remote access policy to define who can make a PPTP connection based on group membership (if you are using the Routing and Remote Access service).
- Configure PPTP packet filters on your firewall interfaces and VPN server interfaces.

For security reasons, you should use PPTP with either MSCHAPv2 or EAP-TLS with PPTP connections. If using EAP-TLS is not feasible, use MSCHAP v2 for secure connections.

For the most secure L2TP VPN connections, do the following:

- Use EAP-TLS with user certificates or smart cards.
- Require the use of the strongest possible encryption.
- Use a remote access policy to define who can make a L2TP connection based on group membership (if you are using the Routing and Remote Access service).
- Configure IPSec or L2TP packet filters on your firewall interfaces and VPN server interfaces.

In some environments, data is so sensitive that it needs to be hidden from the majority of corporate users. Finance data or human resources data are examples of this type of data. Corporations can store extremely sensitive data servers on a separate network segment that is connected to the corporate network by a VPN server. Authorized users use a VPN connection to connect to the VPN server, and then they can access the protected resources. All communication across the VPN connection is encrypted in order to ensure data confidentiality. Users who do not have authorization to establish a VPN connection with the VPN server cannot see the hidden server or its resources.

For information about creating a VPN connection, see Make a virtual private network (VPN) connection.

# Network protocol support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network protocol support**
Because virtual private networks (VPNs) support the most common network protocols, it is easy for clients on Ethernet and TCP/IP networks to use VPNs. Any network protocol supported by remote access is inherently supported in a VPN. This means that you can remotely run applications that are dependent upon particular network protocols. In turn, this translates to reduced costs in setting up and maintaining VPN connections.

For information about creating a VPN connection, see Make a virtual private network (VPN) connection.

**Notes**

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.
- IPX cannot be configured for incoming or dial-up connections in the Windows Server 2003 family.
- IPX cannot be configured on a computer running a member of the Windows Server 2003 family that is being used as a VPN server.

# IP address security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP address security**
If you are using a private IP address range (for example, 10.0.0.0-8, 172.16.0.0-12, 192.168.0.0-16), you can find a path through the Internet by providing valid, public IP addresses for the virtual private network (VPN) client and the VPN server. The original IP packets, those with the private addresses of your source and destination, are encapsulated using Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP) during data transfer over non-secure networks. Because the IP packets are also encrypted, the private addresses are protected. Internet users see only the public IP addresses of the VPN client and the VPN server. In this case, and whenever you use private IP addresses, there are no administrative costs associated with changing IP addresses for remote access.

**Note**

- It is possible to use a nonregistered IP address range (a range of private IP addresses that overlaps with the valid, public IP addresses that are used on the Internet), but it is not recommended.

For information about creating a VPN connection, see Make a virtual private network (VPN) connection.

# VPN administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN administration**

By using access policies in the remote access server infrastructure, your system administrator can configure virtual private network (VPN) user settings that significantly tighten security measures, such as required levels of data and password encryption and authentication. These requirements can be applied to individual users, or to a group (or profile) of similar users. For example, a system administrator can configure remote access dial-in settings to define a group profile such that all users who are assigned to that profile are required to authenticate by using the Extensible Authentication Protocol (EAP) and encrypt their data by using strong (128-bit) data encryption. By assigning a user to the group profile, these security measures are automatically required when anyone in that group profile connects to the remote access server.

You should configure your VPN connection to match the policies that your network administrator established. For information about configuring connections, see Configure a connection to a remote network.

# Using Home and Small Office Networking Features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using home and small office networking features**

This section covers:

- Protecting your home or small office network using Internet Connection Firewall
- Connecting to the Internet in a home or small office network

# Protecting your home or small office network using Internet Connection Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Protecting your home or small office network using Internet Connection Firewall**

Home and small office networking offers computers that are directly connected to the Internet the added security of firewall protection. Internet Connection Firewall (ICF) checks all communications that cross the connection between your network and the Internet and is selective about which responses from the Internet it allows.

**Recommended topologies**

You can enable Internet Connection Sharing (ICS) on only one Internet connection on your network, and you should protect this connection by enabling ICF. ICF can check only the communications that cross the Internet connection on which it is enabled. These types of network topologies are safe and the most recommended.




**Topologies that are not recommended**

You should avoid topologies with multiple Internet connections. If you must have multiple direct Internet connections on your network, you should ensure that ICF is enabled on each direct Internet connection in order to protect your network, as shown in the following illustration. However, because ICF works on a per-connection basis, this topology is still not a recommended topology because there is no central point of administration through which you can ensure the continuous protection of all Internet connections.




Likewise, providing Internet connectivity to your network by connecting your network hub directly to the Internet causes similar vulnerabilities and is not a recommended topology.




Enabling ICF on this type of network topology disrupts some network communications and provides protection only for the computer on which it is enabled. The other computers have direct connections to the Internet through the hub and are not protected.

**Connection icons**

When ICF has been enabled on a network connection, the following network connection icon appears in Network Connections: . To check to see if ICF is enabled, or to enable ICF, see Enable or disable ICF. When ICS is enabled on a network connection, it appears in Network Connections as the shared Internet connection icon. To check to see if ICS is enabled, or to enable ICS, see Enable ICS.

**Recommendations for using ICF**
You should enable ICF on every Internet connection. You should not enable ICF on virtual private network (VPN) connections or on network client computers because ICF interferes with some types of network activities such as file and printer sharing. For a similar reason, ICS prohibits you from enabling ICF on the ICS host computer's private connection, the connection that connects the ICS host to the ICS client computers. A firewall on this location would completely obstruct network communications. You should not enable ICF on any connection that does not directly connect to the Internet. ICF is not needed if an Internet connection already has a firewall or proxy server running.

**ICF and notification messages**
Because ICF inspects all incoming communications, some programs, especially e-mail programs, may behave differently when ICF is enabled. Some e-mail programs periodically poll their e-mail server for new mail, and some e-mail programs wait for notification from their e-mail server. Outlook Express, for example, automatically checks for new e-mail when its timer tells it to do so. When new e-mail is present, Outlook Express prompts the user with a notification about the new e-mail. ICF does not affect the behavior of this program, because the request for new e-mail notification originates from inside the firewall. The firewall makes an entry in a table noting the outbound communication to the mail server. When the mail server returns the response for new e-mail, the firewall finds an associated entry in the table and allows the communication to pass, then the user receives notification that a new e-mail has arrived.

Outlook 2000, however, is connected to a Microsoft Exchange server that uses a remote procedure call (RPC) to send new e-mail notifications to client computers. Outlook 2000 does not automatically check for new e-mail when it is connected to an Exchange server. The Exchange server notifies Outlook 2000 when new e-mail arrives. Because the RPC notification is initiated from the Exchange server that is outside the firewall, and not by Outlook 2000, which is inside the firewall, ICF cannot find the corresponding entry in the table, and the RPC messages are not allowed to cross from the Internet into the home or small office network. The RPC notification message is dropped. Users can send and receive e-mail but need to manually check the inbox for new e-mail.

**Advanced ICF settings**
The ICF security logging feature provides a way to create a security log of firewall activity. ICF is capable of logging both traffic that is permitted and traffic that is rejected. For example, incoming echo requests from the Internet, by default, are not allowed by the firewall. If the Internet Control Message Protocol (ICMP) **Allow incoming echo request** is not enabled, then the inbound request fails, and a log entry that records the failed inbound attempt is generated.

ICMP allows you to modify the behavior of the firewall by enabling various ICMP options, such as **Allow incoming echo request**, **Allow incoming timestamp request**, **Allow incoming router request** and **Allow redirect**. Brief descriptions of these options are provided on the **ICMP** tab. For information about ICMP, see Internet Control Message Protocol (ICMP).

If ICF is enabled on two or more connections on a single computer, the settings for the ICMP options are per connection. If you make a setting or change a setting in the ICMP options on any connection on which ICF is enabled, that setting is not applied to the other ICF firewalls on that computer. For more information, see Enable or disable Internet Control Message Protocol requests for ICF.

You can set the allowable size of the ICF security log to prevent the potential overflow that could be caused by denial-of-service attacks. Event logging is generated in the Extended Log File Format as established by the World Wide Web Consortium (W3C). For more information about ICF security logging, see Internet Connection Firewall security log.

If ICF is enabled on two or more connections on a single computer, the ICF settings are global. If you select a setting or change a setting in **Logging Options** or **Services** on any one of the connections on which ICF is enabled, that setting is applied to all of the ICF firewalls on that computer. For more information, see Internet Connection Firewall security log.

### Note

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Connecting to the Internet in a home or small office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connecting to the Internet in a home or small office network**

With the Internet Connection Sharing (ICS) feature of Network Connections, you can connect your home or small office network to the Internet. For example, you might have a home network that connects to the Internet by using a dial-up connection. By enabling ICS on the computer that uses the dial-up connection, you provide network address translation (NAT), addressing (DHCP Allocator), and name resolution (DNS Proxy) services for all of the computers on your network.

For more information, see Enable ICS, Internet Connection Sharing and network address translation, and Incoming connections.

After ICS is enabled, and users verify that they are all joined to the same networking workgroup as the ICS host, home or small office network users can use programs such as Internet Explorer and Outlook Express as if they were directly connected to the Internet service provider (ISP). The ICS host computer connects to the ISP and creates the connection so that the user can reach the specified Web address or resource.

ICS is intended for use in a home or small office where network configuration and the Internet connection are managed by the computer where the shared connection resides. It is assumed that this computer is the only Internet connection on the network, the only gateway to the Internet, and that ICS sets up all internal network addresses.

You might need to configure services on the ICS host computer to work properly across the Internet. The Web services that you provide must be configured so that Internet users can access them. For example, if you are hosting a Web server on your home network, and you want Internet users to be able to connect to it, you must configure the ICS host computer to permit traffic from the Web on to your network. For information, see Manage service definitions for ICF or ICS.

ICS requires two connections in order to work: one public and one private. The private connection, typically a LAN adapter, connects the ICS host computer to the computers on your home or small office network. The public connection, typically a DSL, cable, or dial-up modem, connects your network to the Internet.

Enable ICS on the public connection of your home or small office network, and before doing so, ensure that the network connection that ICS will use as the private connection connects only to your home or small office network. If you have only one connection to your network, that connection is automatically selected as the private side of your shared Internet connection when you enable ICS. If you have two or more connections to your network, you must do one of the following in order to enable ICS:

- Select one connection to be the private side of your shared Internet connection. For information about how to select the private connection, see Enable ICS.

- Set up Network Bridge to include all of the private connections to your network, so that Network Bridge serves as the private side of your shared Internet connection. If you set up Network Bridge to include all of the connections to your network, Network Bridge is automatically selected as your private connection when you enable ICS. If you do not set up Network Bridge to include all of the connections to your network, you can still select Network Bridge as your private connection.




For information about how to select the private connection, see Enable ICS. For more information about how to bridge connections, see Set up Network Bridge.

**Adding ICS client computers to a home or small office network**

When ICS is enabled on the Internet connection of a computer running Windows XP; Windows Server 2003, Standard Edition;

or Windows Server 2003, Enterprise Edition, you can add a client running Windows 98; Windows 98 Second Edition; Windows Millennium Edition; Windows XP; Windows Server 2003, Standard Edition; or Windows Server 2003, Enterprise Edition to your home or small office network in the following ways:

- Add a computer running Windows Server 2003, Standard Edition or Windows Server 2003, Enterprise Edition to your home or small office network by adding it to the home or small office network workgroup. For more information, see Join a workgroup.

- Add a computer running Windows XP to your home or small office network by running the Windows XP Network Setup Wizard on it.

- Add a computer running Windows 98, Windows 98 Second Edition, or Windows Millennium Edition to your home or small office network by running a Windows Server 2003, Standard Edition or Windows Server 2003, Enterprise Edition installation CD on it, and then running the setup program Netsetup.exe. For more information, see Configure client computers for ICS.

You can also add a computer running Windows 98, Windows 98 Second Edition, or Windows Millennium Edition to your home or small office network if you have a Windows XP installation CD. Run Network Setup Wizard on the computer running Windows 98, Windows 98 Second Edition, or Windows Millennium Edition. The wizard provides an opportunity to create a setup disk that can be run on computers running Windows 98, Windows 98 Second Edition, or Windows Millennium Edition.

Network Setup Wizard cannot be run on computers running Windows 2000. You can provide ICS addressing and name resolution services to a computer running Windows 2000 by adding it to your home and small office network workgroup.

**Notes**

- When you enable ICS, the adapter connected to the home or small office network is given a new static IP address configuration. Consequently, when you enable ICS, TCP/IP connections that are established between any small office or home office computer and the ICS host computer are lost and need to be reestablished.

- If your home or small office users need to access a corporate network that is connected to the Internet by a tunnel server from a home or small office network on which ICS is enabled, they need to create a virtual private network (VPN) connection to tunnel from the computer on the home or small office network on which ICS is enabled to the corporate tunnel server on the Internet. The VPN connection is authenticated and secure, and creating the tunneled connection allocates proper IP addresses, DNS server addresses, and WINS server addresses for the corporate network. For more information about creating a VPN connection, see Make a virtual private network (VPN) connection.

**Items configured by ICS**

When you enable ICS, certain protocols, services, interfaces, and routes are configured automatically. One important configuration is that the ICS host computer automatically becomes a Dynamic Host Configuration Protocol (DHCP) allocator for the home or small office network. DHCP distributes IP addresses to users when they start up.

**Caution**

- If you enable ICS on any of your private network adapters, the home or small office network DHCP allocator might grant IP addresses to users outside of your network, causing problems on their networks. Additionally, if you enable ICS on a private network adapter on a computer that is directly attached to the Internet, you might expose your network to computers outside your network, potentially allowing external attacks.

- Transmissions may be disrupted if you enable ICS on an existing network with other domain controllers running Windows Server 2003, Standard Edition, DNS servers, gateways, DHCP servers, or systems configured for static IP addresses. If you are running Windows Server 2003, Standard Edition and one or more of these components exist, you must use NAT to achieve the same results and to protect your network against data loss. For more information, see Deploying network address translation.

The following table describes each of the protocols, services, interfaces, and routes are configured automatically when you enable ICS.

| Configured item | Action |
| --- | --- |
| IP address 192.168.0.1 | Configured with a subnet mask of 255.255.255.0 on the LAN adapter that is connected to the home or small office network |
| Autodial feature | Enabled |
| Static default IP route | Created when the dial-up connection is established |
| ICS service | Started |
| DHCP allocator | Enabled with the default range of 192.168.0.0 and a subnet mask of 255.255.255.0. Unique addresses in the range of 192.168.0.2 to 192.168.0.254 are allocated to private network clients |
| DNS proxy | Enabled |

**Notes**

- You cannot modify the default configuration of ICS. This means you cannot configure items such as disabling the DHCP allocator or modifying the range of private IP addresses that are distributed, disabling the DNS proxy, configuring a range of public IP addresses, or configuring inbound mappings. If you want to modify any of these items, you must use NAT.

- For information about how to enable ICS, see Enable ICS.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Administering Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Network Connections**

This section covers:

- Managing network connections from the command line
- Logging remote access connections
- Configuring Network and Dial-up Connections

# Managing network connections from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing network connections from the command line**

You can display or change settings for the Server service or the Workstation service while the service is running with Net config. **Net config** allows you to display the configurable services that are running, or to display and change settings for a service. Additional Net services commands are available.

Netsh overview allows you to, either locally or remotely, display or modify the network configuration of a computer that is currently running. Netsh also provides a scripting feature that allows you to run a group of commands in batch mode against a specified computer. Netsh can also save a configuration script in a text file for archival purposes or to help you configure other servers.

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

# Logging remote access connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logging remote access connections**

You can enable logging to automatically record configuration and activity details of remote access connections. If you enable logging, you can review connection errors and pertinent connection information in a log file. This option is especially helpful when you need to determine why a remote access connection failed. If a connection fails, you can generate a diagnostic report based on the log file and send the report to a support engineer.

If you enable logging, you can view a diagnostic report, save the report to a file, or send the report to an e-mail address that you specify. You can also clear previous logs so that the report contains details of the most recent connection only. If you do not need detailed connection information but want to run the report more quickly, you can generate a summary report that contains configuration and installation information only. You can use this information to troubleshoot configurations.

If you do not enable logging and a connection fails, you are prompted to enable logging. After you enable logging, you can dial the remote access connection again to log errors for the failed connection.

For more information, see Enable logging.

# Configuring Network and Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring network and dial-up connections**
This section covers:

- Configuring a network or dial-up connection
- Optional networking components
- Client for Microsoft Networks
- File and Printer Sharing for Microsoft Networks
- Mobile users and certificates
- Configuring Dial-up Connections
- Using Terminal and Scripting
- Configuring Authentication and Data Encryption
- Automating a Network or Dial-up Connection

# Configuring a network or dial-up connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring a network or dial-up connection**

You can change a network or dial-up connection by modifying its settings. Dial-up settings, such as the phone number of the connection, the number of redial attempts, and so on, are defined for each connection. These pre-connection and post-connection settings do not modify or affect the settings of other connections. For example, you may have a dial-up connection that makes 10 attempts to redial a server that is frequently busy. You may have a second dial-up connection that makes 3 attempts to redial another server that is more readily available. The redial setting of the first connection never causes the second connection to redial more than its own established attempts of 3. There is no need to make changes to any settings outside of the settings of the dial-up connection.

Changes to some networking settings of a connection can affect other connections. For example, if you add the IPX protocol to the list of available network components for one connection, the IPX protocol also is available to other connections on the same computer.

You can modify connection settings when you are connected. You can also rename a connection when it is connected. However, the connection may need to be reconnected in order to save the changes. In a dial-up connection, modifying settings displays the message **Since this connection is currently active, some settings will not take affect until the next time you dial it**. In a local area connection, changed settings are stored automatically.

Each connection is configured with general settings that represent the basic information needed to connect successfully. These options are found on the **General** tab. For example, a local area connection only needs to list its network adapter. For a dial-up connection, the **General** tab includes required information, such as the device used to dial, the area code, the phone number, and the country/region code to dial.

For outgoing connections, you can set additional configuration options on the **Options**, **Security**, **Networking**, and **Advanced** tabs. You can set additional configuration options for an incoming connection on the **Users** and **Networking** tabs. Local area connections have additional configuration options in the **Authentication** and **Advanced** tabs.

You can configure advanced settings from **Advanced Settings** on the **Advanced** menu of the Network Connections folder. By modifying these settings, you can improve performance. For example, network providers and protocols are accessed in the order specified in **Advanced Settings**. If your LAN connection is enabled to access NetWare and Microsoft Windows networks that use TCP/IP and IPX, but your primary connection is to a Microsoft Windows network that uses TCP/IP, you can move **Microsoft Windows Network** to the top of the **Network Providers** list on the **Provider Order** tab, and move **Internet Protocol (TCP/IP)** to the top of the **File and Printer Sharing for Microsoft Networks** binding on the **Adapters and Bindings** tab. By changing your provider order, and by changing the order of protocols bound to those providers, you can improve performance. For more information about modifying the order of network providers, see Modify the network provider order. For more information about modifying the protocol binding order, see Modify the protocol bindings order.

The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

For more information about configuring connections, see Configure a connection to a remote network.

# Optional networking components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Optional networking components**

You can add or remove optional networking components. Optional networking components support network operations and are not automatically installed. Some of these components include the Route Listening Service, Simple TCP/IP Services, SNMP Services, and Print Services for UNIX.

For information about installing optional networking components, see Add optional networking components.

# Client for Microsoft Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Client for Microsoft Networks**
The Client for Microsoft Networks component allows a computer to access resources on a Microsoft network.

This component is installed and enabled by default for all VPN connections. However, this component needs to be enabled for PPPoE connections. You should not need to configure the component unless you are running software from The Open Group, such as Distributed Computing Environment (DCE) server or client kits. If you are, contact your system administrator for the correct DCE Cell Directory Service network address.

The Client for Microsoft Networks component is the equivalent of the Workstation service in Windows NT 4.0.

For information about configuring the Client for Microsoft Networks component, see Configure Client for Microsoft Networks.

# File and Printer Sharing for Microsoft Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File and Printer Sharing for Microsoft Networks**
The File and Printer Sharing for Microsoft Networks component allows other computers on a network to access resources on your computer by using a Microsoft network.

This component is installed and enabled by default for all VPN connections. However, this component needs to be enabled for PPPoE and dial-up connections. It is enabled per connection and is necessary to share local folders.

The File and Printer Sharing for Microsoft Networks component is the equivalent of the Server service in Windows NT 4.0.

For information about configuring the File and Printer Sharing for Microsoft Networks component, see Configure File and Printer Sharing for Microsoft Networks. For information about sharing folders, see Share a folder or drive.

**Note**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

# Mobile users and certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Mobile users and certificates**

If you are a mobile user, you need to enable the use of certificates on your computer. Unless your system administrator preconfigures your computer with machine and user certificates before you receive it, you must connect to your corporate network by using conventional, password-based authentication methods to get your machine and user certificates. At the time you connect, you join your computer to the corporate domain, obtain certificates, and set certificates policy. The next time you connect to the corporate network, you can use certificate-based authentication methods such as EAP.

To enable the use of certificates on a computer, do the following:

- Connect to the corporate network by using a dial-up or PPTP network connection, and authentication protocols such as MS-CHAP, or MS-CHAP v2. When you connect, your computer name joins the corporate domain and receives machine certificates.

- Request a user certificate. For more information about requesting certificates, see Request a certificate, Submit a user certificate request via the Web, and Requesting certificates.

- Create another connection that uses certificate-based authentication, and then connect again by using certificate-based authentication methods such as EAP or IPSec. For more information about enabling a connection to use a smart card or other certificate, see Enable smart card or other certificate authentication.

You can avoid these steps by having your system administrator load machine and user certificates before you take your laptop out onto the road.

**Note**

- For machine certificates to be received automatically, you must implement public key policies (by using Group Policy) that autoenroll computers to receive certificates. If the machine certificate process is not automated, you must log on to your computer with administrative credentials, install Certificate Services, and request a computer certificate. For more information, see Manage certificates for a computer and Manage certificates for your user account.

# Configuring Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring dial-up connections**
This section covers:

- Configuring dial-up options
- Configuring redial options
- Configuring multiple device dialing

# Configuring dial-up options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring dial-up options**

There are many dialing features that you can use to streamline the dial-up connection process. For instructions about configuring a dialing feature, click the feature in the following table.

| Dialing feature | Function |
|---|---|
| Configure client callback options | Reduces phone charges by having the remote access server dial back to your location after you initially connect. |
| Configure autodial | Remembers network addresses to connection destinations, which allows them to be automatically dialed when referenced from an application or from the command prompt. |
| Enable multiple device dialing | Combines multiple physical links, which increases the aggregate bandwidth of your connection, and dynamically dials and drops multilinked lines as needed. |
| Configure phone number modifiers | Manually inserts characters to force dialing behaviors (pause, suspend dialing until input provided, and so on). |
| Configure redial options | Specifies the number of times your connection automatically redials when you fail to connect to your remote access server. You can also set this feature to reconnect automatically when an inactive link is dropped. |
| Assign multiple phone numbers to a connection | Assigns more than one destination phone number to your dial-up connection. This is useful if you have a pool of phone numbers to connect to. |
| Enable operator-assisted calls or manual dialing | Overrides automatic dial-up situations where intervention is required. For example, when you call through a manually operated switchboard. |

**Note**

- You can configure dial-up options that apply to all connections within a location (for example, dialing 9 for an outside line, disabling call waiting, using pulse dialing instead of tone dialing, and so on) in Phone and Modem Options. For more information, see Create a new location.

# Configuring redial options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring redial options**

With redial options, you can determine the number of times your connection automatically redials when you fail to connect to a remote access server. You can also reconnect automatically when a link is dropped.

By default, a dial-up connection is set up to redial 3 times (at 1 minute intervals) if it cannot reach its remote access server. Also, by default the **Idle time before hanging up** is set to **Never**.

For information about configuring redial options, see Configure redial options.

# Configuring multiple device dialing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring multiple device dialing**

The Multilink Control Protocol allows dialing over multiple ISDN, X.25, or modem lines. The feature combines multiple physical links into a logical bundle and the resulting aggregate link increases your connection bandwidth. To dial multiple devices, both your connection and your remote access server must have Multilink enabled.

Network Connections can dynamically control the use of multilinked lines. By allocating lines only as they are required, thereby eliminating excess bandwidth, you can realize a significant efficiency advantage. You can configure the conditions under which extra lines are dialed, and underused lines are disconnected, through Network Connections settings.

For more information, see Enable multiple device dialing.

**Note**

- If you use Multilink to dial a server that requires callback, then only one of your multilinked devices is called back. This is because you can only store one number in a user account. Therefore, only one device connects and all other devices fail to complete the connection, and your connection loses Multilink functionality. You can avoid this problem:
    - If the multilinked phonebook entry uses a standard modem configuration, and the remote access server that your connection is calling uses more than one line for the same number.
    - If the multilinked phonebook entry is ISDN with two channels that have the same phone number.

# Using Terminal and Scripting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Terminal and scripting**

If you are connecting to a remote access server that is not running Windows, you may need to provide logon information in a separate logon Terminal window. Your connection must be configured to accommodate this requirement.

You can use one of two possible scripting methods to automate the remote Terminal logon process: Switch.inf scripts or Windows scripting. Switch.inf is an older scripting method that was developed for earlier versions of Windows NT. You may find the Windows 95 scripting language easier to use than Switch.inf.

This section covers:

- Using the Terminal feature for the remote logon process
- Automating the Dial-up Logon Process by Using Switch.inf Scripts

For information on automating the dial-up logon process by using Windows scripting, see "Using Windows Scripting with Network Connections" at the Microsoft Windows Resource Kits Web site.

# Using the Terminal feature for the remote logon process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the Terminal feature for the remote logon process**

For a PPP or SLIP server, if the remote computer you connect to requires that you log on with a terminal-emulation screen, you must configure the script settings for that network connection to use a Terminal logon window. With this logon process, after your connection connects to the remote system, a character-based window displays the logon sequence from the remote computer. You can use this window for logging on to the remote computer. For more information, see Use the Terminal feature to log on to a remote computer. Alternatively, you can automate this manual logon process as described in Automating the Dial-up Logon Process by Using Switch.inf Scripts.

Some commercial networks present a large menu of available services before you log on. On SLIP servers, typically running on UNIX, you may go through an extensive sequence of commands that updates files, collects data about you, or configures your SLIP connection during your logon process. On a newer PPP server, you may be prompted for only your user name and password before you are given a connection.

**Note**

- If the remote computer is a remote access server running Windows, you do not need to use a Terminal logon window. Instead, the logon process is completely automated for you.

# Automating the Dial-up Logon Process by Using Switch.inf Scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automating the dial-up logon process by using Switch.inf scripts**
To automate the logon process, you can use the Switch.inf file (or Pad.inf on X.25 networks) instead of the manual Terminal window described in Using the Terminal feature for the remote logon process.

Automated scripts are especially useful when a constant connection to a remote computer is needed. If your network connection is configured to use a script, and a remote connection fails, your connection automatically redials the number and reestablishes the connection. Scripts also save time if you frequently log on to a remote system and do not want to manually log on each time.

The Switch.inf file provides a generic script that you can use with little or no modification. You can try Switch.inf without modification and if it does not work, then copy and modify the generic script to match the logon sequence of the remote computer to which you want to connect.

The script language is also designed to communicate with other devices, including modems. If you are unfamiliar with modem scripts, scripting can be difficult to understand. You may find it easier to copy, then modify, one of the generic sample scripts.

This section covers:

- Using device command (.inf) files
- Creating Switch.inf Scripts for Network and Dial-up Connections
- Stepping through an example Switch.inf script
- Getting through two-second gaps and large blocks of text in Switch.inf
- Using commands and carriage returns in Switch.inf
- Troubleshooting logon scripts

# Using device command (.inf) files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using device command (.inf) files**
Network Connections uses the Switch.inf and Pad.inf files to automate the interactive logon process with servers and to initialize X.25 packet assemblers/disassemblers (PADs). These device command (.inf) files, located in the *systemroot*\System32\Ras folder, use the same commands and syntax to communicate with devices or computers automatically.

**Switch.inf**
You can use the commands in this file when you connect to the network through a security host or a remote server that requires you to log on interactively (for example, a UNIX server). You can create a script by typing the name of the device in brackets and then adding one or more commands following the device name. After you create the script, activate the script (on the **Security** tab) for each connection that will use the script. You can use the scripts before you dial (for example, modem pools) or after you dial (for example, using security hosts or actually logging on to a remote server). For more information, see Automating the Dial-up Logon Process by Using Switch.inf Scripts.

**Pad.inf**
You can use the commands in this file when you connect to the network through an X.25 PAD. Most supported PADs have scripts that are provided in Pad.inf. If necessary, you can create a script by typing the name of the device within brackets and then adding one or more commands following the device name. After you create the script, activate the script (on the **Security** tab) for each connection that will use the script.

### Caution

- You should not edit Switch.inf or Pad.inf without a basic understanding of system configuration files.

# Creating Switch.inf Scripts for Network and Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating Switch.inf scripts for network and dial-up connections**
The Switch.inf file, located in the *systemroot*\System32\Ras folder, is like a set of small batch files (scripts) contained in one file. The Switch.inf file contains a different script for each intermediary device or online service that a connection calls.

A Switch.inf script has six elements:

- Section headers
- Comment lines
- Commands
- Responses
- Response keywords
- Macros

# Section headers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Section headers**

Section headers divide the Switch.inf file into individual scripts. A section header marks the beginning of a script for a particular remote computer and must not exceed 31 characters. The text of a section header appears in Network Connections when you activate the script. The section header is enclosed in square brackets. For example:

```
[Route 66 Logon Process]
```

For information about activating a Switch.inf script, see Activate a logon script. For more information about Switch.inf scripts, see Automating the Dial-up Logon Process by Using Switch.inf Scripts.

# Comment lines

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Comment lines**

Comment lines must have a semicolon (;) in column one and can appear anywhere in the file. Comment lines contain information for people who maintain the Switch.inf file. For example:

```
; This script was created by Martha on September 29, 1996.
```

For more information about activating a Switch.inf script, see Activate a logon script. For more information about Switch.inf scripts, see Automating the Dial-up Logon Process by Using Switch.inf Scripts.

# Commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Commands**

Each line in a script is a command from your local computer to a remote computer or a response from a remote computer to your local computer. Each command or response is a stream of data or text. For example, the following command sends a user name (Martha) and a carriage return (the macro <cr>) to the remote computer:

```
COMMAND=Martha<cr>
```

The commands and responses must be in the exact order the remote computer expects them. Branching statements, such as GOTO or IF, are not supported.

The required sequence of commands and responses for a specific remote device should be in the documentation for that device or, if you are connecting to a commercial service, from the support staff of that service. If the exact sequence is not available, you can activate the generic script provided with Network Connections and modify it to match the logon sequence of the remote computer as described in Troubleshooting logon scripts.

You can use the COMMAND= statement in two additional ways:

- `COMMAND= NoResponse`

  This is the default behavior and causes an approximate two-second delay. This is useful when an intermediate device requires a delay.

- `COMMAND=` *string*

  In this case, *string* is not followed by a carriage return (<cr>). This is useful when a device requires slow input. Instead of receiving the whole command string, the device requires characters to be sent one-by-one.

The following example shows a script for an intermediary device that is so slow, it is able to receive and process only one character of the command PPP at a time:

```
COMMAND=P
NoResponse

COMMAND=P
NoResponse

COMMAND=P
NoResponse
```

**Note**

- To activate a script.inf file, see Activate a logon script.

# Responses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Responses**

A response is sent from the remote device or computer. To write an automatic script, you must know the responses you will receive from the remote device. If a gap of two or more seconds occurs between characters, the received text is sent as a response. This gap is the only cue that a response is over. For more information, see Getting through two-second gaps and large blocks of text in Switch.inf.

For information about activating a Switch.inf script, see Activate a logon script. For more information about Switch.inf scripts, see Automating the Dial-up Logon Process by Using Switch.inf Scripts.

# Response keywords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Response keywords**

The keyword in a response line specifies what to do with the responses you receive from the remote computer, as follows:

- `OK=` *remote computer response<macro>*

  The script continues to the next line if the response or macro is encountered.

- `LOOP=` *remote computer response<macro>*

  The script returns to the previous line if the response or macro is encountered.

- CONNECT=remote computer response<macro>

  Used at the end of a successful modem script. Not generally useful with the Switch.inf file.

- `ERROR=` *remote computer response<macro>*

  Causes Network Connections to display a generic error message if the response is encountered. Useful for notifying you when the remote computer reports a specific error.

- `ERROR_DIAGNOSTICS=` *remote computer response<diagnostics>*

  Causes Network Connections to display the specific cause for an error returned by the device. Not all devices report specific errors. Use **ERROR=** if your device does not return specific errors that can be identified with Network Connections diagnostics.

- `NoResponse`

  Used when no response comes from the remote device.

A connection always expects a response from the remote device and waits until a response is received unless a NoResponse statement follows the COMMAND= line. If there is no statement for a response following a COMMAND= line, the COMMAND= line executes and stops the script at that point.

For information about activating a Switch.inf script, see Activate a logon script. For more information about Switch.inf scripts, see Automating the Dial-up Logon Process by Using Switch.inf Scripts.

# Macros

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Macros**

Macros are enclosed in angle brackets (<>) and perform a variety of special functions, as follows:

- `<cr>`

  Inserts a carriage return.

- `<lf>`

  Inserts a line feed.

- `<match>` " *string* "

  Reports a match if the string enclosed in quotation marks is found in the device response. Each character in the string is matched according to upper and lower case. For example, <match> "Smith" matches Denise Smith and Jeff Smith, but not SMITH.

- `<?>`

  Inserts a wildcard character, for example, CO<?><?>2 matches COOL2 or COAT2, but not COOL3.

- `<hXX>`

  (XX are hexadecimal digits)

  Allows any hexadecimal character to appear in a string--including the zero byte, <h00>.

- `<ignore>`

  Ignores the rest of a response from the macro on.

- `<diagnostics>`

  Passes specific error information from a device to remote access, which displays the specific error to you. Otherwise, a nonspecific error message appears.

For information about activating a Switch.inf script, see Activate a logon script. For more information about Switch.inf scripts, see Automating the Dial-up Logon Process by Using Switch.inf Scripts.

# Stepping through an example Switch.inf script

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Stepping through an example Switch.inf script**

This section describes each part of the generic script provided in the Switch.inf file.

Every script must start with a command to the remote computer, followed by one or more response lines. This initial command might be simply to wait for the remote computer to initialize and send its logon banner. The default initial command is to wait two seconds for the logon banner. It appears as follows in the Switch.inf file:

```
COMMAND=
```

If the response (the logon banner from the remote computer) is the following:

```
Welcome to A. Datum Corporation. Please enter your login:
```

then the corresponding response line in the Switch.inf file should be:

```
OK=<match>"Please enter your login:"
```

If the remote computer sends the "Please enter your login:" string, then this indicates that everything is correct. You respond by sending a command with the characters in your user name and the carriage return, as follows:

```
COMMAND=Martha<cr>
```

If the response from the remote computer is the following:

```
Please enter your password:
```

then the corresponding response line in the Switch.inf file should be:

```
OK=<match>"Please enter your password:"
```

To send your password, you send the following command:

```
COMMAND=mUs3naB<cr>
```

On many PPP computers, this script automatically logs you on.

For information about activating a Switch.inf script, see Activate a logon script. For more information about Switch.inf scripts, see Automating the Dial-up Logon Process by Using Switch.inf Scripts.

# Getting through two-second gaps and large blocks of text in Switch.inf

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Getting through two-second gaps and large blocks of text in Switch.inf**

If the computer you are connecting to has a two-second gap in the data stream response to your computer, the remote access server assumes that the gap is the end of the response. These gaps can occur anywhere, even between words, and are detected by using modem diagnostics. For more information, see Log and view modem commands.

If you write a script that fails for no apparent reason, you can use modem diagnostics, and then consult *systemroot*\ModemLog_*model*.txt to see if a response ends in the middle of a word. *Model* is the name of the modem as it appears in the list of installed modems in **Modem Properties**. If a response ends in the middle of a word, your script must account for the two-second gap. A simple way to do this is to include the following command:

```
COMMAND=<cr>
```

You can skip to the end of large blocks of text that contain multiple gaps by using the LOOP= keyword and by matching text at the end of a block. For example:

```
COMMAND=<cr>
OK=<match>"Enter the service to start:"
LOOP=<ignore>
```

In this example, the network connection sends a null command (waits two seconds). The network connection then waits for the "Enter the service to start:" message. If this is a long block of text, the connection does not find the string. The connection then moves to the **LOOP** command. The **LOOP** command causes Network Connections to return to the line above, and the connection waits for the words "Enter the service to start:" in the second response. In this manner, you can loop though long blocks of text until you reach the text of the desired prompt.

For information about activating a Switch.inf script, see Activate a logon script. For more information about Switch.inf scripts, see Automating the Dial-up Logon Process by Using Switch.inf Scripts.

# Using commands and carriage returns in Switch.inf

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using commands and carriage returns in Switch.inf**

In Switch.inf scripts, you usually must include <cr>, which indicates a carriage return, at the end of a command. The carriage return causes the remote computer to process the command immediately. If you do not include <cr>, the remote computer may not recognize the command.

In other situations, you cannot use <cr> because the remote computer accepts a command without a carriage return and requires time to process the command. This situation mainly applies when you are sending a series of commands without expecting a response.

For more information about activating a Switch.inf script, see Activate a logon script. For more information about Switch.inf scripts, see Automating the Dial-up Logon Process by Using Switch.inf Scripts.

# Troubleshooting logon scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting logon scripts**

You can log all information passed between your network connection, the modem, and the remote device, including errors reported by the remote device.

You can view modem commands by enabling logging. After you enable logging, commands sent to the modem are captured in the file *systemroot*\ModemLog_*model*.txt. *Model* is the name of the modem as it appears in the list of installed modems. For instructions about logging modem commands, see Log and view modem commands.

# Configuring Authentication and Data Encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring authentication and data encryption**
This section covers:

- Dial-up connection authentication and data encryption
- VPN connection authentication and data encryption

# Dial-up connection authentication and data encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up connection authentication and data encryption**

The **Typical (recommended settings)** security options that you select on the **Security** tab result in a predefined set of authentication methods and encryption requirements that are negotiated with the server during a PPP exchange.

The following table shows the authentication and data encryption methods that you can use with each combination of **Validate my identity as follows** and **Require data encryption (disconnect if none)** selections. You can also view these settings by making your identity validation and data encryption requirement selections in **Typical (recommended) settings**, and then clicking **Settings** in **Advanced (custom) settings**.

You may individually enable, configure, and disable these combinations of security settings by using **Advanced (custom settings)**, but this requires a knowledge of security protocols.

For more information about a specific authentication or data encryption method, click the method in the table. For information about configuring a connection, see Configure a connection to a remote network.

| Validate my identity as follows | Require data encryption | Authentication methods negotiated | Encryption enforcement |
|---|---|---|---|
| Allow unsecured password | No | Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), Shiva Password Authentication Protocol (SPAP), Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) | Optional encryption (connect even if no encryption) |
| Require secured password | No | Challenge Handshake Authentication Protocol (CHAP), Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) | Optional encryption (connect even if no encryption) |
| Require secured password | Yes | Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) | Require encryption (disconnect if server declines) |
| Smart card | No | Extensible Authentication Protocol (EAP) | Optional encryption (connect even if no encryption) |
| Smart card | Yes | Extensible Authentication Protocol (EAP) | Require encryption (disconnect if server declines) |

**Notes**

- Data is only encrypted if MS-CHAP, MS-CHAP v2, or EAP-TLS authentication is negotiated. These are the only authentication protocols that generate their own initial encryption keys, which are required for encryption.
- Microsoft Point-to-Point Encryption (MPPE) encrypts data in PPP-based dial-up connections. Strong (128-bit key) and

standard (40-bit key) MPPE encryption schemes are supported.

- MS-CHAP v2 and EAP are mutual authentication protocols, which means that both the client and the server prove their identities. If your connection is configured to use either MS-CHAP v2 or EAP as its only authentication method, and the server that you are connecting to does not provide proof of its identity, your connection disconnects. Previously, servers could skip authentication and simply accept the call. This change ensures that you can configure a connection to connect to the expected server.

# VPN connection authentication and data encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN connection authentication and data encryption**

The **Typical (recommended settings)** security options that you select on the **Security** tab result in a predefined set of authentication methods and encryption requirements that are negotiated with the server during a PPP exchange.

The following tables show the authentication and data encryption methods that you can use with each combination of **Validate my identity as follows** and **Require data encryption (disconnect if none)** selections. You can also view these settings by making your identity validation and data encryption requirement selections in **Typical (recommended settings)**, and then clicking **Settings** in **Advanced (custom settings)**.

You may individually enable, configure, and disable these combinations of security settings, by using **Advanced (details for all possible settings)**, but this requires a knowledge of security protocols.

For more information about a specific authentication or data encryption method, click the method in the table. For information about configuring a connection, see Configure a connection to a remote network.

**Point-to-Point Tunneling Protocol (PPTP) remote access server**

| Validate my identity as follows | Require data encryption | Authentication methods negotiated | Encryption enforcement |
|---|---|---|---|
| Require secured password | No | Challenge Handshake Authentication Protocol (CHAP), Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) | Optional encryption (connect even if no encryption) |
| Require secured password | Yes | Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) | Require encryption (disconnect if server declines) |
| Smart card | No | Extensible Authentication Protocol (EAP) | Optional encryption (connect even if no encryption) |
| Smart card | Yes | Extensible Authentication Protocol (EAP) | Require encryption (disconnect if server declines) |

**Layer Two Tunneling Protocol (L2TP) remote access server**

| Validate my identity as follows | Require data encryption | Authentication methods negotiated | Encryption enforcement |
|---|---|---|---|
| Require secured password | No | Challenge Handshake Authentication Protocol (CHAP), Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) | Optional encryption (connect even if no encryption) |
| Require secured password | Yes | Challenge Handshake Authentication Protocol (CHAP), Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) | Require encryption (disconnect if server declines) |
| Smart card | No | Extensible Authentication Protocol (EAP) | Optional encryption (connect even if no encryption) |

| | | | |
|---|---|---|---|
| Smart card | Yes | Extensible Authentication Protocol (EAP) | Require encryption (disconnect if server declines) |

**Notes**

- In L2TP VPN connections, data is encrypted by using Internet Protocol Security (IPSec) Encryption.
- Microsoft Point-to-Point Encryption (MPPE) encrypts data in PPTP VPN connections. Strong (128-bit key) and standard (40-bit key) MPPE encryption schemes are supported.
- Data is only encrypted by MPPE if MS-CHAP, MS-CHAP v2, or EAP-TLS authentication is negotiated. These are the only authentication protocols that generate their own initial encryption keys. MPPE requires common client and server keys as generated by these types of authentication.
- MS-CHAP v2 and EAP-TLS are mutual authentication protocols, which means that both the client and the server prove their identities. If your connection is configured to use either MS-CHAP v2 or EAP-TLS as its only authentication method, and the server that you are connecting to does not provide proof of its identity, your connection disconnects. Previously, servers could skip authenticating themselves to clients and simply accept the call. This change ensures that you can configure a connection to connect to the expected server.

# Automating a Network or Dial-up Connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automating a network or dial-up connection**
You can automate the connection process for a network or dial-up connection by using the autodial feature. With autodial enabled, network addresses are mapped to connection destinations, which allows a destination connection to be automatically dialed when a remote resource is referenced.

You can also automate the connection process for any Microsoft client by using a simple batch file and the **rasdial** command or by using a custom application that recognizes remote access.

This section covers:

- Autodial
- Using rasdial

# Autodial

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Autodial**
Autodial maps and maintains network addresses to connection destinations, which allows the destinations to be automatically dialed when referenced, whether from an application or from the command prompt. A network address can be an Internet host name, an IP address, or a NetBIOS server name. When a more specific mapping is not available in the database, autodial dials the default Internet connection.

The following example describes how autodial works:

1. You are not connected to your ISP, and you click an Internet address that is embedded in a word processing document.

2. If you have a connection that is set as the default Internet connection, autodial dials this connection, and then you can access the Internet address that you are trying to reach.

3. If you do not have a connection that is set as the default Internet connection, you are prompted to select a connection to use to reach your ISP. Autodial dials with this connection, and then you can access the Internet address that you are trying to reach.

4. The next time that you are not connected to your ISP and click the Internet address in a word processing document, a similar process occurs. If you selected a specific connection, that connection is automatically dialed. Otherwise, the default connection is dialed.

Autodial maps the resources that are required to initially make a connection but does not map resources once the connection is made.

There are two possibilities when autodial attempts to make a connection:

- If you are not connected to a network, autodial attempts to create a connection whenever an application references a remote address.

- If you are connected to a network, autodial attempts to create a connection only for those addresses that it has previously learned. Incorrectly typed server or Internet host names do *not* cause an autodial attempt.

For more information, see Configure autodial.

**Note**

- Autodial is disabled by default in all members of the Windows Server 2003 family. To enable autodial, the Remote Access Auto Connection Manager service needs to be started. For more information, see Start Remote Access Auto Connection Manager.

# Using rasdial

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using rasdial**

You can create a simple batch file that connects to a dial-up network connection by using the **rasdial** command. You can also use **rasdial** to connect to or disconnect from a network connection from a command line. Unlike autodial, the **rasdial** command does not automatically detect your network connection status before dialing a connection and is not dependent on the Remote Access Auto Connection Manager service.

For example, you might want to automatically connect to a dial-up connection during off-peak hours in order to transfer large data files. You can create a batch file or a series of batch files to automate this process. You would use the **rasdial** command to start and end the dial-up connection. There are additional configurable parameters that are associated with **Rasdial**, such as what phone number to use, the callback number, the phone book file, prefix and suffix.

For more information about using the **rasdial** command, see Rasdial. For more information about batch files, see Using batch files.

# Network Connections Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section covers:

- Troubleshooting network and dial-up connections
- Error messages in Network Connections (Windows Server 2003)

# Troubleshooting network and dial-up connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting network and dial-up connections**
What problem are you having?

- The modem does not work.
- Unable to connect to a remote access server.
- When trying to connect, a message is received that says the remote access server is not responding.
- The modem always connects at a lower bits-per-second (bps) rate than specified.
- The sessions with a remote access server on the network keep getting dropped.
- Connections are disconnecting abnormally.
- When trying to connect, a hardware error is received.
- Conflicts between serial ports are causing connection problems.
- When trying to connect by using ISDN, a "No Answer" message is received.
- Connections made by using X.25 fail.
- Connections through PPTP fail.
- Connections configured by using PPP or TCP/IP are failing.
- Connections made by using Internet Connection Sharing (ICS) are failing.
- When using a local area network connection, there is no response.
- When using a laptop to connect to an ISP, some or all of the applications do not run properly.
- Incoming connection clients cannot see resources beyond the incoming connection computer.
- The smart card reader does not work.
- The smart card does not work.
- When you try to establish a remote access connection using a smart card, you receive an error.

**The modem does not work.**
**Cause:** The modem is incompatible.

**Solution:** If you have access to another computer with an Internet connection, check the list of compatible modems at Support resources.

**Cause:** The modem is not connected properly or is turned off.

**Solution:** Verify that the modem is connected properly to the correct port on your computer. If the modem is external, verify that the power is on.

**See Also:** Modems Troubleshooting.

**Unable to connect to a remote access server.**
**Cause:** The remote access server is not running.

**Solution:** Ask your system administrator to verify that the remote access server is running. If the server is down, the administrator needs to check the error and audit logs to see why the service stopped. After the problem is fixed, restart the service. If the service is running, the administrator needs to check whether other remote access clients can connect properly. If other clients can connect, the problem may be specific to your workstation.

**Cause:** You do not have a valid user account, or you do not have remote access permission.

**Solution:** Verify with your system administrator that your user account has been created, and that you have remote access permission.

**Cause:** You dialed the wrong number, or you dialed the correct number but forgot to dial an external line-access number, such as 9.

**Solution:** Verify that the dialed number is correct.

**Cause:** Your modem cannot negotiate with the modem of the server.

**Solution:** Try using the same type of modem as the server.

**Cause:** The modem cabling is faulty.

**Solution:** Do not use the 9-to-25-pin converters that are included with most mouse hardware because some of them do not carry modem signals. To be safe, you should use a converter made especially for this purpose.

**Cause:** The telephone line (for example, in your hotel room) does not accommodate your modem speed.

**Solution:** Select a lower bits-per-second (bps) rate (or call the hotel manager to request a direct line).

**See also:** Change the maximum modem port speed.

**Cause:** The line you are trying to use is digital.

**Solution:** Most modems work only with analog phone lines. Verify that you have analog phone lines installed or, if you have digital phone lines installed, verify that the servers and clients have digital modems.

**Cause:** Your modem settings need to be changed because of a remote access server change.

**Solution:** Verify the modem settings with your system administrator.

**When trying to connect, a message is received that says the remote access server is not responding.**
**Cause:**

- At higher bits-per-second (bps) rates, your modem is incompatible with the modem of the server.
- There is a lot of static on the phone line, which prevents a modem from connecting at a higher bps rate.
- Switch equipment between the client and server that prevents the two modems from negotiating at a higher bps rate.

**Solution:** Adjust the speed of your modem to a lower bit-per-seconds (bps) rate.

**See also:** Change the maximum modem port speed.

**Cause:** The remote access server is not running.

**Solution:** Ask your system administrator to verify that the server is running.

**The modem always connects at a lower bits-per-second (bps) rate than specified.**
**Cause:** The modem and telephone line are not operating correctly. Excessive static on the telephone line causes sessions to be dropped.

**Solution:** You can use modem diagnostics to confirm correct modem operation.

**See also:** Log and view modem commands.

**Cause:** The remote access server is not functioning properly.

**Solution:** Ask your system administrator to verify that the destination server is running properly. Try connecting to the same server from another workstation. If other workstations are having the same problem, there may be problems with server applications or hardware. If not, the problem is specific to your workstation.

**Cause:** The quality of your line is insufficient.

**Solution:** Contact your telephone company to verify the quality of your line.

**Cause:** The line you are dialing is affecting the speed.

**Solution:** If you can connect to your remote access server by using more than one number, try another number and see if the speed improves.

**Cause:** Your modem software needs to be updated.

**Solution:** Check with your modem manufacturer for modem software updates.

**The sessions with a remote access server on the network keep getting dropped.**
**Cause:** Call waiting is disrupting your connection.

**Solution:** Verify that the phone has call waiting. If so, disable call waiting and try calling again.

**See also:** Change dialing properties for a location.

**Cause:** The remote access server disconnected you because of inactivity.

**Solution:** Reattempt dial-up connection again.

**Cause:** Someone picked up the phone. Picking up the phone automatically disconnects you.

**Solution:** Reattempt dial-up connection again.

**Cause:** Your modem cable is loose.

**Solution:** Verify that the modem cable is connected properly.

**Cause:** Your modem software needs to be updated.

**Solution:** Check with your modem manufacturer for modem software updates.

**Connections are disconnecting abnormally.**
**Cause:** The remote access server is not performing optimally.

**Solution:** Ask your system administrator to verify that the remote access server is running properly.

**Cause:**

- Your modem is unable to negotiate correctly with the modem of the remote access server.
- The serial port of the computer cannot keep up with the speed you have selected.

**Solution:** Try to connect at a lower initial port speed.

**See also:** Change the maximum modem port speed.

**Cause:** Your modem software needs to be updated.

**Solution:** Check with your modem manufacturer for modem software updates.

**When trying to connect, a hardware error is received.**
**Cause:** The modem is turned off.

**Solution:** Verify that the modem is turned on. If the modem is turned off, turn it on and redial.

**Cause:** Your modem is not functioning properly.

**Solution:** Enable modem logging to test the modem.

**See also:** Log and view modem commands.

**Cause:** Your cable is incompatible.

**Solution:** If your modem does not communicate through Network Connections, the cable that attaches your modem to the computer is probably incompatible. You need to install a compatible cable.

**Conflicts between serial ports are causing connection problems.**
**Cause:** The serial ports are conflicting.

**Solution:** Com1 and Com3 share interrupt request (IRQ) 4. Com2 and Com4 share IRQ 3. As a result, for serial communications, you cannot use COM1 and COM3 simultaneously, or COM2 and COM4 simultaneously. For example, you cannot use Network Connections on COM1 and Terminal on COM3.

This rule applies if you are using the mouse in addition to other serial communications programs such as Network Connections or the Terminal program. The rule does not apply if you are using an intelligent serial adapter.

**See Also:** Configure a connection to a remote network; Change the COM port for a modem.

**When trying to connect by using ISDN, a "No Answer" message is received.**
**Cause:** The remote access server did not answer because it is turned off or the modem is not connected.

**Solution:** Contact your system administrator.

**Cause:** The line is busy.

**Solution:** Try calling later, or contact your system administrator.

**Cause:** There is a problem with the ISDN hardware.

**Solution:** Verify that the ISDN adapters are installed and configured correctly.

**Cause:** Your phone number is not configured correctly.

**Solution:** In some cases, each B channel on an ISDN line has its own number, while in other cases both B channels share a single number. Your telephone company can tell you how many numbers your ISDN line has.

**Cause:** If you are located in the United States or Canada, your Service Profile Identifier (SPID) is configured incorrectly. The SPID normally consists of the phone number with additional digits added to the beginning, the end, or both. The SPID helps the switch understand what type of equipment is attached to the line and routes calls to appropriate devices on the line. If an ISDN channel requires a SPID, but it is not entered correctly, then the device cannot place or accept calls.

**Solution:** Verify that the SPID is entered correctly.

**Cause:** A poor line condition (for example, too much static) interrupted your connection.

**Solution:** Wait a few minutes and try dialing again.

**Cause:** You did not enable line-type negotiation, or a connection cannot be made with the line type you selected.

**Solution:** Enable line-type negotiation.

**Cause:** Your ISDN switching facility is busy.

**Solution:** Try again later.

**Cause:** Your DigiBoard card is too old.

**Solution:** If you do not have the latest PCIMAC-ISA DigiBoard card, serial number A14308 or greater, contact DigiBoard for a replacement.

**Connections made by using X.25 fail.**
**Cause:** The dial-up PAD is configured with the wrong X.3 parameters or serial settings.

**Solution:** If the remote access server is running and you cannot connect to it directly through an X.25 smart card or an external PAD, modify the dial-up PAD X.3 parameters or serial settings. Ask your system administrator for the correct settings.

**See also:** Configure a connection to a remote network.

**Cause:** New Pad.inf entries are incorrect.

**Solution:** You can check other Pad.inf entries for direct connections and external PADs, and view the comments that go with them. You may need a line analyzer or a terminal program to see the response for the PAD. For dial-up PAD entries, you can use an entry in Pad.inf as an example, paying attention to the comments that go with the example.

**Cause:** Your modem is incompatible.

**Solution:** If the modem that connects to a dial-up PAD connects at a lower speed than it should, replace the modem with a compatible one.

**See also:** The modem does not work with Network Connections.

**Cause:** The line for the remote access server is congested. If a connection has been established, but the network drives are disconnecting, and you are dropping sessions or getting network errors, the cause may be congestion on the leased line for the remote access server.

For example, four clients connecting at 9600 bps (through dial-up PADs) require a 38,400-bps (four times 9600) leased line on the server end. If the leased line does not have adequate bandwidth, it can cause time-outs and degrade performance for connected clients. This example assumes that Routing and Remote Access is using all the bandwidth. If Routing and Remote Access is sharing the bandwidth, fewer connections can be made.

**Solution:** Your system administrator needs to verify that the speed of the leased line can support all the COM ports at all speeds clients use to dial in.

**Connections through PPTP fail.**
**Cause:** TCP/IP connectivity problems are keeping you from connecting to the PPTP server.

**Solution:** You or your system administrator can use the **ipconfig** and **ping** commands to verify a connection to your server.

**Cause:** A Winsock Proxy client is active.

**Solution:** A VPN connection cannot operate with an active Winsock Proxy client. Winsock Proxy immediately redirects packets to the proxy server before they can be processed by a virtual network connection for encapsulation. To fix this, you need to disable the Winsock Proxy client.

**Cause:** You do not have the appropriate remote access permissions and domain rights on the remote access server.

**Solution:** Contact your system administrator.

**Cause:** If you are using the TCP/IP protocol, you do not have a unique TCP/IP address.

**Solution:** Contact your system administrator.

**Cause:** Name resolution problems are keeping you from resolving names to IP addresses.

**Solution:** Specify a valid IP address for your connection.

**See Also:** Configure TCP/IP settings.

**Connections configured by using PPP or TCP/IP are failing.**
**Cause:** The server does not support LCP extensions.

**Solution:** If you cannot connect to a server by using PPP, or the remote computer terminates your connection, the server may not support LCP extensions. In Network Connections, clear the **Enable LCP extensions** check box.

**Cause:** IP header compression is keeping TCP/IP from running. If you successfully connect to a remote server by using PPP, but TCP/IP does not work, the problem may be IP header compression.

**Solution:** Try to reconnect after turning off IP header compression.

**See Also:** Stop requesting LCP extensions in PPP; Turn off IP header compression in PPP.

**Connections made by using Internet Connection Sharing (ICS) are failing.**
**Cause:** The wrong LAN network adapter is shared.

**Solution:** A computer with ICS needs two connections. One connection, typically a LAN adapter, connects to the computers on the home or small office network and the other connection connects the home or small office network to the Internet. You need to ensure that ICS is enabled on the connection that connects your home or small office network to the Internet.

**Cause:** TCP/IP is not installed on home network computers.

**Solution:** By default, the TCP/IP protocol is installed on computers running Windows NT version 4.0, Windows 98, Windows 2000, Windows XP, and Windows Server 2003 operating systems. If users on your home or small office network are running operating systems other than these, verify that TCP/IP is installed on their computers.

**Cause:** If users on your home or small office network cannot reach the Internet, TCP/IP is incorrectly configured on their home or small office network computers.

**Solution:** Verify that the following TCP/IP settings are established on home or small office network local area connections:

- IP Address: Obtain an IP address automatically (through DHCP)
- DNS server: Obtain DNS server address automatically
- Default gateways: None specified

For computers running Windows 9*x* or Windows NT version 4.0, you can find the TCP/IP settings in Network in Control Panel.

**Cause:** The Internet Connection Sharing (ICS) service is not started.

**Solution:** Use Event Viewer to verify that the ICS service is started.

**See also:** Event Viewer.

**Cause:** The Internet Connection Sharing (ICS) host computer is not properly configured for name resolution.

**Solution:** You may need to configure the WINS or DNS name resolution services on the computer. If computers on the home or small office network cannot resolve names to IP addresses, you can check the name resolution configuration of the ICS host computer by using the **ipconfig** command. There are two ways that your ISP can configure name resolution:

- Statically assigned name servers

  You must manually configure the TCP/IP protocol with the IP address (or addresses) of the name servers provided by the ISP. If you have statically assigned name servers, you can use the **ipconfig** command at any time to get the IP addresses of your configured name servers.

- Dynamically assigned name servers

  Manual configuration is not required. The IP addresses of the name servers provided by the ISP are dynamically assigned whenever you dial the ISP. If you have dynamically assigned name servers, you must run the **ipconfig** command *after* a connection to the ISP has been made.

**Cause:** If you cannot play a game across the Internet, the protocol used by the application is not translatable.

**Solution:** Try running the application from the Internet Connection Sharing host computer. If the application works there but not at a workstation on the home or small office network, then the application may not be translatable.

**Cause:** If you cannot play a game across the Internet, the application is not configured on the Internet Connection Sharing host computer.

**Solution:** Verify that the application, including port numbers, is configured correctly.

**See also:** Manage service definitions for ICF or ICS.

**Cause:** If Internet users cannot see services on your home or small office network, such as a Web server, the service is not configured correctly.

**Solution:** Verify that the service, including port numbers and TCP/IP addresses, is configured correctly.

**Cause:** If users on your home or small office network cannot reach Internet sites by using friendly names, there is a DNS resolution problem.

**Solution:** Have users on your home or small office network use fully qualified domain names or IP addresses when accessing Internet resources.

**See Also:** Manage service definitions for ICF or ICS; Connecting to the Internet in a home or small office network; Configure TCP/IP to use WINS; Configure TCP/IP to use DNS; Test a TCP/IP configuration by using the ping command.

**Notes**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**When using a local area network connection, there is no response.**
**Cause:** There may be problems with your network adapter.

**Solution:** Try the following:

- Check the appearance of the local area connection icon. Depending on the status of the local area connection, the icon appears in different ways in the Network Connections folder. Also, if the local area connection media is disconnected (for example, the cable is unplugged), a status icon is displayed in the taskbar.

- Use Device Manager to verify that your network adapter is working correctly.

**See also:** Using local area connections

**Cause:** The LAN cable may not be plugged into the network adapter.

**Solution:** Check to make sure the LAN cable is inserted into the network adapter.

**When using a laptop to connect to an ISP, some or all of the applications do not run properly.**
**Cause:** The WinSock Proxy Client may be preventing your applications from running properly when you use the ISP connection.

**Solution:** If you are a mobile user and use your laptop in your corporate environment, you may need to disable the Microsoft WinSock Proxy Client (WSP Client in Control Panel) when you use the same computer to dial to an ISP or other network. For example, if you use a laptop in your office and use the same computer to connect to an ISP or other network from your home, you may have problems running all of your applications when you use the ISP connection. (For example, your applications may not be able to find the resources or servers they need.) If this is the case, you should disable the Microsoft WinSock Proxy Client (WSP Client in Control Panel) to run the applications that you typically run when you use your laptop in the corporate office.

**Incoming connection clients cannot see resources beyond the incoming connection computer.**
**Cause:** If the addresses that are being allocated to incoming clients are not a subset of the network to which the incoming connection computer is attached, you must create a route to the incoming client computers on the intranet computers.

**Solution:** Reconfigure your range of IP addresses that are being allocated to incoming clients so that it is a subset of the network to which the incoming connections computer is attached. If you cannot do this, then configure your intranet hosts with the IP address of the incoming connections computer as a default gateway.

- If your intranet hosts are configured to obtain an IP address automatically and a Dynamic Host Configuration Protocol (DHCP) server is present, you can configure your DHCP server to assign the default gateway.

- If your intranet hosts are configured to obtain an IP address automatically and a DHCP server is not present (you are using the Automatic Private IP Addressing feature), then you must manually configure all of your intranet hosts with an IP address, subnet mask, and default gateway.

**See Also:** To configure a default gateway, Incoming connections and IP addresses.

**Cause:** The calculated range of addresses allocated to connecting clients is larger than the range that you configured.

**Solution:** Most TCP/IP networks use subnets in order to effectively manage routed IP addresses. The closest matching subnet is calculated for the range that you specified in **From** and **To**. The range of addresses in the closest matching subnet may

exceed the range that you specified. Unless the addresses specified in **From** and **To** are subnet boundaries, the range based on the calculated subnet will be larger than the range that you specified. To avoid this, specify a range that falls on subnet boundaries. For example, if you are using the 10.0.0.0 private network ID for your intranet, a range that falls on subnet boundaries is 10.0.1.168 to 10.0.1.175. Or, if you are using the 192.168.0.0 private network ID for your intranet, a range that falls on subnet boundaries is 192.168.1.0 to 192.168.1.255.

**See Also:** Configure an incoming connection to use TCP/IP.

**The smart card reader does not work.**
**Cause:** The smart card reader is not installed.

**Solution:** Install the smart card reader and make sure that it is properly attached to the port.

**Cause:** The smart card reader is not installed properly.

**Solution:** Try the following:

- Make sure that the smart card reader appears in Device Manager.
- Ensure that you have the latest drivers installed for your smart card reader. To do so, right-click the smart card reader in Device Manager and click **Update driver**. If this does not work, check Windows Update to verify whether you are using the most recent drivers.
- If the drivers for your smart card reader are not loaded correctly or have become corrupted, you can reinstall them. To do this, right-click the smart card reader in Device Manager, and click **Uninstall**. Right-click any device and then click **Scan for hardware changes**.
- If the solutions given above do not work, uninstall the smart card reader drivers using Device Manager and then unplug the smart card reader from its port. Next, plug the smart card reader back into its appropriate port. Windows will automatically reinstall the drivers for the newly detected hardware.
- If Windows Update does not offer drivers for your smart card reader, contact the supplier for your smart card reader to verify whether you have the appropriate drivers.
- Try a different smart card reader.

**Notes**

- The above-mentioned causes can generate the following errors:
    - 0x80090022: Provider could not perform the action since the context was acquired as silent.
    - Error 0x80090019: The keyset is not defined.
    - Error 764: No smart card reader is installed.
    - An internal error has been detected, but the source is unknown.
    - Error 0x8010002e: Cannot find a smart card reader. For customized troubleshooting information for this connection, click Help.
- If you are using cryptographic service provider (CSP) software from a supplier other than Microsoft, the software might not be installed properly or might be corrupted. If this happens, you may see the following error messages:
    - 0x80090022: Provider could not perform the action since the context was acquired as silent.
    - Error 0x80090019: The keyset is not defined.
    - An internal error has been detected, but the source is unknown.

**Cause:** The smart card reader might be physically damaged.

**Solution:** Try the following:

- Check whether the smart card reader has an LED or a display that tells you whether the device is functioning properly. Verify using the manual supplied by your vendor.

- Try a different smart card reader.

**The smart card does not work.**
**Cause:** The smart card might be inserted incorrectly.

**Solution:** Try the following:

- The smart card might be inserted upside down in the smart card reader. Make sure you insert it with the correct side up.

- If this does not work, see the solutions for "The smart card reader does not work."

**Cause:** The smart card might be physically damaged.

**Solution:** Try the following:

- See the solutions for "The smart card reader does not work."

- If the smart card reader is working properly and the smart card is inserted correctly, then the smart card itself could be damaged. Try a different smart card.

- Contact your system administrator for additional support.

**When you try to establish a remote access connection using a smart card, you receive an error.**
**Cause:** Multiple users might be logged on to the same computer.

**Solution:** Try the following:

- Log off any additional users that are logged on at the same time using Fast User Switching.

- If more than one user is logged on with administrative credentials, the following error message appears: "Local Security Authority cannot be contacted." Log off additional users that have logged on with administrative credentials.

- If more than one user is logged on using local user credentials, the following error message appears: "Cannot find a smart card reader (Error 0x8010002e). For customized troubleshooting information for this connection, click Help." Log off additional users that have logged on using local user credentials.

# Error messages in Network Connections (Windows Server 2003)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Error messages**

You may receive one or more of the following errors when you use Network Connections. To see more information about an error, click the error.

- 600 An operation is pending.
- 601 An invalid port handle was detected.
- 602 The specified port is already open.
- 603 The caller's buffer is too small.
- 604 Incorrect information was specified.
- 605 The port information cannot be set.
- 606 The specified port is not connected.
- 607 An invalid event was detected.
- 608 A device was specified that does not exist.
- 609 The device type does not exist.
- 610 The buffer is invalid.
- 611 The route is not available.
- 612 The route is not allocated.
- 613 Invalid compression specified.
- 614 Out of buffers.
- 615 The port was not found.
- 616 An asynchronous request is pending.
- 617 The port or device is already disconnecting.
- 618 The port is not open.
- 619 A connection to the remote computer could not be established, so the port used for this connection was closed.
- 620 There are no endpoints.
- 621 Cannot open the phone book file.
- 622 Cannot load the phone book file.

- 623 Cannot find the phone book entry.
- 624 Cannot write the phone book file.
- 625 Invalid information found in the phone book.
- 626 Cannot load a string.
- 627 Cannot find key.
- 628 The connection was terminated by the remote computer before it could be completed.
- 629 The data link was terminated by the remote machine.
- 630 The port was disconnected due to hardware failure.
- 631 The port was disconnected by the user.
- 632 The structure size is incorrect.
- 633 The port is already in use or is not configured for Remote Access dialout.
- 634 Cannot register your computer on the remote network.
- 635 Unknown error.
- 636 The wrong device is attached to the port.
- 637 The string could not be converted.
- 638 The request has timed out.
- 639 No asynchronous net available.
- 640 A NetBIOS error has occurred.
- 641 The server cannot allocate NetBIOS resources needed to support the client.
- 642 One of your NetBIOS names is already registered on the remote network.
- 643 A network adapter at the server failed.
- 644 You will not receive network message popups.
- 645 Internal authentication error.
- 646 The account is not permitted to log on at this time of day.
- 647 The account is disabled.
- 648 The password has expired.
- 649 The account does not have Remote Access permission.
- 650 The Remote Access server is not responding.

- 651 Your modem (or other device) has reported an error.
- 652 Unrecognized response from the device.
- 653 A macro required by the device was not found in the device .INF file section.
- 654 A command or response in the device .INF file section refers to an undefined macro.
- 655 The <message> macro was not found in the device .INF file section.
- 656 The <defaultoff> macro in the device .INF file section contains an undefined macro.
- 657 The device .INF file could not be opened.
- 658 The device name in the device .INF or media .INI file is too long.
- 659 The media .INI file refers to an unknown device name.
- 660 The device .INF file contains no responses for the command.
- 661 The device .INF file is missing a command.
- 662 Attempted to set a macro not listed in the device .INF file section.
- 663 The media .INI file refers to an unknown device type.
- 664 Cannot allocate memory.
- 665 The port is not configured for Remote Access.
- 666 Your modem (or other device) is not functioning.
- 667 Cannot read the media .INI file.
- 668 The connection dropped.
- 669 The usage parameter in the media .INI file is invalid.
- 670 Cannot read the section name from the media .INI file.
- 671 Cannot read the device type from the media .INI file.
- 672 Cannot read the device name from the media .INI file.
- 673 Cannot read the usage from the media .INI file.
- 674 Cannot read the maximum connection BPS rate from the media .INI file.
- 675 Cannot read the maximum carrier BPS rate from the media .INI file.
- 676 The line is busy.
- 677 A person answered instead of a modem.

- 678 The remote computer did not respond.
- 679 Cannot detect carrier.
- 680 There is no dial tone.
- 681 General error reported by device.
- 691 Access denied because user name and/or password is invalid on the domain.
- 692 Hardware failure in port or attached device.
- 695 The state machines are not started.
- 696 The state machines are already started.
- 697 The response looping did not complete.
- 699 The device response caused buffer overflow.
- 700 The expanded command in the device .INF file is too long.
- 701 The device moved to a BPS rate not supported by the COM driver.
- 703 The connection needs information from you, but the application does not allow user interaction.
- 704 The callback number is invalid.
- 705 The authorization state is invalid.
- 707 There was an error related to the X.25 protocol.
- 708 The account has expired.
- 709 Error changing password on domain.
- 710 Serial overrun errors were detected while communicating with your modem.
- 711 A configuration error on this computer is preventing this connection.
- 712 Biplex port is initializing. Wait a few seconds and redial.
- 713 No active ISDN lines are available.
- 714 Not enough ISDN channels are available to make the call.
- 715 Too many errors occurred because of poor phone line quality.
- 716 The Remote Access IP configuration is unusable.
- 717 No IP addresses are available in the static pool of Remote Access IP addresses.
- 718 The connection was terminated because the remote computer did not respond in a timely manner.
- 719 PPP was terminated by remote machine.

- 720 A connection to the remote computer could not be established. You might need to change the network settings for this connection.
- 721 The remote computer did not respond.
- 722 The PPP packet is invalid.
- 723 The phone number, including prefix and suffix, is too long.
- 726 The IPX protocol cannot be used for dial-out on more than one port at a time.
- 728 Cannot find an IP adapter bound to Remote Access.
- 729 SLIP cannot be used unless the IP protocol is installed.
- 731 The protocol is not configured.
- 732 The PPP negotiation is not converging.
- 733 A connection to the remote computer could not be completed.
- 734 The PPP link control protocol was terminated.
- 735 The requested address was rejected by the server.
- 736 The remote computer terminated the control protocol.
- 737 Loopback detected.
- 738 The server did not assign an address.
- 739 The remote server cannot use the encrypted password.
- 740 The TAPI devices configured for Remote Access failed to initialize or were not installed correctly.
- 741 The local computer does not support encryption.
- 742 The remote server does not support encryption.
- 743 The remote server requires encryption.
- 752 A syntax error was encountered while processing a script.
- 753 The connection could not be disconnected because it was created by the multi-protocol router.
- 754 The system could not find the multi-link bundle.
- 755 The system cannot perform automated dial because this entry has a custom dialer specified.
- 756 This connection is already being dialed.
- 757 Remote access services could not be started automatically. Additional information is provided in the event log.
- 758 Internet Connection Sharing is already enabled on the connection.

- 760 An error occurred while routing capabilities were being enabled.
- 761 An error occurred while Internet Connection Sharing was being enabled for the connection.
- 763 Internet Connection Sharing cannot be enabled. There are two or more LAN connections in addition to the connection to be shared.
- 764 No smart card reader is installed.
- 765 Internet Connection Sharing cannot be enabled. A LAN connection is already configured with the IP address required for automatic IP addressing.
- 766 A certificate could not be found. Connections that use the L2TP protocol over IPSec require the installation of a machine certificate, also known as a computer certificate.
- 767 Internet Connection Sharing (ICS) cannot be enabled. The LAN connection selected on the private network has more than one IP address configured. Reconfigure the LAN connection with a single IP address before enabling ICS.
- 768 The connection attempt failed because of failure to encrypt data.
- 769 The specified destination is not reachable.
- 770 The remote computer rejected the connection attempt.
- 771 The connection attempt failed because the network is busy.
- 772 The remote computer's network hardware is incompatible with the type of call requested.
- 773 The connection attempt failed because the destination number has changed.
- 774 The connection attempt failed because of a temporary failure. Try connecting again.
- 775 The call was blocked by the remote computer.
- 776 The call could not be connected because the destination has invoked the Do Not Disturb feature.
- 777 The connection attempt failed because the modem (or other connecting device) on the remote computer is out of order.
- 778 It was not possible to verify the identity of the server.
- 779 To dial out using this connection you must use a smart card.
- 780 An attempted function is not valid for this connection.
- 781 The encryption attempt failed because no valid certificate was found.
- 782 Network address translation (NAT) is currently installed as a routing protocol and must be removed before enabling Internet Connection Sharing.
- 783 Internet Connection Sharing (ICS) cannot be enabled. The LAN connection selected as the private network is either not present, or it is disconnected from the network. Please ensure that the LAN adapter is connected before enabling ICS.
- 784 You cannot dial using this connection at logon time, because it is configured to use a user name different than the one on the smart card. If you want to use it at logon time, you must configure it to use the user name on the smart card.

- 785 You cannot dial using this connection at logon time, because it is not configured to use a smart card. If you want to use it at logon time, you must edit the properties of this connection so that it uses a smart card.
- 786 The L2TP connection attempt failed because there is no valid machine certificate on your computer for security authentication.
- 787 The L2TP connection attempt failed because the security layer could not authenticate the remote computer.
- 788 The L2TP connection attempt failed because the security layer could not negotiate compatible parameters with the remote computer.
- 789 The L2TP connection attempt failed because the security layer encountered a processing error during initial negotiations with the remote computer.
- 790 The L2TP connection attempt failed because certificate validation on the remote computer failed.
- 791 The L2TP connection attempt failed because security policy for the connection was not found.
- 792 The L2TP connection attempt failed because security negotiation timed out.
- 793 The L2TP connection attempt failed because an error occurred while negotiating security.
- 794 The Framed Protocol RADIUS attribute for this user is not PPP.
- 795 The Tunnel Type RADIUS attribute for this user is not correct.
- 796 The Service Type RADIUS attribute for this user is neither Framed nor Callback Framed.
- 797 A connection to the remote computer could not be established because the modem was not found or was busy.
- 798 A certificate could not be found that can be used with this Extensible Authentication Protocol.
- 799 Internet Connection Sharing (ICS) cannot be enabled due to an IP address conflict on the network. ICS requires that the host be configured to use 192.168.0.1. Ensure that no other client on the network is configured to use 192.168.0.1.
- 800 Unable to establish the VPN connection. The VPN server may be unreachable, or security parameters may not be configured properly for this connection.
- 801 This connection is configured to validate the identity of the access server, but Windows cannot verify the digital certificate sent by the server.
- 802 The card supplied was not recognized. Please check that the card is inserted correctly, and fits tightly.

**600**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**601**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**602**
The COM port that the network connection attempted to use is already in use by another active network connection or another process (for example, a phone line monitoring program, such as a fax program). Close the application that is blocking the use of the COM port.

**603**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**604**
The remote access phonebook file and the current Network Connections configuration are probably inconsistent. If you changed your communications equipment (such as your serial port or modem), make sure you reconfigure Network Connections. See Configure a connection to a remote network.

**605**
The remote access phonebook file and the current Network Connections configuration are probably inconsistent. If you changed your communications equipment (such as your serial port or modem), make sure you reconfigure Network Connections. See Configure a connection to a remote network.

If the error persists, delete and re-create the connection in Network Connections.

**606**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**607**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**608**
The remote access phonebook file and the current Network Connections configuration are probably inconsistent. If you changed your communications equipment (such as your serial port or modem), make sure you reconfigure Network Connections. See Configure a connection to a remote network.

If the error persists, delete and re-create the connection in Network Connections.

**609**
The remote access phonebook file and the current Network Connections configuration are probably inconsistent. If you changed your communications equipment (such as your serial port or modem), make sure you reconfigure Network

Connections. See Configure a connection to a remote network.

If the error persists, delete and re-create the connection in Network Connections.

**610**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**611**
Your network configuration is probably incorrect. Restart your computer to make sure that all recent configuration changes have taken effect. If the error persists, check Event Viewer for detailed warnings or errors.

**612**
Your network configuration is probably incorrect. Restart your computer to make sure that all recent configuration changes have taken effect. This error may also occur if your computer is running low on resources. If the error persists, check Event Viewer for detailed warnings or errors.

**613**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**614**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

If the error persists, check Event Viewer for detailed warnings or errors.

**615**
The remote access phonebook file and the current Network Connections configuration are probably inconsistent. If you changed your communications equipment (such as your serial port or modem), make sure you reconfigure Network Connections. See Configure a connection to a remote network.

If the error persists, delete and re-create the connection in Network Connections.

**616**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**617**
Wait for Network Connections to finish disconnecting.

**618**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**619**
There are several possible reasons why a connection to the remote computer could not be established:

- The remote computer might have been too busy. Wait a few minutes and try the connection again.
- If you are trying to establish a dial-up connection, you might have tried to redial before the modem fully disconnected. Wait a short time and try your call again.
- If you are trying to establish a connection by using a modem, the modem might not be functioning properly. For more information, see Troubleshooting modems.
- If you are using a device such as a router, a hub, or a network adapter for network address translation (NAT), the device might not be functioning properly. If the device provides firewall capabilities, the device might be blocking the connection. Consult the documentation for the device.

**620**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

If the error persists, check Event Viewer for detailed warnings or errors.

**621**
Network Connections uses the Rasphone.pbk file located in the systemroot\System32\Ras folder. Make sure the file is in this folder, and then restart Network Connections.

**622**
Network Connections uses the Rasphone.pbk file located in the systemroot\System32\Ras folder. Make sure the file is in this folder, and then restart Network Connections.

**623**
Network Connections located the phone book but cannot find the specified connection entry. This error should not occur unless another application is using the Network Connections phone book and has specified an incorrect connection entry.

**624**
Network Connections uses the Rasphone.pbk file located in the systemroot\System32\Ras folder. Make sure your disk is not full and that you have permission to change this file.

**625**
The phonebook file, Rasphone.pbk, may be corrupt. Delete this file from the systemroot\System32\Ras folder, and then restart Network Connections to create a new one.

**626**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.

- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

If the error persists, check Event Viewer for detailed warnings or errors.

**627**
One of the Network Connections configuration files probably contains invalid information. If you are using a modem that is not supported, install a supported modem.

**628**
The connection could not be completed because the connection attempt was terminated by the remote computer. Try the following:

- If you are trying to establish a dial-up connection, make sure that the phone number you are dialing is correct. Try dialing this phone number from your telephone. For more information, see Configure a connection to a remote network.

- If you are trying to establish a connection by using a modem, you might need to change the modem settings. For more information, see Troubleshooting modems.

- If you are trying to establish a connection to a virtual private network (VPN), you might need to change the security settings for this type of connection. For more information, see VPN connection authentication and data encryption, or contact the system administrator for the VPN to which you are trying to connect.

**629**
The connection link disconnected for one of the following reasons:

- An unrecoverable phone line error.

- A noisy line.

- Disconnection by the system administrator.

- Failure to negotiate correctly with the modem on the remote computer at the selected speed.

To reestablish the connection, click **Redial**. You can also enable automatic redialing on the **Options** tab for the connection. If this error persists, reduce the modem speed for the connection to 9600 bps, and then try to redial. See Change the maximum modem port speed.

You can try to connect to another computer to determine if the problem is related to the particular computer you are calling. Also, you can try connecting to the original computer from another phone line.

If the problem persists, contact the system administrator of the remote computer.

**630**
The link disconnected for one of the following reasons:

- An unrecoverable error occurred in your modem (or other communication device).

- An unrecoverable error occurred on your communications port.

- Your modem cable is unplugged.

To diagnose and correct the problem, do the following:

- Make sure your modem is on and your cable is securely attached.

- Make sure your modem is functioning properly. For instructions on testing your modem, see Test a modem.

**631**
An action on your computer disconnected the line. Redial.

**632**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

If the error persists, check Event Viewer for detailed warnings or errors.

**633**
If this is a dial-up connection, the COM port that the network connection attempted to use is already in use by another active network connection or another process (for example, a phone line monitoring program, such as a fax program). Close the application that is blocking the use of the COM port. For more information about troubleshooting modem problems, see Troubleshooting modems.

If this is a virtual private network (VPN) connection, the VPN device that the network connection attempted to use cannot be opened. If the problem persists, contact your system administrator.

If this is a broadband connection, the broadband device that the network connection attempted to use cannot be opened. Make sure that the required hardware, such as a cable modem, is functioning properly. For more information about required hardware, contact your broadband service provider.

**634**
The remote access server cannot register the name of your computer on the network. This generally occurs with the NetBIOS protocol, but it may also occur with TCP/IP or IPX protocols. Typically, the error occurs when an address is already in use on the network. Contact your system administrator.

**635**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click Start, and then click Help and Support.

If the error persists, check Event Viewer for detailed warnings or errors.

**636**
Your hardware configuration and the configuration settings for your connection are probably inconsistent with each other. If you changed your communications equipment (such as a serial port or modem), make sure you reconfigure the connection. See To configure a connection to a remote network.

**637**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**638**
An internal error occurred.

- Check the networking configurations for your computer.

- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**639**
Your NetBIOS network configuration is probably incorrect. Restart your computer to make sure that all recent configuration changes have taken effect. If the error persists, check Event Viewer for detailed warnings or errors.

If you call for help, record the Fail Code (a NetBIOS error code) and give it to the technical support staff.

**640**
The modem cannot negotiate a connection at the speed that is set. Set the modem to a lower initial speed, and then redial. See Change the maximum modem port speed. Also, try disabling modem compression and software compression. See Configure a modem for a dial-up connection or Disable software compression.

If you still cannot make a connection, try adding the **NWLink IPX/SPX/NetBIOS Compatible Transport Protocol** to the connection. See Add a network component.

**641**
Ask your system administrator to increase the resource capacity of the remote access server, or use **Services** in **Computer Management** to stop nonessential services on your computer, such as messenger services and Network DDE.

**642**
Another computer with the same computer name is already logged on to the remote network. Each computer on the network must be registered with a unique name. Verify the following:

- A computer with the same computer name as yours is not located on the network you are attempting to connect to.
- Your computer is not already physically attached to the network that you are attempting to connect to.

**643**
Report this error to your system administrator.

**644**
Another computer that is already connected to the network is using your messaging name. Messages addressed to you are sent to that computer. If you want to receive messages at your remote workstation, you must log off your office computer before you dial in to the network the next time.

This error does not affect messaging through Microsoft Outlook, Microsoft Outlook Express, or Microsoft Exchange.

**645**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**646**
Your account is configured for limited access to the network. If you need to access the network at a different time of day than what is presently configured, ask your system administrator to change the configuration.

**647**
The user account is disabled. This may be due to repeated failed logon attempts or because the system administrator disabled the account for security reasons. Ask your system administrator to enable your account in **Local Users and Groups**.

**648**
If you connect through Network Connections, you are automatically prompted to change your password. If you connect by using the **rasdial** command, you can change your password by doing the following:

- Press CTRL+ALT+DEL.

- Click **Change Password**, and then follow the instructions on the screen.

If you belong to an administrators group and your password expires, another administrator must change your password. You cannot change the password yourself.

**649**
The account does not have permission to connect for one of the following reasons:

- You have a valid account on the domain you selected, but your account does not have permission to access the network remotely. Ask your system administrator to enable dial-up permission for your user account or to enable dial-up permission in Routing and Remote Access.

- Your account has expired, been disabled, or is locked out, or your dial-up access has been locked out.

- You are either attempting to connect outside of allowable server logon hours or attempting to connect outside of the times allowed for dial-up access, or the policy that applies to your account may not allow dial-up access.

- A caller ID rule may prevent you from connecting. For example, your account may be required to dial in from a specific number.

- The remote computer may allow only local accounts to connect.

- An authentication protocol may be required that your computer cannot negotiate, or your computer may be attempting to use a protocol that is not authorized by the policy on the remote computer.

If you have an account with dial-in permission on another domain, do the following to use your account on that domain:

1. Right-click the connection, and then click **Properties**.

2. On the **Options** tab, select the **Prompt for name and password, certificate, etc.** and **Include Windows logon domain** check boxes.

3. On the **Security** tab, clear the **Automatically use my Windows logon name and password (and domain if any)** check box, and then click **OK**.

4. Double-click the connection, and then click **Dial**.

5. Specify the appropriate user name, password, and domain.

**650**
One of the following may have caused this error:

- The remote access server is not running. Contact your system administrator to make sure the server is running.

- The line may be too noisy, or your modem may have failed to negotiate correctly with the modem of the remote access server at the selected speed. For either of these possibilities, lower the initial speed (bps) of the modem, and then redial. See Change the maximum modem port speed.

- The serial cable of your modem may need replacing.

- The authentication settings for the connection may be incorrect. Contact your system administrator to make sure that your authentication settings meet those required by the remote access server.

- The software compression and modem hardware compression for the remote access server may be enabled at the same time. Generally, the remote access server software compression is enabled and hardware compression is disabled.

**651**
Your modem (or other device) reported an error.

If this is a dial-up connection and you are using a supported external modem, turn off and then restart the modem. Close and restart Network Connections, and then redial. If your modem is not supported by Network Connections, switch to a supported modem. Make sure you correctly configured your modem for remote access. To check your configuration, see Configure a connection to a remote network.

If this is a virtual private network (VPN) connection, you may have specified an incorrect TCP/IP address in the connection configuration, or the server that you are trying to reach may not be available. To determine if the server is available, contact your system administrator.

**652**
Your modem (or other device) returned a message that is not listed in one or more of your script files (Pad.inf, Switch.inf, or FileName.scp).

If you are using a supported external modem, turn off and restart the modem, and then redial.

If the problem persists, try connecting at a lower initial speed. See Change the maximum modem port speed.

**653**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**654**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**655**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**656**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**657**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**658**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**659**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**660**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**661**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**662**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**663**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**664**
The system is out of memory. Close some programs, and then redial.

**665**
If the device is already being used for another connection, hang up that connection.

Modify the connection to use another device. See Configure a connection to a remote network.

**666**
Your modem (or other device) is not responding for one of the following reasons:

- The external modem is turned off.
- The modem is not securely connected to your computer. Make sure the cable is securely fastened to both the modem and the computer.
- The serial cable does not meet the specifications required by Network Connections.
- The modem has experienced a hardware problem. Turn off the modem, wait for 20 seconds, and then restart the modem.

**667**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**668**
Redial the connection. If you continue to receive this message, reduce the initial speed of the modem and turn off advanced modem features, such as software compression. See Change the maximum modem port speed or Disable software compression.

Also, you should confirm that the modem is compatible.

If you are dialing manually, make sure you are connected before you click **Done**. You are dialing manually if you selected **Operator-Assisted Dialing** on the **Advanced** menu in the Network Connections folder.

**669**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**670**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**671**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**672**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**673**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**674**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**675**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**676**
Redial the number. You can automate redialing on the **Options** tab in the properties of the connection. See Configuring redial options. If this is a virtual private network (VPN) connection, check the host name or IP address of your destination server, and try to connect again. Also, contact your system administrator to verify that the remote server is running.

**677**
A modem or other device did not pick up the phone. Check the number and dial again. If this is a virtual private network (VPN) connection, check the host name or IP address of your destination server, and try to connect again.

**678**
The connection could not be completed because the remote computer did not answer the connection request. Try the following:

- If you are trying to establish a dial-up connection, make sure that the phone number you are dialing is correct. Try dialing this phone number from your telephone. For more information, see Configure a connection to a remote network.
- If you are trying to establish a connection by using a modem, make sure the modem is functioning properly. For more information, see Troubleshooting modems.
- If you are trying to establish a connection to a virtual private network (VPN), make sure that the host name or the IP address of your destination server is correct, and try again to establish the connection. To view the properties of a connection, right-click the connection, click **Properties**, and review the information on the **General** tab. For more information, contact the system administrator for the VPN to which you are trying to connect.
- If you are trying to establish a connection to a VPN, network congestion or a network device (such as a firewall) might be blocking your attempt to make a connection. For more information, contact the system administrator for the VPN to which you are trying to connect.
- If you are trying to establish a broadband connection, make sure the service name is correct and try again to establish the connection. To view the properties of a connection, right-click the connection, click **Properties**, and review the information on the **General** tab. For more information, contact your broadband service provider.
- If you are trying to establish a broadband connection, make sure that the necessary hardware (such as a cable modem) is installed and functioning properly. For more information, contact your broadband service provider.

**679**
A modem or other device did not pick up the phone. Many modems return this error if the remote modem does not pick up the phone. Check the number and dial again. If this is a virtual private network (VPN) connection, check the host name or IP address of your destination server, and try to connect again. You can automate redialing on the **Options** tab in the properties of the connection. See Configuring redial options.

**680**
A phone line may be incorrectly connected to your modem or disconnected from your modem. The number may require a prefix, such as 9, to access external lines, or the number may be too long.

Make sure that the phone line is plugged into the correct socket in the modem. Also, make sure you added any special access numbers, such as the prefix 9 followed by a comma, to connect to an outside line (for example: 9,555-0100). For more information, see Configure phone number modifiers.

Check that there is no stutter tone, which indicates voicemail messages, on your phone line.

Many modems cannot dial more than 34 digits. When longer numbers are encountered, these modems split them into two or more parts, of which only the first part is dialed. This occurs with USRobotics and Multitech modems. The only known solution for this problem is to use another brand of modem.

For more information, see Technical support options.

**681**
One of the Network Connections configuration files probably contains information that is not valid.

If your modem is not supported by Network Connections, switch to a supported modem.

**691**
Either your user account is not registered with the domain listed, your password is expired, or you mistyped the information. If you do not specify a domain, the remote access server attempts to verify your user name and password on the domain of which it is a member.

Carefully retype your user name, password, and domain. If you are unsure of this information, contact your system administrator.

**692**
Your modem (or other device) is not responding because of the following:

- The modem is turned off, malfunctioning, or is not securely connected to your computer.

To resolve this problem, do the following:

- Reset your modem. For more information, consult your modem documentation.
- If you are using an external modem, ensure that you are using a proper serial cable and that it is connected securely. You may want to try replacing the modem cable. Also, if you are using an adapter to connect your external modem to your serial port, ensure that the adapter is wired properly for modem communication. For example, a 9-to-25 pin adapter for a mouse does not work properly with a serial network connection.
- Test the serial port or multiport adapter, and replace it if necessary.
- Ensure that the handshaking options for the modem are configured properly. Consult your modem documentation for information about different handshaking options available for your modem.
- If your modem is not supported by Network Connections, switch to a supported modem.
- Verify that another application, such as HyperTerminal, is not using the communications port. If the port is being used, subsequently starting Network Connections may cause this message to appear.
- For more information, see Technical support options.

**695**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**696**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**697**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**699**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**700**
Commands in script files cannot exceed 256 characters. Break up the command into multiple commands. See Automating the Dial-up Logon Process by Using Switch.inf Scripts.

**701**
Your modem tried to connect at a speed that the serial port cannot interpret. Reset your initial speed to the next lowest standard bps rate. See Change the maximum modem port speed.

**703**
The server you are attempting to connect to requires user interaction. The rasdial command, or the application you are using that is attempting to dial, does not support user interaction.

If possible, try to connect by using an existing connection in the Network Connections folder.

If you are using scripting to connect, try using a dial-up connection that is configured with the Terminal feature. The Terminal window may enable the user interaction that is required. See Use the Terminal feature to log on to a remote computer.

**704**
The callback number that you specified for the client is not valid. Check to make sure the callback number has been typed correctly, including area code.

**705**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**707**
The X.25 connection returned an error. Ask your X.25 provider to interpret the diagnostic information that is provided.

**708**
Ask your system administrator to reactivate your account.

**709**
Try changing your password again. If you continue to receive this message, report it to your system administrator.

**710**
Lower the modem's initial speed (bits per second), and then redial. See Change the maximum modem port speed.

**711**
A connection could not be established because a required service was not started on your computer. Make sure the following services are started on your computer and try again to establish the connection:

- Plug and Play

- Remote Access Auto Connection Manager

For information about how to start a service, see Start, stop, pause, resume, or restart a service.

**712**
This error can occur on a connection that is configured to receive calls (biplex connection). It happens when you dial out at the same time the server is initializing the connection for receiving calls. Network Connections redials in a few seconds.

**713**
Make sure that your ISDN line is plugged in correctly and that the terminating resistors are installed correctly (see the documentation for your ISDN card), and then redial. If you still receive this error, contact customer service for the ISDN card vendor or your ISDN telephone company.

**714**
All of the available ISDN channels are busy. Typically, ISDN provided in the BRI (Basic Rate Interface) level of service offers two channels on which voice or data calls can be made. If both channels are in use, Network Connections is not able to dial out. Hang up one channel and then redial.

**715**
Too many asynchronous errors occurred on your phone line during authentication. Try again. If the problem persists, lower your baud rate and disable any modem features that are set. See Change the maximum modem port speed and Configure an installed modem.

**716**
There is a problem with the TCP/IP configuration on the remote access server. For example, your connection is requesting a TCP/IP address from the server that is not usable.

Contact your system administrator to verify the TCP/IP settings on your server.

**717**
Try to use an IP address that does not cause a conflict on the remote network. If possible, use DHCP to avoid addressing conflicts. See Configure TCP/IP settings.

**718**
A Point to Point Protocol (PPP) connection was initiated but could not be completed because the remote computer did not respond in a timely manner or it did not give a valid response. This can occur for any of several reasons:

- If you are trying to establish a connection to a virtual private network (VPN), try entering your domain name and your user name in **User name** as follows: DOMAIN\username, and try again to establish the connection.

- If you are trying to establish a connection to a remote network, network congestion or a network device (such as a firewall) might be blocking your attempt to make a connection. Try your connection later, or contact the administrator of the remote network for more information.

- You might need to install updated drivers for your modem, your network card, or another connection device. For more information, contact your hardware manufacturer.

**719**
A PPP conversation started, but it was terminated at the request of the remote computer. An error probably occurred on the server.

**720**
A connection could not be established because the Point to Point Protocol (PPP) settings on your computer and on the remote computer do not match. Try the following:

- Make sure that TCP/IP is installed and enabled for this type of connection. For more information, see Enable a network component.

- Make sure that all protocols needed to establish this type of connection are installed. For more information, see PPP operation and protocols.

- Contact the administrator of the remote computer for more information about establishing this type of connection.

**721**
The connection could not be established because the remote computer did not respond to the connection request. Try the following:

- Make sure that your modem is functioning properly. For more information, see Troubleshooting modems.
- Make sure that TCP/IP is installed and configured correctly for this connection. For more information, see Connect to the Internet by using PPPoE.
- This connection might require a Terminal window. For more information on enabling a Terminal window, see Use the Terminal feature to log on to a remote computer.
- If you are trying to establish a virtual private network (VPN) connection, make sure that the host name or the IP address of your destination server is correct and try the connection again. To check this information, right-click the connection, click **Properties**, and review the information on the **General** tab.

**722**
A PPP packet was received that was not in a valid format. Your system administrator may need to log PPP events to a file to solve the problem. For more information, see Support resources.

**723**
The maximum length of the phone number, including prefix and suffix, is 128 characters.

**726**
Only one connection can be used to dial out by using the IPX protocol.

**728**
There is a problem with the TCP/IP configuration. Restart your computer to ensure that all recent configuration changes have taken effect.

**729**
Verify that TCP/IP is installed. For more information, see Add a network component.

**731**
Specific information was requested on a PPP control protocol that is not configured. To identify the PPP negotiation problem, ask your system administrator to enable PPP logging by using the Netsh.exe utility.

**732**
PPP parameter negotiation did not succeed because the local and remote computers did not agree on a common set of parameters. Contact your system administrator to verify the configuration of network protocols such as TCP/IP or IPX. Also, your system administrator can enable PPP logging by using the Netsh.exe utility, which may help identify the PPP negotiation problem.

**733**
The connection could not be completed because your computer and the remote computer could not negotiate PPP control protocols. Try the following:

- Make sure that all network components necessary for the connection are installed and configured for the connection type. For more information, see Add a network component and Configure a connection to a remote network.
- Enable or disable multi-link negotiation for the connection. To configure multi-link, open the **Network Connections** folder, right-click the icon for the connection, click **Properties**, click the **Networking** tab, and click **Settings**.

**734**
The PPP link control protocol was terminated. There are several possible causes. Try the following:

- Your Internet service provider might have experienced a problem. Contact your Internet service provider.
- The security protocols on the server to which you are trying to connect are incompatible with your security settings. In **Network Connections**, right-click the connection you are using. Click **Properties**, and click the **Security** tab. In **Validate my identity as follows**, select **Allow unsecured password**.

- LCP extensions are not supported. In **Network Connections**, right-click the connection you are using. Click **Properties**, and click the **Networking** tab. Click **Settings**, and clear the **Enable LCP extension** check box.

**735**
Your connection is configured to request a specific IP address. Either the server is not configured to permit clients that request specific IP addresses, or the specific IP address may already be in use by another client. If possible, use DHCP to avoid addressing conflicts. See Configure TCP/IP settings.

**736**
A PPP network control protocol conversation started, but it was terminated at the request of the remote computer. An error probably occurred on the server. See your system administrator. Also, your system administrator can enable PPP logging by using the Netsh.exe utility, which may help to identify the PPP negotiation problem.

**737**
The local and remote computers that are involved in the PPP conversation are the same computer. This usually means that a device (for example, a modem) on the link is echoing characters back. Try resetting these devices.

For other servers from other vendors, it may indicate that the remote computer is attempting a teletypewriter (TTY) logon before connecting. You can configure the connection for a post-connect terminal on the **Security** tab.

**738**
The server cannot assign an IP address to your computer from the pool of allocated addresses. Ask your system administrator to configure the connection to use a specific IP address on the client that does not cause a conflict on the remote network. Or, if possible, use DHCP to avoid addressing conflicts. See Configure TCP/IP settings.

**739**
This error occurs only when making a PPP connection to a server that is not running Windows. The standard PPP authentication protocols, used for interoperating with other servers, require that the password is available in plaintext form. Servers that are running Windows store only the encrypted form for security.

**740**
The TAPI devices configured for Network Connections failed to initialize or are not installed correctly. Restart your computer to make sure that all recent configuration changes have taken effect. If the error persists, check Event Viewer for detailed warnings or errors.

**741**
This error may be resolved by either selecting or clearing the **Require data encryption (disconnect if none)** check box located on the **Security** tab. Check with your system administrator regarding client-side security settings.

This error may also occur if your password is longer than 14 characters, you selected the **Require data encryption (disconnect if none)** check box on the **Security** tab, and you selected the **Require secured password** check box in **Validate my identity as follows**. Use a password less than 14 characters, clear the Require data encryption (disconnect if none) check box, and clear the **Allow unsecured password** check box in **Validate my identity as follows**.

See Configure a connection to a remote network and VPN connection authentication and data encryption.

**742**
The required level of encryption is not available on the remote computer. One computer is probably using 128-bit encryption technology while the other is using 40-bit or 56-bit encryption. Contact your system administrator to determine your level of encryption and to verify that the server is using the same level of encryption. For more information, see Data encryption.

**743**
The remote access server requires encryption, but your connection is not enabled for encryption. See Configure identity authentication and data encryption settings.

**751**
An invalid callback number has been specified. Only the characters 0 to 9, T, P, W, (, ), -, @, and space are allowed in the number. Enter the correct number and redial.

**752**
A syntax error was encountered while processing a script. See Automating the Dial-up Logon Process by Using Switch.inf Scripts and Troubleshooting logon scripts.

**753**
This connection is being used by Routing and Remote Access. Use the Routing and Remote Access service to terminate the connection. The connection appears in either the list of routing interfaces or the list of remote access clients.

**754**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.

For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

**755**
Contact your system administrator.

**756**
Wait and try again later.

**757**
An internal error occurred.

- Check the networking configurations for your computer.
- Make sure all devices are functional.
- For more information about diagnosing a problem, click **Start**, and then click **Help and Support**.

If the error persists, check the Event Viewer for detailed warnings or errors.

**758**
This error may occur if an application enables Internet Connection Sharing while the properties of the connection are open, and then you try to enable the feature. Close and then reopen the properties of the connection. The **Enable Internet Connection Sharing** check box on the **Advanced** tab should be selected.

**760**
An internal error occurred while attempting to enable routing. Contact your system administrator.

**761**
An internal error occurred while Internet Connection Sharing was being enabled. Contact your system administrator.

**763**
Close and reopen the properties of the connection. The **Enable Internet Connection Sharing** check box on the **Advanced** tab should be selected. Select a LAN connection from the list.

**764**
Install a smart card reader.

**765**
Internet Connection Sharing configures the LAN connection of the home network or small office network with a static address. Another network adapter in the system is configured with the same address. Change the static address of the network adapter before enabling Internet Connection Sharing.

**766**
Contact your system administrator. If you are a member of an Active Directory domain and need to request a certificate, see Request a certificate. If you are not a member of an Active Directory domain or you need to request a certificate from the Internet, see Submit a user certificate request via the Web.

**767**
Remove the additional static IP addresses, or reconfigure the LAN connection for DHCP. See Configure TCP/IP settings.

**768**
If encryption is not necessary, disable encryption and try again. See Configure identity authentication and data encryption settings.

**769**
Either you specified a destination address that is not valid, or your remote access server is down. Check the destination address and try again.

**770**
The remote computer may be rejecting your connection because of a caller ID rule or another hardware setting.

**771**
Wait and try again later.

**772**
Contact your system administrator.

**773**
Enter the correct number and try again.

**774**
Wait and try again later. This error may indicate a timing issue or a problem with the remote computer.

**775**
The remote computer may be blocking your connection because of a caller ID rule, a scheduling policy that allows access only during specific times, or another setting.

**776**
Contact your system administrator.

**777**
Contact your system administrator.

**778**
In the authentication process, the server identifies itself to your computer, and your computer identifies itself to the server. When this error occurs, your computer is not recognizing the server.

**779**
Install a smart card reader. See Install a smart card reader on a computer.

**780**
An internal error occurred. Contact your system administrator. You may need to configure the connection for EAP.

**781**
This L2TP/IPSec connection requires a valid certificate to be installed on your computer, and no such certificate was found. Try the following:

- If you are a member of an Active Directory domain, request a certificate from the appropriate certification authority. For more information, see Request a certificate.
- If you want to request a certificate from the Internet, see Submit a user certificate request via the Web.
- If you want more information about certificates, see Smart card and other certificate authentication.
- If this connection requires a preshared key, contact your network administrator.

**782**
Remove the Network Address Translation (NAT) routing protocol. See Delete an IP routing protocol.

**783**
Before you enable Internet Connection Sharing, verify that the LAN adapter and its associated cabling are securely connected.

**784**
Configure the connection to use the same name as the smart card. Open the properties of the connection. On the **Security** tab, click **Advanced (custom settings)** and **Settings**. Click **Smart Card or other Certificate (encryption enabled)**, click **Properties**, and then clear the **Use a different user name for the connection** check box.

**785**
Configure the connection to use a smart card. See Enable smart card or other certificate authentication.

**786**
Contact the administrator of the remote network. Your connection may be failing for one or more of the following reasons:

- Your computer's certificate is invalid, is expired, or lacks the preshared key that it needs.

- None of your computer's certificates are trusted by the server.
- Your computer is using a preshared key that is not valid. Try using another pre-shared key.

If you are a member of an Active Directory domain and need to request a certificate, see Request a certificate. If you are not a member of an Active Directory domain or you need to request a certificate from the Internet, see Submit a user certificate request via the Web. For more information about remote access security, see How security works at connection.

**787**
Contact your system administrator. Your connection may be failing for one or more of the following reasons:

- Your user name and password credentials are incorrect.
- The certificate your computer is using expired.
- The certificate the server is using expired.

If you are a member of an Active Directory domain and need to request a certificate, see Request a certificate. If you are not a member of an Active Directory domain or you need to request a certificate from the Internet, see Submit a user certificate request via the Web.

**788**
Contact your system administrator. The current configuration of L2TP parameters is not compatible with the Microsoft implementation of L2TP.

**789**
Contact your system administrator.

**790**
Contact your system administrator. Your connection may be failing for one or more of the following reasons:

- The server you are reaching does not have a certificate.
- The server you are reaching does not have the correct certificate.
- Your connection is not configured to use the correct certificate.

If you are a member of an Active Directory domain and need to request a certificate, see Request a certificate. If you are not a member of an Active Directory domain or you need to request a certificate from the Internet, see Submit a user certificate request via the Web. For more information about remote access security, see How security works at connection.

**791**
Contact your system administrator. Your connection may be failing for one or more of the following reasons:

- The remote server may be offline.
- The server you are reaching does not have the correct credentials.
- Your connection is not configured with the correct credentials.

**792**
Check your connection for connectivity to the Internet. If the connection is working, contact your system administrator. Your connection may be failing for one or more of the following reasons:

- The remote server may be offline.
- The server you are reaching does not have the correct credentials.
- Your connection is not configured with the correct credentials.

- Your connection is configured to use a different security method than the remote computer. Try changing the type of VPN connection. To change the type of VPN connection, right-click the connection, click **Properties**, and review the information on the **Networking** tab.

**793**
Contact your system administrator. Your connection may be failing for one or more of the following reasons:

- The remote server may be offline.
- The server you are reaching does not have the correct credentials.
- Your connection is not configured with the correct credentials.
- Your computer is using a preshared key that is not valid. Try using another pre-shared key.

For more information about remote access security, see How security works at connection.

**794**
An attribute on the remote network must be configured to allow this connection. Contact the system administrator of the remote network.

**795**
An attribute on the remote network must be configured to allow this connection. Contact the system administrator of the remote network.

**796**
An attribute on the remote network must be configured to allow this connection. Contact the system administrator of the remote network.

**797**
This connection requires a modem. Try the following:

- Close all other programs and try again to establish the connection. Another program might be using the modem or the communication port to which the modem is attached.
- Make sure that a modem has been installed on your computer. For more information, see Install a modem.
- Make sure that the modem is functioning correctly. For more information, see Troubleshooting modems.

**798**
Contact your system administrator. If you are a member of an Active Directory domain and need to request a certificate, see To request a certificate. If you are not a member of an Active Directory domain or you need to request a certificate from the Internet, see Submit a user certificate request via the Web.

**799**
Check the IP address settings on all computers connected to the local network, and reconfigure them if necessary. See Configure TCP/IP settings.

**800**
The connection to the VPN server could not be established using either PPTP or L2TP. One or more of the following circumstances might have caused this error:

- The PPTP server is not reachable. The server name you supplied might not exist, or a network component (such as a firewall) might be blocking PPTP traffic.
- Verify whether the PPTP server name for this connection is correct. To check your configuration, see Configure a connection to a remote network.
- Check whether Internet Connection Firewall has been enabled for this connection. For more information, see Set up and

Manage Internet Connection Firewall (ICF).

- Contact your system administrator for additional support.

The L2TP/IPSec server is not reachable. The server name you supplied might not exist. Your connection configuration might prevent IPSec from establishing a secure channel with the server. A network component such as Network Address Translation (NAT) or a firewall might be blocking IPSec traffic.

- Verify whether the L2TP server name is correct and that the connection you are trying to use has been configured correctly for your computer. To check your configuration, see Configure a connection to a remote network.
- Verify whether you have a current machine certificate or that the preshared key is correct. For more information, see Enable smart card or other certificate authentication.
- Check whether Internet Connection Firewall has been enabled for this connection. For more information, see Set up and Manage Internet Connection Firewall (ICF).
- Contact your system administrator for additional support.

**801**
However, the digital certificate the server sent cannot be verified. One or more of the following circumstances might have caused this error.

- The root of the certificate chain that sent by the server is not available for the local system to verify.
- You have specified the server and the trusted root certification authority for your local system. The client computer cannot find or authenticate the root certification authority, which is probably missing and thus generating an error.
- If the server is not part of a domain, it might not have a root certification authority.

To resolve this problem:

- Ensure that you obtain the root certificate and distribute it to any user and computer that needs to establish trust in the root authority.
- Using Group Policy, verify whether the policy settings for your trusted root certification authority are appropriate. For more information, see Policies to establish trust of root certification authorities.
- Contact your system administrator for additional support.

**802**
The smart card that you are using could not be recognized. One or more of the following circumstances might have caused this error.

- The smart card is not inserted properly. For example, it might be inserted upside down.
- The smart card reader is not installed properly.
- If you are using cryptographic service provider (CSP) software, it might be corrupted or not installed properly.

To resolve this problem:

- Insert the smart card properly so that it fits tightly in the smart card reader.
- Insert that the smart card with the correct side up. See instructions on the smart card reader to ensure that the smart card is inserted with the correct side facing up.

- Verify that the smart card reader is installed properly. For more information, see Install a smart card reader on a computer.
- Ensure that you are using the correct CSP and that it is installed properly. For more information, see Set up a smart card for user logon or Key type and cryptographic service provider type.
- Contact your system administrator for additional support.

# Connect clients to a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connect clients to a printer**


- Connect to a printer published in Active Directory
- Connect to a printer on a network
- Connect to a printer with a browser
- Connect Windows NT 4.0, Windows 95, and Windows 98 clients
- Connect MS-DOS, Windows 3.x, and OS/2 clients
- Connect Macintosh clients to a printer
- Print from an LPR client to a printer
- Connect to a printer on a NetWare server
- Connect to a printer on a Macintosh server

# User and Group Accounts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User and group accounts**
This section covers:

- Managing Local Users and Groups
- Managing Domain Users and Groups

# Managing Local Users and Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing local users and groups**
This section covers:

- Local Users and Groups overview
- Manage Local Users
- Manage Local Groups

# Local Users and Groups overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local Users and Groups overview**

Local Users and Groups is located in Computer Management, a collection of administrative tools that you can use to manage a single local or remote computer. You can use Local Users and Groups to secure and manage user accounts and groups stored locally on your computer. A local user or group account can be assigned permissions and rights on a particular computer and that computer only. Local Users and Groups is available on the following client and server operating systems:

- Client computers running Microsoft® Windows® 2000 Professional or Windows XP Professional
- Member servers running a product in the Microsoft Windows 2000 Server family or the Windows Server 2003 family
- Stand-alone servers running a product in the Microsoft Windows 2000 Server family or the Windows Server 2003 family

Using Local Users and Groups you can limit the ability of users and groups to perform certain actions by assigning them rights and permissions. A right authorizes a user to perform certain actions on a computer, such as backing up files and folders or shutting down a computer. A permission is a rule associated with an object (usually a file, folder, or printer), and it regulates which users can have access to the object and in what manner.

You cannot use Local Users and Groups to view local user and group accounts once a member server has been promoted to a domain controller. However, you can use Local Users and Groups on a domain controller to target remote computers (that are not domain controllers) on the network. Use Active Directory Users and Computers to manage users and groups in Active Directory.

For more information, see:

- Local user accounts
- Default local groups
- Default security settings for groups
- Active Directory Users and Computers (Windows Server 2003)

# Manage Local Users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage local users**

- Create a local user account
- Reset the password for a local user account
- Disable or activate a local user account
- Delete a local user account
- Rename a local user account
- Assign a logon script to a local user account
- Assign a home folder to a local user account

# Manage Local Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage local groups**

- Create a local group
- Add a member to a local group
- Identify members of a local group
- Delete a local group
- Allow members of the Anonymous Logon group to be members of the Everyone group on a local computer

# Managing Domain Users and Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing domain users and groups**
This section covers:

- User and computer accounts
- Manage Users
- Manage Groups
- Manage Computers

# User and computer accounts

Updated: July 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User and computer accounts**

Active Directory user accounts and computer accounts represent a physical entity such as a computer or person. User accounts can also be used as dedicated service accounts for some applications.

User accounts and computer accounts (as well as groups) are also referred to as security principals. Security principals are directory objects that are automatically assigned security IDs (SIDs), which can be used to access domain resources. A user or computer account is used to:

- Authenticate the identity of a user or computer.

  A user account enables a user to log on to computers and domains with an identity that can be authenticated by the domain. For information about authentication, see Access control in Active Directory. Each user who logs on to the network should have his or her own unique user account and password. To maximize security, you should avoid multiple users sharing one account.

- Authorize or deny access to domain resources.

  Once the user has been authenticated, the user is authorized or denied access to domain resources based on the explicit permissions assigned to that user on the resource. For more information, see Security information for Active Directory.

- Administer other security principals.

  Active Directory creates a foreign security principal object in the local domain to represent each security principal from a trusted external domain. For more information about foreign security principals, see When to create an external trust.

- Audit actions performed using the user or computer account.

  Auditing can help you monitor account security. For more information about auditing, see Auditing overview.

**User accounts**

The Users container located in Active Directory Users and Computers displays the three built-in user accounts: Administrator, Guest, and HelpAssistant. These built-in user accounts are created automatically when you create the domain.

Each built-in account has a different combination of rights and permissions. The Administrator account has the most extensive rights and permissions over the domain, while the Guest account has limited rights and permissions. The table below describes each default user account on domain controllers running Windows Server 2003.

| Default user account | Description |
|---|---|

| | |
|---|---|
| Administrator account | The Administrator account has full control of the domain and can assign user rights and access control permissions to domain users as necessary. This account must be used only for tasks that require administrative credentials. It is recommended that you set up this account with a strong password. For more information, see Strong passwords. For additional security considerations for accounts with administrative credentials, see Active Directory Best practices.<br><br>The Administrator account is a default member of the Administrators, Domain Admins, Enterprise Admins, Group Policy Creator Owners, and Schema Admins groups in Active Directory. The Administrator account can never be deleted or removed from the Administrators group, but it can be renamed or disabled. Because the Administrator account is known to exist on many versions of Windows, renaming or disabling this account will make it more difficult for malicious users to try and gain access to it. For more information about how to rename or disable a user account, see Rename a local user account or Disable or enable a user account.<br><br>The Administrator account is the first account created when you set up a new domain using the Active Directory Installation Wizard.<br><br>• **Important** When the Administrator account is disabled, it can still be used to gain access to a domain controller using Safe Mode. |
| Guest account | The Guest account is used by people who do not have an actual account in the domain. A user whose account is disabled (but not deleted) can also use the Guest account. The Guest account does not require a password.<br><br>You can set rights and permissions for the Guest account just like any user account. By default, the Guest account is a member of the built-in Guests group and the Domain Guests global group, which allows a user to log on to a domain. The Guest account is disabled by default, and it is recommended that it stay disabled. |
| HelpAssistant account (installed with a Remote Assistance session) | The primary account used to establish a Remote Assistance session. This account is created automatically when you request a Remote Assistance session and has limited access to the computer. The HelpAssistant account is managed by the Remote Desktop Help Session Manager service and will be automatically deleted if no Remote Assistance requests are pending. For more information about Remote Assistance, see Administering Remote Assistance. |

**Securing user accounts**

If built-in account rights and permissions are not modified or disabled by a network administrator, they could be used by a malicious user (or service) to illegally log on to a domain using the Administrator or Guest identity. A good security practice for protecting these accounts is to rename or disable them. Because it retains its security ID (SID), a renamed user account retains all its other properties, such as its description, password, group memberships, user profile, account information, and any assigned permissions and user rights.

To obtain the security of user authentication and authorization, create an individual user account for each user who will participate on your network by using Active Directory Users and Computers. Each user account (including the Administrator and Guest account) can then be added to a group to control the rights and permissions assigned to the account. Using accounts and groups that are appropriate for your network ensures that users logging on to a network can be identified and can access only the permitted resources.

You can help defend your domain from attackers by requiring strong passwords and implementing an account lockout policy. Strong passwords reduce the risk of intelligent guessing and dictionary attacks on passwords. For more information, see Strong passwords and Password Best practices for passwords.

An account lockout policy decreases the possibility of an attacker compromising your domain through repeated logon attempts. This is because an account lockout policy determines how many failed logon attempts a user account can have before it is disabled. For more information, see Apply or modify account lockout policy.

For more information about securing user accounts, see Securing Active Directory.

**Account options**

Each Active Directory user account has a number of account options that determine how someone logging on with that particular user account is authenticated on the network. You can use the following options to configure password settings and security-specific information for user accounts:

| Account option | Description |
| --- | --- |
| User must change password at next logon | Forces a user to change their password the next time the user logs on to the network. Use this option when you want to ensure that the user will be the only person to know their password. |
| User cannot change password | Prevents a user from changing their password. Use this option when you want to maintain control over a user account, such as for a guest or temporary account. |
| Password never expires | Prevents a user password from expiring. It is recommended that Service accounts should have this option enabled and should use strong passwords. For more information about strong passwords, see Strong passwords. |
| Store passwords using reversible encryption | Allows a user to log on to a Windows network from Apple computers. If a user is not logging on from an Apple computer, this option should not be used. For more information, see Store passwords using reversible encryption. |
| Account is disabled | Prevents a user from logging on with the selected account. Many administrators use disabled accounts as templates for common user accounts. For more information, see Disable or enable a user account. |
| Smart card is required for interactive logon | Requires that a user possess a smart card to log on to the network interactively. The user must also have a smart card reader attached to their computer and a valid personal identification number (PIN) for the smart card. When this option is selected, the password for the user account is automatically set to a random and complex value. For more information about smart cards, see Logging on to a computer with a smart card and Authentication process. |
| Account is trusted for delegation | Allows a service running under this account to perform operations on behalf of other user accounts on the network. A service running under a user account (otherwise known as a service account) that is trusted for delegation can impersonate a client to gain access to resources on the computer where the service is running or on other computers. In a forest set to the Windows Server 2003 functional level, this setting is found on the **Delegation** tab, and is only available for accounts that have been assigned service principal names (SPNs), as set using the **setspn** command from the Windows Support Tools. This is a security-sensitive capability and should be cautiously assigned. For more information, see Allow a user to be trusted for delegation and Delegating authentication.<br><br>This option is only available on domain controllers running Windows Server 2003 where the domain functionality is set to Windows 2000 mixed or Windows 2000 native. On domain controllers running Windows Server 2003 where the domain functional level is set to Windows Server 2003, the **Delegation** tab is used to configure delegation settings. The **Delegation** tab only appears for accounts which have an assigned SPN. For more information about domain functionality, see Domain and forest functionality. For more information about configuring delegation in a Windows Server 2003 domain, see Allow a user to be trusted for delegation. |
| Account is sensitive and cannot be delegated | Allows control over a user account, such as for a guest or temporary account. This option can be used if this account cannot be assigned for delegation by another account. |
| Use DES encryption types for this account | Provides support for the Data Encryption Standard (DES). DES supports multiple levels of encryption, including MPPE Standard (40-bit), MPPE Standard (56-bit), MPPE Strong (128-bit), IPSec DES (40-bit), IPSec 56-bit DES, and IPSec Triple DES (3DES). For more information about DES encryption, see Data encryption. |
| Do not require Kerberos preauthentication | Provides support for alternate implementations of the Kerberos protocol. Domain controllers running Windows 2000 or Windows Server 2003 can use other mechanisms to synchronize time. Because preauthentication provides additional security, use caution when enabling this option. For more information about Kerberos, see Kerberos V5 authentication. |

**InetOrgPerson accounts**

Active Directory provides support for the InetOrgPerson object class and its associated attributes defined in RFC 2798. The InetOrgPerson object class is used in several non-Microsoft LDAP and X.500 directory services to represent people within an organization.

Support for InetOrgPerson makes migrations from other LDAP directories to Active Directory more efficient. The InetOrgPerson object is derived from the user class and can be used as a security principal just like the user class. For

information about creating an inetOrgPerson user account, see Create a new user account.

When the domain functional level has been set to Windows Server 2003, you can set the userPassword attribute on InetOrgPerson and user objects as being the effective password just like you can with the *unicodePwd* attribute.

**Computer accounts**

Every computer running Windows NT, Windows 2000, Windows XP, or a server running Windows Server 2003 that joins a domain has a computer account. Similar to user accounts, computer accounts provide a means for authenticating and auditing computer access to the network and to domain resources. Each computer account must be unique.

**Note** Computers running Windows 95 and Windows 98 do not have advanced security features and are not assigned computer accounts.

User and computer accounts can be added, disabled, reset, and deleted using Active Directory Users and Computers. A computer account can also be created when you join a computer to a domain. For more information about user and computer accounts, see Active Directory naming and Object names.

When the domain functional level has been set to Windows Server 2003, a new lastLogonTimestamp attribute is used to track the last logon time of a user or computer account. This attribute is replicated within the domain and can provide you with important information regarding the history of a user or computer.

# Manage Users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage users**

- Create a new user account
- Reset a user password
- Copy a user account
- Move a user account
- Set logon hours
- Disable or enable a user account
- Map a certificate to a user account
- Change a user's primary group
- Delete a user account
- Rename a user account

# Rename a user account

Updated: August 9, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a user account**
**To rename a user account using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.
3. **Where?**
   - Active Directory Users and Computers/*domain node*/Users

   Or, click the folder that contains the user account.
4. In the details pane, right-click the user account that you want to rename, and then click **Rename**.
5. Type the new full name for the user account, and then press ENTER. The **Rename User** dialog box appears.
6. The name that you type appears as the full name. You may edit the name, if necessary. When you are satisfied with the new full name, press the Tab key to go to the next box, or click the appropriate box.
7. In **First name**, type the user's first name.
8. In **Last name**, type the user's last name.
9. In **Display name**, type the name for the user to be displayed in Active Directory.
10. In **User logon name**, type the user's logon name, click the user principal name (UPN) suffix in the drop-down list, and then click **Next**.
11. In **User logon name (pre-Windows 2000)**, type a name for the user that is unique for the Active Directory forest.
12. Click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**To rename a user account using a command line**

1. Open a command prompt.
2. Type the following command, and then press ENTER:

**Dsmod user** *UserDN* **[-upn** *UPN* **] [-fn** *FirstName* **] [-mi** *Initial* **] [-ln** *LastName* **] [-dn** *DisplayName* **] [-email** *EmailAddress* **]**

| Parameter | Description |
| --- | --- |
| *UserDN* | Specifies the distinguished names of the users that you want to modify. A user distinguished name **(UserDN**) is required. If values are omitted, they are optained through standard input (**stdin**) to support piping of output from another command to input of this command. An example of a user distinguished name is CN=CGreen,OU=IT,DC=Contoso,DC=Com. |
| **-upn** *UPN* | Specifies the new user principal names for the user objects that you want to modify, for example, cgreen@contoso.com. |
| **-fn** *FirstName* | Specifies the first names of the user names that you want to modify. |
| **-mi** *Initial* | Specifies the middle initials of the user objects that you want to modify. |
| **-ln** *LastName* | Specifies the last names of the user objects that you want to modify. |
| **-dn** *DisplayName* | Specifies the display names for the user objects that you want to modify. The display name of a user object is the name that you want to see when the user communicates using the account. The display name is also the name that other users search for when they look up this user object in Active Directory Domain Services (AD DS), Exchange, or any directory search tool. |
| **-email** *EmailAddress* | Specifies the e-mail addresses of the user object that you want to modify, which is often the same as the UPN. |

Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- If a value that you supply contains spaces, use quotation marks around the text, for example, "CN=Mike Danseglio,CN=Users,DC=Microsoft,DC=Com".

- If you supply multiple values for a parameter, use spaces to separate the values, for example, a list of distinguished names.

- To view the complete syntax for this command and for information about entering more user account information, at a command prompt, type the following command, and then press ENTER:

  **dsmod /?**

**See Also**
**Concepts**
Working with MMC console files
User and computer accounts
Object names
Install the Active Directory Schema snap-in

# Manage Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage groups**


- Create a new group
- Add a member to a group
- Allow anonymous users to be members of the Everyone security group on a domain controller
- Convert a group to another group type
- Change group scope
- Delete a group
- Find groups in which a user is a member
- Assign user rights to a group in Active Directory

# Manage Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage computers**

- Create a new computer account
- Add a computer account to a group
- Delete a computer account
- Manage a remote computer
- Move a computer account
- Reset a computer account
- Disable or enable a computer account

# Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security**
This section discusses Microsoft® Windows Server 2003 family security; how you can ensure that data access, storage, and exchange complies with your security needs.

This section covers:

- New features in security
- New ways to do familiar Security tasks
- Best practices for security
- Security overview
- Authentication
- Access Control
- Authorization Manager
- Security Configuration Manager
- Software Restriction Policies
- Auditing Security Events
- Encrypting File System
- Public Key Infrastructure
- Internet Protocol Security (IPSec)
- Windows Time Service
- Security Resources
- Security Troubleshooting

# New features in security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features in security**

| New features | More information |
|---|---|
| **Authorization Manager**. Authorization Manager gives application developers a flexible framework for integrating role-based access control into their applications, and it lets administrators who use those applications provide access in a natural, intuitive way: through assigned user roles relating to job function. | Authorization Manager |
| **Stored user names and passwords**. This utility provides secure storage for user names and credentials that are needed to access network or Internet resources. | Stored User Names and Passwords |
| **Software restriction policies**. This new security policy enables administrators to prevent software applications from running based on a hash algorithm of the software, the software's associated file path, the software publisher's certificate, or the Internet zone that hosts the software. | Software Restriction Policies |
| **Changes to existing technologies** | |
| **Certification authoritities**. Numerous improvements and added features are included for certification authorities. | New features in Certificate Services |
| **Constrained delegation**. With this new security feature, you can specify what services you want to trust for delegation for a server. | Delegating authentication |
| **Effective permissions tool**. This tool calculates the permissions that are granted to a specified user or group. | Effective Permissions tool |
| **Encrypting File System (EFS)** no longer requires a recovery agent. | Encrypting File System |
| **Everyone membership**. The built-in Everyone group includes Authenticated Users and Guests, but no longer includes members of the Anonymous group. | Differences in default security settings |
| **Operation-based auditing**. Operation based auditing offers more descriptive auditing events and provides you with the opportunity to choose what operations you want to audit when auditing object accesses. | Operation-based auditing on files or folders |
| **Reapplying security defaults**. This procedure enables you to easily reapply the default security settings for the Windows Server 2003 family. | Reapply default security settings |

# New ways to do familiar Security tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**
The following table lists common tasks for managing domain security. For most tasks, the user interface is different in Windows Server 2003 family than in Windows NT 4.0.

### Note

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0, use | In this version of Windows, use |
|---|---|---|
| Edit user rights for domain groups and users that apply to domain controllers | User Manager for Domains | Active Directory Users and Computers. For more information, see Define a security template. |
| Edit account policies for users in a domain | User Manager for Domains | Active Directory Users and Computers. For more information, see Define a security template. |
| Edit audit policies for groups and users that apply to domain controllers or specific computers | User Manager for Domains | Active Directory Users and Computers. For more information, see Define or modify auditing policy settings for an event category. |
| Establish domain trust relationships | User Manager for Domains | Active Directory Domains and Trusts. |
| Promote member server to a domain controller | Not available | Active Directory Installation Wizard |
| Assign permissions to files, folders, and registry keys | Appropriate object manager and **Properties** page | Appropriate object manager and **Properties** page. For more information, see the Access Control How To... page for Access Control. |
| Audit security events | User Manager for Domains | Appropriate object manager and **Properties** page. For more information, see Apply or modify auditing policy settings for a local file or folder. |
| View security log | Event Viewer | Event Viewer. For more information, see View the security log. |
| Create users and groups | User Manager for Domains | Active Directory Users and Computers. For more information, see Create a new user account and Create a new group. |
| Manage a Windows NT 4.0 domain | User Manager for Domains | User Manager for Domains |

# Best practices for security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Restrict physical access to computers, especially domain controllers, to trusted personnel**

- Physical access to a server is a high security risk. Physical access to a server by an intruder could result in unauthorized data access or modification as well as installation of hardware or software designed to circumvent security. To maintain a secure environment, you must restrict physical access to all servers and network hardware.

**For administrative tasks, use the principle of least privilege**

- Using the principle of least privilege, Administrators should use an account with restrictive permissions to perform routine, nonadministrative tasks and use an account with broader permissions only when performing specific administrative tasks.

- To accomplish this without logging off and back on, log on with a regular user account and use the Runas command to run the tools that require the broader permissions.

  For more information, see Run a program with administrative credentials and Runas.

**Define groups and their membership**

- When deciding on the default level of computer access that end users will have, the determining factor is the installed base of applications that need to be supported. If your organization uses only applications that belong to the Windows Logo for Software program, then you can make all your end users members of the Users group. If not, you may have to make your end users part of the Power Users group or relax the Users group's permissions security settings. Both are less secure options.

  Running legacy programs on Windows 2000 or Windows XP Professional often requires you to modify access to certain system settings. The same default permissions that allow Power Users to run legacy programs also make it possible for a Power User to gain additional privileges on the system, even complete administrative control. Therefore, it is important to deploy Windows 2000 or Windows XP Professional programs in the Windows Logo program for Software in order to achieve maximum security without sacrificing program functionality.

  For more information, see the Windows Logo program for Software page at the Microsoft Web site.

- All domain administrators (including child domain administrators) who can modify the system software on domain controllers in a forest must be equally trusted.

- Domain administrators and enterprise administrators should be the only users allowed to link a Group Policy object to an organizational unit. This is the default setting.

  Authenticated users only need Read and Apply Group Policy object permissions.

**Secure data on computers**

- Ensure that the system files and the registry are protected using strong access control lists.

  For more information, see Access Control.

- Use Syskey to provide additional protection of the security account manager (SAM), especially on your domain controllers.

For more information, see The system key utility.

**Use strong passwords throughout your organization**

- Most authentication methods require the user to provide a password to prove their identity. These passwords are normally chosen by the user, who may want a simple password that is easily remembered. In most cases, these passwords are weak and may be easily guessed or determined by an intruder. Weak passwords can circumvent this security element and become the weak point of an otherwise strong security environment. Strong passwords tend to be more difficult for an intruder to discern and, as a result, help provide an effective defense of your organization's resources.

  For more information, see Strong passwords.

**Do not download or run programs that come from untrusted sources**

- programs can contain instructions to violate security in a number of ways including data theft, denial of service and data destruction. These malicious programs often masquerade as legitimate software and can be difficult to identify. To avoid these programs, you should only download and run software that is guaranteed authentic and obtained from a trusted source. You should also ensure a current virus scanner is installed and functioning in case this type of software does inadvertently wind up on your computer.

**Keep virus scanners up to date**

- Virus scanners frequently identify infected files by scanning for a signature, which is a known component of a previously identified virus. The scanners keep these virus signatures in a signature file, which is usually stored on the local hard disk. Because new viruses are discovered frequently, this file should also be updated frequently for the virus scanner to easily identify all current viruses.

**Keep all software patches up to date**

- Software patches provide solutions to known security issues. Check software provider Web sites periodically to see if there are new patches available for software used in your organization.

  For more information and for the latest recommendations on security, see Security.

In order to subscribe to the product Security Notification Service, see "product Security Notification" at the Microsoft Web site

For more information, see:

- Best practices for permissions and user rights
- Auditing Security Events Best practices
- Encrypting File System best practices
- Software Restriction Policies Best Practices
- Certificate Services Best practices
- Certificate Templates Best practices

For more information on creating and maintaining security policy, see:

- Best practices for Security Configuration and Analysis

- Best practices for Security Templates
- Best practices for Security Settings

# Security overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security overview**
The primary features of the Microsoft® Windows Server 2003 family security model are user authentication and access control. The Active Directory® directory service ensures that administrators can manage these features easily and efficiently. The following sections describe these features of the security model.

**Authentication**

- *Interactive logon* confirms the user's identification to the user's local computer or Active Directory account.

- *Network authentication* confirms the user's identification to any network service that the user is attempting to access. To provide this type of authentication, the security system includes these authentication mechanisms: Kerberos V5, public key certificates, Secure Sockets Layer/Transport Layer Security (SSL/TLS), Digest, and NTLM (for compatibility with Windows NT® 4.0 systems).

  Single sign-on makes it possible for users to access resources over the network without having to repeatedly supply their credentials. For the Windows Server 2003 family, users need to only authenticate once to access network resources; subsequent authentication is transparent to the user.

Authentication in the Windows Server 2003 family also includes two-factor authentication, such as smart cards. For more information, see Smart Cards.

For more information, see Authentication protocols overview.

**Object-based access control**
Along with user authentication, administrators are allowed to control access to resources or objects on the network. To do this, administrators assign security descriptors to objects that are stored in Active Directory. A security descriptor lists the users and groups that are granted access to an object and the specific permissions assigned to those users and groups. A security descriptor also specifies the various access events to be audited for an object. Examples of objects include files, printers, and services. By managing properties on objects, administrators can set permissions, assign ownership, and monitor user access.

Not only can administrators control access to a specific object, they can also control access to a specific attribute of that object. For example, through proper configuration of an object's security descriptor, a user could be allowed to access a subset of information, such as employees' names and phone numbers but not their home addresses. For more information, see Access Control.

For more information, see "Authorization and access control" at the Microsoft Windows Resource Kits Web site.

**Security policy**
You can control security on your local computer or on multiple computers by controlling password policies, account lockout policies, Kerberos policies, auditing policies, user rights, and other policies. To create a systemwide policy, you can use security templates, apply templates using Security Configuration and Analysis or edit policies on the local computer, organizational unit, or domain. For more information, see Security Configuration Manager.

**Auditing**
Monitoring the creation or modification of objects gives you a way to track potential security problems, helps to ensure user accountability, and provides evidence in the event of a security breach. For more information, see Auditing overview.

**Active Directory and security**
Active Directory provides protected storage of user account and group information by using access control on objects and user credentials. Because Active Directory stores not only user credentials but also access control information, users who log on to the network obtain both authentication and authorization to access system resources. For example, when a user logs on to the network, the security system authenticates the user with information stored in Active Directory. Then, when the user attempts to access a service on the network, the system checks the properties defined in the discretionary access control list (DACL) for that service.

Because Active Directory allows administrators to create group accounts, administrators can manage system security more

efficiently. For example, by adjusting a file's properties, an administrator can permit all users in a group to read that file. In this way, access to objects in Active Directory is based on group membership.

**Data protection**
Stored data (online or offline) can be protected using:

- **Encrypting File System (EFS)**. EFS uses public key encryption to encrypt local NTFS data. For more information, see Storing Data Securely.

- **Digital signatures**. Digital signatures sign software components, ensuring their validity. For more information, see Digital signatures.

**Network data protection**
Network data within your site (local network and subnets) is secured by the authentication protocol. For an additional level of security, you can also choose to encrypt network data within a site. Using Internet Protocol security, you can encrypt all network communication for specific clients, or for all clients in a domain.

Network data passing in and out of your site (across intranets, extranets, or an Internet gateway) can be secured using the following utilities:

- Internet Protocol Security (IPSec) Encryption. A suite of cryptography-based protection services and security protocols.

- **Routing and Remote Access**. Configures remote access protocols and routing. For more information, see Routing Overview.

- **Internet Authentication Service (IAS)**. Provides security and authentication for dial-in users. For more information, see Internet Authentication Service.

**Public key infrastructure**
Public key encryption is an important aspect of security. For more information, see Deploying a Public Key Infrastructure.

**Trusts**
The Windows Server 2003 family supports domain trusts and forest trusts.

Domain trust allows a user to authenticate to resources in another domain. For more information on establishing and managing domain trust relationships, see Trust direction.

In a Windows Server 2003 forest, administrators can create a forest to extend two-way transitivity beyond the scope of a single forest to a second Windows Server 2003 forest. For more information, see Forest trusts.

For more information about authorization and access control, see "Domain Trust" at the Microsoft Windows Resource Kits Web site.

# Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication**


- Authentication protocols
- Smart Cards
- Stored User Names and Passwords
- Passwords

# Authentication protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication protocols**

This section covers the authentication protocols.

- For help with specific tasks, see Authentication Protocol How To....
- For general background information, see Authentication Protocol Concepts.

# Authentication Protocol How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Configure Kerberos policy
- Allow a computer to be trusted for delegation for specific services
- Allow a user to be trusted for delegation for specific services
- Allow a computer to be trusted for delegation
- Allow a user to be trusted for delegation
- Map an account from a trusted non-Windows kerberos realm to a user account

# Configure Kerberos policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Kerberos policy**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Do one of the following:
   - To edit an existing Group Policy object (GPO), in **Group Policy Object Links**, click the GPO that you want to edit, then click **Edit**.
   - To create a new GPO, click **New**, type the name of the new GPO, and then click **Edit**.
5. In the console tree, click **Kerberos Policy**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Account Policies/Kerberos Policy
6. In the details pane, double-click the Kerberos policy you want to modify.
7. Modify the policy, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Any modifications to Kerberos policy will affect all the computers in the domain.
- The Kerberos policies in the Windows Server 2003 family refer to the ticket-granting ticket as a **user ticket**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Kerberos Policy
Account and local policies

# Allow a computer to be trusted for delegation for specific services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow a computer to be trusted for delegation for specific services**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Computers**.

   **Where?**

   - *DomainName*/Computers

3. In the details pane, right-click the computer you want to trust for delegation and then click **Properties**.
4. On the **Delegation** tab, click **Trust this computer for delegation to specified services only**.
5. Do one of the following:

   - Confirm that **Use Kerberos only** is selected.
   - Click **Use any authentication protocol**.

6. Click **Add** and, in **Add Services**, click **Users and Computers**.
7. In **Enter the object names to select (examples)**, type the name of the user or computer that the computer will be trusted to delegate for, and then click **OK**.
8. In **Add Services**, click the service or services that will be trusted for delegation and click **OK**. Repeat this step as necessary.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- If you cannot see the **Delegation** tab, do one or both of the following:

  - Register a Service Principal Name (SPN) for the computer account using the **Setspn** utility in the support tools that are on your CD. Delegation is only intended to be used by service accounts, which should have registered SPNs, as opposed to a regular user account which typically does not have SPNs.
  - Raise the functional level of your domain to Windows Server 2003 . For more information, see Related Topics.

- Constrained delegation, delegation of authentication for only specified services, can only be enabled on a member of the

Windows Server 2003 family.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delegating authentication
Allow a computer to be trusted for delegation
Domain and forest functionality
Raise the domain functional level

# Allow a user to be trusted for delegation for specific services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow a user to be trusted for delegation for specific services**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.
3. Right-click the user you want to be trusted for delegation, and click **Properties**.
4. Click the **Delegation** tab and click **Trust this user for delegation to specified services only**.
5. Select **Use Kerberos only** (default) or select **Use any authentication protocol**, where the initial client-to-server authentication protocol can be any protocol supported by the Windows Server 2003 local security authority. Typically, this could be NTLM, SSL client certificate mapping or Digest.
6. Click **Add** and, in **Add Services** click **Users and Computers**.
7. In **Select Users or Computers**, enter the name of the user or computer that the user will be trusted to delegate for.
8. In **Add Services**, select the service or services that will be trusted for delegation and click **OK**. Repeat as necessary.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- If you cannot see the **Delegation** tab, do one or both of the following:
    - Register a Service Principal Name (SPN) for the user account with the **Setspn** utility in the support tools on your CD. Delegation is only intended to be used by service accounts, which should have registered SPNs, as opposed to a regular user account which typically does not have SPNs.
    - Raise the functional level of your domain to Windows Server 2003 . For more information, see Related Topics.
- Constrained delegation, which is trusting a computer for delegation of specified services, can only be enabled for a computer in a Windows Server 2003 domain.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delegating authentication
Allow a user to be trusted for delegation
Domain and forest functionality

Raise the domain functional level

# Allow a computer to be trusted for delegation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow a computer to be trusted for delegation**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Computers**.

   **Where?**

   - *DomainName*/Computers

3. In the details pane, right-click the computer you want to be trusted for delegation, and click **Properties**.
4. Do one of the following.

   - If in a Windows Server 2003 domain, on the **Delegation** tab, click **Trust this computer for delegation to any service (Kerberos only)**, and then click **OK**.
   - If in a Windows 2000 native domain, click the **General** tab, click **Trust this computer for delegation**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- If you cannot see the **Delegation** tab, do one or both of the following:

  - Register a Service Principal Name (SPN) for the computer account using the **Setspn** utility in the support tools that are on your CD. Delegation is only intended to be used by service accounts, which should have registered SPNs, as opposed to a regular user account, which typically does not have SPNs.
  - Raise the functional level of your domain to Windows Server 2003 . For more information, see Related Topics.

- For security reasons, do not allow servers on the enterprise network to perform delegation at will on behalf of any network connection.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delegating authentication

Allow a computer to be trusted for delegation for specific services
Domain and forest functionality
Raise the domain functional level

# Allow a user to be trusted for delegation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow a user to be trusted for delegation**

- In a Windows Server 2003 domain
- In a Windows 2000 native domain

**In a Windows Server 2003 domain**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - *DomainName*/Users

3. In the details pane, right-click the user you want to be trusted for delegation, and click **Properties**.
4. Click the **Delegation** tab, select the **Account is trusted for delegation** check box, and then click **OK**.

**Note**

- If you cannot see the **Delegation** tab, do one or both of the following:
  - Register a Service Principal Name (SPN) for the user account with the **Setspn** utility in the support tools on your CD. Delegation is only intended to be used by service accounts, which should have registered SPNs, as opposed to a regular user account which typically does not have SPNs.
  - Raise the functional level of your domain to Windows Server 2003 . For more information, see Related Topics.

**In a Windows 2000 native domain**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.
3. In the details pane, right-click the user you want to be trusted for delegation, and click **Properties**.
4. Click the **Accounts** tab, select the **Account is trusted for delegation** check box, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- For security reasons, do not allow servers on the enterprise network to perform delegation at will on behalf of any network connection.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delegating authentication
Allow a user to be trusted for delegation for specific services
Domain and forest functionality
Raise the domain functional level

# Map an account from a trusted non-Windows kerberos realm to a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To map an account from a trusted non-Windows kerberos realm to a user account**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Advanced Features**.
3. In the console tree, click **Users**.

   **Where?**

   - Active Directory User and Computers/*Domain Name*/Users

4. In the details pane, right-click the user account to which you want to map a certificate, and then click **Name Mappings**.
5. On the **Kerberos Names** tab, click **Add**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Mapping certificates to user accounts
Map a certificate to a user account

# Authentication Protocol Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

For more information on authentication protocols, see the following:

- Authentication protocols overview
- Introduction to authentication
- Kerberos V5 authentication
- Enabling Kerberos V5 authentication
- Delegating authentication
- Interoperability with RFC-1510 Kerberos implementations
- SSL/TLS authentication
- NTLM authentication
- Passport authentication

# Authentication protocols overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication protocols overview**

Authentication is a fundamental aspect of system security. It confirms the identity of any user trying to log on to a domain or access network resources. Windows Server 2003 family authentication enables single sign-on to all network resources. With single sign-on, a user can log on to the domain once, using a single password or smart card, and authenticate to any computer in the domain.

**Authentication types**

When attempting to authenticate a user, several industry-standard types of authentication may be used, depending on a variety of factors. The following table lists the types of authentication that the Windows Server 2003 family supports.

| Authentication protocols | Description |
|---|---|
| Kerberos V5 authentication | A protocol that is used with either a password or a smart card for interactive logon. It is also the default method of network authentication for services. |
| SSL/TLS authentication | A protocol that is used when a user attempts to access a secure Web server. |
| NTLM authentication | A protocol that is used when either the client or server uses a previous version of Windows. |
| Digest authentication | Digest authentication transmits credentials across the network as an MD5 hash or message digest. |
| Passport authentication | Passport authentication is a user-authentication service which offers single sign-in service. |

For more information on authentication types, see "Logon and Authentication" at the Microsoft Windows Resource Kits Web site.

# Introduction to authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introduction to authentication**
A key feature of authentication in the Windows Server 2003 family is its support of single sign-on. Single sign-on allows a user to log on to the domain once, using a single password, and authenticate to any computer in the domain.

Single sign-on provides two main security benefits:

- For a user, the use of a single password or smart card reduces confusion and improves work efficiency.
- For administrators, the amount of administrative support required for domain users is reduced, because the administrator only needs to manage one account per user.

Authentication, including single sign-on, is implemented as a two-part process: interactive logon and network authentication. Successful user authentication depends on both of these processes.

**Interactive logon**
Interactive logon confirms the user's identification to either a domain account or a local computer. This process is different, depending on the type of user account:

- With a domain account, a user logs on to the network with a password or smart card, using single sign-on credentials stored in the Active Directory directory service. By logging in with a domain account, an authorized user can access resources in the domain and any trusting domains. If a password is used to log on to a domain account, Kerberos V5 is used for authentication. If a smart card is used instead, Kerberos V5 authentication is used with certificates.
- With a local computer account, a user logs on to a local computer, using credentials stored in Security Account Manager (SAM), which is the local security account database. Any workstation or member server can store local user accounts, but those accounts can only be used for access to that local computer.

**Network authentication**
Network authentication confirms the user's identification to any network service that the user is attempting to access. To provide this type of authentication, the security system supports many different authentication mechanisms, including Kerberos V5, Secure Socket Layer/Transport Layer Security (SSL/TLS), and, for compatibility with Windows NT 4.0, NTLM.

Users who use a domain account do not see network authentication. Users who use a local computer account must provide credentials (such as a user name and password) every time they access a network resource. By using the domain account, the user has credentials that can be used for single sign-on.

For more information about authentication, see "Logon and Authentication" at the Microsoft Windows Resource Kits Web site.

# Kerberos V5 authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Kerberos V5 authentication**
Kerberos V5 is the primary security protocol for authentication within a domain. The Kerberos V5 protocol verifies both the identity of the user that is requesting authentication as well as server providing the requested authentication. This dual verification is also known as *mutual authentication*.

**An overview of how Kerberos V5 works**
The Kerberos V5 authentication mechanism issues tickets for accessing network services. These tickets contain encrypted data, including an encrypted password, that confirms the user's identity to the requested service. Except for entering a password or smart card credentials, the entire authentication process is invisible to the user.

An important service within Kerberos V5 is the Key Distribution Center (KDC). The KDC runs on each domain controller as part of the Active Directory directory service, which stores all client passwords and other account information.

The Kerberos V5 authentication process works as follows:

1. The user on a client system, using a password or a smart card, authenticates to the KDC.

2. The KDC issues a special ticket-granting ticket to the client. The client system uses this TGT to access the ticket-granting service (TGS), which is part of the Kerberos V5 authentication mechanism on the domain controller.

3. The TGS then issues a service ticket to the client.

4. The client presents this service ticket to the requested network service. The service ticket proves both the user's identity to the service and the service's identity to the user.

The Kerberos V5 services are installed on each domain controller, and a Kerberos client is installed on each workstation and server.

Every domain controller acts as a KDC. A client uses a Domain Name Service (DNS) lookup to locate the nearest available domain controller. That domain controller then functions as the preferred KDC for that user during the user's logon session. If the preferred KDC becomes unavailable, the system locates an alternate KDC to provide authentication.

For more information on how Kerberos V5 provides authentication, see Kerberos Authentication Technical Reference in the Windows Server 2003 Technical Reference.

For more information about Kerberos, see "Logon and Authentication" at the Microsoft Windows Resource Kits Web site.

# Enabling Kerberos V5 authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling Kerberos V5 authentication**

The Kerberos V5 authentication protocol is enabled by default for all computers joined to a Windows Server 2003 or Windows 2000 domain during installation. Kerberos provides single sign-on to resources within a domain, and to resources residing in trusted domains.

You can control certain aspects of Kerberos configuration through Kerberos security settings which are part of the account policies. For example, you can set the lifetime of users' Kerberos V5 tickets. As an administrator, you can use the default Kerberos policy or you can change it to suit the needs of your environment.

Successful authentication using Kerberos V5 requires that both the client and server computers must run the Windows 2000, Windows Server 2003 family, or Windows XP Professional operating system.

If a client system tries to authenticate to a server running another version of the Windows operating system, the NTLM protocol will be used as the authentication mechanism.

**Note**

- Computers using Kerberos for authentication must have their time settings synchronized with a common time service within five minutes or authentication fails. Computers running members of the Windows Server 2003 family, Windows XP Professional, or Windows 2000 automatically update the current time, using the domain controller as a network time service.

For more information about configuring Kerberos policy, see Kerberos Policy and Account Policies.

For more information about Kerberos, see "Logon and Authentication" at the Microsoft Windows Resource Kits Web site.

# Delegating authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Delegating authentication**

*Delegation* is the act of allowing a service to impersonate a user account or computer account in order to access resources throughout the network. When a service is *trusted for delegation*, that service can impersonate a user to use other network services.

For more information about delegation, see "Designing an Authentication Strategy" at the Microsoft Windows Resource Kits Web site.

To set up delegation on a computer, the following conditions must be met:

- The account doing the delegation must be set to **Trusted for delegation to any service** or **Trusted for delegation to specified services only**.
- The account that the service is delegating for must not have the **Account is sensitive and cannot be delegated** option chosen.
- An administrator must have the Enable computer and user accounts to be trusted for delegation privilege on the computer in order to enable delegation.

For more information about how to allow users or computers to be trusted for delegation, see Authentication Protocol How To....

**Differences in Windows Server 2003 family delegation**

| Windows 2000 | Members of the Windows Server 2003 family |
|---|---|
| Delegation requires the Kerberos authentication protocol. | Delegation can be used when the client authenticates to the service using protocols other then Kerberos because of a new feature of the Windows Server 2003 family, Protocol Transition. |
| The Key Distribution Center (KDC) embeds a copy of the user's ticket-granting ticket (TGT) inside the service ticket that is sent to the server. With this TGT, the server that has been trusted for delegation can request service tickets for the user to any other service on the network. | A ticket-granting ticket (TGT) is not needed. With protocol transition, a service can use a service ticket for delegation. |
| Delegation allows the server to connect to any resource in the domain on behalf of the client. | An administrator can specify which Service Principal Names (SPNs) an account is able to delegate to with constrained delegation. |

**Constrained delegation**

Constrained delegation is a new option and can only be used on members of the Windows Server 2003 family. With this option, the administrator can specify which Service Principal Names (SPNs) an account is able to delegate to. A service can be trusted for delegation, but that trust can be limited to a select group of services explicitly specified by a domain administrator.

In a Windows Server 2003 domain, there are three options to consider in trusting a computer for delegation:

- **Do not trust this computer for delegation**. This option was available in Windows 2000 and is the default setting for members of the Windows Server 2003 family.
- **Trust this computer for delegation to any service (Kerberos only)**. This option was available in Windows 2000. This setting is only as secure as the service and the actions of all its administrators. When this option is selected on a computer, all services under the Local System account on the computer will be trusted for delegation. This means an administrator on that computer may install any service and then that service will have the ability to access any network resource by impersonating a user.

- **Trust this computer for delegation to specified services only**. This is a new feature for the Windows Server 2003 family, and it is commonly referred to as constrained delegation. With constrained delegation, the administrator can specify which Service Principal Names (SPNs) this account is able to delegate to. If an application requires delegation for best functionality, this is the most secure option. Delegation for specified services allows an administrator to choose what services on the network can be delegated to by choosing a specific service or computer account. By only allowing delegation to specific services, an administrator can control what network resources the service or computer can use.

For more information on functional levels, see Domain and forest functionality.

For information on about configuring constrained delegation, see

- Allow a computer to be trusted for delegation for specific services
- Allow a user to be trusted for delegation for specific services

For information on about configuring delegation, see

- Allow a computer to be trusted for delegation
- Allow a user to be trusted for delegation

# Interoperability with RFC-1510 Kerberos implementations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interoperability with RFC-1510 Kerberos implementations**
Kerberos V5 is designed for interoperability with other security services that are based on the RFC-1510 Kerberos reference implementation.

Windows Server 2003 family supports two types of Kerberos V5 interoperability:

- A trust relationship can be established between a domain and an MIT-based Kerberos realm. This means that a client in a Kerberos realm can authenticate to an Active Directory directory service domain to access network resources in that domain.

- Within a domain, UNIX clients and servers can have Active Directory accounts and therefore obtain authentication from a domain controller.

Kerberos V5 interoperates with other RFC-1510 implementations of the Kerberos protocol in the following ways:

- A domain controller can provide authentication for client systems running implementations of RFC-1510 Kerberos, including clients running an operating system other than Windows XP Professional or Windows 2000.

- Windows XP Professional systems can authenticate to a Kerberos V5 server within a realm, with single sign-on to both the server and a local Windows XP Professional account.

- Client applications for Win32 and operating systems other than Windows XP Professional that are based on the General Security Service Application program Interface (GSS API) can obtain session tickets for services within a domain.

For more information, see the Step-by-Step Guide to Kerberos 5 (krb5 1.0) Interoperability on the Microsoft Web site and the Kerberos Network Authentication Service (V5).

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# SSL/TLS authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SSL/TLS authentication**

For conceptual information about certificates and Transport Layer Security (TLS)/Secure Socket Layer (SSL) connections and authentication, see Certificates Technical Reference and TLS/SSL Technical Reference in the Windows Server 2003 Technical Reference.

With Secure Socket Layer/Transport Layer Security (TLS/SSL), you have the ability to control what ciphers or cryptographic algorithms are allowed. By default, all ciphers can be used; however, you can control what ciphers your computer can use. For more information, see article Q216482, "How to Control the Ciphers for SSL and TLS," in the Microsoft Knowledge Base (http://www.microsoft.com/).

For instructions on how to map certificates to user accounts in Internet Information Services (IIS), see "Mapping Client Certificates to User Accounts" in the Internet Information Services (IIS) 6.0 documentation.

**Note**

- You must have IIS installed to have access to the documentation.

# NTLM authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NTLM authentication**

In a networking environment, NTLM is used as the authentication protocol for transactions between two computers where one or both computers are running Windows NT 4.0 or earlier. Networks with this configuration are referred to as mixed-mode, the default setting in the Windows Server 2003 family.

As examples, the following configurations would use NTLM as the authentication mechanism:

- A Windows 2000 or Windows XP Professional client authenticating to a Windows NT 4.0 domain controller.
- A Windows NT 4.0 Workstation client authenticating to a Windows 2000 or Windows Server 2003 domain controller.
- A Windows NT 4.0 Workstation client authenticating to a Windows NT 4.0 domain controller.
- Users in a Windows NT 4.0 domain authenticating to a Windows 2000 or a domain running Windows Server 2003 family.
- A client running Windows 95, Windows 98, or Windows Millennium Edition that is authenticating to any domain controller.

In addition, NTLM is the authentication protocol for computers that are not participating in a domain, such as stand-alone servers and workgroups.

For more information about NTLM and NTLMv2, see "Logon and Authentication" at the Microsoft Windows Resource Kits Web site. (http://www.microsoft.com/)

# Passport authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Passport authentication**
Passport authentication is a centralized service provided by Microsoft that offers single logon and core profile services for member Web sites. This benefits Web site users by making it unnecessary to log on again to access additional protected resources or Web sites. You can make your business Web site compatible with Passport authentication and authorization by using the Passport authentication provider.

The Microsoft .NET Passport Home Web page provides information for:

- Consumers, who want to manage their .NET Passports and change their passwords.
- Businesses, which can license and use .NET Passport on their Web sites.

For more information, see the Microsoft .NET Passport Home Web page at the Microsoft Web site.

Web site administrators can use the Microsoft Passport Manager Administration utility to change registry settings that affect the default behavior of the Passport Manager object. For information about the Passport Manager Administration utility, see the MSDN Web site. For information about the Microsoft .NET Passport Software Development Kit, see the MSDN Web site.

**Note**

- To start the Passport Manager Administration utility, click **Start**, click **Run**, type **msppcnfg.exe**, and then click **OK**.

# Smart Cards

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Smart cards**
Smart cards are a tamper-resistant and portable way to provide security solutions for tasks such as client authentication, logging on to a Microsoft® Windows Server 2003 family domain, code signing and securing e-mail.

- Before deploying smart cards, see Checklist: Deploying smart cards for logging on to Windows.
- For help with specific tasks, see Smart Card How To....
- For general background information, see Smart Card Concepts.

# Checklist: Deploying smart cards for logging on to Windows

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Deploying smart cards for logging on to Windows**

| Step | Reference |
|---|---|
| Review concepts about smart cards. | Smart Card Smart Card Concepts |
| **Set up certification authorities (CAs) and the certification hierarchy** | |
| If you have not already done so, follow the checklist for deploying certification authorities and public key infrastructure (PKI) on an intranet. | Checklist: Creating a certification hierarchy with an offline root certification authority |
| **For each domain** | |
| Set the security permissions of the Smart Card User, Smart Card Logon, and Enrollment Agent certificate templates to allow smart card users to enroll for certificates. | Allow subjects to request a certificate that is based on the template |
| **For each certification authority issuing smart card certificates** | |
| Set up the certification authority to issue smart card certificates. | Prepare a certification authority to issue smart card certificates |
| If you did not do so in the previous step, set up a certification authority to issue Enrollment Agent certificates. | To establish the certificate types that an enterprise certification authority can issue |
| **Set up a smart card enrollment station** | |
| Install a smart card reader. | Install a smart card reader on a computer |
| Get an Enrollment Agent certificate for the persons who will be setting up smart cards for users. | Prepare a smart card certificate enrollment station |
| **For each smart card user** | |
| Set up smart cards. | Set up a smart card for user logon |
| Install a smart card reader on the smart card user's computer. | Install a smart card reader on a computer |
| Review smart card logon procedure. | Log on to a computer with a smart card |

# Smart Card How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to ...**

- Install a smart card reader on a computer
- Log on to a computer with a smart card
- Administer Smart Cards

# Install a smart card reader on a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a smart card reader on a computer**

1. Depending on the type of reader you have purchased, attach your reader to an available serial or USB port or insert the PC card reader into an available PCMCIA slot.

2. If you connected a serial port reader, restart the computer to detect it and then log on to the computer.

3. Do one of the following:

   - If the device driver for the smart card reader is available as part of Windows Server 2003 family, the installation of the driver takes place without any prompting by the user. This may take a few minutes.

   - If the device driver for the smart card reader is not available as part of Windows Server 2003 family, the Add Hardware Wizard starts. Follow the directions for installing the device driver software.

### Important

- Smart card readers generally come with setup instructions from the manufacturer. If your reader comes with instructions, you should use those. If not, then use these general procedures.

### Notes

- You can confirm that the installation was successful by confirming that the reader appears in the **Device Manager** tree.

- If the smart card reader is not installed automatically or the Add Hardware Wizard does not start automatically, then your smart card reader may not be compliant with Plug and Play specifications. You should contact the smart card reader manufacturer for the device driver and the instructions on how to install and configure the device.

- When using the Add Hardware Wizard to install a smart card reader, you need the device driver. Contact the manufacturer of the smart card reader. Alternatively, your administrator may provide the driver.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Smart card readers
Supported smart cards

# Log on to a computer with a smart card

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To log on to a computer with a smart card**

1. At the Windows Server 2003 family logon screen, insert your smart card in the smart card reader.
2. Type the personal identification number (PIN) for the smart card when prompted by your computer.

## Notes

- You must prepare the smart card by creating the appropriate credentials before using it to log on to the computer. For more information on preparing a smart card, see Related Topics.
- If the PIN you enter is recognized as legitimate, this logs you on to the computer and to the Windows Server 2003 family domain, based on the permissions assigned to your user account by the domain administrator.
- If you enter the incorrect PIN for a smart card several times in a row, you will be unable to log on to the computer using that smart card. The number of allowable invalid logon attempts before lockout occurs varies according to the smart card manufacturer. Contact your administrator for assistance.
- If you insert the smart card backwards or upside down, the smart card reader will not work. However, you may be prompted for a PIN which will not work. Follow the directions for your specific model of card reader to understand which way to insert the smart card.
- Smart card logon only works for computers that are joined to a domain.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Deploying smart cards for logging on to Windows
Install a smart card reader on a computer
Supported smart cards
Smart card readers
Understanding smart cards

# Administer Smart Cards

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer smart cards**

- Prepare a certification authority to issue smart card certificates
- Prepare a smart card certificate enrollment station
- Set up a smart card for user logon
- Configure a smart card user to temporarily log on to the network without a smart card

# Prepare a certification authority to issue smart card certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prepare a certification authority to issue smart card certificates**

1. Confirm that the proper security permissions are set on the Smart Card Logon, Smart Card User, and Enrollment Agent certificate templates.

2. Log on with administrator rights to the certification authority (CA) you will use to issue smart card certificates.

3. Open Certification Authority.

4. In the console tree, click **Certificate Template**.

   **Where?**

   - Certification Authority (*Computer*)/*CA_Name*/Certificate Templates

5. Do one of the following:

   - If the CA will issue certificates that are only for logging on to Windows with smart cards, click the **Smart Card Logon** certificate template, and then click **OK**.

   - If the CA will issue certificates that can be used to log on to Windows with smart cards and other uses, click the **Smart Card User** certificate template, and then click **OK**.

6. On the **Action** menu, point to **New**, and then click **Certificate to Issue**.

7. Click the **Enrollment Agent** certificate template, and then click **OK**.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

- The security permission setting of a certificate template indicates who is allowed to request a certificate of that type.

- The Enrollment Agent certificate does not have to be issued from the same certification authority (CA) that will issue certificates for smart cards, as takes place in this procedure, but the issuing CA for the Enrollment Agent certificate must be a trusted enterprise CA in the domain. In that case, make sure that there is an enterprise CA in your domain that is capable of issuing Enrollment Agent certificates. To do this, follow the above steps for the CA that will issue certificates that will be used for logging on to Windows with a smart card and other uses. Then follow the above steps to issue an enrollment agent certificate.

- This procedure only applies to enterprise CAs.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Enrolling for a smart card certificate
Install a smart card reader on a computer
Set up a smart card for user logon
Prepare a smart card certificate enrollment station
Log on to a computer with a smart card
Working with MMC console files

# Prepare a smart card certificate enrollment station

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prepare a smart card certificate enrollment station**

1. On the computer that you will use to set up smart cards, install a smart card reader, following the manufacturer's instructions.
2. Log on as the user or administrator who will be installing certificates on smart cards.
3. On the taskbar, click the **Start** button, click **Run**, type **mmc**, and then click **OK**.
4. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
5. In **Snap-in**, double-click **Certificates**. If you are logged on as a user, the Certificates snap-in automatically loads.
   - If you are logged on as an Administrator, click **My user account**, and then click **Finish**.
6. Click **Close** and then click **OK**.
7. Double-click **Certificates - Current User**.
8. In the console tree, click **Personal**.

   **Where?**

   - Certificates - Current User/Personal
9. On the **Action** menu, point to **All Tasks**, and then click **Request New Certificate**.
10. In the Certificate Request Wizard, click the **Enrollment Agent** certificate template and type a friendly name and a description for the certificate.
11. When prompted by the Certificate Request Wizard, click **Install Certificate**.

**Notes**

- The Web enrollment station can also be used instead of the method described in this procedure. For more information, see Related Topics.
- Before requesting smart card logon certificates for users, a smart card administrator must have an Enrollment Agent certificate available to generate smart card certificate requests on the behalf of others. That is the purpose of this procedure.
- To perform this procedure, you must have the appropriate security permissions to administer the Enrollment Agent certificate template. See Related Topics for more information about Enrollment Agent certificates and enrolling for a smart card certificate.
- The Enrollment Agent certificate can be installed on a smart card, if you want. In that case, you must use the smart card manufacturer's cryptographic service provider (CSP) when requesting the certificate. (Click the **Advanced Options** button in the Certificate Request Wizard to select a smart card CSP for the Enrollment Agent certificate.)

- These tasks can be performed on any computer running Windows XP or a member of the Windows Server 2003 family that you want to use as a smart card certificate enrollment station.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enrolling for a smart card certificate
Prepare a certification authority to issue smart card certificates
Set up a smart card for user logon
Log on to a computer with a smart card
Manage certificates for your user account
Working with MMC console files

# Set up a smart card for user logon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up a smart card for user logon**

1. Log on as an enrollment agent for the domain where the user's account is located.
2. Open Internet Explorer.
3. In **Address**, type the address of the certification authority (CA) that issues smart card certificates, and then press ENTER.
4. Click **Request a certificate** and then click **advanced certificate request**.
5. Click **Request a certificate for a smart card on behalf of another user using the smart card certificate enrollment station**. If you are prompted to accept the smart card signing certificate, click **Yes**.
6. On the **Smart Card Certificate Enrollment Station** Web page, in **Certificate Template**, do one of the following:
   - Click **Smart Card Logon** if you want to use the smart card for logging on to Windows only.
   - Click **Smart Card User** if you want to use the smart card for secure e-mail as well as logging on to Windows.
7. In **Certification Authority**, click the name of the CA you want to issue the smart card certificate.
8. In **Cryptographic Service Provider**, select the cryptographic service provider (CSP) of the smart card's manufacturer.
9. In **Administrator Signing Certificate**, click the Enrollment Agent certificate that will sign the enrollment request.
10. In **User To Enroll**, click **Select User**, select the appropriate user account, and then click **Enroll**.
11. When prompted by the system, insert the smart card into the smart card reader on your computer, click **OK**, and then, when prompted by the system, enter the personal identification number (PIN) for the smart card.
12. (Optional) If the smart card you are setting up has a previously installed certificate on it, a message appears, asking whether you want to replace the existing credentials on the card. Click **Yes**.
13. After the certificate is installed on the smart card, the CA Web page will give you the option of viewing the certificate you just installed or beginning a new smart card certificate request.

**Notes**

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- In the first step, anyone in the domain who has an Enrollment Agent certificate and has security permissions to issue smart card certificates is considered an enrollment agent.
- The address of the certification server is the name of the server followed by **/Certsrv**. For example, in order to connect to the CA on a server named SmartcardCA, you would connect to: http://SmartcardCA/Certsrv

  Be sure to use the name of the server that the CA is installed on, not the CA name itself. In many cases, these names will be different.

- If you have no Enrollment Agent certificate available, see Related Topics.

- Users that will be logging on to computers running the Windows 2000 operating system must have a smart card enrolled from a computer running Windows 2000. Users that will log on to computers running Windows XP or Windows Server 2003, Standard Edition can have a smart card enrolled from a computer running any of these operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enrolling for a smart card certificate
Supported smart cards
Smart card readers
Prepare a certification authority to issue smart card certificates
Prepare a smart card certificate enrollment station
Certification authority Web enrollment services

# Configure a smart card user to temporarily log on to the network without a smart card

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a smart card user to temporarily log on to the network without a smart card**

1. Open Active Directory Users and Computers.
2. In the console tree, do one of the following.
   - Click **Users**.

     **Where?**

   - Active Directory Users and Computers/*Domain Node*/Users
   - Click the folder that contains the user account that you want.
3. Right-click the user account, and then click **Properties**.
4. Click the **Account** tab.
5. In the **Account options** list, do the following:
   - Clear the **Smart card is required for interactive logon** check box.
   - Select the **User must change password at next logon** check box.
6. Right-click the user account, and then click **Reset Password**.
7. Provide a new password for this user.
8. Inform the user of the new password and that the password must be changed the next time he or she logs on to the network.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- The **User must change password at next logon** option is not required for this procedure to work, but it ensures that only the user has the correct password.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Logging on to a computer with a smart card
Set up a smart card for user logon
Create a new forest

# Smart Card Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general information about smart cards:

- Smart cards overview
- Understanding smart cards
- Using Smart Cards

# Smart cards overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Smart cards overview**
Smart cards are a tamper-resistant and portable way to provide security solutions for tasks such as logging on to a Microsoft® Windows Server 2003 family domain, client authentication, code signing, and securing e-mail.

Support for cryptographic smart cards is a key feature of the public key infrastructure (PKI) that Microsoft has integrated into Microsoft® Windows® XP and the Windows Server 2003 family.

Smart cards provide:

- Tamper-resistant storage for protecting private keys and other forms of personal information.
- Isolation of security-critical computations involving authentication, digital signatures, and key exchange from other parts of the computer that do not have a "need to know." These operations are all performed on the smart card.
- Portability of credentials and other private information between computers at work, home, or on the road.

For more information, see Understanding smart cards and Using Smart Cards.

For smart card administrative and user procedures, see Smart Card How To....

For the smart card deployment checklist, see Checklist: Deploying smart cards for logging on to Windows.

# Understanding smart cards

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding smart cards**

Logging on to a network with a smart card provides a strong form of authentication because it uses cryptography-based identification and proof of possession when authenticating a user to a domain.

For example, if a malicious person obtains a user's password, that person can assume the user's identity on the network simply through use of the password. Many people choose passwords they can remember easily, which makes passwords inherently weak and open to attack.

In the case of smart cards, that same malicious person would have to obtain both the user's smart card and the personal identification number (PIN) to impersonate the user. This combination is obviously more difficult to attack because an additional layer of information is needed to impersonate a user. An additional benefit is that, after a small number of unsuccessful PIN inputs occur consecutively, a smart card is locked, making a dictionary attack against a smart card extremely difficult. (Note that a PIN does not have to be a series of numbers; it can also use other alphanumeric characters.) Smart cards are also resistant to undetected attacks because the card needs to be obtained by the malicious person, which is relatively easy for a user to know about.

For more conceptual information, see Using Smart Cards.

For smart card administrative and user procedures, see Smart Card How To....

# Using Smart Cards

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using smart cards**

This section provides general information about using smart cards:

- Smart card readers
- Supported smart cards
- Enrolling for a smart card certificate
- Logging on to a computer with a smart card
- Using smart cards with system tools

# Smart card readers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Smart card readers**

It is strongly recommended that only smart card readers that have been tested by the Microsoft Windows Hardware Quality Labs and have obtained the Windows-compatible logo be installed on Windows XP and Windows Server 2003 family computers.

Smart card readers that are not Plug and Play-compliant are not recommended on computers running Windows XP or the Windows Server 2003 family. If you are using such a reader, then you must obtain installation instructions (including associated device driver software) directly from the manufacturer of the smart card reader. Microsoft neither supports nor recommends the use of non-Plug and Play smart card readers.

For instructions on installing smart card readers, refer to the manufacturer's materials that are provided with the reader, or see Install a smart card reader on a computer.

The following smart card readers are supported by Windows XP and Windows Server 2003 family. Their drivers are installed only upon detection of the corresponding Plug and Play smart card reader hardware.

| Brand | Smart card reader | Interface | Device driver |
|---|---|---|---|
| American Express | GCR435 | USB | Grclass.sys |
| Bull | SmarTLP3 | Serial | Bulltlp3.sys |
| Compaq | Serial reader | Serial | grserial.sys |
| Gemplus | GCR410P | Serial | Grserial.sys |
| Gemplus | GPR400 | PCMCIA | Gpr400.sys |
| Gemplus | GemPC430 | USB | Grclass.sys |
| Hewlett Packard | ProtectTools | Serial | Scr111.sys |
| Litronic | 220P | Serial | Lit220p.sys |
| Schlumberger | Reflex 20 | PCMCIA | Pscr.sys |
| Schlumberger | Reflex 72 | Serial | Scmstcs.sys |
| Schlumberger | Reflex Lite | Serial | Scr111.sys |
| SCM Microsystems | SCR111 | Serial | Scr111.sys |
| SCM Microsystems | SCR200 | Serial | Scmstcs.sys |
| SCM Microsystems | SCR120 | PCMCIA | Pscr.sys |
| SCM Microsystems | SCR300 | USB | Stcusb.sys |
| Systemneeds | External | Serial | Scr111.sys |
| Omnikey AG | 2010 | Serial | Sccmn50m.sys |
| Omnikey AG | 2020 | USB | Sccmusbm.sys |
| Omnikey AG | 4000 | PCMCIA | Cmbp0wdm.sys |

# Supported smart cards

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Supported smart cards**

When you install Windows XP or the Windows Server 2003 family, support for the following smart cards is included in the default installation.

| Manufacturer | Model |
|---|---|
| Gemplus | GemSAFE 4k |
| Gemplus | GemSAFE 8k |
| Infineon | SICRYPT v2 |
| Schlumberger | Cryptoflex 4k |
| Schlumberger | Cryptoflex 8k |
| Schlumberger | Cyberflex Access 16k |

You do not need to configure anything on the client or server to use either of these cards. Cryptographic smart cards can only be obtained directly from the respective companies and not from Microsoft Corporation.

While support for these cards is included in Windows XP and the Windows Server 2003 family, other RSA-based cryptographic smart cards will also work with Windows XP and the Windows Server 2003 family, provided the card vendor has developed a cryptographic service provider (CSP) for the card using CryptoAPI and the Smart Card Software Developer's Kit, which is available through Microsoft Developer's Network (MSDN).

Smart card personal identification numbers (PINs) can be changed anytime the CSP displays the private key PIN dialog box. PIN management is the responsibility of the card CSP and the user. Windows XP and the Windows Server 2003 family do not manage PINs. For PIN management tools and instructions, contact the smart card vendor.

# Enrolling for a smart card certificate

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enrolling for a smart card certificate**

A domain user cannot enroll for a Smart Card Logon certificate (which provides authentication) or a Smart Card User certificate (which provides authentication plus other uses of the smart card cryptography) unless a system administrator has granted the user access rights to the certificate template that is stored in Active Directory. Enrollment for a smart card certificate must be a controlled procedure, in the same manner that employee badges are controlled for purposes of identification and physical access.

The recommended method for enrolling users for smart card-based certificates and keys is through the smart card enrollment station that is integrated with Certificate Services in Windows Server 2003, Standard Edition and Windows Server 2003, Enterprise Edition.

When an enterprise certification authority (CA) is installed, the installation includes the Smart Card Enrollment station. This allows an administrator to act on behalf of a user to request and install a Smart Card Logon certificate or Smart Card User certificate on the user's smart card. Prior to using the Smart Card Enrollment station, the smart card issuer must have obtained a signing certificate based on the Enrollment Agent certificate template. The signing certificate signs the certificate request that is generated on behalf of the smart card recipient.

By default, only domain administrators are granted permission to request a certificate based on the Enrollment Agent template. A user other than a domain administrator can be granted permission to enroll for an Enrollment Agent certificate by means of Active Directory Sites and Services.

**Important**

- Once someone has an Enrollment Agent certificate, they can enroll for a certificate and generate a smart card on behalf of anyone in the organization. The resulting smart card could then be used to log on to the network and impersonate the real user.

Because of the powerful capability of the Enrollment Agent certificate, it is strongly recommended that your organization maintain very strong security policies over who has one. A scenario to minimize risk of Enrollment Agent certificate misuse would be to have one subordinate CA with very tight administrative controls in your organization that is only used to issue Enrollment Agent certificates. Once the initial Enrollment Agent certificates have been issued, the administrator of the CA can disable the issuance of Enrollment Agent certificates until they are needed again. By restricting the administrators who can operate the CA service on the subordinate CA, the service can be kept online for the generation and distribution of certificate revocation lists (CRLs) if necessary. Other CAs in the hierarchy can still issue Enrollment Agent certificates if their policy settings are changed, but you can determine whether inappropriate Enrollment Agent certificates are issued by checking the Issued Certificates log for each CA regularly.

For information on how to change the security permissions for a certificate template, see Allow subjects to request a certificate that is based on the template.

The enrollment station does not provide any card personalization functions, such as creating a file structure or setting the personal identification number (PIN), because those are card-specific functions and can only be done using specialized software provided by the smart card manufacturer.

For instructions on enrolling users for smart card certificates, see Set up a smart card for user logon. Note that users that log on to computers running a Windows 2000 operating system must have a smart card enrolled from a computer running Windows 2000. Users that log on to computers running Windows XP or Windows Server 2003, Standard Edition can have a smart card enrolled from a computer running any of these operating systems.

# Logging on to a computer with a smart card

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logging on to a computer with a smart card**

To log on to a computer with a smart card, users do not need to type CTRL+ALT+DEL. They simply insert the smart card into the smart card reader, and the computer prompts them for their personal identification number (PIN) instead of their user name and password (and, if applicable, domain).

For smart card logon procedures, see Log on to a computer with a smart card. For a list of procedures to enable logging on with a smart card, see Checklist: Deploying smart cards for logging on to Windows.

# Using smart cards with system tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using smart cards with system tools**

In previous versions of Windows, smart cards were primarily used for user logon. In Windows XP and the Windows Server 2003 family, smart cards can be used with a number of system tools.

For making network connections, the user interface feature for mapping network drives can now use smart cards. The corresponding command-line tool **Net** also has the ability to use smart cards to provide credentials. For more information about **Net**, see Net use. For more information about mapping network drives, see To assign a drive letter to a network computer or folder.

In a secure environment, administrators will log on with normal user accounts for most of their regular work. When administrative functions need to be performed, a separate user account's credentials must be used. This is easily achieved with the **Runas** command. Smart cards can be used to provide the credentials to the **Runas** command, providing further enhanced levels of security. For more information about **Runas**, see Runas.

Domain administrators may also use smart cards to store their credentials. One operation that requires very secure credentials is installing Active Directory. This operation can accept credentials supplied from smart cards as well as the traditional user name and password. For more information, see Using the Active Directory Installation Wizard.

# Stored User Names and Passwords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Stored User Names and Passwords**
Stored User Names and Passwords is a feature of Microsoft® Windows® XP and the Windows Server 2003 family that stores user names and passwords for servers. This allows a user to connect to servers using user names and passwords that are different than those used to log on to the network. The user can store these for later reuse.

- Before using Stored User Names and Passwords, see Checklist: Using Stored User Names and Passwords safely.
- For tips about using Stored User Names and Passwords, see Stored User Name and Password Best practices.
- For help with specific tasks, see Stored User Name and Password How To....
- For general background information, see Stored User Name and Password Concepts.
- For problem-solving instructions, see Stored User Name and Password Troubleshooting.

# Checklist: Using Stored User Names and Passwords safely

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Using Stored User Names and Passwords safely**

| Step | Reference |
| --- | --- |
| Ensure strong passwords are used when creating accounts for remote resources | Strong passwords |
| Ensure strong passwords are used when creating local computer or domain accounts | Strong passwords |
| Secure your computer whenever not in use | Strong passwords |
| Use a different strong password for all resources | Strong passwords |

# Stored User Name and Password Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Protect your account**

- Storing user names and passwords as part of your user account information means that anyone who has access to your account can access that stored information. Users must protect their account from misuse. Locking or logging off of computers when they not attended, providing password protection to screen savers, and ensuring the use of strong passwords are elements that contribute to the protection of the account.

  For more information on locking a computer, see To lock a computer in a domain environment.

**Use strong passwords for all accounts**

- The risk of using weak passwords is not removed by Stored User Names and Passwords. To protect against this, use strong passwords to protect against an intruder guessing or stealing the password.

  For more information, see Strong passwords.

**Change passwords regularly**

- Although a strong password can help protect against intruders, it is possible to eventually guess or steal the password of a resource. For this reason, passwords should be changed periodically. This not only minimizes chances that an intruder can guess or steal a password through lengthy password-guessing processes, but also minimizes damages when a password is compromised without the user's knowledge.

  For more information, see Change an existing password.

**Use different passwords for individual accounts**

- Stored User Names and Passwords allows the user to supply a different user name and password to each computer that is used. This provides for additional security on a per-computer basis. Using different passwords for each computer helps ensure that one guessed or stolen password does not weaken security. If an intruder is able to determine one password that is stored by Stored User Names and Passwords and all other passwords are different, the intruder is limited to the damage they can do with that single password.

**Store user names and passwords only when appropriate**

- User name and passwords that are used to access extremely sensitive resources must be protected carefully. In these cases, store the user names and passwords only for the specific logon session by choosing the appropriate option in **Logon Information Properties**.

# Stored User Name and Password How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Store a new user name and password
- Store default Microsoft .NET Passport information
- View and modify stored user names and passwords
- Delete stored user names and passwords
- Change a domain password

# Store a new user name and password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To store a new user name and password**

1. Open Stored User Names and Passwords.
2. Click **Add**, and then type the information in the spaces that are provided.

## Notes

- You must be logged on as a computer administrator to complete this procedure.
- To open Stored User Names and Passwords, click **Start**, click **Control Panel**, and then double-click **Stored User Names and Passwords**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View and modify stored user names and passwords
Delete stored user names and passwords

# Store default Microsoft .NET Passport information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To store default Microsoft .NET Passport information**

1. Open .NET Passport Wizard.
2. Provide the information requested by the .NET Passport Wizard.

## Notes

- You must be logged on as a computer administrator to complete this procedure.
- You may also encounter the .NET Passport Wizard when accessing Web sites that support Passport information.
- You can only store one .NET Passport user name and password per user at a time.
- Passports can only be created or associated with your user account through this wizard. You can only change the .NET Passport properties on the .NET Passport Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View and modify stored user names and passwords
Delete stored user names and passwords

# View and modify stored user names and passwords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view and modify stored user names and passwords**

1. Open Stored User Names and Passwords.
2. Click the entry in the list that you want to work with, and then click **Properties**.

## Notes

- You must be logged on as a computer administrator to complete this procedure.
- To open Stored User Names and Passwords, click **Start**, click **Control Panel**, and then double-click **Stored User Names and Passwords**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Store a new user name and password
Delete stored user names and passwords

# Delete stored user names and passwords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete stored user names and passwords**

1. Open Stored User Names and Passwords.
2. Click the entry that you want, and then click **Remove**.

## Notes

- You must be logged on as a computer administrator to complete this procedure.
- To open Stored User Names and Passwords, click **Start**, click **Control Panel**, and then double-click **Stored User Names and Passwords**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Store a new user name and password
View and modify stored user names and passwords

# Change a domain password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a domain password**

1. Open Stored User Names and Passwords.
2. Click an entry that is in the domain that you want to change the password for, and then click **Properties**.
3. Click **Change**, and then supply the necessary information.

## Notes

- You must be logged on as a computer administrator to complete this procedure.
- To open Stored User Names and Passwords, click **Start**, click **Control Panel**, and then double-click **Stored User Names and Passwords**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Store a new user name and password
Delete stored user names and passwords

# Stored User Name and Password Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Stored User Names and Passwords:

- Stored User Names and Passwords overview
- Benefits of Stored User Names and Passwords
- How Stored User Names and Passwords works

# Stored User Names and Passwords overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Stored User Names and Passwords overview**

When you log on to a computer running an operating system in the Microsoft® Windows Server 2003 family, you can supply a user name and password. This becomes your default security context for connecting to other computers on networks and over the Internet. However, this user name and password may not provide access to all of your resources. Stored User Names and Passwords provides a way to store these as a part of your profile.

There may be cases where you want to use different names and passwords for connecting to different resources. Examples could include:

- You want to log on to your computer with a standard account, but connect to certain computers as an administrator for maintenance and troubleshooting reasons.
- You work at home and want to use your work user name and password to connect to work-related servers.
- Your account is in a domain and you need access to computers in an untrusted domain.
- You want to access Web sites with user names and passwords that are specific to each of those sites.

For example, administrators may log on to the network using their standard user name and password but need to connect to a remote server with administrative access to perform specific functions. In this case, the user must be able to supply a different user name and password for this connection. The user may also want to store this user name and password for reuse at a later date. This is the function of Stored User Names and Passwords.

A user may also need to connect to secure Web servers using a specific user name and password. Stored User Names and Passwords allows users to connect to different Web servers using supplied user names and passwords and store them for future reuse. The user names and passwords can be either specific to a unique Web server or they can be generic so that they will be supplied when the user attempts to log on to a secured Web server.

Stored User Names and Passwords also stores saved information as part of a user's profile. This means that these user names and passwords will travel with the user from computer to computer anywhere on the network.

# Benefits of Stored User Names and Passwords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Benefits of Stored User Names and Passwords**

The benefits of Stored User Names and Passwords include:

- The user has a single sign-on experience.
- The user does not need to log off and on to supply multiple user names and passwords for different computers.
- The user can store any number of user names and passwords to be reused later.
- The user can store user names and passwords in a user's profile to provide portability and privacy of the user names and passwords.
- The user can supply different strong passwords for a variety of resources and store them in a single and easy-to-use location.

# How Stored User Names and Passwords works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Stored User Names and Passwords works**

Stored User Names and Passwords obtains its information in two ways: explicit creation and learning from the user. When users enter a user name and password for a target computer or domain, that information is stored and used when the users attempt to log on to an appropriate computer. If no stored information is available and users supply a user name and password, they have the ability to save the information. If the user decides to save the information, Stored User Names and Passwords then receives and stores it.

When Windows XP or a member of the Windows Server 2003 family attempts to connect to a new computer on the network, it supplies the current user name and password to the computer. If this is not sufficient to provide access, Stored User Names and Passwords will attempt to supply the necessary user name and password. All stored user names and passwords will be examined, from most specific to least specific as appropriate to the resource, and the connection will be attempted with those user names and passwords in that order. Because user names and passwords are read and applied in order from most to least specific, no more than one user name and password can be stored for each individual target or domain.

Stored User Names and Passwords also allows the user to save the supplied user name and password for reuse. This information is stored in a secure part of the user's profile and cannot be accessed by other users. If the user is configured to use a single profile across the enterprise, the stored user names and passwords will be retained wherever the user logs on to the network.

# Stored User Name and Password Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- My computer connects to computers with the wrong access level or account.
- My computer cannot connect to resources that were previously available.
- My computer obtains incorrect access while using a shared user account.

**My computer connects to computers with the wrong access level or account.**
**Cause:** A user name and password was stored for this account that is providing more or less access to resources than expected.

**Solution:** Remove the stored user name and password by following the steps in Delete stored user names and passwords.

**My computer cannot connect to resources that were previously available.**
**Cause:** A user name and password was stored for this account that has expired or has changed its password without updating the stored information.

**Solution:** Correct the stored user name and password by following the steps in View and modify stored user names and passwords, then try to access the resource again.

**My computer obtains incorrect access while using a shared user account.**
**Cause:** A user name and password was stored for this shared user account that provides access to this resource.

**Solution:** Remove the stored user name and password by following the steps in Delete stored user names and passwords. To ensure this problem does not reoccur, disable Stored User Names and Passwords for this shared user account by using a Group Policy object. For more information on creating Group Policy objects, see Group Policy (pre-GPMC).

# Passwords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Passwords**

Passwords provide the first line of defense against unauthorized access to your domain and local computers.

- For tips about using passwords, see Password Best practices.
- For help with specific tasks, see Password How To....
- For general background information, see Password Concepts.
- For problem-solving instructions, see Password Troubleshooting.

# Password Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Create an extensive defense model.**

- Educate your users about how to best protect their accounts from unauthorized attacks. For more information, see Encourage your users to follow best practices for password protection.

- Use the system key utility (Syskey) on computers throughout your network. The system key utility uses strong encryption techniques to secure account password information that is stored in the Security Accounts Manager (SAM) database.

  For more information about the system key utility, see The system key utility. For information about how to create or update a system key, see Create or update a system key.

- Define password policy that ensures that every user is following the password guidelines that you decide are appropriate. For more information, see Define password policy so that all user accounts are protected with strong passwords.

- Consider whether implementing account lockout policy is appropriate for your organization. For information about the risks involved in applying account lockout policy, see Be cautious when defining account lockout policy.

**Encourage your users to follow best practices for password protection.**

- Always use strong passwords. For more information, see Strong passwords.

- If passwords must be written down on a piece of paper, store the paper in a secure place and destroy it when it is no longer needed.

- Never share passwords with anyone.

- Use different passwords for all user accounts.

- Change passwords immediately if they may have been compromised.

- Be careful about where passwords are saved on computers. Some dialog boxes, such as those for remote access and other telephone connections, present an option to save or remember a password. Selecting this option poses a potential security threat.

**Define password policy so that all user accounts are protected with strong passwords.**

- Define the **Enforce password history** policy setting so that several previous passwords are remembered. With this policy setting, users cannot use the same password when their password expires.

- Define the **Maximum password age** policy setting so that passwords expire as often as necessary for your environment, typically, every 30 to 90 days. With this policy setting, if an attacker cracks a password, the attacker only has access to the network until the password expires.

- Define the **Minimum password age** policy setting so that passwords cannot be changed until they are more than a certain number of days old. This policy setting works in combination with the **Enforce password history** policy setting. If a minimum password age is defined, users cannot repeatedly change their passwords to get around the **Enforce password history** policy setting and then use their original password. Users must wait the specified number of days to change their passwords.

- Define a **Minimum password length** policy setting so that passwords must consist of at least a specified number of characters. Long passwords--seven or more characters--are usually stronger than short ones. With this policy setting, users cannot use blank passwords, and they have to create passwords that are a certain number of characters long.

- Enable the **Password must meet complexity requirements** policy setting. This policy setting checks all new passwords to ensure that they meet basic strong password requirements. For a full list of these requirements, see Passwords must meet complexity requirements.

- For information about how to apply or modify these policy settings, see Apply or modify password policy. For information about each of these policy settings, see Password Policy.

**Be cautious when defining account lockout policy.**

- Account lockout policy should not be applied haphazardly. While you increase the probability of thwarting an unauthorized attack on your organization with account lockout policy, you can also unintentionally lock out authorized users, which can be quite costly for your organization.

- If you decide to apply account lockout policy, set the **Account lockout threshold** policy setting to a high enough number that authorized users are not locked out of their user accounts simply because they mistype a password.

- Authorized users can be locked out if they change their passwords on one computer, but not on another computer. The computer that is still using the old password will continuously attempt to authenticate the user with the wrong password, and it will eventually lock out the user account. This might be a costly consequence of defining account lockout policy, because the authorized users cannot access network resources until their accounts are restored. This issue does not exist for organizations that only use domain controllers that are running Windows Server 2003 family operating systems.

- For more information about account lockout policy, see Account lockout policy overview. For information about how to apply or modify account lockout policy, see Apply or modify account lockout policy.

# Password How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Change an existing password
- Reset a password
- Reset a local user account password with a password reset disk
- Create a password reset disk
- Apply or modify password policy
- Apply or modify account lockout policy
- Create or update a system key

# Change an existing password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change an existing password**

1. Press CTRL+ALT+DEL.
2. Click **Change Password**.
3. In **User name**, type the user name of the account that you want to change the password for.
4. In **Old Password**, type the existing password.
5. In **New Password** and in **Confirm New Password**, type the new password, and then click **OK**.

**Note**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Strong passwords
Password Best practices
Passwords

# Reset a password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reset a password**

- For a local user account
- For a domain user account, and you are on a member server or a workstation that is joined to the domain
- For a domain user account, and you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed

**For a local user account**

1. Open Computer Management (Local).
2. In the console tree, double-click **Local Users and Groups**.

   **Where?**

   - Computer Management (Local)/System Tools/Local Users and Groups

3. In the console tree, click **Users**.
4. In the details pane, right-click the user name, and then click **Set Password**.
5. Read the warning message, and then, if you want to continue, click **Proceed**.
6. In **New password** and in **Confirm password**, type a new password, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. Power Users can reset passwords of accounts they create. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**For a domain user account, and you are on a member server or a workstation that is joined to the domain**

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Active Directory Users and Computers**, and then click **Add**.
4. Click **Close**, and then click **OK**.
5. In the console tree, click **Users**.

**Where?**

- Active Directory Users and Computers/*Domain name*/Users

6. In the details pane, right-click the user name, and then click **Reset Password**.

7. Type a new password in **New password** and in **Confirm password**, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.

#### For a domain user account, and you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed

1. Open Active Directory Users and Computers.

2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*Domain name*/Users

3. In the details pane, right-click the user name, and then click **Reset Password**.

4. Type a new password in **New password** and in **Confirm password**, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

### Caution

- After a user's password is reset, some types of information are no longer accessible, including the following:

  - E-mail that is encrypted with the user's public key

  - Internet passwords that are saved on the computer

  - Files that the user has encrypted

  To prevent data loss when you reset a local user account password, use a password reset disk to reset the password. For more information about password reset disks, see Related Topics.

- Because data can be lost when passwords are reset, use this procedure only if there is no other way to access the account.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Password reset disks
Create a password reset disk
Reset a local user account password with a password reset disk
Passwords

# Create a password reset disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a password reset disk**

1. Press CTRL+ALT+DEL, and then click **Change Password**.
2. In **User name**, type the user name of the account that you want to create a password reset disk for.
3. In **Log on to**, click your *LocalComputerName*.
4. Click **Backup**.
5. Follow the steps in the Forgotten Password Wizard until the procedure is complete. Store the password reset disk in a safe and secure place.

## Notes

- Password reset disks can only be used for local computer accounts, not for domain accounts. For information on how to reset passwords for domain accounts, see "Reset a password" in Related Topics.
- You cannot change your password and create a password reset disk at the same time. If you type a new password in **New Password** and in **Confirm New Password** before you click **Backup**, the new password information is not saved. When the wizard prompts you for your current user account password, you must type the old password. You can change your password after you create the password reset disk.
- You can change your password any time after you create a password reset disk. You do not have to create a new password reset disk if you change your password or if the password is reset manually.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Passwords
Password reset disks
Reset a password

# Reset a local user account password with a password reset disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reset a local user account password with a password reset disk**

1. Start your computer, or log off from it.
2. At the **Welcome to Windows** dialog box, press CTRL+ALT+DEL. In **User name**, type your user name, and then click **OK**. The **Logon Failed** dialog box appears.
3. Click **Reset**, and then insert your password reset disk into drive A.
4. Follow the instructions in the Password Reset Wizard to create a new password.
5. Log on to your computer with the new password, and then store your password reset disk in a safe place, in case you need it to reset your password in the future. It is not necessary to make a new password reset disk when you change your password.

## Notes

- Follow the steps in this procedure if you are logging on to a local or workgroup environment. For information about how to reset a forgotten password if you are in a domain environment, see "Reset a password" in Related Topics.
- A password reset disk must have been created for the local account for the **Logon Failed** dialog box to appear. If a password reset disk has not been created, the **Logon Message** dialog box appears, informing you that your logon information is incorrect. If a password reset disk has not been created, and the user account password has been forgotten, you must reset the password to access the user account. For more information about resetting passwords, see "Reset a password" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Passwords
Create a password reset disk
Reset a local user account password with a password reset disk
Password reset disks
Reset a password

# Apply or modify password policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To apply or modify password policy**

- For a local computer
- For a domain, and you are on a member server or a workstation that is joined to the domain
- For a domain, and you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed

**For a local computer**

1. Open Local Security Settings.
2. In the console tree, click **Password Policy**.

   **Where?**

   - Security Settings/Account Policies/Password Policy

3. In the details pane, right-click the policy setting that you want, and then click **Properties**.
4. Select the options that you want, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**For a domain, and you are on a member server or a workstation that is joined to the domain**

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. In **Select Group Policy Object**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit--or create a new one, click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree, click **Password Policy**.

**Where?**

- *Group Policy Object* [*computer name*] Policy/Computer Configuration/Windows Settings/Security Settings/Account Policies/Password Policy

8. In the details pane, right-click the policy setting that you want, and then click **Properties**.
9. If you are defining this policy setting for the first time, select the **Define this policy setting** check box.
10. Select the options that you want, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.

#### For a domain, and you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or organizational unit that you want to set Group Policy for.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Click an entry in **Group Policy Object Links** to select an existing Group Policy object (GPO), and then click **Edit**. You can also click **New** to create a new GPO, and then click **Edit**.
5. In the console tree, click **Password Policy**.

   **Where?**

   - *Group Policy Object* [*computer name*] Policy/Computer Configuration/Windows Settings/Security Settings/Account Policies/Password Policy

6. In the details pane, right-click the policy setting that you want, and then click **Properties**.
7. If you are defining this policy setting for the first time, select the **Define this policy setting** check box.
8. Select the options that you want, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Note**

- Password policy ensures that all users are creating passwords that adhere to the guidelines set by administrators. For more information about each of the password policy settings, see "Password policy" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Password Best practices
Passwords
Password Policy

# Apply or modify account lockout policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To apply or modify account lockout policy**

- For a local computer
- For a domain, and you are on a member server or a workstation that is joined to a domain
- For a domain, and you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed

**For a local computer**

1. Open Local Security Settings.
2. In the console tree, click **Account Lockout Policy** (console tree location is: Security Settings/Account Policies/Account Lockout Policy).
3. In the details pane, right-click the policy setting that you want, and then click **Properties**.
4. Select the options that you want, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**For a domain, and you are on a member server or a workstation that is joined to a domain**

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. In **Select Group Policy Object**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object in the appropriate domain, site, or organizational unit--or create a new one, click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree, click **Account Lockout Policy** (console tree location is: *Group Policy Object* [*computer name*] Policy/Computer Configuration/Windows Settings/Security Settings/Account Policies/Account Lockout Policy).
8. In the details pane, right-click the policy setting that you want, and then click **Properties**.
9. If you are defining this policy setting for the first time, select the **Define this policy setting** check box.

10. Select the options that you want, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.

#### For a domain, and you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the domain or organizational unit that you want to set Group Policy for.

3. Click **Properties**, and then click the **Group Policy** tab.

4. Click an entry in **Group Policy Object Links** to select an existing Group Policy object (GPO), and then click **Edit**. You can also click **New** to create a new GPO, and then click **Edit**.

5. In the console tree, click **Account Lockout Policy** (console tree location is: *Group Policy Object* [*computer name*] Policy/Computer Configuration/Windows Settings/Security Settings/Account Policies/Account Lockout Policy).

6. In the details pane, right-click the policy setting that you want, and then click **Properties**.

7. If you are defining this policy setting for the first time, select the **Define this policy setting** check box.

8. Select the options that you want, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

### Caution

- When account lockout policy is enabled, there is a risk of locking out legitimate users. To avoid locking out legitimate users who have simply mistyped or misremembered their passwords, set the account lockout threshold to a high number.

  Also, if a password is changed on one computer, but the user is logged on to another computer with the old password, the computer with the old password continuously attempts to authenticate the user by using the old password, and it eventually locks out the user account. This issue does not exist for organizations that only use domain controllers that are running Windows Server 2003 family operating systems.

### Note

- For more information about each of the account lockout policy settings, see "Account Lockout Policy" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Account lockout policy overview
Password Best practices
Passwords
Account Lockout Policy

# Create or update a system key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create or update a system key**

1. Click **Start**, click **Run**, type **syskey**, and then click **OK**.
2. Click **Encryption Enabled**, and then click **Update**.
3. Click the option that you want, and then click **OK**.

   For information about each option, see "The system key utility" in Related Topics.

## Notes

- Different administrative credentials are required to perform this procedure, depending on your environment:
  - If you create or update a system key for your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
  - If you create or update a system key for a domain controller: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- If the disk that stores the system key is lost, or if the system key password is forgotten, you cannot start the computer without restoring the registry to the state it was in before the system key was used. For more information on restoring the registry, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Passwords
The system key utility
Restore the registry

# Password Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Strong passwords
- Password reset disks
- Preventing data loss from password resets
- Account lockout policy overview
- The system key utility

# Strong passwords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Strong passwords**

The role that passwords play in securing an organization's network is often underestimated and overlooked. Passwords provide the first line of defense against unauthorized access to your organization. The Microsoft® Windows Server 2003 family has a new feature that checks the complexity of the password for the Administrator account during setup of the operating system. If the password is blank or does not meet complexity requirements, the **Windows Setup** dialog box appears, warning you of the dangers of not using a strong password for the Administrator account. If you leave this password blank, you will not be able to access this account over the network.

Weak passwords provide attackers with easy access to your computers and network, while strong passwords are considerably harder to crack, even with the password-cracking software that is available today. Password-cracking tools continue to improve, and the computers that are used to crack passwords are more powerful than ever. Password-cracking software uses one of three approaches: intelligent guessing, dictionary attacks, and brute-force automated attacks that try every possible combination of characters. Given enough time, the automated method can crack any password. However, strong passwords are much harder to crack than weak passwords. A secure computer has strong passwords for all user accounts.

A weak password:

- Is no password at all.
- Contains your user name, real name, or company name.
- Contains a complete dictionary word. For example, *Password* is a weak password.

A strong password:

- Is at least seven characters long.
- Does not contain your user name, real name, or company name.
- Does not contain a complete dictionary word.
- Is significantly different from previous passwords. Passwords that increment (*Password1*, *Password2*, *Password3* ...) are not strong.
- Contains characters from each of the following four groups:

| Group | Examples |
|---|---|
| Uppercase letters | A, B, C ... |
| Lowercase letters | a, b, c ... |
| Numerals | 0, 1,2, 3, 4, 5, 6, 7, 8, 9 |
| Symbols found on the keyboard (all keyboard characters not defined as letters or numerals) | ` ~ ! @ # $ % ^ & * ( ) _ + - = { } \| \ : " ; ' < > ? , . / |

An example of a strong password is *J*p2leO4>F*.

A password can meet most of the criteria of a strong password but still be rather weak. For example, *Hello2U!* is a relatively weak password even though it meets most of the criteria for a strong password and also meets the complexity requirements of password policy. *H!elZl2o* is a strong password because the dictionary word is interspersed with symbols, numbers, and other letters. It is important to educate users about the benefits of using strong passwords and to teach them how to create passwords that are actually strong.

You can create passwords that contain characters from the extended ASCII character set. Using extended ASCII characters increases the number of characters that you can choose when you create a password. As a result, it might take more time for password-cracking software to crack passwords that contain these extended ASCII characters than it does to crack other passwords. Before using extended ASCII characters in your password, test them thoroughly to make sure that passwords containing extended ASCII characters are compatible with the applications that your organization uses. Be especially cautious about using extended ASCII characters in passwords if your organization uses several different operating systems.

You can find extended ASCII characters in Character Map. Some extended ASCII characters should not be used in passwords. Do not use a character if a keystroke is not defined for it in the lower-right corner of the **Character Map** dialog box. For more information about how to use Character Map, see Using Character Map.

Examples of passwords that contain characters from the extended ASCII character set are *kUµ!¶0o* and *Wf©$0k#»g¤5ªrd*.

You can implement a password policy setting that enforces password complexity requirements. For more information about this policy setting, see Passwords must meet complexity requirements. For information about how to apply a password policy, see Apply or modify password policy.

Windows passwords can be up to 127 characters long. However, if you are on a network that also has computers running Windows 95 or Windows 98, consider using passwords that are not longer than 14 characters. Windows 95 and Windows 98 support passwords of up to 14 characters. If your password is longer, you might not be able to log on to your network from those computers.

For more information about passwords, see Passwords.

# Password reset disks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Password reset disks**

It is not unusual for users to forget their passwords to local user accounts from time to time, especially when they use strong passwords. Before the advent of password reset disks, the only way for administrators to restore a forgotten local user account password was to manually reset the user's password. In the process, the following information was lost:

- E-mail that was encrypted with the user's public key
- Internet passwords that were saved on the computer
- Files that the user had encrypted

Password reset disks offer another solution to the problem of a forgotten password for a local user account. If users create password reset disks for their local accounts before they forget their passwords, they can reset the passwords without losing valuable data that was lost previously with administrative password resets.

When you create a password reset disk, a public key and private key pair are created. The private key is stored on a disk: the password reset disk. The public key encrypts the local user account password. If users forget their passwords, they can insert the password reset disk, which contains the private key, and decrypt the current password. The Forgotten Password Wizard prompts the user for a new password, which is then encrypted with the public key. Data is not lost because, basically, the user is simply changing a password. It is essential that password reset disks be stored in secured locations. For more information on how to create a password reset disk, see Create a password reset disk.

For more information about encryption, see Encryption. For more information about passwords, see Passwords.

# Preventing data loss from password resets

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Preventing data loss from password resets**
To keep protected information secure, after a user's password is reset, some types of information are no longer accessible, including the following:

- E-mail that is encrypted with the user's public key
- Internet passwords that are saved on the computer
- Files that the user has encrypted

To avoid such data loss, do not reset a user's password. When a new local user account is created, have the user create a password reset disk. Then, if the user forgets the password, the password reset disk can be used to reset the password without data loss. If a user forgets the password to a domain user account, the password must be reset manually.

For more information about how to create a password reset disk, see Create a password reset disk.

# Account lockout policy overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Account lockout policy overview**

Account lockout policy disables a user account if an incorrect password is entered a specified number of times over a specified period. These policy settings help you to prevent attackers from guessing users' passwords, and they decrease the likelihood of successful attacks on your network. For more information about how to implement account lockout policy, see Apply or modify account lockout policy.

Before you enable account lockout policy, it is important to realize that there is a risk of unintentionally locking authorized users out of their accounts. Such a result can be quite costly for your organization, because locked-out users cannot access their user accounts until the account unlocks automatically after a specified amount of time or until you unlock the accounts for them.

Authorized users can lock themselves out of an account by mistyping their password or by remembering it incorrectly, or by changing their password on a computer while they are logged on to another computer. The computer with the incorrect password continuously tries to authenticate the user, and because the password it is using to authenticate is incorrect, the user account is eventually locked out. This issue does not exist for organizations that only use domain controllers that are running Windows Server 2003 family operating systems. To avoid locking out authorized users, set the account lockout threshold to a high number. Remember, however, that the scenario in which a computer continuously tries to authenticate a user with an incorrect password is very similar to the behavior that is employed by password-cracking software. Setting the account lockout threshold high enough that the authorized user will not be locked out in this situation may inadvertently allow unauthorized access to your network by hackers. For more information about the account lockout threshold, see Account lockout threshold.

For more information about passwords and their policies, see Passwords.

# The system key utility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The system key utility**

Password information for user accounts is stored in the Security Accounts Manager (SAM) database of the registry on workstations and member servers. On domain controllers, password information is stored in directory services. It is not unusual for password-cracking software to target the SAM database or directory services to access passwords for user accounts. The system key utility (Syskey) provides an extra line of defense against password-cracking software. It uses strong encryption techniques to secure account password information that is stored in the SAM database or in directory services. Cracking encrypted account passwords is more difficult and time consuming than cracking nonencrypted account passwords. For more information on encryption, see Encryption.

There are three system key options in the **Startup Key** dialog box that are designed to meet the needs of different environments, as described in the following table.

| System key option | Relative security level | Description |
|---|---|---|
| **System Generated Password, Store Startup Key Locally** | Secure | Uses a computer-generated random key as the system key and stores an encrypted version of the key on the local computer. This option provides strong encryption of password information in the registry, and it enables the user to restart the computer without the need for an administrator to enter a password or insert a disk. |
| Administrator generated password, **Password Startup** | More secure | Uses a computer-generated random key as the system key and stores an encrypted version of the key on the local computer. The key is also protected by an administrator-chosen password. Users are prompted for the system key password when the computer is in the initial startup sequence. The system key password is not stored anywhere on the computer. |
| **System Generated Password, Store Startup Key on Floppy Disk** | Most secure | Uses a computer-generated random key and stores the key on a floppy disk. The floppy disk that contains the system key is required for the system to start, and it must be inserted at a prompt during the startup sequence. The system key is not stored anywhere on the computer. |

For information on how to implement the system key utility on a computer, see Create or update a system key.

Use of the system key utility is optional. If the disk that contains the system key is lost, or if the password is forgotten, you cannot start the computer without restoring the registry to the state it was in before the system key was used. For information on restoring the registry, see Restore the registry.

# Password Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- I changed the password policy, but the changes have not taken effect.
- When I define a policy setting in account lockout policy, the Suggested Value Changes dialog box opens, suggesting that I change my other policy settings.
- I am trying to log on to Windows 95, Windows 98, or another product, and the password is not working.

**I changed the password policy, but the changes have not taken effect.**
**Cause:** Policy settings have not been updated.

**Solution:** Click **Start**, click **Run**, type **gpupdate**, and then click **OK**. This refreshes policy settings, and it updates the changes to password policy on the computer.

**See also:** Gpupdate

**When I define a policy setting in account lockout policy, the Suggested Value Changes dialog box opens, suggesting that I change my other policy settings.**
**Cause:** **Account lockout duration** must be greater than or equal to the time defined in **Reset account lockout counter after**. If you define the policy settings so that the **Reset account lockout counter after** time is greater than **Account lockout duration**, the **Suggested Value Changes** dialog box appears, reminding you to change the policy settings so that they adhere to the policy setting guidelines.

**Solution:** In the **Suggested Value Changes** dialog box, click **OK** to change the policy settings to the suggested values, or click **Cancel**, and then define the policy settings yourself.

**See also:** Account Lockout Policy; Apply or modify account lockout policy

**I am trying to log on to Windows 95, Windows 98, or another product, and the password is not working.**
**Cause:** The password is longer than 14 characters. Windows 95 or Windows 98 computers only recognize passwords that are 14 characters or less.

**Solution:** Log on to another computer and change the password so that it is 14 characters or less.

**See also:** Change an existing password; Strong passwords

**Cause:** The system that you are using does not support unusual characters, such as spaces or characters in the extended ASCII character set.

**Solution:** Log on to another computer and change the password so that it does not contain these special characters.

**See also:** Change an existing password; Strong passwords

# Access Control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Access control**
In order to secure a computer and its resources, you must take into consideration what rights users will have. You can secure a computer or multiple computers by granting users or groups specific user rights. You can secure an object, such as a file or folder, through assigning permissions to allow users or groups to perform specific actions on that object.

- For tips about access control, see Access Control Best Practices.
- For help with specific tasks, see Access Control How To....
- For general background information, see Access Control Concepts.

# Access Control Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- Best practices for permissions and user rights
- Best practices for assigning permissions on Active Directory objects

# Best practices for permissions and user rights

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for permissions and user rights**

**Assign permissions to groups rather than to users**

- Because it is inefficient to maintain user accounts directly, assigning permissions on a user basis should be the exception.

**Deny permissions should be used for certain special cases**

- Use Deny permissions to exclude a subset of a group which has Allowed permissions.
- Use Deny to exclude one special permission when you have already granted full control to a user or group.

**Use security templates**

- Rather than set individual permissions, use security templates whenever possible. For information on security templates, see Security Templates.

**If possible, avoid changing the default permission entries on file system objects, particularly on system folders and root folders**

- Changing default permissions can cause unexpected access problems or reduce security.

**Never deny the Everyone group access to an object**

- If you deny everyone permission to an object, that includes administrators. A better solution would be to remove the Everyone group, as long as you give other users, groups, or computers permissions to that object.

**Assign permissions to an object as high on the tree as possible and then apply inheritance to propagate the security settings through the tree**

- You can quickly and effectively apply access control settings to all children or a subtree of a parent object. By doing this, you gain the greatest breadth of effect with the least effort. The permission settings you establish should be adequate for the majority of users, groups, and computers.

**Privileges can sometimes override permissions**

- Privileges and permissions may disagree, and you should know what happens if they do. For more information, see Privileges.

**For permissions on Active Directory objects, make sure you understand the best practices specific to Active Directory objects**

- Active Directory has its own set of best practices regarding permissions. For more information, see Best practices for assigning permissions on Active Directory objects

  For more information, see Assign user rights for your local computer and Set, view, change, or remove permissions on files and folders.

  **Notes**

  - Inherited Deny permissions do not prevent access to an object if the object has an explicit Allow permission entry.

- Explicit permissions take precedence over inherited permissions, even inherited Deny permissions.

# Best practices for assigning permissions on Active Directory objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for assigning permissions on Active Directory objects**
For information on security and access control in Active Directory, see Securing Active Directory and Access control in Active Directory.

**If possible, avoid changing the default permissions on Active Directory objects**

- Changing default permissions could cause unexpected access problems or reduce security.

**Avoid granting Full Control permissions over an object or organizational unit**

- Granting someone Full Control allows them to take ownership of an object and modify the permissions on it. If someone has Full Control on a container, then they can Take Ownership of, and have Full Control over, all objects in that container. As far as possible, instead of allowing Full Control, give only the permissions needed by the user.

**Minimize the number of access control entries that apply to child objects**

- When using the Apply Onto option to control inheritance, be aware that not only will the specified objects inherit that access control entry (ACE) but *all* child objects will also receive a copy of that ACE. If there are enough objects that will get copies of this ACE, then that increased amount of data can cause performance problems on your network.

**When possible, assign the same set of permissions to multiple objects**

- Access control lists (ACLs) in the Windows Server 2003 family feature single-instancing: If multiple objects have identical access control lists (ACLs), then Active Directory will only store one instance of the ACL. For more information on how inheritance works for Active Directory objects, see Changing inherited permissions.

**When possible, assign access rights on a broad level rather than assigning individual user rights**

- Minimizing the number of access control entries will improve performance.
- Allow "Read All Properties" or "Write All Properties" rather than individual properties.
- Allow Read or Write access to property sets rather than individual properties.

  A property set is a collection of attributes. For example, the Personal Information property set includes the attribute's address, personal title, and so on. By setting access on the property set, you have automatically set access on all the attributes contained in that property set.
- Allow "Create All Child Objects" or "Delete All Child Objects," rather than specifying individual child objects.
- Allow "All Extended Rights" rather than allowing the individual extended rights.
- Allow "All Validated Writes" rather than allowing the individual validated rights.

**Note**

- There are times when assigning individual permissions is appropriate. For example, if a Help Desk operator needs Write access to two properties of a User object, then it makes sense to use one access control entry (ACE) for each of them, rather than trying use only one ACE by granting write access to the entire user object.

**Whenever possible, assign permissions to groups rather than users**

- Use groups to delegate authority over the directory. If one set of users needs Read permissions, and another set of users needs Change permissions, then create one group for each set of users and assign the permission to the group. If multiple global groups need the same access, create a local group containing the global groups and assign permissions to the local group.

  For more information, see Active Directory object permissions.

  **Caution**

  - The operation of your domain may become impaired if the total size of all the Active Directory permissions approaches the disk storage or processing speed capacity of your domain controller. Minimize the number of individual permissions that are propagated to child objects.

## Notes

- Inherited Deny permissions do not prevent access to an object if the object has an explicit Allow permission entry.

- Explicit permissions take precedence over inherited permissions, even inherited Deny permissions.

# Access Control How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Set, View, Change, or Remove Permissions on an Object
- Secure a file or folder
- Apply or modify permission entries for objects using Group Policy
- Take ownership of a file or folder
- View effective permissions on files and folders

# Set, View, Change, or Remove Permissions on an Object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set, view, change, or remove permissions on an object**

- Set, view, change, or remove permissions on files and folders
- Set, view, change, or remove special permissions
- Set permissions on a shared resource
- Assign permissions to a registry key
- Set or remove permissions for a printer
- Change the permissions a user or group has to a connection
- Authorize WMI users and set permissions

# Set, view, change, or remove permissions on files and folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set, view, change, or remove permissions on files and folders**

1. Open Windows Explorer.
2. Right-click the file or folder for which you want to set permissions, click **Properties**, and then click the **Security** tab.
3. Do one of the following:
   - To set permissions for a group or user that does not appear in the **Group or user names** box, click **Add**. Type the name of the group or user you want to set permissions for and then click **OK**.
   - To change or remove permissions from an existing group or user, click the name of the group or user.
4. Do one of the following:
   - To allow or deny a permission, in the **Permissions for *User or Group*** box, select the **Allow** or **Deny** check box.
   - To remove the group or user from the **Group or user names** box, click **Remove**.

**Notes**

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- In the Windows Server 2003 family, the Everyone group no longer includes Anonymous Logon.
- You can only set file and folder permissions on drives formatted to use NTFS.
- To change permissions, you must be the owner or have been granted permission to do so by the owner.
- Groups or users that are granted Full Control for a folder can delete files and subfolders within that folder, regardless of the permissions that protect the files and subfolders.
- If the check boxes under **Permissions for *User or Group*** are shaded or if the **Remove** button is unavailable, then the file or folder has inherited permissions from the parent folder. For more information on how inheritance affects files and folders, see Related Topics.
- When adding a new user or group, by default, this user or group will have **Read & Execute**, **List Folder Contents**, and **Read** permissions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View effective permissions on files and folders
Set, view, change, or remove special permissions

File and folder permissions
How inheritance affects file and folder permissions
Command-line reference A-Z
Permissions

# Set, view, change, or remove special permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set, view, change, or remove special permissions**

1. Right-click the object you want to set advanced or special permissions on, click **Properties**, and then click the **Security** tab.

2. Click **Advanced**, and then do one of the following:

| To | Do this |
|---|---|
| Set special permissions for an additional group or user | • Click **Add**.<br>• In **Enter than object name to select (examples)**, type the name of the user or group, and then click **OK**. |
| View or change special permissions for an existing group or user | Click the name of the group or user, and then click **Edit**. |
| Remove an existing group or user and its special permissions | Click the name of the group or user and then click **Remove**. If the **Remove** button is unavailable, clear the **Inherit from parent the permission entries that apply to child objects. Include these with entries explicitly defined here** check box, and then click **Remove**. |

3. In the **Permissions** box, select or clear the appropriate **Allow** or **Deny** check boxes.

4. In **Apply onto**, click the folders or subfolders you would like these permissions to be applied to.

5. To configure security so that the subfolders and files will not inherit these permissions, clear the **Apply these permissions to objects and/or containers within this container only** check box.

6. Click **OK**, and then, in **Advanced Security Settings for *ObjectName***, click **OK**.

**Caution**

- If you select the **Replace permission entries on all child objects with entries shown here that apply to child objects** check box, then all subfolders and files will have all of their permission entries reset to those inheritable from the parent object. Once you click **Apply** or **OK**, you cannot undo this operation by clearing the check box.

**Notes**

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

- The Everyone group no longer includes the Anonymous Logon permission.

- If you select the **Inherit from parent the permission entries that apply to child objects. Include these with entries explicitly defined here** check box, then this file or folder will inherit permission entries from the parent object.

- You can only set permissions only on drives formatted to use NTFS.

- If the check boxes under **Permissions** are shaded, the permissions are inherited from the parent folder.

- To change permissions, you must be the owner or have been granted permission to do so by the owner.

- Groups or users that have been granted Full Control for a folder can delete files and subfolders within that folder, regardless of the permissions that protect the files and subfolders.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Permissions
Permissions for files and folders
Best practices for assigning permissions on Active Directory objects
Selecting where to apply permissions

# Set permissions on a shared resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions on a shared resource**

- Using Shared Folders
- Using Windows Explorer

**Using Shared Folders**

1. To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. In the details pane, right-click the shared resource that you want to set permissions for, and then click **Properties**.
4. On the **Share Permissions** tab, make any of the following changes, and then click **OK**:

   - To assign permissions to a user or group for a shared resource, click **Add**. In the **Select Users, Computers, or Groups** dialog box, look for or type the user or group name, and then click **OK**.
   - To revoke access to the shared resource, click **Remove**.
   - To set individual permissions for the user or group, in the **Permissions for***group or user* box, select the **Allow** or **Deny** check boxes.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using Windows Explorer**

1. Open Windows Explorer.
2. Right-click the shared folder or drive that you want to set permissions for, and then click **Sharing and Security**.
3. On the **Sharing** tab, click **Permissions**, make any of the following changes, and then click **OK**:

   - To assign permissions to a user or group for a shared resource, click **Add**. In the **Select Users, Computers, or Groups** dialog box, look for or type the user or group name, and then click **OK**.
   - To revoke access to a shared resource, click **Remove**.
   - To set individual permissions for the user or group, in the **Permissions for** *group or user* box, select the **Allow** or **Deny** check boxes.

### Note

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

### Important

- Share permissions apply only to users who gain access to the resource over the network. They do not apply to users who log on locally, such as on a terminal server. In these cases, use access control on the NTFS file system to set permissions. For more information, see Related Topics.

### Notes

- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- You can use Shared Folders to manage shared resources on both local and remote computers. For information about how to connect to another computer, see Related Topics. With Windows Explorer and the command line, you can manage shared resources on your local computer only.

- When permissions have been assigned both to the shared resource and at the file system level, the more restrictive permission always applies.

- It is usually easier to assign permissions to groups and then add users to groups, rather than assigning identical permissions to individual users.

- If you change permissions on special shared resources, such as ADMIN$, the default settings may be restored when the Server service is stopped and restarted or when the computer is restarted. Note that this does not apply to user-created shared resources whose share name ends in $. For more information about special shared resources, see Related Topics.

- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "How to configure file sharing in Windows XP," in the Microsoft Knowledge Base.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Special shared resources
Set, view, change, or remove permissions on files and folders
Microsoft Management Console
Connect to another computer
**Other Resources**
Net session
Send a console message to all connected computers
Send a console message to a computer that you are managing

# Assign permissions to a registry key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign permissions to a registry key**

1. Open Registry Editor.
2. Click the key to which you want to assign permissions.
3. On the **Edit** menu, click **Permissions**.
4. Assign an access level to the selected key as follows:
   - To grant the user permission to read the key contents, but not save any changes made to the file, under Permissions for *name*, for **Read**, select the **Allow** check box.
   - To grant the user permission to open, edit, and take ownership of the selected key, under Permissions for *name*, for **Full Control**, select the **Allow** check box.
   - To grant the user special permission in the selected key, click **Advanced**.
5. If you are assigning permissions to a subkey and you want the inheritable permissions assigned to the parent key to apply to the subkey also, click **Advanced** and select the **Inherit from parents the permission entries that apply to child objects. Include these with entries explicitly defined here** check box.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- You must have appropriate permissions to make changes to a registry key. To maintain security when making changes to a registry key for which you need administrative credentials, log in as a member of the Users group and run Regedit as an administrator by right-clicking the **Regedit** icon, clicking **Run as**, and clicking an account in the local Administrators group. The **Regedit** icon does not appear by default from the **Start** menu. To access the icon, open the Windows or WINNT folder on your computer.
- If you own a registry key, you can specify the users and groups that can open that key. To determine who can open your registry keys, you need to assign permissions to them. You can add or remove users or groups from those authorized to access your registry keys at any time.
- The **Special Permissions** check boxes indicate whether custom permissions have been set for this key, but you cannot set special permissions by clicking these check boxes. Click **Advanced** to set special permissions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Audit activity on a registry key
Grant full control of a registry key
Add users or groups to the Permissions list
Remove a user or group from the Permissions list
Assign special access to a registry key

# Set or remove permissions for a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set or remove permissions for a printer**

1. Open Printers and Faxes.
2. Right-click the printer for which you want to set permissions, click **Properties**, and then click the **Security** tab.
3. Do one of the following:
   - To change or remove permissions from an existing user or group, click the name of the user or group.
   - To set up permissions for a new user or group, click **Add**. In **Select Users, Computers, or Groups**, type the name of the user or group you want to set permissions for, and then click **OK** to close the dialog box.
4. In **Permissions**, click **Allow** or **Deny** for each permission you want to allow or deny, if necessary. Or, to remove the user or group from the permissions list, click **Remove**.

## Notes

- To change device settings, you must have the Manage Printers permission. For information about printing security permissions, see Related Topics.
- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- To view or change the underlying permissions that make up Print, Manage Printers, and Manage Documents, click the **Advanced** button.
- A printer must be shared in order for the permission settings to affect the users and groups listed.
- You can also view the permissions assigned to you by clicking the group you belong to on the **Security** tab. For information on finding out what group you belong to, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Share your printer
Identify members of a local group
Assigning printer permissions

# Change the permissions a user or group has to a connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the permissions a user or group has to a connection**

1. Open Terminal Services Configuration.
2. In the console tree, click **Connections**.
3. In the details pane, right-click the connection for which you want to change permissions, and then click **Properties**.
4. On the **Permissions** tab, click **Advanced** to open the **Advanced Security Settings** dialog box.
5. In **Permission Entries**, select the user or group for which you want to change permissions, and then click **Edit...** to open the **Permission Entry** dialog box.
6. In **Permissions**, select or clear, as appropriate, the **Allow** or **Deny** check boxes next to the permissions you want to set for the group.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Terminal Services Configuration, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Terminal Services Configuration**.
- You must use the Remote Desktop Users group to control remote access to Terminal Server and Remote Desktop for Administration.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add users and groups to permissions lists
Remove a user or group from permissions lists

# Authorize WMI users and set permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To authorize WMI users and set permissions**

1. Open WMI Control.
2. In the console tree, right-click **WMI Control**, and then click **Properties**.
3. Click the **Security** tab.
4. Select the namespace for which you want to give a user or group access, and then click **Security**.
5. In the **Security** dialog box, click **Add**.
6. In the **Select Users, Computers, or Groups** dialog box, enter the name of the object (user or group) that you want to add. Click the **Check Names** button to verify your entry and then click **OK**. You might have to change the location or use the **Advanced** button to query for objects. See the dialog box help for more details.
7. In the **Security** dialog box, under **Permissions**, select the permissions to allow or deny the new user or group.

| Level | Description |
|---|---|
| Execute Methods | Allows methods exported from the WMI classes or instances to be run. |
| Full Write | Allows full read, write, and delete access to all WMI objects, classes, and instances. |
| Partial Write | Allows write access to static WMI objects. |
| Provider Write | Allows write access to objects that are provided by providers. |
| Enable Account | Allows read access to WMI objects. |
| Remote Enable | Allows remote access to the namespace. |
| Read Security | Allows read-only access to WMI security information. |
| Edit Security | Allows read and write access to WMI security information. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the WMI Control console, click **Start**, click **Run**, type **wmimgmt.msc**, and then click **OK**.
- You can set permissions on a remote computer or a local computer. To access a remote computer, right-click **WMI Control**, click **Connect to another computer**, click **Another computer**, and then type the name of the computer to which you want to connect. If you are using WMI Control from the Computer Management console, right-click the **Computer Management** node to connect to the other computer.
- On computers running Windows 95, Windows 98, or Windows ME, all users have full control locally. Security settings are only relevant for remote connection to a computer running Windows 95, Windows 98, or Windows ME.
- You can delete a user's or group's authorization to access WMI services by selecting that user or group and clicking

**Remove**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
WMI Control Overview
Modify permissions or delete authorized users
Managing WMI security
Working with MMC console files
Permissions for files and folders

# Secure a file or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To secure a file or folder**

1. Open Windows Explorer, and then locate the file or folder you want to secure.
2. Right-click the file or folder, click **Properties**, and then click the **Security** tab.
3. In **Group or user names**, click the group or user that you want to change permissions for.
4. For users and groups that do not need the permission, clear the appropriate **Allow** check box, particularly for **Full Control**.
5. For the users and groups that need to use the file or folder, select the appropriate **Allow** check box for the permissions that are needed.
6. Repeat steps 3-4, as needed.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- The process of securing a file, folder, or other computer resource is sometimes called *lockdown*.
- Files and folders usually are sufficiently locked down by default.
- Incorrect permissions can cause security vulnerabilities and access problems, so locking down should be done cautiously and only when security considerations require it.
- Share permissions, not just NTFS permissions, can prevent access to a file or folder. Share permissions are not included in effective permissions calculations.
- An administrator can take ownership of any file on the computer.
- In the Windows Server 2003 family, the Everyone group no longer includes the Anonymous Logon group.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ownership
Take ownership of files or other objects
Restore files and directories

# Apply or modify permission entries for objects using Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To apply or modify permission entries for objects using Group Policy**

1. Open Microsoft Management Console.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. In the Group Policy Wizard, on the **Select Group Policy Object** page, click **Browse**.
5. In **Browse for a Group Policy object**, do one of the following:

   - To use an existing non-local Group Policy object (GPO), select a Group Policy object in the appropriate domain, site, or organizational unit, click **OK**, and then click **Finish**.
   - To create a new non-local Group Policy object, in the **Domain, OUs and linked Group Policy Objects** box, right-click in an empty area, click **New**, type the name of the new Group Policy object, click **OK**, and then click **Finish**.

   **Note**

   - If the **Browse for a Group Policy object** dialog box has only one tab, and it has the label **Computers**, then you only have access to local Group Policy objects and you cannot complete this procedure. This probably means the computer is not joined to an Active Directory domain.

6. Click **Close**, and then click **OK**.
7. Do one or more of the following:

| To modify permission entries on | Do this |
|---|---|
| System services | • In the console tree, click **System Services**.<br><br>**Where?**<br><br>  • Computer Configuration/Windows Settings/Security Settings/System Services<br><br>• In the details pane, right-click the service you want to change, click **Properties**, select the **Define this policy setting** check box, and then click **Edit Security**. |

| | |
|---|---|
| Registry keys | • In the console tree, right-click **Registry**.<br><br>**Where?**<br><br>  • Computer Configuration/Windows Settings/Security Settings/Registry<br><br>• Click **Add Key**, in **Select Registry Key**, click the key that you want to change, and then click **OK**. |
| Files or folders | • In the console tree, right-click **File System**.<br><br>**Where?**<br><br>  • Computer Configuration/Windows Settings/Security Settings/File System<br><br>• Click **Add File**, in **Add a file or folder**, click the file or folder that you want to change, and then click **OK**. |

8. Do one of the following:

   - To set permissions for a group or user that does not appear in the **Group or user names** box, click **Add**. Type the name of the group or user you want to set permissions for, and then click **OK**.
   - To change or remove permissions from an existing group or user, click the name of the group or user.

9. Do one of the following:

   - To allow a permission, in the **Permissions for *User or Group*** box, select the **Allow** check box.
   - To deny a permission, in the **Permissions for *User or Group*** box, select the **Deny** check box.
   - To remove the group or user from **Group or user names**, click **Remove**.

**Notes**

- You must be logged on as a computer administrator to complete this procedure.
- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- The security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes.
- To complete this procedure, the computer must be joined to an Active Directory domain, and you must have permission to edit the Group Policy object.
- These settings are refreshed on an ongoing basis, so configuring many access control lists through Group Policy may increase the load on the network.
- In the Windows Server 2003 family, the Everyone group no longer includes Anonymous Logon.
- You can set file and folder permissions only on drives that are formatted to use NTFS.

- To change permissions, you must be the owner or have been granted permission to do so by the owner.
- For more information about permissions on other objects, see Permissions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View effective permissions on files and folders
Set, view, change, or remove special permissions
File and folder permissions
How inheritance affects file and folder permissions
Access control overview

# Take ownership of a file or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To take ownership of a file or folder**

1. Open Windows Explorer, and then locate the file or folder you want to take ownership of.
2. Right-click the file or folder, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, and then click the **Owner** tab.
4. In the **Change owner to** box, do one of the following:
   - To change the owner to a user or group that is not listed, double-click **Other users and groups** and, in **Enter the object name to select (examples)**, type the name of the user or group, and then click **OK**.
   - To change the owner to a user or group that is listed, click the new owner.
5. (Optional) To change the owner of all subcontainers and objects within the tree, select the **Replace owner on subcontainers and objects** check box.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- You can transfer ownership in two ways:
  - The current owner can grant the **Take ownership** permission to others, allowing those users to take ownership at any time. A user granted the **Take ownership** permission can take ownership of the object or assign ownership to any group that the user is a member of.
  - A user who has the **Restore files and directories** privilege can double-click **Other users and groups** and choose any user or group to assign ownership to.
- An administrator can take ownership of any file on the computer.
- In the Windows Server 2003 family, the Everyone group no longer includes the Anonymous Logon group.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ownership
Take ownership of files or other objects
Restore files and directories

# View effective permissions on files and folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view effective permissions on files and folders**

1. Open Windows Explorer, and then locate the file or folder for which you would like to view effective permissions.
2. Right-click the file or folder, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, click the **Effective Permissions** tab, and then click **Select**.
4. In **Enter the object name to select (examples)**, enter the name of a user or group and then click **OK**. The selected check boxes indicate the effective permissions of the user or group for that file or folder.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- If the specified object grants access to the Everyone group, the Authenticated Users group or the Local Users group, then the effective rights will always include those permissions, except when the specified user or group is the Anonymous group. Windows XP and the Windows Server 2003 family no longer include the Anonymous users in the Everyone group.
- The **Effective Permissions** tab shows information that is calculated from the existing permissions entries. Therefore, the information displayed on that page is read-only and does not support changing a user's permissions by selecting or clearing permission check boxes.
- Permissions are set only on drives formatted to use NTFS.
- **Share permissions** are not part of the effective permissions calculation. Access to shares can be denied though share permissions even when access is allowed through NTFS permissions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Effective Permissions tool
Security identifiers

# Access Control Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Access Control:

- Access control overview
- Understanding Access Control
- Using Access Control
- Access Control Resources

# Access control overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Access control overview**
Access control is the process of authorizing users, groups, and computers to access objects on the network. Key concepts that make up access control are permissions, user rights, and object auditing.

**Permissions**
Permissions define the type of access granted to a user or group for an object or object property. For example, the Finance group can be granted Read and Write permissions for a file named Payroll.dat.

Permissions are applied to any secured objects such as files, Active Directory® objects, or registry objects. Permissions can be granted to any user, group, or computer. It is a good practice to assign to groups.

You can assign permissions for objects to:

- Groups, users, and Security identifiers in the domain.
- Groups and users in that domain and any trusted domains.
- Local groups and users on the computer where the object resides.

The permissions attached to an object depend on the type of object. For example, the permissions that can be attached to a file are different from those that can be attached to a registry key. Some permissions, however, are common to most types of objects. These common permissions are:

- Read permissions
- Modify permissions
- Change owner
- Delete

When you set up permissions, you specify the level of access for groups and users. For example, you can let one user read the contents of a file, let another user make changes to the file, and prevent all other users from accessing the file. You can set similar permissions on printers so that certain users can configure the printer and other users can only print from it.

If you need to change the permissions on an individual object, you can simply start the appropriate tool and change the properties for that object. For example, to change the permissions on a file, you can run Windows Explorer, right-click the file name, and click **Properties**. On the **Security** tab, you can change permissions on the file. For more information, see Permissions.

**Ownership of objects**

An owner is assigned to an object when that object is created. By default, the owner is the creator of the object. No matter what permissions are set on an object, the owner of the object can always change the permissions on an object. For more information, see Ownership.

**Inheritance of permissions**

Inheritance allows administrators to easily assign and manage permissions. This feature automatically causes objects within a container to inherit all the inheritable permissions of that container. For example, the files within a folder, when created, inherit the permissions of the folder. Only permissions marked to be inherited will be inherited.

**User rights**
User rights grant specific privileges and logon rights to users and groups in your computing environment. For information on user rights, see the following: For more information, see User rights.

**Object auditing**

You can audit users' access to objects. You can then view these security-related events in the security log with the Event Viewer. For more information, see Auditing overview.

For more information about authorization and access control, see the Windows Server 2003 Technical Reference Web site. For information about authorization strategy, see "Designing a Resource Authorization Strategy" at the Microsoft Windows Server 2003 Deployment Kit Web site.

# Understanding Access Control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding access control**
This section covers:

- Permissions and security descriptors
- User rights
- Explicit vs. inherited permissions
- Ownership
- Effective Permissions tool
- How inheritance affects file and folder permissions
- Security identifiers

# Permissions and security descriptors

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Permissions and security descriptors**

Every container and object on the network has a set of access control information attached to it. Known as a security descriptor, this information controls the type of access allowed to users and groups. The security descriptor is automatically created along with the container or object that is created. A typical example of an object with a security descriptor is a file.

Permissions are defined within an object's security descriptor. Permissions are associated with, or assigned to, specific users and groups. For example, for the file Temp.dat, the Administrator group might be assigned read, write, and delete permissions, while the Operator group might be assigned Read and Write permissions only.

Each assignment of permissions to a user or group is known as a permission entry, which is a type of access control entry (ACE). The entire set of permission entries in a security descriptor is known as a permission set or access control list (ACL). Thus, for a file named Temp.dat, the permission set includes two permission entries, one for the Administrator group and one for the Operator group.

# User rights

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User rights**

Administrators can assign specific rights to group accounts or to individual user accounts. These rights authorize users to perform specific actions, such as logging on to a system interactively or backing up files and directories. User rights are different from permissions because user rights apply to user accounts, and permissions are attached to objects. For information on permissions, see How inheritance affects file and folder permissions.

User rights define capabilities at the local level. Although user rights can apply to individual user accounts, user rights are best administered on a group account basis. This ensures that a user logging on as a member of a group automatically inherits the rights associated with that group. By assigning user rights to groups rather than individual users, you simplify the task of user account administration. When users in a group all require the same user rights, you can assign the set of user rights once to the group, rather than repeatedly assigning the same set of user rights to each individual user account.

User rights that are assigned to a group are applied to all members of the group while they remain members. If a user is a member of multiple groups, the user's rights are cumulative, which means that the user has more than one set of rights. The only time that rights assigned to one group might conflict with those assigned to another is in the case of certain logon rights. In general user rights assigned to one group do not conflict with the rights assigned to another group. To remove rights from a user, the administrator simply removes the user from the group. In this case, the user no longer has the rights assigned to that group.

There are two types of user rights: privileges, such as the right to back up files and directories, and logon rights, such as the right to log on to a system locally.

For more information, see Privileges, Logon rights and Security Configuration Manager tools.

# Explicit vs. inherited permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Explicit vs. inherited permissions**

There are two types of permissions: explicit permissions and inherited permissions.

- Explicit permissions are those that are set by default when the object is created, by user action.
- Inherited permissions are those that are propagated to an object from a parent object. Inherited permissions ease the task of managing permissions and ensure consistency of permissions among all objects within a given container.

By default, objects within a container inherit the permissions from that container when the objects are created. For example, when you create a folder called MyFolder, all subfolders and files created within MyFolder automatically inherit the permissions from that folder. Therefore, MyFolder has explicit permissions, while all subfolders and files within it have inherited permissions.

**Notes**

- Inherited Deny permissions do not prevent access to an object if the object has an explicit Allow permission entry.
- Explicit permissions take precedence over inherited permissions, even inherited Deny permissions.

For more information on inherited permissions, see How inheritance affects file and folder permissions.

# Ownership

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ownership**
Every object has an owner, whether in an NTFS volume or Active Directory. The owner controls how permissions are set on the object and to whom permissions are granted.

**Important**

- An administrator who needs to repair or change permissions on a file must begin by taking ownership of the file.

By default, in the Windows Server 2003 family, the owner is the Administrators group. The owner can always change permissions on an object, even when denied all access to the object.

Ownership can be taken by:

- An administrator. By default, the Administrators group is given the **Take ownership of files or other objects** user right.
- Anyone or any group who has the **Take ownership** permission on the object in question.
- A user who has the **Restore files and directories** privilege.

Ownership can be transferred in the following ways:

- The current owner can grant the **Take ownership** permission to another user, allowing that user to take ownership at any time. The user must actually take ownership to complete the transfer.
- An administrator can take ownership.
- A user who has the **Restore files and directories** privilege can double-click **Other users and groups** and choose any user or group to assign ownership to.

For more information, see Take ownership of a file or folder and Restore files and directories.

# Effective Permissions tool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Effective Permissions tool**

If you would like to find out what permissions a user or group has on an object, you can use the Effective Permissions tool. It calculates the permissions that are granted to the specified user or group. The calculation takes the permissions in effect from group membership into account, as well as any permissions inherited from the parent object. It looks up all domain and local groups in which the user or group is a member.

The Everyone group will always be included, as long as the selected user or group is not a member of the Anonymous Logon group. In the Windows Server 2003 family, the Everyone group no longer includes Anonymous Logon.

**Important**

- The Effective Permissions tool only produces an approximation of the permissions that a user has. The actual permissions the user has may be different, since permissions can be granted or denied based on how a user logs on. This logon-specific information cannot be determined by the Effective Permissions tool, since the user is not logged on; therefore, the effective permissions it displays reflect only those permissions specified by the user or group and not the permissions specified by the logon.

  For example, if a user is connected to this computer via a file share, then the logon for that user is marked as a *network logon*. Permissions can be granted or denied to the well-known security ID (SID) Network which the connected user receives, so a user has different permissions when logged on locally than when logged on over a network. For information about shared folders and share permissions, see Shared Folders and Share permissions.

  For information about granting access for effective permissions, see article Q331951 in the Microsoft Knowledge Base.

**Factors that are used to determine effective permissions**

- Global group membership
- Local group membership
- Local permissions
- Local privileges

**Factors that are not used to determine effective permissions**

The following well-known SIDs that are available in the Windows Server 2003 family are not used to determine effective permissions:

- Anonymous Logon
- Batch, Creator Group
- Dialup
- Enterprise Domain Controllers
- Interactive
- Network
- Proxy

- Restricted
- Remote
- Service
- System
- Terminal Server User
- Other Organization
- This Organization

Also, share permissions are not part of the effective permissions calculation. Access to shares can be denied through share permissions even when access is allowed through NTFS permissions.

**Factors that are not used for objects that are accessed remotely**

- Local group membership
- Local privileges
- Share permissions

**Retrieving effective permissions**
Accurate retrieval of the above information requires permission to read the membership information. If the specified user or group is a domain object, you must have permission to read the object's group information on the domain. Here are some relevant default domain permissions:

- Domain administrators have permission to read membership information on all objects.
- Local administrators on a workstation or stand-alone server cannot read membership information for a domain user.
- Authenticated domain users can only read membership information when the domain is in Pre-Windows 2000 compatibility mode.

For more information, see View effective permissions on files and folders and Security identifiers.

# How inheritance affects file and folder permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How inheritance affects file and folder permissions**

After you set permissions on a parent folder, new files and subfolders that are created in the folder inherit these permissions. If you do not want them to inherit permissions, select **This folder only** in **Apply onto** when you set up special permissions for the parent folder. In cases where you want to prevent only certain files or subfolders from inheriting permissions, right-click the file or subfolder, click **Properties**, click the **Security** tab, click **Advanced**, and then clear the **Inherit from parent the permission entries that apply to child objects. Include these with entries explicitly defined here.** check box.

If the check boxes appear shaded, the file or folder has inherited permissions from the parent folder. There are three ways to make changes to inherited permissions:

- Make the changes to the parent folder, and then the file or folder will inherit these permissions.
- Select the opposite permission (**Allow** or **Deny**) to override the inherited permission.
- Clear the **Inherit from parent the permission entries that apply to child objects. Include these with entries explicitly defined here.** check box. You can then make changes to the permissions or remove the user or group from the permissions list. However, the file or folder will no longer inherit permissions from the parent folder.

In most cases, **Deny** overrides **Allow** unless a folder is inheriting conflicting settings from different parents. In that case, the setting inherited from the parent closest to the object in the subtree will have precedence.

Only inheritable permissions are inherited by child objects. When setting permissions on the parent object, you can decide whether folders or subfolders can inherit them with **Apply onto**. For more information, see Selecting where to apply permissions.

You can find out what permissions a user or group has on an object with the Effective Permissions tool.

For more information, see Set, view, change, or remove permissions on files and folders.

# Security identifiers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security identifiers**

Security identifiers (SIDs) are numeric values that identify a user or group. For each access control entry (ACE), there exists a SID that identifies the user or group for whom access is allowed, denied, or audited.

**Well-known security identifiers (special identities)**

| Well-known SID | Description |
|---|---|
| Anonymous Logon (S-1-5-7) | A user who has connected to the computer without supplying a user name and password. |
| Authenticated Users (S-1-5-11) | Includes all users and computers whose identities have been authenticated. Authenticated Users does not include Guest even if the Guest account has a password. |
| Batch (S-1-5-3) | Includes all users who have logged on through a batch queue facility such as task scheduler jobs. |
| Creator Owner (S-1-3-0) | A placeholder in an inheritable access control entry (ACE). When the ACE is inherited, the system replaces this SID with the SID for the object's current owner. |
| Creator Group (S-1-3-1) | A placeholder in an inheritable ACE. When the ACE is inherited, the system replaces this SID with the SID for the primary group of the object's current owner. |
| Dialup (S-1-5-1) | Includes all users who are logged on to the system through a dial-up connection. |
| Everyone (S-1-1-0) | On computers running Windows Server 2003 operating systems, Everyone includes Authenticated Users and Guest. On computers running earlier versions of the operating system, Everyone includes Authenticated Users and Guest plus Anonymous Logon.<br><br>For more information, see Differences in default security settings. |
| Interactive (S-1-5-4) | Includes all users logging on locally or through a Remote Desktop connection. |
| Local System (S-1-5-18) | A service account that is used by the operating system. |
| Network (S-1-5-2) | Includes all users who are logged on through a network connection. Access tokens for interactive users do not contain the Network SID. |
| Self (or Principal Self) (S-1-5-10) | A placeholder in an ACE on a user, group, or computer object in Active Directory. When you grant permissions to Principal Self, you grant them to the security principal represented by the object. During an access check, the operating system replaces the SID for Principal Self with the SID for the security principal represented by the object. |
| Service (S-1-5-6) | A group that includes all security principals that have logged on as a service. Membership is controlled by the operating system. |
| Terminal Server Users (S-1-5-13) | Includes all users who have logged on to a Terminal Services server that is in Terminal Services version 4.0 application compatibility mode. |

| | |
|---|---|
| Other Organization (S-1-5-1000) | Causes a check to ensure that a user from another forest or domain is allowed to authenticate to a particular service. |
| This Organization (S-1-5-15) | Added by the authentication server to the authentication data of a user, provided the Other Organization SID is not already present. |

For more information on the Administrators (S-1-5-32-544), Power Users (S-1-5-32-547), Backup operators (S-1-5-32-551), and Users (S-1-5-32-545) groups, see Default security settings for groups.

For more information, see Well-Known Security Identifiers at the Microsoft Resource Kits Web site.

# Using Access Control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using access control**
This section covers:

- Permissions
- Permissions for files and folders
- Active Directory object permissions
- Selecting where to apply permissions
- Changing inherited permissions
- Privileges
- Logon rights

# Permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Permissions**

Each type of object is controlled by an object manager. There is a different object manager for each type of object. The object types, their object managers, and the tools you use to manage these objects are as follows:

| Object Type | Permission Descriptions | Procedures |
|---|---|---|
| Files and folders | Permissions for files and folders | Set, view, change, or remove permissions on files and folders<br><br>or<br><br>Set, view, change, or remove special permissions |
| Shares | Share permissions | Set permissions on a shared resource |
| Registry keys | Maintain Registry Security | Add users or groups to the Permissions list |
| Services | Services permissions | Local computer security for file system, registry, system services |
| Printer | Assigning printer permissions | Set or remove permissions for a printer |
| Terminal Services Connections | Managing Permissions on Connectionsor Controlling connection access | Manage Connection Permissions |
| WMI object | Authorize WMI users and set permissions | Modify permissions or delete authorized users |
| Active Directory objects | Active Directory object permissions.<br><br>Best practices for assigning permissions on Active Directory objects. | |

# Permissions for files and folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Permissions for files and folders**

Folder permissions include **Full Control**, **Modify**, **Read & Execute**, **List Folder Contents**, **Read**, and **Write**. For information about these permissions, see File and folder permissions. Each of these permissions consists of a logical group of special permissions which are listed and defined below.

| Permission | Description |
|---|---|
| **Traverse Folder/Execute File** | **For folders**: Traverse Folder allows or denies moving through folders to reach other files or folders, even if the user has no permissions for the traversed folders. (Applies to folders only.) Traverse folder takes effect only when the group or user is not granted the **Bypass traverse checking** user right in the Group Policy snap-in. (By default, the Everyone group is given the **Bypass traverse checking** user right.)<br><br>**For files:** Execute File allows or denies running program files. (Applies to files only).<br><br>Setting the Traverse Folder permission on a folder does not automatically set the Execute File permission on all files within that folder. |
| **List Folder/Read Data** | List Folder allows or denies viewing file names and subfolder names within the folder. List Folder only affects the contents of that folder and does not affect whether the folder you are setting the permission on will be listed. (Applies to folders only.)<br><br>Read Data allows or denies viewing data in files. (Applies to files only.) |
| **Read Attributes** | Allows or denies viewing the attributes of a file or folder, such as read-only and hidden. Attributes are defined by NTFS. |
| **Read Extended Attributes** | Allows or denies viewing the extended attributes of a file or folder. Extended attributes are defined by programs and may vary by program. |
| **Create Files/Write Data** | Create Files allows or denies creating files within the folder. (Applies to folders only).<br><br>Write Data allows or denies making changes to the file and overwriting existing content. (Applies to files only.) |
| **Create Folders/Append Data** | Create Folders allows or denies creating folders within the folder. (Applies to folders only.)<br><br>Append Data allows or denies making changes to the end of the file but not changing, deleting, or overwriting existing data. (Applies to files only.) |
| **Write Attributes** | Allows or denies changing the attributes of a file or folder, such as read-only or hidden. Attributes are defined by NTFS.<br><br>The Write Attributes permission does not imply creating or deleting files or folders, it only includes the permission to make changes to the attributes of a file or folder. In order to allow (or deny) create or delete operations, see **Create Files/Write Data**, **Create Folders/Append Data**, **Delete Subfolders and Files**, and **Delete**. |
| **Write Extended Attributes** | Allows or denies changing the extended attributes of a file or folder. Extended attributes are defined by programs and may vary by program.<br><br>The Write Extended Attributes permission does not imply creating or deleting files or folders, it only includes the permission to make changes to the attributes of a file or folder. In order to allow (or deny) create or delete operations, see **Create Files/Write Data**, **Create Folders/Append Data**, **Delete Subfolders and Files**, and **Delete**. |

| | |
|---|---|
| **Delete Subfolders and Files** | Allows or denies deleting subfolders and files, even if the Delete permission has not been granted on the subfolder or file. (Applies to folders.) |
| **Delete** | Allows or denies deleting the file or folder. If you do not have Delete permission on a file or folder, you can still delete it if you have been granted Delete Subfolders and Files on the parent folder. |
| **Read Permissions** | Allows or denies reading permissions of the file or folder, such as Full Control, Read, and Write. |
| **Change Permissions** | Allows or denies changing permissions of the file or folder, such as Full Control, Read, and Write. |
| **Take Ownership** | Allows or denies taking ownership of the file or folder. The owner of a file or folder can always change permissions on it, regardless of any existing permissions that protect the file or folder. |
| **Synchronize** | Allows or denies different threads to wait on the handle for the file or folder and synchronize with another thread that may signal it. This permission applies only to multithreaded, multiprocess programs. |

For more information, see Set, view, change, or remove special permissions and Set, view, change, or remove permissions on files and folders.

**Note**

- You will not be able to access an encrypted file without the Encrypting File System (EFS) key, even if you have the necessary permissions. For information about EFS, see Encrypting File System.

# Permissions on a file server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Permissions on a file server**

Access on a shared folder is determined through two sets of permission entries; the permissions set on the share (called *share permissions*) and the permissions set on the folder (called *NTFS file and folder permissions*). Share permissions are often used for managing computers with FAT32 file systems, or other computers that don't use the NTFS file system.

Share Permissions and NTFS Permissions are independent in the sense that neither changes the other. The final access permissions on a shared folder are determined by taking into consideration both the Share permission and the NTFS permission entries. The more restrictive permissions are then applied.

The following table suggests permissions that a security-conscious administrator could grant to the Users group for certain shared folder types. The recommended permissions have been tested, and work correctly; but there are alternative approaches. For example, some experienced administrators prefer always to set share permissions to Full Control for Everyone, and to rely entirely on NTFS permissions to restrict access.

| Folder type | Share permissions | NTFS permissions |
|---|---|---|
| **Public folder**. A folder that can be accessed by everyone. | Grant Change permission to the Users group. | Grant Modify permission to the Users group. |
| **Drop folder**. A folder where users can drop confidential reports or homework assignments that only the group manager or instructor can read. | Grant the Change permission to the Users group.<br><br>Grant the Full Control permission to the group manager. | Grant the Write permission for the users' group that is applied to **This Folder only**. (This is an option available on the **Advanced** page.) For more information, see Set, view, change, or remove special permissions.<br><br>If each user needs to have certain permissions to the files that he or she dropped, you can create a permission entry for the Creator Owner Security identifiers (SID) and apply it to **Subfolder and files only**. For example, you can grant the Read and Write permission to the Creator Owner SID on the drop folder and apply it to all subfolders and files. This grants the user who dropped or created the file (the Creator Owner) the ability to read and write to the file. The Creator Owner can then access the file through the **Run** command using *****ServerName******DropFolder******FileName*.<br><br>Grant the Full Control permission for the group manager. |
| **Application folder**. A folder containing applications that can be run over the network. | Grant Read permission for the Users group. | Grant Read, Read and Execute, and List Folder Content permissions to the Users group. |
| **Home folders**. Individual folders for each user. Only the user has access to the folder. | Grant the Full Control permission to each user on their respective folder. | Grant the Full Control permission to each user for their respective folder. |

**Notes**

- Granting a user Full Control NTFS permission on a folder enables that user to take ownership of the folder unless the user is restricted in some other way. Be cautious in granting Full Control.
- If you want to manage folder access by using NTFS permissions exclusively, set Share permissions to Full Control for

Everyone. This frees you from having to think about Share permissions, but NTFS permissions are more complex than Share permissions, so using NTFS permissions correctly requires deeper understanding on your part. For more information on NTFS permissions, search for the term "NTFS permissions" on TechNet on the Microsoft Web site.

- NTFS permissions affect access both locally and remotely. NTFS permissions apply regardless of protocol. Share permissions, by contrast, apply only to network shares. Share permissions do not restrict access to any local user, or to any terminal server user, of the computer on which you have set Share permissions. Thus, Share permissions do not provide privacy between users on a computer used by several users, nor on a terminal server accessed by several users.

- By default, Everyone does not include Anonymous, so permissions applied to Everyone do not affect Anonymous. This default behavior is new for the Windows Server 2003 family.

For more information, see Set, view, change, or remove permissions on files and folders and Share a folder or drive.

# Active Directory object permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory object permissions**

- Extended rights
- Validated writes
- Property sets

**Extended rights**

These permissions are for a special operation on an Active Directory object, where the operation is not necessarily related to read or write access to a particular attribute on that object.

| Extended right | Description |
|---|---|
| **Applies to Computers and Users** | |
| Receive As<br><br>**Applies to**: Computer and User | Exchange right: allows receiving mail as a given mailbox. |
| Send As<br><br>**Applies to**: User, computer, and inetOrgPerson | Exchange right: allows sending mail as the mailbox. |
| Send To<br><br>**Applies to**: Group | Exchange right: allows sending to a mailbox. |
| Change Password | Permits changing password on user account. You do need to know the original password to be able to make the change. |
| Reset Password<br><br>**Applies to**: User, computer, and inetOrgPerson | Permits resetting password on user account. With this permission, when you reset the password, you do not need to know the original password. |
| **Applies to Domain-DNS** | |
| Replicating Directory Changes<br><br>**Applies to**: Domain-DNS | Extended right needed to replicate changes from a given NC. |
| Add/Remove Replica In Domain<br><br>**Applies to**: Domain-DNS | Extended right needed to do a replica install. |

| | |
|---|---|
| Add GUID<br><br>**Applies to**: Domain-DNS | Extended right needed at the NC root to add an object with a specific GUID. |
| Change PDC<br><br>**Applies to**: Domain-DNS | Extended right needed to change the PDC. |
| Manage Replication Topology<br><br>**Applies to**: Domain-DNS | Extended right needed to update the replication topology for a given NC. |
| Replication Synchronization<br><br>**Applies to**: Domain-DNS, DMD, Configuration | Extended right needed to synchronize replication from a given computer. |
| Generate Resultant Set of Policy (Logging)<br><br>**Applies to**: Domain-DNS, Organizational-Unit | The user who has the rights on an organizational unit or domain will be able to generate logging mode Resultant Set of Policy (RSoP) data for the users/computers within the organizational unit. |
| Generate Resultant Set of Policy (Planning)<br><br>**Applies to**: Domain-DNS, Organizational-Unit | The user who has the rights on an organizational unit or domain will be able to generate planning mode RSoP data for the users or computers within the organizational unit. |
| **Applies to NTDS-DSA** | |
| Allocate Rids | Extended right needed to request RID pool. |
| Do Garbage Collection<br><br>**Applies to**: NTDS-DSA | Extended right to force the directory service to do garbage collection. Control right to force the directory service to do garbage collection. |
| Recalculate Hierarchy<br><br>**Applies to**: NTDS-DSA | Extended right to force the directory service to recalculate the hierarchy. |
| **Applies to a Group Policy container** | |
| Apply Group Policy | Extended right to determine if a Group Policy object applies or not. For a Group Policy object to apply to a user group or computer, the **Apply Group Policy** and **Read** permissions must be set. |
| **Applies to Site** | |
| Open Connector Queue | Allows opening connector queue. |

**Validated writes**

A validated write is different from a write permission which does not allow for any value checking. The value checking, or *validation*, ensures that the value conforms to required semantics, is within a legal range of values, or undergoes some other special checking that would not be performed for a simple low-level write to the property.

| **Validated write** | **Description** |
|---|---|

| | |
|---|---|
| Add/remove self as member<br><br>**Applies to**: Group | Validated write permission to enable one to add or remove one's own account from membership of a group. |
| Validated write to DNS host name<br><br>**Applies to**: Computer | Validated write permission to enable setting of a DNS host name attribute that is compliant with the computer name and domain name. |
| Validated write to service principal name **Applies to**: Computer | Validated write permission to enable setting of the SPN attribute which is compliant to the DNS host name of the computer. |

**Property sets**
A property set consists of a group of related properties (or attributes). Granting access rights to a property set rather than to individual properties greatly improves performance and simplifies security management.

**Domain Password**
Password and account lockout properties for the domain are stored in the Directory Service as attributes of the domain object. These properties can also be managed through the user interface using the **Domain Security Policy** Group Policy object, the values are then sychronized to the Directory Service. Password policies as well as all account policies are domainwide and applied to all members of the domain.

**Applies to:** Domain

- lockOutObservationWindow
- lockoutDuration
- lockoutThreshold
- maxPwdAge
- minPwdAge
- minPwdLength
- Pwd-Properties

**Other Domain Parameters (for use by the security account manager, also called SAM)**
Property set permitting control to a list of domain attributes.

**Applies to:** Domain

- domainReplica
- forceLogoff
- modifiedCount
- oEMInformation
- serverRole
- serverState
- uASCompat

**E-mail Information**
Property set that contains user attributes that describe user e-mail information.

**Applies to:** Group, User

**General Information**
Property set containing a set of user attributes that constitute general user information.

**Applies to:** User

- Display Name
- adminDescription
- codePage
- CountryCode
- ObjectSid
- primaryGroupID
- sAMAccountName
- sAMAccountType
- sDRightsEffective
- showInAdvancedViewOnly
- sIDHistory
- UID
- comment

**Membership**
Property set containing user attributes that describe group membership information.

**Applies to:** User

- memberOf
- member

**Personal Information**
Property set containing user attributes that describe personal user information.

**Applies to:** Computer, Contact, User

- streetAddress
- homePostalAddress
- assistant
- info
- country/region name
- facsimileTelephoneNumber (fax number)

- International-ISDN-Number
- Locality-Name
- MSMQ-Digests
- mSMQSignCertificates
- Personal-Title
- Phone-Fax-Other
- Phone-Home-Other
- Phone-Home-Primary
- otherIpPhone
- ipPhonenumber
- primaryInternationalISDNNumber Phone-ISDN-Primary
- Phone-Mobile-Other (otherMobile)
- Phone-Mobile-Primary
- Phone-Office-Other (otherTelephone)
- Phone-Pager-Other
- Phone-Pager-Primary
- physicalDeliveryOfficeName
- thumbnailPhoto (Picture)
- postalCode
- preferredDeliveryMethod
- registeredAddress
- State-Or-Province-Name
- Street-Address
- telephoneNumber
- teletexTerminalIdentifier
- telexNumber
- primaryTelexNumber
- userCert

- User-Shared-Folder
- User-Shared-Folder-Other
- userSMIMECertificate
- x121Address
- X509-Cert

**Public Information**
Property set containing user attributes that describe user public information.

**Applies to:** Computer, User

- Additional-Information notes
- Allowed-Attributes
- allowedAttributesEffective
- allowedChildClasses
- allowedChildClassesEffective
- altSecurityIdentities
- Common-Name (cn)
- company
- department
- description
- displayNamePrintable
- division
- E-mail-Addresses
- givenName
- initials
- legacyExchangeDN
- manager
- msDS-Approx-Immed-Subordinates
- msDS-Auxiliary-Classes
- distinguishedName (Obj-Dist-Name)

- Object-Category
- Object-Class
- Object-Guid
- Organization-Name
- Organizational-Unit-Name
- otherMailbox
- Proxy-Addresses
- RDN name
- Reports (directReports)
- servicePrincipalName
- showInAddressBook
- Surname
- System-Flags
- Text-Country/Region
- Title
- userPrincipalName

**RAS Information**
System Internal: Do not use or modify this right.

**Applies to:** User

- msNPAllowDialin
- msNPCallingStationID
- msRADIUSCallbackNumber
- msRADIUSFramedIPAddress
- msRADIUSFramedRoute
- msRADIUSServiceType
- tokenGroups
- Token-Groups-Global-And-Universal

**User Account Restrictions**
Property set containing user attributes that describe account restrictions.

**Applies to:** Computer, User

- accountExpires
- pwdLastSet
- userAccountControl
- userParameters
- tokenGroupsNoGCAcceptable

**User Logon**
Property set containing user attributes that describe user logon information.

**Applies to:** User

- badPwdCount
- homeDirectory
- homeDrive
- lastLogoff
- Last-Logon
- lastLogonTimestamp
- logonCount
- logonHours
- logonWorkstation
- profilePath

**Web Information Contact**
Property set containing user attributes that describe user web related information.

**Applies to:** Web Information Contact, User

- WWWHomePage
- WWW-Page-Other url

# Selecting where to apply permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Selecting where to apply permissions**

When you set advanced permissions on files and folders, you can use **Permission Entry for *File or Folder Name*** to set which folders or subfolders will inherit permissions. In this dialog box, **Apply onto** lists the locations where you can apply permissions. How these permissions are applied depends on whether you select the **Apply these permissions to objects and/or containers within this container only** check box. By default, this check box is cleared.

**Important**

- For Active Directory objects, not only will the specified objects in the Apply onto field inherit the access control entries but ALL child objects will receive a copy of that ACE. The child objects not specified in the Apply onto field will not utilize the ACE whose copy they receive but if there are enough objects that will get copies of this ACE, then that increased amount of data can cause serious performance problems to your network. For more information, see Best practices for assigning permissions on Active Directory objects in Help.

For more information, see How inheritance affects file and folder permissions.

- When the Apply these permissions to objects and/or containers within this container only check box is cleared
- When the Apply these permissions to objects and/or containers within this container only check box is selected

**When the Apply these permissions to objects and/or containers within this container only check box is cleared**

| Apply onto | Applies permissions to current folder | Applies permissions to subfolders in current folder | Applies permissions to files in current folder | Applies permissions to all subsequent subfolders | Applies permissions to files in all subsequent subfolders |
|---|---|---|---|---|---|
| This folder only | x | | | | |
| The folder, subfolders and files | x | x | x | x | x |
| This folder and subfolders | x | x | | x | |
| This folder and files | x | | x | | x |
| Subfolders and files only | | x | x | x | x |
| Subfolders only | | x | | x | |
| Files only | | | x | | x |

**When the Apply these permissions to objects and/or containers within this container only check box is selected**

| Apply onto | Applies permissions to current folder | Applies permissions to subfolders in current folder | Applies permissions to files in current folder | Applies permissions to all subsequent subfolders | Applies permissions to files in all subsequent subfolders |
|---|---|---|---|---|---|

| | | | | | |
|---|---|---|---|---|---|
| This folder only | **x** | | | | |
| The folder, subfolders and files | **x** | **x** | **x** | | |
| This folder and subfolders | **x** | **x** | | | |
| This folder and files | **x** | | **x** | | |
| Subfolders and files only | | **x** | **x** | | |
| Subfolders only | | **x** | | | |
| Files only | | | **x** | | |

For more information, see Set, view, change, or remove special permissions.

# Changing inherited permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Changing inherited permissions**

**Inheritance for all objects**
If the check boxes are shaded when you view the permissions of an object, the object has inherited permissions from the parent object. There are three ways to make changes to inherited permissions:

- Make the changes to the parent object, and then the child object will inherit these permissions.
- Select the opposite permission (Allow or Deny) to override the inherited permission.
- Clear the **Inherit from parent the permission entries that apply to child objects. Include these with entries explicitly defined here** check box. Then you can make changes to the permissions or remove users or groups from the **Permissions** list. However, the object will no longer inherit permissions from the parent object.

### Notes

- Inherited Deny permissions do not prevent access to an object if the object has an explicit Allow permission entry.
- Explicit permissions take precedence over inherited permissions, even inherited Deny permissions. For more information, see Explicit vs. inherited permissions.

On the **Advanced** page, in **Permission entries**, the **Apply To** column lists what folders or subfolders a permission is applied to. The **Inherited From** column lists where the permissions have been inherited from.

You can use the **Apply onto** field to select the folders or subfolders you would like permissions to be applied to. For more information, see Explicit vs. inherited permissions.

If the **Special Permissions** entry in **Permissions for *User or Group*** is shaded, it does not imply that this permission has been inherited. This means that a special permission has been selected. For more information, see Set, view, change, or remove special permissions and Selecting where to apply permissions.

**Inheritance for Active Directory objects**
For Active Directory objects, when using the **Apply Onto** option to control inheritance, be aware that not only will the objects specified in the **Apply onto** field inherit that access control entry (ACE) but *all* child objects will also receive a copy of that ACE. The child objects that are not specified in the **Apply onto** field will receive copies of the ACE but will not enforce it. If there are enough objects that will get copies of this ACE, then that increased amount of data can cause serious performance problems to your network.

If you assign permissions to a parent object and want child objects to inherit these permission entries, you can keep performance optimal by making sure all the child objects have identical access control lists (ACLs). In the Windows Server 2003 family single-instancing allows for the Active Directory to store only one copy of all identical ACLs. By creating ACLs that many objects can utilize, you can preserve the performance of your network.

# Privileges

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Privileges**

To ease the task of user account administration, you should assign privileges primarily to group accounts, rather than to individual user accounts. When you assign privileges to a group account, users are automatically assigned those privileges when they become a member of that group. This method of administering privileges is far easier than assigning individual privileges to each user account when the account is created.

The following table lists and describes the privileges that can be granted to a user.

| Privilege | Description |
|---|---|
| Act as part of the operating system | This user right allows a process to impersonate any user without authentication. The process can therefore gain access to the same local resources as that user.<br><br>Processes that require this privilege should use the *LocalSystem* account, which already includes this privilege, rather than using a separate user account with this privilege specially assigned. If your organization only uses servers that are members of the Windows Server 2003 family, you do not need to assign this privilege to your users. However, if your organization uses servers running Windows 2000 or Windows NT 4.0, you might need to assign this privilege to use applications that exchange passwords in plaintext.<br><br>**Default:** Local System. |
| Add workstations to a domain | This security setting determines which groups or users can add workstations to a domain.<br><br>This security setting is valid only on domain controllers. By default, any authenticated user has this right and can create up to 10 computer accounts in the domain.<br><br>Adding a computer account to the domain allows the computer to participate in Active Directory-based networking. For example, adding a workstation to a domain enables that workstation to recognize accounts and groups that exist in Active Directory.<br><br>**Default:** Authenticated Users on domain controllers. |
| Adjust memory quotas for a process | This privilege determines who can change the maximum memory that can be consumed by a process.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>**Default:** Administrators. |
| Back up files and directories | This user right determines which users can bypass file and directory, registry, and other persistent object permissions for the purposes of backing up the system.<br><br>**Default:** Administrators and Backup Operators. |

| | |
|---|---|
| Bypass traverse checking | This user right determines which users can traverse directory trees even though the user may not have permissions on the traversed directory. This privilege does not allow the user to list the contents of a directory, only to traverse directories.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>**Default:**<br><br>• On workstations and servers:<br>  • Administrators<br>  • Backup Operators<br>  • Power Users<br>  • Users<br>  • Everyone<br><br>• On domain controllers:<br>  • Administrators<br>  • Authenticated Users |
| Change the system time | This user right determines which users and groups can change the time and date on the internal clock of the computer. Users that are assigned this user right can affect the appearance of event logs. If the system time is changed, events that are logged will reflect this new time, not the actual time that the events occurred.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>**Default:**<br><br>• On workstations and servers:<br>  • Administrators<br>  • Power Users<br><br>• On domain controllers:<br>  • Administrators<br>  • Server Operators |
| Create a pagefile | Allows the user to create and change the size of a pagefile. This is done by specifying a paging file size for a particular drive under **Performance Options** on the **Advanced** tab of **System properties**.<br><br>**Default setting**: Administrators |

| | |
|---|---|
| Create a token object | Allows a process to create a token which it can then use to get access to any local resources when the process uses NtCreateToken() or other token-creation APIs.<br><br>It is recommended that processes requiring this privilege use the LocalSystem account, which already includes this privilege, rather than using a separate user account with this privilege specially assigned<br><br>**Default setting**: No one |
| Create global objects | This security setting determines which accounts are allowed to create global objects in a terminal services session.<br><br>**Default:** Administrators and Local System. |
| Create permanent shared objects | Allows a process to create a directory object in the Windows Server 2003 family and Windows XP Professional object manager. This privilege is useful to kernel-mode components that extend the object namespace. Components that are running in kernel mode already have this privilege inherently; it is not necessary to assign them the privilege.<br><br>**Default setting**: No one |
| Debug programs | This user right determines which users can attach a debugger to any process or to the kernel. Developers who are debugging their own applications to not need to be assigned this user right. Developers who are debugging new system components will need this user right to be able to do so. This user right provides complete access to sensitive and critical operating system components.<br><br>**Default setting**:<br><br>● Administrators<br><br>● Local System |
| Enable computer and user accounts to be trusted for delegation | This security setting determines which users can set the **Trusted for Delegation** setting on a user or computer object.<br><br>The user or object that is granted this privilege must have write access to the account control flags on the user or computer object. A server process running on a computer (or under a user context) that is trusted for delegation can access resources on another computer using the delegated credentials of a client, as long as the account of the client does not have the **Account cannot be delegated** account control flag set.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>**Default setting**: On domain controllers:<br><br>● Administrators |
| Force shutdown from a remote system | This security setting determines which users are allowed to shut down a computer from a remote location on the network. Misuse of this user right can result in a denial of service.<br><br>This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.<br><br>**Default:**<br><br>● On workstations and servers: Administrators.<br><br>● On domain controllers: Administrators, Server Operators. |

| | |
|---|---|
| Generate security audits | This security setting determines which accounts can be used by a process to add entries to the security log. The security log is used to trace unauthorized system access. Misuse of this user right can result in the generation of many auditing events, potentially hiding evidence of an attack or causing a denial of service if the **Audit: Shut down system immediately if unable to log security audits** security policy setting is enabled. For more information, see Audit: Shut down system immediately if unable to log security audits.<br><br>**Default:** Local System. |
| Impersonate a client after authentication | This security setting determines which accounts are allowed to impersonate other accounts.<br><br>**Default:** Administrators and Service. |
| Increase scheduling priority | This security setting determines which accounts can use a process with Write property access to another process to increase the execution priority assigned to the other process. A user with this privilege can change the scheduling priority of a process through the Task Manager user interface.<br><br>**Default:** Administrators. |
| Load and unload device drivers | This user right determines which users can dynamically load and unload device drivers or other code in to kernel mode. This user right does not apply to Plug and Play device drivers. It is recommended that you do not assign this privilege to other users. Instead, use the `StartService()` API.<br><br>**Default setting**: Administrators. It is recommended that you not assign this privilege to any other user. Device drivers run as trusted (or highly privileged) programs. |
| Lock pages in memory | This security setting determines which accounts can use a process to keep data in physical memory, which prevents the system from paging the data to virtual memory on disk. Exercising this privilege could significantly affect system performance by decreasing the amount of available random access memory (RAM).<br><br>**Default:** None. Certain system processes have the privilege inherently. |
| Manage auditing and security log | This security setting determines which users can specify object access auditing options for individual resources, such as files, Active Directory objects, and registry keys.<br><br>This security setting does not allow a user to enable file and object access auditing in general. For such auditing to be enabled, the Audit object access setting in Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policies must be configured.<br><br>You can view audited events in the security log of the Event Viewer. A user with this privilege can also view and clear the security log.<br><br>**Default:** Administrators. |
| Modify firmware environment values | This security setting determines who can modify firmware environment values. Firmware environment variables are settings stored in the nonvolatile RAM of non-x86-based computers. The effect of the setting depends on the processor.<br><br>• On x86-based computers, the only firmware environment value that can be modified by assigning this user right is the Last Known Good Configuration setting, which should only be modified by the system.<br><br>• On Itanium-based computers, boot information is stored in nonvolatile RAM. Users must be assigned this user right to run **bootcfg.exe** and to change the Default Operating System setting on **Startup and Recovery** in **System properties**.<br><br>• On all computers, this user right is required to install or upgrade Windows.<br><br>**Default setting**:<br><br>• Administrators<br><br>• Local System |

| | |
|---|---|
| profile a single process | This security setting determines which users can use performance monitoring tools to monitor the performance of nonsystem processes.<br><br>**Default:** Administrators, Power users, Local System. |
| profile system performance | This security setting determines which users can use performance monitoring tools to monitor the performance of system processes.<br><br>**Default:** Administrators, Local System. |
| Remove computer from docking station | This security setting determines whether a user can undock a portable computer from its docking station without logging on.<br><br>If this policy is enabled, the user must log on before removing the portable computer from its docking station. If this policy is disabled, the user may remove the portable computer from its docking station without logging on.<br><br>**Default:** Disabled. |
| Replace a process level token | Determines which user accounts can initiate a process to replace the default token associated with a started subprocess.<br><br>This user right is defined in the Default Domain Controller Group Policy object and in the local security policy of workstations and servers.<br><br>**Default setting**: Local Service and Network Service. |
| Restore files and directories | This security setting determines which users can bypass file, directory, registry, and other persistent objects permissions when restoring backed up files and directories and determines which users can set any valid security principal as the owner of an object.<br><br>Specifically, this user right is similar to granting the following permissions to the user or group in question on all files and folders on the system:<br><br>● Traverse Folder/Execute File<br><br>● Write<br><br>**Default:**<br><br>● Workstations and servers: Administrators, Backup Operators.<br><br>● Domain controllers: Administrators, Backup Operators, Server Operators. |
| Shut down the system | This security setting determines which users who are logged on locally to the computer can shut down the operating system using the **Shut Down** command. Misuse of this user right can result in a denial of service.<br><br>**Default:**<br><br>● Workstations: Administrators, Backup Operators, Power Users, Users.<br><br>● Servers: Administrators, Backup Operators, Power Users.<br><br>● Domain controllers: Account Operators, Administrators, Backup Operators, Server Operators, Print Operators. |
| Synchronize directory service data | This security setting determines which users and groups have the authority to synchronize all directory service data. This is also known as Active Directory synchronization.<br><br>**Defaults:** None. |

| | |
|---|---|
| Take ownership of files or other objects | This security setting determines which users can take ownership of any securable object in the system, including Active Directory objects, files and folders, printers, registry keys, processes, and threads.<br><br>**Default setting**: Administrators |

Some privileges can override permissions set on an object. For example, a user logged on to a domain account as a member of the Backup Operators group has the right to perform backup operations for all domain servers. However, this requires the ability to read all files on those servers, even files on which their owners have set permissions that explicitly deny access to all users, including members of the Backup Operators group. A user right--in this case, the right to perform a backup--takes precedence over all file and directory permissions.

For more information, see Logon rights, User rights and Security Configuration Manager tools.

**Note**

- At a command prompt, you can type **whoami /priv** to see your privileges. For more information about the **whoami** command, see Whoami.

# Logon rights

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logon rights**

The following table lists and describes logon rights.

| Logon Right | Description |
|---|---|
| Access this computer from a network | This user right determines which users and groups are allowed to connect to the computer over the network. Terminal Services are not affected by this user right.<br>**Default:**<br>• On workstations and servers:<br>  • Administrators<br>  • Backup Operators<br>  • Power Users<br>  • Users<br>  • Everyone<br>• On domain controllers:<br>  • Administrators<br>  • Authenticated Users<br>  • Everyone |
| Allow log on locally | This logon right determines which users can interactively log on to this computer. Logons initiated by pressing CTRL+ALT+DEL on the attached keyboard requires the user to have this logon right. Additionally this logon right may be required by some service or administrative applications that can log on users. If you define this policy for a user or group, you must also give the **Administrators** group this right.<br>**Default:**<br>• On workstations and servers: Administrators, Backup Operators, Power Users, Users, and Guest.<br>• On domain controllers: Account Operators, Administrators, Backup Operators, Print Operators, and Server Operators. |
| Allow log on through Terminal Services | This security setting determines which users or groups have permission to log on as a Terminal Services client.<br>**Default:**<br>• On workstation and servers: Administrators, Remote Desktop Users.<br>• On domain controllers: Administrators. |

| | |
|---|---|
| Deny access to this computer from network | This security setting determines which users are prevented from accessing a computer over the network. This policy setting supersedes the Access this computer from the network policy setting if a user account is subject to both policies.<br><br>**Default:** No one. |
| Deny log on as a batch job | This security setting determines which accounts are prevented from being able to log on as a batch job. This policy setting supersedes the Log on as a batch job policy setting if a user account is subject to both policies.<br><br>**Default:** None. |
| Deny logon as a service | This security setting determines which service accounts are prevented from registering a process as a service. This policy setting supersedes the Log on as a service policy setting if an account is subject to both policies.<br><br>**Note**<br><br>• This security setting does not apply to the System, Local Service, or Network Service accounts.<br><br>**Default:** None. |
| Deny log on locally | This security setting determines which users are prevented from logging on at the computer. This policy setting supersedes the Allow log on locally policy setting if an account is subject to both policies.<br><br>**Important**<br><br>• If you apply this security policy to the Everyone group, no one will be able to log on locally.<br><br>**Default:** None. |
| Deny log on through Terminal Services | This security setting determines which users and groups are prohibited from logging on as a Terminal Services client.<br><br>**Default:** None. |
| Log on as a batch job | This security setting allows a user to be logged on by means of a batch-queue facility.<br><br>For example, when a user submits a job by means of the task scheduler, the task scheduler logs that user on as a batch user rather than as an interactive user.<br><br>**Note**<br><br>• In Windows 2000 Server, Windows 2000 Professional, Windows Server 2003 and Windows XP Professional, the Task Scheduler automatically grants this right as necessary.<br><br>**Default:** Local System. |
| Log on as a service | This security setting determines which service accounts can register a process as a service.<br><br>**Default:** None. |

For more information, see Privileges and Security Configuration Manager tools.

**Note**

- The default settings listed above are for Windows XP Professional and the Windows Server 2003 family.

# Access Control Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

For more information about security, see the Microsoft Windows Resource Kits Web site.

For more information about authorization and access control, see the Microsoft Windows Resource Kits Web site.

For more information about authentication, see "Logon and Authentication" at the Microsoft Windows Resource Kits Web site.

For more information about security groups, see "Designing an Authorization Strategy" at the Microsoft Windows Resource Kits Web site.

For the latest recommendations on security, see the Microsoft Security page on the Microsoft Web site.

For information about the Microsoft Security Notification Service, see the Product Security Notification page on the Microsoft Web site.

# Authorization Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authorization Manager**
Authorization Manager provides a flexible framework for integrating role-based access control into applications. It enables administrators who use those applications to provide access through assigned user roles that relate to job functions. Authorization Manager applications store authorization policy in the form of authorization stores that are stored in Active Directory or XML files and apply authorization policy at runtime.

- Before using Authorization Manager, see Checklist: Before you start using Authorization Manager.
- For help with specific tasks, see Authorization Manager How To....
- For general background information, see Authorization Manager Concepts.

# Checklist: Before you start using Authorization Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Before you start using Authorization Manager**

| Step | References |
|---|---|
| Review fundamental security concepts. | Security |
| Review related tools such as the Delegation of Control Wizard for comparison with Authorization Manager. | Delegating administration; Authorization Manager and role-based administration overview |
| Read background material about Authorization Manager. | Authorization Manager Concepts |
| Read about the difference between administrator mode and developer mode in Authorization Manager. | Authorization Manager and role-based administration overview; Set Authorization Manager options |
| Learn about authorization stores and applications, and about the difference between authorization stores that are stored as XML files and authorization stores that are stored in Active Directory. | Authorization stores and applications |
| (For developers) Review the latest information about Authorization Manager and authorization-related application programming interfaces (APIs) in the Platform SDK. | MSDN Library at the Microsoft Web site; TechNet at the Microsoft Web site; Platform SDK at the Microsoft Web site |
| Before you can use Authorization Manager for authorization stores that are stored in Active Directory, you must raise the Active Directory domain functional level to Windows Server 2003 .<br><br>**Caution**<br><br>• Raising the domain functional level is an irreversible step. | Domain and forest functionality; Raise the domain functional level; Raising domain and forest functional levels |
| Open Authorization Manager. | Start Authorization Manager |

# Authorization Manager How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Start Authorization Manager
- Set Authorization Manager options
- Assign an application group to a role
- Assign a Windows user or group to a role
- Add an authorization rule
- Refresh Authorization Manager
- Reload an authorization store
- Work with Authorization Stores
- Work with Applications
- Work with Role, Task, and Operation Definitions
- Work with Scopes

# Start Authorization Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Authorization Manager**

- Open Authorization Manager.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To open Authorization Manager at a command prompt, type:

  **azman.msc**

  To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Authorization Manager opens without a default authorization store. To use Authorization Manager, you need to create or open an authorization store. For more information, see Related Topics.
- If you want to configure your own Microsoft Management Console (MMC) and include Authorization Manager as part of that console, see Related Topics. If you create and save your own MMC console, it opens to the configuration at the time you saved the file. It can open to an authorization store.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Work with Authorization Stores
Authorization Manager How To...
Author an MMC console file

# Set Authorization Manager options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set Authorization Manager options**

1. Open Authorization Manager.
2. In the console tree, right-click **Authorization Manager**, and then click **Options**.
3. In the **Options** dialog box, click either **Developer mode** or **Administrator mode**.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- In developer mode, you can create, deploy, and maintain applications. With developer mode, you have unrestricted access to all features.
- In administrator mode, you can deploy and maintain applications. With administrator mode, you have access most Authorization Manager features. In administrator mode, you cannot do the following:
    - Create new authorization stores
    - Create new applications
    - Define or modify operations
    - Change an application name
    - Change the version information for an application
- Administrator mode is the default. Before you can use administrator mode, you need to provide an application that supports roles, includes all of the necessary operation and task definitions, includes its own authorization store, and is ready for use in Authorization Manager. After you set up an application in developer mode, you do not have to use Authorization Manager in developer mode for that application. You can work in authorization mode instead.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...

# Assign an application group to a role

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign an application group to a role**

1. Open Authorization Manager.

2. In the console tree, right-click the role to which you want to assign the application group, and then click **Assign Application Groups**. The role can be at the application or scope level.

   **Application level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Role Assignments/*Role*

   **Scope level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/*Scope*/Roles Assignments/*Role*

3. In the **Add Groups** dialog box, select the check boxes next to the application groups that you want to add.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**

- This is one of the most common Authorization Manager procedures for administrators.

- To carry out this procedure, a role assignment needs to exist. For more information about creating a role assignment, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create a role assignment within an application
Create a role assignment within a scope

# Assign a Windows user or group to a role

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign a Windows user or group to a role**

1. Open Authorization Manager.

2. In the console tree, right-click the role to which you want to assign a user or group, and then click **Assign Windows Users and Groups**. The role can be at the application or scope level.

   **Application level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Role Assignments/*Role*

   **Scope level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/*Scope*/Roles Assignments/*Role*

3. In the **Select Users, Computers, or Groups** dialog box, type the name of the user or group that you want to assign to the role.

4. Click **Check Names**, and then click **OK**.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**

- This is one of the most common Authorization Manager procedures for administrators.

- To carry out this procedure, a role assignment needs to exist. For more information about creating a role assignment, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create a role assignment within an application
Create a role assignment within a scope

# Add an authorization rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an authorization rule**

1. Open Authorization Manager.

2. In the console tree, right-click the role definition or task definition to which you want to add an authorization rule, and then click **Properties**. The role or task definition can be at the application or scope level.

   **Application level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Definitions/Role Definitions (or Task Definitions)/*Role definition* (or *Task definition*)

   **Scope level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/*Scope*/Definitions/Role Definitions (or Task Definitions)/*Role definition* (or *Task definition*)

3. Click the **Definition** tab, and then click **Authorization Rule**.

4. In the **Authorization Rule** dialog box, in **Script path**, type the complete path to the authorization rule, including the file name extension.

5. Under **Script type**, do one of the following:

   - If your script has a .vbs file name extension, click **VBScript**.

   - If your script has a .js file name extension, click **JScript**.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**

- To perform this procedure, you need a role definition or a task definition. For more information about creating a role assignment, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager

Authorization Manager How To...
Create a role assignment within an application
Create a role assignment within a scope

# Refresh Authorization Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To refresh Authorization Manager**

1. Open Authorization Manager.
2. In the console tree, right-click the item that you want to refresh, and then click **Refresh**.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- This procedure updates the user interface from the authorization store cache. If more than one process is modifying an authorization store and you refresh an item in Authorization Manager, the refreshed item might remain unsynchronized with the authorization store. To synchronize Authorization Manager with an actual authorization store, not just the cached copy, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Reload an authorization store

# Reload an authorization store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reload an authorization store**

1. Open Authorization Manager.
2. In the console tree, right-click the authorization store that you want to reload, and then click **Reload**.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- This procedure updates Authorization Manager from the authorization store, not from the cache. If you want do a faster update that might be less thorough, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Refresh Authorization Manager

# Work with Authorization Stores

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with authorization stores**

- Create an authorization store
- Edit the properties of an authorization store
- Create a group within an authorization store
- Edit the properties of a group within an authorization store
- Delegate an authorization store
- Audit an authorization store

# Create an authorization store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an authorization store**

1. Open Authorization Manager.
2. In the console tree, right-click **Authorization Manager**, and then click **Options**.
3. In the **Options** dialog box, click **Developer mode**.
4. In Authorization Manager, right-click **Authorization Manager**, and then click **New Authorization Store**.
5. In the **New Authorization Store** dialog box, click either **Active Directory** or **XML file**.
6. In **Store name**, type the authorization store name or click **Locations** to find the authorization store.
7. (Optional) In **Description**, type a description for the new authorization store.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you must be working in developer mode.
- If you are using an authorization store that is stored in Active Directory, use the LDAP name (for example, CN=myStore,CN=Program Data,DN=nwtraders,DN=com). If you are using an XML-based authorization store, use a path and file name that is valid at runtime (for example, C:\AuthStores\MyStore.xml).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Set Authorization Manager options
Authorization stores and applications

# Edit the properties of an authorization store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit the properties of an authorization store**

1. Open Authorization Manager.
2. In the console tree, right-click the authorization store that you want to edit, and then click **Properties**.
3. In the **Properties** dialog box, modify the settings that you want to change.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to have access to an authorization store. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store

# Create a group within an authorization store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a group within an authorization store**

1. Open Authorization Manager.
2. In the console tree, right-click **Groups**, and then click **New Application Group**.

   **Where?**

   - Authorization Manager/*Authorization store*/Groups

3. In the **New Application Group** dialog box, in **Name**, type a name for the group.
4. In **Description**, type a description for the group.
5. In **Group type**, click **Basic** or **LDAP query**.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need access to an authorization store. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store

# Edit the properties of a group within an authorization store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit the properties of a group within an authorization store**

1. Open Authorization Manager.
2. In the console tree, right-click the group that you want to edit, and then click **Properties**.

   **Where?**

   - Authorization Manager/*Authorization store*/Groups/*Group*

3. In the **Properties** dialog box, modify the settings that you want to change.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to have access to an authorization store. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store

# Delegate an authorization store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate an authorization store**

1. Open Authorization Manager.
2. In the console tree, right-click **Authorization Manager**, and then click **Open Authorization Store**.
3. In the **Open Authorization Store** dialog box, click **Active Directory**.
4. Click **Browse**, find the authorization store that you want, double-click it, and then click **OK**.
5. In the console tree, right-click the authorization store, and then click **Properties**.
6. In the **Properties** dialog box, click the **Security** tab.
7. Under **Authorization Manager user role**, click **Administrator**.
8. Under **Users and groups that are assigned to this role**, click **Add** or **Remove** to add or remove users and groups to which you want to assign the Administrator role.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- XML-based authorization stores, applications, and scopes do not support delegation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...

# Audit an authorization store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To audit an authorization store**

1. Open Authorization Manager.
2. In the console tree, right-click the authorization store that you want to audit, and then click **Properties**.
3. In the **Properties** dialog box, click the **Auditing** tab, and then do one or both of the following:
   - To audit runtime application initialization, select the **Runtime application initialization auditing** check box.
   - To audit authorization store changes, select the **Authorization store change auditing** check box.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To audit an authorization store, you must have specific privileges and security settings. For more information, see Related Topics.
- To perform this procedure, you need to have access to an authorization store. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Authorization Manager auditing
Auditing Security Events
Create an authorization store

# Work with Applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with applications**

- Create an application
- Edit the properties of an application
- Create a group within an application
- Edit the properties of a group within an application
- Create a role assignment within an application
- Edit the properties of a role assignment within an application
- Delegate an application
- Audit an application

# Create an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an application**

1. Open Authorization Manager.
2. In the console tree, right-click **Authorization Manager**, and then click **Options**.
3. In the **Options** dialog box, click **Developer mode**, and then click **OK**.
4. In the console tree, right-click the authorization store for which you want to create an application, and then click **New Application**.
5. In the **New Application** dialog box, in **Name**, type a name for the application.
6. In **Description**, type a description for the application.
7. (Optional) In **Version information**, type the application version information.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need an authorization store. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store

# Edit the properties of an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit the properties of an application**

1. Open Authorization Manager.
2. In the console tree, right-click **Authorization Manager**, click **Options**.
3. In the **Options** dialog box, click **Developer mode**, and then click **OK**.
4. In the console tree, right-click the application that you want to edit, and then click **Properties**.

   **Where?**

   - Authorization Manager/*Authorization store/Application*

5. In the **Properties** dialog box, modify the settings that you want to edit.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an existing authorization store that contains an application. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application

# Create a group within an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a group within an application**

1. Open Authorization Manager.
2. In the console tree, right-click **Groups**, and then click **New Application Group**.

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Groups

3. In the **New Application Group** dialog box, in **Name**, type a name for the group.
4. In **Description**, type a description for the group.
5. Under **Group type**, click **Basic** or **LDAP query**.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an authorization store that contains an application. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Create an authorization store

# Edit the properties of a group within an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit the properties of a group within an application**

1. Open Authorization Manager.
2. In the console tree, double-click the group that you want to edit.

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Groups/*Group*

3. In the **Properties** dialog box, modify the settings that you want to change.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an authorization store that contains an application, and the application must contain a group. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application

# Create a role assignment within an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a role assignment within an application**

1. Open Authorization Manager.
2. In the console tree, right-click **Role Assignments**, and then click **Assign Roles**.

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Role Assignments

3. In the **Add Role** dialog box, select the check boxes for the role definitions that you want to add.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- If you have not defined any roles, the **Add Role** dialog box will not contain any role definitions.
- To perform this procedure, you need to use an authorization store that contains an application. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create a role definition
Create an authorization store
Create an application

# Edit the properties of a role assignment within an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit the properties of a role assignment within an application**

1. Open Authorization Manager.
2. In the console tree, right-click the role assignment that you want to edit, and then click **Properties**.

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Role Assignments/*Role assignment*

3. In the **Properties** dialog box, modify the settings that you want to change.
4. (Optional) To edit the role definition upon which the role assignment is based, do the following:

   - In the **Properties** dialog box, click **Show Definition**.
   - In the **Definition Properties** dialog box, modify the settings that you want to change.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an authorization store that contains an application. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application

# Delegate an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate an application**

1. Open Authorization Manager.
2. In the console tree, right-click the authorization store that contains the application that you want to delegate, and then click **Properties**.
3. In the **Properties** dialog box, click the **Security** tab.
4. On the **Security** tab, under **Authorization Manager user role**, click **Delegated Users**.
5. Under **Users and groups that are assigned to this role**, click **Add** or **Remove** to add or remove users and groups.
6. In Authorization Manager, right-click the application that you want to delegate, and then click **Properties**.

   **Where?**

   - Authorization Manager/*Authorization store/Application*

7. In the **Properties** dialog box, click the **Security** tab.
8. On the **Security** tab, under **Authorization Manager user role**, click **Administrator**.
9. Under **Users and groups that are assigned to this role**, click **Add** or **Remove** to add or remove users and groups to which you want to assign the Administrator role.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- XML-based authorization stores, applications, and scopes do not support delegation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...

# Audit an application

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To audit an application**

1. Open Authorization Manager.
2. In the console tree, right-click the application that you want to audit, and then click **Properties**.

   **Where?**

   - Authorization Manager/*Authorization store/Application*

3. In the **Properties** dialog box, click the **Auditing** tab, and then do any of the following:

   - To audit runtime client context and access checks, select the **Runtime client context and access check auditing** check box.
   - To audit authorization store changes, select the **Authorization store change auditing** check box. This check box only exists for authorization stores that are stored in Active Directory.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To audit your applications, you must have specific privileges and security settings. For more information, see Related Topics.
- To perform this procedure, you need to use an authorization store that contains an application. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Auditing Security Events
Create an authorization store
Create an application

# Work with Role, Task, and Operation Definitions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with role, task, and operation definitions**

- Create an operation definition
- Edit an operation definition
- Create a task definition
- Edit a task definition
- Create a role definition
- Edit a role definition

# Create an operation definition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an operation definition**

1. Open Authorization Manager.
2. In the console tree, right-click **Authorization Manager**, and then click **Options**.
3. In the **Options** dialog box, click **Developer mode**, and then click **OK**.
4. In the console tree, right-click **Operation Definitions**, and then click **New Operation Definition**.

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Definitions/Operation Definitions

5. In the **New Operation** dialog box, in **Name**, type a name for the operation definition.
6. In **Description**, type a description for the operation definition.
7. In **Operation number**, type the specific number from 0 to 2147483647 that the application uses to identify the operation.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you must be working in developer mode. For more information about developer mode, see Related Topics.
- To perform this procedure, you need to use an authorization store that contains an application. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Set Authorization Manager options
Create an authorization store
Create an application

# Edit an operation definition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit an operation definition**

1. Open Authorization Manager.
2. In the console tree, right-click **Authorization Manager**, and then click **Options**.
3. In the **Options** dialog box, click **Developer mode**, and then click **OK**.
4. In the console tree, double-click the operation definition that you want to edit.

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Definitions/Operation Definitions/*Operation definition*

5. In the **Definition Properties** dialog box, modify the settings that you want to change.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you must be working in developer mode. For more information about developer mode, see Related Topics.
- To perform this procedure, you need to use an authorization store that contains an application. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Set Authorization Manager options
Create an application
Create an authorization store

# Create a task definition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a task definition**

1. Open Authorization Manager.
2. In the console tree, right-click **Task Definitions**, and then click **New Task Definition**. The Task Definitions folder can be at the application or scope level.

   **Application level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Definitions/Task Definitions

   **Scope level**

   **Where?**

   - Authorization Manager/Authorization store/*Application*/*Scope*/Definitions/Task Definitions

3. In the **New Task** dialog box, in **Name**, type a name for the task.
4. In **Description**, type a description of the task.
5. (Optional) Specify any tasks, operations, and authorization rules that you want to include.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an authorization store that contains an application. The application can contain a scope. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application

# Edit a task definition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit a task definition**

1. Open Authorization Manager.

2. In the console tree, double-click the task definition that you want to edit. The task definition can be at the application or scope level.

   **Application level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Definitions/Task Definitions/*Task definition*

   **Scope level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/*Scope*/Definitions/Task Definitions/*Task definition*

3. In the **Definition Properties** dialog box, modify the settings that you want to change.

4. Click the **Definition** tab, and then select any tasks, operations, and authorization rules that you want to include.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**

- To perform this procedure, you need to use an authorization store that contains an application, and the application must contain a scope. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application

# Create a role definition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a role definition**

1. Open Authorization Manager.
2. In the console tree, right-click **Role Definitions**, and then click **New Role Definition**. The Role Definitions folder can be at the application or scope level.

   **Application level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/Definitions/Role Definitions

   **Scope level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/*Scope*/Definitions/Role Definitions

3. In the **New Role** dialog box, in **Name**, type a name for the role.
4. In **Description**, type a description of the role.
5. (Optional) Specify any operations, tasks, roles, and authorization rules that you want to include.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an authorization store that contains an application. The authorization store can also contain an application that contains a scope. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create a task definition
Create a scope
Create an application
Create an authorization store

# Edit a role definition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit a role definition**

1. Open Authorization Manager.

2. In the console tree, double-click the role definition that you want to edit. The role definition can be at the application or scope level.

   **Application level**

   **Where?**

   - Authorization Manager/*Authorization Store*/*Application*/Definitions/Role Definitions/*Role definition*

   **Scope level**

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/*Scope*/Definitions/Role Definitions/*Role definition*

3. In the **Definition Properties** dialog box, modify the settings that you want to change.

4. Click the **Definition** tab, and then select any operations, tasks, roles, and authorization rules that you want to include.

**Notes**

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**

- To perform this procedure, you need to use an authorization store that contains an application, or both an application and a scope. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application
Create a scope

# Work with Scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Work with scopes**

- Create a scope
- Edit the properties of a scope
- Create a role assignment within a scope
- Edit the properties of a role assignment within a scope
- Delegate a scope
- Audit a scope

# Create a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a scope**

1. Open Authorization Manager.
2. In the console tree, right-click the application within which you want to create the scope, and then click **New Scope**.

   **Where?**

   - Authorization Manager/*Authorization store/Application*

3. In the **New Scope** dialog box, in **Name**, type a name for the scope.
4. (Optional) In **Description**, type a description for the scope, and then click **OK**.

### Important

- The scope name must be recognized by the application, and the application must resolve the name at the time of the access check. Following are some examples of scope names: "*.txt", "C:\temp", and "Computers in Sales Department." The scope name must contain only printable characters. It must contain at most 65536 characters, and it must be unique (that is, it cannot match the name of another scope currently used by the application).

### Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an authorization store that contains an application. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application

# Edit the properties of a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit the properties of a scope**

1. Open Authorization Manager.
2. In the console tree, right-click the scope that you want to edit, and then click **Properties**.

   **Where?**

   - Authorization Manager/*Authorization store/Application/Scope*

3. In the **Properties** dialog box, modify the settings that you want to change.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an authorization store that contains an application, and the application must contain a scope. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application
Create a scope

# Create a role assignment within a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a role assignment within a scope**

1. Open Authorization Manager.
2. In the console tree, right-click **Role Assignments**, and then click **Assign Role**.

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/*Scope*/Role Assignments

3. In the **Add Role** dialog box, select the check boxes next to the role definitions that you want add.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- If you have not defined roles, the **Add Role** dialog box will be empty.
- To perform this procedure, you need to use an authorization store that contains an application, and the application must contain a scope. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application
Create a scope
Create a role definition

# Edit the properties of a role assignment within a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit the properties of a role assignment within a scope**

1. Open Authorization Manager.
2. In the console tree, right-click the role assignment that you want to edit, and then click **Properties**.

   **Where?**

   - Authorization Manager/*Authorization store*/*Application*/*Scope*/Role Assignments/*Role assignment*

3. In the **Properties** dialog box, modify the settings that you want to change.
4. (Optional) To edit the role definition upon which the role assignment is based, do the following:

   - In the **Properties** dialog box, click **Show Definition**.
   - In the **Definition Properties** dialog box, modify the settings that you want to change.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an authorization store that contains an application, and the application must contain a scope. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application
Create a scope

# Delegate a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate a scope**

1. Open Authorization Manager.
2. In the console tree, right-click the authorization store that contains the parent application of the scope that you want to delegate, and then click **Properties**.
3. In the **Properties** dialog box, click the **Security** tab.
4. On the **Security** tab, under **Authorization Manager user role**, click **Delegated Users**.
5. Under **Users and groups that are assigned to this role**, click **Add** or **Remove** to add or remove users and groups.
6. In the console tree, right-click the application that contains the scope that you want to delegate, and then click **Properties**.

   **Where?**

   - Authorization Manager\Authorization store\Application

7. In the **Properties** dialog box, click the **Security** tab.
8. On the **Security** tab, under **Authorization Manager user role**, click **Delegated Users**.
9. Under **Users and groups that are assigned to this role**, click **Add** or **Remove** to add or remove users and groups.
10. In the console tree, right-click the scope that you want to delegate, and then click **Properties**.

    **Where?**

    - Authorization Manager\Authorization store\Application\Scope

11. In the **Properties** dialog box, click the **Security** tab.
12. On the **Security** tab, under **Authorization Manager user role**, click **Administrator**.
13. Under **Users and groups that are assigned to this role**, click **Add** or **Remove** to add or remove users and groups to which you want to assign the Administrator role.

**Notes**

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- To perform this procedure, you need to use an authorization store that contains an application, and the application must contain a scope. For more information, see Related Topics.
- XML-based authorization stores, applications, and scopes do not support delegation.

- If a scope within an authorization store that is stored in Active Directory contains task definitions that include authorization rules or role definitions that include authorization rules, you cannot delegate the scope.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Create an authorization store
Create an application
Create a scope

# Audit a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To audit a scope**

1. Open Authorization Manager.
2. In the console tree, right-click the scope that you want to audit, and then click **Properties**.

   **Where?**

   - Authorization Manager/*Authorization store/Application/Scope*

3. In the **Properties** dialog box, click the **Auditing** tab.
4. On the **Auditing** tab, select the **Authorization store change auditing** check box. This check box does not exit for XML-based authorization stores.

## Notes

- To open Authorization Manager, click **Start**, click **Run**, and then type: **azman.msc**
- XML-based authorization stores do not support auditing at the scope level.
- If you configure auditing at the application level or the authorization store level, this is indicated on the **Auditing** tab in the **Properties** dialog box for the scope.
- To audit your applications, you must have specific privileges and security settings. For more information, see Related Topics.
- To perform this procedure, you need to use an authorization store that contains an application, and the application must contain a scope. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Before you start using Authorization Manager
Authorization Manager How To...
Auditing Security Events
Create an authorization store
Create an application
Create a scope

# Authorization Manager Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Authorization Manager:

- Authorization Manager and role-based administration overview
- Using Authorization Manager

# Authorization Manager and role-based administration overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authorization Manager and role-based administration overview**

Network administrators, Information Technology specialists, and others responsible for computer-related infrastructure, are most effective when they help people do their jobs. The role-based management model enables you to assign users to roles. The settings that authorize users for specific roles are made automatically, by means of scripts. The scripts, called authorization rules, enable you to apply fine-grained control over the mapping between access control and the structure of your organization.

Role-based administration is often used to facilitate authorization and computer configuration. Authorization and computer configuration are two categories of roles that you can manage by using role-based administration.

- **Authorization roles** are based on a user's job function. You can use authorization roles to authorize access, to delegate administrative privileges, or to manage interaction with computer-based resources. For example, you might define a Treasurer role that includes the right to authorize expenditures and audit account transactions. Authorization Manager enables administrators to implement this type of role-based administration through applications. For more information about applications, see Authorization stores and applications.

- **Computer configuration roles** are based on a computer's function. You can use computer configuration roles to select features that you want to install, to enable services, and to select options. For example, server roles might be defined for Web servers, domain controllers, file servers, and custom server configurations that are appropriate to your organization.

**Using developer mode and administrator mode in Authorization Manager**

With Authorization Manager, you can use the following two modes:

- **Developer mode.** In developer mode, you can create, deploy, and maintain applications. You have unrestricted access to all of the Authorization Manager features.

- **Administrator mode.** This is the default mode. In administrator mode, you can deploy and maintain applications. You have access to all Authorization Manager features, but you cannot create new applications or define operations.

Applications that support roles usually create an authorization store, or use an existing authorization store, with pre-defined operations and tasks. In that case developer mode need not be used.

When you use developer mode, it is recommended that you run Authorization Manager in developer mode only until the authorization store, application, and other necessary objects are created and configured. After you initially set up Authorization Manager, run Authorization Manager in administrator mode. For more information about using developer or administrator mode, see Set Authorization Manager options.

**Comparing Authorization Manager to other management tools**

Authorization Manager is capable of implementing multiple configuration and permission changes at once. There are other management tools available with the Microsoft® Windows Server 2003 family of operating systems that are comparable to Authorization Manager.

- **ACL Editor.** The access control list (ACL) editor for access control on Active Directory directory service and Windows objects sets access control policy. Authorization Manager differs from the ACL editor by facilitating the use of role-based access control in applications that support roles. Sometimes you might not know precisely which permissions you need to set for authority to be delegated. Authorization Manager takes the guess work out of this process. For more information about access control, see Access Control.

- **Delegation of Control Wizard.** The Delegation of Control Wizard also sets multiple permissions automatically. Authorization Manager provides more scope and configuration options for role-supporting applications. For more information about delegation, see Delegating administration.

For more information about Authorization Manager, see Using Authorization Manager.

# Using Authorization Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Authorization Manager**

This section covers:

- Authorization rules
- Groups in Authorization Manager
- Authorization stores and applications
- Scopes in Authorization Manager
- Roles, tasks, and operations
- Authorization Manager auditing
- Authorization Manager delegation

# Authorization rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authorization rules**

Authorization rules are scripts, written in VBScript or JScript, that you can include in role definitions and task definitions. An authorization rule determines whether the role or task is allowed. For information about role definitions and task definitions, see Roles, tasks, and operations.

By using authorization rules, you can base authorization decisions on any conditions that a script can test. These include privileges and permissions, as well as time of day, billable expense limits, account balances, and other criteria.

With Authorization Manager, you can use authorization rules, but it is not designed for writing or debugging authorization rules. You can write your scripts in a text editor (for example, Notepad), in an integrated development environment like Visual Studio .NET, or in another application of your choice.

Authorization rules are usually written by developers. For more information about authorization-related application programming interfaces (APIs), see Authorization at the Microsoft Web site.

**VBScript example**

The following is a VBScript authorization rule that always grants permission:

```
AzBizRuleContext.BusinessRuleResult = True
```

For more information about VBScript, see VBScript at the Microsoft Web site.

**JScript example**

The following is a JScript authorization rule that always grants permission:

```
AzBizRuleContext.BusinessRuleResult = true;
```

For more information about JScript, see JScript at the Microsoft Web site.

# Groups in Authorization Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Groups in Authorization Manager**

In Authorization Manager, recipients of authorization policy are represented by the following different kinds of groups:

- **Windows Users and Groups**--These groups includes users, computers, and built-in groups for security principals. Windows users and groups are not used only in Authorization Manager. They are used throughout Windows.
- **LDAP query groups**--Membership in these groups is dynamically calculated as needed from LDAP queries. An LDAP query group is a type of application group.
- **Basic Application Groups**--These groups are defined in terms of LDAP query groups, Windows users and groups, and other basic application groups. A basic application group is a type of application group.
- **Application groups**--These groups include basic application groups and LDAP query groups as special cases. Application groups are specific to Authorization Manager role-based administration. An application group is a group of users, computers, or other security principals. An application group is not a group of applications.

**Windows Users and Groups**

For more information about groups in Active Directory, see Understanding Groups. For more information about security principals that are not stored in Active Directory, see Local Users and Groups.

**LDAP query groups**

In Authorization Manager, you can use LDAP queries to find objects in Active Directory and other LDAP compliant directories. For example, the following query finds everyone except Andy:

```
(&(objectCategory=person)(objectClass=user)(!cn=andy))
```

The following query finds all members of the DogLovers alias at northwindtraders.com:

```
(memberOf=CN=DogLovers,OU=Distribution Lists,DC=nwtraders,DC=com)
```

You can use an LDAP query to specify an LDAP query group by typing the desired LDAP query in the space provided on the **Query** tab of the **Properties** dialog box of the application group.

**Basic application groups**

Basic application groups are specific to role-based administration.

To define basic application group membership, you need to:

1. Define who is a member.
2. Define who is not a member.

Both of these steps are carried out in the same way: First, you specify zero or more Windows Users and Groups, previously defined basic application groups, or LDAP query groups.

Second, the membership of the basic application group is calculated by removing any nonmembers from the group. Authorization Manager does this automatically at runtime.

Nonmembership in a basic application group takes precedence over membership.

Circular membership definitions are not allowed, and result in the error message "Cannot add *GroupName*. The following problem occurred: A loop has been detected."

**Application groups**

When you create a new application group, you need to determine whether you want it to be an LDAP query group or a basic application group. For Authorization Manager role-based applications, any authorization you can do with Windows users and

groups can also be done with application groups. Application groups are more dynamic and flexible.

For more information about creating application groups, see Create a group within an authorization store and Create a group within an application.

For information about editing the properties of an application group, see Edit the properties of a group within an application and Edit the properties of a group within an authorization store.

# Authorization stores and applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authorization stores and applications**

With Authorization Manager, you can provide authorization services to administrators that you support by creating Authorization Manager applications that access authorization stores.

In Authorization Manager, there is no default authorization store and there is no default application. To create an authorization store, you must work in the Authorization Manager developer mode. For more information about working in developer mode, see Set Authorization Manager options.

You can store authorization stores in either XML files or Active Directory. The following table compares the two types.

| Authorization store type | Delegation support | Authorization store is specified by | Windows support |
|---|---|---|---|
| Active Directory | Supported at the authorization store, application, and scope levels | A URL, beginning with the protocol prefix MSLDAP:// or an LDAP distinguished name (for example, CN=myStore,CN=Program Data,DN=nwtraders,DN=com) | Windows Server 2003 domain functional level Active Directory domain only<br>**Important**<br>• In Windows 2000, Active Directory does not support authorization stores. |
| XML | Not supported<br>The XML file is secured as a whole by its NTFS file system access control entries (ACEs). | A URL beginning with the protocol prefix MSXML:// or a Path (for example, C:\Temp\MyStore.xml or \\ServerName\ShareName\MyStore.xml) | Any NTFS partition |

An application is specific to an authorization store, and it is always located directly under its parent authorization store in Authorization Manager. For more information about creating an application, see Create an application. Scopes, roles, tasks, and operations are always specific to an application. For more information, see Scopes in Authorization Manager and Roles, tasks, and operations.

**Using application groups**

An application group is a group of users of an Authorization Manager application. You can create application groups at any of the three levels in the Authorization Manager console. The following table lists the different Authorization Manager levels where you can create application groups.

| Level | Application group can be used in |
|---|---|
| Authorization store | The authorization store, and applications and scopes underneath it |
| Application | The application, and scopes underneath it |
| Scope | The scope |

For more information about application groups, see Groups in Authorization Manager.

**Delegating authorization stores and applications**

Authorization stores that are stored in Active Directory, and the applications that they contain, support delegation. For more information about performing delegation, see Delegate an authorization store and Delegate an application.

**Note**

- XML-based authorization stores, applications, and scopes do not support delegation.

# Scopes in Authorization Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scopes in Authorization Manager**

A scope is a subdivision within an application that separates resources from other resources that are used by that application. You do not have to use scopes. A scope can be a folder, an Active Directory container, a file-masked collection of files (for example, *.doc), a URL, or any item that can be accessed by the application and its underlying authorization store. You can use scopes to prevent unintended resource sharing and to support auditing and delegation.

If you have Authorization Manager groups, role assignments, role definitions, or task definitions that you do not want to apply to an entire application, you can create them at the scope level. The application that contains the scope must be able to recognize the scope name. For example, file-based applications might have scope names that include file names or paths. Web-based applications might have URL-based scope names. Registry applications might have scope names based on registry hives, and Active Directory scope names could specify organizational units. You cannot define operations at the scope level.

**Auditing scopes**

You cannot perform Authorization Manager runtime auditing on scopes. You can perform Authorization Manager authorization store change auditing on scopes contained in authorization stores that are stored in Active Directory. For more information about auditing, see Authorization Manager auditing.

**Delegating scopes**

Scopes within authorization stores that are stored in Active Directory support delegation. For more information about delegating a scope, see Delegate a scope.

XML-based authorization stores do not support delegated administration. If a scope within an authorization store that is stored in Active Directory contains task definitions that include authorization rules, or role definitions that include authorization rules, you cannot delegate the scope.

# Roles, tasks, and operations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Roles, tasks, and operations**

A role is a set of permissions that a user must have to do a job. Well-designed roles should correspond to a job category or responsibility (for example, receptionist, hiring manager, or archivist) and be named accordingly. With Authorization Manager, you can adds users to a role to authorize them for the job.

A task is a collection of operations, and sometimes other tasks. Well-designed tasks are inclusive enough to represent work items that are recognizable (for example, "change password" or "submit expense").

An operation is a set of permissions that you associate with system-level or API-level security procedures like WriteAttributes or ReadAttributes. You use operations as building blocks for tasks.

You can define roles, tasks, and operations only in developer mode, not administrator mode. To set developer mode, see Set Authorization Manager options.

**Role definitions**

The role definitions that are appropriate depends on the structure and goals of your organization. Roles support inheritance from other roles. To define a role, you specify a non-arbitrary name, a friendly description, and some lower-level tasks, roles, and operations that are part of it. This provides a mechanism for role inheritance. For example, a Helpdesk role might include a Product Support role. You can specify an authorization rule, which may be either VBScript or JScript. For more information, see VBScript at the Microsoft Web site and JScript at the Microsoft Web site.

If there are several authorization rules associated with a role definition (for example, it has several subroles and tasks), the authorization rules run synchronously. In Authorization Manager, the order has no effect on authorization.

**Role assignments**

A role assignment is a virtual container for application groups whose members are authorized for the role. A role assignment is based on a single role definition, and a single role definition can be the basis of many role assignments.

The most common procedure that administrators carry out is the assignment of application groups, or Windows users and groups, to a role. For step-by-step instructions, see Assign an application group to a role or Assign a Windows user or group to a role.

**Task definitions**

A task definition is smaller than a role definition and can be used to define roles and other tasks.

With Authorization Manager, you associate tasks with roles in an intuitive way. For example, the Recruiter role might include the Interview task. Tasks, like roles, are defined in a way that is appropriate to the organization. To define a task, you specify a name, a friendly description, and some lower-level tasks and operations that are part of it. You can also specify a VBScript or JScript authorization rule.

For more information about creating or editing task definitions, see Create a task definition and Edit a task definition.

**Operation definitions**

Operations are small computer-level actions that are used to define tasks and usually are not relevant to an administrator. You define operations only in developer mode.

You can set operation definitions at the application level, but not at the authorization store level or the scope level. An operation definition includes a name, a description, and an operation number. The operation number X must be an integer from zero to 2147483647 (that is, 0 ≤ X ≤ 231 - 1). The operation number is used by the application to identify the operation, so entering a wrong operation number will cause a bug in the application. For more information about creating or editing operation definitions, see Create an operation definition or Edit an operation definition.

# Authorization Manager auditing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authorization Manager auditing**
You need the **Generate security audits**privilege and the **Manage auditing and security log** privilege to do auditing. For more information about these privileges, see Privileges.

With Authorization Manager, you can use two kinds of auditing: runtime auditing and authorization store change auditing.

**Runtime auditing**
When you enable runtime auditing, applications generate audits when they use policy that is defined in the authorization store. You can configure auditing to log successes, failures, or both. Runtime auditing records client context and access checks. Access checks are based on the AccessCheck method described in the Authorization section of the Platform SDK. For more information about authorization-related application programming interfaces (APIs), see Authorization at the Microsoft Web site.

**Authorization store change auditing**
When you enable authorization store change auditing, audits are generated every time the authorization store is modified. The audit logs all events, successes and failures.

For authorization store change auditing, Authorization Manager supports the NTFS file system (for XML-based authorization stores) and Active Directory. The APIs upon which Authorization Manager is based are flexible enough to allow for the use of the Windows registry, Exchange Server, and SQL Server.

**Auditing availability**
The availability of auditing depends on the following:

- Whether the authorization store is based on Active Directory or XML.
- Whether auditing is applied at the authorization store level, the application level, or the scope level.

The following table describes the availability of the two auditing types.

| Level | Runtime auditing is available in | Authorization store change auditing is available in |
|---|---|---|
| Authorization store | • XML<br>• Active Directory | • XML<br>• Active Directory |
| Application | • XML<br>• Active Directory | • Active Directory |
| Scope | Not available | • Active Directory |

To use auditing, you have to select the appropriate check box on the **Auditing** tab. To enable runtime auditing, click the **Runtime application initialization auditing** check box. To enable authorization store change auditing, click the **Runtime client context and access check auditing** check box.

If the authorization store is based on XML, you also have to specify object access auditing. If the authorization store is based on Active Directory, you also have to specify directory service access auditing.

**Important**

- By default, object access auditing is turned off. Auditing for XML-based authorization stores only works if you turn on object access auditing. To do this, you need to use Group Policy at the domain, domain controller, or other applicable

organizational unit level in Active Directory, or through local security policy. If the XML-based authorization store is located on a domain controller, the **Default Domain Controllers Policy** Group Policy object (GPO) is the most appropriate place to turn on object access auditing. If the XML-based authorization store is located on a workstation or member server, edit the local Group Policy object for that computer to set local security policy, or edit another Group Policy object that applies to the computer through Active Directory. For more information about Active Directory, see Active Directory.

To enable object access auditing, double-click **Audit object access**.

**Where?**

- *Group Policy object name* (for example: Default Domain Controllers Policy)/Computer Configuration/Windows Settings/Security Settings/Local Policies/Audit Policy/Audit object access

Select the **Define these policy settings** check box, select the **Success** check box, and then select the **Failure** check box.

**Notes**

- If you are editing the local Group Policy object, the **Define these policy settings** check box does not appear in the Group Policy Object Editor. It only appears if you are editing Group Policy objects that are stored in Active Directory.

- If the **Success** and **Failure** auditing check boxes are unavailable, the **Define these policy settings** check box has probably been selected through security policy that is acting at a higher level in Active Directory. In this situation, you need to find out where the **Define these policy settings** check box is selected and clear the check box. To find this setting, look in the Group Policy objects (GPOs) that affect this computer.

**Important**

- By default, directory service access auditing is turned off. To use auditing for authorization stores that are stored in Active Directory, you must turn on directory service access auditing. You do this through Group Policy at the domain, domain controllers, or other applicable organizational unit level in Active Directory. The **Default Domain Controllers Policy** Group Policy object is usually the most appropriate place to turn on directory service-based access auditing. For more information about Group Policy, see Group Policy (pre-GPMC).

  To enable directory service access auditing, double-click **Audit directory service access** in Group Policy.

  **Where?**

  - *Group Policy object name* (for example: Default Domain Controllers Policy)/Computer Configuration/Windows Settings/Security Settings/Local Policies/Audit Policy/Audit directory service access

  **Notes**

  - If you are editing the local Group Policy object, the **Define these policy settings** check box does not appear in Group Policy. It only appears if you are editing Group Policy objects that are stored in Active Directory.

  - If the **Success** and **Failure** auditing check boxes are unavailable, the **Define these policy settings** check box has probably been selected through a security policy that is acting at a higher level in Active Directory. In this situation, you need to find out where the **Define these policy settings** check box is selected and clear the check box. To find this setting, look in the Group Policy objects (GPOs) that affect this computer.

After editing the Group Policy objects, run the **gpupdate** command to ensure that the changes take effect immediately. For more information, see Gpupdate.

**Auditing event logging**

All Authorization Manager audits are written to the **Security** section of the Event Viewer. For more information about Event

Viewer, see Event Viewer.

**Auditing that is enabled by inheritance**

Any auditing obtained through inheritance takes place regardless of the local setting. For example, in the case of an authorization store that is stored in Active Directory, auditing policy can be inherited from a parent organizational unit in Active Directory. In the case of an XML-based authorization store, audit policy on the folder containing the XML file is applicable.

For more information about auditing, see Auditing Security Events.

# Authorization Manager delegation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authorization Manager delegation**

Using delegation, you can assign responsibility to someone else. Delegation helps balance workloads, and it can help with security by assuring that individuals do not have more rights than they need.

Authorization stores that are stored in Active Directory support delegation.

XML-based authorization stores do not support delegation at any level. The XML file is secured as a whole by its NTFS file system access control entries (ACEs).

A scope within an authorization store that is stored in Active Directory cannot be delegated if the scope contains task definitions that include authorization rules, or role definitions that include authorization rules.

For more information about delegating authorization stores, see Delegate an authorization store. For more information about delegating applications, see Delegate an application. For more information about delegating a scope, see Delegate a scope.

For more information about delegation, see Delegating administration.

# Security Configuration Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Configuration Manager**
The Security Configuration Manager tool set allows you to create, apply and edit the security for your local computer, organizational unit, or domain.

- Before configuring security, see Checklist: Securing computers using Security Configuration Manager.
- For tips about using Security Configuration Manager, see Best practices for Security Configuration and Analysis, Best practices for Security Templates, and Best practices for Security Settings.
- For help with specific tasks, see Security Configuration Manager How To....
- For general background information, see Security Configuration Manager Concepts.
- For problem-solving instructions, see Security Troubleshooting.

| Security Configuration Manager tool | Description |
|---|---|
| Security Templates | Defines a security policy in a template. These templates can be applied to Group Policy or to your local computer. |
| Security Settings extension to Group Policy | Edits individual security settings on a domain, site, or organizational unit. |
| Local Security Policy | Edits individual security settings on your local computer. |
| Secedit | Automates security configuration tasks at a command prompt. |

# Checklist: Securing computers using Security Configuration Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Securing computers using Security Configuration Manager**

| **Preparing to set up security** | |
|---|---|
| **Step** | **Reference** |
| Review security best practices. | Best practices for security |
| Review the descriptions of individual settings. | Security Setting Descriptions |
| Review default security settings. | Default security settings for groups, Default groups and Differences in default security settings |
| Have Users only use applications belonging to the Windows Logo Program for Software. | Applications that belong to the Windows Logo Program for Software are guaranteed to run for members of the Users group. For other applications to run, you might need to elevate the privileges of your users to Power Users or relax the User permissions security settings. Both are less secure configurations for your organization. For more information, see the Windows Logo Program for Software page at the Microsoft Web site. |
| **Analyze your current settings** | |
| **Step** | **Reference** |
| Review **Security Templates.** | Security Templates |
| Select a predefined security template or create a new security template to use for your computers. | Predefined security templates or Define a security template |
| Perform a security analysis on the computer. | Analyze system security |
| **Modify security settings** | |
| **Step** | **Reference** |
| Apply the appropriate template. | Apply a security template to local policy or Import a security template to a Group Policy object |
| (Optional) Modify security settings on your local computer | Edit local security settings |

| | |
|---|---|
| (Optional) Modify security settings for a domain or organizational unit. | Edit security settings on a Group Policy object |

# Security Configuration Manager Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- Best practices for Security Configuration and Analysis
- Best practices for Security Templates
- Best practices for Security Settings

# Best practices for Security Configuration and Analysis

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for Security Configuration and Analysis**
Security Configuration and Analysis is a tool which can be used to analyze and configure computer security. For general security issues, see Best practices for security.

**If you require frequent analysis of a large number of computers, as with a domain-based infrastructure, use the Secedit.exe command line tool.**
For more information, see Automating security configuration tasks.

**Only use Security Configuration and Analysis to configure security areas not affected by Group Policy settings.**
This includes areas such as security on local files and folders, registry keys, and system services. Otherwise, Group Policy settings will override the local settings. For more information, see Group Policy.

**Do not use Security Configuration and Analysis when you are configuring security for a domain or an organizational unit.**
Otherwise, you must configure each client individually. In that case, you can:

- Use Security Templates to create a template and apply it to the appropriate Group Policy object. For more information, see Import a security template to a Group Policy object.

- Use the Security Settings extension to Group Policy to edit individual security settings on a Group Policy object. For more information, see Edit security settings on a Group Policy object.

# Best practices for Security Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for Security Templates**

**Use templates properly**

- Do not apply predefined or newly-created security templates to your computer or network without testing to ensure that the right level of application functionality is maintained.

- Never edit the Setup security.inf template, since it gives you the option to reapply the default security settings. If you ever remove a security template from a Group Policy object, appropriately reapply the Setup security.inf to restore all default settings.

- Do not apply the Setup security.inf template through Group Policy. The Setup security.inf template should only be applied to the local computer through Secedit or Security Configuration and Analysis. It is preferable to apply it in parts using the Secedit command-line tool.

  For more information, see Automating security configuration tasks.

- Do not apply the Compatible template to domain controllers. For example, do not import the Compatible template to the Default Domain or Default Domain Controller Group Policy object.

**Use caution when modifying predefined templates**

- Instead of modifying a predefined template, customize the predefined template and then save the changes under a different template name. Since these templates were designed for specific needs, having the original template will always give you the option of using it.

  For more information, see Predefined security templates.

**Choose the appropriate default level of computer access**

- When deciding on the default level of computer access that end users will have, the determining factor is the installed base of applications that need to be supported. If users only use applications that belong to the Windows Logo Program for Software, then you can make all your end users members of the Users group. If not, you may have to make your end users part of the Power users group so that they have the appropriate privileges to use the application, which is less secure.

  For more information about the Windows Logo Program for Software, see the Windows Logo Program for Software page at the Microsoft Web site.

For more information, see Best practices for Security Configuration and Analysis and Best practices for Security Settings.

# Best practices for Security Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Be cautious when creating new policies**

- Always test a newly-created policy on a test computer before applying it to your network.

**Be aware that more than one policy can be applied to a computer**

- Because of this, there can be conflicts in security policy settings. The order of precedence from highest precedence to lowest precedence is:

  1. Organizational unit
  2. Domain
  3. Local computer

  For more information, see Applying security settings.

- You can use Resultant Set of Policy to find out what policies apply to a certain computer. For more information , see Resultant Set of Policy.

**Keep in mind that there can only be one account policy in a domain: the Default Domain Policy**
For more information, see Account and local policies.

**Apply templates appropriately**

- Do not apply the Compatible template to Domain Controller computers. For example, do not import the Compatible template to Default Domain Policy or Default Domain Controller Policy.

- Do not apply the Setup security template through Group Policy.

**Use the correct tools for configuring local policy**

- For local security policy, use the Local Security Policy shortcut for editing and fine-tuning security policy. Use Security Templates to create a local policy and then use Security Configuration and Analysis to apply the policy.

# Security Configuration Manager How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Analyze and configure security
- Modify local security policy
- Modify security settings

# Analyze and configure security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Analyze and configure security**

- Add Security Configuration and Analysis to an MMC console
- Add Security Templates to an MMC console
- Analyze system security
- Configure local computer security
- Import a security template to a Group Policy object
- Import a security template
- Export a security template
- Edit the analysis database
- Customize a predefined security template
- Define a security template
- Set a new security template path

# Add Security Configuration and Analysis to an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add Security Configuration and Analysis to an MMC console**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. Do one of the following:
   - To create a new console, on the **File** menu, click **New**.
   - To open an existing console, on the **File** menu, click **Open**, click the console that you want to open, and then click **Open**.
3. On the **File** menu, click **Add/Remove Snap-in**, and then, in **Add/Remove Snap-in**, click **Add**.
4. Click **Security Configuration and Analysis,** click **Add**, click **Close**, and then click **OK**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- For more information, see Working with MMC console files.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Analyze system security
Configure local computer security

# Add Security Templates to an MMC console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add Security Templates to an MMC console**

1. Click **Start**, click **Run**, then type **mmc**, then click **OK**.
2. (Optional) To open an existing console, on the **File** menu, click **Open**, click the console you want to open, and then click **Open**.
3. On the **File** menu, click **Add/Remove Snap-in**.
4. In **Add/Remove Snap-in**, click **Add**.
5. In **Available Standalone Snap-ins**, click **Security Templates**, click **Add**, click **Close**, and then click **OK**.
6. On the **File** menu, click **Save**.

**Note**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Define a security template
Security Templates
Customize a predefined security template
Apply a security template to local policy
Import a security template to a Group Policy object
Working with MMC console files

# Analyze system security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To analyze system security**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Security Configuration and Analysis.
2. In the console tree, right-click **Security Configuration and Analysis**, and then click **Open Database**.

   **Where?**

   - *ConsoleRoot*/Security Configuration and Analysis

3. In **Open database**, do one of the following:

   - To create a new database, in **File name**, type a file name, and then click **Open**.
   - To open an existing database, click a database, and then click **Open**.

4. If you are creating a new database, in **Import Template**, click a template, and then click **Open**.
5. In the details pane, right-click **Security Configuration and Analysis**, and then click **Analyze Computer Now**.
6. Do one of the following:

   - To use the default log, in **Error log file path**, click **OK**.
   - To specify a different log, in **Error log file path**, type a valid path and file name, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Security Configuration and Analysis, click **Start**, click **Run**, type **mmc**, and then click **OK**. On the **File** menu, click **Open**, click the console that you want to open, and then click **Open**. Then, in the console tree, click **Security Configuration and Analysis**.
- Security Configuration and Analysis displays the different security areas as they are analyzed. Once this is complete, you can check the log file or review the results. For more information on reviewing the analysis results, see Related Topics.
- To view the log file, in the console tree, right-click **Security Configuration and Analysis**, and then click **View Log File**.

- The default path for the log file is:

  *systemroot*\Documents and Settings*UserAccount*\My Documents\Security\Logs*DatabaseName*.log

#### Using a command line

1. Open Command Prompt.

2. Type:

   **secedit /analyze /db** *FileName*.sdb [**/cfg** *FileName*] [**/overwrite**] [**/log** *FileName*] [**/quiet**]

ArgumentDescription

**/db** *FileName*Specifies the database used to perform the analysis.

**/cfg** *FileName*Specifies a security template to import into the database prior to performing the analysis. Security templates are created using the Security Templates snap-in.

**/log** *FileName*Specifies a file in which to log the status of the configuration process. If not specified, configuration data is logged in Scesrv.log, which is located in the %windir%\Security\Logs folder.

**/quiet**Specifies that the analysis process should take place without further comments.

Example: secedit /analyze /db hisecws.sdb

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **secedit /?**

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Security Configuration and Analysis to an MMC console
Analyzing security and viewing results
Import a security template
Configure local computer security
Working with MMC console files
Automating security configuration tasks

# Configure local computer security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure local computer security**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Security Configuration and Analysis.
2. In the console tree, right-click **Security Configuration and Analysis**, and then click **Open Database**.

   **Where?**

   - *ConsoleRoot*/Security Configuration and Analysis

3. In **Open database**, do one of the following:

   - To create a new database, in **File name**, type a file name, and then click **Open**.
   - To open an existing database, click a database, and then click **Open**.

4. If you are creating a new database, in **Import Template**, click a template, and then click **Open**.
5. In the console tree, right-click **Security Configuration and Analysis**, and then click **Configure Computer Now**.
6. Do one of the following:

   - To use the default log in **Error log file path**, click **OK**.
   - To specify a different log, in **Error log file path**, type a valid path and file name.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Security Configuration and Analysis, click **Start**, click **Run**, type **mmc**, and then click **OK**. On the **File** menu, click **Open**, click the console that you want to open, and then click **Open**. In the console tree, click **Security Configuration and Analysis**.
- To check the log file, right-click **Security Configuration and Analysis**, and then click **View Log File**.
- The default path for the log file is:

  *systemroot*\Documents and Settings*UserAccount*\My Documents\Security\Logs\

**Using a command line**

1. Open Command Prompt.

2. Type:

   **secedit/configure/DB***FileName* [**/CFG** *FileName*][**/overwrite**][**/areas** *Area1 Area2*...] [**/log** *LogPath*] [**/quiet**]

| Argument | Description |
|---|---|
| **/DB** *FileName* | Specifies the database used to perform the security configuration. |
| **/CFG** *FileName* | Specifies a security template to import into the database prior to configuring the computer. Security templates are created using the Security Templates snap-in. |
| **/overwrite** | Specifies that the database should be emptied prior to importing the security template. If this parameter is not specified, the settings in the security template are accumulated in the database. If this parameter is not specified and there are conflicting settings in the database and the template that is being imported, the template settings take precedence. |
| **/areas** *Area1 Area2*... | Specifies the security areas to be applied to the system. If this parameter is not specified, all security settings that are defined in the database are applied to the system. To configure multiple areas, separate each area by a space. The following security areas are supported:<br><br>SECURITYPOLICY - Includes account policies, audit policies, event log settings, and security options.<br><br>GROUP_MGMT - Includes Restricted Group settings.<br><br>USER_RIGHTS - Includes user rights assignment.<br><br>REGKEYS - Includes registry permissions.<br><br>FILESTORE - Includes file system permissions.<br><br>SERVICES - Includes system service settings. |
| **/log** *LogPath* | Specifies a file in which to log the status of the configuration process. If not specified, configuration data is logged in the scesrv.log file, which is located in the %windir%\Security\Logs folder. |
| **/quiet** | Specifies that the configuration process should take place without prompting the user. |

Examples:

- secedit /configure /db hisecws.sdb /cfg

- hisecws.inf /overwrite /log hisecws.log

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- The default path for the log file is:

  %windir%\Security\Logs\Scesrv.log

- **Secedit /refreshpolicy** has been replaced with **gpupdate**. For information on how to refresh security settings, see Gpupdate.

- To view the complete syntax for this command, at a command prompt, type:

  **secedit /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Automating security configuration tasks
Command-line reference A-Z

# Import a security template to a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import a security template to a Group Policy object**

1. Perform one of these steps

| If | Do this |
|---|---|
| You are on a workstation or server which is joined to a domain and you would like to import a security template for a Group Policy object. | • Click **Start**, point to **Run**, type **mmc** and click **OK**.<br>• On the **File** menu, click **Add/Remove snap-in**.<br>• In **Add/Remove Snap-in**, click **Add**, in **Add Standalone Snap-in**, double-click **Group Policy Object Editor**.<br>• In **Select Group Policy Object**, click **Browse**, select the policy object you would like to modify, click **OK**, and then click **Finish**.<br>• Click **Close** and then click **OK**. |
| You are on a domain controller and you would like to import a security template for a domain or organizational unit. | • Open Active Directory Users and Computers.<br>• In the console tree, right-click the domain or organizational unit you want to set Group Policy for.<br>• Click **Properties**, and then click the **Group Policy** tab.<br>• Click **Edit** to open the Group Policy object you want to edit or click **New** to create a new Group Policy object, and then click **Edit**. |

2. In the Group Policy console tree, right-click **Security Settings.**

   **Where?**

   - *Group Policy Object* Policy/Computer Configuration/Windows Settings/Security Settings

3. Click **Import Policy**, click the security template you want to import, and then click **Open**.

4. (Optional) If you want to clear the database of any previously stored security templates, select the **Clear this database before importing** check box.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using

Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- Do not import Setup security.inf to a Group Policy object.

- The security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Apply a security template to local policy
Working with MMC console files
Security Setting Descriptions
Security Templates

# Import a security template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import a security template**

1. Open Security Configuration and Analysis.
2. In the console tree, right-click **Security Configuration and Analysis**, and then click **Import Template**.

   **Where?**

   - *ConsoleRoot*/Security Configuration and Analysis

3. (Optional) To clear the database of any template that has been previously stored, select the **Clear this database before importing** check box.
4. Click a template file, and then click **Open**.
5. Repeat these steps for each template that you want to merge into the database.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Security Configuration and Analysis, click **Start**, click **Run**, type **mmc**, and then click **OK**. On the **File** menu, click **Open**, click the console that you want to open, and then click **Open**. In the console tree, click **Security Configuration and Analysis**. If Security Templates is not in the console, for instructions on installing it, see Related Topics.
- Importing templates to a personal database only affects the analysis database and does not affect your system settings.
- If you import more than one template, the database will merge the various templates to create one composite template, resolving conflicts in order of import. The last template that is imported takes precedence.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Security Configuration and Analysis to an MMC console
Importing and exporting security templates
Analyze system security
Configure local computer security
Export a security template
Working with MMC console files

# Export a security template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export a security template**

1. Open Security Configuration and Analysis.
2. In the console tree, right-click **Security Configuration and Analysis**, and then click **Export Template**.

   **Where?**

   - *ConsoleRoot*/Security Configuration and Analysis

3. Type a valid file name and a path to the location where your template will be saved.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Security Configuration and Analysis, click **Start**, click **Run**, type **mmc**, and then click **OK**. On the **File** menu, click **Open**, click the console that you want to open, and then click **Open**. In the console tree, click **Security Configuration and Analysis**. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Security Configuration and Analysis to an MMC console
Importing and exporting security templates
Analyze system security
Configure local computer security
Import a security template
Working with MMC console files

# Edit the analysis database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit the analysis database**

1. Open Security Configuration and Analysis.
2. In the console tree, right-click **Security Configuration and Analysis**, click **Analyze Computer Now**, and then click **OK**.

   **Where?**

   - *ConsoleRoot*/Security Configuration and Analysis

3. After Security Configuration and Analysis has finished analyzing the computer, in the console tree, click the group of settings that contains the security attribute you want to edit.
4. In the details pane, double-click the security attribute that you want to edit.
5. Confirm that the **Define this policy in the database** check box is selected.
6. Configure the security attribute setting as you want it to be recorded in the database, and then click **OK**.
7. Repeat the last four steps for each security attribute that you want to edit.
8. When you are finished modifying the database settings, in the console tree, right-click **Security Configuration and Analysis**, and then click **Save**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Security Configuration and Analysis, click **Start**, click **Run**, type **mmc**, and then click **OK**. On the **File** menu, click **Open**, click the console that you want to open, and then click **Open**. Then, in the console tree, click **Security Configuration and Analysis**.
- The changes that you make will not affect system security. They will only affect the analysis database.
- For more information, see Security Setting Descriptions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Security Configuration and Analysis to an MMC console
Analyze system security
Configure local computer security
Working with MMC console files

# Customize a predefined security template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To customize a predefined security template**

1. Open **Security Templates**.
2. In the console tree, double-click the default path folder (*systemroot*\Security\Templates), and in the details pane, right-click the predefined template you want to modify.
3. Click **Save As**, type a new file name for the security template, and click **Save**.
4. In the console tree, double-click the new security template to display the security policies, and navigate until the security attribute you want to modify appears in the details pane.
5. In the details pane, right-click the security attribute and click **Properties**.
6. Select the **Define this policy setting in the template** check box, make your changes, then click **OK**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Security Templates, click **Start**, and click **Run**, then type **mmc** and click **OK**. On the **File** menu, click **Open**, click the console that you want to open, then click **Open**. Then, in the console tree, click **Security Templates**.
- Only administrators have the permissions to make any changes to the default security template folder, which is:

  *systemroot*\Security\Templates

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Security Templates to an MMC console
Predefined security templates
Define a security template
Working with MMC console files
Security Setting Descriptions
Security Templates
Apply a security template to local policy
Import a security template to a Group Policy object
Security Templates

# Define a security template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define a security template**

1. Open **Security Templates**.
2. Right-click the folder where you want to store the new template and click **New Template**.
3. In **Template name**, type the name for your new security template.
4. In **Description**, type a description of your new security template, and then click **OK**.
5. In the console tree, double-click the new security template to display the security areas and navigate until the security setting you want to configure is in the details pane.
6. In the details pane, right-click the security setting you want to configure and click **Properties**.
7. Select the **Define this policy setting in the template** check box, edit the settings, and then click **OK**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Security Templates, click **Start**, click **Run**, type **mmc**, and then click **OK**. On the **File** menu, click **Open**, click the console that you want to open, then click **Open**. In the console tree, click **Security Templates**.
- Only administrators have the permissions to make any changes to the default security template folder, which is:

  *systemroot*\Security\Templates

  If you would like to create a new folder for new security templates, see Set a new security template path.
- For more information on the different policy settings, see Security Setting Descriptions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Security Templates to an MMC console
Security Templates
Customize a predefined security template
Set a new security template path
Working with MMC console files
Apply a security template to local policy
Import a security template to a Group Policy object
Security Templates

# Set a new security template path

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set a new security template path**

1. Open **Security Templates**.
2. In the console tree, right-click **Security Templates**, click **New Template Search Path**, and select the new location.

   A folder with the path of the new location appears in the console tree.
3. (Optional) To add a description to the new template path (in addition to the name), in the console tree, right-click the new template folder, click **Set Description**, and then type the description you want to add.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Security Templates, click **Start**, and click **Run**, then type **mmc** and click **OK**. On the **File** menu, click **Open**, click the MMC console that you want to open, then click **Open**. Then, in the console tree, click **Security Templates**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Security Templates to an MMC console
Customize a predefined security template
Security Templates
Define a security template
Working with MMC console files

# Modify local security policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modify local security policy**

- Edit local security settings
- Apply a security template to local policy
- Assign user rights for your local computer
- Reapply default security settings

# Edit local security settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit local security settings**

1. Open Local Security Settings.
2. Do one of the following:
   - To edit **Password Policy** or **Account Lockout Policy**, in the console tree, click **Account Policies**.
   - To edit an **Audit Policy**, **User Right Assignment**, or **Security Options**, in the console tree, click **Local Policies**.
3. In the console tree, click the folder that contains the policy you want to modify, and then, in the details pane, double-click the policy that you want to modify.
4. Make the changes you want and click **OK**.
5. To change other policies, repeat the three previous steps.

## Notes

- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.
- If you would like to edit the security settings for file system, registry, and system services on the local computer, you must first edit the changes in a security template and then configure the computer using Security Configuration and Analysis.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Local Security Policy
Account and local policies
Checklist: Securing computers using Security Configuration Manager
Security Setting Descriptions
Assign user rights for your local computer

# Apply a security template to local policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To apply a security template to local policy**

1. Open **Security Configuration and Analysis**.
2. In the console tree, right-click **Security Configuration and Analysis**, and then click **Open Database**.

   **Where?**

   - *ConsoleRoot*/**Security Configuration and Analysis**

3. In **File name**, type a file name and then click **Open**.
4. (Optional) If you are creating a new policy, in **Import Template**, do one of the following:

   - To clear the database of any previously imported templates, select the **Clear this database before importing** check box.
   - To append this template to any previously imported templates, confirm that the **Clear this database before importing** check box is cleared. (The last template that was imported has precedence over any conflicting values.)

5. In **Import Template**, click a template, and then click **Open**.
6. In the console tree, right-click **Security Configuration and Analysis**, and then click **Configure Computer Now**.
7. Do one of the following:

   - To use the default error log, in **Error log file path**, click **OK**.
   - To specify a different error log, in **Error log file path**, type a valid path and file name.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Security Configuration and Analysis, click **Start**, and click **Run**, then type **mmc** and click **OK**. On the **File** menu, click **Open**, click the console that you want to open, and then click **Open**. Then, in the console tree, click **Security Configuration and Analysis**.
- To check the log file, right click **Security Configuration and Analysis**, and then click **View Log File**.
- The default path for the log file is:

  *systemroot*\Documents and Settings*UserAccount*\My Documents\Security\Logs\

- The security settings on the local computer take effect immediately.

- To import more than one template to your computer policy, see Import a security template.

- If you import more than one template, the database will merge the various templates to create one composite template, resolving conflicts in order of import. The last template that is imported takes precedence over any conflicting settings.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add Security Configuration and Analysis to an MMC console
Import a security template to a Group Policy object
Working with MMC console files
Configure local computer security
Security Templates

# Assign user rights for your local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign user rights for your local computer**

1. Open Local Security Settings.
2. In the console tree, click **User Rights Assignment**.

   **Where?**

   - Security Settings/Local Policies/User Rights Assignments

3. In the details pane, double-click the user right you want to change.
4. In ***UserRight* Properties**, click **Add User or Group**.
5. Add the user or group and click **OK**.

### Note

- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Account and local policies
User rights
Logon rights
About user rights
Privileges
Security Setting Descriptions
Local Security Policy

# Reapply default security settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reapply default security settings**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Security Configuration and Analysis.
2. In the console tree, right-click **Security Configuration and Analysis**, and then click **Open Database**.

   **Where?**

   - *ConsoleRoot*/Security Configuration and Analysis

3. In **File name**, type the file name, and then click **Open**.
4. Do one of the following:

   - For a domain controller, in the console tree, right-click **Security Configuration and Analysis**, click **Import Template**, and then click **DC security**.
   - For other computers, in the console tree, right-click **Security Configuration and Analysis**, click **Import Template**, and then click **setup security**.

5. Select the **Clear this database before importing** check box, and then click **Open**.
6. In the console tree, right-click **Security Configuration and Analysis**, and then click **Configure Computer Now**.
7. Do one of the following:

   - To use the default log specified in **Error log file path**, click **OK**.
   - To specify a different log, in **Error log file path**, type a valid path and file name, and then click **OK**.

8. When the configuration is done, right-click **Security Configuration and Analysis**, and then click **View Log File**.

**Important**

- Applying the entire setup security template is a drastic measure that should be avoided. Instead, use the **secedit** command-line tool to apply default settings for specific areas. See the **Using a command line** section of this procedure.

**Notes**

- Different permissions are required to perform this procedure, depending on the environment in which you reapply default security settings:

- If you reapply default security settings to your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- If you reapply default security settings to a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Security Configuration and Analysis, click **Start**, click **Run**, type **mmc**, and then click **OK**. On the **File** menu, click **Open**, click the console that you want to open, and then click **Open**. In the console tree, click **Security Configuration and Analysis**.

- The default path for the log file is:

  *systemroot*\Documents and Settings*UserAccount*\My Documents\Security\Logs\

- When you reapply default security settings, all settings that are defined in Setup security.inf are set as the template specifies, but other settings that are not defined in the template may persist. For more information, see Applying security settings.

**Using a command line**

- Open Command Prompt.

- For a server or workstation, type:

  **secedit/configure/DB***FileName***/CFG "%windir%\Security\Templates\Setup security.inf"** [**/overwrite**][**/areas** *Area1 Area2*...] [**/log** *LogPath*] [**/quiet**]

  For a domain controller, type:

  **secedit/configure/DB***FileName***/CFG "%windir%\Security\Templates\DC security.inf"** [**/overwrite**][**/areas** *Area1 Area2*...] [**/log** *LogPath*] [**/quiet**]

| Argument | Description |
|---|---|
| **/DB** *FileName* | Required. Provides the path to a database that contains the security template that should be applied. To create a new database, type the database file name and path. |
| **/CFG "%windir%\Security\Templates\Setup security.inf"** | Specifies the Setup Security.inf template that contains the default security settings. |
| **/overwrite** | Specifies that the database should be emptied prior to importing the security template. If this parameter is not specified, the settings in the security template are accumulated into the database. If this parameter is not specified and there are conflicting settings in the database and the template being imported, the template settings win. |

| | |
|---|---|
| **/areas** *Area1 Area2* | Specifies the security areas to be applied to the system. If this parameter is not specified, all security settings defined in the database are applied to the system. To configure multiple areas, separate each area by a space. The following security areas are supported:<br><br>SECURITYPOLICY - Includes account policies, audit policies, event log settings, and security options.<br><br>GROUP_MGMT - Includes Restricted Group settings.<br><br>USER_RIGHTS - Includes user rights assignment.<br><br>REGKEYS - Includes registry permissions.<br><br>FILESTORE - Includes file system permissions.<br><br>SERVICES - Includes System Service settings. |
| **/log** *LogPath* | Specifies a file in which to log the status of the configuration process. If not specified, configuration data is logged in the Scesrv.log file, which is located in the %windir%\Security\Logs folder. |
| **/quiet** | Specifies that the configuration process should take place without prompting the user. |

**Important**

- It is advisable to apply **Setup security** in parts using the **Areas** parameter, so you can have control over which parts you are restoring.

**Notes**

- Different permissions are required to perform this procedure, depending on the environment in which you reapply default security settings:

    - If you reapply default security settings to your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

    - If you reapply default security settings to a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **secedit /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Automating security configuration tasks
Command-line reference A-Z

# Modify security settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modify security settings**

- Edit security settings on a Group Policy object
- Permit users to log on locally to a domain controller

# Edit security settings on a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit a security setting on a Group Policy object**

Choose the appropriate environment for which you want to edit a security setting:

- For your local computer
- For a Group Policy object, when you are on a workstation or server that is joined to a domain.
- For a Group Policy object, when you are on a domain controller or on a workstation that has the Windows Server 2003Administration Tools Pack installed.
- For only domain controllers, when you are on a domain controller.

**For your local computer**

1. Open Local Security Settings.
2. In the console tree, click **Security Settings**.
3. Do one of the following:
   - To edit **Password Policy** or **Account Lockout Policy**, click **Account Policies**.
   - To edit an **Audit Policy**, a **User Right Assignment**, or **Security Options**, click **Local Policies**.
4. Double-click the security setting in the details pane that you want to modify.
5. Modify the security setting, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**For a Group Policy object, when you are on a workstation or server that is joined to a domain.**

1. On the taskbar, click **Start**, point to **Run**, type **mmc**, and then click **OK**.
2. In the console, on the **File** menu, click **Add/Remove snap-in**.
3. In **Add/Remove Snap-in**, click **Add**, and then, in **Add Standalone Snap-in**, double-click **Group Policy Object Editor**.
4. In **Select Group Policy Object**, click **Browse**, browse to the policy object you would like to modify, and then click **Finish**.
5. Click **Close**, and then click **OK**.

6. In the console tree, click **Security Settings**.

   **Where?**

   - GroupPolicyObject [ComputerName] Policy\Computer Configuration\Windows Settings\Security Settings

7. Do one of the following:

   - To edit **Password Policy**, **Account Lockout Policy**, or **Kerberos Policy**, in the details pane, double-click **Account Policies**.
   - To edit **Audit Policy**, **User Rights Assignment**, or **Security Options**, in the details pane, double-click **Local Policies**.
   - To edit event log settings, on the console tree, click **Event Log**.

8. In the details pane, double-click the security setting that you want to modify.
9. (Optional) If this security setting has not yet been defined, select the **Define these policy settings** check box.
10. Modify the security setting and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

**For a Group Policy object, when you are on a domain controller or on a workstation that has the Windows Server 2003Administration Tools Pack installed.**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the Group Policy object for which you want to edit security settings.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Do one of the following:

   - To edit an existing Group Policy object, click **Edit**.
   - To create a new Group Policy object, click **New**, and then click **Edit**.

5. In the console tree, click **Security Settings**.

   **Where?**

   - GroupPolicyObject [ComputerName] Policy\Computer Configuration\Windows Settings\Security Settings

6. Do one of the following:

- To edit **Password Policy**, **Account Lockout Policy**, or **Kerberos Policy**, click **Account Policies**.
- To edit **Audit Policy**, **User Rights Assignment**, or **Security Options**, click **Local Policies**.
- To edit Event log settings, click **Event Log**.

7. Double-click the security setting in the details pane that you want to modify.
8. (Optional) If this security setting has not yet been defined, select the **Define these policy settings** check box.
9. Modify the security setting and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**For only domain controllers, when you are on a domain controller.**

1. Open Domain Controller Security Policy.
2. In the console tree, click **Security Settings**.

   **Where?**

   - GroupPolicyObject [ComputerName] Policy\Computer Configuration\Windows Settings\Security Settings

3. Do one of the following:

   - To edit **Password Policy**, **Account Lockout Policy**, or **Kerberos Policy**, in the console tree, double-click **Account Policies**.
   - To edit **Audit Policy**, **User Rights Assignment**, or **Security Options**, in the console tree, click **Local Policies**.
   - To edit event log settings, in the console tree, click **Event Log**.

4. In the details pane, double-click the security setting that you want to modify.
5. (Optional) If this security setting has not yet been defined, select the **Define these policy settings** check box.
6. Modify the security setting, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Domain Controller Security Policy, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and

then double-click **Domain Controller Security Policy**.

### Notes

- Always test a newly-created policy on a test organizational unit before applying it to your network.
- When you change a security setting and click **OK**, that setting will take effect in the next refresh of settings.
- The security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Applying security settings
Best practices for Security Settings

# Permit users to log on locally to a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To permit users to log on locally to a domain controller**

1. Open Domain Controller Security Policy.
2. In the console tree, click **User Rights Assignment.**

   **Where?**

   - *GroupPolicyObjectName* [*DomainControllerName*] Policy/Computer Configuration/Windows Settings/Security Settings/Local Policies/User Rights Assignment

3. In the details pane, double-click **Allow log on locally**.
4. If this security setting has not yet been defined, select the **Define these policy settings** check box, and click **Add User or Group**.
5. In **Add user or group**, specify the user or group who will be granted permission to log on locally, and then click **OK** twice.

## Notes

- To open Domain Controller Security Policy, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Domain Controller Security Policy**.
- You can also explicitly deny users or groups the permission to log on locally by configuring the "Deny log on locally" user right. For more information, see Related Topics.
- Setting changes are applied every five minutes on a domain controller. Every 16 hours, there is a forced refresh on the settings, regardless of any changes.
- User Rights are defined by default in the Default Domain Controller Group Policy object, which is associated with the Domain Controllers organizational unit. As a result, all domain controllers have the same User Rights policy.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Logon rights
Deny log on locally
Allow log on locally
User Rights Assignment

# Security Configuration Manager Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Security Configuration Manager:

- Security Configuration Manager overview
- Understanding Security Configuration Manager
- Using Security Configuration Manager

# Security Configuration Manager overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Configuration Manager overview**

- Security Configuration and Analysis overview
- Security Templates overview
- Local Security Policy overview
- Security settings overview

# Security Configuration and Analysis overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Configuration and Analysis overview**
Security Configuration and Analysis is a tool for analyzing and configuring local system security.

**Security analysis**
The state of the operating system and applications on a computer is dynamic. For example, you may need to temporarily change security levels so that you can immediately resolve an administration or network issue. However, this change can often go unreversed. This means that a computer may no longer meet the requirements for enterprise security.

Regular analysis enables an administrator to track and ensure an adequate level of security on each computer as part of an enterprise risk management program. An administrator can tune the security levels and, most importantly, detect any security flaws that may occur in the system over time.

Security Configuration and Analysis enables you to quickly review security analysis results. It presents recommendations alongside of current system settings and uses visual flags or remarks to highlight any areas where the current settings do not match the proposed level of security. Security Configuration and Analysis also offers the ability to resolve any discrepancies that analysis reveals.

**Security configuration**
Security Configuration and Analysis can also be used to directly configure local system security. Through its use of personal databases, you can import security templates that have been created with Security Templates and apply these templates to the local computer. This immediately configures the system security with the levels specified in the template.

For more information, see Security Configuration and Analysis.

# Security Templates overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Templates overview**
With the Security Templates snap-in for Microsoft Management Console, you can create a security policy for your computer or for your network. It is a single point of entry where the full range of system security can be taken into account. The Security Templates snap-in does not introduce new security parameters, it simply organizes all existing security attributes into one place to ease security administration.

Importing a security template to a Group Policy object eases domain administration by configuring security for a domain or organizational unit at once. For more information, see Import a security template to a Group Policy object.

In order to apply a security template to your local computer, you can use Security configuration and analysis or the Secedit command-line tool. For more information, see Configure local computer security.

Security templates can be used to define:

- Account Policies
    - Password Policy
    - Account Lockout Policy
    - Kerberos Policy
- Local Policies
    - Audit Policy
    - User Rights Assignment
    - Security Options
- Event log: Application, system, and security Event Log settings
- Restricted Groups: Membership of security-sensitive groups
- System Services security settings: Startup and permissions for system services
- Registry security settings: Permissions for registry keys
- File System security settings: Permissions for folders and files

Each template is saved as a text-based .inf file. This enables you to copy, paste, import, or export some or all of the template attributes. With the exceptions of Internet Protocol security and public key policies, all security attributes can be contained in a security template.

**New and predefined templates**
With each member of the Microsoft® Windows Server 2003 family or with the Microsoft® Windows® XP operating system, there are a set of predefined templates created for different levels of security to suit your organization.

There are several predefined templates that can help you to secure your system based on your needs, these templates are for:

- Reapplying default settings.
- Implementing a highly secure environment.

- Implementing a less secure but more compatible environment.
- Securing the systemroot.

For more information, see Predefined security templates.

The **Setup security.inf** template allows you to reapply default security settings. This template is created during setup for each computer and must be applied locally.

You can create a new security template with your own preferences or use one of the predefined security templates. Before making any changes to your security settings you should understand what the default settings of your system are and what they mean.

For more information about Security Templates, see Security Templates.

# Local Security Policy overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local security policy overview**
A security policy is a combination of security settings that affect the security on a computer. You can use your local security policy to edit account policies and local policies on your local computer.

With the local security policy, you can control:

- Who accesses your computer.
- What resources users are authorized to use on your computer.
- Whether or not a user or group's actions are recorded in the event log.

**How policy is applied to a computer that is joined to a domain**
If your local computer is joined to a domain, you are subject to obtaining security policy from the domain's policy or from the policy of any organizational unit that you are a member of. If you are getting policy from more than one source, then any conflicts are resolved in this order of precedence, from highest to lowest:

- Organizational unit policy
- Domain policy
- Site policy
- Local computer policy

When you modify the security settings on your local computer using the local security policy, then you are directly modifying the settings on your computer. Therefore, the settings take effect immediately, but this may only be temporary. The settings will actually remain in effect on your local computer until the next refresh of Group Policy security settings, when the security settings that are received from Group Policy will override your local settings wherever there are conflicts. The security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes.

For more information, see Local Security Policy.

# Security settings overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security Settings overview**

Organizational units, domains, and sites are linked to Group Policy objects. The Security Settings tool allows you change the security configuration of the Group Policy object, in turn, affecting multiple computers. With Security Settings, you can modify the security settings of many computers, depending on the Group Policy object you modify, from just one computer joined to a domain.

Security settings or security policies are rules that are configured on a computer or multiple computers for protecting resources on a computer or network. Security settings can control:

- How users are authenticated to a network or computer.
- What resources users are authorized to use.
- Whether or not a user's or group's actions are recorded in the event log.
- Group membership.

For more information, see Security Setting Descriptions.

**Using Security Settings**

You can change the security configuration on multiple computers in two ways:

- Create a security policy using a security template with Security Templates and then import it through Security Settings to a Group Policy object. For more information, see Import a security template to a Group Policy object.
- Change a few select settings with Security Settings. For more information, see Edit security settings on a Group Policy object.

# Understanding Security Configuration Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Security Configuration Manager**
This section covers:

- Account and local policies
- Event log security settings
- Security settings for file system, registry, and system services
- Local computer security for file system, registry, system services

# Account and local policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Account and local policies**
All security policies are computer-based policies. This section explains account policies and local policies.

**Account policies**
Account policies are defined on computers, yet they affect how user accounts can interact with the computer or domain. Account policies contain three subsets:

- Password Policy. Used for domain or local user accounts. Determines settings for passwords, such as enforcement and lifetimes.
- Account Lockout Policy. Used for domain or local user accounts. Determines the circumstances and length of time that an account will be locked out of the system.
- Kerberos Policy. Used for domain user accounts. Determines Kerberos-related settings, such as ticket lifetimes and enforcement. Kerberos policies do not exist in Local Computer Policy.

For domain accounts, the account policy must be defined in the Default Domain Policy Group Policy object (GPO) or in a new GPO that is linked to the root of the domain and given precedence over the Default Domain Policy GPO, which is enforced by the domain controllers that make up the domain. If more than one GPO containing account policy settings is linked at the domain level, the domain's account policy consists of the cumulative policy settings from all the domain-linked GPOs.

A domain controller always obtains the account policy from a GPO linked to the domain, which by default is the Default Domain Policy GPO. This behavior occurs even if a different account policy is applied to the organizational unit (OU) that contains the domain controller. By default, workstations and servers joined to a domain (such as member computers) also receive the same account policy for their local accounts. However, local account policies for member computers can be differentiated from the domain account policy by defining an account policy for the OU that contains the member computers.

**Important**

- Modifying the Default Domain Policy is not recommended. If you need to set some account policy that differs from that in the Default Domain Policy GPO, you can create a new GPO and link to the root of domain, set the policy you want to use, and assign it higher precedence than Default Domain Policy GPO.

There are two policies in Security Options that also behave like account policies. These are:

- Network Access: Allow anonymous SID/NAME translation
- Network Security: Force Logoff when Logon Hours expire

**Local policies**
These policies apply to a computer and contain these subsets:

- Auditing Policy. Determines whether security events are logged into the Security log on the computer. Also determines whether to log successful attempts, failed attempts or both. (The Security log is part of Event Viewer.)
- User Rights Assignment. Determines which users or groups have logon rights or privileges on the computer.
- Security Options. Enables or disables security settings for the computer, such as digital signing of data, Administrator and Guest account names, floppy drive and CD-ROM access, driver installation, and logon prompts.

Because a computer can have more than one policy applied to it, there can be conflicts in security policy settings. The order of precedence from highest precedence to lowest precedence is organizational unit, domain, and local computer. For more

information, see Applying security settings.

# Event log security settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Event log security settings**

This security area defines attributes that are related to the application, security, and system logs: maximum log size, access rights for each log, and retention settings and methods. The application log records events generated by programs; the security log records security events, including logon attempts, object access, and changes to security, depending on what is audited; and the system log records operating system events. For more information, see Event Viewer.

For more information on event log policies, see Settings for Event Logs.

# Security settings for file system, registry, and system services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security settings for file system, registry, and system services**

**File system and registry keys**
You can configure security attributes using Security Configuration Manager for all existing files, folders, and registry keys in the local file system:

- **Propagate inheritable permissions to all subfolders and files (or subkeys)**: Any child objects of this object inherit the parental object's security. provided the child object is not protected from inheritance. If there is a conflict, the explicit permissions on the child object will override the permissions inherited from the parent object.

- **Replace existing permissions on all subfolders and files (or subkeys) with inheritable permissions**: The parent object's permission entries will override any security settings on the child object, regardless of the child object's settings. The child object will have identical Access Control Entries as the parent object.

  **Caution**

  - If you select this option and then click **Apply**, you will override all child objects' security settings. You will not be able to reset the child objects' security settings by clearing the check box.

- **Do not allow permissions on this file or folder (or key) to be replaced**: Use this setting if you do not want to configure or analyze security for this object or any of its child objects.

**System services**
You can select this if the startup mode is:

- **Automatic**: Automatically starts when computer is restarted.

- **Manual**: Does not start unless someone starts it.

- **Disabled**: The service cannot be started.

### Notes

- You can also edit the security properties of the object: what user or group accounts have permission to read, write, delete, or execute inheritance settings or auditing and ownership permission by clicking **Edit Security**

- If you choose to set system service startup to **Automatic**, perform adequate testing to verify that the services can start without user intervention.

- You should track the system services used on a computer. For performance optimization, set unnecessary or unused services to Manual.

For more information on file and folder permissions, see Permissions for files and folders.

# Local computer security for file system, registry, system services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local computer security for file system, registry, system services**
In order to edit security settings on the file system, registry, or system services for the local computer, you must:

1. Create or edit a security template. For more information, see Define a security template.

2. Apply it to the local computer using Security Configuration and Analysis. For more information, see Configure local computer security

For more information, see Security settings for file system, registry, and system services.

# Using Security Configuration Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Security Configuration Manager**

This section covers:

- Applying security settings
- Predefined security templates
- Importing and exporting security templates
- Analyzing security and viewing results
- Resolving security discrepancies
- Automating security configuration tasks
- Security Setting Descriptions

# Applying security settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Applying security settings through Group Policy**

**When policy is applied**
Once you have edited the security settings, the settings are refreshed on the computers in the organizational unit linked to your Group Policy object:

- When a computer is restarted the settings on that computer will be refreshed.
- The security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes.
- To force a computer to refresh its security settings as well as all Group Policy settings, see the Gpupdate command line tool.

**Precedence of policy when more than one policy is applied to a computer**
For security settings which are defined by more than one policy, the following order of precedence, from highest to lowest, is observed:

- Organizational Unit Policy
- Domain Policy
- Site Policy
- Local computer Policy

For example, a workstation that is joined to a domain will have its local security settings overridden by the domain policy wherever there is a conflict. Likewise, if the same workstation is a member of an Organizational Unit, the settings applied from the Organizational Unit's policy will override both the domain and local settings. If the workstation is a member of more than one Organizational Unit, then the Organizational Unit that immediately contains the workstation has the highest order of precedence.

**Notes**

- Use Resultant Set of Policies tool to find out what policies are applied and in what order to a computer. For more information , see Resultant Set of Policy.
- For domain accounts, there can be only one account policy which includes password policies, account lockout policies, and kerberos policies. For more information, see Account and local policies.

**Persistence in security settings**
Security settings may still persist even if a setting is no longer defined in the policy that originally applied it.

Persistence in security settings occurs when:

- The setting has not been previously defined for the computer.
- The setting is for a registry object.
- The settings is for a file system object.

All settings applied through local policy or a Group Policy object are stored in a local database on your computer. Whenever a security settings is modified, the computer saves the security setting value to the local database which retains a history of all the settings that have been applied to the computer. If a policy first defines a security setting and then no longer defines that setting, then the setting takes on the previous value in the database. If a previous value does not exist in the database then the setting does not revert to anything and remains defined as is. This behavior is sometimes called "tattooing."

Registry and file settings will maintain the values applied through policy until that setting is set to other values.

**Filtering security settings based on group membership**

You can also decide what users or groups will or will not have a Group Policy object applied to them regardless of what computer they have logged onto by denying them either the **Apply Group Policy** or **Read** permission on that Group Policy object. Both of these permissions are needed to apply Group Policy. For more information, see Group Policy (pre-GPMC).

# Predefined security templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Predefined security templates**

The predefined security templates are provided as a starting point for creating security policies that are customized to meet different organizational requirements. You can customize the templates with the Security Templates snap-in. Once you customize the predefined security templates, you can use them to configure security on an individual computer or thousands of computers. You can configure individual computers with the Security Configuration and Analysis snap-in, the Secedit command-line tool, or by importing the template into Local Security Policy. You can configure multiple machines by importing a template into Security Settings extension to Group Policy, which is an extension to Group Policy. You can also use a security template as a baseline for analyzing a system for potential security holes or policy violations by using the Security Configuration and Analysis snap-in. By default, the predefined security templates are stored in:

*systemroot*\Security\Templates

- **Default security (Setup security.inf)**

  The Setup security.inf template is created during installation for each computer. It can vary from computer to computer, based on whether the installation was a clean installation or an upgrade. Setup security.inf represents the default security settings that are applied during installation of the operating system, including the file permissions for the root of the system drive. It can be used on servers and client computers; it cannot be applied to domain controllers. You can apply portions of this template for disaster recovery purposes.

  Setup security.inf should never be applied using Group Policy. It contains a large amount of data and can seriously degrade performance if it is applied through Group Policy, because policy is periodically refreshed and a large amount of data would move through the domain.

  It is advisable to apply the Setup security.inf template in parts. Since the Secedit command-line tool gives you this option, it is advisable to use it.

  For more information, see Automating security configuration tasks.

- **Domain controller default security (DC security.inf)**

  This template is created when a server is promoted to a domain controller. It reflects file, registry, and system service default security settings. Reapplying it resets these areas to the default values, but it may overwrite permissions on new files, registry keys and system services created by other applications. It can be applied using the Security Configuration and Analysis snap-in or the Secedit command-line tool.

- **Compatible (Compatws.inf)**

  Default permissions for workstations and servers are primarily granted to three local groups: Administrators, Power Users, and Users. Administrators have the most privileges while Users have the least. Because of this, you can significantly improve the security, reliability, and total cost of system ownership by:

  - Making sure that end users are members of the Users group.
  - Deploying applications that can be run successfully by members of the Users group.

  People with User privileges can successfully run applications that take part in the Windows Logo program for Software. However, Users may not be able to run applications that do not meet the requirements of the program. If other applications must be supported, there are two options:

  - Allow members of the Users group to be members of the Power Users group.

- Relax the default permissions that are granted to the Users group.

Since Power Users have inherent capabilities, such as creating users, groups, printers, and shares, some administrators would rather relax the default User permissions than allow end users to be members of the Power Users group. This is precisely what the Compatible template is for. The Compatible template changes the default file and registry permissions that are granted to Users in a manner that is consistent with the requirements of most applications that do not belong to the Windows Logo program for Software. Additionally, since it is assumed that the administrator that is applying the Compatible template does not want end users to be Power Users, the Compatible template also removes all members of the Power Users group. For more information, see Default security settings for groups.

For more information on the Windows Logo program for Software, see the Microsoft Web site.

The Compatible template should not be applied to domain controllers. For example, do not import the Compatible template to the Default Domain policy or Default Domain Controller policy.

- **Secure (Secure*.inf)**

  The Secure templates define enhanced security settings that are least likely to impact application compatibility. For example, the Secure templates define stronger password, lockout, and audit settings.

  Additionally, the Secure templates limit the use of LAN Manager and NTLM authentication protocols by configuring clients to send only NTLMv2 responses and configuring servers to refuse LAN Manager responses.

  - In order to apply Securews.inf to a member computer, all of the domain controllers that contain the accounts of all users that log on to the client must run Windows NT 4.0Service Pack 4 or higher.

  - In order to apply Securews.inf to a member computer that is joined to a domain which contains domain controllers running Windows NT 4.0, the clocks between the domain controllers running Windows NT 4.0and the member computers must be within 30 minutes of each other.

  - If a client is configured with Securews.inf, it will not be able to connect to servers that only use the LAN Manager authentication protocol or that run Windows NT 4.0prior to Service Pack 4 using a local account defined on the target server.

  - If a client is configured with Securews.inf, it will not be able to connect to servers running Windows 2000or Windows NT 4.0using a local account defined on the target server unless the clock on the target server is within 30 minutes of the clock on the client.

  - If a client is configured with Securews.inf, it will not be able to connect to a computer running Windows XPor above using a local account defined on the target server unless the clock on the target server is within 20 hours of the clock on the client.

  - If a client is configured with Securews.inf, then it will not be able to connect to servers running LAN Manager that are running in share-level security mode.

  - If a server is configured with Securews.inf, then a user with a local account on that server will not be able to connect to it from a client computer running only LAN Manager that is using that local account.

  - If a server running Windows 2000is configured with securews.inf, then a client with a local account on that server that is also configured to use NTLMv2 authentication will not be able to connect unless the clocks on the two machines are within 30 minutes of each other.

  - If a server running Windows XPis configured with securews.inf, then a client with a local account on that server that is also configured to use NTLMv2 authentication will not be able to connect unless the clocks on the two machines are within 20 hours of each other.

  - If a domain controller is configured with Securedc.inf, then a user with an account in that domain will not be able

to connect to any member server from a client computer running only LAN Manager using that domain account.

- Computers that run LAN Manager include Windows for Workgroupsas well as Windows 95and Windows 98platforms that do not have the Active Directory Client Extensions Pack installed. If the Active Directory Client Extensions Pack is installed on Windows 95or Windows 98, then those clients will be able to use NTLMv2. Windows Millinnium Editionsupports NTLMv2 without additional modification.

- Computers that run Windows NTService Pack 4 or higher can be configured to send only NTLMv2 responses by setting **HKLM\System\CurrentControlSet\Control\LSA\LMCompatibilityLevel** to equal 3 or higher.

- Computers that run Windows NTService Pack 4 or higher can be configured to send only NTLMv2 responses by specifying this preference in the Network security: LAN Manager authentication level security option.

The Secure templates also provide further restrictions for anonymous users by preventing anonymous users (such as users from untrusted domains) from:

- Enumerating account names and shares.

- Performing SID-to-name or name-to-SID translations.

Finally, the Secure templates enable server-side Server Message Block (SMB) packet signing, which is disabled by default for servers. Since client-side SMB packet signing is enabled by default, SMB packet signing will always be negotiated when workstations and servers are operating at the Secure level.

- **Highly Secure (hisec*.inf)**

  The Highly Secure templates are supersets of the secure templates that impose further restrictions on the levels of encryption and signing that are required for authentication and for the data that flows over secure channels and between SMB clients and servers. For example, while the Secure templates cause servers to refuse LAN Manager responses, the Highly Secure templates cause servers to refuse both LAN Manager and NTLM responses. While the Secure template enables server-side SMB packet signing, the Highly Secure template requires it. Additionally, the Highly Secure templates require strong encryption and signing for the secure channel data that constitutes domain-to-member and domain-to-domain trust relationships. relationships. Thus, in order to apply hisecws.inf to a member machine:

  - All of the domain controllers that contain the accounts of all users that will logon to the client must be running Windows NT 4.0Service Pack 4 or higher.

  - All of the domain controllers for the domain that the client is joined to must run Windows 2000or later.

  - Clients that are configured with Hisecws.inf cannot connect to computers that only run LAN Manager or computers running Windows NT 4.0, Service Pack 3 or earlier using a local account defined on the target server.

  - Clients that are configured with Hisecws.inf cannot connect to servers running Windows 2000or Windows NT 4.0Service Pack 4 using a local account defined on the target server unless the clock on the target server is within 30 minutes of the clock on the client.

  - Clients that are configured with Hisecws.inf cannot connect to computers running Windows XPor above using a local account defined on the target computer unless the clock on the target computer is within 20 hours of the clock on the client.

  - Clients that are configured with Hisecws.inf cannot connect to LAN Manager servers operating in share-level security mode.

  - In order to apply Hisecdc.inf to a domain controller, all of the domain controllers in all trusted or trusting domains must run Windows 2000or Windows Server 2003 family operating systems.

  - If a server is configured with Hisecws.inf, then a user with a local account on that server will not be able to

connect to it from a client that does not support NTLMv2.

- If a server is configured with Hisecws.inf, then a client with a local account on that server will not be able to connect to the server unless the client's computer is configured to send NTLMv2 responses.

- If a server is configured with Hisecws.inf, then all clients that want to use SMB to connect to that server must enable client-side SMB packet signing. All computers running Windows 2000and Windows XPoperating systems enable client-side SMB packet signing by default.

- If a domain controller is configured with Hisecdc.inf, then a user with an account in that domain cannot connect to member servers using that domain user account if the connection is being attempted from a client that only uses the LAN Manager authentication protocol.

- If a domain controller is configured with Hisecdc.inf, then a user with an account in that domain will not be able to connect to member servers using that domain account unless:

Both the client and target server are running Windows 2000or above and can use Kerberos-based authentication rather than LAN Manager-based authentication.

The client is configured to send NTLMv2 responses.

- If a domain controller is configured with Hisecdc.inf, then Lightweight Directory Access protocol (LDAP) clients will not be able to bind with the Active Directory LDAP server unless data signing is negotiated. BIND requests using ldap_simple_bind or ldap_simple_bind_s are rejected. By default, all Microsoft LDAP clients that ship with Windows XPwill request data signing if Transport Layer Security/Secure Sockets Layer (TLS/SSL) is not already being used. If TLS/SSL is being used, then data signing is considered to be negotiated.

- Clients that do not support NTLMv2 include Windows for Workgroups, Windows NTclients prior to Service Pack 4, and Windows 95and Windows 98platforms that do not have the Active Directory Client Extensions Pack installed.

- Computers that run Windows NTService Pack 4 or higher can be configured to send only NTLMv2 responses by setting **HKLM\System\CurrentControlSet\Control\LSA\LMCompatibilityLevel**=3 or higher.

- Computers that run Windows 2000or higher can be configured to send only NTLMv2 responses by specifying this preference in the Network security: LAN Manager authentication level security option.

  In addition to further restrictions on the use of LAN Manager protocols and the requirements for encryption and signing of secure channel and SMB traffic, the Highly Secure templates also limit the use of cached logon data, such as data stored by Winlogon and Stored User Names and Passwords.

  Finally, Hisecws.inf uses restricted group settings to:

  - Remove all members of the Power Users group.

  - Ensure that only Domain Admins and the local Administrator account are members of the local Administrators group.

  Hisecws.inf defines these group restrictions under the assumption that users only use applications that belong to the Windows Logo program for Software. With those applications in place, neither the insecure Compatible template nor the insecure Power Users group is needed. Instead, users can run those applications successfully under the secure context of a normal member of the Users group, which is defined by the default security settings of the file system and registry. Members of the Administrators group can still use any application that they want.

  For information about the Windows Logo program for Software, see the Windows Logo program for Software page at the Microsoft Web site.

- **System root security (Rootsec.inf)**

Rootsec.inf specifies the root permissions. By default, Rootsec.inf defines these permissions for the root of the system drive. This template can be used to reapply the root directory permissions if they are inadvertently changed, or the template can be modified to apply the same root permissions to other volumes. As specified, the template does not overwrite explicit permissions that are defined on child objects; it propagates only the permissions that are inherited by child objects.

- **No Terminal Server user SID (Notssid.inf)**

  The default file system and registry access control lists that are on servers grant permissions to a Terminal Server security identifier (SID). The Terminal Server SID is used only when Terminal Server is running in application compatibility mode. If Terminal Server is not being used, this template can be applied to remove the unnecessary Terminal Server SIDs from the file system and registry locations. However, removing the access control entry for the Terminal Server SID from these default file system and registry locations does not increase the security of the system. Instead of removing the Terminal Server SID, simply run Terminal Server in Full Security mode. When running in Full Security mode, the Terminal Server SID is not used.

- **Internet Explorer (iesacls.inf)**

  The iesacls.inf template allows you to audit Internet Explorer use. This can be useful in conjunction with a plan to regularly review the event logs.

- **Optional Components (ocfiles*.inf)**

  These templates from previous Windows versions are no longer used or needed. Setup Security.inf replaces them.

**Caution**

- These security templates are constructed with the assumption that they will be applied to computers that use the default security settings. In other words, these templates incrementally modify the default security settings, if they are on the computer. They do *not* install the default security settings before performing the modifications.

- Predefined security templates should not be applied to production systems without testing to ensure that the right level of application functionality is maintained for your network and system architecture.

Security template settings can be viewed through Security Templates. The *.inf files can also be viewed as text files. These files are located in:

%windir%\Security\Templates

You cannot secure Windows XP Professionalsystems that are installed on file allocation table (FAT) file systems.

For more information, see Customize a predefined security template, Import a security template to a Group Policy object, Configure local computer security, and Reapply default security settings.

# Importing and exporting security templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Importing and exporting security templates**

Security Configuration and Analysis provides the ability to import and export security templates into or from a database.

If you have made any changes to the analysis database, you can save those settings by exporting them into a template. The export feature provides the ability to save the analysis database settings as a new template file. This template file can then be used to analyze or configure a system or it can be imported to a Group Policy object.

For more information, see Import a security template and Export a security template.

# Analyzing security and viewing results

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Analyzing security and viewing results**

Security Configuration and Analysis performs security analysis by comparing the current state of system security against an *analysis database*. During creation, the analysis database uses at least one security template. If you choose to import more than one security template, the database will merge the various templates and create one composite template. It resolves conflicts in order of import; the last template that is imported takes precedence.

**Viewing security analysis results**

Security Configuration and Analysis displays the analysis results by security area, using visual flags to indicate problems. It displays the current system and base configuration settings for each security attribute in the security areas. To change the analysis database settings, right-click the entry and click **Properties**.

| Visual flag | Meaning |
|---|---|
| Red X | The entry is defined in the analysis database and on the system, but the security setting values do not match. |
| Green check | The entry is defined in the analysis database and on the system and the setting values match. |
| Question mark | The entry is not defined in the analysis database and, therefore, was not analyzed.<br><br>If an entry is not analyzed, it may be that it was not defined in the analysis database or that the user who is running the analysis may not have sufficient permission to perform analysis on a specific object or area. |
| Exclamation point | This item is defined in the analysis database, but does not exist on the actual system. For example, there may be a restricted group that is defined in the analysis database but does not actually exist on the analyzed system. |
| No highlight | The item is not defined in the analysis database or on the system. |

If you choose to accept the current settings, the corresponding value in the base configuration is modified to match them. If you change the system setting to match the base configuration, the change will be reflected when you configure the system with Security Configuration and Analysis.

To avoid continued flagging of settings that you have investigated and determined to be reasonable, you can modify the base configuration. The changes are made to a copy of the template.

For more information, see Analyze system security, Resolving security discrepancies, and Security Setting Descriptions.

# Resolving security discrepancies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resolving security discrepancies**

You can resolve discrepancies between analysis database and system settings by:

- Accepting or changing some or all of the values that are flagged or not included in the configuration, if you determine that the local system security levels are valid due to the context (or role) of that computer. These attribute values are then updated in the base and applied to the system when you click **Configure Computer Now**. For more information, see Edit the analysis database.

- Configuring the system to the analysis database values, if you determine the system is not in compliance with valid security levels. For more information, see Configure local computer security.

- Importing a more appropriate template for the role of that computer into the database as the new base configuration and applying it to the system. For more information, see Import a security template.

Changes to the analysis database are made to the stored template in the database, not to the security template file. The security template file will only be modified if you either return to Security Templates and edit that template or export the stored configuration to the same template file.

You should use **Configure Computer Now** only to modify security areas *not* affected by Group Policy settings, such as security on local files and folders, registry keys, and system services. Otherwise, when the Group Policy settings are applied, it will take precedence over local settings—such as account policies. In general, do not use **Configure Computer Now** when you are analyzing security for domain-based clients, since you will have to configure each client individually. In this case, you should return to Security Templates, modify the template, and reapply it to the appropriate Group Policy object.

# Automating security configuration tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automating security configuration tasks**
By calling the Secedit.exe tool at a command prompt from a batch file or automatic task scheduler, you can use it to automatically create and apply templates, and analyze system security. You can also run it dynamically from a command prompt.

Secedit.exe is useful when you have multiple computers on which security must be analyzed or configured, and you need to perform these tasks during off-hours.

- For procedural information about using the Secedit.exe tool to perform security tasks, see Secedit.
- For information about defining Security Templates, see Security Templates.
- For more information about security configuration and analysis, see Security Configuration and Analysis.

**Note**

- The **secedit /refreshpolicy** command has been replaced with the **gpupdate** command. For information about how to refresh security settings, see Gpupdate.

# Security Setting Descriptions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security setting descriptions**

To locate this security setting in the console tree in Microsoft Management Console (MMC), see Computer Configuration\Windows Settings\Security Settings

| Security area | Description |
|---|---|
| Account Policies | Password Policy, Account Lockout Policy, and Kerberos Policy |
| Local Policies | Auditing Policy, User Rights Assignment, and Security Options |
| Event log | Application, system, and security Event Log settings |
| Restricted Groups | Membership of security-sensitive groups |
| System Services security settings | Startup and permissions for system services |
| Registry security settings | Permissions for registry keys |
| File System security settings | Permissions for folders and files |

You can configure the security settings that are described in this section on one computer, or on many computers, by using the Security Configuration Manager tools. The Security Configuration Manager tools consist of:

- Security Templates
- Security Configuration and Analysis
- The Secedit.exe command-prompt tool
- Local Security Policy
- The Security Settings Extension to Group Policy

To set or to modify individual security settings on individual computers, use Local Security Policy. To define security settings that are enforced on any number of computers, use the Security Settings Extension to Group Policy. To apply several settings in a batch, use Security Templates to define the settings, and then apply those settings by using Security Configuration and Analysis or Secedit.exe, or import the template that contains your settings into Local Policy or Group Policy.

# Account Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Account policies**
All security policies are computer-based policies. Account policies are defined on computers, yet they affect how user accounts can interact with the computer or domain. Account policies contain three subsets:

- Password Policy
- Account Lockout Policy
- Kerberos Policy

For domain accounts, the account policy must be defined in the Default Domain Policy Group Policy object (GPO) or in a new GPO that is linked to the root of the domain and given precedence over the Default Domain Policy GPO, which is enforced by the domain controllers that make up the domain. If more than one GPO containing account policy settings is linked at the domain level, the domain's account policy consists of the cumulative policy settings from all the domain-linked GPOs.

A domain controller always obtains the account policy from a GPO linked to the domain, which by default is the Default Domain Policy GPO. This behavior occurs even if a different account policy is applied to the organizational unit (OU) that contains the domain controller. By default, workstations and servers joined to a domain (such as member computers) also receive the same account policy for their local accounts. However, local account policies for member computers can be differentiated from the domain account policy by defining an account policy for the OU that contains the member computers.

**Important**

- Modifying the Default Domain Policy is not recommended. If you need to set some account policy that differs from that in the Default Domain Policy GPO, you can create a new GPO and link to the root of domain, set the policy you want to use, and assign it higher precedence than Default Domain Policy GPO.

For more information, see Account and local policies.

# Password Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Password policy**
Password policies are used for domain accounts or local user accounts. They determine settings for passwords, such as enforcement and lifetimes. This section covers:

- Enforce password history
- Maximum password age
- Minimum password age
- Minimum password length
- Passwords must meet complexity requirements
- Store passwords using reversible encryption

# Enforce password history

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enforce password history**

**Description**
This security setting determines the number of unique new passwords that have to be associated with a user account before an old password can be reused. The value must be between 0 and 24 passwords.

This policy enables administrators to enhance security by ensuring that old passwords are not reused continually.

**Default**:

- 24 on domain controllers.
- 0 on stand-alone servers.

**Note**

- By default, member computers follow the configuration of their domain controllers.

To maintain the effectiveness of the password history, do not allow passwords to be changed immediately after they were just changed by also enabling the Minimum password age security policy setting. For information about the minimum password age security policy setting, see Minimum password age.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

For specific instructions about how to configure password policy settings, see Apply or modify password policy.

For more information, see:

- Password Best practices for passwords
- Password Policy
- Security Configuration Manager tools

# Maximum password age

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maximum password age**

**Description**
This security setting determines the period of time (in days) that a password can be used before the system requires the user to change it. You can set passwords to expire after a number of days between 1 and 999, or you can specify that passwords never expire by setting the number of days to 0. If the maximum password age is between 1 and 999 days, the Minimum password age must be less than the maximum password age. If the maximum password age is set to 0, the minimum password age can be any value between 0 and 998 days.

**Note**

- It is a security best practice to have passwords expire every 30 to 90 days, depending on your environment. This way, an attacker has a limited amount of time in which to crack a user's password and have access to your network resources.

**Default:** 42.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

For specific instructions about how to configure password policy settings, see Apply or modify password policy.

For more information, see:

- Password Best practices for passwords
- Password Policy
- Security Configuration Manager tools

# Minimum password age

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Minimum password age**

**Description**
This security setting determines the period of time (in days) that a password must be used before the user can change it. You can set a value between 1 and 998 days, or you can allow changes immediately by setting the number of days to 0.

The minimum password age must be less than the Maximum password age, unless the maximum password age is set to 0, indicating that passwords will never expire. If the maximum password age is set to 0, the minimum password age can be set to any value between 0 and 998.

Configure the minimum password age to be more than 0 if you want Enforce password history to be effective. Without a minimum password age, users can cycle through passwords repeatedly until they get to an old favorite. The default setting does not follow this recommendation, so that an administrator can specify a password for a user and then require the user to change the administrator-defined password when the user logs on. If the password history is set to 0, the user does not have to choose a new password. For this reason, Enforce password history is set to 1 by default.

**Default**:

- 1 on domain controllers.
- 0 on stand-alone servers.

**Note**

- By default, member computers follow the configuration of their domain controllers.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

For specific instructions about how to configure password policy settings, see Apply or modify password policy.

For more information, see:

- Password Best practices for passwords
- Password Policy
- Security Configuration Manager tools

# Minimum password length

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Minimum password length**

**Description**
This security setting determines the least number of characters that a password for a user account may contain. You can set a value of between 1 and 14 characters, or you can establish that no password is required by setting the number of characters to 0.

**Default**:

- 7 on domain controllers.
- 0 on stand-alone servers.

**Note**

- By default, member computers follow the configuration of their domain controllers.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

For specific instructions about how to configure password policy settings, see Apply or modify password policy.

For more information, see:

- Password Best practices for passwords
- Password Policy
- Security Configuration Manager tools

# Passwords must meet complexity requirements

Updated: October 1, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

**Password must meet complexity requirements**

**Description**
This security setting determines whether passwords must meet complexity requirements. Complexity requirements are enforced when passwords are changed or created.

If this policy is enabled, passwords must meet the following minimum requirements when they are changed or created:

- Passwords must not contain the user's entire samAccountName (Account Name) value or entire displayName (Full Name) value. Both checks are not case sensitive:
    - The samAccountName is checked in its entirety only to determine whether it is part of the password. If the samAccountName is less than three characters long, this check is skipped.
    - The displayName is parsed for delimiters: commas, periods, dashes or hyphens, underscores, spaces, pound signs, and tabs. If any of these delimiters are found, the displayName is split and all parsed sections (tokens) are confirmed not to be included in the password. Tokens that are less than three characters in length are ignored, and substrings of the tokens are not checked. For example, the name "Erin M. Hagens" is split into three tokens: "Erin," "M," and "Hagens." Because the second token is only one character long, it is ignored. Therefore, this user could not have a password that included either "erin" or "hagens" as a substring anywhere in the password.
- Passwords must contain characters from three of the following five categories:
    - Uppercase characters of European languages (A through Z, with diacritic marks, Greek and Cyrillic characters)
    - Lowercase characters of European languages (a through z, sharp-s, with diacritic marks, Greek and Cyrillic characters)
    - Base 10 digits (0 through 9)
    - Nonalphanumeric characters: ~!@#$%^&*_-+=`|\(){}[]:;"'<>,.?/
    - Any Unicode character that is categorized as an alphabetic character but is not uppercase or lowercase. This includes Unicode characters from Asian languages.

| **Note** |
| --- |
| A given character can satisfy only one category. The GetStringTypeW API (http://go.microsoft.com/fwlink/?LinkId=205607) is used to test whether each character in the password is uppercase, lowercase, or alphanumeric. |

To create custom password filters, see Password Filters (http://go.microsoft.com/fwlink/?LinkId=205613).

**Default**:

- Enabled on domain controllers.
- Disabled on stand-alone servers.

| **Note** |
| --- |

| |
|---|
| By default, member computers inherit the password policy configuration from the domain. This password policy configuration is either defined at the domain-level or at any OU that has the member computer within its scope of management. If you define a password policy at an OU by using a GPO and that OU contains computer accounts, the computers that belong to those computer accounts inherit the password policy defined at the OU and apply it to the local users on the server. |

**Configuring this security setting**

You can configure this security setting by opening the appropriate policy and expanding the console tree: Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

For specific instructions about how to configure password policy settings, see Apply or modify password policy.

For more information, see:

- Passwords Technical Overview
- Password policy
- Strong passwords
- Password Best practices
- Security Configuration Manager tools

**Change History**

| Date | Revision |
|---|---|
| November 15, 2010 | The description was revised to be more precise and accurate. The precise attribute names were added and the character categories were updated. |
| November 30, 2010 | The description "Far eastern" was changed to "Asian." |
| Oct 1, 2012 | The Note about how domain computers obtain password policy was revised. |

# Store passwords using reversible encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Store passwords using reversible encryption**

**Description**
This security setting determines whether the operating system stores passwords using reversible encryption.

This policy provides support for applications that use protocols that require knowledge of the user's password for authentication purposes. Storing passwords using reversible encryption is essentially the same as storing plaintext versions of the passwords. For this reason, this policy should never be enabled unless application requirements outweigh the need to protect password information.

This policy is required when using Challenge-Handshake Authentication Protocol (CHAP) authentication through remote access or Internet Authentication Services (IAS). It is also required when using Digest Authentication in Internet Information Services (IIS).

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\

For specific instructions about how to configure password policy settings, see Apply or modify password policy.

For more information, see:

- Password Best practices for passwords
- Password Policy
- Security Configuration Manager tools

# Account Lockout Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Account lockout policy**
Account lockout policies are used for domain accounts or local user accounts. They determine the circumstances and length of time that an account will be locked out of the system. This section covers:

- Account lockout duration
- Account lockout threshold
- Reset account lockout counter after

# Account lockout duration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Account lockout duration**

**Description**
This security setting determines the number of minutes a locked-out account remains locked out before automatically becoming unlocked. The available range is from 0 minutes through 99,999 minutes. If you set the account lockout duration to 0, the account will be locked out until an administrator explicitly unlocks it.

If an account lockout threshold is defined, the account lockout duration must be greater than or equal to the reset time.

**Default:** None, because this policy setting only has meaning when an Account lockout threshold is specified.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy\

For specific instructions about how to configure account lockout policy settings, see Apply or modify account lockout policy.

For more information, see:

- Password Best practices for account lockout policy
- Account Lockout Policy
- Security Configuration Manager tools

# Account lockout threshold

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Account lockout threshold**

**Description**
This security setting determines the number of failed logon attempts that causes a user account to be locked out. A locked-out account cannot be used until it is reset by an administrator or until the lockout duration for the account has expired. You can set a value between 0 and 999 failed logon attempts. If you set the value to 0, the account will never be locked out.

Failed password attempts against workstations or member servers that have been locked using either CTRL+ALT+DELETE or password-protected screen savers count as failed logon attempts.

**Default:** 0.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy\

For specific instructions about how to configure account lockout policy settings, see Apply or modify account lockout policy.

For more information, see:

- Password Best practices for account lockout policy
- Account Lockout Policy
- Security Configuration Manager tools

# Reset account lockout counter after

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Reset account lockout counter after**

**Description**
This security setting determines the number of minutes that must elapse after a failed logon attempt before the failed logon attempt counter is reset to 0 bad logon attempts. The available range is 1 minute to 99,999 minutes.

If an account lockout threshold is defined, this reset time must be less than or equal to the Account lockout duration.

**Default:** None, because this policy setting only has meaning when an Account lockout threshold is specified.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy\

For specific instructions about how to configure account lockout policy settings, see Apply or modify account lockout policy.

For more information, see:

- Password Best practices for account lockout policy
- Account Lockout Policy
- Security Configuration Manager tools

# Kerberos Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Kerberos policy**
Kerberos policies are used for domain user accounts. They determine Kerberos-related settings, such as ticket lifetimes and enforcement. Kerberos policies do not exist in Local Computer Policy. This section covers:

- Enforce user logon restrictions
- Maximum lifetime for service ticket
- Maximum lifetime for user ticket
- Maximum lifetime for user ticket renewal
- Maximum tolerance for computer clock synchronization

# Enforce user logon restrictions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enforce user logon restrictions**

**Description**
This security setting determines whether the Kerberos V5 Key Distribution Center (KDC) validates every request for a session ticket against the user rights policy of the user account. Validation of each request for a session ticket is optional, because the extra step takes time and it may slow network access to services.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Kerberos Policy
- Configure Kerberos policy
- Security Configuration Manager tools

# Maximum lifetime for service ticket

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maximum lifetime for service ticket**

**Description**
This security setting determines the maximum amount of time (in minutes) that a granted session ticket can be used to access a particular service. The setting must be greater than 10 minutes and less than or equal to the setting for Maximum lifetime for user ticket.

If a client presents an expired session ticket when it requests a connection to a server, the server returns an error message. The client must request a new session ticket from the Kerberos V5 Key Distribution Center (KDC). Once a connection is authenticated, however, it no longer matters whether the session ticket remains valid. Session tickets are used only to authenticate new connections with servers. Ongoing operations are not interrupted if the session ticket that is used to authenticate the connection expires during the connection.

**Default:** 600 minutes (10 hours).

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Kerberos Policy
- Configure Kerberos policy
- Security Configuration Manager tools

# Maximum lifetime for user ticket

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maximum lifetime for user ticket**

**Description**
This security setting determines the maximum amount of time (in hours) that a user's ticket-granting ticket (TGT) may be used. When a user's TGT expires, a new one must be requested or the existing one must be "renewed."

**Default:** 10 hours.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Kerberos Policy
- Configure Kerberos policy
- Security Configuration Manager tools

# Maximum lifetime for user ticket renewal

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maximum lifetime for user ticket renewal**

**Description**
This security setting determines the period of time (in days) during which a user's ticket-granting ticket (TGT) may be renewed.

**Default:** 7 days.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Kerberos Policy
- Configure Kerberos policy
- Security Configuration Manager tools

# Maximum tolerance for computer clock synchronization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maximum tolerance for computer clock synchronization**

**Description**
This security setting determines the maximum time difference (in minutes) that Kerberos V5 tolerates between the time on the client clock and the time on the domain controller running Windows Server 2003 that provides Kerberos authentication.

To prevent "replay attacks," Kerberos V5 uses time stamps as part of its protocol definition. For time stamps to work properly, the clocks of the client and the domain controller need to be in sync as much as possible. In other words, both computers must be set to the same time and date. Because the clocks of two computers are often out of sync, administrators can use this policy to establish the maximum acceptable difference to Kerberos V5 between a client clock and domain controller clock. If the difference between a client clock and the domain controller clock is less than the maximum time difference that is specified in this policy, any time stamp that is used in a session between the two computers is considered to be authentic.

**Important**

- This setting is not persistent. If you configure this setting and then restart the computer, this setting reverts to the default value. For more information, see the Microsoft Knowledge Base.

**Default:** 5 minutes.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Kerberos Policy
- Configure Kerberos policy
- Security Configuration Manager tools

# Local Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local policies**

These policies apply to computers, and they contain these subsets:

- Auditing Policy. Determines whether security events are logged into the Security log on the computer. Also determines whether to log successful attempts, failed attempts or both. (The Security log is part of Event Viewer.)
- User Rights Assignment. Determines which users or groups have logon rights or privileges on the computer.
- Security Options. Enables or disables security settings for the computer, such as digital signing of data, Administrator and Guest account names, floppy drive and CD-ROM access, driver installation, and logon prompts.

# Auditing Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Auditing policy**
Before you implement auditing policy, you must decide which event categories you want to audit. The auditing settings that you choose for the event categories define your auditing policy. On member servers and workstations that are joined to a domain, auditing settings for the event categories are undefined by default. On domain controllers, auditing is turned on by default. By defining auditing settings for specific event categories, you can create an auditing policy that suits the security needs of your organization.

The event categories that you can choose to audit are:

- Audit account logon events
- Audit account management
- Audit directory service access
- Audit logon events
- Audit object access
- Audit policy change
- Audit privilege use
- Audit process tracking
- Audit system events

For more information about auditing, see Auditing Security Events.

# Audit account logon events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Vista

**Audit account logon events**

**Description**
This security setting determines whether to audit each instance of a user logging on to or logging off from another computer in which this computer is used to validate the account. Account logon events are generated when a domain user account is authenticated on a domain controller. The event is logged in the domain controller's security log. Logon events are generated when a local user is authenticated on a local computer. The event is logged in the local security log. Account logoff events are not generated.

If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit the event type at all. Success audits generate an audit entry when an account logon attempt succeeds. Failure audits generate an audit entry when an account logon attempt fails.

To set this value to **No auditing**, in the **Properties** dialog box for this policy setting, select the **Define these policy settings** check box and clear the **Success** and **Failure** check boxes.

If success auditing for account logon events is enabled on a domain controller, an entry is logged for each user who is validated against that domain controller, even though the user is actually logging on to a workstation that is joined to the domain.

**Default**: Success.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

For specific instructions about how to configure auditing policy settings, see Define or modify auditing policy settings for an event category.

| Account Logon Events | Description |
|---|---|
| 672 | An authentication service (AS) ticket was successfully issued and validated. |
| 673 | A ticket granting service (TGS) ticket was granted. |
| 674 | A security principal renewed an AS ticket or TGS ticket. |
| 675 | Preauthentication failed. This event is generated on a Key Distribution Center (KDC) when a user types in an incorrect password. |
| 676 | Authentication ticket request failed. This event is not generated in Windows XP or in the Windows Server 2003 family. |
| 677 | A TGS ticket was not granted. This event is not generated in Windows XP or in the Windows Server 2003 family. |
| 678 | An account was successfully mapped to a domain account. |
| 681 | Logon failure. A domain account logon was attempted. This event is not generated in Windows XP or in the Windows Server 2003 family. |
| 682 | A user has reconnected to a disconnected terminal server session. |
| 683 | A user disconnected a terminal server session without logging off. |

For more information about security events, see Security Events on the Microsoft Windows Resource Kits Web site.

For more information, see:

- Audit logon events
- Auditing Policy
- Auditing Security Events Best practices for auditing
- Security Configuration Manager tools

# Audit account management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Vista

**Audit account management**

**Description**
This security setting determines whether to audit each event of account management on a computer. Examples of account management events include:

- A user account or group is created, changed, or deleted.
- A user account is renamed, disabled, or enabled.
- A password is set or changed.

If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit the event type at all. Success audits generate an audit entry when any account management event succeeds. Failure audits generate an audit entry when any account management event fails. To set this value to **No auditing**, in the **Properties** dialog box for this policy setting, select the **Define these policy settings** check box and clear the **Success** and **Failure** check boxes.

**Default**:

- Success on domain controllers.
- No auditing on member servers

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

For specific instructions about how to configure auditing policy settings, see Define or modify auditing policy settings for an event category.

| Account Management Events | Description |
|---|---|
| 624 | A user account was created. |
| 627 | A user password was changed. |
| 628 | A user password was set. |
| 630 | A user account was deleted. |
| 631 | A global group was created. |
| 632 | A member was added to a global group. |
| 633 | A member was removed from a global group. |
| 634 | A global group was deleted. |
| 635 | A new local group was created. |
| 636 | A member was added to a local group. |

| | |
|---|---|
| 637 | A member was removed from a local group. |
| 638 | A local group was deleted. |
| 639 | A local group account was changed. |
| 641 | A global group account was changed. |
| 642 | A user account was changed. |
| 643 | A domain policy was modified. |
| 644 | A user account was auto locked. |
| 645 | A computer account was created. |
| 646 | A computer account was changed. |
| 647 | A computer account was deleted. |
| 648 | A local security group with security disabled was created.<br>**Note**<br>• SECURITY_DISABLED in the formal name means that this group cannot be used to grant permissions in access checks. |
| 649 | A local security group with security disabled was changed. |
| 650 | A member was added to a security-disabled local security group. |
| 651 | A member was removed from a security-disabled local security group. |
| 652 | A security-disabled local group was deleted. |
| 653 | A security-disabled global group was created. |
| 654 | A security-disabled global group was changed. |
| 655 | A member was added to a security-disabled global group. |
| 656 | A member was removed from a security-disabled global group. |
| 657 | A security-disabled global group was deleted. |
| 658 | A security-enabled universal group was created. |
| 659 | A security-enabled universal group was changed. |
| 660 | A member was added to a security-enabled universal group. |
| 661 | A member was removed from a security-enabled universal group. |
| 662 | A security-enabled universal group was deleted. |
| 663 | A security-disabled universal group was created. |
| 664 | A security-disabled universal group was changed. |
| 665 | A member was added to a security-disabled universal group. |
| 666 | A member was removed from a security-disabled universal group. |
| 667 | A security-disabled universal group was deleted. |

| | |
|---|---|
| 668 | A group type was changed. |
| 684 | Set the security descriptor of members of administrative groups.<br>**Note**<br>• Every 60 minutes on a domain controller a background thread searches all members of administrative groups (such as domain, enterprise, and schema administrators) and applies a fixed security descriptor on them. This event is logged. |
| 685 | Name of an account was changed. |

For more information about security events, see Security Events on the Microsoft Windows Resource Kits Web site.

For more information, see:

- Auditing Policy
- Auditing Security Events Best practices
- Security Configuration Manager tools

# Audit directory service access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Vista

**Audit directory service access**

**Description**
This security setting determines whether to audit the event of a user accessing an Active Directory object that has its own system access control list (SACL) specified.

By default, this value is set to no auditing in the Default Domain Controller Group Policy object (GPO), and it remains undefined for workstations and servers where it has no meaning.

If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit the event type at all. Success audits generate an audit entry when a user successfully accesses an Active Directory object that has a SACL specified. Failure audits generate an audit entry when a user unsuccessfully attempts to access an Active Directory object that has a SACL specified. To set this value to **No auditing**, in the **Properties** dialog box for this policy setting, select the **Define these policy settings** check box and clear the **Success** and **Failure** check boxes.

Note that you can set a SACL on an Active Directory object by using the **Security** tab in that object's **Properties** dialog box. This is the same as Audit object access, except that it applies only to Active Directory objects and not to file system and registry objects.

**Default**:

- Success on domain controllers.
- Undefined for a member computer.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

For specific instructions about how to configure auditing policy settings, see Define or modify auditing policy settings for an event category.

There is only one directory service access event, which is identical to the Object Access security event message 566.

| Directory service access events | Description |
|---|---|
| 566 | A generic object operation took place. |

For more information about security events, see Security Events on the Microsoft Windows Resource Kits Web site.

For more information, see:

- Auditing Policy
- Auditing Security Events Best practices for auditing
- Security Configuration Manager tools

# Audit logon events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Vista

**Audit logon events**

**Description**
This security setting determines whether to audit each instance of a user logging on to or logging off from a computer.

Account logon events are generated on domain controllers for domain account activity and on local computers for local account activity. If both account logon and logon audit policy categories are enabled, logons that use a domain account generate a logon or logoff event on the workstation or server, and they generate an account logon event on the domain controller. Additionally, interactive logons to a member server or workstation that use a domain account generate a logon event on the domain controller as the logon scripts and policies are retrieved when a user logs on. For more information about account logon events, see Audit account logon events.

If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit the event type at all. Success audits generate an audit entry when a logon attempt succeeds. Failure audits generate an audit entry when a logon attempt fails.

To set this value to **No auditing**, in the **Properties** dialog box for this policy setting, select the **Define these policy settings** check box and clear the **Success** and **Failure** check boxes.

**Default**: Success.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

For specific instructions about how to configure auditing policy settings, see Define or modify auditing policy settings for an event category.

| Logon Events | Description |
|---|---|
| 528 | A user successfully logged on to a computer. For information about the type of logon, see the Logon Types table below. |
| 529 | Logon failure. A logon attempt was made with an unknown user name or a known user name with a bad password. |
| 530 | Logon failure. A logon attempt was made user account tried to log on outside of the allowed time. |
| 531 | Logon failure. A logon attempt was made using a disabled account. |
| 532 | Logon failure. A logon attempt was made using an expired account. |
| 533 | Logon failure. A logon attempt was made by a user who is not allowed to log on at this computer. |
| 534 | Logon failure. The user attempted to log on with a type that is not allowed. |
| 535 | Logon failure. The password for the specified account has expired. |
| 536 | Logon failure. The Net Logon service is not active. |
| 537 | Logon failure. The logon attempt failed for other reasons.<br>**Note**<br>• In some cases, the reason for the logon failure may not be known. |

| | |
|---|---|
| 538 | The logoff process was completed for a user. |
| 539 | Logon failure. The account was locked out at the time the logon attempt was made. |
| 540 | A user successfully logged on to a network. |
| 541 | Main mode Internet Key Exchange (IKE) authentication was completed between the local computer and the listed peer identity (establishing a security association), or quick mode has established a data channel. |
| 542 | A data channel was terminated. |
| 543 | Main mode was terminated.<br>**Note**<br>• This might occur as a result of the time limit on the security association expiring (the default is eight hours), policy changes, or peer termination. |
| 544 | Main mode authentication failed because the peer did not provide a valid certificate or the signature was not validated. |
| 545 | Main mode authentication failed because of a Kerberos failure or a password that is not valid. |
| 546 | IKE security association establishment failed because the peer sent a proposal that is not valid. A packet was received that contained data that is not valid. |
| 547 | A failure occurred during an IKE handshake. |
| 548 | Logon failure. The security ID (SID) from a trusted domain does not match the account domain SID of the client. |
| 549 | Logon failure. All SIDs corresponding to untrusted namespaces were filtered out during an authentication across forests. |
| 550 | Notification message that could indicate a possible denial-of-service attack. |
| 551 | A user initiated the logoff process. |
| 552 | A user successfully logged on to a computer using explicit credentials while already logged on as a different user. |
| 682 | A user has reconnected to a disconnected terminal server session. |
| 683 | A user disconnected a terminal server session without logging off.<br>**Note**<br>• This event is generated when a user is connected to a terminal server session over the network. It appears on the terminal server. |

When event 528 is logged, a logon type is also listed in the event log. The following table describes each logon type.

| **Logon type** | **Logon title** | **Description** |
|---|---|---|
| 2 | Interactive | A user logged on to this computer. |
| 3 | Network | A user or computer logged on to this computer from the network. |
| 4 | Batch | Batch logon type is used by batch servers, where processes may be executing on behalf of a user without their direct intervention. |
| 5 | Service | A service was started by the Service Control Manager. |
| 7 | Unlock | This workstation was unlocked. |

| | | |
|---|---|---|
| 8 | NetworkCleartext | A user logged on to this computer from the network. The user's password was passed to the authentication package in its unhashed form. The built-in authentication packages all hash credentials before sending them across the network. The credentials do not traverse the network in plaintext (also called cleartext). |
| 9 | NewCredentials | A caller cloned its current token and specified new credentials for outbound connections. The new logon session has the same local identity, but uses different credentials for other network connections. |
| 10 | RemoteInteractive | A user logged on to this computer remotely using Terminal Services or Remote Desktop. |
| 11 | CachedInteractive | A user logged on to this computer with network credentials that were stored locally on the computer. The domain controller was not contacted to verify the credentials. |

For more information about security events, see Security Events on the Microsoft Windows Resource Kits Web site.

For more information, see:

- Auditing Policy
- Auditing Security Events Best practices for auditing
- Security Configuration Manager tools

# Audit object access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Vista

**Audit object access**

**Description**
This security setting determines whether to audit the event of a user accessing an object--for example, a file, folder, registry key, printer, and so forth--that has its own system access control list (SACL) specified.

If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit the event type at all. Success audits generate an audit entry when a user successfully accesses an object that has an appropriate SACL specified. Failure audits generate an audit entry when a user unsuccessfully attempts to access an object that has a SACL specified.

To set this value to **No auditing**, in the **Properties** dialog box for this policy setting, select the **Define these policy settings** check box and clear the **Success** and **Failure** check boxes.

Note that you can set a SACL on a file system object using the **Security** tab in that object's **Properties** dialog box.

**Default**: No auditing.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

For specific instructions about how to configure auditing policy settings, see Define or modify auditing policy settings for an event category.

| Object Access Events | | Description |
|---|---|---|
| 560 | Access was granted to an already existing object. | |
| 562 | A handle to an object was closed. | |
| 563 | An attempt was made to open an object with the intent to delete it.<br>**Note**<br>• This is used by file systems when the FILE_DELETE_ON_CLOSE flag is specified in Createfile(). | |
| 564 | A protected object was deleted. | |
| 565 | Access was granted to an already existing object type. | |
| 567 | A permission associated with a handle was used.<br>**Note**<br>• A handle is created with certain granted permissions (Read, Write, and so on). When the handle is used, up to one audit is generated for each of the permissions that was used. | |
| 568 | An attempt was made to create a hard link to a file that is being audited. | |
| 569 | The resource manager in Authorization Manager attempted to create a client context. | |

| | |
|---|---|
| 570 | A client attempted to access an object.<br>**Note**<br>• An event will be generated for every attempted operation on the object. |
| 571 | The client context was deleted by the Authorization Manager application. |
| 572 | The administrator manager initialized the application. |
| 772 | The certificate manager denied a pending certificate request. |
| 773 | Certificate Services received a resubmitted certificate request. |
| 774 | Certificate Services revoked a certificate. |
| 775 | Certificate Services received a request to publish the certificate revocation list (CRL). |
| 776 | Certificate Services published the certificate revocation list (CRL). |
| 777 | A certificate request extension was made. |
| 778 | One or more certificate request attributes changed. |
| 779 | Certificate Services received a request to shutdown. |
| 780 | Certificate Services backup started. |
| 781 | Certificate Services backup completed. |
| 782 | Certificate Services restore started. |
| 783 | Certificate Services restore completed. |
| 784 | Certificate Services started. |
| 785 | Certificate Services stopped. |
| 786 | The security permissions for Certificate Services changed. |
| 787 | Certificate Services retrieved an archived key. |
| 788 | Certificate Services imported a certificate into its database. |
| 789 | The audit filter for Certificate Services changed. |
| 790 | Certificate Services received a certificate request. |
| 791 | Certificate Services approved a certificate request and issued a certificate. |
| 792 | Certificate Services denied a certificate request. |
| 793 | Certificate Services set the status of a certificate request to pending. |
| 794 | The certificate manager settings for Certificate Services changed. |
| 795 | A configuration entry changed in Certificate Services. |
| 796 | A property of Certificate Services changed. |
| 797 | Certificate Services archived a key. |
| 798 | Certificate Services imported and archived a key. |

| | |
|---|---|
| 799 | Certificate Services published the CA certificate to Active Directory. |
| 800 | One or more rows have been deleted from the certificate database. |
| 801 | Role separation enabled. |

For more information about security events, see Security Events on the Microsoft Windows Resource Kits Web site.

For more information, see:

- Audit directory service access
- Auditing Policy
- Auditing Security Events Best practices
- Security Configuration Manager tools

# Audit policy change

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Vista

**Audit policy change**

**Description**
This security setting determines whether to audit every incident of a change to user rights assignment policies, audit policies, or trust policies.

If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit the event type at all. Success audits generate an audit entry when a change to user rights assignment policies, audit policies, or trust policies is successful. Failure audits generate an audit entry when a change to user rights assignment policies, audit policies, or trust policies fails.

To set this value to **No auditing**, in the **Properties** dialog box for this policy setting, select the **Define these policy settings** check box and clear the **Success** and **Failure** check boxes.

**Default**:

- Success on domain controllers.
- No auditing on member servers.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

For specific instructions about how to configure auditing policy settings, see Define or modify auditing policy settings for an event category.

| Policy Change Events | Description |
|---|---|
| 608 | A user right was assigned. |
| 609 | A user right was removed. |
| 610 | A trust relationship with another domain was created. |
| 611 | A trust relationship with another domain was removed. |
| 612 | An audit policy was changed. |
| 613 | An Internet Protocol security (IPSec) policy agent started. |
| 614 | An IPSec policy agent was disabled. |
| 615 | An IPSec policy agent changed. |
| 616 | An IPSec policy agent encountered a potentially serious failure. |
| 617 | A Kerberos policy changed. |
| 618 | Encrypted Data Recovery policy changed. |
| 620 | A trust relationship with another domain was modified. |
| 621 | System access was granted to an account. |

| | |
|---|---|
| 622 | System access was removed from an account. |
| 623 | Per user auditing policy was set for a user.<br><br>For information about per user selective auditing, see Per-user selective auditing. |
| 625 | Per user audit policy was refreshed. |
| 768 | A collision was detected between a namespace element in one forest and a namespace element in another forest.<br><br>**Note**<br><br>• When a namespace element in one forest overlaps a namespace element in another forest, it can lead to ambiguity in resolving a name belonging to one of the namespace elements. This overlap is also called a collision. Not all parameters are valid for each entry type. For example, fields such as DNS name, NetBIOS name, and SID are not valid for an entry of type 'TopLevelName'. |
| 769 | Trusted forest information was added.<br><br>**Note**<br><br>• This event message is generated when forest trust information is updated and one or more entries are added. One event message is generated per added, deleted, or modified entry. If multiple entries are added, deleted, or modified in a single update of the forest trust information, all the generated event messages have a single unique identifier called an operation ID. This allows you to determine that the multiple generated event messages are the result of a single operation. Not all parameters are valid for each entry type. For example, parameters such as DNS name, NetBIOS name and SID are not valid for an entry of type "TopLevelName". |
| 770 | Trusted forest information was deleted.<br><br>**Note**<br><br>• See note for event 769. |
| 771 | Trusted forest information was modified.<br><br>**Note**<br><br>• See note for event 769. |
| 805 | The event log service read the security log configuration for a session. |

For more information about security events, see Security Events in the Microsoft Windows Resource Kits on the Microsoft Web site.

For more information, see:

- Auditing Policy
- Auditing Security Events Best practices
- Security Configuration Manager tools

# Audit privilege use

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Vista

**Audit privilege use**

**Description**
This security setting determines whether to audit each instance of a user exercising a user right.

If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit this type of event at all. Success audits generate an audit entry when the exercise of a user right succeeds. Failure audits generate an audit entry when the exercise of a user right fails.

To set this value to **No auditing**, in the **Properties** dialog box for this policy setting, select the **Define these policy settings** check box and clear the **Success** and **Failure** check boxes.

**Default**: No auditing.

Audits are not generated for use of the following user rights, even if success audits or failure audits are specified for **Audit privilege use**. Enabling auditing of these user rights tend to generate many events in the security log which may impede your computer's performance. To audit the following user rights, enable the **FullPrivilegeAuditing** registry key.

- Bypass traverse checking
- Debug programs
- Create a token object
- Replace process level token
- Generate security audits
- Back up files and directories
- Restore files and directories

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

For specific instructions about how to configure auditing policy settings, see Define or modify auditing policy settings for an event category.

| Privilege use events | Description |
|---|---|
| 576 | Specified privileges were added to a user's access token.<br>**Note**<br>• This event is generated when the user logs on. |

| | |
|---|---|
| 577 | A user attempted to perform a privileged system service operation. |
| 578 | Privileges were used on an already open handle to a protected object. |

For more information about security events, see Security Events on the Microsoft Windows Resource Kits Web site.

For more information, see:

- Auditing Policy
- Auditing Security Events Best practices for auditing
- Security Configuration Manager tools

# Audit process tracking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Vista

**Audit process tracking**

**Description**
This security setting determines whether to audit detailed tracking information for events such as program activation, process exit, handle duplication, and indirect object access.

If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit the event type at all. Success audits generate an audit entry when the process being tracked succeeds. Failure audits generate an audit entry when the process being tracked fails.

To set this value to **No auditing**, in the **Properties** dialog box for this policy setting, select the **Define these policy settings** check box and clear the **Success** and **Failure** check boxes.

**Default**: No auditing

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

For specific instructions about how to configure auditing policy settings, see Define or modify auditing policy settings for an event category.

| Event | Event Description |
|---|---|
| 592 | A new process was created. |
| 593 | A process exited. |
| 594 | A handle to an object was duplicated. |
| 595 | Indirect access to an object was obtained. |
| 596 | A data protection master key was backed up.<br>**Note**<br>• The master key is used by the CryptProtectData and CryptUnprotectData routines, and Encrypting File System (EFS). The master key is backed up each time a new one is created. (The default setting is 90 days.) The key is usually backed up to a domain controller. |
| 597 | A data protection master key was recovered from a recovery server. |
| 598 | Auditable data was protected. |
| 599 | Auditable data was unprotected. |
| 600 | A process was assigned a primary token. |
| 601 | A user attempted to install a service. |
| 602 | A scheduler job was created. |

For more information about security events, see Security Events on the Microsoft Windows Resource Kits Web site.

For more information, see:

- Auditing Policy

- Auditing Security Events Best practices
- Security Configuration Manager tools

# Audit system events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Vista

**Audit system events**

**Description**
This security setting determines whether to audit when a user restarts or shuts down the computer or when an event occurs that affects either the system security or the security log.

If you define this policy setting, you can specify whether to audit successes, audit failures, or not audit the event type at all. Success audits generate an audit entry when a system event is executed successfully. Failure audits generate an audit entry when a system event is attempted unsuccessfully.

To set this value to **No auditing**, in the **Properties** dialog box for this policy setting, select the **Define these policy settings** check box and clear the **Success** and **Failure** check boxes.

**Default**:

- Success on domain controllers.
- No auditing on member servers.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy\

For specific instructions about how to configure auditing policy settings, see Define or modify auditing policy settings for an event category.

| System Event Messages | Description |
|---|---|
| 512 | Windows is starting up. |
| 513 | Windows is shutting down. |
| 514 | An authentication package was loaded by the Local Security Authority. |
| 515 | A trusted logon process has registered with the Local Security Authority. |
| 516 | Internal resources allocated for the queuing of security event messages have been exhausted, leading to the loss of some security event messages. |
| 517 | The audit log was cleared. |
| 518 | A notification package was loaded by the Security Accounts Manager. |
| 519 | A process is using an invalid local procedure call (LPC) port in an attempt to impersonate a client and reply or read from or write to a client address space. |
| 520 | The system time was changed.<br>**Note**<br>• This audit normally appears twice. |

For more information about security events, see Security Events on the Microsoft Windows Resource Kits Web site.

For more information, see:

- Auditing Policy
- Auditing Security Events Best practices
- Security Configuration Manager tools

# User Rights Assignment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User rights assignment**
This section covers:

- Access this computer from the network
- Act as part of the operating system
- Add workstations to domain
- Adjust memory quotas for a process
- Allow log on locally
- Allow log on through Terminal Services
- Back up files and directories
- Bypass traverse checking
- Change the system time
- Create a pagefile
- Create a token object
- Create global objects
- Create permanent shared objects
- Debug programs
- Deny access to this computer from the network
- Deny log on as a batch job
- Deny log on as a service
- Deny log on locally
- Deny log on through Terminal Services
- Enable computer and user accounts to be trusted for delegation
- Force shutdown from a remote system
- Generate security audits
- Impersonate a client after authentication
- Increase scheduling priority

- Load and unload device drivers
- Lock pages in memory
- Log on as a batch job
- Log on as a service
- Manage auditing and security log
- Modify firmware environment values
- Perform volume maintenance tasks
- Profile single process
- Profile system performance
- Remove computer from docking station
- Replace a process level token
- Restore files and directories
- Shut down the system
- Synchronize directory service data
- Take ownership of files or other objects

# Access this computer from the network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Access this computer from the network**

**Description**
This user right determines which users and groups are allowed to connect to the computer over the network. Terminal Services are not affected by this user right.

**Default:**

- On workstations and servers:
    - Administrators
    - Backup Operators
    - Power Users
    - Users
    - Everyone
- On domain controllers:
    - Administrators
    - Authenticated Users
    - Everyone

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Act as part of the operating system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Act as part of the operating system**

**Description**
This user right allows a process to impersonate any user without authentication. The process can therefore gain access to the same local resources as that user.

Processes that require this privilege should use the *LocalSystem* account, which already includes this privilege, rather than using a separate user account with this privilege specially assigned. If your organization only uses servers that are members of the Windows Server 2003 family, you do not need to assign this privilege to your users. However, if your organization uses servers running Windows 2000 or Windows NT 4.0, you might need to assign this privilege to use applications that exchange passwords in plaintext.

**Caution**

- Assigning this user right can be a security risk. Only assign this user right to trusted users.

**Default:** Local System.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and opening the following folder in the console tree: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Add workstations to domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Add workstations to domain**

**Description**
This security setting determines which groups or users can add workstations to a domain.

This security setting is valid only on domain controllers. By default, any authenticated user has this right and can create up to 10 computer accounts in the domain.

Adding a computer account to the domain allows the computer to participate in Active Directory-based networking. For example, adding a workstation to a domain enables that workstation to recognize accounts and groups that exist in Active Directory.

**Default:** Authenticated Users on domain controllers.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Note**

- Users who have the Create Computer Objects permission on the Active Directory computers container can also create computer accounts in the domain. The distinction is that users with permissions on the container are not restricted to the creation of only 10 computer accounts. In addition, computer accounts that are created by means of **Add workstations to domain** have Domain Administrators as the owner of the computer account, while computer accounts that are created by means of permissions on the computers container have the creator as the owner of the computer account. If a user has permissions on the container and also has the **Add workstations to domain** user right, the computer is added, based on the computer container permissions rather than on the user right.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Adjust memory quotas for a process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adjust memory quotas for a process**

**Description**
This privilege determines who can change the maximum memory that can be consumed by a process.

This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.

**Note**

- This privilege is useful for system tuning, but it can be misused, for example, in a denial-of-service attack.

**Default:** Administrators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Allow log on locally

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Allow log on locally**

**Description**
This logon right determines which users can interactively log on to this computer. Logons initiated by pressing CTRL+ALT+DEL sequence on the attached keyboard requires the user to have this logon right. Additionally this logon right may be required by some service or administrative applications that can log on users. If you define this policy for a user or group, you must also give the **Administrators** group this right.

**Default:**

- On workstations and servers: Administrators, Backup Operators, Power Users, Users, and Guest.
- On domain controllers: Account Operators, Administrators, Backup Operators, Print Operators, and Server Operators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Deny log on locally
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Allow log on through Terminal Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Allow log on through Terminal Services**

**Description**
This security setting determines which users or groups have permission to log on as a Terminal Services client.

**Default:**

- On workstation and servers: Administrators, Remote Desktop Users.
- On domain controllers: Administrators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- This setting does not have any effect on Windows 2000 computers that have not been updated to Service Pack 2.

For more information, see:

- Deny log on through Terminal Services
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Back up files and directories

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Back up files and directories**

**Description**
This user right determines which users can bypass file and directory, registry, and other persistent object permissions for the purposes of backing up the system.

Specifically, this user right is similar to granting the following permissions to the user or group in question on all files and folders on the system:

- Traverse Folder/Execute File
- List Folder/Read Data
- Read Attributes
- Read Extended Attributes
- Read Permissions

**Caution**

- Assigning this user right can be a security risk. Since there is no way to be sure that a user is backing up data, stealing data, or copying data to be distributed, only assign this user right to trusted users.

**Default:** Administrators and Backup Operators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Restore files and directories
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Bypass traverse checking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Bypass traverse checking**

**Description**
This user right determines which users can traverse directory trees even though the user may not have permissions on the traversed directory. This privilege does not allow the user to list the contents of a directory, only to traverse directories.

This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.

**Default:**

- On workstations and servers:
    - Administrators
    - Backup Operators
    - Power Users
    - Users
    - Everyone
- On domain controllers:
    - Administrators
    - Authenticated Users

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Change the system time

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Change the system time**

**Description**
This user right determines which users and groups can change the time and date on the internal clock of the computer. Users that are assigned this user right can affect the appearance of event logs. If the system time is changed, events that are logged will reflect this new time, not the actual time that the events occurred.

This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.

**Default:**

- On workstations and servers:
    - Administrators
    - Power Users
- On domain controllers:
    - Administrators
    - Server Operators

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Create a pagefile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create a pagefile**

**Description**
This user right determines which users and groups can call an internal application programming interface (API) to create a page file. This user right is used internally by the operating system and usually does not need to be assigned to any users.

For information about how to specify a paging file size for a given drive, see Change the size of the virtual memory paging file.

**Default:** Administrators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Create a token object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create a token object**

**Description**
This security setting determines which accounts can be used by processes to create a token that can then be used to get access to any local resources when the process uses an internal application programming interface (API) to create an access token.

This user right is used internally by the operating system. Unless it is necessary, do not assign this user right to a user, group, or process other than *Local System*.

**Caution**

- Assigning this user right can be a security risk. Do not assign this user right to any user, group, or process that you do not want to take over the system.

**Default:** Local System.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Create global objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create global objects**

**Description**
This user right is required for a user account to create global objects during Terminal Services sessions. Users can still create session-specific objects without being assigned this user right.

**Caution**

- Assigning this user right can be a security risk. Assign this user right only to trusted users.

**Default:**

- Administrators
- Local System
- Services

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Create permanent shared objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create permanent shared objects**

**Description**
This user right determines which accounts can be used by processes to create a directory object in the Windows 2000 Server, Windows 2000 Professional, Windows XP Professional, and Windows Server 2003 family object manager.

This user right is used internally by the operating system and is useful to kernel-mode components that extend the Windows 2000 Server, Windows 2000 Professional, Windows XP Professional, and Windows Server 2003 family object namespace. Because components that are running in kernel mode already have this user right assigned to them, it is not necessary to specifically assign it.

**Default:** Local System.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Debug programs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Debug programs**

**Description**
This user right determines which users can attach a debugger to any process or to the kernel. Developers who are debugging their own applications to not need to be assigned this user right. Developers who are debugging new system components will need this user right to be able to do so. This user right provides complete access to sensitive and critical operating system components.

### Caution

- Assigning this user right can be a security risk. Only assign this user right to trusted users.

### Default:

- Administrators
- Local System

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Deny access to this computer from the network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deny access to this computer from the network**

**Description**
This security setting determines which users are prevented from accessing a computer over the network. This policy setting supersedes the Access this computer from the network policy setting if a user account is subject to both policies.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Deny log on as a batch job

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deny log on as a batch job**

**Description**
This security setting determines which accounts are prevented from being able to log on as a batch job. This policy setting supersedes the Log on as a batch job policy setting if a user account is subject to both policies.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Deny log on as a service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deny log on as a service**

**Description**
This security setting determines which service accounts are prevented from registering a process as a service. This policy setting supersedes the Log on as a service policy setting if an account is subject to both policies.

**Note**

- This security setting does not apply to the System, Local Service, or Network Service accounts.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Deny log on locally

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deny log on locally**

**Description**
This security setting determines which users are prevented from logging on at the computer. This policy setting supersedes the Allow log on locally policy setting if an account is subject to both policies.

**Important**

- If you apply this security policy to the Everyone group, no one will be able to log on locally.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Deny log on through Terminal Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deny log on through Terminal Services**

**Description**
This security setting determines which users and groups are prohibited from logging on as a Terminal Services client.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- This setting does not have any effect on Windows 2000 computers that have not been updated to Service Pack 2.

For more information, see:

- Allow log on through Terminal Services
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Enable computer and user accounts to be trusted for delegation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enable computer and user accounts to be trusted for delegation**

**Description**
This security setting determines which users can set the **Trusted for Delegation** setting on a user or computer object.

The user or object that is granted this privilege must have write access to the account control flags on the user or computer object. A server process running on a computer (or under a user context) that is trusted for delegation can access resources on another computer using delegated credentials of a client, as long as the client account does not have the **Account cannot be delegated** account control flag set.

This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.

**Caution**

- Misuse of this user right, or of the **Trusted for Delegation** setting, could make the network vulnerable to sophisticated attacks using Trojan horse programs that impersonate incoming clients and use their credentials to gain access to network resources.

**Default:** Administrators on domain controllers.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Allow a computer to be trusted for delegation
- Allow a user to be trusted for delegation
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Force shutdown from a remote system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Force shutdown from a remote system**

**Description**
This security setting determines which users are allowed to shut down a computer from a remote location on the network. Misuse of this user right can result in a denial of service.

This user right is defined in the Default Domain Controller Group Policy object (GPO) and in the local security policy of workstations and servers.

**Default:**

- On workstations and servers: Administrators.
- On domain controllers: Administrators, Server Operators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Shut down the system
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Generate security audits

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Generate security audits**

**Description**
This security setting determines which accounts can be used by a process to add entries to the security log. The security log is used to trace unauthorized system access. Misuse of this user right can result in the generation of many auditing events, potentially hiding evidence of an attack or causing a denial of service if the **Audit: Shut down system immediately if unable to log security audits** security policy setting is enabled. For more information about the Audit: Shut down system immediately if unable to log security audits security policy setting, see Audit: Shut down system immediately if unable to log security audits

**Default:** Local System.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Manage auditing and security log
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Impersonate a client after authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Impersonate a client after authentication**

**Description**
Assigning this privilege to a user allows programs running on behalf of that user to impersonate a client. Requiring this user right for this kind of impersonation prevents an unauthorized user from convincing a client to connect (for example, by remote procedure call (RPC) or named pipes) to a service that they have created and then impersonating that client, which can elevate the unauthorized user's permissions to administrative or system levels.

### Caution

- Assigning this user right can be a security risk. Only assign this user right to trusted users.

### Default:

- Administrators
- Service

### Note

- By default, services that are started by the Service Control Manager have the built-in Service group added to their access tokens. Component Object Model (COM) servers that are started by the COM infrastructure and that are configured to run under a specific account also have the Service group added to their access tokens. As a result, these services get this user right when they are started.

  In addition, a user can also impersonate an access token if any of the following conditions exist.

  - The access token that is being impersonated is for this user.
  - The user, in this logon session, created the access token by logging on to the network with explicit credentials.
  - The requested level is less than Impersonate, such as Anonymous or Identify.

  Because of these factors, users do not usually need this user right.

  For more information, search for "SeImpersonatePrivilege" in the Microsoft Platform SDK.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Increase scheduling priority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Increase scheduling priority**

**Description**
This security setting determines which accounts can use a process with Write Property access to another process to increase the execution priority assigned to the other process. A user with this privilege can change the scheduling priority of a process through the Task Manager user interface.

**Default:** Administrators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Load and unload device drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Load and unload device drivers**

**Description**
This user right determines which users can dynamically load and unload device drivers or other code in to kernel mode. This user right does not apply to Plug and Play device drivers. It is recommended that you do not assign this privilege to other users. Instead, use the `StartService()` API.

### Caution

- Assigning this user right can be a security risk. Do not assign this user right to any user, group, or process that you do not want to take over the system.

**Default:** Administrators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Lock pages in memory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Lock pages in memory**

**Description**
This security setting determines which accounts can use a process to keep data in physical memory, which prevents the system from paging the data to virtual memory on disk. Exercising this privilege could significantly affect system performance by decreasing the amount of available random access memory (RAM).

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Log on as a batch job

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Log on as a batch job**

**Description**
This security setting allows a user to be logged on by means of a batch-queue facility.

For example, when a user submits a job by means of the task scheduler, the task scheduler logs that user on as a batch user rather than as an interactive user.

**Note**

- In Windows 2000 Server, Windows 2000 Professional, Windows XP Professional, and Windows Server 2003, the Task Scheduler automatically grants this right as necessary. The Task Scheduler injects this right for the user into the computer's effective policy immediately after the task is scheduled. If there is a contradictory group policy defining the **Logon as a batch job** user right, the effective policy settings will be overwritten by the group policy only upon policy refresh.

**Default:** Local System.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Deny log on as a batch job
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Log on as a service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Log on as a service**

**Description**
This security setting determines which service accounts can register a process as a service.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Deny log on as a service
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Manage auditing and security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage auditing and security log**

**Description**
This security setting determines which users can specify object access auditing options for individual resources, such as files, Active Directory objects, and registry keys.

This security setting does not allow a user to enable file and object access auditing in general. For such auditing to be enabled, the Audit object access setting in Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policies must be configured.

You can view audited events in the security log of the Event Viewer. A user with this privilege can also view and clear the security log.

**Default:** Administrators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Modify firmware environment values

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modify firmware environment values**

**Description**
This security setting determines who can modify firmware environment values. Firmware environment variables are settings stored in the nonvolatile RAM of non-x86-based computers. The effect of the setting depends on the processor.

- On x86-based computers, the only firmware environment value that can be modified by assigning this user right is the Last Known Good Configuration setting, which should only be modified by the system.
- On x64-based computers, boot information is stored in nonvolatile RAM. Users must be assigned this user right to run **bootcfg.exe** and to change the Default Operating System setting on **Startup and Recovery** in **System Properties**.
- On all computers, this user right is required to install or upgrade Windows.

**Note**

- This security setting does not affect who can modify the system environment variables and user environment variables that are displayed on the **Advanced** tab of **System Properties**. For information about how to modify these variables, see Add or change the values of environment variables.

**Default:** Administrators, Local System.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Bootcfg
- Assign user rights for your local computer
- Security Configuration Manager tools

# Perform volume maintenance tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Perform volume maintenance tasks**

**Description**

This security setting determines which users and groups can run maintenance tasks on a volume, such as remote defragmentation.

Use caution when assigning this user right. Users with this user right can explore disks and extend files in to memory that contains other data. When the extended files are opened, the user might be able to read and modify the acquired data.

**Default:** Administrators

**Configuring this security setting**

You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignments\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Profile single process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Profile single process**

**Description**
This security setting determines which users can use performance monitoring tools to monitor the performance of nonsystem processes.

**Default:** Administrators, Power users, Local System.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Profile system performance
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Profile system performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Profile system performance**

**Description**
This security setting determines which users can use performance monitoring tools to monitor the performance of system processes.

**Default:** Administrators, Local System.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Profile single process
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Remove computer from docking station

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remove computer from docking station**

**Description**
This security setting determines whether a user can undock a portable computer from its docking station without logging on.

If this policy is enabled, the user must log on before removing the portable computer from its docking station. If this policy is disabled, the user may remove the portable computer from its docking station without logging on.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Replace a process level token

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Replace a process level token**

**Description**
This security setting determines which user accounts can call the `CreateProcessAsUser()` application programming interface (API) so that one service can start another. An example of a process that uses this user right is Task Scheduler. For information about Task Scheduler, see Task Scheduler overview.

**Default:** Network Service, Local System.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Restore files and directories

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restore files and directories**

**Description**
This security setting determines which users can bypass file, directory, registry, and other persistent objects permissions when restoring backed up files and directories, and determines which users can set any valid security principal as the owner of an object.

Specifically, this user right is similar to granting the following permissions to the user or group in question on all files and folders on the system:

- Traverse Folder/Execute File
- Write

### Caution

- Assigning this user right can be a security risk. Since users with this user right can overwrite registry settings, hide data, and gain ownership of system objects, only assign this user right to trusted users.

### Default:

- Workstations and servers: Administrators, Backup Operators.
- Domain controllers: Administrators, Backup Operators, Server Operators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Back up files and directories
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Shut down the system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shut down the system**

**Description**
This security setting determines which users who are logged on locally to the computer can shut down the operating system using the **Shut Down** command. Misuse of this user right can result in a denial of service.

**Default:**

- Workstations: Administrators, Backup Operators, Power Users, Users.
- Servers: Administrators, Backup Operators, Power Users.
- Domain controllers: Administrators, Backup Operators, Server Operators, Print Operators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Force shutdown from a remote system
- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Synchronize directory service data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Synchronize directory service data**

**Description**
This security setting determines which users and groups have the authority to synchronize all directory service data. This is also known as Active Directory synchronization.

**Defaults:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Take ownership of files or other objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Take ownership of files or other objects**

**Description**
This security setting determines which users can take ownership of any securable object in the system, including Active Directory objects, files and folders, printers, registry keys, processes, and threads.

**Caution**

- Assigning this user right can be a security risk. Since owners of objects have full control of them, only assign this user right to trusted users.

**Default:** Administrators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- User Rights Assignment
- Assign user rights for your local computer
- Security Configuration Manager tools

# Security Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security options**
This section covers:

- Accounts: Administrator account status
- Accounts: Guest account status
- Accounts: Limit local account use of blank passwords to console logon only
- Accounts: Rename administrator account
- Accounts: Rename guest account
- Audit: Audit the access of global system objects
- Audit: Audit the use of Backup and Restore privilege
- Audit: Shut down system immediately if unable to log security audits
- Devices: Allow undock without having to log on
- Devices: Allowed to format and eject removable media
- Devices: Prevent users from installing printer drivers
- Devices: Restrict CD-ROM access to locally logged-on user only
- Devices: Restrict floppy access to locally logged-on user only
- Devices: Unsigned driver installation behavior
- Domain controller: Allow server operators to schedule tasks
- Domain controller: LDAP server signing requirements
- Domain controller: Refuse machine account password changes
- Domain member: Digitally encrypt or sign secure channel data (always)
- Domain member: Digitally encrypt secure channel data (when possible)
- Domain member: Digitally sign secure channel data (when possible)
- Domain member: Maximum machine account password age
- Domain member: Require strong (Windows 2000 or later) session key
- Domain member: Disable machine account password changes
- Interactive logon: Do not display last user name

- Interactive logon: Do not require CTRL+ALT+DEL
- Interactive logon: Message text for users attempting to log on
- Interactive logon: Message title for users attempting to log on
- Interactive logon: Number of previous logons to cache (in case domain controller is not available)
- Interactive logon: Prompt user to change password before expiration
- Interactive logon: Require Domain Controller authentication to unlock
- Interactive logon: Require smart card
- Interactive logon: Smart card removal behavior
- Microsoft network client: Digitally sign communications (always)
- Microsoft network client: Digitally sign communications (if server agrees)
- Microsoft network client: Send unencrypted password to connect to third-party SMB servers
- Microsoft network server: Amount of idle time required before suspending session
- Microsoft network server: Digitally sign communications (always)
- Microsoft network server: Digitally sign communications (if client agrees)
- Microsoft network server: Disconnect clients when logon hours expire
- Network access: Allow anonymous SID/Name translation
- Network access: Do not allow anonymous enumeration of SAM accounts
- Network access: Do not allow anonymous enumeration of SAM accounts and shares
- Network access: Do not allow storage of credentials or .NET Passports for network authentication
- Network access: Let Everyone permissions apply to anonymous users
- Network access: Named pipes that can be accessed anonymously
- Network access: Remotely accessible registry paths
- Network access: Remotely accessible registry paths and subpaths
- Network access: Restrict anonymous access to Named Pipes and Shares
- Network access: Shares that can be accessed anonymously
- Network access: Sharing and security model for local accounts
- Network security: Do not store LAN Manager hash value on next password change

- Network security: Force logoff when logon hours expire
- Network security: LAN Manager authentication level
- Network security: LDAP client signing requirements
- Network security: Minimum session security for NTLM SSP based (including secure RPC) clients
- Network security: Minimum session security for NTLM SSP based (including secure RPC) servers
- Recovery console: Allow automatic administrative logon
- Recovery console: Allow floppy copy and access to all drives and all folders
- Shutdown: Allow system to be shut down without having to log on
- Shutdown: Clear virtual memory pagefile
- System cryptography: Force strong key protection for user keys stored on the computer
- System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing
- System objects: Default owner for objects created by members of the Administrators group
- System objects: Require case insensitivity for non-Windows subsystems
- System objects: Strengthen default permissions of internal system objects (e.g., Symbolic Links)
- System settings: Optional subsystems
- System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies

# Accounts: Administrator account status

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accounts: Administrator account status**

**Description**
This security setting determines whether the local Administrator account is enabled or disabled.

### Notes

- If you try to reenable the Administrator account after it has been disabled, and if the current Administrator password does not meet the password requirements, you cannot reenable the account. In this case, an alternative member of the Administrators group must reset the password on the Administrator account. For information about how to reset a password, see Reset a password.

- Disabling the Administrator account can become a maintenance issue under certain circumstances. For example, in a domain environment, if the secure channel that constitutes your join fails for any reason, and there is no other local Administrator account, you must restart in Safe Mode to fix the problem that is causing your join status to be broken.

- Under Safe Mode boot, the Administrator account is always enabled, regardless of this setting.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Options

- Security Configuration Manager tools

# Accounts: Guest account status

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accounts: Guest account status**

**Description**
This security setting determines if the Guest account is enabled or disabled.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Note**

- If the Guest account is disabled and the security option **Network Access: Sharing and Security Model for local accounts** is set to **Guest Only**, network logons, such as those performed by the Microsoft Network Server (SMB Service), will fail.

For more information, see:

- Network access: Sharing and security model for local accounts
- Security Configuration Manager tools

# Accounts: Limit local account use of blank passwords to console logon only

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accounts: Limit local account use of blank passwords to console logon only**

**Description**
This security setting determines whether local accounts that are not password protected can be used to logon from locations other than the physical computer console. If enabled, then local accounts that are not password protected will only be able to log on at the computer's keyboard. This security setting does not apply to guest accounts.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

### Caution

- Computers that are not in physically secure locations should always enforce strong password policies for all local user accounts. Otherwise, anyone with physical access to the computer can log on using a user account that does not have a password. This is especially important for portable computers.
- If you apply this security policy to the Everyone group, no one will be able to log on through terminal services.

### Notes

- This setting does not affect logons that use domain accounts.
- It is possible for applications that use remote interactive logons to bypass this setting.

For more information, see:

- Security Configuration Manager tools

# Accounts: Rename administrator account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accounts: Rename administrator account**

**Description**
This security setting determines whether a different account name is associated with the security identifier (SID) for the account **Administrator**. Because the Administrator account is known to exist on all computers running Windows 2000 Server, Windows 2000 Professional, Windows XP Professional, and the Windows Server 2003 family, renaming the account makes it slightly more difficult for unauthorized persons to guess this privileged user name and password combination.

**Default:** Administrator.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Accounts: Rename guest account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accounts: Rename guest account**

**Description**
This security setting determines whether a different account name is associated with the security identifier (SID) for the account "Guest." Because the Guest account is known to exist on all computers running Windows 2000 Server, Windows 2000 Professional, Windows XP Professional, and the Windows Server 2003 family, renaming the account makes it slightly more difficult for unauthorized persons to guess this user name and password combination.

**Default:** Guest.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Audit: Audit the access of global system objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Audit: Audit the access of global system objects**

**Description**
This security setting determines whether to audit the access of global system objects.

If this policy is enabled, it causes system objects, such as mutexes (mutual exclusive), events, semaphores (locking mechanisms used inside resource managers or resource dispensers), and DOS devices, to be created with a default system access control list (SACL). If the Audit object access audit policy is also enabled, access to these system objects is audited.

**Note**

- When configuring this security setting, changes will not take effect until you restart Windows.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Audit: Audit the use of Backup and Restore privilege

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Audit: Audit the use of Backup and Restore privilege**

**Description**
This security setting determines whether to audit the use of all user privileges, including Backup and Restore, when the Audit privilege use policy is in effect. Enabling this option when the **Audit privilege use** policy is also enabled generates an audit event for every file that is backed up or restored.

If you disable this policy, then use of the Backup or Restore privilege is not audited even when **Audit privilege use** is enabled.

**Note**

- When configuring this security setting, changes will not take effect until you restart Windows.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Audit: Shut down system immediately if unable to log security audits

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Audit: Shut down system immediately if unable to log security audits**

**Description**
This security setting determines whether the system shuts down if it is unable to log security events.

If this security setting is enabled, it causes the system to stop if a security audit cannot be logged for any reason. Typically, an event fails to be logged when the security audit log is full and the retention method that is specified for the security log is either **Do Not Overwrite Events** or **Overwrite Events by Days**.

If the security log is full and an existing entry cannot be overwritten, and this security option is enabled, the following Stop error appears:

**STOP: C0000244 {Audit Failed}**

**An attempt to generate a security audit failed.**

To recover, an administrator must log on, archive the log (optional), clear the log, and reset this option as desired. Until this security setting is reset, no users, other than a member of the Administrators group will be able to log on to the system, even if the security log is not full.

**Note**

- When configuring this security setting, changes will not take effect until you restart Windows.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Auditing Security Events
- Security Configuration Manager tools

# Devices: Allow undock without having to log on

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devices: Allow undock without having to log on**

**Description**
This security setting determines whether a portable computer can be undocked without having to log on. If this policy is enabled, logon is not required and an external hardware eject button can be used to undock the computer. If disabled, a user must log on and have the **Remove computer from docking station** privilege to undock the computer.

**Default**: Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Caution**

- Disabling this policy may tempt users to try and physically remove the laptop from its docking station using methods other than the external hardware eject button. Since this may cause damage to the hardware, this setting, in general, should only be disabled on laptop configurations that are physically securable.

For more information, see:

- Security Configuration Manager tools
- Remove computer from docking station

# Devices: Allowed to format and eject removable media

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devices: Allowed to format and eject removable media**

**Description**
This security setting determines who is allowed to format and eject removable NTFS media. This capability can be given to:

- Administrators
- Administrators and Power Users
- Administrators and Interactive Users

**Default**: Administrators.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Devices: Prevent users from installing printer drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devices: Prevent users from installing printer drivers**

**Description**
For a computer to print to a network printer, the driver for that network printer must be installed on the local computer. This security setting determines who is allowed to install a printer driver as part of adding a network printer. If this setting is enabled, only Administrators and Power Users can install a printer driver as part of adding a network printer. If this setting is disabled, any user can install a printer driver as part of adding a network printer.

### Default:

- Enabled on servers.
- Disabled on workstations.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

### Notes

- This setting does not affect the ability to add a local printer.
- This setting does not affect Power Users or Administrators.

For more information, see:

- Security Configuration Manager tools

# Devices: Restrict CD-ROM access to locally logged-on user only

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devices: Restrict CD-ROM access to locally logged-on user only**

**Description**
This security setting determines whether a CD-ROM is accessible to both local and remote users simultaneously.

If this policy is enabled, it allows only the interactively logged-on user to access removable CD-ROM media. If this policy is enabled and no one is logged on interactively, the CD-ROM can be accessed over the network.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Devices: Restrict floppy access to locally logged-on user only

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devices: Restrict floppy access to locally logged-on user only**

**Description**
This security setting determines whether removable floppy media are accessible to both local and remote users simultaneously.

If this policy is enabled, it allows only the interactively logged-on user to access removable floppy media. If this policy is enabled and no one is logged on interactively, the floppy can be accessed over the network.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Devices: Unsigned driver installation behavior

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devices: Unsigned driver installation behavior**

**Description**
This security setting determines what happens when an attempt is made to install a device driver (by means of Setup API) that has not been tested by the Windows Hardware Quality Lab (WHQL).

The options are:

- **Silently succeed**
- **Warn but allow installation**
- **Do not allow installation**

**Default: Warn but allow installation**.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Domain controller: Allow server operators to schedule tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain controller: Allow server operators to schedule tasks**

**Description**
This security setting determines if Server Operators are allowed to submit jobs by means of the AT schedule facility.

### Note

- This security setting only affects the AT schedule facility; it does not affect the Task Scheduler facility.

**Default:** This policy is not defined, which means that the system treats it as disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Domain controller: LDAP server signing requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain controller: LDAP server signing requirements**

**Description**
This security setting determines whether the LDAP server requires signing to be negotiated with LDAP clients, as follows:

- **None:** Data signing is not required in order to bind with the server. If the client requests data signing, the server supports it.

- **Require signature:** Unless TLS\SSL is being used, the LDAP data signing option must be negotiated.

**Default:** Not defined, which has the same effect as **None**.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Caution**

- If you set the server to **Require Signature,** you must also set the client. Not setting the client results in loss of connection with the server.

**Notes**

- This setting does not have any impact on ldap_simple_bind or ldap_simple_bind_s. No Microsoft LDAP clients that are shipped with Windows XP Professional use ldap_simple_bind or ldap_simple_bind_s to talk to a domain controller.

- If signing is required, then ldap_simple_bind and ldap_simple_bind_s requests are rejected. No Microsoft LDAP clients running Windows XP Professional or the Windows Server 2003 family use ldap_simple_bind or ldap_simple_bind_s to bind to directory service.

For more information, see:

- Security Configuration Manager tools

# Domain controller: Refuse machine account password changes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain controller: Refuse machine account password changes**

## Description

This security setting determines whether domain controllers will refuse requests from member computers to change computer account passwords. By default, member computers change their computer account passwords every 30 days. If enabled, the domain controller will refuse computer account password change requests.

If it is enabled, this setting does not allow a domain controller to accept *any* changes to a computer account's password.

**Default:** Disabled.

**Configuring this security setting**

You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Domain member: Maximum machine account password age
- Domain member: Disable machine account password changes
- Security Configuration Manager tools

# Domain member: Digitally encrypt or sign secure channel data (always)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain member: Digitally encrypt or sign secure channel data (always)**

**Description**
This security setting determines whether all secure channel traffic initiated by the domain member must be signed or encrypted.

When a computer joins a domain, a computer account is created. After that, when the system starts, it uses the computer account password to create a secure channel with a domain controller for its domain. This secure channel is used to perform operations such as NTLM pass through authentication, LSA SID\Name Lookup etc.

This setting determines whether or not all secure channel traffic initiated by the domain member meets minimum security requirements. Specifically it determines whether all secure channel traffic initiated by the domain member must be signed or encrypted. If this policy is enabled, then the secure channel will not be established unless either signing or encryption of all secure channel traffic is negotiated. If this policy is disabled, then encryption and signing of all secure channel traffic is negotiated with the Domain Controller in which case the level of signing and encryption depends on the version of the Domain Controller and the settings of the following two policies:

- Domain member: Digitally encrypt secure channel data (when possible)
- Domain member: Digitally sign secure channel data (when possible)

**Default**: Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- In order to take advantage of this policy on member workstations and servers, all domain controllers that constitute the member's domain must be running Windows NT 4.0 Service Pack 6 or higher.
- In order to take advantage of this policy on domain controllers, *All* domain controllers in the same domain, as well as all trusted domains, must be running Windows NT 4.0 Service Pack 6 or higher.

**Notes**

- If this policy is enabled, the policy **Domain member: Digitally sign secure channel data (when possible)** is assumed to be enabled regardless of its current setting. This ensures that the domain member attempts to negotiate at least signing of the secure channel traffic.
- If this policy is enabled, the policy **Domain member: Digitally sign secure channel data (when possible)** is assumed to be enabled regardless of its current setting. This ensures that the domain member attempts to negotiate at least signing of the secure channel traffic.
- Logon information transmitted over the secure channel is always encrypted regardless of whether encryption of ALL other secure channel traffic is negotiated or not.

For more information, see:

- Domain member: Digitally encrypt secure channel data (when possible)
- Domain member: Digitally sign secure channel data (when possible)
- Domain member: Require strong (Windows 2000 or later) session key
- Security Configuration Manager tools

# Domain member: Digitally encrypt secure channel data (when possible)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain member: Digitally encrypt secure channel data (when possible)**

**Description**
This security setting determines whether a domain member attempts to negotiate encryption for all secure channel traffic that it initiates.

When a computer joins a domain, a computer account is created. After that, when the system starts, it uses the computer account password to create a secure channel with a domain controller for its domain. This secure channel is used to perform operations such as NTLM passthrough authentication, LSA SID\Name Lookup etc.

This setting determines whether or not the domain member attempts to negotiate encryption for all secure channel traffic that it initiates. If enabled, the domain member will request encryption of all secure channel traffic. If the domain controller supports encryption of all secure channel traffic (Windows NT 4.0 Service Pack 6 or higher), then all secure channel traffic will be encrypted. Otherwise only logon information transmitted over the secure channel will be encrypted. If this setting is disabled, then the domain member will not attempt to negotiate secure channel encryption.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- There is no known reason for disabling this setting. Besides unnecessarily reducing the potential confidentiality level of the secure channel, disabling this setting may unnecessarily reduce secure channel throughput, because concurrent API calls that use the secure channel are only possible when the secure channel is signed or encrypted.

**Note**

- Domain controllers are also domain members and establish secure channels with other domain controllers in the same domain as well as domain controllers in trusted domains.

For more information, see:

- Domain member: Digitally encrypt or sign secure channel data (always)
- Domain member: Digitally sign secure channel data (when possible)
- Domain member: Require strong (Windows 2000 or later) session key
- Security Configuration Manager tools

# Domain member: Digitally sign secure channel data (when possible)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain member: Digitally sign secure channel data (when possible)**

**Description**
This security setting determines whether a domain member attempts to negotiate signing for all secure channel traffic that it initiates.

When a computer joins a domain, a computer account is created. After that, when the system starts, it uses the computer account password to create a secure channel with a domain controller for its domain. This secure channel is used to perform operations such as NTLM pass through authentication, LSA SID\Name Lookup etc.

This setting determines whether or not the domain member attempts to negotiate signing for all secure channel traffic that it initiates. If enabled, the domain member will request signing of all secure channel traffic. If the Domain Controller supports signing of all secure channel traffic (Windows NT 4.0 Service Pack 6 or higher), then all secure channel traffic will be signed which ensures that it cannot be tampered with in transit.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Notes**

- If the policy **Domain member: Digitally encrypt or sign secure channel data (always) is enabled**, then this policy is assumed to be enabled regardless of its current setting.
- Domain controllers are also domain members and establish secure channels with other domain controllers in the same domain as well as domain controllers in trusted domains.

For more information, see:

- Domain member: Digitally encrypt or sign secure channel data (always)
- Domain member: Digitally encrypt secure channel data (when possible)
- Domain member: Require strong (Windows 2000 or later) session key
- Security Configuration Manager tools

# Domain member: Maximum machine account password age

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain member: Maximum machine account password age**

**Description**
This security setting determines how often a domain member will attempt to change its computer account password.

**Default:** 30 days.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- This setting applies to Windows 2000 computers, but it is not available through the Security Configuration Manager tools on these computers.

For more information, see:

- Domain controller: Refuse machine account password changes
- Domain member: Disable machine account password changes
- Security Configuration Manager tools

# Domain member: Require strong (Windows 2000 or later) session key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain member: Require strong (Windows 2000 or later) session key**

**Description**
This security setting determines whether 128-bit key strength is required for encrypted secure channel data.

When a computer joins a domain, a computer account is created. After that, when the system starts, it uses the computer account password to create a secure channel with a domain controller within the domain. This secure channel is used to perform operations such as NTLM passthrough authentication, LSA SID\Name Lookup, and so on.

Based on what version of Windows is running on the domain controller that the domain member is communicating with and the settings of the following parameters:

- Domain member: Digitally encrypt or sign secure channel data (always)
- Domain member: Digitally encrypt secure channel data (when possible)

some or all of the information that is transmitted over the secure channel will be encrypted. This policy setting determines whether or not 128-bit key strength is required for the secure channel information that is encrypted.

If this setting is enabled, then the secure channel will not be established unless 128-bit encryption can be performed. If this setting is disabled, then the key strength is negotiated with the domain controller.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- In order to take advantage of this policy on member workstations and servers, all domain controllers that constitute the member's domain must be running Windows 2000 or later.
- In order to take advantage of this policy on domain controllers, all domain controllers in the same domain as well as all trusted domains must run Windows 2000 or later.

For more information, see:

- Domain member: Digitally encrypt or sign secure channel data (always)
- Domain member: Digitally encrypt secure channel data (when possible)
- Domain member: Digitally sign secure channel data (when possible)
- Security Configuration Manager tools

# Domain member: Disable machine account password changes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain member: Disable machine account password changes**

**Description**
Determines whether a domain member periodically changes its computer account password. If this setting is enabled, the domain member does not attempt to change its computer account password. If this setting is disabled, the domain member attempts to change its computer account password as specified by the setting for **Domain Member: Maximum age for machine account password**, which by default is every 30 days.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

### Notes

- This security setting should not be enabled. Computer account passwords are used to establish secure channel communications between members and domain controllers and, within the domain, between the domain controllers themselves. Once it is established, the secure channel is used to transmit sensitive information that is necessary for making authentication and authorization decisions.

- This setting should not be used in an attempt to support dual-boot scenarios that use the same computer account. If you want to dual-boot two installations that are joined to the same domain, give the two installations different computer names.

For more information, see:

- Domain controller: Refuse machine account password changes

- Domain member: Maximum machine account password age

- Security Configuration Manager tools

# Interactive logon: Do not display last user name

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive logon: Do not display last user name**

**Description**
This security setting determines whether the name of the last user to log on to the computer is displayed in the Windows logon screen.

If this policy is enabled, the name of the last user to successfully log on is not displayed in the **Log On to Windows** dialog box.

If this policy is disabled, the name of the last user to log on is displayed.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Interactive logon: Do not require CTRL+ALT+DEL

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive logon: Do not require CTRL+ALT+DEL**

**Description**
This security setting determines whether pressing CTRL+ALT+DEL is required before a user can log on.

If this policy is enabled on a computer, a user is not required to press CTRL+ALT+DEL to log on. Not having to press CTRL+ALT+DEL leaves users susceptible to attacks that attempt to intercept the users' passwords. Requiring CTRL+ALT+DEL before users log on ensures that users are communicating by means of a trusted path when entering their passwords.

If this policy is disabled, any user is required to press CTRL+ALT+DEL before logging on to Windows (unless they are using a smart card for Windows logon).

**Default:**

- Disabled on workstations and servers that are joined to a domain.
- Enabled on stand-alone workstations.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Interactive logon: Message text for users attempting to log on

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive logon: Message text for users attempting to log on**

**Description**
This security setting specifies a text message that is displayed to users when they log on.

This text is often used for legal reasons, for example, to warn users about the ramifications of misusing company information or to warn them that their actions may be audited.

**Default**: No message.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree to: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\.

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Caution**

- Windows XP and Windows Server 2003 add support for configuring logon text that can exceed 512 characters in length and that can also contain carriage return/line feed (CR/LF) sequences. If you do not use CR/LF sequences in your logon text, the maximum number of characters that you can add is 512. If you use CR/LF sequences, you can use a maximum of four lines of 512 characters each for a total of 2,048 characters. However, Windows 2000–based client computers cannot interpret and display logon text that is created by computers running Windows XP or Windows Server 2003. You must use a Windows 2000–based computer to create a logon text policy that applies to Windows 2000–based computers. If you inadvertently create a logon text policy by using a computer running Windows XP Professional or Windows Server 2003, and you discover that it does not display properly on Windows 2000–based computers, do the following:

    - Undefine the setting.

    - Redefine the setting by using a Windows 2000–based computer.

      You cannot simply change a Windows XP Professional or a Windows Server 2003–defined logon message policy by using a Windows 2000–based computer. The setting must be undefined first.

For more information, see:

- Interactive logon: Message title for users attempting to log on

- Security Configuration Manager tools

# Interactive logon: Message title for users attempting to log on

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive logon: Message title for users attempting to log on**

**Description**
This security setting allows the specification of a title to appear in the title bar of the window that contains the Interactive logon: Message text for users attempting to log on.

**Default**: No message.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Interactive logon: Number of previous logons to cache (in case domain controller is not available)

Updated: June 17, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive logon: Number of previous logons to cache (in case domain controller is not available)**

**Description**
Determines the number of users who can have cached credentials on the computer.

All previous users' logon information is cached locally so that, in the event that a domain controller is unavailable during subsequent logon attempts, they are able to log on . If a domain controller is unavailable and a user's logon information is cached, the user is prompted with a message that reads as follows:

*Windows cannot connect to a server to confirm your logon settings. You have been logged on using previously stored account information. If you changed your account information since you last logged on to this computer, those changes will not be reflected in this session.*

If a domain controller is unavailable and a user's logon information is not cached, the user is prompted with this message:

*The system cannot log you on now because the domain <DOMAIN_NAME> is not available.*

In this policy setting, a value of 0 disables logon caching. Any value above 50 only caches 50 logon attempts.

**Default:** 10

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Interactive logon: Prompt user to change password before expiration

Updated: June 29, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive logon: Prompt user to change password before expiration**

**Description**
Determines how far in advance (in days) users are warned that their password is about to expire. With this advance warning, the user has time to construct a password that is sufficiently strong.

**Default:** 14 days.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Interactive logon: Require Domain Controller authentication to unlock

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive logon: Require Domain Controller authentication to unlock**

**Description**
Logon information must be provided to unlock a locked computer. For domain accounts, this security setting determines whether a domain controller must be contacted to unlock a computer. If this setting is disabled, a user can unlock the computer using cached credentials. If this setting is enabled, a domain controller must authenticate the domain account that is being used to unlock the computer.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- This setting applies to Windows 2000 computers, but it is not available through the Security Configuration Manager tools on these computers.

For more information, see:

- Security Configuration Manager tools

# Interactive logon: Require smart card

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive logon: Require smart card**

**Description**
This security setting requires users to log on to a computer using a smart card.

The options are:

- **Enabled**. Users can only log on to the computer using a smart card.
- **Disabled**. Users can log on to the computer using any method.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- This setting will apply to any computers running Windows 2000 through changes in the registry, but the security setting is not viewable through the Security Configuration Manager tool set.

For more information, see:

- Security Configuration Manager tools
- Smart Cards

# Interactive logon: Smart card removal behavior

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interactive logon: Smart card removal behavior**

**Description**
This security setting determines what happens when the smart card for a logged-on user is removed from the smart card reader.

The options are:

- **No Action**
- **Lock Workstation**
- **Force Logoff**
- **Disconnect if a remote Terminal Services session**

If you click **Lock Workstation** in the **Properties** dialog box for this policy, the workstation is locked when the smart card is removed, allowing users to leave the area, take their smart card with them, and still maintain a protected session.

If you click **Force Logoff** in the **Properties** dialog box for this policy, the user is automatically logged off when the smart card is removed.

If you click **Disconnect if a remote Terminal Services session**, removal of the smart card disconnects the session without logging the user off. This allows the user to insert the smart card and resume the session later, or at another smart card reader-equipped terminal, without having to log on again.

**Default:** No action specified.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Microsoft network client: Digitally sign communications (always)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft network client: Digitally sign communications (always)**

**Description**
This security setting determines whether packet signing is required by the SMB client component.

The server message block (SMB) protocol provides the basis for Microsoft file and print sharing and many other networking operations, such as remote Windows administration. To prevent man-in-the-middle attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. This policy setting determines whether SMB packet signing must be negotiated before further communication with an SMB server is permitted.

If this setting is enabled, the Microsoft network client will not communicate with a Microsoft network server unless that server agrees to perform SMB packet signing. If this policy is disabled, SMB packet signing is negotiated between the client and server.

**Default**: Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- For this policy to take effect on computers running Windows 2000, client-side packet signing must also be enabled. To enable client-side SMB packet signing, set **Microsoft network client: Digitally sign communications (if server agrees)**.
- Computers that have this policy set will not be able to communicate with computers that do not have server-side packet signing enabled. By default, server-side packet signing is enabled only on domain controllers running Windows 2000 and later.
- Server-side packet signing can be enabled on computers running Windows 2000 and later by setting **Microsoft network server: Digitally sign communications (if client agrees)**
- Server-side packet signing can be enabled on computers running Windows NT 4.0 Service Pack 3 and later by setting the following registry value to 1:

  **HKLM\System\CurrentControlSet\Services\LanManServer\Parameters\EnableSecuritySignature**

- Server-side packet signing cannot be enabled on computers running Windows 95 or Windows 98.

**Notes**

- All Windows operating systems support both a client-side SMB component and a server-side SMB component. To take advantage of SMB packet signing, both the client-side SMB component and server-side SMB component that are involved in a communication must have SMB packet signing either enabled or required. On Windows 2000 and later operating systems, enabling or requiring packet signing for client and server-side SMB components is controlled by the following four policy settings:
  - **Microsoft network client: Digitally sign communications (always)** - Controls whether or not the client-side SMB component requires packet signing.

    - **Microsoft network client: Digitally sign communications (if server agrees)** - Controls whether or not the client-side SMB component has packet signing enabled.

    - **Microsoft network server: Digitally sign communications (always)** - Controls whether or not the server-side SMB component requires packet signing.

    - **Microsoft network server: Digitally sign communications (if client agrees)** - Controls whether or not the server-side SMB component has packet signing enabled.

- If server-side SMB signing is required, a client will not be able to establish a session with that server, unless it has client-side SMB signing enabled. By default, client-side SMB signing is enabled on workstations, servers, and domain controllers. Similarly, if client-side SMB signing is required, that client will not be able to establish a session with servers that do not have packet signing enabled. By default, server-side SMB signing is enabled only on domain controllers.

- If server-side SMB signing is enabled, SMB packet signing will be negotiated with clients that have client-side SMB signing enabled.

- Using SMB packet signing can impose up to a 15 percent performance hit on file service transactions.

For more information, see:

- Microsoft network client: Digitally sign communications (if server agrees)

- Microsoft network server: Digitally sign communications (if client agrees)

- Microsoft network server: Digitally sign communications (always)

- Security Configuration Manager tools

# Microsoft network client: Digitally sign communications (if server agrees)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft network client: Digitally sign communications (if server agrees)**

**Description**
This security setting determines whether the SMB client attempts to negotiate SMB packet signing.

The server message block (SMB) protocol provides the basis for Microsoft file and print sharing and many other networking operations, such as remote Windows administration. To prevent man-in-the-middle attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. This policy setting determines whether the SMB client component attempts to negotiate SMB packet signing when it connects to an SMB server.

If this setting is enabled, the Microsoft network client will ask the server to perform SMB packet signing upon session setup. If packet signing has been enabled on the server, packet signing will be negotiated. If this policy is disabled, the SMB client will never negotiate SMB packet signing.

**Default**: Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

### Notes

- All Windows operating systems support both a client-side SMB component and a server-side SMB component. To take advantage of SMB packet signing, both the client-side SMB component and server-side SMB component that are involved in a communication must have SMB packet signing either enabled or required. On Windows 2000 and later, enabling or requiring packet signing for client and server-side SMB components is controlled by the following four policy settings:

    - **Microsoft network client: Digitally sign communications (always)** - Controls whether or not the client-side SMB component requires packet signing.

    - **Microsoft network client: Digitally sign communications (if server agrees)** - Controls whether or not the client-side SMB component has packet signing enabled.

    - **Microsoft network server: Digitally sign communications (always)** - Controls whether or not the server-side SMB component requires packet signing.

    - **Microsoft network server: Digitally sign communications (if client agrees)** - Controls whether or not the server-side SMB component has packet signing enabled.

- If server-side SMB signing is required, a client will not be able to establish a session with that server unless it has client-side SMB signing enabled. By default, client-side SMB signing is enabled on workstations, servers, and domain controllers.

- Similarly, if client-side SMB signing is required, that client will not be able to establish a session with servers that do not have packet signing enabled. By default, server-side SMB signing is enabled only on domain controllers.

- If server-side SMB signing is enabled, SMB packet signing will be negotiated with clients that have client-side SMB signing enabled.

- Using SMB packet signing can degrade performance up to 15 percent on file service transactions.

For more information, see:

- Microsoft network client: Digitally sign communications (always)
- Microsoft network server: Digitally sign communications (if client agrees)
- Microsoft network server: Digitally sign communications (always)
- Security Configuration Manager tools

# Microsoft network client: Send unencrypted password to connect to third-party SMB servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft network client: Send unencrypted password to connect to third-party SMB servers**

**Description**
If this security setting is enabled, the Server Message Block (SMB) redirector is allowed to send plaintext passwords to non-Microsoft SMB servers that do not support password encryption during authentication.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Microsoft network server: Amount of idle time required before suspending session

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft network server: Amount of idle time required before suspending a session**

**Description**
This security setting determines the amount of continuous idle time that must pass in a Server Message Block (SMB) session before the session is suspended due to inactivity.

Administrators can use this policy to control when a computer suspends an inactive SMB session. If client activity resumes, the session is automatically reestablished.

For this policy setting, a value of 0 means to disconnect an idle session as quickly as is reasonably possible. The maximum value is 99999, which is 208 days; in effect, this value disables the policy.

**Default:**

- 15 minutes for servers.
- Undefined for workstations.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Microsoft network server: Digitally sign communications (always)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft network server: Digitally sign communications (always)**

**Description**
This security setting determines whether packet signing is required by the SMB server component.

The server message block (SMB) protocol provides the basis for Microsoft file and print sharing and many other networking operations, such as remote Windows administration. To prevent "man-in-the-middle" attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. This policy setting determines whether SMB packet signing must be negotiated before further communication with an SMB client is permitted.

If this setting is enabled, the Microsoft network server will not communicate with a Microsoft network client unless that client agrees to perform SMB packet signing. If this setting is disabled, SMB packet signing is negotiated between the client and server.

**Default**:

- Disabled for member servers.
- Enabled for domain controllers.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Notes**

- All Windows operating systems support both a client-side SMB component and a server-side SMB component. To take advantage of SMB packet signing, both the client-side SMB component and server-side SMB component that are involved in a communication must have SMB packet signing either enabled or required. On Windows 2000 and later, enabling or requiring packet signing for client and server-side SMB components is controlled by the following four policy settings:
    - Microsoft network client: Digitally sign communications (always) - Controls whether or not the client-side SMB component requires packet signing.
    - Microsoft network client: Digitally sign communications (if server agrees) - Controls whether or not the client-side SMB component has packet signing enabled.
    - Microsoft network server: Digitally sign communications (always) - Controls whether or not the server-side SMB component requires packet signing.
    - Microsoft network server: Digitally sign communications (if client agrees) - Controls whether or not the server-side SMB component has packet signing enabled.
- If server-side SMB signing is required, a client will not be able to establish a session with that server unless it has client-side SMB signing enabled. By default, client-side SMB signing is enabled on workstations, servers, and domain controllers.
- Similarly, if client-side SMB signing is required, that client will not be able to establish a session with servers that do not have packet signing enabled. By default, server-side SMB signing is enabled only on domain controllers.

- If server-side SMB signing is enabled, SMB packet signing will be negotiated with clients that have client-side SMB signing enabled.

- Using SMB packet signing can degrade performance up to 15 percent on file service transactions.

**Important**

- For this policy to take effect on computers running Windows 2000, server-side packet signing must also be enabled. To enable server-side SMB packet signing, set the following policy:

  Microsoft network server: Digitally sign communications (if server agrees)

- For servers running Windows 2000 Server or Windows Server 2003 to negotiate signing with client computers running Windows NT 4.0 or Windows 98, the following registry value must be set to 1 on the server:

  HKLM\System\CurrentControlSet\Services\lanmanserver\parameters\enableW9xsecuritysignature

- Computers that have this policy set will not communicate with computers that do not have client-side packet signing enabled. Client-side packet signing can be enabled on computers running Windows 2000 and later by setting the following policy:

For more information, see:

- Microsoft network client: Digitally sign communications (always)

- Microsoft network client: Digitally sign communications (if server agrees)

- Microsoft network server: Digitally sign communications (if client agrees)

- Security Configuration Manager tools

# Microsoft network server: Digitally sign communications (if client agrees)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft network server: Digitally sign communications (if client agrees)**

**Description**
This security setting determines whether the SMB server will negotiate SMB packet signing with clients that request it.

The server message block (SMB) protocol provides the basis for Microsoft file and print sharing and many other networking operations, such as remote Windows administration. To prevent man-in-the-middle attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. This policy setting determines whether the SMB server will negotiate SMB packet signing when an SMB client requests it.

If this setting is enabled, the Microsoft network server will negotiate SMB packet signing as requested by the client. That is, if packet signing has been enabled on the client, packet signing will be negotiated. If this policy is disabled, the SMB client will never negotiate SMB packet signing.

**Default**: Enabled on domain controllers only.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

### Important

- Tor Windows 2000 servers to negotiate signing with Windows NT 4.0 clients, the following registry value must be set to 1 on the server running Windows 2000:
  HKLM\System\CurrentControlSet\Services\lanmanserver\parameters\enableW9xsecuritysignature

### Notes

- All Windows operating systems support both a client-side SMB component and a server-side SMB component. To take advantage of SMB packet signing, both the client-side SMB component and server-side SMB component that are involved in a communication must have SMB packet signing either enabled or required. For Windows 2000 and above, enabling or requiring packet signing for client and server-side SMB components is controlled by the following four policy settings:

    - Microsoft network client: Digitally sign communications (always) - Controls whether or not the client-side SMB component requires packet signing.

    - Microsoft network client: Digitally sign communications (if server agrees) - Controls whether or not the client-side SMB component has packet signing enabled.

    - Microsoft network server: Digitally sign communications (always) - Controls whether or not the server-side SMB component requires packet signing.

    - Microsoft network server: Digitally sign communications (if client agrees) - Controls whether or not the server-side SMB component has packet signing enabled.

- If server-side SMB signing is required, a client will not be able to establish a session with that server unless it has client-side SMB signing enabled. By default, client-side SMB signing is enabled on workstations, servers, and domain controllers.

- Similarly, if client-side SMB signing is required, that client will not be able to establish a session with servers that do not have packet signing enabled. By default, server-side SMB signing is enabled only on domain controllers.

- If server-side SMB signing is enabled, SMB packet signing will be negotiated with clients that have client-side SMB signing enabled.

- Using SMB packet signing can impose up to a 15 percent performance hit on file service transactions.

For more information, see:

- Microsoft network client: Digitally sign communications (always)

- Microsoft network client: Digitally sign communications (if server agrees)

- Microsoft network server: Digitally sign communications (always)

- Security Configuration Manager tools

# Microsoft network server: Disconnect clients when logon hours expire

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft network server: Disconnect clients when logon hours expire**

**Description**
This security setting determines whether to disconnect users who are connected to the local computer outside their user account's valid logon hours. This setting affects the Server Message Block (SMB) component.

When this policy is enabled, it causes client sessions with the SMB Service to be forcibly disconnected when the client's logon hours expire.

If this policy is disabled, an established client session is allowed to be maintained after the client's logon hours have expired.

**Default:** Undefined.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Set logon hours
- Security Configuration Manager tools

# Network access: Allow anonymous SID/Name translation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Allow anonymous SID/name translation**

**Description**
This security setting determines if an anonymous user can request security identifier (SID) attributes for another user.

If this policy is enabled, a user with knowledge of an administrator's SID could contact a computer that has this policy enabled and use the SID to get the administrator's name.

**Default**:

- Disabled on workstations and member servers.
- Enabled on domain controllers.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Account and local policies
- Security Configuration Manager tools

# Network access: Do not allow anonymous enumeration of SAM accounts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Do not allow anonymous enumeration of SAM accounts**

**Description**
This security setting determines what additional permissions will be granted for anonymous connections to the computer.

Windows allows anonymous users to perform certain activities, such as enumerating the names of domain accounts and network shares. This is convenient, for example, when an administrator wants to grant access to users in a trusted domain that does not maintain a reciprocal trust.

This security option allows additional restrictions to be placed on anonymous connections as follows:

- **None. Rely on default permissions.**
- **Do not allow enumeration of SAM accounts.** This option replaces Everyone with Authenticated Users in the security permissions for resources.

**Default:**

- Enabled on workstation.
- Disabled on server.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- This policy has no impact on domain controllers.

For more information, see:

- Network access: Let Everyone permissions apply to anonymous users
- Security Configuration Manager tools

# Network access: Do not allow anonymous enumeration of SAM accounts and shares

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Do not allow anonymous enumeration of SAM accounts and shares**

**Description**
This security setting determines whether anonymous enumeration of SAM accounts and shares is allowed.

Windows allows anonymous users to perform certain activities, such as enumerating the names of domain accounts and network shares. This is convenient, for example, when an administrator wants to grant access to users in a trusted domain that does not maintain a reciprocal trust. If you do not want to allow anonymous enumeration of SAM accounts and shares, then enable this policy.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Network access: Let Everyone permissions apply to anonymous users
- Security Configuration Manager tools

# Network access: Do not allow storage of credentials or .NET Passports for network authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Do not allow storage of credentials or .NET Passports for network authentication**

**Description**
This security setting determines whether Stored User Names and Passwords saves passwords, credentials, or .NET Passports for later use when it gains domain authentication.

If it is enabled, this setting prevents the Stored User Names and Passwords from storing passwords and credentials.

**Note**

- When configuring this security setting, changes will not take effect until you restart Windows.

For more information about Stored User Names and Passwords, see Stored User Names and Passwords.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools
- Stored User Names and Passwords

# Network access: Let Everyone permissions apply to anonymous users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Let Everyone permissions apply to anonymous users**

**Description**
This security setting determines what additional permissions are granted for anonymous connections to the computer.

Windows allows anonymous users to perform certain activities, such as enumerating the names of domain accounts and network shares. This is convenient, for example, when an administrator wants to grant access to users in a trusted domain that does not maintain a reciprocal trust. By Default, the Everyone security identifier (SID) is removed from the token created for anonymous connections. Therefore, permissions granted to the Everyone group do not apply to anonymous users. If this option is set, anonymous users can only access those resources for which the anonymous user has been explicitly given permission.

If this policy is enabled, the Everyone SID is added to the token that is created for anonymous connections. In this case, anonymous users are able to access any resource for which the Everyone group has been given permissions.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Network access: Do not allow anonymous enumeration of SAM accounts and shares
- Security Configuration Manager tools

# Network access: Named pipes that can be accessed anonymously

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Named pipes that can be accessed anonymously**

**Description**
This security setting determines which communication sessions (pipes) will have attributes and permissions that allow anonymous access.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Network access: Remotely accessible registry paths

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Remotely accessible registry paths**

**Description**
This security setting determines which registry paths can be accessed over the network, regardless of the users or groups listed in the access control list (ACL) of the **winreg** registry key. For information about the **winreg** registry key, see article Q153183, "How to Restrict Access to the Registry from a Remote Computer," in the Microsoft Knowledge Base.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

### Note

- This security setting is not available on earlier versions of Windows. The security setting that appears on computers running Windows XP, "Network access: Remotely accessible registry paths" corresponds to the "Network access: Remotely accessible registry paths and subpaths" security option on members of the Windows Server 2003 family. For more information, see Network access: Remotely accessible registry paths and subpaths.

### Default:

- **System\CurrentControlSet\Control\ProductOptions**
- **System\CurrentControlSet\Control\Server Applications**
- **Software\Microsoft\Windows NT\CurrentVersion**

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Network access: Remotely accessible registry paths and subpaths

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Remotely accessible registry paths and subpaths**

**Description**
This security setting determines which registry paths and subpaths can be accessed over the network, regardless of the users or groups listed in the access control list (ACL) of the **winreg** registry key. For information about the **winreg** registry key, see article Q153183, "How to Restrict Access to the Registry from a Remote Computer," in the Microsoft Knowledge Base.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Note**

- On Windows XP, this security setting was called "Network access: Remotely accessible registry paths." If you configure this setting on a member of the Windows Server 2003 family that is joined to a domain, this setting is inherited by computers running Windows XP, but will appear as the "Network access: Remotely accessible registry paths" security option. For more information, see Network access: Remotely accessible registry paths.

**Default:**

- **System\CurrentControlSet\Control\Print\Printers**
- **System\CurrentControlSet\Services\Eventlog**
- **Software\Microsoft\OLAP Server**
- **Software\Microsoft\Windows NT\CurrentVersion\Print**
- **Software\Microsoft\Windows NT\CurrentVersion\Windows**
- **System\CurrentControlSet\Control\ContentIndex**
- **System\CurrentControlSet\Control\Terminal Server**
- **System\CurrentControlSet\Control\Terminal Server\UserConfig**
- **System\CurrentControlSet\Control\Terminal Server\DefaultUserConfiguration**
- **Software\Microsoft\Windows NT\CurrentVersion\Perflib**
- **System\CurrentControlSet\Services\SysmonLog**
- **System\CurrentControlSet\Services\CertSvc**
- **System\CurrentControlSet\Services\Wins**

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer

Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Network access: Restrict anonymous access to Named Pipes and Shares

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Restrict anonymous access to Named Pipes and Shares**

**Description**
When enabled, this security setting restricts anonymous access to shares and pipes to the settings for:

- Network access: Named pipes that can be accessed anonymously
- Network access: Shares that can be accessed anonymously

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For information on implementing this policy, see:

- Security Configuration Manager tools

# Network access: Shares that can be accessed anonymously

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Shares that can be accessed anonymously**

**Description**
This security setting determines which network shares can accessed by anonymous users.

**Default:** None specified.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Network access: Sharing and security model for local accounts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access: Sharing and security model for local accounts**

**Description**
This security setting determines how network logons using local accounts are authenticated. If this setting is set to **Classic**, network logons that use local account credentials authenticate by using those credentials. If this setting is set to **Guest only**, network logons that use local accounts are automatically mapped to the Guest account. The Classic model allows fine control over access to resources. By using the Classic model, you can grant different types of access to different users for the same resource. By using the Guest only model, you can have all users treated equally. All users authenticate as Guest, and they all receive the same level of access to a given resource, which can be either Read Only or Modify.

There are two models available:

- Classic: Local users authenticate as themselves.
- Guest only: Local users authenticate as Guest.

**Default:** Guest only on Windows XP Professional. Classic on the Windows Server 2003 family and Windows XP Professional computers joined to a domain.

**Important**

- With the Guest only model, any user who can access your computer over the network (including anonymous Internet users) can access your shared resources. You must use the Internet Connection Firewall (ICF), Windows Firewall, or other similar device to protect your computer from unauthorized access. Similarly, with the Classic model, local accounts must be password protected; otherwise, those user accounts can be used by anyone to access shared system resources.
- This setting only affects computers running Windows XP Professional which are not joined to a domain.
- This policy will have no impact on computers running Windows 2000.
- Windows Firewall is not included in the original release of Windows Server 2003. Internet Connection Firewall is included only in the original release of Windows Server 2003, Standard Edition and the original release of the 32-bit version of Windows Server 2003, Enterprise Edition.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Notes**

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- This setting does not affect interactive logons that are performed remotely by using such services as Telnet or Terminal Services.
- When the computer is not joined to a domain, this setting also modifies the **Sharing** and **Security** tabs in the Windows Explorer to correspond to the sharing and security model that is being used.

For more information, see:

- Security Configuration Manager tools

# Network security: Do not store LAN Manager hash value on next password change

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network security: Do not store LAN Manager hash value on next password change**

**Description**
This security setting determines if, at the next password change, the LAN Manager (LM) hash value for the new password is stored. The LM hash is relatively weak and prone to attack, as compared with the cryptographically stronger Windows NT hash. Since the LM hash is stored on the local computer in the security database the passwords can be compromised if the security database is attacked. For more information on cryptographic hashes of passwords, see Microsoft NTLM .

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- Windows 2000 Service Pack 2 (SP2) and above offer compatibility with authentication to previous versions of Windows, such as Microsoft Windows NT 4.0.

- This setting can affect the ability of computers running Windows 2000 Server, Windows 2000 Professional, Windows XP, and the Windows Server 2003 family to communicate with computers running Windows 95 and Windows 98.

For more information, see:

- Security Configuration Manager tools

# Network security: Force logoff when logon hours expire

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network security: Force logoff when logon hours expire**

**Description**
This security setting determines whether to disconnect users who are connected to the local computer outside their user account's valid logon hours. This setting affects the Server Message Block (SMB) component.

When this policy is enabled, it causes client sessions with the SMB server to be forcibly disconnected when the client's logon hours expire.

If this policy is disabled, an established client session is allowed to be maintained after the client's logon hours have expired.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Note**

- This security setting behaves as an account policy. For domain accounts, there can be only one account policy. The account policy must be defined in the Default Domain Policy, and it is enforced by the domain controllers that make up the domain. A domain controller always pulls the account policy from the Default Domain Policy Group Policy object (GPO), even if there is a different account policy applied to the organizational unit that contains the domain controller. By default, workstations and servers that are joined to a domain (for example, member computers) also receive the same account policy for their local accounts. However, local account policies for member computers can be different from the domain account policy by defining an account policy for the organizational unit that contains the member computers. Kerberos settings are not applied to member computers.

For more information, see:

- Account and local policies
- Set logon hours
- Security Configuration Manager tools

# Network security: LAN Manager authentication level

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network security: LAN Manager authentication level**

**Description**
This security setting determines which challenge/response authentication protocol is used for network logons. This choice affects the level of authentication protocol used by clients, the level of session security negotiated, and the level of authentication accepted by servers as follows:

- **Send LM & NTLM responses:** Clients use LM and NTLM authentication and never use NTLMv2 session security; domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Send LM & NTLM - use NTLMv2 session security if negotiated:** Clients use LM and NTLM authentication and use NTLMv2 session security if the server supports it; domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Send NTLM response only:** Clients use NTLM authentication only and use NTLMv2 session security if the server supports it; domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Send NTLMv2 response only:** Clients use NTLMv2 authentication only and use NTLMv2 session security if the server supports it; domain controllers accept LM, NTLM, and NTLMv2 authentication.
- **Send NTLMv2 response only\refuse LM:** Clients use NTLMv2 authentication only and use NTLMv2 session security if the server supports it; domain controllers refuse LM (accept only NTLM and NTLMv2 authentication).
- **Send NTLMv2 response only\refuse LM & NTLM:** Clients use NTLMv2 authentication only and use NTLMv2 session security if the server supports it; domain controllers refuse LM and NTLM (accept only NTLMv2 authentication).

**Important**

- This setting can affect the ability of computers running Windows 2000 Server, Windows 2000 Professional, Windows XP Professional, and the Windows Server 2003 family to communicate with computers running Windows NT 4.0 and earlier over the network. For example, at the time of this writing, computers running Windows NT 4.0 SP4 and earlier did not support NTLMv2. Computers running Windows 95 and Windows 98 did not support NTLM. For more information about how these settings apply to previous versions of Windows, see Knowledge Base Search

**Default:**

- **Send NTLM response only.**

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Network security: LDAP client signing requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network security: LDAP client signing requirements**

**Description**
This security setting determines the level of data signing that is requested on behalf of clients issuing LDAP BIND requests, as follows:

- **None:** The LDAP BIND request is issued with the options that are specified by the caller.
- **Negotiate signing:** If Transport Layer Security/Secure Sockets Layer (TLS\SSL) has not been started, the LDAP BIND request is initiated with the LDAP data signing option set in addition to the options specified by the caller. If TLS\SSL has been started, the LDAP BIND request is initiated with the options that are specified by the caller.
- **Require signature:** This is the same as **Negotiate signing**. However, if the LDAP server's intermediate saslBindInprogress response does not indicate that LDAP traffic signing is required, the caller is told that the LDAP BIND command request failed.

**Caution**

- If you set the server to **Require signature,** you must also set the client. Not setting the client results in a loss of connection with the server.

**Note**

- This setting does not have any impact on ldap_simple_bind or ldap_simple_bind_s. No Microsoft LDAP clients that are shipped with Windows XP Professional use ldap_simple_bind or ldap_simple_bind_s to talk to a domain controller.

**Default:** Negotiate signing.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Network security: Minimum session security for NTLM SSP based (including secure RPC) clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network security: Minimum session security for NTLM SSP based (including secure RPC) clients**

**Description**
This security setting allows a client to require the negotiation of message confidentiality (encryption), message integrity, 128-bit encryption, or NTLMv2 session security. These values are dependent on the LAN Manager Authentication Level security setting value.

- **Require message integrity**. The connection will fail if message integrity is not negotiated. The integrity of a message can be assessed through message signing. Message signing proves that the message has not been tampered with by attaching a cryptographic signature which identifies the sender and is a numeric representation of the contents of the message. This signature ensures that the message has not been tampered with.

- **Require message confidentiality**. The connection will fail if encryption is not negotiated. Encryption converts data into a form that is not readable by anyone until decrypted.

- **Require NTLMv2 session security**. The connection will fail if the NTLMv2 protocol is not negotiated. For more information about NTLM and NTLMv2, see Logon and Authentication; at the Microsoft Windows Resource Kits Web site.

- **Require 128-bit encryption**. The connection will fail if strong encryption (128-bit) is not negotiated.

**Default:** No requirements.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- This setting will apply to any computers running Windows 2000 through changes in the registry but the security setting will not be viewable through the Security Configuration Manager tool set.

For more information, see:

- Security Configuration Manager tools

- Network security: LAN Manager authentication level

# Network security: Minimum session security for NTLM SSP based (including secure RPC) servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network security: Minimum session security for NTLM SSP based (including secure RPC) servers**

**Description**
This security setting allows a server to require the negotiation of message confidentiality (encryption), message integrity, 128-bit encryption, or NTLMv2 session security. These values are dependent on the LAN Manager Authentication Level security setting value.

- **Require message integrity**. The connection will fail if message integrity is not negotiated. The integrity of a message can be assessed through message signing. Message signing proves that the message has not been tampered with by attaching a cryptographic signature which identifies the sender and is a numeric representation of the contents of the message. This signature ensures that the message has not been tampered with.

- **Require message confidentiality**. The connection will fail if encryption is not negotiated. Encryption converts data into a form that is not readable by anyone until decrypted.

- **Require NTLMv2 session security**. The connection will fail if the NTLMv2 protocol is not negotiated. For more information about NTLM and NTLMv2, see Logon and Authentication; at the Microsoft Windows Resource Kits Web site.

- **Require 128-bit encryption**. The connection will fail if strong encryption (128-bit) is not negotiated.

**Default:** No requirements.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Important**

- This setting will apply to any computers running Windows 2000 through changes in the registry but the security setting will not be viewable through the Security Configuration Manager tool set.

For more information, see:

- Security Configuration Manager tools

- Network security: LAN Manager authentication level

# Recovery console: Allow automatic administrative logon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovery console: Allow automatic administrative logon**

**Description**
This security setting determines if the password for the Administrator account must be given before access to the system is granted. If this option is enabled, the Recovery Console does not require you to provide a password, and it automatically logs on to the system.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Recovery console: Allow floppy copy and access to all drives and all folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovery console: Allow floppy copy and access to all drives and all folders**

**Description**
Enabling this security option makes the Recovery Console SET command available, which allows you to set the following Recovery Console environment variables:

- **AllowWildCards:** Enable wildcard support for some commands (such as the DEL command).
- **AllowAllPaths:** Allow access to all files and folders on the computer.
- **AllowRemovableMedia:** Allow files to be copied to removable media, such as a floppy disk.
- **NoCopyPrompt:** Do not prompt when overwriting an existing file.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Shutdown: Allow system to be shut down without having to log on

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shutdown: Allow system to be shut down without having to log on**

**Description**
This security setting determines whether a computer can be shut down without having to log on to Windows.

When this policy is enabled, the **Shut Down** command is available on the Windows logon screen.

When this policy is disabled, the option to shut down the computer does not appear on the Windows logon screen. In this case, users must be able to log on to the computer successfully and have the Shut down the system user right before they can perform a system shutdown.

**Default**

- Enabled on workstations.
- Disabled on servers.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# Shutdown: Clear virtual memory pagefile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shutdown: Clear virtual memory pagefile**

**Description**
This security setting determines whether the virtual memory pagefile is cleared when the system is shut down.

Virtual memory support uses a system pagefile to swap pages of memory to disk when they are not used. On a running system, this pagefile is opened exclusively by the operating system, and it is well protected. However, systems that are configured to allow booting to other operating systems might have to make sure that the system pagefile is wiped clean when this system shuts down. This ensures that sensitive information from process memory that might go into the pagefile is not available to an unauthorized user who manages to directly access the pagefile.

When this policy is enabled, it causes the system pagefile to be cleared upon clean shutdown. If you enable this security option, the hibernation file (hiberfil.sys) is also zeroed out when hibernation is disabled on a portable computer system.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# System cryptography: Force strong key protection for user keys stored on the computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Cryptography: Force strong key protection for user keys stored on the computer**

**Description**
This security setting determines if users' private keys require a password to be used.

The options are:

- **User input is not required when new keys are stored and used**
- **User is prompted when the key is first used**
- **User must enter a password each time they use a key**

  For more information, see Public Key Infrastructure.

**Default:** Not defined.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System cryptography: Use FIPS compliant algorithms for encryption, hashing and signing**

**Description**
This security setting determines if the Transport Layer Security/Secure Sockets Layer (TL/SS) Security Provider supports only the TLS_RSA_WITH_3DES_EDE_CBC_SHA cipher suite. In effect, this means that the provider only supports the Transport Layer Security (TLS) protocol as a client and as a server (if applicable). It uses only the Triple DES encryption algorithm for the TLS traffic encryption, only the Rivest, Shamir, and Adleman (RSA) public key algorithm for the TLS key exchange and authentication, and only the Secure Hashing Algorithm 1 (SHA-1) for the TLS hashing requirements.

For Encrypting File System Service (EFS), it supports only the Triple Data Encryption Standard (DES) encryption algorithm for encrypting file data supported by the NTFS file system. By default, EFS uses the Advanced Encryption Standard (AES) algorithm with a 256-bit key in the Windows Server 2003 family and DESX algorithm in Windows XP for encrypting file data. For information about EFS, see Encrypting File System.

For Terminal Services, it supports only the Triple DES encryption algorithm for encrypting terminal services network communication. For information about Terminal Services, see Terminal Services.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

**Note**

- The Federal Information Processing Standard (FIPS) 140-1 is a security implementation designed for certifying cryptographic software. FIPS 140-1 validated software is required by the U.S. Government and requested by other prominent institutions.

For more information, see:

- Encrypting File System
- Security Configuration Manager tools

# System objects: Default owner for objects created by members of the Administrators group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System objects: Default owner for objects created by members of the Administrators group**

## Description

This security setting determines which users and groups have the authority to run volume maintenance tasks, such as Disk Cleanup and Disk Defragmenter.

**Default:** Administrators group (on servers).

**Configuring this security setting**

You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# System objects: Require case insensitivity for non-Windows subsystems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System objects: Require case insensitivity for non-Windows subsystems**

**Description**
This security setting determines whether case insensitivity is enforced for all subsystems. The Win32 subsystem is case insensitive. However, the kernel supports case sensitivity for other subsystems, such as POSIX.

If this setting is enabled, case insensitivity is enforced for all directory objects, symbolic links, and IO objects, including file objects. Disabling this setting does not allow the Win32 subsystem to become case sensitive.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# System objects: Strengthen default permissions of internal system objects (e.g., Symbolic Links)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System objects: Strengthen default permissions of internal system objects (e.g., Symbolic Links)**

**Description**
This security setting determines the strength of the default discretionary access control list (DACL) for objects.

Active Directory maintains a global list of shared system resources, such as DOS device names, mutexes, and semaphores. In this way, objects can be located and shared among processes. Each type of object is created with a default DACL that specifies who can access the objects and what permissions are granted.

If this policy is enabled, the default DACL is stronger, allowing users who are not administrators to read shared objects but not allowing these users to modify shared objects that they did not create.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# System settings: Optional subsystems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System settings: Optional subsystems**

**Description**
This security setting determines which subsystems are used to support your applications. With this security setting, you can specify as many subsytems to support as your environment demands.

**Default:** POSIX.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Security Configuration Manager tools

# System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**

**Description**
This security setting determines if digital certificates are processed when a user or process attempts to run software with an .exe file name extension. This security settings is used to enable or disable certificate rules, a type of software restriction policies rule. With software restriction policies, you can create a certificate rule that will allow or disallow software that is signed by Authenticode to run, based on the digital certificate that is associated with the software. In order for certificate rules to take effect, you must enable this security setting.

When certificate rules are enabled, software restriction policies will check a certificate revocation list (CRL) to make sure the software's certificate and signature are valid. This may decrease performance when start signed programs. You can disable this feature. On **Trusted Publishers Properties**, clear the **Publisher** and **Timestamp** check boxes. For more information, see Set trusted publisher options.

**Default:** Disabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\

For specific instructions about how to configure this security setting, see Enable certificate rules.

For more information, see:

- Software Restriction Policies
- Create a certificate rule
- Security levels and additional rules
- Best practices for security
- Security Options
- Security Configuration Manager tools

# Settings for Event Logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Settings for event logs**
This section covers:

- Maximum application log size
- Maximum security log size
- Maximum system log size
- Prevent local guests group from accessing application log
- Prevent local guests group from accessing security log
- Prevent local guests group from accessing system log
- Retain application log
- Retain security log
- Retain system log
- Retention method for application log
- Retention method for security log
- Retention method for system log

# Maximum application log size

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maximum application log size**

**Description**
This security setting specifies the maximum size of the application event log, which has a maximum of 4 GB.

### Notes

- Log file sizes must be a multiple of 64 KB. If you enter a value that is not a multiple of 64 KB, Event Viewer will set the log file size to a multiple of 64 KB.
- This setting does not appear in the Local Computer Policy object.
- Event Log size and log wrapping should be defined to match the business and security requirements you determined when designing your enterprise security plan. Consider implementing these Event Log settings at the site, domain, or organizational unit level, to take advantage of Group Policy settings.

**Default:** For the Windows Server 2003 family, 16 MB; for Windows XP Professional Service Pack 1, 8 MB; for Windows XP Professional, 512 KB.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Maximum security log size

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maximum security log size**

**Description**
This security setting specifies the maximum size of the security event log, which has a maximum size of 4 GB.

### Notes

- Log file sizes must be a multiple of 64 KB. If you enter a value that is not a multiple of 64 KB, Event Viewer will set the log file size to a multiple of 64 KB.
- This setting does not appear in the Local Computer Policy object.
- Event Log size and log wrapping should be defined to match the business and security requirements you determined when designing your enterprise security plan. Consider implementing these Event Log settings at the site, domain, or organizational unit level, to take advantage of Group Policy settings.

**Default:** For the Windows Server 2003 family, 16 MB; for Windows XP Professional Service Pack 1, 8 MB; for Windows XP Professional, 512 KB.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Maximum system log size

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maximum system log size**

**Description**
This security setting specifies the maximum size of the system event log, which has a maximum size of 4 GB.

### Notes

- Log file sizes must be a multiple of 64 KB. If you enter a value that is not a multiple of 64 KB, Event Viewer will set the log file size to a multiple of 64 KB.
- This setting does not appear in the Local Computer Policy object.
- Event Log size and log wrapping should be defined to match the business and security requirements you determined when designing your Enterprise Security Plan. Consider implementing these Event Log settings at the site, domain, or organizational unit level, to take advantage of Group Policy settings.

**Default:** For the Windows Server 2003 family, 16 MB; for Windows XP Professional Service Pack 1, 8 MB; for Windows XP Professional, 512 KB.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Prevent local guests group from accessing application log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Prevent local guests group from accessing application log**

**Description**
This security setting determines if guests are prevented from accessing the application event log.

### Notes

- This setting does not appear in the Local Computer Policy object.
- This security setting affects only computers running Windows 2000, Windows XP and Windows Server 2003.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Prevent local guests group from accessing security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Prevent local guests group from accessing security log**

**Description**
This security setting determines if guests are prevented from accessing the security event log.

### Notes

- This setting does not appear in the Local Computer Policy object.
- This security setting affects only computers running Windows 2000, Windows XP and Windows Server 2003.
- A user must possess the Manage auditing and security log user right to access the security log.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Prevent local guests group from accessing system log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Prevent local guests group from accessing system log**

**Description**
This security setting determines if guests are prevented from accessing the system event log.

### Notes

- This setting does not appear in the Local Computer Policy object.
- This security setting affects only computers running Windows 2000, Windows XP and Windows Server 2003.

**Default:** Enabled.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Retain application log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Retain application log**

**Description**
This security setting determines the number of days' worth of events to be retained for the application log if the retention method for the application log is **By Days**.

Set this value only if you archive the log at scheduled intervals and you make sure that the Maximum application log size is large enough to accommodate the interval.

### Notes

- This setting does not appear in the Local Computer Policy object.
- A user must possess the Manage auditing and security log user right to access the security log.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Retain security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Retain security log**

**Description**
This security setting determines the number of days' worth of events to be retained for the security log if the retention method for the security log is **By Days**.

Set this value only if you archive the log at scheduled intervals and you make sure that the Maximum security log size is large enough to accommodate the interval.

### Note

- This setting does not appear in the Local Computer Policy object.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Retain system log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Retain system log**

**Description**
This security setting determines the number of days' worth of events to be retained for the system log if the retention method for the system log is **By Days**.

Set this value only if you archive the log at scheduled intervals and you make sure that the Maximum system log size is large enough to accommodate the interval.

### Note

- This setting does not appear in the Local Computer Policy object.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Retention method for application log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Retention method for application log**

**Description**
This security setting determines the "wrapping" method for the application log.

If you do not archive the application log, in the **Properties** dialog box for this policy, select the **Define this policy setting** check box, and then click **Overwrite events as needed**.

If you archive the log at scheduled intervals, in the **Properties** dialog box for this policy, select the **Define this policy setting** check box, and then click **Overwrite events by days** and specify the appropriate number of days in the **Retain application log** setting. Make sure that the Maximum application log size is large enough to accommodate the interval.

If you must retain all the events in the log, in the **Properties** dialog box for this policy, select the **Define this policy setting** check box, and then click **Do not overwrite events (clear log manually)**. This option requires that the log be cleared manually. In this case, when the maximum log size is reached, new events are discarded.

**Note**

- This setting does not appear in the Local Computer Policy object.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Retention method for security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Retention method for security log**

**Description**
This security setting determines the "wrapping" method for the security log.

If you do not archive the security log, in the **Properties** dialog box for this policy, select the **Define this policy setting** check box, and then click **Overwrite events as needed**.

If you archive the log at scheduled intervals, in the **Properties** dialog box for this policy, select the **Define this policy setting** check box, and then click **Overwrite events by days** and specify the appropriate number of days in the retain security log setting. Make sure that the Maximum security log size is large enough to accommodate the interval.

If you must retain all the events in the log, in the **Properties** dialog box for this policy, select the **Define this policy setting** check box, and then click **Do not overwrite events (clear log manually)**. This option requires that the log be cleared manually. In this case, when the maximum log size is reached, new events are discarded.

### Note

- This setting does not appear in the Local Computer Policy object.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Retention method for system log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Retention method for system log**

**Description**
This security setting determines the "wrapping" method for the system log.

If you do not archive the system log, in the **Properties** dialog box for this policy, select the **Define this policy setting** check box, and then click **Overwrite events as needed**.

If you archive the log at scheduled intervals, in the **Properties** dialog box for this policy, select the **Define this policy setting** check box, and then click **Overwrite events by days** and specify the appropriate number of days in the **Retain system log** setting. Make sure that the Maximum system log size is large enough to accommodate the interval.

If you must retain all the events in the log, in the **Properties** dialog box for this policy, select the **Define this policy setting** check box, and then click **Do not overwrite events (clear log manually)**. This option requires that the log be cleared manually. In this case, when the maximum log size is reached, new events are discarded.

**Note**

- This setting does not appear in the Local Computer Policy object.

**Default:** None.

**Configuring this security setting**
You can configure this security setting by opening the appropriate policy and expanding the console tree as such: Computer Configuration\Windows Settings\Security Settings\Event Log\

For specific instructions about how to configure security policy settings, see Edit security settings on a Group Policy object.

For more information, see:

- Event Viewer
- Security Configuration Manager tools

# Restricted Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restricted Groups**
To locate this security setting in the console tree in Microsoft Management Console (MMC), see Computer Configuration\Windows Settings\Security Settings\Restricted Groups

**Description**
This security setting allows an administrator to define two properties for security-sensitive groups ("restricted" groups).

The two properties are **Members** and **Member Of**. The **Members** list defines who belongs and who does not belong to the restricted group. A group can be a member of groups other than those listed in the **Member Of** section. The **Member Of** section simply ensures that the restricted group is added to the groups listed in **Member Of**. It does not remove the group from other groups of which it is a member.

When a Restricted Groups Policy is enforced, any current member of a restricted group that is not on the **Members** list is removed. Any user on the **Members** list who is not currently a member of the restricted group is added.

You can use Restricted Groups policy to control group membership. Using the policy, you can specify what members are part of a group. Any members that are not specified in the policy are removed during configuration or refresh. In addition, the reverse membership configuration option ensures that each Restricted Group is a member of only those groups that are specified in the **Member Of** column.

For example, you can create a Restricted Groups policy to only allow specified users (for example, Alice and John) to be members of the Administrators group. When policy is refreshed, only Alice and John will remain as members of the Administrators group.

There are two ways to apply Restricted Groups policy:

- Define the policy in a security template, which will be applied during configuration on your local computer.
- Define the setting on a Group Policy object (GPO) directly, which means that the policy goes into effect with every refresh of policy. The security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes.

**Default:** None specified.

**Caution**

- If a Restricted Groups policy is defined and Group Policy is refreshed, any current member not on the Restricted Groups policy members list is removed. This can include default members, such as administrators.

**Notes**

- Restricted Groups should be used primarily to configure membership of local groups on workstation or member servers.
- An empty **Members** list means that the restricted group has no members; an empty **Member Of** list means that the groups to which the restricted group belongs are not specified.

For more information, see:

- Security Configuration Manager tools
- To Create a Restricted Groups policy in a security template

# System Services security settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Services security settings**
To locate this security setting in the console tree in Microsoft Management Console (MMC), see Computer Configuration\Windows Settings\Security Settings\System Services

**Description**
Allows an administrator to define the startup mode (manual, automatic, or disabled) as well as the access permissions (Start, Stop, or Pause) for all system services.

**Default:** Undefined.

### Notes

- This setting does not appear in the Local Computer Policy object.
- If you choose to set system service startup to Automatic, perform adequate testing to verify that the services can start without user intervention.
- For performance optimization, set unnecessary or unused services to Manual.

For more information, see:

- Security Configuration Manager tools
- Services Snap-in

# Registry security settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Registry security settings**
To locate this security setting in the console tree in Microsoft Management Console (MMC), see Computer Configuration\Windows Settings\Security Settings\Registry

**Description**
Allows an administrator to define access permissions (on discretionary access control lists (DACLs)) and audit settings (on system access control lists (SACLs)) for registry keys using Security Configuration Manager.

**Default:** Undefined.

**Note**

- This setting does not appear in the Local Computer Policy object.

For more information, see:

- Security Configuration Manager tools
- The Registry
- Access Control
- Auditing Security Events

# File System security settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File System security settings**
To locate this security setting in the console tree in Microsoft Management Console (MMC), see Computer Configuration\Windows Settings\Security Settings\File System

**Description**
Allows an administrator to define access permissions (on discretionary access control lists (DACLs)) and audit settings (on system access control lists (SACLs)) for file system objects using Security Configuration Manager.

**Default:** Undefined.

**Note**

- This setting does not appear in the Local Computer Policy object.

For more information, see:

- Security Configuration Manager tools
- Access Control
- Auditing Security Events

# Public Key Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Public key policies**
Using public key policies in a Windows Server 2003 domain, an administrator can automatically issue certificates to computers, manage encrypted data recovery agents, create certificate trust lists, or automatically establish trust of certification authorities.

- For help with specific tasks, see Public Key Policies How To ...
- For general background information, see Public Key Policies Concepts

# Public Key Policies How To ...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Add a trusted root certification authority to a Group Policy object
- Delete a trusted root certification authority from a Group Policy object
- Prevent users from selecting new root certification authorities with a Group Policy
- Prevent users from trusting third-party root certification authorities with a Group Policy
- Create an automatic certificate request for computers in a Group Policy object
- Delete an automatic certificate request for computers in a Group Policy object
- Create a certificate trust list for a Group Policy object
- Edit a certificate trust list for a Group Policy object
- Delete a certificate trust list for a Group Policy object
- Add a recovery agent for the local computer
- Add a recovery agent for a domain
- Change the recovery policy for the local computer
- Create a recovery policy for a domain
- Not require a recovery agent for a domain
- Disable Encrypting File System for a domain

# Add a trusted root certification authority to a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a trusted root certification authority to a Group Policy object**

1. Open the Group Policy object (GPO) that you want to edit.

2. In the console tree, click **Trusted Root Certification Authorities**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Trusted Root Certification Authorities

3. On the **Action** menu, point to **All Tasks**, and then click **Import**.

   This starts the Certificate Import Wizard, which guides you through the process of importing a root certificate and installing it as a trusted root certification authority (CA) for this GPO.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a GPO, see Group Policy (pre-GPMC).

- This procedure does not apply to Local Policy objects.

- You can import a trusted root certificate from a PKCS #12 file (.pfx, .p12), a PKCS #7 file (.spc, .p7b), a certificate file (.cer, .crt), or a Microsoft serialized certificate store file (.sst).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Public Key Policies How To ...
Automatic certificate request settings
Automatic certificate request policy
Working with MMC console files
Deploying a Public Key Infrastructure

# Delete a trusted root certification authority from a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a trusted root certification authority from a Group Policy object**

1. Open the Group Policy object (GPO) that you want to edit.
2. In the console tree, click **Trusted Root Certification Authorities**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Trusted Root Certification Authorities

3. In the details pane, click the root certificate that you want to delete.
4. On the **Action** menu, click **Delete**.
5. If you are sure that you want to delete this certificate, click **Yes**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a GPO, see Group Policy (pre-GPMC).
- This procedure does not apply to Local Policy objects.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Public Key Policies How To ...
Automatic certificate request settings
Automatic certificate request policy
Working with MMC console files
Deploying a Public Key Infrastructure

# Prevent users from selecting new root certification authorities with a Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prevent users from selecting new root certification authorities with a Group Policy**

1. Open the Group Policy object (GPO) that you want to edit.

2. In the console tree, right-click **Trusted Root Certification Authorities**, and then click **Properties**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Trusted Root Certification Authorities

3. Clear the **Allow users to select new root certification authorities (CAs) to trust** check box.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a GPO, see Group Policy (pre-GPMC).

- This procedure does not apply to Local Policy objects.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Public Key Policies How To ...
Automatic certificate request settings
Automatic certificate request policy
Working with MMC console files
Deploying a Public Key Infrastructure

# Prevent users from trusting third-party root certification authorities with a Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prevent users from trusting third-party root certification authorities with a Group Policy**

1. Open the Group Policy object (GPO) that you want to edit.
2. In the console tree, right-click **Trusted Root Certification Authorities**, and click **Properties**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Trusted Root Certification Authorities

3. In **Client computers can trust the following certificate stores**, click **Enterprise Root Certification Authorities**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a GPO, see Group Policy (pre-GPMC).
- This procedure does not apply to Local Policy objects.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Public Key Policies How To ...
Automatic certificate request settings
Automatic certificate request policy
Working with MMC console files
Deploying a Public Key Infrastructure

# Create an automatic certificate request for computers in a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an automatic certificate request for computers in a Group Policy object**

1. Open the Group Policy object (GPO) that you want to edit.

2. In the console tree, click **Automatic Certificate Request Settings**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies

   - Automatic Certificate Request Settings

3. On the **Action** menu, point to **New**, and then click **Automatic Certificate Request**.

   This starts the Automatic Certificate Request Setup Wizard. Follow the steps in the wizard to create an automatic certificate request for computers that have this GPO applied to them.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a GPO, see Group Policy (pre-GPMC).

- This procedure does not apply to Local Policy objects.

- To use the Automatic Certificate Request Setup Wizard, you must know the certificate template that you want to use. A certificate based on the selected template will be requested automatically at the first occurrence of any of the following: a user logs on, Group Policy is refreshed, or a computer joins the domain and is subject to a Group Policy setting.

- The Automatic Certificate Request Setup Wizard asks which certification authority (CA) it should query when the wizard runs on Windows 2000. The wizard will not prompt you when it runs on Windows XP or the Windows Server 2003 family.

- There can only be one certificate request for each certificate type per GPO.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Public Key Policies How To ...
Automatic certificate request settings
Automatic certificate request policy

Working with MMC console files
Deploying a Public Key Infrastructure

# Delete an automatic certificate request for computers in a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an automatic certificate request for computers in a Group Policy object**

1. Open the Group Policy object (GPO) that you want to edit.
2. In the console tree, click **Automatic Certificate Request Settings**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Automatic Certificate Request Settings

3. In the details pane, click the automatic certificate request that you want to delete.
4. On the **Action** menu, click **Delete**.
5. If you are sure that you want to delete this automatic certificate request, click **Yes**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a GPO, see Group Policy (pre-GPMC).
- This procedure does not apply to Local Policy objects.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Public Key Policies How To ...
Automatic certificate request settings
Automatic certificate request policy
Working with MMC console files
Deploying a Public Key Infrastructure

# Create a certificate trust list for a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a certificate trust list for a Group Policy object**

1. Open the Group Policy object (GPO) that you want to edit.

2. In the console tree, click **Enterprise Trust**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Enterprise Trust

3. On the **Action** menu, point to **New**, and then click **Certificate Trust List**.

   This starts the Certificate Trust List Wizard. Follow the steps in the wizard to create a certificate trust list (CTL) for the GPO.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a GPO, see Group Policy (pre-GPMC).

- This procedure does not apply to Local Policy objects.

- To use the Certificate Trust List Wizard, you need the following information:

  - The certificate purposes for this CTL. For more information about CTL purposes, see Using enterprise trust policy.

  - The certificates that you want to add to the CTL. You can add certificates from the certificate store on the local computer or from a file. Acceptable file formats from which you can import a certificate are an X.509 v3 certificate file (.cer, .crt), a PKCS #7 file (.spc, .p7b), or a Microsoft serialized certificate store (.sst).

  - A certificate that has the Trust List Signing purpose and its associated private key in your personal certificate store. For general instructions on how to request a certificate, see Request a certificate.

  - As an option, you can enter a time stamp service URL to time-stamp the CTL.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Public Key Policies How To ...
Automatic certificate request settings
Automatic certificate request policy

Working with MMC console files
Deploying a Public Key Infrastructure

# Edit a certificate trust list for a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit a certificate trust list for a Group Policy object**

1. Open the Group Policy object (GPO) that you want to edit.
2. In the console tree, click **Enterprise Trust**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Enterprise Trust

3. In the details pane, click the certificate trust list (CTL) that you want to edit.
4. On the **Action** menu, point to **All Tasks**, and then click **Edit**.

   This starts the Certificate Trust List Wizard. Follow the steps in the wizard to edit the CTL for the GPO.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a GPO, see Group Policy (pre-GPMC).
- This procedure does not apply to Local Policy objects.
- To use the Certificate Trust List Wizard, you need the following information:
  - The certificate purposes for this CTL. For more information about CTL purposes, see Using enterprise trust policy.
  - The certificates that you want to add to the CTL. You can add certificates from the certificate store on the local computer or from a file. Acceptable file formats from which you can import a certificate are an X.509 v3 certificate file (.cer, .crt), a PKCS #7 file (.spc, .p7b) or a Microsoft serialized certificate store (.sst).
  - A certificate that has the Trust List Signing purpose and its associated private key in your personal certificate store. For general instructions on how to request a certificate, see Request a certificate.
  - As an option, you can enter a time stamp service URL to time-stamp the CTL.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Public Key Policies How To ...
Automatic certificate request settings

Automatic certificate request policy
Working with MMC console files
Deploying a Public Key Infrastructure

# Delete a certificate trust list for a Group Policy object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a certificate trust list for a Group Policy object**

1. Open the Group Policy object (GPO) that you want to edit.
2. In the console tree, click **Enterprise Trust**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Enterprise Trust

3. In the details pane, right-click the certificate trust list (CTL) that you want to delete.
4. On the **Action** menu, click **Delete**.
5. Click **Yes** if you are sure that you want to delete this CTL.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a GPO, see Group Policy (pre-GPMC).
- This procedure does not apply to Local Policy objects.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Public Key Policies How To ...
Automatic certificate request settings
Automatic certificate request policy
Working with MMC console files
Deploying a Public Key Infrastructure

# Add a recovery agent for the local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a recovery agent for the local computer**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Add Standalone Snap-in**, click **Group Policy Object Editor**, and then click **Add**.
4. Under **Group Policy Object**, make sure that Local Computer is displayed, and then click **Finish**.
5. Click **Close**, and then click **OK**.
6. In **Local Computer Policy**, click **Public Key Policies**.

   **Where?**

   - Local Computer Policy/Computer Configuration/Windows Settings/Security Settings/Public Key Policies

7. In the details pane, right-click **Encrypting File System**.
8. Click **Add Data Recovery Agent**, and then follow the instructions in the Add Recovery Agent Wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- Be prepared to provide the wizard with the user name for a user with a published recovery certificate. Alternatively, you can use the wizard to browse for .cer files that contain information about the recovery agent you are adding.
- Adding a recovery agent from a file identifies the user as USER_UNKNOWN. This is because the name is not stored in the file.
- Before you can add or create a recovery agent, you must configure Group Policy on your computer. For more information about using Group Policy, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovering data
Add a snap-in to a new MMC console for a local computer
Public Key Policies

# Add a recovery agent for a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a recovery agent for a domain**

1. Open Active Directory Users and Computers.
2. Right-click the domain whose recovery policy you want to change, and then click **Properties**.
3. Click the **Group Policy** tab.
4. Right-click the recovery policy you want to change, and then click **Edit**.
5. In the console tree, click **Encrypting File System**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Encrypting File System

6. In the details pane, right-click, then click **Add Data Recovery Agent**.
7. Follow the instructions of the wizard to complete this procedure.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- This operation can be performed on any sites, domains or organizational units within an Active Directory forest.
- Adding a recovery agent from a file identifies the user as USER_UNKNOWN. This is because the name is not stored in the file.
- Adding a recovery agent from Active Directory requires that File Recovery certificates are published in Active Directory. However, the default EFS File Recovery certificate template does not publish these certificates. This can be changed by copying the default EFS File Recovery certificate template to create a new template and configuring it to **Publish certificate in Active Directory**. For more information on modifying certificate templates, see Related Topics.
- Before you can add or create a recovery agent, you must configure Group Policy on your computer. For more information about using Group Policy, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Create a new certificate template
Encrypting File System overview
Recovering data

# Change the recovery policy for the local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the recovery policy for the local computer**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Add Standalone Snap-in**, click **Group Policy Object Editor**, and then click **Add**.
4. Under **Group Policy Object**, make sure that **Local Computer** is displayed, click **Finish**.
5. Click **Close**, and then click **OK**.
6. In **Local Computer Policy**, click **Public Key Policies**.

   **Where?**

   - Local Computer Policy/Computer Configuration/Windows Settings/Security Settings/Public Key Policies

7. Right-click **Encrypting File System**, and then do one of the following:

   - To designate a user as an additional recovery agent using the Add Recovery Agent Wizard, click **Add Data Recovery Agent**.
   - To allow EFS to work without recovery agents, point to **All Tasks** and then click **Do Not Require Data Recovery Agents**.
   - To delete this EFS policy and every recovery agent, point to **All Tasks** and then click **Delete Policy**. If you select this option, users can still encrypt files on this computer. Note that this option will not appear unless there is an EFS policy on the computer.

### Important

- Before changing the recovery policy in any way, you should first back up the recovery keys to a floppy disk.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- There is no default recovery agent on a standalone computer. A file recovery certificate can be created by running cipher.exe /r, and the **Add Data Recovery Agent** option can be used to import this certificate into the EFS policy. Fore more information on cipher.exe, see Related Topics.
- You can make changes to the File Recovery certificate by right-clicking the certificate and then clicking **Properties**. For example, you can give the certificate a friendly name and enter a text description.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Cipher
Recovering data
Back up default recovery keys to a floppy disk
Add a snap-in to a new MMC console for a local computer

# Create a recovery policy for a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a recovery policy for a domain**

1. Open Active Directory Users and Computers.
2. Right-click the domain whose recovery policy you want to change, and then click **Properties**.
3. Click the **Group Policy** tab.
4. Right-click the recovery policy you want to change, and then click **Edit**.
5. In the console tree, click **Encrypting File System**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Encrypting File System

6. Do one of the following:

   - If you need to create a certificate to use as the Encrypting File System (EFS) recovery certificate, right-click **Encrypting File System** and click **Create Data Recovery Agent**.
   - If you want to use an existing certificate, click **Add Data Recovery Agent** and follow the instructions that are given by the wizard to complete this procedure.

**Important**

- Before changing the recovery policy in any way, you should first back up the recovery keys to a floppy disk.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- In a domain, a default recovery policy is implemented for the domain when the first domain controller is set up. The first domain administrator is issued the self-signed certificate, which designates the domain administrator as the recovery agent. To change the default recovery policy for a domain, log on to the first domain controller as an administrator.
- If you choose **Create Data Recovery Agent** during this procedure, the domain controller will contact a Windows Server 2003 family certification authority (CA) to request a certificate based on the EFS Recovery Agent certificate template. If this template is unavailable or does not allow you to obtain a certificate, Windows displays the following message: "Windows cannot create a data recovery agent."

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovering data
Back up default recovery keys to a floppy disk
Certificate Templates

# Not require a recovery agent for a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To not require a recovery agent for a domain**

1. Open Active Directory Users and Computers.
2. Right-click the domain whose recovery policy you want to change, and then click **Properties**.
3. Click the **Group Policy** tab.
4. Right-click the recovery policy you want to change, and then click **Edit**.
5. In the console tree, click **Encrypting File System**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Encrypting File System

6. Right-click **Encrypting File System**, point to **All Tasks**, then click **Do Not Require Data Recovery Agents**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- This operation can be performed on any sites, domains or organizational units within an Active Directory forest.
- If EFS policy is currently defined, you must select **Delete Policy** before you are able to follow this procedure.
- This procedure allows EFS to continue working without using recovery agents. To disable EFS, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable Encrypting File System for a domain
Modify a Certificate Template
Cipher
Encrypt a file or folder
Decrypt a file or folder
Encrypting File System overview
Recovering data

# Disable Encrypting File System for a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable Encrypting File System for a domain**

1. Open Active Directory Users and Computers.
2. Right-click the domain whose recovery policy you want to change, and then click **Properties**.
3. Click the **Group Policy** tab.
4. Right-click the recovery policy you want to change, and then click **Edit**.
5. In the console tree, click **Encrypting File System**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Encrypting File System

6. Right-click **Encrypting File System**, then click **Properties**.
7. Clear the **Allow users to encrypt files using Encrypting File System (EFS)** check box.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- This operation can be performed on any sites, domains or organizational units within an Active Directory forest.
- This procedure completely disables EFS. To allow EFS to continue working without using recovery agents, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable Encrypting File System for the local computer
Modify a Certificate Template
Cipher
Encrypt a file or folder
Decrypt a file or folder
Encrypting File System overview
Recovering data
Not require a recovery agent for a domain

# Public Key Policies Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Public Key Policies overview
- Understanding Public Key Policies
- Using Public Key Policies

# Public Key Policies overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Public key policies overview**

You can use the public key policy settings in Group Policy to:

- Have computers automatically submit a certificate request to an enterprise certification authority and install the issued certificate. This is useful for ensuring that computers have the certificates that they need to perform public key cryptographic operations in your organization, for example, for Internet Protocol security (IPSec) or for client authentication. For more information about certificate autoenrollment for computers, see Automatic certificate request settings.

- Create and distribute a certificate trust list (CTL). A certificate trust list is a signed list of root certification authority certificates that an administrator considers reputable for designated purposes such as client authentication or secure e-mail. For example, if you want to trust a certification authority's certificates for IPSec, but not for client authentication, you can implement that trust relationship with a certificate trust list. For more information about certificate trust lists, see Enterprise trust policy.

- Establish common trusted root certification authorities. You can use this policy setting to make computers and users subject to common root certification authorities (in addition to the ones that they already trust individually). It is not necessary to use this policy setting for certification authorities in a domain, because they are already trusted by all users and computers in the domain. This policy is primarily for establishing trust in a root certification authority that is not a part of your organization. For more information about root certification authorities, see Policies to establish trust of root certification authorities.

- Add encrypted data recovery agents, and change the encrypted data recovery policy settings. For more information about this policy setting, see Recovering data. For a general overview of Encrypting File System (EFS), see Encrypting File System overview.

It is not necessary for you to use these public key policy settings in Group Policy to deploy a public key infrastructure in your organization. However, these settings give you additional flexibility and control when you establish trust in certification authorities, issue certificates to computers, and deploy EFS across a domain.

# Understanding Public Key Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding public key policies**

You can create public key policies for distribution to computers or users within a site, domain, or organizational unit. To do this, use Group Policy, with the Security Settings extension enabled. This section covers:

- Automatic certificate request settings
- Policies to establish trust of root certification authorities
- Enterprise trust policy
- Recovering data

# Automatic certificate request settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automatic certificate request settings**

Certificate enrollment is the process of requesting, receiving, and installing a certificate. By using automatic certificate settings in public key policies, you can have computers that are associated with a Group Policy object (GPO) automatically enroll for certificates. This can save you the step of explicitly enrolling for computer-related certificates for each computer.

After you establish an automatic certificate request, the actual certificate requests occur the next time the computers associated with the GPO log on to the network.

For more information, see Automatic certificate request policy and Public Key Policies overview.

# Policies to establish trust of root certification authorities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policies to establish trust of root certification authorities**

When a client presents a certificate to a host, the host has to trust the certificate of the root certification authority (CA) in the certification path to accept the certificate as a valid credential. You might want to establish trust automatically in specific root CAs for groups of users or computers.

You can use the public key policy in Group Policy to establish common trusted root CAs for the users and computers associated with a Group Policy object (GPO). When you apply the GPO to a site, domain, or organizational unit, the corresponding computers inherit the policy. These computers then trust the root CAs whose certificates you imported into the trusted root certification authority policy.

You have the option of designating trusted CAs by using either the trusted root certification authority policy or the enterprise trust policy. Use the following guidelines in determining which policy to use:

- If your organization has its own Windows 2000 root CAs and uses Active Directory, you do not need to use the Group Policy mechanism to distribute the root certificates.
- If your organization has its own root CAs that are installed on servers not running Windows 2000, use the trusted root certification authority policy to distribute your organization's root certificates. For more information, see Trusted root certification authority policy.
- If your organization does not have its own CAs, use the enterprise trust policy to create certificate trust lists (CTLs) to establish your organization's trust of external root CAs. For more information, see Enterprise trust policy.

For more information, see Public Key Policies overview.

# Enterprise trust policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enterprise trust policy**
A certificate trust list (CTL) enables you to control trust of the purpose and of the validity period of certificates issued by external certification authorities (CAs).

Typically, a certification authority can issue certificates for a wide variety of purposes, such as secure e-mail or client authentication. But there might be situations in which you want to limit the trust of certificates issued by a particular certification authority, especially if the CA is outside your organization. In these situations, creating a CTL and using it through Group Policy can be useful.

Suppose, for example, a certification authority named "My CA" is capable of issuing certificates for server authentication, client authentication, code signing, and secure e-mail. However, you only want to trust certificates issued by My CA for the purpose of client authentication. You can create a CTL and limit the purpose for which you trust certificates issued by My CA so that they are only valid for client authentication. Any certificates issued for another purpose by My CA are not accepted for use by any computer or user in the scope of the Group Policy object (GPO) to which the CTL is applied.

There can be multiple CTLs in an organization. Because the uses and trusts of certificates for particular domains or organizational units might be different, you can create separate CTLs to reflect these uses and assign particular CTLs to particular GPOs.

Through the use of Group Policy in your organization, you have the option of designating trust in CAs by using either the trusted root certification authority policy or the enterprise trust policy (CTLs). Use the following guidelines in determining which policy to use:

- If your organization has its own root CAs and uses Active Directory, you do not need to use the Group Policy mechanism to distribute those root certificates.

- If your organization has its own root CAs that are not installed on servers, you should use the trusted root certification authority policy to distribute your organization's root certificates. For more information, see Trusted root certification authority policy.

- If your organization does not have its own CAs, use the enterprise trust policy to create CTLs to establish your organization's trust of external root CAs. For more information, see Using enterprise trust policy.

For more information, see Public Key Policies overview.

# Recovering data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovering data**

Data recovery is important when you need to be able to recover data encrypted by an employee after the employee leaves, or when the user loses the private key. Data recovery is available through the Encrypting File System (EFS) as a part of the overall security policy for the system. For example, if you should ever lose your file encryption certificate and associated private key through disk failure, arson, or any other reason, the person who is the designated recovery agent can recover the data. In a business environment, an organization can recover data encrypted by an employee after the employee leaves.

**Recovery policy**

EFS uses recovery policies to provide built-in data recovery. A *recovery policy* is a type of public key policy that provides for one or more user accounts to be designated as recovery agents.

A recovery policy is configured locally for stand-alone computers. For computers that are part of a network, a recovery policy is configured at the domain, organizational unit, or individual computer level, and applies to all Windows XP and Windows Server 2003 family-based computers that the policy applies to. A certification authority (CA) issues recovery certificates, and you use Certificates in Microsoft Management Console (MMC) to manage them.

In a domain, the Windows Server 2003 family implements a default recovery policy for the domain when the first domain controller is set up. The self-signed certificate is issued to the domain administrator. That certificate designates the domain administrator as the recovery agent. To change the default recovery policy for a domain, log on to the first domain controller as an administrator. Additional recovery agents can be added to this policy and the original recovery agent can be removed at any time.

Because the Windows XP and Windows Server 2003 family security subsystems handle enforcing, replicating, and caching of the recovery policy, users can implement file encryption on a system that is temporarily offline, such as a portable computer. This process is similar to logging on to their domain account using cached credentials. For more information, see Change the recovery policy for the local computer and Create a recovery policy for a domain.

**Recovery agents**

A recovery agent is an individual authorized to decrypt data that was encrypted by another user. Recovery agents do not need any other permissions to function in this role. Recovery agents are useful, for example, when employees leave the company and their remaining data needs to be decrypted. Before you can add a recovery agent for a domain, you must ensure that each recovery agent has been issued an X.509v3 certificate.

Each recovery agent has a special certificate and associated private key that allows data recovery wherever the recovery policy applies. If you are the recovery agent, you should be sure to use the **Export** command in Certificates in MMC to back up the recovery certificate and the associated private key to a secure location. After backing them up, you should use Certificates in MMC to delete the recovery certificate. Then, when you need to perform a recovery operation for a user, you should first restore the recovery certificate and associated private key using the **Import** command from Certificates in MMC. After recovering the data, you should again delete the recovery certificate. You do not have to repeat the export process.

To add recovery agents for a domain, you add their certificates to the existing recovery policy. For steps on how to add recovery agents to a domain, see Add a recovery agent for a domain.

### Note

- Recovery agent information that has been added and removed is not automatically updated on existing EFS files. The information in these files is updated the next time the file is accessed. New files always use the current recovery agent information.

# Using Public Key Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using public key policies**

This section covers:

- Automatic certificate request policy
- Trusted root certification authority policy
- Using enterprise trust policy
- Encrypted data recovery agents policy

# Automatic certificate request policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automatic certificate request policy**

To create an automatic certificate request, you need the following information:

- The certificate type that you want computers that are associated with the Group Policy object (GPO) to automatically enroll for. Computers can automatically request only certificate types that have a purpose associated with computers, not with users. Examples of computer-related certificates include:
    - Computer certificates
    - Internet Protocol security (IPSec) certificates
    - Web Server certificates
- The certification authority (CA) that can issue the certificate.

  You can select from a predefined list of CAs. The CA that you select must be capable of issuing a certificate of the type and purpose that you have specified.

For more information, see Automatic certificate request settings, Create an automatic certificate request for computers in a Group Policy object, and Delete an automatic certificate request for computers in a Group Policy object.

# Trusted root certification authority policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Trusted root certification authority policy**

To establish a trusted root certification authority using Group Policy, the Group Policy object (GPO) that you create must have access to the root certificate. This requires that you import a copy of the root certification authority certificate.

You can do this by using the procedure in Add a trusted root certification authority to a Group Policy object.

For a root certification authority certificate to be imported, the root certificate must be in a PKCS #12 file, in a PKCS #7 file, or in binary-encoded X.509 v3 certificate files. For more information about these file formats, see Importing and exporting certificates.

For more information on trusted root certification authority policies, see Policies to establish trust of root certification authorities and Delete a trusted root certification authority from a Group Policy object.

# Using enterprise trust policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using enterprise trust policy**

You can assign a certificate trust list (CTL) to a Group Policy object (GPO) either by creating a new CTL or by importing an existing one. If you are creating a new CTL, you must sign the CTL with a certificate issued for this purpose by a certification authority (CA).

Creating a CTL also requires:

- The root certificates that you want to include in the CTL.
- The purposes for which you want to trust the root certificates in the CTL.

For more information, see Enterprise trust policy, Create a certificate trust list for a Group Policy object, Edit a certificate trust list for a Group Policy object, and Delete a certificate trust list for a Group Policy object.

# Encrypted data recovery agents policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encrypted data recovery agents policy**

You can manage encrypted data recovery policy for computers in a domain by using the following procedures:

- Add a recovery agent for a domain
- Create a recovery policy for a domain

You can manage encrypted data recovery policy for a local computer by using the following procedures:

- Add a recovery agent for the local computer
- Change the recovery policy for the local computer

For more information, see Encrypting File System overview and Recovering data.

# Software Restriction Policies

Updated: November 22, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This product help topic links to checklists, best practices, procedures, and troubleshooting for Windows Server 2003.

**Software restriction policies**

| **Tip** |
| --- |
| For updated information about Software Restriction Policies, including how it works and troubleshooting for Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003, see Software Restriction Policies. |

You can use software restriction policies to identify software and to control its ability to run on your local computer, organizational unit, domain, or site.

- Before implementing software restriction policies to prevent an e-mail virus from running on your computer, see Checklist: Using software restriction policies to help protect your computer against an e-mail virus.
- For tips about using software restriction policies, see Software Restriction Policies Best Practices.
- For help with specific tasks, see Software Restriction Policies How To....
- For general background information, see Software Restriction Policies Concepts.
- For problem-solving instructions, see Software Restriction Policies Troubleshooting.

# Checklist: Using software restriction policies to help protect your computer against an e-mail virus

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Using software restriction policies to help protect your computer against an e-mail virus**

| Step | Reference |
|---|---|
| Review best practices for software restriction policies. | Software Restriction Policies Best Practices |
| Open Software Restriction Policies. | Open Software Restriction Policies |
| If you have not previously defined software restriction policies, create new software restriction policies. | Create new software restriction policies |
| Create a path rule for the folder that your e-mail program uses to run e-mail attachments, and then set the security level to **Disallowed**. | Create a path rule |
| Specify the file types to which the rule applies. | Add or delete a designated file type |
| Modify policy settings so that they apply to the users and groups that you want:<br>• Specify users or groups to which you do not want the Group Policy object's (GPO) policy settings to apply.<br>• Exclude local administrators from the software restriction policies of a specific policy setting in Group Policy and still have the rest of Group Policy apply to the administrators. | Group Policy (pre-GPMC); Prevent software restriction policies from applying to local administrators |

# Software Restriction Policies Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Do not modify the default domain policy.**

- If you do not edit the default domain policy, you always have the option of reapplying the default domain policy if something goes wrong with your customized domain policy.

**Create a separate Group Policy object for software restriction policies.**

- If you create a separate Group Policy object (GPO) for software restriction policies, you can disable software restriction policies in an emergency without disabling the rest of your domain policy.

  For more information, see Group Policy (pre-GPMC).

**If you experience problems with applied policy settings, restart Windows in Safe Mode.**

- Software restriction policies do not apply when Windows is started in Safe Mode. If you accidentally lock down a workstation with software restriction policies, restart the computer in Safe Mode, log on as a local administrator, modify the policy, run **gpupdate**, restart the computer, and then log on normally.

  For more information about restarting the computer in Safe Mode, see Start the computer in Safe Mode. For more information about **gpupdate**, see Gpupdate.

**Use caution when defining a default setting of Disallowed.**

- When you define a default setting of **Disallowed**, all software is disallowed except for software that has been explicitly allowed. Any file that you want to open has to have a software restriction policies rule that allows it to open.

- To protect administrators from locking themselves out of the system, when the default security level is set to **Disallowed**, four registry path rules are automatically created. You can delete or modify these registry path rules; however, this is not recommended.

  For more information, see Setting the default security level to Disallowed.

**For best security, use access control lists in conjunction with software restriction policies.**

- Users might try to circumvent software restriction policies by renaming or moving disallowed files or by overwriting unrestricted files. As a result, it is recommended that you use access control lists (ACLs) to deny users the access necessary to perform these tasks. For information about access control, see Access Control.

**Test new policy settings thoroughly in test environments before applying the policy settings to your domain.**

- New policy settings might act differently than originally expected. Testing diminishes the chance of encountering a problem when you deploy policy settings across your network.

- You can set up a test domain, separate from your organization's domain, in which to test new policy settings. You can also test the policy settings by creating a test GPO and linking it to a test organizational unit. When you have thoroughly tested the policy settings with test users, you can link the test GPO to your domain.

- Do not set programs or files to **Disallowed** without testing to see what the effect may be. Restrictions on certain files can seriously affect the operation of your computer or network.

- Information that is entered incorrectly or typing mistakes can result in a policy setting that does not perform as expected. Testing new policy settings before applying them can prevent unexpected behavior.

**Filter user policy settings based on membership in security groups.**

- You can specify users or groups for which you do not want a policy setting to apply by clearing the **Apply Group Policy** and **Read** check boxes, which are located on the **Security** tab of the properties dialog box for the GPO.

- When the Read permission is denied, the policy setting is not downloaded by the computer. As a result, less bandwidth is consumed by downloading unnecessary policy settings, which enables the network to function more quickly. To deny the Read permission, select **Deny** for the **Read** check box, which is located on the **Security** tab of the properties dialog box for the GPO.

  For more information, see Group Policy (pre-GPMC).

**Do not link to a GPO in another domain or site.**

- Linking to a GPO in another domain or site can result in poor performance.

# Software Restriction Policies How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Create new software restriction policies
- Open Software Restriction Policies
- Apply software restriction policies to DLLs
- Prevent software restriction policies from applying to local administrators
- Enable certificate rules
- Create a certificate rule
- Create a hash rule
- Create an Internet zone rule
- Create a path rule
- Create a registry path rule
- Add or delete a designated file type
- Change the default security level of software restriction policies
- Set trusted publisher options

# Create new software restriction policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create new software restriction policies**

1. Open Software Restriction Policies.

2. On the **Action** menu, click **New Software Restriction Policies**.

## Notes

- Different administrative credentials are required to perform this procedure, depending on your environment:

    - If you create new software restriction policies for your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

    - If you create new software restriction policies for a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.

- If software restriction policies have already been created for a Group Policy object (GPO), the **New Software Restriction Policies** command does not appear on the **Action** menu. To delete the software restriction policies that are applied to a GPO, in the console tree, right-click **Software Restriction Policies**, and then click **Delete Software Restriction Policies**. When you delete software restriction policies for a GPO, you also delete all software restriction policies rules for that GPO. After you delete software restriction policies, you can create new software restriction policies for that GPO.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Software restriction policies overview
Software Restriction Policies

# Open Software Restriction Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open Software Restriction Policies**

- For your local computer
- For a domain, site, or organizational unit, and you are on a member server or on a workstation that is joined to a domain
- For a domain or organizational unit, and you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed
- For a site, and you are on a domain controller or on a workstation that has the Administration Tools Pack installed

**For your local computer**

1. Open Local Security Settings.
2. In the console tree, click **Software Restriction Policies**.

   **Where?**

   - Security Settings/Software Restriction Policies

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**For a domain, site, or organizational unit, and you are on a member server or on a workstation that is joined to a domain**

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. In **Select Group Policy Object**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit--or create a new one, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree, click **Software Restriction Policies**.

   **Where?**

- *Group Policy Object* [*ComputerName*] Policy/Computer Configuration or

  User Configuration/Windows Settings/Security Settings/Software Restriction Policies

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.

**For a domain or organizational unit, and you are on a domain controller or on a workstation that has the Windows Server 2003 Administration Tools Pack installed**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the Group Policy object (GPO) that you want to open software restriction policies for.

3. Click **Properties**, and then click the **Group Policy** tab.

4. Click **Edit** to open the GPO that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.

5. In the console tree, click **Software Restriction Policies**.

   **Where?**

   - *Group Policy Object* [*ComputerName*] Policy/Computer Configuration or

     User Configuration/Windows Settings/Security Settings/Software Restriction Policies

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**For a site, and you are on a domain controller or on a workstation that has the Administration Tools Pack installed**

1. Open Active Directory Sites and Services.

2. In the console tree, right-click the site that you want to set Group Policy for.

   **Where?**

   - Active Directory Sites and Services [*Domain_Controller_Name.Domain_Name*]/Sites/Site

3. Click **Properties**, and then click the **Group Policy** tab.

4. Click an entry in **Group Policy Object Links** to select an existing Group Policy object (GPO), and then click **Edit**. You can also click **New** to create a new GPO, and then click **Edit**.

5. In the console tree, click **Software Restriction Policies**.

   **Where**

   - *Group Policy Object* [*ComputerName*] Policy/Computer Configuration or

     User Configuration/Windows Settings/Security Settings/Software Restriction Policies

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

## Notes

- To set policy settings that will be applied to computers, regardless of which users log on to them, click **Computer Configuration**.

- To set policy settings that will be applied to users, regardless of which computer they log on to, click **User Configuration**.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a certificate rule
Create a hash rule
Create a path rule
Create an Internet zone rule
Software Restriction Policies
Windows Server 2003 Administration Tools Pack

# Apply software restriction policies to DLLs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To apply software restriction policies to DLLs**

1. Open Software Restriction Policies.
2. In the details pane, double-click **Enforcement**.
3. Under **Apply software restriction policies to the following**, click **All software files**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- By default, software restriction policies do not check dynamic-link libraries (DLLs). Checking DLLs can decrease system performance, because software restriction policies must be evaluated every time a DLL is loaded. However, you may decide to check DLLs if you are concerned about receiving a virus that targets DLLs. If the default security level is set to **Disallowed**, and you enable DLL checking, you must create software restriction policies rules that allow each DLL to run. For more information about rules, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Software Restriction Policies
Security levels and additional rules

# Prevent software restriction policies from applying to local administrators

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prevent software restriction policies from applying to local administrators**

1. Open Software Restriction Policies.
2. In the details pane, double-click **Enforcement**.
3. Under **Apply software restriction policies to the following users**, click **All users except local administrators**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- It may be necessary to create a new software restriction policy setting for the Group Policy object (GPO) if you have not already done so. For information about how to create new software restriction policies, see Related Topics.
- If it is common for users to be members of the local Administrators group on their computers in your organization, you may not want to enable this option.
- If you are defining a software restriction policy setting for your local computer, use this procedure to prevent local administrators from having software restriction policies applied to them. If you are defining a software restriction policy setting for your network, filter user policy settings based on membership in security groups through Group Policy. For more information about Group Policy, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Create new software restriction policies
Software Restriction Policies
Group Policy (pre-GPMC)

# Enable certificate rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable certificate rules**

- For your local computer
- For a Group Policy object, and you are on a server that is joined to a domain
- For a Group Policy object, and you are on a domain controller or a on workstation that has the Windows Server 2003 Administration Tools Pack installed
- For only domain controllers, and you are on a domain controller or on a workstation that has the Administration Tools Pack installed

**For your local computer**

1. Open Local Security Settings.
2. In the console tree, click **Security Options**.

   **Where?**

   - Security Settings/Local Policies/Security Options
3. In the details pane, double-click **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**.
4. Do one of the following, and then click **OK**:

   - To enable certificate rules, click **Enabled**.
   - To disable certificate rules, click **Disabled**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**For a Group Policy object, and you are on a server that is joined to a domain**

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.

4. In **Select Group Policy Object**, click **Browse**.

5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit--or create a new one, and then click **Finish**.

6. Click **Close**, and then click **OK**.

7. In the console tree, click **Security Options**.

   **Where?**

   - *GroupPolicyObject* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/Local Policies/Security Options

8. In the details pane, double-click **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**.

9. If this policy setting has not yet been defined, select the **Define these policy settings** check box.

10. Do one of the following, and then click **OK**:

    - To enable certificate rules, click **Enabled**.

    - To disable certificate rules, click **Disabled**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.

**For a Group Policy object, and you are on a domain controller or a on workstation that has the Windows Server 2003 Administration Tools Pack installed**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the Group Policy object (GPO) for which you want to enable certificate rules.

3. Click **Properties**, and then click the **Group Policy** tab.

4. Click **Edit** to open the GPO that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.

5. In the console tree, click **Security Options**.

   **Where?**

   - *GroupPolicyObject* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/Local Policies/Security Options

6. In the details pane, double-click **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**.

7. If this policy setting has not yet been defined, select the **Define these policy settings** check box.

8. Do one of the following, and then click **OK**:

    - To enable certificate rules, click **Enabled**.

    - To disable certificate rules, click **Disabled**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**For only domain controllers, and you are on a domain controller or on a workstation that has the Administration Tools Pack installed**

1. Open Domain Controller Security Settings.

2. In the console tree, click **Security Options**.

    **Where?**

    - *GroupPolicyObject* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/Local Policies/Security Options

3. In the details pane, double-click **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**.

4. If this policy setting has not yet been defined, select the **Define these policy settings** check box.

5. Do one of the following, and then click **OK**:

    - To enable certificate rules, click **Enabled**.

    - To disable certificate rules, click **Disabled**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Domain Controller Security Policy, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Domain Controller Security Policy**.

**Note**

- You must perform this procedure before certificate rules can take effect.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Create a certificate rule
Security levels and additional rules
Registry Editor

# Create a certificate rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a certificate rule**

1. Open Software Restriction Policies.
2. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Certificate Rule**.
3. Click **Browse**, and then select a certificate or signed file.
4. In **Security level**, click either **Disallowed** or **Unrestricted**.
5. In **Description**, type a description for this rule, and then click **OK**.

## Notes

- Different administrative credentials are required to perform this procedure, depending on your environment:
    - If you create a certificate rule for your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
    - If you create a certificate rule for a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- It might be necessary to create a new software restriction policy setting for the Group Policy object (GPO) if you have not already done so. For information about how to create new software restriction policies, see Related Topics.
- Certificate rules are not enabled by default. For information about how to enable certificate rules, see "Enable certificate rules" in Related Topics.
- The only file types that are affected by certificate rules are those that are listed in **Designated file types** in the details pane for Software Restriction Policies. There is one list of designated file types that is shared by all rules. For more information, see "Add or delete a designated file type" in Related Topics.
- For software restriction policies to take effect, users must update policy settings by logging off from and logging on to their computers.
- When more than one software restriction policies rule is applied to policy settings, there is a precedence of rules for handling conflicts. For more information, see "Precedence of software restriction policies" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Open Software Restriction Policies
Enable certificate rules
Create new software restriction policies
Security levels and additional rules
Software Restriction Policies
Add or delete a designated file type
Precedence of software restriction policies rules
Certificates overview
Security with Certificates

# Create a hash rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a hash rule**

1. Open Software Restriction Policies.
2. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Hash Rule**.
3. Click **Browse** to find a file, or paste a precalculated hash in **File hash**.
4. In **Security level**, click either **Disallowed** or **Unrestricted**.
5. In **Description**, type a description for this rule, and then click **OK**.

**Notes**

- Different administrative credentials are required to perform this procedure, depending on your environment:
    - If you create a hash rule on your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
    - If you create a hash rule on a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- It may be necessary to create a new software restriction policy setting for the Group Policy object (GPO) if you have not already done so. For information about how to create new software restriction policies, see Related Topics.
- A hash rule can be created for a virus or a Trojan horse to prevent them from running.
- If you want other people to use a hash rule so that a virus cannot run, calculate the hash of the virus by using software restriction policies, and then e-mail the hash value to the other people. Never e-mail the virus itself.
- If a virus has been sent through e-mail, you can also create a path rule to prevent execution of e-mail attachments. For more information, see "Create a path rule" in Related Topics.
- A file that is renamed or moved to another folder results in the same hash. Any change to the file itself results in a different hash.
- The only file types that are affected by hash rules are those that are listed in **Designated File Types** in the details pane for Software Restriction Policies. There is one list of designated file types that is shared by all rules. For more information, see "Add or delete a designated file type" in Related Topics.
- For software restriction policies to take effect, users must update policy settings by logging off from and logging on to their computers.
- When more than one software restriction policies rule is applied to policy settings, there is a precedence of rules for

handling conflicts. For more information see, "Precedence of software restriction policies" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Create new software restriction policies
Software Restriction Policies
Security levels and additional rules
Create a path rule
Add or delete a designated file type
Precedence of software restriction policies rules

# Create an Internet zone rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an Internet zone rule**

1. Open Software Restriction Policies.
2. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Internet Zone Rule**.
3. In **Internet zone**, click an Internet zone.
4. In **Security level**, click either **Disallowed** or **Unrestricted**, and then click **OK**.

## Notes

- Different administrative credentials are required to perform this procedure, depending on your environment:
  - If you create an Internet zone rule for your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
  - If you create an Internet zone rule for a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- It may be necessary to create a new software restriction policy setting for the Group Policy object (GPO) if you have not already done so. For information about how to create new software restriction policies, see Related Topics.
- Zone rules only apply to files with an .msi file type, which are Windows Installer packages.
- For software restriction policies to take effect, users must update policy settings by logging off from and logging on to their computers.
- When more than one software restriction policies rule is applied to policy settings, there is a precedence of rules for handling conflicts. For more information see, "Precedence of software restriction policies" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Create new software restriction policies
Software Restriction Policies
Security levels and additional rules
Windows Installer overview
Precedence of software restriction policies rules

# Create a path rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a path rule**

1. Open Software Restriction Policies.
2. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Path Rule**.
3. In **Path**, type a path, or click **Browse** to find a file or folder.
4. In **Security level**, click either **Disallowed** or **Unrestricted**.
5. In **Description**, type a description for this rule, and then click **OK**.

**Caution**

- On certain folders, such as the Windows folder, setting the security level to **Disallowed** can adversely affect the operation of your operating system. Make sure that you do not disallow a crucial component of the operating system or one of its dependent programs.

**Notes**

- Different administrative credentials are required to perform this procedure, depending on your environment:
    - If you create a path rule for your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
    - If you create a path rule for a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- It may be necessary to create new software restriction policies for the Group Policy object (GPO) if you have not already done so. For information about how to create new software restriction policies, see Related Topics.
- If you create a path rule for software with a security level of **Disallowed**, users can still run the software by copying it to another location.
- The wildcard characters that are supported by the path rule are * and ?.
- You can use environment variables, such as %programfiles% or %systemroot%, in the path rule.
- If you want to create a path rule for software when you do not know where it is stored on a computer but you have its registry key, you can create a registry path rule. For more information about how to create a registry path rule, see Related Topics.
- To prevent users from executing e-mail attachments, you can create a path rule for your e-mail program's attachment directory that prevents users from running e-mail attachments.

- The only file types that are affected by path rules are those that are listed in **Designated File Types** in the details pane for Software Restriction Policies. There is one list of designated file types that is shared by all rules. For more information, see "Add or delete a designated file type" in Related Topics.

- For software restriction policies to take effect, users must update policy settings by logging off from and logging on to their computers.

- When more than one software restriction policies rule is applied to policy settings, there is a precedence of rules for handling conflicts. For more information, see "Precedence of software restriction policies" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Create new software restriction policies
Create a registry path rule
Security levels and additional rules
Software Restriction Policies
Add or delete a designated file type
Precedence of software restriction policies rules
Set, view, change, or remove permissions on files and folders

# Create a registry path rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a registry path rule**

1. Click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. In the console tree, right-click the registry key that you want to create a rule for, and then click **Copy Key Name**. Note the value name in the details pane.
3. Open Software Restriction Policies.
4. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Path Rule**.
5. In **Path**, paste the registry key name, followed by the value name.
6. Enclose the registry path in percent signs (%), for example, **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\PlatformSDK\Directories\InstallDir%**.
7. In **Security level**, click either **Disallowed** or **Unrestricted**.
8. In **Description**, type a description for this rule, and then click **OK**.

## Notes

- Different administrative credentials are required to perform this procedure, depending on your environment:
  - To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
  - If you create a registry path rule for a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- It may be necessary to create a new software restriction policy setting for the Group Policy object (GPO) if you have not already done so. For information about how to create new software restriction policies, see Related Topics.
- Format the registry path as follows: **%[Registry Hive]\[Registry Key Name]\[Value Name]%**
- The name of the registry hive has to be written out; abbreviations do not work. For example, **HKCU** cannot be substituted for **HKEY_CURRENT_USER**.
- The registry path rule can contain a suffix after the closing percent sign. Do not use a backslash (\) in the suffix. For example, **%HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\Shell Folders\Cache%OLK*** is valid.
- The only file types that are affected by path rules are those that are listed in **Designated File Types** in the details pane for Software Restriction Policies. There is one list of designated file types that is shared by all rules. For more information,

see "Add or delete a designated file type" in Related Topics.

- To refresh software restriction policies, you must log off from and then log on to your computer.
- When more than one software restriction policies rule is applied to policy settings, there is a precedence of rules for handling conflicts. For more information, see "Precedence of software restriction policies" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Create new software restriction policies
Create a path rule
Software Restriction Policies
Add or delete a designated file type
Precedence of software restriction policies rules
**Other Resources**
Registry Editor overview

# Add or delete a designated file type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add or delete a designated file type**

1. Open Software Restriction Policies.
2. In the details pane, double-click **Designated File Types**.
3. Do one of the following:
   - To add a file type, in **File extension**, type the file name extension, and then click **Add**.
   - To delete a file type, in **Designated file types**, click the file type, and then click **Remove**.

## Notes

- Different administrative credentials are required to perform this procedure, depending on the environment in which you add or delete a designated file type:
  - If you add or delete a designated file type for your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
  - If you add or delete a designated file type for a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- It may be necessary to create a new software restriction policy setting for the Group Policy object (GPO) if you have not already done so. For information about how to create new software restriction policies, see Related Topics.
- The list of designated file types is shared by all rules for both Computer Configuration and User Configuration for a GPO.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Create new software restriction policies
Software Restriction Policies

# Change the default security level of software restriction policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the default security level of software restriction policies**

1. Open Software Restriction Policies.
2. In the details pane, double-click **Security Levels**.
3. Right-click the security level that you want to set as the default, and then click **Set as default**.

### Caution

- In certain directories, setting the default security level to **Disallowed** can adversely affect your operating system. For more information about setting your default security setting to **Disallowed**, see Related Topics.

### Notes

- Different administrative credentials are required to perform this procedure, depending on the environment for which you change the default security level of software restriction policies:
- If you change the default security level of software restriction policies on your local computer: To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- If you change the default security level of software restriction policies on a computer that is joined to a domain: To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- It may be necessary to create a new software restriction policy setting for this Group Policy object (GPO) if you have not already done so. For information about how to create new software restriction policies, see Related Topics.
- In the details pane, the current default security level is indicated by a black circle with a check mark in it. If you right-click the current default security level, the **Set as default** command does not appear in the menu.
- Software restriction policies rules are created to specify exceptions to the default security level. When the default security level is set to **Unrestricted**, rules can specify software that is not allowed to run. When the default security level is set to **Disallowed**, rules can specify software that is allowed to run.
- At installation, the default security level of software restriction policies on all files on your system is set to **Unrestricted**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Setting the default security level to Disallowed
Software Restriction Policies
Open Software Restriction Policies
Create new software restriction policies

# Set trusted publisher options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set trusted publisher options**

1. Open Software Restriction Policies.
2. Double-click **Trusted Publishers**.
3. Click the users that you want to decide which certificates will be trusted, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Software Restriction Policies, see "Open Software Restriction Policies" in Related Topics.
- It may be necessary to create a new software restriction policy setting for the Group Policy object (GPO) if you have not already done so. For information about how to create new software restriction policies, see Related Topics.
- You can use this procedure to determine whether users, local administrators, or enterprise administrators can add trusted publishers. For example, you can use this tool to prevent users from making trust decisions about publishers of ActiveX controls.
- Local computer administrators have the right to specify trusted publishers on the local computer, while enterprise administrators have the right to specify trusted publishers on an organizational-unit level.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open Software Restriction Policies
Create new software restriction policies
Software Restriction Policies

# Software Restriction Policies Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Software restriction policies overview
- Applying software restriction policies
- Security levels and additional rules
- Precedence of software restriction policies rules
- Setting the default security level to Disallowed
- System events for software restriction policies

# Software restriction policies overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Software restriction policies overview**

Software restriction policies address the need to regulate unknown or untrusted software. With the rise in the use of networks, the Internet, and e-mail for business computing, users find themselves exposed to new software in a variety of ways. Users must constantly make decisions about running unknown software. Viruses and Trojan horses often intentionally misrepresent themselves to trick users into running them. It is difficult for users to make safe choices about what software they should run.

With software restriction policies, you can protect your computing environment from untrusted software by identifying and specifying what software is allowed to run. You can define a default security level of **Unrestricted** or **Disallowed** for a Group Policy object (GPO) so that software is either allowed or not allowed to run by default. You can make exceptions to this default security level by creating software restriction policies rules for specific software. For example, if the default security level is set to **Disallowed**, you can create rules that allow specific software to run. The types of rules are as follows:

- Hash rules
- Certificate rules
- Path rules (including registry path rules)
- Internet zone rules

For more information about each of these rules, see Security levels and additional rules.

Software restriction policies consist of the default security level and all the rules that apply to a GPO. Software restriction policies can be applied across a domain, to local computers, or to individual users. Software restriction policies provide a number of ways to identify software, and they provide a policy-based infrastructure to enforce decisions about whether the identified software can run. With software restriction policies, when users execute software programs, they must adhere to the guidelines that are set up by administrators.

With software restriction policies, you can:

- Control the ability of software to run on your system. For example, if you are concerned about users receiving viruses through e-mail, you can apply a policy setting that does not allow certain file types to run in the e-mail attachment directory of your e-mail program.
- Permit users to run only specific files on multiuser computers. For example, if you have multiple users on your computers, you can set up software restriction policies in such a way that users do not have access to any software except the specific files that are necessary for their work.
- Decide who can add trusted publishers to your computer.
- Control whether software restriction policies affect all users or just certain users on a computer.
- Prevent any files from running on your local computer, organizational unit, site, or domain. For example, if your system has a known virus, you can use software restriction policies to stop a computer from opening the file that contains the virus.

**Important**

- Software restriction policies should not be used as a replacement for antivirus software.

For more information, see:

- Software Restriction Policies
- Create a certificate rule
- Create a hash rule
- Create an Internet zone rule
- Create a path rule
- Create a registry path rule

# Applying software restriction policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Applying software restriction policies**

Software restriction policies use rules to identify and control how software runs. You can identify a software program by its hash, certificate, or path or by the Internet zone where it resides. After the software is identified, you can decide whether or not to allow it to run. For more information, see Security levels and additional rules.

Software restriction policies can apply to computers or users, depending on whether you choose to modify settings in Computer Configuration or in User Configuration. For more information, see Open Software Restriction Policies

You apply software restriction policies through Group Policy. You apply the policy settings to a Group Policy object (GPO), which is linked to your local computer, site, domain, or organizational unit. If more than one policy setting applies, the policy settings are in the following order of precedence, from lowest to highest:

- Local computer policy
- Site policy
- Domain policy
- Organizational-unit policy

For example, for a workstation that is joined to a domain, the domain policy settings override the local security policy settings for the workstation's software restriction policies wherever there is a conflict. Likewise, if the same workstation is a member of an organizational unit, the policy settings that are applied from the organizational unit policy override both the domain policy settings and the local policy settings wherever there is a conflict. If the workstation is a member of more than one organizational unit, the organizational unit that immediately contains the workstation has the highest precedence. For more information, see Group Policy (pre-GPMC).

To check the precedence of policy when more than one GPO is applied, you can use Resultant Set of Policy (RSoP). For more information, see Resultant Set of Policy.

Each policy setting is refreshed when you restart the computer. When you modify the policy settings, they are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The policy settings are also refreshed every 16 hours, whether or not there are any changes. You can refresh policy settings by running the command-line utility **gpupdate** and then logging off from and logging on to your computer. For more information, see Gpupdate.

# Security levels and additional rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security levels and additional rules**

For software restriction policies, the security level options are:

- **Unrestricted**, which allows software to run with the full rights of the user who is logged on to the computer.
- **Disallowed**, which does not allow the software to run.

You can use the following rules to make exceptions to the default security level for software restriction policies. For example, if the default security level is **Disallowed**, you can create a path rule that makes an exception to the default security level and allows a software program to run.

More than one rule can be applied to software. When this happens, the rule with the highest precedence determines if the software will run or not run. For information about the precedence of these rules, see Precedence of software restriction policies rules.

**Hash rule**

A hash is a series of bytes with a fixed length that uniquely identifies a software program or file. The hash is computed by a hash algorithm. When a hash rule is created for a software program, software restriction policies calculate a hash of the program. When a user tries to open a software program, a hash of the program is compared to existing hash rules for software restriction policies. The hash of a software program is always the same, regardless of where the program is located on the computer. However, if a software program is altered in any way, its hash also changes, and it no longer matches the hash in the hash rule for software restriction policies.

For example, you can create a hash rule and set the security level to **Disallowed** to prevent users from running a certain file. A file can be renamed or moved to another folder and still result in the same hash. However, any changes to the file itself also change its hash value and allow the file to bypass restrictions.

For more information, see Create a hash rule.

**Certificate rule**

Software restriction policies can also identify software by its signing certificate. You can create a certificate rule that identifies software and then allows or does not allow the software to run, depending on the security level. For example, you can use certificate rules to automatically trust software from a trusted source in a domain without prompting the user. You can also use certificate rules to run files in disallowed areas of your operating system.

Certificate rules are not enabled by default. For information about how to enable certificate rules, see Enable certificate rules.

For information on how to create a certificate rule, see Create a certificate rule. For more information about certificates, see Certificates.

**Path rule**

A path rule identifies software by its file path. For example, if you have a computer that has a default security level of **Disallowed**, you can still grant unrestricted access to a specific folder for each user. You can create a path rule by using the file path and setting the security level of the path rule to **Unrestricted**. Some common paths for this type of rule are %userprofile%, %windir%, %appdata%, %programfiles%, and %temp%. You can also create registry path rules that use the registry key of the software as its path.

Because these rules are specified by the path, if a software program is moved, the path rule no longer applies.

For more information about how to create a path rule, see Create a path rule. For more information on how to create a registry path rule, see Create a registry path rule.

For information about the precedence of path rules when more than one path rule is applied to a file or folder, see Precedence of software restriction policies rules.

**Internet zone rule**

Zone rules apply only to Windows Installer packages. For more information, see Windows Installer overview.

A zone rule can identify software from a zone that is specified through Internet Explorer. These zones are Internet, Intranet, Restricted sites, Trusted sites, and My Computer.

For more information, see Create an Internet zone rule.

# Precedence of software restriction policies rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Precedence of software restriction policies rules**

You can apply several software restriction policies rules to the same software. The rules are applied in the following order of precedence, from highest to lowest:

- Hash rule
- Certificate rule
- Path rule
- Internet zone rule

For example, if you create a hash rule with a security level of **Unrestricted** for a software program that resides in a folder that has a path rule assigned to it with a security level of **Disallowed**, the program will run. The hash rule takes precedence over the path rule.

If two path rules are assigned to the same object, the more specific rule takes precedence. For example, if there is a path rule for C:\Windows\ with a security level of **Disallowed**, but there is also a path rule for C:\Windows\System32\ with a security level of **Unrestricted**, the more specific path rule takes precedence. Software programs in C:\Windows\ will not run, but programs in C:\Windows\System32\ will run.

If two identical rules with differing security levels are applied to software, the more conservative rule takes precedence. For example, if two hash rules--one with a security level of **Disallowed** and one with a security level of **Unrestricted**--are applied to the same software program, the rule with a security level of **Disallowed** takes precedence, and the program will not run.

For more information, see:

- Software Restriction Policies
- Create a hash rule
- Create a certificate rule
- Create a path rule
- Create an Internet zone rule

# Setting the default security level to Disallowed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting the default security level to Disallowed**

You can set software restriction policies so that the operating system has one of two default security levels: **Unrestricted** and **Disallowed**. When you first create software restriction policies, a default setting of **Unrestricted** is automatically defined. You can then apply rules that restrict the execution of specified software programs. When you set the default security level to **Disallowed**, most software is restricted. There are four registry path rules that are automatically created to prevent you from completely locking down your system, including locking yourself out of it. You must then apply rules for all other software that you want to run.

Software restriction policies are applied through the operating system. Certain software programs must be set to **Unrestricted** for the operating system to function at all. The following tips can help you decide which programs to set to **Unrestricted** if you decide to use a default security level of **Disallowed**:

- When the default security level is set to **Disallowed**, four registry path rules are created automatically. They are as follows:

  - **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SystemRoot%**
  - **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SystemRoot%*.exe**
  - **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SystemRoot%\System32*.exe**
  - **%HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\ProgramFilesDir%**

  These registry path rules are created as a safeguard against locking yourself and all users out of the system. Only advanced users should consider modifying or deleting these rules.

- Startup items are put in **HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Run** in the registry. If you want them to run, make sure that you create a registry path rule for them. For information about creating a registry path rule, see Create a registry path rule.

- If a computer must run logon scripts, create a path rule that allows them. For information about creating a path rule, see Create a path rule.

- Explore your software programs when you create rules for them. Many programs start other programs to perform certain tasks. For example, Microsoft Word starts Microsoft Clip Organizer to manage clip art. For optimal performance, it is essential that all supporting programs be allowed to run.

For more information about software restriction policies, see Software Restriction Policies.

# System events for software restriction policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System events for software restriction policies**

When users try to run software that has been disallowed by software restriction policies or by software restriction policies rules, events are generated and stored in the application log, as indicated in the following table. You can view these events with the Event Viewer. For more information about Event Viewer, see Event Viewer.

| Event ID | What it means |
|---|---|
| 865 | A user attempted to run software that is disallowed by the default security level. |
| 866 | A user attempted to run software that is disallowed by a path rule. |
| 867 | A user attempted to run software that is disallowed by a certificate rule. |
| 868 | A user attempted to run software that is disallowed by an Internet zone rule or a hash rule. |

For more information about the event log, see Settings for Event Logs. For more information about software restriction policies, see Software Restriction Policies.

# Software Restriction Policies Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Users receive a message that says "Windows cannot open this program because it has been prevented by a software restriction policy. For more information, open Event Viewer or contact your system administrator." Or, on the command line, a message says "The system cannot execute the specified program."
- Modified software restriction policies are not taking effect.
- You have added a rule to software restriction policies, and you cannot log on to your computer.
- A new policy setting is not applying to a specific file name extension.

**Users receive a message that says "Windows cannot open this program because it has been prevented by a software restriction policy. For more information, open Event Viewer or contact your system administrator." Or, on the command line, a message says "The system cannot execute the specified program."**
**Cause:** The default security level (or a rule) was created so that the software program is set as **Disallowed**, and as a result it will not start.

**Solution:** Look in the event log for an in-depth description of the message. The event log message indicates what software program is set as **Disallowed** and what rule is applied to the program.

**See also:** Software Restriction Policies

**Modified software restriction policies are not taking effect.**
**Cause:** Software restriction policies that are specified in a domain through Group Policy override any policy settings that are configured locally. This may imply that there is a policy setting from the domain that is overriding your policy setting. For more information, see Group Policy (pre-GPMC).

**Cause:** Group Policy may not have refreshed its policy settings. Group Policy applies changes to policy settings periodically; therefore, it is likely that the policy changes that were made in the directory have not yet been refreshed.

**Solution:** You can refresh policy settings with the command-line utility **gpupdate** or by logging off from and then logging back on to your computer. For best results, run **gpupdate**, and then log off from and log back on to your computer. For more information, see Gpupdate.

The security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes.

**Cause:** The local computer on which you changed software restriction policies for the network cannot contact a domain controller.

**Solution:** To successfully modify policy for a network, the computer on which you modify software restriction policies must be able to contact a domain controller to update policy for the network. If the computer on which you are working cannot contact the domain controller, software restriction policies are not applied. If your computer is a member of a domain, local software restriction policies are not applied unless the computer can contact the domain controller to ensure that network policy does not override local policy.

**You have added a rule to software restriction policies, and you cannot log on to your computer.**
**Cause:** Your computer accesses many software programs and files when it starts. It is possible to inadvertently set one of these programs or files to **Disallowed**. Because the computer cannot access the program or file, it cannot start properly.

**Solution:** Start the computer in Safe Mode, log on as a local administrator, and then change software restriction policies to allow the program or file to run.

**A new policy setting is not applying to a specific file name extension.**
**Cause:** The file name extension is not in the list of designated file types for software restriction policies.

**Solution:** Add the file name extension to the list of supported file types.

For more information, see Add or delete a designated file type.

# Auditing Security Events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Auditing security events**
You can set up audit policy so that user or system activity in specified event categories is recorded. You can monitor security-related activity, such as who accesses an object, if a user logs on to or logs off from a computer, or if changes are made to an auditing policy setting.

- Before defining auditing policy settings in the object access event category, see Checklist: Setting up object access auditing.
- For tips about auditing, see Auditing Security Events Best practices.
- For help with specific tasks, see Auditing Security Events How To....
- For general background information, see Auditing Security Events Concepts.

# Checklist: Setting up object access auditing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Setting up object access auditing**

| Step | Reference |
| --- | --- |
| Read the auditing overview. | Auditing overview |
| Enable auditing of the object access event category. | Define or modify auditing policy settings for an event category |
| Apply auditing policy to a specified object. | Apply or modify auditing policy settings for a local file or folder; Apply or modify auditing policy settings for an object using Group Policy |

# Auditing Security Events Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Create an audit plan before implementing audit policy.**

- Decide what type of information you want to gain by collecting audit events:

  - If you are interested in intrusion detection--tracking the attempts of users to gain access to areas for which they are not authorized--you can collect failure audits. But enabling failure audits can be a risk to your organization. If users attempt to access a resource for which they are not authorized, they can create so many failure audits that the security log becomes full, and the computer cannot collect any more audits. If you have the **Audit: Shut down system immediately if unable to log security audits** policy setting enabled, users can start a denial-of-service attack through your audit policy.

  - If you are interested in forensics--using the audit log to determine exactly what happens in your organization--you can collect a combination of success and failure audits.

- Consider the resources that you have available for collecting and reviewing an audit log. Audit events take up space on your computers, and they take up your time and the time of people in your organization. Do not audit events that do not really interest you.

**Collect and archive security logs across your organization.**

- If an intrusion occurs, isolate and preserve the security log entries. These entries can be valuable during an investigation of the intrusion.

- An audit trail can contain information about changes that are made to your computer or to other computers on the network. If intruders gain administrator rights and permissions, or if administrators abuse their rights and permissions, they can clear the security log, leaving you without a trail of their actions. If you use a tool that regularly collects and saves security log entries across your organization, even if intruders or administrators clear the local security log, you are more likely to be able to trace the actions of intruders or administrators. Microsoft Operations Manager is an example of such a tool. For the most recent tools available, see Search Microsoft.com at the Microsoft Web site.

**Audit success and failure events in the system event category.**

- By auditing success and failure events in the system event category, you can notice unusual activity that may indicate that an intruder is attempting to gain access to your computer or your network.

- The number of audits that are generated when this setting is enabled tends to be relatively low, and the quality of information that is gained from the events tends to be relatively high.

  For information about the system event category, see Audit system events.

  For information about how to enable auditing in the system event category, see Define or modify auditing policy settings for an event category.

**Audit success events in the policy change event category on domain controllers.**

- If an event is logged in the policy change event category, someone has changed the Local Security Authority (LSA) security policy configuration.

- If you use Group Policy to edit your auditing policy settings, it is not necessary to audit events in the policy change event

category on member servers.

- If you decide to audit failure events in the policy change event category, you can see if unauthorized users or attackers are trying to change policy settings, including security policy settings. Although this can be helpful for intrusion detection, the increase in resources that is required and the possibility of a denial-of-service attack usually outweigh the benefits.

  For information about the policy change event category, see Audit policy change.

  For information about how to enable auditing in the policy change event category, see Define or modify auditing policy settings for an event category.

**Audit success events in the account management event category.**

- By auditing success events in the account management event category, you can verify changes that are made to account properties and group properties.

- If you decide to audit failure events in the account management event category, you can see if unauthorized users or attackers are trying to change account properties or group properties. Although this can be helpful for intrusion detection, the increase in resources that is required and the possibility of a denial-of-service attack usually outweigh the benefits.

  For information about the account management event category, see Audit account management.

  For information about how to enable auditing in the account management event category, see Define or modify auditing policy settings for an event category.

**Audit success events in the logon event category.**

- By auditing success events in the logon event category, you have a record of when each user logs on to or logs off from a computer. If an unauthorized person steals a user's password and logs on to a computer, you can determine when the breach of security occurred.

- If you decide to audit failure events in the logon event category, you can see if unauthorized users or attackers are trying to log on to a computer. Although this can be helpful for intrusion detection, the possibility of a denial-of-service attack increases with the auditing of failure events in the logon event category. If you have failure auditing enabled in these event categories, users from outside of your organization can fill the security log or cause events to be overwritten by continuously attempting to log on to your network with incorrect usernames or passwords. If you also have the **Audit: Shut down system immediately if unable to log security audits** policy setting enabled, the consequences of logging failure events in these categories can be more severe - the user could cause a denial of service if the security log is filled.

  For information about how to enable auditing in the logon event category, see Define or modify auditing policy settings for an event category.

**Audit success events in the account logon event category on domain controllers.**

- By auditing success events in the account logon event category, you can see when users log on to or log off from the domain.

- It is not necessary to audit events in the account logon event category on member servers.

- If you decide to audit failure events in the account logon event category, you can see if unauthorized users or attackers are trying to log on to your network. Although this can be helpful for intrusion detection, the possibility of a denial-of-service attack increases with the auditing of failure events in the account logon event category. If you have failure auditing enabled in these event categories, users from outside of your organization can fill the security log or cause events to be overwritten by continuously attempting to log on to your network with incorrect usernames or passwords. If you also have the **Audit: Shut down system immediately if unable to log security audits** policy setting enabled, the consequences of logging failure events in these categories can be more severe - the user could cause a denial of service

if the security log is filled.

For information about the account logon event category, see Audit account logon events.

For information about how to enable auditing in the account logon event category, see Define or modify auditing policy settings for an event category.

**Be specific when setting up auditing on an object.**

- To audit object access, you must enable the **Audit object access** policy setting and then edit the system access control list (SACL) that is associated with the object. For more information, see Checklist: Setting up object access auditing.

- For best system performance, minimize the number of entries in the SACL for an object. One entry in a SACL that contains 1000 users does not degrade system performance as much as 1000 separate entries.

- To reduce the volume of events that is generated and to maximize the effectiveness of each event, only audit the actions that really interest you. For example, if you are interested in users reading a file, do not audit Full Control.

  For information about the object access event category, see Audit object access.

  For information about operation-based auditing, see Operation-based auditing on files or folders.

  For information about how to enable auditing in the object access event category, see Define or modify auditing policy settings for an event category.

**Set an appropriate size for the security log.**

- It is important that the size of the security log be configured appropriately, based on the number of events that your auditing policy settings generate.

  For information about managing event logs, see Event Viewer.

  For information about settings for the event logs, see Settings for Event Logs.

# Auditing Security Events How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Define or modify auditing policy settings for an event category
- Apply or modify auditing policy settings for a local file or folder
- Apply or modify auditing policy settings for an object using Group Policy
- View the security log

# Define or modify auditing policy settings for an event category

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define or modify auditing policy settings for an event category**

- For your local computer
- For only domain controllers, when you are on a domain controller or on a workstation that hasWindows Server 2003Administration Tools Pack installed
- For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed
- For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain

**For your local computer**

1. Open Local Security Settings.
2. In the console tree, click **Audit Policy**.

   **Where?**

   - Security Settings/Local Policies/Audit Policy

3. In the details pane, double-click an event category that you want to change the auditing policy settings for.
4. Do one or both of the following, and then click **OK**.

   - To audit successful attempts, select the **Success** check box.
   - To audit unsuccessful attempts, select the **Failure** check box.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**For only domain controllers, when you are on a domain controller or on a workstation that has Windows Server 2003 Administration Tools Pack installed**

1. Open Domain Controller Security Policy.
2. In the console tree, click **Audit Policy**.

   **Where?**

- Computer Configuration/Windows Settings/Security Settings/Local Policies/Audit Policy

3. In the details pane, double-click an event category that you want to change the auditing policy settings for.

4. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.

5. Do one or both of the following, and then click **OK**.

   - To audit successful attempts, select the **Success** check box.

   - To audit unsuccessful attempts, select the **Failure** check box.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Domain Controller Security Policy, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Domain Controller Security Policy**.

**For a domain or organizational unit, when you are on a domain controller or on a workstation that has Administration Tools Pack installed**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

3. Click **Properties**, and then click the **Group Policy** tab.

4. Click **Edit** to open the Group Policy object (GPO) that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.

5. In the console tree, click **Audit Policy**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Local Policies/Audit Policy

6. In the details pane, double-click an event category that you want to change the auditing policy settings for.

7. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.

8. Do one or both of the following, and then click **OK**.

   - To audit successful attempts, select the **Success** check box.

   - To audit unsuccessful attempts, select the **Failure** check box.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**For a domain or organizational unit, when you are on a member server or on a workstation that is joined to a domain**

1. Open Microsoft Management Console (MMC).

2. In the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.

3. Click **Group Policy Object Editor**, and then click **Add**.

4. On the **Select Group Policy Object** page in the Group Policy Wizard, click **Browse**.

5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit--or create a new one, click **OK**, and then click **Finish**.

6. Click **Close**, and then click **OK**.

7. In the console tree, click **Audit Policy**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Local Policies/Audit Policy

8. In the details pane, double-click an event category that you want to change the auditing policy settings for.

9. If you are defining auditing policy settings for this event category for the first time, select the **Define these policy settings** check box.

10. Do one or both of the following, and then click **OK**.

    - To audit successful attempts, select the **Success** check box.

    - To audit unsuccessful attempts, select the **Failure** check box.

    **Notes**

    - To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

    - To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.

**Notes**

- To audit object accesses, enable auditing of the object access event category by following the steps above. Then, enable auditing on the specific object. For information about how to enable auditing on an object, see "Apply or modify auditing

policy settings for a local file or folder" or "Apply or modify auditing policy settings for an object using Group Policy" in Related Topics.

- After your audit policy is configured, events will be recorded in the security log. Open the security log to view these events. For information about the security log, see "Use the security log" in Related Topics.

- The default auditing policy setting for domain controllers is **No Auditing**. This means that even if auditing is enabled in the domain, the domain controllers do not inherit auditing policy locally. If you want domain auditing policy to apply to domain controllers, you must modify this policy setting.

**See Also**
**Concepts**
Auditing Policy
Auditing Security Events
Auditing overview
Apply or modify auditing policy settings for a local file or folder
Apply or modify auditing policy settings for an object using Group Policy
Use the Security Log

# Apply or modify auditing policy settings for a local file or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To apply or modify auditing policy settings for a local file or folder**

1. Open Windows Explorer.
2. Right-click the file or folder that you want to audit, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, and then click the **Auditing** tab.
4. Do one of the following:
   - To set up auditing for a new user or group, click **Add**. In **Enter the object name to select**, type the name of the user or group that you want, and then click **OK**.
   - To remove auditing for an existing group or user, click the group or user name, click **Remove**, click **OK**, and then skip the rest of this procedure.
   - To view or change auditing for an existing group or user, click its name, and then click **Edit**.
5. In the **Apply onto** box, click the location where you want auditing to take place.
6. In the **Access** box, indicate what actions you want to audit by selecting the appropriate check boxes:
   - To audit successful events, select the **Successful** check box.
   - To stop auditing successful events, clear the **Successful** check box.
   - To audit unsuccessful events, select the **Failed** check box.
   - To stop auditing unsuccessful events, clear the **Failed** check box.
   - To stop auditing all events, click **Clear All**.
7. If you want to prevent subsequent files and subfolders of the original object from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.

**Important**

- Before setting up auditing for files and folders, you must enable object access auditing by defining auditing policy settings for the object access event category. If you do not enable object access auditing, you will receive an error message when you set up auditing for files and folders, and no files or folders will be audited. For more information about how to enable object access auditing, see "Define or modify auditing policy settings for an event category" in Related Topics.

**Notes**

- You must be logged on as a member of the Administrators group or you must have been granted the **Manage auditing**

**and security log** right in Group Policy to perform this procedure.

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- For information about how to audit local registry keys, see "Audit activity on a registry key" in Related Topics.
- After object access auditing is enabled, view the security log in Event Viewer to review the results of your changes.
- You can set up file and folder auditing only on NTFS drives.
- If you see the following:
    - In the **Auditing Entry for *File or Folder*** dialog box, in the **Access** box, the check boxes are unavailable ...
    - In the **Advanced Security Settings for *File or Folder*** dialog box, the **Remove** button is unavailable ...

  auditing has been inherited from the parent folder.
- Because the security log is limited in size, select the files and folders to be audited carefully. Also, consider the amount of disk space that you want to devote to the security log. The maximum size for the security log is defined in Event Viewer.

**See Also**
**Concepts**
How inheritance affects file and folder auditing
Selecting where to apply auditing entries
Auditing Security Events
Apply or modify auditing policy settings for an object using Group Policy
Define or modify auditing policy settings for an event category
Audit activity on a registry key

# Apply or modify auditing policy settings for an object using Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To apply or modify auditing policy settings for an object using Group Policy**

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. On the **Select Group Policy Object** page in the Group Policy Wizard, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy object (GPO) in the appropriate domain, site, or organizational unit--or create a new one, click **OK**, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. Do one or more of the following:

| To audit | Do this |
|---|---|
| System services | In the console tree, click **System Services**.<br>Where?<br>• Computer Configuration/Windows Settings/Security Settings/System Services<br>In the details pane, right-click the service that you want to apply or modify auditing policy settings for, and then click **Properties**.<br>If it is not already selected, select the **Define this policy setting** check box, and then select the appropriate setting.<br>Click **Edit security**. |
| Registry keys | In the console tree, click **Registry**.<br>Where?<br>• Computer Configuration/Windows Settings/Security Settings/Registry<br>If you want to add a registry key to this GPO to audit, right-click **Registry**, and then click **Add Key**. Browse to the key that you want, and then click **OK**.<br>If you want to apply or modify auditing settings on a registry key that has already been added to this GPO, in the details pane, right-click the registry key, click **Properties**, and then click **Edit Security**. |

| | |
|---|---|
| Files or folders | In the console tree, click **File System**.<br><br>Where?<br><br>• Computer Configuration/Windows Settings/Security Settings/File System<br><br>If you want to add a file or folder to this GPO to audit, right-click **File System**, and then click **Add File**. Browse to the file that you want, or make a new folder, and then click **OK**.<br><br>If you want to apply or modify auditing settings on a file or folder that has already been added to this GPO, in the details pane, right-click the file or folder, click **Properties**, and then click **Edit Security**. |

8. Click **Advanced**, and then click the **Auditing** tab.

9. Do one of the following:

   - To set up auditing for a new user or group, click **Add**. In **Name**, type the name of the user or group that you want, and then click **OK**.
   - To view or change auditing for an existing group or user, click the name that you want, and then click **Edit**.
   - To remove auditing for an existing group or user, click the name that you want, click **Remove**, click **OK**, and then skip the rest of this procedure.

10. Select the appropriate entry in the **Apply onto** list.

11. In the **Access** box, indicate what actions you want to audit by selecting the appropriate check boxes:

    - To audit successful events, select the **Successful** check box.
    - To stop auditing successful events, clear the **Successful** check box.
    - To audit unsuccessful events, select the **Failed** check box.
    - To stop auditing unsuccessful events, clear the **Failed** check box.
    - To stop auditing all events, click **Clear All**.

12. If you want to prevent files and subfolders in the tree from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.

**Important**

- Before setting up auditing for files and folders, you must enable object access auditing by defining auditing policy settings for the object access event category. If you do not enable object access auditing, you will receive an error message when you set up auditing for files and folders, and no files or folders will be audited. For more information about how to enable object access auditing, see "Define or modify auditing policy settings for an event category" in Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Microsoft Management Console, click **Start**, click **Run**, type **mmc**, and then click **OK**.
- For more information on selecting where to apply auditing entries, see Related Topics.
- You can set up file and folder auditing only on NTFS drives.
- If you see the following:
    - In the **Auditing Entry for *File or Folder*** dialog box, in the **Access** box, the check boxes are unavailable ...
    - In the **Advanced Security Settings for *File or Folder*** dialog box, the **Remove** button is unavailable ...

  auditing is inherited from the parent folder.
- After object access auditing is enabled, view the security log in Event Viewer to review the results of your changes.
- Because the security log is limited in size, select the files and folders to be audited carefully. Also, consider the amount of disk space that you want to devote to the security log. The maximum size for the security log is defined in Event Viewer.

**See Also**
**Concepts**
Auditing Security Events
View the security log
Define or modify auditing policy settings for an event category
Apply or modify auditing policy settings for a local file or folder
Apply or modify auditing policy settings for an object using Group Policy
Selecting where to apply auditing entries

# View the security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the security log**

1. Open Event Viewer.
2. In the console tree, click **Security**. The details pane lists individual security events.
3. If you want to see more details about a specific event, in the details pane, double-click the event.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- For more information about security events, see "Security Events" at the Microsoft Windows Resource Kits Web site

**See Also**
**Concepts**
Auditing Security Events
Auditing Policy
Settings for Event Logs
Understanding event logging options

# Auditing Security Events Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about auditing:

- Auditing overview
- Viewing security logs
- Operation-based auditing on files or folders
- Auditing settings on objects
- How inheritance affects file and folder auditing
- Selecting where to apply auditing entries
- Per-user selective auditing

# Auditing overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Auditing overview**
Establishing audit policy is an important facet of security. Monitoring the creation or modification of objects gives you a way to track potential security problems, helps to ensure user accountability, and provides evidence in the event of a security breach.

The most common types of events to be audited are:

- Access to objects, such as files and folders.
- Management of user accounts and group accounts.
- Users logging on to and logging off from the system.

When you implement audit policy:

- Specify the categories of events that you want to audit. Examples of event categories are user logon, user logoff, and account management. The event categories that you select constitute your audit policy. For more information about each event category, see Auditing Policy.
- Set the size and behavior of the security log. You can view the security log with Event Viewer. For more information about the security log, see Viewing security logs.
- If you want to audit directory service access or object access, determine which objects you want to monitor access of and what type of access you want to monitor. For example, if you want to audit any attempts by users to open a particular file, you can configure auditing policy settings in the object access event category so that both successful and failed attempts to read a file are recorded.

  For more information about how to set up object access auditing, see:

  - Apply or modify auditing policy settings for a local file or folder
  - Apply or modify auditing policy settings for an object using Group Policy

For more information about auditing, see Auditing Security Events.

# Viewing security logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing security logs**

Auditing a local object creates an entry in the security log. The security entries that appear in the security log depend on the auditing categories that you select for your audit policy. For events that are related to object access, the entries in the security log also depend on the auditing policy settings that are defined for each object. For more information about auditing, see Auditing Security Events.

For example, if your audit policy specifies the auditing of files and folders, and a file's properties specify that failed deletions of that file are to be audited, each failed attempt by a user to delete the file appears in the security log.

You can view the security log by using Event Viewer.

**Size of the security log**

It is important that you set the size of the security log appropriately. Because the security log is limited in size, choose which events you audit carefully. Also, consider the amount of disk space that you are willing to devote to the security log. The maximum amount is defined in Event Viewer.

For information on what to do when the security log is full, see Audit: Shut down system immediately if unable to log security audits.

For more information about security events, see "Security Events" at the Microsoft Windows Resource Kits Web site.

# Operation-based auditing on files or folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Operation-based auditing on files or folders**

Operation-based auditing is a new feature in the Windows Server 2003 family. In earlier versions of Windows, information that was gained from object access auditing was not as detailed as it is with operation-based auditing. While you could determine that a user attempted to access an object, there was no way to be sure that the object was accessed in every way that was documented in the audit event. With operation-based auditing, you can audit operations on files and folders. This means that you can audit certain operations, such as Write, as well as the accessing of objects. Operation-based auditing is enabled when object access auditing is enabled on a file or folder. Object access events are recorded, along with operations such as Write, in the security log.

To enable operation-based auditing, you must:

- Enable the **Audit object access** policy setting. For more information, see Define or modify auditing policy settings for an event category.
- Apply audit policy to a specific folder. For more information, see:
  - Apply or modify auditing policy settings for a local file or folder
  - Apply or modify auditing policy settings for an object using Group Policy

Operation-based audits are categorized as object audits, and they are logged as event number 567 in the security log. They are generated the first time an operation is performed. Only files and folders can be set up to generate operation-based audits.

# Auditing settings on objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Auditing settings on objects**

Each object has a set of security information, or security descriptor, attached to it. Part of the security descriptor specifies the groups or users that can access an object and the types of access (permissions) that are granted to those groups or users. This part of the security descriptor is known as a discretionary access control list (DACL).

A security descriptor for an object also contains auditing information. This auditing information is known as a system access control list (SACL). More specifically, a SACL specifies the following:

- The group or user accounts to audit when they access the object.
- The operations to be audited for each group or user, for example, modifying a file.
- A Success or Failure attribute for each access event, based on the permissions that are granted to each group and user in the object's DACL.

You can apply auditing to an object, and, through inheritance, the auditing can apply to any child objects. For example, if you want to audit failed access to a folder, this auditing event can be inherited by all files within the folder. For more information, see How inheritance affects file and folder auditing.

To audit files and folders, you must be logged on as a member of the Administrators group.

For more information about auditing, see Auditing Security Events.

# How inheritance affects file and folder auditing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How inheritance affects file and folder auditing**

After you set up auditing on a parent folder, new files and subfolders that are created in that folder inherit auditing. If you do not want them to inherit auditing, in the **Apply onto** box in the **Auditing Entry for *File or Folder*** dialog box, click **This folder only** when you set up auditing for the parent folder. In cases where you want to prevent only certain files or subfolders from inheriting auditing, right-click the file or subfolder that you want to prevent from inheriting auditing, click **Properties**, click the **Security** tab, click **Advanced**, click the **Auditing** tab, and then clear the check box that is labeled **Inherit from parent the auditing entries that apply to child objects. Include these with entries explicitly defined here**.

If the check boxes in **Auditing Entry for *File or Folder*** are unavailable or if, in **Access Control Settings**, the **Remove** button is unavailable, auditing has been inherited from the parent folder.

To make changes to inherited auditing:

- Make the changes to the parent folder, and then the file or folder will inherit auditing.
- On the **Auditing** tab in the **Advanced Security Settings for *File or Folder*** dialog box, clear the check box that is labeled **Inherit from parent the auditing entries that apply to child objects. Include these with entries explicitly defined here**. You can then click **Edit** or **Remove** to make changes or remove auditing. However, the file or folder will no longer inherit auditing changes from the parent folder.

For more information, see Auditing Security Events.

# Selecting where to apply auditing entries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Selecting where to apply auditing entries**

The **Auditing Entry** dialog box appears when you audit files or folders. In this dialog box, the **Apply onto** list displays the locations where you can apply auditing entries. How these auditing entries are applied depends on whether the **Apply these auditing entries to objects and/or containers within this container only** check box is selected. By default, this check box is cleared.

Click either **Cleared** or **Selected** below to see how auditing entries are applied, depending on whether this check box is cleared or selected. When this check box is:

- Cleared ...
- Selected ...

**Cleared ...**

| Apply onto | Audits the current folder | Audits subfolders in the current folder | Audits files in the current folder | Audits all subsequent subfolders | Audits files in all subsequent subfolders |
|---|---|---|---|---|---|
| This folder only | x | | | | |
| This folder, subfolders, and files | x | x | x | x | x |
| This folder and subfolders | x | x | | x | |
| This folder and files | x | | x | | x |
| Subfolders and files only | | x | x | x | x |
| Subfolders only | | x | | x | |
| Files only | | | x | | x |

**Selected ...**

| Apply onto | Audits the current folder | Audits subfolders in the current folder | Audits files in the current folder | Audits all subsequent subfolders | Audits files in all subsequent subfolders |
|---|---|---|---|---|---|
| This folder only | x | | | | |
| This folder, subfolders, and files | x | x | x | | |
| This folder and subfolders | x | x | | | |
| This folder and files | x | | x | | |

| | | | | | |
|---|---|---|---|---|---|
| Subfolders and files only | | **x** | **x** | | |
| Subfolders only | | **x** | | | |
| Files only | | | **x** | | |

For information about how to set up auditing on objects, see Checklist: Setting up object access auditing.

For more information about auditing, see Auditing Security Events.

# Per-user selective auditing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Per-user selective auditing**

For information about per-user selective auditing, see "Security Auditing" at the Microsoft Windows Resource Kits Web site.

For information about auditing, see Auditing Security Events.

# Encrypting File System

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encrypting File System**

With Encrypting File System (EFS), you can encrypt files and directories that are stored on a disk.

- Before using EFS with a server cluster, see Checklist: Configuring Encrypting File System in a server cluster.
- For tips about using EFS, see Encrypting File System best practices.
- For help with specific tasks, see Encrypting File System How To ....
- For general background information, see Encrypting File System Concepts.
- For problem-solving instructions, see Encrypting File System Troubleshooting.

# Checklist: Configuring Encrypting File System in a server cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring Encrypting File System in a server cluster**
Using Encrypting File System (EFS) on a server cluster requires several tasks to be performed before EFS files can be safely encrypted and decrypted on the server cluster.

| | Step | Reference |
|---|---|---|
| | Create a server cluster | Checklist: Planning and creating a server cluster |
| | Configure the computer accounts for the server cluster to be trusted for delegation<br><br>**Important**<br><br>• All cluster node computer accounts, as well as the virtual server computer account, must be trusted for delegation. | Allow a computer to be trusted for delegation |
| | Configure users of this cluster for roaming user profiles to ensure the private keys of the users are available to each node of the server cluster | Create a roaming user profile |
| | Create a file share for storage of encrypted files | Create a cluster-managed file share |
| | Configure the cluster to support Kerberos authentication | Enable Kerberos authentication for virtual servers |
| | Create or copy files on the cluster server file share | |
| | Encrypt files on the cluster server file share | Encrypt a file or folder |

# Encrypting File System best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### Encrypting File System best practices

#### Ensure files intended for encryption are created and remain encrypted

- Encrypt folders before creating sensitive files in them for maximum security. Doing this causes the files to be created as encrypted and their data is never written to the disk as plaintext.
- Encrypt the My Documents folder if you save most of your documents to the My Documents folder. This ensures that your personal documents are encrypted by default. For Roaming User Profiles, this should only be done when the My Documents folder is redirected to a network location.
- Encrypt folders instead of individual files so that, if a program creates temporary files during editing, these are encrypted as well.

#### Manage private keys to ensure file security

- The designated recovery agent should export the data recovery certificate and private key to disk, secure them in a safe place, and delete the data recovery private key from the system. In this way, the only person who can recover data for the system is the person who has physical access to the data recovery private key.
- The number of designated recovery agents should be kept to a minimum. This exposes fewer keys to cryptographic attack and provides a higher level of assurance that encrypted data will not be decrypted inappropriately.
- Use Microsoft Certificate Services to manage Encrypting File System (EFS) and Data Recovery Agent (DRA) certificates and private keys.

### Caution

- When configuring Certificate Services and using a custom certificate template to issue EFS certificates, do not select the **Prompt the user during enrollment and require user input when the private key is used** option. This option prevents EFS from using the private key for encryption or decryption.

#### Provide security and reliability of data at all times

- Encrypt sensitive data on computers that are members of a domain. This protects against compromise of data though offline cryptographic attacks.
- Use Internet Protocol security (IPSec) to ensure that data remains encrypted as it is transmitted over the network. EFS can be used in conjunction with Web Distributed Authoring and Versioning (WebDAV) to store encrypted data on the Internet. In addition, EFS can be used with Server Message Block (SMB) signing to ensure that the transmission and reception of EFS files across a network is not altered in any way.
- Back up the entire server that stores server-based encrypted data regularly. This ensures that, in case of data recovery, the profiles that include decryption keys can be restored.

# Encrypting File System How To ...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Encrypt or Decrypt Files
- Copy or Move Encrypted Files or Folders
- Recover Files or Folders
- Add users to or remove users from a file
- Back up an EFS certificate with the private key

# Encrypt or Decrypt Files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encrypt or decrypt files**

- Encrypt a file or folder
- Decrypt a file or folder
- Enable a remote server for file encryption
- Encrypt a file or folder on a remote computer

# Encrypt a file or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To encrypt a file or folder**

1. Open Windows Explorer.
2. Right-click the file or folder that you want to encrypt, and then click **Properties**.
3. On the **General** tab, click **Advanced**.
4. Select the **Encrypt contents to secure data** check box, and then click **OK**.
5. In the **Properties** dialog box, click **OK**, and then do one of the following:
   - To encrypt a file and the parent folder, in the **Encryption Warning** dialog box, click **Encrypt the file and the parent folder**.
   - To encrypt a file only, in the **Encryption Warning** dialog box, click **Encrypt the file only**.
   - To encrypt a folder only, in the **Confirm Attribute Changes** dialog box, click **Apply changes to this folder only**.
   - To encrypt a folder and its subfolders and files, in the **Confirm Attribute Changes** dialog box, click **Apply changes to this folder, subfolders and files**.

**Notes**

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- You can only encrypt files and folders on NTFS file system volumes.
- Files or folders that are compressed cannot also be encrypted. If you encrypt a compressed file or folder, that file or folder will be uncompressed.
- Files marked with the System attribute cannot be encrypted, nor can files in the systemroot folder.
- If you choose to encrypt the parent folder when you encrypt a single file, all files and subfolders that are added to the folder in the future will be encrypted when they are added.
- If you choose to encrypt all files and subfolders when you encrypt a folder, all files and subfolders currently in the folder are encrypted, as well as any files and subfolders that are added to the folder in the future. If you choose to encrypt the folder only, all files and subfolders currently in the folder are not encrypted. However, any files and subfolders that are added to the folder in the future are encrypted when they are added.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Decrypt a file or folder

# Decrypt a file or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To decrypt a file or folder**

1. Open Windows Explorer.
2. Right-click the encrypted file or folder, and then click **Properties**.
3. On the **General** tab, click **Advanced**.
4. Clear the **Encrypt contents to secure data** check box.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- When you decrypt a folder, you are asked if you want all files and subfolders within the folder to be decrypted as well. If you choose to decrypt the folder only, the encrypted files and folders within the decrypted folder remain encrypted. However, new files and folder that you create within the decrypted folder will not be automatically encrypted.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Encrypt a file or folder
Encrypting and decrypting data

# Enable a remote server for file encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable a remote server for file encryption**

1. Open Active Directory Users and Computers, locate the remote server computer account.
2. Right-click the remote server computer account, and then click **Properties**.
3. Click the **Delegation** tab, select the **Trust this computer for delegation to specified services only** option.
4. Ensure that Use **Kerberos** only is selected. Click **Add**.
5. Click **Users or Computers**. Use **Select Users or Computers** to locate and add the account for a local domain controller.
6. From the list of available services, select the **protectedstorage** and the **cifs** services.
7. Repeat steps 5, and 6 for each domain controller computer account in the domain.
8. Once all domain controller computer objects have been added, click OK on the remote computer's **Properties** dialog box.

| **Important** |
| --- |
| The list must be updated when domain controllers are added or removed from the domain. |

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- When encrypting files on a WebDAV server, the computer does not need to be trusted for delegation.
- You cannot configure a computer in another forest for encryption, even if there is a trust relationship established.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Encrypt a file or folder on a remote computer
Encrypting and decrypting data
User and computer accounts
Delegating authentication

# Encrypt a file or folder on a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To encrypt a file or folder on a remote computer**

1. Open Windows Explorer.
2. On the **Tools** menu, click **Map Network Drive**, and then follow the instructions in the **Map Network Drive** dialog box.
3. Right-click the file or folder that you want to encrypt, and then click **Properties**.
4. On the **General** tab, click **Advanced**.
5. Select the **Encrypt contents to secure data** check box.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- You can only encrypt files and folders on NTFS file system volumes.
- Files or folders that are compressed cannot also be encrypted. If you encrypt a compressed file or folder, that file or folder will be uncompressed.
- Files that are marked with the System attribute or are in the systemroot directory cannot be encrypted.
- In a domain environment, remote encryption is not enabled by default. To enable encryption for a specific computer, your network administrator can make that computer trusted for delegation. For more information, consult your network administrator.
- When you encrypt a folder, you are asked if you want all files and subfolders within the folder to be encrypted as well. If you choose to do so, all future files and subfolders that are added to the folder are automatically encrypted.
- Programs that create temporary work files can compromise file encryption security. When working with such programs, encrypt the folder rather than the individual files.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable a remote server for file encryption
Encrypting and decrypting data

# Copy or Move Encrypted Files or Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Copy or move encrypted files or folders**

- Copy an encrypted file or folder
- Move an encrypted file or folder

# Copy an encrypted file or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy an encrypted file or folder**

1. Open Windows Explorer.
2. Click the encrypted file or folder that you want to copy.
3. On the **Edit** menu, click **Copy**.
4. Open the folder or disk where you want to store the copy.
5. On the **Edit** menu, click **Paste**.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- An encrypted file or folder is decrypted if you copy it to a volume that is not an NTFS file system volume. The exception to this is when copying to a WebDAV folder on a server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Move an encrypted file or folder
Encrypting File System overview
Encrypting and decrypting data

# Move an encrypted file or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move an encrypted file or folder**

1. Open Windows Explorer.
2. Click the encrypted file or folder that you want to move.
3. On the **Edit** menu, click **Cut**.
4. Open the folder where you want to move the file or folder.
5. On the **Edit** menu, click **Paste**.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- An encrypted file or folder is decrypted if you move it to a volume that is not an NTFS file system volume.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Copy an encrypted file or folder
Encrypting File System overview

# Recover Files or Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recover files or folders**

- Recover an encrypted file or folder if you are a designated recovery agent
- Recover an encrypted file or folder without your file encryption certificate
- Back up default recovery keys to a floppy disk
- Add a recovery agent for the local computer
- Change the recovery policy for the local computer
- Add a recovery agent for a domain
- Create a recovery policy for a domain
- Not require a recovery agent for a domain
- Disable Encrypting File System for the local computer
- Disable Encrypting File System for a domain

# Recover an encrypted file or folder if you are a designated recovery agent

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To recover an encrypted file or folder if you are a designated recovery agent**

1. Use Backup or another backup tool to restore a user's backup version of the encrypted file or folder to the computer where your file recovery certificate and recovery key are located.
2. Open Windows Explorer.
3. Right-click the file or folder and then click **Properties**.
4. On the **General** tab, click **Advanced**.
5. Clear the **Encrypt contents to secure data** check box.
6. Make a backup version of the decrypted file or folder and return the backup version to the user.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- You can return the backup version of the decrypted file or folder to the user as an e-mail attachment, on a floppy disk, or on a network share.
- An alternate procedure would involve physically transporting the recovery agent's private key and certificate, importing the private key and certificate, decrypting the file or folder, and then deleting the imported private key and certificate. This procedure exposes the private key more than the procedure above but does not require any backup or restore operations or file transportation.
- If you are the recovery agent, use the **Export** command from Certificates in Microsoft Management Console (MMC) to export the file recovery certificate and private key to a floppy disk. Keep the floppy disk in a secure location. Then, if the file recovery certificate or private key on your computer is ever damaged or deleted, you can use the **Import** command from Certificates in MMC to replace the damaged or deleted certificate and private key with the ones you have backed up on the floppy disk.
- For more information about using Certificates in MMC, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovering data
Recover an encrypted file or folder without your file encryption certificate
Add a recovery agent for the local computer

# Recover an encrypted file or folder without your file encryption certificate

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To recover an encrypted file or folder without the file encryption certificate**

1. Open Backup.
2. Use Backup to make a copy of the file in case of loss or damage.
3. Send the original encrypted file to the designated recovery agent.
4. Have the recovery agent use their recovery certificate and private key to decrypt the file.
5. Have the recovery agent send the decrypted file back to you, using any file transfer method that you want.

## Notes

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
- There is no default recovery agent on a local computer unless the computer is in an Active Directory domain environment. In an Active Directory domain environment, the administrator that initially logged on to the first domain controller is the default recovery agent.
- Sending the file to the designated recovery agent can be done in a number of ways, including backing up the file to tape or floppy disk.
- Files backed up using Backup or any other backup tool retain their encryption while in their backup storage location. The original files can be decrypted or modified without affecting the encrypted state of the backup copies.
- You can recover an encrypted file or folder yourself if you have kept a backup copy of your file encryption certificate and private key in a .pfx file format on a floppy disk. Use the **import** command from Certificates in Microsoft Management Console (MMC) to import the .pfx file from the floppy disk into the **Personal** store.
- For more information about using Certificates in MMC, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovering data

# Back up default recovery keys to a floppy disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up default recovery keys to a floppy disk**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Add Standalone Snap-in**, click **Certificates**, and then click **Add**.
4. Click **My user account**, and then click **Finish**.
5. Click **Close**, and then click **OK**.
6. Double-click **Certificates - Current User**, double-click **Personal**, and then double-click **Certificates**.
7. Click the certificate that displays the words **File Recovery** in the **Intended Purposes** column.
8. Right-click the certificate, point to **All Tasks**, and then click **Export**.
9. Follow the instructions in the Certificate Export Wizard to export the certificate and associated private key to a .pfx file format.

## Notes

- This operation must be performed by the recovery agent account that has the recovery certificate and private key in their private store.
- In a home or non-domain environment, there is no default recovery agent. To configure a recovery agent, see Related Topics.
- Before making any changes to the default recovery policy, be sure to secure the default recovery private key. The default recovery keys in a domain are stored on the first domain controller for the domain. The domain administrator is the default recovery agent.
- For more information about using Certificates in MMC, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Certificates
Recovering data
Add a recovery agent for the local computer

# Disable Encrypting File System for the local computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable Encrypting File System for the local computer**

1. Open Registry Editor.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the **Last Known Good Configuration** startup option if you encounter problems after manual changes have been applied.

2. In Registry Editor, navigate to the following registry key:

   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\EFS**

3. Click **Edit**, point to **New**, and then click **DWORD Value**.

4. In the text box that appears, type **EfsConfiguration**, and then press **Enter**.

5. Double-click the **EfsConfiguration** registry key you just created, and then change **Value Data** to 1.

6. Click **OK**, and then close Registry Editor.

7. Restart the computer.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable Encrypting File System for a domain
Recovering data

# Add users to or remove users from a file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add users to or remove users from a file**

1. Open Windows Explorer.
2. Right-click the encrypted file that you want to change, and then click **Properties**.
3. On the **General** tab, click **Advanced**.
4. In **Advanced Attributes**, click **Details**.
5. To add a user to this file, click **Add**, and then do one of the following:
   - To add a user whose certificate is on this computer, click the certificate and then click **OK**.
   - To view a certificate on this computer before adding it to the file, click the certificate and then click **View Certificate**.
   - To add a user from the Active Directory, click **Find User**, then locate the user in the list and click **OK**.
6. To remove a user from this file, click the user name and then click **Remove**.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- You cannot add groups to file encryption access.
- All users that are added to the file must have a certificate located on the computer. For details on how to copy the certificate to the computer, see Related Topics.
- All users that can decrypt the file must also have access to read the file. NTFS permissions must be set properly to allow this access. If a user is denied access through NTFS permissions, they cannot read the encrypted file and cannot decrypt the data. For more information on setting permissions on files, see Related Topics.
- When a user is added to a file and their certificate is imported, the certificate's chain is validated to a trusted root CA. The certificate is then stored in the user's *Other People* certificate store.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Cipher
Recover an encrypted file or folder without your file encryption certificate
Decrypt a file or folder
Encrypting File System overview
Import and Export Certificates

# Back up an EFS certificate with the private key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up an EFS certificate with the private key**

- To back up the EFS certificate currently in use
- To back up any EFS certificate

**To back up the EFS certificate currently in use**

1. Open Certificate Export Wizard.
2. In the Certificate Export Wizard, click **Next**, and then click **Yes, export the private key**. (This option will appear only if the private key is marked as exportable and you have access to the private key.)
3. Under **Export File Format**, ensure that the **Enable strong protection (requires IE 5.0, NT 4.0 SP4 or above)** check box is selected, and then click **Next**.
4. In **Password**, type a password to encrypt the private key you are exporting. In **Confirm password**, type the same password again, and then click **Next**.
5. In **File name**, type a file name and path for the PKCS #12 file that will store the exported certificate and private key, click **Next**, and then click **Finish**.

**To back up any EFS certificate**

1. Open Certificates - Current User.
2. In the console tree, click **Certificates**.

   **Where?**

   - Certificates - Current User/Personal/Certificates

3. In the details pane, click the certificate that has **Encrypting File System** listed in the **Intended Purposes** column.
4. On the **Action** menu, point to **All Tasks**, and then click **Export**.
5. In the Certificate Export Wizard, click **Next**, and then click **Yes, export the private key**. (This option will appear only if the private key is marked as exportable and you have access to the private key.)
6. Under **Export File Format**, ensure that the **Enable strong protection (requires IE 5.0, NT 4.0 SP4 or above)** check box is selected, and then click **Next**.
7. In **Password**, type a password to encrypt the private key you are exporting. In **Confirm password**, type the same password again, and then click **Next**.
8. In **File name**, type a file name and path for the PKCS #12 file that will store the exported certificate and private key, click **Next**, and then click **Finish**.

**Notes**

- To open Certificates, click **Start**, click **Run**, type **certmgr.msc**, and then click **OK**.

- If more than one certificate has an intended purpose of Encrypting File System, each certificate should be backed up individually. This will help to ensure that all encrypted files can be accessed when necessary.

- Strong protection (also known as *iteration count*) is enabled by default in the Certificate Export Wizard when you export a certificate with its associated private key.

  Strong protection is not compatible with older programs, so you must clear the **Enable strong protection** check box if you are going to use the private key with any browser earlier than Microsoft Internet Explorer 5.0.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Display certificate stores in Logical Store mode
Import a certificate

# Encrypting File System Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Encrypting File System overview
- Using Encrypting File System

# Encrypting File System overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encrypting File System overview**

Encrypting File System (EFS) provides the core file encryption technology used to store encrypted files on NTFS file system volumes. Once you encrypt a file or folder, you work with the encrypted file or folder just as you do with any other files and folders.

Encryption is transparent to the user that encrypted the file. This means that you do not have to manually decrypt the encrypted file before you can use it. You can open and change the file as you normally do.

Using EFS is similar to using permissions on files and folders. Both methods can be used to restrict access to data. However, an intruder who gains unauthorized physical access to your encrypted files or folders will be prevented from reading them. If the intruder tries to open or copy your encrypted file or folder he receives an access denied message. Permissions on files and folders does not protect against unauthorized physical attacks.

You encrypt or decrypt a folder or file by setting the encryption property for folders and files just as you set any other attribute such as read-only, compressed, or hidden. If you encrypt a folder, all files and subfolders created in the encrypted folder are automatically encrypted. It is recommended that you encrypt at the folder level.

You can also encrypt or decrypt a file or folder using the **cipher** command. For more information, see Cipher.

When you work with encrypted files and folders, keep in mind the following information:

- Only files and folders on NTFS volumes can be encrypted. Because WebDAV works with NTFS, NTFS is required when encrypting files over WebDAV (Web distributed authoring and versioning).
- Files or folders that are compressed cannot also be encrypted. If the user marks a file or folder for encryption, that file or folder will be uncompressed.
- Encrypted files can become decrypted if you copy or move the file to a volume that is not an NTFS volume.
- Moving unencrypted files into an encrypted folder will automatically encrypt those files in the new folder. However, the reverse operation will not automatically decrypt files. Files must be explicitly decrypted.
- Files marked with the System attribute cannot be encrypted, nor can files in the systemroot directory structure.
- Encrypting a folder or file does not protect against deletion or listing files or directories. Anyone with the appropriate permissions can delete or list encrypted folders or files. For this reason, using EFS in combination with NTFS permissions is recommended.
- You can encrypt or decrypt files and folders located on a remote computer that has been enabled for remote encryption. However, if you open the encrypted file over the network, the data that is transmitted over the network by this process is not encrypted. Other protocols, such as SSL/TLS (Secure Socket Layer/Transport Layer Security) or Internet Protocol security (IPSec) must be used to encrypt data over the wire. WebDAV, however, is able to encrypt the file locally and transmit it in encrypted form.

# Using Encrypting File System

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Encrypting File System**

This section covers:

- Encrypting and decrypting data
- Recovering data
- Managing certificates

# Encrypting and decrypting data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encrypting and decrypting data**
With Encrypting File System (EFS) you can store data securely. EFS does this by encrypting data in selected NTFS file system files and folders.

Because EFS is integrated with the file system, it is easy to manage, difficult to attack, and transparent to the user. This is particularly useful for securing data on computers that may be vulnerable to theft, such as mobile computers.

Files and folders cannot be encrypted or decrypted on FAT volumes. Also, EFS is designed to store data securely on local computers. As such, it does not support the secure transmission of files over a network. Other technologies, such as Internet Protocol Security (IPSec), can be used in conjunction with EFS to provide a larger solution. For more information, see Internet Protocol Security (IPSec).

**Using encryption keys**
Once a user has specified that a file be encrypted, the actual process of data encryption and decryption is completely transparent to the user. The user does not need to understand this process. However, the following explanation of how data encryption and decryption works might be useful for administrators.

*Encryption* of files works as follows:

- Each file has a unique *file encryption key*, which is later used to decrypt the file's data.
- The file encryption key is itself encrypted--it is protected by the user's public key corresponding to the user's EFS certificate.
- The file encryption key is also protected by the public key of each additional EFS user that has been authorized to decrypt the file and each recovery agent.

For steps on how to encrypt a file or folder, see Encrypt a file or folder

The EFS certificate and private key used can be issued by a number of sources, including automatically-generated certificates, certificates created by Microsoft certification authorities (CAs), or third-party CAs. For more information about third-party certificates and EFS, see article Q273856, "Third-Party Certificate Authority Support for Encrypting File System," in the Microsoft Knowledge Base.

*Decryption* of files works as follows:

- To decrypt a file, the file encryption key must first be decrypted. The file encryption key is decrypted when the user has a private key that matches the public key.
- The original user may not be the only person that can decrypt the file encryption key. Other designated users or recovery agents can also decrypt the file encryption key, by using their own private key.

For steps on how to decrypt a file or folder, see Decrypt a file or folder

Private keys are securely held in a protected key store, and not in the Security Account Manager (SAM) or in a separate directory. For more information, see the Microsoft TechNet on the Microsoft Web site.

**Storing encrypted files on remote servers**
If users in your Windows XP or Windows Server 2003 family computing environment want to store encrypted files on remote servers, it is useful to know the following:

- Windows XP and the Windows Server 2003 family support the storage of encrypted files on remote servers.
- Users can use EFS remotely only when both computers are members of the same Windows Server 2003 family forest.

- Encrypted data is not encrypted when in transit over the network, but only when stored on disk. The exceptions to this are when your system includes Internet Protocol security (IPSec) or Web Distributed Authoring and Versioning (WebDAV). IPSec encrypts data while it is transported over a Transmission Control Protocol/Internet Protocol (TCP/IP) network. If the file is encrypted before being copied or moved to a WebDAV folder on a server, it will remain encrypted during the transmission and while it is stored on the server.

- Encrypted files are not accessible from Macintosh clients.

- Storing EFS certificates and private keys on smartcards is not currently supported.

- Strong private key protection for EFS private keys is not currently supported.

Before users can encrypt files that reside on a remote server, an administrator must designate the remote server as *trusted for delegation*. This allows all users with files on that server to encrypt those files. For more information, see Enable a remote server for file encryption and Encrypt a file or folder on a remote computer.

# Recovering data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovering data**

Data recovery is important when you need to be able to recover data encrypted by an employee after the employee leaves, or when the user loses the private key. Data recovery is available through the Encrypting File System (EFS) as a part of the overall security policy for the system. For example, if you should ever lose your file encryption certificate and associated private key through disk failure, arson, or any other reason, the person who is the designated recovery agent can recover the data. In a business environment, an organization can recover data encrypted by an employee after the employee leaves.

**Recovery policy**

EFS uses recovery policies to provide built-in data recovery. A *recovery policy* is a type of public key policy that provides for one or more user accounts to be designated as recovery agents.

A recovery policy is configured locally for stand-alone computers. For computers that are part of a network, a recovery policy is configured at the domain, organizational unit, or individual computer level, and applies to all Windows XP and Windows Server 2003 family-based computers that the policy applies to. A certification authority (CA) issues recovery certificates, and you use Certificates in Microsoft Management Console (MMC) to manage them.

In a domain, the Windows Server 2003 family implements a default recovery policy for the domain when the first domain controller is set up. The self-signed certificate is issued to the domain administrator. That certificate designates the domain administrator as the recovery agent. To change the default recovery policy for a domain, log on to the first domain controller as an administrator. Additional recovery agents can be added to this policy and the original recovery agent can be removed at any time.

Because the Windows XP and Windows Server 2003 family security subsystems handle enforcing, replicating, and caching of the recovery policy, users can implement file encryption on a system that is temporarily offline, such as a portable computer. This process is similar to logging on to their domain account using cached credentials. For more information, see Change the recovery policy for the local computer and Create a recovery policy for a domain.

**Recovery agents**

A recovery agent is an individual authorized to decrypt data that was encrypted by another user. Recovery agents do not need any other permissions to function in this role. Recovery agents are useful, for example, when employees leave the company and their remaining data needs to be decrypted. Before you can add a recovery agent for a domain, you must ensure that each recovery agent has been issued an X.509v3 certificate.

Each recovery agent has a special certificate and associated private key that allows data recovery wherever the recovery policy applies. If you are the recovery agent, you should be sure to use the **Export** command in Certificates in MMC to back up the recovery certificate and the associated private key to a secure location. After backing them up, you should use Certificates in MMC to delete the recovery certificate. Then, when you need to perform a recovery operation for a user, you should first restore the recovery certificate and associated private key using the **Import** command from Certificates in MMC. After recovering the data, you should again delete the recovery certificate. You do not have to repeat the export process.

To add recovery agents for a domain, you add their certificates to the existing recovery policy. For steps on how to add recovery agents to a domain, see Add a recovery agent for a domain.

### Note

- Recovery agent information that has been added and removed is not automatically updated on existing EFS files. The information in these files is updated the next time the file is accessed. New files always use the current recovery agent information.

# Managing certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing certificates**

Encrypting File System (EFS) uses public key cryptography to encrypt the contents of files. The keys that are used are obtained from the certificate of the user and any additional users and designated recovery agents configured. Because the certificates may also contain private key information, they must be managed correctly.

Certificates that are used by EFS can be obtained from a certification authority (CA) or created automatically by the computer. When obtaining an EFS certificate from a CA, the cryptographic service provider (CSP) and the appropriate object identifier (also known as an OID) must be referenced by the certificate. EFS can use either a base or enhanced CSP. If these two attributes are not set correctly in the certificate, EFS is unable to use it.

The certificate and private key of all designated recovery agents should be exported to removable disk and stored securely until needed. For more information, see Export a certificate with the private key. When exporting the certificate and private key, ensure that, in the **Intended Purposes** column, the selected certificate includes **Encrypting File System** and that you have the associated private key.

Certificates and private keys can be used on multiple computers. If your Windows XP network is configured to use roaming profiles, the certificates will be available on any computer you logon to. Otherwise, the certificates and private keys must be exported and imported manually. To manually export a certificate and private key, use the above procedure. To import this information, see Import a certificate.

**Verification of certificate validity**

Certificates are not expected to be valid indefinitely. Over time, an attacker can determine the corresponding private key and render the data encrypted with that key vulnerable. Because of this, certificates have a validity period that defines the length of time during which they can be considered valid. Once the validity period expires, a new certificate must be obtained to encrypt new data. However, the existing certificate and private key are normally retained so older data can still be decrypted.

Certificates can also be revoked by the issuing CA, even when they are still within their validity period. There are a number of reasons why a certificate could become untrustworthy, including compromise of the certificate subject's private key or discovery that a certificate was obtained fraudulently. When a certificate is revoked, it is placed on a certificate revocation list (CRL) maintained by the CA.

When a file is encrypted, EFS checks the validity period on the certificate of the user as well as the recovery agent. When a new user is added to an existing file, EFS checks for both the revocation of the certificate being added and the chaining of the certificate to a trusted root CA. If the certificate is found to be invalid (either because of expiration, revocation or inability to chain) the certificate is not used and the user is typically notified.

# Encrypting File System Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The Advanced button is unavailable.
- When encrypting a file, a message appears: "Recovery policy configured for this system contains invalid recovery certificate" or "ERROR_BAD_RECOVERY_POLICY."
- The "Access denied" message appears when opening an encrypted file.
- The message "Key not valid for use in specified state" appears when encrypting a file on a remote computer.
- The message "You do not have permission to request a certificate based on the selected certificate template" appears when using theCreate Data Recovery Agentoption to create an EFS data recovery agent.

**The Advanced button is unavailable.**
**Cause:** Encrypting File System (EFS) only works on files and folders on NTFS file system volumes. If the folder or file you are trying to encrypt is on a FAT or FAT32 volume, the **Advanced** button does not appear in the properties of that folder or file.

**Solution:** Convert the volume to NTFS with the **convert** utility.

1. Open Command Prompt.
2. Type:

   **convert** *drive* **/fs:ntfs**

   where *drive* is the drive letter of the intended drive.

**See also:** Convert.

**When encrypting a file, a message appears: "Recovery policy configured for this system contains invalid recovery certificate" or "ERROR_BAD_RECOVERY_POLICY."**
**Cause:** The Encrypting File System (EFS) recovery policy that is implemented on this computer contains one or more EFS recovery agent certificates that have expired. These certificates cannot be used.

**Solution:** Either renew the existing certificates or generate new certificates for the EFS recovery agents and reapply the recovery agent policy with those certificates.

**See also:** Requesting certificates or Renewing certificates.

**The "Access denied" message appears when opening an encrypted file.**
**Cause:** The file was encrypted by Encrypting File System (EFS) using a public key certificate, and the associated private key for this certificate is not available on this computer.

**Solution:** Locate the private key for the appropriate certificate and import it onto this computer using the Certificates snap-in.

**See also:** Import a certificate and **Encrypting File System** at the Microsoft Windows XP Resource Kits Web site.

**The message "Key not valid for use in specified state" appears when encrypting a file on a remote computer.**
**Cause:** Both computers must reside in the same Active Directory domain to use Encrypting File System (EFS) on remote shares. The remote computer must also be trusted for delegation in the Active Directory.

**Solution:** Join both computers to the same Active Directory domain and trust the remote computer for delegation.

**See also:** Join a domain and Enable a remote server for file encryption.

**The message "You do not have permission to request a certificate based on the selected certificate template" appears when using theCreate Data Recovery Agentoption to create an EFS data recovery agent.**

**Cause:** EFS is unable to connect to a certification authority (CA) to obtain a new certificate and store it in Active Directory.

**Solution:** Install a certification authority to issue certificates for EFS.

**See also:** Certificate Services.

# Public Key Infrastructure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Public Key Infrastructure**
A *public key infrastructure*, often shortened to PKI, is a system of digital certificates, certification authorities (CAs) and other registration authorities (RAs) that verify and authenticate the validity of each party involved in an electronic transaction through the use of public key cryptography.

- Certificates
- Certificate Services
- Certificate Templates

# Certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificates**
This topic discusses certificates and the Certificates snap-in, which enables you to manage certificates for a user, computers, or services.

- For help with specific tasks, see Certificates How To ....
- For general background information, see Certificates Concepts.
- For problem-solving instructions, see Certificates Troubleshooting.

# Certificates How To ...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Add the Certificates Snap-in to the MMC Console
- Set Display Options of the Certificates Snap-in
- Request Certificates
- Manage an Individual Certificate
- Import and Export Certificates
- Configure credential roaming
- Use Windows 2000 Certificate Services Web Pages
- Use Windows Server 2003 Certificate Services Web Pages
- Manage Certificate Trusts

# Add the Certificates Snap-in to the MMC Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Add the Certificates snap-in to the MMC console**

- Manage certificates for your user account
- Manage certificates for a computer
- Manage certificates for a service

# Manage certificates for your user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage certificates for your user account**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Snap-in**, double-click **Certificates**, and then:
   - If you are logged on as an administrator, click **My user account**, and then click **Finish**.
   - If you are logged on as a user, Certificates automatically loads.
4. Click **Close**.

   **Certificates - Current User** appears on the list of selected snap-ins for the new console.
5. If you have no more snap-ins to add to the console, click **OK**.
6. To save this console, on the **File** menu, click **Save**.

## Notes

- To switch between managing certificates for your user account, a computer, or a service, you must have separate instances of Certificates added into the console.
- For additional help with creating and using MMC consoles, on the **Help** menu, click **Help Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Save an MMC console file
Starting the Certificates snap-in
Manage certificates for a computer
Manage certificates for a service

# Manage certificates for a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage certificates for a computer**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Snap-in**, double-click **Certificates**, click **Computer account**, and then click **Next**.
4. Do one of the following:
   - To manage certificates for the local computer, click **Local computer**, and then click **Finish**.
   - To manage certificates for a remote computer, click **Another computer** and type the name of the computer, or click **Browse** to select the computer name, and then click **Finish**.
5. Click **Close**.

   **Certificates (*Computer Name*)** appears on the list of selected snap-ins for the new console.
6. If you have no more snap-ins to add to the console, click **OK**.
7. To save this console, on the **File** menu, click **Save**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To manage certificates for another computer, you can either create another instance of Certificates in the console, or right-click **Certificates (*Computer Name*)** and click **Connect to Another Computer**.
- For more information on using Microsoft Management Console (MMC), on the **Help** menu, click **Help Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Save an MMC console file
Starting the Certificates snap-in
Manage certificates for your user account
Manage certificates for a service

# Manage certificates for a service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage certificates for a service**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Snap-in**, double-click **Certificates**, click **Service account**, and then click **Next**.
4. Do one of the following:
   - To manage certificates for services on your local computer, click **Local computer**, and then click **Next**.
   - To manage certificates for a remote computer, click **Another computer** and type the name of the computer, or click **Browse** to select the computer name, and then click **Next**.
5. Click the service for which you are managing certificates.
6. Click **Finish**, and then click **Close**.

   **Certificates - Service (*Service Name*) on *Computer Name*** appears on the list of selected snap-ins for the new console.
7. If you have no more snap-ins to add to the console, click **OK**.
8. To save this console, on the **File** menu, click **Save**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To manage certificates for a service on another computer, you can either create another instance of Certificates in the console or right-click **Certificates - Service (*Service Name*) on *Computer Name*** and click **Connect to Another Computer**.
- For additional help on using Microsoft Management Console (MMC), click **Help Topics** on the Help menu.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Save an MMC console file
Starting the Certificates snap-in
Manage certificates for your user account

Manage certificates for a computer

# Set Display Options of the Certificates Snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set display options of the Certificates snap-in**

- Display certificate stores in Logical Store mode
- Display certificate stores in Purpose mode
- Display certificate stores storage structure
- Display archived certificates

# Display certificate stores in Logical Store mode

Updated: February 28, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To display certificate stores in Logical Store mode**

1. Open the Certificates console for the user, computer, or service you want to manage.

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. With the **Certificates** <*console type*>, where <console type> is Current User, Computer, or Service, selected in the console pane, click the **View** menu, and then click **Options**.

3. Under **Organize view mode by**, click **Logical Certificate Stores**.

**Note**

- When you look at the contents of a certificate store in Logical Store mode, you will occasionally see what appears to be two copies of the same certificate in the store. This occurs because the same certificate is stored in separate physical stores under a logical store. When the contents of the physical certificates stores are combined into one logical store view, both instances of the same certificate are displayed.

  You can verify this by setting the view option to show the physical certificate stores and then noting that the certificate is stored in separate physical stores under the same logical store. You can verify that it is the same certificate by comparing the serial numbers.

  For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Certificate stores
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Purpose mode
Display archived certificates
Display certificate stores storage structure

# Display certificate stores in Purpose mode

Updated: January 21, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To display certificate stores in Purpose mode**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. On the **View** menu, click **Options**.

3. Under **Organize view mode by**, click **Certificate purpose**.

   **Note**

   For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Certificate stores
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode
Display archived certificates

# Display certificate stores storage structure

Updated: January 21, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To display the certificate stores storage structure**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. On the **View** menu, click **Options**.

   **Note**

   For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

3. Under **Organize view mode by**, click **Logical Certificate Stores**.

4. Under **Show the following**, select the **Physical certificate stores** check box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Purpose mode
Display certificate stores in Logical Store mode
Display archived certificates

# Display archived certificates

Updated: January 24, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Display archived certificates**

Archived are certificates that have expired or have been renewed. In many cases, it is good practice to retain archived certificates instead of deleting them. For example, you would need to keep an archived certificate to verify digital signatures on old documents signed using the key on the now-expired or renewed certificate.

**To display archived certificates**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. Ensure you have **Certificates** <*console type*> selected. Where <console type> is Current User, Computer, or Service, depending upon the category of certificates you want to view.

3. click the **View** menu, and then click **Options**.

4. In the **View Options** dialog box, in **Show the following**, select the **Archived certificates**.

**Tip**

For more information on the **Certificates console** and changing **View Options**, see Certificates Console (http://go.microsoft.com/fwlink/?LinkID=209851).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode

# Request Certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Request certificates**

- Request a certificate
- Request a certificate with the same key
- Renew a certificate with a new key
- Renew a certificate with the same key
- Request Certificates

# Request a certificate

Updated: January 21, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To request a certificate**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. In the console tree, select the correct store that has the certificate you want to request. For example, under **Certificates – Current User**, expand **Personal**, click **Certificates**.

   **Note**

   For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

3. In the details pane, click the certificate associated with the public key you want to associate a new certificate.

4. On the **Action** menu, point to **All Tasks**, and then click **Request New Certificate** to start the Certificate Request Wizard.

5. In the Certificate Request Wizard, provide the following information:

   - Click the type of certificate you want to request.
   - (Optional) If you have selected the **Advanced** check box, click the cryptographic service provider (CSP) you are using.
   - (Optional) If you have selected the **Advanced** check box, click the key length (measured in bits) of the public key associated with the certificate.
   - (Optional) If you have selected the **Advanced** check box and you want to enable strong private key protection, select the **Enable strong private key protection** check box. Enabling strong private key protection will ensure that you are prompted for a password every time the private key is used. This is useful if you want to make sure that the private key is not used without your knowledge.
   - (Optional) If you have selected the **Advanced** check box and if you have more than one certification authority (CA) available, choose the name of the certification authority that will issue the certificate.
   - A friendly name for your new certificate.

6. After the Certificate Request Wizard has successfully finished, click **OK**.

**Note**

- You can use this procedure to request certificates from an enterprise certification authority only. To request certificates from a stand-alone certification authority, you need to request certificates via Web pages. A Windows certification authority has its Web pages located at http://*servername*/certsrv, where *servername* is the name of the server hosting the certification authority.

- You can use this procedure to request certificates based on templates where the subject name is generated by Windows only. To request a certificate based on a template configured to obtain the subject name from the subject, you need to request certificates via Web pages.

- If the requested certificate type requires approval before it is issued, you need to retrieve the completed certificate via Web pages. A Windows certification authority has its Web pages located at http:// *servername*/certsrv, where *servername* is the name of the server hosting the certification authority.

- In order to request a Digital Signature Standard (DSS) certificate from an enterprise CA, you must select the User Signature Only certificate template in the Certificate Request Wizard.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Requesting certificates
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode
Request a certificate with the same key
Certificates and certification authorities
Using Windows 2000 Certificate Services Web pages

# Request a certificate with the same key

Updated: January 24, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To request a certificate with the same key**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. In the console tree, select the correct store that has the certificate you want to request. For example, under **Certificates – Current User**, expand **Personal**, click **Certificates**.

   **Note**

   For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

3. In the details pane, click the certificate associated with the public key you want to associate a new certificate.

4. On the **Action** menu, point to **All Tasks**, and then click **Request Certificate with Same Key** to start the Certificate Request Wizard.

5. In the Certificate Request Wizard, provide the following information:

   - Click the type of certificate you want to request.
   - (Optional) If you have selected the **Advanced** check box and you have more than one certification authority available, choose the name of the certification authority that will issue the certificate.
   - Type a friendly name for your new certificate.

6. After the Certificate Request Wizard has successfully finished, click **OK**.

**Warning**

- To perform this procedure, the view mode must be organized by Logical Certificate Stores. For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).
- You can use this procedure to request certificates from an enterprise certification authority only. To request certificates from a stand-alone certification authority, you must request certificates by way of Web pages. On Microsoft Certification Authority (CA) installations, Web pages of a Windows certification authority are located at http://*servername*/certsrv, where *servername* is the name of the server hosting the certification authority, by default.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Request a certificate
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode
Find a certificate
Certificates and certification authorities
Using Windows 2000 Certificate Services Web pages

# Renew a certificate with a new key

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To renew a certificate with a new key**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. Confirm that you are in Certificate purpose view. For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

3. In the console tree, under **Personal**, click **Certificates**.

4. In the details pane, click the certificate you are renewing.

5. On the **Action** menu, point to **All Tasks**, and then click **Renew Certificate with New Key** to open the Certificate Renewal Wizard.

6. In the **Certificate Renewal Wizard**, do one of the following:

   - Use the default values to renew the certificate.

   - (For advanced users only) Provide your own certificate renewal settings. You need to know the cryptographic service provider (CSP) and the certification authority (CA) issuing the certificate.

     You need to select the key length (measured in bits) of the public key associated with the certificate.

     You can also choose to enable strong private key protection. Enabling strong private key protection ensures that you are prompted for a password every time the private key is used. This is useful if you want to make sure that the private key is not used without your knowledge.

7. After the Certificate Renewal Wizard has successfully finished, click **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage certificates for your user account
Manage certificates for a computer

Manage certificates for a service
Display certificate stores in Logical Store mode
Display archived certificates
Renew a certificate with the same key
Using Windows 2000 Certificate Services Web pages

# Renew a certificate with the same key

Updated: February 14, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To renew a certificate with the same key**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. In the console tree, select the correct store that has the certificate you want to request. For example, under **Certificates – Current User**, expand **Personal**, click **Certificates**.

   **Note**

   For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

3. In the details pane, click the certificate you are renewing.

4. On the **Action** menu, point to **All Tasks**, and then click **Renew Certificate with Same Key** to start the Certificate Renewal Wizard.

5. In the Certificate Renewal Wizard, do one of the following:

   - Use the default values to renew the certificate.
   - Provide your own certificate renewal settings. You need to know the certification authority issuing the certificate.

6. After the Certificate Renewal Wizard has successfully finished, click **OK**.

**Note**

- To perform this procedure, the view mode must be organized by Logical Certificate Stores. For more information, see Related Topics.
- Once renewed, the old certificate will be archived. For information on how to view archived certificates, see Related Topics.
- You can use this procedure to request certificates from an enterprise certification authority only. To request certificates from a stand-alone certification authority, you need to request certificates via Web pages. A Windows certification authority has its Web pages located at http://*ServerName*/certsrv, where *ServerName* is the name of the server that hosts the certification authority.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode
Display archived certificates
Renew a certificate with a new key
Using Windows 2000 Certificate Services Web pages

# Trigger automatic enrollment for certificates

Updated: February 28, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To trigger automatic enrollment for certificates**

1. Open the Certificates console for the user or computer certificate console you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account or computer account, see the appropriate article from the following list

   - Manage certificates for your user account
   - Manage certificates for a computer
   - For more information on the Certificates console and changing View Options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkID=209851).

2. On the **Action** menu, point to **All Tasks**, and then click **Automatically Enroll Certificates**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Requesting certificates
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Certificate autoenrollment
Certificate Services example implementation: Establishing autoenrollment for user certificates

# Manage an Individual Certificate

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage an individual certificate**

- View a certificate
- Find a certificate
- Move a certificate
- Copy a certificate
- Delete a certificate
- Modify the properties of a certificate

# View a certificate

Updated: January 24, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view a certificate**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. In the Console pane, select the container that holds the certificate you want to view (for example, under **Certificates – Current User**, expand **Personal**, and then select **Certificates**). For information on the different **Certificates** console view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

3. In the details pane, double-click the certificate.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Viewing certificate information
Find a certificate
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display archived certificates

# Find a certificate

Updated: January 24, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a certificate**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. On the console tree, click **Certificates** - *Certificate Holder*, where Certificate Holder is a Certificates instance (user, computer or service). For more information regarding Certificate console View Options, see Certificates Console(http://go.microsoft.com/fwlink/?LinkID=209851)

3. On the **Action** menu, click **Find Certificates**.

4. Enter the following criteria to find certificates:

   - The certificate store to search
   - Text to search for
   - The certificate field to search for the specified text. You can look for text in the **Issued By**, **Issued To**, **MD5 Hash**, **Serial Number**, and **SHA1 Hash** fields.

5. Click **Find Now**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Certificate stores

# Move a certificate

Updated: January 24, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move a certificate**

1. Open the certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. Under the appropriate logical store (e.g. **Personal**, **Trusted Root Certification Authorities**, **Intermediate Certification Authorities**, and so on) where the certificate is stored, click **Certificates**.

3. In the details pane, click the certificate you are moving. (To select more than one certificate, hold down CTRL and click each certificate.)

4. On the **Action** menu, click **Cut**.

5. Click the logical certificate store where you want to move the certificate.

6. On the **Action** menu, click **Paste**.

**Note**

- To perform this procedure, the view mode must be organized by Logical Certificate Stores. For more information, see Display certificate stores in Logical Store mode.

- You can only move an object within certificate stores in Certificates. For example, you cannot move an object to a folder in Windows Explorer. See Related Topics for information about importing and exporting certificates.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Find a certificate
Certificate stores
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode
Export a certificate

Import a certificate
Copy a certificate

# Copy a certificate

Updated: February 14, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy a certificate**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. In the console tree, select the correct store that has the certificate you want to request. For example, under **Certificates – Current User**, expand **Personal**, click **Certificates**.

   **Note**

   For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

3. In the details pane, click the certificate you are copying. (To select multiple certificates, hold down CTRL and click each certificate.)

4. On the **Action** menu, click **Copy**.

5. Click the logical certificate store where you want to copy the certificate.

6. On the **Action** menu, click **Paste**.

**Note**

1. To perform this procedure, the view mode must be organized by Logical Certificate Stores. For instructions on changing to Logical Certificate Stores view mode, see Display certificate stores in Logical Store mode.

2. You can only copy a certificate within certificate stores in Certificates. For example, you cannot copy a certificate to a folder in Windows Explorer. See Related Topics for information about importing and exporting certificates.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Find a certificate

Certificate stores
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode
Export a certificate
Import a certificate
Move a certificate

# Delete a certificate

Updated: January 21, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a certificate**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. In the console tree, select the correct store that has the certificate you want to request. For example, under **Certificates – Current User**, expand **Personal**, click **Certificates**.

   **Note**

   For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

3. In the details pane, click the certificate (or certificates) you want to delete. (To select multiple certificates, hold down CTRL and click each certificate.)

4. On the **Action** menu, click **Delete**.

5. Click **Yes** if you are sure you want to permanently delete the certificate.

**Note**

- In some instances, there is not a **Certificates** folder in the console tree. In that case, navigate in the console tree until the certificate you want appears in the details pane, and then continue the procedure.
- You might want to back up the certificate by exporting it before you delete it. For the procedure to export a certificate, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Find a certificate
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service

Display archived certificates
Export a certificate

# Modify the properties of a certificate

Updated: January 24, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify the properties of a certificate**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. On the **Action** menu, click **Properties**.

3. To change the friendly name of the certificate, in **Friendly name**, type the new name.

4. To change the description of the certificate, in **Description**, type the new description.

5. To enable all purposes for the certificate, based on the certificate type, in **Certificate purposes**, click **Enable all purposes for this certificate**.

6. To disable all purposes for the certificate, in **Certificate purposes**, click **Disable all purposes for this certificate**.

7. To designate specific purposes for the certificate, in **Certificate purposes**, click **Enable only the following purposes** and select the appropriate check boxes.

**Note**

- If you are logged on as a user, you can only modify the properties of certificates in the **Personal** store.
- To modify the properties of a certificate, double-click the certificate, click the **Details** tab, and then click **Edit Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Find a certificate
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display archived certificates
View a certificate

# Import and Export Certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Import and export certificates**

- Import a certificate
- Export a certificate
- Export a certificate with the private key
- View the certificates in a PKCS #7 file

# Import a certificate

Updated: February 28, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import a certificate**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   - Manage certificates for your user account
   - Manage certificates for a computer
   - Manage certificates for a service
   - For more information on the Certificates console and changing View Options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkID=209851).

2. Select a certificate store for the certificate type you want to import. For example, you might select the **Personal** store under **Certificates (Local Computer)**.

3. On the **Action** menu, point to **All Tasks** and then click **Import** to start the Certificate Import Wizard.

4. Type the file name containing the certificate to be imported. (You can also click **Browse** and navigate to the file.)

5. If it is a PKCS #12 file, do the following:

   - Type the password used to encrypt the private key.
   - (Optional) If you want to be able to use strong private key protection, select the **Enable strong private key protection** check box.
   - (Optional) If you want to back up or transport your keys at a later time, select the **Mark key as exportable** check box.

6. Do one of the following:

   - If the certificate should be automatically placed in a certificate store based on the type of certificate, click **Automatically select the certificate store based on the type of certificate**.
   - If you want to specify where the certificate is stored, select **Place all certificates in the following store**, click **Browse**, and choose the certificate store to use.

**Note**

- To perform this procedure, the view mode must be organized by Logical Certificate Stores. For more information on the **Certificates console** and changing **View Option**s, see Certificates Console (http://go.microsoft.com/fwlink/?LinkID=209851).

- You should only import certificates obtained from trusted sources. Importing an altered or unreliable certificate could compromise the security of any system component that uses the imported certificate.

- You can import a certificate into any logical store. In most cases, you will import certificates into the Personal store or the Trusted Root Certification Authorities store, depending on whether the certificate is intended for you or if it is a root CA certificate.

- Enabling strong private key protection will ensure that you are prompted for a password every time the private key is used. This is useful if you want to make sure that the private key is not used without your knowledge.

- To open the Certificate Import Wizard, you can open Windows Explorer. Then, in the cases of DER (.cer), Base-64 (.cer), and PKCS #7 (.p7b) files, right-click the file and click **Install Certificate**. In the case of a PKCS #12 (.pfx) file, right-click the file and click **Install PFX**.

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

- The file from which you import certificates will remain intact after you have completed importing the certificates. You can use Windows Explorer to delete the file if it is no longer needed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Importing and exporting certificates
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode
Export a certificate
Move a certificate

# Export a certificate

Updated: January 21, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export a certificate**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. In the console tree, select the correct store that has the certificate you want to request. For example, under **Certificates – Current User**, expand **Personal**, click **Certificates**.

   **Note**

   For information on the different view options, see Certificates Console (http://go.microsoft.com/fwlink/?LinkId=209851).

3. In the details pane, click the certificate you want to manage.

4. On the **Action** menu, point to **All Tasks**, and then click **Export**.

5. In the **Certificate Export Wizard**, click **No, do not export the private key**. (This option will appear only if the private key is marked as exportable and you have access to the private key.)

6. Provide the following information in the **Certificate Export Wizard**:

   - Click the file format you want to use to store the exported certificate: a DER-encoded file, a Base64-encoded file, or a PKCS #7 file.

   - If you are exporting the certificate to a PKCS #7 file, you also have the option to include all certificates in the certification path.

**Warning**

- For the procedure to export a certificate with the private key, see Related Topics.

- After the **Certificate Export Wizard** is finished, the certificate will remain in the certificate store in addition to being in the new file. If you want to remove the certificate from the certificate store, you must delete it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Importing and exporting certificates
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode
Export a certificate with the private key
Import a certificate
Delete a certificate

# Export a certificate with the private key

Updated: January 24, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export a certificate with the private key**

1. Open the Certificates console for the user, computer, or service you want to manage.

   **Tip**

   For instructions on creating a Microsoft Management Console (MMC) that allows you to manage the certificates of a user account, computer account, or service account, see the appropriate article from the following list

   a. Manage certificates for your user account

   b. Manage certificates for a computer

   c. Manage certificates for a service

2. In the console pane, select the certificate store and container holding the certificate that you want to export.

   **Tip**

   For more information on the **Certificates console** and changing **View Option**s, see Certificates Console (http://go.microsoft.com/fwlink/?LinkID=209851).

3. In the details pane, click the certificate you want to export.

4. On the **Action** menu, point to **All Tasks**, and then click **Export**.

5. In the Certificate Export Wizard, click **Yes, export the private key**. (This option will appear only if the private key is marked as exportable and you have access to the private key.)

6. Under **Export File Format**, do one or all of the following, and then click **Next**.

   - To include all certificates in the certification path, select the **Include all certificates in the certification path if possible** check box.
   - To enable strong protection, select the **Enable strong protection (requires IE 5.0, NT 4.0 SP4 or above)** check box.
   - To delete the private key if the export is successful, select the **Delete the private key if the export is successful** check box.

7. In **Password**, type a password to encrypt the private key you are exporting. In **Confirm password**, type the same password again, and then click **Next**.

8. In **File name**, type a file name and path for the PKCS #12 file that will store the exported certificate and private key, click **Next**, and then click **Finish**.

**Note**

- If a certificate was issued from a Windows Server 2003 certification authority, the private key for that certificate is only exportable if the certificate request was made via the Advanced Certificate Request certification authority Web page with the **Mark keys as exportable** check box selected, or if the certificate is for EFS (Encrypting File System) or EFS recovery.

- Strong protection (also known as iteration count) is enabled by default in the Certificate Export Wizard when you export a certificate with its associated private key.

  Strong protection is not compatible with older programs, so you need to clear the **Enable strong protection** option if you are going to use the private key with any browser earlier than Microsoft Internet Explorer 5.

- After the **Certificate Export Wizard** is finished, the certificate will remain in the certificate store in addition to being in the newly-created file. If you want to remove the certificate from the certificate store, you will need to delete it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Importing and exporting certificates
Manage certificates for your user account
Manage certificates for a computer
Manage certificates for a service
Display certificate stores in Logical Store mode
Export a certificate
Import a certificate
Delete a certificate

# View the certificates in a PKCS #7 file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the certificates in a PKCS #7 file**

1. Open Windows Explorer.
2. Locate the PKCS #7 file in which you want to view the certificates.
3. In the details pane, right-click the PKCS #7 file and then click **Open**. **Certificates** loads.
4. In the console tree, double-click the folder containing the PKCS #7 file, and then click **Certificates**. The certificates contained in the PKCS #7 file are displayed in the details pane.

## Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- A PKCS #7 file typically has a .p7b file name extension, but this is not always the case. As with any other data file, the creator of the file has control over the name and whether or not the .p7b extension is used.
- For procedures to import and export a PKCS #7 file, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Import a certificate
Export a certificate

# Configure credential roaming

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2


Checklist: Credential roaming

Configure credential roaming Group Policy

# Checklist: Credential roaming

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Credential roaming**

| Step | Reference |
| --- | --- |
| Review concepts about credential roaming. | Using credential roaming |
| Review credential roaming best practices. | Credential roaming best practices |
| Review Active Directory schema concepts and practices. | Extending the schema |
| Extend the Active Directory schema in a test environment by using the credential roaming schema extensions. After verifying the Active Directory schema extension in a test environment, implement the change in the production environment. | Credential roaming schema changes |
| Create the credential roaming administrative template. | Credential roaming administrative template |
| Configure credential roaming Group Policy. | Configure credential roaming Group Policy |

# Configure credential roaming Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

To configure credential roaming Group Policy, you need to:

- Install the Credential Roaming administrative template.
- Set credential roaming Group Policy options.

**Install the DIMS administrative template**

To install the CredentialRoaming.adm template:

1. Open the **Active Directory Users and Computers** MMC snap-in.
2. In the console tree, highlight the domain node, right-click and select **Properties**.
3. With the domain properties window open, select the **Group Policy** tab.
4. Highlight the domain policy in the domain policy objects area and click **Edit**.
5. In the console tree, expand **User Configuration\Administrative Templates**.
6. Right-click **Administrative Templates**. Choose **Add/Remove Templates**.
7. Click **Add**.
8. Navigate to the folder where the CredentialRoaming.adm file is located. Select the CredentialRoaming.adm file, click **Open**, and then click **Close**. **Credential Roaming** will now appear as a template under **Administrative Templates**.

**Setting credential roaming Group Policy options**

To configure credential roaming Group Policy options:

1. Open the **Group Policy Object Editor**, and go to:

   ```
   User Configuration)/Administrative Templates/ Credential Roaming
   ```

2. In the details pane, select and right-click **X.509 Certificate and Key Roaming**. Click **Properties**.
3. Under **X.509 Certificate and Key Roaming Settings**, click **Enabled**.
4. In the **Enable Credential Roaming** drop-down list, select the type of credential roaming policy appropriate for your organization:

   - **Roam all using strict arbitration rules**
   - **Roam all x.509 certificates and keys**
   - **Roam Encryption only with strict arbitration**
   - **Roam only encryption certificates and keys**

   **Note**

- In most circumstances, **Roam all x.509 certificates and keys** will be the preferred option. The strict arbitration options can be used when a user may be issued multiple certificates or keys for a similar purpose but with different settings, and you want one specific group of settings to win out.

5. If appropriate for your organization, you can change the default values for the following options:

   - **Maximum tombstone credentials lifetime in days:**
   - **Maximum number of roaming credentials per user:**
   - **Maximum size (in bytes) of a roaming credential:**

6. When you are done, click **OK**.
7. Close the **Group Policy Object Editor**.
8. Click **OK** to accept the new domain Group Policy settings properties and close **Active Directory Users and Computers**.

**Notes**

- You must be a member of the Domain Admins group to define domain Group Policy options.
- New domain Group Policy settings are not applied to clients until after Group Policy is refreshed — typically every eight hours.
- Credential roaming is supported for clients running Windows Server 2003 with Service Pack 1 (SP1). Credential roaming is supported on domain controllers running Windows 2000 Service Pack 3 (SP3) or later or Windows Server 2003. Mixed environments containing domain controllers running Windows 2000 SP3 or later and Windows Server 2003 are also supported. The forest functional level can be either Windows 2000 or Windows Server 2003.
- An update was recently made available that enables credential roaming on clients running Windows XP Professional SP2 and an update is available for Windows Server 2003 (SP1).

  For the best performance and security, it is recommended that the servers be upgraded to Windows Server 2003 SP1. For other best practices, see Credential roaming best practices.
- For an explanation of the different credential roaming Group Policy settings and to create the CredentialRoaming.adm template, see Credential roaming administrative template.

**See Also**
**Concepts**
Open Group Policy from Active Directory Users and Computers
Open Group Policy from Active Directory Sites and Services
Ways to open Group Policy Object Editor
Edit the local Group Policy object
Add or remove an Administrative Template (.adm file)

# Use Windows 2000 Certificate Services Web Pages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Windows 2000 Certificate Services Web pages**

- Submit a user certificate request via the Web
- Submit an advanced certificate request via the Web
- Request a certificate using a PKCS #10 or PKCS #7 file
- Save a certificate request to a PKCS #10 file
- Check on a pending certificate request
- Retrieve a certification authority certificate
- Retrieve a certificate revocation list

# Submit a user certificate request via the Web

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To submit a user certificate request via the Web**

1. Open Internet Explorer.
2. In **Address**, type http://*ServerName*/certsrv, where *ServerName* is the name of the Windows 2000 Web server where the certification authority (CA) you want to access is located.
3. Click **Request a certificate**.
4. In **Request a Certificate**, click the type of certificate you want to request.
5. Under ***Certificate Type* -Identifying Information**, do one of the following:
   - If you see the message "No further identifying information is required. To complete your certificate, press submit," click **Submit**.
   - Enter your identifying information for the certificate request, and then click **Submit**.
6. Do one of the following:
   - If you see the **Certificate Pending** Web page, see Related Topics below for the procedure to check on a pending certificate.
   - If you see the **Certificate Issued** Web page, click **Install this certificate**.
7. If you are finished using the Certificate Services Web pages, close Internet Explorer.

## Notes

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- Running Web enrollment on a Windows 2000 Web server acting as a Web enrollment agent for a Windows Server 2003 CA may not provide the full functionality of the CA, because it was not designed to fully integrate with the Windows Server 2003 CA. Limits may occur with smart card enrollment or enrollment for a certificate based on a Version 2 certificate template. To provide full functionality through the Web enrollment agent, run it on a Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows 2000 Certificate Services Web pages
Check on a pending certificate request
Submit an advanced certificate request via the Web
Request a certificate using a PKCS #10 or PKCS #7 file

Use Windows Server 2003 Certificate Services Web Pages

# Submit an advanced certificate request via the Web

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To submit an advanced certificate request via the Web**

1. Open Internet Explorer.
2. In **Address**, type **http://***ServerName***/certsrv**, where *ServerName* is the name of the Windows 2000 Web server where the certification authority that you want to access is located.
3. Click **Request a certificate**.
4. Click **advanced certificate request**.
5. Click **Creae and submit a request to this CA**.
6. Fill in any identifying information requested and any other options you require and click **Submit**.
7. Do one of the following:
   - If you see the **Certificate Pending** Web page, see Related Topics below for the procedure to check on a pending certificate.
   - If you see the **Certificate Issued** Web page, click **Install this certificate**.
8. If you are finished using the Certificate Services Web pages, close Internet Explorer.

**Notes**

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- A Windows XP computer cannot request a certificate on behalf of another subject when using the Windows 2000 certification authority Web pages.
- Running Web enrollment on a Windows 2000 Web server that is acting as a Web enrollment agent for a Windows Server 2003 family CA may not provide the full functionality of the CA, because it was not designed to fully integrate with the Windows Server 2003 family CA. Limits may occur with smart card enrollment or enrollment for a certificate based on a Version 2 certificate template. To provide full functionality through the Web enrollment agent, run it on a Web server that is running a Windows Server 2003 operating system.
- Using the **Advanced Certificate Request** Web page, you can set the following options for each certificate requested:
  - Certificate template (from enterprise certification authorities) or **Intended purposes** (from stand-alone certification authorities). Indicates what applications the public key in the certificate can be used for, such as client authentication or e-mail.
  - Cryptographic service provider (CSP).
  - **Key size.** The length, in bits, of the public key on the certificate. In general, the longer the key, the more secure it is.
  - Hash algorithm

- **Key usage.** How the private key can be used. "Exchange" means that the private key can be used to enable the exchange of sensitive information. "Signature" means that the private key can be used only to create a digital signature. "Both" means that the key can be used for both exchange and signature functions.

- **Create a new key set or use an existing key set.** You can use an existing public and private key pair stored on your computer or create a new public and private key pair for a certificate. For more information about the issues of reusing a key as opposed to generating a new key, see the resources on Certificates Resources.

- **Enable strong private key protection.** When you enable strong private key protection, you are prompted for a password every time the private key needs to be used.

- **Mark keys as exportable.** When you mark keys as exportable, you can save the public and private key to a PKCS #12 file. This is useful if you change computers and want to move the key pair, or if you want to remove the key pair and secure them in another location.

- **Use the local machine store.** Select this option if the computer will need access to the private key associated with the certificate when other users are logged on. Select this option when requesting certificates intended to be issued to computers (such as Web servers) instead of certificates issued to people.

- **Save the request to PKCS #10 file.** This is useful if the certification authority is unavailable for processing certificate requests online. For information on creating and submitting a certificate request using a PKCS #10 file, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows 2000 Certificate Services Web pages
Save a certificate request to a PKCS #10 file
Check on a pending certificate request
Submit a user certificate request via the Web
Request a certificate using a PKCS #10 or PKCS #7 file
Use Windows Server 2003 Certificate Services Web Pages

# Request a certificate using a PKCS #10 or PKCS #7 file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To request a certificate using a PKCS #10 or PKCS #7 file**

1. Open Internet Explorer.
2. In **Address**, type http://*servername*/certsrv, where *servername* is the name of the Windows 2000 Web server where the certification authority (CA) you want to access is located.
3. Click **Request a certificate**, and then click **advanced certificate request**.
4. Click **Submit a certificate request using a base64-encoded CMC or PKCS #10 file, or submit a renewal request by using a base-64-encoded PKCS #7 file**.
5. Do one of the following:
   - Open Notepad. On the **File** menu, click **Open**. Select the PKCS #10 or PKCS #7 file and click **Open**. On the **Edit** menu, click **Select all**, and then, on the **Edit** menu, click **Copy**. On the Web page, click in the **Saved request** scroll box. On the **Edit** menu, click **Paste** to paste the contents of certificate request into the scroll box.
   - If your Web browser security settings do not prohibit a Web page from accessing your disk, you can click **Browse for a file to insert** to locate the file you want to use for the certificate request. If you get a warning about the ActiveX control, click **Yes** to allow it to run, then click **Browse**. After locating and selecting the file you want to use for the certificate request, click **Read!**. On the Web page, click **Read!** to paste the contents of the file into the scroll box. See the note about using **Browse**.
6. If you are connected to an enterprise CA, choose the certificate template you want to use.
7. Click **Submit**.
8. Do one of the following:
   - If you see the **Certificate Pending** Web page, see Related Topics below for the procedure to check on a pending certificate.
   - If you see the **Certificate Issued** Web page, click **Download certificate**. Choose to save the file to your hard disk, and then import the certificate into your certificate store. For the procedure to import a certificate, see Related Topics.

**Notes**

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- To open Notepad, click **Start**, point to **All programs**, point to **Accessories**, and then click **Notepad**.
- In general, you use a PKCS #10 file to submit a request for a new certificate and a PKCS #7 file to submit a request to renew an existing certificate. Submitting requests with files is useful when the certificate requester is unable to submit a request online to the certification authority.
- If you use a PKCS #7 file to submit a request to renew an existing certificate, the renewal request file must be generated by a program. This procedure will not work with a PKCS #7 file that merely contains an exported certificate.

- You might need to make http://*servername* a trusted site for Internet Explorer in order to browse for a file on the computer's disk drive. To make http://*servername* a trusted site, in Internet Explorer, click **Tools**, then point to **Internet Options**, point to **Security**, point to **Trusted Sites**, and click **Sites**. Type **http://*servername***, and click **OK**.

- If you submit the request and immediately get a message asking you if you want to submit the request even though it does not contain a "BEGIN" or "END" tag, click **OK**.

- For the procedure to export a certificate to create a PKCS #7 file, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows 2000 Certificate Services Web pages
Save a certificate request to a PKCS #10 file
Export a certificate
Import a certificate
Importing and exporting certificates
Use Windows Server 2003 Certificate Services Web Pages

# Save a certificate request to a PKCS #10 file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save a certificate request to a PKCS #10 file**

1. Open Internet Explorer.
2. In **Address**, type http://*servername*/certsrv, where *servername* is the name of the Windows 2000 Web server where the certification authority you want to access is located.
3. Click **Request a certificate**, and then click **Next**.
4. Click **Advanced request**, and then click **Next**.
5. Click **Submit a certificate request to this CA using a form**, and then click **Next**.
6. In **Certificate Template** and **Key Options**, select the options which you want to use for the certificate request.
7. Under **Additional Options**, select the **Save request to a PKCS #10 file** check box.
8. In **File name**, type a file name and click **Save**.
9. If you are finished using the Certificate Services Web pages, close Internet Explorer.

## Notes

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- The file is saved to your computer's desktop.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows 2000 Certificate Services Web pages
Request a certificate using a PKCS #10 or PKCS #7 file
Use Windows Server 2003 Certificate Services Web Pages

# Check on a pending certificate request

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check on a pending certificate request**

1. Open Internet Explorer.
2. In **Address**, type http://*servername*/certsrv, where *servername* is the name of the Windows 2000 Web server where the certification authority (CA) you want to access is located.
3. Click **Check on a pending certificate**, and then click **Next**.
4. If there are no pending certificate requests, you will see a message to that effect. Otherwise, select the certificate request you want to check and **Next**, and check the pending certificate request.
5. If you are finished using the Certificate Services Web pages, close Internet Explorer.

## Notes

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- The available status messages are:
  - **Still pending.** You must wait for the administrator of the certification authority to issue the certificate. To remove the certificate request, click **Remove**.
  - **Issued.** To install the certificate, click **Install this certificate**.
  - **Denied.** Contact the administrator of the certification authority for further information.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Submit a user certificate request via the Web
Submit an advanced certificate request via the Web
Request a certificate using a PKCS #10 or PKCS #7 file
Using Windows 2000 Certificate Services Web pages
Use Windows Server 2003 Certificate Services Web Pages

# Retrieve a certification authority certificate

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To retrieve a certification authority certificate**

1. Open Internet Explorer.
2. In **Address**, type http://*servername*/certsrv, where *servername* is the name of the Windows 2000 Web server where the certification authority (CA) you want to access is located.
3. Click **Retrieve the CA certificate or certificate revocation list**, and then click **Next**.
4. Do one of the following:
   - If you want to trust all the certificates issued by this certification authority (CA), click **Install this CA certification path**.
   - If the certification authority has been renewed, you will have the choice of which version of the CA certificate you want to download.
5. Under **Choose file to download**, click the CA certificate you want to download, and then click **Download CA Certificate**.
6. In **File Download**, click **Open this file from its current location**, and then click **OK**.
7. When the **Certificate** dialog box appears, click **Install certificate**.
8. In the Certificate Import Wizard, click **Automatically select the certificate store based on the type of certificate**.
9. If you are finished using the Certificate Services Web pages, close Internet Explorer.

## Notes

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- Choosing **Install this CA certification path** is useful when you want to establish trust in a subordinate CA, but you do not currently have the certificate of the root CA of the certification hierarchy in your certificate store.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows 2000 Certificate Services Web pages
Certificates and certification authorities
Use Windows Server 2003 Certificate Services Web Pages

# Retrieve a certificate revocation list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To retrieve a certificate revocation list**

1. Open Internet Explorer.
2. In **Address**, type http://*servername*/certsrv, where *servername* is the name of the Windows 2000 Web server where the certification authority (CA) you want to access is located.
3. Click **Retrieve the CA certificate or certificate revocation list (CRL)**, and then click **Next**.
4. Click **Download latest certificate revocation list**.
5. When the **File Download** dialog box appears, click **Save this file to disk**, and then click **OK**. Select a folder on your computer in which to store the .crl file, click **Save**, and then click **Close**.
6. Open Windows Explorer and locate the .crl file you just saved.
7. Right-click the .crl file and click **Install CRL**.
8. When the Certificate Import Wizard opens, click **Automatically select the certificate store based on the type of certificate**.
9. After completing the Certificate Import Wizard, close Internet Explorer if you are finished using the Certificate Services Web pages.

## Notes

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows 2000 Certificate Services Web pages
Certificates and certification authorities
Use Windows Server 2003 Certificate Services Web Pages

# Use Windows Server 2003 Certificate Services Web Pages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Windows Server 2003 Certificate Services Web pages**

- Submit a user certificate request via the Web to a Windows Server 2003 CA
- Submit an advanced certificate request via the Web to a Windows Server 2003 CA
- Request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file
- Save a certificate request to a Windows Server 2003 CA in a PKCS #10 file
- Check on a pending certificate request to a Windows Server 2003 CA
- Retrieve a certification authority certificate from a Windows Server 2003 CA
- Retrieve a certificate revocation list from a Windows Server 2003 CA

# Submit a user certificate request via the Web to a Windows Server 2003 CA

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To submit a user certificate request via the Web to a Windows Server 2003 CA**

1. Open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the Web server running Windows Server 2003 where the certification authority (CA) you want to access is located.
3. Click **Request a certificate**.
4. On **Request a Certificate**, click **User Certificate**.
5. On the User Certificate - Identifying Information page, do one of the following:
   - Read the message "No further identifying information is required. To complete your certificate, press submit".
   - Enter your identifying information for the certificate request.
6. (Optional) Click **More Options** to specify the cryptographic service provider (CSP) and whether you want to enable strong private key protection. (This means that you will receive a prompt every time that the private key associated with the certificate is used.)
7. Click **Submit**.
8. Do one of the following:
   - If you see the **Certificate Pending** Web page, see Related Topics below for the procedure to check on a pending certificate.
   - If you see the **Certificate Issued** Web page, click **Install this certificate**.
9. If you are finished using the Certificate Services Web pages, close Internet Explorer.

### Note

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows Server 2003 Certificate Services Web pages
Check on a pending certificate request to a Windows Server 2003 CA
Submit an advanced certificate request via the Web to a Windows Server 2003 CA

Request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file
Use Windows 2000 Certificate Services Web Pages

# Submit an advanced certificate request via the Web to a Windows Server 2003 CA

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To submit an advanced certificate request via the Web to a Windows Server 2003 CA**

1. Open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the Web server running Windows Server 2003 where the certification authority that you want to access is located.
3. Click **Request a certificate**.
4. Click **advanced certificate request**.
5. Click **Create and submit a certificate request to this CA**.
6. Fill in any identifying information requested and any other options you require.
7. Click **Submit**.
8. Do one of the following:
   - If you see the **Certificate Pending** Web page, see Related Topics below for the procedure to check on a pending certificate.
   - If you see the **Certificate Issued** Web page, click **Install this certificate**.
9. If you are finished using the Certificate Services Web pages, close Internet Explorer.

**Notes**

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- Using the **Advanced Certificate Request** Web page, you can set the following options for each certificate requested:
  - **The Certificate template** (from an enterprise certification authorities) or **Intended purposes** (from a stand-alone certification authorities). Indicates what applications the public key in the certificate can be used for, such as client authentication or e-mail.
  - **The Cryptographic service provider (CSP).** A CSP is responsible for creating keys, destroying them, and using them to perform a variety of cryptographic operations. Each CSP provides a different implementation of the CryptoAPI. Some provide stronger cryptographic algorithms, while others use hardware components, such as smart cards
  - **Key size.** The length, in bits, of the public key on the certificate. In general, the longer the key, the more secure it is.
  - **The Hash algorithm.** A good hash algorithm makes it computationally infeasible to construct two independent inputs that have the same hash. Typical hash algorithms include MD2, MD4, MD5, and SHA-1.

- **Key usage.** How the private key can be used. "Exchange" means that the private key can be used to enable the exchange of sensitive information. "Signature" means that the private key can be used only to create a digital signature. "Both" means that the key can be used for both exchange and signature functions.

- **Create a new key set or use an existing key set.** You can use an existing public and private key pair stored on your computer or create a new public and private key pair for a certificate. For more information that compares reusing a key and generating a new key, see the resources on Certificates Resources.

- **Enable strong private key protection.** When you enable strong private key protection, you will be prompted for a password every time the private key needs to be used.

- **Mark keys as exportable.** When you mark keys as exportable, you can save the public and private key to a PKCS #12 file. This is useful if you change computers and want to move the key pair, or if you want to remove the key pair and secure them in another location.

- **Use the local machine store.** Select this option if the computer will need access to the private key associated with the certificate when other users are logged on. Select this option when requesting certificates intended to be issued to computers (such as Web servers) instead of certificates issued to people.

- **Save the request to PKCS #10 file.** This is useful if the certification authority is unavailable for processing certificate requests online. For information on creating and submitting a certificate request using a PKCS #10 file, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows Server 2003 Certificate Services Web pages
Save a certificate request to a Windows Server 2003 CA in a PKCS #10 file
Check on a pending certificate request to a Windows Server 2003 CA
Submit a user certificate request via the Web to a Windows Server 2003 CA
Request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file
Use Windows 2000 Certificate Services Web Pages

# Request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file**

1. Open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the Web server running Windows Server 2003 where the certification authority (CA) you want to access is located.
3. Click **Request a certificate** and then click **advanced certificate request**.
4. Click **Submit a certificate request using a base-64-encoded CMC or PKCS #10 file, or submit a renewal request by using a base-64-encoded PKCS #7 file**.
5. Do one of the following:
   - In Notepad, click **File**, click **Open**, select the PKCS #10 or PKCS #7 file, click **Edit**, click **Select all**, click **Edit**, and click **Copy**. On the Web page, click in the **Saved request** scroll box. Click **Edit** and then click **Paste** to paste the contents of certificate request into the scroll box.
   - Click **Browse for a file to insert** to locate the file you want to use for the certificate request. If you get a warning about the ActiveX control, click **Yes** to allow it to run, then click the **Browse** button. After locating and selecting the file you want to use for the certificate request, click **Open**. On the Web page, click **Read!** to paste the contents of the file into the scroll box. See the note about using **Browse**.
6. If you are connected to an enterprise CA, choose the certificate template you want to use. By default, the appropriate template is named **Subordinate Certification Authority**.
7. If you have any attributes to add to the certificate request, enter them into **Additional Attributes**.
8. Click **Submit**.
9. Do one of the following:
   - If you see the **Certificate Pending** Web page, see Related Topics below for the procedure to check on a pending certificate.
   - If you see the **Certificate Issued** Web page, click **Download certificate chain**. Choose to save the file to your hard disk, and then import the certificate into your certificate store. For the procedure to import a certificate, see Related Topics.

**Notes**

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- To open Notepad, click **Start**, point to **All programs**, point to **Accessories**, and then click **Notepad**.
- In general, you use a PKCS #10 file to submit a request for a new certificate and a PKCS #7 file to submit a request to renew an existing certificate. Submitting requests with files is useful when the certificate requester is unable to submit a

request online to the certification authority.

- You might need to make http://*servername* a trusted site for Internet Explorer in order to browse for a file on the computer's disk drive. To make http://*servername* a trusted site, in Internet Explorer, click **Tools**, then point to **Internet Options**, point to **Security**, point to **Trusted Sites**, and click **Sites**. Type **http://*servername***, and click **OK**.

- If you submit the request and immediately get a message asking you if you want to submit the request even though it does not contain a BEGIN or END tag, click **OK**.

- For the procedure to export a certificate to create a PKCS #7 file, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows Server 2003 Certificate Services Web pages
Save a certificate request to a Windows Server 2003 CA in a PKCS #10 file
Export a certificate
Import a certificate
Importing and exporting certificates
Use Windows 2000 Certificate Services Web Pages

# Save a certificate request to a Windows Server 2003 CA in a PKCS #10 file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save a certificate request to a Windows Server 2003 CA in a PKCS #10 file**

1. Open Internet Explorer.
2. In Internet Explorer, open http://*servername*/certsrv, where *servername* is the name of the Web server running Windows Server 2003 where the certification authority you want to access is located.
3. Click **Request a certificate**, and then click **advanced certificate request**.
4. Click **Create and submit a request to this CA**.
5. Enter any identifying information requested and any other options you require.
6. Under **Additional Options**, select the **Save request to a PKCS #10 file** check box.
7. In the **Full path name** field, type a path and file name.
8. Click **Submit**. The file is saved to your computer's desktop.
9. If you are finished using the Certificate Services Web pages, close Internet Explorer.

**Note**

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows Server 2003 Certificate Services Web pages
Request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file
Use Windows 2000 Certificate Services Web Pages

# Check on a pending certificate request to a Windows Server 2003 CA

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check on a pending certificate request to a Windows Server 2003 CA**

1. Open Internet Explorer.
2. In Internet Explorer, open http://*servername*/certsrv, where *servername* is the name of the Web server running Windows Server 2003 where the certification authority you want to access is located.
3. Click **View the status of a pending certificate request**.
4. If there are no pending certificate requests, you will see a message to that effect. Otherwise, select the certificate request you want to check, and click **Next**.
5. Check the pending certificate requests:
   - **Still pending.** You must wait for the administrator of the certification authority to issue the certificate. To remove the certificate request, click **Remove**.
   - **Issued.** To install the certificate, click **Install this certificate**.
   - **Denied.** Contact the administrator of the certification authority for further information.
6. If you are finished using the Certificate Services Web pages, close Internet Explorer.

**Note**

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.

Related Topics

Submit a user certificate request via the Web to a Windows Server 2003 CA

Submit an advanced certificate request via the Web to a Windows Server 2003 CA

Request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file

Using Windows Server 2003 Certificate Services Web pages

Use Windows 2000 Certificate Services Web Pages

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Retrieve a certification authority certificate from a Windows Server 2003 CA

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To retrieve a certification authority certificate from a Windows Server 2003 CA**

1. Open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the Web server running Windows Server 2003 where the certification authority (CA) you want to access is located.
3. Click **Download a CA certificate, certificate chain, or CRL**.
4. Do one of the following:
   - If you want to trust all the certificates issued by this certification authority (CA), click **Install this CA certificate chain**.
   - If the certification authority has been renewed, you will have the choice of which version of the CA certificate you want to download.
5. Select the encoding method you want for the CRL: **DER** or **Base 64**.
6. Under **CA Certificate**, click the CA certificate you want to download, and then click **Download CA certificate** or click **Download CA certificate chain**.
7. In **File Download**, click **Open this file from its current location**, and then click **OK**.
8. When the **Certificate** dialog box appears, click **Install this certificate**.
9. In the Certificate Import wizard, click **Automatically select the certificate store based on the type of certificate**.
10. If you are finished using the Certificate Services Web pages, close Internet Explorer.

## Notes

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- Choosing **Install this CA certification path** is useful when you want to establish trust in a subordinate CA, but you do not currently have the certificate of the root CA of the certification hierarchy in your certificate store.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows Server 2003 Certificate Services Web pages
Certificates and certification authorities
Use Windows 2000 Certificate Services Web Pages

# Retrieve a certificate revocation list from a Windows Server 2003 CA

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To retrieve a certificate revocation list from a Windows Server 2003 CA**

1. Open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the Web server running Windows Server 2003 where the certification authority (CA) you want to access is located.
3. Click **Download a CA certificate, certificate chain, or CRL**.
4. Click the encoding method you want for the CRL, **DER** or **Base 64**.
5. Do one of the following:
   - Click **Download CA certificate**.
   - Click **Download CA certificate chain**.
   - Click **Download latest base CRL**.
   - Click **Download latest delta CRL**. (You need to have the latest base CRL already installed for the delta CRL to be functional.)
6. When the **File Download** dialog box appears, click **Save**. Select a folder on your computer in which to store the .crl file, and click **Save**.
7. Open Windows Explorer and locate the .crl file you just saved.
8. Right-click the .cer or .crl file and click **Install Certificate** or **Install CRL**, and then click **Next**.
9. When the Certificate Import Wizard opens, click **Automatically select the certificate store based on the type of certificate**.
10. After completing the Certificate Import Wizard, close Internet Explorer if you are finished using the Certificate Services Web pages.

### Note

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows Server 2003 Certificate Services Web pages

Certificates and certification authorities
Use Windows 2000 Certificate Services Web Pages

# Manage Certificate Trusts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage certificate trusts**

- Turn on automatic updating of trusted root authority certificates
- Turn off automatic updating of trusted root authority certificates
- Disable trust of third-party root CAs for a Windows Server 2003 domain
- Disable trust of user-selected root CAs for a Windows Server 2003 domain

# Turn on automatic updating of trusted root authority certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To turn on automatic updating of trusted root authority certificates**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. Select the **Update Root Certificates** check box.
4. Click **Next**, and then click **Finish**.

### Note

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Turn off automatic updating of trusted root authority certificates
Automatic updating of trusted root authorities

# Turn off automatic updating of trusted root authority certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To turn off automatic updating of trusted root authority certificates**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. Clear the **Update Root Certificates** check box.
4. Click **Next**, and then click **Finish**.

### Note

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Turn on automatic updating of trusted root authority certificates
Automatic updating of trusted root authorities

# Disable trust of third-party root CAs for a Windows Server 2003 domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable trust of third-party root CAs for a Windows Server 2003 domain**

1. Open the Group Policy object that you want to edit.

2. In the console tree, right-click **Trusted Root Certification Authorities**, and click **Properties**.

   **Where?**

   - *Policy Object Name*/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Trusted Root Certification Authorities

3. In **Client computers can trust the following certificate stores**, select the **Enterprise Root Certification Authorities** option.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a Group Policy object, see Related Topics.

- This procedure does not apply to Local Policy objects.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing trust of third-party certification authorities
Certificate stores
Managing trust of user-selected certification authorities
Group Policy

# Disable trust of user-selected root CAs for a Windows Server 2003 domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable trust of user-selected root CAs for a Windows Server 2003 domain**

1. On the domain controller, open Active Directory Users and Computers.
2. Click the domain name and, on the **Action** menu, click **Properties**.
3. Click the **Group Policy** tab and then click the **Edit** button.
4. In the console tree, click **Trusted Root Certification Authorities**.

   **Where?**

   - Default Domain Policy/Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Trusted Root Certification Authorities

5. On the **Action** menu, click **Properties**.
6. Clear the **Allow users to select new root certification authorities (CAs) to trust** check box and then click **OK**.
7. Close the Group Policy snap-in.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- When you disable the trust of user-selected CAs, these are the consequences:
  - When the users access any secure Web sites that are validated by these user-selected root CAs, they will receive security alert warnings informing them that the sites are not trusted.
  - In Internet Explorer, under **Certificates**, any previous user-selected root CAs will no longer be listed under Trusted Root Certification Authorities.
- If you don't want to disable trust of user-selected CA certificates for an entire domain, you can follow steps 4-7 for any Group Policy object that applies to all the computers for which you do want to have this trust disabled.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing trust of user-selected certification authorities
Certificate stores
Managing trust of third-party certification authorities

# Certificates Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about the Certificates snap-in.

- Certificates overview
- Understanding Certificates
- Using Certificates
- Certificates Resources

# Certificates overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificates overview**

**Certificates**
A public key certificate, usually just called a certificate, is a digitally-signed statement that binds the value of a public key to the identity of the person, device, or service that holds the corresponding private key. Most certificates in common use are based on the X.509v3 certificate standard. For references to extensive information about public key cryptography, see Certificates Resources.

Certificates can be issued for a variety of functions such as Web user authentication, Web server authentication, secure e-mail (using Secure/Multipurpose Internet Mail Extensions, also called S/MIME), Internet Protocol security (IPSec), Transport Layer Security (TLS), and code signing. Certificates are also issued from one certification authority (CA) to another in order to establish a certification hierarchy.

The entity that receives the certificate is the *subject* of the certificate. The issuer and signer of the certificate is a certification authority.

Typically, certificates contain the following information:

- The subject's public key value.
- The subject's identifier information, such as the name and e-mail address.
- The validity period (the length of time that the certificate is considered valid).
- Issuer identifier information.
- The digital signature of the issuer, which attests to the validity of the binding between the subject's public key and the subject's identifier information.

A certificate is valid only for the period of time specified within it; every certificate contains *Valid From* and *Valid To* dates, which set the boundaries of the validity period. Once a certificate's validity period has passed, a new certificate must be requested by the subject of the now-expired certificate.

In instances where it becomes necessary to undo the binding that is asserted in a certificate, a certificate can be revoked by the issuer. Each issuer maintains a certificate revocation list that can be used by programs when checking the validity of any given certificate.

One of the main benefits of certificates is that hosts no longer have to maintain a set of passwords for individual subjects who need to be authenticated as a prerequisite to access. Instead, the host merely establishes trust in a certificate issuer.

When a host, such as a secure Web server, designates an issuer as a trusted root authority, the host implicitly trusts the policies that the issuer has used to establish the bindings of certificates it issues. In effect, the host trusts that the issuer has verified the identity of the certificate subject. A host designates an issuer as a trusted root authority by placing the issuer's self-signed certificate, which contains the issuer's public key, into the trusted root certification authority certificate store of the host computer. Intermediate or subordinate certification authorities are trusted only if they have a valid certification path from a trusted root certification authority.

For more information about certificates, see Understanding Certificates.

**Certificate uses**
Because certificates are generally used to establish identity and create trusts for the secure exchange of information, certification authorities (CAs) can issue certificates to people, to devices (such as computers), and to services running on computers (such as IPSec).

In some cases, computers must be able to exchange information with a high degree of confidence in the identity of the other device, service, or person involved in the transaction. In some cases, people need to exchange information with a high degree

of confidence in the identity of the other person, computer, or service involved in the transaction. Applications and services that run on computers also frequently need to verify that they are accessing information from a trusted source.

In circumstances where two entities--such as devices, persons or applications or services--attempt to establish identity and trust, the fact that both entities trust the same certification authority allows the bond of identity and trust to be established between them. Once a certificate subject has presented a certificate issued by a trusted CA, the entity attempting to establish trust can proceed with an information exchange by storing the certificate subject's certificate in its own certificate store, and, where applicable, using the public key contained in the certificate to encrypt a session key so that all subsequent communications to the certificate subject are secure.

For example, when using the Internet for online banking, it is important to know that your Web browser is communicating directly and securely with your bank's Web server. Your Web browser must be able to achieve Web server authentication before a safe transaction can occur. That is, the Web server must be able to prove its identity to your Web browser before the transaction can proceed. Microsoft® Internet Explorer uses Secure Sockets Layer (SSL) to encrypt messages and transmit them securely across the Internet, as do most other modern Web browsers and Web servers.

When you connect using an SSL-enabled browser to an online banking Web server that has a server certificate from a certification authority such as Verisign, the following events occur:

- You access your bank's secured online banking login Web page using your Web browser. If you use Internet Explorer, a locked-padlock icon will appear in the lower right corner of the browser status bar to indicate that the browser is connected to a secure Web site. Other browsers depict secure connections in other ways.

- The bank's Web server automatically sends a server certificate to your Web browser.

- To authenticate the Web server, your Web browser checks the certificate store on your computer. If the certification authority that issued the certificate to your bank is trusted, the transaction can proceed, and the bank certificate is stored in your certificate store.

- To encrypt all communications with the bank Web server, your Web browser creates a unique session key. Your Web browser encrypts the session key with the bank Web server certificate so that only the bank Web server can read messages sent by your browser. (Some of these messages will contain your login name and password and other sensitive information, so this level of security is necessary.)

- The secure session is established, and sensitive information can be sent between your Web browser and the bank's Web server in a secure manner.

For more information, see Security with Certificates.

Certificates can also be used to verify the authenticity of software code that you download from the Internet, install from your company intranet, or purchase on CD-ROM and install on your computer. Unsigned software--software that does not have a valid software publisher's certificate--can pose a risk to your computer and the information you store on your computer.

When software is signed with a valid certificate from a trusted certification authority, you know that the software code has not been tampered with and can be safely installed on your computer. During software installation, you are prompted to verify that you trust the software manufacturer (for example, Microsoft Corporation). You may also be offered the option to always trust software content from that particular software manufacturer. If you choose to trust content from the manufacturer, their certificate goes into your certificate store and other software installations of their products can occur with a circumstance of predefined trust. In the circumstance of predefined trust, you can install software from the manufacturer without being prompted to indicate whether they are trusted; the certificate on your computer states that you trust the manufacturer of the software.

As with other certificates, those certificates used to verify the authenticity of software and the identity of a software publisher can have other purposes. For example, when the Certificates snap-in is set to view certificates by purpose, the Code Signing folder might contain a certificate issued to Microsoft Windows Hardware Compatibility by the Microsoft Root Authority. This single certificate has three purposes:

- It ensures that the software came from the software publisher.

- It protects software from alteration after publication.

- It provides Windows hardware driver verification.

**Certificate use in organizations**
Many organizations install their own certification authorities and issue certificates to internal devices, services, and employees to create a more secure computing environment. Large organizations may have multiple certification authorities, set up in a hierarchy that leads to a root certification authority. Thus, an employee of an organization may have a multitude of certificates in their certificate store that have been issued by a variety of internal certification authorities, all of whom share a trust connection via the certification path to the root certification authority.

When an employee logs in to the organization's network from home using a virtual private network (VPN), the VPN server can present a server certificate to establish its identity. Because the corporate root authority is trusted and the corporate root certification authority issued the certificate of the VPN server, the client computer can proceed with the connection and the employee knows his computer is actually connected to his organization's VPN server.

The VPN server must also be able to authenticate the VPN client before data can be exchanged over the VPN connection. Either computer-level authentication occurs with the exchange of computer certificates or user-level authentication occurs though the use of a Point-to-Point Protocol (PPP) authentication method. For L2TP (Layer 2 Tunneling Protocol)/IPSec connections, computer certificates are required for both the client and server.

For more information, see Authentication of VPN clients, Computer certificates for L2TP/IPSec VPN connections, and Remote access VPN design considerations.

The client computer certificate may serve multiple purposes, most of which are based in authentication, allowing the client to use many organizational resources without the need for individual certificates for each resource. For example, the client certificate might allow VPN connectivity as well as access to the company store intranet site, to product servers, and to the human resources database where employee data is stored.

The VPN server certificate might also serve multiple purposes. The same certificate might have the purpose of verifying the identity of e-mail servers, Web servers, or application servers. The certification authority that issues the certificate determines the number of purposes for each certificate.

**Certificates issued to persons**
You can purchase a certificate from a commercial certification authority, such as Verisign, to send personal e-mail messages that are encrypted for security or digitally signed to prove authenticity.

Once you have purchased a certificate and you use it to digitally sign an e-mail message, the message recipient can verify that the message has not been altered during transit and that the message came from you--assuming, of course, that the message recipient trusts the certification authority that issued your certificate.

When you encrypt an e-mail message, nobody can read the message while it is in transit, and only the message recipient can decrypt and read the message.

**Certificates and applications**
Most e-mail clients allow you to automatically sign or encrypt e-mail messages or to encrypt or sign messages on an individual basis. Microsoft applications that allow e-mail messages to be digitally signed or encrypted are Microsoft Outlook® 2000, Microsoft Outlook Express and Microsoft Outlook 98.

Many Windows applications use certificates. Here are some links and information about how certificates can be used with Microsoft Internet Information Services (IIS), Microsoft Outlook 2003, Microsoft Outlook Express, and Microsoft Internet Explorer.

### Internet Information Services certificates information

IIS certificate storage is integrated with the CryptoAPI storage. The Certificates snap-in provides a single point of entry that lets administrators store, back up, and configure server certificates.

IIS comes with three security task wizards that simplify most of the security tasks necessary to maintain a secure Web site. You can use the Web Server Certificate Wizard to manage Secure Sockets Layer (SSL) features in IIS and server certificates. Certificates are used in negotiating a secure link between your server and a user's browser. You can use the CTL Wizard to manage certificate trust lists (CTLs). Certificate trust lists are lists of trusted certification authorities for each Web site or virtual directory. You can use the IIS Permissions Wizard to assign Web and NTFS access permissions to Web sites, virtual directories, and files on your server.

For more information about the use of certificates in IIS, see the Microsoft Web site.

**Outlook 2003 certificates information**

- "Set Up Security for Internet E-mail Messages" on the Microsoft Web site.

**Internet Explorer certificates information:**

- Internet Explorer 6.0 Resource Kit on the Microsoft Web site.

Other Microsoft resources on certificates include Knowledge Base articles:

- "Description of Symmetric and Asymmetric Encryption" on the Microsoft Web site.
- "Using a Certificate Authority for the Encrypting File Service" on the Microsoft Web site.
- "Best Practices for Encrypting File System" on the Microsoft Web site.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

**Certificate import and export**
If you have certificates you want to use on other computers, you can export your certificates for import to other computers. For more information, see Importing and exporting certificates and Import and Export Certificates.

**The certificate store**
There are four basic sources for the certificates found in the certificate stores on your computer:

- The certificate is included with your installation of Windows Server 2003 and comes on the Windows Server 2003 CD.
- You use an application such as an Internet browser to engage in a SSL session, during which certificates are stored on your computer after establishment of trust.
- You explicitly choose to accept a certificate, as when you install software or receive an encrypted or digitally signed e-mail from others.
- You request a certificate from a certification authority, such as a certificate needed to access specific organizational resources.

Over time, as you use the Internet and connect to servers that use certificates for authentication and other purposes, the certificate store on your computer will gain entries. Some certificates may have one purpose, such as server authentication or client authentication, while other certificates may serve multiple purposes. The certification authority that issues the certificate determines the purpose or purposes the certificate may serve.

Although Internet Explorer and Windows store certificates in the same certificate stores, they allow different views of the certificate stores. For more information, see Set Display Options of the Certificates Snap-in in Certificates Help. To better understand how the Certificates snap-in displays the certificate store, see Certificate stores.

By using the Certificates snap-in, you can view certificates by logical store or by purpose. If you view certificates by purpose, a certificate with multiple purposes will appear listed in every folder that defines a purpose for which the certificate can be used.

**The Certificates snap-in**
You can use the Certificates snap-in to manage certificates for users, computers, or services.

Users and administrators can use the Certificates snap-in to request new certificates from Windows 2000 and enterprise certification authorities. In addition, users can find, view, import, and export certificates from within certificate stores. However, in most cases, users do not have to personally manage their certificates and their certificate stores. That can be accomplished by administrators, by policy settings, and through programs that use certificates.

Administrators are the primary users of the Certificates snap-in and, as such, they are able to perform a wide variety of

certificate management tasks in their personal certificate store as well as the certificate stores for any computer or service that they have the right to administer.

For more information, see Requesting certificates

# Understanding Certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding certificates**

This section covers the following topics:

- Certificates and certification authorities
- Generating encryption keys and certificate requests
- Certificate stores
- Security with Certificates

# Certificates and certification authorities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificates and certification authorities**

When a certificate is presented to an entity as a means of identifying the certificate holder (the *subject* of the certificate), it is useful only if the entity receiving the certificate trusts the issuer, which is often referred to as the certification authority (CA).

When you *trust* a certification authority, that means you have confidence that the certification authority has the proper policies in place when evaluating certificate requests and will deny certificates to any entity that does not meet those policies. In addition, you trust that the certification authority will revoke certificates that should no longer be considered valid by publishing an up-to-date certificate revocation list. Certificate revocation lists are considered valid until they expire. So even if the CA publishes a new certificate revocation list with newly revoked certificates listed, all clients that have an old certificate revocation list will not look for, nor retrieve the new one until the old one expires or is deleted. Clients can use the CA Web pages to manually retrieve the most current certificate revocation list if necessary.

For Windows Server 2003 users, computers, and services, trust in a certification authority is established when you have a copy of the root certificate in the trusted root certification authorities store, as well as having a valid certification path, meaning that none of the certificates in the certification path has been revoked or has had its validity period expire. The certification path includes every certificate issued to each CA in the certification hierarchy from a subordinate CA to the root CA. For example, for a root CA, the certification path is one certificate, its own self-signed certificate. For a subordinate CA, just below the root CA in the hierarchy, its certification path is two certificates: its own certificate and the root CA certificate.




If your organization is using Active Directory, then trust in your organization's certification authorities will typically be established automatically, based on decisions and settings made by the system administrator.

A related concept with which you should be familiar is certificate store inheritance. If you place a root CA certificate into the computer's trusted root certification authorities store or enterprise trust store, then any user of the computer will see that certificate in their own user trusted root certification authorities store or enterprise trust store even though the root certificate is actually in the computer's store. Essentially, users will trust any CA that their computer trusts. Certificate store inheritance does not work the other way around: certificates in the user's trusted root certification authorities store and enterprise trust store are not inherited by the computer.

If your organization is using the version of Certificate Services installed with the Windows Server 2003 family to run its certification authority, the certification authority is one of two types: enterprise or stand-alone. The differences between the two standard types of Windows Server 2003 certification authorities for certificate users and requesters are summarized below.

**Enterprise certification authority**

An enterprise certification authority depends upon Active Directory being present.

You can use the Certificate Request Wizard (which is started from within the Certificates snap-in), as well as certification authority Web pages, to request certificates from an enterprise certification authority.

An enterprise certification authority offers different types of certificates to a requester based on the certificates it is configured to issue as well as the security permissions of the requester. An enterprise certification authority uses information available in

Active Directory to help verify the requester's identity. An enterprise certification authority publishes its certificate revocation list to Active Directory as well as to a shared directory.

**Stand-alone certification authority**
A stand-alone certification authority is less automated for a user than an enterprise certification authority because it does not depend on the use of Active Directory.

By default, users can request certificates from a stand-alone certification authority only by using Web pages.

Stand-alone certification authorities that do not use Active Directory will generally have to request that the certificate requester provide more complete identifying information. A stand-alone certification authority makes its certificate revocation list available from a shared folder, or from Active Directory, if it is available.

### Notes

- For more information about the Certificate Services policy and exit modules, which determine whether a Windows Server 2003 certification authority is enterprise or stand-alone, see Policy and exit modules.

- It is possible to customize a Windows Server 2003 certification authority with a policy module that is different from the one installed with the Windows Server 2003 operating system. In that case, the certification authority is technically neither a stand-alone nor an enterprise certification authority, but may have qualities of one or the other.

For more information, see Certification Authorities, Certificate Services overview, Requesting certificates, Using Windows 2000 Certificate Services Web pages, Certificate stores, and Certificates Resources

# Generating encryption keys and certificate requests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Generating encryption keys and certificate requests**

How you request and receive a certificate depends upon the policies and processes of the certification authority (CA) that issues the certificate. For example, some certification authorities have Web pages to which you can submit a reques or, if you are in an organization that has deployed Active Directory and Certificate Services, you can use the Certificates snap-in to request a certificate, if your computer is a member of a domain and you are authorized to request certificates.

In any case, when you generate a request for a new certificate, the information in that request is first passed from the requesting program to CryptoAPI. CryptoAPI will pass the proper data to a program known as a cryptographic service provider (CSP) that is installed on your computer or on a device that is accessible to your computer. If the CSP is software-based, it will generate a public key and a private key, often referred to as a key pair, on your computer. If the CSP is hardware-based, such as a smart card CSP, it will instruct a piece of hardware to generate the key pair.

After the keys are generated, a software CSP encrypts and then secures the private key. A smart card CSP stores the private key on a smart card and the smart card controls access to the key. The public key is sent to the certification authority, along with the certificate requester information. Once the CA verifies the certificate request according to its policies, it will use its own private key to create a digital signature in the certificate and then issue it to the requester. The certificate requester will then be presented with the certificate from the CA and the option to install it in the appropriate certificate store on the computer or hardware device.

For more information about keys and certificates, see Certificates Resources.

For more information, see Requesting certificates.

# Certificate stores

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate stores**

Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, store a certificate locally on the computer or device that requested it or, in the case of a user, on the computer or device that the user used to request it. The storage location is called the certificate store. A certificate store will often have numerous certificates, possibly issued from a number of a different certification authorities.

Using the Certificates snap-in, you can display the certificate store for a user, a computer, or a service according to the purpose for which the certificates were issued or by using their logical storage categories. When you display certificates according to their storage categories, you can also choose to display the physical stores, showing the certificate storage hierarchy. (This is recommended for advanced users only.)

If you have the user rights to do so, you can import or export certificates from any of the folders in the certificate store. Additionally, if the private key associated with a certificate is marked as available for export, you can export both into a PKCS #12 file.

Windows can also publish certificates to Active Directory. Publishing a certificate in Active Directory enables all users or computers with adequate permissions to retrieve the certificate as needed.

Certificates can be displayed by purpose or by logical stores, as shown in the following table. Displaying certificates by logical stores is the Certificates default. (Note that the list of certificate purpose stores does not include all the possible purpose stores.)

| Display by | Folder name | Contents |
|---|---|---|
| Logical Store | Personal | Certificates associated with private keys to which you have access. These are the certificates that have been issued to you, or to the computer or service for which you are managing certificates.<br><br>Note to administrators: Computers in a Windows Server 2003 Active Directory domain can have certificates automatically placed in this store through the use of Group Policy-based autoenrollment. For more information, see Automatic certificate request settings. |
| | Trusted Root Certification Authorities | Implicitly trusted certification authorities. Includes all of the certificates in the Third-Party Root Certification Authorities store plus root certificates from your organization and Microsoft.<br><br>If you are an administrator and want to add third-party certification authority certificates to this store for all computers in a Windows Server 2003 Active Directory domain, you can use Group Policy to distribute trusted root certificates to your organization. For more information, see Trusted root certification authority policy. |
| | Enterprise Trust | A container for certificate trust lists. A certificate trust list provides a mechanism for trusting self-signed root certificates from other organizations and limiting the purposes for which these certificates are trusted. For more information about Enterprise trust see Enterprise trust policy. |
| | Intermediate Certification Authorities | Certificates issued to subordinate certification authorities. |
| | Trusted People | Certificates issued to people or end entities that are explicitly trusted. Most often these are self-signed certificates or certificates explicitly trusted in an application such as Microsoft Outlook. |
| | Other People | Certificates issued to people or end entities that are implicitly trusted. These certificates must be part of a trusted certification hierarchy. Most often these are cached certificates for services like Encrypting File System, where certificates are used for creating authorization for decrypting an encrypted file. |
| | Trusted Publishers | Certificates from certification authorities that are trusted by Software Restriction policies. |

| | | |
|---|---|---|
| | Disallowed Certificates | These are certificates that you have explicitly decided not to trust using either Software Restriction policy or by clicking "Do not trust this certificate" when the decision is presented to you in mail or a Web browser. |
| | Third-Party Root Certification Authorities | Trusted root certificates from certification authorities other than Microsoft and your organization. |
| | Certificate Enrollment Requests | Pending or rejected certificate requests. |
| | Active Directory User Object | Certificates associated with your user object and published in Active Directory. |
| Purpose | Server Authentication | Certificates that server programs use to authenticate themselves to clients. |
| | Client Authentication | Certificates that client programs use to authenticate themselves to servers. |
| | Code Signing | Certificates associated with key pairs used to sign active content. |
| | Secure Email | Certificates associated with key pairs used to sign e-mail messages. |
| | Encrypting File System | Certificates associated with key pairs that encrypt and decrypt the symmetric key used for encrypting and decrypting data by Encrypting File System (EFS). |
| | File Recovery | Certificates associated with key pairs that encrypt and decrypt the symmetric key used for recovering encrypted data by Encrypting File System (EFS). |

When you look at the contents of a certificate store in Logical Store mode, you will occasionally see what appears to be two copies of the same certificate in the store. This occurs because the same certificate is stored in separate physical stores under a logical store. When the contents of the physical certificates stores are combined into one logical store view, both instances of the same certificate are displayed.

You can verify this by setting the view option to show the physical certificate stores and then noting that the certificate is stored in separate physical stores under the same logical store. You can verify that it is the same certificate by comparing the serial numbers.

For more information, see Generating encryption keys and certificate requests, Importing and exporting certificates, Display certificate stores in Purpose mode, Display certificate stores in Logical Store mode, Display archived certificates, Display certificate stores storage structure

# Security with Certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security with certificates**

Certificates can be used for:

- Authentication, which verifies the identity of someone or something.
- Privacy, which ensures that information is only available to the intended audience.
- Encryption, which disguises information so that unauthorized readers are unable to decipher it.
- Digital signatures, which provide nonrepudiation and message integrity.

These security services can be important to the security of your communications.

- Certificates and Authentication
- Privacy
- Encryption
- Digital signatures

# Certificates and Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication**

Authentication is crucial to secure communication. Users must be able to prove their identity to those with whom they communicate and must be able to verify the identity of others. Authentication of identity on a network is complex because the communicating parties do not physically meet as they communicate. This can allow an unethical person to intercept messages or to impersonate another person or entity.

The digital certificate is a common credential that provides a means to verify identity. Certificates use cryptographic techniques to address the problem of the lack of physical contact between those communicating. Using these techniques limits the possibility of an unethical person intercepting, altering, or counterfeiting messages. These cryptographic techniques make certificates difficult to modify. Thus, it is difficult for an entity to impersonate someone else.

The data in a certificate includes the public cryptographic key from the certificate subject's public and private key pair. A message signed with its sender's private key can be verified by the message's recipient as authentic by using the sender's public key. This key can be found on a copy of the sender's certificate. Verifying a signature with a public key from a certificate proves that the signature was produced using the certificate subject's private key. If the sender has been vigilant and has kept the private key secret, the receiver can be confident in the identity of the message sender.

A few of the ways certificates are used to provide authentication are:

- Authentication of a user to a secure Web site via the Transport Layer Security (TLS) or the Secure Sockets Layer (SSL) protocol.
- Authentication of a server to a user via TLS.
- Logging on to a Windows Server 2003 domain.

For more information about authentication and certificates, see Certificates Resources.

# Privacy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Privacy**

Communications on a network, such as the Internet, are subject to possible monitoring by unknown and, perhaps, malicious users. Public networks are treacherous for unencrypted sensitive information because anyone can access the network and analyze the data being transmitted between two points. Even private local area networks (LANs) are vulnerable to determined efforts by intruders to acquire physical access to the network. Consequently, if sensitive information is transmitted between computing devices on any type of network, users will almost certainly want to use some sort of encryption to keep their data private.

Public key encryption is not used to encrypt large amounts of data. Instead, data is typically protected with secret key encryption and, in turn, that secret key is encrypted with the public key of the recipient of the data. The encrypted secret key will then be transmitted to the recipient along with encrypted data itself. The recipient will use the private key to decrypt the secret key. The secret key will then be used to decrypt the message itself.

Certificates enable privacy of transmitted data using a number of different methods. Some of the commonly used privacy-enabling protocols that use certificates are:

- Secure Multipurpose Internet Mail Extensions (S/MIME)
- Transport Layer Security (TLS)
- Internet Protocol security (IPSec). For more information, see Introducing IPSec.

For more information about privacy and certificates, see Certificates Resources.

# Encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encryption**

Encryption can be thought of as locking something valuable into a strong box with a key. Conversely, decryption can be compared to opening the box and retrieving the valuable item. On computers, sensitive data in the form of e-mail messages, files on a disk, and files being transmitted across the network can be encrypted using a key. Encrypted data and the key used to encrypt data are both unintelligible.

Typically, public key encryption is not used to encrypt large amounts of data. However, public key cryptography does offer an efficient method to send someone the secret key that is used when a symmetric encryption operation is performed on a large amount of data.

As an example, suppose Bob wants to send Alice a large number of encrypted files. For performance reasons, he will use a symmetric encryption key algorithm, such as Data Encryption Standard (DES), to encrypt the data. To send the encrypted data and the DES secret key needed to decrypt the data securely, Bob will encrypt the secret key with Alice's public key obtained from her certificate. Because her public key was used to encrypt the secret key, Alice, using her private key, will be the only one able to decrypt the DES secret key and thus decrypt the DES-encrypted data.

For more information, see Encrypting File System overview. For more information about encryption and certificates, see Certificates Resources.

# Digital signatures

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Digital signatures**

A digital signature is a way to ensure the integrity and origin of data. A digital signature provides strong evidence that the data has not been altered since it was signed and it confirms the identity of the person or entity who signed the data. This enables the important security features of *integrity* and *nonrepudiation*, which are essential for secure electronic commerce transactions.

Digital signatures are typically used when data is distributed in plaintext, or unencrypted form. In these cases, while the sensitivity of the message itself may not warrant encryption, there could be a compelling reason to ensure that the data is in its original form and has not been sent by an impostor because, in a distributed computing environment, plaintext can conceivably be read or altered by anyone on the network with the proper access, whether authorized or not.

For more information about digital signatures and certificates, see Certificates Resources.

# Using Certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using certificates**

This section covers the following topics:

- Starting the Certificates snap-in
- Requesting certificates
- Renewing certificates
- Certificate autoenrollment
- Importing and exporting certificates
- Using credential roaming
- Viewing certificate information
- Using Windows 2000 Certificate Services Web pages
- Using Windows Server 2003 Certificate Services Web pages
- Managing Certificate Trusts

# Starting the Certificates snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Starting the Certificates snap-in**

You use the Microsoft Management Console (MMC) Certificates snap-in to manage public key certificates.

If you are the administrator of the computer, when you add Certificates to your MMC console, you have the option of managing certificates that are issued to one of the following:

- Yourself (**My user account** option)
- Your computer (**Computer account** option)
- Local services (**Service account** option)

If you are a user, when you add Certificates to your MMC console, you can manage certificates only for your user account.

A separate instance of Certificates must be loaded for each entity for which you are managing certificates.

For more information, see Manage certificates for your user account, Manage certificates for a computer, Manage certificates for a service, and Certificates overview.

# Requesting certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Requesting certificates**

Certificate requests must be made by the user, computer, or service that has access to the private key associated with the public key that will be part of the certificate. Depending upon the public key policies established by your system administrator, machines and services can automatically request certificates without user intervention. In addition, administrators can request smart card user certificates and smart card certificates for logging on to the system on behalf of other users by using their enrollment agent certificate.

There are two primary ways to explicitly request certificates in a Windows Server 2003 operating system.

**Request certificates using the Certificate Request Wizard**

When you request certificates from a Windows Server 2003 enterprise certification authority, you can use the Certificate Request Wizard located in the Certificates snap-in. This wizard guides you through the following steps:

- Selecting the certification authority to which you will submit the request.

  Only enterprise certification authorities that are available in your Windows domain will be able to issue certificates using the Certificate Request Wizard.

- Selecting the appropriate certificate template to use for the new certificate.

  Certificate templates are predefined configurations that provide common settings for the certificate request. Certificate templates describe the purpose for which the requested certificate is to be used. The list of certificate templates that is available to you is determined by the certificate types which the certification authority is configured to issue and whether you have been granted the access rights to the certificate template by the system administrator.

- (Optional) Using **Advanced Options** in the Certificate Request Wizard to select the cryptographic service provider (CSP) for the key pair associated with the certificate request.

Only Basic EFS (Encrypting File System) and EFS Recovery Agent certificates have their associated private keys marked as available for export when you use the Certificate Request Wizard. If you want to request another type of certificate and have its private key available for export to a PKCS #12 file, you will need to use the **Advanced request** option on the Windows Server 2003 Certificate Services Web pages.

For instructions on opening and using the Certificate Request Wizard, see Request a certificate.

You can also use the Certificate Request Wizard to request a new certificate from an enterprise certification authority by using an existing key pair that is already associated with another certificate. For more information, see Request a certificate with the same key.

**Request certificates using the Windows Server 2003 Certificate Services Web pages**

Each certification authority that is installed on a computer running Windows Server 2003 has Web pages that users can access to submit basic and advanced certificate requests. By default, these pages are located at http://*servername*/certsrv, where *servername* is the name of the computer running Windows Server 2003.

When you request certificates from a Windows Server 2003 stand-alone certification authority, you use the Certificate Services Web pages. Web pages can also be used to request certificates from Windows Server 2003 enterprise certification authorities if you want to set optional request features that are not available in the Certificate Request Wizard, such as marking the keys as exportable, setting key length, choosing the hash algorithm, or saving the request to a PKCS #10 file.

For more information on using Certificate Services Web pages, see Using Windows Server 2003 Certificate Services Web pages.

**Processing certificate requests**

When you submit a certificate request to a Windows Server 2003 enterprise certification authority, it is immediately processed, as opposed to being set to "pending." The certificate request will either immediately fail or be granted. If it is granted, the certificate is issued, and you will be prompted to install it.

When you submit a certificate request to a Windows Server 2003 stand-alone certification authority, it will either be immediately processed or, by default, it will be considered pending until the administrator of the certification authority approves or rejects the request. In the case of a pending request, the certificate requester will have to use the Certificate Services Web pages to check the status of pending certificates. For more information, see Check on a pending certificate request to a Windows Server 2003 CA.

For more information about certificates and certification authorities, see Certificates and certification authorities.

# Renewing certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Renewing certificates**

Every certificate has a validity period. After the end of the validity period, the certificate is no longer considered an acceptable or usable credential. The Certificates snap-in enables you to renew a certificate issued from a Windows Server 2003 enterprise certification authority before or after the end of its validity period by using the Certificate Renewal Wizard.

You can either renew the certificate with the same key set you used before, or you can renew a certificate with a new key set. For information comparing reusing a key set and generating a new key set for certificate renewal, see Certificates Resources.

Before you renew a certificate, you will need to know:

- The issuing certification authority.
- (Optional) If you want a new public key and private key pair for the certificate, the cryptographic service provider (CSP) that should be used to generate the key pair.

For more information, see Renew a certificate with a new key and Renew a certificate with the same key.

In addition, you can renew certificates issued from both Windows Server 2003 enterprise certification authorities and Windows Server 2003 stand-alone certification authorities with the Certificate Services Web pages by pasting in the contents of a PKCS #7 file. For more information, see Request a certificate using a PKCS #10 or PKCS #7 file.

# Certificate autoenrollment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate autoenrollment**

Administrators of a network frequently require certificates users and computers to hold one or more certificates. The manual certificate enrollment process for a user can be somewhat confusing and time-consuming. The administrator may not have control of the enrollment process, so many users or computers could get useless or incorrect certificates. A process for obtaining, storing and updating the certificates for subjects without intervention is ideal. That is the purpose of certificate autoenrollment.

Certificate autoenrollment is a process that allows clients to automatically submit certificate requests to a certification authority and retrieve and store issued certificates. It is controlled by the administrator so the certificates are obtained by the appropriate subject with the correct settings. Periodically the client computer will check for any autoenrollment tasks that may be required and perform them. This includes requesting new certificates, retrieving fulfilled requests from a certification authority, or renewing an existing certificate.

Although the process is automatic, starting the process manually may be desired. The Certificates snap-in has a menu option to trigger certificate autoenrollment for the current user or local computer immediately.

For more information, see Trigger automatic enrollment for certificates and Certificate Services example implementation: Establishing autoenrollment for user certificates

# Importing and exporting certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Importing and exporting certificates**
The Certificates snap-in provides administrative tools to export and import certificates, including their certification paths and private keys, if needed. You can export certificates to or import certificates from PKCS #12 files, PKCS #7 files, and binary-encoded X.509 certificate files.

**Importing a certificate**
You might want to import a certificate:

- To install a certificate that was sent to you in a file by another user, computer, or certification authority.
- To restore a damaged or lost certificate that you previously backed up.
- To install a certificate and its associated private key from a computer that the certificate holder was previously using.

When you import a certificate, you copy the certificate from a file that uses a standard certificate storage format to a certificate store for your user account or your computer account.

**Exporting a certificate**
You might want to export a certificate:

- To back up a certificate.
- To back up a certificate and its associated private key.
- To copy a certificate for use on another computer.
- To remove a certificate and its private key from the certificate holder's current computer for installation on another computer.

When you export a certificate, you are copying the certificate from its certificate store to a file that uses a standard certificate storage format.

**Standard certificate file formats**
You can import and export certificates in the following formats:

- **Personal Information Exchange (PKCS #12)**

  The Personal Information Exchange format (PFX, also called PKCS #12) enables the transfer of certificates and their corresponding private keys from one computer to another or from a computer to removable media.

  PKCS #12 (Public Key Cryptography Standard #12) is an industry format that is suitable for transport or backup and restoration of a certificate and its associated private key. This can be between products from the same vendor or different vendors.

  To use the PKCS #12 format, the cryptographic service provider (CSP) must recognize the certificate and keys as exportable. If a certificate was issued from a Windows Server 2003 or Windows 2000 certification authority, the private key for that certificate is only exportable if one of the following is true:

  - The certificate is for EFS (Encrypting File System) or EFS recovery.
  - The certificate was requested through the Advanced Certificate Request certification authority Web page with the **Mark keys as exportable** check box selected.

Because exporting a private key might expose it to unintended parties, the PKCS #12 format is the only format supported in the Windows Server 2003 family for exporting a certificate and its associated private key.

- **Cryptographic Message Syntax Standard (PKCS #7)**

  The PKCS #7 format enables the transfer of a certificate and all the certificates in its certification path from one computer to another, or from a computer to removable media. PKCS #7 files typically use the .p7b extension, and are compatible with the ITU-T X.509 standard. PKCS #7 allows for attributes, such as countersignatures, to be associated with signatures, and attributes such as signing time can be authenticated along with message content. For more information on PKCS #7, see the PKCS #7 page at the RSA Labs Web site.

- **DER Encoded Binary X.509**

  DER (Distinguished Encoding Rules) for ASN.1, as defined in ITU-T Recommendation X.509, is a more restrictive encoding standard than the alternative BER (Basic Encoding Rules) for ASN.1, as defined in ITU-T Recommendation X.209, upon which DER is based. Both BER and DER provide a platform-independent method of encoding objects such as certificates and messages for transmission between devices and applications.

  During certificate encoding, most applications use DER because a portion of the certificate (the CertificationRequest's CertificationRequestInfo) must be DER-encoded to be signed.

  This format might be used by certification authorities that are not on computers running Windows Server 2003 , so it is supported for interoperability. DER certificate files use the .cer extension.

  For more information, see the document "ITU-T Recommendation X.509, Information Technology--Open Systems Interconnection--The Directory: Authentication Framework," at the International Telecommunication Union (ITU) Web site.

- **Base64 Encoded X.509**

  This is an encoding method developed for use with Secure/Multipurpose Internet Mail Extensions (S/MIME), which is a popular, standard method for transferring binary attachments over the Internet. Base64 encodes files into ASCII text format, making corruption less likely as the files are sent through Internet gateways, while S/MIME provides some cryptographic security services for electronic messaging applications, including non-repudiation of origin using digital signatures, privacy and data security using encryption, authentication, and message integrity.

  The MIME (Multipurpose Internet Mail Extensions) specification (RFC 1341 and successors) defines a mechanism for encoding arbitrary binary information for transmission by electronic mail.

  Because all MIME-compliant clients can decode Base64 files, this format might be used by certification authorities that are not on computers running Windows Server 2003 , so it is supported for interoperability. Base64 certificate files use the .cer extension.

  For more information, see the document "RFC 2633 S/MIME Version 3 Message Specification, 1999," at the Internet Engineering Task Force (IETF) Web site.

**Choosing an export format**

If you are exporting certificates to be imported onto a computer running Windows, PKCS #7 format is the preferred export format, primarily because this format preserves the chain of certification authorities, or the certification path, of any certificate that includes countersignatures associated with signatures.

If you are exporting certificates for import onto a computer running another operating system, it is possible that the PKCS #7 format is supported. If it is not supported, the DER Encoded Binary format or the Base64 Encoded format are provided for interoperability.

For more information, see Import a certificate, Export a certificate, Export a certificate with the private key, and View the certificates in a PKCS #7 file.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Using credential roaming

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

A growing number of organizations and users rely on certificates to secure critical transactions such as signing, encrypting, and decrypting e-mail and authenticating identity on wireless networks.

But many organizations require users to log onto and use more than one computer. Thus, a typical user of multiple computers who has been issued certificates for signing e-mail, encrypting and decrypting documents, and authentication might have separate sets of certificates and private keys on each computer.

Credential roaming enables using a single set of certificates and private keys on multiple computers in a manner that is extremely secure, easy to implement and manage, and transparent to the user.

Credential roaming only supports x.509 v3 certificates and RSA or DSA key pairs that are stored in the user's credential store. Only credentials in the user's MY and REQUEST stores are roamed. Credentials that use non-Microsoft cryptographic service providers (CSPs) or are located in other certificate stores are not supported.

**Note**

- Version 1 x.509 certificates are not stored in the user's credential store and cannot be used with credential roaming. The Microsoft Exchange Cryptographic Provider v1.0 does not use the default key container and cannot be used for credential roaming. In addition, keys stored on smart cards and within hardware security modules (HSMs) cannot be used with credential roaming.

Credential roaming is supported for clients running Windows Server 2003 with Service Pack 1 (SP1). Credential roaming is supported on domain controllers running Windows 2000 with SP3 or later or Windows Server 2003. Mixed environments containing domain controllers running Windows 2000 with SP3 or later and Windows Server 2003 are also supported. The forest functional level can be either Windows 2000 or Windows Server 2003.

- An update was recently made available that enables credential roaming on clients running Windows XP Professional (SP2) and an update is available for Windows Server 2003 SP1.

**Note**

- For the best performance and security, it is recommended that the servers be upgraded to Windows Server 2003 SP1. For other best practices, see Credential roaming best practices.

**See Also**
**Concepts**
Understanding credential roaming
Credential roaming schema changes
Credential roaming administrative template
Credential roaming best practices

# Understanding credential roaming

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Credential roaming, enabled allows organizations to store certificates and private keys in Active Directory apart from application state or configuration information.

**How credential roaming works**
Credential roaming uses existing logon, Group Policy, and auto-enrollment mechanisms. These mechanisms make it possible to safely and securely download certificates and keys to a local computer whenever a user logs on and, if desired, remove them when the user logs off. In addition, the integrity of these credentials is maintained under any conditions, such as when certificates are updated and when users log onto more than one computer at a time.

The following steps describe how digital credential roaming works.

1. A user logs onto a client computer that is connected to an Active Directory domain.

2. As part of the logon process, credential roaming Group Policy is applied to the user's computer.

3. If this is the first time that credential roaming is being used, the certificates in the user's store on the client computer are copied to Active Directory.

4. If the user already has certificates in Active Directory, the certificates in the user's certificate store on the client computer are compared to the certificates stored for the user in Active Directory.

5. If the certificates in the user's certificate store are current, then no further action is taken. However, if more recent certificates for the user are stored in Active Directory, then these credentials are copied to the client computer.

6. After the user's certificate store on the client computer has been updated, then outstanding certificate auto-enrollment requests are processed.

   **Note**

   - Newly issued certificates are stored in the certificate store on the client and replicated to the domain controller for the domain.

7. Later, when the user logs onto another client computer connected to the domain, the same Group Policy is applied and credentials are once again replicated from Active Directory. Credential roaming synchronizes and resolves any conflicts between certificates and private keys from any number of client computers that the user logs onto, as well as Active Directory.

   **Important**

   - In multidomain environments and domains with multiple domain controllers, credentials may not be immediately available when a user logs on to the network using one domain controller shortly after being issued a certificate on a computer that validates the user's identity against a different domain controller. The credentials will only become available after replication has been completed between the two domains or domain controllers.

8. On subsequent logons, the certificate in the local certificate and key stores are compared to the certificates and keys in Active Directory:

   - If the certificates and keys are current, no further action is taken.

   - If the certificates and keys need to be renewed, the certification authority (CA) is contacted from the laptop or

workstation, the renewal process occurs and the updated certificates and keys are replicated to Active Directory.

9. When the user logs onto the other computer, the updated certificates and keys are replicated to the certificate and key stores on this computer as well.

10. When the user's certificate expires, the old certificate is automatically archived in the client's profile and in Active Directory.

Credential roaming is triggered any time a private key or certificate in the user's local certificate store changes, whenever the user locks or unlocks their computer, and whenever Group Policy is refreshed.

All Credential roaming-related communication between components on the local computer and between the local computer and Active Directory is signed and encrypted.

# Credential roaming schema changes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The following schema attributes and attribute values will be used in Active Directory for storing user credentials and data relating to credential roaming:

- **ms-PKI-DPAPIMasterKey.** This multivalued attribute contains master key files and information for DPAPI. The following objects will be roamed and contained within this attribute:

- **All master key files.** There can be multiple master key files. New master key files can be created every 90 days by the domain. Master key files must be maintained and roamed in perpetuity.

- **The Preferred file that specifies the master key to be used for encryption.** This attribute is updated every time a new master key is created.

- **ms-PKI-AccountCredentials.** This multivalued attribute contains binary blobs of encrypted credential objects from the Stored User Names and Passwords store, private keys, certificates and requests. Each binary blob stored in Active Directory may contain a delete flag with a timestamp that persists for 60 days to ensure that all clients delete the object.

- **ms-PKI-RoamingTimeStamp.** This attribute is used by credential roaming to record the time of the latest change to the user object in Active Directory.

These schema attributes and attribute values must be added to Active Directory before credential roaming is deployed. Copy the following text into Notepad and Save it as CredRoam.ldf (be sure to use the .ldf extension), then refer to Preparing Active Directory to store user certificates and keys in order to configure your schema extensions.

```
# ----------------------------------------------------------------------
#   Copyright 2004 Microsoft Corporation
#
#   MODULE:     credroam.ldf
#   ABSTRACT:   add key roaming Active Directory schemas & attributes
# ----------------------------------------------------------------------

#   define property set Private-Information
# ----------------------------------------------------------------------

dn: cn=Private-Information,CN=Extended-Rights,CN=Configuration,DC=X
changetype: add
cn: Private-Information
objectClass: controlAccessRight
displayName: Private Information
appliesTo: 4828cc14-1437-45bc-9b07-ad6f015e5f28
appliesTo: bf967aba-0de6-11d0-a285-00aa003049e2
rightsGUID: 91e647de-d96f-4b70-9557-d63ff4f3ccd8
validAccesses: 48

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

# ----------------------------------------------------------------------
#   define schemas for:
#       ms-PKI-RoamingTimeStamp
#       ms-PKI-DPAPIMasterKeys
#       ms-PKI-AccountCredentials
#
#   NOTE: searchFlags 128 (CONFIDENTIAL_DATA) is implemented in an AD
#       server-side update such that the attributes having this bit
#       will not be readable except for SELF.
# ----------------------------------------------------------------------
```

```
dn: cn=ms-PKI-RoamingTimeStamp,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaadd
objectClass: attributeSchema
lDAPDisplayName: msPKIRoamingTimeStamp
adminDisplayName: MS-PKI-RoamingTimeStamp
adminDescription: Time stamp for last change to roaming tokens
attributeId: 1.2.840.113556.1.4.1892
attributeSyntax: 2.5.5.10
oMSyntax: 4
schemaIDGUID:: rOQXZvGiq0O2DBH70frPBQ==
attributeSecurityGUID:: 3kfmkW/ZcEuVV9Y/9PPM2A==
searchFlags: 128

dn: cn=ms-PKI-DPAPIMasterKeys,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaadd
objectClass: attributeSchema
lDAPDisplayName: msPKIDPAPIMasterKeys
adminDisplayName: MS-PKI-DPAPIMasterKeys
adminDescription: Storage of encrypted DPAPI Master Keys for user
attributeId: 1.2.840.113556.1.4.1893
attributeSyntax: 2.5.5.7
linkID: 2046
oMSyntax: 127
oMObjectClass:: KoZIhvcUAQEBCw==
schemaIDGUID:: IzD5szmSfE+5nGdF2Hrbwg==
attributeSecurityGUID:: 3kfmkW/ZcEuVV9Y/9PPM2A==
searchFlags: 128

dn: cn=ms-PKI-AccountCredentials,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaadd
objectClass: attributeSchema
lDAPDisplayName: msPKIAccountCredentials
adminDisplayName: MS-PKI-AccountCredentials
adminDescription: Storage of encrypted user credential token blobs for roaming
attributeId: 1.2.840.113556.1.4.1894
attributeSyntax: 2.5.5.7
linkID: 2048
oMSyntax: 127
oMObjectClass:: KoZIhvcUAQEBCw==
schemaIDGUID:: RKffuNwx8U6sfIS69++dpw==
attributeSecurityGUID:: 3kfmkW/ZcEuVV9Y/9PPM2A==
searchFlags: 128

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-
# ----------------------------------------------------------------------
#   adding the attributes to User class.
# ----------------------------------------------------------------------

dn: CN=User,CN=Schema,CN=Configuration,DC=X
changetype: modify
add: mayContain
mayContain: msPKIRoamingTimeStamp
mayContain: msPKIDPAPIMasterKeys
mayContain: msPKIAccountCredentials
-

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-
```

**See Also**

**Concepts**

Best practices for Active Directory Schema

LDIFDE

**Other Resources**

# Credential roaming administrative template

Updated: October 2, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

The Credential Roaming administrative template allows you to configure the following credential roaming options:

- **Not Configured**. The default setting.
- **Enabled**. Enables credential roaming.
- **Disabled**. Prevents credential roaming from being enabled locally.

You can also configure specific Credential Roaming settings:

- **Maximum tombstone credentials lifetime in days:** Allows you to define how long a roaming credential will remain in Active Directory for a certificate or key that has been deleted locally.
- **Maximum number of roaming credentials per user:** Allows you to define a maximum number of certificates and keys that can be used with credential roaming.
- **Maximum size (in bytes) of a roaming credential:** Allows you to restrict roaming for credentials that exceed a defined size.

To create the Credential Roaming administrative template, copy the following code, paste it into a text file, and save the file with the name DIMS.adm or a name of your choice.

```
CLASS USER
CATEGORY  !!CertSvcClient
KEYNAME "Software\Policies\Microsoft\Cryptography\AutoEnrollment"
POLICY !!CredentialRoaming

EXPLAIN !!CredentialRoaming_Explain
VALUENAME "DIMSRoaming"
VALUEON NUMERIC 1

PART !!CredentialRoaming_Box TEXT
END PART

PART !!CredentialRoaming_TombstoneValue NUMERIC REQUIRED
VALUENAME "DIMSRoamingTombstoneDays"
MIN 1 MAX 3650 DEFAULT 60 SPIN 30
END PART

PART !!CredentialRoaming_MaxNumTokens NUMERIC REQUIRED
VALUENAME "DIMSRoamingMaxNumTokens"
MIN 1 MAX 10000 DEFAULT 2000 SPIN 100
END PART

PART !!CredentialRoaming_MaxTokenSize NUMERIC REQUIRED
VALUENAME "DIMSRoamingMaxTokenSize"
MIN 1 MAX 100000 DEFAULT 65535 SPIN 1000
END PART

END POLICY
END CATEGORY

[strings]
CertSvcClient="Certificate Services Client"
CredentialRoaming_Explain="This policy setting specifies the behavior for user X.509
certificates, requests, and key roaming.\n\n
User certificates and keys will be roamed and synchronized between the local user
profile on the desktop and the user object in
```

```
Active Directory when a user logs on interactively.\n\nIf you enable this policy
setting, all X.509 certificates, keys, and enrollment
requests will be uploaded and synchronized with the user object in Active Directory.
You should also enable folder exclusion policies
for roaming user profiles to avoid any conflicts in the use of multiple roaming
technologies.\n\nIf you disable this policy setting,
all future synchronization and roaming will cease, but no keys or certificates will be
deleted from the local user profile or
Active Directory user object.\n\nIf you do not configure this policy setting, user
certificate and key roaming will not be performed.\n\n
Note: Folder exclusion policy settings may be configured in the user profiles section
of the System administrative template.\n\n"
DisableAll="None"
CredentialRoaming="X.509 certificate and key roaming"
CredentialRoaming_Box="Specific Credential Roaming settings:"
CredentialRoaming_TombstoneValue="Maximum tombstone credentials lifetime in days:"
CredentialRoaming_MaxNumTokens="Maximum number of roaming credentials per user:"
CredentialRoaming_MaxTokenSize="Maximum size (in bytes) of a roaming credential:"
```

**See Also**

**Concepts**

The role of Administrative Templates

# Credential roaming best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certificates and private keys are important elements of an organization's security infrastructure. Therefore, you need to consider the following best practices and other considerations when planning and implementing credential roaming:

- Credential roaming is supported for clients running Windows Server 2003 with Service Pack 1 (SP1). Credential roaming is supported on domain controllers running Windows 2000 Service Pack 3 or later or Windows Server 2003. Mixed environments containing domain controllers running Windows 2000 SP3 or later and Windows Server 2003 are also supported. The forest functional level can be either Windows 2000 or Windows Server 2003.

- An update was recently made available that enables credential roaming on clients running Windows XP Professional (SP2) and an update is available for Windows Server 2003 SP1.

  **Note**

  - For the best performance and security, it is recommended that the servers be upgraded to Windows Server 2003 SP1.

- Organizations should not enable both roaming user profiles and credential roaming for the same user. Although the two capabilities are not mutually exclusive, unexpected results can occur if both are used by a given user. For example, if certificates and private keys are stored in both Active Directory and as part of their roaming user profile, the user might use the less current versions.

  **Note**

  - For more information about configuring roaming user profiles, see Implementing Roaming User Profiles in the Windows Server 2003 Deployment Kit.

- Credential roaming does not provide complete protection for certificates and private keys if they are accidentally deleted by a user. Use certificate and key archival and recovery to protect credentials from accidental deletion. For more information, see Key archival and recovery.

- Credentials such as enrollment agent certificates that require very strong protection of private keys should not be used in conjunction with credential roaming; instead, these credentials should be stored on smart cards and usage of the smart card should be required each time the user logs on.

- Organizations that are concerned about keys being left on computers that might be stolen and subject to brute force attacks should seriously consider whether to implement credential roaming for their most sensitive users and data. If the benefits of credential roaming outweigh the potential risks, they can configure the user's profile by using Group Policy to delete on logoff the user's profile that includes the certificates and private keys. For more information, see Using User Profiles. Alternately, require smart cards to be used for every logon.

- If Encrypting File System (EFS) is used, it is recommended that the data recovery certificate (DRA) and private key not be imported into the end user's account profile. Importing DRA certificates into the user's account profile would unintentionally roam the DRA key and certificate to the user's account in Active Directory. It is a best practice to use a dedicated DRA or separate user account for the purposes of EFS data recovery. In addition, DRA keys should never be imported into a non-administrator account.

- Organizations that are concerned about an administrator resetting a user password and logging onto the system as the user to steal or use their keys should use smart cards or strong private key protection to prevent this from happening. Credential roaming is not supported for clients who are migrated from one domain to another. The user must export all keys and certificates before the account is migrated and then reimport all keys and certificates after the account is migrated to the new domain.

# Viewing certificate information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing certificate information**

Using Certificates, you can view information about the fields, extensions, and properties that define an issued certificate by double-clicking any certificate displayed in the certificate store.

| To view this | Use this |
|---|---|
| Overview of the certificate information | **General** tab |
| Supported uses of the certificate | **General** tab |
| Entity to which the certificate was issued | **General** tab |
| Validity period of the certificate | **General** tab |
| Path of certificate to a trusted root certification authority | **Certification path** tab |
| Root certificate | **Certification path** tab |
| Subordinate certification authorities' certificates | **Certification path** tab |
| Whether root certification authority is trusted or not | **Certification path** tab |
| Certificate status | **Certification path** tab |

The **Details** dialog box provides the following information about the certificate.

| Certificate field | Description |
|---|---|
| Version | The X.509 version number. |
| Serial number | The unique serial number that the issuing certification authority assigns to the certificate. The serial number is unique for all certificates issued by a given certification authority. |
| Signature algorithm | The hash algorithm that the certification authority uses to digitally sign the certificate. |
| Issuer | Information regarding the certification authority that issued the certificate. |
| Valid from | The beginning date for the period in which the certificate is valid. |
| Valid to | The final date for the period in which the certificate is valid. |
| Subject | The name of the individual or certification authority to whom the certificate is issued. If the issuing certification authority exists on a domain member server in your enterprise, this will be a distinguished name within the enterprise. Otherwise, this may be a full name and e-mail name or some other personal identifier. |
| Public key | The public key type and length associated with the certificate. |
| Thumbprint algorithm | The hash algorithm that generates a digest of data (or thumbprint) for digital signatures. |
| Thumbprint | The digest (or thumbprint) of the certificate data. |
| Friendly name | (Optional) A friendly, or common, name for the name in the Subject field. |

| | |
|---|---|
| Enhanced key usage | (Optional) The purposes for which this certificate can be used. |

There are additional possible X.509v3 extensions that could be in a certificate. If they are present, they will be displayed as well.

For the procedure that shows how to see the properties of a certificate, see View a certificate.

# Using Windows 2000 Certificate Services Web pages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Windows 2000 Certificate Services Web pages**

By default, every certification authority that is hosted on a server running Windows 2000 has Web pages available for users and administrators. These are used to perform a variety of tasks related to requesting certificates. These Web pages are located at http://*servername*/certsrv, where *servername* is the name of the server running Windows 2000 that hosts the certification authority. The **certsrv** portion of the URL should always be in lowercase letters; otherwise, users may have trouble checking and retrieving pending certificates. Users can access these Web pages using Microsoft Internet Explorer version 5.0 and greater or Netscape Navigator version 3.01 and greater.

The Web pages are the *only* way for users to request certificates from stand-alone certification authorities. They are optional for users who want to request certificates from enterprise certification authorities. For more information about certification authorities, see Certificates and certification authorities.

If you have been granted access permissions, you can perform the following tasks from these Web pages:

- Request a basic certificate. For more information, see Submit a user certificate request via the Web.
- Request a certificate with advanced options. For more information, see Submit an advanced certificate request via the Web.

  This gives you greater control over the certificate request. Some of the user-selectable options that are available in an advanced certificate request include:

  - **Cryptographic service provider (CSP) options.** The name of the cryptographic service provider, the key size (512, 1024, and so on), the hash algorithm (SHA/RSA, SHA/DSA, MD2, MD5) and the key spec (exchange or signature).
  - **Key generation options.** Create a new key set or use an existing key set, mark the keys as exportable, enable strong key protection, and use the local computer store to generate the key.
  - **Additional options.** Save the request to a PKCS #10 file or add any specific attributes you want to add to the certificate.

  Netscape clients cannot use the advanced option Web pages.

- Check on a pending certificate request. If you have submitted a certificate request to a stand-alone certification authority, you will need to check the status of the pending request to see if the certification authority has issued the certificate. If the certificate has been issued, it will be available for you to install it. For more information, see Check on a pending certificate request.
- Retrieve the certification authority's certificate to place in your trusted root store. For more information, see Retrieve a certification authority certificate.
- Retrieve the current certificate revocation list. For more information, see Retrieve a certificate revocation list.
- Submit a certificate request using a PKCS #10 file or a PKCS #7 file. For more information, see Request a certificate using a PKCS #10 or PKCS #7 file.
- (For administrators only) Request a certificate for a smart card on behalf of another user. For more information, see Set up a smart card for user logon.

# Using Windows Server 2003 Certificate Services Web pages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Windows Server 2003 Certificate Services Web pages**

By default, every Windows Server 2003 certification authority has Web pages available for users and administrators. These can be used to perform a variety of tasks related to requesting certificates. These Web pages are located at http://*servername*/certsrv, where *servername* is the name of the server that hosts the certification authority. The **certsrv** portion of the URL should always be in lowercase letters; otherwise, users may have trouble checking and retrieving pending certificates. Users can access these Web pages using Microsoft Internet Explorer version 5.0 and greater, or Netscape Navigator version 3.01 and greater.

The Web pages are the *only* way for users to request certificates from stand-alone certification authorities. They are optional for users who want to request certificates from enterprise certification authorities. For more information about certification authorities, see Certificates and certification authorities.

If you have been granted access permissions, you can perform the following tasks from these Web pages:

- Request a basic certificate. For more information, see Submit a user certificate request via the Web to a Windows Server 2003 CA.

- Request a certificate with advanced options. For more information, see Submit an advanced certificate request via the Web to a Windows Server 2003 CA.

  This gives you greater control over the certificate request. Some of the user-selectable options available in an advanced certificate request include:

  - **Cryptographic service provider (CSP) options.** The name of the cryptographic service provider, the key size (512, 1024, and so on), the hash algorithm (SHA/RSA, SHA/DSA, MD2, MD5) and the key spec (exchange or signature).

  - **Key generation options.** Create a new key set or use an existing key set, mark the keys as exportable, enable strong key protection, and use the local computer store to generate the key.

  - **Additional options.** Save the request to a PKCS #10 file or add any specific attributes you want to add to the certificate.

  Netscape clients cannot use the advanced option Web pages.

- Check on a pending certificate request to a Windows Server 2003 CA. If you have submitted a certificate request to a stand-alone certification authority, you will need to check the status of the pending request to see if the certification authority has issued the certificate. If the certificate has been issued, it will be available for you to install it. For more information, see Check on a pending certificate request to a Windows Server 2003 CA.

- Retrieve the certification authority's certificate to place in your trusted root store or install the entire certificate chain in your certificate store. For more information, see Retrieve a certification authority certificate from a Windows Server 2003 CA.

- Retrieve the current base certificate revocation list (CRL) and delta CRL. For more information, see Retrieve a certificate revocation list from a Windows Server 2003 CA.

- Submit a certificate request using a PKCS #10 file or a PKCS #7 file. For more information, see Request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file.

- (For administrators only) Request a certificate for a smart card on behalf of another user. For more information, see Set up a smart card for user logon.

# Managing Certificate Trusts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing certificate trusts**

This section covers:

- Automatic updating of trusted root authorities
- Managing trust of third-party certification authorities
- Managing trust of user-selected certification authorities

# Automatic updating of trusted root authorities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automatic updating of trusted root authorities**

You typically use a certificate when you use a secure Web site or when you send and receive secure e-mail. Theoretically, anyone can issue certificates, but to have truly secure transactions, certificates must be issued by a trusted entity or organization. Microsoft has included a list--in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; and other products--of companies and organizations that it considers to be trusted authorities.

Typically, when you are presented with a certificate issued by an authority that is not in the trusted authority list that is provided with your browser or operating system, you are asked if you want to establish trust in the certification authority (CA) which issued the certificate. Many users do not want to establish trust to an authority in this way, since they have limited resources to verify the trustworthiness and issuing policies of the CA.

In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, you can use the Update Root Certificates function for this. Update Root Certificates is turned on by default. With this feature turned on, if you are presented with a certificate issued by an untrusted root authority, your computer will contact the Windows Update Web site to see if Microsoft has added the CA to its list of trusted authorities. If it has been added to the Microsoft list of trusted authorities, its certificate will automatically be added to your trusted certificate store.

For more information, see Turn off automatic updating of trusted root authority certificates, Turn on automatic updating of trusted root authority certificates, and Certificate stores.

# Managing trust of third-party certification authorities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing trust of third-party certification authorities**

By default, when you install Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, a large number of certification authority (CA) certificates are listed in the Third-Party Root Certification Authorities physical store.

To view the certificates in this store, you can follow the procedures in Manage certificates for a computer and Display certificate stores storage structure. In the console tree, click **Certificates -*Certificate Holder***, click **Third-Party Root Certification Authorities**, and then click **Certificates**.

The Third-Party Root Certification Authorities store contains the trusted root certification authorities (CAs) from companies outside your organization. Administrators can disable the trust of these CAs by selecting **Trusted Root Certification Authorities only** option in Group Policy. Then, when users access any secure Web site that is validated by one of the third-party root CAs, they receive a security alert message that informs them that the site is not trusted.

To disable trust in third-party root CAs for a domain (or for any Group Policy object), see Disable trust of third-party root CAs for a Windows Server 2003 domain.

For more information, see Managing trust of user-selected certification authorities and Certificate stores.

# Managing trust of user-selected certification authorities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing trust of user-selected certification authorities**

For users, all trusted root certificates are stored under Trusted Root Certification Authorities.

You can view the certificates in this store by following the procedures in Manage certificates for your user account and Display certificate stores storage structure. In the console tree, click **Certificates - Current User**, then click **Trusted Root Certification Authorities**, and then click **Certificates**.

Administrators can use Group Policy to disable the trust of user-selected root certification authorities (CAs). Then, when network users access any secure Web site that is validated by one of these user-selected root CAs, they receive a security alert message that informs them that the site is not trusted. To disable trust in user-selected root CAs for a domain (or for any Group Policy object), see Disable trust of user-selected root CAs for a Windows Server 2003 domain.

For more information, see Managing trust of third-party certification authorities and Certificate stores.

# Certificates Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

The following books provide detailed information about certificates, certification authorities, public key encryption, and public key infrastructure (PKI):


- Austin, Thomas, *PKI:A Wiley Tech Brief*. John Wiley & Sons, December 2000.
- Feghhi, Jalal and Jalil Feghhi, *Secure Networking with Windows 2000 and Trust Services*. Addison-Wesley Pub Co., February 2001.
- Housley, R.P., and T. Polk. 2001. *Planning for PKI*. New York, New York: John Wiley & Sons, Inc.
- Howard, Michael, *Designing Secure Web-Based Applications for Microsoft Windows 2000*. Microsoft Press, August 2000.
- Kaufman, C., Perlman, R., and M. Speciner. 1995. *Network Security: Private Communication in a Public World*. Upper Saddle River, New Jersey: Prentice-Hall Inc.
- Schneier, Bruce, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. John Wiley & Sons, December 1995.


The following resources have in-depth information about planning, deploying, and using a public key infrastructure with Windows Certificate Services:

- Microsoft Security on the Microsoft Web site.

For updated technical information, see:

- For more information about certificate services architecture, see "Designing a Public Key Infrastructure" at the Microsoft Windows Resource Kits Web site.
- Microsoft TechNet on the Microsoft Web site.
- Updated technical information on the Microsoft Web site.

# Certificates Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The certificate template name is a long string of alphanumeric characters in the Certificates snap-in.
- When trying to use the Certification Authority Web enrollment pages, the following error appears: "The proper version of the ActiveX control failed to download and install. You may not have sufficient permissions. Please ask your system administrator for assistance."

**The certificate template name is a long string of alphanumeric characters in the Certificates snap-in.**
**Cause:** The string is an object identifier (also known as an OID). Each certificate template is represented by an object identifier. The snap-in has not received the full name of the template from the certification authority.

**Solution:** Close and reopen the Certificates snap-in.

**See also:** Add the Certificates Snap-in to the MMC Console.

**When trying to use the Certification Authority Web enrollment pages, the following error appears: "The proper version of the ActiveX control failed to download and install. You may not have sufficient permissions. Please ask your system administrator for assistance."**
**Cause:** The user who is accessing the Web pages is not a member of the Administrators or Power Users group on the local computer. When a newer version of the Web enrollment software is available on the certification authority (CA), the client computer must install that software. The user must be a member of the Administrators or Power Users group to install the software.

**Solution:** Log on to the network as a user who is a member of the Administrators or Power Users group to access the Web enrollment pages and download the newer version of the software.

# Certificate Services

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Certificate Services provides customizable services for issuing and managing certificates that are used in software security systems that employ public key technology. Certificate Services is available on computers running Microsoft® Windows Server® 2003, Standard Edition; Microsoft® Windows Server™ 2003, Enterprise Edition; and Microsoft® Windows Server™ 2003, Datacenter Edition.

**Tip**

Overviews for more recent versions of certificate services are available:

- Windows Server® 2012: What's New in AD CS?
- Windows Server® 2008 and Windows Server® 2008 R2: Active Directory Certificate Services.
- Community Overview: Active Directory Certificate Services (AD CS) Overview.

- Before deploying certification authorities and a public key infrastructure, see Checklist: Creating a certification hierarchy with an offline root certification authority.
- Before configuring certification authorities and clients for certificate autoenrollment, see Checklist: Configuring certificate autoenrollment.
- Before removing a deployed certification authority from your enterprise, see Checklist: Decommissioning a certification authority.
- For tips about using Certificate Services, see Certificate Services Best practices.
- For help with specific tasks, see Certificate Services How To....
- For general background information, see Certificate Services Concepts.
- For problem-solving instructions, see Certificate Services Troubleshooting.

# Checklist: Creating a certification hierarchy with an offline root certification authority

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

| **Note** |
| --- |
| A revised version of this information has been placed on the TechNet Wiki in the Offline Root Certification Authority (CA) article (http://social.technet.microsoft.com/wiki/contents/articles/2900.aspx). The Checklist section (http://social.technet.microsoft.com/wiki/contents/articles/2900.aspx#Checklist_Creating_a_certification_hierarchy_with_an_offline_root_certification_authority) is the replacement for the information below. |

**Checklist: Creating a certification hierarchy with an offline root certification authority**

This checklist is provided for the cases where the root certification authority (CA) is not connected to your organization's network. You might choose to have an isolated, offline root CA for security reasons in order to protect it from possible attacks by intruders by way of the network.

Administrators of public key infrastructures (PKIs) must provide certificate verifiers with online certificate revocation checking. This checklist helps them to set up functional certificate revocation checking for certificates issued by an offline root CA.

| Step | Reference |
| --- | --- |
| **Review concepts** | |
| Review public key infrastructure concepts. | Deploying a Public Key Infrastructure |
| Review certificates concepts. | Certificates Concepts |
| Review concepts about certification authorities. | Certificate Services Concepts |
| **Set up the offline root certification authority** | |
| Plan the certification hierarchy. | Certification authority hierarchies |
| Set up a server that runs Windows that you will use for the root certification authority. The server should not be a member of any domain, should be disconnected from the network, and should be physically secure. The server should also have Internet Information Services (IIS) installed as part of the setup process, although this is not required. | Checklist: Performing a new installation |
| Plan the renewal strategy you are going to use for the root certification authority | Renewing a certification authority |
| Log on to the server as the administrator and install Certificate Services to create a stand-alone root certification authority. | Install a stand-alone root certification authority |
| **Prepare the offline root certification authority to issue certificates** | |
| On the new root CA, change the default action upon receipt of a certificate request so that all requested certificates are set to pending. This is to ensure only authorized requests are issued by the top-level CA. | Set the default action upon receipt of a certificate request |
| On the new root CA, change the URL location of the certificate revocation list (CRL) distribution point to a location of your choice that is accessible to all users in you organization's network. It is possible to enter multiple URLs. It is necessary to do this because the offline root CA's default CRL Distribution Points (CDPs) are not accessible to users on the network and, if they are left unchanged, certificate revocation checking will fail. | Specify certificate revocation list distribution points in issued certificates |
| On the new root CA, change the URL location of the authority information access (AIA) distribution points to a location of your choice that is accessible to all users in you organization's network. It is possible to enter multiple URLs. It is necessary to do this because the offline root CA's default AIA points are not accessible to users on the network and, if they are left unchanged, certificate chain verification will fail. | Specify CA certificate access points in issued certificates |
| Schedule the publication of the certificate revocation list. Since publishing the CRL from an offline CA has the administrative overhead of having to physically copy the CRL to a server on the network, you may want to have a lengthy validity period. | Schedule the publication of the certificate revocation list |
| On the root certification authority, publish the certificate revocation list. | Manually publish the certificate revocation list |
| In Windows Explorer on the root CA, locate the certificate revocation list you just published. The CRL's default location is:<br><br>**** *systemroot* **\system32\CertSrv\CertEnroll** *CAname* **.crl**<br><br>Right-click the CRL file and send it to a drive that has portable storage media. | |
| Retrieve the certification authority's certificate and save it to a drive that has portable storage media. | Retrieve a certification authority certificate |

| | | |
|---|---|---|
| | Copy the certificate revocation list file and the CA certificate to every URL location that you specified as a CRL distribution point in the root CA's policy settings. | |
| | Copy the CA certificate file to every URL location that you specified as an authority information access distribution point in the root CA's policy settings. | |
| | **If you are deploying your PKI in an Active Directory directory service environment** | |
| | Publish the root certificate to the enterprise root store and add the certificate to the customary Authority Information Access (AIA) points in the directory. You need to use certutil.exe. You can also use this command to put the CA certificate from a third party root CA into Active Directory. | At a command prompt, type: **certutil-dspublish-f**.*Crt File Name***RootCA** |
| | Publish the CRL to the customary location in Active Directory. To do this, use certutil.exe. You can also use this command to put the CRL from a third-party root CA into Active Directory. | From the command line, type: **certutil-dspublish-f**.*Crl File Name* |
| | **To create a online certification authority that is subordinate to an offline root certification authority** | |
| | Set up a server running Windows to use for the subordinate certification authority | Checklist: Performing a new installation |
| | Install subordinate certification authorities, as required by your planned certification hierarchy. These can be stand-alone certification authorities or, if you are using Active Directory, enterprise certification authorities. During setup for each subordinate CA, choose to save the CA certificate request to a file, which will be a PKCS #10 request. | Install a stand-alone subordinate certification authority;Install an enterprise subordinate certification authority |
| | Copy the CA certificate request file from the subordinate certification authority to some portable storage media. Take the CA certificate request to the root certification authority. | |
| | Using the Certificates Microsoft Management Console (MMC) on the offline CA, submit the certificate request (*requestfilename*) to the CA and copy the new certificate (*newcertname*) to the portable storage media. | Manage certificates for a computer |
| | Take the portable storage media back to the subordinate certification authority. In Windows Explorer, locate the certificate and certification path files you just copied, then right-click each file and choose **Install Certificate**. Have the Certificate Import Wizard automatically place the certificates in stores based on the type of certificate. | Import a certificate |

# Checklist: Configuring certificate autoenrollment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring certificate autoenrollment**

| Step | Reference |
|---|---|
| Install and configure an enterprise certification authority (CA) | Install an enterprise root certification authority; Install an enterprise subordinate certification authority |
| Create a new certificate template | Create a new certificate template |
| Configure the certificate template to provide the intended functionality for the certificate | Modify a Certificate Template |
| Configure the certificate template to allow autoenrollment | Planning for autoenrollment deployment |
| Grant Enroll and Autoenroll permissions for the certificate template to intended users | Allow subjects to request a certificate that is based on the template |
| Configure the certification authority to issue certificates based on the template | Add a certificate template to a certification authority |
| Configure Active Directory directory service users to request certificates with autoenrollment | Certificate Services example implementation: Establishing autoenrollment for user certificates |
| (Optional) Retrieve updated Group Policy information before scheduled replication occurs | Gpupdate |

**Notes**

- Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition is required to configure version 2 certificate templates for autoenrollment requests. However, autoenrollment manages certificates or pending certificate requests that are based on any version of certificate template.

# Checklist: Decommissioning a certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Decommissioning a certification authority**

Once a certification authority (CA) is configured and operating in an enterprise, it becomes an important security resource. The certification authority is required by existing clients to renew certificates, by new clients to issue certificates, and by users of certificates to verify the trustworthiness of issued certificates. If a CA is removed from an enterprise without following the appropriate procedure, functionality may be impaired and extensive cleanup work may be required to restore clients to functionality.

| | Step | Reference |
|---|---|---|
| | Perform a backup of the certification authority (CA) database to ensure recoverability of data at a later date. | Back up a certification authority |
| | Deny all pending certificate requests currently stored on the CA. | Review pending certificate requests |
| | (Optional) Allow the Active Directory directory service to replicate the certificate denials and allow Group Policy to inform the clients that the certificate request is denied. The denied clients are then informed that their request was denied so they can remove the request from their list of outstanding requests. | Replication overview |
| | Revoke the CA's certificate from its parent CA. When revoking the CA certificate, specify the reason as "Cessation of Operation." | Revoke an issued certificate |
| | Manually publish a new certificate revocation list (CRL) to ensure the CRL contains revocation information about the CA that was recently revoked. | Manually publish the certificate revocation list |
| | Uninstall the Certificate Services component. | Uninstall a certification authority |
| | Remove remaining information about this CA from Active Directory. | At a command prompt, type **certutil.exe -dsdel *CAName*** and press ENTER, where *CAName* is the name of the CA you are removing. |

# Certificate Services Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate Services Best practices**

**Plan your public key infrastructure (PKI) before deploying certification authorities (CAs)**

- For information that will assist you with planning a PKI, see Certificate Services Concepts in Certificate Services and Certificates Resources in Public Key Infrastructure.

**Place database and transaction log files on separate hard drives**

- As in many databases, the certification authority's database is a file on the hard drive. In addition to this file, other files serve as the transaction logs and receive all modifications to the database before the changes are made. Because these files may be accessed frequently and simultaneously, it is best to keep the database and transaction logs on separate hard drives or high-performance disk configurations, such as striped volumes.

  For more information, see Disk Management.

**Keep the root certification authority offline and secure its signing key by hardware and keep it in a vault to minimize potential for key compromise**

- For more information, see Checklist: Creating a certification hierarchy with an offline root certification authority.

**If you are going to use a custom policy module for a Microsoft CA, install Certificate Services using stand-alone policy and then replace stand-alone policy with your custom policy**

- Replacing enterprise policy on a certification authority with a custom policy is not supported and will have unpredictable results.

**When changing security permissions for the certification authority (CA), always use the Certification Authority snap-in**

- Setting permissions using other tools (such as the Active Directory Sites and Services snap-in) may create problems for users who attempt to access and request certificates from the certification authority.

  For more information, see Set security permissions and delegate control of a certification authority.

**Do not issue certificates to users or computers directly from the root certification authority**

- Instead, deploy at least a two-level CA hierarchy comprised of Root-Issuer CAs to provide flexibility and insulate the root certification authority from attempts to compromise its private key by malicious individuals.

  For more information, see Certification authority hierarchies and Establishing a certification hierarchy.

**Back up the CA database, the CA certificate, and the CA keys**

- This is essential to protect against the loss of critical data. The CA should be backed up on a regular basis (daily, weekly, monthly), based on the number of certificates issued over the same interval. The more certificates issued, the more frequently you should back up the CA. Full backups should be used to provide the fastest recovery and most reliable data redundancy possible.

  For more information, see Backing up and restoring a certification authority.

**Ensure that key lifetimes are long enough to avoid renewal issues**

- A certification authority can issue a certificate with a lifetime equal to or less than the remaining lifetime on its own CA certificate. When the CA certificate has less valid time remaining than the renewal overlap period on issued certificates, the issued certificates are very short lived. In this case, it is possible that clients will be unable to renew their certificates. To avoid this, ensure the CA certificate is renewed long enough before the certificates it issues become short lived.

  For more information, see Renew a root certification authority and Renew a subordinate certification authority.

**Review the concepts of security permissions and access control, since enterprise certification authorities issue certificates based on the security permissions of the certificate requester**

- For general information about access control and security permissions in the Windows Server 2003 family, see Access control overview. For the procedure to set access control on a CA, see Set security permissions and delegate control of a certification authority.

  For the procedure to set access control on certificate templates, see Controlling enrollment access to certificate templates.

**Use Secure Sockets Layer (SSL) when using Web-based certificate enrollment**

- Web enrollment pages provide a simple but insecure method for users to obtain certificates from the CA. The communication between the client and the Web server could be forged or altered in route, compromising security. To mitigate the risk of this occurring, you can use Secure Sockets Layer (SSL) communications between the client and the Web server. This will provide both confidentiality and integrity of the certificate request and response.

  For more information, see Configuring SSL on your server and Set up certification authority Web enrollment support.

# Certificate Services How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Set Up a Certification Authority
- Administer a Certification Authority

# Set Up a Certification Authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up a certification authority**

- Install an enterprise root certification authority
- Install an enterprise subordinate certification authority
- Install a stand-alone root certification authority
- Install a stand-alone subordinate certification authority
- Uninstall a certification authority
- Set up certification authority Web enrollment support
- Perform qualified subordination

# Install an enterprise root certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install an enterprise root certification authority**

1. Log on as a member of both the Enterprise Admins group and the root domain's Domain Admins group.
2. If any Windows 2000 enterprise certification authorities (CAs) currently exist or have ever existed in your enterprise, open Certificate Templates and, when prompted to install new certificate templates, click **OK**.
3. Open Add or Remove Programs in Control Panel.
4. Click **Add/Remove Windows Components**.
5. In the Windows Components Wizard, select the **Certificate Services** check box. A dialog box appears to inform you that the computer cannot be renamed and that the computer cannot be joined to or removed from a domain after Certificate Services is installed. Click **Yes**, and then click **Next**.
6. Click **Enterprise root CA**.
7. (Optional) Select the **Use custom settings to generate the key pair and CA certificate** check box, and then click **Next** to specify the following.

| To set this | Do this |
|---|---|
| Cryptographic service provider (CSP) | In **CSP**, click the CSP that you want to use. The default is the Microsoft Strong Cryptographic Provider. Certificate Services does support third party CSPs, but you must refer to the CSP vendor's documentation for information about using their CSP with Certificate Services. |
| Hash algorithm | In **Hash algorithm**, click the hash algorithm you want to use. The default is SHA-1. |
| Use an existing key | Select the **Use existing key** check box, click **Import**, and then, in **Open PFX File**, type the file name and password of the public and private key pair. This is helpful if you are relocating or restoring a previously installed certification authority (CA). Note that, when using an existing key, a new certificate is generated.<br>**Important**<br>• Be sure that you select an existing key that you know to be uncompromised and trustworthy. Using a key that may be compromised or untrusted could cause this CA and all its issued certificates to be insecure. |
| Key length | In **Key length**, type or select a key length. The default key length using the Microsoft Strong Cryptographic Provider is 2048 bits. Default key lengths for other CSPs vary. In general, the longer the key length, the more secure the key is. Also, longer key lengths require more system resources for operations such as signing, encryption, and chain verification. For a root CA, you should use a key length of at least 2048 bits. This option is not available if you use existing keys. |
| Allow this CSP to interact with the desktop | Select the **Allow this CSP to interact with the desktop** check box. Without this option, system services cannot interact with the desktop of the user who is currently logged on. |
| Import | Click **Import**. This imports an existing key in the PKCS #12 PFX format. |

| | |
|---|---|
| View certificate | Click **View certificate**. This allows you to view the certificate that you select or generate during installation. |

When you are done, click **Next**.

8. In **Common name for this CA**, type the common name of the certification authority.

9. In **Validity period**, specify the validity duration for the root CA and click **Next**.

10. Specify the storage locations of the certificate database, the certificate database log, and the shared folder and click **Next**.

11. If Internet Information Services (IIS) is running, you will see a request to stop the service before proceeding with the installation and click **OK**.

12. If prompted, type the path to the Certificate Services installation files.

**Notes**

- The preliminary information that is supplied during setup, such as the name of the certification authority, cannot be changed after the CA setup is complete.

- Before installing a certification authority, you should configure the computer's domain settings first, such as joining a domain or promoting a server to a domain controller. These settings cannot be changed once the certification authority is installed.

- If Active Server Pages are not enabled through Internet Information Services, you will be prompted to activate them. The Web interface for the certification authority requires running Active Server Pages.

- The enterprise root CA selection requires that the host computer be a member of a domain and that it use the Active Directory directory service. The administrator who is installing an enterprise CA must have Write permission to Active Directory.

- If you have Write permission to Active Directory, then specifying the shared folder is optional, and is not typically done for enterprise certification authorities.

- The validity duration you choose for the CA will determine when the CA "expires." For information about renewing CAs, see Related Topics.

- If you installed the enterprise certification authority as an Enterprise Admin or delegated user, then you must use the Enterprise Admin or delegated user account when you uninstall the enterprise certification authority.

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

- Once the certification authority is installed, add certificate templates to the certification authority and configure the certification authority to allow subjects to request a certificate that is based on a template. For more information, see Related Topics.

- For more information on any installation settings, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Enterprise certification authorities
Install an enterprise subordinate certification authority
Renewing a certification authority
Add a certificate template to a certification authority
Install new templates and upgrade existing templates

# Install an enterprise subordinate certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install an enterprise subordinate certification authority**

1. Log on as a member of both the Enterprise Admins group and the root domain's Domain Admins group.
2. If any Windows 2000 enterprise certification authorities (CAs) currently exist or have ever existed in your enterprise, open Certificate Templates and when prompted to install new certificate templates, click **OK**.
3. Open Add or Remove Programs in Control Panel.
4. Click **Add/Remove Windows Components**.
5. In the Windows Components Wizard, select the **Certificate Services** check box. A dialog box appears to inform you that the computer cannot be renamed and that the computer cannot be joined to or removed from a domain after Certificate Services is installed. Click **Yes**, and then click **Next**.
6. Click **Enterprise subordinate CA**.
7. (Optional) Select the **Use custom settings to generate the key pair and CA certificate** check box, and then click **Next** to specify the following.

| To set this | Do this |
|---|---|
| Cryptographic service provider (CSP) | In **CSP**, click the CSP that you want to use. The default is the Microsoft Strong Cryptographic Provider. Certificate Services does support third party CSPs, but you must refer to the CSP vendor's documentation for information about using their CSP with Certificate Services. |
| Hash algorithm | In **Hash algorithm**, click the hash algorithm you want to use. The default is SHA-1. |
| Use an existing key | Select the **Use existing key** check box, click **Import**, and then, in **Open PFX File**, type the file name and password of the public and private key pair. This is helpful if you are relocating or restoring a previously installed certification authority (CA). Note that, when using an existing key, a new certificate is generated.<br>**Important**<br>• Be sure that you select an existing key that you know to be uncompromised and trustworthy. Using a key that may be compromised or untrusted could cause this CA and all its issued certificates to be insecure. |
| Key length | In **Key length**, type or select a key length. The default key length using the Microsoft Strong Cryptographic Provider is 2048 bits. Default key lengths for other CSPs vary. In general, the longer the key length, the more secure the key is. Also, longer key lengths require more system resources for operations such as signing, encryption, and chain verification. For a root CA, you should use a key length of at least 2048 bits. This option is not available if you use existing keys. |
| Allow this CSP to interact with the desktop | Select the **Allow this CSP to interact with the desktop** check box. Without this option, system services cannot interact with the desktop of the user who is currently logged on. |
| Import | Click **Import**. This imports an existing key in the PKCS #12 PFX format. |

| | |
|---|---|
| View certificate | Click **View certificate**. This allows you to view the certificate that you select or generate during installation. |

When you are done, click **Next**.

8. Type the common name (CN) of the CA and click **Next**.

9. Specify the storage locations of the certificate database, the certificate database log, and the shared folder. Click **Next**.

10. Obtain the certificate for the subordinate CA. For instructions on how to do this, see Notes.

11. If Internet Information Services is running, the system will request that you stop the service before proceeding with the installation. Click **OK**.

12. If prompted, type the path to the Certificate Services installation files.

13. Once the certification authority is installed, add certificate templates to the certification authority and configure the certification authority to allow subjects to request a certificate that is based on a template.

    For more information, see Related Topics.

**Notes**

- The name of the CA and the other initial information cannot be changed after the CA setup is complete.

- If Active server Pages are not enabled through Internet Information Services, you will be prompted to activate them. The Web interface for the certification authority requires running Active server Pages.

- To obtain the certificate for a subordinate CA, you must submit a certificate request to a parent CA. The procedure for doing so differs depending on whether the parent CA is available online.

    - If a parent CA is available online:

Click **Send the request directly to a CA already on the network**.

In **Computer Name**, type the name of the computer on which the parent CA is installed.

In **Parent CA**, click the name of the parent CA.

- If a parent CA is not available online:

Click **Save the request to a file**.

In **Request file**, type the path and file name of the file that will store the request.

Obtain this subordinate CA's certificate from the parent CA.

The procedure for doing this will be unique to the parent CA. At a minimum, the parent CA should provide a file containing the subordinate CA's newly issued certificate and, preferably, its full certification path. For the procedure to submit a certificate request using a file to a Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition CA, see Related Topics.

If you get a subordinate CA certificate that does *not* include the full certification path, the new subordinate CA you are installing must be able to build a valid CA chain when it starts. Thus you must install the parent CA's certificate in the Intermediate Certification Authorities certificate store of the computer (if the parent CA is not a root CA), as well as the certificates of any other intermediate CA in the chain, and you must install the certificate of the root CA in the chain into the Trusted Root Certification Authorities store. These certificates should be installed in the certificate store before you install the CA certificate on the subordinate CA you have just set up.

Open Certification Authority.

In the console tree, click the name of the CA.

**Where?**

- Certification Authority (*Computer*)/*CA name*

On the **Action** menu, point to **All Tasks**, and then click **Install CA Certificate**.

Locate the certificate file received from the parent certification authority, click this file, and then click **Open**.

If the root CA is untrusted, click **OK** to trust the root CA's certificate.

Once the certificate is installed, start Certificate Services.

To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

- The enterprise subordinate CA selection requires that the host computer be a member of a domain and that it use the Active Directory directory service. The administrator who is installing an enterprise CA must have Write permission to Active Directory.

- If you have Write permission to Active Directory, then specifying the shared folder is optional, and is not typically done for enterprise CAs.

- To open Certification Authority, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Certification Authority**.

- If you installed the enterprise subordinate certification authority as an Enterprise Admin or delegated user, then you must use the Enterprise Admin or delegated user account when you uninstall the enterprise subordinate certification authority.

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enterprise certification authorities
Install an enterprise root certification authority
Request a certificate using a PKCS #10 or PKCS #7 file
Install new templates and upgrade existing templates
Add a certificate template to a certification authority
Working with MMC console files

# Install a stand-alone root certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a stand-alone root certification authority**

1. Log on to the system as an Administrator, or if you have the Active Directory directory service, log on to the system as a Domain Administrator.

2. Click **Start**, point to **Settings**, and then click **Control Panel**.

3. Double-click **Add or Remove Programs** and then click **Add/Remove Windows Components**.

4. In the Windows Components Wizard, select the **Certificate Services** check box. A dialog box appears to inform you that the computer cannot be renamed and that the computer cannot be joined to or removed from a domain after Certificate Services is installed. Click **Yes**, and then click **Next**.

5. Click **Stand-alone root CA**.

6. (Optional) Select the **Use custom settings to generate the key pair and CA certificate** check box, and then click **Next** to specify the following.

| To set this | Do this |
|---|---|
| Cryptographic service provider (CSP) | In **CSP**, click the CSP that you want to use. The default is the Microsoft Strong Cryptographic Provider. Certificate Services does support third party CSPs, but you must refer to the CSP vendor's documentation for information about using their CSP with Certificate Services. |
| Hash algorithm | In **Hash algorithm**, click the hash algorithm you want to use. The default is SHA-1. |
| Use an existing key | Select the **Use existing key** check box, click **Import**, and then, in **Open PFX File**, type the file name and password of the public and private key pair. This is helpful if you are relocating or restoring a previously installed certification authority (CA). Note that, when using an existing key, a new certificate is generated.<br>**Important**<br>• Be sure that you select an existing key that you know to be uncompromised and trustworthy. Using a key that may be compromised or untrusted could cause this CA and all its issued certificates to be insecure. |
| Key length | In **Key length**, type or select a key length. The default key length using the Microsoft Strong Cryptographic Provider is 2048 bits. Default key lengths for other CSPs vary. In general, the longer the key length, the more secure the key is. Also, longer key lengths require more system resources for operations such as signing, encryption, and chain verification. For a root CA, you should use a key length of at least 2048 bits. This option is not available if you use existing keys. |
| Allow this CSP to interact with the desktop | Select the **Allow this CSP to interact with the desktop** check box. Without this option, system services cannot interact with the desktop of the user who is currently logged on. |
| Import | Click **Import**. This imports an existing key in the PKCS #12 PFX format. |
| View certificate | Click **View certificate**. This allows you to view the certificate that you select or generate during installation. |

When you are done, click **Next**.

7. Type the common name of the certification authority. None of this information can be changed after the CA setup is complete.

8. In **Validity period**, specify the validity duration for the root CA. See the note below about considerations when setting this value. Click **Next**.

9. Specify the storage locations of the certificate database, the certificate database log, and the shared folder. Click **Next**.

10. If Internet Information Services (IIS) is running, you will receive a request to stop the service before proceeding with the installation. Click **OK**.

11. If prompted, type the path to the Certificate Services installation files.

### Notes

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

- If Active Directory is available and you have Write permission to Active Directory, then specifying the shared folder is optional.

- If Active Server Pages are not enabled through Internet Information Services, you will be prompted to activate them. The Web interface for the certification authority requires running Active Server Pages.

- The validity duration you choose for the CA will determine when the CA "expires." For information about renewing CAs, see Related Topics.

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

- For more information on any installation settings, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Stand-alone certification authorities
Installing and configuring a certification authority
Install a stand-alone subordinate certification authority
Renewing a certification authority

# Install a stand-alone subordinate certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a stand-alone subordinate certification authority**

1. Log on to the system as an Administrator, or if you have the Active Directory directory service, log on to the system as a Domain Administrator.

2. Open Add or Remove Programs in Control Panel.

3. Click **Add/Remove Windows Components**.

4. In the Windows Components Wizard, select the **Certificate Services** check box. A dialog box appears to inform you that the computer cannot be renamed and that the computer cannot be joined to or removed from a domain after Certificate Services is installed. Click **Yes**, and then click **Next**.

5. Click **Stand-alone subordinate CA**.

6. (Optional) Select the **Use custom settings to generate the key pair and CA certificate** check box, and then click **Next** to specify the following.

| To set this | Do this |
|---|---|
| Cryptographic service provider (CSP) | In **CSP**, click the CSP that you want to use. The default is the Microsoft Strong Cryptographic Provider. Certificate Services does support third party CSPs, but you must refer to the CSP vendor's documentation for information about using their CSP with Certificate Services. |
| Hash algorithm | In **Hash algorithm**, click the hash algorithm you want to use. The default is SHA-1. |
| Use an existing key | Select the **Use existing key** check box, click **Import**, and then, in **Open PFX File**, type the file name and password of the public and private key pair. This is helpful if you are relocating or restoring a previously installed certification authority (CA). Note that, when using an existing key, a new certificate is generated.<br>**Important**<br>• Be sure that you select an existing key that you know to be uncompromised and trustworthy. Using a key that may be compromised or untrusted could cause this CA and all its issued certificates to be insecure. |
| Key length | In **Key length**, type or select a key length. The default key length using the Microsoft Strong Cryptographic Provider is 2048 bits. Default key lengths for other CSPs vary. In general, the longer the key length, the more secure the key is. Also, longer key lengths require more system resources for operations such as signing, encryption, and chain verification. For a root CA, you should use a key length of at least 2048 bits. This option is not available if you use existing keys. |
| Allow this CSP to interact with the desktop | Select the **Allow this CSP to interact with the desktop** check box. Without this option, system services cannot interact with the desktop of the user who is currently logged on. |
| Import | Click **Import**. This imports an existing key in the PKCS #12 PFX format. |
| View certificate | Click **View certificate**. This allows you to view the certificate that you select or generate during installation. |

When you are done, click **Next**.

7. Type in the common name (CN) and distinguished name suffix of the CA. None of this information can be changed after the CA setup is complete. Click **Next**.

8. Specify the storage locations of the certificate database, the certificate database log, and the shared folder. Click **Next**.

9. Obtain the certificate for the subordinate CA. For instructions on how to do this, see Notes.

10. If Internet Information Services is running, the system will request that you stop the service before proceeding with the installation. Click **OK**.

11. If prompted, type the path to the Certificate Services installation files.

**Notes**

- If Active Server Pages (ASP) are not enabled through Internet Information Services, you will be prompted to activate them. The Web interface for the certification authority requires running Active Server Pages.

- To obtain the certificate for a subordinate CA, you must submit a certificate request to a parent CA. The procedure for doing so differs depending on whether the parent CA is available online.

    - If a parent CA is available online:

Click **Send the request directly to a CA already on the network**.

In **Computer Name**, type the name of the computer on which the parent CA is installed.

In **Parent CA**, click the name of the parent CA.

- If a parent CA is not available online:

Click **Save the request to a file**.

In **Request file**, type the path and file name of the file that will store the request.

Obtain this subordinate CA's certificate from the parent CA.

The procedure for doing this will be unique to the parent CA. At a minimum, the parent CA should provide a file containing the subordinate CA's newly-issued certificate and, preferably, its full certification path. For the procedure to submit a certificate request using a file to a Microsoft CA, see Related Topics.

If you get a subordinate CA certificate that does *not* include the full certification path, the new subordinate CA you are installing must be able to build a valid CA chain when it starts. Because of this, you must install the parent CA's certificate in the Intermediate Certification Authorities certificate store of the computer (if the parent CA is not a root CA), as well as the certificates of any other intermediate CA in the chain, and you must install the certificate of the root CA in the chain into the Trusted Root Certification Authorities store. These certificates should be installed in the certificate store before you install the CA certificate on the subordinate CA you have just set up.

Open Certification Authority.

In the console tree, click the name of the CA.

**Where?**

Certification Authority (*Computer*)/*CA name*

On the **Action** menu, point to **All Tasks**, and then click **Install CA Certificate**.

Locate the certificate file received from the parent certification authority, click this file, and then click **Open**.

If the root CA is untrusted, click **OK** to trust the root CA's certificate.

Once the certificate is installed, start the Certificate Services service.

- If Active Directory is available and you have Write permission to Active Directory, then specifying the shared folder is optional.

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Stand-alone certification authorities
Installing and configuring a certification authority
Install a stand-alone root certification authority
Request a certificate using a PKCS #10 or PKCS #7 file
Working with MMC console files
Request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file

# Uninstall a certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To uninstall a certification authority**

1. Log on to the system as the user who installed the certification authority.
2. Open Add or Remove Programs in Control Panel.
3. Click **Add/Remove Windows Components**.
4. In the Windows Components Wizard, clear the **Certificate Services** check box, and then click **Next**.
5. If Internet Information Services is running, the system requests that you stop the service before proceeding with the uninstall process. If this happens, click **OK**.

## Notes

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- You should back up the entire server before uninstalling the certification authority (CA).
- When you uninstall a CA, the following information is left on the server:
  - The CA database
  - The CA public and private keys
  - The CA's certificates in the Personal store
  - The CA's certificates in the shared folder, if a shared folder was specified during Certificate Services setup
  - The CA chain's root certificate in the Trusted Root Certification Authorities store
  - The CA chain's intermediate certificates in the Intermediate Certification Authorities store
  - The CA's certificate revocation list (CRL)

  This information is kept on the server by default, in case you are uninstalling and then reinstalling the CA. For example, you might uninstall and reinstall if you wanted to change a stand-alone CA to an enterprise CA.
- If you installed the enterprise or enterprise subordinate certification authority as an Enterprise Admin or delegated user, then you must use the Enterprise Admin or delegated user account when you uninstall the certification authority.
- If you are permanently decommissioning the CA before its expected expiration date, then the CA certificate should be revoked from its parent CA for a certificate revocation reason of "Cease of operation". If the CA is a self-signed root CA, then all of the certificates that have not expired should be revoked and a CRL generated with the same reason. This will indicate that the certificates are no longer valid because the CA has been decommissioned.
- When uninstalling an enterprise CA, it is important that it be uninstalled properly to ensure that their CA enrollment object is removed from the Active Directory directory service. Failure to do so may result in Active Directory clients continuing to attempt to enroll against that CA. If an enterprise CA cannot be uninstalled normally, Active Directory tools

should be used to manually remove the CA objects from Active Directory.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing and configuring a certification authority

# Set up certification authority Web enrollment support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up certification authority Web enrollment support**

1. Log on to the system as an Administrator, or if you have the Active Directory directory service, log on to the system as a Domain Administrator.

2. Open Add or Remove Programs in Control Panel.

3. Click **Add/Remove Windows Components**.

4. In the Windows Components Wizard, select the **Certificate Services** check box. A dialog box appears to inform you that the computer cannot be renamed and that the computer cannot be joined to or removed from a domain after Certificate Services is installed. Click **Yes**, and then click **Details**.

5. Clear the **Certificate Services CA** check box, verify that the **Certificate Services Web Enrollment Support** check box remains selected, click **OK**, and then click **Next**.

6. In **Computer Name**, type in the name of the computer on which the certification authority (CA) is installed, the CA for which these Web enrollment pages will be used. The name of the CA will appear in the list. Click **Next**.

7. If Internet Information Services is running, the system will request that you stop the service before proceeding with the installation. If this happens, click **Yes**.

8. If prompted, type in the path to the Certificate Services installation files.

**Notes**

- If you installed the CA Web enrollment pages on a computer that is not a CA, the computer account must be trusted for delegation in Active Directory. For more information on trusting a computer for delegation, see Related Topics.

- CA Web enrollment pages are installed by default when you set up Certificate Services on a server that has Internet Information Services (IIS) installed. This procedure is necessary only if you want to install the CA Web enrollment pages on a server that is separate from the server on which the CA is installed.

- If you installed the CA Web enrollment pages before installing IIS, the required virtual roots are not created. You can create the virtual roots after installing IIS by typing **certutil -vroot** at a command prompt. The command **certutil -vroot** does not install the Web enrollment pages. It creates the IIS virtual roots which point to the Web enrollment pages, CA certificate, CRLs, and enrollment controls (xenroll.dll and scrdenrl.dll).

  For more information on creating or deleting the related set of Certificate Services Web server virtual roots and file shares, see Certutil tasks for configuring a Certification Authority (CA).

- When you initially install IIS, the service is installed in a highly secure mode. Because IIS only serves static content by default, you must enable features such as ASP. To enable ASP and allow the CA Web enrollment pages to work correctly, see Enabling and Disabling Dynamic Content in Microsoft Internet Information Services.

- Installation of the Web enrollment pages configures the computer as a registration authority (RA). This computer is also known as a "CA Web proxy" or a "Web enrollment station."

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
User and computer accounts
Certutil tasks for configuring a Certification Authority (CA)
Certification authority Web enrollment services
Installing and configuring a certification authority
Using Windows 2000 Certificate Services Web pages
Using Windows Server 2003 Certificate Services Web pages
Use Windows Server 2003 Certificate Services Web Pages

# Perform qualified subordination

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To perform qualified subordination**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Create the information file for the qualified subordinate certification authority.
2. Save the information file as CAPolicy.inf in the systemroot directory on the computer where you are installing the qualified subordinate certification authority.

   For more information, see Notes.
3. Install the qualified subordinate certification authority.

## Notes

- The information file should contain the certificate extensions needed to specify the qualified subordination constraints for the qualified subordinate certification authority you are installing.
- When you set up a qualified subordinate CA, you can add a CA policy statement to the CA certificate that is created during setup or a CA certificate renewal, in the form of a text file (CAPolicy.inf) or a pointer to a Web site. The CA policy can contain all the qualified subordination extensions and provide legal and other pertinent information about the CA and its issuing policies, limitations of liability, and so on. An end user will see this CA policy statement when they view the CA certificate and click Issuer Statement.

  For more information about CAPolicy.inf, see "Certificate Services" at the Microsoft Windows Resource Kits Web site.
- When qualified subordinate CA certificates and cross-certificates are revoked, they should also be deleted from the Active Directory directory service. This deletion is not required for security reasons, but is considered a best practice, as it also removes the certificates from the Local Machine certificates store of client computers.

**Using a command line**

1. Open Command Prompt.
2. Type the following to construct a qualified subordination request from an existing CA certificate or from an existing request:

   **certreq-policy**[*Options*] [*RequestFileIn* [*PolicyFileIn*[*RequestFileOut* [*PKCS10FileOut*]]]]
3. Type the following to sign the qualified subordination request:

   **certreq-sign-cert***CertId*[*RequestFileIn*[*RequestFileOut*]]
4. Type the following to submit the request to the certification authority:

   **certreq**[**-submit**] **-config***CAComputerName******CAName*[*RequestFileIn* [*CertFileOut*[*CertChainFileOut* [*FullResponseFileOut*]]]]

5. Install the enterprise qualified subordinate certification authority.

   For more information, see Related Topics.

| Value | Description |
|---|---|
| **certreq** | Requests certificates from a certification authority (CA) from the command prompt. |
| **-policy** | Required. Sets the policy for a request. |
| **-sign** | Required. Signs a cross-certification or qualified subordination request. |
| **-cert** *CertId* | Option. Specifies the signing certificate by common name, serial number, or by SHA-1 key or certificate hash. |
| **-attrib** *AttributeString* | Option. Specifies the Name and Value string pairs, separated by a colon. |
| **-binary** | Option. Formats output files as binary instead of Base64-encoded. |
| **-config** *CAComputerName* **** *CAName* | Option. Processes the operation using the CA that is identified in the configuration string specified by *CAComputerName******CAName*.<br>**Notes**<br>• Using **-config -** processes the operation using the default CA.<br>• If no **-config** option is used, the **Select Certification Authority** dialog box displays a list of all available CAs. |
| **-crl** | Option. Includes certificate revocation lists (CRLs) in the Base64-encoded PKCS #7 file, *CertChainFileOut*, or to the Base64-encoded file, *RequestFileOut*. |
| **-rpc** | Option. Instructs Certificate Services to use the RPC server connection instead of DCOM. |
| *RequestFileIn* | Specifies the Base64-encoded or binary input file name. May be a PKCS #10 certificate request, PKCS #7 certificate renewal request, KeyGen tag format certificate request, or a CMC request. |
| *CertFileOut* | Specifies to send output to either a binary or a Base64-encoded X.509 file. |
| *CertChainFileOut* | Specifies to send output to either a binary or a Base64-encoded PKCS #7 file. |
| *FullResponseFileOut* | Specifies to send output to either a binary or a Base64-encoded Full Response file. For **-submit**, this is the name of the qualified subordinate certificate file (for example, *QualifiedSub.cer*). |
| *PolicyFileIn* | Requires Certificate Services to get input from an .inf file containing a textual representation of extensions that are used to qualify a request. |
| *RequestFileOut* | Requires Certificate Services to send output to a Base64-encoded file. *RequestFileOut* for Step 2 will become *RequestFileIn* for Step 3. |
| *PKCS10FileOut* | Requires Certificate Services to send output to a Base64-encoded PKCS #10 file. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at a command prompt, type one of the following:
  - **certreq -?**

- **certreq -policy -?**
- **certreq -sign -?**
- **certreq -submit -?**

- Signing the qualified subordination request may require Enterprise Administrator credentials. This is a best practice for issuing signing certificates for qualified subordination.
- The certificate used to sign the qualified subordination request is created using the qualified subordination template. Enterprise Admins will have to sign the request or grant user permissions for the individuals that will sign the certificate.
- In Step 3, if you omit **-cert** *CertId*, you are prompted to supply the name of the signing certificate.
- When you install a qualified subordinate CA using the **certreq** command-line tool, the information (.inf) file can have any name. This is different from when you are using an information file as an issuer policy statement. In that case, the information file must be named *CAPolicy.inf*, as described in the Windows interface procedure above.
- When you sign the CMC request, you may need to have multiple personnel sign this request, depending on the assurance level that is associated with the qualified subordination.
- In some cases, if you omit values at the command line, you will be prompted to supply the necessary information.
- Once the qualified subordination request is signed, it is automatically published to Active Directory.
- If the parent CA of the qualified subordinate CA you are installing is offline, you must obtain the CA certificate for the qualified subordinate CA from the offline parent. If the parent CA is online, specify the CA certificate for the qualified subordinate CA during the Certificate Services Installation Wizard.
- When qualified subordinate CA certificates and cross-certificates are revoked, they should also be deleted from Active Directory. This deletion is not required for security reasons, but is considered a best practice as it also removes the certificates from the Local Machine certificates store of client computers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing and configuring a certification authority
Install an enterprise subordinate certification authority
Install a stand-alone subordinate certification authority
Command-line reference A-Z
Qualified subordination
Qualified subordination overview

# Administer a Certification Authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer a certification authority**

- Start or stop the certification authority service
- Set security permissions and delegate control of a certification authority
- View the certification authority certificate
- Back up a certification authority
- Restore a certification authority from a backup copy
- Set security for access to certification authority Web pages
- Configure certificate manager restrictions
- Configure event auditing
- Renew a root certification authority
- Renew a subordinate certification authority
- Publish certificates in a foreign Active Directory forest
- Send e-mail when a certification event occurs
- Use the Certification Authority snap-in
- Manage Certificate Revocation
- Manage Certificate Requests on a Stand-alone Certification Authority
- Manage Certificate Templates for an Enterprise Certification Authority
- Manage Key Archival and Recovery
- Change Policy Settings of a Certification Authority
- Change the Policy or Exit Modules of a Certification Authority
- Manage Role-based Administration

# Start or stop the certification authority service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start or stop the certification authority service**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, point to **All Tasks**, and click **Start Service** to start the service or click **Stop Service** to stop the service.

**Notes**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- You can also start or stop Certificate Services if you open Computer Management, double-click **Services and Applications**, and then click **Services**. In the details pane, click **Certificate Services**. On the **Action** menu, click **Stop** or **Start**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **net start certsvc** to start the service or type **net stop certsvc** to stop the service.

| Value | Description |
|---|---|
| start | Start the specified service. |
| stop | Stop the specified service |
| certsvc | Specifies the Certification Authority service as the service to act on. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **net /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Role-based administration
Working with MMC console files
Administer the certification authority on another computer
Administer the certification authority on this computer

# Set security permissions and delegate control of a certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set security permissions and delegate control of a certification authority**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. Click the **Security** tab and specify the security permissions.

   For general information about setting access control security permissions, see Related Topics.

**Note**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Access control overview
Role-based administration
Using role separation

# View the certification authority certificate

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the certification authority certificate**

1. Log on to the system as any CA role holder.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **General** tab, click **View Certificate**.

**Note**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files

# Back up a certification authority

Updated: January 19, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up a certification authority**

**Warning**

The procedures discussed in this article demonstrate how to back up only the certification authority (CA) database and private key. If you want to back up the CA database, private key, and CA configuration, see the backup procedures in AD CS Migration: Migrating the Certification Authority (http://technet.microsoft.com/library/ee126140.aspx). To back up all the configuration data for a computer, run a System State Data backup (http://technet.microsoft.com/library/cc781353.aspx).

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Log on to the system as a Backup Operator or a Certification Authority Administrator.
2. Open the Certification Authority console.
3. In the console tree, click the name of the certification authority (CA).
   - Certification Authority (*Computer*)/*CA name*
4. On the **Action** menu, point to **All Tasks**, and then click **Back Up CA**.
5. Follow the instructions in the Certification Authority Backup Wizard.

**Note**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- This procedure is useful if you want to back up a CA database and private key, without backing up the entire server or CA configuration on which the CA is installed.

  In general, you should use Backup to back up and restore both the CA and the server. For more information about backing up a computer running a Windows Server 2003 operating system, see Related Topics.
- By default, members of the local Administrators group of CA computer can also back up the CA.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **certutil** -backup <*BackupFolder*>

| Value | Description |
| --- | --- |
| backup | Specifies a full backup of the CA's database. |
| *BackupDirectory* | Specifies the path to use to store the backup data. |

**Tip**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at a command prompt, type:

  **certutil -backup -?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Role-based administration
Working with MMC console files
Restore a certification authority from a backup copy
Using Network Monitor
Backup
Checklists for Backing Up and Restoring Data
Certutil tasks for backing up and restoring certificates

# Restore a certification authority from a backup copy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore a certification authority from a backup copy**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Log on to the system as a Backup Operator or a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, point to **All Tasks**, and click **Restore CA**.
5. Follow the instructions in the Certification Authority Restore Wizard.

**Notes**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- This procedure is useful if you want to restore a CA without restoring the entire server on which the CA is installed. To back up the CA for later restoration using this procedure, see Related Topics.

  In general, you should use Backup to back up and restore both the CA and the server. For more information about backing up a computer running a Windows Server 2003 operating system, see Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **certutil** -restore *BackupDirectory*

| Value | Description |
|---|---|
| restore | Specifies a full restore of the CA's data. |
| *BackupDirectory* | Specifies the path where the backup data is located. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command**

**prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **certutil -restore -?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Role-based administration
Working with MMC console files
Back up a certification authority
Using Network Monitor
Backup
Checklists for Backing Up and Restoring Data
Certutil tasks for backing up and restoring certificates

# Set security for access to certification authority Web pages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set security for access to certification authority Web pages**

1. Log on to the system as an Internet Information Services Administrator.
2. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Internet Information Services**.
3. In the console tree, right-click **CertSrv**, and click **Properties**.

   **Where?**

   - Internet Information Services/*Computer Name*/Web Sites/Default Web Site/CertSrv

4. On the **Directory Security** tab, under **Anonymous access and authentication control**, click **Edit**.
5. Clear all check boxes *except***Integrated Windows authentication**.

### Important

- An enterprise certification authority (CA) requires that the certificate requester be authenticated by the page so that it can determine the correct information to put in the certificate. If you don't have authentication set for the Web pages in an enterprise CA, then the pages will fail to generate a certificate or, if a certificate is generated, it will be useless. For this reason, integrated Windows authentication is set by default on enterprise CAs. This procedure is provided so that you can confirm the default setting or fix an erroneous change to the IIS Directory Security settings for an enterprise certification authority.

### Notes

- You do not need to perform this procedure on a stand-alone certification authority.
- If you cannot locate CertSrv in the console tree, confirm that Certificate Services is installed.
- If Certificate Services is installed and the CertSrv virtual directory doesn't exist, run **certutil -vroot** from the command prompt to create it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Certification authority Web enrollment services

# Configure certificate manager restrictions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure certificate manager restrictions**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Certificate Manager Restrictions** tab, click **Restrict certificate managers**.
6. In **Available certificate managers**, select the certificate manager subject or group you want.
7. Click **Add** to add groups, users or computers to manage for the selected certificate manager.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- If a CA is running on a member server and the Restrict certificate managers property is enabled, then the member server needs to be added to the Pre-Windows 2000 Compatible Access built-in group of every domain from which it will receive certificate requests. Once added to these groups, the administrator of the CA is allowed to issue a certificate for subjects in those domains.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Role-based administration
Certificate Services example implementation: Key archival and recovery
Key archival and recovery
Establishing key options and key archival

# Configure event auditing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure event auditing**

1. Log on to the system as a Certification Authority Administrator or a CA Auditor.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Auditing** tab, click the events you want to audit.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- To audit events, the computer must also be configured for auditing of object access. Also, the auditor must perform this operation if the certification authority has been configured to enforce role-based administration. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Role-based administration
Working with MMC console files
Add a certification authority auditor
Audit object access
View the security log

# Renew a root certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To renew a root certification authority**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, point to **All Tasks**, and click **Renew CA Certificate**.
5. Do one of the following:

   - If you want to generate a new public and private key pair for the certification authority's certificate, click **Yes**.
   - If you want to reuse the current public and private key pair for the certification authority's certificate, click **No**.

**Note**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Renewing a certification authority
Renew a subordinate certification authority

# Renew a subordinate certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To renew a subordinate certification authority**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, point to **All Tasks**, and click **Renew CA Certificate**.
5. Do one of the following:

   - If you want to generate a new public and private key pair for the certification authority's certificate, click **Yes**.
   - If you want to reuse the current public and private key pair for the certification authority's certificate, click **No**.

6. Get the CA certificate from the parent CA. For more information, see Notes.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- To obtain the certificate for a subordinate CA, you must submit a certificate request to a parent CA. The procedure for doing so differs depending on whether the parent CA is available online.

  - If a parent CA is available online
  - If a parent CA is not available online

**If a parent CA is available online**

1. Click **Send the request directly to a CA already on the network**.
2. In **Computer Name**, type the name of the computer on which the parent CA is installed.
3. In **Parent CA**, click the name of the parent CA.

#### If a parent CA is not available online

1. Click **Save the request to a file**.

2. In **Request file**, type the path and file name of the file that will store the request.

3. Obtain this subordinate CA's certificate from the parent CA.

   The procedure for doing this will be unique to the parent CA. At a minimum, the parent CA should provide a file containing the subordinate CA's newly issued certificate and, preferably, its full certification path. For the procedure to submit a certificate request using a file to a Microsoft CA, see Related Topics.

   If you get a subordinate CA certificate that does *not* include the full certification path, the new subordinate CA you are installing must be able to build a valid CA chain when it starts. Thus you must install the parent CA's certificate in the Intermediate Certification Authorities certificate store of the computer (if the parent CA is not a root CA), as well as the certificates of any other intermediate CA in the chain, and you must install the certificate of the root CA in the chain into the Trusted Root Certification Authorities store. These certificates should be installed in the certificate store before you install the CA certificate on the subordinate CA you have just set up.

4. Open Certification Authority.

5. In the console tree, click the name of the CA.

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

6. On the **Action** menu, point to **All Tasks**, and then click **Install CA Certificate**.

7. Locate the certificate file received from the parent certification authority, click this file, and then click **Open**.

#### Using a command line

1. Open Command Prompt.

2. Type:

   **certutil** -renewcert

| Value | Description |
|---|---|
| renewcert | Instructs the CA to renew its certificate. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **certutil -renewcert -?**

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files
Renewing a certification authority
Renew a root certification authority
Request a certificate using a PKCS #10 or PKCS #7 file

# Publish certificates in a foreign Active Directory forest

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To publish certificates in a foreign Active Directory forest**

1. Log on to the system as an administrator.
2. Open Command Prompt.
3. Type:

   **certutil -setreg CA\AlternatePublishDomains +"***DomainName***"**

| Value | Description |
|---|---|
| **certutil** | Specifies the name of the command-line program. |
| **-setreg** | Modifies the registry. |
| **CA\AlternatePublishDomains** | Indicates the registry value that contains a list of Active Directory domains that receive certificate publication. |
| **+** | Indicates that if there are current entries stored in this registry value, append this entry to those. |
| *DomainName* | Specifies the fully-qualified DNS domain name of the foreign Active Directory domain to publish certificates in. |

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- A forest trust relationship must exist between the CA's forest and the forest that certificates will be published in before publication in the foreign domain. For more information, see Related Topics.
- The default permissions in a trust across forests do not allow certificate publication from a foreign forest. To grant the appropriate permissions, use the Delegation of Control Wizard on the foreign Active Directory's root domain to make the following settings:

| Setting | Value |
|---|---|
| Selected users and groups | **Certificate Publishers** group in the foreign Active Directory domain |
| Tasks to delegate | **Create a custom task to delegate** |

| | |
|---|---|
| Delegate control of: | Under **Only the following objects in the folder**, select **Contact objects** |
| Show these permissions | Click **Property-specific**, then select **Read userCertificate** and **Write userCertificate** |

For more information on running the Delegation of Control Wizard, see Related Topics.

- The default membership of the foreign Active Directory's Cert Publishers group does not include the source CA. The source CA's Active Directory computer account must be added to the foreign Cert Publishers group before publishing across forests will work. Because computer accounts cannot be found across trusts across forests by browsing, ensure you know the name of the CA to add.

- Publication of certificates in a foreign Active Directory requires the e-mail property in the source forest's User object match the e-mail property in the destination forest's Contact object. The e-mail property must be populated on User objects for this to happen. For more information on changing properties of a User object, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **certutil -setreg -?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start or stop the certification authority service
Forest trusts
Create a forest trust
**Other Resources**
Delegate control

# Send e-mail when a certification event occurs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send e-mail when a certification event occurs**

1. Open Command Prompt.
2. Type:

   **certutil -setreg exit\smtp\smtpserver***ServerName*

   and then type:

   **certutil -setreg exit\smtp\eventfilter +***Event*

| Value | Description |
|---|---|
| **certutil** | Specifies the name of the command-line program. |
| **-setreg** | Modifies the registry. |
| **exit\smtp\smtpserver** | Indicates the registry value that contains the name of the Simple Mail Transfer Protocol (SMTP) server. |
| **exit\smtp\eventfilter** | Indicates the registry value that contains the list of events that the certification authority (CA) should monitor. When any of these events occur, the CA will send e-mail. |
| **+** | Indicates that, if there are current entries stored in this registry value, this entry should be appended to them. |
| *Event* | Specifies the event to add to the list of events for the CA to monitor. An event can be any value in the following table: |

| Event value | Description |
|---|---|
| ExitEvent_CertIssued | Specifies the action of issuing a certificate. |
| ExitEvent_CertPending | Specifies the action of a certificate request being received by the CA and set to Pending. |
| ExitEvent_CertDenied | Specifies the action of a certificate request being received by the CA and that request being denied. |
| ExitEvent_CertRevoked | Specifies the action of a revocation of an existing certificate. |
| ExitEvent_CRLIssued | Specifies the action of a certificate revocation list being issued. |
| ExitEvent_Startup | Specifies the action of the certification authority starting. |
| ExitEvent_Shutdown | Specifies the action of the certification authority shutting down. |

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- When the ExitEvent_CRLIssued, ExitEvent_Startup and ExitEvent_Shutdown events occur, the CA does not know where to send the e-mail, as there is no user associated with this event. Therefore, an e-mail address must be configured when using these events. To configure the e-mail address to send e-mail to when these events occur, type the following certutil.exe commands at a command prompt:

  **certutil -setreg exit\smtp\CRLIssued\To***E-mailString*

  **certutil -setreg exit\smtp\Startup\To***E-mailString*

  **certutil -setreg exit\smtp\Shutdown\To***E-mailString*

  *E-mailString* specifies an e-mail address or a string of e-mail addresses that are separated by semicolons.

- If the SMTP server is not set to accept anonymous connections, the CA must be configured to provide a user name and password when it connects. To configure the CA to authenticate with the SMTP server, type the following **certutil.exe** commands at a command prompt:

  **certutil -setreg exit\smtp\SMTPAuthenticate 1**

  **certutil -setsmtpinfo***UserName*

  *UserName* specifies the user name of a valid account on the SMTP server. Certutil will prompt you to provide the password for this user name.

- To view the complete syntax for this command, at a command prompt, type:

  **certutil -setreg -?**

  For more information about configuring a certification authority for e-mail, see "Designing a Public Key Infrastructure" at the Microsoft Windows Resource Kits Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start or stop the certification authority service

# Use the Certification Authority snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use the Certification Authority snap-in**

- Administer the certification authority on this computer
- Administer the certification authority on another computer
- Set display filters for the Certification Authority snap-in
- Customize the display of columns in Certification Authority

# Administer the certification authority on this computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To administer the certification authority on this computer**

1. Log on to the system as a Certification Authority Administrator.

2. Open Certification Authority.

**Note**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Save an MMC console file
Administer the certification authority on another computer

# Administer the certification authority on another computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To administer the certification authority on another computer**

1. Log on to the system as a Certification Authority Administrator of the remote CA.
2. Open Certification Authority.
3. On the **Action** menu, click **Retarget Certification Authority**.
4. Click **Another computer** and type the name of the computer.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- As an alternate way to manage a CA on another computer, perform the following steps:

  1. Log on to the system as an Administrator.
  2. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
  3. On the **File** menu, click **Add/Remove Snap-in**. Click **Add**.
  4. In **Snap-in**, double-click **Certification Authority**.
  5. Click **Another computer** and type the name of the computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Save an MMC console file
Working with MMC console files
Administer the certification authority on this computer

# Set display filters for the Certification Authority snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set display filters for the Certification Authority snap-in**

1. Log on to the system as any CA role holder.
2. Open Certification Authority.
3. In the console tree, double-click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. Click any of the displayed folders except **Certificate templates**.
5. On the **View** menu, click **Filter**.
6. For each of the selection criteria:

   - Click **Add**.
   - In **Field**, click the field on which to filter.
   - In **Operation**, click the operation to qualify the filter value for this field.
   - In **Value**, type the qualification value.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- Display filters are essentially a way to make simple queries against a CA's database. By using filters, you can restrict the items displayed in the details pane of the Certification Authority snap-in to items that meet a set of criteria you establish. For example, you can create a filter that will display in the **Issued Certificates** folder only those certificates that were effective after a specific date.
- To remove a filter, click it in the **Filter** dialog box, and then click **Remove**.
- To remove all existing filters, in the **Filter** dialog box, click **Reset**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files

Customize the display of columns in Certification Authority

# Customize the display of columns in Certification Authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To customize the display of columns in Certification Authority**

1. Log on to the system as any CA role holder.
2. Open Certification Authority.
3. In the console tree, double-click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. Click any of the displayed folders, such as **Revoked Certificates** or **Issued Certificates**.
5. On the **View** menu, click **Add/Remove Columns**.
6. Do any of the following until you have made all of the changes that you want.

| To | Do this |
|---|---|
| Add a new column to the details pane | In **Available Columns**, click the column you want to add to the details pane. Click **Add**. |
| Remove a column from the details pane | In **Displayed Columns**, click the column you want to remove from the details pane. Click **Remove**. |
| Change the order of columns in the details pane | In **Displayed Columns**, click the column you want to move up or down in the order of the details pane. Click **Move Up** or **Move Down** as appropriate. |
| Restore the default columns | Click **Restore Defaults**. |

**Notes**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- To sort the results, click the column heading in the details pane. You can only sort results based on columns (indicating the fields) that are indexed in the database of the certification authority. There will be a small arrow on the column heading you are sorting by.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Set display filters for the Certification Authority snap-in

# Manage Certificate Revocation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage certificate revocation**

- Specify certificate revocation list distribution points in issued certificates
- Configure CRL and delta CRL overlap period
- Specify online certificate status protocol responder in issued certificates
- Configure an offline root certification authority to support certificate revocation with Active Directory
- Revoke an issued certificate
- Schedule the publication of the certificate revocation list
- Schedule the publication of the delta certificate revocation list
- Manually publish the certificate revocation list
- View the certificate revocation list

# Specify certificate revocation list distribution points in issued certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can add, remove, or modify certificate revocation list distribution points (CDPs) in issued certificates by using the following procedure. However, modifying the URL for a CDP only affects newly issued certificates. Previously issued certificates will continue to reference the original location.

**To specify certificate revocation list distribution points in issued certificates**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the certification authority.

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Extensions** tab, confirm that **Select extension** is set to **CRL Distribution Point (CDP)**.
6. Do one or more of the following. (The list of CRL distribution points is in the **Specify locations from which users can obtain a certificate revocation list (CRL)** box.)

| To | Do this |
|---|---|
| Add a new certificate revocation list (CRL) distribution point. | Click **Add**, type the name of the new CRL distribution point, and click **OK**. |
| Remove a CRL distribution point from the list. | Click the CRL distribution point, and then click **Remove** and click **OK**. |
| Indicate that you want to use a URL as a CRL distribution point. | Click the CRL distribution point, select the **Include in the CDP extension of issued certificates** check box, and then click **OK**. |
| Indicate that you do not want to use a URL as a CRL distribution point. | Click the CRL distribution point, clear the **Include in the CDP extension of issued certificates** check box, and then click **OK**. |
| Indicate that you want to use a URL as a delta CRL distribution point. | Click the CRL distribution point, select the **Publish Delta CRLs to this location** check box, and then click **OK**. |
| Indicate that you do not want to use a URL as a delta CRL distribution point. | Click the CRL distribution point, clear the **Publish Delta CRLs to this location** check box, and then click **OK**. |
| Indicate that you want to publish this location in CRLs to point clients to a delta CRL. | Click the CRL distribution point, select the **Include in CRLs. Clients use this to find Delta CRL locations.** check box, and then click **OK**. |
| Indicate that you do not want to publish this location in CRLs to point clients to a Delta CRL. | Click the CRL distribution point, clear the **Include in CRLs. Clients use this to find Delta CRL locations.** check box, and then click **OK**. |

7. Click **Yes** to stop and restart the Certificate Services service.

**Notes**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

- Certificate revocation list URLs can be either HTTP, FTP, LDAP, or FILE addresses. You can use the following variables when specifying the address of the CRL.

| Variable | Value |
|---|---|
| CAName | The name of the certification authority |
| CAObjectClass | The object class identifier for a certification authority, used when publishing to an LDAP URL |
| CATruncatedName | The "sanitized" name of the certification authority, truncated to 32 characters with a hash on the end |
| CDPObjectClass | The object class identifier for CRL distribution points, used when publishing to an LDAP URL |
| CertificateName | The renewal extension of the certification authority |
| ConfigurationContainer | The location of the Configuration container in Active Directory |
| CRLNameSuffix | Inserts a name suffix at the end of the file name when publishing a CRL to a file or URL location |
| DeltaCRLAllowed | When a delta CRL is published, this replaces the CRLNameSuffix with a separate suffix to distinguish the delta CRL |
| ServerDNSName | The DNS name of the certification authority server |
| ServerShortName | The NetBIOS name of the certification authority server |

- To stop and restart the Certificate Services service, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Certificate revocation
Revoke an issued certificate
Schedule the publication of the certificate revocation list
Manually publish the certificate revocation list
View the certificate revocation list

# Configure CRL and delta CRL overlap period

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure CRL and delta CRL overlap period**

1. Open Command Prompt.
2. Type:

   **certutil -setreg ca\CRLOverlapUnits***Value*

   **certutil -setreg ca\CRLOverlapPeriod***Units*

   **certutil -setreg ca\CRLDeltaOverlapUnits** *Value*

   **certutil -setreg ca\DeltaOverlapPeriod***Units*
3. Open Certification Authority.
4. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*
5. On the **Action** menu, point to **All Tasks**, and click **Stop Service** to stop the service.
6. On the **Action** menu, point to **All Tasks**, and click **Start Service** to start the service.

| Value | Description |
|---|---|
| **certutil** | Specifies the name of the command-line program. |
| **-setreg** | Modifies the registry. |
| **ca\CRLOverlapUnits** | Indicates the registry value that stores the value for the CRL overlap setting. |
| **ca\CRLDeltaOverlapUnits** | Indicates the registry value that stores the value for the delta CRL overlap setting. |
| *Value* | Provides the numerical value to set this option to. |
| **ca\CRLOverlapPeriod** | Indicates the registry value that stores the value for the CRL overlap unit type setting. |
| **ca\DeltaOverlapPeriod** | Indicates the registry value that stores the value for the delta CRL overlap unit type setting. |
| *Units* | Provides the type of units for the overlap period. Valid values are **Minutes** and **Hours**. |

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have

been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

- The maximum value for either the CRL or delta CRL overlap period is 12 hours.

- The overlap period for CRLs is the amount of time at the end of a published CRLs lifetime that a client can use to obtain a new CRL before the old CRL is considered unusable. The default setting for this value is 10% of the CRL lifetime. Because some environments may require longer periods to replicate a CRL, this setting can be configured manually.

- When both a base CRL and delta CRL have been recently published, a revoked certificate may appear in both. This is because the newer delta CRL may still point at the older base CRL while the new base CRL is being replicated. Having the certificate appear in both CRLs ensures the revocation information is available.

- If your environment is not configured to issue delta CRLs, the settings for **CRLDeltaOverlapUnits** and **DeltaOverlapPeriod** will have no effect.

- To view the complete syntax for this command, at a command prompt, type:

  **certutil -setreg -?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manually publish the certificate revocation list
Schedule the publication of the delta certificate revocation list
Revoking certificates and publishing CRLs
Start or stop the certification authority service

# Specify online certificate status protocol responder in issued certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify online certificate status protocol (OCSP) responder in issued certificates**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the certification authority.

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Extensions** tab, under **Select extension**, click **Authority Information Access (AIA)**.
6. Do one or more of the following. (The list of OCSP responders is in the **Specify locations from which users can obtain a certificate revocation list (CRL)** box.)

| To | Do this |
|---|---|
| Add a new OCSP responder. | Click **Add**, type in the URL of the new CRL distribution point, and then click **OK**. |
| Remove a OCSP responder from the list. | Click the OCSP responder, click **Remove**, and then click **OK**. |
| Indicate that you do not want to use a URL as an OCSP responder. | Click the OCSP responder, clear the **Include in the online certificate status protocol (OCSP) extension** check box, and then click **OK**. |
| Indicate that you want to use a URL as an OCSP responder. | Click the OCSP responder, select the **Include in the online certificate status protocol (OCSP) extension** check box, and then click **OK**. |

7. On the **Action** menu, point to **All Tasks**, click **Stop Service**, and on the **Action** menu, point to **All Tasks**, and then click **Start Service**.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- OCSP responder list URLs must be formatted as HTTP addresses.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files
Certificate revocation
Revoke an issued certificate
Schedule the publication of the certificate revocation list
Manually publish the certificate revocation list
View the certificate revocation list
Specify CA certificate access points in issued certificates

# Configure an offline root certification authority to support certificate revocation with Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure an offline root certification authority to support certificate revocation with Active Directory**

1. Log on to the system as a Certification Authority Administrator.
2. Open Command Prompt.
3. Type the following, and then press ENTER.

   **certutil -setreg ca\DSConfigDN "CN=Configuration,** *DNpath***"**

4. Type the following, and then press ENTER.

   **certutil -setreg ca\DSDomainDN "***DNpath***"**

5. Open Certification Authority.
6. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

7. On the **Action** menu, point to **All Tasks**, and click **Stop Service** to stop the service.
8. On the **Action** menu, point to **All Tasks**, and click **Start Service** to start the service.

| Value | Description |
|---|---|
| **certutil** | Specifies the name of the command-line program. |
| **-setreg** | Modifies the registry. |
| **ca\DSConfigDN** | Indicates the registry value for the distinguished name value of the path to the Configuration container in the Active Directory directory service. |
| **ca\DSDomainDN** | Indicates the registry value for the distinguished name value of the Active Directory domain. |
| *DNpath* | Specifies the distinguished name value in Active Directory. For more information, see Related Topics. |

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

- To ensure clients have the ability to locate the certificate revocation list (CRL), the certification authority (CA) must be configured with the appropriate LDAP path to retrieve this CRL from Active Directory. This is done by selecting the **Include in the CDP extension of issued certificates** check box when configuring the CDP. For more information, see Related Topics.

- This procedure should be used as part of a larger plan to deploy an offline root CA. For more information about deploying an offline root CA, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **certutil -setreg -?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Creating a certification hierarchy with an offline root certification authority
Specify certificate revocation list distribution points in issued certificates
Start or stop the certification authority service
Active Directory naming

# Revoke an issued certificate

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To revoke an issued certificate**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Log on to the system as a Certification Authority Administrator or Certificate Manager.
2. Open Certification Authority.
3. In the console tree, click **Issued Certificates**

   **Where?**

   - Certification Authority (*Computer*)/*CA name*/Issued Certificates

4. In the details pane, click the certificate you want to revoke.
5. On the **Action** menu, point to **All Tasks**, and click **Revoke Certificate**.
6. Select the reason for revoking the certificate and click **Yes**.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- The certificate is marked as revoked and is moved to the Revoked Certificates folder. The revoked certificate will appear on the certificate revocation list (CRL) the next time it is published.
- Certificates revoked with the reason code "Certificate Hold" can be unrevoked, left on "Certificate Hold" until they expire, or have their revocation reason code changed. "Certificate Hold" is the only revocation reason that will allow you to unrevoke the certificate. It is useful if the status of the certificate is questionable and is meant to provide some flexibility to the CA administrator.
- To unrevoke a certificate revoked with the reason code "Certificate Hold," at a command prompt on the CA, type:

  **certutil -revoke** *CertificateSerialNumber* **unrevoke**

- To identify the certificate serial number, in the Revoked Certificates folder, in the details pane, double-click the revoked certificate, and then click the **Details** tab.
- To change the reason code for a certificate previously revoked with the reason code "Certificate Hold," type the appropriate command at a command prompt on the CA.

| New reason code | Command |
|---|---|

| | |
|---|---|
| Unspecified | **certutil -revoke***CertificateSerialNumber***0** |
| Key Compromise | **certutil -revoke***CertificateSerialNumber***1** |
| CA Compromise | **certutil -revoke***CertificateSerialNumber***2** |
| Affiliation Changed | **certutil -revoke***CertificateSerialNumber***3** |
| Superseded | **certutil -revoke***CertificateSerialNumber***4** |
| Cessation of Operation | **certutil -revoke***CertificateSerialNumber***5** |

- You can also unrevoke a certificate in the Certification Authority by right-clicking the certificate you want, clicking **All Tasks**, then clicking **Unrevoke Certificate**. This certificate must be revoked for the reason of "Certificate Hold".

**Using a command line**

1. Open Command Prompt.

2. Type:

   **certutil** -revoke *SerialNumberReasonCode*

| Value | Description |
|---|---|
| revoke | Specifies the revocation of an existing certificate. |
| *SerialNumber* | Specifies the serial number of the certificate for revocation. |
| *ReasonCode* | Specifies the reason code for this certificate revocation. For values, see Notes. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **certutil -revoke -?**

- The valid reason codes for certificate revocation are:

| Reason forrevoking a certificate | Reason code |
|---|---|
| Unspecified | **0** |
| Key Compromise | **1** |
| CA Compromise | **2** |
| Affiliation Changed | **3** |
| Superseded | **4** |
| Cessation of Operation | **5** |
| Certificate Hold | **6** |

- Certificates that are revoked with the reason code "Certificate Hold" can be unrevoked, left on "Certificate Hold" until they expire, or have their revocation reason code changed. "Certificate Hold" is the only revocation reason that will allow you to unrevoke the certificate. It is useful if the status of the certificate is questionable and is meant to provide some

flexibility to the CA administrator.

- To unrevoke a certificate that is revoked with the reason code "Certificate Hold," at a command prompt on the CA, type:

  **certutil -revoke** *CertificateSerialNumber* **unrevoke**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Specify certificate revocation list distribution points in issued certificates
Certificate revocation
Schedule the publication of the certificate revocation list
Manually publish the certificate revocation list
View the certificate revocation list

# Schedule the publication of the certificate revocation list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To schedule the publication of the certificate revocation list**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click **Revoked Certificates**.

   **Where?**

   - Certification Authority (*Computer*)/*CA name/*Revoked Certificates

4. On the **Action** menu, click **Properties**.
5. In **CRL publication interval**, type the increment and click the unit of time to use for the automatic publishing of the certificate revocation list (CRL).

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- The certificate revocation list (CRL) is published by default in:

  *systemroot*\system32\CertSrv\CertEnroll\

- If the computer is a domain member and has permission to write to the Active Directory directory service, then the CRL is also published to Active Directory.
- The publishing period for a CRL is not the same as the validity period for a CRL. By default, the validity period of a CRL exceeds the publishing period of a CRL by 10% (up to a 12-hour maximum) to allow for directory replication. For more information about revoking certificates and publishing CRLs, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Revoking certificates and publishing CRLs
Certificate revocation
Specify certificate revocation list distribution points in issued certificates
Revoke an issued certificate
Manually publish the certificate revocation list
View the certificate revocation list
Schedule the publication of the delta certificate revocation list

# Schedule the publication of the delta certificate revocation list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To schedule the publication of the delta certificate revocation list**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click **Revoked Certificates**.

   **Where?**

   - Certification Authority (*Computer*)/*CA name*/Revoked Certificates

4. On the **Action** menu, click **Properties**.
5. Select the **Publish Delta CRLs** check box.
6. In **Publication interval**, type the increment and click the unit of time to use for the automatic publishing of the delta certificate revocation list (CRL).

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- The delta certificate revocation list (CRL) is published by default in:

  *Systemroot*\system32\CertSrv\CertEnroll\

  If the computer is a domain member and has permission to write to Active Directory, then the delta CRL is also published to Active Directory.

- The publishing period for a CRL is not the same as the validity period for a CRL. By default, the validity period of a CRL exceeds the publishing period of a CRL by 10% (up to a 12-hour maximum) to allow for directory replication. For more information about revoking certificates and publishing CRLs, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Revoking certificates and publishing CRLs
Certificate revocation
Specify certificate revocation list distribution points in issued certificates
Revoke an issued certificate
Manually publish the certificate revocation list
View the certificate revocation list
Schedule the publication of the certificate revocation list

# Manually publish the certificate revocation list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manually publish the certificate revocation list**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click **Revoked Certificates**.

   **Where?**

   - Certification Authority (*Computer*)/*CA name*/Revoked Certificates

4. On the **Action** menu, point to **All Tasks**, and click **Publish**.
5. Select **New CRL** to overwrite the previously-published certificate revocation list (CRL), or select **Delta CRL only** to publish a current delta CRL.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- Clients that have a cached copy of the previously-published CRL or delta CRL will continue using it until its validity period has expired, even though a new CRL has been published. Manually publishing a CRL does not affect cached copies of CRLs that are still valid; it only makes a new CRL available for systems that do not have a valid CRL.
- See Related Topics for the procedure clients can use to get the most recent CRL published by the certification authority (CA) even if they still have a valid CRL cached.
- By default, on the server on which the CA is installed, the CRL and delta CRL are published in:

  *Systemroot*\system32\CertSrv\CertEnroll\

- If the Active Directory directory service is available, they are also published to Active Directory.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **certutil** -crl

| Value | Description |
| --- | --- |
| crl | Specifies that a full certificate revocation list will be published. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **certutil -crl -?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure CRL and delta CRL overlap period
Working with MMC console files
Revoking certificates and publishing CRLs
Certificate revocation
Revoke an issued certificate
Specify certificate revocation list distribution points in issued certificates
Schedule the publication of the certificate revocation list
View the certificate revocation list
Retrieve a certificate revocation list

# View the certificate revocation list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the certificate revocation list**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click **Revoked Certificates**.

   **Where?**

   - Certification Authority (*Computer*)/*CA name*/Revoked Certificates

4. On the **Action** menu, click **Properties**.
5. Click the **View CRLs** tab.
6. Click **View CRL** to view the certificate revocation list or **View Delta CRL** to view the delta certificate revocation list.
7. Click the **Revocation List** tab.

**Note**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Revoking certificates and publishing CRLs
Certificate revocation
Revoke an issued certificate
Schedule the publication of the certificate revocation list
Manually publish the certificate revocation list

# Manage Certificate Requests on a Stand-alone Certification Authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage certificate requests on a stand-alone certification authority**

- Review pending certificate requests
- Set the default action upon receipt of a certificate request

# Review pending certificate requests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To review pending certificate requests**

1. Log on to the system as a Certification Authority Administrator or Certificate Manager.
2. Open Certification Authority.
3. In the console tree, click **Pending Requests**.

   **Where?**

   - Certification Authority (Computer)/CA name/Pending Requests

4. In the details pane, examine the certificate request by noting the values for requester name, requester e-mail address, and any other fields that you consider critical information for issuing the certificate.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- Failed certificate requests can also be issued using the same procedure in the **Failed Requests** container. However, not all of the security requirements for the certificate can be verified when issuing a failed request. Caution should be used.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Role-based administration
Working with MMC console files
Check on a pending certificate request
Set the default action upon receipt of a certificate request
Processing certificate requests

# Set the default action upon receipt of a certificate request

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the default action upon receipt of a certificate request**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Policy Module** tab, click **Properties**.
6. Click the option you want:

| To | Do this |
|---|---|
| Have the certification authority administrator review every certificate request before issuing a certificate. | Click **Set the certificate request status to pending**. |
| Have the certification authority issue certificates based on the configuration of the certificate template | Click **Follow the settings in the certificate template, if applicable. Otherwise, automatically issue the certificate.**. |

7. Stop and restart the Certificate Services service.

**Caution**

- In most cases, for security reasons, it is strongly recommended that all incoming certificate requests to a stand-alone CA be marked as "pending." Unlike enterprise certification authorities, stand-alone CAs do not use the Active Directory directory service, even if it is available, to verify that an individual or computer is authorized to be issued a certificate from the CA automatically. For stand-alone CA's, the CA administrator is responsible for verifying the identity of the certificate requestor.

**Notes**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- If you change the setting from **Set the certificate request status to pending** to **Always issue the certificate**, this will apply only to certificate requests submitted to the CA after the default action has been changed. If there are pending requests held by the CA, these requests will stay pending until the CA administrator issues the certificates or denies the requests.
- To review pending certificate requests, see Related Topics.
- To stop and restart the Certificate Services service, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Review pending certificate requests
Processing certificate requests
Start or stop the certification authority service
Stand-alone certification authorities

# Manage Certificate Templates for an Enterprise Certification Authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage certificate templates for an enterprise certification authority**


- Install new templates and upgrade existing templates
- Create a new certificate template
- Delete a certificate template
- Add a certificate template to a certification authority
- Allow subjects to request a certificate that is based on the template
- Re-enroll all certificate holders
- Remove a certificate template from a certification authority
- Modify a Certificate Template

# Install new templates and upgrade existing templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install new templates and upgrade existing templates**

1. Open Certificate Templates.

2. When prompted to install new certificate templates, click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.

- If this procedure has already been performed in your enterprise, you will not receive a prompt when you open Certificate Templates.

- This procedure must be performed before any Windows 2000 enterprise certification authorities are upgraded to Windows Server 2003 family certification authorities. If any Windows 2000 enterprise certification authorities currently exist or have ever existed in your enterprise, you must perform this procedure to ensure the certificate templates are added and updated correctly.

- This procedure must be performed on a certification authority (CA) in the certification hierarchy. Performing this procedure on a different computer may not work correctly.

- Performing this procedure is important to apply the proper security permissions on the existing certificate templates as well as create new certificate templates. If you do not perform this procedure before upgrading your certification authorities to the Windows Server 2003 family, you will be prompted when opening the Certification Authority snap-in.

- You must prepare the Active Directory schema with Adprep.exe for the upgrade before installing the new certificate templates. For more information on Adprep.exe, see Related Topics.

**See Also**
**Concepts**
Version 2 certificate templates
Upgrading from a Windows 2000 domain

# Create a new certificate template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new certificate template**

1. Open Certificate Templates.
2. Right-click the template to copy from and then click **Duplicate Template**.
3. Type a new name for this certificate template.
4. Make any necessary changes and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- New certificate templates must be created by copying an existing template and using the existing template's properties as the default for the new template. Copy the existing certificate template closest to the configuration of the intended new template to minimize the work necessary.
- Although Version 2 templates can be created and duplicated in Windows Server 2003, Standard Edition, certificates based on Version 2 templates can only be issued by a certification authority running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition.

**See Also**
**Concepts**
Delete a certificate template
Supersede templates
Modify a Certificate Template

# Delete a certificate template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a certificate template**

1. Open Certificate Templates
2. Right-click the template, and then click **Delete**.

### Caution

- When you delete a certificate template, certificates can no longer be issued based on it. If you are using Enterprise CAs, this will affect all CAs in the forest. Certificate templates cannot be recovered once they are deleted.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.

**See Also**
**Concepts**
Create a new certificate template
Remove a certificate template from a certification authority

# Add a certificate template to a certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a certificate template to a certification authority**

1. Open **Certification Authority**
2. Right-click the **Certificate Templates** container, click **New**, then click **Certificate Template to Issue**.
3. Select the certificate template and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press **Enter**.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.

**See Also**
**Concepts**
Create a new certificate template
Deploying Certificate Templates
Version 2 certificate templates

# Allow subjects to request a certificate that is based on the template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow subjects to request a certificate that is based on the template**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Security** tab, add the groups, computers, or users that you want.
4. In **Group or user names**, click one of the new objects, and then, on **Permissions for *ObjectName***, under the **Allow** column, select the **Read** and **Enroll** check boxes.
5. Repeat the previous step for each new object.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- The **Autoenroll** permission must also be set if the subject will be using client autoenrollment to obtain certificates. For more information on autoenrollment, see Related Topics.
- To disallow subjects from requesting a certificate based on a template, clear the **Read** and **Enroll** check boxes using the same steps as above.

**See Also**
**Concepts**
Add a certificate template to a certification authority
Re-enroll all certificate holders
Allowing for autoenrollment

# Re-enroll all certificate holders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To re-enroll all certificate holders**

1. Open Certificate Templates.
2. Right-click the template that you want to use and then click **Reenroll All Certificate Holders**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is used when a critical change is made to the certificate template and you want all subjects that hold a certificate that is based on this template to re-enroll as quickly as possible. The next time the subject verifies the version of the certificate against the version of the template on the certification authority, the subject will re-enroll.

**See Also**
**Concepts**
Modify a Certificate Template
Planning for certificate templates

# Remove a certificate template from a certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a certificate template from a certification authority**

1. Open Certification Authority
2. In the console tree, click **Certificate Templates**.
3. In the details pane, right-click the certificate template that you want, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certification Authority, click **Start**, click **Run**, type **certsrv.msc**, and then press Enter.

**See Also**
**Concepts**
Create a new certificate template
Deploying Certificate Templates
Delete a certificate template

# Modify a Certificate Template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modify a certificate template**

- Supersede templates
- Configure a certificate template for key archival and recovery
- Configure a certificate template for client autoenrollment
- Modify enrollment policy
- Create an application policy
- Create an issuance policy
- Add an application policy
- Add an issuance policy
- Make certificate or application policy critical
- Change key usage
- Change basic constraints

# Supersede templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To supersede templates**

1. Open Certificate Templates
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. Click the **Superseded Templates** tab.
4. Click **Add**.
5. Click one or more templates to supersede, then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- This procedure is used when a template replaces another template and you want subjects of certificates based on the old template to obtain new certificates based on the new template. It forces subjects to obtain a new certificate before their configured renewal period.

**See Also**
**Concepts**
Version 2 certificate templates
Re-enroll all certificate holders
Deploying Certificate Templates

# Configure a certificate template for key archival and recovery

Updated: November 7, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a certificate template for key archival and recovery**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Request Handling** tab, select the **Archive subject's encryption private key** check box.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- In addition to this procedure, the certification authority must be configured to archive keys. For more information, see Related Topics.
- Clients must be re-enrolled to receive a certificate that is based on the changed template if they already have a valid certificate that is based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Establishing key options and key archival
Key archival
Re-enroll all certificate holders
Key archival and recovery
Manage Key Archival and Recovery
**Other Resources**
Active Directory Certificate Services PKI - Key Archival and Management

# Configure a certificate template for client autoenrollment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a certificate template for client autoenrollment**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Security** tab, add the users or groups that you want and, under **Allow**, select the **Read**, **Enroll**, and **Autoenroll** check boxes.
4. (Optional) To automatically enroll the subject without any more user input, on the **Request Handling** tab, click **Enroll subject without requiring any user input**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- In addition to this procedure, other settings must be configured to allow client autoenrollment. For more information, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.
- When autoenrollment retrieves a certificate that is configured to be stored on a smart card, you will receive a popup message when the certificate is ready to be stored on the smart card. If the message refers to a type of smart card that you do not have, click **Cancel** until the correct smart card type appears. You then provide the smart card PIN and the certificate will be stored on the smart card.
- When autoenrollment enrolls for a certificate that requires user interaction for the enrollment process, you will receive a **Certificate Enrollment** popup message and an icon will appear on your taskbar. Clicking that icon or message begins the certificate autoenrollment process.

**See Also**
**Concepts**
Certificate Services example implementation: Establishing autoenrollment for user certificates
Version 2 certificate templates
Planning for autoenrollment deployment
Allowing for autoenrollment
Re-enroll all certificate holders

# Modify enrollment policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify enrollment policy**

1. Open Certificate Templates
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. Click the **Issuance Requirements** tab.
4. Provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create an application policy
Re-enroll all certificate holders

# Create an application policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an application policy**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Application Policies**, and then click **Edit**.
4. In **Edit Application Policies Extension**, click **Add**.
5. In **Add Application Policy**, click **New**.
6. Provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create a new object identifier
Re-enroll all certificate holders
Application policies
Version 2 certificate templates

# Create an issuance policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an issuance policy**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Issuance Policies**, and then click **Edit**.
4. In **Edit Issuance Policies Extension**, click **Add**.
5. In **Add Issuance Policy**, click **New**.
6. Provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create a new object identifier
Re-enroll all certificate holders
Application policies
Version 2 certificate templates

# Add an application policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an application policy**

1. Open **Certificate Templates**
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Application Policies** and then click **Edit**.
4. In **Edit Application Policies Extension**, click **Add**.
5. In **Add Application Policy**, click the application policy that you want, then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press **Enter**.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create an application policy
Create a new object identifier
Re-enroll all certificate holders
Version 2 certificate templates

# Add an issuance policy

Updated: November 2, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an issuance policy**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Issuance Policies**, and then click **Edit**.
4. In **Edit Issuance Policies Extension**, click **Add**.
5. In **Add Issuance Policy**, click **New**.
6. Provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Issuance Policies are also known as Certificate Policies.
- You can optionally fill out the **CPS Location** with a URI. For example, the Microsoft CPS is located at http://www.microsoft.com/pki/rms/cps/. For more information, see Creating Certificate Policies and Certificate Practice Statements (http://go.microsoft.com/fwlink/?LinkId=205512).
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create a new object identifier
Re-enroll all certificate holders
Version 2 certificate templates
Issuance policies

# Make certificate or application policy critical

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make certificate or application policy critical**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Issuance Policies** or **Application Policies** and then click **Edit**.
4. Select the **Make this extension critical** check box.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press ENTER.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Certificate Policies are also known as Issuance Policies.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create a new object identifier
Re-enroll all certificate holders
Version 2 certificate templates
Issuance policies
Application policies

# Change key usage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change key usage**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, in the **Extensions included in this template** box, click **Key Usage**, and then click **Edit**.
4. Make any necessary changes.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Changing key usage is affected by the **Purpose** set for this certificate template on the Extensions tab.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Add an application policy
Add an issuance policy
Configuring key usage
Re-enroll all certificate holders
Version 2 certificate templates

# Change basic constraints

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change basic constraints**

1. Open Certificate Templates
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Basic Constraints** and then click **Edit**.
4. In **Edit Basic Constraints Extension**, provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure only works with certificate templates that issue certificates that sign other certificates, such as Cross Certification Authority and Root Certification Authority.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.

**See Also**
**Concepts**
Renew a subordinate certification authority
Renew a root certification authority
Version 2 certificate templates
Create an application policy
Create an issuance policy
Add an application policy
Add an issuance policy

# Manage Key Archival and Recovery

Updated: November 5, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage key archival and recovery**

Configure a certificate template for key archival and recovery

- Identify a key recovery agent
- Enable key archival
- Certutil tasks for key archival and recovery

**See Also**

**Other Resources**

Active Directory Certificate Services PKI - Key Archival and Management

# Configure a certificate template for key archival and recovery

Updated: November 7, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a certificate template for key archival and recovery**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Request Handling** tab, select the **Archive subject's encryption private key** check box.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- In addition to this procedure, the certification authority must be configured to archive keys. For more information, see Related Topics.
- Clients must be re-enrolled to receive a certificate that is based on the changed template if they already have a valid certificate that is based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Establishing key options and key archival
Key archival
Re-enroll all certificate holders
Key archival and recovery
Manage Key Archival and Recovery
**Other Resources**
Active Directory Certificate Services PKI - Key Archival and Management

# Identify a key recovery agent

Updated: November 5, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To identify a key recovery agent**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Recovery Agents** tab, click **Archive the key**.
6. In the **Number of recovery agents to use** box, type the number of key recovery agents that will be used to encrypt the archived key.
7. Click **Add** to add key recovery agent certificates.

**Notes**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- To add key recovery agents, the **Number of recovery agents to use** must be 1 or more.
- If the **Number of recovery agents to use** value exceeds the number of recovery agent certificates with the status of "Valid," enrollment requests that require key archival will fail.
- This procedure configures a certification authority to archive private keys when issuing certificates based on templates that have key archival configured. For more information about configuring certificate templates for key archival, see Related Topics.
- When the recovery agent certificates are added to the CA, a status is displayed for each certificate. Status can be one of the following values and causes:

| Status | Cause |
|---|---|
| Expired | The certificate's expiration date has passed so the certificate cannot be used. |
| Invalid | The certificate may be malformed or causes and error when loading. |
| Not found | The certificate was configured but cannot be located by the CA. |
| Not loaded | The certificate was configured but has not yet been loaded by the CA. |
| Revoked | The certificate has been revoked and cannot be used. |

| | |
|---|---|
| Untrusted | The root CA for this certificate is not trusted by the CA. |
| Valid | The certificate has been loaded by the CA and is operating normally. |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Certificate Services example implementation: Key archival and recovery
Key archival and recovery
Establishing key options and key archival
**Other Resources**
Active Directory Certificate Services PKI - Key Archival and Management

# Enable key archival

Updated: November 5, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable key archival**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. Click the **Recovery Agents** tab and click **Archive the key**.
6. In the **Number of recovery agents to use** box, specify the number of recovery agents required to recover an archived key, and then add the recovery agents that you want to the **Key recovery agent certificates** list.

**Note**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Certificate Services example implementation: Key archival and recovery
Key archival and recovery
Working with MMC console files
Access control overview
Identify a key recovery agent
**Other Resources**
Active Directory Certificate Services PKI - Key Archival and Management

# Certutil tasks for key archival and recovery

Updated: November 5, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certutil tasks for key archival and recovery**

You can use **certutil** to retrieve and recover archived keys.

To view the syntax for a specific task, click a task:

- To retrieve an archived private key recovery blob
- To recover an archived private key

**To retrieve an archived private key recovery blob**

**Syntax**

**certutil -getkey** [**-f**] [**-gmt**] [**-seconds**] [**-v**] *SearchToken* [*RecoveryBlobOutFile*]

**Parameters**

**-getkey**
Retrieves the archived private key.

**-f**
Overwrites existing files or keys.

**-gmt**
Displays time as Greenwich mean time.

**-seconds**
Displays time with seconds and milliseconds.

**-v**
Specifies verbose output.

***SearchToken***
Specifies the keys and certificates that you want to recover.

***RecoveryBlobOutFile***
Specifies the output file containing a certificate chain and an associated private key, still encrypted to one or more key recovery agent (KRA) certificates.

**-?**
Displays a list of **certutil** commands.

**Remarks**

- *SearchToken* can be a certificate common name, a certificate serial number, a certificate Secure Hash Algorithm (SHA-1) hash, a requester name (that is, domain\user), or a user principal name (UPN) (that is, domain@user).

**To recover an archived private key**

**Syntax**

**certutil -recoverkey** [**-f**] [**-user**] [**-gmt**] [**-seconds**] [**-split**] [**-v**] [**-p***Password*] *RecoveryBlobInFile* [*PFXOutFile*] [*RecipientIndex*]]

**Parameters**

**-recoverkey**
Recovers the archived private key.

**-f**
Overwrites existing files or keys.

**-user**
Uses the **HKEY_CURRENT_USER** keys or certificate store.

**-gmt**
Displays time as Greenwich mean time.

**-seconds**
Displays time with seconds and milliseconds.

**-split**
Splits the embedded Abstract Syntax Notation One (ASN.1) elements, and saves them to files.

**-v**
Specifies verbose output.

**-p *Password***
Specifies a password.

***RecoveryBlobInFile***
Specifies the input file that contained the recovery blob retrieved from the CA.

***PFXOutFile***
Specifies the file where you want to save the recovered key and associated PKCS #12 certificate.

***Password***
Encrypts *PFXOutFile* to a password.

***RecipientIndex***
Specifies the index of the key recovery agent (KRA) certificate to be used for decrypting the private key blob. If you do not specify this parameter, **certutil** tries all of the KRA certificates.

**-?**
Displays a list of **certutil** commands.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Command-line reference A-Z
Command shell overview
**Other Resources**
Active Directory Certificate Services PKI - Key Archival and Management

# Change Policy Settings of a Certification Authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Change policy settings of a certification authority**

- Specify CA certificate access points in issued certificates
- Publish certificates to the file system

# Specify CA certificate access points in issued certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify CA certificate access points in issued certificates**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Extensions** tab, click **Select extension**, and then click **Authority Information Access (AIA)**.
6. Specify the locations from which users can obtain the certificate for this CA.

| To | Do this |
|---|---|
| Add a URL that will be published as part of any certificate issued by a CA. | Click **Add**, then type a URL where users can obtain the CA's certificate. |
| Remove an authority information URL from the list on issued certificates. | Click the URL and then click **Remove**. |
| Indicate that that you do not want to use a URL as an authority information access point in certificates without removing it from the list. | Clear the **Include in the AIA extension of issued certificates** check box. |
| Indicate that a URL can now be used as an authority information access point. | Select the **Include in the AIA extension of issued certificates** check box. |
| Indicate that a URL can now be used for online certificate status protocol (OCSP). | Select the **Include in the online certificate status protocol (OCSP) extension** check box. |
| Indicate that that you do not want to use a URL for online certificate status protocol (OCSP) in certificates without removing it from the list. | Clear the **Include in the online certificate status protocol (OCSP) extension** check box. |

7. Stop and restart the Certificate Services service.

**Notes**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- Authority information access URLs can be either HTTP, FTP, LDAP, or FILE addresses. You can use the following variables when specifying the address of the authority information access point:

| Variable | Value |
|---|---|

| | |
|---|---|
| CAName | The name of the certification authority. |
| CAObjectClass | The object class identifier for a certification authority, used when publishing to an LDAP URL |
| CATruncatedName | The "sanitized" name of the certification authority, truncated to 32 characters with a hash on the end |
| CDPObjectClass | The object class identifier for CRL distribution points, used when publishing to an LDAP URL |
| CertificateName | The renewal extension of the certification authority |
| ConfigurationContainer | The location of the Configuration container in Active Directory |
| CRLNameSuffix | Inserts a name suffix at the end of the file name when publishing a CRL to a file or URL location |
| DeltaCRLAllowed | When a delta CRL is published, this replaces the CRLNameSuffix with a separate suffix to distinguish the delta CRL |
| ServerDNSName | The DNS name of the certification authority server |
| ServerShortName | The NetBIOS name of the certification authority server |

- To stop and restart the Certificate Services service, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify online certificate status protocol responder in issued certificates
Working with MMC console files
Start or stop the certification authority service
Configuring the policy and exit modules

# Publish certificates to the file system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To publish certificates to the file system**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Exit Module** tab, click **Properties**.
6. On the **Publication Settings** tab, select the **Allow certificates to be published to the file system** check box, and then click **OK**.
7. Stop and restart the Certification Authority service.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- To stop and restart Certificate Services, see Related Topics.
- To request a certificate and have it published to the file system, certificate requestors must include a **certfile:true** attribute in their request. After the certificate is issued, it is copied to *FileName*.cer, where *FileName* is the request ID of the certificate request.
  - The file is written in the CertEnroll share on the CA.
  - You cannot overwrite files using this method.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Start or stop the certification authority service
Policy and exit modules

# Change the Policy or Exit Modules of a Certification Authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Change the policy or exit modules of a certification authority**

- Select a different policy module
- Select a different exit module

# Select a different policy module

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select a different policy module**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Policy Module** tab, click **Select**.
6. Click the new policy module, and then click **OK**.
7. If the policy module has its own configuration interface, you can configure it by clicking **Properties**.
8. Stop and restart the Certificate Services service.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- To register a new policy module so that it is available in the list of policy modules, refer to the *Microsoft Platform Software Development Kit*.
- Only one policy module can be active.
- To stop and restart the Certificate Services service, at the command prompt, type:

  **net stop certsvc**

  **net start certsvc**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Policy and exit modules
Select a different exit module

# Select a different exit module

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select a different exit module**

1. Log on to the system as a Certification Authority Administrator.
2. Open Certification Authority.
3. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

4. On the **Action** menu, click **Properties**.
5. On the **Exit Module** tab, click **Add**
6. Click the new exit module and then click **OK**.
7. (Optional) If the exit module has its own configuration interface, to configure it, click **Properties**.
8. Stop and restart the Certificate Services service.

## Notes

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.
- To register a new exit module so that it is available in the list of exit modules, refer to the *Microsoft Platform Software Development Kit*.
- To stop and restart the Certificate Services service, at a command prompt, type:

  **net stop certsvc**

  **net start certsvc**

- Multiple exit policies can be enabled at the same time.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Policy and exit modules
Select a different policy module

# Manage Role-based Administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing role-based administration**

- Add a certification authority administrator
- Add a certificate manager
- Add a certification authority auditor
- Add a certification authority backup operator
- Enable role separation
- Disable role separation

# Add a certification authority administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a certification authority administrator**

1. Open Certification Authority.
2. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

3. On the **Action** menu, click **Properties**.
4. Click the **Security** tab, in **Group or user names**, confirm that **Administrators** is selected, and then click **Add**.
5. In **Permissions for Administrators**, type the name of the group or user you want to make a CA Administrator, and then click **OK**.
6. In **Permissions for** *Added User or Group*, for **Manage CA**, select the **Allow** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable role separation
Disable role separation
Add a certificate manager
Role-based administration
Add a certification authority auditor
Add a certification authority backup operator
Access control overview

# Add a certificate manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a certificate manager**

1. Open Certification Authority.
2. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

3. On the **Action** menu, click **Properties**.
4. Click the **Security** tab, click **Add** and, in **Enter the object names to select**, type the user or group you want to make a Certificate Manager, and then click **OK**.
5. In the **Permissions for***Added User or Group* box, for **Issue and Manage Certificates**, select the **Allow** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable role separation
Disable role separation
Add a certification authority administrator
Role-based administration
Add a certification authority auditor
Add a certification authority backup operator
Access control overview

# Add a certification authority auditor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a certification authority auditor**

1. Open Local Security Settings.
2. In the console tree, click **User Rights Assignment**.

   **Where?**

   - Security Settings/Local Policies/User Rights Assignments

3. In the details pane, double-click **Manage auditing and security log**.
4. Click **Add User or Group**.
5. In **Enter the object names to select (examples)**, type the name of the user or group you want to make an auditor, then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable role separation
Disable role separation
Add a certification authority administrator
Add a certificate manager
Role-based administration
Add a certification authority backup operator

# Add a certification authority backup operator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a certification authority backup operator**

1. Open Local Security Settings.
2. In the console tree, click **User Rights Assignment**.

   **Where?**

   - Security Settings/Local Policies/User Rights Assignments

3. In the details pane, double-click **Back up file and directories**.
4. Click **Add User or Group** and, in **Enter the object names to select**, type the name of the user or group to add as a backup operator, and then click **OK**.
5. In the details pane, double-click **Restore file and directories**.
6. Click **Add User or Group** and, in **Enter the object names to select (examples)**, type the name of the user or group to add as a backup operator, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable role separation
Disable role separation
Add a certification authority administrator
Add a certificate manager
Role-based administration
Add a certification authority auditor

# Enable role separation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable role separation**

**Caution**

- Before enabling role separation, make certain that no user is assigned to more than one role. If a user is assigned to more than one role and role separation is enabled, then the user will not be permitted to perform any activity for the Certificate Services service. For more information, see Role-based administration.

1. Open Command Prompt.
2. Type:

   **certutil -setreg ca\RoleSeparationEnabled 1**

3. Open Certification Authority.
4. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*

5. On the **Action** menu, point to **All Tasks**, and click **Stop Service** to stop the service.
6. On the **Action** menu, point to **All Tasks**, and click **Start Service** to start the service.

| Value | Description |
|---|---|
| **certutil** | Specifies the name of the command-line program. |
| **-setreg** | Modifies the registry. |
| **ca\RoleSeparationEnabled** | Indicates the registry value for role separation. |
| **1** | Specifies a value for the registry entry. A value of 1 enables role separation. |

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-

click **Certification Authority**.

- Once role separation is enabled, a local Administrator will not be able to perform any role-based administration of the CA. This is because the local Administrator account has both the CA Backup Operator and CA Auditor user rights.

- You can assign certification authority roles for role-based administration on servers running any version of the Windows Server 2003 family, but you can only enable role separation on servers running Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, including the 64-bit versions of Windows Server 2003, Enterprise Edition and 64-bit versions of Windows Server 2003, Datacenter Edition.

- To display the role separation setting, type:

  **certutil -getreg ca\RoleSeparationEnabled**

- To view the complete syntax for this command, at a command prompt, type:

  **certutil -setreg -?**

- To stop and restart the Certificate Services service at the command prompt, type:

  **net stop certsvc**

  **net start certsvc**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable role separation
Add a certification authority administrator
Add a certificate manager
Role-based administration
Add a certification authority auditor
Add a certification authority backup operator
Start or stop the certification authority service

# Disable role separation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable role separation**

1. Open Command Prompt.
2. Type:

   **certutil -delreg ca\RoleSeparationEnabled**
3. Open Certification Authority.
4. In the console tree, click the name of the certification authority (CA).

   **Where?**

   - Certification Authority (*Computer*)/*CA name*
5. On the **Action** menu, point to **All Tasks**, and click **Stop Service** to stop the service.
6. On the **Action** menu, point to **All Tasks**, and click **Start Service** to start the service.

| Value | Description |
|---|---|
| **certutil** | Specifies the name of the command-line program. |
| **-delreg** | Modifies the registry. |
| **ca\RoleSeparationEnabled** | Indicates the registry value for role separation. |

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To display the role separation setting, type the following at a command prompt:

  **certutil -getreg ca\RoleSeparationEnabled**
- To view the complete syntax for this command, at a command prompt, type:

  **certutil -delreg -?**

- To stop and restart the Certificate Services service at the command prompt, type:

  **net stop certsvc**

  **net start certsvc**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable role separation
Add a certification authority administrator
Add a certificate manager
Role-based administration
Add a certification authority auditor
Add a certification authority backup operator
Start or stop the certification authority service

# Certificate Services Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**
Certificate Services and the Certification Authority snap-in provide a way for organizations to set up and manage a certification authority (CA) on a server for use in a public key infrastructure (PKI).

- Deploying a Public Key Infrastructure
- Certificate Services overview
- New features in Certificate Services
- Understanding Certificate Services
- Using Certificate Services
- Certificates Resources

# Deploying a Public Key Infrastructure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Public key infrastructure**

Computer networks are no longer closed systems in which a user's mere presence on the network can serve as proof of identity. In this age of information interconnection, an organization's network may consist of intranets, Internet sites, and extranets--all of which are potentially susceptible to access by unauthorized individuals who intend to maliciously view or alter an organization's digital information assets.

There are many potential opportunities for unauthorized access to information on networks. A person can attempt to monitor or alter information streams such as e-mail, electronic commerce transactions, and file transfers. Your organization may work with partners on projects of limited scope and duration, with employees whom you know nothing about, but who, nonetheless, must be given access to some of your information resources. If your users have a multitude of passwords to remember for accessing different secure systems, they may choose weak or common passwords to more easily remember them. This not only provides an intruder with a password that is easy to crack, but also one that will provide access to multiple secure systems and stored data.

How can a system administrator be sure of the identity of a person accessing information and, given that identity, control which information that person has access to? Additionally, how can a system administrator easily and securely distribute and manage identification credentials across an organization? These are issues that can be addressed with a well-planned public key infrastructure.

A public key infrastructure (PKI) is a system of digital certificates, certification authorities (CAs) and other registration authorities (RAs) that verify and authenticate the validity of each party that is involved in an electronic transaction through the use of public key cryptography. Standards for PKIs are still evolving, even as they are being widely implemented as a necessary element of electronic commerce. For detailed information about planning a PKI and using public key cryptography, see Certificates Resources on public key infrastructure.

There are a number of reasons why an organization may choose to deploy a PKI using Windows:

- **Strong security.** You can have strong authentication with smart cards. You can also maintain confidentiality and the integrity of transmitted data on public networks by using Internet Protocol security and the confidentiality of stored data by using EFS (encrypting file system).

- **Simplified administration.** Your organization can issue certificates and in conjunction with other technologies eliminate the use of passwords. You can revoke certificates as necessary and publish certificate revocation lists (CRLs). There is the ability to use certificates to scale trust relationships across an enterprise. You can also take advantage of Certificate Services integration with the Active Directory directory service and policy. The capability to map certificates to user accounts is also available.

- **Additional opportunities.** You can exchange files and data securely over public networks, such as the Internet. You have the ability to implement secure e-mail using Secure Multipurpose Internet Mail Extensions (S/MIME) and secure Web connections using Secure Sockets Layer (SSL) or Transport Layer Security (TLS). You can also implement security enhancements to Wireless Networking.

The Windows Server 2003 family has features to help your organization implement a public key infrastructure include:

- **Certificates.** A certificate is basically a digital statement issued by an authority that vouches for the identity of the certificate holder. A certificate binds a public key to the identity of the person, computer, or service who holds the corresponding private key. Certificates are used by a variety of public key security services and applications that provide authentication, data integrity and secure communications across networks such as the Internet.

  The standard certificate format used by Windows certificate-based processes is X.509v3. An X.509 certificate includes information about the person or entity to whom the certificate is issued, information about the certificate, plus optional information about the certification authority issuing the certificate. Subject information may include the entity's name, the public key, and the public-key algorithm.

Users can manage certificates using the Certificates MMC console. For more information, see Certificates overview. Users can also allow certificate autoenrollment to manage their certificates automatically. For more information, see Planning for autoenrollment deployment.

- **Certificate Services**. On Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, Certificate Services is the component that is used to create and manage certification authorities (CAs). A CA is responsible for establishing and vouching for the identity of certificate holders. A CA also revokes certificates if they should no longer be considered valid and publishes certificate revocation lists (CRLs) to be used by certificate verifiers. The simplest PKI design has only one root CA. In practice, however, the majority of organizations deploying a PKI will use a number of CAs, organized into certification hierarchies.

  Administrators can manage Certificate Services using the Certification Authority MMC console. For more information, see Certificate Services overview.

- **Certificate templates.** Certificates are issued by the CA based on information provided in the certificate request and settings contained in a certificate template. A certificate template is the set of rules and settings that are applied against incoming certificate requests. For each type of certificate that an enterprise CA can issue, a certificate template must be configured.

  Certificate templates are customizable in Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition enterprise CAs, and are stored in Active Directory for use by all CAs in the forest. This allows the administrator to choose one or more of the default templates installed with Certificate Services or to create templates that are customized for specific tasks or roles.

  Administrators can manage Certificate Templates using the Certificate Templates MMC console. For more information, see Certificate Templates.

- **Certificate autoenrollment.** Autoenrollment allows the administrator to configure subjects to automatically enroll for certificates, retrieve issued certificates, and renew expiring certificates without requiring subject interaction. This requires no knowledge by the subject of any certificate operations, unless the certificate template is configured to interact with the subject or the CSP requires interaction (such as with a smart card CSP). This greatly simplifies the experience of the client with certificates and minimizes administrative tasks.

  Administrators can configure autoenrollment through configuration of Certificate Templates and CA settings. For more information see Planning for autoenrollment deployment.

- **Web enrollment pages.** These are a separate component of Certificate Services. These Web pages are installed by default when you set up a CA and allow certificate requesters to submit certificate requests using a Web browser. Additionally, the CA Web pages can be installed on servers running Windows that do not have a certification authority installed. In this case, the Web pages are used to direct certificate requests to a CA that, for whatever reason, you do not want requesters to directly access. If you choose to create custom Web pages for your organization to access a CA, the Web pages provided in your operating system can be used as samples. Refer to the *Microsoft Platform Software Development Kit* for information about customizing Certificate Services and CA Web pages.

- **Smart card support.** Windows supports logon via certificates on smart cards, as well as the use of smart cards to store certificates and private keys. They can be used for Web authentication, secure e-mail, wireless networking and other public key cryptography-related activities. For more information, see Smart Cards.

- **Public key policies.** You can use Group Policy in Windows to distribute certificates to subjects automatically, establish common trusted certification authorities, as well as managing recovery policies for EFS (Encrypting File System). For more information, see Public Key Policies overview

# Certificate Services overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate Services overview**

Certificate Services provides customizable services for issuing and managing certificates used in software security systems employing public key technologies. For background information about public key cryptography and the benefits of having a public key infrastructure (PKI), see Deploying a Public Key Infrastructure.

You can use Certificate Services in the Windows® operating system to create a certification authority (CA) which will receive certificate requests, verify the information in the request and the identity of the requester, issue certificates, revoke certificates, and publish a certificate revocation list (CRL).

For more information about CAs, see Certification Authorities

Certificate Services can also be used to:

- Enroll users for certificates from the CA using the Web or the Certificates Microsoft Management Console (MMC) snap-in, or transparently through autoenrollment.
- Use certificate templates to help simplify the choices a certificate requester has to make when requesting a certificate, depending upon the policy used by the CA.
- Take advantage of the Active Directory® directory service for publishing trusted root certificates, publishing issued certificates, and publishing CRLs.
- Implement the ability to log on to a Windows operating system domain using a smart card.

**Certificate policy**

A certificate policy is a set of instructions or rules that are used when processing certificate requests, issuing certificates, revoking certificates, and publishing CRLs. These instructions are a combination of administrative policy and configuration settings on the CA.

When you install Certificate Services, the CA is configured with a default set of rules and settings. These define CA-specific settings such as the CA's certificate, its default issuance behavior, and its key recovery agents. The CA may also install a number of preconfigured certificate templates, which define what information a certificate request must have and how to process incoming requests for a certificate based on that template. The combination of applying CA settings and certificate template settings, plus the defined administrative guidelines, results in the certificate policy that governs the operation of a CA.

**Processing certificate requests**

A user can request certificates using Microsoft® Internet Explorer 5.0 or later or a browser such as Netscape Navigator 4.7*x* or later. In addition, a user can use the Certificates snap-in to request a certificate from an enterprise CA or an administrator can configure certificate autoenrollment to transparently request and install certificates for users.

When a user initiates a certificate request, a cryptographic service provider (CSP) on their computer generates a public key and private key pair for the user. The user's public key is sent with their necessary identifying information to the CA. If the user's identifying information meets the CA criteria for granting a request, the CA generates the certificate, which is retrieved by the client application and stored locally. For more information about certificates, see Understanding Certificates.

**Security considerations for CAs**

CAs are valuable resources, and you should provide them with a high degree of protection. Specific actions that should be considered include:

- **Physical protection.** Since CAs represent highly-trusted entities within an enterprise, you should protect them from tampering, depending on the inherent value of the certification made by the CA. Physical isolation of the CA server, in a facility accessible only to security administrators, can dramatically reduce the possibility of such attacks.
- **Restoration.** A CA might be lost if there is a hardware failure. This can create a number of administrative and operational problems, and can prevent revocation of existing certificates. Certificate Services supports the backup of a CA

using Backup so that it can be restored at a later time. This is an important part of the overall CA management process.

- **Key management.** The CA's keys are its most valuable asset, because the private key provides the basis for trust in the certification process. Cryptographic hardware modules can provide tamper-resistant key storage and isolate the cryptographic operations from other software running on the server. This reduces the likelihood of a CA key being compromised. Certificate Services supports cryptographic service providers (CSPs) from other sources, but the documentation included with Windows is specific to the software CSPs that are included with Windows. If you use a CSP from another source, you should confirm with the vendor that the CSP can work with Certificate Services.

**Customizing Certificate Services**

Certificate Services includes programmable interfaces so that developers can create support for additional transports, policies, and certificate properties and formats. Refer to the *Microsoft Platform Software Development Kit* for information about customizing Certificate Services.

For more information about Certificate Services, see Understanding Certificate Services

# New features in Certificate Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features in Certificate Services**
Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition have a number of new features and improvements related to Certificate Services and public key infrastructure (PKI). A few new procedures are also provided to demonstrate certificate template editing features and certificate autoenrollment for users and computers.

**New PKI features**
Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition implement new PKI features and improvements.

- **Editable certificate templates and the Certificate Templates MMC snap-in**

  Certificate templates were available in Windows 2000 Certificate Services, but they could not be modified or changed. There is a new Certificate Templates MMC snap-in that enables administrators to:

  - Create a new certificate template by duplicating and renaming an existing template.
  - Modify template properties such as certificate validity period, renewal period, cryptographic service provider (CSP), key size, key archival
  - Establish and apply enrollment policies, issuance policies and application policies.

    For example, a user can enroll themselves for a certificate or an enrollment agent certificate can enroll for a certificate on their behalf. A significant issue with this scenario is that an enrollment agent can enroll for any user in the enterprise. This means the enrollment agent certificate is very powerful and only very trusted people may have access to it.
  - Allow for autoenrollment for certificates based on the template.
  - Set access control on certificate templates to establish which users or computers can enroll and autoenroll for certificates.

  There are now two different versions of certificate templates for Windows Server operating systems. Windows 2000 clients and certificate services can only use version 1 templates. Windows XP and Windows Server 2003 family clients and certificate services will support both version 1 and version 2 templates. Version 1 templates are read-only, but when they are duplicated, the duplicate template is an editable version 2 template.

- **Certificate autoenrollment and autorenewal for all subjects**

  In Windows 2000, it was possible to autoenroll for EFS certificates and computer certificates, however, autoenrollment for users was not possible. The new autoenrollment feature improves both the user and computer enrollment experience. A member of the Enterprise Admins group can specify the types of certificates that any entity should automatically be issued. The enterprise administrator controls autoenrollment by setting security permissions on certificate templates using the Certificate Templates snap-in. A Windows XP or Windows Server 2003 family client then accesses the templates in the Active Directory directory service and, if access has been granted, then it will enroll for those certificates.

  For an example of establishing autoenrollment for user certificates, see Certificate Services example implementation: Establishing autoenrollment for user certificates.

  Autorenewal is a new feature similar to autoenrollment and the same mechanism on the templates is used to control who can autorenew a certificate. Every certificate in the certificate store that has a template extension can potentially be

autorenewed by the system. This means that applications no longer need to worry about certificates expiring.

- **Delta CRLs**

  Many applications require up-to-date certificate revocation status information. This requires the certification authority (CA) to frequently publish a new certificate revocation list (CRL). A CRL is the entire list of revoked certificates, so for a CA with a large amount of issued certificates, this can become a very large list. Even if there are no changes, a CA has to republish the entire list so that applications have the latest information, which involves a lot of repetition. Frequent publication of large objects will in turn generates a large amount of replication traffic.

  Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition have a new feature called delta CRLs, an option in RFC 2459. Delta CRLs are CRLs which contain the list of changes in revocation status from a full "base" CRL. Because delta CRLs are a list of changes and not a restatement of the entire CRL, they are typically much smaller and generate significantly less replication traffic than large base CRLs.

- **Role-based administration**

  Certificate Services in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition allows for the separation of roles for the management and maintenance of a CA. Role separation is not enforced by default, but you can select to enforce assigned roles.

  For more information, see Role-based administration.

- **Key archival and recovery**

  You can configure a CA to archive the keys associated with the certificates it issues. If necessary, you can then recover lost keys through the use of a key recovery agent certificate.

  For more information, see Certificate Services example implementation: Key archival and recovery.

- **Event auditing**

  Event auditing provides the ability to log most events that occur on a CA. This can be useful for monitoring the activities of a CA or the administrative functions, such as certificate issuance and role changes.

  For more information, see Configure event auditing.

- **Qualified subordination**

  Qualified subordination is an extension of standard CA subordination that allows you to:

  - Define the namespaces for which a subordinate CA will issue certificates.
  - Specify the acceptable uses of certificates issued by a qualified subordinate CA.
  - Enable a certificate to be used in separate certification hierarchies.

  For more information, see Qualified subordination.

- **New Certutil.exe commands**

  **certutil -dspublish** [**cert**|**crl**]

  Publishes the CA certificate or the certificate revocation list to Active Directory

  For examples of using **certutil -getkey** and **certutil -recoverkey** to recover archived keys, see Certificate Services example implementation: Key archival and recovery.

# Understanding Certificate Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Certificate Services**
This section covers:

- Processing certificate requests
- Certificate revocation
- Key archival and recovery
- Requesting certificate templates
- Policy and exit modules
- Certification authority hierarchies
- Certification authority database
- Certification authority Web enrollment services
- Interoperability in Certificate Services
- Role-based administration
- Certification Authorities

# Processing certificate requests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Processing certificate requests**

A certification authority (CA) performs the following operations when processing a certificate request:

- **Request reception.** The certificate request is sent by the client application (such as the Certificate Request Wizard in the Certificates snap-in), which formats it into a PKCS #10 or CMC format request and submits it to the CA. On receipt, the CA server engine stores the private key in its certificate database if it is configured for key recovery.

- **Request approval.** The CA server engine calls the CA policy module, which queries the request properties, decides whether the request is authorized or not, and sets optional certificate properties. If the request is denied or set to Pending, the requestor is notified.

- **Certificate formation.** If the request is approved, the CA server engine takes the request and any properties requested by the policy module, and then builds a complete certificate.

- **Certificate publication.** The CA server engine stores the completed certificate in its certificate database and notifies the intermediary application of the request status. If the exit module has requested it, the server engine notifies it of a certificate issuance event. This allows the exit module to perform further operations, such as sending e-mail to the subject or the CA administrator. Meanwhile, the client application gets the issued certificate from the certificate database and saves it in its own local certificate store.

# Certificate revocation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate revocation**
Revocation of a certificate invalidates a certificate as a trusted security credential prior to the natural expiration of its validity period. There are a number of reasons why a certificate, as a security credential, could become untrustworthy prior to its expiration. Examples include:

- Compromise, or suspected compromise, of the certificate subject's private key.
- Compromise, or suspected compromise, of a certification authority's private key.
- Discovery that a certificate was obtained fraudulently.
- Change in the status of the certificate subject as a trusted entity.
- Change in the name of the certificate subject.

A public key infrastructure (PKI) depends on distributed verification of credentials in which there is no need for direct communication with the central trusted entity that vouches for the credentials. This creates a need to distribute certificate revocation information to individuals, computers, and applications attempting to verify the validity of certificates. The need for revocation information and its timeliness will vary, according to the application and its implementation of certificate revocation checking.

To effectively support certificate revocation, the client must determine whether the certificate is valid or has been revoked. To support a variety of scenarios, Certificate Services supports industry-standard methods of certificate revocation. These include publication of certificate revocation lists (CRLs) and delta CRLs in several locations for clients to access, including the Active Directory directory service, Web servers, and network file shares.

CRLs are complete, digitally-signed lists of unexpired certificates that have been revoked. This list is retrieved by clients who can then cache it (based on the configured lifetime of the CRL) and use it to verify certificates presented for use. Because CRLs can get large, depending on the size of the certification authority, delta CRLs can also be published. Delta CRLs contain only the certificates revoked since the last base CRL was published. This allows clients to retrieve the smaller delta CRL and quickly build a complete list of revoked certificates. The use of delta CRLs also allows more frequent publishing, because the size of the delta CRL usually does not require as much overhead as a full CRL.

For conceptual information about using certificate revocation in Certificate Services, see Revoking certificates and publishing CRLs. For procedures to manage certificate revocation, see Manage Certificate Revocation.

# Key archival and recovery

Updated: November 5, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key archival and recovery**
Windows Server 2003, Enterprise Edition can be configured to archive the private key of specific certificates when they are issued. This private key archive allows the key to be recovered at a later time if the private key is lost. This process is implemented in two separate phases: key archival and key recovery.

| **Tip** |
| --- |
| For updated information, see Active Directory Certificate Services PKI - Key Archival and Management. |

**Key archival**
The process of obtaining a certificate includes the subject locating the appropriate certificate template, gathering the information required by that template, and supplying it to a certification authority. This information normally contains information such as the subject name, public key and supported cryptographic algorithms. When key archival is configured, the subject will also provide their private key to the certification authority. The certification authority stores that private key in its database until you want to perform key recovery.

By default, the private key of issued certificates is not archived. This is because the storage of the private key in multiple locations, by definition, allows more attacks against it. To configure a certification authority to issue certificates with key archival enabled, see Certificate Services example implementation: Key archival and recovery.

**Key recovery**
Subjects can lose their private key in a variety of ways such as accidental deletion or deliberate misuse. An administrator may also want to recover the key of a particular subject to access data protected by that key. Key recovery can be used whenever the key archival process has stored the subject's private key.

The key recovery process requires an administrator to retrieve the encrypted certificate and private key and then a key recovery agent (KRA) to submit to the certification authority. When a correctly signed key recovery request is received, the subject's certificate and private key are provided to the requestor. The requestor would then use the key as appropriate or securely transfer the key to the subject for continued use. No recertification or rekeying is necessary, as the private key is not necessarily compromised.

For more information, see Manage Key Archival and Recovery.

# Requesting certificate templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate templates**

A certificate template profiles certificates based on their intended use. When requesting a certificate from a Microsoft certification authority (CA), the certificate requester will, depending on their access rights, be able to select from a variety of certificate types that are based on certificate templates, such as User and Basic EFS. The certificate template saves users from low-level, technical decisions about the type of certificate that they need. Instead, they can rely on the judgment of their administrators and use the template name that indicates the purpose of the certificate. If none of the preset certificate templates meet your needs, you can create new certificate templates and customize them for a variety of different uses.

For more information, see Certificate Templates.

# Policy and exit modules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policy and exit modules**

One powerful administrative feature of Certificate Services is the ability to control and customize the behavior of the certification authority (CA) through the use of policy and exit modules.

Policy modules can determine whether a certificate request should be automatically approved, denied, or marked as Pending. Exit modules provide an opportunity to perform post-processing after a certificate is issued.

Certificate Services comes with one exit module (Certxds.dll) and one policy module (Certpdef.dll). The policy module includes two separate policies: enterprise and stand-alone. To compare a CA using enterprise policy and a CA using stand-alone policy, see Enterprise certification authorities and Stand-alone certification authorities.

As a CA administrator, you can replace these default modules with your own custom policy and exit modules or commercial policy and exit modules. In addition, if you have upgraded to Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition Certificate Services from an earlier version of Certificate Services, you will have the option of using the policy module you have been using prior to upgrading. It will either be listed as a *legacy* policy module when you look at the properties of the CA or with its original name, depending on how it was created.

The policy module provided with the Microsoft CA performs the following function:

- **Determines the default action of a certification authority upon receiving a certificate request.** Upon receiving a certificate request, a certification authority can either automatically issue a certificate or hold it as Pending until an administrator reviews the request.

  Microsoft certification authorities can either issue a certificate automatically upon receiving a request or hold the request as Pending. In the majority of instances, the administrator of a stand-alone CA will want to have all incoming certificate requests set to Pending. Otherwise, because the stand-alone CA does not verify the identity of requesters via the Active Directory directory service, there is no way to verify the identity and validity of the certificate requester.

The CA can only have one policy module loaded at a time. The Windows 2000 CA policy module contained a great deal of functionality that has been integrated into the core certification authority functionality. This allows the policy module to be more easily replaced without losing functionality.

The exit module that is provided with the Microsoft CA performs the following functions:

- **Sends e-mail when a certification event occurs.** If the CA is configured to send e-mail when specific events occur, the exit module will do so.

- **Allows certificate publication to the file system.** If the certificate request specifies a location to publish the certificate in the file system, the exit module will do so.

Please note that this is not an exhaustive list of the functions of the exit module. Unlike the policy module, multiple exit modules can be used by a CA simultaneously.

To configure the settings of the default policy and exit modules, see Configuring the policy and exit modules. To configure e-mail sending options, see Send e-mail when a certification event occurs.

**Customizing Certificate Services policy and exit modules**

Programmable interfaces are included in Certificate Services for developers to create customized policy modules. For more information, refer to the *Microsoft Platform Software Development Kit*.

If you have created a customized policy module using the guidelines in the *Microsoft Platform Software Development Kit* and you want to change the policy module, see Select a different policy module.

If you have created a customized exit module using the guidelines in the *Microsoft Platform Software Development Kit* and you want to change or add an exit module, see Select a different exit module.

# Certification authority hierarchies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certification authority hierarchies**

The Microsoft public key infrastructure (PKI) supports a hierarchical certification authority (CA) model. A certification hierarchy provides scalability, ease of administration, and consistency with a growing number of commercial and other CA products.

In its simplest form, a certification hierarchy consists of a single CA. However, in general, a hierarchy will contain multiple CAs with clearly defined parent-child relationships. In this model, the child subordinate certification authorities are certified by their parent CA-issued certificates, which bind a certification authority's public key to its identity. The CA at the top of a hierarchy is referred to as the root authority, or root CA. The child CAs of the root CAs are called subordinate certification authorities (CAs).

In Windows XP and the Windows Server 2003 family, if you trust a root CA (by having its certificate in your Trusted Root Certification Authorities certificate store), you trust every subordinate CA in the hierarchy, unless a subordinate CA has had its certificate revoked by the issuing CA or has an expired certificate. Thus, any root CA is a very important point of trust in an organization and should be secured and maintained accordingly.

Verification of certificates thus requires trust in only a small number of root CAs. At the same time, it provides flexibility in the number of certificate-issuing subordinate CAs. There are several practical reasons for supporting multiple subordinate CAs, including:

- **Usage.** Certificates may be issued for a number of purposes, such as secure e-mail and network authentication. The issuing policy for these uses may be distinct, and separation provides a basis for administering these polices.

- **Organizational divisions.** There may be different policies for issuing certificates, depending upon an entity's role in the organization. Again, you can create subordinate CAs to separate and administer these policies.

- **Geographic divisions.** Organizations may have entities at multiple physical sites. Network connectivity between these sites may dictate a requirement for multiple subordinate CAs to meet usability requirements.

- **Load balancing.** If your PKI is going to support the issuing of a large number of certificates, having only one certification authority issue and manage all of these certificates can result in considerable network load for that single certification authority. Using multiple subordinate certification authorities to issue the same kind of certificates divides the network load between certification authorities.

- **Backup and fault tolerance.** Multiple certification authorities increase the possibility that your network will always have operational certification authorities available to service users.

Such a certification authority hierarchy also provides administrative benefits, including:

- Flexible configuration of the CA security environment to tailor the balance between security and usability, such as key strength, physical protection, and protection against network attacks. For example, you may choose to employ special purpose cryptographic hardware on a root CA, operate it in a physically secure area, or operate it offline. These may be unacceptable for subordinate CAs, due to cost or usability considerations.

- The ability to "turn off" a specific portion of the CA hierarchy without affecting the established trust relationships. For example, you can easily shut down and revoke an issuing CA certificate that is associated with a specific geographic site without affecting other parts of the organization.

For more information, see Establishing a certification hierarchy

# Certification authority database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certification authority database**
When you install a certification authority (CA), you create a certification authority database on the server. This database should be located on an NTFS file system partition on the server's disk drives to provide the best security possible for the database file. You specify the locations for the database during the setup of a CA. By default, the database is located in:

*systemroot*\system32\certlog

The name of the database file is based on the name of the CA, with an .edb extension.

The Certification Authority MMC console provides a view into the certification authority database on a CA and the ability to administer the database.

A certification authority database stores:

- Every certificate issued by the CA.
- Every private key archived by the CA.
- Every certificate revoked by the CA.
- Every certificate request received by the CA, regardless of the disposition.

You also specify the location of the certificate database log during Certificate Services setup. The certificate database log keeps a record of every transaction involving the certificate database. Certificate database logs are used when restoring the CA from a backup. If a CA is restored from a backup that is one month old, then the certification database logs that have been generated since the last backup can be replayed against the database to restore it to its most current state. When you back up a certification authority, the existing certificate database logs will be truncated in size, because they will no longer be needed to restore the certificate database to its most current state.

For more information about CA backup and restore, see Backing up and restoring a certification authority.

# Certification authority Web enrollment services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certification authority Web enrollment services**

A set of certification authority (CA) Web pages is provided with Certificate Services in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. These Web pages provide a simple user interface to perform many of the common user tasks on a CA, and they are installed by default when you set up Certificate Services.

These Web enrollment pages allow you to connect to the CA via a Web browser and perform common tasks, such as requesting certificates from a CA, requesting the CA's certificate, submitting a certificate request using a PKCS #10 file, retrieving the CA's certificate revocation list (CRL) or performing smart card certificate enrollment. For a stand-alone CA, the Web pages are the primary way a certificate requester can interface with the CA, since the Certificates snap-in cannot be used to request certificates from a stand-alone CA. Enterprise CAs can accept certificate requests via the Certificates snap-in or the Web enrollment pages.

You can install Web enrollment pages for a CA on a Windows server 2003 family server that is not a CA in order to divert Web-based traffic from the actual CA. For more information about installing the Web pages on a server that is not a CA, see Set up certification authority Web enrollment support. You can also configure your Web server to use Secure Sockets Layer (SSL) encryption to protect the network traffic that contains the enrollment request and response. This is especially important when clients are using Windows 95, Windows 98, Windows Millennium Edition and Windows 2000 because of their handling of certificates.

For more information about using the Web enrollment pages, see Using Windows 2000 Certificate Services Web pages or Using Windows Server 2003 Certificate Services Web pages.

# Interoperability in Certificate Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interoperability in Certificate Services**
By using industry-standard X.509v3 certificate formats and open interfaces, Certificate Services operates with many products and technologies that support the use of public key cryptography and public key infrastructure (PKI).

**Internet standards**
PKI standards for the Internet are still evolving as this is being written. However, Certificate Services has been designed to adhere to existing PKI interoperability standards established by the Internet Engineering Task Force (IETF). The IETF working group charged with defining the basis for an interoperable PKI is PKIX. For more information on PKIX, see the PKIX Web page. The Web page also has a link to RFC 2459, "Internet Public Key Infrastructure X.509 Certificate and CRL Profile, Part 1," which is the specification for the basic architecture of PKI.

**Public Key Cryptography Standards**
There is a set of *de facto* cryptographic message standards called Public Key Cryptography Standards (PKCS) which are developed and maintained by RSA Laboratories. (For more information, see the Web page at RSA Laboratories.)

PKCS provides a basic, but well-understood framework for interoperability. The standards that are most relevant to PKI and ones that are used by Certificate Services are PKCS #7, Cryptographic Message Syntax Standard, PKCS #10, Certification Request Syntax Standard, and PKCS #12, Personal Information Exchange Syntax Standard.

**Web servers**
On a corporate intranet or on the Internet, Web servers, such as those using Microsoft Internet Information Services (IIS), can perform client authentication for secure communications using certificates generated by Certificate Services. Certificate Services can also generate server certificates used by IIS and other Web servers to provide server authentication to assure clients that they are communicating with the intended entity.

**Web browsers**
Certificate Services can issue certificates that are used by Web browsers that support client authentication, such as Microsoft Internet Explorer 5.01 or later. Certificate Services also supports Web enrollment, where a Web browser uses Web pages to request and retrieve a certificate.

**Windows**
Windows enables you to map certificates to users and groups. This mapping is automatic if the certificate is issued from an enterprise certification authority (CA). You can then use standard Windows administrative tools, such as security permission sets, to implement Internet and intranet security requirements taking advantage of the relationship established between domain users and the certificates issued to them.

For more information about mapping certificates in Windows see Mapping certificates to user accounts.

### Note

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Role-based administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Role-based administration**

You can use role-based administration to organize your certification authority (CA) administrators into separate, predefined CA roles, each with its own set of tasks. Roles are assigned using each user's security settings. You assign a role to a user by assigning that user the specific security settings that are associated with the role. A user that has one type of permission, such as Manage CA permission, can perform specific CA tasks that a user with another type of permission, such as Issue and Manage Certificates permission, cannot perform. For more information, see Role explanation.

**Note**

- Role-based administration is supported by both Windows server 2003 enterprise and stand-alone certification authorities. For information about using Windows 2000 server certification authorities and role-based administration, see Windows 2000 roles andWindows server 2003role-based administration.

**Role explanation**

Role-based administration involves CA roles, users, and groups. To assign a role to a user or group, you must assign the role's corresponding security permissions, group memberships, or user rights to the user or group. These security permissions, group memberships, and user rights are used to distinguish which users have which roles. The following table describes the CA roles of role-based administration and the groups relevant to role-based administration.

| Roles and groups | Security permission | Description |
|---|---|---|
| CA Administrator | **Manage CA** permission | Configure and maintain the CA. This is a CA role and includes the ability to assign all other CA roles and renew the CA certificate. |
| Certificate Manager | **Issue and Manage Certificates** permission | Approve certificate enrollment and revocation requests. This is a CA role. This role is sometimes referred to as CA Officer. |
| Backup Operator | **Back up file and directories** and **Restore file and directories** permissions | Perform system backup and recovery. This is an operating system role. |
| Auditor | **Manage auditing and security log** permission | Configure, view, and maintain audit logs. This is an operating system role. |
| Enrollees | Authenticated Users | Enrollees are clients who are authorized to request certificates from the CA. This is not a CA role. |

All CA roles are assigned and modified by local Administrators, Enterprise Admins, and Domain Admins. Local Administrators, Enterprise Admins, and Domain Admins are CA Administrators by default on an Enterprise CA. Only local Administrators are CA Administrators by default on a stand-alone CA. If the stand-alone CA is logged on to an Active Directory domain, Domain Admins are also CA Administrators.

The CA Administrator and Certificate Manager roles can be assigned to both Active Directory users or local users in the Security Accounts Manager (SAM) of the local computer, which is the local security account database. As a best practice, assign roles to group accounts instead of individual user accounts.

Only CA Administrator, Certificate Manager, Auditor, and Backup Operator are CA roles. The other users described in the table are relevant to role-based administration and should be understood before assigning CA roles.

Only CA Administrators and Certificate Managers are assigned using the Certification Authority Microsoft Management Console (MMC) snap-in. Other roles, users, and groups are specified in their related consoles. To change the roles of a user, you must change the user's security permissions, group membership, or user rights. For information about assigning roles, see Assigning roles.

When key archival is configured, the subject obtaining a certificate from a CA will provide their private key to the CA. The CA stores that private key in its database until key recovery. Only a Certificate Manager can get the encrypted private key blob out

of the CA database, which is then passed on to key recovery agents (KRAs). For more information, see Key archival and recovery.

**Roles and activities**

Each CA role has a specific list of CA administration tasks associated with it. The following table lists all the CA administration tasks along with the roles in which they are performed.

| Activity | CA Administrator | Certificate Manager | Auditor | Backup Operator | Local Administrator | Notes |
|---|---|---|---|---|---|---|
| Install CA | | | | | X | |
| Configure policy and exit module | X | | | | | |
| Stop and start the Certificate Services service | X | | | | | |
| Configure extensions | X | | | | | |
| Configure roles | X | | | | | |
| Renew CA keys | | | | | X | |
| Define key recovery agents | X | | | | | |
| Configure Certificate Managers restrictions | X | | | | | |
| Delete single row in database | X | | | | | |
| Delete multiple rows in database (bulk deletion) | X | X | | | | The user must be both a CA Administrator and a Certificate Manager. This does not work when Role Separation is enforced. |
| Enable role separation | | | | | X | |
| Issue and approve certificates | | X | | | | |
| Deny certificates | | X | | | | |
| Revoke certificates | | X | | | | |
| Reactivate certificates placed on hold | | X | | | | |
| Renew certificates | | X | | | | |
| Enable, publish, or configure CRL schedule | X | | | | | |

| | | | | | | |
|---|---|---|---|---|---|---|
| Recover archived key | | X | | | | Only a Certificate Manager can retrieve the encrypted key data structure from the database. The private key of a valid Key Recovery Agent is required to decrypt the key data structure and generate a PKCS#12 file. |
| Configure audit parameters | | | X | | | By default, the local Administrator holds the system audit user right. |
| Audit logs | | | X | | | By default, the local Administrator holds the system audit user right. |
| Back up system | | | | X | | By default, the local Administrator holds the system backup user right. |
| Restore system | | | | X | | By default, the local Administrator holds the system backup user right. |
| Read CA database | X | X | X | X | | By default, the local Administrator holds the system audit and backup user rights. |
| Read CA configuration information | X | X | X | X | | By default, the local Administrator holds the system audit and backup user rights. |

**Notes**

- Enrollees are allowed to read CA properties and certificate request lists (CRLs), and can request certificates. On an Enterprise CA, a user must have Read and Enroll permissions on the certificate template to request a certificate. For more information, see Allow subjects to request a certificate that is based on the template.
- CA Administrators, Certificate Managers, Auditor, and Backup Operators have implicit Read permissions.
- An Auditor holds the system audit user right.
- A Backup Operator holds the system backup user right. In addition, the Backup Operator has the ability to start and stop the Certificate Services service.

**Assigning roles**

The CA Administrator for a CA assigns users to the separate roles of role-based administration by giving each user the security settings required by a role. The CA Administrator can assign a user to more than one role, but the CA is more secure when each user belongs to one role only. When each CA role belongs to one user only, then there are fewer CA tasks that can be compromised if a user's account becomes compromised.

For more information on assigning roles, see Manage Role-based Administration.

**Administrator concerns**

The default installation setting for a stand-alone CA is to have members of the local Administrators security group as CA Administrators. The default installation setting for an Enterprise CA is to have local Administrators, Enterprise Admins, and Domain Admins as CA Administrators. To limit the power of any of these accounts, they should be removed from the CA Administrator and Certificate Manager roles once all CA roles are assigned.

As a best practice, group accounts that have been assigned CA Administrator or Certificate Manager roles should not be members of the local Administrators security group. Also, CA roles should only be assigned to group accounts and not individual user accounts.

**Note**

- Membership in the local Administrators group on the CA is required to renew the CA certificate. Members of this group are considered to be all powerful on the CA, with administrative authority over all other CA roles.

**Windows 2000 roles and Windows server 2003 role-based administration**

Windows 2000 server CA administration is changed in Windows Server 2003, Standard Edition CA role-based administration. Windows 2000 server administrators can perform any activity on a Windows 2000 server CA, but once CA roles are assigned using Windows server 2003 role-based administration, its administrators are subject to its roles. An administrator that could perform all tasks on a Windows 2000 server CA will only be able to perform the tasks associated with his or her role on the Windows server 2003 CA. After upgrading a Windows 2000 server CA to a Windows server 2003 CA, its administrators need to be assigned to the roles defined in the role-based administration for the Windows server 2003 CA.

**Certificate Services service level functionality**

Role-based administration applies to the Certificate Services service only. Any services other than Certificate Services service are unaffected by the assignment of CA roles.

**Role separation**

The separation of CA roles can be enforced using *role separation*. Once enforced, role separation allows a user to be assigned only a single role. If a user is assigned to more than one role and attempts to perform an operation on the CA, the operation is denied. For this reason, before role separation is enabled, a user should be assigned only one CA role.

This feature is valuable for large enterprises where the separation of roles ensures that the compromise of a user's account does not compromise the entire CA administered by the user. For more information, see Using role separation.

# Certification Authorities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certification authorities**
A certification authority (CA) is an entity entrusted to issue certificates to individuals, computers, or organizations that affirm the identity and other attributes of the certificate subject to other entities.

- Types of certification authorities
- Enterprise certification authorities
- Stand-alone certification authorities
- Qualified subordination

# Types of certification authorities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Types of certification authorities**

A certification authority (CA) accepts a certificate request, verifies the requester's information according to the policy of the CA, and then uses its private key to apply its digital signature to the certificate. The CA then issues the certificate to the subject of the certificate for use as a security credential within a public key infrastructure (PKI). A CA is also responsible for revoking certificates and publishing a certificate revocation list (CRL).

A CA can be an outside entity, such as VeriSign, or it can be a CA that you create for use by your organization by installing Microsoft Certificate Services. Each CA can have distinct proof-of-identity requirements for certificate requesters, such as a Windows Server 2003 family domain account, employee badge, driver's license, notarized request, or physical address. Identification checks such as this often warrant an onsite CA, so that organizations can validate their own employees or members.

Microsoft enterprise CAs use a person's user account credentials as proof of identity. In other words, if you are logged on to a Windows Server 2003 family domain and request a certificate from an enterprise CA, the CA knows that you are who the Active Directory directory service says you are.

Every CA also has a certificate to confirm its own identity, issued by another trusted CA or, in the case of root CAs, issued by itself. It is important to remember than anyone can create a CA. The real question revolves around whether you, as a user or an administrator, trust that CA and, by extension, the policies and procedures that CA has in place for confirming the identity of the entities issued certificates by that CA.

For more information about certificates, see Certificates overview.

**Root and subordinate certification authorities**

A *root CA*, sometimes called a root authority, is meant to be the most trusted type of CA in an organization's PKI. Typically, both the physical security and the certificate issuance policy of a root CA are more rigorous than those for subordinate CAs; if the root CA is compromised or issues a certificate to an unauthorized entity, then any certificate-based security in your organization is suddenly vulnerable. While root CAs can be used to issue certificates to end users for such tasks as sending secure e-mail, in most organizations they will only be used to issue certificates to other CAs, called subordinate CAs.

A *subordinate CA* is a CA that has been certified by another CA in your organization. Typically, a subordinate CA will issue certificates for specific uses, such as secure e-mail, Web-based authentication, or smart card authentication. Subordinate CAs can also issue certificates to other, more subordinate CAs. Together, a root CA, the subordinate CAs that have been certified by the root, and subordinate CAs that have been certified by other subordinate CAs form a certification hierarchy.

For more information about certification hierarchies, see Certification authority hierarchies.

**Enterprise and stand-alone certification authorities**

This version of Certificate Services supports the installation of stand-alone CAs and enterprise CAs. For information about the operational characteristics of enterprise CAs and stand-alone CAs, see Enterprise certification authorities and Stand-alone certification authorities.

# Enterprise certification authorities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enterprise certification authorities**

The Enterprise Administrator can install Certificate Services to create an enterprise certification authority (CA). Enterprise CAs can issue certificates for purposes such as digital signatures, secure e-mail using S/MIME (Secure Multipurpose Internet Mail Extensions), authentication to a secure Web server using Secure Sockets Layer (SSL) or Transport Layer Security (TLS), and logging on to a Windows Server 2003 family domain using a smart card.

An enterprise CA has the following features:

- An enterprise CA requires the Active Directory directory service.
- When you install an enterprise root CA, it uses Group Policy to propagate its certificate to the Trusted Root Certification Authorities certificate store for all users and computers in the domain. You must be a Domain Administrator or be an administrator with write access to Active Directory to install an enterprise root CA.
- Certificates can be issued for logging on to a Windows Server 2003 family domain using smart cards.
- The enterprise exit module publishes user certificates and the certificate revocation list (CRL) to Active Directory. In order to publish certificates to Active Directory, the server that the CA is installed on must be a member of the Certificate Publishers group. This is automatic for the domain the server is in, but the server must be delegated the proper security permissions to publish certificates in other domains. For more information about the exit module, see Policy and exit modules.

An enterprise CA uses certificate types, which are based on a certificate template. The following functionality is possible when you use certificate templates:

- Enterprise CAs enforce credential checks on users during certificate enrollment. Each certificate template has a security permission set in Active Directory that determines whether the certificate requester is authorized to receive the type of certificate they have requested.
- The certificate subject name can be generated automatically from the information in Active Directory or supplied explicitly by the requestor.
- The policy module adds a predefined list of certificate extensions to the issued certificate. The extensions are defined by the certificate template. This reduces the amount of information a certificate requester has to provide about the certificate and its intended use.

For more information, compare Stand-alone certification authorities. For more information about certificate templates, see Requesting certificate templates. For general information about CAs, see Certification Authorities.

# Stand-alone certification authorities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Stand-alone certification authorities**

You can install Certificate Services to create a stand-alone certification authority (CA). Stand-alone CAs can issue certificates for purposes such as digital signatures, secure e-mail using S/MIME (Secure Multipurpose Internet Mail Extensions), and authentication to a secure Web server using Secure Sockets Layer (SSL) or Transport Layer Security (TLS).

A stand-alone CA has the following characteristics:

- Unlike an enterprise CA, a stand-alone CA does not require the use of the Active Directory directory service. Stand-alone CAs are primarily intended to be used as Trusted Offline Root CAs in a CA hierarchy or when extranets and the Internet are involved. Additionally, if you want to use a custom policy module for a CA, you would first install a stand-alone CA and then replace the stand-alone policy module with your custom policy module.

- When submitting a certificate request to a stand-alone CA, a certificate requester must explicitly supply all identifying information about themselves and the type of certificate that is wanted in the certificate request. (This does not need to be done when submitting a request to an enterprise CA, since the enterprise user's information is already in Active Directory and the certificate type is described by a certificate template). The authentication information for requests is obtained from the local computer's Security Accounts Manager database.

- By default, all certificate requests sent to the stand-alone CA are set to Pending until the administrator of the stand-alone CA verifies the identity of the requester and approves the request. This is done for security reasons, because the certificate requester's credentials are not verified by the stand-alone CA.

- Certificate templates are not used.

- No certificates can be issued for logging on to a Windows Server 2003 family domain using smart cards, but other types of certificates can be issued and stored on a smart card.

- The administrator has to explicitly distribute the stand-alone CA's certificate to the domain user's trusted root store or users must perform that task themselves.

When a stand-alone CA uses Active Directory, it has these additional features:

- If a member of the Domain Administrators group or an administrator with write access to Active Directory, installs a stand-alone root CA, it is automatically added to the Trusted Root Certification Authorities certificate store for all users and computers in the domain. For this reason, if you install a stand-alone root CA in an Active Directory domain, you should *not* change the default action of the CA upon receiving certificate requests (which marks requests as Pending). Otherwise, you will have a trusted root CA that automatically issues certificates without verifying the identity of the certificate requester.

- If a stand-alone CA is installed by a member of the Domain Administrators group of the parent domain of a tree in the enterprise, or by an administrator with write access to Active Directory, then the stand-alone CA will publish its CA certificate and the certificate revocation list (CRL) to Active Directory.

For more information, compare Enterprise certification authorities. For general information about CAs, see Certification Authorities.

# Qualified subordination

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Qualified subordination**

Qualified subordination is an extension of standard certification authority (CA) subordination that allows you to:

- Define the namespaces for which a subordinate CA will issue certificates.
- Specify the acceptable uses of certificates issued by a qualified subordinate CA.
- Create trust between separate certification hierarchies.

For more information, see Qualified subordination overview.

**Qualified subordination properties**

Qualified subordination properties are defined in the CA certificate request for a subordinate CA. The CA certificate that is used to install a qualified subordinate CA contains extensions where the different qualified subordination properties are listed. The following table describes the extensions in the CA certificate that are used to install a qualified subordinate CA:

| Extension | Description |
|---|---|
| Name constraints | The range of namespaces that are permitted or excluded by the qualified subordinate CA and its subordinates. |
| Policies | The list of acceptable policies for a certificate, included in the certificates issued by the qualified subordinate CA. These policies are identified in a certificate by object identifiers (also known as OIDs) and include both Issuance policies and Application policies policies. The object identifiers that are used are enterprise-dependent (or dependent within the scope of the object identifiers you use). The object identifier for an application in one PKI deployment must be different from another. |
| Mapping policies between trust hierarchies | Determines how a policy object identifier from one trust hierarchy is mapped to a policy object identifier of another trust hierarchy. |
| Policy constraints | Defines the certification path length required and allowed for policy identifiers and policy mapping. |

All features of qualified subordinate CAs are new to the Windows Server 2003 family and are not available on Windows 2000 Server. To use these new features, you must use a Windows Server 2003 certification authority.

**Note**

- The constraints and policy types listed above are defined in RFC 2459. For more information on RFCs, see the RFC Editor Web site.

**Intranet and extranet qualified subordination**

Qualified subordination may be used in both intranet and extranet environments. Intranet-qualified subordination involves a parent CA issuing a CA certificate to a qualified subordinate CA within a company's intranet. For example, you might install a qualified subordinate CA to issue certificates for a specific network domain in your network. Extranet qualified subordination establishes trust between CAs in separate networks.

Extranet qualified subordination is also commonly called *cross-certification* because trust is being established across separate certification hierarchies. For Windows Server 2003 certification authorities, *cross-certification* is synonymous with *qualified subordination*.

**Name constraints**

Name constraints specify a range of permitted and excluded namespaces that a qualified subordinate CA and its subordinates will use when evaluating certificate requests. A user requesting a certificate from a qualified subordinate CA is identifiable by one or more of the namespaces required by the certificate template for which the request is made. Name constraints are defined using the namespaces used in your network, such as DNS domain names. By permitting or excluding a namespace, the qualified subordinate CA permits or exclude all subjects that are identifiable by that namespace from receiving one of its certificates.

Name constraints apply to subtrees of a namespace. A permitted name constraint defines the namespace within which the CA may issue certificates. An excluded name constraint is used to exclude a subspace of a permitted namespace. For example, a CA with the permitted name constraint example.microsoft.com may have an excluded name constraint of TestDomain.example.microsoft.com. In this case, the CA and its subordinates will issue certificates for users within the example.microsoft.com namespace but not for users within the TestDomain.example.microsoft.com namespace.

**Name constraint definition and usage**

When the CA certificate for the qualified subordinate CA is created, name constraints are defined to determine what namespaces the qualified subordinate CA will permit or exclude. When the installed qualified subordinate CA receives a certificate request from an entity such as a user, name constraints determine whether the qualified subordinate CA will issue the certificate.

When someone submits a certificate request for a qualified subordinate CA to the parent CA, the parent CA uses the certificate path to verify that the qualified subordinate CA will not permit names that are excluded by any other CA in the certification path. If the name constraints defined in the CA certificate request permit names that a CA in the certification path excludes, then the certificate request is denied. If the permitted namespaces defined in the CA certificate request do not conflict with the permitted namespaces of the CAs in the certification path, then the CA certificate request is successful.

For a certificate request that is submitted to a qualified subordinate CA to be successful, all namespaces identified with the user in the certificate request must be within permitted namespaces and none may be within an excluded namespace.

The following name forms are allowed as name constraints. The RFCs listed may be found at the RFC Editor Web site.

| Name form | Resource |
|---|---|
| Directory name (for example, Active Directory distinguished name) | Object names |
| DNS | RFC 1024 |
| E-mail | RFC 822 |
| URI/URL | RFC 1738 |
| User principal name (UPN) | For more information, see Active Directory naming. |
| IP address | IPv4 RFC 791; IPv6 RFC 1883 |

**Policies**

Certificate policies describe the policies according to which a certificate was issued and define how it may be used. There are two types of policies used in qualified subordination:

- Issuance policies
- Application policies

Both policy types can be included in all types of certificates, including CA certificates and user certificates. This availability is in contrast to name constraints, policy mapping, and policy constraints, which are only used in creating the CA certificate for a qualified subordinate CA.

The policies that are available to a certificate template are first defined in the CA certificate for the qualified subordinate CA that will use certificate templates. Each policy listed has a corresponding object identifier (also known as an OID) identifying the policy. The object identifiers used are restricted to that policy for the entire public key infrastructure (PKI). The object identifiers assigned to policies in one PKI may be mapped to the policies in a separate PKI to facilitate trust. For more information, see Mapping policies between trust hierarchies.

The issuance and application policies that are placed on a qualified subordinate CA must be within the set of policies object

identifiers defined for all the CAs in the certification path of the qualified subordinate CA. When a qualified CA certificate request is evaluated, the policy object identifiers contained in the request are evaluated against the policy object identifiers used in the entire certification path to determine that the policy object identifiers contained in the request are a subset of the object identifiers used in the path. For a qualified subordinate CA to use issuance policies, policies must be present at every CA in the certificate hierarchy between the root CA and itself. For this reason, policy cannot be removed from a certification hierarchy, but can only be constrained when installing a qualified subordinate CA.

**Issuance policies**
The issuance policies that are contained in a certificate identify the assurance level of the certificate's subject when that subject presents the certificate to another entity. *Assurance* describes the extent to which a CA trusts a subject's identification; for example, whether the user is recognized by domain authentication only or if the user has passed multiple authentication stages. Assurance levels are defined for a PKI and issuance policies are used to represent them in certificates.

When a subject uses a certificate that contains an issuance policy, security entities can identify the assurance level of that subject and react accordingly. The security credentials that these assurance levels represent depend on how the policies are defined in the PKI.

For more information, see Issuance policies.

**Application policies**
Application policies are included in a certificate because they are required by the applications that will receive these certificates. Some applications have policy requirements and maintain a list of the policies that they will accept. When an application receives signed information, it compares the application policy object identifiers in the received certificate to its list or required policies. For example, a remote access server may require that a Client Authentication application policy be present in a certificate that is presented by a client performing a remote logon.

Application policies are similar to the Extended Key Usage extension in a certificate, as both use one or more object identifiers to prescribe how the public key in a certificate must be used. Extended Key Usage is still supported by Windows Server 2003 CAs to provide for PKIs that use this extension, but application policies replace them.

For more information, see Application policies.

**Policy mapping**
Policy mapping allows you to map a certificate policy object identifier in one trust hierarchy to a policy object identifier used in another trust hierarchy in your intranet or extranet. The policy mappings are defined in the CA certificate request for a qualified subordinate CA. Each pair of policies that is configured on the qualified subordinate CA indicates that it considers its trust hierarchy's policies equivalent to those in the other trust hierarchy.

For an example of policy mapping, see Qualified subordination overview.

When mapping the policies used in an external trust hierarchy, the qualified subordinate CA that contains the policy mappings is said to *subordinate* the CA in the external trust hierarchy whose policies it has mapped. This subordination does not require the administrator of the external trust hierarchy's CA to perform any task. This allows you to set up a trust relationship with another trust hierarchy by subordinating a single subordinate CA or the root CA of the external PKI.

For more information, see Mapping policies between trust hierarchies.

**Policy constraints**
Policy constraints restrict certificate validation in two ways:

- They require that each CA certificate in a certification path contain a specific policy.
- They specify the number of CA certificates that are permitted in a certification path before policy mapping is no longer allowed.

Policy constraints are applied to both issuance and application policies. For more information, see Policy constraints.

**Unexpected trust**
When mapping policies between trust hierarchies in two different companies, there could be a problem of unexpected trust if a third company maps its policy object identifier to the mapped policies of one of the two companies. In this case, the third company is trusted by both of the other companies, which may not be wanted. In addition, either of the companies could make modifications to their certification hierarchies that would have the same effect. Policy constraints resolve this unexpected trust by restricting the certification path depth in which policy mapping may be used.

For more information, see Mapping policies between trust hierarchies; Policy constraints

**Notes**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.
- Qualified subordination is performed by the parent CA on any subordinate CA and can therefore by applied to any type of CA, such as Windows NT 4.0, Windows 2000 Server, or other CAs.
- Qualified subordinate CA certificate renewal is the same as standard subordinate CA renewal.
- When qualified subordinate CA certificates and cross-certificates are revoked, they should also be deleted from the Active Directory directory service. This deletion is not required for security reasons, but is considered a best practice, because it also removes the certificates from the Local Machine certificates store of client computers.

# Using Certificate Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Certificate Services**
This section covers:

- Installing and configuring a certification authority
- Configuring the policy and exit modules
- Establishing a certification hierarchy
- Backing up and restoring a certification authority
- Revoking certificates and publishing CRLs
- Mapping certificates to user accounts
- Renewing a certification authority
- Managing Certificate Services from a command prompt
- Certificate Services Example Implementations
- Using Qualified Subordination

# Installing and configuring a certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing and configuring a certification authority**
This section will give you additional information you need to install and configure a certification authority (CA).

**Planning the installation of a certification authority**
Before installing Certificate Services, you should plan the deployment of certification authorities (CAs) and Deploying a Public Key Infrastructure in your organization.

For more information about deploying a PKI, see Certificates Resources in public key infrastructure and Microsoft Windows Deployment and Resource Kits.

**Ways to install Certificate Services to create a certification authority**
If you are installing Certificate Services, there are a number of situations in which you may want to set up a CA. (The most typical is listed first here.)

- **After base setup has completed**:

  To install Certificate Services on a server that already has Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition installed, open **Add or Remove Programs** in **Control Panel**. (To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.) After you select Certificate Services for installation, the Certificate Services Installation Wizard guides you through the installation process.

  For the procedures to set up a certification authority (CA) as part of the base setup, see Set Up a Certification Authority.

- **As part of** Windows Server 2003 **family base setup**:

  Although Certificate Services is a service and is included with the Windows Server 2003 family, it is not installed as part of the initial installation process by default. (You do not want every file server to be a CA.)

  To install Certificate Services during the initial base installation of the Windows Server 2003 family, you must select it from the optional components list that is displayed during setup. Certificate Services will not actually be installed until you log on to the server after setup has completed. Then, a message will prompt you to complete the setup of the CA.

  This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Setup options and information**
When you set up Certificate Services, you will be prompted for the following information:

**Certification authority type selection**
During the Certificate Services installation, you can choose to set up any of the following types of certification authority (CA):

| Certification Authority Type | Description |
|---|---|
| Enterprise root CA | An enterprise root CA is a top-level CA in a certification hierarchy. An enterprise root CA requires the Active Directory directory service. It self-signs its own CA certificate and uses Group Policy to publish that certificate to the Trusted Root Certification Authorities store of all servers and workstations in the domain. Normally, an enterprise root CA does not directly provide resources for user and computer certificates, but is the foundation for a certificate hierarchy. For more information, see Enterprise certification authorities |

| | |
|---|---|
| Enterprise subordinate CA | An enterprise subordinate CA must obtain its CA certificate from another CA. An enterprise subordinate CA requires Active Directory. You use enterprise subordinate CAs when you want to take advantage of Active Directory, certificate templates, and smart card logon to Windows XP and computers running Windows Server 2003 family operating systems |
| Stand-alone root CA | A stand-alone root CA is a top-level CA in a certification hierarchy. The stand-alone root CA may or may not be a member of a domain and, therefore, does not require Active Directory. However, it will use Active Directory if it exists for publishing certificates and certificate revocation lists. Since a stand-alone root CA does not require Active Directory, it can easily be disconnected from the network and placed in a secure area, which is useful when creating a secure offline root CA. For more information, see Stand-alone certification authorities |
| Stand-alone subordinate CA | A stand-alone subordinate CA must obtain its CA certificate from another CA. The stand-alone subordinate CA may or may not be a member of a domain and, therefore, does not require Active Directory. However, it will use Active Directory if it exists for publishing certificates and certificate revocation lists. It must obtain its CA certificate from another CA. |

**Note**

- Certificate Services should not be installed on any node in a server cluster as this configuration may not allow Certificate Services to run properly and is therefore not supported.

**Installing Web enrollment support**

You may want to install a certification authority infrastructure within a defined security boundary and only allow specific computers to communicate with your certification authority. This means clients cannot connect directly to a CA to request and retrieve certificates. Or you may want to set up a CA that supplies certificates to Internet-based subjects. You may want to place a layer of insulation between these clients and the CA. Web enrollment allows these scenarios to be implemented easily.

Web enrollment is the component of a Microsoft CA that allows clients to submit certificate issuance and retrieval requests based on Web pages. It is installed when you install a Microsoft CA. It can also be installed on a non-CA computer to provide a Web front-end to your CA infrastructure on the back-end. In this configuration, clients make requests with the Web pages and only the computer that has Web enrollment installed needs to communicate with your CA.

For more information on installing Web enrollment, see Set up certification authority Web enrollment support.

**Public and private key pair: cryptographic service providers, key lengths, hash algorithms**

If you select the **Use custom settings to generate the key pair and CA certificate** check box when you choose the type of CA to install, you can select the cryptographic service provider (CSP) to use. The CSP generates the public key and private key pair and performs cryptographic operations on behalf of the CA.

You can set the key length for the public key cryptographic keys that the CA uses to sign certificates. In general, the longer the key length, the more secure it is. Note that a longer key will take more time to generate during setup and longer to transmit on the network when building certificate trust lists.

You can also choose the message hash algorithm used by the CA, as well as specifying the use of existing cryptographic keys instead of generating new ones.

**Note**

- A CA may be installed using a smart card CSP. Once a CA is installed using a smart card CSP, the smart card must be available and inserted with the personal identification number (PIN) in order for the CA to start and perform its operations. This is a higher security option for CA operations than other options, but a disadvantage of using smart card CSPs is that CA performance may be slower and key length may be smaller.

**Certification authority identifying information**

The following are some guidelines for completing the CA identifying information in Certificate Services setup:

| Field | Description |
|---|---|

| | |
|---|---|
| CA Name | The name you want to give to the CA. You can enter a string using any Unicode character. The name of the CA will also be the common name of the CA's distinguished name in Active Directory.<br><br>When special characters exist in the CA name, a sanitized CA name is used for operations that are unable to use the unmodified CA name. A CA's sanitized name is the name of the CA with all special characters encoded in a form that will allow them to be used for file names, CryptoAPI key container names, and Active Directory object names. Special characters are those characters that cannot be used in one or more of these names; the list includes all characters which are not ASCII characters and many ASCII punctuation characters.<br><br>Further, Active Directory object names are limited to 64 characters by the Lightweight Directory Access Protocol (LDAP) standard. To accommodate this limit, Active Directory object names are constructed by truncating the sanitized name and appending a hash computed over the truncated part of the sanitized name. The distinguished name suffix field is automatically filled with the distinguished name of the Active Directory domain. If you edit this distinguished name, you must conform with the LDAP standard.<br><br>Type **certutil.exe** at a command prompt without arguments to see the sanitized name for all of the published CAs. Type **certutil.exe -v -ds** to see all of the CA-related Active Directory names. The first column is a truncated CA name with the hash appended (the actual Active Directory object's container name, with special characters reverted back to their original form). A second column is displayed only if the truncated, sanitized name does not match the truncated CA name, and it is the actual Active Directory object's container name. |
| Distinguished name suffix | The X.500 distinguished name suffix that will be appended to the CA name. |
| Validity period | The length of time the CA's certificate should be valid. The CA will obtain a certificate that is valid for this length of time and use that certificate's private key for signing issued certificates and certificate revocation lists. |

**Database and configuration storage**

Certificate Services uses local storage for its database, configuration data, backup data and logging data. You can specify locations for the database and log during CA setup. By default, the certificates issued by a CA are stored in:

*Systemroot*\system32\certlog

For best performance, the database and log files should be kept on separate physical disk drives, preferably on separate disk controllers. This will maximize disk throughput and allow the CA to perform better.

You also have the option of specifying a shared folder when setting up a CA. The shared folder acts as a location where computer users can find information about certification authorities. This option is only useful if you are installing a stand-alone CA and do not use Active Directory.

If the host computer for Certificate Services is a member of a domain, information about the CA is automatically published to the Active Directory. However, the Active Directory does not act as the database for Certificate Services. This function remains with the local computer.

### Note

- It is considered unsafe to install Certificate Services on an FAT file system because FAT does not support domain-based security. Always install Certificate Services on a NTFS file system to take advantage of support for a variety of features including Active Directory, which is needed for domains, user accounts, and other important security features.

**(Optional) Creating an issuer policy statement for the CA**

When you set up a CA, you can add a CA policy statement to the CA certificate that is created during setup or a CA certificate renewal, in the form of text or a pointer to a Web site. The CA policy statement gives legal and other pertinent information about the CA and its issuing policies, limitations of liability, and so on. An end user will see this CA policy statement when they view the CA certificate and click **Issuer Statement**.

Here are the steps you need to follow to attach a policy statement to a CA's certificate:

- The policy statement file must be installed *before* you set up Certificate Services. This file, named CAPolicy.inf, must be placed in the systemroot directory. URLs in CAPolicy.inf should use the replaceable parameter syntax that also appears in the **Extensions** tab. If you don't use the replaceable parameter syntax, then CDP and AIA extensions in renewed CA root certificates may point to the same location that was specified in the previous CA root certificate. To see the syntax for replaceable parameters, see Specify certificate revocation list distribution points in issued certificates.

- The first two lines of CAPolicy.inf must be:

  **[Version]**

  **Signature="$Windows NT$"**

- The next two lines of the file list the name of the policies for this CA. Multiple policies can be listed in **Policies=** if they are separated by commas. The name *LegalPolicy* is used here as an example, but the name can be whatever the CA administrator chooses when creating CAPolicy.inf:

  **[CAPolicy]**

  **Policies=***LegalPolicy*

- For each policy, you need to provide a user-defined object identifier (also known as an OID) and either the text you want displayed as the policy statement or a URL pointer to the policy statement. (The URL can be in the form of an HTTP, FTP, file, or LDAP address.) Continuing on with the example, if you are going to have text in the policy statement, the next three lines of CAPolicy.inf will read:

  **[***LegalPolicy***]**

  **OID=***1.1.1.1.1.1.1.1.1*

  **Notice="***Legal policy statement text***"**

  If you are going use a URL to host the CA policy statement, the next three lines would instead read:

  **[***LegalPolicy***]**

  **OID=***1.1.1.1.1.1.1.1.1*

  **URL="http://***CompanyWebSite/CAPolicy/default.asp***"**

  (Please note that the "OID=" entry used in this example is arbitrary and is shown only for illustrative purposes.)

  In addition:

  - Multiple URLs are supported
  - Multiple notices are supported
  - Mixed notices and URLs are supported
  - URLs with spaces or text with spaces must be surrounded by quotes

  An example of multiple notices and URLs in a policy section might be:

  **[***LegalPolicy***]**

  **OID=***1.1.1.1.1.1.1.1.1*

  **URL = "http://***http.microsoft.com/somewhere/default.asp* **"**

  **URL = "ftp://***ftp.microsoft.com/somewhere else/default.asp***"**

  **Notice = "***Legal policy statement text.***"**

- You can specify CRL Distribution Points (CDPs) in CAPolicy.inf. Note that any CDP in CAPolicy.inf will take precedence for certificate verifiers over the CDP's specified in the CA policy module. If you want to specify the CDP using CAPolicy.inf, a sample of the syntax is:

  **[CRLDistributionPoint]**

  **URL="http://***CompanyWebSite/Public/myCA.crl***"**

  Some additional notes about the CRL section:

  - Multiple URLs are supported.
  - HTTP, file, FTP, and LDAP URLs are supported.
  - This section can be used during CA setup or CA certificate renewal.
  - This section is only used if you are setting up a root CA or renewing a root CA certificate. Subordinate CA CRL distribution point extensions are determined by the CA which issued the subordinate CA's certificate.
  - URLs with spaces must be surrounded by quotes.
  - If no URLs are specified (in other words, if the **[CRLDistributionPoint]** section is empty), the default CRL Distribution Point extension will be suppressed.

- You can specify the authority information access points in CAPolicy.inf. The syntax is:

  **[AuthorityInformationAccess]**

  **URL="http://***CompanyWebSite/Public/myCA.crl***"**

  Some additional notes about the authority information access section:

  - Multiple URLs are supported.
  - HTTP, file, FTP, and LDAP URLs are supported.
  - This section can be used during CA setup or CA certificate renewal.
  - This section is only used if you are setting up a root CA or renewing a root CA certificate. Subordinate CA authority information access extensions are determined by the CA which issued the subordinate CA's certificate.
  - If no URLs are specified (in other words, if the **[AuthorityInformationAccess]** section is empty), the default Authority Information Access extension will be suppressed.

- Another optional section of CApolicy.inf is **[EnhancedKeyUsage]**, which is used to specify Enhanced Key Usage extension object identifiers (also known as OIDs).

  - Multiple object identifiers are supported.
  - This section can be used during CA setup or CA certificate renewal.
  - This section is only used if you are setting up a root CA or renewing a root CA certificate. The Enhanced Key Usage extension for a subordinate CA is determined by the CA which issued the subordinate CA's certificate.

  An example of this section is:

  **[EnhancedKeyUsage]**

**OID=***1.2.3.4.5*

**OID=***1.2.3.4.6*

- Another optional section of CApolicy.inf is **[certsrv_server]**, which is used to specify renewal key length, the renewal validity period, and the certificate revocation list (CRL) validity period for a CA that is being renewed or installed.

  An example would be:

  **[certsrv_server]**

  **RenewalKeyLength=***2048*

  **RenewalValidityPeriod=***Years*

  **RenewalValidityPeriodUnits=***5*

  **CRLPeriod=***Days*

  **CRLPeriodUnits=***2*

  **CRLDeltaPeriod=***Hours*

  **CRLDeltaPeriodUnits=***4*

  - **RenewalKeyLength** sets key size for renewal only. This is only used when CA renewal is generating new keys.
  - **RenewalValidityPeriod** and **RenewalValidityPeriodUnits** establish the lifetime of new root CA certificate when renewing the old CA certificate.
  - **CRLPeriod** and **CRLPeriodUnits** establish the validity period for the full CRL, while **CRLDeltaPeriod** and **CRLDeltaPeriodUnits** establish the validity period for the delta CRL.
  - **RenewalKeyLength**, **RenewalValidityPeriod** and **RenewalValidityPeriodUnits** are only processed when the CA is being renewed.
  - **LoadDefaultTemplates** is not processed when the CA is a subordinate CA.

# Configuring the policy and exit modules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the policy and exit modules**
The administrator of a certification authority (CA) can configure settings in the default policy and exit modules provided with Certificate Services by using the Certification Authority snap-in.

**Policy module setting you can change:**

- **The default action of the certification authority upon receiving a valid certificate request.** You can specify whether a stand-alone CA will hold incoming certificate requests as Pending or automatically issue the certificate. In most cases, for security reasons, it is recommended that all incoming certificate requests to a stand-alone CA be marked as Pending.

- For the procedure to change the default action of a certification authority upon receipt of a certificate request, see Set the default action upon receipt of a certificate request.

**Exit module settings you can change:**

- **Allow certificate publication to the file system.** You can select whether to allow the publishing of certificates to the file system. Actual publication will only occur if the certificate request specifies a file system location where the certificate is to be published.

- For the procedure to allow or disallow the publishing of certificates to the file system, see Publish certificates to the file system.

- **Sends e-mail when a certification event occurs.** You can configure the CA to send e-mail when a certification event occurs, such as the issuance of a certificate or when a certificate request is set to pending.

- To configure e-mail sending options, see Send e-mail when a certification event occurs.

For more information about policy and exit modules, see Policy and exit modules.

If you want to replace the policy module provided with Certificate Services with a custom policy module or a policy module developed for Certificate Server 1.0 and Windows NT 4.0, you must first register the policy module .dll file by typing **rgsrv32 *yourmodule*.dll** at a command prompt and then follow the procedure in Select a different policy module.

# Establishing a certification hierarchy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing a certification hierarchy**
The first step to establish a certification hierarchy is that a Domain Admin installs a root certification authority (CA).

**Using a Microsoft certification authority as your root authority**
The installation process for a Certificate Services root authority generates a root CA certificate containing the CA's public key and the digital signature created using the root's private key. If the root authority is installed using Certificate Services on a server that has access to the Active Directory directory service, the root authority's certificate will automatically be placed in all users' Trusted Root Certification Authorities certificate store, thereby establishing forest-wide trust in the root CA.

**Using a certification authority outside of your organization as your root authority**
If your organization is using a non-Microsoft certification authority from outside your organization as the root authority, you will need to obtain the root certificate and distribute it to any user and computer that needs to establish trust in the root authority. One way to distribute this type of root certificate to Windows XP and Windows Server 2003 family users is to use a certificate trust list (CTL) by way of Group Policy. For more information about distributing third party root certificates using Group Policy, see Enterprise trust policy.

**Using a non-Microsoft certification authority inside your organization as your root authority**
If your organization is using its own non-Microsoft certification authority (CA) as the root authority, you will need to obtain the root certificate and distribute it to any user and computer that needs to establish trust in the CA root authority. One way to distribute a root certificate to Windows XP and Windows Server 2003 family users is to use the Trusted Root Certification Authorities policy setting via Group Policy. For more information about distributing non-Microsoft root certificates using Group Policy, see Policies to establish trust of root certification authorities.

**Subordinate certification authorities**
After trust in a root authority has been established, you can install certification authorities (CAs) that are subordinate to the root CA as well as installing subordinate CAs that are subordinate to other subordinate CAs. By doing this, you can create a chain of parent-child relationships between CAs that serve different functions in an organization's public key infrastructure (PKI). The only significant difference in the installation process between a root CA and a subordinate CA is that a certificate request is generated for submission to another CA by a subordinate CA instead of creating a self-signed certificate. This request may be routed automatically to online CAs located via Active Directory, or routed manually, if offline. In either case, the resulting certificate must be installed on the new subordinate CA before it can begin operation.

Note that there is a relationship between the enterprise CAs and the Windows Server 2003 family domain trust model, but this does not imply a direct mapping between CA trust relationships and domain trust relationships. There is nothing that prevents a single CA from servicing entities in multiple domains or even entities outside the domain boundary. Similarly, a given domain may have multiple enterprise and root CAs.

For a brief overview of certification hierarchies, see Certification authority hierarchies.

For extensive information about planning certification hierarchies, refer to the Microsoft Windows Deployment and Resource Kits.

# Backing up and restoring a certification authority

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backing up and restoring a certification authority**

The purpose of backup and restore operations is to protect the certification authority (CA) and its operational data from accidental loss due to hardware or storage media failure. The recommended method to back up a CA is to use Backup (included with the operating system) to back up the entire server, including the system state which contains the CA's data.

| **Warning** |
| --- |
| The procedures discussed in this article demonstrate how to back up only the certification authority (CA) database and private key. If you want to back up the CA database, private key, and CA configuration, see the backup procedures in AD CS Migration: Migrating the Certification Authority (http://technet.microsoft.com/library/ee126140.aspx). To back up all the configuration data for a computer, run a System State Data backup (http://technet.microsoft.com/library/cc781353.aspx). |

It is also possible to back up and restore a CA using the Certification Authority snap-in, but this backup method is intended for use only in special cases where you do not want to back up the entire server on which the CA is installed. Using the Certification Authority snap-in, you can back up and restore the following types of information:

- Private key, and the certificate that the CA uses for digitally signing
- Certificate database

The public key and private key are backed up or restored using the PKCS #12 PFX format.

The Backup or Restore Wizard will ask you to supply a password when backing up the public and private keys and CA certificate. This password will be needed to restore the CA. For procedures on backing up a CA using the Certification Authority snap-in, see Back up a certification authority.

After performing the initial full backup of the CA, you can do incremental backups from that point on. When restoring, you will need to restore the full backup first and then each incremental backup in the order that they were created. For procedures on restoring a CA using the Certification Authority snap-in, see Restore a certification authority from a backup copy.

**Important**

- In general, you should use **Backup** to back up and restore both the CA and the server. For more information about backing up a Windows Server 2003 family computer, see Backup.

Upon restoring a CA, the Internet Information Services (IIS) metabase must also be restored if it has been damaged or lost. If a damaged or missing IIS metabase is not restored, IIS will fail to start, and that will result in Certificate Services Web pages failing to load. The IIS snap-in is used to back up the IIS metabase. Backup should be used to back up the IIS Web content pages and the CA. An alternative method is to recreate the IIS metabase and then use the **certutil.exe -vroot** command at a command line to reconfigure the IIS server to support the CA Web pages.

When restoring a CA, if the database logs are not manually deleted before the restore, the CA will be restored to the point in time of the restore--the database logs will be replayed, and changes made since the last backup will be applied to the database. (The default location of the database logs is *systemroot*\system32\certlog). If the database logs are manually deleted before the restore, the CA will be restored to the point in time that the backup was performed.

For more information on backup, restore, and verification of a CA backup, see Certutil tasks for backing up and restoring certificates.

# Revoking certificates and publishing CRLs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Revoking certificates and publishing CRLs**

You can use the Certification Authority snap-in to revoke a certificate, to administer certificate revocation list (CRL) publication, and to specify the CRL Distribution Points (CDPs) published in every certificate issued by the certification authority (CA).

**Revoking a certificate**

To help maintain the integrity of an organization's public key infrastructure (PKI), the administrator of a CA must revoke a certificate if the subject of certificate leaves the organization, if the certificate subject's private key has been compromised, or if some other security-related event dictates that it is no longer desirable to have a certificate considered "valid." When a certificate is revoked by a CA, it is added to that CA's certificate revocation list (CRL). This may take the form of a new CRL or a delta CRL, which is a small CRL that lists the certificates that have been revoked since the last full CRL.

For conceptual information about the role of certificate revocation and revocation checking in a PKI, see Certificate revocation. For the procedure to revoke a certificate, see Revoke an issued certificate. For instructions on how to display the current CRL, see View the certificate revocation list.

**Scheduling certificate revocation list (CRL) publication**

One of the features of Certificate Services is that every CA automatically publishes an updated CRL after an interval of time specified by the administrator of the CA. This interval of time is known as the CRL "publish period." After the initial setup of a CA, the CRL publish period is set to one week (based upon the local computer's time, starting from the date when the CA is first installed). For the procedure to change a certification authority's CRL publishing interval, see Schedule the publication of the certificate revocation list. The CRL and delta CRL publish periods can be scheduled independently.

A CA administrator should understand that there is a difference between a CRL publish period and the validity period of a CRL. The validity period of a CRL is the period of time that the CRL is considered authoritative by a verifier of a certificate. As long as the verifier of a certificate has a valid CRL in its local cache, it will not attempt to retrieve another CRL from the CA which publishes it.

The publish period of a CRL is established by the CA administrator. However, the validity period of the CRL is extended from the publish period to allow for Active Directory replication. By default, Certificate Services extends the publish period by 10% (up to a maximum of 12 hrs) to establish the validity period. So, for example, if a CA is publishing a CRL every 24 hours, the validity period is set to 26.4 hours.

Additionally, there is a clock skew of an additional 10 minutes added to the validity period on either side of the publish period, so a CRL will be valid 10 minutes before the beginning of its publish period to account for variances in computer clock settings.

There are registry entries which allow an administrator to control the variance between publish period and validity period to allow for slower directory replication. For information about these registry entries, see the Microsoft Windows Deployment and Resource Kits.

When determining the schedule for CRL publishing, there is a trade-off between performance and security. The more often CRLs are published, the more current clients can be with the list of revoked CRLs and the less likely they are to accept recently-revoked certificates. However, frequent publishing can negatively impact network and client performance. To help mitigate this trade off, delta CRLs can be used if the environment supports them.

**Using delta certificate revocation lists**

CRLs can become very long on large CAs that have experienced significant amounts of certificate revocation. This can become a burden for clients to download frequently. To help minimize frequent downloads of lengthy CRLs, delta CRLs can be published. This allows the client to download the most current delta CRL and combine that with the most current base CRL to have a complete list of revoked certificates. Because the client will normally have the CRL cached locally, the use of delta CRLs can potentially improve performance.

To use delta CRLs, the client application must be aware of and explicitly use delta CRLs for revocation checking. If the client does not use delta CRLs, it will retrieve the CRL from the CA every time it refreshes its cache, regardless of whether a delta CRL exists or not. For this reason, you should verify that the intended applications use delta CRLs and configure the CA accordingly. If the clients do not support the use of delta CRLs, you should either not configure the CA to publish delta CRLs or configure it so CRLs and delta CRLs are published at the same interval. This would still allow future applications that support delta CRLs to use them, while providing current CRLs to all applications. Note that all applications that use CryptoAPI in Windows XP and the Windows Server 2003 family use delta CRLs.

**Publishing a CRL before the next scheduled publish period**
You can also publish a CRL on demand at any time, such as when a valuable certificate becomes compromised. Choosing to publish a CRL outside the established schedule resets the scheduled publication period to begin at that time. In other words, if you manually publish a CRL in the middle of a scheduled publish period, the CRL publish period is restarted.

It is important to realize that clients that have a cached copy of the previously published CRL will continue using it until its validity period has expired, even though a new CRL has been published. Manually publishing a CRL does not affect cached copies of CRLs that are still valid; it only makes a new CRL available for systems that do not have a cached copy of a valid CRL.

For the procedure to force the immediate publication of a CRL or delta CRL, see Manually publish the certificate revocation list.

**CRL distribution points and the CRL file name**
Every certificate that is issued by a Microsoft certification authority has CRL distribution points as part of its content. A CRL distribution point provides a certificate verifier with the network location where it can retrieve the current copy of the CRL or delta CRL. For procedures to designate the certificate revocation list distribution points in certificates, see Specify certificate revocation list distribution points in issued certificates.

By default, CRL and delta CRL files are published on the CA in the following location:

*Systemroot*\System32\Certsrv\Certenroll

The format of the CRL file name is the "sanitized name" of the CA plus, in parentheses, the "key id" of the CA (if the CA certificate has been renewed with a new key) and a .crl extension. Refer to the table for some sample CRL file names based on sample renewal histories of a CA:

| Scenario | Name of CRL file |
|---|---|
| A certification authority named "MyCA" that has never had its CA certificate renewed | myca.crl |
| A certification authority named "MyCA" that has been renewed once with the same key | myca.crl |
| A delta CRL for a certification authority named "MyCA" that has been renewed once with the same key | myca+.crl |
| A certification authority named "MyCA" that has been renewed once with a new key | myca(1).crl |
| A certification authority named "MyCA" that has been renewed twice with a new key | myca(2).crl |

See Installing and configuring a certification authority for information about a CA's "sanitized name."

See Renewing a certification authority for more information about CA renewal.

**Enabling Netscape-compatible Web-based revocation checking**
To enable Netscape-compatible Web-based revocation check extensions to be added to every certificate, run the following **certutil** command from the command prompt on the certification authority:

**certutil -SetReg Policy\RevocationType +AspEnable**

Then, stop and start the Certification Authority service. Certificates issued by the certification authority after it is restarted will contain the extension.

**Important**

- Because clients that use this type of revocation checking will access the CA through Web pages, Internet Information Services must have access to read the revocation information from the CA. If IIS does not have at least read-only access to the CA, access is denied by the CA and the revocation information will not be available.

For more information on the **certutil** command, see Certutil.

# Mapping certificates to user accounts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Mapping certificates to user accounts**

It is possible to *map* (or create an association from) a certificate that has been issued to a user to the user's account. A server application can then use public key cryptography technology to authenticate the user using this certificate. If the user is authenticated, then the user's account is logged on. The end result is the same as if the user provided a user ID and password, yet the process is much more manageable.

Traditionally, computer systems have used a centralized accounts database to manage users, their user rights, and their access controls. This technique has worked well and is well understood. However, as systems become more and more distributed--with hundreds of thousands to millions of users--this form of centralized control becomes unwieldy. The problems range from trying to verify an account against a database located on the other side of the Internet to administering a lengthy list of users.

Public key certificates can help simplify these problems. Certificates can be widely distributed, issued by numerous parties, and can be verified by simply examining the certificate, without having to refer to a centralized database. However, existing operating systems and administration tools can only deal with accounts, not certificates. The simple solution--one that maintains the advantages of both certificates and user accounts-- is to create a mapping between a certificate and a user account. This allows the operating system to continue using accounts while the larger "system" and the user use certificates.

In this model, when a user presents a certificate, the system looks at the mapping to determine which user account should be logged on. (Note that this should not be confused with logging on with a smart card. Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition support logging on with a smart card using account mapping that is automatic.)

**Types of mapping**

In most cases, a certificate is mapped to a user account in one of two ways: a single certificate is mapped to a single user account (one-to-one mapping) or multiple certificates are mapped to one user account (many-to-one mapping).

**User principal name mapping**

User principal name mapping is a special case of one-to-one mapping. To use user principal name mapping, you must use the Active Directory directory service. With user principal name mapping, the user principal name is used to find the user's account in Active Directory and log it onto the network or host. The user principal name looks very much like an e-mail name, and is unique within a Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition domain. Enterprise certification authorities (CAs) place the user principal name of the certificate holder into each certificate. Thus, for accessing a secure IIS server or logging on to Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition with a smart card, the mapping of user names to accounts is automatic on these certificates.

For more information on using Active Directory for certificate mapping, see Map a certificate to a user account.

**One-to-one mapping**

One-to-one mapping maps a single user certificate to a single user account. For example, imagine you want to provide a Web page to your employees that will allow them to view and modify their deductions, manage their health care, and a number of other benefits options. This page should work over the Internet and should be secure. As a solution, you decide to use certificates and certificate mapping on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. You can either issue certificates to each of your employees from your own certificate service, or you can have your employees get certificates from a certification authority approved by your company. You can then take these user certificates and map them to the employee's user account. This allows a user to connect to the Web page using SSL (Secure Sockets Layer) or TLS (Transport Layer Security) from anywhere by providing his or her client certificate. The user then logs onto his or her own user account and normal access controls can be applied.

**Many-to-one**

Many-to-one mapping maps many certificates to a single user account. For example, you have a partnership with an agency that provides temporary workers for your job openings. You would like to allow the agency personnel to view Web pages that describe current job openings that only company employees can see. The agency has its own certification authority that it uses to issue certificates to its employees. After installing the agency certification authority's root certificate as a trusted root in your enterprise, you can set a rule that maps all certificates issued by that certification authority to a single Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition account. You can then set the access rights of the account so this account can access that Web page.

# Renewing a certification authority

Updated: November 20, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Renewing certification authorities**

**Why certification authorities need to be renewed**
A certification authority may need to be renewed for either of the following reasons:

- Change in the policy of certificates issued by the CA
- Expiration of the CA's issuing certificate

**Certification authority validity periods**
Every certificate issued by a certification authority (CA) has a validity period. The validity period is the range of time where the certificate can be accepted as an authoritative credential of the identity of the subject of the certificate. This is, of course, assuming the certificate is not revoked before the validity period ends and that the issuing CA is trusted. The primary purpose of the validity period is to limit the time in which a certificate is exposed to the possibility of being compromised.

Since a CA is really just another entity that has been issued a certificate--either issued by itself (in the case of a root CA) or issued by a parent (in the case of a subordinate CA)--every CA has a built-in expiration date, as determined by the end of the validity period on its CA certificate. This is not meant to imply that the lifetime of a CA is equivalent to the validity period of its CA certificate, only that a CA cannot issue certificates if it does not have a valid certificate of its own. The lifetime of a CA includes the validity periods of all its CA certificates--past and present. With these considerations in mind, an organization with a public key infrastructure (PKI) has to plan for the renewal of every certificate that is issued to a CA in the certification hierarchy in order to maintain the existing trust relationships in the PKI and to extend the lifetimes of CAs.

**The validity period of a CA and the validity period of the certificates it issues**
Before discussing the renewal of a CA certificate, it is helpful to understand how the end of the validity period of a CA affects the validity period of the certificates it issues. Certificate Services enforces a rule that a CA never issues a certificate to be valid beyond the expiration date of its own certificate. Therefore, when a CA's certificate reaches the end of its validity period, all certificates it has issued will also expire. This way, if the CA is purposely not renewed and the CA reaches the end of its lifetime, a PKI administrator can be assured that all the certificates that the now-expired CA has issued can no longer be used as valid security credentials. In other words, there will be no "orphaned" certificates that are still within their validity period but which have been issued by a CA that is no longer valid.

To illustrate how a decreasing validity period works, consider the following scenario: An organization installs a root CA with a certificate validity period of five years. The organization then uses this root CA to issue certificates with a validity period of two years to subordinate CAs. For the first three years every certificate issued to a subordinate CA by the root CA will continue to have a validity period of two years. After three years, when there is less than two years left in the validity period of the root CA certificate, Certificate Services begins to reduce the validity period of the certificates issued by the root CA so that they do not exceed the end of the CA's certificate's expiration date. Therefore, after four years, the CA issues subordinate CA certificates that are valid for one year. After 4.5 years, issued subordinate CA certificates have a validity period of only six months.

Certificate Services allows for the following maximum validity periods that are based on the type of certificate. These validity periods are configurable by the Domain Admin or Enterprise Admin when using version 2 certificate templates:

| Type of certificate | Maximum validity period |
|---|---|
| Root CA | Specified during setup of Certificate Services |
| Subordinate CA, Internet Protocol Security, Enrollment Agent, Domain Controller | Up to 5 years, never more than the Root CA's validity period |
| All other certificates | 1 year, never more than the Issuing CA's validity period |

**Planning for the renewal of a CA**
Since a CA that is approaching the end of its own validity period issues certificates valid for shorter and shorter periods of time,

you need to have a plan in place to renew the CA well before it expires in order to avoid issuing certificates of a very short validity period.

One renewal consideration: When you renew a CA, you have the option of reusing its existing key pair or generating a new key pair. When you generate a new key pair for a CA that is being renewed, a new certificate revocation list (CRL) distribution point is also created. This is to ensure that the key used to sign a certificate issued by the CA also matches the key used to sign the CRL. For more information about how renewing a CA with a new key affects certificate revocation and the name of CRLs, see Revoking certificates and publishing CRLs.

The following strategies are presented as possible approaches an organization could take when planning for CA renewal. You may need to adapt these to your specific situation.

- **Renewing a root certification authority**

  Since the root CA is the most critical CA in the hierarchy, you may prefer to have a strategy here that reduces the need to renew the root certificate often.

  The first consideration is choosing the key length of the root's public key and private key pair during setup of the root authority. By using a long key length, which is generally more secure against brute force attack than a shorter key length, you increase the length of time that the CA can use the same private key and have reasonable confidence that it has not been compromised. The second consideration is establishing the validity period of the root certificate itself. In general, you will want to create a root certificate that has a shorter validity period than the estimated lifetime of the key.

  Keeping these considerations in mind, a possible strategy for setting up a root CA that reduces the need for frequent renewal would be to create a 4096-bit RSA key during Certificate Services setup. Given the current state of computer technology, one estimate is that a 4096-bit private key is reasonably secure from a brute force attack for 15-20 years. After creating the 4096-bit key in Certificate Services setup, you could create a root certificate using the 4096-bit key that is valid for five years. Afterwards, the CA's certificate should be renewed every four years (one year before the expiration of the validity period), each time with a certificate validity of five years. Of course, every time the CA certificate is renewed, the administrator will have to assess whether the same key, given current computer technology and other security considerations, can be used with confidence for the next five years.

- **Renewing a subordinate certification authority**

  **Intermediate CAs**

  Within an organization's certification hierarchy, some subordinate CAs may be "intermediate" CAs . In other words, they do not issue certificates to end users or computers; they only issue certificates to other subordinate CAs that are below them in the certification hierarchy. For these intermediate CAs, given their limited role and the relative ease of monitoring their certificate database logs, a possible strategy may be to renew them a year to six months before they expire and reuse the existing key.

  **Issuing CAs**

  For a subordinate CA that issues certificates to end users and computers, referred as an issuing CA, a preferable strategy may be to renew the CA's certificate regularly with a new key 6-12 months before the end of the CA's validity period, for example. This makes an attack on any one key less valuable, since any compromised key would have a relatively limited lifetime.

  The other advantage of renewing an issuing CA using a new key is CRL management. When a CA is renewed with a new key, it begins to publish a separate CRL for the revoked certificates it has issued using the new key. It also will continue to publish a CRL for certificates signed with the old key, for as long as those revoked certificates validity period is still in effect. This can be advantageous especially if a CA is issuing a large number of certificates and, potentially, revoking a large number of those certificates. This reduces the possibility of having one huge CRL that certificate verifiers have to download. Instead, issue a number of certificates with one key, renew the CA early into its validity period with a new key, then issue a number of certificates using the new key. In this way, you will reduce the size of the CRL that the certificate verifier has to download when presented with a certificate from an issuing CA.

For more detailed information about planning for renewal and renewing Microsoft certification authorities, refer to the Microsoft Windows Deployment and Resource Kits.

**Procedures to renew CAs**

For the procedure to renew a root CA, see Renew a root certification authority.

For the procedure to renew a subordinate or intermediate CA, see Renew a subordinate certification authority.

# Managing Certificate Services from a command prompt

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Certificate Services from a command prompt**
There are two utilities you can use to manage certificates at a command prompt:

- Certutil is a useful tool for administrators who are managing CAs and troubleshooting certificate-related issues.
- Certreq requests certificates from a certification authority (CA).

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

# Certificate Services Example Implementations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate Services example implementations**

This section covers:

- Certificate Services example implementation: Establishing autoenrollment for user certificates
- Certificate Services example implementation: Key archival and recovery

# Certificate Services example implementation: Establishing autoenrollment for user certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate Services example implementation: Establishing autoenrollment for user certificates**
These tasks are required for using key recovery on a Microsoft certification authority (CA):

1. Duplicate the User template and use the copy for autoenrollment
2. Configure an enterprise certification authority to issue the Autoenrolled User certificate
3. Establish policy for autoenrollment of Domain Users

## Prerequisites

Before doing these tasks, for the purposes of this walkthrough:

- The Windows Server 2003 domain controller must also be configured as an enterprise subordinate or root CA. The CA is referred to by the name **EntCA** in this walkthrough.

### Duplicate the User template and use the copy for autoenrollment

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Under **Snap-in**, double-click **Certificate Templates**, click **Close**, and then click **OK**.
4. In the console tree, click **Certificate Templates**. All of the certificate templates will be displayed in the details pane.
5. In the details pane, click the **User** template.
6. On the **Action** menu, click **Duplicate Template**.
7. In the **Display Name** field, type **Autoenrolled User**.
8. Make sure that the **Publish Certificate in Active Directory** check box is selected.
9. Click the **Security** tab.
10. In the **Group or user names** field, click **Domain Users**.
11. In the **Permissions for Domain Users** list, select the **Enroll** and **Autoenroll** permission check boxes and then click **OK**.

    Autoenrolling **Domain Users** is an example. These permission settings are variable, depending on who you want to autoenroll for these certificates.

### Configure an enterprise certification authority to issue the Autoenrolled User certificate

1. Open Certification Authority.
2. In the console tree, click **Certificate Templates**.

**Where?**

- Certification Authority/*CA name*/Certificate Templates

3. On the **Action** menu, point to **New**, and then click **Certificate to Issue**.

4. Click **Autoenrolled User** and click **OK**.

**Note**

- To open Certification Authority, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Certification Authority**.

**Establish policy for autoenrollment of Domain Users**

You must be using the Windows Server 2003 Active Directory schema for this procedure to work.

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.

2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.

3. Under **Snap-in**, double-click **Active Directory Users and Computers**, click **Close**, and then click **OK**.

4. In the console tree, click the domain name.

   **Where?**

   - Active Directory Users and Computers/*Domain name*

5. On the **Action** menu, click **Properties**.

6. Click the **Group Policy** tab and click **Edit**.

7. In the console tree, click **Public Key Policies** under **User Configuration**.

   **Where?**

   - Default Domain Policy/User Configuration/Windows Settings/Security Settings/Public Key Policies

8. In the details pane, double-click **Autoenrollment Settings**.

9. Click **Enroll certificates automatically**.

10. Select the **Renew expired certificates, update pending certificates, and remove revoked certificates** check box.

11. Select the **Update certificates that use certificate templates** check box and click **OK**.

12. Click **File**, click **Exit**, and then click **OK**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and

then double-click **Active Directory Users and Computers**.

- Unless users refresh their security policy using the **GPUpdate** command, it will take a number of hours for the autoenrollment policy to take effect.

- When autoenrollment retrieves a certificate that is configured to be stored on a smart card, you will receive a message when the certificate is ready to be stored on the smart card. If the message refers to a type of smart card that you do not have, click **Cancel** until the correct smart card type appears. You then provide the smart card PIN and the certificate will be stored on the smart card.

- When autoenrollment enrolls for a certificate that requires user interaction for the enrollment process, you will receive a **Certificate Enrollment** message and an icon appears on your taskbar. Clicking that icon or message begins the certificate autoenrollment process.

# Certificate Services example implementation: Key archival and recovery

Updated: November 5, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate Services example implementation: Key archival and recovery**

Here are the tasks required for using key recovery with a Microsoft certification authority (CA):

1. Create a key recovery agent account
2. Acquire the key recovery agent certificate
3. Configure the certification authority to allow key recovery
4. Create a new certificate template that allows key archiving
5. Acquire an user certificate that has an archived key
6. Perform key recovery
7. Import the recovered private key

### Note

- Steps 6 and 7 are not required to configure key archival. They are used when a private key is lost and key recovery must take place.

**Prerequisites**

Before doing these tasks, for the purposes of this walkthrough:

- You must have a Windows Server 2003 domain controller. A fictitious domain, adventure-works.com, is used in this walkthrough.
- The Windows Server 2003 domain controller must also be configured as an enterprise root or subordinate CA. The CA is referred to by the name EntCA in this walkthrough.

**Task 1--Creating a Key Recovery Agent account**

Configure and add the Key Recovery Agent certificate template as a template that can be issued by the enterprise CA.

### Important

- Perform the following procedures at the domain controller.

1. Verify who can enroll the Key Recovery Agent template

a. Log on as Administrator@adventure-works.com.

b. Click **Start**, click **Run**, type **certtmpl.msc** and then press Enter.

c. In the console tree, click **Certificate Templates**.

d. In the details pane, right-click **Key Recovery Agent** and click **Properties**.

e. In Key Recovery Agent **Properties**, click the **Security** tab.

By default, the security groups can enroll the Key Recovery Agent certificate template are Domain Administrators and

Enterprise Administrators.

f. If you want another recovery agent, click **Add** to add the user and grant that user **Read** and **Enroll** permissions.

2. Change the default issuance behavior of the Key Recovery Agent template

a. In Key Recovery Agent **Properties**, click the **Issuance Requirements** tab.

b. Clear the **CA certificate manager approval** check box and click **OK**.

c. Close **Certificate Templates**.

3. Configure the EntCA certification authority to issue Key Recovery Agent certificates.

a. On the **Administrative Tools** menu, click **Certification Authority**.

b. In the console tree, double-click **EntCA**, and then click **Certificate Templates**.

c. Right-click **Certificate Templates**, then click **New**, and then click **Certificate to Issue**.

d. In **Select Certificate Template**, click **Key Recovery Agent**, and then click **OK**.

#### Task 2--Acquiring the Key Recovery Agent certificate

In this series of steps, you will acquire a Key Recovery Agent Certificate for the purpose of recovering private keys.

**Important**

- Perform the following procedures at the domain controller.

1. Create an MMC console with the Certificates snap-in loaded.

a. Ensure that you are logged on as administrator@adventure-works.com.

b. On the taskbar, click the **Start** button, and then click **Run**.

c. In **Run**, type **mmc**, and then click **OK**.

d. On the **File** menu, click **Add/Remove Snap-in**.

e. In **Add/Remove Snap-in**, click **Add**.

f. In the **Add Stand-alone Snap-in**, click **Certificates**, and then click **Add**.

g. In **Certificates**, click **User account** and then click **Finish**.

h. Click **Close**, and then click **OK**.

2. Acquire a Key Recovery Agent certificate.

a. In the newly-created MMC console, in the console tree, double-click **Certificates (Current User)**.

b. In the console tree, right-click **Personal**, click **All Tasks**, and then click **Request New Certificate**.

c. In the Certificate Request Wizard, click **Next**.

d. In **Certificate Template**, select **Key Recovery Agent**, and then click **Next**.

e. In the **Certificate Friendly Name and Description**, in **Friendly name**, type **Key Recovery**, and then click **Next**.

f. In **Completing the Certificate Request Wizard**, click **Finish**.

g. In the console tree, double-click **Personal** and then click the Certificates folder.

h. Ensure that a certificate with the friendly name of "Key Recovery" exists.

i. Close the console without saving changes.

#### Task 3--Configuring the CA to allow key recovery

In this series of steps, you configure the EntCA enterprise CA to use the Recovery Agent certificate acquired in Task 2. The CA must load the public key for the Key Recovery Agent to be used for encrypting the recovery data.

**Important**

- Perform the following procedures at the domain controller.

1. Configure the Recovery Agent to be the Administrator's Key Recovery Agent certificate.

a. Ensure that you are logged on as Administrator@adventure-works.com.

b. In **Administrative Tools**, open **Certification Authority**.

c. In the console tree, click **EntCA**.

d. Right-click **EntCA**, and then click **Properties**.

e. On EntCA **Properties**, on the **Recovery Agents** tab, click **Archive the key** and then click **Add**.

f. In **Key Recovery Agent Selection**, click the certificate that is displayed, and then click **OK**.

g. When prompted to restart the CA, click **Yes**.

2. Open the Certificates console, focused on the local computer.

a. On the taskbar, click the **Start** button, and then click **Run**.

b. In **Run**, type **mmc**, and then click **OK**.

c. On the **File** menu, click **Add/Remove Snap-in**.

d. In **Add/Remove Snap-in**, click **Add**.

e. In **Add Standalone Snap-in**, click **Certificates**, and then click **Add**.

f. In **Certificates**, click **Computer account** and then click **Next**.

g. In **Select Computer**, click **Local Computer**, and then click **Finish**.

h. Click **Close**, and then click **OK**.

3. Verify the installation of the KRA certificate.

a. In the console tree, double-click **Certificates (Local Computer)**, double-click **KRA**, and then click **Certificates**.

b. In the details pane, double-click the certificate.

The intended use of the certificate is Key Recovery Agent and the certificate is issued to Administrator. This procedure ensures that the Key Recovery Agent has been successfully configured, and then click **OK**.

d. Close the console without saving changes.

**Task 4--Creating a new certificate template that allows key archiving**

In this series of steps, you will define a new template that allows Key Archival by using the Certificate Templates console. This will allow key recovery in the domain in the event that the private key is lost or corrupted at the client computer.

**Important**

- Perform the following procedures at the domain controller.

1. Open the Certificate Templates console.

a. Log on as Administrator@adventure-works.com.

b. On the taskbar, click the **Start** button, and then click **Run**.

c. In **Run**, type **mmc**, and then click **OK**.

d. On the **File** menu, click **Add/Remove Snap-in**.

e. In **Add/Remove Snap-in**, click **Add**.

f. In **Add Standalone Snap-in**, click **Certificate Templates**, and then click **Add**.

g. Click **Close**, and then click **OK**.

2. Create a duplicate of the Users certificate template with the following properties:

- Template is named "Archive User"
- Key archival is enabled

a. In the console tree, click **Certificate Templates**.

b. In the details pane, right-click the **User** template, and click **Duplicate Template**.

c. In the **Properties** of New Template dialog box, in the **General** tab, in the **Display name** box, type **Archive User**.

d. In the **Request Handling** tab, enable the **Archive subject's private key** option.

The archive key option makes it possible for the Key Recovery Agent to recover the private key from the certificate store.

e. Click the **Security** tab.

Domain Administrators and Domain Users can enroll for this certificate. These permissions were copied from the Users certificate template.

f. Click **OK**.

g. Close the console without saving changes.

**Task 5--Acquiring a User certificate that has an archived key**

In this series of tasks, you configure the EntCA certification authority to issue Archive User certificates. Using a newly created account, you will acquire an Archive User certificate and record the certificate's serial number for later usage.

**Important**

- Perform the following procedures at the domain controller.

1. Configure EntCA to issue the new Archive User certificate template.

a. Ensure that you are logged on as Administrator@adventure-works.com, with a password of **password**.

b. From **Administrative Tools**, open **Certification Authority**.

c. In the console tree, double-click **EntCA**, and then click **Certificates Templates**.

d. Right-click Certificate Templates, click **New**, and then click **Certificate to Issue**.

e. In **Select Certificate Template**, click **Archive User** and then click **OK**.

The Archive User certificate template now appears in the details pane.

f. Close **Certification Authority**.

2. Create a new user account

a. In **Administrative Tools**, open **Active Directory Users and Computers**.

b. Double-click adventure-works.com.

c. In **Users**, create a user account with the following properties:

- First name: Mike
- Last name: Danseglio
- User logon name: Mikedan

- Password: password
- Member of: Server Operators
- E-mail: Mikedan

d. Click **Next**, and then click **Finish**.

e. Double-click the **Mikedan** account, and then select the **Member of** tab.

f. Click **Add**, in **Select Groups**, type **Server Operators**, click **Check Names**, and then click **OK**.

g. Click **OK** to close **Properties**.

h. Close **Active Directory Users and Computers**.

i. Close all open windows and log off of the computer.

3. Log in to the network as Mikedan@adventure-works.com and open the Certificates console.

**Note**

- Mikedan was added to the Server Operators group so he can log on locally to the domain controller.

a. Log on as Mikedan@adventure-works.com.

b. On the taskbar, click the **Start** button, and then click **Run**.

c. In **Run**, type **mmc**, and then click **OK**.

d. From the **Console** menu, click **Add/Remove Snap-in**.

e. In **Add/Remove Snap-in**, click **Add**.

f. In **Add Stand-alone Snap-in**, click **Certificates**, click **Add**, and then click **OK**.

4. Use the Certificates MMC to acquire an Archive User certificate

a. In the newly-created MMC console, in the console tree, double-click **Certificates (Current User)**.

b. In the console tree, right-click **Personal**, click **All Tasks**, and then click **Request New Certificate**.

c. In the Certificate Request Wizard, click **Next**.

d. On **Certificate Template**, select the **Archive User** certificate, and then click **Next**.

e. On **Certificate Friendly Name and Description**, in **Friendly name**, type **Archive User**, and then click **Next**.

f. On **Completing the Certificate Request Wizard**, click **Finish**.

g. On **Certificate Request Wizard**, click **Install Certificate** and then click **OK**.

h. Double-click **Personal**, and then click **Certificates**.

i. In the details pane, double-click the certificate with the friendly name of **Archive User**.

j. In **Certificate**, click the **Details** tab.

Note that the certificate template that used to generate this certificate was ArchiveUser, then click **OK**.

k. Close the new console without saving changes

l. Close all windows and log off of the computer.

**Task 6--Performing a Key Recovery**

In this series of tasks, you will perform a key recovery by using Certutil.exe.

For more information on Certutil, see Certutil.

**Important**

- Perform the following procedures at the domain controller.

1. Log on to the network as Administrator and ensure that the private key is still recoverable by viewing the **Archived Key** column in the Certification Authority console.

a. Log on as Administrator@adventure-works.com.

b. From **Administrative Tools**, open **Certification Authority**.

c. In the console tree, double-click **EntCA**, and then click **Issued Certificates**.

d. From the **View** menu, click **Add/Remove Columns**.

e. In **Add/Remove Columns**, in **Available Column**, select **Archived Key**, and then click **Add**. Archived Key should now appear in **Displayed Columns**.

f. Click **OK** and then, in the details pane, scroll to the right and confirm that the last issued certificate to Mikedan has a **Yes** value in the Archived Key column.

**Note**

- A certificate template must have been modified so that the Archive bit and Mark Private Key as Exportable attributes were enabled. The private key is only recoverable if there is data in the Archived Key column.

g. Double-click the **Archive User** certificate.

h. Click the **Details** tab

Write down the serial number of the certificate. (Do not include spacing between digit pairs.) This is required for recovery.

The serial number will be a hexadecimal string that is 20 characters long. The serial number of the private key is the same as the serial number of the certificate. For the purposes of this walkthrough, the serial number will be referred to as *serialnumber*.

i. Click **OK**.

j. Close **Certification Authority**.

2. Recover the private key into a BLOB output file by using certutil.exe.

a. On the taskbar, click the **Start** button, click **Run**, type **cmd**, then click **OK**.

This opens a command prompt window.

b. Type **cd ** and then press ENTER.

c. Ensure that you are in the **c:** directory.

d. At the command prompt, type:

**Certutil -getkey** *serialnumber* **outputblob**

e. At the command prompt, type **dir outputblob**

**Note**

- If the file *outputblob* does not exist, you probably typed the serial number incorrectly for the certificate.

The *outputblob* file is a PKCS#7 file containing the KRA certificates and the user certificate and chain. The inner content is an encrypted PKCS#7 containing the private key (encrypted to the KRA certificates).

3. Recover the original private/public key pair using Certutil.exe

a. On the taskbar, click the **Start** button, click **Run**, type **cmd**, then click **OK**.

This opens a command prompt window.

b. At a command prompt, type:

**Certutil -recoverkey outputblob Mikedan.pfx**

c. When prompted, enter the following information:

Enter new password: **password**

Confirm new password: **password**

d. Type **exit**, and then press ENTER.

e. Close all windows and log off as the current user.

**Task 7--Importing the recovered private key**

In this series of tasks, you will restore the recovered private key in Mike's certificate store by importing the Mikedan.pfx file.

**Important**

- Perform the following procedures at the domain controller.

1. Log on as Mikedan@adventure-works.com and start the Certificates mmc.

a. Log on as Mikedan@adventure-works.com, with a password of *password*.

b. On the taskbar, click the **Start** button, and then click **Run**.

c. In **Run**, type **mmc**, and then click **OK**.

d. On the **File** menu, click **Add/Remove Snap-in**.

e. In **Add/Remove Snap-in**, click **Add**.

f. In **Add Standalone Snap-in**, click **Certificates**, click **Add**,and then click **OK**.

2. Delete all certificates issued by EntCA to simulate a re-installed computer.

a. Right-click **Certificates (Current User)**, and then click **Find Certificates**.

b. In **Find Certificates**, in **Contains**, type **EntCA** and then click **Find Now**.

c. In **Find Certificates**, on the **Edit** menu, click **Select All**.

d. In **Find Certificates**, on the **File** menu, click **Delete**.

e. In Certificates, click **Yes**.

f. In Root Certificate Store, click **Yes**.

g. Close **Find Certificates**.

3. Import the certificate at c:\Mikedan.pfx and let the certificates be placed automatically.

a. In the console tree, right-click **Personal** and then click **All Tasks** and then click **Import**.

b. In the Certificate Import Wizard, click **Next**.

c. On **Files to Import**, in the **File name** box, type **c:\Mikedan.pfx**, and then click **Next**.

d. In **Password**, type **password** and then click **Next**.

e. On **Certificate Store**, click **Automatically select the certificate store based on the type of certificate** and then click **Next**.

f. On **Completing the Certificate Import Wizard**, click **Finish**

g. If the **Root Certificate Store** dialog box appears, click **Yes**.

h. In **Certificate Wizard Import**, click **OK**.

Two certificates were imported. The Archive User certificate for Mike is located in the Personal certificates store and the EntCA certificate is located in the Trusted Root Certification Authorities store.

4. Verify the serial number of the imported certificate.

a. In the console tree, double-click **Personal** and then click **Certificates**.

b. Double-click the certificate

c. In **Certificate**, click the **Details** tab. Verify that the serial number matches the original.

d. Close all open windows and log off of the network.

**See Also**
**Other Resources**
Active Directory Certificate Services PKI - Key Archival and Management

# Using Qualified Subordination

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using qualified subordination**

This section covers:

- Qualified subordination overview
- Name constraints
- Mapping policies between trust hierarchies
- Policy constraints

# Qualified subordination overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Qualified subordination overview**
Qualified subordination allows you to place certificate issuance constraints on subordinate CAs and to place usage constraints on the certificates they issue. With qualified subordination, you can focus subordinate CAs according to specific certification needs and administer your public key infrastructure (PKI) more efficiently.

**Certificate-based trust between trust hierarchies**
Qualified subordination also allows you to establish trust between CAs in separate trust hierarchies. This type of trust relationship is also called *cross-certification*. With this trust relationship, qualified subordination is not limited to subordinate CAs. Trust between hierarchies may be established using a subordinate CA in one hierarchy and the root CA in another hierarchy.

**Additional levels of security controls**
Qualified subordination extends your trust hierarchy by allowing you to place additional trust conditions within and between the domains in your PKI. With qualified subordination, the qualified subordinate CAs in your trust hierarchy can each have different rules governing how they will issue certificates and how their certificates may be used. For more information on qualified subordination, see Qualified subordination.

**Qualified subordination extensions**
Each qualified subordination constraint adds an additional trust condition and more narrowly defines the scope of a CA. These conditions are defined in the certificate extensions of a qualified subordinate CA certificate. When you install a qualified subordinate CA, you define the following extensions in its CA certificate:

| Extension | Description |
|---|---|
| Name Constraints | Restricts the namespaces that are permitted or excluded by the qualified subordinate CA and its subordinates when issuing certificates. |
| Policies | Defines the list of acceptable issuance and application policies for certificate usage. These policies are identified in the certificate by object identifiers (also known as object identifiers). |
| Policy Mapping | Allows a policy from one domain to be mapped onto a policy of another domain. |
| Policy Constraints | Restricts the subordination levels in a certificate hierarchy to which a policy is applied. These extensions are used in conjunction with issuance and application policies only. |

All constraints that are placed on a qualified subordinate CA are defined in the complete request for the CA certificate that is used to install the qualified subordinate CA. The complete request is a CMC certificate request that has been generated using an information file (.inf file). All qualified subordinate CA policies, constraints, and mappings can be defined using an information file (.inf). In order to modify the policies and constraints placed on a qualified subordinate CA after installation, the administrator of that qualified subordinate CA's parent CA must reissue the CA certificate for the qualified subordinate CA, thereby reinstalling it. For more information, see Perform qualified subordination.

The constraints and policy types listed above are defined in RFC 2459. For more information on RFCs, see the RFC Editor Web site.

For more information, see Name constraints, Policy constraints, Mapping policies between trust hierarchies, and Policy qualifiers.

**Qualified subordinate CA scenarios**
The following sections show how qualified subordination can be used.

**Name constraints scenarios**
The Humongous Insurance Company has two Active Directory directory service domains hosted in different geographic locations, one domain in North America and one domain in South America. The company's PKI administrator wants the qualified subordinate CAs that are deployed in each domain to issue certificates for users in their local domain only. The two Active Directory domains are *northamerica.humongousinsurance.com* and *southamerica.humongousinsurance.com*.

The administrator performs the following steps:

1. Installs a qualified subordinate CA for the domain *northamerica.humongousinsurance.com* with the following permitted name constraints:

| Name constraint format | Permitted name constraint | Example |
|---|---|---|
| Distinguished name | DC=northamerica,DC=humongousinsurance,DC=com | CN=Mike Danseglio,CN=Users,DC=northamerica,DC=humongousinsurance,DC=com |

| | | |
|---|---|---|
| DNS domain name | .northamerica.humongousinsurance.com | hostname.northamerica.humongousinsurance.com |

1. Installs a qualified subordinate CA for the domain *southamerica.humongousinsurance.com* with the following permitted name constraints:

| **Name constraint format** | **Permitted name constraint** | **Example** |
|---|---|---|
| Distinguished name | DC=southamerica,DC=humongousinsurance,DC=com | CN=Don Funk,CN=Users,DC=southamerica,DC=humongousinsurance,DC=com |
| DNS domain name | .southamerica.humongousinsurance.com | hostname.southamerica.humongousinsurance.com |

As a result, the certificate issuance behavior of each qualified subordinate CA is restricted to users in their respective domains. If a user in *northamerica.humongousinsurance.com* requests a certificate from the qualified subordinate CA for *southamerica.humongousinsurance.com*, the request is denied because the namespace specified in a permitted name constraint automatically excludes any other namespaces.

For more information, see Name constraints.

**Policy scenarios**

Humongous Insurance wants to use application and issuance policies to configure the following PKI scenarios:

- Only allow remote network logon using smart card authentication.
- Create a certificate that allows selected users to make purchases of up to one million dollars.

**Only allow remote network logon using smart card authentication**

To restrict remote network logon to only allow smart card authentication, the company performs the following steps:

1. Configures their remote access server to require smart card authentication.
2. Installs and configures a qualified subordinate CA to issue smart card logon certificates containing a specific application policy that is required by the remote access server.
3. Distributes smart cards containing the specific application policy that is required by the remote access server.

When a user attempts to log in, the following steps occur:

1. Users attempt to log on remotely over a dial-up connection using their smart cards.
2. The remote access server software checks to see if the smart card contains the required application policy.
3. If the smart card contains the required application policy, then the remote access server software forwards the logon request to the nearest domain controller.
4. The domain controller grants the user access to the network.

**Create a certificate that allows selected users to make purchases of up to one million dollars**

To create a certificate that allows selected users to make purchases of up to one million dollars, the company performs the following steps:

1. Creates a signing certificate for the company's accounting department that contains a specific issuance policy object identifier (also known as an OID).
2. When the company creates a Million Dollar Purchase certificate template, it adds an *issuance requirement* for the Accounting department's specific issuance policy object identifier.
3. The user creates a million dollar purchase request using the template.
4. The accounting department manager is notified that there is a pending certificate request which she must accept or reject.
5. The accounting department manager signs the Million Dollar Purchase certificate request with the signing certificate that contains the specific issuance policy object identifier.

6. The signed request is sent to the company's CA.

7. The company's CA receives the signed request, determines that the issuance requirement for the certificate template has been met by identifying the accounting department's specific issuance policy object identifier, and issues the Million Dollar Purchase certificate to the user.

**Policy mapping scenario**
Two companies, Humongous Insurance and Northwind Traders, enter into a partnership and want to establish a trusted communication path where the $1 million issuance policy of Humongous Insurance is considered equal with the two million dollar issuance policy of Northwind Traders.

The CA administrator for the Humongous Insurance domain *humongousinsurance.com* performs the following steps:

1. Obtains the two million dollar issuance policy object identifier for Northwind Traders under a very secure channel.

2. Defines the policy mapping between Humongous Insurance's one million dollar issuance policy and Northwind Traders' two million dollar issuance policy.

3. Installs a qualified subordinate CA in the Humongous Insurance trust hierarchy containing the policy mapping.

The policy mapping results in allowing purchase orders signed in the Northwind Traders trust hierarchy to be validated by the trust hierarchy of Humongous Insurance.

For more information, see Mapping policies between trust hierarchies; Policy constraints

**Enterprise and stand-alone qualified subordinate CAs**
Windows Server 2003 qualified subordinate CAs can be installed as enterprise qualified subordinate CAs or stand-alone qualified subordinate CAs. The difference between the two types of qualified subordinate CAs is how they identify users requesting certificates. Enterprise qualified subordinate CAs require the use of Active Directory and stand-alone qualified subordinate CAs do not.

Enterprise qualified subordinate CAs identify users by some or all of their Active Directory user object attributes. When an enterprise qualified subordinate CA receives a certificate request from a user, it identifies the user according to the authentication credentials that are used by that user during network logon and uses that information to query the Active Directory and obtain the user's user object attributes.

When users request certificates from a stand-alone qualified subordinate CA, the user fills out a Web page form where the user specifies information such as a name, e-mail address, company name, department, and country/region. This information is included in the fields of the certificate request that is sent to the stand-alone qualified subordinate CA, which uses this information to determine certificate issuance. A stand-alone CA cannot publish any issued certificates to Active Directory other than its own CA certificate. However, the more common scenario is to have a stand-alone root CA with an enterprise qualified subordinate CA. For more information, see Stand-alone certification authorities.

**Verification paths for qualified subordinate CAs**
Each qualified subordinate CA's policies and constraints are defined according to a subset of its certification path's policies and constraints. Using qualified subordination, boundaries are established within the certification hierarchy to ensure that all CAs in a hierarchy's paths issue certificates within certain constraints. These constraints may be maintained or more narrowly defined the further the certification hierarchy branches from its trust anchor (root CA).

The following graphic shows how qualified subordinate CAs can only issue CA certificates according to a subset of the rules permitted by their certification path. In the graphic, policies are used as an example and are represented as numbers to demonstrate a subset.



**Qualified subordination certification hierarchy restrictions**

Qualified subordinate CAs CA 1, CA 2, and CA 3 have defined policies from which the qualified subordinate CAs below them may select policies. With each level of the certificate hierarchy, fewer of the acceptable policies are used and the qualified subordinate CAs are more limited in scope.

No trust model will permit the qualified subordinate CA to allow certificate issuance behavior that its certificate path does not allow. To enforce this rule, the following process is followed when a parent CA receives a CA certificate request:

1. The CA verifies the certificate request against all the constraints placed on all the CAs from its root CA down its certification hierarchy to itself.

2. At each intersection in the certification hierarchy, the issuing qualified subordinate CA compares the information contained in the certificate request with those set at the existing CA.

3. If there are no conflicts between the information in the certificate request and those placed on the CAs in the certification hierarchy, then the certificate is issued. If there is a conflict, the certificate request is denied.

**Notes**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

- In addition to the new constraints discussed in this topic, *basic constraints* can be used to control how a CA issues certificates. Basic constraints are available in both Windows 2000 and Windows Server 2003 CAs. For more information, see Basic constraints and Change basic constraints.

# Name constraints

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Name constraints**
You can use qualified subordinate certification authority (CA) name constraints to restrict a qualified subordinate CA's certificate issuance authority to single or multiple namespaces used within your network or public key infrastructure (PKI), such as those used as Active Directory domains. By constraining the qualified subordinate CA's certificate issuance according to one or more namespaces, you will have greater control over which users and computers receive certificates from your qualified subordinate CAs. For more information, see Qualified subordination.

You can use the following naming and addressing formats to constrain the certificate issuance activities of qualified subordinate CAs:

- Directory name (for example, an Active Directory distinguished name)
- DNS domain name
- E-mail name
- User principal name (UPN)
- Universal Resource Identifier (URI)
- Internet Protocol address

The qualified subordinate CA's name constraints define namespaces within which all user names identified in the certificate requests it receives must be located. Name constraints are organized as *permitted* or *excluded* name constraints. If the same name formats used as name constraints are present in the certificate requests received by the qualified subordinate CA, it evaluates them according to its permitted and excluded name constraints.

When using qualified subordination in your PKI deployment, you can leverage the well-planned naming and addressing methods used in your network to support the goals of your PKI. In a sense, you can imprint your trust model onto these existing methods, assigning CAs to different domains and configuring them to issue certificates only to users and computers in those domains. When you deploy a qualified subordinate CA to a network domain, you will need to consider the namespaces and addressing methods used by the resources of that domain, and possibly subdomains and other domains, in order to configure the qualified subordinate CA to issue certificates according to your administrative and security goals.

**Choosing permitted and excluded name constraints**
When choosing a qualified subordinate CA's name constraints, you need to decide on the users that the qualified subordinate CA will receive certificate requests from. Typically, the qualified subordinate CA will be deployed to an Active Directory domain in your network and be responsible for issuing certificates to the users and computers within that domain and possibly its subdomains. The namespaces used by these domains, including Active Directory names, DNS namespaces, and other namespaces, can be used to define the permitted and excluded name constraints of the qualified subordinate CA.

Excluded name constraints are used to exclude a child namespace of a permitted namespace. For example, your qualified subordinate CA may have a permitted name constraint for the DNS namespace *.microsoft.com*, but you do not want it to issue certificates for the subdomain *.betatest.microsoft.com* because that subdomain may be experimenting with software that should not be certified within your PKI. An excluded name constraint of *.betatest.microsoft.com* removes the qualified subordinate CA's ability to issue certificates to users and computers within the namespace *.betatest.microsoft.com*.

If an excluded name constraint is used to exclude a parent namespace of a permitted name constraint, it will exclude the entire permitted namespace. For example, if a permitted name constraint is *.example.microsoft.com* and the excluded name constraint is *.microsoft.com*, then the entire *.microsoft.com* namespace is excluded.

It is also important to note that excluded namespaces take precedence over permitted namespaces, guaranteeing that the qualified subordinate CA will not issue a certificate to a user within an excluded namespace even if the user is also within a permitted namespace. For example, a user could be within the permitted Active Directory namespace *.microsoft.com* but also within the excluded DNS namespace *.widgets.microsoft.com*. In this example, the excluded DNS namespace overrides the

permitted Active Directory namespace and the qualified subordinate CA does not fulfill the user's certificate request.

When choosing permitted and excluded name constraints, you should be aware of how the qualified subordinate CA processes certificate requests. When the qualified subordinate CA is comparing the names identified in a certificate request to its name constraints, all names in a certificate request must be within permitted namespaces and none may be within an excluded namespace. For more information on how a qualified subordinate CA using name constraints when processing certificate requests.

**Directory name constraints**

Active Directory allows you to use distinguished names as name constraints for your qualified subordinate CAs. Distinguished names identify users and resources on the network in Active Directory. This allows you to constrain a qualified subordinate CA to permit or exclude users in the Active Directory using the users' distinguished names. Active Directory also uses distinguished names to create and reference groups of objects in the directory, such as users and computers. The distinguished names of these object groups can also be used as name constraints, allowing you to constrain a qualified subordinate CA to permit and exclude certificate issuance for entire groups in the directory.

The following table describes types of distinguished names that are used as permitted and excluded names.

| Directory object | Distinguished name constraint examples |
|---|---|
| User | `CN=Users,DC=example,DC=microsoft,DC=com` |
| Computer | `CN=Computers,DC=microsoft,DC=com` |
| Domain controller | `CN=Servers,CN=Sites,CN=Configuration,DC=microsoft,DC=com` |
| Contact | `OU=Contacts,DC=microsoft,DC=com` |
| Group | `OU=My Group,DC=microsoft,DC=com` |
| Organizational Unit | `OU=ITServices,DC=microsoft,DC=com` |

X.500 directory names can also be used as name constraints. For example, **o=cities,st=Redmond,c=US**.

**DNS name constraints**

You can apply the DNS namespaces used by your network for name resolution as name constraints for a qualified subordinate CA. When the qualified subordinate CA receives a certificate request, it will compare the DNS name associated with the computer requesting the certificate to its DNS name constraints and decide whether or not to issue a certificate. The subordinate CA will apply name constraints to all entities within the DNS namespaces specified as DNS name constraints. The DNS domain names accepted by the subordinate CA must follow the standard DNS naming conventions as specified in RFCs 1034 and 1035. The permitted and excluded DNS name constraints are specified using the information file (.inf) that is used to install the qualified subordinate CA.

There are two different ways to specify a DNS name constraint:

- **DNS host name**. Specifies a specific DNS host name as a permitted or excluded name constraint, such as `host1.example.microsoft.com`.

- **DNS namespace**. Specifies a DNS name space wherein all DNS host names are permitted or excluded, such as `.example.microsoft.com`.

To specify that a DNS name constraint should apply to all entities within a DNS namespace and not host names containing part of the name, you must use a period (.) after the leftmost label in the domain name, such as `.example.microsoft.com`. This period indicates that all subjects with a DNS name ending in the domain `.example.microsoft.com` are subject to the DNS name constraint, but the host name `host1example.microsoft.com` is not a match because `host1example` contains no period separator between *host1* and *example*.

When the subordinate CA compares the DNS name in a certificate request with its name constraints, it performs a string comparison. The string comparison attempts to locate a match within any part of the names. If the qualified subordinate CA's DNS name constraint is `example.microsoft.com` and the certificate request contains a DNS host name `host1example.microsoft.com`, the qualified subordinate CA will consider those names as a match. Using the same DNS name constraint (`example.microsoft.com`), the domain name `www.example.microsoft.com` is also considered a match.

With DNS name constraints involving domain names, any subdomain of a name constraint satisfies the constraint. If you apply

an excluded DNS name constraint for `.example.microsoft.com` (notice the leftmost period), then names such as `www.example.microsoft.com` and `us.sales.example.microsoft.com` are excluded also. Using the same excluded DNS name constraint, names such as `it.microsoft.com` and `accounts.microsoft.com` would not be excluded.

The following table demonstrates the two different DNS name constraint evaluation methods used by a qualified subordinate CA.

| Certificate request's DNS information | DNS constraint | Qualified subordinate CA evaluation |
|---|---|---|
| host1example.microsoft.com | .example.microsoft.com | No match |
| host1.example.microsoft.com | .example.microsoft.com | Match |
| www.example.microsoft.com | .example.microsoft.com | Match |
| us.sales.example.microsoft.com | .example.microsoft.com | Match |

When creating DNS constraints, you should consider the hierarchies of the DNS namespaces to avoid conflicts between permitted and excluded DNS constraints. If you are permitting and excluding DNS domains for a qualified subordinate CA, the excluded domain names override the permitted names. If you specify the permitted DNS name constraint `www.example.microsoft.com` and also specify the excluded DNS name constraint `.microsoft.com`, the permitted name will be cancelled out by the excluded name because the permitted domain name is a subdomain of the excluded domain name.

A qualified subordinate CA's DNS constraints cannot permit DNS names excluded by its parent CA. For example, if the parent CA excludes the DNS name `.microsoft.com`, then its qualified subordinate CA cannot permit the subdomain `www.example.microsoft.com`.

**URI constraints**

You can specify Uniform Resource Identifiers (URIs) as name constraints for a qualified subordinate CA. URIs are used to identify resources on the Internet using identifiers such as URL, FTP, HTTP, telnet, mailto, news, and gopher. The URI naming conventions supported by the qualified subordinate CA must follow the syntax specified in RFC 2396.

When validating the URI names in a certificate request, the qualified subordinate CA ignores the protocol element in the URI, such as http:// or ftp://, and uses the domain or host names only. For example, if the URI constraint `URL=http://.example.microsoft.com` is compared to the name `ftp://host1.example.microsoft.com`, then the comparison is actually between the URI constraint `.example.microsoft.com` and the name `host1.example.microsoft.com`.

Similar to DNS constraints, URI constraints may specify a host or domain. When specifying a domain, URI name constraints must begin with a period (.), such as `.example.microsoft.com`. This URI constraint also applies to all subdomains of the domain constraint, such as `abc.example.microsoft.com` or `xyz.example.microsoft.com`. Using the same example, if the URI name in a certificate request does not contain a period to the left of example, then the constraint does not apply. For example, the certificate request URI name `example.microsoft.com` (notice that there is no beginning period) does not satisfy the name constraint `.example.microsoft.com`.

When specifying a host, the beginning period is not used, such as `host1example.microsoft.com`. Without the beginning period, the name constraint is satisfied only when the subject in the certificate request has the leftmost label (`host1example`). For example, the name constraint `example.microsoft.com` would be satisfied by the certificate request name `host1example.microsoft.com` but not by `host1.example.microsoft.com`.

The most frequent application of URI name constraints is the validation of certificate requests for subjects wanting to access to a secure Web server. For example, if the Web server certificate request has a name of http://www.microsoft.com and the qualified subordinate CA's permitted URI name constraint is `URL=http://www.microsoft.com`, then this URI constraint would permit the certificate request for the Web server certificate.

**E-mail and UPN constraints**

Both e-mail and UPN naming conventions follow a similar format that must be followed when you format name constraints. E-mail extensions and UPNs used as name constraints follow the naming conventions specified in RFC 822.

You can specify e-mail and UPN name constraints for an individual subject, such as mikedan@example.microsoft.com, or you can specify constraints for all subjects whose e-mail names or UPNs end in a specific name, such as `@example.microsoft.com`. Typically, you will specify e-mail or UPN name constraints for all subjects whose e-mail addresses and UPNs end in a specific name.

The difference in specifying these two name constraints is the at sign (@). If the name constraint contains any characters to the left of the @, the qualified subordinate CA regards the name constraint as applying to a single subject. Conversely, if there are

no characters to the left of the @, then the qualified subordinate CA regards the name constraint as applying to all subjects whose e-mail addresses or UPNs end in the name constraint.

| E-mail and UPN constraints | Certificate name value | Qualified subordinate CA's evaluation |
|---|---|---|
| host1@example.microsoft.com | host1@example.microsoft.com | Match |
| @example.microsoft.com | host1@example.microsoft.com | Match |
| host1@example.microsoft.com | mikedan@example.microsoft.com | No match |

**Note**

- If UPN and e-mail permitted name constraints are not defined when defining the permitted name constraints for a qualified subordinate CA certificate request, then the UPN and e-mail permitted name constraints are automatically entered by the issuing CA's policy module when it creates the qualified subordinate CA certificate. The UPN and e-mail permitted name constraints are left unspecified and therefore permit all names. If you want to apply UPN and e-mail permitted name constraints, you must specify the UPN and e-mail permitted constraints when defining the permitted name constraints for a qualified subordinate CA certificate request.

**IP address constraints**

You can specify IP addresses and IP address ranges as name constraints on a qualified subordinate CA. IP address name constraints follow the formatting conventions specified in RFCs 791 (IPv4) and 1883 (IPv6). The IP addresses that are contained in the certificate requests that are made to a qualified subordinate CA will be compared to the IP addresses in the qualified subordinate CA's name constraints.

**Name constraint processing rules**

For a certificate request to be successful, all names that are identified in a certificate request must be within permitted name spaces and none may be within an excluded namespace. Each permitted and excluded name constraint is processed separately.

Name constraints are processed by the qualified subordinate CA according to the following rules:

- Constraints apply only when the namespace types specified as name constraints are present in the certificate request.
- A certificate request is successful if all names in the certificate request succeed in matching the permitted name constraint.
- A certificate request is unsuccessful if all names in the certificate request succeed in matching the excluded name constraint.
- A certificate request containing names that can not be matched with either the permitted or excluded names is considered excluded and is failed.
- Comparison of the name identifying the subject entity with the qualified subordinate CA's name constraints is performed using Unicode. The comparison is not case sensitive.

To find a match, the qualified subordinate CA starts at the rightmost part of the name in the certificate request and compares that name to its list of permitted and excluded name constraints. For example, when the CA receives a certificate request with the DNS name example.microsoft.com it will attempt to find a name constraint match starting with .com, then microsoft.com, and then example.microsoft.com. The exceptions to this are distinguished names (DNs), where the qualified subordinate CA starts from the leftmost name, and IP addresses, where it attempts to find a match starting with the network address and then the host address.

For a certificate request to be accepted by the qualified subordinate CA, each name in the certificate request should match at least one of the qualified subordinate CA's name constraints of the same name type. If there are two permitted distinguished names on the qualified subordinate CA, then the certificate request subject's name should match at least one of these distinguished names. No comparisons are performed between different types of names.

Excluded names take precedence over permitted names. If a permitted DNS name constraint allows users from the DNS domain `.example.microsoft.com`, but the same name is listed as an excluded name constraint, then users from the DNS domain `.example.microsoft.com` are excluded and the qualified subordinate CA will fail their certificate requests. This

processing rule requires that you pay close attention to the hierarchy of the different name conventions. For example, if a permitted DNS name constraint permits users from the DNS domain `.example.microsoft.com`, but the parent domain `.microsoft.com` is listed as an excluded name constraint, then users from the DNS domain `.example.microsoft.com` are excluded because `.example.microsoft.com` is a child domain of `.microsoft.com`.

# Mapping policies between trust hierarchies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Mapping policies between trust hierarchies**

To establish trust between users and computers in different trust hierarchies, you can map the issuance and application policies of each trust hierarchy. The qualified subordinate CA that contains this mapping is called the *issuer CA* and the subordinate CA whose policies have been mapped is called the *subject CA*. In mapping some or all of the subject CA's policies to the issuer's policies, the issuer CA effectually subordinates the subject CA. The result of this mapping is that users and computers in the issuer CA trust hierarchy who attempt to validate certificates from users and computers in the subject CA trust hierarchy can use their own certification path to perform validation. The separate trust hierarchy can be within the same intranet or in separate PKI environments over an extranet.

Mapping policies across trust hierarchies only requires that you install the issuer CA in your trust hierarchy. All the policies that are used in the subject CA trust hierarchy are mapped to the policies in your trust hierarchy, where you install the qualified subordinate CA. For this reason, you can establish a trust relationship, in fact subordinate, with a CA in another company by obtaining its policy object identifiers (also known as OIDs) and mapping them to the policy object identifiers that are used in your trust hierarchy.

The following steps are performed when using policy mapping for both issuance and application policies:

1. Determine the trust hierarchy with which you want to establish a trust relationship.

2. Establish the equivalence between the assurance levels that are used between the two trust hierarchies involved in the trust relationship.

3. Obtain the issuance and application policy object identifiers that are used in both trust hierarchies involved in the trust relationship.

4. Map the issuance and application policy object identifiers in the separate trust hierarchies and define their policy constraints in the CA certificate request for the qualified subordinate CA you are installing in your trust hierarchy.

5. Install the qualified subordinate CA with the policies, policy mappings, and policy constraints in your trust hierarchy.

When a user in the issuer CA trust hierarchy receives a file signed by a user in the subject CA trust hierarchy, the application of the user in your domain will attempt to validate the certificate using the validation path in your trust hierarchy. Because the object identifier in the certificate maps to an object identifier in your trust hierarchy, the certificate is validated. When a user in the subject CA trust hierarchy receives a file signed by a user in the issuer CA trust hierarchy, the subject CA user's application will attempt to validate the certificate using the validation path in its own certification hierarchy. Because the certificate contains the object identifiers used in the subject CA trust hierarchy, the certificate is validated.

Before defining the policy mapping, you need to obtain all the object identifiers that will correspond between the domains in the trust relationship. If you are creating a qualified subordinate CA for the purpose of establishing a trust relationship between trust hierarchies in different companies or PKIs, you will need to obtain the object identifiers used in the subject CA trust hierarchy by obtaining a certificate used in that hierarchy or from the CA administrator of that hierarchy. This information should only be passed under the most secure channels.

Issuance policy object identifiers in your trust hierarchy can only be mapped to issuance policy object identifiers in the foreign trust hierarchy. Similarly, application policy object identifiers in your trust hierarchy can only be mapped to application policy object identifiers in the foreign trust hierarchy.

For more information, see Qualified subordination; Qualified subordination overview; Name constraints; Issuance policies; Application policies; Policy constraints; Policy qualifiers; and Certificate Templates.

**Constraining policy mapping**

When a qualified subordinate CA uses policy mapping, you can use policy constraints to set the number of subordinate CAs below it in the certificate hierarchy that may map to its policies. An *inhibit policy mapping* policy constraint of 3 will inhibit policy mapping to subordinate CAs four levels below the qualified subordinate CA. A subordinate CA 5 levels below the qualified subordinate CA will not be allowed to define policy mappings.

**Unexpected trust**

When mapping policies between trust hierarchies in two different companies, the problem of unexpected trust could arise should a third company map its policy object identifier to the mapped object identifier of one of the two companies. For example:

- The secure e-mail policy used by *humongousinsurance.com* is mapped to the secure e-mail policy in *wideworldimporters.com*.

- A third company, *northwindtraders.com*, maps its secure e-mail policy to the policy of wideworldimporters.com.

- This allows a user in northwindtraders.com to send signed e-mail to humongousinsurance.com that is considered valid because the e-mail application in humongousinsurance.com sees its own valid policy object identifier in the certificate issued by the CA in northwindtraders.com.

If humongousinsurance.com wants to prevent any additional policy mapping by northwindtraders.com from affecting its domain, it could place an *inhibit policy mapping* of 0 on its issuer CA. The CA for northwindtraders.com is considered a subordinate of the CA for wideworldimporters.com, and is thereby two levels below the issuer CA for humongousinsurance.com. Consequently, an *inhibit policy mapping* of 0 would cause the qualified subordinate CA for humongousinsurance.com to only process the policy mapping in certificates from subjects directly below it in its trust model (wideworldimporters.com) and refuse any certificates issued below that level (northwindtraders.com).

# Policy constraints

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policy constraints**
Qualified subordinate certification authority (CA) policy constraints allow you to constrain the path validation of your certification hierarchy by:

- Defining how far below a qualified subordinate CA in your certification hierarchy that a policy is required.
- Defining how far below a subordinate CA in your certification hierarchy that policy mapping is permitted.

There are two policy constraint settings:

- **Require explicit policy**. Determines the maximum number of additional certificates that may appear below the qualified subordinate CA in the certification hierarchy before a policy is required.
- **Inhibit policy mapping**. Determines the maximum number of additional certificates that may appear below the qualified subordinate CA in the certification hierarchy before policy mapping is no longer permitted.

Policy constraints are set separately for issuance and application policies. The policy constraints for a qualified subordinate CA's issuance policies do not have to be the same as the *Require explicit policy* and *Inhibit policy mapping* settings for its application policies.

When defining policy constraints, the numbering system used begins with 0. For example, if a qualified subordinate CA policy has an inhibit policy mapping of 0, then that policy can only be mapped one level below the qualified subordinate CA in the trust hierarchy.

For more information, see Qualified subordination; Qualified subordination overview; Name constraints; Mapping policies between trust hierarchies; Policy qualifiers; and Certificate Templates.

**Note**

- Policy constraints cannot be specified in certificate templates, but must be specified in the information file (.inf) used to create the CA certificate that, in turn, is used to install the qualified subordinate CA.

# Certificates Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

The following books provide detailed information about certificates, certification authorities, public key encryption, and public key infrastructure (PKI):


- Austin, Thomas, *PKI:A Wiley Tech Brief*. John Wiley & Sons, December 2000.
- Feghhi, Jalal and Jalil Feghhi, *Secure Networking with Windows 2000 and Trust Services*. Addison-Wesley Pub Co., February 2001.
- Housley, R.P., and T. Polk. 2001. *Planning for PKI*. New York, New York: John Wiley & Sons, Inc.
- Howard, Michael, *Designing Secure Web-Based Applications for Microsoft Windows 2000*. Microsoft Press, August 2000.
- Kaufman, C., Perlman, R., and M. Speciner. 1995. *Network Security: Private Communication in a Public World*. Upper Saddle River, New Jersey: Prentice-Hall Inc.
- Schneier, Bruce, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. John Wiley & Sons, December 1995.


The following resources have in-depth information about planning, deploying, and using a public key infrastructure with Windows Certificate Services:

- Microsoft Security on the Microsoft Web site.

For updated technical information, see:

- For more information about certificate services architecture, see "Designing a Public Key Infrastructure" at the Microsoft Windows Resource Kits Web site.
- Microsoft TechNet on the Microsoft Web site.
- Updated technical information on the Microsoft Web site.

# Certificate Services Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Clients do not automatically enroll for certificates after autoenrollment is configured.
- General troubleshooting for a certification authority (CA).
- Error when accessing the certification authority Web pages.
- Web pages on an enterprise certification authority (CA) either don't generate certificates or generate certificates that are not valid.
- Can't log into certification authority (CA) Web pages.
- Certification authority (CA) Web enrollment pages that are installed on a remote server other than the CA do not work.
- Internet Information Services (IIS) 4.0 and Certificate Server 1.0 cannot enroll for certificates using Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition Certificate Services Web pages.
- A user tries to log on with the smart card and receives this message: "The system cannot log you on to this domain because the system's computer account in its primary domain is missing or the password on the account is incorrect."
- After renewing a CA, computers are no longer automatically enrolling for certificates from that CA.
- When trying to enroll for a certificate from a computer or account belonging to a child domain of the domain where the CA is located, the following error appears: "No template could be found. There are no CAs from which you have permission to request a certificate, or an error occurred while accessing the Active Directory."

**Clients do not automatically enroll for certificates after autoenrollment is configured.**
**Cause:** The Group Policy information used for autoenrollment has not yet replicated to the client computers. By default, this information can take up to 120 minutes to replicate to all computers.

**Solution:** Wait for Group Policy to complete replication or use the Gpupdate tool to force replication to occur immediately.

**General troubleshooting for a certification authority (CA).**
Check the event log of the server. It often logs more detailed errors than you will notice from doing the procedure you're attempting.

For more information about Certificate Services logging, see article Q305018, "How to Change the Event Logging Level for Certificate Services," in the Microsoft Knowledge Base.

**Error when accessing the certification authority Web pages.**
**Cause:** The user accessing the Web pages is not a member of the Administrators or Power Users groups on the local computer. When a newer version of the Web enrollment software is available on the CA, the client computer must install that software. The user must be a member of the Administrators or Power Users group to install the software.

**Solution:** Log on as a user that is a member of the Administrators or Power Users group to access the Web enrollment pages and download the newer version of the software.

**Cause:** Web pages aren't installed on the certification authority (CA).

**Solution:** From the command prompt on the CA, run **certutil -vroot** to install the Web enrollment pages.

**Cause:** Web pages don't have *execute script* permissions.

**Solution:** From Internet Information Services (IIS), open the CertSrv folder and confirm that there are *execute script* permissions on the folder. The CertSrv folder is:

*Systemroot*/System32/Certsrv

**Web pages on an enterprise certification authority (CA) either don't generate certificates or generate certificates that are not valid.**
**Cause:** For an enterprise CA, Web pages require that the user be authenticated. If the pages are set to allow anonymous connections, then the CA will either not generate certificates or will generate certificates that are not valid.

**Solution:** See Set security for access to certification authority Web pages.

**Can't log into certification authority (CA) Web pages.**
**Cause:** If the Web pages negotiate basic authentication with the browser, then your Windows 2000 user account must have the user right to log on to the server.

**Solution:** By default, only domain administrators have this user right. You will need to change this default security permission setting for the server hosting the CA if you use the Netscape browser.

**Certification authority (CA) Web enrollment pages that are installed on a remote server other than the CA do not work.**
**Cause:** The pages are set to use NTLM authentication instead of basic authentication.

**Solution:** If the Web pages are located on a different server than the CA, then you must set the pages to use basic authentication rather than NTLM if the CA is an enterprise CA. You should also use Secure Sockets Layers (SSL) to secure these pages to protect the passwords. See the Internet Information Services (IIS) documentation to change these settings.

**Internet Information Services (IIS) 4.0 and Certificate Server 1.0 cannot enroll for certificates using Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition Certificate Services Web pages.**
**Cause:** IIS 4.0 and Certificate Server 1.0 cannot process a certification path, rather than just a single certificate.

**Solution:** Use the command-line utility Certreq to process the enrollment request.

**A user tries to log on with the smart card and receives this message: "The system cannot log you on to this domain because the system's computer account in its primary domain is missing or the password on the account is incorrect."**
**Cause:** The computer account may be disabled or the CA that issued the smart card certificate is not trusted by the computer.

**Solution:**

1. Make sure that the computer account is enabled in the domain.

2. Use the Certificates snap-in to verify that the root CA's certificate is in the **Trusted Root Certification Authorities** store on the user's computer.

3. Use the Certificates snap-in to verify that the domain controller has been issued a domain controller certificate that can be verified to a trusted root.

**After renewing a CA, computers are no longer automatically enrolling for certificates from that CA.**
**Cause:** CA renewal requires that all automatic certificate enrollment objects that enroll for certificates from that CA be recreated.

**Solution:** See Create an automatic certificate request for computers in a Group Policy object.

**When trying to enroll for a certificate from a computer or account belonging to a child domain of the domain where the CA is located, the following error appears: "No template could be found. There are no CAs from which you have permission to request a certificate, or an error occurred while accessing the Active Directory."**
**Cause:** You don't have the proper security permissions set on the certificate templates.

**Solution:** Modify the security permissions for the certificate templates to include the child domain accounts from which you want to allow enrollment. To set access control for certificate templates, see Controlling enrollment access to certificate templates. Some access control caches must time out after making changes to security permissions, so you have to wait a short period of time before seeing the new security permissions replicate through the network.

# Certificate Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate Templates**
Certificate Templates allow customization of certificates issued by Certificate Services, including both how certificates are issued and what they contain.

Although all issued certificates are based on a certificate template, certificate templates can only be issued by an enterprise certification authority running on Microsoft® Windows Server™ 2003, Enterprise Edition, or Microsoft® Windows Server™ 2003, Datacenter Edition.

- For tips about using Certificate Templates, see Certificate Templates Best practices.
- For help with specific tasks, see Certificate Templates How To....
- For general background information, see Certificate Templates Concepts.
- For problem-solving instructions, see Certificate Templates Troubleshooting.

**Note**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

# Certificate Templates Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Do not delete the Certificate Publishers security group**

- The Certificate Publishers security group contains each certification authority's computer account and is used when publishing certificate templates to Active Directory. If this group is removed, the certification authority may not publish certificates to Active Directory correctly. To avoid this, the group should not be deleted and its membership should not be modified.

**Do not exceed the certificate lifetime of the issuing certification authority**

- Certificate lifetimes work as a subset of the certification authority's (CA) certificate lifetime. All certificates, including the CA certificate, have an expiration date that they are no longer valid after. As a result, a certificate cannot be issued with a lifetime that exceeds the lifetime of the issuing CA. Issuing such a certificate would allow it to be valid for longer than the issuing CA certificate, which is impossible. A CA will therefore continue to issue certificates until the CA's certificate expires or until the requested template's renewal period is greater than the CA's certificate remaining lifetime.

**Plan certificate templates before deployment**

- Certificates can be issued to subjects in many ways, including manual enrollment, autoenrollment and Web enrollment. In addition, there are many certificate strategies including issuing one all-inclusive certificate to all subjects and issuing several application-specific certificates to subjects as needed. Because there are so many options, planning should be done well in advance of certificate deployment.

**Upgrade the certificate templates in Active Directory before upgrading from Windows 2000**

- Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, certification authorities use newer versions of certificate templates than those used by Windows 2000. These templates must be upgraded before the certification authorities are upgraded to ensure proper operation.

  For more information, see Install new templates and upgrade existing templates.

**Duplicate new templates from existing templates closest in function to the intended template**

- New certificate templates are duplicated from existing templates. Many settings are copied from the original template. Because of this, duplicating one template to another of a totally different type may carry over some unintended settings. When duplicating a template, examine the subject type of the original template and ensure that you duplicate one that has a similar function to that of the intended template. Although most settings for certificate templates can be edited once the template is duplicated, the subject type cannot be changed.

  For a list of certificate templates and their subject types, see Default Certificate Templates.

# Certificate Templates How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install new templates and upgrade existing templates
- Create a new certificate template
- Delete a certificate template
- Add a certificate template to a certification authority
- Allow subjects to request a certificate that is based on the template
- Re-enroll all certificate holders
- Remove a certificate template from a certification authority
- Modify a Certificate Template

# Install new templates and upgrade existing templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install new templates and upgrade existing templates**

1. Open Certificate Templates.

2. When prompted to install new certificate templates, click **OK**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.

- If this procedure has already been performed in your enterprise, you will not receive a prompt when you open Certificate Templates.

- This procedure must be performed before any Windows 2000 enterprise certification authorities are upgraded to Windows Server 2003 family certification authorities. If any Windows 2000 enterprise certification authorities currently exist or have ever existed in your enterprise, you must perform this procedure to ensure the certificate templates are added and updated correctly.

- This procedure must be performed on a certification authority (CA) in the certification hierarchy. Performing this procedure on a different computer may not work correctly.

- Performing this procedure is important to apply the proper security permissions on the existing certificate templates as well as create new certificate templates. If you do not perform this procedure before upgrading your certification authorities to the Windows Server 2003 family, you will be prompted when opening the Certification Authority snap-in.

- You must prepare the Active Directory schema with Adprep.exe for the upgrade before installing the new certificate templates. For more information on Adprep.exe, see Related Topics.

**See Also**
**Concepts**
Version 2 certificate templates
Upgrading from a Windows 2000 domain

# Create a new certificate template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new certificate template**

1. Open Certificate Templates.
2. Right-click the template to copy from and then click **Duplicate Template**.
3. Type a new name for this certificate template.
4. Make any necessary changes and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- New certificate templates must be created by copying an existing template and using the existing template's properties as the default for the new template. Copy the existing certificate template closest to the configuration of the intended new template to minimize the work necessary.
- Although Version 2 templates can be created and duplicated in Windows Server 2003, Standard Edition, certificates based on Version 2 templates can only be issued by a certification authority running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition.

**See Also**
**Concepts**
Delete a certificate template
Supersede templates
Modify a Certificate Template

# Delete a certificate template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a certificate template**

1. Open Certificate Templates
2. Right-click the template, and then click **Delete**.

### Caution

- When you delete a certificate template, certificates can no longer be issued based on it. If you are using Enterprise CAs, this will affect all CAs in the forest. Certificate templates cannot be recovered once they are deleted.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.

**See Also**
**Concepts**
Create a new certificate template
Remove a certificate template from a certification authority

# Add a certificate template to a certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a certificate template to a certification authority**

1. Open **Certification Authority**
2. Right-click the **Certificate Templates** container, click **New**, then click **Certificate Template to Issue**.
3. Select the certificate template and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press **Enter**.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.

**See Also**
**Concepts**
Create a new certificate template
Deploying Certificate Templates
Version 2 certificate templates

# Allow subjects to request a certificate that is based on the template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow subjects to request a certificate that is based on the template**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Security** tab, add the groups, computers, or users that you want.
4. In **Group or user names**, click one of the new objects, and then, on **Permissions for *ObjectName***, under the **Allow** column, select the **Read** and **Enroll** check boxes.
5. Repeat the previous step for each new object.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- The **Autoenroll** permission must also be set if the subject will be using client autoenrollment to obtain certificates. For more information on autoenrollment, see Related Topics.
- To disallow subjects from requesting a certificate based on a template, clear the **Read** and **Enroll** check boxes using the same steps as above.

**See Also**
**Concepts**
Add a certificate template to a certification authority
Re-enroll all certificate holders
Allowing for autoenrollment

# Re-enroll all certificate holders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To re-enroll all certificate holders**

1. Open Certificate Templates.
2. Right-click the template that you want to use and then click **Reenroll All Certificate Holders**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is used when a critical change is made to the certificate template and you want all subjects that hold a certificate that is based on this template to re-enroll as quickly as possible. The next time the subject verifies the version of the certificate against the version of the template on the certification authority, the subject will re-enroll.

**See Also**
**Concepts**
Modify a Certificate Template
Planning for certificate templates

# Remove a certificate template from a certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a certificate template from a certification authority**

1. Open Certification Authority
2. In the console tree, click **Certificate Templates**.
3. In the details pane, right-click the certificate template that you want, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certification Authority, click **Start**, click **Run**, type **certsrv.msc**, and then press Enter.

**See Also**
**Concepts**
Create a new certificate template
Deploying Certificate Templates
Delete a certificate template

# Modify a Certificate Template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modify a certificate template**

- Supersede templates
- Configure a certificate template for key archival and recovery
- Configure a certificate template for client autoenrollment
- Modify enrollment policy
- Create an application policy
- Create an issuance policy
- Add an application policy
- Add an issuance policy
- Make certificate or application policy critical
- Change key usage
- Change basic constraints

# Supersede templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To supersede templates**

1. Open Certificate Templates
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. Click the **Superseded Templates** tab.
4. Click **Add**.
5. Click one or more templates to supersede, then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- This procedure is used when a template replaces another template and you want subjects of certificates based on the old template to obtain new certificates based on the new template. It forces subjects to obtain a new certificate before their configured renewal period.

**See Also**
**Concepts**
Version 2 certificate templates
Re-enroll all certificate holders
Deploying Certificate Templates

# Configure a certificate template for key archival and recovery

Updated: November 7, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a certificate template for key archival and recovery**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Request Handling** tab, select the **Archive subject's encryption private key** check box.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- In addition to this procedure, the certification authority must be configured to archive keys. For more information, see Related Topics.
- Clients must be re-enrolled to receive a certificate that is based on the changed template if they already have a valid certificate that is based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Establishing key options and key archival
Key archival
Re-enroll all certificate holders
Key archival and recovery
Manage Key Archival and Recovery
**Other Resources**
Active Directory Certificate Services PKI - Key Archival and Management

# Configure a certificate template for client autoenrollment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a certificate template for client autoenrollment**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Security** tab, add the users or groups that you want and, under **Allow**, select the **Read**, **Enroll**, and **Autoenroll** check boxes.
4. (Optional) To automatically enroll the subject without any more user input, on the **Request Handling** tab, click **Enroll subject without requiring any user input**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- In addition to this procedure, other settings must be configured to allow client autoenrollment. For more information, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.
- When autoenrollment retrieves a certificate that is configured to be stored on a smart card, you will receive a popup message when the certificate is ready to be stored on the smart card. If the message refers to a type of smart card that you do not have, click **Cancel** until the correct smart card type appears. You then provide the smart card PIN and the certificate will be stored on the smart card.
- When autoenrollment enrolls for a certificate that requires user interaction for the enrollment process, you will receive a **Certificate Enrollment** popup message and an icon will appear on your taskbar. Clicking that icon or message begins the certificate autoenrollment process.

**See Also**
**Concepts**
Certificate Services example implementation: Establishing autoenrollment for user certificates
Version 2 certificate templates
Planning for autoenrollment deployment
Allowing for autoenrollment
Re-enroll all certificate holders

# Modify enrollment policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify enrollment policy**

1. Open Certificate Templates
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. Click the **Issuance Requirements** tab.
4. Provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create an application policy
Re-enroll all certificate holders

# Create an application policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an application policy**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Application Policies**, and then click **Edit**.
4. In **Edit Application Policies Extension**, click **Add**.
5. In **Add Application Policy**, click **New**.
6. Provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create a new object identifier
Re-enroll all certificate holders
Application policies
Version 2 certificate templates

# Create an issuance policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an issuance policy**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Issuance Policies**, and then click **Edit**.
4. In **Edit Issuance Policies Extension**, click **Add**.
5. In **Add Issuance Policy**, click **New**.
6. Provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create a new object identifier
Re-enroll all certificate holders
Application policies
Version 2 certificate templates

# Add an application policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an application policy**

1. Open **Certificate Templates**
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Application Policies** and then click **Edit**.
4. In **Edit Application Policies Extension**, click **Add**.
5. In **Add Application Policy**, click the application policy that you want, then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press **Enter**.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create an application policy
Create a new object identifier
Re-enroll all certificate holders
Version 2 certificate templates

# Add an issuance policy

Updated: November 2, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an issuance policy**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Issuance Policies**, and then click **Edit**.
4. In **Edit Issuance Policies Extension**, click **Add**.
5. In **Add Issuance Policy**, click **New**.
6. Provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Issuance Policies are also known as Certificate Policies.
- You can optionally fill out the **CPS Location** with a URI. For example, the Microsoft CPS is located at http://www.microsoft.com/pki/rms/cps/. For more information, see Creating Certificate Policies and Certificate Practice Statements (http://go.microsoft.com/fwlink/?LinkId=205512).
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create a new object identifier
Re-enroll all certificate holders
Version 2 certificate templates
Issuance policies

# Make certificate or application policy critical

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make certificate or application policy critical**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Issuance Policies** or **Application Policies** and then click **Edit**.
4. Select the **Make this extension critical** check box.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press ENTER.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Certificate Policies are also known as Issuance Policies.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Create a new object identifier
Re-enroll all certificate holders
Version 2 certificate templates
Issuance policies
Application policies

# Change key usage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change key usage**

1. Open Certificate Templates.
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, in the **Extensions included in this template** box, click **Key Usage**, and then click **Edit**.
4. Make any necessary changes.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.
- Changing key usage is affected by the **Purpose** set for this certificate template on the Extensions tab.
- Clients must be re-enrolled to receive a certificate based on the changed template if they already have a valid certificate based on the old template. For more information about re-enrolling clients, see Related Topics.

**See Also**
**Concepts**
Add an application policy
Add an issuance policy
Configuring key usage
Re-enroll all certificate holders
Version 2 certificate templates

# Change basic constraints

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change basic constraints**

1. Open Certificate Templates
2. In the details pane, right-click the certificate template that you want to change, and then click **Properties**.
3. On the **Extensions** tab, click **Basic Constraints** and then click **Edit**.
4. In **Edit Basic Constraints Extension**, provide the requested information.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Certificate Templates, click **Start**, click **Run**, type **certtmpl.msc**, and then press Enter.
- This procedure only works with certificate templates that issue certificates that sign other certificates, such as Cross Certification Authority and Root Certification Authority.
- This procedure is applicable to version 2 templates. For more information about version 2 templates, see Related Topics.

**See Also**
**Concepts**
Renew a subordinate certification authority
Renew a root certification authority
Version 2 certificate templates
Create an application policy
Create an issuance policy
Add an application policy
Add an issuance policy

# Certificate Templates Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Certificate Templates:

- Certificate Templates overview
- Understanding Certificate Templates
- Deploying Certificate Templates
- Using Certificate Templates
- Certificates Resources

# Certificate Templates overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate Templates overview**

Certificate templates are an integral part of an enterprise certification authority (CA). They are an important element of the certificate policy for an environment, which is the set of rules and formats for certificate enrollment, use, and management.

When a CA receives a request for a certificate, groups of rules and settings must be applied to that request to perform the requested function, such as certificate issuance or renewal. These rules can be simple or complex and may apply to all users or specific groups of users. Certificate templates are the sets of rules and settings that are configured on a CA to be applied against incoming certificate requests. Certificate templates also give instructions to the client on how to create and submit a valid certificate request.

Although all issued certificates are based on a certificate template, certificate templates can only be issued by an enterprise CA running on Microsoft® Windows Server™ 2003, Enterprise Edition, or Microsoft® Windows Server™ 2003, Datacenter Edition. The templates are stored in Active Directory for use by every CA in the forest. This allows the CA to always have access to the current standard template and ensures homogenous application of certificate policy across the forest.

# Understanding Certificate Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Certificate Templates**
This section covers:

- Key type and cryptographic service provider type
- Certificate template enrollment policies
- Validity and renewal periods
- Subject name
- Certificate Template Extensions

# Key type and cryptographic service provider type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key type and cryptographic service provider type**

Certificates contain public key information that is used to encrypt or verify the digital signature of information. Clients store this certificate in their certificate store and also store data that indicates which cryptographic service provider (CSP) stores the associated private key. This CSP could store the private key in memory, on disk, or on a hardware key store, such as a smart card. This allows the client to perform any public-key cryptography action based on the key pair. However, keys are created differently, depending on their purpose. Some keys will work for encrypting data but not signing data and vice versa. This is why key type and cryptographic service provider type must be configured correctly.

**Key type**

When a public/private key pair is generated, several types of keys can be created. Keys can be created to allow their use with encryption, digital signatures, or both. Certificate templates can be configured for a key purpose of encryption, signature, signature and encryption, or signature and smartcard logon. This setting is labeled **Purpose** in the Certificate Templates console.

**Cryptographic service provider type**

Cryptographic service providers (CSPs) are hardware and software components of Windows operating systems that provide generic cryptographic functions. These CSPs can be written to provide a variety of encryption and signature algorithms. Each of the CSPs that are configured to be used by a certificate template can potentially support different cryptographic algorithms and, therefore, different key lengths. This means that certificate templates must be configured to support one or more CSPs. Selecting specific CSPs allows the administrator to control what algorithms and key lengths are used with this certificate. Windows Server 2003 family includes a number of CSPs, and others can be added for enhanced functionality.

# Certificate template enrollment policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate template enrollment policies**

Certificates can be requested by any subject that has at least read and enroll permissions for the corresponding certificate template. In some cases, the administrator may want to place some restrictions on the process that occurs after a certificate request is made. This gives the administrator control of what certificates are issued and how the issuance process is implemented. This type of restriction is known as an *enrollment policy*. Enrollment policies can contain requirements for certificate manager approval, multiple authorized signature requirements, and whether application and issuance policies should be implemented for this certificate. For information about modifying enrollment policies, see Modify enrollment policy.

# Validity and renewal periods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Validity and renewal periods**

Certificate-based cryptography uses public-key cryptography to protect and sign data. Over time, evildoers can obtain data protected with the public key and attempt to derive the private key from it. Given enough time and resources, this private key could be compromised, effectively rendering all protected data unprotected. Also, over time, the names guaranteed by a certificate may need to be changed. Because a certificate is a binding between a name and a public key, when either of these change, the certificate should be renewed.

**Validity periods**

Certificates are enabled for a specific length of time, which is the *validity period*. This time is expressed in a length of time beginning from when a certificate is issued. When that length of time is reached, the certificate is no longer valid and cannot be trusted. Because an expired certificate can cause problems, certificates can be renewed to extend their validity period.

**Renewal periods**

A *renewal period* is the amount of time prior to the end of the validity period when the subject will renew the certificate using autoenrollment. Renewing the certificate during this interval ensures that last-minute requests for certificate renewal can be serviced before certificate expiration to allow uninterrupted use of the certificate.

# Subject name

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Subject name**

The holder of the private key associated with a certificate is known as the *subject*. This can be a user, a program, or virtually any object or service. Because the subject can vary greatly depending on who or what it is, some flexibility is needed when providing the subject name in the certificate request. Windows can either build the subject name automatically or request it from the subject manually. If it automatically provides the subject name, Windows obtains the information from Active Directory. You can configure this process to include or exclude information useful in the environment. If configured to manually provide the subject name, the subject supplies that information in the certificate request, for example by using the Web-based enrollment pages.

# Certificate Template Extensions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate Template extensions**

This section covers:

- Issuance policies
- Application policies
- Key usage
- Key archival
- Basic constraints

# Issuance policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Issuance policies**

A certification authority (CA) processes each certificate request by a defined set of rules. The certification authority may issue some certificates with no proof of identification and require subjects of another type to submit some proof. This provides different levels of assurance for different certificates. These levels of assurance are represented in certificates as *issuance policies*.

An *issuance policy* (also known as a *certificate policy*) is a group of administrative rules that is implemented when issuing certificates. They are represented in a certificate by an object identifier (also known as an OID) that is defined at the certification authority. This object identifier is included in the issued certificate. When a subject presents its certificate, it can be examined by the target to verify the issuance policy and determine if that level of issuance policy is sufficient to perform the requested action.

# Application policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application policies**

Application policies give you the important ability to decide which certificates can be used for certain purposes. This allows you to issue certificates widely without being concerned that they are misused for an unintended purpose.

Application policies are settings that inform a target that the subject holds a certificate that can be used to perform a specific task. They are represented in a certificate by an object identifier (also known as an OID) that is defined for a given application. This object identifier is included in the issued certificate. When a subject presents its certificate, it can be examined by the relying party to verify the application policy and determine if it can perform the requested action.

Application policies are sometimes called *extended key usage* or *enhanced key usage*. Because some implementations of public key infrastructure (PKI) applications may not understand application policies, both application policies and enhanced key usage sections appear in certificates issued by a Microsoft CA.

# Key usage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key usage**

A certificate enables the subject to perform a specific task. To help control the usage of a certificate outside its intended purpose, restrictions are automatically placed on certificates. Key usage is a restriction method and determines what a certificate can be used for. This allows the administrator to issue certificates that can only be used for specific tasks or certificates that be used for a broad range of functions.

# Key archival

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key archival**
When subjects lose their private keys, any information that was persistently encrypted with the corresponding public key is inaccessible. To help avoid this, Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, have the ability to archive a subject's keys in its database when certificates are issued. These keys are encrypted and stored by the certification authority (CA). If a subject loses its keys, the information can be retrieved from the database and provided to the subject. This allows the encrypted information to be recovered instead of lost.

**Important**

- Key archival and key recovery are only available on Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, running an enterprise CA.

# Basic constraints

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Basic constraints**

Certification authorities (CAs) must have a certificate before they can issue certificates. They use the private key associated with this certificate to digitally sign issued certificates. When a CA obtains a certificate from another CA, the parent CA may want to control whether that certificate can be used to issue certificates to other certificate servers. This is a basic constraint.

Basic constraints are used to ensure a certificate is only used in certain applications. An example is the path length that can be specified as a basic constraint. A path length is the maximum number of CA certificates above this one in a certification path. This path length constraint is used to ensure that CA certificates can only issue end entity certificates, not CA certificates. This is used to ensure that some certificates are used by a CA and some are used by other subjects.

# Deploying Certificate Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying Certificate Templates**
This section covers:

- Planning for autoenrollment deployment
- Certificate templates deployment
- Planning for certificate templates

# Planning for autoenrollment deployment

Updated: January 18, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for autoenrollment deployment**

Autoenrollment is a useful feature of certification services in Windows XP and Windows Server 2003, Standard Edition. Autoenrollment allows the administrator to configure subjects to automatically enroll for certificates, retrieve issued certificates, and renew expiring certificates without requiring subject interaction. The subject does not need to be aware of any certificate operations, unless you configure the certificate template to interact with the subject.

To properly configure subject autoenrollment, the administrator must plan the appropriate certificate template or templates to use. Several settings in the certificate template directly affect the behavior of subject autoenrollment.

- On the **Request Handling** tab of the selected certificate template, the selection of an autoenrollment user interaction setting will affect autoenrollment:

| Setting | Affect on autoenrollment behavior |
|---|---|
| Enroll subject without requiring any user input | This setting will allow "silent" autoenrollment without requiring the user to take any action. This setting is preferred when clients require certificates but may not be aware that they are using them. |
| Prompt the user during enrollment | The user will receive a message and may need to take an action when enrollment is performed. This action may be necessary when the certificate is intended for a smart card, which would require the user to provide their personal identification (PIN). |
| Prompt the user during enrollment and require user input when the private key is used | This setting prompts the user both during enrollment and whenever the private key is used. This is the most interactive autoenrollment behavior, as it requires the user to confirm all use of the private key. It is also the setting that provides the highest level of user awareness regarding key usage.<br>**Caution**<br>• This setting is provided to the client during certificate enrollment. The client should follow the configuration setting, but the setting is not enforced by the certification authority. |

- On the **Request Handling** tab of the selected certificate template, the number of CSPs that are selected in the **CSP Selection** list box changes the behavior of autoenrollment for smart card templates. If you select more than one smart card CSP, users may receive more than one dialog box when Windows XP retrieves the autoenrolled certificate and begins to install it on the smart card. It is recommended that only one CSP be selected from this list for each template.

- On the **Subject Name** tab of the selected certificate template, selecting the **Supply in the request** option disables autoenrollment based on this template. Enabling the option prompts the subject to interactively create the subject name in the request, which will not work with autoenrollment.

- On the **Issuance Requirements** tab of the selected certificate template, selecting **This number of authorized signatures** and making the value greater than **1** disables subject autoenrollment based on this template.

- On the **Issuance Requirements** tab of the selected certificate template, selecting **This number of authorized signatures** and setting the value to **1** requires the requester to sign the request with a private key from a valid certificate in their certificate store. This certificate must contain the application and issuance policies that are specified in the **Application policy** and *Issuance policies* lists on the same tab. If an appropriate certificate exists in the requester's certificate store, autoenrollment signs the request with this certificate's private key and obtains and installs the requested certificate automatically.

- On the **Issuance Requirements** tab of the selected certificate template, the **Valid existing certificate** option may affect subject autoenrollment. This option tells the CA that the subject does not need to meet issuance requirements when

renewing a valid certificate. Subjects who may have been unable to autoenroll for the initial certificate may be able to use autoenrollment to renew that certificate.

- On the **General** tab of the selected certificate template, the **Validity period** and **Renewal period** settings specify the valid lifetime of the certificate and how long before the end of its lifetime autoenrollment will request a renewal. Because the validity period can be very short and the renewal period may overlap, autoenrollment will not renew a certificate until at least 80% of the certificate lifetime has expired or when the renewal interval period specified on the template is reached, with preference given to the shortest remaining certificate validity time period. For example:

    - For a certificate with a validity period of 10 days and a renewal period set to 3 days, autoenrollment will not renew the certificate until 8 days (80%) has passed.

    - For a certificate with a validity period of 10 days and a renewal period set to 1 day, autoenrollment will not renew the certificate until 9 days (90%) has passed.

  This prevents autoenrollment from endlessly renewing a certificate due to misconfigured validity and renewal period settings.

**Note**

1. Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, is required to configure version 2 certificate templates for autoenrollment requests. However, autoenrollment manages certificates or pending certificate requests based on any version of certificate template.

2. By default, certificate autoenrollment occurs on an eight hour cycle. Also, there is a built-in tolerance for a one-hour clock skew for checking whether a certificate is expired.

# Certificate templates deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate templates deployment**

When deploying a Windows Server 2003 family enterprise certification authority (CA), certificate templates are stored in Active Directory. This centralizes the location for use by all CAs in the enterprise, as well as simplifies replication and security management. This also allows the CA to upgrade current certificate templates in Active Directory during installation. Note that this requires the root domain's Domain Admins group to have full control access to all certificate templates or for this access to have been granted to another user or group.

Once you have planned and created the appropriate certificate templates, they will be replicated automatically to all domain controllers in the enterprise. This replication normally takes approximately eight hours to complete. Because of this interval, you should create the certificate template and allow it to replicate before deploying it to clients. This is best accomplished during an idle time in your environment. Configuring and using certificates before replication is completed can have unwanted effects.

# Planning for certificate templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for certificate templates**

This section covers:

- Certificate templates design considerations
- Configuring a certification authority to support certificate template options

# Certificate templates design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate template design considerations**
Before deploying certificates, you must plan in order to ensure that appropriate certificates are issued to subjects. The certificates must provide the capability to perform the tasks expected by the subject but must not allow the subject to use the certificate for unintended or malicious tasks. Some questions that must be considered at this stage are:

- Who or what needs certificates?
- What should I allow the subjects to do with the certificates?
- How long will the certificates be valid for?
- How many certificates could each subject possibly have?
- How will the subjects obtain the certificates?

Each of these questions is further discussed in the following sections.

**Name considerations**
The holder of the private key associated with a certificate is known as the *subject*. This can be a user, a program, or virtually any object or service. Because the subject can vary greatly depending on who or what it is, you need some flexibility when providing the subject name in the certificate request. Windows can either build the subject name automatically or request it from the subject. If Windows automatically provides the subject name, it obtains the information from Active Directory. You can configure this process to include or exclude information that is useful in the environment. If it is configured to manually provide the subject name, the subject supplies that information in the certificate request using the Web-based enrollment pages.

**Certificate lifetime**
Certificate-based cryptography uses public key cryptography to protect and sign data. Over time, evildoers can obtain data that is protected with the public key and attempt to derive the private key from it. Given enough time and resources, this private key could be compromised, effectively rendering all protected data unprotected. Also, over time, the names guaranteed by a certificate may need to be changed. Because a certificate is a binding between a name and a public key, when either of these change, the certificate should be renewed.

**Number of certificates per subject**
Because subjects may have many uses for their certificates, an administrative decision must be made: whether to issue many specific certificates or a few broad certificates. This decision depends on the environment and the level of administration that you want and the possible effects on the subjects, as well as the effects of multiple certificates on applications that will use them.

One strategy of certificate administration creates a number of granular templates, one for each job function, such as file encryption or code signing. Subjects can then enroll for each certificate as needed for the appropriate function. This allows subjects to start with few certificates and only obtain new certificates that they need over time. The drawback to this strategy is that the subject may end up with a large number of certificates and private keys that become harder to manage over time.

A different strategy creates a few broad certificate templates that encompass job functions for most common groups of subjects. For example, if most employees use their certificates for mail signing and encryption as well as file encryption and IP Security, create one template that allows all those functions in the same certificate. This allows most subjects to obtain a single, all-purpose certificate. The drawback to this strategy is that there is no granular control of the usage of the certificates. The administrator cannot decide that subgroups cannot encrypt mail or use IP Security without modifying the template or changing the strategy.

**Cryptographic service providers and key size**
One or more cryptographic service providers (CSPs) are identified as usable by a template. This allows the administrator to control what types of cryptography subjects can use within an enterprise. This is very useful when security is of paramount importance. Because subjects use the CSP for both portions of any cryptographic service--either encryption and decryption or

signing and confirming signature--it is necessary to ensure that all subjects can use the same CSP. Otherwise, subjects may not be able to communicate with each other or perform the necessary job tasks if the key types between their CSPs are not compatible.

The easiest way to do this is to configure each certificate template to identify exactly one CSP. Which CSP should be identified is up to the administrator and depends on the level of security required, the intended purposes of the certificate, and the presence of security hardware, such as smart cards. In addition, new CSPs can easily be installed and used after the certificate template is deployed.

Each CSP provides one or more cryptographic algorithms for encryption or digital signature. Because each algorithm is different and its strength relies directly on the length of the key it uses, the minimum key size is specified in the certificate template. As a general rule, larger keys provide more protection over shorter keys for the same algorithm. However, larger keys take longer to generate and use. Care should be given to select a minimum key size that ensures the necessary amount of protection without impacting performance.

When selecting the CSP to associate with a certificate template, the use of smart cards is also a consideration. Each type of smart card has an associated CSP that must be selected for the smart card to be used. If the correct smart card CSP is not associated with the template, the smart card will not be recognized and the template will fail.

**Obtaining certificates**

There are many ways that subjects can obtain certificates. The two general categories of these ways are *manual* and *automatic*. The subject can either manually obtain a certificate using a Web page or Microsoft Management Console (MMC) snap-in, or they can obtain it via autoenrollment that has been configured by the administrator. While autoenrollment requires a bit more initial configuration, it is effective for deploying required certificates to users and computers without the subject performing additional steps. For more information on autoenrollment, see Planning for autoenrollment deployment.

# Configuring a certification authority to support certificate template options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring a certification authority to support certificate template options**
Certificate templates are useful configuration elements of a Windows Server 2003, Enterprise Edition, certification authority (CA). Many of the benefits of editable templates work in conjunction with a properly configured certification authority. Although the default configuration of a certification authority will support most of the functions available, there may be deployment considerations which require reconfiguration of the certification authority.

**Key archival and recovery**
When you want key archival and recovery on a certification authority, three complimentary configuration settings must be made.

- The certificate template must be configured to allow key archival. For more information, see Configure a certificate template for key archival and recovery.
- One or more key recovery agents must be identified on the certification authority and key recovery agent certificates must be issued to those agents. For more information, see Certificate Services example implementation: Key archival and recovery and Identify a key recovery agent.
- Key archival must be configured on the certification authority. For more information, see Enable key archival.

**Cryptographic service provider**
When a cryptographic service provider (CSP) is selected for a certificate template, the selected CSP must be installed on the client computer (or device, if not a computer) as well as the computer configuring the certificate template. If the CSP is not installed, it will not be available for subject requests. This will render the certificate template useless until the configuration is corrected. The CSP must also be installed on the computer where the certificate template is edited.

When selecting CSPs for a certificate template, the intended use of the certificate must also be considered. The intended functionality of the certificate must be provided by the CSP for the template to be useful. For example, if a template was created for clients to use with Encrypting File System (EFS) and the CSP selected is Microsoft Base DSS and Diffie-Hellman Cryptographic Provider, which does not provide encryption functionality, the issued certificates will not be able to perform that function.

# Certificate Template Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

For more information on Certificate Templates or Certificate Services, consult the following references:


- Austin, Thomas, *PKI:A Wiley Tech Brief*. John Wiley & Sons, December 2000.
- Feghhi, Jalal and Jalil Feghhi, *Secure Networking with Windows 2000 and Trust Services*. Addison-Wesley Pub Co., February 2001.
- Housley, RP., and T. Polk. 2001. *Planning for PKI*. New York, New York: John Wiley & Sons, Inc.
- Howard, Michael, *Designing Secure Web-Based Applications for Microsoft Windows 2000*. Microsoft Press, August 2000.
- Kaufman, CA., Perilla, R., and M. Speciner. 1995. *Network Security: Private Communication in a Public World*. Upper Saddle River, New Jersey: Parametrically Inc.
- Schneier, Bruce, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. John Wiley & Sons, December 1995.


The following resources have in depth information about planning, deploying and using a public key infrastructure with Windows Certificate Services:

- Microsoft Security on the Microsoft Web site.

For updated technical information, see:

- For more information about certificate services architecture, see "Designing a Public Key Infrastructure" at the Microsoft Windows Resource Kits Web site.
- Microsoft TechNet on the Microsoft Web site.
- Updated technical information on the Updated technical information.

# Using Certificate Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Certificate Templates**
This section covers:

- Version 1 certificate templates
- Version 2 certificate templates
- Default Certificate Templates
- Controlling enrollment access to certificate templates
- Certificate templates and Active Directory
- Configuring a Certificate Template

# Version 1 certificate templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Version 1 certificate templates**

Microsoft certification authorities (CAs) support two types of certificate templates: version 1 and version 2. Version 1 templates are provided for backwards compatibility and support many general needs for subject certification. They are created by default when a certification authority is installed and cannot be modified or removed. They can be duplicated, creating a version 2 template that can be modified.

Windows Server 2003, Standard Edition, supports only version 1 templates. Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, support both version 1 and version 2 templates. For more information on version 2 templates, see Version 2 certificate templates.

# Version 2 certificate templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Version 2 certificate templates**

Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, certification authorities support two types of certificate templates: version 1 and version 2. Version 2 templates are new to the Windows Server 2003 family. They allow customization of most settings in the template. Several preconfigured version 2 templates are supplied in the default configuration, and more can be added as necessary. This allows complete configuration flexibility for administrators.

Version 2 templates are only available as part of a certification authority that is installed as an enterprise certification authority. For that reason, they require Active Directory. Although Version 2 templates can be created and duplicated in the Windows Server 2003 family, certificates that are based on Version 2 templates can only be issued by a certification authority that is running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition.

# Default Certificate Templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Default certificate templates**

Windows Server 2003 family certification authorities come with a number of preconfigured certificate templates that are designed to meet the needs of most organizations. These templates are:

| Name | Description | Key Usage | Subject Type | Published to Active Directory? | Template Version |
|---|---|---|---|---|---|
| Administrator | Allows trust list signing and user authentication | Signature and encryption | User | Yes | 1 |
| Authenticated Session | Allows subject to authenticate to a Web server | Signature | User | No | 1 |
| Basic EFS | Used by Encrypting File System (EFS) to encrypt data | Encryption | User | Yes | 1 |
| CA Exchange | Used to store keys that are configured for private key archival | Encryption | Computer | No | 2 |
| CEP Encryption | Allows the holder to act as a registration authority (RA) for simple certificate enrollment protocol (SCEP) requests | Encryption | Computer | No | 1 |
| Code Signing | Used to digitally sign software | Signature | User | No | 1 |
| Computer | Allows a computer to authenticate itself on the network | Signature and encryption | Computer | No | 1 |
| Cross-Certification Authority | Used for cross-certification and qualified subordination. For more information, see Qualified subordination overview. | Signature | CrossCA | Yes | 2 |
| Directory E-mail Replication | Used to replicate e-mail within Active Directory | Signature and encryption | DirEmailRep | Yes | 2 |
| Domain Controller | All-purpose certificates used by domain controllers | Signature and encryption | DirEmailRep | Yes | 1 |
| Domain Controller Authentication | Used to authenticate Active Directory computers and users | Signature and encryption | Computer | No | 2 |
| EFS Recovery Agent | Allows the subject to decrypt files previously encrypted with EFS | Encryption | User | No | 1 |
| Enrollment Agent | Used to request certificates on behalf of another subject | Signature | User | No | 1 |
| Enrollment Agent (Computer) | Used to request certificates on behalf of another computer subject | Signature | Computer | No | 1 |
| Exchange Enrollment Agent (Offline request) | Used to request certificates on behalf of another subject and supply the subject name in the request | Signature | User | No | 1 |

| | | | | | |
|---|---|---|---|---|---|
| Exchange Signature Only | Used by Microsoft Exchange Key Management Service to issue certificates to Exchange users for digitally signing e-mail | Signature | User | No | 1 |
| Exchange User | Used by Microsoft Exchange Key Management Service to issue certificates to Exchange users for encrypting email | Encryption | User | Yes | 1 |
| IPSEC | Used by IP Security (IPSec) to digitally sign, encrypt, and decrypt network communication | Signature and encryption | Computer | No | 1 |
| IPSEC (Offline request) | Used by IP Security (IPSec) to digitally sign, encrypt and decrypt network communication when the subject name is supplied in the request | Signature and encryption | Computer | No | 1 |
| Key Recovery Agent | Recovers private keys that are archived on the certification authority. For more information, see Key archival and recovery. | Encryption | KRA | No | 2 |
| RAS and IAS Server | Enables RAS and IAS servers to authenticate their identity to other computers | Signature and Encryption | Computer | No | 2 |
| Root Certification Authority | Used to prove the identity of the root certification authority | Signature | CA | No | 1 |
| Router (Offline request) | Used by a router when requested through Simple Certificate Enrollment Protocol (SCEP) from a certification authority that holds a CEP Encryption certificate | Signature and encryption | Computer | No | 1 |
| Smartcard Logon | Allows the holder to authenticate using a smart card | Signature and encryption | User | No | 1 |
| Smartcard User | Allows the holder to authenticate and protect e-mail using a smart card | Signature and encryption | User | Yes | 1 |
| Subordinate Certification Authority | Used to prove the identity of the root certification authority. It is issued by the parent or root certification authority | Signature | CA | No | 1 |
| Trust List Signing | Allows the holder to digitally sign a trust list | Signature | User | No | 1 |
| User | Used by users for e-mail, EFS and client authentication | Signature and encryption | User | Yes | 1 |
| User Signature Only | Allows users to digitally sign data | Signature | User | No | 1 |
| Web Server | Proves the identity of a Web server | Signature and encryption | Computer | No | 1 |
| Workstation Authentication | Enables client computers to authenticate their identity to servers | Signature and encryption | Computer | No | 2 |

For more information on certificate templates, see Certificate Templates. For more information on key usage, see Key usage. For more information on template versions and publishing to Active Directory, see Using Certificate Templates.

# Controlling enrollment access to certificate templates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Controlling enrollment access to certificate templates**

Certificate templates are published on a server. Each contains an access control list (ACL) which defines what specific operations a subject can do with the certificate.

| Setting | Description |
|---|---|
| Full Control | The selected group or user can perform any action on this template. |
| Read | The selected group or user can read this template. |
| Write | The selected group or user can modify this template. |
| Enroll | The selected group or user can submit a certificate issuance or renewal request based on this template. |
| Autoenroll | The selected group or user can submit a certificate request based on this template by way of autoenrollment. This option will not work unless the Enroll option is also selected. |

The most common use of certificates is for subject enrollment with autoenrollment permitted. In this case, the subject must be granted **Read**, **Enroll** and **Autoenroll** permissions. If autoenrollment is not wanted but manual or Web-based enrollment is, granting the **Read** and **Enroll** permissions is appropriate. When subjects already hold a certificate, they only need **Read** and **Enroll** permissions to renew that certificate, whether they use autoenrollment or not.

**Write** and **Full Control** permissions should be restricted to CA managers to ensure the templates are not improperly configured.

# Certificate templates and Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Certificate templates and Active Directory**

A Windows Server 2003 family certification authority can be installed as an enterprise server or a stand-alone server. When installed as an enterprise server, the certification authority is integrated into Active Directory. This allows the version 2 certificate templates to be replicated to all other enterprise certification authorities via the normal Active Directory replication process.

Certificate templates must be available from all certification authorities in the enterprise simultaneously in order to work correctly. As a result, creation or change of a certificate template must replicate completely before that certificate template can be used. This replication can take up to eight hours to complete, depending on your Active Directory implementation. Configuring and using certificates before replication is completed can have unwanted effects.

# Configuring a Certificate Template

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring a certificate template**
This section covers:

- Configuring the subject name
- Configuring the certificate purpose
- Configuring certificate publishing in Active Directory
- Establishing key options and key archival
- Establishing issuance policies
- Establishing application policies
- Establishing authentication and signature requirements for enrollment
- Allowing for autoenrollment
- Configuring key usage

# Configuring the subject name

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the subject name**

When establishing a certificate template, the subject name must be defined. This is included in the issued certificate template and must uniquely identify the subject. The subject name can either be built automatically during the certificate request or explicitly defined by the subject and included in the certificate request.

There are a number of options that can be included with the subject name, as well as specific configuration settings for the subject name. The formats and options are:

**Subject name formats**

| Format | Description |
|---|---|
| None | Does not enforce any name format for this field |
| Common name | The certification authority creates the subject name from the common name (CN) obtained from Active Directory. These should be unique within a domain, but may not be unique within an enterprise. |
| Fully distinguished name | The certification authority creates the subject name from the fully distinguished name obtained from Active Directory. This guarantees that the name is unique within an enterprise. |
| Include e-mail name in subject name | If the e-mail name field is populated in the Active Directory user object, that e-mail name will be included with either the common name or fully distinguished name as part of the subject name. |

**Alternate subject name options**

| Field | Description | Useful for subject types |
|---|---|---|
| E-mail name | If the e-mail name field is populated in the Active Directory user object, that e-mail name will be used. | User |
| DNS name | The fully qualified domain name (FQDN) of the subject that requested the certificate. | Computer |
| User principal name (UPN) | The user principal name is part of the Active Directory user object and will be used. | User |
| Service principal name (SPN) | The service principal name is part of the Active Directory computer object and will be used. | Computer |

**Notes**

- If the Subject Name option is set to **Supply in the request**, one or more Issuance Requirements should be set for the template. If no Issuance Requirements are set, subjects are able to request and obtain certificates that contain any subject name. This would allow subjects to impersonate other subjects easily.

- A subject cannot request a certificate with a different subject name from the subject name of the requestor. That name is obtained through security authentication. The only subject that can request a certificate of this type is one who holds a certificate based on the Enrollment Agent template. That subject can request certificates on behalf of any other subject.

# Configuring the certificate purpose

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the certificate purpose**

The certificate purpose defines the intended primary use of the certificate. The certificate purpose can be one of three settings:

| Purpose | Uses |
|---|---|
| Encryption | Used for encrypting data. A certificate with this purpose will contain cryptographic keys for encryption and decryption. |
| Signature | Used to digitally sign data. A certificate with this purpose will contain cryptographic keys for signing data only. |
| Signature and encryption | Used to either digitally sign data, digitally encrypt data, or both. This certificate purpose covers all primary uses of a certificate's cryptographic key. |
| Signature and smartcard logon | Used for initial logon with a smart card and to digitally sign data. A certificate with this purpose cannot be used for data encryption. |

# Configuring certificate publishing in Active Directory

Updated: November 14, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring certificate publishing in Active Directory**

A Microsoft certification authority (CA) can add certificates that have been issued to Active Directory subjects to the appropriate Active Directory object. This provides other users of Active Directory with the ability to easily locate and use the subject's certificate. There are two settings that affect the way this feature works:

- **Publish certificate in Active Directory**. When a subject obtains a certificate based on this template, the issued certificate will be added to that subject's Active Directory object.

| **Note** |
| --- |
| This setting indicates the certificate issued based on the certificate template should be published to the Active Directory Domain Services (AD DS) database. When this setting is enabled, the user or computer object in the AD DS database is updated with the certificate of the user or computer respectively. The private key is not published to the AD DS database. For both computer and user certificates, the userCertificate attribute of the AD DS object is updated with the certificate. The CA must have write permission to the AD DS database user and computer objects to make this update. The permission to write to the computer and user objects in the AD DS database is granted to CAs through their membership in the **Cert Publishers** group by default. This setting is typically only used with user certificates. When a user's certificate is published in the AD DS database, other users can search the AD DS database to find the certificate of that user. The certificate can then be used to encrypt email or files to the user whose certificate is published in the AD DS database. |

- **Do not automatically re-enroll if a duplicate certificate exists in Active Directory**. When the subject attempts to enroll for a certificate based on this template, computers running Windows XP or Windows Server 2003 will check to see if a duplicate certificate exists in Active Directory. If one does, autoenrollment will not submit a re-enrollment request. This allows certificates to be renewed but prevents multiple duplicate certificates from being issued.

# Establishing key options and key archival

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing key options and key archival**

Windows Server 2003, Standard Edition, establishes public and private key options when issuing certificates to subjects. It also configures management and storage of the private key at that time. There are a number of settings that control the generation and management of certificate-based keys.

| Setting | Description |
|---|---|
| Archive subject's encryption private key | If the issuing Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, certification authority is configured for key archival, the subject's private key will be archived. |
| Include symmetric algorithms allowed by the subject | When the subject requests the certificate, they can supply a list of supported symmetric algorithms. This option allows the issuing certification authority to include those algorithms in the certificate, even if they are not recognized or supported by that server. The algorithms are commonly used by applications like Encrypting File System (EFS) or secure e-mail. |
| Minimum key size | This specifies the minimum size, in bits, of the key that will be generated for this certificate. |
| Cryptographic service providers | This is a list of cryptographic service providers (CSPs) that are installed on the certification authority. Selecting one or more CSPs configures the certificate to only work with those CSPs. If you do not select a CSP, the certificate works with any installed CSP. |
| Allow private key to be exported | When this option is specified, the subject can export their private key for backup or transportation. |

Note that some of these settings are not enforced by the issuing certification authority and may be ignored.

# Establishing issuance policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing issuance policies**

An issuance policy describes the conditions under which a certificate is issued. This provides a level of assurance that the subject's certificate request was verified in a specific way.

One or more issuance policies may be selected for a certificate template. Because these issuance policies are simply text labels with an associated object identifier, there are no options associated with them. The only special issuance policy is **All issuance policies**, which indicates that this policy includes all others. This is normally reserved for certificates held by certification authorities.

# Establishing application policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing application policies**

Application policies, which are called *enhanced key usage* in Windows 2000, define how a certificate can be used. Specifically, application policies are object identifiers that are included in the issued certificate. When a subject uses that certificate, Windows Server 2003, Standard Edition, reads the application policies to determine if this certificate can be used for the current intended purpose. If the intended purpose is included in the application policies, the operation continues. If not, the operation fails.

Because these application policies are simply text labels with an associated object identifier, there are no options associated with them. The only special application policy is **All application policies**, which indicates that this policy includes all others. This is normally reserved for certificates held by certification authorities.

# Establishing authentication and signature requirements for enrollment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing authentication and signature requirements for enrollment**

When subjects submit certificate requests to a certification authority, the request can either be automatically approved or placed into a "pending" state. A pending state is normally used for certificates that require a higher level of assurance and consequently require more administration and further verification of the request. There are a number of settings that can configure the authentication and signature requirements for issuance certificates that are based on a template.

| Setting | Description |
|---|---|
| CA certificate manager approval | All certificates are placed into the pending container for a certificate manager to issue or deny. |
| This number of authorized signatures | This setting requires the certificate request to be digitally signed by one or more subjects before it can be issued. This enables several other configuration parameters. |
| Policy type required in signature | The signatures that are required for issuing a certificate must contain either a specific application policy, issuance policy, or both. This is how the certification authority determines whether the signature is appropriate for authorizing the issuance of the subject's certificate. This option is enabled when **This number of authorized signatures** is set. |
| Application policy | Specifies the application policy that is checked for when signing a certificate request. This option is enabled when **Policy type required in signature** is set to either **Application policy** or **Both application and issuance policy**. |
| Issuance policy | Specifies which issuance policies will be checked when signing a certificate request. This option is enabled when **Policy type required in signature** is set to either **Issuance policy** or **Both application and issuance policy**. |

# Allowing for autoenrollment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Allowing for autoenrollment**

You can use autoenrollment so that subjects automatically enroll for certificates, retrieve issued certificates, and renew expiring certificates without subject interaction. For certificate templates, the intended subjects must have Read, Enroll and Autoenroll permissions before the subjects can enroll. To ensure that unintended subjects cannot request a certificate based on this template, you must identify those unintended subjects and explicitly configure the Deny permission for them. This acts as a safeguard, further ensuring that they cannot even present an unacceptable request to the certification authority. Note that Read permission does not allow enrollment or autoenrollment, it only allows the subject to view the certificate template.

Renewal of existing certificates requires only the Enroll permission for the requesting subject. Certificates obtained in any way, including autoenrollment and manual requests, can be renewed automatically. These types of renewals do not require Autoenroll permission, even if they are renewed automatically.

For more information, see Planning for autoenrollment deployment.

# Configuring key usage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring key usage**

Windows Server 2003, Standard Edition, establishes public and private keys when it issues certificates to subjects. The usage of these keys is configured by key usage and application policy. The key usage configuration is a basic constraint on the broad types of operations that can be performed with this certificate. There are several options, depending on what purpose has been configured for this certificate, either Signature or Encryption.

**Signature**

| Setting | Description |
|---|---|
| Digital signature | Data can be digitally signed by this certificate. This is the most basic signature operation. |
| Signature is proof of origin (nonrepudiation) | Data signed by this certificate can be traced back to the subject that provided the digital signature. This provides nonrepudiation of signatures and is useful in certificate-based transactions. |
| Certificate signing | This certificate can be used to sign other certificates. This option is normally granted to certificate managers and certificates used by certification authorities. |
| Certificate revocation list signing | This certificate can be used to sign certificate revocation lists (CRLs). This option is normally granted to certificate managers and certificates used by certification authorities. |

**Encryption**

| Setting | Description |
|---|---|
| Allow key exchange without key encryption (key agreement) | This option configures the subject to use a key agreement protocol, such as Diffie-Hellman, to establish a symmetric key that can be used to encrypt and decrypt data between the subject and the intended target. |
| Allow key exchange only with key encryption (key encipherment) | This option will configure the subject to establish a symmetric key with a target by generating a symmetric key, encrypting it, and sending it to the target for decryption. |
| Allow encryption of user data | With this option, the subject can use the established symmetric key to encrypt and decrypt actual application data with the same key used for key establishment. This option is only available when **Allow key exchange only with key encryption (key encipherment)** is selected. |

For more information on application policy, see Establishing application policies.

# Certificate Templates Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The Certificate Templates Microsoft Management Console (MMC) does not list any templates after prompting to install new certificate templates.
- Certificates are not being issued to clients.
- Certificates are issued to subjects, but cryptographic operations with those certificates fail.
- Domain controllers are not obtaining a domain controller certificate.
- Clients are unable to obtain certificates via autoenrollment.
- Names of certificate templates in the snap-in are inconsistent between views or windows.
- The private key cannot be exported from smart card certificates, even whenAllow private key to be exportedis selected in the certificate template.
- The certificate template is modified, but some certification authorities (CAs) still have the unmodified version.
- The private key is not being archived even though I selected theArchive subject's encryption private keyoption and configured the CA to require key recovery.
- Autoenrollment is prompting me to renew a certificate that isn't mine, and I have certificates in my private certificate store that I didn't put there.

**The Certificate Templates Microsoft Management Console (MMC) does not list any templates after prompting to install new certificate templates.**
**Cause:** The certificate templates have not yet replicated to the certification authority (CA) that the computer is connected to. This replication is part of Active Directory replication.

**Solution:** Wait for the certificate templates to replicate and then reopen the Certificate Templates MMC.

**Certificates are not being issued to clients.**
**Cause:** The certification authority (CA) issuing certificate has a shorter remaining lifetime than the template overlap period configured for the request certificate template. This means the issued certificate would be immediately eligible for reenrollment. Instead of issuing and endlessly renewing this certificate the certificate request is not processed.

**Solution:** Renew the issuing certificate used by the CA.

**Certificates are issued to subjects, but cryptographic operations with those certificates fail.**
**Cause:** Cryptographic service provider does not match Key Usage settings.

**Solution:** Confirm that you set the cryptographic service provider in the template to one that supports the type of cryptographic operation that the certificate will be used for.

**See also:** Key type and cryptographic service provider type; Modify a Certificate Template

**Domain controllers are not obtaining a domain controller certificate.**
**Cause:** Autoenrollment is turned off by way of Group Policy on domain controllers. Domain controllers obtain their certificates through autoenrollment.

**Solution:** Enable autoenrollment for domain controllers.

**See also:** Modify a Certificate Template

**Cause:** The default Automatic Certificate Request setting for Domain Controllers has been removed from the Default Domain Controllers Policy.

**Solution:** Create a new Automatic Certificate Request in the Default Domain Controllers Policy for the Domain Controller certificate template.

**See also:** Automatic certificate request settings

**Clients are unable to obtain certificates via autoenrollement.**
**Cause:** Security permissions must be set to allow intended subjects to both enroll and autoenroll on the certificate template. Both permissions are required to enable autoenrollment.

**Solution:** Modify the access control list on the certificate template to grant **Read**, **Enroll** and **Autoenroll** permissions for the subjects that you want.

**See also:** Allow subjects to request a certificate that is based on the template

**Names of certificate templates in the snap-in are inconsistent between views or windows.**
**Cause:** Active Directory Sites and Services is being used to view the certificate templates. This tool may not provide as accurate a display as Certificate Templates.

**Solution:** Use the Certificate Templates snap-in to administer certificate templates.

**See also:** Modify a Certificate Template

**The private key cannot be exported from smart card certificates, even when Allow private key to be exported is selected in the certificate template.**
**Cause:** Smart cards do not allow private keys to be exported once they are written to the smart card.

**Solution:** None

**See also:** Smart Cards

**The certificate template is modified, but some certification authorities (CAs) still have the unmodified version.**
**Cause:** Certificate templates are replicated between CAs with the Active Directory replication process. Because this replication is not instantaneous, there may be a short delay before the new version of the template is available on all CAs.

**Solution:** Wait until the modified template is replicated to all CAs. To display the certificate templates that are available on the CA, use the Certutil.exe command.

**See also:** Understanding Sites and Replication and Certutil.

**The private key is not being archived even though I selected the Archive subject's encryption private key option and configured the CA to require key recovery.**
**Cause:** Private keys will not be archived when the key usage for the certificate template is set to **Signature**. This is because the digital signature usage requires the key to not be recoverable.

**Solution:** None

**See also:** Establishing key options and key archival and Key archival and recovery.

**Autoenrollment is prompting me to renew a certificate that isn't mine, and I have certificates in my private certificate store that I didn't put there.**
**Cause:** When using the smart card enrollment station on the administrator's computer to renew or change the certificate stored on the smart card, the certificate from the smart card is copied to the administrator's private certificate store. This certificate may be processed by autoenrollment and prompt you to begin the renewal process.

**Solution:** Click **Start** to begin the autoenrollment renewal process. Because the certificate is not yours, the autoenrollment process will disappear at that point. If you want to remove the certificates from your private store, they can be deleted manually.

**See also:** Delete a certificate.

# Internet Protocol Security (IPSec)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol security (IPSec)**
Internet Protocol security (IPSec) is a framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services.

- For tips about using IPSec, see IPSec Best practices.
- For help with specific tasks, see IPSec How To....
- For general background information, see IPSec Concepts.
- For problem-solving instructions, see IPSec Troubleshooting.

### Note

- Portions of IPSec-related services were jointly developed by Microsoft® and Cisco Systems, Inc.

# IPSec Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Establish an IPSec deployment plan**.

  The deployment plan should address the following considerations: which deployment scenarios (such as server-to-server or remote access) require the use of IPSec, what level of security you require for each scenario, which types of data to secure, which computers to secure, which physical links to secure, who will manage IPSec policies, and how you will provide ongoing support and troubleshooting for end users after IPSec is deployed. The deployment plan should also address broader organizational security issues, such as the vulnerability of the network to specific types of attacks, the use of Active Directory, and how security is applied to Group Policy objects.

  For more information about IPSec security planning, see Establishing an IPSec security plan.

  For more information about Active Directory, see Active Directory Overview.

- **Create and test IPSec policies for each deployment scenario**.

  Before deploying IPSec in a production environment, test the IPSec policies in a realistic lab environment. To obtain realistic performance data, run standard workloads on programs. During initial tests, view packet contents with Network Monitor, or use Authentication Header (AH) or Encapsulating Security Payload (ESP) with null encryption to view packet contents for test environments.

- **Do not use preshared keys**.

  For enhanced security, the use of preshared key authentication is not recommended because it is a relatively weak authentication method. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

  For more information, see Preshared key authentication.

- **Do not use Diffie-Hellman Group 1 (low)**.

  For enhanced security, do not use Diffie-Hellman Group 1, which provides 768 bits of keying strength. For maximum security, use Group 2048 (high), which provides 2,048 bits of keying strength. Use Group 2 (medium), which provides 1,024 bits of keying strength when required for interoperability with Windows 2000 and Windows XP. Strong Diffie-Hellman groups combined with longer key lengths increase the computational difficulty of determining a secret key.

  For more information, see Key exchange methods.

  **Note**

  - Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.

- **Use the Triple Data Encryption Standard (3DES) algorithm for stronger encryption**.

  For enhanced security, when configuring key exchange security methods for IPSec policies, use 3DES, which is a stronger encryption algorithm than DES.

  For information about how to configure key exchange security methods, see Create key exchange security methods.

  **Note**

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

- **Create and assign a persistent IPSec policy for failsafe security**.

  To enhance security, create and assign a persistent IPSec policy, so that computers can be secured if a local IPSec policy or an Active Directory-based IPSec policy cannot be applied. When you create and assign a persistent policy, it is applied before a local policy or an Active Directory-based policy, and it remains in effect regardless of whether the local policy or the Active Directory-based policy is applied (for example, an IPSec policy will not be applied if it is corrupted).

  **Note**

  - You cannot configure this feature in the IP Security Policy Management console. To configure this feature, you must use the Netsh IPSec command-line tool. For more information, see Netsh commands for Internet Protocol security.

- **For computers connected to the Internet, do not send the name of the certification authority (CA) with certificate requests**.

  When certificate authentication is used to establish trust between IPSec peers, each IPSec peer sends to the other peer a list of trusted root CAs from which it accepts a certificate for authentication. Each of these CA names is sent as a certificate request payload (CRP), and it must be sent before trust is established. Although transmitting this list aids in connectivity by facilitating the selection of a CA, it can expose sensitive information about the trust relationships of a computer, such as the name of the company that owns the computer and the domain membership of the computer (if an internal public key infrastructure is being used), to an attacker. Therefore, to secure computers that are connected to the Internet, enable the option to exclude the CA name from the certificate request.

- **For computers connected to the Internet, do not use Kerberos as an authentication method**.

  When Kerberos V5 authentication is used, during main mode negotiation, each IPSec peer sends its computer identity in unencrypted format to the other peer. The computer identity is unencrypted until encryption of the entire identity payload takes place during the authentication phase of the main mode negotiation. An attacker can send an Internet Key Exchange (IKE) packet that causes the responding IPSec peer to expose its computer identity and domain membership. To secure computers that are connected to the Internet, certificate authentication is recommended.

  For more information, see Authentication methods.

- **For computers that are connected to the Internet, do not allow unsecured communication**.

  If you configure a filter action to negotiate IPSec security, ensure that the following options are disabled in order to secure computers that are connected to the Internet:

  - **Accept unsecured communication, but always respond using IPSec**. This option allows initial incoming unsecured traffic (for example, TCP SYN packets) but requires outgoing traffic to be protected. This option should be disabled to prevent denial-of-service attacks.

  - **Allow unsecured communication with non-IPSec-aware computers**. This option allows unsecured communications with computers that cannot negotiate the use of IPSec or process IPSec-secured communications, and it is appropriate only in environments where IPSec-secured communication is not necessary.

    For more information, see Filter action.

- **Restrict the use of administrative credentials in your organization**.

  Members of the local Administrators group can view and modify the IPSec policy settings on their computers. For this reason, and as a general security best practice, ensure that end-users in your organization use the principle of least privilege.

  For information about general best practices for security, see Best practices for security for security.

- **When applying the same IPSec policy to computers running different versions of the Windows operating system, test the policy thoroughly**.

  Several features in the Windows Server 2003 family implementation of IPSec are not provided in Windows 2000 or in Windows XP. These new features include:

  - A stronger Diffie-Hellman group (2,048-bit Diffie-Hellman key exchange).
  - The ability for IPSec policy filters to allow logical addresses for local IP configuration.
  - The removal of default traffic exemptions from IPSec filtering (only Internet Key Exchange traffic is now exempt by default).
  - The ability to exclude the name of the certification authority (CA) from a certificate request.

  To ensure that the same IPSec policy functions as expected on computers running the Windows Server 2003 family and on computers running Windows XPor Windows 2000, test the policy thoroughly on all relevant operating systems before deployment.

  For more information about new features in IPSec, see New features for IPSec.

- **Use the** Windows Server 2003 **IP Security Policy Management console to manage the IPSec policies that use the new features that are available only in the** Windows Server 2003 **family implementation of IPSec**.

  If you plan to apply IPSec policies that use any of the new features that are available only in the Windows Server 2003 family implementation of IPSec, do not use the Windows XP or the Windows 2000 version of the IP Security Policy Management console to manage these policies. The settings in the earlier versions of the IP Security Policy Management console will override the settings in the Windows Server 2003 family IPSec policy, and the new features will not be functional.

- **Use Terminal Services to remotely manage and monitor IPSec on computers running different versions of the Windows operating system**.

  Remote management and monitoring of IPSec is supported only for computers running the same version of the Windows operating system. To remotely manage and monitor IPSec on a computer that is running a different version of Windows than the version of Windows that is running on your computer, use Terminal Services. For example, if your computer is running the Windows Server 2003 family, and you plan to remotely manage and monitor IPSec on computers running Windows 2000 or Windows XP, use Terminal Services to gain remote access to these computers.

  If your computer is running the Windows Server 2003 family, and you plan to manage and monitor IPSec on computers that are also running the Windows Server 2003 family, you can run the IPSec Policy Management console and IP Security Monitor, or you can use the Netsh command-line tool remotely.

# IPSec How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Start the IP Security Policy Management snap-in
- Define IPSec Policies
- Assign IPSec Policy
- Manage IPSec Policy
- Use Resultant Set of Policy (RSoP) to View IPSec Policy Assignments
- Monitor IPSec Activity

# Start the IP Security Policy Management snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the IP Security Policy Management snap-in**

- To start the IP Security Policy Management snap-in from the Microsoft Management Console:

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **IP Security Policy Management**, and then click **Add**.
4. Select the computer for which you want to manage IPSec policies:

| To | Do this |
|---|---|
| Manage only the computer on which this console is running | Click **Local computer** |
| Manage IPSec policies for any domain members | Click **The Active Directory domain of which this computer is a member**. |
| Manage IPSec policies for a domain of which the computer that is running this console is not a member | Click **Another Active Directory Domain**. |
| Manage a remote computer | Click **Another computer**. |

5. Click **Finish**, click **Close**, and then click **OK**.

- To access the IP Security Policy Management snap-in from Group Policy (Active Directory):

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

   **Where?**

   - Active Directory Users and Computers [*DomainControllerName.DomainName*]/*Domain*/*OrganizationalUnit*/*ChildOrganizationalUnit*...

3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the Group Policy object that you want to edit. Or, click **New** to create a new Group Policy object, and then click **Edit**.
5. In the Group Policy console tree, click **IP Security Policies on Active Directory**.

   **Where?**

   - *PolicyName* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/IP Security Policies on Active Directory

- To access the IP Security Policy Management snap-in from Local Computer Policy:

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. Click **Finish**, click **Close**, and then click **OK**.
5. In the Group Policy console tree, click **IP Security Policies on Local Computer**.

   **Where?**

   - Local Computer Policy/Computer Configuration/Windows Settings/Security Settings/IP Security Policies on Local Computer

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To save console settings, on the **File** menu, click **Save**.
- To view the saved console, see Related Topics.
- You cannot administer Active Directory-based IPSec policy from a computer running Windows XP Home Edition.
- To manage policies in a remote domain, you must be using a computer that is a member of a domain that is trusted by the remote domain. You cannot configure policies in a remote domain from a computer that is a member of a workgroup (also known as a stand-alone computer).
- To revise your initial choice, you can start the Microsoft Management Console, add this snap-in again, and save the console again. If you require multiple configurations, you can save the console with another name.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add, edit, or remove IPSec policies
Policy-based security
Open a saved MMC console for local and remote computers
Save an MMC console file

# Define IPSec Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define IPSec policies**

- Add, edit, or remove IPSec policies
- Define IPSec Rules
- Define IPSec Key Exchange Settings

# Add, edit, or remove IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.

2. Choose whether to add a new policy, or edit or remove an existing one:

| To | Do this |
|---|---|
| Add a new policy | In the console tree, click **IP Security Policies on *Name***. Next, on the **Action** menu, click **Create IP Security Policy**. Follow the instructions in the IP Security Policy Wizard until the **Properties** dialog box for your new policy appears. |
| Modify an existing policy | Double-click the policy that you want to modify. |
| Remove a policy | Click the policy that you want to remove, and on the **Action** menu, click **Delete**. |

3. If you are adding or modifying a policy, click the **General** tab, and then in **Name**, type a unique name.

4. In **Description**, type a description of the security policy, such as which groups or domains it affects.

5. To specify how often the IPSec Policy Agent checks for updates, type a value in **Check for policy changes every numberminute(s)**.

6. If you have special requirements for the security of the key exchange, click **Settings**.

7. Click the **Rules** tab, and create or modify any necessary rules for the policy.

8. Activate or deactivate rules as needed.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- For information about defining IPSec key exchange settings, see Related Topics.

- For information about configuring IPSec rules, see Related Topics.

- For information about how to activate or deactivate rules, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec rules
Activate or deactivate IPSec rules
Working with MMC console files
Define IPSec Key Exchange Settings

# Define IPSec Rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define IPSec rules**

- Add, edit, or remove IPSec rules
- Activate or deactivate IPSec rules
- Define IPSec authentication methods
- Specify IPSec connection types
- Specify an IPSec tunnel
- Define IP Filter Lists
- Define Filter Actions

# Add, edit, or remove IPSec rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IPSec rules**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. To add a rule, decide whether you want to use the Create IP Security Rule Wizard or add the rule manually:
   - To add a rule by using the Create IP Security Rule Wizard, confirm that the **Use Add Wizard** check box is selected, click **Add**, and then follow the instructions.
   - To add a rule manually, confirm that the **Use Add Wizard** check box is cleared, click **Add**, and then define settings on the **IP Filter List**, **Filter Action**, **Authentication Methods**, **Tunnel Setting**, and **Connection Type** tabs.
4. To edit a rule, select the rule that you want to edit, click **Edit**, and then modify the rule properties as needed.
5. To remove a rule, select the rule that you want to remove, and then click **Remove**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- New rules are automatically applied to the policy being edited or created.
- The rules for a policy are displayed in IP Security Policies in reverse alphabetical order, based on the name of the filter list selected for each rule. There is no method to specify an order in which to apply the rules in a policy. The IPSec driver automatically orders the rules based on the most specific filter list to the least specific filter list. For example, the IPSec driver places a rule containing a filter list that specifies individual IP addresses and TCP ports above a rule containing a filter list that specifies all addresses on a subnet.
- The default response rule is automatically added to every new IPSec policy. If you do not want this rule to be part of your policy, you can deactivate it by clearing the check box next to **<Dynamic>**. The default response rule cannot be removed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC

Activate or deactivate IPSec rules
Add, edit, or remove filter actions
Define IPSec authentication methods
IPSec Policy Rules
Working with MMC console files

# Activate or deactivate IPSec rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate or deactivate IPSec rules**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. On the **Rules** tab, do one of the following:
   - To activate a rule, select the check box next to the rule that you want to activate.
   - To deactivate a rule, clear the check box next to the rule that you want to deactivate.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Assign or unassign IPSec policy in Group Policy
Assign or unassign IPSec policy on a computer
Working with MMC console files

# Define IPSec authentication methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define IPSec authentication methods**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify.
4. On the **Authentication Methods** tab, click **Add**. Or, if you are reconfiguring an existing method, click the authentication method, and then click **Edit**.
5. Select the authentication method that you want to add or modify:
   - To use the Kerberos V5 security protocol for authentication services, click **Active Directory default (Kerberos V5 protocol)** if this rule applies to computers that are validated by a Windows 2000 or Windows Server 2003 Active Directory domain or a trusted Active Directory domain.
   - To use a public key certificate for authentication services, click **Use a certificate from this certification authority (CA)**, and then click **Browse** to select a root CA.
   - To prevent the name of the CA from being sent with the certificate request, select the **Exclude the CA name from the certificate request** check box.
   - To enable certificate to account mapping, select the **Enable certificate to account mapping** check box.
   - To specify your own key that will be used for authentication, click **Use this string (preshared key)**.
6. To delete an authentication method or change its preference order, do one of the following:
   - To delete the selected method, click **Remove**.
   - To move the selected method up one level, click **Move up**. Repeat until the method is at the required preference level.
   - To move the selected method down one level, click **Move down**. Repeat until the method is at the required preference level.

**Important**

- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a

computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- The Kerberos V5 protocol authentication method is not supported on computers running Windows XP Home Edition.

- If you choose to use a certificate for authentication, you must select a CA (most commonly the root CA for your installed computer certificate). You cannot leave this field blank.

- For preshared key authentication, each IPSec peer must use the same preshared key, or authentication will fail.

- If filters in more than one rule apply to traffic between the same two IP addresses, then the authentication methods list should be identical in both rules. Otherwise, security requests that are initiated by one rule might match the filter of the other rule, but have a different authentication method. In this case, the negotiation will fail. After a rule has initiated a successful main mode security association between two IP addresses, other rules that are filtering different traffic will not renegotiate the authentication and master key (the main mode security association). Instead, these other rules will use the same main mode security association to negotiate security (algorithms and session keys) for traffic that matches their particular filter.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Authentication methods
Working with MMC console files

# Specify IPSec connection types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify IPSec connection types**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify.
4. On the **Connection Type** tab, select the type of network connections to which this rule will apply:
   - To apply this rule to all network connections you have created on the computer, click **All network connections**.
   - To apply this rule to all LAN connections you have created on the computer, click **Local area network (LAN)**.
   - To apply this rule to any remote VPN or dial-up connections you have created on the computer, click **Remote access**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Connection type
Working with MMC console files

# Specify an IPSec tunnel

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify an IPSec tunnel**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify.
4. On the **Tunnel Setting** tab, specify the computer that will be the tunnel endpoint:

| To | Do this |
|---|---|
| Disable tunneling for this rule | Click **This rule does not specify an IPSec tunnel**. |
| Use tunneled communications to a specific tunnel endpoint | Click **The tunnel endpoint is specified by this IP address** and type the IP address of the tunnel endpoint. |

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- Because you cannot mirror filters for tunneled traffic, you must configure two rules. One rule is used for the outbound traffic and the other is used for inbound traffic. For the outbound traffic rule, the tunnel endpoint is the IP address of the IPSec peer on the other end of the tunnel. For the inbound traffic rule, the tunnel endpoint is an IP address configured on the local computer.
- IPSec tunnels are not supported for remote access VPN scenarios. L2TP/IPSec or PPTP should be used for remote access VPN connections. IPSec tunnels are used primarily for interoperability with other routers, gateways, or end-systems that do not support L2TP/IPSec or PPTP connections. IPSec tunnel mode is supported as an advanced feature, used only in gateway-to-gateway (also known as router-to-router) tunneling scenarios and for server-to-server or server-to-gateway configurations.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Virtual private networking with IPSec
Tunnel endpoint

# Define IP Filter Lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define IP filter lists**

- Add, edit, or remove IP filter lists
- Add, edit, or remove IPSec filters
- Configure advanced IPSec filter settings
- Set an active IP filter list for a rule
- Manage multiple IP filter lists

# Add, edit, or remove IP filter lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IP filter lists**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that contains the IP filter list that you want to modify.
4. On the **IP Filter List** tab, do one of the following:
   - To add a new IP filter list, click **Add**.
   - To modify an existing IP filter list, click the filter list that you want to modify, and then click **Edit**.
   - To remove an IP filter list, click the filter list that you want to remove, and then click **Remove**.
5. If you are adding or modifying an IP filter list, type a unique name for the list.
6. Type a description to specify which computers and traffic types the IP filter list encompasses (for example, "Matches all TCP packets between this computer and any other computer").
7. Create or modify the filters as needed.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- For procedures on creating filters and advanced filter settings, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Manage multiple IP filter lists
Add, edit, or remove IPSec filters
Configure advanced IPSec filter settings
Filter list

# Add, edit, or remove IPSec filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IPSec filters**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that contains the IP filter list you want to modify.
4. On the **IP Filter List** tab, double-click the IP filter list that contains the IPSec filter you want to modify.
5. In the **IP Filter List** dialog box, do one of the following:
   - To add a filter, click **Add**.
   - To modify an existing filter, select the filter that you want to modify, and then click **Edit**.
   - To remove an existing filter, select the filter that you want to remove, and then click **Remove**.
6. If you are adding or editing a filter, in **IP Filter Properties**, the **Addresses** tab, and then select the **Source Address**:

| Select | To secure packets from |
|---|---|
| **My IP Address** | All IP addresses on the computer for which you are configuring this filter. |
| **Any IP Address** | Any computer. |
| **A Specific DNS Name** | The Domain Name System (DNS) name that you specify in **Host name**. The DNS name is resolved to its IP addresses, and then filters are automatically created for the resolved IP addresses. This option is only available when creating new filters. |
| **A Specific IP Address** | The IP address that you specify in **IP Address**. |
| **A Specific IP Subnet** | The IP subnet, as defined by the IP address that you specify in **IP Address** and the subnet mask that you specify in **Subnet Mask**. |
| **DNS Servers <dynamic>** | The DNS server(s) for the computer for which you are configuring this filter. Changes in DNS server addresses are automatically detected, and the filter is updated as needed. |
| **WINS Servers <dynamic>** | The WINS server(s) for the computer for which you are configuring this filter. Changes in WINS server addresses are automatically detected, and the filter is updated as needed. |
| **DHCP Server <dynamic>** | The DHCP server for the computer for which you are configuring this filter. Changes in DHCP server addresses are automatically detected, and the filter is updated as needed. |
| **Default Gateway <dynamic>** | The default gateway for the computer for which you are configuring this filter. Changes in default gateway server addresses are automatically detected, and the filter is updated as needed. |

7. Click **Destination Address**, and repeat step 6 for the destination address.

8. Select the appropriate **Mirrored** setting:

| To | Do this |
|---|---|
| Automatically create two filters based on the filter settings, one for traffic to the destination and one for traffic from the destination | Select the **Mirrored** check box. |
| Create a single filter based on the filter settings | Clear the **Mirrored** check box. |
| Create a filter for an IPSec tunnel | Clear the **Mirrored** check box. For IPSec tunnels, you must create two filter lists: one list describes the traffic to be sent through the tunnel (outbound traffic) and another describes the traffic to be received through the tunnel (inbound). Then, create two rules that use the inbound and outbound filter lists in your policy. |

9. On the **Description** tab, in **Description**, type a description for this filter (for example, specify to what computers and traffic types it applies).

10. If you require additional IP filtering by a specific protocol or port number, on the **Protocol** tab, configure advanced filter settings.

**Caution**

Filters are the most important part of IPSec policy for a computer which is protected by IPSec. If you do not specify the proper filters in either client or server policies, or if the IP addresses change before the policy's filters are updated, security might not be provided.

- When adding a new static IP address to a protected computer:
    - Modify the IPSec policy filters on all clients and servers that make security requests to the protected computer. Ensure that those clients have updated their policy before adding the new address.
    - Inspect the policy being used on the protected computer. If the filters specify static IP addresses for local connections, after adding the new IP address to the interface, edit and save the new filter list to include the new static IP address. **My IP Address** filters will be automatically updated when the new static IP address is added.
- If the protected computer is a Web server and your clients use a proxy server, make sure that communications over all network paths are secured by IPSec:
    - Between the Web server and all clients that directly connect
    - Between the Web server and the proxy server
    - Between the proxy server and all clients of the proxy server
- The following filter is internally defined to permit (not secure) Internet Key Exchange (IKE) traffic:
    - Source Address = Any
    - Destination Address = Any
    - Protocol = UDP
    - Source port = 500

- Destination port = 500

**Important**

- In the Windows Server 2003 family, IPSec provides new source and destination address options for defining filters. In addition, the default behavior for traffic exemptions has changed. If you require previously exempted traffic types, you must create filters to permit this traffic. For instructions on how to configure permit filters for previously exempted traffic types and information about other IPSec special considerations, see Related Topics.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- To configure IP protocol, TCP port, or UDP port settings for a filter, see Related Topics.

- For IPSec tunnel rules, only address-based filters are supported. Protocol-specific and port-specific filters are not supported. Tunnel filters should not be mirrored.

- Filters are applied in the order of most-specific filters first. Filters are not applied in the order in which they appear in the list.

- If an outbound packet does not match any filter, it is sent unsecured.

- If an inbound packet does not match any filter, it is permitted.

- If an IKE security request is received, the source IP address of the request is used to find a matching filter. The security action and tunnel setting that is associated with that filter determines the IKE response.

- The **A specific DNS Name** option is used when creating filters to create IP address-based filters by resolving a DNS name to its IP addresses. The use of a computer name is for the one-time resolution of DNS names to IP addresses when creating the filter. The computer name is not used after the filters are created.

- All filters used in tunnel rules are matched first, before end-to-end transport filters are matched.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Special IPSec considerations
Start the IP Security Policy Management snap-in
Open MMC
Configure advanced IPSec filter settings
Filter list
Working with MMC console files
Virtual private networking with IPSec

# Configure advanced IPSec filter settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure advanced IPSec filter settings**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that contains the IP filter list you want to modify.
4. On the **IP Filter List** tab, double-click the IP filter list that contains the IPSec filter you want to modify.
5. In the **IP Filter List** dialog box, click the filter that you want to modify, and then click **Edit**.
6. In **IP Filter Properties**, click the **Protocol** tab, and then specify settings for **Select a protocol type**:

| To filter | Do this |
|---|---|
| Any packets sent on any protocol | Click **Any**. |
| Traffic sent on a specific protocol | Click the protocol. |
| Packets sent on a custom protocol | Click **Other**, and then type the protocol number. |

7. If you chose TCP or UDP, you can also filter packets by the source port.

   - To filter packets sent on any port used by the protocol type that you selected, click **From any port**.
   - To filter packets sent on a specific port used by the protocol type that you selected, click **From this port**, and then type the port number.

   If you chose TCP or UDP, you can also filter packets by the destination port.

   - To filter packets received on any port used by the protocol type that you selected, click **To any port**.
   - To only filter packets received on the port number that you specify, click **To this port**, and then type the port number.

**Caution**

- An erroneous port number might cause communication to be blocked or unsecured. If you are unsure of the correct value, see the Windows Resource Kits in Related Topics.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- For IPSec tunneling, only address-based filters are supported. Protocol-specific and port-specific filters are not supported.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec filters
Filter list
Working with MMC console files
IPSec Resources

# Set an active IP filter list for a rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set an active IP filter list for a rule**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that contains the IP filter list you want to modify.
4. On the **IP Filter List** tab, select the option that precedes the name of the IP filter list that you want to designate as the active filter list for this rule.

   The IP filter list you selected is now the active filter list for the rule.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- If the option preceding the name of an IP filter list is not selected, the IP filter list will not be set as active for the rule.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec filters
Configure advanced IPSec filter settings
Manage multiple IP filter lists
Working with MMC console files

# Manage multiple IP filter lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage multiple IP filter lists**

1. In **IP Security Policies**, in the console tree, click **IP Security Policies on *Name***.
2. On the **Action** menu, click **Manage IP filter lists and filter actions**.
3. On the **Manage IP Filters Lists** tab, click **Add**. Or, if you are modifying an existing list, double-click that list.
4. Define the filters for the filter list.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec filters
Configure advanced IPSec filter settings
Set an active IP filter list for a rule
Filter list
Working with MMC console files

# Define Filter Actions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define filter actions**

- Add, edit, or remove filter actions
- Add, edit, or remove IPSec security methods
- Configure custom IPSec security methods
- Select a filter action for a rule
- Manage multiple filter actions

# Add, edit, or remove filter actions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove filter actions**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify, and then click the **Filter Action** tab.
4. To add a filter action, decide whether you want to use the IP Security Filter Action Wizard or add the filter action manually:
   - To add a filter action by using the IP Security Filter Action Wizard, confirm that the **Use Add Wizard** check box is selected, click **Add**, and then follow the instructions.
   - To add a filter action manually, confirm that the **Use Add Wizard** check box is cleared, click **Add**, and then define settings on the **Security Methods** and **General** tabs.
5. To modify an existing filter action, select the filter action that you want to modify, and then click **Edit**.
6. To remove a filter action, select the filter action that you want to remove, and then click **Remove**.
7. If you are adding or modifying a filter action, choose a filter action type:
   - Click **Permit** to allow receiving or sending of packets in plaintext (unsecured traffic). Security will not be requested for these packets.
   - Click **Block** to discard packets. Security will not be requested for these packets.
   - Click **Negotiate security** to use the list of security methods in **Security method preference order** to provide security for packets that match this filter. Security requests will be accepted for these packets.
8. If you chose **Negotiate security**, add new security methods or edit the existing ones for the filter action.
9. If you do not want to block incoming, unsecured communications, but you want to ensure that all outgoing communications and subsequent two-way communications are secured, select the **Accept unsecured communication, but always respond using IPSec** check box.
10. To enable communication with other computers that do not support IPSec, and ensure that communication continues if there is no response to a request for IPSec negotiation, select the **Allow unsecured communication with non-IPSec-aware computers** check box. After the initial IPSec negotiation has failed, IPSec negotiation will be retried at five minute intervals.
11. To guarantee that no master keys or master keying material will be reused to generate the session key, select the **Use session key perfect forward secrecy (PFS)** check box.
12. On the **General** tab, in **Name**, type a unique name.
13. In **Description**, type a description. For example, you might type what security levels this filter action represents.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- The filter action types **Permit** and **Block** and the options to **Allow unsecured communication with non-IPSec-aware computers** and **Accept unsecured communication, but always respond using IPSec** are not available in the default response rule.

- Session key perfect forward secrecy (PFS) renegotiates new master key keying material during every session rekey operation. This is the most secure setting. However, it adds overhead and increases the time required to complete a rekey operation. Both initiator and responder must be configured with the same setting for successful negotiation.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Select a filter action for a rule
Add, edit, or remove IPSec security methods
Filter action
Working with MMC console files

# Add, edit, or remove IPSec security methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IPSec security methods**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify, and then click the **Filter Action**tab.
4. Double-click the filter action that you want to modify.
5. On the **Security Methods** tab, click **Negotiate security**, and then do one of the following:
   - To add a new security method, click **Add**.
   - To modify an existing security method, click the security method that you want to modify, and then click **Edit**.
   - To remove a security method, click the security method that you want to remove, and then click **Remove**.
6. If you are adding or modifying a security method, on the **Security Method** tab, select the security level:
   - To use the ESP protocol to provide data confidentiality (encryption) with the triple Data Encryption Standard (3DES) algorithm, data integrity and authentication with the Secure Hash Algorithm 1 (SHA1) integrity algorithm, and default key lifetimes (100MB, 1 hour), click **Integrity and Encryption**.
   - To use the ESP protocol to provide data integrity and authentication with the SHA1 integrity algorithm and default key lifetimes (100MB, 1 hour), click **Integrity only**. ESP is not configured to provide data confidentiality (encryption).
   - Click **Custom**, and then click **Settings** to configure a custom security method or key lifetimes. For more information, see Related Topics.
7. Repeat this procedure to configure additional security methods for the filter action.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should

only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 server operating system support 3DES and do not require installation of the High Encryption Pack.

- Click **Move up** to move the selected security method up one level. Repeat until the security method is at the required preference level.

- Click **Move down** to move the selected security method down one level. Repeat until the security method is at the required preference level.

- For information about how key lifetimes affect security, see **Filter action** in Related Topics.

- When negotiating security, the initiator proposes this list of security methods to the responder. The responder agrees to the first method (from the top down) that is found anywhere in its own security methods list. If none of the proposed security methods match an entry in the responder's list, then the responder fails the negotiation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Configure custom IPSec security methods
Define IPSec Key Exchange Settings
Filter action
Working with MMC console files

# Configure custom IPSec security methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure custom IPSec security methods**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify, and then click the **Filter Action** tab.
4. Double-click the filter action that you want to modify.
5. On the **Security Methods** tab, do one of the following:
   - To add a new security method, click **Add**.
   - To modify an existing security method, click the security method that you want to modify, and then click **Edit**.
   - To remove a security method, click the security method that you want to remove, and then click **Remove.**
6. If you are adding or modifying a security method, on the **Security Method** tab, click **Custom**, and then click **Settings**.
7. To provide integrity for both the packet's addressing information (IP header) and the data, select the **Data and address integrity without encryption (AH)** check box, and then, in **Integrity algorithm**, click one of the following:
   - **MD5** to use the Message Digest 5 (MD5) integrity algorithm, which uses a 128-bit key.
   - **SHA1** to use the Secure Hash Algorithm 1 (SHA1) integrity algorithm, which uses a 160-bit key.
8. To provide both integrity and encryption (confidentiality) for the data, select the **Data integrity and encryption (ESP)** check box, and then, in **Integrity algorithm**, click one of the following:
   - **<None>** to use no data integrity.

     If you have enabled **AH**, you can click **None** for the ESP integrity algorithm to improve performance.
   - **MD5** to use the MD5 integrity algorithm, which uses a 128-bit key.
   - **SHA1** to use the SHA1 integrity algorithm, which uses a 160-bit key.
9. If you selected the **Data integrity and encryption (ESP)** check box, in **Encryption algorithm**, click one of the following:
   - **<None>** to use no encryption.
   - **DES** to use Data Encryption Standard (DES) with a 56-bit key.
   - **3DES** to use the triple DES (3DES) with three 56-bit keys.

10. To change the default session key lifetime settings, do one of the following:

    - To specify the number of kilobytes of data transferred before a new session key is generated, select the **Generate a new key every** check box, and then type a value in **Kbytes**.

    - To specify the number of seconds that elapse before a new session key is generated, select the **Generate a new key every** check box, and then type a value in **seconds**.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- You cannot choose **<None>** for both the ESP integrity and the encryption algorithm.

- Use caution when specifying lifetimes. The default internal lifetimes should work in most environments and be interoperable with other products. If your security requirements demand shorter lifetime values, you can explicitly set appropriate lifetime values in each security method. The lowest values that are explicitly set on initiator and responder will be used.

- Using shorter lifetime values does not increase the security with which data is protected. Short key lifetimes will only decrease the amount of data that is revealed if an attacker discovers one encryption key. Instead of adjusting the lifetime values lower, consider using a stronger encryption algorithm to protect data, such as 3DES.

- Before changing the default lifetimes, you should understand how lifetimes and rekeying work in detail. For more information, see **Filter action** in Related Topics.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec security methods
Filter action
Working with MMC console files

# Select a filter action for a rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select a filter action for a rule**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify.
4. On the **Filter Action** tab, select the option that precedes the name of the filter action that you want to designate as the active filter action for this rule.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- If the option preceding the name of a filter action is not selected, the filter action will not be set as active for the rule.
- If you have not removed them, you can activate one of the three predefined filter actions: **Permit**, **Request Security (Optional)**, or **Require Security**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove filter actions
Add, edit, or remove IPSec security methods
Working with MMC console files

# Manage multiple filter actions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage multiple filter actions**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, and then click **Manage IP filter lists and filter actions**.
4. On the **Manage Filter Actions** tab, do one of the following:
   - To add a new filter action, click **Add**.
   - To modify an existing filter action, click the filter action that you want to modify, and then click **Edit**.
   - To remove a filter action, click the filter action that you want to remove, and then click **Remove**.
5. If you are adding or modifying a filter action, define settings on the **Security Methods** and **General** tabs as needed.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- For more information about defining settings for filter actions, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove filter actions
Add, edit, or remove IPSec security methods
Configure custom IPSec security methods
Filter action
Working with MMC console files

# Define IPSec Key Exchange Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define IPSec key exchange settings**

- Configure key exchange settings
- Create key exchange security methods

# Configure key exchange settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure key exchange settings**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Click the **General** tab, and then click **Settings**.
4. To force reauthentication and the negotiation of new master key keying material each time a new session key is required, click **Master key perfect forward secrecy (PFS)**.
5. If you require a different setting, type a value in **Authenticate and generate a new key after every number minutes**, which will cause a reauthentication and new master key generation at that interval.
6. If you require a different setting, type a value in **Authenticate and generate a new key after every number sessions** to set a maximum limit on the number of times a master key or its base keying material can be reused to generate the session key. This limit, when reached, will force a reauthentication and new master key generation. If you have enabled **Master key perfect forward secrecy (PFS)**, the number of sessions is set to 1 and is not configurable.
7. If you have special requirements for security methods on the master key exchange, click **Methods**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- Portions of IPSec-related services were jointly developed by Microsoft and Cisco Systems, Inc.
- Use **Master key perfect forward secrecy (PFS)** where it is required for interoperability. The default setting is disabled, which should work in most environments and be interoperable with most other products.
- A session limit of zero (0) will cause rekeys to be determined only by the time setting.
- Enabling master key perfect forward secrecy (PFS) might impact performance because each quick mode requires a new main mode negotiation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in

Open MMC
Create key exchange security methods
Key exchange settings
Working with MMC console files

# Create key exchange security methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create key exchange security methods**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Click the **General** tab, click **Settings**, and then click **Methods**.
4. In the **Key Exchange Security Methods** dialog box, do one of the following:
   - To add a new key exchange security method, click **Add**.
   - To modify an existing key exchange security method, click the security method that you want to modify, and then click **Edit**.
   - To remove a key exchange security method, click the security method that you want to remove, and then click **Remove**.
5. If you are adding or modifying a key exchange security method, in the **IKE Security Algorithms** dialog box, select an **Integrity algorithm**:
   - Click **MD5** to use a 128-bit key (faster).
   - Click **SHA1** to use a 160-bit key (stronger).
6. Select an **Encryption algorithm**:
   - Click **3DES** to use the triple Data Encryption Standard (3DES) algorithm and three unique 56-bit keys.
   - Click **DES** to use the DES algorithm and a single 56-bit key. Use this option if you are required to connect to computers that do not have 3DES or if you do not need the higher security and overhead of 3DES. For more information, see Notes.
7. Select a **Diffie-Hellman group** to set the length of base keying material used to generate the actual keys:
   - Click **Low (1)** to generate 768 bits of master key keying material.
   - Click **Medium (2)** to generate 1,024 bits of master key keying material (stronger).
   - Click **High (2048)** to generate 2,048 bits of master key keying material (strongest).

**Important**

- For enhanced security, do not use Diffie-Hellman Group 1 (low). For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP . When you use a stronger group, the secret key that is derived from the Diffie-Hellman exchange has greater strength. For more

information about Diffie-Hellman groups and key exchange methods, see Related Topics.

### Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- Portions of IPSec-related services were jointly developed by Microsoft and Cisco Systems, Inc.

- A key exchange security method is a combination of three settings (integrity algorithm, encryption algorithm, and Diffie-Hellman group). The initiator and the responder must have a method in common (one that uses the same settings) for negotiations to succeed.

- Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

- Click **Move up** to move a selected security method up one level. Repeat until the security method is at the required preference level.

- Click **Move down** to move a selected security method down one level. Repeat until the security method is at the required preference level.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Configure key exchange settings
Key exchange methods
Working with MMC console files
Key management and protection

# Assign IPSec Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Assign IPSec policy**

- Assign or unassign IPSec policy in Group Policy
- Assign or unassign IPSec policy on a computer

# Assign or unassign IPSec policy in Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign or unassign IPSec policy in Group Policy**

- For Active Directory-based Group Policy:

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

   **Where?**

   - Active Directory Users and Computers [*DomainControllerName.DomainName*]/*Domain/OrganizationalUnit/ChildOrganizationalUnit...*

3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the Group Policy object that you want to edit. Or, click **New** to create a new Group Policy object, and then click **Edit**.
5. In the Group Policy console tree, click **IP Security Policies on Active Directory**.

   **Where?**

   - *PolicyName* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/IP Security Policies on Active Directory

6. In the details pane, click the IPSec policy that you want to assign or unassign, and then do one of the following:

   a. To assign the policy, click the **Action** menu, and then click **Assign**.

   b. To unassign the policy, click the **Action** menu, and then click **Unassign**.

- For local computer policy:

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. Click **Finish**, click **Close**, and then click **OK**.
5. In the Group Policy console tree, click **IP Security Policies on Local Computer**.

   **Where?**

   - Local Computer Policy/Computer Configuration/Windows Settings/Security Settings/IP Security Policies on Local

Computer

6. In the details pane, click the IPSec policy that you want to assign or unassign, and then do one of the following:

   a. To assign the policy, click the **Action** menu, and then click **Assign**.

   b. To unassign the policy, click the **Action** menu, and then click **Un-assign**.

**Important**

- An IPSec policy might remain active even after the IPSec policy or Group Policy object to which it is assigned has been deleted. Therefore, you should unassign the IPSec policy before you delete either the policy or the Group Policy object. To prevent problems, use the following procedure:

1. Unassign the IPSec policy in the Group Policy object.

2. Wait 24 hours to ensure that the change is propagated.

3. Delete the IPSec policy or Group Policy object.

   If you delete the IPSec policy or Group Policy object without following this procedure, computers in the Active Directory container to which the IPSec policy is assigned might treat the IPSec policy as if it cannot be located and continue to use a cached copy.

**Notes**

- To manage Active Directory-based IPSec policy, you must be a member of the Domain Admins group in Active Directory, or you must have permission to edit Group Policy objects (for information about Group Policy object editing, see Related Topics). To manage local or remote IPSec policy for a computer, you must be a member of the Administrators group on the local or remote computer. For more information, see Default local groups and Default groups.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- You cannot administer Active Directory-based IPSec policy from a computer running Windows XP Home Edition.

- The settings will take effect the next time Group Policy is refreshed.

- If a policy is currently assigned and you assign a new policy, the currently assigned policy is automatically unassigned.

- You cannot assign a policy from the IP Security Policies on Active Directory console; you can only configure policy. To assign policy for Active Directory containers, you must use **IP Security Policies on Active Directory** within the Group Policy console.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add, edit, or remove IPSec policies
Creating, modifying, and assigning IPSec policies
Working with MMC console files
Group Policy (pre-GPMC)

# Assign or unassign IPSec policy on a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign or unassign IPSec policy on a computer**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. In the details pane, click the IPSec policy that you want to assign or unassign, and then do one of the following:
   - To assign the policy, click the **Action** menu, and then click **Assign**.
   - To unassign the policy, click the **Action** menu, and then click **Un-assign**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- If a policy is currently assigned and you assign a new policy, the currently assigned policy is automatically unassigned.
- If the computer is a member of a Windows 2000 or Windows Server 2003 domain and IPSec policy is assigned at the domain level, then the assigned local computer policy is not used.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Assign or unassign IPSec policy in Group Policy
Add, edit, or remove IPSec policies
Working with MMC console files

# Manage IPSec Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage IPSec policy**

- Restart the IPSec services
- Restore default IPSec policies
- Export IPSec policies
- Import IPSec policies
- Refresh the list of IPSec policies
- Delete an IPSec policy
- Rename an IPSec policy
- Export the list of IPSec policies

# Restart the IPSec services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To restart the IPSec service

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Services.
2. Right-click **IPSec services**.
3. Click **Restart**.

### Important

- Stopping and restarting the IPSec service can disconnect all computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. For information about how to troubleshoot and resolve loss of connectivity in this situation, see Related Topics.

### Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.
- You might need to restart the IPSec service in order to do any of the following:
    - Clear old security associations (SAs).
    - Clear soft SAs (unsecured negotiations).
    - Force an IPSec policy to download from Active Directory to domain clients.
- If the IPSec service does not start, then use Event Viewer to determine possible causes of failure. For more information, see Related Topics.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **net stop policyagent & net start policyagent**

**Important**

- Stopping and restarting the IPSec service can disconnect all computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. For information about how to troubleshoot and resolve loss of connectivity in this situation, see Related Topics.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.

- On a computer running the Routing and Remote Access service and either Windows 2000 Server or a Windows Server 2003 operating system, you must use the following sequence when restarting the IPSec service:

  1. Stop the Routing and Remote Access service using the **net stop remoteaccess** command.

  2. Stop and restart the IPSec service using the **net stop policyagent & net start policyagent** command.

  3. Start the Routing and Remote Access service using the **net start remoteaccess** command.

- You might need to restart the IPSec service in order to do any of the following:

  - Clear old security associations (SAs).

  - Clear soft SAs (unsecured negotiations).

  - Force an IPSec policy to download from Active Directory to domain clients.

- If the IPSec service does not start, then use Event Viewer to determine possible causes of failure. For more information, see Related Topics.

- For more information about the **net start** and **net stop** commands, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
IPSec troubleshooting tools
IPSec Policy Agent service
Event Viewer overview
Net services commands

# Restore default IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore default IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, point to **All Tasks**, and then click **Restore Default Policies**.
4. When prompted, click **Yes**.

### Important

- Restoring default IPSec policies will overwrite any changes to the original default policies, even when the names of the default policies have been changed. If you have modified the default policies and you do not want those modifications to be overwritten, do not restore the default policies.

### Notes

To perform this procedure, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- The option to restore default IPSec policies is only available for computer-based IPSec policy. You cannot restore the default IPSec policies for Active Directory-based IPSec policy.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Assign or unassign IPSec policy in Group Policy
Assign or unassign IPSec policy on a computer
Working with MMC console files

# Export IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, point to **All Tasks**, and then click **Export Policies**.
4. Type the path and the name of the file to which you want the policy information saved, and then click **Save**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- To enhance interoperablity in a mixed environment with computers running Windows 2000, limit the name of the file to which you want the policy information saved to 60 characters.
- To export IPSec policies, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Import IPSec policies
Export the list of IPSec policies
Working with MMC console files

# Import IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, point to **All Tasks**, and then click **Import Policies**.
4. Type the path and the name of the file from which you are importing the IPSec policy information, and then click **Open**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- Importing IPSec policies replaces existing IPSec policies with the same name.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Export IPSec policies
Export the list of IPSec policies
Working with MMC console files

# Refresh the list of IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To refresh the list of IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, and then click **Refresh**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- This procedure is only required to update the list of policies for any additions or deletions made by using another management tool.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Working with MMC console files

# Delete an IPSec policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an IPSec policy**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Click the policy that you want to remove, click **Action**, and then click **Delete**.
3. When prompted, click **Yes**.

### Important

- For Active Directory-based IPSec policy, an IPSec policy might remain active even after the IPSec policy has been deleted. Therefore, you should unassign the IPSec policy before you delete the policy. To prevent problems, use the following procedure:

  1. Unassign the IPSec policy in the Group Policy object.
  2. Wait 24 hours to ensure that the change is propagated.
  3. Delete the IPSec policy.

  If you delete the IPSec policy without following this procedure, computers in the Active Directory container to which the IPSec policy is assigned might treat the IPSec policy as if it cannot be located and continue to use a cached copy.

### Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Refresh the list of IPSec policies
Working with MMC console files

# Rename an IPSec policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename an IPSec policy**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Click the policy that you want to rename, click **Action**, and then click **Rename**.
3. Type a new name, and then press **Enter**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Refresh the list of IPSec policies
Working with MMC console files

# Export the list of IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export the list of IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, and then click **Export List**.
4. Type the file name to which you want to save the selected IPSec policy or the displayed list of IPSec policies, and then click **Save**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- This procedure does not export the IPSec policies, only a text list of the policies displayed in the details pane.
- To export IPSec policies, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Import IPSec policies
Export IPSec policies
Working with MMC console files

# Use Resultant Set of Policy (RSoP) to View IPSec Policy Assignments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Resultant Set of Policy (RSoP) to view IPSec policy assignments**

- Use RSoP to view IPSec policy assignments for a computer
- Use RSoP to view IPSec policy assignments for members of a Group Policy container

# Use RSoP to view IPSec policy assignments for a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use RSoP to view IPSec policy assignments for a computer**

1. Open Resultant Set of Policy.

   An RSoP query runs in logging mode. For more information, see Related Topics.

2. In the console tree, click **IP Security Policies**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/IP Security Policies

   The list of all IPSec policies is displayed in the details pane.

3. To view detailed policy settings for the IPSec policy that is being applied, in the details pane, right-click the policy with the highest precedence, and then click **Properties**.

## Notes

- To open RSoP, click **Start**, click **Run**, type **rsop.msc**, and then click **OK**.
- To run an RSoP query in logging mode, see Related Topics.
- The RSoP console displays detailed IPSec policy settings for only the policy that is being applied. The detailed policy settings include the filter rules, filter actions, authentication methods, tunnel endpoints, and connection type.
- The details pane in the RSoP console provides two ways to indicate which IPSec policy is being applied:
  - The lower the number in the **Precedence** column, the higher the precedence of the IPSec policy. The IPSec policy with a precedence number of 1 is the policy that is being applied.
  - In the **Name** column, the folder icon that precedes the name of the IPSec policy that is being applied contains a green dot.
- The settings of the IPSec policy with the highest precedence apply in their entirety; they are not merged with the settings of IPSec policies that are assigned at higher levels of the Active Directory hierarchy. For more information about using RSoP to view IPSec policy assignments, see Related Topics.
- If a local user overrides domain-level IPSec policies with a local IPSec policy, RSoP does not detect the application of the local IPSec policy. As a result, no IPSec policies will be displayed in the RSoP console when a logging mode query is run.
- To hide or display columns in the RSoP console, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Run RSoP

Access RSoP data for an existing computer and user (logging mode)

RSoP logging mode

Hide or display columns in an MMC console

Using Resultant Set of Policy to view IPSec policy assignments

# Use RSoP to view IPSec policy assignments for members of a Group Policy container

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use RSoP to view IPSec policy assignments for members of a Group Policy container**

1. On a domain controller, open Resultant Set of Policy (RSoP) as a Microsoft Management Console (MMC). For more information, see Related Topics.

2. On the **Action** menu, click **Generate RSoP Data**.

3. On the **Welcome to the Resultant Set of Policy Wizard** page, click **Next**.

4. On the **Mode Selection** page, click **Planning mode**, and then follow the instructions in the wizard to run a planning mode query. For more information, see Related Topics.

5. In the console tree, click **IP Security Policies**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/IP Security Policies

   The list of all IPSec policies is displayed in the details pane.

6. To view detailed policy settings for the IPSec policy that is being applied, in the details pane, right-click the policy with the highest precedence, and then click **Properties**.

**Notes**

- The RSoP console displays detailed IPSec policy settings for only the policy that is being applied. The detailed policy settings include the filter rules, filter actions, authentication methods, tunnel endpoints, and connection type.

- The details pane in the RSoP console provides two ways to indicate which IPSec policy is being applied:

  - The lower the number in the **Precedence** column, the higher the precedence of the IPSec policy. The IPSec policy with a precedence number of 1 is the policy that is being applied.

  - In the **Name** column, the folder icon that precedes the name of the IPSec policy that is being applied contains a green dot.

- The settings of the IPSec policy with the highest precedence apply in their entirety; they are not merged with the settings of IPSec policies that are assigned at higher levels of the Active Directory hierarchy. For more information about using RSoP to view IPSec policy assignments, see Related Topics.

- You can run an RSoP planning mode query only on a domain controller (when you run a planning mode query, you must explicitly specify the domain controller name). However, you can specify any IPSec client as the target for the query, provided you have the appropriate permissions to do so.

- To hide or display columns in the RSoP console, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open RSoP as an MMC snap-in
Simulate RSoP data for a computer and user (planning mode)
Hide or display columns in an MMC console
Using Resultant Set of Policy to view IPSec policy assignments

# Monitor IPSec Activity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor IPSec activity**

- Add the IP Security Monitor snap-in
- Add a computer to the IP Security Monitor console
- View details about active IPSec policies
- View IP security statistics
- View main mode details
- View quick mode details
- Find a matching specific filter
- Change the automatic update refresh interval
- Enable DNS name resolution

# Add the IP Security Monitor snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add the IP Security Monitor snap-in**

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **IP Security Monitor**, and then click **Add**.
4. Click **Close**, and then click **OK**.

## Notes

- To save the console settings, on the **File** menu, click **Save**.
- To view the saved console, see Related Topics.
- To revise your initial choice, you can start the Microsoft Management Console, add this snap-in again, and save the console again. If you require multiple configurations, you can save the console with another name.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor

# Add a computer to the IP Security Monitor console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a computer to the IP Security Monitor console**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. In the console tree, right-click **IP Security Monitor**, and then click **Add Computer**.
3. In the **Add Computer** dialog box:
   - For the local computer, click **This computer**.
   - For a remote computer, click **The following computer**, and then type the name of the remote computer. Or, click **Browse** to find it on the network.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.
- If the IPSec service is not started on the computer that is being monitored, the server icon is displayed as a stopped service. To refresh IP Security Monitor after the IPSec service on that computer have been restarted, right-click the computer and then click **Reconnect**.
- IP Security Monitor can be used to monitor IPSec only on computers running Windows XP or a Windows Server 2003 operating system. To monitor IPSec on a computer running Windows 2000, use the **IPSecmon** command at a Windows 2000 command prompt.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# View details about active IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view details about active IPSec policies**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.

2. Double-click **Active Policy**.

   **Where?**

   - IP Security Monitor/*ServerName*/Active Policy

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- For information about adding the IP Security Monitor snap-in, see Related Topics.

- For information about the items displayed in **Active Policy**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# View IP security statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view IP security statistics**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Click a server name.

   **Where?**

   - IP Security Monitor/*ServerName*

3. Do one of the following:

   - To view main mode statistics, in the console tree, double-click **Main Mode**, and then double-click **Statistics**.
   - To view quick mode statistics, in the console tree, double-click **Quick Mode**, and then double-click **Statistics**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.
- For more information about the items displayed in **Statistics**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in
Viewing main mode and quick mode statistics in IP Security Monitor

# View main mode details

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view main mode details**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Double-click **Main Mode**.

   **Where?**

   - IP Security Monitor/*ServerName*/Main Mode

3. To view main mode details, do one of the following:

   - For generic filter details, in the console tree, click **Generic Filters**, and then in the details pane, double-click a filter.
   - For specific filter details, in the console tree, click **Specific Filters**. To find a specific main mode filter, right-click **Specific Filters**, and then click **Find matching filters**.
   - For IKE policy details, in the console tree, click **IKE Policies**, and then in the details pane, double-click a policy.
   - For statistical details, in the console tree, click **Statistics**.
   - For security association details, in the console tree, click **Security Associations**, and then in the details pane, double-click a security association.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# View quick mode details

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view quick mode details**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Double-click **Quick Mode**.

   **Where?**

   - IP Security Monitor/*ServerName*/Quick Mode

3. To view quick mode details, do one of the following:

   - For generic filter details, in the console tree, click **Generic Filters**, and then in the details pane, double-click a filter.
   - For specific filter details, in the console tree, click **Specific Filters**. To find a specific quick mode filter, right-click **Specific Filters**, and then click **Find matching filters**.
   - For negotiation policy details, in the console tree, click **Negotiation Policies**, and then in the details pane, double-click a policy.
   - For statistical details, in the console tree, click **Statistics**.
   - For security association details, in the console tree, click **Security Associations**, and then in the details pane, double-click a security association.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.
- Quick mode details do not display **Block** or **Permit** filters.
- Although there are two separate quick mode SAs established (one for the inbound traffic and one for the outbound traffic), IP Security Monitor only displays a single quick mode SA.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file

Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# Find a matching specific filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a matching specific filter**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. For a specific main mode filter:
   a. Double-click **Main Mode**.(console tree location is: IP Security Monitor/*ServerName*/Main Mode).
   b. In the console tree, right-click **Specific Filters**, and then click **Find matching filters**.
   c. In the **Find Matching Main Mode Filters** dialog box, configure the filter settings, and then click **Search**.
3. For a specific quick mode filter:
   a. Double-click **Quick Mode** (console tree location is: IP Security Monitor/*ServerName*/Quick Mode).
   b. In the console tree, right-click **Specific Filters**, and then click **Find matching filters**.
   c. In the **Find Matching Quick Mode Filters** dialog box, configure the filter settings, and then click **Search**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# Change the automatic update refresh interval

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the automatic update refresh interval**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Right-click a server name.

   **Where?**

   - IP Security Monitor/*ServerName*

3. Click **Properties**.
4. Configure the automatic update refresh interval, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# Enable DNS name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable DNS name resolution**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Right-click a server name.

   **Where?**

   - IP Security Monitor/*ServerName*

3. Click **Properties**.
4. Select the **Enable DNS name resolution** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.
- The **Enable DNS name resolution** check box specifies whether you want to use Domain Name System (DNS) reverse queries and NetBIOS Adapter Status messages to resolve the DNS or NetBIOS computer name that is associated with an IP address when displaying sources or destinations. To resolve the DNS computer name from its IP address, your DNS infrastructure must have the appropriate reverse domains and pointer records configured, either manually or with DNS dynamic update. To resolve the NetBIOS computer name from its IP address, NetBIOS over TCP/IP must be enabled on the computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# IPSec Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Internet Protocol security (IPSec):

- New features for IPSec
- Internet Protocol Security Overview
- Understanding Internet Protocol Security
- Using Internet Protocol Security
- IPSec Resources

# New features for IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features for IPSec**
Internet Protocol security (IPSec) provides the following new features for enhanced security, scalability, and availability, and ease of deployment and administration.

### Note

- Because several new IPSec features are available only in the Windows Server 2003 family, as a best practice, if you plan to apply the same IPSec policy to computers running the Windows Server 2003 family and to computers running Windows 2000 or Windows XP, test the policy thoroughly on all relevant operating systems before deployment. For additional best practices for using IPSec, see IPSec Best practices.

**IP Security Monitor**
In Windows 2000, IP Security Monitor was implemented as an executable program (IPSecmon.exe). In Windows XP and the Windows Server 2003 family, IP Security Monitor is implemented as a Microsoft Management Console (MMC) console and includes enhancements that allow you to:

- Monitor IPSec information for your local computer and for remote computers.
- View details about active IPSec policies, including the name, description, date last modified, store, path, organizational unit, and Group Policy object name.
- View main mode and quick mode generic filters and specific filters.
- View main mode and quick mode statistics. For information about the statistics displayed in IP Security Monitor, see Viewing main mode and quick mode statistics in IP Security Monitor.
- View main mode and quick mode security associations.
- Customize refresh rates, and use DNS name resolution for filter and security association output.
- Search for specific main mode or quick mode filters that match any source or destination IP address, a source or destination IP address on your local computer, or a specific source or destination IP address.

For information about adding the IP Security Monitor snap-in, see Add the IP Security Monitor snap-in. For information about additional IPSec troubleshooting tools, see IPSec troubleshooting tools.

**Stronger cryptographic master key (Diffie-Hellman)**
For enhanced security, IPSec now supports the use of a 2048-bit Diffie-Hellman key exchange. With a stronger Diffie-Hellman group, the secret key that is derived from the Diffie-Hellman exchange has greater strength. Strong Diffie-Hellman groups combined with longer key lengths increase the computational difficulty of determining a secret key.

This feature is provided only with the Windows Server 2003 family.

For more information about Diffie-Hellman groups and key exchange, see Key exchange methods and Key management and protection.

**Command-line management with Netsh**
Using commands in the Netsh IPSec context, you can configure static or dynamic IPSec main mode settings, quick mode settings, rules, and configuration parameters. To enter the Netsh IPSec context, type **netsh -c ipsec** at a command prompt. The Netsh IPSec context replaces the Ipsecpol.exe tool, which is provided with the Windows 2000 Server Resource Kit. You can use this feature to script and automate IPSec configuration. For more information, see IPSec troubleshooting tools.

This feature is provided only with the Windows Server 2003 family.

For more information about Netsh commands for IPSec, see Netsh commands for Internet Protocol security.

**Computer startup security**
For enhanced security, IPSec now provides stateful filtering of network traffic during computer startup. With stateful filtering, only the following traffic is permitted during computer startup: the outbound traffic that the computer initiates during startup, the inbound traffic that is sent in response to the outbound traffic, and DHCP traffic. As an alternative to stateful filtering, you can specify that all inbound and outbound traffic be blocked until an IPSec policy is applied. If you use stateful filtering, or if you specify that traffic be blocked during computer startup, you can also specify the traffic types that you want to exempt from IPSec filtering during computer startup.

### Note

- You cannot configure this feature in the IP Security Policy Management console. To configure this feature, you must use the Netsh IPSec command-line tool.

**Persistent policy for enhanced security**
You can now create and assign a persistent IPSec policy to secure a computer if a local IPSec policy or an Active Directory-based IPSec policy cannot be applied. When you create and assign a persistent policy, it is applied before the local policy or the Active Directory-based policy is applied, and it remains in effect regardless of whether the local policy or the Active Directory-based policy is applied (for example, an IPSec policy will not be applied if it is corrupted).

### Note

- You cannot configure this feature in the IP Security Policy Management console. To configure this feature, you must use the Netsh commands for IPSec.

**Removed default traffic exemptions**
In Windows 2000 and Windows XP, by default, all broadcast, multicast, Internet Key Exchange (IKE), Kerberos, and Resource Reservation Protocol (RSVP) traffic is exempt from IPSec filtering. To significantly improve security, in the Windows Server 2003 family, only IKE traffic (which is required for establishing IPSec-secured communication) is exempt from IPSec filtering. All other traffic types are now matched against IPSec filters, and you can configure, block, or permit filter actions specifically for multicast and broadcast traffic (IPSec does not negotiate security associations for multicast and broadcast traffic). For more information, see Special IPSec considerations.

**IPSec certificate to account mapping for network access control**
With the Windows Server 2003 family, if you use either Kerberos V5 or certificate authentication, you can set restrictions on which computers are allowed to connect. This functionality allows you to use IPSec to allow or deny any of the following access to a server running Windows Server 2003 :

- Computers that are members of a specific domain.
- Computers that have a certificate from a specific issuing certification authority.
- A specific group of computers.
- A specific computer.

When you enable certificate to account mapping in IPSec, the IKE protocol associates (maps) a computer certificate to a computer account in an Active Directory domain or forest, and then retrieves an access token, which includes the list of the user rights that are assigned to the computer. You can restrict access by configuring Group Policy security settings and assigning either the **Access this computer from the network** user right or the **Deny access to this computer from the network** user right to individual or multiple computers, as needed.

For more information about certificate to account mapping for IPSec, see Authentication methods.

For information about configuring security settings in Group Policy, see Security settings overview.

### Notes

- Certificate to account mapping for IPSec is provided only with the Windows Server 2003 family.

- Certificate mapping is not supported across forests.

**Ability to exclude the name of the certification authority (CA) from certificate requests**
For enhanced security, when you use certificate authentication to establish trust between IPSec peers, you can now exclude the name of the CA from the certificate request. When you exclude the name of the CA from the certificate request, you prevent the potential disclosure of sensitive information about the trust relationships of a computer, such as name of the company that owns the computer and the domain membership of the computer (if an internal public key infrastructure is being used), to an attacker.

**IPSec policy filters allow logical addresses for local IP configuration**
You can now use the IP Security Policy Management console to configure the source or the destination address fields that the local IPSec policy will interpret as the addresses for the DHCP server, the DNS servers, the WINS servers, and the default gateway. As a result, IPSec policies can now automatically accommodate changes in the IP configuration of the server, by using either DHCP or static IP configurations.

This feature is provided only with the Windows Server 2003 family.

**IPSec functionality over network address translation (NAT)**
IPSec Encapsulating Security Payload (ESP) packets can now pass through NATs that allow User Datagram Protocol (UDP) traffic. The IKE protocol automatically detects the presence of a NAT and uses UDP-ESP encapsulation to allow IPSec traffic to pass through the NAT. This functionality is an implementation of the Internet Engineering Task Force (IETF) IP Security Working Group standard for IPSec.

NATs are widely used for Internet Connection Sharing (ICS) and in locations that provide public Internet access (such as hotels and airports) and that are likely to be used by telecommuters. In addition, some Internet service providers (ISPs) use a centralized NAT to connect their clients to the Internet.

IPSec functionality over NAT enables IPSec-secured connections to be established in the following common deployment scenarios:

- Layer Two Tunneling Protocol (L2TP)/IPSec virtual private network (VPN) clients that are behind NATs can establish IPSec-secured connections over the Internet to their corporate network, using IPSec ESP transport mode.

- Servers running Routing and Remote Access can establish gateway-to-gateway IPSec tunnels when one of the servers running Routing and Remote Access is behind a NAT.

- Clients and servers can send IPSec-secured TCP and UDP packets to other clients or servers using IPSec ESP transport mode, when one or both of the computers are behind a NAT. For example, a program running on a server on a perimeter network can be IPSec-secured when it is used to make connections to the corporate network.

### Note

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

For more information about VPNs and IPSec, see Virtual private networking with IPSec. For more information about IPSec transport mode and tunnel mode, see Transport mode and Tunnel mode.

**Improved IPSec integration with Network Load Balancing**
Improved IPSec integration with Network Load Balancing allows a Network Load Balancing group of servers to provide highly available IPSec-based VPN services. Network Load Balancing can accurately track IPSec-secured sessions, and the IPSec IKE protocol can detect when an IPSec-secured session is being established with a cluster server and quickly recover from a failover. Additionally, Network Load Balancing can now maintain IPSec-secured connections to the correct Network Load Balancing host, even when the number of hosts in the cluster (and the algorithm used to map clients to hosts) changes. Because the IKE protocol automatically detects the Network Load Balancing service, no additional configuration is required to use this feature.

This feature is provided only with Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.

For more information about Network Load Balancing, see Introduction to Network Load Balancing and How Network Load

Balancing works.

**IPSec support for RSoP**

To enhance IPSec deployment and troubleshooting, IPSec now provides an extension to the Resultant Set of Policy (RSoP) console. RSoP is an addition to Group Policy that you can use to view existing IPSec policy assignments for a computer or for members of a Group Policy container. To view IPSec policy assignments for a computer, run an RSoP logging mode query. To view IPSec policy assignments for members of a Group Policy container, run an RSoP planning mode query. For information about how to view IPSec policy assignments, see Use Resultant Set of Policy (RSoP) to View IPSec Policy Assignments.

After you run an RSoP logging mode query or an RSoP planning mode query, you can view detailed settings (the filter rules, filter actions, authentication methods, tunnel endpoints, and connection types that were specified when the IPSec policy was created) for the IPSec policy that is being applied.

For more information about using RSoP with IPSec, see Using Resultant Set of Policy to view IPSec policy assignments. For general information about RSoP, see RSoP overview.

# Internet Protocol Security Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol security overview**

This section covers:

- Security information for IPSec
- Internet Protocol security defined
- Introducing IPSec
- Security issues with IP
- Protection against attacks
- Policy-based security
- Simplified deployment

# Security information for IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for IPSec**
Before deploying IPSec in your organization, consider the following security issues:

- 3DES and computers running Microsoft® Windows® 2000
- Authentication method
- Firewall packet filtering
- Protected traffic (using AH or ESP)
- Unprotected traffic (using the default response rule)
- Diffie-Hellman groups (using Diffie-Hellman Group 1, 2, or 2048)

**3DES and computers running Windows 2000**
IPSec policies allow the choice of a strong encryption algorithm, 3DES, which provides stronger encryption than DES for higher security. Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to perform the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, then the 3DES setting is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should use DES only as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP and the Windows Server 2003 family support 3DES and do not require installation of the High Encryption Pack.

**Authentication method**
If the computers in your organization are part of an Active Directory® domain, IPSec main mode authentication can be accomplished with the default authentication method (Kerberos V5). You do not have to deploy public key certificates for intranet communication. However, computers running Windows XP Home Edition do not support the Kerberos V5 authentication method. Also, if you have computers connected to the Internet, it is recommended that you do not use Kerberos V5 as an authentication method. When Kerberos authentication is used, during main mode negotiation, each IPSec peer sends its computer identity in unencrypted format to the other peer. The computer identity is unencrypted until encryption of the entire identity payload takes place during the authentication phase of the main mode negotiation. An attacker can send an Internet Key Exchange (IKE) packet that causes the responding IPSec peer to expose its computer identity and domain membership. To secure computers that are connected to the Internet, certificate authentication is therefore recommended.

For enhanced security, the use of preshared key authentication is not recommended because it is a relatively weak authentication method. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing.

For more information, see Authentication methods and Preshared key authentication.

**Firewall packet filtering**
For a firewall, security gateway, router, or any other server or device that is connected to the Internet and providing packet filtering capabilities to computers on a perimeter network (also known as a demilitarized zone or DMZ), special filtering must be enabled on that computer to ensure that packets secured with IPSec are allowed to be forwarded to computers on the perimeter network. The firewall or other device should normally allow the following types of traffic to pass through it:

- IP Protocol ID of 50 (0x32) for IPSec Encapsulating Security Payload (ESP) traffic
- IP Protocol ID of 51 (0x33) for IPSec Authentication Header (AH) traffic
- UDP port 500 (0x1F4) for Internet Key Exchange (IKE) negotiation traffic

Most packet filtering software allows you to be more specific. You can define separate packet filters for inbound traffic (inbound filters), for outbound traffic (outbound filters), and for each interface. Additionally, you can specify IP addresses for the IPSec computers on the perimeter network. The most restrictive packet filters for IPSec traffic exchanged with a single IPSec computer on the perimeter network are outlined in the following sections.

**Filters on the Internet interface**

Configure the following inbound packet filters on the Internet interface of the firewall to allow the specific types of traffic:

- Destination IP address of the IPSec computer's perimeter network interface and UDP destination port of 500 (0x1F4)

  This filter allows IKE traffic to be sent to the IPSec computer on the perimeter network.

- Destination IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 50 (0x32)

  This filter allows IPSec ESP traffic to be sent to the IPSec computer on the perimeter network.

- Destination IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 51 (0x33)

  This filter allows IPSec AH traffic to be sent to the IPSec computer on the perimeter network.

Configure the following outbound packet filters on the Internet interface of the firewall to allow the specific types of traffic:

- Source IP address of the IPSec computer's perimeter network interface and UDP source port of 500 (0x1F4)

  This filter allows IKE traffic to be sent from the IPSec computer on the perimeter network.

- Source IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 50 (0x32)

  This filter allows IPSec ESP traffic to be sent from the IPSec computer on the perimeter network.

- Source IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 51 (0x33)

  This filter allows IPSec AH traffic to be sent from the IPSec computer on the perimeter network.

**Filters on the perimeter network interface**

Configure the following inbound packet filters on the perimeter network interface of the firewall to allow the specific types of traffic:

- Source IP address of the IPSec computer's perimeter network interface and UDP source port of 500 (0x1F4)

  This filter allows IKE traffic to be sent from the IPSec computer on the perimeter network.

- Source IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 50 (0x32)

  This filter allows IPSec ESP traffic to be sent from the IPSec computer on the perimeter network.

- Source IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 51 (0x33)

  This filter allows IPSec AH traffic to be sent from the IPSec computer on the perimeter network.

Configure the following outbound packet filters on the perimeter network interface of the firewall to allow the specific types of traffic:

- Destination IP address of the IPSec computer's perimeter network interface and UDP destination port of 500 (0x1F4)

  This filter allows IKE traffic to be sent to the IPSec computer on the perimeter network.

- Destination IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 50 (0x32)

This filter allows IPSec ESP traffic to be sent to the IPSec computer on the perimeter network.

- Destination IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 51 (0x33)

  This filter allows IPSec AH traffic to be sent to the IPSec computer on the perimeter network.

**Protected traffic (using AH or ESP)**
When deciding whether to use AH or ESP to protect your IP traffic, consider the following:

- AH provides both data authentication and integrity services through the computation and inclusion of a keyed hash for each packet. With AH, the hash calculation includes the entire IP packet and header. Some fields that are allowed to change in transit are excluded. Packet replay protection services are provided through the inclusion of a sequence number for each packet.

- ESP provides data authentication and integrity services through the computation and inclusion of a keyed hash for each packet. With ESP, the hash calculation only includes the ESP header, trailer, and payload. The IP header is not protected with the hash. ESP provides data confidentiality services by encrypting the ESP payload with the DES (Data Encryption Standard) or 3DES (triple DES) encryption algorithms. Packet replay protection services are provided through the inclusion of a sequence number for each packet.

It is less CPU-intensive to verify the hash for each packet than it is to encrypt and decrypt each packet. If performance is a significant consideration, you can use AH to protect most of your traffic. When confidentiality is required, you can use ESP instead. For example, you can use AH to protect traffic on your intranet and ESP for traffic that is sent over the Internet.

**Note**

- This performance consideration assumes that you are not using IPSec offload network adapters in your organization. Offload network adapters perform cryptographic calculations, such as the calculation and verification of the hash and the encryption and decryption services, on the adapter itself, enhancing performance for IPSec-secured traffic.

**Unprotected traffic (using the default response rule)**
You can easily implement an IPSec deployment scenario that secures traffic from a specific set of servers, as follows:

1. Place the servers in an organizational unit and assign an IPSec policy to a Group Policy object of that organizational unit with a rule that requires secured traffic between the servers and any other computer. Within the filter action for the rule, enable the secure servers to accept unsecured communication but always respond using IPSec.

2. Place the client computers in an organizational unit and assign an IPSec policy to a Group Policy object of that organizational unit that uses only the default response rule. For more information, see IPSec Policy Rules.

In this configuration, the following traffic between the client computers and the secure servers is unsecured:

- The initial request (for example, a TCP SYN packet) sent by the client computer to establish communication is unsecured because the client does not have a rule in its policy to initiate secured communication with the secure server.

- Although ongoing communications are secured after secured communication is negotiated between the client and secure server computers, the quick mode security association (SA) that provides security settings for traffic can eventually time out and be removed from both the secure server and the client computer. This occurs when a TCP connection exists and no traffic is sent between the client computer and the secure server for an extended period of time. The dynamic filter created on the client computer by the default response rule is also removed after the dynamic filter has been timed out.

  If the secure server sends new data on the existing connection, it renegotiates the quick mode SA before sending this data to the client computer because a rule exists on the secure server to secure traffic between itself and all other computers. However, if the client computer sends new data after the quick mode SAs and dynamic filter time out, the new data is sent unsecured because the client does not have a rule to initiate secure communications with the secure

server. For TCP connections, one or multiple TCP segments can be sent unsecured by the client computer. To prevent client computers from sending unsecured data to secure servers, you must configure your client computer IPSec policy with additional rules that initiate secured communications to secure servers.

**Diffie-Hellman groups (using Diffie-Hellman group 1, 2, or 2048)**

Diffie-Hellman groups are used to determine the length of the base prime numbers that are used during the key exchange process. Group 2048 (high) is stronger (more secure) than Group 2 (medium), which is stronger than Group 1 (low). Group 1 provides 768 bits of keying strength, Group 2 provides 1,024 bits, and Group 2048 provides 2,048 bits.

Strong Diffie-Hellman groups combined with longer key lengths increase the computational difficulty of determining a secret key.

**Important**

- For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.

**Note**

- Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.

For more information about Diffie-Hellman groups and key exchange, see Key exchange methods and Key management and protection.

# Internet Protocol security defined

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol security defined**

Internet Protocol security (IPSec) is a framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services. The Microsoft® Windows® 2000 , Windows XP, and the Windows Server 2003 family implementations of IPSec are based on standards developed by the Internet Engineering Task Force (IETF) IPSec working group.

IPSec is the long-term direction for secure networking. It provides aggressive protection against private network and Internet attacks through end-to-end security. The only computers that must know about IPSec protection are the sender and receiver in the communication. In Windows XP and the Windows Server 2003 family, IPSec provides the ability to protect communication between workgroups, local area network computers, domain clients and servers, branch offices (which might be physically remote), extranets, and roving clients.

# Introducing IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introducing IPSec**

IPSec is the long-term direction for secure networking. It provides a key line of defense against private network and Internet attacks.

IPSec has two goals:

1. To protect the contents of IP packets.

2. To provide a defense against network attacks through packet filtering and the enforcement of trusted communication.

Both goals are met through the use of cryptography-based protection services, security protocols, and dynamic key management. This foundation provides both the strength and flexibility to protect communications between private network computers, domains, sites, remote sites, extranets, and dial-up clients. It can even be used to block receipt or transmission of specific traffic types.

IPSec is based on an end-to-end security model, establishing trust and security from a source IP to a destination IP address. The IP address itself does not necessarily have to be considered an identity, rather the system behind the IP address has an identity that is validated through an authentication process. The only computers that must know about the traffic being secured are the sending and receiving computers. Each computer handles security at its respective end, with the assumption that the medium over which the communication takes place is not secure. Any computers that only route data from source to destination are not required to support IPSec, unless firewall-type packet filtering or network address translation is being done between the two computers. This model allows IPSec to be successfully deployed for the following enterprise scenarios:

- Local area network (LAN): client/server and peer-to-peer

- Wide area network (WAN): router-to-router and gateway-to-gateway

- Remote access: dial-up clients and Internet access from private networks

Typically both sides require IPSec configuration (called an IPSec policy), to set options and security settings that will allow two systems to agree on how to secure traffic between them. The Microsoft® Windows® 2000 , Windows XP, and the Windows Server 2003 family implementations of IPSec is based on industry standards developed by the Internet Engineering Task Force (IETF) IPSec working group. Portions of IPSec-related services were jointly developed by Microsoft and Cisco Systems, Inc.

# Security issues with IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security issues with IP**
Without security, both public and private networks are susceptible to unauthorized monitoring and access. Internal attacks might be a result of minimal or nonexistent intranet security. Risks from outside the private network originate from connections to the Internet and extranets. Password-based user access controls alone do not protect data transmitted across a network.

**Common types of network attacks**
Without security measures and controls in place, your data might be subjected to an attack. Some attacks are passive in that information is only monitored. Other attacks are active and information is altered with intent to corrupt or destroy the data or the network itself. Your networks and data are vulnerable to any of the following types of attacks if you do not have a security plan in place.

**Eavesdropping**
In general, the majority of network communications occur in a plaintext (unencrypted) format, which allows an attacker who has gained access to data paths in a network to monitor and interpret (read) the traffic. When an attacker is eavesdropping on communications, it is referred to as sniffing or snooping. The ability of an eavesdropper to monitor the network is generally the biggest security problem that administrators face in an enterprise. Without strong encryption services that are based on cryptography, data can be read by others as it traverses the network.

**Data modification**
After an attacker has read data, the next logical step is often to modify it. An attacker can modify the data in the packet without the knowledge of the sender or receiver. Even if you do not require confidentiality for all communications, you do not want any of your messages to be modified in transit. For example, if you are exchanging purchase requisitions, you do not want the items, amounts, or billing information to be changed.

**Identity spoofing (IP address spoofing)**
Most networks and operating systems use the IP address to identify a computer as being valid on a network. In some cases, it is possible for an IP address to be falsely used. This is known as identity spoofing. An attacker might use special programs to construct IP packets that appear to originate from valid addresses inside an organization intranet.

After gaining access to the network with a valid IP address, the attacker can modify, reroute, or delete data. The attacker can also conduct other types of attacks, as described in the following sections.

**Password-based attacks**
A commonality among most operating systems and network security plans is password-based access control. Access to both a computer and network resources are determined by a user name and password.

Historically, many versions of operating system components have not always protected identity information as it was passed through the network for validation. This might allow an eavesdropper to determine a valid user name and password and use it to gain access to the network by posing as a valid user.

When an attacker finds and accesses a valid user account, the attacker has the same rights as the actual user. For example, if the user has administrative rights, the attacker can create additional accounts for access at a later time.

After gaining access to a network with a valid account, an attacker can do any of the following:

- Obtain lists of valid user and computer names and network information.
- Modify server and network configurations, including access controls and routing tables.
- Modify, reroute, or delete data.

**Denial-of-service attack**
Unlike a password-based attack, the denial-of-service attack prevents normal use of a computer or network by valid users.

After gaining access to a network, an attacker can do any of the following:

- Distract information systems staff so that they do not immediately detect the intrusion. This gives an attacker the opportunity to make additional attacks.
- Send invalid data to applications or network services, causing applications or services to close or operate abnormally.
- Send a flood of traffic until a computer or an entire network is shut down.
- Block traffic, which results in a loss of access to network resources by authorized users.

**Man-in-the-middle attack**
As the name indicates, a man-in-the-middle attack occurs when someone between two users who are communicating is actively monitoring, capturing, and controlling the communication without the knowledge of the users. For example, an attacker can negotiate encryption keys with both users. Each user then sends encrypted data to the attacker, who can decrypt the data. When computers are communicating at low levels of the network layer, the computers might not be able to determine with which computers they are exchanging data.

**Compromised-key attack**
A key is a secret code or number required to encrypt, decrypt, or validate secured information. Although determining a key is a difficult and resource-intensive process for an attacker, it is possible. After an attacker determines a key, that key is referred to as a compromised key.

An attacker uses the compromised key to gain access to a secured communication without the sender or receiver being aware of the attack. With the compromised key, the attacker can decrypt or modify data. The attacker can also attempt to use the compromised key to compute additional keys, which might allow access to other secured communications.

**Sniffer attack**
A sniffer is an application or device that can read, monitor, and capture network data exchanges and packets. If the packets are not encrypted, a sniffer provides a full view of the data that is inside of the packet. Even encapsulated (tunneled) packets can be opened and read if they are not encrypted.

Using a sniffer, an attacker can do the following:

- Analyze a network and access information, eventually causing the network to stop responding or become corrupted.
- Read private communications.

**Application-layer attack**
An application-layer attack targets application servers by causing a fault in a server's operating system or applications. This results in the attacker gaining the ability to bypass normal access controls. The attacker takes advantage of this situation, gaining control of an application, system, or network, and can do any of the following:

- Read, add, delete, or modify data or an operating system.
- Introduce a virus that uses computers and software applications to copy viruses throughout the network.
- Introduce a sniffer program to analyze the network and gain information that can eventually be used to cause the network to stop responding or become corrupted.
- Abnormally close data applications or operating systems.
- Disable other security controls to enable future attacks.

# Protection against attacks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Protection against attacks**
IPSec protects data so that an attacker finds it extremely difficult or impossible to interpret it. The level of protection provided is determined by the strength of the security levels specified in your IPSec policy structure.

IPSec has a number of features that significantly reduce or prevent the following attacks:

- Sniffers (lack of confidentiality)

  The Encapsulating Security Payload (ESP) protocol in IPSec provides data confidentiality by encrypting the payload of IP packets.

- Data modification

  IPSec uses cryptography-based keys, shared only by the sending and receiving computers, to create a cryptographic checksum for each IP packet. Any modification to the packet data alters the checksum, which indicates to the receiving computer that the packet was modified in transit.

- Identity spoofing, password-based, and application-layer attacks

  IPSec allows the exchange and verification of identities without exposing that information to interpretation by an attacker. Mutual verification (authentication) is used to establish trust between the communicating systems and only trusted systems can communicate with each other. After identities are established, IPSec uses cryptography-based keys, shared only by the sending and receiving computers, to create a cryptographic checksum for each IP packet. The cryptographic checksum ensures that only the computers that have knowledge of the keys could have sent each packet.

- Man-in-the-middle attacks

  IPSec combines mutual authentication with shared, cryptography-based keys.

- Denial-of-service attacks

  IPSec uses IP packet filtering methodology as the basis for determining whether communication is allowed, secured, or blocked, according to the IP address ranges, IP protocols, or even specific TCP and UDP ports.

# Policy-based security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policy-based security**

Although stronger security methods based on cryptography have become necessary to fully protect communication, they can often greatly increase administrative overhead. To reduce overhead, IPSec uses policy-based administration.

IPSec policies, rather than application programming interfaces (APIs), are used to configure IPSec security services. The policies provide variable levels of protection for most traffic types in most existing networks.

You can configure IPSec policies to meet the security requirements of a computer, application, organizational unit, domain, site, or global corporation. You can use the IP Security Policy Management console provided in Microsoft® Windows® XP and the Windows Server 2003 family to define IPSec policies for computers through Active Directory® (for domain members) or on the local computer (for computers that do not belong to domains).

# Simplified deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Simplified deployment**
To achieve secure communications with a low cost of ownership, Microsoft® Windows® XP and the Windows Server 2003 family simplify the deployment of IPSec with the following features:

- Integration with the security framework in Windows 2000 and the Windows Server 2003 family.

  IPSec uses the secure domain in Windows 2000 and the Windows Server 2003 family as a trust model. By default, IPSec policies use the Active Directory® default authentication method (Kerberos V5 authentication) to identify and trust communicating computers. Computers that are members of a Windows 2000 or a Windows Server 2003 domain and are in trusted domains can easily establish IPSec-secured communications.

- Centralized IPSec policy administration through Active Directory

  IPSec policies can be assigned through Group Policy configuration of Active Directory domains and organizational units. This allows the IPSec policy to be assigned at the domain, site, or organizational unit level, eliminating the administrative overhead of configuring each computer separately.

- Transparency of IPSec to users and applications

  Integration at the IP layer (layer 3) provides security for any protocol in the TCP/IP suite of protocols. You do not need separate security for each protocol in the TCP/IP suite of protocols, because applications using TCP/IP pass the data to the IP protocol layer, where it is secured.

- Flexible security configuration

  The security services within each policy can be customized to meet the majority of security requirements for the network and data traffic.

- Automatic key management

  Internet Key Exchange (IKE) services dynamically exchange and manage cryptographic keys between communicating computers.

- Automatic security negotiation

  Internet Key Exchange (IKE) services dynamically negotiate a common set of security settings between communicating computers, eliminating the need for both computers to have identically configured policies.

- Public key infrastructure support

  Using public key certificates for authentication is supported. This allows trust and secure communication for computers that do not belong to a trusted Windows 2000 or Windows Server 2003 domain, non-Microsoft operating systems, computers that have membership in untrusted domains, and instances in which computer access must be restricted to a smaller group than domain authentication allows.

- Preshared key support

  If authentication using the Kerberos V5 protocol or public key certificates is not possible, a preshared authentication key can be configured.

**Important**

- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

**Note**

- The integration with the security framework in Windows 2000 or the Windows Server 2003 family, centralized IPSec policy administration through Active Directory, and the use of the Kerberos V5 protocol as described here do not apply to computers running Windows XP Home Edition. You cannot administer Active Directory-based IPSec policy from computers running Windows XP Home Edition because these computers cannot join Active Directory domains.

# Understanding Internet Protocol Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Internet Protocol security**

This section introduces the key concepts and terminology you should review before you deploy IPSec. This section covers:

- The IPSec process
- Virtual private networking with IPSec
- IPSec Security Services
- IPSec Protocol Types
- IPSec Components
- Description of an IPSec Policy

# The IPSec process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The IPSec process**

This topic provides an overview of IPSec concepts that are central to understanding the IPSec process, including IPSec policy configuration and the Internet Key Exchange (IKE) protocol. In addition, this topic describes how IPSec network traffic processing works, using two intranet computers as an example.

**IPSec Policy Configuration**

In Windows 2000, Windows XP, and the Windows Server 2003 family, IPSec is implemented primarily as an administrative tool that you can use to enforce security policies on IP network traffic. A security policy is a set of packet filters that define network traffic as it is recognized at the IP layer. A filter action defines the security requirements for the network traffic. A filter action can be configured to: **Permit**, **Block**, or **Negotiate security** (negotiate IPSec).

IPSec filters are inserted into the IP layer of the computer TCP/IP networking protocol stack so that they can examine (filter) all inbound or outbound IP packets. Except for a brief delay required to negotiate a security relationship between two computers, IPSec is transparent to end-user applications and operating system services.

A collective set of IPSec security settings is known as an IPSec policy. Windows 2000, Windows XP, and the Windows Server 2003 family provide a graphical user interface and several command-line tools that you can use to configure an IPSec policy, and then assign it to a computer.

To ensure that IPSec communication is successful and that IPSec meets the security requirements of your organization, you must carefully design, configure, coordinate, and manage IPSec policies. In many organizations, one administrator might be responsible for configuring and managing IPSec policies for many, if not all, computers.

**Internet Key Exchange (IKE) security associations**

The IKE protocol is designed to securely establish a trust relationship between each computer, to negotiate security options, and dynamically generate shared, secret cryptographic keying material. The agreement of security settings associated with keying material is called a security association, also known as an SA. These keys will provide authenticity, integrity, and optionally, encryption of IP packets that are sent using the security association. IKE negotiates two types of security associations:

- A main mode security association (the IKE security association that is used to protect the IKE negotiation itself).
- IPSec security associations (the security associations that are used to protect application traffic).

You can configure IPSec policy settings for both types of security associations.

The IPSec service interprets an IPSec policy, expanding it into the components that it needs to control the IKE negotiation. The IPSec policy contains one definition of a packet filter. The packet filter is interpreted in two ways: one uses only the address and identity information to allow IKE to establish a main mode SA (the IKE security association); the other allows IKE to establish the IPSec security associations (also known as quick mode security associations).

**IPSec network traffic processing**

The following illustration shows how IPSec works in terms of the IPSec components for two intranet computers.




For simplicity, this example is of an intranet in which two computers have an active IPSec policy.

1. Alice, using a data application on Computer A, sends an application IP packet to Bob on Computer B.
2. The IPSec driver on Computer A checks its outbound IP filter lists and determines that the packets should be secured.

3. The action is to negotiate security, so the IPSec driver notifies IKE to begin negotiations.

4. The IKE service on Computer A completes a policy lookup, using its own IP address as the source and the IP address of Computer B as the destination. The main mode filter match determines the main mode settings that Computer A proposes to Computer B. Computer A sends the first IKE message in main mode, using UDP source port 500, destination port 500. IKE packets receive special processing by the IPSec driver to bypass filters.

5. Computer B receives an IKE main mode message requesting secure negotiation. It uses the source IP address and the destination IP address of the UDP packet to perform a main mode policy lookup, to determine which security settings to agree to. Computer B has a main mode file that matches, and so replies to begin negotiation of the main mode SA.

6. Computer A and Computer B now negotiate options, exchange identities, verify trust in those identities (authentication), and generate a shared master key. They have now established an IKE main mode SA. Computer A and Computer B must mutually trust each other.

7. Computer A then performs an IKE quick mode policy lookup, using the full filter to which the IPSec driver matched the outbound packet. Computer A selects the quick mode security settings and proposes them, and the quick mode filter, to Computer B.

   Computer B also performs an IKE quick mode policy lookup, using the filter description offered by Computer A. Computer B selects the security settings required by its policy and compares those settings to those offered by computer A. Computer B accepts one set of options and completes the remainder of the IKE quick mode negotiation to create a pair of IPSec security associations.

8. One IPSec SA is inbound and one IPSec SA is outbound. The IPSec SAs are identified by a Security Parameter Index (SPI), which is inserted into the IPSec header of each packet sent.

9. The IPSec driver on Computer A uses the outbound SA to sign and, if required, encrypt the packets. If the network adapter can perform hardware offload of IPSec cryptographic functions, the IPSec driver formats the packets, but does not perform the IPSec cryptographic functions.

10. The IPSec driver passes the packets to the network adapter driver, indicating whether the adapter must perform the IPSec cryptographic functions. The network adapter transmits the packets into the network.

11. The network adapter driver at Computer B receives the encrypted packets from the network. The SPI is used by the receiver of an IPSec packet to find the corresponding IPSec security association, with the cryptographic keys required to verify and decrypt the packets. If the network adapter can decrypt the packets in hardware, it verifies whether it can recognize the SPI. If it cannot decrypt the packets in hardware, or if it cannot recognize the SPI, it passes the packets up to the IPSec driver.

12. The IPSec driver on Computer B uses the inbound SA SPI to retrieve the keys required to validate authentication and integrity and, if required, to decrypt the packets.

13. The IPSec driver converts the packets from IPSec format back to standard IP packet format. It passes the validated and decrypted IP packets to the TCP/IP driver, which passes them to the receiving application on Computer B.

14. The IPSec SAs continue to provide very strong, transparent protection for application data traffic. The IPSec SAs are automatically refreshed by an IKE quick mode negotiation for as long as the application sends and receives data. When the application stops sending and receiving data, the IPSec SAs become idle and are deleted.

15. Typically, the IKE main mode SA is not deleted. By default, the main mode SA has a lifetime of 8 hours. You can configure the main mode SA lifetime to as short as 5 minutes to a maximum of 48 hours. Whenever more traffic is sent, a new quick mode is negotiated automatically to create two new IPSec SAs to protect application traffic. This process is rapid, because the main mode SA already exists. If a main mode SA expires, it is automatically renegotiated as needed.

Any routers or switches in the path between the communicating computers simply forward the encrypted IP packets to their

destination. However, if there is a firewall, security router, or proxy server in the path, IPSec and IKE traffic might not be forwarded. These devices must be configured to allow IPSec and IKE protocol packets through. If the IPSec packets are not encrypted, the firewall or security router can still inspect the TCP or UDP ports or other content in the packets. If the contents of these packets are modified after they are sent, the receiving IPSec computer detects the modification and discards the packets.

However, the Windows Server 2003 family implementation of IPSec provides support for a new Internet specification that allows IPSec packets to be modified by a network address translator (NAT). IPSec Encapsulating Security Payload (ESP) packets can pass through NATs that allow UDP traffic. The IKE protocol automatically detects the presence of a NAT and uses User Datagram Protocol-ESP (UDP-ESP) encapsulation to allow IPSec traffic to pass through the NAT.

# Virtual private networking with IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private networking with IPSec**

Tunneling is the entire process of encapsulation, routing, and decapsulation. Tunneling wraps, or encapsulates, the original packet inside a new packet. This new packet might have new addressing and routing information, which enables it to travel through a network. When tunneling is combined with data confidentiality, the original packet data (as well as the original source and destination) is not revealed to those listening to traffic on the network. After the encapsulated packets reach their destination, the encapsulation is removed, and the original packet header is used to route the packet to its final destination.

The tunnel itself is the logical data path through which the encapsulated packets travel. To the original source and destination peer, the tunnel is usually transparent and appears as just another point-to-point connection in the network path. The peers are unaware of any routers, switches, proxy servers, or other security gateways between the tunnel's beginning point and the tunnel's end point. When tunneling is combined with data confidentiality, it can be used to provide a virtual private network (VPN).

The encapsulated packets travel through the network inside the tunnel. In this example, the network is the Internet. The gateway might be an edge gateway that stands between the outside Internet and the private network. The edge gateway can be a router, firewall, proxy server, or other security gateway. Also, two gateways can be used inside the private network to protect traffic across untrusted parts of the network.

**Tunnel types that use IPSec**

In Windows XP and the Windows Server 2003 family, two types of tunneling are provided that use IPSec:

- **The combination of the Layer Two Tunneling Protocol (L2TP) and IPSec (L2TP/IPSec)**

  L2TP is used to tunnel the data across a shared or public network such as the Internet, and IPSec Encapsulating Security Payload (ESP) is used to encrypt the data. L2TP/IPSec can be used to tunnel IP or Internetwork Packet Exchange (IPX) traffic.

- **IPSec in tunnel mode**

  When IPSec is used in tunnel mode, IPSec itself provides encapsulation for IP traffic only. The primary reason for using IPSec tunnel mode is interoperability with other routers, gateways, or end-systems that do not support L2TP/IPSec or PPTP VPN tunneling. IPSec tunnel mode is supported only in gateway-to-gateway tunneling scenarios and for certain server-to-server or server-to-gateway configurations as an advanced feature. IPSec tunnel mode is not supported for remote access VPN scenarios. L2TP/IPSec or PPTP should be used for remote access VPN connections.

  For information about IPSec tunnel mode and configuring IPSec tunnels, see Tunnel mode.

**How L2TP and IPSec work**

L2TP and IPSec are combined to provide both tunneling and security for IP, IPX, and other protocol packets across any IP network. IPSec can also perform tunneling without L2TP, but it is only recommended for interoperability, when one of the gateways does not support L2TP or PPTP.

L2TP encapsulates original packets first inside a PPP frame (performing compression when possible) and then inside a UDP message using port 1701. Because the UDP message is an IP payload, L2TP uses IPSec transport mode to secure the tunnel. This design supports RFC 3193 ("Securing L2TP Using IPSec"), and it is compatible with Windows XP and Windows 2000. The IPSec Internet Key Exchange (IKE) protocol negotiates security for the L2TP tunnel using certificate-based or preshared key authentication. If IPSec main mode and quick mode security associations are successfully established, L2TP negotiates the tunnel, including compression and user authentication options, and performs PPP-based user authentication.




The original packet, shown here as the **PPP payload**, includes the original and ultimate source and destination addresses used

on the private network. The outer IP header, shown as **IP header**, contains the source and destination addresses of the VPN client and server on the public network. The L2TP header includes L2TP tunnel control information. The PPP header identifies the protocol of the original packet (for example, IP or IPX).

### Note

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

#### Common L2TP/IPSec scenarios

Two common scenarios for L2TP/IPSec are securing communications between remote access clients and the corporate network across the Internet and securing communications between branch offices.

#### Remote access clients with L2TP/IPSec

A common requirement is securing communications between remote access clients and the corporate network across the Internet. Such a client might be a sales consultant who spends most of the time traveling, or an employee working from a home office.




In this figure, the remote gateway is a server that provides edge security for the corporate intranet. The remote client represents a roaming user who requires regular access to network resources and information. An ISP is used as an example to demonstrate the path of communication when the client uses an ISP to access the Internet. L2TP is combined with IPSec to provide a simple, efficient way to build the tunnel and protect the data across the Internet.

For more information about configuring L2TP-based client remote access VPN connections, see Setting Up Remote Access VPNs, Make a virtual private network (VPN) connection, and Virtual Private Network (VPN) Connections.

#### Connecting branch offices with L2TP/IPSec

A large corporation often has multiple sites that require communication--for example, a corporate office in New York and a sales office in Washington. In this case, L2TP is combined with IPSec to provide the VPN connection and protect the data between the sites.

In this figure, the router running the Windows Server 2003 family provides edge security. The routers might have a leased line, dial-up, or other type of Internet connection. The IPSec SA and L2TP tunnel runs between the routers only and provides protected communication through the Internet.

For more information about configuring L2TP-based router-to-router VPN connections, see Deploying Router-to-Router VPNs.

# IPSec Security Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec security services**

IPSec provides a high level of security by using cryptography-based services. Cryptography allows information to be transmitted securely by hashing (to provide integrity and authentication) and encrypting (to provide confidentiality) the information. IPSec provides or uses the following security services:

- Security properties
- Public key certificate-based authentication
- Preshared key authentication
- Data integrity with hash functions
- Data confidentiality with encryption
- Key management and protection

# Security properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security properties**

IPSec provides the following properties for secured communications:

- Anti-replay

  Ensures the uniqueness of each IP packet. Anti-replay is also called replay prevention. Anti-replay ensures that data captured by an attacker cannot be reused or replayed to establish a session or gain information illegally. This protects against attempts to intercept a message and then use the identical message to illegally gain access to resources, possibly months later.

- Integrity

  Protects data from unauthorized modification in transit, ensuring that the data received is exactly the same as the data sent. Hash functions sign each packet with a cryptographic checksum using a shared, secret key. Only the sender and receiver have the key that is used to calculate the checksum. If the packet contents have changed, the checksum verification fails and the packet is discarded.

- Confidentiality (encryption)

  Ensures that data is only disclosed to intended recipients. This is achieved by encrypting the data before transmission, ensuring that the data cannot be read during transmission, even if the packet is monitored or intercepted. Only the communicating computers with the shared, secret key are able to read the data after it has been decrypted. This property is optional and is dependent on IPSec policy settings.

- Authentication

  Verifies that a message could only have been sent from a computer that has knowledge of the shared, secret key. The sender includes (with the message) a message authentication code with a calculation that includes the shared, secret key. The receiver performs the same calculation and, if the receiver's calculation does not match the message authentication code that is included in the message, the message is discarded. The message authentication code is the same as the cryptographic checksum that is used for integrity.

# Public key certificate-based authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Public key certificate-based authentication**

An effectively implemented public key infrastructure, in which security credentials can be presented without compromising them in the process, resolves many security problems. IPSec works with your public key infrastructure to enable certificate-based authentication of computers.

A public key certificate is one type of authentication that reliably provides identity verification. Public key certificates are like digital passports. They are used to verify the identities of non-Microsoft operating system computers, stand-alone computers, clients that are not members of a trusted domain, or computers that are not running the Kerberos V5 authentication protocol (the default authentication method in Windows XP and Windows Server 2003 family) and Routing and Remote Access.

# Preshared key authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Preshared key authentication**

IPSec can use preshared keys for authentication. Preshared means that the parties agree on a shared, secret key that is used for authentication in an IPSec policy.

During security negotiation, information is encrypted before transmission by using a session key. The session key is created by using a Diffie-Hellman calculation and the shared, secret key. Information is decrypted on the receiving end using the same key. One IPSec peer authenticates the other peer's packet by decryption and verification of the hash inside the packet (the hash inside the packet is a hash of the preshared key). If authentication fails, the packet is discarded.

The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext in the registry. In Active Directory, preshared keys are stored in readable hexadecimal format.

Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment. For more information about IPSec authentication methods, see Authentication methods and Public key certificate-based authentication.

# Data integrity with hash functions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data integrity with hash functions**

Hash message authentication codes (HMAC) sign packets to verify that the information received is exactly the same as the information sent. This is called integrity. HMACs provide integrity through a keyed hash, the result of a mathematical calculation on a message using a hash function (algorithm) combined with a shared, secret key. A hash is commonly described as a signature on the packet. However, a hash differs from a digital signature. A hash uses a secret, shared key, and a digital signature uses public key technology and the private key of the sending computer. A digital signature provides nonrepudiation, and a hash does not. Nonrepudiation ensures that a communication can be proven to have originated from a specific person whose identity can be verified. It also ensures that the communication actually occurred.

Hash functions are also called one-way functions because it is easy to determine the hash from the message but mathematically infeasible to determine the message from the hash. Conversely, in two-way functions, the original message can be determined from its converted form. Encryption and decryption schemes are examples of two-way functions.

The hash is a cryptographic checksum or message integrity code (MIC) that each party must compute to verify the message. For example, the sending computer uses a hash function and shared key to compute the checksum for the message, including it with the packet. The receiving computer must perform the same hash function on the received message and shared key and compare it to the original (included in the packet from the sender). If the message has changed in transit, the hash values are different and the packet is rejected.

For integrity, you can choose between two hash functions when setting policy:

- MD5

  Message Digest 5 (MD5) is based on RFC 1321. MD5 completes four passes over the data blocks, using a different numeric constant for each word in the message on each pass. The number of 32-bit constants used during the MD5 computation ultimately produces a 128-bit hash that is used for the integrity check.

- SHA1

  Secure Hash Algorithm 1 (SHA1) was developed by the National Institute of Standards and Technology as described in Federal Information Processing Standard (FIPS) PUB 180-1. The SHA process is closely modeled after MD5. The SHA1 computation results in a 160-bit hash that is used for the integrity check. Because longer hash lengths provide greater security, SHA is stronger than MD5.

# Data confidentiality with encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data confidentiality with encryption**
IPSec in Windows XP and the Windows Server 2003 family uses the US Data Encryption Standard (DES) to provide confidentiality (data encryption). The DES algorithm was published in 1977 by the US National Bureau of Standards.

DES is a block cipher that uses a 56-bit key. A block cipher is an encryption algorithm that operates on a fixed size block of data. DES encrypts data in 64-bit blocks using a 64-bit key. The key appears to be a 64-bit key, but one bit in each of the 8 bytes is used for error checking, resulting in 56 bits of usable key.

Cipher block chaining (CBC) is also used to hide patterns of identical blocks of data within a packet. An initialization vector (an initial random number) is used as the first random block to encrypt and decrypt a block of data. Different random blocks are used in conjunction with the secret key to encrypt each successive block. This ensures that identical sets of unsecured data (plaintext) result in unique, encrypted data blocks.

The Windows XP and the Windows Server 2003 family implementation of IPSec supports the use of:

- DES

  Used when the high security and overhead of 3DES are not necessary.

- 3DES

  Used when high security is required. 3DES processes each block three times, using a unique 56-bit key each time:

  1. Encryption on the block with key 1
  2. Decryption on the block with key 2
  3. Encryption on the block with key 3

     This process is reversed if the computer is decrypting a packet.

**Note**

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

# Key management and protection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key management and protection**
A key is a secret code or number that is required to read, modify, or verify secured data. Keys are used in conjunction with algorithms (a mathematical process) to secure data.

**Key management**
Windows XP and the Windows Server 2003 family automatically handle key generation and implement the following keying properties:

- Variable key lengths

  With IPSec, you can select integrity and encryption algorithms that provide a range of key lengths. Every time the length of a key is increased by one bit, the number of possible keys doubles, making it exponentially more difficult to determine the key. Shorter key lengths are processed more quickly than longer key lengths. Longer key lengths provide a higher level of security, but are more resource-intensive to process.

- Key material generation: The Diffie-Hellman algorithm

  To enable secure communication, two computers must be able to gain the same shared key (session key), without sending the key across a network and compromising the secret.

  The Diffie-Hellman algorithm predates Rivest-Shamir-Adleman (RSA) cryptographic algorithms and offers better performance. It is one of the oldest and most secure algorithms used for key exchange. The two peers publicly exchange keying information, which Windows XP and the Windows Server 2003 family additionally protect with a hash function signature. Neither peer ever exchanges the actual key; however, after their exchange of keying material, each can generate the identical shared key.

  The Diffie-Hellman keying material that is exchanged by the two parties can be based on 768, 1,024, or 2,048 bits of keying material, known as Diffie-Hellman groups 1, 2, and 2048, respectively. The strength of the Diffie-Hellman group is proportional to the strength of the key that is computed from the Diffie-Hellman exchange. Strong Diffie-Hellman groups combined with longer key lengths increase the computational difficulty of determining a secret key.

  IPSec uses the Diffie-Hellman algorithm to provide the keying material for all other encryption keys. Diffie-Hellman does not provide authentication. In the Windows XP and the Windows Server 2003 family implementation of IPSec, identities are authenticated after the Diffie-Hellman exchange takes place, providing protection against man-in-the-middle attacks.

- Dynamic rekeying

  IPSec policy controls how often a new key is generated during communication, by using a method called dynamic rekeying. The communication is sent in blocks, and each block of data is secured with a different key. This prevents an attacker who has obtained part of a communication, and the corresponding session keys, from obtaining the remainder of the communication. This on-demand security negotiation and automatic key management service is provided using IKE, as defined in RFC 2409, "The Internet Key Exchange (IKE)."

  IPSec policy allows you to control how often a new key is generated. If no values are configured, keys are regenerated automatically at default intervals.

**Key protection**
The following features enhance the base prime numbers (keying material) and the strength of the keys for the master and session keys.

- Key lifetimes

  Key lifetimes determine when a new key is generated, rather than how it is generated. A key lifetime allows you to force

an automatic key regeneration after a specific interval of either kilobytes (KB) or seconds, whichever occurs first. This ensures that even if an attacker is able to decipher part of a communication protected by one key, new keys protect the remainder of the communication. Whenever a key lifetime is reached, the security association (SA) is also renegotiated and the key is refreshed or regenerated. For this reason, key lifetimes are also referred to as SA lifetimes. You can specify key lifetimes for the master key and for session keys.

The master key lifetime corresponds to the IKE SA lifetime created by the IKE main mode (phase I) negotiation. It is configured in terms of time and number of quick mode negotiations, and it applies to all security rules in the IPSec policy. The IKE SA and master key typically have a long lifetime (the default value is eight hours). Master keys are much more resource-intensive to generate than session keys, due to the fact that reauthentication and additional Diffie-Hellman exchanges are required.

The session keys correspond to the IPSec SAs that are used to protect program traffic. The session keys and IPSec SAs are quickly derived from the master key by the IKE quick mode (phase II) negotiation. Thus the session keys are used to protect the data, and they have lifetimes based on the amount of data sent and the amount of time elapsed since the key started being used. Typically, session keys have shorter lifetimes than master keys. The default values are 100,000 KB (approximately 100 megabytes) or one hour. When the session keys are refreshed, new IPSec SAs replace the old ones. The IPSec SA session keys must be refreshed before either the data or time lifetime expires, or traffic is discarded. A session key will be deleted if the IPSec SAs become idle (by default, SAs become idle after five minutes). Because the master key lives longer than the session key, it allows new IPSec SAs and session keys to be established quickly. To prevent data loss, the quick mode SA session key is generated shortly before the main mode SA expires. The IKE SA and corresponding master key are not deleted when session keys and IPSec SAs are deleted, unless the specified number of quick mode SA negotiations has elapsed.

For example, if a communication takes one hour (3,600 seconds), and if you specify the minimum session key lifetime of five minutes (300 seconds), more than 12 keys are generated to complete the communication. If a 100 megabytes (MB) file is transferred over a fast corporate LAN using an IPSec security association with a 100 MB and one hour lifetime, at least one, if not two, session rekeys occur.

**Notes**

- To provide maximum protection against cryptographic attacks, the amount of data encrypted using Data Encryption Standard (DES) or Triple DES (3DES) by a single session key should not exceed 100 MB. In a high-security environment, consider using session key lifetimes for MD5 or Secure Hash Algorithm (SHA-1) that are shorter then the default values, to provide stronger protection against future cryptographic attacks.

- In certain scenarios, there might be practical operational limits on key lifetimes. The shorter the key lifetimes, the more CPU cycles are spent refreshing keys. This might limit either the number of active IPSec SAs that a computer can maintain or the throughput performance of IPSec-secured traffic.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

- Session key refresh limit

  Repeated rekeying off of a session key can compromise the Diffie-Hellman shared secret. Because of this, you might want to set a limit to the number of quick mode session keys that can be derived from a main mode negotiation.

  If you have enabled master key perfect forward secrecy (PFS), the session key limit is not used. Setting a session key limit to 1 is identical to enabling master key PFS. If both a master key lifetime and a session limit are specified, whichever limit is reached first causes a new main mode negotiation. By default, IPSec policy does not specify a session limit.

- Diffie-Hellman groups

Diffie-Hellman groups are used to determine the length of the base prime numbers (key material) for the Diffie-Hellman exchange. The cryptographic strength of any key derived from a Diffie-Hellman exchange depends, in part, on the strength of the Diffie-Hellman group on which the prime numbers are based.

Diffie-Hellman Group 2048 (high) is stronger than Group 2 (medium), which is stronger than Group 1 (low). When a stronger group is used, the key that is derived from a Diffie-Hellman exchange is stronger and more difficult for an attacker to break.

IKE negotiates which specific group to use, ensuring that there are not any negotiation failures that result from a mismatched Diffie-Hellman group between the two peers.

If session key PFS is enabled, a new Diffie-Hellman key is negotiated during the first quick mode SA negotiation. This new key removes the dependency of the session key on the Diffie-Hellman exchange that is performed for the master key.

Both the initiator and responder must have session key PFS enabled, or negotiation fails.

The Diffie-Hellman group is the same for both the main mode and quick mode SA negotiations. When session key PFS is enabled, even though the Diffie-Hellman group is set as part of the main mode SA negotiation, it affects any rekeys during session key establishment.

**Important**

- For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.

  For more information about Diffie-Hellman groups and key exchange, see Key exchange methods.

**Note**

- Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.

- Perfect forward secrecy

  Unlike key lifetimes, PFS determines how a new key is generated, rather than when it is generated. Specifically, PFS ensures that the compromise of a single key permits access only to data that is protected by it, not necessarily to the entire communication. To achieve this, PFS ensures that a key used to protect a transmission cannot be used to generate additional keys. In addition, if the key that was used was derived from specific keying material, that material cannot be used to generate other keys.

  Master key PFS requires a reauthentication and is resource intensive. When it is enabled, IKE must reauthenticate identities, increasing overhead for domain controllers when Kerberos V5 is used for authentication. It requires a new main mode negotiation for every quick mode negotiation that occurs.

  Session key PFS can be used without a reauthentication and is less resource intensive. Session key PFS results in a Diffie-Hellman exchange to generate new keying material. It requires only four messages and no authentication.

  Unlike session key PFS, master key PFS does not have to be enabled on both peers. This is because master key PFS is not part of the SA negotiation. If the responder requires PFS, and the quick mode SA of the initiator expires, the responder rejects the message from the initiator and requires a new main mode negotiation. The initiator expires the main mode SA and renegotiates. PFS can be individually set for both the master and session keys.

# IPSec Protocol Types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec protocol types**
IPSec protocols provide data and identity protection for each IP packet by adding their own security protocol header to each packet.

For information about AH and ESP when used in transport and tunnel mode, see the following topics:

- Transport mode
- Tunnel mode

# Transport mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Transport mode**
Transport mode is the default mode for IPSec, and it is used for end-to-end communications (for example, for communications between a client and a server). When transport mode is used, IPSec encrypts only the IP payload. Transport mode provides the protection of an IP payload through an AH or ESP header. Typical IP payloads are TCP segments (containing a TCP header and TCP segment data), a UDP message (containing a UDP header and UDP message data), or an ICMP message (containing an ICMP header and ICMP message data).

**Authentication Header transport mode**
Authentication Header (AH) provides authentication, integrity, and anti-replay protection for the entire packet (both the IP header and the data payload carried in the packet). It does not provide confidentiality, which means that it does not encrypt the data. The data is readable, but protected from modification. AH uses keyed hash algorithms to sign the packet for integrity. For more information, see Data integrity with hash functions.

For example, Alice on Computer A sends data to Bob on Computer B. The IP header, the AH header, and the data are protected with integrity. This means Bob can be certain it was really Alice who sent the data and that the data was unmodified.

Integrity and authentication are provided by the placement of the AH header between the IP header and the IP payload, as shown in the following illustration.

AH is identified in the IP header with an IP protocol ID of 51. AH can be used alone or combined with the Encapsulating Security Payload (ESP) protocol.

The AH header contains the following fields:

- Next Header

  Identifies the IP payload by using the IP protocol ID. For example, a value of 6 represents TCP.

- Length

  Indicates the length of the AH header.

- Security Parameters Index (SPI)

  Used in combination with the destination address and the security protocol (AH or ESP) to identify the correct security association for the communication. The receiver uses this value to determine with which security association the packet is identified.

- Sequence Number

  Provides anti-replay protection for the packet. The sequence number is a 32-bit, incrementally increasing number (starting from 1) that indicates the packet number sent over the security association for the communication. The sequence number cannot repeat for the life of the quick mode security association. The receiver checks this field to verify that a packet for a security association with this number has not already been received. If one has been received, the packet is rejected.

- Authentication Data

  Contains the integrity check value (ICV), also known as the message authentication code, which is used to verify both message authentication and integrity. The receiver calculates the ICV value and checks it against this value (which is calculated by the sender) to verify integrity. The ICV is calculated over the IP header, the AH header, and the IP payload.

**Packet signature with the AH header**
AH signs the entire packet for integrity, with the exception of some fields in the IP header which might change in transit (for example, the Time to Live and Type of Service fields). If another IPSec header is being used in addition to AH, the AH header is inserted before any other IPSec headers. The AH packet signature is shown in the following illustration.




**Encapsulating Security Payload transport mode**
Encapsulating Security Payload (ESP) provides confidentiality (in addition to authentication, integrity, and anti-replay protection) for the IP payload. ESP in transport mode does not sign the entire packet. Only the IP payload (not the IP header) is protected. ESP can be used alone or in combination with AH.

For example, Alice on Computer A sends data to Bob on Computer B. The IP payload is encrypted and signed for integrity. Upon receipt, after the integrity verification process is complete, the data payload in the packet is decrypted. Bob can be certain that it was Alice who sent the data, the data is unmodified, and no one else was able to read it.

ESP is identified in the IP header with the IP protocol ID of 50. As shown in the following illustration, the ESP header is placed before the IP payload, and an ESP trailer and ESP authentication trailer is placed after the IP payload.




The ESP header contains the following fields:

- Security Parameters Index

  Identifies the correct security association for the communication when used in combination with the destination address and the security protocol (AH or ESP). The receiver uses this value to determine the security association with which this packet should be identified.

- Sequence Number

  Provides anti-replay protection for the packet. The sequence number is a 32-bit, incrementally increasing number (starting from 1) that indicates the packet number sent over the quick mode security association for the communication. The sequence number cannot repeat for the life of the quick mode security association. The receiver checks this field to verify that a packet for a security association with this number has not already been received. If one has been received, the packet is rejected.

The ESP trailer contains the following fields:

- Padding

  Padding of 0 to 255 bytes is used to ensure that the encrypted payload with the padding bytes are on byte boundaries required by encryption algorithms.

- Padding Length

  Indicates the length of the Padding field in bytes. The receiver uses this field to remove padding bytes after the encrypted payload with the padding bytes has been decrypted.

- Next Header

  Identifies the type of data in the payload, such as TCP or UDP.

The ESP authentication trailer contains the following field:

- Authentication Data

  Contains the integrity check value (ICV), also known as the message authentication code, which is used to verify both

message authentication and integrity. The receiver calculates the ICV value and checks it against this value (which is calculated by the sender) to verify integrity. The ICV is calculated over the ESP header, the payload data, and the ESP trailer.

**Packet signature and encryption**

As shown in the following illustration, ESP provides protection for IP payloads. The signed portion of the packet indicates where the packet has been signed for integrity and authentication. The encrypted portion of the packet indicates what information is protected with confidentiality.




The IP header is not signed and is not necessarily protected from modification. To provide data integrity and authentication for the IP header, use ESP and AH.

# Tunnel mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tunnel mode**

When IPSec tunnel mode is used, IPSec encrypts the IP header and the payload, whereas transport mode only encrypts the IP payload. Tunnel mode provides the protection of an entire IP packet by treating it as an AH or ESP payload. With tunnel mode, an entire IP packet is encapsulated with an AH or ESP header and an additional IP header. The IP addresses of the outer IP header are the tunnel endpoints, and the IP addresses of the encapsulated IP header are the ultimate source and destination addresses.

IPSec tunnel mode is useful for protecting traffic between different networks, when traffic must pass through an intermediate, untrusted network. Tunnel mode is primarily used for interoperability with gateways, or end-systems that do not support L2TP/IPSec or PPTP connections. You can use tunnel mode in the following configurations:

- Gateway-to-gateway
- Server-to-gateway
- Server-to-server

**AH tunnel mode**

As shown in the following illustration, AH tunnel mode encapsulates an IP packet with an AH and IP header and signs the entire packet for integrity and authentication.




**ESP tunnel mode**

As shown in the following illustration, ESP tunnel mode encapsulates an IP packet with both an ESP and IP header and an ESP authentication trailer.




The signed portion of the packet indicates where the packet has been signed for integrity and authentication. The encrypted portion of the packet indicates what information is protected with confidentiality.

Because a new header for tunneling is added to the packet, everything that comes after the ESP header is signed (except for the ESP authentication trailer) because it is now encapsulated in the tunneled packet. The original header is placed after the ESP header. The entire packet is appended with an ESP trailer before encryption occurs. Everything that follows the ESP header, except for the ESP authentication trailer, is encrypted. This includes the original header which is now considered to be part of the data portion of the packet.

The entire ESP payload is then encapsulated within the new tunnel header, which is not encrypted. The information in the new tunnel header is used only to route the packet from origin to tunnel endpoint.

If the packet is being sent across a public network, it is routed to the IP address of the gateway for the receiving intranet. The gateway decrypts the packet, discards the ESP header, and uses the original IP header to route the packet to the intranet computer.

ESP and AH can be combined when tunneling, providing both confidentiality for the tunneled IP packet and integrity and authentication for the entire packet.

**Using IPSec tunnels**

IPSec tunnels provide security for IP traffic only. The tunnel is configured to protect traffic between either two IP addresses or two IP subnets. If the tunnel is used between two computers instead of two gateways, the IP address outside the AH or ESP payload is the same as the IP address inside the AH or ESP payload. In Windows XP and the Windows Server 2003 family, IPSec does not support protocol-specific or port-specific tunnels. You can configure tunnels by using the IP Security Policy Management and Group Policy consoles to configure and enable two rules:

1. A rule for the outbound traffic for the tunnel.

   The rule for the outbound traffic is configured with a filter list that describes the traffic to be sent through the tunnel and a tunnel endpoint of an IP address configured on the IPSec tunnel peer (the computer or router on the other side of the tunnel).

2. A rule for the inbound traffic for the tunnel.

   The rule for the inbound traffic is configured with a filter list that describes the traffic to be received through the tunnel and a tunnel endpoint of a local IP address (the computer or router on the local side of the tunnel).

Additionally, filter actions, authentication methods, and other settings need to be specified for each rule.

For conceptual information about IPSec policy tunnel settings, see Tunnel endpoint. For information about configuring an IPSec tunnel, see Specify an IPSec tunnel. For information about how tunneling is used for virtual private networking, see Virtual private networking with IPSec.

# IPSec Components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec components**

The following topics describe IPSec components that are installed with Windows XP and the Windows Server 2003 family:

- IPSec Policy Agent service
- Internet Key Exchange
- IPSec driver

# IPSec Policy Agent service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec Policy Agent service**

The purpose of the IPSec Policy Agent is to retrieve policy information and pass it to other IPSec components that require this information to perform security services, as shown in the following illustration.




The IPSec Policy Agent is a service that resides on each computer running a Windows Server 2003 operating system, appearing as **IPSec services** in the list of system services in the Services console. The IPSec Policy Agent:

- Retrieves the appropriate IPSec policy (if one has been assigned) from Active Directory, if the computer is a domain member, or from the local registry, if the computer is not a member of a domain.
- Polls for changes in policy configuration.
- Sends the assigned IPSec policy information to the IPSec driver.

If the computer is a member of a domain, policy retrieval occurs when the system starts, at the interval specified in the IPSec policy, and at the default Winlogon polling interval. You can also manually poll Active Directory for policy using the **gpupdate /target:computer** command.

The following are additional aspects of IPSec policy behavior for a computer that is a member of a domain:

- If IPSec policy information is centrally configured for computers that are domain members, the IPSec policy information is stored in Active Directory and cached in the local registry of the computer to which it applies.
- If the computer is temporarily not connected to the domain and policy is cached, new policy information for that computer replaces old, cached information when the computer reconnects to the domain.
- If the computer is a stand-alone computer or a member of a domain that is not using Active Directory for policy storage, IPSec policy is stored in the local registry.

The IPSec Policy Agent starts automatically at system start time. If there are no IPSec policies in Active Directory or the registry, or if the IPSec Policy Agent cannot connect to Active Directory, the IPSec Policy Agent waits for policy to be assigned or activated.

# Internet Key Exchange

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Key Exchange**
Before secured data can be exchanged, a security agreement between the two computers must be established. In this security agreement, called a security association (SA), both agree on how to exchange and protect information, as shown in the following illustration.




To build this agreement between the two computers, the IETF has established a standard method of security association and key exchange resolution named Internet Key Exchange (IKE) which:

- Centralizes security association management, reducing connection time.
- Generates and manages shared, secret keys that are used to secure the information.

This process not only protects communication between computers, it also protects remote computers that request secure access to a corporate network. In addition, this process works whenever the negotiation for the final destination computer (endpoint) is performed by a security gateway.

**Security association (SA) defined**
A security association (SA) is the combination of a negotiated key, security protocol, and security parameters index (SPI), which together define the security used to protect the communication from sender to receiver. The SPI is a unique, identifying value in the SA that is used to distinguish among multiple security associations that exist at the receiving computer. For example, multiple associations might exist if a computer is securely communicating with multiple computers at the same time. This is a common occurrence when the computer is a file server or a remote access server that serves multiple clients. In these situations, the receiving computer uses the SPI to determine which SA is used to process the incoming packets.

**Phase I or main mode SA**
In order to ensure successful and secure communication, IKE performs a two-phase operation. Confidentiality and authentication are ensured during each phase by the use of encryption and authentication algorithms that are agreed upon by the two computers during security negotiations. With the duties split between two phases, key creation can be rapidly accomplished.

During the first phase, the two computers establish a secure, authenticated channel. This is called the phase I or main mode SA. IKE automatically provides necessary identity protection during this exchange.

**Phase I or main mode negotiation**
The following are the steps that comprise a main mode negotiation.

1. Policy negotiation

   The following four mandatory parameters are negotiated as part of the main mode SA:

   - The encryption algorithm (DES or 3DES)
   - The integrity algorithm (MD5 or SHA1)
   - The Diffie-Hellman group to be used for the base keying material: Group 1 (768 bits of keying material) Group 2 (1,024 bits), or Group 2048 (2,048 bits)
   - The authentication method (Kerberos V5, certificate, or preshared key authentication)

     If certificates or preshared keys are used for authentication, the computer identity is protected. If Kerberos V5 authentication is used, the computer identity is unencrypted until encryption of the entire identity payload takes

place during authentication.

**Important**

For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.

For more information about Diffie-Hellman groups, see Key exchange methods. For more information about preshared key authentication, see Preshared key authentication.

2. Diffie-Hellman exchange (of public values)

   At no time are actual keys exchanged. Only the base information required by the Diffie-Hellman key determination algorithm to generate the shared, secret key is exchanged. After this exchange, the IKE service on each computer generates the master key that is used to protect authentication.

3. Authentication

   To prevent a successful man-in-the-middle attack, the computers attempt to authenticate the Diffie-Hellman key exchange. Without successful authentication, communication will not proceed. The master key is used, in conjunction with the negotiation algorithms and methods, to authenticate identities. The entire identity payload is hashed and encrypted using the keys generated from the Diffie-Hellman exchange in the second step. The payload includes the identity type (for authentication), port, and protocol. IPSec uses the following identity types for authentication: For certificate authentication, the certificate distinguished name and general name; for Kerberos V5 and preshared key authentication, IPv4 addresses, the fully qualified domain name (FQDN) of the computer, and FQDN of the user. The identity payload, regardless of which authentication method is used, is protected from both modification and interpretation.

The sender presents an offer for a potential security association to the receiver. The responder cannot modify the offer. Should the offer be modified, the initiator rejects the responder's message. The responder sends either a reply accepting the offer or a reply with alternatives.

Messages sent during this phase have an automatic retry cycle that is repeated five times. If a response is received before the retry cycle ends, standard SA negotiation begins. If allowed by IPSec policy, unsecured communications will begin after a brief interval. If unsecured communications begin, after five minutes of idle time (during which no messages are sent), secured communication negotiation is attempted the next time messages are sent. If messages are sent continuously, the communication remains unsecured during the lifetime set for the main mode policy. After the policy time has elapsed, a new secured communication negotiation attempt is made.

There is no preset limit to the number of exchanges that can take place. The number of SAs established is only limited by system resources. When estimating the number of SAs that can be established without significantly degrading computer performance, consider the CPU processing strength and RAM of the computer, the lifetime of the SA, and how much traffic is being sent over the SAs.

For more information about unsecured communications, see Filter action.

**Phase II or quick mode SA**
In this phase, SAs are negotiated on behalf of the IPSec driver.

**Phase II or quick mode negotiation**
The following are the steps that comprise a quick mode negotiation.

1. Policy negotiation occurs.

   The IPSec computers exchange the following requirements for securing the data transfer:

   - The IPSec protocol (AH or ESP)
   - The hash algorithm for integrity and authentication (MD5 or SHA1)

- The algorithm for encryption, if requested (3DES or DES)

  A common agreement is reached, and two SAs are established. One SA is for inbound communication and the other is for outbound communication.

2. Session key material is refreshed or exchanged.

   IKE refreshes the keying material and new shared keys are generated for packet integrity, authentication, and encryption (if negotiated). If rekeying is required, either a second Diffie-Hellman exchange (as described in main mode negotiation) occurs, or a refresh of the original Diffie-Hellman key is used.

3. The SAs and keys, along with the SPI, are passed to the IPSec driver.

The second negotiation of security settings and keying material (for the purpose of securing data) is protected by the main mode SA. As the first phase provided identity protection, the second phase provides protection by refreshing the keying material prior to sending data. IKE can accommodate a key exchange payload for an additional Diffie-Hellman exchange if a rekey is necessary--that is, master key perfect forward secrecy (PFS) is enabled. Otherwise, IKE refreshes the keying material from the Diffie-Hellman exchange completed in main mode.

Quick mode results in a pair of security associations, each with its own SPI and key. One SA is used for inbound communication, and the other for outbound communication.

**Notes**

- Although there are two separate quick mode SAs established, IP Security Monitor only displays a single quick mode SA.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

The retry algorithm for a message is similar to the process described in main mode negotiation. However, if this process times out for any reason during the second or higher negotiation off of the same main mode SA, a renegotiation of the main mode SA is attempted. If a message for this phase is received without an established main mode SA, it is rejected.

Using a single main mode SA for multiple quick mode SA negotiations increases the speed of the process. As long as the main mode SA does not expire, renegotiation and reauthentication are not necessary. The number of quick mode SA negotiations that can be performed is determined by IPSec policy settings.

**Note**

- Excessive rekeying off of the same main mode SA might make the shared, secret key vulnerable to a known plaintext attack. A known plaintext attack is a sniffer attack in which the attacker attempts to determine the encryption key from encrypted data based on known plaintext.

**SA lifetimes**

The main mode SA is cached to allow multiple quick mode SA negotiations (unless master key PFS is enabled). When a key lifetime is reached for the master or session key, the SA is renegotiated. In addition, the key is refreshed or regenerated.

When the default time-out period elapses for the main mode SA, or the master or session key lifetime is reached, a delete message is sent to the responder. The IKE delete message tells the responder to expire the main mode SA. This prevents additional new quick mode SAs from being created from the expired main mode SA. IKE does not expire the quick mode SA, because only the IPSec driver contains the number of seconds or bytes that have passed to reach the key lifetime.

Use caution when setting very different key lifetimes for master and session keys. For example, setting a master key lifetime of

eight hours and a session key lifetime of two hours might leave a quick mode SA in place for almost two hours after the main mode SA has expired. This occurs when the quick mode SA is generated shortly before main mode SA expiration.

It is generally recommended that all of the IKE settings (for example, master key PFS and key lifetime) and security methods remain at their defaults to avoid unnecessary administrative overhead. This provides a standard (medium) level of security. If your security plan calls for a high level of security, you should consider modifying the default security methods.

# IPSec driver

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec driver**

The IPSec driver receives the active IP filter list from the IPSec Policy Agent, as shown in the following illustration, and then attempts to match every inbound and outbound packet against the filters in the list.




When a packet matches a filter, it applies the filter action. When a packet does not match any filters, the packet is passed back without modification to the TCP/IP driver to be received or transmitted.

If the filter action permits transmission, the packet is received or sent with no modifications. If the action blocks transmission, the packet is discarded. If the action requires the negotiation of security, main mode and quick mode SAs are negotiated.

The negotiated quick mode security association (SA) and keys are used with both outbound and inbound processing. The IPSec driver stores all current quick mode SAs in a database. The IPSec driver uses the Security Parameters Index (SPI) field to match the correct SA with the correct packet.

When an outbound IP packet matches the IP filter list with an action to negotiate security, the IPSec driver queues the packet and then notifies Internet Key Exchange (IKE), which begins security negotiations with the destination IP address of that packet. If several outbound packets are going to the same destination and match the same filter before IKE has finished the negotiation, then only the last packet sent is saved.

After a successful negotiation is completed, the IPSec driver on the sending computer:

1. Receives the SA containing the session key from IKE.
2. Locates the outbound SA in its database, and inserts the SPI from the SA into the AH or ESP header.
3. Signs the packets (and encrypts them if confidentiality is required).
4. Sends the packets to the IP layer to be forwarded to the destination computer.

If the negotiation failed, the IPSec driver discards the packet.

When an IPSec-secured inbound packet matches a filter in the IP filter list, the IPSec driver:

1. Receives the session key, SA, and SPI from IKE.
2. Locates the inbound SA in its database by destination address and SPI.
3. Checks the signature and, if required, decrypts the packets.
4. Matches the IP packet searches for a matching filter in the filter list to ensure that no traffic, other than what was agreed upon during the negotiation, is received.
5. Sends packets to the TCP/IP driver to pass to the receiving application.

When an unsecured IP packet is received, the IPSec driver searches for a matching filter in the filter list. If a match occurs and the filter action for that filter either requires IP security or blocks the packet, then the packet is discarded.

**Note**

- The IPSec driver matches all inbound unsecured packets with the list of filters that specify IPSec tunnels first, and then matches the packet with all filters that specify end-to-end (transport) filters. The IPSec driver does not filter certain types of IP packets. For more information, see Add, edit, or remove IPSec filters.

# Description of an IPSec Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Description of an IPSec policy**
An IPSec policy consists of:

- General IPSec Policy Settings

  Settings that apply regardless of which rules are configured. These settings determine the name of the policy, its description for administrative purposes, key exchange settings, and key exchange methods.

- IPSec Policy Rules

  One or more IPSec rules that determine which types of traffic IPSec must examine, how traffic is treated, how to authenticate an IPSec peer, and other settings.

After the policies are created, they can be assigned and applied at the domain, site, organizational unit, and local level. For more information, see Creating, modifying, and assigning IPSec policies.

# General IPSec Policy Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**General IPSec policy settings**

The general settings for an IPSec policy consist of the following types:

- Name

  The name for the policy.

- Description

  Optional text that describes the purpose of the IPSec policy. It is recommended that you complete and update this description to provide a summary of the settings and rules for the policy.

- Policy change poll interval

  The number of minutes between consecutive polls for changes in Active Directory-based IPSec policies. This polling does not detect a change in domain or organizational unit membership, or the assigning or unassigning of a new policy. These events are detected when the Winlogon service polls for changes in Group Policy, which occurs by default every 90 minutes.

- Key exchange settings

  The way in which new keys are derived and how often they are renewed.

- Key exchange methods

  The ways in which identities are protected during the key exchange.

The default key exchange settings and methods are configured to work for most IPSec deployments. Unless there are special security requirements, default settings should not have to be changed.

# Key exchange settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key exchange settings**
The strength of the keys which protect the key exchange phase is enhanced by the following features:

**Key lifetimes**
Lifetime settings determine when a new key is generated. Any time a key lifetime is reached, the associated SA is also renegotiated. The process of generating new keys at intervals is called dynamic rekeying or key regeneration. Lifetimes allow you to force the generation of a new key after a specific interval. For example, if the communication takes 100 minutes and you specify the key lifetime as 10 minutes, 10 keys will be generated (one every 10 minutes) during the exchange. Using multiple keys ensures that if an attacker manages to gain the key to one part of a communication, the entire communication is not compromised. Automatic key regeneration is provided by default. You can override the defaults and either specify a master key lifetime in minutes or by the number of session keys, or enable master key perfect forward secrecy (PFS).

Caution should be taken when setting very different key lifetimes, since they also determine the lifetime of the SA. For example, setting a master key lifetime of 8 hours (480 minutes) and a session key lifetime (set within a filter action) of 2 hours might result in a quick mode SA that is in place for almost 2 hours after the main mode SA has expired. This has the potential to occur if the new quick mode SA is generated just before the main mode SA expires.

- **Session key limit**

  Repeated rekeying off of the same master key might eventually compromise the key. For example, if Alice on Computer A sends a message to Bob on Computer B, and then a few minutes later sends another message to Bob, the same master key material can be reused since an SA was recently established with that computer. If you want to limit the number of times this reuse occurs, you can specify a session key limit.

  If you decide to enable master key perfect forward secrecy (PFS), the session key limit will be set to 1. Master key PFS forces key regeneration each time. If you specify both a master key lifetime in minutes, and a session key limit, whichever interval is reached first will start a new key.

- **Master key perfect forward secrecy (PFS)**

  Master key PFS determines how a new session key is generated. Enabling master key PFS ensures that the master key keying material cannot be used to generate more than one session key.

  Master key PFS should be used with caution because it requires reauthentication and might impact performance. It is not required to be enabled on both peers.

For information about how to configure key exchange settings and methods, see Configure key exchange settings and Create key exchange security methods.

# Key exchange methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key exchange methods**
In addition to key properties, you can set security methods for main mode IKE negotiation. For example, you can specify which algorithms are used for integrity and confidentiality. The same algorithms that are available for quick mode security methods are available for main mode IKE negotiation: MD5 and SHA1 for integrity; DES and 3DES for confidentiality.

**Diffie-Hellman groups**
In addition to setting integrity and confidentiality algorithms, you can specify which Diffie-Hellman groups to use. Diffie-Hellman groups are used to determine the length of the base prime numbers used during the key exchange process. The cryptographic strength of any key derived depends, in part, on the strength of the Diffie-Hellman group upon which the prime numbers are based.

Group 2048 (high) is stronger (more secure) than Group 2 (medium), which is stronger than Group 1 (low). Group 1 provides 768 bits of keying strength, Group 2 provides 1024 bits, and Group 2048 provides 2048 bits. If mismatched groups are specified on each peer, negotiation fails. The group cannot be switched during the negotiation.

The Diffie-Hellman group is configured as part of the main mode key exchange settings. New session keys generated during quick mode are derived from the Diffie-Hellman main mode master key material, unless master key or session key perfect forward secrecy (PFS) is enabled. If either master key or session key PFS is enabled, a new Diffie-Hellman exchange is performed to obtain new master key keying material for each new session key that is required. The difference between master key PFS and session key PFS is that master key PFS requires a reauthentication of the main mode SA in addition to the Diffie-Hellman exchange.

**Important**

- For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.

**Notes**

- Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.
- For a standard level of security, it is generally recommended that all of the IKE settings (for example, master key PFS and key lifetime) and security methods remain at their defaults. For more information about IKE settings, see Internet Key Exchange.
- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

For information about how to configure key exchange settings and methods, see Configure key exchange settings and Create key exchange security methods.

# IPSec Policy Rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec policy rules**

An IPSec policy consists of one or more rules that determine IPSec behavior. IPSec rules are configured on the **Rules** tab in the properties of an IPSec policy. Each IPSec rule contains the following configuration items:

- Filter list

  A single filter list is selected that contains one or more predefined packet filters that describe the types of traffic to which the configured filter action for this rule is applied. The filter list is configured on the **IP Filter List** tab in the properties of an IPSec rule within an IPSec policy.

- Filter action

  A single filter action is selected that includes the type of action required (Permit, Block, or Negotiate Security) for packets that match the filter list. For the Negotiate Security filter action, the negotiation data contains one or more security methods that are used (in order of preference) during IKE negotiations and other IPSec settings. Each security method determines the security protocol (such as AH or ESP), the specific cryptographic and hashing algorithms, and session key regeneration settings used. The filter action is configured on the **Filter Action** tab in the properties of an IPSec rule within an IPSec policy.

- Authentication methods

  One or more authentication methods are configured (in order of preference) and used for authentication of IPSec peers during main mode negotiations. The available authentication methods are the Kerberos V5 protocol, use of a certificate issued from a specified certification authority, or a preshared key. The negotiation data is configured on the **Authentication Methods** tab in the properties of an IPSec rule within an IPSec policy.

  **Important**

  - The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

- Tunnel endpoint

  Specifies whether the traffic is tunneled and, if it is, the IP address of the tunnel endpoint. For outbound traffic, the tunnel endpoint is the IP address of the IPSec tunnel peer. For inbound traffic, the tunnel endpoint is a local IP address. The tunnel endpoint is configured on the **Tunnel Setting** tab in the properties of an IPSec rule within an IPSec policy. For more information, see Tunnel mode.

- Connection type

  Specifies whether the rule applies to local area network (LAN) connections, dial-up connections, or both. The connection type is configured on the **Connection Type** tab in the properties of an IPSec rule within an IPSec policy.

The rules for a policy are displayed in IP Security Policies in reverse alphabetical order based on the name of the filter list selected for each rule. There is no method for specifying an order in which to apply the rules in a policy. The IPSec driver automatically orders the rules based on the most specific to the least specific filter list. For example, the IPSec driver would apply a rule containing a filter list that specified individual IP addresses and TCP ports before a rule containing a filter list that specified all addresses on a subnet.

**Default response rule**

The default response rule, which can be used for all policies, has the IP filter list of **<Dynamic>** and the filter action of **Default Response** when the list of rules is viewed with the IP Security Policy Management console. The default response rule cannot be deleted, but it can be deactivated. It is activated by default for all policies.

The default response rule is used to ensure that the computer responds to requests for secure communication. If an active policy does not have a rule defined for a computer that is requesting secure communication, then the default response rule is applied and security is negotiated. For example, when Computer A communicates securely with Computer B, and Computer B does not have an inbound filter defined for Computer A, the default response rule is used.

Security methods and authentication methods can be configured for the default response rule. The filter list of **<Dynamic>** indicates that the filter list is not configured, but that filters are created automatically based on the receipt of IKE negotiation packets. The filter action of **Default Response** indicates that the action of the filter (Permit, Block, or Negotiate security) cannot be configured. Negotiate Security will be used. However, you can configure:

- The security methods and their preference order on the **Security Methods** tab.
- The authentication methods and their preference order on the **Authentication Methods** tab.

# Filter list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Filter list**

When traffic matches a source, destination, and type of IP traffic in a filter, security negotiations are initiated. This type of IP packet filtering enables a network administrator to precisely define what IP traffic is secured. Each IP filter list contains one or multiple filters, which define IP addresses and traffic types. One IP filter list can be used for multiple communication scenarios.

IPSec requires both an inbound and outbound filter between the computers specified in the filter list. Inbound filters apply to incoming traffic. Outbound filters apply to traffic leaving a computer towards a destination. For example, if Computer A wants to securely exchange data with Computer B:

- The active IPSec policy on Computer A must have a filter for any outbound packets to Computer B. Source=A and Destination=B.
- The active IPSec policy on Computer A must have a filter for any inbound packets from Computer B. Source=B and Destination=A.

Each peer must also have the reverse filter:

- The active IPSec policy on Computer B must have a filter for any inbound packets from Computer A. Source=A and Destination=B.
- The active IPSec policy on Computer B must have a filter for any outbound packets to Computer A. Source=B and Destination=A.

**Filter settings**

Each filter defines a particular subset of inbound or outbound network traffic that should be secured. You must have a filter to cover any traffic to which the associated rule applies. A filter contains the following settings:

- **The source and destination address of the IP packet.** You can specify any IP address assigned to the IPSec peer, a single IP address, IP addresses by DNS name, or groups of addresses to specify IP subnets.
- **The protocol over which the packet is being transferred.** By default, all protocols in the TCP/IP protocol suite are selected. However, you can specify an individual protocol for this filter to meet special requirements, including custom protocols.
- **The source and destination port for TCP and UDP.** By default, all TCP and UDP ports are selected, but you can select a specific TCP or UDP port.

For information about how to configure IP filter lists, see Add, edit, or remove IP filter lists.

Special filters might be required for certain types of traffic, such as SNMP traffic or security gateways. For more information, see Special IPSec considerations.

# Filter action

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Filter action**

A filter action defines the security requirements for the data transmission. A filter action can be configured to:

- Permit traffic (**Permit**)

  IPSec passes this traffic without modification or the requirement for security. This is appropriate for traffic from specific computers that are not IPSec-capable. Be sure to limit the IP filter list to a minimal scope when using this type of filter action, so you do not let traffic through which should be secured.

- Block traffic (**Block**)

  IPSec silently discards this traffic. Be sure to use an IP filter list that appropriately defines the scope of traffic when using a blocking filter action so that you do not block valid computers from communicating.

- Negotiate IPSec (**Negotiate security**)

  IPSec requires the negotiation of security associations and the sending or receiving of IPSec-secured traffic. After you choose to negotiate IPSec, you can also configure the following:

  - Security methods and their order

    For more information, see IPSec security methods.

  - Allowance of initial incoming unsecured traffic (**Accept unsecured communication, but always respond using IPSec**)

    IPSec allows an incoming packet that matches the configured filter list to be unsecured (not protected by IPSec). However, the outgoing response to the incoming packet must be protected. This setting is useful when you are using the default response rule for clients. When a group of servers are configured with a rule that secures communications with any IP address and accepts unsecured communication, responding with only secured communications, enabling the default response rule on client computers ensures that the clients will respond to the server request to negotiate security. This option should be disabled for secure computers connected to the Internet to prevent denial of service attacks.

  - Enabling of communication with non-IPSec-enabled computers (**Allow unsecured communication with non-IPSec-aware computer**)

    IPSec falls back to unsecured communication, if necessary. Again, you might need to limit the IP filter list to a minimal scope. Otherwise, if negotiation fails for any reason, any communications affected by the rule in which this filter action resides could result in data being sent without security. If you are concerned about unsecured communication, you might consider disabling this setting. However, communication with computers that cannot initiate IPSec, such as legacy systems, might be blocked. This option should be disabled for secure computers connected to the Internet.

  - Generation of session keys from new keying material (**Use session key perfect forward secrecy (PFS)**)

    This determines whether existing master key keying material can be used to derive a new session key. Enabling session key perfect forward secrecy (PFS) ensures that master key keying material cannot be used to derive more than one session key. When session key PFS is enabled, a new Diffie-Hellman key exchange is performed to generate new master key keying material before the new session key is created. Session key PFS does not require main mode reauthentication and uses less resources than master key PFS.

For information about how to configure filter actions, see Add, edit, or remove filter actions.

**IPSec security methods**

Each security method defines the security requirements of any communications to which the associated rule applies. Creating multiple security methods increases the chance that a common method can be found between two computers. The IKE component reads the list of security methods in descending order and sends a list of allowed security methods to the other peer. The first method in common is selected. Typically, the most secure methods are placed at the top of the list and the least secure methods are placed at the bottom of the list.

**Predefined security methods**

The following security methods are predefined:

- Encryption and Integrity

  Uses the ESP protocol to provide data confidentiality (encryption) with the triple Data Encryption Standard (3DES) algorithm, data integrity and authentication with the Secure Hash Algorithm 1 (SHA1) integrity algorithm, and default key lifetimes (100MB, 1 hour). If you require both data and addressing (IP header) protection, you can create a custom security method. If you do not require encryption, use **Integrity only**.

- Integrity only

  Uses the ESP protocol to provide data integrity and authentication with the SHA1 integrity algorithm and default key lifetimes (100MB, 1 hour). In this configuration, ESP does not provide data confidentiality (encryption). This is appropriate when your security plan calls for standard levels of security.

**Custom security methods**

If the predefined **Encryption and Integrity** or **Integrity only** settings do not meet your security requirements, you can specify custom security methods. For example, you can use custom methods when encryption and address integrity, stronger algorithms, or key lifetimes need to be specified. When configuring a custom security method, you can configure the following:

- Security protocols

  Both AH (**Data and address integrity without encryption**) and ESP (**Data integrity and encryption**) can be enabled in a custom security method when you require IP header integrity and data encryption. If you chose to enable both, you do not need to specify an integrity algorithm for ESP. The algorithm that you select for AH provides integrity.

- Integrity algorithm

  - Message Digest 5 (MD5), which uses a 128-bit key.
  - Secure Hash Algorithm 1 (SHA1), which uses a 160-bit key.

- Encryption algorithm

  - 3DES is the most secure of the DES combinations, and somewhat slower in performance. 3DES processes each block three times, using three unique 56-bit keys.
  - DES is used when the higher security and overhead of 3DES is not necessary. DES uses a single 56-bit key.

- Session key settings

  Session key settings determine when a new key is generated, rather than how it is generated. You can specify a lifetime in kilobytes, seconds, or both. For example, if the communication takes 10,000 seconds and you specify the key lifetime as 1000 seconds, 10 keys will be generated to complete the transfer. This ensures that, even if an attacker manages to determine one session key and decipher part of a communication, deciphering the entire communication is not possible. By default, new session keys are generated for every 100 MB of data or every hour. Note that any time a key lifetime is reached, the quick mode SA is also renegotiated in addition to the key refresh or regeneration.

**Note**

- When IPSec provides data confidentiality between a Windows XP or Windows Server 2003 family-based computer (or a Windows 2000-based computer that has the High Encryption Pack or Service Pack 2 or higher installed) and another computer that does not use 3DES, your list of security methods must include DES. Otherwise, the computer that you are trying to communicate with might not be able to reach a common security agreement with your computer. If confidentiality of the data is not required, you can select ESP format with an integrity algorithm only and set the encryption to **<None>**. Alternately, you can select an AH format with an integrity algorithm.

For information about how to configure security methods, see Add, edit, or remove IPSec security methods.

# Authentication methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication methods**
Each rule defines a list of authentication methods. Each authentication method defines the requirements for how identities are verified in communications to which the associated rule applies. The two peers must have at least one common authentication method or communication will fail. Creating multiple authentication methods increases the chance that a common method between two computers can be found.

Only one authentication method can be used between a pair of computers, regardless of how many are configured. If you have multiple rules that apply to the same pair of computers, you must configure the authentication methods list in those rules to enable the pair to use the same method. For example, if a rule between a pair of computers specifies only Kerberos for authentication and filters only TCP data and, in another rule, specifies only certificates for authentication and filters only UDP data, authentication will fail.

**Overview of authentication methods**
For authentication, IPSec allows you to use the Kerberos V5 protocol, certificate-based authentication, or preshared key authentication.

- The Kerberos V5 security protocol is the default authentication technology. This method can be used for any clients running the Kerberos V5 protocol (whether or not the clients are running Windows 2000, Windows XP Professional, or a Windows Server 2003 operating system) that are members of the same or trusted domains.

- A public key certificate should be used in situations that include Internet access, remote access to corporate resources, external business partner communications, or computers that do not run the Kerberos V5 security protocol. This requires that at least one trusted certification authority (CA) and associated certificate have been configured. Computers running Windows 2000, Windows XP, or a Windows Server 2003 operating system support X.509 Version 3 certificates, including computer certificates generated by commercial CAs.

- A preshared secret key can be specified. It is simple to use and does not require the client to run the Kerberos V5 protocol or have a public key certificate. Both parties must manually configure IPSec to use this preshared key.

  **Important**

  - The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

## Notes

- The ability to authenticate using Kerberos V5 is not supported on computers running Windows XP Home Edition or computers running any other version of Windows 2000, Windows XP, or a Windows Server 2003 operating system that are not a member of an Active Directory domain.

- For information about how to configure authentication methods, see Define IPSec authentication methods.

**IPSec certificate to account mapping**
When you use either Kerberos V5 or certificate authentication, you can set restrictions on which computers are allowed to connect. When you enable IPSec certificate to account mapping, the IKE protocol associates (maps) a computer certificate to a computer account in an Active Directory domain or forest, and then retrieves an access token, which includes the list of user rights assigned to the computer. You can restrict access by configuring Group Policy security settings and assigning either the **Access this computer from the network** user right or the **Deny access to this computer from the network** user right to individual or multiple computers as needed.

When a Windows 2000 or a Windows Server 2003 family certification authority is used for computer certificate autoenrollment, the certificates are automatically stored in Active Directory as a property of the computer account. Alternatively, you can use a non-Microsoft certification authority in which to enroll computers. If you do not use computer certificate autoenrollment, you must manage certificate to account mapping manually.

**Notes**

- If you use the Kerberos V5 protocol for authentication, the IKE protocol verifies access controls, but only if IKE is a responder. Therefore, for access controls to be applied, you must configure the IPSec policy on the client computer to initiate IKE to the server. If you use certificate authentication, access controls are applied if either the server or the client computer initiates IKE to the server.

- For information about configuring security settings in Group Policy, see Security settings overview.

# Tunnel endpoint

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tunnel endpoint**

As discussed in Virtual private networking with IPSec, IPSec can perform layer 3 tunneling for scenarios in which Layer Two Tunneling Protocol (L2TP) cannot be used. If you are using L2TP for remote communications, no tunnel configuration is needed because the client and server VPN components of Windows XP and the Windows Server 2003 family automatically create the appropriate rules to secure L2TP traffic.

To create a layer 3 tunnel using IPSec, use the IP Security Policy Management or Group Policy consoles to configure and enable the following two rules for the appropriate policy:

1. A rule for the outbound traffic for the tunnel.

   The rule for the outbound traffic is configured with both a filter list that describes the traffic to be sent across the tunnel and a tunnel endpoint of an IP address configured on the IPSec tunnel peer (the computer or router on the other side of the tunnel).

2. A rule for the inbound traffic for the tunnel.

   The rule for the inbound traffic is configured with both a filter list that describes the traffic to be received across the tunnel and a tunnel endpoint of a local IP address (the computer or router on this side of the tunnel).

For each rule, you must also specify filter actions, authentication methods, and other settings.

For information about how to configure IPSec tunnel endpoints, see Specify an IPSec tunnel. For more information about IPSec tunnel mode, see Tunnel mode.

For advanced information about configuring IPSec tunnels, see article Q252735, "How to Configure IPSec Tunneling in Windows 2000," in the Microsoft Knowledge Base.

# Connection type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection type**
For each IPSec rule, you can define which connection types on your computer are affected. The connection types apply to all connections in Network Connections on the computer for which you are configuring IPSec policy.

Each rule has one connection type setting:

- **All Network Connections**. The rule applies to communications sent through any of the network connections that you have configured on the computer.
- **Local Area Network (LAN)**. The rule only applies to communications sent through LAN connections that you have configured on the computer.
- **Remote Access**. The rule only applies to communications sent through any virtual private network (VPN) or dial-up connections that you have configured on the computer.

For information about how to configure connection types, see Specify IPSec connection types.

For information about configuring network connections on a computer, see Configuring a network or dial-up connection.

# Using Internet Protocol Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Internet Protocol security**

This section covers:

- Managing IPSec from the command line
- Designing and Planning IPSec Policies
- Viewing IPSec Statistics and Details About Active IPSec Policies

# Managing IPSec from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing IPSec from the command line**
The Netsh commands for Internet Protocol security (IPSec) provide a fully equivalent alternative to the console-based management and diagnostic capabilities provided by the IP Security Policy Management and IP Security Monitor consoles. You can use Netsh commands for IPSec to script IPSec policy creation, display details about IPSec policies, and change the IPSec configuration for troubleshooting. In addition, administering IPSec from the command line is useful when you want to extend the security and manageability of IPSec. For example, you can use Netsh commands for IPSec to enable IPSec driver event logging, set default traffic exemptions, and configure computer startup security.

For a comprehensive **netsh ipsec** command reference (including syntax and parameters), see Netsh commands for Internet Protocol security.

### Notes

- For information about how to use Netsh commands for IPSec to change the IPSec configuration for troubleshooting, see the section "Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family" in IPSec troubleshooting tools.

- You can use the Netsh commands for IPSec to configure IPSec policies only on computers running members of the Windows Server 2003 family. To use the command line to configure IPSec policies on computers running Windows XP, use Ipseccmd.exe. To use the command line to configure IPSec policies on computers running Windows 2000, use Ipsecpol.exe, which is provided with the Windows 2000 Server Resource Kit. For more information, see IPSec troubleshooting tools.

# Designing and Planning IPSec Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Designing and planning IPSec policies**
This section provides information about planning and deploying IPSec in your organization. You should also review IPSec Best practices for general information about IPSec deployment.

For information on specific topics, see:

- Establishing an IPSec security plan
- Default IPSec policies
- Creating, modifying, and assigning IPSec policies
- Special IPSec considerations
- IPSec troubleshooting tools

# Establishing an IPSec security plan

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing an IPSec security plan**

Whether for a large domain or a small workgroup, implementing IPSec means finding a balance between making information easily available to the largest number of users, and protecting sensitive information from unauthorized access.

Finding the proper balance requires:

- Assessing the risk and determining the appropriate level of security for your organization.
- Identifying valuable information.
- Defining security policies that use your risk management criteria and protect the identified information.
- Determining how the policies can best be implemented within the existing organization.
- Ensuring that management and technology requirements are in place.
- Providing all users with both secure and efficient access to the appropriate resources, according to their needs.

Security considerations are also influenced by the way the computer will be used. For example, the required security might differ, depending on whether the computer is a domain controller, Web server, remote access server, file server, database server, or intranet or remote client. The security framework in Windows 2000, Windows XP, and Windows Server 2003 family is designed to fulfill the most stringent security requirements. However, software alone might be less effective without careful planning and assessment, effective security guidelines, enforcement, auditing, and sensible security policy design and assignment.

There is no exact definition of the measures that define standard security. These can vary widely, depending on an organization's policies and infrastructures. The following security levels can be considered a general basis for planning your IPSec deployment:

- Minimal security

  Computers do not exchange sensitive data. IPSec is not active by default. No administrative action to disable IPSec is required.

- Standard security

  Computers, especially file servers, are used to store valuable data. Security must be balanced so it does not become a barrier to users trying to perform their tasks. Windows XP and the Windows Server 2003 family provide example IPSec policies that secure data, but do not necessarily require the highest level of security: **Client (Respond Only)** and **Server (Request Security)**. Use these as a basis for your IPSec policies. These, or similar custom policies, optimize efficiency without compromising security.

- High security

  Computers that contain highly sensitive data are at risk for data theft, accidental or malicious disruption of the system (especially in remote dial-up scenarios), or any public network communications. **Secure Server (Require Security)**, a default policy, requires IPSec protection for all traffic being sent or received (except initial inbound communication) with stronger security methods. Unsecured communication with a non-IPSec-aware computer is not allowed.

**From planning to policy configuration**

After you have identified the types of traffic that need to be secured and the level of security required, you can begin policy configuration. Windows XP and Windows Server 2003 IPSec policy configuration is the translation of your security

requirements to one or more IPSec policies, only one of which can be assigned at each of the following levels: domain, site, organizational unit, or local levels. Each IPSec policy consists of one or more IPSec rules. Each IPSec rule consists of:

- A selected filter list.
- A selected filter action.
- Selected authentication methods.
- A selected connection type.
- A selected tunnel setting.

You can configure IPSec policy in the following ways:

- Create a new policy and define the set of rules for the policy, adding filter lists and filter actions as required.

  In this method, an IPSec policy is created first and then rules are added and configured. Filter lists (specifying traffic types) and filter actions (specifying how the traffic is treated) are added during rule creation.

  For more information, see Add, edit, or remove IPSec policies and Add, edit, or remove IPSec rules.

- Create the set of filter lists and filter actions, and then create the policies and add rules that combine the filter lists with filter actions.

  In this method, the filter lists and the filter actions are configured first. Next, IPSec policies are created and rules are added that combine the appropriate filter list with the appropriate filter action. Additionally, authentication methods, connection types, and tunnel settings are specified.

  For information about how to individually configure IP filter lists and filter actions, see Manage multiple IP filter lists and Manage multiple filter actions.

For either configuration method, after the IPSec policies have been created, they must be assigned. For more information, see Creating, modifying, and assigning IPSec policies.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- The integration of IPSec policies with Active Directory described here does not apply to computers running Windows XP Home Edition. You cannot administer Active Directory-based IPSec policy from a computer running Windows XP Home Edition because these computers cannot join Active Directory domains.

# Default IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Default IPSec policies**

Windows XP and the Windows Server 2003 family provide a set of predefined IPSec filter lists and filter actions and default policies. This is only intended to provide an example. It is not designed for operational use without modification. A knowledgeable IPSec network administrator or advanced user should design new, custom polices for operational use.

Computers that are connected to the Internet must be carefully protected from network attacks with custom policies. The default policies are defined as examples for intranets because they allow a computer to receive unsecured traffic. The default policies are designed for computers that are members of an Active Directory domain.

**Predefined filter lists**

The following predefined filter lists are provided as examples for use in the default policies:

- **All ICMP Traffic**

  A filter list for all ICMP traffic (IP protocol 1) sent and received between this computer and all other computers.

- **All IP Traffic**

  A filter list for all IP traffic sent and received between this computer and all other computers. By default, all Internet Key Exchange (IKE) traffic is excluded. All other traffic types are matched against IPSec filters, however, and you can configure, block, or permit filter actions specifically for multicast and broadcast traffic (IPSec does not negotiate security associations for multicast and broadcast traffic).

**Predefined filter actions**

The following predefined filter actions are provided as examples for use in the default policies:

- **Permit**

  Traffic is permitted (the security method is set to **Permit**).

- **Request Security (Optional)**

  Traffic is secured (the security method is set to **Negotiate security**), however, the following exceptions apply:

  - Initial incoming unsecured communication is allowed. This behavior is known as inbound passthrough.

  - If security negotiation fails after 3 seconds, unsecured communication to a non-IPSec-aware computer is allowed. This filter action secures traffic to IPSec-aware computers, but allows unsecured communication if it is required. This behavior is known as fallback to clear.

  In order to avoid the potential for denial-of-service attacks, this filter action must not be used on the Internet.

  The following table lists the security methods (in preference order) for the **Request Security (Optional)** predefined filter action.

| Type | AH integrity | ESP confidentiality | ESP integrity | Key lifetimes (KB/sec) |
|---|---|---|---|---|
| Custom | None | 3DES | SHA1 | 100000/900 |
| Custom | None | DES | SHA1 | 100000/900 |
| Custom | SHA1 | None | None | 100000/300 |

| | | | | |
|---|---|---|---|---|
| Custom | MD5 | None | None | 100000/300 |

The security negotiation for traffic using this filter action attempts ESP encryption (first 3DES, and then DES) with SHA1 integrity. If those security methods cannot be negotiated with the peer, AH integrity (first SHA1, and then MD5) only is attempted.

- **Require Security**

  Traffic is secured (the security method is set to **Negotiate security**), however, initial incoming unsecured communication is allowed (inbound passthrough is enabled). Unsecured communication with non-IPSec-aware computers is not allowed (fallback to clear is disabled). For example, an incoming connection request packet is allowed. The computer's reply to the connection request matches the outbound filter, causing a request for security to be sent to the sender of the unsecured traffic. If security negotiation is successful, traffic in both directions is required to be secure. If security negotiation is not successful, outgoing traffic is not allowed. The computer might continue to receive incoming unsecured traffic and will continue to attempt to negotiate security.

  In order to avoid the potential for denial-of-service attacks, this filter action must not be used on the Internet.

  The following table lists the security methods (in preference order) for the **Require Security** predefined filter action.

| Type | AH integrity | ESP confidentiality | ESP integrity | Key lifetimes (KB/sec) |
|---|---|---|---|---|
| Custom | None | 3DES | SHA1 | 100000/900 |
| Custom | None | 3DES | MD5 | 100000/900 |
| Custom | None | DES | SHA1 | 100000/900 |
| Custom | None | DES | MD5 | 100000/900 |

The security negotiation for traffic using this filter action will only attempt to negotiate ESP encryption (first 3DES, and then DES) with differing ESP authentication methods (first SHA1, and then MD5).

The most secure filter action is one that negotiates security and neither receives unsecured traffic (inbound passthrough is disabled) nor allows communication with non-IPSec-aware computers (fallback to clear is disabled). This is known as a lockdown filter action. When a server is using this filter action, clients must initiate secured communication before sending any IP traffic. The IPSec policy for client computers must be configured with additional rules that force the client to negotiate security with the server. A client configured with policy that consists of only the default response rule is not able to communicate with a locked down server.

**Predefined rules**

The only predefined rule that is available on all policies is the default response rule. The default response rule is automatically created in a new policy and cannot be created manually. It is used to respond to security requests when no other rule applies. For example, a client policy that enables a client to respond to security requests but does not secure any traffic by default, would have only the default response rule active. In some cases, it is enabled by default. You can disable the default response rule on the **Rules** tab in properties for a policy. You can also modify the authentication method and security methods of the default response rule for specific needs. For more information, see the section "Default response rule" in IPSec Policy Rules.

**Default policies**

The following are provided as examples for use as default policies:

- **Client (Respond Only)**

  This is an example policy for computers that secure communication on request. For example, intranet clients might not require IPSec except when requested by another computer. This policy enables the computer on which it is active to respond to requests for secured communications. This policy contains the default response rule, which creates dynamic IPSec filters for inbound and outbound traffic based on the requested protocol and port traffic for the communication that is being secured.

  This policy has the following settings:

**First rule (default response rule)**

- IP Filter List: <Dynamic>
- Filter Action: Default Response
- Authentication: Kerberos
- Tunnel Setting: None
- Connection Type: All

- **Server (Request Security)**

  This is an example policy for computers that should secure communications in most cases, allowing unsecured communication with non-IPSec-aware computers. With this policy, the computer accepts unsecured traffic, but always attempts to secure additional communications by requesting security from the original sender. This policy allows the entire communication to be unsecured if the other computer is not IPSec-enabled. For example, communication to specific servers can be secure while allowing the server to communicate in an unsecured manner to accommodate a mixture of clients (some of which support IPSec and some of which do not).

  This policy has a rule to request security for all IP traffic, a rule to permit ICMP traffic, and the default response rule to respond to requests for security from other computers.

  This policy has the following settings:

  **First rule**

  - IP Filter List: All IP Traffic
  - Filter Action: Request Security (Optional)
  - Authentication: Kerberos
  - Tunnel Setting: None
  - Connection Type: All

  **Second rule**

  - IP Filter List: All ICMP Traffic
  - Filter Action: Permit
  - Authentication: N/A
  - Tunnel Setting: None
  - Connection Type: All

  **Third rule (default response rule)**

  - IP Filter List: <Dynamic>
  - Filter Action: Default Response

- Authentication: Kerberos
- Tunnel Setting: None
- Connection Type: All

- **Secure Server (Require Security)**

  This is an example policy for computers on an intranet that require secure communications, such as a server that transmits highly sensitive data. Administrators can use this IPSec policy as an example to create their own custom IPSec policy for production use. The filters used in this policy require all outbound communication to be secured, allowing only the initial inbound communication request to be unsecured.

  This policy has a rule to require security for all IP traffic, a rule to permit ICMP traffic, and the default response rule to respond to requests for security from other computers.

  **First rule**

  - IP Filter List: All IP Traffic
  - Filter Action: Require Security
  - Authentication: N/A
  - Tunnel Setting: None
  - Connection Type: All

  **Second rule**

  - IP Filter List: All ICMP Traffic
  - Filter Action: Permit
  - Authentication: Kerberos
  - Tunnel Setting: None
  - Connection Type: All

  **Third rule (default response rule)**

  - IP Filter List: <Dynamic>
  - Filter Action: Default Response
  - Authentication: Kerberos
  - Tunnel Setting: None
  - Connection Type: All

**Notes**

- In the Windows Server 2003 family, IPSec provides new source and destination address options for defining filters. In addition, the default behavior for traffic exemptions has changed. If you require previously exempted traffic types, you

must create filters to permit this traffic. For more information about the new source and destination options provided by IPSec and for instructions on how to configure permit filters for several previously exempted traffic types, see Special IPSec considerations.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

# Creating, modifying, and assigning IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating, modifying, and assigning IPSec policies**

IPSec policies can be applied to local computers, domains, sites, organizational units, or any Group Policy object in Active Directory. You can create, modify, and assign IPSec policies by using the IP Security Policy Management console that is available in Microsoft Management Console (MMC) or by using the Netsh commands for IPSec that are available at the command line.

Your organization's IPSec policies should be based on your written guidelines for secure operations. Policies can store multiple rules, so that one policy can be used to affect multiple types of traffic.

The Windows 2000, Windows XP, and Windows Server 2003 family include example IPSec policies. These examples are provided only to show how policies can be designed. They are not designed as default policies for securing a client or a server in a production environment. Instead, you should design and create custom IPSec policies that apply to the specific scenarios in your security plan.

### Notes

- For information about how to use the IP Security Policy Management console, see Define IPSec Policies.
- For information about how to use Netsh, see Netsh commands for Internet Protocol security.
- Although you can assign local IPSec policies for a computer by using the IP Security Policy Management console, you must use Group Policy to assign IPSec policies for groups of computers in Active Directory domains. For more information, see Assign or unassign IPSec policy in Group Policy.

**IPSec driver modes**

To understand how IPSec policies are processed and applied, it is important to understand the modes in which the IPSec driver operates. The IPSec driver operates in computer startup mode and operational mode and, optionally, in a diagnostics mode for troubleshooting.

### Computer startup mode

During computer startup, the Windows operating system loads the IPSec driver to perform in computer startup mode. Computer startup mode is used until the IPSec Policy Agent sets the IPSec driver into an operational mode. The IPSec driver can perform in any one of the following computer startup modes:

- **Permit.** When the IPSec driver operates in permit mode, it does not process any IP packets, and, therefore, no IPSec security is provided. If an IPSec policy has never been assigned to a computer, permit mode is the default mode for the IPSec driver.
- **Stateful.** When the IPSec driver operates in stateful mode, it allows all outbound traffic, and it automatically creates inbound permit filters in response to the outbound traffic. All other inbound unicast, broadcast, and multicast packets are dropped. If an IPSec policy has been assigned to a computer, the IPSec Policy Agent sets the stateful mode for the IPSec driver by default.
- **Block.** When the IPSec driver operates in block mode, it discards all inbound and outbound IP packets, except for those that match specific filters that you can configure to be used during block mode. By default, all inbound and outbound DHCP traffic is permitted so that an IP address can be obtained. You can configure filters for block mode by using the **netsh ipsec dynamic set config bootexemptions** command.

To change the default mode for the IPSec driver during computer startup, use the **netsh ipsec static set config bootmode** command.

### Operational mode

After the IPSec service starts, the IPSec Policy Agent can set the IPSec driver to any one of the following operational modes. The operational mode that is chosen does not change the mode that the IPSec driver uses during computer startup.

- **Secure.** When the IPSec driver operates in secure mode, IPSec policy filters are enforced for standard IPSec operation. The IPSec Policy Agent sets the driver into secure mode after it applies persistent policy (if a persistent policy has been configured) and before it applies the Active Directory-based policy or the local policy. If no persistent policy has been configured and assigned, secure mode does not provide IPSec protection until either the assigned Active Directory-based policy or the local policy is applied. When the Active Directory-based policy or the local policy changes, the IPSec Policy Agent updates the IPSec filter list in the IPSec driver. If no IPSec policy is assigned, the IPSec Policy Agent deletes the filters from the IPSec driver and no IPSec protection is provided in this mode.

- **Permit.** When the IPSec driver operates in permit mode, it does not process any IP packets, and, therefore, no IPSec protection is provided. Permit mode is used when the IPSec service is manually stopped.

- **Block.** When the IPSec driver operates in this mode, the exemptions that might apply during computer startup are not applied. Instead, all inbound and outbound traffic is blocked. This mode is used to enhance security, in the event that the IPSec Policy Agent fails to apply a persistent policy, if a persistent policy has been configured.

The operational mode can be configured only by the IPSec Policy Agent. You cannotuse the Netsh commands for IPSec to configure this mode.

### Diagnostic mode

To record all inbound and outbound dropped packets and other packet processing errors in the Event Viewer System log for troubleshooting, you can enable IPSec driver event logging and specify the level of logging to provide. The types of packet processing errors that the IPSec driver records in the System log depend on the level of logging that is provided. IPSec driver logs can record inbound and outbound per-packet drop events during computer startup mode and operational mode.

IPSec driver event logging is disabled by default, and it should not be used for extended periods. Depending on the logging level that you set, many events might be generated that fill the System log very quickly. For more information, see the section "Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family" in IPSec troubleshooting tools.

#### Application methods for IPSec policy

IPSec policies are provided in six ways, in the order described in the following sections.

#### IPSec startup policy

When a computer loads the Windows operating system during startup, the IPSec driveris loaded at the same time as the TCP/IP driver. The following table describes thebehavior of the IPSec driver based on the startup type of the IPSec service.

| IPSec service startup type | IPSec driver behavior |
|---|---|
| Disabled | The IPSec driver loads in permit mode, meaning that no IPSec security is applied. In this mode, the IPSec driver does not filter any traffic. |
| Manual | The IPSec driver loads in permit mode, meaning that no IPSec security is applied. In this mode, the IPSec driver does not filter any traffic. |
| Automatic | The IPSec driver loads in a startup mode configured by the IPSec Policy Agent. If no IPSec policy is assigned, by default the IPSec driver loads in permit mode. If IPSec policy is assigned, the IPSec Policy Agent configures the IPSec driver to load in stateful mode.<br><br>After the IPSec service starts, persistent policy (if it has been configured) is applied, and the IPSec driver is set to secure mode (normal operation). |

When the IPSec driver is loaded, it reads the registry to determine the actions that it must perform during computer startup. These actions include:

- Whether to apply stateful filtering, permit all traffic, or block all traffic.

- Whether to exempt specific traffic types from filtering during computer startup.
- Whether to enable IPSec diagnostic logging and, if so, which level of logging to provide.
- Whether to modify the default IPSec traffic exemption. By default, in the Windows Server 2003 family, RSVP, Kerberos, multicast, and broadcast traffic are not exempt from IPSec filtering. Only IKE traffic is exempt.

The IPSec driver remains in a startup mode until the IPSec Policy Agent sets it to an operational mode. After the IPSec driver is set to secure operation mode, the startup stateful mode settings are discarded. You can use the Netsh commands for IPSec to configure IPSec driver settings. For more information, see the **netsh ipsec dynamic set config** section in Netsh commands for Internet Protocol security.

**Persistent policy**

In the Windows Server 2003 family, you can use the Netsh commands for IPSec to add, modify, and maintain IPSec policy settings permanently, after you run the appropriate commands. This type of policy is known as a persistent policy. You can configure persistent policies to extend existing Active Directory-based or local IPSec policies, override Active Directory-based or local IPSec policies, and enhance security during computer startup. Persistent policies enhance security by providing a secure transition from computer startup to Active Directory-based IPsec policy enforcement.

Persistent policies provide the following two key benefits:

- You can configure a persistent policy as your most restrictive IPSec policy. For example, the persistent policy can block all inbound traffic to the Internet IP address of a computer, except for ICMP traffic and communication with a particular management station. To complement the persistent policy, you can design an Active Directory-based policy or a local policy with additional rules that include more specific filters to allow additional communication types.
- You can selectively enforce security rules that are required only for specific servers, due to unique network adapter configurations or program needs. For example, you can enforce security rules specifying that:
    - A specific port always be blocked.
    - Traffic that uses a specific network adapter be permitted without IPSec security, or only with IPSec security.
    - Unsecured traffic to a specific monitoring station be permitted.

Persistent policies, if configured, are stored in the local registry. The persistent store is loaded by the IPSec Policy Agent during service startup. After a persistent policy is successfully applied, the IPSec Policy Agent sets the IPSec driver to secure mode. If an error occurs when a persistent policy is applied, the IPSec Policy Agent removes the persistent policy settings from the IPSec driver and sets the IPSec driver to block operation mode for security. You can update a persistent policy at any time, as long as the IPSec service is running. However, changes in persistent policy are not active immediately. You must restart the IPSec service to load the new persistent policy settings.

If you have configured Active Directory-based policies, you can use a persistent policy as a tool to require that traffic to Active Directory always be secured by IPSec, including the retrieval of Active Directory-based IPSec policies. When an Active Directory-based or local policy is applied, those policy settings are added to the persistent policy settings.

To configure persistent policies, you must use the **netsh ipsec static set store location=persistent** command. To view the IPSec policy that is currently being applied, you can use either IP Security Monitor or the **netsh ipsec dynamic show** commands. For information about how to use Netsh, see Netsh commands for Internet Protocol security.

**Important**

- To provide maximum protection against attacks during computer startup, it is highly recommended that you configure and assign a persistent policy. If you do not configure a persistent policy, no policy can be enforced by the IPSec driver until the IPSec Policy Agent retrieves and applies the assigned Active Directory-based or local IPSec policy.

**Active Directory-based policy**

A Group Policy object defines access, configuration, and usage settings for accounts and resources. IPSec policies can be assigned to the Group Policy object of a site, domain, or organizational unit. When the IPSec policy is applied to one of the

Group Policy objects for the Active Directory object, the IPSec policy is propagated to all of the computer accounts that are affected by that Group Policy object.

You can manage Active Directory-based policy by using the IP Security Policy Management console or by using the **netsh ipsec static set store location=domain** command. IPSec policies that are configured and assigned for the domain take precedence over the local, active IPSec policy on a computer, when that computer is a member of the domain. Active Directory-based policy overrides any local IPSec policy that is assigned, and it adds to the persistent IPSec policy that has already been applied by the IPSec Policy Agent, if a persistent policy has been configured.

**Important**

- Several features in the Windows Server 2003 family implementation of IPSec are not provided in Windows 2000 or in Windows XP. To ensure that the same IPSec policy functions as expected on computers running the Windows Server 2003 family and on computers running Windows 2000 or Windows XP, test the policy thoroughly on all relevant operating systems before deployment. For more information about new features in IPSec, see New features for IPSec.

- If you plan to apply IPSec policies that use the new features that are available only in the Windows Server 2003 family implementation of IPSec, do not use the Windows 2000 or the Windows XP version of the IP Security Policy Management console to manage these policies. The settings in the earlier versions of the IP Security Policy Management console will override the settings in the Windows Server 2003 family IPSec policy, and the new features will not be functional.

When assigning an IPSec policy in Active Directory, consider the following:

- The list of all IPSec policies is available to assign at any level in the Active Directory hierarchy. However, only a single IPSec policy can be assigned at a specific level in Active Directory.

- An IPSec policy that is assigned to an organizational unit in Active Directory takes precedence over a domain-level policy for members of that organizational unit.

- An IPSec policy that is assigned to the lowest-level organizational unit in the domain hierarchy overrides an IPSec policy that is assigned to a higher-level organizational unit, for member computers of that organizational unit.

- An organizational unit inherits the policy of its parent organizational unit unless either policy inheritance is explicitly blocked or policy is explicitly assigned.

- IPSec policies from different organizational units are never merged.

- The highest possible level of the Active Directory hierarchy should be used to assign policies to reduce the amount of configuration and administration required.

- An IPSec policy might remain active even after the Group Policy object to which it is assigned has been deleted. Because of this, you should unassign the IPSec policy before you delete the policy object. To prevent problems, use the following procedure:

  1. Unassign the IPSec policy in the Group Policy object.

  2. Wait 24 hours to ensure that the change is propagated.

  3. Delete the Group Policy object.

     If you delete the Group Policy object without following this procedure, computers in the Active Directory container to which the IPSec policy is assigned treat the IPSec policy as if it cannot be located and continue to use a cached copy.

- You can use Group Policy backup and restore tools to store information about which policy is assigned to a Group Policy

object. To ensure consistency, when you use these tools, make sure that you include IPSec policies. Because these tools do not store the IPSec policies themselves, however, you cannot use them to back up and restore the IPSec policies themselves. You can only back up and restore information about whether the IPSec policies are assigned to Group Policy objects and, if so, to which Group Policy objects. To back up and restore all of the IPSec policies in the IPSec container in Active Directory, use the Export Policies and Import Policies commands in the IP Security Policy Management console. For more information, see Export IPSec policies and Import IPSec policies.

- Before assigning an IPSec policy to a Group Policy object, verify the Group Policy settings that are required for the IPSec policy. For example, if an IPSec policy requires certificate authentication, assign the Group Policy settings that allow computers to enroll for certificates (usually one or two days before you assign the IPSec policy that requires use of the computer certificate). In addition, you should test the certificate enrollment process and resolve any errors before assigning the IPSec policy.

The IPSec Policy Agent on a computer running Windows XP Professional or a Windows Server 2003 operating system polls Active Directory for updates to the assigned IPSec policy. This polling does not detect a change in domain or organizational unit membership or the assigning or unassigning of a new IPSec policy. These events are detected when the Winlogon service polls for changes in Group Policy, which occurs by default every 90 minutes. The Winlogon service discovers these changes, notifies the IPSec Policy Agent, and the IPSec policy changes are applied.

For more information about Group Policy, see Group Policy overview.

### Notes

- The integration of IPSec policies with Active Directory described here does not apply to computers running Windows XP Home Edition because you cannot administer Active Directory-based IPSec policy from a computer running Windows XP Home Edition.

- Careful planning is required when you design Active Directory-based IPSec policies for computers that are used the internal network and on other public or partner networks.

- By default, IPSec policies are read from Active Directory when the connection to Active Directory is interpreted by Group Policy to be a slow link.

#### Directory cache policy

If you assign an Active Directory-based IPSec policy, a copy of the current policy is maintained in a cache in the local registry. If the computer to which the policy is assigned cannot connect to the domain, the cached copy of the Active Directory-based policy is applied. This cached policy adds to the settings of the persistent policy, if a persistent policy has been configured. Note that you cannot configure or manage the cached copy of an Active Directory-based IPsec policy.

#### Local computer policy

Each computer running Windows XP or a Windows Server 2003 operating system has exactly one local Group Policy object, often called the local computer policy. In using this local Group Policy object, Group Policy settings can be stored on individual computers regardless of whether they are members of an Active Directory domain. The local Group Policy object can be overwritten by Group Policy objects associated with sites, domains, or organizational units in an Active Directory environment. On a network without an Active Directory domain (a network that lacks a Windows 2000 or a Windows Server 2003 domain controller), the local Group Policy object settings determine IPSec behavior because they are not overwritten by other Group Policy objects.

You can manage local IPSec policies by using the IP Security Policy Management console or by using the **netsh ipsec static set store location=local** command. The settings of a local IPSec policy are added to the persistent policy, if a persistent policy has been configured. If an Active Directory-based IPSec policy is assigned, and if the computer is connected to an Active Directory domain, the settings of the Active Directory-based policy are applied instead.

For information about how to assign IPSec policy locally, see Assign or unassign IPSec policy on a computer.

#### Netsh dynamic mode policy

You can use the **netsh ipsec dynamic** commands to create, modify, and assign IPSec rules that take effect immediately and are not stored. If the IPSec service is stopped, either administratively or when the computer is restarted, these rules are lost. However, if you use the **netsh ipsec dynamic set config** commands to change the IPSec configuration, those settings are not lost.

# Special IPSec considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Special IPSec considerations**
The following special IPSec considerations can help simplify the administration of IPSec policies.

**Recommended IPSec scenarios and IPSec scenarios that are not recommended**
Following are recommended scenarios and scenarios that are not recommended for the Windows Server 2003 family implementation of IPSec.

**Recommended IPSec scenarios**
The Windows Server 2003 family implementation of IPSec is recommended for use as follows:

- Packet filtering. IPSec provides limited firewall capabilities for end-systems. You can also use IPSec with Internet Connection Firewall, Windows Firewall, and Routing and Remote Access to permit or block inbound or outbound traffic.

- Securing host-to-host traffic on specific paths. You can use IPSec to provide mutual authentication and cryptographic protection for traffic between servers or other static IP addresses or subnets. For example, IPSec can secure traffic between domain controller sites or forests, or between Web servers and database servers.

- Securing traffic to servers. You can use IPSec to require mutual authentication for all client computers that access a server. In addition, you can set restrictions on which computers are allowed to connect to a server running Windows Server 2003 family.

- Using L2TP/IPSec tunneling for VPN connections. You can use the combination of the Layer Two Tunneling Protocol (L2TP) and IPSec (L2TP/IPSec) for all virtual private network VPN scenarios.

- Using IPSec in tunnel mode for gateway-to-gateway tunnels. For interoperability with other routers, gateways, or end-systems that do not support L2TP/IPSec or PPTP VPN tunneling, you can use IPSec in tunnel mode for gateway-to-gateway tunnels.

### Notes

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**IPSec scenarios that are not recommended**
When using IPSec authentication and data protection features, the policy management model in Windows 2000, Windows XP, and Windows Server 2003 family lends itself best to server-to-server and client-to-server scenarios where one end-point has a static address. In large network deployments, when policy involves dynamic addresses at both end–systems, and in some mobility cases, policy management complexity can make IPSec deployment difficult. Because of this, IPSec is not recommended for use as follows:

- Securing communication between domain members and their domain controllers. Because of the complexity of the IPSec policy configuration and management required for this scenario, using IPSec to secure communication between domain members and their domain controllers is not recommended.

- Securing all traffic in a network. Configuring IPSec communications between all client computers and all servers in order to secure many or all computers in a network is difficult. IPSec cannot negotiate security for multicast and broadcast traffic. Traffic from real-time communications, applications that require ICMP, and peer-to-peer applications might be incompatible with IPSec. For these reasons, using IPSec to secure all traffic in a network is not recommended.

Additionally, the IPSec protocol and implementation has characteristics that require special consideration in other scenarios:

- Securing traffic over wireless 802.11 networks. IPSec policies are not optimized for configuration for mobile clients. Because of this, IPSec is not recommended as the only method to secure traffic sent over wireless 802.11 networks. Instead, it is recommended that you use IEEE 802.1X authentication. 802.1X enhances security and ease of deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting. In cases where clients roam between access points on the same network, IPSec can be used in combination with 802.11 and 802.1x. In cases where roaming causes the client IP address to change, IPSec security associations become not valid and new security associations are renegotiated.

- Using IPSec in tunnel mode for remote access VPN connections. Using IPSec in tunnel mode is not recommended for remote access VPN scenarios. Instead, use L2TP/IPSec or PPTP for remote access VPN connections.

**Managing policies that use new features available only in the Windows Server 2003 family**
The following features are available only in the Windows Server 2003 family implementation of IPSec:

- Support for new source and destination address options for defining filters. For example, you can now create:

    - Filters with a source address and destination address of **Any IP Address**.

    - Filters that use a non-standard subnet class as a source address or a destination address.

    - Filters that use a multicast or a broadcast address as a destination address.

    - Filters with a source address or a destination address that corresponds to the local IP configuration for the DNS servers, WINS servers, and the DHCP server and default gateway for the computer to which the filter applies.

- The ability to exclude the name of the certification authority (CA) from certificate requests.

- Certificate-to-account mapping for network access control.

- Support for root CA names that are represented in extended ASCII character sets.

- A stronger Diffie-Hellman group (2048-bit Diffie-Hellman key exchange).

Computers running Windows XP and Windows 2000 cannot use these new policy settings. If you plan to apply Active Directory-based IPSec policies that use features that are only available in the Windows Server 2003 family, use the Windows Server 2003 family version of the IP Security Policy Management console to manage them. If you use the Windows XP or Windows 2000 version of the IP Security Policy Management console to manage these policies, the earlier versions of the IP Security Policy Management console will delete the new feature settings.

In addition, if you want to apply the same IPSec policy to computers running different versions of Windows, test the policy to ensure that it functions as expected on both servers running Windows Server 2003 and computers running Windows XP or Windows 2000.

For information about new features available with the Windows Server 2003 family implementation of IPSec, see New features for IPSec. For more information about best practices for using IPSec, see IPSec Best practices.

**IP filters**
Consider the following when configuring IP filters:

- Use general filters when you want a single filter to apply to a group of computers. For example, when configuring the filter, use **My IP address**, **Any IP Address**, or a subnet address instead of specifying a specific computer's source and destination IP addresses.

- Define filters that allow you to group and secure traffic from logically associated segments or subnets within your network.

- Be aware that when you are viewing the IPSec policy, the order in which the filters apply is not in the order displayed. When the IPSec Policy Agent reads an IPSec policy, filters are processed into one ordered list, sorted from most to least

specific. You can use the IP Security Monitor console to view filters sorted by weight. Filters that have the same weight might appear in any order relative to each other.

During policy updates, the IPSec Policy Agent recalculates the sorted order and updates the IPSec driver with the changes only. If a filter is deleted, or if the filter action is changed for a filter, all security associations related to that filter are deleted. As a result, some packets are lost when a new security association is negotiated. TCP connections, however, should immediately recover by retransmitting the lost packets.

- If you are configuring an IPSec policy to secure Kerberos traffic, IPSec can only negotiate security associations if your authentication method does not use Kerberos. If Kerberos is required for authentication between domain members, special consideration must be given to the traffic required by Kerberos between domain members and domain controllers, and the following requirements must be met:

  - If a domain member and a domain controller are members of different domains, a two-way trust must exist between the domains or Internet Key Exchange (IKE) rekeys will fail.

  - If a client computer or a server in the domain uses Kerberos authentication to negotiate IPSec with all peers, you must create mirrored permit filters for all Kerberos-dependent traffic to all domain controllers. It is recommended that you create permit filters to exempt all traffic to each domain controller IP address in the domain member's domain. You do not need to create permit filters to exempt all traffic to all trusted domain controller IP addresses. Instead, create permit filters to exempt only Kerberos traffic and domain controller site location (LDAP) queries to all domain controller IP addresses in the trust path between the computers. If DNS service is not provided by the domain controller (for which you already permit all traffic), you might also need to exempt DNS traffic to the DNS server IP addresses so that DNS can be used for domain controller discovery.

    To exempt these traffic types in the IPSec policy, create the following filters, each with a permit filter action:

| **Filter** | **Filter Action** |
|---|---|
| All traffic to my domain controllers | • Mirrored<br>• Source Address = My IP Address<br>• Destination Address = **A specific IP Address** (where **A specific IP Address** is the static IP address of each domain controller in the domain)<br>• Protocol = **Any**<br>• This filter permits LDAP site location queries and ICMP, DNS, and Kerberos traffic. In addition, it permits the remote procedure call (RPC) traffic and Group Policy downloads required to manage Active Directory-based IPSec policies. |
| Kerberos | • Mirrored<br>• Source Address = My IP Address<br>• Destination Address = **A specific IP Address** (where **A specific IP Address** is the static IP address of each domain controller in the domain)<br>• Protocol = **Any**<br>• This filter permits LDAP site location queries and ICMP, DNS, and Kerberos traffic. In addition, it permits the remote procedure call (RPC) traffic and Group Policy downloads required to manage Active Directory-based IPSec policies. |

| | |
|---|---|
| Domain controller site location (LDAP) queries | • Mirrored<br>• Source Address = **My IP Address**<br>• Destination Address = **A specific IP Address** (where **A specific IP Address** is the static IP address of each domain controller in trusted domains)<br>• Protocol = **UDP**<br>• Source port = **Any**<br>• Destination port = **88**<br>• After you create this filter, create another filter with the same settings, but specify **TCP** for the protocol. |
| DNS client queries | • Mirrored<br>• Source Address = **My IP Address**<br>• Destination Address = **DNS Servers <dynamic>** (or **A specific IP Address**, where the IP address is the static IP address of each DNS server in the network)<br>• Protocol = **UDP**<br>• Source port = **Any**<br>• Destination port = **53**<br>• After you create this filter, create another filter with the same settings, but specify **TCP** for the protocol. |

**Caution**

- To minimize the security issues, you should allow these exemptions only for the specific IP addresses of the domain controllers and the DNS servers.

When designing an IPSec policy to secure or block ICMP traffic, ensure that you plan carefully and test the policy before implementing it in a production environment. Be aware that ICMP is required for the following scenarios:

- Outbound TCP connections that traverse networks with different link frame sizes require path maximum transmission unit (PMTU) discovery. For PMTU discovery to function correctly, inbound ICMP traffic from routers and other gateways must be received.

- The **tracert** command requires ICMP traffic to determine the path taken to a destination. If IPSec protects ICMP traffic, however, **tracert** output shows only the destination computer--not the path of intermediate routers.

- File sharing and other applications use ICMP traffic to determine whether a destination IP address can be reached.

- Network logon and Group Policy measure ICMP response times to each domain controller IP address. This measurement is performed to both aid in the selection of a domain controller to service a network logon and determine whether to apply Group Policy settings over slow links.

In some cases, even the three-second delay before a soft SA (an unsecured connection) is established might impact upper-

layer communication that requires fast ICMP responses.

- In Windows 2000 and Windows XP, by default, broadcast, multicast, Kerberos, Internet Key Exchange (IKE), and Resource Reservation Protocol (RSVP) traffic are exempt from filter matches. In the Windows Server 2003 family, broadcast, multicast, and Kerberos, are not exempt from filter matches by default (only IKE traffic is exempt). In additon, Resource Reservation Protocol (RSVP) has been turned off. Broadcast and multicast packets will be dropped if they match a filter with a filter action to negotiate security. By default, the Windows Server 2003 family provides limited support for filtering broadcast and multicast traffic. A filter with a source address of **Any IP Address** will match multicast and broadcast addresses. A filter with a source address of **Any IP Address** and a destination address of **Any IP Address** will match inbound and outbound multicast addresses. You can use such a filter to block all traffic. One-way filters that would be used to block or permit specific multicast or broadcast traffic, however, are not supported.

  As a result of the change in default exemption behavior for the Windows Server 2003 family implementation of IPSec, you should verify the behavior of IPSec policies designed for Windows 2000 or Windows XP, and determine whether to configure explicit permit filters to permit specific traffic types. To restore the default Windows 2000 and Windows XP behavior for IPSec policies, you can use the **netsh ipsec dynamic set config command**, or you can modify the registry. For more information, see the section "Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family" in IPSec troubleshooting tools.

  **Caution**

  - The Windows 2000 and Windows XP default exemption settings for IPSec are designed for corporate LAN environments with a low risk of attack. For this reason, you should only use the Windows 2000 and Windows XP default exemption settings when necessary for troubleshooting, in low-risk environments, or when you cannot solve program compatibility issues by configuring explicit filters in IPSec policies.

**Safe Mode with Networking**

If you start a computer in Safe Mode with Networking, consider the following:

- The IPSec service cannot apply filters from an IPSec policy or negotiate IP security when the computer is in Safe Mode with Networking. As a result, when Directory Services Restore Mode operations are in progress, communication does not occur with remote domain controllers or network backup devices that enforce an IPSec policy to require IPSec.

- When a computer loads the Windows operating system during startup, the IPSec driver is loaded at the same time as the TCP/IP driver. If an IPSec policy is assigned to the computer and the IPSec service is configured for automatic startup, by default, the IPSec driver startup mode is changed to perform stateful filtering. When stateful filtering is performed, all outbound traffic is permitted, but all inbound traffic is blocked, except for DHCP traffic and traffic that is sent in response to the outbound traffic initiated during computer startup. (The IPSec driver automatically creates inbound filters to permit responses to outbound traffic.) The IPSec service itself, however, does not start. As a result, no persistent, local, or Active Directory-based IPSec policies can be applied.

- When the computer is in Safe Mode with Networking, you can allow the IPSec driver to continue to provide stateful filtering protection by configuring inbound exemptions for specific traffic. Or, you can force the IPSec driver to permit all traffic.

- To configure inbound exemptions for specific traffic, restart the computer normally (not in Safe Mode), use the **netsh ipsec dynamic set config bootexemptions** command to specify the traffic types to exempt during computer startup, and then restart the computer. The new exemptions will take effect when the computer is restarted and will remain in effect until they are deleted or changed.

- If you allow the IPSec driver to configure inbound exemptions for specific traffic, keep in mind that unless you configure the appropriate exemptions, the following will occur:

  - Inbound connections will be blocked to network services, such as Remote Desktop Connection and Remote Assistance, which open ports for receiving connections. To allow clients to use Remote Desktop Connection or Remote Assistance, you must configure an exemption to allow inbound traffic to TCP port 3389. To configure this exemption, use the following Netsh command:

**netsh ipsec dynamic set config bootexemptions TCP:any:3389:inbound**

- Some outbound TCP connections, such as those requiring PMTU discovery, might not work. To enable TCP PMTU discovery, you must configure an exemption to allow inbound ICMP traffic. To configure this exemption, use the following Netsh command:

  **netsh ipsec dynamic set config bootexemptions ICMP:inbound**

  **Important**

- Modifying IPSec traffic exemptions from startup security (that is, modifying the **bootexemptions=** parameter) will overwrite all previous exemptions from startup security.

For more information about how to specify traffic types to exempt during computer startup, see Netsh commands for Internet Protocol security.

To force the IPsec driver to permit all inbound traffic while the computer is in Safe Mode with Networking, you must either configure the IPSec service for manual startup or disable it, and then restart the computer in Safe Mode with Networking.

**Filter actions**
Consider the following special requirements when configuring IP filter actions:

- If there is specific traffic that you always want to block, such as traffic to and from either not valid addresses or specific TCP ports, create a blocking filter action with the **Security Method** set to **Block**.

- If there is specific traffic that you always want to be unsecured, create a permit filter action with the **Security Method** set to **Permit**.

- When you are configuring custom security methods for encrypted traffic, you should only set the ESP confidentiality selection to **<None>** when a higher layer protocol provides data encryption.

- When you are planning Internet scenarios, such as IPSec tunneling between branch offices, consider using a list of security methods that specifies high levels of security (for example, 3DES only), short key lifetimes (less than 50 MB), and perfect forward secrecy (PFS) for the session keys. This helps protect against known key attacks.

**Note**

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

**VPN connections using L2TP**
When you are using Layer Two Tunneling Protocol (L2TP) for your remote access or router-to-router VPN connections, certificates are the default authentication method for main mode security associations. Computer certificates must be installed either on both the VPN server and remote access client (for remote access connections), or on both the calling router and answering router (for router-to-router VPN connections). The simplest method for installing computer certificates is to configure your Windows 2000 and Windows Server 2003 domains to auto-enroll computer certificates for domain members. For more information, see Computer certificates for L2TP/IPSec VPN connections.

When you are using the certificate authentication method for L2TP connections, the list of certification authorities (CAs) is not configurable. Instead, each computer in the L2TP connection sends a list of root CAs to its IPSec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Computer A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPSec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPSec peer, Computer B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPSec

main mode negotiation fails. Therefore, ensure one of the following before attempting an L2TP connection:

- Both the VPN client and VPN server were issued computer certificates from the same CA.
- Both the VPN client and VPN server were issued computer certificates from CAs that follow a certificate chain up to the same root CA.

In general, the VPN client must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN server trusts. Additionally, the VPN server must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN client trusts.

For example, if the VPN client was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies the VPN server during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the VPN server does not have a valid computer certificate issued from a CA that follows a certificate chain to either CertAuth1 or CertAuth2, IPSec main mode negotiation fails.

A single CA commonly issues computer certificates to all computers in an organization. Because of this, all computers within the organization both have computer certificates from a single CA and request certificates for authentication from the same single CA.

Windows XP and Windows Server 2003 family VPN clients also support the use of preshared key authentication to create L2TP/IPSec connections.

**Important**

- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing or as an interim measure in a production environment, when your public key infrastructure (PKI) is being deployed to obtain computer certificates, or when preshared key authentication is required by your VPN server. For long-term use in production environment, it is recommended instead that you use certificates or Kerberos V5.

For a remote access VPN client running Windows XP or a Windows Server 2003 operating system, you can enable preshared key authentication and configure a preshared key in **IPSec settings** on the **Security** tab in the properties of a VPN connection in Network Connections.

**Domain controllers and DHCP, DNS, or WINS servers**
Before enabling IPSec for a computer that functions as a domain controller or DHCP, DNS, or WINS server, determine whether all of the clients are IPSec-enabled. If the IPSec policy assigned to the domain controller or DHCP, DNS, or WINS server is not configured to allow fall back to clear, which permits unsecured communication with non-IPSec-capable clients, secure negotiation will fail and access to these network services will be blocked.

**SNMP**
If you are using Simple Network Management Protocol (SNMP) and your SNMP management systems or agents include computers running Windows 2000, Windows XP, or a Windows Server 2003 operating system that are configured with IPSec policy, you can add a rule to the appropriate IPSec policy to prevent SNMP messages from being secured. The filter list must specify the source and destination addresses of the SNMP management systems and agents for all UDP traffic that uses ports 161 and 162. The filter action should be set to **Permit**, which passes through any traffic that matches the IP filter list.

# IPSec troubleshooting tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting tools**

This topic provides information about the following IPSec troubleshooting tasks and the troubleshooting tools that you can use to perform each of these tasks:

- Viewing IPSec policy assignment information
- Viewing details about the active IPSec policy and IPSec statistics
- Viewing details about the active IPSec policy and IPSec statistics in IP Security Monitor
- Verifying that security auditing is enabled
- Viewing IPSec-related events in Event Viewer
- Enabling detailed tracing for Internet Key Exchange (IKE) negotiations
- Viewing IPSec and other network communication with Network Monitor
- Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family
- Using Ipseccmd.exe to manage and monitor IPSec on computers running Windows XP
- Using Netdiag.exe to display IPSec information and to test and view network configuration
- Disabling TCP/IP and IPSec hardware acceleration

Details about each of these subjects are provided in the following sections.

**Notes**

- For general recommended practices that can help you enhance security and minimize the potential for problems when deploying IPSec, see IPSec Best practices.
- For the Windows Server 2003 family, the Netsh commands for IPSec replace and enhance the functionality provided by Netdiag.exe and Ipseccmd.exe. For information about how to use Netsh to script IPSec policy creation and to monitor IPSec activity, see Netsh commands for Internet Protocol security.
- Ipseccmd.exe is a support tool that is provided only on the Windows XP CD. To install this tool, run the Setup.exe file from the Support\Tools folder. You can use this tool to manage and monitor IPSec policies only on computers running Windows XP. For more information about Ipseccmd.exe, see the Microsoft Tech Net Web site.

Netdiag.exe is available in the Windows Server 2003 family, Windows XP, and in Windows 2000, as follows:

- In the Windows Server 2003 family, run the Suptools.msi file from the \Support\Tools folder on the Windows Server 2003 family CD (select the **Complete** setup option). Although you can still use this version of Netdiag.exe to obtain basic networking information that is not IPSec-specific, the Netsh commands for IPSec replace all IPSec-specific functionality.
- In Windows XP, run the Setup.exe file from the \Support\Tools folder on the Windows XP CD (select the **Complete** setup option).

- For Windows 2000, an updated version of Netdiag.exe is available for download from the Web. For more information, see "Windows 2000 Resource Kits," at the Microsoft Windows Resource Kits Web site.

For more information about Netdiag.exe, see Using Netdiag.exe to display IPSec information and to test and view network configuration in this topic.

**Viewing IPSec policy assignment information**

Viewing the name of an active IPSec policy and the name of the Group Policy object to which the active IPSec policy is assigned can be useful for troubleshooting policy precedence issues. The following table summarizes the tools that you can use to view this information for the Windows Server 2003 family, Windows XP, and Windows 2000.

| On computers running | Use these tools to view the name of the active IPSec policy | Use these tools to view the name of the Group Policy object to which the active IPSec policy is assigned |
|---|---|---|
| Windows Server 2003 family | • IP Security Monitor console.<br><br>For more information, see View details about active IPSec policies.<br><br>• Netsh, **netsh ipsec static show gpoassignedpolicy** command.<br><br>For more information, see Netsh commands for Internet Protocol security. | • Resultant Set of Policy (RSoP) console.<br><br>For more information, see Use Resultant Set of Policy (RSoP) to View IPSec Policy Assignments.<br><br>• Netsh, **netsh ipsec static show gpoassignedpolicy** command. |
| Windows XP | • IP Security Policy Management console (for local policies only).<br><br>• Network Connectivity Tester (Netdiag.exe), **netdiag /test:ipsec** command.<br><br>For more information, see Using Netdiag.exe to display IPSec information and to test and view network configuration in this topic. | • Netdiag.exe, **netdiag /test:ipsec** command. |
| Windows 2000 | • Netdiag.exe, **netdiag /test:ipsec** command.<br><br>• In the properties for the relevant network connection, **TCP/IP Properties/Advanced/Options/IPSec**.<br><br>The assigned IPSec policy that is displayed in **TCP/IP Properties** is global. It is not specific to the connection. | • Netdiag.exe, **netdiag /test:ipsec** command.<br><br>• Gpresult.exe (Group Policy Results).<br><br>• Gpotool.exe (Group Policy Verification Tool).<br><br>Gpresult.exe and Gpotool.exe are both available for download from the Web. For more information, see "Windows 2000 Resource Kits," at the Microsoft Windows Resource Kits Web site. |

**Notes**

- To view all IPSec policies that are available (but not necessarily assigned or applied) to computers, use the IP Security Policy Management console. For information about policy precedence and IPSec policy behavior in an Active Directory environment, see the section "Active Directory-based policy" in Creating, modifying, and assigning IPSec policies.

- For the Windows Server 2003 family, to determine which IPSec policies are assigned but are not being applied to IPSec clients, use the RSoP console.

- The Windows XP implementation of the RSoP console does not support the display of IPSec policies. In addition, the **gpresult /scope computer** command does not display the Group Policy object that contains an IPSec policy assignment. For these reasons, you should use Netdiag.exe to view IPSec policy assignment information on computers running Windows 2000 or Windows XP. The **netdiag /test:ipsec** command displays the Group Policy object that contains the IPSec policy assignment, and the organizational unit to which the Group Policy object is assigned.

- Gpotool.exe allows you to monitor the health of Group Policy objects on domain controllers. You can use this tool to check the consistency and replication of Group Policy objects and to display Group Policy object properties. Gpotool.exe is available only for computers running Windows 2000.

**Viewing details about the active IPSec policy and IPSec statistics**

The IP Security Policy Management console, Netdiag.exe, Gpresult.exe, and Gpotool.exe allow you to determine which IPSec policy has been assigned through Group Policy. After you verify this information, you might need to view details about the assigned IPSec policy and IPSec statistics (for example, filters, filter actions, and active security associations). The following table summarizes the tools that you can use to view an active IPSec policy and IPSec statistics for the Windows Server 2003 family, Windows XP, and Windows 2000.

| **On computers running** | **Use these tools to display the active IPSec policy, IPSec statistics, or both** | **Notes** |
|---|---|---|
| Windows Server 2003 family | • IP Security Monitor console.<br><br>For more information, see View IP security statistics<br><br>• Netsh, **netsh ipsec dynamic show all** command. | • To view details about the active IPSec policy and IPSec statistics on a local or remote computer, you must be a member of the Administrators group on that computer. |
| Windows XP | • IP Security Monitor console.<br><br>• IPseccmd.exe, **ipseccmd show all** command. | • To view details about the active IPSec policy and IPSec statistics for a local or remote computer, you must be a member of the Administrators group on that computer. |
| Windows 2000 | • Netdiag.exe, **netdiag /test:ipsec /v /debug** command.<br><br>• Ipsecmon.exe. | • To use the **/debug** option, you must be logged on as a member of the Administrators group on that computer. In addition, to view details about Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory.<br><br>• Ipsecmon.exe only displays active outbound quick mode security associations (SAs). |

**Note**

- To remotely monitor IPSec on a computer that is running a different version of Windows than your computer, use Remote Desktop Connection. For example, if your computer is running the Windows Server 2003 family and you plan to

remotely monitor IPSec on computers running Windows 2000 or Windows XP, use Remote Desktop Connection to gain remote access to these computers. For information about Remote Desktop Connection, see Remote Desktop Connection.

**Viewing details about the active IPSec policy and IPSec statistics in IP Security Monitor**

In Windows XP and the Windows Server 2003 family, IP Security Monitor is implemented as a Microsoft Management Console (MMC) snap-in, and it includes enhancements that allow you to view details about an active IPSec policy that is applied by the domain or locally, as well as quick mode and main mode statistics, and active IPSec SAs. IP Security Monitor also enables you to search for specific main mode or quick mode filters. To troubleshoot complex IPSec policy designs, you can use IP Security Monitor to search for all matches for filters of a specific traffic type.

- For procedures on how to use IP Security Monitor, see Monitor IPSec Activity.
- For detailed information about IP security statistics, see Viewing main mode and quick mode statistics in IP Security Monitor.

### Notes

- You can use IP Security Monitor to monitor only computers that are running Windows XP or the Windows Server 2003 family. In addition, the computer that is being monitored must run the same version of the Windows operating system as the computer on which IP Security Monitor is running. To monitor IPSec on a computer that is running Windows 2000, use the **ipsecmon** command at the Windows 2000 command prompt on the computer that is being monitored.
- For remote monitoring, you can use IP Security Monitor only to monitor computers that are running the same version of the Windows operating system. To remotely monitor IPSec on a computer that is running a different version of Windows than your computer, use Remote Desktop Connection. For information about Remote Desktop Connection, see Remote Desktop Connection.
- If your computer is running the Windows Server 2003 family, and you plan to monitor IPSec on computers that are also running the Windows Server 2003 family, you can run IP Security Monitor or use the Netsh command-line tool remotely.

**Verifying that security auditing is enabled**

You can use Local Security Policy settings (for a local computer) or Group Policy Object Editor (for a domain) to verify that security auditing is enabled, so you can ensure that the success and failure of IKE negotiations is recorded. Auditing for IKE is supported in Windows 2000, Windows XP, and the Windows Server 2003 family (IKE uses the Logon Events category). In the Windows Server 2003 family, you can also enable auditing for the security policy database (SPD). SPD uses the Policy Change category. For more information, see Define or modify auditing policy settings for an event category.

**Viewing IPSec-related events in Event Viewer**

You can use Event Viewer to view the following IPSec-related events:

**IKE events (negotiation success and failure) in the security log.**

To view these events, enable success or failure auditing for the **Audit logon events** audit policy for your domain or local computer. The IKE event category is also used for auditing user logon events in services other than IPSec. For more information, see Define or modify auditing policy settings for an event category.

When you enable success or failure auditing for the **Audit logon events** audit policy, IPSec records the success or failure of each main mode and quick mode negotiation and the establishment and termination of each negotiation as separate events. However, enabling this type of auditing can cause the security log to fill with IKE events. For example, for servers that are connected to the Internet, attacks on the IKE protocol can cause the security log to fill with IKE events. IKE events can also fill the security log for servers that use IPSec to secure traffic to many clients. To avoid this, you can disable auditing for IKE events in the security log by modifying the registry.

- To disable auditing of IKE events in the security log, do the following:

  1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Audit\DisableIKEAudits** registry setting to a value of **1**.

     The **DisableIKEAudits** key does not exist by default and must be created.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

1. Restart the computer, or stop and then restart the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

**Note**

- Stopping and restarting the IPSec service can disconnect all of the computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. If you restart the IPSec service immediately, TCP-based communication might resume, due to the retransmit behavior of TCP, after new IKE and IPSec SAs are established. For more information, see "Stopping and restarting the IPSec service," later in this topic.

**IPSec policy change events in the Security log.**

To view these events, enable success or failure auditing for the **Audit policy change** audit policy for your domain or local computer. For more information, see Define or modify auditing policy settings for an event category.

**IPSec driver per-packet drop events in the System log.**

In Windows 2000, Windows XP, and the Windows Server 2003 family, you can enable packet event logging for the IPSec driver by modifying the registry. The IPSec driver reads the registry during computer startup. In the Windows Server 2003 family, you can enable packet event logging for the IPSec driver by using the **Netsh ipsec** command-line tool. To enable logging of dropped inbound and outbound packets, specify a value of **7**. For information about the other levels of IPSec driver event logging, see Notes.

- To enable IPSec driver logging of dropped inbound and outbound packets by modifying the registry, do the following:

  1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSec\EnableDiagnostics DWORD** registry setting to a value of **7**.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

  1. Restart the computer.

- To enable IPSec driver logging of dropped inbound and outbound packets by using the Netsh IPSec command-line tool in the Windows Server 2003 family, do the following:

  1. At the command prompt, type: **netsh ipsec dynamic set config ipsecdiagnostics 7**

  2. Restart the computer.

You can also change the interval for writing IPSec driver packet events to the System log. By default, the IPSec driver writes events to the System log once an hour or after a threshold for the number of events has been reached. For troubleshooting, you should set this interval to the minimum value, 60 seconds.

- To change the interval for writing IPSec driver packet events by modifying the registry, do the following:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSec\LogInterval DWORD** registry setting to **60 decimal**.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

1. Restart the computer.

- To change the interval for writing IPSec driver packet events by using the **Netsh ipsec** command-line tool in the Windows Server 2003 family, do the following:

  1. At the command prompt, type **netsh ipsec dynamic set config ipseclogginterval 60**

  2. Restart the computer.

The IPSec driver reads the registry during computer startup.

By default, packet event logging for the IPSec driver is disabled (that is, the value is set to **0**). When you enable packet event logging for the IPSec driver, you can specify any of the following values, to enable different levels of logging:

| Logging Level | Description |
|---|---|
| **1** | When **1** is specified, bad SPI packets (the total number of packets for which the Security Parameters Index or SPI was incorrect), IKE negotiation failures, IPSec processing failures, packets received with invalid packet syntax, and other errors are recorded in the System log. Unauthenticated hashes (with the exception of the "Clear text received when should have been secured" event) are logged as well. |
| **2** | When **2** is specified, inbound per-packet drop events are recorded in the System log. |
| **3** | When **3** is specified, level 1 and level 2 logging are performed. In addition, unexpected clear text events (packets that are sent or received in plaintext) are also recorded. |
| **4** | When **4** is specified, outbound per-packet drop events are recorded in the System log. |
| **5** | When **5** is specified, level 1 and level 4 logging are performed. |
| **6** | When **6**is specified, level 2 and level 4 logging are performed. |
| **7** | When **7** is specified, all levels of logging are performed. |

For more information about IPSec driver event logging, see the Microsoft Windows Resource Kits Web site.

You cannot configure audit policies on a computer running Windows XP Home Edition. However, success and failure auditing for the **Audit logon events** and **Audit policy change** audit policies for the local computer are enabled by default.

Enabling auditing in the security log or IPSec driver diagnostics in the system log can cause these logs to fill with events quickly. Before you perform either of these tasks, you should do the following:

- Ensure that the size of the log is at least 10 megabytes (MB).

- Save the existing log to a file.

- Clear all events in the log so that the log is empty.

- Consider disabling auditing of IKE events in the security log by modifying the registry, if your computer hosts many simultaneous network connections.

**Enabling detailed tracing for Internet Key Exchange (IKE) negotiations**

Enabling audit logging for IKE events and viewing the events in Event Viewer is the fastest and simplest way to troubleshoot failed main mode or quick mode negotiations. However, some scenarios might require a more detailed analysis of the IKE main mode negotiation and quick mode negotiations for troubleshooting. You can enable tracing for IKE negotiations if the audit failure events do not provide enough information. The IKE tracing log is a very detailed log intended for troubleshooting IKE interoperability under controlled circumstances. Expert knowledge of the ISAKMP RFC 2408 and IKE RFC 2409 is required to interpret this log.

The IKE tracing log appears as the *systemroot*\Debug\Oakley.log file. The log has a fixed size of 50,000 lines and will overwrite as necessary. A new Oakley.log file is created each time the IPSec service is started and the previous version of the Oakley.log file is saved as Oakley.log.sav. When the Oakley.log file becomes full, the current file is saved as Oakley.log.bak, and a new Oakley.log file is created.

Because many IKE negotiations can occur simultaneously, you should minimize the number of negotiations and log for as short a period of time as possible to capture a more easily interpreted log.

**Enabling and disabling the IKE tracing log in the** Windows Server 2003 **family**

In the Windows Server 2003 family, you can enable or disable the IKE tracing log dynamically while the IPSec service is running by doing the following:

- To enable the IKE tracing log, type the following at the command prompt:

  **netsh ipsec dynamic set config ikelogging 1**

  This command creates the IKE tracing log file if it does not exist. If the file does exist, it appends logging information to the existing file.

- To disable the IKE tracing log, type the following at the command prompt:

  **netsh ipsec dynamic set config ikelogging 0**

**Enabling the IKE tracing log in Windows 2000 and** Windows XP

In Windows 2000 and Windows XP, you must enable IKE tracing by modifying the registry. For the changes to take effect, you must also stop and restart the IPSec service:

- To enable the IKE tracing log in Windows XP and Windows 2000, do the following:

  1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLoggingDWORD** registry setting to a value of **1**.

     The Oakley key does not exist by default and must be created.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

  1. Stop and start the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

**Stopping and restarting the IPSec service**

Stopping and restarting the IPSec service can disconnect all of the computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. However, if you restart the IPSec service immediately, TCP-based communication might resume, due to the retransmit behavior of TCP, after new IKE and IPSec SAs are established. If you leave the IPSec service stopped, any clients that used the default response rule to establish security will be unable to communicate with this computer for two hours.

To avoid losing Terminal Services connectivity for computers that are using IPSec over a Terminal Server session, you must stop and restart the IPSec service by using a single command line:

- To stop and restart the IPSec service for computers over a Terminal Server session

    - At the command prompt, type the following:

      **net stop policyagent & net start policyagent**

If you are restarting the IPSec service while an L2TP/IPSec VPN tunnel is connected, the tunnel will lose connectivity and must be reconnected when the IPSec service is restarted. A PPTP tunnel, however, does not use IPSec and therefore will stay connected if you stop the IPSec service.

If you are restarting the IPSec service on a computer that is running the Windows Server 2003 family or the Windows Server 2003 family and that is also running the Routing and Remote Access service, any IPSec configuration for L2TP will be lost and the L2TP tunnels will be disconnected. Therefore, you must stop and restart the Routing and Remote Access service, as well as the IPSec service.

- To stop and restart the IPSec service and the Routing and Remote Access service, do the following:

    1. Stop the Routing and Remote Access service using the **net stop remoteaccess** command.
    2. Stop the IPSec service by using the **net stop policyagent** command.
    3. Start the IPSec service by using the **net start policyagent** command.
    4. Start the Routing and Remote Access service using the **net start remoteaccess** command.

**Viewing IPSec and other network communication with Network Monitor**

You can install and use Network Monitor to view IPSec and other network communication. Note that the version of Network Monitor that is provided with the Windows Server 2003 family can be used only to view the network traffic that is sent to or from the computer on which it is installed. To view network traffic that is sent to or from another computer and is routed through your computer (using the Routing and Remote Access service), you must use the Network Monitor component that is provided with Microsoft Systems Management Server.

The Network Monitor component that is provided with the Windows Server 2003 family includes parsers for the ISAKMP (IKE), AH, and ESP protocols. The Network Monitor parsers for ESP can parse inside the ESP packet only if null-encryption is being used and the full ESP packet is captured. Network Monitor cannot parse the encrypted portions of IPSec-secured ESP traffic when encryption is performed in software. However, if encryption is being performed by an IPSec hardware offload network adapter, the ESP packets are decrypted when Network Monitor captures them and as a result, can be parsed and interpreted into the upper-layer protocols. If you need to diagnose ESP software-encrypted communication, you must disable ESP encryption and use ESP-null encryption by changing the IPSec policy on both computers.

### Notes

- For information about how to install the version of Network Monitor that is provided with the Windows Server 2003 family, see Install Network Monitor.

- For information about Systems Management Server, go to the Microsoft Web site.

**Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family**

For computers running the Windows Server 2003 family, you can use the Netsh commands for IPSec to script IPSec policy creation, display details about IPSec policies, and change the IPSec configuration for troubleshooting.

For troubleshooting, you can use the Netsh commands for IPSec to change the following settings:

#### Disable Internet Key Exchange (IKE) certificate revocation list (CRL) checking

By default, in Windows 2000 CRLs are not checked during IKE certificate authentication. In Windows XP and the Windows Server 2003 family, CRLs are checked during IKE certificate authentication, but a fully successful check is not required for the certificate to be accepted. In some cases, failures during CRL processing might cause IKE to not accept the certificate. Or, the delay required for CRL checking might delay IKE negotiation enough to cause the connection attempt to time-out. To determine whether certificate authentication will be successful without CRL checking, you can disable IKE CRL checking. To do this, type the following at the command prompt:

**netsh ipsec dynamic set config strongcrlcheck 0**

This setting takes effect immediately. It does not require a restart.

#### Enable IPSec driver event logging

To record all inbound and outbound dropped packets and other packet processing errors in the Event Viewer System log, you can set the IPSec driver event logging level to **7**. To do this, type the following at the command prompt:

**netsh ipsec dynamic set config ipsecdiagnostics 7**

This setting will not take effect until you restart your computer.

### Notes

- If you configured an IPSec filter action on a computer to block traffic, setting the IPSec driver event logging level to **7** on that computer might generate many events that fill the System log very quickly. Accordingly, you should set the IPSec driver event logging level to **7** only for testing.

- All packets that are dropped by IPSec contribute to the **Datagrams Received Discarded** or **Datagrams Outbound Discarded** System Monitor counters for the IP object. IPSec itself does not provide a counter object. You can monitor IPSec performance counters manually by using IP Security Monitor.

#### Permit inbound and outbound traffic during computer startup

To permit inbound and outbound traffic during computer startup (before the IPSec service starts), you can use the following Netsh command:

**netsh ipsec dynamic set config bootmode permit**

This setting will not take effect until you restart your computer.

### Note

- If you start the computer in Safe Mode with Networking, the IPSec service will not start and cannot retrieve IPSec policy settings from Active Directory or the local registry, and persistent policy cannot be applied. As a result, inbound and outbound traffic will continue to be permitted until the computer is no longer in Safe Mode.

#### Exempt all broadcast, multicast, IKE, Kerberos, and RSVP traffic from IPSec filtering

In Windows 2000 and Windows XP, by default, the IPSec driver exempts all broadcast, multicast, IKE, Kerberos, and RSVP traffic from IPSec filtering. In the Windows Server 2003 family, only IKE traffic is exempt from IPSec filtering, and you can configure, block, or permit filter actions specifically for multicast and broadcast traffic (IPSec does not negotiate SAs for multicast and broadcast traffic). To restore the IPSec driver to the default Windows 2000 and Windows XP filtering behavior, you can use the following Netsh command:

**netsh ipsec dynamic set config ipsecexempt 0**

This setting will not take effect until you restart your computer.

### Caution

- The Windows 2000 and Windows XP default exemption settings for IPSec are designed for corporate LAN environments with a low risk of attack. For this reason, you should use only the Windows 2000 and Windows XP default exemption settings when necessary for troubleshooting, in low-risk environments, or when you cannot solve program compatibility issues by configuring explicit filters in IPSec policies.

You can also exempt all broadcast, multicast, IKE, Kerberos, and RSVP traffic from IPSec filtering by modifying the registry as follows:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSEC\NoDefaultExempt DWORD** registry setting to a value of **0**.

2. The **NoDefaultExempt** key does not exist by default and must be created.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer

3. Restart the computer.

#### Using Ipseccmd.exe to manage and monitor IPSec on computers running Windows XP

For computers running Windows XP, you can use the Ipseccmd.exe command-line tool to script IPSec policy creation and to display IPSec policy assignments, active SAs, and detailed IPSec policy settings, including filters, filter actions, and authentication methods. Ipseccmd.exe replaces and enhances the functionality provided by the Ipsecpol.exe (used in Windows 2000 for scripting IPSec policy creation) and the Netdiag.exe **/test:ipsec /v /debug** command (used in Windows 2000 to display configuration information and SAs for an active IPSec policy).

To install Ipseccmd.exe on a computer running Windows XP, run the Setup.exe file from the \Support\Tools folder on the Windows XP CD.

**Notes**

- Ipseccmd.exe is provided only with Windows XP, and you can use this tool to manage and monitor IPSec policies only on computers running Windows XP.

- To display all IPSec policy settings and statistics for diagnostics, use the **ipseccmd show all** command.

#### Using Netdiag.exe to display IPSec information and to test and view network configuration

You can use the Netdiag.exe command-line tool to display information about IPSec policies and statistics, report network configuration, and test basic networking capabilities and domain-based functionality. The following table describes how to install Netdiag.exe and provides information about functional changes relevant to IPSec for each version of this tool.

| On computers running | Use this method to install | Notes |
|---|---|---|
| Windows Server 2003 family | For the Windows Server 2003 family, Netdiag.exe is available for installation from the Windows Server 2003 family CD (run the Suptools.msi file from the \Support\Tools folder, and choose the **Complete** setup option. | • Although you can still use Netdiag.exe to obtain basic networking information that is not IPSec-specific, for the Windows Server 2003 family, the **netdiag /test:ipsec** option has been removed, and the Netsh commands for IPSec replace all IPSec-specific functionality. To view details about IPSec policies, use either the **netsh ipsec static show** command or the **netsh ipsec dynamic show** command. |

| | | |
|---|---|---|
| Windows XP | For Windows XP, Netdiag.exe is available for installation from the Windows XP CD (run the Setup.exe file from the \Support\Tools folder, and choose the **Complete** setup option). | • Displays information about IPSec policy assignments, including the name of the active IPSec policy, the name of the Group Policy object that assigned the policy, and the policy path.<br><br>• The display of detailed IPSec policy information and IPSec statistics (provided by the **/debug** and **/v** options in Windows 2000) is not supported in Windows XP. Instead, to view details about IPSec policies, use the Ipseccmd.exe **show** command. |
| Windows 2000 | For Windows 2000, an updated version of Netdiag.exe is available for download from the Web. For information, see "Windows 2000 Resource Kits," at the Microsoft Windows Resource Kits Web site. | • Netdiag.exe was enhanced in Windows 2000 Service Pack 1 to display the number of bytes offloaded by the IPSec driver to a network adapter that is capable of IPSec hardware offload. The hardware offload statistics are Offloaded Bytes Sent and Offloaded Bytes Received.<br><br>• To use the **/debug** option to view details about Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory. |

You can use the **netdiag /v /l** command to obtain networking information that is not IPSec-specific, for any platform. Typical uses of these options include:

- Reporting the IP configuration and routing configuration for a computer with one command.
- Testing WINS and DNS name resolution and consistency.
- Reporting the build version of a computer and the hotfixes that are installed on that computer.
- Testing the validity of domain membership, whether domain members can successfully contact domain controllers, and trust relationships.

The **/l** option generates a Netdiag.log file. This file is written to the folder in which Netdiag.exe is run.

**Notes**

- To run all Netdiag.exe commands, you must be a member of the Administrators group on the local computer.
- Each version of Netdiag.exe is customized to run on a different version of the Windows operating system (that is, Windows 2000, Windows XP, or the Windows Server 2003 family). You can only run Netdiag.exe on a computer that is running the Windows operating system for which Netdiag.exe has been specifically designed.
- If you are running a version of Netdiag.exe that is designed for Windows 2000 or Windows XP in a Windows Server 2003 domain, Netdiag.exe might not report the correct operating system version of the domain.
- If you run Netdiag.exe on a computer that is assigned an IPSec policy that affects most inbound and outbound IP traffic on that computer, Netdiag.exe might not report results correctly the first time it is run. This is because it might take a few seconds for IPSec to establish SAs to the many remote computers for which Netdiag.exe tests connectivity. If this occurs, wait five seconds, and then run Netdiag.exe again.

**Disabling TCP/IP and IPSec hardware acceleration**

Network adapters can accelerate IPSec processing by performing hardware offload of IPSec cryptographic functions (IPSec offload) and the calculation of TCP checksums (checksum offload). Such adapters might also process large TCP segments for very fast transmission (large-send offload). By default, when a network adapter driver that can perform hardware offload is enabled during the Plug and Play initialization process, the driver advertises this capability to TCP/IP and IPSec. TCP/IP and

IPSec then offload tasks to the network adapter driver as appropriate.

If you use network adapters to perform hardware offload of IPSec cryptographic functions, you might need to verify that hardware acceleration is not causing problems with the packet processing that is performed by the network adapter. To do so, you can:

- Disable hardware offload functions in the network adapter driver (if this function is supported by the driver). This task does not require you to restart the computer.
- Disable TCP/IP and IPSec hardware acceleration. This task requires you to restart the computer.
- Disable only IPSec hardware acceleration. This task requires you to restart the computer.

**Disabling hardware offload functions in the network adapter driver**

To determine whether you can disable hardware offload functions in the network adapter driver, do one of the following:

- Open Device Manager, click **Network Adapters**, right-click the network adapter whose properties you want to view, and then click **Properties**. For more information, see Change network adapter settings
- Open Network Connections, right-click the network adapter that you want to view, click **Properties**, and then, on the **General** tab, click **Configure**.

**Disabling TCP/IP and IPSec hardware acceleration**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

To disable both TCP/IP and IPSec hardware acceleration by modifying the registry, do the following:

1. Set the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\DisableTaskOffload** DWORD registry setting to a value of **1**.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

2. Restart the computer.

If the **DisableTaskOffload** DWORD registry setting is set to a value of **0**, and if the network adapter advertises hardware offload capabilities, IPSec and TCP/IP will attempt to offload the appropriate functions to the network adapter.

**Disable only IPSec hardware acceleration**

To disable only IPSec hardware acceleration by modifying the registry, do the following:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSEC\EnableOffload** DWORD registry setting to a value of **0**.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

2. Restart the computer.

**Note**

- To help minimize the potential for hardware acceleration to cause problems with packet processing performed by the network adapter, make sure that you use the latest network adapter driver available from the vendor.

# Viewing IPSec Statistics and Details About Active IPSec Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing IPSec statistics and details about active IPSec policies**
This section covers:

- Viewing main mode and quick mode statistics in IP Security Monitor
- Viewing details about active IPSec policies in IP Security Monitor
- Using Resultant Set of Policy to view IPSec policy assignments

# Viewing main mode and quick mode statistics in IP Security Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing main mode and quick mode statistics in IP Security Monitor**

**Main Mode (IKE) Statistics**
The following table describes the items in the list of main mode (IKE) statistics in IP Security Monitor.

For information about how to view these statistics, see View IP security statistics.

| Main Mode (IKE) Statistic | Description |
|---|---|
| Active Acquire | The number of pending requests to initiate an Internet Key Exchange (IKE) negotiation in order to establish a security association (SA) between IPSec peers. The Active Acquire statistic includes the outstanding request and the number of any queued requests. Under a heavy load, the number of active acquires is 1 plus the number of requests that are queued by IKE for processing. |
| Active Receive | The number of IKE messages received that are queued for processing. |
| Acquire Failures | The total number of acquire outbound requests that have failed since the IPSec service was last started. Acquires are requests to establish SAs between IPSec peers. |
| Receive Failures | The total number of errors that have occurred during the process of receiving IKE messages since the IPSec service was last started. |
| Send Failures | The total number of errors that have occurred during the process of sending IKE messages since the IPSec service was last started. The number of Send Failures typically increases for computers that establish SAs over temporary network connections, such as dial-up connections, virtual private network tunnels, and wireless connections. |
| Acquire Heap Size | The number of entries in the acquire heap. The acquire heap stores successful acquires. Acquires are outbound requests to establish SAs between IPSec peers. |
| Receive Heap Size | The number of entries in the IKE receive buffers. The receive buffers store incoming IKE messages. |
| Authentication Failures | The total number of identity authentication (Kerberos, certificate, and preshared key) failures that have occurred during main mode negotiation since the IPSec service was last started. If you are having difficulty communicating securely, attempt the communication and determine whether the number of Authentication Failures increases. If it does, check your authentication settings for either an unmatched authentication method or an incorrect authentication method configuration (for example, the use of preshared keys that do not match). |
| Negotiation Failures | The total number of negotiation failures that have occurred during main mode or quick mode negotiation since the IPSec service was last started. If you are having difficulty communicating securely, attempt the communication and determine whether the number of Negotiation Failures increases. If it does, check your authentication and security method settings for an unmatched authentication method, an incorrect authentication method configuration (for example, the use of preshared keys that do not match), or unmatched security methods or settings. |
| Invalid Cookies Received | The total number of cookies that could not be matched with an active main mode SA since the IPSec service was last started. A cookie is a value contained in a received IKE message that is used to help identify the corresponding main mode SA. |
| Total Acquire | The total number of requests that have been submitted to IKE since the IPSec service was last started to establish an SA. This number includes acquires that result in soft SAs. |

| | |
|---|---|
| Total Get SPI | The total number of requests that have been submitted by IKE to the IPSec driver to obtain a unique Security Parameters Index (SPI) since the IPSec service was last started. The SPI matches inbound packets with SAs. |
| Key Additions | The total number of outbound quick mode SAs that have been added by IKE to the IPSec driver since the IPSec service was last started. |
| Key Updates | The total number of inbound quick mode SAs that have been added by IKE to the IPSec driver since the IPSec service was last started. |
| Get SPI Failures | The total number of failed requests that have been submitted by IKE to the IPSec driver to obtain a unique SPI since the IPSec service was last started. |
| Key Addition Failures | The total number of failed outbound quick mode SA addition requests that have been submitted by IKE to the IPSec driver since the IPSec service was last started. |
| Key Update Failures | The total number of failed inbound quick mode SA addition requests that have been submitted by IKE to the IPSec driver since the IPSec service was last started. |
| ISADB List Size | The number of main mode state entries. This number includes successfully negotiated main modes, main mode negotiations in progress, and main mode negotiations that failed or expired and have not yet been deleted. |
| Connection List Size | The number of quick mode negotiations that are in progress. |
| IKE Main Mode | The total number of successful SAs that have been created during main mode negotiations since the IPSec service was last started. |
| IKE Quick Mode | The total number of successful SAs that have been created during quick mode negotiations since the IPSec service was last started. |
| Soft Associations | The total number of SAs formed with computers that have not responded to main mode negotiation attempts since the IPSec service was last started. Although these computers did not respond to main mode negotiation attempts, IPSec policy allowed communications with the computers. Soft SAs are not secured by IPSec. |
| Invalid Packets Received | The total number of invalid IKE messages that have been received since the IPSec service was last started. This number includes IKE messages with invalid header fields, incorrect payload lengths, and incorrect values for the responder cookie. Invalid IKE messages are commonly caused by retransmitted IKE messages or an unmatched preshared key between the IPSec peers. |

**Quick Mode (IPSec) Statistics**

The following table describes the items in the list of quick mode (IPSec) statistics in IP Security Monitor.

| **Quick Mode (IPSec) Statistic** | **Description** |
|---|---|
| Active Security Associations | The number of active quick mode SAs. |
| Offloaded Security Associations | The number of active quick mode SAs offloaded to hardware. Certain network adapters can accelerate IPSec processing by performing hardware offload of IPSec cryptographic functions. |
| Pending Key Operations | The number of IPSec key exchange operations that are in progress but are not yet completed. |
| Key Additions | The total number of keys for quick mode SA negotiations that have been successfully added since the computer was last started. |
| Key Deletions | The total number of keys for quick mode SAs that have been successfully deleted since the computer was last started. |

| | |
|---|---|
| Rekeys | The total number of successful rekey operations for quick mode SAs since the computer was last started. |
| Active Tunnels | The number of active IPSec tunnels. |
| Bad SPI Packets | The total number of packets for which the SPI has been incorrect since the computer was last started. If the SPI is incorrect, it might mean that the inbound SA has expired and a packet using the old SPI has recently arrived. This number is likely to increase if rekey intervals are short and there are a large number of SAs. A large number of packets with bad SPIs that are received within a short amount of time might indicate a packet spoofing attack. |
| Packets Not Decrypted | The total number of packets that could not be decrypted since the computer was last started. A packet might not be decrypted if it fails a validation check. |
| Packets Not Authenticated | The total number of packets for which data could not be verified (for which the integrity hash verification failed) since the computer was last started. Increases in this number might indicate an IPSec packet spoofing or modification attack or packet corruption by network devices. |
| Packets With Replay Detection | The total number of packets that have contained an invalid sequence number since the computer was last started. Increases in this number might indicate a network problem or replay attack. |
| Confidential Bytes Sent | The total number of bytes that have been sent using the Encapsulating Security Payload (ESP) protocol (excluding non-encrypted ESP) since the computer was last started. |
| Confidential Bytes Received | The total number of bytes that have been received using the ESP protocol (excluding non-encrypted ESP) since the computer was last started. |
| Authenticated Bytes Sent | The total number of authenticated bytes that have been sent using the Authentication Header (AH) protocol or the ESP protocol since the computer was last started. |
| Authenticated Bytes Received | The total number of authenticated bytes that have been received using the AH protocol or the ESP protocol since the computer was last started. |
| Transport Bytes Sent | The total number of bytes that have been sent using IPSec transport mode since the computer was last started. |
| Transport Bytes Received | The total number of bytes that have been received using IPSec transport mode since the computer was last started. |
| Bytes Sent in Tunnels | The total number of bytes that have been sent using IPSec tunnel mode since the computer was last started. |
| Bytes Received in Tunnels | The total number of bytes that have been received using IPSec tunnel mode since the computer was last started. |
| Offloaded Bytes Sent | The total number of bytes that have been sent using IPSec hardware offload since the computer was last started. |
| Offloaded Bytes Received | The total number of bytes that have been received using IPSec hardware offload since the computer was last started. |

# Viewing details about active IPSec policies in IP Security Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing details about active IPSec policies in IP Security Monitor**

IP Security Monitor allows you to view details about the active IPSec policies that are applied to the domain, the local computer, or a remote computer. This topic describes the type of details that IP Security Monitor displays.

For information about how to view details about active IPSec policies, see View details about active IPSec policies. For information about creating, assigning, and assigning IPSec policies, see Creating, modifying, and assigning IPSec policies.

**Active IPSec policies that are applied by the domain**

The following table describes the items that are displayed when you use IP Security Monitor to view details about an active IPSec policy that is applied by the domain.

| Active IPSec policy item | Description |
|---|---|
| **Policy Name** | Specifies the name of the active IPSec policy. |
| **Description** | Optional text that describes the purpose of the active IPSec policy. |
| **Policy Last Modified** | Not applicable (this information is displayed only for active IPSec policies that are applied to a computer). |
| **Policy Store** | Specifies the storage location for the active IPSec policy. For active IPSec policies that are applied by the domain, this location appears as "Domain Store." |
| **Policy Path** | Specifies the LDAP path to the active IPSec policy. |
| **Organizational Unit** | Specifies the organizational unit to which the Group Policy object is applied. |
| **Group Policy Object Name** | Specifies the Group Policy object to which the active IPSec policy is applied. |

**Active IPSec policies that are applied to a computer**

The following table describes the items that are displayed when you use IP Security Monitor to view details about active IPSec policies that are applied to a computer.

| Active IPSec policy item | Description |
|---|---|
| **Policy Name** | Specifies the name of the active IPSec policy. |
| **Description** | Optional text that describes the purpose of the active IPSec policy. |
| **Policy Last Modified** | Specifies the date and the time that the active IPSec policy was last modified. |
| **Policy Store** | Specifies the storage location of the active IPSec policy. For active IPSec policies that are applied to a computer, this location appears as "Local Store." |
| **Policy Path** | Not applicable (this information is only displayed for active IPSec policies that are applied by the domain). |
| **Organizational Unit** | Not applicable (this information is only displayed for active IPSec policies that are applied by the domain). |

| | |
|---|---|
| **Group Policy Object Name** | Not applicable (this information is only displayed for active IPSec policies that are applied to the domain). |

**Note**

- You cannot use IP Security Monitor to add, edit, or remove IPSec policies. For information about how to add, edit, or remove IPSec policies, see Add, edit, or remove IPSec policies.

# Using Resultant Set of Policy to view IPSec policy assignments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Resultant Set of Policy to view IPSec policy assignments**
Resultant Set of Policy (RSoP) is an addition to Group Policy that you can use to view IPSec policy assignments for a computer or for members of a Group Policy container. This information can help you troubleshoot policy precedence issues and plan your deployment.

To view IPSec policy assignments in RSoP, you must first open the RSoP MMC console, and then run a query. RSoP provides two types of queries: logging mode queries (for viewing IPSec policy assignments for a computer) and planning mode queries (for viewing IPSec policy assignments for members of a Group Policy container).

**Logging mode queries**
You can run an RSoP logging mode query to view all of the IPSec policies that are assigned to an IPSec client. The query results display the precedence of each IPSec policy assignment, so that you can quickly determine which IPSec policies are assigned but are not being applied and which IPSec policy is being applied. The RSoP console also displays detailed settings (that is, the filter rules, filter actions, authentication methods, tunnel endpoints, and connection type) for the IPSec policy that is being applied.

When you run a logging mode query, RSoP retrieves policy information from the Windows Management Instrumentation (WMI) repository on the target computer, and then displays this information in the RSoP console. In this way, RSoP provides a view of the policy settings that are being applied to a computer at a given time.

For information about how to run RSoP logging mode queries for IPSec policies, see Use RSoP to view IPSec policy assignments for a computer. For general information about RSoP logging mode queries, see RSoP logging mode. For general information about RSoP, see RSoP overview.

**Planning mode queries**
You can run an RSoP planning mode query to view all of the IPSec policies that are assigned to members of a Group Policy container. For example, a planning mode query can be useful if you are planning a company reorganization and you want to move computers from one organizational unit to a new organizational unit. By supplying the appropriate information and then running a planning mode query, you can determine which IPSec policies are assigned but are not being applied to the new organizational unit and which IPSec policy is being applied. In this way, you can identify which policy would be applied if you were to move the computers to the new organizational unit. As with logging mode queries, when you run a planning mode query, the RSoP console displays detailed policy settings for the IPSec policy that is being applied.

When you run a planning mode query, RSoP retrieves the names of the target user, computer, and domain controller from the WMI repository on the domain controller. WMI then uses the Group Policy Data Access Service (GPDAS) to create the policy settings that would be applied to the target computer, based on the RSoP query settings that you entered. RSoP reads the policy settings from the WMI repository on the domain controller, and then displays this information in the RSoP console user interface.

For information about how to run RSoP planning mode queries for IPSec policies, see Use RSoP to view IPSec policy assignments for members of a Group Policy container. For general information about RSoP planning mode queries, see RSoP planning mode.

### Note

- You can run an RSoP planning mode query only on a domain controller (when you run a planning mode query, you must explicitly specify the domain controller name). However, you can specify any IPSec client as the target for the query, provided you have the appropriate permissions to do so.

**Assigning and processing IPSec policy in Group Policy**
IPSec policies can be assigned from and stored in Active Directory, as part of Group Policy, or they can be assigned and stored locally, on a computer. When a computer is joined to an Active Directory domain, the domain-level IPSec policy applies. If a computer is not joined to an Active Directory domain, the local IPSec policy applies.

Group Policy settings are contained in Group Policy objects, which are linked with specific Active Directory objects (sites, domains, and organizational units). When an IPSec policy is assigned to a Group Policy object for an Active Directory object (such as an organizational unit), that IPSec policy is propagated to any computer accounts that are affected by the Group Policy

object.

Multiple Group Policy objects, each of which can contain an IPSec policy, can be assigned to a computer account. When multiple IPSec policies are assigned, the last policy that is processed is the policy that is applied (that is, the last policy takes the highest precedence and overrides the settings of any IPSec policy assignments that were processed earlier).

Policy precedence is based on the Group Policy inheritance model. The policy used is the policy assigned at the lowest level of the domain hierarchy for the domain container of which the computer is a member. For example, if there are IPSec policies that are configured for both the domain and for an organizational unit within the domain, the computers that are members of the domain use the domain IPSec policies. The computers that are members of the organizational unit within the domain use the organizational unit IPSec policies. If there are multiple organizational units, members of each organizational unit use the IPSec policy assigned to the organizational unit that is closest in level to their container in the Active Directory hierarchy. If no IPSec policies are configured for Active Directory, or if a computer is not connected to an Active Directory domain, the local IPSec policy is used.

**IPSec policy information displayed in the RSoP console**
The RSoP console simplifies the task of determining which IPSec policy is being applied by displaying the following information for each Group Policy object that contains an IPSec policy assignment: the name of the IPSec policy, the name of the Group Policy object that the IPSec policy is assigned to, the IPSec policy precedence (the lower the number, the higher the precedence), and the name of the site, domain, and organizational unit to which the Group Policy object containing the IPSec policy applies (that is, the scope of management for the Group Policy object).

For information about policy precedence and IPSec policy behavior in an Active Directory environment, see the section "Active Directory-based policy" in Creating, modifying, and assigning IPSec policies. For general information about Group Policy, see Group Policy overview.

### Note

- The settings of the IPSec policy with the highest precedence apply in their entirety; they are not merged with the settings of IPSec policies that are applied at higher levels of the Active Directory hierarchy.

# IPSec Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

**Windows Resource Kits**
The Windows Resource Kits contain extensive technical information on TCP/IP and IPSec. For more information, see Using the Windows Deployment and Resource Kits.

**IPSec-related Web sites**
The following Web sites contain links that relate to IPSec:

- The IPSec Working Group page contains links to RFCs, Internet Drafts, and other IPSec-related links at the Internet Engineering Task Force Web site.
- For information and the latest recommendations on security, see Microsoft Security at the Microsoft Web site.
- For information about the standards and technology for IPSec, as well as other computer security topics, see the National Institute of Standards and Technology Web site.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# IPSec Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Traffic is not being secured.
- IPSec SA negotiation is failing.
- IPSec SA negotiation is failing when using certificate authentication.
- Some computers are unable to communicate with secure computers.
- Local computer IPSec policy is not being used.
- Unable to assign or unassign an IP Security policy.
- DES is being used when 3DES is needed.
- Unsecured traffic is being sent to a secure server.
- Event Viewer is displaying error messages related to IPSec.
- A policy location error message is displayed.
- Event Viewer is displaying "Bad SPI" error message.
- IP Security Monitor displays the error message "The IPSec server is unavailable or incompatible with the IPSec monitor."

**Traffic is not being secured.**
**Cause:** An IPSec security association (SA) that secures traffic (also known as a hard SA) is not being established.

**Solution:** This might be caused by either filter list misconfiguration or the enabling of unsecured communication with non-IPSec-aware computers. Examine the policies that are assigned to both IPSec peers and determine, based on the rules, filter lists, and filter action settings, whether IPSec requires traffic to be secured. Check to determine whether the **Allow unsecured communication with non-IPSec-aware computer** option on the **Security Methods** tab of the appropriate filter action is enabled. Use audit logging to determine the cause of the failed negotiation.

**See also:** IPSec troubleshooting tools; Add, edit, or remove IPSec security methods

**IPSec SA negotiation is failing.**
**Cause:** An IPSec security association (SA) that secures traffic (also known as a hard SA) is not being established. If unsecured communication with non-IPSec-aware computers is enabled, the negotiation failure is expected.

Negotiation failure between two IPSec-aware computers commonly occurs when the default policies are used on two IPSec peers and one or both are not members of an Active Directory domain. Computers running Windows XP Home Edition do not support Active Directory domain membership or the Kerberos V5 authentication method. Default policies use the Kerberos V5 authentication method, which does not work for computers that are not Active Directory domain members.

**Solution:** A negotiation failure between two IPSec-aware computers is most often caused by a mismatched key exchange method, authentication method, or security method. Examine the policies that are assigned to both IPSec peers and determine, based on the advanced settings, rules, filter lists, and filter action settings, whether IPSec should be securing the traffic. If it should, determine whether there is a mismatched key exchange, authentication method, or security method that is preventing a successful negotiation of a hard SA. Use audit logging to determine the cause of the failed negotiation.

**See also:** IPSec troubleshooting tools; Add, edit, or remove IPSec security methods; Define IPSec authentication methods;

Create key exchange security methods

**IPSec SA negotiation is failing when using certificate authentication.**
**Cause:** The correct computer certificates are not installed on the IPSec peers.

**Solution:** Install the correct computer certificates on the IPSec peers.

**Cause:** An incorrect certification authority (CA) is selected in the authentication method.

**Solution:** For the rules of each IPSec peer, select the CA that issued the computer certificate that is installed on the other IPSec peer.

**Cause:** The certificate chain from the issuing CA to the root CA is broken.

**Solution:** Install a computer certificate on each IPSec peer that follows a valid certificate chain up to a trusted root CA.

**Cause:** The computer certificate on one or both of the IPSec peers has expired.

**Solution:** Install a current computer certificate on each IPSec peer.

**See also:** IPSec troubleshooting tools; Authentication methods; Define IPSec authentication methods; Special IPSec considerations; Using Certificates

**Some computers are unable to communicate with secure computers.**
**Cause:** Secure computers are using a lockdown policy that requires client computers to initiate secure communications. The computers that are unable to communicate are not configured with a rule to initiate secure communication.

**Solution:** Either allow unsecured incoming connection requests on the secure servers or configure additional rules for the computers that are unable to communicate so that they can initiate secure communication with the secure computers.

**Cause:** Client computers are not IPSec-capable and secure computers are not allowing unsecured communications.

**Solution:** Either install client computers that are IPSec-capable or reconfigure the secure computers to allow communication with computers that are not IPSec-capable.

**Cause:** Some computers are configured with a policy that has a key exchange, authentication method, or security method that is mismatched with the policy configured for the secure computers.

**Solution:** Examine the policies that are assigned to both IPSec peers and determine, based on the advanced settings, rules, filter lists, and filter action settings, whether IPSec should be securing the traffic. If it should, determine whether there is a mismatched key exchange, authentication method, or security method that is preventing a successful SA negotiation. Use audit logging to determine the cause of the failed negotiation.

**See also:** Add, edit, or remove filter actions; IPSec troubleshooting tools; Add, edit, or remove IPSec security methods; Define IPSec authentication methods; Create key exchange security methods

**Local computer IPSec policy is not being used.**
**Cause:** The computer is a member of an Active Directory domain and is receiving Active Directory-based IPSec policy.

**Solution:** All domain members receive Active Directory-based IPSec policy, which overrides local IPSec policy settings. To prevent this, remove the computer from the domain.

**Cause:** The IPSec service is not started.

**Solution:** Start or restart the IPSec service.

**See also:** Restart the IPSec services

**Unable to assign or unassign an IP Security policy.**
**Cause:** You cannot assign or unassign policies when IP Security Policies are used for Active Directory.

**Solution:** To assign Active Directory-based policies, use IP Security Policies within Group Policy.

**See also:** Assign or unassign IPSec policy in Group Policy

**DES is being used when 3DES is needed.**
**Cause:** The IPSec peer is a computer running Windows 2000 without either the Windows 2000 High Encryption Pack or Windows 2000 Service Pack 2 (or later) installed.

**Solution:** Install either the Windows 2000 High Encryption Pack or Windows 2000 Service Pack 2 (or later) on the computer

running Windows 2000.

**Unsecured traffic is being sent to a secure server.**
**Cause:** The client computer is using the default response rule.

**Solution:** When the client computer is configured with the default response rule, it sends unsecured packets when it either initiates communication or restarts communication after a delay long enough for the quick mode SA and dynamic filter list to time out. To prevent this, configure additional rules for client computers that always initiate secure communication to secure servers.

**See also:** Security information for IPSec

**Event Viewer is displaying error messages related to IPSec.**
**Cause:** There might be a problem with either the IPSec Policy Agent or the Oakley service.

**Solution:** View the system and security event logs for messages that might help in determining the origin of the problem.

**A policy location error message is displayed.**
**Cause:** This error might occur if the IPSec Policy Agent is attempting to retrieve policy from an incorrect storage location.

**Solution:** Examine both the IPSec Policy Agent events in Event Viewer and the IPSec Policy Agent log for messages that indicate from what location policy is being retrieved (the local registry or the domain). If **HKLM\SOFTWARE\Policies\Microsoft\Windows\IPSec\GPTIPSecPolicy\DSIPSecPolicyFlags** is set to 1, then IPSec policy is obtained from Active Directory. Otherwise, it is obtained from the registry. The local computer warns a user that domain policy is in effect when it starts. If the policy is supposed to be Active Directory-based, but cannot be obtained, you can delete **HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Group Policy\History** and then restart the computer.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**See also:** View Event Logs.

**Event Viewer is displaying "Bad SPI" error message.**
**Cause:** This error might occur if a key lifetime value is set too low, or if the SA has expired but the sender continues to transmit data to the receiver.

**Solution:** To determine and correct the problem, use IP Security Monitor to examine the number of rekeys. If the number of rekeys is very large compared to the amount of time the connections have been active, set the key lifetimes in the policy to be longer. Good values for high-traffic Ethernet connections are greater than both 50 megabytes and five minutes. This event is usually logged when an SA has been deleted on one IPSec peer, and the other IPSec peer sends secure traffic on the SA. This is usually a temporary situation that occurs on systems that are busy and fixes itself.

**IP Security Monitor displays the error message "The IPSec server is unavailable or incompatible with the IPSec monitor."**
**Cause:** You are using the IP Security Monitor console to monitor IPSec on a computer running Windows 2000.

**Solution:** The IP Security Monitor console can only be used to monitor IPSec on computers running Windows XP or a Windows Server 2003 operating system. To monitor IPSec on a computer running Windows 2000, use the **IPSecmon** command at the Windows 2000 command prompt on the computer that is being monitored.

**See also:** IPSec troubleshooting tools

# Windows Time Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Time Service**
Windows Time Service is a component of Windows XP and the Windows Server 2003 family that automates the synchronization of computers on a network.

- For tips about using Windows Time Service, see Windows Time Service Best practices.
- For help with specific tasks, see Windows Time Service How To....
- For general background information, see Windows Time Service Concepts.
- For problem-solving instructions, see Windows Time Service Troubleshooting.

# Windows Time Service Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Synchronize your domain with an accurate external time source**

  Synchronized clocks are necessary on all computers on the network for the proper operation of many services. You must also ensure that the synchronized time is correct. By default, the computers on the network obtain the time from their authenticating domain controller. The domain controllers on a network obtain their time from their domain's primary domain controller (PDC) operations master or an explicitly configured time source. Because of this, you must ensure that the PDC operations master in each domain is set to the correct time.

  There are many computers on the Internet that provide accurate time information. One such source is the National Institute of Standards and Technology (NIST), which provides the NIST Network Time Service.

- **Check Event Viewer**

  When Windows Time Service encounters errors, it writes events to the system log, so check Event Viewer periodically for messages regarding Windows Time Service. For more information, see Event Viewer.

# Windows Time Service How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Synchronize the client time with a time server
- Synchronize an internal time server with an external time source

**See Also**
**Other Resources**
Videos about the Windows Time Service

# Synchronize the client time with a time server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To synchronize the client time with a time server**

1. Open Command Prompt.
2. Type:

   **w32tm /resync**

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure only works on computers that are joined to a domain. For information about synchronizing computers that are not joined to a domain, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Synchronize an internal time server with an external time source
W32tm

# Synchronize an internal time server with an external time source

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To synchronize an internal time server with an external time source**

1. Open Command Prompt.

2. Type the following, where *PeerList* is a comma-separated list of Domain Name System (DNS) names or Internet protocol (IP) addresses of the desired time sources:

   **w32tm /config /syncfromflags:manual /manualpeerlist:***PeerList*

3. Type:

   **w32tm /config /update**

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The most common use of this procedure is to synchronize the internal network's authoritative time source with a very precise external time source. However, this procedure can be run on any computer running Windows XP or a member of the Windows Server 2003 family.

- If the computer cannot reach the servers, the procedure fails and an entry is written to the Event Log.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Synchronize the client time with a time server
W32tm

# Windows Time Service Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Windows Time Service:

- Windows Time Service overview

# Windows Time Service overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Time Service overview**

Computers keep the time on their internal clocks, which allows them to perform any function that requires the date or time. However, for scheduling purposes, the clocks must be set to the correct date and time, and they must be synchronized with the other clocks in the network. Without some other method in place, these clocks must be set manually.

With time synchronization, computers set their clocks automatically to match another computer's clock. One computer maintains very accurate time, and then all other computers set their clocks to match that computer. In this way, you can set accurate time on all computers.

**How Windows Time Service works**

In the Windows Server 2003 family, Windows Time Service automatically synchronizes the local computer's time with other computers on the network. The time source for this synchronization varies, depending on whether the computer is joined to an Active Directory® domain or a workgroup.

When the computers are part of a workgroup, you must manually configure the time synchronization settings. One or more computers might be identified as a locally reliable time source by configuring Windows Time Service on those computers to use a known accurate time source, either by using special hardware or a time source available on the Internet. Then, all other workgroup computers can be configured manually to synchronize their time with these local time sources.

When the computers belong to an family Active Directory domain, Windows Time Service configures itself automatically, using Windows Time Service that is available on domain controllers. Windows Time Service configures a domain controller within its domain as a reliable time source and synchronizes itself periodically with this source. These settings can be modified or overwritten, depending on specific needs.

**Benefits of Windows Time Service**

To function correctly, a number of Windows Server 2003 family components rely on accurate and synchronized time. For example, without clocks that are synchronized to the correct time on all computers, the Kerberos V5 authentication protocol might falsely interpret logon requests as intrusion attempts and deny access to users.

Another important benefit of time synchronization is the ability to correlate events on different computers in an enterprise. With synchronized clocks on all of your computers, you ensure that you can correctly analyze events that happen in sequence on multiple computers for success or failure.

# Windows Time Service Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Clients are not synchronized with servers.
- Windows Time service generates event ID 7023 after upgrading to Windows Server 2003 with Service Pack 1 (SP1).

**Clients are not synchronized with servers.**
**Cause:** Windows Time Service has been stopped manually.

**Solution:** Start Windows Time Service. To start Windows Time Service, follow these steps:

1. Open Command Prompt.
2. Type:

   **net start w32time**

**See also:** Start, stop, pause, resume, or restart a service; Configure how a service is started

**Windows Time service generates event ID 7023 after upgrading to Windows Server 2003 with Service Pack 1 (SP1).**
**Cause:** The LocalService account has not been granted "Change the system time" permissions. Windows Server 2003 with SP1 changes the startup configuration of the Windows Time service from LocalSystem to LocalService. Therefore, the startup account that the Windows Time service uses must have "Change the system time" permissions.

**Solution:** Use Domain Controller Security Policy to grant "Change the system time" permissions to the LocalService account. Alternatively, change the logon account of the Windows Time Service. For more information about performing these solutions, see article Q892501 "The Windows Time service may generate event ID 7023 after you upgrade to Windows Server 2003 Service Pack 1," in the Microsoft Knowledge Base.

**See also:** Change the system time; Permit users to log on locally to a domain controller

# Security Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

For more information about security, see the Windows Security Collection.

For more information about security groups, see "Designing an Authorization Strategy" at the Microsoft Windows Resource Kits Web site.

For the latest recommendations about security, see the Microsoft Security page on the Microsoft Web site.

For information about the Microsoft Security Notification Service, see Product Security Notification Web site.

# Security Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

- Securing Data Troubleshooting
- Stored User Name and Password Troubleshooting
- Password Troubleshooting
- Software Restriction Policies Troubleshooting
- Encrypting File System Troubleshooting
- Certificates Troubleshooting
- Certificate Services Troubleshooting
- Certificate Templates Troubleshooting
- Windows Time Service Troubleshooting

# Security Configuration Wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Security Configuration Wizard (SCW) is an attack surface reduction tool for members of the Microsoft® Windows® Server 2003 family with Service Pack 1 (SP1). SCW determines the minimum functionality required for a server's role or roles, and disables functionality that is not required.

For the latest information about SCW, see "Security Configuration Wizard for Windows Server 2003" at the Microsoft Windows Server 2003 Web site. (http://go.microsoft.com/fwlink/?LinkID=42434)

SCW does the following:

- Disables unneeded services.
- Blocks unused ports.
- Allows further address or security restrictions for ports that are left open.
- Prohibits unnecessary IIS web extensions, if applicable.
- Reduces protocol exposure to server message block (SMB), LanMan, and Lightweight Directory Access Protocol (LDAP).
- Defines a high signal-to-noise audit policy.

SCW guides you through the process of creating, editing, applying, or rolling back a security policy based on the selected roles of the server. The security policies that are created with SCW are XML files that, when applied, configure services, network security, specific registry values, audit policy, and if applicable, Internet Information Services (IIS).

**Note**

- In some cases, you must be connected to the Internet to use the links in SCW Help. If your computer is not connected to the Internet, you can find the same topic in Help and Support Center by searching for the link text. To open Help and Support Center, click **Start**, and then click **Help and Support Center**.

You can install SCW on computers running a member of the Windows Server 2003 family with SP1.

- For recommendations about how to get the most out of SCW, see Security Configuration Wizard Best Practices.
- For information about how to perform certain tasks, see Security Configuration Wizard How To.
- For general background information, see Security Configuration Wizard Concepts.

You do not need to run SCW to help secure your computer running Windows Small Business Server 2003. Instead, Windows Small Business Server 2003 uses the default settings in Setup and in the Configure E-mail and Internet Connection Wizard to help secure your server.

If you have not already run the Configure E-mail and Internet Connection Wizard, you should run it to help secure your server.

To start the Configure E-mail and Internet Connection Wizard on the computer running Windows Small Business Server 2003:

- Click **Start** and then click **Server Management**.
- In the console tree, click **Internet and E-mail**.
- In the details pane, click **Connect to the Internet**.

# Security Configuration Wizard Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best Practices**
This section tells how to get the most out of Security Configuration Wizard (SCW).

**Identify and target similar servers**
SCW helps to reduce the attack surface of servers by creating a security policy that is specifically designed for their specific roles. Administrators can simplify policy authoring and distribution by identifying groups of servers that perform the same, or similar, tasks. Here are ways you can do this:

- **Author one policy for a group of servers**. SCW authors a security policy based on the roles, tasks, and functions performed by a server. Others servers that perform the same, or very similar, functions can be configured with the same security policy. Administrators can use SCW once to author a security policy, save it, and apply it to all servers that perform the job function.

- **Group similar servers in one organizational unit (OU)**. The SCW transform operation can apply a security policy to a domain or OU by using Group Policy. To simplify policy distribution, an administrator could group servers that perform similar job functions, and use the same security policy, into a single OU. A new security policy can be distributed quickly and easily to the server OU by using the SCW transform operation.

- **Create policies for similar platforms**. For services or ports specific to 64-bit computers, create the policies on a 64-bit computer. Then deploy these policies to other 64-bit computers only (not 32-bit computers) to ensure the services are properly identified and configured.

It is highly recommended that the prototype server from which the security policy will be created matches the target servers to be configured at the service level. The security policy disables any service on the server that is contained in the Security Configuration Database but was not present on the prototype server when the policy was created. For example, if the DCOM Server Process Launcher service is listed in the Security Configuration Database, but is not present on the prototype server, the security policy created based on the prototype server will set the DCOM Server Process Launcher state to disabled. When you apply the security policy to other servers, the DCOM Server Process Launcher service will be disabled on those servers. You can configure unnecessary services in SCW (you can disable the service or leave the startup mode of the service unchanged), but only services that are not in the Security Configuration Database, and therefore are not defined in the security policy that you create with SCW.

**Test new security policies offline before deployment**

- The settings configured in the new security policies may cause compatibility issues with applications or services. Therefore, thoroughly test new security policies in a test environment before applying the policies to production servers.

**Create one complete security policy**

- SCW should be used to author a single security policy that contains all desired security settings for a server. This will simplify configuration, rollback, and analysis. For simple configuration and rollback, a single security policy for a machine, or set of machines, is much easier to understand and update than a series of policies. If a security policy defines all the desired settings for a server, a compliance report can be generated by executing one scan. This makes analysis using the scwcmd /analyze command easy. For more information about scwcmd, see Security Configuration Wizard command-line tool.

**Organize similar servers into organizational units (OUs) in Active Directory**

- Grouping servers by OUs in Active Directory domains facilitates the application of security policy through Group Policy.

# Security Configuration Wizard How To

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2


Install Security Configuration Wizard

Install Security Configuration Wizard unattended

Start Security Configuration Wizard

Create, edit, apply, or roll back a security policy

View Change Logs

Remove Security Configuration Wizard from a computer

# Install Security Configuration Wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

1. Open **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. Select the **Security Configuration Wizard** check box, click **Next**, and then click **Finish**.

**Notes**

- Security Configuration Wizard (SCW) is an optional component in Windows Server 2003 SP1. To install SCW, you must first install Service Pack 1 for Windows Server 2003.
- To perform this procedure, you must be a member of the Administrators group on the local computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open **Add or Remove Programs**, click **Start**, click **Control Panel**, and then double-click **Add or Remove Programs**.

# Install Security Configuration Wizard unattended

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

1. Before initiating the unattended installation, add the following entry to the **[Components]** section of Unattend.txt: **SCW=On**

To apply an SCW policy during the unattended installation, also carry out the following steps.

1. Create the policy file on a server that already has SCW installed.
2. Either create a Cmdlines.txt file (or modify the existing one) so that it has a **[Commands]** section containing the following line: **scwcmd configure /p:SCWPolicy.xml**
3. Copy the Cmdlines.txt file and the previously created policy file to the **$OEM$** directory. The **$OEM$** directory is specified by the OEMFILESPATH entry in the **[Unattended]** section of Unattend.txt.

**Note**

- For detailed instructions about running unattended installation, see Automating and Customizing Installations on the Microsoft Web Site.

# Start Security Configuration Wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

1. Open **Administrative Tools**.
2. Double-click **Security Configuration Wizard.**

### Notes

- To open Administrative Tools, click **Start**, click **Control Panel**, and then click **Administrative Tools**.
- Security Configuration Wizard (SCW) can be used to author policy and to perform configuration or rollback for one server at a time. The SCW executable is called SCW.exe. To start SCW from the command prompt, type **SCW**.
- The SCW command-line tool can be used for multi-machine configuration, analysis, rollback, and Group Policy object (GPO) policy management. The SCW command-line tool is called scwcmd.exe. To start SCW from the command prompt, type **SCWcmd**. To see SCW command-line help, type scwcmd.exe /?

# Create, edit, apply, or roll back a security policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

1. Start the Security Configuration Wizard.
2. On the **Welcome to the Security Configuration Wizard** page, click **Next**.
3. On the **Configuration Action** page, do one of the following:
   - To create a security policy, click **Create a new security policy**.
   - To edit a security policy, click **Edit an existing security policy**.
   - To apply a security policy, click **Apply an existing security policy**.
   - To roll back a security policy, click **Rollback the last applied security policy**.
4. If you chose to edit or apply an existing security policy, in Existing security policy file, enter the security policy path and file name, or browse to the security policy, and click Next.
5. Complete the Security Configuration Wizard.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- For information about how to open Security Configuration Wizard (SCW), see Related Topics.
- To apply or roll back a security policy that you have created with SCW, SCW must be installed on the selected server. When creating or editing a security policy, you can select a server that does not have SCW installed. In either case, the selected server must always be a computer running a member of the Windows Server 2003 family with Service Pack 1.
- Do not end the SCW task while policy configuration is in progress as this can leave the server in a partially configured state. A policy should instead be rolled back after it has been applied to the system.

**See Also**
**Concepts**
Starting Security Configuration Wizard
Security Configuration Wizard How To
Security Configuration Wizard

# View Change Logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

- At the command prompt, type scwcmd view /x:file_name.xml, where file_name is the name of the change log. Example: scwcmd view /x:AnalysisLog10.xml

**Note**

- When you use Security Configuration Wizard (SCW) to apply a policy, SCW creates a change log file to document the changes. Change logs are stored in the following directory:

  %windir%\security\msscw\ChangeConfigurationLogs

# Remove Security Configuration Wizard from a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

1. Open **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. Clear the **Security Configuration Wizard** check box, click **Next**, and then click **Finish**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open **Add or Remove Programs**, click **Start**, click **Control Panel**, and then double-click **Add or Remove Programs**.
- When Security Configuration Wizard (SCW) is uninstalled, all files contained in the %Windir%\Security\Msscw folder are removed. This includes all policy files, log files, and rollback files that were created after SCW was installed.

# Security Configuration Wizard Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2


Security Configuration Wizard Overview

Security Configuration Wizard command-line tool

Extending the Security Configuration Database

Rolling back a security policy

Understanding Security Configuration Wizard

# Security Configuration Wizard Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Security Configuration Wizard (SCW) is an attack surface reduction tool for computers running a member of the Windows Server 2003 family with Service Pack 1 (SP1). SCW guides you through the process of creating a security policy, based on the roles performed by a given server. Once a policy is created, it can be edited or applied to one or more similarly configured servers. Applied policies can be rolled back in order to undo changes that have caused problems. To edit, apply, or roll back a security policy, the policy must have been created with SCW.

SCW can only be installed on computers running a member of the Windows Server 2003 family with SP1. Also, the server that you are configuring must be running a member of the Windows Server 2003 family with SP1.

SCW is composed of the following sections:

- Role-based service configuration
- Network security
- Registry settings
- Audit policy
- Internet Information Services (IIS)

You will not see the IIS section if you have not selected the Web Server role on the Select Server Roles page of SCW.

SCW helps reduce the attack surface of a configured server, and you can use other security features to further secure your computers. For more information, see Security on the Microsoft Web Site.

The security policy that you create is based on the roles that are installed on the selected server. The selected server can be a server to which you want to apply the security policy, or you can use the selected server to create a security policy, and then apply the security policy to a group of servers that perform similar roles as the selected server. For example, if you want to create a security policy for all of the file servers in your organization, you can choose one representative file server as the selected server, use that server to create the security policy, and then apply the security policy to each file server in your organization. For more information about selecting a server, see Select Server Roles.

While using SCW to configure a policy, if you are not sure of the correct selection for your environment, it is recommended that you skip the associated sections by clicking Next and then configure the rest of the security policy. You can configure the skipped sections later, once you are familiar with the technology. SCW does not configure skipped sections. These sections are left undefined until you edit the policy and configure those sections. However, you cannot skip the Role-Based Service Configuration section.

**Deploying SCW security policies by means of Group Policy**

SCW policies are XML files that can contain settings for services, Windows Firewall, Internet Protocol security (IPsec), registry values, audit policy, IIS, and other settings that are imported from existing security templates. These SCW XML policy files are not natively supported by the Group Policy infrastructure. Instead, SCW XML policy files must be transformed into formats that the various Group Policy extensions understand. The SCW command line tool, Scwcmd.exe, provides the Transform operation. For help with this tool, at a command prompt, type **scwcmd transform**. In short, the administrator provides the Transform operation with the name of the SCW policy file to be transformed, along with the name of a Group Policy object (GPO) to be created. The Transform operation then creates the GPO and defines the settings for these Group Policy extensions:

- **Security Settings** contains service settings, registry values, audit policy and security template settings that were imported into the SCW XML policy.
- **IP Security Policies** contains the IPsec configuration that is defined in the SCW policy.
- **Windows Firewall** contains Windows Firewall settings that are defined in the SCW policy.

After the GPO is created, the administrator must manually link the GPO to the target organizational unit (OU) by using Active Directory Users and Computers, or by using the Group Policy Management Console.

**Notes**

- For more information about Group Policy, see Configure Group Policy. For more information about scwcmd.exe, see Security Configuration Wizard command-line tool. For more information about OUs, see Organizational units. For the latest information about Windows Firewall performance considerations, see Windows Firewall on the Microsoft Web Site.
- You must be logged on with an account that is a member of the Domain Admins group in order to perform the Transform operation.
- The Transform operation must be performed from a server that is joined to the domain where the Group Policy object will apply.
- Any IIS settings that are defined in the SCW policy will be lost during the Transform operation because Group Policy does not have an IIS extension.
- The OU hierarchy in Active Directory must be organized along server roles in order to take advantage of SCW-generated policies. There should be a separate OU for each server role supported by your organization.
- Applying SCW-generated policies to servers running versions of Windows earlier than Windows Server 2003 with SP1 is not supported. Use the WMI filtering capabilities of Group Policy to ensure that SCW-generated policies do not apply to servers running versions of Windows earlier than Windows Server 2003 with SP1. Alternatively, place only servers running Windows Server 2003 with SP1 into the OUs where the transformed SCW policies will apply.

**Roles**

*Server roles* are the main functions performed by the servers in your organization. Examples of server roles are file server, domain controller, and Web server roles. *Client features* use services provided by other servers. Examples of client features are the domain member and Dynamic Host Configuration Protocol (DHCP) client.

A role is defined by the services, ports, and IIS requirements that are necessary to perform the role. The specific services, inbound ports, and settings that are required vary for each role. As a result, security policies created by SCW are specific to servers that perform the same roles as the selected server. If you apply a file server security policy to a Web server, the services, inbound ports, and settings that are required to perform the Web server role are disabled and the file server security policy settings are enabled. In some cases, security policies are specific to the selected server and must be modified before applying them to other servers.

To determine which services are required for a server to perform a certain role, on the **Select Server Roles, Select Client Features**, and **Select Administration and Other Options pages**, click the triangle between the check box and the role name. Information about the role appears below its name, including a brief description and the required services.

SCW analyzes the selected server to determine which roles it performs. The scanning and analysis process is guided by the *Security Configuration Database*. The Security Configuration Database also contains information about services that are required to perform a role, as well as ports that must be open, if applicable.

For more information about the Security Configuration Database, see Security Configuration Database.

By using the information in the Security Configuration Database, SCW determines which roles the selected server is capable of performing with the software that is currently installed and which roles the server is currently configured to perform. If a server is currently configured to perform a specific role, that role is automatically selected by default, and the role appears in the **Installed server roles** view on the **Select Server Roles** page, and the same is true for the **Select Client Features** and **Select Administration and Other Options** pages. If a server is likely performing a specific role, then that installed role is automatically selected by default.

Selecting roles from the **All roles** view will have no effect on the selected server with the software that is currently installed. These views are useful if you plan to install other roles on the selected server or if you plan to apply the security policy to multiple servers that perform similar roles, but have slight differences.

For more information, see Security Configuration Wizard.

# Security Configuration Wizard command-line tool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Security Configuration Wizard (SCW) includes the Scwcmd.exe command-line tool. You can use Scwcmd for the following tasks:

- Configure one or many servers with an SCW-generated policy.
- Analyze one or many servers with an SCW-generated policy.
- View analysis results in HTML format.
- Roll back SCW policies.
- Transform an SCW-generated policy into native files that are supported by Group Policy.
- Register a Security Configuration Database extension with SCW.

When you use scwcmd to configure, analyze, or roll back a policy on a remote server, SCW is required to be installed on the remote server.

**Getting command-line help for Scwcmd**
To get basic help on the Scwcmd tool, at a command prompt, type **scwcmd**.

**Viewing formatted analysis results with the scwcmd /view command**
The directory %windir%\security\msscw\transformfiles contains .xsl transform files. These are applied to the .xml policy file for the rendering process when you view analysis results through the scwcmd /view command.

# Extending the Security Configuration Database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This is an advanced topic for developers with XML expertise and for IT professionals.

You can create a new Security Configuration Database file or custom applications in your environment.

**Caution**

- When customizing the Security Configuration Database, do not make changes to W2K3.xml or to any other XML files that are included with Security Configuration Wizard (SCW).

You can create a new Security Configuration Database so that it is customized for your environment. For example, you may be creating a security policy and your organization uses custom applications that require specific services to be enabled or ports to be open. You can customize the Security Configuration Database to contain this information. For more information about the Security Configuration Database, see Security Configuration Database.

**Caution**

- When customizing the Security Configuration Database, do not make changes to W2K3.xml.

Create your own XML file and save it to **%Systemroot%\Security\Msscw\KBs**. To customize an existing role, in your custom Security Configuration Database, create a new role with the same role type and name as the role in W2K3.xml. For example, if you want to customize the file server role, in your custom XML file, you would create a role with Server and File. You can customize the entries for <Selected>, <Services>, and <Ports>. Then, open KBReg.xml and replace <**Custom** /> with <**Custom**>Custom_KB_Name.**xml** </**Custom**>.

For more information about the Security Configuration Database, see Security Configuration Database.

# Rolling back a security policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use Security Configuration Wizard (SCW) to roll back a security policy after it is applied.

You are not limited to one rollback. However, to avoid administrative overhead, do not roll back a policy more often than you must. Make multiple changes during each SCW session, rather than running through SCW several times and creating a separate policy each time.

If, after you roll back a security policy, some services do not behave as expected, check whether services that were started manually were stopped. If you apply a security policy that disables manually-started services and then roll back this policy, those services might need to be restarted manually.

# Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory**
Active Directory® is a directory service that stores information about objects on a network and makes this information available to users and network administrators.

- Before deploying Active Directory, see Active Directory Checklists.
- To find administrative tool equivalents that have changed from Windows NT® 4.0 Server, see New ways to do familiar Active Directory tasks.
- For tips about using Active Directory, see Active Directory Best practices.
- For help with specific tasks, see Active Directory How To....
- For general background information, see Active Directory Concepts
- For problem-solving instructions, see Active Directory Troubleshooting.

# Active Directory Checklists

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

- Checklist: Creating a new forest
- Checklist: Creating a new child domain
- Checklist: Creating a new domain tree
- Checklist: Creating an additional domain controller in an existing domain
- Checklist: Creating a group
- Checklist: Configuring a site
- Checklist: Optimizing intersite replication
- Checklist: Creating a forest trust

# Checklist: Creating a new forest

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a new forest**

| Step | Reference |
|---|---|
| (Optional) Review concepts and key DNS implications to consider before setting up a new forest. | Creating a new forest |
| (Optional) Review concepts about security and other options available when using the Active Directory Installation Wizard. | Using the Active Directory Installation Wizard |
| Verify that the server on which you will be installing Active Directory has an NTFS partition. | Reformatting or converting a partition to use NTFS |
| (Optional) Review the role of a domain controller. | Domain controllers |
| (Optional) Review concepts on Domain Name System (DNS) namespace planning and how Active Directory is integrated with DNS. | Namespace planning for DNS; Configure a DNS server for use with Active Directory; DNS integration |
| Verify that you are a member of the Administrators group on the server on which you will be installing Active Directory. | Default groups |
| Verify that DNS is properly configured before installing Active Directory. | Checklist: Verifying DNS before installing Active Directory |
| Create the domain controller. | Create a new forest |

# Checklist: Creating a new child domain

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a new child domain**

| Step | Reference |
| --- | --- |
| (Optional) Review concepts about adding new child domains. | Creating a new child domain |
| (Optional) Review concepts about security and other options available when using the Active Directory Installation Wizard. | Using the Active Directory Installation Wizard |
| Verify that the server on which you will be installing Active Directory has an NTFS partition. | Reformatting or converting a partition to use NTFS |
| Verify that you are a member of the Enterprise Admins group or the Domain Admins group of the parent domain. | Default groups |
| Verify that DNS is properly configured before installing Active Directory. | Checklist: Verifying DNS before installing Active Directory |
| Create the domain controller. | Create a new child domain |

# Checklist: Creating a new domain tree

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a new domain tree**

| Step | Reference |
|---|---|
| (Optional) Review concepts about adding new domain trees. | Creating a new domain tree |
| (Optional) Review concepts about security and other options available when using the Active Directory Installation Wizard. | Using the Active Directory Installation Wizard |
| Verify that the server on which you will be installing Active Directory has an NTFS partition. | Reformatting or converting a partition to use NTFS |
| Verify that you are a member of the Enterprise Admins group or the Domain Admins group in the forest root domain. | Default groups |
| Verify that DNS is properly configured before installing Active Directory. | Checklist: Verifying DNS before installing Active Directory |
| Create the domain controller. | Create a new domain tree |

# Checklist: Creating an additional domain controller in an existing domain

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating an additional domain controller in an existing domain**

| Step | Reference |
|---|---|
| (Optional) Review concepts about creating additional domain controllers over the network or by using backup media. | Creating an additional domain controller |
| (Optional) Review concepts about security and other options available when using the Active Directory Installation Wizard. | Using the Active Directory Installation Wizard |
| Verify that the server on which you will be installing Active Directory has an NTFS partition. | Reformatting or converting a partition to use NTFS |
| (Optional) Review the role of a domain controller. | Domain controllers |
| Verify that you are a member of the Domain Admins group in the domain where you will be adding the domain controller. | Default groups |
| Verify that DNS is properly configured before installing Active Directory. | Checklist: Verifying DNS before installing Active Directory |
| Configure antivirus scanning exceptions. | Managing Antivirus Software on Active Directory Domain Controllers |
| Create the domain controller. | Create an additional domain controller |

# Checklist: Creating a group

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a group**

| Step | Reference |
|---|---|
| (Optional) Review concepts about groups. | Groups; Understanding Groups |
| Identify which type of group (security or distribution) you want to create. | Group types |
| Identify which scope (universal, global, or domain local) you want to create for the group. | Group scope |
| Use the appropriate administrative tool to create the group:<br>• To create a local group on a workstation or member server, use Local Users and Groups in Computer Management.<br>• To create a group in a Windows Server 2003 domain, use Active Directory Users and Computers. | Create a local group; Create a new group |

# Checklist: Configuring a site

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring a site**

| Step | Reference |
|---|---|
| (Optional) Review sites and replication concepts. | Sites overview; Replication overview; Managing replication |
| Create a new site. | Create a site |
| Associate a subnet with a site. | Associate a subnet with a site |
| Connect the site. | Create a site link;<br><br>Add a site to a site link |
| Customize replication. | Checklist: Optimizing intersite replication |

# Checklist: Optimizing intersite replication

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Optimizing intersite replication**

| Step | Reference |
| --- | --- |
| (Optional) Review sites and replication concepts. | Sites overview; Replication overview; Managing replication |
| Create site links. | Create a site link |
| Configure site link cost. | Configure site link cost; Configure site link replication availability |

# Checklist: Creating a forest trust

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a forest trust**

| Task | Reference |
|---|---|
| (Optional) Review the various trust types and understand forest trust concepts. | Trust types; Forest trusts; When to create a forest trust; Accessing resources across forests; Routing name suffixes across forests |
| Ensure that Domain Name System (DNS) is properly set up.<br><br>If there is a root DNS server that can be made the root DNS server for both of the forest DNS namespaces, make it the root server by ensuring that the root zone contains delegations for each of the DNS namespaces. Also, update the root hints of all DNS servers with the new root DNS server.<br><br>If there is no shared root DNS server and the root DNS servers for each forest DNS namespace are running a member of the Windows Server 2003 family, configure DNS conditional forwarders in each DNS namespace to route queries for names in the other namespace.<br><br>If there is no shared root DNS server, and the root DNS servers for each forest DNS namespace are not running a member of the Windows Server 2003 family, configure DNS secondary zones in each DNS namespace to route queries for names in the other namespace. | Verify a zone delegation using the nslookup command; Update root hints on the DNS server; Configure a DNS server to use forwarders; Add a secondary server for an existing zone; Namespace planning for DNS |
| Set the forest functional level in both forests to Windows Server 2003. | Raise the forest functional level |
| Create a forest trust. | Create a forest trust |

# New ways to do familiar Active Directory tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**
The following table lists common tasks for configuring Active Directory. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0.

### Note

- There are no significant user interface differences between Windows 2000 and the Windows Server 2003 family of products for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Install a domain controller | Windows Setup | The Active Directory Installation Wizard.<br><br>For more information, see Using the Active Directory Installation Wizard;Create a new forest;Create a new domain tree;Create a new child domain;Create an additional domain controller. |
| Manage users | User Manager | Active Directory Users and Computers.<br><br>For more information, see Manage Users. |
| Manage groups | User Manager | Active Directory Users and Computers. For more information, see Manage Groups. |
| Manage computer accounts | Server Manager | Active Directory Users and Computers. For more information, see Manage Computers. |
| Add a computer to a domain | Server Manager | Active Directory Users and Computers.<br><br>For more information, see Manage Computers. |
| Create or manage trust relationships | User Manager | Active Directory Domains and Trusts. For more information, see Manage Trusts. |
| Manage account policy | User Manager | Active Directory Users and Computers.<br><br>For more information, see Account and local policies. |
| Manage user rights | User Manager | Active Directory Users and Computers. Edit the Group Policy object for the domain or organizational unit containing the computers to which the user rights apply. For more information, see Edit security settings on a Group Policy object. |
| Manage audit policy | User Manager | Active Directory Users and Computers.<br><br>Edit the Group Policy object assigned to the Domain Controllers organizational unit. For more information, see Edit security settings on a Group Policy object. |

# Active Directory Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **As a security best practice, it is recommended that you do not log on to your computer with administrative credentials.**

  When you are logged on to your computer without administrative credentials, you can use Run as to accomplish administrative tasks.

  For more information, see Why you should not run your computer as an administrator and Using Run as.

- **To further secure Active Directory, it is recommended that you implement the following security guidelines:**

  - Rename or disable the Administrator account (and guest account) in each domain to prevent attacks on your domains. For more information, see User and computer accounts.

  - Physically secure all domain controllers in a locked room. For more information, see Domain controllers and Securing Active Directory.

  - Manage the security relationship between two forests and simplify security administration and authentication across forests. For more information, see Forest trusts.

  - To provide additional protection for the Active Directory schema, remove all users from the Schema Admins group, and add a user to the group only when schema changes need to be made. Once the change has been made remove the user from the group.

  - Restrict user, group, and computer access to shared resources and to filter Group Policy settings. For more information, see Group types.

  - Avoid disabling the use of signed or encrypted LDAP traffic for Active Directory administrative tools. For more information, see Connecting to domain controllers running Windows 2000.

  - Some default user rights assigned to specific default groups may allow members of those groups to gain additional rights in the domain, including administrative rights. Therefore, your organization must equally trust all personnel that are members of the Enterprise Admins, Domain Admins, Account Operators, Server Operators, Print Operators and Backup Operators groups. For more information about these groups, see Default groups.

  - Use global groups or universal groups instead of domain local groups when specifying permissions on domain directory objects replicated to the global catalog. For more information, see Global catalog replication.

    For general security information about Active Directory, see Security information for Active Directory and Securing Active Directory.

- **Establish as a site every geographic area that requires fast access to the latest directory information.**

  Establishing areas that require immediate access to up-to-date Active Directory information as separate sites will provide the resources required to meet your needs.

  For more information, see Create a site.

- **Place at least one domain controller in every site, and make at least one domain controller in each site a global catalog.**

Sites that do not have their own domain controllers and at least one global catalog are dependent on other sites for directory information and are less efficient.

For more information, see Enable or disable a global catalog.

- **Perform regular backups of domain controllers in order to preserve all trust relationships within that domain.**

  For more information, see Domain controllers.

# Active Directory How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Manage Users, Groups, and Computers
- Publish Resources
- Manage Domains
- Manage Operations Master Roles
- Manage Trusts
- Manage Forests
- Manage Sites
- Manage COM+ Partitions and Partition Sets
- Search Active Directory

# Manage Users, Groups, and Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage users, groups, and computers**

- Manage Users
- Manage Groups
- Manage Computers

# Manage Users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage users**

- Create a new user account
- Reset a user password
- Copy a user account
- Move a user account
- Set logon hours
- Disable or enable a user account
- Map a certificate to a user account
- Change a user's primary group
- Delete a user account
- Rename a user account

# Create a new user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new user account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the folder in which you want to add a user account.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder*

3. Point to **New**, and then click **User**.
4. In **First name**, type the user's first name.
5. In **Initials**, type the user's initials.
6. In **Last name**, type the user's last name.
7. Modify **Full name** to add initials or reverse order of first and last names.
8. In **User logon name**, type the user logon name, click the UPN suffix in the drop-down list, and then click **Next**.

   If the user will use a different name to log on to computers running Windows 95, Windows 98, or Windows NT, then you can change the user logon name as it appears in **User logon name (pre-Windows 2000)** to the different name.

9. In **Password** and **Confirm password**, type the user's password, and then select the appropriate password options.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To add a user, you can also click

  on the toolbar.

- To add a user, you can also copy any previously created user account. For more information, see Related Topics.

- A new user account with the same name as a previously deleted user account does not automatically assume the permissions and group memberships of the previously deleted account because the security ID (SID) for each account is unique. To duplicate a deleted user account, all permissions and memberships must be manually recreated.

- For interoperability with other directory services, you can create an InetOrgPerson user object. To create a new inetOrgPerson, in step three, click **InetOrgPerson** instead of **User**. For more information about InetOrgPerson, see User and computer accounts in Related Topics.

- When creating a new user, the full name attribute is created in the *FirstNameLastName* format by default. The full name attribute also governs the display name format that is shown in the global address list. You can change the display name format by using ADSI Edit. If you do so, this will also change the full name format. For more information, see article Q250455, "How to Change Display Names of Active Directory Users" in the Microsoft Knowledge Base.

- Windows NT 4.0 and earlier domains allow the use of a period (.) at the end of a user logon name as long as the user logon name does not consist solely of period characters. Windows Server 2003 domains do not allow the use of a period or multiple periods at the end of a user logon name.

#### Using a command line

1. Open Command Prompt.

2. Type:

   **dsadd user***UserDN* [**-samid***SAMName*] **-pwd** {*Password*|*****}

| Value | Description |
|---|---|
| *UserDN* | Specifies the distinguished name of the user object to be added. |
| *SAMName* | Specifies the Security Accounts Manager (SAM) name as the unique SAM account name for this user (for example, Linda). If not specified, dsadd will attempt to create the SAM account name using up to the first 20 characters from the common name (CN) value of UserDN. |
| *Pwd* | Specifies the password to be used for the user account. If set to *, you are prompted for a user password. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To add a user, you can also copy any previously created user account. For more information, see Related Topics.

- A new user account with the same name as a previously deleted user account does not automatically assume the permissions and group memberships of the previously deleted account because the security ID (SID) for each account is unique. To duplicate a deleted user account, all permissions and memberships must be manually recreated.

- To view the complete syntax for this command, and for information on entering more user account information, at a command prompt, type:

  **dsadd user /?**

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsadd
Find a user account
User and computer accounts
Copy a user account
Object names
Install Windows Support Tools
Using online resources

# Reset a user password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reset a user password**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Users

   Or, click the folder that contains the user account.
3. In the details pane, right-click the user whose password you want to reset, and then click **Reset Password**.
4. Type and confirm the password.
5. If you want to require the user to change this password at the next logon process, select the **User must change password at next logon** check box.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Services that are authenticated with a user account must be reset if the password for the service's user account is changed.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmod user***UserDN***-pwd***NewPassword*
3. If you want to require the user to change this password at the next logon process, Type:

   **dsmod user***UserDN***-mustchpwd** {**yes**|**no**}

| Value | Description |
| --- | --- |
| *UserDN* | Specifies the distinguished name of the user for which the password will be reset. |
| *NewPassword* | Specifies the password that will replace the current user password. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If a password is not assigned, then the first time the user tries to log on (using a blank password), a logon message will appear that says "You are required to change your password at first logon" and after user has changed the password, the logon process continues.

- Services that are authenticated with a user account must be reset if the password for the service's user account is changed.

- To view the complete syntax for this command, at a command prompt, type:

  **dsmod user /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find a user account
Command-line reference A-Z
Directory service command-line tools

# Copy a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy a user account**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Users

   Or, click the folder that contains the user account.
3. In the details pane, right-click the user account that you want to copy, and then click **Copy**.
4. In **First name**, type the user's first name.
5. In **Last name**, type the user's last name.
6. Modify **Full name** to add initials or reverse order of first and last names.
7. In **User logon name**, type the user logon name, click the UPN suffix in the drop-down list, and then click **Next**.

   If the user will use a different name to log on to computers running Windows 95, Windows 98, or Windows NT, then you can change the user logon name as it appears in **User logon name (pre-Windows 2000)** to the different name.
8. In **Password** and **Confirm password**, type the user's password, and then select the appropriate password options.

   If the user account from which the new user account was copied was disabled, click **Account is disabled** to enable the new account.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- After creating the user account, edit the user account properties to enter additional user account information.
- Only the most commonly used attributes, such as logon hours, workstation restrictions, and account expiration restrictions are transferred to the copied user by default.
- You can modify which default attributes are carried over to a newly copied user or specify additional attributes that will be copied to the new user. To do this, open the Active Directory Schema snap-in, view the desired attribute properties, and select (or clear) the **Attribute is copied when duplicating user** check box. You can only modify or add attributes that are instances of the user class. If the Active Directory Schema snap-in is not installed, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
User and computer accounts
Object names
Install the Active Directory Schema snap-in

# Move a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move a user account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Users

   Or, click the folder that contains the user account.
3. In the details pane, right-click the user you want to move, and then click **Move**.
4. In the **Move** dialog box, click the folder to which you want to move the user account.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- You cannot use Active Directory Users and Computers to move user accounts between domains. To move a user account between domains, use Movetree, one of the Active Directory support tools. For more information, see Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmove***UserDN* [**-newparent** ParentDN]

| Value | Description |
|---|---|
| *UserDN* | Specifies the distinguished name of the user object to be moved. |

| | |
|---|---|
| *ParentDN* | Specifies the new location to which you want to move the user object. The new location is specified as the distinguished name of the new parent node. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, and for information on entering more user account information, at a command prompt, type:

  **dsmove /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find a user account
Active Directory support tools

# Set logon hours

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set logon hours**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Users

   Or, click the folder that contains the user account.
3. Right-click the user account, and then click **Properties**.
4. On the **Account** tab, click **Logon Hours**, and then set the permitted or denied logon hours for the user.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To modify the logon hours for multiple users, press and hold down CTRL, and then click each user. Right-click the selected users, and then click **Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
User and computer accounts
Working with MMC console files

# Disable or enable a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable or enable a user account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Users

   Or, click the folder that contains the user account.
3. In the details pane, right-click the user.
4. Depending on the status of the account, do one of the following:

   - To disable, click **Disable Account**.
   - To enable, click **Enable Account**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To prevent a particular user from logging on for security reasons, you can disable user accounts rather than deleting user accounts.
- By creating disabled user accounts with common group memberships, you can use disabled user accounts as account templates to simplify user account creation. For more information, see Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmod user***UserDN***-disabled** {**yes**|**no**}

| Value | Description |
| --- | --- |
| *UserDN* | Specifies the distinguished name of the user object to be disabled or enabled. |
| {**yes**\|**no**} | Specifies whether the user account is disabled for log on (**yes**) or not (**no**). |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **dsmod user /?**

- To prevent a particular user from logging on for security reasons, you can disable user accounts rather than deleting user accounts.

- By creating disabled user accounts with common group memberships, you can use disabled user accounts as account templates to simplify user account creation. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
User and computer accounts
Copy a user account
Command-line reference A-Z
Directory service command-line tools

# Map a certificate to a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To map a certificate to a user account**

1. Open Active Directory Users and Computers.
2. On the **View** menu, select **Advanced Features**.
3. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers\domain node\Users

   Or, click the folder that contains the user account.
4. In the details pane, click the user account to which you want to map a certificate.
5. On the **Action** menu, click **Name Mappings**.
6. In the **Security Identity Mapping** dialog box, on the **X.509 Certificates** tab, click **Add**.
7. Type the name and path of the .cer file that contains the certificate you want to map to this user account, and then click **Open**.
8. Do one of the following:

| To | Do this |
|---|---|
| Map the certificate to one account (one-to-one mapping) | Confirm that both the **Use Issuer for alternate security identity** and the **Use Subject for alternate security identity** check boxes are selected. |
| Map any certificate that has the same subject to the user account, regardless of the issuer of the certificate (many-to-one mapping) | Clear the **Use Issuer for alternate security identity** check box, and confirm that the **Use Subject for alternate security identity** check box is selected. |
| Map any certificate that has the same issuer to the user account, regardless of the subject of the certificate (many-to-one mapping) | Clear the **Use Subject for alternate security identity** check box, and confirm that the **Use Issuer for alternate security identity** check box is selected. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- The certificate you are mapping to a user account must be in Distinguished Encoding Rules (DER) or Base64 encoded binary format. For instructions on exporting an existing certificate to a .cer file, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Mapping certificates to user accounts
Export a certificate
Importing and exporting certificates
Map an account from a trusted non-Windows kerberos realm to a user account

# Change a user's primary group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a user's primary group**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Users

   Or, click the folder that contains the user account.
3. In the details pane, right-click the user you want to change, and then click **Properties**.
4. On the **Member of** tab, click the group that you want to set as the user's primary group, and then click **Set Primary Group**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- The user's primary group applies only to users who log on to the network through Services for Macintosh or to who run POSIX-compliant applications. Unless you are using these services, there is no need to change the primary group from Domain Users, which is the default value.
- Setting the user's primary group membership to a value other than Domain Users may adversely impact performance as all users in the domain are members of Domain Users. If the user's primary group is set to another group, it may cause the group membership to exceed the supported maximum number of members.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Accounts for Macintosh clients
Services for Macintosh

# Delete a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a user account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Users

   Or, click the folder that contains the user account.
3. In the details pane, right-click the user account, and then click **Delete**.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Once a user account has been deleted, all permissions and memberships associated with that user account are permanently deleted. Because the security ID (SID) for each account is unique, a new user account with the same name as a previously deleted user account does not automatically assume the permissions and memberships of the previously deleted account. To duplicate a deleted user account, all permissions and memberships must be manually recreated.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsrm***ObjectDN*

| Value | Description |
|---|---|
| *ObjectDN* | Specifies the distinguished name of the user object to be deleted. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the

Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Once a user account has been deleted, all permissions and memberships associated with that user account are permanently deleted. Because the security ID (SID) for each account is unique, a new user account with the same name as a previously deleted user account does not automatically assume the permissions and memberships of the previously deleted account. To duplicate a deleted user account, all permissions and memberships must be manually recreated.

- To view the complete syntax for this command, at a command prompt, type:

  **dsrm /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find a user account
Command-line reference A-Z
Directory service command-line tools

# Manage Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage groups**


- Create a new group
- Add a member to a group
- Allow anonymous users to be members of the Everyone security group on a domain controller
- Convert a group to another group type
- Change group scope
- Delete a group
- Find groups in which a user is a member
- Assign user rights to a group in Active Directory

# Create a new group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new group**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the folder in which you want to add a new group.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder*

3. Point to **New**, and then click **Group**.
4. Type the name of the new group.

   By default, the name you type is also entered as the pre–Windows 2000 name of the new group.

5. In **Group scope**, click one of the options.
6. In **Group type**, click one of the options.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To add a group, you can also click the folder in which you want to add the group, and then click

  on the toolbar.

- If the domain in which you are creating the group is set to the domain functional level of Windows 2000 mixed, you can select only the **Security** group type with **Domain local** or **Global** scopes. For more information, see Related Topics.
- In Active Directory Users and Computers, by default, the name that you type is also entered as the pre–Windows 2000 name of the new group.
- A group name cannot be identical to any other group name in the domain.
- A group name (CN) can contain up to 64 uppercase or lowercase characters.

- A group name (CN) cannot consist solely of spaces.
- A group name (pre–Windows 2000) (**samAccountName** object attribute) can contain up to 256 uppercase of lowercase characters except for the following:

  " / \ [ ] : ; | = , + * ? < >

- A group name (pre–Windows 2000) (**samAccountName** object attribute) cannot consist solely of periods or spaces.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsadd group** *GroupDN* **–samid** *SAMName* **–secgrp** *yes* | *no* **–scope** *l* | *g* | *u*

| Value | Description |
|---|---|
| *GroupDN* | Specifies the distinguished name of the group object that you want to add. |
| *SAMName* | Specifies to use the Security Accounts Manager (SAM) name as the unique SAM account name for this group (for example, operators). |
| *yes* \| *no* | Specifies whether the group you want to add is a security group (yes) or a distribution group (no). |
| *l* \| *g* \| *u* | Specifies whether the scope of the group you want to add is domain local (l), global (g), or universal (u). |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- If the domain in which you are creating the group is set to the domain functional level of Windows 2000 mixed, you can select only the **Security** group type with **Domain local** or **Global** scopes. For more information, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dsadd group /?**

- A group name cannot be identical to any other group name in the domain.
- A group name (CN) can contain up to 64 uppercase or lowercase characters.
- A group name (CN) cannot consist solely of spaces.
- A group name (pre–Windows 2000) (**samAccountName** object attribute) can contain up to 256 uppercase of lowercase characters except for the following:

  " / \ [ ] : ; | = , + * ? < >

- A group name (pre–Windows 2000) (**samAccountName** object attribute) cannot consist solely of periods or spaces.

| Note |
| --- |
| When you use **net group** to create a new group account (**net group <group_name> /add /domain**, for example, **net group Group1 /add /domain**), if you specify a group name that is longer than 64 characters, the directory service sets the group's CN to the automatically generated **objectSID** of the newly created group account and the **samAccountName** object attribute assumes the name that you specify in the **net group** command. For more information about **net group** name restrictions, see NetGroupAdd (http://go.microsoft.com/fwlink/?LinkID=159751). |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Group scope
Group types
Object names
Command-line reference A-Z
Dsadd

# Add a member to a group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a member to a group**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click the folder that contains the group to which you want to add a member.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder that contains the group*

3. In the details pane, right-click the group, and then click **Properties**.
4. On the **Members** tab, click **Add**.
5. In **Enter the object names to select**, type the name of the user, group, or computer that you want to add to the group, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- In addition to users and computers, membership in a particular group can include contacts and other groups.
- To add members to a group, you can click the members you want to add, click

  on the toolbar, and then type the name of the group into which they will be added. You can also drag a member object to a group, or right-click the object, and then click **Add to a group**.

- When you administer a domain, security principals in the parent domain or other trusted domains are not visible on the **Member Of** tab of a domain users properties. The only domain accounts that you can add or view are the present domain groups. Only domain groups in the present domain are shown, even if the member belongs to other trusted domain groups.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **dsmod group***GroupDN***-addmbr***MemberDN*

| Value | Description |
|---|---|
| *GroupDN* | Specifies the distinguished name of the group object to which you want to add the object. |
| *MemberDN* | Specifies the distinguished name of the object that you want to add to the group. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- In addition to users and computers, a group can contain contacts and other groups.

- To view the complete syntax for this command, at a command prompt, type:

  **dsmod group /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding Groups
Finding directory information
Group scope
Dsmod
Working with MMC console files
Find a group
Command-line reference A-Z

# Allow anonymous users to be members of the Everyone security group on a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow members of the Anonymous Logon group to be members of the Everyone group on a domain controller**

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Domain Security Policy**.
2. In the console tree, click **Security Options**.

   **Where?**

   - Security Settings/Local Policies/Security Options

3. In the details pane, right-click **Network access: Let Everyone permissions apply to anonymous users**, and then click **Properties**.
4. Select the **Define this policy setting** check box, and then click **Enabled**.

**Note**

- By default, members of the Anonymous Logon group are not included in the Everyone group.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using the Active Directory Installation Wizard
Special identities

# Convert a group to another group type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To convert a group to another group type**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click the folder that contains the group that you want to convert to another group type.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder that contains the group*

3. In the details pane, right-click the group, and then click **Properties**.
4. On the **General** tab, under **Group type**, click the group type.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To convert a group, the domain functional level must be set to Windows 2000 native or higher. Groups cannot be converted when the domain functional level is set to Windows 2000 mixed.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmod group** *GroupDN* **-secgrp** {**yes**|**no**}

| Value | Description |
|---|---|
| *GroupDN* | Specifies the distinguished name of the group object for which you want to change the group type. |
| {**yes**\|**no**} | Specifies that the group type is set to security group (**yes**) or distribution group (**no**). |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the

Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To convert a group, the domain functional level must be set to Windows 2000 native or higher. Groups cannot be converted when the domain functional level is set to Windows 2000 mixed.

- To view the complete syntax for this command, at a command prompt, type:

  **dsmod group /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsmod
Group types
Groups
Understanding Groups
Command-line reference A-Z
Domain and forest functionality
Working with MMC console files

# Change group scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change group scope**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click the folder that contains the group for which you want to change the group scope.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder that contains the group*

3. In the details pane, right-click the group, and then click **Properties**.
4. On the **General** tab, under **Group scope**, click the group scope.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- You can only change group scopes when the domain functional level is set to Windows 2000 native or higher.
- Changing the scope of a group from universal to domain local can only be done on a global catalog server. An error is returned if the domain controller is not a global catalog server.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmod group***GroupDN***-scope***L*|*G*|*U*

| Value | Description |
|---|---|
| *GroupDN* | Specifies the distinguished names (DNs) of the group object to which the scope will be changed. |

| | |
|---|---|
| *L\|G\|U* | Specifies that the scope of the group be set to local, global or universal. If the domain is still in Windows 2000 mixed, then the universal scope will not be supported. Also, it is not possible to convert a domain local group to global group or vice versa. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can only change group scopes when domain functional level is set to Windows 2000 native or higher.

- Changing the scope of a group from universal to domain local can only be done on a global catalog server. An error is returned if the domain controller is not a global catalog server.

- To view the complete syntax for this command, at a command prompt, type:

  **dsmod group /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group scope
Groups
Directory service command-line tools
Command-line reference A-Z
Domain and forest functionality
Working with MMC console files

# Delete a group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a group**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click the folder that contains the group you want to delete.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder that contains the group*

3. In the details pane, right-click the group, and then click **Delete**.

### Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Deleting a group is a permanent operation.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsrm***GroupDN*

| Value | Description |
|---|---|
| *GroupDN* | Specifies the distinguished name of the group object to be deleted. |

### Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command**

**prompt**.

- Deleting a group is a permanent operation.
- To view the complete syntax for this command, at a command prompt, type:

  **dsrm /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsrm
Working with MMC console files
Find a group
Command-line reference A-Z

# Find groups in which a user is a member

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find groups in which a user is a member**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Users

   Or, click the folder that contains the user account.
3. In the details pane, right-click a user account, and then click **Properties**.
4. Click the **Member Of** tab.

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- The **Member Of** tab for a user displays a list of groups in the domain where the user's account is located. Active Directory does not display groups that reside in trusted domains where the user is a member.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsget user***UserDN***-memberof**

| Value | Description |
|---|---|
| *UserDN* | Specifies the distinguished name of the user object for which you want to display group membership. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **dsget user /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsget
Working with MMC console files
Find a user account
Find a group
Command-line reference A-Z

# Assign user rights to a group in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign user rights to a group in Active Directory**

1. Open Domain Controller Security Policy.
2. In the console tree, click **User Rights Assignment**.

   **Where?**

   - Security Settings/Local Policies/User Rights Assignment

3. In the details pane, double-click the user right that you want to assign.
4. Click **Add User or Group**.

   If the button appears dimmed, select the **Define these policy settings** check box.

5. Type the name of the group to which you want to assign this right.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Domain Controller Security Policy, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Domain Controller Security Policy**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group types
Default groups
User Rights Assignment

# Manage Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage computers**

- Create a new computer account
- Add a computer account to a group
- Delete a computer account
- Manage a remote computer
- Move a computer account
- Reset a computer account
- Disable or enable a computer account

# Create a new computer account

Updated: June 5, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

**To create a new computer account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click **Computers**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Computers

   Or, right-click the folder in which you want to add the computer.
3. Point to **New**, and then click **Computer**.
4. Type the computer name.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- By default, members of the Account Operators group can create computer accounts in the Computers container and in new organizational units.
- By default, Authenticated Users in a domain are assigned the **Add workstations to a domain** user right and can create up to 10 computer accounts in the domain. For more information, see Related Topics.
- There are two additional ways to give a user or group permission to add a computer to the domain: Use a Group Policy object to assign the Add Computer User permission, or, on the organizational unit, assign the user or group the Create Computer Objects permission.
- To view or change a computer name and to modify the domain in which it belongs, click **Start**, right-click **My Computer**, click **properties**, click the **Computer Name** tab, and then click **Change**.
- If the computer using the account you are creating is a pre-Windows 2000 computer, select the **Assign this computer account as a pre-Windows 2000 computer** check box.

- If the computer using the account you are creating is a Windows NT backup domain controller, select **Assign this computer account as a backup domain controller** check box.

- By default, a newly created computer account is assigned an "Account Ops-FC" access control entry (ACE) that gives members of the Account Operators group full control over that computer account. If a server that is represented by this computer account is promoted to a domain controller, the computer account retains this "Account Ops-FC" ACE, and therefore, members of the Account Operators group will have full control on this domain controller, which is not a recommended configuration. For example, members of Account Operators group will be able to reset this domain controller's password. Because the Account Operators group has significant power in the domain, we recommend that you add members to it with caution. For a detailed description of the Account Operators group, see Default groups (http://go.microsoft.com/fwlink/?LinkID=131422). To modify permissions for Account Operators on a computer account, you can use Active Directory Users and Computers snap-in and complete the following steps:

  1. To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

  2. In the console tree, right-click the computer account that represents the server that you want to promote to a domain controller status, and then click **Properties**.

  3. On the **Security** tab, select **Account Operators** in the **Group or user names** list, and then modify permissions according to the specifications of your environment.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **dsadd computer** *ComputerDN*

| Value | Description |
|---|---|
| *ComputerDN* | Specifies the distinguished name of the computer you want to add. The distinguished name specifies the directory location. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- By default, members of the Account Operators group can create computer accounts in the Computers container and in new organizational units.

- By default, Authenticated Users in a domain are assigned the **Add workstations to a domain** user right and can create up to 10 computer accounts in the domain. For more information, see Related Topics.

- There are two additional ways to give a user or group permission to add a computer to the domain: Use a Group Policy object to assign the Add Computer User permission, or, on the organizational unit, assign the user or group the Create Computer Objects permission.

- To view the complete syntax for this command, at a command prompt, type:

**dsadd computer /?**

- To modify the properties of a computer account, use the **dsmod computer** command.

- By default, a newly created computer account is assigned an "Account Ops-FC" access control entry (ACE) that gives members of the Account Operators group full control over that computer account. If a server that is represented by this computer account is promoted to a domain controller, the computer account retains this "Account Ops-FC" ACE, and therefore, members of the Account Operators group will have full control on this domain controller, which is not a recommended configuration. For example, members of Account Operators group will be able to reset this domain controller's password. Because the Account Operators group has significant power in the domain, we recommend that you add members to it with caution. For a detailed description of the Account Operators group, see Default groups (http://go.microsoft.com/fwlink/?LinkID=131422). To modify permissions for Account Operators on a computer account, you can use Active Directory Users and Computers snap-in and complete the following steps:

  1. To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

  2. In the console tree, right-click the computer account that represents the server that you want to promote to a domain controller status, and then click **Properties**.

  3. On the **Security** tab, select **Account Operators** in the **Group or user names** list, and then modify permissions according to the specifications of your environment.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsadd
Find a computer account
Active Directory integration
Object names
Domain controllers
Add workstations to domain
Privileges

# Add a computer account to a group

Updated: June 5, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

**To add a computer account to a group**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Computers**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Computers

   Or, click the folder in which the computer is located.
3. In the details pane, right-click the computer, and then click **Properties**.
4. On the **Member Of** tab, click **Add**.
5. In **Enter the object names to select**, type the name of a group that you want this computer to be a member of, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- When viewing users, computers, and groups as containers, you can add a computer to a group by dragging it onto the desired group. For more information about view options in Active Directory Users and Computers, see Related Topics.
- By adding a computer to a group, you can assign permissions to all of the computer accounts in that group and filter Group Policy settings on all accounts in that group.

**Using a command line**

1. Open Command Prompt.
2. Type:

```
Dsmod group GroupDN -addmbr ComputerDN
```

For example, to add computer1 to Group A in the Contoso domain, type

```
Dsmod "CN=Group A,OU=Domain Groups,DC=Contoso,DC=com" -addmbr
"CN=Computer1,OU=Domain Computers,DC=Contoso,DC=com"
```

| Value | Description |
|---|---|
| *GroupDN* | Specifies the distinguished names of the group object to which you want to add the computer object. |
| *ComputerDN* | Specifies the distinguished name of the computer object to be added to the group. The distinguished name specifies the directory location. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- By adding a computer to a group, you can assign permissions to all of the computer accounts in that group and filter Group Policy settings on all accounts in that group.

- To view the complete syntax for this command, at a command prompt, type:

  **dsmod group /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsmod
Find a computer account
Command-line reference A-Z
Working with MMC console files
Managing Active Directory from MMC

# Delete a computer account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a computer account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Computers**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Computers

   Or, click the folder in which the computer is located.

3. In the details pane, right-click the computer, and then click **Delete**.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Once a computer account has been deleted, all permissions and memberships associated with that computer account are permanently deleted. Because the security ID (SID) for each account is unique, a new computer account with the same name as a previously deleted computer account does not inherit the permissions and memberships of the previously deleted account. To duplicate a deleted computer account, all permissions and memberships must be manually recreated.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsrm computer***ComputerDN*

| Value | Description |
|---|---|
| *ComputerDN* | Specifies the distinguished name of the computer you want to add. The distinguished name specifies the directory location. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Once a computer account has been deleted, all permissions and memberships associated with that computer account are permanently deleted. Because the security ID (SID) for each account is unique, a new computer account with the same name as a previously deleted computer account does not inherit the permissions and memberships of the previously deleted account. To duplicate a deleted computer account, all permissions and memberships must be manually recreated.

- To view the complete syntax for this command, at a command prompt, type:

  **dsrm computer /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find a computer account
Join a workgroup

# Manage a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage a remote computer**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Computers**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Computers

   Or, click the folder that contains the computer you want to manage.
3. In the details pane, right-click the computer, and then click **Manage**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Clicking **Manage** starts Computer Management from which you can administer remote computers. You must have administrative credentials on the local computer to view certain information or to modify computer properties using Computer Management.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find a computer account
Computer Management Overview

# Move a computer account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move a computer account**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Computers**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Computers

   Or, click the folder that contains the computer that you want to move.
3. In the details pane, right-click the computer, and then click **Move**.
4. In the **Move** dialog box, click the domain node.
5. Click the folder to which you want to move the computer, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- You cannot use Active Directory Users and Computers to move computer accounts between domains. To move a computer account between domains, use Movetree, one of the Active Directory support tools. For more information, see Related Topics.
- Members of the Account Operators group can move computer accounts to organizational units but not to default containers, such as Builtin or Computers. However, Account Operators cannot move computer accounts into the Domain Controllers organizational unit but can move computer accounts from the Domain Controllers organizational unit.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find a computer account
Active Directory support tools

# Reset a computer account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reset a computer account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Computers**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Computers

   Or, click the folder that contains the computer that you want to reset.
3. In the details pane, right-click the computer, and then click **Reset Account**.

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmod computer***ComputerDN***-reset**

| Value | Description |
|---|---|
| *ComputerDN* | Specifies the distinguished names of one or more computer objects that you want to reset. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command**

**prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **dsmod computer /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsmod
Working with MMC console files
User and computer accounts
Command-line reference A-Z

# Disable or enable a computer account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable or enable a computer account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Computers**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/Computers

   Or, click the folder that contains the computer that you want to enable or disable.
3. In the details pane, right-click the desired computer, and then do one of the following:

   - To disable, click **Disable Account**.
   - To enable, click **Enable Account**.

## Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- When you disable a computer account, the computer cannot authenticate to the domain until it has been enabled.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmod computer***ComputerDN***-disabled** {**yes**|**no**}

| Value | Description |
|---|---|
| *ComputerDN* | Specifies the distinguished names of the computer object that you want to disable or enable. |

| | |
|---|---|
| {**yes**\|**no**} | Specifies whether the computer is disabled for log on (**yes**) or not (**no**). |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- When you disable a computer account, the computer cannot authenticate to the domain until it has been enabled.

- To view the complete syntax for this command, at a command prompt, type:

  **dsmod computer /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsmod
Working with MMC console files
User and computer accounts
Command-line reference A-Z

# Publish Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Publish resources**

- Publish a shared folder
- Manually publish a printer in Active Directory

# Publish a shared folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To publish a shared folder**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the folder in which you want to add a shared folder.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder*

3. Point to **New**, and then click **Shared Folder**.
4. Type the name of the folder and the network path.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Shared Folders overview

# Manually publish a printer in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manually publish a printer in Active Directory**

1. Open **Active Directory Users and Computers**.
2. In **Name**, right-click the container object folder in which you want to publish the printer.
3. Point to **New**, and then click **Printer**.
4. Type the UNC name of the printer that you want to publish in the directory.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To perform this procedure, you must be a member of the Print Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Group Policy settings for printers determine whether printers published with a print server running Windows NT 4.0 or an earlier operating system should be pruned if the print server becomes unavailable.
- You can also use the **Pubprn.vbs** script, provided in the System32 folder, to publish printers with print servers running Windows NT 4.0 or an older operating system.
- The printer you want to publish must be shared.
- You must have directory service administration rights to publish printers.
- You must have the Manage Printers permission for the printer you want to share or to publish.
- You can create printer names that contain spaces, special characters, and more than eight characters. However, some clients do not recognize or correctly handle such names. The entire qualified name of a shared printer (for example, \\PRINTER2\PSCRIPT) must contain 32 or fewer characters.

**See Also**
**Concepts**
Group Policy and printing
Publish a printer in Active Directory

# Manage Domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage domains**

- Create a new forest
- Create a new domain tree
- Create a new child domain
- Manage a different domain
- Remove a domain
- Manage the domain using a different domain controller
- Add user principal name suffixes
- Delete extinct domain metadata
- Raise the domain functional level
- Assign, change, or remove permissions on Active Directory objects or attributes
- Modify the AdminSDHolder container
- Manage Organizational Units
- Manage Domain Controllers

# Create a new forest

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new forest**

1. Click **Start**, click **Run**, and then type **dcpromo** to start the Active Directory Installation Wizard.
2. On the **Operating System Compatibility** page, read the information and then click **Next**.

   If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.
3. On the **Domain Controller Type** page, click **Domain controller for a new domain**, and then click **Next**.
4. On the **Create New Domain** page, click **Domain in a new forest**, and then click **Next**.
5. On the **New Domain Name** page, type the full DNS name for the new domain, and then click **Next**.
6. On the **NetBIOS Domain Name** page, verify the NetBIOS name, and then click **Next**.
7. On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location, and then click **Next**.
8. On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location, and then click **Next**.
9. On the **DNS Registration Diagnostics** page, verify if an existing DNS server will be authoritative for this forest or, if necessary, choose to install and configure DNS on this server by clicking **Install and configure the DNS server on this computer, and set this computer to use this DNS server as its preferred DNS server**, and then click **Next**.
10. On the **Permissions** page, select one of the following:

    - **Permissions compatible with pre-Windows 2000 Server operating systems**
    - **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**
11. Review the **Summary** page, and then click **Next** to begin the installation.
12. Restart the computer.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- The server on which you install Active Directory using this procedure will be the first domain controller in the forest root domain.
- Before installing Active Directory, you will need to consider pre-Windows 2000 compatible security levels and identify the DNS name of the domain. For more information, see the checklist in Related Topics.

- The wizard options on the **Permissions** page affect application compatibility with computers running pre-Windows 2000 and Windows Server 2003 operating systems and are not related to domain functionality. For more information about permissions, see Related Topics.

- The Active Directory Installation Wizard allows Active Directory domain names up to 64 characters or up to 155 bytes. Although the limit of 64 characters is usually reached before the limit of 155 bytes, the opposite could be true if the name contains Unicode characters consuming three bytes. These limits do not apply to computer names.

- You cannot install Active Directory on a computer running Windows Server 2003, Web Edition, but you can join the computer to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Creating a new forest
Domains
Domain controllers
Using Smart Cards
Using the Active Directory Installation Wizard
Domain and forest functionality
Overview of Windows Server 2003, Web Edition

# Create a new domain tree

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new domain tree**

1. Click **Start**, click **Run**, and then type **dcpromo** to start the Active Directory Installation Wizard.

2. On the **Operating System Compatibility** page, read the information and then click **Next**.

   If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.

3. On the **Domain Controller Type** page, click **Domain controller for a new domain**, and then click **Next**.

4. On the **Create New Domain** page, click **Domain tree in an existing forest**.

5. On the **Network Credentials** page, type the user name, password, and user domain of the user account you want to use for this operation, and then click **Next**.

   The user account must be a member of the Enterprise Admins group.

6. On the **New Domain Tree** page, type the full DNS name for the new domain, and then click **Next**.

7. On the **NetBIOS Domain Name** page, verify the NetBIOS name, and then click **Next**.

8. On the **Database and Log Folders** page, type the location in which to install the database and log folders, or click **Browse** to choose a location, and then click **Next**.

9. On the **Shared System Volume** page, type the location in which to install the Sysvol folder, or click **Browse** to choose a location, and then click **Next**.

10. On the **DNS Registration Diagnostics** page, verify if an existing DNS server will be authoritative for this forest or, if necessary, choose to install and configure DNS on this server by clicking **Install and configure the DNS server on this computer, and set this computer to use this DNS server as its preferred DNS server**, and then click **Next**.

11. On the **Permissions** page, select one of the following:

    - **Permissions compatible with pre-Windows 2000 Server operating systems**

    - **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**

12. On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password that you want to assign to the Administrator account for this server, and then click **Next**.

    Use this password when starting the computer in Directory Services Restore Mode.

13. Review the **Summary** page, and then click **Next** to begin the installation.

14. Restart the computer.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Before installing Active Directory, you will need to consider pre-Windows 2000 compatible security levels and identify the DNS name of the domain. For more information, see the checklist in Related Topics.

- When a new domain tree is created in an existing forest, a two-way, transitive tree root trust is established by default.

- The wizard options on the **Permissions** page affect application compatibility with computers running pre-Windows 2000 and Windows Server 2003 operating systems and are not related to domain functionality. For more information about permissions, see Related Topics.

- You can also use a smart card to verify administrative credentials. For more information about smart cards, see Related Topics.

- The Active Directory Installation Wizard allows Active Directory domain names up to 64 characters or up to 155 bytes. Although the limit of 64 characters is usually reached before the limit of 155 bytes, the opposite could be true if the name contains Unicode characters consuming three bytes. These limits do not apply to computer names.

- You cannot install Active Directory on a computer running Windows Server 2003, Web Edition, but you can join the computer to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Creating a new domain tree
Domains
Domain controllers
Using Smart Cards
Using the Active Directory Installation Wizard
Domain and forest functionality
Overview of Windows Server 2003, Web Edition

# Create a new child domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new child domain**

1. Click **Start**, click **Run**, and then type **dcpromo** to start the Active Directory Installation Wizard.

2. On the **Operating System Compatibility** page, read the information and then click **Next**.

   If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.

3. On the **Domain Controller Type** page, click **Domain controller for a new domain**, and then click **Next**.

4. On the **Create New Domain** page, click **Child domain in an existing domain tree**, and then click **Next**.

5. On the **Network Credentials** page, type the user name, password, and user domain of the user account you want to use for this operation, and then click **Next**.

   The user account must be a member of the Enterprise Admins group.

6. On the **Child Domain Installation** page, verify the parent domain and type the new child domain name, and then click **Next**.

7. On the **NetBIOS Domain Name** page, verify the NetBIOS name, and click **Next**.

8. On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location, and then click **Next**.

9. On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location, and then click **Next**.

10. On the **DNS Registration Diagnostics** page, verify the DNS configuration settings are accurate, and then click **Next**.

11. On the **Permissions** page, select one of the following:

    - **Permissions compatible with pre-Windows 2000 server operating systems**

    - **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**

12. On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password that you want to assign to the Administrator account for this server, and then click **Next**.

    Use this password when starting the computer in Directory Services Restore Mode.

13. Review the **Summary** page, and then click **Next** to begin the installation.

14. Restart the computer.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in

Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- The server on which you install Active Directory using this procedure will be the first domain controller in a new child domain.

- Before installing Active Directory, you will need to consider pre-Windows 2000 compatible security levels and identify the DNS name of the domain. For more information, see the checklist in Related Topics.

- When a child domain is added to an existing tree domain, a two-way, transitive parent and child trust is established by default.

- The wizard options on the **Permissions** page affect application compatibility with computers running pre-Windows 2000 and Windows Server 2003 operating systems and are not related to domain functionality. For more information about permissions, see Related Topics.

- You can also use a smart card to verify administrative credentials. For more information about smart cards, see Related Topics.

- The Active Directory Installation Wizard allows Active Directory domain names up to 64 characters or up to 155 bytes. Although the limit of 64 characters is usually reached before the limit of 155 bytes, the opposite could be true if the name contains Unicode characters consuming three bytes. These limits do not apply to computer names.

- You cannot install Active Directory on a computer running Windows Server 2003, Web Edition, but you can join the computer to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Creating a new child domain
Domains
Domain controllers
Using Smart Cards
Using the Active Directory Installation Wizard
Domain and forest functionality
Trusts
Overview of Windows Server 2003, Web Edition

# Manage a different domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage a different domain**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click **Active Directory Users and Computers**, and then click **Connect to Domain**.
3. Type the domain name.

   Or, click **Browse**, and then select a domain from the list.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials. However, you must have the appropriate administrative credentials to manage the domain you select.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- If you select the **Save this domain setting for the current console** option, Active Directory Users and Computers will attempt to connect to the domain each time it is opened.
- The domain name includes all of the domain's parent domains. For example, the domain *widgets* could have a full name of *widgets.microsoft.com*.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Domain controllers

# Remove a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a domain**

1. On the last domain controller in the domain, click **Start**, and then click **Run**.
2. In **Open**, type **dcpromo** to open the Active Directory Installation Wizard, and then click **Next**.
3. On the **Remove Active Directory** page, select the **This server is the last domain controller in the domain** check box, and then continue to follow the wizard.

## Caution

- Removing a domain will result in the permanent loss of any data contained in that domain. This includes all user, group, and computer accounts.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To remove an Active Directory domain, you must first demote all other domain controllers associated with the domain. For information about demoting domain controllers, see Related Topics.
- If the domain has child domains, it cannot be deleted.
- If this is the last domain in the forest, demoting this domain controller will also delete the forest.
- If the domain controller currently holds the domain naming master or schema master operations roles, transfer the role to domain controller in a different domain before removing the domain. For information about locating operations masters and transferring operations master roles, see Related Topics.
- If the domain controller holds the last replica of one or more application directory partitions, you must first remove the last application directory partition replicas from the domain controller before you can delete the domain. You can use the Active Directory Installation Wizard to remove all application directory partition replicas from the domain controller or you can manually remove the partitions using the Ntdsutil command-line tool. For information about application directory partitions, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domains
Enable or disable a global catalog
Operations master roles
Domain controllers
Application directory partitions

Add or remove an application directory partition replica
Identify operations master roles
Transfer operations master roles

# Manage the domain using a different domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage the domain using a different domain controller**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click **Active Directory Users and Computers**, and then click **Connect to Domain Controller**.
3. Click a domain controller in the list.

   Or, in **Enter the name of another domain controller**, type the name of a domain controller.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials. However, you must have the appropriate administrative credentials to manage the domain you select.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- By default, domain controllers are created in the Domain Controllers organizational unit when they are first configured.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find a computer account

# Add user principal name suffixes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding user principal name suffixes**
**To add user principal name suffixes**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click **Active Directory Domains and Trusts**, and then click **Properties**.
3. On the **UPN Suffixes** tab, type an alternative user principal name (UPN) suffix for the forest, and then click **Add**.
4. Repeat step 3 to add additional alternative UPN suffixes.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- UPN suffixes should conform to DNS conventions for valid characters and syntax.
- UPN suffixes are stored in the uPNSuffixes property of the following LDAP leaf: LDAP://cn=partitions,cn=configuration,dc=<*domain*>,dc=<*domainsuffix*>.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Active Directory naming
Unicode character support
Routing name suffixes across forests

# Delete extinct domain metadata

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete extinct domain metadata**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command, type:

   **metadata cleanup**
4. At the metadata cleanup command, type:

   **connection**
5. At the connection command, type:

   **connect to server***Server*
6. At the connection command, type:

   **quit**
7. At the metadata cleanup command, type:

   **select operation target**
8. At the select operation target command, type:

   **list sites**
9. A numbered list of sites will be displayed, type:

   **select site***SiteNumber*
10. At the select operation target command, type:

    **list domains in site**
11. A numbered list of domains in the selected site will be displayed, type:

    **select domain***DomainNumber*

    **quit**
12. At the metadata cleanup command, type:

    **remove selected domain**

| Value | Description |
| --- | --- |
| *Server* | The DNS name of a domain controller. |
| *SiteNumber* | The number in the displayed list that corresponds to the site that you want to connect to. |
| *DomainNumber* | The number in the displayed list that corresponds to the domain that you want to clean up. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To clean metadata you must first select the operation target. When selecting the operation target for a server, you must first select the site and then the domain.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ntdsutil
Administration tools for the Active Directory schema
Managing Active Directory from the command line

# Raise the domain functional level

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To raise the domain functional level**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain for which you want to raise functionality, and then click **Raise Domain Functional Level**.
3. In **Select an available domain functional level**, do one of the following:
   - To raise the domain functional level to Windows 2000 native, click **Windows 2000 native**, and then click **Raise**.
   - To raise domain functional level to Windows Server 2003, click Windows Server 2003, and then click **Raise**.

**Caution**

- If you have or will have any domain controllers running Windows NT 4.0 and earlier, then do not raise the domain functional level to Windows 2000 native. Once the domain functional level is set to Windows 2000 native, it cannot be changed back to Windows 2000 mixed.
- If you have or will have any domain controllers running Windows NT 4.0 and earlier or Windows 2000, then do not raise the domain functional level to Windows Server 2003. Once the domain functional level is set to Windows Server 2003, it cannot be changed back to Windows 2000 mixed or Windows 2000 native.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group in the domain for which you want to raise functionality or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- You can also raise the domain functional level by right-clicking a domain displayed in Active Directory Users and Computers, and then clicking **Raise Domain Functional Level**.
- The current domain functional level is displayed under **Current domain functional level** in the **Raise Domain Functional Level** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domain and forest functionality
Raising domain and forest functional levels

# Assign, change, or remove permissions on Active Directory objects or attributes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign, change, or remove permissions on Active Directory objects or attributes**

1. Open Active Directory Users and Computers.
2. On the **View** menu, select **Advanced Features**.
3. Right-click the object for which you want to assign, change, or remove permissions, and then click **Properties**.
4. On the **Security** tab, click **Advanced** to view all of the permission entries that exist for the object.
5. Do one or more of the following:
   - To assign new permissions on an object or attribute, click **Add**. Type the name of the group, computer, or user that you want to add, and then click **OK**. In the **Permission Entry for *ObjectName*** dialog box, on the **Object** and **Properties** tabs, select or clear the **Allow** or **Deny** check boxes, as appropriate.
   - To change existing permissions on an object or attribute, click a permission entry, and then click **Edit**. On the **Object** and **Properties** tabs, select or clear the **Allow** or **Deny** check boxes, as appropriate.
   - To remove existing permissions from an object or attribute, click a permission entry, and then click **Remove**.

### Important

- Before adding access control permissions on Active Directory objects or properties, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Access control overview
Permissions
Best practices for assigning permissions on Active Directory objects

# Modify the AdminSDHolder container

Updated: April 17, 2007

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This procedure provides steps for modifying the AdminSDHolder container object, which is a built-in container object for each Active Directory domain that has a specific access control list (ACL) set on it. You can use the Dsacls tool to help prevent administrators from making unauthorized changes to the membership of administrative groups in an Active Directory domain.

| **Important** |
| --- |
| This procedure will not completely prevent unauthorized changes to membership of administrative groups in an Active Directory domain. It is always recommended that you have auditing in place that alerts you when any changes are made to the membership of administrative groups in your Active Directory domain. For more information, see "Additional references." |

**To modify the AdminSDHolder container**

1. Download and install the Windows Server 2003 Service Pack 1 (SP1) Support Tools.

   Dsacls.exe is available as part of the Windows Support Tools. For more information, see article 89277 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=44321).

2. In the domain that you want to modify, create a new domain local group, for example, AdminDenyGrp.

   For more information about creating a new domain local group, see Create a new group.

3. For each administrator that you want to help prevent from making changes to administrative groups in the domain, add the administrator's name to the group that you created.

   For more information about adding a member to a group, see Add a member to a group.

4. Open a command prompt with your administrative credentials.

5. At the command prompt, type the following, and then press ENTER:

   ```
   dsacls <AdmSDHolderDN> /D <DomName>\<AdminDenyGrp>:WP;"<Attribute>"
   ```

6. Apply the parameters in the following table as appropriate for your deployment.

| Parameter | Description |
| --- | --- |
| <AdmSDHolderDN> | Specifies the distinguished name of the AdminSDHolder object for the domain. For example, where the domain is CORPHQ, as in corphq.contoso.com, the distinguished name is the following:<br><br>CN=AdminSDHolder,CN=System,DC=corphq,C=contoso,DC=com |
| <DomName> | Specifies the name of the trusting domain. For example, where the full distinguished name is corphq.contoso.com, the domain name is specified as "CORPHQ" here. |
| <AdminDenyGrp> | Specifies the actual name of the group that you created in step 2 in this procedure. |
| <Attribute> | Specifies the attribute on the AdminSDHolder container that will be modified when the command is run. For this procedure, specify the value as "member" (including the quotation marks). This is the name of the attribute that Active Directory uses to store the membership data that is associated with a group. |

When you run this command, it modifies the AdminSDHolder container for the domain that you specified to deny the specified group the ability to write the **member** attribute for group objects.

For example, the following command line uses actual values that might appear in the **dsacls** command in which **/D** indicates that you are adding a DENY access control entry (ACE), and **WP** indicates the **Write** property for the **member** attribute.

```
dsacls CN=AdminSDHolder,CN=System,DC=corphq,C=contoso,DC=com /D
CORPHQ\AdminDenyGrp:WP;"member"
```

To view the complete syntax for this command, and for information about entering user account information, at a command prompt, type the following, and then press ENTER:

```
dsacls /?
```

**Additional considerations**

| **Important** |
| --- |
| Before you add access control permissions on Active Directory objects or properties, see "Additional references." |

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To complete this procedure, you must have the Dsacls command-line tool installed.

- You can use this procedure to put other restrictions in place for protected objects. For example, the following command puts DENY-Delete & DENY-Modify Permissions ACEs on protected objects for the specified user or group object (ServiceAccount):

  ```
  DSACLS DNofAdminSDHolderContainer /D CONTOSO\ServiceAccount:SDWD
  ```

  The following command puts a DENY-Reset Password ACE on protected objects for the specified user or group object (ServiceAccount):

  ```
  DSACLS DNofAdminSDHolderContainer /D CONTOSO\ServiceAccount:CA;"Reset Password"
  ```

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**Additional references**

- Access control overview

- Permissions

- Best practices for assigning permissions on Active Directory objects

- HOW TO: Audit Active Directory Objects in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkID=86852)

- Auditing Security Events Best practices (http://go.microsoft.com/fwlink/?LinkID=20182)

# Manage Organizational Units

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage organizational units**

- Create a new organizational unit
- Delegate control of an organizational unit
- Delete an organizational unit
- Move an organizational unit

# Create a new organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new organizational unit**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the folder in which you want to add an organizational unit.

   **Where?**

   - Active Directory Users and Computers/*domain node/folder*

3. Point to **New**, and then click **Organizational Unit**.
4. Type the name of the organizational unit.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To add an organizational unit, you can also click the domain node or folder in which you want to add the organizational unit, and then click

  on the toolbar.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsadd ou***OrganizationalUnitDN*

| Value | Description |
|---|---|
| *OrganizationalUnitDN* | Specifies the distinguished name of the organizational unit to be added. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using

Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **dsadd ou /?**

- To modify the properties of an organizational unit, use the **dsmod ou** command.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsadd
Working with MMC console files
Organizational units
Find an organizational unit
Object names
Command-line reference A-Z
Dsmod

# Delegate control of an organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate control of an organizational unit**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the organizational unit for which you want to delegate control.

   **Where?**

   - Active Directory Users and Computers/*domain node/organizational unit*

3. Click **Delegate control** to start the Delegation of Control Wizard, and then follow the instructions.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find an organizational unit
Delegating administration

# Delete an organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an organizational unit**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the organizational unit in which you want to delete.

   **Where?**

   - Active Directory Users and Computers/*domain node*/*organizational unit*

3. Click **Delete**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- If the selected organizational unit contains other objects, Active Directory Users and Computers prompts you to either continue or cancel the deletion. If you continue the deletion, all objects within the organizational unit will also be deleted.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsrm***OrganizationalUnitDN*

| Value | Description |
|---|---|
| *OrganizationalUnitDN* | Specifies the distinguished name of the organizational unit to be deleted. |

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Attempting to delete an organizational unit containing other objects will result in an error, unless the **-subtree** option is used. With the **-subtree** option, **dsrm** deletes both the selected organizational unit and any objects it contains.

- To view the complete syntax for this command, at a command prompt, type:

  **dsrm /?**

  or see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsrm
Working with MMC console files
Find an organizational unit
Command-line reference A-Z

# Move an organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move an organizational unit**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the organizational unit that you want to move.

   **Where?**

   - Active Directory Users and Computers/*domain node/organizational unit*

3. Click **Move**, and then click the folder to which you want to move the organizational unit.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Active Directory Users and Computers can only move organizational units within the same domain. To move organizational units between domains, use Movetree. For more information about Movetree and other Active Directory support tools, see Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmov***OuDN* [**-newparent** ParentDN] **-pwd** {*Password*|*****}

| Value | Description |
|---|---|
| *OuDN* | Specifies the distinguished name of the organizational unit object to be moved. |
| *ParentDN* | Specifies the new location to which you want to move the organizational unit object. The new location is specified as the distinguished name of the new parent node. |
| *Pwd* | Specifies to use either a password or a * to log on to a remote server. If you type *, then you are prompted for a password. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- Active Directory Users and Computers can only move organizational units within the same domain. To move organizational units between domains, use Movetree. For more information about Movetree and other Active Directory support tools, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Organizational units
Find an organizational unit
Active Directory support tools

# Manage Domain Controllers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage domain controllers**

- Create an additional domain controller
- Demote a domain controller
- Rename a domain controller
- Disable signed or encrypted LDAP traffic

# Create an additional domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an additional domain controller**

1. Click **Start**, click **Run**, and then type **dcpromo /adv** to open the Active Directory Installation Wizard with the option to create an additional domain controller from restored backup files.

2. On the **Operating System Compatibility** page, read the information and then click **Next**.

   If this is the first time you have installed Active Directory on a server running Windows Server 2003, click **Compatibility Help** for more information.

3. On the **Domain Controller Type** page, click **Additional domain controller for an existing domain**, and then click **Next**.

4. On the **Copying Domain Information** page, do one of the following:

   - Click **Over the network**, and then click **Next**.

   - Click **From these restored backup files**, and type the location of the restored backup files, or click **Browse** to locate the restored files, and then click **Next**.

5. On the **Network Credentials** page, type the user name, password, and user domain of the user account you want to use for this operation, and then click **Next**.

   The user account must be a member of the Domain Admins group for the target domain.

6. On the **Database and Log Folders** page, type the location in which you want to install the database and log folders, or click **Browse** to choose a location, and then click **Next**.

7. On the **Shared System Volume** page, type the location in which you want to install the Sysvol folder, or click **Browse** to choose a location, and then click **Next**.

8. On the **Directory Services Restore Mode Administrator Password** page, type and confirm the password that you want to assign to the Administrator account for this server, and then click **Next**.

   Use this password when starting the computer in Directory Services Restore Mode.

9. Review the **Summary** page, and then click **Next** to begin the installation.

10. Restart the computer.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- The **/adv** switch is only necessary when you want to create a domain controller from restored backup files. It is not required when creating an additional domain controller over the network.

- In step 3, when choosing the option to copy domain information over the network, all directory data for the domain in

which this domain controller will be a member will be copied over your network connection. You will have the option to cancel non-critical replication, if necessary.

- In step 3, when choosing the option to copy domain information from restored backup files, you will need to first back up the System State data of a domain controller running Windows Server 2003 from the domain in which this member server will become an additional domain controller. Then, the System State backup must be restored locally on the server on which you are installing Active Directory. To do this using Backup, choose the option **Restore files to: Alternate location**. For more information about restoring backups, see Related Topics.

- If a domain controller that was backed up contained an application directory partition, the application directory partition will not be restored on the new domain controller. For information about how to manually create an application directory partition on a new domain controller, see Related Topics.

- If the domain controller from which you restored the System State data was a global catalog, you will have the option to make this new domain controller a global catalog.

- You can also use a smart card to verify administrative credentials. For more information about smart cards, see Related Topics.

- You cannot install Active Directory on a computer running Windows Server 2003, Web Edition, but you can join the computer to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Creating an additional domain controller in an existing domain
Using the Active Directory Installation Wizard
Domain controllers
Using Smart Cards
Create or delete an application directory partition
Restore System State data
Authoritative, primary, and normal restores
Overview of Windows Server 2003, Web Edition

# Demote a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To demote a domain controller**

1. On a domain controller, click **Start**, and then click **Run**.
2. In **Open**, type **dcpromo** to open the Active Directory Installation Wizard, and then click **Next**.
3. On the **Remove Active Directory** page, click **Next**, and then continue to follow the wizard.

### Caution

- Before you complete this procedure, verify that this domain controller is not the only global catalog and that it does not hold an operations master role.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- If this domain controller is a global catalog, ensure that another global catalog is available to users before demoting it. For information about configuring domain controllers to host the global catalog, see Related Topics.
- If this domain controller currently holds one or more operations master roles, transfer the operations master roles to another domain controller before demoting it. For information about locating operations masters and transferring operations master roles, see Related Topics.
- If this domain controller is the last domain controller in the domain, demoting this domain controller will remove this domain from the forest. If this is the last domain in the forest, demoting this domain controller will also delete the forest. For information about removing domains and forests, see Related Topics.
- If this domain controller holds the last replica of one or more application directory partitions, you must first remove the last application directory partition replicas from this domain controller before you can demote it. You can use the Active Directory Installation Wizard to remove all application directory partition replicas from this domain controller or you can manually remove them using the Ntdsutil command-line tool. For information about application directory partitions, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domain controllers
Enable or disable a global catalog
Operations master roles
Domains
Application directory partitions
Add or remove an application directory partition replica

# Rename a domain controller

Updated: April 23, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2


To rename a domain controller

To rename a domain controller in a domain that contains a single domain controller


**To rename a domain controller**

1. Open Command Prompt.

2. Type:

   **netdom computername** *CurrentComputerName***/add:***NewComputerName*

   This command will update the service principal name (SPN) attributes in Active Directory for this computer account and register DNS resource records for the new computer name. The SPN value of the computer account must be replicated to all domain controllers for the domain and the DNS resource records for the new computer name must be distributed to all the authoritative DNS servers for the domain name. If the updates and registrations have not occurred prior to removing the old computer name, then some clients may be unable to locate this computer using the new or old name.

3. Ensure the computer account updates and DNS registrations are completed, then type:

   **netdom computername** *CurrentComputerName* **/makeprimary:***NewComputerName*

4. Restart the computer.

5. From the command prompt, type:

   **netdom computername** *NewComputerName* **/remove:***OldComputerName*

| Value | Description |
|---|---|
| *CurrentComputerName* | The current, or primary, computer name or IP address of the computer you are renaming. |
| *NewComputerName* | The new name for the computer. The *NewComputerName* must be a fully qualified domain name (FQDN). The primary DNS suffix specified in the FQDN for *NewComputerName* must be the same as the primary DNS suffix of *CurrentComputerName* or it must be contained in the list of allowed DNS suffixes specified in the **msDS-AllowedDNSSuffixes** attribute of the **domainDns** object. |
| *OldComputerName* | The old name of renamed computer. The *OldComputerName* must be a fully qualified domain name (FQDN). |

**Important**

- To rename a domain controller using the Netdom tool, the domain functional level must be set to Windows Server 2003. For more information, see Related Topics.

- Renaming a domain controller requires that you first provide a FQDN as a new computer name for the domain controller. All of the computer accounts for the domain controller must contain the updated SPN attribute and all the authoritative DNS servers for the domain name must contain the host (A) resource record for the new computer name. Both the old and new computer names are maintained until you remove the old computer name. This ensures that there will be no interruption in the ability of clients to locate or authenticate to the renamed domain controller, except when the domain controller is restarted.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This command-line method requires the Netdom Windows support tool. For information about installing Windows support tools, see Related Topics.

- The corresponding nTFRSMember or msDFSR-Member object is not renamed automatically, but the reference attributes are correctly set so SYSVOL replication is not impacted. The only potential problem with not renaming these objects is that if another domain controller is created at a later date with the same NetBIOS name of the old domain controller, then a conflict can occur as described in KB article 316826. After the rename is complete, you can optionally rename the nTFRSMember or msDFSR-Member object as part of cleanup.

- If the domain controller belongs to a group with a Group Policy enabled on its primary DNS suffix, the string specified in the Group Policy is used as the primary DNS suffix. The local setting is used only if the Group Policy is disabled or unspecified.

- By default, the primary DNS suffix portion of a computer's FQDN is the same as the name of the Active Directory domain to which the computer is joined. To allow different primary DNS suffixes, a domain administrator can create a restricted list of allowed suffixes by creating the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is managed by the domain administrator using Active Directory Service Interfaces (ADSI) or Lightweight Directory Access Protocol (LDAP). For more information about programming interfaces and directory access protocol, see Related Topics.

- Domain controller locator (Locator) DNS resource records are registered by the domain controller after the renamed domain controller has been restarted. The records that are registered are available on the domain controller in the systemroot\System32\Config\Netlogon.dns file.

- To enumerate the names with which the computer is currently configured, at a command prompt, type:

  **netdom computername** *ComputerName***/enumerate:***{AlternateNames | PrimaryName | AllNames}*

- You can also specify a parameter that will use administrator credentials required to modify the computer account in Active Directory. If this parameter is not specified, Netdom uses the credentials of the user currently logged on. For more information, see the Netdom command-line help.

- If you rename a domain controller through the **System Properties** dialog box instead of using the Netdom tool, DNS and Active Directory replication latency may delay the ability of clients to locate or authenticate to the renamed domain controller. The length of this latency depends on your network design and the replication topology of your organization.

**To rename a domain controller in a domain that contains a single domain controller**

1. Install a Windows Server 2003 member server in the domain.

2. On the new server, create an additional domain controller by installing Active Directory.

3. After Active Directory is installed, enable the global catalog on the new domain controller.

4. Transfer the operations master roles from the domain controller that you want to rename to the new domain controller. Note that you must transfer the roles, do not seize them.

5. Verify that the new domain controller is functioning correctly by doing the following:

a. Verify authentications and global catalog searches.

b. Run Dcdiag.exe against the domain controller.

c. Perform any other appropriate tests to verify that the new domain controller can provide all of the domain functions of the first domain controller.

6. Verify that the \sysvol and \netlogon drives are shared on the new domain controller by doing the following:

   a. On the new domain controller, open Command Prompt.

   b. Type:

   **Net share**

   a. In the list that is generated, verify the existence of Sysvol and Netlogon.

7. Uninstall Active Directory from the domain controller that you want to rename to be a member server.

8. Rename the member server.

9. Install Active Directory on the renamed server to create an additional domain controller.

10. Transfer the operations master roles back to the renamed domain controller.

11. Enable the global catalog on the renamed domain controller.

For information about performing the individual steps in this procedure, see the topics under "See Also."

**Note**

- You can also rename a domain controller in a domain that contains a single domain controller by using My Computer. However, doing so will result in a service interruption to clients. Always perform a system state backup before you rename a domain controller. For more information about backing up system state data, see the topics under "See Also."

**See Also**
**Concepts**
Renaming domain controllers
Rename a computer
Domain and forest functionality
How to install and upgrade the operating system
Create an additional domain controller
Enable or disable a global catalog
Transfer operations master roles
Demote a domain controller
Install Windows Support Tools
Programming interfaces
Directory access protocol
Back up System State data

# Disable signed or encrypted LDAP traffic

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To disable signed or encrypted LDAP traffic

1. Open Registry Editor.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. In Registry Editor, navigate to the following registry key:

   **Where?**

   - HKEY_LOCAL_MACHINE/SOFTWARE/Microsoft/Windows/CurrentVersion/AdminDebug

3. Click **Edit**, point to **New**, and then click **DWORD Value**.

4. In the text box that appears, type **ADsOpenObjectFlags**, and then press ENTER.

5. Double-click the **ADsOpenObjectFlags** registry key you just created, and then change the **Value Data** to one of the following values below:

| Value data (Hexadecimal) | Disables |
|---|---|
| 1 | Signing |
| 2 | Encryption |
| 3 | Encryption and Signing |

### Caution

- This procedure will disable the use of signed or encrypted LDAP traffic for some Active Directory administrative tools. It is recommended that you avoid disabling this feature. For more information, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Connecting to domain controllers running Windows 2000

Directory access protocol

# Manage Operations Master Roles

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage operations master roles**

- Identify operations master roles
- Transfer operations master roles
- Seize operations master roles

# Identify operations master roles

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Identify operations master roles**


- Identify the RID master
- Identify the PDC emulator
- Identify the infrastructure master
- Identify the domain naming master
- Identify the schema master

# Identify the RID master

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To identify the RID master

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Active Directory Users and Computers.
2. Right-click the domain node, and then click **Operations Masters**.
3. On the **RID** tab, under **Operations masters**, view the current RID master.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Every domain has only one RID master. To identify the RID master in a different domain, target the appropriate domain before clicking **Operations Masters**.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **dsquery server -hasfsmo rid**

| Value | Description |
|---|---|
| **-hasfsmo** | Finds the domain controller (server object) that currently holds the requested operations masters role. |
| **rid** | Requests the RID master of the current domain. |

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Directory service command-line tools
Working with MMC console files
Operations master roles
Responding to operations master failures
Command-line reference A-Z

# Identify the PDC emulator

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To identify the PDC emulator

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Active Directory Users and Computers.
2. Right-click the domain node, and then click **Operations Masters**.
3. On the **PDC** tab, under **Operations masters**, view the operations masters that will serve as the PDC emulator.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Every domain has only one PDC emulator master. To identify the PDC emulator in a different domain, target the appropriate domain before clicking **Operations Masters**.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **dsquery server -hasfsmo pdc**

| Value | Description |
|---|---|
| **-hasfsmo** | Finds the domain controller (server object) that currently holds the requested operations masters role. |
| **pdc** | Requests the PDC emulator master of the current domain. |

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Directory service command-line tools
Working with MMC console files
Operations master roles
Responding to operations master failures
Command-line reference A-Z

# Identify the infrastructure master

Updated: October 13, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To identify the infrastructure master**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. Right-click the domain node, and then click **Operations Masters**.
3. On the **Infrastructure** tab, under **Operations masters**, view the name of the current infrastructure master.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Every domain has only one infrastructure master. To identify the infrastructure master in a different domain, target the appropriate domain before clicking **Operations Masters**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsquery server –hasfsmo infr**

| Value | Description |
|---|---|
| **-hasfsmo** | Finds the domain controller (server object) that currently holds the requested operations masters role. |
| **infr** | Requests the infrastructure master of the forest. |

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Directory service command-line tools
Working with MMC console files
Operations master roles
Responding to operations master failures
Command-line reference A-Z

# Identify the domain naming master

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To identify the domain naming master**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. Right-click **Active Directory Domains and Trusts**, and then click **Operations Master**.
3. Under **Domain naming operations master**, view the current domain naming operations master.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- Every forest has only one domain naming master. To identify the domain naming master in a different forest, target the appropriate forest before clicking **Operations Master**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsquery server -hasfsmo name**

| Value | Description |
|---|---|
| **-hasfsmo** | Finds the domain controller (server object) that currently holds the requested operations master role. |
| **name** | Requests the domain naming master of the forest. |

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Directory service command-line tools
Working with MMC console files
Operations master roles
Responding to operations master failures
Command-line reference A-Z

# Identify the schema master

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To identify the schema master**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open the Active Directory Schema snap-in.
2. In the console tree, right-click **Active Directory Schema**, and then click **Operations Master**.
3. Under **Current Schema Master**, view the current schema operations master.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open the Active Directory Schema snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Schema**.
- Each forest has only one schema master. To identify the schema master in a different forest, target the appropriate forest before clicking **Operations Master**.
- If the Active Directory Schema snap-in is not installed, see Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

```
dsquery server -hasfsmo schema
```

| Value | Description |
|---|---|
| -hasfsmo | Finds the domain controller (server object) that currently holds the requested operations master role. |
| schema | Requests the schema master of the forest. |

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Directory service command-line tools
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Operations master roles
Responding to operations master failures
Command-line reference A-Z

# Transfer operations master roles

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Transfer operations master roles**


- Transfer the RID master role
- Transfer the PDC emulator role
- Transfer the infrastructure master role
- Transfer the domain naming master role
- Transfer the schema master role

# Transfer the RID master role

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To transfer the RID master role

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Active Directory Users and Computers.
2. In the console tree, right-click **Active Directory Users and Computers**, and then click **Connect to Domain Controller**.
3. In **Enter the name of another domain controller**, type the name of the domain controller you want to hold the RID master role.

   Or, click the domain controller in the list of available domain controllers.
4. In the console tree, right-click **Active Directory Users and Computers**, point to **All Tasks**, and then click **Operations Masters**.
5. Click the **RID** tab, and then click **Change**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To transfer a role to another domain controller, you may first need to change the focus of Active Directory Users and Computers to the target domain controller. To do this, right-click **Active Directory Users and Computers**, click **Connect to Domain Controller**, and then click the target domain controller.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **roles**
4. At the fsmo maintenance command prompt, type:

   **connection**

5. At the server connections command prompt, type:

   **connect to server***DomainController*

6. At the server connections command prompt, type:

   **quit**

7. At the fsmo maintenance command prompt, type:

   **transfer RID master**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to assign the new operations master role. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at the **ntdsutil** command prompt, type:

  **roles** ENTER

  **?** ENTER

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Manage the domain using a different domain controller
Operations master roles
Ntdsutil
Command-line reference A-Z

# Transfer the PDC emulator role

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To transfer the PDC emulator role**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click **Active Directory Users and Computers**, and then click **Connect to Domain Controller**.
3. In **Enter the name of another domain controller**, type the name of the domain controller you want to hold the PDC emulator role.

   Or, click the domain controller in the list of available domain controllers.
4. In the console tree, right-click **Active Directory Users and Computers**, point to **All Tasks**, and then click **Operations Masters**.
5. Click the **PDC** tab, and then click **Change**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To transfer a role to another domain controller, you may first need to change the focus of Active Directory Users and Computers to the target domain controller. To do this, right-click **Active Directory Users and Computers**, click **Connect to Domain Controller**, and then click the target domain controller.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **roles**
4. At the fsmo maintenance command prompt, type:

   **connection**

5. At the server connections command prompt, type:

   **connect to server***DomainController*

6. At the server connections command prompt, type:

   **quit**

7. At the fsmo maintenance command prompt, type:

   **transfer PDC**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to assign the new operations master role. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at the **ntdsutil** command prompt, type:

  **roles** ENTER

  **?** ENTER

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Manage the domain using a different domain controller
Operations master roles
Command-line reference A-Z
Ntdsutil

# Transfer the infrastructure master role

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To transfer the infrastructure master role

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Active Directory Users and Computers.
2. In the console tree, right-click **Active Directory Users and Computers**, and then click **Connect to Domain Controller**.
3. In **Enter the name of another domain controller**, type the name of the domain controller you want to hold the infrastructure master role.

   Or, click the domain controller in the list of available domain controllers.
4. In the console tree, right-click **Active Directory Users and Computers**, point to **All Tasks**, and then click **Operations Masters**.
5. On the **Infrastructure** tab, click **Change**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To transfer a role to another domain controller, you may first need to change the focus of Active Directory Users and Computers to the target domain controller. To do this, right-click **Active Directory Users and Computers**, click **Connect to Domain Controller**, and then click the target domain controller.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **roles**
4. At the fsmo maintenance command prompt, type:

   **connection**

5. At the server connections command prompt, type:

   **connect to server***DomainController*

6. At the server connections command prompt, type:

   **quit**

7. At the fsmo maintenance command prompt, type:

   **transfer infrastructure master**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to assign the new operations master role. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at the **ntdsutil** command prompt, type:

  **roles** ENTER

  **?** ENTER

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Manage the domain using a different domain controller
Operations master roles
Command-line reference A-Z
Ntdsutil

# Transfer the domain naming master role

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To transfer the domain naming master role

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click **Active Directory Domains and Trusts**, and then click **Connect to Domain Controller**.
3. In **Enter the name of another domain controller**, type the name of the domain controller you want to hold the domain naming master role.

   Or, click the domain controller in the list of available domain controllers.
4. In the console tree, right-click **Active Directory Domains and Trusts**, and then click **Operations Master**.
5. Click **Change**.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- To transfer a role to another domain controller, you may first need to change the focus of Active Directory Domains and Trusts to the target domain controller. To do this, right-click **Active Directory Domains and Trusts**, click **Connect to Domain Controller**, and then click the target domain controller.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **roles**
4. At the fsmo maintenance command prompt, type:

   **connection**

5. At the server connection command prompt, type:

   **connect to server***DomainController*

6. At the server connection command prompt, type:

   **quit**

7. At the fsmo maintenance command prompt, type:

   **transfer domain naming master**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to assign the new operations master role. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at the **ntdsutil** command prompt, type:

  **roles** ENTER

  **?** ENTER

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Manage the domain using a different domain controller
Operations master roles
Command-line reference A-Z
Ntdsutil

# Transfer the schema master role

Updated: January 20, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To transfer the schema master role

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open the **Active Directory Schema** snap-in.

   If you have not already installed the Active Directory Schema snap-in, please see Install the Active Directory Schema snap-in (http://go.microsoft.com/fwlink/?LinkID=209652).

2. In the console tree, right-click **Active Directory Schema** and then click **Change Domain Controller**.

3. Click **Specify Name** and type the name of the domain controller that you want to hold the schema master role.

4. In the console tree, right-click **Active Directory Schema**, and then click **Operations Master**.

5. Click **Change**.

### Notes

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The Active Directory Schema snap-in must be connected to the schema master to perform this procedure. The Active Directory Schema snap-in connects to the schema master by default when it is started.
- If the Active Directory Schema snap-in is not installed, see Related Topics.

#### Using a command line

1. Open Command Prompt.

2. Type:

   **ntdsutil**

3. At the ntdsutil command prompt, type:

   **roles**

4. At the fsmo maintenance command prompt, type:

   **connection**

5. At the server connections command prompt, type:

**connect to server***DomainController*

6. At the server connections command prompt, type:

   **quit**

7. At the fsmo maintenance command prompt, type:

   **transfer schema master**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to transfer the role. |

**Notes**

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open the Active Directory Schema snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Schema**.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at the **ntdsutil** command prompt, type:

  **roles** ENTER

  **?** ENTER

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Operations master roles
Command-line reference A-Z
Ntdsutil

# Seize operations master roles

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Seize operations master roles**


- Seize the RID master role
- Seize the PDC emulator role
- Seize the infrastructure master role
- Seize the domain naming master role
- Seize the schema master role

# Seize the RID master role

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To seize the RID master role**

1. Open Command Prompt.
2. Type:

   **ntdsutil**

3. At the ntdsutil command prompt, type:

   **roles**

4. At the fsmo maintenance command prompt, type:

   **connections**

5. At the server connections command prompt, type:

   **connect to server***DomainController*

6. At the server connections prompt, type:

   **quit**

7. At the fsmo maintenance command prompt, type:

   **seize RID master**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to assign the new operations master role. Use only for recovery purposes. |

**Caution**

- Do not seize the RID master role if you can transfer it instead. Seizing the RID master is a drastic step that should be considered only if the current operations master will never be available again. For more information about transferring operations master roles, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Before seizing the RID master, use the **repadmin** command-line tool to verify whether the new operations master has received any updates performed by the previous role holder, and then remove the current operations master from the network. For more information about operations master roles and the Active Directory command-line tools, see Related

Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Identify the RID master
Active Directory support tools
Operations master roles
Transferring operations master roles
Responding to operations master failures
Transfer the RID master role
Ntdsutil

# Seize the PDC emulator role

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To seize the PDC emulator role**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **roles**
4. At the fsmo maintenance command prompt, type:

   **connections**
5. At the server connections command prompt, type:

   **connect to server***DomainController*
6. At the server connections prompt, type:

   **quit**
7. At the fsmo maintenance command prompt, type:

   **seize PDC**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to assign the new operations master role. Use only for recovery purposes. |

**Caution**

- Do not seize the PDC emulator role if you can transfer it instead. Seizing the PDC emulator is a drastic step that should be considered only if the current operations master will never be available again. For more information about transferring operations master roles, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Before seizing the PDC emulator role, use the **repadmin** command-line tool to verify whether the new operations master has received any updates performed by the previous role holder, and then remove the current operations master from the network. For more information about operations master roles and the Active Directory command-line tools, see

Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Identify the PDC emulator
Active Directory support tools
Operations master roles
Transferring operations master roles
Responding to operations master failures
Transfer the PDC emulator role
Ntdsutil

# Seize the infrastructure master role

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To seize the infrastructure master role**

1. Open Command Prompt.
2. Type:

   **ntdsutil**

3. At the ntdsutil command prompt, type:

   **roles**

4. At the fsmo maintenance command prompt, type:

   **connections**

5. At the server connections command prompt, type:

   **connect to server***DomainController*

6. At the server connections prompt, type:

   **quit**

7. At the fsmo maintenance command prompt, type:

   **seize infrastructure master**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to assign the new operations master role. Use only for recovery purposes. |

**Caution**

- Do not seize the infrastructure master role if you can transfer it instead. Seizing the infrastructure master role is a drastic step that should be considered only if the current operations master will never be available again. For more information about transferring operations master roles, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Before seizing the infrastructure master role, use the **repadmin** command-line tool to verify whether the new operations master has received any updates performed by the previous role holder, and then remove the current operations master from the network. For more information about operations master roles and the Active Directory command-line tools, see

Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Identify the infrastructure master
Active Directory support tools
Operations master roles
Transferring operations master roles
Responding to operations master failures
Transfer the infrastructure master role
Ntdsutil

# Seize the domain naming master role

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To seize the domain naming master role**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **roles**
4. At the fsmo maintenance command prompt, type:

   **connections**
5. At the server connections command prompt, type:

   **connect to server** *DomainController*
6. At the server connections prompt, type:

   **quit**
7. At the fsmo maintenance command prompt, type:

   **seize domain naming master**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to assign the new operations master role. Use only for recovery purposes. |

**Caution**

- Do not seize the domain naming master role if you can transfer it instead. Seizing the domain naming master role is a drastic step that should be considered only if the current operations master will never be available again. For more information about transferring operations master roles, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Before seizing the domain naming master role, use the **repadmin** command-line tool to verify whether the new operations master has received any updates performed by the previous role holder, and then remove the current operations master from the network. For more information about operations master roles and the Active Directory

command-line tools, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Identify the domain naming master
Working with MMC console files
Active Directory support tools
Operations master roles
Transferring operations master roles
Responding to operations master failures
Transfer the domain naming master role
Ntdsutil

# Seize the schema master role

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To seize the schema master role**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **roles**
4. At the fsmo maintenance command prompt, type:

   **connections**
5. At the server connections command prompt, type:

   **connect to server** *DomainController*
6. At the server connections prompt, type:

   **quit**
7. At the fsmo maintenance command prompt, type:

   **seize schema master**

| Value | Description |
|---|---|
| *DomainController* | The domain controller to which you want to assign the new operations master role. Use only for recovery purposes. |

**Caution**

- Do not seize the schema master role if you can transfer it instead. Seizing the schema master role is a drastic step that should be considered only if the current operations master will never be available again. For more information about transferring operations master roles, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Before seizing the schema master role, use the **repadmin** command-line tool to verify whether the new operations master has received any updates performed by the previous role holder, and then remove the current operations master from the network. For more information about operations master roles and the Active Directory command-line tools, see

Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Identify the schema master
Active Directory support tools
Operations master roles
Transferring operations master roles
Responding to operations master failures
Transfer the schema master role
Ntdsutil

# Manage Trusts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage trusts**

- Verify a trust
- Remove a trust
- Create a shortcut trust
- Create an external trust
- Create a realm trust
- Select the scope of authentication for users
- Manage Forest Trusts

# Verify a trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify a trust**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that contains the trust you want to verify, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, click the trust to be verified, and then click **Properties**.
4. Click **Validate**.
5. Do one of the following, and then click **OK**:

   - Click **No, do not validate the incoming trust**.

     If you choose this option, it is recommended that you repeat this procedure for the reciprocal domain.

   - Click **Yes, validate the incoming trust**.

     If you choose this option, you must type a user account and password with administrative credentials for the reciprocal domain.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- You can verify trusts for shortcut, external, and forest trusts but not realm trusts.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **netdom trust***TrustingDomainName***/d:***TrustedDomainName***/verify**

| Value | Description |
|---|---|

| | |
|---|---|
| *TrustingDomainName* | Specifies the DNS name of the trusting domain in the trust that is being verified. |
| *TrustedDomainName* | Specifies the DNS name of the domain that is trusted in the trust that is being verified. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can verify trusts for shortcut, external, and forest trusts but not realm trusts.

- This command-line method requires the Netdom Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **netdom trust | more**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domain controllers
Runas
Using Run as
Forest trusts
Trust types
Trust direction
Install Windows Support Tools

# Remove a trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a trust**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain that contains the trust you want to remove, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, click the trust to be removed, and then click **Remove**.
4. Do one of the following, and then click **OK**:

   - Click **No, remove the trust from the local domain only**.

     If you choose this option, it is recommended that you repeat this procedure for the reciprocal domain.

   - Click **Yes, remove the trust from both the local domain and the other domain**.

     If you choose this option, you must type a user account and password with administrative credentials for the reciprocal domain.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- It is not possible to revoke the default two-way, transitive trusts between domains in a forest. Explicitly created shortcut trusts can be deleted.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **netdom trust***TrustingDomainName***/d:***TrustedDomainName***/remove/UserD:***User***/PasswordD:****

| Value | Description |
|---|---|
| *TrustingDomainName* | Specifies the DNS name of the trusting domain in the trust that is being removed. |

| | |
|---|---|
| *TrustedDomainName* | Specifies the DNS name of the domain that is trusted in the trust that is being removed. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This command-line method requires the Netdom Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **netdom trust | more**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domain controllers
Runas
Using Run as
Forest trusts
Trust types
Trust direction
Install Windows Support Tools

# Create a shortcut trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a shortcut trust**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain you want to establish a shortcut trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the DNS name (or NetBIOS name) of the domain, and then click **Next**.
5. On the **Direction of Trust** page, do one of the following:

   - To create a two-way, shortcut trust, click **Two-way**.

     Users in this domain and users in the specified domain can can use this trust path.

   - To create a one-way, incoming shortcut trust, click **One-way:incoming**.

     Users in the specified domain will not be able to use this trust path.

   - To create a one-way, outgoing shortcut trust, click **One-way:outgoing**.

     Users in this domain will not be able to use this trust path.

6. Continue to follow the wizard.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- If you have the appropriate administrative credentials for each domain, you can create both sides of a shortcut trust at the same time by clicking **Both this domain and the specified domain** on the **Sides of Trust** page. For more information, see Related Topics.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **netdom trust***TrustingDomainName***/d:***TrustedDomainName***/add**

| Value | Description |
|---|---|
| *TrustingDomainName* | Specifies the DNS name (or NetBIOS name) of the trusting domain in the trust being created. |
| *TrustedDomainName* | Specifies the DNS name (or NetBIOS name) of the domain that will be trusted in the trust being created. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This command-line method requires the Netdom Windows support tool. For information about installing Windows support tools, see Related Topics.

- Other switches can be used to assign a password or determine the direction of the trust. For example, to make the above trust a two-way, transitive trust you would use the following syntax:

  **netdom trust***TrustingDomainName***/d:***TrustedDomainName***/add/twoway**

- To view the complete syntax for this command, at a command prompt, type:

  **netdom trust | more**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
When to create a shortcut trust
Trust types
Trust direction
Runas
Using Run as
Install Windows Support Tools

# Create an external trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an external trust**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain you want to establish a trust with, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the DNS name (or NetBIOS name) of the domain, and then click **Next**.
5. On the **Trust Type** page, click **External trust**, and then click **Next**.
6. On the **Direction of Trust** page, do one of the following:

   - To create a two-way, external trust, click **Two-way**.

     Users in this domain and users in the specified domain can access resources in either domain.

   - To create a one-way, incoming external trust, click **One-way:incoming**.

     Users in the specified domain will not be able to access any resources in this domain.

   - To create a one-way, outgoing external trust, click **One-way:outgoing**.

     Users in this domain will not be able to access any resources in the specified domain.

7. Continue to follow the wizard.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- If you have the appropriate administrative credentials for each domain, you can create both sides of an external trust at the same time by clicking **Both this domain and the specified domain** on the **Sides of Trust** page. For more information, see Related Topics.
- If you want to allow users from the specified domain to obtain access to all the resources in this domain, click **Allow**

**authentication for all resources in the local domain** on the **Outgoing Trust Properties** page. This option should be used when both domains belong to the same organization.

- If you want to restrict users in the specified domain from obtaining access to any of the resources in this domain, click **Allow authentication only for selected resources in the local domain** on the **Outgoing Trust Properties** page. This option should be used when each domain belongs to a separate organization.

#### Using a command line

1. Open Command Prompt.

2. Type:

   **netdom trust***TrustingDomainName***/d:***TrustedDomainName***/add**

| Value | Description |
|---|---|
| *TrustingDomainName* | Specifies the DNS name (or NetBIOS name) of the trusting domain in the trust being created. |
| *TrustedDomainName* | Specifies the DNS name (or NetBIOS name) of the domain that will be trusted in the trust being created. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This command-line method requires the Netdom Windows support tool. For information about installing Windows support tools, see Related Topics.

- Other switches can be used to assign a password or determine the direction of the trust. For example, to make a two-way, transitive trust, you can use the following syntax:

  **netdom trust***TrustingDomainName***/d:***TrustedDomainName***/add/twoway**

- To view the complete syntax for this command, at a command prompt, type:

  **netdom trust | more**

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
When to create an external trust
Trust types
Trust direction
Runas
Using Run as
Install Windows Support Tools
Accessing resources across domains

# Create a realm trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a realm trust**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain you want to administer, and then click **Properties**.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the realm name for the target realm, and then click **Next**.
5. On the **Trust Type** page, select the **Realm trust** option, and then click **Next**.
6. On the **Transitivity of Trust** page, do one of the following:
   - To form a trust relationship with the domain and the specified realm, click **Nontransitive**, and then click **Next**.
   - To form a trust relationship with the domain and the specified realm and all trusted realms, click **Transitive**, and then click **Next**.
7. On the **Direction of Trust** page, do one of the following:
   - To create a two-way, realm trust, click **Two-way**.

     Users in this domain and users in the specified realm can access resources in either domain or realm.
   - To create a one-way, incoming realm trust, click **One-way:incoming**.

     Users in the specified realm will not be able to access any resources in this domain.
   - To create a one-way, outgoing realm trust, click **One-way:outgoing**.

     Users in this domain will not be able to access any resources in the specified realm.
8. Continue to follow the wizard.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and

then double-click **Active Directory Domains and Trusts**.

#### Using a command line

1. Open Command Prompt.

2. Type:

   **netdom trust***TrustingDomainName***/d:***TrustedDomainName***/add/realm/PasswordT:***NewRealmTrustPassword*

| Value | Description |
|---|---|
| *TrustingDomainName* | Specifies the DNS name of the trusting domain in the new realm trust. |
| *TrustedDomainName* | Specifies the DNS name of the domain that will be trusted in the new realm trust. |
| *NewRealmTrustPassword* | Specifies the trust password for the new realm trust. This password must match the one used in the Kerberos realm. |

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This command-line method requires the Netdom Windows support tool. For information about installing Windows support tools, see Related Topics.

- Other switches can be used to assign a password or determine the direction of the trust. For example, to make the above trust a two-way, transitive trust you would use the following syntax:

  **netdom trust***TrustingDomainName***/d:***TrustedDomainName***/add/realm/twoway**

- To view the complete syntax for this command, at a command prompt, type:

  **netdom trust | more**

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
When to create a realm trust
Trust types
Trust direction
Runas
Using Run as
Install Windows Support Tools

# Select the scope of authentication for users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select the scope of authentication for users**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain you want to administer, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, do one of the following:
   - To select the scope of authentication for users authenticating through an external trust, click the external trust that you want to administer, and then click **Properties**. On the **Authentication** tab, click either **Domain-wide authentication** or **Selective authentication**.
   - To select the scope of authentication for users authenticating through a forest trust, click the forest trust that you want to administer, and then click **Properties**. On the **Authentication** tab, click either **Forest-wide authentication** or **Selective authentication**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To perform this procedure for a forest trust, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- For an external trust, if you select **Selective authentication**, you need to manually enable permissions on the local domain and on the resource to which you want users in the external domain to have access.
- For a forest trust, if you select **Selective authentication**, you need to manually enable permissions on each domain and resource in the local forest to which you want users in the second forest to have access.
- You can use selective authentication only on external and forest trusts. For more information about selective authentication, see Security Considerations for Trusts.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Accessing resources across forests
Accessing resources across domains
Forest trusts
Create an external trust
Trust types

Trust direction

# Manage Forest Trusts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage forest trusts**

- Create a forest trust
- Change the routing status of a name suffix
- Enable or disable an existing name suffix from routing
- Exclude name suffixes from routing to a local forest

# Create a forest trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a forest trust**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the forest root domain, and then click **Properties**.
3. On the **Trust** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type the DNS name (or NetBIOS name) of another forest, and then click **Next**.
5. On the **Trust Type** page, click **Forest trust**, and then click **Next**.
6. On the **Direction of Trust** page, do one of the following:

   - To create a two-way, forest trust, click **Two-way**.

     Users in this forest and users in the specified forest can access resources in either forest.

   - To create a one-way, incoming forest trust, click **One-way:incoming**.

     Users in the specified forest will not be able to access any resources in this forest.

   - To create a one-way, outgoing forest trust, click **One-way:outgoing**.

     Users in this forest will not be able to access any resources in the specified forest.

7. Continue to follow the wizard.

**Important**

- To successfully create a forest trust, your environment will need to be set up properly. For more information, see the checklist for creating a forest trust in Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- If you are a member of the Incoming Forest Trust Builders group, you can create one-way, incoming forest trusts to this forest.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- On a domain controller running Windows Server 2003, the forest root domain is the first domain listed in the console tree in Active Directory Domains and Trusts. This configuration is not provided on domain controllers running Windows 2000.

- If you have the appropriate administrative credentials for each forest, you can create both sides of a forest trust at the same time by clicking **Both this domain and the specified domain** on the **Sides of Trust** page. For more information, see Related Topics.

- If you want users from the specified forest to have access to all computers in the local forest, on the **Outgoing Trust Properties** page, click **Forest-wide authentication**. This option is preferred when both forests belong to the same organization.

- If you want to selectively limit authentication to particular users and groups from the specified forest, on the **Outgoing Trust Properties** page, click **Selective authentication**. This option is preferred if the specified forest belongs to a separate organization.

- In addition to creating new trusts, you can modify existing trusts by clicking the **Trust** tab.

- The command-line tool Dsmod.exe does not support the addition of security principals in one forest to groups that are located in another forest when both forests are joined by a forest trust. You can use the Active Directory Users & Computers snap-in to add security principals across a forest trust.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Creating a forest trust
Forest trusts
When to create a forest trust
Routing name suffixes across forests
Accessing resources across forests
Trust types
Trust direction
Command-line reference A-Z

# Change the routing status of a name suffix

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the routing status of a name suffix**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain you want to administer, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, click the forest trust that you want to administer, and then click **Properties**.
4. On the **Name Suffix Routing** tab, under **Name suffixes in the *x.x* forest**, click the suffix for which you want to modify the routing status, and then click **Edit**.
5. In **Existing name suffixes in *x.x***, click the suffix that you want to modify, and then click **Enable** or **Disable**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- You cannot enable a name suffix that is in conflict. If the conflict is with a local UPN name suffix, you must remove the local user principal name (UPN) name suffix before you can enable the routing name. If the conflict is with a name that is claimed by another trust partner, you must disable the name in the other trust before it can be enabled for this trust.
- To view a log of name suffixes, Domain Name System (DNS) names, NetBIOS names and the status that is associated with this trust, click **Save As**. This log can assist you in troubleshooting authentication problems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Forest trusts
Trust types
Trust direction
Exclude name suffixes from routing to a local forest
Routing name suffixes across forests

# Enable or disable an existing name suffix from routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable an existing name suffix from routing**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain you want to administer, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, click the forest trust that you want to administer, and then click **Properties**.
4. Click the **Name Suffix Routing** tab, under **Name suffixes in the *x.x* forest**, do one of the following:
   - To enable a name suffix, click the suffix that you want to enable, and then click **Enable**. If the **Enable** button appears dimmed, then the name suffix is already enabled.
   - To disable a name suffix, click the suffix that you want to disable, and then click **Disable**. If the **Disable** button appears dimmed, then the name suffix is already disabled.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- When you disable a name suffix, all children of that Domain Name System (DNS) name will also be disabled.
- You cannot enable a name suffix that is in conflict. If the conflict is with a local user principal name (UPN) name suffix, you must remove the local UPN name suffix before you can enable the routing name. If the conflict is with a name that is claimed by another trust partner, you must disable the name in the other trust before it can be enabled for this trust.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Forest trusts
Trust types
Trust direction
Routing name suffixes across forests

# Exclude name suffixes from routing to a local forest

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To exclude name suffixes from routing to a local forest**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain node for the domain you want to administer, and then click **Properties**.
3. On the **Trusts** tab, under either **Domains trusted by this domain (outgoing trusts)** or **Domains that trust this domain (incoming trusts)**, click the forest trust that you want to administer, and then click **Properties**.
4. On the **Name Suffix Routing** tab, under **Name suffixes in the *x.x* forest**, click the unique name suffix for which you want to exclude the routing status, and then click **Edit**.
5. In **Name suffixes to exclude from routing to *x.x***, click **Add**, type a Domain Name System (DNS) name suffix that is subordinate to the unique name suffix, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- When you exclude a name suffix, all children of that DNS name will also be excluded.
- To view a log of name suffixes, DNS names, NetBIOS names and the status that is associated with this trust, click **Save As**. This log can assist you in troubleshooting authentication problems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Forest trusts
Trust types
Trust direction
Routing name suffixes across forests

# Manage Forests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage forests**

- Raise the forest functional level
- Manage Application Directory Partitions
- Manage the Schema

# Raise the forest functional level

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To raise the forest functional level**

1. Open Active Directory Domains and Trusts.
2. In the console tree, right-click the **Active Directory Domains and Trusts** node, and then click **Raise Forest Functional Level**.
3. In **Select an available forest functional level**, click Windows Server 2003, and then click **Raise**.

### Caution

- If you have or will have any domain controllers running Windows NT 4.0 and earlier or Windows 2000, then do not raise the forest functional level to Windows Server 2003. Once the forest functional level has been set to Windows Server 2003, it cannot be changed back to Windows 2000.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Domains and Trusts, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Domains and Trusts**.
- Windows 2000 native is the minimum domain functionality requirement for raising the forest functional level to Windows Server 2003.
- To raise the forest functional level, you must first upgrade (or demote) all existing Windows 2000 domain controllers within your forest.
- If you are unable to raise the forest functional level, you can click **Save As** in the **Raise Forest Functional Level** dialog box to save a log file that specifies which domain controllers in the forest still need to be upgraded from Windows NT or Windows 2000.
- The current forest functional level is displayed under **Current forest functional level** in the **Raise Forest Functional Level** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domain and forest functionality
Raising domain and forest functional levels
Raise the domain functional level

# Manage Application Directory Partitions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage application directory partitions**

- Create or delete an application directory partition
- Add or remove an application directory partition replica
- Display application directory partition information
- Set an application directory partition reference domain
- Prepare a cross-reference object
- Set a notification delay

# Create or delete an application directory partition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create or delete an application directory partition**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **domain management**
4. At the domain management command prompt, type:

   **connection**
5. At the server connections command prompt, type:

   **connect to server** *ServerName*
6. At the server connections command prompt, type:

   **quit**
7. At the domain management command prompt, do one of the following:

   - To create an application directory partition, type:

     **create nc** *ApplicationDirectoryPartition DomainController*
   - To delete an application directory partition, type:

     **delete nc** *ApplicationDirectoryPartition*

| Value | Description |
|---|---|
| *ServerName* | The DNS name of the domain controller to which you want to connect. |
| *ApplicationDirectoryPartition* | The distinguished name of the application directory partition that you want to create or delete. For example, the distinguished name of the application directory partition test.microsoft.com is dc=test,dc=Microsoft,dc=com. |
| *DomainController* | The DNS name of the domain controller on which to create or delete the application directory partition. Type NULL to create the application directory partition on the domain controller to which you are currently connected. |

**Caution**

- Removing the last replica of an application directory partition will result in the permanent loss of any data contained in the partition. For more information, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ntdsutil
Application directory partitions
Application directory partitions and domain controller demotion

# Add or remove an application directory partition replica

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add or remove an application directory partition replica**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **domain management**
4. At the domain management command prompt, type:

   **connection**
5. At the server connections command prompt, type:

   **connect to server** *DomainController*
6. At the server connections command prompt, type:

   **quit**
7. At the domain management command prompt, do one of the following:

   - To add an application directory partition replica, type:

     **add nc replica** *ApplicationDirectoryPartition DomainController*
   - To remove an application directory partition replica, type:

     **remove nc replica** *ApplicationDirectoryPartition DomainController*

| Value | Description |
|---|---|
| *DomainController* | The DNS name of the domain controller on which to add or remove the replica of the application directory partition. |
| *ApplicationDirectoryPartition* | The distinguished name for the application directory partition of which you want to add or remove a replica. For example, the distinguished name of the application directory partition test.microsoft.com is dc=test,dc=microsoft,dc=com. |

**Caution**

- Removing the last replica of an application directory partition will result in the permanent loss of any data contained in the partition.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in

Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The value NULL can be used for the *DomainController* parameter of the **add nc replica** and **remove nc replica** commands if you are adding or removing the application directory partition replica on the domain controller to which you are currently connected.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ntdsutil
Application directory partitions
Application directory partitions and domain controller demotion

# Display application directory partition information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To display application directory partition information**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command prompt, type:

   **domain management**
4. At the domain management command prompt, type:

   **connection**
5. At the connection command prompt, type:

   **connect to server** *DomainController*
6. At the connection command prompt, type:

   **quit**
7. At the domain management command prompt, do one or more of the following:

   - To show the distinguished names of known directory partitions, type:

     **list**
   - To show the reference domain and replication delays for an application directory partition, type:

     **list nc information** *DistinguishedName*
   - To show the list of domain controllers in the replica set for an application directory partition, type:

     **list nc replicas** *DistinguishedName*

| Value | Description |
|---|---|
| *DomainController* | The DNS name of the domain controller for which you want to display application directory partition information. |
| *DistinguishedName* | The distinguished name of the application directory partition you want information about. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command**

**prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ntdsutil
Application directory partitions

# Set an application directory partition reference domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set an application directory partition reference domain**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command, type:

   **domain management**
4. At the domain management command, type:

   **connection**
5. At the connection command, type:

   **connect to server** *DomainController*
6. At the connection command, type:

   **quit**
7. At the domain management command, type:

   **set nc reference domain** *ApplicationDirectoryPartition ReferenceDomain*

| Value | Description |
|---|---|
| *DomainController* | The DNS name of the domain controller on which the application directory partition resides. |
| *ApplicationDirectoryPartition* | The distinguished name of the application directory partition for which you want to set the reference domain. For example, the distinguished name of the application directory partition test.microsoft.com is dc=test, dc=microsoft, dc=com. |
| *ReferenceDomain* | The distinguished name of the domain that you want to be the reference domain for this application directory partition. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Ntdsutil

Application directory partitions

Managing application directory partitions

# Prepare a cross-reference object

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prepare a cross-reference object**

1. Open Command Prompt.
2. Type:

   **ntdsutil**
3. At the ntdsutil command, type:

   **domain management**
4. At the domain management command, type:

   **connection**
5. At the connection command, type:

   **connect to server** *DomainController*
6. At the connection command, type:

   **quit**
7. At the domain management command, type:

   **Precreate** *ObjectNameDomainController*

| Value | Description |
|---|---|
| *DomainController* | The DNS name of the domain controller on which the object will reside. |
| *ObjectName* | The distinguished name of the object that you want to create. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure can be used for either preparing a domain for use with the Active Directory Installation Wizard or preparing an application directory partition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Ntdsutil

Managing application directory partitions

# Set a notification delay

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set a notification delay**

1. Open Command Prompt.
2. Type:

   **ntdsutil**

3. At the ntdsutil command, type:

   **domain management**

4. At the domain management command, type:

   **connection**

5. At the connection command, type:

   **connect to server** *DomainController*

6. At the connection command, type:

   **quit**

7. At the domain management command, type:

   **set nc replicate notification delay** *DirectoryPartition DelayInSeconds AdditionalDelayInSeconds*

| Value | Description |
|---|---|
| *DirectoryPartition* | The distinguished name of the directory partition for which you want to set notification delays. |
| *DelayInSeconds* | The number of seconds to delay before sending the change notification to the first replication partner. |
| *AdditionalDelayInSeconds* | The number of seconds to delay before sending subsequent change notifications to the remaining replication partners. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure can be applied to any directory partition including application directory partitions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ntdsutil
Managing application directory partitions

# Manage the Schema

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage the schema**

- Install the Active Directory Schema snap-in
- Add a member to the Schema Admins group
- Remove a member from the Schema Admins group
- Modify schema permissions
- Add an attribute to the global catalog
- Index an attribute in Active Directory
- Index an attribute for a containerized search
- View a schema class or attribute definition
- Modify an existing schema class or attribute definition
- Add a new schema class or attribute definition
- Deactivate a class or attribute
- Reactivate a class or attribute
- Reload the schema
- Enable schema changes on a schema master running Windows 2000

# Install the Active Directory Schema snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install the Active Directory Schema snap-in**

1. Open Command Prompt.
2. Type:

   **regsvr32 schmmgmt.dll**

   This command will register Schmmgmt.dll on your computer. For more information about using **regsvr32**, see Related Topics.
3. Click **Start**, click **Run**, type **mmc /a**, and then click **OK**.
4. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
5. Under **Available Standalone Snap-ins**, double-click **Active Directory Schema**, click **Close**, and then click **OK**.
6. To save this console, on the **File** menu, click **Save**.
7. In **Save in**, point to the systemroot\system32 directory.
8. In **File name**, type **schmmgmt.msc**, and then click **Save**.
9. To create a shortcut on your **Start** menu:

   - Right-click **Start**, click **Open All Users**, double-click the **programs** folder, and then double-click the **Administrative Tools** folder.
   - On the **File** menu, point to **New**, and then click **Shortcut**.
   - In the Create Shortcut Wizard, in **Type the location of the item**, type **schmmgmt.msc**, and then click **Next**.
   - On the **Select a Title for the program** page, in **Type a name for this shortcut**, type **Active Directory Schema**, and then click **Finish**.

**Caution**

- Modifying the schema is an advanced operation best performed by experienced programmers and system administrators. For detailed information about modifying the schema, see the Active Directory programmer's Guide at the Microsoft Web site.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- You can also run the Active Directory Schema snap-in from a computer running Windows XP Professional. Simply install the Windows Server 2003 Administration Tools Pack on the computer, and then complete step 9 above.

- The Windows Server 2003 Administration Tools Pack cannot be installed on computers running Windows 2000 Professional or Windows 2000 Server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack
Regsvr32
View a schema class or attribute definition
Identify the schema master

# Add a member to the Schema Admins group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a member to the Schema Admins group**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users** in the forest root domain.

   **Where?**

   - Active Directory Users and Computers/*forest root domain*/Users

     Or, click the folder that contains the user account.
3. In the details pane, right-click the user account that you want to add, and then click **Properties**.
4. Click the **Member Of** tab, and then click **Add**.
5. In **Enter the object names to select**, type **Schema Admins**, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To add users and computers to a group, you can also select the users and computers to be added, click

  on the toolbar, and then click the group into which they will be added.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmod group** *GroupDN* **-addmbr** *MemberDN*

| Value | Description |
|---|---|

| | |
|---|---|
| *GroupDN* | Specifies the distinguished name of the Schema Admins group in your root domain.<br><br>For example, if the name of your root domain is example.microsoft.com, you would type: **"cn=schema admins,cn=users,dc=example,dc=microsoft,dc=com"** |
| *MemberDN* | Specifies the distinguished name of the user you want to add to the Schema Admins group.<br><br>For example, if you want to add a user named "Mary Baker" to the Schema Admins group, you would type: **"cn=mary baker,cn=users,dc=example,dc=microsoft,dc=com"** |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt type:

  **dsmod group /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find a group
Group types
Dsmod

# Remove a member from the Schema Admins group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a member from the Schema Admins group**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users** in the forest root domain.

   **Where?**

   - Active Directory Users and Computers/*forest root domain*/Users

3. In the details pane, right-click the **Schema Admins** group, and then click **Properties**.
4. On the **Members** tab, click the member you want to delete, and then click **Remove**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsmod group***GroupDN***-rmmbr***MemberDN*

| Value | Description |
|---|---|
| *GroupDN* | Specifies the distinguished name of the Schema Admins group in your root domain.<br><br>For example, if the name of your root domain is example.microsoft.com, you would type: **"cn=schema admins,cn=users,dc=example,dc=microsoft,dc=com"** |
| *MemberDN* | Specifies the distinguished name of the user you want to remove from the Schema Admins group.<br><br>For example, if you want to remove a user named "Mary Baker" from the Schema Admins group, you would type: **"cn=mary baker,cn=users,dc=example,dc=microsoft,dc=com"** |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt type:

  **dsmod group /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Find a group
Dsmod

# Modify schema permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify schema permissions**

1. Open the Active Directory Schema snap-in.
2. In the console tree, right-click **Active Directory Schema**, and then click **Permissions**.
3. On the **Security** tab, under **Group or user names**, click the group whose permissions you want to change.
4. Under **Permissions for Administrators**, select **Allow** or **Deny** for the permissions you want to change.

## Notes

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The Active Directory Schema snap-in must be connected to the schema master to perform this procedure. The Active Directory Schema snap-in connects to the schema master by default when it is started.
- If the Active Directory Schema snap-in is not installed, see Other Resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in

# Add an attribute to the global catalog

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an attribute to the global catalog**

1. Open the Active Directory Schema snap-in.
2. In the console tree, click **Attributes**.

   **Where?**

   - Active Directory Schema/Attributes

3. In the details pane, right-click the attribute that you want to add to the global catalog, and then click **Properties**.
4. Select the **Replicate this attribute to the Global Catalog** check box.

### Caution

- If the forest functional level is not set to Windows Server 2003, adding a new attribute to the global catalog will cause a full synchronization of the global catalog. In other words, when you add a new attribute to the global catalog attribute set, all attributes that were previously part of the global catalog in addition to the new attribute are immediately synchronized throughout the forest. For more information, see Related Topics.

### Important

- As a best practice, it is strongly recommended that you use global groups or universal groups instead of domain local groups when specifying permissions on domain objects replicated to the global catalog. For more information, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The Active Directory Schema snap-in must be connected to the schema master to perform this procedure. The Active Directory Schema snap-in connects to the schema master by default when it is started.
- If the Active Directory Schema snap-in is not installed, see Other Resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Global catalog replication

Domain and forest functionality

# Index an attribute in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To index an attribute in Active Directory**

1. Open the Active Directory Schema snap-in.
2. In the console tree, click **Attributes**.

   **Where?**

   - Active Directory Schema/Attributes

3. In the details pane, right-click the attribute that you want to index, and then click **Properties**.
4. Click **Index this attribute in the Active Directory**.

## Notes

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The Active Directory Schema snap-in must be connected to the schema master to perform this procedure. The Active Directory Schema snap-in connects to the schema master by default when it is started.
- If the Active Directory Schema snap-in is not installed, see Other Resources.
- Indexing an attribute improves the performance of queries run against the attribute.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Index an attribute for a containerized search

# Index an attribute for a containerized search

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To index attributes for a containerized search**

1. Open the Active Directory Schema snap-in.
2. In the console tree, click **Attributes**.

   **Where?**

   - Active Directory Schema/Attributes

3. In the details pane, right-click the attribute that you want to index, and then click **Properties**.
4. Select the **Index this attribute for containerized searches in the Active Directory** check box.

## Notes

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The Active Directory Schema snap-in must be connected to the schema master to perform this procedure. The Active Directory Schema snap-in connects to the schema master by default when it is started.
- If the Active Directory Schema snap-in is not installed, see Other Resources.
- To complete this procedure, you must first index the attribute. For more information, see Other Resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the Active Directory Schema snap-in
Index an attribute in Active Directory

# View a schema class or attribute definition

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view a schema class or attribute definition**

1. Open the Active Directory Schema snap-in.
2. In the console tree, click **Active Directory Schema**.
3. Do one of the following:
   - To view a class definition, in the console tree, click **Classes**. In the details pane, right-click the class for which you want to view the definition, and then click **Properties**.
   - To view an attribute definition, in the console tree, click **Attributes**. In the details pane, right-click the attribute for which you want to view the definition, and then click **Properties**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- If the Active Directory Schema snap-in is not installed, see Other Resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Global catalog replication
Domain and forest functionality

# Modify an existing schema class or attribute definition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify an existing schema class or attribute definition**

1. Open the Active Directory Schema snap-in.
2. In the console tree, click **Active Directory Schema**.
3. Do one of the following:
   - To modify a class definition, in the console tree, click **Classes**. In the details pane, right-click the class you want to modify, and then click **Properties**.
   - To modify an attribute definition, in the control tree, click **Attributes**. In the details pane, right-click the attribute you want to modify, and then click **Properties**.

## Notes

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The Active Directory Schema snap-in must be connected to the schema master to perform this procedure. The Active Directory Schema snap-in connects to the schema master by default when it is started.
- If the Active Directory Schema snap-in is not installed, see Other Resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Global catalog replication
Domain and forest functionality

# Add a new schema class or attribute definition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a new schema class or attribute definition**

1. Open the Active Directory Schema snap-in.
2. In the console tree, click **Active Directory Schema**.
3. Do one of the following:
   - To add a class definition, in the console tree, right-click **Classes**, click **Create Class**, and then follow the instructions.
   - To add an attribute definition, in the console tree, right-click **Attributes**, click **Create Attribute**, and then follow the instructions.

## Notes

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The Active Directory Schema snap-in must be connected to the schema master to perform this procedure. The Active Directory Schema snap-in connects to the schema master by default when it is started.
- If the Active Directory Schema snap-in is not installed, see Other Resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Global catalog replication
Domain and forest functionality
Schema classes and attributes

# Deactivate a class or attribute

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To deactivate a class or attribute**

1. Open the Active Directory Schema snap-in.
2. In the console tree, click **Active Directory Schema**.
3. Do one of the following:
   - To deactivate a class, in the console tree, click **Classes**. In the details pane, right-click the class you want to deactivate, and then click **Properties**.
   - To deactivate an attribute, in the console tree, click **Attributes**. In the details pane, right-click the attribute you want to deactivate, and then click **Properties**.
4. On the **General** tab, clear the **Class is active** or **Attribute is active** check box, as appropriate.

## Notes

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The Active Directory Schema snap-in must be connected to the schema master to perform this procedure. The Active Directory Schema snap-in connects to the schema master by default when it is started.
- If the Active Directory Schema snap-in is not installed, see Other Resources.
- The status of an attribute or class is displayed in the **Status** column of the details pane. Once a class or attribute has been deactivated, it is considered *defunct*.
- Default schema attributes or classes in the base schema cannot be deactivated. Only attributes or classes added as extensions to the base schema can be deactivated.
- Classes and attributes added to the base schema can be deactivated without raising the forest functional level. However, they can be redefined only in forests with the forest functional level set to Windows Server 2003.
- To view defunct classes or attributes, click **Classes** or **Attributes**, and then on the **View** menu, click **Defunct Objects**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Reactivate a class or attribute

Deactivating a class or attribute
Domain and forest functionality

# Reactivate a class or attribute

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reactivate a class or attribute**

1. Open the Active Directory Schema snap-in.
2. In the console tree, click **Active Directory Schema**.
3. Do one of the following:
   - To reactivate a defunct class, in the console tree, click **Classes**. In the details pane, right-click the class you want to reactivate, and then click **Properties**.
   - To reactivate a defunct attribute, in the console tree, click **Attributes**. In the details pane, right-click the attribute you want to reactivate, and then click **Properties**.
4. On the **General** tab, select the **Class is active** or **Attribute is active** check box, as appropriate.

## Notes

- To perform this procedure, you must be a member of the Schema Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- The Active Directory Schema snap-in must be connected to the schema master to perform this procedure. The Active Directory Schema snap-in connects to the schema master by default when it is started.
- If the Active Directory Schema snap-in is not installed, see Other Resources.
- The status of an attribute or class is displayed in the **Status** column of the details pane. Once a class or attribute has been deactivated, it is considered *defunct*.
- To view defunct classes or attributes, click **Classes** or **Attributes**, and then on the **View** menu, click **Defunct Objects**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Deactivating a class or attribute
Domain and forest functionality

# Reload the schema

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reload the schema**

1. Open the Active Directory Schema snap-in.
2. In the console tree, right-click **Active Directory Schema**, and then click **Reload the Schema**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain), the Schema Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- If the Active Directory Schema snap-in is not installed, see Other Resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Install the Active Directory Schema snap-in
Schema
Schema classes and attributes

# Enable schema changes on a schema master running Windows 2000

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable schema changes on a schema master running Windows 2000**

1. Open Registry Editor.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. In Registry Editor, navigate to the **Parameters** registry key:

   **Where?**

   - HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Services/NTDS/Parameters

3. On the **Edit** menu, click **New**, and then click **DWORD Value**.

4. Type **Schema Update Allowed**, and then press ENTER.

5. Double-click the **Schema Update Allowed** registry key that you just created, and then, in **Value data**, type **1** to enable schema changes.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

- To disable schema changes, in the **Schema Update Allowed** registry key, under **Value data**, type **0**.

- This procedure should only be performed on a domain controller that is running Windows 2000 and that also holds the schema master role.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domain and forest functionality
Raising domain and forest functional levels
Operations master roles

# Manage Sites

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage sites**

- Configure Site Settings
- Configure Active Directory Server Settings
- Configure Replication Between Sites

# Configure Site Settings

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure site settings**

- Create a site
- Rename a site
- Delete a site
- Create a subnet
- Associate a subnet with a site
- Delete a subnet
- Select another licensing computer
- Cache universal group memberships
- Delegate control of a site
- Connect to a forest
- Connect to a domain controller
- Reveal the services node

# Create a site

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a site**

1. Open Active Directory Sites and Services.
2. Right-click the **Sites** folder, and then click **New Site**.
3. In **Name**, type the name of the new site.
4. Click a site link object, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- It is recommended that you use legal Domain Name System (DNS) names when you create new site names; otherwise, your site will only be accessible where a Microsoft DNS server is used. For more information about DNS names, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Associate a subnet with a site
Select another licensing computer
Namespace planning for DNS
Sites overview
When to establish a single or separate sites

# Rename a site

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a site**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click the site that you want to rename, and then click **Rename**.

   **Where?**

   - Active Directory Sites and Services/Sites/*site that you want to rename*

3. Type the new site name.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- It is recommended that you use legal Domain Name System (DNS) names when you create new site names; otherwise, your site will only be accessible where a Microsoft DNS server is used. For more information about DNS names, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Associate a subnet with a site
Select another licensing computer
Namespace planning for DNS
Sites overview
When to establish a single or separate sites

# Delete a site

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a site**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click the site that you want to delete, and then click **Delete**.

   **Where?**

   - Active Directory Sites and Services/Sites/*site that you want to delete*

### Important

- Before a site can be deleted, any domain controller objects within the site must be deleted or moved to another site.

### Notes

- To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sites overview
Delete a subnet

# Create a subnet

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a subnet**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click **Subnets**, and then click **New Subnet**.

   **Where?**

   - Active Directory Sites and Services/Sites/Subnets

3. In **Address**, type the subnet address.
4. In **Mask**, type the subnet mask that describes the range of addresses included in this subnet.
5. Under **Select a site object for this subnet**, click the site to associate with this subnet, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sites overview
IP addressing
Associate a subnet with a site
When to establish a single or separate sites

# Associate a subnet with a site

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To associate a subnet with a site**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click the subnet with which you want to associate the site, and then click **Properties**.

   **Where?**

   - Active Directory Sites and Services/Sites/Subnets/*subnet*

3. In **Site**, click the site with which to associate this subnet.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sites overview
IP addressing

# Delete a subnet

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a subnet**

1. Open Active Directory Sites and Services.

2. In the console tree, right-click the subnet that you want to delete, and then click **Delete**.

   **Where?**

   - Active Directory Sites and Services/Sites/Subnets/*subnet that you want to delete*

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sites overview
Delete a site

# Select another licensing computer

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select another licensing computer**

1. Open Active Directory Sites and Services.
2. In the console tree, click the site you want to configure.

   **Where?**

   - Active Directory Sites and Services/Sites/*site*

3. In the details pane, right-click **Licensing Site Settings**, and then click **Properties**.
4. On the **Licensing Settings** tab, click **Change**.
5. In **Enter the object name to select**, type the computer name of the alternative licensing computer, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- The licensing server collects licensing information from within the site for use by the licensing administrative tool. For more information, see Related Topics.
- Active Directory will automatically configure a licensing computer for each site. For optimal performance, the licensing computer for a site should reside within that same site. A licensing computer does not have to be a domain controller.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the site license server for a site
Licensing replication
Find the site license server for a site
Sites overview

# Cache universal group memberships

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To cache universal group memberships**

1. Open Active Directory Sites and Services.
2. In the console tree, click the site in which you want to enable universal group membership caching.

   **Where?**

   - Active Directory Sites and Services/Sites/*the site in which you want to enable universal group membership caching*

3. In the details pane, right-click **NTDS Site Settings**, and then click **Properties**.
4. Select the **Enable Universal Group Membership Caching** check box.
5. In **Refresh cache from**, click a site from which this site will refresh its cache, or accept **<Default>** to refresh the cache from the nearest site that has a global catalog.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
The role of the global catalog
Global catalogs and sites
Global catalog replication

# Delegate control of a site

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate control of a site**

1. Open Active Directory Sites and Services.
2. Right-click the container whose control you want to delegate, and then click **Delegate Control** to start the Delegation of Control Wizard.
3. Follow the instructions in the Delegation of Control Wizard.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- In Active Directory Sites and Services, you can delegate control for the subnets, intersite transports, sites, and server containers. Delegating control of an object allows you to specify who has permissions to access or modify that object or its child objects.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delegating administration
Sites overview
**Other Resources**
Delegate control

# Connect to a forest

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a forest**

1. Open Active Directory Sites and Services.
2. Right-click **Active Directory Sites and Services**, and then click **Connect to Forest**.
3. In **Root domain**, type the root domain of the forest, and then click **OK**.

   This is the root domain of the first domain tree created in the forest.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- Active Directory Sites and Services views a single forest. Change forests to view and configure Active Directory Sites and Services settings for a forest of your choosing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domain and forest functionality

# Connect to a domain controller

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a domain controller**

1. Open Active Directory Sites and Services.
2. Right-click **Active Directory Sites and Services**, and then click **Connect to Domain Controller**.
3. In **Domain**, type a domain name, or click **Browse** to select from the list of available domains.
4. In **Enter the name of another domain controller**, click a domain controller from the list of available domain controllers, or accept the default value **Any Writable Domain Controller** to have a domain controller automatically selected.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- Active Directory Sites and Services views and configures a single copy of the directory data. Change the domain controller to select an alternate domain controller whose directory data will read and modify.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domain controllers
Domains

# Reveal the services node

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reveal the services node**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click **Active Directory Sites and Services**, point to **View**, and then click **Show Services Node**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- Show Services Node is a persistent setting.
- Managing objects located in the Services container from within Active Directory Sites and Services provides limited functionality. It is recommended that you use other appropriate administrative tools, including ADSI Edit, for managing objects located in the Services container. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Active Directory support tools
Publishing resources
**Other Resources**
Services

# Configure Active Directory Server Settings

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure server settings**

- Select a query policy
- Enable or disable a global catalog
- Designate a preferred bridgehead server
- Move a domain controller between sites
- Check the replication topology
- Delete extinct server metadata

# Select a query policy

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select a query policy**

1. Open Active Directory Sites and Services.

2. In the console tree, click the domain controller whose query policy you want to change.

   **Where?**

   - Active Directory Sites and Services/Sites/*site that contains the domain controller whose query policy you want to change*/Servers/*domain controller whose query policy you want to change*

3. In the details pane, right-click **NTDS Settings**, and then click **Properties**.

4. On the **General** tab, in **Query Policy**, click a query policy.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the domain of the selected domain controller) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- A query policy prevents specific Lightweight Directory Access Protocol (LDAP) operations from adversely impacting the performance of the domain controller and also makes the domain controller more resilient to denial-of-service attacks.

- To create or modify query policies, use the Ntdsutil command-line tool.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Ntdsutil
Ntdsutil

# Enable or disable a global catalog

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable a global catalog**

1. Open Active Directory Sites and Services.
2. In the console tree, click the domain controller where you want to enable or disable the global catalog.

   **Where?**

   - Active Directory Sites and Services/Sites/*site that contains the domain controller that you want to disable or enable*/Servers/*domain controller*

3. In the details pane, right-click **NTDS Settings**, and then click **Properties**.
4. Select the **Global Catalog** check box to enable the global catalog, or clear the check box to disable the global catalog.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the domain of the selected domain controller) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- The Builtin Administrator account in the domain can always log on to the domain, even when a global catalog is not available.
- Enabling a global catalog can cause additional replication traffic.
- A domain controller on which you enable the global catalog does not advertise itself in Domain Name System (DNS) as a global catalog until all partial domain directory partitions have been received.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Windows Server 2003 Administration Tools Pack Overview
Global catalogs and sites
The role of the global catalog

# Designate a preferred bridgehead server

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To designate a preferred bridgehead server**

1. Open Active Directory Sites and Services.

2. In the console tree, right-click the domain controller that you want to make a preferred bridgehead server, and then click **Properties**.

   **Where?**

   - Active Directory Sites and Services/Sites/*site that contains the domain controller that you want to make a preferred bridgehead server*/Servers/*the domain controller that you want to make a preferred bridgehead server*

3. On the **General** tab, click the intersite transport or transports for which this computer will be a preferred bridgehead server, and then click **Add**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the domain of the selected domain controller) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- Assigning a bridgehead server manually can cause replication problems because it prevents the Knowledge Consistency Checker (KCC) from making an automatic failover assignment if a manually selected bridgehead server fails. For more information, see Related Topics.

- You can establish preferred bridgehead servers for one or more protocols (such as IP or SMTP).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing replication
Managing directory replication

# Move a domain controller between sites

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move a domain controller between sites**

1. Open Active Directory Sites and Services.

2. In the console tree, right-click the domain controller that you want to move to a different site, and then click **Move**.

   **Where?**

   - Active Directory Sites and Services/Sites/*site that contains the domain controller you want to move*/Servers/*domain controller to be moved*

3. In the **Move Server** dialog box, click the site to which you want to move the domain controller.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the domain of the selected domain controller) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- This procedure can be used to move member servers, as well as domain controllers, between sites.

- Before you move a domain controller to a different site, verify that the IP address of the domain controller is included in the range of IP addresses defined in a subnet of that site.

- You can also drag server objects to move them.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sites overview
Active Directory server roles
**Other Resources**
Locating a domain controller

# Check the replication topology

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check the replication topology**

1. Open Active Directory Sites and Services.

2. In the console tree, click the domain controller you want to use to check replication topology.

   **Where?**

   - Active Directory Sites and Services/Sites/*site that contains the domain controller whose replication topology you want to check*/Servers/*server you want to use to check replication topology*

3. In the details pane, right-click **NTDS Settings**, point to **All Tasks**, and then click **Check Replication Topology**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- Active Directory runs the intersite topology generator on one domain controller in a site to consider the cost of intersite connections and to check if any previously available domain controllers are no longer available or if new domain controllers have been added. Then, the Knowledge Consistency Checker (KCC) uses this information to add or remove connection objects as needed for efficient replication. This process does not affect manually created connection objects. Only one domain controller per site assumes the role of the intersite topology generator.

- To identify the intersite topology generator in each site, you can use the Dcdiag Windows Support tool available on your operating system installation CD.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Replication between sites
Replication overview
How replication works
Install Windows Support Tools

# Delete extinct server metadata

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete extinct server metadata**

1. Open Command Prompt.

2. Type:

   **ntdsutil**

3. At the **ntdsutil** command, type:

   **metadata cleanup**

4. Perform metadata cleanup as follows:

   - If you are performing metadata cleanup by using the version of Ntdsutil.exe that is included with Windows Server 2003 Service Pack 1 (SP1), at the **metadata cleanup** command, type:

     **remove selected server***ServerName*

     Or

     **remove selected server***ServerName1***on** *ServerName2*

     | Value | Description |
     |---|---|
     | *ServerName*, *ServerName1* | The distinguished name of the domain controller whose metadata you want to remove, in the form cn=*ServerName*,cn=Servers,cn=*SiteName*, cn=Sites,cn=Configuration,dc=*ForestRootDomain*. |
     | *ServerName2* | The DNS name of the domain controller to which you want to connect and from which you want to remove server metadata. |

   - If you are performing metadata cleanup by using the version of Ntdsutil.exe that is included with Windows Server 2003 with no service pack, perform metadata cleanup as follows:

   a. At the metadata cleanup command, type:

      **connection**

   b. At the **server connections** command, type:

      **connect to server** *Server*

   c. At the connection command, type:

      **quit**

   d. At the metadata cleanup command, type:

      **select operation target**

   e. At the select operation target command, type:

**list sites**

f. A numbered list of sites will be displayed, type:

**select site** *SiteNumber*

g. At the select operation target command, type:

**list domains in site**

h. A numbered list of domains in the selected site will be displayed, type:

**select domain** *DomainNumber*

i. At the select operation target command, type:

**list servers in site**

j. A numbered list of servers in a domain and site will be displayed, type:

**select server** *ServerNumber*

k. At the select operation target command, type:

**quit**

l. At the metadata cleanup command, type:

**remove selected server**

| Value | Description |
|---|---|
| *Server* | The DNS name of a domain controller that you want to connect to. |
| *SiteNumber* | The number associated to the site of the server that you want to clean up, which displays in the list. |
| *DomainNumber* | The number associated to the domain of the server that you want to clean up, which displays in the list. |
| *ServerNumber* | The number associated to the server that you want to clean up, which displays in the list. |

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group (in the domain of the selected domain controller) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can use this procedure to delete from the directory the metadata of a domain controller that could not be cleanly decommissioned.

- Do not delete the metadata of existing domain controllers.

- To clean up metadata if you are using the version of Ntdsutil.exe that is included with Windows Server 2003 with no service pack installed, you must first select the operation target. When selecting the operation target for a server to clean up, you must first select the site, and then the domain, of that server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ntdsutil
Using Ntdsutil
Administration tools for the Active Directory schema
Managing Active Directory from the command line

# Configure Replication Between Sites

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure replication between sites**


- Create a site link
- Delete a site link
- Create a site link bridge
- Delete a site link bridge
- Configure site link replication availability
- Configure site link cost
- Configure site link replication frequency
- Ignore replication schedules
- Enable or disable site link bridges
- Add a site to a site link
- Manually add connections
- Force replication over a connection

# Create a site link

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a site link**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click the intersite transport protocol you want the site link to use, and then click **New Site Link**.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*inter-site transport protocol you want the site link to use*

3. In **Name**, type the name to be given to the link.
4. Click two or more sites to connect, and then click **Add**.
5. Configure the site link's cost, schedule, and replication frequency. See Related Topics for information.

### Important

- If you create a site link that uses SMTP, you must have an enterprise certification authority (Enterprise CA) available and SMTP must be installed on all domain controllers that will use the site link.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- Creating a site link between two or more sites is a way to influence replication topology. By creating a site link, you provide Active Directory with information about what connections are available, which ones are preferred, and how much bandwidth is available. Active Directory uses this information to choose times and connections for replication that will afford the best performance.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure site link cost
Configure site link replication availability
Configure site link replication frequency
Replication overview
Add a site to a site link
Installing and configuring a certification authority

# Delete a site link

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a site link**

1. Open Active Directory Sites and Services.

2. In the console tree, click the inter-site transport folder that contains the site link you want to delete.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*inter-site transport folder that contains the site link you want to delete*

3. In the details pane, right-click the site link you want to delete, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- Before deleting a site link, you should ensure that its member sites will still belong to at least one existing site link.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Replication overview
Managing replication
Add a site to a site link

# Create a site link bridge

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a site link bridge**

1. Open Active Directory Sites and Services.

2. In the console tree, right-click the inter-site transport folder for which you want to create a new site link bridge, and then click **New Site Link Bridge**.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*intersite transport for which you want to create a new site link bridge*

3. In **Name**, type a name for the site link bridge.

4. Click two or more site links to be bridged, and then click **Add**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- If **Bridge all site links** is enabled (default setting), this procedure is redundant and will have no effect.

- A site link bridge creates a chain of site links across which domain controllers from different sites in the site links can directly communicate.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable site link bridges
Replication overview
Replication between sites
Managing replication

# Delete a site link bridge

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a site link bridge**

1. Open Active Directory Sites and Services.

2. In the console tree, click the inter-site transport folder that contains the site link bridge you want to delete.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*inter-site transport folder that contains the site link bridge you want to delete*

3. In the details pane, right-click the site link bridge that you want to delete, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Replication overview
Managing replication

# Configure site link replication availability

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure site link replication availability**

1. Open Active Directory Sites and Services.
2. In the console tree, click the inter-site transport folder that contains the site link whose schedule you want to adjust.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*inter-site transport folder that contains the site link whose schedule you want to adjust*

3. In the details pane, right-click the site link whose schedule you want to adjust, and then click **Properties**.
4. Click **Change Schedule**.
5. Select the block of time you want to schedule, and then click **Replication Not Available** or **Replication Available**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- Replication availability schedules can be temporarily ignored by selecting **Ignore schedules** on the appropriate inter-site transport folder. For more information, see Related Topics.
- By default, Simple Mail Transfer Protocol (SMTP) replication does not utilize site link availability information.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure site link replication frequency
How replication works
Managing replication
Ignore replication schedules

# Configure site link cost

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure site link cost**

1. Open Active Directory Sites and Services.
2. In the console tree, click the inter-site transport folder that contains the type of link you are going to configure.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*inter-site transport folder that contains the type of link you are going to configure*

3. In the details pane, right-click the site link whose cost you want to set, and then click **Properties**.
4. In **Cost**, enter a value for the cost of replication.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- The Active Directory Knowledge Consistency Checker (KCC) considers a site link with a higher cost to be less desirable than a site link with a lower cost. When multiple paths for replication are available, the KCC prefers the site link with the lower cost.
- You cannot apply costs directly to site link bridges.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing replication
Replication between sites
Replication overview

# Configure site link replication frequency

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure site link replication frequency**

1. Open Active Directory Sites and Services.
2. In the console tree, click the inter-site transport folder that contains the site link you want to configure.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*inter-site transport folder that contains the site link you want to configure*

3. In the details pane, right-click the site link whose replication frequency you want to set, and then click **Properties**.
4. In **Replicate every**, type the number of minutes between replications.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- The default site link replication frequency is 180 minutes.
- The **Replicate every** value will be processed as the nearest multiple of 15 ranging from a minimum of 15 to a maximum of 10,080 minutes, which corresponds to one week.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure site link replication availability
Managing replication
Replication overview

# Ignore replication schedules

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To ignore replication schedules**

1. Open Active Directory Sites and Services.

2. In the console tree, right-click the intersite transport protocol that contains the site link schedules you want ignored, and then click **Properties**.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*the intersite transport protocol for which you want replication schedules ignored*

3. Select the **Ignore schedules** check box.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- When **Ignore schedules** is selected, any site link availability schedules that have been configured are ignored, and all site links become always available for replication. When **Ignore schedules** is cleared, any previously configured site link availability schedules go back into effect. If no site link availability schedules have been configured, selecting **Ignore schedules** has no effect.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure site link replication availability
Managing replication
Replication overview

# Enable or disable site link bridges

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable site link bridges**

1. Open Active Directory Sites and Services.

2. In the console tree, right-click the intersite transport folder (such as **IP** or **SMTP**) for which you want to enable or disable site link bridges, and then click **Properties**.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*intersite transport for which you want to enable or disable link bridges*

3. Do one of the following:

   - To enable site link bridges, select the **Bridge all site links** check box.

   - To disable site link bridges, clear the **Bridge all site links** check box.

### Important

- By default, all site links are bridged. For more information about disabling this option and its affects on intersite replication, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Replication overview
Managing replication

# Add a site to a site link

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a site to a site link**

1. Open Active Directory Sites and Services.

2. In the console tree, click the inter-site transport folder that contains the site link to which you are adding the site.

   **Where?**

   - Active Directory Sites and Services/Sites/Inter-Site Transports/*inter-site transport folder that contains the site link to which you are adding the site*

3. In the details pane, right-click the site link to which you want to add the site, and then click **Properties**.

4. Click the site you want to add to this site link, and then click **Add**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the forest root domain) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Replication overview
Sites overview
Managing replication

# Manually add connections

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manually add connections**

1. Open Active Directory Sites and Services.

2. In the console tree, right-click **NTDS Settings** for the domain controller for which you want to manually add a connection, and then click **New Active Directory Connection**.

   **Where?**

   - Active Directory Sites and Services/Sites/*site that contains the domain controller for which you want to manually add a connection*/Servers/*domain controller for which you want to manually add a connection*/NTDS Settings

3. In the **Find Domain Controllers** dialog box, click the domain controller that you want to include in the connection object.

4. In the **New Object-Connection** dialog box, type a name for the new connection object.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- Active Directory automatically creates and deletes connection objects under normal conditions.

- If you are certain that a connection is required and that you want it to persist until it is manually removed, create a connection manually.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Replication overview
Troubleshooting replication

# Force replication over a connection

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To force replication over a connection**

1. Open Active Directory Sites and Services.

2. In the console tree, click **NTDS Settings** for the server that you want to force replication.

   **Where?**

   - Active Directory Sites and Services/Sites/*site that contains the connection over which you want to replicate directory information*/Servers/*server that you want to force replication*/NTDS Settings

3. In the details pane, right-click the connection over which you want to replicate directory information, and then click **Replicate Now**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group (in the domain of the selected domain controller) or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority.

- If you want to synchronize the configuration and schema directory partitions on a domain controller in a child domain, you must be a member of the Domain Admins groups in the forest root domain or the Enterprise Admins group in the forest.

- As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- You can also use the **repadmin** and **replmon** command-line tools, available on your operating system installation CD, to force replication. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Replication overview
Active Directory support tools

# Manage COM+ Partitions and Partition Sets

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage COM+ partitions and partition sets**

- Create a COM+ partition in Active Directory
- Create a COM+ partition set in Active Directory
- Map a user or organizational unit to a COM+ partition set
- View COM+ partitions and partition sets

# Create a COM+ partition in Active Directory

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a COM+ partition in Active Directory**

1. Open Active Directory Users and Computers.
2. On the **View** menu, select **Advanced Features**.
3. In the console tree, right-click **ComPartitions**, and then click **New Partition**.

   **Where?**

   - Active Directory Users and Computers\domain node\System\ComPartitions

4. Under **Partition Name**, type a name for the new partition.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Once you have created a COM+ partition using Active Directory Users and Computers, you must then create an empty partition on a application server and associate it with the appropriate COM+ partition in Active Directory. For more information, see Related Topics.
- When creating a new COM+ partition, a partition ID is automatically generated and displayed under **Partition ID** in the **Properties** dialog box. If you want to specify a different GUID for the new COM+ partition, delete the existing partition ID and type a new one.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a COM+ partition set in Active Directory
Map a user or organizational unit to a COM+ partition set
View COM+ partitions and partition sets
Managing COM+ partitions in Active Directory
Managing COM+ partition sets in Active Directory

# Create a COM+ partition set in Active Directory

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a COM+ partition set in Active Directory**

1. Open Active Directory Users and Computers.
2. On the **View** menu, select **Advanced Features**.
3. In the console tree, right-click **ComPartitionSets**, and then click **New Partition Set**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/System/ComPartitionSets

4. Under **Name** and **Description**, type the name and description of the new partition set, and then click **Next**.
5. On the **Configure Partitions for New Set** page, click the partitions that you want to include in the partition set, click **Add**, and then click **Next**.
6. On the **Configure Users for New Set** page, click **Add User** and **Add Org Unit** to add the users and organizational units that are to be mapped to the partition set, and then click **Finish**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- By using COM+ partition sets, you can associate multiple users or organizational units with a COM+ partition set at one time instead of having to map multiple user identities or organizational units. A COM+ partition set is available to users or organizational units that are mapped to it.
- When creating a new COM+ partition set, a partition ID is automatically generated and displayed under **Partition Set ID** in the **Properties** dialog box. If you want to specify a different GUID for the new COM+ partition set, delete the existing partition ID and type a new one.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a COM+ partition in Active Directory
Map a user or organizational unit to a COM+ partition set
View COM+ partitions and partition sets
Managing COM+ partitions in Active Directory
Managing COM+ partition sets in Active Directory

# Map a user or organizational unit to a COM+ partition set

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To map a user or organizational unit to a COM+ partition set**

1. Open Active Directory Users and Computers.

2. On the **View** menu, select **Advanced Features**.

3. Do one of the following:

   **At the user or organizational unit**

   a. In the console tree, right-click the user name or organizational unit that you want to map to a partition set, and then click **Properties**.

   b. On the **COM+** tab, under **Partition Set**, click the partition set to which you want to map the user or organizational unit.

   **At the COM+ partition set**

   a. In the console tree, click **ComPartitionSets**.

   **Where?**

   - Active Directory Users and Computers/*domain node*/System/ComPartitionSets

   a. Right-click the COM+ partition set to which you want to map a user or organizational unit, and then click **Configure Set Members**.

   b. In the **Add/Remove Users** dialog box, click **Add User** and **Add Org Unit** to add the users and organizational units that you want to map to the partition set, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- By using COM+ partition sets, you can associate multiple users or organizational units with a COM+ partition set at one time instead of having to map multiple user identities or organizational units. A COM+ partition set is available to users or organizational units that are mapped to it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Create a COM+ partition in Active Directory
Create a COM+ partition set in Active Directory
View COM+ partitions and partition sets
Managing COM+ partitions in Active Directory
Managing COM+ partition sets in Active Directory

# View COM+ partitions and partition sets

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view COM+ partitions and partition sets**

1. Open Active Directory Users and Computers.
2. On the **View** menu, select **Advanced Features**.
3. Do one of the following:
   - To view partitions, in the console tree, click **ComPartitions**.

     **Where?**

     - Active Directory Users and Computers/*domain node*/System/ComPartitions
   - To view partition sets, in the console tree, click **ComPartitionSets**.

     **Where?**

     - Active Directory Users and Computers/*domain node*/System/ComPartitionSets

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To view the properties of a COM+ partition or partition set, right-click it, and then click **Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a COM+ partition in Active Directory
Create a COM+ partition set in Active Directory
Map a user or organizational unit to a COM+ partition set
Managing COM+ partitions in Active Directory
Managing COM+ partition sets in Active Directory

# Search Active Directory

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Search Active Directory**

- Find a user account
- Find a contact
- Find a printer
- Find a shared folder
- Find a group
- Find a computer account
- Find an organizational unit
- Find a domain controller
- Perform a custom search
- Create a saved query
- Export or import a saved query

# Find a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a user account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. If you want to search the entire domain, in the console tree, right-click the domain node, and then click **Find**.

   Or, if you know which organizational unit the user is in, in the console tree, right-click the organizational unit, and then click **Find**.
3. In **Name**, type the name of the user you want to find.
4. Click **Find Now**.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- You can search using partial search criteria. For example, *B* will return all objects whose name begins with the letter *B*, such as Backup Operators.
- Use the **Advanced** tab for more powerful search options.
- To find a user, you can also click

  on the toolbar.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsquery user***parameter*

| Value | Description |
|---|---|
| *parameter* | For the list of parameters, see Dsquery. |

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **dsquery user /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Finding directory information
Command-line reference A-Z

# Find a contact

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a contact**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. If you want to search the entire domain, in the console tree, right-click the domain node, and then click **Find**.

   Or, if you know which organizational unit the contact is in, in the console tree, right-click the organizational unit, and then click **Find**.
3. On the **Users, Contacts, and Groups** tab, in **Name**, type the name of the contact you want to find.
4. Click **Find Now**.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- You can search using partial search criteria. For example, *B* will return all objects whose name begins with the letter *B*, such as Backup Operators.
- Use the **Advanced** tab for more powerful search options.
- To find a contact, you can also click

  on the toolbar.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsquery contact***parameter*

| Value | Description |
|---|---|
| *parameter* | For the list of parameters, see Dsquery. |

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **dsquery contact /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Finding directory information

# Find a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a printer**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain node, and then click **Find**.
3. In the **Find** list, click **Printers**.
4. On the **Printers** tab, in **Name**, type the name of the printer you want to find.
5. Click **Find Now**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- You can also click the **Features** tab to find a printer with particular features, or click the **Advanced** tab for more powerful search options.
- To find printers, you can also click

  on the toolbar.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Finding directory information

# Find a shared folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a shared folder**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain node, and then click **Find**.
3. In the **Find** list, click **Shared Folders**.
4. On the **Shared Folders** tab, in **Named**, type the name of the folder you want to find.
5. Click **Find Now**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Use the **Advanced** tab for more powerful search options.
- To find shared folders, you can also click

  on the toolbar.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Finding directory information

# Find a group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a group**

- Using the Windows interface
- Using the Windows interface

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain node, and then click **Find**.
3. On the **Users, Contacts, and Groups** tab, in **Name**, type the name of the group you want to find.
4. Click **Find Now**.

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- You can search using partial search criteria. For example, *B* will return all objects whose name begins with the letter *B*, such as Backup Operators.
- Use the **Advanced** tab for more powerful search options.
- To find groups, you can also click

  on the toolbar.

**Using the Windows interface**

1. Open Command Prompt.
2. Type:

   **dsquery group***parameter*

| Value | Description |
|---|---|
| *parameter* | There are several attributes of a group that you can search by using this command. For the list of parameters, see Dsquery. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Directory service command-line tools
Working with MMC console files
Groups
Command-line reference A-Z

# Find a computer account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a computer account**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. If you want to search the entire domain, in the console tree, right-click the domain node, and then click **Find**.

   Or, if you know which organizational unit the computer is in, in the console tree, right-click the organizational unit, and then click **Find**.
3. In **Find**, click **Computers**.
4. In **Name**, type the name of the computer you want to find.
5. To find only domain controllers, in **Role**, click **Domain Controller**.

   Or, to find only workstations and servers (not domain controllers), in **Role**, click **Workstations and Servers**.
6. Click **Find Now**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To configure more powerful search options, click the **Advanced** tab.
- To find a computer, you can also click

  on the toolbar.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsquery computer-name***Name*

| Value | Description |
|---|---|

| | |
|---|---|
| *Name* | Searches for computers whose name attributes (value of CN attribute) matches *Name*. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **dsquery computer /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsquery
Working with MMC console files
Finding directory information
Command-line reference A-Z

# Find an organizational unit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find an organizational unit**

Using the Windows interface

Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain node, and then click **Find**.
3. In the **Find** list, click **Organizational Units**.
4. Type the name of the organizational unit you want to find, and then click **Find Now**.

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Use the **Advanced** tab for more powerful search options.
- To find organizational units, you can also click

  on the toolbar.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsquery ou***parameter*

| Value | Description |
|---|---|
| *parameter* | For the list of parameters, see Dsquery. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at a command prompt, type:

  **dsquery ou /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dsquery
Working with MMC console files
Organizational units
Command-line reference A-Z

# Find a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a domain controller**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click any node or folder, and then click **Find**.

   Or, if you know which folder contains the domain controller, right-click that folder to narrow the search to that folder.
3. In **Find**, click **Computers** and in **Role**, click **Domain Controller**.

   If you do not know which folder contains the domain controller, in **In**, click the domain you want to search or click **Entire Directory**.
4. Click **Find Now**.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Use the **Advanced** tab for more powerful search options.
- To find a computer, you can also click

  on the toolbar.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsquery server***Parameter*

| Value | Description |
|---|---|
| *Parameter* | There are several attributes of a server that you can search by using this command. For the list of parameters, see Dsquery. |

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice,

consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Directory service command-line tools
Working with MMC console files
Finding directory information
Command-line reference A-Z

# Perform a custom search

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To perform a custom search**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain node, and then click **Find**.
3. In **Find**, click **Custom Search**.
4. Click **Field**, point to the kind of object for which to search, and then click the attribute of that object for which you want to specify a search value.
5. In **Condition**, click the condition for your search.
6. In **Value**, type the attribute value for which the search will apply the condition.
7. Click **Add** to add this search condition to the custom search.
8. Repeat steps 4 through 7 until you have added all of the desired search conditions.
9. Click **Find Now**.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- Use the **Advanced** tab to enter a query using LDAP syntax. For more information about LDAP, see Related Topics.
- To start a search, you can also click

  on the toolbar.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dsquery** *****_Parameter_

| Value | Description |
| --- | --- |
| *Parameter* | There are several attributes that you can search by using this command. For the list of parameters, see Dsquery. For more information about LDAP searches, see Related Topics. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Directory service command-line tools
Working with MMC console files
Directory access protocol
Command-line reference A-Z

# Create a saved query

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a saved query**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click **Saved Queries** or any of its subfolders in which you want to save a query.

   **Where?**

   - Active Directory Users and Computers/Saved Queries

3. Point to **New**, and then click **Query**.
4. In **Name**, type a query name.
5. In **Description**, type a query description.
6. Click **Browse** to define the container from which to begin your search.

   To search all subcontainers of the selected container, select the **Include subcontainers** check box.

7. Click **Define Query** to define your query.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To organize your saved queries, you can create subfolders by right-clicking **Saved Query**, pointing to **New**, and then clicking **Folder**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using saved queries
Export or import a saved query
Working with MMC console files

# Export or import a saved query

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export or import a saved query**

1. Open Active Directory Users and Computers.
2. Do one of the following:
   - To export a saved query, in the Saved Queries folder, or a subfolder of Saved Queries, right-click the query you want to export, and then click **Export Query Definition**. Name the file, and then click **Open**.
   - To import a saved query, right-click **Saved Queries**, or a subfolder of Saved Queries, and then click **Import Query Definition**. Browse to the saved query, click it, and then click **Open**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using saved queries
Create a saved query
Working with MMC console files

# Active Directory Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Active Directory®:

- Active Directory Overview
- Understanding Active Directory
- Deploying Active Directory
- Administering Active Directory
- Active Directory Resources

# Active Directory Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory overview**
A directory is a hierarchical structure that stores information about objects on the network. A directory service, such as Active Directory Domain Services (AD DS), provides the methods for storing directory data and making this data available to network users and administrators. For example, AD DS stores information about user accounts, such as names, passwords, phone numbers, and so on, and enables other authorized users on the same network to access this information.

This section covers:

- Introduction to Active Directory
- New features for Active Directory
- Security information for Active Directory

# Introduction to Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introduction to Active Directory**

Active Directory can be installed on servers running Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. Active Directory stores information about objects on the network and makes this information easy for administrators and users to find and use. Active Directory uses a structured data store as the basis for a logical, hierarchical organization of directory information.

This data store, also known as the directory, contains information about Active Directory objects. These objects typically include shared resources such as servers, volumes, printers, and the network user and computer accounts. For more information about the Active Directory data store, see Directory data store.

Security is integrated with Active Directory through logon authentication and access control to objects in the directory. With a single network logon, administrators can manage directory data and organization throughout their network, and authorized network users can access resources anywhere on the network. Policy-based administration eases the management of even the most complex network. For more information about Active Directory security, see Security overview.

Active Directory also includes:

- A set of rules, the schema, that defines the classes of objects and attributes contained in the directory, the constraints and limits on instances of these objects, and the format of their names. For more information about the schema, see Schema.
- A global catalog that contains information about every object in the directory. This allows users and administrators to find directory information regardless of which domain in the directory actually contains the data. For more information about the global catalog, see The role of the global catalog.
- A query and index mechanism, so that objects and their properties can be published and found by network users or applications. For more information about querying the directory, see Finding directory information.
- A replication service that distributes directory data across a network. All domain controllers in a domain participate in replication and contain a complete copy of all directory information for their domain. Any change to directory data is replicated to all domain controllers in the domain. For more information about Active Directory replication, see Replication overview.
- Support for Active Directory client software, which makes many features on Microsoft® Windows® 2000 Professional or Windows XP Professional available to computers running Windows 95, Windows 98, and Windows NT® Server 4.0. To client computers not running Active Directory client software, the directory will appear like a Windows NT directory. For more information about client software, see Active Directory clients.

**Note**

- You cannot install Active Directory on a server running Windows Server 2003, Web Edition, but you can join the server to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

# New features for Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New Active Directory features in Windows Server 2003 with Service Pack 1 (SP1)**

The following list summarizes the Active Directory features that are new since the original release of Windows Server 2003.

- **Directory service backup reminders.** A new event message, event ID 2089, provides the backup status of each directory partition that a domain controller stores, including application directory partitions and Active Directory Application Mode (ADAM) partitions. If halfway through the tombstone lifetime a partition has not been backed up, this event is logged in the Directory Service event log and continues daily until the partition is backed up.

- **Added replication security and fewer replication errors.** Replication metadata for domain controllers from which Active Directory has been removed is no longer retained by default, although a waiting period can be configured. This change improves replication security and eliminates replication error messages that are caused by failed attempts to replicate with decommissioned domain controllers. For more information about preserving replication metadata, see How the Active Directory Replication Model Works.

- **Install from Media improvement for installing DNS servers.** Install from Media improvements make it easier to create a new domain controller that is a Domain Name System (DNS) server by providing a new option to include application directory partitions in the backup media that is used to install the new domain controller. This option eliminates the requirement for replication of the DomainDNSZones and ForestDNSZones application directory partitions before the DNS server is operational.

- **Enhancements for replication and DNS testing.** The Dcdiag.exe command-line tool, which is available in Windows Support Tools, provides new reporting on the overall health of replication with respect to Active Directory security. This test provides a summary of results, along with detailed information for each domain controller that is tested and a diagnosis of any security errors. Dcdiag.exe also has new DNS tests for connectivity, service availability, forwarders and root hints, delegation, dynamic update, locator record registrations, external name resolution, and enterprise infrastructure. These tests can be performed on one domain controller or on all domain controllers in a forest. For more information about using Dcdiag.exe, see Windows Support Tools Help.

- **Support for running domain controllers in virtual machines.** On a single physical server that is running Windows Server 2003 and Microsoft Virtual Server 2005, you can install multiple Windows Server 2003 or Windows 2000 Server domain controllers in separate virtual machines. This platform is well suited for test environments. By using virtual machines, you can effectively host multiple domains, multiple domain controllers for the same domain, or even multiple forests on one physical server that is running a single operating system. Windows Server 2003 SP1 also provides protection against directory corruption that can result from improper backup and restore of domain controller images. For more information about running domain controllers in virtual machines, see Running Domain Controllers in Virtual Server 2005.

- **Operations master health and status reporting.** If an operation that requires a domain controller that holds an operations master role (also known as flexible single-master operations (FSMO)) cannot be performed, events are now logged in the Directory Service event log. Events identify role holders that do not exist, exist but are not available, or are available but have not replicated recently with the contacting domain controller. For more information about operations masters, see How Operations Masters Work.

- **Extended storage of deleted objects.** The default period that a copy of a deleted object is retained in Active Directory, called the tombstone lifetime, is extended from 60 days to 180 days. Longer tombstone lifetime decreases the chance that a deleted object remains in the local directory of a disconnected domain controller beyond the time when the object is permanently deleted from online domain controllers. The tombstone lifetime is not changed automatically when you upgrade to Windows Server 2003 with SP1, but you can change the tombstone lifetime manually after the upgrade. New forests that are installed with Windows Server 2003 with SP1 have a default tombstone lifetime of 180 days. For more information about tombstone lifetime, see How the Data Store Works.

- **Improved domain controller name resolution.** In response to DNS name resolution failures that may be encountered

during location of replication partners and global catalog servers, domain controllers running Windows Server 2003 with SP1 request other variations of the server name that might be registered, which results in fewer failures due to DNS delays and misconfiguration. For more information about DNS name resolution, see How DNS Support for Active Directory Works.

- **Improved server metadata removal.** The Ntdsutil.exe command-line tool for managing the Active Directory database has new functionality that makes it easier to remove domain controller metadata. Preliminary steps, such as connecting to a server, domain, and site, are no longer required. You simply specify the server to remove. You can also specify the server on which to perform the deletion. Metadata removal is now more comprehensive: in addition to Active Directory replication metadata, the tool now removes File replication service (FRS) metadata and operations master metadata. If an operations master role is assigned to the server that is being removed, the tool attempts to transfer the role to an appropriate domain controller. For more information, see Delete extinct server metadata.

- **Improved security to protect confidential attributes.** To prevent Read access to confidential attributes, such as a Social Security number, while allowing Read access to other object attributes, you can designate specific attributes as confidential by setting a search flag on the respective attributeSchema object. By default, only domain administrators have Read access to confidential attributes, but this access can be delegated. For more information about access to attributes, see How Security Descriptors and Access Control Lists Work.

- **Retention of SID history on tombstones.** The **sIDHistory** attribute has been added to the set of attributes that are retained on an object tombstone when the object is deleted. If a tombstoned object is reactivated (undeleted), the **sIDHistory** attribute is now restored with the object. For more information about tombstones, see How the Data Store Works.

- **Adprep.exe improvements for Windows 2000 Server upgrades.** The Adprep tool has been improved to reduce the impact of FRS synchronization that results from updating SYSVOL files during upgrade. Adprep is used to upgrade the Windows 2000 Server schema to the Windows Server 2003 schema and to update some forest- and domain-specific configuration, including SYSVOL, that is required for a Windows Server 2003 domain controller to be operational. The tool now allows performing SYSVOL operations in a separate step when the domain is prepared for upgrade. A new switch, **/gpprep**, has been added to accommodate the SYSVOL updates, which can be performed at a convenient time following the upgrade. The **adprep /domainprep** command, which formerly performed both directory and SYSVOL updates, now updates only the directory. Adprep also now detects third-party schema extensions that block an upgrade, identifies the blocking extensions, and recommends fixes. Microsoft Exchange schema objects are also detected so that the Exchange schema can be prepared appropriately to accommodate inetOrgPerson naming. For more information about Adprep.exe, see Adprep.

- **Improved authoritative restore.** The **authoritative restore** option in Ntdsutil now locates backlinks for all objects that are authoritatively restored, including links that were created before implementation of the Windows Server 2003 or Windows Server 2003 interim forest functional level, in which linked-value replication (LVR) functionality was introduced. For example, suppose that a user object is restored and the user belongs to group G1, which was created before the forest functional level was raised, and the user also belongs to group G2, which was created after the forest functional level was raised. During authoritative restore of the user object, the member attribute of G2 is updated, but not the member attribute of G1. Ntdsutil now creates a text file that identifies the authoritatively restored objects and uses this file to create an LDAP Data Interchange Format (LDIF) file that can be used to restore all backlinks for pre-LVR groups in this domain. In the example, when this LDIF file is run after authoritative restore, the restored user is added to group G1. A new option in **authoritative restore** also allows you to generate an LDIF file that you can use to restore links in other domains in which a restored object has backlinks.

**New Active Directory features in Windows Server 2003**

With the new Active Directory features available in Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, more efficient administration of Active Directory is available to you.

The following list summarizes the Active Directory features that are available by default on any domain controller running Windows Server 2003.

- **Multiple selection of user objects.** Modify common attributes of multiple user objects at one time.

- **Drag-and-drop functionality.** Move Active Directory objects from container to container by dragging one or more objects to a desired location in the domain hierarchy. You can also add objects to group membership lists by dragging one or more objects (including other group objects) to the target group.

- **Efficient search capabilities.** Search functionality is object-oriented and provides an efficient search that minimizes network traffic associated with browsing objects. For more information, see Finding directory information.

- **Saved queries.** Save commonly used search parameters for reuse in Active Directory Users and Computers. For more information, see Using saved queries.

- **Active Directory command-line tools.** Run new directory service commands for administration scenarios. For more information, see Managing Active Directory from the command line.

- **InetOrgPerson class.** The **inetOrgPerson** class has been added to the base schema as a security principal and can be used in the same manner as the user class. The **userPassword** attribute can also be used to set the account password. For more information, see User and computer accounts.

- **Application directory partitions.** Configure the replication scope for application-specific data among domain controllers. For example, you can control the replication scope of Domain Name System (DNS) zone data stored in Active Directory so that only specific domain controllers in the forest participate in DNS zone replication. For more information, see Application directory partitions.

- **Ability to add additional domain controllers using backup media.** Reduce the time it takes to add an additional domain controller in an existing domain by using backup media. For more information, see Using the Active Directory Installation Wizard.

- **Universal group membership caching.** Prevent the need to locate a global catalog across a wide area network (WAN) when logging on by storing universal group membership information on an authenticating domain controller. For more information, see Global catalogs and sites.

- **Secure LDAP traffic.** Active Directory administrative tools sign and encrypt all Lightweight Directory Access Protocol (LDAP) traffic by default. Signing LDAP traffic guarantees that the packaged data comes from a known source and that it has not been tampered with. For more information, see Connecting to domain controllers running Windows 2000.

- **Active Directory quotas.** Quotas can be specified in Active Directory to control the number of objects a user, group, or computer can own in a given directory partition. Domain Administrators and Enterprise Administrators are exempt from quotas.

#### New domain- and forest-wide Active Directory features

New domain- or forest-wide Active Directory features can be enabled only when all domain controllers in a domain or forest are running Windows Server 2003 and the domain functionality or forest functionality has been set to Windows Server 2003. For more information about domain and forest functionality settings, see Domain and forest functionality.

The following list summarizes the domain- and forest-wide Active Directory features that can be enabled when either a domain or forest functional level has been raised to Windows Server 2003.

- **Domain controller rename tool.** Rename domain controllers without first demoting them. For more information, see Renaming domain controllers.

- **Domain rename.** Rename any Windows Server 2003 domain. You can change the NetBIOS name or DNS name of any child, parent, tree, or forest root domain. For more information, see Renaming domains.

- **Different location option for user and computer accounts.** You can now redirect the default location for user accounts and computer accounts that are created by the following application programming interfaces (APIs): NetUserAdd, NetGroupAdd, and NetJoinDomain. You can redirect the location of the accounts from the Users and Computers containers to organizational units (OUs) where Group Policy settings can be applied. For more information, see Redirect the Users and Computers Containers.

- **Forest trusts.** Create a forest trust to extend two-way transitivity beyond the scope of a single forest to a second forest. For more information, see Forest trusts.

- **Forest restructuring.** Move existing domains to other locations in the domain hierarchy. For more information, see Renaming domains.

- **Defunct schema objects.** Deactivate unnecessary classes or attributes from the schema. For more information, see Deactivating a class or attribute.

- **Dynamic auxiliary classes.** Provides support for dynamically linking auxiliary classes to individual objects, and not just to entire classes of objects. In addition, auxiliary classes that have been attached to an object instance can subsequently be removed from the instance.

- **Global catalog replication improvements.** Preserves the synchronization state of the global catalog when an administrative action results in an extension of the partial attribute set. This minimizes the replication traffic as a result of a partial attribute set extension by only transmitting attributes that were added. For more information, see Global catalog replication.

- **Replication enhancements.** Linked-value replication allows individual group members to be replicated across the network instead of treating the entire group membership as a single unit of replication. For more information about linked-value replication, see How replication works. In addition, new spanning tree algorithms make replication more efficient, as well as more scalable across a larger number of domains and sites in both Windows 2000 and Windows Server 2003 forests. For more information, see Replication overview.

- **User access control to resources between domains or forests.** Block users in a domain or forest from accessing resources in another domain or forest, and then allow selective access by setting the Allow to authenticate access control entry (ACE) on a local resource for the user or group object. For more information, see Accessing resources across domains or Accessing resources across forests.

# Security information for Active Directory

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for Active Directory**

Active Directory® provides a secure directory environment for your organization using built-in logon authentication and user authorization, which are core features of the Local Security Authority (LSA). Logon authentication and user authorization are available by default and provide immediate protection for network access and network resources.

For information about additional security tools and features that you can use to further secure Active Directory, see Securing Active Directory and Active Directory Best practices.

**Protecting access to the network**

Active Directory requires confirmation of the identity of a user before allowing access to the network, a process known as authentication. Users only need to provide a single sign-on to the domain (or to trusted domains) to gain access to the network. Once Active Directory confirms the identity of the user, the LSA on the authenticating domain controller generates an access token that determines what level of access that user has on network resources. For more information about the authentication process, see Access control in Active Directory and Certificates and Authentication.

Active Directory supports a number of secure Internet-standard protocols and authentication mechanisms used to prove identity upon logon, including Kerberos V5, X.509 v3 certificates, smart cards, public key infrastructure (PKI) and Lightweight Directory Access Protocol (LDAP) using Secure Sockets Layer (SSL).

Authentication between domains occurs through trusts. A trust is a relationship established between two or more domains to allow users in one domain to be authenticated by a domain controller in another domain.

Trust relationships can be transitive or nontransitive but must always be present in order for users in one domain to access shared resources in another domain. For more information, see Trusts.

In addition to securing network access through authentication, Active Directory helps to protect shared resources by facilitating user authorization. Once a user logon has been authenticated by Active Directory, the user rights assigned to the user through security groups and the permissions assigned on the shared resource will determine if the user will be authorized to access that resource. This authorization process protects shared resources from unauthorized access and permits access to only authorized users or groups.

For more information about user rights and security groups, see Group types. For more information about the user rights assigned to default groups, see Default groups. For more information about authorization, see "Designing an Authorization Strategy" at the Microsoft Windows Resource Kits Web site.

For more information about authentication, see "Logon and Authentication" at the Microsoft Windows Resource Kits Web site. For more information about authorization and access control, see "Authorization and Access Control" at the Microsoft Windows Resource Kits Web site.

# Understanding Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Active Directory**
Active Directory is an implementation of Internet standard directory and naming protocols. It uses a database engine for transactional support and supports a variety of application programming interface standards.

This section covers:

- Securing Active Directory
- Access control in Active Directory
- Active Directory naming
- Directory data store
- Directory access protocol
- User and computer accounts
- Object names
- Organizational units
- Active Directory server roles
- Active Directory clients
- Understanding Domains and Forests
- Understanding Groups
- Understanding Trusts
- Understanding Sites and Replication
- Understanding the Global Catalog
- Interoperating with DNS and Group Policy
- Understanding Schema

# Securing Active Directory

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing Active Directory**
Active Directory provides a secure directory environment for your organization using built-in logon authentication and user authorization. To further secure Active Directory once it has been deployed, consider the following precautions and recommendations.

For general security information about Active Directory, see Security information for Active Directory.

For more information about authentication, see "Logon and Authentication" at the Microsoft Windows Resource Kits Web Site. For more information about authorization, see "Authorization and Access Control" at the Microsoft Windows Resource Kits Web Site.

**Important**

- Physical access to a domain controller can provide a malicious user unauthorized access to encrypted passwords. Therefore, it is recommended that all domain controllers be locked in a secured room with limited public access. In addition, you should limit membership in the Enterprise Admins, Domain Admins, Account Operators, Server Operators, Print Operators, and Backup Operators groups to trusted personnel in your organization. For more information about domain controllers and groups, see Domain controllers and Default groups.

| To | Use |
|---|---|
| Manage the security relationship between two forests and simplify security administration and authentication across forests. | Forest trusts<br><br>See Forest trusts. |
| Force domain users to use strong passwords. | Group Policy<br><br>See Strong passwords. |
| Enable audit policy. Auditing event logs can notify you of actions that could pose a security risk. | Group Policy<br><br>See Auditing Policy. |
| Assign user rights to new security groups so you can specifically define a user's administrative role in the domain. | Group Policy<br><br>See Group types. |
| Enforce account lockouts on user accounts and decrease the possibility of an attacker compromising your domain through repeated logon attempts. | Group Policy<br><br>See User and computer accounts. |
| Enforce password history on user accounts and decrease the possibility of an attacker compromising your domain. | Group Policy<br><br>See Enforce password history. |
| Enforce minimum and maximum password ages on user accounts and decrease the possibility of an attacker compromising your domain. | Group Policy<br><br>See Minimum password age, Maximum password age. |
| Verify and authenticate the validity of each user through the use of public key cryptography. | Public key infrastructure<br><br>See Deploying a Public Key Infrastructure. |
| Promote a secure operating environment by running your computer without administrative credentials except when required. | Run as<br><br>See Using Run as. |

| | |
|---|---|
| Restrict user, group, and computer access to shared resources and filter Group Policy settings. | Security groups<br><br>See Group types. |
| Prevent attacks from malicious users who might try to grant elevated user rights to another user account. | SID filtering<br><br>See "Using Security Identifier (SID) Filtering to Prevent Elevation of Privilege Attacks" at the Microsoft Web Site. |
| Provide tamper-resistant user authentication and e-mail security. | Smart cards<br><br>See Smart cards overview. |
| Use strong encryption techniques to secure account password information on local computers, member servers, or domain controllers. | Syskey<br><br>See The system key utility. |

# Access control in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Access control in Active Directory**
Administrators can use access control to manage user access to shared resources for security purposes. In Active Directory, access control is administered at the object level by setting different levels of access, or permissions, to objects, such as Full Control, Write, Read, or No Access. Access control in Active Directory defines how different users can use Active Directory objects. By default, permissions on objects in Active Directory are set to the most secure setting.

The elements that define access control permissions on Active Directory objects include security descriptors, object inheritance, and user authentication.

**Security descriptors**
Access control permissions are assigned to shared objects and Active Directory objects to control how different users can use each object. A shared object, or shared resource, is an object that is intended to be used over a network by one or more users and includes files, printers, folders, and services. Both shared objects and Active Directory objects store access control permissions in security descriptors.

A security descriptor contains two access control lists (ACLs) used to assign and track security information for each object: the discretionary access control list (DACL) and the system access control list (SACL).

- **Discretionary access control lists (DACLs).** DACLs identify the users and groups that are assigned or denied access permissions on an object. If a DACL does not explicitly identify a user, or any groups that a user is a member of, the user will be denied access to that object. By default, a DACL is controlled by the owner of an object or the person who created the object, and it contains access control entries (ACEs) that determine user access to the object.

- **System access control lists (SACLs).** SACLs identify the users and groups that you want to audit when they successfully access or fail to access an object. Auditing is used to monitor events related to system or network security, to identify security breaches, and to determine the extent and location of any damage. By default, a SACL is controlled by the owner of an object or the person who created the object. A SACL contains access control entries (ACEs) that determine whether to record a successful or failed attempt by a user to access a object using a given permission, for example, Full Control and Read.

  For more information about auditing, see Auditing settings on objects.

To view DACLs and SACLs on Active Directory objects using Active Directory Users and Computers, on the **View** menu, click **Advanced Features** to access the **Security** tab for each object. For more information, see Assign, change, or remove permissions on Active Directory objects or attributes. You can also use the DSACLS support tool to manage access control lists in Active Directory. For more information, see Active Directory support tools.

By default, DACLs and SACLs are associated with every Active Directory object, which reduces attacks to your network by malicious users or any accidental mistakes made by domain users. However, if a malicious user obtains a user name and password of any account with administrative credentials to Active Directory, your forest will be vulnerable to attack. For this reason, consider renaming or disabling the default administrator account and following the best practices as described in Active Directory Best practices.

**Object inheritance**
By default, Active Directory objects inherit ACEs from the security descriptor located in their parent container object. Inheritance enables the access control information defined at a container object in Active Directory to apply to the security descriptors of any subordinate objects, including other containers and their objects. This eliminates the need to apply permissions each time a child object is created. If necessary, you can change the inherited permissions. However, as a best practice, avoid changing the default permissions or inheritance settings on Active Directory objects. For more information, see Best practices for assigning permissions on Active Directory objects and Changing inherited permissions.

**User authentication**
Active Directory also authenticates and authorizes users, groups, and computers to access objects on the network. The Local Security Authority (LSA) is the security subsystem responsible for all interactive user authentication and authorization services on a local computer. The LSA is also used to process authentication requests made through the Kerberos V5 protocol or NTLM

protocol in Active Directory. For more information about Kerberos authentication, see Kerberos V5 authentication. For more information about NTLM authentication, see NTLM authentication.

Once the identity of a user has been confirmed in Active Directory, the LSA on the authenticating domain controller generates a user access token and associates a security ID (SID) with the user.

- **Access token.** When a user is authenticated, LSA creates a security access token for that user. An access token contains the user's name, the groups to which that user belongs, a SID for the user, and all of the SIDs for the groups to which the user belongs. If you add a user to a group after the user access token has been issued, the user must log off and log on again before the access token will be updated.

- **Security ID (SID).** Active Directory automatically assigns SIDs to security principal objects at the time they are created. Security principals are accounts in Active Directory that can be assigned permissions such as computer, group, or user accounts. Once a SID is issued to the authenticated user, it is attached to the access token of the user.

The information in the access token is used to determine a user's level of access to objects whenever the user attempts to access them. The SIDs in the access token are compared with the list of SIDs that make up the DACL for the object to ensure that the user has sufficient permission to access the object. This is because the access control process identifies user accounts by SID rather than by name.

**Important**

- When a domain controller provides an access token to a user, the access token only contains information about membership in domain local groups if the domain local groups are local to the domain controller's domain. For domain directory objects replicated in the global catalog, this fact requires certain security considerations. For more information, see Global catalog replication.

For more information about authentication, see "Logon and Authentication" at the Microsoft Windows Resource Kits Web site.

For more information about access control and permissions, see Access control overview. For more information about authorization and access control, see "Authorization and Access Control" at the Microsoft Windows Resource Kits Web site.

For information about additional security measures that can be implemented to protect Active Directory, see Securing Active Directory and Security information for Active Directory.

# Active Directory naming

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory naming**
Active Directory domain names are usually the full Domain Name System (DNS) name of the domain. However, for backward compatibility, each domain also has a pre-Windows 2000 name for use by computers running pre-Windows 2000 operating systems. The pre-Windows 2000 domain name can be used to log on to a Windows Server 2003 domain from computers running pre-Windows 2000 operating systems using the *DomainName\UserName* format. This same format can also be used to log on to a Windows Server 2003 domain from computers running Windows 2000, Windows XP, or servers running Windows Server 2003. Users can also log on to computers running Windows 2000, Windows XP, or servers running Windows Server 2003 using the user principal name (UPN) associated with their user account.

**User accounts**
In Active Directory, each user account has a user logon name, a pre-Windows 2000 user logon name (security account manager account name), and a UPN suffix. The administrator enters the user logon name and selects the UPN suffix when creating the user account. Active Directory suggests a pre-Windows 2000 user logon name using the first 20 bytes of the user logon name. Administrators can change the pre-Windows 2000 logon name at any time.

In Active Directory, each user account has a UPN based on IETF RFC 822, *Standard for the Format of ARPA Internet Text Messages*. The UPN is composed of the user logon name and the UPN suffix joined by the @ sign.

**Note**

- Do not add the @ sign to the user logon name or to the UPN suffix. Active Directory automatically adds it when it creates the UPN. A UPN that contains more than one @ sign is invalid.

**Important**

- Windows NT 4.0 and earlier domains allowed the use of a period (.) at the end of a user logon name as long as the user logon name did not consist solely of period characters. Windows Server 2003 domains do not allow the use of a period or multiple periods at the end of a user logon name.

The second part of the UPN, the UPN suffix, identifies the domain in which the user account is located. This UPN suffix can be the DNS domain name, the DNS name of any domain in the forest, or it can be an alternative name created by an administrator and used just for log on purposes. This alternative UPN suffix does not need to be a valid DNS name.

In Active Directory, the default UPN suffix is the DNS name of the domain in which user account created. In most cases, this is the domain name registered as the enterprise domain on the Internet. Using alternative domain names as the UPN suffix can provide additional logon security and simplify the names used to log on to another domain in the forest.

For example, if your organization uses a deep domain tree, organized by department and region, domain names can get quite long. The default user UPN for a user in that domain might be *sales.westcoast.microsoft.com*. The logon name for a user in that domain would be *user@sales.westcoast.microsoft.com*. Creating a UPN suffix of "microsoft" would allow that same user to log on using the much simpler logon name of *user@microsoft*. For more information about user accounts, see User and computer accounts and Object names.

You can add or remove UPN suffixes using Active Directory Domains and Trusts. For more information, see Add user principal name suffixes.

**Computer accounts**
Each computer account created in Active Directory has a relative distinguished name, a pre-Windows 2000 computer name (security account manager account name), a primary DNS suffix, a DNS host name, and a service principal name (SPN). The administrator enters the computer name when creating the computer account. This computer name is used as the Lightweight Directory Access Protocol (LDAP) relative distinguished name.

Active Directory suggests the pre-Windows 2000 name using the first 15 bytes of the relative distinguished name. The administrator can change the pre-Windows 2000 name at any time.

The DNS name for a host is called a full computer name and is a DNS fully qualified domain name (FQDN). The full computer

name is a concatenation of the computer name (the first 15 bytes of the SAM account name of the computer account without the "$" character) and the primary DNS suffix (the DNS domain name of the domain in which the computer account exists). It is listed on the **Computer Name** tab in System Properties in Control Panel.

By default, the primary DNS suffix portion of the FQDN for a computer must be the same as the name of the Active Directory domain where the computer is located. To allow different primary DNS suffixes, a domain administrator may create a restricted list of allowed suffixes by creating the msDS-AllowedDNSSuffixes attribute in the domain object container. This attribute is created and managed by the domain administrator using Active Directory Service Interfaces (ADSI) or the Lightweight Directory Access Protocol (LDAP).

For more information, see Object names and User and computer accounts.

The service principal name (SPN) is a multivalue attribute. It is usually built from the DNS name of the host. The SPN is used in the process of mutual authentication between the client and the server hosting a particular service. The client finds a computer account based on the SPN of the service to which it is trying to connect. The SPN can be modified by members of the Domain Admins group.

# Directory data store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Directory data store**

The Active Directory directory service uses a data store for all directory information. This data store is often referred to as the directory. The directory contains information about objects such as users, groups, computers, domains, organizational units, and security policies. This information can be published for use by users and administrators.

The directory is stored on domain controllers and can be accessed by network applications or services. A domain can have one or more domain controllers. Each domain controller has a copy of the directory for the entire domain in which it is located. Changes made to the directory on one domain controller are replicated to other domain controllers in the domain, domain tree, or forest. Active Directory uses four distinct directory partition types to store and copy different types of data. Directory partitions contain domain, configuration, schema, and application data. This storage and replication design provides directory information to users and administrators throughout the domain.

Directory data is stored in the Ntds.dit file on the domain controller. It is recommended that you store this file on an NTFS partition. For more information about the tool used to manage the Active Directory database and log files, see **Files** in Ntdsutil. Private data is stored securely, and public directory data is stored on a shared system volume where it can be replicated to other domain controllers in the domain. For more information about replication, see Replication overview.

Directory data replicated between domain controllers includes the following:

- **Domain data**

  The domain data holds information about objects within a domain. This is information such as e-mail contacts, user and computer account attributes, and published resources that are of interest to administrators and users.

  For example, when a user account is added to your network, a user account object and attribute data are stored in the domain data. When changes to your organization's directory objects occur, such as object creation, deletion, or attribute modification, this data is stored in the domain data.

- **Configuration data**

  The configuration data describes the topology of the directory. This configuration data includes a list of all domains, trees, and forests and the locations of the domain controllers and global catalogs.

- **Schema data**

  The schema is the formal definition of all object and attribute data that can be stored in the directory. Domain controllers running Windows Server 2003 include a default schema that defines many object types, such as user and computer accounts, groups, domains, organizational units, and security policies. Administrators and programmers can extend the schema by defining new object types and attributes or by adding new attributes for existing objects. Schema objects are protected by access control lists, ensuring that only authorized users can alter the schema.

  For more information, see Schema.

- **Application data**

  Data stored in the application directory partition is intended to satisfy cases where information needs to be replicated but not necessarily on a global scale. Application directory partitions are not part of the directory data store by default; they must be created, configured, and managed by the administrator.

  For more information, see Application directory partitions.

**Note**

- If a domain controller is also a global catalog, it stores a subset of the directory data for all other domains in the forest.

For more information about domain controllers, see Domain controllers. For more information about the global catalog, see The role of the global catalog.

**Quotas and directory partitions**

Quotas, a new feature with domain controllers running Windows Server 2003 , determine the number of objects that can be owned in a given directory partition by a security principal. (The owner of an object is usually, but not always, the creator of the object.) Quotas can help prevent the denial of service that can occur if a security principal accidentally, or intentionally, creates objects until the affected domain controller runs out of storage space.

Quotas are specified and administered for each directory partition separately. The schema partition, however, has no quotas. On a given directory partition, you can assign quotas for any security principal, including users, inetOrgPersons, computers, and groups. Members of the Domain Admins and Enterprise Admins groups are exempt from quotas. In some cases, a security principal might be covered by multiple quotas. For example, a user might be assigned an individual quota, and also belong to one or more security groups that also have quotas assigned to them. In such cases, the effective quota is the maximum of the quotas assigned to the security principal.

If a security principal is not assigned a quota either directly or through a group membership, a default quota on the partition governs the security principal. If you do not explicitly set the default quota on a given partition, the default quota of that partition is unlimited (ie, there is no limit).

Tombstone objects owned by a security principal are also counted as part of the quota consumption of that security principal. For each partition, you can specify a tombstone quota factor to determine the percentage weight given to a tombstone object in quota accounting. For example, if the tombstone quota factor for a given partition is set to 25 (or 25%), then a tombstone object on the partition is counted as 0.25 (or ¼) of a normal object. If a quota of 100 is specified for a user on this partition, then the user could own a maximum of 100 normal objects, or a maximum 400 tombstone objects. The default tombstone quota factor for each partition is initially set to 100 (or 100%), meaning that normal and tombstones objects are weighted equally.

The following example illustrates how quotas can be used. Consider the domain "sales.northwindtraders.com." Because this domain supports a lot of printing activity, the domain contains several print servers that each support 1,000 or more print queues. Initially, the default quota of the sales.northwindtraders.com domain partition is set to unlimited. To help control the number of objects created and owned, the administrator specifies a default quota of 500. Now, each user can own a maximum of 500 objects on the partition. Because print queues are directory objects that are created and owned by the respective print servers, the new default quota of 500 limits each print server to 500 print queues. To remove this constraint, the administrator creates a group called "Print Servers" and adds the computer account of each print server to the group. The administrator then specifies a quota of 2,000 for the Print Servers group. Now, each print server can support its original number of print queues, while the default quota continues to prevent excess object creation by all other security principals.

Only domain controllers running Windows Server 2003 can enforce quotas. Quotas are enforced only on originating directory operations; quotas are not enforced on replicated operations. In order for quotas to be fully effective for any given directory partition, all domain controllers that contain a writable copy of that partition must be running Windows Server 2003 . Therefore, for quotas to be effective on a domain directory partition, all domain controllers in that domain must be running Windows Server 2003 . For quotas to be effective on the configuration partition, all domain controllers in the forest must by running Windows Server 2003 .

For information about creating, modifying, and querying quotas, default quotas, and tombstone quota factors, see Dsadd, Dsmod, and Dsquery.

# Directory access protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Directory access protocol**
Active Directory clients must communicate with domain controllers when logging on to the network and when searching for shared resources. Access to domain controllers and global catalogs is performed using the Lightweight Directory Access Protocol (LDAP).

**Lightweight Directory Access Protocol**
LDAP is a communication protocol designed for use on TCP/IP networks. LDAP defines how a directory client can access a directory server and how the client can perform directory operations and share directory data. LDAP standards are established by working groups of the Internet Engineering Task Force (IETF). Active Directory implements the LDAP attribute draft specifications and the IETF standards for LDAP versions 2 and 3.

As its name implies, LDAP is designed as an efficient method for accessing directory services without the complexity of other directory service protocols. Because LDAP defines what operations can be performed to query and modify information in a directory and how information in a directory can be securely accessed, you can use LDAP to find or enumerate directory objects and to query or administer Active Directory.

**LDAP and interoperability**
LDAP is an open Internet standard. By using LDAP, Active Directory enables interoperability with other vendor directory services. Active Directory support for LDAP includes an LDAP provider object as part of Active Directory Service Interfaces (ADSI). ADSI supports the C-binding application programming interfaces for LDAP. Other directory service applications can be easily modified to access information in Active Directory by using ADSI and LDAP.

For more information about LDAP, see the Internet Engineering Task Force at the Internet Engineering Task Force Web site.

### Note

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# User and computer accounts

Updated: July 26, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User and computer accounts**

Active Directory user accounts and computer accounts represent a physical entity such as a computer or person. User accounts can also be used as dedicated service accounts for some applications.

User accounts and computer accounts (as well as groups) are also referred to as security principals. Security principals are directory objects that are automatically assigned security IDs (SIDs), which can be used to access domain resources. A user or computer account is used to:

- Authenticate the identity of a user or computer.

  A user account enables a user to log on to computers and domains with an identity that can be authenticated by the domain. For information about authentication, see Access control in Active Directory. Each user who logs on to the network should have his or her own unique user account and password. To maximize security, you should avoid multiple users sharing one account.

- Authorize or deny access to domain resources.

  Once the user has been authenticated, the user is authorized or denied access to domain resources based on the explicit permissions assigned to that user on the resource. For more information, see Security information for Active Directory.

- Administer other security principals.

  Active Directory creates a foreign security principal object in the local domain to represent each security principal from a trusted external domain. For more information about foreign security principals, see When to create an external trust.

- Audit actions performed using the user or computer account.

  Auditing can help you monitor account security. For more information about auditing, see Auditing overview.

**User accounts**

The Users container located in Active Directory Users and Computers displays the three built-in user accounts: Administrator, Guest, and HelpAssistant. These built-in user accounts are created automatically when you create the domain.

Each built-in account has a different combination of rights and permissions. The Administrator account has the most extensive rights and permissions over the domain, while the Guest account has limited rights and permissions. The table below describes each default user account on domain controllers running Windows Server 2003.

| Default user account | Description |
|---|---|

| | |
|---|---|
| Administrator account | The Administrator account has full control of the domain and can assign user rights and access control permissions to domain users as necessary. This account must be used only for tasks that require administrative credentials. It is recommended that you set up this account with a strong password. For more information, see Strong passwords. For additional security considerations for accounts with administrative credentials, see Active Directory Best practices.<br><br>The Administrator account is a default member of the Administrators, Domain Admins, Enterprise Admins, Group Policy Creator Owners, and Schema Admins groups in Active Directory. The Administrator account can never be deleted or removed from the Administrators group, but it can be renamed or disabled. Because the Administrator account is known to exist on many versions of Windows, renaming or disabling this account will make it more difficult for malicious users to try and gain access to it. For more information about how to rename or disable a user account, see Rename a local user account or Disable or enable a user account.<br><br>The Administrator account is the first account created when you set up a new domain using the Active Directory Installation Wizard.<br><br>• **Important** When the Administrator account is disabled, it can still be used to gain access to a domain controller using Safe Mode. |
| Guest account | The Guest account is used by people who do not have an actual account in the domain. A user whose account is disabled (but not deleted) can also use the Guest account. The Guest account does not require a password.<br><br>You can set rights and permissions for the Guest account just like any user account. By default, the Guest account is a member of the built-in Guests group and the Domain Guests global group, which allows a user to log on to a domain. The Guest account is disabled by default, and it is recommended that it stay disabled. |
| HelpAssistant account (installed with a Remote Assistance session) | The primary account used to establish a Remote Assistance session. This account is created automatically when you request a Remote Assistance session and has limited access to the computer. The HelpAssistant account is managed by the Remote Desktop Help Session Manager service and will be automatically deleted if no Remote Assistance requests are pending. For more information about Remote Assistance, see Administering Remote Assistance. |

**Securing user accounts**

If built-in account rights and permissions are not modified or disabled by a network administrator, they could be used by a malicious user (or service) to illegally log on to a domain using the Administrator or Guest identity. A good security practice for protecting these accounts is to rename or disable them. Because it retains its security ID (SID), a renamed user account retains all its other properties, such as its description, password, group memberships, user profile, account information, and any assigned permissions and user rights.

To obtain the security of user authentication and authorization, create an individual user account for each user who will participate on your network by using Active Directory Users and Computers. Each user account (including the Administrator and Guest account) can then be added to a group to control the rights and permissions assigned to the account. Using accounts and groups that are appropriate for your network ensures that users logging on to a network can be identified and can access only the permitted resources.

You can help defend your domain from attackers by requiring strong passwords and implementing an account lockout policy. Strong passwords reduce the risk of intelligent guessing and dictionary attacks on passwords. For more information, see Strong passwords and Password Best practices for passwords.

An account lockout policy decreases the possibility of an attacker compromising your domain through repeated logon attempts. This is because an account lockout policy determines how many failed logon attempts a user account can have before it is disabled. For more information, see Apply or modify account lockout policy.

For more information about securing user accounts, see Securing Active Directory.

**Account options**

Each Active Directory user account has a number of account options that determine how someone logging on with that particular user account is authenticated on the network. You can use the following options to configure password settings and security-specific information for user accounts:

| Account option | Description |
| --- | --- |
| User must change password at next logon | Forces a user to change their password the next time the user logs on to the network. Use this option when you want to ensure that the user will be the only person to know their password. |
| User cannot change password | Prevents a user from changing their password. Use this option when you want to maintain control over a user account, such as for a guest or temporary account. |
| Password never expires | Prevents a user password from expiring. It is recommended that Service accounts should have this option enabled and should use strong passwords. For more information about strong passwords, see Strong passwords. |
| Store passwords using reversible encryption | Allows a user to log on to a Windows network from Apple computers. If a user is not logging on from an Apple computer, this option should not be used. For more information, see Store passwords using reversible encryption. |
| Account is disabled | Prevents a user from logging on with the selected account. Many administrators use disabled accounts as templates for common user accounts. For more information, see Disable or enable a user account. |
| Smart card is required for interactive logon | Requires that a user possess a smart card to log on to the network interactively. The user must also have a smart card reader attached to their computer and a valid personal identification number (PIN) for the smart card. When this option is selected, the password for the user account is automatically set to a random and complex value. For more information about smart cards, see Logging on to a computer with a smart card and Authentication process. |
| Account is trusted for delegation | Allows a service running under this account to perform operations on behalf of other user accounts on the network. A service running under a user account (otherwise known as a service account) that is trusted for delegation can impersonate a client to gain access to resources on the computer where the service is running or on other computers. In a forest set to the Windows Server 2003 functional level, this setting is found on the **Delegation** tab, and is only available for accounts that have been assigned service principal names (SPNs), as set using the **setspn** command from the Windows Support Tools. This is a security-sensitive capability and should be cautiously assigned. For more information, see Allow a user to be trusted for delegation and Delegating authentication.<br><br>This option is only available on domain controllers running Windows Server 2003 where the domain functionality is set to Windows 2000 mixed or Windows 2000 native. On domain controllers running Windows Server 2003 where the domain functional level is set to Windows Server 2003, the **Delegation** tab is used to configure delegation settings. The **Delegation** tab only appears for accounts which have an assigned SPN. For more information about domain functionality, see Domain and forest functionality. For more information about configuring delegation in a Windows Server 2003 domain, see Allow a user to be trusted for delegation. |
| Account is sensitive and cannot be delegated | Allows control over a user account, such as for a guest or temporary account. This option can be used if this account cannot be assigned for delegation by another account. |
| Use DES encryption types for this account | Provides support for the Data Encryption Standard (DES). DES supports multiple levels of encryption, including MPPE Standard (40-bit), MPPE Standard (56-bit), MPPE Strong (128-bit), IPSec DES (40-bit), IPSec 56-bit DES, and IPSec Triple DES (3DES). For more information about DES encryption, see Data encryption. |
| Do not require Kerberos preauthentication | Provides support for alternate implementations of the Kerberos protocol. Domain controllers running Windows 2000 or Windows Server 2003 can use other mechanisms to synchronize time. Because preauthentication provides additional security, use caution when enabling this option. For more information about Kerberos, see Kerberos V5 authentication. |

**InetOrgPerson accounts**

Active Directory provides support for the InetOrgPerson object class and its associated attributes defined in RFC 2798. The InetOrgPerson object class is used in several non-Microsoft LDAP and X.500 directory services to represent people within an organization.

Support for InetOrgPerson makes migrations from other LDAP directories to Active Directory more efficient. The InetOrgPerson object is derived from the user class and can be used as a security principal just like the user class. For

information about creating an inetOrgPerson user account, see Create a new user account.

When the domain functional level has been set to Windows Server 2003, you can set the userPassword attribute on InetOrgPerson and user objects as being the effective password just like you can with the *unicodePwd* attribute.

**Computer accounts**

Every computer running Windows NT, Windows 2000, Windows XP, or a server running Windows Server 2003 that joins a domain has a computer account. Similar to user accounts, computer accounts provide a means for authenticating and auditing computer access to the network and to domain resources. Each computer account must be unique.

**Note** Computers running Windows 95 and Windows 98 do not have advanced security features and are not assigned computer accounts.

User and computer accounts can be added, disabled, reset, and deleted using Active Directory Users and Computers. A computer account can also be created when you join a computer to a domain. For more information about user and computer accounts, see Active Directory naming and Object names.

When the domain functional level has been set to Windows Server 2003, a new lastLogonTimestamp attribute is used to track the last logon time of a user or computer account. This attribute is replicated within the domain and can provide you with important information regarding the history of a user or computer.

# Object names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Object names**

Every object in Active Directory is an instance of a class defined in the schema. Each class has attributes that ensure:

- Unique identification of each object (instance of a class) in a directory data store
- Backward compatibility with security IDs used in Windows NT 4.0 and earlier
- Compatibility with LDAP standards for directory object names

For more information about schema, classes, and attributes, see Schema.

Each object in Active Directory can be referenced by several different names. Active Directory creates a relative distinguished name and a canonical name for each object based upon information that was provided when the object was created or modified. Each object can also be referenced by its distinguished name, which is derived from the relative distinguished name of the object and all of its parent container objects.

- The LDAP relative distinguished name uniquely identifies the object within its parent container. For example, the LDAP relative distinguished name of a computer named my computer is CN=mycomputer. Relative distinguished names must be unique in that users cannot have the same name within an organizational unit.
- The LDAP distinguished name is globally unique. For example, the distinguished name of a computer named mycomputer in the MyOrganizationalUnit organizational unit in the microsoft.com domain is CN=mycomputer, OU=MyOrganizationalUnit, DC=microsoft, DC=com.
- The canonical name is constructed the same way as the distinguished name, but it is represented using a different notation. The canonical name of the computer in the previous example would be Microsoft.com/MyOrganizationalUnit/mycomputer.

Security principal objects are Active Directory objects that are assigned security IDs (SIDs) and can be used to log on to the network and can be assigned access to domain resources. An administrator needs to provide names for security principal objects (user accounts, computer accounts, and groups) that are unique within a domain.

Consider what occurs when a new user account is added to your directory. You provide a name the user must use to log on to the network, the name of the domain that contains the user account, and other descriptive data, such as first name, last name, telephone number and so on (called attributes). All this information is recorded in the directory.

The names of security principal objects can contain all Unicode characters except the special LDAP characters defined in RFC 2253. This list of special characters includes: a leading space; a trailing space; and any of the following characters: # , + " \ < > ;

Security principal names must conform to the following guidelines:

| Type of account name | Maximum size | Special limitations |
|---|---|---|
| User account | Computers running Windows Server 2003 and Windows 2000 can use a user principal name (UPN) for a user account. Computers running Windows NT 4.0 and earlier are limited to 20 characters or 20 bytes depending upon the character set; individual characters may require more than one byte. | A user account cannot consist solely of periods (.) or spaces, or end in a period. Any leading periods or spaces are cropped. Use of the @ symbol is not supported with the logon format for Windows NT 4.0 and earlier, which is *DomainName**UserName*. Windows 2000 logon names are unique to the domain and Windows Server 2003 logon names are unique within the forest. |

| | | |
|---|---|---|
| Computer account | NetBIOS = 15 characters, or 15 bytes depending upon the character set; individual characters may require more than one byte.<br><br>DNS = 63 characters or 63 bytes depending upon the character set and 255 characters for a fully qualified domain name (FQDN) individual characters may require more than one byte. | A computer account cannot consist solely of numbers, periods (.), or spaces. Any leading periods or spaces are cropped. |
| Group account | 63 characters, or 63 bytes depending upon the character set; individual characters may require more than one byte. | A group account cannot consist solely of numbers, periods (.), or spaces. Any leading periods or spaces are cropped. |

**Note**

- If the administrator changes the default security settings, then it is possible to use computer names containing more than 15 characters. For more information, see Active Directory naming.

From the information provided by the person who creates the security principal object, Active Directory generates a security ID (SID), and a globally unique ID used to identify the security principal. Active Directory also creates an LDAP relative distinguished name, based on the security principal name. An LDAP distinguished name and a canonical name are derived from the relative distinguished name and the names of the domain and container contexts in which the security principal object is created.

If your organization has several domains, it is possible to use the same user name or computer name in different domains. The SID, globally unique ID, LDAP distinguished name, and canonical name generated by Active Directory will uniquely identify each user, computer, or group in the forest. If the security principal object is renamed or moved to a different domain, the SID, LDAP relative distinguished name, LDAP distinguished name, and canonical name will change, but the globally unique ID generated by Active Directory will not change.

Security principal objects, such as user accounts, may be renamed, moved, or contained within a nested domain hierarchy. To reduce the effect of renaming, moving, or assigning user account names within a nested domain hierarchy, Active Directory provides a method for simplifying user logon names. For information about user logon names, see Active Directory naming and Add user principal name suffixes, and User and computer accounts.

# Organizational units

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Organizational units**

A particularly useful type of directory object contained within domains is the organizational unit. Organizational units are Active Directory containers into which you can place users, groups, computers, and other organizational units. An organizational unit cannot contain objects from other domains.

An organizational unit is the smallest scope or unit to which you can assign Group Policy settings or delegate administrative authority. Using organizational units, you can create containers within a domain that represent the hierarchical, logical structures within your organization. You can then manage the configuration and use of accounts and resources based on your organizational model. For more information about Group Policy settings, see Group Policy (pre-GPMC).




As shown in the figure, organizational units can contain other organizational units. A hierarchy of containers can be extended as necessary to model your organization's hierarchy within a domain. Using organizational units will help you minimize the number of domains required for your network.

You can use organizational units to create an administrative model that can be scaled to any size. A user can have administrative authority for all organizational units in a domain or for a single organizational unit. An administrator of an organizational unit does not need to have administrative authority for any other organizational units in the domain. For more information about delegating administrative authority, see Delegating administration.

# Active Directory server roles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory server roles**
Computers that function as servers within a domain can have one of two roles: member server or domain controller. A server that is not in a domain is a stand-alone server.

**Member servers**
A member server is a computer that:

- Runs an operating system in the Windows 2000 Server family or the Windows Server 2003 family.
- Belongs to a domain.
- Is not a domain controller.

A member server does not process account logons, participate in Active Directory replication, or store domain security policy information.

Member servers typically function as the following types of servers: file servers, application servers, database servers, Web servers, certificate servers, firewalls, and remote access servers. For more information about server roles, see Server roles.

The following security-related features are common to all member servers:

- Member servers adhere to Group Policy settings that are defined for the site, domain, or organizational unit.
- Access control for resources that are available on a member server.
- Member server users have assigned user rights.
- Member servers contain a local security account database, the Security Accounts Manager (SAM).

**Domain controllers**
A domain controller is a computer that:

- Runs an operating system in the Windows 2000 Server family or the Windows Server 2003 family.
- Uses Active Directory to store a read-write copy of the domain database, participate in multimaster replication, and authenticate users.

Domain controllers store directory data and manage communication between users and domains, including user logon processes, authentication, and directory searches. Domain controllers synchronize directory data using multimaster replication, ensuring consistency of information over time. For more information about multimaster replication, see Replication overview.

Active Directory supports multimaster replication of directory data between all domain controllers in a domain; however, multimaster replication is not appropriate for some directory data replication. In this case, a domain controller, called the operations master, will process data. In an Active Directory forest, there are at least five different operations master roles that are assigned to one or more domain controllers. For more information about operations masters, see Operations master roles.

As the needs of your computing environment change, you might want to change the role of a server. Using the Active Directory Installation Wizard, you can install Active Directory on a member server to make it a domain controller, or you can remove Active Directory from a domain controller to make it a member server. For more information about domain controllers, see Domain controllers.

**Note**

- You cannot install Active Directory on a computer running Windows Server 2003, Web Edition, but you can join the

computer to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

# Active Directory clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory clients**
With the Active Directory client, many of the Active Directory features available on Windows 2000 Professional or Windows XP Professional are available to computers running Windows 95, Windows 98, and Windows NT 4.0.

- **Site awareness**

  You can log on to the domain controller that is closest to the client in the network. For more information, see Locating a domain controller.

- **Active Directory Service Interfaces (ADSI)**

  You can use scripting to Active Directory. ADSI also provides a common programming API to Active Directory programmers.

- **Distributed File System (DFS)**

  You can access DFS shared resources on servers running Windows 2000 and Windows Server 2003.

- **NTLM version 2 authentication**

  You can use the improved authentication features in NTLM version 2.

  For more information about enabling NTML version 2, see article Q239869, "How to Enable NTLM 2 Authentication," in the Microsoft Knowledge Base.

- **Active Directory Windows Address Book (WAB) property pages**

  You can change properties, such as phone number and address, on user object pages.

- **Active Directory search capability**

  From the **Start** button, you can locate printers and people in a Windows 2000 Server or Windows Server 2003 domain.

  For information about publishing printers in Active Directory, see article Q234619, "Publishing a Printer in Windows 2000 Active Directory," in the Microsoft Knowledge Base.

Windows 2000 Professional and Windows XP Professional provide functionality that the Active Directory client on Windows 95, Windows 98, and Windows NT 4.0 does not, including Kerberos V5 support; Group Policy or IntelliMirror support; and service principal name (SPN), or mutual authentication. You can take advantage of these additional features by upgrading to Windows 2000 Professional or Windows XP Professional.

To install the Active Directory client, see the Active Directory client page at the Microsoft Web site.

# Understanding Domains and Forests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Domains and Forests**

This section covers:

- Domain controllers
- Renaming domain controllers
- Connecting to domain controllers running Windows 2000
- Domains
- Renaming domains
- Domain and forest functionality
- Raising domain and forest functional levels
- Application directory partitions
- Application directory partitions and domain controller demotion
- Managing application directory partitions
- Operations master roles
- Transferring operations master roles
- Responding to operations master failures

# Domain controllers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain controllers**
When you create the first domain controller in your organization, you are also creating the first domain, the first forest, the first site, and installing Active Directory. Domain controllers running Windows Server 2003 store directory data and manage user and domain interactions, including user logon processes, authentication, and directory searches. Domain controllers are created by using the Active Directory Installation Wizard. For more information, see Using the Active Directory Installation Wizard.

**Note**

- You cannot install Active Directory on a computer running Windows Server 2003, Web Edition, but you can join the computer to an Active Directory domain as a member server. For more information about Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition.

When using domain controllers in your organization, you will want to think about how many domain controllers you'll need, the physical security of those domain controllers, a plan for backing up the domain data, and upgrading domain controllers.

**Determining the number of domain controllers you need**
A small organization using a single local area network (LAN) might need only one domain with two domain controllers for high availability and fault tolerance. A larger organization with many network locations will need one or more domain controllers in each site to provide high availability and fault tolerance.

If your network is divided into sites, it is often good practice to put at least one domain controller in each site to enhance network performance. When users log on to the network, a domain controller must be contacted as part of the logon process. If clients must connect to a domain controller located in a different site, the logon process can take a long time. For more information, see Replication between sites.

By creating a domain controller in each site, user logons are processed more efficiently within the site. For information about how to create additional domain controllers, see Create an additional domain controller.

To optimize network traffic, you can also configure domain controllers to receive directory replication updates only during off-peak hours. For information about how to schedule site replication, see Configure site link replication availability.

The best network performance is available when the domain controller at a site is also a global catalog. This way, the server can fulfill queries about objects in the entire forest. However, enabling many domain controllers as global catalogs can increase the replication traffic on your network. For more information about the global catalog, see The role of the global catalog. For more information about adding global catalogs to sites, see Global catalogs and sites.

In domains with more than one domain controller, do not enable the domain controller holding the infrastructure master role as a global catalog. For more information, see Operations master roles.

**Physical security**
Physical access to a domain controller can provide a malicious user unauthorized access to encrypted passwords. For this reason, it is recommended that all domain controllers in your organization be locked in a secured room with limited public access. You can use additional security measures such as Syskey for extra protection on domain controllers. For more information about Syskey, see The system key utility.

**Backing up domain controllers**
You can back up domain directory partition data and data from other directory partitions by using Backup, which is included with the Windows Server 2003 family, from any domain controller in a domain. By using the backup tool on a domain controller, you can:

- Back up Active Directory while the domain controller is online.
- Back up Active Directory using batch file commands.
- Back up Active Directory to removable media, an available network drive, or a file.

- Back up other system and data files.

When you use the backup tool on a domain controller it will automatically back up all of the system components and all of the distributed services upon which Active Directory is dependent. This dependent data, which includes Active Directory, is known collectively as the System State data.

On a domain controller running Windows Server 2003, the System State data consists of the system startup files; the system registry; the class registration database of COM+ (an extension to the Component Object Model); the SYSVOL directory; Certificate Services database (if installed); Domain Name System (if installed); Cluster service (if installed); and Active Directory. It is recommended that you regularly back up System State data.

For general information about the System State, see System State data. For more information about how to back up the System State, see Back up System State data. For more information about how to restore a System State backup, see Restore System State data.

You can install Active Directory on a server running Windows Server 2003 by using a restored backup taken from a domain controller running Windows Server 2003. For more information, see Creating an additional domain controller.

**Upgrading domain controllers**

On domain controllers running Windows NT 4.0, you will first need to upgrade the primary domain controller (PDC) to successfully upgrade the domain. Once the PDC has been upgraded, you can upgrade the backup domain controllers (BDCs). For more information, see Upgrading from a Windows NT domain.

If you currently have a Windows 2000 forest that does not have any domain controllers running Windows Server 2003, you will need to prepare the forest and the target domain before you can upgrade domain controllers running Windows 2000. For more information, see Upgrading from a Windows 2000 domain.

# Renaming domain controllers

Updated: April 23, 2013

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Renaming domain controllers**

The ability to rename domain controllers running Windows Server 2003 provides you with the flexibility to make changes in a Windows Server 2003 domain whenever the need arises. Rename a domain controller to:

- Restructure your network for organizational and business needs.
- Make management and administrative control easier.

When you rename a domain controller, you must ensure that there will be no interruption in the ability of clients to locate or authenticate to the renamed domain controller, except when the domain controller is restarted. The recommended practice for renaming a domain controller without interruption to clients is to use the Netdom tool. To rename a domain controller using the Netdom tool, the domain functional level must be set to Windows Server 2003. For more information about renaming a domain controller, see Rename a domain controller.

The **System Properties** dialog box can also be used to rename a domain controller, and it does not require the functional level to be raised to Windows Server 2003. Using this dialog box may result in a service interruption to clients. For more information about functional levels, see Domain and forest functionality.

The new name of the domain controller is automatically updated to Domain Name System (DNS) and Active Directory. Once the new name propagates to DNS and Active Directory, clients are then capable of locating and authenticating to the renamed domain controller. DNS and Active Directory replication latency may delay client ability to locate or authenticate to the renamed domain controller. The length of time this takes depends on specifics of your network and the replication topology of your particular organization.

During replication latency, clients may not be able to access the newly renamed domain controller. This might be acceptable for clients that try to locate and authenticate to a particular domain controller since other domain controllers should be available to process the authentication request.

| **Note** |
| --- |
| The corresponding nTFRSMember or msDFSR-Member object is not renamed automatically, but the reference attributes are correctly set so SYSVOL replication is not impacted. The only potential problem with not renaming these objects is that if another domain controller is created at a later date with the same NetBIOS name of the old domain controller, then a conflict can occur as described in KB article 316826. After the rename is complete, you can optionally rename the nTFRSMember or msDFSR-Member object as part of cleanup. |

# Connecting to domain controllers running Windows 2000

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connecting to domain controllers running Windows 2000**
When you need to connect to a domain controller running Windows 2000 from a domain controller running Windows Server 2003, a number of Active Directory administrative tools are available, such as Active Directory Domains and Trusts.

The following Windows Server 2003 Active Directory administrative tools will sign and encrypt all LDAP traffic by default:

- Active Directory Users and Computers
- Active Directory Sites and Services
- Active Directory Domains and Trusts
- Active Directory Schema
- ADSI Edit
- Dsrm.exe
- Dsmove.exe
- Dsadd.exe
- Dsmod.exe
- Dsget.exe
- Dsquery.exe

Signing LDAP traffic guarantees that the packaged data comes from a known source and that it has not been tampered with. The Active Directory administrative tools in Windows 2000 do not sign and encrypt all LDAP traffic by default. In order to maintain a secure network, it is strongly recommended that you upgrade all domain controllers running Windows 2000 to Service Pack 3 or later.

You can use the corresponding Active Directory administrative tools in Windows 2000 to manage Windows 2000 domain controllers that do not have the Windows 2000 Server Service Pack 3 or later installed. However, traffic is not signed and encrypted by LDAP on domain controllers running Windows 2000.

Although it is not recommended, you can disable encrypted and signed LDAP traffic used with Active Directory administrative tools on a server running Windows Server 2003 or on a computer running Windows XP Professional that has the Windows Server 2003 Administration Tools Pack installed. For more information, see Disable signed or encrypted LDAP traffic.

# Domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domains**

Domains are units of replication. All of the domain controllers in a particular domain can receive changes and replicate those changes to all other domain controllers in the domain. Each domain in Active Directory is identified by a Domain Name System (DNS) domain name and requires one or more domain controllers. If your network requires more than one domain, you can easily create multiple domains.

One or more domains that share a common schema and global catalog are referred to as a forest. The first domain in a forest is referred to as the forest root domain. For more information about forests, see Creating a new forest. If multiple domains in the forest have contiguous DNS domain names, then the structure is referred to as a domain tree. For more information, see Active Directory naming and Creating a new domain tree.

A single domain can span multiple physical locations or sites and can contain millions of objects. Site structure and domain structure are separate and flexible. A single domain can span multiple geographical sites, and a single site can include users and computers belonging to multiple domains. For more information, see Sites overview.

A domain provides several benefits:

- **Organizing objects.**

  You do not need to create separate domains merely to reflect your company's organization of divisions and departments. Within a domain, you can use organizational units for this purpose. Using organizational units helps you manage the accounts and resources in the domain. You can then assign Group Policy settings and place users, groups, and computers into the organizational units. Using a single domain greatly simplifies administrative overhead. For more information, see Organizational units.

- **Publishing resources and information about domain objects.**

  A domain stores only the information about objects located in that domain, so by creating multiple domains, you are partitioning or segmenting the directory to better serve a disparate user base. When using multiple domains, you can scale the Active Directory directory service to accommodate your administrative and directory publishing requirements. For more information, Publishing resources.

- **Applying a Group Policy object to the domain consolidates resource and security management.**

  A domain defines a scope or unit of policy. A Group Policy object (GPO) establishes how domain resources can be accessed, configured, and used. These policies are applied only within the domain and not across domains. For more information about applying GPOs, see Group Policy (pre-GPMC).

- **Delegating authority eliminates the need for a number of administrators with broad administrative authority.**

  Using delegated authority in conjunction with Group Policy objects and group memberships enables you to assign an administrator rights and permissions to manage objects in an entire domain or in one or more organizational units within the domain. For more information about delegating administrative control, see Delegating administration.

- **Security policies and settings (such as user rights and password policies) do not cross from one domain to another.**

  Each domain has its own security policies and trust relationships with other domains. However, the forest is the final security boundary. For more information, see Creating a new forest.

- **Each domain stores only the information about the objects located in that domain.**

  By partitioning the directory this way, Active Directory can scale to very large numbers of objects.

**Creating a domain**
You create a domain by creating the first domain controller for a domain. To do this, install Active Directory on a member server running Windows Server 2003 by using the Active Directory Installation Wizard. The wizard uses the information that you provide to create the domain controller and create the domain within the existing domain structure of your organization. Depending on the existing domain structure, the new domain could be the first domain in a new forest, the first domain in a new domain tree, or a child domain of an existing domain tree. For more information, see Creating a new forest, Creating a new domain tree, and Creating a new child domain.

A domain controller provides the Active Directory directory service to network users and computers, stores directory data, and manages user and domain interactions, including user logon processes, authentication, and directory searches. Every domain must contain at least one domain controller. For more information, see Domain controllers.

After you create the first domain controller for a domain, you can create additional domain controllers in an existing domain for fault tolerance and high availability of the directory. For more information, see Creating an additional domain controller.

**Planning for multiple domains**
Some reasons to create more than one domain are:

- Different password requirements between departments or divisions
- Massive numbers of objects
- Decentralized network administration
- More control of replication

Although using a single domain for an entire network has several advantages, to meet additional scalability, security, or replication requirements you may consider creating one or more domains for your organization. Understanding how directory data is replicated between domain controllers will help you plan the number of domains needed by your organization. For more information about replication, see How replication works.

**Removing a domain**
In order to remove a domain, you must first remove Active Directory from all of the domain controllers associated with that domain. Once Active Directory has been removed from the last domain controller the domain will be removed from the forest and all of the information in that domain will be deleted. A domain can only be removed from the forest if it has no child domains. If this is the last domain in the forest, removing this domain will also delete the forest.

For more information about how to remove a domain, see Remove a domain.

### Caution

- Removing a domain will result in the permanent loss of amy data contained in that domain. This includes all user, group, and computer accounts.

Before removing Active Directory from a domain controller, you should first remove any application directory partitions from that domain controller. For more information, see Application directory partitions and Create or delete an application directory partition.

**Trust relationships between domains**
Trust relationships are automatically created between adjacent domains (parent and child domains) when a domain is created in Active Directory. In a forest, a trust relationship is automatically created between the forest root domain and any tree root domains or child domains that are subordinate to the forest root domain. Because these trust relationships are transitive, users and computers can be authenticated between any domains in the forest. For more information about trust relationships, see Trust transitivity.

When upgrading a Windows NT domain to a Windows Server 2003 domain, the existing one-way trust relationship between that domain and any other domains remains intact. This includes all trusts with other Windows NT domains. If you are creating a new Windows Server 2003 domain and want trust relationships with any Windows NT domains, you must create external trusts with those domains. For more information about external trusts, see When to create an external trust.

# Renaming domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Renaming domains**

The ability to rename a domain provides you with the flexibility to make important changes to your forest structure and namespace as the needs of your organization change. Renaming domains can accommodate acquisitions, mergers, name changes, or reorganizations. Domain rename allows you to:

- Change the Domain Name System (DNS) and NetBIOS names of any domain in the forest (including the forest root domain).
- Restructure the position of any domain in the forest (except the forest root domain).

You can only rename domains in a forest where all of the domain controllers are running Windows Server 2003 and the forest functional level has been raised to Windows Server 2003. For more information, see Domain and forest functionality.

**Forest restructuring**

Using domain rename, you can also restructure the hierarchy of domains in your forest so that a domain residing in one domain tree can be moved to another domain tree. Restructuring a forest allows you to move a domain anywhere within the forest in which it resides (except the forest root domain). This includes the ability to move a domain so that it becomes the root of its own domain tree.

You can use the domain rename utility (Rendom.exe) to rename or restructure a domain. A domain rename will affect every domain controller in your forest and is a multistep process that requires a detailed understanding of the operation. For more information about this process and to download the Rendom.exe tool, see the Windows Server 2003 Active Directory Domain Rename Tools page at the Microsoft Web site.

# Domain and forest functionality

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain and forest functionality**

Domain and forest functionality, introduced in Windows Server 2003 Active Directory, provides a way to enable domain- or forest-wide Active Directory features within your network environment. Different levels of domain functionality and forest functionality are available depending on your environment.

If all domain controllers in your domain or forest are running Windows Server 2003 and the functional level is set to Windows Server 2003, all domain- and forest-wide features are available. When Windows NT 4.0 or Windows 2000 domain controllers are included in your domain or forest with domain controllers running Windows Server 2003, Active Directory features are limited. For more information about how to enable domain- or forest-wide features, see Raising domain and forest functional levels.

The concept of enabling additional functionality in Active Directory exists in Windows 2000 with mixed and native modes. Mixed-mode domains can contain Windows NT 4.0 backup domain controllers and cannot use Universal security groups, group nesting, and security ID (SID) history capabilities. When the domain is set to native mode, Universal security groups, group nesting, and SID history capabilities are available. Domain controllers running Windows 2000 Server are not aware of domain and forest functionality.

**Domain functionality**

Domain functionality enables features that will affect the entire domain and that domain only. Four *domain* functional levels are available: Windows 2000 mixed (default), Windows 2000 native, Windows Server 2003 interim, and Windows Server 2003. By default, domains operate at the Windows 2000 mixed functional level.

The following table lists the domain functional levels and their corresponding supported domain controllers.

| Domain functional level | Domain controllers supported |
|---|---|
| **Windows 2000 mixed** (default) | Windows NT 4.0<br>Windows 2000<br>Windows Server 2003 family |
| **Windows 2000 native** | Windows 2000<br>Windows Server 2003 family |
| **Windows Server 2003 interim** | Windows NT 4.0<br>Windows Server 2003 family |
| **Windows Server 2003** | Windows Server 2003 family |

Once the domain functional level has been raised, domain controllers running earlier operating systems cannot be introduced into the domain. For example, if you raise the domain functional level to Windows Server 2003, domain controllers running Windows 2000 Server cannot be added to that domain.

The following table describes the domain-wide features that are enabled for three of the domain functional levels. For information about the Windows Server 2003 interim functional level, see Upgrading from a Windows NT domain.

| Domain feature | Windows 2000 mixed | Windows 2000 native | Windows Server 2003 |
|---|---|---|---|
| **Domain controller rename tool**<br>For more information, see Renaming domain controllers. | Disabled | Disabled | Enabled |

| | | | |
|---|---|---|---|
| **Different location option for user and computer accounts**<br><br>For more information about how to redirect the default location for user and computer accounts, see Redirect the Users and Computers Containers. | Disabled | Disabled | Enabled |
| **Update logon timestamp**<br><br>For more information about the lastLogonTimestamp attribute, see User and computer accounts. | Disabled | Disabled | Enabled |
| **User password on InetOrgPerson object**<br><br>For more information about InetOrgPerson objects, see User and computer accounts. | Disabled | Disabled | Enabled |
| **Universal Groups**<br><br>For more information, see Group types and Group scope. | Enabled for distribution groups.<br><br>Disabled for security groups. | Enabled<br><br>Allows both security and distribution groups. | Enabled<br><br>Allows both security and distribution groups. |
| **Group Nesting**<br><br>For more information, see Nesting groups. | Enabled for distribution groups.<br><br>Disabled for security groups, except for domain local security groups that can have global groups as members. | Enabled<br><br>Allows full group nesting. | Enabled<br><br>Allows full group nesting. |
| **Converting Groups**<br><br>For more information, see Converting groups. | Disabled<br><br>No group conversions allowed. | Enabled<br><br>Allows conversion between security groups and distribution groups. | Enabled<br><br>Allows conversion between security groups and distribution groups. |
| **SID history** | Disabled | Enabled<br><br>Allows migration of security principals from one domain to another. | Enabled<br><br>Allows migration of security principals from one domain to another. |

**Forest functionality**

Forest functionality enables features across all the domains within your forest. Three forest functional levels are available: Windows 2000 (default), Windows Server 2003 interim, and Windows Server 2003 . By default, forests operate at the Windows 2000 functional level. You can raise the forest functional level to Windows Server 2003 .

The following table lists the forest functional levels and their corresponding supported domain controllers:

| Forest functional level | Domain controllers supported |
|---|---|
| **Windows 2000** (default) | Windows NT 4.0<br><br>Windows 2000<br><br>Windows Server 2003 family |
| **Windows Server 2003 interim** | Windows NT 4.0<br><br>Windows Server 2003 family |

| | |
|---|---|
| **Windows Server 2003** | Windows Server 2003 family |

Once the forest functional level has been raised, domain controllers running earlier operating systems cannot be introduced into the forest. For example, if you raise the forest functional level to Windows Server 2003, domain controllers running Windows 2000 Server cannot be added to the forest.

If you are upgrading your first Windows NT 4.0 domain so that it becomes the first domain in a new Windows Server 2003 forest, you can set the domain functional level to Windows Server 2003 interim. For more information, see Upgrading from a Windows NT domain.

The following table describes the forest-wide features that are enabled for the Windows 2000 and Windows Server 2003 forest functional levels.

| Forest feature | Windows 2000 | Windows Server 2003 |
|---|---|---|
| **Global catalog replication improvements**<br><br>For more information, see Global catalog replication. | Enabled if both replication partners are running Windows Server 2003.<br><br>Otherwise, disabled. | Enabled |
| **Defunct schema objects**<br><br>For more information, see Deactivating a class or attribute. | Disabled | Enabled |
| **Forest trusts**<br><br>For more information, see Forest trusts. | Disabled | Enabled |
| **Linked value replication**<br><br>For more information, see How replication works. | Disabled | Enabled |
| **Domain rename**<br><br>For more information, see Renaming domains. | Disabled | Enabled |
| **Improved Active Directory replication algorithms**<br><br>For more information, see Replication overview. | Disabled | Enabled |
| **Dynamic auxiliary classes**.<br><br>For more information, see New features for Active Directory. | Disabled | Enabled |
| **InetOrgPerson objectClass change**<br><br>For more information about InetOrgPerson objects, see User and computer accounts. | Disabled | Enabled |

# Raising domain and forest functional levels

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Raising domain and forest functional levels**

When Active Directory is installed on a server running Windows Server 2003, a set of basic Active Directory features is enabled by default. For more information about the new default features, see New features for Active Directory.

In addition to the basic Active Directory features on individual domain controllers, there are new domain- and forest-wide Active Directory features available when *all* domain controllers in a domain or forest are running Windows Server 2003.

To enable the new domain-wide features, all domain controllers in the domain must be running Windows Server 2003, and the domain functional level must be raised to Windows Server 2003 . For information about raising the domain functional level, see Raise the domain functional level.

To enable new forest-wide features, all domain controllers in the forest must be running Windows Server 2003, and the forest functional level must be raised to Windows Server 2003 . Before raising the *forest* functional level to Windows Server 2003 , verify that all *domains* in the forest are set to the *domain* functional level of Windows 2000 native or Windows Server 2003 . Note that domains that are set to the domain functional level of Windows 2000 native will automatically be raised to Windows Server 2003 at the same time the *forest* functional level is raised to Windows Server 2003 . For information about raising the forest functional level, see Raise the forest functional level.

# Application directory partitions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application directory partitions**
An application directory partition is a directory partition that is replicated only to specific domain controllers. A domain controller that participates in the replication of a particular application directory partition hosts a replica of that partition. Only domain controllers running Windows Server 2003 can host a replica of an application directory partition.

Applications and services can use application directory partitions to store application-specific data. Application directory partitions can contain any type of object, except security principals. TAPI is an example of a service that stores its application-specific data in an application directory partition.

Application directory partitions are usually created by the applications that will use them to store and replicate data. For testing and troubleshooting purposes, members of the Enterprise Admins group can manually create or manage application directory partitions using the Ntdsutil command-line tool.

One of the benefits of an application directory partition is that, for redundancy, availability, or fault tolerance, the data in it can be replicated to different domain controllers in a forest. The data can be replicated to a specific domain controller or any set of domain controllers anywhere in the forest. This differs from a domain directory partition in which data is replicated to all domain controllers in that domain. Storing application data in an application directory partition instead of in a domain directory partition may reduce replication traffic because the application data is only replicated to specific domain controllers. Some applications may use application directory partitions to replicate data only to servers where the data will be locally useful.

Another benefit of application directory partitions is that applications or services that use Lightweight Directory Access Protocol (LDAP) can continue using it to access and store their application data in Active Directory. For more information, see Managing application directory partitions.

**Application directory partition naming**
An application directory partition is part of the overall forest namespace just like a domain directory partition. It follows the same Domain Name System (DNS) and distinguished names naming conventions as a domain directory partition. An application directory partition can appear anywhere in the forest namespace that a domain directory partition can appear.

There are three possible application directory partition placements within your forest namespace:

- A child of a domain directory partition.
- A child of an application directory partition.
- A new tree in the forest.

If you created an application directory partition called *example1* as a child of the microsoft.com domain, the DNS name of the application directory partition would be example1.microsoft.com. The distinguished name of the application directory partition would be dc=example1, dc=microsoft, dc=com.

If you then created an application directory partition called *example2* as a child of example1.microsoft.com, the DNS name of the application directory partition would be example2.example1.microsoft.com and the distinguished name would be dc=example2, dc=example1, dc=microsoft, dc=com.

If the domain microsoft.com was the root of the only domain tree in your forest, and you created an application directory partition with the DNS name of *example1* and the distinguished name of dc=example1, this application directory partition is not in the same tree as the microsoft.com domain. This application directory partition would be the root of a new tree in the forest.

Domain directory partitions cannot be children of an application directory partition. For example, if you created an application directory partition with the DNS name of example1.microsoft.com, you could not create a domain with the DNS name of domain.example1.microsoft.com.

**Application directory partition replication**
The Knowledge Consistency Checker (KCC) automatically generates and maintains the replication topology for all application

directory partitions in the enterprise. When an application directory partition has replicas in more than one site, those replicas follow the same intersite replication schedule as the domain directory partition.

### Notes

- Objects stored in an application directory partition are never replicated to the global catalog as read-only replicas. Any domain controller running Windows Server 2003 can hold a replica, including global catalogs.

- Additionally, if an application requests data through the global catalog port, that query will not return any objects from an application directory partition, even if the computer hosting the application directory partition is also hosting the global catalog. This is done so that LDAP queries to different global catalogs will not return inconsistent results because the application directory partition is replicated to only one of the global catalogs.

**Security descriptor reference domain**

Every container and object on the network has a set of access control information attached to it. Known as a security descriptor, this information controls the type of access allowed by users, groups, and computers. If the object or container is not assigned a security descriptor by the application or service that created it, then it is assigned the default security descriptor for that object class as defined in the schema. This default security descriptor is ambiguous in that it may assign members of the Domain Admins group read permissions to the object, but it does not specify to what domain the domain administrators belong. When this object is created in a domain naming partition, that domain naming partition is used to specify which Domain Admins group actually is assigned read permission. For example, if the object is created in mydomain.microsoft.com then members of the mydomain Domain Admins group would be assigned read permission.

When an object is created in an application directory partition, the definition of the default security descriptor is difficult because an application directory partition can have replicas on different domain controllers belonging to different domains. Because of this potential ambiguity, a default security descriptor reference domain is assigned when the application directory partition is created.

The default security descriptor reference domain defines what domain name to use when an object in the application directory partition needs to define a domain value for the default security descriptor. The default security descriptor reference domain is assigned at the time of creation.

If the application directory partition is a child of a domain directory partition, by default, the parent domain directory partition becomes the security descriptor reference domain. If the application directory partition is a child object of another application directory partition, the security descriptor reference domain of the parent application directory partition becomes the reference domain of the new, child, application directory partition. If the new application directory partition is created as the root of a new tree, then the forest root domain is used as the default security descriptor reference domain.

You can manually specify a security reference domain using Ntdsutil. However, if you plan to change the default security descriptor reference domain of a particular application directory partition, you should do so before creating the first instance of that partition. To do this, you must prepare the cross-reference object and change the default security reference domain before completing the application directory partition creation process. For information about precreating the cross-reference object, see Prepare a cross-reference object. For information about changing the default security reference domain, see Set an application directory partition reference domain.

# Application directory partitions and domain controller demotion

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application directory partitions and domain controller demotion**
If a domain controller holds a replica of an application directory partition, then you must remove the domain controller from the replica set of the application directory partition or delete the application directory partition before you can demote the domain controller.

If a domain controller holds the *last* replica of a particular application directory partition, then you must delete the application directory partition before you can demote the domain controller.

The Active Directory Installation Wizard will not remove a replica or delete an application directory partition programmatically. You must decide when it is safe to delete the last replica of a particular partition.

Before deleting the last replica of an application directory partition, identify the applications that use the application directory partition, determine if it is safe to delete the last replica, identify the partition deletion tool provided by the application, and then remove the application directory partition by using the tool provided or by using the Ntdsutil command-line tool.

**Identify the applications that use the application directory partition**
To determine what application directory partitions are hosted on a computer, refer to the list on the first page of the Active Directory Installation Wizard. If the list does not provide enough information to identify the programs using a particular application directory partition, you may be able to identify them in one of the following ways:

- Speak to a member of the Enterprise Admins group.
- Consult the network change control records for your organization.
- Use LDP or ADSI Edit to view the data contained in the partition. For more information about these tools, see the Active Directory Programmer's Guide at the Microsoft Web site.

**Determine if it is safe to delete the last replica**
Removing the last replica of an application directory partition will result in the permanent loss of any data contained in the partition. If you have identified the applications using the application directory partition, consult the documentation provided with those applications to determine if there is any reason to keep the data. If the applications that use the application directory partition are out of service, it is probably safe to remove the partition.

If it is not safe to delete the last replica, or if you cannot determine whether or not it is safe, and you must demote the domain controller holding the last replica of a particular application directory partition, follow these steps: Add a replica of the partition on another domain controller, force the replication of the contents of the application directory partition to the domain controller holding the new replica, and then remove the replica of the partition on the domain controller to be demoted. For more information, see Add or remove an application directory partition replica.

**Identify the partition deletion tool provided by the application**
Most applications that create application directory partitions provide a utility to remove the partitions. When possible, always delete an application directory partition using the utility provided. For example, to delete a TAPI partition, use the Tapicfg.exe command-line tool. For more information about TAPI and removing TAPI application directory partitions, see Telephony.

**Remove the application directory partition using the tool provided or use Ntdsutil**
Refer to the application's documentation for information about removing application directory partitions that were created and used by that application.

### Caution

- If possible, use the application's tool for managing its application directory partitions. The application may keep other data in addition to Active Directory managed data for the application directory partitions. By using Ntdsutil, the two sets of data could cause a conflict.

If you cannot identify the application that created the application directory partition, or if your application does not provide a means to delete application directory partitions that it created, you can use the Ntdsutil command-line tool. To do this, see Create or delete an application directory partition.

For information about demoting a domain controller, see Demote a domain controller. For general information about application directory partitions, see Application directory partitions.

# Managing application directory partitions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing application directory partitions**
You can use the following tools to create, delete, or manage application directory partitions.

- application-specific tools from the application vendor
- Ntdsutil command-line tool
- LDP
- Active Directory Service Interfaces (ADSI)

For information about creating and managing application directory partitions with ADSI, see Active Directory Service Interfaces (ADSI) at the Microsoft Web site. For information about LDP, see Administration tools for the Active Directory schema.

For information about the Ntdsutil command-line tool, see Ntdsutil.

The following provides information about using Ntdsutil to create and manage application directory partitions.

**Creating an application directory partition**
When you create an application directory partition, you are creating the first instance of this partition. You can create an application directory partition by using the *create nc* option in the **domain management** menu of Ntdsutil. When creating an application directory partition using LDP or ADSI, provide a description in the description attribute of the domain DNS object that indicates the specific application that will use the partition. For example, if the application directory partition will be used to store data for a Microsoft accounting program, the description could be Microsoft accounting application. Ntdsutil does not facilitate the creation of a description. For more information, see Create or delete an application directory partition.

**Deleting an application directory partition**
When you delete an application directory partition, you are removing all replicas of that partition from your forest. You can delete an application directory partition by using the *delete nc* command in the **domain management** menu of Ntdsutil. The deletion process will need to replicate to all domain controllers that contain a replica of the application directory partition before the deletion process is complete. Any data that is contained in the application directory partition will be lost.

The Active Directory Promotion Wizard (Dcpromo) cannot demote a domain controller if that domain controller holds a copy of an application directory partition. For more information, see Create or delete an application directory partition.

**Adding and removing a replica of an application directory partition**
An application directory partition replica is an instance of an partition on another domain controller. The information in the application directory partition is replicated between the domain controllers. Application directory partition replicas are created for either redundancy or data access purposes. You can add a replica of an application directory partition by using the *add nc replica* command in the **domain management** menu of Ntdsutil. You can remove an application directory partition replica by using the *delete nc replica* command in the **domain management** menu of Ntdsutil. For more information, see Add or remove an application directory partition replica.

**Setting application directory partition reference domain**
The security descriptor reference domain defines a domain name for the default security descriptor for objects in the application directory partition. By default, the security descriptor reference domain is the parent domain of the application directory partition. If the application directory partition is a child of another application directory partition, the default security descriptor reference domain is the security descriptor reference domain of the parent application directory partition. If the application directory partition has no parent, the forest root domain becomes the default security descriptor reference domain. You can use Ntdsutil to change the default security descriptor reference domain. For more information, see Set an application directory partition reference domain.

**Setting replication notification delays**
Changes made to a particular directory partition on a particular domain controller are replicated to the other domain controllers that contain that directory partition. The domain controller on which the change was made notifies its replication partners that it has a change. You can configure how long the domain controller will wait to send the change notification to its

first replication partner. You can also configure how long it waits to send the subsequent change notification to its remaining replication partners. These delays can be set for any directory partition (including domain directory partitions) on a particular domain controller. For more information, see Set a notification delay.

**Displaying application directory partition information**

Any domain controller that holds a replica of a particular directory partition (including application directory partitions) is said to be a member of the replica set for that directory partition. You can use Ntdsutil to list the domain controllers that are members of a particular replica set for an application directory partition. An addition of a domain controller to the replica set attribute on the cross-reference object does not create the replica, but it will display when the **list nc replica** command is used in Ntdsutil. The creation of the instance must replicate before the creation of the replica is complete. For more information, see Display application directory partition information.

**Delegating the creation of application directory partitions**

There are two things that happen when creating an application directory partition:

- Creation of the cross-reference object.
- Creation of the application directory partition root node.

Normally only members of the Enterprise Admins group can create an application directory partition. However, it is possible for a member of the Enterprise Admins group to prepare a cross-reference object for the application directory partition and to delegate the rest of the process to someone with more limited permissions.

The cross-reference object for an application directory partition holds several valuable pieces of information, including the domain controllers that are to have a replica of this partition and the security descriptor reference domain. The partition root node is the Active Directory object at the root of the partition

The Enterprise Admin can create the cross-reference object then delegate to a person or group with less permissions the right to create the application directory partition root node. Both creation of the cross-reference object and the application directory partition root node can be accomplished using Ntdsutil.

After using Ntdsutil to create the cross-reference object, the enterprise administrator must modify the cross-reference object's access control list to allow the delegated administrator to modify this cross-reference. This will allow the delegated administrator to create the application directory partition and modify the list of domain controllers that holds replicas of this application directory partition. The delegated administrator must use the names of the application directory partition and the domain controller name that were specified during the precreation process. For more information, see Prepare a cross-reference object.

For more information about application directory partitions, see Application directory partitions.

# Operations master roles

Updated: October 24, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

**Operations master roles**
Active Directory supports multimaster replication of the directory data store between all domain controllers (DC) in the domain, so all domain controllers in a domain are essentially peers. However, some changes are impractical to perform in using multimaster replication, so, for each of these types of changes, one domain controller, called the *operations master*, accepts requests for such changes.

In every forest, there are at least five operations master roles that are assigned to one or more domain controllers. Forest-wide operations master roles must appear only once in every forest. Domain-wide operations master roles must appear once in every domain in the forest.

### Note

- The operations master roles are sometimes called flexible single master operations (FSMO) roles.

**Forest-wide operations master roles**
Every forest must have the following roles:

- Schema master
- Domain naming master

These roles must be unique in the forest. This means that throughout the entire forest there can be only one schema master and one domain naming master.

**Schema master**
The schema master domain controller controls all updates and modifications to the schema. To update the schema of a forest, you must have access to the schema master. There can be only one schema master in the entire forest.

**Domain naming master**
The domain controller holding the domain naming master role controls the addition or removal of domains in the forest. There can be only one domain naming master in the entire forest.

### Note

- Any domain controller running Windows Server 2003 can hold the role of the domain naming master. A domain controller running Windows 2000 Server that holds the role of domain naming master must also be enabled as a global catalog server.

**Domain-wide operations master roles**
Every domain in the forest must have the following roles:

- Relative ID (RID) master
- Primary domain controller (PDC) emulator master
- Infrastructure master

These roles must be unique in each domain. This means that each domain in the forest can have only one RID master, PDC emulator master, and infrastructure master.

**RID master**
The RID master allocates sequences of relative IDs (RIDs) to each of the various domain controllers in its domain. At any time,

there can be only one domain controller acting as the RID master in each domain in the forest.

Whenever a domain controller creates a user, group, or computer object, it assigns the object a unique security ID (SID). The SID consists of a *domain* SID, which is the same for all SIDs created in the domain, and a RID, which is unique for each SID created in the domain.

To move an object between domains (using Movetree.exe), you must initiate the move on the domain controller acting as the RID master of the domain that currently contains the object.

**PDC emulator master**
The PDC emulator master processes password changes from client computers and replicates these updates to all domain controllers throughout the domain. At any time, there can be only one domain controller acting as the PDC emulator master in each domain in the forest.

The PDC emulator role is used in the following ways:

- To provide consistent password experience for users across sites (can be turned off with AvoidPdcOnWan registry parameter) - The PDC emulator is used as a reference DC to double-check incorrect passwords and it also receives new password changes. When the PDC is reachable, users can use a new password immediately and consistently across the environment. For more information about AvoidPdcOnWan see, New Password Change and Conflict Resolution Functionality in Windows (http://go.microsoft.com/fwlink/?LinkId=202451)

- As a preferred point of administration for services (examples are Group Policy and Distributed File System, DFS) - For more information about DFS see, DFS Management (http://go.microsoft.com/fwlink/?LinkId=202456). For more information about DCs and Group Policy see, Specifying a Domain Controller for Editing Group Policy (http://go.microsoft.com/fwlink/?LinkId=202457).

- As a point of contact for applications hard-coded to the PDC (often written for Windows NT 4.0 and older domains) - The legacy API often used for this is NetGetDcName. It is strongly suggested to change applications to use the new API to locate DCs. DsGetDcName by default does not target the PDC, and has more options that allows you to pick the type of DC needed to perform the operation. For more information about locating DCs from applications see, DSGetDcName Function (http://go.microsoft.com/fwlink/?LinkId=202455).

- As a default time server for all other DCs in the domain - The time server configuration of a PDC requires manual consideration and should be reviewed when you change the owner of the PDC role.

The domain controller configured with the PDC emulator role supports two authentication protocols:

- The Kerberos V5 protocol

- The NTLM protocol

| **Note** |
| --- |
| For instructions on how to configure the PDC emulator to synchronize with an external time source, see Synchronize the Time Server for the Domain Controller with an External Source (http://go.microsoft.com/fwlink/?LinkId=122513). |

**Infrastructure master**
At any time, there can be only one domain controller acting as the infrastructure master in each domain. The infrastructure master is responsible for updating references from objects in its domain to objects in other domains. The infrastructure master compares its data with that of a global catalog. Global catalogs receive regular updates for objects in all domains through replication, so the global catalog data will always be up to date. If the infrastructure master finds data that is out of date, it requests the updated data from a global catalog. The infrastructure master then replicates that updated data to the other domain controllers in the domain.

**Important**

- Unless there is only one domain controller in the domain, the infrastructure master role should not be assigned to the domain controller that is hosting the global catalog. If the infrastructure master and global catalog are on the same domain controller, the infrastructure master will not function. The infrastructure master will never find data that is out of

date, so it will never replicate any changes to the other domain controllers in the domain.

In the case where all of the domain controllers in a domain are also hosting the global catalog, all of the domain controllers will have the current data and it does not matter which domain controller holds the infrastructure master role.

The infrastructure master is also responsible for updating the group-to-user references whenever the members of groups are renamed or changed. When you rename or move a member of a group (and that member resides in a different domain from the group), the group may temporarily appear not to contain that member. The infrastructure master of the group's domain is responsible for updating the group so it knows the new name or location of the member. This prevents the loss of group memberships associated with a user account when the user account is renamed or moved. The infrastructure master distributes the update via multimaster replication.

There is no compromise to security during the time between the member rename and the group update. Only an administrator looking at that particular group membership would notice the temporary inconsistency.

For information about transferring operations master roles, see Transferring operations master roles. For information about what to do when an operations master fails, see Responding to operations master failures.

# Transferring operations master roles

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Transferring operations master roles**

Transferring an operations master role means moving it from one domain controller to another with the cooperation of the original role holder. Depending upon the operations master role to be transferred, you perform the role transfer using one of the three Active Directory consoles in Microsoft Management Console (MMC).

| Role | Console in MMC |
|---|---|
| Schema master | Active Directory Schema |
| Domain naming master | Active Directory Domains and Trusts |
| RID master | Active Directory Users and Computers |
| PDC emulator master | Active Directory Users and Computers |
| Infrastructure master | Active Directory Users and Computers |

For general information about operations master roles, see Operations master roles.

**Note**

- The operations master roles are sometimes called flexible single master operations (FSMO) roles.

For procedures describing the transfer of operations master roles, see:

- Transfer the schema master role
- Transfer the domain naming master role
- Transfer the RID master role
- Transfer the PDC emulator role
- Transfer the infrastructure master role

# Responding to operations master failures

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2

**Responding to operations master failures**

Some of the operations master roles are crucial to the operation of your network. Others can be unavailable for quite some time before their absence becomes a problem. Generally, you will notice that a single master operations role holder is unavailable when you try to perform some function controlled by the particular operations master.

If an operations master is not available due to computer failure or network problems, you can seize the operations master role. This is also referred to as forcing the transfer of the operations master role. Do not seize the operations master role if you can transfer it instead. For more information, see Transferring operations master roles.

### Note

- The operations master roles are sometimes called flexible single master operations (FSMO) roles.

Before forcing the transfer, first determine the cause and expected duration of the computer or network failure. If the cause is a networking problem or a server failure that will be resolved soon, wait for the role holder to become available again. If the domain controller that currently holds the role has failed, you must determine if it can be recovered and brought back online.

In general, seizing an operations master role is a drastic step that should be considered only if the current operations master will never be available again. The decision depends upon the role and how long the particular role holder will be unavailable. The impact of various role holder failures is discussed in the following topics.

**Schema master failure**

Temporary loss of the schema master is not visible to network users. It will not be visible to network administrators either, unless they are trying to modify the schema or install an application that modifies the schema during installation.

If the schema master will be unavailable for an unacceptable length of time, you can seize the role to the standby operations master. However, seizing this role is a drastic step that you should take only when the failure of the schema master is permanent.

### Important

- A domain controller whose schema master role has been seized must *never* be brought back online.

For procedures on how to seize the schema master role, see Seize the schema master role.

**Domain naming master failure**

Temporary loss of the domain naming master is not visible to network users. It will not be visible to network administrators either, unless they are trying to add a domain to the forest or remove a domain from the forest.

If the domain naming master will be unavailable for an unacceptable length of time, you can seize the role to the standby operations master. However, seizing this role is a drastic step that you should take only when the failure of the domain naming master is permanent.

### Important

- A domain controller whose domain naming master role has been seized must *never* be brought back online.

For procedures on how to seize the domain naming master role, see Seize the domain naming master role.

**RID master failure**

Temporary loss of the RID master is not visible to network users. It will not be visible to network administrators either, unless they are creating objects and the domain in which they are creating the objects runs out of relative IDs (RIDs).

If the RID master will be unavailable for an unacceptable length of time, you can seize the role to the operations master. However, seizing this role is a drastic step that you should take only when the failure of the RID master is permanent.

**Important**

- A domain controller whose RID master role has been seized must *never* be brought back online.

For procedures on how to seize the RID master role, see Seize the RID master role.

**PDC emulator master failure**

The severity of a PDC outage depends on your Service Level Agreement (SLA) and the actual behavior and configuration of the environment. For example, inconsistent password change behavior may affect users beyond what your SLAs allow, or the lack of time synchronization may cause resource access failures.

Also, in smaller environments, it may happen that the PDC as the first server in the domain is the only DNS or Global Catalog Server, or is the only domain controller (DC) with a valid SYSVOL in case other DCs did not successfully initiate or maintain SYSVOL replication. The PDC role holder may also be the target for regular file server access. When this is done for folder redirection or logon script activities, it may also affect users when logging on and while they work.

Other than the conditions described above, there is no direct dependency of the domain members on the PDC role holder. However, you might be using applications that are coded to contact the PDC only. You should try to avoid having this single point of failure.

Often, these applications were written for Windows NT 3.x and 4.0 deployments where the PDC was the only writable DC. However, since Active Directory, all DCs except Read-Only DCs are writable. The DsGetDcName API allows you to pick the right type; similar options are available in AD API interfaces like ADSI (ADS_READONLY_SERVER) or the .NET runtime.

The loss of the primary domain controller (PDC) emulator master may affect network users. Therefore, when the PDC emulator master is not available, you may need to immediately seize the role.

For procedures on how to seize the PDC emulator role, see Seize the PDC emulator role.

**Infrastructure master failure**

Temporary loss of the infrastructure master is not visible to network users. It will not be visible to network administrators either, unless they have recently moved or renamed a large number of accounts.

If the infrastructure master will be unavailable for an unacceptable length of time, you can seize the role to a domain controller that is not a global catalog but is well connected to a global catalog (from any domain), ideally in the same site as the current global catalog. When the original infrastructure master is returned to service, you can transfer the role back to the original domain controller.

For procedures on how to seize the infrastructure master role, see Seize the infrastructure master role.

For general information about operations masters, see Operations master roles.

# Understanding Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding groups**
This section covers:

- Groups
- Group scope
- Group types
- Default groups
- Nesting groups
- Special identities
- Where groups can be created

# Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Groups**

A group is a collection of user and computer accounts, contacts and other groups that can be managed as a single unit. Users and computers that belong to a particular group are referred to as group members.

Using groups can simplify administration by assigning a common set of permissions and rights to many accounts at once, rather than assigning permissions and rights to each account individually. For an overview of permissions and rights, see Access control overview.

Groups can be either directory-based or local to a particular computer. Groups in Active Directory are directory objects that reside within a domain and organizational unit container objects. Active Directory provides a set of default groups upon installation, and also allows the option to create groups. For more information, see Default groups.

Local groups, which exist on local computers and not in Active Directory, are discussed in Default local groups.

Groups in Active Directory allow you to:

- Simplify administration by assigning permissions on a shared resource to a group, rather than to individual users. This assigns the same access on the resource to all members of that group.
- Delegate administration by assigning user rights once to a group through Group Policy, and then adding necessary members to the group that you want to have the same rights as the group.
- Create e-mail distribution lists. For more information, see Group types.

Groups are characterized by their scope and their type. The scope of a group determines the extent to which the group is applied within a domain or forest. For information about group scope, see Group scope. The group type determines whether a group can be used to assign permissions from a shared resource (for security groups) or if a group can be used for e-mail distribution lists only (for distribution groups). For information about security groups and distribution groups, see Group types.

There are also groups for which you cannot modify or view the memberships. These groups are referred to as special identities and are used to represent different users at different times, depending on the circumstances. For example, the Everyone group represents all current network users, including guests and users from other domains. For more information, see Special identities.

# Group scope

Updated: October 1, 2014

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2

**Group scope**

Any group, whether it is a security group or a distribution group, is characterized by a scope that identifies the extent to which the group is applied in the domain tree or forest. The boundary, or reach, of a group scope is also determined by the domain functional level setting of the domain in which it resides. There are three group scopes: universal, global, and domain local.

The following table describes the differences between the scopes of each group.

| Group scope | Group can include as members... | Group can be assigned permissions in... | Group scope can be converted to... |
|---|---|---|---|
| Universal | • Accounts from any domain within the forest in which this Universal Group resides<br>• Global groups from any domain within the forest in which this Universal Group resides<br>• Universal groups from any domain within the forest in which this Universal Group resides | Any domain or forest | • Domain local<br>• Global (as long as no other universal groups exist as members) |
| Global | • Accounts from the same domain as the parent global group<br>• Global groups from the same domain as the parent global group | Member permissions can be assigned in any domain | Universal (as long as it is not a member of any other global groups) |
| Domain local | • Accounts from any domain<br>• Global groups from any domain<br>• Universal groups from any domain<br>• Domain local groups but only from the same domain as the parent domain local group | Member permissions can be assigned only within the same domain as the parent domain local group | Universal (as long as no other domain local groups exist as members) |

**Note**

The information in this table implies that the domain functional level is set to either Windows 2000 native or Windows Server 2003. When the domain functional level is set to Windows 2000 mixed or Windows Server 2003 interim, security groups with universal scope cannot be created, although distribution groups with universal scope are still permitted.

**When to use groups with domain local scope**

Groups with domain local scope help you define and manage access to resources within a single domain. For example, to give five users access to a particular printer, you can add all five user accounts in the printer permissions list. If, however, you later want to give the five users access to a new printer, you must again specify all five accounts in the permissions list for the new

printer.

With a little planning, you can simplify this routine administrative task by creating a group with domain local scope and assigning it permission to access the printer. Put the five user accounts in a group with global scope, and then add this group to the group having domain local scope. When you want to give the five users access to a new printer, assign the group with domain local scope permission to access the new printer. All members of the group with global scope automatically receive access to the new printer.

**When to use groups with global scope**

Use groups with global scope to manage directory objects that require daily maintenance, such as user and computer accounts. Because groups with global scope are not replicated outside their own domain, you can change accounts in a group having global scope frequently without generating replication traffic to the global catalog. For more information about groups and replication, see How replication works.

Although rights and permissions assignments are valid only within the domain in which they are assigned, by applying groups with global scope uniformly across the appropriate domains, you can consolidate references to accounts with similar purposes. This simplifies and rationalizes group management across domains. For example, in a network with two domains, Europe and UnitedStates, if you have a group with global scope called GLAccounting in the UnitedStates domain, create a group called GLAccounting in the Europe domain (unless the accounting function does not exist in the Europe domain).

It is strongly recommended that you use global groups or universal groups instead of domain local groups when you specify permissions on domain directory objects that are replicated to the global catalog. For more information, see Global catalog replication.

| **Note** |
| --- |
| When the domain functional level is set to Windows 2000 mixed, members of global groups can include only accounts from the same domain. |

**When to use groups with universal scope**

Use groups with universal scope to consolidate groups that span domains. To do this, add the accounts to groups with global scope, and then nest these groups within groups that have universal scope. When you use this strategy, any membership changes in the groups that have global scope do not affect the groups with universal scope.

For example, in a network with two domains, Europe and UnitedStates, and a group that has global scope called GLAccounting in each domain, create a group with universal scope called UAccounting that has as its members the two GLAccounting groups, UnitedStates\GLAccounting and Europe\GLAccounting. The UAccounting group can then be used anywhere in the enterprise. Any changes in the membership of the individual GLAccounting groups will not cause replication of the UAccounting group.

If the forest functional level is Windows Server 2003 or higher, changes to the membership of universal groups are replicated to each global catalog server using linked-value replication. This means that only the changed membership is replicated, rather than the entire group. If the forest functional level is lower than Windows Server 2003, you should not change the membership of a group with universal scope frequently because any changes to these group memberships cause the entire membership of the group to be replicated to every global catalog in the forest. For more information about universal groups and replication, see Global catalogs and sites. For more information about linked value replication, see How the Active Directory Replication Model Works.

| **Note** |
| --- |
| When the domain functional level is set to Windows 2000 mixed, you cannot create security groups with universal scope. |

**Changing group scope**

When you create a new group, by default the new group is configured as a security group with global scope, regardless of the current domain functional level. Although changing a group scope is not allowed in domains with a domain functional level of Windows 2000 mixed, the following conversions are allowed in domains with the domain functional level of Windows 2000 native or Windows Server 2003:

- **Global to universal.** This conversion is allowed only if the group that you want to change is not a member of another global scope group.

- **Domain local to universal.** This conversion is allowed only if the group that you want to change does not have another domain local group as a member.

- **Universal to global.** This conversion is allowed only if the group that you want to change does not have another universal group as a member.

- **Universal to domain local.** There are no restrictions for this operation.

For more information, see Change group scope.

**Groups on client computers and stand-alone servers**

Some group features, such as universal groups, group nesting, and the distinction between security groups and distribution groups, are available only on Active Directory domain controllers and member servers. Group accounts on Windows 2000 Professional, Windows XP Professional, Windows 2000 Server, and stand-alone servers running Windows Server 2003 work the same way as in Windows NT 4.0:

- Only local groups can be created locally on the computer.

- A local group that is created on one of these computers can be assigned permissions only on that one computer.

For more information, see Default local groups.

# Group types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group types**
Groups are used to collect user accounts, computer accounts, and other group accounts into manageable units. Working with groups instead of with individual users helps simplify network maintenance and administration.

There are two types of groups in Active Directory: distribution groups and security groups. You can use distribution groups to create e-mail distribution lists and security groups to assign permissions to shared resources.

**Distributions groups**
Distribution groups can be used only with e-mail applications (such as Exchange) to send e-mail to collections of users. Distribution groups are not security-enabled, which means that they cannot be listed in discretionary access control lists (DACLs). If you need a group for controlling access to shared resources, create a security group.

**Security groups**
Used with care, security groups provide an efficient way to assign access to resources on your network. Using security groups, you can:

- **Assign user rights to security groups in Active Directory**

  User rights are assigned to security groups to determine what members of that group can do within the scope of a domain (or forest). User rights are automatically assigned to some security groups at the time Active Directory is installed to help administrators define a person's administrative role in the domain. For example, a user who is added to the Backup Operators group in Active Directory has the ability to backup and restore files and directories located on each domain controller in the domain.

  This is possible because by default, the user rights Back up files and directories and Restore files and directories are automatically assigned to the Backup Operators group. Therefore, members of this group inherit the user rights assigned to that group. For more information about user rights, see User rights. For more information about the user rights assigned to security groups, see Default groups.

  You can assign user rights to security groups, using Group Policy, to help delegate specific tasks. You should always use discretion when assigning delegated tasks because an untrained user assigned too many rights on a security group can potentially cause significant harm to your network. For more information, see Delegating administration. For more information about assigning user rights to groups, see Assign user rights to a group in Active Directory.

- **Assign permissions to security groups on resources**

  Permissions should not be confused with user rights. Permissions are assigned to the security group on the shared resource. Permissions determine who can access the resource and the level of access, such as Full Control. Some permissions set on domain objects are automatically assigned to allow various levels of access to default security groups such as the Account Operators group or the Domain Admins group. For more information about permissions, see Access control in Active Directory.

  Security groups are listed in DACLs that define permissions on resources and objects. When assigning permissions for resources (file shares, printers, and so on), administrators should assign those permissions to a security group rather than to individual users. The permissions are assigned once to the group, instead of several times to each individual user. Each account added to a group receives the rights assigned to that group in Active Directory and the permissions defined for that group at the resource.

Like distribution groups, security groups can also be used as an e-mail entity. Sending an e-mail message to the group sends the message to all the members of the group.

**Converting between security and distribution groups**
A group can be converted from a security group to a distribution group, and vice versa, at any time, but only if the domain functional level is set to Windows 2000 native or higher. No groups can be converted while the domain functional level is set

to Windows 2000 mixed.

For specific procedural information, see Convert a group to another group type. For information about domain functionality, see Domain and forest functionality.

**Note**

- Although a contact can be added to a security group as well as to a distribution group, contacts cannot be assigned rights and permissions. Contacts in a group can be sent e-mail.

# Default groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Default groups**

Default groups, such as the Domain Admins group, are security groups that are created automatically when you create an Active Directory domain. You can use these predefined groups to help control access to shared resources and delegate specific domain-wide administrative roles. For information about default groups stored on *local* computers, see Default local groups.

Many default groups are automatically assigned a set of user rights that authorize members of the group to perform specific actions in a domain, such as logging on to a local system or backing up files and folders. For example, a member of the Backup Operators group has the right to perform backup operations for all domain controllers in the domain.

When you add a user to a group, the user receives all the user rights assigned to the group and all the permissions assigned to the group on any shared resources. For more information about user rights and permissions, see Group types.

You can manage groups by using the Active Directory Users and Computers snap-in in Microsoft Management Console (MMC). Default groups are located in the **Builtin** container and the **Users** container.

The **Builtin** container default groups contain groups that are defined with domain local scope. You can move groups in and out of this container, but you cannot move the default groups in this container to another location or to another domain.

The **Users** container default groups contain groups that are defined with global scope and groups that are defined with domain local scope. You can add the groups in this container to other groups and you can move the default groups in this container to other organizational units (OUs) or containers, but you cannot move the group to another domain.

As a security best practice, it is recommended that members of default groups with broad administrative access use the **Run as** command to perform administrative tasks. For more information, see Using Run as. For information about security best practices, see Active Directory Best practices. For information about additional security measures that can be used to protect Active Directory, see Securing Active Directory.

**Note**

- To learn what group you need to be a member of to perform a particular procedure, many procedural topics under **How To** in Help and Support Center provide a note that identifies this information.

**Groups in the Builtin container**

The following table provides descriptions of the default groups located in the Builtin container and lists the assigned user rights for each group. For complete descriptions of the user rights listed in the table, see User Rights Assignment. For information about editing these rights, see Edit security settings on a Group Policy object.

| Group | Description | Default user rights |
|---|---|---|
| Account Operators | Members of this group can create, modify, and delete accounts for users, groups, and computers located in the Users or Computers containers and organizational units in the domain, except the Domain Controllers organizational unit. Members of this group do not have permission to modify the Administrators or the Domain Admins groups, nor do they have permission to modify the accounts for members of those groups. Members of this group can log on locally to domain controllers in the domain and shut them down. Because this group has significant power in the domain, add users with caution. | Allow log on locally; Shut down the system. |

| | | |
|---|---|---|
| Administrators | Members of this group have full control of all domain controllers in the domain. By default, the Domain Admins and Enterprise Admins groups are members of the Administrators group. The Administrator account is also a default member. Because this group has full control in the domain, add users with caution. | Access this computer from the network; Adjust memory quotas for a process; Back up files and directories; Bypass traverse checking; Change the system time; Create a pagefile; Debug programs; Enable computer and user accounts to be trusted for delegation; Force a shutdown from a remote system; Increase scheduling priority; Load and unload device drivers; Allow log on locally; Manage auditing and security log; Modify firmware environment values; Profile single process; Profile system performance; Remove computer from docking station; Restore files and directories; Shut down the system; Take ownership of files or other objects. |
| Backup Operators | Members of this group can back up and restore all files on domain controllers in the domain, regardless of their own individual permissions on those files. Backup Operators can also log on to domain controllers and shut them down. This group has no default members. Because this group has significant power on domain controllers, add users with caution. | Back up files and directories; Allow log on locally; Restore files and directories; Shut down the system. |
| Guests | By default, the Domain Guests group is a member of this group. The Guest account (which is disabled by default) is also a default member of this group. | No default user rights. |
| Incoming Forest Trust Builders (only appears in the forest root domain) | Members of this group can create one-way, incoming forest trusts to the forest root domain. For example, members of this group residing in Forest A can create a one-way, incoming forest trust from Forest B. This one-way, incoming forest trust allows users in Forest A to access resources located in Forest B. Members of this group are granted the permission Create Inbound Forest Trust on the forest root domain. This group has no default members. For more information about creating forest trusts, see Create a forest trust. | No default user rights. |
| Network Configuration Operators | Members of this group can make changes to TCP/IP settings and renew and release TCP/IP addresses on domain controllers in the domain. This group has no default members. | No default user rights. |
| Performance Monitor Users | Members of this group can monitor performance counters on domain controllers in the domain, locally and from remote clients without being a member of the Administrators or Performance Log Users groups. | No default user rights. |
| Performance Log Users | Members of this group can manage performance counters, logs and alerts on domain controllers in the domain, locally and from remote clients without being a member of the Administrators group. | No default user rights. |
| Pre-Windows 2000 Compatible Access | Members of this group have read access on all users and groups in the domain. This group is provided for backward compatibility for computers running Windows NT 4.0 and earlier. By default, the special identity Everyone is a member of this group. For more information about special identities, see Special identities. Add users to this group only if they are running Windows NT 4.0 or earlier. | Access this computer from the network; Bypass traverse checking. |

| | | |
|---|---|---|
| Print Operators | Members of this group can manage, create, share, and delete printers connected to domain controllers in the domain. They can also manage Active Directory printer objects in the domain. Members of this group can log on locally to domain controllers in the domain and shut them down. This group has no default members. Because members of this group can load and unload device drivers on all domain controllers in the domain, add users with caution. | Allow log on locally; Shut down the system. |
| Remote Desktop Users | Members of this group can remotely log on to domain controllers in the domain. This group has no default members. | No default user rights. |
| Replicator | This group supports directory replication functions and is used by the File Replication service on domain controllers in the domain. This group has no default members. Do not add users to this group. | No default user rights. |
| Server Operators | On domain controllers, members of this group can log on interactively, create and delete shared resources, start and stop some services, back up and restore files, format the hard disk, and shut down the computer. This group has no default members. Because this group has significant power on domain controllers, add users with caution. | Back up files and directories; Change the system time; Force shutdown from a remote system; Allow log on locally; Restore files and directories; Shut down the system. |
| Users | Members of this group can perform most common tasks, such as running applications, using local and network printers, and locking the server. By default, the Domain Users group, Authenticated Users, and Interactive are members of this group. Therefore, any user account created in the domain becomes a member of this group. | No default user rights. |

**Groups in the Users container**

The following table provides a description of the default groups located in the Users container and lists the assigned user rights for each group. For complete descriptions of the user rights listed in the table, see User Rights Assignment. For information about editing these rights, see Edit security settings on a Group Policy object.

| Group | Description | Default user rights |
|---|---|---|
| Cert Publishers | Members of this group are permitted to publish certificates for users and computers. This group has no default members. | No default user rights. |
| DnsAdmins (installed with DNS) | Members of this group have administrative access to the DNS Server service. This group has no default members. | No default user rights. |
| DnsUpdateProxy (installed with DNS) | Members of this group are DNS clients that can perform dynamic updates on behalf of other clients, such as DHCP servers. This group has no default members. | No default user rights. |

| | | |
|---|---|---|
| Domain Admins | Members of this group have full control of the domain. By default, this group is a member of the Administrators group on all domain controllers, all domain workstations, and all domain member servers at the time they are joined to the domain. By default, the Administrator account is a member of this group. Because the group has full control in the domain, add users with caution. | Access this computer from the network; Adjust memory quotas for a process; Back up files and directories; Bypass traverse checking; Change the system time; Create a pagefile; Debug programs; Enable computer and user accounts to be trusted for delegation; Force a shutdown from a remote system; Increase scheduling priority; Load and unload device drivers; Allow log on locally; Manage auditing and security log; Modify firmware environment values; Profile single process; Profile system performance; Remove computer from docking station; Restore files and directories; Shut down the system; Take ownership of files or other objects. |
| Domain Computers | This group contains all workstations and servers joined to the domain. By default, any computer account created becomes a member of this group automatically. | No default user rights. |
| Domain Controllers | This group contains all domain controllers in the domain. | No default user rights. |
| Domain Guests | This group contains all domain guests. | No default user rights. |
| Domain Users | This group contains all domain users. By default, any user account created in the domain becomes a member of this group automatically. This group can be used to represent all users in the domain. For example, if you want all domain users to have access to a printer, you can assign permissions for the printer to this group (or add the Domain Users group to a local group, on the print server, that has permissions for the printer). | No default user rights. |
| Enterprise Admins (only appears in the forest root domain) | Members of this group have full control of all domains in the forest. By default, this group is a member of the Administrators group on all domain controllers in the forest. By default, the Administrator account is a member of this group. Because this group has full control of the forest, add users with caution. | Access this computer from the network; Adjust memory quotas for a process; Back up files and directories; Bypass traverse checking; Change the system time; Create a pagefile; Debug programs; Enable computer and user accounts to be trusted for delegation; Force shutdown from a remote system; Increase scheduling priority; Load and unload device drivers; Allow log on locally; Manage auditing and security log; Modify firmware environment values; Profile single process; Profile system performance; Remove computer from docking station; Restore files and directories; Shut down the system; Take ownership of files or other objects. |
| Group Policy Creator Owners | Members of this group can modify Group Policy in the domain. By default, the Administrator account is a member of this group. Because this group has significant power in the domain, add users with caution. | No default user rights. |

| | | |
|---|---|---|
| IIS_WPG (installed with IIS) | The IIS_WPG group is the Internet Information Services (IIS) 6.0 worker process group. Within the functioning of IIS 6.0 are worker processes that serve specific namespaces. For example, www.microsoft.com is a namespace served by one worker process, which can run under an identity added to the IIS_WPG group, such as MicrosoftAccount. This group has no default members. | No default user rights. |
| RAS and IAS Servers | Servers in this group are permitted access to the remote access properties of users. | No default user rights. |
| Schema Admins (only appears in the forest root domain) | Members of this group can modify the Active Directory schema. By default, the Administrator account is a member of this group. Because this group has significant power in the forest, add users with caution. | No default user rights. |

**See Also**

**Concepts**

Using Run as

Create a shortcut using the runas command

Why you should not run your computer as an administrator

**Other Resources**

Run a program with administrative credentials

# Nesting groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Nesting groups**

Using nesting, you can add a group as a member of another group. You nest groups to consolidate member accounts and reduce replication traffic.

Nesting options depend on whether the domain functionality of your Windows Server 2003 domain is set to Windows 2000 native or Windows 2000 mixed.

By default, when you nest a group within another group, user rights are inherited. For example, if you make Group_1 a member of Group_2, users in Group_1 have the same permissions as the users in Group_2.

Groups in domains set to the Windows 2000 native functional level or distribution groups in domains set to the Windows 2000 mixed functional level can have the following members:

- Groups with universal scope can have the following members: accounts, computer accounts, other groups with universal scope, and groups with global scope from any domain.
- Groups with global scope can have the following members: accounts from the same domain and other groups with global scope from the same domain.
- Groups with domain local scope can have the following members: accounts, groups with universal scope, and groups with global scope, all from any domain. This group can also have as members other groups with domain local scope from within the same domain.

Security groups in domains set to the Windows 2000 mixed functional level are restricted to the following types of membership:

- Groups with global scope can have as their members only accounts.
- Groups with domain local scope can have as their members other groups with global scope and accounts.

Security groups with universal scope cannot be created in domains with the domain functional level set to Windows 2000 mixed because universal scope is supported only in domains where the domain functional level is set to Windows 2000 native or Windows Server 2003.

| **Note** |
| --- |
| You cannot add the default groups that are located in the Builtin container as members to other groups. However, you can add other groups as members to the default groups that are located in the Builtin container. |

For more information about domain functionality, see Domain and forest functionality.

# Special identities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Special identities**
In addition to the groups in the Users and Builtin containers, servers running Windows Server 2003 include several special identities. For convenience, these identities are generally referred to as groups. These special groups do not have specific memberships that can be modified, but they can represent different users at different times, depending on the circumstances. The special groups are:

- Anonymous Logon

  Represents users and services that access a computer and its resources through the network without using an account name, password, or domain name. On computers running Windows NT and earlier, the Anonymous Logon group is a default member of the Everyone group. On computers running a member of the Windows Server 2003 family, the Anonymous Logon group is not a member of the Everyone group by default.

- Everyone

  Represents all current network users, including guests and users from other domains. Whenever a user logs on to the network, the user is automatically added to the Everyone group.

- Network

  Represents users currently accessing a given resource over the network (as opposed to users who access a resource by logging on locally at the computer where the resource is located). Whenever a user accesses a given resource over the network, the user is automatically added to the Network group.

- Interactive

  Represents all users currently logged on to a particular computer and accessing a given resource located on that computer (as opposed to users who access the resource over the network). Whenever a user accesses a given resource on the computer to which they are currently logged on, the user is automatically added to the Interactive group.

Although the special identities can be assigned rights and permissions to resources, the memberships cannot be modified or viewed. Group scopes do not apply to special identities. Users are automatically assigned to these special identities whenever they log on or access a particular resource.

For general information about groups, see Groups.

# Where groups can be created

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Where groups can be created**

In Active Directory, groups are created in domains. You use Active Directory Users and Computers to create groups. With the necessary permissions, groups can be created in the root domain of the forest, in any other domain in the forest, or in an organizational unit.

Besides the domain in which it is created, a group is also characterized by its scope. The scope of a group determines:

- The domain from which members can be added
- The domain in which the rights and permissions assigned to the group are valid

For more information about group scopes, see Group scope.

Choose the particular domain or organizational unit where you create a group based on the administration required for the group. For example, if your directory has multiple organizational units, each of which has a different administrator, you may want to create groups with global scope within those organizational units so those administrators can manage group membership for users in their respective organizational units. If groups are required for access control outside the organizational unit, the groups within the organizational unit can be nested into groups with universal scope (or other groups with global scope) that can be used elsewhere in the forest.

If the domain functional level is set to Windows 2000 native or higher, the domain contains a hierarchy of organizational units and administration is delegated to administrators at each organizational unit, it may be more efficient to nest groups with global scope. For example, if the organizational unit OU1 contained organizational units OU2 and OU3, a group with global scope in OU1 could have as its members groups with global scope in OU2 and OU3. In OU1, the administrator could add or remove group members from OU1, and the administrators of OU2 and OU3 could add or remove group members for accounts from their own organizational units without having administrative rights for the group with global scope in OU1.

**Note**

- Groups can be moved within a domain. However, only groups with universal scope can be moved from one domain to another. The rights and permissions assigned to a group with universal scope are lost when the group is moved to another domain and new assignments must be made.

For information about the tools used to move groups between domains, see Using the Windows Deployment and Resource Kits.

# Understanding Trusts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding trusts**
This section covers:


- Trusts
- Trust types
- Trust direction
- Trust transitivity
- When to create an external trust
- When to create a shortcut trust
- When to create a realm trust
- When to create a forest trust
- Forest trusts
- Accessing resources across domains
- Accessing resources across forests
- Routing name suffixes across forests

# Trusts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Trusts**
A trust is a relationship established between domains that enables users in one domain to be authenticated by a domain controller in the other domain. Trust relationships in Windows NT are different than in Windows 2000 and Windows Server 2003 operating systems.

**Trusts in Windows NT**
In Windows NT 4.0 and earlier, trusts are limited to two domains and the trust relationship is one-way and nontransitive. In the following figure, the nontransitive, one-way trust is shown by the straight arrow pointing to the trusted domain.




**Trusts in Windows Server 2003 and Windows 2000 Server operating systems**
All trusts in a Windows 2000 and Windows Server 2003 forest are transitive, two-way trusts. Therefore, both domains in a trust relationship are trusted. As shown in the following figure, this means that if Domain A trusts Domain B and Domain B trusts Domain C, then users from Domain C can access resources in Domain A (when assigned the proper permissions). Only members of the Domain Admins group can manage trust relationships.




**Trust protocols**
A domain controller running Windows Server 2003 authenticates users and applications using one of two protocols: Kerberos V5 or NTLM. The Kerberos V5 protocol is the default protocol for computers running Windows 2000, Windows XP Professional, or Windows Server 2003. If any computer involved in a transaction does not support Kerberos V5, the NTLM protocol will be used.

With the Kerberos V5 protocol, the client requests a ticket from a domain controller in its account domain to the server in the trusting domain. This ticket is issued by an intermediary trusted by the client and the server. The client presents this trusted ticket to the server in the trusting domain for authentication. For more information, see Kerberos V5 authentication.

When a client tries to access resources on a server in another domain using NTLM authentication, the server containing the resource must contact a domain controller in the client account domain to verify the account credentials.

**Trusted domain objects**
Trusted domain objects (TDOs) are objects that represent each trust relationship within a particular domain. Each time a trust is established a unique TDO is created and stored (in the System container) in its domain. Attributes such as a trust transitivity, type, and the reciprocal domain names are represented in a TDO.

Forest trust TDOs store additional attributes to identify all of the trusted namespaces from its partner forest. These attributes include domain tree names, user principal name (UPN) suffixes, service principal name (SPN) suffixes, and security ID (SID) namespaces.

For more information about domain trusts, see "Domain Trust" at the Microsoft Windows Resource Kits Web site. For more information about trust relationships, see "Designing an Authorization Strategy" at the Microsoft Windows Resource Kits Web

site.

# Trust types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Trust types**

Communication between domains occurs through trusts. Trusts are authentication pipelines that must be present in order for users in one domain to access resources in another domain. Two default trusts are created when using the Active Directory Installation Wizard. There are four other types of trusts that can be created using the New Trust Wizard or the Netdom command-line tool.

**Default trusts**

By default, two-way, transitive trusts are automatically created when a new domain is added to a domain tree or forest root domain using the Active Directory Installation Wizard. The two default trust types are defined in the following table.

| Trust type | Transitivity | Direction | Description |
|---|---|---|---|
| Parent and child | Transitive | Two-way | By default, when a new child domain is added to an existing domain tree, a new parent and child trust is established. Authentication requests made from subordinate domains flow upward through their parent to the trusting domain. For information about creating a new child domain, see Create a new child domain. |
| Tree-root | Transitive | Two-way | By default, when a new domain tree is created in an existing forest, a new tree-root trust is established. For information about creating a new domain tree, see Create a new domain tree. |

**Other trusts**

Four other types of trusts can be created using the New Trust Wizard or the Netdom command-line tool: external, realm, forest, and shortcut trusts. These trusts are defined in the following table.

| Trust type | Transitivity | Direction | Description |
|---|---|---|---|
| External | Nontransitive | One-way or two-way | Use external trusts to provide access to resources located on a Windows NT 4.0 domain or a domain located in a separate forest that is not joined by a forest trust. For more information, see When to create an external trust. |
| Realm | Transitive or nontransitive | One-way or two-way | Use realm trusts to form a trust relationship between a non-Windows Kerberos realm and a Windows Server 2003 domain. For more information, see When to create a realm trust. |
| Forest | Transitive | One-way or two-way | Use forest trusts to share resources between forests. If a forest trust is a two-way trust, authentication requests made in either forest can reach the other forest. For more information, see When to create a forest trust. |
| Shortcut | Transitive | One-way or two-way | Use shortcut trusts to improve user logon times between two domains within a Windows Server 2003 forest. This is useful when two domains are separated by two domain trees. For more information, see When to create a shortcut trust. |

When creating external, shortcut, realm, or forest trusts, you have the option to create each side of the trust separately or both sides of a trust simultaneously.

If you choose to create each side of the trust separately, then you will need to run the New Trust Wizard twice--once for each domain. When creating trusts using the method, you will need to supply the same trust password for each domain. As a security best practice, all trust passwords should be strong passwords. For more information, see Strong passwords.

If you choose to create both sides of the trust simultaneously, you will need to run the New Trust Wizard once. When you choose this option, a strong trust password is automatically generated for you.

You will need the appropriate administrative credentials for each domain between which you will be creating a trust.

Netdom.exe can also be used to create trusts. For more information about Netdom, see Active Directory support tools.

For more information about trusts, see Trust transitivity and Trust direction.

# Trust direction

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Trust direction**

The trust type and its assigned direction will impact the trust path used for authentication. A trust path is a series of trust relationships that authentication requests must follow between domains. Before a user can access a resource in another domain, the security system on domain controllers running Windows Server 2003 must determine whether the trusting domain (the domain containing the resource the user is trying to access) has a trust relationship with the trusted domain (the user's logon domain). To determine this, the security system computes the trust path between a domain controller in the trusting domain and a domain controller in the trusted domain. In the following figure, trust paths are indicated by arrows showing the direction of the trust (this is a one-way trust):




All domain trust relationships have only two domains in the relationship: the trusting domain and the trusted domain.

**One-way trust**

A one-way trust is a unidirectional authentication path created between two domains. This means that in a one-way trust between Domain A and Domain B, users in Domain A (trusted domain) can access resources in Domain B (trusting domain). However, users in Domain B cannot access resources in Domain A. Some one-way trusts can be a nontransitive trust or a transitive trust depending on the type of trust being created. For more information about trust types, see Trust types.

**Two-way trust**

All domain trusts in a Windows Server 2003 forest are two-way, transitive trusts. When a new child domain is created, a two-way, transitive trust is automatically created between the new child domain and the parent domain. In a two-way trust, both domains that are involved in a trust relationship trust each other. This means that authentication requests can be passed between the two domains in both directions. Some two-way relationships can be nontransitive or transitive depending on the type of trust being created. For more information, see Trust types.

A Windows Server 2003 domain can establish a one-way or two-way trust with:

- Windows Server 2003 domains in the same forest
- Windows Server 2003 domains in a different forest
- Windows NT 4.0 domains
- Kerberos V5 realms

For more information Kerberos V5, see Kerberos V5 authentication.

# Trust transitivity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Trust transitivity**
Transitivity determines whether a trust can be extended outside of the two domains with which it was formed. A transitive trust can be used to extend trust relationships with other domains and a nontransitive trust can be used to deny trust relationships with other domains.

**Transitive trusts**
Each time you create a new domain in a forest, a two-way, transitive trust relationship is automatically created between the new domain and its parent domain. If child domains are added to the new domain, the trust path flows upward through the domain hierarchy extending the initial trust path created between the new domain and its parent domain.

Transitive trust relationships flow upward through a domain tree as it is formed, creating transitive trusts between all domains in the domain tree.

Authentication requests follow these trust paths, so accounts from any domain in the forest can be authenticated at any other domain in the forest. With a single logon process, accounts with the proper permissions can access resources in any domain in the forest. For more information, see Authentication protocols overview.




The diagram displays that all domains in the Domain A tree and all domains in the Domain 1 tree have transitive trust relationships by default. As a result, users in the Domain A tree can access resources in domains in the Domain 1 tree and users in the Domain 1 tree can access resources in the Domain A tree, when the proper permissions are assigned at the resource.

In addition to the default transitive trusts established in a Windows Server 2003 forest, using the New Trust Wizard, you can manually create the following transitive trusts.

- **Shortcut trust.** A transitive trust between a domain in the same domain tree or forest used to shorten the trust path in a large and complex domain tree or forest.
- **Forest trust.** A transitive trust between a forest root domain and a second forest root domain.
- **Realm trust.** A transitive trust between an Active Directory domain and an Kerberos V5 realm. For more information about Kerberos V5 realms, see Kerberos V5 authentication.

For more information about trust types, see Trust types.

**Nontransitive trust**
A nontransitive trust is restricted by the two domains in the trust relationship and does not flow to any other domains in the forest. A nontransitive trust can be a two-way trust or a one-way trust.

Nontransitive trusts are one-way by default, although you can also create a two-way relationship by creating two one-way trusts. In summary, nontransitive domain trusts are the only form of trust relationship possible between:

- A Windows Server 2003 domain and a Windows NT domain
- A Windows Server 2003 domain in one forest and a domain in another forest (when not joined by a forest trust)

Using the New Trust Wizard, you manually create the following nontransitive trusts:

- **External trust.** A nontransitive trust created between a Windows Server 2003 domain and a Windows NT domain or a Windows 2000 domain or Windows Server 2003 domain in another forest.

  When you upgrade a Windows NT domain to a Windows Server 2003 domain, all existing Windows NT trusts are preserved intact. All trust relationships between Windows Server 2003 domains and Windows NT domains are nontransitive.

- **Realm trust.** A nontransitive trust between an Active Directory domain and an Kerberos V5 realm. For more information about Kerberos V5 realms, see Kerberos V5 authentication.

For more information about trust types, see Trust types.

# When to create an external trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**When to create an external trust**

You can create an external trust to form a one-way or two-way, nontransitive trust with domains outside of your forest. External trusts are sometimes necessary when users need access to resources located in a Windows NT 4.0 domain or in a domain located within a separate forest that is not joined by a forest trust, as shown in the figure.




When a trust is established between a domain in a particular forest and a domain outside of that forest, security principals from the external domain can access resources in the internal domain. Active Directory creates a foreign security principal object in the internal domain to represent each security principal from the trusted external domain. These foreign security principals can become members of domain local groups in the internal domain. Domain local groups can have members from domains outside of the forest.

Directory objects for foreign security principals are created by Active Directory and should not be manually modified. You can view foreign security principal objects from Active Directory Users and Computers by enabling advanced features. For information about enabling advanced features, see To view advanced features.

In domains with the functional level set to Windows 2000 mixed, it is recommended that you delete external trusts from a domain controller running Windows Server 2003. External trusts to Windows NT 4.0 or 3.51 domains can be deleted by authorized administrators on the domain controllers running Windows NT 4.0 or 3.51. However, only the trusted side of the relationship can be deleted on the domain controllers running Windows NT 4.0 or 3.51. The trusting side of the relationship (created in the Windows Server 2003 domain) is not deleted, and although it will not be operational, the trust will continue to display in Active Directory Domains and Trusts. To remove the trust completely, you will need to delete the trust from a domain controller running Windows Server 2003 in the trusting domain. If an external trust is inadvertently deleted from a domain controller running Windows NT 4.0 or 3.51, you will need to recreate the trust from any domain controller running Windows Server 2003 in the trusting domain.

For more information about how to create an external trust, see Create an external trust.

**Securing external trusts**

To improve the security of Active Directory forests, domain controllers running Windows Server 2003 and Windows 2000 Service Pack 4 (or higher) enable security identifier (SID) filter quarantining on all new outgoing external trusts by default.

By applying SID filter quarantining to outgoing external trusts, you prevent malicious users who have domain administrator level access in the trusted domain from granting, to themselves or other user accounts in their domain, elevated user rights to the trusting domain.

When a malicious user can grant unauthorized user rights to another user it is known as an elevation of privilege attack. For more information about SID filtering and how to further mitigate an elevation of privilege attack, see MS02-001: Forged SID could result in elevated privileges in Windows 2000 (http://go.microsoft.com/fwlink/?LinkId=102075).

**How SID filter quarantining works**

When security principals are created in a domain, the domain SID is included in the security principal's SID to identify the domain in which it was created. The domain SID is an important characteristic of a security principal because the Windows security subsystem uses it to verify the security principal's authenticity.

In a similar fashion, outgoing external trusts created from the trusting domain use SID filter quarantining to verify that

incoming authentication requests made from security principals in the trusted domain contain SIDs of security principals from the trusted domain only. This is done by comparing the SIDs of the incoming security principal to the domain SID of the trusted domain. If any of the security principal SIDs include a domain SID other than the one from the trusted domain, the trust removes the offending SID.

SID filtering ensures that any misuse of the SID history attribute on security principals (including inetOrgPerson) in the trusted forest cannot pose a threat to the integrity of the trusting forest.

The SID history attribute can be useful to domain administrators when they migrate user and group accounts from one domain to another. Domain administrators can add SIDs from an old user or group account to the SID history attribute of the new, migrated account. By doing this, domain administrators give the new account the same level of access to resources as the old account.

If domain administrators could not use the SID history attribute in this way, they would have to track down and reapply permissions for the new account on each network resource that the old account had access to.

**Understanding the threat**
If not for SID filtering on outgoing external trusts, a malicious user with administrative credentials residing in the trusted domain could sniff network authentication requests from the trusting domain to obtain the SID information of a user who has full access to resources in the trusting domain, such as a domain administrator.

After obtaining the domain administrators SID from the trusting domain, a malicious user with administrative credentials can add that SID to a user account's SID history attribute in the trusted domain and attempt to gain full access to the trusting domain and the resources within that domain. In this scenario, a malicious user who has domain administrator credentials in the trusted domain is a threat to the entire trusting forest.

SID filtering neutralizes the threat of malicious users in the trusted domain from using the SID history attribute to gain elevated privileges.

**Impact of SID filter quarantining**
SID filter quarantining on external trusts can affect your existing Active Directory infrastructure in the following two areas:

- SID history data that contains SIDs from any domain other than the trusted domain will be removed from authentication requests made from the trusted domain. This will result in access being denied to resources that have the user's old SID.

- Universal group access control strategy between forests will require changes.

When SID filter quarantining is enabled, users who use SID history data for authorization to resources in the trusting domain no longer have access to those resources.

If you typically assign universal groups from a trusted forest to access control lists (ACLs) on shared resources in the trusting domain, SID filter quarantining will have a major impact on your access control strategy.

Because universal groups must adhere to the same SID filter quarantining guidelines as other security principal objects (that is, the universal group object SID must also contain the domain SID), you should verify that any universal groups that are assigned to shared resources in the trusting domain were created in the trusted domain.

If the universal group in the trusted forest was not created in the trusted domain, even though it may contain users from the trusted domain as members, authentication requests made from members of that universal group will be filtered and discarded.

Therefore, before assigning access to resources in the trusting domain for users in the trusted domain, you should confirm that the universal group containing the trusted domain users was created in the trusted domain.

**Disabling SID Filter quarantining**
Although it is not recommended, you can disable SID filter quarantining for an external trust by using the Netdom.exe tool. You should consider disabling SID filter quarantining only in the following situations:

- You have the same level of trust for all administrators who have physical access to domain controllers in the trusted domain as the administrators in the trusting domain.

- You have a strict requirement to assign universal groups to resources in the trusting domain that were not created in the trusted domain.

- Users have been migrated to the trusted domain with their SID histories preserved, and you want to grant them access to resources in the trusting domain based on the SID history attribute.

Only domain administrators can disable SID filtering. To disable SID filter quarantining for the trusting domain, type the following syntax at a command-prompt:

**Netdom trust** *TrustingDomainName* **/domain:** *TrustedDomainName* **/quarantine:**No **/userD:***domainadministratorAcct***/passwordD:***domainadminpwd*

To enable SID filter quarantining, set the **/quarantine:** command-line option to **Yes**. For more information about Netdom.exe, see Active Directory support tools.

You can enable or disable SID filter quarantining only from the trusting side of the trust. If the trust is a two-way trust, you can also disable SID filter quarantining in the trusted domain by using the domain administrator's credentials for the trusted domain and reversing the *TrustingDomainName* and *TrustedDomainName* values in the command-line syntax.

**Notes**

- To further secure your forest, you should consider enabling SID filter quarantining on all existing external trusts that were created by domain controllers running Windows 2000 Service Pack 3 (or earlier). You can do this by using Netdom.exe to enable SID filtering on existing external trusts, or by recreating these external trusts from a domain controller running Windows Server 2003 or Windows 2000 Service Pack 4 (or later).

- You cannot turn off the default behavior that enables SID filter quarantining for newly created external trusts.

- External trusts created from domain controllers running Windows 2000 Service Pack 3 (or earlier) do not enforce SID filter quarantining by default.

- Domain controllers running Windows NT Server 4.0 do not take part in the trust creation process when existing domain controllers in the same domain are running Windows 2000 or Windows Server 2003.

- You can enable or disable SID filter quarantining only for trusts that extend beyond forest boundaries such as external and forest trusts. For more information about SID filtering and forest trusts, see Forest trusts.

**Allowing SID history to traverse forest trusts**

If you are migrating users from one domain to another in different forests, you may want to allow the migrated users to access resources in their original forest by using their migrated (SID history) credentials. The default SID filtering that is applied to forest trusts prevents user-resource-access requests from traversing the trusts with the credentials of the original domain. If you want to make it possible for users to use the credentials that were migrated from their original domain, you can allow SID history to traverse forest trusts by using the **netdom** command.

Only domain administrators or enterprise administrators can modify SID filtering settings. To allow SID history credentials to traverse a trust relationship between two forests, type a command using the following syntax at a command prompt, and then press ENTER:

**Netdom trust** *TrustingDomainName* **/domain:** *TrustedDomainName* **/enablesidhistory:**Yes **/usero:***domainadministratorAcct***/passwordo:***domainadminpwd*

To re-enable the default SID filtering setting across forest trusts, set the **/enablesidhistory:** command-line option to **No**. For more information about Netdom, see "Domain and Forest Trust Tools and Settings."

| **Note** |
| --- |
| The same security considerations for removing SID filter quarantining from external trusts apply to allowing SID history to traverse forest trusts. |

# When to create a shortcut trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**When to create a shortcut trust**

Shortcut trusts are one-way or two-way, transitive trusts that can be used when administrators need to optimize the authentication process. Authentication requests must first travel a trust path between domain trees, and in a complex forest this can take time, which can be reduced with shortcut trusts. A trust path is the series of domain trust relationships that must be traversed in order to pass authentication requests between any two domains. For more information about trust paths, see Trust direction.

Shortcut trusts are necessary when many users in a domain regularly log on to other domains in a forest. For example, using the following figure as an example, you could form a shortcut trust between domain B and domain D or domain A and domain 1 and so on.




Shortcut trusts effectively shorten the path traveled for authentication's made between domains located in two separate trees.

For more information about how to create a shortcut trust, see Create a shortcut trust.

**Using one-way trusts**

A one-way, shortcut trust established between two domains located in separate domain trees can reduce the time needed to fulfill authentication requests, but in only one direction. For example, when a one-way, shortcut trust is established between domain A and domain B, authentication requests made in domain A to domain B can utilize the new one-way trust path. However, authentication requests made in domain B to domain A will still need to travel the longer trust path.

**Using two-way trusts**

A two-way, shortcut trust established between two domains located in separate domain trees will reduce the time needed to fulfill authentication requests originating in either domain. For example, when a two-way trust is established between domain A and domain B, authentication requests made from either domain to the other can utilize the new, two-way trust path.

# When to create a realm trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**When to create a realm trust**

A realm trust can be established between any non-Windows Kerberos V5 realm and a Windows Server 2003 domain. This trust relationship allows cross-platform interoperability with security services based on other Kerberos V5 versions such as UNIX and MIT implementations. Realm trusts can switch from nontransitive to transitive and back. Realm trusts can also be either one-way or two-way.

For more information about how to create a realm trust, see Create a realm trust.

# When to create a forest trust

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**When to create a forest trust**

A forest trust can only be created between a forest root domain in one Windows Server 2003 forest and a forest root domain in another Windows Server 2003 forest. Creating a forest trust between two Windows Server 2003 forests provides a one-way or two-way, transitive trust relationship between every domain residing within each forest. Forest trusts are useful for Application Service Providers, companies undergoing mergers or acquisitions, collaborative business extranets, and companies seeking a solution for administrative autonomy.

For more information about creating forest trusts, see Checklist: Creating a forest trust.

| **Note** |
| --- |
| You should not enable SID filter quarantining on forest trusts, that is, by using the **netdom** command with the **/quarantine:yes** option. However, if you have migrated users from one Windows Server 2003 forest to another and the migrated users need access to resources in the former domain, you can relax the default SID filtering that is applied to a forest trust by using the **netdom** command with the **/enablesidhistory:yes** option. Using that command on a forest trust reduces the level of SID filtering on the forest trust. So, ensure that you trust the administrators of the trusted domain, as well as their security practices. |

**Using one-way forest trusts**

A one-way, forest trust between two forests allows members of the trusted forest to utilize resources located in the trusting forest. However, the trust operates in only one direction. For example, when a one-way, forest trust is created between forest A (the trusted forest) and forest B (the trusting forest), members of forest A can access resources located in forest B, but members of forest B cannot access resources located in forest A using the same trust.

**Using two-way forest trusts**

A two-way, forest trust between two forests allows members from either forest to utilize resources located in the other forest; domains in each respective forest trust domains in the other forest implicitly. For example, when a two-way, forest trust is established between forest A and forest B, members of forest A can access resources located in forest B, and members of forest B can access resources in forest A, using the same trust.

# Forest trusts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Forest trusts**

In a Windows Server 2003 forest, you can link two disjoined Windows Server 2003 forests together to form a one-way or two-way, transitive trust relationships. A two-way, forest trust is used to form a transitive trust relationship between every domain in both forests.

Forest trusts can provide the following benefits:

- Simplified management of resources across two Windows Server 2003 forests by reducing the number of external trusts necessary to share resources.
- Complete two-way trust relationships with every domain in each forest.
- Use of user principal name (UPN) authentication across two forests.
- Use of both the Kerberos V5 and NTLM authentication protocols to improve the trustworthiness of authorization data transferred between forests.
- Flexibility of administration. Administrative tasks can be unique to each forest.

Forest trusts can only be created between two forests and cannot be implicitly extended to a third forest. This means that if a forest trust is created between forest 1 and forest 2, and a forest trust is also created between forest 2 and forest 3, forest 1 will not have an implicit trust with forest 3. For more information about the requirements needed for a forest trust, see When to create a forest trust.

**Notes**

- In a Windows 2000 forest, if users in one forest need to access resources in another forest, an administrator can create an external trust relationship between the two domains. External trusts can be one-way or two-way and are nontransitive, and therefore, limit the ability for trust paths to extend to other domains. For more information about external trusts, see Trust types.
- All trusts in Windows Server 2003 Active Directory use security identifier (SID) filtering to some degree. External trusts are quarantined by default, which prevents any domain SIDs other than those of the quarantined trusted domain from traversing the trust relationship. SID filtering is used to prevent attacks from malicious users who might try to grant elevated user rights to another user account. SID filtering on forest trusts does not prevent migrations to domains within the same forest from using SID history and will not affect your universal group access control strategy. For more information about SID filtering, see When to create an external trust.

**Managing a multiple forest environment**

Forest trusts help you to manage a segmented Active Directory infrastructure within your organization by providing support for accessing resources and other objects across multiple forests. For more information about accessing resources across multiple forests, see Accessing resources across forests.

Because each forest is administered separately, adding additional forests to your organization increases your organization's management needs. For more information, see Creating a new forest.

Reasons to create multiple forests in your organization include:

- To secure data within each forest. Sensitive data can be protected so that only users within that forest can access it.
- To isolate directory replication within each forest. Schema changes, configuration changes, and the addition of new domains to a forest only have forest-wide impact within that forest, not on a trusting forest.

**Delegating forest-wide administrative control**
Active Directory data that is stored in the schema and configuration containers is replicated to every domain controller in the forest. Since changes to the schema and configuration containers will affect all domains in the forest, administrative control for forest-wide changes should be entrusted to highly trained or experienced administrators. All domain data contained in the forest root domain should also be regarded as highly sensitive data.

The following groups provide forest-wide administrative control in each forest:

- Enterprise Admins
- Domain Admins (in the forest root domain)
- Schema Admins

Since membership in any of these groups can affect forest-wide behavior, add users with caution. As a security best practice, avoid adding users from another forest to any of these forest-wide administrative groups. For more information about these groups, see Default groups.

**Synchronizing data across forests**
You can synchronize global address lists (GALs) and objects across forests using Microsoft Metadirectory Services (MMS) or another supported synchronization tool. Common data types that need synchronization across forests include:

- GALs (Exchange)
- Public folders
- Directory objects

Synchronizing this data across forests will help end users view address lists and other data the same way as they do when viewing this information within their own forest.

For more information about MMS, see Microsoft Metadirectory Services.

# Accessing resources across domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing resources across domains**

Since two or more Active Directory domains within the same forest are implicitly connected by two-way, transitive trusts, authentication requests made from one domain to another are successfully routed in order to provide a seamless coexistence of resources across domains. Users can only gain access to resources in other domains after first being authenticated in their own domain.

Domain controllers running Windows 2000 or Windows Server 2003 authenticate users and applications using one of two authentication protocols: Kerberos V5 or NTLM. For more information about Kerberos authentication, see Kerberos V5 authentication. For more information about NTLM authentication, see NTLM authentication. For more information about the Active Directory authentication process, see Access control in Active Directory.

Once authenticated, the user can attempt to gain access to resources from any domain in the forest using the Active Directory authorization process. For more information about the Active Directory authorization process, see Security information for Active Directory.

To access a shared resource in another domain using Kerberos, a user's workstation first requests a ticket from a domain controller in its account domain to the server (hosting the resource) in the trusting domain. This ticket is then issued by an intermediary trusted by the workstation and the server. The workstation presents this trusted ticket to the server in the trusting domain for authentication.

The following figure and corresponding steps provide a detailed description of the Kerberos authentication process that is used when a computer running Windows 2000 Professional, Windows 2000 Server, Windows XP Professional, or a member of the Windows Server 2003 family attempts to access resources from a server located in another domain.




1. User1 logs on to Workstation1 using credentials from the child1.microsoft.com domain. As part of this process, the authenticating domain controller issues User1 a ticket-granting ticket (TGT). This ticket is required to be authenticated to resources. The user then attempts to access a shared resource (\\fileserver1\share) on a file server located in the child2 domain.

2. Workstation1 contacts the Key Distribution Center (KDC) on a domain controller in its domain (ChildDC1) and requests a service ticket for the FileServer1 service principal name (SPN).

3. ChildDC1 does not find the SPN in its domain database and queries the global catalog to see if any domains in the forest contain this SPN. The global catalog sends the requested information back to ChildDC1.

4. ChildDC1 sends the referral to Workstation1.

5. Workstation1 contacts a domain controller in ForestRootDC1 (its parent domain) for a referral to a domain controller (ChildDC2) in the Child2 domain. ForestRootDC1 sends the referral to Workstation 1.

6. Workstation1 contacts the KDC on ChildDC2 and negotiates the ticket for the user to gain access to FileServer1.

7. Now that Workstation1 has a service ticket, it sends the service ticket to FileServer1, which reads the user's security credentials and constructs an access token accordingly.

Each domain has its own set of security policies that do not cross from one domain to another. For more information, see Domains.

**Planning your access control strategy for multiple domains**
It is recommended that you carefully plan the most efficient access control strategy for your organization's resource needs. Your design and implementation of security groups throughout each domain in a single forest will be an important factor to consider during your planning. For information about planning an access control strategy for multiple forests, see Accessing resources across forests.

It is important to understand the following security group concepts before you begin the planning process:

- **Security groups.** User rights can be applied to groups in Active Directory while permissions can be assigned to security groups on member servers hosting a resource. For more information, see Group types.

- **Group nesting.** The ability to nest security groups is dependent on group scopes and domain functionality. For more information, see Nesting groups.

- **Group scope.** Group scope helps determine the domain-wide and forest-wide access boundaries of security groups. For more information, see Group scope.

- **Domain functionality.** The domain functional level of the trusting and trusted domains can affect group functionality such as group nesting. For more information, see Domain and forest functionality.

Once you have gained a thorough understanding of security group concepts, determine the resource needs of each department and geographical division to assist you with the planning effort.

**Best practices for controlling access to shared resources across domains**
By carefully using domain local, global, and universal groups, administrators can more effectively control access to resources located in other domains. Consider the following best practices:

- Organize domain users based on administrative needs, such as their locations or departments, and then create a global group, and add the appropriate user accounts as members.

  For example, add all employee user accounts in the Sales department to the Sales Department global group, and add all employee user accounts in the Accounting Department to the Accounting Department global group.

- Create a domain local group, and add all global groups from the other domains that need the same access to a resource in your domain.

  For example, employees in the Sales Department and Accounting Department global groups located in DomainA use similar print resources located in DomainB. To make future administration changes more flexible, create a domain local group called Print Resources in DomainB, and add as members both the Sales Department and Accounting Department global groups in DomainA.

- Assign the required permissions on the shared resource to the domain local group.

  For example, assign permissions to the Print Resources domain local group located in DomainB so that its members, including the Sales Department and Accounting Department groups from DomainA, can access the printer located in DomainB.

**Selective authentication between domains in an external trust**

Using Active Directory Domains and Trusts, you can determine the scope of authentication between two domains that are joined by an external trust. You can set selective authentication differently for outgoing and incoming external trusts. With selective trusts, administrators can make flexible access control decisions between external domains. For more information about how to set selective authentication, see Select the scope of authentication for users.

If you use domain-wide authentication on the incoming external trust, users in the second domain would have the same level of access to resources in the local domain as users who belong to the local domain. For example, if DomainA has an incoming external trust from DomainB and domain-wide authentication is used, any user from DomainB would be able to access any resource in DomainA (assuming that they have the required permissions).

If you set selective authentication on an incoming external trust, you need to manually assign permissions on each resource to which you want users in the second domain to have access. To do this, set a control access right *Allowed to authenticate* on an object for that particular user or group from the external domain.

When a user authenticates across a trust with the **Selective authentication** option enabled, an *Other Organization*security ID (SID) is added to the user's authorization data. The presence of this SID prompts a check on the resource domain to ensure that the user is allowed to authenticate to the particular service. Once the user is authenticated, if the *Other Organization* SID is not already present, then the server adds the *This Organization* SID. Only one of these special SIDs can be present in an authenticated user's context. For more detailed information about how selective authentication works, see Security Considerations for Trusts.

Administrators in each domain can add objects from one domain to access control lists (ACLs) on shared resources in the other domain. You can use the ACL editor to add or remove objects residing in one domain to ACLs on resources in the other domain. For more information about how to set permissions on resources, see Set permissions on a shared resource.

For information about setting authentication restrictions for multiple forests, see Accessing resources across forests.

# Accessing resources across forests

Updated: February 21, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing resources across forests**

When two Windows Server 2003 forests are connected by a forest trust, authentication requests made using the Kerberos V5 or NTLM protocols can be routed between forests to provide access to resources in both forests. For more information about routing authentication requests across forests, see Routing name suffixes across forests.

Before authentication protocols can follow the forest trust path, the service principal name (SPN) of the resource computer must be resolved to a location in the other forest. An SPN can be one of the following:

- Domain Name System (DNS) name of a host
- DNS name of a domain
- Distinguished name of a service connection point object

When a workstation in one forest attempts to access data on the resource computer in another forest, Kerberos contacts the domain controller for a service ticket to the SPN of the resource computer. Once the domain controller queries the global catalog and identifies that the SPN is not in the same forest as the domain controller, the domain controller sends a referral for its parent domain back to the workstation. At that point, the workstation queries the parent domain for the service ticket and follows the referral chain until it gets to the domain where the resource is located.

The following figure and corresponding steps provide a detailed description of the Kerberos authentication process that is used when computers running Windows 2000 Professional, Windows XP Professional, Windows 2000 Server, or a member of the Windows Server 2003 family attempt to access resources from a computer located in another forest.




1. User1 logs on to Workstation1 using credentials from the child.microsoft.com domain. The user then attempts to access a shared resource on FileServer1 located in the msn.com forest.

2. Workstation1 contacts the Key Distribution Center (KDC) on a domain controller in its domain (ChildDC1) and requests a service ticket for the FileServer1 SPN.

3. ChildDC1 does not find the SPN in its domain database and queries the global catalog to see if any domains in the microsoft.com forest contain this SPN. Since a global catalog is limited to its own forest, the SPN is not found. The global catalog then checks its database for information about any forest trusts that are established with its forest, and, if found, it compares the name suffixes listed in the forest trust trusted domain object (TDO) to the suffix of the target SPN to find a match. Once a match is found, the global catalog provides a routing hint back to ChildDC1.

4. ChildDC1 sends a referral for its parent domain back to Workstation1.

5. Workstation1 contacts a domain controller in ForestRootDC1 (its parent domain) for a referral to a domain controller (ForestRootDC2) in the forest root domain of the msn.com forest.

6. Workstation1 contacts ForestRootDC2 in the msn.com forest for a service ticket to the requested service.

7. ForestRootDC2 contacts its global catalog to find the SPN, and the global catalog finds a match for the SPN and sends it back to ForestRootDC2.

8. ForestRootDC2 then sends the referral to child.msn.com back to Workstation1.

9. Workstation1 contacts the KDC on ChildDC2 and negotiates the ticket for User1 to gain access to FileServer1.

10. Now that workstation1 has a service ticket, it sends the server service ticket to FileServer1, which reads User1's security credentials and constructs an access token accordingly.

When a forest trust is first established, each forest collects all of the trusted namespaces in its partner forest and stores the information in a TDO. Trusted namespaces include domain tree names, user principal name (UPN) suffixes, service principal name (SPN) suffixes, and security ID (SID) namespaces used in the other forest. TDO objects are replicated to the global catalog. For more information about TDOs, see Trusts.

**Routing hints**

Routing hints are only used when all traditional authentication channels (local domain controller and then global catalog) have failed to produce a location of a SPN. Routing hints help direct authentication requests toward the destination forest.

When an SPN cannot be located in the domain from where the network logon request originated or from the global catalog database, the global catalog checks the forest trust TDO for trusted name suffixes located in the other forest that might match the suffix in the SPN. If a match is found, the forest root domain returns a routing hint back to the original source computer so that it can continue the SPN location process in the other forest.

### Notes

- Routing hints can only reference trusted name suffixes that are listed in the TDO for its forest trust. They do not verify the name suffix before sending the hint back to the original source computer.

- Accessing NetBIOS names and Kerberos delegation across forest trusts is not supported. NTLM is fully supported and cannot be disabled.

**Planning your access control strategy for multiple forests**

It is recommended that you carefully plan the most efficient access control strategy for your organization's resource needs. Your design and implementation of security groups throughout each forest will be an important factor to consider during your planning. For information about planning an access control strategy for multiple domains, see Accessing resources across domains.

It is important to understand the following security group concepts before you begin the planning process:

- **Security groups.** User rights can be applied to groups in Active Directory while permissions can be assigned to security groups on member servers hosting a resource. For more information, see Group types.

- **Group nesting.** The ability to nest security groups is dependent on groups scopes and domain functionality. For more information, see Nesting groups.

- **Group scope.** Group scope helps determine the domain-wide and forest-wide access boundaries of security groups. For more information, see Group scope.

- **Domain functionality.** The domain functional level of the trusting and trusted domains can affect group functionality such as group nesting. For more information, see Domain and forest functionality.

Once you have gained a thorough understanding of security group concepts, determine the resource needs of each department and geographical division to assist you with the planning effort.

**Best practices for using security groups across forests**

By carefully using domain local, global, and universal groups, administrators can more effectively control access to resources located in other forests. Consider the following best practices:

- To represent the sets of users who need access to the same types of resources, create role-based global groups in every domain and forest that contains these users. For example, users in the Sales Department in ForestA require access to an order-entry application that is a resource in ForestB. Account Department users in ForestA require access to the same application, but these users are in a different domain. In ForestA, create the global group SalesOrder and add users in the Sales Department to the group. Create the global group AccountsOrder and add users in the Accounting Department to that group.

- To group the users from one forest who require similar access to the same resources in a different forest, create universal groups that correspond to the global group roles. For example, in ForestA, create a universal group called SalesAccountsOrders and add the global groups SalesOrder and AccountsOrder to the group.

  | **Note** |
  | --- |
  | Universal groups are not available as security groups in Windows 2000 Server mixed-mode domains or in Windows Server 2003 domains that have a domain functional level of Windows 2000 mixed. They are available as distribution groups. |

- To assign permissions to resources that are to be accessed by users from a different forest, create resource-based domain local groups in every domain and use these groups to assign permissions on the resources in that domain. For example, in ForestB, create a domain local group called OrderEntryApp. Add this group to the access control list (ACL) that allows access to the order entry application, and assign appropriate permissions.

- To implement access to a resource across a forest, add universal groups from trusted forests to the domain local groups in the trusting forests. For example, add the SalesAccountsOrders universal group from ForestA to the OrderEntryApp domain local group in ForestB.

When a new user account needs access to a resource in a different forest, add the account to the respective global group in the domain of the user. When a new resource needs to be shared across forests, add the appropriate domain local group to the ACL for that resource. In this way, access is enabled across forests for resources on the basis of group membership.

For more information, see Set permissions on a shared resource.

**Selective authentication between forests**

Using Active Directory Domains and Trusts, you can determine the scope of authentication between two forests that are joined by a forest trust. You can set selective authentication differently for outgoing and incoming forest trusts. With selective trusts, administrators can make flexible forest-wide access control decisions. For more information about how to set selective authentication, see Select the scope of authentication for users.

If you use forest-wide authentication on an incoming forest trust, users from the outside forest have the same level of access to resources in the local forest as users who belong to the local forest. For example, if ForestA has an incoming forest trust from ForestB and forest-wide authentication is used, users from ForestB would be able to access any resource in ForestA (assuming they have the required permissions).

If you decide to set selective authentication on an incoming forest trust, you need to manually assign permissions on each computer in the domain as well as the resources to which you want users in the second forest to have access. To do this, set a control access right *Allowed to authenticate* on the computer object that hosts the resource in Active Directory Users and Computers in the second forest. Then, allow user or group access to the particular resources you want to share.

When a user authenticates across a trust with the **Selective authentication** option enabled, an *Other Organization* security ID (SID) is added to the user's authorization data. The presence of this SID prompts a check on the resource domain to ensure that the user is allowed to authenticate to the particular service. Once the user is authenticated, then the server to which he authenticates adds the *This Organization* SID if the *Other Organization* SID is not already present. Only one of these special

SIDs can be present in an authenticated user's context. For more detailed information about how selective authentication works, see Security Considerations for Trusts.

Administrators in each forest can add objects from one forest to access control lists (ACLs) on shared resources in the other forest. You can use the ACL editor to add or remove objects residing in one forest to ACLs on resources in another forest. For more information about how to set permissions on resources, see Set permissions on a shared resource.

For information about setting authentication restrictions for external domains, see Accessing resources across domains.

# Routing name suffixes across forests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing name suffixes across forests**

Name suffix routing is a mechanism used to manage how authentication requests are routed across Windows Server 2003 forests that are joined together by forest trusts. To simplify administration of authentication requests, when a forest trust is initially created, all unique name suffixes are routed by default. A unique name suffix is a name suffix within a forest, such as a user principal name (UPN) suffix, service principal name (SPN) suffix, or DNS forest or domain tree name, that is not subordinate to any other name suffix. For example, the DNS forest name microsoft.com is a unique name suffix within the microsoft.com forest.

Forests can contain multiple unique name suffixes, and all children of unique name suffixes are routed implicitly. In Active Directory Domains and Trusts, name suffixes appear with an asterisk (*) at the beginning because of this. For example, if your forest uses *.microsoft.com as a unique name suffix, then authentication requests for all children of microsoft.com (*.child.microsoft.com) will be routed because the child domains are part of the microsoft.com name suffix.

If a forest trust exists between two forests, then name suffixes that do not exist in one forest can be used to route authentication requests to a second forest. When a new child name suffix (*.child.widgets.com) is added to a unique name suffix (*.widgets.com), the child name suffix will inherit the routing configuration of the unique name suffix to which it belongs. Any new unique name suffixes that are created after a forest trust has been established will be visible in the forest trust **Properties** dialog box after you verify the trust. However, routing for those new unique name suffixes will be disabled by default. For more information about how to verify a trust, see Verify a trust.

When a duplicate name suffix is detected, the routing for the newest name suffix will be disabled by default. For more information about how to route name suffixes, see Enable or disable an existing name suffix from routing. Administrators can use the forest trust **Properties** dialog box to manually prevent authentication requests for specific name suffixes from being routed to a forest.

### Notes

- Do not add the @ sign to the UPN suffix or a user name. When authentication requests are routed to a trusted forest, all characters before the first @ symbol are interpreted as the user name and everything after the first @ symbol as the UPN suffix.

- The Local Security Authority (LSA) will block routing to any UPN suffix that is not a valid DNS name. For example, adding an @ symbol to a UPN suffix will cause it to automatically be disabled.

**Collision detection**

When two Windows Server 2003 forests are linked by a forest trust, there is a possibility that unique name suffixes existing in one forest might collide with unique name suffixes located in the second forest. Collision detection guarantees that each name suffix will only be routed to a single forest.

Active Directory Domains and Trusts will detect a name suffix conflict when:

- The same Domain Name System (DNS) name is already in use.

- The same NetBIOS name is already in use.

- A domain security ID (SID) conflicts with another name suffix SID.

When a name suffix in a forest conflicts with a new forest trust partner or when a name suffix in an existing forest trust conflicts with a new forest trust partner, the name will be disabled in the new trust. For example, a conflict will occur if one forest is named widgets.com and the second forest is named sales.widgets.com. Despite the name suffix conflict, routing will still work for any other unique name suffixes in the second forest.

For another example, assume that the msn.com forest needs to establish a two-way, forest trust with the widgets.com forest. Both msn.com and widgets.com have identical UPN suffixes of microsoft.com. During the creation of the two-way, forest trust, the New Trust Wizard will detect and display the conflict between the two UPN name suffixes and then create the forest trust.

If Active Directory Domains and Trusts detects a name suffix conflict with a trust partner domain, access to that domain from outside the forest may be denied. However, access to the conflicted domain from within its forest will function normally.

For example, if the domain widgets.com exists in both the msn.com and microsoft.com forests, users within the msn.com forest will be able to access resources in widgets.com domain that resides in the msn.com forest. However, users in the msn.com forest will be denied access to resources in the widgets.com domain located in the microsoft.com forest.

Conflicts will be listed in Active Directory Domains and Trusts, in the forest trust **Properties** dialog box, on the **Name Suffix Routing** tab. If a name suffix conflict is detected during the creation of the forest trust, then the New Trust Wizard will prompt you to save a log file of the conflicts. You can also save a log file after a trust has been created. For more information about how to save this log file, see Change the routing status of a name suffix.

If a NetBIOS or domain SID conflict exists, Active Directory Domains and Trusts identifies the name suffixes routing status as enabled or disabled with exceptions, as appropriate. The details about where the conflict exists are listed in the log file.

You can also use the Netdom command-line tool to list all routed names and to enable and disable routing for NetBIOS names and SIDs. For example, a forest named microsoft.com has a forest trust with widgets.com and you also need to add a forest trust to msn.com. Both widgets.com and msn.com have child domains with the NetBIOS name of SALES (and DNS names of USsales.widgets.com and sales.msn.com).

After you create the new trust to msn.com, routing to the SALES domain name in msn.com will be disabled. If you need to use the name SALES to route to the msn.com forest but don't need to use it to the widgets.com forest, you can use Netdom to disable SALES in widgets.com, and then enable it in msn.com.

For information about Netdom, see Active Directory support tools.

# Understanding Sites and Replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding sites and replication**
This section covers:

- Sites overview
- Replication overview
- How replication works
- Replication within a site
- Replication between sites
- When to establish a single or separate sites
- Bandwidth
- Managing replication

# Sites overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sites overview**

Sites in Active Directory® represent the physical structure, or topology, of your network. Active Directory uses topology information, stored as site and site link objects in the directory, to build the most efficient replication topology. You use Active Directory Sites and Services to define sites and site links. A site is a set of well-connected subnets. Sites differ from domains; sites represent the physical structure of your network, while domains represent the logical structure of your organization.

**Using sites**

Sites help facilitate several activities within Active Directory, including:

- **Replication.** Active Directory balances the need for up-to-date directory information with the need for bandwidth optimization by replicating information within a site more frequently than between sites. You can also configure the relative cost of connectivity between sites to further optimize replication. For more information, see Replication between sites and Managing replication.

- **Authentication.** Site information helps make authentication faster and more efficient. When a client logs on to a domain, it first searches its local site for a domain controller to authenticate against. By establishing multiple sites, you can ensure that clients authenticate against domain controllers nearest to them, reducing authentication latency and keeping traffic off WAN connections.

- **Active Directory-enabled services.** Active Directory-enabled services can leverage site and subnet information to enable clients to locate the nearest server providers more easily. For information about services, see Services.

**Defining sites using subnets**

In Active Directory, a site is a set of computers well-connected by a high-speed network, such as a local area network (LAN). All computers within the same site typically reside in the same building, or on the same campus network. A single site consists of one or more Internet Protocol (IP) subnets. Subnets are subdivisions of an IP network, with each subnet possessing its own unique network address. A subnet address groups neighboring computers in much the same way that postal codes group neighboring postal addresses. The following figure shows several clients within a subnet that defines an Active Directory site.




Sites and subnets are represented in Active Directory by site and subnet objects, which you create through Active Directory Sites and Services. Each site object is associated with one or more subnet objects.

For information about creating sites, see Create a site.

For information about creating subnets, see Create a subnet.

For information about subnets, see "Introduction to TCP/IP" at the Microsoft Windows Resource Kits Web site.

**Assigning computers to sites**

Computers are assigned to sites based on their Internet Protocol (IP) address and subnet mask. Site assignment is handled differently for clients and member servers than for domain controllers. For a client, site assignment is dynamically determined by its IP address and subnet mask during logon. For a domain controller, site membership is determined by the location of its associated server object in Active Directory. For more information, see "Active Directory Replication" at the Microsoft Windows Resource Kits Web site.

For information about associating subnets with sites, see Associate a subnet with a site.

For information about establishing single or multiple sites, see When to establish a single or separate sites.

**Understanding sites and domains**

In Active Directory, sites map the physical structure of your network, while domains map the logical or administrative structure of your organization. This separation of physical and logical structure provides the following benefits:

- You can design and maintain the logical and physical structures of your network independently.
- You do not have to base domain namespaces on your physical network.
- You can deploy domain controllers for multiple domains within the same site. You can also deploy domain controllers for the same domain in multiple sites.




Multiple sites in a single domain




Single site with multiple domains

For more information about domains, see Domains.

# Replication overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Replication overview**
Except for very small networks, directory data must reside in more than one place on the network to be equally useful to all users. Through replication, the Active Directory® directory service maintains replicas of directory data on multiple domain controllers, ensuring directory availability and performance for all users. Active Directory uses a multimaster replication model, allowing you to make directory changes at any domain controller, not just at a designated primary domain controller. Active Directory relies on the concept of sites to help keep replication efficient, and on the Knowledge Consistency Checker (KCC) to automatically determine the best replication topology for the network.

**Organizing data for replication**
Data is stored on each domain controller in the directory store, which is divided logically into specific directory partitions. Each partition stores a different type of directory data, either domain data, forest schema data, forest configuration data, or application data. All domain controllers within a forest hold a replica of the schema and configuration partitions for that forest, and all domain controllers within a particular domain hold a replica of the domain partition for their domain. Application directory partitions hold directory data specific to a particular application and can be stored by domain controllers belonging to different domains. Changes to each directory partition are replicated to all other domain controllers that hold a copy of that partition. For more information, see Directory data store.

Replication also ensures the availability of the global catalog throughout the entire forest. The global catalog is a searchable directory store containing data about every object in all domains. The global catalog is stored by domain controllers for which the global catalog has been enabled. For more information, see Global catalog replication.

**Improving replication efficiency with sites**
To help make replication more efficient, Active Directory relies on sites. Sites, defined as groups of well-connected computers, determine how directory data is replicated. Active Directory replicates directory information within a site more frequently than among sites. This way, the best connected domain controllers, those most likely to need particular directory information, receive replicated updates first. The domain controllers in other sites also receive the changes, but less frequently, reducing network bandwidth consumption. For more information, see How replication works and Sites overview.

**Determining the replication topology**
The Knowledge Consistency Checker (KCC), a process running on each domain controller, automatically identifies the most efficient replication topology for your network, based on information you provide about your network in Active Directory Sites and Services. The KCC regularly recalculate the replication topology to adjust for any network changes that have occurred. The KCC of one domain controller within each site (the intersite topology generator) determines the intersite replication topology. For more information about the KCC, see Active Directory Replication Technologies.

**Replication enhancements in the Windows Server 2003 family**
The Microsoft® Windows Server 2003 family includes enhancements to make replication both more efficient, as well as more scalable across a larger number of domains and sites. These include refinements in memory usage, enhancements to the Windows 2000 spanning tree algorithm, a completely new spanning tree algorithm for Windows Server 2003 forests, and a new load balancing tool.

In a forest set to the Windows 2000 functional level, the replication enhancements provide gains in replication efficiency and scalability, even when sites and domains contain domain controllers running Windows 2000. If a site contains at least one domain controller running Windows Server 2003, then a domain controller running Windows Server 2003 assumes the intersite topology generator role for the site, allowing the enhancements to take effect.

In a forest set to the Windows Server 2003 functional level, the new Windows Server 2003 spanning tree algorithm goes into effect for larger gains in both efficiency and scalability. For example, using the original spanning tree algorithm from Windows 2000, one domain can contain up to 300 sites. With the new Windows Server 2003 algorithm, one domain can contain up to at least 3,000 sites. In the new algorithm, the intersite topology generator in each site uses a randomized selection process to determine the bridgehead servers for the site. This selection process more evenly distributes the bridgehead replication workload among domain controllers in a site, resulting in much better efficiency (particularly in hub sites with a number of domain controllers). By default, the randomized selection process takes place only when new connection objects are added to the site. However, a new tool, called **adlb.exe**, can be run to rebalance the load each time changes occur in the topology or in the number of domain controllers in the site. In addition, **adlb** can stagger schedules so that the outbound replication load for each server is spread out evenly across time. For more information about **adlb** and to download the tool, see the "Windows Server 2003 Active Directory Branch Office Planning and Deployment Guide."

# How replication works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How replication works**
To keep directory data on all domain controllers consistent and up to date, Active Directory replicates directory changes on a regular basis. Replication occurs over standard network protocols, uses change tracking information to prevent unnecessary replication, and uses linked value replication to improve efficiency.

**Transferring replication data**
Active Directory uses remote procedure call (RPC) over Internet Protocol (IP) to transfer replication data between domain controllers. RPC over IP is used for both intersite and intrasite replication. To keep data secure while in transit, RPC over IP replication uses both authentication (using the Kerberos V5 authentication protocol) and data encryption.

When a direct or reliable IP connection is not available, replication between sites can be configured to use the Simple Mail Transfer Protocol (SMTP). However, SMTP replication functionality is limited, and requires an enterprise certification authority (CA). SMTP can only be used to replicate the configuration, schema and application directory partitions, and does not support the replication of domain directory partitions. For more information, see "Active Directory Replication" at the Microsoft Windows Resource Kits Web site.

**Preventing unnecessary replication**
Once a domain controller has processed a directory change from another domain controller successfully, it should not try to replicate those changes back to the domain controller that sent the change. In addition, a domain controller should avoid sending updates to another domain controller if the target domain controller has already received that same update from a different replication partner. To prevent such unnecessary replication, Active Directory uses change tracking information stored in the directory. For information about change tracking, see "Active Directory Replication" at the Microsoft Windows Resource Kits Web site.

**Resolving conflicting changes**
It is possible for two different users to make changes to the exact same object property and to have these changes applied at two different domain controllers in the same domain before replication of either change occurs. In such a case, both changes are replicated as new changes, creating a conflict. To resolve this conflict, domain controllers that receive these conflicting changes examine the attribute data contained within the changes, each of which holds a version and a timestamp. Domain controllers will accept the change with the higher version and discard the other change. If the versions are identical, domain controllers will accept the change with the more recent timestamp.

**Improving replication efficiency**
Introduced in the Windows Server 2003 family, linked value replication allows individual values of a multivalued attribute to be replicated separately. In Windows 2000, when a change was made to a member of a group (one example of a multivalued attribute with linked values) the entire group had to be replicated. With linked value replication, only the group member that has changed is replicated, and not the entire group. To enable linked value replication, you must raise the forest functional level to Windows Server 2003 . For more information about forest functional levels, see Domain and forest functionality. For more information about multivalued attributes, see Schema.

For more information about how replication works, see "Active Directory Replication" at the Microsoft Windows Resource Kits Web site.

# Replication within a site

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Replication within a site**

Active Directory handles replication within a site, or intrasite replication, differently than replication between sites because bandwidth within a site is more readily available. The Active Directory Knowledge Consistency Checker (KCC) builds the intrasite replication topology using a bidirectional ring design. Intrasite replication is optimized for speed, and directory updates within a site occur automatically on the basis of change notification. Unlike replication data travelling between sites, directory updates replicated within a site are not compressed.

For information about intersite replication, see Replication between sites and How replication works.

**Building the intrasite replication topology**

The Knowledge Consistency Checker (KCC) on each domain controller automatically builds the most efficient replication topology for intrasite replication, using a bidirectional ring design. This bidirectional ring topology attempts to create at least two connections to each domain controller (for fault tolerance) and no more than three hops between any two domain controllers (to reduce replication latency). To prevent connections of more than three hops, the topology can include shortcut connections across the ring. The KCC updates the replication topology regularly.

### Note

- The KCC actually creates a separate replication topology for each directory partition (schema, configuration, domain, application). Within a single site, these topologies are usually identical for all partitions hosted by the same set of the domain controllers.

**Determining when intrasite replication occurs**

Directory updates made within a site are likely to have the most direct impact on local clients, so intrasite replication is optimized for speed. Replication within a site occurs automatically on the basis of change notification. Intrasite replication begins when you make a directory update on a domain controller. By default, the source domain controller waits 15 seconds and then sends an update notification to its closest replication partner. If the source domain controller has more than one replication partner, subsequent notifications go out by default at 3 second intervals to each partner. After receiving notification of a change, a partner domain controller sends a directory update request to the source domain controller. The source domain controller responds to the request with a replication operation. The 3 second notification interval prevents the source domain controller from being overwhelmed with simultaneous update requests from its replication partners.

For some directory updates in a site, the 15 second waiting time does not apply and replication occurs immediately. Known as urgent replication, this immediate replication applies to critical directory updates, including the assigning of account lockouts and changes in the account lockout policy, the domain password policy, or the password on a domain controller account.

For more information about intrasite replication, see "Active Directory Replication" at the Microsoft Windows Resource Kits Web site.

# Replication between sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Replication between sites**
Active Directory handles replication between sites, or intersite replication, differently than replication within sites because bandwidth between sites is usually limited. The Active Directory Knowledge Consistency Checker (KCC) builds the intersite replication topology using a least-cost spanning tree design. Intersite replication is optimized for bandwidth efficiency, and directory updates between sites occur automatically based on a configurable schedule. Directory updates replicated between sites are compressed to preserve bandwidth.

For information about intrasite replication, see Replication within a site.

For information about site design, see "Designing the Site Topology" at the Microsoft Windows Resource Kits Web site.

**Building the intersite replication topology**
Active Directory automatically builds the most efficient intersite replication topology using information you provide (through Active Directory Sites and Services) about your site connections. The directory stores this information as site link objects. One domain controller per site, called the intersite topology generator, is assigned to build the topology. A least-cost spanning tree algorithm is used to eliminate redundant replication paths between sites. The intersite replication topology is updated regularly to respond to any changes that occur in the network. You can control intersite replication through the information you provide when you create your site links. For more information, see Managing replication.

**Determining when intersite replication occurs**
Active Directory preserves bandwidth between sites by minimizing the frequency of replication and by allowing you to schedule the availability of site links for replication. By default, intersite replication across each site link occurs every 180 minutes (3 hours). You can adjust this frequency to match your specific needs. Be aware that increasing this frequency increases the amount of bandwidth used by replication. In addition, you can schedule the availability of site links for use by replication. By default, a site link is available to carry replication traffic 24 hours a day, 7 days a week. You can limit this schedule to specific days of the week and times of day. You can, for example, schedule intersite replication so that it only occurs after normal business hours. For more information, see Configure site link replication frequency and Configure site link replication availability.

### Notes

- With certain restrictions, you can use the Simple Mail Transfer Protocol (SMTP) for replicating to sites that do not have a direct or reliable Internet Protocol (IP) connection. For more information, see "Active Directory Replication" at the Microsoft Windows Resource Kits Web site.

- Intersite replication through a firewall or virtual private network requires some special considerations. For more information, see Active Directory at the Microsoft Web site.

# When to establish a single or separate sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**When to establish a single or separate sites**
You can optimize the replication efficiency and reduce the administrative overhead of your network by establishing sites appropriately. The most effective number of sites depends on the physical design of your network. When you first create a new forest, a single, default Active Directory site (called Default-Site-First-Name) is created that represents your entire network. A forest or domain consisting of a single site can be very efficient for a single location network connected completely by high-speed bandwidth. If your forest or domain contains multiple geographic locations that communicate over low-speed wide area network (WAN) connections, establishing multiple sites gives you more detailed control of replication behavior, reduces authentication latency, and reduces network traffic on the WAN.

For information about sites, see Sites overview.

For information about designing a site topology, see "Designing the Site Topology" at the Microsoft Windows Resource Kits Web site.

**Why bandwidth is important**
Within a site, bandwidth affects how efficiently replication can work. The frequency with which intrasite replication occurs requires high-speed bandwidth to function most effectively. So before you a create new site, you should make sure that high-speed bandwidth connects all computers within the site candidate. Any area where domain controllers are connected by 10 megabits per second (Mbps) or more of bandwidth is a good site candidate.

**When to establish a single site**
If you have a single LAN consisting of a single subnet, or if your network contains multiple subnets connected by a high-speed backbone, establishing a single site replication topology can provide the following benefits:

- Simplified replication management
- Prompt directory updates between all domain controllers

A single site topology allows all replication on your network to occur as intrasite replication, which requires no manual replication configuration. A single site design also allows all domain controllers to remain very current with respect to directory changes, because directory updates are replicated almost immediately. For more information, see Replication within a site. For information about creating sites, see Create a site.

**When to establish multiple sites**
When your network consists of multiple geographic locations connected by a WAN, establishing separate sites for each location provides the following benefits:

- Efficient use of WAN bandwidth for replication
- Detailed control of replication behavior
- Reduction in authentication latency

Physically separate network locations typically communicate over WAN connections, which are most often characterized by low-speed bandwidth. By creating a separate site for each physical location on your network, you ensure that domain controllers communicating over WAN connections use intersite replication, which is specifically designed for efficiency on low bandwidth connections. For more information, see Replication between sites.

With multiple sites, you have more detailed control of replication behavior through several configurable intersite replication settings. These settings include the relative cost of different replication paths, the domain controllers associated with each site, the subnets associated with each site, the frequency of directory update transfers, and the availability of connections for use by replication. For more information, see Managing replication.

A network client logging on to a domain must contact a domain controller as part of the authentication process, first looking within its own site. If the site includes two or more physically separate network locations, the client may authenticate against a

domain controller across a WAN connection. Authentication across a WAN introduces a delay in the authentication process. By placing physically separate network locations into separate Active Directory sites, you can ensure that clients first attempt to authenticate against a domain controller in their own site.

# Bandwidth

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Bandwidth**

Bandwidth is the most important practical consideration affecting intrasite replication on your network. Although sites are conveniently defined by subnets, it is important to understand that the reason for this choice is that subnets are typically well-connected. This means that subnets are generally, but not always, an effective way to determine sites.

To organize your sites effectively, consider replication requirements and available connectivity. Find a balance that ensures domain controllers in the same site are sufficiently well-connected to handle the frequent exchange of directory information, while not exacting what you consider to be excessive costs (such as high financial expense or compromised network performance).

When considering bandwidth, you may want to combine several subnets into one site, or split one subnet among several sites, such as in the following instances:

- You have very poorly connected computers, including several domain controllers in one subnet. Create several smaller subnets that are better connected than they were as one large subnet. For this to be effective, each new subnet should contain a domain controller.

- You have several subnets that are all well-connected or you have fewer domain controllers than subnets, but you want several subnets to be in a single site. To do this, add multiple subnets to one site.

For more information about how bandwidth may affect the way you configure sites, see Replication between sites.

# Managing replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing replication**
Active Directory relies on site configuration information to manage and optimize the process of replication. Active Directory provides automatic configuration of these settings in some cases. In addition, you can configure site-related information for your network using Active Directory Sites and Services. Configurable information includes settings for site links, site link bridges, and bridgehead servers.

**Configuring site links**
You can use site link settings to control replication between sites. Configurable settings include the relative cost of each site link, the frequency of replication on each site link, and the schedule availability of each site link for replication. For information about site links, see Replication between sites.

**Site link cost**
The cost of a site link determines the relative preference of the Active Directory Knowledge Consistency Checker (KCC) for using a site link in the replication topology. The higher the cost of the site link, the lower will be the KCC's preference for using the site link. For example, if you have two site links, site link A and site link B, and you set the cost of site link A to 150 and the cost of site link B to 200, the KCC will prefer to use site link A in the replication topology. By default, the cost of a newly created site link is 100. For information about setting site link cost, see Configure site link cost. For information about the KCC, see Replication overview.

**Replication frequency**
The replication frequency of a site link determines how often replication occurs over that site link. By default, the replication frequency for a site link is 180 minutes, meaning that replication occurs over that site link every 180 minutes, or three hours. Using Active Directory Sites and Services, you can set the replication frequency from 15 minutes to 10,080 minutes (one week). A site link must be available for any replication to occur. If a site link is not available when the number of minutes between replication updates has passed, no replication will occur. For more information, see Configure site link replication frequency.

**Site link availability**
The availability schedule for a site link determines during which hours or days of the week a site link can be used for replication. By default, a site link is always available for replication, 24 hours a day and 7 days a week. You can change this schedule, for example, to exclude business hours during which your network is busy handling other types of traffic. Or, you can exclude particular days on which you do not want replication to occur. Scheduling information is ignored by site links that use the Simple Mail Transfer Protocol (SMTP) for replication. For more information, see Configure site link replication availability.

**Configuring site link bridges**
By default, all site links are bridged, or transitive. This allows any two sites that are not connected by an explicit site link to communicate directly, through a chain of intermediary site links and sites. One advantage to bridging all site links is that your network is easier to maintain because you do not need to create a site link to describe every possible path between pairs of sites.

Generally, you can leave automatic site link bridging enabled. However, you might want to disable automatic site link bridging and create site link bridges manually just for specific site links, in the following cases:

- Your network is not fully routed (not every domain controller can directly communicate with every other domain controller).
- You have a network routing or security policy in place that prevents every domain controller from being able to directly communicate with every other domain controller.
- Your Active Directory design includes a large number of sites. For more information, see the Active Directory Branch Office Planning Guide at the Microsoft Web site.

For more information about site link bridges and their affects on replication, see "Active Directory Replication" at the Microsoft Windows Resource Kits Web site.

For information about disabling automatic site link bridging, see Enable or disable site link bridges.

For information about creating a site link bridge manually, see Create a site link bridge.

**Configuring preferred bridgehead servers**

When the KCC constructs the intersite replication topology, it automatically assigns one or more bridgehead servers for each site to ensure that directory changes only need to be replicated across a site link one time. It is recommended that you allow the KCC to make the bridgehead server assignments. You can make the bridgehead server assignments manually through Active Directory Sites and Services. However, doing so can potentially disrupt replication if one of your manually assigned bridgehead servers becomes unavailable. For more information, see "Active Directory Replication" at the Microsoft Windows Resource Kits Web site.

For information about manually configuring a bridgehead server, see Designate a preferred bridgehead server.

# Understanding the Global Catalog

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding the global catalog**
This section covers:

- The role of the global catalog
- Global catalog replication
- Customizing the global catalog
- Global catalogs and sites

# The role of the global catalog

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The role of the global catalog**

A global catalog is a domain controller that stores a copy of all Active Directory objects in a forest. The global catalog stores a full copy of all objects in the directory for its host domain and a partial copy of all objects for all other domains in the forest, as shown in the following figure.




The partial copies of all domain objects included in the global catalog are those most commonly used in user search operations. These attributes are marked for inclusion in the global catalog as part of their schema definition. Storing the most commonly searched upon attributes of all domain objects in the global catalog provides users with efficient searches without affecting network performance with unnecessary referrals to domain controllers.

You can manually add or remove other object attributes to the global catalog by using the Active Directory Schema snap-in. For more information, see Customizing the global catalog.

A global catalog is created automatically on the initial domain controller in the forest. You can add global catalog functionality to other domain controllers or change the default location of the global catalog to another domain controller. For more information, see Enable or disable a global catalog.

A global catalog performs the following directory roles:

- **Finds objects**

  A global catalog enables user searches for directory information throughout all domains in a forest, regardless of where the data is stored. Searches within a forest are performed with maximum speed and minimum network traffic.

  When you search for people or printers from the Start menu or choose the Entire Directory option within a query, you are searching a global catalog. Once you enter your search request, it is routed to the default global catalog port 3268 and sent to a global catalog for resolution. For more information, see Finding directory information and "Finding information in Active Directory" at the Microsoft Windows Resource Kits Web site.

- **Supplies user principal name authentication**

  A global catalog resolves user principal names (UPNs) when the authenticating domain controller does not have knowledge of the account. For example, if a user's account is located in example1.microsoft.com and the user decides to log on with a user principal name of user1@example1.microsoft.com from a computer located in example2.microsoft.com, the domain controller in example2.microsoft.com will be unable to find the user's account, and will then contact a global catalog to complete the logon process. For more information, see Active Directory naming.

- **Supplies universal group membership information in a multiple domain environment**

Unlike global group memberships, which are stored in each domain, universal group memberships are only stored in a global catalog. For example, when a user who belongs to a universal group logs on to a domain that is set to the Windows 2000 native domain functional level or higher, the global catalog provides universal group membership information for the user's account at the time the user logs on to the domain.

If a global catalog is not available when a user logs on to a domain set to the functional level of Windows 2000 native or higher, the computer will use cached credentials to log on the user if the user has logged on to the domain previously. If the user has not logged on to the domain previously, the user can only log on to the local computer. However, if a user logs on as the Administrator in the domain (Builtin Administrator account), the user can always log on to the domain, even when a global catalog is not available.

For more information about universal groups, see Group scope. For more information about universal groups and replication, see Global catalog replication and Global catalogs and sites.

**Note**

- When there is only one domain in a forest, it is not necessary for users to obtain universal group memberships from a global catalog when logging on. This is because Active Directory can detect that there are no other domains in the forest and will prevent a query to the global catalog for this information.

- **Validates object references within a forest**

  A global catalog is used by domain controllers to validate references to objects of other domains in the forest. When a domain controller holds a directory object with an attribute containing a reference to an object in another domain, this reference is validated using a global catalog.

# Global catalog replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Global catalog replication**

Replication of the global catalog ensures that users throughout the forest have fast access to information about every object in the forest. The default attributes that make up the global catalog provide a baseline of the most commonly searched attributes. These attributes are replicated to the global catalog as part of normal Active Directory replication.

The replication topology for the global catalog is generated automatically by the Knowledge Consistency Checker (KCC). However, the global catalog is replicated only to other domain controllers that have been designated as global catalogs. Global catalog replication is affected both by the attributes marked for inclusion in the global catalog, and by universal group memberships.

**Adding attributes**

Active Directory defines a base set of attributes for each object in the directory. Each object and some of its attributes (such as universal group memberships) are stored in the global catalog. Using the Active Directory Schema snap-in, you can specify additional attributes to be kept in the global catalog.

In Windows 2000 forests, extending the partial attribute set causes a full synchronization of all object attributes stored in the global catalog (for all domains in the forest). In a large, multi-domain forest, this synchronization can cause significant network traffic. Between domain controllers enabled as global catalogs that are running Windows Server 2003, only the newly added attribute is replicated. For more information about adding attributes to the global catalog, see Customizing the global catalog.

**Preventing unpredictable access to global catalog data**

Special security consideration should be given when specifying permissions on domain data that is also replicated to the global catalog. When a user connects to a global catalog, an impersonation token is created for the user, which is used in subsequent access control decisions on the global catalog. The user's universal, global and domain local group memberships are represented in this token. However, only domain local groups from the domain that the domain controller hosting the global catalog (to which the user has connected) belongs to and of which the user is a member show up in the user's token. Domain local groups in the user's domain (and in other domains) of which the user is a member do not show up in the access token.

A global catalog stores a replicated, read-only copy of all objects in the forest and a partial set of each object's attributes, including the security descriptor for each object. The security descriptor contains a discretionary access control list (DACL), which specifies permissions on the object. When a user connects to a global catalog and tries to access an object, an access check is performed based on the user's token and the object's DACL. Any permissions specified in the object's DACL for domain local groups that are not from the domain that the domain controller hosting the global catalog (to which the user has connected) belongs to, will be ineffective because only domain local groups from the global catalog's domain of which the user is a member are represented in the user's access token. As a result, a user may be denied access when access should have been granted, or allowed access when access should have been denied.

As a best practice, you should avoid using domain local groups when assigning permissions on Active Directory objects, or be aware of the implications if you do use them. To prevent unauthorized access to global catalog data, use global groups or universal groups instead. For information about global and universal groups, see Group scope.

**How universal groups affect global catalog replication**

Groups with universal scope, and their members, are listed exclusively in the global catalog. Groups with global or domain local scope are also listed in the global catalog, but their members are not. This reduces the size of the global catalog and the replication traffic associated with keeping the global catalog up to date. You can improve network performance by using groups with global or domain local scope for directory objects that will change frequently.

When you first create a universal group, you do so from any domain that is set to the domain functional level of Windows 2000 or higher. The universal group resides in the domain directory partition in which it was created and is also replicated to the global catalog. Updates to the group membership are thereafter replicated to both the domain and the global catalog.

For more information about domain functional levels, see Domain and forest functionality.

# Customizing the global catalog

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Customizing the global catalog**

There may be instances where you will need to customize the global catalog to include additional attributes. However, you will want to carefully consider your options as changes to attributes can impact network traffic. By default, the global catalog contains an object's most common attributes for every object in the entire forest, which applications and users can query. For example, you can find a user by first name, last name, e-mail address, or other common properties of a user account.

To add additional searchable attributes to the global catalog, administrators can use the Active Directory Schema snap-in. For more information about adding additional attributes to the global catalog, see Add an attribute to the global catalog.

When determining whether or not to add an attribute to the global catalog, consider only adding additional attributes that are frequently queried and referenced by users or applications across the enterprise. Also consider how frequently an attribute gets updated during replication. Attributes that are stored in the global catalog are replicated to every global catalog in the forest. The smaller the attribute, the lower the impact of that replication. If the attribute is large, but very seldom changes, it will have a smaller replication impact than a small attribute that changes often.

For more information about attribute definitions, see the Active Directory Programmer's Guide at the Microsoft Web site.

**Important**

- It is strongly recommended that you use global groups or universal groups instead of domain local groups when specifying permissions on domain directory partition objects replicated to the global catalog. For more information, see Global catalog replication.

**Notes**

- In Windows 2000, adding a new attribute to the global catalog causes a full synchronization of all of the domain data from all of the domains in the forest. In a large, multi-domain Windows 2000 forest, this synchronization can cause significant network traffic. Between domain controllers enabled as global catalogs that are running Windows Server 2003, only the newly added attribute is replicated.

- When deciding whether or not to include an attribute in the global catalog, remember that you are trading increased replication and increased disk storage on global catalogs for potentially faster query performance.

# Global catalogs and sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Global catalogs and sites**

To optimize network performance in a multiple site environment, consider adding global catalogs for select sites. In a single site environment, a single global catalog is usually sufficient to cover common Active Directory queries. The following table will help you determine whether your multiple site environment will benefit using additional global catalogs.

| Use a global catalog when | Advantage | Disadvantage |
|---|---|---|
| A commonly used application in the site utilizes port 3268 to resolve global catalog queries. | Performance improvement | Additional network traffic due to replication |
| A slow or unreliable WAN connection is used to connect to other sites. Use the same failure and load distribution rules that you used for individual domain controllers to determine whether additional global catalog servers are necessary in each site. | Fault tolerance | Additional network traffic due to replication |
| Users in the site belong to a Windows 2000 domain running in native mode. In this case, all users must obtain universal group membership information from a global catalog server. If a global catalog is not located within the same site all logon requests must be routed over your WAN connection to a global catalog located in another site.<br><br>If a domain controller running Windows Server 2003 in the site has universal group membership caching enabled, then all users will obtain a current cached listing of their universal group memberships. | Fast user logons | Additional network traffic due to replication |

**Note**

- Network traffic related to global catalog queries generally use more network resources than normal directory replication traffic.

**Universal group membership caching**

Due to available network bandwidth and server hardware limitations, it may not be practical to have a global catalog in smaller branch office locations. For these sites, you can deploy domain controllers running Windows Server 2003, which can store universal group membership information locally.

Information is stored locally once this option is enabled and a user attempts to log on for the first time. The domain controller obtains the universal group membership for that user from a global catalog. Once the universal group membership information is obtained, it is cached on the domain controller for that site indefinitely and is periodically refreshed. The next time that user attempts to log on, the authenticating domain controller running Windows Server 2003 will obtain the universal group membership information from its local cache without the need to contact a global catalog.

By default, the universal group membership information contained in the cache of each domain controller will be refreshed every 8 hours. To refresh the cache, domain controllers running Windows Server 2003 will send a universal group membership confirmation request to a designated global catalog. Up to 500 universal group memberships can be updated at once. Universal group membership caching can be enabled using Active Directory Sites and Services. Universal group membership caching is site specific and requires that all domain controllers running Windows Server 2003 be located in that site to participate. For more information about how to enable this option, see Cache universal group memberships.

The following list summarizes potential benefits for caching universal group memberships in branch office locations:

- Faster logon times since authenticating domain controllers no longer need to access a global catalog to obtain universal group membership information.

- No need to upgrade hardware of existing domain controllers to handle the extra system requirements necessary for

hosting a global catalog.

- Minimized network bandwidth usage since a domain controller will not have to handle replication for all of the objects located in the forest.

**Note**

- You might want to continue using a global catalog in branch office locations if an application in a site is sending global catalog queries to port 3268. Universal group membership caching does not intercept calls made to port 3268.

# Interoperating with DNS and Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interoperating with DNS and Group Policy**

This section covers:

- DNS integration
- Group Policy integration

# DNS integration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS integration**

Active Directory is integrated with DNS in the following ways:

- Active Directory and Domain Name System (DNS) have the same hierarchical structure.

  Although separate and implemented differently for different purposes, an organization's namespace for DNS and Active Directory have an identical structure. For example, microsoft.com is a DNS domain and an Active Directory domain. For more information, see Namespace planning for DNS.

- DNS zones can be stored in Active Directory.

  If you are using the Windows Server 2003 DNS Server service, primary zone files can be stored in Active Directory for replication to other Active Directory domain controllers. For more information, see Active Directory integration.

- Active Directory uses DNS as a locator service, resolving Active Directory domain, site, and service names to an IP address.

  To log on to an Active Directory domain, an Active Directory client queries their configured DNS server for the IP address of the LDAP service running on a domain controller for a specified domain. For more information about how Active Directory clients rely on DNS, see Locating a domain controller.

**Note**

- You can use **Dcdiag.exe** and **Netdiag.exe** to troubleshoot client computers that cannot locate a domain controller. These tools can help determine both server and client DNS misconfigurations. For more information, see article Q265706, "DCDiag/NetDiag Facilitate Join and DC Creation" in the Microsoft Knowledge Base. For a brief description of support tools, see Active Directory support tools.

While Active Directory is integrated with DNS and shares the same namespace structure, it is important to distinguish the difference between them:

- DNS is a name resolution service.

  DNS clients send DNS name queries to their configured DNS server. The DNS server receives the name query and either resolves the name query through locally stored files or consults another DNS server for resolution. DNS does not require Active Directory to function.

- Active Directory is a directory service

  Active Directory provides an information repository and services to make information available to users and applications. Active Directory clients send queries to domain controllers using the Lightweight Directory Access Protocol (LDAP). In order to locate a domain controller, an Active Directory client queries DNS. Active Directory requires DNS to function.

For a checklist on deploying DNS for Active Directory, see Checklist: Deploying DNS for Active Directory.

For information about configuring DNS servers for Active Directory, see Configure a DNS server for use with Active Directory.

# Group Policy integration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy integration**

Group Policy can be used to define default settings that will be automatically applied to user and computer accounts in Active Directory. Policy settings can be used to manage desktop appearance, assign scripts, redirect folders from local computers to network locations, determine security options and control what software can be installed on particular computers and what software is available to particular groups of users.

Here are a few examples of how Group Policy settings can be used in Active Directory:

- Set the minimum password length and the maximum length of time that a password will remain valid. This can be configured for an entire domain.
- Administrators can automatically install an application on every computer in a particular domain or on all computers assigned to a particular group in a particular site. For example, you could automatically install Microsoft Outlook on every computer in the domain and automatically install Microsoft Excel only on those computers belonging to the Accounting group in a particular site.
- Logon, logoff, startup, and shutdown scripts can be assigned based on the locations of the computer and user accounts in Active Directory.
- If members of a particular group often use different computers, administrators can install the necessary applications on each of those computers.
- Any user's My Documents folder can be redirected to a network location. Users can then gain access to their documents from any computer on the network.

Policy settings are stored in Group Policy objects. Group Policy settings from more than one Group Policy object can be applied to a particular site, domain, or organizational unit. For example, if a site contains three domains, one Group Policy object could control computer configurations for the entire site. A separate policy for each domain could determine specific security settings for the computers in each domain. If each domain contains an Accounting and a Marketing organizational unit, additional Group Policy objects could determine what software is installed on the computers used by the Accounting and Marketing groups throughout the entire site.

This ability to automatically configure and secure computers throughout your organization by selectively applied Group Policy objects is a very powerful administrative tool. For more information about controlling software installation with Group Policy and how to create a Group Policy object, see Group Policy (pre-GPMC).

You can use security groups to filter how Group Policy settings are applied to collections of users and computers belonging to a particular site, domain, or organizational unit. For more information about security groups, see Group types. For general information about Group Policy, see Group Policy overview.

# Understanding Schema

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding schema**
This section covers:

- Schema
- Schema classes and attributes
- Schema object names
- Deactivating a class or attribute
- Extending the schema

# Schema

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Schema**
The Active Directory schema contains the definitions for all objects in the directory. Every new directory object you create is validated against the appropriate object definition in the schema before being written to the directory. The schema is made up of object classes and attributes. The base (or default) schema contains a rich set of object classes and attributes to meet the needs of most organizations, and is modeled after the International Standards Organization (ISO) X.500 standard for directory services. Because it is extensible, you can modify and add classes and attributes to the base schema. However, you should carefully consider each change you make, because extending the schema affects the entire network. For more information, see Extending the schema.

**How directory objects are defined**
In the schema, an object class represents a category of directory objects, such as users, printers, or application programs, that share a set of common characteristics. The definition for each object class contains a list of the schema attributes that can be used to describe instances of the class. For example, the User class has attributes such as givenName, surname, and streetAddress. When you create a new user in the directory the user becomes an instance of the User class, and the information you enter about the user becomes instances of the attributes. For more information, see Schema classes and attributes.

**How the schema is stored**
Each forest can contain only one schema, which is stored in the schema directory partition. The schema directory partition, along with the configuration directory partition, is replicated to all domain controllers in a forest. However, a single domain controller, the schema master, controls the structure and content of the schema. For more information about the schema master, see Operations master roles.

**Schema cache**
To improve performance on schema operations (such as new object validation), each domain controller holds a copy of the schema in memory (in addition to the copy it holds on disk). This cached version is automatically updated (after a small time interval) each time you update the schema. Or, you can reload the updated schema to cache manually for immediate effect. For more information, see Reload the schema.

**Securing the schema**
Like every object in Active Directory, schema objects are protected from unauthorized use by access control lists (ACLs). By default, only members of the Schema Admins group have write access to the schema. So, to extend the schema you must be a member of the Schema Admins group. The only default member of the Schema Admins group is the administrator account in the root domain of the forest. You should restrict membership in the Schema Admins group because extending the schema improperly can have serious consequences to your network. For more information, see Access control in Active Directory and Default groups.

For more information about the schema, see "Active Directory Schema" at the Microsoft Windows Resource Kits Web site and at the Microsoft MSDN Web site.

# Schema classes and attributes

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Schema classes and attributes**

Every directory object you create is an instance of an object class contained in the schema. Each object class contains a list of associated attributes that determine the information the object can contain. Classes and attributes are defined independently, so that a single attribute can be associated with multiple classes. All schema classes and attributes are defined by the classSchema and attributeSchema objects, respectively.

**Classes**

ClassSchema objects are used to define classes in the schema. A classSchema object provides the template for building directory objects of that class. Examples of classSchema include *User* and *Server*. A classSchema object contains, among other things, the following information:

- Class type (structural, abstract, or auxiliary)
- Common name and Lightweight Directory Access Protocol (LDAP) display name
- Lists of the "must contain" and "may contain" attributes for instances of the object
- Relative distinguished name attribute
- A list of possible parent classes

**Class types**

Three different types of classes exist in the schema:

| Class type | Purpose |
|---|---|
| **Structural** | Used to instantiate objects (users, servers and so on) in the directory. |
| **Abstract** | Provides templates for deriving structural classes |
| **Auxiliary** | Contains predefined lists of attributes that can be included in structural and abstract classes. |

**Note**

- With the Windows Server 2003 family, the inetOrgPerson class is now a part of base schema. This class can be used as a security principal in the same manner as the user class.

**Attributes**

AttributeSchema objects are used to define attributes in the schema. An attributeSchema object determines the allowable contents and syntax for instances of that attribute in the directory. Examples of attributeSchema include *User-Principal-Name* and *Telex-Number*. An attributeSchema object contains, among other things, the following information:

- Common name and LDAP display name
- Syntax rules
- Data constraints (single versus multivalued, minimum, and maximum values)
- Whether and how the attribute is indexed

**Single and multivalued attributes**

Attributes can be single-valued or multivalued. An instance of a single-valued attribute can can only contain a single value. An

instance of a multivalued attribute can contain multiple values of uniform syntax. A multivalued attribute stores no information about ordering of the attributes it contains. Each value of a multivalue attribute must be unique.

**Indexed attributes**

Both multivalued and single valued attributes can be indexed to help improve the performance of queries on that attribute. (Indexing does not apply to classes.) Attributes are marked for indexing based on their schema definition. Indexing an attribute also allows users to use wildcards (*) as prefixes and suffixes when specifying a search string. When you mark an attribute as indexed, all instances of the attribute are added to the index, not just the instances that are members of a particular class. Indexing attributes, particularly multivalued attributes, can negatively affect replication and object creation time, as well as directory database size. So, it is recommended that you only index commonly used attributes. For more information, see Index an attribute in Active Directory.

For more information about the schema, see "Active Directory Schema" at the Microsoft Windows Resource Kits Web site and the Microsoft MSDN Web site.

# Schema object names

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Schema object names**
When extending the schema, you need to know how to reference schema objects. Both class and attribute schema objects can be referenced in several ways:

- Lightweight Directory Access Protocol (LDAP) display name
- common name
- object identifier

**LDAP display name**
The Active Directory Schema snap-in and other administrative tools display the LDAP display name of objects. Programmers and system administrators use LDAP display names to reference objects programmatically. The LDAP display name typically consists of two or more words combined. When the name consists of multiple words, subsequent words in the name are identified using capitalization. Example LDAP display names are mailAddress and machinePasswordChangeInterval. The LDAP display name is guaranteed to be unique for each object.

**Common name**
The common name is a more readable version of the LDAP display name. The common names of the two attributes used in the previous example are SMTP-Mail-Address and Machine-Password-Change-Interval. Common names are guaranteed to be unique within a container.

**Object identifier**
An object identifier (also known as OID) is issued by an issuing authority such as the International Organization for Standardization (ISO) or the American National Standards Institute (ANSI). For example, the object identifier for the SMTP-Mail-Address attribute is 1.2.840.113556.1.4.786. Every object identifier must be unique. For more information about ISO, see the International Organization for Standardization Web site.

When extending your schema, you can register new object identifiers through Microsoft. For more information, see the Microsoft Web site.

**Schema object naming rules**
To help standardize schema naming conventions, Microsoft strongly suggests that schema extensions adhere to naming rules for both the LDAP Display Name and the Common Name. To qualify as "Certified for Windows," an application that extends the schema must adhere to these naming rules. For more information, see the Active Directory chapter of the certification program documentation at the Microsoft Web site. See also the Active Directory Programmer's Guide at the Microsoft Web site.

**Message queuing aliases**
A message queuing alias is an object in Active Directory that can be used to reference queues which might not be listed in Active Directory. For example, a queuing alias can be used to reference a private queue within the context of a distribution list. You can create a queuing alias by using Active Directory Users and Computers. Using queuing aliases provides the following benefits:

- When a queuing alias object is deleted, the alias is automatically removed from all distribution lists that reference the alias.
- A queue referenced by a queuing alias can be changed without changing the alias.
- Queuing aliases can be used to reference a queue not listed in the directory service, including private queues or queues from another organization.
- Queuing aliases can be used to send http messages by referencing the destination queue using a direct format name.

A queuing alias object has a single attribute, a format name that references a queue. Queuing aliases can contain public,

private, and direct format names. The format name for the queue cannot exceed 255 characters. For more information, see Using Message Queuing.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Deactivating a class or attribute

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deactivating a class or attribute**
Domain controllers running Windows Server 2003 do not permit the deletion of classes or attributes, but they can be deactivated if they are no longer needed or if there was an error in the original definition. A deactivated class or attribute is considered *defunct*. A defunct class or attribute is unavailable for use; however, it is easily reactivated.

If your forest has been raised to the Windows Server 2003 functional level, you can reuse the object identifier (governsId and attributeId values), the ldapDisplayName, and the schemaIdGUID that were associated with the defunct class or attribute. This allows you to change the object identifier associated with a particular class or attribute. The only exception to this is that an attribute used as a rdnAttId of a class continues to own its attributeId, ldapDisplayName, and schemaIdGuid values even after being deactivated (for example, those values cannot be reused).

If your forest has been raised to the Windows Server 2003 functional level, you can deactivate a class or attribute and then redefine it. For example, the Unicode String syntax of an attribute called SalesManager could be changed to Distinguished Name. Since Active Directory does not permit you to change the syntax of an attribute after it has been defined in the schema, you can deactivate the SalesManager attribute and create a new SalesManager attribute that reuses the same object identifier and LDAP display name as the old attribute, but with the desired attribute syntax. You must rename the deactivated attribute before it can be redefined.

For more information about forest functionality, see Domain and forest functionality.

**Notes**

- Classes and attributes added to the base schema can be deactivated without raising the forest functional level. However, they can be redefined only in forests raised to the Windows Server 2003 functional level or higher.
- Default schema attributes or classes in the base schema cannot be deactivated if bit 4 of the systemFlags attribute is set to 1. Only attributes or classes where bit 4 of the systemFlags attribute is equal to zero can be deactivated.

Object identifiers must be unique. Mistyping a number when entering an object identifier can cause a conflict between the invalid object identifier and a legitimate object identifier that is registered by an application. When installing an application that adds an attribute or class, the application may need to use the mistyped object identifier for one of its legitimate schema extensions. You can correct object identifier conflicts in forests that are set to the Windows Server 2003 functional level. To correct the conflict, deactivate the attribute or class with the incorrect object identifier. The application installation program will then be able to create the new attribute or class using the correct object identifier.

Before deactivating a class, consider the following:

- You can deactivate a class only if that class is not specified as a subClassOf, auxiliaryClass, systemAuxiliaryClass, possSuperiors, or systemPossSuperiors of any existing active class.
- You cannot use a defunct class in definitions of new classes, and you cannot add it to existing class definitions.
- You cannot create objects that are instances of the defunct class or modify existing instances of the class. However, the instances of the defunct class become available again for modification when the defunct class is reactivated.

Before deactivating an attribute, consider the following:

- You can deactivate an attribute only if the attribute is not specified as a systemMustContain, mustContain, systemMayContain, mayContain, or rdnAttId of any existing active class.
- You cannot use a defunct attribute in definitions of new classes, and you cannot add it to existing class definitions.
- You cannot read, modify, or delete instances of the defunct attribute present in existing objects. However, the instances of the defunct attribute become available when the defunct attribute is reactivated.

- To purge the directory of instances of an attribute you must delete the instances before deactivating the attribute.

Classes and attributes can be deactivated programmatically (recommended) or by using the Active Directory Schema snap-in. To deactivate a class or attribute using the Active Directory Schema snap-in, see Deactivate a class or attribute. For information about programmatically deactivating classes and attributes, see the Active Directory Programmer's Guide at the Microsoft Web site.

**Reactivating a class**
A defunct class can be reactivated. However, a defunct class cannot be reactivated unless the attributes referenced in its mustContain, systemMustContain, mayContain, and systemMayContain properties are active and the classes referenced in its subClassOf, auxiliaryClass, systemAuxiliaryClass, possSuperiors, and systemPossibleSuperiors properties are also active.

Further, a defunct class cannot be reactivated if the values of the following attributes collide with an already active attribute or class: ldapDisplayName, attributeId, governsId, or schemaIdGuid.

To reactive a defunct class, see Reactivate a class or attribute.

**Reactivating an attribute**
An defunct attribute can be reactivated. However, a defunct attribute cannot be reactivated if the values of the following attributes collide with an already active attribute or class: ldapDisplayName, attributeId, governsId, schemaIdGuid, or mapiId.

To reactive a defunct attribute, see Reactivate a class or attribute.

# Extending the schema

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Extending the schema**

When the set of classes and attributes in the base Active Directory schema do not meet your needs, you can extend the schema by modifying or adding classes and attributes. You should only extend the schema when absolutely necessary. The easiest way to extend the schema is through the Schema Microsoft Management Console (MMC) snap-in. You should always develop and test your schema extensions in a test lab before moving them to your production network.

**Before extending the schema**

Before extending the schema, review these key points:

| | |
|---|---|
| Check the base schema first | Verify that no existing class or attribute in the base schema meets your application or data needs. For the complete set of classes and attributes in the base schema, see the Microsoft Web site. |
| Review schema documentation | For detailed information about extending the schema, see the Active Directory Programmer's Guide at the Microsoft Web site and "Active Directory Schema" at the Microsoft Windows Resource Kits Web site. |
| Schema modifications are global | When you extend the schema, the changes apply to every domain controller in the entire forest. |
| Schema classes related to the system cannot be modified | You cannot modify default system classes (those classes required for Windows to run) within the schema. However, directory-enabled applications that modify the schema may add new classes that you can modify. |
| Schema extensions are not reversible | Attributes or classes cannot be removed after creation. At best, they can be modified or deactivated. For more information, see Deactivating a class or attribute. |
| Obtain valid object identifiers | Every class and attribute in the schema must have a unique and valid object identifier (also known as OID). Do not create arbitrary object identifiers or recycle old object identifiers. For information about obtaining valid object identifiers, see Schema object names. |
| Document your changes | If you do decide to extend the schema, be sure to document your changes. |

**How to extend the schema**

You can modify the schema through graphical user interface (GUI) tools, command-line tools, and through scripting. The easiest way to modify the schema is by using the Active Directory Schema snap-in in Microsoft Management Console (MMC), which is a GUI tool for schema management. For information about installing the Active Directory Schema snap-in, see Install the Active Directory Schema snap-in. Modifying the schema through scripting requires programming knowledge and familiarity with the Active Directory Service Interfaces (ADSI). For more information, see the Active Directory Programmer's Guide and Extending the Schema at the Microsoft MSDN Web site.

For more information about schema administration tools, see Administration tools for the Active Directory schema.

For more information about extending the schema, see Modify an existing schema class or attribute definition and Add a new schema class or attribute definition. For information about deactivation and reactivation, see Deactivating a class or attribute, Deactivate a class or attribute and Reactivate a class or attribute.

**Using a test forest**

A very simple way to avoid damaging or costly schema mistakes in your production forest is to first test your schema extensions on a test forest. By using a test environment, you can identify any potential problems in your plan before they affect your users and your production environment.

After making schema changes in a test forest, you can reinstall the default schema by demoting each domain controller in the test forest to which the schema changes have replicated. Then, use the Active Directory Installation Wizard to reinstall Active Directory on the servers. This procedure is practical only in a test environment.

# Deploying Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying Active Directory**
This section covers:

- Deployment resources
- Using the Active Directory Installation Wizard
- Creating an additional domain controller
- Creating a new domain tree
- Creating a new child domain
- Creating a new forest
- Upgrading from Windows NT or Windows 2000

# Deployment resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deployment resources**

To deploy domain controllers running Windows Server 2003, it is recommended that you use the *Microsoft Windows Server 2003 Deployment Kit*, which is available online at the Microsoft Windows Resource Kits Web site.

The online deployment kit includes case studies and the necessary information for deploying Windows Server 2003 Active Directory in networks that have domain controllers running Windows NT 4.0 or Windows 2000. For complete information, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003" at the Microsoft Windows Resource Kits Web site.

If your network includes branch offices, see "Designing the Site Topology" at the Microsoft Windows Resource Kits Web site. This online guide helps plan Active Directory deployment when there are branch office sites connected with slow network links.

For information about monitoring Active Directory, see "Monitoring Active Directory" at the Microsoft Windows Resource Kits Web site.

# Using the Active Directory Installation Wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the Active Directory Installation Wizard**
The Active Directory Installation Wizard installs and configures domain controllers, which provide network users and computers access to the Active Directory directory service. You can install Active Directory on any member server (except those with restrictive license agreements) using the Active Directory Installation Wizard. Using the wizard, you will define one of the following roles for the new domain controller:

- New forest (also a new domain)

  For a checklist about creating a new forest, see Checklist: Creating a new forest.

- New child domain

  For a checklist about creating a child domain, see Checklist: Creating a new child domain.

- New domain tree in an existing forest

  For a checklist about creating a new domain tree, see Checklist: Creating a new domain tree.

- An additional domain controller in an existing domain.

  For a checklist about creating an additional domain controller, see Checklist: Creating an additional domain controller in an existing domain.

Before using the Active Directory Installation Wizard, consider DNS configuration and support for existing applications.

**DNS configuration**
By default, the Active Directory Installation Wizard attempts to locate an authoritative DNS server for the new domain from its list of configured DNS servers that will accept a dynamic update of a service (SRV) resource record. If found, all the appropriate records for the domain controller are automatically registered with the DNS server after the domain controller is restarted.

If a DNS server that can accept dynamic updates is not found, either because the DNS server does not support dynamic updates or dynamic updates are not enabled for the domain, then the Active Directory Installation Wizard will take the following steps to ensure that the installation process is completed with the necessary registration of the SRV resource records:

1. The DNS service is installed on the domain controller and is automatically configured with a zone based on the Active Directory domain.

   For example, if the domain that you chose for your first domain in the forest is example.microsoft.com, then a zone rooted at the DNS domain name of example.microsoft.com is added and configured to use the DNS Server service on the new domain controller.

2. A text file containing the appropriate DNS resource records for the domain controller is created.

   The file called Netlogon.dns is created in the systemroot\System32\Config folder and contains all the records needed to register the resource records of the domain controller. Netlogon.dns is used by the Net Logon service and supports Active Directory on servers running non-Windows Server 2003 DNS.

   If you are using a DNS server that supports the SRV resource record but does not support dynamic updates (such as a UNIX-based DNS server or a Windows NT DNS server), you can import the records in Netlogon.dns into the appropriate primary zone file to manually configure the primary zone on that server to support Active Directory.

If no DNS servers are available on the network, you can choose the option to automatically install and configure a local DNS

server when you install Active Directory using the Active Directory Installation Wizard. The DNS server will be installed on the server on which you are running the wizard, and the server's preferred DNS server setting will be configured to use the new local DNS server.

Before running the Active Directory Installation Wizard, ensure that the authoritative DNS zone allows dynamic updates and that the DNS server hosting the zone supports the DNS SRV resource record. For more information, see Checklist: Verifying DNS before installing Active Directory.

For more information, see Configure a DNS server for use with Active Directory. For general information about DNS integration with Active Directory, see DNS integration.

**Support for existing applications**

On servers running Windows NT 4.0 and earlier, read access for user and group information is assigned to anonymous users so that existing applications and some non-Microsoft applications function correctly.

On servers running Windows 2000 and Windows Server 2003, members of the Anonymous Logon group have read access to this information only when the group is added to the Pre-Windows 2000 Compatible Access group.

Using the Active Directory Installation Wizard, you can choose if you want the Anonymous Logon group and the Everyone security groups to be added to the Pre-Windows 2000 Compatible Access group by selecting the **Permissions compatible with pre-Windows 2000 Server operating systems** option. To prevent members of the Anonymous Logon group from gaining read access to user and group information, choose the **Permissions compatible only with Windows Server 2003 operating systems** option.

When upgrading a domain controller from Windows 2000 to a Windows Server 2003 operating system, if the Everyone security group is already a member of the pre-Windows 2000 Compatible Access security group (indicating backward compatibility settings), the Anonymous Logon security group will be added as a member of the pre-Windows 2000 Compatible Access security group during the upgrade.

You can manually switch between the backward compatible and high-security settings on Active Directory objects by adding the Anonymous Logon security group to the pre-Windows 2000 Compatible Access security group using Active Directory Users and Computers. For more information about adding members to a group, see Add a member to a group. For more information about default groups, see Default groups and Special identities.

**Note**

- If you select the **Permissions compatible only with Windows Server 2003 operating systems** check box when installing Active Directory and find that your applications are not functioning correctly, try resolving the problem by manually adding the special group Everyone to the Pre-Windows 2000 Compatible Access security group, and then restarting the domain controllers in the domain. Once you have upgraded to applications compatible with the Windows Server 2003 family, you should return to the more secure Windows Server 2003 operating system configuration by removing the Everyone group from the Pre-Windows 2000 Compatible Access security group and restarting the domain controllers in the affected domain.

# Creating an additional domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating additional domain controllers**

If you already have one domain controller in a domain, you can add additional domain controllers to the domain to improve the availability and reliability of network services. Adding additional domain controllers can help provide fault tolerance, balance the load of existing domain controllers, and provide additional infrastructure support to sites.

More than one domain controller in a domain makes it possible for the domain to continue to function if a domain controller fails or must be disconnected. Multiple domain controllers can also improve performance by making it easier for clients to connect to a domain controller when logging on to the network. You can add additional domain controllers over the network or from backup media.

Before adding domain controllers you should thoroughly understand Active Directory and the requirements necessary to set up additional domain controllers in an existing domain. For more information, see Checklist: Creating an additional domain controller in an existing domain and Create an additional domain controller.

**Using backup media to create additional domain controllers**

With Windows 2000, the only way you can create an additional domain controller in an existing domain is by replicating the entire directory database to the new domain controller. With low network bandwidth or a large directory database, this replication can take hours or days to complete. With servers running Windows Server 2003, you can create an additional domain controller using a *restored backup* taken from a domain controller running Windows Server 2003. This backup can be stored on any backup media (tape, CD, or DVD) or a shared resource.

Using restored backup files to create an additional domain controller will greatly reduce the network bandwidth used when installing Active Directory over a shared resource; however, network connectivity is still necessary so that all new objects and recent changes to existing objects are replicated to the new domain controller.

It is recommended that you use the most recent backup available. Older backups require more network bandwidth for replication. The backup used cannot be older than the tombstone lifetime of the domain, which is set to a default value of 60 days (180 days in a forest that is created on a server running Windows Server 2003 with Service Pack 1 [SP1]).

If a domain controller that was backed up contained an application directory partition, it will not be restored on the new domain controller. To manually create an application directory partition on a new domain controller, see Create or delete an application directory partition.

When adding an additional domain controller using backup media, a System State backup taken only from a domain controllers running Windows Server 2003 can be used once it has been restored. For more information about how to restore a System State backup, see Restore System State data.

For general information about restoring backups, see Authoritative, primary, and normal restores.

# Creating a new domain tree

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating a new domain tree**

Create a new domain tree only when you need to create a domain whose DNS namespace is not related to the other domains in the forest. This means that the name of the tree root domain (and all of its children) does not have to contain the full name of the parent domain. A forest can contain one or more domain trees. For information about how to create a new domain tree, see Create a new domain tree.

Before you create a new domain tree, consider whether another forest is necessary. Multiple forests provide isolation of the schema and configuration directory partitions, separate security boundaries, administrative autonomy, and the flexibility to use an independent namespace design for each forest. Adding an additional forest does, however, increase the complexity and cost of implementing and maintaining your deployment. Therefore, the decision to create a new forest should not be made lightly. Create a forest where you have a specific deployment requirement. For example, create a forest when you need a separate security boundary or if you need to isolate schema changes. These requirements are not met by the creation of a new domain tree. For more information, see either Creating a new forest. or "Design Considerations for Delegation of Administration in Active Directory" on the Microsoft Web site.

# Creating a new child domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating a new child domain**

Create a new child domain when you want to create a domain that shares a contiguous namespace with one or more domains. This means that the name of the new domain contains the full name of the parent domain. For example, sales.microsoft.com would be a child domain of microsoft.com. As a best practice, you create new domains as children of the forest root domain.

You can create a new child domain by creating a new domain under a parent domain using the Active Directory Installation Wizard. For information about how to create a new child domain, see Create a new child domain.

After you create the child domain, you can create additional domain controllers in the child domain for fault tolerance and high availability of Active Directory. For more information, see Creating an additional domain controller.

# Creating a new forest

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating a new forest**

When you create the first domain controller in your organization, you are creating the first domain (also called the *forest root* domain) and the first forest.

The top-level Active Directory container is called a forest. A forest consists of one or more domains that share a common schema and global catalog. An organization can have multiple forests.

A forest is the security and administrative boundary for all objects that reside within the forest. In contrast, a domain is the administrative boundary for managing objects, such as users, groups, and computers. In addition, each domain has individual security policies and trust relationships with other domains.

Multiple domain trees within a single forest do not form a contiguous namespace; that is, they have noncontiguous DNS domain names. Although trees in a forest do not share a namespace, a forest does have a single root domain, called the forest root domain. The forest root domain is, by definition, the first domain created in the forest. The Enterprise Admins and Schema Admins groups are located in this domain. By default, members of these two groups have forest-wide administrative credentials.

**When to create a new forest**

A first step in the Active Directory design process is to determine how many forests your organization needs. For most organizations, a single forest design is the preferred model and the simplest to administer. However, a single forest may not be practical for every organization.

With a single forest, users do not need to be aware of directory structure because all users see a single directory through the global catalog. When adding a new domain to a forest, no additional trust configuration is required because all domains in a forest are connected by two-way, transitive trusts. In a forest with multiple domains, configuration changes need be applied only once to update all domains.

However, there are scenarios in which you might want to create more than one forest:

- **When upgrading a Windows NT domain to a Windows Server 2003 forest.** You can upgrade a Windows NT domain to become the first domain in a new Windows Server 2003 forest. To do this, you must first upgrade the primary domain controller in that domain. Then, you can upgrade backup domain controllers, member servers, and client computers at any time.

  You can also keep a Windows NT domain and create a new Windows Server 2003 forest by installing Active Directory on a member server running Windows Server 2003. For more information, see Upgrading from a Windows NT domain.

- **To provide administrative autonomy.** You can create a new forest when you need to segment your network for purposes of administrative autonomy. Administrators who currently manage the IT infrastructure for autonomous divisions within the organization may want to assume the role of forest owner and proceed with their own forest design. However, in other situations, potential forest owners may choose to merge their autonomous divisions into a single forest to reduce the cost of designing and operating their own Active Directory or to facilitate resource sharing. Another alternative is to provide for some delegation of administrative authority that enables the benefits of both approaches. For more information, see "Best Practice Active Directory Design for Managing Windows Networks" on the Microsoft Web site or "Design Considerations for Delegation of Administration in Active Directory", also available on the Microsoft Web site.

**Operations master roles in a new forest**

When you create the first forest in your organization, all five operations master roles are automatically assigned to the first domain controller in the forest. As new child domains are added to the forest, the first domain controller in each of the new child domains is automatically assigned the following roles:

- Relative identifier master

- Primary domain controller (PDC) emulator

- Infrastructure master

Because there can be only one schema master and one domain naming master in a forest, these roles remain in the forest root domain. In an Active Directory forest with only one domain and one domain controller, that domain controller owns all the operations master roles. For more information, see Operations master roles.

**Adding new domains to your forest**
A domain stores only the information about objects located in that domain, so by creating multiple domains within a new forest, you are partitioning or segmenting Active Directory to better serve a disparate user base.

The easiest domain structure to administer is a single domain within a single forest. When planning, you should start with a single domain and only add additional domains when the single domain model no longer meets your needs. For more information about creating domains, see Domains.

**Before creating a new forest**
Active Directory requires DNS to function and both share the same hierarchical domain structure. For example, microsoft.com is a DNS domain and an Active Directory domain. Because of the reliance that Active Directory has on DNS you must thoroughly understand Active Directory and DNS concepts before creating a new forest. For more information, see Checklist: Creating a new forest.

# Upgrading from Windows NT or Windows 2000

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrading from Windows NT or Windows 2000**
This section covers:

- Upgrading overview
- Upgrading from a Windows NT domain
- Upgrading from a Windows 2000 domain
- Converting groups

# Upgrading overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrading overview**

The Active Directory® directory service is compatible with Microsoft® Windows NT® and supports a mix of operations that support domain controllers running Windows NT 4.0, Windows® 2000, and Windows Server 2003. This allows you to upgrade domains and computers at your own pace, based on your organization's needs.

Active Directory supports the NTLM protocol used by Windows NT. This enables authorized users and computers from a Windows NT domain to log on and access resources in Windows 2000 or Windows Server 2003 domains. To clients running Windows 95, Windows 98, or Windows NT that are not running Active Directory client software, a Windows 2000 or Windows Server 2003 domain appears to be a Windows NT 4.0 domain. For more information, see Active Directory clients.

The upgrade to Active Directory can be gradual and performed without interrupting operations. If you follow domain upgrade recommendations, it should never be necessary to take a domain offline to upgrade domain controllers, member servers, or workstations. When upgrading a Windows NT domain, you must upgrade the primary domain controller first. You can upgrade member servers and workstations at any time after this. For more information, see Upgrading from a Windows NT domain.

Active Directory allows upgrading from any Windows NT 4.0 domain model and supports both centralized and decentralized models. The typical master or multiple-master domain model can be easily upgraded to an Active Directory forest.

# Upgrading from a Windows NT domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrading from a Windows NT domain**
See the following recommended steps for upgrading from a Windows NT domain.

**Plan and implement a namespace and DNS infrastructure**
Because Domain Name System (DNS) is required for Active Directory, ensure that you have designed a DNS and Active Directory namespace and have either configured DNS servers or are planning to use the Active Directory Installation Wizard to automatically install the DNS service. For more information, see DNS integration and Namespace planning for DNS.

**Determine forest functionality**
Forest functionality determines the type of Active Directory features that can be enabled within the scope of a single forest. Each forest functional level has a set of specific minimum requirements for the versions of operating systems that domain controllers throughout the forest can run. For example, the Windows Server 2003 forest functional level requires that all domain controllers run a product in the Windows Server 2003 family.

When you are upgrading your first Windows NT domain to become the first Windows Server 2003 domain, it is recommended that you set the forest functional level to the Windows Server 2003 interim forest functional level, which you will be prompted to set during the upgrade. This level contains all of the features used in the Windows 2000 forest functional level, including the following two important Active Directory replication enhancements:

- Improved replication efficiency and scalability. For more information, see Replication overview.
- Linked value replication for more efficient replication of group memberships. For more information, see How replication works.

For more information, see New features for Active Directory.

The Windows Server 2003 interim functional level is an option only when upgrading the first Windows NT domain to a new forest and can be manually configured after the upgrade. For more information about how to manually set this functional level, see "Upgrading Windows NT 4.0 Domains to Windows Server 2003" at the Microsoft Windows Resource Kits Web site. The Windows Server 2003 interim forest functional level only supports domain controllers running Windows Server 2003 and Windows NT, not domain controllers running Windows 2000. You cannot install Active Directory on servers running Windows 2000 in a forest with the functional level set to Windows Server 2003 interim. For more information about forest functionality, see Domain and forest functionality.

**Upgrade the primary domain controller**
The first server you must upgrade is the primary domain controller (PDC) running Windows NT 4.0 or earlier. During the upgrade, the Active Directory Installation Wizard requires that you choose to join an existing domain tree or forest or create a new domain tree or forest. If you decide to join an existing domain tree, you must provide a reference to the desired parent domain. For more information, see Checklist: Creating a new forest, Checklist: Creating a new child domain, or Checklist: Creating a new domain tree.

The Active Directory Installation Wizard installs all the necessary components on the domain controller, such as the directory data store and the Kerberos V5 protocol authentication software. Once the Kerberos V5 protocol is installed, the installation process starts the authentication service and the ticket granting service.

If this is a new child domain, then a transitive trust relationship is established with the parent domain. Eventually, the domain controller from the parent domain copies all schema and configuration information to the new child domain controller. The existing Security Accounts Manager (SAM) objects will be copied from the registry to the new data store. These objects are security principals.

During the upgrade, container objects are created to contain the accounts and groups from the Windows NT domain. The container objects are named Users, Computers, and Builtin and are displayed as folders in Active Directory Users and Computers. User accounts and predefined groups are routed to the Users container. Computer accounts are routed to the Computers container. Built-in groups are routed to the Builtin container.

Existing Windows NT 4.0 and earlier groups are located in different folders depending on the nature of the group. Windows NT

4.0 and earlier built-in local groups (such as Administrators and Server Operators) are located in the Builtin container. Windows NT 4.0 and earlier global groups (such as Domain Admins), any user-created local groups, and global groups are located in the Users container.

The upgraded PDC can synchronize security principal changes to remaining backup domain controllers (BDCs) running Windows NT 4.0 or earlier. The upgraded PDC is recognized as the domain master by BDCs running Windows NT Server 4.0 or earlier.

The upgraded domain controller is a fully functional member of the forest. The new domain is added to the domain and site structure and all domain controllers receive the notification that a new domain has joined the forest.

**Upgrade any remaining backup domain controllers**
Once you have upgraded the PDC running Windows NT 4.0 or earlier, you can proceed to upgrade all remaining BDCs. During the upgrade process, you might want to remove one BDC from the network to guarantee a backup if any problems develop. This BDC will store a secure copy of your current domain database.

If any problems arise during the upgrade, you can remove all domain controllers running Windows Server 2003 from the production environment, and then bring the BDC back into your network and make it the new PDC. This new PDC will then replicate its data throughout the domain so that the domain is returned to its previous state.

The only drawback to this method is that all changes that were made while the safe BDC was offline are lost. To minimize this loss, you could periodically turn the safe BDC on and off again (when the domain is in a stable state) during the upgrade process to update its safe copy of the directory.

If a domain controller running Windows Server 2003 becomes unavailable and no other domain controllers running Windows Server 2003 exist in the domain, a BDC running Windows NT can be promoted to a PDC to fill the role for the offline domain controller running Windows Server 2003.

When upgrading Windows NT 4.0 and earlier domains, only one domain controller running Windows Server 2003 can create security principals (users, groups, and computer accounts). This single domain controller is configured as a PDC emulator master. The PDC emulator master emulates a Windows NT 4.0 or earlier PDC. For more information about the PDC emulator role, see Operations master roles.

**Complete the upgrade of the domain**
After you have upgraded all existing Windows NT 4.0 and earlier primary and backup domain controllers to a Windows Server 2003 operating system, and you have no plans to use Windows NT 4.0 and earlier domain controllers, you can raise the domain functional level from Windows 2000 mixed to Windows 2000 native. For more information about how to raise the domain functional level, see Raise the domain functional level.

Several things happen when you raise the domain functional level to Windows 2000 native:

- Domain controllers no longer support NTLM replication.
- The domain controller that is emulating the PDC operations master cannot synchronize data with a BDC running Windows NT 4.0 or earlier.
- Domain controllers running Windows NT 4.0 and earlier cannot be added to the domain. You can add new domain controllers running Windows 2000 or Windows Server 2003.
- Users and computers using previous versions of Windows begin to benefit from the transitive trusts of Active Directory and can access resources anywhere in the forest with the appropriate permissions. Although previous versions of Windows do not support the Kerberos V5 protocol, the pass-through authentication provided by the domain controllers allows users and computers to be authenticated in any domain in the forest. This enables users or computers to access resources in any domain in the forest for which they have the appropriate permissions.

Other than the enhanced access to any other domains in the forest, clients will not be aware of any changes in the domain.

After upgrading a Windows NT 4.0 domain to an Active Directory domain, it is recommended that you delete and recreate all previously existing trusts with Active Directory domains. Even though the domain has been upgraded, the trust remains a Windows NT 4.0 trust. Internet protocol security (IPSec) does not work over a Windows NT 4.0 trust.

**Install Active Directory client software on older client computers**
Computers running Active Directory client software can use Active Directory features, such as authentication, to access

resources in the domain tree or forest and to query the directory. By default, client computers running Windows 2000 Professional and Windows XP Professional have the client software built in and can access Active Directory resources normally.

Computers running previous versions of Windows (Windows 95, Windows 98, and Windows NT) require installation of the Active Directory client software before access to Active Directory resources is available. Without the client software, previous versions of Windows can only access the domain as if it were a Windows NT 4.0 and earlier domain, finding only those resources available through Windows NT 4.0 and earlier one-way trusts. For more information about the Active Directory client software, see Active Directory clients.

# Upgrading from a Windows 2000 domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrading from a Windows 2000 domain**
Before upgrading a domain controller running Windows 2000 to Windows Server 2003, or installing Active Directory on the first domain controller running Windows Server 2003, ensure that your server, your forest, and your domain are ready. This process involves checking the upgrade compatibility of your server, preparing the forest, and preparing the domain, all of which can be accomplished using command-line tools.

To check the upgrade compatibility of the server and to copy any updated Setup files, run the **winnt32** command, using the **winnt32 /checkupgradeonly** command. For more information about the **winnt32** command, see Winnt32.

To prepare the forest, run the **adprep** command on the schema operations master, using the **adprep /forestprep** command. Running **adprep** on the schema master will update the schema, which will in turn replicate to all of the other domain controllers in the forest. The time it takes to replicate these changes will vary, depending on the replication schedule for your organization.

After the changes made by **adprep /forestprep** have replicated to the infrastructure operations master, you are ready to prepare the domain. To prepare the domain, run **adprep** on the infrastructure operations master, this time using the **adprep /domainprep** command.

On Windows Server 2003 domain controllers with no service pack installed **adprep /domainprep** can cause significant replication traffic due to synchronization of the Group Policy objects (GPOs) in the SYSVOL shared folder.

Windows Server 2003 with Service Pack 1 (SP1) includes a new **adprep** command, **adprep /domainprep /gpprep**, which you can run separately from **adprep /domainprep** and which you can run anytime to synchronize the GPOs that are located in the SYSVOL shared folder.

Therefore, if you are upgrading a Windows 2000 Server environment to Windows Server 2003 with SP1, prepare the forest using **adprep /forestprep** and prepare each domain using **adprep /domainprep**. Then, when network conditions are optimal, use **adprep /domainprep /gpprep** to add inheritable access control entries (ACEs) to the GPOs in the SYSVOL shared folder and to synchronize the contents of the SYSVOL shared folder among the domain controllers in the domain.

The Adprep command-line tool is available from the I386 directory on the operating system installation CD.

**Important**

- You cannot upgrade domain controllers running Windows 2000 to Windows Server 2003 or Windows Server 2003 with SP1, or add domain controllers running Windows Server 2003 or Windows Server 2003 with SP1 to Windows 2000 domains, until you have used **adprep** to prepare the forest and the domains within the forest.

For more information about the enhancements to Adprep.exe in Windows Server 2003 with SP1, see article 324392, “Enhancements to Adprep.exe in Windows Server 2003 Service Pack 1 and in hotfix 324392,” in the Microsoft Knowledge Base.

For more information about how to prepare your forest and domains using Adprep.exe, see "Overview: Upgrading Windows 2000 Domain Controllers to Windows Server 2003" in article 325379, “How to Upgrade Windows 2000 Domain Controllers to Windows Server 2003,” in the Microsoft Knowledge Base.

**See Also**
**Concepts**
Checklist: Performing an upgrade
Operations master roles
Adprep

# Converting groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Converting groups**

When you upgrade a primary domain controller (PDC) running Windows NT Server 4.0 or earlier to a server running Windows Server 2003, existing Windows NT groups are converted in the following ways:

- Windows NT local groups are converted to domain local groups on servers running Windows Server 2003.
- Windows NT global groups are converted to global groups on servers running Windows Server 2003.

Domain member computers running Windows NT can continue to display and access the converted groups. The groups appear to these clients as Windows NT 4.0 local and global groups. However, a Windows NT client cannot display members of groups or modify the member properties when that membership violates Windows NT group rules. For example, when a Windows NT client views the members of a global group on a server running Windows Server 2003, it does not view any other groups that are members of that global group.

**Converted groups and Microsoft Exchange**

Exchange allows users to arrange e-mail addresses in groups and distribution lists. When Exchange servers are upgraded to Active Directory, the Exchange distribution lists are converted to distribution groups with universal scope. The administrator can convert the group to a security group later, if desired, by using Active Directory Users and Computers to change the group properties. The messaging application programming interface (MAPI) enables computers running previous version Exchange clients to view the converted distribution group.

**Using converted groups with servers running Windows Server 2003**

Client computers that do not run Active Directory client software identify groups with universal scope on servers running Windows Server 2003 as having global scope instead. When viewing the members of a group with universal scope, the Windows NT client can only view and access group members that conform to the membership rules of global groups on servers running Windows Server 2003.

In a Windows Server 2003 domain that is set to a domain functional level of Windows 2000 native, all the domain controllers must be running on servers running Windows Server 2003. However, the domain can contain member servers that run Windows NT Server 4.0. These servers view groups with universal scope as having global scope and can assign groups with universal scope rights and permissions and place them in local groups.

In a Windows Server 2003 domain, a Windows NT Server 4.0 member server running Windows NT administrative tools cannot access domain local groups. However, you can work around this by using a server running Windows Server 2003 and using its Windows Server 2003 Administration Tools Pack to access the server running Windows NT Server 4.0. You can use the Administration Tools Pack to display the domain local groups and assign to them permissions to resources on the server running Windows NT Server 4.0.

# Administering Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Active Directory**
This section covers:

- Using Run as
- Using saved queries
- Managing Active Directory from MMC
- Managing Active Directory from the command line
- Finding directory information
- Administering other domains
- Delegating administration
- Publishing resources
- Managing COM+ partitions in Active Directory
- Managing COM+ partition sets in Active Directory

# Using Run as

Updated: April 20, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Run as**

Logging onto your computer using administrative credentials can pose a security risk to your computer and network. Therefore, as a security best practice, it is recommended that you do not log on to your computer with administrative credentials. Instead, you can use Run as to accomplish administrative tasks without having to log on to your computer with administrative credentials. For information about the security risks associated with being logged on as an administrator, see Why you should not run your computer as an administrator.

Using Run as, you can open and run a program using a different account and security context than the one you are logged on with. So, you can log on using a regular user account, then, using Run as, open an administrative program in the context of an administrative account. The administrative context is only used for that specific program and is only available until that program is closed.

It is especially important for domain administrators to use Run as to accomplish administrative tasks. Running your computer as a domain administrator in Active Directory makes your domain (and forest) vulnerable to Trojan horses and other security risks that target the logon sequence.

You can use Run as through the user interface or as a command-line tool.

- User interface

  The Run as feature built into the user interface is a shortcut that is accessed on the right-click menu for some programs (.exe), some Control Panel (.cpl) items, and Microsoft Management Console (MMC) (.msc) consoles. Run as will prompt you for a user account and password before starting a program, Control Panel item, or MMC console. Some programs and administrative tasks, such as upgrading the operating system or configuring system parameters, do not support Run as. These tasks require an interactive logon. For more information about how to use the Run as feature, see Run a program with administrative credentials.

- Command-line tool

  In addition to the built-in Run as feature, the **runas** command provides the same capabilities. For more information about the **runas** command, see Runas. You can also create a custom shortcut using the **runas** command. For more information about creating shortcuts, see Create a shortcut using the runas command.

| **Tip** |
| --- |
| For more options and examples, see the TechNet Wiki topic How to Run with Alternate Credentials and Open Elevated Command Prompts (http://social.technet.microsoft.com/wiki/contents/articles/how-to-run-with-alternate-credentials-and-open-elevated-command-prompts.aspx). |

# Using saved queries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using saved queries**

Saved queries provide a quick and consistent way for administrators to access a common set of directory objects on which to perform specific tasks or to monitor. Saved queries use predefined LDAP strings to search only the specified domain partition, and searches can be narrowed down to a single container object. You can also create a customized saved query that contains a LDAP search filter. For more information about saving queries, see Create a saved query. For more information about LDAP search filters, see the Internet Engineering Task Force Web site and search for "Representation of LDAP Search Filters."

Active Directory Users and Computers provides a Saved Queries folder in which administrators can create, edit, save, and organize saved queries. Before saved queries, administrators were required to create custom ADSI scripts that would perform a query on common objects. This was an often lengthy process that required knowledge of how ADSI utilizes LDAP search filters to resolve a query.

All queries located in the Saved Queries folder are stored in Active Directory Users and Computers (dsa.msc). Once you have successfully created your customized set of queries you can copy the .msc file to other domain controllers running Windows Server 2003 (located in the same domain) and use the same set of saved queries. You can also export saved queries to an .xml file and import them into other Active Directory User and Computer consoles located on domain controllers running Windows Server 2003 (within the same domain). For more information about saving console files, see Saving consoles.

You can save queries to search for disabled user or computer accounts, number of days since the last user logon, users with non-expiring Passwords, and many other commonly used queries. Once a saved query is executed and the desired objects are displayed, each object can then be modified directly through a query results screen. You can then select multiple objects and perform a task on them. For example, you can use a drag-and-drop operation to move two or more displayed objects onto a group.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Managing Active Directory from MMC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Active Directory from MMC**
The Active Directory administrative tools simplify directory service administration. You can use the standard tools or, using Microsoft Management Console (MMC), create custom tools that focus on single management tasks. You can combine several tools into one console. You can also assign custom tools to individual administrators with specific administrative responsibilities. For information about MMC, see Working with MMC console files.

The Active Directory administrative tools can only be used from a computer with access to a domain. The following Active Directory administrative tools are available on the Administrative Tools menu:

- Active Directory Users and Computers
- Active Directory Domains and Trusts
- Active Directory Sites and Services

You can also remotely administer Active Directory from a computer that is not a domain controller, such as a computer running Windows XP Professional. To do this, you must install the Windows Server 2003 Administration Tools Pack. For more information, see Windows Server 2003 Administration Tools Pack Overview.

The Active Directory Schema snap-in is an Active Directory administrative tool for managing the schema. It is not available by default on the Administrative Tools menu, and must be added manually. For more information, see Install the Active Directory Schema snap-in.

For advanced administrators and network support specialists, there are many command-line tools that can be used to configure, manage, and troubleshoot Active Directory. For more information, see Active Directory support tools.

You can also create scripts that use Active Directory Service Interfaces (ADSI). Several sample scripts are supplied on the operating system installation media. For more information about the sample scripts, see Using the Windows Deployment and Resource Kits. For more information about using ADSI, see Programming interfaces.

**Customizing how data is displayed in Active Directory administrative tools and snap-ins**
The Active Directory administrative tools, such as Active Directory Users and Computers, and the Windows shell extensions use display specifiers to dynamically create context menu items and property pages. Display specifiers permit localization of class and attribute names, context menus, and property pages, and also support new classes and attributes. You can add and modify classes and attributes in the schema and extend both the administrative tools and the Windows shell in many ways by modifying attributes in display specifiers. For more information the Active Directory schema and display specifiers, see the Active Directory programmer's Guide at the Microsoft Web site.

**Using Active Directory Users and Computers**
You can change how directory objects are displayed in Active Directory Users and Computers by selecting commands on the **View** menu of the console. Menu commands include the ability to toggle features on and off, such as the console tree, description bar, status bar, large icons, small icons, and so on.

When you start Active Directory Users and Computers and expand the domain node, several containers are displayed in the console tree. If you have just created a domain controller, the containers that are displayed by default are:

- **Builtin:** Contains objects that define the default built-in groups, such as Account Operators or Administrators.
- **Computers:** Contains Windows 2000, Windows XP, and Windows Server 2003 computer objects, including computer accounts that were originally created using application programming interfaces (APIs) that could not use Active Directory. Computer objects are moved to the Computer container when Windows NT domains are upgraded to Windows 2000 or a Windows Server 2003 operating system.
- **Domain Controllers:** Contains computer objects for domain controllers running Windows 2000 or Windows Server 2003.

- **Users:** Contains user accounts and groups that were originally created using APIs that could not use Active Directory. User accounts and groups are moved to the Users container when Windows NT domains are upgraded to Windows 2000 or a Windows Server 2003 operating system. You can use the Windows NT 4.0 User Manager (Usrmgr) tool to modify users and groups created using the APIs that could not use Active Directory.

When you select **Advanced Features** on the **View** menu, two additional folders are displayed in the console:

- **LostAndFound:** Contains objects whose containers were deleted at the same time that the object was created. If an object has been created in or moved to a location that is missing after replication, the lost object is added to the LostAndFound container. The LostAndFoundConfig container in the configuration directory partition serves the same purpose for forest-wide objects.

- **System:** Contains built-in system settings for the various system service containers and objects. For more information about the System container, see Using the Windows Deployment and Resource Kits.

When you select **Filter options** on the **View** menu, you can show all objects, show only selected objects, configure the number of items that can be displayed for each folder, or create custom filters using object attributes and LDAP queries.

**Starting Active Directory MMC consoles from the command-line**
Active Directory MMC consoles, including Active Directory Users and Computers (dsa.msc), Active Directory Domains and Trusts (domain.msc) and Active Directory Sites and Services (dssite.msc), provide command-line options that allow you to start a console focused on a particular domain or domain controller. The command-line options support both fully qualified domain names and NetBIOS names.

The command-line options are:

- **/domain=** *FullyQualifiedDomainName*

- **/domain=** *NetBIOSDomainName*

- **/server=** *FullyQualifiedDomainControllerName*

- **/server=** *NetBIOSDomainControllerName*

You can use these command-line options to run the Active Directory MMC consoles directly from the command line, or you can create a shortcut to start a console and add the appropriate command-line options to the shortcut. You can also use the command-line options with any custom consoles that you create. For more information about creating and saving console files, see Windows interface administrative tool reference A-Z: Microsoft Management Console.

Command-line examples:

- To start Active Directory Users and Computers focused on domain1, type:

  **dsa.msc /domain=***domain1*

- To start Active Directory Users and Computers focused on server1, type:

  **dsa.msc /server=***server1***.***domain1*

- To start Active Directory Sites and Services focused on server1, type:

  **dssite.msc /server=***server1***.***domain1*

- To start Active Directory Domains and Trusts focused on server1, type:

  **domain.msc /server=***server1***.***domain1*

**Notes**

- Do not use both a /domain and /server command-line option at the same time.
- The /domain options can only be used with Active Directory Users and Computers.

# Managing Active Directory from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Active Directory from the command line**

The following command-line tools can be used to manage Active Directory.

| Name | Description |
|---|---|
| CSVDE | Import and export Active Directory data using comma-separated format. |
| Dsadd | Add users, groups, computers, contacts, and organizational units to Active Directory. |
| Dsmod | Modify an existing object of a specific type in the directory. The types of objects that can be modified are: users, groups, computers, servers, contacts, and organizational units. |
| Dsrm | Remove objects of the specified type from Active Directory. |
| Dsmove | Rename an object without moving it in the directory tree, or move an object from its current location in the directory to a new location within a single domain controller. (For cross-domain moves, use the Movetree command-line tool.) |
| Dsquery | Query and find a list of objects in the directory using specified search criteria. Use in a generic mode to query for any type of object or in a specialized mode to query for for selected object types. The specific types of objects that can be queried through this command are: computers, contacts, subnets, groups, organizational units, sites, servers and users. |
| Dsget | Display selected attributes of specific object types in Active Directory. Attributes of the following object types can be viewed: computers, contacts, subnets, groups, organizational units, servers, sites, and users. |
| LDIFDE | Ceate, modify, and delete directory objects. This tool can also be used to extend the schema, export Active Directory user and group information to other applications or services, and populate Active Directory with data from other directory services. |
| Ntdsutil | General purpose Active Directory management tool. Use Ntdsutil to perform database maintenance of Active Directory, to manage single master operations, and remove metadata left behind by domain controllers that were removed from the network without being properly uninstalled. |

Additional command-line tools are available in the Support tools folder of the product compact disc and from the Resource Kit

For information about other command-line tools, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

# Finding directory information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Finding directory information**

Active Directory is designed to provide information to queries about directory objects from both users and programs. Administrators and users can easily search for and find information in the directory by using the Search command on the Start menu. Client programs can access information in Active Directory by using Active Directory Service Interfaces (ADSI).

One of the principal benefits of Active Directory is its rich store of information about network objects. Information published in Active Directory about users, computers, files, and printers is available to network users. This availability is controlled by security permissions to view information. For information about publishing information in the directory, see Publishing resources.

Everyday tasks on a network involve communication with other users and connection to published resources. These tasks require finding names and addresses to send mail or connect to shared resources. In this respect, Active Directory functions as a shared address book for the enterprise. For example, you can find a user by first name, last name, e-mail name, office location, or other properties of that person's user account. Finding information is optimized by use of the global catalog. For more information, see The role of the global catalog.

**Restricting directory information access**

In some cases, such as for security or privacy reasons, you may want to restrict access to certain directory information. Access control permissions provide you with detailed control over the visibility of information stored in Active Directory. Using permissions, you can ensure that only users who need particular directory information have access to it. For information on assigning permissions to Active Directory objects or properties, see Assign, change, or remove permissions on Active Directory objects or attributes. In addition, see Best practices for assigning permissions on Active Directory objects.

**Efficient search tools**

Administrators can use the advanced Find dialogs in Active Directory Users and Computers to perform management tasks with greater efficiency and to easily customize and filter data retrieved from the directory. For more information, see Search Active Directory.

Additionally, administrators can add objects to groups quickly and with minimal network impact by utilizing browse-less queries to help find likely members. For more information, see Add a member to a group.

# Administering other domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering other domains**
In an organization with more than one domain, it is often necessary to administer a domain other than the one to which you are currently logged on. For example, when creating trusts, the trust must be created in both the trusting and the trusted domain. You can do this through:

- Cooperation with those who have administrative credentials in another domain
- Logging on with a user account with the necessary permissions
- Using Run as to start an administrative tool targeted on the particular domain (recommended)

A secure method of controlling administrative access to a particular domain is to tightly control the number of accounts that are added to the Domain Admins group for that domain and the number of people aware of those accounts. Only members of the Domain Admins group (or Enterprise Admins group) can make broad administrative changes to the domain.

For example, if a domain administrator in one domain wanted to establish a shortcut trust with another domain, the domain administrator could establish the trust relationship in that domain by communicating with the domain administrator of the other domain, agreeing on a common strong password for the trust, and having the administrator of the other domain create the trust in that domain.

A more convenient, but less secure method of administering more than one domain is to logon interactively with a user account that has administrative credentials in both domains. For example, user accounts that are members of the Enterprise Admins security group have permission to administer every domain in the forest. Logging on to a computer with a user account that has broad administrative access is not recommended. For more information, see Why you should not run your computer as an administrator.

Run as, when used with the appropriate user name and password, can be used to securely administer domains in any forest in your organization. For more information, see Using Run as.

**Notes**

- It is a security best practice to never log on with an account that has more rights and permissions than you need for the task you are currently performing.
- It is also a security best practice to limit the scope of a particular administrator's rights and permissions to include only those tasks that the administrator commonly performs.
- You cannot administer a Windows Server 2003 domain or forest using the Windows 2000 Administrative Tools Pack. For more information, see Windows Server 2003 Administration Tools Pack Overview.

# Delegating administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Delegating administration**

By delegating administration, you can assign a range of administrative tasks to the appropriate users and groups. You can assign basic administrative tasks to regular users or groups, and leave domain-wide and forest-wide administration to members of the Domain Admins and Enterprise Admins groups. By delegating administration, you can allow groups within your organization to take more control of their local network resources. You also help secure your network from accidental or malicious damage by limiting the membership of administrator groups.

You can delegate administrative control to any level of a domain tree by creating organizational units within a domain and delegating administrative control for specific organizational units to particular users or groups.

To decide what organizational units you want to create, and which organizational units should contain accounts or shared resources, consider the structure of your organization.

For example, you might want to create an organizational unit that enables you to assign a user the administrative control for all user and computer accounts in all branches of a department, such as Human Resources. Or, you might want to assign a user administrative control only to some resources within a department, for example, computer accounts. Another possible delegation of administrative control would be to assign a user the administrative control for the Human Resources organizational unit, but not any organizational units contained within the Human Resources organizational unit.

Active Directory defines specific permissions and user rights that can be used for the purposes of delegating or restricting administrative control. Using a combination of organizational units, groups, and permissions, you can define the most appropriate administrative scope for a particular person, which could be an entire domain, all organizational units within a domain, or a single organizational unit.

Administrative control can be assigned to a user or group by using the Delegation of Control Wizard or through the Authorization Manager console. Both of these tools allow you to assign rights or permissions to particular users or groups.

For example, a user can be given permissions to modify the Owner Of Accounts property, without being assigned the permissions to delete accounts in that organizational unit. The Delegation of Control Wizard, as its name suggests, allows you to delegate administrative tasks using a wizard that steps you through the entire process. The Authorization Manager is a Microsoft Management Console (MMC) console that also allows you to delegate administration. The Authorization Manager provides greater flexibility than the Delegation of Control Wizard, at the cost of greater complexity.

For more information about using the Delegation of Control Wizard, see To delegate control. For more information about using Authorization Manager, see Authorization Manager.

**Delegating administration safely**

Delegate administration carefully and document all your delegated assignments. Before you delegate any tasks, ensure adequate training for users who will be assigned administrative control of objects. To review Active Directory security concepts, including permissions and inheritance, see Access control in Active Directory.

For more information about delegating administration safely, see Designing the Active Directory Logical Structure.

**Customizing MMC consoles for specific groups**

You can use Microsoft Management Console (MMC) options to create a limited-use version of a snap-in such as Active Directory Users and Computers. This allows administrators to control the options available to groups to whom you have delegated administrative responsibilities by restricting access to operations and areas within that customized console.

For example, suppose you delegate the Manage Printers right to the PrintManagers group in the Manufacturing organizational unit. To simplify administration, you can create a custom console for use by members of the PrintManagers group containing only the Manufacturing organizational unit and restrict the scope of the console using the console modes.

For more information about customizing a console and using console modes, see Using custom consoles and Console access options.

This type of delegation is also enhanced by the Group Policy settings available for MMC. These settings enable the administrator to establish which MMC snap-ins can be run by a particular user. The settings can be inclusive, allowing a set of snap-ins to run, or exclusive, restricting the set of snap-ins to run.

For more information about Group Policy settings for MMC, see Setting Group Policy in MMC.

**Using Group Policy to publish and assign customized consoles**
Using Group Policy, you can distribute a customized console to specific groups in one of two modes: publishing or assigning. Publishing a customized console advertises the console to the members of a group specified in the Group Policy setting by adding the console to the list of available programs in Add or Remove Programs. The next time the members of the group open Add or Remove Programs they have the option to install the new console. Assigning (as opposed to publishing) a console forces the console to be automatically installed for all specified accounts.

To publish or assign a console, create or modify a Group Policy object and then apply it to the appropriate group of users. Then, use the Software Installation extension of the Group Policy snap-in to either publish or assign the console.

For more information, see Group Policy (pre-GPMC).

**Important**

- The console must be packaged before using the Software Installation snap-in. You can use a tool such as Windows Installer to package the customized console. Once this has been accomplished you can configure the Software Installation snap-in to publish or assign the newly created package. For more information about how to package an application, see Using the Windows Deployment and Resource Kits.

- If the customized console you are packaging uses a snap-in that is not installed on the destination workstation or server for the published or assigned user, you will need to include the snap-in file and the registration of the file in the package. You can either create a separate package that contains the snap-in or add the snap-in during the creation of the customized console package so that it will be properly installed on the computer every time a user installs the console package.

**Note**

- If a user is logged on to their computer at the time that a Group Policy object is applied to their account, the user will not see the published or assigned console until they log off and then log on again.

For more information about other Active Directory security issues, see Security overview.

# Publishing resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Publishing resources**
To help users locate the network resources they need, you can publish searchable information about these resources in Active Directory. Resources that can be published include users, computers, printers, shared folders, and network services. Some commonly used directory information, like user and computer names, is published by default when the objects are created. Other directory information, such as information about shared folders, must be published manually. For information about searching the directory, see Finding directory information.

Using access control permissions, you can control which users and groups can search and view published information. Access control permissions give you detailed control over directory information at both the resource and the property levels. For example, you can use permissions to prevent a particular group from viewing any user information published in the directory. Or, you can use permissions to allow that group to view user names, but no other user information. For an overview of permissions, see Access control overview. For information about assigning permissions to Active Directory objects and properties, see Assign, change, or remove permissions on Active Directory objects or attributes.

**Publishing users and computers**
User and computer accounts are added to the directory using Active Directory Users and Computers and are automatically published to the directory upon creation. Commonly used account information, such as account names, is published by default. Other information, such as account security information, is only visible to administrators. For information about creating new accounts, see Create a new user account and Create a new computer account.

**Publishing shared printers**
You can also publish information about shared printers in Active Directory. Information about printers shared from Windows NT must be published manually. Information about printers shared from the Windows Server 2003 family or the Windows 2000 Server family is published to the directory automatically when you create a shared printer. You use Active Directory Users and Computers to manually publish shared printer information. For more information about publishing printers, see Manually publish a printer in Active Directory.

**Publishing shared folders**
To help users find shared folders more easily, you can publish information about shared folders in Active Directory. You use Active Directory Users and Computers to manually publish shared folder information. For more information about publishing shared folders, see Publish a shared folder.

**Publishing services**
A service is an application that makes data or operations available to network users. Publishing a service in Active Directory enables users and administrators to move from a machine-centric view of the network to a service-centric view. By publishing a service, rather than computers or servers, administrators can focus on managing the service regardless of which computer is providing the service or where the computer is located.

Some services, such as Certificate Services, are automatically published in Active Directory when they are installed. Other services can be published in the directory using programming interfaces. For more information, see Programming interfaces. Administrators can manage published services using Active Directory Sites and Services. For more information about services and how to publish them, see the Service Publication page at the Microsoft Web site. (http://msdn.microsoft.com/)

**Categories of service information**
Binding and configuration information are the two types of service information frequently published using Active Directory:

- Binding information allows clients to connect to services that do not have well-known bindings and that conform to a service-centric model. By publishing the bindings for these kinds of services, connections can be automatically established with services. Machine-centric services are typically handled on a service-by-service basis and should not be published to the directory.

- Configuration information can be common across client applications. Publishing this sort of information allows you to distribute current configuration information for these applications to all clients in the domain. The configuration information is accessed by client applications as needed. This eases application configuration for users and gives you more control over application behaviors.

**Characteristics of service information**

Service information that you publish to the directory is most effective if it has the following characteristics:

- **Useful to many clients.** Information that is useful to a small set of clients or that is useful only in certain areas of the network should not be published. If not widely used, this information wastes network resources, since it is published to every domain controller in the domain.

- **Relatively stable and unchanging.** Although there may be exceptions to this rule, it generally makes sense to publish only service information that changes less frequently than two replication intervals. For intrasite replication, the maximum replication period is fifteen minutes, and for intersite replication, the maximum replication period is configured based on the replication interval of the site link used for the replication. Object properties that change more frequently create excessive demands on network resources. Property values may be out of date until updates are published, which can take as long as the maximum replication period. Consequently, having properties out of date for that period of time must not create unacceptable conditions.

  For example, some network services select a valid TCP port for use each time they are started. After selecting the port, the service updates Active Directory with this information, which is stored as the service connection point. Clients access the service connection point when they want to use the service, but if the new service connection point has not been replicated when the client requests it, the client will receive an outdated port, rendering the service temporarily inaccessible.

- **Well-defined, reasonable properties.** Information that is of a consistent form is easier for services to use. The information should be relatively small in size.

# Managing COM+ partitions in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing COM+ partitions in Active Directory**
COM+ partitions stored in Active Directory are used to map a local COM+ partition, which stores the actual COM+ *application*, to specific users or organizational units within your enterprise. COM+ applications are groups of COM components developed to work together to make use of COM+ services such as queuing, role-based security, and so on.

There are two types of COM+ partitions: COM+ partitions stored in Active Directory and local COM+ partitions stored on application servers. Using COM+ partitions stored in Active Directory, you can assign domain users and entire organizational units to applications stored in local COM+ partitions. Local COM+ partitions are application containers used to manage multiple instances of COM+ applications on a single application server.

A local COM+ partition can only store one instance of an application. In other words, if you needed to make two or more versions of the same application (AppX 1.0 and AppX 2.0) available to domain users, then you would need to create two separate local COM+ partitions on an application server and associate (or link) them with two separate COM+ partitions in Active Directory. For more information about associating COM+ partitions, see Managing COM+ partition sets in Active Directory.

Each COM+ partition is configured and managed separately, according to the specific needs of its users.

**Note**

- COM+ partitions in Active Directory are not available from domain controllers running Windows 2000.

# Managing COM+ partition sets in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing COM+ partition sets in Active Directory**
COM+ partition *sets* stored in Active Directory can contain one or more COM+ partitions and are used to assign one or more applications, residing on an application server, to domain users or organizational units. For more information about COM+ partitions, see Managing COM+ partitions in Active Directory.

Each COM+ partition set defines the COM+ partitions for which a domain user is permitted to access.

COM+ partition sets provide the following advantages to both administrators and application programmers:

- Administering distributed applications becomes much easier because you can tailor access of a particular set of domain users to a specific set of applications.

- Security policies can be applied to domain users and organizational units within each partition set.

Before assigning COM+ partition sets, you must first logically group one or more COM+ partitions into a single COM+ partition set. Once a COM+ partition set has been defined, you can make applications available throughout the domain by mapping COM+ partition sets to domain users or organizational units. To do this, you must perform tasks on both the domain controller where Active Directory resides and on the application server where the COM+ application is installed.

**Tasks to perform on a domain controller**
You must first create a COM+ partition within Active Directory by using Active Directory Users and Computers or, programmatically, by using Active Directory Service Interfaces (ADSI). For more information, see Create a COM+ partition in Active Directory. Once you have created and configured your COM+ partitions in Active Directory, you can use Component Services on an application server to associate a local COM+ partition with a COM+ partition stored in Active Directory.

After the COM+ partition is created and you have associated it with a local COM+ partition, you need to create a COM+ partition set. When creating a COM+ partition set, you can define the COM+ partitions that comprise that set. Creating a COM+ partition set is done either by using Active Directory Users and Computers or, programmatically, by using Active Directory Service Interfaces (ADSI). For more information, see Create a COM+ partition set in Active Directory.

Finally, you map domain users or organizational units to the newly created COM+ partition set. You can associate multiple users or organizational units with a COM+ partition set at one time instead of having to map multiple user identities or organizational units. For more information, see Map a user or organizational unit to a COM+ partition set.

**Tasks to perform on Application servers**
Using Component Services you can link a single COM+ application stored locally on an application server to a single COM+ partition stored in Active Directory. You can do this by creating a local COM+ partition on the application server, through Component Services, and changing the local COM+ partition ID to reflect the COM+ partition ID located in Active Directory.

Each COM+ partition in Active Directory has a general description, a unique name, and a unique partition ID. The partition ID is what associates the local COM+ partition to the COM+ partition defined in Active Directory. The COM+ partitions defined in Active Directory are intended to be unique across an enterprise. A local COM+ partition is necessary in order to store the COM+ application on an application server, which may also have multiple versions of a specific application installed.

You can define a local COM+ partition ID on the application server and then later revise the partition ID in Active Directory to match it. Or, you can define a COM+ partition ID in Active Directory and then later revise the partition ID on the application server to match the one in Active Directory.

**Note**

- Local COM+ partitions are sometimes referred to as Empty Partitions in the Component Services user interface.

# Active Directory Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This section covers:

- Web resources
- Active Directory support tools
- Programming interfaces

# Web resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Web resources**

The following items provide links to more information on the Web related to Active Directory:

- For information about deploying Active Directory, see "Identifying an Active Directory Deployment Strategy" at the Microsoft Windows Resource Kits Web site.
- For information about domain trust, see "Domain Trust" at the Microsoft Windows Resource Kits Web site.
- For information about establishing domain trusts, see "Designing an Authorization Strategy" at the Microsoft Windows Resource Kits Web site.
- For information about monitoring Active Directory, see "Monitoring Active Directory" at the Microsoft Windows Resource Kits Web site.
- For information about delegating enterprise administrator operations, see "Delegating Enterprise Administrator Operations" at the Microsoft Windows Resource Kits Web site.
- For development information about Active Directory, see the Active Directory Programmer's Guide at the Microsoft Web site.
- For development information about Lightweight Directory Access Protocol (LDAP), see the Lightweight Directory Access Protocol API at the Microsoft Web site.
- For information about obtaining valid object identifiers, see Active Directory Naming Registration at the Microsoft Web site.
- For information about Request for Comments (RFCs) and Internet drafts, see the Internet Engineering Task Force Web site.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Active Directory support tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory support tools**

Several additional tools that can be used to configure, manage, and debug Active Directory are available as command-line tools. These tools are known as the Support Tools and are available on the installation CD in the \Support\Tools folder.

**Note**

You can find downloadable versions of the Support Tools at the following locations:

- Windows Server 2003 Service Pack 1 32-bit Support Tools (http://go.microsoft.com/fwlink/?LinkId=70775)
- Windows Server 2003 Service Pack 2 32-bit Support Tools (http://go.microsoft.com/fwlink/?LinkId=168625)

**List and description of tools**

In addition, the Active Directory Migration Tool (ADMT) is available to help you migrate user accounts, groups, and computer accounts from Windows NT 4.0 domains to Active Directory domains. The Active Directory Migration Tool is a Microsoft Management Console (MMC) snap-in and is available on the installation compact disk in the \i386\ADMT folder.

| Tool | Description |
|---|---|
| Movetree | Move objects from one domain to another. |
| SIDWalk | Set the access control lists on objects previously owned by accounts that were moved, orphaned, or deleted. |
| LDP | Allows LDAP operations to be performed against Active Directory. This tool has a graphical user interface (GUI). |
| Dnscmd | Enables administrator to check presence of domain controller locator records in DNS, add or delete such records and perform configuration of DNS servers, zones and records. |
| DSACLS | View or modify the access control lists of directory objects. |
| Netdom | Batch management of trusts, joining computers to domains, verifying trusts and secure channels. |
| NETDiag | Check end to end network and distributed services functions. |
| NLTest | Check that the locator and secure channel are functioning. |
| Repadmin | Check replication consistency between replication partners, monitor replication status, display replication metadata, force replication events and knowledge consistency checker recalculation. |
| Replmon | Display replication topology, monitor replication status (including group policies), force replication events and knowledge consistency checker recalculation. This tool has a graphical user interface (GUI). |
| DSAStat | Compare directory information on domain controllers and detect differences. |
| ADSI Edit | A Microsoft Management Console (MMC) snap-in used to view all objects in the directory (including schema and configuration information), modify objects and set access control lists on objects. |
| SDCheck | Check access control list propagation and replication for specified objects in the directory. This tool enables an administrator to determine if access control lists are being inherited correctly and if access control list changes are being replicated from one domain controller to another. |
| ACLDiag | Determine whether a user has been assigned or denied access to a directory object. It can also be used to reset access control lists to their default state. |

| | |
|---|---|
| DFSUtil | Command-line utility for managing all aspects of Distributed File System (DFS), checking the configuration concurrency of DFS servers, and displaying the DFS topology. |
| Dcdiag | Analyzes the state of domain controllers in a forest or enterprise and reports any problems to assist in troubleshooting. |
| Active Directory Migration Tool (ADMT) | A Microsoft Management Console (MMC) snap-in used to migrate user accounts, groups, and computer accounts from Windows NT 4.0 domains to Active Directory domains (available on the installation compact disk in the \i386\ADMT folder). |

For more information, see Install Windows Support Tools and Using the Windows Deployment and Resource Kits.

# Programming interfaces

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Programming interfaces**

Active Directory Service Interfaces (ADSI) provides a simple, powerful, object-oriented interface to Active Directory. ADSI makes it easy for programmers and administrators to create directory programs by using high-level tools, such as Microsoft Visual Basic, without having to worry about the underlying differences between the different namespaces.

ADSI enables you to build or buy programs that give you a single point of access to multiple directories in your network environment, whether those directories are based on LDAP or another protocol. ADSI is fully scriptable for ease of use by administrators.

For more information about ADSI, see the Active Directory Programmer's Guide at the Microsoft Web site and the Using the Windows Deployment and Resource Kits.

Active Directory also provides support for Messaging Application Programming Interface (MAPI), so legacy MAPI programs will continue to work with Active Directory.

In addition, Active Directory supports the LDAP C API as a lower-level interface for C programmers. For more information, see the Lightweight Directory Access Protocol (LDAP) API at the Microsoft Web site.

# Active Directory Troubleshooting

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

- Troubleshooting Active Directory
- Troubleshooting replication
- Troubleshooting schema
- Troubleshooting trusts

# Troubleshooting Active Directory

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting Active Directory**
What problem are you having?

- Cannot add or remove a domain.
- Cannot create security principals in Active Directory.
- Changes to group memberships are not taking effect.
- Clients without Active Directory client software installed fail to authenticate.
- Receiving "Domain not found," "Server not available," or "RPC server is unavailable" error messages.
- User is unable to log on locally to a domain controller.
- Cannot enable auditing.
- Cannot remove Active Directory from a domain controller.
- Cannot connect to a domain controller running Windows 2000.

**Cannot add or remove a domain.**
**Cause:** The domain naming master is not available. This may be caused by a network connectivity problem or an Active Directory Installation Wizard failure. It may also be due to a failure of the computer holding the domain naming master role. Or, the user who is attempting to add or remove the domain does not have the necessary administrative credentials.

**Solution:** Identify the computer holding the domain naming master role by using the command **netdom query fsmo** and repair or replace the domain naming master computer. It may be necessary to seize the domain naming master role. Or, resolve the network connectivity problem. If this does not help solve the issue, see article Q223787, "Flexible Single Master Operation Transfer and Seizure Process," in the Microsoft Knowledge Base.

**See also:** Diagnose Connections; Operations master roles; Responding to operations master failures; Transfer the domain naming master role; Seize the domain naming master role

**Cannot create security principals in Active Directory.**
**Cause:** The RID master is not available or failed to replicate. This may be caused by a network connectivity problem or may be due to a failure of the computer holding the RID master role. This behavior can also occur when the **Access this computer from the network** user right is not assigned to the appropriate groups on the RID master. Or, the user who is attempting to create the security principal does not have the necessary administrative credentials.

**Solution:** Identify the computer holding the RID master role by using the command **netdom query fsmo** and repair or replace the computer holding the RID master role. It may be necessary to seize the RID master role. Or, resolve the network connectivity problem. If this does not help solve the issue, see articles Q223787, "Flexible Single Master Operation Transfer and Seizure Process," and Q248410 "Err Msg: The Account-Identifier Allocator Failed To..." in the Microsoft Knowledge Base.

**See also:** Access this computer from the network; Diagnose Connections; Operations master roles; Responding to operations master failures; Transfer the RID master role; Seize the RID master role

**Changes to group memberships are not taking effect.**
**Cause:** The infrastructure master is not available. This may be caused by a network connectivity problem. It may also be due to a failure of the computer holding the infrastructure master role. Or, the user who is attempting to change group membership does not have the necessary administrative credentials.

**Solution:** Identify the computer holding the infrastructure master role by using the command **netdom query fsmo** and

repair or replace the computer holding the infrastructure master role. It may be necessary to seize the infrastructure master role.Or, resolve the network connectivity problem. If this does not help solve the issue, see article Q223787, "Flexible Single Master Operation Transfer and Seizure Process," in the Microsoft Knowledge Base.

**See also:** Diagnose Connections; Operations master roles; Responding to operations master failures; Transfer the infrastructure master role; Seize the infrastructure master role

**Clients without Active Directory client software installed fail to authenticate.**
**Cause:** The end user's password has expired and the primary domain controller (PDC) emulator master is not available. This may be caused by a network connectivity problem. It may also be due to a failure of the computer holding the PDC emulator master role.

**Solution:** Identify the computer holding the PDC emulator master role by using the command **netdom query fsmo** and repair or replace the computer holding the PDC emulator master role. It may be necessary to seize the PDC emulator master role. Or, resolve the network connectivity problem. If this does not help solve the issue, see articles Q223787, "Flexible Single Master Operation Transfer and Seizure Process," and Q239869, "How to Enable NTLM 2 Authentication for Windows 95, Windows 98, Windows NT Server 4.0, and Windows 2000 Advanced Server," in the Microsoft Knowledge Base.

**See also:** Diagnose Connections; Operations master roles; Responding to operations master failures; Transfer the PDC emulator role; Seize the PDC emulator role

**Receiving "Domain not found," "Server not available," or "RPC server is unavailable" error messages.**
**Cause:** Name registration or name resolution problem.

**Solution:** Verify that Domain Name System (DNS) is available and functioning correctly. Run the **Netdiag /debug** command on the server in question. This will evaluate the registration of NetBIOS, DNS, and services. If this does not help solve the issue, see article Q265706, "DCDiag/NetDiag Facilitate Join and DC Creation," in the Microsoft Knowledge Base.

**See also:** Active Directory support tools; Install Windows Support Tools

**User is unable to log on locally to a domain controller.**
**Cause:** The ability to log on locally to a domain controller is controlled by security policies, which are established in Group Policy settings.

**Solution:** In the default Domain Controller Policy object, assign the particular user or group the **Allow log on locally** user right. If this does not help solve the issue, see article Q234237, "Assign "Log On locally" Rights to Windows Domain Controller," in the Microsoft Knowledge Base.

**See also:** Allow log on locally; Logon rights; Permit users to log on locally to a domain controller

**Cannot enable auditing.**
**Cause:** For servers running Windows Server 2003, auditing must be enabled in the Group Policy object. Or, the user who is attempting to enable auditing does not have the necessary administrative credentials.

**Solution:** Run GPEdit.msc to enable auditing or monitoring of system events.

**See also:** Define or modify auditing policy settings for an event category

**Cannot remove Active Directory from a domain controller.**
**Cause:** The NTDS Settings object could not be properly removed. Or, the user who is attempting to remove the domain controller does not have the necessary administrative credentials.

**Solution:** Use the Ntdsutil command-line tool to manually remove the NTDS Settings object. If this does not help solve the issue, see article Q216498, "Remove Data in Active Directory After an Unsuccessful Demotion," in the Microsoft Knowledge Base.

**See also:** Delete extinct server metadata

**Cannot connect to a domain controller running Windows 2000.**
**Cause:** You are trying to connect to a domain controller running Windows 2000 that does not have Service Pack 3 or later installed.

**Solution:** Upgrade domain controllers running Windows 2000 to Service Pack 3 or later.

**See also:** Connecting to domain controllers running Windows 2000

# Troubleshooting replication

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting replication**

For more detailed replication troubleshooting information than is available here, and for additional information about functionality in the version of Dcdiag.exe that is included in Windows Support Tools that ship with Windows Server 2003 with Service Pack 1 (SP1), see Troubleshooting Active Directory Replication Problems on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=60980).

What problem are you having?

- Monitoring replication.
- Replication between sites is slow.
- Received Event ID 1311 in the directory service log.
- Received Event ID 1265 with error "DNS Lookup Failure," or "RPC server is unavailable" in the directory service log. Or, received "DNS Lookup Failure" or "Target account name is incorrect" from the repadmin command.
- Received Event ID 1265 "Access denied," in directory service log. Or, received "Access denied" from the repadmin command.
- Received "Access denied" from Active Directory Sites and Services when manual replication was attempted.
- Unable to connect to a domain controller running Windows 2000 from the Active Directory Sites and Services snap-in.
- Search for new and updated information about replication. Or, your question does not match any of those listed above.

**Monitoring replication.**

- **Cause:** You should monitor replication regularly to help you identify and fix problems before they grow.
- **Solution:** Regular monitoring is the key to good replication maintenance. Repadmin.exe and dcdiag.exe (both part of the Windows Support Tools) and the directory service event log (accessible through the Event Viewer) are the primary tools for monitoring replication.

  Repadmin is a command-line tool that report failures on a replication link between two replication partners. The following **repadmin** example displays the replication partners and any replication link failures for Server1 on the microsoft.com domain:

  **repadmin /showrepl** *server1***.microsoft.com**

  For a complete list of **repadmin** options, use the **?** option:

  **repadmin /?**

  Dcdiag is a command-line tool that can check the DNS registration of a domain controller, check to see that the security descriptors (SIDs) on the naming context heads have appropriate permissions for replication, analyze the state of domain controllers in a forest or enterprise, and more. The following **dcdiag** example checks for any replication errors between domain controllers:

  **dcdiag /test:replications**

  For a complete list of dcdiag options, use the **?** option:

**dcdiag /?**

The directory service log reports replication errors that occur after a replication link has been established. For information about viewing the directory service log, see View an event log.

Large enterprises may also want to use the Microsoft Operations Manager for automated monitoring of large numbers of domain controllers. For more information, see Active Directory Management Pack Technical Reference for Microsoft Operations Manager 2005 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=38341).

**See also:** Event Viewer; Install Windows Support Tools; Technical support options

**Replication between sites is slow.**

- **Cause:** The time required to replicate directory data between domain controllers is known as the replication latency. Replication latency can vary greatly, depending on the number of domain controllers, the number of sites, the available bandwidth between sites, replication frequency, and more.

- **Solution:**

  - Monitoring replication regularly is a good way to determine the normal replication latency on your network. With this knowledge, you can more easily determine if a problem is occurring. For more information, see the "Monitoring Replication" troubleshooting topic above.

  - Review the directory service log for any recent replication errors. Also, run **repadmin /showrepl** and review any resulting errors.

  - A good site topology design is important for replication efficiency. For information about site topology design guidelines, see When to establish a single or separate sites and Designing the Site Topology on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4724).

  - A number of algorithm enhancements have been made to replication in the Windows Server 2003 operating systems to improve replication efficiency and scalability. Some of these enhancements take effect in a forest set to Windows 2000 functional level, while others require the Windows Server 2003 functional level. You will gain the greatest improvement from these enhancements by upgrading your forest to Windows Server 2003 functional level. Adlb.exe is a tool that can help improve replication efficiency even further in forests set to the Windows Server 2003 functional level. For more information about Adlb, see the Windows Server 2003 Active Directory Branch Office Planning and Deployment Guide on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=28523). For more information about forest functionality, see Domain and forest functionality.

**See also:** Replication overview; Replication between sites; Managing replication; Bandwidth; Checklist: Optimizing intersite replication

**Received Event ID 1311 in the directory service log.**

- **Cause:** The replication configuration information in Active Directory Sites and Services does not accurately reflect the physical topology of the network.

  Common causes of Event ID 1311 include:

  - One or more domain controllers are offline.

  - Bridgehead servers are online but experiencing errors replicating a required naming context between Active Directory sites.

  - Preferred bridgehead servers defined by administrators are online but do not host the required naming contexts.

- One or more sites are not contained in site links.
- Site links contain all sites but the site links are not all site links are interconnected.
- Preferred bridgeheads defined by the administrator are offline.

- **Solution:** To resolve an error in the configuration of replication:
  - Make sure all sites belong to at least one site link. For more information, see Add a site to a site link.
  - Make sure that the combination of site links you have created allows a path between all domain controllers containing a replica of a given directory partition. For example, if a directory partition is held by domain controllers in both Site A and Site C, make sure that Site A and Site C belong to a common site link, or that an intermediary site exists that has at least one site link in common with Site A and at least one site link in common with Site B.
  - Make sure that you have cleared the **Bridge all site links** check box in Active Directory Sites and Services if your network is not fully routed. Or, if your network is fully routed and you have cleared the **Bridge all site links** check box, you may need to select it again to allow full replication of a directory partition. For more information, see Enable or disable site link bridges.
  - If you have manually assigned preferred bridgehead servers, make sure these servers are not offline. (It is generally recommended that you allow Active Directory to select bridgehead servers automatically.)
  - Use Ping.exe and Network Monitor to verify connectivity through WAN links and across routers. For more information about Network Monitor, see Network Monitor overview.
  - You can also search the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4441) for new and updated information about Event ID 1311.

**See also:** Create a site link; Add a site to a site link; Enable or disable site link bridges; Install Windows Support Tools

**Received Event ID 1265 with error "DNS Lookup Failure," or "RPC server is unavailable" in the directory service log. Or, received "DNS Lookup Failure" or "Target account name is incorrect" from the repadmin command.**

- **Cause:** These messages are often the result of DNS problems. Active Directory replication depends on the following:
  - Each domain controller in the forest must register its CNAME record for the name *DsaGuid*._msdcs.*ForestName*. DsaGuid is the GUID of the NTDS Settings object of the domain controller (visible in Active Directory Sites and Services as the **DNS alias** property of the server object's NTDS settings). This record usually belongs to the _msdcs.*ForestName* zone or, if that zone does not exist, the *ForestName* zone.
  - Each advertising domain controller in the forest must register its A record in the appropriate zone for each domain in the forest.
  - The A record must map to the current IP address of the respective domain controller.
  - The records must have replicated to the DNS servers used by direct replication partners.
  - Each DNS zone must have the proper delegations to the child zones.
  - The IP configuration of the domain controllers must contain correct preferred and alternate DNS servers.

  DNS errors that are reported by the directory service log or by **repadmin /showrepl** mean that the destination domain controller could not resolve the GUID-based DNS name of its source replication partner.

- **Solution:** Do the following:

1. Verify CNAME and A records. At a command prompt, type the following:

   **dcdiag /test:connectivity**

2. If the CNAME and A records are missing, restart netlogon. At a command prompt, type the following:

   **net start netlogon**

3. Again, verify CNAME and A records, by repeating step 1.

4. If the records are still missing, verify IP configuration. Verify that the preferred and alternate DNS servers specified in the IP configuration of the source and destination domain controllers are correct.

5. If the client is configured correctly, verify that the zone is dynamic. At a command prompt, type the following:

   **dcdiag /test:registerindns /dnsdomain**

6. To verify that name resolution is the cause of the problem, ping the GUID-based name of the domain controller where replication failed. If it works, the next replication cycle should not return this error.

7. If the ping fails, further DNS troubleshooting is required. For more information, see Troubleshooting Domain Name System on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=62177).

**See also:** Nslookup; Ping; Troubleshooting DNS; Install Windows Support Tools

**Received Event ID 1265 "Access denied," in directory service log. Or, received "Access denied" from the repadmin command.**

- **Cause:** This error can occur if the local domain controller failed to authenticate against its replication partner when creating the replication link or when trying to replicate over an existing link. This typically happens when the domain controller has been disconnected from the rest of the network for a long time and its computer account password is not synchronized with its computer account password stored in the directory of its replication partner.

- **Solution:** Do the following:

1. Stop the Key Distribution Center (KDC) service using **net stop KDC**.

2. Purge the ticket cache on the local domain controller.

3. Reset the domain controller's account password on the primary domain controller (PDC) emulator master using **netdom /resetpwd**. (Netdom.exe is available in Windows Support Tools).

4. Synchronize the domain directory partition of the replication partner with the PDC emulator master

5. Manually force replication between the replication partner and the PDC emulator master.

6. Start the KDC on the local domain controller:

   **net start KDC**

**See also:** User and computer accounts; Net start; Install Windows Support Tools

**Received "Access denied" from Active Directory Sites and Services when manual replication was attempted.**

- **Cause:** Using Active Directory Sites and Services to force replication initiates replication on all common directory partitions between the replication partners. However, a user can only force manual replication for containers on which

they have been assigned the **Replication Synchronization** permission. The replication of other directory partitions will fail, causing the "Access Denied" error.

- **Solution:** The **repadmin** or **replmon** command-line tools from Windows Support Tools can be used to manually force the replication of a specific directory partition.

  Replication synchronization is a special permission. For more information about special permissions, see Set, view, change, or remove special permissions and Active Directory object permissions.

**See also:** Install Windows Support Tools; Force replication over a connection; Active Directory support tools

**Unable to connect to a domain controller running Windows 2000 from the Active Directory Sites and Services snap-in.**

- **Cause:** You are trying to connect to a domain controller running Windows 2000 that does not have Service Pack 3 or later installed.

- **Solution:** Upgrade domain controllers running Windows 2000 to Service Pack 3 or later.

**See also:** Connecting to domain controllers running Windows 2000; Managing Active Directory from MMC

**Search for new and updated information about replication. Or, your question does not match any of those listed above.**

- **Cause:** New and updated information is regularly published on the Microsoft Web site.

- **Solution:** Visit the following links for the latest information:

  - Searching the Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4441)

    Search the Microsoft Knowledge Base of technical support information and self-help tools for Microsoft products.

  - Product Support Services on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=281)

    Search FAQs by product, browse the product support newsgroups, and contact Microsoft Support.

  - Active Directory Collection in the Windows Server 2003 Technical Reference on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=4549)

    View detailed technical information about Active Directory replication and other technologies, including troubleshooting Active Directory replication.

  - Microsoft Windows Server TechCenter on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=34403)

    Search for troubleshooting information, service packs, patches, and downloads for your system. View the Technical Library for the latest product information, including deployment, operations, and technical reference.

  - Windows Server Community on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=832)

    The official online community for enthusiasts of the Windows server operating systems.

**See also:** Technical support options; Install Windows Support Tools; Using the Windows Deployment and Resource Kits

# Troubleshooting schema

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting schema**
What problem are you having?

- Unable to modify or extend the schema.
- Unable to add attributes to a class.
- Unable to connect to a domain controller running Windows 2000 from the Active Directory Schema snap-in.
- Unable to find or run the Active Directory Schema snap-in.
- Search for new and updated information. Or, your question does not match any of those listed above.

**Unable to modify or extend the schema.**
**Cause:** The schema operations master is not available. The domain controller that currently holds the schema master role is offline or network connectivity between you and the schema master is inoperable.

**Solution:** Restore the availability of the schema master by fixing the network connectivity problem or by transferring or seizing the schema master role.

**See also:** Diagnose Connections; Operations master roles; Responding to operations master failures; Transfer the schema master role; Seize the schema master role

**Cause:** The user trying to modify the schema does not belong to the Schema Admins group.

**Solution:** Add the correct user account to the Schema Admins group.

**See also:** Add a member to the Schema Admins group

**Cause:** There are syntax errors in the LDAP import format (LDIF) script.

**Solution:** Debug and fix the LDIF script then try the modification again.

**See also:** The Active Directory Programmer's Guide at the Microsoft Web site; Internet Engineering Task Force Web site

**Cause:** The domain controller holding the schema master role is running Windows 2000 and has not been enabled for schema changes.

**Solution:** Enable schema changes on the domain controller. For more information, see Enable schema changes on a schema master running Windows 2000.

**See also:** Domain and forest functionality; Operations master roles; The Active Directory Programmer's Guide at the Microsoft Web site

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here..

**Unable to add attributes to a class.**
**Cause:** The attributes cannot be associated with a class because the schema cache has not been updated.

**Solution:** Ensure that the schema cache is updated after an attribute is added before attempting to add the attribute to a class. For more information, see Reload the schema.

**Cause:** You are attempting to define a class on a domain controller that is not the schema operations master and the new attribute has not replicated in.

**Solution:** Always perform schema updates on the schema operations master.

**See also:** Using the Windows Deployment and Resource Kits; Extending the schema; The Active Directory Programmer's Guide at the Microsoft Web site.

**Unable to connect to a domain controller running Windows 2000 from the Active Directory Schema snap-in.**
**Cause:** You are trying to connect to a domain controller running Windows 2000 that does not have Service Pack 3 or later installed.

**Solution:** Upgrade domain controllers running Windows 2000 to Service Pack 3 or later.

**See also:** Connecting to domain controllers running Windows 2000

**Unable to find or run the Active Directory Schema snap-in.**
**Cause:** The Active Directory Schema snap-in has not been installed or registered.

**Solution:** For instructions on installing and registering the Active Directory Schema snap-in, see Install the Active Directory Schema snap-in.

**See also:** Managing Active Directory from MMC

**Search for new and updated information. Or, your question does not match any of those listed above.**
**Cause:** You can find new and updated information on the Microsoft Web site.

**Solution:** The Microsoft Web site contains the latest product support information. For information about where to look, see Using online resources.

For the latest Active Directory development information, see the Active Directory Programmer's Guide at the Microsoft Web site.

**See also:** Managing Active Directory from MMC; Using the Windows Deployment and Resource Kits; Technical support options

# Troubleshooting trusts

Updated: May 1, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting Trusts**
What problem are you having?

- Clients are unable to access resources in a domain outside of the forest.
- Trust errors between servers or workstations.
- Trust errors between Windows NT 4.0 and Active Directory domains.
- After upgrading a Windows NT 4.0 domain with existing trusts to Active Directory domains, you encounter various trust-related problems.
- Cannot connect to a domain controller running Windows 2000.

**Clients are unable to access resources in a domain outside of the forest.**
**Cause:** A failure has occurred on the external trust between the domains.

**Solution:** Reset and verify the trust between the domains. The PDC emulator master must be available for a trust to be successfully reset.

**See also:** Verify a trust; Operations master roles; When to create an external trust

**Trust errors between servers or workstations.**
**Cause:** Incorrect time synchronization between domain controllers or workstations, the server could be down, or the trust relationship could be broken.

**Solution:** Run Netdom to verify, reset, or establish the trust between computers. This command-line tool performs batch management of trusts, verifies trusts and secures channels between computers, and can join computers to domains.

**See also:** Install Windows Support Tools

**Trust errors between Windows NT 4.0 and Active Directory domains.**
**Cause:** Automatic trust password resets for the trust may not reach the PDC emulator master role holder.

**Solution:** Run Netdom to verify, reset, or establish trust between computers. This command-line tool performs batch management of trusts, verifies trusts and secures channels between computers, and can join computers to domains. If this does not help solve the issue, see article Q317178, "Windows NT 4.0 Domain Updates Trust Account Password on Non-PDC," in the Microsoft Knowledge Base.

**After upgrading a Windows NT 4.0 domain with existing trusts to Active Directory domains, you encounter various trust-related problems.**
**Cause:** When the domain has been upgraded, the existing trusts to Active Directory domains remain Windows NT 4.0 trusts. Internet Protocol Security (IPSec) cannot work over a Windows NT 4.0 trust. Or, trusts to other domains in the forest are no longer available.

**Solution:** After upgrading a Windows NT 4.0 domain to an Active Directory domain, it is recommended that you delete and recreate all previously existing trusts with Active Directory domains. If this does not solve the issue, see article Q275221, "Trusts Unavailable on Backup Domain Controllers After Upgrading the Windows NT Primary Domain Controller," in the Microsoft Knowledge Base.

**See also:** Install Windows Support Tools; Upgrading from a Windows NT domain

**Cannot connect to a domain controller running Windows 2000.**
**Cause:** You are trying to connect to a domain controller running Windows 2000 that does not have Service Pack 3 or later installed.

**Solution:** Upgrade domain controllers running Windows 2000 to Service Pack 3 or later.

**See also:** Connecting to domain controllers running Windows 2000

# Network Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network services**
This section covers:

- Managing Core Network Services
- Managing Remote Connections
- Configuring Communication Links
- Configuring Communication Services
- Managing Bandwidth
- Integrating with Other Networks

# Managing Core Network Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing core network services**
This section covers:

- TCP/IP Core Networking Service
- IP Version 6
- Internet Protocol Security (IPSec)
- DNS
- DHCP
- WINS

# TCP/IP Core Networking Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP**
Transmission Control Protocol/Internet Protocol (TCP/IP) is a suite of standard protocols for connecting computers and building networks. TCP/IP is network protocol software provided with Microsoft® Windows® operating systems that implements and supports the TCP/IP protocol suite:

- Before you begin using TCP/IP with this computer, see Checklist: Configuring TCP/IP.
- To find features that have been moved in the Windows Server 2003 family, see New ways to do familiar tasks with TCP/IP.
- For tips about using TCP/IP, see Best practices for TCP/IP.
- For help with specific tasks, see TCP/IP How To....
- For general background information, see TCP/IP Concepts.
- For problem-solving instructions, see Troubleshooting TCP/IP.

The Windows Server 2003 family includes the latest version of Internet Protocol (IP), known as IP version 6. For more information, see IP Version 6.

# Checklist: Configuring TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring TCP/IP**

| Step | Reference |
|---|---|
| Review TCP/IP concepts such as IP addressing, subnet masks, and default gateways. | Understanding TCP/IP |
| Review TCP/IP name resolution concepts and options for resolving NetBIOS and DNS names to IP addresses. | Name Resolution |
| Determine how computers on your network will obtain IP addresses. | DHCP |
| Determine how computers on your network will register and resolve NetBIOS names. | WINS |
| Determine how computers on your network will resolve and register DNS names. | DNS |
| Determine how this computer will be used with your TCP/IP network. | Configuring TCP/IP |
| Verify that network adapter hardware is installed, or install hardware if needed. | Hardware overview |
| Verify that network adapter hardware is installed, or install hardware if needed. | Configuring devices |
| Determine whether you will configure TCP/IP properties manually, dynamically, or automatically. | Configure TCP/IP |
| If you decide to manually configure TCP/IP, obtain and record the following TCP/IP information from your network administrator:<br>• IP address<br>• Subnet mask<br>• Default gateways<br>• DNS servers<br>• DNS domain<br>• WINS servers | |
| If you decide to manually configure TCP/IP, configure IP address, subnet mask, and default gateway settings. | Configure TCP/IP for static addressing |
| If you decide to manually configure TCP/IP, use WINS to obtain and resolve NetBIOS names. | Configure TCP/IP to use WINS |
| If you decide to manually configure TCP/IP, use DNS to obtain and resolve DNS names. | Configure TCP/IP to use DNS |
| (Optional) Configure additional or secondary methods to obtain and resolve NetBIOS names. | Advanced configuration |

| | |
|---|---|
| (Optional) Configure additional or secondary methods to obtain and resolve DNS names. | Advanced configuration |
| Use the **ipconfig /all** command to verify all current TCP/IP network configuration values. | Ipconfig |
| Use the **ping** command to verify that you have TCP/IP installed and working properly. | Test a TCP/IP configuration by using the ping command |
| Use the **nslookup** command to verify that you are able to obtain and resolve DNS names. | TCP/IP utilities; Verify DNS registration for domain controllers using the nslookup command |
| Review Internet Protocol security (IPSec) concepts, and then configure IPSec for TCP/IP. | Internet Protocol Security (IPSec) |

The Windows Server 2003 family includes the latest version of Internet Protocol (IP), known as IP version 6. For more information, see IP Version 6.

# New ways to do familiar tasks with TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring TCP/IP. The graphical user interface (GUI) for performing these tasks is different from the way it was in Windows NT Server 4.0 or Windows 2000.

Please choose your upgrade path:

- Windows NT 4.0 to a Windows Server 2003 operating system
- Windows 2000 to a product in a Windows Server 2003 operating system

**Windows NT 4.0 to a Windows Server 2003 operating system**

| If you want to | In Windows NT 4.0 | In this version of Windows |
|---|---|---|
| Configure TCP/IP | Use the **Protocols** tab of **Network** in Control Panel. | Open Network Connections and use the **Properties** of a connection. For more information, see Configure TCP/IP. |
| Add or remove TCP/IP | Use the **Protocols** tab of **Network** in Control Panel. | TCP/IP is installed by default and cannot be added or removed through the GUI. |

**Windows 2000 to a product in a Windows Server 2003 operating system**

| If you want to | In Windows 2000 | In this version of Windows |
|---|---|---|
| Configure TCP/IP | Use the **Properties** of a connection in Network and Dial-up Connections. | Open Network Connections and use the **Properties** of a connection. For more information, see Configure TCP/IP. |
| Add or remove TCP/IP | Use the **Properties** of a connection in Network and Dial-up Connections. | TCP/IP is installed by default and cannot be added or removed through the GUI. |

# Best practices for TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **If your network does not already use TCP/IP, develop a comprehensive IP addressing plan for your network.**

  You need to know the following information to help make a workable plan:

  1. How many physically separate network segments are contained within your network?
  2. How many host systems on each network segment will use TCP/IP?
  3. Will your network be connected to the Internet?

     - If so, will your network be directly connected to the Internet by using a router (which uses public addresses allocated by your Internet service provider), or will your network be indirectly connected to the Internet by using network address translation (NAT) or an application layer gateway such as a proxy server (which uses private addresses)?
     - If not, then it is technically possible to use any IP addressing scheme. However, it is highly recommended that you use private addresses so that an eventual connection to the Internet does not force you to renumber your network.

For more information about connecting a home or small office network to the Internet, see Internet Connection Sharing and network address translation.

- **If you are using private IP addressing, use addresses from the private address ranges reserved by the Internet Assigned Numbers Authority (IANA).**

  For more information, see Numbering your network.

- **Configure a default gateway on only one network adapter.**

  This practice reduces confusion and assures the results you intended. For information about configuring multiple gateways, see Default gateways.

# TCP/IP How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Configure TCP/IP
- Manage Routes
- Diagnose Connections

# Configure TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure TCP/IP**


- Install Simple TCP/IP Services
- Configure TCP/IP for dynamic addressing
- Configure TCP/IP for automatic addressing
- Configure TCP/IP for an alternate configuration
- Disable automatic address configuration
- Configure TCP/IP for static addressing
- Configure TCP/IP to use DNS
- Configure TCP/IP to use WINS
- Configure TCP/IP to use TCP/IP filtering

# Install Simple TCP/IP Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Simple TCP/IP Services**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. In **Components**, click **Networking Services**, and then click **Details**.
4. In **Networking Services**, select the **Simple TCP/IP Services** check box, and then click **OK**.
5. Click **Next**.
6. If prompted, type the path where the distribution files are located, and then click **OK**.
7. Click **Finish**, and then click **Close**.

## Notes

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Simple TCP/IP Services provide support for the optional TCP/IP protocol services listed in the following table.

| Protocol | Description | RFC |
|---|---|---|
| Character Generator (CHARGEN) | Sends data made up of the set of 95 printable ASCII characters. Useful as a debugging tool for testing or troubleshooting line printers. | 864 |
| Daytime | Returns messages containing the day of the week, month, day, year, current time (in *hh:mm:ss* format), and time-zone information. Some programs can use the output from this service for debugging or monitoring variations in system clock time or on a different host. | 867 |
| Discard | Discards all messages received on this port without response or acknowledgment. Can serve as a null port for receiving and routing TCP/IP test messages during network setup and configuration or, in some cases, can be used by programs as a message discard function. | 863 |
| Echo | Echoes back data from any messages it receives on this server port. Echo can be useful as a network debugging and monitoring tool. | 862 |
| Quote of the Day (QUOTE) | Returns a quotation as one or more lines of text in a message. Quotations are taken at random from the following file: *systemroot*\System32\Drivers\Etc\Quotes. A sample quote file is installed with the Simple TCP/IP Services. If this file is missing, the quote service fails. | 865 |

All of these protocol services are classified as elective Internet standards and are defined and described in the specified Request for Comments (RFC) document listed in the table. For more information about these protocol services, see the RFCs.

- Do not install Simple TCP/IP Services unless you specifically need this computer to support communication with other systems that use these protocol services.

- After Simple TCP/IP Services are installed, you cannot enable or disable the separate services individually.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
TCP/IP RFCs

# Configure TCP/IP for dynamic addressing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure TCP/IP for dynamic addressing**

1. Open Network Connections.
2. Right-click the network connection that you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Obtain an IP address automatically**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- These instructions are only required if a static TCP/IP configuration was previously used. By default, computers running Windows Server 2003 operating systems attempt to obtain the TCP/IP configuration from a DHCP server on your network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure TCP/IP for automatic addressing
Configure TCP/IP for static addressing
TCP/IP configuration items
TCP/IP configuration methods
Configure TCP/IP for an alternate configuration

# Configure TCP/IP for automatic addressing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure TCP/IP for automatic addressing**

1. Open Network Connections.
2. Right-click the network connection that you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Obtain an IP address automatically**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- By default, TCP/IP is configured to use DHCP addressing. If a DHCP server is unavailable, then servers running Windows Server 2003 can use Automatic Private IP Addressing (APIPA) or the alternate configuration method to automate Internet Protocol (IP) configuration of network connections. For more information about these new features, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable automatic address configuration
Configure TCP/IP for dynamic addressing
Configure TCP/IP for static addressing
TCP/IP configuration items
TCP/IP configuration methods
Configure TCP/IP for an alternate configuration
New features for TCP/IP

# Configure TCP/IP for an alternate configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure TCP/IP for an alternate configuration**

1. Open Network Connections.
2. Right-click the network connection that you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. On the **General** tab, click **Obtain an IP address automatically**.
5. On the **Alternate Configuration** tab, click **User configured**, and then type appropriate values for the following:
   - IP address
   - Subnet mask
   - Default gateway
   - Preferred and alternate DNS server
   - Preferred and alternate WINS server

**Notes**

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The alternate configuration is used for computers that travel between a network with static IP addresses and one or more DHCP networks.
- Computers running Windows Server 2003 can use Automatic Private IP Addressing (APIPA) or the alternate configuration method to automate Internet Protocol (IP) configuration of network connections. For more information about these new features, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
New features for TCP/IP
Disable automatic address configuration
Configure TCP/IP for dynamic addressing
Configure TCP/IP for static addressing
TCP/IP configuration items

TCP/IP configuration methods
Configure TCP/IP for automatic addressing

# Disable automatic address configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable automatic address configuration**

1. Open Registry Editor.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the **Last Known Good Configuration** startup option if problems are encountered after manual changes have been applied.

2. In Registry Editor, navigate to the following registry key:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters**

3. Click the Parameters folder, and then click **Edit**.

4. Click **New**, and then click **DWORD Value**.

5. Type the following to create a new entry:

   **IPAutoconfigurationEnabled**

6. Right-click the new entry, and then click **Modify**.

7. In **Edit DWORD Value**, under **Value data**, ensure that the setting is **O**.

8. Exit Registry Editor, and restart the computer.

## Notes

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

- If the **IPAutoconfigurationEnabled** entry is not present, a default value of **1** is assumed, which indicates that APIPA can be used.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure TCP/IP for automatic addressing
Configure TCP/IP for dynamic addressing
Configure TCP/IP for static addressing
TCP/IP configuration items

TCP/IP configuration methods
New features for TCP/IP

# Configure TCP/IP for static addressing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure TCP/IP for static addressing**

1. Open Network Connections.
2. Right-click the network connection you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Use the following IP address**, and do one of the following:
   - For a local area connection, in **IP address**, **Subnet mask**, and **Default gateway**, type the IP address, subnet mask, and default gateway addresses.
   - For all other connections, in **IP address**, type the IP address.
5. Click **Use the following DNS server addresses**.
6. In **Preferred DNS server** and **Alternate DNS server**, type the primary and secondary DNS server addresses.

To configure advanced static address settings for a local area connection, click **Advanced**, and do one or more of the following:

To configure additional IP addresses:

1. On the **IP Settings** tab, in **IP addresses**, click **Add**.
2. In **TCP/IP Address**, type an IP address in **IP address** and a subnet mask in **Subnet mask**, and then click **Add**.
3. Repeat steps 1 and 2 for each IP address you want to add, and then click **OK**.

To configure additional default gateways:

1. On the **IP Settings** tab, in **Default gateways**, click **Add**.
2. In **TCP/IP Gateway Address**, type the IP address of the default gateway in **Gateway**. To manually configure a default route metric, clear the **Automatic metric** check box and type a metric in **Metric**.
3. Click **Add**.
4. Repeat steps 1 through 3 for each default gateway you want to add, and then click **OK**.

To configure a custom metric for this connection, clear the **Automatic metric** check box, and then type a metric value in **Interface metric**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure TCP/IP for automatic addressing
Configure TCP/IP for dynamic addressing
IP addressing
Subnet masks
Default gateways
TCP/IP configuration items
Advanced configuration
Numbering your network
Configure TCP/IP to use DNS
Configure TCP/IP to use WINS

# Configure TCP/IP to use DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure TCP/IP to use DNS**

1. Open Network Connections.
2. Right-click the network connection you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. If you want to obtain DNS server addresses from a DHCP server, click **Obtain DNS server address automatically**.
5. If you want to manually configure DNS server addresses, click **Use the following DNS server addresses**, and in **Preferred DNS server** and **Alternate DNS server**, type the preferred DNS server and alternate DNS server IP addresses.
6. To configure advanced DNS properties, click **Advanced**, click the **DNS** tab, and do one or more of the following:

- To configure an additional DNS server IP address:

Under **DNS server addresses, in order of use**, click **Add**.

In **TCP/IP DNS server**, type the IP address of the DNS server, and then click **Add**.

- To modify the resolution behavior for unqualified DNS names, do the following:

To resolve an unqualified name by appending the primary DNS suffix and the DNS suffix of each connection (if configured), click **Append primary and connection specific DNS suffixes**. If you also want to search the parent suffixes of the primary DNS suffix up to the second level domain, select the **Append parent suffixes of the primary DNS suffix** check box.

To resolve an unqualified name by appending the suffixes from a list of configured suffixes, click **Append these DNS suffixes (in order)**, and then click **Add** to add suffixes to the list.

- To configure a connection-specific DNS suffix, type the DNS suffix in **DNS suffix for this connection**.
- To modify DNS dynamic update behavior, do the following:

To use a DNS dynamic update to register the IP addresses of this connection and the primary domain name of the computer, select the **Register this connection's addresses in DNS** check box. This option is enabled by default. The primary domain name of the computer is the primary DNS suffix appended to the computer name and can be viewed as the full computer name on the **Computer Name** tab (available in System in Control Panel).

To use a DNS dynamic update to register the IP addresses and the connection-specific domain name of this connection, select the **Use this connection's DNS suffix in DNS registration** check box. This option is disabled by default. The connection-specific domain name of this connection is the DNS suffix for this connection appended to the computer name.

**Notes**

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- To completely disable DNS dynamic update for all names on the computer, clear the **Register this connection's addresses in DNS** and **Use this connection's DNS suffix in DNS registration** check boxes for all connections in Network Connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Host name resolution
TCP/IP configuration items
Advanced configuration
Configure TCP/IP for static addressing
Configure TCP/IP to use WINS
Dynamic update

# Configure TCP/IP to use WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure TCP/IP to use WINS**

1. Open Network Connections.
2. Right-click the network connection you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Advanced**, click the **WINS** tab, and then click **Add**.
5. In **TCP/IP WINS server**, type the IP address of the WINS server, and then click **Add**.
6. Repeat steps 4 and 5 for each WINS server IP address you want to add, and then click **OK**.
7. To enable the use of the Lmhosts file to resolve remote NetBIOS names, select the **Enable LMHOSTS lookup** check box. This option is enabled by default.
8. To specify the location of the file that you want to import into the Lmhosts file, click **Import LMHOSTS**, and then select the file in the **Open** dialog box.
9. To enable or disable NetBIOS over TCP/IP, do the following:
    - To enable the use of NetBIOS over TCP/IP, click **Enable NetBIOS over TCP/IP**.
    - To disable the use of NetBIOS over TCP/IP, click **Disable NetBIOS over TCP/IP**.
    - To have the DHCP server determine whether NetBIOS over TCP/IP is enabled or disabled, click **Default**. For more information, see Notes.

**Notes**

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you are using a DHCP server to allocate WINS server IP addresses, you do not need to add WINS server addresses.
- If you disable NetBIOS over TCP/IP, you might not be able to connect to computers that are running operating systems other than Windows 2000, Windows XP, or a Windows Server 2003 operating system.
- If you disable NetBIOS over TCP/IP, you cannot use broadcast-based NetBIOS name resolution to resolve computer names to IP addresses for computers on the same network segment. If your computers are on the same network segment, and NetBIOS over TCP/IP is disabled, you must install a DNS server and either have the computers register with DNS (or manually configure DNS records) or configure entries in the local Hosts file for each computer.
- The DHCP server determines whether NetBIOS over TCP/IP is enabled or disabled based on the value of the Microsoft vendor-specific DHCP option **Disable NetBIOS over TCP/IP (NetBT)**. This option either enables (set to **1**) or disables (set to **2**) NetBIOS over TCP/IP on the DHCP client computer. If the DHCP server does not send the **Disable NetBIOS**

**over TCP/IP (NetBT)** DHCP option during DHCP lease negotiation or TCP/IP is manually configured, NetBIOS over TCP/IP is enabled. For more information about the **Disable NetBIOS over TCP/IP (NetBT)** DHCP option, see Related Topics.

- You can no longer configure the NetBIOS scope ID on the **WINS** tab (previously available in Windows NT 4.0). To configure the NetBIOS scope ID, set the following registry value to the name of the scope ID that you want to use:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\ScopeID**

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

- For more information about DHCP options, see "DHCP Options" at the Microsoft Windows Resource Kits Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
NetBIOS name resolution
TCP/IP configuration items
Advanced configuration
Configure TCP/IP to use DNS
Configure TCP/IP for static addressing

# Configure TCP/IP to use TCP/IP filtering

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure TCP/IP to use TCP/IP filtering**

1. Open Network Connections.
2. Right-click the network connection you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Advanced**.
5. Click **Options**, click **TCP/IP filtering**, and then click **Properties**.
6. Do one of the following:
   - To enable TCP/IP filtering for all adapters, select the **Enable TCP/IP Filtering (All adapters)** check box.
   - To disable TCP/IP filtering for all adapters, clear the **Enable TCP/IP Filtering (All adapters)** check box.
7. Based on your requirements for TCP/IP filtering, configure TCP ports, UDP ports, or IP protocols for the allowed traffic.

## Notes

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Security features
Advanced configuration

# Manage Routes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage routes**

- View the IP routing table
- Add a static IP route
- Remove a static IP route
- Enable the Routing and Remote Access service

# View the IP routing table

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the IP routing table**

1. Open Command Prompt.
2. At the command prompt, type:

   **route print**

**Note**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a static IP route
Remove a static IP route
IP routing
The IP routing table
Route

# Add a static IP route

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a static IP route**

1. Open Command Prompt.

2. At the command prompt, type:

   **route add***destination***mask***subnetmaskgateway***metric***costmetric***if***interface*

   where:

| Static IP route entry | Definition |
|---|---|
| *destination* | Specifies either an IP address or host name for the network or host. |
| *subnetmask* | Specifies a subnet mask to be associated with this route entry. If *subnetmask* is not specified, 255.255.255.255 is used. |
| *gateway* | Specifies either an IP address or host name for the gateway or router to use when forwarding. |
| *costmetric* | Assigns an integer cost metric (ranging from 1 through 9,999) to be used in calculating the fastest, most reliable, and/or least expensive routes. If *costmetric* is not specified, 1 is used. |
| *interface* | Specifies the interface to be used for the route that uses the interface number. If an interface is not specified, the interface to be used for the route is determined from the gateway IP address. |

   For example, to add a static route to the 10.0.0.0 network that uses a subnet mask of 255.0.0.0, a gateway of 192.168.0.1, and a cost metric of 2, you type the following at a command prompt:

   **route add 10.0.0.0 mask 255.0.0.0 192.168.0.1 metric 2**

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To make a static route persistent, you can either enter **route add** commands in a batch file that is run during system startup or use the **-p** option when adding routes.

- Routes added by using the **-p** option are stored in the registry under the following key:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip \Parameters\PersistentRoutes**

- All symbolic names used for *destination* or *gateway* are looked up in the network and computer name database files (Networks and Hosts), which are stored in the local *systemroot*\System32\Drivers\Etc folder.

- If a route addition fails, you can use the **tracert** command to verify that the gateway specified is directly reachable from the same subnet as this computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View the IP routing table
Remove a static IP route
IP routing
The IP routing table
TCP/IP database files
Route

# Remove a static IP route

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a static IP route**

1. Open Command Prompt.

2. At the command prompt, type:

   **route delete***destination*

   where:

   *destination* specifies either an IP address or network name for the remote TCP/IP host computer or network.

   For example, to delete a static route to the 10.0.0.0 network, you type the following at a command prompt:

   **route delete 10.0.0.0**

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- All symbolic names used for *destination* are looked up in the network database file Networks, which is stored in the local *systemroot*\System32\Drivers\Etc folder.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View the IP routing table
Add a static IP route
IP routing
The IP routing table
TCP/IP database files
Route

# Enable the Routing and Remote Access service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable the Routing and Remote Access service**

1. If this server is a member of an Active Directory domain and you are not a domain administrator, instruct your domain administrator to add the computer account of this server to the **RAS and IAS Servers** security group in the domain of which this server is a member. The domain administrator can add the computer account to the **RAS and IAS Servers** security group by using Active Directory Users and Computers or with the **netsh ras add registeredserver** command. If this server is using local authentication or is authenticating against a RADIUS server, skip this step.

2. Open Routing and Remote Access.

3. By default, the local computer is listed as a server.

   To add another server, in the console tree, right-click **Server Status**, and then click **Add Server**.

   In the **Add Server** dialog box, click the applicable option, and then click **OK**.

4. In the console tree, right-click the server you want to enable, and then click **Configure and Enable Routing and Remote Access**.

5. Follow the instructions in the Routing and Remote Access Server Setup Wizard.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Installing the router
Installing the remote access server
Netsh commands for remote access

# Diagnose Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Diagnose connections**

- Test a TCP/IP configuration by using the ping command
- Test TCP/IP connections by using the ping and net view commands
- Trace a path by using the tracert command
- View the NetBIOS name table by using the nbtstat command
- Release and refresh NetBIOS names by using the nbtstat command
- View the Address Resolution Protocol (ARP) cache
- Add a static ARP cache entry
- View current TCP/IP protocol and connection statistics

# Test a TCP/IP configuration by using the ping command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test a TCP/IP configuration by using the ping command**

1. To quickly obtain the TCP/IP configuration of a computer, open Command Prompt, and then type **ipconfig**. From the display of the **ipconfig** command, ensure that the network adapter for the TCP/IP configuration you are testing is not in a **Media disconnected** state. For more information, see Notes.

2. At the command prompt, ping the loopback address by typing **ping 127.0.0.1**.

3. Ping the IP address of the computer.

4. Ping the IP address of the default gateway.

   If the **ping** command fails, verify that the default gateway IP address is correct and that the gateway (router) is operational.

5. Ping the IP address of a remote host (a host that is on a different subnet).

   If the **ping** command fails, verify that the remote host IP address is correct, that the remote host is operational, and that all of the gateways (routers) between this computer and the remote host are operational.

6. Ping the IP address of the DNS server.

   If the **ping** command fails, verify that the DNS server IP address is correct, that the DNS server is operational, and that all of the gateways (routers) between this computer and the DNS server are operational.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The **ipconfig** command is the command-line equivalent to the **winipcfg** command, which is available in Windows 95, Windows 98, Windows 98 Second Edition, and Windows Millennium Edition. Computers running Windows XP or Windows Server 2003 operating systems do not include a graphical equivalent of the **winipcfg** command, however, you can get equivalent functionality for viewing and renewing an IP address by opening Network Connections, right-clicking a network connection, clicking **Status**, and then clicking the **Support** tab.

- If the **ipconfig** command displays **Media disconnected**, the network cable is not plugged in to the network adapter.

- If the **ping** command is not found or the command fails, you can use Event Viewer to check the system log and look for problems reported by Setup or the Internet Protocol (TCP/IP) service.

- The **ping** command uses Internet Control Message Protocol (ICMP) Echo Request and Echo Reply messages. Packet filtering policies on routers, firewalls, or other types of security gateways might prevent the forwarding of this traffic.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Troubleshooting TCP/IP
Test TCP/IP connections by using the ping and net view commands
Ipconfig
Ping
Command-line utilities

# Test TCP/IP connections by using the ping and net view commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test TCP/IP connections by using the ping and net view commands**

1. To test TCP/IP connectivity by using the Ping command, use the display of the **ipconfig** command to ensure that your network adapter is not in a **Media disconnected** state. For more information, see Notes.

2. Open Command Prompt, and then ping the desired host using its IP address.

   If the **ping** command fails with a "Request timed out" message, verify that the host IP address is correct, that the host is operational, and that all of the gateways (routers) between this computer and the host are operational.

3. To test host name resolution by using the **ping** command, ping the desired host using its host name.

   If the **ping** command fails with an "Unable to resolve target system name" message, verify that the host name is correct and that the host name can be resolved by your DNS server.

4. To test TCP/IP connectivity by using the **net view** command, open Command Prompt, and then type **net view ***ComputerName*. The **net view** command lists the file and print shares by establishing a temporary connection. If there are no file or print shares on the specified computer, the **net view** command displays a "There are no entries in the list" message.

   If the **net view** command fails with a "System error 53 has occurred" message, verify that *ComputerName* is correct, that the computer is operational, and that all of the gateways (routers) between this computer and the computer are operational.

   If the **net view** command fails with a "System error 5 has occurred. Access is denied." message, verify that you are logged on using an account that has permission to view the shares on the remote computer.

   To further troubleshoot this connectivity problem, do the following:

   - Use the **ping** command to ping *ComputerName*.

     If the **ping** command fails with an "Unable to resolve target system name" message, then *ComputerName* cannot be resolved to its IP address.

   - Use the **net view** command and the IP address of the computer, as follows:

     **net view ***IPAddress*

     If the **net view** command succeeds, then *ComputerName* is being resolved to the wrong IP address.

     If the **net view** command fails with a "System error 53 has occurred" message, the remote computer might not be running the File and Printer Sharing for Microsoft Networks service.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The **ipconfig** command is the command-line equivalent to the **winipcfg** command, which is available in Windows 95,

Windows 98, and Windows Millennium Edition. Computers running Windows XP or a Windows Server 2003 operating system do not include a graphical equivalent of the **winipcfg** command, however, you can get equivalent functionality for viewing and renewing an IP address by opening Network Connections, right-clicking a network connection, clicking **Status**, and then clicking the **Support** tab.

- If the **ipconfig** command displays **Media disconnected**, the network cable is not plugged in to the network adapter.

- If the **ipconfig** command only yields the message **Windows IP Configuration**, and then returns you to the command prompt, you should check to see if your network adapter has been disabled. For more information, see **To make a local area connection** in Related Topics

- If the **ping** command is not found or the command fails, you can use Event Viewer to check the system log and look for problems reported by Setup or the Internet Protocol (TCP/IP) service.

- The **ping** command uses Internet Control Message Protocol (ICMP) Echo Request and Echo Reply messages. Packet filtering policies on routers, firewalls, or other types of security gateways might prevent the forwarding of this traffic.

- The procedure for testing TCP/IP connections using the **net view** command assumes that you are using only the TCP/IP protocol suite. If additional protocols are installed, such as the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol, **net view** might work correctly when there are problems resolving names and making connections with TCP/IP because **net view** attempts a connection using all installed protocols. Even if the **net view** connection with TCP/IP fails, **net view** might successfully connect to the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting TCP/IP
Test a TCP/IP configuration by using the ping command
Make a virtual private network (VPN) connection
Ping
Net use
Command-line utilities

# Trace a path by using the tracert command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To trace a path by using the tracert command**

- Open Command Prompt, and type the following:

  **tracert***HostName*

  Or, type **tracert** *IPAddress*

  where *HostName* or *IPAddress* is the host name or IP address, respectively, of the remote computer.

  For example, to trace the path from this computer to www.microsoft.com, type the following at a command prompt:

  **tracert www.microsoft.com**

  If you do not want the tracert command to resolve and display the names of all routers in the path, use the **-d** parameter. This expedites the display of the path. For example, to trace a path from this computer to www.microsoft.com without displaying the router names, type the following at a command prompt:

  **tracert -d www.microsoft.com**

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The **tracert** command traces the path by sending Internet Control Message Protocol (ICMP) Echo Request and Echo Reply messages (similar to the **ping** command) to produce command-line report output about each router that is crossed and the roundtrip time (RTT) for each hop. Packet filtering policies on routers, firewalls, or other types of security gateways might prevent forwarding of this traffic.

- If **tracert** is unsuccessful, you can use the command output to help determine at which intermediate router forwarding failed or was slowed.

- For details about packet forwarding and packet loss at each router and link in the path, use the **pathping** command.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting TCP/IP
Test a TCP/IP configuration by using the ping command
Tracert
Pathping
Command-line utilities

# View the NetBIOS name table by using the nbtstat command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the NetBIOS name table by using the nbtstat command**

1. Open Command Prompt.

2. At the command prompt, type:

   **nbtstat -n**

   The NetBIOS local name table for this computer is displayed as command-line output. The name type is listed to indicate whether each name is a unique name or a group name. Also, the status for each name is given, indicating whether the name is registered or unregistered with the network.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The **nbtstat** command displays statistics and names information used by NetBIOS over TCP/IP (NetBT) connections. The **-n** option is used to list only local NetBIOS names for this computer. You can also use **nbtstat** to:

  - List the NetBIOS name table of a remote computer.

  - Display the contents of the NetBIOS names cache of this computer or a remote computer.

  - Manually load or reload the NetBIOS names cache by using entries in the Lmhosts file with a **#PRE** option.

  - List NetBIOS over TCP/IP sessions statistics.

  - Release and refresh NetBIOS names in WINS.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Command-line utilities
TCP/IP database files
NetBIOS name resolution
Troubleshooting TCP/IP
Release and refresh NetBIOS names by using the nbtstat command
Nbtstat
Command-line utilities

# Release and refresh NetBIOS names by using the nbtstat command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To release and refresh NetBIOS names by using the nbtstat command**

1. Open Command Prompt.

2. At the command prompt, type:

   **nbtstat -RR**

   The progress of the release and refresh process is displayed as command-line output. This information indicates if all local NetBIOS names currently registered in WINS for this computer have released and renewed their registration with a WINS server.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- NetBIOS names are registered with WINS and are normally released when the computer shuts down properly. If the computer did not properly shut down or could not contact a WINS server during shutdown and release its names, you can use this command to force a refresh and update in WINS of the local names for this computer. This is useful for mobile or portable computers that are moved between different locations on a network.

- You can also use the **nbtstat** command to:

  - List the NetBIOS names registered to this computer or a remote computer.

  - Display the contents of the NetBIOS names cache of this computer or a remote computer.

  - Manually load or reload the NetBIOS names cache by using entries in the Lmhosts file with a **#PRE** option.

  - List NetBIOS over TCP/IP sessions statistics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Command-line utilities
TCP/IP database files
NetBIOS name resolution
Troubleshooting TCP/IP
View the NetBIOS name table by using the nbtstat command
Nbtstat
Command-line utilities

# View the Address Resolution Protocol (ARP) cache

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the Address Resolution Protocol (ARP) cache**

1. Open Command Prompt.

2. At the command prompt, type **arp -a**.

   For example, if you recently used the **ping** command to test and verify connectivity from this computer to a host computer at IP address 10.0.0.99, the ARP cache displays the following entry:

```
Interface: 10.0.0.1 on Interface 0x1
  Internet Address      Physical Address      Type
  10.0.0.99             00-e0-98-00-7c-dc     dynamic
```

   In this example, the cache entry indicates that the remote host computer at 10.0.0.99 resolved to a media access control address of 00-e0-98-00-7c-dc assigned in the network adapter hardware of the remote computer. The media access control address is the address that your computer used to physically communicate with this remote TCP/IP host on your network.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To minimize ARP broadcast traffic on your network, a cache of hardware-to-software address mappings is maintained for future use. This cache contains the following two types of entries:

  - **Dynamic ARP cache entries**

    These entries are added and deleted automatically during normal use of TCP/IP sessions with remote computers. Dynamic entries age and expire from the cache if not reused within 2 minutes. If a dynamic entry is reused within 2 minutes, it may remain in the cache and age up to a maximum cache life of 10 minutes before being removed or requiring cache renewal by using the ARP broadcast process.

  - **Static ARP cache entries**

    These entries are added manually by using the **arp** command with the **-s** option. Static entries remain in the ARP cache until either the computer is restarted, or the interface is reinitiated. Examples of some of the events that cause an interface to be reinitiated are:

1. The interface is **Disabled**, and then **Enabled**.

2. The **Repair** feature is used.

3. The commands **ipconfig /release** and **ipconfig /renew** are given.

4. The computer awakes from a suspend or hibernation state.

5. You removed and replugged a network cable or moved a computer out of and back into wireless range.

6. ARP cache entries can be deleted manually by using the **arp** command with the **-d** option.

```
Example: arp -d 192.168.0.1
Removes ARP mappings for 192.168.0.1 on all interfaces.
```

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting TCP/IP
Address Resolution Protocol (ARP)
Add a static ARP cache entry
Arp
Command-line utilities

# Add a static ARP cache entry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a static ARP cache entry**

1. Open Command Prompt.

2. At the command prompt, type:

   **arp -s***ip_addressmac_address*

   where:

| | |
|---|---|
| *ip_address* | Specifies the IP address of a local (on the same subnet) TCP/IP node. |
| *mac_address* | Specifies the media access control address for a network adapter installed and used on the local TCP/IP node. |

   For example, to add a static ARP entry for a local TCP/IP node with an IP address of 10.0.0.200 that resolves to a media access control address of 00-10-54-CA-E1-40, type the following at a command prompt:

   **arp -s 10.0.0.200 00-10-54-CA-E1-40**

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Static ARP entries can help speed access for hosts that are frequently used.

- Static ARP entries are valid only until your computer is restarted. To make static ARP cache entries persistent, you can add **arp** commands to a batch file that is run at system startup.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting TCP/IP
Address Resolution Protocol (ARP)
View the Address Resolution Protocol (ARP) cache
Arp
Command-line utilities

# View current TCP/IP protocol and connection statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view current TCP/IP protocol and connection statistics**

1. Open Command Prompt.
2. At the command prompt, type:

   **netstat**

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view **netstat** command-line options, type:

  **netstat /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting TCP/IP
Command-line utilities
Core Protocols of TCP/IP
Netstat

# TCP/IP Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about TCP/IP for Windows-based networks:

- TCP/IP Overview
- Understanding TCP/IP
- Using TCP/IP
- TCP/IP Resources

# TCP/IP Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP overview**
This section covers:

- Security features
- Security information for TCP/IP
- TCP/IP defined
- TCP/IP background
- The TCP/IP model
- TCP/IP utilities
- New features for TCP/IP

# Security features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security features**
TCP/IP incorporates security features that provide protection of the TCP/IP data as it is sent on the network and configuration of the types of local host traffic that are processed.

**Internet Protocol security**
Internet Protocol security (IPSec) is a set of Internet standards that uses cryptographic security services to provide the following:

- Confidentiality

  IPSec traffic is encrypted. Captured IPSec traffic is unintelligible without knowledge of the encryption key.

- Authentication

  IPSec traffic is digitally signed with the shared encryption key so that the receiver can verify that it was sent by the IPSec peer.

- Data integrity

  IPSec traffic contains a cryptographic checksum that incorporates the encryption key. The receiver can verify that the packet was not modified in transit.

For more information about IPSec, see Internet Protocol Security (IPSec).

**TCP/IP filtering**
With TCP/IP filtering, a feature known as TCP/IP Security in Microsoft® Windows NT® 4.0, you can specify exactly which types of incoming TCP/IP traffic are processed for each IP interface. This feature is designed to isolate the traffic that is processed by Internet or intranet servers in the absence of other TCP/IP filtering provided by the Routing and Remote Access service or other TCP/IP programs or services. TCP/IP filtering is disabled by default.

TCP/IP filtering is a set of filters for inbound local host TCP/IP traffic. Local host traffic is traffic that is processed by the host because the destination IP address of inbound TCP/IP traffic is addressed to an assigned interface addresses, appropriate subnet broadcast addresses, or a multicast address. TCP/IP filtering does not apply to routed traffic that is forwarded between interfaces.

With TCP/IP filtering, you can confine local host inbound TCP/IP traffic based on the:

- Destination TCP port
- Destination UDP port
- IP protocol

For information about configuring TCP/IP filtering, see Configure TCP/IP to use TCP/IP filtering.

# Security information for TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for TCP/IP**

TCP/IP internetworks are susceptible to a variety of possible attacks, from passive attacks (such as eavesdropping) to active attacks (such as denial-of-service attacks). For more information, see Security issues with IP and Security information for IPv6.

It is important to follow best practices for security when using IPv4 on your network. For more information, see Best practices for security.

The following are known security issues for IPv4:

- **The installation of an unauthorized router can cause the reconfiguration of clients and the rerouting of IPv4 traffic.**

  To communicate with IPv4 nodes on other network segments, IPv4 hosts use a default router, also called a default gateway. The IP address of the default router is statically configured or assigned through Dynamic Host Configuration Protocol (DHCP).

  A malicious user with physical access to the IPv4 network can cause a denial-of-service attack on IPv4 hosts by installing an unauthorized router on the network segment. The unauthorized IPv4 router can reroute link traffic and disrupt other network services.

  **Recommendation:**

  - Ensure that unauthorized persons do not have physical or wireless access to your network.

- **A computer can seize the IPv4 address of a computer or device on the same subnet, causing other network devices to create an incorrect entry in their ARP cache.**

  Nodes on an IPv4 subnet use Address Resolution Protocol(ARP) to resolve IP addresses used by TCP/IP-based software to media access control (MAC) addresses used by LAN hardware. The resolved link-layer address becomes an entry in the ARP cache on a node. If an IPv4 node spoofs the address of another node, it can cause other computers on the subnet to add a false entry to their ARP cache. All of the traffic that is intended for the computer with the spoofed IPv4 address instead goes to the computer of the attacker. In addition, the attacker can send traffic that will appear to have been generated by the spoofed computer.

  **Recommendation:**

  - Ensure that unauthorized persons do not have physical or wireless access to your network.

# TCP/IP defined

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP defined**
TCP/IP for is:

- Networking software based on industry-standard networking protocols.
- A routable, enterprise networking protocol that supports the connection of your Microsoft® Windows®-based computer to both LAN and WAN environments.
- Core technologies and utilities for connecting your Windows-based computer with dissimilar systems and sharing information.
- A foundation for gaining access to global Internet services, such as the World Wide Web and File Transfer Protocol (FTP) servers.
- A robust, scalable, cross-platform, client/server framework.

TCP/IP provides basic TCP/IP utilities that enable Windows-based computers to connect and share information with other Microsoft and non-Microsoft systems, including:

- Internet hosts
- Apple Macintosh systems
- IBM mainframes
- UNIX systems
- Open VMS systems
- Windows Server 2003 family
- Windows XP
- Windows 2000
- Windows NT®
- Windows Millennium Edition
- Windows 98
- Windows 95
- Windows for Workgroups
- LAN Manager
- Network-ready printers, such as HP LaserJet series printers that use HP JetDirect cards

**Note**

- The Windows Server 2003 family includes the latest version of Internet Protocol (IP), known as IP version 6. For more information, see IP Version 6.

# TCP/IP background

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP background**

Transmission Control Protocol/Internet Protocol (TCP/IP) is an industry-standard suite of protocols designed for large-scale internetworks that span LAN and WAN environments.

As the following timeline shows, the origins of TCP/IP began in 1969, when the U.S. Department of Defense (DoD) commissioned the Advanced Research Projects Agency Network (ARPANET).




The ARPANET was the result of a resource-sharing experiment. The purpose was to provide high-speed network communication links between various supercomputers located at various regional sites within the United States.

Early protocols such as Telnet (for virtual terminal emulation) and File Transfer Protocol (FTP) were first developed to specify basic utilities needed for sharing information across the ARPANET. As the ARPANET grew in size and scope, two other important protocols appeared:

- In 1974, Transmission Control Protocol (TCP) was introduced as a draft specification that described how to build a reliable, host-to-host data transfer service over a network.

- In 1981, Internet Protocol (IP) was introduced in draft form and described how to implement an addressing standard and route packets between interconnected networks.

On January 1, 1983, ARPANET began to require standard use of the TCP and IP protocols for all network traffic and essential communication. From this date forward, ARPANET started to become more widely known as the Internet and its required protocols started to become more widely known as the *TCP/IP protocol suite*.

The TCP/IP protocol suite is implemented in a variety of TCP/IP software offerings available for use with many computing platforms. Today, TCP/IP software remains widely in use on the Internet and is used often for building large routed private internetworks.

### Note

- The Microsoft® Windows Server 2003 family includes the latest version of the Internet Protocol (IP), known as IP version 6. For more information, see IP Version 6.

# The TCP/IP model

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The TCP/IP model**

TCP/IP is based on a four-layer reference model. All protocols that belong to the TCP/IP protocol suite are located in the top three layers of this model.

As shown in the following illustration, each layer of the TCP/IP model corresponds to one or more layers of the seven-layer Open Systems Interconnection (OSI) reference model proposed by the International Standards Organization (ISO).




The types of services performed and protocols used at each layer within the TCP/IP model are described in more detail in the following table.

| Layer | Description | Protocols |
|---|---|---|
| Application | Defines TCP/IP application protocols and how host programs interface with transport layer services to use the network. | HTTP, Telnet, FTP, TFTP, SNMP, DNS, SMTP, X Windows, other application protocols |
| Transport | Provides communication session management between host computers. Defines the level of service and status of the connection used when transporting data. | TCP, UDP, RTP |
| Internet | Packages data into IP datagrams, which contain source and destination address information that is used to forward the datagrams between hosts and across networks. Performs routing of IP datagrams. | IP, ICMP, ARP, RARP |
| Network interface | Specifies details of how data is physically sent through the network, including how bits are electrically signaled by hardware devices that interface directly with a network medium, such as coaxial cable, optical fiber, or twisted-pair copper wire. | Ethernet, Token Ring, FDDI, X.25, Frame Relay, RS-232, v.35 |

For more information about ARP, IP, ICMP, IGMP, UDP, and TCP, see Understanding TCP/IP.

**Note**

- The OSI reference model is not specific to TCP/IP. It was developed by the ISO in the late 1970s as a framework for describing all functions required of an open interconnected network. It is a widely known and accepted reference model in the data communications field and is used here only for comparison purposes.

# TCP/IP utilities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP utilities**

To assist with the management of TCP/IP, there are three types of TCP/IP-based utilities:

- Connectivity utilities that you can use to interact with and use resources on a variety of Microsoft® and non-Microsoft hosts, such as UNIX systems.
- Diagnostic utilities that you can use to detect and resolve networking problems.
- TCP/IP server software that provides printing and publishing services to TCP/IP-based Microsoft Windows® clients.

The following tables describe the TCP/IP-based utilities.

**Connectivity utilities**

| Utility | Description |
|---|---|
| Ftp | Transfers files of any size between computers running Windows Server 2003 operating systems and any computer running File Transfer protocol (FTP) server software. For more information, see Ftp. |
| Lpr | Sends print jobs to remote UNIX printers managed by Line Printer Daemon (LPD) print server software. For more information, see Lpr. |
| Rcp | Copies files between computers running Windows Server 2003 operating systems and computers running Remote Copy protocol (RCP) server software. For more information, see Rcp. |
| Rexec | Executes processes on remote computers. For more information, see Rexec. |
| Rsh | Runs commands on a computer running Remote Shell (RSH) server software. For more information, see Rsh. |
| Telnet | Uses terminal-based login to remotely access network devices that are running Telnet server software. For more information, see Telnet. |
| Tftp | Transfers small files between computers running Windows Server 2003 operating systems and computers running Trivial File Transfer protocol (TFTP) server software. For more information, see Tftp. |

**Diagnostic utilities**

| Utility | Description |
|---|---|
| Arp | Displays and modifies the Address Resolution protocol (ARP) cache. This cache is a local table used to resolve IP addresses to media access control addresses used on the local network. For more information, see Arp and Address Resolution Protocol (ARP). |
| Hostname | Returns the host name of the local computer. For more information, see Hostname. |
| Ipconfig | Displays the current TCP/IP configuration. Also used to manually release and renew TCP/IP configurations assigned by a DHCP server. For more information, see Ipconfig and Command-line utilities. |
| Lpq | Obtains print queue status information from computers running Line Printer Daemon (LPD) print server software. For more information, see Lpq. |
| Nbtstat | Displays the local NetBIOS name table, a table of NetBIOS names registered by local programs, and the NetBIOS name cache, a local cache listing of NetBIOS computer names that have been resolved to IP addresses. For more information, see Nbtstat and Release and refresh NetBIOS names by using the nbtstat command. |

| | |
|---|---|
| Netsh | Displays and administers TCP/IP protocol settings on the local computer or a remote computer. For more information, see Netsh commands for Interface IP. |
| Netstat | Displays TCP/IP protocol session information. For more information, see Netstat. |
| Nslookup | Checks records, domain host aliases, domain host services, and operating system information by querying DNS servers. For more information, see Nslookup and Verify DNS registration for domain controllers using the nslookup command. |
| Ping | Verifies configurations and tests IP connectivity. For more information, see Ping and Test a TCP/IP configuration by using the ping command. |
| Route | Displays or modifies the local routing table. For more information, see Route. |
| Tracert | Traces the route a packet takes to a destination. For more information, see Tracert and Trace a path by using the tracert command. |
| Pathping | Traces the route a packet takes to a destination and displays information on packet losses for each router in the path. Pathping can also be used to troubleshoot Quality of Service (QoS) connectivity. For more information, see Using the pathping command. |

**Server-based software**

| **Server software** | **Description** |
|---|---|
| TCP/IP Printing service | The TCP/IP Printing service offers an LPD printing service. This service allows UNIX computers to send print jobs to computers running Windows Server 2003 operating systems by using the Line Printer Remote (Lpr) utility for printing. |
| Internet Information Services | Internet Information Services (IIS) offers Web, Gopher, and FTP server software for providing TCP/IP-based publishing services. IIS is provided to offer an enterprise-ready Web server to use with an unlimited number of simultaneous connections. |
| Peer Web Services | Peer Web Services is provided with Windows XP Professional to provide Web-publishing services similar to those offered by IIS for serving fewer than 10 simultaneous connections. |

# New features for TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features for TCP/IP**
TCP/IP includes the following new features for Microsoft® Windows Server 2003 family:

- IGMP version 3
- Alternate configuration
- Automatic determination of the interface metric
- IP version 6

**IGMP version 3**
Internet Group Management Protocol (IGMP) version 3, described in the Internet draft entitled "Internet Group Management Protocol, version 3," provides source-based multicast group membership reporting. Hosts can request to receive multicast traffic from specified sources or from all but a specific set of sources. Source-specific reporting prevents multicast-enabled routers from delivering multicast traffic to a subnet where there are no listening hosts. For more information, see Internet Group Management Protocol (IGMP).

**Alternate configuration**
Alternate configuration enables a computer to use an alternate, manually configured IP address configuration in the absence of a Dynamic Host Configuration Protocol (DHCP) server. You can use an alternate configuration when a computer is used on more than one network, at least one of the networks does not have a DHCP server, and automatic configuration is not wanted.

For example, if you have a laptop computer that you use both at the office and at home, it is useful to configure TCP/IP for an alternate configuration. At the office, the laptop uses a DHCP-allocated TCP/IP configuration. At home, where there is no DHCP server present, the laptop automatically uses the alternate configuration, which provides easy access to home network devices and the Internet and allow seamless operation on both networks, without the manual reconfiguration of TCP/IP settings.

Without an alternate configuration, TCP/IP uses Automatic Private IP Addressing (APIPA).

For more information, see Configure TCP/IP for an alternate configuration.

**Automatic determination of the routing metric**
By default, TCP/IP automatically determines the routing metric for the default gateway of each interface, based on the speed of its associated interface. If interfaces of different speeds are configured to use the same default gateway, the fastest interface has the lowest routing metric for its default route and is used to forward traffic to its default gateway. If there are multiple interfaces of the fastest speed, the interface that is listed first in the binding order is used to forward traffic to its default gateway. For more information, see The IP routing table.

Automatic determination of the interface metric is enabled by default through the **Automatic metric** check box on the **IP Settings** tab and when you manually configure default gateways in **Advanced TCP/IP Settings**. For more information, see Default gateways.

**IP version 6**
The Windows Server 2003 family includes the latest version of Internet Protocol (IP), known as IP version 6. For more information, see IP Version 6.

**Configuration Features**
TCP/IP includes several features that simplify configuration on a single subnet and optimize TCP performance in high-bandwidth network environments.

These features include support for:

- Automatic and alternate private address configuration
- Large TCP windows

- Better roundtrip time (RTT) estimation
- Selective acknowledgments
- ICMP router discovery
- DNS caching
- Disabling NetBIOS over TCP/IP

**Automatic and alternate private address configuration**
Use APIPA to automate TCP/IP address configuration for single-subnet networks that do not contain a DHCP server.

By default, computers running Windows Server 2003 operating systems first try to contact a DHCP server on the network to dynamically obtain configuration for each installed network connection.

- If a DHCP server is reached and the leased configuration is successful, TCP/IP configuration is completed.
- If a DHCP server is not reached, the computer checks for an alternate configuration. If an alternate configuration exists, then its settings are used to configure TCP/IP.
- If an alternate configuration does not exist, TCP/IP uses APIPA to automatically configure TCP/IP. When you use APIPA, a server running Windows Server 2003 determines an address in the reserved IP address range from 169.254.0.1 through 169.254.255.254. This address is used as a temporary IP address configuration until a DHCP server is located. The subnet mask is set to 255.255.0.0.

The APIPA range of IP addresses is reserved by the Internet Assigned Numbers Authority (IANA). Any IP addresses within this range are not used on the Internet.

APIPA eliminates IP address configuration for single-network small office or home office networks that are not connected to the Internet. For information about disabling APIPA, see Disable automatic address configuration.

**Large TCP windows**
Window size reflects the maximum number of packets that can be sent without waiting for positive acknowledgment. Large TCP windows improve TCP/IP performance when large amounts of data are in transit between the sender and receiver. In typical TCP-based communication, the maximum window size is usually fixed at the onset of connection and limited to 64 kilobytes.

With large window support, you can dynamically recalculate and scale the actual window size by using a TCP option as needed during longer sessions. With this option, more data packets are in transit on the network at one time, which increases throughput.

By default, computers running Windows Server 2003 operating systems only accept client requests for the large TCP windows option from TCP1323Opts-enabled computers that they are connecting with. TCP1323Opts-enabled computers make requests for the large TCP windows option during the initial three-way handshake. If you want your computer to be able to make requests for the large TCP windows option, you must enable TCP1323Opts in the registry. For more information about large TCP windows, see RFC 1323, "TCP Extensions for High Performance."

**Better RTT estimation**
TCP uses roundtrip time (RTT) to estimate the amount of time that is needed for roundtrip communication between a sender and a receiver. Servers running Windows Server 2003 support the use of the RFC 1323 TCP timestamp option to improve the way that RTT is estimated. By calculating more accurate RTT information more often, TCP uses better estimates for setting retransmission timers, which helps improve overall TCP speed and performance.

Improvements in estimating RTT help significantly over longer roundtrip network links, such as WANs that span continents or that use either wireless or satellite communication links.

By default, computers running Windows Server 2003 operating systems only accept client requests for the TCP timestamp option from TCP1323Opts-enabled computers that they are connecting with. TCP1323Opts-enabled computers make requests for the TCP timestamp option during the initial three-way handshake. If you want your computer to be able to make requests

for the TCP timestamp option, you must enable TCP1323Opts in the registry. For more information about TCP timestamp, see RFC 1323, "TCP Extensions for High Performance."

**Selective acknowledgments**
In typical TCP-based communication, acknowledgments are cumulative. TCP only acknowledges segments received that are contiguous with previously acknowledged segments. Noncontiguous segments--segments received out of sequence--are not explicitly acknowledged. TCP requires that segments are received and acknowledged within a brief time period, or the missing segment, and all subsequent segments that follow it, must be retransmitted.

Selective acknowledgments are a recent TCP option that allow the receiver to selectively notify and request that a sender resend only data that is actually missing. This results in smaller amounts of data that require retransmission and better use of network bandwidth.

For more information about selective acknowledgments, see RFC 2018, "TCP Selective Acknowledgment Options."

**ICMP router discovery**
Internet Control Message Protocol (ICMP) router discovery is the use of ICMP messages to discover the default gateway on a network segment when a default gateway is not manually configured or assigned by using DHCP. ICMP router discovery consists of two ICMP messages: the router solicitation and the router advertisement. A router solicitation is sent by a host to discover the routers on the network. A router advertisement is sent by a router in response to a router solicitation and periodically to notify hosts on the network that the router is still available. ICMP router discovery is disabled by default on TCP/IP for host computers running Windows XP or Windows Server 2003 operating systems, unless the host receives the perform router discovery option from a DHCP server.

You can configure a server running Windows Server 2003 and the Routing and Remote Access service to support ICMP router discovery as a router. For more information, see Routing Overview.

ICMP router discovery is described in RFC 1256, "ICMP Router Discovery Messages".

**DNS caching**
The Domain Name System (DNS) resolver in TCP/IP caches DNS name queries. You can view and flush the contents of the DNS cache by using the Ipconfig utility. For more information, see Client features and Flush and reset a client resolver cache using the ipconfig command.

**Disabling NetBIOS over TCP/IP**
With computers running Windows Server 2003 operating systems, you can disable NetBIOS over TCP/IP (NetBT) for each network connection. This feature is intended for computers that only use DNS name registration and resolution techniques and communicate by using the Client for Microsoft Networks and the File and Print Sharing for Microsoft Networks components with other computers where NetBT is disabled. Examples of disabling NetBT include computers in specialized or secured roles for your network, such as an edge proxy server or bastion host in a firewall environment, where NetBT support is not required or desired.

The following are considerations for disabling NetBT on computers running Windows Server 2003 operating systems:

- The computer no longer listens for traffic to the NetBIOS datagram service at User Datagram Protocol (UDP) port 138, the NetBIOS name service at UDP port 137, and the NetBIOS session service at Transmission Control Protocol (TCP) port 139.

- TCP/IP-based connections that use the Client for Microsoft Networks and the File and Print Sharing for Microsoft Networks components are only possible to other computers that have NetBT disabled. This affects the ability to browse the network to see network computers and to connect to file shares and network printers.

- NetBIOS name resolution techniques such as WINS, local subnet broadcasts, and the Lmhosts file are no longer used. All name resolution occurs through DNS queries and the Hosts file.

- If the computer needs to participate in WINS as a client, it must have NetBT enabled on at least one network connection.

- If a server running Windows Server 2003 needs to run the WINS service, it must have NetBT enabled on at least one private network connection.

For example, consider disabling NetBT if you have a server computer that has a connection to a private network and a connection to an external network, such as the Internet. In this case, NetBT is not required for the Internet connection. By

disabling NetBT on only the Internet connection, the dual-homed computer continues to function as either a WINS server or client for the internal network, and WINS clients are still serviced for connections made by using other physical network adapters installed on the computer.

You can disable NetBT on the **WINS** tab in the properties of the TCP/IP protocol. For more information, see Configure TCP/IP to use WINS. You can also disable NetBT through DHCP by using a Microsoft vendor-specific DHCP option. For more information, see "Dynamic Host Configuration Protocol" at the Microsoft Resource Kits Web site.

**Additional feature information**

For information about obtaining RFC documents, see TCP/IP RFCs.

# Understanding TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding TCP/IP**
This section covers:

- Core Protocols of TCP/IP
- IP Addressing and Routing
- Name Resolution

# Core Protocols of TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Core protocols of TCP/IP**
This section covers:

- Address Resolution Protocol (ARP)
- Internet Protocol (IP)
- Internet Control Message Protocol (ICMP)
- Internet Group Management Protocol (IGMP)
- User Datagram Protocol (UDP)
- Transmission Control Protocol (TCP)

# Address Resolution Protocol (ARP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Address Resolution Protocol (ARP)**
Address Resolution Protocol (ARP) is a required TCP/IP standard defined in RFC 826, "Address Resolution Protocol (ARP)." ARP resolves IP addresses used by TCP/IP-based software to media access control addresses used by LAN hardware. ARP provides the following protocol services to hosts located on the same physical network:

- Media access control addresses are obtained by using a network broadcast request in the form of the question "What is the media access control address for a device that is configured with the enclosed IP address?"

- When an ARP request is answered, both the sender of the ARP reply and the original ARP requester record each other's IP address and media access control address as an entry in a local table called the ARP cache for future reference.

**Hardware addressing**
Hardware built for use on LANs must contain a unique address programmed into the device by the manufacturer. For Ethernet and Token Ring LAN hardware, this address is known as a media access control address.

Each media access control address identifies the device within its own physical network with a 6-byte number programmed into read-only memory (ROM) on each physical hardware device, such as a network adapter. Media access control addresses are typically displayed in hexadecimal (for example, 00-AA-00-3F-89-4A).

Authority and registration of media access control addresses are overseen by the Institute of Electrical and Electronics Engineers (IEEE). Currently, the IEEE registers and assigns unique numbers for the first three bytes of the media access control address to individual manufacturers. Each manufacturer can then assign the last three bytes of the media access control address to individual network adapters.

**How ARP resolves media access control addresses for local traffic**
The following illustration shows how ARP resolves IP addresses to hardware addresses for hosts on the same local network.




In this example, two TCP/IP hosts, Hosts A and B, are both located on the same physical network. Host A is assigned the IP address of 10.0.0.99 and Host B is assigned the IP address of 10.0.0.100.

When Host A tries to communicate with Host B, the following steps resolve Host B's software-assigned address (10.0.0.100) to Host B's hardware-assigned media access control address:

1. Based on the contents of the routing table on Host A, IP determines that the forwarding IP address to be used to reach Host B is 10.0.0.100. Host A then checks its own local ARP cache for a matching hardware address for Host B.

2. If Host A finds no mapping in the cache, it broadcasts an ARP request frame to all hosts on the local network with the question "What is the hardware address for 10.0.0.100?" Both hardware and software addresses for the source, Host A, are included in the ARP request.

   Each host on the local network receives the ARP request and checks for a match to its own IP address. If a host does not find a match, it discards the ARP request.

3. Host B determines that the IP address in the ARP request matches its own IP address and adds a hardware/software address mapping for Host A to its local ARP cache.

4. Host B sends an ARP reply message containing its hardware address directly back to Host A.

5. When Host A receives the ARP reply message from Host B, it updates its ARP cache with a hardware/software address mapping for Host B.

Once the media access control address for Host B has been determined, Host A can send IP traffic to Host B by addressing it to Host B's media access control address.

**How ARP resolves media access control addresses for remote traffic**
ARP is also used to forward IP datagrams to local routers for destinations that are not on the local network. In this situation, ARP resolves the media access control address of a router interface on the local network.

The following illustration shows how ARP resolves IP addresses to hardware addresses for two hosts on different physical networks connected by a common router.




In this example, Host A is assigned an IP address of 10.0.0.99 and Host B uses an IP address of 192.168.0.99. Router interface 1 is on the same physical network as Host A and uses the IP address 10.0.0.1. Router interface 2 is on the same physical network as Host B and uses the IP address 192.168.0.1.

When Host A tries to communicate with Host B, the following steps resolve Router interface 1's software-assigned address (10.0.0.1) to its hardware-assigned media access control address:

1. Based on the contents of the routing table on Host A, IP determines that the forwarding IP address to be used to reach host B is 10.0.0.1, the IP address of its default gateway. Host A then checks its own local ARP cache for a matching hardware address for 10.0.0.1.

2. If Host A finds no mapping in the cache, it broadcasts an ARP request frame to all hosts on the local network with the question "What is the hardware address for 10.0.0.1?" Both hardware and software addresses for the source, Host A, are included in the ARP request.

   Each host on the local network receives the ARP request and checks for a match to its own IP address. If a host does not find a match, it discards the ARP request.

3. The router determines that the IP address in the ARP request matches its own IP address and adds a hardware/software address mapping for Host A to its local ARP cache.

4. The router then sends an ARP reply message containing its hardware address directly back to Host A.

5. When Host A receives the ARP reply message from the router, it updates its ARP cache with a hardware/software address mapping for 10.0.0.1.

Once the media access control address for Router interface 1 has been determined, Host A can send IP traffic to Router interface 1 by addressing it to the Router interface 1 media access control address. The router then forwards the traffic to Host B through the same ARP process as discussed in this section.

**The ARP cache**
To minimize the number of broadcasts, ARP maintains a cache of IP address-to-media access control address mappings for future use. The ARP cache can contain both dynamic and static entries. Dynamic entries are added and removed automatically

over time. Static entries remain in the cache until the computer is restarted.

Each dynamic ARP cache entry has a potential lifetime of 10 minutes. New entries added to the cache are timestamped. If an entry is not reused within 2 minutes of being added, it expires and is removed from the ARP cache. If an entry is used, it receives two more minutes of lifetime. If an entry keeps getting used, it receives an additional two minutes of lifetime up to a maximum lifetime of 10 minutes.

You can view the ARP cache by using the **arp** command. To view the ARP cache, type **arp -a** at a command prompt. To view **arp** command-line options, type **arp /?** at a command prompt.

**Note**

- There is a separate ARP cache for each network adapter.

# Internet Protocol (IP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol (IP)**
IP is a required TCP/IP standard defined in RFC 791, "Internet Protocol (IP)." IP is a connectionless, unreliable datagram protocol primarily responsible for addressing and routing packets between hosts.

Connectionless means that a session is not established before exchanging data. Unreliable means that delivery is not guaranteed. IP always makes a best-effort attempt to deliver a packet. An IP packet might be lost, delivered out of sequence, duplicated, or delayed. IP does not attempt to recover from these types of errors. The acknowledgment of packets delivered and the recovery of lost packets is the responsibility of a higher-layer protocol, such as TCP.

An IP packet, also known as an IP datagram, consists of an IP header and an IP payload, as shown in the following illustration.

The IP header contains the following fields for addressing and routing:

| IP header field | Function |
|---|---|
| Source IP address | The IP address of the original source of the IP datagram. |
| Destination IP address | The IP address of the final destination of the IP datagram. |
| Time-to-Live (TTL) | Designates the number of network segments on which the datagram is allowed to travel before being discarded by a router. The TTL is set by the sending host and is used to prevent packets from endlessly circulating on an IP internetwork. When forwarding an IP packet, routers are required to decrease the TTL by at least 1. |

**Note**

- The Windows Server 2003 family includes the latest version of Internet Protocol (IP), known as IP version 6. For more information, see IP Version 6.

# Internet Control Message Protocol (ICMP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Control Message Protocol (ICMP)**
Internet Control Message Protocol (ICMP) is a required TCP/IP standard defined in RFC 792, "Internet Control Message Protocol (ICMP)." With ICMP, hosts and routers that use IP communication can report errors and exchange limited control and status information.

ICMP messages are usually sent automatically in one of the following situations:

- An IP datagram cannot reach its destination.
- An IP router (gateway) cannot forward datagrams at the current rate of transmission.
- An IP router redirects the sending host to use a better route to the destination.

ICMP messages are encapsulated and sent within IP datagrams, as shown in the following illustration.




Different types of ICMP messages are identified in the ICMP header. Because ICMP messages are carried in IP datagrams, they are unreliable.

The most common ICMP messages are listed and described in the following table.

| ICMP message | Description |
|---|---|
| Echo request | Determines whether an IP node (a host or a router) is available on the network. |
| Echo reply | Replies to an ICMP echo request. |
| Destination unreachable | Informs the host that a datagram cannot be delivered. |
| Source quench | Informs the host to lower the rate at which it sends datagrams because of congestion. |
| Redirect | Informs the host of a preferred route. |
| Time exceeded | Indicates that the Time-to-Live (TTL) of an IP datagram has expired. |

You can use the **ping** command to send ICMP echo request messages and record the receipt of ICMP echo reply messages. With these messages, you can detect network or host communication failures and troubleshoot common TCP/IP connectivity problems.

For more information about ICMP, see RFC 792, "Internet Control Message Protocol (ICMP)." For more information about obtaining RFCs, see TCP/IP RFCs

# Internet Group Management Protocol (IGMP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Group Management Protocol (IGMP)**
The use of IP multicasting in TCP/IP networks is defined as a TCP/IP standard in RFC 1112, "Internet Group Management Protocol (IGMP)." In addition to defining address and host extensions for how IP hosts support multicasting, this RFC also defines the Internet Group Management Protocol (IGMP) version 1. RFC 2236, "Internet Group Management Protocol (IGMP), version 2" defines IGMP version 2. Both versions of IGMP provide a protocol to exchange and update information about host membership in specific multicast groups. Additionally, the Windows Server 2003 family supports IGMP version 3, described in the Internet draft entitled "Internet Group Management Protocol, version 3." With IGMP version 3, hosts can specify interest in receiving multicast traffic from specified sources or from all but a specific set of sources.

**What is IP multicasting?**
Multicast IP traffic is sent to a single address but is processed by multiple hosts. IP multicasting is similar to a newsletter subscription. When only subscribers receive the newsletter when it is published, only host computers that belong to the multicast group receive and process IP traffic sent to the group's reserved IP address. The set of hosts listening on a specific IP multicast address is called a multicast group.

Other important aspects of IP multicasting include the following:

- Group membership is dynamic, allowing hosts to join and leave the group at any time.
- The ability of hosts to join multicast groups is performed through the sending of IGMP messages.
- Groups are not limited by size and members can be spread out across multiple IP networks (if connecting routers support the propagation of IP multicast traffic and group membership information).
- A host can send IP traffic to the group's IP address without belonging to the corresponding group.

**Multicast addressing**
IP multicast addresses are reserved and assigned from within the Class D address range from 224.0.0.0 through 239.255.255.255. The following table is a partial list of well-known Class D addresses that are reserved for IP multicasting and registered with the Internet Assigned Numbers Authority (IANA).

| IP multicast address | Description |
|---|---|
| 224.0.0.0 | Base address (reserved). |
| 224.0.0.1 | The All Hosts multicast group that contains all systems on the same network segment. |
| 224.0.0.2 | The All Routers multicast group that contains all routers on the same network segment. |
| 224.0.0.5 | The Open Shortest Path First (OSPF) AllSPFRouters address. Used to send OSPF routing information to all OSPF routers on a network segment. |
| 224.0.0.6 | The OSPF AllDRouters address. Used to send OSPF routing information to OSPF designated routers on a network segment. |
| 224.0.0.9 | The RIP Version 2 group address. Used to send RIP routing information to all RIP v2 routers on a network segment. |
| 224.0.1.24 | WINS server group address. Used to support autodiscovery and dynamic configuration of replication for WINS servers. For more information, see WINS replication overview. |

For a full and current listing of additional IP addresses that are reserved for multicasting, see the Internet Assigned Numbers Authority (IANA) Web site.

A single IP address within the Class D reserved range identifies each multicast group. Each group's reserved IP address is

shared by all host members of the group who listen and receive any IP messages sent to the group's IP address.

IP multicast addresses are mapped to a reserved set of media access control multicast addresses.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

**IGMP messages**

IGMP is used to exchange membership status information between IP routers that support multicasting and members of multicast groups. Host membership in a multicast group is reported by individual member hosts, and membership status is periodically polled by multicast routers.

IGMP message types are described in the following table.

| IGMP message type | Description |
|---|---|
| Host membership report | Sent when a host joins a multicast group to declare membership in a specific host group. IGMP host membership report messages are also sent in response to an IGMP host membership query sent by a router. For an IGMP version 3 host membership report message, the host can specify interest in receiving multicast traffic from specified sources or from all but a specific set of sources. Source-specific reporting prevents multicast-enabled routers from delivering multicast traffic to a subnet where there are no listening hosts. |
| Host membership query | Used by a multicast router to periodically poll a network for group members. For an IGMP version 3 host membership query message, the router can query the host's interest in receiving multicast traffic from a specified list of sources. |
| Leave group | Sent by a host when they leave a host group and are the last member of that group on the network segment. |

IGMP messages are encapsulated and sent within IP datagrams as shown in the following illustration.



# User Datagram Protocol (UDP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**User Datagram Protocol (UDP)**
User Datagram Protocol (UDP) is a TCP/IP standard defined in RFC 768, "User Datagram Protocol (UDP)." UDP is used by some programs instead of TCP for fast, lightweight, unreliable transportation of data between TCP/IP hosts.

UDP provides a connectionless datagram service that offers best-effort delivery, which means that UDP does not guarantee delivery or verify sequencing for any datagrams. A source host that needs reliable communication must use either TCP or a program that provides its own sequencing and acknowledgment services.

UDP messages are encapsulated and sent within IP datagrams, as shown in the following illustration.




**UDP ports**
UDP ports provide a location for sending and receiving UDP messages. A UDP port functions as a single message queue for receiving all datagrams intended for the program specified by each protocol port number. This means UDP-based programs can receive more than one message at a time.

The server side of each program that uses UDP listens for messages arriving on their well-known port number. All UDP server port numbers less than 1,024 (and some higher numbers) are reserved and registered by the Internet Assigned Numbers Authority (IANA).

Each UDP server port is identified by a reserved or well-known port number. The following table shows a partial list of well-known UDP server port numbers that are used by standard UDP-based programs.

| UDP port number | Description |
|---|---|
| 53 | DNS name queries |
| 69 | Trivial File Transfer Protocol (TFTP) |
| 137 | NetBIOS name service |
| 138 | NetBIOS datagram service |
| 161 | Simple Network Management Protocol (SNMP) |
| 520 | Routing Information Protocol (RIP) |

For an updated, complete list of all currently registered well-known UDP ports, see the Internet Assigned Numbers Authority (IANA) Web site.

**Note**

- Copyright

**UDP and TCP**
In general, differences in how UDP and TCP deliver data are similar to the differences between a telephone call and a postcard. TCP works like a telephone call by verifying that the destination is available and ready to communicate. UDP works like a postcard--messages are small and delivery is likely, but not always assured.

UDP is typically used by programs that transmit small amounts of data at one time or have real-time requirements. In these

situations, the low overhead and multicasting capabilities of UDP (for example, one datagram, many recipients) are better suited than TCP.

UDP contrasts directly with the services and features provided by TCP. The following table compares differences in how TCP/IP communication is handled depending on whether UDP or TCP is used for transporting data.

| **UDP** | **TCP** |
|---|---|
| Connectionless service; no session is established between hosts. | Connection-oriented service; a session is established between hosts. |
| UDP does not guarantee or acknowledge delivery, or sequence data. | TCP guarantees delivery through the use of acknowledgments and sequenced delivery of data. |
| Programs that use UDP are responsible for providing any reliability needed to transport data. | Programs that use TCP are provided assurance of reliable data transport. |
| UDP is fast, has low overhead requirements, and can support point-to-point and point-to-multipoint communication. | TCP is slower, has higher overhead requirements, and only supports point-to-point communication. |

Both UDP and TCP use ports to identify communications for each TCP/IP program.

# Transmission Control Protocol (TCP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Transmission Control Protocol (TCP)**
Transmission Control Protocol (TCP) is a required TCP/IP standard defined in RFC 793, "Transmission Control Protocol (TCP)," that provides a reliable, connection-oriented packet delivery service. The Transmission Control Protocol:

- Guarantees delivery of IP datagrams.
- Performs segmentation and reassembly of large blocks of data sent by programs.
- Ensures proper sequencing and ordered delivery of segmented data.
- Performs checks on the integrity of transmitted data by using checksum calculations.
- Sends positive messages depending on whether data was received successfully. By using selective acknowledgments, negative acknowledgments for data not received are also sent.
- Offers a preferred method of transport for programs that must use reliable session-based data transmission, such as client/server database and e-mail programs.

**How TCP works**
TCP is based on point-to-point communication between two network hosts. TCP receives data from programs and processes this data as a stream of bytes. Bytes are grouped into segments that TCP then numbers and sequences for delivery.

Before two TCP hosts can exchange data, they must first establish a session with each other. A TCP session is initialized through a process known as a three-way handshake. This process synchronizes sequence numbers and provides control information that is needed to establish a virtual connection between both hosts.

Once the initial three-way handshake completes, segments are sent and acknowledged in a sequential manner between both the sending and receiving hosts. A similar handshake process is used by TCP before closing a connection to verify that both hosts are finished sending and receiving all data.

TCP segments are encapsulated and sent within IP datagrams, as shown in the following illustration.




**TCP ports**
TCP ports use a specific program port for delivery of data sent by using Transmission Control Protocol (TCP). TCP ports are more complex and operate differently from UDP ports.

While a UDP port operates as a single message queue and the network endpoint for UDP-based communication, the final endpoint for all TCP communication is a unique connection. Each TCP connection is uniquely identified by dual endpoints.

Each single TCP server port is capable of offering shared access to multiple connections because all TCP connections are uniquely identified by two pairs of IP address and TCP ports (one address/port pairing for each connected host).

TCP programs use reserved or well-known port numbers, as shown in the following illustration.




The server side of each program that uses TCP ports listens for messages arriving on their well-known port number. All TCP server port numbers less than 1,024 (and some higher numbers) are reserved and registered by the Internet Assigned Numbers Authority (IANA).

The following table is a partial list of some well-known TCP server ports used by standard TCP-based programs.

| TCP port number | Description |
|---|---|
| 20 | FTP server (data channel) |
| 21 | FTP server (control channel) |
| 23 | Telnet server |
| 53 | Domain Name System zone transfers |
| 80 | Web server (HTTP) |
| 139 | NetBIOS session service |

For an updated, complete list of all currently registered well-known TCP ports, see the see the Internet Assigned Numbers Authority (IANA) Web site.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# IP Addressing and Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP addressing and routing**
This section covers:

- IP addressing
- Subnet masks
- IP routing
- The IP routing table

# IP addressing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP addressing**

Each TCP/IP host is identified by a logical IP address. This address is unique for each host that communicates by using TCP/IP. Each 32-bit IP address identifies a location of a host system on the network in the same way that a street address identifies a house on a city street.

Just as a street address has a standard two-part format (a street name and a house number), each IP address is separated internally into two parts--a network ID and a host ID:

- The network ID, also known as a network address, identifies a single network segment within a larger TCP/IP internetwork (a network of networks). All the systems that attach and share access to the same network have a common network ID within their full IP address. This ID is also used to uniquely identify each network within the larger internetwork.

- The host ID, also known as a host address, identifies a TCP/IP node (a workstation, server, router, or other TCP/IP device) within each network. The host ID for each device identifies a single system uniquely within its own network.

Here is an example of a 32-bit IP address:

```
10000011 01101011 00010000 11001000
```

To make IP addressing easier, IP addresses are expressed in dotted decimal notation. The 32-bit IP address is segmented into four 8-bit octets. The octets are converted to decimal (base-10 numbering system) and separated by periods. Therefore, the previous IP address example is 131.107.16.200 when converted to dotted decimal notation.

For more information about dotted decimal notation and the conversion of numbers from binary to decimal, see Converting binary to decimal.

The following illustration shows a sample view of an IP address (131.107.16.200) as it is divided into network and host ID sections. The network ID portion (131.107) is indicated by the first two numbers of the IP address. The host ID portion (16.200) is indicated by the last two numbers of the IP address.




**Notes**

- Because IP addresses identify devices on a network, a unique IP address must be assigned to each device on the network.

- In general, most computers have only a single network adapter installed and therefore require only a single IP address. If a computer has multiple network adapters installed, each adapter needs its own IP address.

**IP address classes**

The Internet community has defined five address classes. Class A, B, and C addresses are used for assignment to TCP/IP nodes.

The class of address defines which bits are used for the network and host ID parts of each address. The address class also defines how many networks and hosts per network can be supported.

The following table uses *w.x.y.z* to designate the four octet values in any given IP address. The table is used to show:

- How the value of the first octet (*w*) of any given IP address effectively indicates the class of address.

- How the octets in an address are divided into network ID and host ID.
- The number of possible networks and hosts per network available for each class.

| **Class** | **Value of *w*** | **Network ID** | **Host ID** | **Number of networks** | **Number of hosts per network** |
|---|---|---|---|---|---|
| A | 1-126 | *w* | *x.y.z* | 126 | 16,777,214 |
| B | 128-191 | *w.x* | *y.z* | 16,384 | 65,534 |
| C | 192-223 | *w.x.y* | *z* | 2,097,152 | 254 |
| D | 224-239 | Reserved for multicast addressing | N/A | N/A | N/A |
| E | 240-254 | Reserved for experimental use | N/A | N/A | N/A |

# Subnet masks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Subnet masks**

Network IDs and host IDs within an IP address are distinguished by using a subnet mask. Each subnet mask is a 32-bit number that uses consecutive bit groups of all ones (1) to identify the network ID and all zeroes (0) to identify the host ID portions of an IP address.

For example, the subnet mask normally used with the IP address 131.107.16.200 is the following 32-bit binary number:

```
11111111 11111111 00000000 00000000
```

This subnet mask number is 16 one-bits followed by 16 zero-bits, indicating that the network ID and host ID sections of this IP address are both 16 bits in length. Normally, this subnet mask is displayed in dotted decimal notation as 255.255.0.0.

The following table displays subnet masks for the Internet address classes.

| Address class | Bits for subnet mask | Subnet mask |
|---|---|---|
| Class A | 11111111 00000000 00000000 00000000 | 255.0.0.0 |
| Class B | 11111111 11111111 00000000 00000000 | 255.255.0.0 |
| Class C | 11111111 11111111 11111111 00000000 | 255.255.255.0 |

Typically, default subnet mask values (as shown in the previous table) are acceptable for most networks with no special requirements and where each IP network segment corresponds to a single physical network.

In some cases, you can use customized subnet masks to implement IP subnetting. With IP subnetting, you can subdivide the default host ID portion of an IP address to specify subnets, which are subdivisions of the original class-based network ID.

By customizing the subnet mask length, you can reduce the number of bits that are used for the actual host ID. For more information about using a custom subnet mask for subnetting your network, see Resource Kit.

**Important**

- To prevent addressing and routing problems, you should make sure all TCP/IP computers on any network segment use the same subnet mask.

# IP routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP routing**

In general terms, routing is the process of forwarding packets between connected networks. For TCP/IP-based networks, routing is part of Internet Protocol (IP) and is used in combination with other network protocol services to provide forwarding capabilities between hosts that are located on separate network segments within a larger TCP/IP-based network.

IP is the mailroom of the TCP/IP protocol, where IP data sorting and delivery take place. Each incoming or outgoing packet is called an IP datagram. An IP datagram contains two IP addresses: the source address of the sending host and the destination address of the receiving host. Unlike hardware addresses, the IP addresses within a datagram remain the same as it travels across a TCP/IP network.

Routing is the primary function of IP. IP datagrams are exchanged and processed on each host by using IP at the Internet layer.

Above the IP layer, transport services on the source host pass data in the form of TCP segments or UDP messages to the IP layer. The IP layer assembles IP datagrams with source and destination address information that is used to route the data through the network. The IP layer then passes datagrams down to the network interface layer. At this layer, data-link services convert IP datagrams into frames for transmission over network-specific media on a physical network. This process happens in reverse order on the destination host.

Each IP datagram contains a source and destination IP address. IP layer services on each host examine the destination address of each datagram, compare this address to a locally maintained routing table, and then decide what further forwarding action to take. IP routers are attached to two or more IP network segments that are enabled to forward packets between them. The following sections discuss IP routers and the use of routing tables in further detail.

**IP routers**

TCP/IP network segments are interconnected by IP routers, which are devices that pass IP datagrams from one network segment to another. This process is known as IP routing and is shown in the following illustration.




IP routers provide the primary means of joining together two or more physically separated IP network segments. All IP routers share two essential characteristics:

- IP routers are multihomed hosts.

  A multihomed host is a network host that uses two or more network connection interfaces to connect to each physically separated network segment.

- IP routers provide packet forwarding for other TCP/IP hosts.

  IP routers are distinct from other hosts that use multihoming in one important way: an IP router must be able to forward IP-based communication between networks for other IP network hosts.

You can implement IP routers by using a variety of possible hardware and software products. Box-based routers--dedicated

hardware devices that run specialized software--are common. In addition, you can use routing solutions that are based on software, such as the Routing and Remote Access service.

For information about IP routing with the Routing and Remote Access service, see Routing.

Regardless of the type of IP routers that you use, all IP routing relies on the use of a routing table to communicate between network segments.

**Routing tables**

TCP/IP hosts use a routing table to maintain knowledge about other IP networks and IP hosts. Networks and hosts are identified by using an IP address and a subnet mask. In addition, routing tables are important because they provide needed information to each local host regarding how to communicate with remote networks and hosts.

For each computer on an IP network, you can maintain a routing table with an entry for every other computer or network that communicates with the local computer. In general, this is not practical, and a default gateway (IP router) is used instead.

When a computer prepares to send an IP datagram, it inserts its own source IP address and the destination IP address of the recipient into the IP header. The computer then examines the destination IP address, compares it to a locally maintained IP routing table, and takes appropriate action based on what it finds. The computer does one of three things:

- It passes the datagram up to a protocol layer above IP on the local host.
- It forwards the datagram through one of its attached network interfaces.
- It discards the datagram.

IP searches the routing table for the route that is the closest match to the destination IP address. The most specific to the least specific route is searched for in the following order:

- A route that matches the destination IP address (host route).
- A route that matches the network ID of the destination IP address (network route).
- The default route.

If a matching route is not found, IP discards the datagram.

For more information, see The IP routing table.

# The IP routing table

Updated: April 13, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The IP routing table**

Every computer that runs TCP/IP makes routing decisions. These decisions are controlled by the IP routing table. To display the IP routing table on computers running Windows Server 2003 operating systems, you can type **route print** at a command prompt.

The following table shows an example of an IP routing table. This example is for a computer running Windows Server 2003, Standard Edition with one 10 megabit/second (Mbit/s) network adapter and the following configuration:

- IP address: 10.0.0.169
- Subnet mask: 255.0.0.0
- Default gateway: 10.0.0.1

| Description | Network destination | Netmask | Gateway | Interface | Metric |
|---|---|---|---|---|---|
| Default route | 0.0.0.0 | 0.0.0.0 | 10.0.0.1 | 10.0.0.169 | 30 |
| Loopback network | 127.0.0.0 | 255.0.0.0 | 127.0.0.1 | 127.0.0.1 | 1 |
| Local network | 10.0.0.0 | 255.0.0.0 | 10.0.0.169 | 10.0.0.169 | 30 |
| Local IP address | 10.0.0.169 | 255.255.255.255 | 127.0.0.1 | 127.0.0.1 | 30 |
| Multicast addresses | 224.0.0.0 | 240.0.0.0 | 10.0.0.169 | 10.0.0.169 | 30 |
| Limited broadcast address | 255.255.255.255 | 255.255.255.255 | 10.0.0.169 | 10.0.0.169 | 1 |

**Note**

- The descriptions in the first column of the preceding table are not actually displayed in the output of the **route print** command.

The routing table is built automatically, based on the current TCP/IP configuration of your computer. Each route occupies a single line in the displayed table. Your computer searches the routing table for an entry that most closely matches the destination IP address.

Your computer uses the default route if no other host or network route matches the destination address included in an IP datagram. The default route typically forwards an IP datagram (for which there is no matching or explicit local route) to a default gateway address for a router on the local subnet. In the previous example, the default route forwards the datagram to a router with a gateway address of 10.0.0.1.

Because the router that corresponds to the default gateway contains information about the network IDs of the other IP subnets within the larger TCP/IP internet, it forwards the datagram to other routers until the datagram is eventually delivered to an IP router that is connected to the specified destination host or subnet within the larger network.

The following sections describe each of the columns displayed in the IP routing table: network destination, netmask, gateway, interface, and metric.

**Network destination**

The network destination is used with the netmask to match the destination IP address. The network destination can range from 0.0.0.0 for the default route through 255.255.255.255 for the limited broadcast, which is a special broadcast address to all hosts on the same network segment.

**Netmask**

The netmask is the subnet mask that is applied to the destination IP address when matching it to the value in the network

destination. When netmask is written in binary, a "1" must match and a "0" need not match. For example, a default route uses a 0.0.0.0 netmask that translates to the binary value 0.0.0.0, so bits need not match. A host route--a route that matches an IP address--uses a 255.255.255.255 netmask that translates to the binary value 11111111.11111111.11111111.11111111, so all of the bits must match.

**Gateway**
The gateway address is the IP address that the local host uses to forward IP datagrams to other IP networks. This is either the IP address of a local network adapter or the IP address of an IP router (such as a default gateway router) on the local network segment.

**Interface**
The interface is the IP address that is configured on the local computer for the local network adapter that is used when an IP datagram is forwarded on the network.

**Metric**
A metric indicates the cost of using a route, which is typically the number of hops to the IP destination. Anything on the local subnet is one hop, and each router crossed after that is an additional hop. If there are multiple routes to the same destination with different metrics, the route with the lowest metric is selected.

For information about adding routes to the IP routing table, see Add a static IP route. For information about deleting routes in the IP routing table, see Remove a static IP route.

**Multihomed hosts**
The following shows the default routing table for a multihomed Windows Server 2003, Standard Edition host with this configuration:

- Network adapter 1 (10 MB)
    - IP address: 10.0.0.169
    - Subnet mask: 255.0.0.0
    - Default gateway: 10.0.0.1
- Network adapter 2 (100 MB)
    - IP address: 192.168.0.200
    - Subnet mask: 255.255.0.0
    - Default gateway: 192.168.0.1

| Adapter | Description | Network destination | Netmask | Gateway | Interface | Metric |
|---|---|---|---|---|---|---|
| 1 | Default route | 0.0.0.0 | 0.0.0.0 | 10.0.0.1 | 10.0.0.169 | 20 |
| 2 | Default route | 0.0.0.0 | 0.0.0.0 | 192.168.0.1 | 192.168.0.200 | 30 |
| 1 | Loopback network | 127.0.0.0 | 255.0.0.0 | 127.0.0.1 | 127.0.0.1 | 1 |
| 1 | Local network | 10.0.0.0 | 255.0.0.0 | 10.0.0.169 | 10.0.0.169 | 20 |
| 1 | Local IP address | 10.0.0.169 | 255.255.255.255 | 127.0.0.1 | 127.0.0.1 | 20 |
| 2 | Local network | 192.168.0.0 | 255.255.0.0 | 192.168.0.200 | 192.168.0.200 | 30 |
| 2 | Local IP address | 192.168.0.200 | 255.255.255.255 | 127.0.0.1 | 127.0.0.1 | 30 |
| 2 | Subnet broadcast | 192.168.0.255 | 255.255.255.255 | 192.168.0.200 | 192.168.0.200 | 30 |
| 1 | Multicast address | 224.0.0.0 | 240.0.0.0 | 10.0.0.169 | 10.0.0.169 | 20 |

| | | | | | | |
|---|---|---|---|---|---|---|
| 2 | Multicast address | 224.0.0.0 | 240.0.0.0 | 192.168.0.200 | 192.168.0.200 | 30 |
| 1 | Limited broadcast | 255.255.255.255 | 255.255.255.255 | 10.0.0.169 | 10.0.0.169 | 1 |
| 2 | Limited broadcast | 255.255.255.255 | 255.255.255.255 | 192.168.0.200 | 192.168.0.200 | 1 |

**Note**

- The descriptions in the first and second columns of the preceding table are not actually displayed in the output of the **route print** command.

For information about enabling IP forwarding on a multihomed computer running a Windows Server 2003 operating system, see Enable the Routing and Remote Access service.

**Notes**

- When you configure a default gateway on each network adapter, you create a 0.0.0.0 route for each network adapter. However, only one default route is actually used. In the previous example, the 10.0.0.169 IP address is the first network adapter in the TCP/IP bindings, and therefore the default route for Network adapter 1 is used. Because only one default gateway is used, you only need to configure one network adapter with a default gateway. This reduces confusion and assures the results you intended.

- If the IP router is a server running Windows Server 2003 and does not have an interface on a given network, it needs a route to get to that network. You can add static routes or use routing protocols that are provided by the Routing and Remote Access service. For information about IP routing with the Routing and Remote Access service, see Routing.

# Name Resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Name resolution**
By default, Windows Server 2003 operating systems use name resolution to translate numerical IP addresses, which are necessary for TCP/IP communications, to computer names, which are easier for users to remember and use than 32-bit IP addresses. With name resolution, you can assign computer names to the IP addresses of the source and destination hosts, and then use the computer name to contact the hosts, rather than using the 32-bit IP addresses. If you use a computer name as an IP alias, ensure that the name is unique and that it resolves to the correct IP address.

There are two types of names to resolve:

- **Host names**

  Host names are used by programs that use the Windows Sockets programming interface, such as Web browsers. For more information, see Host name resolution.

- **NetBIOS names**

  NetBIOS names are used by network programs or services that use the NetBIOS programming interface, such as Client for Microsoft Networks and File and Printer Sharing for Microsoft Networks. For more information, see NetBIOS name resolution.

# Host name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Host name resolution**

Host name resolution means successfully mapping a host name to an IP address. A host name is an alias that is assigned to an IP node to identify it as a TCP/IP host. The host name can be up to 255 characters long and can contain alphabetic and numeric characters, hyphens, and periods. You can assign multiple host names to the same host.

Windows Sockets (Winsock) programs, such as Internet Explorer and the FTP utility, can use one of two values for the destination to which you want to connect: the IP address or a host name. When the IP address is specified, name resolution is not needed. When a host name is specified, the host name must be resolved to an IP address before IP-based communication with the desired resource can begin.

Host names can take various forms. The two most common forms are a nickname and a domain name. A nickname is an alias to an IP address that individual people can assign and use. A domain name is a structured name in a hierarchical namespace called the Domain Name System (DNS). An example of a domain name is www.microsoft.com.

Nicknames are resolved through entries in the Hosts file, which is stored in the *systemroot*\System32\Drivers\Etc folder. For more information, see TCP/IP database files.

Domain names are resolved by sending DNS name queries to a configured DNS server. The DNS server is a computer that stores domain name-to-IP address mapping records or has knowledge of other DNS servers. The DNS server resolves the queried domain name to an IP address and sends the result back.

You are required to configure your computers with the IP address of your DNS server in order to resolve domain names. You must configure Active Directory-based computers running Windows XP Professional or Windows Server 2003 operating systems with the IP address of a DNS server.

For more information, see DNS defined.

# NetBIOS name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NetBIOS name resolution**

NetBIOS name resolution means successfully mapping a NetBIOS name to an IP address. A NetBIOS name is a 16-byte address that is used to identify a NetBIOS resource on the network. A NetBIOS name is either a unique (exclusive) or group (nonexclusive) name. When a NetBIOS process is communicating with a specific process on a specific computer, a unique name is used. When a NetBIOS process is communicating with multiple processes on multiple computers, a group name is used.

An example of a process that uses a NetBIOS name is the File and Printer Sharing for Microsoft Networks service on a computer running Windows XP Professional. When your computer starts up, this service registers a unique NetBIOS name based on the name of your computer. The exact name used by the service is the 15-character computer name plus a 16th character of 0x20. If the computer name is not 15 characters long, it is padded with spaces up to 15 characters.

When you attempt to make a file-sharing connection to a computer by using its computer name, the File and Printer Sharing for Microsoft Networks service on the file server you specify corresponds to a specific NetBIOS name. For example, when you attempt to connect to a computer called CORPSERVER, the NetBIOS name corresponding to the File and Printer Sharing for Microsoft Networks service on that computer is:

```
CORPSERVER      [20]
```

Note the use of spaces to pad the computer name. Before you can establish a file and print sharing connection, a TCP connection must be created. In order for a TCP connection to be established, the NetBIOS name "CORPSERVER [20]" must be resolved to an IP address.

The exact mechanism by which NetBIOS names are resolved to IP addresses depends on the NetBIOS node type that is configured for the node. RFC 1001, "protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods," defines the NetBIOS node types, as listed in the following table.

| Node type | Description |
|---|---|
| B-node(broadcast) | B-node uses broadcast NetBIOS name queries for name registration and resolution. B-node has two major problems: (1) Broadcasts disturb every node on the network, and (2) Routers typically do not forward broadcasts, so only NetBIOS names on the local network can be resolved. |
| P-node (peer-peer) | P-node uses a NetBIOS name server (NBNS), such as a WINS server, to resolve NetBIOS names. P-node does not use broadcasts; instead, it queries the name server directly. |
| M-node (mixed) | M-node is a combination of B-node and P-node. By default, an M-node functions as a B-node. If an M-node is unable to resolve a name by broadcast, it queries a NBNS using P-node. |
| H-node(hybrid) | H-node is a combination of P-node and B-node. By default, an H-node functions as a P-node. If an H-node is unable to resolve a name through the NBNS, it uses a broadcast to resolve the name. |

Computers running Windows Server 2003 operating systems are B-node by default and become H-node when they are configured with a WINS server. Those computers can also use a local database file called Lmhosts to resolve remote NetBIOS names. The Lmhosts file is stored in the *systemroot*\System32\Drivers\Etc folder. For more information, see TCP/IP database files.

It is highly recommended that you configure Windows-based computers with the IP address of your WINS server in order for remote NetBIOS names to be resolved. You must configure Active Directory-based computers, such as Windows XP Professionaland Windows Server 2003 operating systems, with the IP address of a WINS server if they are to communicate with computers running Windows NT, Windows 95, Windows 98, Windows 2000, or Windows Millennium Edition that are not Active Directory-based.

For more information about WINS, see WINS defined.

# Using TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using TCP/IP**

This section covers:

- TCP/IP configuration items
- TCP/IP configuration methods
- Configuring TCP/IP
- Numbering your network
- Default gateways
- Advanced configuration
- Managing TCP/IP from the command line

# TCP/IP configuration items

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP configuration items**
For TCP/IP to function properly, you need to configure the following:

- IP address
- Subnet mask
- Default gateway
- DNS server
- WINS server

**IP address**
You must configure each interface on each TCP/IP node (host or router) with a unique IP address that is correct for the attached network segment. The IP address is a required configuration item.

For more information, see IP addressing.

**Subnet mask**
You must configure each interface on each TCP/IP node (host or router) with a subnet mask that, when combined with the IP address, yields the network ID. All IP interfaces on the same network segment must use the same network ID. Therefore, all IP interfaces on the same network segment must use the same subnet mask. The subnet mask is a required configuration item.

For more information, see Subnet masks.

**Default gateway**
To communicate with TCP/IP nodes on other network segments, you must configure at least one interface with the IP address of a default gateway (a local router that forwards remote TCP/IP traffic to its destination).

You do not need to configure a default gateway for a network that consists of a single network segment.

For more information, see IP routing.

**DNS server**
A DNS server can resolve domain names to IP addresses. When a TCP/IP host is configured with the IP address of a DNS server, the TCP/IP host sends DNS name queries to the DNS server for resolution. A DNS server is required for computers operating in an Active Directory-based network environment.

You do not need to configure a DNS server for a network that consists of a single network segment.

For more information, see Host name resolution.

A server running Windows Server 2003 can be a DNS server. For more information, see DNS.

**WINS server**
A WINS server can store and resolve NetBIOS names to IP addresses. When a TCP/IP host is configured with the IP address of a WINS server, the TCP/IP host registers its own NetBIOS names with the WINS server and sends NetBIOS name queries to the WINS server for resolution. A WINS server is highly recommended when there is more than one network segment in your network and when you have computers that are not based on Active Directory (for example, computers running Windows NT 4.0, Windows 95, Windows 98, and Windows Millennium Edition).

You do not need to configure a WINS server for a network that consists of a single network segment.

For more information, see NetBIOS name resolution.

A server running Windows Server 2003 can be a WINS server. For more information, see WINS.

**IPSec**

To provide security for TCP/IP, you need to configure Internet Protocol security (IPSec). IPSec is a framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services. It provides aggressive protection against private network and Internet attacks through end-to-end security. In Windows XP and the Windows Server 2003 family, IPSec provides the ability to protect communication between workgroups, local area network computers, domain clients and servers, branch offices (which might be physically remote), extranets, and roving clients.

For more information, see Internet Protocol Security (IPSec).

# TCP/IP configuration methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP configuration methods**
You can configure TCP/IP on servers running Windows Server 2003 by using the following methods:

- Automatic configuration
- Dynamic configuration
- Alternate configuration
- Manual configuration

If you are using either a dynamic or manual configuration, find out more about developing an effective TCP/IP network numbering plan in Numbering your network.

**Automatic configuration**
TCP/IP uses Automatic Private IP Addressing (APIPA) by default to provide automatic configuration, using the IP address range 169.254.0.1 to 169.254.255.254 and the subnet mask 255.255.0.0. There is no automatic configuration of a default gateway, DNS server, or WINS server since APIPA is designed for networks that consist of a single network segment and that are not connected to the Internet.

For more information, see Configure TCP/IP for automatic addressing.

**Dynamic configuration**
By using DHCP, TCP/IP configuration is done dynamically and automatically when the computer is started. Dynamic configuration requires the configuration of a DHCP server. By default, computers running Windows Server 2003 operating systems are DHCP clients. By properly configuring the DHCP server, TCP/IP hosts can obtain IP address, subnet mask, default gateway, DNS server, NetBIOS node type, and WINS server configuration information. Dynamic configuration (using DHCP) is recommended for medium-to-large TCP/IP networks.

For more information, see Configure TCP/IP for dynamic addressing.

**Alternate configuration**
Alternate configuration enables a computer to use an alternate, manually configured IP address configuration in the absence of a DHCP server. You can use an alternate configuration when a computer is used on more than one network, at least one of the networks does not have a DHCP server, and automatic configuration is not wanted.

For example, if you have a laptop computer that you use both at the office and at home, it is useful to configure TCP/IP for an alternate configuration. At the office, the laptop uses a DHCP-allocated TCP/IP configuration. At home, where there is no DHCP server present, the laptop automatically uses the alternate configuration, which provides easy access to home network devices and the Internet and allow seamless operation on both networks, without the manual reconfiguration of TCP/IP settings.

Without an alternate configuration, TCP/IP uses APIPA by default.

For more information, see Configure TCP/IP for an alternate configuration.

**Manual configuration**
By manually configuring the properties of the TCP/IP protocol through the properties of a network connection (in Network Connections), you can assign an IP address, subnet mask, default gateway, DNS server, and WINS server. Manual configuration is required in a network with multiple network segments when no DHCP server is present.

For more information, see Configure TCP/IP for static addressing.

# Configuring TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring TCP/IP**

TCP/IP is installed by default. Before you configure TCP/IP you need to know:

- Whether your network supports a *dynamic*TCP/IP configuration, which uses the Dynamic Host Configuration Protocol (DHCP).

  A dynamic configuration is the simplest to complete and can be used if there is a DHCP server on your network. The DHCP server provides IP address, subnet mask, default gateway (IP router), DNS domain name, DNS server, and WINS server configuration information. TCP/IP is configured to use a dynamic configuration by default.

- Whether your network requires a *manual*TCP/IP configuration.

  Some existing networks either do not use DHCP or require that TCP/IP configuration is performed manually for each computer.

- Whether your computer uses an *alternate*TCP/IP configuration.

  For example, laptop computers are sometimes connected to a corporate network that provides dynamic addressing. At other times, laptops require an alternate, manual configuration for home networks.

For more information about configuration methods, see TCP/IP configuration methods.

If you do not use DHCP to configure TCP/IP dynamically, obtain the following information from your network administrator to configure TCP/IP manually:

- The IP address and subnet mask for each network adapter that is installed on the computer.

- The IP address for the local default gateway (IP router).

- The name of your DNS domain and the IP addresses of the DNS servers on the network.

  For Active Directory deployment, the primary DNS suffix (also known as the primary domain name) of a computer is, by default, the name of the Active Directory domain that the computer has joined. If you want your computer's primary DNS suffix to be different from the name of the Active Directory domain that the computer has joined, see Change the DNS suffix of your computer.

- The IP addresses for the WINS servers, if WINS services are available on your network.

For more information about configuration items, see TCP/IP configuration items.

For instructions about configuring TCP/IP, see Configure TCP/IP.

For information and instructions about securing TCP/IP, see Internet Protocol Security (IPSec)

# Numbering your network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Numbering your network**

For private TCP/IP networks that are not directly or indirectly connected to the Internet, you can use any range of valid IP addresses from Class A, B, or C.

For private TCP/IP networks that are indirectly connected to the Internet by using a network address translator (NAT) or an application layer gateway such as a proxy server, the Internet Assigned Numbers Authority (IANA) recommends that you use the private IP addresses shown in the following table.

| Private network ID | Subnet mask | Range of IP addresses |
|---|---|---|
| 10.0.0.0 | 255.0.0.0 | 10.0.0.1 - 10.255.255.254 |
| 172.16.0.0 | 255.240.0.0 | 172.16.0.1 - 172.31.255.254 |
| 192.168.0.0 | 255.255.0.0 | 192.168.0.1 - 192.168.255.254 |

Numbers in these ranges are reserved by IANA for private use on TCP/IP networks and are not used on the Internet.

Usually, for security reasons, you should not connect more than a few TCP/IP systems within your network directly to the Internet. For any host systems on your network that connect to the Internet, you must obtain the use of registered IP addresses from your Internet service provider (ISP).

**Notes**

- If your ISP cannot provide sufficient public addresses to accommodate your needs and you are located in the United States, you can apply directly to the American Registry for Internet Numbers (ARIN) to obtain an IANA-assigned IP network address. For more information about registering a public address, see the American Registry for Internet Numbers Web site.
- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.
- For more information about private IP address numbering for private networks, see RFC 1918, "Address Allocation for Private Internets."

# Default gateways

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Default gateways**

Default gateways serve an important role in TCP/IP networking. They provide a default route for TCP/IP hosts to use when communicating with other hosts on remote networks.

The following illustration shows the role played by two default gateways (IP routers) for two networks: Network 1 and Network 2.




In order for Host A on Network 1 to communicate with Host B on Network 2, Host A first checks its routing table to see if a specific route to Host B exists. If there is no specific route to Host B, Host A forwards its TCP/IP traffic for Host B to its own default gateway, IP Router 1.

The same principle applies if Host B is sending to Host A. Without a specific route to Host A, Host B forwards any TCP/IP traffic destined for Host A to its own default gateway, IP Router 2.

**Why gateways work**

Default gateways are important to make IP routing work efficiently. In most cases, the router that acts as the default gateway for TCP/IP hosts--either a dedicated router or a computer that connects two or more network segments--maintains knowledge of other networks in the larger network and how to reach them.

TCP/IP hosts rely on default gateways for most of their communication needs with hosts on remote network segments. In this way, individual hosts are freed of the burden of having to maintain extensive and continuously updated knowledge about individual remote IP network segments. Only the router that acts as the default gateway needs to maintain this level of routing knowledge to reach other remote network segments in the larger internetwork.

If the default gateway fails, communication beyond the local network segment may be impaired. To prevent this, you can use the **Advanced TCP/IP Settings** dialog box (in Network Connections) for each connection to specify multiple default gateways. You can also use the **route** command to manually add routes to the routing table for heavily used hosts or networks.

**Using multiple gateways**

If you have multiple interfaces and you configure a default gateway for each interface, TCP/IP by default automatically calculates an interface metric that is based on the speed of the interface. The interface metric becomes the metric of the default route in the routing table for the configured default gateway. The interface with the highest speed has the lowest metric for its default route. The result is that whenever multiple default gateways are configured on multiple interfaces, the fastest interface will be used to forward traffic to its default gateway.

If multiple interfaces of the same speed have the same lowest interface metric, then, based upon the binding order, the default gateway of the first network adapter is used. The default gateway for the second network adapter is used when the first is unavailable.

In previous versions of TCP/IP, multiple default gateways all had a default route metric set to 1, and the default gateway used depended on the order of the interfaces. This sometimes caused difficulty in determining which default gateway the TCP/IP protocol was using.

The automatic determination of the interface metric is enabled by default through the **Automatic metric** check box on the **IP Settings** tab on the advanced properties of the **Internet Protocol (TCP/IP)** protocol.

You can disable the automatic determination of the interface metric and type a new value for the interface metric.

# Advanced configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Advanced configuration**
Beyond the configuration of an IP address, subnet mask, default gateway, DNS server, and WINS server, you can configure TCP/IP for:

- Advanced IP settings
- Advanced DNS settings
- Advanced WINS settings
- Advanced options

**Advanced IP settings**
Advanced IP settings include:

- **Multiple IP addresses**

  You can use multiple IP addresses per network connection for multiple IP numbering schemes, such as public addresses used for the Internet and private addresses, and for multiple logical IP networks on the same physical network segment.

- **Multiple default gateways**

  You can use TCP/IP to detect downed routers by using a feature called dead gateway detection. If multiple default gateways are configured, a failing TCP connection updates the IP routing table with the next default gateway in the list.

For more information about configuring advanced IP settings, see Configure TCP/IP for static addressing.

**Advanced DNS settings**
Advanced DNS settings include:

- **Multiple DNS servers**

  If multiple DNS servers are configured, and TCP/IP fails to receive any response from the current DNS server, TCP/IP switches to the next DNS server.

- **Unqualified name resolution**

  For unqualified names, you can configure TCP/IP to either:

  - Append the primary and connection-specific DNS suffixes to the unqualified name for DNS queries.
  - Append a series of configured DNS suffixes to the unqualified name for DNS queries.

- **Connection-specific DNS suffixes**

  You can configure each connection in Network Connections to have its own DNS suffix, in addition to the primary DNS suffix that is configured for the computer on the **Computer Name** tab (available in System in Control Panel).

- **DNS dynamic update behavior**

  If you have DNS servers that support DNS dynamic update (such as the Windows 2000 DNS service or the Windows

Server 2003 DNS service), you can enable the DNS dynamic update of the domain name and IP addresses for the computer. If you configure a connection-specific DNS suffix, you can also enable the DNS dynamic update of the domain name and IP addresses for the connection.

For more information about configuring advanced DNS settings, see Configure TCP/IP to use DNS.

**Advanced WINS settings**
Advanced WINS settings include:

- **Multiple WINS servers**

  If multiple WINS servers are configured, and TCP/IP fails to receive any response from the current WINS server, TCP/IP switches to the next WINS server.

- **Enabling and disabling the use of the Lmhosts file**

  If the Lmhosts file is enabled, TCP/IP parses the Lmhosts file found in the *systemroot*\System32\Drivers\Etc folder during NetBIOS name resolution. By default, the use of the Lmhosts file is enabled. For more information, see TCP/IP database files.

- **Enabling and disabling the use of NetBIOS over TCP/IP**

  You can enable or disable the use of NetBIOS over TCP/IP. If the use of NetBIOS over TCP/IP is disabled, NetBIOS programs cannot run over TCP/IP. Therefore, you may not be able to connect to computers that are running an operating system other than Windows 2000, Windows XP, or Windows Server 2003 operating systems. By default, the use of NetBIOS over TCP/IP is specified by the settings obtained from a DHCP server.

For more information about configuring advanced WINS settings, see Configure TCP/IP to use WINS.

**Note**

- You can no longer configure the NetBIOS scope ID on the **WINS** tab. To configure the NetBIOS scope ID, set the following registry value to the name of the scope ID that you want to use:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\ScopeID**

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Advanced options**
Advanced options include:

- **Enabling TCP/IP filtering and specifying TCP/IP traffic**

  If TCP/IP filtering is enabled, you can specify what types of TCP/IP traffic are processed. For more information, see Security features and Configure TCP/IP to use TCP/IP filtering.

# Managing TCP/IP from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing TCP/IP from the command line**

In addition to using the properties of **Internet Protocol (TCP/IP)** in Network Connections to manage TCP/IP configuration, you can manage TCP/IP configuration by using the Netsh commands for Interface IP.

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

# TCP/IP Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This section covers:

- TCP/IP RFCs
- TCP/IP references
- Resource Kit
- Converting binary to decimal
- Resource: TCP/IP Utilities
- TCP/IP database files
- Using Performance Logs and Alerts

# TCP/IP RFCs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP RFCs**

The standards for TCP/IP are published in a series of documents called Requests for Comments (RFCs). RFCs are an evolving series of reports, proposals for protocols, and protocol standards that describe the internal workings of TCP/IP and the Internet.

Although TCP/IP standards are always published as RFCs, not all RFCs specify standards. RFCs are authored by individuals who voluntarily write and submit a draft proposal for a new protocol or specification to the Internet Engineering Task Force (IETF) and other working groups. Submitted drafts are first reviewed by a technical expert, a task force, or an RFC editor, and then assigned a status.

If a draft passes this initial review stage, it is circulated to the larger Internet community for a period of further comment and review and assigned an RFC number. This RFC number remains constant.

If changes are made to the proposed specification, drafts that are revised or updated are circulated by using a new RFC (a number higher than the original RFC number) to identify more recent documents.

There are five status assignments for RFCs in the standards process, as shown in the following table.

| Status | Description |
|---|---|
| Standard protocol | An official standard protocol of the Internet. |
| Draft standard protocol | Under active consideration and review to become a standard protocol. |
| Proposed standard protocol | A protocol that in the future might become a standard protocol. |
| Experimental protocol | A protocol designed for experimental purposes. An experimental protocol is not intended for operational use. |
| Informational protocol | A protocol developed by another standards organization that is included for the convenience of the Internet community. |
| Historic protocol | Protocols that have been superceded or obsoleted by other protocols. |

**Related RFCs for TCP/IP**

The following table shows the RFCs supported by the TCP/IP protocol and supporting services.

| RFC number | Title |
|---|---|
| 768 | User Datagram Protocol (UDP) |
| 783 | Trivial File Transfer Protocol (TFTP) |
| 791 | Internet Protocol (IP) |
| 792 | Internet Control Message Protocol (ICMP) |
| 793 | Transmission Control Protocol (TCP) |
| 816 | Fault Isolation and Recovery |
| 826 | Address Resolution Protocol (ARP) |
| 854 | Telnet Protocol (TELNET) |
| 862 | Echo Protocol (ECHO) |

| | |
|---|---|
| 863 | Discard Protocol (DISCARD) |
| 864 | Character Generator Protocol (CHARGEN) |
| 865 | Quote of the Day Protocol (QUOTE) |
| 867 | Daytime Protocol (DAYTIME) |
| 894 | IP over Ethernet |
| 919 | Broadcasting Internet Datagrams |
| 922 | Broadcasting Internet Datagrams in the Presence of Subnets |
| 950 | Internet Standard Subnetting Procedure |
| 959 | File Transfer Protocol (FTP) |
| 1001 | Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods |
| 1002 | Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Detailed Specifications |
| 1009 | Requirements for Internet Gateways |
| 1034 | Domain Names - Concepts and Facilities |
| 1035 | Domain Names - Implementation and Specification |
| 1042 | IP over Token Ring |
| 1055 | A Nonstandard for Transmission of IP Datagrams Over Serial Lines: SLIP |
| 1065 | Structure and Identification of Management Information for TCP/IP-based Internets |
| 1112 | Internet Group Management Protocol (IGMP) |
| 1122 | Requirements for Internet Hosts - Communication Layers |
| 1123 | Requirements for Internet Hosts - Application and Support |
| 1144 | Compressing TCP/IP Headers for Low-Speed Serial Links |
| 1157 | Simple Network Management Protocol (SNMP) |
| 1179 | Line Printer Daemon Protocol |
| 1188 | IP over FDDI |
| 1191 | Path MTU Discovery |
| 1201 | IP over ARCNET |
| 1256 | ICMP Router Discovery Messages |
| 1323 | TCP Extensions for High Performance |
| 1332 | PPP Internet Protocol Control Protocol (IPCP) |
| 1518 | An Architecture for IP Address Allocation with CIDR |
| 1519 | Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy |
| 1534 | Interoperation Between DHCP and BOOTP |
| 1542 | Clarifications and Extensions for the Bootstrap Protocol |

| | |
|---|---|
| 1552 | PPP Internetwork Packet Exchange Control Protocol (IPXCP) |
| 1661 | The Point-to-Point Protocol (PPP) |
| 1662 | PPP in HDLC-like Framing |
| 1748 | IEEE 802.5 MIB using SMIv2 |
| 1749 | IEEE 802.5 Station Source Routing MIB using SMIv2 |
| 1812 | Requirements for IP Version 4 Routers |
| 1828 | IP Authentication using Keyed MD5 |
| 1829 | ESP DES-CBC Transform |
| 1851 | ESP Triple DES-CBC Transform |
| 1852 | IP Authentication using Keyed SHA |
| 1878 | Variable Length Subnet Table For IPv4 |
| 1886 | DNS Extensions to Support IP Version 6 |
| 1994 | PPP Challenge Handshake Authentication Protocol (CHAP) |
| 1995 | Incremental Zone Transfer in DNS |
| 1996 | A Mechanism for Prompt DNS Notification of Zone Changes |
| 2018 | TCP Selective Acknowledgment Options |
| 2085 | HMAC-MD5 IP Authentication with Replay Prevention |
| 2104 | HMAC: Keyed Hashing for Message Authentication |
| 2131 | Dynamic Host Configuration Protocol (DHCP) |
| 2136 | Dynamic Updates in the Domain Name System (DNS UPDATE) |
| 2181 | Clarifications to the DNS Specification |
| 2236 | Internet Group Management Protocol, Version 2 |
| 2308 | Negative Caching of DNS Queries (DNS NCACHE) |
| 2401 | Security Architecture for the Internet Protocol |
| 2402 | IP Authentication Header |
| 2406 | IP Encapsulating Security Payload (ESP) |
| 2581 | TCP Congestion Control |

**Obtaining RFCs**

You can obtain RFCs from the RFC Editor Web site. This Web site is currently maintained by members of the Information Sciences Institute (ISI), who publish a classified listing of all RFCs. RFCs are classified as one of the following: Approved Internet standard, proposed Internet standard (circulated in draft form for review), Internet best practices, or For Your Information (FYI) document.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# TCP/IP references

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP references**

For more information about the TCP/IP protocol suite, consult the following references:


- Comer, D. *Internetworking with TCP/IP, Volume I: Principles, Protocols, and Architecture*. Third edition. Englewood Cliffs, NJ: Prentice Hall, 1995.

- Lee, T. and J. Davies. *Microsoft Windows 2000 TCP/IP Protocols and Services Technical Reference*. Redmond, WA: Microsoft Press, 2000.

- Stevens, W.R. *TCP/IP Illustrated, Volume 1: The Protocols*. Reading, MA: Addison-Wesley, 1994.

# Resource Kit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource Kit**

The Windows Server 2003 Resource Kit contains extensive technical information on TCP/IP. See the following chapters:

- "Introduction to TCP/IP" at the Microsoft Resource Kits Web site.
- "Windows Server 2003 TCP/IP" at the Microsoft Windows Resource Kits Web site.
- "TCP/IP Troubleshooting" at the Microsoft Resource Kits Web site.

For more information, see Using the Windows Deployment and Resource Kits.

# Converting binary to decimal

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Converting binary to decimal**

Because all IP addresses and subnet mask values are composed of a standard-length 32-bit data field, they are viewed and interpreted by computers as a single binary numeric string, such as:

10000011 01101011 00000111 00011011

To communicate IP addresses simply and enter them in configurations quickly, you can use *dotted decimal notation* to convert IP address numbers from binary format.

With dotted decimal notation, each 32-bit address number is viewed as four distinct groupings of 8 bits. Each of the four groupings of 8 consecutive bits is referred to as an *octet*.

The first octet uses the first 8 bits (bit places 1-8), the second octet uses the next 8 bits (bit places 9-16), followed by the third octet (bit places 17-24) and fourth octet (bit places 25-32). Periods are used to separate the four octets that are described as separate decimal number values in the IP address.

The following table shows the scientific notation for each bit place in an individual octet and the equivalent decimal value.

**Equivalent decimal values for bit places in a single IP address octet**

| **Octet** | **1st bit** | **2nd bit** | **3rd bit** | **4th bit** | **5th bit** | **6th bit** | **7th bit** | **8th bit** |
|---|---|---|---|---|---|---|---|---|
| Scientific notation | 27 | 26 | 25 | 24 | 23 | 22 | 21 | 20 |
| Decimal notation | 128 | 64 | 32 | 16 | 8 | 4 | 2 | 1 |

For example, if the first bit place is filled with a bit value of 1, the equivalent decimal value is 128. Where the bit value is 0, the equivalent decimal value is 0 as well.

If all the bit places in an octet are filled with ones (1), the highest equivalent decimal value is 255. If all the bit places are filled with zeroes (0), the lowest equivalent decimal value is 0.

To see how each octet in an IP address is converted from an 8-digit binary number to an equivalent decimal number from 0 through 255, a brief example follows.

The following binary string shows the first octet in an IP address:

10000011

In this 8-digit binary number the first, seventh, and eighth bit places are filled with ones. All other bit places are filled with zeroes. By using the previous table as a reference, you can perform simple addition of each bit place's decimal equivalent value to find the decimal sum for this octet string, as follows:

1st bit (128) + 7th bit (2) + 8th bit (1) = octet total (131)

Because the sum is 131, the first octet of the example IP address is 131. After using this method on the other octets, the final result of the conversion is the dotted decimal equivalent: 131.107.7.27.

# Resource: TCP/IP Utilities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP utilities**

This section covers:

- Diagnostic utilities
- Command-line utilities

# Diagnostic utilities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Diagnostic utilities**

The following table lists TCP/IP diagnostic utilities. You can use these utilities to help identify or resolve TCP/IP networking problems.

| Utility | Description |
|---|---|
| Arp | Allows viewing and editing of the Address Resolution Protocol (ARP) cache. The ARP cache contains software-to-hardware address mappings that are used for sending data on the local network. For more information, see Arp. |
| Hostname | Displays the host name of the computer as used for authentication by the Rcp, Rsh, and Rexec utilities. For more information, see Hostname. |
| Ipconfig | Displays current TCP/IP configuration values. Also used to manually release and renew a TCP/IP configuration lease assigned by a DHCP server and to reset DNS name registrations. For more information, see Ipconfig and Command-line utilities. |
| Lpq | Displays print queue status information on computers running Line Printer Daemon (LPD) print server software. For more information, see Lpq. |
| Nbtstat | Checks the state of current NetBIOS over TCP/IP connections, updates the Lmhosts cache, and determines the registered names and scope ID. For more information, see Nbtstat. |
| Netsh | Displays and administers TCP/IP protocol settings on either the local computer or a remote computer. For more information, see Netsh commands for Interface IP. |
| Netstat | Displays protocol statistics and information on current TCP/IP connections. For more information, see Netstat. |
| Nslookup | Checks records, domain host aliases, domain host services, and operating system information by querying DNS server. For more information, see Nslookup and Verify DNS registration for domain controllers using the nslookup command. |
| Ping | Verifies whether TCP/IP is configured correctly and tests connectivity to other host systems. For more information, see Ping and Test a TCP/IP configuration by using the ping command. |
| Route | Allows viewing and editing of the local IP routing table. For more information, see Route. |
| Tracert | Traces the network route taken by an IP datagram to its destination. For more information, see Tracert and Trace a path by using the tracert command. |
| Pathping | Traces the route a packet takes to a destination and displays information on packet losses for each router in the path. Pathping can also be used to troubleshoot Quality of Service (QoS) connectivity. For more information, see Using the pathping command. |

The following table lists additional administrative utilities you can use for TCP/IP troubleshooting:

| Administrative tool | Description |
|---|---|
| SNMP service | Provides statistical information to SNMP management systems. |
| Event Viewer | Tracks errors and events. |
| Performance Logs and Alerts | Analyzes TCP/IP core protocol performance when the SNMP service is installed. |
| Network Monitor | Captures and analyzes TCP/IP to and from computers running Windows Server 2003 operating systems. |

| | |
|---|---|
| Registry Editor | Allows viewing and editing of registry settings. |

# Command-line utilities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Command-line utilities**
This section covers:

- Viewing configuration by using ipconfig /all
- Viewing configuration by using the Status feature
- Refreshing configuration by using ipconfig /renew
- Managing DNS and DHCP class IDs by using ipconfig
- Testing connections by using ping
- Troubleshooting hardware addresses by using arp
- Troubleshooting NetBIOS names by using nbstat
- Displaying connection statistics by using netstat
- Tracing network connections by using tracert
- Testing routers by using pathping

**Viewing configuration by using ipconfig /all**
When you troubleshoot a TCP/IP networking problem, begin by checking the TCP/IP configuration on the computer that is experiencing the problem. You can use the **ipconfig** command to get host computer configuration information, including the IP address, subnet mask, and default gateway.

**Note**

- For clients running Windows 95, Windows 98, and Windows Millennium Edition, use the **winipcfg** command instead of **ipconfig**.

When you use the **ipconfig** command with the **/all** option, a detailed configuration report is produced for all interfaces, including any configured serial ports. With **ipconfig /all**, you can redirect command output to a file and paste the output into other documents. You can also use this output to confirm the TCP/IP configuration of each computer on the network or to further investigate TCP/IP network problems.

For example, if a computer is configured with an IP address that is a duplicate of an existing IP address, the subnet mask appears as 0.0.0.0.

The following example shows the output of the **ipconfig /all** command on a computer that running Windows XP Professional and is configured to use the DHCP server for automatic TCP/IP configuration, and WINS and DNS servers for name resolution.

```
 IP Configuration

        Node Type . . . . . . . . . . : Hybrid
        IP Routing Enabled. . . . . . : No
        WINS proxy Enabled. . . . . . : No

Ethernet adapter Local Area Connection:

        Host Name . . . . . . . . . . : client1.microsoft.com
        DNS Servers . . . . . . . . . : 10.1.0.200
```

```
Description . . . . . : 3Com 3C90x Ethernet Adapter
Physical Address. . . . . : 00-60-08-3E-46-07
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . : Yes
IP Address. . . . . : 192.168.0.112
Subnet Mask . . . . . : 255.255.0.0
Default Gateway . . . . . : 192.168.0.1
DHCP Server . . . . . : 10.1.0.50
Primary WINS Server . . . . : 10.1.0.101
Secondary WINS Server . . . : 10.1.0.102
Lease Obtained. . . . . : Wednesday, September 02, 1998 10:32:13 AM
Lease Expires . . . . . : Friday, September 18, 1998 10:32:13 AM
```

If no problems appear in the TCP/IP configuration, the next step is testing the ability to connect to other host computers on the TCP/IP network.

**Viewing configuration by using the Status feature**
An alternate method of viewing configuration is available through the Status feature of a network connection. For more information, see View the status of a local area connection.

**Refreshing configuration by using ipconfig /renew**
When you troubleshoot a TCP/IP networking problem, begin by checking the TCP/IP configuration on the computer that is experiencing the problem. If the computer is DHCP-enabled and is using a DHCP server to obtain configuration, you can initiate a refresh of the lease by using the **ipconfig /renew** command.

When you use **ipconfig /renew**, all network adapters on the computer that uses DHCP (except those that are manually configured) try to contact a DHCP server and renew their existing configuration or obtain a new configuration.

You can also use the **ipconfig** command with the **/release** option to immediately release the current DHCP configuration for a host.

### Note

- For Windows 95, Windows 98, and Windows Millennium Edition DHCP-enabled clients, use the **release** and **renew** options of the **winipcfg** command instead of **ipconfig /release** and **ipconfig /renew** to perform manual release or renewal of the IP configuration lease for a client.

**Repair feature**
As an alternative to **ipconfig** you can use **Repair** to renew LAN or high-speed Internet connection IP settings. **Repair** performs a series of commands that repair a connection. The commands that are invoked by **Repair** are listed below with their command-line equivalents:

| Repair | Command-line equivalent |
|---|---|
| Checks whether DHCP is enabled and, if enabled, issues a broadcast renew to refresh the IP address | No command line equivalent available |
| Flushes the ARP cache | **arp -d *** |
| Flush the NetBIOS cache | **nbtstat -R** |
| Flushes the DNS cache | **ipconfig /flushdns** |
| Re-registers with WINS | **nbtstat -RR** |
| Re-registers with DNS | **ipconfig /registerdns** |

**Important**

- **Repair** uses a broadcast renew and will cause a computer to accept any lease from any DHCP server that is on the network. In contrast, a unicast renew (**ipconfig /renew**) will only renew the existing lease from the last DHCP server from which the client got a lease.

For more information, see Repair a LAN or high-speed Internet connection.

**Managing DNS and DHCP class IDs by using ipconfig**
You can also use the **ipconfig** command to:

- Display or reset the DNS cache.

  For more information, see Flush and reset a client resolver cache using the ipconfig command.

- Refresh registered DNS names.

  For more information, see Renew DNS client registration using the ipconfig command.

- Display the DHCP class IDs for an adapter.

  For more information, see Show DHCP class ID information at a client computer.

- Set the DHCP class IDs for an adapter.

  For more information, see Set DHCP class ID information at a client computer.

**Testing connections by using ping**
The **ping** command helps to verify IP-level connectivity. When troubleshooting, you can use **ping** to send an ICMP echo request to a target host name or IP address. Use **ping** whenever you need to verify that a host computer can connect to the TCP/IP network and network resources. You can also use **ping** to isolate network hardware problems and incompatible configurations.

It is usually best to verify that a route exists between the local computer and a network host by first using the **ping** command and the IP address of the network host to which you want to connect. Try pinging the IP address of the target host to see if it responds, as follows:

**ping** *IP_address*

You should perform the following steps when using **ping**:

1. Ping the loopback address to verify that TCP/IP is configured correctly on the local computer.

   **ping** 127.0.0.1

2. Ping the IP address of the local computer to verify that it was added to the network correctly.

   **ping***IP_address_of_local_host*

3. Ping the IP address of the default gateway to verify that the default gateway is functioning and that you can communicate with a local host on the local network.

   **ping***IP_address_of_default_gateway*

4. Ping the IP address of a remote host to verify that you can communicate through a router.

   **ping***IP_address_of_remote_host*

The **ping** command uses Windows Sockets-style name resolution to resolve a computer name to an IP address, so if pinging by address succeeds, but pinging by name fails, then the problem lies in address or name resolution, not network connectivity. For more information, see Troubleshooting hardware addresses by using arp.

If you cannot use **ping** successfully at any point, confirm that:

- The computer was restarted after TCP/IP was configured.

- The IP address of the local computer is valid and appears correctly on the **General** tab of the **Internet protocol (TCP/IP) properties** dialog box.

- IP routing is enabled and the link between routers is operational.

You can use different options with the **ping** command to specify the size of packets to use, how many packets to send, whether to record the route used, what Time-to-Live (TTL) value to use, and whether to set the "don't fragment" flag. You can type **ping -?** to see these options.

The following example illustrates how to send two pings, each 1,450 bytes in size, to IP address 131.107.8.1:

```
C:\>ping -n 2 -l 1450 131.107.8.1
Pinging 131.107.8.1 with 1450 bytes of data:

Reply from 131.107.8.1: bytes=1450 time<10ms TTL=32
Reply from 131.107.8.1: bytes=1450 time<10ms TTL=32

Ping statistics for 131.107.8.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate roundtrip times in milliseconds:
    Minimum = 0ms, Maximum =  10ms, Average =  2ms
```

By default, **ping** waits 4,000 milliseconds (4 seconds) for each response to be returned before displaying the "Request Timed Out" message. If the remote system being pinged is across a high-delay link, such as a satellite link, responses may take longer to be returned. You can use the **-w** (wait) option to specify a longer time-out.

**Troubleshooting hardware addresses by using arp**

The Address Resolution protocol (ARP) allows a host to find the media access control address of a host on the same physical network, given the IP address of the host. To make ARP efficient, each computer caches IP-to-media access control address mappings to eliminate repetitive ARP broadcast requests.

You can use the **arp** command to view and modify the ARP table entries on the local computer. The **arp** command is useful for viewing the ARP cache and resolving address resolution problems.

For more information, see View the Address Resolution Protocol (ARP) cache and Add a static ARP cache entry.

**Troubleshooting NetBIOS names by using nbtstat**

NetBIOS over TCP/IP (NetBT) resolves NetBIOS names to IP addresses. TCP/IP provides many options for NetBIOS name resolution, including local cache lookup, WINS server query, broadcast, DNS server query, and Lmhosts and Hosts file lookup.

Nbtstat is a useful tool for troubleshooting NetBIOS name resolution problems. You can use the **nbtstat** command to remove or correct preloaded entries:

- **nbtstat -n** displays the names that were registered locally on the system by programs such as the server and redirector.

- **nbtstat -c** shows the NetBIOS name cache, which contains name-to-address mappings for other computers.

- **nbtstat -R** purges the name cache and reloads it from the Lmhosts file.

- **nbtstat -RR** releases NetBIOS names registered with a WINS server and then renews their registration.

- **nbtstat -a** *name* performs a NetBIOS adapter status command against the computer specified by *name*. The adapter status command returns the local NetBIOS name table for that computer plus the media access control address of the adapter.

- **nbtstat -S** lists the current NetBIOS sessions and their status, including statistics, as shown in the following example:

```
NetBIOS connection table

Local name State     In/out Remote Host    Input   Output
---------------------------------------------------------------
CORP1 <00> Connected Out    CORPSUP1<20>   6MB     5MB
```

```
    CORP1 <00> Connected Out    CORPPRINT<20>  108KB   116KB
    CORP1 <00> Connected Out    CORPSRC1<20>   299KB   19KB
    CORP1 <00> Connected Out    CORPEMAIL1<20> 324KB   19KB
    CORP1 <03> Listening
```

**Displaying connection statistics by using netstat**

You can use the **netstat** command to display protocol statistics and current TCP/IP connections. The **netstat -a** command displays all connections, and **netstat -r** displays the route table plus active connections. The **netstat -o** command displays process IDs so you can view the owner of the port for each connection. The **netstat -e** command displays Ethernet statistics, and **netstat -s** displays per-protocol statistics. If you use **netstat -n**, addresses and port numbers are not converted to names. The following shows sample output for **netstat**:

```
C:\>netstat -e
Interface Statistics

                       Received      Sent
Bytes                  3995837940    47224622
Unicast packets        120099        131015
Non-unicast packets    7579544       3823
Discards               0             0
Errors                 0             0
Unknown protocols      363054211

C:\>netstat -n -o

Active Connections

  proto  Local Address          Foreign Address        State           PID
  TCP    172.31.71.152:1136     157.54.2.84:389        CLOSE_WAIT      180
  TCP    172.31.71.152:2730     172.31.71.99:139       ESTABLISHED     4
  TCP    172.31.71.152:3110     157.54.2.84:389        CLOSE_WAIT      364
  TCP    172.31.71.152:3796     172.30.236.233:1479    ESTABLISHED     1128
  TCP    172.31.71.152:3800     172.30.236.233:1740    ESTABLISHED     1128
  TCP    172.31.71.152:3815     172.30.236.233:1479    ESTABLISHED     908
  TCP    172.31.71.152:3819     172.30.236.233:1740    ESTABLISHED     908
  TCP    172.31.71.152:4034     172.31.16.197:139      TIME_WAIT       0
  TCP    172.31.71.152:4037     157.54.4.183:445       TIME_WAIT       0
  TCP    172.31.71.152:4043     157.60.218.11:119      TIME_WAIT       0
  TCP    172.31.71.152:4044     157.60.218.11:119      TIME_WAIT       0
  TCP    172.31.71.152:4045     157.60.218.11:119      TIME_WAIT       0
  TCP    172.31.71.152:4046     157.60.218.11:119      TIME_WAIT       0
  TCP    172.31.71.152:4047     157.60.218.11:119      TIME_WAIT       0
  TCP    172.31.71.152:4048     157.60.218.11:119      TIME_WAIT       0
  TCP    172.31.71.152:4049     157.60.218.11:119      TIME_WAIT       0
  TCP    172.31.71.152:4050     157.60.218.11:119      TIME_WAIT       0

C:\>netstat -a

Active Connections

 proto Local Address      Foreign Address       State
 TCP  CORP1:1572       172.16.48.10:nbsession   ESTABLISHED
 TCP  CORP1:1589       172.16.48.10:nbsession   ESTABLISHED
 TCP  CORP1:1606       172.16.105.245:nbsession ESTABLISHED
 TCP  CORP1:1632       172.16.48.213:nbsession  ESTABLISHED
 TCP  CORP1:1659       172.16.48.169:nbsession  ESTABLISHED
 TCP  CORP1:1714       172.16.48.203:nbsession  ESTABLISHED
 TCP  CORP1:1719       172.16.48.36:nbsession   ESTABLISHED
 TCP  CORP1:1241       172.16.48.101:nbsession  ESTABLISHED
 UDP  CORP1:1025       *:*
 UDP  CORP1:snmp       *:*
 UDP  CORP1:nbname     *:*
 UDP  CORP1:nbdatagram *:*
 UDP  CORP1:nbname     *:*
 UDP  CORP1:nbdatagram *:*

C:\>netstat -s
IP Statistics

 Packets Received             = 5378528
 Received Header Errors       = 738854
```

```
 Received Address Errors      = 23150
 Datagrams Forwarded          = 0
 Unknown protocols Received   = 0
 Received Packets Discarded   = 0
 Received Packets Delivered   = 4616524
 Output Requests              = 132702
 Routing Discards             = 157
 Discarded Output Packets     = 0
 Output Packet No Route       = 0
 Reassembly Required          = 0
 Reassembly Successful             = 0
 Reassembly Failures               = 0
 Datagrams Successfully Fragmented = 0
 Datagrams Failing Fragmentation   = 0
 Fragments Created                 = 0

ICMP Statistics
                          Received  Sent
 Messages                 693       4
 Errors                   0         0
 Destination Unreachable  685       0
 Time Exceeded            0         0
 Parameter problems       0         0
 Source Quenches          0         0
 Redirects                0         0
 Echoes                   4         0
 Echo Replies             0         4
 Timestamps               0         0
 Timestamp Replies        0         0
 Address Masks            0         0
 Address Mask Replies     0         0

TCP Statistics

 Active Opens                 = 597
 Passive Opens                = 135
 Failed Connection Attempts   = 107
 Reset Connections            = 91
 Current Connections          = 8
 Segments Received            = 106770
 Segments Sent                = 118431
 Segments Retransmitted       = 461

UDP Statistics

 Datagrams Received   = 4157136
 No Ports             = 351928
 Receive Errors       = 2
 Datagrams Sent       = 13809
```

**Tracing network connections by using tracert**

Tracert (Trace Route) is a route-tracing utility that is used to determine the path that an IP datagram takes to reach a destination. The **tracert** command uses the IP Time-to-Live (TTL) field and ICMP error messages to determine the route from one host to another through a network.

**How tracert works**

The Tracert diagnostic utility determines the route taken to a destination by sending Internet Control Message protocol (ICMP) echo packets with varying IP Time-to-Live (TTL) values to the destination. Each router along the path is required to decrement the TTL on a packet by at least 1 before forwarding it. When the TTL on a packet reaches 0, the router should send an "ICMP Time Exceeded" message back to the source computer.

Tracert determines the route by sending the first echo packet with a TTL of 1 and incrementing the TTL by 1 on each subsequent transmission until the target responds or the maximum TTL is reached. The route is determined by examining the "ICMP Time Exceeded" messages sent back by intermediate routers. Some routers silently drop packets with expired TTLs and are invisible to the Tracert utility.

The **tracert** command prints out an ordered list of the near-side interface of the routers in the path that returned the "ICMP Time Exceeded" message. If the **-d** option is used, the Tracert utility does not perform a DNS lookup on each IP address.

In the following example, the packet must travel through two routers (10.0.0.1 and 192.168.0.1) to get to host 172.16.0.99. The

default gateway of the host is 10.0.0.1 and the IP address of the router on the 192.168.0.0 network is 192.168.0.1.

```
C:\>tracert 172.16.0.99 -d
Tracing route to 172.16.0.99 over a maximum of 30 hops
1     2 ms     3 ms     2 ms  10.0.0.1
2    75 ms    83 ms    88 ms  192.168.0.1
3    73 ms    79 ms    93 ms  172.16.0.99
Trace complete.
```

**Troubleshooting with tracert**

You can use the **tracert** command to determine where a packet stopped on the network. In the following example, the default gateway has determined that there is not a valid path for the host on 192.168.10.99. There is probably a router configuration problem or the 192.168.10.0 network does not exist (a bad IP address).

```
C:\>tracert 192.168.10.99
Tracing route to 192.168.10.99 over a maximum of 30 hops
1  10.0.0.1  reports: Destination net unreachable.
Trace complete.
```

The Tracert utility is useful for troubleshooting large networks where several paths can be taken to arrive at the same point.

**Tracert command-line options**

The **tracert** command supports several options, as shown in the following table.

**tracert** [**-d**] [**-h** *maximum_hops*] [**-j** *host-list*] [**-w** *timeout*] *target_name*

| Option | Description |
|---|---|
| **-d** | Specifies that IP addresses are not resolved to host names. |
| **-h** *maximum_hops* | Specifies the number of hops to allow in tracing a route to the host named in *target_name*. |
| **-j** *host-list* | Specifies the list of router interfaces in the path taken by the Tracert utility packets. |
| **-w** *timeout* | Waits the number of milliseconds specified by *timeout* for each reply. |
| *target_name* | Name or IP address of the target host. |

For more information, see Trace a path by using the tracert command.

**Testing routers by using pathping**

The **pathping** command is a route tracing tool that combines features of the **ping** and **tracert** commands with additional information that neither of those tools provides. The **pathping** command sends packets to each router on the way to a final destination over a period of time, and then computes results based on the packets returned from each hop. Since the command shows the degree of packet loss at any given router or link, it is easy to determine which routers or links might be causing network problems. A number of options are available, as shown in the following table.

| Option | Name | Function |
|---|---|---|
| **-n** | Hostnames | Does not resolve addresses to host names. |
| **-h** | Maximum hops | Maximum number of hops to search for target. |
| **-g** | Host-list | Loose source route along host list. |
| **-p** | Period | Number of milliseconds to wait between pings. |
| **-q** | Num_queries | Number of queries per hop. |
| **-w** | Time-out | Waits this many milliseconds for each reply. |
| **-i** | address | Use the specified source address. |
| **-4** | IPv4 | Force pathping to use IPv4. |
| **-6** | IPv6 | Force pathping to use IPv6. |

The default number of hops is 30, and the default wait time before a time-out is 3 seconds. The default period is 250 milliseconds, and the default number of queries to each router along the path is 100.

The following is a typical **pathping** report. The compiled statistics that follow the hop list indicate packet loss at each individual router.

```
D:\>pathping -n server1

Tracing route to server1 [10.54.1.196]
over a maximum of 30 hops:
  0  172.16.87.35
  1  172.16.87.218
  2  192.168.52.1
  3  192.168.80.1
  4  10.54.247.14
  5  10.54.1.196

Computing statistics for 125 seconds...
            Source to Here   This Node/Link
Hop  RTT    Lost/Sent = Pct  Lost/Sent = Pct  Address
  0                                           172.16.87.35
                                 0/ 100 =  0%   |
  1   41ms     0/ 100 =  0%      0/ 100 =  0%  172.16.87.218
                                13/ 100 = 13%   |
  2   22ms    16/ 100 = 16%      3/ 100 =  3%  192.168.52.1
                                 0/ 100 =  0%   |
  3   24ms    13/ 100 = 13%      0/ 100 =  0%  192.168.80.1
                                 0/ 100 =  0%   |
  4   21ms    14/ 100 = 14%      1/ 100 =  1%  10.54.247.14
                                 0/ 100 =  0%   |
  5   24ms    13/ 100 = 13%      0/ 100 =  0%  10.54.1.196

Trace complete.
```

When **pathping** is run, you first see the results for the route as it is tested for problems. This is the same path that is shown by the **tracert** command. The **pathping** command then displays a busy message for the next 125 seconds (this time varies by the hop count). During this time, **pathping** gathers information from all the routers previously listed and from the links between them. At the end of this period, it displays the test results.

The two rightmost columns--**This Node/Link Lost/Sent=Pct** and **Address**--contain the most useful information. The link between 172.16.87.218 (hop 1), and 192.168.52.1 (hop 2) is dropping 13 percent of the packets. All other links are working normally. The routers at hops 2 and 4 also drop packets addressed to them (as shown in the **This Node/Link** column), but this loss does not affect their forwarding path.

The loss rates displayed for the links (marked as a **|** in the rightmost column) indicate losses of packets being forwarded along the path. This loss indicates link congestion. The loss rates displayed for routers (indicated by their IP addresses in the rightmost column) indicate that those routers' CPUs might be overloaded. These congested routers might also be a factor in end-to-end problems, especially if packets are forwarded by software routers.

# TCP/IP database files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP database files**

The following table lists the UNIX-style database files that are stored in the *systemroot*\System32\Drivers\Etc folder corresponding to TCP/IP for computers running Windows Server 2003 operating systems.

| File name | Use |
|---|---|
| Hosts | Provides host name-to-IP address resolution for Windows Sockets programs. |
| Lmhosts | Provides NetBIOS name-to-IP address resolution for Windows-based networking. |
| Networks | Provides network name-to-network ID resolution for Windows Sockets programs. |
| Protocols | Provides protocol name-to-protocol ID resolution for Windows Sockets programs. |
| Services | Provides service name-to-port ID resolution for Windows Sockets programs. |

To troubleshoot any of these files on a local computer, do the following:

1. Make sure the format of entries in each file matches the format defined in the sample file originally provided with TCP/IP.
2. Check for spelling errors.
3. Check for invalid IP addresses and identifiers.

**Note**

- By default, the Lmhosts file does not exist. There is an Lmhosts.sam sample file that you can use as a basis for creating an Lmhosts file.

# Using Performance Logs and Alerts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Performance Logs and Alerts**

You can use the Performance Logs and Alerts to view many different TCP/IP-related counters. Because the Performance Logs and Alerts access statistics that are gathered by the SNMP service, the SNMP service must be installed on Windows-based computers where TCP/IP statistics are to be monitored. Counters for Performance Logs and Alerts are available for network adapters, IP, ICMP, UDP, and TCP.

With Performance Logs and Alerts, you can monitor counters from various systems from a single management window. You can also set alert levels for the counters being monitored.

# Troubleshooting TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Getting 169.254.y.z address.
- Host name resolution is failing.
- Unable to add a default gateway.
- System error 53 occurred.
- Cannot connect to a specific server.
- Long connect times when using Lmhosts for name resolution.
- TCP/IP connection to a host appears to have stopped.
- Cannot ping across a local router when using a dial-up connection.
- A "The registry subkey already exists" message is displayed.

**Getting 169.254.y.z address.**
**Cause:** Through Automatic Private IP Addressing (APIPA), an IP address in the range 169.254.0.1 to 169.254.255.254 with a subnet mask of 255.255.0.0 is automatically assigned when TCP/IP is configured to obtain an IP address automatically, a Dynamic Host Configuration Protocol (DHCP) server is not available, and an alternate configuration is not specified. APIPA is designed to provide automatic IP addressing on single segment networks.

**Solution:** Disable automatic address configuration, determine and fix the cause of the computer's inability to contact the DHCP server, or specify an alternate configuration.

**See also:** New features for TCP/IP; Configure TCP/IP for an alternate configuration; Disable automatic address configuration

**Host name resolution is failing.**
**Cause:** Host name resolution techniques are failing to resolve a host name to its IP address.

**Solution:** The Windows Server 2003 family includes several basic TCP/IP tools that you can use to troubleshoot address and name resolution problems, including the **ping** command. If you receive the correct response when you use the **ping** command with an IP address but an incorrect response when you use **ping** with the host name or NetBIOS name, there is a name resolution problem.

When you use TCP/IP tools such as the **ping** command, the Hosts file or a DNS server is used for name resolution. You can find the Hosts file in the *systemroot*\System32\Drivers\Etc folder. This file is *not* dynamic; you must manually add entries. The format of the file is as follows:

**IP address  Friendly name**

```
172.16.48.10   server1   # Remarks are denoted with a #.
```

When you use a Hosts file for name resolution, the following steps occur:

1. A user at Computer A enters a command by using the host name of Computer B.
2. The Hosts file on Computer A (in the *systemroot*\System32\Drivers\Etc folder) is parsed. When the host name of Computer B is found, it is resolved to an IP address.

The following problems can occur because of errors related to the Hosts file:

- The Hosts file or the DNS server does not contain the particular host name.
- The host name in the Hosts file or in the command is misspelled.
- An invalid IP address is entered for the host name in the Hosts file.
- The Hosts file contains multiple entries for the same host on separate lines; if so, the first entry from the top of the file is used.
- A mapping for a computer name-to-IP address was mistakenly added to the Hosts file, rather than the Lmhosts file.

For DNS server host name resolution, verify that the correct DNS servers are configured and in the proper order. You can use the **ipconfig /all** command to check the current TCP/IP configuration and the **ping** command to check connectivity with your DNS servers. You can use the **nslookup** command to send DNS name queries to your primary DNS server.

If you use a DNS server for host name resolution, the name resolution steps are similar. Computer A sends a DNS name query to its configured DNS server. The DNS server resolves the host name of Computer B to an IP address and sends the results back to Computer A.

**See also:** Host name resolution; Ipconfig; Ping; Nslookup; Verify DNS registration for domain controllers using the nslookup command; TCP/IP database files

**Unable to add a default gateway.**
**Cause:** The IP address of the default gateway is not on the same IP network ID as your IP address.

**Solution:** Find out whether the default gateway is located on the same logical network as the network adapter of the computer by comparing the IP address of the default gateway with the network IDs of any of the network adapters of the computer.

For example, a computer with a single network adapter that is configured with an IP address of 192.168.0.33 and a subnet mask of 255.255.0.0 requires that the default gateway be of the form 192.168.*y.z* because the network ID portion of the IP interface is 192.168.0.0.

**See also:** IP addressing; Subnet masks; Default gateways

**System error 53 occurred.**
**Cause:** System error 53 is returned if name resolution fails for a particular computer name when the **net use** command is used.

**Solution:** If the computer is on the local subnet, confirm that the name is spelled correctly and that the target computer is running TCP/IP as well. If the computer is not on the local subnet, be sure that its name and IP address mapping are available in the Lmhosts file or the WINS database. If all TCP/IP elements appear to be installed properly, use the **ping** command with the remote computer to verify that its TCP/IP software is working.

**See also:** Test a TCP/IP configuration by using the ping command; Test TCP/IP connections by using the ping and net view commands; Net use; Ping

**Cannot connect to a specific server.**
**Cause:** Either NetBIOS name resolution is failing to resolve the name or the wrong IP address is being resolved.

**Solution:** Use the **nbtstat -n** command on the server to determine what names the server registered on the network. The computer name of the computer you are trying to connect to should be on the displayed list. If the name is not listed, try one of the other unique computer names displayed by **nbtstat**.

If the name being used by a remote computer is the same as the name displayed by the **nbtstat -n** command, then make sure that the remote computer has an entry for the server name on the WINS server or in its Lmhosts file.

**See also:** NetBIOS name resolution; Test TCP/IP connections by using the ping and net view commands; Net use; TCP/IP database files; Nbtstat

**Long connect times when using Lmhosts for name resolution.**
**Cause:** The Lmhosts file is parsed sequentially to locate entries without the #PRE option.

**Solution:** Place frequently used entries near the top of the file and place the #PRE entries near the bottom. If an entry is added to the end of a large Lmhosts file, mark the entry in Lmhosts as a preloaded entry by using the #PRE option. Then use the **nbtstat -R** command to update the local name cache immediately.

**See also:** TCP/IP database files; Nbtstat; View the NetBIOS name table by using the nbtstat command; Release and refresh NetBIOS names by using the nbtstat command

**TCP/IP connection to a host appears to have stopped.**
**Cause:** Either data is blocked in TCP and UDP queues or there are network or user-level software delay problems.

**Solution:** Use the **netstat -a** command to show the status of all activity on TCP and UDP ports on the local computer.

The state of a good TCP connection is usually established with 0 bytes in the send and receive queues. If data is blocked in either queue, or if the state is irregular, there is probably a problem with the connection. If not, you are probably experiencing a network or user-level software delay.

**See also:** Netstat; View current TCP/IP protocol and connection statistics

**Cannot ping across a local router when using a dial-up connection.**
**Cause:** This problem occurs if the **Use default gateway on remote network** check box is selected on the **General** tab in the **Advanced TCP/IP Settings** dialog box (in the **Internet Protocol (TCP/IP)** dialog box on the **General** or **Networking** tab for the properties of the dial-up connection).

This feature adds a default route to the IP routing table that is used instead of the current default route for the duration of the connection. The new default route causes network traffic for IP addresses that are not resolved by other entries in the routing table to be sent across the dial-up connection. This new default route is needed when you dial an Internet service provider (ISP) and use Internet utilities, such as a Web browser or FTP, to access Internet resources.

**Solution:** If you want to ping or otherwise connect to computers in a remote network segment across a local router while you are connected as a dial-up client to an ISP, then use the **route add** command to add the route of the network segment you are attempting to reach by specifying the destination network and subnet mask of the remote network segment and the IP address of the local router.

**See also:** Route

**A "The registry subkey already exists" message is displayed.**
**Cause:** The registry keys for a component being installed already exist.

**Solution:** Ensure that all the components of a given TCP/IP service are properly removed, and then remove the appropriate registry subkeys by using the Registry Editor.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**TCP/IP registry keys**

If you removed TCP/IP and its related service components, you must also remove the following registry subkeys:

- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NetBT**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dhcp**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LmHosts**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\AdapterName\Parameters\Tcpip**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\NetBT**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tcpip**

- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\TcpipCU**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\WinSock**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\WinSock2**

**SNMP service registry keys**

If you removed the SNMP service components, you must also remove the following registry subkeys:

- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\RFC1156Agent**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Snmp**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Snmp**

**TCP/IP Printing service registry keys**

If you removed the TCP/IP Printing service components, you must also remove the following registry subkeys:

- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Lpdsvc**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\TcpPrint**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LPDSVC**

**Simple TCP/IP Services registry keys**

If you removed the Simple TCP/IP Services components, you must also remove the following registry subkeys:

- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\SimpTcp**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\SimpTcp**

**DHCP service registry keys**

If you removed the DHCP service components, you must also remove the following registry subkeys:

- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\DhcpMibAgent**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\DhcpServer**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DhcpServer**

**WINS service registry keys**

If you removed the WINS service components, you must also remove the following registry subkeys:

- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Wins**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\WinsMibAgent**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Wins**

**DNS service registry keys**

If you removed the DNS service components, you must also remove the following registry subkeys:

- **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Dns**
- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dns**

For more advanced TCP/IP troubleshooting information, see Resource Kit.

# IP Version 6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP version 6**

IP version 6 (IPv6) is a suite of standard protocols that is the next generation of network layer protocols for the Internet.

- Before you begin using IPv6 with this computer, see Checklist: Installing and configuring IPv6.
- For help with specific tasks, see IPv6 How To....
- For general background information, see IPv6 Concepts.
- For problem-solving instructions, see Troubleshooting IPv6.

# Checklist: Installing and configuring IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring IPv6**

| Step | Reference |
|---|---|
| Review IPv6 concepts, such as core protocols, addressing, and routing. | Understanding IPv6 |
| Review IPv6 name resolution concepts for resolving DNS names to IP addresses. | Name Resolution |
| Verify that network adapter hardware is installed, or install hardware if needed. | Installing devices overview |
| Install IPv6 on this computer. | Install IPv6 |
| (Manual configuration only) Configure IPv6 address and prefix settings. | Configure IPv6 with manual addresses |
| Use the **ping** command to verify that you have IPv6 installed and running. | Test an IPv6 configuration by using the ping command |
| Use the **ping** command to verify that you can communicate with another IPv6 node. | Test IPv6 connectivity by using the ping command |

# IPv6 How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**
This section provides instructions for installing and configuring the IPv6 protocol.

- Install and Configure IPv6
- Managing Routes with IPv6
- Diagnose IPv6 Connections

# Install and Configure IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and configure IPv6**

- Install IPv6
- Remove IPv6
- Configure IPv6

# Install IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install IPv6**

1. Open Network Connections.
2. Right-click any local area connection, and then click **Properties**.
3. Click **Install**.
4. In the **Select Network Component Type** dialog box, click **Protocol**, and then click **Add**.
5. In the **Select Network Protocol** dialog box, click **Microsoft TCP/IP version 6**, and then click **OK**.
6. Click **Close** to save changes to your network connection.

### Note

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
IPv6 configuration items
IPv6 configuration methods
Remove IPv6
IPv6 utilities

# Remove IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove IPv6**

1. Open Network Connections.
2. Right-click any local area connection, and then click **Properties**.
3. Click **Microsoft TCP/IP version 6** in the list of installed components, and then click **Uninstall**.
4. In the **Uninstall Microsoft TCP/IP version 6** dialog box, click **Yes**.
5. Click **Close** to save changes to your network connection.

### Note

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install IPv6
Configure IPv6
IPv6 utilities

# Configure IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure IPv6**

- Configure IPv6 with manual addresses
- Configure interface attributes

# Configure IPv6 with manual addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure IPv6 with manual addresses**

1. Open Command Prompt.
2. At the command prompt, type **netsh**, and then press ENTER.
3. Type **interface ipv6**, and then press ENTER.
4. Type **add address** [**interface=**]*String* [**address=**]*IPv6Address*

Where:

**[ interface=] *String***
Specifies the name for the interface.

**[ address=] *IPv6Address***
Specifies the IPv6 address.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Additional parameters are available for this command. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
IPv6 configuration items
IPv6 configuration methods
IPv6 utilities
Netsh commands for Interface IPv6

# Configure interface attributes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure interface attributes**

1. Open Command Prompt.
2. At the command prompt, type **netsh**, and then press ENTER.
3. Type **interface ipv6**, and then press ENTER.
4. Type **set interface** [**interface=**]*String* [[**forwarding=**]{**enabled** | **disabled**}] [[**advertise=**]{**enabled** | **disabled**}] [[**mtu=**]*Integer*] [[**siteid=**]*Integer*] [[**metric=**]*Integer*] [[**firewall=**]{**enabled** | **disabled**}] [[**siteprefixlength=**]*Integer*] [[**store=**]{**active** | **persistent**}]

Where:

**[ interface=] *String***
Specifies the interface name.

**[[ forwarding=]{ enabled| disabled}]**
Specifies whether packets arriving on this interface can be forwarded to other interfaces. The default is disabled.

**[[ advertise=]{ enabled| disabled}]**
Specifies whether Router Advertisements are sent on this interface. The default is disabled.

**[[ mtu=] *Integer*]**
Specifies the maximum transmission unit (MTU) of this interface. If **mtu** is not specified, the default MTU of the link is used.

**[[ siteid=] *Integer*]**
Specifies the site scope zone identifier. The site identifier is used to distinguish among interfaces belonging to different administrative regions that use site-local addressing.

**[[ metric=] *Integer*]**
Specifies the interface metric, which is added to route metrics for all routes over the interface.

**[[ firewall=]{ enabled| disabled}]**
Specifies whether to operate in firewall mode.

**[[ siteprefixlength=] *Integer*]**
Specifies the default length of the global prefix for the entire site.

**[[ store=]{ active| persistent}]**
If **active** is specified, the change only lasts until the computer is restarted. If **persistent** is specified, the change is permanent. The default is **Persistent**.

**Note**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

IPv6 routing
Neighbor Discovery (ND)
IPv6 utilities
Netsh commands for Interface IPv6

# Managing Routes with IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage routes**

- View the IPv6 routing table
- Add an IPv6 route
- Remove an IPv6 route
- Enable IPv6 forwarding

# View the IPv6 routing table

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the IPv6 routing table**

1. Open Command Prompt.
2. At the command prompt, type **netsh**, and then press ENTER.
3. Type **interface ipv6**, and then press ENTER.
4. Type **show routes**, and then press ENTER

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Additional parameters are available for this command. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
IPv6 routing
IPv6 utilities
The IPv6 routing table
Add an IPv6 route
Remove an IPv6 route
Enable IPv6 forwarding
Netsh commands for Interface IPv6

# Add an IPv6 route

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an IPv6 route**

1. Open Command Prompt.
2. At the command prompt, type **netsh**, and then press ENTER.
3. Type **interface ipv6**, and then press ENTER.
4. Type **add route** [**prefix=**]*IPv6Address***/***Integer* [[**interface=**]*String*] [[**nexthop=**]*IPv6Address*] [[**siteprefixlength=**]*Integer*] [[**metric=**]*Integer*] [[**publish=**]{**no** | **yes** | **immortal**}] [[**validlifetime=**]{*Integer* | **infinite**}] [[**preferredlifetime=**]{*Integer* | **infinite**}] [[**store=**]{**active** | **persistent**}]

Where:

**[ prefix=] *IPv6Address/Integer***
Required. Specifies the prefix for which to add a route. *Integer* specifies the prefix length.

**[[ interface=] *String*]**
Specifies an interface name or index.

**[[ nexthop=] *IPv6Address*]**
Specifies the gateway address, if the prefix is not on-link.

**[[ siteprefixlength=] *Integer*]**
Specifies the prefix length for the entire site, if the prefix is not on-link.

**[[ metric=] *Integer*]**
Specifies the route metric.

**[[ publish=]{ no| yes| immortal}]**
Specifies whether routes are advertised (**yes**), advertised with an infinite lifetime (**immortal**), or not advertised (**no**) in Route Advertisements. The default selection is **no**.

**[[ validlifetime=]{ *Integer*| infinite}]**
Specifies the lifetime over which the route is valid. The default value is **infinite**.

**[[ preferredlifetime=]{ *Integer*| infinite}]**
Specifies the lifetime over which the route is preferred. The default value is **infinite**.

**[[ store=]{ active| persistent}]**
Specifies whether the change lasts only until the next boot (**active**) or is persistent (**persistent**). The default selection is **persistent**.

**Note**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

IPv6 routing
IPv6 utilities
The IPv6 routing table
Remove an IPv6 route
Enable IPv6 forwarding
View the IPv6 routing table
Netsh commands for Interface IPv6

# Remove an IPv6 route

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove an IPv6 route**

1. Open Command Prompt.
2. At the command prompt, type **netsh**, and then press ENTER.
3. Type **interface ipv6**, and then press ENTER.
4. Type **show routes**

   to obtain the route prefix and the interface index of the interface over which the addresses for the route prefix are reachable.
5. To delete a route, type**delete route** [**prefix=**]*IPv6Address***/***Integer* [**interface=**]*String*

Where:

**[ prefix=] *IPv6Address/Integer***
Specifies the prefix for which to delete a route. *IPv6Address* is an IPv6 address and *Integer* is the prefix length of the route to delete.

**[ interface=] *String***
Specifies the interface name or index.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Additional parameters are available for this command. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
IPv6 routing
IPv6 utilities
The IPv6 routing table
Add an IPv6 route
Enable IPv6 forwarding
View the IPv6 routing table
Netsh commands for Interface IPv6

# Enable IPv6 forwarding

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable IPv6 forwarding**

1. Open Command Prompt.
2. At the command prompt, type **netsh**, and then press ENTER.
3. Type **interface ipv6**, and then press ENTER.
4. Type **set interface** [**interface=**]*String* [**forwarding=**]**enabled**

Where:

**[ interface=] *String***
Specifies the interface name or index.

**[ forwarding=] enabled**
Specifies whether packets arriving on this interface can be forwarded to other interfaces. The default is disabled.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- You can also send Router Advertisement messages by adding the **advertise** parameter to the command, as follows:

  **set interface** [**interface=**]*String* [**forwarding=**]**enabled** [**advertise=**]**enabled**

- Additional parameters are available for this command. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
IPv6 routing
IPv6 utilities
The IPv6 routing table
Add an IPv6 route
Remove an IPv6 route
View the IPv6 routing table
Configure interface attributes

# Diagnose IPv6 Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Diagnose connections**

- Test an IPv6 configuration by using the ping command
- Test IPv6 connectivity by using the ping command
- Trace an IPv6 path by using the tracert command
- View the interface configuration
- View the neighbor cache
- View the destination cache

# Test an IPv6 configuration by using the ping command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test an IPv6 configuration by using the ping command**

1. To obtain the IPv6 configuration for a computer, open Command Prompt, and then type **netsh interface ipv6 show interface**. If IPv6 is not installed, see Related Topics.

2. At the command prompt, ping the loopback address by typing **ping ::1**.

   If the **ping** command fails, verify that the ::1 address is assigned to the interface named Loopback Pseudo-Interface.

3. Use the following command to ping a link-local IPv6 address of the computer:

   **ping***Address***%***ZoneID*

   where *Address* is the link-local address and *ZoneID* is the interface index for the interface to which the link-local address is assigned. A link-local address begins with FE80.

   If the **ping** command fails, verify the address and interface index.

4. Use the following command to ping the link-local address of another host on your link (also known as a subnet):

   **ping***Address***%***ZoneID*

   where *Address* is the link-local address of the other host and *ZoneID* is the interface index for the interface from which you want to send the ping packets.

   If the **ping** command fails, verify the link-local address of the other host and the zone ID.

### Note

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Control Message Protocol for IPv6 (ICMPv6)
IPv6 Name resolution
Single subnet with link-local addresses
IPv6 traffic between nodes on different subnets of an IPv6 internetwork
Troubleshooting IPv6
IPv6 utilities
Trace an IPv6 path by using the tracert command
View the interface configuration
Install IPv6
Test IPv6 connectivity by using the ping command

# Test IPv6 connectivity by using the ping command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test IPv6 connectivity by using the ping command**

1. To obtain the IPv6 configuration for a computer, open Command Prompt, and then type **netsh interface ipv6 show interface***InterfaceName*, where *InterfaceName* is the name of an interface on your computer. For example, if you have an interface named Local Area Connection, type:

   **netsh interface ipv6 show interface "Local Area Connection"**

   If IPv6 is not installed, see Related Topics.

2. Use one of the following commands to ping another IPv6 node:

   - To ping the link-local address of another node on your link (also known as a subnet):

     **ping** *Address***%***ZoneID*

     where *Address* is the link-local address of the other node and *ZoneID* is typically the interface index for the interface from which you want to send ping packets. You can obtain the interface index from the display of the **netsh interface ipv6 show interface** command.

     If the **ping** command fails, verify the link-local address of the other node and the zone ID.

   - To ping the site-local address of another node:

     **ping** *Address***%***ZoneID*

     where *Address* is the site-local address of the other node and *ZoneID* is the site identifier (Zone ID for Site) from the display of the **netsh interface ipv6 show interface level=verbose** command. If you are not using site identifiers or have only a single site, the **%***ZoneID* portion of the command is not required.

     If the **ping** command fails, verify the site-local address of the other node and the zone ID.

   - To ping the global address of another node:

     **ping***Address*

     where *Address* is the global address of the other node.

     If the **ping** command fails, verify the global address of the other node.

   - To ping another node by name:

     **ping -6***Name*

     where *Name* is a name that can be resolved to an IPv6 address through entries in the local hosts file, or through AAAA resource records that are present in your Domain Name System (DNS) infrastructure. When you identify the target host by name rather than IPv6 address, it is necessary to include the **-6** parameter to ensure that IPv6 addresses are used.

     If the **ping** command fails, verify that the name can be resolved to an IPv6 address.

**Note**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Control Message Protocol for IPv6 (ICMPv6)
IPv6 Name resolution
Single subnet with link-local addresses
IPv6 traffic between nodes on different subnets of an IPv6 internetwork
Troubleshooting IPv6
IPv6 utilities
Trace an IPv6 path by using the tracert command
View the interface configuration
Install IPv6
Test an IPv6 configuration by using the ping command

# Trace an IPv6 path by using the tracert command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To trace a path by using the tracert command**

- Open Command Prompt, and type the following:

  **tracert -d -6***HostName*

  Or, type **tracert***IPv6Address***%***ZoneID*

where:

- *HostName* is the host name of the remote computer.
- *IPv6Address* is the IPv6 address of the remote computer.
- *ZoneID* is the zone identifier (ID) for the destination address. The zone ID for link-local destination addresses is typically the interface index of the interface from which you want to send **tracert** packets. You can obtain the interface index from the display of the **netsh interface ipv6 show interface** command. The zone ID for site-local destination addresses is the Zone ID for Site from the display of the **netsh interface ipv6 show interface level=verbose** command. If you are not using site identifiers or have only a single site, the **%***ZoneID* portion of the command for site-local addresses is not required. The **%***ZoneID* portion of the command is not required for global destination addresses.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- For IPv6, the **tracert** command traces the path taken by IPv6 packets from this computer to another remote computer. The **tracert** command uses ICMPv6 Echo Request messages (similar to the **ping** command) to produce command-line report information about each router that is crossed and the roundtrip time for each hop.
- If **tracert** is unsuccessful, you can use the command-line report information to determine at which intermediate router forwarding either failed or was slowed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Control Message Protocol for IPv6 (ICMPv6)
IPv6 Name resolution
Single subnet with link-local addresses
IPv6 traffic between nodes on different subnets of an IPv6 internetwork
Troubleshooting IPv6
IPv6 utilities
Test an IPv6 configuration by using the ping command
View the interface configuration

# View the interface configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the interface configuration**

1. Open Command Prompt.
2. At the command prompt, type **netsh**, and then press ENTER.
3. Type **interface ipv6**, and then press ENTER.
4. Type **show interface** [**interface=**]*String*

Where:

**[ interface=] *String***
Specifies the interface name or index.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Additional parameters are available for this command. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
IPv6 interface identifiers
Single subnet with link-local addresses
Troubleshooting IPv6
IPv6 utilities
Configure interface attributes
Netsh commands for Interface IPv6

# View the neighbor cache

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the neighbor cache**

1. Open Command Prompt.
2. At the command prompt, type **netsh**, and then press ENTER.
3. Type **interface ipv6**, and then press ENTER.
4. Type **show neighbors**, and then press ENTER.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Additional parameters are available for this command. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Neighbor Discovery (ND)
IPv6 test lab tasks
IPv6 utilities
Troubleshooting IPv6
View the destination cache
Netsh overview
Netsh commands for Interface IPv6

# View the destination cache

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the destination cache**

1. Open Command Prompt.
2. At the command prompt, type **netsh**, and then press ENTER.
3. Type **interface ipv6**, and then press ENTER.
4. Type **show destinationcache**, and then press ENTER.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Additional parameters are available for this command. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Neighbor Discovery (ND)
IPv6 test lab tasks
IPv6 utilities
Troubleshooting IPv6
View the neighbor cache
Netsh overview
Netsh commands for Interface IPv6

# IPv6 Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about IPv6:

- IPv6 Overview
- Understanding IPv6
- Administering IPv6
- IPv6 Resources

# IPv6 Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 overview**

This section covers:

- Security information for IPv6
- Introduction to IPv6
- IPv6 features
- IPv6 protocol features
- Security features for IPv6

# Security information for IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for IPv6**

The use of IPv6 is currently supported only when IPv4 is also installed. TCP/IP internetworks are susceptible to a variety of possible attacks, from passive attacks (such as eavesdropping) to active attacks (such as denial-of-service attacks). For more information, see Security issues with IP.

It is important to follow best practices for security when using IPv6 on your network. For more information, see Best practices for security.

Following is known security information for IPv6:

- **The installation of an unauthorized router can cause reconfiguration of clients and rerouting of IPv6 traffic**

  Configuration of IPv6 addresses is based on the receipt of Router Advertisement messages from IPv6 routers. In addition, to communicate with IPv6 nodes on other network segments, IPv6 must use a default router. A default router is automatically assigned based on the receipt of a router advertisement.

  Malicious users with physical access to the IPv6-enabled network can cause a denial of service attack on IPv6 hosts by installing an unauthorized IPv6 router on the network segment. The unauthorized IPv6 router can reconfigure IPv6 clients, set itself as the default router, reroute link traffic, and disrupt other network services.

  **Recommendations:**

  - Ensure that unauthorized persons do not have physical or wireless access to your network.

- **Internet Connection Firewall (ICF) and Basic Firewall cannot filter or block IPv6 traffic**

  ICF, available on the 32-bit editions of Microsoft® Windows Server® 2003, Standard Edition and Windows Server 2003, Enterprise Edition, is used to restrict the traffic that is allowed to enter your network from the Internet. Because ICF can only filter IPv4 traffic, IPv6 traffic might circumvent the firewall and enter your network.

  Basic Firewall is a component of Routing and Remote Access that you can enable for any public interface on a computer running both Routing and Remote Access and a member of the Windows Server 2003 family. Because Basic Firewall can only filter IPv4 traffic, IPv6 traffic circumvent the firewall and enter your network.

  **Recommendations:**

  - Use firewall software that can filter and block IPv6 traffic if you are running IPv6 on your network.

- **On-link computers can take control of another IPv6 address, causing on-link devices to create an incorrect entry in their neighbor cache**

  Nodes on an IPv6 link use address resolution to resolve a neighboring node's IPv6 address to its link-layer address (equivalent to ARP in IPv4). The resolved link-layer address becomes an entry in a node's neighbor cache (equivalent to the ARP cache in IPv4). If an IPv6 node uses another node's address, it can cause other computers on the link to add a false entry to their neighbor cache. All traffic intended for the original computer instead goes to the attacker's computer, and the attacker can appear to send traffic from the original computer.

  **Recommendations:**

  - Ensure that unauthorized persons do not have physical or wireless access to your network.

- **When native IPv6 connectivity is not present, spoofing off-link IPv6 source addresses is easier**

The common defense against IP source address spoofing is source ingress and egress filtering on packets by routers. Because traffic between hosts on the same link does not cross a router, this protective filtering is not used.

While this threat applies to native connectivity (in both IPv4 and IPv6) where on-link hosts can spoof off-link addresses in communication with other hosts on the link, the threat is greater when native connectivity does not exist. When native IPv6 connectivity is not present, encapsulation technologies such as ISATAP and 6to4 are used.

Because the virtual link used for packet encapsulation spans a large logical area of the network (for example, in the case of 6to4, the entire IPv4 Internet), an attacker can be anywhere on the IPv4 Internet and still spoof off-link addresses.

**Recommendations:**

- Ensure that native IPv6 connectivity is present.

## Note

- When the computer running IPv6 is started, IPv6 logs event ID 3100 in the system event log.

### Additional information

- This implementation of IPSec for IPv6 is not recommended for use in a production environment because it relies on static keying and has no provisions for updating keys upon sequence number reuse.
- When you manually configure Security Parameters Indexes (SPIs) for IPSec for IPv6, always use random numbers. Do not use sequential numbers for SPIs, or you will compromise the security of your IPSec for IPv6 policies.
- The IPv6 protocol for the Windows Server 2003 family does not support the use of IPSec Encapsulating Security Payload (ESP) encryption. However, the use of ESP with NULL encryption is supported. Although NULL encryption uses the ESP header, only data origin authentication and data integrity services are provided.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

For additional information about IPv6 security, see Security features for IPv6 and Using IPSec between two local link hosts.

# Introduction to IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introduction to IPv6**

The current version of the Internet Protocol (known as IP version 4 or IPv4) has not been substantially changed since RFC 791 was published in 1981. IPv4 has proven to be robust, easily implemented and interoperable, and has stood the test of scaling an internetwork to a global utility the size of today's Internet. This is a tribute to its initial design.

However, the initial design did not anticipate:

- The recent exponential growth of the Internet and the impending exhaustion of the IPv4 address space.

  IPv4 addresses have become relatively scarce, forcing some organizations to use a network address translator (NAT) to map multiple private addresses to a single public IP address. While NATs promote reuse of the private address space, they do not support standards-based network layer security or the correct mapping of all higher layer protocols and can create problems when connecting two organizations that use the private address space.

  Additionally, the rising prominence of Internet-connected devices and appliances assures that the public IPv4 address space will eventually be depleted.

- The growth of the Internet and the ability of Internet backbone routers to maintain large routing tables.

  Because of the way in which IPv4 network IDs have been and are currently allocated, there are routinely over 70,000 routes in the routing tables of Internet backbone routers. The current IPv4 Internet routing infrastructure is a combination of both flat and hierarchical routing.

- The need for simpler configuration.

  Most current IPv4 implementations must be configured either manually or through a stateful address configuration protocol such as Dynamic Host Configuration Protocol (DHCP). With more computers and devices using IP, there is a need for a simpler and more automatic configuration of addresses and other configuration settings that do not rely on the administration of a DHCP infrastructure.

- The requirement for security at the IP level.

  Private communication over a public medium like the Internet requires encryption services that protect the data sent from being viewed or modified in transit. Although a standard now exists for providing security for IPv4 packets (known as Internet Protocol security or IPSec), this standard is optional and proprietary solutions are prevalent.

- The need for better support for real-time delivery of data (also known as quality of service).

  While standards for quality of service (QoS) exist for IPv4, real-time traffic support relies on the IPv4 Type of Service (TOS) field and the identification of the payload, typically using a UDP or TCP port. Unfortunately, the IPv4 TOS field has limited functionality and has different interpretations. In addition, payload identification using a TCP and UDP port is not possible when the IPv4 packet payload is encrypted.

To address these concerns, the Internet Engineering Task Force (IETF) has developed a suite of protocols and standards known as IP version 6 (IPv6). This new version, previously named IP-The Next Generation (IPng), incorporates the concepts of many proposed methods for updating the IPv4 protocol. IPv6 is intentionally designed for minimal impact on upper and lower layer protocols by avoiding the arbitrary addition of new features.

# IPv6 features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 features**
The following are the features of the IPv6 protocol:

- New header format
- Large address space
- Efficient and hierarchical addressing and routing infrastructure
- Stateless and stateful address configuration
- Built-in security
- Better support for quality of service (QoS)
- New protocol for neighboring node interaction
- Extensibility

The following sections discuss each of these new features in detail.

**New header format**
The IPv6 header has a new format that is designed to minimize header overhead. This is achieved by moving both nonessential fields and option fields to extension headers that are placed after the IPv6 header. The streamlined IPv6 header provides more efficient processing at intermediate routers.

IPv4 headers and IPv6 headers are not interoperable and the IPv6 protocol is not backward compatible with the IPv4 protocol. A host or router must use an implementation of both IPv4 and IPv6 in order to recognize and process both header formats. The new IPv6 header is only twice as large as the IPv4 header, even though IPv6 addresses are four times as large as IPv4 addresses.

For more information, see Internet Protocol version 6 (IPv6).

**Large address space**
IPv6 has 128-bit (16-byte) source and destination addresses. Although 128 bits can provide over 3.4×1038 possible combinations, the large address space of IPv6 has been designed to allow for multiple levels of subnetting and address allocation from the Internet backbone to the individual subnets within an organization.

Although only a small percentage of possible addresses are currently allocated for use by hosts, there are plenty of addresses available for future use. With a much larger number of available addresses, address-conservation techniques, such as the deployment of NATs, are no longer necessary.

For more information, see IPv6 Addressing.

**Efficient and hierarchical addressing and routing infrastructure**
IPv6 global addresses used on the IPv6 portion of the Internet are designed to create an efficient, hierarchical, and summarizable routing infrastructure that addresses the common occurrence of multiple levels of Internet service providers. On the IPv6 Internet, backbone routers have much smaller routing tables.

For more information, see Unicast IPv6 addresses.

**Stateless and stateful address configuration**
To simplify host configuration, IPv6 supports both stateful address configuration, such as address configuration in the presence of a DHCP server, and stateless address configuration (address configuration in the absence of a DHCP server). With stateless address configuration, hosts on a link automatically configure themselves with IPv6 addresses for the link (link-local

addresses) and with addresses that are derived from prefixes advertised by local routers. Even in the absence of a router, hosts on the same link can automatically configure themselves with link-local addresses and communicate without manual configuration.

For more information, see IPv6 address autoconfiguration.

**Built-in security**
Support for IPSec is an IPv6 protocol suite requirement. This requirement provides a standards-based solution for network security needs and promotes interoperability between different IPv6 implementations.

For more information, see Security features for IPv6.

**Better support for quality of service (QoS)**
New fields in the IPv6 header define how traffic is handled and identified. Traffic identification, by using a Flow Label field in the IPv6 header, allows routers to identify and provide special handling for packets that belong to a flow. A flow is a series of packets between a source and destination. Because the traffic is identified in the IPv6 header, support for QoS can be easily achieved even when the packet payload is encrypted with IPSec.

**New protocol for neighboring node interaction**
The Neighbor Discovery protocol for IPv6 is a series of Internet Control Message Protocol for IPv6 (ICMPv6) messages that manage the interaction of neighboring nodes (that is, nodes on the same link). Neighbor Discovery replaces Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and ICMPv4 Redirect messages with efficient multicast and unicast messages and provides additional functionality.

For more information, see Neighbor Discovery (ND).

**Extensibility**
IPv6 can be extended for new features by adding extension headers after the IPv6 header. Unlike the IPv4 header, which can only support 40 bytes of options, the size of IPv6 extension headers is only constrained by the size of the IPv6 packet.

# IPv6 protocol features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 protocol features**

The IPv6 protocol for the Microsoft® Windows Server 2003 family includes the following features:

- 6to4 tunneling
- Intrasite Automatic Tunnel Addressing Protocol
- Temporary addresses
- Site prefixes in router advertisements
- DNS support
- IPSec support
- Application support
- RPC support
- Static router support
- Command-line tool support
- 6over4 tunneling

**6to4 tunneling**

6to4 is a tunneling technique that is described in RFC 3056. 6to4 hosts do not require any manual configuration and create 6to4 addresses by using standard autoconfiguration. 6to4 uses the global address prefix of 2002:*WWXX*:*YYZZ*::/48, where *WWXX*:*YYZZ* is the colon-hexadecimal representation of a public IPv4 address (*w.x.y.z*) that is assigned to a site or host. WWXX:YYZZ is the Next Level Aggregator (NLA) portion of a 6to4 address.

6to4 allows IPv6-enabled sites and hosts to communicate by using unicast IPv6 the Internet. IPv6 sites and hosts can use their 6to4 address prefix and the Internet to communicate without obtaining an IPv6 global address prefix from an Internet service provider and connecting to the IPv6 Internet.

For more information, see IPv6 traffic between nodes in different sites across the Internet (6to4).

**Intrasite Automatic Tunnel Addressing Protocol**

Intrasite Automatic Tunnel Addressing Protocol (ISATAP) is an address assignment and tunneling mechanism for communication between IPv6/IPv4 nodes within an IPv4 site. It is described in the Internet draft titled "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)" (draft-ietf-ngtrans-isatap-0*x*.txt). For more information, see IPv6 traffic between nodes on different subnets of an IPv4 internetwork.

**Temporary addresses**

To provide a level of anonymity when accessing Internet resources, the 64-bit interface identifier of a global IPv6 address is derived by using random numbers to create a temporary global address.

For more information, see IPv6 interface identifiers.

**Site prefixes in router advertisements**

Published on-link prefixes can be configured with a site prefix length, as described in the Internet draft titled "Site prefixes in Neighbor Discovery" (draft-ietf-ipngwg-site-prefixes-0*x*.txt). You can use the **netsh interface ipv6 add route** command to include a site prefix length with the address prefix. For more information, see Netsh commands for Interface IPv6.

When a prefix information option that specifies a site prefix is received, an entry is created in the site prefix table. You can view this table by using the **netsh interface ipv6 show siteprefixes** command. The site prefix table is used to remove inappropriate site-local addresses from those that are returned by the **getaddrinfo()** Windows sockets function.

**DNS support**
Processing for Domain Name System (DNS) IPv6 host records (known as AAAA or quad-A resource records), as defined in RFC 1886, "DNS Extensions to support IP version 6," is supported by the DNS resolver (client) in Microsoft® Windows® XP and the Windows Server 2003 family, and DNS Server service in the Windows Server 2003 family and Windows 2000. All DNS traffic is sent over IPv4 or IPv6.

For more information, see IPv6 Name resolution.

**IPSec support**
Processing for the Authentication Header (AH) using the Message Digest 5 (MD5) hash, and the Encapsulating Security Payload (ESP) using the NULL ESP header and the MD5 hash, is supported. There is no support for ESP data encryption.

### Caution

- This implementation of IPSec for IPv6 is not recommended for use in a production environment because it relies on static keying and has no provisions for updating keys upon sequence number reuse.

For an example configuration, see Using IPSec between two local link hosts.

**Application support**
Applications and components that are provided with the Windows Server 2003 family and support the use of IPv6 include:

- Distributed Component Object Model (DCOM)
- Distributed File System (DFS)
- Domain Name System (DNS) server and client
- File and print sharing
- File Transfer Protocol (FTP) client
- Internet Explorer
- Microsoft .NET Framework
- Network Monitor
- Telnet Server and Telnet Client
- Windows Media Services
- Windows Sockets

For more information, see IPv6 applications.

**RPC support**
Remote Procedure Call (RPC) functions, which are used to forward application function calls to a remote system across the network, can be used over IPv6. A typical use for RPC is remote administration.

**Static router support**
A computer running a member of the Windows Server 2003 family can act as a static IPv6 router that forwards IPv6 packets between interfaces based on the contents of the IPv6 routing table. You can configure static routes with the **netsh interface ipv6 add route** command. There is currently no support for IPv6 routing protocols.

A computer running a member of the Windows Server 2003 family can send router advertisements. The contents of router advertisements are automatically derived from the published routes in the routing table. Nonpublished routes are used for

routing but are not sent in router advertisements. Router advertisements always contain a source link-layer address option and an MTU option. The value for the MTU option is taken from the sending interface's current link MTU. You can change this value with the following Netsh command:

**netsh interface ipv6 set interface** [**interface=**]*String* [[**mtu=**]*Integer*]

A computer running a member of the Windows Server 2003 family will only advertise itself as a default router (by using a router advertisement with a router lifetime other than zero) if there is a default route that is configured to be published.

For more information and examples, see Netsh commands for Interface IPv6, IPv6 traffic between nodes on different subnets of an IPv6 internetwork, and IPv6 test lab tasks.

**Command-line tool support**

The following command-line tools are provided with the Windows Server 2003 family and support the use and configuration of IPv6:

- Netsh commands for Interface IPv6
- Netsh commands for Interface Portproxy
- Ipconfig
- Pathping
- Ping
- Route
- Tracert
- Netstat

The Netstat command-line tool provided with the Windows Server 2003 family and Windows XP includes support for IPv6-related statistics. There are four categories of statistics:

1. IPv6 statistics
2. ICMPv6 statistics
3. TCP connections over IPv6 statistics
4. UDP traffic over IPv6 statistics

**6over4 tunneling**

6over4, also known as IPv4 multicast tunneling, is a tunneling technique that is described in RFC 2529. 6over4 allows IPv6 and IPv4 nodes to communicate by using IPv6 over an IPv4 infrastructure. 6over4 uses the IPv4 infrastructure as a multicast-capable link. In order for 6over4 to work correctly, the IPv4 infrastructure must be IPv4 multicast-enabled.

For more information about configuring a 6over4 interface, see the command **add 6over4tunnel** in Netsh commands for Interface IPv6.

# Security features for IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security features for IPv6**

The IPv6 protocol for Microsoft® Windows Server 2003 family incorporates Internet Protocol security (IPSec), which provides protection of IPv6 data as it is sent over the network. IPSec is a set of Internet standards that uses cryptographic security services to provide the following:

- Confidentiality

  IPSec traffic is encrypted. Captured IPSec traffic cannot be deciphered without the encryption key.

- Authentication

  IPSec traffic is digitally signed with the shared encryption key so that the receiver can verify that it was sent by the IPSec peer.

- Data integrity

  IPSec traffic contains a cryptographic checksum that incorporates the encryption key. The receiver can verify that the packet was not modified in transit.

For an example of configuring IPSec for IPv6, see Using IPSec between two local link hosts.

The IPv6 protocol for Windows Server 2003 family also provides support for temporary addresses. Temporary addresses provide a level of anonymity when accessing Internet resources. For more information about temporary addresses, see the section "Temporary address interface identifiers" in IPv6 interface identifiers.

**Caution**

- This implementation of IPSec for IPv6 is not recommended for use in a production environment because it relies on static keying and has no provisions for updating keys upon sequence number reuse.

**Notes**

- IPSec in IPv6 is separate from and not interoperable with IPSec for the TCP/IP protocol. IPSec policies that are configured with the IP Security Policies or Group Policy snap-ins have no effect on IPv6 traffic. For more information about IPSec for the TCP/IP protocol, see Internet Protocol Security (IPSec).

- IPSec in the IPv6 protocol for Windows Server 2003 family does not support the use of data encryption for data confidentiality.

- IPSec in the IPv6 protocol for Windows Server 2003 family does not support the use of Internet Key Exchange (IKE) to negotiate security associations (SAs). IPSec policies, SAs, and keys must be manually configured.

# Understanding IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding IPv6**
This section covers:

- IPv6 Name resolution
- Core Protocols of IPv6
- IPv6 Addressing and Routing
- IPv6 and Network Applications

# IPv6 Name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Name resolution**

While IPv6 is designed to work with the 128-bit IPv6 addresses of the source and the destination hosts, computer users are likely to experience difficulty in using and remembering the IPv6 addresses of the computers with which they want to communicate. Unique names, which are easier to remember, can be used instead.

If a name is used as an alias for an IPv6 address, you need to ensure that the name is unique and that it resolves to the correct IPv6 address. The IPv6 protocol for the Windows Server 2003 family can use host names to resolve a name to an IPv6 address. Host names are used by programs that use Windows Sockets.

Host name resolution is successfully mapping a host name to an IPv6 address. A host name is an alias that is assigned to an IPv6 node, identifying it as an IPv6 host. The host name can be up to 255 characters long and can contain alphabetic and numeric characters, hyphens, and periods. You can assign multiple host names to the same host.

Windows Sockets (Winsock) programs can use one of two values for the destination to which you want to connect: the IPv6 address or a host name. When the IPv6 address is specified, name resolution is not required. When a host name is specified, the host name must be resolved to an IPv6 address before IPv6-based communication with a resource can begin.

Host names can take various forms. The two most common forms are a nickname and a domain name. A nickname is an alias for an IPv6 address that individuals can assign and use. A domain name is a structured name in a hierarchical namespace named Domain Name System (DNS). An example of a domain name is www.microsoft.com.

Nicknames or domain names are resolved through entries in the Hosts file, which is stored in the *systemroot*\System32\Drivers\Etc folder. For IPv6 name-to-address entries, the IPv6 address is written by using standard colon-hexadecimal format. For more information, see Expressing IPv6 addresses and TCP/IP database files.

Domain names are resolved by sending DNS name queries to a configured DNS server, which is a computer that either stores domain name-to-IPv6 address mapping records or has records of other DNS servers. The DNS server resolves the queried domain name to an IPv6 address and returns the results. The DNS client in Windows XP and the Windows Server 2003 family supports the processing of AAAA (quad-A) resource records. All DNS queries and responses are sent by using IPv6 and IPv4. DNS name devolution for fully qualified domain names is also supported. For more information, see DNS defined.

# Core Protocols of IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Core protocols of IPv6**
This section covers:

- Internet Protocol version 6 (IPv6)
- Internet Control Message Protocol for IPv6 (ICMPv6)
- Multicast Listener Discovery (MLD)
- Neighbor Discovery (ND)

# Internet Protocol version 6 (IPv6)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol version 6 (IPv6)**
IPv6 is defined in RFC 2460, "Internet Protocol, Version 6 (IPv6) Specification." IPv6 is a connectionless, unreliable datagram protocol used primarily for addressing and routing packets between hosts.

Connectionless means that a session is not established before exchanging data. Unreliable means that delivery is not guaranteed. IPv6 always makes a best-effort attempt to deliver a packet. An IPv6 packet might be lost, delivered out of sequence, duplicated, or delayed. IPv6 does not attempt to recover from these types of errors. The acknowledgment of packets delivered and the recovery of lost packets are done by a higher-layer protocol, such as TCP.

An IPv6 packet, also known as an IPv6 datagram, consists of an IPv6 header and an IPv6 payload, as shown in the following illustration.




The IPv6 header contains the following fields for addressing and routing.

| IPv6 header field | Function |
|---|---|
| Source Address | The IPv6 address of the original source of the IPv6 packet. |
| Destination Address | The IPv6 address of the intermediate or final destination of the IPv6 packet. |
| Hop Limit | The number of network segments on which the packet is allowed to travel before being discarded by a router. The Hop Limit is set by the sending host and is used to prevent packets from endlessly circulating on an IPv6 internetwork. When forwarding an IPv6 packet, IPv6 routers are required to decrease the Hop Limit by 1 and to discard the IPv6 packet when the Hop Limit is 0. |

**Note**

- IPv6 is a rapidly evolving standard. The RFCs referenced might have been made obsolete by newer RFCs.

# Internet Control Message Protocol for IPv6 (ICMPv6)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Control Message Protocol for IPv6 (ICMPv6)**
Internet Control Message Protocol for IPv6 (ICMPv6) is a required IPv6 standard defined in RFC 2463, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification." With ICMPv6, hosts and routers that use IPv6 communication can report errors and send simple echo messages.

The ICMPv6 protocol also provides a framework for the following:

- Multicast Listener Discovery (MLD)

  MLD is a series of three ICMPv6 messages that replace version 2 of the Internet Group Management Protocol (IGMP) for IPv4 to manage subnet multicast membership. For more information, see Multicast Listener Discovery (MLD).

- Neighbor Discovery (ND)

  Neighbor Discovery is a series of five ICMPv6 messages that manage node-to-node communication on a link. Neighbor Discovery replaces Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and the ICMPv4 Redirect message and provides additional functions. For more information, see Neighbor Discovery (ND).

ICMPv6 messages are usually sent automatically when an IPv6 packet cannot reach its destination.

ICMPv6 messages are encapsulated and sent as the payload of IPv6 packets, as shown in the following illustration.




Different types of ICMPv6 messages are identified in the ICMPv6 header. Because ICMPv6 messages are carried in IPv6 packets, they are unreliable.

ICMPv6 messages that are not related to MLD or ND are listed and described in the following table.

| ICMPv6 message | Description |
|---|---|
| Destination Unreachable | An error message that informs the sending host that a packet cannot be delivered. |
| Packet Too Big | An error message that informs the sending host that the packet is too large to forward. |
| Time Exceeded | An error message that informs the sending host that the Hop Limit of an IPv6 packet has expired. |
| Parameter Problem | An error message that informs the sending host that an error was encountered in processing the IPv6 header or an IPv6 extension header. |
| Echo Request | An informational message that is used to determine whether an IPv6 node is available on the network. |
| Echo Reply | An informational message that is used to reply to the ICMPv6 Echo Request message. |

You can use the **ping** command to send ICMPv6 Echo Request messages and record the receipt of ICMPv6 Echo Reply messages. With ping, you can detect network or host communication failures and troubleshoot common IPv6 connectivity problems. For more information, see Test IPv6 connectivity by using the ping command.

You can use the **tracert** command to send ICMPv6 Echo Request messages with incrementally increasing values in the Hop Limit field. Tracert will trace and display the path taken by IPv6 packets between a source and destination, allowing you to

troubleshoot common IPv6 routing problems. For more information, see Trace an IPv6 path by using the tracert command.

**Note**

- IPv6 is a rapidly evolving standard. The RFCs referenced might have been made obsolete by newer RFCs.

# Multicast Listener Discovery (MLD)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multicast Listener Discovery (MLD)**
The use of multicasting in IP networks is defined as a TCP/IP standard in RFC 1112, "Internet Group Management Protocol (IGMP)." This RFC defines address and host extensions for the way in which IP hosts support multicasting. The same concepts originally developed for the current version of IP, known as IP version 4 (IPv4), also apply to IPv6.

**What is multicasting?**
Multicast traffic is sent to a single address but is processed by multiple hosts. Multicasting is similar to a newsletter subscription. As only subscribers receive the newsletter when it is published, only host computers that belong to the multicast group receive and process traffic sent to the group's reserved address. The set of hosts listening on a specific multicast address is called a multicast group.

Other important aspects of multicasting include the following:

- Group membership is dynamic, allowing hosts to join and leave the group at any time.
- The joining of multicast groups is performed through the sending of group membership messages. In IPv6, Multicast Listener Discovery (MLD) messages are used to determine group membership on a network segment, also known as a link or subnet.
- Groups are not limited by size and members can be spread out across multiple network segments (if connecting routers support the forwarding of multicast traffic and group membership information).
- A host can send traffic to the group's address without belonging to the corresponding group.

**IPv6 multicast addressing**
IPv6 multicast addresses are reserved and assigned from the Format Prefix 1111 1111 (0xFF). The following table is a partial list of IPv6 multicast addresses that are reserved for IPv6 multicasting and registered with the Internet Assigned Numbers Authority (IANA).

| IPv6 multicast address | Description |
|---|---|
| FF02::1 | The all-nodes address used to reach all nodes on the same link. |
| FF02::2 | The all-routers address used to reach all routers on the same link. |
| FF02::4 | The all-Distance Vector Multicast Routing Protocol (DVMRP) routers address used to reach all DVMRP multicast routers on the same link. |
| FF02::5 | The all-Open Shortest Path First (OSPF) routers address used to reach all OSPF routers on the same link. |
| FF02::6 | The all-OSPF designated routers address used to reach all OSPF designated routers on the same link. |
| FF02::1:FF*XX:XXXX* | The solicited-node address used in the address resolution process to resolve the IPv6 address of a link-local node to its link-layer address. The last 24 bits (*XX:XXXX*) of the solicited-node address are the last 24 bits of an IPv6 unicast address. |

For a full and current listing of additional IPv6 addresses that are reserved for multicasting, see the Information Sciences Institute Web site document: Internet Protocol Version 6 Multicast Addresses.

A single IPv6 multicast address identifies each multicast group. Each group's reserved IPv6 address is shared by all host members of the group who listen and receive any IPv6 messages sent to the group's address.

IPv6 multicast addresses are mapped to a reserved set of media access control (MAC) multicast addresses. For information about the mapping of IPv6 multicast addresses to Ethernet MAC addresses, see RFC 2464, "Transmission of IPv6 Packets over Ethernet Networks."

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

**MLD messages**

MLD is used to exchange membership status information between IPv6 routers that support multicasting and members of multicast groups on a network segment. Host membership in a multicast group is reported by individual member hosts, and membership status is periodically polled by multicast routers. MLD is defined in RFC 2710, "Multicast Listener Discovery (MLD) for IPv6."

MLD message types are described in the following table.

| MLD message type | Description |
|---|---|
| Multicast Listener Query | Sent by a multicast router to poll a network segment for group members. Queries can be general (requesting group membership for all groups), or specific (requesting group membership for a specific group). |
| Multicast Listener Report | Sent by a host when it joins a multicast group, or in response to an MLD Multicast Listener Query sent by a router. |
| Multicast Listener Done | Sent by a host when it leaves a host group and might be the last member of that group on the network segment. |

MLD messages are sent as ICMPv6 messages.

**Note**

- IPv6 is a rapidly evolving standard. The RFCs referenced might have been made obsolete by newer RFCs.

# Neighbor Discovery (ND)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Neighbor Discovery (ND)**
IPv6 Neighbor Discovery (ND) is a set of messages and processes that determine relationships between neighboring nodes. ND replaces Address Resolution Protocol (ARP), Internet Control Message Protocol (ICMP) Router Discovery, and ICMP Redirect, which are used in IPv4, and provides additional functionality. ND is described in RFC 2461, "Neighbor Discovery for IP Version 6 (IPv6)."

ND is used by hosts to:

- Discover neighboring routers.
- Discover addresses, address prefixes, and other configuration parameters.

ND is used by routers to:

- Advertise their presence, host configuration parameters, and on-link prefixes.
- Inform hosts of a better next-hop address to forward packets for a specific destination.

ND is used by nodes to:

- Both resolve the link-layer address of a neighboring node to which an IPv6 packet is being forwarded and determine when the link-layer address of a neighboring node has changed.
- Determine whether IPv6 packets can be sent to and received from a neighbor.

The following table lists and describes ND processes.

| Process | Description |
|---|---|
| Router discovery | The process by which a host discovers the local routers on an attached link (equivalent to ICMPv4 Router Discovery) and automatically configures a default router (equivalent to a default gateway in IPv4). |
| Prefix discovery | The process by which a host discovers the network prefixes for local destinations. |
| Parameter discovery | The process by which a host discovers additional operating parameters, including the link maximum transmission unit (MTU) and the default hop limit for outbound packets. |
| Address autoconfiguration | The process for configuring IP addresses for interfaces in either the presence or absence of a stateful address configuration server such as Dynamic Host Configuration Protocol version 6 (DHCPv6). For more information, see IPv6 address autoconfiguration. |
| Address resolution | The process by which a node resolves a neighboring node's IPv6 address to its link-layer address (equivalent to ARP in IPv4). The resolved link-layer address becomes an entry in a node's neighbor cache (equivalent to the ARP cache in IPv4). You can use the **netsh interface ipv6 show neighbors** command to view the contents of the neighbor cache on a computer running the Windows Server 2003 family and Windows XP. The neighbor cache displays the interface identifier for the neighbor cache entry, the neighboring node's IPv6 address, the corresponding link-layer address, and the state of the neighbor cache entry. For more information, see View the neighbor cache and Netsh commands for Interface IPv6. |

| | |
|---|---|
| Next-hop determination | The process by which a node determines the IPv6 address of the neighbor to which a packet is being forwarded based on the destination address. The forwarding or next-hop address is either the destination address of the packet being sent or the address of a neighboring router. The resolved next hop address for a destination becomes an entry in a node's destination cache (also known as a route cache). You can use the **netsh interface ipv6 show destinationcache** command to view the contents of the destination cache on a computer running the Windows Server 2003 family and Windows XP. The destination cache displays the destination address, the interface identifier and next-hop address, the interface identifier and address used as a source address when sending to the destination, and the path MTU for the destination. For more information, see View the destination cache and Netsh commands for Interface IPv6. |
| Neighbor unreachability detection | The process by which a node determines that IPv6 packets cannot be sent to and received from a neighboring node. After the link-layer address for a neighbor has been determined, the state of the entry in the neighbor cache is tracked. If the neighbor is no longer receiving and sending back packets, the neighbor cache entry is eventually removed. Neighbor unreachability detection provides a mechanism for IPv6 to determine that neighboring hosts or routers are no longer available on the local network segment. |
| Duplicate address detection | The process by which a node determines that an address considered for use is not already in use by a neighboring node (equivalent to the use of gratuitous ARP frames in IPv4). |
| Redirect function | The process by which a router informs a host of a better first-hop IPv6 address to reach a destination (equivalent to the function of the IPv4 ICMP Redirect message). |

**Note**

- IPv6 is a rapidly evolving standard. The RFCs referenced might have been made obsolete by newer RFCs.

# IPv6 Addressing and Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 addressing and routing**
This section covers:

- IPv6 address autoconfiguration
- IPv6 routing
- The IPv6 routing table
- IPv6 Addressing

# IPv6 address autoconfiguration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 address autoconfiguration**
A highly useful aspect of IPv6 is its ability to automatically configure itself without the use of a stateful configuration protocol, such as Dynamic Host Configuration Protocol for IPv6 (DHCPv6). By default, an IPv6 host can configure a link-local address for each interface. By using router discovery, a host can also determine the addresses of routers, additional addresses, and other configuration parameters. Included in the Router Advertisement message is an indication of whether a stateful address configuration protocol should be used.

Address autoconfiguration can only be performed on multicast-capable interfaces. Address autoconfiguration is described in RFC 2462, "IPv6 Stateless Address Autoconfiguration."

**Autoconfigured address states**
Autoconfigured addresses are in one or more of the following states:

- Tentative

  The address is in the process of being verified as unique. Verification occurs through duplicate address detection.

- Preferred

  An address for which uniqueness has been verified. A node can send and receive unicast traffic to and from a preferred address. The period of time that an address can remain in the tentative and preferred states is included in the Router Advertisement message.

- Deprecated

  An address that is still valid, but its use is discouraged for new communication. Existing communication sessions can continue to use a deprecated address. A node can send and receive unicast traffic to and from a deprecated address.

- Valid

  An address from which unicast traffic can be sent and received. The valid state covers both the preferred and deprecated states. The amount of time that an address remains in the tentative and valid states is included in the Router Advertisement message. The valid lifetime must be greater than or equal to the preferred lifetime.

- Invalid

  An address for which a node can no longer send or receive unicast traffic. An address enters the invalid state after the valid lifetime expires.

The relationship between the states of an autoconfigured address and the preferred and valid lifetimes is shown in the following illustration.




**Note**

- With the exception of link-local addresses, address autoconfiguration is only specified for hosts. Routers must obtain address and configuration parameters through another means (for example, manual configuration).

**Types of autoconfiguration**
There are three types of autoconfiguration:

1. Stateless

   Configuration of addresses is based on the receipt of Router Advertisement messages. These messages include stateless address prefixes and require that hosts not use a stateful address configuration protocol.

2. Stateful

   Configuration is based on the use of a stateful address configuration protocol, such as DHCPv6, to obtain addresses and other configuration options. A host uses stateful address configuration when it receives Router Advertisement messages that do not include address prefixes and require that the host use a stateful address configuration protocol. A host will also use a stateful address configuration protocol when there are no routers present on the local link.

3. Both

   Configuration is based on receipt of Router Advertisement messages. These messages include stateless address prefixes and require that hosts use a stateful address configuration protocol.

For all autoconfiguration types, a link-local address is always configured.

**Autoconfiguration process**
The address autoconfiguration process for an IPv6 node occurs as follows:

1. A tentative link-local address is derived, based on the link-local prefix of FE80::/64 and the 64-bit interface identifier.

2. Duplicate address detection is performed to verify the uniqueness of the tentative link-local address.

3. If duplicate address detection fails, manual configuration must be performed on the node.

4. If duplicate address detection succeeds, the tentative link-local address is assumed to be unique and valid. The link-local address is initialized for the interface. The corresponding solicited-node multicast link-layer address is registered with the network adapter.

For an IPv6 host, address autoconfiguration continues as follows:

1. The host sends a Router Solicitation message.

2. If no Router Advertisement messages are received, then the host uses a stateful address configuration protocol to obtain addresses and other configuration parameters. The IPv6 protocol for the Windows Server 2003 family and Windows XP does not support the use of a stateful address configuration protocol.

3. If a Router Advertisement message is received, the configuration information that is included in the message is set on the host.

4. For each stateless autoconfiguration address prefix that is included:

   - The address prefix and the appropriate 64-bit interface identifier are used to derive a tentative address.

   - Duplicate address detection is used to verify the uniqueness of the tentative address.

If the tentative address is in use, the address is not initialized for the interface.

If the tentative address is not in use, the address is initialized. This includes setting the valid and preferred lifetimes based on

information included in the Router Advertisement message.

1. If it is specified in the Router Advertisement message, the host uses a stateful address configuration protocol to obtain additional addresses or configuration parameters.

**Notes**

- IPv6 is a rapidly evolving standard. The RFCs referenced might have been made obsolete by newer RFCs.

- Standards for stateful address configuration for IPv6 are in development. The IPv6 protocol for the Windows Server 2003 family and Windows XP does not support stateful address configuration or the DHCPv6 protocol.

# IPv6 routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 routing**

Routing is the process of forwarding packets between connected network segments. For IPv6-based networks, routing is the part of IPv6 that provides forwarding capabilities between hosts that are located on separate segments within a larger IPv6-based network.

IPv6 is the mailroom in which IPv6 data sorting and delivery occur. Each incoming or outgoing packet is called an IPv6 packet. An IPv6 packet contains both the source address of the sending host and the destination address of the receiving host. Unlike link-layer addresses, IPv6 addresses in the IPv6 header typically remain the same as the packet travels across an IPv6 network.

Routing is the primary function of IPv6. IPv6 packets are exchanged and processed on each host by using IPv6 at the Internet layer.

Above the IPv6 layer, transport services on the source host pass data in the form of TCP segments or UDP messages down to the IPv6 layer. The IPv6 layer creates IPv6 packets with source and destination address information that is used to route the data through the network. The IPv6 layer then passes packets down to the link layer, where IPv6 packets are converted into frames for transmission over network-specific media on a physical network. This process occurs in reverse order on the destination host.

IPv6 layer services on each sending host examine the destination address of each packet, compare this address to a locally maintained routing table, and then determine what additional forwarding is required. IPv6 routers are attached to two or more IPv6 network segments that are enabled to forward packets between them.

**IPv6 routers**

IPv6 network segments, also known as links or subnets, are connected by IPv6 routers, which are devices that pass IPv6 packets from one network segment to another. This process is known as IPv6 routing and is shown in the following illustration.




IPv6 routers provide the primary means for joining together two or more physically separated IPv6 network segments. All IPv6 routers have the following characteristics:

- IPv6 routers are physically multihomed hosts.

  A physically multihomed host is a network host that uses two or more network connection interfaces to connect to each physically separated network segment.

- IPv6 routers provide packet forwarding for other IPv6 hosts.

  IPv6 routers are distinct from other hosts that use multihoming. An IPv6 router must be able to forward IPv6-based communication between networks for other IPv6 network hosts.

You can implement IPv6 routers by using a variety of hardware and software products, including a computer running a member of the Windows Server 2003 family with the IPv6 protocol. Routers that are dedicated hardware devices running specialized software are common. Regardless of the type of IPv6 routers that you use, all IPv6 routing relies on a routing table to communicate between network segments.

**Routing tables**

IPv6 hosts use a routing table to maintain information about other IPv6 networks and IPv6 hosts. Network segments are identified by using an IPv6 network prefix and prefix length. In addition, routing tables provide important information for each

local host regarding how to communicate with remote networks and hosts.

For each computer on an IPv6 network, you can maintain a routing table with an entry for every other computer or network that communicates with that local computer. In general, this is not practical, and a default router is used instead.

Before a computer sends an IPv6 packet, it inserts its source IPv6 address and the destination IPv6 address (for the recipient) into the IPv6 header. The computer then examines the destination IPv6 address, compares it to a locally maintained IPv6 routing table, and takes appropriate action. The computer does one of three things:

- It passes the packet to a protocol layer above IPv6 on the local host.
- It forwards the packet through one of its attached network interfaces.
- It discards the packet.

IPv6 searches the routing table for the route that is the closest match to the destination IPv6 address. The most specific to the least specific route is determined in the following order:

1. A route that matches the destination IPv6 address (a host route with a 128-bit prefix length).
2. A route that matches the destination with the longest prefix length.
3. The default route (the network prefix ::/0).

If a matching route is not found, the destination is determined to be an on-link destination.

For more information, see The IPv6 routing table.

# The IPv6 routing table

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The IPv6 routing table**

Every computer that runs IPv6 determines how to forward packets based on the contents of the IPv6 routing table. To display the IPv6 routing table on a computer that is running a member of the Windows Server 2003 family or Windows XP, type **netsh interface ipv6 show routes** at the Command Prompt.

Entries in the IPv6 routing table consist of:

- An address prefix
- The interface over which packets matching the address prefix are sent
- A forwarding or next-hop address
- A preference value used to select between multiple routes with the same prefix
- The lifetime of the route
- The specification of whether the route is published (advertised in a Routing Advertisement)
- The specification of how the route is aged
- The route type

The IPv6 routing table is built automatically, based on the current IPv6 configuration of your computer. When forwarding IPv6 packets, the routing table is searched by your computer for an entry that is the most specific match to the destination IPv6 address. A route for the link-local prefix (FE80::/64) is not displayed.

The default route (a route with a prefix of ::/0) is typically used to forward an IPv6 packet to a default router on the local link. Because the router that corresponds to the default router contains information about the network prefixes of the other IPv6 subnets within the larger IPv6 internetwork, it forwards the packet to other routers until it is eventually delivered to the destination.

For information about managing IPv6 routes, see Managing Routes with IPv6.

# IPv6 Addressing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 addressing**
This section covers:

- IPv6 address space
- Expressing IPv6 addresses
- Unicast IPv6 addresses
- Multicast IPv6 addresses
- Anycast IPv6 addresses
- Addresses for hosts and routers
- IPv6 interface identifiers

# IPv6 address space

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 address space**

The most obvious distinguishing feature of IPv6 is its use of much larger addresses. The size of an address in IPv6 is 128 bits, which is four times larger than an address in IPv4. A 32-bit address space allows for 232 or 4,294,967,296 possible addresses. A 128-bit address space allows for 2128 or 340,282,366,920,938,463,463,374,607,431,768,211,456 (3.4 × 1038) possible addresses.

In the late 1970s when the IPv4 address space was designed, it was unimaginable that it could be exhausted. However, due to changes in technology and an allocation practice that did not anticipate the recent explosion of hosts on the Internet, the IPv4 address space was consumed to the point that by 1992, it was clear a replacement would be necessary.

With IPv6, it is even harder to conceive that the IPv6 address space will be consumed. To help put this number in perspective, a 128-bit address space provides 655,570,793,348,866,943,898,599 (6.5 × 1023) addresses for every square meter of the Earth's surface.

It is important to note that the decision to make the IPv6 address 128 bits in length was not so that every square meter of the Earth could have 6.5 x 1023 addresses. Rather, the relatively large size of the IPv6 address is designed to be subdivided into hierarchical routing domains that reflect the topology of the modern-day Internet. The use of 128 bits provides multiple levels of hierarchy and flexibility in designing hierarchical addressing and routing that is currently lacking on the IPv4-based Internet.

The architecture of IPv6 addressing is described in RFC 2373, "IP Version 6 Addressing Architecture."

**Current allocation**

Similar to the way in which the IPv4 address space is divided, the IPv6 address space is divided based on the value of high order bits in the address. The high order bits and their fixed values are known as a Format Prefix (FP).

The following table shows the allocation of the IPv6 address space by FPs.

| Allocation | Format Prefix (FP) | Fraction of the address space |
|---|---|---|
| Reserved | 0000 0000 | 1/256 |
| Reserved for NSAP allocation | 0000 001 | 1/128 |
| Aggregatable global unicast addresses | 001 | 1/8 |
| Link-local unicast addresses | 1111 1110 10 | 1/1024 |
| Site-local unicast addresses | 1111 1110 11 | 1/1024 |
| Multicast addresses | 1111 1111 | 1/256 |

The remainder of the IPv6 address space is unassigned.

The current set of unicast addresses that can be used with IPv6 nodes consists of aggregatable global unicast addresses, link-local unicast addresses, and site-local unicast addresses. These represent only 15 percent of the entire IPv6 address space. For more information, see Unicast IPv6 addresses.

**Note**

- IPv6 is a rapidly evolving standard. The RFCs referenced might have been made obsolete by newer RFCs.

# Expressing IPv6 addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Expressing IPv6 addresses**

IPv4 addresses are represented in dotted-decimal format. This 32-bit address is divided along 8-bit boundaries. Each set of 8 bits is converted to its decimal equivalent and separated by periods. For IPv6, the 128-bit address is divided along 16-bit boundaries, and each 16-bit block is converted to a 4-digit hexadecimal number and separated by colons. The resulting representation is called colon-hexadecimal.

The following is an IPv6 address in binary form:

```
0010000111011010000000011010011000000000000000000010111100111011
0000001010101010000000001111111111111110001010001001110001011010
```

The 128-bit address is divided along 16-bit boundaries, as follows:

```
0010000111011010   0000000011010011   0000000000000000   0010111100111011
0000001010101010   0000000011111111   1111111000101000   1001110001011010
```

Each 16-bit block is converted to hexadecimal and delimited with colons. The result is:

```
21DA:00D3:0000:2F3B:02AA:00FF:FE28:9C5A
```

IPv6 representation can be further simplified by removing the leading zeros within each 16-bit block. However, each block must have at least a single digit. With leading zero suppression, the address representation becomes:

```
21DA:D3:0:2F3B:2AA:FF:FE28:9C5A
```

**Compressing zeros**

Some types of addresses contain long sequences of zeros. To further simplify the representation of IPv6 addresses, a contiguous sequence of 16-bit blocks set to 0 in the colon-hexadecimal format can be compressed to :: (known as double-colon).

For example, the link-local address of FE80:0:0:0:2AA:FF:FE9A:4CA2 can be compressed to FE80::2AA:FF:FE9A:4CA2. The multicast address of FF02:0:0:0:0:0:0:2 can be compressed to FF02::2. Zero compression can only be used to compress a single contiguous series of 16-bit blocks expressed in colon-hexadecimal notation. You cannot use zero compression to include part of a 16-bit block. For example, you cannot express FF02:30:0:0:0:0:0:5 as FF02:3::5.

To determine how many 0 bits are represented by the ::, you can count the number of blocks in the compressed address, subtract this number from 8, and then multiply the result by 16. For example, in the address FF02::2, there are two blocks (the FF02 block and the 2 block). The number of bits expressed by the :: is 96 (96 = (8 - 2) × 16).

Zero compression can only be used once in a given address. Otherwise, you could not determine the number of 0 bits represented by each instance of a double-colon (::).

**IPv6 prefixes**

The prefix is the part of the address that indicates the bits that have fixed values or are the bits of the network identifier. Prefixes for IPv6 routes and subnet identifiers are expressed in the same way as Classless Inter-Domain Routing (CIDR) notation for IPv4. An IPv6 prefix is written in *address/prefix-length* notation. For example, 21DA:D3::/48 is a route prefix and 21DA:D3:0:2F3B::/64 is a subnet prefix.

### Note

- IPv4 implementations commonly use a dotted decimal representation of the network prefix known as the subnet mask. A subnet mask is not used in IPv6. Only prefix-length notation is supported.

# Unicast IPv6 addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Unicast IPv6 addresses**
A unicast address identifies a single interface within the scope of the type of unicast address. With the appropriate unicast routing topology, packets addressed to a unicast address are delivered to a single interface. The following types of addresses are unicast IPv6 addresses:

- Aggregatable global unicast addresses
- Link-local addresses
- Site-local addresses
- Special addresses
- Compatibility addresses
- NSAP addresses

**Aggregatable global unicast addresses**
Aggregatable global unicast addresses, identified by the Format Prefix (FP) of 001, are equivalent to public IPv4 addresses. They are globally routable and reachable on the IPv6 Internet. Aggregatable global unicast addresses are also known as global addresses.

As the name implies, aggregatable global unicast addresses are designed to be aggregated or summarized to produce an efficient routing infrastructure. Unlike the current IPv4-based Internet, which has a mixture of both flat and hierarchical routing, the IPv6-based Internet has been designed from its foundation to support efficient, hierarchical addressing and routing. The scope, the region of the IPv6 internetwork over which the address is unique, of an aggregatable global unicast address is the entire IPv6 Internet.




The fields in the aggregatable global unicast address are described as follows:

- TLA ID

  The TLA ID field indicates the Top Level Aggregation Identifier (TLA ID) for the address. The size of this field is 13 bits. The TLA identifies the highest level in the routing hierarchy. TLAs are administered by IANA and allocated to local Internet registries that, in turn, allocate individual TLA IDs to large, global Internet service providers (ISPs). A 13-bit field allows up to 8,192 different TLA IDs. Routers in the highest level of the IPv6 Internet routing hierarchy (called default-free routers) do not have a default route--only routes with 16-bit prefixes that correspond to the allocated TLAs.

- Res

  The Res field is reserved for future use in expanding the size of either the TLA ID or the NLA ID. The size of this field is 8 bits.

- NLA ID

  The NLA ID field indicates the Next Level Aggregation Identifier (NLA ) for the address. The NLA ID is used to identify a specific customer site. The size of this field is 24 bits. The NLA ID allows an ISP to create multiple levels of addressing hierarchy to organize addressing and routing and to identify sites. The structure of the ISP's network is not visible to the default-free routers.

- SLA ID

  The SLA ID field indicates the Site Level Aggregation Identifier (SLA ID) for the address. The SLA ID is used by an individual organization to identify subnets within its site. The size of this field is 16 bits. The organization can use these 16 bits within its site to create 65,536 subnets or multiple levels of addressing hierarchy and an efficient routing infrastructure. With 16 bits of subnetting flexibility, an aggregatable global unicast prefix assigned to an organization is equivalent to that organization being allocated an IPv4 Class A network ID (assuming that the last octet is used for identifying nodes on subnets). The structure of the customer's network is not visible to the ISP.

- Interface ID

  The Interface ID field indicates the interface of a node on a specific subnet. The size of this field is 64 bits.

The following illustration shows how the fields within the aggregatable global unicast address create a three-level topological structure.




The public topology is the collection of larger and smaller ISPs that provide access to the IPv6 Internet. The site topology is the collection of subnets within an organization's site. The interface identifier identifies a specific interface on a subnet within an organization's site. For more information about aggregatable global unicast addresses, see RFC 2374, "An IPv6 Aggregatable Global Unicast Address Format."

**Local-use unicast addresses**

There are two types of local-use unicast addresses:

1. Link-local addresses, which are used between on-link neighbors and for Neighbor Discovery processes.

2. Site-local addresses, which are used between nodes that communicate with other nodes in the same site.

**Link-local addresses**

Link-local addresses, identified by the FP of 1111 1110 10, are used by nodes when communicating with neighboring nodes on the same link. For example, on a single link IPv6 network with no router, link-local addresses are used to communicate between hosts on the link. Link-local addresses are equivalent to Automatic Private IP Addressing (APIPA) IPv4 addresses (using the 169.254.0.0/16 prefix). The scope of a link-local address is the local link. A link-local address is required for Neighbor Discovery processes and is always automatically configured, even in the absence of all other unicast addresses. For more information about the address autoconfiguration process for link-local addresses, see IPv6 address autoconfiguration.

Link-local addresses always begin with FE80. With the 64-bit interface identifier, the prefix for link-local addresses is always FE80::/64. An IPv6 router never forwards link-local traffic beyond the link.

For information about how to use link-local addresses, see Single subnet with link-local addresses.

**Site-local addresses**

Site-local addresses, identified by the FP of 1111 1110 11, are equivalent to the IPv4 private address space (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16). For example, private intranets that do not have a direct, routed connection to the IPv6 Internet can use site-local addresses without conflicting with aggregatable global unicast addresses. Site-local addresses are not reachable from other sites, and routers must not forward site-local traffic outside of the site. Site-local addresses can be used in addition to aggregatable global unicast addresses. The scope of a site-local address is the site (the organization internetwork).

Unlike link-local addresses, site-local addresses are not automatically configured and must be assigned either through stateless or stateful address configuration processes. For more information, see IPv6 address autoconfiguration.

The first 48-bits are always fixed for site-local addresses, beginning with FEC0::/48. After the 48 fixed bits is a 16-bit subnet identifier (Subnet ID field) that provides 16 bits with which you can create subnets within your organization. With 16 bits, you can have up to 65,536 subnets in a flat subnet structure, or you can subdivide the high-order bits of the Subnet ID field to create a hierarchical and aggregatable routing infrastructure. After the Subnet ID field is a 64-bit Interface ID field that identifies a specific interface on a subnet.

The aggregatable global unicast address and site-local address share the same structure beyond the first 48 bits of the address. In aggregatable global unicast addresses, the SLA ID identifies the subnet within an organization. For site-local addresses, the Subnet ID performs the same function. Because of this, you can assign a specific subnet number to identify a subnet that is used for both site-local and aggregatable global unicast addresses.

**Special addresses**
The following are special IPv6 addresses:

- Unspecified address

  The unspecified address (0:0:0:0:0:0:0:0 or ::) is used only to indicate the absence of an address. It is equivalent to the IPv4 unspecified address of 0.0.0.0. The unspecified address is typically used as a source address for packets that are attempting to verify the uniqueness of a tentative address. The unspecified address is never assigned to an interface or used as a destination address.

- Loopback address

  The loopback address (0:0:0:0:0:0:0:1 or ::1) is used to identify a loopback interface, enabling a node to send packets to itself. It is equivalent to the IPv4 loopback address of 127.0.0.1. Packets addressed to the loopback address are never sent on a link or forwarded by an IPv6 router.

**Compatibility addresses**
To aid in the migration from IPv4 to IPv6 and facilitate the coexistence of both types of hosts, the following addresses are defined:

- IPv4-compatible address

  The IPv4-compatible address, 0:0:0:0:0:0:*w.x.y.z* or ::*w.x.y.z* (where *w.x.y.z* is the dotted decimal representation of a public IPv4 address), is used by dual-stack nodes that are communicating with IPv6 over an IPv4 infrastructure. Dual-stack nodes are nodes with both IPv4 and IPv6 protocols. When the IPv4-compatible address is used as an IPv6 destination, IPv6 traffic is automatically encapsulated with an IPv4 header and sent to the destination by using the IPv4 infrastructure.

- IPv4-mapped address

  The IPv4-mapped address, 0:0:0:0:0:FFFF:*w.x.y.z* or ::FFFF:*w.x.y.z*, is used to represent an IPv4-only node to an IPv6 node. It is used only for internal representation. The IPv4-mapped address is never used as a source or destination address for an IPv6 packet. The IPv6 protocol does not support the use of IPv4-mapped addresses.

- 6to4 address

  The 6to4 address is used for communicating between two nodes running both IPv4 and IPv6 over the Internet. The 6to4 address is formed by combining the prefix 2002::/16 with the 32 bits of the public IPv4 address of the node, forming a 48-bit prefix. For example, for the IPv4 address of 131.107.0.1, the 6to4 address prefix is 2002:836B:1::/48. For more information about 6to4, see IPv6 traffic between nodes in different sites across the Internet (6to4).

**NSAP addresses**
To provide a means for mapping Network Service Access Point (NSAP) addresses to IPv6 addresses, NSAP addresses use the FP of 0000001 and map the last 121 bits of the IPv6 address to an NSAP address. For more information about the four types of NSAP address mappings, see RFC 1888, "OSI NSAPs and IPv6."

### Note

- IPv6 is a rapidly evolving standard. The RFCs referenced might have been made obsolete by newer RFCs.

# Multicast IPv6 addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multicast IPv6 addresses**

A multicast address identifies multiple interfaces. With the appropriate multicast routing topology, packets addressed to a multicast address are delivered to all interfaces that are identified by the address.

IPv6 multicast addresses have the Format Prefix (FP) of 1111 1111. An IPv6 address is simple to classify as multicast because it always begins with FF. Multicast addresses cannot be used as source addresses.

Beyond the FP, multicast addresses include additional structure to identify their flags, scope, and multicast group, as shown in the following illustration.




The fields in the multicast address are as follows:

- Flags

  The Flags field indicates flags that are set on the multicast address. The size of this field is 4 bits. As of RFC 2373, the only flag defined is the Transient (T) flag. The T flag uses the low-order bit of the Flags field. When set to 0, the T flag indicates that the multicast address is a permanently-assigned (well-known) multicast address allocated by the Internet Assigned Numbers Authority (IANA). When set to 1, the T flag indicates that the multicast address is a transient (not permanently assigned) multicast address.

- Scope

  The Scope field indicates the scope of the IPv6 internetwork for which the multicast traffic is intended. The size of this field is 4 bits. In addition to information provided by multicast routing protocols, routers use the multicast scope to determine whether multicast traffic can be forwarded.

- The following scopes are defined in RFC 2373:

| Scope field value | Scope |
|---|---|
| 1 | Node-local |
| 2 | Link-local |
| 5 | Site-local |
| 8 | Organization-local |
| E | Global |

  For example, traffic with the multicast address of FF02::2 has a link-local scope. An IPv6 router never forwards this traffic beyond the local link.

- Group ID

  The Group ID field identifies the multicast group and is unique within the scope. The size of this field is 112 bits. Permanently assigned group IDs are independent of the scope. Transient group IDs are only relevant to a specific scope. Multicast addresses from FF01:: through FF0F:: are reserved, well-known addresses.

To identify all nodes for the node-local and link-local scopes, the following multicast addresses are defined:

- FF01::1 (node-local scope all-nodes address)

- FF02::1 (link-local scope all-nodes address)

To identify all routers for the node-local, link-local, and site-local scopes, the following multicast addresses are defined:

- FF01::2 (node-local scope all-routers address)
- FF02::2 (link-local scope all-routers address)
- FF05::2 (site-local scope all-routers address)

With 112 bits in the Group ID, it is possible to have 2112 group IDs. However, because of the way in which IPv6 multicast addresses are mapped to Ethernet multicast MAC addresses, RFC 2373 recommends assigning the Group ID from the low order 32 bits of the IPv6 multicast address and setting the remaining original group ID bits to 0. By using only the low-order 32 bits in the group ID, each group ID maps to a unique Ethernet multicast MAC address.

**Solicited-node address**
The solicited-node address facilitates efficient querying of network nodes during address resolution. In IPv4, the ARP Request frame is sent to the MAC-level broadcast, disturbing all nodes on the network segment, including those that are not running IPv4. IPv6 uses the Neighbor Solicitation message to perform address resolution. However, instead of using the local-link scope all-nodes address as the Neighbor Solicitation message destination, which would disturb all IPv6 nodes on the local link, the solicited-node multicast address is used. The solicited-node multicast address consists of the prefix FF02::1:FF00:0/104 and the last 24-bits of the IPv6 address that is being resolved.

For example, for the node with the link-local IPv6 address of FE80::2AA:FF:FE28:9C5A, the corresponding solicited-node address is FF02::1:FF28:9C5A. To resolve the FE80::2AA:FF:FE28:9C5A address to its link layer address, a node sends a Neighbor Solicitation message to the solicited-node address of FF02::1:FF28:9C5A. The node that is using the address of FE80::2AA:FF:FE28:9C5A is listening for multicast traffic at the solicited-node address and, for interfaces that correspond to a physical network adapter, has registered the corresponding multicast address with the network adapter.

The result of using the solicited-node multicast address is that address resolution, which commonly occurs on a link, is not required to use a mechanism that disturbs all network nodes. In fact, very few nodes are disturbed during address resolution. In practice, because of the relationship between the Ethernet MAC address, the IPv6 interface ID, and the solicited-node address, the solicited-node address acts as a pseudo-unicast address for very efficient address resolution.

# Anycast IPv6 addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Anycast IPv6 addresses**

An anycast address identifies multiple interfaces. With the appropriate routing topology, packets addressed to an anycast address are delivered to a single interface (the nearest interface that is identified by the address). The nearest interface is defined as being closest in terms of routing distance. A multicast address is used for one-to-many communication, with delivery to multiple interfaces. An anycast address is used for one-to-one-of-many communication, with delivery to a single interface.

In order to facilitate delivery to the nearest anycast group member, the routing infrastructure must be aware of the interfaces that are assigned anycast addresses and their distances in terms of routing metrics. At present, anycast addresses are used only as destination addresses and are assigned only to routers. Anycast addresses are assigned from the unicast address space. The scope of an anycast address is the scope of the type of unicast address from which the anycast address is assigned.

**Subnet-router anycast address**

The subnet-router anycast address is both predefined and required. It is created from the subnet prefix for a given interface. To construct the subnet-router anycast address, the bits in the subnet prefix are fixed at their appropriate values and the remaining bits are set to 0. All router interfaces attached to a subnet are assigned the subnet-router anycast address for that subnet. The subnet-router anycast address is used for communication with one of multiple routers attached to a remote subnet.

# Addresses for hosts and routers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Addresses for hosts and routers**

An IPv4 host with a single network adapter typically has a single IPv4 address assigned to that adapter. An IPv6 host, however, usually has multiple IPv6 addresses, even with a single interface.

An IPv6 host is assigned the following unicast addresses:

- A link-local address for each interface
- Unicast addresses for each interface (which could be a site-local address and one or multiple global addresses)
- The loopback address (::1) for the loopback interface

Typical IPv6 hosts are logically multihomed because they have at least two addresses with which they can receive packets. Each host has a link-local address for local link traffic and a routable site-local or global address.

Additionally, each host is listening for traffic on the following multicast addresses:

- The node-local scope all-nodes address (FF01::1)
- The link-local scope all-nodes address (FF02::1)
- The solicited-node address for each unicast address on each interface
- The multicast addresses of joined groups on each interface

An IPv6 router is assigned the following unicast addresses:

- A link-local address for each interface
- Unicast addresses for each interface (which might be a site-local address and one or multiple aggregatable global unicast addresses)
- The loopback address (::1) for the loopback interface

An IPv6 router is assigned the following anycast addresses:

- A subnet-router anycast address for each subnet
- Additional anycast addresses (optional)

Additionally, each router is listening for traffic on the following multicast addresses:

- The node-local scope all-nodes address (FF01::1)
- The node-local scope all-routers address (FF01::2)
- The link-local scope all-nodes address (FF02::1)
- The link-local scope all-routers address (FF02::2)
- The site-local scope all-routers address (FF05::2)

- The solicited-node address for each unicast address on each interface
- The addresses of joined groups on each interface

# IPv6 interface identifiers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 interface identifiers**

The last 64 bits of an IPv6 address are the interface identifier that is unique to the 64-bit prefix of the IPv6 address. The ways in which an interface identifier is determined are as follows:

- RFC 2373 states that all unicast addresses that use the prefixes 001 through 111 must also use a 64-bit interface identifier that is derived from the Extended Unique Identifier (EUI)-64 address.
- RFC 3041 describes a randomly-generated interface identifier that changes over time to provide a level of anonymity.
- An interface identifier that is assigned during stateful address autoconfiguration (for example, through DHCPv6). DHCPv6 standards are currently being defined. The IPv6 protocol for the Windows Server 2003 family and Windows XP does not support stateful address configuration or DHCPv6.
- An interface identifier that is manually configured.

**EUI-64 address-based interface identifiers**

The 64-bit EUI-64 address is defined by the Institute of Electrical and Electronic Engineers (IEEE). EUI-64 addresses are either assigned to a network adapter or derived from IEEE 802 addresses.

**IEEE 802 addresses**

Traditional interface identifiers for network adapters use a 48-bit address called an IEEE 802 address. This address consists of a 24-bit company ID (also called the manufacturer ID), and a 24-bit extension ID (also called the board ID). The combination of the company ID, which is uniquely assigned to each manufacturer of network adapters, and the board ID, which is uniquely assigned to each network adapter at the time of assembly, produces a globally unique 48-bit address. This 48-bit address is also called the physical, hardware, or media access control (MAC) address.




Defined bits within the IEEE 802 address are the following:

- Universal/Local (U/L)

  The U/L bit is the seventh bit of the first byte and is used to determine whether the address is universally or locally administered. If the U/L bit is set to 0, the IEEE, through the designation of a unique company ID, has administered the address. If the U/L bit is set to 1, the address is locally administered. The network administrator has overridden the manufactured address and specified a different address.

- Individual/Group (I/G)

  The I/G bit is the low order bit of the first byte and is used to determine whether the address is an individual address (unicast) or a group address (multicast). When set to 0, the address is a unicast address. When set to 1, the address is a multicast address.

For a typical 802.x network adapter address, both the U/L and I/G bits are set to 0, corresponding to a universally administered, unicast MAC address.

**IEEE EUI-64 addresses**

The IEEE EUI-64 address represents a new standard for network interface addressing. The company ID is still 24-bits in length, but the extension ID is 40 bits, creating a much larger address space for a network adapter manufacturer. The EUI-64 address uses the U/L and I/G bits in the same way as the IEEE 802 address.




**Mapping IEEE 802 addresses to EUI-64 addresses**

To create an EUI-64 address from an IEEE 802 address, the 16 bits of 11111111 11111110 (0xFFFE) are inserted into the IEEE 802 address between the company ID and the extension ID. The following illustration shows the conversion of an IEEE 802 address to an EUI-64 address.




**Mapping EUI-64 addresses to IPv6 interface identifiers**

To obtain the 64-bit interface identifier for IPv6 unicast addresses, the U/L bit in the EUI-64 address is complemented (if it is a 1, it is set to 0; and if it is a 0, it is set to 1). The following illustration shows the conversion for a universally administered, unicast EUI-64 address.




To obtain an IPv6 interface identifier from an IEEE 802 address, you must first map the IEEE 802 address to an EUI-64 address, and then complement the U/L bit. The following illustration shows the conversion process for a universally administered, unicast IEEE 802 address.




**IEEE 802 address conversion example**

Host A has the Ethernet MAC address of 00-AA-00-3F-2A-1C. First, it is converted to EUI-64 format by inserting FF-FE between the third and fourth bytes, yielding 00-AA-00-FF-FE-3F-2A-1C. Next, the U/L bit, which is the seventh bit in the first byte, is complemented. The first byte in binary form is 00000000. When the seventh bit is complemented, it becomes 00000010 (0x02). The final result is 02-AA-00-FF-FE-3F-2A-1C which, when converted to colon-hexadecimal notation, becomes the interface identifier 2AA:FF:FE3F:2A1C. As the result, the link-local address that corresponds to the network adapter with the MAC address of 00-AA-00-3F-2A-1C is FE80::2AA:FF:FE3F:2A1C.

### Note

- When complementing the U/L bit, add 0x2 to the first byte if the EUI-64 address is universally administered, and subtract 0x2 from the first byte if the EUI-64 address is locally administered.

**Temporary address interface identifiers**

In today's IPv4-based Internet, a typical Internet user connects to an Internet service provider (ISP) and obtains an IPv4 address by using the Point-to-Point Protocol (PPP) and the Internet Protocol Control Protocol (IPCP). Each time the user connects, a different IPv4 address might be obtained. Because of this, it is difficult to track a user's traffic on the Internet on the basis of IP address.

For IPv6-based dial-up connections, the user is assigned a 64-bit prefix after the connection is made through router discovery and stateless address autoconfiguration. If the interface identifier is always based on the EUI-64 address (as derived from the static IEEE 802 address), it is possible to identify the traffic of a specific node regardless of the prefix, making it easy to track a specific user and their use of the Internet. To address this concern and provide a level of anonymity, an alternative IPv6 interface identifier that is randomly generated and changes over time is described in RFC 3041.

The initial interface identifier is generated by using random numbers. For IPv6 systems that cannot store any historical information for generating future interface identifier values, a new random interface identifier is generated each time the IPv6 protocol is initialized. For IPv6 systems that have storage capabilities, a history value is stored and, when the IPv6 protocol is initialized, a new interface identifier is created through the following process:

1. Retrieve the history value from storage and append the interface identifier based on the EUI-64 address of the adapter.

2. Compute the Message Digest-5 (MD5) one-way encryption hash over the quantity in step 1.

3. Save the last 64 bits of the MD5 hash computed in step 2 as the history value for the next interface identifier computation.

4. Take the first 64 bits of the MD5 hash computed in Step 2 and set the seventh bit to zero. The seventh bit corresponds to the U/L bit which, when set to 0, indicates a locally administered interface identifier. The result is the interface identifier.

The resulting IPv6 address, based on this random interface identifier, is known as a temporary address. Temporary addresses are generated for public address prefixes that use stateless address autoconfiguration. Temporary addresses are used for the lower of the following values of the valid and preferred lifetimes:

- The lifetimes included in the Prefix Information option in the received Router Advertisement message.

- Local default values of 1 week for valid lifetime and 1 day for preferred lifetime.

After the valid lifetime of a temporary address expires, a new interface identifier and temporary address are generated.

The IPv6 protocol for the Windows Server 2003 family does not create temporary addresses for global address prefixes by default. You can modify this default setting with the **netsh interface ipv6 set privacy state=enabled** command. For more information, see Netsh commands for Interface IPv6.

# IPv6 and Network Applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 and network applications**
This section covers:

- Support for Windows Sockets
- IPv6 applications

# Support for Windows Sockets

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Support for Windows Sockets**

Changes to existing Windows Sockets functions for IPv6 are documented in RFC 2553, "Basic Socket Interface Extensions for IPv6." The IPv6 protocol for the Windows Server 2003 family supports RFC 2553 with the following exceptions:

- The header files listed in the RFC do not apply to the Windows Server 2003 family. The Windows Server 2003 family header files are Winsock2.h and Ws2tcpip.h.
- The Windows Server 2003 family socket structure does not include the **sa_len** member. Therefore, section 3.3 applies, but section 3.4 does not.
- IPv4-mapped addresses, as described in section 3.7, are not supported.
- The interface identification functions described in Section 4 are not supported.
- Sections 6.1 through 6.3 and 6.6 are not supported because they are in the process of being made obsolete by the Internet Engineering Task Force (IETF).

**Tools for IPv6 application development**

For Winsock IPv6 application development, the following are required:

- Microsoft Platform SDK (January 2000 or later)
- Microsoft Visual C++ 6.0 or later
- IPv6 protocol for the Windows Server 2003 family

# IPv6 applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 applications**

The Windows Server 2003 family includes the following IPv6-enabled applications and components:

- FTP client

  The File Transfer Protocol (FTP) client, Ftp.exe, can be used to establish FTP sessions with IPv4 and IPv6 FTP servers.

- Telnet client

  The Telnet client, Telnet.exe, can be used to establish Telnet sessions with IPv4 and IPv6 Telnet servers.

- Internet Explorer

  The new Internet extensions dynamic link library, Wininet.dll, enables Web browsers to access IPv6-enabled Web servers. For example, Wininet.dll is used by Microsoft Internet Explorer to make connections with a Web server to view Web pages. Internet Explorer uses IPv6 to download Web pages when the Domain Name System (DNS) query for the name of the Web server in the URL returns an IPv6 address.

  URLs that use the format for literal IPv6 addresses described in RFC 2732, "Format for Literal IPv6 Addresses in URLs," are not supported by the version of Internet Explorer provided with Windows 2000.

  For a list of IPv6-only Web sites, see the IPv6 Web site. Accessing IPv6-only Web sites assumes that you can connect to the IPv6 Internet. For more information about connecting to the IPv6 Internet, see Connecting to the IPv6 Internet.

Additionally, the Remote Procedure Call (RPC) components of Windows Server 2003 family are IPv6-enabled and the versions of Network Monitor provided with the Windows Server 2003 family, Windows 2000, and Systems Management Server 2.0 support the parsing of IPv6 headers.

**Notes**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

- Internet Explorer cannot browse IPv6 Web sites if it is configured to use a proxy server. When Internet Explorer is configured to use a proxy server, all name resolution requests for Web sites are forwarded to the proxy server. Until the proxy server is IPv6-enabled, proxy-based requests for local or remote IPv6 Web pages are unsuccessful. For information about how to configure Internet Explorer to operate without the use of a proxy server, see Internet Explorer Help.

# Administering IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering IPv6**
This section covers:

- IPv6 configuration items
- IPv6 configuration methods
- IPv6 Configurations
- Setting up an IPv6 Test Lab

# IPv6 configuration items

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 configuration items**
You can configure the following for the IPv6 protocol:

- IPv6 address
- Default router
- DNS server

**IPv6 address**
By default, link-local addresses are automatically configured for each interface on each IPv6 node (host or router) with a unique link-local IPv6 address. If you want to communicate with IPv6 nodes that are not on attached links, the host must have additional site-local or global unicast addresses. Additional addresses for hosts are either obtained from router advertisements sent by a router or assigned manually. Additional addresses for routers must be assigned manually.

For more information, see Unicast IPv6 addresses, Configure IPv6 with manual addresses, and IPv6 address autoconfiguration.

**Default router**
To communicate with IPv6 nodes on other network segments, IPv6 must use a default router. A default router is automatically assigned based on the receipt of a router advertisement. Alternately, you can add a default route to the IPv6 routing table. You do not need to configure a default router for a network that consists of a single network segment.

For more information, see IPv6 address autoconfiguration and Add an IPv6 route.

**DNS server**
You can use a Domain Name System (DNS) server to resolve host names to IPv6 addresses. When an IPv6 host is configured with the address of a DNS server, the host sends DNS name queries to the server for resolution. AAAA (quad-A) resource records, which are stored on your DNS servers, enable mapping from a host name to its IPv6 address.

To enable DNS name resolution, configure an IPv6 router with forwarding enabled and a global prefix that is advertised to clients. You can do this by using the **netsh interface ipv6 add route** and **netsh interface ipv6 set interface** commands. For more information, see Add an IPv6 route and Enable IPv6 forwarding.

By default, DNS is configured to allow DNS dynamic updates. You can either leave dynamic update enabled when you use IPv6 with DNS, or you can manually add DNS records for IPv6 clients.

**Configuring clients with the DNS server address**
To provide communication between DNS clients and servers, you can configure the clients with the IPv6 address of the DNS server, or you can configure your DNS server with one of the three default DNS server IPv6 addresses that are automatically configured on all IPv6 clients.

You can configure clients with the IPv6 address of the DNS server by using the **netsh interface ipv6 add dns** command at each client computer or in a logon script that is run each time a client logs on to the network.

To configure the DNS server with one of the three IPv6 addresses that are available on IPv6 client computers by default, use the **netsh interface ipv6 add address** command. The three default DNS server addresses are:

- FEC0:0:0: FFFF::1
- FEC0:0:0: FFFF::2
- FEC0:0:0:FFFF::3

If your DNS server is on a different subnet than your IPv6 clients, configure a static route to the DNS server on any IPv6 router that is available on the DNS server's subnet.

**Configuring the DNS server to listen on IPv6**

You can configure the DNS server to listen for DNS name registration and resolution requests over IPv6. When your DNS server is configured to listen on both IPv4 and IPv6:

- Devices that function over IPv6 but not IPv4 will function properly with your DNS server.
- Computers and other devices that are configured with both IPv4 and IPv6 use IPv6 by default.

To configure your DNS server to listen over IPv6, do the following:

1. Install Windows Support Tools.

   For more information, see Install Windows Support Tools

2. Open Command Prompt.

3. Type the following command:

   **dnscmd /config /EnableIPv6 1**

4. Restart the DNS Server service.

   For more information, see Start or stop a DNS server.

For more information, see IPv6 Name resolution, DNS, and Netsh commands for Interface IPv6.

# IPv6 configuration methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 configuration methods**

The IPv6 protocol can use the following configuration methods:

- Automatic configuration using stateless addresses
- Manual configuration

**Automatic configuration**

The IPv6 protocol supports address autoconfiguration for stateless addresses as defined in RFC 2462, "IPv6 Stateless Address Autoconfiguration." IPv6 nodes (hosts and routers) automatically create unique link-local addresses for all LAN interfaces that appear to be Ethernet or FDDI interfaces. IPv6 hosts use received Router Advertisement messages to automatically configure:

- A default router.
- The default setting for the Hop Limit field in the IPv6 header.
- The determination of whether the node should use a stateful address configuration protocol, such as Dynamic Host Configuration Protocol for IPv6 (DHCPv6), for addresses and other configuration requirements. The IPv6 protocol for the Windows Server 2003 family does not currently support DHCPv6 or any other stateful address configuration protocol.
- The timers used in Neighbor Discovery processes.
- The maximum transmission unit (MTU) of the local link.
- The list of network prefixes that are defined for the link. Each network prefix contains both the IPv6 network prefix and its valid and preferred lifetimes. If indicated, a network prefix is combined with the interface identifier to create a stateless IPv6 address configuration for the receiving interface. A network prefix also defines the range of addresses for nodes on the local link.

For more information, see IPv6 address autoconfiguration.

Additionally, the IPv6 protocol for the Windows Server 2003 family automatically configures:

- 6to4 addresses on a 6to4 tunneling interface for all public IPv4 addresses that are assigned to the computer (if a Router Advertisement is not received and the computer is configured with a public IPv4 address).
- Intrasite Automatic Tunnel Addressing Protocol (ISATAP) addresses on the Automatic Tunneling Pseudo-Interface for all IPv4 addresses that are assigned to the computer.
- Routes to off-link prefixes (if the off-link address prefix is advertised by a router running a member of the Windows Server 2003 family).

**Manual configuration**

You can manually configure IPv6 addresses and routes by using the Netsh commands for Interface IPv6 command-line tool. Manual configuration might be required in a network that has multiple IPv6 network segments within which routers are not configured to send router advertisements.

For more information, see Netsh commands for Interface IPv6, Configure IPv6 with manual addresses, and Add an IPv6 route.

# IPv6 Configurations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 configurations**
This section covers:

- Single subnet with link-local addresses
- Using IPSec between two local link hosts
- IPv6 traffic between nodes on different subnets of an IPv6 internetwork
- IPv6 traffic between nodes on different subnets of an IPv4 internetwork
- IPv6 traffic between nodes in different sites across the Internet (6to4)
- Connecting to the IPv6 Internet
- Using IPv4-compatible addresses

# Single subnet with link-local addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Single subnet with link-local addresses**
This configuration requires only the installation of the IPv6 protocol on at least two nodes on the same network segment (also known as a link or subnet) with no intermediate routers.

The following illustration shows the configuration of two nodes on a single subnet using link-local addresses.




By default, the IPv6 protocol for the Windows Server 2003 family and Windows XP configures link-local addresses for each interface that corresponds to installed Ethernet network adapters. Link-local addresses have the prefix of FE80::/64. The last 64 bits of the IPv6 address is known as the interface identifier. It is derived from the 48-bit MAC address of the network adapter. For more information, see IPv6 interface identifiers.

To create the IPv6 interface identifier from the 48-bit (6-byte) Ethernet MAC address:

- The hexadecimal digits 0xFF-FE are inserted between the third and fourth bytes of the MAC address.
- The Universal/Local bit (the second low-order bit of the first byte of the MAC address) is complemented. If it is a 1, it is set to 0; and if it is a 0, it is set to 1.

For example, for the MAC address of 00-60-08-52-F9-D8:

- The hexadecimal digits 0xFF-FE are inserted between 0x08 (the third byte) and 0x52 (the fourth byte) of the MAC address, forming the 64-bit address of 00-60-08-FF-FE-52-F9-D8.
- The Universal/Local bit, the second low-order bit of 0x00 (the first byte) of the MAC address, is complemented. The second low-order bit of 0x00 is 0 which, when complemented, becomes 1. The result is that for the first byte, 0x00 becomes 0x02.

As a result, the IPv6 interface identifier that corresponds to the Ethernet MAC address of 00-60-08-52-F9-D8 is 02-60-08-FF-FE-52-F9-D8.

The link-local address of a node is the combination of the prefix FE80::/64 and the 64-bit interface identifier expressed in colon-hexadecimal notation. As a result, the link-local address of this example node, with the prefix of FE80::/64 and the interface identifier 02-60-08-FF-FE-52-F9-D8, is FE80::260:8FF:FE52:F9D8.

You can view your link local address by using **netsh interface ipv6 show interface**. When you run this command without any additional parameters, a simple list is provided as output:

```
Idx  Met   MTU    State         Name
---  ----  -----  ------------  -----
  5     0   1500  Connected     Local Area Connection
  4     0   1500  Connected     Local Area Connection 2
  3     1   1280  Connected     6to4 Tunneling Pseudo-Interface
  2     1   1280  Connected     Automatic Tunneling Pseudo-Interface
  1     0   1500  Connected     Loopback Pseudo-Interface
```

For more detailed output, you can designate a connection name as a command parameter (for example, **netsh interface ipv6 show interface "Local Area Connection 2"**). Using the **interface** parameter (with either an interface name or an interface index number) results in the following output, which includes the link local address for the connection:

```
Interface 4: Local Area Connection 2

Addr Type  DAD State  Valid Life   Pref. Life   Address
---------  ---------- ------------ ------------ ----------------------------
Temporary  Preferred    6d4h41m33s     4h38m46s 3ffe:2900:d005:f282:4063:32a8:5c81:62f2
Temporary  Deprecated   5d4h44m24s           0s 3ffe:2900:d005:f282:cd74:3dd7:857b:b57
Temporary  Deprecated   4d4h47m16s           0s 3ffe:2900:d005:f282:d880:d193:f2dd:d929
Temporary  Deprecated    3d4h50m7s           0s 3ffe:2900:d005:f282:7482:2f05:8129:54ba
Temporary  Deprecated   2d4h52m59s           0s 3ffe:2900:d005:f282:d530:25de:57b:7ee0
Temporary  Deprecated   1d4h55m50s           0s 3ffe:2900:d005:f282:c58c:4290:22c6:7b3b
Temporary  Deprecated     4h58m42s           0s 3ffe:2900:d005:f282:8464:acf0:8393:cf6
Public     Preferred  29d23h57m19s  6d23h57m19s fec0::f282:2c0:4fff:fe19:bad3
Public     Preferred  29d23h57m19s  6d23h57m19s 3ffe:2900:d005:f282:2c0:4fff:fe19:bad3
Link       Preferred      infinite     infinite fe80::2c0:4fff:fe19:bad3

Connection Name          : Local Area Connection 2
GUID                     : {433F15CA-E3FD-4DE4-B3FF-7EF4B30CA4E7}
State                    : Connected
Metric                   : 0
Link MTU                 : 1500 bytes
True Link MTU            : 1500 bytes
Current Hop Limit        : 64
Reachable Time           : 4h43m20s
Base Reachable Time      : 8h20m
Retransmission Interval  : 16m40s
DAD Transmits            : 1
DNS Suffix               : example.microsoft.com
Zone ID for Link         : 4
Zone ID for Site         : 1
Uses Neighbor Discovery  : Yes
Sends RA's               : No
Forwards Packets         : No
Link-Layer Address       : 00-c0-4f-19-ba-d3
```

Interface 4 is an interface that corresponds to an installed Ethernet adapter with a link-local address of FE80::2C0:4FFF:FE19:BAD3.

**Testing connectivity between two link-local hosts**

You can perform a simple ping (an exchange of ICMPv6 Echo Request and Echo Reply messages) with IPv6 between two link-local hosts by completing the following steps:

1. Install IPv6 on two computers. For more information, see Install IPv6.

2. Use **netsh interface ipv6 show interface** on Host A to obtain the link-local address and the interface index for the Ethernet interface.

   For example, the link-local address of Host A is FE80::210:5AFF:FEAA:20A2, and the interface index of the Ethernet interface is 4.

3. Use **netsh interface ipv6 show interface** on Host B to obtain the link-local address and the interface index for the Ethernet interface.

   For example, the link-local address of Host B is FE80::260:97FF:FE02:6EA5, and the interface index for the Ethernet interface is 3.

4. From Host A, use Ping.exe to ping Host B.

   For example, **ping FE80::260:97FF:FE02:6EA5%4**

When you specify a link-local or site-local destination address, you can specify the zone ID to make the zone (an area of the network of a specific scope) of the traffic specific.

- For example, on a computer with multiple Ethernet adapters that are connected to separate links, each Ethernet adapter is assigned a link-local address. Destination link-local addresses in this configuration are ambiguous because a specific link-local address can be assigned to multiple nodes located on the links that are reachable from all of the installed

Ethernet adapters. To define the area of the network for which the destination is intended, the zone ID is used to indicate the Ethernet adapter over which traffic is sent and received. In the IPv6 protocol, the zone ID for link-local addresses is typically the interface index from the display of the **netsh interface ipv6 show interface** command. The interface index is defined locally on each IPv6 host. Because of this, the interface index used by Host A to reach Host B might not be the same as the interface index used by Host B to reach Host A.

- When using site-local addresses, it is possible to be connected to multiple sites. In this case, each site is assigned a site identifier. To define the area of the network for which the destination is intended, the zone ID is used to indicate the site identifier. In the IPv6 protocol, the zone ID is the Zone ID for Site from the display of the **netsh interface ipv6 show interface level=verbose** command. If you are connected only to a single site, the default site identifier is 1 and the zone ID does not have to be specified. The site identifier is defined locally on each IPv6 host. Because of this, the site identifier used by Host A to reach Host B might not be the same as the site identifier used by Host B to reach Host A.

The notation that is used to specify the zone ID with an address is *Address***%***ZoneID*.

For additional information about configurations, see IPv6 Configurations.

For information about using IPv6 in a test lab, see Setting up an IPv6 Test Lab.

For more information about the **netsh interface ipv6 show interface** command, see Netsh commands for Interface IPv6.

# Using IPSec between two local link hosts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using IPSec between two local link hosts**

**Caution**

- This implementation of IPSec for IPv6 is not recommended for use in a production environment because it relies on static keying and has no provisions for updating keys upon sequence number reuse.
- When you manually configure Security Parameters Indexes (SPIs), always use random numbers. Do not use sequential numbers for SPIs, or you will compromise the security of your IPSec for IPv6 policies.
- The IPv6 protocol for the Windows Server 2003 family does not support the use of IPSec Encapsulating Security Payload (ESP) encryption. However, the use of ESP with NULL encryption is supported. Although NULL encryption uses the ESP header, only data origin authentication and data integrity services are provided.

This configuration creates an IPSec security association (SA) between two hosts on the same subnet. The SA performs authentication by using the Authentication Header (AH) and the Message Digest 5 (MD5) hashing algorithm. In this example, the configuration secures all traffic between two neighboring hosts. Host 1 has the link-local address of FE80::2AA:FF:FE53:A92C, and Host 2 has the link-local address of FE80::2AA:FF:FE92:D0F1.

1. On Host 1, create blank security association (.sad) and security policy (.spd) files by using the **ipsec6 s** command. In this example, the Ipsec6.exe command is **ipsec6 s test**. This creates two files with blank entries for manually configuring security associations (Test.sad) and security policies (Test.spd).

2. On Host 1, edit the .spd file, adding a security policy that secures all traffic between Host 1 and Host 2.

   The following table shows the security policy entry that is added to Test.spd before the first entry (the first entry in Test.spd is not modified):

| .spd file field name | Example value |
|---|---|
| Policy | 2 |
| RemoteIPAddr | - FE80::2AA:FF:FE92:D0F1 |
| LocalIPAddr | - * |
| Protocol | - * |
| RemotePort | - * |
| LocalPort | - * |
| IPSecProtocol | AH |
| IPSecMode | TRANSPORT |
| RemoteGWIPAddr | * |
| SABundleIndex | NONE |
| Direction | BIDIRECT |
| Action | APPLY |
| InterfaceIndex | 0 |

Type a semicolon at the end of the entry configuring this security policy. Policy entries must be placed in decreasing numerical order.

3. On Host 1, edit the .sad file, adding SA entries to secure all traffic between Host 1 and Host 2. Two security associations must be created, one for traffic to Host 2 and one for traffic from Host 2.

   The following table shows the first SA entry that is added to Test.sad (for traffic to Host 2):

| .sad file field name | Example value |
|---|---|
| SAEntry | 2 |
| SPI | 3001 |
| SADestIPAddr | FE80::2AA:FF:FE92:D0F1 |
| DestIPAddr | POLICY |
| SrcIPAddr | POLICY |
| Protocol | POLICY |
| DestPort | POLICY |
| SrcPort | POLICY |
| AuthAlg | HMAC-MD5 |
| KeyFile | Test.key |
| Direction | OUTBOUND |
| SecPolicyIndex | 2 |

   Type a semicolon at the end of the entry configuring this SA.

   The following table shows the second SA entry that is added to Test.sad (for traffic from Host 2):

| .sad file field name | Example value |
|---|---|
| SAEntry | 1 |
| SPI | 3000 |
| SADestIPAddr | FE80::2AA:FF:FE53:A92C |
| DestIPAddr | POLICY |
| SrcIPAddr | POLICY |
| Protocol | POLICY |
| DestPort | POLICY |
| SrcPort | POLICY |
| AuthAlg | HMAC-MD5 |
| KeyFile | Test.key |
| Direction | INBOUND |
| SecPolicyIndex | 2 |

Type a semicolon at the end of the entry configuring this SA. SA entries must be placed in decreasing numerical order.

4. On Host 1, create a file that contains data used to create and validate the Message Digest 5 (MD5) keyed hash on each IPSec-protected packet that is exchanged with Host 2. In this example, a text file is used. Test.key is created with the contents **This is a test**. There are no extra characters, spaces, or lines.

   The IPv6 protocol supports only manually configured keys for quick mode SAs (also known as IPSec or Phase II SAs), because main mode negotiation through Internet Key Exchange (IKE) is not performed. Manual keys are configured by creating files that contain either the text or binary data of the manual key. In this example, the same key for the SAs is used in both directions. You can use different keys for inbound and outbound SAs by creating different key files and referencing them with the **KeyFile** field in the .sad file.

5. On Host 2, use the **ipsec6 s** command to create blank security association (.sad) and security policy (.spd) files. In this example, the Ipsec6.exe command is **ipsec6 s test**. This creates two files with blank entries for manually configuring security associations (Test.sad) and security policies (Test.spd).

   To simplify the example, the same file names for the .sad and .spd files are used on Host 2. You can choose to use different file names on each host.

6. On Host 2, edit the .spd file, adding a security policy that secures all traffic between Host 2 and Host 1.

   The following table shows the security policy entry that is added to Test.spd before the first entry (the first entry in Test.spd is not modified):

| **.spd file field name** | **Example value** |
|---|---|
| Policy | 2 |
| RemoteIPAddr | - FE80::2AA:FF:FE53:A92C |
| LocalIPAddr | - * |
| Protocol | - * |
| RemotePort | - * |
| LocalPort | - * |
| IPSecProtocol | AH |
| IPSecMode | TRANSPORT |
| RemoteGWIPAddr | * |
| SABundleIndex | NONE |
| Direction | BIDIRECT |
| Action | APPLY |
| InterfaceIndex | 0 |

   Type a semicolon at the end of the entry configuring this security policy. Policy entries must be placed in decreasing numerical order.

7. On Host 2, edit the .sad file, adding SA entries to secure all traffic between Host 2 and Host 1. Two security associations must be created: one for traffic to Host 1 and one for traffic from Host 1.

   The following table shows the first SA entry that is added to Test.sad (for traffic to Host 1):

| .sad file field name | Example value |
|---|---|
| SAEntry | 2 |
| SPI | 3000 |
| SADestIPAddr | FE80::2AA:FF:FE53:A92C |
| DestIPAddr | POLICY |
| SrcIPAddr | POLICY |
| Protocol | POLICY |
| DestPort | POLICY |
| SrcPort | POLICY |
| AuthAlg | HMAC-MD5 |
| KeyFile | Test.key |
| Direction | OUTBOUND |
| SecPolicyIndex | 2 |

Type a semicolon at the end of the entry configuring this SA.

The following table shows the second SA entry that is added to Test.sad (for traffic from Host 1):

| .sad file field name | Example value |
|---|---|
| SAEntry | 1 |
| SPI | 3001 |
| SADestIPAddr | FE80::2AA:FF:FE92:D0F1 |
| DestIPAddr | POLICY |
| SrcIPAddr | POLICY |
| Protocol | POLICY |
| DestPort | POLICY |
| SrcPort | POLICY |
| AuthAlg | HMAC-MD5 |
| KeyFile | Test.key |
| Direction | INBOUND |
| SecPolicyIndex | 2 |

Type a semicolon at the end of the entry configuring this SA. SA entries must be placed in decreasing numerical order.

8. On Host 2, create a text file that contains a text string that is used to authenticate the SAs created with Host 1. In this example, Test.key is created with the contents **This is a test**. There are no extra characters, spaces, or lines.

9. On Host 1, use the **ipsec6 l** command to add the configured security policies and SAs from the .spd and .sad files. In this example, the **ipsec6 l test** command is run on Host 1.

10. On Host 2, use the **ipsec6 l** command to add the configured security policies and SAs from the .spd and .sad files. In this example, the **ipsec6 l test** command is run on Host 2.

11. On Host 2, use the **ping** command to ping Host 1's link-local address.

    If you use Network Monitor to capture the traffic, you should see the exchange of ICMPv6 Echo Request and Echo Reply messages, with an Authentication Header (AH) listed between the IPv6 header and the ICMPv6 header.

For additional information about configurations, see IPv6 Configurations.

For information about using IPv6 in a test lab, see Setting up an IPv6 Test Lab.

# IPv6 traffic between nodes on different subnets of an IPv6 internetwork

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 traffic between nodes on different subnets of an IPv6 internetwork**

This configuration requires three computers (two hosts and one router) and additional router configuration beyond the installation of the IPv6 protocol. There are two separate network segments (also known as links or subnets) and an IPv6-capable router that forwards IPv6 packets between hosts on the segments.

The following illustration shows the configuration of two hosts on separate network segments that are connected by a router.




By default, the IPv6 protocol for the Windows Server 2003 family configures link-local IP addresses for each interface that correspond to installed Ethernet network adapters. Link-local addresses have the prefix of FE80::/64. The last 64 bits of the IPv6 address are the interface identifier, as derived from the 48-bit MAC address of the network adapter. For more information, see IPv6 interface identifiers. With link-local addresses, Host A and Host B can communicate with the router computer, but not with each other.

In this configuration, additional site-local prefixes are advertised by the router. The site-local prefixes are used by Host A and Host B to automatically configure site-local addresses that are derived from the 48-bit MAC address of the network adapter. After Host A and Host B have site-local addresses, they can communicate with each other.

To obtain the interface index numbers of the two Ethernet adapters on the router computer, type the following command for each adapter:

**netsh interface ipv6 show interface** [**interface=**]*String*

Where *String* is the name or index of the interface, (for example, **Local Area Connection**). Subnet 1 is the network segment to which Host A is attached. Subnet 2 is the network segment to which Host B is attached.

After you have obtained the interface index numbers, type the following commands on the router computer:

**netsh interface ipv6 set interface** [**interface=**]*Subnet1InterfaceIndex* [**forwarding=**]**enabled** [**advertise=**]**enabled**

**netsh interface ipv6 set interface** [**interface=**]*Subnet2InterfaceIndex* [**forwarding=**]**enabled** [**advertise=**]**enabled**

**netsh interface ipv6 add route** [**prefix=**]**FEC0:0:0:1::/64** [**interface=**]*Subnet1InterfaceIndex* [**publish=**]**yes**

**netsh interface ipv6 add route** [**prefix=**]**FEC0:0:0:2::/64** [**interface=**]*Subnet2InterfaceIndex* [**publish=**]**yes**

Where:

- *Subnet1InterfaceIndex* is the interface index of the adapter on the router computer that is attached to Subnet 1.
- *Subnet2InterfaceIndex* is the interface index of the adapter on the router computer that is attached to Subnet 2.

For example, if the Subnet 1 and Subnet 2 interface indexes of the router computer are 4 and 3 respectively, the commands are:

**netsh interface ipv6 set interface 4 enabled enabled**

**netsh interface ipv6 set interface 3 enabled enabled**

**netsh interface ipv6 add route FEC0:0:0:1::/64 4 publish=yes**

**netsh interface ipv6 add route FEC0:0:0:2::/64 3 publish=yes**

Wait about 30 seconds for both the router computer to advertise new site-local prefixes on Subnets 1 and 2 and for Hosts A and B to automatically configure site-local addresses based on these prefixes.

On Host A, type the **netsh interface ipv6 show interface** command to check for a new IPv6 address for the Ethernet adapter that is based on the site-local prefix of FEC0:0:0:1::/64. On Host B, type the **netsh interface ipv6 show interface** command to check for a new IPv6 address for the Ethernet adapter that is based on the site-local prefix of FEC0:0:0:2::/64.

On Host A, use the **ping** command and the site-local address of Host B to ping Host B. For example, if the Host B site-local address is FEC0::2:260:97FF:FE02:6EA5, the command is **ping FEC0::2:260:97FF:FE02:6EA5**.

For more information about Netsh commands, see Netsh overview and Netsh commands for Interface IPv6.

For additional information about configurations, see IPv6 Configurations.

For information about using IPv6 in a test lab, see Setting up an IPv6 Test Lab.

# IPv6 traffic between nodes on different subnets of an IPv4 internetwork

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 traffic between nodes on different subnets of an IPv4 internetwork**

The IPv6 protocol for the Windows Server 2003 family and Windows XP provides the following methods for communicating between IPv6 nodes on different subnets of an IPv4 internetwork:

- Intrasite Automatic Tunnel Addressing Protocol addresses
- 6over4

  6over4 requires that the IPv4 internetwork be multicast-capable. Because most IPv4 networks are not multicast-capable, 6over4 is rarely used. For more information about 6over4, see IPv6 protocol features and RFC 2529.

- 6to4

  While 6to4 is primarily designed to allow communication between separate IPv6-enabled 6to4 sites, 6to4 hosts that are using the IPv6 protocol can also use 6to4 addresses and 6to4 tunneling to communicate across the IPv4 Internet. For more information, see IPv6 traffic between nodes in different sites across the Internet (6to4).

- IPv4-compatible addresses

  IPv4-compatible addresses derived from IPv4 public addresses provide a method for connecting IPv6 hosts or sites over the existing IPv4 Internet infrastructure. For more information, see Using IPv4-compatible addresses.

In all of these cases, while IPv6 traffic is being carried as the payload of an IPv4 packet (treating the IPv4 infrastructure as an IPv6 link-layer), it remains IPv6 traffic. Applications that use the addresses associated with these methods are using the same Windows Sockets functions as if global IPv6 addresses and an IPv6 infrastructure were being used.

**Using Intrasite Automatic Tunnel Addressing Protocol addresses**

Another address assignment and tunneling mechanism that can be used for communication between IPv6/IPv4 nodes on an IPv4 network is described in the Internet draft titled "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)" (draft-ietf-ngtrans-isatap-O*x*.txt). These addresses are called Intrasite Automatic Tunnel Addressing Protocol (ISATAP) addresses. ISATAP addresses use the locally administered interface ID ::0:5EFE:*w.x.y.z* where:

- The 0:5EFE portion is formed from the combination of both the organizational unit identifier that is assigned to the Internet Assigned Numbers Authority (IANA) (00-00-5E) and a type that indicates an embedded IPv4 address (FE).
- The *w.x.y.z* portion is any unicast IPv4 address, which includes both public and private addresses.

The ISATAP interface ID can be combined with any 64-bit prefix that is valid for IPv6 unicast addresses. This includes the link-local address prefix (FE80::/64), site-local prefixes, and global prefixes.

Like IPv4-compatible addresses, 6over4 addresses, and 6to4 addresses, ISATAP addresses contain an embedded IPv4 address that is used to determine either the source or destination IPv4 addresses within the IPv4 header when ISATAP-addressed IPv6 traffic is sent across an IPv4 network.

By default, the IPv6 protocol for the Windows Server 2003 family and Windows XP automatically configures the ISATAP address of FE80::5EFE:*w.x.y.z* on the Automatic Tunneling Pseudo-Interface for each IPv4 address that is assigned to the node. This link-local ISATAP address allows two hosts to communicate over an IPv4 network by using each other's ISATAP addresses.

For example, Host A is configured with the IPv4 address of 10.40.1.29 and Host B is configured with the IPv4 address of 192.168.41.30. When the IPv6 protocol for the Windows Server 2003 family and Windows XP is started, Host A is automatically configured with the ISATAP address of FE80::5EFE:10.40.1.29 and Host B is automatically configured with the

ISATAP address of FE80::5EFE:192.168.41.30. When Host A sends IPv6 traffic to Host B by using Host B's ISATAP address, the source and destination addresses for the IPv4 and IPv6 headers are listed in the following table.

| Field | Value |
|---|---|
| Source address in IPv6 header | FE80::5EFE:10.40.1.29 |
| Destination address in IPv6 header | FE80::5EFE:192.168.41.30 |
| Source address in IPv4 header | 10.40.1.29 |
| Destination address in IPv4 header | 192.168.41.30 |

To test connectivity, you can use the **ping** command. For example, Host A would use the following command to ping Host B by using its link-local ISATAP address:

**ping FE80::5EFE:192.168.41.30%2**

The **%***ZoneID* portion of the command is used to specify the interface index of the interface from which traffic is sent. In this case, **%2** specifies interface 2, which is the interface index assigned to the Automatic Tunneling Pseudo-Interface on Host A.

The use of link-local ISATAP addresses allows IPv6/IPv4 hosts on an IPv4 intranet to communicate with each other, but not with other IPv6 hosts outside of the site. To communicate outside of the site, the following additional configuration is required:

- A host must receive a router advertisement from the site border router that contains a global address prefix. The site border router is placed between the intranet and either the IPv4 Internet or the IPv6 Internet. A site border router is most often a 6to4 router that is connected to the Internet. Upon receiving the router advertisement, additional ISATAP addresses that are based on the global prefix are automatically added.

  For example, if the site is connected to the IPv6 Internet and Host A receives the global prefix of 3FFE:FFF::/64 in a router advertisement, the ISATAP address of 3FFE:FFFF::5EFE:10.40.1.29 is automatically configured. Without a global address prefix and an IPv6 Internet connection, a site can use a 6to4-based global address prefix and connect to other 6to4 sites, 6to4 hosts, and the IPv6 Internet, by using the IPv4 Internet. If the site is using the 6to4 address prefix of 2002:836B:1:5::/64 (based on the public address of 131.107.0.1 and an SLA ID of 5), the ISATAP address of 2002:836B:1:5:0:5EFE:10.40.1.29 is automatically configured.

  ISATAP router addresses are auto-configured in Windows XP when IPv6 queries for "_ISATAP" and receives responses from ISATAP routers. For members of the Windows Server 2003 family and for Windows XP Service Pack 1, IPv6 queries for "ISATAP" (without the underscore character) and auto-configures ISATAP router addresses when responses from ISATAP routers are received.

- A host must have a default route, pointing to an ISATAP address that corresponds to the intranet interface of the site border router.

  For example, if the intranet interface of the site border router is configured with the IPv4 address of 172.16.0.1, Host A must be configured with a default route (::/0), which uses the ISATAP address of FE80::5EFE:172.16.0.1 as the next-hop address. As the result, all IPv6 traffic that matches this default route (as the closest matching route) is encapsulated and forwarded to the site border router. The site border router then forwards the traffic. If the site border router is a 6to4 router, it encapsulates the IPv6 traffic and forwards it on the Internet.

  For more information, see Add an IPv6 route.

For additional information about configurations, see IPv6 Configurations. For information about using IPv6 in a test lab, see Setting up an IPv6 Test Lab.

# IPv6 traffic between nodes in different sites across the Internet (6to4)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 traffic between nodes in different sites across the Internet (6to4)**

6to4 is a tunneling technique that is described in RFC 3056. When 6to4 is used, IPv6 traffic is encapsulated with an IPv4 header before it is sent over the IPv4 Internet.

6to4 uses the global address prefix of 2002:*WWXX*:*YYZZ*::/48, where *WWXX*:*YYZZ* is both the Next Level Aggregation ID (NLA ID) portion of a global address and the colon-hexadecimal representation of a public IPv4 address (*w.x.y.z*) that is assigned to the site or host. The complete 6to4 address of a 6to4 host is 2002:*WWXX*:*YYZZ*:[*SLAID*]:[*InterfaceID*].

RFC 3056 defines the following:

- 6to4 host

  An IPv6 host that is configured with at least one 6to4 address.

- 6to4 router

  An IPv4/IPv6 router that forwards 6to4-addressed traffic both between the 6to4 hosts within a site and to other 6to4 routers or 6to4 relay routers on the IPv6 Internet.

- 6to4 relay router

  An IPv4/IPv6 router that forwards 6to4-addressed traffic between 6to4 routers on the Internet and hosts on the IPv6 Internet.

When you use 6to4 hosts, an IPv6 routing infrastructure within 6to4 sites, a 6to4 router at site boundaries, and a 6to4 relay router, the following types of communication are possible:

1. A 6to4 host can communicate with another 6to4 host within the same site.

   This type of communication is available through the IPv6 routing infrastructure, which provides reachability to all hosts within the site.

2. A 6to4 host can communicate with 6to4 hosts in other sites across the IPv4 Internet.

   This type of communication occurs when a 6to4 host forwards IPv6 traffic, which is destined to a 6to4 host in another site, to the local site 6to4 router. The local site 6to4 router encapsulates the IPv6 traffic with an IPv4 header and sends it to the 6to4 router at the destination site on the Internet. At the destination site, the 6to4 router removes the IPv4 header and forwards the IPv6 packet to the appropriate 6to4 host by using the IPv6 routing infrastructure of the destination site.

3. A 6to4 host can communicate with hosts on the IPv6 Internet.

   This type of communication occurs when a 6to4 host forwards IPv6 traffic, which is destined for an IPv6 Internet host, to the local site 6to4 router. The local site 6to4 router encapsulates the IPv6 traffic with an IPv4 header and sends it to a 6to4 relay router that is connected to both the IPv4 Internet and the IPv6 Internet. The 6to4 relay router removes the IPv4 header and forwards the IPv6 packet to the appropriate IPv6 Internet host by using the IPv6 routing infrastructure of the IPv6 Internet.

All of these types of communication use IPv6 traffic without the requirement of obtaining either a direct connection to the IPv6 Internet or an IPv6 global address prefix from an Internet service provider (ISP).

Support for 6to4 hosts and routers is provided in the IPv6 Helper service that is included with the IPv6 protocol for Windows

Server 2003 family and Windows XP. The IPv6 Helper service:

- Automatically configures 6to4 addresses on the interface (named 6to4 Pseudo-Interface or interface index 3) that is used for all public IPv4 addresses that are assigned to interfaces on the computer. For more information, see Single subnet with link-local addresses.

- Automatically creates a 2002::/16 route that forwards all 6to4 traffic with the 6to4 Pseudo-Interface (interface index 3). All traffic forwarded by this host to 6to4 destinations is encapsulated with an IPv4 header.

- Automatically performs a Domain Name System (DNS) query for the name 6to4.ipv6.microsoft.com to obtain the IPv4 address of the Microsoft 6to4 relay router on the Internet. You can use the **netsh interface ipv6 6to4 set relay** command to specify which DNS name to query. For more information, see IPv6 utilities.

With the IPv6 Helper service configuration, computers are automatically configured as a 6to4 host when:

- The host is running the IPv6 protocol.

- The host is configured with an IPv4 public IP address.

- The host does not receive a Router Advertisement from an IPv6 router.

A 6to4 host can perform its own tunneling to reach either 6to4 hosts in other sites or hosts on the IPv6 Internet.

For more information about connecting to the IPv6 Internet, see Connecting to the IPv6 Internet.

If Internet Connection Sharing (ICS) is enabled on an interface that is assigned a public IPv4 address, the IPv6 Helper service:

- Enables routing on the private interface.

- Sends Router Advertisements that contain 6to4 address prefixes, which are based on the public IPv4 address of the public interface. The SLA ID in the 6to4 address prefix is set to the interface ID of the interface on which the advertisements are sent.

By enabling ICS, you can use a computer running the IPv6 protocol for Windows XP as a 6to4 router, which is capable of both encapsulating and forwarding 6to4 traffic to other 6to4 hosts or sites on the Internet, and forwarding IPv6 Internet traffic to a 6to4 relay router on the Internet.

The following illustration shows how 6to4 is used to communicate between two 6to4 sites.




Each site uses a computer with ICS enabled on the public interface to create a 6to4 router. Host computers on the private network segments receive the Router Advertisement that is sent by their site's 6to4 router and contains a 6to4 address prefix. As the result, two 6to4 hosts can communicate by using 6to4 addresses over the Internet.

For additional information about configurations, see IPv6 Configurations.

For information about using IPv6 in a test lab, see Setting up an IPv6 Test Lab.

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Connecting to the IPv6 Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connecting to the IPv6 Internet**

The simplest way to connect to the IPv6 Internet is to use 6to4, which is included with the IPv6 protocol for the Windows Server 2003 family. You can use 6to4 as either a 6to4 host or as a 6to4 router by enabling Internet Connection Sharing (ICS) on a computer that is connected to the Internet. The IPv6 Helper service automatically configures itself with the appropriate 6to4 addresses and uses a specific 6to4 relay router on the Internet. For more information, see IPv6 traffic between nodes in different sites across the Internet (6to4). The following illustration shows the configuration of a host that is using 6to4 to communicate on the IPv6 Internet through a 6to4 relay router.




With 6to4, you can ping other computers on the IPv6 Internet (for example, **ping -6 ipv6.research.microsoft.com**). The IPv6 Web site provides a list of IPv6-accessible servers that communicate on the IPv6 Internet. Additionally, the registry available at this site contains the names of other IPv6 Internet computers.

Some IPv6 Internet sites might be inaccessible. You might also experience connectivity problems. The **tracert -d***Address* command might be helpful in either case. The **-d** parameter prevents reverse DNS lookup on intermediate router addresses. For more information about the IPv6 Internet, see the 6bone Web site.

For additional information about configurations, see IPv6 Configurations.

For information about using IPv6 in a test lab, see Setting up an IPv6 Test Lab.

**Notes**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Using IPv4-compatible addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using IPv4-compatible addresses**

IPv4-compatible addresses, derived from IPv4 public addresses, provide a method for connecting IPv6 hosts or sites over the existing IPv4 Internet infrastructure. IPv6 traffic, when used with IPv4-compatible addresses, does not require the addition of IPv6 routers. Its traffic is encapsulated with an IPv4 header.

The following illustration shows the configuration of two nodes on separate subnets using IPv4-compatible addresses to communicate across an IPv4 router.




When IPv4-compatible IPv6 addresses are enabled, the IPv6 protocol for the Windows Server 2003 family and Windows XP automatically configures IPv4-compatible addresses for IPv4 public addresses on the Automatic Tunneling Pseudo-Interface (interface ID 2). An IPv4-compatible address takes the form ::*w.x.y.z*, where *w.x.y.z* is an IPv4 public address assigned to an interface on the computer. When IPv4-compatible addresses are enabled, the IPv6 protocol also automatically creates a ::/96 route that forwards all IPv4-compatible addressed traffic with the Automatic Tunneling Pseudo-Interface (interface ID 2). All traffic forwarded by this host to IPv4-compatible destinations is encapsulated with an IPv4 header.

IPv4-compatible addresses are disabled by default. To enable IPv4-compatible addresses, open Command Prompt and type:

**netsh interface ipv6 set state v4compat=enabled**

When you send traffic to an IPv4-compatible address, the traffic is sent from an IPv4-compatible address and is encapsulated with an IPv4 header. The Protocol field is set to 41 in the IPv4 header to indicate that the payload is an IPv6 packet. The IPv4 header allows the traffic to travel across an IPv4 infrastructure. The IPv4 addresses that are embedded in the source and destination IPv4-compatible addresses of the IPv6 header become the IPv4 source and destination addresses in the IPv4 header.

For example, when Host A (which is configured with the IPv4 address of 131.107.41.17) uses IPv4-compatible addresses to send IPv6 traffic to Host B (which is configured with the IPv4 address of 157.60.15.93), the source and destination addresses for the IPv4 and IPv6 headers are listed in the following table.

| Field | Value |
|---|---|
| Source address in IPv6 header | ::131.107.41.17 |
| Destination address in IPv6 header | ::157.60.15.93 |
| Source address in IPv4 header | 131.107.41.17 |
| Destination address in IPv4 header | 157.60.15.93 |

The IPv4 routing infrastructure forwards the packet from Host A to Host B, based on the IPv4 destination address of 157.60.15.93. After it is received at Host B, the IPv4 packet payload (the IPv6 packet) is passed to the IPv6 protocol.

To test connectivity, use the **ping** command. For example, Host A would use the following command to ping Host B by using its IPv4-compatible address:

**ping ::157.60.15.93**

For more information, see Add an IPv6 route.

For additional information about configurations, see IPv6 Configurations.

For information about using IPv6 in a test lab, see Setting up an IPv6 Test Lab.

# Setting up an IPv6 Test Lab

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up an IPv6 test lab**

This section provides information about how you can use five computers to create a test lab to configure and test the IPv6 protocol. These instructions are designed to walk you through a set of tasks, exposing you to the IPv6 protocol and its associated functionality. Beyond the set of tasks, these instructions allow you to create a functioning IPv6 configuration. You can use this configuration to learn about and experiment with IPv6 features and functionality, to aid in developing applications for IPv6 or porting existing IPv4 applications.

This section covers:

- Setting up the IPv6 infrastructure
- IPv6 test lab tasks

**Note**

- The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

# Setting up the IPv6 infrastructure

Updated: April 13, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up the infrastructure**

The infrastructure for the IPv6 test lab network consists of five computers performing the following services:

- A computer running Windows Server 2003, Standard Edition, that is used as a Domain Name System (DNS) server. This computer is named DNS1.
- A computer running Windows XP Professional that is used as a client. This computer is named CLIENT1.
- A computer running Windows Server 2003, Standard Edition, that is used as a router. This computer is named ROUTER1.
- A computer running Windows Server 2003, Standard Edition, that is used as a router. This computer is named ROUTER2.
- A computer running Windows XP Professional that is used as a client. This computer is named CLIENT2.

**Note**

- You can configure the computers named ROUTER1 and ROUTER2 with Windows XP Professional. The computer named DNS1, however, must be configured with Windows Server 2003, Standard Edition.

The following illustration shows the configuration of the IPv6 test lab.




There are three network segments:

- A network segment known as Subnet 1 that uses the private IP network ID of 10.0.1.0/24 and subnet ID of 2001:DB8:0:1::/64.
- A network segment known as Subnet 2 that uses the private IP network ID of 10.0.2.0/24 and subnet ID of 2001:DB8:0:2::/64.
- A network segment known as Subnet 3 that uses the private IP network ID of 10.0.3.0/24 and subnet ID of 2001:DB8:0:3::/64.

All computers on each subnet are connected to a separate common hub or Layer 2 switch. The two router computers, ROUTER1 and ROUTER2, have two network adapters installed.

For the IPv4 configuration, each computer is manually configured with the appropriate IP address, subnet mask, default gateway, and DNS server IP address. Dynamic Host Configuration protocol (DHCP) and Windows Internet Name Service (WINS) servers are not used. For the IPv6 configuration, link-local addresses are used initially.

The following sections describe how each of the computers in the test lab are configured. To reconstruct this test lab, please configure the computers in the order presented.

**Note**

- The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

**DNS1**

DNS1 is a computer running Windows Server 2003, Standard Edition. It is providing DNS server services for the testlab.microsoft.com DNS domain. To configure DNS1 for this service, perform the following steps:

1. Install Windows Server 2003, Standard Edition, as a stand-alone server. Set the Administrator password.

2. After restarting, log on as Administrator.

3. Install the IPv6 protocol. For more information, see Install IPv6.

4. Install the Domain Name System (DNS) server service. During installation of DNS, define a forward lookup zone named testlab.microsoft.com that allows dynamic updates. For more information, see Install a DNS server. You can also create a forward lookup zone after DNS installation. For more information, see Add a forward lookup zone

5. Configure the TCP/IP protocol with the IP address of 10.0.1.2, the subnet mask of 255.255.255.0, and the default gateway of 10.0.1.1. For more information, see Configure TCP/IP for static addressing.

**Note**

- When you install DNS using Configure Your server, the Configure Your DNS server Wizard runs after the DNS service is installed. The Configure Your DNS server Wizard prompts you to define either a forward or reverse lookup zone. To complete the test lab tasks, you must define a forward lookup zone that allows dynamic updates. The definition of a reverse lookup zone is optional.

**CLIENT1**

CLIENT1 is a computer running Windows XP Professional that is being used as a client. To configure CLIENT1 as a client computer, perform the following steps:

1. On CLIENT1, install Windows XP as a workgroup computer. Set the Administrator password.

2. After restarting, log on as Administrator.

3. Install the IPv6 protocol. For more information, see Install IPv6.

4. Configure the TCP/IP protocol with the IP address of 10.0.1.3, the subnet mask of 255.255.255.0, a default gateway of 10.0.1.1, and the DNS server IP address of 10.0.1.2. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

**ROUTER1**

ROUTER1 is a computer running Windows Server 2003, Standard Edition, that is being used as a router between Subnet 1 and Subnet 2. To configure ROUTER1 as a router, perform the following steps:

1. On ROUTER1, install Windows Server 2003, Standard Edition, as a workgroup computer. Set the Administrator password.

2. After restarting, log on as Administrator.

3. Install the IPv6 protocol. For more information, see Install IPv6.

4. For the interface on Subnet 1, configure the TCP/IP protocol with the IP address of 10.0.1.1, the subnet mask of 255.255.255.0, and the DNS server IP address of 10.0.1.2. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

5. For the interface on Subnet 2, configure the TCP/IP protocol with the IP address of 10.0.2.1, the subnet mask of 255.255.255.0, and a default gateway of 10.0.2.2. For more information, see Configure TCP/IP for static addressing.

6. Enable IP forwarding using Routing and Remote Access. For more information, see Enable the Routing and Remote Access service.

**ROUTER2**

ROUTER2 is a computer running Windows Server 2003, Standard Edition, that is being used as a router between Subnet 2 and Subnet 3. To configure ROUTER2 as a router, perform the following steps:

1. On ROUTER2, install Windows Server 2003, Standard Edition, as a workgroup computer. Set the Administrator password.

2. After restarting, log on as Administrator.

3. Install the IPv6 protocol. For more information, see Install IPv6.

4. For the interface on Subnet 2, configure the TCP/IP protocol with the IP address of 10.0.2.2, the subnet mask of 255.255.255.0, and a default gateway of 10.0.2.1. For more information, see Configure TCP/IP for static addressing.

5. For the interface on Subnet 3, configure the TCP/IP protocol with the IP address of 10.0.3.1, and the subnet mask of 255.255.255.0. For more information, see Configure TCP/IP for static addressing.

6. Enable IP forwarding using Routing and Remote Access. For more information, see Enable the Routing and Remote Access service.

**CLIENT2**

CLIENT2 is a computer running Windows XP that is being used as a client. To configure CLIENT2 as a client computer, perform the following steps:

1. On CLIENT2, install Windows XP as a workgroup computer. Set the Administrator password.

2. After restarting, log on as Administrator.

3. Install the IPv6 protocol. For more information, see Install IPv6.

4. Configure the TCP/IP protocol with the IP address of 10.0.3.2, the subnet mask of 255.255.255.0, and a default gateway of 10.0.3.1. For more information, see Configure TCP/IP for static addressing.

5. Verify the integrity of the IPv4 routing infrastructure by pinging 10.0.1.3 from the CLIENT2 computer.

# IPv6 test lab tasks

Updated: April 13, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 test lab tasks**

The following tasks are designed to take you through the common IPv6 configurations by using the test lab infrastructure configured in Setting up the IPv6 infrastructure.

- Link-local ping
- Creating a static routing infrastructure
- Using name resolution

To complete these tasks, you must use the Netsh commands for Interface IPv6 at the command prompt. For information about how to use Netsh, see Netsh overview. For the complete IPv6 Netsh reference, see Netsh commands for Interface IPv6.

**Note**

- The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

**Link-local ping**

To ping a host using link-local addresses and view the entries created in the neighbor and route caches, complete the following steps:

1. On DNS1, type the command **netsh interface ipv6 show address** to obtain the link-local address of the interface named **Local Area Connection**.

2. On CLIENT1, type the command **netsh interface ipv6 show address** to obtain the link-local address and interface index of the interface named **Local Area Connection**.

3. On CLIENT1, type the following command to ping the link-local address of DNS1:

   **ping***DNS1LinkLocalAddress***%***InterfaceIndex*

   For example, if the link-local address of DNS1 is FE80::2AA:FF:FE9D:10C5, and the interface index for the **Local Area Connection** interface on CLIENT1 is 3, the command is:

   **ping FE80::2AA:FF:FE9D:10C5%3**

4. On CLIENT1, type the following command:

   **netsh interface ipv6 show neighbors**

   to view the entry in the CLIENT1 neighbor cache for DNS1. For more information, see Neighbor Discovery (ND)

5. On CLIENT1, type the following command:

   **netsh interface ipv6 show destinationcache**

   to view the entry in the CLIENT1 destination cache for DNS1. For more information, see Neighbor Discovery (ND).

6. On CLIENT1, type the following command:

   **netsh interface ipv6 show routes**

   to view the entries in the CLIENT1 routing table. For more information, see The IPv6 routing table

**Creating a static routing infrastructure**

To configure a static IPv6 routing infrastructure so that all test lab nodes are reachable by using IPv6 traffic, complete the following steps:

1. On ROUTER1, type the **netsh interface ipv6 show address** command to obtain the interface indexes of the interfaces connected to Subnet 1 Connection, Subnet 2 Connection, and their link-local addresses.

2. On ROUTER2, type the **netsh interface ipv6 show address** command to obtain the interface indexes of the interfaces connected to Subnet 2 Connection, Subnet 3 Connection, and their link-local addresses.

3. On ROUTER1, type the following commands:

   **netsh interface ipv6 set interface"Subnet 1 Connection"forwarding=enabledadvertise=enabled**

   **netsh interface ipv6 set interface"Subnet 2 Connection"forwarding=enabledadvertise=enabled**

   **netsh interface ipv6 add route2001:DB8:0:1::/64interface="Subnet 1 Connection"publish=yes**

   **netsh interface ipv6 add route2001:DB8:0:2::/64"Subnet 2 Connection"publish=yes**

   **netsh interface ipv6 add route::/0"Subnet 2 Connection"nexthop=***ROUTER2AddressOnSubnet2***publish=yes**

   In the preceding command, *ROUTER2AddressOnSubnet2* represents the link-local address assigned to the Subnet 2 Connection interface on ROUTER2.

   For example, if the Subnet 2 Connection interface on ROUTER2 is FE80::2AA:FF:FE87:4D5C, the last command is typed as follows:

   **netsh interface ipv6 add route ::/0 "Subnet 2 Connection"nexthop=fe80::2aa:ff:fe87:4d5c publish=yes**

4. On ROUTER2, type the following commands:

   **netsh interface ipv6 set interface"Subnet 2 Connection"forwarding=enabledadvertise=enabled**

   **netsh interface ipv6 set interface"Subnet 3 Connection"forwarding=enabledadvertise=enabled**

   **netsh interface ipv6 add route2001:DB8:0:2::/64"Subnet 2 Connection"publish=yes**

   **netsh interface ipv6 add route2001:DB8:0:3::/64"Subnet 3 Connection"publish=yes**

   **netsh interface ipv6 add route::/0"Subnet 2 Connection"nexthop=***ROUTER1AddressOnSubnet2***publish=yes**

   In the preceding command, *ROUTER1AddressOnSubnet2* represents the link-local address assigned to the Subnet 2 Connection interface on ROUTER1.

For example, if the link-local address of the Subnet 2 Connection interface on ROUTER1 is FE80::2AA:FF:FE9A:203F, the last command should be typed as follows:

- **netsh interface ipv6 add route ::/0 "Subnet 2 Connection" nexthop=fe80::2aa:ff:fe9a:203f publish=yes**

On CLIENT1, type the command **netsh interface ipv6 show address** to obtain the address beginning with 2001:DB8

assigned to the interface named **Local Area Connection**.

On CLIENT2, type the following commands:

- **ping** *CLIENT1Address*
- **tracert -d** *CLIENT1Address*

**Using name resolution**

To view DNS AAAA (quad-A) records and test the use of DNS with IPv6, do the following:

1. On DNS1, open the DNS console. Under the server DNS1, click **Forward Lookup Zones**. In the right pane, open **testlab.microsoft.com**. In the right pane, you can view both IPv4 A records and IPv6 AAAA records. To view the quad-A record properties for one of the clients on your network, double-click the quad-A record.

   For example, for the CLIENT2 address of 2001:DB8::3:260:8FF:FE52:F9D8, the contents of the AAAA resource record are:

   Name: client2.testlab.microsoft.com

   Address: 2001:DB8::3:260:8FF:FE52:F9D8

2. On CLIENT1, type the following command:

   **ping -6 client2.testlab.microsoft.com**

   The name client2.testlab.microsoft.com is resolved to its address by sending a DNS query to DNS1.

3. On CLIENT2, create the following entry in the Hosts file (located in the *systemroot*\System32\Drivers\Etc folder):

   **client1***Client1Address*

4. On CLIENT2, type the following command:

   **ping -6 client1**

   The name client1.testlab.microsoft.com is resolved to its address by using the local Hosts file.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

# IPv6 Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This section covers:

- IPv6 RFCs and Internet drafts
- IPv6 references
- IPv6 utilities

# IPv6 RFCs and Internet drafts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 RFCs and Internet drafts**

The following table displays the RFCs and Internet drafts that are implemented by the IPv6 protocol for the Windows Server 2003 family and Windows XP.

| RFC number or Internet draft | Title |
|---|---|
| 1752 | The Recommendation for the IP Next Generation Protocol |
| 1828 | IP Authentication Using Keyed MD5 |
| 1886 | DNS Extensions to support IP version 6 |
| 1933 | Transition Mechanisms for IPv6 Hosts and Routers |
| 1981 | Path MTU Discovery for IP version 6 |
| 2185 | Routing Aspects of IPv6 Transition |
| 2373 | IP Version 6 Addressing Architecture |
| 2374 | An IPv6 Aggregatable Global Unicast Address Format |
| 2401 | Security Architecture for the Internet Protocol |
| 2402 | IP Authentication Header |
| 2403 | The Use of HMAC-MD5-96 within ESP and AH (implemented for AH only) |
| 2404 | The Use of HMAC-SHA-1-96 within ESP and AH (implemented for AH only) |
| 2406 | IP Encapsulating Security Payload (ESP) |
| 2428 | FTP Extensions for IPv6 and NATs |
| 2460 | Internet Protocol, Version 6 (IPv6) Specification |
| 2461 | Neighbor Discovery for IP Version 6 (IPv6) |
| 2462 | IPv6 Stateless Address Autoconfiguration |
| 2463 | Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification |
| 2464 | Transmission of IPv6 Packets over Ethernet Networks |
| 2467 | Transmission of IPv6 Packets over FDDI Networks |
| 2526 | Reserved IPv6 Subnet Anycast Addresses |
| 2529 | Transmission of IPv6 over IPv4 Domains without Explicit Tunnels |
| 2553 | Basic Socket Interface Extensions for IPv6 |
| 2710 | Multicast Listener Discovery (MLD) for IPv6 (implemented for host only) |

| | |
|---|---|
| 2711 | IPv6 Router Alert Option (implemented for host only) |
| 2732 | Format for Literal IPv6 Addresses in URLs |
| 3041 | Privacy Extensions for Stateless Address Autoconfiguration in IPv6 |
| 3056 | Connection of IPv6 Domains via IPv4 Clouds |
| Internet draft | An Extension of Format for IPv6 Scoped Addresses |
| Internet draft | Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) |
| Internet draft | Default Address Selection for IPv6 |
| Internet draft | IP Version 6 Addressing Architecture |
| Internet draft | IP Version 6 Scoped Address Architecture |
| Internet draft | Mobility Support in IPv6 (host only) |
| Internet draft | Routing of Scoped Addresses in the Internet Protocol Version 6 (IPv6) |
| Internet draft | Site Prefixes in Neighbor Discovery |
| Internet draft | The UDP Lite Protocol |

**Obtaining RFCs and Internet drafts**
You can obtain RFCs and Internet drafts from the Internet Engineering Task Force Web site.

### Note

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# IPv6 references

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 references**

For more information about the IPv6 protocol suite, you can consult the following references:


- Davies, J. *Understanding IPv6*. Redmond, WA: Microsoft Press, 2002.
- Huitema, C. *IPv6: The New Internet Protocol*. Second edition. Upper Saddle River, NJ: Prentice Hall, 1998.
- Miller, M. *Implementing IPv6: Supporting the Next Generation of Protocols*. Second edition. Foster City, CA: M&T Books, 2000.

# IPv6 utilities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPv6 utilities**
The following IPv6 command-line utilities are available for making configuration changes.

**Netsh commands for Interface IPv6**
You can configure the IPv6 protocol with the interface IPv6 context of Netsh.exe, which is used for both configuring and displaying the configuration of interfaces, addresses, caches, and routes.

Netsh commands for Interface IPv6 is a set of commands, each with its own set of parameters. In addition, you can use commands that are available in the **netsh interface IPv6 6to4** context to both configure or display the configuration of the IPv6 Helper service on either a 6to4 host or a 6to4 router.

For more information about using the IPv6 Helper service, see IPv6 traffic between nodes in different sites across the Internet (6to4).

For more information about Netsh, see Netsh overview.

For the complete Netsh interface IPv6 command reference, see Netsh commands for Interface IPv6.

**Ipsec6.exe**
You can use Ipsec6.exe to configure IPSec policies and security associations.

Ipsec6.exe has multiple commands, each with its own set of parameters.

- **ipsec6 sp** [*Interface*]

  Displays the active security policies. Alternately, displays the active security policies for a specific interface.

- **ipsec6 sa**

  Displays the active security associations.

- **ipsec6 l** *FileNameWithNoExtension*

  Loads the security policies from *FileName*.spd, and the security associations from *FileName*.sad.

- **ipsec6 s** *FileNameWithNoExtension*

  Saves the current security policies to *FileName*.spd, and the current security associations to *FileName*.sad. You can use this command to create files that are used to configure security policy and security associations. When there are no security policies or security associations, this command creates *FileName*.spd for security policies and *FileName*.sad for security associations. You can use these files as templates to configure the desired security policies or security associations by modifying them with a text editor.

- **ipsec6 d** [{**sp** | **sa**}] [*Index*]

  Deletes the security policies (using the **sp** parameter) or security associations (using the **sa** parameter) from the list of active security policies and security associations, as specified by index number. You can use **ipsec6 sp** or **ipsec6 sa** to display the index number.

- **ipsec6 m** [{**on** | **off**}]

  Specifies whether binding updates that are used for mobile IPv6 are protected by IP security. This is enabled by default.

**Caution**

- This implementation of IPSec for IPv6 is not recommended for use in a production environment because it relies on static keying and has no provisions for updating keys upon sequence number reuse.

- When you manually configure Security Policies Identifiers (SPIs), always use random numbers. Do not use sequential numbers for SPIs, or you will compromise the security of your IPSec for IPv6 policies.

- The IPv6 protocol for the Windows Server 2003 family does not support the use of IPSec Encapsulating Security Payload (ESP) encryption. However, the use of ESP with NULL encryption is supported. Although NULL encryption uses the ESP header, only data origin authentication and data integrity services are provided.

# Troubleshooting IPv6

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- A router is not advertising itself as a default router.
- A default route is not present on the host.
- Address has unexpected interface identifier.
- Internet Explorer is not connecting when literal IPv6 addresses are used in the URL.
- Off-link routes are in the routing table.
- Ping fails Echo Request messages when specifying a link-local destination.
- IPSec traffic is not encrypted.
- IPSec in IPv6 is not operating according to local or domain-based IPSec policies.
- Tunneled traffic is not reaching the destination.
- No IPv4-compatible address is configured on my Automatic Tunneling Pseudo-Interface.
- Unable to reach other 6to4 sites or the IPv6 Internet by using the 6to4 router.

**A router is not advertising itself as a default router.**
**Cause:** A computer running a member of the Windows Server 2003 family that is being used as an IPv6 router will not advertise itself as a default router unless it is configured with a default route (::/0) that is configured to be published.

**Solution:** Add a default route to the router computer and configure it to be published with the **netsh interface ipv6 add route** command.

**See also:** Netsh commands for Interface IPv6; Add an IPv6 route; IPv6 utilities; Setting up an IPv6 Test Lab

**A default route is not present on the host.**
**Cause:** A computer running a member of the Windows Server 2003 family that is being used as an IPv6 router will not advertise itself as a default router unless it is configured with a default route (::/0) that is configured to be published.

**Solution:** Add a default route to the router computer and configure it to be published with the **netsh interface ipv6 add route** command.

**See also:** Netsh commands for Interface IPv6; Add an IPv6 route; IPv6 utilities; Setting up an IPv6 Test Lab

**Address has unexpected interface identifier.**
**Cause:** A local router is configured to advertise a global prefix. By default, a temporary address that is based on the global prefix is automatically configured.

**Solution:** Use the **netsh interface ipv6 set privacy state=disabled** command to disable temporary addresses.

**See also:** Netsh commands for Interface IPv6; IPv6 interface identifiers; IPv6 utilities

**Internet Explorer is not connecting when literal IPv6 addresses are used in the URL.**
**Cause:** The version of the Internet Explorer provided with Windows Server 2003 family does not support the format for literal IPv6 addresses in URLs that is described in RFC 2732, "Format for Literal IPv6 Addresses."

**Solution:** Create DNS AAAA resource records that resolve Web server names to an IPv6 address and then use Web server names in the URL.

**See also:** IPv6 applications; IPv6 Name resolution

**Off-link routes are in the routing table.**
**Cause:** When a router running a member of the Windows Server 2003 family sends a Router Advertisement message, it includes Route Information options for directly attached subnets.

**Solution:** None. This is intended.

**See also:** IPv6 routing

**Ping fails Echo Request messages when specifying a link-local destination.**
**Cause:** The zone ID is not specified.

**Solution:** This is a common problem when a link-local destination address is used. Link-local addresses are often configured automatically for multiple interfaces. To specify the exact interface over which to send Echo Request messages, use the **ping***Address***%***ZoneID* syntax where *ZoneID* is the interface index for the interface over which the **ping** traffic is sent.

**See also:** Single subnet with link-local addresses; IPv6 utilities

**IPSec traffic is not encrypted.**
**Cause:** The IPv6 protocol for Windows Server 2003 family does not support the use of IPSec Encapsulating Security Payload (ESP) encryption. However, the use of ESP with NULL encryption is supported. Although NULL encryption uses the ESP header, only data origin authentication and data integrity services are provided.

**Solution:** None. This is a current limitation of the IPv6 protocol for Windows Server 2003 family and Windows XP.

**See also:** Security features for IPv6; Using IPSec between two local link hosts

**IPSec in IPv6 is not operating according to local or domain-based IPSec policies.**
**Cause:** IPSec support for IPv4 traffic is separate from IPSec support for IPv6 traffic. Local or domain-based IPSec policies configured with the IP Security Policies or Group Policy snap-ins are for IPv4 traffic only. These policies have no effect on IPv6 traffic.

**Solution:** None. This is a current limitation of the IPv6 protocol for Windows Server 2003 family and Windows XP.

**See also:** Security features for IPv6; Using IPSec between two local link hosts

**Tunneled traffic is not reaching the destination.**
**Cause:** Routers or firewalls are dropping IPv4 traffic that has the IP Protocol field value set to 41.

**Solution:** All IPv6 traffic that is encapsulated (tunneled) inside of an IPv4 header has the IPv4 Protocol field in the header set to 41. IPv6 tunneled traffic includes traffic that uses IPv4-compatible addresses, 6over4 addresses, ISATAP addresses, and 6to4 addresses. To allow IPv6 tunneled traffic to be forwarded, configure your routers or firewalls to pass IPv4 traffic that has the Protocol field set to 41.

**See also:** Security features for IPv6; Using IPSec between two local link hosts

**No IPv4-compatible address is configured on my Automatic Tunneling Pseudo-Interface.**
**Cause:** You do not have a public IPv4 address assigned to any of your interfaces or IPv4-compatible addresses are not enabled.

**Solution:** Because IPv4-compatible addresses have a global scope and are globally unique, IPv4-compatible addresses that are derived from private addresses are not allowed. Private addresses for the IPv6 protocol for Windows Server 2003 family and Windows XP are defined by the following ranges: 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16, and 169.254.0.0/16. Configure a public IPv4 address on one of your interfaces using the **netsh interface ipv6 set state** command.

**See also:** Netsh commands for Interface IPv6

**Unable to reach other 6to4 sites or the IPv6 Internet by using the 6to4 router.**
**Cause:** Your firewall or Internet router is dropping IPv4 traffic that has the IP Protocol field value set to 41.

**Solution:** All IPv6 traffic that is encapsulated (tunneled) inside of an IPv4 header has the IPv4 Protocol field in the header set to 41. IPv6 tunneled traffic includes traffic that uses IPv4-compatible addresses, 6over4 addresses, and 6to4 addresses. To allow IPv6 tunneled traffic to be forwarded, configure your firewall or Internet router to pass IPv4 traffic that has the Protocol field set to 41.

**Cause:** You are unable to resolve the DNS name 6to4.ipv6.microsoft.com.

**Solution:** By default, the IPv6 Helper service attempts to first resolve the name 6to4.ipv6.microsoft.com to its IPv4 addresses and then choose a relay router. If you cannot resolve the name 6to4.ipv6.microsoft.com, a relay router is not configured and you cannot reach any locations on the IPv6 Internet. Type **ping 6to4.ipv6.microsoft.com** to determine whether you can resolve the name 6to4.ipv6.microsoft.com.

**Cause:** You do not have the correct route.

**Solution:** View the display of the **netsh interface ipv6 show routes** command. You should see a **2002::/16** route that uses the 6to4 Tunneling Pseudo-Interface to send all traffic.

**Cause:** You do not have the correct address.

**Solution:** Verify that you have a 6to4 address assigned before attempting to reach a destination. For example, without a 6to4 address, you might be using a link-local address to reach a 6to4 global address. Use the display of the **netsh interface ipv6 show interface** command to determine your address configuration.

# Internet Protocol Security (IPSec)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol security (IPSec)**
Internet Protocol security (IPSec) is a framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services.

- For tips about using IPSec, see IPSec Best practices.
- For help with specific tasks, see IPSec How To....
- For general background information, see IPSec Concepts.
- For problem-solving instructions, see IPSec Troubleshooting.

### Note

- Portions of IPSec-related services were jointly developed by Microsoft® and Cisco Systems, Inc.

# IPSec Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Establish an IPSec deployment plan**.

  The deployment plan should address the following considerations: which deployment scenarios (such as server-to-server or remote access) require the use of IPSec, what level of security you require for each scenario, which types of data to secure, which computers to secure, which physical links to secure, who will manage IPSec policies, and how you will provide ongoing support and troubleshooting for end users after IPSec is deployed. The deployment plan should also address broader organizational security issues, such as the vulnerability of the network to specific types of attacks, the use of Active Directory, and how security is applied to Group Policy objects.

  For more information about IPSec security planning, see Establishing an IPSec security plan.

  For more information about Active Directory, see Active Directory Overview.

- **Create and test IPSec policies for each deployment scenario**.

  Before deploying IPSec in a production environment, test the IPSec policies in a realistic lab environment. To obtain realistic performance data, run standard workloads on programs. During initial tests, view packet contents with Network Monitor, or use Authentication Header (AH) or Encapsulating Security Payload (ESP) with null encryption to view packet contents for test environments.

- **Do not use preshared keys**.

  For enhanced security, the use of preshared key authentication is not recommended because it is a relatively weak authentication method. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

  For more information, see Preshared key authentication.

- **Do not use Diffie-Hellman Group 1 (low)**.

  For enhanced security, do not use Diffie-Hellman Group 1, which provides 768 bits of keying strength. For maximum security, use Group 2048 (high), which provides 2,048 bits of keying strength. Use Group 2 (medium), which provides 1,024 bits of keying strength when required for interoperability with Windows 2000 and Windows XP. Strong Diffie-Hellman groups combined with longer key lengths increase the computational difficulty of determining a secret key.

  For more information, see Key exchange methods.

  **Note**

  - Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.

- **Use the Triple Data Encryption Standard (3DES) algorithm for stronger encryption**.

  For enhanced security, when configuring key exchange security methods for IPSec policies, use 3DES, which is a stronger encryption algorithm than DES.

  For information about how to configure key exchange security methods, see Create key exchange security methods.

  **Note**

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

- **Create and assign a persistent IPSec policy for failsafe security**.

  To enhance security, create and assign a persistent IPSec policy, so that computers can be secured if a local IPSec policy or an Active Directory-based IPSec policy cannot be applied. When you create and assign a persistent policy, it is applied before a local policy or an Active Directory-based policy, and it remains in effect regardless of whether the local policy or the Active Directory-based policy is applied (for example, an IPSec policy will not be applied if it is corrupted).

  **Note**

  - You cannot configure this feature in the IP Security Policy Management console. To configure this feature, you must use the Netsh IPSec command-line tool. For more information, see Netsh commands for Internet Protocol security.

- **For computers connected to the Internet, do not send the name of the certification authority (CA) with certificate requests**.

  When certificate authentication is used to establish trust between IPSec peers, each IPSec peer sends to the other peer a list of trusted root CAs from which it accepts a certificate for authentication. Each of these CA names is sent as a certificate request payload (CRP), and it must be sent before trust is established. Although transmitting this list aids in connectivity by facilitating the selection of a CA, it can expose sensitive information about the trust relationships of a computer, such as the name of the company that owns the computer and the domain membership of the computer (if an internal public key infrastructure is being used), to an attacker. Therefore, to secure computers that are connected to the Internet, enable the option to exclude the CA name from the certificate request.

- **For computers connected to the Internet, do not use Kerberos as an authentication method**.

  When Kerberos V5 authentication is used, during main mode negotiation, each IPSec peer sends its computer identity in unencrypted format to the other peer. The computer identity is unencrypted until encryption of the entire identity payload takes place during the authentication phase of the main mode negotiation. An attacker can send an Internet Key Exchange (IKE) packet that causes the responding IPSec peer to expose its computer identity and domain membership. To secure computers that are connected to the Internet, certificate authentication is recommended.

  For more information, see Authentication methods.

- **For computers that are connected to the Internet, do not allow unsecured communication**.

  If you configure a filter action to negotiate IPSec security, ensure that the following options are disabled in order to secure computers that are connected to the Internet:

  - **Accept unsecured communication, but always respond using IPSec**. This option allows initial incoming unsecured traffic (for example, TCP SYN packets) but requires outgoing traffic to be protected. This option should be disabled to prevent denial-of-service attacks.

  - **Allow unsecured communication with non-IPSec-aware computers**. This option allows unsecured communications with computers that cannot negotiate the use of IPSec or process IPSec-secured communications, and it is appropriate only in environments where IPSec-secured communication is not necessary.

    For more information, see Filter action.

- **Restrict the use of administrative credentials in your organization**.

  Members of the local Administrators group can view and modify the IPSec policy settings on their computers. For this reason, and as a general security best practice, ensure that end-users in your organization use the principle of least privilege.

  For information about general best practices for security, see Best practices for security for security.

- **When applying the same IPSec policy to computers running different versions of the Windows operating system, test the policy thoroughly**.

  Several features in the Windows Server 2003 family implementation of IPSec are not provided in Windows 2000 or in Windows XP. These new features include:

  - A stronger Diffie-Hellman group (2,048-bit Diffie-Hellman key exchange).
  - The ability for IPSec policy filters to allow logical addresses for local IP configuration.
  - The removal of default traffic exemptions from IPSec filtering (only Internet Key Exchange traffic is now exempt by default).
  - The ability to exclude the name of the certification authority (CA) from a certificate request.

  To ensure that the same IPSec policy functions as expected on computers running the Windows Server 2003 family and on computers running Windows XPor Windows 2000, test the policy thoroughly on all relevant operating systems before deployment.

  For more information about new features in IPSec, see New features for IPSec.

- **Use the** Windows Server 2003 **IP Security Policy Management console to manage the IPSec policies that use the new features that are available only in the** Windows Server 2003 **family implementation of IPSec**.

  If you plan to apply IPSec policies that use any of the new features that are available only in the Windows Server 2003 family implementation of IPSec, do not use the Windows XP or the Windows 2000 version of the IP Security Policy Management console to manage these policies. The settings in the earlier versions of the IP Security Policy Management console will override the settings in the Windows Server 2003 family IPSec policy, and the new features will not be functional.

- **Use Terminal Services to remotely manage and monitor IPSec on computers running different versions of the Windows operating system**.

  Remote management and monitoring of IPSec is supported only for computers running the same version of the Windows operating system. To remotely manage and monitor IPSec on a computer that is running a different version of Windows than the version of Windows that is running on your computer, use Terminal Services. For example, if your computer is running the Windows Server 2003 family, and you plan to remotely manage and monitor IPSec on computers running Windows 2000 or Windows XP, use Terminal Services to gain remote access to these computers.

  If your computer is running the Windows Server 2003 family, and you plan to manage and monitor IPSec on computers that are also running the Windows Server 2003 family, you can run the IPSec Policy Management console and IP Security Monitor, or you can use the Netsh command-line tool remotely.

# IPSec How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Start the IP Security Policy Management snap-in
- Define IPSec Policies
- Assign IPSec Policy
- Manage IPSec Policy
- Use Resultant Set of Policy (RSoP) to View IPSec Policy Assignments
- Monitor IPSec Activity

# Start the IP Security Policy Management snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the IP Security Policy Management snap-in**

- To start the IP Security Policy Management snap-in from the Microsoft Management Console:

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **IP Security Policy Management**, and then click **Add**.
4. Select the computer for which you want to manage IPSec policies:

| To | Do this |
|---|---|
| Manage only the computer on which this console is running | Click **Local computer** |
| Manage IPSec policies for any domain members | Click **The Active Directory domain of which this computer is a member**. |
| Manage IPSec policies for a domain of which the computer that is running this console is not a member | Click **Another Active Directory Domain**. |
| Manage a remote computer | Click **Another computer**. |

5. Click **Finish**, click **Close**, and then click **OK**.

- To access the IP Security Policy Management snap-in from Group Policy (Active Directory):

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

   **Where?**

   - Active Directory Users and Computers [*DomainControllerName.DomainName*]/*Domain*/*OrganizationalUnit*/*ChildOrganizationalUnit*...

3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the Group Policy object that you want to edit. Or, click **New** to create a new Group Policy object, and then click **Edit**.
5. In the Group Policy console tree, click **IP Security Policies on Active Directory**.

   **Where?**

   - *PolicyName* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/IP Security Policies on Active Directory

- To access the IP Security Policy Management snap-in from Local Computer Policy:

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. Click **Finish**, click **Close**, and then click **OK**.
5. In the Group Policy console tree, click **IP Security Policies on Local Computer**.

   **Where?**

   - Local Computer Policy/Computer Configuration/Windows Settings/Security Settings/IP Security Policies on Local Computer

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To save console settings, on the **File** menu, click **Save**.
- To view the saved console, see Related Topics.
- You cannot administer Active Directory-based IPSec policy from a computer running Windows XP Home Edition.
- To manage policies in a remote domain, you must be using a computer that is a member of a domain that is trusted by the remote domain. You cannot configure policies in a remote domain from a computer that is a member of a workgroup (also known as a stand-alone computer).
- To revise your initial choice, you can start the Microsoft Management Console, add this snap-in again, and save the console again. If you require multiple configurations, you can save the console with another name.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add, edit, or remove IPSec policies
Policy-based security
Open a saved MMC console for local and remote computers
Save an MMC console file

# Define IPSec Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define IPSec policies**

- Add, edit, or remove IPSec policies
- Define IPSec Rules
- Define IPSec Key Exchange Settings

# Add, edit, or remove IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.

2. Choose whether to add a new policy, or edit or remove an existing one:

| To | Do this |
|---|---|
| Add a new policy | In the console tree, click **IP Security Policies on *Name***. Next, on the **Action** menu, click **Create IP Security Policy**. Follow the instructions in the IP Security Policy Wizard until the **Properties** dialog box for your new policy appears. |
| Modify an existing policy | Double-click the policy that you want to modify. |
| Remove a policy | Click the policy that you want to remove, and on the **Action** menu, click **Delete**. |

3. If you are adding or modifying a policy, click the **General** tab, and then in **Name**, type a unique name.

4. In **Description**, type a description of the security policy, such as which groups or domains it affects.

5. To specify how often the IPSec Policy Agent checks for updates, type a value in **Check for policy changes every numberminute(s)**.

6. If you have special requirements for the security of the key exchange, click **Settings**.

7. Click the **Rules** tab, and create or modify any necessary rules for the policy.

8. Activate or deactivate rules as needed.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- For information about defining IPSec key exchange settings, see Related Topics.

- For information about configuring IPSec rules, see Related Topics.

- For information about how to activate or deactivate rules, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec rules
Activate or deactivate IPSec rules
Working with MMC console files
Define IPSec Key Exchange Settings

# Define IPSec Rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define IPSec rules**

- Add, edit, or remove IPSec rules
- Activate or deactivate IPSec rules
- Define IPSec authentication methods
- Specify IPSec connection types
- Specify an IPSec tunnel
- Define IP Filter Lists
- Define Filter Actions

# Add, edit, or remove IPSec rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IPSec rules**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. To add a rule, decide whether you want to use the Create IP Security Rule Wizard or add the rule manually:
   - To add a rule by using the Create IP Security Rule Wizard, confirm that the **Use Add Wizard** check box is selected, click **Add**, and then follow the instructions.
   - To add a rule manually, confirm that the **Use Add Wizard** check box is cleared, click **Add**, and then define settings on the **IP Filter List**, **Filter Action**, **Authentication Methods**, **Tunnel Setting**, and **Connection Type** tabs.
4. To edit a rule, select the rule that you want to edit, click **Edit**, and then modify the rule properties as needed.
5. To remove a rule, select the rule that you want to remove, and then click **Remove**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- New rules are automatically applied to the policy being edited or created.
- The rules for a policy are displayed in IP Security Policies in reverse alphabetical order, based on the name of the filter list selected for each rule. There is no method to specify an order in which to apply the rules in a policy. The IPSec driver automatically orders the rules based on the most specific filter list to the least specific filter list. For example, the IPSec driver places a rule containing a filter list that specifies individual IP addresses and TCP ports above a rule containing a filter list that specifies all addresses on a subnet.
- The default response rule is automatically added to every new IPSec policy. If you do not want this rule to be part of your policy, you can deactivate it by clearing the check box next to **<Dynamic>**. The default response rule cannot be removed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC

Activate or deactivate IPSec rules
Add, edit, or remove filter actions
Define IPSec authentication methods
IPSec Policy Rules
Working with MMC console files

# Activate or deactivate IPSec rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate or deactivate IPSec rules**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. On the **Rules** tab, do one of the following:
   - To activate a rule, select the check box next to the rule that you want to activate.
   - To deactivate a rule, clear the check box next to the rule that you want to deactivate.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Assign or unassign IPSec policy in Group Policy
Assign or unassign IPSec policy on a computer
Working with MMC console files

# Define IPSec authentication methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define IPSec authentication methods**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify.
4. On the **Authentication Methods** tab, click **Add**. Or, if you are reconfiguring an existing method, click the authentication method, and then click **Edit**.
5. Select the authentication method that you want to add or modify:
   - To use the Kerberos V5 security protocol for authentication services, click **Active Directory default (Kerberos V5 protocol)** if this rule applies to computers that are validated by a Windows 2000 or Windows Server 2003 Active Directory domain or a trusted Active Directory domain.
   - To use a public key certificate for authentication services, click **Use a certificate from this certification authority (CA)**, and then click **Browse** to select a root CA.
   - To prevent the name of the CA from being sent with the certificate request, select the **Exclude the CA name from the certificate request** check box.
   - To enable certificate to account mapping, select the **Enable certificate to account mapping** check box.
   - To specify your own key that will be used for authentication, click **Use this string (preshared key)**.
6. To delete an authentication method or change its preference order, do one of the following:
   - To delete the selected method, click **Remove**.
   - To move the selected method up one level, click **Move up**. Repeat until the method is at the required preference level.
   - To move the selected method down one level, click **Move down**. Repeat until the method is at the required preference level.

**Important**

- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a

computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- The Kerberos V5 protocol authentication method is not supported on computers running Windows XP Home Edition.

- If you choose to use a certificate for authentication, you must select a CA (most commonly the root CA for your installed computer certificate). You cannot leave this field blank.

- For preshared key authentication, each IPSec peer must use the same preshared key, or authentication will fail.

- If filters in more than one rule apply to traffic between the same two IP addresses, then the authentication methods list should be identical in both rules. Otherwise, security requests that are initiated by one rule might match the filter of the other rule, but have a different authentication method. In this case, the negotiation will fail. After a rule has initiated a successful main mode security association between two IP addresses, other rules that are filtering different traffic will not renegotiate the authentication and master key (the main mode security association). Instead, these other rules will use the same main mode security association to negotiate security (algorithms and session keys) for traffic that matches their particular filter.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Authentication methods
Working with MMC console files

# Specify IPSec connection types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify IPSec connection types**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify.
4. On the **Connection Type** tab, select the type of network connections to which this rule will apply:
   - To apply this rule to all network connections you have created on the computer, click **All network connections**.
   - To apply this rule to all LAN connections you have created on the computer, click **Local area network (LAN)**.
   - To apply this rule to any remote VPN or dial-up connections you have created on the computer, click **Remote access**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Connection type
Working with MMC console files

# Specify an IPSec tunnel

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify an IPSec tunnel**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify.
4. On the **Tunnel Setting** tab, specify the computer that will be the tunnel endpoint:

| To | Do this |
|---|---|
| Disable tunneling for this rule | Click **This rule does not specify an IPSec tunnel**. |
| Use tunneled communications to a specific tunnel endpoint | Click **The tunnel endpoint is specified by this IP address** and type the IP address of the tunnel endpoint. |

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- Because you cannot mirror filters for tunneled traffic, you must configure two rules. One rule is used for the outbound traffic and the other is used for inbound traffic. For the outbound traffic rule, the tunnel endpoint is the IP address of the IPSec peer on the other end of the tunnel. For the inbound traffic rule, the tunnel endpoint is an IP address configured on the local computer.
- IPSec tunnels are not supported for remote access VPN scenarios. L2TP/IPSec or PPTP should be used for remote access VPN connections. IPSec tunnels are used primarily for interoperability with other routers, gateways, or end-systems that do not support L2TP/IPSec or PPTP connections. IPSec tunnel mode is supported as an advanced feature, used only in gateway-to-gateway (also known as router-to-router) tunneling scenarios and for server-to-server or server-to-gateway configurations.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Virtual private networking with IPSec
Tunnel endpoint

# Define IP Filter Lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define IP filter lists**

- Add, edit, or remove IP filter lists
- Add, edit, or remove IPSec filters
- Configure advanced IPSec filter settings
- Set an active IP filter list for a rule
- Manage multiple IP filter lists

# Add, edit, or remove IP filter lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IP filter lists**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that contains the IP filter list that you want to modify.
4. On the **IP Filter List** tab, do one of the following:
   - To add a new IP filter list, click **Add**.
   - To modify an existing IP filter list, click the filter list that you want to modify, and then click **Edit**.
   - To remove an IP filter list, click the filter list that you want to remove, and then click **Remove**.
5. If you are adding or modifying an IP filter list, type a unique name for the list.
6. Type a description to specify which computers and traffic types the IP filter list encompasses (for example, "Matches all TCP packets between this computer and any other computer").
7. Create or modify the filters as needed.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- For procedures on creating filters and advanced filter settings, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Manage multiple IP filter lists
Add, edit, or remove IPSec filters
Configure advanced IPSec filter settings
Filter list

# Add, edit, or remove IPSec filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IPSec filters**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that contains the IP filter list you want to modify.
4. On the **IP Filter List** tab, double-click the IP filter list that contains the IPSec filter you want to modify.
5. In the **IP Filter List** dialog box, do one of the following:
   - To add a filter, click **Add**.
   - To modify an existing filter, select the filter that you want to modify, and then click **Edit**.
   - To remove an existing filter, select the filter that you want to remove, and then click **Remove**.
6. If you are adding or editing a filter, in **IP Filter Properties**, the **Addresses** tab, and then select the **Source Address**:

| Select | To secure packets from |
|---|---|
| **My IP Address** | All IP addresses on the computer for which you are configuring this filter. |
| **Any IP Address** | Any computer. |
| **A Specific DNS Name** | The Domain Name System (DNS) name that you specify in **Host name**. The DNS name is resolved to its IP addresses, and then filters are automatically created for the resolved IP addresses. This option is only available when creating new filters. |
| **A Specific IP Address** | The IP address that you specify in **IP Address**. |
| **A Specific IP Subnet** | The IP subnet, as defined by the IP address that you specify in **IP Address** and the subnet mask that you specify in **Subnet Mask**. |
| **DNS Servers <dynamic>** | The DNS server(s) for the computer for which you are configuring this filter. Changes in DNS server addresses are automatically detected, and the filter is updated as needed. |
| **WINS Servers <dynamic>** | The WINS server(s) for the computer for which you are configuring this filter. Changes in WINS server addresses are automatically detected, and the filter is updated as needed. |
| **DHCP Server <dynamic>** | The DHCP server for the computer for which you are configuring this filter. Changes in DHCP server addresses are automatically detected, and the filter is updated as needed. |
| **Default Gateway <dynamic>** | The default gateway for the computer for which you are configuring this filter. Changes in default gateway server addresses are automatically detected, and the filter is updated as needed. |

7. Click **Destination Address**, and repeat step 6 for the destination address.

8. Select the appropriate **Mirrored** setting:

| To | Do this |
|---|---|
| Automatically create two filters based on the filter settings, one for traffic to the destination and one for traffic from the destination | Select the **Mirrored** check box. |
| Create a single filter based on the filter settings | Clear the **Mirrored** check box. |
| Create a filter for an IPSec tunnel | Clear the **Mirrored** check box. For IPSec tunnels, you must create two filter lists: one list describes the traffic to be sent through the tunnel (outbound traffic) and another describes the traffic to be received through the tunnel (inbound). Then, create two rules that use the inbound and outbound filter lists in your policy. |

9. On the **Description** tab, in **Description**, type a description for this filter (for example, specify to what computers and traffic types it applies).

10. If you require additional IP filtering by a specific protocol or port number, on the **Protocol** tab, configure advanced filter settings.

**Caution**

Filters are the most important part of IPSec policy for a computer which is protected by IPSec. If you do not specify the proper filters in either client or server policies, or if the IP addresses change before the policy's filters are updated, security might not be provided.

- When adding a new static IP address to a protected computer:
    - Modify the IPSec policy filters on all clients and servers that make security requests to the protected computer. Ensure that those clients have updated their policy before adding the new address.
    - Inspect the policy being used on the protected computer. If the filters specify static IP addresses for local connections, after adding the new IP address to the interface, edit and save the new filter list to include the new static IP address. **My IP Address** filters will be automatically updated when the new static IP address is added.
- If the protected computer is a Web server and your clients use a proxy server, make sure that communications over all network paths are secured by IPSec:
    - Between the Web server and all clients that directly connect
    - Between the Web server and the proxy server
    - Between the proxy server and all clients of the proxy server
- The following filter is internally defined to permit (not secure) Internet Key Exchange (IKE) traffic:
    - Source Address = Any
    - Destination Address = Any
    - Protocol = UDP
    - Source port = 500

- Destination port = 500

**Important**

- In the Windows Server 2003 family, IPSec provides new source and destination address options for defining filters. In addition, the default behavior for traffic exemptions has changed. If you require previously exempted traffic types, you must create filters to permit this traffic. For instructions on how to configure permit filters for previously exempted traffic types and information about other IPSec special considerations, see Related Topics.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- To configure IP protocol, TCP port, or UDP port settings for a filter, see Related Topics.

- For IPSec tunnel rules, only address-based filters are supported. Protocol-specific and port-specific filters are not supported. Tunnel filters should not be mirrored.

- Filters are applied in the order of most-specific filters first. Filters are not applied in the order in which they appear in the list.

- If an outbound packet does not match any filter, it is sent unsecured.

- If an inbound packet does not match any filter, it is permitted.

- If an IKE security request is received, the source IP address of the request is used to find a matching filter. The security action and tunnel setting that is associated with that filter determines the IKE response.

- The **A specific DNS Name** option is used when creating filters to create IP address-based filters by resolving a DNS name to its IP addresses. The use of a computer name is for the one-time resolution of DNS names to IP addresses when creating the filter. The computer name is not used after the filters are created.

- All filters used in tunnel rules are matched first, before end-to-end transport filters are matched.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Special IPSec considerations
Start the IP Security Policy Management snap-in
Open MMC
Configure advanced IPSec filter settings
Filter list
Working with MMC console files
Virtual private networking with IPSec

# Configure advanced IPSec filter settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure advanced IPSec filter settings**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that contains the IP filter list you want to modify.
4. On the **IP Filter List** tab, double-click the IP filter list that contains the IPSec filter you want to modify.
5. In the **IP Filter List** dialog box, click the filter that you want to modify, and then click **Edit**.
6. In **IP Filter Properties**, click the **Protocol** tab, and then specify settings for **Select a protocol type**:

| To filter | Do this |
|---|---|
| Any packets sent on any protocol | Click **Any**. |
| Traffic sent on a specific protocol | Click the protocol. |
| Packets sent on a custom protocol | Click **Other**, and then type the protocol number. |

7. If you chose TCP or UDP, you can also filter packets by the source port.

   - To filter packets sent on any port used by the protocol type that you selected, click **From any port**.
   - To filter packets sent on a specific port used by the protocol type that you selected, click **From this port**, and then type the port number.

   If you chose TCP or UDP, you can also filter packets by the destination port.

   - To filter packets received on any port used by the protocol type that you selected, click **To any port**.
   - To only filter packets received on the port number that you specify, click **To this port**, and then type the port number.

**Caution**

- An erroneous port number might cause communication to be blocked or unsecured. If you are unsure of the correct value, see the Windows Resource Kits in Related Topics.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- For IPSec tunneling, only address-based filters are supported. Protocol-specific and port-specific filters are not supported.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec filters
Filter list
Working with MMC console files
IPSec Resources

# Set an active IP filter list for a rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set an active IP filter list for a rule**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that contains the IP filter list you want to modify.
4. On the **IP Filter List** tab, select the option that precedes the name of the IP filter list that you want to designate as the active filter list for this rule.

   The IP filter list you selected is now the active filter list for the rule.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- If the option preceding the name of an IP filter list is not selected, the IP filter list will not be set as active for the rule.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec filters
Configure advanced IPSec filter settings
Manage multiple IP filter lists
Working with MMC console files

# Manage multiple IP filter lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage multiple IP filter lists**

1. In **IP Security Policies**, in the console tree, click **IP Security Policies on *Name***.
2. On the **Action** menu, click **Manage IP filter lists and filter actions**.
3. On the **Manage IP Filters Lists** tab, click **Add**. Or, if you are modifying an existing list, double-click that list.
4. Define the filters for the filter list.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec filters
Configure advanced IPSec filter settings
Set an active IP filter list for a rule
Filter list
Working with MMC console files

# Define Filter Actions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define filter actions**

- Add, edit, or remove filter actions
- Add, edit, or remove IPSec security methods
- Configure custom IPSec security methods
- Select a filter action for a rule
- Manage multiple filter actions

# Add, edit, or remove filter actions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove filter actions**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify, and then click the **Filter Action** tab.
4. To add a filter action, decide whether you want to use the IP Security Filter Action Wizard or add the filter action manually:
   - To add a filter action by using the IP Security Filter Action Wizard, confirm that the **Use Add Wizard** check box is selected, click **Add**, and then follow the instructions.
   - To add a filter action manually, confirm that the **Use Add Wizard** check box is cleared, click **Add**, and then define settings on the **Security Methods** and **General** tabs.
5. To modify an existing filter action, select the filter action that you want to modify, and then click **Edit**.
6. To remove a filter action, select the filter action that you want to remove, and then click **Remove**.
7. If you are adding or modifying a filter action, choose a filter action type:
   - Click **Permit** to allow receiving or sending of packets in plaintext (unsecured traffic). Security will not be requested for these packets.
   - Click **Block** to discard packets. Security will not be requested for these packets.
   - Click **Negotiate security** to use the list of security methods in **Security method preference order** to provide security for packets that match this filter. Security requests will be accepted for these packets.
8. If you chose **Negotiate security**, add new security methods or edit the existing ones for the filter action.
9. If you do not want to block incoming, unsecured communications, but you want to ensure that all outgoing communications and subsequent two-way communications are secured, select the **Accept unsecured communication, but always respond using IPSec** check box.
10. To enable communication with other computers that do not support IPSec, and ensure that communication continues if there is no response to a request for IPSec negotiation, select the **Allow unsecured communication with non-IPSec-aware computers** check box. After the initial IPSec negotiation has failed, IPSec negotiation will be retried at five minute intervals.
11. To guarantee that no master keys or master keying material will be reused to generate the session key, select the **Use session key perfect forward secrecy (PFS)** check box.
12. On the **General** tab, in **Name**, type a unique name.
13. In **Description**, type a description. For example, you might type what security levels this filter action represents.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- The filter action types **Permit** and **Block** and the options to **Allow unsecured communication with non-IPSec-aware computers** and **Accept unsecured communication, but always respond using IPSec** are not available in the default response rule.

- Session key perfect forward secrecy (PFS) renegotiates new master key keying material during every session rekey operation. This is the most secure setting. However, it adds overhead and increases the time required to complete a rekey operation. Both initiator and responder must be configured with the same setting for successful negotiation.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Select a filter action for a rule
Add, edit, or remove IPSec security methods
Filter action
Working with MMC console files

# Add, edit, or remove IPSec security methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove IPSec security methods**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify, and then click the **Filter Action**tab.
4. Double-click the filter action that you want to modify.
5. On the **Security Methods** tab, click **Negotiate security**, and then do one of the following:
   - To add a new security method, click **Add**.
   - To modify an existing security method, click the security method that you want to modify, and then click **Edit**.
   - To remove a security method, click the security method that you want to remove, and then click **Remove**.
6. If you are adding or modifying a security method, on the **Security Method** tab, select the security level:
   - To use the ESP protocol to provide data confidentiality (encryption) with the triple Data Encryption Standard (3DES) algorithm, data integrity and authentication with the Secure Hash Algorithm 1 (SHA1) integrity algorithm, and default key lifetimes (100MB, 1 hour), click **Integrity and Encryption**.
   - To use the ESP protocol to provide data integrity and authentication with the SHA1 integrity algorithm and default key lifetimes (100MB, 1 hour), click **Integrity only**. ESP is not configured to provide data confidentiality (encryption).
   - Click **Custom**, and then click **Settings** to configure a custom security method or key lifetimes. For more information, see Related Topics.
7. Repeat this procedure to configure additional security methods for the filter action.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should

only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 server operating system support 3DES and do not require installation of the High Encryption Pack.

- Click **Move up** to move the selected security method up one level. Repeat until the security method is at the required preference level.

- Click **Move down** to move the selected security method down one level. Repeat until the security method is at the required preference level.

- For information about how key lifetimes affect security, see **Filter action** in Related Topics.

- When negotiating security, the initiator proposes this list of security methods to the responder. The responder agrees to the first method (from the top down) that is found anywhere in its own security methods list. If none of the proposed security methods match an entry in the responder's list, then the responder fails the negotiation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Configure custom IPSec security methods
Define IPSec Key Exchange Settings
Filter action
Working with MMC console files

# Configure custom IPSec security methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure custom IPSec security methods**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify, and then click the **Filter Action** tab.
4. Double-click the filter action that you want to modify.
5. On the **Security Methods** tab, do one of the following:
   - To add a new security method, click **Add**.
   - To modify an existing security method, click the security method that you want to modify, and then click **Edit**.
   - To remove a security method, click the security method that you want to remove, and then click **Remove.**
6. If you are adding or modifying a security method, on the **Security Method** tab, click **Custom**, and then click **Settings**.
7. To provide integrity for both the packet's addressing information (IP header) and the data, select the **Data and address integrity without encryption (AH)** check box, and then, in **Integrity algorithm**, click one of the following:
   - **MD5** to use the Message Digest 5 (MD5) integrity algorithm, which uses a 128-bit key.
   - **SHA1** to use the Secure Hash Algorithm 1 (SHA1) integrity algorithm, which uses a 160-bit key.
8. To provide both integrity and encryption (confidentiality) for the data, select the **Data integrity and encryption (ESP)** check box, and then, in **Integrity algorithm**, click one of the following:
   - **<None>** to use no data integrity.

     If you have enabled **AH**, you can click **None** for the ESP integrity algorithm to improve performance.
   - **MD5** to use the MD5 integrity algorithm, which uses a 128-bit key.
   - **SHA1** to use the SHA1 integrity algorithm, which uses a 160-bit key.
9. If you selected the **Data integrity and encryption (ESP)** check box, in **Encryption algorithm**, click one of the following:
   - **<None>** to use no encryption.
   - **DES** to use Data Encryption Standard (DES) with a 56-bit key.
   - **3DES** to use the triple DES (3DES) with three 56-bit keys.

10. To change the default session key lifetime settings, do one of the following:

    - To specify the number of kilobytes of data transferred before a new session key is generated, select the **Generate a new key every** check box, and then type a value in **Kbytes**.

    - To specify the number of seconds that elapse before a new session key is generated, select the **Generate a new key every** check box, and then type a value in **seconds**.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- You cannot choose **<None>** for both the ESP integrity and the encryption algorithm.

- Use caution when specifying lifetimes. The default internal lifetimes should work in most environments and be interoperable with other products. If your security requirements demand shorter lifetime values, you can explicitly set appropriate lifetime values in each security method. The lowest values that are explicitly set on initiator and responder will be used.

- Using shorter lifetime values does not increase the security with which data is protected. Short key lifetimes will only decrease the amount of data that is revealed if an attacker discovers one encryption key. Instead of adjusting the lifetime values lower, consider using a stronger encryption algorithm to protect data, such as 3DES.

- Before changing the default lifetimes, you should understand how lifetimes and rekeying work in detail. For more information, see **Filter action** in Related Topics.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec security methods
Filter action
Working with MMC console files

# Select a filter action for a rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select a filter action for a rule**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Double-click the rule that you want to modify.
4. On the **Filter Action** tab, select the option that precedes the name of the filter action that you want to designate as the active filter action for this rule.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- If the option preceding the name of a filter action is not selected, the filter action will not be set as active for the rule.
- If you have not removed them, you can activate one of the three predefined filter actions: **Permit**, **Request Security (Optional)**, or **Require Security**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove filter actions
Add, edit, or remove IPSec security methods
Working with MMC console files

# Manage multiple filter actions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage multiple filter actions**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, and then click **Manage IP filter lists and filter actions**.
4. On the **Manage Filter Actions** tab, do one of the following:
   - To add a new filter action, click **Add**.
   - To modify an existing filter action, click the filter action that you want to modify, and then click **Edit**.
   - To remove a filter action, click the filter action that you want to remove, and then click **Remove**.
5. If you are adding or modifying a filter action, define settings on the **Security Methods** and **General** tabs as needed.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- For more information about defining settings for filter actions, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove filter actions
Add, edit, or remove IPSec security methods
Configure custom IPSec security methods
Filter action
Working with MMC console files

# Define IPSec Key Exchange Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define IPSec key exchange settings**

- Configure key exchange settings
- Create key exchange security methods

# Configure key exchange settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure key exchange settings**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Click the **General** tab, and then click **Settings**.
4. To force reauthentication and the negotiation of new master key keying material each time a new session key is required, click **Master key perfect forward secrecy (PFS)**.
5. If you require a different setting, type a value in **Authenticate and generate a new key after every number minutes**, which will cause a reauthentication and new master key generation at that interval.
6. If you require a different setting, type a value in **Authenticate and generate a new key after every number sessions** to set a maximum limit on the number of times a master key or its base keying material can be reused to generate the session key. This limit, when reached, will force a reauthentication and new master key generation. If you have enabled **Master key perfect forward secrecy (PFS)**, the number of sessions is set to 1 and is not configurable.
7. If you have special requirements for security methods on the master key exchange, click **Methods**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- Portions of IPSec-related services were jointly developed by Microsoft and Cisco Systems, Inc.
- Use **Master key perfect forward secrecy (PFS)** where it is required for interoperability. The default setting is disabled, which should work in most environments and be interoperable with most other products.
- A session limit of zero (0) will cause rekeys to be determined only by the time setting.
- Enabling master key perfect forward secrecy (PFS) might impact performance because each quick mode requires a new main mode negotiation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in

Open MMC
Create key exchange security methods
Key exchange settings
Working with MMC console files

# Create key exchange security methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create key exchange security methods**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Double-click the policy that you want to modify.
3. Click the **General** tab, click **Settings**, and then click **Methods**.
4. In the **Key Exchange Security Methods** dialog box, do one of the following:
   - To add a new key exchange security method, click **Add**.
   - To modify an existing key exchange security method, click the security method that you want to modify, and then click **Edit**.
   - To remove a key exchange security method, click the security method that you want to remove, and then click **Remove**.
5. If you are adding or modifying a key exchange security method, in the **IKE Security Algorithms** dialog box, select an **Integrity algorithm**:
   - Click **MD5** to use a 128-bit key (faster).
   - Click **SHA1** to use a 160-bit key (stronger).
6. Select an **Encryption algorithm**:
   - Click **3DES** to use the triple Data Encryption Standard (3DES) algorithm and three unique 56-bit keys.
   - Click **DES** to use the DES algorithm and a single 56-bit key. Use this option if you are required to connect to computers that do not have 3DES or if you do not need the higher security and overhead of 3DES. For more information, see Notes.
7. Select a **Diffie-Hellman group** to set the length of base keying material used to generate the actual keys:
   - Click **Low (1)** to generate 768 bits of master key keying material.
   - Click **Medium (2)** to generate 1,024 bits of master key keying material (stronger).
   - Click **High (2048)** to generate 2,048 bits of master key keying material (strongest).

**Important**

- For enhanced security, do not use Diffie-Hellman Group 1 (low). For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP . When you use a stronger group, the secret key that is derived from the Diffie-Hellman exchange has greater strength. For more

information about Diffie-Hellman groups and key exchange methods, see Related Topics.

### Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

- Portions of IPSec-related services were jointly developed by Microsoft and Cisco Systems, Inc.

- A key exchange security method is a combination of three settings (integrity algorithm, encryption algorithm, and Diffie-Hellman group). The initiator and the responder must have a method in common (one that uses the same settings) for negotiations to succeed.

- Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

- Click **Move up** to move a selected security method up one level. Repeat until the security method is at the required preference level.

- Click **Move down** to move a selected security method down one level. Repeat until the security method is at the required preference level.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Configure key exchange settings
Key exchange methods
Working with MMC console files
Key management and protection

# Assign IPSec Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Assign IPSec policy**

- Assign or unassign IPSec policy in Group Policy
- Assign or unassign IPSec policy on a computer

# Assign or unassign IPSec policy in Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign or unassign IPSec policy in Group Policy**

- For Active Directory-based Group Policy:

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

   **Where?**

   - Active Directory Users and Computers [*DomainControllerName.DomainName*]/*Domain/OrganizationalUnit/ChildOrganizationalUnit...*

3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the Group Policy object that you want to edit. Or, click **New** to create a new Group Policy object, and then click **Edit**.
5. In the Group Policy console tree, click **IP Security Policies on Active Directory**.

   **Where?**

   - *PolicyName* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/IP Security Policies on Active Directory

6. In the details pane, click the IPSec policy that you want to assign or unassign, and then do one of the following:

   a. To assign the policy, click the **Action** menu, and then click **Assign**.

   b. To unassign the policy, click the **Action** menu, and then click **Unassign**.

- For local computer policy:

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Group Policy Object Editor**, and then click **Add**.
4. Click **Finish**, click **Close**, and then click **OK**.
5. In the Group Policy console tree, click **IP Security Policies on Local Computer**.

   **Where?**

   - Local Computer Policy/Computer Configuration/Windows Settings/Security Settings/IP Security Policies on Local

Computer

6. In the details pane, click the IPSec policy that you want to assign or unassign, and then do one of the following:

   a. To assign the policy, click the **Action** menu, and then click **Assign**.

   b. To unassign the policy, click the **Action** menu, and then click **Un-assign**.

**Important**

- An IPSec policy might remain active even after the IPSec policy or Group Policy object to which it is assigned has been deleted. Therefore, you should unassign the IPSec policy before you delete either the policy or the Group Policy object. To prevent problems, use the following procedure:

1. Unassign the IPSec policy in the Group Policy object.

2. Wait 24 hours to ensure that the change is propagated.

3. Delete the IPSec policy or Group Policy object.

   If you delete the IPSec policy or Group Policy object without following this procedure, computers in the Active Directory container to which the IPSec policy is assigned might treat the IPSec policy as if it cannot be located and continue to use a cached copy.

**Notes**

- To manage Active Directory-based IPSec policy, you must be a member of the Domain Admins group in Active Directory, or you must have permission to edit Group Policy objects (for information about Group Policy object editing, see Related Topics). To manage local or remote IPSec policy for a computer, you must be a member of the Administrators group on the local or remote computer. For more information, see Default local groups and Default groups.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- You cannot administer Active Directory-based IPSec policy from a computer running Windows XP Home Edition.

- The settings will take effect the next time Group Policy is refreshed.

- If a policy is currently assigned and you assign a new policy, the currently assigned policy is automatically unassigned.

- You cannot assign a policy from the IP Security Policies on Active Directory console; you can only configure policy. To assign policy for Active Directory containers, you must use **IP Security Policies on Active Directory** within the Group Policy console.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add, edit, or remove IPSec policies
Creating, modifying, and assigning IPSec policies
Working with MMC console files
Group Policy (pre-GPMC)

# Assign or unassign IPSec policy on a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign or unassign IPSec policy on a computer**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. In the details pane, click the IPSec policy that you want to assign or unassign, and then do one of the following:
   - To assign the policy, click the **Action** menu, and then click **Assign**.
   - To unassign the policy, click the **Action** menu, and then click **Un-assign**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- If a policy is currently assigned and you assign a new policy, the currently assigned policy is automatically unassigned.
- If the computer is a member of a Windows 2000 or Windows Server 2003 domain and IPSec policy is assigned at the domain level, then the assigned local computer policy is not used.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Assign or unassign IPSec policy in Group Policy
Add, edit, or remove IPSec policies
Working with MMC console files

# Manage IPSec Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage IPSec policy**

- Restart the IPSec services
- Restore default IPSec policies
- Export IPSec policies
- Import IPSec policies
- Refresh the list of IPSec policies
- Delete an IPSec policy
- Rename an IPSec policy
- Export the list of IPSec policies

# Restart the IPSec services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To restart the IPSec service

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Services.
2. Right-click **IPSec services**.
3. Click **Restart**.

### Important

- Stopping and restarting the IPSec service can disconnect all computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. For information about how to troubleshoot and resolve loss of connectivity in this situation, see Related Topics.

### Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.
- You might need to restart the IPSec service in order to do any of the following:
    - Clear old security associations (SAs).
    - Clear soft SAs (unsecured negotiations).
    - Force an IPSec policy to download from Active Directory to domain clients.
- If the IPSec service does not start, then use Event Viewer to determine possible causes of failure. For more information, see Related Topics.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **net stop policyagent & net start policyagent**

**Important**

- Stopping and restarting the IPSec service can disconnect all computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. For information about how to troubleshoot and resolve loss of connectivity in this situation, see Related Topics.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.

- On a computer running the Routing and Remote Access service and either Windows 2000 Server or a Windows Server 2003 operating system, you must use the following sequence when restarting the IPSec service:

  1. Stop the Routing and Remote Access service using the **net stop remoteaccess** command.

  2. Stop and restart the IPSec service using the **net stop policyagent & net start policyagent** command.

  3. Start the Routing and Remote Access service using the **net start remoteaccess** command.

- You might need to restart the IPSec service in order to do any of the following:

  - Clear old security associations (SAs).

  - Clear soft SAs (unsecured negotiations).

  - Force an IPSec policy to download from Active Directory to domain clients.

- If the IPSec service does not start, then use Event Viewer to determine possible causes of failure. For more information, see Related Topics.

- For more information about the **net start** and **net stop** commands, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
IPSec troubleshooting tools
IPSec Policy Agent service
Event Viewer overview
Net services commands

# Restore default IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore default IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, point to **All Tasks**, and then click **Restore Default Policies**.
4. When prompted, click **Yes**.

### Important

- Restoring default IPSec policies will overwrite any changes to the original default policies, even when the names of the default policies have been changed. If you have modified the default policies and you do not want those modifications to be overwritten, do not restore the default policies.

### Notes

To perform this procedure, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- The option to restore default IPSec policies is only available for computer-based IPSec policy. You cannot restore the default IPSec policies for Active Directory-based IPSec policy.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Assign or unassign IPSec policy in Group Policy
Assign or unassign IPSec policy on a computer
Working with MMC console files

# Export IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, point to **All Tasks**, and then click **Export Policies**.
4. Type the path and the name of the file to which you want the policy information saved, and then click **Save**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- To enhance interoperablity in a mixed environment with computers running Windows 2000, limit the name of the file to which you want the policy information saved to 60 characters.
- To export IPSec policies, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Import IPSec policies
Export the list of IPSec policies
Working with MMC console files

# Import IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, point to **All Tasks**, and then click **Import Policies**.
4. Type the path and the name of the file from which you are importing the IPSec policy information, and then click **Open**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- Importing IPSec policies replaces existing IPSec policies with the same name.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Export IPSec policies
Export the list of IPSec policies
Working with MMC console files

# Refresh the list of IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To refresh the list of IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, and then click **Refresh**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- This procedure is only required to update the list of policies for any additions or deletions made by using another management tool.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Working with MMC console files

# Delete an IPSec policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an IPSec policy**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Click the policy that you want to remove, click **Action**, and then click **Delete**.
3. When prompted, click **Yes**.

### Important

- For Active Directory-based IPSec policy, an IPSec policy might remain active even after the IPSec policy has been deleted. Therefore, you should unassign the IPSec policy before you delete the policy. To prevent problems, use the following procedure:

  1. Unassign the IPSec policy in the Group Policy object.
  2. Wait 24 hours to ensure that the change is propagated.
  3. Delete the IPSec policy.

  If you delete the IPSec policy without following this procedure, computers in the Active Directory container to which the IPSec policy is assigned might treat the IPSec policy as if it cannot be located and continue to use a cached copy.

### Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Refresh the list of IPSec policies
Working with MMC console files

# Rename an IPSec policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename an IPSec policy**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Click the policy that you want to rename, click **Action**, and then click **Rename**.
3. Type a new name, and then press **Enter**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Add, edit, or remove IPSec policies
Refresh the list of IPSec policies
Working with MMC console files

# Export the list of IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export the list of IPSec policies**

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, click **IP Security Policies on *Name***.
3. Click **Action**, and then click **Export List**.
4. Type the file name to which you want to save the selected IPSec policy or the displayed list of IPSec policies, and then click **Save**.

## Notes

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.
- To create a console containing IP Security Policies, start the IP Security Policies snap-in. To open a saved console file, open MMC. For more information, see Related Topics.
- This procedure does not export the IPSec policies, only a text list of the policies displayed in the details pane.
- To export IPSec policies, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the IP Security Policy Management snap-in
Open MMC
Import IPSec policies
Export IPSec policies
Working with MMC console files

# Use Resultant Set of Policy (RSoP) to View IPSec Policy Assignments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Resultant Set of Policy (RSoP) to view IPSec policy assignments**

- Use RSoP to view IPSec policy assignments for a computer
- Use RSoP to view IPSec policy assignments for members of a Group Policy container

# Use RSoP to view IPSec policy assignments for a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use RSoP to view IPSec policy assignments for a computer**

1. Open Resultant Set of Policy.

   An RSoP query runs in logging mode. For more information, see Related Topics.

2. In the console tree, click **IP Security Policies**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/IP Security Policies

   The list of all IPSec policies is displayed in the details pane.

3. To view detailed policy settings for the IPSec policy that is being applied, in the details pane, right-click the policy with the highest precedence, and then click **Properties**.

## Notes

- To open RSoP, click **Start**, click **Run**, type **rsop.msc**, and then click **OK**.
- To run an RSoP query in logging mode, see Related Topics.
- The RSoP console displays detailed IPSec policy settings for only the policy that is being applied. The detailed policy settings include the filter rules, filter actions, authentication methods, tunnel endpoints, and connection type.
- The details pane in the RSoP console provides two ways to indicate which IPSec policy is being applied:
  - The lower the number in the **Precedence** column, the higher the precedence of the IPSec policy. The IPSec policy with a precedence number of 1 is the policy that is being applied.
  - In the **Name** column, the folder icon that precedes the name of the IPSec policy that is being applied contains a green dot.
- The settings of the IPSec policy with the highest precedence apply in their entirety; they are not merged with the settings of IPSec policies that are assigned at higher levels of the Active Directory hierarchy. For more information about using RSoP to view IPSec policy assignments, see Related Topics.
- If a local user overrides domain-level IPSec policies with a local IPSec policy, RSoP does not detect the application of the local IPSec policy. As a result, no IPSec policies will be displayed in the RSoP console when a logging mode query is run.
- To hide or display columns in the RSoP console, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Run RSoP

Access RSoP data for an existing computer and user (logging mode)

RSoP logging mode

Hide or display columns in an MMC console

Using Resultant Set of Policy to view IPSec policy assignments

# Use RSoP to view IPSec policy assignments for members of a Group Policy container

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use RSoP to view IPSec policy assignments for members of a Group Policy container**

1. On a domain controller, open Resultant Set of Policy (RSoP) as a Microsoft Management Console (MMC). For more information, see Related Topics.

2. On the **Action** menu, click **Generate RSoP Data**.

3. On the **Welcome to the Resultant Set of Policy Wizard** page, click **Next**.

4. On the **Mode Selection** page, click **Planning mode**, and then follow the instructions in the wizard to run a planning mode query. For more information, see Related Topics.

5. In the console tree, click **IP Security Policies**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/IP Security Policies

   The list of all IPSec policies is displayed in the details pane.

6. To view detailed policy settings for the IPSec policy that is being applied, in the details pane, right-click the policy with the highest precedence, and then click **Properties**.

**Notes**

- The RSoP console displays detailed IPSec policy settings for only the policy that is being applied. The detailed policy settings include the filter rules, filter actions, authentication methods, tunnel endpoints, and connection type.

- The details pane in the RSoP console provides two ways to indicate which IPSec policy is being applied:

  - The lower the number in the **Precedence** column, the higher the precedence of the IPSec policy. The IPSec policy with a precedence number of 1 is the policy that is being applied.

  - In the **Name** column, the folder icon that precedes the name of the IPSec policy that is being applied contains a green dot.

- The settings of the IPSec policy with the highest precedence apply in their entirety; they are not merged with the settings of IPSec policies that are assigned at higher levels of the Active Directory hierarchy. For more information about using RSoP to view IPSec policy assignments, see Related Topics.

- You can run an RSoP planning mode query only on a domain controller (when you run a planning mode query, you must explicitly specify the domain controller name). However, you can specify any IPSec client as the target for the query, provided you have the appropriate permissions to do so.

- To hide or display columns in the RSoP console, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open RSoP as an MMC snap-in
Simulate RSoP data for a computer and user (planning mode)
Hide or display columns in an MMC console
Using Resultant Set of Policy to view IPSec policy assignments

# Monitor IPSec Activity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor IPSec activity**

- Add the IP Security Monitor snap-in
- Add a computer to the IP Security Monitor console
- View details about active IPSec policies
- View IP security statistics
- View main mode details
- View quick mode details
- Find a matching specific filter
- Change the automatic update refresh interval
- Enable DNS name resolution

# Add the IP Security Monitor snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add the IP Security Monitor snap-in**

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **IP Security Monitor**, and then click **Add**.
4. Click **Close**, and then click **OK**.

## Notes

- To save the console settings, on the **File** menu, click **Save**.
- To view the saved console, see Related Topics.
- To revise your initial choice, you can start the Microsoft Management Console, add this snap-in again, and save the console again. If you require multiple configurations, you can save the console with another name.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor

# Add a computer to the IP Security Monitor console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a computer to the IP Security Monitor console**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. In the console tree, right-click **IP Security Monitor**, and then click **Add Computer**.
3. In the **Add Computer** dialog box:
   - For the local computer, click **This computer**.
   - For a remote computer, click **The following computer**, and then type the name of the remote computer. Or, click **Browse** to find it on the network.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.
- If the IPSec service is not started on the computer that is being monitored, the server icon is displayed as a stopped service. To refresh IP Security Monitor after the IPSec service on that computer have been restarted, right-click the computer and then click **Reconnect**.
- IP Security Monitor can be used to monitor IPSec only on computers running Windows XP or a Windows Server 2003 operating system. To monitor IPSec on a computer running Windows 2000, use the **IPSecmon** command at a Windows 2000 command prompt.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# View details about active IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view details about active IPSec policies**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.

2. Double-click **Active Policy**.

   **Where?**

   - IP Security Monitor/*ServerName*/Active Policy

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- For information about adding the IP Security Monitor snap-in, see Related Topics.

- For information about the items displayed in **Active Policy**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# View IP security statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view IP security statistics**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Click a server name.

   **Where?**

   - IP Security Monitor/*ServerName*

3. Do one of the following:

   - To view main mode statistics, in the console tree, double-click **Main Mode**, and then double-click **Statistics**.
   - To view quick mode statistics, in the console tree, double-click **Quick Mode**, and then double-click **Statistics**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.
- For more information about the items displayed in **Statistics**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in
Viewing main mode and quick mode statistics in IP Security Monitor

# View main mode details

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view main mode details**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Double-click **Main Mode**.

   **Where?**

   - IP Security Monitor/*ServerName*/Main Mode

3. To view main mode details, do one of the following:

   - For generic filter details, in the console tree, click **Generic Filters**, and then in the details pane, double-click a filter.
   - For specific filter details, in the console tree, click **Specific Filters**. To find a specific main mode filter, right-click **Specific Filters**, and then click **Find matching filters**.
   - For IKE policy details, in the console tree, click **IKE Policies**, and then in the details pane, double-click a policy.
   - For statistical details, in the console tree, click **Statistics**.
   - For security association details, in the console tree, click **Security Associations**, and then in the details pane, double-click a security association.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# View quick mode details

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view quick mode details**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Double-click **Quick Mode**.

   **Where?**

   - IP Security Monitor/*ServerName*/Quick Mode

3. To view quick mode details, do one of the following:

   - For generic filter details, in the console tree, click **Generic Filters**, and then in the details pane, double-click a filter.
   - For specific filter details, in the console tree, click **Specific Filters**. To find a specific quick mode filter, right-click **Specific Filters**, and then click **Find matching filters**.
   - For negotiation policy details, in the console tree, click **Negotiation Policies**, and then in the details pane, double-click a policy.
   - For statistical details, in the console tree, click **Statistics**.
   - For security association details, in the console tree, click **Security Associations**, and then in the details pane, double-click a security association.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.
- Quick mode details do not display **Block** or **Permit** filters.
- Although there are two separate quick mode SAs established (one for the inbound traffic and one for the outbound traffic), IP Security Monitor only displays a single quick mode SA.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file

Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# Find a matching specific filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a matching specific filter**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. For a specific main mode filter:
   a. Double-click **Main Mode**.(console tree location is: IP Security Monitor/*ServerName*/Main Mode).
   b. In the console tree, right-click **Specific Filters**, and then click **Find matching filters**.
   c. In the **Find Matching Main Mode Filters** dialog box, configure the filter settings, and then click **Search**.
3. For a specific quick mode filter:
   a. Double-click **Quick Mode** (console tree location is: IP Security Monitor/*ServerName*/Quick Mode).
   b. In the console tree, right-click **Specific Filters**, and then click **Find matching filters**.
   c. In the **Find Matching Quick Mode Filters** dialog box, configure the filter settings, and then click **Search**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# Change the automatic update refresh interval

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the automatic update refresh interval**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Right-click a server name.

   **Where?**

   - IP Security Monitor/*ServerName*

3. Click **Properties**.
4. Configure the automatic update refresh interval, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# Enable DNS name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable DNS name resolution**

1. Create a console containing IP Security Monitor. Or, open a saved console file containing IP Security Monitor.
2. Right-click a server name.

   **Where?**

   - IP Security Monitor/*ServerName*

3. Click **Properties**.
4. Select the **Enable DNS name resolution** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For information about adding the IP Security Monitor snap-in, see Related Topics.
- The **Enable DNS name resolution** check box specifies whether you want to use Domain Name System (DNS) reverse queries and NetBIOS Adapter Status messages to resolve the DNS or NetBIOS computer name that is associated with an IP address when displaying sources or destinations. To resolve the DNS computer name from its IP address, your DNS infrastructure must have the appropriate reverse domains and pointer records configured, either manually or with DNS dynamic update. To resolve the NetBIOS computer name from its IP address, NetBIOS over TCP/IP must be enabled on the computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor IPSec Activity
Open a saved MMC console for local and remote computers
Save an MMC console file
Viewing main mode and quick mode statistics in IP Security Monitor
Viewing details about active IPSec policies in IP Security Monitor
Add the IP Security Monitor snap-in

# IPSec Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Internet Protocol security (IPSec):

- New features for IPSec
- Internet Protocol Security Overview
- Understanding Internet Protocol Security
- Using Internet Protocol Security
- IPSec Resources

# New features for IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features for IPSec**
Internet Protocol security (IPSec) provides the following new features for enhanced security, scalability, and availability, and ease of deployment and administration.

### Note

- Because several new IPSec features are available only in the Windows Server 2003 family, as a best practice, if you plan to apply the same IPSec policy to computers running the Windows Server 2003 family and to computers running Windows 2000 or Windows XP, test the policy thoroughly on all relevant operating systems before deployment. For additional best practices for using IPSec, see IPSec Best practices.

**IP Security Monitor**
In Windows 2000, IP Security Monitor was implemented as an executable program (IPSecmon.exe). In Windows XP and the Windows Server 2003 family, IP Security Monitor is implemented as a Microsoft Management Console (MMC) console and includes enhancements that allow you to:

- Monitor IPSec information for your local computer and for remote computers.
- View details about active IPSec policies, including the name, description, date last modified, store, path, organizational unit, and Group Policy object name.
- View main mode and quick mode generic filters and specific filters.
- View main mode and quick mode statistics. For information about the statistics displayed in IP Security Monitor, see Viewing main mode and quick mode statistics in IP Security Monitor.
- View main mode and quick mode security associations.
- Customize refresh rates, and use DNS name resolution for filter and security association output.
- Search for specific main mode or quick mode filters that match any source or destination IP address, a source or destination IP address on your local computer, or a specific source or destination IP address.

For information about adding the IP Security Monitor snap-in, see Add the IP Security Monitor snap-in. For information about additional IPSec troubleshooting tools, see IPSec troubleshooting tools.

**Stronger cryptographic master key (Diffie-Hellman)**
For enhanced security, IPSec now supports the use of a 2048-bit Diffie-Hellman key exchange. With a stronger Diffie-Hellman group, the secret key that is derived from the Diffie-Hellman exchange has greater strength. Strong Diffie-Hellman groups combined with longer key lengths increase the computational difficulty of determining a secret key.

This feature is provided only with the Windows Server 2003 family.

For more information about Diffie-Hellman groups and key exchange, see Key exchange methods and Key management and protection.

**Command-line management with Netsh**
Using commands in the Netsh IPSec context, you can configure static or dynamic IPSec main mode settings, quick mode settings, rules, and configuration parameters. To enter the Netsh IPSec context, type **netsh -c ipsec** at a command prompt. The Netsh IPSec context replaces the Ipsecpol.exe tool, which is provided with the Windows 2000 Server Resource Kit. You can use this feature to script and automate IPSec configuration. For more information, see IPSec troubleshooting tools.

This feature is provided only with the Windows Server 2003 family.

For more information about Netsh commands for IPSec, see Netsh commands for Internet Protocol security.

**Computer startup security**
For enhanced security, IPSec now provides stateful filtering of network traffic during computer startup. With stateful filtering, only the following traffic is permitted during computer startup: the outbound traffic that the computer initiates during startup, the inbound traffic that is sent in response to the outbound traffic, and DHCP traffic. As an alternative to stateful filtering, you can specify that all inbound and outbound traffic be blocked until an IPSec policy is applied. If you use stateful filtering, or if you specify that traffic be blocked during computer startup, you can also specify the traffic types that you want to exempt from IPSec filtering during computer startup.

### Note

- You cannot configure this feature in the IP Security Policy Management console. To configure this feature, you must use the Netsh IPSec command-line tool.

**Persistent policy for enhanced security**
You can now create and assign a persistent IPSec policy to secure a computer if a local IPSec policy or an Active Directory-based IPSec policy cannot be applied. When you create and assign a persistent policy, it is applied before the local policy or the Active Directory-based policy is applied, and it remains in effect regardless of whether the local policy or the Active Directory-based policy is applied (for example, an IPSec policy will not be applied if it is corrupted).

### Note

- You cannot configure this feature in the IP Security Policy Management console. To configure this feature, you must use the Netsh commands for IPSec.

**Removed default traffic exemptions**
In Windows 2000 and Windows XP, by default, all broadcast, multicast, Internet Key Exchange (IKE), Kerberos, and Resource Reservation Protocol (RSVP) traffic is exempt from IPSec filtering. To significantly improve security, in the Windows Server 2003 family, only IKE traffic (which is required for establishing IPSec-secured communication) is exempt from IPSec filtering. All other traffic types are now matched against IPSec filters, and you can configure, block, or permit filter actions specifically for multicast and broadcast traffic (IPSec does not negotiate security associations for multicast and broadcast traffic). For more information, see Special IPSec considerations.

**IPSec certificate to account mapping for network access control**
With the Windows Server 2003 family, if you use either Kerberos V5 or certificate authentication, you can set restrictions on which computers are allowed to connect. This functionality allows you to use IPSec to allow or deny any of the following access to a server running Windows Server 2003 :

- Computers that are members of a specific domain.
- Computers that have a certificate from a specific issuing certification authority.
- A specific group of computers.
- A specific computer.

When you enable certificate to account mapping in IPSec, the IKE protocol associates (maps) a computer certificate to a computer account in an Active Directory domain or forest, and then retrieves an access token, which includes the list of the user rights that are assigned to the computer. You can restrict access by configuring Group Policy security settings and assigning either the **Access this computer from the network** user right or the **Deny access to this computer from the network** user right to individual or multiple computers, as needed.

For more information about certificate to account mapping for IPSec, see Authentication methods.

For information about configuring security settings in Group Policy, see Security settings overview.

### Notes

- Certificate to account mapping for IPSec is provided only with the Windows Server 2003 family.

- Certificate mapping is not supported across forests.

**Ability to exclude the name of the certification authority (CA) from certificate requests**
For enhanced security, when you use certificate authentication to establish trust between IPSec peers, you can now exclude the name of the CA from the certificate request. When you exclude the name of the CA from the certificate request, you prevent the potential disclosure of sensitive information about the trust relationships of a computer, such as name of the company that owns the computer and the domain membership of the computer (if an internal public key infrastructure is being used), to an attacker.

**IPSec policy filters allow logical addresses for local IP configuration**
You can now use the IP Security Policy Management console to configure the source or the destination address fields that the local IPSec policy will interpret as the addresses for the DHCP server, the DNS servers, the WINS servers, and the default gateway. As a result, IPSec policies can now automatically accommodate changes in the IP configuration of the server, by using either DHCP or static IP configurations.

This feature is provided only with the Windows Server 2003 family.

**IPSec functionality over network address translation (NAT)**
IPSec Encapsulating Security Payload (ESP) packets can now pass through NATs that allow User Datagram Protocol (UDP) traffic. The IKE protocol automatically detects the presence of a NAT and uses UDP-ESP encapsulation to allow IPSec traffic to pass through the NAT. This functionality is an implementation of the Internet Engineering Task Force (IETF) IP Security Working Group standard for IPSec.

NATs are widely used for Internet Connection Sharing (ICS) and in locations that provide public Internet access (such as hotels and airports) and that are likely to be used by telecommuters. In addition, some Internet service providers (ISPs) use a centralized NAT to connect their clients to the Internet.

IPSec functionality over NAT enables IPSec-secured connections to be established in the following common deployment scenarios:

- Layer Two Tunneling Protocol (L2TP)/IPSec virtual private network (VPN) clients that are behind NATs can establish IPSec-secured connections over the Internet to their corporate network, using IPSec ESP transport mode.

- Servers running Routing and Remote Access can establish gateway-to-gateway IPSec tunnels when one of the servers running Routing and Remote Access is behind a NAT.

- Clients and servers can send IPSec-secured TCP and UDP packets to other clients or servers using IPSec ESP transport mode, when one or both of the computers are behind a NAT. For example, a program running on a server on a perimeter network can be IPSec-secured when it is used to make connections to the corporate network.

### Note

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

For more information about VPNs and IPSec, see Virtual private networking with IPSec. For more information about IPSec transport mode and tunnel mode, see Transport mode and Tunnel mode.

**Improved IPSec integration with Network Load Balancing**
Improved IPSec integration with Network Load Balancing allows a Network Load Balancing group of servers to provide highly available IPSec-based VPN services. Network Load Balancing can accurately track IPSec-secured sessions, and the IPSec IKE protocol can detect when an IPSec-secured session is being established with a cluster server and quickly recover from a failover. Additionally, Network Load Balancing can now maintain IPSec-secured connections to the correct Network Load Balancing host, even when the number of hosts in the cluster (and the algorithm used to map clients to hosts) changes. Because the IKE protocol automatically detects the Network Load Balancing service, no additional configuration is required to use this feature.

This feature is provided only with Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition.

For more information about Network Load Balancing, see Introduction to Network Load Balancing and How Network Load

Balancing works.

**IPSec support for RSoP**

To enhance IPSec deployment and troubleshooting, IPSec now provides an extension to the Resultant Set of Policy (RSoP) console. RSoP is an addition to Group Policy that you can use to view existing IPSec policy assignments for a computer or for members of a Group Policy container. To view IPSec policy assignments for a computer, run an RSoP logging mode query. To view IPSec policy assignments for members of a Group Policy container, run an RSoP planning mode query. For information about how to view IPSec policy assignments, see Use Resultant Set of Policy (RSoP) to View IPSec Policy Assignments.

After you run an RSoP logging mode query or an RSoP planning mode query, you can view detailed settings (the filter rules, filter actions, authentication methods, tunnel endpoints, and connection types that were specified when the IPSec policy was created) for the IPSec policy that is being applied.

For more information about using RSoP with IPSec, see Using Resultant Set of Policy to view IPSec policy assignments. For general information about RSoP, see RSoP overview.

# Internet Protocol Security Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol security overview**

This section covers:

- Security information for IPSec
- Internet Protocol security defined
- Introducing IPSec
- Security issues with IP
- Protection against attacks
- Policy-based security
- Simplified deployment

# Security information for IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for IPSec**
Before deploying IPSec in your organization, consider the following security issues:

- 3DES and computers running Microsoft® Windows® 2000
- Authentication method
- Firewall packet filtering
- Protected traffic (using AH or ESP)
- Unprotected traffic (using the default response rule)
- Diffie-Hellman groups (using Diffie-Hellman Group 1, 2, or 2048)

**3DES and computers running Windows 2000**
IPSec policies allow the choice of a strong encryption algorithm, 3DES, which provides stronger encryption than DES for higher security. Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to perform the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, then the 3DES setting is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should use DES only as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP and the Windows Server 2003 family support 3DES and do not require installation of the High Encryption Pack.

**Authentication method**
If the computers in your organization are part of an Active Directory® domain, IPSec main mode authentication can be accomplished with the default authentication method (Kerberos V5). You do not have to deploy public key certificates for intranet communication. However, computers running Windows XP Home Edition do not support the Kerberos V5 authentication method. Also, if you have computers connected to the Internet, it is recommended that you do not use Kerberos V5 as an authentication method. When Kerberos authentication is used, during main mode negotiation, each IPSec peer sends its computer identity in unencrypted format to the other peer. The computer identity is unencrypted until encryption of the entire identity payload takes place during the authentication phase of the main mode negotiation. An attacker can send an Internet Key Exchange (IKE) packet that causes the responding IPSec peer to expose its computer identity and domain membership. To secure computers that are connected to the Internet, certificate authentication is therefore recommended.

For enhanced security, the use of preshared key authentication is not recommended because it is a relatively weak authentication method. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing.

For more information, see Authentication methods and Preshared key authentication.

**Firewall packet filtering**
For a firewall, security gateway, router, or any other server or device that is connected to the Internet and providing packet filtering capabilities to computers on a perimeter network (also known as a demilitarized zone or DMZ), special filtering must be enabled on that computer to ensure that packets secured with IPSec are allowed to be forwarded to computers on the perimeter network. The firewall or other device should normally allow the following types of traffic to pass through it:

- IP Protocol ID of 50 (0x32) for IPSec Encapsulating Security Payload (ESP) traffic
- IP Protocol ID of 51 (0x33) for IPSec Authentication Header (AH) traffic
- UDP port 500 (0x1F4) for Internet Key Exchange (IKE) negotiation traffic

Most packet filtering software allows you to be more specific. You can define separate packet filters for inbound traffic (inbound filters), for outbound traffic (outbound filters), and for each interface. Additionally, you can specify IP addresses for the IPSec computers on the perimeter network. The most restrictive packet filters for IPSec traffic exchanged with a single IPSec computer on the perimeter network are outlined in the following sections.

**Filters on the Internet interface**

Configure the following inbound packet filters on the Internet interface of the firewall to allow the specific types of traffic:

- Destination IP address of the IPSec computer's perimeter network interface and UDP destination port of 500 (0x1F4)

  This filter allows IKE traffic to be sent to the IPSec computer on the perimeter network.

- Destination IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 50 (0x32)

  This filter allows IPSec ESP traffic to be sent to the IPSec computer on the perimeter network.

- Destination IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 51 (0x33)

  This filter allows IPSec AH traffic to be sent to the IPSec computer on the perimeter network.

Configure the following outbound packet filters on the Internet interface of the firewall to allow the specific types of traffic:

- Source IP address of the IPSec computer's perimeter network interface and UDP source port of 500 (0x1F4)

  This filter allows IKE traffic to be sent from the IPSec computer on the perimeter network.

- Source IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 50 (0x32)

  This filter allows IPSec ESP traffic to be sent from the IPSec computer on the perimeter network.

- Source IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 51 (0x33)

  This filter allows IPSec AH traffic to be sent from the IPSec computer on the perimeter network.

**Filters on the perimeter network interface**

Configure the following inbound packet filters on the perimeter network interface of the firewall to allow the specific types of traffic:

- Source IP address of the IPSec computer's perimeter network interface and UDP source port of 500 (0x1F4)

  This filter allows IKE traffic to be sent from the IPSec computer on the perimeter network.

- Source IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 50 (0x32)

  This filter allows IPSec ESP traffic to be sent from the IPSec computer on the perimeter network.

- Source IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 51 (0x33)

  This filter allows IPSec AH traffic to be sent from the IPSec computer on the perimeter network.

Configure the following outbound packet filters on the perimeter network interface of the firewall to allow the specific types of traffic:

- Destination IP address of the IPSec computer's perimeter network interface and UDP destination port of 500 (0x1F4)

  This filter allows IKE traffic to be sent to the IPSec computer on the perimeter network.

- Destination IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 50 (0x32)

This filter allows IPSec ESP traffic to be sent to the IPSec computer on the perimeter network.

- Destination IP address of the IPSec computer's perimeter network interface and IP Protocol ID of 51 (0x33)

  This filter allows IPSec AH traffic to be sent to the IPSec computer on the perimeter network.

**Protected traffic (using AH or ESP)**
When deciding whether to use AH or ESP to protect your IP traffic, consider the following:

- AH provides both data authentication and integrity services through the computation and inclusion of a keyed hash for each packet. With AH, the hash calculation includes the entire IP packet and header. Some fields that are allowed to change in transit are excluded. Packet replay protection services are provided through the inclusion of a sequence number for each packet.

- ESP provides data authentication and integrity services through the computation and inclusion of a keyed hash for each packet. With ESP, the hash calculation only includes the ESP header, trailer, and payload. The IP header is not protected with the hash. ESP provides data confidentiality services by encrypting the ESP payload with the DES (Data Encryption Standard) or 3DES (triple DES) encryption algorithms. Packet replay protection services are provided through the inclusion of a sequence number for each packet.

It is less CPU-intensive to verify the hash for each packet than it is to encrypt and decrypt each packet. If performance is a significant consideration, you can use AH to protect most of your traffic. When confidentiality is required, you can use ESP instead. For example, you can use AH to protect traffic on your intranet and ESP for traffic that is sent over the Internet.

**Note**

- This performance consideration assumes that you are not using IPSec offload network adapters in your organization. Offload network adapters perform cryptographic calculations, such as the calculation and verification of the hash and the encryption and decryption services, on the adapter itself, enhancing performance for IPSec-secured traffic.

**Unprotected traffic (using the default response rule)**
You can easily implement an IPSec deployment scenario that secures traffic from a specific set of servers, as follows:

1. Place the servers in an organizational unit and assign an IPSec policy to a Group Policy object of that organizational unit with a rule that requires secured traffic between the servers and any other computer. Within the filter action for the rule, enable the secure servers to accept unsecured communication but always respond using IPSec.

2. Place the client computers in an organizational unit and assign an IPSec policy to a Group Policy object of that organizational unit that uses only the default response rule. For more information, see IPSec Policy Rules.

In this configuration, the following traffic between the client computers and the secure servers is unsecured:

- The initial request (for example, a TCP SYN packet) sent by the client computer to establish communication is unsecured because the client does not have a rule in its policy to initiate secured communication with the secure server.

- Although ongoing communications are secured after secured communication is negotiated between the client and secure server computers, the quick mode security association (SA) that provides security settings for traffic can eventually time out and be removed from both the secure server and the client computer. This occurs when a TCP connection exists and no traffic is sent between the client computer and the secure server for an extended period of time. The dynamic filter created on the client computer by the default response rule is also removed after the dynamic filter has been timed out.

  If the secure server sends new data on the existing connection, it renegotiates the quick mode SA before sending this data to the client computer because a rule exists on the secure server to secure traffic between itself and all other computers. However, if the client computer sends new data after the quick mode SAs and dynamic filter time out, the new data is sent unsecured because the client does not have a rule to initiate secure communications with the secure

server. For TCP connections, one or multiple TCP segments can be sent unsecured by the client computer. To prevent client computers from sending unsecured data to secure servers, you must configure your client computer IPSec policy with additional rules that initiate secured communications to secure servers.

**Diffie-Hellman groups (using Diffie-Hellman group 1, 2, or 2048)**

Diffie-Hellman groups are used to determine the length of the base prime numbers that are used during the key exchange process. Group 2048 (high) is stronger (more secure) than Group 2 (medium), which is stronger than Group 1 (low). Group 1 provides 768 bits of keying strength, Group 2 provides 1,024 bits, and Group 2048 provides 2,048 bits.

Strong Diffie-Hellman groups combined with longer key lengths increase the computational difficulty of determining a secret key.

**Important**

- For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.

**Note**

- Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.

For more information about Diffie-Hellman groups and key exchange, see Key exchange methods and Key management and protection.

# Internet Protocol security defined

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol security defined**

Internet Protocol security (IPSec) is a framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services. The Microsoft® Windows® 2000 , Windows XP, and the Windows Server 2003 family implementations of IPSec are based on standards developed by the Internet Engineering Task Force (IETF) IPSec working group.

IPSec is the long-term direction for secure networking. It provides aggressive protection against private network and Internet attacks through end-to-end security. The only computers that must know about IPSec protection are the sender and receiver in the communication. In Windows XP and the Windows Server 2003 family, IPSec provides the ability to protect communication between workgroups, local area network computers, domain clients and servers, branch offices (which might be physically remote), extranets, and roving clients.

# Introducing IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introducing IPSec**

IPSec is the long-term direction for secure networking. It provides a key line of defense against private network and Internet attacks.

IPSec has two goals:

1. To protect the contents of IP packets.

2. To provide a defense against network attacks through packet filtering and the enforcement of trusted communication.

Both goals are met through the use of cryptography-based protection services, security protocols, and dynamic key management. This foundation provides both the strength and flexibility to protect communications between private network computers, domains, sites, remote sites, extranets, and dial-up clients. It can even be used to block receipt or transmission of specific traffic types.

IPSec is based on an end-to-end security model, establishing trust and security from a source IP to a destination IP address. The IP address itself does not necessarily have to be considered an identity, rather the system behind the IP address has an identity that is validated through an authentication process. The only computers that must know about the traffic being secured are the sending and receiving computers. Each computer handles security at its respective end, with the assumption that the medium over which the communication takes place is not secure. Any computers that only route data from source to destination are not required to support IPSec, unless firewall-type packet filtering or network address translation is being done between the two computers. This model allows IPSec to be successfully deployed for the following enterprise scenarios:

- Local area network (LAN): client/server and peer-to-peer

- Wide area network (WAN): router-to-router and gateway-to-gateway

- Remote access: dial-up clients and Internet access from private networks

Typically both sides require IPSec configuration (called an IPSec policy), to set options and security settings that will allow two systems to agree on how to secure traffic between them. The Microsoft® Windows® 2000 , Windows XP, and the Windows Server 2003 family implementations of IPSec is based on industry standards developed by the Internet Engineering Task Force (IETF) IPSec working group. Portions of IPSec-related services were jointly developed by Microsoft and Cisco Systems, Inc.

# Security issues with IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security issues with IP**
Without security, both public and private networks are susceptible to unauthorized monitoring and access. Internal attacks might be a result of minimal or nonexistent intranet security. Risks from outside the private network originate from connections to the Internet and extranets. Password-based user access controls alone do not protect data transmitted across a network.

**Common types of network attacks**
Without security measures and controls in place, your data might be subjected to an attack. Some attacks are passive in that information is only monitored. Other attacks are active and information is altered with intent to corrupt or destroy the data or the network itself. Your networks and data are vulnerable to any of the following types of attacks if you do not have a security plan in place.

**Eavesdropping**
In general, the majority of network communications occur in a plaintext (unencrypted) format, which allows an attacker who has gained access to data paths in a network to monitor and interpret (read) the traffic. When an attacker is eavesdropping on communications, it is referred to as sniffing or snooping. The ability of an eavesdropper to monitor the network is generally the biggest security problem that administrators face in an enterprise. Without strong encryption services that are based on cryptography, data can be read by others as it traverses the network.

**Data modification**
After an attacker has read data, the next logical step is often to modify it. An attacker can modify the data in the packet without the knowledge of the sender or receiver. Even if you do not require confidentiality for all communications, you do not want any of your messages to be modified in transit. For example, if you are exchanging purchase requisitions, you do not want the items, amounts, or billing information to be changed.

**Identity spoofing (IP address spoofing)**
Most networks and operating systems use the IP address to identify a computer as being valid on a network. In some cases, it is possible for an IP address to be falsely used. This is known as identity spoofing. An attacker might use special programs to construct IP packets that appear to originate from valid addresses inside an organization intranet.

After gaining access to the network with a valid IP address, the attacker can modify, reroute, or delete data. The attacker can also conduct other types of attacks, as described in the following sections.

**Password-based attacks**
A commonality among most operating systems and network security plans is password-based access control. Access to both a computer and network resources are determined by a user name and password.

Historically, many versions of operating system components have not always protected identity information as it was passed through the network for validation. This might allow an eavesdropper to determine a valid user name and password and use it to gain access to the network by posing as a valid user.

When an attacker finds and accesses a valid user account, the attacker has the same rights as the actual user. For example, if the user has administrative rights, the attacker can create additional accounts for access at a later time.

After gaining access to a network with a valid account, an attacker can do any of the following:

- Obtain lists of valid user and computer names and network information.
- Modify server and network configurations, including access controls and routing tables.
- Modify, reroute, or delete data.

**Denial-of-service attack**
Unlike a password-based attack, the denial-of-service attack prevents normal use of a computer or network by valid users.

After gaining access to a network, an attacker can do any of the following:

- Distract information systems staff so that they do not immediately detect the intrusion. This gives an attacker the opportunity to make additional attacks.
- Send invalid data to applications or network services, causing applications or services to close or operate abnormally.
- Send a flood of traffic until a computer or an entire network is shut down.
- Block traffic, which results in a loss of access to network resources by authorized users.

**Man-in-the-middle attack**
As the name indicates, a man-in-the-middle attack occurs when someone between two users who are communicating is actively monitoring, capturing, and controlling the communication without the knowledge of the users. For example, an attacker can negotiate encryption keys with both users. Each user then sends encrypted data to the attacker, who can decrypt the data. When computers are communicating at low levels of the network layer, the computers might not be able to determine with which computers they are exchanging data.

**Compromised-key attack**
A key is a secret code or number required to encrypt, decrypt, or validate secured information. Although determining a key is a difficult and resource-intensive process for an attacker, it is possible. After an attacker determines a key, that key is referred to as a compromised key.

An attacker uses the compromised key to gain access to a secured communication without the sender or receiver being aware of the attack. With the compromised key, the attacker can decrypt or modify data. The attacker can also attempt to use the compromised key to compute additional keys, which might allow access to other secured communications.

**Sniffer attack**
A sniffer is an application or device that can read, monitor, and capture network data exchanges and packets. If the packets are not encrypted, a sniffer provides a full view of the data that is inside of the packet. Even encapsulated (tunneled) packets can be opened and read if they are not encrypted.

Using a sniffer, an attacker can do the following:

- Analyze a network and access information, eventually causing the network to stop responding or become corrupted.
- Read private communications.

**Application-layer attack**
An application-layer attack targets application servers by causing a fault in a server's operating system or applications. This results in the attacker gaining the ability to bypass normal access controls. The attacker takes advantage of this situation, gaining control of an application, system, or network, and can do any of the following:

- Read, add, delete, or modify data or an operating system.
- Introduce a virus that uses computers and software applications to copy viruses throughout the network.
- Introduce a sniffer program to analyze the network and gain information that can eventually be used to cause the network to stop responding or become corrupted.
- Abnormally close data applications or operating systems.
- Disable other security controls to enable future attacks.

# Protection against attacks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Protection against attacks**
IPSec protects data so that an attacker finds it extremely difficult or impossible to interpret it. The level of protection provided is determined by the strength of the security levels specified in your IPSec policy structure.

IPSec has a number of features that significantly reduce or prevent the following attacks:

- Sniffers (lack of confidentiality)

  The Encapsulating Security Payload (ESP) protocol in IPSec provides data confidentiality by encrypting the payload of IP packets.

- Data modification

  IPSec uses cryptography-based keys, shared only by the sending and receiving computers, to create a cryptographic checksum for each IP packet. Any modification to the packet data alters the checksum, which indicates to the receiving computer that the packet was modified in transit.

- Identity spoofing, password-based, and application-layer attacks

  IPSec allows the exchange and verification of identities without exposing that information to interpretation by an attacker. Mutual verification (authentication) is used to establish trust between the communicating systems and only trusted systems can communicate with each other. After identities are established, IPSec uses cryptography-based keys, shared only by the sending and receiving computers, to create a cryptographic checksum for each IP packet. The cryptographic checksum ensures that only the computers that have knowledge of the keys could have sent each packet.

- Man-in-the-middle attacks

  IPSec combines mutual authentication with shared, cryptography-based keys.

- Denial-of-service attacks

  IPSec uses IP packet filtering methodology as the basis for determining whether communication is allowed, secured, or blocked, according to the IP address ranges, IP protocols, or even specific TCP and UDP ports.

# Policy-based security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Policy-based security**

Although stronger security methods based on cryptography have become necessary to fully protect communication, they can often greatly increase administrative overhead. To reduce overhead, IPSec uses policy-based administration.

IPSec policies, rather than application programming interfaces (APIs), are used to configure IPSec security services. The policies provide variable levels of protection for most traffic types in most existing networks.

You can configure IPSec policies to meet the security requirements of a computer, application, organizational unit, domain, site, or global corporation. You can use the IP Security Policy Management console provided in Microsoft® Windows® XP and the Windows Server 2003 family to define IPSec policies for computers through Active Directory® (for domain members) or on the local computer (for computers that do not belong to domains).

# Simplified deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Simplified deployment**
To achieve secure communications with a low cost of ownership, Microsoft® Windows® XP and the Windows Server 2003 family simplify the deployment of IPSec with the following features:

- Integration with the security framework in Windows 2000 and the Windows Server 2003 family.

  IPSec uses the secure domain in Windows 2000 and the Windows Server 2003 family as a trust model. By default, IPSec policies use the Active Directory® default authentication method (Kerberos V5 authentication) to identify and trust communicating computers. Computers that are members of a Windows 2000 or a Windows Server 2003 domain and are in trusted domains can easily establish IPSec-secured communications.

- Centralized IPSec policy administration through Active Directory

  IPSec policies can be assigned through Group Policy configuration of Active Directory domains and organizational units. This allows the IPSec policy to be assigned at the domain, site, or organizational unit level, eliminating the administrative overhead of configuring each computer separately.

- Transparency of IPSec to users and applications

  Integration at the IP layer (layer 3) provides security for any protocol in the TCP/IP suite of protocols. You do not need separate security for each protocol in the TCP/IP suite of protocols, because applications using TCP/IP pass the data to the IP protocol layer, where it is secured.

- Flexible security configuration

  The security services within each policy can be customized to meet the majority of security requirements for the network and data traffic.

- Automatic key management

  Internet Key Exchange (IKE) services dynamically exchange and manage cryptographic keys between communicating computers.

- Automatic security negotiation

  Internet Key Exchange (IKE) services dynamically negotiate a common set of security settings between communicating computers, eliminating the need for both computers to have identically configured policies.

- Public key infrastructure support

  Using public key certificates for authentication is supported. This allows trust and secure communication for computers that do not belong to a trusted Windows 2000 or Windows Server 2003 domain, non-Microsoft operating systems, computers that have membership in untrusted domains, and instances in which computer access must be restricted to a smaller group than domain authentication allows.

- Preshared key support

  If authentication using the Kerberos V5 protocol or public key certificates is not possible, a preshared authentication key can be configured.

**Important**

- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

**Note**

- The integration with the security framework in Windows 2000 or the Windows Server 2003 family, centralized IPSec policy administration through Active Directory, and the use of the Kerberos V5 protocol as described here do not apply to computers running Windows XP Home Edition. You cannot administer Active Directory-based IPSec policy from computers running Windows XP Home Edition because these computers cannot join Active Directory domains.

# Understanding Internet Protocol Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Internet Protocol security**

This section introduces the key concepts and terminology you should review before you deploy IPSec. This section covers:

- The IPSec process
- Virtual private networking with IPSec
- IPSec Security Services
- IPSec Protocol Types
- IPSec Components
- Description of an IPSec Policy

# The IPSec process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The IPSec process**

This topic provides an overview of IPSec concepts that are central to understanding the IPSec process, including IPSec policy configuration and the Internet Key Exchange (IKE) protocol. In addition, this topic describes how IPSec network traffic processing works, using two intranet computers as an example.

**IPSec Policy Configuration**

In Windows 2000, Windows XP, and the Windows Server 2003 family, IPSec is implemented primarily as an administrative tool that you can use to enforce security policies on IP network traffic. A security policy is a set of packet filters that define network traffic as it is recognized at the IP layer. A filter action defines the security requirements for the network traffic. A filter action can be configured to: **Permit**, **Block**, or **Negotiate security** (negotiate IPSec).

IPSec filters are inserted into the IP layer of the computer TCP/IP networking protocol stack so that they can examine (filter) all inbound or outbound IP packets. Except for a brief delay required to negotiate a security relationship between two computers, IPSec is transparent to end-user applications and operating system services.

A collective set of IPSec security settings is known as an IPSec policy. Windows 2000, Windows XP, and the Windows Server 2003 family provide a graphical user interface and several command-line tools that you can use to configure an IPSec policy, and then assign it to a computer.

To ensure that IPSec communication is successful and that IPSec meets the security requirements of your organization, you must carefully design, configure, coordinate, and manage IPSec policies. In many organizations, one administrator might be responsible for configuring and managing IPSec policies for many, if not all, computers.

**Internet Key Exchange (IKE) security associations**

The IKE protocol is designed to securely establish a trust relationship between each computer, to negotiate security options, and dynamically generate shared, secret cryptographic keying material. The agreement of security settings associated with keying material is called a security association, also known as an SA. These keys will provide authenticity, integrity, and optionally, encryption of IP packets that are sent using the security association. IKE negotiates two types of security associations:

- A main mode security association (the IKE security association that is used to protect the IKE negotiation itself).
- IPSec security associations (the security associations that are used to protect application traffic).

You can configure IPSec policy settings for both types of security associations.

The IPSec service interprets an IPSec policy, expanding it into the components that it needs to control the IKE negotiation. The IPSec policy contains one definition of a packet filter. The packet filter is interpreted in two ways: one uses only the address and identity information to allow IKE to establish a main mode SA (the IKE security association); the other allows IKE to establish the IPSec security associations (also known as quick mode security associations).

**IPSec network traffic processing**

The following illustration shows how IPSec works in terms of the IPSec components for two intranet computers.




For simplicity, this example is of an intranet in which two computers have an active IPSec policy.

1. Alice, using a data application on Computer A, sends an application IP packet to Bob on Computer B.
2. The IPSec driver on Computer A checks its outbound IP filter lists and determines that the packets should be secured.

3. The action is to negotiate security, so the IPSec driver notifies IKE to begin negotiations.

4. The IKE service on Computer A completes a policy lookup, using its own IP address as the source and the IP address of Computer B as the destination. The main mode filter match determines the main mode settings that Computer A proposes to Computer B. Computer A sends the first IKE message in main mode, using UDP source port 500, destination port 500. IKE packets receive special processing by the IPSec driver to bypass filters.

5. Computer B receives an IKE main mode message requesting secure negotiation. It uses the source IP address and the destination IP address of the UDP packet to perform a main mode policy lookup, to determine which security settings to agree to. Computer B has a main mode file that matches, and so replies to begin negotiation of the main mode SA.

6. Computer A and Computer B now negotiate options, exchange identities, verify trust in those identities (authentication), and generate a shared master key. They have now established an IKE main mode SA. Computer A and Computer B must mutually trust each other.

7. Computer A then performs an IKE quick mode policy lookup, using the full filter to which the IPSec driver matched the outbound packet. Computer A selects the quick mode security settings and proposes them, and the quick mode filter, to Computer B.

   Computer B also performs an IKE quick mode policy lookup, using the filter description offered by Computer A. Computer B selects the security settings required by its policy and compares those settings to those offered by computer A. Computer B accepts one set of options and completes the remainder of the IKE quick mode negotiation to create a pair of IPSec security associations.

8. One IPSec SA is inbound and one IPSec SA is outbound. The IPSec SAs are identified by a Security Parameter Index (SPI), which is inserted into the IPSec header of each packet sent.

9. The IPSec driver on Computer A uses the outbound SA to sign and, if required, encrypt the packets. If the network adapter can perform hardware offload of IPSec cryptographic functions, the IPSec driver formats the packets, but does not perform the IPSec cryptographic functions.

10. The IPSec driver passes the packets to the network adapter driver, indicating whether the adapter must perform the IPSec cryptographic functions. The network adapter transmits the packets into the network.

11. The network adapter driver at Computer B receives the encrypted packets from the network. The SPI is used by the receiver of an IPSec packet to find the corresponding IPSec security association, with the cryptographic keys required to verify and decrypt the packets. If the network adapter can decrypt the packets in hardware, it verifies whether it can recognize the SPI. If it cannot decrypt the packets in hardware, or if it cannot recognize the SPI, it passes the packets up to the IPSec driver.

12. The IPSec driver on Computer B uses the inbound SA SPI to retrieve the keys required to validate authentication and integrity and, if required, to decrypt the packets.

13. The IPSec driver converts the packets from IPSec format back to standard IP packet format. It passes the validated and decrypted IP packets to the TCP/IP driver, which passes them to the receiving application on Computer B.

14. The IPSec SAs continue to provide very strong, transparent protection for application data traffic. The IPSec SAs are automatically refreshed by an IKE quick mode negotiation for as long as the application sends and receives data. When the application stops sending and receiving data, the IPSec SAs become idle and are deleted.

15. Typically, the IKE main mode SA is not deleted. By default, the main mode SA has a lifetime of 8 hours. You can configure the main mode SA lifetime to as short as 5 minutes to a maximum of 48 hours. Whenever more traffic is sent, a new quick mode is negotiated automatically to create two new IPSec SAs to protect application traffic. This process is rapid, because the main mode SA already exists. If a main mode SA expires, it is automatically renegotiated as needed.

Any routers or switches in the path between the communicating computers simply forward the encrypted IP packets to their

destination. However, if there is a firewall, security router, or proxy server in the path, IPSec and IKE traffic might not be forwarded. These devices must be configured to allow IPSec and IKE protocol packets through. If the IPSec packets are not encrypted, the firewall or security router can still inspect the TCP or UDP ports or other content in the packets. If the contents of these packets are modified after they are sent, the receiving IPSec computer detects the modification and discards the packets.

However, the Windows Server 2003 family implementation of IPSec provides support for a new Internet specification that allows IPSec packets to be modified by a network address translator (NAT). IPSec Encapsulating Security Payload (ESP) packets can pass through NATs that allow UDP traffic. The IKE protocol automatically detects the presence of a NAT and uses User Datagram Protocol-ESP (UDP-ESP) encapsulation to allow IPSec traffic to pass through the NAT.

# Virtual private networking with IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private networking with IPSec**

Tunneling is the entire process of encapsulation, routing, and decapsulation. Tunneling wraps, or encapsulates, the original packet inside a new packet. This new packet might have new addressing and routing information, which enables it to travel through a network. When tunneling is combined with data confidentiality, the original packet data (as well as the original source and destination) is not revealed to those listening to traffic on the network. After the encapsulated packets reach their destination, the encapsulation is removed, and the original packet header is used to route the packet to its final destination.

The tunnel itself is the logical data path through which the encapsulated packets travel. To the original source and destination peer, the tunnel is usually transparent and appears as just another point-to-point connection in the network path. The peers are unaware of any routers, switches, proxy servers, or other security gateways between the tunnel's beginning point and the tunnel's end point. When tunneling is combined with data confidentiality, it can be used to provide a virtual private network (VPN).

The encapsulated packets travel through the network inside the tunnel. In this example, the network is the Internet. The gateway might be an edge gateway that stands between the outside Internet and the private network. The edge gateway can be a router, firewall, proxy server, or other security gateway. Also, two gateways can be used inside the private network to protect traffic across untrusted parts of the network.

**Tunnel types that use IPSec**

In Windows XP and the Windows Server 2003 family, two types of tunneling are provided that use IPSec:

- **The combination of the Layer Two Tunneling Protocol (L2TP) and IPSec (L2TP/IPSec)**

  L2TP is used to tunnel the data across a shared or public network such as the Internet, and IPSec Encapsulating Security Payload (ESP) is used to encrypt the data. L2TP/IPSec can be used to tunnel IP or Internetwork Packet Exchange (IPX) traffic.

- **IPSec in tunnel mode**

  When IPSec is used in tunnel mode, IPSec itself provides encapsulation for IP traffic only. The primary reason for using IPSec tunnel mode is interoperability with other routers, gateways, or end-systems that do not support L2TP/IPSec or PPTP VPN tunneling. IPSec tunnel mode is supported only in gateway-to-gateway tunneling scenarios and for certain server-to-server or server-to-gateway configurations as an advanced feature. IPSec tunnel mode is not supported for remote access VPN scenarios. L2TP/IPSec or PPTP should be used for remote access VPN connections.

  For information about IPSec tunnel mode and configuring IPSec tunnels, see Tunnel mode.

**How L2TP and IPSec work**

L2TP and IPSec are combined to provide both tunneling and security for IP, IPX, and other protocol packets across any IP network. IPSec can also perform tunneling without L2TP, but it is only recommended for interoperability, when one of the gateways does not support L2TP or PPTP.

L2TP encapsulates original packets first inside a PPP frame (performing compression when possible) and then inside a UDP message using port 1701. Because the UDP message is an IP payload, L2TP uses IPSec transport mode to secure the tunnel. This design supports RFC 3193 ("Securing L2TP Using IPSec"), and it is compatible with Windows XP and Windows 2000. The IPSec Internet Key Exchange (IKE) protocol negotiates security for the L2TP tunnel using certificate-based or preshared key authentication. If IPSec main mode and quick mode security associations are successfully established, L2TP negotiates the tunnel, including compression and user authentication options, and performs PPP-based user authentication.




The original packet, shown here as the **PPP payload**, includes the original and ultimate source and destination addresses used

on the private network. The outer IP header, shown as **IP header**, contains the source and destination addresses of the VPN client and server on the public network. The L2TP header includes L2TP tunnel control information. The PPP header identifies the protocol of the original packet (for example, IP or IPX).

### Note

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

#### Common L2TP/IPSec scenarios

Two common scenarios for L2TP/IPSec are securing communications between remote access clients and the corporate network across the Internet and securing communications between branch offices.

#### Remote access clients with L2TP/IPSec

A common requirement is securing communications between remote access clients and the corporate network across the Internet. Such a client might be a sales consultant who spends most of the time traveling, or an employee working from a home office.




In this figure, the remote gateway is a server that provides edge security for the corporate intranet. The remote client represents a roaming user who requires regular access to network resources and information. An ISP is used as an example to demonstrate the path of communication when the client uses an ISP to access the Internet. L2TP is combined with IPSec to provide a simple, efficient way to build the tunnel and protect the data across the Internet.

For more information about configuring L2TP-based client remote access VPN connections, see Setting Up Remote Access VPNs, Make a virtual private network (VPN) connection, and Virtual Private Network (VPN) Connections.

#### Connecting branch offices with L2TP/IPSec

A large corporation often has multiple sites that require communication--for example, a corporate office in New York and a sales office in Washington. In this case, L2TP is combined with IPSec to provide the VPN connection and protect the data between the sites.

In this figure, the router running the Windows Server 2003 family provides edge security. The routers might have a leased line, dial-up, or other type of Internet connection. The IPSec SA and L2TP tunnel runs between the routers only and provides protected communication through the Internet.

For more information about configuring L2TP-based router-to-router VPN connections, see Deploying Router-to-Router VPNs.

# IPSec Security Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec security services**

IPSec provides a high level of security by using cryptography-based services. Cryptography allows information to be transmitted securely by hashing (to provide integrity and authentication) and encrypting (to provide confidentiality) the information. IPSec provides or uses the following security services:

- Security properties
- Public key certificate-based authentication
- Preshared key authentication
- Data integrity with hash functions
- Data confidentiality with encryption
- Key management and protection

# Security properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security properties**

IPSec provides the following properties for secured communications:

- Anti-replay

  Ensures the uniqueness of each IP packet. Anti-replay is also called replay prevention. Anti-replay ensures that data captured by an attacker cannot be reused or replayed to establish a session or gain information illegally. This protects against attempts to intercept a message and then use the identical message to illegally gain access to resources, possibly months later.

- Integrity

  Protects data from unauthorized modification in transit, ensuring that the data received is exactly the same as the data sent. Hash functions sign each packet with a cryptographic checksum using a shared, secret key. Only the sender and receiver have the key that is used to calculate the checksum. If the packet contents have changed, the checksum verification fails and the packet is discarded.

- Confidentiality (encryption)

  Ensures that data is only disclosed to intended recipients. This is achieved by encrypting the data before transmission, ensuring that the data cannot be read during transmission, even if the packet is monitored or intercepted. Only the communicating computers with the shared, secret key are able to read the data after it has been decrypted. This property is optional and is dependent on IPSec policy settings.

- Authentication

  Verifies that a message could only have been sent from a computer that has knowledge of the shared, secret key. The sender includes (with the message) a message authentication code with a calculation that includes the shared, secret key. The receiver performs the same calculation and, if the receiver's calculation does not match the message authentication code that is included in the message, the message is discarded. The message authentication code is the same as the cryptographic checksum that is used for integrity.

# Public key certificate-based authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Public key certificate-based authentication**

An effectively implemented public key infrastructure, in which security credentials can be presented without compromising them in the process, resolves many security problems. IPSec works with your public key infrastructure to enable certificate-based authentication of computers.

A public key certificate is one type of authentication that reliably provides identity verification. Public key certificates are like digital passports. They are used to verify the identities of non-Microsoft operating system computers, stand-alone computers, clients that are not members of a trusted domain, or computers that are not running the Kerberos V5 authentication protocol (the default authentication method in Windows XP and Windows Server 2003 family) and Routing and Remote Access.

# Preshared key authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Preshared key authentication**

IPSec can use preshared keys for authentication. Preshared means that the parties agree on a shared, secret key that is used for authentication in an IPSec policy.

During security negotiation, information is encrypted before transmission by using a session key. The session key is created by using a Diffie-Hellman calculation and the shared, secret key. Information is decrypted on the receiving end using the same key. One IPSec peer authenticates the other peer's packet by decryption and verification of the hash inside the packet (the hash inside the packet is a hash of the preshared key). If authentication fails, the packet is discarded.

The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext in the registry. In Active Directory, preshared keys are stored in readable hexadecimal format.

Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment. For more information about IPSec authentication methods, see Authentication methods and Public key certificate-based authentication.

# Data integrity with hash functions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data integrity with hash functions**

Hash message authentication codes (HMAC) sign packets to verify that the information received is exactly the same as the information sent. This is called integrity. HMACs provide integrity through a keyed hash, the result of a mathematical calculation on a message using a hash function (algorithm) combined with a shared, secret key. A hash is commonly described as a signature on the packet. However, a hash differs from a digital signature. A hash uses a secret, shared key, and a digital signature uses public key technology and the private key of the sending computer. A digital signature provides nonrepudiation, and a hash does not. Nonrepudiation ensures that a communication can be proven to have originated from a specific person whose identity can be verified. It also ensures that the communication actually occurred.

Hash functions are also called one-way functions because it is easy to determine the hash from the message but mathematically infeasible to determine the message from the hash. Conversely, in two-way functions, the original message can be determined from its converted form. Encryption and decryption schemes are examples of two-way functions.

The hash is a cryptographic checksum or message integrity code (MIC) that each party must compute to verify the message. For example, the sending computer uses a hash function and shared key to compute the checksum for the message, including it with the packet. The receiving computer must perform the same hash function on the received message and shared key and compare it to the original (included in the packet from the sender). If the message has changed in transit, the hash values are different and the packet is rejected.

For integrity, you can choose between two hash functions when setting policy:

- MD5

  Message Digest 5 (MD5) is based on RFC 1321. MD5 completes four passes over the data blocks, using a different numeric constant for each word in the message on each pass. The number of 32-bit constants used during the MD5 computation ultimately produces a 128-bit hash that is used for the integrity check.

- SHA1

  Secure Hash Algorithm 1 (SHA1) was developed by the National Institute of Standards and Technology as described in Federal Information Processing Standard (FIPS) PUB 180-1. The SHA process is closely modeled after MD5. The SHA1 computation results in a 160-bit hash that is used for the integrity check. Because longer hash lengths provide greater security, SHA is stronger than MD5.

# Data confidentiality with encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data confidentiality with encryption**
IPSec in Windows XP and the Windows Server 2003 family uses the US Data Encryption Standard (DES) to provide confidentiality (data encryption). The DES algorithm was published in 1977 by the US National Bureau of Standards.

DES is a block cipher that uses a 56-bit key. A block cipher is an encryption algorithm that operates on a fixed size block of data. DES encrypts data in 64-bit blocks using a 64-bit key. The key appears to be a 64-bit key, but one bit in each of the 8 bytes is used for error checking, resulting in 56 bits of usable key.

Cipher block chaining (CBC) is also used to hide patterns of identical blocks of data within a packet. An initialization vector (an initial random number) is used as the first random block to encrypt and decrypt a block of data. Different random blocks are used in conjunction with the secret key to encrypt each successive block. This ensures that identical sets of unsecured data (plaintext) result in unique, encrypted data blocks.

The Windows XP and the Windows Server 2003 family implementation of IPSec supports the use of:

- DES

  Used when the high security and overhead of 3DES are not necessary.

- 3DES

  Used when high security is required. 3DES processes each block three times, using a unique 56-bit key each time:

  1. Encryption on the block with key 1
  2. Decryption on the block with key 2
  3. Encryption on the block with key 3

     This process is reversed if the computer is decrypting a packet.

**Note**

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

# Key management and protection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key management and protection**
A key is a secret code or number that is required to read, modify, or verify secured data. Keys are used in conjunction with algorithms (a mathematical process) to secure data.

**Key management**
Windows XP and the Windows Server 2003 family automatically handle key generation and implement the following keying properties:

- Variable key lengths

  With IPSec, you can select integrity and encryption algorithms that provide a range of key lengths. Every time the length of a key is increased by one bit, the number of possible keys doubles, making it exponentially more difficult to determine the key. Shorter key lengths are processed more quickly than longer key lengths. Longer key lengths provide a higher level of security, but are more resource-intensive to process.

- Key material generation: The Diffie-Hellman algorithm

  To enable secure communication, two computers must be able to gain the same shared key (session key), without sending the key across a network and compromising the secret.

  The Diffie-Hellman algorithm predates Rivest-Shamir-Adleman (RSA) cryptographic algorithms and offers better performance. It is one of the oldest and most secure algorithms used for key exchange. The two peers publicly exchange keying information, which Windows XP and the Windows Server 2003 family additionally protect with a hash function signature. Neither peer ever exchanges the actual key; however, after their exchange of keying material, each can generate the identical shared key.

  The Diffie-Hellman keying material that is exchanged by the two parties can be based on 768, 1,024, or 2,048 bits of keying material, known as Diffie-Hellman groups 1, 2, and 2048, respectively. The strength of the Diffie-Hellman group is proportional to the strength of the key that is computed from the Diffie-Hellman exchange. Strong Diffie-Hellman groups combined with longer key lengths increase the computational difficulty of determining a secret key.

  IPSec uses the Diffie-Hellman algorithm to provide the keying material for all other encryption keys. Diffie-Hellman does not provide authentication. In the Windows XP and the Windows Server 2003 family implementation of IPSec, identities are authenticated after the Diffie-Hellman exchange takes place, providing protection against man-in-the-middle attacks.

- Dynamic rekeying

  IPSec policy controls how often a new key is generated during communication, by using a method called dynamic rekeying. The communication is sent in blocks, and each block of data is secured with a different key. This prevents an attacker who has obtained part of a communication, and the corresponding session keys, from obtaining the remainder of the communication. This on-demand security negotiation and automatic key management service is provided using IKE, as defined in RFC 2409, "The Internet Key Exchange (IKE)."

  IPSec policy allows you to control how often a new key is generated. If no values are configured, keys are regenerated automatically at default intervals.

**Key protection**
The following features enhance the base prime numbers (keying material) and the strength of the keys for the master and session keys.

- Key lifetimes

  Key lifetimes determine when a new key is generated, rather than how it is generated. A key lifetime allows you to force

an automatic key regeneration after a specific interval of either kilobytes (KB) or seconds, whichever occurs first. This ensures that even if an attacker is able to decipher part of a communication protected by one key, new keys protect the remainder of the communication. Whenever a key lifetime is reached, the security association (SA) is also renegotiated and the key is refreshed or regenerated. For this reason, key lifetimes are also referred to as SA lifetimes. You can specify key lifetimes for the master key and for session keys.

The master key lifetime corresponds to the IKE SA lifetime created by the IKE main mode (phase I) negotiation. It is configured in terms of time and number of quick mode negotiations, and it applies to all security rules in the IPSec policy. The IKE SA and master key typically have a long lifetime (the default value is eight hours). Master keys are much more resource-intensive to generate than session keys, due to the fact that reauthentication and additional Diffie-Hellman exchanges are required.

The session keys correspond to the IPSec SAs that are used to protect program traffic. The session keys and IPSec SAs are quickly derived from the master key by the IKE quick mode (phase II) negotiation. Thus the session keys are used to protect the data, and they have lifetimes based on the amount of data sent and the amount of time elapsed since the key started being used. Typically, session keys have shorter lifetimes than master keys. The default values are 100,000 KB (approximately 100 megabytes) or one hour. When the session keys are refreshed, new IPSec SAs replace the old ones. The IPSec SA session keys must be refreshed before either the data or time lifetime expires, or traffic is discarded. A session key will be deleted if the IPSec SAs become idle (by default, SAs become idle after five minutes). Because the master key lives longer than the session key, it allows new IPSec SAs and session keys to be established quickly. To prevent data loss, the quick mode SA session key is generated shortly before the main mode SA expires. The IKE SA and corresponding master key are not deleted when session keys and IPSec SAs are deleted, unless the specified number of quick mode SA negotiations has elapsed.

For example, if a communication takes one hour (3,600 seconds), and if you specify the minimum session key lifetime of five minutes (300 seconds), more than 12 keys are generated to complete the communication. If a 100 megabytes (MB) file is transferred over a fast corporate LAN using an IPSec security association with a 100 MB and one hour lifetime, at least one, if not two, session rekeys occur.

**Notes**

- To provide maximum protection against cryptographic attacks, the amount of data encrypted using Data Encryption Standard (DES) or Triple DES (3DES) by a single session key should not exceed 100 MB. In a high-security environment, consider using session key lifetimes for MD5 or Secure Hash Algorithm (SHA-1) that are shorter then the default values, to provide stronger protection against future cryptographic attacks.

- In certain scenarios, there might be practical operational limits on key lifetimes. The shorter the key lifetimes, the more CPU cycles are spent refreshing keys. This might limit either the number of active IPSec SAs that a computer can maintain or the throughput performance of IPSec-secured traffic.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

- Session key refresh limit

  Repeated rekeying off of a session key can compromise the Diffie-Hellman shared secret. Because of this, you might want to set a limit to the number of quick mode session keys that can be derived from a main mode negotiation.

  If you have enabled master key perfect forward secrecy (PFS), the session key limit is not used. Setting a session key limit to 1 is identical to enabling master key PFS. If both a master key lifetime and a session limit are specified, whichever limit is reached first causes a new main mode negotiation. By default, IPSec policy does not specify a session limit.

- Diffie-Hellman groups

Diffie-Hellman groups are used to determine the length of the base prime numbers (key material) for the Diffie-Hellman exchange. The cryptographic strength of any key derived from a Diffie-Hellman exchange depends, in part, on the strength of the Diffie-Hellman group on which the prime numbers are based.

Diffie-Hellman Group 2048 (high) is stronger than Group 2 (medium), which is stronger than Group 1 (low). When a stronger group is used, the key that is derived from a Diffie-Hellman exchange is stronger and more difficult for an attacker to break.

IKE negotiates which specific group to use, ensuring that there are not any negotiation failures that result from a mismatched Diffie-Hellman group between the two peers.

If session key PFS is enabled, a new Diffie-Hellman key is negotiated during the first quick mode SA negotiation. This new key removes the dependency of the session key on the Diffie-Hellman exchange that is performed for the master key.

Both the initiator and responder must have session key PFS enabled, or negotiation fails.

The Diffie-Hellman group is the same for both the main mode and quick mode SA negotiations. When session key PFS is enabled, even though the Diffie-Hellman group is set as part of the main mode SA negotiation, it affects any rekeys during session key establishment.

**Important**

- For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.

  For more information about Diffie-Hellman groups and key exchange, see Key exchange methods.

**Note**

- Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.

- Perfect forward secrecy

  Unlike key lifetimes, PFS determines how a new key is generated, rather than when it is generated. Specifically, PFS ensures that the compromise of a single key permits access only to data that is protected by it, not necessarily to the entire communication. To achieve this, PFS ensures that a key used to protect a transmission cannot be used to generate additional keys. In addition, if the key that was used was derived from specific keying material, that material cannot be used to generate other keys.

  Master key PFS requires a reauthentication and is resource intensive. When it is enabled, IKE must reauthenticate identities, increasing overhead for domain controllers when Kerberos V5 is used for authentication. It requires a new main mode negotiation for every quick mode negotiation that occurs.

  Session key PFS can be used without a reauthentication and is less resource intensive. Session key PFS results in a Diffie-Hellman exchange to generate new keying material. It requires only four messages and no authentication.

  Unlike session key PFS, master key PFS does not have to be enabled on both peers. This is because master key PFS is not part of the SA negotiation. If the responder requires PFS, and the quick mode SA of the initiator expires, the responder rejects the message from the initiator and requires a new main mode negotiation. The initiator expires the main mode SA and renegotiates. PFS can be individually set for both the master and session keys.

# IPSec Protocol Types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec protocol types**
IPSec protocols provide data and identity protection for each IP packet by adding their own security protocol header to each packet.

For information about AH and ESP when used in transport and tunnel mode, see the following topics:

- Transport mode
- Tunnel mode

# Transport mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Transport mode**
Transport mode is the default mode for IPSec, and it is used for end-to-end communications (for example, for communications between a client and a server). When transport mode is used, IPSec encrypts only the IP payload. Transport mode provides the protection of an IP payload through an AH or ESP header. Typical IP payloads are TCP segments (containing a TCP header and TCP segment data), a UDP message (containing a UDP header and UDP message data), or an ICMP message (containing an ICMP header and ICMP message data).

**Authentication Header transport mode**
Authentication Header (AH) provides authentication, integrity, and anti-replay protection for the entire packet (both the IP header and the data payload carried in the packet). It does not provide confidentiality, which means that it does not encrypt the data. The data is readable, but protected from modification. AH uses keyed hash algorithms to sign the packet for integrity. For more information, see Data integrity with hash functions.

For example, Alice on Computer A sends data to Bob on Computer B. The IP header, the AH header, and the data are protected with integrity. This means Bob can be certain it was really Alice who sent the data and that the data was unmodified.

Integrity and authentication are provided by the placement of the AH header between the IP header and the IP payload, as shown in the following illustration.

AH is identified in the IP header with an IP protocol ID of 51. AH can be used alone or combined with the Encapsulating Security Payload (ESP) protocol.

The AH header contains the following fields:

- Next Header

  Identifies the IP payload by using the IP protocol ID. For example, a value of 6 represents TCP.

- Length

  Indicates the length of the AH header.

- Security Parameters Index (SPI)

  Used in combination with the destination address and the security protocol (AH or ESP) to identify the correct security association for the communication. The receiver uses this value to determine with which security association the packet is identified.

- Sequence Number

  Provides anti-replay protection for the packet. The sequence number is a 32-bit, incrementally increasing number (starting from 1) that indicates the packet number sent over the security association for the communication. The sequence number cannot repeat for the life of the quick mode security association. The receiver checks this field to verify that a packet for a security association with this number has not already been received. If one has been received, the packet is rejected.

- Authentication Data

  Contains the integrity check value (ICV), also known as the message authentication code, which is used to verify both message authentication and integrity. The receiver calculates the ICV value and checks it against this value (which is calculated by the sender) to verify integrity. The ICV is calculated over the IP header, the AH header, and the IP payload.

**Packet signature with the AH header**
AH signs the entire packet for integrity, with the exception of some fields in the IP header which might change in transit (for example, the Time to Live and Type of Service fields). If another IPSec header is being used in addition to AH, the AH header is inserted before any other IPSec headers. The AH packet signature is shown in the following illustration.




**Encapsulating Security Payload transport mode**
Encapsulating Security Payload (ESP) provides confidentiality (in addition to authentication, integrity, and anti-replay protection) for the IP payload. ESP in transport mode does not sign the entire packet. Only the IP payload (not the IP header) is protected. ESP can be used alone or in combination with AH.

For example, Alice on Computer A sends data to Bob on Computer B. The IP payload is encrypted and signed for integrity. Upon receipt, after the integrity verification process is complete, the data payload in the packet is decrypted. Bob can be certain that it was Alice who sent the data, the data is unmodified, and no one else was able to read it.

ESP is identified in the IP header with the IP protocol ID of 50. As shown in the following illustration, the ESP header is placed before the IP payload, and an ESP trailer and ESP authentication trailer is placed after the IP payload.




The ESP header contains the following fields:

- Security Parameters Index

  Identifies the correct security association for the communication when used in combination with the destination address and the security protocol (AH or ESP). The receiver uses this value to determine the security association with which this packet should be identified.

- Sequence Number

  Provides anti-replay protection for the packet. The sequence number is a 32-bit, incrementally increasing number (starting from 1) that indicates the packet number sent over the quick mode security association for the communication. The sequence number cannot repeat for the life of the quick mode security association. The receiver checks this field to verify that a packet for a security association with this number has not already been received. If one has been received, the packet is rejected.

The ESP trailer contains the following fields:

- Padding

  Padding of 0 to 255 bytes is used to ensure that the encrypted payload with the padding bytes are on byte boundaries required by encryption algorithms.

- Padding Length

  Indicates the length of the Padding field in bytes. The receiver uses this field to remove padding bytes after the encrypted payload with the padding bytes has been decrypted.

- Next Header

  Identifies the type of data in the payload, such as TCP or UDP.

The ESP authentication trailer contains the following field:

- Authentication Data

  Contains the integrity check value (ICV), also known as the message authentication code, which is used to verify both

message authentication and integrity. The receiver calculates the ICV value and checks it against this value (which is calculated by the sender) to verify integrity. The ICV is calculated over the ESP header, the payload data, and the ESP trailer.

**Packet signature and encryption**

As shown in the following illustration, ESP provides protection for IP payloads. The signed portion of the packet indicates where the packet has been signed for integrity and authentication. The encrypted portion of the packet indicates what information is protected with confidentiality.




The IP header is not signed and is not necessarily protected from modification. To provide data integrity and authentication for the IP header, use ESP and AH.

# Tunnel mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tunnel mode**

When IPSec tunnel mode is used, IPSec encrypts the IP header and the payload, whereas transport mode only encrypts the IP payload. Tunnel mode provides the protection of an entire IP packet by treating it as an AH or ESP payload. With tunnel mode, an entire IP packet is encapsulated with an AH or ESP header and an additional IP header. The IP addresses of the outer IP header are the tunnel endpoints, and the IP addresses of the encapsulated IP header are the ultimate source and destination addresses.

IPSec tunnel mode is useful for protecting traffic between different networks, when traffic must pass through an intermediate, untrusted network. Tunnel mode is primarily used for interoperability with gateways, or end-systems that do not support L2TP/IPSec or PPTP connections. You can use tunnel mode in the following configurations:

- Gateway-to-gateway
- Server-to-gateway
- Server-to-server

**AH tunnel mode**

As shown in the following illustration, AH tunnel mode encapsulates an IP packet with an AH and IP header and signs the entire packet for integrity and authentication.




**ESP tunnel mode**

As shown in the following illustration, ESP tunnel mode encapsulates an IP packet with both an ESP and IP header and an ESP authentication trailer.




The signed portion of the packet indicates where the packet has been signed for integrity and authentication. The encrypted portion of the packet indicates what information is protected with confidentiality.

Because a new header for tunneling is added to the packet, everything that comes after the ESP header is signed (except for the ESP authentication trailer) because it is now encapsulated in the tunneled packet. The original header is placed after the ESP header. The entire packet is appended with an ESP trailer before encryption occurs. Everything that follows the ESP header, except for the ESP authentication trailer, is encrypted. This includes the original header which is now considered to be part of the data portion of the packet.

The entire ESP payload is then encapsulated within the new tunnel header, which is not encrypted. The information in the new tunnel header is used only to route the packet from origin to tunnel endpoint.

If the packet is being sent across a public network, it is routed to the IP address of the gateway for the receiving intranet. The gateway decrypts the packet, discards the ESP header, and uses the original IP header to route the packet to the intranet computer.

ESP and AH can be combined when tunneling, providing both confidentiality for the tunneled IP packet and integrity and authentication for the entire packet.

**Using IPSec tunnels**

IPSec tunnels provide security for IP traffic only. The tunnel is configured to protect traffic between either two IP addresses or two IP subnets. If the tunnel is used between two computers instead of two gateways, the IP address outside the AH or ESP payload is the same as the IP address inside the AH or ESP payload. In Windows XP and the Windows Server 2003 family, IPSec does not support protocol-specific or port-specific tunnels. You can configure tunnels by using the IP Security Policy Management and Group Policy consoles to configure and enable two rules:

1. A rule for the outbound traffic for the tunnel.

   The rule for the outbound traffic is configured with a filter list that describes the traffic to be sent through the tunnel and a tunnel endpoint of an IP address configured on the IPSec tunnel peer (the computer or router on the other side of the tunnel).

2. A rule for the inbound traffic for the tunnel.

   The rule for the inbound traffic is configured with a filter list that describes the traffic to be received through the tunnel and a tunnel endpoint of a local IP address (the computer or router on the local side of the tunnel).

Additionally, filter actions, authentication methods, and other settings need to be specified for each rule.

For conceptual information about IPSec policy tunnel settings, see Tunnel endpoint. For information about configuring an IPSec tunnel, see Specify an IPSec tunnel. For information about how tunneling is used for virtual private networking, see Virtual private networking with IPSec.

# IPSec Components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec components**

The following topics describe IPSec components that are installed with Windows XP and the Windows Server 2003 family:

- IPSec Policy Agent service
- Internet Key Exchange
- IPSec driver

# IPSec Policy Agent service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec Policy Agent service**

The purpose of the IPSec Policy Agent is to retrieve policy information and pass it to other IPSec components that require this information to perform security services, as shown in the following illustration.




The IPSec Policy Agent is a service that resides on each computer running a Windows Server 2003 operating system, appearing as **IPSec services** in the list of system services in the Services console. The IPSec Policy Agent:

- Retrieves the appropriate IPSec policy (if one has been assigned) from Active Directory, if the computer is a domain member, or from the local registry, if the computer is not a member of a domain.
- Polls for changes in policy configuration.
- Sends the assigned IPSec policy information to the IPSec driver.

If the computer is a member of a domain, policy retrieval occurs when the system starts, at the interval specified in the IPSec policy, and at the default Winlogon polling interval. You can also manually poll Active Directory for policy using the **gpupdate /target:computer** command.

The following are additional aspects of IPSec policy behavior for a computer that is a member of a domain:

- If IPSec policy information is centrally configured for computers that are domain members, the IPSec policy information is stored in Active Directory and cached in the local registry of the computer to which it applies.
- If the computer is temporarily not connected to the domain and policy is cached, new policy information for that computer replaces old, cached information when the computer reconnects to the domain.
- If the computer is a stand-alone computer or a member of a domain that is not using Active Directory for policy storage, IPSec policy is stored in the local registry.

The IPSec Policy Agent starts automatically at system start time. If there are no IPSec policies in Active Directory or the registry, or if the IPSec Policy Agent cannot connect to Active Directory, the IPSec Policy Agent waits for policy to be assigned or activated.

# Internet Key Exchange

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Key Exchange**
Before secured data can be exchanged, a security agreement between the two computers must be established. In this security agreement, called a security association (SA), both agree on how to exchange and protect information, as shown in the following illustration.




To build this agreement between the two computers, the IETF has established a standard method of security association and key exchange resolution named Internet Key Exchange (IKE) which:

- Centralizes security association management, reducing connection time.
- Generates and manages shared, secret keys that are used to secure the information.

This process not only protects communication between computers, it also protects remote computers that request secure access to a corporate network. In addition, this process works whenever the negotiation for the final destination computer (endpoint) is performed by a security gateway.

**Security association (SA) defined**
A security association (SA) is the combination of a negotiated key, security protocol, and security parameters index (SPI), which together define the security used to protect the communication from sender to receiver. The SPI is a unique, identifying value in the SA that is used to distinguish among multiple security associations that exist at the receiving computer. For example, multiple associations might exist if a computer is securely communicating with multiple computers at the same time. This is a common occurrence when the computer is a file server or a remote access server that serves multiple clients. In these situations, the receiving computer uses the SPI to determine which SA is used to process the incoming packets.

**Phase I or main mode SA**
In order to ensure successful and secure communication, IKE performs a two-phase operation. Confidentiality and authentication are ensured during each phase by the use of encryption and authentication algorithms that are agreed upon by the two computers during security negotiations. With the duties split between two phases, key creation can be rapidly accomplished.

During the first phase, the two computers establish a secure, authenticated channel. This is called the phase I or main mode SA. IKE automatically provides necessary identity protection during this exchange.

**Phase I or main mode negotiation**
The following are the steps that comprise a main mode negotiation.

1. Policy negotiation

   The following four mandatory parameters are negotiated as part of the main mode SA:

   - The encryption algorithm (DES or 3DES)
   - The integrity algorithm (MD5 or SHA1)
   - The Diffie-Hellman group to be used for the base keying material: Group 1 (768 bits of keying material) Group 2 (1,024 bits), or Group 2048 (2,048 bits)
   - The authentication method (Kerberos V5, certificate, or preshared key authentication)

     If certificates or preshared keys are used for authentication, the computer identity is protected. If Kerberos V5 authentication is used, the computer identity is unencrypted until encryption of the entire identity payload takes

place during authentication.

**Important**

For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.

For more information about Diffie-Hellman groups, see Key exchange methods. For more information about preshared key authentication, see Preshared key authentication.

2. Diffie-Hellman exchange (of public values)

   At no time are actual keys exchanged. Only the base information required by the Diffie-Hellman key determination algorithm to generate the shared, secret key is exchanged. After this exchange, the IKE service on each computer generates the master key that is used to protect authentication.

3. Authentication

   To prevent a successful man-in-the-middle attack, the computers attempt to authenticate the Diffie-Hellman key exchange. Without successful authentication, communication will not proceed. The master key is used, in conjunction with the negotiation algorithms and methods, to authenticate identities. The entire identity payload is hashed and encrypted using the keys generated from the Diffie-Hellman exchange in the second step. The payload includes the identity type (for authentication), port, and protocol. IPSec uses the following identity types for authentication: For certificate authentication, the certificate distinguished name and general name; for Kerberos V5 and preshared key authentication, IPv4 addresses, the fully qualified domain name (FQDN) of the computer, and FQDN of the user. The identity payload, regardless of which authentication method is used, is protected from both modification and interpretation.

The sender presents an offer for a potential security association to the receiver. The responder cannot modify the offer. Should the offer be modified, the initiator rejects the responder's message. The responder sends either a reply accepting the offer or a reply with alternatives.

Messages sent during this phase have an automatic retry cycle that is repeated five times. If a response is received before the retry cycle ends, standard SA negotiation begins. If allowed by IPSec policy, unsecured communications will begin after a brief interval. If unsecured communications begin, after five minutes of idle time (during which no messages are sent), secured communication negotiation is attempted the next time messages are sent. If messages are sent continuously, the communication remains unsecured during the lifetime set for the main mode policy. After the policy time has elapsed, a new secured communication negotiation attempt is made.

There is no preset limit to the number of exchanges that can take place. The number of SAs established is only limited by system resources. When estimating the number of SAs that can be established without significantly degrading computer performance, consider the CPU processing strength and RAM of the computer, the lifetime of the SA, and how much traffic is being sent over the SAs.

For more information about unsecured communications, see Filter action.

**Phase II or quick mode SA**
In this phase, SAs are negotiated on behalf of the IPSec driver.

**Phase II or quick mode negotiation**
The following are the steps that comprise a quick mode negotiation.

1. Policy negotiation occurs.

   The IPSec computers exchange the following requirements for securing the data transfer:

   - The IPSec protocol (AH or ESP)
   - The hash algorithm for integrity and authentication (MD5 or SHA1)

- The algorithm for encryption, if requested (3DES or DES)

  A common agreement is reached, and two SAs are established. One SA is for inbound communication and the other is for outbound communication.

2. Session key material is refreshed or exchanged.

   IKE refreshes the keying material and new shared keys are generated for packet integrity, authentication, and encryption (if negotiated). If rekeying is required, either a second Diffie-Hellman exchange (as described in main mode negotiation) occurs, or a refresh of the original Diffie-Hellman key is used.

3. The SAs and keys, along with the SPI, are passed to the IPSec driver.

The second negotiation of security settings and keying material (for the purpose of securing data) is protected by the main mode SA. As the first phase provided identity protection, the second phase provides protection by refreshing the keying material prior to sending data. IKE can accommodate a key exchange payload for an additional Diffie-Hellman exchange if a rekey is necessary--that is, master key perfect forward secrecy (PFS) is enabled. Otherwise, IKE refreshes the keying material from the Diffie-Hellman exchange completed in main mode.

Quick mode results in a pair of security associations, each with its own SPI and key. One SA is used for inbound communication, and the other for outbound communication.

**Notes**

- Although there are two separate quick mode SAs established, IP Security Monitor only displays a single quick mode SA.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

The retry algorithm for a message is similar to the process described in main mode negotiation. However, if this process times out for any reason during the second or higher negotiation off of the same main mode SA, a renegotiation of the main mode SA is attempted. If a message for this phase is received without an established main mode SA, it is rejected.

Using a single main mode SA for multiple quick mode SA negotiations increases the speed of the process. As long as the main mode SA does not expire, renegotiation and reauthentication are not necessary. The number of quick mode SA negotiations that can be performed is determined by IPSec policy settings.

**Note**

- Excessive rekeying off of the same main mode SA might make the shared, secret key vulnerable to a known plaintext attack. A known plaintext attack is a sniffer attack in which the attacker attempts to determine the encryption key from encrypted data based on known plaintext.

**SA lifetimes**

The main mode SA is cached to allow multiple quick mode SA negotiations (unless master key PFS is enabled). When a key lifetime is reached for the master or session key, the SA is renegotiated. In addition, the key is refreshed or regenerated.

When the default time-out period elapses for the main mode SA, or the master or session key lifetime is reached, a delete message is sent to the responder. The IKE delete message tells the responder to expire the main mode SA. This prevents additional new quick mode SAs from being created from the expired main mode SA. IKE does not expire the quick mode SA, because only the IPSec driver contains the number of seconds or bytes that have passed to reach the key lifetime.

Use caution when setting very different key lifetimes for master and session keys. For example, setting a master key lifetime of

eight hours and a session key lifetime of two hours might leave a quick mode SA in place for almost two hours after the main mode SA has expired. This occurs when the quick mode SA is generated shortly before main mode SA expiration.

It is generally recommended that all of the IKE settings (for example, master key PFS and key lifetime) and security methods remain at their defaults to avoid unnecessary administrative overhead. This provides a standard (medium) level of security. If your security plan calls for a high level of security, you should consider modifying the default security methods.

# IPSec driver

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec driver**

The IPSec driver receives the active IP filter list from the IPSec Policy Agent, as shown in the following illustration, and then attempts to match every inbound and outbound packet against the filters in the list.




When a packet matches a filter, it applies the filter action. When a packet does not match any filters, the packet is passed back without modification to the TCP/IP driver to be received or transmitted.

If the filter action permits transmission, the packet is received or sent with no modifications. If the action blocks transmission, the packet is discarded. If the action requires the negotiation of security, main mode and quick mode SAs are negotiated.

The negotiated quick mode security association (SA) and keys are used with both outbound and inbound processing. The IPSec driver stores all current quick mode SAs in a database. The IPSec driver uses the Security Parameters Index (SPI) field to match the correct SA with the correct packet.

When an outbound IP packet matches the IP filter list with an action to negotiate security, the IPSec driver queues the packet and then notifies Internet Key Exchange (IKE), which begins security negotiations with the destination IP address of that packet. If several outbound packets are going to the same destination and match the same filter before IKE has finished the negotiation, then only the last packet sent is saved.

After a successful negotiation is completed, the IPSec driver on the sending computer:

1. Receives the SA containing the session key from IKE.
2. Locates the outbound SA in its database, and inserts the SPI from the SA into the AH or ESP header.
3. Signs the packets (and encrypts them if confidentiality is required).
4. Sends the packets to the IP layer to be forwarded to the destination computer.

If the negotiation failed, the IPSec driver discards the packet.

When an IPSec-secured inbound packet matches a filter in the IP filter list, the IPSec driver:

1. Receives the session key, SA, and SPI from IKE.
2. Locates the inbound SA in its database by destination address and SPI.
3. Checks the signature and, if required, decrypts the packets.
4. Matches the IP packet searches for a matching filter in the filter list to ensure that no traffic, other than what was agreed upon during the negotiation, is received.
5. Sends packets to the TCP/IP driver to pass to the receiving application.

When an unsecured IP packet is received, the IPSec driver searches for a matching filter in the filter list. If a match occurs and the filter action for that filter either requires IP security or blocks the packet, then the packet is discarded.

**Note**

- The IPSec driver matches all inbound unsecured packets with the list of filters that specify IPSec tunnels first, and then matches the packet with all filters that specify end-to-end (transport) filters. The IPSec driver does not filter certain types of IP packets. For more information, see Add, edit, or remove IPSec filters.

# Description of an IPSec Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Description of an IPSec policy**
An IPSec policy consists of:

- General IPSec Policy Settings

  Settings that apply regardless of which rules are configured. These settings determine the name of the policy, its description for administrative purposes, key exchange settings, and key exchange methods.

- IPSec Policy Rules

  One or more IPSec rules that determine which types of traffic IPSec must examine, how traffic is treated, how to authenticate an IPSec peer, and other settings.

After the policies are created, they can be assigned and applied at the domain, site, organizational unit, and local level. For more information, see Creating, modifying, and assigning IPSec policies.

# General IPSec Policy Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**General IPSec policy settings**

The general settings for an IPSec policy consist of the following types:

- Name

  The name for the policy.

- Description

  Optional text that describes the purpose of the IPSec policy. It is recommended that you complete and update this description to provide a summary of the settings and rules for the policy.

- Policy change poll interval

  The number of minutes between consecutive polls for changes in Active Directory-based IPSec policies. This polling does not detect a change in domain or organizational unit membership, or the assigning or unassigning of a new policy. These events are detected when the Winlogon service polls for changes in Group Policy, which occurs by default every 90 minutes.

- Key exchange settings

  The way in which new keys are derived and how often they are renewed.

- Key exchange methods

  The ways in which identities are protected during the key exchange.

The default key exchange settings and methods are configured to work for most IPSec deployments. Unless there are special security requirements, default settings should not have to be changed.

# Key exchange settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key exchange settings**
The strength of the keys which protect the key exchange phase is enhanced by the following features:

**Key lifetimes**
Lifetime settings determine when a new key is generated. Any time a key lifetime is reached, the associated SA is also renegotiated. The process of generating new keys at intervals is called dynamic rekeying or key regeneration. Lifetimes allow you to force the generation of a new key after a specific interval. For example, if the communication takes 100 minutes and you specify the key lifetime as 10 minutes, 10 keys will be generated (one every 10 minutes) during the exchange. Using multiple keys ensures that if an attacker manages to gain the key to one part of a communication, the entire communication is not compromised. Automatic key regeneration is provided by default. You can override the defaults and either specify a master key lifetime in minutes or by the number of session keys, or enable master key perfect forward secrecy (PFS).

Caution should be taken when setting very different key lifetimes, since they also determine the lifetime of the SA. For example, setting a master key lifetime of 8 hours (480 minutes) and a session key lifetime (set within a filter action) of 2 hours might result in a quick mode SA that is in place for almost 2 hours after the main mode SA has expired. This has the potential to occur if the new quick mode SA is generated just before the main mode SA expires.

- **Session key limit**

  Repeated rekeying off of the same master key might eventually compromise the key. For example, if Alice on Computer A sends a message to Bob on Computer B, and then a few minutes later sends another message to Bob, the same master key material can be reused since an SA was recently established with that computer. If you want to limit the number of times this reuse occurs, you can specify a session key limit.

  If you decide to enable master key perfect forward secrecy (PFS), the session key limit will be set to 1. Master key PFS forces key regeneration each time. If you specify both a master key lifetime in minutes, and a session key limit, whichever interval is reached first will start a new key.

- **Master key perfect forward secrecy (PFS)**

  Master key PFS determines how a new session key is generated. Enabling master key PFS ensures that the master key keying material cannot be used to generate more than one session key.

  Master key PFS should be used with caution because it requires reauthentication and might impact performance. It is not required to be enabled on both peers.

For information about how to configure key exchange settings and methods, see Configure key exchange settings and Create key exchange security methods.

# Key exchange methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Key exchange methods**
In addition to key properties, you can set security methods for main mode IKE negotiation. For example, you can specify which algorithms are used for integrity and confidentiality. The same algorithms that are available for quick mode security methods are available for main mode IKE negotiation: MD5 and SHA1 for integrity; DES and 3DES for confidentiality.

**Diffie-Hellman groups**
In addition to setting integrity and confidentiality algorithms, you can specify which Diffie-Hellman groups to use. Diffie-Hellman groups are used to determine the length of the base prime numbers used during the key exchange process. The cryptographic strength of any key derived depends, in part, on the strength of the Diffie-Hellman group upon which the prime numbers are based.

Group 2048 (high) is stronger (more secure) than Group 2 (medium), which is stronger than Group 1 (low). Group 1 provides 768 bits of keying strength, Group 2 provides 1024 bits, and Group 2048 provides 2048 bits. If mismatched groups are specified on each peer, negotiation fails. The group cannot be switched during the negotiation.

The Diffie-Hellman group is configured as part of the main mode key exchange settings. New session keys generated during quick mode are derived from the Diffie-Hellman main mode master key material, unless master key or session key perfect forward secrecy (PFS) is enabled. If either master key or session key PFS is enabled, a new Diffie-Hellman exchange is performed to obtain new master key keying material for each new session key that is required. The difference between master key PFS and session key PFS is that master key PFS requires a reauthentication of the main mode SA in addition to the Diffie-Hellman exchange.

**Important**

- For enhanced security, do not use Diffie-Hellman Group 1. For maximum security, use Group 2048 whenever possible. Use Group 2 when required for interoperability with Windows 2000 and Windows XP.

**Notes**

- Diffie-Hellman Group 2048 is provided only with the Windows Server 2003 family.
- For a standard level of security, it is generally recommended that all of the IKE settings (for example, master key PFS and key lifetime) and security methods remain at their defaults. For more information about IKE settings, see Internet Key Exchange.
- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

For information about how to configure key exchange settings and methods, see Configure key exchange settings and Create key exchange security methods.

# IPSec Policy Rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPSec policy rules**

An IPSec policy consists of one or more rules that determine IPSec behavior. IPSec rules are configured on the **Rules** tab in the properties of an IPSec policy. Each IPSec rule contains the following configuration items:

- Filter list

  A single filter list is selected that contains one or more predefined packet filters that describe the types of traffic to which the configured filter action for this rule is applied. The filter list is configured on the **IP Filter List** tab in the properties of an IPSec rule within an IPSec policy.

- Filter action

  A single filter action is selected that includes the type of action required (Permit, Block, or Negotiate Security) for packets that match the filter list. For the Negotiate Security filter action, the negotiation data contains one or more security methods that are used (in order of preference) during IKE negotiations and other IPSec settings. Each security method determines the security protocol (such as AH or ESP), the specific cryptographic and hashing algorithms, and session key regeneration settings used. The filter action is configured on the **Filter Action** tab in the properties of an IPSec rule within an IPSec policy.

- Authentication methods

  One or more authentication methods are configured (in order of preference) and used for authentication of IPSec peers during main mode negotiations. The available authentication methods are the Kerberos V5 protocol, use of a certificate issued from a specified certification authority, or a preshared key. The negotiation data is configured on the **Authentication Methods** tab in the properties of an IPSec rule within an IPSec policy.

  **Important**

  - The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

- Tunnel endpoint

  Specifies whether the traffic is tunneled and, if it is, the IP address of the tunnel endpoint. For outbound traffic, the tunnel endpoint is the IP address of the IPSec tunnel peer. For inbound traffic, the tunnel endpoint is a local IP address. The tunnel endpoint is configured on the **Tunnel Setting** tab in the properties of an IPSec rule within an IPSec policy. For more information, see Tunnel mode.

- Connection type

  Specifies whether the rule applies to local area network (LAN) connections, dial-up connections, or both. The connection type is configured on the **Connection Type** tab in the properties of an IPSec rule within an IPSec policy.

The rules for a policy are displayed in IP Security Policies in reverse alphabetical order based on the name of the filter list selected for each rule. There is no method for specifying an order in which to apply the rules in a policy. The IPSec driver automatically orders the rules based on the most specific to the least specific filter list. For example, the IPSec driver would apply a rule containing a filter list that specified individual IP addresses and TCP ports before a rule containing a filter list that specified all addresses on a subnet.

**Default response rule**

The default response rule, which can be used for all policies, has the IP filter list of **<Dynamic>** and the filter action of **Default Response** when the list of rules is viewed with the IP Security Policy Management console. The default response rule cannot be deleted, but it can be deactivated. It is activated by default for all policies.

The default response rule is used to ensure that the computer responds to requests for secure communication. If an active policy does not have a rule defined for a computer that is requesting secure communication, then the default response rule is applied and security is negotiated. For example, when Computer A communicates securely with Computer B, and Computer B does not have an inbound filter defined for Computer A, the default response rule is used.

Security methods and authentication methods can be configured for the default response rule. The filter list of **<Dynamic>** indicates that the filter list is not configured, but that filters are created automatically based on the receipt of IKE negotiation packets. The filter action of **Default Response** indicates that the action of the filter (Permit, Block, or Negotiate security) cannot be configured. Negotiate Security will be used. However, you can configure:

- The security methods and their preference order on the **Security Methods** tab.
- The authentication methods and their preference order on the **Authentication Methods** tab.

# Filter list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Filter list**

When traffic matches a source, destination, and type of IP traffic in a filter, security negotiations are initiated. This type of IP packet filtering enables a network administrator to precisely define what IP traffic is secured. Each IP filter list contains one or multiple filters, which define IP addresses and traffic types. One IP filter list can be used for multiple communication scenarios.

IPSec requires both an inbound and outbound filter between the computers specified in the filter list. Inbound filters apply to incoming traffic. Outbound filters apply to traffic leaving a computer towards a destination. For example, if Computer A wants to securely exchange data with Computer B:

- The active IPSec policy on Computer A must have a filter for any outbound packets to Computer B. Source=A and Destination=B.
- The active IPSec policy on Computer A must have a filter for any inbound packets from Computer B. Source=B and Destination=A.

Each peer must also have the reverse filter:

- The active IPSec policy on Computer B must have a filter for any inbound packets from Computer A. Source=A and Destination=B.
- The active IPSec policy on Computer B must have a filter for any outbound packets to Computer A. Source=B and Destination=A.

**Filter settings**

Each filter defines a particular subset of inbound or outbound network traffic that should be secured. You must have a filter to cover any traffic to which the associated rule applies. A filter contains the following settings:

- **The source and destination address of the IP packet.** You can specify any IP address assigned to the IPSec peer, a single IP address, IP addresses by DNS name, or groups of addresses to specify IP subnets.
- **The protocol over which the packet is being transferred.** By default, all protocols in the TCP/IP protocol suite are selected. However, you can specify an individual protocol for this filter to meet special requirements, including custom protocols.
- **The source and destination port for TCP and UDP.** By default, all TCP and UDP ports are selected, but you can select a specific TCP or UDP port.

For information about how to configure IP filter lists, see Add, edit, or remove IP filter lists.

Special filters might be required for certain types of traffic, such as SNMP traffic or security gateways. For more information, see Special IPSec considerations.

# Filter action

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Filter action**

A filter action defines the security requirements for the data transmission. A filter action can be configured to:

- Permit traffic (**Permit**)

  IPSec passes this traffic without modification or the requirement for security. This is appropriate for traffic from specific computers that are not IPSec-capable. Be sure to limit the IP filter list to a minimal scope when using this type of filter action, so you do not let traffic through which should be secured.

- Block traffic (**Block**)

  IPSec silently discards this traffic. Be sure to use an IP filter list that appropriately defines the scope of traffic when using a blocking filter action so that you do not block valid computers from communicating.

- Negotiate IPSec (**Negotiate security**)

  IPSec requires the negotiation of security associations and the sending or receiving of IPSec-secured traffic. After you choose to negotiate IPSec, you can also configure the following:

  - Security methods and their order

    For more information, see IPSec security methods.

  - Allowance of initial incoming unsecured traffic (**Accept unsecured communication, but always respond using IPSec**)

    IPSec allows an incoming packet that matches the configured filter list to be unsecured (not protected by IPSec). However, the outgoing response to the incoming packet must be protected. This setting is useful when you are using the default response rule for clients. When a group of servers are configured with a rule that secures communications with any IP address and accepts unsecured communication, responding with only secured communications, enabling the default response rule on client computers ensures that the clients will respond to the server request to negotiate security. This option should be disabled for secure computers connected to the Internet to prevent denial of service attacks.

  - Enabling of communication with non-IPSec-enabled computers (**Allow unsecured communication with non-IPSec-aware computer**)

    IPSec falls back to unsecured communication, if necessary. Again, you might need to limit the IP filter list to a minimal scope. Otherwise, if negotiation fails for any reason, any communications affected by the rule in which this filter action resides could result in data being sent without security. If you are concerned about unsecured communication, you might consider disabling this setting. However, communication with computers that cannot initiate IPSec, such as legacy systems, might be blocked. This option should be disabled for secure computers connected to the Internet.

  - Generation of session keys from new keying material (**Use session key perfect forward secrecy (PFS)**)

    This determines whether existing master key keying material can be used to derive a new session key. Enabling session key perfect forward secrecy (PFS) ensures that master key keying material cannot be used to derive more than one session key. When session key PFS is enabled, a new Diffie-Hellman key exchange is performed to generate new master key keying material before the new session key is created. Session key PFS does not require main mode reauthentication and uses less resources than master key PFS.

For information about how to configure filter actions, see Add, edit, or remove filter actions.

**IPSec security methods**

Each security method defines the security requirements of any communications to which the associated rule applies. Creating multiple security methods increases the chance that a common method can be found between two computers. The IKE component reads the list of security methods in descending order and sends a list of allowed security methods to the other peer. The first method in common is selected. Typically, the most secure methods are placed at the top of the list and the least secure methods are placed at the bottom of the list.

**Predefined security methods**

The following security methods are predefined:

- Encryption and Integrity

  Uses the ESP protocol to provide data confidentiality (encryption) with the triple Data Encryption Standard (3DES) algorithm, data integrity and authentication with the Secure Hash Algorithm 1 (SHA1) integrity algorithm, and default key lifetimes (100MB, 1 hour). If you require both data and addressing (IP header) protection, you can create a custom security method. If you do not require encryption, use **Integrity only**.

- Integrity only

  Uses the ESP protocol to provide data integrity and authentication with the SHA1 integrity algorithm and default key lifetimes (100MB, 1 hour). In this configuration, ESP does not provide data confidentiality (encryption). This is appropriate when your security plan calls for standard levels of security.

**Custom security methods**

If the predefined **Encryption and Integrity** or **Integrity only** settings do not meet your security requirements, you can specify custom security methods. For example, you can use custom methods when encryption and address integrity, stronger algorithms, or key lifetimes need to be specified. When configuring a custom security method, you can configure the following:

- Security protocols

  Both AH (**Data and address integrity without encryption**) and ESP (**Data integrity and encryption**) can be enabled in a custom security method when you require IP header integrity and data encryption. If you chose to enable both, you do not need to specify an integrity algorithm for ESP. The algorithm that you select for AH provides integrity.

- Integrity algorithm

  - Message Digest 5 (MD5), which uses a 128-bit key.
  - Secure Hash Algorithm 1 (SHA1), which uses a 160-bit key.

- Encryption algorithm

  - 3DES is the most secure of the DES combinations, and somewhat slower in performance. 3DES processes each block three times, using three unique 56-bit keys.
  - DES is used when the higher security and overhead of 3DES is not necessary. DES uses a single 56-bit key.

- Session key settings

  Session key settings determine when a new key is generated, rather than how it is generated. You can specify a lifetime in kilobytes, seconds, or both. For example, if the communication takes 10,000 seconds and you specify the key lifetime as 1000 seconds, 10 keys will be generated to complete the transfer. This ensures that, even if an attacker manages to determine one session key and decipher part of a communication, deciphering the entire communication is not possible. By default, new session keys are generated for every 100 MB of data or every hour. Note that any time a key lifetime is reached, the quick mode SA is also renegotiated in addition to the key refresh or regeneration.

**Note**

- When IPSec provides data confidentiality between a Windows XP or Windows Server 2003 family-based computer (or a Windows 2000-based computer that has the High Encryption Pack or Service Pack 2 or higher installed) and another computer that does not use 3DES, your list of security methods must include DES. Otherwise, the computer that you are trying to communicate with might not be able to reach a common security agreement with your computer. If confidentiality of the data is not required, you can select ESP format with an integrity algorithm only and set the encryption to **<None>**. Alternately, you can select an AH format with an integrity algorithm.

For information about how to configure security methods, see Add, edit, or remove IPSec security methods.

# Authentication methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication methods**
Each rule defines a list of authentication methods. Each authentication method defines the requirements for how identities are verified in communications to which the associated rule applies. The two peers must have at least one common authentication method or communication will fail. Creating multiple authentication methods increases the chance that a common method between two computers can be found.

Only one authentication method can be used between a pair of computers, regardless of how many are configured. If you have multiple rules that apply to the same pair of computers, you must configure the authentication methods list in those rules to enable the pair to use the same method. For example, if a rule between a pair of computers specifies only Kerberos for authentication and filters only TCP data and, in another rule, specifies only certificates for authentication and filters only UDP data, authentication will fail.

**Overview of authentication methods**
For authentication, IPSec allows you to use the Kerberos V5 protocol, certificate-based authentication, or preshared key authentication.

- The Kerberos V5 security protocol is the default authentication technology. This method can be used for any clients running the Kerberos V5 protocol (whether or not the clients are running Windows 2000, Windows XP Professional, or a Windows Server 2003 operating system) that are members of the same or trusted domains.

- A public key certificate should be used in situations that include Internet access, remote access to corporate resources, external business partner communications, or computers that do not run the Kerberos V5 security protocol. This requires that at least one trusted certification authority (CA) and associated certificate have been configured. Computers running Windows 2000, Windows XP, or a Windows Server 2003 operating system support X.509 Version 3 certificates, including computer certificates generated by commercial CAs.

- A preshared secret key can be specified. It is simple to use and does not require the client to run the Kerberos V5 protocol or have a public key certificate. Both parties must manually configure IPSec to use this preshared key.

  **Important**

  - The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing and that you use certificates or Kerberos V5 instead in a production environment.

## Notes

- The ability to authenticate using Kerberos V5 is not supported on computers running Windows XP Home Edition or computers running any other version of Windows 2000, Windows XP, or a Windows Server 2003 operating system that are not a member of an Active Directory domain.

- For information about how to configure authentication methods, see Define IPSec authentication methods.

**IPSec certificate to account mapping**
When you use either Kerberos V5 or certificate authentication, you can set restrictions on which computers are allowed to connect. When you enable IPSec certificate to account mapping, the IKE protocol associates (maps) a computer certificate to a computer account in an Active Directory domain or forest, and then retrieves an access token, which includes the list of user rights assigned to the computer. You can restrict access by configuring Group Policy security settings and assigning either the **Access this computer from the network** user right or the **Deny access to this computer from the network** user right to individual or multiple computers as needed.

When a Windows 2000 or a Windows Server 2003 family certification authority is used for computer certificate autoenrollment, the certificates are automatically stored in Active Directory as a property of the computer account. Alternatively, you can use a non-Microsoft certification authority in which to enroll computers. If you do not use computer certificate autoenrollment, you must manage certificate to account mapping manually.

**Notes**

- If you use the Kerberos V5 protocol for authentication, the IKE protocol verifies access controls, but only if IKE is a responder. Therefore, for access controls to be applied, you must configure the IPSec policy on the client computer to initiate IKE to the server. If you use certificate authentication, access controls are applied if either the server or the client computer initiates IKE to the server.

- For information about configuring security settings in Group Policy, see Security settings overview.

# Tunnel endpoint

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tunnel endpoint**

As discussed in Virtual private networking with IPSec, IPSec can perform layer 3 tunneling for scenarios in which Layer Two Tunneling Protocol (L2TP) cannot be used. If you are using L2TP for remote communications, no tunnel configuration is needed because the client and server VPN components of Windows XP and the Windows Server 2003 family automatically create the appropriate rules to secure L2TP traffic.

To create a layer 3 tunnel using IPSec, use the IP Security Policy Management or Group Policy consoles to configure and enable the following two rules for the appropriate policy:

1. A rule for the outbound traffic for the tunnel.

   The rule for the outbound traffic is configured with both a filter list that describes the traffic to be sent across the tunnel and a tunnel endpoint of an IP address configured on the IPSec tunnel peer (the computer or router on the other side of the tunnel).

2. A rule for the inbound traffic for the tunnel.

   The rule for the inbound traffic is configured with both a filter list that describes the traffic to be received across the tunnel and a tunnel endpoint of a local IP address (the computer or router on this side of the tunnel).

For each rule, you must also specify filter actions, authentication methods, and other settings.

For information about how to configure IPSec tunnel endpoints, see Specify an IPSec tunnel. For more information about IPSec tunnel mode, see Tunnel mode.

For advanced information about configuring IPSec tunnels, see article Q252735, "How to Configure IPSec Tunneling in Windows 2000," in the Microsoft Knowledge Base.

# Connection type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection type**
For each IPSec rule, you can define which connection types on your computer are affected. The connection types apply to all connections in Network Connections on the computer for which you are configuring IPSec policy.

Each rule has one connection type setting:

- **All Network Connections**. The rule applies to communications sent through any of the network connections that you have configured on the computer.
- **Local Area Network (LAN)**. The rule only applies to communications sent through LAN connections that you have configured on the computer.
- **Remote Access**. The rule only applies to communications sent through any virtual private network (VPN) or dial-up connections that you have configured on the computer.

For information about how to configure connection types, see Specify IPSec connection types.

For information about configuring network connections on a computer, see Configuring a network or dial-up connection.

# Using Internet Protocol Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Internet Protocol security**

This section covers:

- Managing IPSec from the command line
- Designing and Planning IPSec Policies
- Viewing IPSec Statistics and Details About Active IPSec Policies

# Managing IPSec from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing IPSec from the command line**
The Netsh commands for Internet Protocol security (IPSec) provide a fully equivalent alternative to the console-based management and diagnostic capabilities provided by the IP Security Policy Management and IP Security Monitor consoles. You can use Netsh commands for IPSec to script IPSec policy creation, display details about IPSec policies, and change the IPSec configuration for troubleshooting. In addition, administering IPSec from the command line is useful when you want to extend the security and manageability of IPSec. For example, you can use Netsh commands for IPSec to enable IPSec driver event logging, set default traffic exemptions, and configure computer startup security.

For a comprehensive **netsh ipsec** command reference (including syntax and parameters), see Netsh commands for Internet Protocol security.

### Notes

- For information about how to use Netsh commands for IPSec to change the IPSec configuration for troubleshooting, see the section "Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family" in IPSec troubleshooting tools.

- You can use the Netsh commands for IPSec to configure IPSec policies only on computers running members of the Windows Server 2003 family. To use the command line to configure IPSec policies on computers running Windows XP, use Ipseccmd.exe. To use the command line to configure IPSec policies on computers running Windows 2000, use Ipsecpol.exe, which is provided with the Windows 2000 Server Resource Kit. For more information, see IPSec troubleshooting tools.

# Designing and Planning IPSec Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Designing and planning IPSec policies**
This section provides information about planning and deploying IPSec in your organization. You should also review IPSec Best practices for general information about IPSec deployment.

For information on specific topics, see:

- Establishing an IPSec security plan
- Default IPSec policies
- Creating, modifying, and assigning IPSec policies
- Special IPSec considerations
- IPSec troubleshooting tools

# Establishing an IPSec security plan

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing an IPSec security plan**

Whether for a large domain or a small workgroup, implementing IPSec means finding a balance between making information easily available to the largest number of users, and protecting sensitive information from unauthorized access.

Finding the proper balance requires:

- Assessing the risk and determining the appropriate level of security for your organization.
- Identifying valuable information.
- Defining security policies that use your risk management criteria and protect the identified information.
- Determining how the policies can best be implemented within the existing organization.
- Ensuring that management and technology requirements are in place.
- Providing all users with both secure and efficient access to the appropriate resources, according to their needs.

Security considerations are also influenced by the way the computer will be used. For example, the required security might differ, depending on whether the computer is a domain controller, Web server, remote access server, file server, database server, or intranet or remote client. The security framework in Windows 2000, Windows XP, and Windows Server 2003 family is designed to fulfill the most stringent security requirements. However, software alone might be less effective without careful planning and assessment, effective security guidelines, enforcement, auditing, and sensible security policy design and assignment.

There is no exact definition of the measures that define standard security. These can vary widely, depending on an organization's policies and infrastructures. The following security levels can be considered a general basis for planning your IPSec deployment:

- Minimal security

  Computers do not exchange sensitive data. IPSec is not active by default. No administrative action to disable IPSec is required.

- Standard security

  Computers, especially file servers, are used to store valuable data. Security must be balanced so it does not become a barrier to users trying to perform their tasks. Windows XP and the Windows Server 2003 family provide example IPSec policies that secure data, but do not necessarily require the highest level of security: **Client (Respond Only)** and **Server (Request Security)**. Use these as a basis for your IPSec policies. These, or similar custom policies, optimize efficiency without compromising security.

- High security

  Computers that contain highly sensitive data are at risk for data theft, accidental or malicious disruption of the system (especially in remote dial-up scenarios), or any public network communications. **Secure Server (Require Security)**, a default policy, requires IPSec protection for all traffic being sent or received (except initial inbound communication) with stronger security methods. Unsecured communication with a non-IPSec-aware computer is not allowed.

**From planning to policy configuration**

After you have identified the types of traffic that need to be secured and the level of security required, you can begin policy configuration. Windows XP and Windows Server 2003 IPSec policy configuration is the translation of your security

requirements to one or more IPSec policies, only one of which can be assigned at each of the following levels: domain, site, organizational unit, or local levels. Each IPSec policy consists of one or more IPSec rules. Each IPSec rule consists of:

- A selected filter list.
- A selected filter action.
- Selected authentication methods.
- A selected connection type.
- A selected tunnel setting.

You can configure IPSec policy in the following ways:

- Create a new policy and define the set of rules for the policy, adding filter lists and filter actions as required.

  In this method, an IPSec policy is created first and then rules are added and configured. Filter lists (specifying traffic types) and filter actions (specifying how the traffic is treated) are added during rule creation.

  For more information, see Add, edit, or remove IPSec policies and Add, edit, or remove IPSec rules.

- Create the set of filter lists and filter actions, and then create the policies and add rules that combine the filter lists with filter actions.

  In this method, the filter lists and the filter actions are configured first. Next, IPSec policies are created and rules are added that combine the appropriate filter list with the appropriate filter action. Additionally, authentication methods, connection types, and tunnel settings are specified.

  For information about how to individually configure IP filter lists and filter actions, see Manage multiple IP filter lists and Manage multiple filter actions.

For either configuration method, after the IPSec policies have been created, they must be assigned. For more information, see Creating, modifying, and assigning IPSec policies.

**Notes**

- To manage Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. To manage local or remote IPSec policies for a computer, you must be a member of the Administrators group on the local or remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. For more information, see Default local groups and Default groups.

- The integration of IPSec policies with Active Directory described here does not apply to computers running Windows XP Home Edition. You cannot administer Active Directory-based IPSec policy from a computer running Windows XP Home Edition because these computers cannot join Active Directory domains.

# Default IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Default IPSec policies**

Windows XP and the Windows Server 2003 family provide a set of predefined IPSec filter lists and filter actions and default policies. This is only intended to provide an example. It is not designed for operational use without modification. A knowledgeable IPSec network administrator or advanced user should design new, custom polices for operational use.

Computers that are connected to the Internet must be carefully protected from network attacks with custom policies. The default policies are defined as examples for intranets because they allow a computer to receive unsecured traffic. The default policies are designed for computers that are members of an Active Directory domain.

**Predefined filter lists**

The following predefined filter lists are provided as examples for use in the default policies:

- **All ICMP Traffic**

  A filter list for all ICMP traffic (IP protocol 1) sent and received between this computer and all other computers.

- **All IP Traffic**

  A filter list for all IP traffic sent and received between this computer and all other computers. By default, all Internet Key Exchange (IKE) traffic is excluded. All other traffic types are matched against IPSec filters, however, and you can configure, block, or permit filter actions specifically for multicast and broadcast traffic (IPSec does not negotiate security associations for multicast and broadcast traffic).

**Predefined filter actions**

The following predefined filter actions are provided as examples for use in the default policies:

- **Permit**

  Traffic is permitted (the security method is set to **Permit**).

- **Request Security (Optional)**

  Traffic is secured (the security method is set to **Negotiate security**), however, the following exceptions apply:

  - Initial incoming unsecured communication is allowed. This behavior is known as inbound passthrough.

  - If security negotiation fails after 3 seconds, unsecured communication to a non-IPSec-aware computer is allowed. This filter action secures traffic to IPSec-aware computers, but allows unsecured communication if it is required. This behavior is known as fallback to clear.

  In order to avoid the potential for denial-of-service attacks, this filter action must not be used on the Internet.

  The following table lists the security methods (in preference order) for the **Request Security (Optional)** predefined filter action.

| Type | AH integrity | ESP confidentiality | ESP integrity | Key lifetimes (KB/sec) |
|---|---|---|---|---|
| Custom | None | 3DES | SHA1 | 100000/900 |
| Custom | None | DES | SHA1 | 100000/900 |
| Custom | SHA1 | None | None | 100000/300 |

| | | | | |
|---|---|---|---|---|
| Custom | MD5 | None | None | 100000/300 |

The security negotiation for traffic using this filter action attempts ESP encryption (first 3DES, and then DES) with SHA1 integrity. If those security methods cannot be negotiated with the peer, AH integrity (first SHA1, and then MD5) only is attempted.

- **Require Security**

  Traffic is secured (the security method is set to **Negotiate security**), however, initial incoming unsecured communication is allowed (inbound passthrough is enabled). Unsecured communication with non-IPSec-aware computers is not allowed (fallback to clear is disabled). For example, an incoming connection request packet is allowed. The computer's reply to the connection request matches the outbound filter, causing a request for security to be sent to the sender of the unsecured traffic. If security negotiation is successful, traffic in both directions is required to be secure. If security negotiation is not successful, outgoing traffic is not allowed. The computer might continue to receive incoming unsecured traffic and will continue to attempt to negotiate security.

  In order to avoid the potential for denial-of-service attacks, this filter action must not be used on the Internet.

  The following table lists the security methods (in preference order) for the **Require Security** predefined filter action.

| Type | AH integrity | ESP confidentiality | ESP integrity | Key lifetimes (KB/sec) |
|---|---|---|---|---|
| Custom | None | 3DES | SHA1 | 100000/900 |
| Custom | None | 3DES | MD5 | 100000/900 |
| Custom | None | DES | SHA1 | 100000/900 |
| Custom | None | DES | MD5 | 100000/900 |

The security negotiation for traffic using this filter action will only attempt to negotiate ESP encryption (first 3DES, and then DES) with differing ESP authentication methods (first SHA1, and then MD5).

The most secure filter action is one that negotiates security and neither receives unsecured traffic (inbound passthrough is disabled) nor allows communication with non-IPSec-aware computers (fallback to clear is disabled). This is known as a lockdown filter action. When a server is using this filter action, clients must initiate secured communication before sending any IP traffic. The IPSec policy for client computers must be configured with additional rules that force the client to negotiate security with the server. A client configured with policy that consists of only the default response rule is not able to communicate with a locked down server.

**Predefined rules**

The only predefined rule that is available on all policies is the default response rule. The default response rule is automatically created in a new policy and cannot be created manually. It is used to respond to security requests when no other rule applies. For example, a client policy that enables a client to respond to security requests but does not secure any traffic by default, would have only the default response rule active. In some cases, it is enabled by default. You can disable the default response rule on the **Rules** tab in properties for a policy. You can also modify the authentication method and security methods of the default response rule for specific needs. For more information, see the section "Default response rule" in IPSec Policy Rules.

**Default policies**

The following are provided as examples for use as default policies:

- **Client (Respond Only)**

  This is an example policy for computers that secure communication on request. For example, intranet clients might not require IPSec except when requested by another computer. This policy enables the computer on which it is active to respond to requests for secured communications. This policy contains the default response rule, which creates dynamic IPSec filters for inbound and outbound traffic based on the requested protocol and port traffic for the communication that is being secured.

  This policy has the following settings:

**First rule (default response rule)**

- IP Filter List: <Dynamic>
- Filter Action: Default Response
- Authentication: Kerberos
- Tunnel Setting: None
- Connection Type: All

- **Server (Request Security)**

  This is an example policy for computers that should secure communications in most cases, allowing unsecured communication with non-IPSec-aware computers. With this policy, the computer accepts unsecured traffic, but always attempts to secure additional communications by requesting security from the original sender. This policy allows the entire communication to be unsecured if the other computer is not IPSec-enabled. For example, communication to specific servers can be secure while allowing the server to communicate in an unsecured manner to accommodate a mixture of clients (some of which support IPSec and some of which do not).

  This policy has a rule to request security for all IP traffic, a rule to permit ICMP traffic, and the default response rule to respond to requests for security from other computers.

  This policy has the following settings:

  **First rule**

  - IP Filter List: All IP Traffic
  - Filter Action: Request Security (Optional)
  - Authentication: Kerberos
  - Tunnel Setting: None
  - Connection Type: All

  **Second rule**

  - IP Filter List: All ICMP Traffic
  - Filter Action: Permit
  - Authentication: N/A
  - Tunnel Setting: None
  - Connection Type: All

  **Third rule (default response rule)**

  - IP Filter List: <Dynamic>
  - Filter Action: Default Response

- Authentication: Kerberos
- Tunnel Setting: None
- Connection Type: All

- **Secure Server (Require Security)**

  This is an example policy for computers on an intranet that require secure communications, such as a server that transmits highly sensitive data. Administrators can use this IPSec policy as an example to create their own custom IPSec policy for production use. The filters used in this policy require all outbound communication to be secured, allowing only the initial inbound communication request to be unsecured.

  This policy has a rule to require security for all IP traffic, a rule to permit ICMP traffic, and the default response rule to respond to requests for security from other computers.

  **First rule**

  - IP Filter List: All IP Traffic
  - Filter Action: Require Security
  - Authentication: N/A
  - Tunnel Setting: None
  - Connection Type: All

  **Second rule**

  - IP Filter List: All ICMP Traffic
  - Filter Action: Permit
  - Authentication: Kerberos
  - Tunnel Setting: None
  - Connection Type: All

  **Third rule (default response rule)**

  - IP Filter List: <Dynamic>
  - Filter Action: Default Response
  - Authentication: Kerberos
  - Tunnel Setting: None
  - Connection Type: All

**Notes**

- In the Windows Server 2003 family, IPSec provides new source and destination address options for defining filters. In addition, the default behavior for traffic exemptions has changed. If you require previously exempted traffic types, you

must create filters to permit this traffic. For more information about the new source and destination options provided by IPSec and for instructions on how to configure permit filters for several previously exempted traffic types, see Special IPSec considerations.

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

# Creating, modifying, and assigning IPSec policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating, modifying, and assigning IPSec policies**

IPSec policies can be applied to local computers, domains, sites, organizational units, or any Group Policy object in Active Directory. You can create, modify, and assign IPSec policies by using the IP Security Policy Management console that is available in Microsoft Management Console (MMC) or by using the Netsh commands for IPSec that are available at the command line.

Your organization's IPSec policies should be based on your written guidelines for secure operations. Policies can store multiple rules, so that one policy can be used to affect multiple types of traffic.

The Windows 2000, Windows XP, and Windows Server 2003 family include example IPSec policies. These examples are provided only to show how policies can be designed. They are not designed as default policies for securing a client or a server in a production environment. Instead, you should design and create custom IPSec policies that apply to the specific scenarios in your security plan.

### Notes

- For information about how to use the IP Security Policy Management console, see Define IPSec Policies.
- For information about how to use Netsh, see Netsh commands for Internet Protocol security.
- Although you can assign local IPSec policies for a computer by using the IP Security Policy Management console, you must use Group Policy to assign IPSec policies for groups of computers in Active Directory domains. For more information, see Assign or unassign IPSec policy in Group Policy.

**IPSec driver modes**

To understand how IPSec policies are processed and applied, it is important to understand the modes in which the IPSec driver operates. The IPSec driver operates in computer startup mode and operational mode and, optionally, in a diagnostics mode for troubleshooting.

### Computer startup mode

During computer startup, the Windows operating system loads the IPSec driver to perform in computer startup mode. Computer startup mode is used until the IPSec Policy Agent sets the IPSec driver into an operational mode. The IPSec driver can perform in any one of the following computer startup modes:

- **Permit.** When the IPSec driver operates in permit mode, it does not process any IP packets, and, therefore, no IPSec security is provided. If an IPSec policy has never been assigned to a computer, permit mode is the default mode for the IPSec driver.
- **Stateful.** When the IPSec driver operates in stateful mode, it allows all outbound traffic, and it automatically creates inbound permit filters in response to the outbound traffic. All other inbound unicast, broadcast, and multicast packets are dropped. If an IPSec policy has been assigned to a computer, the IPSec Policy Agent sets the stateful mode for the IPSec driver by default.
- **Block.** When the IPSec driver operates in block mode, it discards all inbound and outbound IP packets, except for those that match specific filters that you can configure to be used during block mode. By default, all inbound and outbound DHCP traffic is permitted so that an IP address can be obtained. You can configure filters for block mode by using the **netsh ipsec dynamic set config bootexemptions** command.

To change the default mode for the IPSec driver during computer startup, use the **netsh ipsec static set config bootmode** command.

### Operational mode

After the IPSec service starts, the IPSec Policy Agent can set the IPSec driver to any one of the following operational modes. The operational mode that is chosen does not change the mode that the IPSec driver uses during computer startup.

- **Secure.** When the IPSec driver operates in secure mode, IPSec policy filters are enforced for standard IPSec operation. The IPSec Policy Agent sets the driver into secure mode after it applies persistent policy (if a persistent policy has been configured) and before it applies the Active Directory-based policy or the local policy. If no persistent policy has been configured and assigned, secure mode does not provide IPSec protection until either the assigned Active Directory-based policy or the local policy is applied. When the Active Directory-based policy or the local policy changes, the IPSec Policy Agent updates the IPSec filter list in the IPSec driver. If no IPSec policy is assigned, the IPSec Policy Agent deletes the filters from the IPSec driver and no IPSec protection is provided in this mode.

- **Permit.** When the IPSec driver operates in permit mode, it does not process any IP packets, and, therefore, no IPSec protection is provided. Permit mode is used when the IPSec service is manually stopped.

- **Block.** When the IPSec driver operates in this mode, the exemptions that might apply during computer startup are not applied. Instead, all inbound and outbound traffic is blocked. This mode is used to enhance security, in the event that the IPSec Policy Agent fails to apply a persistent policy, if a persistent policy has been configured.

The operational mode can be configured only by the IPSec Policy Agent. You cannotuse the Netsh commands for IPSec to configure this mode.

### Diagnostic mode

To record all inbound and outbound dropped packets and other packet processing errors in the Event Viewer System log for troubleshooting, you can enable IPSec driver event logging and specify the level of logging to provide. The types of packet processing errors that the IPSec driver records in the System log depend on the level of logging that is provided. IPSec driver logs can record inbound and outbound per-packet drop events during computer startup mode and operational mode.

IPSec driver event logging is disabled by default, and it should not be used for extended periods. Depending on the logging level that you set, many events might be generated that fill the System log very quickly. For more information, see the section "Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family" in IPSec troubleshooting tools.

#### Application methods for IPSec policy

IPSec policies are provided in six ways, in the order described in the following sections.

#### IPSec startup policy

When a computer loads the Windows operating system during startup, the IPSec driveris loaded at the same time as the TCP/IP driver. The following table describes thebehavior of the IPSec driver based on the startup type of the IPSec service.

| IPSec service startup type | IPSec driver behavior |
|---|---|
| Disabled | The IPSec driver loads in permit mode, meaning that no IPSec security is applied. In this mode, the IPSec driver does not filter any traffic. |
| Manual | The IPSec driver loads in permit mode, meaning that no IPSec security is applied. In this mode, the IPSec driver does not filter any traffic. |
| Automatic | The IPSec driver loads in a startup mode configured by the IPSec Policy Agent. If no IPSec policy is assigned, by default the IPSec driver loads in permit mode. If IPSec policy is assigned, the IPSec Policy Agent configures the IPSec driver to load in stateful mode.<br><br>After the IPSec service starts, persistent policy (if it has been configured) is applied, and the IPSec driver is set to secure mode (normal operation). |

When the IPSec driver is loaded, it reads the registry to determine the actions that it must perform during computer startup. These actions include:

- Whether to apply stateful filtering, permit all traffic, or block all traffic.

- Whether to exempt specific traffic types from filtering during computer startup.
- Whether to enable IPSec diagnostic logging and, if so, which level of logging to provide.
- Whether to modify the default IPSec traffic exemption. By default, in the Windows Server 2003 family, RSVP, Kerberos, multicast, and broadcast traffic are not exempt from IPSec filtering. Only IKE traffic is exempt.

The IPSec driver remains in a startup mode until the IPSec Policy Agent sets it to an operational mode. After the IPSec driver is set to secure operation mode, the startup stateful mode settings are discarded. You can use the Netsh commands for IPSec to configure IPSec driver settings. For more information, see the **netsh ipsec dynamic set config** section in Netsh commands for Internet Protocol security.

**Persistent policy**

In the Windows Server 2003 family, you can use the Netsh commands for IPSec to add, modify, and maintain IPSec policy settings permanently, after you run the appropriate commands. This type of policy is known as a persistent policy. You can configure persistent policies to extend existing Active Directory-based or local IPSec policies, override Active Directory-based or local IPSec policies, and enhance security during computer startup. Persistent policies enhance security by providing a secure transition from computer startup to Active Directory-based IPsec policy enforcement.

Persistent policies provide the following two key benefits:

- You can configure a persistent policy as your most restrictive IPSec policy. For example, the persistent policy can block all inbound traffic to the Internet IP address of a computer, except for ICMP traffic and communication with a particular management station. To complement the persistent policy, you can design an Active Directory-based policy or a local policy with additional rules that include more specific filters to allow additional communication types.
- You can selectively enforce security rules that are required only for specific servers, due to unique network adapter configurations or program needs. For example, you can enforce security rules specifying that:
    - A specific port always be blocked.
    - Traffic that uses a specific network adapter be permitted without IPSec security, or only with IPSec security.
    - Unsecured traffic to a specific monitoring station be permitted.

Persistent policies, if configured, are stored in the local registry. The persistent store is loaded by the IPSec Policy Agent during service startup. After a persistent policy is successfully applied, the IPSec Policy Agent sets the IPSec driver to secure mode. If an error occurs when a persistent policy is applied, the IPSec Policy Agent removes the persistent policy settings from the IPSec driver and sets the IPSec driver to block operation mode for security. You can update a persistent policy at any time, as long as the IPSec service is running. However, changes in persistent policy are not active immediately. You must restart the IPSec service to load the new persistent policy settings.

If you have configured Active Directory-based policies, you can use a persistent policy as a tool to require that traffic to Active Directory always be secured by IPSec, including the retrieval of Active Directory-based IPSec policies. When an Active Directory-based or local policy is applied, those policy settings are added to the persistent policy settings.

To configure persistent policies, you must use the **netsh ipsec static set store location=persistent** command. To view the IPSec policy that is currently being applied, you can use either IP Security Monitor or the **netsh ipsec dynamic show** commands. For information about how to use Netsh, see Netsh commands for Internet Protocol security.

**Important**

- To provide maximum protection against attacks during computer startup, it is highly recommended that you configure and assign a persistent policy. If you do not configure a persistent policy, no policy can be enforced by the IPSec driver until the IPSec Policy Agent retrieves and applies the assigned Active Directory-based or local IPSec policy.

**Active Directory-based policy**

A Group Policy object defines access, configuration, and usage settings for accounts and resources. IPSec policies can be assigned to the Group Policy object of a site, domain, or organizational unit. When the IPSec policy is applied to one of the

Group Policy objects for the Active Directory object, the IPSec policy is propagated to all of the computer accounts that are affected by that Group Policy object.

You can manage Active Directory-based policy by using the IP Security Policy Management console or by using the **netsh ipsec static set store location=domain** command. IPSec policies that are configured and assigned for the domain take precedence over the local, active IPSec policy on a computer, when that computer is a member of the domain. Active Directory-based policy overrides any local IPSec policy that is assigned, and it adds to the persistent IPSec policy that has already been applied by the IPSec Policy Agent, if a persistent policy has been configured.

**Important**

- Several features in the Windows Server 2003 family implementation of IPSec are not provided in Windows 2000 or in Windows XP. To ensure that the same IPSec policy functions as expected on computers running the Windows Server 2003 family and on computers running Windows 2000 or Windows XP, test the policy thoroughly on all relevant operating systems before deployment. For more information about new features in IPSec, see New features for IPSec.

- If you plan to apply IPSec policies that use the new features that are available only in the Windows Server 2003 family implementation of IPSec, do not use the Windows 2000 or the Windows XP version of the IP Security Policy Management console to manage these policies. The settings in the earlier versions of the IP Security Policy Management console will override the settings in the Windows Server 2003 family IPSec policy, and the new features will not be functional.

When assigning an IPSec policy in Active Directory, consider the following:

- The list of all IPSec policies is available to assign at any level in the Active Directory hierarchy. However, only a single IPSec policy can be assigned at a specific level in Active Directory.

- An IPSec policy that is assigned to an organizational unit in Active Directory takes precedence over a domain-level policy for members of that organizational unit.

- An IPSec policy that is assigned to the lowest-level organizational unit in the domain hierarchy overrides an IPSec policy that is assigned to a higher-level organizational unit, for member computers of that organizational unit.

- An organizational unit inherits the policy of its parent organizational unit unless either policy inheritance is explicitly blocked or policy is explicitly assigned.

- IPSec policies from different organizational units are never merged.

- The highest possible level of the Active Directory hierarchy should be used to assign policies to reduce the amount of configuration and administration required.

- An IPSec policy might remain active even after the Group Policy object to which it is assigned has been deleted. Because of this, you should unassign the IPSec policy before you delete the policy object. To prevent problems, use the following procedure:

  1. Unassign the IPSec policy in the Group Policy object.

  2. Wait 24 hours to ensure that the change is propagated.

  3. Delete the Group Policy object.

     If you delete the Group Policy object without following this procedure, computers in the Active Directory container to which the IPSec policy is assigned treat the IPSec policy as if it cannot be located and continue to use a cached copy.

- You can use Group Policy backup and restore tools to store information about which policy is assigned to a Group Policy

object. To ensure consistency, when you use these tools, make sure that you include IPSec policies. Because these tools do not store the IPSec policies themselves, however, you cannot use them to back up and restore the IPSec policies themselves. You can only back up and restore information about whether the IPSec policies are assigned to Group Policy objects and, if so, to which Group Policy objects. To back up and restore all of the IPSec policies in the IPSec container in Active Directory, use the Export Policies and Import Policies commands in the IP Security Policy Management console. For more information, see Export IPSec policies and Import IPSec policies.

- Before assigning an IPSec policy to a Group Policy object, verify the Group Policy settings that are required for the IPSec policy. For example, if an IPSec policy requires certificate authentication, assign the Group Policy settings that allow computers to enroll for certificates (usually one or two days before you assign the IPSec policy that requires use of the computer certificate). In addition, you should test the certificate enrollment process and resolve any errors before assigning the IPSec policy.

The IPSec Policy Agent on a computer running Windows XP Professional or a Windows Server 2003 operating system polls Active Directory for updates to the assigned IPSec policy. This polling does not detect a change in domain or organizational unit membership or the assigning or unassigning of a new IPSec policy. These events are detected when the Winlogon service polls for changes in Group Policy, which occurs by default every 90 minutes. The Winlogon service discovers these changes, notifies the IPSec Policy Agent, and the IPSec policy changes are applied.

For more information about Group Policy, see Group Policy overview.

### Notes

- The integration of IPSec policies with Active Directory described here does not apply to computers running Windows XP Home Edition because you cannot administer Active Directory-based IPSec policy from a computer running Windows XP Home Edition.

- Careful planning is required when you design Active Directory-based IPSec policies for computers that are used the internal network and on other public or partner networks.

- By default, IPSec policies are read from Active Directory when the connection to Active Directory is interpreted by Group Policy to be a slow link.

#### Directory cache policy

If you assign an Active Directory-based IPSec policy, a copy of the current policy is maintained in a cache in the local registry. If the computer to which the policy is assigned cannot connect to the domain, the cached copy of the Active Directory-based policy is applied. This cached policy adds to the settings of the persistent policy, if a persistent policy has been configured. Note that you cannot configure or manage the cached copy of an Active Directory-based IPsec policy.

#### Local computer policy

Each computer running Windows XP or a Windows Server 2003 operating system has exactly one local Group Policy object, often called the local computer policy. In using this local Group Policy object, Group Policy settings can be stored on individual computers regardless of whether they are members of an Active Directory domain. The local Group Policy object can be overwritten by Group Policy objects associated with sites, domains, or organizational units in an Active Directory environment. On a network without an Active Directory domain (a network that lacks a Windows 2000 or a Windows Server 2003 domain controller), the local Group Policy object settings determine IPSec behavior because they are not overwritten by other Group Policy objects.

You can manage local IPSec policies by using the IP Security Policy Management console or by using the **netsh ipsec static set store location=local** command. The settings of a local IPSec policy are added to the persistent policy, if a persistent policy has been configured. If an Active Directory-based IPSec policy is assigned, and if the computer is connected to an Active Directory domain, the settings of the Active Directory-based policy are applied instead.

For information about how to assign IPSec policy locally, see Assign or unassign IPSec policy on a computer.

#### Netsh dynamic mode policy

You can use the **netsh ipsec dynamic** commands to create, modify, and assign IPSec rules that take effect immediately and are not stored. If the IPSec service is stopped, either administratively or when the computer is restarted, these rules are lost. However, if you use the **netsh ipsec dynamic set config** commands to change the IPSec configuration, those settings are not lost.

# Special IPSec considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Special IPSec considerations**
The following special IPSec considerations can help simplify the administration of IPSec policies.

**Recommended IPSec scenarios and IPSec scenarios that are not recommended**
Following are recommended scenarios and scenarios that are not recommended for the Windows Server 2003 family implementation of IPSec.

**Recommended IPSec scenarios**
The Windows Server 2003 family implementation of IPSec is recommended for use as follows:

- Packet filtering. IPSec provides limited firewall capabilities for end-systems. You can also use IPSec with Internet Connection Firewall, Windows Firewall, and Routing and Remote Access to permit or block inbound or outbound traffic.

- Securing host-to-host traffic on specific paths. You can use IPSec to provide mutual authentication and cryptographic protection for traffic between servers or other static IP addresses or subnets. For example, IPSec can secure traffic between domain controller sites or forests, or between Web servers and database servers.

- Securing traffic to servers. You can use IPSec to require mutual authentication for all client computers that access a server. In addition, you can set restrictions on which computers are allowed to connect to a server running Windows Server 2003 family.

- Using L2TP/IPSec tunneling for VPN connections. You can use the combination of the Layer Two Tunneling Protocol (L2TP) and IPSec (L2TP/IPSec) for all virtual private network VPN scenarios.

- Using IPSec in tunnel mode for gateway-to-gateway tunnels. For interoperability with other routers, gateways, or end-systems that do not support L2TP/IPSec or PPTP VPN tunneling, you can use IPSec in tunnel mode for gateway-to-gateway tunnels.

### Notes

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**IPSec scenarios that are not recommended**
When using IPSec authentication and data protection features, the policy management model in Windows 2000, Windows XP, and Windows Server 2003 family lends itself best to server-to-server and client-to-server scenarios where one end-point has a static address. In large network deployments, when policy involves dynamic addresses at both end–systems, and in some mobility cases, policy management complexity can make IPSec deployment difficult. Because of this, IPSec is not recommended for use as follows:

- Securing communication between domain members and their domain controllers. Because of the complexity of the IPSec policy configuration and management required for this scenario, using IPSec to secure communication between domain members and their domain controllers is not recommended.

- Securing all traffic in a network. Configuring IPSec communications between all client computers and all servers in order to secure many or all computers in a network is difficult. IPSec cannot negotiate security for multicast and broadcast traffic. Traffic from real-time communications, applications that require ICMP, and peer-to-peer applications might be incompatible with IPSec. For these reasons, using IPSec to secure all traffic in a network is not recommended.

Additionally, the IPSec protocol and implementation has characteristics that require special consideration in other scenarios:

- Securing traffic over wireless 802.11 networks. IPSec policies are not optimized for configuration for mobile clients. Because of this, IPSec is not recommended as the only method to secure traffic sent over wireless 802.11 networks. Instead, it is recommended that you use IEEE 802.1X authentication. 802.1X enhances security and ease of deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting. In cases where clients roam between access points on the same network, IPSec can be used in combination with 802.11 and 802.1x. In cases where roaming causes the client IP address to change, IPSec security associations become not valid and new security associations are renegotiated.

- Using IPSec in tunnel mode for remote access VPN connections. Using IPSec in tunnel mode is not recommended for remote access VPN scenarios. Instead, use L2TP/IPSec or PPTP for remote access VPN connections.

**Managing policies that use new features available only in the Windows Server 2003 family**
The following features are available only in the Windows Server 2003 family implementation of IPSec:

- Support for new source and destination address options for defining filters. For example, you can now create:

    - Filters with a source address and destination address of **Any IP Address**.

    - Filters that use a non-standard subnet class as a source address or a destination address.

    - Filters that use a multicast or a broadcast address as a destination address.

    - Filters with a source address or a destination address that corresponds to the local IP configuration for the DNS servers, WINS servers, and the DHCP server and default gateway for the computer to which the filter applies.

- The ability to exclude the name of the certification authority (CA) from certificate requests.

- Certificate-to-account mapping for network access control.

- Support for root CA names that are represented in extended ASCII character sets.

- A stronger Diffie-Hellman group (2048-bit Diffie-Hellman key exchange).

Computers running Windows XP and Windows 2000 cannot use these new policy settings. If you plan to apply Active Directory-based IPSec policies that use features that are only available in the Windows Server 2003 family, use the Windows Server 2003 family version of the IP Security Policy Management console to manage them. If you use the Windows XP or Windows 2000 version of the IP Security Policy Management console to manage these policies, the earlier versions of the IP Security Policy Management console will delete the new feature settings.

In addition, if you want to apply the same IPSec policy to computers running different versions of Windows, test the policy to ensure that it functions as expected on both servers running Windows Server 2003 and computers running Windows XP or Windows 2000.

For information about new features available with the Windows Server 2003 family implementation of IPSec, see New features for IPSec. For more information about best practices for using IPSec, see IPSec Best practices.

**IP filters**
Consider the following when configuring IP filters:

- Use general filters when you want a single filter to apply to a group of computers. For example, when configuring the filter, use **My IP address**, **Any IP Address**, or a subnet address instead of specifying a specific computer's source and destination IP addresses.

- Define filters that allow you to group and secure traffic from logically associated segments or subnets within your network.

- Be aware that when you are viewing the IPSec policy, the order in which the filters apply is not in the order displayed. When the IPSec Policy Agent reads an IPSec policy, filters are processed into one ordered list, sorted from most to least

specific. You can use the IP Security Monitor console to view filters sorted by weight. Filters that have the same weight might appear in any order relative to each other.

During policy updates, the IPSec Policy Agent recalculates the sorted order and updates the IPSec driver with the changes only. If a filter is deleted, or if the filter action is changed for a filter, all security associations related to that filter are deleted. As a result, some packets are lost when a new security association is negotiated. TCP connections, however, should immediately recover by retransmitting the lost packets.

- If you are configuring an IPSec policy to secure Kerberos traffic, IPSec can only negotiate security associations if your authentication method does not use Kerberos. If Kerberos is required for authentication between domain members, special consideration must be given to the traffic required by Kerberos between domain members and domain controllers, and the following requirements must be met:

  - If a domain member and a domain controller are members of different domains, a two-way trust must exist between the domains or Internet Key Exchange (IKE) rekeys will fail.

  - If a client computer or a server in the domain uses Kerberos authentication to negotiate IPSec with all peers, you must create mirrored permit filters for all Kerberos-dependent traffic to all domain controllers. It is recommended that you create permit filters to exempt all traffic to each domain controller IP address in the domain member's domain. You do not need to create permit filters to exempt all traffic to all trusted domain controller IP addresses. Instead, create permit filters to exempt only Kerberos traffic and domain controller site location (LDAP) queries to all domain controller IP addresses in the trust path between the computers. If DNS service is not provided by the domain controller (for which you already permit all traffic), you might also need to exempt DNS traffic to the DNS server IP addresses so that DNS can be used for domain controller discovery.

    To exempt these traffic types in the IPSec policy, create the following filters, each with a permit filter action:

| **Filter** | **Filter Action** |
|---|---|
| All traffic to my domain controllers | • Mirrored<br>• Source Address = My IP Address<br>• Destination Address = **A specific IP Address** (where **A specific IP Address** is the static IP address of each domain controller in the domain)<br>• Protocol = **Any**<br>• This filter permits LDAP site location queries and ICMP, DNS, and Kerberos traffic. In addition, it permits the remote procedure call (RPC) traffic and Group Policy downloads required to manage Active Directory-based IPSec policies. |
| Kerberos | • Mirrored<br>• Source Address = My IP Address<br>• Destination Address = **A specific IP Address** (where **A specific IP Address** is the static IP address of each domain controller in the domain)<br>• Protocol = **Any**<br>• This filter permits LDAP site location queries and ICMP, DNS, and Kerberos traffic. In addition, it permits the remote procedure call (RPC) traffic and Group Policy downloads required to manage Active Directory-based IPSec policies. |

| | |
|---|---|
| Domain controller site location (LDAP) queries | • Mirrored<br>• Source Address = **My IP Address**<br>• Destination Address = **A specific IP Address** (where **A specific IP Address** is the static IP address of each domain controller in trusted domains)<br>• Protocol = **UDP**<br>• Source port = **Any**<br>• Destination port = **88**<br>• After you create this filter, create another filter with the same settings, but specify **TCP** for the protocol. |
| DNS client queries | • Mirrored<br>• Source Address = **My IP Address**<br>• Destination Address = **DNS Servers <dynamic>** (or **A specific IP Address**, where the IP address is the static IP address of each DNS server in the network)<br>• Protocol = **UDP**<br>• Source port = **Any**<br>• Destination port = **53**<br>• After you create this filter, create another filter with the same settings, but specify **TCP** for the protocol. |

**Caution**

- To minimize the security issues, you should allow these exemptions only for the specific IP addresses of the domain controllers and the DNS servers.

When designing an IPSec policy to secure or block ICMP traffic, ensure that you plan carefully and test the policy before implementing it in a production environment. Be aware that ICMP is required for the following scenarios:

- Outbound TCP connections that traverse networks with different link frame sizes require path maximum transmission unit (PMTU) discovery. For PMTU discovery to function correctly, inbound ICMP traffic from routers and other gateways must be received.

- The **tracert** command requires ICMP traffic to determine the path taken to a destination. If IPSec protects ICMP traffic, however, **tracert** output shows only the destination computer--not the path of intermediate routers.

- File sharing and other applications use ICMP traffic to determine whether a destination IP address can be reached.

- Network logon and Group Policy measure ICMP response times to each domain controller IP address. This measurement is performed to both aid in the selection of a domain controller to service a network logon and determine whether to apply Group Policy settings over slow links.

In some cases, even the three-second delay before a soft SA (an unsecured connection) is established might impact upper-

layer communication that requires fast ICMP responses.

- In Windows 2000 and Windows XP, by default, broadcast, multicast, Kerberos, Internet Key Exchange (IKE), and Resource Reservation Protocol (RSVP) traffic are exempt from filter matches. In the Windows Server 2003 family, broadcast, multicast, and Kerberos, are not exempt from filter matches by default (only IKE traffic is exempt). In additon, Resource Reservation Protocol (RSVP) has been turned off. Broadcast and multicast packets will be dropped if they match a filter with a filter action to negotiate security. By default, the Windows Server 2003 family provides limited support for filtering broadcast and multicast traffic. A filter with a source address of **Any IP Address** will match multicast and broadcast addresses. A filter with a source address of **Any IP Address** and a destination address of **Any IP Address** will match inbound and outbound multicast addresses. You can use such a filter to block all traffic. One-way filters that would be used to block or permit specific multicast or broadcast traffic, however, are not supported.

  As a result of the change in default exemption behavior for the Windows Server 2003 family implementation of IPSec, you should verify the behavior of IPSec policies designed for Windows 2000 or Windows XP, and determine whether to configure explicit permit filters to permit specific traffic types. To restore the default Windows 2000 and Windows XP behavior for IPSec policies, you can use the **netsh ipsec dynamic set config command**, or you can modify the registry. For more information, see the section "Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family" in IPSec troubleshooting tools.

  **Caution**

  - The Windows 2000 and Windows XP default exemption settings for IPSec are designed for corporate LAN environments with a low risk of attack. For this reason, you should only use the Windows 2000 and Windows XP default exemption settings when necessary for troubleshooting, in low-risk environments, or when you cannot solve program compatibility issues by configuring explicit filters in IPSec policies.

**Safe Mode with Networking**

If you start a computer in Safe Mode with Networking, consider the following:

- The IPSec service cannot apply filters from an IPSec policy or negotiate IP security when the computer is in Safe Mode with Networking. As a result, when Directory Services Restore Mode operations are in progress, communication does not occur with remote domain controllers or network backup devices that enforce an IPSec policy to require IPSec.

- When a computer loads the Windows operating system during startup, the IPSec driver is loaded at the same time as the TCP/IP driver. If an IPSec policy is assigned to the computer and the IPSec service is configured for automatic startup, by default, the IPSec driver startup mode is changed to perform stateful filtering. When stateful filtering is performed, all outbound traffic is permitted, but all inbound traffic is blocked, except for DHCP traffic and traffic that is sent in response to the outbound traffic initiated during computer startup. (The IPSec driver automatically creates inbound filters to permit responses to outbound traffic.) The IPSec service itself, however, does not start. As a result, no persistent, local, or Active Directory-based IPSec policies can be applied.

- When the computer is in Safe Mode with Networking, you can allow the IPSec driver to continue to provide stateful filtering protection by configuring inbound exemptions for specific traffic. Or, you can force the IPSec driver to permit all traffic.

- To configure inbound exemptions for specific traffic, restart the computer normally (not in Safe Mode), use the **netsh ipsec dynamic set config bootexemptions** command to specify the traffic types to exempt during computer startup, and then restart the computer. The new exemptions will take effect when the computer is restarted and will remain in effect until they are deleted or changed.

- If you allow the IPSec driver to configure inbound exemptions for specific traffic, keep in mind that unless you configure the appropriate exemptions, the following will occur:

  - Inbound connections will be blocked to network services, such as Remote Desktop Connection and Remote Assistance, which open ports for receiving connections. To allow clients to use Remote Desktop Connection or Remote Assistance, you must configure an exemption to allow inbound traffic to TCP port 3389. To configure this exemption, use the following Netsh command:

**netsh ipsec dynamic set config bootexemptions TCP:any:3389:inbound**

- Some outbound TCP connections, such as those requiring PMTU discovery, might not work. To enable TCP PMTU discovery, you must configure an exemption to allow inbound ICMP traffic. To configure this exemption, use the following Netsh command:

  **netsh ipsec dynamic set config bootexemptions ICMP:inbound**

  **Important**

- Modifying IPSec traffic exemptions from startup security (that is, modifying the **bootexemptions=** parameter) will overwrite all previous exemptions from startup security.

For more information about how to specify traffic types to exempt during computer startup, see Netsh commands for Internet Protocol security.

To force the IPsec driver to permit all inbound traffic while the computer is in Safe Mode with Networking, you must either configure the IPSec service for manual startup or disable it, and then restart the computer in Safe Mode with Networking.

**Filter actions**
Consider the following special requirements when configuring IP filter actions:

- If there is specific traffic that you always want to block, such as traffic to and from either not valid addresses or specific TCP ports, create a blocking filter action with the **Security Method** set to **Block**.

- If there is specific traffic that you always want to be unsecured, create a permit filter action with the **Security Method** set to **Permit**.

- When you are configuring custom security methods for encrypted traffic, you should only set the ESP confidentiality selection to **<None>** when a higher layer protocol provides data encryption.

- When you are planning Internet scenarios, such as IPSec tunneling between branch offices, consider using a list of security methods that specifies high levels of security (for example, 3DES only), short key lifetimes (less than 50 MB), and perfect forward secrecy (PFS) for the session keys. This helps protect against known key attacks.

**Note**

- Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use the 3DES algorithm. If a computer running Windows 2000 receives a 3DES setting, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the 3DES setting in the security method is set to the weaker DES, to provide some level of confidentiality for communication, rather than blocking all communication. However, you should only use DES as a fallback option if not all computers in your environment support the use of 3DES. Computers running Windows XP or a Windows Server 2003 operating system support 3DES and do not require installation of the High Encryption Pack.

**VPN connections using L2TP**
When you are using Layer Two Tunneling Protocol (L2TP) for your remote access or router-to-router VPN connections, certificates are the default authentication method for main mode security associations. Computer certificates must be installed either on both the VPN server and remote access client (for remote access connections), or on both the calling router and answering router (for router-to-router VPN connections). The simplest method for installing computer certificates is to configure your Windows 2000 and Windows Server 2003 domains to auto-enroll computer certificates for domain members. For more information, see Computer certificates for L2TP/IPSec VPN connections.

When you are using the certificate authentication method for L2TP connections, the list of certification authorities (CAs) is not configurable. Instead, each computer in the L2TP connection sends a list of root CAs to its IPSec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Computer A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPSec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPSec peer, Computer B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPSec

main mode negotiation fails. Therefore, ensure one of the following before attempting an L2TP connection:

- Both the VPN client and VPN server were issued computer certificates from the same CA.
- Both the VPN client and VPN server were issued computer certificates from CAs that follow a certificate chain up to the same root CA.

In general, the VPN client must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN server trusts. Additionally, the VPN server must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN client trusts.

For example, if the VPN client was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies the VPN server during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the VPN server does not have a valid computer certificate issued from a CA that follows a certificate chain to either CertAuth1 or CertAuth2, IPSec main mode negotiation fails.

A single CA commonly issues computer certificates to all computers in an organization. Because of this, all computers within the organization both have computer certificates from a single CA and request certificates for authentication from the same single CA.

Windows XP and Windows Server 2003 family VPN clients also support the use of preshared key authentication to create L2TP/IPSec connections.

**Important**

- The use of preshared key authentication is not recommended because it is a relatively weak authentication method. Preshared key authentication creates a master key that is less secure (that might produce a weaker form of encryption) than certificates or the Kerberos V5 protocol. In addition, preshared keys are stored in plaintext. Preshared key authentication is provided for interoperability purposes and to adhere to IPSec standards. It is recommended that you use preshared keys only for testing or as an interim measure in a production environment, when your public key infrastructure (PKI) is being deployed to obtain computer certificates, or when preshared key authentication is required by your VPN server. For long-term use in production environment, it is recommended instead that you use certificates or Kerberos V5.

For a remote access VPN client running Windows XP or a Windows Server 2003 operating system, you can enable preshared key authentication and configure a preshared key in **IPSec settings** on the **Security** tab in the properties of a VPN connection in Network Connections.

**Domain controllers and DHCP, DNS, or WINS servers**
Before enabling IPSec for a computer that functions as a domain controller or DHCP, DNS, or WINS server, determine whether all of the clients are IPSec-enabled. If the IPSec policy assigned to the domain controller or DHCP, DNS, or WINS server is not configured to allow fall back to clear, which permits unsecured communication with non-IPSec-capable clients, secure negotiation will fail and access to these network services will be blocked.

**SNMP**
If you are using Simple Network Management Protocol (SNMP) and your SNMP management systems or agents include computers running Windows 2000, Windows XP, or a Windows Server 2003 operating system that are configured with IPSec policy, you can add a rule to the appropriate IPSec policy to prevent SNMP messages from being secured. The filter list must specify the source and destination addresses of the SNMP management systems and agents for all UDP traffic that uses ports 161 and 162. The filter action should be set to **Permit**, which passes through any traffic that matches the IP filter list.

# IPSec troubleshooting tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting tools**

This topic provides information about the following IPSec troubleshooting tasks and the troubleshooting tools that you can use to perform each of these tasks:

- Viewing IPSec policy assignment information
- Viewing details about the active IPSec policy and IPSec statistics
- Viewing details about the active IPSec policy and IPSec statistics in IP Security Monitor
- Verifying that security auditing is enabled
- Viewing IPSec-related events in Event Viewer
- Enabling detailed tracing for Internet Key Exchange (IKE) negotiations
- Viewing IPSec and other network communication with Network Monitor
- Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family
- Using Ipseccmd.exe to manage and monitor IPSec on computers running Windows XP
- Using Netdiag.exe to display IPSec information and to test and view network configuration
- Disabling TCP/IP and IPSec hardware acceleration

Details about each of these subjects are provided in the following sections.

**Notes**

- For general recommended practices that can help you enhance security and minimize the potential for problems when deploying IPSec, see IPSec Best practices.
- For the Windows Server 2003 family, the Netsh commands for IPSec replace and enhance the functionality provided by Netdiag.exe and Ipseccmd.exe. For information about how to use Netsh to script IPSec policy creation and to monitor IPSec activity, see Netsh commands for Internet Protocol security.
- Ipseccmd.exe is a support tool that is provided only on the Windows XP CD. To install this tool, run the Setup.exe file from the Support\Tools folder. You can use this tool to manage and monitor IPSec policies only on computers running Windows XP. For more information about Ipseccmd.exe, see the Microsoft Tech Net Web site.

Netdiag.exe is available in the Windows Server 2003 family, Windows XP, and in Windows 2000, as follows:

- In the Windows Server 2003 family, run the Suptools.msi file from the \Support\Tools folder on the Windows Server 2003 family CD (select the **Complete** setup option). Although you can still use this version of Netdiag.exe to obtain basic networking information that is not IPSec-specific, the Netsh commands for IPSec replace all IPSec-specific functionality.
- In Windows XP, run the Setup.exe file from the \Support\Tools folder on the Windows XP CD (select the **Complete** setup option).

- For Windows 2000, an updated version of Netdiag.exe is available for download from the Web. For more information, see "Windows 2000 Resource Kits," at the Microsoft Windows Resource Kits Web site.

For more information about Netdiag.exe, see Using Netdiag.exe to display IPSec information and to test and view network configuration in this topic.

**Viewing IPSec policy assignment information**

Viewing the name of an active IPSec policy and the name of the Group Policy object to which the active IPSec policy is assigned can be useful for troubleshooting policy precedence issues. The following table summarizes the tools that you can use to view this information for the Windows Server 2003 family, Windows XP, and Windows 2000.

| On computers running | Use these tools to view the name of the active IPSec policy | Use these tools to view the name of the Group Policy object to which the active IPSec policy is assigned |
|---|---|---|
| Windows Server 2003 family | • IP Security Monitor console.<br><br>For more information, see View details about active IPSec policies.<br><br>• Netsh, **netsh ipsec static show gpoassignedpolicy** command.<br><br>For more information, see Netsh commands for Internet Protocol security. | • Resultant Set of Policy (RSoP) console.<br><br>For more information, see Use Resultant Set of Policy (RSoP) to View IPSec Policy Assignments.<br><br>• Netsh, **netsh ipsec static show gpoassignedpolicy** command. |
| Windows XP | • IP Security Policy Management console (for local policies only).<br><br>• Network Connectivity Tester (Netdiag.exe), **netdiag /test:ipsec** command.<br><br>For more information, see Using Netdiag.exe to display IPSec information and to test and view network configuration in this topic. | • Netdiag.exe, **netdiag /test:ipsec** command. |
| Windows 2000 | • Netdiag.exe, **netdiag /test:ipsec** command.<br><br>• In the properties for the relevant network connection, **TCP/IP Properties/Advanced/Options/IPSec**.<br><br>The assigned IPSec policy that is displayed in **TCP/IP Properties** is global. It is not specific to the connection. | • Netdiag.exe, **netdiag /test:ipsec** command.<br><br>• Gpresult.exe (Group Policy Results).<br><br>• Gpotool.exe (Group Policy Verification Tool).<br><br>Gpresult.exe and Gpotool.exe are both available for download from the Web. For more information, see "Windows 2000 Resource Kits," at the Microsoft Windows Resource Kits Web site. |

**Notes**

- To view all IPSec policies that are available (but not necessarily assigned or applied) to computers, use the IP Security Policy Management console. For information about policy precedence and IPSec policy behavior in an Active Directory environment, see the section "Active Directory-based policy" in Creating, modifying, and assigning IPSec policies.

- For the Windows Server 2003 family, to determine which IPSec policies are assigned but are not being applied to IPSec clients, use the RSoP console.

- The Windows XP implementation of the RSoP console does not support the display of IPSec policies. In addition, the **gpresult /scope computer** command does not display the Group Policy object that contains an IPSec policy assignment. For these reasons, you should use Netdiag.exe to view IPSec policy assignment information on computers running Windows 2000 or Windows XP. The **netdiag /test:ipsec** command displays the Group Policy object that contains the IPSec policy assignment, and the organizational unit to which the Group Policy object is assigned.

- Gpotool.exe allows you to monitor the health of Group Policy objects on domain controllers. You can use this tool to check the consistency and replication of Group Policy objects and to display Group Policy object properties. Gpotool.exe is available only for computers running Windows 2000.

**Viewing details about the active IPSec policy and IPSec statistics**

The IP Security Policy Management console, Netdiag.exe, Gpresult.exe, and Gpotool.exe allow you to determine which IPSec policy has been assigned through Group Policy. After you verify this information, you might need to view details about the assigned IPSec policy and IPSec statistics (for example, filters, filter actions, and active security associations). The following table summarizes the tools that you can use to view an active IPSec policy and IPSec statistics for the Windows Server 2003 family, Windows XP, and Windows 2000.

| **On computers running** | **Use these tools to display the active IPSec policy, IPSec statistics, or both** | **Notes** |
|---|---|---|
| Windows Server 2003 family | • IP Security Monitor console.<br><br>For more information, see View IP security statistics<br><br>• Netsh, **netsh ipsec dynamic show all** command. | • To view details about the active IPSec policy and IPSec statistics on a local or remote computer, you must be a member of the Administrators group on that computer. |
| Windows XP | • IP Security Monitor console.<br><br>• IPseccmd.exe, **ipseccmd show all** command. | • To view details about the active IPSec policy and IPSec statistics for a local or remote computer, you must be a member of the Administrators group on that computer. |
| Windows 2000 | • Netdiag.exe, **netdiag /test:ipsec /v /debug** command.<br><br>• Ipsecmon.exe. | • To use the **/debug** option, you must be logged on as a member of the Administrators group on that computer. In addition, to view details about Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory.<br><br>• Ipsecmon.exe only displays active outbound quick mode security associations (SAs). |

**Note**

- To remotely monitor IPSec on a computer that is running a different version of Windows than your computer, use Remote Desktop Connection. For example, if your computer is running the Windows Server 2003 family and you plan to

remotely monitor IPSec on computers running Windows 2000 or Windows XP, use Remote Desktop Connection to gain remote access to these computers. For information about Remote Desktop Connection, see Remote Desktop Connection.

**Viewing details about the active IPSec policy and IPSec statistics in IP Security Monitor**

In Windows XP and the Windows Server 2003 family, IP Security Monitor is implemented as a Microsoft Management Console (MMC) snap-in, and it includes enhancements that allow you to view details about an active IPSec policy that is applied by the domain or locally, as well as quick mode and main mode statistics, and active IPSec SAs. IP Security Monitor also enables you to search for specific main mode or quick mode filters. To troubleshoot complex IPSec policy designs, you can use IP Security Monitor to search for all matches for filters of a specific traffic type.

- For procedures on how to use IP Security Monitor, see Monitor IPSec Activity.
- For detailed information about IP security statistics, see Viewing main mode and quick mode statistics in IP Security Monitor.

### Notes

- You can use IP Security Monitor to monitor only computers that are running Windows XP or the Windows Server 2003 family. In addition, the computer that is being monitored must run the same version of the Windows operating system as the computer on which IP Security Monitor is running. To monitor IPSec on a computer that is running Windows 2000, use the **ipsecmon** command at the Windows 2000 command prompt on the computer that is being monitored.
- For remote monitoring, you can use IP Security Monitor only to monitor computers that are running the same version of the Windows operating system. To remotely monitor IPSec on a computer that is running a different version of Windows than your computer, use Remote Desktop Connection. For information about Remote Desktop Connection, see Remote Desktop Connection.
- If your computer is running the Windows Server 2003 family, and you plan to monitor IPSec on computers that are also running the Windows Server 2003 family, you can run IP Security Monitor or use the Netsh command-line tool remotely.

**Verifying that security auditing is enabled**

You can use Local Security Policy settings (for a local computer) or Group Policy Object Editor (for a domain) to verify that security auditing is enabled, so you can ensure that the success and failure of IKE negotiations is recorded. Auditing for IKE is supported in Windows 2000, Windows XP, and the Windows Server 2003 family (IKE uses the Logon Events category). In the Windows Server 2003 family, you can also enable auditing for the security policy database (SPD). SPD uses the Policy Change category. For more information, see Define or modify auditing policy settings for an event category.

**Viewing IPSec-related events in Event Viewer**

You can use Event Viewer to view the following IPSec-related events:

**IKE events (negotiation success and failure) in the security log.**

To view these events, enable success or failure auditing for the **Audit logon events** audit policy for your domain or local computer. The IKE event category is also used for auditing user logon events in services other than IPSec. For more information, see Define or modify auditing policy settings for an event category.

When you enable success or failure auditing for the **Audit logon events** audit policy, IPSec records the success or failure of each main mode and quick mode negotiation and the establishment and termination of each negotiation as separate events. However, enabling this type of auditing can cause the security log to fill with IKE events. For example, for servers that are connected to the Internet, attacks on the IKE protocol can cause the security log to fill with IKE events. IKE events can also fill the security log for servers that use IPSec to secure traffic to many clients. To avoid this, you can disable auditing for IKE events in the security log by modifying the registry.

- To disable auditing of IKE events in the security log, do the following:

  1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Audit\DisableIKEAudits** registry setting to a value of **1**.

     The **DisableIKEAudits** key does not exist by default and must be created.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

1. Restart the computer, or stop and then restart the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

**Note**

- Stopping and restarting the IPSec service can disconnect all of the computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. If you restart the IPSec service immediately, TCP-based communication might resume, due to the retransmit behavior of TCP, after new IKE and IPSec SAs are established. For more information, see "Stopping and restarting the IPSec service," later in this topic.

**IPSec policy change events in the Security log.**

To view these events, enable success or failure auditing for the **Audit policy change** audit policy for your domain or local computer. For more information, see Define or modify auditing policy settings for an event category.

**IPSec driver per-packet drop events in the System log.**

In Windows 2000, Windows XP, and the Windows Server 2003 family, you can enable packet event logging for the IPSec driver by modifying the registry. The IPSec driver reads the registry during computer startup. In the Windows Server 2003 family, you can enable packet event logging for the IPSec driver by using the **Netsh ipsec** command-line tool. To enable logging of dropped inbound and outbound packets, specify a value of **7**. For information about the other levels of IPSec driver event logging, see Notes.

- To enable IPSec driver logging of dropped inbound and outbound packets by modifying the registry, do the following:

  1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSec\EnableDiagnostics DWORD** registry setting to a value of **7**.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

  1. Restart the computer.

- To enable IPSec driver logging of dropped inbound and outbound packets by using the Netsh IPSec command-line tool in the Windows Server 2003 family, do the following:

  1. At the command prompt, type: **netsh ipsec dynamic set config ipsecdiagnostics 7**

  2. Restart the computer.

You can also change the interval for writing IPSec driver packet events to the System log. By default, the IPSec driver writes events to the System log once an hour or after a threshold for the number of events has been reached. For troubleshooting, you should set this interval to the minimum value, 60 seconds.

- To change the interval for writing IPSec driver packet events by modifying the registry, do the following:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSec\LogInterval DWORD** registry setting to **60 decimal**.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

1. Restart the computer.

- To change the interval for writing IPSec driver packet events by using the **Netsh ipsec** command-line tool in the Windows Server 2003 family, do the following:

  1. At the command prompt, type **netsh ipsec dynamic set config ipseclogginterval 60**

  2. Restart the computer.

The IPSec driver reads the registry during computer startup.

By default, packet event logging for the IPSec driver is disabled (that is, the value is set to **0**). When you enable packet event logging for the IPSec driver, you can specify any of the following values, to enable different levels of logging:

| Logging Level | Description |
|---|---|
| **1** | When **1** is specified, bad SPI packets (the total number of packets for which the Security Parameters Index or SPI was incorrect), IKE negotiation failures, IPSec processing failures, packets received with invalid packet syntax, and other errors are recorded in the System log. Unauthenticated hashes (with the exception of the "Clear text received when should have been secured" event) are logged as well. |
| **2** | When **2** is specified, inbound per-packet drop events are recorded in the System log. |
| **3** | When **3** is specified, level 1 and level 2 logging are performed. In addition, unexpected clear text events (packets that are sent or received in plaintext) are also recorded. |
| **4** | When **4** is specified, outbound per-packet drop events are recorded in the System log. |
| **5** | When **5** is specified, level 1 and level 4 logging are performed. |
| **6** | When **6**is specified, level 2 and level 4 logging are performed. |
| **7** | When **7** is specified, all levels of logging are performed. |

For more information about IPSec driver event logging, see the Microsoft Windows Resource Kits Web site.

You cannot configure audit policies on a computer running Windows XP Home Edition. However, success and failure auditing for the **Audit logon events** and **Audit policy change** audit policies for the local computer are enabled by default.

Enabling auditing in the security log or IPSec driver diagnostics in the system log can cause these logs to fill with events quickly. Before you perform either of these tasks, you should do the following:

- Ensure that the size of the log is at least 10 megabytes (MB).

- Save the existing log to a file.

- Clear all events in the log so that the log is empty.

- Consider disabling auditing of IKE events in the security log by modifying the registry, if your computer hosts many simultaneous network connections.

**Enabling detailed tracing for Internet Key Exchange (IKE) negotiations**

Enabling audit logging for IKE events and viewing the events in Event Viewer is the fastest and simplest way to troubleshoot failed main mode or quick mode negotiations. However, some scenarios might require a more detailed analysis of the IKE main mode negotiation and quick mode negotiations for troubleshooting. You can enable tracing for IKE negotiations if the audit failure events do not provide enough information. The IKE tracing log is a very detailed log intended for troubleshooting IKE interoperability under controlled circumstances. Expert knowledge of the ISAKMP RFC 2408 and IKE RFC 2409 is required to interpret this log.

The IKE tracing log appears as the *systemroot*\Debug\Oakley.log file. The log has a fixed size of 50,000 lines and will overwrite as necessary. A new Oakley.log file is created each time the IPSec service is started and the previous version of the Oakley.log file is saved as Oakley.log.sav. When the Oakley.log file becomes full, the current file is saved as Oakley.log.bak, and a new Oakley.log file is created.

Because many IKE negotiations can occur simultaneously, you should minimize the number of negotiations and log for as short a period of time as possible to capture a more easily interpreted log.

**Enabling and disabling the IKE tracing log in the** Windows Server 2003 **family**

In the Windows Server 2003 family, you can enable or disable the IKE tracing log dynamically while the IPSec service is running by doing the following:

- To enable the IKE tracing log, type the following at the command prompt:

  **netsh ipsec dynamic set config ikelogging 1**

  This command creates the IKE tracing log file if it does not exist. If the file does exist, it appends logging information to the existing file.

- To disable the IKE tracing log, type the following at the command prompt:

  **netsh ipsec dynamic set config ikelogging 0**

**Enabling the IKE tracing log in Windows 2000 and** Windows XP

In Windows 2000 and Windows XP, you must enable IKE tracing by modifying the registry. For the changes to take effect, you must also stop and restart the IPSec service:

- To enable the IKE tracing log in Windows XP and Windows 2000, do the following:

  1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLoggingDWORD** registry setting to a value of **1**.

     The Oakley key does not exist by default and must be created.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

  1. Stop and start the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

**Stopping and restarting the IPSec service**

Stopping and restarting the IPSec service can disconnect all of the computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. However, if you restart the IPSec service immediately, TCP-based communication might resume, due to the retransmit behavior of TCP, after new IKE and IPSec SAs are established. If you leave the IPSec service stopped, any clients that used the default response rule to establish security will be unable to communicate with this computer for two hours.

To avoid losing Terminal Services connectivity for computers that are using IPSec over a Terminal Server session, you must stop and restart the IPSec service by using a single command line:

- To stop and restart the IPSec service for computers over a Terminal Server session

    - At the command prompt, type the following:

      **net stop policyagent & net start policyagent**

If you are restarting the IPSec service while an L2TP/IPSec VPN tunnel is connected, the tunnel will lose connectivity and must be reconnected when the IPSec service is restarted. A PPTP tunnel, however, does not use IPSec and therefore will stay connected if you stop the IPSec service.

If you are restarting the IPSec service on a computer that is running the Windows Server 2003 family or the Windows Server 2003 family and that is also running the Routing and Remote Access service, any IPSec configuration for L2TP will be lost and the L2TP tunnels will be disconnected. Therefore, you must stop and restart the Routing and Remote Access service, as well as the IPSec service.

- To stop and restart the IPSec service and the Routing and Remote Access service, do the following:

    1. Stop the Routing and Remote Access service using the **net stop remoteaccess** command.
    2. Stop the IPSec service by using the **net stop policyagent** command.
    3. Start the IPSec service by using the **net start policyagent** command.
    4. Start the Routing and Remote Access service using the **net start remoteaccess** command.

**Viewing IPSec and other network communication with Network Monitor**

You can install and use Network Monitor to view IPSec and other network communication. Note that the version of Network Monitor that is provided with the Windows Server 2003 family can be used only to view the network traffic that is sent to or from the computer on which it is installed. To view network traffic that is sent to or from another computer and is routed through your computer (using the Routing and Remote Access service), you must use the Network Monitor component that is provided with Microsoft Systems Management Server.

The Network Monitor component that is provided with the Windows Server 2003 family includes parsers for the ISAKMP (IKE), AH, and ESP protocols. The Network Monitor parsers for ESP can parse inside the ESP packet only if null-encryption is being used and the full ESP packet is captured. Network Monitor cannot parse the encrypted portions of IPSec-secured ESP traffic when encryption is performed in software. However, if encryption is being performed by an IPSec hardware offload network adapter, the ESP packets are decrypted when Network Monitor captures them and as a result, can be parsed and interpreted into the upper-layer protocols. If you need to diagnose ESP software-encrypted communication, you must disable ESP encryption and use ESP-null encryption by changing the IPSec policy on both computers.

### Notes

- For information about how to install the version of Network Monitor that is provided with the Windows Server 2003 family, see Install Network Monitor.

- For information about Systems Management Server, go to the Microsoft Web site.

**Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family**

For computers running the Windows Server 2003 family, you can use the Netsh commands for IPSec to script IPSec policy creation, display details about IPSec policies, and change the IPSec configuration for troubleshooting.

For troubleshooting, you can use the Netsh commands for IPSec to change the following settings:

#### Disable Internet Key Exchange (IKE) certificate revocation list (CRL) checking

By default, in Windows 2000 CRLs are not checked during IKE certificate authentication. In Windows XP and the Windows Server 2003 family, CRLs are checked during IKE certificate authentication, but a fully successful check is not required for the certificate to be accepted. In some cases, failures during CRL processing might cause IKE to not accept the certificate. Or, the delay required for CRL checking might delay IKE negotiation enough to cause the connection attempt to time-out. To determine whether certificate authentication will be successful without CRL checking, you can disable IKE CRL checking. To do this, type the following at the command prompt:

**netsh ipsec dynamic set config strongcrlcheck 0**

This setting takes effect immediately. It does not require a restart.

#### Enable IPSec driver event logging

To record all inbound and outbound dropped packets and other packet processing errors in the Event Viewer System log, you can set the IPSec driver event logging level to **7**. To do this, type the following at the command prompt:

**netsh ipsec dynamic set config ipsecdiagnostics 7**

This setting will not take effect until you restart your computer.

### Notes

- If you configured an IPSec filter action on a computer to block traffic, setting the IPSec driver event logging level to **7** on that computer might generate many events that fill the System log very quickly. Accordingly, you should set the IPSec driver event logging level to **7** only for testing.

- All packets that are dropped by IPSec contribute to the **Datagrams Received Discarded** or **Datagrams Outbound Discarded** System Monitor counters for the IP object. IPSec itself does not provide a counter object. You can monitor IPSec performance counters manually by using IP Security Monitor.

#### Permit inbound and outbound traffic during computer startup

To permit inbound and outbound traffic during computer startup (before the IPSec service starts), you can use the following Netsh command:

**netsh ipsec dynamic set config bootmode permit**

This setting will not take effect until you restart your computer.

### Note

- If you start the computer in Safe Mode with Networking, the IPSec service will not start and cannot retrieve IPSec policy settings from Active Directory or the local registry, and persistent policy cannot be applied. As a result, inbound and outbound traffic will continue to be permitted until the computer is no longer in Safe Mode.

#### Exempt all broadcast, multicast, IKE, Kerberos, and RSVP traffic from IPSec filtering

In Windows 2000 and Windows XP, by default, the IPSec driver exempts all broadcast, multicast, IKE, Kerberos, and RSVP traffic from IPSec filtering. In the Windows Server 2003 family, only IKE traffic is exempt from IPSec filtering, and you can configure, block, or permit filter actions specifically for multicast and broadcast traffic (IPSec does not negotiate SAs for multicast and broadcast traffic). To restore the IPSec driver to the default Windows 2000 and Windows XP filtering behavior, you can use the following Netsh command:

**netsh ipsec dynamic set config ipsecexempt 0**

This setting will not take effect until you restart your computer.

### Caution

- The Windows 2000 and Windows XP default exemption settings for IPSec are designed for corporate LAN environments with a low risk of attack. For this reason, you should use only the Windows 2000 and Windows XP default exemption settings when necessary for troubleshooting, in low-risk environments, or when you cannot solve program compatibility issues by configuring explicit filters in IPSec policies.

You can also exempt all broadcast, multicast, IKE, Kerberos, and RSVP traffic from IPSec filtering by modifying the registry as follows:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSEC\NoDefaultExempt DWORD** registry setting to a value of **0**.

2. The **NoDefaultExempt** key does not exist by default and must be created.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer

3. Restart the computer.

#### Using Ipseccmd.exe to manage and monitor IPSec on computers running Windows XP

For computers running Windows XP, you can use the Ipseccmd.exe command-line tool to script IPSec policy creation and to display IPSec policy assignments, active SAs, and detailed IPSec policy settings, including filters, filter actions, and authentication methods. Ipseccmd.exe replaces and enhances the functionality provided by the Ipsecpol.exe (used in Windows 2000 for scripting IPSec policy creation) and the Netdiag.exe **/test:ipsec /v /debug** command (used in Windows 2000 to display configuration information and SAs for an active IPSec policy).

To install Ipseccmd.exe on a computer running Windows XP, run the Setup.exe file from the \Support\Tools folder on the Windows XP CD.

**Notes**

- Ipseccmd.exe is provided only with Windows XP, and you can use this tool to manage and monitor IPSec policies only on computers running Windows XP.

- To display all IPSec policy settings and statistics for diagnostics, use the **ipseccmd show all** command.

#### Using Netdiag.exe to display IPSec information and to test and view network configuration

You can use the Netdiag.exe command-line tool to display information about IPSec policies and statistics, report network configuration, and test basic networking capabilities and domain-based functionality. The following table describes how to install Netdiag.exe and provides information about functional changes relevant to IPSec for each version of this tool.

| On computers running | Use this method to install | Notes |
|---|---|---|
| Windows Server 2003 family | For the Windows Server 2003 family, Netdiag.exe is available for installation from the Windows Server 2003 family CD (run the Suptools.msi file from the \Support\Tools folder, and choose the **Complete** setup option. | • Although you can still use Netdiag.exe to obtain basic networking information that is not IPSec-specific, for the Windows Server 2003 family, the **netdiag /test:ipsec** option has been removed, and the Netsh commands for IPSec replace all IPSec-specific functionality. To view details about IPSec policies, use either the **netsh ipsec static show** command or the **netsh ipsec dynamic show** command. |

| | | |
|---|---|---|
| Windows XP | For Windows XP, Netdiag.exe is available for installation from the Windows XP CD (run the Setup.exe file from the \Support\Tools folder, and choose the **Complete** setup option). | • Displays information about IPSec policy assignments, including the name of the active IPSec policy, the name of the Group Policy object that assigned the policy, and the policy path.<br><br>• The display of detailed IPSec policy information and IPSec statistics (provided by the **/debug** and **/v** options in Windows 2000) is not supported in Windows XP. Instead, to view details about IPSec policies, use the Ipseccmd.exe **show** command. |
| Windows 2000 | For Windows 2000, an updated version of Netdiag.exe is available for download from the Web. For information, see "Windows 2000 Resource Kits," at the Microsoft Windows Resource Kits Web site. | • Netdiag.exe was enhanced in Windows 2000 Service Pack 1 to display the number of bytes offloaded by the IPSec driver to a network adapter that is capable of IPSec hardware offload. The hardware offload statistics are Offloaded Bytes Sent and Offloaded Bytes Received.<br><br>• To use the **/debug** option to view details about Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory. |

You can use the **netdiag /v /l** command to obtain networking information that is not IPSec-specific, for any platform. Typical uses of these options include:

- Reporting the IP configuration and routing configuration for a computer with one command.
- Testing WINS and DNS name resolution and consistency.
- Reporting the build version of a computer and the hotfixes that are installed on that computer.
- Testing the validity of domain membership, whether domain members can successfully contact domain controllers, and trust relationships.

The **/l** option generates a Netdiag.log file. This file is written to the folder in which Netdiag.exe is run.

**Notes**

- To run all Netdiag.exe commands, you must be a member of the Administrators group on the local computer.
- Each version of Netdiag.exe is customized to run on a different version of the Windows operating system (that is, Windows 2000, Windows XP, or the Windows Server 2003 family). You can only run Netdiag.exe on a computer that is running the Windows operating system for which Netdiag.exe has been specifically designed.
- If you are running a version of Netdiag.exe that is designed for Windows 2000 or Windows XP in a Windows Server 2003 domain, Netdiag.exe might not report the correct operating system version of the domain.
- If you run Netdiag.exe on a computer that is assigned an IPSec policy that affects most inbound and outbound IP traffic on that computer, Netdiag.exe might not report results correctly the first time it is run. This is because it might take a few seconds for IPSec to establish SAs to the many remote computers for which Netdiag.exe tests connectivity. If this occurs, wait five seconds, and then run Netdiag.exe again.

**Disabling TCP/IP and IPSec hardware acceleration**

Network adapters can accelerate IPSec processing by performing hardware offload of IPSec cryptographic functions (IPSec offload) and the calculation of TCP checksums (checksum offload). Such adapters might also process large TCP segments for very fast transmission (large-send offload). By default, when a network adapter driver that can perform hardware offload is enabled during the Plug and Play initialization process, the driver advertises this capability to TCP/IP and IPSec. TCP/IP and

IPSec then offload tasks to the network adapter driver as appropriate.

If you use network adapters to perform hardware offload of IPSec cryptographic functions, you might need to verify that hardware acceleration is not causing problems with the packet processing that is performed by the network adapter. To do so, you can:

- Disable hardware offload functions in the network adapter driver (if this function is supported by the driver). This task does not require you to restart the computer.
- Disable TCP/IP and IPSec hardware acceleration. This task requires you to restart the computer.
- Disable only IPSec hardware acceleration. This task requires you to restart the computer.

**Disabling hardware offload functions in the network adapter driver**

To determine whether you can disable hardware offload functions in the network adapter driver, do one of the following:

- Open Device Manager, click **Network Adapters**, right-click the network adapter whose properties you want to view, and then click **Properties**. For more information, see Change network adapter settings
- Open Network Connections, right-click the network adapter that you want to view, click **Properties**, and then, on the **General** tab, click **Configure**.

**Disabling TCP/IP and IPSec hardware acceleration**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

To disable both TCP/IP and IPSec hardware acceleration by modifying the registry, do the following:

1. Set the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\DisableTaskOffload** DWORD registry setting to a value of **1**.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

2. Restart the computer.

If the **DisableTaskOffload** DWORD registry setting is set to a value of **0**, and if the network adapter advertises hardware offload capabilities, IPSec and TCP/IP will attempt to offload the appropriate functions to the network adapter.

**Disable only IPSec hardware acceleration**

To disable only IPSec hardware acceleration by modifying the registry, do the following:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSEC\EnableOffload** DWORD registry setting to a value of **0**.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

2. Restart the computer.

**Note**

- To help minimize the potential for hardware acceleration to cause problems with packet processing performed by the network adapter, make sure that you use the latest network adapter driver available from the vendor.

# Viewing IPSec Statistics and Details About Active IPSec Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing IPSec statistics and details about active IPSec policies**
This section covers:

- Viewing main mode and quick mode statistics in IP Security Monitor
- Viewing details about active IPSec policies in IP Security Monitor
- Using Resultant Set of Policy to view IPSec policy assignments

# Viewing main mode and quick mode statistics in IP Security Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing main mode and quick mode statistics in IP Security Monitor**

**Main Mode (IKE) Statistics**
The following table describes the items in the list of main mode (IKE) statistics in IP Security Monitor.

For information about how to view these statistics, see View IP security statistics.

| Main Mode (IKE) Statistic | Description |
|---|---|
| Active Acquire | The number of pending requests to initiate an Internet Key Exchange (IKE) negotiation in order to establish a security association (SA) between IPSec peers. The Active Acquire statistic includes the outstanding request and the number of any queued requests. Under a heavy load, the number of active acquires is 1 plus the number of requests that are queued by IKE for processing. |
| Active Receive | The number of IKE messages received that are queued for processing. |
| Acquire Failures | The total number of acquire outbound requests that have failed since the IPSec service was last started. Acquires are requests to establish SAs between IPSec peers. |
| Receive Failures | The total number of errors that have occurred during the process of receiving IKE messages since the IPSec service was last started. |
| Send Failures | The total number of errors that have occurred during the process of sending IKE messages since the IPSec service was last started. The number of Send Failures typically increases for computers that establish SAs over temporary network connections, such as dial-up connections, virtual private network tunnels, and wireless connections. |
| Acquire Heap Size | The number of entries in the acquire heap. The acquire heap stores successful acquires. Acquires are outbound requests to establish SAs between IPSec peers. |
| Receive Heap Size | The number of entries in the IKE receive buffers. The receive buffers store incoming IKE messages. |
| Authentication Failures | The total number of identity authentication (Kerberos, certificate, and preshared key) failures that have occurred during main mode negotiation since the IPSec service was last started. If you are having difficulty communicating securely, attempt the communication and determine whether the number of Authentication Failures increases. If it does, check your authentication settings for either an unmatched authentication method or an incorrect authentication method configuration (for example, the use of preshared keys that do not match). |
| Negotiation Failures | The total number of negotiation failures that have occurred during main mode or quick mode negotiation since the IPSec service was last started. If you are having difficulty communicating securely, attempt the communication and determine whether the number of Negotiation Failures increases. If it does, check your authentication and security method settings for an unmatched authentication method, an incorrect authentication method configuration (for example, the use of preshared keys that do not match), or unmatched security methods or settings. |
| Invalid Cookies Received | The total number of cookies that could not be matched with an active main mode SA since the IPSec service was last started. A cookie is a value contained in a received IKE message that is used to help identify the corresponding main mode SA. |
| Total Acquire | The total number of requests that have been submitted to IKE since the IPSec service was last started to establish an SA. This number includes acquires that result in soft SAs. |

| | |
|---|---|
| Total Get SPI | The total number of requests that have been submitted by IKE to the IPSec driver to obtain a unique Security Parameters Index (SPI) since the IPSec service was last started. The SPI matches inbound packets with SAs. |
| Key Additions | The total number of outbound quick mode SAs that have been added by IKE to the IPSec driver since the IPSec service was last started. |
| Key Updates | The total number of inbound quick mode SAs that have been added by IKE to the IPSec driver since the IPSec service was last started. |
| Get SPI Failures | The total number of failed requests that have been submitted by IKE to the IPSec driver to obtain a unique SPI since the IPSec service was last started. |
| Key Addition Failures | The total number of failed outbound quick mode SA addition requests that have been submitted by IKE to the IPSec driver since the IPSec service was last started. |
| Key Update Failures | The total number of failed inbound quick mode SA addition requests that have been submitted by IKE to the IPSec driver since the IPSec service was last started. |
| ISADB List Size | The number of main mode state entries. This number includes successfully negotiated main modes, main mode negotiations in progress, and main mode negotiations that failed or expired and have not yet been deleted. |
| Connection List Size | The number of quick mode negotiations that are in progress. |
| IKE Main Mode | The total number of successful SAs that have been created during main mode negotiations since the IPSec service was last started. |
| IKE Quick Mode | The total number of successful SAs that have been created during quick mode negotiations since the IPSec service was last started. |
| Soft Associations | The total number of SAs formed with computers that have not responded to main mode negotiation attempts since the IPSec service was last started. Although these computers did not respond to main mode negotiation attempts, IPSec policy allowed communications with the computers. Soft SAs are not secured by IPSec. |
| Invalid Packets Received | The total number of invalid IKE messages that have been received since the IPSec service was last started. This number includes IKE messages with invalid header fields, incorrect payload lengths, and incorrect values for the responder cookie. Invalid IKE messages are commonly caused by retransmitted IKE messages or an unmatched preshared key between the IPSec peers. |

**Quick Mode (IPSec) Statistics**

The following table describes the items in the list of quick mode (IPSec) statistics in IP Security Monitor.

| **Quick Mode (IPSec) Statistic** | **Description** |
|---|---|
| Active Security Associations | The number of active quick mode SAs. |
| Offloaded Security Associations | The number of active quick mode SAs offloaded to hardware. Certain network adapters can accelerate IPSec processing by performing hardware offload of IPSec cryptographic functions. |
| Pending Key Operations | The number of IPSec key exchange operations that are in progress but are not yet completed. |
| Key Additions | The total number of keys for quick mode SA negotiations that have been successfully added since the computer was last started. |
| Key Deletions | The total number of keys for quick mode SAs that have been successfully deleted since the computer was last started. |

| | |
|---|---|
| Rekeys | The total number of successful rekey operations for quick mode SAs since the computer was last started. |
| Active Tunnels | The number of active IPSec tunnels. |
| Bad SPI Packets | The total number of packets for which the SPI has been incorrect since the computer was last started. If the SPI is incorrect, it might mean that the inbound SA has expired and a packet using the old SPI has recently arrived. This number is likely to increase if rekey intervals are short and there are a large number of SAs. A large number of packets with bad SPIs that are received within a short amount of time might indicate a packet spoofing attack. |
| Packets Not Decrypted | The total number of packets that could not be decrypted since the computer was last started. A packet might not be decrypted if it fails a validation check. |
| Packets Not Authenticated | The total number of packets for which data could not be verified (for which the integrity hash verification failed) since the computer was last started. Increases in this number might indicate an IPSec packet spoofing or modification attack or packet corruption by network devices. |
| Packets With Replay Detection | The total number of packets that have contained an invalid sequence number since the computer was last started. Increases in this number might indicate a network problem or replay attack. |
| Confidential Bytes Sent | The total number of bytes that have been sent using the Encapsulating Security Payload (ESP) protocol (excluding non-encrypted ESP) since the computer was last started. |
| Confidential Bytes Received | The total number of bytes that have been received using the ESP protocol (excluding non-encrypted ESP) since the computer was last started. |
| Authenticated Bytes Sent | The total number of authenticated bytes that have been sent using the Authentication Header (AH) protocol or the ESP protocol since the computer was last started. |
| Authenticated Bytes Received | The total number of authenticated bytes that have been received using the AH protocol or the ESP protocol since the computer was last started. |
| Transport Bytes Sent | The total number of bytes that have been sent using IPSec transport mode since the computer was last started. |
| Transport Bytes Received | The total number of bytes that have been received using IPSec transport mode since the computer was last started. |
| Bytes Sent in Tunnels | The total number of bytes that have been sent using IPSec tunnel mode since the computer was last started. |
| Bytes Received in Tunnels | The total number of bytes that have been received using IPSec tunnel mode since the computer was last started. |
| Offloaded Bytes Sent | The total number of bytes that have been sent using IPSec hardware offload since the computer was last started. |
| Offloaded Bytes Received | The total number of bytes that have been received using IPSec hardware offload since the computer was last started. |

# Viewing details about active IPSec policies in IP Security Monitor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing details about active IPSec policies in IP Security Monitor**

IP Security Monitor allows you to view details about the active IPSec policies that are applied to the domain, the local computer, or a remote computer. This topic describes the type of details that IP Security Monitor displays.

For information about how to view details about active IPSec policies, see View details about active IPSec policies. For information about creating, assigning, and assigning IPSec policies, see Creating, modifying, and assigning IPSec policies.

**Active IPSec policies that are applied by the domain**

The following table describes the items that are displayed when you use IP Security Monitor to view details about an active IPSec policy that is applied by the domain.

| Active IPSec policy item | Description |
|---|---|
| **Policy Name** | Specifies the name of the active IPSec policy. |
| **Description** | Optional text that describes the purpose of the active IPSec policy. |
| **Policy Last Modified** | Not applicable (this information is displayed only for active IPSec policies that are applied to a computer). |
| **Policy Store** | Specifies the storage location for the active IPSec policy. For active IPSec policies that are applied by the domain, this location appears as "Domain Store." |
| **Policy Path** | Specifies the LDAP path to the active IPSec policy. |
| **Organizational Unit** | Specifies the organizational unit to which the Group Policy object is applied. |
| **Group Policy Object Name** | Specifies the Group Policy object to which the active IPSec policy is applied. |

**Active IPSec policies that are applied to a computer**

The following table describes the items that are displayed when you use IP Security Monitor to view details about active IPSec policies that are applied to a computer.

| Active IPSec policy item | Description |
|---|---|
| **Policy Name** | Specifies the name of the active IPSec policy. |
| **Description** | Optional text that describes the purpose of the active IPSec policy. |
| **Policy Last Modified** | Specifies the date and the time that the active IPSec policy was last modified. |
| **Policy Store** | Specifies the storage location of the active IPSec policy. For active IPSec policies that are applied to a computer, this location appears as "Local Store." |
| **Policy Path** | Not applicable (this information is only displayed for active IPSec policies that are applied by the domain). |
| **Organizational Unit** | Not applicable (this information is only displayed for active IPSec policies that are applied by the domain). |

| | |
|---|---|
| **Group Policy Object Name** | Not applicable (this information is only displayed for active IPSec policies that are applied to the domain). |

**Note**

- You cannot use IP Security Monitor to add, edit, or remove IPSec policies. For information about how to add, edit, or remove IPSec policies, see Add, edit, or remove IPSec policies.

# Using Resultant Set of Policy to view IPSec policy assignments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Resultant Set of Policy to view IPSec policy assignments**
Resultant Set of Policy (RSoP) is an addition to Group Policy that you can use to view IPSec policy assignments for a computer or for members of a Group Policy container. This information can help you troubleshoot policy precedence issues and plan your deployment.

To view IPSec policy assignments in RSoP, you must first open the RSoP MMC console, and then run a query. RSoP provides two types of queries: logging mode queries (for viewing IPSec policy assignments for a computer) and planning mode queries (for viewing IPSec policy assignments for members of a Group Policy container).

**Logging mode queries**
You can run an RSoP logging mode query to view all of the IPSec policies that are assigned to an IPSec client. The query results display the precedence of each IPSec policy assignment, so that you can quickly determine which IPSec policies are assigned but are not being applied and which IPSec policy is being applied. The RSoP console also displays detailed settings (that is, the filter rules, filter actions, authentication methods, tunnel endpoints, and connection type) for the IPSec policy that is being applied.

When you run a logging mode query, RSoP retrieves policy information from the Windows Management Instrumentation (WMI) repository on the target computer, and then displays this information in the RSoP console. In this way, RSoP provides a view of the policy settings that are being applied to a computer at a given time.

For information about how to run RSoP logging mode queries for IPSec policies, see Use RSoP to view IPSec policy assignments for a computer. For general information about RSoP logging mode queries, see RSoP logging mode. For general information about RSoP, see RSoP overview.

**Planning mode queries**
You can run an RSoP planning mode query to view all of the IPSec policies that are assigned to members of a Group Policy container. For example, a planning mode query can be useful if you are planning a company reorganization and you want to move computers from one organizational unit to a new organizational unit. By supplying the appropriate information and then running a planning mode query, you can determine which IPSec policies are assigned but are not being applied to the new organizational unit and which IPSec policy is being applied. In this way, you can identify which policy would be applied if you were to move the computers to the new organizational unit. As with logging mode queries, when you run a planning mode query, the RSoP console displays detailed policy settings for the IPSec policy that is being applied.

When you run a planning mode query, RSoP retrieves the names of the target user, computer, and domain controller from the WMI repository on the domain controller. WMI then uses the Group Policy Data Access Service (GPDAS) to create the policy settings that would be applied to the target computer, based on the RSoP query settings that you entered. RSoP reads the policy settings from the WMI repository on the domain controller, and then displays this information in the RSoP console user interface.

For information about how to run RSoP planning mode queries for IPSec policies, see Use RSoP to view IPSec policy assignments for members of a Group Policy container. For general information about RSoP planning mode queries, see RSoP planning mode.

### Note

- You can run an RSoP planning mode query only on a domain controller (when you run a planning mode query, you must explicitly specify the domain controller name). However, you can specify any IPSec client as the target for the query, provided you have the appropriate permissions to do so.

**Assigning and processing IPSec policy in Group Policy**
IPSec policies can be assigned from and stored in Active Directory, as part of Group Policy, or they can be assigned and stored locally, on a computer. When a computer is joined to an Active Directory domain, the domain-level IPSec policy applies. If a computer is not joined to an Active Directory domain, the local IPSec policy applies.

Group Policy settings are contained in Group Policy objects, which are linked with specific Active Directory objects (sites, domains, and organizational units). When an IPSec policy is assigned to a Group Policy object for an Active Directory object (such as an organizational unit), that IPSec policy is propagated to any computer accounts that are affected by the Group Policy

object.

Multiple Group Policy objects, each of which can contain an IPSec policy, can be assigned to a computer account. When multiple IPSec policies are assigned, the last policy that is processed is the policy that is applied (that is, the last policy takes the highest precedence and overrides the settings of any IPSec policy assignments that were processed earlier).

Policy precedence is based on the Group Policy inheritance model. The policy used is the policy assigned at the lowest level of the domain hierarchy for the domain container of which the computer is a member. For example, if there are IPSec policies that are configured for both the domain and for an organizational unit within the domain, the computers that are members of the domain use the domain IPSec policies. The computers that are members of the organizational unit within the domain use the organizational unit IPSec policies. If there are multiple organizational units, members of each organizational unit use the IPSec policy assigned to the organizational unit that is closest in level to their container in the Active Directory hierarchy. If no IPSec policies are configured for Active Directory, or if a computer is not connected to an Active Directory domain, the local IPSec policy is used.

**IPSec policy information displayed in the RSoP console**
The RSoP console simplifies the task of determining which IPSec policy is being applied by displaying the following information for each Group Policy object that contains an IPSec policy assignment: the name of the IPSec policy, the name of the Group Policy object that the IPSec policy is assigned to, the IPSec policy precedence (the lower the number, the higher the precedence), and the name of the site, domain, and organizational unit to which the Group Policy object containing the IPSec policy applies (that is, the scope of management for the Group Policy object).

For information about policy precedence and IPSec policy behavior in an Active Directory environment, see the section "Active Directory-based policy" in Creating, modifying, and assigning IPSec policies. For general information about Group Policy, see Group Policy overview.

### Note

- The settings of the IPSec policy with the highest precedence apply in their entirety; they are not merged with the settings of IPSec policies that are applied at higher levels of the Active Directory hierarchy.

# IPSec Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

**Windows Resource Kits**
The Windows Resource Kits contain extensive technical information on TCP/IP and IPSec. For more information, see Using the Windows Deployment and Resource Kits.

**IPSec-related Web sites**
The following Web sites contain links that relate to IPSec:

- The IPSec Working Group page contains links to RFCs, Internet Drafts, and other IPSec-related links at the Internet Engineering Task Force Web site.
- For information and the latest recommendations on security, see Microsoft Security at the Microsoft Web site.
- For information about the standards and technology for IPSec, as well as other computer security topics, see the National Institute of Standards and Technology Web site.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# IPSec Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Traffic is not being secured.
- IPSec SA negotiation is failing.
- IPSec SA negotiation is failing when using certificate authentication.
- Some computers are unable to communicate with secure computers.
- Local computer IPSec policy is not being used.
- Unable to assign or unassign an IP Security policy.
- DES is being used when 3DES is needed.
- Unsecured traffic is being sent to a secure server.
- Event Viewer is displaying error messages related to IPSec.
- A policy location error message is displayed.
- Event Viewer is displaying "Bad SPI" error message.
- IP Security Monitor displays the error message "The IPSec server is unavailable or incompatible with the IPSec monitor."

**Traffic is not being secured.**
**Cause:** An IPSec security association (SA) that secures traffic (also known as a hard SA) is not being established.

**Solution:** This might be caused by either filter list misconfiguration or the enabling of unsecured communication with non-IPSec-aware computers. Examine the policies that are assigned to both IPSec peers and determine, based on the rules, filter lists, and filter action settings, whether IPSec requires traffic to be secured. Check to determine whether the **Allow unsecured communication with non-IPSec-aware computer** option on the **Security Methods** tab of the appropriate filter action is enabled. Use audit logging to determine the cause of the failed negotiation.

**See also:** IPSec troubleshooting tools; Add, edit, or remove IPSec security methods

**IPSec SA negotiation is failing.**
**Cause:** An IPSec security association (SA) that secures traffic (also known as a hard SA) is not being established. If unsecured communication with non-IPSec-aware computers is enabled, the negotiation failure is expected.

Negotiation failure between two IPSec-aware computers commonly occurs when the default policies are used on two IPSec peers and one or both are not members of an Active Directory domain. Computers running Windows XP Home Edition do not support Active Directory domain membership or the Kerberos V5 authentication method. Default policies use the Kerberos V5 authentication method, which does not work for computers that are not Active Directory domain members.

**Solution:** A negotiation failure between two IPSec-aware computers is most often caused by a mismatched key exchange method, authentication method, or security method. Examine the policies that are assigned to both IPSec peers and determine, based on the advanced settings, rules, filter lists, and filter action settings, whether IPSec should be securing the traffic. If it should, determine whether there is a mismatched key exchange, authentication method, or security method that is preventing a successful negotiation of a hard SA. Use audit logging to determine the cause of the failed negotiation.

**See also:** IPSec troubleshooting tools; Add, edit, or remove IPSec security methods; Define IPSec authentication methods;

Create key exchange security methods

**IPSec SA negotiation is failing when using certificate authentication.**
**Cause:** The correct computer certificates are not installed on the IPSec peers.

**Solution:** Install the correct computer certificates on the IPSec peers.

**Cause:** An incorrect certification authority (CA) is selected in the authentication method.

**Solution:** For the rules of each IPSec peer, select the CA that issued the computer certificate that is installed on the other IPSec peer.

**Cause:** The certificate chain from the issuing CA to the root CA is broken.

**Solution:** Install a computer certificate on each IPSec peer that follows a valid certificate chain up to a trusted root CA.

**Cause:** The computer certificate on one or both of the IPSec peers has expired.

**Solution:** Install a current computer certificate on each IPSec peer.

**See also:** IPSec troubleshooting tools; Authentication methods; Define IPSec authentication methods; Special IPSec considerations; Using Certificates

**Some computers are unable to communicate with secure computers.**
**Cause:** Secure computers are using a lockdown policy that requires client computers to initiate secure communications. The computers that are unable to communicate are not configured with a rule to initiate secure communication.

**Solution:** Either allow unsecured incoming connection requests on the secure servers or configure additional rules for the computers that are unable to communicate so that they can initiate secure communication with the secure computers.

**Cause:** Client computers are not IPSec-capable and secure computers are not allowing unsecured communications.

**Solution:** Either install client computers that are IPSec-capable or reconfigure the secure computers to allow communication with computers that are not IPSec-capable.

**Cause:** Some computers are configured with a policy that has a key exchange, authentication method, or security method that is mismatched with the policy configured for the secure computers.

**Solution:** Examine the policies that are assigned to both IPSec peers and determine, based on the advanced settings, rules, filter lists, and filter action settings, whether IPSec should be securing the traffic. If it should, determine whether there is a mismatched key exchange, authentication method, or security method that is preventing a successful SA negotiation. Use audit logging to determine the cause of the failed negotiation.

**See also:** Add, edit, or remove filter actions; IPSec troubleshooting tools; Add, edit, or remove IPSec security methods; Define IPSec authentication methods; Create key exchange security methods

**Local computer IPSec policy is not being used.**
**Cause:** The computer is a member of an Active Directory domain and is receiving Active Directory-based IPSec policy.

**Solution:** All domain members receive Active Directory-based IPSec policy, which overrides local IPSec policy settings. To prevent this, remove the computer from the domain.

**Cause:** The IPSec service is not started.

**Solution:** Start or restart the IPSec service.

**See also:** Restart the IPSec services

**Unable to assign or unassign an IP Security policy.**
**Cause:** You cannot assign or unassign policies when IP Security Policies are used for Active Directory.

**Solution:** To assign Active Directory-based policies, use IP Security Policies within Group Policy.

**See also:** Assign or unassign IPSec policy in Group Policy

**DES is being used when 3DES is needed.**
**Cause:** The IPSec peer is a computer running Windows 2000 without either the Windows 2000 High Encryption Pack or Windows 2000 Service Pack 2 (or later) installed.

**Solution:** Install either the Windows 2000 High Encryption Pack or Windows 2000 Service Pack 2 (or later) on the computer

running Windows 2000.

**Unsecured traffic is being sent to a secure server.**
**Cause:** The client computer is using the default response rule.

**Solution:** When the client computer is configured with the default response rule, it sends unsecured packets when it either initiates communication or restarts communication after a delay long enough for the quick mode SA and dynamic filter list to time out. To prevent this, configure additional rules for client computers that always initiate secure communication to secure servers.

**See also:** Security information for IPSec

**Event Viewer is displaying error messages related to IPSec.**
**Cause:** There might be a problem with either the IPSec Policy Agent or the Oakley service.

**Solution:** View the system and security event logs for messages that might help in determining the origin of the problem.

**A policy location error message is displayed.**
**Cause:** This error might occur if the IPSec Policy Agent is attempting to retrieve policy from an incorrect storage location.

**Solution:** Examine both the IPSec Policy Agent events in Event Viewer and the IPSec Policy Agent log for messages that indicate from what location policy is being retrieved (the local registry or the domain). If **HKLM\SOFTWARE\Policies\Microsoft\Windows\IPSec\GPTIPSecPolicy\DSIPSecPolicyFlags** is set to 1, then IPSec policy is obtained from Active Directory. Otherwise, it is obtained from the registry. The local computer warns a user that domain policy is in effect when it starts. If the policy is supposed to be Active Directory-based, but cannot be obtained, you can delete **HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Group Policy\History** and then restart the computer.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**See also:** View Event Logs.

**Event Viewer is displaying "Bad SPI" error message.**
**Cause:** This error might occur if a key lifetime value is set too low, or if the SA has expired but the sender continues to transmit data to the receiver.

**Solution:** To determine and correct the problem, use IP Security Monitor to examine the number of rekeys. If the number of rekeys is very large compared to the amount of time the connections have been active, set the key lifetimes in the policy to be longer. Good values for high-traffic Ethernet connections are greater than both 50 megabytes and five minutes. This event is usually logged when an SA has been deleted on one IPSec peer, and the other IPSec peer sends secure traffic on the SA. This is usually a temporary situation that occurs on systems that are busy and fixes itself.

**IP Security Monitor displays the error message "The IPSec server is unavailable or incompatible with the IPSec monitor."**
**Cause:** You are using the IP Security Monitor console to monitor IPSec on a computer running Windows 2000.

**Solution:** The IP Security Monitor console can only be used to monitor IPSec on computers running Windows XP or a Windows Server 2003 operating system. To monitor IPSec on a computer running Windows 2000, use the **IPSecmon** command at the Windows 2000 command prompt on the computer that is being monitored.

**See also:** IPSec troubleshooting tools

# DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS**
Domain Name System (DNS) is the name resolution protocol for TCP/IP networks, such as the Internet. A DNS server hosts the information that enables client computers to resolve memorable, alphanumeric DNS names to the IP addresses that computers use to communicate with each other.

- Before you begin using this computer as a DNS server, see DNS Checklists.
- To find features that have been moved in the Microsoft® Windows Server 2003 family, see New ways to do familiar DNS tasks.
- For tips about using DNS, see DNS best practices.
- For help with specific tasks, see DNS How To....
- For general background information, see DNS Concepts.
- For problem-solving instructions, see Troubleshooting DNS.

# DNS Checklists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

- Checklist: Installing a DNS server
- Checklist: Configuring a DNS server
- Checklist: Deploying DNS
- Checklist: Deploying DNS for Active Directory
- Checklist: Verifying DNS before installing Active Directory
- Checklist: Securing your DNS infrastructure

# Checklist: Installing a DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a DNS server**

| Step | Reference |
|---|---|
| If you are new to working with Domain Name System (DNS), first learn the basic DNS concepts and then examine how to use DNS-related tools to deploy and maintain the client and server features of DNS. | DNS Overview |
| Learn in more detail about how DNS works and understand more about DNS-related features. | Understanding DNS; Deploying DNS; Administering DNS |
| Before installing a DNS server, determine if the server is now a domain controller or whether it will be promoted to one later. | Server planning for DNS; Interoperability issues |
| Install the first DNS server on your network.<br><br>If you want to install Active Directory, the Active Directory Installation Wizard will automatically prompt you to install a DNS server. Otherwise, you can install the DNS Server service as a Windows component. | Install a DNS server |
| Select the checklist that best describes your deployment needs and follow it to complete your DNS installation. | Checklist: Configuring a DNS server; Checklist: Deploying DNS; Checklist: Deploying DNS for Active Directory |

# Checklist: Configuring a DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring a DNS server**

| Step | Reference |
|---|---|
| If you have not installed a DNS server, first follow the DNS server installation checklist, and then return to this checklist to configure your new DNS server. | Checklist: Installing a DNS server |
| **For small office/home office.** If you are setting up a home network with multiple machines connecting to the Internet, have the following on hand:<br>• The name of the zone for which this server will be authoritative<br>• The IP address of your ISP's DNS server | Understanding zones and zone transfer; Dynamic update; Using forwarders |
| **For medium office/large office.** If the DNS server will service requests from clients on a local network that may or may not be available from the internet, have the following on hand:<br>• The name of the zone for which this server will be authoritative<br>• Possibly:<br>  • The name of a reverse lookup zone<br>  • The IP address of a DNS server used for forwarding | Server planning for DNS; Zone planning for DNS; Reverse lookup; Using forwarders; Active Directory integration |
| Before configuring your DNS server, learn DNS server best practices, server features, and the concepts involved in configuring primary and secondary servers. | DNS best practices; Server features; Configuring a new primary server; Using secondary servers |
| Configure your DNS server. | Configure a new DNS server |
| Verify the DNS server's configuration. | Verifying server configuration |

# Checklist: Deploying DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Deploying DNS**

| Step | Reference |
|---|---|
| Review best practices and all planning issues specific to deploying Domain Name System (DNS) servers on your network when not additionally deploying Active Directory. | Deploying DNS; DNS best practices |
| Plan your namespace by specifying a domain name standard for your organization. | Namespace planning for DNS |
| Plan how to divide up your DNS domain name and network address space into forward and reverse lookup zones as needed. | Zone planning for DNS; Reverse lookup |
| Plan your DNS servers, such as determining how many servers you need to use and where to put them on your network. | Server planning for DNS |
| (Optional) Test and evaluate server performance using DNS performance optimizing and monitoring tools. | Monitoring and Optimizing Servers |
| Review migration and interoperability issues, if appropriate. | Migrating servers; Interoperability issues |
| If clients on your network need to be able to resolve external DNS names, consider the use of forwarders for your network. | Using forwarders; Configure a DNS server to use forwarders |
| Register your organization's domain name with the Internet domain name registrar. This name is used to form the DNS names your organization uses on the Internet.<br><br>**Internet DNS**<br><br>**Name:** _____._____<br><br>**Note**<br><br>• Registering a parent or second-level DNS domain name for your organization with an active Internet domain name registration authority is only required if you plan to use this name externally on the Internet. | Namespace planning for DNS; Interoperability issues |
| Determine how to manage DNS client resolver configurations. | Managing Clients |
| Begin adding and configuring zones for your first DNS server. | Configure a new DNS server; Add a forward lookup zone; Add a reverse lookup zone |
| Install and configure additional DNS servers if they are to be used as secondary servers for your initial zones. | Install a DNS server; Add a secondary server for an existing zone; Using secondary servers |
| (For primary zones only) If the previous step was used, add additional secondary servers for your existing zones to the name server (NS) record to make them authoritative for the zone. | Specify other DNS servers as authoritative for a zone |
| (Optional) Install any additional zones (and DNS servers) for subdomains you want to delegate away from your first DNS server and its zones. | Configure a new DNS server; Configuring a new primary server |

| | |
|---|---|
| If the previous step was used, add delegations in parent zones for any subdomains added based on the previous step.<br><br>For example, if adding sub.example.microsoft.com as a new DNS domain, you would need to add to it a delegation at the example.microsoft.com zone. | Create a zone delegation; Delegating zones |
| Add DNS resource records--such as A, PTR, CNAME, and MX records--as needed to complete zone configurations. | Managing resource records; Add a resource record to a zone; Resource records reference |
| Enable or disable dynamic updates for zones as needed. | Allow dynamic updates; Dynamic update |
| (Optional) Enable WINS lookup for selected zones. | Enable DNS to use WINS resolution; WINS lookup integration; Using WINS lookup |
| (Optional) If you want to push updates to secondary servers for a zone, configure DNS notify at the primary server. | Create and manage a notify list for a zone; Understanding zones and zone transfer |
| Use the monitoring features of the DNS console, such as simple or recursive query testing, to verify that DNS servers are operating correctly. | Monitor Servers |
| Troubleshoot common DNS-related problems, if needed. | Troubleshooting DNS |

# Checklist: Deploying DNS for Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Deploying DNS for Active Directory**

| Step | Reference |
|---|---|
| Review best practices for planning your deployment and use of Active Directory with Domain Name System (DNS). | DNS best practices; Active Directory Best practices |
| Understand how Active Directory integration works for DNS servers. | Active Directory integration |
| Consider additional advanced namespace design issues that can affect DNS namespace planning and use before finalizing your Active Directory namespace plans.<br><br>Obtain additional information resources on DNS namespace design available from either the Microsoft Web site or Microsoft Press books before for planning your network. | Namespace planning for DNS |
| Review all planning issues specific to deploying DNS servers on your network for supporting Active Directory. | Deploying DNS |
| Plan your namespace by specifying a domain name standard for your organization. | Namespace planning for DNS |
| If necessary, register your organization's domain name with the Internet domain name registrar. This name is then used to form the basis for DNS names your organization uses on the Internet and might also form the basis for names used internally on your intranet as well.<br><br>**Internet DNS**<br><br>**Name: _____._____**<br><br>**Note**<br><br>• Registering a parent or second-level DNS domain name for your organization with an active Internet domain name registration authority is only required if you plan to use this name externally on the Internet. | Namespace planning for DNS; Namespace planning for DNS; Interoperability issues |
| Plan how to divide up your DNS domain name and network address space into forward and reverse lookup zones as needed. | Zone planning for DNS; Reverse lookup |
| Plan your DNS servers, such as determining how many servers you need to use and where to put them on your network.<br><br>For Active Directory, keep in mind that by default server computers promoted to be domain controllers also operate as DNS servers. | Server planning for DNS |
| (Optional) Test and evaluate server performance before finalizing your server plans. To perform DNS server evaluation, use related performance and event monitoring tools. | Monitoring and Optimizing Servers |
| Review additional migration and interoperability issues, if appropriate. | Migrating servers; Interoperability issues |
| If clients on your network need to be able to resolve external DNS names, consider whether you need to configure and use forwarders at DNS servers on your network. | Using forwarders; Configure a DNS server to use forwarders |
| Determine how to manage DNS client resolver configurations. | Managing Clients |
| If needed, add additional zones for any additional subdomains you need to add and use on your network. | Add a forward lookup zone; Add a reverse lookup zone |

| | |
|---|---|
| (Optional) Modify zone configurations to enable dynamic updates and change zone types to integrate storage of zones in Active Directory. | Allow dynamic updates; Change the zone type; Allow only secure dynamic updates |
| Install and configure domain controllers (or optionally, additional non-integrated DNS servers) and domains if they are to be used for hosting your zones. | Configure a DNS server for use with Active Directory; Install a DNS server; Configure a new DNS server |
| As needed, add delegations in parent zones for any subdomains added based on the previous step.<br><br>For example, if adding sub.example.microsoft.com as a new Active Directory domain, you would need to add to it a delegation at the example.microsoft.com zone. | Create a zone delegation; Delegating zones |
| (For standard primary zones only) Add new DNS servers you need to host your zones to the name server (NS) record at the primary zone/server to make them authoritative for the zone. | Specify other DNS servers as authoritative for a zone |
| Add DNS resource records--such as A, PTR, CNAME, and MX records--as needed to complete zone configurations. | Managing resource records; Add a resource record to a zone; Resource records reference |
| (Optional) Enable WINS lookup for selected zones. | Enable DNS to use WINS resolution |
| (Optional) Install and configure the DNS servers to act as secondary servers for your DNS deployment, if needed. | Install a DNS server; Add a secondary server for an existing zone; Using secondary servers |
| Use the monitoring features of the DNS console, such as simple or recursive query testing, to verify that DNS servers are operating correctly. | Monitor Servers |
| Troubleshoot DNS-related problems, if needed. | Troubleshooting DNS |

# Checklist: Verifying DNS before installing Active Directory

Published: January 21, 2005

Updated: January 13, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Verifying DNS before installing Active Directory**

| Step | Reference |
|---|---|
| Verify that a network connection on the server on which you are installing Active Directory is configured with a DNS server to query for domain names.<br><br>If you are installing Active Directory on a server in an *existing forest,* then the network connection of the server must be configured with one or more DNS servers. If not, then configure the network connection with the IP addresses of one or more DNS servers.<br><br>If you are installing Active Directory to create a *new forest,* and no DNS servers are available on the network, then the Active Directory Installation Wizard can automatically install and configure a local DNS server. The DNS server will be installed on the server where you are running the wizard and the server's preferred DNS server setting will be configured to use the new local DNS server. If you install a local DNS server using the Active Directory Installation Wizard, you may stop using this checklist. | Configure TCP/IP to use DNS; Create a new forest |
| Verify DNS resource records.<br><br>If you are installing Active Directory on a server in an existing forest, verify that the appropriate service (SRV) resource records and the corresponding address (A) resource records exist in DNS.<br><br>• Additional domain controller in an existing domain: _ldap._tcp.dc._mcdcs.*ActiveDirectoryDomainDNSName*<br><br>• First domain controller in a new child domain: _ldap._tcp.dc._mcdcs.*ParentActiveDirectoryDomainDNSName*<br><br>• First domain controller in a new tree: _ldap._tcp.dc._mcdcs.*ForestRootDomainDNSName*<br><br>If the required SRV and A resource records do not exist in DNS, you can add these records. | Verify DNS registration for domain controllers using the nslookup command; Add a host (A) resource record to a zone; Add a resource record to a zone |

| | |
|---|---|
| Verify dynamic update and service (SRV) resource record.<br><br>Verify that the authoritative DNS zone for the domain controller locator (Locator) DNS resource records allows dynamic updates and that the DNS server hosting the zone supports the DNS service (SRV) resource record. To locate the authoritative DNS zone, you will need to locate the primary DNS server hosting that zone.<br><br>If the authoritative DNS zone is not configured to allow dynamic updates, configure the zone to allow dynamic updates.<br><br>**Notes**<br><br>• The authoritative DNS zone for the Locator DNS resource records will have the DNS name of the Active Directory domain or the name of a parent zone. For example, if the DNS name of the Active Directory domain is example.microsoft.com., then the authoritative DNS zone could be one of the following:<br>  • example.microsoft.com<br>  • microsoft.com<br>  • com<br><br>• If you have configured DNS to not dynamically register these resource records in DNS, then these resource records must be added manually. The list of resource records that should be registered by a domain controller are stored in the following location:<br><br>systemroot\System32\Config\Netlogon.dns | Allow dynamic updates |

**Note**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**See Also**
**Concepts**
Checklist: Installing a DNS server

# Checklist: Securing your DNS infrastructure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Securing your DNS infrastructure**

| Step | Reference |
|---|---|
| To prevent anyone outside of your company from obtaining internal network information, use separate DNS servers for internal and Internet name resolution. Your internal DNS namespace should be hosted on DNS servers behind the firewall for your network. Your external, Internet DNS presence should be managed by a DNS server in a perimeter network (also known as DMZ, demilitarized zone, or screened subnet). To provide Internet name resolution for internal hosts, you can have your internal DNS servers use a forwarder to send external queries to your external DNS server. | Using forwarders; "Windows Server 2003 DNS" at the Microsoft Windows Resource Kits Web site |
| To prevent anyone outside of your company from obtaining information about your internal DNS namespace, configure your external router and firewall to only allow DNS traffic between your internal and external DNS servers.<br><br>For the DNS servers in your network that are exposed to the Internet, restrict DNS zone transfers to either DNS servers identified in the zone by name server (NS) resource records or to specific DNS servers in your network.<br><br>**Note**<br><br>• If you are using Microsoft Internet Security and Acceleration (ISA) Server, then you may use block filters to define the traffic allowed through the ISA Server. | Modify zone transfer settings; Microsoft Internet Security and Acceleration (ISA) Server Web site |
| If the server running the DNS Server service is a multihomed computer, then restrict the DNS Server service to only listen on the interface IP address used by its DNS clients and internal servers. For example, a server acting as proxy server may have two network interface cards, one for the intranet and one for the Internet. If that server is also running the DNS Server service, you can configure the service to only listen for DNS traffic on the IP address used by the intranet network interface card. | Restrict a DNS server to listen only on selected addresses |
| If the server running DNS Server service is a domain controller, then use Active Directory access control lists (ACLs) to secure access control of the DNS Server service. | Modify security for the DNS Server service on a domain controller; Best practices for assigning permissions on Active Directory objects; Best practices for permissions and user rights; Understanding Groups |
| Use Active Directory-integrated DNS zones. DNS zones stored in Active Directory can take advantage of Active Directory security features, such as secure dynamic update and the ability to apply Active Directory security settings to DNS servers, zones, and resource records. | Change the zone type; Active Directory integration; Allow only secure dynamic updates; Modify security for a directory-integrated zone; Modify security for a resource record |
| If a DNS zone is not stored in Active Directory, then secure the DNS zone file by modifying permissions on the DNS zone file or on the folder where the zone files are stored. The zone file or folder permissions should be configured to only allow Full Control to the System group. By default, zone files are stored in the systemroot\System32\Dns folder. | Set, view, change, or remove permissions on files and folders |
| Secure the DNS registry keys. The DNS registry keys can be found in the following registry location:<br><br>HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\ | Add users or groups to the Permissions list |

| | |
|---|---|
| Disable recursion on DNS servers that do not respond to DNS clients directly and are not configured with forwarders. A DNS server only requires recursion if it responds to recursive queries from DNS clients or is configured with a forwarder. DNS servers use iterative queries to communicate with each other. | Disable recursion on the DNS server |
| Secure the caches of all DNS servers against names pollution. | Secure server cache against names pollution |
| If you have a private internal DNS namespace, then configure the root hints on your internal DNS servers to only point to the DNS servers hosting your internal root domain and not the DNS servers hosting the Internet root domain. | Update root hints on the DNS server; Updating root hints |

# New ways to do familiar DNS tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring Domain Name System (DNS). The user interface for performing these tasks is different in this version of Windows from the way it was in Microsoft® Windows® NT version 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Install a DNS server on this computer | **Network** in Control Panel | Open Windows Components wizard.<br><br>For more information, see Install a DNS server. |
| Start DNS Manager | **Start** menu shortcut to **DNS Manager**, listed in **Administrative Tools** | Open the DNS console.<br><br>For more information, see Open the DNS console. |
| Start or stop a DNS server | **Services** in Control Panel | Use the DNS console to start or stop a DNS server.<br><br>For more information, see Start or stop a DNS server. |
| Add a remote server to DNS Manager. | **Server** menu in **DNS Manager** | Add the remote server to the DNS console.<br><br>For more information, see Add a server to the DNS console. |

# DNS best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Enter the correct e-mail address of the responsible person for each zone you add to or manage on a DNS server.**

  This field is used by applications to notify DNS administrators for a variety of reasons. For example, query errors, incorrect data returned in a query, and security problems are a few ways in which this field can be used. While most Internet e-mail addresses contain the at sign (**@**) when used in e-mail applications, this symbol must be replaced with a period (**.**) when entering an e-mail address for this field. For example, instead of "administrator@microsoft.com", you would use "administrator.microsoft.com".

  For more information on configuring the responsible person for a zone, see Modify the start of authority (SOA) record for a zone.

- **Be conservative in adding alias records to zones.**

  Avoid using CNAME resource records (RRs) where they are not needed to alias a host name used in a host (A) resource record. Also, ensure that any alias names you use are not used in other RRs.

  DNS allows an owner name of a CNAME resource record to be used as the owner name of the other types of resource records, such as NS, MX, and TXT resource records.

  For more information, see Managing resource records.

- **When designing your DNS network use standard guidelines and, wherever possible, follow preferred practices for managing your DNS infrastructure.**

  DNS was designed to provide a level of fault tolerance for resolving names. If possible, you should have at least two name servers hosting each zone.

- **If you are using Active Directory, use directory-integrated storage for your DNS zones for increased security, fault tolerance, simplified deployment and management.**

  By integrating zones, you can simplify network planning. For example, domain controllers for each of your Active Directory domains correspond in a direct one-to-one mapping to DNS servers. This can simplify planning and troubleshooting DNS and Active Directory replication problems because the same server computers are used in both topologies.

  If you are using directory-integrated storage for your zones, you may select from the different replication scopes that replicate your DNS zone data throughout the directory. If your DNS infrastructure must support Windows 2000 DNS servers, you will use the directory-integrated storage method that replicates DNS zone data to all domain controllers in a domain. If your DNS infrastructure is composed of DNS servers running Windows Server 2003 only, you may also select from replication scopes that replicate your DNS zone data to all DNS servers in the Active Directory forest, all DNS servers in a specified Active Directory domain, or all domain controllers specified in a custom replication scope.

  For more information on directory-integrated DNS zone storage and replication options, see DNS zone replication in Active Directory.

  Any DNS server hosting a directory-integrated zone is a primary DNS server for that zone. This enables a multimaster model where multiple DNS servers may update the same zone data. A multimaster model eliminates a single point of failure associated with a conventional single-master DNS topology, where updates may only be done to a single DNS server for a given zone.

One of the important benefits of directory integration is the support for secure dynamic update of the names within a zone. For more information, see Dynamic update.

- **Consider the use of secondary zones to assist in off-loading DNS query traffic wherever it makes sense.**

  Secondary servers can be used as backups for DNS clients. This allows you to use secondary servers as a means to load balance DNS query traffic on your network and reserve your DNS-enabled primary servers for use only by those clients that need them to perform dynamic registration and updates of their A and PTR RRs.

  For more information about the use of secondary or caching-only servers, see Using secondary servers and Using caching-only servers.

- **If you are planning a large DNS design, such as for a large Internet service provider (ISP) that supports the use of DNS, review the following Request for Comments (RFC) documents published by the Internet Engineering Task Force (IETF).**

| RFC | Title |
|---|---|
| 1912 | Common DNS Operational and Configuration Errors |
| 2182 | Selection and Operation of Secondary DNS Servers |
| 2219 | Use of DNS Aliases for Network Services |

You can obtain these RFCs from the RFC Editor Web site. This Web site is currently maintained by members of the Information Sciences Institute (ISI) who publish a classified listing of all RFCs. RFCs are classified as one of the following: approved Internet standards, proposed Internet standards (circulated in draft form for review), Internet best practices, or For Your Information (FYI) documents.

To see the most current RFCs related to actively defining the DNS standard, see DNS Extension (dsnext) at the Internet Engineering Task Force Web site, which discusses dynamic update, notify, zone transfers, security, and adjustments to address the requirements of IPv6 addressing. Also, see DNS operations (dnsop) at the Internet Engineering Task Force Web site, which provides guidelines for the implementation of DNS by administrators of DNS domains.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# DNS How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install and Configure Servers
- Install and Configure Clients
- Manage Servers
- Optimize Servers
- Monitor Servers
- Add and Remove Zones
- Configure Zone Properties
- Manage Zones
- Manage Resource Records
- Use Aging and Scavenging

# Install and Configure Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and configure servers**

- Install a DNS server
- Configure a DNS server for use with Active Directory
- Verify DNS registration for domain controllers using the nslookup command
- Configure a new DNS server
- Modify security for the DNS Server service on a domain controller
- Add a secondary server for an existing zone
- Install a caching-only DNS server
- Restrict a DNS server to listen only on selected addresses
- Configure a DNS server to use forwarders
- Create the default DNS application directory partitions
- Create a DNS application directory partition
- Enlist a DNS server in a DNS application directory partition
- Remove a DNS server from a DNS application directory partition

# Install a DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a DNS server**

1. Open Windows Components Wizard.
2. In **Components**, select the **Networking Services** check box, and then click **Details**.
3. In **Subcomponents of Networking Services**, select the **Domain Name System (DNS)** check box, click **OK**, and then click **Next**.
4. If prompted, in **Copy files from**, type the full path to the distribution files, and then click **OK**.

   Required files are copied to your hard disk.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.
- It is recommended that you manually configure the computer to use a static IP address. If the DNS server is configured to use DHCP-assigned dynamic addresses, when the DHCP server assigns a new IP address to the DNS server, the DNS clients configured to use that DNS server's previous IP address will be unable to resolve the previous IP address and locate the DNS server.
- After you install a DNS server, you can decide how to administer it and its zones. Although you can use a text editor to make changes to server boot and zone files, this method is not recommended. The DNS console and the DNS command-line tool, dnscmd, simplify maintenance of these files and should be used whenever possible. Once you begin using console-based or command-line management of these files, manually editing them is not recommended. For more information, see Related Topics.
- DNS zones stored in Active Directory can be administered using the DNS console or the dnscmd command-line tool only. These zones cannot be administered using a text editor.
- If you uninstall a DNS server hosting Active Directory-integrated zones, these zones will be saved or deleted according to their storage type. For all storage types, the zone data is stored on other domain controllers or DNS servers and will not be deleted unless the DNS server that you uninstall in the last DNS server hosting that zone.
- If you uninstall a DNS server hosting standard DNS zones, the zone files will remain in the systemroot\system32\Dns directory, but they will not be reloaded if the DNS server is reinstalled. If you create a new zone with the same name as an old zone, the old zone file is replaced with the new zone file.
- When writing DNS server boot and zone data to text files, DNS servers use the Berkeley Internet Name Domain (BIND) file format recognized by legacy BIND 4 servers, not the more recent BIND 8 format.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure a DNS server for use with Active Directory
Configure a new DNS server
DNS Checklists
Server administration using Dnscmd
Open the DNS console

# Configure a DNS server for use with Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a DNS server for use with Active Directory**

- When Active Directory is installed using the Active Directory Installation Wizard, the option to automatically install and configure a local DNS server for use is provided.

  To install Active Directory on this computer, use the Active Directory Installation Wizard. For more information, see Related Topics.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- This procedure only applies to server computers used as domain controllers. If member servers are used as DNS servers, they are not integrated with Active Directory.

- If you choose the Active Directory Installation Wizard option to automatically install and configure a local DNS server, the DNS server is installed on the computer where you are running the wizard and the computer's preferred DNS server setting is configured to use the new local DNS server. You will also want to configure any other computers that will join this domain to use this DNS server's IP address as their preferred DNS server.

- For more information about setting up DNS for Active Directory, see Using the Windows Deployment and Resource Kits.

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Installing a DNS server
DNS requirements for installing Active Directory
Checklist: Deploying DNS for Active Directory
Active Directory integration
To configure a computer's preferred DNS server
Security information for DNS
Microsoft Windows XP Professional Resource Kit Documentation
Microsoft Windows Deployment and Resource Kits
Microsoft Windows Server 2003 Deployment Kit

# Verify DNS registration for domain controllers using the nslookup command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify DNS registration for domain controllers using the nslookup command**

1. Open Command Prompt.

2. Type:

   **nslookup**

3. After the previous command completes, at the nslookup (">") prompt type:

   **set q=***rr_type*

4. After the previous command completes, type:

   **_ldap._tcp.dc._msdcs.***Active_Directory_domain_name*

5. Review the output of the previous SRV query and determine if further action is needed based on whether the previous query succeeded or failed:

   - If the query succeeded, review the registered SRV RRs returned in the query to determine if all domain controllers for your Active Directory domain are included and registered using valid IP addresses.

   - If the query failed, continue troubleshooting dynamic update or DNS server related issues to determine the exact cause of the problem.

| Value | Description |
|---|---|
| **nslookup** | The name of the command-line program. |
| **_ldap._tcp.dc._msdcs.** *Active_Directory_domain_name* | The DNS name configured for use with your Active Directory domain and any of its associated domain controllers.<br><br>For example, if the DNS domain name of your Active Directory domain is example.microsoft.com, type:<br><br>**_ldap._tcp.dc._msdcs.example.microsoft.com**. |
| **set q=** | The command to send the query to the root server. |
| *rr_type* | The resource record (RR) type to apply as a filter for subsequent lookups.<br><br>For example, in this instance, because you want to limit subsequent name queries to filter and return only service location (SRV) RRs that use a specified name, type:<br><br>**set q=srv** |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command**

**prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **nslookup**, press **Enter** and then type **help**

- In some cases, when performing the above procedure, you might see several time-outs reported. This happens when reverse lookup is not configured for DNS servers servicing the same DNS domain as your Active Directory domain.

- The following is an example of command-line output for an Nslookup session, used to verify service location (SRV) resource records that are registered by domain controllers. In this example, the two domain controllers are dc1 and dc2 and are registered for the "example.microsoft.com" domain.

```
C:\nslookup
Default Server:  dc1.example.microsoft.com
Address:  10.0.0.14
 set type=srv
 _ldap._tcp.dc._msdcs.example.microsoft.com
Server:  dc1.example.microsoft.com
Address:  10.0.0.14
_ldap._tcp.dc._msdcs.example.microsoft.com   SRV service location:
          priority       = 0
          weight         = 0
          port           = 389
          svr hostname   = dc1.example.microsoft.com
_ldap._tcp.dc._msdcs.example.microsoft.com   SRV service location:
          priority       = 0
          weight         = 0
          port           = 389
          svr hostname   = dc2.example.microsoft.com
dc1.example.microsoft.com     internet address = 10.0.0.14
dc2.example.microsoft.com     internet address = 10.0.0.15
```

- The **nslookup** command is a standard command-line tool provided in most DNS service implementations. It offers the ability to perform query testing of DNS servers and obtain detailed responses as the command output. This information is useful in troubleshooting name resolution problems, verifying that resource records (RRs) are added or updated correctly in a zone, and debugging other server-related problems.

- Verify that resource records used to register services and critical hosts, such as domain controllers, are correctly added to zones.

  In some cases, you might need to manually add or verify registration of the service location (SRV) resource records used to support Windows Server 2003 domain controllers.

  To add the SRV resource records that have been created for a domain controller, open and view the Netlogon.dns file, created by the Active Directory Installation wizard when a server computer is promoted to a domain controller. It can be found at:

  systemroot\System32\Config\Netlogon.dns

- The resource records used in this file are listed in RFC-compliant text-file format. When verifying these records, look for the following records:

  _ldap._tcp.*Active_Directory_domain_name* IN SRV 0 0 389 *ldap_server_name*
  _ldap._tcp.dc._msdcs.*Active_Directory_domain_name* IN SRV 0 0 389 *domain_controller_name*

In some cases, you might need to modify the Lightweight Directory Access Protocol (LDAP) server name if you are using a non-domain controller as an LDAP server for your network.

- The Net Logon service on each domain controller registers, as appropriate, a number of different DNS resource records with DNS servers. To learn more about these records and how Net Logon updates DNS, obtain additional technical information on DNS available from the Microsoft Web site. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting DNS servers
Nslookup subcommands
Nslookup
Using online resources
DNS RFCs

# Configure a new DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a new DNS server**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. If needed, add and connect to the applicable server in the console.
3. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*Applicable DNS server*

4. On the **Action** menu, click **Configure a DNS Server**.
5. Follow the instructions in the Configure a DNS Server Wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- If the DNS server is running locally, you do not need to perform step 2.
- As a best practice, use the checklist for installing a new DNS server. For more information, see Related Topics.
- When you finish configuring the server, you might need to complete additional tasks, such as enabling dynamic updates for its zones or adding resource records to its zones.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/Config** {*ZoneName*|**..AllZones**} *Property* {**1**|**0**}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |

| | |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/Config** | Specifies the configuration command. |
| {*ZoneName*\|**..AllZones**} | Specifies the name of the zone to be configured. To apply the configuration for all zones hosted by the specified DNS server, type **..AllZones**. |
| *Property* | Specifies the server property or zone property to be configured. There are different properties available for servers and zones. For a list of the available properties, at the command prompt, type: **dnscmd /Config /help**. |
| {**1**\|**0**} | Sets configuration options to either **1** (on) or **0** (off). Note that some server and zone properties must be reset as part of a more complex operation. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/Config/help**

- As a best practice, use the checklist for installing a new DNS server provided in the online Help. For more information, see Related Topics.
- When you finish configuring the server, you might need to complete additional tasks, such as enabling dynamic updates for its zones or adding resource records to its zones.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

DNS Checklists
Install a DNS server
Add a server to the DNS console
Install Windows Support Tools

# Modify security for the DNS Server service on a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify security for the DNS Server service on a domain controller**

1. Open DNS.
2. In the console tree, right-click the applicable server, and then click **Properties**.
3. On the **Security** tab, modify the list of member users or groups that are allowed to administer the applicable server.

## Notes

- To perform this procedure, you must be a member of the DnsAdmins or the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Active Directory access control lists (ACLs) are only supported for the DNS Server service when it is running on a domain controller.
- The security settings determine who can administer the server, but do not affect the ACLs for the zones and resource records hosted on the server. To apply security settings for DNS zones and resource records, see Related Topics.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure a DNS server for use with Active Directory
Modify security for a directory-integrated zone
Modify security for a resource record
Best practices for permissions and user rights
Best practices for permissions and user rights

# Add a secondary server for an existing zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a secondary server for an existing zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*Applicable DNS server*

3. On the **Action** menu, click **New Zone**.
4. Follow the instructions in the New Zone Wizard.

   When adding the zone, select **Secondary zone** as the zone type.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- If the DNS server is running locally, you do not need to perform step 2.
- In order to add a secondary server for an existing zone, you need to have network access to the server acting as the master server for this server and its use of the zone. The master server acts as the source for zone data. It is contacted periodically to assist in renewing the zone and to transfer zone updates whenever they are needed.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName* **/ZoneAdd** *ZoneName* **/Secondary** *MasterIPaddress*...[**/file** *FileName*]

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |

| | |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/ZoneAdd** | Required. Adds a zone. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the secondary zone you are adding. The zone name must be the same as the primary zone from which the secondary zone is created. |
| **/Secondary** | Required. Adds a secondary zone type. |
| *MasterIPaddress...* | Required. Specifies one or more IP addresses for the master servers of the secondary zone, from which it copies zone data. |
| **/file** | Specifies the command to use a file. |
| *FileName* | Specifies the name of the file to use for creating the secondary zone. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/ZoneAdd /help**

- To add a secondary server for an existing zone, you need to have network access to the server acting as the master server for this server and its use of the zone. The master server acts as the source for zone data. It is contacted periodically to assist in renewing the zone and to transfer zone updates whenever they are needed.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Checklist: Installing a DNS server
Install a DNS server
Add a server to the DNS console
Using secondary servers
Understanding zones and zone transfer
Install Windows Support Tools

# Install a caching-only DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a caching-only DNS server**

1. To install a caching-only DNS server, install a DNS server on the server computer.
2. Do not configure the DNS server (as you might normally) to load any zones.
3. Verify server root hints are configured or updated correctly.

   For more information, see Related Topics.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- Caching-only DNS servers do not host any zones and are not authoritative for a particular domain. They are DNS servers that build a local server cache of names learned while performing recursive queries on behalf of their clients. This information is then available from its cache when answering subsequent client queries.
- A caching-only DNS server can be valuable at a site where DNS functionality is needed locally but it is not administratively desirable to create a separate domain or zone for that location.
- It is strongly recommended that, when operating the computer as a DNS server, you manually configure TCP/IP and use a static IP address.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using caching-only servers
Install a DNS server
Open the DNS console
Add a server to the DNS console
Update root hints on the DNS server

# Restrict a DNS server to listen only on selected addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restrict a DNS server to listen only on selected addresses**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. On the **Interfaces** tab, click **Only the following IP addresses**.
5. In **IP address**, type an IP address for the DNS server to be enabled for use, and then click **Add**.
6. As needed, repeat the previous step to specify other server IP addresses to be enabled for use by this DNS server.

   If you need to remove an IP address from the list, click it and then click **Remove**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- By default, the DNS Server service listens for DNS message communications on all configured IP addresses for the server computer.
- Server IP addresses that are added here need to be statically managed. If later you change or remove addresses specified here from TCP/IP configurations maintained at this server, update this list accordingly.
- After you update or revise the list of restricted interfaces, you need to stop and restart the DNS server to apply the new list.
- Restricting the DNS Server service to only listen on specific IP addresses is an effective security measure because only hosts on the same network subnet, or hosts with a router that connects them to that same segment, will have access to the server.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **dnscmd** *ServerName* **/ResetListenAddresses** [*ListenAddress* ...]

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/ResetListenAddresses** | Required. Resets the IP addresses of the interfaces on which the DNS server listens. |
| *ListenAddress...* | Specifies one or more IP addresses for the interfaces on which you want the DNS server to listen. By default, the DNS Server service listens for DNS message communications on all configured IP addresses for the server computer. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd** *ServerName* **/ResetListenAddresses /help**

- Server IP addresses that are added here need to be statically managed. If later you change or remove addresses specified here from TCP/IP configurations maintained at this server, update this list accordingly.

- After you update or revise the list of restricted interfaces, you need to stop and restart the DNS server to apply the new list.

- Restricting the DNS Server service to only listen on specific IP addresses is an effective security measure because only hosts on the same network subnet, or hosts with a router that connects them to that same segment, will have access to the server.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |

| | |
|---|---|
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start or stop a DNS server
Configuring multihomed servers
Install Windows Support Tools
Security information for DNS
Securing the DNS Server service
Securing DNS clients

# Configure a DNS server to use forwarders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a DNS server to use forwarders**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open the DNS snap-in.
2. In the console tree, click the applicable Domain Name System (DNS) server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. On the **Forwarders** tab, click **Edit**.
5. Type the IP address for the fully qualified domain name (FQDN) of a forwarder, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- By default, the DNS server will wait 5 seconds for a response from one forwarder IP address before trying another forwarder IP address. In **Number of seconds before forward queries time out**, you can change the number of seconds the DNS server will wait. When the server has exhausted all forwarders, it will attempt standard recursion.
- If you want the DNS server to only use forwarders and not attempt any further recursion if the forwarders fail, select the **Do not use recursion for this domain** check box.

  You can disable recursion for the DNS server so that it will not perform recursion on any query. If you disable recursion on the DNS server, you will not be able to use forwarders on the same server. For more information about disabling recursion on the DNS server, see Related Links.

- Do not enter a forwarder's IP address more than once in a DNS server's forwarders list because it is a more reliable or geographically closer server. If one of the forwarders is preferred, that forwarder should be ordered first in the series of forwarder IP addresses.
- Problems associated with forwarders often result from inefficient configurations and overuse.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **dnscmd***ServerName***/ZoneAdd***ZoneName***/Forwarder***MasterIPaddress* ... [**/TimeOut** *Time*] [**/Slave**]

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/ZoneAdd** | Required. Adds a zone. |
| *ZoneName* | Required. Specifies the FQDN of the zone. |
| **/Forwarder** | Required. Specifies the command to configure a forwarder. When configuring forwarders on DNS servers running on Active Directory domain controllers, you must use **/DsForwarder** in place of **/Forwarder**. **/DsForwarder** will replicate the forwarder setting to all DNS servers running on domain controllers in an Active Directory domain. |
| *MasterIPaddress*... | Required. Specifies a space-separated list of one or more IP addresses of the DNS servers where queries for *ZoneName* are forwarded. You may specify a list of space-separated IP addresses. |
| **/TimeOut** | Specifies the timeout setting. The timeout setting is the number of seconds before unsuccessful forward queries time out. |
| *Time* | Specifies the value for the **/TimeOut** parameter. The value is in seconds. The default timeout is 5 seconds. |
| **/Slave** | Determines whether or not the DNS server uses recursion when querying for the domain name specified by *ZoneName*. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Links.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/ZoneAdd/help**

- To view a zone added for use as only a conditional forwarder, use the following command:

  **dnscmd***ServerName***/ZoneInfo***ZoneName*

- To reset the forwarder IP addresses for a conditional forwarder domain name, type:

  **dnscmd***ServerName***/ZoneResetMasters***ZoneName* [**/Local**] [*ServerIPs*]

  The **/Local** parameter sets the local master list for Active Directory–integrated forwarders, and the *ServerIPs* parameter is the list of one or more IP addresses of master servers for the zone. Master servers may include DNS servers that host primary or secondary copies of the zone, but they should not include DNS server IP addresses in such a way that two

DNS servers hosting copies of a zone use each other as master servers. Such a configuration would make the forwardering path cyclical.

- To reset the standard, nonconditional forwarder for a DNS server, type:

  **dnscmd***ServerName***/ResetForwarders** [*IPAddress* ...] [ **/**[**No**]**Slave** ] [**/TimeOut** *Time*]

  The parameter *IPAddress* is the IP address where the DNS server will forward unsolvable DNS queries. The **/Slave** parameter sets the DNS server as a subordinate server. The **/NoSlave** parameter (default setting) sets the DNS server as a nonsubordinate server, meaning that it will perform recursion. The **/Timeout** and *Time* parameters are described in the table above.

- You cannot use a domain name in a conditional forwarder if the DNS server hosts a primary, secondary, or stub zone for that domain name. For example, if a DNS server is authoritative for the domain name *example.microsoft.com* (hosts the primary zone for that domain name), you cannot configure that DNS server with a conditional forwarder for *example.microsoft.com*.

- Problems associated with forwarders often result from inefficient configurations and overuse.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
How DNS query works
Using forwarders
Disable recursion on the DNS server
Install Windows Support Tools
Directing queries through forwarders

# Create the default DNS application directory partitions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create the default DNS application directory partitions**

Using the Windows interface

Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable DNS server.
   - **Where?**
   - **DNS**/*applicable DNS server*
3. Click **Create Default Application Directory Partitions**.
4. Follow the instructions to create the DNS application directory partitions.

**Notes**

- By default, only members of the Enterprise Admins group can create a DNS application directory partition.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- By default, the DNS Server service will attempt to locate and create the default DNS application directory partitions in Active Directory. If the DNS Server service is unable to do this, the administrator can manually create the application directory partitions using this procedure.
- If the default DNS application directory partitions are currently available in Active Directory, the option to create the default application directory partitions in the DNS console will not be available.
- The following table describes the options available when creating the DNS default application directory partitions.

| Option | Partition name | Description |
|---|---|---|
| Create a single application directory partition that stores DNS zone data and replicates that data to all DNS servers in the domain | DomainDnsZones.*DnsDomainName* | DNS application directory partition for each domain in the forest. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controllers in the domain. |
| Create a single application directory partition that stores DNS zone data and replicates that data to all DNS servers in the forest | ForestDnsZones.*DnsForestName* | DNS application directory partition for the entire forest. It contains all the DNS servers running on the domain controllers in the forest. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controllers in the forest. |

**Notes**

- By default, the Net Logon service registers domain controller locator (Locator) DNS resource records for any application directory partitions hosted on a domain controller in the same manner as it registers domain controller locator (Locator) DNS resource records for any domain hosted on a domain controller. Once the default DNS application directory partitions are created, Net Logon will register domain controller locator (Locator) DNS resource records on behalf of the domain controller hosting the default DNS application directory partitions.

- For more information about creating and deleting an application directory partition, see Related Topics.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **dnscmd** ServerName **/CreateBuiltinDirectoryPartitions** {**/Domain**|**/Forest**|**/AllDomains**}

| **Value** | **Description** |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/CreateBuiltinDirectoryPartitions** | Required. Creates a default application directory partition. |
| {**/Domain**\|**/Forest**\|**/AllDomains**} | Required. Specifies which default application directory partition to create. Do one of the following:<br><br>To create a default domain-wide DNS application directory partition for the Active Directory domain where the specified DNS server is located, type **/Domain**.<br><br>To create a default forest-wide DNS application directory partition for the Active Directory forest where the specified DNS server is located, type **/Forest**.<br><br>To create a default domain-wide DNS application directory partitions on a DNS server in each domain in the Active Directory forest where the user running this command is logged on, type **/AllDomains**.<br><br>The *ServerName* parameter is ignored for **/AllDomains**. The computer on which this command is run must be joined to a domain in the forest where you want to create all of the default domain-wide application directory partitions. |

**Notes**

- By default, only members of the Enterprise Admins group can create a DNS application directory partition.

- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

- **dnscmd /CreateDirectoryPartition /?**

- By default, the DNS Server service will attempt to locate and create the default DNS application directory partitions in Active Directory. If the DNS Server service is unable to do this, the administrator can manually create the application directory partitions using this procedure.

- If the default DNS application directory partitions are currently available in Active Directory, the option to create the default application directory partitions in the DNS console will not be available.

- The following table describes the options available when creating the DNS default application directory partitions.

| Option | Partition name | Description |
|---|---|---|
| Create a single application directory partition that stores DNS zone data and replicates that data to all DNS servers in the domain | DomainDnsZones.*DnsDomainName* | DNS application directory partition for each domain in the forest. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controllers in the domain. |
| Create a single application directory partition that stores DNS zone data and replicates that data to all DNS servers in the forest | ForestDnsZones.*DnsForestName* | DNS application directory partition for the entire forest. It contains all the DNS servers running on the domain controllers in the forest. DNS zones stored in this application directory partition are replicated to all DNS servers running on domain controllers in the forest. |

**Notes**

- By default, the Net Logon service registers domain controller locator (Locator) DNS resource records for any application directory partitions hosted on a domain controller in the same manner as it registers domain controller locator (Locator) DNS resource records for any domain hosted on a domain controller. Once the default DNS application directory partitions are created, Net Logon will register domain controller locator (Locator) DNS resource records on behalf of the domain hosting the default DNS application directory partitions.

- For more information about creating and deleting an application directory partition, see Related Topics.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add or remove an application directory partition replica
Change zone replication scope
Create or delete an application directory partition
DNS zone replication in Active Directory
Enlist a DNS server in a DNS application directory partition
Install Windows Support Tools

Remove a DNS server from a DNS application directory partition
Security information for DNS

# Create a DNS application directory partition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a DNS application directory partition**

1. Open Command Prompt.
2. Type:

   **dnscmd** *ServerName* **/CreateDirectoryPartition***FQDN*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/CreateDirectoryPartition** | Required. Creates a DNS application directory partition. |
| *FQDN* | Required. Specifies the name of the new DNS application directory partition. You must use a DNS fully qualified domain name (FQDN). |

**Notes**

- By default, only members of the Enterprise Admins group can create a DNS application directory partition.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /CreateDirectoryPartition /?**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DNS zone replication in Active Directory
Enlist a DNS server in a DNS application directory partition
Remove a DNS server from a DNS application directory partition
Change zone replication scope
Install Windows Support Tools
Create or delete an application directory partition
Add or remove an application directory partition replica
Security information for DNS

# Enlist a DNS server in a DNS application directory partition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enlist a DNS server in a DNS application directory partition**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/EnlistDirectoryPartition***FQDN*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/EnlistDirectoryPartition** | Required. Enlists a DNS server in a DNS application directory partition. |
| *FQDN* | Required. Specifies the fully qualified domain name (FQDN) of the DNS application directory partition. |

**Notes**

- To perform this procedure, you must be a member of the DnsAdmins or the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /EnlistDirectoryPartition /?**
- For more information about creating and deleting an application directory partition, see Related Topics.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |

| | |
|---|---|
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DNS zone replication in Active Directory
Create the default DNS application directory partitions
Remove a DNS server from a DNS application directory partition
Change zone replication scope
Install Windows Support Tools
Create or delete an application directory partition
Add or remove an application directory partition replica
Security information for DNS

# Remove a DNS server from a DNS application directory partition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a DNS server from a DNS application directory partition**

1. Open Command Prompt.

2. Type:

   **dnscmd***ServerName***/UnenlistDirectoryPartition***FQDN*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/UnenlistDirectoryPartition** | Required. Removes a DNS server from a DNS application directory partition. |
| *FQDN* | Required. Specifies the fully qualified domain name (FQDN) of the DNS application directory partition from which you are removing the DNS server specified by *ServerName*. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /UnenlistDirectoryPartition /?**

- For more information about creating and deleting an application directory partition, see Related Topics.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |

| | |
|---|---|
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DNS zone replication in Active Directory
Create the default DNS application directory partitions
Enlist a DNS server in a DNS application directory partition
Change zone replication scope
Install Windows Support Tools
Create or delete an application directory partition
Add or remove an application directory partition replica

# Install and Configure Clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and configure clients**

- Configure DNS for static clients
- Enable DNS for DHCP-enabled clients
- Configure the primary DNS suffix for a client computer
- Preload the client resolver cache
- Display and view a client resolver cache using the ipconfig command
- Flush and reset a client resolver cache using the ipconfig command
- Renew DNS client registration using the ipconfig command

# Configure DNS for static clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure DNS for static clients**

- To configure DNS for clients with statically configured IP addresses, you likely need to configure the following:

  1. DNS host name (or names) for the client computer.
  2. Primary and alternate DNS servers that the client uses to assist in resolving DNS domain names.
  3. A list of DNS suffixes to be appended for use in completing unqualified DNS names, which are used for searching and submitting DNS queries at the client for resolution.
  4. Connection-specific dynamic update and registration behavior, such as whether specific network adapters installed at the client dynamically register their configured IP addresses with a DNS server.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- For more information about how to configure DNS for static clients not running Windows XP, see the applicable TCP/IP or DNS documentation provided by the appropriate vendor for these clients.
- By default, the DNS client on Windows XP does not attempt dynamic update over a Remote Access Service (RAS) or virtual private network connection. To modify this configuration, you can modify the advanced TCP/IP settings of the particular network connection or modify the registry. For more information, see Related Topics.
- By default, the DNS client does not attempt dynamic update of top-level domain (TLD) zones. Any zone named with a single-label name is considered a TLD zone, for example, com, edu, blank, my-company. To configure the DNS client to allow the dynamic update of TLD zones, you can use the **Update Top Level Domain Zones** policy setting or modify the registry.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing Clients
Configure TCP/IP to use DNS

# Enable DNS for DHCP-enabled clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable DNS for DHCP-enabled clients**

- To configure DNS for clients with dynamically configured IP addresses provided by a DHCP server, you generally need to configure the following at either the DHCP server or applicable clients:

  1. **DNS host name (or names) for the client computer.**

     For DHCP clients, this must be set at the client computer or assigned during unattended setup.

  2. **Primary and alternate DNS servers that the client uses to assist in resolving DNS domain names.**

     For DHCP clients, this can be set by assigning the DNS server option (option 6) and providing a configured list of ordered IP addresses for the DNS servers that the client is configured to use.

  3. **A list of DNS suffixes to be appended for use in completing unqualified DNS names used for searching and submitting DNS queries at the client for resolution.**

     For DHCP clients, this can be set by assigning the DNS domain name option (option 15) and providing single DNS suffix for the client to append and use in searches. To configure additional DNS suffixes, configure TCP/IP manually for DNS configuration.

  4. **Connection-specific dynamic update and registration behavior, such as whether specific network adapters installed at the client dynamically register their configured IP addresses with a DNS server.**

     For DHCP clients, the default is for client connections to register their configured IP addresses with a DNS server. To modify this behavior at the client, configure TCP/IP manually for DNS configuration.

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- For more information on how to configure other DNS for DHCP clients, see the applicable TCP/IP or DNS documentation provided by the appropriate vendor.

- By default, the DNS client on Windows XP does not attempt dynamic update over a Remote Access Service (RAS) or virtual private network connection. To modify this configuration, you can modify the advanced TCP/IP settings of the particular network connection or modify the registry. For more information, see Related Topics.

- By default, the DNS client does not attempt dynamic update of top-level domain (TLD) zones. Any zone named with a single-label name is considered a TLD zone, for example, com, edu, blank, my-company. To configure the DNS client to allow the dynamic update of TLD zones, you can use the **Update Top Level Domain Zones** policy setting or modify the registry.

- For more information about DHCP options, see "DHCP Options" at the Microsoft Windows Resource Kits Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Assign a scope-based option
Managing Clients
Configure TCP/IP to use DNS
Securing the DNS Server service
Securing DNS clients

# Configure the primary DNS suffix for a client computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the primary DNS suffix for a client computer**

1. Open System in Control Panel.
2. Click the **Computer Name** tab.

   This tab displays the computer name, the workgroup or domain to which it belongs, and a brief description of the computer.
3. Click **Change**, and then click **More**.
4. In **DNS Suffix and NetBIOS Computer Name**, do the following:

   For **Primary DNS suffix of this computer**, specify the DNS suffix to be appended to the name of this computer when completing its fully qualified domain name (FQDN).
5. After applying these changes, restart the computer to initialize it with its new DNS domain name.
6. If the computer has been previously installed and configured as a DNS server, verify that zone authority records are updated.

   These include the start of authority (SOA) and name server (NS) resource records, substituting the new FQDN to replace the single label name previously in use. For more information, see Related Topics.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open System, click **Start**, point to **Settings**, and then click **Control Panel**. In Control Panel, double-click **System**.
- For more information about how to configure the primary DNS suffix for other clients and servers, see the applicable TCP/IP or DNS documentation provided by the appropriate vendor for your other clients.
- By default, the primary DNS suffix portion of a computer's FQDN is the same as the name of the Active Directory domain to which the computer is joined. To allow different primary DNS suffixes, a domain administrator can create a restricted list of allowed suffixes by modifying the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is managed by the domain administrator using Active Directory Service Interfaces (ADSI) or the Lightweight Directory Access Protocol (LDAP).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring DNS client settings
Configure TCP/IP to use DNS
Active Directory naming

# Preload the client resolver cache

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To preload the client resolver cache**

1. At the client computer, open Command Prompt.

2. At the command prompt, type the following command:

   **notepad %systemroot%\system32\drivers\etc\hosts**

3. Using the default entry in the file (a mapping for the local host to the loopback IP address, 127.0.0.1), add additional host name-to-address mappings on separate lines to be preloaded into the resolver cache of the client. For example, you might add:

   10.0.0.1  host-a  host-a.example.microsoft.com

4. On the **File** menu, click **Save**, and then **Exit**.

5. As an option, you can verify that your changes have been updated in the resolver cache by viewing its contents.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To open Notepad, click **Start**, point to **All programs**, point to **Accessories**, and then click **Notepad**.

- Entries you add are always answered first from the local resolver cache and not sent to the DNS query when queries are made locally to resolve these names to host (A) resource records.

- Every line in the Hosts file should contain an IP address followed by one or more host names. For example, you could add a line such as the following with an IP address (10.0.0.1) that maps to more than one DNS host name:

  10.0.0.1  host-a  host-a.example.microsoft.com  host-b.example2.microsoft.com

  Likewise, a single DNS host name can correspond to more than one IP address if each of the addresses is mapped and used in separate lines. For example, you could add lines for the following multi-homed or multi-addressable DNS host computer:

  10.0.0.1  host-a.example.microsoft.com

  10.0.0.2  host-a.example.microsoft.com

  10.0.0.3  host-a.example.microsoft.com

- When multiple names or IP addresses are used in the file, the DNS Client service must be running for all entries to be returned or used in answering queries. If the DNS Client service is not running, only the first entry in the file is used to resolve the query.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
How DNS query works
Display and view a client resolver cache using the ipconfig command
Client features

# Display and view a client resolver cache using the ipconfig command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To display and view a client resolver cache using the ipconfig command**

1. Open Command Prompt.

2. Type:

   **ipconfig /displaydns**

| Value | Description |
|---|---|
| **ipconfig** | The name of the command-line program. |
| **/displaydns** | The command to display a client resolver cache. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **ipconfig /help**

- To pause the display of the command output to one screen at time, type **ipconfig /displaydns|more**.

- The **ipconfig /displaydns** command provides you with a means to view the contents of the DNS client resolver cache, which includes entries preloaded from the local Hosts file, as well as any recently obtained resource records for name queries resolved by the system. This information is used by the DNS Client service to quickly resolve frequently queried names before it queries its configured DNS servers.

- When the **ipconfig /displaydns** command is used to display current resolver cache contents, the resultant output generally includes the local host and loopback IP address (127.0.0.1) mappings. This is because these mappings typically exist in the default (unmodified) contents of the local Hosts file.

- After you can add host mapping entries to the local Hosts file and save the file, these entries are added to the displayed output of this command. For more information, see Related Topics.

- The resolver cache can also support negative caching of unresolved or non-valid DNS names. These entries are added by the DNS Client service when it receives a negative answer from a DNS server for a queried name. The negative result is cached for a short period of time so that it is not again queried, which could cause query performance problems.

  During DNS troubleshooting, you can flush and reset the cache to discard negative entries from the cache and any other dynamically added entries that were not preloaded. For more information, see Related Topics.

- Although the **ipconfig** command is provided for earlier versions of Windows, the **/displaydns** option is only available for use on computers running Windows 2000, Windows XP, or Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Preload the client resolver cache
Flush and reset a client resolver cache using the ipconfig command
Client features
DNS tools
Ipconfig

# Flush and reset a client resolver cache using the ipconfig command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To flush and reset a client resolver cache using the ipconfig command**

1. Open Command Prompt.
2. Type:

   **ipconfig /flushdns**

| Value | Description |
|---|---|
| **ipconfig** | The name of the command-line program. |
| **/flushdns** | The command to flush and reset a client resolver cache. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at a command prompt, type:

  **ipconfig /help**
- The **ipconfig /flushdns** command provides you with a means to flush and reset the contents of the DNS client resolver cache. During DNS troubleshooting, if necessary, you can use this procedure to discard negative cache entries from the cache, as well as, any other dynamically added entries.
- Resetting the cache does not eliminate entries that are preloaded from the local Hosts file. To eliminate those entries from the cache, remove them from this file. For more information, see Related Topics.
- Although the **ipconfig** command is provided for earlier versions of Windows, the **/flushdns** option is only available for use on computers running Windows 2000, Windows XP, or Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Preload the client resolver cache
Display and view a client resolver cache using the ipconfig command
Client features
DNS tools
Ipconfig

# Renew DNS client registration using the ipconfig command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To renew DNS client registration using the ipconfig command**

1. Open Command Prompt.

2. Type:

   **ipconfig /registerdns**

| Value | Description |
|---|---|
| **ipconfig** | The name of the command-line program. |
| **/registerdns** | The command to renew DNS client registration. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **ipconfig /help**

- An additional command to **/registerdns** is to type:

  **ipconfig /registerdns** [*adapter*]

  Where *adapter* is the name of a specific network adapter installed on the computer for which you want to renew or update registrations.

- The **ipconfig /registerdns** command provides you with a means to manually initiate dynamic registration for the DNS names and IP addresses configured at a computer. This option can assist in troubleshooting a failed DNS name registration or in resolving a dynamic update problem between a client and the DNS server without restarting the client.

- By default, the **ipconfig /registerdns** command refreshes all DHCP address leases and registers all related DNS names configured and used by the client computer.

- To learn the names of adapters that you can optionally specify with this command, first type the **ipconfig** command by itself (that is, do not specify any additional parameters). The command output displays all adapters by name that are available for use at the computer.

- Although the **ipconfig** command is provided for earlier versions of Windows, the **/registerdns** option is only available for use on computers running Windows 2000, Windows XP, or Windows Server 2003 operating systems.

- On computers running Windows 2000, Windows XP, or Windows Server 2003 operating systems, the DHCP Client service is used to perform dynamic registrations and updates, regardless of whether the computer uses a DHCP server or static configuration to obtain its IP address.

- If you are troubleshooting a failed DNS dynamic registration for a client computer and its DNS names, it might help to verify that the cause is not related to one of the following commonly known causes for such failures:

    1. The zone where the client requires update or registration is not able to accept dynamic updates.

    2. The DNS servers that the client is configured to use do not support or recognize the DNS dynamic update protocol.

    3. The primary (or directory-integrated) DNS server for the zone refused the update request. This can most likely occur because the client is not permitted under current zone or resource records security sufficient access rights to update its own name.

    4. The server or zone is not available because of other problems, such as a network or server failure.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dynamic update
Allow dynamic updates
Modify security for a directory-integrated zone
Modify security for a resource record
Allow only secure dynamic updates
Troubleshooting dynamic updates
Ipconfig

# Manage Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage servers**


- Open the DNS console
- Start or stop a DNS server
- Add a server to the DNS console
- Remove a server from the DNS console
- Manually update server data files
- Change the boot method used by the DNS server
- Change the name-checking method used by the DNS server
- Restrict NS resource record registration
- Allow NS record creation for specific domain controllers
- Restrict DNS resource records updated by Netlogon

# Open the DNS console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open the DNS console**

- Open DNS.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The DNS console is an administrative tool for managing DNS servers running Windows Server 2003 family operating systems only. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage Servers
DNS Checklists
Managing DNS Servers

# Start or stop a DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start or stop a DNS server**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*Applicable DNS server*

3. On the **Action** menu, point to **All Tasks** and then click one of the following:

   - To start the service, click **Start**.
   - To stop the service, click **Stop**.
   - To interrupt the service, click **Pause**.
   - To stop and then automatically restart the service, click **Restart**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- After you pause or stop the service, on the **Action** menu, in **All Tasks**, you can click **Resume** to immediately resume service.
- When using registry-based configuration, changes are applied to DNS servers only when the DNS Server service is re-initialized. In these cases, if a DNS value is manually changed directly in the registry, the DNS Server service must always be restarted for the new value to be used.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage Servers
DNS Checklists
Managing DNS Servers

# Add a server to the DNS console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a server to the DNS console**

1. Open DNS.
2. On the **Action** menu, click **Connect To DNS Server**.
3. In **Connect to DNS Server**, click either:
   - **This computer**, if the server you want to connect to and manage is located on the same computer you are using to manage it.
   - **The following computer**, if the server you want to connect to and manage is located on a remote computer.

   If you choose to connect to a remote server, specify either its DNS computer name or its IP address.
4. Select the **Connect to the specified computer now** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The DNS console is a Microsoft Management Console (MMC) administrative tool for managing DNS servers running Windows Server 2003 operating systems only. For more information, see Related Topics.
- If you use the Windows Server 2003 DNS console to administer a Windows 2000 DNS server, any new features will not be available when viewing the Windows 2000 DNS server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage Servers
DNS Checklists
Managing DNS Servers

# Remove a server from the DNS console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a server from the DNS console**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Delete**.
4. When prompted to confirm you want to delete this server from the list, click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a server to the DNS console
Start or stop a DNS server

# Manually update server data files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manually update server data files**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Update Server Data Files**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- For standard primary zones, this procedure causes the DNS server to immediately write its in-memory changes out to disk for storage with the zone file. Normally these changes are only written at predefined update intervals and when the DNS server is shut down.
- For Active Directory-integrated zones, this procedure does not apply. To update Active Directory-integrated zones, see the command-line procedure below.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **Dnscmd** *ServerName* **/ZoneUpdateFromDs** *ZoneName*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/ZoneUpdateFromDs** | Required. Updates the zone file with data from Active Directory. |

| | |
|---|---|
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone you are updating. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/ZoneUpdateFromDs/help**

- The command-line procedure updates Active Directory-integrated zones only. For standard zones, see the Windows interface procedure above.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

Related Topics

Manage Zones

Configure Zone Properties

Install Windows Support Tools

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Change the boot method used by the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the boot method used by the DNS server**

1. Open DNS.
2. In the console tree, right-click the applicable DNS server, then click **Properties**.
3. Click the **Advanced** tab.
4. In the **Load zone data on startup** list, select **From registry**, **From file**, or **From Active Directory and registry**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- By default, DNS servers use information stored in the registry to initialize for service and load any zone data for use at the server. As added options, you can configure the DNS server to boot from a file or, in Active Directory environments, you can supplement local registry data with zone data retrieved for directory-integrated zones stored in the Active Directory database. If you use the file method, the file used must be a text file named Boot, located on this computer in the systemroot\Windows\System32\Dns folder.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Optimize Servers
Manage Servers
Modifying server defaults

# Change the name-checking method used by the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the name-checking method used by the DNS server**

1. Open DNS.
2. In the console tree, right-click the applicable DNS server, then click **Properties**.
3. Click the **Advanced** tab.
4. In the **Name checking** list, select **Strict RFC (ANSI)**, **Non RFC (ANSI)**, **Multibyte (UTF8)**, or **All names**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The DNS Server service supports different possible methods for checking the names it receives and processes during normal operations:
  - **Strict RFC (ANSI)** This method strictly enforces RFC-compliant naming rules for all DNS names that the server processes. Names that are not RFC-compliant are treated as erred data by the server.
  - **Non RFC (ANSI)** This method allows names that are not RFC-compliant to be used with the DNS server, such as names that use ASCII characters but are not compliant with RFC host naming requirements.
  - **Multibyte (UTF8)** This method allows names that use the Unicode 8-bit translation encoding scheme, which is a proposed RFC draft, to be used with the DNS server.

    By default, the server uses **Multibyte (UTF8)** to check names.
  - **All names** Allows **Non RFC (ANSI)**, **Strict RFC (ANSI)**, and **Multibyte (UTF8)** naming conventions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checking names and zone data
Optimize Servers
Manage Servers
Modifying server defaults
DNS RFCs

# Restrict NS resource record registration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restrict NS resource record registration**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Registry Editor.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. In Registry Editor, navigate to the following registry key:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters

3. Add the following REG_DWORD value:

   **DisableNSRecordsAutoCreation**

4. Assign a value of **0x1**.

   The REG_DWORD value is a local DNS server setting and applies to DNS zones for which this DNS server is authoritative.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- This procedure restricts NS resource records registered for Active Directory domain controllers only.
- To configure the DNS server to automatically add NS resource records corresponding to itself when loading a zone, you may assign a value of **0x0** or enter no value (default setting). This setting has the same effect as not creating the **DisableNSRecordsAutoCreation** registry entry.
- If you have configured the registry to restrict the DNS server from registering NS resource records for authoritative zones, any existing NS resource records for the authoritative zones located on the DNS server are automatically deleted.
- Regardless of the settings of these registry entries, query responses sent to DNS clients from the authoritative DNS server will indicate that the responses are from an authoritative DNS server.
- The registry key entry described here does not exist by default and must be created and configured according to this procedure.

**Using a command line**

1. Open Command Prompt.

   **Caution**

   - In this procedure you will be editing the registry. Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. Type:

   **dnscmd***ServerName***/Config/DisableNSRecordsAutoCreation 0x1**

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/Config** | Specifies the configuration command. |
| **/DisableNSRecordsAutoCreation** | Determines the local DNS server configuration for registering NS resource records for authoritative zones. |
| **0x1** | Specifies that the DNS server specified in *ServerName* should not add NS resource records for authoritative zones.<br><br>To specify that the DNS server should add NS resource records for all its authoritative zones, type a value of **0x0**. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- This procedure restricts NS resource records registered for Active Directory domain controllers only.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /config /?**

- The DWORD value is a local DNS server setting and applies to authoritative DNS zones hosted on this DNS server.

- Regardless of the settings above, query responses sent to DNS clients from authoritative DNS servers and selected domain controllers will indicate that the responses are from authoritative DNS servers.

- To configure the DNS server to automatically add NS resource records corresponding to itself when loading a zone, you may assign a value of **0x0** or enter no value (default setting). This setting has the same effect as not creating the **DisableNSRecordsAutoCreation** registry entry.

- If you have configured the registry to restrict the DNS server from registering NS resource records for authoritative zones, any existing NS resource records for the authoritative zones located on the DNS server are automatically deleted.

- Regardless of a NS resource record registration setting, query responses sent to DNS clients from the authoritative DNS server will indicate that the responses are authoritative.

- The registry key entries described here do not exist by default and must be created and configured using this procedure.

**Formatting legend**

| **Format** | **Meaning** |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing authority records
Allow NS record creation for specific domain controllers
Install Windows Support Tools
Security information for DNS

# Allow NS record creation for specific domain controllers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow NS resource record creation for specific domain controllers**

1. Open Command Prompt.

   **Important**

   - This procedure applies to domain controller name server (NS) resource records in Active Directory-integrated DNS zones that are hosted on DNS servers configured to not add these resource records for their authoritative zones. For more information, see Related Topics.

2. Type:

   **dnscmd***ServerName* **/Config** *ZoneName* **/AllowNSRecordsAutoCreation** *IpAddresses...*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/Config** | Required. Specifies the configuration command. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| **/AllowNSRecordsAutoCreation** | Required. Specifies that domain controllers entered for *Value* will add their names to NS resource records for the zone specified in *ZoneName*. NS resource records that were previously registered for this zone are not affected. Therefore, you must remove them manually if you do not want them. |
| *IpAddresses...* | Required. Specifies the IP addresses of the domain controllers that will add their names in NS resource records for the zone specified in *ZoneName*. Type a space-separated list of the IP addresses of the DNS servers. For example, 10.0.0.0 172.16.0.0 192.168.0.0. |

**Additional considerations**

- To perform this procedure, you must be a member of the DnsAdmins or the Domain Admins group in Active Directory. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /Config /?**

- If any domain controllers in the specified zone are not listed for *IpAddresses...*, their names will be deleted from the NS resource records for the zone specified in *ZoneName*.

  To specify that all domain controllers are allowed to add their names to NS resource records for the zone, or to clear the list of allowed DNS server IP addresses, type the command and omit *IpAddresses...*:

  **dnscmd***ServerName***/Config***ZoneName***/AllowNSRecordsAutoCreation**

- Regardless of the settings above, query responses sent to DNS clients from authoritative DNS servers and selected domain controllers will indicate that the responses are from authoritative DNS servers.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restrict NS resource record registration

# Restrict DNS resource records updated by Netlogon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restrict the DNS resource records updated by the Net Logon service**

1. Open Registry Editor.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. In Registry Editor, navigate to the following registry key:

   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**

3. Add the following multi-string value (REG_MULTI_SZ) value:

   **DnsAvoidRegisterRecords**

4. In this value, specify the list of data corresponding to the DNS resource records that *should not* be registered for this domain controller by the Net Logon service. The list of data include:

| Data Value | Resource Record Type | DNS Resource Record |
|---|---|---|
| LdapIpAddress | A | <*DnsDomainName*> |
| Ldap | SRV | _ldap._tcp.<*DnsDomainName*> |
| LdapAtSite | SRV | _ldap._tcp.<*SiteName*>._sites.<*DnsDomainName*> |
| Pdc | SRV | _ldap._tcp.pdc._msdcs.<*DnsDomainName*> |
| Gc | SRV | _ldap._tcp.gc._msdcs.<*DnsForestName*> |
| GcAtSite | SRV | _ldap._tcp.<*SiteName*>._sites.gc._msdcs.<*DnsForestName*> |
| DcByGuid | SRV | _ldap._tcp.<*DomainGuid*>.domains._msdcs.<*DnsForestName*> |
| GcIpAddress | A | gc._msdcs.<*DnsForestName*> |
| DsaCname | CNAME | <*DsaGuid*>._msdcs.<*DnsForestName*> |
| Kdc | SRV | _kerberos._tcp.dc._msdcs.<*DnsDomainName*> |
| KdcAtSite | SRV | _kerberos._tcp.<*SiteName*>._sites.dc._msdcs.<*DnsDomainName*> |
| Dc | SRV | _ldap._tcp.dc._msdcs.<*DnsDomainName*> |
| DcAtSite | SRV | _ldap._tcp.<*SiteName*>._sites.dc._msdcs.<*DnsDomainName*> |
| Rfc1510Kdc | SRV | _kerberos._tcp.<*DnsDomainName*> |
| Rfc1510KdcAtSite | SRV | _kerberos._tcp.<*SiteName*>._sites.<*DnsDomainName*> |
| GenericGc | SRV | _gc._tcp.<*DnsForestName*> |

| | | |
|---|---|---|
| GenericGcAtSite | SRV | _gc._tcp.<*SiteName*>._sites.<*DnsForestName*> |
| Rfc1510UdpKdc | SRV | _kerberos._udp.<*DnsDomainName*> |
| Rfc1510Kpwd | SRV | _kpasswd._tcp.<*DnsDomainName*> |
| Rfc1510UdpKpwd | SRV | _kpasswd._udp.<*DnsDomainName*> |

**Important**

- This procedure restricts DNS resource records registered by the Net Logon service for Active Directory domain controllers only.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

- Restart of the Net Logon service is *not* required to make the changes to this value effective. If the **DnsAvoidRegisterRecords** registry key is created or modified while the Net Logon service is stopped or within the first 15 minutes after it is started, then appropriate DNS updates may take place with a short delay; however, the delay is no later than 15 minutes after the Net Logon service starts.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restrict NS resource record registration
Allow NS record creation for specific domain controllers
Security information for DNS

# Optimize Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Optimize servers**


- Enable or disable fast transfer format during zone transfers
- Prevent loading of a zone when bad data is found
- Disable round-robin rotation for multihomed names
- Disable local subnet prioritization for multihomed names
- Restore server default preferences
- Disable recursion on the DNS server
- Update root hints on the DNS server
- Secure server cache against names pollution
- Clear the server names cache
- Modify DNSSEC configuration
- Modify EDNS0 configuration
- Modify UDP message size

# Enable or disable fast transfer format during zone transfers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable fast transfer format during zone transfers**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

You can enable or disable fast transfer format during zone transfers using the Windows interface.

### To enable or disable fast transfer format during zone transfers using the Windows interface

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. Do one of the following:
   - To enable the fast transfer format (the default), in the **Server options** list, clear the **BIND secondaries** check box, and then click **OK**.
   - To disable the fast transfer format, in the **Server options** list, select the **BIND secondaries** check box, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The fast transfer format optimizes zone transfers between Windows-based DNS servers and other DNS server implementations, and it is enabled by default. Zone transfers between Windows-based DNS servers always use the fast transfer format.
- DNS servers running versions of the Berkeley Internet Name Domain (BIND) server implementation prior to version 4.9.4 do not support the fast transfer format. You should enable the **Bind secondaries** option if you are transferring zones to BIND servers running versions earlier than 4.9.4.

**Using a command line**

You can enable or disable fast transfer format during zone transfers using a command line.

### The procedure title

1. Open Command Prompt.

2. Type:

   **dnscmd** *ServerName* **/Config /BindSecondaries** {**1**|**0**}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/Config** | Specifies the configuration command. |
| **/BindSecondaries** | Specifies use of fast transfer format used by legacy Berkeley Internet Name Domain (BIND) servers. |
| {**1**\|**0**} | To disable fast transfer format when transferring a zone to legacy BIND DNS servers, type **1** (on). To enable fast transfer format, type **0** (off). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd***ServerName***/Config/help**

- The fast transfer format optimizes zone transfers between Windows-based DNS servers and other DNS server implementations, and it is enabled by default. Zone transfers between Windows-based DNS servers always use the fast transfer format.

- DNS servers running versions of the BIND server implementation earlier than version 4.9.4 do not support the fast transfer format. You should set **BindSecondaries** to 1 if you are transferring zones to BIND servers running versions earlier than 4.9.4.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Interoperability issues
Tuning advanced server parameters
Install Windows Support Tools

# Prevent loading of a zone when bad data is found

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prevent loading of a zone when bad data is found**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, select the **Fail on load if bad zone data** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Tuning advanced server parameters
Checking names and zone data
Interoperability issues

# Disable round-robin rotation for multihomed names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable round-robin rotation for multihomed names**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, clear the **Enable round robin** check box, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure..
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/Config/RoundRobin** {**1**|**0**}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/Config** | Specifies the configuration command. |
| **/RoundRobin** | Configures round robin rotation. |

| | |
|---|---|
| {**1**\|**0**} | To enable round robin, type **1** (on). To disable round robin, type **0** (off). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/Config/help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure Zone Properties
Configuring round robin
Install Windows Support Tools

# Disable local subnet prioritization for multihomed names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable local subnet prioritization for multihomed names**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, clear the **Enable netmask ordering** check box, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/Config/LocalNetPriority** {**1**|**0**}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/Config** | Specifies the configuration command. |
| **/LocalNetPriority** | Configures netmask ordering. |

| | |
|---|---|
| {**1**\|**0**} | To enable netmask ordering, type **1** (on). To disable netmask ordering, type **0** (off). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/Config/help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure Zone Properties
Prioritizing local subnets
Install Windows Support Tools

# Restore server default preferences

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore server default preferences**

1. Open DNS.
2. In the console tree, right-click the applicable DNS server, and then click **Properties**.
3. Click the **Advanced** tab.
4. Click **Reset to Default**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Clicking **Reset to Default** configures the DNS server with the initial configuration it had following installation. These setting are displayed in the table below.

| Property | Setting |
|---|---|
| Disable recursion | Off |
| BIND secondaries | On |
| Fail on load if bad zone data | Off |
| Enable round robin | On |
| Enable netmask ordering | On |
| Secure cache against pollution | On |
| Name checking | Multibyte (UTF8) |
| Load zone data on startup | From Active Directory and registry |
| Enable automatic scavenging of stale records | Off |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modifying server defaults
Tuning advanced server parameters

# Disable recursion on the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable recursion on the DNS server**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable DNS server, then click **Properties**.
3. Click the **Advanced** tab.
4. In **Server options**, select the **Disable recursion** check box, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- If you disable recursion on the DNS server, you will not be able to use forwarders on the same server.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/Config/NoRecursion** {**1**|**0**}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/Config** | Required. Specifies the configuration command. |
| **/NoRecursion** | Required. Specifies the command to disable recursion. |
| {**1**\|**0**} | Required. To disable recursion, type **1** (off). To enable recursion, type **0** (on). By default, recursion is enabled. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/Config/help**

- If you disable recursion on the DNS server, you will not be able to use forwarders on the same server.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
How DNS query works
Optimize Servers
Manage Servers
Modifying server defaults
Install Windows Support Tools
Security information for DNS
Securing the DNS Server service

# Update root hints on the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To update root hints on the DNS server**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Root Hints** tab.
5. Modify server root hints as follows:

   - To add a root server to the list, click **Add**, and then specify the name and IP address of the server to be added to the list.
   - To modify a root server in the list, click **Edit**, and then specify the name and IP address of the server to be modified in the list.
   - To remove a root server from the list, select it in the list, and then click **Remove**.
   - To copy root hints from a DNS server, click **Copy from server**, and then specify the IP address of the DNS server from which you want to copy a list of root servers to use in resolving queries. These root hints will not overwrite any existing root hints.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Updating root hints
How DNS query works
Security information for DNS
Securing the DNS Server service

# Secure server cache against names pollution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To secure server cache against names pollution**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Server options**, select the **Secure cache against pollution** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The **Secure cache against pollution** option is enabled by default.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
How DNS query works
Tuning advanced server parameters
Security information for DNS
Securing the DNS Server service

# Clear the server names cache

Published: January 21, 2005

Updated: November 4, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To clear the server names cache**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Clear Cache**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type the following command and then press ENTER:

   **Dnscmd** *ServerName* **/clearcache**

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/clearcache** | Required. Specifies the command to clear the DNS server cache. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
How DNS query works
Tuning advanced server parameters
Install Windows Support Tools

# Modify DNSSEC configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify DNSSEC configuration**

1. Open Registry Editor.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. In Registry Editor, navigate to the following registry subkey:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters

3. Add the following DWORD entry:

   **EnableDnsSec**

4. Do one of the following:

   - To exclude DNSSEC resource records in query responses other than responses to requests for SIG, KEY or NXT resource records, assign a value of **0x0**. Appropriate resource records will be included in responses to requests for SIG, KEY, or NXT resource records only.

   - To include the DNSSEC resource records in all query responses (according to RFC 2535), assign a value of **0x2**.

   - To include DNSSEC resource records only in cases where the original client query contained the OPT resource record (according to RFC 2671), assign a value of **0x1** or do not create the value at all. The DNS server behaves the same if the value is **0x1** or if the entry does not appear in the registry.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

- The value of the registry entry **EnableDnsSec** determines whether the DNS server will include or exclude DNSSEC resource records when it receives queries.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using DNS Security Extensions (DNSSEC)

# Modify EDNS0 configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify EDNS0 configuration**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Registry Editor.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. In Registry Editor, navigate to the following registry subkey:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters

3. Add the following DWORD entry:

   **EDNSCacheTimeout**

4. To change the cache timeout, type a value in seconds between 3600 (1 hour) and 15724800 (182 days).

5. In the same registry subkey (Parameters), add the following DWORD entry:

   **EnableEDNSProbes**

6. To configure the DNS server to include an OPT resource record only in response to EDNS0 requests containing OPT resource records, type **0x1** (DWORD).

7. Restart DNS server.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- The value of the registry key **EDNSCacheTimeout** determines how long the DNS server will keep information about the EDNS versions supported by other DNS servers that have responded to a query with a OPT resource record.

**Using a command line**

- Open Command Prompt.

- Type one of the following:
    - **dnscmd***ServerName***/Config/EDNSCacheTimeout***Value***
    - **dnscmd***ServerName***/Config/EnableEDNSProbes***Value***

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/Config** | Required. Specifies the command to configure the DNS server. |
| **/EDNSCacheTimeout** | Required. Specifies the length of time the DNS server remembers the EDNS parameters remote servers report. |
| **/EnableEdnsProbes** | Required. Specifies whether or not the DNS server probes other DNS servers to determine if they support EDNS. |
| Value | Required. For **/EDNSCacheTimeout**, type a value in between 3600 (1 hour) and 15724800 (182 days). For **/EnableEDNSProbes**, type 1 to configure the DNS server to probe other DNS servers and determine if they support EDNS. Type 0 to configure the DNS server to not probe remote servers for EDNS support. If you type 0, the DNS server will continue to use EDNS if other servers request it. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /Config help**
- For information about the current registry setting, type one of the following:
    - **dnscmd /Info /EDNSCacheTimeout**
    - **dnscmd /Info /EnableEDNSProbes**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |

| | |
|---|---|
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Extension Mechanisms for DNS (EDNS0)
Modify UDP message size

# Modify UDP message size

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To modify UDP message size

1. Open Registry Editor.

   **Caution**

   - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

2. In Registry Editor, navigate to the following registry key:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters

3. Add the following DWORD entry:

   **MaximumUdpPacketSize**

4. Type a maximum UDP packet size value in bytes.

   The default value is 1280 bytes. The value must be between 512 and 16384 in decimal format (200 and 4000 in hexadecimal format).

5. Restart DNS server.

### Caution

- When configuring the UDP packet size to be larger than 512 bytes, remember UDP packets must travel through devices other than UDP hosts, such as routers, and these devices may not support UDP packets larger than 512 bytes. It is recommended that you establish the maximum UDP packet length support for all devices, and the path's MTU, if possible, and configure your UDP hosts according to this maximum.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

- For information on discovering the maximum transmission unit (MTU) of an arbitrary Internet path, see Request for Comment (RFC) 1191, "Path MTU Discovery."

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

#### See Also

**Concepts**

User Datagram Protocol (UDP)
Troubleshooting TCP/IP
Monitoring DNS server performance
Troubleshooting DNS servers
DNS RFCs

# Monitor Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor servers**

- Select and enable debug logging options on the DNS server
- Disable debug logging options on the DNS server
- Test a simple query on the DNS server
- Test a recursive query on the DNS server
- Enable automatic query testing on the DNS server
- View the DNS server system event log
- View a DNS server debug log file
- Verify DNS server responsiveness using the nslookup command

# Select and enable debug logging options on the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select and enable debug logging options on the DNS server**

1. Open DNS.
2. In the console tree, right-click the applicable DNS server, then click **Properties**.
3. Click the **Debug Logging** tab.
4. Select **Log packets for debugging**, and then select the events that you want the DNS server to record for debug logging.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- To set the debug logging options, you must first select **Log packets for debugging**.
- To get useful debug logging output you need to select a **Packet direction**, a **Transport protocol** and at least one more option.
- In addition to selecting events for the DNS debug log file, you can specify the file name, location, and maximum file size for the file.
- Using debug logging options slows DNS server performance. For this reason, all debug logging options are disabled by default.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable debug logging options on the DNS server
Using server debug logging options

# Disable debug logging options on the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable debug logging options on the DNS server**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Debug Logging** tab.
5. Clear the **Log packets for debugging** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Select and enable debug logging options on the DNS server
Using server debug logging options

# Test a simple query on the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test a simple query on the DNS server**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Monitoring** tab.
5. Select the **A simple query against this DNS server** check box.
6. Click **Test Now**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Results of the query test appear in **Test results**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Verifying server configuration
Test a recursive query on the DNS server
Enable automatic query testing on the DNS server
How DNS query works

# Test a recursive query on the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test a recursive query on the DNS server**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Monitoring** tab.
5. Select the **A recursive query to other DNS servers** check box.
6. Click **Test Now**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Results of the query test appear in the **Test results** list box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable automatic query testing on the DNS server
Test a simple query on the DNS server
How DNS query works

# Enable automatic query testing on the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable automatic query testing on the DNS server**

1. Open DNS.
2. In the console tree, click the applicable DNS server.

   **Where?**

   - DNS/*applicable DNS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Monitoring** tab.
5. Select the type of testing to be used during automatic query testing. You can select one of both of the following:

   - **A simple query against this DNS server**
   - **A recursive query to other DNS servers**

6. Select the **Perform automatic testing at the following interval** check box.
7. Set the **Test interval** to be used.

   The query tests that you select are performed at regular intervals based on the value of the interval you specify. The default polling interval is 1 minute.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Results of automated query tests appear in **Test results** and are updated after each test interval.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Test a recursive query on the DNS server
Test a simple query on the DNS server
How DNS query works

# View the DNS server system event log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the DNS server system event log**

1. Open DNS.
2. In the console tree, click **DNS Events**.

   **Where?**

   - DNS/*applicable DNS server*/Event Viewer/DNS Events

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- If the DNS server for which you want to view the log is located on another computer, in the console tree, click **DNS**, and then on the **Action** menu, click **Connect to DNS Server**. Click **The following computer**, and then specify the name or IP address of the remote computer.

Related Topics

DNS server log reference

DNS tools

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# View a DNS server debug log file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view a DNS server debug log file**

1. Stop the DNS Server service.
2. Open WordPad.
3. On the **File** menu, click **Open**.
4. In **Open**, for **File name**, specify the path to the DNS server debug log file.

   By default, if the applicable DNS server is running locally, the file and path are as follows:

   systemroot\System32\Dns\Dns.log
5. After you specify the correct path and file, click **Open** to view the log file.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open WordPad, click **Start**, point to **All programs**, point to **Accessories**, and then click **WordPad**.
- To stop the DNS Server service, see Related Topics.
- The location of the DNS.log file is managed using the DNS console. To specify the name and location of the DNS.log file, see Related Topics.
- By default, the Dns.log file is empty if you have not previously enabled debug logging options.
- Debug logging slows DNS server performance and should only be enabled for temporary use.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start or stop a DNS server
View the DNS server system event log
Using server debug logging options
DNS tools
Select and enable debug logging options on the DNS server
Disable debug logging options on the DNS server

# Verify DNS server responsiveness using the nslookup command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Verify DNS server responsiveness**
**To verify DNS server responsiveness**

1. Open a command prompt, type the following command, and then press ENTER:

   **nslookup 127.0.0.1** *server_ip_address*

2. If the server is responding, the name "localhost" is returned.

| Value | Description |
|---|---|
| *server_ip_address* | The Internet Protocol (IP) address of the DNS server. For example, if the IP address of your DNS server is 10.0.0.1, type:<br><br>**nslookup 127.0.0.1 10.0.0.1** |

**Note**

In the previous procedure, the syntax for the nslookup command is: **nslookup***[-option] host server* . This command can be entered from any computer that is running a Microsoft Windows operating system and has network connectivity to the DNS server you wish to query. Only the host entry is required for the command. However, if an IP address or hostname for the server is not supplied, then the default DNS server specified in TCP/IP properties will be queried.

When you enter 127.0.0.1 as the host, this IP address will automatically resolve to the name **localhost** if the DNS Server service is running at the IP address that you specify as the *server*.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting DNS servers
Nslookup subcommands
Nslookup

# Add and Remove Zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Add and remove zones**

- Add a forward lookup zone
- Add a reverse lookup zone
- Add a stub zone
- Delete a zone
- Pause a zone
- Start a zone

# Add a forward lookup zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a forward lookup zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click a DNS server, and then click **New Zone** to open the New Zone Wizard.
3. Follow the instructions to create a new primary, secondary, or stub zone.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/ZoneAdd***ZoneName* {**/Primary**|**/DsPrimary**|**/Secondary**|**/Stub**|**/DsStub**} [**/file** *FileName*] [**/load**] [**/a** *AdminEmail*] [**/DP** *FQDN*]

| Value | Description |
|---|---|
| **dsncmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneAdd** | Required. Adds a zone. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| **/Primary**\|**/DsPrimary**\|**/Secondary**\|**/Stub**\|**/DsStub** | Required. Specifies the type of zone. **/DsPrimary** and **/DsStub** specify an Active Directory-integrated zone type. |
| **/file** | Required for **/Primary**. Specifies a file for the new zone. This parameter is invalid for the **/DsPrimary** zone type. |
| *FileName* | Required for **/Primary**. Specifies the name of the zone file. This parameter is invalid for the **/DsPrimary** zone type. |

| | |
|---|---|
| **/load** | Loads an existing file for the zone. Loads an existing file for the zone. If this parameter is not specified, default zone records are created automatically. This parameter does not apply to **/DsPrimary**. |
| **/a** | Adds an administrator e-mail address for the zone. |
| *AdminEmail* | Specifies the administrator email name for the zone. |
| **/DP** | Adds the zone to an application directory partition. You may also use one of the following:<br>• **/DP /domain** For domain directory partition (replicates to all DNS servers in the domain).<br>• **/DP /forest** For forest directory partition (replicates to all DNS server in the forest).<br>• **/DP /legacy** For legacy directory partition (replicates to all domain controllers in the domain). This setting supports domains using legacy Windows 2000 Server domain controllers. |
| *FQDN* | Specifies the fully qualified domain name of the directory partition. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/ZoneAdd/help**

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Understanding zones and zone transfer
Configure Zone Properties
Manage Zones
Managing Zones
Install Windows Support Tools
Security information for DNS
Securing DNS zones

# Add a reverse lookup zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a reverse lookup zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click a DNS server, and then click **New Zone** to open the New Zone Wizard.
3. Follow the instructions to create a new reverse lookup zone.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd** *ServerName* **/ZoneAdd** *ZoneName* {**/Primary**|**/DsPrimary**} [**/file** *FileName*] [**/load**] [**/a** *AdminEmail*] [**/DP** *FQDN*]

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneAdd** | Required. Adds a zone. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the in-addr.arpa domain for the zone. For example, **20.1.168.192.in-addr.arpa.**. |
| **/Primary**\|**/DsPrimary** | Required. Specifies the type of zone. To specify an Active Directory-integrated zone, type **/DsPrimary**. |
| **/file** | Required for **/Primary**. Specifies a file for the new zone. This parameter is invalid for the **/DsPrimary** zone type. |
| *FileName* | Required for **/Primary**. Specifies the name of the zone file. This parameter is invalid for the **/DsPrimary** zone type. |

| | |
|---|---|
| **/load** | Loads an existing file for the zone. Loads an existing file for the zone. If this parameter is not specified, default zone records are created automatically. This parameter does not apply to **/DsPrimary**. |
| **/a** | Adds an administrator e-mail address for the zone. |
| *AdminEmail* | Specifies the administrator e-mail name for the zone. |
| **/DP** | Adds the zone to an application directory partition. You may also use one of the following:<br>• **/DP /domain** For domain directory partition (replicates to all DNS servers in the domain).<br>• **/DP /forest** For forest directory partition (replicates to all DNS server in the forest).<br>• **/DP /legacy** For legacy directory partition (replicates to all domain controllers in the domain). This setting supports domains using legacy Windows 2000 Server domain controllers. |
| *FQDN* | Specifies the fully qualified domain name of the directory partition. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /ZoneAdd /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing reverse lookup zones
Reverse lookup

Understanding zones and zone transfer
Configure Zone Properties
Manage Zones
Managing Zones
Install Windows Support Tools

# Add a stub zone

Published: January 21, 2005

Updated: August 25, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a stub zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click a DNS server, and then click **New Zone** to open the New Zone Wizard.
3. Follow the instructions to create a new stub zone.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The stub zone cannot be hosted on a DNS server that is authoritative for the same zone.
- If you choose to integrate the stub zone into Active Directory (using Active Directory as the stub zone's storage method), you have the option to specify that the DNS server hosting the stub zone use a local list of master servers when updating the stub zone's resource records, rather than have the DNS server use the master servers list stored in Active Directory. If you want to use a local master servers list, you will need the IP addresses of the local master servers.
- If you want the DNS server hosting a stub zone to use a local list of master servers, see Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName* **/ZoneAdd** *ZoneName* {**/Stub**|**/DsStub**} *MasterIPaddress...* [**/file** *FileName*] [**/load**] [**/DP** *FQDN*]

ValueDescription

**dnscmd** Specifies the name of the command-line tool.

*ServerName* Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.)

**/ZoneAdd** Required. Adds a zone.

*ZoneName* Required. Specifies the fully qualified domain name (FQDN) of the zone.

**/Stub**|**/DsStub**Required. Specifies the type of zone. To specify an Active Directory-integrated stub zone, type **/DsStub**.

*MasterIPaddress...* Required. Specifies one or more IP addresses for the master servers of the stub zone, from which it copies zone data.

**/file** Adds a file for the new zone.

*FileName* Specifies the name of the zone file.

**/load** Loads an existing file for the zone. If this parameter is not specified, default zone records are created automatically.

**/DP** Adds the zone to an application directory partition. You may also use one of the following:

- **/DP /domain** For domain directory partition (replicates to all DNS servers in the domain).
- **/DP /forest** For forest directory partition (replicates to all DNS server in the forest).
- **/DP /legacy** For legacy directory partition (replicates to all domain controllers in the domain). This setting supports domains using legacy domain controllers running Windows 2000 Server.

*FQDN* Specifies the fully qualified domain name of the directory partition.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/ZoneAdd /help**

- The stub zone cannot be hosted on a DNS server that is authoritative for the same zone.
- If you choose to integrate the stub zone into Active Directory (using Active Directory as the stub zone's storage method), you have the option to specify that the DNS server hosting the stub zone use a local list of master servers when updating the stub zone's resource records, rather than have the DNS server use the master servers list stored in Active Directory. If you want to use a local master servers list, you will need the IP addresses of the local master servers.
- If you want the DNS server hosting a stub zone to use a local list of master servers, see Related Topics.

**Formatting legend**

**FormatMeaning**

ItalicInformation that the user must supply

BoldElements that the user must type exactly as shown

Ellipsis (...)Parameter that can be repeated several times in a command line

Between brackets ([])Optional items

Between braces ({}); choices separated by pipe (|). Example: {even|odd}Set of choices from which the user must choose only one

Courier fontCode or program output

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Reload or transfer a stub zone
Configure a stub zone for local master servers
Understanding zones and zone transfer
Using stub zones
Active Directory-integrated stub zones
Install Windows Support Tools

# Delete a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To delete a zone

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open DNS.
2. In the console tree, click the applicable zone.

   **Where?**

   - DNS/*applicable DNS server*/Forward Lookup Zones (*or* Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Delete**.
4. When asked to confirm that you want to delete the zone, click **OK**.

### Caution

- Deleting an Active Directory-integrated zone effectively deletes the zone and eliminates its use at all other DNS servers using the same directory store of zone data.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- This procedure is most often used to delete a secondary copy of a zone, although it can also be used to delete a primary zone.
- Deleting a standard primary zone is usually unnecessary, unless you are redesigning your DNS namespace and the zone is no longer needed or used. In most cases, you can change the zone type if you only want to modify the zone. For more information, see Related Topics.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/ZoneDelete***ZoneName* [**/DsDel**] [**/f**]

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneDelete** | Required. Specifies the command to delete the zone specified by *ZoneName*. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone you are deleting. |
| **/DsDel** | Deletes a the zone from Active Directory. |
| **/f** | Performs the command without asking for confirmation. If you omit this parameter, you are prompted to confirm the deletion of the resource record. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /ZoneDelete /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Pause a zone
Start a zone
Understanding zones and zone transfer
Change the zone type
Configure Zone Properties

Manage Zones
Managing Zones
Install Windows Support Tools

# Pause a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To pause a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable zone.

   **Where?**

   - DNS/*applicable DNS server*/Forward Lookup Zones (*or* Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, click **Pause**, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- By default, zones are started when created or loaded at the server. Once you use this procedure to pause a zone, you must restart the zone before it is available for servicing clients or zone updates.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/ZonePause***ZoneName*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZonePause** | Required. Pauses the zone. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

**Formatting legend**

| **Format** | **Meaning** |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start a zone
Understanding zones and zone transfer
Configure Zone Properties
Manage Zones
Managing Zones

# Start a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable zone.

   **Where?**

   - DNS/*applicable DNS server*/Forward Lookup Zones (*or* Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, click **Start**, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- By default, zones are started when created or loaded at the server. Only zones that have previously been paused need to be restarted.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/ZoneResume***ZoneName*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneResume** | Required. Resumes the hosting of the zone by the DNS server. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone resuming operation. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- By default, zones are started when created or loaded at the server. Only zones that have previously been paused need to be restarted.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Pause a zone
Understanding zones and zone transfer
Configure Zone Properties
Manage Zones
Managing Zones
Install Windows Support Tools

# Configure Zone Properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure zone properties**

- Change the zone type
- Change a zone file name
- Change zone replication scope
- Modify the start of authority (SOA) record for a zone
- Modify zone transfer settings
- Create and manage a notify list for a zone
- Create a zone delegation
- Verify a zone delegation using the nslookup command
- Configure a stub zone for local master servers
- Specify other DNS servers as authoritative for a zone
- Update the master server for a secondary zone
- Enable DNS to use WINS resolution
- Verify WINS as the source for answering a DNS query

# Change the zone type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the zone type**

- Using the Windows interface
- Using a command line

### To change the zone type using the Windows interface

1. Open DNS.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, note the current zone type, and then click **Change**.
4. In **Change Zone Type**, select a zone type other than the current one, and then click **OK**.

**Additional considerations**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- You can select from **Primary zone**, **Secondary zone**, or **Stub zone**. When selecting the secondary or stub zone types, you must specify the IP address of another DNS server to be used as the source for obtaining updated information for the zone.
- If the DNS server computer is operating as a domain controller, the option to change the zone type to **Active Directory-integrated** is available. This option is not otherwise available. When this zone type is selected for use, zone data is stored and replicated as part of the Active Directory database.

| **Note** |
| --- |
| You cannot change the zone type (primary, secondary, or stub) and the method for storing the zone at the same time. You must perform the two operations separately. |

- Changing a zone from secondary to primary type can affect other zone activities, including management of dynamic updates and zone transfers and the use of DNS notify lists to notify other servers about changes in the zone. For more information, see Related Topics.
- Changing a zone from stub to primary type or vice versa is not recommended due to the purpose of stub zones.
- Changing DNS zone type or storage can be time-consuming for large zones.

### To change the zone type using a command line

1. Open Command Prompt.
2. Type the following command, and then press ENTER:

   **dnscmd** *ServerName* **/ZoneResetType** *ZoneName Property* [*MasterIPaddress...*] [**/file** *FileName*] {**/OverWrite_Mem**|**/OverWrite_Ds**|**/DirectoryPartition** *FQDN*}

| Value | Description |
| --- | --- |
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the Domain Name System (DNS) host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of zone. |
| *Property* | Required. One of the following zone types:<br>• **/Primary**<br>Standard primary zone. The **/file***FileName* option is required.<br>• **/DsPrimary**<br>Active Directory Domain Services (AD DS)–integrated primary zone. If the zone is not already a primary zone, you must convert it to a primary zone (using **/Primary**) before you use this option to integrate the zone with AD DS.<br>• **/Secondary**<br>Secondary zone. You must specify at least one *MasterIPaddress*.<br>• **/Stub**<br>Stub zone. You must specify at least one *MasterIPaddress*. If the zone is an AD DS–integrated primary zone, you must use **/DsStub** to convert it to an AD DS–integrated stub zone before using this option.<br>• **/DsStub**<br>Active Directory-integrated stub zone. You must specify at least one *MasterIPaddress*. If the zone is not already a stub zone, you must convert it to a stub zone (using **/Stub**) before using this option to integrate the zone with AD DS. |
| **/file** *FileName* | Required for **/Primary**. Specifies the name of a file for the new zone. This parameter is not valid for the **/DsPrimary** zone type. |
| *MasterIPaddress...* | Required for **/Secondary**, **/Stub** and **/DsStub**. Specifies one or more IP addresses for the master servers of the secondary or stub zone, from which it copies zone data. |
| **/OverWrite_Mem\|/OverWrite_Ds\|**<br>**/DirectoryPartition** *FQDN* | **/OverWrite_Mem** overwrites exisiting DNS data using the data in Active Directory. **/OverWrite_Ds** overwrites Active Directory data with data in DNS. **/DirectoryPartition** stores the new zone in the application directory partition specified by *FQDN*, such as DomainDnsZones.corp.example.microsoft.com. |

| Format | Meaning |
| --- | --- |
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |

| | |
|---|---|
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Additional considerations**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/ZoneResetType/help**

- You can select from primary, secondary or stub zone. When selecting the secondary or stub zone type, you need to specify the IP address of another DNS server to be used as the source for obtaining updated information for the zone.
- If the DNS server computer is operating as a domain controller, the option to change the zone type to **Active Directory-integrated** is available. This option is not otherwise available. When this zone type is selected for use, zone data is stored and replicated as part of the Active Directory database.
- Changing a zone from secondary to primary type can affect other zone activities, including management of dynamic updates and zone transfers, and the use of DNS notify lists to notify other servers about changes in the zone.
- Changing a zone from stub to primary type or vice versa is not recommended due to the purpose of stub zones.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**Additional references**

- Active Directory integration
- Dynamic update
- Pause a zone
- Understanding zones and zone transfer
- Configure Zone Properties
- Manage Zones
- Managing Zones
- Using stub zones
- Install Windows Support Tools

# Change a zone file name

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a zone file name**

1. Open DNS.
2. In the console tree, click the applicable zone.

   **Where?**

   - DNS/*applicable DNS server*/Forward Lookup Zones (*or* Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, in the **Zone file name** text box, type the new file name for this zone.
5. Click **OK** when you have finished entering the new zone file name.

### Caution

- If the zone file name is changed, be sure to update **Zone file name** on other DNS servers that maintain this zone. Otherwise, subsequent zone transfers and updates might fail. This can occur in the following situations:
  - The zone type is primary on this server.
  - The zone type is secondary on this server and this server acts as a source or master server for this zone to other DNS servers that host secondary copies of this zone.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The name of the zone file changes, not the name of the zone. You can use Windows Explorer to view or verify the new zone file name.
- The zone file name is not used for Active Directory-integrated zones because these zones store zone data in the Active Directory database and not a text file on the DNS server computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Migrating servers
Interoperability issues

Configure Zone Properties
Manage Zones
Managing Zones

# Change zone replication scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change zone replication scope**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, note the current zone replication type, and then click **Change**.
4. Select a replication scope for the zone.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Only Active Directory-integrated primary and stub forward lookup zones can change their replication scope. Secondary forward lookup zones cannot change their replication scope.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/ZoneChangeDirectoryPartition***ZoneNameNewPartitionName***

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneChangeDirectoryPartition** | Required. Changes a zone's replication scope. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |

| | |
|---|---|
| *NewPartitionName* | Required. The FQDN of the DNS application directory partition where the zone will be stored. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /ZoneChangeDirectoryPartition /?**
- Only Active Directory-integrated primary forward lookup zones and Active Directory-integrated stub zones can change their replication scope. Secondary forward lookup zones cannot change their replication scope.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DNS zone replication in Active Directory
Create the default DNS application directory partitions
Enlist a DNS server in a DNS application directory partition
Install Windows Support Tools
Create or delete an application directory partition
Add or remove an application directory partition replica
Security information for DNS

# Modify the start of authority (SOA) record for a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify the start of authority (SOA) record for a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. Click the **Start of Authority (SOA)** tab.
4. As needed, modify properties for the start of authority (SOA) record.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The settings applied for the start of authority (SOA) record affect how zone transfers are made between servers. For more information, see Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/RecordAdd***ZoneNameNodeName* [**/Aging**] [*Ttl*] **SOA** *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Adds or modifies a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |

| | |
|---|---|
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this parameter is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time to Live (TTL) setting for the resource record. (The default TTL is defined in SOA resource record). |
| **SOA** | Required. Specifies the type of resource record you are modifying. |
| *PrimSvr* | Required. Specifies the FQDN name of the server that is the primary source for information about the zone. For example, *nameserver.place.example.microsoft.com*. |
| *Admin* | Required. Specifies the name of the DNS administrator for the zone. For example, *postmaster.nameserver.place.example.microsoft.com*. |
| *Serial#* | Required. Specifies the version information for the zone. |
| *Refresh* | Required. Specifies the refresh interval for the zone. The standard setting is 3600 (one hour). |
| *Retry* | Required. Specifies the retry interval for the zone. The standard setting is 600 (ten minutes). |
| *Expire* | Required. Specifies the expire interval for the zone. The standard setting is 86400 (one day). |
| *MinTTL* | Required. Specifies the minimum Time to Live (TTL) value. This is the length of time used by other DNS servers to determine how long to cache information for a record in the zone before expiring and discarding it. The standard setting is 3600 (one hour). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

- To modify any specific SOA resource record's values using dnscmd, you must specify all of the SOA values (*PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*).

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |

| | |
|---|---|
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing authority records
Managing general properties
Understanding zones and zone transfer
Install Windows Support Tools

# Modify zone transfer settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify DNS zone transfer settings**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. Right-click a DNS zone, and then click **Properties**.
3. On the **Zone Transfers** tab, do one of the following:
   - To disable zone transfers, clear the **Allow zone transfers** check box.
   - To allow zone transfers, select the **Allow zone transfers** check box.
4. If you allowed zone transfers, do one of the following:
   - To allow zone transfers to any server, click **To any server**.
   - To allow zone transfers only to the DNS servers listed on the **Name Servers** tab, click **Only to servers listed on the Name Servers tab**.
   - To allow zone transfers only to specific DNS servers, click **Only to the following servers**, and then add the IP address of one or more DNS servers.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- To improve the security of your DNS infrastructure, zone transfers should only be allowed for either the DNS servers in the name server (NS) resource records for a zone or for specified DNS servers. If you allow any DNS server to perform a zone transfer, you are allowing internal network information to be transferred to any host that can contact your DNS server.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/ZoneResetSecondaries***ZoneName* {**/NoXfr** | **/NonSecure** | **/SecureNs** | **/SecureList** [*SecondaryIPAddress*...]}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of zone. |
| **/NoXfr** | Disables zone transfers for the zone. |
| **/NonSecure** | Permits zone transfers to any DNS server. |
| **/SecureNs** | Permits zone transfers only to DNS servers listed in the zone using name server (NS) resource records. |
| **/SecureList** | Permits zone transfers only to DNS servers specified by *SecondaryIPAddress*. |
| *SecondaryIPAddress* | Required, if **/SecureList** is specified. A list of one or more IP addresses for DNS servers that are permitted to obtain zone transfers. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /ZoneResetSecondaries /?**

- To improve the security of your DNS infrastructure, zone transfers should only be allowed for either the DNS servers in the NS resource records for a zone or for specified DNS servers. If you allow any DNS server to perform a zone transfer, you are allowing internal network information to be transferred to any host that can contact your DNS server.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Create and manage a notify list for a zone
Security information for DNS
Checklist: Securing your DNS infrastructure
Command-line reference A-Z

# Create and manage a notify list for a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create and manage a notify list for a zone**

1. Open DNS.
2. In the console tree, click the applicable zone.

   **Where?**

   - DNS/*applicable DNS server*/Forward Lookup Zones (*or* Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Properties**.
4. Click the **Zone Transfers** tab.
5. Click **Notify**.
6. Verify that the **Automatically notify** check box is checked.
7. Select the method to be used for creating a list for notifying other DNS servers when changes to the zone occur. Your options are:

   - Use the default, **Servers listed on the Name Servers tab**, to permit only those servers that appear by IP address on the **Name Servers** tab to be included in the notify list.
   - Select **The following servers** if you want to specify a different notify list to be used instead.

8. If you selected **The following servers** in the previous step, add or remove server IP addresses to form the notify list as needed:

   - To add a server to the notify list, type its IP address in the **IP address** field and click **Add**.
   - To remove a server from the notify list, click the server IP address in the list box and click **Remove**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Changes to the notify list properties are only available on primary zones. For secondary zones, these properties are read-only. DNS Notify is an RFC-compliant extension of the DNS standard defined in RFC 1996, "A Mechanism for Prompt Notification of Zone Changes."
- By default, the DNS server will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding zones and zone transfer
DNS RFCs

# Create a zone delegation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a zone delegation**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open the DNS console.
2. In the console tree, right-click the applicable subdomain, and then click **New Delegation**.
3. Follow the instructions provided in the **New Delegation Wizard** to finish creating the new delegated domain.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- All domains (or subdomains) that appear as part of the applicable zone delegation must be created in the current zone prior to performing delegation as described here. As necessary, use the DNS console to first add domains to the zone before completing this procedure. For more information, see Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/RecordAdd***ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **NS** {*HostName*|*FQDN*}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Specifies the command to add a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is manually updated or removed. |

| | |
|---|---|
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time-To-Live (TTL) setting for the resource record. (The default TTL is defined in SOA resource record). |
| **NS** | Required. Specifies that you are adding a name server (NS) resource record to the zone specified in *ZoneName*. |
| *HostName\|FQDN* | Required. Specifies the host name or FQDN of the new authoritative server. |

See the following examples:

**dnscmd dnssvr1.contoso.com /recordadd test A 10.0.0.5**

**dnscmd /recordadd test.contoso.com test MX 10 mailserver.test.contoso.com**

For more information, see Dnscmd Syntax.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a new domain to a zone

Delegating zones
Troubleshooting zone problems
Specify other DNS servers as authoritative for a zone
Install Windows Support Tools

# Verify a zone delegation using the nslookup command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify a zone delegation using the nslookup command**

1. Open Command Prompt.

2. Type:

   **nslookup***RootServerIpAddress*

3. Then type:

   **nslookup**

4. At the next prompt, type:

   **set norecurse**

5. At the next prompt, type:

   **set q=NS**

6. Type the fully qualified domain name (FQDN) for the failed name.

   Use the trailing period (**.**) when entering the name. If zone delegations are set correctly, a list of name server (NS) resource records for delegated servers should be returned in the response.

7. If the NS query response contains no names or IP addresses for delegated servers, type **q=ns** and query again using the FQDN for the parent zone of the failed name.

   For example, if the failed name you used in the previous step was example.microsoft.com, query for microsoft.com.

8. If the response contains NS resource records, but no host (A) resource records, type **set recurse** and query individually for any of the A resource records of servers listed in the NS resource records.

   If, for each NS resource record you encounter in a zone, you do not find at least one valid IP address in an A resource record, you have a broken delegation.

9. Either fix the broken delegation or retry the delegation test described in the previous step using a different IP address.

   If more than one A resource record or IP address is found, use it to repeat the delegation test described in the previous step. To fix a delegation, add or update an A resource record in the parent zone with a valid IP address for a correct DNS server for the delegated zone.

| Value | Description |
|---|---|
| **nslookup** | The name of the command-line tool. |
| *root_server_ip_address* | The IP address of a valid root server for your network. |
| **set norecursion** | A command to instruct the root server to not perform recursion on your query. |
| **set q=NS** | The command to send the query for NS resource records to the root server. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **nslookup**, press **Enter** and then type **help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delegating zones
Understanding zones and zone transfer

# Configure a stub zone for local master servers

Published: January 21, 2005

Updated: February 17, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

By default, the list of master servers is stored in Active Directory. You can also specify a **local list of master servers**, which is a list of master DNS servers stored in the Windows Registry on the local server. If it exists, this list overrides the list of master servers stored in Active Directory.

By specifying a local list of master servers, an administrator can allow a server to have a different list of master servers than other servers that are hosting the same Active Directory integrated zone.

**To configure a stub zone to use local master servers**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the stub zone, and then click **Properties**.
3. On the **General** tab, under **IP address**, modify the list to display the IP addresses of the local master servers that you want the DNS server to use when loading and updating the stub zone.

   Ensure that the IP addresses of the local master servers are for only those authoritative DNS servers that should be queried to update the records of the stub zone on this server.
4. Select the **Use the list above as a local list of masters** check box, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the DnsAdmins or the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- If the local list of master servers is cleared at a later date, the master servers list from Active Directory is applied and the local list of master servers is deleted.
- When modifications to the master servers list are made and applied on a domain controller hosting the stub zone, the list of master servers for the stub zone is updated in Active Directory.
- The DNS server will keep the master servers list from Active Directory stored in memory.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/ZoneResetMasters***ZoneName* [**/Local**] [*MasterIPaddress*...]

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| **/Local** | Configures the local master list for Active Directory-integrated zones. |
| *MasterIPaddress...* | List of one or more IP addresses of master servers for this zone. Master servers may include the server hosting the primary zone or servers hosting other secondary copies for the zone. To clear the local list of masters, type the command without entering any IP addresses. Ensure that the IP addresses of the local master servers are for only those authoritative DNS servers that should be queried to update the records of the stub zone on this server. |

**Notes**

- To perform this procedure, you must be a member of the DnsAdmins or the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/ZoneResetMasters/help**
- If the local list of master servers is cleared at a later date, the master servers list from Active Directory is applied and the local list of master servers is deleted.
- When modifications to the master servers list are made and applied on a domain controller hosting the stub zone, the list of master servers for the stub zone is updated in Active Directory.
- The DNS server will keep the master servers list from Active Directory stored in memory.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a stub zone
Reload or transfer a stub zone
Using stub zones
Understanding zones and zone transfer
Install Windows Support Tools

# Specify other DNS servers as authoritative for a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify other DNS servers as authoritative for a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. Click the **Name Servers** tab.
4. Click **Add**.
5. Specify additional DNS servers by their names and IP addresses, and then click **Add** to add them to the list.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- To add a name server to the list of authoritative servers for the zone, you must specify both the server's IP address and its DNS name. When entering names, click **Resolve** to resolve the name to its IP address prior to adding it to the list.
- DNS servers specified using this procedure are added to those server IP addresses already present for the existing name server (NS) resource record for the zone. Typically, you might only need to perform this procedure at the primary zone when adding DNS servers to act as secondary servers and also to specify that these servers are known to be authoritative when answering queries for zone data.
- DNS servers automatically add and perform initial configuration of the NS resource record for each new primary zone added to the server.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **NS** {*HostName*|*DomainName*}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |

| | |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Specifies the command to add a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time-To-Live (TTL) setting for the resource record. (The default TTL is defined in SOA resource record). |
| **NS** | Required. Specifies that you are adding a name server (NS) resource record to the zone specified in *ZoneName*. |
| *HostName\|FQDN* | Required. Specifies the host name or FQDN of the new authoritative server. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd/RecordAdd/help**

- DNS servers specified using this procedure are added to those server IP addresses already present for the existing name server (NS) resource record for the zone. Typically, you might only need to perform this procedure at the primary zone when adding DNS servers to act as secondary servers and also to specify that these servers are known to be authoritative when answering queries for zone data.

- DNS servers automatically add and perform initial configuration of the NS resource record for each new primary type added to the server.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |

| | |
|---|---|
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a secondary server for an existing zone
Using secondary servers
Install Windows Support Tools

# Update the master server for a secondary zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To update the master server for a secondary zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable secondary zone, and then click **Properties**.
3. On the **General** tab, in **IP address**, specify the IP address for a new master server, and then click **Add** to update the list.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/ZoneResetMasters***ZoneName* [**/Local**] *MasterIPaddress...*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneResetMasters** | Required. Updates the master servers for a secondary zone. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone you are updating. |
| **/Local** | Specifies the local master list for Active Directory-integrated zones. |
| *MasterIPaddress...* | Required. Specifies the IP addresses of the master servers to be used by the DNS server when updating the specified secondary zones. If you do not specify *ServerIPs*, you are requesting the DNS server to reset the value to an empty list. The request may be denied because a zone must always have at least one master server. *MasterIPaddress...* is required to clear the local master list for a zone. |

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group

might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /ZoneResetMasters /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a secondary server for an existing zone
Using secondary servers
Specify other DNS servers as authoritative for a zone
Install Windows Support Tools

# Enable DNS to use WINS resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable DNS to use WINS resolution**

1. Open DNS.
2. In the console tree, right-click the applicable zone, then click **Properties**.
3. Do one of the following:
   - If the applicable zone is a forward lookup zone, on the **WINS** tab, select the **Use WINS forward lookup** check box. In **IP address**, type the IP address of a WINS server to be used for resolution of names not found in DNS, and then click **Add**.
   - If the applicable zone is a reverse lookup zone, on the **WINS-R** tab, select the **Use WINS-R lookup** check box. In **Domain to append to returned name**, type a name.
4. Select the **Do not replicate this record** check box for this WINS record, if applicable.

   If you are replicating this zone between DNS servers that do not recognize the WINS or WINS-R resource records, click this check box. This prevents these records from being replicated to these other servers during zone transfers. If this zone will be used in performing zone transfers to BIND servers, this is a critical option as BIND will not recognize WINS records.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- When this option is used, specified WINS servers configured in this procedure are used for final referral of names not found in the applicable zone.
- Optionally, click **Advanced** to adjust advanced WINS lookup parameters.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Reverse lookup
Using WINS lookup
WINS lookup integration
Managing reverse lookup zones
Domain and forest functionality
WINS

# Verify WINS as the source for answering a DNS query

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify WINS as the source for answering a DNS query**

1. Open Command Prompt.

2. Type:

   **nslookup**

3. After the previous command completes, at the nslookup ("") prompt type:

   **set debug**

4. Next, either type:

   **set querytype=a**

   if you are testing for a WINS forward lookup, or:

   **set querytype=ptr**

   if you are testing for a WINS-R reverse lookup.

   Respectively, these two commands can be used to set the query type to filter either by host (A) or pointer (PTR) resource records as appropriate for researching either a forward or reverse lookup.

5. Based on whether you are verifying possible WINS sourcing for either a forward or reverse lookup, type the appropriate fully qualified domain name (FQDN).

   For example, if the forward lookup you are tracing is for a domain name host-a.example.microsoft.com, type:

   **host-a.example.microsoft.com.**

   If the reverse lookup you are tracing is for an IP address 10.0.0.1, type:

   **1.0.0.10.in-addr.arpa.**

6. In the response, note whether the server answered authoritatively or non-authoritatively, and note the Time-To-Live (TTL) value.

7. If the server answered authoritatively, repeat the same query you performed in step 4.

8. In the response, note whether the TTL value decreased with the second query answer or if it remained consistent with the TTL value specified in the first query answer.

   If the TTL value decreased for an authoritatively answered query, the source of the query answer is a WINS server.

9. To leave debug mode and return to the command prompt, type **exit**.

| Value | Description |
|---|---|

| | |
|---|---|
| **nslookup** | The name of the command-line program. |
| **set debug** | Enables the **nslookup** command to operate in debug mode, providing extended information in the command output.<br><br>This mode is required to view query response information about whether the source for a query answer is:<br><br>• *authoritative* (from a DNS zone or WINS server database)<br><br>• *non-authoritative* (cached data from previous queries made by the DNS server or loaded from root hints) |
| **set querytype** | Changes the type of information query. More information about types can be found in Request For Comment (RFC) 1035. |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **nslookup**, press **Enter** and then type **help**

- Normally, when a DNS server answers a query from its authoritative zone data, it uses the set minimum or default TTL for the zone or the record-specific TTL value (if one is configured). In so doing, TTLs are decreased in answers the server returns if based on non-authoritative data, such as a cached record at the server.

  WINS lookups present an exceptional case, where an answer received back from a WINS server is cached by the DNS server but is also considered to be authoritative data. In this case, the WINS sourced data is returned to clients as authoritative but ages while in the DNS server names cache, causing the TTL used by the server to decrease over time.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting DNS servers
WINS lookup integration
Using WINS lookup
Nslookup subcommands
Nslookup

# Manage Zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage zones**

- Allow dynamic updates
- Allow only secure dynamic updates
- Initiate a zone transfer at a secondary server
- Reload or transfer a stub zone
- Adjust the refresh interval for a zone
- Adjust the retry interval for a zone
- Adjust the expire interval for a zone
- Modify security for a directory-integrated zone

# Allow dynamic updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow dynamic updates**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone, and then click **Properties**.
3. On the **General** tab, verify that the zone type is either **Primary** or **Active Directory-integrated**.
4. In **Dynamic Updates**, click **Nonsecure and secure**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Dynamic update is an RFC-compliant extension to the DNS standard. The DNS update process is defined in RFC 2136, "Dynamic Updates in the Domain Name System (DNS UPDATES)."

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd** *ServerName* **/Config** {*ZoneName*|**..AllZones**} **/AllowUpdate** {**1**|**0**}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/Config** | Required. Specifies the configuration command. |
| *ZoneName*\|**..AllZones** | Required. Specifies the fully qualified domain name (FQDN) of the zone. To configure all zones hosted on the specified DNS server to allow dynamic updates, type **..AllZones**. |
| **/AllowUpdate** | Required. Specifies the allow update command. |
| **1**\|**0** | Configures dynamic update. To allow dynamic updates, enter a value of 1. To not allow dynamic updates, enter a value of 0. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /Config /help**

- Dynamic update is an RFC-compliant extension to the DNS standard. The DNS update process is defined in RFC 2136, "Dynamic Updates in the Domain Name System (DNS UPDATES)."

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dynamic update
Allow only secure dynamic updates
Install Windows Support Tools
DNS RFCs

# Allow only secure dynamic updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow only secure dynamic updates**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone and click **Properties**.
3. On the **General** tab, verify that the zone type is **Active Directory-integrated**.
4. In **Dynamic Updates**, click **secure only**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Secure dynamic update is supported only for Active Directory-integrated zones. If the zone type is configured differently, you must change the zone type and directory integrate the zone prior to securing it for DNS dynamic updates.
- Dynamic update is an RFC-compliant extension to the DNS standard. The DNS update process is defined in RFC 2136, "Dynamic Updates in the Domain Name System (DNS UPDATE)."
- By default, the DNS server will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/Config** {*ZoneName*|**..AllZones**} **/AllowUpdate 2**

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |

| | |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/Config** | Required. Specifies the configuration command. |
| *ZoneName*\|**..AllZones** | Required. Specifies the fully qualified domain name (FQDN) of the zone. To configure all zones hosted on the specified DNS server to allow dynamic updates, type **..AllZones**. |
| **/AllowUpdate** | Required. Specifies the allow update command. |
| **2** | Required. Configures server to allow secure update. If you exclude the **2**, the zone will be set to perform standard dynamic updates only. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /Config /help**
- Dynamic update is an RFC-compliant extension to the DNS standard. The DNS update process is defined in RFC 2136, "Dynamic Updates in the Domain Name System (DNS UPDATE)."
- By default, the DNS server will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Dynamic update
Change the zone type
Allow dynamic updates
Using DNS servers with DHCP
Install Windows Support Tools
DNS RFCs
Security information for DNS
Securing DNS zones

# Initiate a zone transfer at a secondary server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To initiate a zone transfer at a secondary server**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone and click **Transfer from master**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- This procedure checks to see if the start of authority (SOA) resource record in the secondary zone is the most recent version of the SOA resoure record in the primary zone. If the SOA resource records are synchronized, then there is no zone transfer. If the SOA resource records are not synchronized, then there is a zone transfer.
- By default, the DNS server will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/ZoneRefresh***ZoneName*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneRefresh** | Required. Updates the secondary zone. |
| *ZoneName* | Required. Specifies the name of the secondary zone to update. |

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /ZoneRefresh /help**

- This procedure checks to see if the start of authority (SOA) resource record in the secondary zone is the most recent version of the SOA resoure record in the primary zone. If the SOA resource records are synchronized, then there is no zone transfer. If the SOA resource records are not synchronized, then there is a zone transfer.

- By default, the DNS server will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a secondary server for an existing zone
Using secondary servers
Specify other DNS servers as authoritative for a zone
Install Windows Support Tools

# Reload or transfer a stub zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reload or transfer a stub zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable stub zone, and do one of the following:
   - To reload the stub zone from storage, click **Reload**.
   - To have the DNS server determine if the serial number in the stub zone's SOA resource record has expired and then perform a zone transfer from the stub zone's master server, click **Transfer from Master**.
   - To perform a zone transfer from the stub zone's master server regardless of the serial number in the stub zone's SOA resource record, click **Reload from Master**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName* {**/ZoneReload**|**/ZoneUpdateFromDs**|**/ZoneRefresh**} *ZoneName*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneReload** | Reloads the stub zone. |
| **/ZoneUpdateFromDs** | Reloads the stub zone from Active Directory. |
| **/ZoneRefresh** | Refreshes the stub zone. The DNS server will determine if the serial number in the stub zone's SOA resource record has expired. If the serial number has expired, the DNS server will perform a zone transfer from the stub zone's master server. |

| | |
|---|---|
| *ZoneName* | Required. Specifies the name of the stub zone you want to reload or refresh. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /ZoneReload /help** or **dnscmd /ZoneUpdateFromDs /help** or **dnscmd /ZoneRefresh /help**.
- There is no **dnscmd** command to perform a zone transfer regardless of the SOA resource record's expiration date. To perform this operation, use the Windows interface procedure.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a stub zone
Configure a stub zone for local master servers
Understanding stub zones
Using stub zones
Install Windows Support Tools

# Adjust the refresh interval for a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To adjust the refresh interval for a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone and click **Properties**.
3. On the **General** tab, verify that the zone type is either **Primary** or **Active Directory-integrated**.
4. Click the **Start of Authority (SOA)** tab.
5. In **Refresh interval**, click a time period in minutes, hours, or days, and type a number in the text box.
6. Click **OK** to save the adjusted interval.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- By default, the refresh interval for each zone is set to 15 minutes. The refresh interval is used to determine how often other DNS servers that load and host the zone must attempt to renew the zone.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneName NodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **SOA** *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Adds or modifies a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |

| | |
|---|---|
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this parameter is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time-To-Live (TTL) setting for the resource record. (The default TTL is defined in SOA resource record). |
| **SOA** | Required. Specifies the type of resource record you are modifying. |
| *PrimSvr* | Required. Specifies the FQDN name of the server that is the primary source for information about the zone. For example, *nameserver.place.example.microsoft.com*. |
| *Admin* | Required. Specifies the name of the DNS administrator for the zone. For example, *postmaster.nameserver.place.example.microsoft.com*. |
| *Serial#* | Required. Specifies the version information for the zone. |
| *Refresh* | Required. Specifies the refresh interval for the zone. The standard setting is 900 (15 minutes). |
| *Retry* | Required. Specifies the retry interval for the zone. The standard setting is 600 (ten minutes). |
| *Expire* | Required. Specifies the expire interval for the zone. The standard setting is 86400 (one day). |
| *MinTTL* | Required. Specifies the minimum Time-To-Live value. This is the length of time used by other DNS servers to determine how long to cache information for a record in the zone before expiring and discarding it. The standard setting is 3600 (one hour). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

- To modify any specific SOA resource record's values using dnscmd, you must specify all of the SOA values (*PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*).

- By default, the refresh interval for each zone is set to 15 minutes. The refresh interval is used to determine how often other DNS servers that load and host the zone must attempt to renew the zone.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |

| | |
|---|---|
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding zones and zone transfer
Managing authority records
Install Windows Support Tools

# Adjust the retry interval for a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To adjust the retry interval for a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone and click **Properties**.
3. On the **General** tab, verify that the zone type is either **Primary** or **Active Directory-integrated**.
4. Click the **Start of Authority (SOA)** tab.
5. In **Retry interval**, click an interval in minutes, hours, or days, and type a number in the text box.
6. Click **OK** to save the adjusted interval.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- By default, the retry interval for each zone is set at 10 minutes. The retry interval is used to determine how often other DNS servers that load and host the zone are to retry a request for update of the zone each time the refresh interval occurs.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneName NodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **SOA** *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Adds or modifies a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |

| | |
|---|---|
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this parameter is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time-To-Live (TTL) setting for the resource record. (The default TTL is defined in SOA resource record). |
| **SOA** | Required. Specifies the type of resource record you are modifying. |
| *PrimSvr* | Required. Specifies the FQDN name of the server that is the primary source for information about the zone. For example, *nameserver.place.example.microsoft.com*. |
| *Admin* | Required. Specifies the name of the DNS administrator for the zone. For example, *postmaster.nameserver.place.example.microsoft.com*. |
| *Serial#* | Required. Specifies the version information for the zone. |
| *Refresh* | Required. Specifies the refresh interval for the zone. The standard setting is 3600 (one hour). |
| *Retry* | Required. Specifies the retry interval for the zone. The standard setting is 600 (ten minutes). |
| *Expire* | Required. Specifies the expire interval for the zone. The standard setting is 86400 (one day). |
| *MinTTL* | Required. Specifies the minimum Time-To-Live value. This is the length of time used by other DNS servers to determine how long to cache information for a record in the zone before expiring and discarding it. The standard setting is 3600 (one hour). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

- To modify any specific SOA resource record's values using dnscmd, you must specify all of the SOA values (*PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*).

- By default, the retry interval for each zone is set at 10 minutes. The retry interval is used to determine how often other DNS servers that load and host the zone are to retry a request for update of the zone each time the refresh interval occurs.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |

| | |
|---|---|
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding zones and zone transfer
Managing authority records
Install Windows Support Tools

# Adjust the expire interval for a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To adjust the expire interval for a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone and click **Properties**.
3. On the **General** tab, verify that the zone type is either **Primary** or **Active Directory-integrated**.
4. Click the **Start of Authority (SOA)** tab.
5. In **Expires after**, click an interval in either minutes, hours, or days, and then type a number in the text box.
6. Click **OK** to save the adjusted interval.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- By default, the expire interval for each zone is set to 1 day. The expire interval is used by other DNS servers configured to load and host the zone to determine when zone data expires if not renewed.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **SOA** *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Adds or modifies a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |

| | |
|---|---|
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace for which the SOA record is added. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this parameter is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time-To-Live (TTL) setting for the resource record. (The default TTL is defined in SOA resource record). |
| **SOA** | Required. Specifies the type of resource record you are modifying. |
| *PrimSvr* | Required. Specifies the FQDN name of the server that is the primary source for information about the zone. For example, *nameserver.place.example.microsoft.com*. |
| *Admin* | Required. Specifies the name of the DNS administrator for the zone. For example, *postmaster.nameserver.place.example.microsoft.com*. |
| *Serial#* | Required. Specifies the version information for the zone. |
| *Refresh* | Required. Specifies the refresh interval for the zone. The standard setting is 3600 (one hour). |
| *Retry* | Required. Specifies the retry interval for the zone. The standard setting is 600 (ten minutes). |
| *Expire* | Required. Specifies the expire interval for the zone. The standard setting is 86400 (one day). |
| *MinTTL* | Required. Specifies the minimum Time-To-Live value. This is the length of time used by other DNS servers to determine how long to cache information for a record in the zone before expiring and discarding it. The standard setting is 3600 (one hour). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

- To modify any specific SOA resource record's values using dnscmd, you must specify all of the SOA values (*PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL*).

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |

| | |
|---|---|
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding zones and zone transfer
Managing authority records
Install Windows Support Tools

# Modify security for a directory-integrated zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify security for a directory-integrated zone**

1. Open DNS.
2. In the console tree, click the applicable zone.

   **Where?**

   - DNS/*applicable DNS server*/Forward Lookup Zones (*or* Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, verify that the zone type is **Active Directory-integrated**.
5. On the **Security** tab, modify the list of member users or groups that are allowed to securely update the applicable zone and reset their permissions as needed.

## Notes

- To perform this procedure, you must be a member of the DnsAdmins or the Domain Admins group in Active Directory. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups and Using Run as.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Secure dynamic updates are only supported for zones stored in Active Directory.
- The security settings determine who can administer the zone, but do not affect dynamic updates to the zone. To apply security settings for dynamic updates, see Related Topics.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dynamic update
Active Directory integration
Modify security for a resource record
Changing inherited permissions
Remote Desktop Connection
Security information for DNS
Securing DNS zones

# Manage Resource Records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage resource records**

- Add a host (A) resource record to a zone
- Add a mail exchanger (MX) resource record to a zone
- Add an alias (CNAME) resource record to a zone
- Add a new domain to a zone
- Add a pointer (PTR) resource record to a reverse zone
- Add a resource record to a zone
- Modify an existing resource record in a zone
- Delete a resource record from a zone
- View unsupported resource records in a zone
- Modify security for a resource record

# Add a host (A) resource record to a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a host (A) resource record to a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable forward lookup zone and click **New Host**.
3. In the **Name** text box, type the DNS computer name for the new host.
4. In the **IP address** text box, type the IP address for the new host.
5. As an option, select the **Create associated pointer (PTR) record** check box to create an additional pointer record in a reverse zone for this host, based on the information you entered in **Name** and **IP address**.
6. Click **Add Host** to add the new host record to the zone.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- PTR resource records created automatically when adding an A resource record to a zone will be deleted automatically if the corresponding A resource record is deleted.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **A** *IPAddress*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Required. Adds a new resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |

| | |
|---|---|
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time-To-Live (TTL) setting for the resource record. (The default TTL is defined in SOA resource record). |
| **A** | Required. Specifies the resource record type of the record you are adding. |
| *IPAddress* | Required. The IP address for the host. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

- PTR resource records are deleted automatically if the corresponding A resource record is deleted.

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing resource records
Resource records reference

Install Windows Support Tools

# Add a mail exchanger (MX) resource record to a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a mail exchanger (MX) resource record to a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable forward lookup zone and click **New Mail Exchanger**.
3. In the **Host or domain** text box, type the domain name for which this record is to be used to deliver mail.
4. In the **Mail server** text box, type the DNS host computer name of the mail exchanger or mail server host that delivers mail for the specified domain name.

   As an option, you can click **Browse** to view the DNS namespace for mail exchanger hosts in this domain that have host (A) records already defined.
5. Adjust the **Mail server priority** as needed for this zone.
6. Click **OK** to add the new record to the zone.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [*Ttl*] **MX** *PreferenceMXServerName*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Adds a new resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone in which you will add the new MX resource record. |

| | |
|---|---|
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| **Ttl** | Specifies the Time-To-Live setting for the resource record. |
| **MX** | Required. Specifies the MX resource record type for the record you are adding. |
| *Preference* | Required. Specifies a numeric value (between 0 and 65535) that indicates the mail exchange server's priority with respect to the other mail exchange servers. Lower numbers are given greater preference. |
| *MXServerName* | Required. Specifies the fully qualified domain name (FQDN) for a mail exchanger. The value entered here must resolve to a corresponding host (A) resource record in this zone. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing resource records
Resource records reference

Install Windows Support Tools

# Add an alias (CNAME) resource record to a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an alias (CNAME) resource record to a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Alias**.
3. In the **Alias name** text box, type the alias name.
4. In the **Fully qualified domain name (FQDN) for target host** text box, type the fully qualified domain name of the DNS host computer for which this alias is to be used.

   As an option, you can click **Browse** to search the DNS namespace for hosts in this domain that have host (A) records already defined.
5. Click **OK** to add the new record to the zone.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] **CNAME** *HostName*|*DomainName*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/RecordAdd** | Required. Specifies the command to add a new resource record. |
| *ZoneName* | Required. Specifies the name of the zone where this CNAME resource record will be added. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |

| | |
|---|---|
| **/Aging** | Specifies that this resource record is aged and scavenged. If this parameter is not used, the resource record remains in the DNS database unless it is manually updated or removed. |
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| *Ttl* | Specifies the Time-To-Live (TTL) setting for the resource record. (The default TTL is defined in SOA resource record). |
| **CNAME** | Required. Specifies the resource record type of the record you are adding. |
| *HostName\|DomainName* | Required. Specifies the FQDN of any valid DNS host or domain name in the namespace. For FQDN's, a trailing period (.) is used to fully qualify the name. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing resource records
Resource records reference
Install Windows Support Tools

# Add a new domain to a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a new domain to a zone**

1. Open DNS.
2. In the console tree, click the applicable zone.

   **Where?**

   - DNS/*applicable DNS server*/Forward Lookup Zones (*or* Reverse Lookup Zones)/*applicable zone*

3. On the **Action** menu, click **New Domain**, and then type the name of the new domain without using periods.
4. Click **OK** to add the new domain to the zone.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a zone delegation
Delegating zones
Troubleshooting zone problems
Specify other DNS servers as authoritative for a zone

# Add a pointer (PTR) resource record to a reverse zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a pointer (PTR) resource record to a reverse zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable reverse lookup zone.
3. On the **Action** menu, click **New Pointer**.
4. In the **Host IP number** text box, type the host IP address octet number.
5. In the **Host name** text box, type the fully qualified domain name for the DNS host computer for which this pointer record is to be used to provide reverse lookup (address-to-name resolution).

   As an option, you can click **Browse** to search the DNS namespace for hosts in this domain that have host (A) records already defined.
6. Click **OK** to add the new record to the zone.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- When creating a new A resource record, there is an option to create an associated PTR resource record automatically. PTR resource records created automatically when adding an A resource record to a zone will be deleted automatically if the corresponding A resource record is deleted.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [**Ttl**] **PTR** *HostName*|*DomainName*

ValueDescription

**dnscmd** Specifies the name of the command-line tool.

*ServerName* Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.).

**/RecordAdd** Required. Adds a new resource record.

*ZoneName* Required. Specifies the fully qualified domain name (FQDN) of the zone.

*NodeName* Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node.

**/Aging** Specifies that this resource record is able to be aged and scavenged. If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is manually updated or removed.

**/OpenAcl** Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record.

**Ttl** Specifies the Time-To-Live setting for the resource record.

**PTR** Required. Specifies the resource record type.

*HostName|DomainName* Required. Specifies the FQDN of a resource record located in the DNS namespace. The host you specify is used as the data for answering reverse lookups based on the address information specified by this pointer (PTR) resource record.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

- PTR resource records are deleted automatically if the corresponding A resource record is deleted.

**Formatting legend**

**FormatMeaning**

ItalicInformation that the user must supply

BoldElements that the user must type exactly as shown

Ellipsis (...)Parameter that can be repeated several times in a command line

Between brackets ([])Optional items

Between braces ({}); choices separated by pipe (|). Example: {even|odd}Set of choices from which the user must choose only one

Courier fontCode or program output

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Resource records reference
Managing resource records
Install Windows Support Tools

# Add a resource record to a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a resource record to a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone and click **Other New Records**.
3. In **Select a resource record type** list box, select the type of resource record you want to add.
4. Click **Create Record**.
5. In **New Resource Record**, enter the information needed to complete the resource record.
6. After you specify all of the necessary information for the resource record, click **OK** to add the new record to the zone.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [**Ttl**] *RRType RRData*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Required. Adds a new resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| **/Aging** | Specifies that this resource record is able to be aged and scavenged. If this command is used, this resource record is able to be aged and scavenged. If this command is not used, the resource record remains in the DNS database unless it is manually updated or removed. |

| | |
|---|---|
| **/OpenAcl** | Specifies that new records are open to modification by any user. Without this parameter, only administrators may modify the new record. |
| **Ttl** | Specifies the Time-To-Live setting for the resource record. |
| *RRType RRData* | Required. Specifies the type of resource record to add, followed by the data to be contained in the resource record. For information about each resource record type see the Resource records reference. |

| Resource record type | Resource record data |
|---|---|
| A | *IPAddress* |
| NS,CNAME,MB,MD,PTR,MF,MG,MR | *HostName\|DomainName* |
| MX,RT,AFSDB | *Preference ServerName* |
| SRV | *Priority Weight Port HostName* |
| SOA | *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL* |
| AAAA | *Ipv6Address* |
| TXT,X25,HINFO,ISDN | *String [String]* |
| MINFO,RP | *MailboxName ErrMailboxName* |
| WKS | *Protocol IPAddress Service...* |
| WINS | *MapFlag LookupTimeout CacheTimeout IPAddress...* |
| WINSR | *MapFlag LookupTimeout CacheTimeout RstDomainName* |

| Value | Description |
|---|---|
| *IPAddress* | Specifies a standard IP address. For example, 255.255.255.255. |
| *ipv6Address* | Specifies a standard IPv6 address. For example, 1:2:3:4:5:6:7:8. |
| *Protocol* | Specifies the transmission protocol: UDP or TCP. |
| *Service* | Specifies a standard service. For example, domain, smtp. |
| *HostName\|DomainName* | Specifies the FQDN of a resource record located in the DNS namespace. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

**Formatting legend**

| Format | Meaning |
|---|---|

| | |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

| | |
|---|---|
| *IPAddress* | Specifies a standard IP address. For example, 255.255.255.255. |
| *ipv6Address* | Specifies a standard IPv6 address. For example, 1:2:3:4:5:6:7:8. |
| *Protocol* | Specifies the transmission protocol: UDP or TCP. |
| *Service* | Specifies a standard service. For example, domain, smtp. |
| *HostName\|DomainName* | Specifies the FQDN of a resource record located in the DNS namespace. |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Resource records reference
Managing resource records
Install Windows Support Tools

# Modify an existing resource record in a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify an existing resource record in a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable zone.
3. In the details pane, right-click the resource record you want to modify, and then click **Properties**.
4. In **Properties**, edit the properties that can be modified.

   If necessary, you can view and modify advanced resource record properties for the DNS console. To display advanced properties, on the **View** menu, click **Advanced**.
5. Click **OK** when you have finished modifying the record.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- When **Advanced** view options are enabled, you can modify additional settings for an existing resource record, such as its record-specific Time to Live (TTL).

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/RecordAdd** *ZoneNameNodeName* [**/Aging**] [**/OpenAcl**] [*Ttl*] *RRType RRData*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordAdd** | Required. Adds a new resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |

| | |
|---|---|
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| *RRType RRData* | Required. Specifies the type of resource record to add, followed by the data to be contained in the resource record. For information about each resource record type see the Resource records reference. |

| Resource record type | Resource record data |
|---|---|
| A | *IPAddress* |
| NS,CNAME,MB,MD,PTR,MF,MG,MR | *HostName\|DomainName* |
| MX,RT,AFSDB | *Preference ServerName* |
| SRV | *Priority Weight Port HostName* |
| SOA | *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL* |
| AAAA | *Ipv6Address* |
| TXT,X25,HINFO,ISDN | *String [String]* |
| MINFO,RP | *MailboxName ErrMailboxName* |
| WKS | *Protocol IPAddress Service...* |
| WINS | *MapFlag LookupTimeout CacheTimeout IPAddress...* |
| WINSR | *MapFlag LookupTimeout CacheTimeout RstDomainName* |

| Value | Description |
|---|---|
| *IPAddress* | Specifies a standard IP address. For example, 255.255.255.255. |
| *ipv6Address* | Specifies a standard IPv6 address. For example, 1:2:3:4:5:6:7:8. |
| *Protocol* | Specifies the transmission protocol: UDP or TCP. |
| *Service* | Specifies a standard service. For example, domain, smtp. |
| *HostName\|DomainName* | Specifies the FQDN of a resource record located in the DNS namespace. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordAdd /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |

| | |
|---|---|
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

Related Topics

Add a resource record to a zone

Resource records reference

Install Windows Support Tools

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Delete a resource record from a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a resource record from a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable zone.
3. In the details pane, right-click the resource record you want to delete, and then click **Delete**.
4. When you are asked to confirm that you want to delete the selected resource record, click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- PTR resource records are deleted automatically if the corresponding A resource record is deleted.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/RecordDelete***ZoneNameNodeNameRRTypeRRData* [**/f**]

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.). |
| **/RecordDelete** | Required. Deletes a resource record. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |
| *NodeName* | Required. Specifies the FQDN of the node in the DNS namespace. You can also type the node name relative to the *ZoneName* or @, which specifies the zone's root node. |
| *RRType RRData* | Required. Specifies the type of resource record (RR) to add, followed by the data to be contained in the resource record. For information about each resource record type see the Resource records reference. |

| Resource record type | Resource record data |
|---|---|
| A | *IPAddress* |
| NS,CNAME,MB,MD,PTR,MF,MG,MR | *HostName\|DomainName* |
| MX,RT,AFSDB | *Preference ServerName* |
| SRV | *Priority Weight Port HostName* |
| SOA | *PrimSvr Admin Serial#\ Refresh Retry Expire MinTTL* |
| AAAA | *Ipv6Address* |
| TXT,X25,HINFO,ISDN | *String [String]* |
| MINFO,RP | *MailboxName ErrMailboxName* |
| WKS | *Protocol IPAddress Service...* |
| WINS | *MapFlag LookupTimeout CacheTimeout IPAddress...* |
| WINSR | *MapFlag LookupTimeout CacheTimeout RstDomainName* |

| Value | Description |
|---|---|
| *IPAddress* | Specifies a standard IP address. For example, 255.255.255.255. |
| *ipv6Address* | Specifies a standard IPv6 address. For example, 1:2:3:4:5:6:7:8. |
| *Protocol* | Specifies the transmission protocol: UDP or TCP. |
| *Service* | Specifies a standard service. For example, domain, smtp. |
| *HostName\|DomainName* | Specifies the FQDN of a resource record located in the DNS namespace. |
| **/f** | Specifies that the command is executed without asking for confirmation. If you omit this parameter, you are prompted to confirm the deletion of the resource record. |

**Important**

- If the parameter *RRData* is not specified, all resource records of the same type are deleted.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /RecordDelete /help**

- If the variable *RRData* is not specified, all resource record types matching the previous criteria are deleted.

- PTR resource records are deleted automatically if the corresponding A resource record is deleted.

**Formatting legend**

| Format | Meaning |
| --- | --- |
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a resource record to a zone
Managing authority records
Install Windows Support Tools

# View unsupported resource records in a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view unsupported resource records in a zone**

1. Open DNS.
2. In the console tree, click the applicable zone.
3. In the details pane, right-click the record you want to view, then click **Properties**.
4. In **Properties**, view properties specific to this record.
5. When you have finished viewing the record, click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The DNS console allows you to view unsupported resource records (RRs) in secondary zones that are obtained from other DNS server implementations, such as DNS servers running versions of BIND. These records are not used by DNS servers running Windows Server 2003 and cannot be managed through the DNS console. These types of records include legacy records, such as mail forwarder (MF) and mail domain (MD) resource records (RRs).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a resource record to a zone
Manage Resource Records
Delete a resource record from a zone
Interoperability issues
Install Windows Support Tools

# Modify security for a resource record

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify security for a resource record**

1. Open DNS.
2. In the console tree, click the applicable zone.
3. In the details pane, click the record you want to view.
4. On the **Action** menu, click **Properties**.
5. On the **Security** tab, modify the list of member users or groups that are allowed to securely update the applicable record and reset their permissions as needed.

## Notes

- To perform this procedure, you must be a member of the DnsAdmins or the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Secure dynamic updates are only supported or configurable for resource records in zones stored in Active Directory.
- Security settings applied to resource records only affect dynamic updates. These security settings do not affect who may administer the zone where these resource records are located. For information on the security settings that affect who may administer a zone, see Related Topics.
- Resource records with the same name share the same resource record security settings. The names of resource records are listed in the **Name** column of the DNS console.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dynamic update
Active Directory integration
Modify security for a resource record
Modify security for a directory-integrated zone
Changing inherited permissions
Remote Desktop Connection
Security information for DNS
Securing DNS resource records

# Use Aging and Scavenging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use aging and scavenging**

- Set aging/scavenging properties for the DNS server
- Set aging/scavenging properties for a zone
- Enable automatic scavenging of stale resource records
- Start immediate scavenging of stale resource records
- View when a zone can start scavenging stale records
- Reset scavenging and aging properties for a specified resource record

# Set aging/scavenging properties for the DNS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set aging/scavenging properties for the DNS server**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable DNS server, and then click **Set Aging/Scavenging for all zones**.
3. Select the **Scavenge stale resource records** check box.
4. Modify other aging and scavenging properties as needed.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- Aging and scavenging properties configured by this procedure act as server defaults that apply only toward Active Directory-integrated zones. For standard primary zones, you must set the appropriate properties at the applicable zone.
- Once you apply changes for server aging/scavenging settings, the DNS console prompts you to confirm. You then have the option to apply your changes to new Active Directory-integrated zones only. If needed, you can also apply your changes to existing Active Directory-integrated zones.
- Regardless of whether the **Scavenge stale resource records** check box is selected as described in step 4, for standard primary zones, this feature is disabled unless manually enabled at the applicable zone.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/Config** {**/ScavengingInterval** *Value*|**/DefaultAgingState** *Value*|**/DefaultNoRefreshInterval** *Value*|**/DefaultRefreshInterval** *Value*}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |

| | |
|---|---|
| **/Config** | Required. Specifies the configuration command. |
| **/ScavengingInterval** | Required. Sets the frequency by which the server will perform scavenging for all scavenging-enabled zones. |
| **/DefaultAgingState** | Required. Sets the default aging configuration for all zones on the server. |
| **/DefaultNoRefreshInterval** | Required. Sets the default No-refresh interval for scavenging-enabled zones. |
| **/DefaultRefreshInterval** | Sets the default Refresh interval for scavenging-enabled zones. |
| *Value* | For **/ScavengingInterval**, type a value in hours. The default is 168 (one week). For **/DefaultAgingState**, type 1 to enable aging for new zones when they are created. Type 0 to disable aging for new zones. For **/DefaultNoRefreshInterval**, type a value in hours. The default is 168 (one week). For **/DefaultRefreshInterval**, type a value in hours. The default is 168 (one week). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /Config /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing the aging and scavenging of server data
Understanding aging and scavenging
Use Aging and Scavenging

Install Windows Support Tools

# Set aging/scavenging properties for a zone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set aging/scavenging properties for a zone**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable zone, then click **Properties**.
3. On the **General** tab, click **Aging**.
4. Select the **Scavenge stale resource records** check box.
5. Modify other aging and scavenging properties as needed.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/Config** {*ZoneName*|**..AllZones**} {**/Aging***Value*|**/RefreshInterval** *Value*|**/NoRefreshInterval** *Value*}

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/Config** | Required. Specifies the configuration command. |
| *ZoneName*\|**..AllZones** | Required. Specifies the name of the zone to which you want to set aging and scavenging. To apply the operation to all zones, use **..AllZones**. |
| **/Aging** | Required. Enables aging for zones. |
| **/RefreshInterval** | Required. Specifies the Refresh interval for a scavenging-enabled zone. |

| | |
|---|---|
| **/NoRefreshInterval** | Required. Specifies the No-refresh interval for a scavenging-enabled zone. |
| *Value* | Required. For **/Aging**, type 1 to enable aging. Type 0 to disable aging. For **/RefreshInterval**, type a value in hours. The default setting is 168 hours (one week). For **/NoRefreshInterval**, type a value in seconds. The standard setting is 3600 (one hour). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /Config /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing the aging and scavenging of server data
Understanding aging and scavenging
Use Aging and Scavenging
Install Windows Support Tools

# Enable automatic scavenging of stale resource records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable automatic scavenging of stale resource records**

1. Open DNS.
2. In the console tree, right-click the applicable DNS server, then click **Properties**.
3. Click the **Advanced** tab.
4. Select the **Enable automatic scavenging of stale records** check box.
5. To adjust the **Scavenging period**, select from the drop-down list an interval in either hours or days, and then type a number in the text box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing the aging and scavenging of server data
Understanding aging and scavenging
Use Aging and Scavenging

# Start immediate scavenging of stale resource records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start immediate scavenging of stale resource records**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, right-click the applicable DNS server, then click **Scavenge Stale Resource Records**.
3. When asked to confirm that you want to scavenge all stale resource records on the server, click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***ServerName***/StartScavenging**

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/StartScavenging** | Required. Initiates resource record scavenging. |

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.

- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /StartScavenging /help**

**Formatting legend**

| Format | Meaning |
| --- | --- |
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing the aging and scavenging of server data
Understanding aging and scavenging
Use Aging and Scavenging
Install Windows Support Tools

# View when a zone can start scavenging stale records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view when a zone can start scavenging stale records**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. On the **View** menu, click **Advanced**.
3. Right-click the applicable zone, then click **Properties**.
4. On the **General** tab, click **Aging**.
5. Under **Refresh interval**, view when the zone is first eligible to be scavenged for stale resource records.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- The start scavenging date and time stamp are used to determine when zone scavenging starts. For more information, see Related Topics.
- After the start scavenging date and time stamp are reached, scavenging can occur only if the **Scavenge stale resource records** check box is selected. If the check box is cleared, scavenging for the zone cannot be performed.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **dnscmd***Server***/ZoneInfo***ZoneName***RefreshInterval**

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line tool. |
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/ZoneInfo** | Required. Displays configuration information. |
| *ZoneName* | Required. Specifies the fully qualified domain name (FQDN) of the zone. |

| | |
|---|---|
| **RefreshInterval** | Required. Specifies the configuration property that displays when the zone is first eligible to be scavenged for stale resource records. The output value is in hours. The default setting is 168 hours (one week). |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /ZoneInfo /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Understanding aging and scavenging
Use Aging and Scavenging
Install Windows Support Tools

# Reset scavenging and aging properties for a specified resource record

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reset scavenging and aging properties for a specified resource record**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open DNS.
2. In the console tree, click the applicable zone.
3. In the details pane, double-click the resource record for which you want to reset scavenging and aging properties.
4. Depending on the how the resource record was originally added to the zone, do one of the following:
    - If the record was added dynamically using dynamic update, you can clear the **Delete this record when it becomes stale** check box to prevent its aging or potential removal during the scavenging process. If dynamic updates to this record continue to occur, the DNS server will always reset this check box so that the dynamically updated record can be deleted.
    - If you added the record statically, you can select **Delete this record when it becomes stale** check box to permit its aging or potential removal during the scavenging process.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DNS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DNS**.
- This procedure is only necessary for resource records that are dynamically registered. For records that you manually add to a zone, a time stamp value of zero always applies to the record, excluding it from the scavenging process.
- Scavenging and aging properties for NS and SOA resource records are reset in the properties for the zone, not the properties for the resource record.

**Using a command line**

1. Open Command Prompt.
2. Type: **dnscmd** *ServerName* **/Config** {*ZoneName*|**..AllZones**} **/ScavengingInterval** *Value*

| Value | Description |
|---|---|
| **dnscmd** | Specifies the name of the command-line program. |

| | |
|---|---|
| *ServerName* | Required. Specifies the DNS host name of the DNS server. You can also type the IP address of the DNS server. To specify the DNS server on the local computer, you can also type a period (.) |
| **/Config** | Required. Specifies the configuration command. |
| *ZoneName*\|**..AllZones** | Required. Specifies the fully qualified domain name (FQDN) of the zone. To configure all zones hosted on the specified DNS server to allow dynamic updates, type **..AllZones**. |
| **/ScavengingInterval** | Required. Sets the scavenging interval. |
| *Value* | Required. The new value for the scavenging interval, specified in hours. The default is 168 (one week). |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- This procedure requires the Dnscmd Windows support tool. For information about installing Windows support tools, see Related Topics.
- To view the complete syntax for this command, at a command prompt, type:

  **dnscmd /Config /help**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing the aging and scavenging of server data
Understanding aging and scavenging
Use Aging and Scavenging
Install Windows Support Tools

# DNS Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Domain Name System (DNS) and the DNS Server service, as well as details about supporting software provided for DNS clients running under Microsoft operating systems.

- DNS Overview
- Understanding DNS
- Deploying DNS
- Administering DNS
- DNS Resources

# DNS Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS overview**

This section covers:

- DNS defined
- DNS tools
- Server features
- Client features
- Security information for DNS
- New features for DNS

# DNS defined

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS defined**

DNS is an abbreviation for Domain Name System, a system for naming computers and network services that is organized into a hierarchy of domains. DNS naming is used in TCP/IP networks, such as the Internet, to locate computers and services through user-friendly names. When a user enters a DNS name in an application, DNS services can resolve the name to other information associated with the name, such as an IP address.

For example, most users prefer a friendly name such as *example.microsoft.com* to locate a computer such as a mail or Web server on a network. A friendly name can be easier to learn and remember. However, computers communicate over a network by using numeric addresses. To make the use of network resources easier, name systems such as DNS provide a way to map the user-friendly name for a computer or service to its numeric address.

The following figure shows a basic use of DNS, which is finding the IP address of a computer based on its name.

In this example, a client computer queries a DNS server, asking for the IP address of a computer configured to use *host-a.example.microsoft.com* as its DNS domain name. Because the DNS server is able to answer the query based on its local database, it replies with an answer containing the requested information, which is a host (A) resource record that contains the IP address information for *host-a.example.microsoft.com*.

The example shows a simple DNS query between a single client and DNS server. In practice, DNS queries can be more involved than this and include additional steps not shown here. For more information, see How DNS query works.

**Note**

- For additional background information about other DNS concepts, see Understanding DNS.

# DNS tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS tools**

There are a number of utilities for administering, monitoring, and troubleshooting both DNS servers and clients. These utilities include:

- The DNS console, which is part of Administrative Tools.
- Command-line utilities, such as Nslookup, which can be used to troubleshoot DNS problems.
- Logging features, such as the DNS server log, which can be viewed using the DNS console or Event Viewer. File-based logs can also be used temporarily as an advanced debugging option to log and trace selected service events.
- Performance monitoring utilities, such as statistical counters to measure and monitor DNS server activity with System Monitor.
- Windows Management Instrumentation (WMI), a standard technology for accessing management information in an enterprise environment.
- Platform Software Developer Kit (SDK).

**The DNS console**

The primary tool that you use to manage DNS servers is the DNS console, which is located in the Administrative Tools folder in the Start menu's Programs folder. The DNS console can be used on its own or as a Microsoft Management Console (MMC) , further integrating DNS administration into your total network management.

The DNS console can only be used after DNS is installed on the server. You can use the DNS console to perform these basic administrative server tasks:

1. Performing initial configuration of a new DNS server.
2. Connecting to and managing a local DNS server on the same computer, or remote DNS servers on other computers.
3. Adding and removing forward and reverse lookup zones as needed.
4. Adding, removing, and updating resource records in zones.
5. Modifying how zones are stored and replicated between servers.
6. Modifying how servers process queries and handle dynamic updates.
7. Modifying security for specific zones or resource records.

In addition, you can also use the DNS console to perform the following tasks:

- Perform maintenance on the server. You can start, stop, pause, or resume the server, or manually update server data files.
- Monitor the contents of the server cache and, as needed, clear it.
- Tune advanced server options.
- Configure and perform aging and scavenging of stale resource records stored by the server.

In addition, you can also operate the DNS console from a workstation to remotely administer DNS servers. For more information, see Windows Server 2003 Administration Tools Pack.

**Important**

- The DNS console can only be used to manage DNS servers running Microsoft® Windows® and cannot be used to manage other DNS servers, such as BIND.

**Notes**

- The DNS console provides new ways to perform familiar DNS administrative tasks previously performed in Microsoft® Windows® NT Server 4.0 using DNS Manager. For more information, see New ways to do familiar DNS tasks.

- To use the DNS console from another non-server computer, such as one running Microsoft® Windows® XP Professional, you must install the Windows Server 2003 Administration Tools Pack

- For information on installing DNS, see Install a DNS server.

**Command-line utilities**

There are several command-line utilities you can use to manage and troubleshoot DNS servers and clients. The following table describes each of these utilities, which can be run either by typing them at a command prompt or by entering them in batch files for scripted use.

| Command | Description |
|---|---|
| Nslookup | Used to perform query testing of the DNS domain namespace. For more information, see Nslookup. |
| Dnscmd | A command-line interface for managing DNS servers. This utility is useful in scripting batch files to help automate routine DNS management tasks, or to perform simple unattended setup and configuration of new DNS servers on your network. For more information, see Server administration using Dnscmd. |
| Ipconfig | This command is used to view and modify IP configuration details used by the computer. Additional command-line options are included with this utility to provide help in troubleshooting and supporting DNS clients. For more information, see Flush and reset a client resolver cache using the ipconfig command or Renew DNS client registration using the ipconfig command. |

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

**Event monitoring utilities**

The Windows Server 2003 family includes two options for monitoring DNS servers:

- **Default logging of DNS server event messages to the DNS server log.**

  DNS server event messages are separated and kept in their own system event log, the DNS server log, which can be viewed using the DNS console or Event Viewer. For more information, see View the DNS server system event log.

  The DNS server log contains events logged by the DNS Server service. For example, when the DNS server starts or stops, a corresponding event message is written to this log. Most additional critical DNS Server service events are also logged here, such as when the server starts but cannot locate initializing data, such as zones or boot information stored in the registry or (in some cases) Active Directory.

  The event types logged by DNS servers can be changed using the DNS console. For more information, see DNS server log reference.

  You can use Event Viewer to view and monitor client-related DNS events. These appear in the System log and are written by the DNS Client service at any computers running Windows (all versions). For more information, see Windows interface administrative tool reference A-Z: Event Viewer.

- **Optional debug options for trace logging to a text file on the DNS server computer.**

You can also use the DNS console to selectively enable additional debug logging options for temporary trace logging to a text-based file of DNS server activity. The file created and used for this feature, Dns.log, is stored in the systemroot\System32\Dns folder.

For more information, see Using server debug logging options.

**Performance monitoring utilities**

Performance monitoring for DNS servers can be done using additional service-specific counters that measure DNS server performance. These counters are accessible through System Monitor, which is provided in the **Performance** console.

When using System Monitor, you can create charts and graphs of server performance trends over time for any of your DNS servers. These can be further studied and analyzed to determine if additional server tuning is needed.

By measuring and reviewing server metrics over a period of time, it is possible to determine performance benchmarks and decide if further adjustments can be made to optimize the system. For more information, see Monitoring DNS server performance.

**Windows Management Instrumentation (WMI)**

WMI is the Microsoft implementation of Web-Based Enterprise Management (WBEM), which is an industry initiative to develop a standard technology for accessing management information in an enterprise environment. WMI uses the Common Information Model (CIM) industry standard to represent systems, applications, networks, devices, and other managed components in an enterprise environment. For more information about Windows Management Instrumentation, see the Microsoft Platform SDK Web site.

**Platform Software Developer Kit (SDK)**

Computers running a product in the Windows Server 2003 family provide functions that enable application programmers to use DNS, such as programmatically making DNS queries, comparing records, and looking up names.

Programmable DNS components are designed for use by C/C++ programmers. Familiarity with networking and with DNS is required. Programmers should be familiar with the IP protocol suite, as well as the DNS protocol and how DNS operates.

### Note

- For more information about manageability, see Management Strategies and Tools.

# Server features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server features**

The Domain Name System (DNS) Server service provides the following:

- **An RFC-compliant DNS server**

  DNS is an open protocol and is standardized by a set of Request for Comments (RFCs). Microsoft supports and complies with these standard specifications.

  For more information, see DNS RFCs.

- **Interoperability with other DNS server implementations**

  Because the DNS Server service is RFC-compliant and can use standard DNS data file and resource record formats, it can successfully work with most other DNS server implementations, such as those that use the Berkeley Internet Name Domain (BIND) software.

  For more information, see Interoperability issues.

- **Support for Active Directory**

  DNS is required for support of the Active Directory® directory service. If you install Active Directory on a server, you can automatically install and configure a DNS server if a DNS server that meets the Active Directory requirements cannot be located.

  First, in the Active Directory Installation Wizard, you specify the DNS name of the Active Directory domain for which you are promoting the server to become a domain controller. Later in the installation process, the wizard tests for the following:

  1. Based on its TCP/IP client configuration, it checks to see whether a preferred DNS server is configured for its use.
  2. If a preferred DNS server is available, it queries to find the primary authoritative server for the DNS name of the Active Directory domain you specified earlier in the wizard.
  3. It then tests to see whether the authoritative primary server can support and accept dynamic updates as described in the dynamic update protocol (RFC 2136).
  4. If, at this point in the process, a supporting DNS server cannot be located to accept updates for the specified DNS domain name you are using with Active Directory, you are provided with the option to install the DNS Server service locally.
  5. If you choose to install the DNS Server service locally, the IP address for the current preferred DNS server is used to configure a forwarder on the local DNS server. This configuration maintains any existing resolution to an Internet Service Provider (ISP).

  In general, the use of the Windows Server 2003 DNS Server service is strongly recommended for the best possible integration and support of Active Directory and enhanced DNS server features. You can, however, use another type of DNS server to support Active Directory deployment.

  When using other types of DNS servers, consider additional issues related to DNS interoperability. For more information, see Interoperability issues.

**Note**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

- **Enhancements to DNS zone storage in Active Directory**

  DNS zones can be stored in the domain or application directory partitions of Active Directory. A partition is a data structure within Active Directory used to distinguish data for different replication purposes. You can specify in which Active Directory partition to store the zone and, consequently, the set of domain controllers between which that zone's data will be replicated.

  For more information, see DNS zone replication in Active Directory.

  **Note**

  - This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

- **Conditional forwarders**

  The DNS Server service extends a standard forwarder configuration with conditional forwarders. A conditional forwarder is a DNS server on a network that is used to forward DNS queries according to the DNS domain name in the query. For example, a DNS server can be configured to forward all the queries it receives for names ending with widgets.example.com to the IP address of a specific DNS server or to the IP addresses of multiple DNS servers.

  For more information, see Understanding forwarders.

- **Stub zones**

  DNS supports a new zone type called a stub zone. A stub zone is a copy of a zone that contains only those resource records necessary to identify the authoritative DNS servers for that zone. You can use a stub zone instead of a secondary zone in situations where replicating all the zone data would be undesirable, such as over a slow network link. Note, however, that this replication efficiency is at the expense of resolution efficiency because the server hosting the stub zone is not authoritative for the zone and so must refer all queries for the zone to other servers.

  For more information, see Understanding stub zones.

- **Enhanced DNS security features**

  DNS provides enhanced security administration for the DNS Server service, the DNS Client service, and DNS data. For more information, see Security information for DNS.

- **Integration with other Microsoft networking services**

  The DNS Server service offers integration with other services and contains features beyond those specified in the RFCs. These include integration with Active Directory, WINS, and DHCP services.

  For more information, see Active Directory integration; WINS lookup integration; Dynamic update.

- **Improved ease of administration**

  The DNS console offers an improved graphical user interface for managing the DNS Server service. Also, there are several configuration wizards for performing common server administration tasks. In addition to the DNS console, other tools are provided to help you better manage and support DNS servers and clients on your network.

  For more information, see DNS tools.

- **RFC-compliant dynamic update protocol support**

  The DNS Server service allows clients to dynamically update resource records, based on the dynamic update protocol (RFC 2136). This improves DNS administration by reducing the time needed to manually manage these records. Computers running the DNS Client service can dynamically register their DNS names and IP addresses.

  For more information, see Dynamic update.

- **Support for incremental zone transfer between servers**

  Zone transfers are used between DNS servers to replicate information about a portion of the DNS namespace. Incremental zone transfer is used to replicate only the changed portions of a zone, conserving network bandwidth.

  For more information, see Understanding zones and zone transfer.

- **Support for new resource record types**

  The DNS Server service includes support for several new resource record (RR) types. These types, which include the service location (SRV) and ATM address (ATMA) RRs, expand the possibilities for using DNS as a names database service.

  For more information, see Resource records reference.

# Client features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Client features**

The Domain Name System (DNS) Client service is used to resolve DNS domain names and implements the following features:

- **System-wide caching**

  Resource records (RRs) from query responses are added to the client cache as applications query DNS servers. This information is then cached for a set Time to Live (TTL) and can be used again to answer subsequent queries.

- **RFC-compliant negative caching support**

  In addition to caching positive query responses from DNS servers (which contain resource record information in the answered reply), the DNS Client service also caches negative query responses. A negative response results when a resource record for the queried name does not exist.

  Negative caching prevents the repeating of additional queries for names that do not exist, which can adversely affect client system performance. Any query information negatively cached is kept for a shorter period of time than is used for positive query responses; by default, no more than 5 minutes. This avoids continued negative caching of stale information if the records later become available.

  Negative caching is a new DNS standard specification defined in RFC 2308. For more information, refer to this RFC. For more information on obtaining RFCs, see TCP/IP RFCs.

- **Avoidance of unresponsive DNS servers**

  The DNS Client service uses a server search list, ordered by preference. This list includes all preferred and alternate DNS servers configured for each of the active network connections on the system.

  The list is arranged based on the following criteria:

  1. Preferred DNS servers are given first priority.
  2. If no preferred DNS servers are available, then alternate DNS servers are used.
  3. Unresponsive servers are removed temporarily from these lists.

**Important**

- The DHCP Client service initiates dynamic registration for client DNS names. For more information, see Dynamic update or Using DNS servers with DHCP.

# Security information for DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for DNS**

Domain Name System (DNS) was originally designed as an open protocol and is therefore vulnerable to attackers. Windows Server 2003 DNS has improved the ability to prevent an attack on your DNS infrastructure through the addition of security features. Before considering which of the security features to use, you should be aware of the common threats to DNS security and the level of DNS security in your organization.

**DNS security threats**

The following are the typical ways in which your DNS infrastructure can be threatened by attackers:

- **Footprinting** is the process by which DNS zone data is obtained by an attacker to provide the attacker with the DNS domain names, computer names, and IP addresses for sensitive network resources. An attacker commonly begins an attack by using this DNS data to diagram, or footprint, a network. DNS domain and computer names usually indicate the function or location of a domain or computer in order to help users remember and identify domains and computers more easily. An attacker takes advantage of the same DNS principle to learn the function or location of domains and computers in the network.

- **Denial-of-service attack** is when an attacker attempts to deny the availability of network services by flooding one or more DNS servers in the network with recursive queries. As a DNS server is flooded with queries, its CPU usage will eventually reach its maximum and the DNS Server service will become unavailable. Without a fully operating DNS server on the network, network services that use DNS will become unavailable to network users.

- **Data modification** is an attempt by an attacker (that has footprinted a network using DNS) to use valid IP addresses in IP packets the attacker has created, thereby giving these packets the appearance of coming from a valid IP address in the network. This is commonly called IP spoofing. With a valid IP address (an IP address within the IP address range of a subnet), the attacker can gain access to the network and destroy data or conduct other attacks.

- **Redirection** is when an attacker is able to redirect queries for DNS names to servers under the control of the attacker. One method of redirection involves the attempt to pollute the DNS cache of a DNS server with erroneous DNS data that may direct future queries to servers under the control of the attacker. For example, if a query were originally made for example.microsoft.com and a referral answer provided a record for a name outside of the microsoft.com domain, such as malicious-user.com, then the DNS server would use the cached data for malicious-user.com to resolve a query for that name. Redirection can be accomplished whenever an attacker has writable access to DNS data, such as with insecure dynamic updates.

**Mitigating DNS security threats**

DNS can be configured to mitigate the common DNS security issues discussed above. The following table lists five main areas on which to concentrate when determining your DNS security.

| DNS security area | Description |
|---|---|
| DNS namespace | Incorporate DNS security into your DNS namespace design. For more information, see Securing DNS deployment. |
| DNS Server service | Review the default DNS Server service security settings and apply Active Directory security features when the DNS Server service is running on a domain controller. For more information, see Securing the DNS Server service. |
| DNS zones | Review the default DNS zone security settings and apply secure dynamic updates and Active Directory security features when the DNS zone is hosted on a domain controller. For more information, see Securing DNS zones. |

| | |
|---|---|
| DNS resource records | Review the default DNS resource record (RR) security settings and apply Active Directory security features when the DNS resource records are hosted on a domain controller. For more information, see Securing DNS resource records. |
| DNS clients | Control the DNS server IP addresses used by DNS clients. For more information, see Securing DNS clients. |

**Three levels of DNS security**
The following three levels of DNS security will help you understand your current DNS configuration and enable you to increase the DNS security of your organization.

**Low-level security**
Low-level security is a standard DNS deployment without any security precautions configured. Only deploy this level of DNS security in network environments where there is no concern for the integrity of your DNS data or in a private network where there is no threat of external connectivity.

- The DNS infrastructure of your organization is fully exposed to the Internet.
- Standard DNS resolution is performed by all DNS servers in your network.
- All DNS servers are configured with root hints pointing to the root servers for the Internet.
- All DNS servers permit zone transfers to any server.
- All DNS servers are configured to listen on all of their IP addresses.
- Cache pollution prevention is disabled on all DNS servers.
- Dynamic update is allowed for all DNS zones.
- User Datagram Protocol (UDP) and Transmission Control Protocol/Internet Protocol (TCP/IP) port 53 is open on the firewall for your network for both source and destination addresses.

**Medium-level security**
Medium-level security uses the DNS security features available without running DNS servers on domain controllers and storing DNS zones in Active Directory.

- The DNS infrastructure of your organization has limited exposure to the Internet.
- All DNS servers are configured to use forwarders to point to a specific list of internal DNS servers when they cannot resolve names locally.
- All DNS servers limit zone transfers to servers listed in the name server (NS) resource records in their zones.
- DNS servers are configured to listen on specified IP addresses.
- Cache pollution prevention is enabled on all DNS servers.
- Nonsecure dynamic update is not allowed for any DNS zones.
- Internal DNS servers communicate with external DNS servers through the firewall with a limited list of source and destination addresses allowed.
- External DNS servers in front of your firewall are configured with root hints pointing to the root servers for the Internet.
- All Internet name resolution is performed using proxy servers and gateways.

**High-level security**

High-level security uses the same configuration as medium-level security and also uses the security features available when the DNS Server service is running on a domain controller and DNS zones are stored in Active Directory. In addition, high-level security completely eliminates DNS communication with the Internet. This is not a typical configuration, but it is recommended whenever Internet connectivity is not required.

- The DNS infrastructure of your organization has no Internet communication by internal DNS servers.
- Your network uses an internal DNS root and namespace, where all authority for DNS zones is internal.
- DNS servers that are configured with forwarders use internal DNS server IP addresses only.
- All DNS servers limit zone transfers to specified IP addresses.
- DNS servers are configured to listen on specified IP addresses.
- Cache pollution prevention is enabled on all DNS servers.
- Internal DNS servers are configured with root hints pointing to the internal DNS servers hosting the root zone for your internal namespace.
- All DNS servers are running on domain controllers. A discretionary access control list (DACL) is configured on the DNS Server service to only allow specific individuals to perform administrative tasks on the DNS server.
- All DNS zones are stored in Active Directory. A DACL is configured to only allow specific individuals to create, delete, or modify DNS zones.
- DACLs are configured on DNS resource records to only allow specific individuals to create, delete, or modify DNS data.
- Secure dynamic update is configured for DNS zones, except the top-level and root zones, which do not allow dynamic updates at all.

# New features for DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features for DNS**
The following new Domain Name System (DNS) features and feature enhancements are available with the Microsoft® Windows Server™ 2003 family.

- **Improved domain controller name resolution**

  In response to DNS name resolution failures that may be encountered during location of replication partners and global catalog servers, domain controllers running Windows Server 2003 with SP1 request other variations of the server name that might be registered, which results in fewer failures due to DNS delays and misconfiguration. For more information about DNS name resolution, see How DNS Support for Active Directory Works on the Microsoft Web site.

- **Conditional forwarders**

  Forward DNS queries according to the DNS domain name in the query using conditional forwarders. For example, a DNS server can be configured to forward all the queries it receives for names ending with widgets.example.com to the IP address of a specific DNS server or to the IP addresses of multiple DNS servers.

  For more information, see Using forwarders.

- **Stub zones**

  Using stub zones, keep a DNS server hosting a parent zone aware of the authoritative DNS servers for its child zone and, thereby, maintain DNS name resolution efficiency.

  For more information, see Understanding stub zones.

- **DNS zone replication in Active Directory**

  Choose from four default replication options for Active Directory-integrated DNS zone data.

  For more information, see DNS zone replication in Active Directory.

- **Enhanced DNS security features**

  DNS provides greater precision in its security administration for the DNS Server service, the DNS Client service, and DNS data.

  For more information, see Security information for DNS.

- **Round robin all resource record (RR) types**

  By default, the DNS Server service will perform round-robin rotation for all resource record (RR) types.

  For more information, see Configuring round robin.

- **Enhanced debug logging**

  Use the enhanced DNS Server service debug logging settings to troubleshoot DNS problems.

  For more information, see Using server debug logging options.

- **DNSSEC**

DNS provides basic support of DNS Security Extensions (DNSSEC) protocol as defined in RFC 2535.

For more information, see Using DNS Security Extensions (DNSSEC).

- **EDNS0**

  Enable DNS requestors to advertise the size of their UDP packets and facilitate the transfer of packets larger than 512 octets, the original DNS restriction for UDP packet size (RFC 1035).

  For more information, see Using Extension Mechanisms for DNS (EDNS0).

- **Control automatic NS resource record registration on a server and a zone basis**

  For more information, see Restrict NS resource record registration and Allow NS record creation for specific domain controllers.

# Understanding DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding DNS**
This section covers:

- DNS domain names
- How DNS query works
- Understanding zones and zone transfer
- Reverse lookup
- Understanding forwarders
- Understanding stub zones
- Dynamic update
- Active Directory integration
- DNS zone replication in Active Directory
- DNS requirements for installing Active Directory
- DNS requirements for joining an Active Directory domain
- Understanding aging and scavenging
- Unicode character support
- WINS lookup integration

# DNS domain names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS domain names**

Domain Name System (DNS) was originally defined in Request for Comments (RFCs) 1034 and 1035. These documents specify elements common to all implementations of DNS-related software, including:

- A DNS domain namespace, which specifies a structured hierarchy of domains used to organize names.
- Resource records, which map DNS domain names to a specific type of resource information for use when the name is registered or resolved in the namespace.
- DNS servers, which store and answer name queries for resource records.
- DNS clients, also known as resolvers, which query servers to look up and resolve names to a type of resource record specified in the query.

For more information about Requests for Comments (RFCs), see DNS RFCs.

**Understanding the DNS domain namespace**

The DNS domain namespace, as shown in the following figure, is based on the concept of a tree of named domains. Each level of the tree can represent either a branch or a leaf of the tree. A branch is a level where more than one name is used to identify a collection of named resources. A leaf represents a single name used once at that level to indicate a specific resource.




The previous figure shows how Microsoft is assigned authority by the Internet root servers for its own part of the DNS domain namespace tree on the Internet. DNS clients and servers use queries as the fundamental method of resolving names in the tree to specific types of resource information. This information is provided by DNS servers in query responses to DNS clients, who then extract the information and pass it to a requesting program for resolving the queried name.

In the process of resolving a name, keep in mind that DNS servers often function as DNS clients, querying other servers in order to fully resolve a queried name. For more information, see How DNS query works.

**How the DNS domain namespace is organized**

Any DNS domain name used in the tree is technically a domain. Most DNS discussions, however, identify names in one of five ways, based on the level and the way a name is commonly used. For example, the DNS domain name registered to Microsoft (microsoft.com.) is known as a second-level domain. This is because the name has two parts (known as labels) that indicate it is located two levels below the root or top of the tree. Most DNS domain names have two or more labels, each of which indicates a new level in the tree. Periods are used in names to separate labels.

In addition to second-level domains, other terms used to describe DNS domain names by their function in the namespace are described in the following table.

| Name type | Description | Example |
|---|---|---|
| The domain root | This is the top of the tree, representing an unnamed level; it is sometimes shown as two empty quotation marks (**""**), indicating a null value. When used in a DNS domain name, it is stated by a trailing period (**.**) to designate that the name is located at the root or highest level of the domain hierarchy. In this instance, the DNS domain name is considered to be complete and points to an exact location in the tree of names. Names stated this way are called fully qualified domain names (FQDNs). | A single period (**.**) or a period used at the end of a name, such as "example.microsoft.com.". |
| Top-level domain | A name of two or three letters used to indicate a country/region or the type of organization using a name. For more information, see Top-level domains. | ".com", which indicates a name registered to a business for commercial use on the Internet. |
| Second-level domain | Variable-length names registered to an individual or organization for use on the Internet. These names are always based upon an appropriate top-level domain, depending on the type of organization or geographic location where a name is used. | "microsoft.com.", which is the second-level domain name registered to Microsoft by the Internet DNS domain name registrar. |
| Subdomain | Additional names that an organization can create that are derived from the registered second-level domain name. These include names added to grow the DNS tree of names in an organization and divide it into departments or geographic locations. | "example.microsoft.com.", which is a fictitious subdomain assigned by Microsoft for use in documentation example names. |
| Host or resource name | Names that represent a leaf in the DNS tree of names and identify a specific resource. Typically, the leftmost label of a DNS domain name identifies a specific computer on the network. For example, if a name at this level is used in a host (A) RR, it is used to look up the IP address of computer based on its host name. | "host-a.example.microsoft.com.", where the first label ("host-a") is the DNS host name for a specific computer on the network. |

**Interpreting a DNS domain name**

DNS has a method of noting and interpreting the fully qualified path to a DNS domain name similar to the way full paths to files or directories are noted or displayed at a command prompt.

For example, a directory tree path helps point to the exact location of a file stored on your computer. For Windows computers, the back slash (****) indicates each new directory that leads to the exact location of a file. For DNS, the equivalent is a period (**.**) indicating each new domain level used in a name.

For example, for a file called Services, the full path to this file as displayed at a Windows command prompt would be:

C:\Windows\System32\Drivers\Etc\Services

To interpret the full path of the file, the name is read in a left-to-right direction from the highest or most general piece of information (drive C, the drive where the file is stored) to its most specific information, the file name "Services". This example shows five separate levels of hierarchy that lead toward the location of the Services file on drive C:

1. The root folder for drive C (C:\).

2. The system root folder where Windows is installed (Windows).

3. A system folder where system components are stored (System32).

4. A subfolder where system device drivers are stored (Drivers).

5. A subfolder where miscellaneous files used by system and network device drivers are stored (Etc).

For DNS, an example of a domain name with multiple levels is the following, a fully qualified domain name (FQDN):

host-a.example.microsoft.com.

Unlike the file name example, a DNS FQDN, when read from left to right, moves from its most specific information (the DNS name for a computer called "host-a") to its highest or most general piece of information (the trailing period (**.**) that indicates the root of the DNS name tree). This example shows the four separate DNS domain levels that lead away from the specific host location of "host-a":

1. The "example" domain, which corresponds to a subdomain where the computer name "host-a" is registered for use.

2. The "microsoft" domain, which corresponds to the parent domain that roots the "example" subdomain.

3. The "com" domain, which corresponds to the top-level domain designated for use by business or commercial organizations that roots the "microsoft" domain.

4. The trailing period (**.**), which is a standard separator character used to qualify the full DNS domain name to the root level of the DNS namespace tree.

**Background on DNS and the Internet**

Domain Name System (DNS) was developed because of the need to provide a name-to-address mapping service for computers on the Internet. Before DNS was introduced in 1987, the practice of mapping friendly computer names to IP addresses was done mainly through the use of a shared static file, known as a Hosts file.

Originally, the Internet was small enough to use one centrally administered file that was published and downloaded using FTP for Internet-connected sites. Periodically, each Internet site would update its copy of the Hosts file, and updated versions of the Hosts file were posted to reflect network changes.

As the number of computers on the Internet grew, it became unworkable to have one centralized authority managing a single Hosts file for all Internet hosts. The file became increasingly larger, which made it harder to maintain and distribute to all sites in a current and updated form.

The DNS standard was developed to provide an alternative to Hosts files. RFCs 1034 and 1035 specify most of the core protocols and have been added to and updated by additional RFCs submitted to the Internet Engineering Task Force (IETF). The IETF reviews and approves new drafts continuously, so the standards for DNS evolve and change as needed.

For more information, see DNS RFCs.

### Notes

- DNS domain names are required to be unique at each level, but individual name labels can be reused in other domains. For example, the name "mailserver" can be used one time only in both the example.microsoft.com and microsoft.com domains.

- For more information on the use of labels, FQDNs, and other DNS host naming requirements, see Namespace planning for DNS.

- Hosts files are supported as a local static file method of mapping DNS domain names for host computers to their IP addresses. When the DNS Client service is started, it preloads any mapped entries added to this file into the local DNS names caches.

- The Hosts file is provided in the systemroot\System32\Drivers\Etc folder. To view or modify this file, you can use Notepad or another text-based editor. For more information about creating or editing the Hosts file, refer to the comments and examples in the file.

# How DNS query works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How DNS query works**

When a DNS client needs to look up a name used in a program, it queries DNS servers to resolve the name. Each query message the client sends contains three pieces of information, specifying a question for the server to answer:

- A specified DNS domain name, stated as a fully qualified domain name (FQDN)
- A specified query type, which can either specify a resource record by type or a specialized type of query operation
- A specified class for the DNS domain name.

  For Windows DNS servers, this should always be specified as the Internet (IN) class.

For example, the name specified could be the FQDN for a computer, such as "host-a.example.microsoft.com.", and the query type specified to look for an address (A) resource record by that name. Think of a DNS query as a client asking a server a two-part question, such as "Do you have any A resource records for a computer named 'hostname.example.microsoft.com.'?" When the client receives an answer from the server, it reads and interprets the answered A resource record, learning the IP address for the computer it asked for by name.

DNS queries resolve in a number of different ways. A client can sometimes answer a query locally using cached information obtained from a previous query. The DNS server can use its own cache of resource record information to answer a query. A DNS server can also query or contact other DNS servers on behalf of the requesting client to fully resolve the name, then send an answer back to the client. This process is known as recursion.

In addition, the client itself can attempt to contact additional DNS servers to resolve a name. When a client does so, it uses separate and additional nonrecursive queries based on referral answers from servers. This process is known as iteration.

In general, the DNS query process occurs in two parts:

- A name query begins at a client computer and is passed to a resolver, the DNS Client service, for resolution.
- When the query cannot be resolved locally, DNS servers can be queried as needed to resolve the name.

Both of these processes are explained in more detail in the following sections.

**Part 1: The local resolver**

The following figure shows an overview of the complete DNS query process.




As shown in the initial steps of the query process, a DNS domain name is used in a program on the local computer. The request is then passed to the DNS Client service for resolution using locally cached information. If the queried name can be resolved, the query is answered and the process is completed.

The local resolver cache can include name information obtained from two possible sources:

- If a Hosts file is configured locally, any host name-to-address mappings from that file are preloaded into the cache when

the DNS Client service is started.

- Resource records obtained in answered responses from previous DNS queries are added to the cache and kept for a period of time.

If the query does not match an entry in the cache, the resolution process continues with the client querying a DNS server to resolve the name.

**Part 2: Querying a DNS server**
As indicated in the previous figure, the client queries a preferred DNS server. The actual server used during the initial client/server query part of the process is selected from a global list. For more information about how this global list is compiled and updated, see Client features.

When the DNS server receives a query, it first checks to see if it can answer the query authoritatively based on resource record information contained in a locally configured zone on the server. If the queried name matches a corresponding resource record in local zone information, the server answers authoritatively, using this information to resolve the queried name.

If no zone information exists for the queried name, the server then checks to see if it can resolve the name using locally cached information from previous queries. If a match is found here, the server answers with this information. Again, if the preferred server can answer with a positive matched response from its cache to the requesting client, the query is completed.

If the queried name does not find a matched answer at its preferred server -- either from its cache or zone information -- the query process can continue, using recursion to fully resolve the name. This involves assistance from other DNS servers to help resolve the name. By default, the DNS Client service asks the server to use a process of recursion to fully resolve names on behalf of the client before returning an answer. In most cases, the DNS server is configured, by default, to support the recursion process as shown in the following figure.




In order for the DNS server to do recursion properly, it first needs some helpful contact information about other DNS servers in the DNS domain namespace. This information is provided in the form of *root hints*, a list of preliminary resource records that can be used by the DNS service to locate other DNS servers that are authoritative for the root of the DNS domain namespace tree. Root servers are authoritative for the domain root and top-level domains in the DNS domain namespace tree. For more information, see Updating root hints.

By using root hints to find root servers, a DNS server is able to complete the use of recursion. In theory, this process enables any DNS server to locate the servers that are authoritative for any other DNS domain name used at any level in the namespace tree.

For example, consider the use of the recursion process to locate the name "host-b.example.microsoft.com." when the client queries a single DNS server. The process occurs when a DNS server and client are first started and have no locally cached information available to help resolve a name query. It assumes that the name queried by the client is for a domain name of which the server has no local knowledge, based on its configured zones.

First, the preferred server parses the full name and determines that it needs the location of the server that is authoritative for the top-level domain, "com". It then uses an iterative (that is, a nonrecursive) query to the "com" DNS server to obtain a referral to the "microsoft.com" server. Next, a referral answer comes from the "microsoft.com" server to the DNS server for "example.microsoft.com".

Finally, the "example.microsoft.com." server is contacted. Because this server contains the queried name as part of its configured zones, it responds authoritatively back to the original server that initiated recursion. When the original server receives the response indicating that an authoritative answer was obtained to the requested query, it forwards this answer back to the requesting client and the recursive query process is completed.

Although the recursive query process can be resource-intensive when performed as described above, it has some performance advantages for the DNS server. For example, during the recursion process, the DNS server performing the recursive lookup

obtains information about the DNS domain namespace. This information is cached by the server and can be used again to help speed the answering of subsequent queries that use or match it. Over time, this cached information can grow to occupy a significant portion of server memory resources, although it is cleared whenever the DNS service is cycled on and off.

**Alternate query responses**
The previous discussion of DNS queries assumes that the process ends with a positive response returned to the client. However, queries can return other answers as well. These are the most common:

- An authoritative answer
- A positive answer
- A referral answer
- A negative answer

An authoritative answer is a positive answer returned to the client and delivered with the authority bit set in the DNS message to indicate the answer was obtained from a server with direct authority for the queried name.

A positive response can consist of the queried RR or a list of RRs (also known as an RRset) that fits the queried DNS domain name and record type specified in the query message.

A referral answer contains additional resource records not specified by name or type in the query. This type of answer is returned to the client if the recursion process is not supported. The records are meant to act as helpful reference answers that the client can use to continue the query using iteration.

A referral answer contains additional data such as resource records (RRs) that are other than the type queried. For example, if the queried host name was "www" and no A RRs for this name were found in this zone but a CNAME RR for "www" was found instead, the DNS server can include that information when responding to the client.

If the client is able to use iteration, it can make additional queries using the referral information in an attempt to fully resolve the name for itself.

A negative response from the server can indicate that one of two possible results was encountered while the server attempted to process and recursively resolve the query fully and authoritatively:

- An authoritative server reported that the queried name does not exist in the DNS namespace.
- An authoritative server reported that the queried name exists but no records of the specified type exist for that name.

The resolver passes the results of the query, in the form of either a positive or negative response, back to the requesting program and caches the response.

**Notes**

- If the resultant answer to a query is too long to be sent and resolved in a single UDP message packet, the DNS server can initiate a failover response over TCP port 53 to answer the client fully in a TCP connected session.
- Disabling the use of recursion on a DNS server is generally done when DNS clients are being limited to resolving names to a specific DNS server, such as one located on your intranet. Recursion might also be disabled when the DNS server is incapable of resolving external DNS names, and clients are expected to fail over to another DNS server for resolution of these names.

  You can disable the use of recursion by configuring in the **Advanced** properties in the DNS console on the applicable server. For more information, see Disable recursion on the DNS server.
- If you disable recursion on the DNS server, you will not be able to use forwarders on the same server.
- By default, DNS servers use several default timings when performing a recursive query and contacting other DNS servers. These are:

- A recursion retry interval of 3 seconds. This is the length of time the DNS service waits before retrying a query made during a recursive lookup.

- A recursion time-out interval of 15 seconds. This is the length of time the DNS service waits before failing a recursive lookup that has been retried.

Under most circumstances, these parameters do not need adjustment. However, if you are using recursive lookups over a slow-speed WAN link, you might be able to improve server performance and query completion by making slight adjustments to the settings. For more information, see Tuning advanced server parameters.

**How iteration works**

Iteration is the type of name resolution used between DNS clients and servers when the following conditions are in effect:

- The client requests the use of recursion, but recursion is disabled on the DNS server.

- The client does not request the use of recursion when querying the DNS server.

An iterative request from a client tells the DNS server that the client expects the best answer the DNS server can provide immediately, without contacting other DNS servers.

When iteration is used, a DNS server answers a client based on its own specific knowledge about the namespace with regard to the names data being queried. For example, if a DNS server on your intranet receives a query from a local client for "www.microsoft.com", it might return an answer from its names cache. If the queried name is not currently stored in the names cache of the server, the server might respond by providing a referral -- that is, a list of NS and A resource records for other DNS servers that are closer to the name queried by the client.

When a referral is made, the DNS client assumes responsibility to continue making iterative queries to other configured DNS servers to resolve the name. For example, in the most involved case, the DNS client might expand its search as far as the root domain servers on the Internet in an effort to locate the DNS servers that are authoritative for the "com" domain. Once it contacts the Internet root servers, it can be given further iterative responses from these DNS servers that point to actual Internet DNS servers for the "microsoft.com" domain. When the client is provided records for these DNS servers, it can send another iterative query to the external Microsoft DNS servers on the Internet, which can respond with a definitive and authoritative answer.

When iteration is used, a DNS server can further assist in a name query resolution beyond giving its own best answer back to the client. For most iterative queries, a client uses its locally configured list of DNS servers to contact other name servers throughout the DNS namespace if its primary DNS server cannot resolve the query.

**How caching works**

As DNS servers process client queries using recursion or iteration, they discover and acquire a significant store of information about the DNS namespace. This information is then cached by the server.

Caching provides a way to speed the performance of DNS resolution for subsequent queries of popular names, while substantially reducing DNS-related query traffic on the network.

As DNS servers make recursive queries on behalf of clients, they temporarily cache resource records (RRs). Cached RRs contain information obtained from DNS servers that are authoritative for DNS domain names learned while making iterative queries to search and fully answer a recursive query performed on behalf of a client. Later, when other clients place new queries that request RR information matching cached RRs, the DNS server can use the cached RR information to answer them.

When information is cached, a Time-To-Live (TTL) value applies to all cached RRs. As long as the TTL for a cached RR does not expire, a DNS server can continue to cache and use the RR again when answering queries by its clients that match these RRs. Caching TTL values used by RRs in most zone configurations are assigned the **Minimum (default) TTL** which is set used in the zone's start of authority (SOA) resource record. By default, the minimum TTL is 3,600 seconds (1 hour) but can be adjusted or, if needed, individual caching TTLs can be set at each RR.

**Notes**

- You can install a DNS server as a caching-only server. For more information, see Using caching-only servers.

- By default, DNS servers use a root hints file, Cache.dns, that is stored in the systemroot\System32\Dns folder on the server computer. The contents of this file are preloaded into server memory when the service is started and contain

pointer information to root servers for the DNS namespace where you are operating DNS servers. For more information about this file or how it is used, see DNS-related files.

# Understanding zones and zone transfer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding zones and zone transfer**
Domain Name System (DNS) allows a DNS namespace to be divided up into zones, which store name information about one or more DNS domains. For each DNS domain name included in a zone, the zone becomes the authoritative source for information about that domain.

**Understanding the difference between zones and domains**
A zone starts as a storage database for a single DNS domain name. If other domains are added below the domain used to create the zone, these domains can either be part of the same zone or belong to another zone. Once a subdomain is added, it can then either be:

- Managed and included as part of the original zone records, or
- Delegated away to another zone created to support the subdomain

For example, the following figure shows the microsoft.com domain, which contains domain names for Microsoft. When the microsoft.com domain is first created at a single server, it is configured as a single zone for all of the Microsoft DNS namespace. If, however, the microsoft.com domain needs to use subdomains, those subdomains must be included in the zone or delegated away to another zone.




In this example, the example.microsoft.com domain shows a new subdomain -- the example.microsoft.com domain -- delegated away from the microsoft.com zone and managed in its own zone. However, the microsoft.com zone needs to contain a few resource records to provide the delegation information that references the DNS servers that are authoritative for the delegated example.microsoft.com subdomain.

If the microsoft.com zone does not use delegation for a subdomain, any data for the subdomain remains part of the microsoft.com zone. For example, the subdomain dev.microsoft.com is not delegated away but is managed by the microsoft.com zone.

**Why zone replication and zone transfers are needed**
Because of the important role that zones play in DNS, it is intended that they be available from more than one DNS server on the network to provide availability and fault tolerance when resolving name queries. Otherwise, if a single server is used and that server is not responding, queries for names in the zone can fail. For additional servers to host a zone, zone transfers are required to replicate and synchronize all copies of the zone used at each server configured to host the zone.

When a new DNS server is added to the network and is configured as a new secondary server for an existing zone, it performs a full initial transfer of the zone to obtain and replicate a full copy of resource records for the zone. For most earlier DNS server implementations, this same method of full transfer for a zone is also used when the zone requires updating after changes are made to the zone. For DNS servers running Windows Server 2003, the DNS service supports *incremental zone transfer*, a revised DNS zone transfer process for intermediate changes.

**Incremental zone transfers**
Incremental zone transfers are described in Request for Comments (RFC) 1995 as an additional DNS standard for replicating DNS zones. For more information about RFCs, see the RFC Editor Web site. When incremental transfers are supported by both a DNS server acting as the source for a zone and any servers that copy the zone from it, it provides a more efficient method of propagating zone changes and updates.

In earlier DNS implementations, any request for an update of zone data required a full transfer of the entire zone database using an AXFR query. With incremental transfer, an alternate query type (IXFR) can be used instead. This allows the secondary server to pull only those zone changes it needs to synchronize its copy of the zone with its source, either a primary or secondary copy of the zone maintained by another DNS server.

With IXFR zone transfers, differences between the source and replicated versions of the zone are first determined. If the zones are identified to be the same version -- as indicated by the serial number field in the start of authority (SOA) resource record of each zone -- no transfer is made.

If the serial number for the zone at the source is greater than at the requesting secondary server, a transfer is made of only those changes to resource records (RRs) for each incremental version of the zone. For an IXFR query to succeed and changes to be sent, the source DNS server for the zone must keep a history of incremental zone changes to use when answering these queries. The incremental transfer process requires substantially less traffic on a network and zone transfers are completed much faster.

**Example: Zone transfer**
A zone transfer might occur during any of the following scenarios:

- When the refresh interval expires for the zone
- When a secondary server is notified of zone changes by its master server
- When the DNS Server service is started at a secondary server for the zone
- When the DNS console is used at a secondary server for the zone to manually initiate a transfer from its master server

Zone transfers are always initiated at the secondary server for a zone and sent to their configured master servers which act as their source for the zone. Master servers can be any other DNS server that loads the zone, such as either the primary server for the zone or another secondary server. When the master server receives the request for the zone, it can reply with either a partial or full transfer of the zone to the secondary server.

As shown in the following figure, zone transfers between servers follow an ordered process. This process varies depending on whether a zone has been previously replicated, or if initial replication of a new zone is being performed.




In this example, the following sequence is performed for a requesting secondary server -- the destination server -- for a zone and its source server, another DNS server that hosts the zone.

1. During new configuration, the destination server sends an initial "all zone" transfer (AXFR) request to the master DNS server configured as its source for the zone.

2. The master (source) server responds and fully transfers the zone to the secondary (destination) server.

   The zone is delivered to the destination server requesting the transfer with its version established by use of a **Serial number** field in the properties for the start of authority (SOA) resource record (RR). The SOA RR also contains a stated refresh interval in seconds (by default, 900 seconds or 15 minutes) to indicate when the destination server should next request to renew the zone with the source server.

3. When the refresh interval expires, an SOA query is used by the destination server to request renewal of the zone from the source server.

4. The source server answers the query for its SOA record.

   This response contains the serial number for the zone in its current state at the source server.

5. The destination server checks the serial number of the SOA record in the response and determines how to renew the zone.

   If the value of the serial number in the SOA response is equal to its current local serial number, it concludes that the zone is the same at both servers and that a zone transfer is not needed. The destination server then renews the zone by resetting its refresh interval based on the value of this field in the SOA response from its source server.

   If the value of the serial number in the SOA response is higher than its current local serial number, it concludes that the zone has been updated and that a transfer is needed.

6. If the destination server concludes that the zone has changed, it sends an IXFR query to the source server, containing its current local value for the serial number in the SOA record for the zone.

7. The source server responds with either an incremental or full transfer of the zone.

   If the source server supports incremental transfer by maintaining a history of recent incremental zone changes for modified resource records, it can answer with an incremental zone transfer (IXFR) of the zone.

   If the source server does not support incremental transfer, or does not have a history of zone changes, it can answer with a full (AXFR) transfer of the zone instead.

**Note**

- For servers running Windows 2000 and Windows Server 2003, incremental zone transfer through IXFR query is supported. For earlier versions of the DNS service and for many other DNS server implementations, incremental zone transfer is not available and only full-zone (AXFR) queries and transfers are used to replicate zones.

**DNS Notify**

Windows-based DNS servers support DNS Notify, an update to the original DNS protocol specification that permits a means of initiating notification to secondary servers when zone changes occur (RFC 1996). DNS notification implements a push mechanism for notifying a select set of secondary servers for a zone when it is updated. Servers that are notified can then initiate a zone transfer as described above to pull zone changes from their master servers and update their local replicas of the zone.

For secondaries to be notified by the DNS server acting as their configured source for a zone, each secondary server must first have its IP address in the notify list of the source server. When using the DNS console, this list is maintained in the **Notify** dialog box, which is accessible from the **Zone Transfer** tab located in *zone***Properties**.

In addition to notifying the listed servers, the DNS console permits you to use the contents of the notify list as a means to restrict or limit zone transfer access to only those secondary servers specified in the list. This can help prevent an undesired attempt by an unknown or unapproved DNS server to pull, or request, zone updates. For more information, see Create and manage a notify list for a zone.

The following is a brief summary of the typical DNS notification process for zone updates:

1. The local zone at a DNS server acting as a master server, a source for the zone to other servers, is updated. When the zone is updated at the master or source server, the serial number field in the SOA RR is also updated, indicating a new local version of the zone.

2. The master server sends a DNS notify message to other servers that are part of its configured notify list.

3. All secondary servers that receive the notify message can then respond by initiating a zone transfer request back to the notifying master server.

The normal zone transfer process can then continue as described in the previous section.

You cannot configure a notify list for a stub zone.

**Important**

- Use DNS notification only to notify servers operating as secondary servers for a zone. For replication of directory-integrated zones, DNS notification is not needed.

  This is because any DNS servers that load a zone from Active Directory automatically poll the directory (as specified by the SOA resource record's Refresh Interval) to update and refresh the zone.

  In these cases, configuring a notify list can actually degrade system performance by causing unnecessary additional transfer requests for the updated zone.

**Notes**

- By default, the DNS server will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Reverse lookup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Reverse lookup**

In most DNS lookups, clients typically perform a forward lookup, which is a search based on the DNS name of another computer as stored in an address (A) resource record. This type of query expects an IP address as the resource data for the answered response.

DNS also provides a reverse lookup process, enabling clients to use a known IP address during a name query and look up a computer name based on its address. A reverse lookup takes the form of a question, such as "Can you tell me the DNS name of the computer that uses the IP address 192.168.1.20?"

DNS was not originally designed to support this type of query. One problem for supporting the reverse query process is the difference in how the DNS namespace organizes and indexes names and how IP addresses are assigned. If the only method to answer the previous question was to search in all domains in the DNS namespace, a reverse query would take too long and require too much processing to be useful.

To solve this problem, a special domain, the in-addr.arpa domain, was defined in the DNS standards and reserved in the Internet DNS namespace to provide a practical and reliable way to perform reverse queries. To create the reverse namespace, subdomains within the in-addr.arpa domain are formed using the reverse ordering of the numbers in the dotted-decimal notation of IP addresses.

This reversed ordering of the domains for each octet value is needed because, unlike DNS names, when IP addresses are read from left to right, they are interpreted in the opposite manner. When an IP address is read from left to right, it is viewed from its most generalized information (an IP network address) in the first part of the address to the more specific information (an IP host address) contained in the last octets.

For this reason, the order of IP address octets must be reversed when building the in-addr.arpa domain tree. The IP addresses of the DNS in-addr.arpa tree can be delegated to companies as they are assigned a specific or limited set of IP addresses within the Internet-defined address classes.

Finally, the in-addr.arpa domain tree, as built into DNS, requires that an additional resource record (RR) type — the pointer (PTR) RR — be defined. This RR is used to create a mapping in the reverse lookup zone that typically corresponds to a host (A) named RR for the DNS computer name of a host in its forward lookup zone.

### Note

- The in-addr.arpa domain applies for use in all TCP/IP networks that are based on Internet Protocol version 4 (IPv4) addressing. The New Zone Wizard automatically assumes that you are using this domain when creating a new reverse lookup zone.

  If you are installing DNS and configuring reverse lookup zones for an Internet Protocol version 6 (IPv6) network, you can specify an exact name in the New Zone wizard. This will permit you to create reverse lookup zones in the DNS console that can be used to support IPv6 networks, which uses a different special domain name, the ip6.arpa domain.

  Additional information is available on IPv6 and DNS, including examples of how to create and use ip6.arpa domain names in Request for Comments (RFC) 3596, "DNS Extensions to support IP version 6." For more information, refer directly to this RFC, which can be obtained from the RFC Editor Web site.

**Example: Reverse query (for IPv4 networks)**

The following figure shows an example of a reverse query initiated by a DNS client (host-b) to learn the name of another host (host-a) based on its IP address, 192.168.1.20.



The reverse query process as shown in this figure occurs in the following steps:

1. The client, "host-b", queries the DNS server for a pointer (PTR) RR that maps to the IP address of 192.168.1.20 for "host-a".

   Because the query is for PTR records, the resolver reverses the address and appends the in-addr.arpa domain to the end of the reverse address. This forms the fully qualified domain name ("20.1.168.192.in-addr.arpa.") for which to be searched in a reverse lookup zone.

2. Once located, the authoritative DNS server for "20.1.168.192.in-addr.arpa" can respond with the PTR record information. This includes the DNS domain name for "host-a", completing the reverse lookup process.

Keep in mind that if the queried reverse name is not answerable from the DNS server, normal DNS resolution (either recursion or iteration) can be used to locate a DNS server that is authoritative for the reverse lookup zone and that contains the queried name. In this sense, the name resolution process used in a reverse lookup is identical to that of a forward lookup.

**Note**

- The DNS console provides a means for you to configure a subnetted reverse lookup "classless" zone when the **Advanced** view is selected. This allows you to configure a zone in the in-addr.arpa domain for a limited set of assigned IP addresses where a nondefault IP subnet mask is used with those addresses. For more information, see Managing reverse lookup zones.

**Inverse queries**

Inverse queries are an outdated practice, originally proposed as part of the DNS standard to look up a host name based on its IP address. They use a nonstandard DNS query operation, and their use is limited to some of the earlier versions of Nslookup, a command-line utility for troubleshooting and testing DNS service.

The DNS service recognizes and accepts inverse query messages, answering them with a fake inverse query response. For DNS servers running in Windows NT Server 4.0, this support is available by default if the server computer has been updated to Service Pack 4 (SP4) or later.

**Notes**

- The configuration of PTR resource records and reverse lookup zones for identifying hosts by reverse query is strictly an optional part of the DNS standard implementation. You are not required to use reverse lookup zones, although for some networked applications, they are used to perform security checks.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Understanding forwarders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding forwarders**

A forwarder is a Domain Name System (DNS) server on a network used to forward DNS queries for external DNS names to DNS servers outside of that network. You can also forward queries according to specific domain names using *conditional forwarders*.

A DNS server on a network is designated as a forwarder by having the other DNS servers in the network forward the queries they cannot resolve locally to that DNS server. By using a forwarder, you can manage name resolution for names outside of your network, such as names on the Internet, and improve the efficiency of name resolution for the computers in your network. For more information about forwarders and conditional forwarders, see Using forwarders.

The following figure illustrates how external name queries are directed using forwarders.




For more information about directing external queries, see Directing queries through forwarders.

Without having a specific DNS server designated as a forwarder, *all* DNS servers can send queries outside of a network using their root hints. As a result, a lot of internal, and possibly critical, DNS information can be exposed on the Internet. In addition to this security and privacy issue, this method of resolution can result in a large volume of external traffic that is costly and inefficient for a network with a slow Internet connection or a company with high Internet service costs.

When you designate a DNS server as a forwarder, you make that forwarder responsible for handling external traffic, thereby limiting DNS server exposure to the Internet. A forwarder will build up a large cache of external DNS information because all of the external DNS queries in the network are resolved through it. In a small amount of time, a forwarder will resolve a good portion of external DNS queries using this cached data and thereby decrease the Internet traffic over the network and the response time for DNS clients.

A DNS server configured to use a forwarder will behave differently than a DNS server that is not configured to use a forwarder. A DNS server configured to use a forwarder behaves as follows:

1. When the DNS server receives a query, it attempts to resolve this query using the primary and secondary zones that it hosts and its cache.

2. If the query cannot be resolved using this local data, then it will forward the query to the DNS server designated as a forwarder.

3. The DNS server will wait briefly for an answer from the forwarder before attempting to contact the DNS servers specified in its root hints.

When a DNS server forwards a query to a forwarder it sends a recursive query to the forwarder. This is different than the iterative query that a DNS server will send to an other DNS server during standard name resolution (name resolution that does not involve a forwarder).

**Conditional forwarders**

A conditional forwarder is a DNS server on a network that is used to forward DNS queries according to the DNS domain name in the query. For example, a DNS server can be configured to forward all the queries it receives for names ending with

widgets.example.com to the IP address of a specific DNS server or to the IP addresses of multiple DNS servers.

**Intranet name resolution**

A conditional forwarder can be used to improve name resolution for domains within your intranet. Intranet name resolution can be improved by configuring DNS servers with forwarders for specific internal domain names. For example, all DNS servers in the domain widgets.example.com could be configured to forward queries for names that end with test.example.com to the authoritative DNS servers for merged.widgets.example.com, thereby removing the step of querying the root servers of example.com, or removing the step of configuring DNS servers in the widgets.example.com zone with secondary zones for test.example.com.

**Internet name resolution**

DNS servers can use conditional forwarders to resolve queries between the DNS domain names of companies that share information. For example, two companies, Widgets Toys and TailspinToys, want to improve how the DNS clients of Widgets Toys resolve the names of the DNS clients of Tailspin Toys. The administrators from Tailspin Toys inform the administrators of Widgets Toys about the set of DNS servers in the Tailspin Toys network where Widgets can send queries for the domain dolls.tailspintoys.com. The DNS servers within the Widgets Toys network are configured to forward all queries for names ending with dolls.tailspintoys.com to the designated DNS servers in the network for Tailspin Toys. Consequently, the DNS servers in the Widgets Toys network do not need to query their internal root servers, or the Internet root servers, to resolve queries for names ending with dolls.tailspintoys.com.

# Understanding stub zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding stub zones**

A stub zone is a copy of a zone that contains only those resource records necessary to identify the authoritative Domain Name System (DNS) servers for that zone. A stub zone is used to resolve names between separate DNS namespaces. This type of resolution may be necessary when a corporate merger requires that the DNS servers for two separate DNS namespaces resolve names for clients in both namespaces.

A stub zone consists of:

- The start of authority (SOA) resource record, name server (NS) resource records, and the glue A resource records for the delegated zone.
- The IP address of one or more master servers that can be used to update the stub zone.

The master servers for a stub zone are one or more DNS servers authoritative for the child zone, usually the DNS server hosting the primary zone for the delegated domain name.

For more information, see Using stub zones.

**Stub zone resolution**

When a DNS client performs a recursive query operation on a DNS server hosting a stub zone, the DNS server uses the resource records in the stub zone to resolve the query. The DNS server sends an iterative query to the authoritative DNS servers specified in the NS resource records of the stub zone as if it were using NS resource records in its cache. If the DNS server cannot find the authoritative DNS servers in its stub zone, the DNS server hosting the stub zone attempts standard recursion using its root hints.

The DNS server will store the resource records it receives from the authoritative DNS servers listed in a stub zone in its cache, but it will not store these resource records in the stub zone itself; only the SOA, NS, and glue A resource records returned in response to the query are stored in the stub zone. The resource records stored in the cache are cached according to the Time-to-Live (TTL) value in each resource record. The SOA, NS, and glue A resource records, which are not written to cache, expire according to the expire interval specified in the stub zone's SOA record, which is created during the creation of the stub zone and updated during transfers to the stub zone from the original, primary zone.

If the query was an iterative query, the DNS server returns a referral containing the servers specified in the stub zone.

**Communication between DNS servers hosting parent and child zones**

A DNS server that has delegated a domain to a child zone on a different DNS server is made aware of new authoritative DNS servers for the child zone only when the resource records for these new DNS servers are added to the parent zone hosted on the DNS server. This is a manual process and requires that the administrators for the different DNS servers communicate often. With stub zones, a DNS server hosting a stub zone for one of its delegated domains can obtain updates of the authoritative DNS servers for the child zone when the stub zone is updated. The update is performed from the DNS server hosting the stub zone and the administrator for the DNS server hosting the child zone does not need to be contacted. This functionality is explained in the following example.

For more information, see Active Directory-integrated stub zones, Understanding zones and zone transfer, and Delegating zones.

# Dynamic update

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dynamic update**
Dynamic update enables DNS client computers to register and dynamically update their resource records with a DNS server whenever changes occur. This reduces the need for manual administration of zone records, especially for clients that frequently move or change locations and use DHCP to obtain an IP address.

The DNS Client and Server services support the use of dynamic updates, as described in Request for Comments (RFC) 2136, "Dynamic Updates in the Domain Name System." The DNS Server service allows dynamic update to be enabled or disabled on a per-zone basis at each server configured to load either a standard primary or directory-integrated zone. By default, the DNS Client service will dynamically update host (A) resource records (RRs) in DNS when configured for TCP/IP. For more information about RFCs, see DNS RFCs.

**How client and server computers update their DNS names**
By default, computers that are statically configured for TCP/IP attempt to dynamically register host (A) and pointer (PTR) resource records (RRs) for IP addresses configured and used by their installed network connections. By default, all computers register records based on their fully qualified domain name (FQDN).

The primary full computer name, a FQDN, is based on the **primary DNS suffix of a computer** appended to its **Computer name**.

Both of these settings are displayed or configured from the **Computer Name** tab in **System properties**. For more information, see View system properties.

**Notes**

- By default, the DNS client on Windows XP does not attempt dynamic update over a Remote Access Service (RAS) or virtual private network (VPN) connection. To modify this configuration, you can modify the advanced TCP/IP settings of the particular network connection or modify the registry. For more information, see Configure TCP/IP to use DNS and the Microsoft Windows Resource Kits Web site.

- By default, the DNS client does not attempt dynamic update of top-level domain (TLD) zones. Any zone named with a single-label name is considered a TLD zone, for example, com, edu, blank, my-company. To configure the DNS client to allow the dynamic update of TLD zones, you can use the **Update Top Level Domain Zones** policy setting or modify the registry.

- By default, the primary DNS suffix portion of a computer's FQDN is the same as the name of the Active Directory domain to which the computer is joined. To allow different primary DNS suffixes, a domain administrator may create a restricted list of allowed suffixes by modifying the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is managed by the domain administrator using Active Directory Service Interfaces (ADSI) or the Lightweight Directory Access Protocol (LDAP).

  For more information, see Programming interfaces and Directory access protocol.

Dynamic updates can be sent for any of the following reasons or events:

- An IP address is added, removed, or modified in the TCP/IP properties configuration for any one of the installed network connections.

- An IP address lease changes or renews with the DHCP server any one of the installed network connections. For example, when the computer is started or if the **ipconfig /renew** command is used.

- The **ipconfig /registerdns** command is used to manually force a refresh of the client name registration in DNS.

- At startup time, when the computer is turned on.

- A member server is promoted to a domain controller.

When one of the previous events triggers a dynamic update, the DHCP Client service (not the DNS Client service) sends updates. This is designed so that if a change to the IP address information occurs because of DHCP, corresponding updates in DNS are performed to synchronize name-to-address mappings for the computer. The DHCP Client service performs this function for all network connections used on the system, including connections not configured to use DHCP.

**Notes**

- The process of how dynamic updates are performed for computers running Windows 2000, Windows XP, or Windows Server 2003 operating systems that use DHCP to obtain their IP address is different than is described in this section. For more information, see Using DNS servers with DHCP.

- The update process described in this section assumes that installation defaults are in effect for computers running Windows 2000, Windows XP, or servers running Windows Server 2003. Specific names and update behavior is tunable where advanced TCP/IP properties are configured to use non-default DNS settings.

  In addition to the full computer name (or primary name) of the computer, additional connection-specific DNS names can be configured and optionally registered or updated in DNS. For more information, see Configuring multiple names or Configure TCP/IP to use DNS.

**Example: How dynamic update works**

Dynamic updates are typically requested when either a DNS name or IP address changes on the computer. For example, suppose a client named "oldhost" is first configured in **System properties** with the following names:

| | |
|---|---|
| **Computer name** | oldhost |
| **DNS domain name of computer** | example.microsoft.com |
| **Full computer name** | oldhost.example.microsoft.com |

In this example, no connection-specific DNS domain names are configured for the computer. Later, the computer is renamed from "oldhost" to "newhost", resulting in the following name changes on the system:

| | |
|---|---|
| **Computer name** | newhost |
| **DNS domain name of computer** | example.microsoft.com |
| **Full computer name** | newhost.example.microsoft.com |

Once the name change is applied in **System properties**, you are prompted to restart the computer. When the computer restarts Windows, the DHCP Client service performs the following sequence to update DNS:

1. **The DHCP Client service sends a start of authority (SOA) type query using the DNS domain name of the computer.**

   The client computer uses the currently configured FQDN of the computer (such as "newhost.example.microsoft.com") as the name specified in this query.

2. **The authoritative DNS server for the zone containing the client FQDN responds to the SOA-type query.**

   For standard primary zones, the primary server (owner) returned in the SOA query response is fixed and static. It always matches the exact DNS name as it appears in the SOA RR stored with the zone. If, however, the zone being updated is directory-integrated, any DNS server loading the zone can respond and dynamically insert its own name as the primary server (owner) of the zone in the SOA query response.

3. **The DHCP Client service then attempts to contact the primary DNS server.**

   The client processes the SOA query response for its name to determine the IP address of the DNS server authorized as

the primary server for accepting its name. It then proceeds to perform the following sequence of steps as needed to contact and dynamically update its primary server:

a. It sends a dynamic update request to the primary server determined in the SOA query response.

   If the update succeeds, no further action is taken.

b. If this update fails, the client next sends an NS-type query for the zone name specified in the SOA record.

c. When it receives a response to this query, it sends an SOA query to the first DNS server listed in the response.

d. After the SOA query is resolved, the client sends a dynamic update to the server specified in the returned SOA record.

   If the update succeeds, no further action is taken.

e. If this update fails, then the client repeats the SOA query process by sending to the next DNS server listed in the response.

4. **Once the primary server is contacted that can perform the update, the client sends the update request and the server processes it.**

   The contents of the update request include instructions to add A (and possibly PTR) RRs for "newhost.example.microsoft.com" and remove these same record types for "oldhost.example.microsoft.com", the name that was previously registered.

   The server also checks to ensure that updates are permitted for the client request. For standard primary zones, dynamic updates are not secured, so any client attempt to update succeeds. For Active Directory-integrated zones, updates are secured and performed using directory-based security settings.

Dynamic updates are sent or refreshed periodically. By default, computers send a refresh once every 7 days. If the update results in no changes to zone data, the zone remains at its current version and no changes are written. Updates result in actual zone changes or increased zone transfer only if names or addresses actually change.

Note that names are not removed from DNS zones if they become inactive or are not updated within the refresh interval (7 days). DNS does not use a mechanism to release or tombstone names, although DNS clients do attempt to delete or update old name records when a new name or address change is applied.

When the DHCP Client service registers A and PTR resource records for a computer, it uses a default caching Time to Live (TTL) of 15 minutes for host records. This determines how long other DNS servers and clients cache a computer's records when they are included in a query response.

**Secure dynamic update**

DNS update security is available only for zones that are integrated into Active Directory. Once you directory-integrate a zone, access control list (ACL) editing features are available in the DNS console so you can add or remove users or groups from the ACL for a specified zone or resource record. For more information, see Modify security for a resource record or Modify security for a directory-integrated zone.

By default, dynamic update security for DNS servers and clients can be handled as follows:

- **DNS clients attempt to use unsecured dynamic update first. If an unsecured update is refused, clients try to use secure update.**

  Also, clients use a default update policy that permits them to attempt to overwrite a previously registered resource record, unless they are specifically blocked by update security.

- **Once a zone becomes Active Directory-integrated, DNS servers running Windows Server 2003 default to allowing only secure dynamic updates.**

  When using standard zone storage, the default for the DNS Server service is to not allow dynamic updates on its zones.

For zones that are either directory-integrated or use standard file-based storage, you can change the zone to allow all dynamic updates which permits all updates to be accepted.

**Important**

- The DHCP Server service can perform proxy registration and update of DNS records for legacy clients that do not support dynamic updates.

  For information, see Using DNS servers with DHCP.

**Notes**

- Dynamic update is a recent additional DNS standard specification, defined in RFC 2136. For more information, refer to the RFC. For information on obtaining RFCs, see TCP/IP RFCs.

- For more information about DNS dynamic updates and secure updates, see Using the Windows Deployment and Resource Kits.

- The dynamic registration of DNS resource records can be restricted with the use of registry entries. For more information, see article Q246804, "How to Enable/Disable Windows 2000 Dynamic DNS Registrations," in the Microsoft Knowledge Base.

# Active Directory integration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory integration**
The DNS Server service is integrated into the design and implementation of Active Directory. Active Directory provides an enterprise-level tool for organizing, managing, and locating resources in a network.

**Note**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

When deploying DNS servers with Active Directory, consider the following:

- **DNS is required for locating domain controllers running Windows Server 2003.**

  The Net Logon service uses new DNS server support to provide registration of domain controllers in your DNS domain namespace.

- **DNS servers running Windows Server 2003 can use Active Directory for storing and replicating your zones.**

  By directory integrating your zones, you can take advantage of DNS features such as secure dynamic updates and record aging and scavenging features.

  For more information, see either Dynamic update or Understanding aging and scavenging.

**How DNS integrates with Active Directory**
When you install Active Directory on a server, you promote the server to the role of a domain controller for a specified domain. When completing this process, you are prompted to specify a DNS domain name for the Active Directory domain for which you are joining and promoting the server.

If during this process, a DNS server authoritative for the domain that you specified either cannot be located on the network or does not support the DNS dynamic update protocol, you are prompted with the option to install a DNS server. This option is provided because a DNS server is required to locate this server or other domain controllers for members of an Active Directory domain. For more information about how computers locate domain controllers using DNS, see the Using the Windows Deployment and Resource Kits.

Once you have installed Active Directory, you have two options for storing and replicating your zones when operating the DNS server at the new domain controller:

- **Standard zone storage, using a text-based file.**

  Zones stored this way are located in .Dns files that are stored in the systemroot\System32\Dns folder on each computer operating a DNS server. Zone file names correspond to the name you choose for the zone when creating it, such as example.microsoft.com.dns if the zone name was "example.microsoft.com."

- **Directory-integrated zone storage, using the Active Directory database.**

  Zones stored this way are located in the Active Directory tree under the domain or application directory partition. Each directory-integrated zone is stored in a *dnsZone* container object identified by the name you choose for the zone when creating it. For more information, see DNS zone replication in Active Directory and Directory data store.

**Benefits of Active Directory integration**
For networks deploying DNS to support Active Directory, directory-integrated primary zones are strongly recommended and provide the following benefits:

- **Multimaster update and enhanced security based on the capabilities of Active Directory.**

  In a standard zone storage model, DNS updates are conducted based upon a single-master update model. In this model, a single authoritative DNS server for a zone is designated as the primary source for the zone.

  This server maintains the master copy of the zone in a local file. With this model, the primary server for the zone represents a single fixed point of failure. If this server is not available, update requests from DNS clients are not processed for the zone.

  With directory-integrated storage, dynamic updates to DNS are conducted based upon a multimaster update model.

  In this model, any authoritative DNS server, such as a domain controller running a DNS server, is designated as a primary source for the zone. Because the master copy of the zone is maintained in the Active Directory database, which is fully replicated to all domain controllers, the zone can be updated by the DNS servers operating at any domain controller for the domain.

  With the multimaster update model of Active Directory, any of the primary servers for the directory-integrated zone can process requests from DNS clients to update the zone as long as a domain controller is available and reachable on the network.

  Also, when using directory-integrated zones, you can use access control list (ACL) editing to secure a dnsZone object container in the directory tree. This feature provides granulated access to either the zone or a specified RR in the zone.

  For example, an ACL for a zone RR can be restricted so that dynamic updates are only allowed for a specified client computer or a secure group such as a domain administrators group. This security feature is not available with standard primary zones.

  Note that when you change the zone type to be directory-integrated, the default for updating the zone changes to allow only secure updates. Also, while you may use ACLs on DNS-related Active Directory objects, ACLs may only be applied to the DNS client service.

- **Zones are replicated and synchronized to new domain controllers automatically whenever a new one is added to an Active Directory domain.**

  Although DNS service can be selectively removed from a domain controller, directory-integrated zones are already stored at each domain controller, so zone storage and management is not an additional resource. Also, the methods used to synchronize directory-stored information offer performance improvement over standard zone update methods, which can potentially require transfer of the entire zone.

- **By integrating storage of your DNS zone databases in Active Directory, you can streamline database replication planning for your network.**

  When your DNS namespace and Active Directory domains are stored and replicated separately, you need to plan and potentially administer each separately. For example, when using standard DNS zone storage and Active Directory together, you would need to design, implement, test, and maintain two different database replication topologies. For instance, one replication topology is needed for replicating directory data between domain controllers, and another topology would be needed for replicating zone databases between DNS servers.

  This can create additional administrative complexity for planning and designing your network and allowing for its eventual growth. By integrating DNS storage, you unify storage management and replication issues for both DNS and Active Directory, merging and viewing them together as a single administrative entity.

- **Directory replication is faster and more efficient than standard DNS replication.**

  Because Active Directory replication processing is performed on a per-property basis, only relevant changes are propagated. This allows less data to be used and submitted in updates for directory-stored zones.

**Notes**

- Only primary zones can be stored in the directory. A DNS server cannot store secondary zones in the directory. It must store them in standard text files. The multimaster replication model of Active Directory removes the need for secondary zones when all zones are stored in Active Directory.

- The DNS Server service includes an option to initialize the DNS Server service by reading parameters stored in the Active Directory database and the server registry. This is the default boot option. For more information, see Modifying server defaults.

- For more information about configuring DNS for Active Directory integration, see Configure a DNS server for use with Active Directory and Checklist: Verifying DNS before installing Active Directory.

# DNS zone replication in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS zone replication in Active Directory**
DNS zones can be stored in the domain or application directory partitions of Active Directory. A partition is a data structure within Active Directory used to distinguish data for different replication purposes. For more information, see Active Directory integration.

**Important**

- If you are changing the storage of a zone from the domain partition to an application directory partition, such as following the promotion of a new Windows Server 2003 domain controller in an existing Windows 2000 domain, the domain controller holding the domain naming master role must be running Windows Server 2003 for the DNS application directory partitions to exist. If you receive an error when changing the storage of a zone from the domain partition to an application directory partition, transfer the domain naming master role to a domain controller running Windows Server 2003, create the default DNS application directory partitions, and try again. For more information about creating the default DNS application directory partitions and transfering the domain naming master role, see Create the default DNS application directory partitions; Transfer the domain naming master role.

The following table describes the available zone replication scopes for Active Directory-integrated DNS zone data.

| Zone replication scope | Description |
|---|---|
| All DNS servers in the Active Directory forest | Replicates zone data to all DNS servers running on domain controllers in the Active Directory forest. Usually, this is the broadest scope of replication. |
| All DNS servers in the Active Directory domain | Replicates zone data to all DNS servers running on domain controllers in the Active Directory domain. This option is the default setting for Active Directory-integrated DNS zone replication in the Windows Server 2003 family. |
| All domain controllers in the Active Directory domain | Replicates zone data to all domain controllers in the Active Directory domain. If you want Windows 2000 DNS servers to load an Active Directory zone, this setting must be selected for that zone. |
| All domain controllers in a specified application directory partition | Replicates zone data according to the replication scope of the specified application directory partition. For a zone to be stored in the specified application directory partition, the DNS server hosting the zone must be enlisted in the specified application directory partition. For more information, see Create a DNS application directory partition and Enlist a DNS server in a DNS application directory partition. |

When deciding which replication option to choose, consider that the broader the replication scope, the greater the network traffic caused by replication. For example, if you choose to have Active Directory-integrated DNS zone data replicated to all DNS servers in the forest, this will produce greater network traffic than replicating the DNS zone data to all DNS servers in a single Active Directory domain in that forest. For information about changing zone replication scope, see Change zone replication scope.

**Notes**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

- Active Directory-integrated DNS zone data stored in an application directory partition is not replicated to the global

catalog for the forest. The domain controller that contains the global catalog can also host application directory partitions, but it will not replicate this data to its global catalog.

Active Directory-integrated DNS zone data stored in a domain partition is replicated to all domain controllers in its Active Directory domain and a portion of this data is stored in the global catalog. This setting is used to support Windows 2000.

- If an application directory partition's replication scope replicates across Active Directory sites, replication will occur with the same intersite replication schedule as used for domain partition data.

- Root hints are stored in the domain-wide application directory partition if the domain functional level is set to Windows Server 2003 . If the domain functional level is set to Windows 2000 mixed or Windows 2000 native, root hints are stored in the domain partition. For more information about functional levels, see Domain and forest functionality.

- By default, the Net Logon service registers domain controller locator (Locator) DNS resource records for the application directory partitions hosted on a domain controller in the same manner as it registers domain controller locator (Locator) DNS resource records for the domain partition hosted on a domain controller.

# DNS requirements for installing Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS requirements for installing Active Directory**

When you install Active Directory on a member server, the member server is promoted to a domain controller. Active Directory uses DNS as the location mechanism for domain controllers, enabling computers on the network to obtain IP addresses of domain controllers.

During the installation of Active Directory, the service (SRV) and address (A) resource records are dynamically registered in DNS, which are necessary for the successful functionality of the domain controller locator (Locator) mechanism.

To find domain controllers in a domain or forest, a client queries DNS for the SRV and A DNS resource records of the domain controller, which provide the client with the names and IP addresses of the domain controllers. In this context, the SRV and A resource records are referred to as Locator DNS resource records.

When adding a domain controller to a forest, you are updating a DNS zone hosted on a DNS server with the Locator DNS resource records and identifying the domain controller. For this reason, the DNS zone must allow dynamic updates (RFC 2136) and the DNS server hosting that zone must support the SRV resource records (RFC 2782) to advertise the Active Directory directory service. For more information about RFCs, see DNS RFCs.

If the DNS server hosting the authoritative DNS zone is not a server running Windows 2000 or Windows Server 2003, contact your DNS administrator to determine if the DNS server supports the required standards. If the server does not support the required standards, or the authoritative DNS zone cannot be configured to allow dynamic updates, then modification is required to your existing DNS infrastructure.

For more information, see Checklist: Verifying DNS before installing Active Directory and Using the Active Directory Installation Wizard.

**Important**

- The DNS server used to support Active Directory must support SRV resource records for the Locator mechanism to function. For more information, see Managing resource records.

- It is recommended that the DNS infrastructure allows dynamic updates of Locator DNS resource records (SRV and A) before installing Active Directory, but your DNS administrator may add these resource records manually after installation.

  After installing Active Directory, these records can be found on the domain controller in the following location: systemroot\System32\Config\Netlogon.dns

# DNS requirements for joining an Active Directory domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS requirements for joining an Active Directory domain**
Active Directory uses Domain Name System (DNS) to locate domain controllers, enabling computers joining the network to obtain a domain controller's IP address, and then begin the process of network authentication.

Computers joining an Active Directory domain must satisfy the following three DNS requirements:

- The computer must be configured with the IP address of a preferred DNS server.
- The **_ldap._tcp.dc._msdcs.***DNSDomainName*service (SRV) resource record must exist in DNS.
- The address (A) resource record for the DNS name of the domain controllers specified in the data field of the **_ldap._tcp.dc._msdcs.***DNSDomainName* SRV resource record must exist in DNS.

**Configuring a computer's preferred DNS server**
You must configure the computer with the static IP address of a preferred DNS server manually or assign the IP address of a preferred DNS server using Dynamic Host Configuration Protocol (DHCP).

For more information, see To configure a computer's preferred DNS server.

**Configuring the required DNS resource records**
Computers joining an Active Directory domain need the following resource records in DNS to locate a domain controller:

- **_ldap._tcp.dc._msdcs.** *DNSDomainName* SRV resource record, which identifies the name of the domain controller that hosts the Active Directory domain.

  *DNSDomainName* is the DNS name of the Active Directory domain the computer is attempting to join.

- A corresponding address (A) resource record that identifies the IP address for the domain controller listed in the **_ldap._tcp.dc._msdcs.***DNSDomainName* SRV resource record.

For more information, see To verify the DNS resource records needed to join an Active Directory domain using nslookup; Configure TCP/IP to use DNS.

# Understanding aging and scavenging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding aging and scavenging**

DNS servers running Windows Server 2003 support aging and scavenging features. These features are provided as a mechanism for performing cleanup and removal of stale resource records (RRs), which can accumulate in zone data over time.

With dynamic update, RRs are automatically added to zones when computers start on the network. However, in some cases, they are not automatically removed when computers leave the network. For example, if a computer registers its own host (A) RR at startup and is later improperly disconnected from the network, its host (A) RR might not be deleted. If your network has mobile users and computers, this situation can occur frequently.

If left unmanaged, the presence of stale RRs in zone data might cause some problems. The following are examples:

- If a large number of stale RRs remain in server zones, they can eventually take up server disk space and cause unnecessarily long zone transfers.
- DNS servers loading zones with stale RRs might use outdated information to answer client queries, potentially causing the clients to experience name resolution problems on the network.
- The accumulation of stale RRs at the DNS server can degrade its performance and responsiveness.
- In some cases, the presence of a stale RR in a zone could prevent a DNS domain name from being used by another computer or host device.

To solve these problems, the DNS Server service has the following features:

- Time stamping, based on the current date and time set at the server computer, for any RRs added dynamically to primary-type zones. In addition, time stamps are recorded in standard primary zones where aging/scavenging is enabled.

  For RRs that you add manually, a time stamp value of zero is used, indicating that they are not affected by the aging process and can remain without limitation in zone data unless you otherwise change their time stamp or delete them.

- Aging of RRs in local data, based on a specified refresh time period, for any eligible zones.

  Only primary type zones that are loaded by the DNS Server service are eligible to participate in this process.

- Scavenging for any RRs that persist beyond the specified refresh period.

  When a DNS server performs a scavenging operation, it can determine that RRs have aged to the point of becoming stale and remove them from zone data. Servers can be configured to perform recurring scavenging operations automatically, or you can initiate an immediate scavenging operation at the server.

  For more information, see either Enable automatic scavenging of stale resource records or Start immediate scavenging of stale resource records.

**Caution**

- By default, the aging and scavenging mechanism for the DNS Server service is disabled. It should only be enabled when all parameters are fully understood. Otherwise, the server could be accidentally configured to delete records that should not be deleted. If a record is accidentally deleted, not only will users fail to resolve queries for that record, but any user can create the record and take ownership of it, even on zones configured for secure dynamic update.

The server uses the contents of each RR-specific time stamp, along with other aging/scavenging properties that you can adjust or configure, to determine when it scavenges records.

**Prerequisites for aging/scavenging**

Before the aging and scavenging features of DNS can be used, several conditions must be met:

1. *Scavenging and aging must be enabled both at the DNS server and on the zone.*

   By default, aging and scavenging of resource records is disabled.

2. *Resource records must either be dynamically added to zones or manually modified to be used in aging and scavenging operations.*

   Typically, only those resource records added dynamically using the DNS dynamic update protocol are subject to aging and scavenging. For more information, see Dynamic update.

   You can, however, enable scavenging for other resource records added through non-dynamic means. For records added to zones in this way, either by loading a text-based zone file from another DNS server or by manually adding them to a zone, a time stamp of zero is set. This makes these records ineligible for use in aging/scavenging operations.

   In order to change this default, you can administer these records individually, to reset and permit them to use a current (non-zero) time stamp value. This enables these records to become aged and scavenged.

   For more information, see Reset scavenging and aging properties for a specified resource record.

### Note

- In the case of changing a zone from standard primary to Active Directory-integrated, you might want to enable scavenging of all existing resource records in the zone. To enable aging for all existing resource records in a zone, you can use the *AgeAllRecords* command, which is available through the **dnscmd** command-line tool. For more information, see Server administration using Dnscmd.

**Aging/scavenging terminology**

The following table indicates new or revised terms that have been introduced to help specifically when discussing aging and scavenging.

| Term | Description |
|---|---|
| Resource record (RR) time stamp | A date and time value used by the DNS server to determine removal of the resource record when it performs aging and scavenging operations. |
| Current server time | The current date and time on the DNS server. This number can be expressed as an exact numeric value at any point in time. |
| No-refresh interval | An interval of time, determined for each zone, as bounded by the following two events:<br>1. The date and time when the record was last refreshed and its time stamp set.<br>2. The date and time when the record next becomes eligible to be refreshed and have its time stamp reset.<br><br>This value is needed to decrease the number of write operations to the Active Directory database. By default, this interval is set to 7 days. It should not be increased to an unreasonably high level, because the benefits of the aging and scavenging feature might either be lost or diminished. |

| | |
|---|---|
| Refresh interval | An interval of time, determined for each zone, as bounded by the following two distinct events:<br>1. The earliest date and time when the record becomes eligible to be refreshed and have its time stamp reset.<br>2. The earliest date and time when the record becomes eligible to be scavenged and removed from the zone database.<br><br>This value should be large enough to allow all clients to refresh their records. By default, this interval is set to 7 days. It should not be increased to an unreasonably high level, because the benefits of the aging and scavenging feature might either be lost or diminished. |
| Start scavenging time | A specific time, expressed as a number. This time is used by the server to determine when a zone becomes available for scavenging. For more information, see Understanding aging and scavenging. |
| Scavenging period | When automatic scavenging is enabled at the server, this period represents the time between repetitions of the automated scavenging process. The default value for this is 7 days. To prevent deterioration of DNS server performance, the minimum allowed value for this is 1 hour. |
| Record refresh | When a DNS dynamic update is processed for a resource record when only the resource record time stamp, and no other characteristics of the record, are revised.<br><br>Refreshes generally occur for the following reasons:<br>1. When a computer is restarted on the network and, if at startup, its name and IP address information are consistent with the same name and address information it used prior to being shut down, it sends a refresh to renew its associated resource records for this information.<br>2. A periodic refresh is sent by the computer while it is running.<br><br>The Windows XP and Windows Server 2003 DNS Client service renews DNS registration of client resource records every 24 hours. When this dynamic update occurs, if the dynamic update request does not cause modification to the DNS database, then it is considered to be a refresh and not a resource record update.<br>3. Other network services make refresh attempts, such as: DHCP servers which renew client address leases, cluster servers which register and update records for a cluster, and the Net Logon service, which can register and update resource records used by Active Directory domain controllers. |
| Record update | When a DNS dynamic update is processed for a resource record where other characteristics of the record in addition to its time stamp are revised.<br><br>Updates generally occur for the following reasons:<br>1. When a new computer is added to the network and, at startup, it sends an update to register its resource records for the first time with its configured zone.<br>2. When a computer with existing records in the zone has a change in IP address, causing updates to be sent for its revised name-to-address mappings in DNS zone data.<br>3. When the Net Logon service registers a new Active Directory domain controller. |

| | |
|---|---|
| Scavenging servers | An optional advanced zone parameter that enables you to specify a restricted list of IP addresses for DNS servers that are enabled to perform scavenging of the zone.<br><br>By default, if this parameter is not specified, all DNS servers that load a directory-integrated zone (also enabled for scavenging) attempt to perform scavenging of the zone. In some cases, this parameter can be useful if it is preferable that scavenging only be performed at some servers loading the directory-integrated zone.<br><br>To set this parameter, you must specify the list of IP addresses for the servers enabled to scavenge the zone in the **ZoneResetScavengeServers** parameter for the zone. This can be done using the **dnscmd** command, a command-line based tool for administering Windows DNS servers. For more information, see Server administration using Dnscmd. |

**When scavenging can start**

Once all prerequisites for enabling the use of scavenging are met, it can start for a server zone when the current server time is greater than the value of the start scavenging time for the zone.

The server sets the time value to start scavenging on a per-zone basis whenever one of the following events occurs:

- Dynamic updates are enabled for the zone.
- A change in the state of the **Scavenge stale resource records** check box is applied. You can use the DNS console to modify this setting at either an applicable DNS server or one of its primary zones.
- The DNS server loads a primary zone enabled to use scavenging.

  This can occur when the server computer is started or when the DNS Server service is started.
- When a zone resumes service after having been paused.

When the previous events occur, the DNS server sets the value of start scavenging time by calculating the following sum:

Current server time + Refresh interval = Start scavenging time

This value is used as a basis of comparison during scavenging operations.

**Example of the aging/scavenging process for a sample record**

To understand the process of aging and scavenging at the server, consider the life span and successive stages of a single resource record, as it is added to a server and zone where this process is in effect and then aged and removed from the database.

1. A sample DNS host, "host-a.example.microsoft.com", registers its host (A) resource record at the DNS server for a zone where aging/scavenging is enabled for use.
2. When registering the record, the DNS server places a time stamp on this record based on current server time.

   After the record time stamp is written, the DNS server does not accept refreshes for this record for the duration of the zone no-refresh interval. It can, however, accept updates prior to that time. For example, if the IP address for "host-a.example.microsoft.com" changes, the DNS server can accept the update. In this case, the server also updates (resets) the record time stamp.
3. Upon expiration of the no-refresh period, the server begins to accept attempts to refresh this record.

   Once the initial no-refresh period ends, the refresh period immediately begins for the record. During this time, the server does not suppress attempts to refresh the record for its remaining life span.
4. During and after the refresh period, if the server receives a refresh for the record, it processes it.

   This resets the time stamp for the record based on the method described in step 2.
5. When subsequent scavenging is performed by the server for the "example.microsoft.com" zone, the record (and all other

zone records) are examined by the server.

Each record is compared to current server time on the basis of the following sum to determine whether the record should be removed:

Record time stamp + **No-refresh interval** for zone + **Refresh interval** for zone

- If the value of this sum is greater than current server time, no action is taken and the record continues to age in the zone.
- If the value of this sum is less than current server time, the record is deleted both from any zone data currently loaded in server memory and also from the applicable DnsZone object store in Active Directory for the directory-integrated "example.microsoft.com" zone.

**Note**

- For more information, see Managing the aging and scavenging of server data.

# Unicode character support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Unicode character support**

Originally, Internet host names were restricted to the character set specified in Request for Comments (RFCs) 952 and 1123. These restrictions include limiting names to using uppercase and lowercase letters (A-Z, a-z), numbers (0-9) and hyphens (-). In addition, the first character of the DNS name can be a number and names must be encoded and represented using US-ASCII-based characters.

These requirement were maintained when DNS was introduced as part of RFC 1035, one of the core DNS standards specifications. For use of DNS in international settings, this requirement has significant limitations where extended character sets are used for local naming standards.

To remove these limitations, Microsoft expands DNS character support beyond the RFC 1035 specification. The DNS service now provides enhanced default support for UTF-8, a Unicode transformation format.

For more information about RFCs, see the RFC Editor Web site.

**What is UTF-8?**

UTF-8 is the recommended character set for protocols evolving beyond the use of ASCII. The UTF-8 protocol provides for support of extended ASCII characters and translation of UCS-2, a 16-bit Unicode character set that encompasses most of the world's writing systems. UTF-8 enables a far greater range of names than can be achieved using ASCII or extended ASCII encoding for character data.

Computers running Windows 2000, Windows XP, and Windows Server 2003 operating systems are UTF-8 aware. This means that when UTF-8 encoded characters are received or used as data by the server, the server can load and store this data in its zones. Although Windows-based DNS servers are UTF-8 aware, they remain compatible with other DNS servers that use traditional US-ASCII data encoding and current DNS standards.

**How the DNS service implements UTF-8**

To provide standards compatibility and interoperability with other DNS implementations, the DNS service uses uniform downcasing of any received character data. In this process, the DNS service converts all uppercase characters used in standard US-ASCII data to lowercase equivalent data for the following reasons:

- To maintain compatibility with current and existing DNS standards.
- To provide interoperability with DNS server implementations that do not recognize or support UTF-8 encoding.

To understand why uniform downcasing was chosen, several related points must first be considered from the current revised Internet standards for DNS. Several key points in the standards pertain directly to how character data is to be handled between DNS servers and other servers and clients. These include the following:

- Any binary string can be used in a DNS name. (RFC 2181)
- DNS servers must be able to compare names in a case-insensitive way. (RFC 1035)
- The original case for character data should be preserved whenever possible as the data is entered into the system. (RFC 1035)

Because case insensitivity is a required part of the core DNS standard and case preservation is an optional recommendation, uniform downcasing was chosen to provide an effective standards-compliant solution. By downcasing UTF-8 encoded names before transmission, other DNS servers (which are not UTF-8 aware) are able to receive and perform successful binary comparisons of the data and obtain the desired results.

**Considerations for interoperability with UTF-8**

The DNS Server service can be configured to allow or disallow the use of UTF-8 characters on a per-server basis. Although other DNS server implementations that are not UTF-8 aware might be able to accept the transfer of a zone containing UTF-8 encoded names, these servers might not be able to write back those names to a zone file or reload those names from a zone

file. Administrators should exercise caution when transferring a zone containing UTF-8 names to a DNS server that is not UTF-8-aware.

Some protocols place restrictions on the characters allowed in a name. In addition, names that are intended to be globally visible (RFC 1958) should contain ASCII-only characters, as recommended in RFC 1123.

**Notes**

- The use of UTF-8 for transformation of Unicode characters is not noticeable for general users. Only where Network Monitor or another similar tool is used to analyze DNS-related traffic over the physical network are UTF-8 encoded characters observable.

- In addition to DNS server support for the UTF-8 encoding format, the client resolver defaults to using the UTF-8 character encoding format.

- Names encoded in UTF-8 format must not exceed the size limits clarified in RFC 2181, which specifies a maximum of 63 octets per label and 255 octets per name. Character count is insufficient to determine size because some UTF-8 characters exceed one octet in length.

- The UTF-8 encoding protocol adapts to use with existing DNS protocol implementations that expect US-ASCII characters because representation of US-ASCII characters in UTF-8 is identical, byte for byte, to the US-ASCII representation. DNS client or server implementations that do not recognize UTF-8 characters always encode names in the US-ASCII format. Those names are correctly interpreted by the DNS Server service.

- The DNS service provides the ability to configure name checking to allow or restrict the use of UTF-8 characters in DNS data. For more information, see Change the name-checking method used by the DNS server.

  By default, multibyte UTF-8 name checking is used, allowing the greatest tolerance when the DNS service processes characters. This is the preferred name-checking method for most privately operated DNS servers that are not providing name service for Internet hosts.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# WINS lookup integration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**WINS lookup integration**

Support for using Windows Internet Name Service (WINS) is provided to look up DNS names that cannot be resolved by querying the DNS domain namespace. To accomplish WINS lookup, two specific resource record types are used and can be enabled for any zones loaded by the DNS service:

- The WINS resource record, which can be enabled to integrate WINS lookup into forward lookup zones
- The WINS-R resource record, which can be enabled to integrate node adapter status request for reverse lookup zones

**The WINS resource record**

The WINS and DNS services are used to provide name resolution for the NetBIOS namespace and the DNS domain namespace, respectively. Although both DNS and WINS can provide a separate and useful name service to clients, WINS is mainly needed to provide support for older clients and programs that require support for NetBIOS naming.

However, the DNS service can work with WINS to provide combined name searches in both namespaces when resolving a DNS domain name not found in zone information. To provide this interoperability, a new record (the WINS record) was defined as part of the zone database file.

The WINS resource record is specific to computers running Windows NT 4.0 and earlier, Windows 2000, and Windows Server 2003 operating systems and can be attached only to the domain of origin for a zone. The presence of a WINS resource record can instruct the DNS service to use WINS to look up any forward queries for host names or names that are not found in the zone database. This functionality is particularly useful for name resolution required by clients that are not WINS-aware (for example, UNIX) for the names of computers not registered with DNS, such as Windows 95 or Windows 98 computers.

**How WINS lookup works**

The following is an example of a DNS client (*host-b*) querying its DNS server in an attempt to look up the address for another computer named *host-a.example.microsoft.com*..




In Step 1, the client queries its preferred DNS server. In Steps 2 through 8, the normal process of recursion proceeds as the preferred DNS server queries other DNS servers in succession on behalf of the client. This process concludes at Step 8, when the DNS server for the example.microsoft.com zone is located through the previous chain of referral answers. At this point in the process, the server contacted is a DNS server either running Windows NT Server 4.0, Windows 2000, or Windows Server 2003 .

When the DNS server for the example.microsoft.com zone receives the query for *host-a*, it looks in its configured zone to see if a matching address (A) resource record (RR) can be found. If no A record is found and the zone is enabled to use WINS lookup, the server does the following:

- The DNS server separates the host part of the name (host-a) from the fully qualified domain name contained in the DNS query.

The host part of the name is the first label in the queried DNS domain name before a period is used in the name.

- The server then sends a NetBIOS name request to the WINS server using the host name, *host-a*.
- If the WINS server can resolve the name, it returns the IP address to the DNS server.
- The DNS server then compiles an A resource record using the IP address resolved through the WINS server and returns this record to the original preferred DNS server that was queried by the requesting client, *host-b*.
- The preferred DNS server then passes the query answer back to the requesting client.

**How WINS reverse lookup works**

There is also a WINS-R record or WINS reverse lookup entry that can be enabled and added to reverse lookup zones. However, because the WINS database is not indexed by IP address, the DNS service cannot send a reverse name lookup to WINS to get the name of a computer given its IP address.

Because WINS does not provide reverse lookup capability, the DNS service instead sends a *node adapter status request* directly to the IP address implied in the DNS reverse query. When the DNS server gets the NetBIOS name from the node status response, it appends the DNS domain name back onto the NetBIOS name provided in the node status response and forwards the result to the requesting client.

### Note

- The WINS and WINS-R resource records are proprietary to the DNS Server service provided by Windows. You can prevent these resource records from being included in zone transfers to other DNS server implementations. For more information, see Using WINS lookup.

# Deploying DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying DNS**
This section covers:

- Namespace planning for DNS
- Zone planning for DNS
- Server planning for DNS
- Migrating servers
- Interoperability issues
- Securing DNS deployment

# Namespace planning for DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Namespace planning for DNS**
Before you begin using DNS on your network, decide on a plan for your DNS domain namespace. Coming up with a namespace plan involves making some decisions about how you intend to use DNS naming and what goals you are trying to accomplish in using DNS. Some questions you might have at this stage include the following:

- Have you previously chosen and registered a DNS domain name for use on the Internet?
- Are you going to set up DNS servers on a private network or the Internet?
- Are you going to use DNS to support your use of Active Directory?
- What naming requirements do you need to follow when choosing DNS domain names for computers?

Each of these questions is further discussed in the following sections.

**Choosing your first DNS domain name**
When setting up DNS servers, it is recommended that you first choose and register a unique parent DNS domain name that can be used for hosting your organization on the Internet, for example, "microsoft.com". This name is a second-level domain within one of the top-level domains used on the Internet. For a list and description of the most common top-level domains used on the Internet, see Top-level domains.

Once you have chosen your parent domain name, you can combine this name with a location or organizational name used within your organization to form other subdomain names. For example, if a subdomain were added, such as the itg.example.microsoft.com domain tree (for resources used by the information technology group at your organization), additional subdomain names could be formed using this name. For instance, a group of programmers working on electronic data interchange (EDI) in this division could have a subdomain named edi.itg.example.microsoft.com. Likewise, another group of workers providing support in this division might use support.itg.example.microsoft.com.

Before you decide on a parent DNS domain name for your organization to use on the Internet, perform a search to see if the name is already registered to another organization or person. The Internet DNS namespace is currently managed by the Internet Network Information Center (InterNIC). In the future, other domain name registrars might also be available.

For more information, see Interoperability issues.

**DNS namespace planning for Active Directory**
If you are using Active Directory, you need to first establish a namespace plan. Before a DNS domain namespace can be properly implemented, the Active Directory structure needs to be available. Therefore, begin with the Active Directory design and support it with the appropriate DNS namespace. Upon review, if you detect unforeseen or undesirable consequences for either of your plans, make revisions as needed.

Active Directory domains are named with DNS names. When choosing DNS names to use for your Active Directory domains, start with the registered DNS domain name suffix that your organization has reserved for use on the Internet, such as microsoft.com and combine this name with either geographical or divisional names used in your organization to form full names for your Active Directory domains.

For example, the test group at Microsoft might call their domain test.example.microsoft.com. This method of naming ensures that each Active Directory domain name is globally unique. And, once employed, this naming method also makes it easy to use existing names as parents for creating additional subdomains and further grow the namespace to accommodate new departments within your organization.

For a small organization using only a single domain or a small multidomain model, planning can be straightforward and follow an approach similar to the previous examples. For a larger organization using a more complex domain model, refer to additional Microsoft resources. For more information, see Using the Windows Deployment and Resource Kits.

**Caution**

- When planning your DNS and Active Directory namespace, it is recommended that you use a differing set of

distinguished names that do not overlap as the basis for your internal and external DNS use. As an example, assuming your organization's parent domain name is "example.microsoft.com"

For internal DNS names usage you could use a name such as "internal.example.microsoft.com"

For external DNS names usage you could use a name such as "external.example.microsoft.com"

By keeping your internal and external namespaces separate and distinct in this way, you enable simplified maintenance of configurations such as domain name filter or exclusion lists.

**Choosing names**

It is strongly recommended that you use only characters in your names that are part of the Internet standard character set permitted for use in DNS host naming. Allowed characters are defined in Request for Comments (RFC) 1123 as follows: all uppercase letters (A-Z), lowercase letters (a-z), numbers (0-9), and the hyphen (-).

For organizations with a prior investment in Microsoft NetBIOS technology, existing computer names might conform to the NetBIOS naming standard. If this is the case, consider revising the names of your computers to the Internet DNS standard.

The process of adjusting your naming conventions might prove to be time consuming. To ease the transition from NetBIOS names to DNS domain names, the DNS Server service includes support for extended ASCII and Unicode characters. However, this additional character support can only be used in a network environment with computers running Windows 2000 or a product in the Windows Server 2003 family. This is because most other DNS resolver client software is based on RFC 1123, the specification that standardizes Internet host naming requirements. If a nonstandard DNS domain name is entered during setup, a warning message appears recommending that a standard DNS name be used instead.

In networks running Windows NT 4.0 and earlier, a NetBIOS name is used to identify a computer running a Windows operating system. In networks with computers running Windows 2000 or a product in the Windows Server 2003 family, a computer can be identified in any of the following ways:

- A NetBIOS computer name, which is optional and is used for interoperability with earlier Windows systems.
- The full computer name. This is the default name for the computer.

In addition to these, a computer might also be identified by the FQDN comprised of the computer (host) name and a connection-specific domain name, where one is configured and applied for a specific network connection on the computer.

The full computer name is a concatenation of both the computer name and the primary DNS suffix for the computer. The DNS domain name for the computer is part of the System properties for the computer and is not related to any specifically installed networking components. However, computers that do not use either networking or TCP/IP do not have a DNS domain name.

The following table is a comparison of NetBIOS and DNS computer names.

| Restriction | DNS in Windows NT 4.0 (standard DNS) | DNS in Windows 2000 and the Windows Server 2003 family | NetBIOS |
|---|---|---|---|
| Characters | Supports RFC 1123, which allows all uppercase letters (A-Z), lowercase letters (a-z), numbers (0-9), and the hyphen (-). | Supports RFC 1123 and UTF-8. You can configure the DNS server to allow or disallow the use of UTF-8 characters. You can do so on a per-server basis. For more information, see Unicode character support. | Not allowed: Unicode characters, numbers, white space, symbols: / \ [ ] : \| < > + = ; , ? and * |
| Host name and FQDN length. | 63 octets per label. 255 bytes per FQDN (254 bytes for the FQDN plus one byte for the terminating dot). | The same as standard DNS with the addition of UTF-8 support. Character count is insufficient to determine size because some UTF-8 characters exceed one octet in length. Domain controllers are limited to 155 bytes for an FQDN. Domain controllers are limited to 155 bytes for an FQDN. | 15 bytes in length. |

To ensure interoperability between NetBIOS and DNS naming in Windows, a new naming parameter called the NetBIOS computer name was introduced. The value of this parameter, which is not required in a Windows 2000 or Windows Server 2003 environment, is derived from the first 15 characters of the DNS full computer name.

When the full computer name is a combination of the computer name and the primary DNS suffix for the computer, the impact of renaming and making the transition from a NetBIOS namespace to a DNS namespace can be minimal. Users continue to focus on the short computer name. If this name is 15 characters or less, it can be made identical to the NetBIOS computer name. The administrator can then also assign a DNS domain name for each computer. This can be done using remote administration tools.

In this way, the **System** properties page can have the following settings on the **Computer Name** tab (or its related dialog boxes):

| | |
|---|---|
| **Computer name** | longer-than-15characters-name |
| **NetBIOS computer name** | longer-than-15c |
| **Primary DNS suffix for this computer** | example.microsoft.com |

**Important**

- By default, the primary DNS suffix portion of a computer's fully qualified domain name (FQDN) must be the same as the name of the Active Directory domain where the computer is located. To allow different primary DNS suffixes, a domain administrator may create a restricted list of allowed suffixes by creating the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is created and managed by the domain administrator using Active Directory Service Interfaces or the Lightweight Directory Access Protocol (LDAP).

  For more information, see Programming interfaces and Directory access protocol.

**Note**

- In addition to the primary DNS domain name for a computer, connection-specific DNS names can be configured. For more information, see Configuring multiple names.

  Because DNS host names are encoded in UTF-8 format, they do not necessarily have only 1 byte per character. ASCII characters are 1 byte each, but the size of extended characters is more than 1 byte.

  Some non-Microsoft resolver software supports only the characters listed in RFC 1123. If you have any non-Microsoft resolver software, that software is probably not able to look up computers with names that have non-standard characters.

  A DNS server that does not support UTF-8 encoding might accept a zone transfer of a zone containing UTF-8 names, but it cannot write back those names to a zone file or reload those names from a zone file. Therefore, you must not transfer a zone that contains UTF-8 characters to a DNS server that does not support them.

**Integration planning: Supporting multiple namespaces**

In addition to support for the internal DNS namespace, many networks require support for resolving external DNS names, such as those used on the Internet. The DNS Server service provides ways to integrate and manage disjointed namespaces where both external and internal DNS names are resolved on your network.

In deciding how to integrate namespaces, determine which of the following scenarios most closely resembles your situation and proposed use of DNS:

- An internal DNS namespace, used only on your own network.

- An internal DNS namespace with referral and access to an external namespace, such as referral or forwarding to a DNS server on the Internet.

- An external DNS namespace, used only on a public network such as the Internet.

If you decide to limit the use of DNS as a name service within your private namespace, there are no restrictions on how it is designed or implemented. For example, you can choose any DNS naming standard, configure DNS servers to be valid root servers for your network's DNS distributed design, or form an entirely self-contained DNS domain tree structure and

hierarchy.

When you start providing either referral to an external DNS namespace or full DNS service on the Internet, you need to consider compatibility between your private and external namespaces. Additionally, Internet service requires the registration of parent domain names for your organization.

For more information on planning secure integration of Internet DNS names referral through your secured private network, see the Microsoft Windows Resource Kits Web site.

**Notes**

- When you add and connect to a DNS server for the first time using the DNS console, the root hints file (Cache.dns) is automatically primed for use on your network. This occurs transparently and does not require that you take any further action for this to be done.

  Depending on how you use DNS, you can update this file in one of the following ways:

  - If you are connected to the Internet, you can either edit or update the local root hints file when the Internet root hints file (Named.root) is updated and released by the owners of the Internet root zone. For a current copy of this file, you can use anonymous FTP to the InterNIC Web site.

  - If you are not connected to the Internet, you can remove the default resource records contained in this file and replace them with NS and A resource records for the DNS authoritative servers at the root domain of your site. For root domain servers on your private network, you can safely remove this file entirely because servers operating at this domain level do not require or use a cache of root hints.

- If you delete the Cache.dns file for a private root domain server when you have configured the DNS Server service using the **From file** method, you also need to remove the **cache** directive from the server boot file before restarting the DNS Server service.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Zone planning for DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Zone planning for DNS**

When first partitioning your namespace into zones, review traffic patterns in your current or proposed network. Although DNS is designed to help reduce broadcast traffic between local subnets, it does create some traffic between servers and clients that should be reviewed. This is particularly true in cases where DNS is used on routed networks.

To review DNS traffic, you can use DNS server statistics or the DNS performance counters provided with System Monitor. For more information, see Verifying server configuration or Monitoring DNS server performance.

In addition to routing traffic, consider the impact of the following common types of DNS-related communication, especially when you are operating across slow-speed links on a wide area network:

- Server-to-server traffic caused both by zone transfers with other DNS servers and by DNS interoperability with other servers (for example, when WINS lookup is enabled).

- Client-to-server traffic caused by both query loads and dynamic updates sent by DNS client computers or DHCP servers providing dynamic updating for earlier-version DNS clients that do not support dynamic updates.

For small, flat namespaces, you might use full replication of all DNS zones to all DNS servers in your network. For large, vertical namespaces, this is neither possible nor recommended. In larger networks, it is often necessary to study, test, analyze, and revise your zone plans based on observed or estimated traffic patterns. After careful analysis, you can partition and delegate your DNS zones based on what is required for providing efficient and fault-tolerant name service to each location or site.

**Notes**

- If you are using a mixture of Windows DNS servers, be aware that zone transfers can sometimes be slow between the DNS Server service and some other DNS server implementations. For more information, see Interoperability issues.

- The DNS Server service supports incremental zone transfers between servers replicating a standard zone. This feature can reduce DNS replication traffic considerations and should be reviewed in zone planning. For more information, see Understanding zones and zone transfer.

- Although it is useful to plan zones, you might also want to explore the use of caching-only servers, which do not host DNS zones. Caching-only servers are a good option at small remote sites that have a stable and minimal use for DNS name service but are located across a WAN where the transfer of a large zone over a slower link can be resource-intensive. For more information, see Using caching-only servers.

# Server planning for DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server planning for DNS**

When planning for your DNS servers, it is important to consider the following:

- Perform capacity planning and review server hardware requirements.
- Determine how many DNS servers you need and their role in your network.

  When deciding the number of DNS servers to use, decide which servers will host primary and secondary copies of zones. Also, if you are using Active Directory, determine whether the server computer will perform as a domain controller or a member server for the domain.
- Decide where you are going to place DNS servers on your network for traffic loads, replication, and fault tolerance.
- Decide whether to use only DNS servers running Windows Server 2003 for all your DNS servers or if you are operating a mixture of Windows and other DNS server implementations.

  For more information about migration and interoperability of the DNS Server service with other DNS server implementations, see Migrating servers or Interoperability issues.

**Server capacity planning**

Planning and deploying DNS servers on your network involves examining several aspects of your network and the capacity requirements for any DNS servers you intend to use in it. Some questions to consider when planning include the following:

- How many zones is the DNS server expected to load and host?
- For each zone the server is loading for service, how large is the zone (based on the size of the zone file or the number of resource records used in the zone)?
- For a multihomed DNS server, how many interfaces are to be enabled for listening to and servicing DNS clients on each of the server's connected subnets?
- How many total or overall DNS query requests from all of its clients is a DNS server expected to receive and service?

In many cases, adding more RAM to a DNS server can provide the most noticeable improvements in performance. This is because the DNS Server service fully loads all of its configured zones into memory at startup. If your server is operating and loading a large number of zones, and dynamic updates occur frequently for zone clients, additional memory can be helpful.

Keep in mind that for typical usage, the DNS server consumes system memory as follows:

- Approximately 4 MB of RAM is used when the DNS server is started without any zones.
- For each addition of zones or resource records to the server, the DNS server consumes additional server memory.
- It is estimated that for the addition of every resource record to a server zone, an average of approximately 100 bytes of server memory is used.

  For example, if a zone containing 1000 resource records is added to a server, it would require approximately 100 KB (kilobytes) of server memory.

In determining your DNS server plans, you can start by reviewing sample DNS server performance test results collected by the DNS development and testing teams. In addition, you can use DNS server-related counters provided for use with monitoring

tools to obtain your own performance measurements. For more information, see Monitoring DNS server performance.

**Important**

- The previous recommendations are not intended to indicate maximum performance or limitations for DNS servers.
- These numbers are approximate and can be influenced by the type of resource records entered in zones, the number of resource records with the same owner name, and the number of zones in use at a specific DNS server.

**Where to place DNS servers**

In general, place your DNS servers at a location on your network that is centrally accessible to your clients. It is often most practical to use a DNS server on each subnet. There are several factors to consider when deciding where a DNS server is needed:

- If you are deploying DNS to support Active Directory, is the DNS server computer also a domain controller or likely to be promoted to one in the future?
- If the DNS server stops responding, are its local clients able to gain access to an alternate DNS server?
- If the DNS server is located on a subnet that is remote to some of its clients, what other DNS servers or name resolution options are available if the routed connection stops responding?

For DNS server installations where the use of Active Directory is an issue, review special interoperability issues and installation details. For more information, see Interoperability issues.

For all DNS server installations including those in which the use of Active Directory is not an issue, the following server placement and planning guidelines can be usefully applied.

For example, if you have a routed local area network and high-speed links that are fairly reliable, you might be able to use one DNS server for a larger, multiple subnetted network area. If you have a high number of client nodes on a single subnet design, you might want to add more than one DNS server to the subnet to provide backup and failover if the preferred DNS server stops responding.

When determining the number of DNS servers you need to use, assess the effect of zone transfers and DNS query traffic on slower links in your network. Although DNS is designed to help reduce broadcast traffic between local subnets, it does create some traffic between servers and clients that should be reviewed, particularly when used in complexly routed LAN or WAN environments.

Consider the effects of zone transfer over slower speed links, like those typically used for a wide area network (WAN) connection. Although the DNS Server service supports incremental zone transfers and DNS clients and servers can cache recently used names, traffic considerations are sometimes still an issue, particularly when DHCP leases are shortened and, as a result, dynamic updates in DNS are performed more frequently. One option for dealing with remote locations on WAN links is to set up a DNS server at these locations to provide caching-only DNS service.

With most installations, you should have at least two server computers hosting each of your DNS zones for fault tolerance. DNS was designed to have two servers for each zone, one as a primary server and the other as a backup or secondary server. When making any final determinations about the number of servers to use, first assess the level of fault tolerance you need for your network.

**Notes**

- When only a single DNS server is used on a small LAN in a single-subnet environment, you can configure the single server to simulate both the primary and secondary servers for a zone.
- For best results and simplified DNS administration, consider using the Windows Server 2003 DNS Server service for all of your DNS servers.
- When administering zone files created by the DNS Server service, it is recommended that you use the DNS console tool to make changes. As an alternate option, it is possible to administer all of your DNS zone files manually using any application that supports file saving in text mode. Whether you decide to edit zone files manually in a text editor or use the DNS console to modify them, select one method for updating zones and use it consistently. This helps prevent having

zone edits overwritten or rejected, which might happen if you change from one method to another.

# Migrating servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Migrating servers**

DNS migration can occur in one of the following ways:

1. Upgrading a computer running the version of DNS included in Windows NT Server 4.0 or Windows 2000 to a server running Windows Server 2003 .

2. Moving zone files from an existing DNS server running another DNS server implementation, such as server running a version of the Berkeley Internet Name Domain (BIND) software.

3. Migrating zones using master-secondary zone transfer from BIND servers to DNS servers running Windows Server 2003 .

**Upgrading DNS servers from Windows NT Server 4.0 and Windows 2000**

Any zones or server configuration files and parameters created and stored using the Windows NT Server 4.0 or Windows 2000 version of the DNS Server service are stored in the same system folder locations. There is no conversion of data required during the upgrade process.

**Moving BIND-based files**

When moving from a BIND DNS server to a DNS server running Windows Server 2003 , however, you need to copy any BIND-created zone or boot files that you intend to use with the DNS Server service. Also, if you continue to use a BIND boot file to provide the initial configuration settings used by the DNS Server service when it is started, you need to change the boot method used by the DNS Server service or rename the zone files from the BIND naming convention to that used by DNS servers running under the DNS Server service provided in Windows operating systems. For more information, see Change the boot method used by the DNS server.

Any zone files created and stored on UNIX DNS servers that use BIND need to be manually copied from those servers to the systemroot\System32\Dns folder on the server running Windows Server 2003 .

If you do not migrate the BIND boot file or specify the BIND name when creating the zones using the DNS console, then you will need to rename these zones. The following table shows examples of how BIND server files are renamed for use with the DNS Server service that is provided on servers running a member of the Windows Server 2003 family.

| Description | UNIX file name | Windows Server 2003  file name |
|---|---|---|
| Boot file | named.boot | Boot |
| Forward lookup zone file | db.*domain_name* | *domain_name*.dns |
| Reverse lookup zone file | db.*IP_network_forward_notation* | *IP_network_reverse_notation*.dns |

For a reverse lookup zone, BIND uses a convention of noting reverse domain names as a forward notation of mapped IP network, such as db.192.1.168 for a reverse lookup zone created for an example IP network address of 192.168.1.0. DNS servers running Windows Server 2003 , however, use the fully qualified domain name (FQDN) for the zone, which includes the in-addr.arpa domain, to complete the file name. In this example, the correct name to use for the same zone in Windows Server 2003  DNS is 168.1.192.in-addr.arpa.dns when copying and renaming the file.

**Notes**

- If you are using the BIND boot file with the DNS Server service after migration, there are other limitations that apply to the use of this file by the DNS Server service. For example, some BIND boot directives are not supported -- in particular, **xfrnets** and other directives provided with versions of BIND, such as version 8.1.1 or later. For more information, see articles Q194513, "The Structure of a Domain Name System Boot File"; Q234144, "DNS Boot File Directives and Configuration for Windows NT 4.0," in the Microsoft Knowledge Base.

- If you are accustomed to manually editing DNS zone files, be aware that the DNS Server service uses RFC-compliant

notation for its supported resource records (RRs). In most cases, the DNS Server service interprets and loads RRs from zone files originally created for BIND DNS servers without any need for file changes. If, however, you have used non-standard record formatting, the DNS Server service can detect these edits and interpret them as bad or errored zone data. For more information about RFCs, see DNS RFCs.

**Migrating from third-party servers using zone transfer**

Another possible scenario for server migration is using zone transfer to migrate third-party server zones to DNS servers running Windows Server 2003 . To accomplish this, you can take the following steps:

1. At your server running Windows Server 2003 , install the DNS Server service.

2. Using the DNS console, at the new server add secondary zones for all of your existing zones hosted at the BIND-based DNS servers.

   Configure the third-party servers as the master servers for each of the secondary zones you need to create.

3. After completing the zone transfers, convert any of the secondary zones to primary zones for the zones that were obtained from primary zones at the third-party servers.

4. Configure the previous primary servers to be secondary servers for the migrated zones or demote the previous primary servers.

5. For the other secondary zones that remain, update the master servers for those zones to use the new primary DNS servers running Windows Server 2003 .

For more information, see Install a DNS server, Add a secondary server for an existing zone, Initiate a zone transfer at a secondary server, Change the zone type, and Update the master server for a secondary zone.

**Note**

- If you choose to continue to use your third-party DNS servers as secondary servers for zones for which your DNS server running Windows Server 2003 is the primary server, you should review interoperability issues related to zone transfer for this configuration. For more information, see Interoperability issues.

# Interoperability issues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interoperability issues**

DNS servers running Windows Server 2003 are compliant with most of the Request for Comments (RFC) specifications used to define the DNS protocol. This provides clear benefits for operating DNS servers in mixed or heterogeneous environments. For more information about RFCs, see DNS RFCs.

The primary benefits for interoperability in these environments include:

- Full interoperability with other DNS server implementations that implement RFC-compliant behavior for DNS name service.
- Use of Windows DNS servers to provide DNS service on the Internet.

For interoperability testing, the Windows Server 2003 DNS development team has tested Windows Server 2003 DNS Server and Client services with the following versions of the Berkeley Internet Name Domain (BIND) DNS server implementation:

- BIND 4.9.7
- BIND 8.1.2
- BIND 8.2
- BIND 9.1.0

Interoperability and configuration issues related to using Windows Server 2003 DNS with other various environments, or when using DNS servers on the Internet, are covered in the following sections.

**DNS considerations for Active Directory**

DNS is required for network resources to locate Active Directory domain controllers. You can elect to set up the DNS Server service as part of the installation of Active Directory using the Active Directory Installation Wizard. When you install Active Directory on a server, the server is promoted to the role of a domain controller.

When you install Active Directory on a server, the Active Directory Installation Wizard provides the option to automatically install the DNS Server service and add new zones locally. When you select the option to install and configure a DNS server using the Active Directory Installation Wizard, zones are created based on the DNS name you have specified in the wizard.

A simple method of redundancy and fault tolerance planning is to have a DNS server running on each domain controller. For each subnet, have two domain controllers running the DNS Server service and hosting Active Directory-integrated zones.

**Zone transfer with BIND and other DNS server implementations**

When transferring a zone between two Windows DNS servers, the DNS Server services can utilize a fast transfer method that uses compression. This method includes multiple resource records (RRs) in each message sent to complete the transfer of the zone between servers.

By default, DNS servers running Windows Server 2003 are configured to transfer a zone using the slower uncompressed transfer format. This enables successful zone transfers to be made with DNS servers that do not support the faster transfer method, such as BIND servers prior to version 4.9.4.

If you clear the **Bind Secondaries** check box in advanced server properties, fast transfers are enabled. For more information, see Enable or disable fast transfer format during zone transfers.

**Supporting Active Directory with other DNS server implementations**

In many large organizations, DNS is already implemented using other solutions, such as UNIX DNS servers that run legacy versions of BIND software. In some cases, these DNS servers are not equipped to support the DNS requirements for deploying Active Directory. This issue can be addressed in one of two ways:

- Upgrade any BIND DNS servers to version 8.1.2 or later of the BIND software to meet the DNS requirements for Active Directory support.

- Use the DNS Server service provided with Windows Server 2003 family, migrating, if possible, any of your current DNS zones to DNS servers running Windows Server 2003.

Although the DNS Server service is recommended to support Active Directory, you can use other DNS server implementations for this purpose. These other implementations should, however, support the following standard specifications:

- The service location (SRV) resource record, as described in the Internet draft, "A DNS RR for specifying the location of services (DNS SRV)."

- Dynamic updates in DNS, as described in RFC 2136.

Support for dynamic updates is recommended but not essential. Support for the SRV resource record is mandatory because it is required to provide basic DNS support to Active Directory. For example, a DNS server that does not support dynamic updates, like that provided with Windows NT Server 4.0 (updated to Service Pack 4 or later) supports the DNS requirements of Active Directory because SRV resource record support was added with Service Pack 4.

Additional manual administration of SRV resource records is needed for DNS configuration support of Active Directory to function properly on a DNS server that does not support dynamic updates. For more information, see Managing resource records.

**Interoperating Windows DNS servers with other DNS server implementations**
If you decide to use the Windows DNS Server service and manage it with a split DNS configuration in which:

- Existing DNS servers for root zones are not to be upgraded or migrated to other DNS solutions.

- Windows Server 2003 DNS is to be deployed and is to provide management of any DNS domain names required to register, update, and support for use with Active Directory.

You can modify your DNS namespace design plans in either of the following ways:

1. Create a single new subdomain in your current DNS domain namespace to root your first Active Directory domain.

   For example, if your organization has registered and is using a second-level domain name, such as microsoft.com, you can create a single subdomain such as example.microsoft.com and use this domain to root the DNS domain namespace used by Active Directory. The DNS Server service is automatically configured to support Active Directory when you install the first domain controller.

   Before you have created a zone for the new subdomain at a computer running the DNS Server service, you can delegate these subdomains away at the primary zone for your second-level domain, such as "microsoft.com." In some cases, you might only need to notify another DNS or UNIX system administrator in your organization to make the delegation for you. For more information, see Delegating zones.

2. Create multiple subdomains based on your DNS second-level domain to support registration of Active Directory in DNS.

   For example, if your organization has a registered second-level DNS domain name already in use (such as microsoft.com), you can create additional subdomains that are delegated to Windows DNS servers and used only for registering DNS names related to Active Directory.

   This method is more complex to implement, but enables less change to your currently deployed DNS infrastructure that is not Windows-based. With this namespace design, you create only those additional subdomains and appropriate zones needed to support your Active Directory deployment. For example, in this configuration, the domain name microsoft.com is both the root DNS and the root Active Directory domain name for your organization.

   For this configuration, you first need to create zones for the following subdomains using the DNS Microsoft Management Console (MMC) console on a DNS server running Windows Server 2003:

child1.microsoft.com child2.microsoft.com

Before these zones are created, you can delegate these subdomains away at the primary zone for your parent or second-level domain name or notify another DNS administrator who manages these zones for your organization to do so. For more information, see Delegating zones.

**Using DNS on the Internet**

To establish a presence on the Internet, an individual or business must first apply for and register a second-level domain name with an authorized DNS domain name registration authority. Your Internet service provider (ISP) can often perform this function and obtain a name on your behalf, usually for an additional fee.

To register your domain name, there are several required tasks, including:

- Selecting and researching a second-level domain name that is not currently registered or in use.

  This can easily be done if you have Internet access by using a WHOIS query engine provided at the Web site for your applicable Internet DNS domain name registrar. Be prepared to select an alternate name if your WHOIS query indicates that your preferred selection is already registered and in use.

- Registering and obtaining at least one IP address valid for use on the Internet.

  This address is needed for the DNS server on the Internet that you want to establish as the host for the primary copy of the zone based on your second-level domain name. In many cases, if you are using an ISP to register a domain name on your behalf, they can specify the IP addresses for one or two of their servers as primary and secondary for the Internet.

  As part of the registration process, an applicant must provide at least two currently active DNS servers that are used on the Internet as the primary and secondary servers designated for the new domain. This requirement is necessary to ensure proper Internet root server configuration and referral for others that query for your registered DNS domain name on the Internet.

  After one IP address has been obtained, you can sometimes arrange (or your ISP can do so on your behalf) to use another company or ISP's DNS server as a secondary server for the zone. If you still need to obtain an IP address directly for use in the United States, a valid IP address can be obtained through the American Registry for Internet Numbers (ARIN). In other countries or regions, you might contact your local Internet service or telephony provider to find out how to register and acquire an IP address if one is needed.

- Complete the registration application form and submit it with your registration fee to the appropriate Internet DNS domain name registration authority. Registrations are typically in force for a finite period of time and must be periodically renewed.

For more information about selecting and registering your Internet DNS name, see the VeriSign Global Registry Services Web site.

**Interoperability planning: Configuring related services**

Windows Server 2003 DNS provides several interoperability options with other TCP/IP services. The following options can be used to reduce the amount of time you need to spend administering your DNS infrastructure:

- **WINS forward and reverse lookups.**

  The DNS Server service provides for the use of WINS lookup. This feature enables configured DNS zones to refer queries not answered from current zone information to a WINS server for further resolution. With this added search of the WINS namespace, both DNS and WINS are used to complete a full search of registered names for a matched response.

  WINS lookup is supported for both forward and reverse lookup zones and can be enabled on a per-zone basis or configured for selected zones. This feature should also be configured to prevent replication or zone transfer of WINS resource records to servers with other DNS implementations that do not recognize the WINS resource records.

- **Dynamic integration with DHCP servers.**

For Windows Server 2003 DNS, the DHCP service provides default support to register and update information for legacy DHCP clients in DNS zones. Legacy clients typically include other Microsoft TCP/IP client computers that were released prior to Windows 2000. The Windows Server 2003 DNS-DHCP integration enables a DHCP client that is unable to dynamically update DNS resource records directly to have this information updated in DNS forward and reverse lookup zones by the DHCP server.

**Notes**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

- Dynamic integration with the DHCP service is only available on DNS servers running Windows 2000 and Windows Server 2003. DNS-DHCP integration is not supported by DHCP servers running under Windows NT Server 4.0 and earlier.

- By default, a zone hosted on a DNS server running Windows Server 2003 will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.

# Securing DNS deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing DNS deployment**

When designing your DNS server deployment, use the following DNS security guidelines:

- **If your network hosts are not required to resolve names on the Internet, eliminate DNS communication with the Internet.** In this DNS design, you can use a private DNS namespace that is hosted entirely in your network. The private DNS namespace is distributed just as the Internet DNS namespace, with your internal DNS servers hosting zones for the root domain and top-level domains.

- **Split the DNS namespace for your organization between internal DNS servers behind the firewall and external DNS servers in front of the firewall.** In this DNS design, your internal DNS namespace is a subdomain of your external DNS namespace. For example, if the Internet DNS namespace for your organization is adatum.com, then the internal DNS namespace for your network is corp.adatum.com.

- **Host your internal DNS namespace on internal DNS servers and host your external DNS namespace on external DNS servers exposed to the Internet.** To resolves queries for external names made by internal hosts, the internal DNS servers forward queries for external names to the external DNS servers. External hosts use only the external DNS servers for Internet name resolution.

- **Configure your packet-filtering firewall to only allow UDP and TCP port 53 communication between your external DNS server and a single internal DNS server.** This will facilitate communication between internal and external DNS servers and prevent any other external computer from gaining access to your internal DNS namespace.

For more information, see Security information for DNS.

# Administering DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering DNS**
This section covers:

- Managing DNS Servers
- Managing Clients
- Managing Zones
- Monitoring and Optimizing Servers

# Managing DNS Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing servers**

This section covers:

- Configuring a new primary server
- Update existing servers
- Using secondary servers
- Securing the DNS Server service
- Using caching-only servers
- Modifying server defaults
- Managing the aging and scavenging of server data
- Using forwarders
- Directing queries through forwarders
- Updating root hints
- Configuring multihomed servers
- Server administration using Dnscmd

# Configuring a new primary server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring a new primary server**

There are several situations in which you might add and configure a new DNS server for your network:

- When you are adding a new DNS server to your network and configuring a new zone for first-time use.
- When you have already created a zone at another DNS server and are adding a new server that also needs to load and provide service for the zone.
- When you already have a DNS server configured with a zone or zones, but need to add a new zone for another domain name, such as a subdomain.

For first-time configuration of a new server and zone, it is best to use the checklists provided with DNS Help to begin your DNS deployment. If you are going to deploy a DNS server without using Active Directory, use Checklist: Deploying DNS. For deploying DNS to support your use of Active Directory, use Checklist: Deploying DNS for Active Directory.

**Adding a primary DNS server for an existing zone**

The primary DNS server for a zone acts as the point of update for the zone. Newly created zones are always this type. You can use primary zones in one of two ways: as standard primary zones or primary zones integrated with Active Directory.

For standard primary-type zones, only a single DNS server can host and load the master copy of the zone. If you create a zone and keep it as a standard primary zone, no additional primary servers for the zone are permitted. Only one server is allowed to accept dynamic updates and process zone changes.

The standard primary model implies a single point of failure. For example, if for any reason the primary server for a zone is unavailable to the network, no dynamic updates to the zone can be made. Note that queries for names in the zone are not affected and can continue uninterrupted, as long as secondary servers for the zone are available to answer them.

You can add more primary servers for a zone by using the directory-integrated storage and replication features of the DNS Server service. To do this, you need to change a zone and integrate it into Active Directory.

You can integrate an existing zone into Active Directory by changing the type of a zone at the originating primary server where it was first created. Once the zone type is changed from **Standard primary** to **Active Directory-integrated**, you can add the zone to other DNS servers by configuring them to use the **From Active Directory or registry** option when they initialize DNS Server service.

When the **From Active Directory or registry** option is selected, other DNS servers operating as part of the Active Directory domain namespace -- such as domain controllers for the Active Directory domain -- can query the directory and automatically load all directory-integrated zones, which are stored in the directory database. No other steps are necessary. With the **From Active Directory or registry** option, any DNS servers operating as part of Active Directory are, by default, also primary servers for directory-integrated zones.

For more information on changing the zone type, see Change the zone type. For more information on changing the server boot method, see Change the boot method used by the DNS server.

**Note**

- When using standard primary-type zones, other DNS servers should never be configured to act as primary servers for an existing zone. This design is not supported and might cause errors or inconsistencies in zone data between servers loading different versions of the same zone.

**Adding a new primary zone to an existing server**

Adding a new primary zone to an existing server can be done whenever additional domains or subdomains are needed in your DNS domain namespace. For example, you might have a zone for a second-level domain such as microsoft.com and want to add a primary zone for the new subdomain, example.microsoft.com.

In this example, you can create the new zone for the subdomain by using the DNS console and running the New Zone Wizard.

When you are done, you need to create a delegation in the parent zone for the new subdomain (such as the microsoft.com zone) to complete the addition of the new subdomain and its primary zone.

Once the new primary zone is added, you can perform other zone configuration tasks as needed to fully configure the zone for use. For more information, see Managing Zones.

# Update existing servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Update existing servers**

Updating the configuration of your DNS servers can be necessary for various reasons, such as to:

- Change the computer (host) name of the server computer.
- Change the primary or parent DNS domain name of the server computer.
- Change the IP address for the server computer.
- Remove a DNS server from the network.
- Change the primary server for a zone (standard primary zones only).

Each of these is covered in the following sections.

**Changing the computer (host) name of the server**

If you need to change the host name of a DNS server (but not its parent or primary DNS domain name), make the following changes in the zone where the server is configured as an authoritative server for the zone:

1. Change the computer name for the server in System properties.
2. Restart the computer to initiate the following DNS dynamic updates:
   a. Add the new host (A) and pointer (PTR) resource records (RRs) for the new name and IP address of the server.
   b. Remove the old A and PTR RRs for the old name and IP address of the server.
3. Update the name server (NS) RR in zones where the server is configured as authoritative to point to the new A RR added in step 2.
4. If the server is the primary server for a standard zone, revise the name in the owner field of the start of authority (SOA) RR for the zone (If the zone is directory-integrated, this step is not needed).
5. Check in the zones to ensure that any delegation records -- NS or A RRs -- that are used, are updated for the new name of the server.

**Changing the primary DNS domain name of the server**

When changing the DNS domain name of a computer, the ability to change the primary DNS domain name of the server can depend on whether or not you are using the computer as a domain controller.

**Important**

- By default, the primary DNS suffix portion of a computer's fully qualified domain name (FQDN) must be the same as the name of the Active Directory domain where the computer is located. To allow different primary DNS suffixes, a domain administrator may create a restricted list of allowed suffixes by creating the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is created and managed by the domain administrator using Active Directory Service Interfaces or Lightweight Directory Access Protocol (LDAP).

  For more information, see Programming interfaces and Directory access protocol.

If you are running the server computer as a domain controller, the primary DNS domain name of the server computer is set to

be the same as the name of the Active Directory domain where the server was promoted to be a domain controller. To change this name, you must first demote the server so that it is no longer a domain controller. For a server changing its parent DNS domain name in this situation, the process of adding and deleting A and PTR RRs for the server is done automatically when joining and leaving an Active Directory domain. In this case, only manual updating of the name server (NS) RRs for the old and new DNS parent domains might be needed.

To change the DNS domain name for a DNS server computer not using Active Directory, such as a domain member server or a stand-alone server, some additional changes and manual administration might be necessary. For example, if the old and new DNS domain names are in two different zones, such as Zone A (the old zone) and Zone B (the new zone), you would need to make the following changes:

1. Change the **Primary DNS suffix for this computer** in System properties and restart the computer.

2. In Zone A, remove the host (A) RR for the server.

3. In Zone B, add the A RR for the server.

4. Update all start of authority (SOA) and name server (NS) RRs in zones containing the name of the renamed DNS server.

**Changing the IP address of the server**

To change the IP address for an existing DNS server, the process is the same as for any DNS host computer. Update the IP address in the server address (A) resource record. If the name of the server does not change, NS and SOA records are not affected.

A common configuration error that occurs when changing an IP address in the A resource record in a zone is that the same information used elsewhere in DNS is not updated. For example, if a DNS server that has a name or address change is used in a parent zone as part of a set of delegation records, the zone records might be updated but the delegation records left unchanged. For this reason, it is good to check the parent zone as well when changing or updating IP addresses for A resource records for DNS servers in a zone. Otherwise, zone delegation can fail.

Also, if you are using a multihomed server, be aware of the events and conditions under which a server running Windows Server 2003 can initiate a dynamic update when connection-specific DNS domain naming is used. For more information, see Configuring multiple names.

**Removing a DNS server from the network**

To remove a DNS server from the network, make the following changes in zones where the server is configured as an authoritative server for the zone:

1. Remove the address (A) resource record for the server.

2. Update the name server (NS) records in zones where the server is configured as authoritative to no longer include the server by name (as it appeared in the A record deleted in Step 1).

3. If the server is the primary server for a standard zone, revise the owner field of the SOA resource record for the zone to point to the new primary DNS server for the zone. (If the zone is directory-integrated, this step is not needed.)

4. Check the parent zone to ensure that any records (NS or A resource records) used for delegation to the zone are revised and no longer point to the removed server.

**Changing the primary server for a zone (standard primary zones only)**

For standard primary zones, it can sometimes be necessary to change the designated primary server for a zone. For example, suppose that the current primary server for a standard primary zone is Server A, and the new primary server for the zone is Server B.

To affect the changeover in status from Server A to Server B, make the following zone changes:

1. Add a new host (A) resource record (RR) for Server B.

2. Update the name server (NS) RR in the zone to remove Server A and include Server B as a configured server of authority that points to the new A RR added in Step 1.

3. Revise the name in the owner field of the start of authority (SOA) RR for the zone from Server A to Server B.

4. Remove the old A RR for Server A.

5. Check the parent zone to ensure that any delegation records (NS or A RRs) used are updated to refer to Server B.

# Using secondary servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using secondary servers**

DNS design specifications recommend that at least two DNS servers be used to host each zone. For standard primary-type zones, a secondary server is required to add and configure the zone to appear to other DNS servers in the network. For directory-integrated primary zones, secondary servers are supported but not required for this purpose. For example, two DNS servers running on domain controllers can be redundant primary servers for a zone, providing the same benefits of adding a secondary server while including additional advantages.

Secondary servers can provide a means to offload DNS query traffic in areas of the network where a zone is heavily queried and used. Additionally, if a primary server is down, a secondary server can provide some name resolution in the zone until the primary server is available.

If you add a secondary server, try to locate it as close as possible to clients that have a high demand for names used in the zone. Also, consider placing secondary servers across a router, either on other subnets (if you use a routed LAN) or across WAN links. This provides a good use of a secondary server as a local backup in cases where an intermediate network link becomes the point of failure between DNS servers and clients that use the zone.

Because a primary server always maintains the master copy of updates and changes to the zone, a secondary server relies on DNS zone transfer mechanisms to obtain its information and keep it current. Issues like zone transfer methods -- either using full or incremental zone transfers -- are more applicable when using secondary servers. For more information, see Understanding zones and zone transfer.

In considering the impact of zone transfers caused by secondary servers, consider their advantage as a backup source of information and measure this against the added cost they impose on your network infrastructure. A simple rule is that for each secondary server you add, network usage (because of added zone replication traffic) increases, and so does the time required to synchronize the zone at all secondary servers.

Secondary servers are most heavily used for forward lookup zones. If you are using reverse lookup zones, you can assume that you do not need to add as many secondary servers for those zones. Typically, a secondary server for a reverse lookup zone is not used outside of the network and subnet that correspond to the reverse zone.

**Note**

- For more information, see Add a secondary server for an existing zone.

# Securing the DNS Server service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing the DNS Server service**

To secure the DNS servers in your network, use the following guidelines:

- **Examine and configure the default DNS Server service settings that affect security.** The following DNS Server service configuration options have security implications for both the standard and the Active Directory-integrated DNS Server service.

| Default setting | Description |
|---|---|
| Interfaces | By default, a DNS Server service that is running on a multihomed computer is configured to listen for DNS queries using all of its IP addresses. Limit the IP addresses the DNS Server service listens on to the IP address used by its DNS clients as their preferred DNS server.<br><br>For more information, see Restrict a DNS server to listen only on selected addresses. |
| Secure cache against pollution | By default, the DNS Server service is secured from cache pollution, which results when DNS query responses contain nonauthoritative or malicious data. The **Secure cache against pollution** option prevents an attacker from successfully polluting the cache of a DNS server with resource records that were not requested by the DNS server. Changing this default setting will reduce the integrity of the responses provided by DNS Server service.<br><br>For more information, see Secure server cache against names pollution. |
| Disable recursion | By default, recursion is not disabled for the DNS Server service. This enables the DNS server to perform recursive queries on behalf of its DNS clients and DNS servers that have forwarded DNS client queries to it. Recursion can be used by attackers to deny the DNS Server service and therefore if a DNS server in your network is not intended to receive recursive queries, it should be disabled.<br><br>For more information, see Disable recursion on the DNS server. |
| Root hints | If you have an internal DNS root in your DNS infrastructure, configure the root hints of internal DNS servers to only point to the DNS servers hosting your root domain and not the DNS servers hosting the Internet root domain. This will prevent your internal DNS servers from sending private information over the Internet when resolving names.<br><br>For more information, see Update root hints on the DNS server and Updating root hints. |

For more information about default DNS Server service settings, see Restore server default preferences.

- **Manage the discretionary access control list (DACL) on DNS servers running on domain controllers.** In addition to the default DNS Server service settings that affect security described above, DNS servers configured as domain controllers use a DACL. The DACL allows you to control the permissions for the Active Directory users and groups that control the DNS Server service.

  The following table lists the default group or user names and permissions for the DNS Server service when it is running on a domain controller.

| Group or user names | Permissions |
|---|---|
| Administrators | Allow: Read, Write, Create All Child objects, Special Permissions |
| Authenticated Users | Allow: Read, Special Permissions |
| Creator Owner | Special Permissions |

| | |
|---|---|
| DnsAdmins | Allow: Read, Write, Create All Child objects, Delete Child objects, Special Permissions |
| Domain Admins | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects |
| Enterprise Admins | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects |
| Enterprise Domain Controllers | Allow: Special Permissions |
| Pre-Windows 2000 Compatible Access | Allow: Special Permissions |
| System | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects |

For more information about assigning permissions on objects, see Best practices for assigning permissions on Active Directory objects, Best practices for permissions and user rights, and Understanding Groups.

When the DNS Server service is running on a domain controller, its DACL may be managed using the Active Directory object MicrosoftDNS. Configuring the DACL on the MicrosoftDNS Active Directory object has the same effect as configuring the DACL on the DNS server in the DNS console, which is the recommended method. Consequently, the security administrators of Active Directory objects and DNS servers should be in direct contact to ensure that the administrators do not reverse each other's security settings.

For more information, see Assign, change, or remove permissions on Active Directory objects or attributes.

- **Always use an NTFS file system for DNS servers running a** Windows Server 2003 **operating system.** The NTFS file system is a more powerful file system than FAT and FAT32 and provides a variety of features including Active Directory, which is needed for domains, user accounts, and other important security features.

  For more information about NTFS, see NTFS and NTFS compared to FAT and FAT32.

For more information, see Security information for DNS.

# Using caching-only servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using caching-only servers**

Although all DNS servers cache queries that they have resolved, caching-only servers are DNS servers that only perform queries, cache the answers, and return the results. They are not authoritative for any domains and the information that they contain is limited to what has been cached while resolving queries.

In determining when to use this kind of server, note that when this server is initially started, it has no cached information. This information is obtained over time as client requests are serviced. However, if you are dealing with a slow-speed WAN link between sites, this option might be ideal because, once the cache is built, traffic decreases. In addition, the caching-only server does not perform zone transfers, which can also be network intensive in WAN environments.

**Note**

- For more information, see Install a caching-only DNS server.

# Modifying server defaults

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modifying server defaults**

The DNS console provides the ability to reconfigure some types of DNS server defaults for instances when defaults are not appropriate for use with a specific DNS server. Some configurable nondefault settings include:

- Selecting an alternate server configuration method to be used when initiating service.
- Disabling the use of recursion.

**Methods for initiating service**

The following methods are available to determine how the DNS Server service initializes and provides service when started.

| Method | Description |
|---|---|
| From registry | Initializes the DNS Server service by reading parameters stored in the registry. |
| From file | Initializes the DNS Server service using an optional boot file, as used by Berkeley Internet Name Domain (BIND) servers.<br><br>To use this file, supply a copy of a boot file from another DNS server running a DNS server that is based on the BIND implementation. On BIND-based DNS servers, this file is typically called the Named.boot file. The format of this file must be the older BIND 4 format, not the more recent BIND 8 boot file format.<br><br>If used, settings and directives in this file are used instead of comparable settings stored in the registry on the DNS server computer. For any parameters not configurable using boot file directives, registry defaults (or stored reconfigured server settings) are applied by the DNS Server service. |
| From Active Directory and registry | Initializes the DNS Server service by reading parameters stored in the Active Directory database and the server registry. This is the default boot option. |

**Notes**

- For more information on how to change the method used to initiate the DNS Server service, see Change the boot method used by the DNS server.
- For more information on Active Directory integration with the DNS Server service, see Active Directory integration.

**Disabling recursion**

By default, recursion is enabled for the DNS Server service, and clients typically request that the server use recursion to resolve a name when sending a query. If recursion is disabled, the DNS Server service always uses referral, regardless of the client request.

In general, DNS servers can answer queries for names outside of their authoritative zones in two ways:

- Servers can send referral answers, which are an immediate response to the requesting client with a list of resource records for other DNS servers it knows about that appear to be closer or more likely to be of help in resolving the queried name.
- Servers can use recursion to query other servers on behalf of the requesting client, attempting to fully resolve the name. Recursive lookups continue until the server receives an authoritative answer for the queried name. The server then forwards this answer in response to the original query from the requesting client.

In most cases, disabling recursion on a DNS server happens when DNS clients are to be limited to resolving names

authoritatively managed on a specific server. For example, this is the case when a DNS server has only DNS names data for an internal network or when the DNS server is incapable of resolving external DNS names (such as Internet DNS names) and clients are expected to retry another DNS server to resolve these names.

**Notes**

- For more information on how to enable or disable the use of recursion with the DNS Server service, see Disable recursion on the DNS server.

- If you disable recursion on the DNS server, you will not be able to use forwarders on the same server.

# Managing the aging and scavenging of server data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing the aging and scavenging of server data**

Aging and scavenging of stale resource records are features of DNS available when deploying your server with primary zones.

Where aging and scavenging are available for use, you can use the DNS console to perform the following related tasks for your DNS servers and any directory-integrated zones that they load:

1. Enable or disable the use of scavenging at a DNS server.
2. Enable or disable the use of scavenging for selected zones at the DNS server.
3. Modify the no-refresh interval, either as a server default or by specifying an overriding value at selected zones.
4. Modify the refresh interval, either as a server default or by specifying an overriding value at selected zones.
5. Specify whether periodic scavenging occurs automatically at the DNS server for any of its eligible zones, and how often these operations are repeated.
6. Manually initiate a single scavenging operation for all eligible zones at the DNS server.
7. View other related properties, such as the time stamp for individual resource records or the start scavenging time for a specified zone.

**Enabling scavenging of stale resource records**

By default, aging and scavenging features are disabled on all DNS servers and any of their zones. Before using these features, you should configure these settings for the applicable server and its directory-integrated zones:

- **Server aging/scavenging properties** for determining the use of these features on a server-wide basis.

  First, you need to enable aging and scavenging of stale resource records for use at the server. These settings are used to determine the affect of zone-level properties for any directory-integrated zones loaded at the server.

  For more information, see Set aging/scavenging properties for the DNS server.

- **Zone aging/scavenging properties** for determining the use of these features on a per-zone basis.

  When zone-specific properties are set for a selected zone, these settings apply only to the applicable zone and its resource records. Unless these zone-level properties are otherwise configured, they inherit their defaults from comparable settings maintained in server aging/scavenging properties.

  For more information, see Set aging/scavenging properties for a zone.

**Modifying no-refresh intervals**

When the no-refresh interval is in effect for a specific resource record, attempts to dynamically refresh its time stamp are suppressed by the DNS server. This aspect of the aging/scavenging mechanism prevents unnecessary refreshes from being processed by the server for aged resource records. These early refresh attempts, if not handled in this way, might otherwise increase Active Directory replication traffic related to processing DNS zone changes.

To ensure that records do not refresh prematurely, keep the **No-refresh interval** comparable in length to the current refresh interval for each resource record. For example, if you increase the **Refresh interval** to a higher value you can similarly increase the no-refresh interval.

In most instances, the default interval of 7 days is sufficient and does not need to be changed.

**Modifying refresh intervals**
When the refresh interval is in effect for a resource record, attempts to dynamically refresh its time stamp are accepted and processed by the DNS server. When setting this interval, it is important that the length of time used be greater than the maximum possible refresh period for any resource records contained within the zone. This period is equal to the maximum amount of time that it might take the record to be refreshed under normal network conditions, based on the specific source generating the record refresh.

For example, the following table shows default refresh periods for various services known to dynamically register and refresh records in DNS:

| Service | Default refresh period |
|---|---|
| Net logon | 24 hours |
| Clustering | 24 hours |
| DHCP client | 24 hours<br>**Note**<br>• The DHCP Client service sends dynamic updates for the DNS records. This includes both computers that obtain a leased IP address using DHCP and those that are configured statically for TCP/IP. |
| DHCP server | 4 days (half of the lease interval, which is 8 days by default)<br>**Note**<br>• Refresh attempts are only made by the DHCP servers configured to perform DNS dynamic updates on behalf of their clients (for example Windows 2000 Server DHCP server and Windows Server 2003 DHCP server). The period is based on the frequency in which DHCP clients renew their IP address leases with the server. Typically, this occurs when 50% of the scope lease time has elapsed. If the DNS default scope lease duration of 8 days is used, then the maximum refresh period for records updated by DHCP servers on behalf of clients is 4 days. |

By default, the refresh interval is 7 days. In most instances, this value is sufficient and does not need to be changed, unless any resource records in the zone are refreshed less often than every 7 days.

**Automated and manually-initiated scavenging**
Although scavenging start time and other factors determine when zones and records are actually eligible for scavenging, you have the option to perform scavenging using either of two methods to initiate it:

- **Automatic scavenging** specifies that aging and scavenging of stale records is to be performed automatically by the server for any eligible zones at a recurring interval specified as the *scavenging period*. When using automatic scavenging, the default scavenging period is 1 day and the minimum allowed value you can use is 1 hour.

  For more information, see Enable automatic scavenging of stale resource records.

- **Manual scavenging** specifies that aging and scavenging of stale records is to be performed as a non-recurring operation for any eligible zones at the server.

  For more information, see Start immediate scavenging of stale resource records.

**Modifying time stamp values**
For resource records that are not dynamically added to DNS zone data, a record time stamp value of zero is applied, preventing these records from aging or removal during scavenging.

You can, however, manually reset record properties to enable any statically entered records to qualify for the aging/scavenging process. If you do this, the record will be deleted based upon the modified time stamp value, at which point, you might need to recreate a record if it is still needed.

For more information, see Reset scavenging and aging properties for a specified resource record.

**Caution**

- Enabling aging and scavenging for use with standard primary zones modifies the format of zone files. This change does not affect zone replication to secondary servers, but the modified file cannot be loaded by other versions of DNS servers.

- Before using information in this topic, be sure that you have reviewed the aging/scavenging process and use of its configurable properties and features. For more information, see Understanding aging and scavenging.

**Note**

- For information on using Active Directory and DNS, see Active Directory integration.

# Using forwarders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using forwarders**

To use forwarders to manage the DNS traffic between your network and the Internet, configure the firewall used by your network to allow only one DNS server to communicate with the Internet. When you have configured the other DNS servers in your network to forward queries they cannot resolve locally to that DNS server it will act as your forwarder. For more information about forwarders, see Understanding forwarders.

**Forwarding sequence**

The order of the IP addresses listed as forwarders on a DNS server determines the sequence in which the IP addresses are used. After the DNS server forwards the query to the forwarder with the first IP address, it waits a short period for an answer from that forwarder (according to the DNS server's time out setting) before resuming the forwarding operation with the next IP address. It continues this process until it receives an affirmative answer from a forwarder.

For example, in the following figure the DNS servers with the first and second forwarder IP addresses do not respond to the DNS server. The DNS server with the third forwarder IP address responds and the query is forwarded to that DNS server.




Unlike conventional resolution, where a roundtrip time (RTT) is associated with each server, the IP addresses in the forwarders list are not ordered according to roundtrip time and must be reordered manually to change preference.

**Conditional forwarders**

Conditional forwarders are DNS servers that only forward queries for specific domain names. Instead of forwarding *all* queries it cannot resolve locally to a forwarder, a conditional forwarder is configured to forward a query to specific forwarders based on the domain name contained in the query. Forwarding according to domain names improves conventional forwarding by adding a name-based condition to the forwarding process.

The conditional forwarder setting for a DNS server consists of the following:

- The domain names for which the DNS server will forward queries.
- One or more DNS server IP addresses for each domain name specified.

When a DNS client or server performs a query operation against a DNS server, the DNS server looks to see if the query can be resolved using its own zone data or the data stored in its cache. If the DNS server is configured to forward for the domain name designated in the query, then the query is forwarded to the IP address of a forwarder associated with the domain name. For example, in the following figure, each of the queries for the domain names is forwarded to a DNS server associated with the domain name.




If the DNS server has no forwarder listed for the name designated in the query, it can attempt to resolve the query using standard recursion. For more information, see Configure a DNS server to use forwarders and How DNS query works.

Conditional forwarders allow you to improve name resolution between internal (private) DNS namespaces that are not part of the DNS namespace of the Internet, such as results from a company merger. By configuring the DNS servers in one internal namespace to forward all queries to the authoritative DNS servers in a second internal namespace, conditional forwarders enable name resolution between the two namespaces without performing recursion on the DNS namespace of the Internet. This enhancement to name resolution also avoids having your DNS servers perform recursion to your internal root for different namespaces within your network.

**Important**

- A DNS server cannot forward queries for the domain names in the zones it hosts. For example, the authoritative DNS server for the zone microsoft.com cannot forward queries according to the domain name microsoft.com. The DNS server authoritative for microsoft.com can forward queries for DNS names that end with example.microsoft.com, if example.microsoft.com is delegated to another DNS server.

**Conditional forwarder domain name length**

When a DNS server configured with a conditional forwarder receives a query for a domain name, it will compare that domain name with its list of domain name conditions and use the longest domain name condition that corresponds to the domain name in the query. For example, in the figure below, the DNS server performs the following conditional forwarding logic to determine how a query for a domain name will be forwarded:

1. The DNS server receives a query for networks.example.microsoft.com.

2. It compares that domain name with both microsoft.com and example.microsoft.com.

3. The DNS server determines that example.microsoft.com is the domain name that more closely matches the domain name query.

4. The DNS server forwards the query to the DNS server with the IP address 172.31.255.255, which is associated with example.microsoft.com.




**Forward-only servers**

When a DNS server configured to use forwarders cannot resolve a query locally, or using its forwarders, the server attempts to resolve the query using standard recursion. A DNS server can also be configured to not perform recursion after forwarders fail. In this configuration, the server does not attempt any further recursive queries to resolve the name. Instead, if it does not get a successful query response from any of the servers configured as forwarders, then it fails the query. A DNS server configured in this manner is called a forward-only DNS server. If all forwarders for a name in the query do not respond to a forward-only

DNS server, that DNS server will not attempt recursion.

A forward-only DNS server is different from a nonrecursive DNS server because it builds up a cache relating to the domain name and will use this cache to attempt to resolve the domain name. A nonrecursive DNS server will not build up a cache relating to the domain, nor will it perform recursion. In both configurations, the DNS servers will attempt to resolve the query using their authoritative data before using their forwarders.

**Note**

- You can disable recursion for the entire DNS server or on a per domain name basis. If you disable recursion on the entire DNS server, you will not be able to use forwarders on that DNS server.

For more information, see Directing queries through forwarders and Disable recursion on the DNS server.

# Directing queries through forwarders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Directing queries through forwarders**
A DNS server that acts as a forwarder may serve other purposes in the DNS environment. For example, it could be a primary or secondary server, or it could be a caching-only server. Consider the activities of all DNS servers in your network when deploying forwarders to ensure that forwarding complements the other DNS server activities and benefits name resolution. For general information about forwarders, see Understanding forwarders.

**Forwarders and delegation**
A DNS server configured with a forwarder and hosting a parent zone will use its delegation information before forwarding queries. If no delegation record exists for the DNS name in the query, then the DNS server will use its forwarders to resolve the query.

**Forwarders and root servers**
A common error when configuring forwarding is to attempt to configure forwarding on root servers for a private DNS namespace. Root servers are authoritative for the root zone of your organization's DNS namespace. The goal of attempting to configure forwarding on root servers for a private DNS namespace is to forward all offsite queries to Internet DNS servers. *Root servers cannot be configured with standard forwarding*. If a root server is queried about any domain name, then it will refer to a DNS server that can answer the question (from its local zones, cache or root hints), or it will respond with a failure (NXDOMAIN), but it cannot be configured to forward to specific servers.

### Note

- A root server can be configured with a conditional forwarder. Conditional forwarding can be used to forward queries between root servers in separate DNS namespaces. For steps to configure internal root servers, see article 294906 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=82900). For more information about configuring conditional forwarding, see article 304491 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=82901).

**Forwarder configuration tips**
Consider the following tips for efficient forwarder configuration and use:

- **Keep forwarder configuration uncomplicated.** For every DNS server configured with a forwarder, queries can be sent to a number of different places. Each forwarder and each conditional forwarder must be administered for the benefit of DNS client queries, and this process can be time consuming. Use forwarders strategically, where they are needed the most, such as resolving offsite queries or sharing information between namespaces.

- **Avoid chaining your forwarders.** If you have configured a DNS server named server-1 to forward queries for widgets.example.com to DNS server server-2, do not configure server-2 to forward queries for widgets.example.com to DNS server server-3. This is an inefficient resolution process and could result in errors if server-3 is accidentally configured to forward queries for widgets.microsoft.com to server-1, in which case you would have a chain.

- **Do not concentrate too great of a load on forwarders.** The recursive queries that forwarders send to the Internet can require a significant amount of time to answer due to the nature of the Internet. With a large number of internal DNS servers using these forwarders for Internet queries, they can experience a substantial concentration of network traffic. If network load will be an issue, use more than one forwarder and distribute the load between them.

- **Do not create inefficient resolution using forwarders.** The DNS server will attempt to forward domain names according to the order in which they are configured in the DNS console. A DNS server in Seattle may be incorrectly configured to forward a query to a server in London instead of another server in Seattle because the server in London is higher in the forwarders list. This will decrease the efficiency of name resolution on the network. Evaluate your network's forwarding configuration periodically to see if there are similar configurations.

For more information, see Configure a DNS server to use forwarders.

# Updating root hints

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Updating root hints**

Root hints are used to prepare servers authoritative for non-root zones so that they can learn and discover authoritative servers that manage domains located at a higher level or in other subtrees of the DNS domain namespace. These hints are essential for servers authoritative at lower levels of the namespace when locating and finding servers under these conditions.

For example, suppose a DNS server (Server A) has a zone called sub.example.microsoft.com. In the process of answering a query for a higher-level domain, such as the example.microsoft.com domain, Server A needs some assistance to locate an authoritative server (such as Server B) for this domain.

In order for Server A to find Server B, or any other servers that are authoritative for the microsoft.com domain, it needs to be able to query the root servers for the DNS namespace. The root servers can then refer Server A to the authoritative servers for the com domain. The servers for the com domain can, in turn, offer referral to Server B or other servers that are authoritative for the microsoft.com domain.

The root hints used by Server A must have helpful hints to the root servers for this process to locate Server B (or another authoritative server) as intended.

To configure and use root hints correctly, first determine how the following applies to your DNS servers:

- Are you using DNS on the Internet or on a private network?
- Is the server used as a root server?

By default, the DNS Server service implements root hints using a file, Cache.dns, stored in the systemroot\System32\Dns folder on the server computer. This file normally contains the NS and A resource records for the Internet root servers. If, however, you are using the DNS Server service on a private network, you can edit or replace this file with similar records that point to your own internal root DNS servers.

Another server configuration in which root hints are treated differently is one in which a DNS server is configured to be used by other DNS servers in an internal namespace as a forwarder for any DNS queries of names managed externally (the Internet, for example). Even though the DNS server used as a forwarder can be located internally on the same network as servers using it as a forwarder, it needs hints for the Internet root servers to work properly and resolve external names.

To update root hints, see Update root hints on the DNS server.

**Notes**

- If you are operating internal root servers, do not use root hints. Instead, delete the Cache.dns file entirely for any of your root servers.
- If a DNS server is configured to know about other DNS servers, such as through a list of DNS servers configured in its client TCP/IP properties for an installed network connection, the DNS Server service is capable of gathering its own root hints during new server configuration. You can use the Configure a DNS Server Wizard provided with the DNS console to accomplish this.

# Configuring multihomed servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring multihomed servers**
For multihomed DNS servers, you can configure the DNS Server service to selectively enable and bind only to IP addresses that you specify using the DNS console. By default, the DNS Server service binds to all IP interfaces configured for the computer.

This can include:

- Any additional IP addresses configured for a single network connection.
- Individual IP addresses configured for each separate connection where more than one network connection is installed on the server computer.

For multihomed DNS servers, you can restrict DNS service for selected IP addresses. When this feature is used, the DNS Server service only listens for and answers DNS requests that are sent to the IP addresses specified on the **Interface** tab in Server properties.

**When to specify interfaces**
By default, the DNS Server service listens on all IP addresses and accepts all client requests sent to its default service port (UDP 53 or TCP 53). Some DNS resolvers (including original version Windows 95 clients) require that the source address of a DNS response be the same as the destination address that was used in the query. If these addresses differ, clients can reject the response. To accommodate these resolvers, you can specify the list of allowed interfaces for the DNS server. When a list is set, the DNS Server service binds sockets only to allowed IP addresses used on the computer.

In addition to providing support for clients that require explicit bindings to be used, specifying interfaces can be useful for other reasons:

1. If, for administrative reasons, you do not want to use some IP addresses or interfaces on a multihomed server computer.
2. If the server computer is configured to use a large number of IP addresses and you do not want the added expense of binding to all of them.

**Additional considerations for multihoming DNS servers**
When configuring additional IP addresses and enabling them for use with a DNS server, consider the following:

- Additional system resources are consumed at the server computer.
- Although DNS provides the means to configure multiple IP addresses for use with any of your installed network adapters, there is no performance benefit for doing so.
- Even if the DNS server is handling multiple zones registered for Internet use, it is not necessary or required by the Internet registration process to have different IP addresses registered for each zone.

Given these considerations:

- Recognize that when adding IP addresses for DNS server use, each additional address might only slightly increase server performance. In instances when a large overall number of IP addresses are enabled for use, server performance can be degraded noticeably.
- In general, when adding network adapter hardware to the server computer, assign only a single primary IP address for each network connection.
- Whenever possible, remove non-essential IP addresses from existing server TCP/IP configurations.

**Note**

- For more information on how to specify interfaces with the DNS Server service, see Restrict a DNS server to listen only on selected addresses.

# Server administration using Dnscmd

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server administration using Dnscmd**

Dnscmd is provided as a command-line tool for managing DNS servers. This tool can be used to script batch files, to help automate the management and update of existing DNS server configurations, or to perform set up and configuration of DNS servers.

For information about installing this tool, see Install Windows Support Tools. For help in using the command, type **dnscmd /?** at a command prompt or use the Windows Support Tools help.

# Managing Clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing clients**
This section covers:

- Configuring DNS client settings
- Configuring multiple names
- Updating DNS client information
- Securing DNS clients

# Configuring DNS client settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring DNS client settings**

DNS configuration involves the following tasks when configuring TCP/IP properties for each computer:

- Setting a DNS computer or host name for each computer. For example, in the fully qualified domain name (FQDN) *wkstn1.example.microsoft.com.*, the DNS computer name is the leftmost label *wkstn1*.
- Setting a primary DNS suffix for the computer, which is placed after the computer or host name to form the FQDN. Using the previous example, the primary DNS suffix would be *example.microsoft.com*.
- Setting a list of DNS servers for clients to use when resolving DNS names, such as a preferred DNS server, and any alternate DNS servers to use if the preferred server is not available.
- Setting the DNS suffix search list or search method to be used by the client when it performs DNS query searches for short, unqualified domain names.

These tasks are discussed in more detail in each of the following sections.

**Setting computer names**

When setting computer names for DNS, it is useful to think of the name as the leftmost portion of a fully qualified domain name (FQDN). For example, in *wkstn1.example.microsoft.com.*, *wkstn1* is the computer name.

You can configure all Windows DNS clients with a computer name based on any of the standard supported characters defined in Request for Comments (RFC) 1123, "Requirements for Internet Hosts -- Application and Support." These characters include the use of:

- Uppercase letters, A through Z
- Lowercase letters, a through z
- Numbers, 0 through 9
- Hyphens (**-**)

If you are supporting both NetBIOS and DNS namespaces on your network, you can use a different computer name within each namespace. It is recommended that wherever possible, however, you try to use computer names that are 15 characters or less and that you follow the RFC 1123 naming requirements outlined above.

By default, the leftmost label in the FQDN for clients equals the NetBIOS computer name, unless this label is 16 or more characters, which is the maximum for NetBIOS names. When the computer name exceeds the maximum length for NetBIOS, the NetBIOS computer name is truncated based on the full label that is specified.

Before configuring computers with varying DNS and NetBIOS names, consider the following implications and their related issues for your deployment:

- If WINS lookup is enabled for zones hosted by your DNS servers, you need to use the same name for both NetBIOS and DNS computer naming. Otherwise, the results of clients attempting to query and resolve the names of these computers will be inconsistent.

If you have an investment in using NetBIOS names to support legacy Microsoft networking technology, it is recommended that you revise NetBIOS computer names used on your network to prepare for migration to a standard DNS-only environment. This prepares your network well for long-term growth and interoperability with future naming requirements. For example, if you use the same computer name for both NetBIOS and DNS resolution, consider converting any special characters such as the underscore (_) in your current NetBIOS names that do not comply with DNS naming standards. While these characters are

permitted in NetBIOS names, they are more often incompatible with traditional DNS host naming requirements and most existing DNS resolver client software.

**Notes**

- Although the use of the underscore (_) in DNS host names or in host (A) resource records has been traditionally prohibited by DNS standards, the use of underscores in service-related names--such as those used for service locator (SRV) resource records--has been proposed to avoid naming collisions in the Internet DNS namespace. For more information, see DNS RFCs.

- In addition to DNS standard naming conventions, Windows Server 2003 DNS supports the use of extended ASCII and Unicode characters. However, since most resolver software written for other platforms (such as UNIX) is based on the Internet DNS standards, this enhanced character support can be used only in private networks with computers running Windows 2000 or Windows Server 2003 DNS.

- The initial setup of DNS and TCP/IP displays a warning to suggest a standard DNS name if a nonstandard DNS name is entered.

- By default, computers and servers use DNS to resolve any name that is greater than 15 characters in length. If the name is less than or equal to 15 characters, then both NetBIOS and DNS name resolution can be attempted and used to resolve the name.

**Setting domain names**

The domain name is used with the client computer name to form the fully qualified domain name (FQDN), known also as the full computer name. In general, the DNS domain name is the remainder of the FQDN that is not used as the unique host name for the computer.

For example, the DNS domain name used for a client computer could be the following: If the FQDN, or **Full computer name**, is wkstn1.example.microsoft.com, the domain name is the example.microsoft.com portion of this name.

DNS domain names have two variations--a DNS name and a NetBIOS name. The full computer name (a fully qualified DNS name) is used during querying and location of named resources on your network. For earlier version clients, the NetBIOS name is used to locate various types of NetBIOS services that are shared on your network.

An example that shows the need for both NetBIOS and DNS names is the Net Logon service. In Windows Server 2003 DNS, the Net Logon service on a domain controller registers its service (SRV) resource records on a DNS server. For Windows NT Server 4.0 and earlier versions, domain controllers register a *DomainName* entry in WINS to perform the same registration and to advertise their availability for providing authentication service to the network.

When a client computer is started on the network, it uses the DNS resolver to query a DNS server for SRV records for its configured domain name. This query is used to locate domain controllers and provide logon authentication for accessing network resources. A client or a domain controller on the network optionally uses the NetBIOS resolver service to query WINS servers, attempting to locate *DomainName* [1C] entries to complete the logon process.

Your DNS domain names should follow the same standards and recommended practices that apply to DNS computer naming described in the previous section. In general, acceptable naming conventions for domain names include the use of letters A through Z, numerals 0 through 9, and the hyphen (**-**). The use of the period (**.**) in a domain name is always used to separate the discrete parts of a domain name, commonly known as *labels*. Each label corresponds to an additional level defined in the DNS namespace tree.

For most computers, the primary DNS suffix configured for the computer can be the same as its Active Directory domain name, although the two values can also be different.

**Important**

- By default, the primary DNS suffix portion of a computer's FQDN must be the same as the name of the Active Directory domain where the computer is located. To allow different primary DNS suffixes, a domain administrator may create a restricted list of allowed suffixes by creating the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is created and managed by the domain administrator using Active Directory Service Interfaces or the Lightweight Directory Access Protocol (LDAP).

  For more information, see Programming interfaces and Directory access protocol.

**Configuring a DNS servers list**

For DNS clients to operate effectively, a prioritized list of DNS name servers must be configured for each computer to use when processing queries and resolving DNS names. In most cases, the client computer contacts and uses its preferred DNS server, which is the first DNS server on its locally configured list. Listed alternate DNS servers are contacted and used when the preferred server is not available. For this reason, it is important that the preferred DNS server be appropriate for continuous client use under normal conditions.

### Notes

- For computers running Windows XP, the DNS server list is used by clients only to resolve DNS names. When clients send dynamic updates, such as when they change their DNS domain name or a configured IP address, they might contact these servers or other DNS servers as needed to update their DNS resource records. For more information, see Dynamic update.

- By default, the DNS client on Windows XP does not attempt dynamic update over a Remote Access Service (RAS) or virtual private network connection. To modify this configuration, you can modify the advanced TCP/IP settings of the particular network connection or modify the registry. For more information, see Configure TCP/IP to use DNS and Microsoft Windows Resource Kits Web site.

- By default, the DNS client does not attempt dynamic update of top-level domain (TLD) zones. Any zone named with a single-label name is considered a TLD zone, for example, com, edu, blank, my-company. To configure the DNS client to allow the dynamic update of TLD zones, you can use the **Update Top Level Domain Zones** policy setting or modify the registry.

- When DNS clients are configured dynamically using a DHCP server, it is possible to have a larger list of provided DNS servers. To provide an IP address list of DNS servers to your DHCP clients, enable option code 6 on the configured options types provided by your DHCP server. For Windows Server 2003 DHCP servers, you can configure a list of up to 25 DNS servers for each client with this option.

- To effectively share the load when multiple DNS servers are provided in a DHCP options-specified list, you can configure a separate DHCP scope that rotates the listed order of DNS and WINS servers provided to clients. For more information, see Configuring scopes.

**Configuring a DNS suffix search list**

For DNS clients, you can configure a DNS domain suffix search list that extends or revises their DNS search capabilities. By adding additional suffixes to the list, you can search for short, unqualified computer names in more than one specified DNS domain. Then, if a DNS query fails, the DNS Client service can use this list to append other name suffix endings to your original name and repeat DNS queries to the DNS server for these alternate FQDNs.

For computers and servers, the following default DNS search behavior is predetermined and used when completing and resolving short, unqualified names.

When the suffix search list is empty or unspecified, the primary DNS suffix of the computer is appended to short unqualified names, and a DNS query is used to resolve the resultant FQDN. If this query fails, the computer can try additional queries for alternate FQDNs by appending any connection-specific DNS suffix configured for network connections.

If no connection-specific suffixes are configured or queries for these resultant connection-specific FQDNs fail, then the client can then begin to retry queries based on systematic reduction of the primary suffix (also known as *devolution*).

For example, if the primary suffix were "example.microsoft.com", the devolution process would be able to retry queries for the short name by searching for it in the "microsoft.com" and "com" domains.

When the suffix search list is not empty and has at least one DNS suffix specified, attempts to qualify and resolve short DNS names is limited to searching only those FQDNs made possible by the specified suffix list. If queries for all FQDNs formed as a result of appending and trying each suffix in the list are not resolved, the query process fails, producing a "name not found" result.

### Notes

- If the domain suffix list is used, clients continue to send additional alternate queries based on different DNS domain

names when a query is not answered or resolved. Once a name is resolved using an entry in the suffix list, unused list entries are not tried. For this reason, it is most efficient to order the list with the most used domain suffixes first.

- Domain name suffix searches are used only when a DNS name entry is not fully qualified. To fully qualify a DNS name, a trailing period (**.**) is entered at the end of the name.

# Configuring multiple names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring multiple names**
Computers running Windows XP and servers running Windows Server 2003 are given DNS names by default. Each computer can have its DNS names configured using one of two possible methods:

- A primary DNS domain name, which applies as the default fully qualified DNS name for the computer and all of its configured network connections.
- A connection-specific DNS domain name, which can be configured as an alternate DNS domain name that applies only for a single network adapter installed and configured on the computer.

Although most computers do not need to support or use more than one name in DNS, support for configuring multiple connection-specific DNS names is sometimes useful. For example, by using multiple names, a user could specify which network connection to use when connecting to a multihomed computer.

**Example: Using connection-specific names**
As shown in the following figure, a multihomed server computer named "host-a" can be named according to both its primary and connection-specific DNS domain names.




In this example, the server computer host-a attaches to two separate subnets -- Subnet 1 and Subnet 2 -- which are also linked at redundant points using two routers for additional paths between each subnet. Given this configuration, host-a provides access as follows through its separately named local area network (LAN) connections:

- The name "host-a.public.example.microsoft.com" provides access using LAN connection 1 over Subnet 1, a lower-speed (10 megabit) Ethernet LAN, for normal access to users who have typical file and print service needs.
- The name "host-a.backup.example.microsoft.com" provides access using LAN connection 2 over Subnet 2, a higher-speed (100 megabit) Ethernet LAN, for reserved access by server applications and administrators who have special needs, such as troubleshooting server networking problems, performing network-based backup, or replicating zone data between servers.

In addition to the connection-specific DNS names, the computer can also be accessible using either of the two LAN connections by specifying its primary DNS domain name, "host-a.example.microsoft.com".

When configured as shown, a computer can register resource records in DNS according to its three distinct names and sets of IP addresses, as shown in the following table:

| DNS name | IP addresses | Description |
|---|---|---|

| | | |
|---|---|---|
| host-a.example.microsoft.com | 10.1.1.11, 10.2.2.22 | Primary DNS name for computer. The computer registers A and PTR resource records for all configured IP addresses under this name in the "example.microsoft.com" zone. |
| host-a.public.example.microsoft.com | 10.1.1.11 | Connection-specific DNS name for LAN connection 1, which registers A and PTR resource records for IP address 10.1.1.11 in the "public.example.microsoft.com" zone. |
| host-a.backup.example.microsoft.com | 10.2.2.22 | Connection-specific DNS name for LAN connection 2, which registers A and PTR resource records for IP address 10.2.2.22 in the "backup.example.microsoft.com" zone. |

**Important**

- By default, the primary DNS suffix portion of a computer's fully qualified domain name (FQDN) must be the same as the name of the Active Directory domain where the computer is located. To allow different primary DNS suffixes, a domain administrator may create a restricted list of allowed suffixes by creating the **msDS-AllowedDNSSuffixes** attribute in the domain object container. This attribute is created and managed by the domain administrator using Active Directory Service Interfaces or the Lightweight Directory Access Protocol (LDAP).

  For more information, see Programming interfaces and Directory access protocol.

**Notes**

- DNS names can be set using remote administration and other remote configuration services, such as DHCP. For a DNS server running Windows Server 2003, the primary DNS domain name can be set using either remote administration or the unattended setup option.

- For connection-specific naming, you can use TCP/IP configuration methods. You can manually configure the DNS domain name for each connection that appears in the Network Connections folder or use a DHCP option type (option code 15). For more information, see Advanced configuration.

- For more information about DHCP options, see "DHCP Options" at the Microsoft Windows Resource Kits Web site.

# Updating DNS client information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Updating DNS client information**

Updating DNS client information on a computer is sometimes necessary, for example, when a computer is added or removed from an Active Directory domain. In most cases, clients perform dynamic updates for the following changes when entering or leaving a domain:

- When a computer is joined to a domain, it adds its A and PTR resource records to DNS.
- When a computer is removed from a domain, it deletes its A and PTR resource records from DNS.

If you have enabled dynamic update for zones, these changes occur automatically. For more information, see Dynamic update.

# Securing DNS clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing DNS clients**

The following DNS client considerations have security implications for DNS clients in a DNS infrastructure:

- **Whenever possible, specify static IP addresses for the preferred and alternate DNS servers used by a DNS client.** If a DNS client is configured to obtain its DNS server addresses automatically, it will obtain them from a DHCP server. While this method of obtaining DNS server addresses is secure, it is only as secure as the DHCP server. By configuring DNS clients with static IP addresses for the preferred and alternate DNS servers, you eliminate one possible avenue of attack.

  For more information, see Configure TCP/IP to use DNS; Enable DNS for DHCP-enabled clients.

- **Control which DNS clients have access to the DNS server.** If a DNS server is configured to listen only on specific IP addresses, then only DNS clients configured to use these IP addresses as preferred and alternate DNS servers will contact the DNS server.

  For more information, see Restrict a DNS server to listen only on selected addresses.

For more information, see Security information for DNS.

# Managing Zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing zones**

This section covers:

- Managing general properties
- Managing authority records
- Managing resource records
- Delegating zones
- Configuring notify lists
- Securing DNS zones
- Securing DNS resource records
- Using stub zones
- Active Directory-integrated stub zones
- Contrasting stub zones and conditional forwarders
- Checking names and zone data
- Using WINS lookup
- Managing reverse lookup zones

# Managing general properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing general properties**

When you add a zone using the DNS console, you can manage these general properties for it:

- Pause or start the zone to interrupt or restore service for it.
- Change or convert the type for the zone.
- Disable or enable dynamic updates for the zone.

For Active Directory-integrated zones, you can enable the use of secure dynamic updates. This enables you to restrict updates to only a specific set of authorized users or systems. When a secure update policy is enabled for the zone, only users, systems, or groups authorized through Active Directory and included in the access control list (ACL) for each directory-integrated zone, are permitted to update the zone or specific resource records used in it.

In addition to these general zone properties, you can configure or manage the following zone properties using the DNS console:

- **Start of authority (SOA) properties** These include properties supported by the SOA resource record, which is used to initialize the zone and indicate zone authority for a DNS domain name (and any of its subdomains not delegated away to other servers) to others in the DNS namespace. This record affects how often the zone must be renewed and transferred by other servers that load the zone and how long clients can cache resource records (RRs) when returned in answered queries for names in the zone. For more information on this record and configuring it, see Managing authority records.

- **Name server (NS) properties** These include all fields supported by the NS resource record (RR) for the zone. The NS RR is used for designating the names of DNS servers authoritative for the zone to others. For more information on this record and configuring it, see Managing authority records.

- **Zone Transfer properties** With this feature, you can configure how the zone permits transfers to be performed.

  You can choose to deny all requests to the server for transfer of this zone, to allow them only for other DNS servers configured on the **Name Servers** tab, or to transfer the zone only to DNS servers that you specify by IP address in a configured list.

  Using **Notify** located here, you can also enable and configure DNS notification for secondary servers of the zone. When notification is used, other servers (either those configured on the **Name Servers** tab or on a list that you specify) are notified of zone changes. These servers can then pull changes by initiating a zone transfer to update the zone.

  For more information, see Configuring notify lists.

  **Note**

  - By default, the DNS server will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.

- **WINS lookup properties** The Windows Internet Name Service (WINS) lookup feature can be used to provide an expanded DNS name resolution path for zones when a queried name is not found in the zone. If WINS lookup is enabled for the zone, a WINS server (or list of WINS servers) can be contacted to assist in resolving a host name in the WINS-managed NetBIOS namespace. This feature is supported only by Microsoft DNS servers. For more information, see Using WINS lookup.

# Managing authority records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing authority records**

Zones are based on a concept of server authority. When a DNS server is configured to load a zone, it uses two types of resource records to determine the authoritative properties of the zone:

- First, the start of authority (SOA) resource record indicates the name of origin for the zone and contains the name of the server that is the primary source for information about the zone. It also indicates other basic properties of the zone.

- Next, the name server (NS) resource record is used to notate which DNS servers are designated as authoritative for the zone. By listing a server in the NS RR, it becomes known to others as an authoritative server for the zone. This means that any server specified in the NS RR is to be considered an authoritative source by others, and is able to answer with certainty any queries made for names included in the zone.

The SOA and NS resource records occupy a special role in zone configuration. They are required records for any zone and are typically the first resource records listed in files. By default, the New Zone Wizard automatically creates these records when a new primary zone is added using the DNS console.

**The SOA resource record**

The start of authority (SOA) resource record is always first in any standard zone. It indicates the DNS server that either originally created it or is now the primary server for the zone. It is also used to store other properties such as version information and timings that affect zone renewal or expiration. These properties affect how often transfers of the zone are done between servers authoritative for the zone.

The SOA resource record contains the following information:

| Field | Description |
|---|---|
| Primary server (owner) | The host name for the primary DNS server for the zone. |
| Responsible person | The e-mail address of the person responsible for administering the zone. A period (**.**) is used instead of an at sign (**@**) in this e-mail name. |
| Serial number | The revision number of the zone file. This number increases each time a resource record in the zone changes. It is important that this value increase each time the zone is changed, so that either partial zone changes or the fully revised zone can be replicated to other secondary servers during subsequent transfers. |
| Refresh interval | The time, in seconds, that a secondary DNS server waits before querying its source for the zone to attempt renewal of the zone. When the refresh interval expires, the secondary DNS server requests a copy of the current SOA record for the zone from its source, which answers this request. The secondary DNS server then compares the serial number of the source server's current SOA record (as indicated in the response) with the serial number in its own local SOA record. If they are different, the secondary DNS server requests a zone transfer from the primary DNS server. The default for this field is 900 seconds (15 minutes). |
| Retry interval | The time, in seconds, a secondary server waits before retrying a failed zone transfer. Normally, this time is less than the refresh interval. The default value is 600 seconds (10 minutes). |
| Expire interval | The time, in seconds, before a secondary server stops responding to queries after a lapsed refresh interval where the zone was not refreshed or updated. Expiration occurs because at this point in time, the secondary server must consider its local data unreliable. The default value is 86,400 seconds (24 hours). |
| Minimum (default) TTL | The default Time-To-Live (TTL) of the zone and the maximum interval for caching negative answers to name queries. The default value is 3,600 seconds (1 hour). |

The following is an example of a default SOA resource record:

```
@   IN  SOA     nameserver.example.microsoft.com.  postmaster.example.microsoft.com. (
                                1            ; serial number
                                3600         ; refresh   [1h]
                                600          ; retry     [10m]
                                86400        ; expire    [1d]
                                3600 )       ; min TTL   [1h]
```

In the example SOA record shown above, the primary or originating server for the zone is shown as nameserver.example.microsoft.com. The e-mail address for the person to contact regarding questions about this zone is postmaster.example.microsoft.com.

**Notes**

- Periods are used to represent e-mail addresses when writing and storing DNS domain names in a zone. In an e-mail application, the previous example address would instead likely appear as postmaster@example.microsoft.com.
- The parentheses used in the SOA resource record as it appears in a zone file are used to enable wrapping of the record over multiple lines of text.
- If an individual TTL value is assigned and applied to a specified resource record used in the zone, it overrides the minimum (default) TTL set in the SOA record.

**The NS resource record**

Name server (NS) resource records can be used to assign authority to specified servers for a DNS domain name in two ways:

- By establishing a list of authoritative servers for the domain so that those servers can be made known to others that request information about this domain (zone).
- By indicating authoritative DNS servers for any subdomains that are delegated away from the zone.

In the case of assigning servers with host names in the same zone, corresponding address (A) resource records are normally used in the zone to resolve the names of specified servers to their IP addresses. For servers that are specified using this RR as part of a zone delegation to a subdomain, the NS resource record usually contains out-of-zone names. For the out-of-zone names to be resolved, A resource records for the specified out-of-zone servers might be needed. When these out-of-zone NS and A records are needed to provide delegation, they are known as *glue records*.

| |
|---|
| **Description:** Used to map a DNS domain name as specified in **owner** to the name of hosts operating DNS servers specified in the *name_server_domain_name* field. |
| **Syntax:** *owner ttl* **IN NS** *name_server_domain_name* |
| **Example:**<br>`example.microsoft.com.  IN NS  nameserver1.example.microsoft.com` |

# Managing resource records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing resource records**

After you create a zone, additional resource records need to be added to it. The most common resource records (RRs) to be added are:

- **Host (A)** For mapping a DNS domain name to an IP address used by a computer.
- **Alias (CNAME)** For mapping an alias DNS domain name to another primary or canonical name.
- **Mail Exchanger (MX)** For mapping a DNS domain name to the name of a computer that exchanges or forwards mail.
- **Pointer (PTR)** For mapping a reverse DNS domain name based on the IP address of a computer that points to the forward DNS domain name of that computer.
- **Service location (SRV)** For mapping a DNS domain name to a specified list of DNS host computers that offer a specific type of service, such as Active Directory domain controllers.
- Other resource records as needed.

**Host (A) resource records**

Host (A) resource records are used in a zone to associate DNS domain names of computers (or hosts) to their IP addresses, and can be added to a zone in several ways:

- You can manually create an A resource record for a static TCP/IP client computer using the DNS console.
- Windows clients and servers use the DHCP Client service to dynamically register and update their own A resource records in DNS when an IP configuration change occurs.
- DHCP-enabled client computers running earlier versions of Microsoft operating systems can have their A resource records registered and updated by proxy if they obtain their IP lease from a qualified DHCP server (only the Windows 2000 and Windows Server 2003 DHCP Server service currently supports this feature).

The host (A) resource record is not required for all computers, but is needed by computers that share resources on a network. Any computer that shares resources and needs to be identified by its DNS domain name, needs to use A resource records to provide DNS name resolution to the IP address for the computer.

Most A RRs that are required in a zone can include other workstations or servers that share resources, other DNS servers, mail servers, and Web servers. These resource records comprise the majority of resource records in a zone database.

For more information, see Resource records reference.

**Alias (CNAME) resource records**

Alias (CNAME) resource records are also sometimes called *canonical names*. These records allow you to use more than one name to point to a single host, making it easy to do such things as host both an FTP server and a Web server on the same computer. For example, the well-known server names (ftp, www) are registered using CNAME RRs that map to the DNS host name, such as "server-1", for the server computer that hosts these services.

CNAME RRs are recommended for use in the following scenarios:

- When a host specified in an A RR in the same zone needs to be renamed.
- When a generic name for a well-known server such as www needs to resolve to a group of individual computers (each with individual A RRs) that provide the same service. For example, a group of redundant Web servers.

When renaming a computer with an existing A RR in the zone, you can use a CNAME RR temporarily, to allow a grace period for users and programs to switch from specifying the old computer name to using the new one. To do this, you need the following:

- For the new DNS domain name of the computer, a new A RR is added to the zone.
- For the old DNS domain name, a CNAME RR is added that points to the new A RR.
- The original A RR for the old DNS domain name (and its associated PTR RR if applicable) is removed from the zone.

When using a CNAME RR for aliasing or renaming a computer, set a temporary limit on how long the record is used in the zone before removing it from DNS. If you forget to delete the CNAME RR and later its associated A RR is deleted, the CNAME RR can waste server resources by trying to resolve queries for a name no longer used on the network.

The most common or popular use of a CNAME RR is to provide a permanent DNS aliased domain name for generic name resolution of a service-based name, such as www.example.microsoft.com to more than one computer or one IP address used in a Web server. For example, the following shows the basic syntax of how a CNAME RR is used.

*alias_name* **IN CNAME** *primary_canonical_name*

In this example, a computer named host-a.example.microsoft.com needs to function as both a Web server named "www.example.microsoft.com." and an FTP server named "ftp.example.microsoft.com." To achieve the intended use for naming this computer, you can add and use the following CNAME entries in the example.microsoft.com zone:

```
host-a     IN  A      10.0.0.20
ftp        IN  CNAME  host-a
www        IN  CNAME  host-a
```

If you later decide to move the FTP server to another computer, separate from the Web server on "host-a", simply change the CNAME RR in the zone for ftp.example.microsoft.com and add an additional A RR to the zone for the new computer hosting the FTP server.

Based on the earlier example, if the new computer were named "host-b.example.microsoft.com", the new and revised A and CNAME RRs would be as follows:

```
host-a     IN  A      10.0.0.20
host-b     IN  A      10.0.0.21
ftp        IN  CNAME  host-b
www        IN  CNAME  host-a
```

For more information, see Resource records reference.

**Mail exchanger (MX) resource records**

The mail exchanger (MX) RR is used by e-mail applications to locate a mail server based on a DNS domain name used in the destination address for the e-mail recipient of a message. For example, a DNS query for the name "example.microsoft.com" could be used to find an MX RR, enabling an e-mail application to forward or exchange mail to a user with the e-mail address "user@microsoft.com."

The MX RR shows the DNS domain name for the computer or computers that process mail for a domain. If multiple MX RRs exist, the DNS Client service attempts to contact mail servers in the order of preference from lowest value (highest priority) to highest value (lowest priority). The following shows the basic syntax for use of an MX RR.

*mail_domain_name* **IN MX** *preference mailserver_host*

By using the MX RRs shown below in the example.microsoft.com zone, mail addressed to user@example.microsoft.com is delivered to user@mailserver0.example.microsoft.com first if possible. If this server is unavailable, the resolver client can then use user@mailserver1.example.microsoft.com instead.

```
@          IN  MX   1    mailserver0
@          IN  MX   2    mailserver1
```

Note that the use of the at sign (**@**) in the records indicates that the mailer DNS domain name is the same as the name of origin (example.microsoft.com) for the zone.

For more information, see Resource records reference.

**Pointer (PTR) resource records**
Pointer (PTR) RRs are used to support the reverse lookup process, based on zones created and rooted in the in-addr.arpa domain. These records are used to locate a computer by its IP address and resolve this information to the DNS domain name for that computer.

PTR RRs can be added to a zone in several ways:

- You can manually create a PTR RR for a static TCP/IP client computer using the DNS , either as a separate procedure or as part of the procedure for creating an A RR.

- Computers use the DHCP Client service to dynamically register and update their PTR RR in DNS when an IP configuration change occurs.

- All other DHCP-enabled client computers can have their PTR RRs registered and updated by the DHCP server if they obtain their IP lease from a qualified server. The Windows 2000 and Windows Server 2003 DHCP Server service provides this capability.

The pointer (PTR) resource record is used only in reverse lookup zones to support reverse lookup. For more information, see Resource records reference.

**Service location (SRV) resource records**
To locate Active Directory domain controllers, service location (SRV) RRs are required. Typically, you can avoid manual administration of the SRV RR when installing Active Directory.

By default, the Active Directory installation wizard attempts to locate a DNS server based on the list of preferred or alternate DNS servers, configured in any of its TCP/IP client properties, for any of its active network connections. If a DNS server that can accept dynamic update of the SRV RR (and other RRs related to registering Active Directory as a service in DNS) is contacted, the configuration process is complete.

If, during the installation, a DNS server that can accept updates for the DNS domain name used to name your Active Directory is not found, the wizard can install a DNS server locally and automatically configure it with a zone to support the Active Directory domain.

For example, if the Active Directory domain that you chose for your first domain in the forest was example.microsoft.com, a zone rooted at the DNS domain name of example.microsoft.com would be added and configured to use with the DNS server running on the new domain controller.

Whether or not you install the DNS Server service locally, a file (Netlogon.dns) is written and created during the Active Directory installation process that contains the SRV RRs and other RRs needed to support the use of Active Directory. This file is created in the systemroot\System32\Config folder.

If you are using a DNS server that fits one of the following descriptions, you should use the records in Netlogon.dns to manually configure the primary zone on that server to support Active Directory.

1. The computer operating your DNS server is running on another platform, such as UNIX, and cannot accept or recognize dynamic updates.

2. A DNS server at this computer that is not the DNS Server service provided with the Windows Server 2003 family is authoritative for the primary zone corresponding to the DNS domain name for your Active Directory domain.

3. The DNS server supports the SRV RR, as defined in the Internet draft, "A DNS RR specifying the location of services (DNS SRV)", but does not support dynamic updates.

   For example, the DNS Server service provided with Windows NT Server 4.0, when updated to Service Pack 4 or later, fits this description.

In the future, the SRV RR might also be used to register and lookup other well-known TCP/IP services on your network if

applications implement and support DNS name queries that specify this record type. For more information, see Resource records reference.

**Other resource records**

Other additional resource records are supported by Windows Server 2003 DNS and are used less frequently in most zones. These additional types of resource records can be added as needed using the DNS console. For more information about supported resource records, see Resource records reference.

# Delegating zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Delegating zones**

DNS provides the option of dividing up the namespace into one or more zones, which can then be stored, distributed, and replicated to other DNS servers. When deciding whether to divide your DNS namespace to make additional zones, consider the following reasons to use additional zones:

- A need to delegate management of part of your DNS namespace to another location or department within your organization.
- A need to divide one large zone into smaller zones for distributing traffic loads among multiple servers, improve DNS name resolution performance, or create a more fault-tolerant DNS environment.
- A need to extend the namespace by adding numerous subdomains at once, such as to accommodate the opening of a new branch or site.

If, for any of these reasons, you could benefit from delegating zones, it might make sense to restructure your namespace by adding additional zones. When choosing how to structure zones, you should use a plan that reflects the structure of your organization.

When delegating zones within your namespace, be aware that for each new zone you create, you will need delegation records in other zones that point to the authoritative DNS servers for the new zone. This is necessary both to transfer authority and to provide correct referral to other DNS servers and clients of the new servers being made authoritative for the new zone.

When a standard primary zone is first created, it is stored as a text file containing all resource record information on a single DNS server. This server acts as the primary master for the zone. Zone information can be replicated to other DNS servers to improve fault tolerance and server performance.

When structuring your zones, there are several good reasons to use additional DNS servers for zone replication:

1. Added DNS servers provide zone redundancy, enabling DNS names in the zone to be resolved for clients if a primary server for the zone stops responding.
2. Added DNS servers can be placed so as to reduce DNS network traffic. For example, adding a DNS server to the opposing side of a low-speed WAN link can be useful in managing and reducing network traffic.
3. Additional secondary servers can be used to reduce loads on a primary server for a zone.

**Example: Delegating a subdomain to a new zone**

As shown in the following figure, when a new zone for a subdomain (example.microsoft.com) is created, delegation from the parent zone (microsoft.com) is needed.




In this example, an authoritative DNS server computer for the newly delegated example.microsoft.com subdomain is named based on a derivative subdomain included in the new zone (ns1.us.example.microsoft.com). To make this server known to others outside of the new delegated zone, two RRs are needed in the microsoft.com zone to complete delegation to the new zone.

These RRs include:

- An NS RR to effect the delegation. This RR is used to advertise that the server named ns1.us.example.microsoft.com is an authoritative server for the delegated subdomain.

- An A RR (also known as a glue record) is needed to resolve the name of the server specified in the NS RR to its IP address. The process of resolving the host name in this RR to the delegated DNS server in the NS RR is sometimes referred to as glue chasing.

**Note**

- When zone delegations are correctly configured, normal zone referral behavior can sometimes be circumvented if you are using forwarders in your DNS server configuration. For more information, see Directing queries through forwarders.

# Configuring notify lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring notify lists**
Notify lists can be created and managed to provide a method for a master server for a zone (such as the primary server or another secondary server) to notify a list of secondary servers when the zone changes.

To create or manage a notify list, you can use the DNS console to add and remove servers by their IP addresses in the notify list for a zone. By default, the notify list is based on the authoritative servers for the zone as listed on the **Name Servers** tab.

**Notes**

- For more information on how DNS notification works, see Understanding zones and zone transfer.
- For more information on configuring and managing notify lists, see Create and manage a notify list for a zone.
- You cannot configure a notify list for a stub zone.
- By default, a DNS server running a Windows Server 2003 operating system will only allow a zone transfer to authoritative DNS servers listed in the name server (NS) resource records for the zone.

# Securing DNS zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing DNS zones**

The following DNS zone configuration options have security implications for both standard and Active Directory-integrated DNS zones:

- **Configure secure dynamic updates.** By default, the **Dynamic updates** setting is not configured to allow dynamic updates. This is the most secure setting as it prevents an attacker from updating DNS zones, but this setting prevents you from taking advantage of the benefits to administration that dynamic update provides. To have computers securely update DNS data, store DNS zones in Active Directory and use the secure dynamic update feature. Secure dynamic update restricts DNS zone updates to only those computers that are authenticated and joined to the Active Directory domain where the DNS sever is located, and to the specific security settings defined in the ACLs for the DNS zone.

  For more information, see Allow only secure dynamic updates.

- **Manage the discretionary access control list (DACL) on the DNS zones stored in Active Directory.** The DACL allows you to control the permissions for the Active Directory users and groups that may control the DNS zones.

  The following table lists the default group or user names and permissions for DNS zones stored in Active Directory.

| Group or user names | Permissions |
|---|---|
| Administrators | Allow: Read, Write, Create All Child objects, Special Permissions |
| Authenticated Users | Allow: Create All Child objects |
| Creator Owner | Special Permissions |
| DnsAdmins | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects, Special Permissions |
| Domain Admins | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects |
| Enterprise Admins | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects |
| Enterprise Domain Controllers | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects, Special Permissions |
| Everyone | Allow: Read, Special Permissions |
| Pre-Windows 2000 Compatible Access | Allow: Special Permissions |
| System | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects |

  For more information, see Modify security for a directory-integrated zone.

  The DNS Server service running on a domain controller that has zones stored in Active Directory stores its zone data in Active Directory using Active Directory objects and attributes. Configuring the DACL on the DNS Active Directory objects has the same effect as configuring the DACL on DNS zones in the DNS console. Consequently, the security administrators of Active Directory objects and DNS data should be in direct contact to ensure that the administrators do not reverse each other's security settings.

  The Active Directory objects and attributes used by DNS zone data stored in Active Directory are described in the following table.

| Object | Description |
| --- | --- |
| DnsZone | Container created when a zone is stored in Active Directory. |
| DnsNode | Leaf object used to map and associate a name in the zone to resource data. |
| DnsRecord | Multivalued attribute of a dnsNode object used to store the resource records associated with the named node object. |
| DnsProperty | Multivalued attribute of a dnsZone object used to store zone configuration information. |

For more information, see Assign, change, or remove permissions on Active Directory objects or attributes.

- **Restrict zone transfers.** By default, the DNS Server service only allows zone information to be transferred to servers listed in the name server (NS) resource records of a zone. This is a secure configuration, but for increased security, this setting should be changed to the option to allow zone transfers to specified IP addresses. Changing this setting to allow zone transfers to any server may expose your DNS data to an attacker attempting to footprint your network.

  For more information, see Modify zone transfer settings.

- **Understand the compromise involved in zone delegation.** When deciding whether to delegate DNS domain names to zones hosted on DNS servers that are administered separately, it is important to consider the security implications of giving multiple individuals the ability to administer the DNS data for your network. DNS zone delegation involves a compromise between the security benefits of having a single authoritative DNS server for all DNS data and the administrative benefits of distributing responsibility for your DNS namespace to separate administrators. This issue is very important when delegating the top-level domains of a private DNS namespace, as those domains contain very sensitive DNS data.

  For more information, see Delegating zones.

**DNS zone data recovery**

If your DNS data has been corrupted you can restore your DNS zone file from the backup folder located in the systemroot/DNS/Backup folder. When a zone is first created, a copy of the zone is added to the backup folder. To recover the zone, copy the original zone file from the backup folder into the systemroot/DNS folder. When you use the New Zone Wizard to create the zone, specify the zone file in the systemroot/DNS folder as the zone file for the new zone. For more information, see Add a forward lookup zone.

**Note**

- This operation only applies to standard zones that are not stored in Active Directory.

In the case of both standard and Active Directory-integrated zones, you should use the system backup feature to establish a standard data recovery procedure for your DNS infrastructure. For more information, see Create an Automated System Recovery set using Backup.

For more information, see Security information for DNS.

# Securing DNS resource records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing DNS resource records**

The following DNS resource record configuration options have security implications for resource records stored in both standard and Active Directory-integrated DNS zones:

- **Manage the discretionary access control list (DACL) on DNS resource records stored in Active Directory.** The DACL allows you to control the permissions for the Active Directory users and groups that may control the DNS resource records. For more information, see Modify security for a resource record.

  The following table lists the default group or user names and permissions for DNS resource records stored in Active Directory.

| Group or user names | Permissions |
|---|---|
| Administrators | Allow: Read, Write, Create All Child objects, Special Permissions |
| Authenticated Users | Allow: Create All Child objects |
| Creator Owner | Special Permissions |
| DnsAdmins | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects, Special Permissions |
| Domain Admins | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects |
| Enterprise Admins | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects |
| Enterprise Domain Controllers | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects, Special Permissions |
| Everyone | Allow: Read, Special Permissions |
| Pre-Windows 2000 Compatible Access | Allow: Special Permissions |
| System | Allow: Full Control, Read, Write, Create All Child objects, Delete Child objects |

For more information, see Security information for DNS.

# Using stub zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using stub zones**

Use stub zones to:

- **Keep delegated zone information current.** By updating a stub zone for one of its child zones regularly, the DNS server hosting both the parent zone and the stub zone will maintain a current list of authoritative DNS servers for the child zone.
- **Improve name resolution.** Stub zones enable a DNS server to perform recursion using the stub zone's list of name servers without needing to query the Internet or internal root server for the DNS namespace.
- **Simplify DNS administration.** By using stub zones throughout your DNS infrastructure, you can distribute a list of the authoritative DNS servers for a zone without using secondary zones. However, stub zones do not serve the same purpose as secondary zones and are not an alternative when considering redundancy and load sharing.

There are two lists of DNS servers involved in the loading and maintenance of a stub zone:

- The list of master servers from which the DNS server loads and updates a stub zone. A master server may be a primary or secondary DNS server for the zone. In both cases, it will have a complete list of the DNS servers for the zone.
- The list of the authoritative DNS servers for a zone. This list is contained in the stub zone using name server (NS) resource records.

When a DNS server loads a stub zone, such as widgets.example.com, it queries the master servers, which can be in different locations, for the necessary resource records of the authoritative servers for the zone widgets.example.com. The list of master servers may contain a single server or multiple servers and can be changed anytime. For more information, see Configure a stub zone for local master servers.

**Stub zone updates**

Stub zone updates involve the following conditions:

- When a DNS server loads a stub zone, it queries the zone's master server for the SOA resource record, NS resource records at the zone's root, and glue A resource records.
- During updates to the stub zone, the master server is queried by the DNS server hosting the stub zone for the same resource record types requested during the loading of the stub zone.
- The Refresh interval of the SOA resource record determines when the DNS server hosting the stub zone will attempt a zone transfer (update).
- If an update fails, the Retry interval of the SOA resource record determines when the update is retried.
- Once the Retry interval has expired without a successful update, the expiration time as specified in the Expires field of the SOA resource record determines when the DNS server stops using the stub zone data.

Use the DNS console in Microsoft Management Console (MMC) to perform the following stub zone update operations:

- **Reload.** Reload the stub zone from the local storage of the DNS server hosting the stub zone.
- **Transfer from master.** Have the DNS server hosting the stub zone determine if the serial number in the stub zone's SOA resource record has expired, and then perform a zone transfer from the stub zone's master server.

- **Reload from master.** Perform a zone transfer from the stub zone's master server regardless of the serial number in the stub zone's SOA resource record.

For more information, see Reload or transfer a stub zone, Add a stub zone, and Understanding stub zones.

# Active Directory-integrated stub zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory-integrated stub zones**
A stub zone can be stored in Active Directory on a DNS server configured as a domain controller. The following parameters are stored in the Active Directory-integrated stub zone:

- Name of the stub zone, such as widgets.example.com.
- Type of zone (stub zone).
- List of master servers.

For more information, see Understanding stub zones and Active Directory integration.

**Updating stub zones**
When Active Directory-integrated stub zones are hosted in separate sites, you can update them using a local list of master servers in each site. For example, an Active Directory-integrated stub zone, widgets.example.com, was loaded in a site in Seattle and propagated to a site in Santiago, and master servers for the stub zone also exist in each of these sites. When the stub zone in Santiago is updated, the domain controller may contact both master servers for resource records for widgets.example.com, but for network traffic reasons, you might want the domain controller in Santiago to use only the master server in Santiago and not the master server in Seattle. To force the domain controller in Santiago to use only the master server in Santiago, you can specify that the stub zone in Santiago be updated using a local list of master servers.

For information about how to use a local list of master servers to update stub zones, see Configure a stub zone for local master servers.

# Contrasting stub zones and conditional forwarders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Contrasting stub zones and conditional forwarders**

There can be some confusion about when to use conditional forwarders instead of stub zones because both DNS features allow a DNS server to respond to a query with a referral for, or by forwarding to, a different DNS server; however, these settings are used for very different purposes. These features have the following objectives:

- A conditional forwarder setting configures the DNS server to forward a query it receives to a DNS server depending on the DNS name contained in the query.
- A stub zone keeps the DNS server hosting a parent zone aware of all the DNS servers authoritative for a child zone.

**Conditional forwarders**

In situations where you want DNS clients in separate networks to resolve each others' names without having to query DNS servers on the Internet, such as in the case of a company merger, you should configure the DNS servers in each network to forward queries for names in the other network. DNS servers in one network will forward names for clients in the other network to a specific DNS server that will build up a large cache of information about the other network. When forwarding in this way, you create a direct point of contact between two networks' DNS servers, reducing the need for recursion.

Stub zones do not provide the same server-to-server benefit because a DNS server hosting a stub zone in one network will reply to queries for names in the other network with a list of *all* authoritative DNS servers for the zone with that name, instead of the specific DNS servers you have designated to handle this traffic. This configuration complicates any type of security settings that you want to establish between specific DNS servers running in each of the networks. For more information, see Understanding forwarders.

**Stub zones**

Stub zones are used when you want a DNS server hosting a parent zone to remain aware of the authoritative DNS servers for one of its child zones. If the stub zone for a child zone is hosted on the same DNS server as the parent zone, the DNS server hosting the stub zone will receive a list of all new authoritative DNS servers for the child zone when it requests an update from the stub zone's master server . This method of updating the DNS server hosting the parent zone maintains a current list of the authoritative DNS servers for the child zone as they are added and removed.

A conditional forwarder is not an efficient method of keeping a DNS server hosting a parent zone aware of the authoritative DNS servers for a child zone. If you used this method, whenever the authoritative DNS servers for the child zone changed, the conditional forwarder setting on the DNS server hosting the parent zone would have to be manually configured with the IP address for each new authoritative DNS server for the child zone. For more information, see Understanding stub zones.

# Checking names and zone data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checking names and zone data**

The following name-checking methods can be used with the DNS Server service and managed using the DNS console.

| Method | Description |
|---|---|
| Strict RFC (ANSI) | Uses strict checking of names according to RFC 1123 Internet host naming specifications. |
| Non RFC (ANSI) | Permits names that are nonstandard and that do not follow RFC 1123 Internet host naming specifications. |
| Multibyte (UTF8) | Permits recognition of characters other than ASCII, including Unicode, which is normally encoded as more than one octet (8 bits) in length.<br><br>With this option, multibyte characters can be transformed and represented using Unicode Transformation Format (UTF-8) support. |
| All names | Permits any naming conventions. |

**Notes**

- For more information on when to use nondefault settings for name checking with the DNS Server service, see Unicode character support.
- For more information about RFCs, see DNS RFCs.

**Checking for bad zone data**

You can configure the DNS Server service to check for bad or unrecognized zone data when the service is started. To enable this option, see Prevent loading of a zone when bad data is found.

**DNS wildcarding**

The DNS Server service can use wildcards to associate multiple possible DNS name destinations with a single computer acting as a gateway for all names based on a wildcarded name. The most common use of wildcards in DNS records is with the mail exchanger (MX) resource record; for example, to configure DNS for an Internet mail gateway host routing mail between your private network and the Internet.

**Basics of using DNS wildcards**

When you are trying to determine where to use a wildcard in DNS, such as an MX resource record used to route and forward mail, note these guidelines:

- The asterisk (*) is the character reserved to indicate the wildcarded name.
- The asterisk is only used once as the leading (leftmost) character in a DNS name. For example, *.example.microsoft.com.
- If used, an asterisk must always match at least one or more whole labels of a name that precede any non-wildcarded (exact) characters in the remaining part of the name. For example, the wildcard name *.example.microsoft.com may apply to one or more descendents of a domain, such as haifa.example.microsoft.com or seattle.us.example.microsoft.com, but it will not apply to the domain name itself, which is example.microsoft.com.
- Wildcards are not used when a query for an exact name produces an exact match.
- The contents of wildcarded resource records conform to normal DNS formats and rules for resource records.

Typically, wildcards are only used when the DNS name queried is an indeterminate name subordinate to a known name. In this

case, the closest match to the wildcarded name is used. For example, if the following MX resource record is used in the example.microsoft.com zone:

*.example.microsoft.com.  IN  MX  10  mailserver1.example.microsoft.com.

This record would allow mail destined for the following possible DNS domain names to be routed to the mail routing host (mailserver1.example.microsoft.com.) specified in the MX resource record:

haifa.example.microsoft.com.

seattle.us.example.microsoft.com.

santiago.sa.example.microsoft.com.

When wildcarding is used, the asterisk (*****) wildcard symbol must match at least one whole label beyond the specified part of the DNS name in which the wildcard is used.

For example, an e-mail message destined to user@example.microsoft.com could not use the previous example MX resource record to be forwarded to its destination. For the e-mail to be forwarded correctly, an additional explicit MX resource record needs to be added in the same zone, such as:

example.microsoft.com.  IN  MX  10  mailserver1.example.microsoft.com.

**Loose RFC wildcarding**
For Windows NT Server 4.0 (prior to Service Pack 4), the DNS Server service was not strictly compliant with the method for using wildcards in a DNS name lookup described in RFC 1034. It only stopped a wildcard lookup when it reached a node in the DNS domain namespace tree that had resource records of the type specified in a query.

For example, if in the example.microsoft.com. zone the following records were used:

```
      *.example.microsoft.com.             IN  MX  10
mailserver1.example.microsoft.com.
      example.microsoft.com.               IN  MX  10
mailserver2.example.microsoft.com.
      host-a.example.microsoft.com.        IN  A       10.0.0.1
      mailserver1.example.microsoft.com.   IN  A       10.0.0.2
      mailserver2.example.microsoft.com.   IN  A       10.0.0.3
```

and a DNS query was made by a computer, Host B, as follows:

**Queried name:** host-a.example.microsoft.com.

**Queried type:** MX

with loose wildcarding in effect, the following sequence occurs between the DNS Server service and Host B:

1. Host B sends the example query to a DNS server configured to use loose RFC wildcarding, such as a DNS server running under Windows NT Server 4.0 prior to Service Pack 4.

2. No MX resource records match the specified name (host-a.example.microsoft.com.).

3. An MX resource record -- the type specified in the query -- is included in the zone and uses a wildcarded name that is the parent of the queried name (*.example.microsoft.com.).

4. The DNS Server service checks the wildcarded MX resource record name (*.example.microsoft.com.) for its mapped host names, finds a matching A resource record for the specified name (mailserver1.example.microsoft.com.) and uses this information to answer Host B.

5. The mail program on Host B can then send mail to mailserver1.example.microsoft.com. by using its IP address.

The loose RFC approach allows you to set up only two MX records and cover queries to the entire zone. However, the two MX resource records need to be added for both the zone root (the explicit name of the zone) and a wildcard node (*****) directly below the zone root. However, if you have individual mail hosts that should receive their own mail directly, then add MX resource records for each of those hosts that refer directly to their own host name, as stated in their registered A resource record.

**Strict RFC wildcarding**

With the strict wildcarding approach as described in RFC 1034, a revised sequence of events occurs when reviewing the previous example. Based on the same specified zone and query information provided, the revised sequence of steps would be used between the DNS Server service and Host B:

1. Host B sends the example query to a DNS server configured to use strict RFC wildcarding.

2. No MX resource records match the specified name (host-a.example.microsoft.com.), but an A resource record of that name is found in the zone.

   Because the queried name was found, no wildcarding is done.

3. The DNS Server service responds to the query with an authoritative negative response that includes no errors or resource records in the answer section of the response.

4. The mail program on Host B then sends a new query for the same name, specifying that A resource records be used for the query type.

5. The DNS Server service answers the A query with an authoritative positive response, which includes the A resource record in the answer section of the response.

6. The mail program on Host B can then send mail to mailserver1.example.microsoft.com. by using its IP address.

The strict RFC approach to wildcarding is easier if your mail hosts are capable of receiving mail directly. The disadvantage to this approach is that you typically need to add more MX resource records to the zone. This can happen if you have mail hosts not capable of receiving mail directly, or if you need to advertise any mail hosts that are not mail servers.

# Using WINS lookup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using WINS lookup**

The DNS Server service provides the ability to use Windows Internet Name Service (WINS) servers to look up names not found in the DNS domain namespace by checking the NetBIOS namespace managed by WINS.

To use WINS lookup integration, two special resource record (RR) types--the WINS and WINS-R resource records--are enabled and added to a zone. When the WINS RR is used, DNS queries that fail to find a matched host (A) RR in the zone are forwarded to WINS servers configured in the WINS RR. For reverse lookup zones, the WINS-R RR can be enabled and used to provide a similar benefit for further resolving a reverse query not answerable in the reverse in-addr.arpa domain.

A good example of when to use WINS lookup is when you are using a mixed-mode client environment consisting of UNIX clients that use only DNS name resolution and earlier-version Microsoft clients that require NetBIOS naming. In these environments, WINS lookup provides a method for permitting UNIX DNS clients to locate your WINS clients by extending DNS host name resolution into the WINS-managed NetBIOS namespace.

The WINS lookup integration feature is only supported by Windows DNS servers. If you use a mixture of Windows and other DNS servers to host a zone, you should enable the **Do not replicate this record** check box option for any primary zones when using the WINS lookup record. This prevents the WINS lookup record from being included in zone transfers to other DNS servers that do not support or recognize this record. If you do not enable the WINS lookup record to be used only on the local server, it can cause data errors or failed zone transfers at servers running other DNS server implementations that replicate the zone.

The following sections discuss WINS lookup interoperability considerations and explain the following:

- How the caching Time-To-Live (TTL) and lookup time-out values are configured for use with the WINS and WINS-R records.
- The format of the WINS and WINS-R resource records as they are used in zone files created by the DNS Server service.

**WINS lookup interoperability**

Typically, WINS lookup provides the best and most predictable results if only Windows DNS servers are used, and is only available directly for use at Windows DNS servers. There are ways, however, that you use and benefit from WINS lookup as an interoperable solution when other DNS servers are deployed.

For example, consider adding a Windows DNS server that hosts a new WINS lookup-enabled zone. When creating and naming the zone, use a subdomain added to your existing DNS namespace that is used just for WINS-specific referrals added to your DNS domain namespace.

For instance, in example.microsoft.com, call the zone wins.example.microsoft.com when you create it. You can then use this new WINS referral zone as the root zone for any of your WINS-aware computers with names that are not found in your other traditional DNS zones.

In order to use the WINS referral zone, you need to specify its domain name (wins.example.microsoft.com) in a DNS suffix search list for your clients. The suffix list is configurable as part of the TCP/IP properties for a client connection and can be updated either manually, using DHCP, or using Group Policy. As long as the name of the WINS referral zone is included in the domain suffix list, any DNS names not resolved in traditional zones can be resolved using the WINS referral subdomain.

Under normal conditions, this should result in recursion from your other DNS servers to the Windows DNS servers that host the WINS-enabled zone. If the queried host names match NetBIOS computer names found in the WINS database, the names are resolved to the IP addresses mapped in WINS data there.

In our example the WINS-enabled zone is only used for WINS lookup, therefore no additional resource records need to be added to it. In general, WINS records could be added to any forward lookup zone.

By using a specific subdomain just for WINS lookup and specifying a static DNS suffix list to be used in resolving and searching for names, you can prevent unusual situations in which DNS queries for different FQDNs resolve to the same WINS client name and IP address. This might easily occur if you add and configure many zones at each level of your namespace and enable each of them to use WINS lookup integration.

For example, suppose you have two zones both configured to use WINS lookup. The zones are rooted and originate at the following DNS domain names:

example1.microsoft.com.

example2.microsoft.com.

With this configuration, a WINS client named HOST-A could be unintentionally resolved using either of the following FQDNs:

host-a.example1.microsoft.com.

host-a.example2.microsoft.com.

#### WINS resource record

**Description:** WINS forward lookup resource record. Used in a zone to provide further resolution of DNS queries for names not found in the zone by using a name query to WINS servers configured and listed with this record. If used, the WINS record only applies for the topmost level within a zone and not for subdomains used in the zone. The various fields that are used with the WINS RR include the following:

- ***owner*** Indicates the owner domain for this record, should always be set to "@" to indicate that the current domain is the same as the zone origin.
- ***class*** Indicates the class for this record, should always be set to "IN" as the Internet class is the only supported class for DNS servers running Windows Server 2003.
- **LOCAL** When used, indicates that the WINS RR is to used only locally at the DNS server and is not to be included during zone replication with other DNS servers. This field corresponds to whether the **Do not replicate this record** check box was selected when configuring WINS lookup at the DNS console. If this check box was cleared, this field will not be included when the record is written to the zone.
- ***lookup_timeout*** The lookup time-out value to be applied for this record.
- ***cache_timeout*** The cache time-out value to be applied for this record.
- ***wins_ip_addresses*** Used to specify one or more IP addresses of WINS servers. At least one IP address of a valid WINS server is required.

**Syntax:** *owner class* **WINS** [LOCAL] [**L***lookup_timeout*] [**C**cache_timeout] wins_ip_addresses

**Example:**

```
@ IN        WINS  10.0.0.1
@ IN        WINS LOCAL L1 C10 10.10.10.1 10.10.10.2 10.10.10.3
```

**Note**

- In the provided example WINS records, the zone root is assumed to be the current origin.

#### WINS-R resource record

**Description:** WINS reverse lookup resource record. Used in a reverse lookup zone to provide further resolution for reverse queries not found in the zone by using a NetBIOS adapter node status query sent to the queried IP addresses. When this record is used, you need to specify the parent domain to be appended to a NetBIOS computer name when a successful reverse lookup occurs. Other fields used in the WINS-R record have similar description and purpose as described previously for their use in the WINS forward lookup record.

**Syntax:** *owner class* **WINS** [LOCAL] [**L***lookup_timeout*] [**C**cache_timeout] *Domain_to_append_to_returned_NetBIOS_names*

**Example:**

```
@ IN       WINS-R LOCAL L1 C10 example.microsoft.com.
@ IN       WINS-R wins.example.microsoft.com.
```

**Note**

- In the provided example WINS-R records, the zone root is assumed to be the current origin.

**Advanced parameters for WINS lookups**

Two following advanced timing parameters are used with the WINS and WINS-R records:

- The Cache timeout value, which indicates to a DNS server how long it should cache any of the information returned in a WINS lookup. By default, this value is set to 15 minutes.

- The Lookup timeout value, which specifies how long to wait before timing out and expiring a WINS lookup performed by the DNS Server service. By default, this value is set to 2 seconds.

These parameters are configured using the **Advanced** button in the zone properties dialog box when you configure the zone. This button appears on either the WINS or WINS-R tab, depending on whether the zone you are configuring is being used for forward lookup or reverse lookup.

If you are using either the WINS or WINS-R resource record, be aware that the minimum Time-To-Live (TTL) set in the SOA record for the zone is not the default TTL used with these records. Instead, when either an IP address or a host name gets resolved with WINS lookup, the information is cached on the DNS server for the amount of time configured for the WINS cache time-out value. If this address is then ever forwarded to another DNS server, the WINS cache time-out value TTL is what is sent. If your WINS data rarely changes, you can increase the default TTL of 15 minutes.

**Notes**

- If you have a zone configured for WINS lookup, then all DNS servers that are authoritative for that zone need to be capable of WINS lookup or you will have intermittent behavior.

- Because you can select that the WINS and WINS-R RRs not be replicated to other DNS servers, you can selectively enable and configure WINS lookup at each of your secondary servers for zones where this feature is used. This is not a standard practice for other types of resource records which are only to be configured at the primary server for the zone.

For more information, see Enable DNS to use WINS resolution.

# Managing reverse lookup zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing reverse lookup zones**

When adding reverse lookup zones, you can simplify the process of adding pointer (PTR) resource records to these zones by:

- Using the **Update associated pointer (PTR) record** option when manually adding host (A) resource records to forward lookup zones.
- Deploying DHCP servers and enabling DHCP-DNS integration features to help manage and update reverse lookup zones for legacy clients.
- Installing Windows XP as the standard operating system in your organization for deploying new client hardware.

  With Windows XP on your client computers, either the client or its Windows Server 2003 DHCP server can be enabled to perform pointer (PTR) resource record registration for the client when its IP address is obtained and configured.

**Subnetted reverse lookup zones**

Default subnet masks (which have either a value of 0 or 255 used in each of the four octets of the masked value) are assumed by the DNS when configuring a new reverse lookup zone. If necessary, you can use custom subnetting to adjust the subnet mask value to a nondefault length.

Subnetted reverse lookup zones are supported when using the Advanced view setting available through Microsoft Management Console (MMC) and the DNS. This allows you to add and configure a reverse lookup for an IP network address that uses nondefault subnet mask lengths, such as 255.255.248.0, which has a nondefault subnet value of 248 in the third octet.

For more information, see Using the Windows Deployment and Resource Kits.

**Notes**

- Subnetted reverse lookup zones are further described in RFC 2317, "Classless IN-ADDR.ARPA delegation." For more information on RFCs or how to obtain them, see TCP/IP RFCs.

  PTR resource records created automatically when adding an A resource record to a zone will be deleted automatically if the corresponding A resource record is deleted.

# Monitoring and Optimizing Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring and optimizing servers**
This section covers:

- Monitoring DNS server performance
- Using server debug logging options
- Verifying server configuration
- Configuring round robin
- Prioritizing local subnets
- Tuning advanced server parameters
- Using DNS Security Extensions (DNSSEC)
- Using Extension Mechanisms for DNS (EDNS0)

# Monitoring DNS server performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring DNS server performance**

Because DNS servers are of critical importance in most environments, monitoring their performance can help in the following ways:

- By providing a useful benchmark for predicting, estimating, and optimizing DNS server performance.
- In troubleshooting DNS servers where server performance has degraded either over time or during periods of peak activity.

As you begin your DNS server monitoring, you can review sample test results for DNS servers that were collected during product development and testing. You can use this information as a starting reference as you begin to monitor your DNS servers to measure and benchmark performance.

A set of DNS server performance counters are provided with System Monitor to measure and monitor various aspects of server activity.

**Sample DNS server performance test results**

The DNS development and testing team compiled some statistics as a profile of DNS server performance during preliminary testing. In testing, two different DNS server hardware configurations were used and overall DNS query and dynamic update activity was measured, along with processor utilization.

The results of each of these tests are listed in the table below.

| Server configuration | Queries/sec | Dynamic updates/sec | Processor utilization |
|---|---|---|---|
| Intel P-III 733 MHz single-processor | 9500 | 1300 | 75% |

During the collection of these measurements (over 4 days), the monitored DNS server was processing both queries and dynamic updates simultaneously. The numbers above reflect this.

For dynamic updates, standard primary type zones were used, not Active Directory-integrated zones. Where directory integration is used for zones, the rate at which dynamic updates can be processed decreases, since the DNS server must additionally write to and rely upon the Active Directory database.

In addition, if a zone is configured to accept only secure dynamic updates, the update rate can also decrease. Network performance might also be a factor in these cases since the directory database may require network activity to process updates.

For additional background information about the two server hardware configurations used during these tests, see the following section.

**Important**

- These measurements are not intended to indicate maximum performance or server limitations for DNS servers running Windows Server 2003. The objective of the tests was to sample typical DNS server performance and obtain a working benchmark based on standard available hardware as a basis to begin server capacity planning.
- Note that these numbers are approximate and can be influenced by the type of resource records entered in zones, the number of resource records with the same owner name, and the number of zones in use at a specific DNS server.

**Background information on sample DNS server tests**

The list of specific hardware that the Windows Server 2003 DNS development and test team used during the described testing included the following:

| Hardware components | Sizing |
|---|---|

| | |
|---|---|
| Number of processors | One |
| Processor | Intel Pentium III 733 MHz |
| Amount of RAM | 256 MB (megabytes) |
| Hard disk drive space | 4 GB (gigabytes) |

These measurements were based on the server computer running a DNS server, with no additional services in use. Where other hardware specifications or software configurations are made when deploying DNS servers, your performance results are likely to vary from those documented here.

For more information on determining how to perform DNS server capacity planning, see Server planning for DNS.

**DNS server performance counters**

DNS server performance counters can be used to measure and monitor various aspects of server activity, such as the following:

- Overall DNS server performance statistics, such as the number of overall queries and responses processed by a DNS server.
- User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) counters, for measuring DNS queries and responses that are processed respectively using either of these transport protocols.
- Dynamic update and secure dynamic update counters, for measuring registration and update activity generated by dynamic clients.
- Memory usage counters, for measuring system memory usage and memory allocation patterns created by operating the server computer as a DNS server running Windows Server 2003.
- Recursive lookup counters, for measuring queries and responses when the DNS Server service uses recursion to look up and fully resolve DNS names on behalf of requesting clients.
- WINS lookup counters, for measuring queries and responses made to WINS servers when the WINS lookup integration features of the DNS Server service are used.
- Zone transfer counters, including specific counters for measuring: all zone transfer (AXFR), incremental zone transfer (IXFR), and DNS zone update notification activity.

Each of these DNS server performance counters is described in more detail in the following sections.

**All zone transfer (AXFR) counters**

| Field | Description |
|---|---|
| AXFR Request Received | The total number of full zone transfer requests received by the DNS Server service when operating as a master server for a zone. |
| AXFR Request Sent | The total number of full zone transfer requests sent by the DNS Server service when operating as a secondary server for a zone. |
| AXFR Response Received | The total number of full zone transfer requests received by the DNS Server service when operating as a secondary server for a zone. |
| AXFR Success Received | The total number of full zone transfers received by the DNS Server service when operating as a secondary server for a zone. |
| AXFR Success Sent | The total number of full zone transfers successfully sent by the DNS Server service when operating as a master server for a zone. |

**DNS server memory counters**

| Field | Description |
| --- | --- |
| Caching Memory | The total amount of system memory in use by the DNS Server service for caching. |
| Database Node Memory | The total amount of system memory in use by the DNS Server service for database nodes. |
| Nbtstat Memory | The total amount of system memory in use by the DNS Server service for Nbtstat. |
| Record Flow Memory | The total amount of system memory in use by the DNS Server service for record flow. |

**Dynamic update counters**

| Field | Description |
| --- | --- |
| Dynamic Update NoOperation | The total number of No-operation/Empty dynamic update requests received by the DNS server. |
| Dynamic Update NoOperation/sec | The average number of No-operation/Empty dynamic update requests received by the DNS server in each second. |
| Dynamic Update Queued | The total number of dynamic updates queued by the DNS server. |
| Dynamic Update Received | The total number of dynamic update requests received by the DNS server. |
| Dynamic Update Received/sec | The average number of dynamic update requests received by the DNS server in each second. |
| Dynamic Update Rejected | The total number of dynamic updates rejected by the DNS server. |
| Dynamic Update TimeOuts | The total number of dynamic update time-outs of the DNS server. |
| Dynamic Update Written to Database | The total number of dynamic updates written to the database by the DNS server. |
| Dynamic Update Written to Database/sec | The average number of dynamic updates written to the database by the DNS server in each second. |

**Incremental zone transfer (IXFR) counters**

| Field | Description |
| --- | --- |
| IXFR Request Received | The total number of incremental zone transfer requests received by the master DNS server. |
| IXFR Request Sent | The total number of incremental zone transfer requests sent by the secondary DNS server. |
| IXFR Response Received | The total number of incremental zone transfer responses received by the secondary DNS server. |
| IXFR Success Received | The total number of successful incremental zone transfers received by the secondary DNS server. |
| IXFR Success Sent | The total number of successful incremental zone transfers of the master DNS server. |
| IXFR TCP Success Received | The total number of successful TCP incremental zone transfers received by the secondary DNS server. |
| IXFR UDP Success Received | The total number of successful UDP incremental zone transfers received by the secondary DNS server. |

**Notification counters**

| Field | Description |
| --- | --- |
| Notify Received | The total number of notifies received by the secondary DNS server. |
| Notify Sent | The total number of notifies sent by the master DNS server. |

**Recursion counters**

| Field | Description |
|---|---|
| Recursive Queries | The total number of recursive queries received by the DNS server. |
| Recursive Queries/sec | The average number of recursive queries received by the DNS server in each second. |
| Recursive Query Failure | The total number of recursive query failures. |
| Recursive Query Failure/sec | The average number of recursive query failures in each second. |
| Recursive TimeOuts | The total number of recursive query sending time-outs. |
| Recursive TimeOuts/sec | The average number of recursive query sending time-outs in each second. |

**Secure dynamic update counters**

| Field | Description |
|---|---|
| Secure Update Failure | The total number of secure updates failed on the DNS server. |
| Secure Update Received | The total number of secure update requests received by the DNS server. |
| Secure Update Received/sec | The average number of secure update requests received by the DNS server in each second. |

**TCP counters**

| Field | Description |
|---|---|
| TCP Message Memory | The total TCP message memory used by the DNS server. |
| TCP Query Received | The total number of TCP queries received by the DNS server. |
| TCP Query Received/sec | The average number of TCP queries received by the DNS server in each second. |
| TCP Response Sent | The total number of TCP responses sent by the DNS server. |
| TCP Response Sent/sec | The average number of TCP responses sent by the DNS server in each second. |

**Total (overall performance) counters**

| Field | Description |
|---|---|
| Total Query Received | The total number of queries received by the DNS server. |
| Total Query Received/sec | The average number of queries received by the DNS server in each second. |
| Total Response Sent | The total number of responses sent by the DNS server. |
| Total Response Sent/sec | The average number of responses sent by the DNS server in each second. |

**UDP counters**

| Field | Description |
|---|---|
| UDP Message Memory | The total UDP message memory used by the DNS server. |
| UDP Query Received | The total number of UDP queries received by the DNS server. |
| UDP Query Received/sec | The average number of UDP queries received by the DNS server in each second. |
| UDP Response Sent | The total number of UDP responses sent by the DNS server. |
| UDP Response Sent/sec | The average number of UDP responses sent by the DNS server in each second. |

**WINS lookup counters**

| Field | Description |
|---|---|
| WINS Lookup Received | The total number of WINS lookup requests received by the server. |
| WINS Lookup Received/sec | The average number of WINS lookup requests received by the server in each second. |
| WINS Response Sent | The total number of WINS lookup responses sent by the server. |
| WINS Response Sent/sec | The average number of WINS lookup responses sent by the server in each second. |
| WINS Reverse Lookup Received | The total number of WINS reverse lookup requests received by the server. |
| WINS Reverse Lookup Received/sec | The average number of WINS reverse lookup requests received by the server in each second. |
| WINS Reverse Response Sent | The total number of WINS Reverse lookup responses sent by the server. |
| WINS Reverse Response Sent/sec | The average number of WINS Reverse lookup responses sent by the server in each second. |

**Zone transfer counters**

| Field | Description |
|---|---|
| Zone Transfer Failure | The total number of failed zone transfers of the master DNS server. |
| Zone Transfer Request Received | The total number of zone transfer requests received by the master DNS server. |
| Zone Transfer SOA Request Sent | The total number of zone transfer SOA requests sent by the secondary DNS server. |
| Zone Transfer Success | The total number of successful zone transfers of the master DNS server. |

# Using server debug logging options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using server debug logging options**

The following DNS debug logging options are available:

- **Direction of packets**

  **Send** Packets sent by the DNS server are logged in the DNS server log file.

  **Receive** Packets received by the DNS server are logged in the log file.

- **Content of packets**

  **Standard queries** Specifies that packets containing standard queries (per RFC 1034) are logged in the DNS server log file.

  **Updates** Specifies that packets containing dynamic updates (per RFC 2136) are logged in the DNS server log file.

  **Notifies** Specifies that packets containing notifications (per RFC 1996) are logged in the DNS server log file.

- **Transport protocol**

  **UDP** Specifies that packets sent and received over UDP are logged in the DNS server log file.

  **TCP** Specifies that packets sent and received over TCP are logged in the DNS server log file.

- **Type of packet**

  **Request** Specifies that request packets are logged in the DNS server log file (a request packet is characterized by a QR bit set to 0 in the DNS message header).

  **Response** Specifies that response packets are logged in the DNS server log file (a response packet is characterized by a QR bit set to 1 in the DNS message header).

- **Enable filtering based on IP address** Provides additional filtering of packets logged in the DNS server log file. This option allows logging of packets sent from specific IP addresses to a DNS server, or from a DNS server to specific IP addresses.

- **File name** Lets you specify the name and location of the DNS server log file.

  For example:

  - *dns.log* specifies that the DNS server log file should be saved as *dns.log* in the systemroot\System32\Dns directory.

  - *temp\dns.log* specifies that the DNS server log file should be saved as *dns.log* in the systemroot\Temp directory

- **Log file maximum size limit** Lets you set the maximum file size for the DNS server log file. When the specified maximum size of the DNS server log file is reached, the DNS server overwrites the oldest packet information with new information. Note: If left unspecified, the DNS server log file's size can take up a large amount of hard disk space.

By default, all debug logging options are disabled. When selectively enabled, the DNS Server service can perform additional trace-level logging of selected types of events or messages for general troubleshooting and debugging of the server.

Debug logging can be resource intensive, affecting overall server performance and consuming disk space. Therefore, it should only be used temporarily when more detailed information about server performance is needed.

**Notes**

- Dns.log contains debug logging activity. By default, it is located in the *windir*\System32\Dns folder.
- To enable and use file-based logging, see Select and enable debug logging options on the DNS server.
- For more information about RFCs, see DNS RFCs.

# Verifying server configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Verifying server configuration**

When using the DNS console, you can perform either manual or automated verification testing of your DNS servers to monitor their ability to process and resolve queries. This feature is accessed through the **Monitoring** tab in DNS server properties.

When using this feature, two types of test are available:

- **A simple query against this DNS server**.

  This type of test specifies that the DNS server perform a simple or iterative query. This test is a localized query using the DNS resolver (client) on the server computer to query the local DNS server (also located on the same computer).

  For more information, see How DNS query works.

- **A recursive query to other DNS servers.**

  This type of test specifies that the DNS server perform a recursive query. This test is similar in its initial query processing to the previous test in that it uses the local DNS resolver (client) to query the local DNS server, also located on the same computer.

  In this test, however, the client asks the server to use recursion to resolve an NS-type query for the root of the DNS domain namespace, stated as a single dot ("."). This type of query should typically require additional recursive processing and can be helpful in verifying that server root hints or zone delegations have been properly set.

  For more information, see How DNS query works.

Once you have selected the tests to be used, you can either click the **Test Now** button to manually perform the tests immediately, or you can perform automatic testing at a specified time interval.

Automated tests are performed when the **Perform automatic testing at the following interval** check box is selected and a test interval is specified in **Test interval**.

Results of all selected tests, whether manually or automatically performed, are displayed in the **Test results** list box. Typically, this information includes:

- The date and time when each query was made.

- Additional status results of the specific test used, such as whether the simple or recursive query failed or succeeded.

# Configuring round robin

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring round robin**
Round robin is a local balancing mechanism used by DNS servers to share and distribute network resource loads. You can use it to rotate all resource record (RR) types contained in a query answer if multiple RRs are found.

By default, DNS uses round robin to rotate the order of RR data returned in query answers where multiple RRs of the same type exist for a queried DNS domain name. This feature provides a simple method for load balancing client use of Web servers and other frequently queried multihomed computers.

If round robin is disabled for a DNS server, the order of the response for these queries is based on a static ordering of RRs in the answer list as they are stored in the zone (either its zone file or Active Directory).

**Example: Round-robin rotation**
A forward lookup-type query (for all A RRs that match a DNS domain name) is made for a multihomed computer (multihomed.example.microsoft.com) that has three IP addresses. Separate A RRs are used to map the host's name to each of these IP addresses in the zone. In the stored example.microsoft.com zone, the RRs appear in this fixed order:

multihomed IN A 10.0.0.1

multihomed IN A 10.0.0.2

multihomed IN A 10.0.0.3

The first DNS client that queries the server to resolve this host's name receives the list in default order. When a second client sends a subsequent query to resolve this name, the list is rotated as follows:

multihomed IN A 10.0.0.2

multihomed IN A 10.0.0.3

multihomed IN A 10.0.0.1

**Restricting round-robin rotation for selected RR types**
By default, DNS will perform round-robin rotation for all RR types. You can specify that certain RR types are not to be round-robin rotated in the registry. There is a registry entry called **DoNotRoundRobinTypes** (REG_SZ) with a string value containing a list of RR types. By modifying this entry, you turn off round-robin rotation for specific RR types. For example, to prevent round-robin rotation for A, PTR, SRV, and NS record types, you would enter the following value for the registry entry:

**a ptr srv ns**

**Restricting round-robin rotation for all RR types**
The default setting for round-robin rotation is contained in the registry entry **RoundRobin** (REG_DWORD). By default, this entry's value is 1, rotating all RR types except those listed in the **DoNotRoundRobinTypes** registry entry. If the value of **RoundRobin** is set to 0, then no RR types will be round-robin rotated.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

### Notes

- The string value for the DoNotRoundRobinTypes registry entry may contain types in numeric (as shown above) of mnemonic formats.

- Both of the round robin registry entries must be created and stored in the following registry subkey:

  HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

- Local subnet priority supersedes the use of round-robin rotation for multihomed names. When enabled, round-robin continues to be a secondary method used to sort multiple RRs returned in a listed answer. For more information on local subnet priority, see Prioritizing local subnets.

# Prioritizing local subnets

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Prioritizing local subnets**

By default, the DNS Server service uses local subnet prioritizing as the method for giving preference to IP addresses on the same network when a client query resolves to a host name that is mapped to more than one IP address. This feature requires that the client application attempt to connect to the host using its closest (and typically fastest) IP address available for connection.

The DNS Server service uses local subnet priority as follows:

1. The DNS Server service determines if local subnet prioritization is needed to order the query response.

   If more than one A resource record (RR) matches the queried host name, the DNS Server service can reorder the records by their subnet location. If the queried host name only matches a single A resource record, or if the IP network address of the client does not match an IP network address for any of the mapped addresses in an answer list of multiple RRs, no prioritizing is necessary.

2. For each RR in the matched answer list, the DNS Server service determines which records (if any) match the subnet location of the requesting client.

3. The DNS Server service reorders the answer list so that A RRs which match the local subnet of the requesting client are placed first in the answer list.

4. Prioritized by subnet order, the answer list is returned to the requesting client.

**Simple example: Local network prioritizing**

A multihomed computer, multihomed.example.microsoft.com, has three A RRs for its three separate host IP addresses in the example.microsoft.com zone. A separate A RR is used for each of the host's addresses, which appear in this order in the zone:

multihomed IN A 192.168.1.27

multihomed IN A 10.0.0.14

multihomed IN A 172.16.20.4

If a DNS client resolver at IP address 10.4.3.2 queries the server for the IP addresses of host multihomed.example.microsoft.com, the DNS Server service notes that the originating IP network address (10.0.0.0) of the client matches the network (class A) portion of the 10.0.0.4 address in the answer list of RRs. The DNS Server service then reorders the addresses in the response as follows:

multihomed IN A 10.0.0.14

multihomed IN A 192.168.1.27

multihomed IN A 172.16.20.4

If the IP address of the requesting client has no local network match with any of the RRs in the answer list, then the list is not prioritized.

**Complex example: Local subnet prioritizing**

If you are working in a network that uses IP subnetting (nondefault subnet masks), there are only a few minor differences. If more than one address matches in the network portion, then the matching addresses are further ordered and the RR having the closest matching subnet address is placed first.

For example, a multihomed computer, multihomed.example.microsoft.com, has four A RRs for four separate host IP addresses in the example.microsoft.com zone. Two of these IP addresses are for nonlocal networks. The other two IP addresses share a common IP network address but, because IP subnetting is used, represent different physical subnetted network connections based on their custom (nondefault) subnet mask value of 255.255.248.0. These example RRs appear in the following order in

the zone:

```
multihomed   IN  A  192.168.1.27

multihomed   IN  A  172.16.22.4

multihomed   IN  A  10.0.0.14

multihomed   IN  A  172.16.31.5
```

If the IP address of the requesting client is 172.16.22.8, both of the IP addresses that match the same IP network as the client, the 172.16.0.0 network, are returned at the top of the answer list to the client. However, in this example, the 172.16.22.4 address is placed ahead of the 172.16.31.5 address because it matches the client IP address down through the 172.16.20.0 subnet address.

The reordered answer list returned by the DNS Server service would be:

```
multihomed   IN  A  172.16.22.4

multihomed   IN  A  172.16.31.5

multihomed   IN  A  192.168.1.27

multihomed   IN  A  10.0.0.14
```

**Notes**

- IP subnetting is imposed by using a custom or nondefault subnet mask value with all IP addresses on your network. For more information, see Subnet masks.

- Local subnet priority supersedes the use of round-robin rotation for multihomed names. When round robin is enabled, however, RRs continue to be rotated using round robin as the secondary method of sorting the response list. For more information on round robin, see Configuring round robin.

# Tuning advanced server parameters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tuning advanced server parameters**

When initialized for service, DNS servers use server configuration settings taken from the parameters stated in a boot information file, the registry, and possibly zone information provided through Active Directory integration.

In most situations, the installation defaults are acceptable and should not require modification. However, when needed, you can use the DNS console to tune the following advanced parameters, accommodating special deployment needs and situations.

For making these adjustments, use the **Advanced** tab in DNS server **Properties**.

| Value | Description |
|---|---|
| Disable recursion | Determines whether or not the DNS server uses recursion. By default, the DNS Server service is enabled to use recursion. For more information, see Modifying server defaults. |
| BIND secondaries | Determines whether to use fast transfer format when transferring a zone to DNS servers running legacy Berkeley Internet Name Domain (BIND) implementations.<br><br>By default, all Windows-based DNS servers use a fast zone transfer format, which uses compression and can include multiple records per TCP message during a connected transfer. This format is also compatible with more recent BIND-based DNS servers that run versions 4.9.4 and later.<br><br>For more information, see Interoperability issues. |
| Fail on load if bad zone data | Sets the DNS server to parse files strictly.<br><br>By default, the DNS Server service logs data errors, ignore any erred data in zone files, and continue to load a zone. This option can be reconfigured using the DNS console so that the DNS Server service logs errors and fails to load a zone file containing records data that is determined to have errors.<br><br>For more information, see Checking names and zone data. |
| Enable round robin | Determines whether the DNS server uses round robin to rotate and reorder a list of resource records (RRs) if multiple RRs exist of the same type exisit for a query answer.<br><br>By default, the DNS Server service uses round robin. For more information, see Configuring round robin. |
| Enable netmask ordering | Determines whether the DNS server reorders A resource records within the same resource record set in its response to a query based on the IP address of the source of the query.<br><br>By default, the DNS Server service uses local subnet priority. For more information, see Prioritizing local subnets. |
| Secure cache against pollution | Determines whether the server attempts to clean up responses to avoid cache pollution. This setting is enabled by default.<br><br>By default, DNS servers use a secure response option that eliminates adding unrelated resource records included in a referral answer to their cache. In most cases, any names added in referral answers are typically cached and help expedite the speed of resolving subsequent DNS queries.<br><br>With this feature, however, the server can determine that referred names are potentially polluting or insecure and discard them. The server determines whether to cache the name offered in a referral on the basis of whether or not it is part of the exact related DNS domain name tree for which the original queried name was made.<br><br>For example, if a query was originally made for "example.microsoft.com" and a referral answer provided a record for a name outside of the "microsoft.com" domain name tree, such as msn.com, then that name would not be cached where this feature is enabled for use. |

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Notes**

- It is recommended that you use the DNS console to configure these and other server parameter values.

- Additional advanced parameters can be configured by registry-based configuration. For more information, see Using the Windows Deployment and Resource Kits.

- DNS server parameter key values are stored in the following location in the registry:

  HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters

- DNS zone parameter key values are stored in the following location in the registry:

  HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\DNS Server

- When using registry-based configuration, changes are applied to DNS servers only when the service is re-initialized. In these cases, if a value is manually changed directly in the registry, the DNS server must always be restarted for the new value to be used.

  For more information, see Start or stop a DNS server.

# Using DNS Security Extensions (DNSSEC)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using DNS Security Extensions (DNSSEC)**
Windows Server 2003 DNS provides basic support of the DNS Security Extensions (DNSSEC) protocol as defined in RFC 2535. The current feature support allows DNS servers to perform as secondary DNS servers for existing DNSSEC-compliant, secure zones. DNS supports the storing and loading of the DNSSEC-specific resource records (RRs). Currently, a DNS server is not capable of signing zones and resource records (creating cryptographic digital signatures) or validating the SIG RRs. The DNSSEC resource records are KEY, SIG, and NXT. For more information about resource records, see Resource records reference. For more information about RFCs, see DNS RFCs.

**Server support for DNSSEC**
When loading a zone containing DNSSEC resource records, the DNS server loads these records along with all other types of resource records contained in the zone. When receiving a zone transfer containing DNSSEC resource records (SIG, KEY, NXT), the DNS server writes these records to the zone storage (zone data file or Active Directory) along with all other resource records.

When a DNS server receives a request or response containing DNSSEC resource records, it does not verify the digital signatures but caches the response and uses it for ensuing queries. When a DNS server receives a request for a resource record in a zone also containing DNSSEC resource records, it attaches the appropriate DNSSEC records to the response.

When a signed zone contains resource records for an owner name including a CNAME resource record for that name, the DNS server will return the DNSSEC resource records associated with the owner name and the CNAME resource record's alias name. The DNS server will not suppress the retrieval of the CNAME resource record, and it will not return a SIG resource record for the canonical name. Rather, it will return the SIG resource record for the alias name.

**Client support for DNSSEC**
The DNS client does not read and store a key for the trusted zone and, consequently, it does not perform any cryptography, authentication, or verification. When a resolver initiates a DNS query and the response contains DNSSEC resource records, programs running on the DNS client will return these records and cache them in the same manner as any other resource records. This is the extent to which Windows XP DNS clients support DNSSEC. When the DNS client receives the SIG RR relating to the RRset, it will not perform an additional query to obtain the associated KEY record or any other DNSSEC records.

Resolvers do not authenticate resource records by verifying the signature information contained in the SIG resource record. The DNS client does not contain any information to indicate which resource records have been authenticated or to what extent they have been authenticated.

When a resolver receives a response or performs a query operation, it does not recognize the checking disabled (CD) query header bit, which in DNSSEC indicates that the data is authenticated by the server according to its policies, or set the authentic data (AD) query header bit, which in DNSSEC indicates that nonauthenticated data is acceptable to the resolver.

For more information, see Modify DNSSEC configuration and DNSSEC overview.

### Note

- If there is a signing agent running on the DNS server running Windows Server 2003 that signs the zone resource records, this DNS server may also be used as a primary server for DNSSEC-compliant secure zones.

# Using Extension Mechanisms for DNS (EDNS0)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Extension Mechanisms for DNS (EDNS0)**
Extension Mechanisms for DNS (EDNS0 as defined in RFC 2671) allow DNS requestors to advertise the size of their UDP packets and facilitate the transfer of packets larger than 512 octets, the original DNS restriction for UDP packet size (RFC 1035). When a DNS server receives a request over the UDP Transport Layer, it identifies the requestor's UDP packet size from the OPT resource record (RR) and scales its response to contain as many resource records as are allowed in the maximum UDP packet size specified by the requestor. For more information about resource records, see Resource records reference. For more information about RFCs, see DNS RFCs.

**EDNS0 UDP responses**
Before a DNS server assumes that the requestor supports EDNS0, the DNS server must receive a query containing an OPT resource record. An OPT record contains no actual DNS data and its contents relate to the UDP Transport Layer message only. The OPT record stores the sender's UDP payload size in its CLASS field and lists the number of octets in the largest UDP payload that the requestor can deliver in the requestor's network.

When the DNS server receives a query containing an OPT record advertising the maximum UDP packet size, it will truncate any UDP response's size larger than the limit specified in the OPT record.

By default, the DNS server includes OPT resource records indicating its UDP maximum in responses to queries containing OPT resource records. For information about modifying the OPT response setting, see Modify EDNS0 configuration.

If the DNS server receives a query that does not contain an OPT resource record, it assumes the requestor's server does not support EDNS0 and will respond to the requestor assuming that the sender does not accept UDP packets larger than 512 octets. In this case, the DNS server will truncate its UDP response size to a maximum of 512 octets.

**EDNS0 UDP queries**
Before the requesting DNS server sends a query, it checks its cache to identify if the responding DNS server supports EDNS0. If the responding DNS server supports EDNS0, the requesting DNS server attaches an OPT resource record to the additional section of the query it sends. (All queries have five sections: header, question, answer, authority, and additional.) If, according to the requesting DNS server's cache, the responding DNS server does not support EDNS0, the requesting DNS server will not attach an OPT resource record to the query before it is sent. For more information, see Modify UDP message size.

**Identifying and caching EDNS0 support**
When the DNS server receives a request or response from a host containing an OPT record, the DNS server caches the EDNS version supported by the host (such as EDNS0). If there is no OPT record in a request or response from a host, the DNS server's cache will indicate that the host does not support EDNS0. If the cache already indicates that the host does support ENDS0, then cache will not be changed.

The default value for how long a host's EDNS0 support information is cached is 86400 (one day, specified in seconds). This value can be modified in the **EDNSCacheTimeout** registry entry, in the following registry subkey:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Dns\Parameters**

The range of values allowed is from 3600 (1 hour) to 15724800 (182 days). For information about modifying the EDNS0 cache setting, see Modify EDNS0 configuration.

The DNS server decides that a host does not support EDNS0 when it requests an OPT resource record and receives a response containing one of the following RCODE values in the header:

| Name | Value | Description |
|---|---|---|
| FORMERR | 1 | Format Error. The name server did not interpret the OPT resource record. |
| SERVFAIL | 2 | Server Failure. The name server did not process the query because of a problem with the name server. |
| NOTIMPL | 4 | Not Implemented. The name server does not support the kind of query requested. |

(The RCODE field, or response code field, is a 4-bit field set in the header section as part of responses.) In this situation (as a requester), the DNS server identifies that the server does not support EDNS0 and caches this information.

**Caution**

- When considering packet sizes, you should take account of the network transmission path's discovered Maximum Transmission Unit (MTU), if this information is available. When configuring the UDP packet size to be larger than 512 octets, remember that the UDP packets must travel through devices other than UDP hosts and these devices, such as routers, may not support UDP packets larger than 512 octets. The maximum UDP packet size should always be compared with the MTU, which in some cases may be smaller. It is recommended that you establish the maximum UDP packet length support for all devices, and configure your UDP hosts for this maximum.

**Note**

- EDNS0 refers to the first RFC for EDNS (RFC 2671). Future RFCs for EDNS will be identified by greater numbers, such as EDNS1.

# DNS Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This section covers:

- Resource records reference
- DNS server log reference
- DNS-related files
- DNS updated technical information
- DNS RFCs
- Top-level domains
- DNS bibliography

# Resource records reference

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource records reference**

A DNS database consists of one or more zone files used by the DNS server. Each zone holds a collection of structured resource records, the following of which are supported by the DNS Server service.

**Format of DNS resource records**

All resource records have a defined format that uses the same top-level fields, as described in the following table.

| Field | Description |
|---|---|
| Owner | Indicates the DNS domain name that owns a resource record. This name is the same as that of the console tree node where a resource record is located. |
| Time to Live (TTL) | For most resource records, this field is optional. It indicates a length of time used by other DNS servers to determine how long to cache information for a record before expiring and discarding it. For example, most resource records created by the DNS Server service inherit the minimum (default) TTL of 1 hour from the start of authority (SOA) resource record which prevents overlong caching by other DNS servers. For an individual resource record, you can specify a record-specific TTL that overrides the minimum (default) TTL inherited from the SOA RR. TTL values of zero (0) can also be used for resource records that contain volatile data not to be cached for later use after the current DNS query in progress is completed.<br>**Note**<br>• In the resource record examples below, the TTL field is omitted wherever it is optional. The TTL field is included in the syntax for each record to indicate where it may be added. |
| Class | Contains standard mnemonic text indicating the class of the resource record. For example, a setting of "IN" indicates that the resource record belongs to the Internet class, which is the only class supported by Windows Server 2003 DNS. This field is required. |
| Type | Contains standard mnemonic text indicating the type of resource record. For example, a mnemonic of "A" indicates that the resource record stores host address information. This field is required. |
| Record-specific data | A required, variable-length field that contains information describing the resource. The format of this information varies according to the type and class of the resource record. |

**A**

**Description:** Host address (A) resource record. Maps a DNS domain name to an Internet Protocol (IP) version 4 32-bit address. For more information, see RFC 1035.

**Syntax:**

*owner class ttl* **A***IP_v4_address*

**Example:**

```
host1.example.microsoft.com.     IN  A  127.0.0.1
```

**AAAA**

**Description:** IPv6 host address (AAAA) resource record. Maps a DNS domain name to an Internet Protocol (IP) version 6 128-bit address. For more information, see RFC 1886.

**Syntax:**

*owner class ttl* **AAAA***IP_v6_address*

**Example:**

```
ipv6_host1.example.microsoft.com.     IN  AAAA  4321:0:1:2:3:4:567:89ab
```

**AFSDB**

**Description:** Andrew File System Database (AFSDB) resource record. Maps a DNS domain name in the *server_host_name* field to the host name for a server computer of a server subtype. The *subtype* field can have either of the following recognized and supported values:

- A value of **1** to indicate that the server is an AFS version 3.0 volume location server for the named AFS cell.
- A value of **2** to indicate that the server is an authenticated name server holding the cell-root directory node for the server that uses either Open Software Foundation's (OSF) DCE authenticated cell-naming system or HP/Apollo's Network Computing Architecture (NCA).

For more information, see RFC 1183.

**Syntax:**

*owner ttl class* **AFSDB***subtype server_host_name*

**Example:**

```
example.microsoft.com.   AFSDB   1 afs-server1.example.microsoft.com
```

**ATMA**

**Description:** Asynchronous Transfer Mode address (ATMA) resource record. Maps a DNS domain name in the *owner* field to an ATM address referenced in the *atm_address* field.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

For more information, see "ATM Names Service" at the MFA Forum site. At the MFA Forum site, download **af-saa-0069.000**.

**Syntax:**

*owner ttl class* **ATMA***atm_address*

**Example:**

```
atm-host   ATMA   47.0079.00010200000000000000.00a03e000002.00
```

**CNAME**

**Description:** Canonical name (CNAME) resource record. Maps an aliased or alternate DNS domain name in the *owner* field to a canonical or primary DNS domain name specified in the *canonical_name* field. The canonical or primary DNS domain name used in the data is required and must resolve to a valid DNS domain name in the namespace.

**Syntax:**

*owner ttl class* **CNAME** *canonical_name*

**Example:**

```
aliasname.example.microsoft.com.   CNAME   truename.example.microsoft.com
```

**HINFO**

**Description:** Host information (HINFO) resource record. Specifies the type of CPU and operating system in the *cpu_type* and *os_type* fields, respectively, for the host DNS domain name in the *owner* field. Well-known CPU and operating system types that are most often used are noted in RFC 1700. This information can be used by application protocols such as FTP, which use special procedures when communicating with computers of a known CPU and operating system type.

**Syntax:**

*owner ttl class* **HINFO** *cpu_type os_type*

**Example:**

```
my-computer-name.example.microsoft.com.    HINFO   INTEL-386  WIN32
```

**ISDN**

**Description:** Integrated Services Digital Network (ISDN) resource record. Maps a DNS domain name to an ISDN telephone number. Telephone numbers used with this record should follow ITU-T E.163/E.164 international telephone numbering standards, which are compatible with current international phone numbering plans already in use. For more information, see RFC 1183.

**Syntax:**

*owner ttl class* **ISDN** *isdn_address sub_address*

**Example:**

```
my-isdn-host.example.microsoft.com.    ISDN   141555555539699 002
```

**KEY**

**Description:** Public key resource record. Contains a public key that is associated with a zone. In full DNSSEC implementation, resolvers and servers use KEY resource records to authenticate SIG resource records received from a signed zones. KEY resource records are signed by the parent zone, allowing a server that knows a parent zone's public key to discover and verify the child zone's key. Name servers or resolvers receiving resource records from a signed zone obtain the corresponding SIG record then retrieve the zone's KEY record. For more information, see RFC 2535.

**Syntax:**

*owner class* **KEY** *protocol digital_signature_algorithm (DSA) public_key*

**Example:**

```
widgets.microsoft.com IN KEY 0x0000 3 0
```

**MB**

**Description:** Mailbox (MB) resource record. Maps a specified domain mailbox name in the *owner* field to a mailbox host name in *mailbox_hostname*. The mailbox host name must be the same as a valid host address (A) resource record already used by a host in the same zone. Additionally, the specified host must have a domain mailbox that accepts mail for the specified owner. For more information, see RFC 1035.

**Syntax:**

*owner ttl class* **MB** *mailbox_hostname*

**Example:**

```
mailbox.example.microsoft.com.    MB   mailhost1.example.microsoft.com
```

**MG**

**Description:** Mail group (MG) resource record. Used to add domain mailboxes, each specified by a mailbox (MB) resource record in the current zone, to the domain mailing group identified by *owner* in this resource record. Names used in the *mailbox_name* field must be identical to valid mailbox (MB) resource records already present in the current zone. For more information, see RFC 1035.

**Syntax:**

*owner ttl class* **MG** *mailbox_name*

**Example:**

```
administrator.example.microsoft.com.    MG   mailbox1.example.microsoft.com
                                         mailbox2.example.microsoft.com
```

**MINFO**

**Description:** Mailbox mail list information (MINFO) resource record. Specifies (in *responsible_mailbox*) a domain mailbox name for a responsible person who maintains a mailing list or mailbox specified in the *owner* field. The *error_mailbox* field can also be used to specify a domain mailbox that receives error messages related to this mailing list or mailbox. Mailboxes specified for responsible contacts and error forwarding must be the same as valid mailbox (MB) records that already exist in the current zone. For more information, see RFC 1035.

**Syntax:**

*owner ttl class* **MINFO** *responsible_mailbox error_mailbox*

**Example:**

```
administrator.example.microsoft.com.    MINFO  resp-mbox.example.microsoft.com err-
mbox.example.microsoft.com
```

**MR**

**Description:** Mailbox renamed (MR) resource record. Specifies a domain mailbox name in *new_renamed_mailbox*, the proper rename of an existing mailbox specified in the *owner* field. An MR resource record is often used as a forwarding entry for a user who has moved to a different mailbox. MR records do not cause additional section processing. For more information, see RFC 1035.

**Syntax:**

*owner ttl class* **MR** *new_renamed_mailbox*

**Example:**

```
old-mbox.example.microsoft.com.    MR  new-renamed-mbox.example.microsoft.com
```

**MX**

**Description:** Mail exchanger (MX) resource record. Provides message routing to a mail exchanger host, as specified in *mail_exchanger_host*, for mail sent to the domain name specified in the **owner** field. A 2-digit preference value indicates preferred ordering if multiple exchanger hosts are specified. Each exchanger host must have a corresponding host (A) address resource record in a valid zone. For more information, see RFC 1035.

**Syntax:**

*owner ttl class* **MX** *preference mail_exchanger_host*

**Example:**

```
example.microsoft.com.    MX  10 mailserver1.example.microsoft.com
```

**NS**

**Description:** Used to map a DNS domain name as specified in **owner** to the name of hosts operating DNS servers specified in the *name_server_domain_name* field.

**Syntax:**

*owner ttl IN* **NS** *name_server_domain_name*

**Example:**

```
example.microsoft.com.    IN NS  nameserver1.example.microsoft.com
```

**NXT**

**Description:** Next resource record. NXT resource records indicate the nonexistence of a name in a zone by creating a chain of all of the literal owner names in that zone. They also indicate what resource record types are present for an existing name. For more information, see RFC 2535.

**Syntax:**

*owner class* **NXT** *next_domain_name last_record_type* **NXT**

**Example:**

```
east.widgets.microsoft.com. IN NXT ftp.widgets.microsoft.com. A NXT
```

**OPT**

**Description:** Option resource record. One OPT resource record can be added to the additional data section of either a DNS request or response. An OPT resource record belongs to a particular transport level message, such as UDP, and not to actual DNS data. Only one OPT resource record is allowed, but not required, per message. For more information, see RFC 2671.

**Syntax:**

*name***OPT** *class ttl rdlen rdata*

**Example:**

```
null  OPT  IN  1280  0  0
```

**PTR**

**Description:**Pointer (PTR) resource record. Points from the name in **owner** to another location in the DNS namespace as specified by *targeted_domain_name*. Often used in special domains such as the in-addr.arpa domain tree to provide reverse lookups of address-to-name mappings. In most cases, each record provides information that points to another DNS domain name location, such as a corresponding host (A) address resource record in a forward lookup zone. For more information, see RFC 1035.

**Syntax:**

*owner ttl class* **PTR** *targeted_domain_name*

**Example:**

```
1.0.0.10.in-addr.arpa.    PTR  host.example.microsoft.com
```

**RP**

**Description:** Responsible Person (RP) resource record. Specifies the domain mailbox name for a responsible person in *mailbox_name*. This name is then mapped to a domain name in *text_record_name* for which (TXT) resource records exist in the same zone. When RP records are used in DNS queries, subsequent queries are used to retrieve associated text (TXT) resource record information. For more information, see RFC 1183.

**Syntax:**

*owner ttl class* **RP** *mailbox_name text_record_name*

**Example:**

```
example.microsoft.com.                RP    admin.example.microsoft.com.  admin-info.example.microsoft.com.
admin-info.example.microsoft.com.     TXT  "Joe Administrator, (555) 555-0110"
```

**RT**

**Description:** Route through (RT) resource record. Provides an intermediate host binding for internal hosts that do not have a direct wide area network (WAN) or external network connection. The RT record is similar to the MX record in that any communication from an internal host is routed through the *intermediate_host* for the destination DNS domain name specified in the *owner* field. A 2-digit preference value is used to set priority if multiple intermediate routing hosts are specified. For each intermediate host specified, a corresponding host (A) address resource record is needed in the current zone. For more information, see RFC 1183.

**Syntax:**

*owner ttl class* **RT** *preference intermediate_host*

**Example:**

```
example.microsoft.com.    RT    2 lan-router.example.microsoft.com
                          RT   10 wan-router.example.microsoft.com
```

**SIG**

**Description:** Signature resource record. Encrypts a RRset to a signer's (RRset's zone owner) domain name and a validity interval. For more information, see RFC 2535.

**Syntax:**

*owner class* **SIG** *ttl signature_expiration signature_inception key_identifier signer_name***{***digital_signature***}**

**Example:**

```
widgets.microsoft.com IN SIG 86400 19700101000000 19700101000000 49292
widgets.microsoft.com.{
AIYADP8d3zYNyQwW2EM4wXVFdslEJcUx/fxkfBeH1El4ixPFhpfHFElxbvKoWmvjDTCmfiYy2X+8XpFjwICHc398kzWsTMKlxovpz2FnCTM=
}
```

**SOA**

**Description:** Start of authority (SOA) resource record. Indicates the name of origin for the zone and contains the name of the server that is the primary source for information about the zone. It also indicates other basic properties of the zone. The SOA resource record is always first in any standard zone. It indicates the DNS server that either originally created it or is now the primary server for the zone. It is also used to store other properties such as version information and timings that affect zone renewal or expiration. These properties affect how often transfers of the zone are done between servers authoritative for the zone. For more information, see Managing authority records.

**Note**

- In the example below, the owner (primary DNS server) is specified as "@" because the domain name is the same as the origin of all data in the zone (example.microsoft.com.). This is a standard notation convention for resource records and is most often seen in the SOA record.

**Syntax:**

*owner class* **SOA** *name_server responsible_person (serial_number refresh_interval retry_interval expiration minimum_time_to_live)*

**Example:**

```
@   IN  SOA     nameserver.example.microsoft.com.  postmaster.example.microsoft.com. (
                                1            ; serial number
                                3600         ; refresh   [1h]
                                600          ; retry     [10m]
                                86400        ; expire    [1d]
                                3600 )       ; min TTL   [1h]
```

**SRV**

**Description:** Service locator (SRV) resource record. Allows multiple servers providing a similar TCP/IP-based service to be located using a single DNS query operation. This record enables you to maintain a list of servers for a well-known server port and transport protocol type ordered by preference for a DNS domain name. For example, in Windows Server 2003 DNS, it provides the means to locate domain controllers that use Lightweight Directory Access Protocol (LDAP) service over TCP port 389.

The purposes of each of the specialized fields used in an SRV resource record are as follows:

- **service** A symbolic name for the desired service. For well-known services, a reserved universal symbolic name such as "_telnet" or "_smtp" is defined in RFC 1700. If a well-known service name is not defined in RFC 1700, a local or user-preferred name can be used instead. Some widely used TCP/IP services, notably the Post Office Protocol (POP), do not have a single universal symbolic name. If RFC 1700 assigns a name for a service indicated in this field, the RFC-defined name is the only name that is legal to use. Only locally defined services can be named locally.

- **protocol** Indicates the transport protocol type. Typically, this is either TCP or UDP, although any transport protocol named in RFC 1700 can be used.

- **name** The DNS domain name referred to by this resource record. The SRV resource record is unique among other DNS record types in that it is not used to perform the search or query.

- **priority** Sets the preference for a host specified in the *target* field. DNS clients that query for SRV resource records attempt to contact the first reachable host of the lowest numbered preference listed here. Although *target* hosts have the same stated preference value, they can be tried in random order. The range of preference values is 0 to 65535.

- **weight** Can be used in addition to preference to provide a load-balancing mechanism where multiple servers are specified in the **target** field and are all set to the same level of preference. When selecting a target server host among those of equal preference, this value can be used to set an added level of preference that can be used to determine the exact order or balancing of selection for the target hosts used in an answered SRV query. When a non-zero value is used, servers of equal preference are tried in proportion to the weight of this value. The range of values is 1 to 65535. If load balancing is not needed, use a value of **0** in this field to make the record easier to read.

- **port** The server port on the **target** host that provides the service indicated in the *service* field. The range of port numbers is 0 to 65535, although the number is often a well-known assigned service port number, as specified in RFC 1700. Unassigned ports can be used as needed.

- **target** Specifies the DNS domain name of the host that provides the type of service being requested. For each host name used, a corresponding host address (A) resource record in the DNS namespace is required. A single period (.) can be used in this field to indicate authoritatively that the requested service specified in this SRV resource record is not available at this DNS domain name.

For more information, see the Internet draft "A DNS RR for specifying the location of services (DNS SRV)."

**Syntax:**

*service.protocol.name ttl class* **SRV** *preference weight port target*

**Example:**

```
_ldap._tcp._msdcs    SRV  0 0   389 dc1.example.microsoft.com
                   SRV 10 0   389 dc2.example.microsoft.com
```

**TXT**

**Description:** Text (TXT) resource record. Maps a DNS domain name specified in the *owner* field to a string of characters in *text_string* serving as descriptive text. For more information, see RFC 1035.

**Syntax:**

*owner ttl class* **TXT** *text_string*

**Example:**

```
example.microsoft.com.    TXT  "This is an example of additional domain name information."
```

**WKS**

**Description:** Well-known service (WKS) resource record. Describes the well-known TCP/IP services supported by a particular protocol on a specific IP address. WKS records provide TCP and UDP availability information for TCP/IP servers. If a server either supports both TCP and UDP for a well-known service or has multiple IP addresses that support a service, then multiple WKS records are used. For more information, see RFC 1035.

**Syntax:**

*owner ttl class* **WKS** *address protocol service_list*

**Example:**

```
example.microsoft.com.    WKS  10.0.0.1 TCP ( telnet smtp ftp )
```

**X25**

**Description:** X.25 (X25) resource record. Maps a DNS domain name in the **owner** field to a Public Switched Data Network (PSDN) address number specified in *psdn_number*. PSDN numbers used with this record should follow the X.121 international numbering plan. For more information, see RFC 1183.

**Syntax:**

*owner ttl class* **X25** *psdn_number*

**Example:**

```
example.microsoft.com.    X25  52204455506
```

# DNS server log reference

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS server log reference**

The following is a partial list of events recorded in the DNS server log. You can view this log by using the DNS console after the DNS Server service is installed on the server computer. For more information, see View the DNS server system event log.

| Event ID | Description |
|---|---|
| 2 | The DNS server has started.<br><br>This message generally appears at startup when either the server computer is started or the DNS Server service is manually started. For more information, see Start or stop a DNS server. |
| 3 | The DNS server has shut down.<br><br>This message generally appears when either the server computer is shut down or the DNS Server service is stopped manually. |
| 408 | The DNS server could not open socket for address *[IPaddress]*. Verify that this is a valid IP address for the server computer.<br><br>To correct the problem, you can do the following:<br><br>• If the specified IP address is not valid, remove it from the list of restricted interfaces for the server and restart the server. For more information, see Restrict a DNS server to listen only on selected addresses.<br><br>• If the specified IP address is no longer valid and was the only address enabled for the DNS server to use, the server might not have started as a result of this configuration error. To correct this problem, delete the following value from the registry and restart the DNS server:<br><br>HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DNS\Parameters\ListenAddress<br><br>• If this is a valid IP address for the server computer, verify that no other application that would attempt to use the same DNS server port (such as another DNS server application) is running. By default, DNS uses TCP port 53. |
| 413 | The DNS server will send requests to other DNS servers on a port other than its default port (TCP port 53).<br><br>This DNS server is multihomed and has been configured to restrict DNS Server service to only some of its configured IP addresses. For this reason, there is no assurance that DNS queries made by this server to other remote DNS servers will be sent using one of the IP addresses enabled for the DNS server.<br><br>This might prevent query answer responses returned by these servers from being received on the DNS port that the server is currently configured to use. To avoid this problem, The DNS server sends queries to other DNS servers using an arbitrary non-DNS port, and the response is received regardless of the IP address used.<br><br>If you want to limit the DNS server to using only its configured DNS port for sending queries to other DNS servers, use the DNS console to perform one of the following changes in server properties configuration on the **Interfaces** tab:<br><br>1. Either select **All IP addresses** to enable the DNS server to listen on all configured server IP addresses<br><br>2. Or, if you continue to select and use **Only the following IP addresses**, limit the **IP address** list to a single server IP address.<br><br>For more information, see Restrict a DNS server to listen only on selected addresses. |

| | |
|---|---|
| 414 | The server computer currently has no primary DNS suffix configured. Its DNS name is currently a single label host name. For example, its currently configured name is "host" rather than "host.example.microsoft.com" or another fully qualified name.<br><br>While the DNS server has only a single label name, default resource records created for its configured zones use only this single label name when mapping the host name for this DNS server. This can lead to incorrect and failed referrals when clients and other DNS servers use these records to locate this server by name.<br><br>In general, the DNS server should be reconfigured with a full DNS computer name appropriate for its domain or workgroup use on your network. |
| 708 | The DNS server did not detect any zones of either primary or secondary type. It will run as a caching-only server but will not be authoritative for any zones.<br><br>For more information, see Using caching-only servers. |
| 3150 | The DNS server wrote a new version of zone *[zonename]* to file *[filename]*. You can view the new version number by clicking the **Record Data** tab.<br><br>This event should only appear if the DNS server is configured to operate as a root server. For more information, see DNS-related files. |
| 6527 | Zone *[zonename]* expired before it could obtain a successful zone transfer or update from a master server acting as its source for the zone. The zone has been shut down.<br><br>This event ID might appear when the DNS server is configured to host a secondary copy of the zone from another DNS server acting as its source or master server. Verify that this server has network connectivity to its configured master server.<br><br>If the problem continues, consider one or more of the following options:<br><br>1. Delete the zone and recreate it, specifying either a different master server, or an updated and corrected IP address for the same master server.<br><br>For more information, see Add and Remove Zones.<br><br>2. If zone expiration continues, consider adjusting the expire interval.<br><br>For more information, see Adjust the expire interval for a zone.<br><br>For more information, see Understanding zones and zone transfer. |

For more information about all other messages not described in this topic, refer directly to the message text itself as it appears in the DNS server event log.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Notes**

- Event messages logged by the DNS Client service are logged in System log, and can also be viewed using Event Viewer at the applicable client computer. For more information, see Windows interface administrative tool reference A-Z: Event Viewer.

- If needed to troubleshoot persistent problems, you can also enable for use advanced debugging options at the DNS server for obtaining a more detailed log of server activity. For more information, see Using server debug logging options or View a DNS server debug log file.

# DNS-related files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS-related files**

The following files relate to using and configuring DNS servers and clients.

| File | Description |
|---|---|
| Boot | BIND boot configuration file. This file is not created by the DNS console. However, as an optional configuration for the DNS Server service, it can be copied from another DNS server running the Berkeley Internet Name Domain (BIND) server implementation of DNS. To use this file with the DNS Server service, you need to click **From file** in Server properties. On BIND servers, this file is often called the "named.boot" file. |
| Cache.dns | Used to preload resource records into the DNS server names cache. DNS servers use this file to help locate root servers on either your network or the Internet.<br><br>By default, this file contains DNS resource records that prime the local cache of the server with the addresses of authoritative root servers for the Internet. If you are setting up a DNS server to resolve Internet DNS names, the information in this file is required unless you enable the use of another DNS server as a forwarder to resolve these names.<br><br>Traffic to the Internet root servers is heavy, but because host names are not usually resolved at this level, the load can be reasonably handled. Instead, the root hints file provides referral information that can be useful during DNS name resolution to redirect a query to other servers that are authoritative for names located beneath the root.<br><br>For DNS servers operating privately on your internal network, the DNS console can learn and replace the contents of this file with internal root servers on your network, provided they are reachable through the network when you are setting up and configuring new DNS servers. It can be updated using the DNS console from the **Root Hints** tab located under the applicable server properties.<br><br>This file preloads the server names cache when it is started. |
| Root.dns | Root zone file. This file can appear at a DNS server if it is configured as a root server for your network. |
| *zone_name*.dns | Used when a standard zone (either primary or secondary) is added and configured for the server. Files of this type are not created or used for primary type zones that are directory-integrated, which are stored in the Active Directory database. |

These files can be found in the systemroot\System32\Dns folder on the server computer.

**Notes**

- For more information about operating a DNS server as a root server, see Updating root hints.
- For more information on hosting your own DNS presence on the Internet, see Interoperability issues.
- For more information about configuring DNS servers on your network to use forwarders, see Using forwarders.
- If you are migrating server files from a BIND server to use with a Windows DNS server, you need to rename these files and copy them to the folder used by the DNS Server service. For more information, see Migrating servers.

# DNS updated technical information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS updated technical information**

Updated technical information for DNS can be obtained by searching technical articles available from Product Support Services at the Microsoft Web site. For more information, see Updated technical information.

# DNS RFCs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS RFCs**

Request for Comments (RFCs) are an evolving series of reports, proposals for protocols, and protocol standards used by the Internet community. Domain Name System (DNS) specifications are based on approved RFCs published by the Internet Engineering Task Force (IETF) and other working groups.

**RFCs for the DNS Server service**

The following RFCs contain specifications used to design and implement the DNS Server and Client services:

| RFC | Title |
|---|---|
| 1034 | Domain Names -- Concepts and Facilities |
| 1035 | Domain Names -- Implementation and Specification |
| 1123 | Requirements for Internet Hosts -- Application and Support |
| 1886 | DNS Extensions to Support IP Version 6 |
| 1995 | Incremental Zone Transfer in DNS |
| 1996 | A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY) |
| 2136 | Dynamic Updates in the Domain Name System (DNS UPDATE) |
| 2181 | Clarifications to the DNS Specification |
| 2308 | Negative Caching of DNS Queries (DNS NCACHE) |
| 2535 | Domain Name System Security Extensions (DNSSEC) |
| 2671 | Extension Mechanisms for DNS (EDNS0) |
| 2782 | A DNS RR for specifying the location of services (DNS SRV) |
| 2930 | Secret Key Establishment for DNS (TKEY RR) |
| 3645 | Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG) |
| 3646 | DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6) |

**Other specifications for DNS**

The following additional specifications are used to design and implement the DNS Server and Client services:

| File name | Title |
|---|---|
| af-saa-0069.000.pdf | ATM Names Service<br><br>This specification is published by the MFA Forum. For more information, you can obtain this specification by downloading it from the MFA Forum site. |

**Obtaining DNS RFCs**

You can obtain RFCs from the RFC Editor Web site. RFCs are classified as one of the following: approved Internet standards, proposed Internet standards (circulated in draft form for review), Internet best practices, or For Your Information (FYI) documents.

## Notes

- RFCs 1034 and 1035 define the original DNS standard protocol to support domain name services in a TCP/IP environment. These RFCs describe the protocols in a detailed manner, emphasizing the underlying ideas and techniques used in all DNS implementations.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Top-level domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Top-level domains**

The following is a list of the eight top-level domains most often used on the Internet to classify organizations by type when they register a second-level domain name for use. For example, microsoft.com (the second-level domain name registered to Microsoft) is registered in the "com" domain, because this is the top-level domain provided for organizations that do business (commerce) on the Internet.

| Top-level name | Description | Used by |
|---|---|---|
| arpa | Owned by Advanced Research Project Agency (ARPA). Used to register reverse mapping of Internet Protocol (IP) version 4 (IPv4) and IP version 6 (IPv6) addresses assigned by the Internet Assigned Number Authority (IANA) to Domain Name System (DNS) domain names for computers that use those addresses on the Internet. | The in-addr.arpa domain, the ip6.arpa domain |
| com | For business and commercial use. | Businesses and corporations |
| edu | For educational use. | Public and private schools, colleges, and universities |
| gov | For use by governmental institutions. | Local, state, and federal government agencies |
| int | Reserved for international use. Originally planned for use in RFC 1886 to register reverse mapping of IPv6 addresses. For more information about RFCs, see DNS RFCs. | |
| mil | For use by military agencies. | Department of Defense (DoD), U.S. Navy, U.S. Army, U.S. Air Force, and other military agencies |
| net | For use by organizations that provide large-scale Internet or telephony-based service. | InterNIC, AT&T, and other large-scale Internet and telephone service providers |
| org | For use by noncommercial, nonprofit organizations. | Churches and charitable institutions |

**Note**

- In addition to these top-level domains which are all three letters in length, the use of two-letter country/region codes (such as "uk" for United Kingdom) are reserved for use. These country/region codes can also be used in conjunction with the three-letter codes listed above to further specify the country or region of origin for organizations that register names in the Internet DNS namespace.

# DNS bibliography

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNS bibliography**

For more information about Domain Name System (DNS) or other DNS server implementations, you can consult the following references:


- *Microsoft Windows Server 2003 Deployment Kit*. Redmond, Washington: Microsoft Corporation.
- *Microsoft Windows 2000 Server Resource Kit*. Redmond, Washington: Microsoft Corporation.
- Lee, T., Davies, J. 2000. *Microsoft Windows 2000 TCP/IP Protocols and Services Technical Reference*. Redmond, Washington: Microsoft Press.
- Larson, M. and C. Liu. 2001. *DNS on Windows 2000*. Sebastopol, California: O'Reilly and Associates, Inc.
- Wong, W. 2000. *Windows 2000 DNS Server*. Berkeley, California: Osborne/McGraw-Hill.
- Albitz, P., Loukides, M. and C. Liu. 1998. *DNS and BIND, Fourth Edition*. Sebastopol, California: O'Reilly and Associates, Inc.
- Comer, D. 1995. *Internetworking with TCP/IP, Vol. 1, Third Edition*. Englewood Cliffs, New Jersey: Prentice Hall.

# Troubleshooting DNS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

This section covers:

- Troubleshooting DNS clients
- Troubleshooting DNS servers
- Troubleshooting dynamic updates
- Troubleshooting zone problems

# Troubleshooting DNS clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting DNS clients**
What problem are you having?

- The DNS client received a "Name not found" error message.
- The DNS client appears to have received a response with stale or incorrect information in it.
- The DNS client appears to be affected by another problem not described above.

**The DNS client received a "Name not found" error message.**
**Cause:** The DNS client computer does not have a valid IP configuration for the network.

**Solution:** Verify that TCP/IP configuration settings for the client computer are correct, particularly those used for DNS name resolution.

To verify a client IP configuration, use the **ipconfig** command. In the command output, verify that the client has a valid IP address, subnet mask, and default gateway for the network where it is attached and being used.

If the client does not have a valid TCP/IP configuration, you can either:

1. For dynamically configured clients, use the **ipconfig /renew** command to manually force the client to renew its IP address configuration with the DHCP server.
2. For statically configured clients, modify the client TCP/IP properties to use valid configuration settings or complete its DNS configuration for the network.

**See also:** Ipconfig; Configure TCP/IP for dynamic addressing; Configure TCP/IP for static addressing; Configure TCP/IP to use DNS; Configuring DNS client settings.

**Cause:** The client was not able to contact a DNS server because of a network or hardware related failure.

**Solution:** Verify that the client computer has a valid and functioning network connection. First, check that related client hardware (cables and network adapters) are working properly at the client using basic network and hardware troubleshooting steps.

If the client hardware appears to be prepared and functioning properly, verify that it can ping other computers on the same network.

**See also:** Test a TCP/IP configuration by using the ping command.

**Cause:** The DNS client cannot contact its configured DNS servers.

**Solution:** If the DNS client has basic connectivity to the network, verify that it can contact a preferred (or alternate) DNS server.

To verify whether a client has basic TCP/IP access to the DNS server, first try pinging the preferred DNS server by its IP address.

For example, if the client uses a preferred DNS server of 10.0.0.1, type **ping 10.0.0.1** at the command prompt on the client computer. If you are not sure what the IP address is for the preferred DNS server, you can observe it by using the **ipconfig** command.

For example, at the client computer, type **ipconfig /all|more** if necessary to pause the display so you can read and note any IP addresses listed in **DNS servers** for the command output.

If no configured DNS servers respond to a direct pinging of their IP address, it indicates that the source of the problem is more likely a network connectivity problem between the client and the DNS servers. If that is the case, follow basic TCP/IP network

troubleshooting steps to fix the problem.

**See also:** Test a TCP/IP configuration by using the ping command.

**Cause:** The DNS server is not running or responding to queries.

**Solution:** If the DNS client can ping the DNS server computer, verify that the DNS server is started and able to listen for and respond to client requests. Try using the **nslookup** command to test whether the server can respond to DNS clients.

**See also:** Verify DNS server responsiveness using the nslookup command; Start or stop a DNS server.

**Cause:** The DNS server the client is using does not have authority for the failed name and cannot locate the authoritative server for this name.

**Solution:** Confirm whether the DNS domain name the client is trying to resolve is one for which its configured DNS servers are authoritative.

For example, if the client is attempting to resolve the name host.example.microsoft.com, verify that the preferred (or an alternate, if one is being used) DNS server queried by the client loads the authoritative zone where a host (A) resource record (RR) for the failed name should exist.

If the preferred server is authoritative for the failed name and loads the applicable zone, determine whether the zone is missing the appropriate RRs. If needed, add the RRs to the zone.

If the preferred server is not authoritative for the failed name, it indicates that configuration errors at the DNS server are the likely cause. As needed, further troubleshoot the problem at the DNS server.

**See also:** Manage Resource Records; Troubleshooting DNS servers.

**The DNS client appears to have received a response with stale or incorrect information in it.**

**Cause:** The DNS server the client is using does not have authority for the failed name and is using stale information from its local DNS database.

**Solution:** Determine whether the DNS server is authoritative for the name and proceed accordingly.

For example, if the client is attempting to resolve the name host.example.microsoft.com, verify that the preferred (or an alternate, if one is being used) DNS server queried by the client loads the authoritative zone where a host (A) resource record (RR) for the failed name should exist.

If the preferred server is authoritative for the name and answered using incorrect data, it indicates that the applicable zone might have outdated or stale information in the applicable RR data. If that is the case, you can add and remove the appropriate RRs in the zone.

Another option where dynamic updates are enabled is to force registration and update at the computer targeted by the query. You can force it to update the registration of its resource records by typing the **ipconfig /registerdns** command at a command prompt.

If the preferred server is not a direct authority for the queried name, it likely answered the query based on information obtained and cached during an earlier recursive lookup. In this case, you might consider clearing the server names cache. This compels the server to use new recursive queries for this RR data and rebuild its cache contents based on current information.

**See also:** Manage Resource Records; Troubleshooting DNS servers; Renew DNS client registration using the ipconfig command; Clear the server names cache.

**Cause:** The preferred DNS server is a secondary server for the zone containing the targeted name and has outdated information.

**Solution:** If the server that answered the client is a secondary server for the zone, the version of the zone in use at that server might be stale and need to be updated more often.

As an immediate solution, you can initiate a zone transfer at the secondary server to its master server to update the zone. You might also consider using any of the following to improve the freshness of secondary zone data in the future:

1. Specify additional master servers for the secondary server to use when refreshing the zone.

2. Adjust the refresh interval on the zone slightly to decrease the length of time that all authoritative servers for the zone can use the zone before they are required to refresh it.

3. Configure a notify list at a master server that acts as the zone source for the secondary server and enable it to notify this server when the zone changes.

**See also:** Initiate a zone transfer at a secondary server; Adjust the refresh interval for a zone; Create and manage a notify list for a zone; Using secondary servers; Understanding zones and zone transfer.

**Cause:** The name queried was specified in error, either through user input or in a stored client configuration.

**Solution:** Verify that the name was correctly specified in the application where the name query originated.

In most cases, incorrect data in a positive query response indicates one of three possibilities:

- An incorrect DNS name was entered at the client by a user.
- A short unqualified name was used at the client and completed by the local resolver using an unintended DNS suffix.
- Resource records specified in the query were not updated correctly at the DNS server.

Confirm that the name was not entered in error by the user. Verify the exact set of characters entered by the user when the original DNS query was made or check in application settings, such as for any Internet mail or Web browser configurations that were made.

If the name used in the initial query was unqualified, and not the fully qualified domain name (FQDN), try using the FQDN instead in the client application and repeating the query. If you do, be sure to include the trailing dot (.) at the end of the name to indicate that the name entered is an exact FQDN.

If the FQDN query succeeds and returns correct data in the response, the most likely cause of the problem is a misconfigured DNS domain suffix search list used in the client resolver settings.

If you are using DNS in an environment that does not support dynamic updates, or generally administer zone data manually, you might also want to verify that the RRs involved in answering the query were not entered incorrectly. View them to ensure the record data stored in the zone is correct or modify it accordingly.

**See also:** Configure TCP/IP to use DNS; Modify an existing resource record in a zone.

**Cause:** The primary zone might have missing or errored data

**Solution:** Verify that the primary server for the zone has complete and accurate data.

The most likely cause for a primary DNS server for a zone to have missing or incomplete data is because of a failed update request. It is possible that support for dynamic update has not been fully implemented or configured. To resolve the problem, review the DNS dynamic update protocol (RFC 2136) and any requirements it has for DNS servers and clients that use it.

For directory-integrated zones, it is also possible that the affected records for the errored query have been updated in Active Directory but not replicated to all DNS servers loading the zone. By default, all DNS servers that load zones from Active Directory poll it at a set interval (typically every 15 minutes), and update the zone for any incremental changes to it. In most cases, a DNS update takes no more than 20 minutes to replicate to all DNS servers used in an Active Directory domain environment using default replication settings and reliable high-speed links.

If you have specifically configured your zones to disable dynamic update, keep in mind that you need to manually add and update most types of resource records used in a zone. If this is the case, use the DNS console to view and update the affected records.

Another possibility for the errored data is in whether WINS lookup integration is enabled and used with the zone. If you are using WINS lookup with your zones, verify that WINS is not the source of the errored data.

**See also:** Troubleshooting dynamic updates; Manage Resource Records; Managing resource records; Verify WINS as the source for answering a DNS query.

**The DNS client appears to be affected by another problem not described above.**
**Cause:** My problem is not described above.

**Solution:** Search TechNet at the Microsoft Web site for the latest technical information that could relate to the problem. If necessary, you can obtain information and instructions that pertain to your problem or issue.

If you are connected to the Internet, the latest operating system updates are available at the Microsoft Web site.

To obtain the latest service pack updates for Windows NT Server, see the Microsoft Web site.

**See also:** DNS updated technical information; DNS; Using the Windows Deployment and Resource Kits.

# Troubleshooting DNS servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting DNS servers**
What problem are you having?

- The DNS server is not responding to clients.
- The DNS server does not resolve names correctly.
- The DNS server appears to be affected by a problem for reasons not described above.

**The DNS server is not responding to clients.**
**Cause:** The DNS server is affected by a network failure.

**Solution:** Verify that the server computer has a valid functioning network connection. First, check that related client hardware (cables and network adapters) are working properly at the client using basic network and hardware troubleshooting steps.

If the server hardware appears to be prepared and functioning properly, check that it has network connectivity by pinging other computers or routers (such as its default gateway) that are used and available on the same network as the affected DNS servers.

**See also:** Test a TCP/IP configuration by using the ping command.

**Cause:** The DNS server is reachable through basic network testing but is not responding to DNS queries from clients.

**Solution:** If the DNS client can ping the DNS server computer, verify that the DNS server is started and able to listen to and respond to client requests. Try using the **nslookup** command to test whether the server can respond to DNS clients.

**See also:** Verify DNS server responsiveness using the nslookup command; Start or stop a DNS server.

**Cause:** The DNS server has been configured to limit service to a specific list of its configured IP addresses. The IP address originally used in testing its responsiveness is not included in this list.

**Solution:** If the server was previously configured to restrict the IP addresses for which it responds to queries, it is possible that the IP address being used by clients to contact it is not in the list of restricted IP addresses permitted to provide service to clients.

Try testing the server for a response again, but specify a different IP address known to be in the restricted interfaces list for the server. If the DNS server responds for that address, add the missing server IP address to the list.

**See also:** Verify DNS server responsiveness using the nslookup command; Restrict a DNS server to listen only on selected addresses.

**Cause:** The DNS server has been configured to disable the use of its automatically created default reverse lookup zones.

**Solution:** Verify that automatically created reverse lookup zones have been created for the server or that advanced configuration changes have not been previously made to the server.

By default, DNS servers automatically create the following three standard reverse lookup zones based on Request for Comments (RFC) recommendations:

These zones are created with common IP addresses covered by these zones that are not useful in a reverse lookup search (0.0.0.0, 127.0.0.1, and 255.255.255.255). By being authoritative for the zones corresponding to these addresses, the DNS service avoids unnecessary recursion to root servers in order to perform reverse lookups on these types of IP addresses.

It is possible, although unlikely, that these automatic zones are not created. This is because disabling the creation of these zones involves advanced manual configuration of the server registry by a user.

To verify that these zones have been created, do the following:

1. Open the DNS console.

2. From the **View** menu, click **Advanced**.

3. In the console tree, click **Reverse Lookup Zones**.

   **Where?**

   - **DNS**/*applicable DNS server*/**Reverse Lookup Zones**

4. In the details pane, verify that the following reverse lookup zones are present:

   - 0.in-addr.arpa
   - 127.in-addr.arpa
   - 255.in-addr.arpa

**See also:** Open the DNS console; DNS RFCs.

**Cause:** The DNS server is configured to use a non-default service port, such as in an advanced security or firewall configuration.

**Solution:** Verify that the DNS server is not using a non-standard configuration.

This is a rare but possible cause. By default, the **nslookup** command sends queries to targeted DNS servers using User Datagram Protocol (UDP) port 53. If the DNS server is located on another network only reachable through an intermediate host (such as a packet-filtering router or proxy server), the DNS server might use a non-standard port to listen for and receive client requests.

If this situation applies, determine whether any intermediate firewall or proxy server configuration is intentionally used to block traffic on well-known service ports used for DNS. If not, you might be able to add such a packet filter onto these configurations to permit traffic to standard DNS ports.

Also, check the DNS server event log to see if Event ID 414 or other critical service-related events have occurred which might indicate why the DNS server is not responding.

**See also:** DNS server log reference; View the DNS server system event log; Microsoft Windows Deployment and Resource Kits.

**The DNS server does not resolve names correctly.**
**Cause:** The DNS server provides incorrect data for queries it successfully answers.

**Solution:** Determine the cause of the incorrect data for the DNS server.

Some of the most likely causes include the following:

- Resource records (RRs) were not dynamically updated in a zone.
- An error was made when manually adding or modifying static resource records in the zone.
- Stale resource records in the DNS server database, left from cached lookups or zone records not updated with current information or removed when they are no longer needed.

To help prevent the most common types of problems, be sure to first review best practices for tips and suggestions on deploying and managing your DNS servers. Also, follow and use the checklists appropriate for installing and configuring DNS servers and clients based on your deployment needs.

If you are deploying DNS for Active Directory, note new directory integration features. These features can cause some differences for DNS server defaults when the DNS database is directory-integrated, that differ from those used with traditional

file-based storage.

Many DNS server problems start with failed queries at a client, so it is often good to start there and troubleshoot the DNS client first.

**See also:** DNS best practices; DNS Checklists; Troubleshooting DNS clients; Modify an existing resource record in a zone; Clear the server names cache; Modifying server defaults.

**Cause:** The DNS server does not resolve names for computers or services outside of your immediate network, such as those located on external networks or the Internet.

**Solution:** The server has a problem based on its ability to correctly perform recursion. Recursion is used in most DNS configurations to resolve names that are not located within the configured DNS domain name used by the DNS servers and clients.

If a DNS server fails to resolve a name for which it is not authoritative, the cause is usually a failed recursive query. Recursive queries are used frequently by DNS servers to resolve remote names delegated to other DNS zones and servers.

For recursion to work successfully, all DNS servers used in the path of a recursive query must be able to respond to and forward correct data. If not, a recursive query can fail for any of the following reasons:

- The recursive query times out before it can be completed.
- A remote DNS server fails to respond.
- A remote DNS server provides incorrect data.

If a server fails a recursive query for a remote name, review the following possible causes to troubleshoot the problem. If you do not understand recursion or the DNS query process, review conceptual topics in Help to better understand the issues involved.

**See also:** How DNS query works.

**Cause:** The DNS server is not configured to use other DNS servers to assist it in resolving queries.

**Solution:** Check whether the DNS server can use both forwarders and recursion.

By default, all DNS servers are enabled to use recursion, although the option to disable its use is configurable using the DNS console to modify advanced server options. The other possibility where recursion might be disabled is if the server is configured to use forwarders and recursion has been specifically disabled for that configuration.

**Note**

- If you disable recursion on the DNS server, you will not be able to use forwarders on the same server.

**See also:** Disable recursion on the DNS server; Configure a DNS server to use forwarders.

**Cause:** Current root hints for the DNS server are not valid.

**Solution:** Check whether server root hints are valid.

If configured and used correctly, root hints always should point to DNS servers authoritative for the zone containing the domain root and top-level domains.

By default, DNS servers are configured to use root hints appropriate to your deployment, based on the following available choices when using the DNS console to configure a server:

1. If the DNS server is installed as the first DNS server for your network, it is configured as a root server.

   For this configuration, root hints are disabled at the server because the server is authoritative for the root zone.

2. If the installed server is an additional DNS server for your network, you can direct the Configure DNS Server Wizard to update its root hints from an existing DNS server on the network.

3. If you do not have other DNS servers on your network but still need to resolve Internet DNS names, you can use the default root hints file which includes a list of Internet root servers authoritative for the Internet DNS namespace.

**See also:** Update root hints on the DNS server; Updating root hints.

**Cause:** The DNS server does not have network connectivity to the root servers.

**Solution:** Test for connectivity to the root servers.

If root hints appear to be configured correctly, verify that the DNS server used in a failed query can ping its root servers by IP address.

If a ping attempt to one root server fails, it might indicate that an IP address for that root server has changed. Reconfiguration of root servers, however, is uncommon.

A more likely cause is a full loss of network connectivity or in some cases, poor network performance on the intermediate network links between the DNS server and its configured root servers. Follow basic TCP/IP network troubleshooting steps to diagnose connections and determine whether this is the problem.

By default, the DNS service uses a recursive time-out of 15 seconds before failing a recursive query. Under normal network conditions, this time-out does not need to be changed. If performance warrants it, however, you can increase this value.

To review additional performance related information on DNS queries, you can enable and use the DNS server debug log file, Dns.log, which can provide extensive information about some types of service-related events.

**See also:** Test a TCP/IP configuration by using the ping command; Using server debug logging options; View a DNS server debug log file; Tuning advanced server parameters.

**Cause:** Other problems exist with updating DNS server data, such as an issue related to zones or dynamic updates.

**Solution:** Determine whether the problem is related to zones. As needed, Troubleshoot any issues in this area, such as possible failure of zone transfer.

**See also:** Troubleshooting dynamic updates; Troubleshooting zone problems.

**The DNS server appears to be affected by a problem for reasons not described above.**
**Cause:** My problem is not described above.

**Solution:** Search TechNet at the Microsoft Web site for the latest technical information that could relate to the problem. If necessary, you can obtain information and instructions that pertain to your problem or issue.

If you are connected to the Internet, the latest operating system updates are available at the Microsoft Web site.

To obtain the latest service pack updates for Windows NT Server, see the Microsoft Web site.

**See also:** DNS updated technical information; DNS; Using the Windows Deployment and Resource Kits.

# Troubleshooting dynamic updates

Published: January 21, 2005

Updated: August 25, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting dynamic updates**
What problem are you having?

- The DNS client is not performing dynamic updates.
- The DNS server is not performing dynamic updates.
- I am having a different problem related to dynamic updates than those described above.

**The DNS client is not performing dynamic updates.**
**Cause:** The client (or its DHCP server) do not support the use of the DNS dynamic update protocol.

**Solution:** Verify that your clients or servers support the DNS dynamic update protocol.

In order for client computers to be registered and updated dynamically with a DNS server, either:

1. Install or upgrade client computers to Windows XP or Windows Server 2003 .
2. Install and use the DHCP Server service on your network to lease client computers.

By default, computers attempt to register and perform dynamic update of their DNS names and IP addresses with a DNS server.

For other types of computers, you can deploy Windows Server 2003 DHCP servers, which can perform proxied registrations and updates as needed for non-dynamic clients.

**Notes**

- By default, the DNS client on Windows XP does not attempt dynamic update over a Remote Access Service (RAS) or virtual private network connection. To modify this configuration, you can modify the advanced TCP/IP settings of the particular network connection or modify the registry. For more information, see Configure TCP/IP to use DNS and the Microsoft Windows Resource Kits Web site.
- By default, the DNS client does not attempt dynamic update of top-level domain (TLD) zones. Any zone named with a single-label name is considered a TLD zone, for example, com, edu, blank, my-company. To configure the DNS client to allow the dynamic update of TLD zones, you can use the **Update Top Level Domain Zones** policy setting or modify the registry.

**Cause:** The client was not able to register with the DNS server because of intermittent problems with either the DNS server or the network.

**Solution:** At the client computer, use the **ipconfig** command as appropriate to retry registration or renewal and update client information with the DNS server.

You can use the **ipconfig /registerdns** command option to manually force a retry of its dynamic registration.

For computers running earlier versions of Windows, you can use the options of the **ipconfig** command to verify, view, or renew the client TCP/IP configuration details as appropriate.

For example, if the client computer obtains its IP address lease from a DHCP server, you might use the **ipconfig /renew** command to force it to renew its lease with the DHCP server. This action would then cause the DHCP server to proxy an update request to its configured DNS server on behalf of the client.

If the DHCP server succeeds in performing the proxied update with the DNS server, the result would be updated DNS host name and IP address information for the client computer in the DNS database.

**See also:** Renew DNS client registration using the ipconfig command.

**Cause:** The client was not able to register and update with the DNS server because of missing or incomplete DNS configuration.

**Solution:** Verify that the client is fully and correctly configured for DNS, and update its configuration as needed.

One common cause of the client failing to update with the DNS server is that it does not have a DNS suffix (either a primary suffix or connection-specific suffix) configured. This might result in the client attempting to register an incorrect or unintended DNS domain name.

For example, the client could be attempting to register its short or unqualified computer or host name as a top-level domain name in the root zone. This happens because without a DNS suffix configured for the client computer, it determines the configured short name of a computer (such as host-a) is its fully qualified domain name (FQDN). This only occurs because the computer name does not have a DNS suffix to append to it and qualify the computer name when registering it for the client in DNS.

To update the DNS configuration for a client, either:

1. Configure a primary DNS suffix at the client computer for static TCP/IP clients.

2. Configure a connection-specific DNS suffix for use at one of the installed network connections at the client computer.

**See also:** Configure the primary DNS suffix for a client computer; Configure TCP/IP to use DNS; Install and Configure Clients; Managing Clients.

**Cause:** The DNS client attempted to update its information with the DNS server but failed because of a problem related to the server.

**Solution:** If a client can reach its preferred and alternate DNS servers as configured, it is likely that the cause of its failed updates can be found elsewhere.

At Windows-based client computers, you can use Event Viewer to check the System log for any event messages that explain why attempts by the client to dynamically update its host (A) or pointer (PTR) resource records failed.

When reviewing messages in the System log, filter or order the display of all messages to view those that specify **DnsApi** as the source for the message. Typically, these messages are related to the performance of DNS activities, such as DNS queries or dynamic updates.

A common reason updates might fail for a mobile client is that the DNS server required to accept and perform the update does not respond when the client starts at a remote location on the network. This could be due to network performance issues or might indicate a problem in the underlying design of your network. Where these issues persist or seem likely, you should review your DNS deployment and modify it accordingly.

**See also:** Deploying DNS; Dynamic update.

**The DNS server is not performing dynamic updates.**
**Cause:** The DNS server does not support dynamic updates.

**Solution:** Verify that the DNS server used by the client can support the DNS dynamic update protocol, as described in RFC 2136.

Only the Windows 2000 and Windows Server 2003 DNS Server service supports dynamic updates. The DNS Server service provided with Windows NT Server 4.0 does not.

If you are using other DNS servers on your network, verify that they are running a DNS server implementation that supports dynamic updates.

**See also:** Dynamic update; DNS RFCs.

**Cause:** The DNS server supports dynamic updates but is not configured to accept them.

**Solution:** Verify that the primary zone where clients require updates is configured to allow dynamic updates.

The default for a new primary zone is to not accept dynamic updates. At the DNS server that loads the applicable primary zone, modify zone properties to allow updates.

**See also:** Allow dynamic updates.

**Cause:** The zone database is not available.

**Solution:** Verify that the zone is available for update.

First, if necessary, verify that the zone exists. For a standard primary zone, verify that the zone file exists at the server and that the zone is not paused. If you are using Active Directory-integrated zones, verify that the DNS server is running as a domain controller and has access to the Active Directory database where zone data is stored.

Secondary zones do not support dynamic updates. If you are trying to determine which server is the primary server for a standard zone, review zone authority records to determine which server is referenced in both the start of authority (SOA) and name server (NS) resource records for the zone. This is the primary server for the zone which can accept dynamic updates to it.

If you need to, you can use the DNS console to change a secondary zone to become a primary zone so that it can accommodate dynamic updates. However, because standard primary zones use a single-master update model, you can only configure one server to accept dynamic updates for the zone.

If you change the zone type at a secondary server so that it becomes the primary server for that zone, you need to either remove the zone or convert it to another zone type (such as a secondary zone) at the original primary server. Otherwise, zone data would become inconsistent and cause additional problems.

If you want to have more than one DNS server be able to update a zone, it is recommended that you change the zone type so that it becomes Active Directory-integrated. To be able to use this zone type, Active Directory must be installed and the server computer must be promoted to a domain controller.

Once the zone is stored in the directory, other domain controllers can load the zone automatically and be allowed to update it when they are running the DNS Server service. This is because Active Directory supports a multiple (or floating) master update model where more than one computer can process updates to the directory database.

**See also:** Change the zone type; Add and Remove Zones; Managing authority records; Active Directory integration.

**Cause:** The DNS server is configured to allow only secure dynamic updates and has a security-related problem.

**Solution:** Verify that zone or resource record security does not block or prevent dynamic updates at the server.

Secure update can be enabled for directory-integrated zones and their resource records. If secure dynamic update is in effect for a directory-integrated zone, then only users, groups, or computers that have write permissions may add new resource records to the zone. If secure dynamic update is in effect for resource records, then only users, groups, or computers that have write permissions can update these resource records. Consequently, security might block or prevent a DNS client (or its DHCP server) from performing an update of its host (A) and pointer (PTR) resource records.

In most cases, secure dynamic update does not prevent new records from being created or added to a zone, but it does restrict who is given default permissions to update or modify records. Where necessary, you can use the access control list (ACL) editing features available for directory-integrated zones to modify security permissions on a zone or its resource records and enable update by another user, group, or computer.

Typically, this is only needed if the computer requesting the update is different from the one that owns the client records and originally created them.

**See also:** Dynamic update; Modify security for a directory-integrated zone; Modify security for a resource record.

**Cause:** The DNS server required to perform the updates is not available on the network.

**Solution:** Verify that the DNS server is available on the network or troubleshoot any further issues as necessary.

**See also:** Troubleshooting DNS servers.

**I am having a different problem related to dynamic updates than those described above.**
**Cause:** My problem is not described above.

**Solution:** Search TechNet at the Microsoft Web site for the latest technical information that could relate to the problem. If necessary, you can obtain information and instructions that pertain to your problem or issue.

If you are connected to the Internet, the latest operating system updates are available at the Microsoft Web site.

To obtain the latest service pack updates for Windows NT Server, see the Microsoft Web site.

**See also:** DNS updated technical information; DNS.

# Troubleshooting zone problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting zone problems**
What problem are you having?

- I am having a problem related to zone transfers.
- I am trying to use a zone delegation but it appears to be broken.
- I am having a different zone problem than the ones described above.

**I am having a problem related to zone transfers.**
**Cause:** The DNS Server service is stopped or the zone is paused.

**Solution:** Verify that the master (source) and secondary (destination) DNS servers involved in completing transfer of the zone are both started and that the zone is not paused at either server.

**See also:** Start or stop a DNS server; Start a zone; Understanding zones and zone transfer.

**Cause:** The DNS servers used during a transfer do not have network connectivity with each other.

**Solution:** Eliminate the possibility of a basic network connectivity problem between the two servers.

Using the **ping** command, ping each DNS server by its IP address from its remote counterpart.

For example, at the source server, use the **ping** command to test IP connectivity with the destination server. At the destination server, repeat the ping test, substituting the IP address for the source server.

Both ping tests should succeed. If not, investigate and resolve intermediate network connectivity issues.

**See also:** Command-line utilities.

**Cause:** The serial number is the same at both the source and destination servers. Because the value is the same at both servers, no zone transfer occurs between the servers.

**Solution:** Using the DNS console, perform the following tasks:

1. Increase the value of the serial number for the zone at the master server (source) to a number greater than the value at the applicable secondary server (destination).
2. Once you have increased the serial number at the master server to a higher value than is used currently at the secondary server, initiate zone transfer at the secondary server.

When working within the DNS console, the zone serial number can be viewed from the **Start of Authority (SOA)** tab in the applicable zone properties. To increase this number in the zone, click **Increment**.

**See also:** Modify the start of authority (SOA) record for a zone; Initiate a zone transfer at a secondary server.

**Cause:** The master server (source) and its targeted secondary server (destination) are having interoperability-related problems.

**Solution:** Investigate possible causes for any problems related to interoperability between DNS servers running Windows Server 2003 and other DNS server implementations, such as an older version of the Berkeley Internet Name Domain (BIND) distribution.

Older BIND servers use an uncompressed zone transfer format. By default, servers running Windows Server 2003 (and later version BIND servers) use a faster compressed format during zone transfers. To accommodate zone transfer with older BIND servers, you need to change advanced server options at your DNS servers running Windows Server 2003.

Another possible interoperability issue is the use and inclusion of WINS forward lookup (WINS) resource records in a zone or their counterpart, the WINS reverse lookup (WINS-R) resource record used for reverse lookup zones. BIND servers do not recognize these records when they are included in zone data being transferred and can flag these records as bad data, possibly failing the zone transfer.

To prevent these records from being used or included in zone transfers to BIND-based servers and other servers that do not recognize them, select **Do not replicate this record** when configuring WINS properties at the applicable zone.

**See also:** Interoperability issues; Enable or disable fast transfer format during zone transfers; Interoperability issues; Using WINS lookup; Enable DNS to use WINS resolution.

**Cause:** The zone has resource records or other data that cannot be interpreted by the DNS server.

**Solution:** Verify that the zone does not contain incompatible data, such as unsupported resource records types or data errors.

In most cases, the DNS Server service supports all resource record (RR) types that are approved and required for Internet standard DNS usage.

Also, verify that the server has not been configured in advance to prevent loading a zone when bad data is found and investigate its method for checking names. These settings can be configured using DNS console.

**See also:** Resource records reference; Prevent loading of a zone when bad data is found; Change the name-checking method used by the DNS server; Checking names and zone data.

**Cause:** Authoritative zone data is incorrect.

**Solution:** If a zone transfer continues to fail, ensure that the zone does not contain nonstandard data.

If you manually edit zone files, be aware that records need to be formatted and used according to standard record usage and formatting guidelines as specified in the Request for Comments (RFCs) for DNS. In most cases, user input and data errors can be avoided if records are added and managed using the DNS console.

To determine if errored zone data is a likely source for a failed zone transfer, look in the DNS server event log for messages. You can also use the **nslookup** command with the **-ls** option to simulate and test a zone transfer, while observing the data returned in a terminates before full transfer of the zone is complete.

**See also:** Nslookup; DNS RFCs.

**I am trying to use a zone delegation but it appears to be broken.**
**Cause:** Zone delegations are not configured correctly.

**Solution:** Review how zone delegations are used and revise your zone configurations as needed.

Zones contain information about DNS domains and subdomains. For each new zone you create, the zone originally begins as a single-node database for one DNS domain. As needed, new subdomain nodes can be added directly below the original (parent) domain and stored as a single zone. Sometimes when new subdomains remain part of the same zone, they are called *sub-zones*.

If used as sub-zones, new subdomains are retained as part of the zone and replicated and updated along with it as a single entity. You can, however, delegate subdomains away and manage them in their own zones. For each subdomain delegated to its own zone, the parent zone needs to have delegation records added to it.

You can use the New Delegation wizard provided in the DNS console to simplify adding these records.

**See also:** Delegating zones; Create a zone delegation; Verify a zone delegation using the nslookup command.

**I am having a different zone problem than the ones described above.**
**Cause:** My problem is not described above.

**Solution:** Search TechNet at the Microsoft Web site for the latest technical information that could relate to the problem. If necessary, you can obtain information and instructions that pertain to your problem or issue.

If you are connected to the Internet, the latest operating system updates are available at the Microsoft Web site.

To obtain the latest service pack updates for Windows NT Server, see the Microsoft Web site.

**See also:** DNS updated technical information; DNS; Using the Windows Deployment and Resource Kits.

# DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP**

Dynamic Host Configuration Protocol (DHCP) is an IP standard designed to reduce the complexity of administering address configurations by using a server computer to centrally manage IP addresses and other related configuration details used on your network. The Microsoft® Windows Server 2003 family provides the DHCP service, which enables the server computer to perform as a DHCP server and configure DHCP-enabled client computers on your network as described in the current DHCP draft standard, RFC 2131.

DHCP includes Multicast Address Dynamic Client Assignment Protocol (MADCAP) which is used to perform multicast address allocation. When registered clients are dynamically assigned IP addresses through MADCAP, they can participate efficiently in the data stream process, such as for real-time video or audio network transmissions.

- Before installing a DHCP or MADCAP server, see DHCP Checklists.
- To find features that have been moved in the Windows Server 2003 family, see New ways to do DHCP tasks.
- For tips about using DHCP, see DHCP Best Practices.
- For help with specific tasks, see DHCP How To....
- For general background information, see DHCP Concepts.
- For problem-solving instructions, see DHCP Troubleshooting.

# DHCP Checklists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

- Checklist: Installing a DHCP server
- Checklist: Administering a DHCP server
- Checklist: Installing a MADCAP server

# Checklist: Installing a DHCP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a DHCP server**

| Step | Reference |
|---|---|
| Review DHCP concepts such as scopes, leases, and options.<br>**Important**<br>• If you are operating in a domain environment under an earlier version of Windows NT Server, it is recommended that you first upgrade your domain controllers to Windows Server 2003 family before installing DHCP on your network. | DHCP terminology |
| Determine the IP address range or ranges for which DHCP will be needed to provide configuration service on your network. | Configuring scopes |
| Manually configure TCP/IP properties for the DHCP server computer. | Configure TCP/IP for static addressing |
| Install the DHCP server. | Install a DHCP server |
| Open the DHCP console. | Open the DHCP console |
| Authorize the DHCP server in Active Directory. | Authorize a DHCP server in Active Directory; Authorizing DHCP servers |
| If the DHCP server is multihomed, disable service bindings for any connection not to be used to listen and provide service to clients. | Selectively set DHCP server bindings for network connections; Multihomed DHCP servers |
| Set DNS dynamic update policy for the DHCP server. | Enable DNS dynamic updates for clients; Using DNS servers with DHCP |
| (Optional) If you are using multiple scopes, configure any default option types for the server that should be inherited by all scopes. Server options are assigned to all DHCP clients that lease an IP address from any scope configured on the server. | Assign a server-based option; Assigning options |
| Create a new scope and assign commonly used options, such as WINS and DNS servers (unless these are already configured as server options). Scope options are assigned only to DHCP clients that lease an IP address from the scope where the option is configured. | Create a new scope; Configuring scopes; Assigning options |
| (Optional) Configure additional scope option types if you have subnets that need option configurations in addition to server options that are already configured. | Assign a scope-based option; Assigning options |
| (Optional) Add reservations for clients that require a reserved IP address assigned by DHCP. Reservations assure that DHCP clients keep the same IP address and have DHCP options refreshed each time the client renews the IP address lease. | Add a client reservation; Using client reservations |

| | |
|---|---|
| (Optional) Create exclusion ranges to prevent some IP addresses from being leased by the DHCP server. If you have network devices (such as printers, servers, or wireless access points) that must be configured with static IP addresses, create an exclusion range within the scope's address range; you can then assign static addresses to devices from the range of addresses excluded from distribution by the DHCP server. | Exclude an address from a scope |
| Activate the scope. | Activate a scope |
| (Optional) For additional fault tolerance, you can implement DHCP clustering support using a server cluster.<br><br>**Note**<br><br>• DHCP clustering servers are implemented using the Clustering service, which is provided only with Windows Server 2003, Enterprise Edition. | Cluster support for DHCP servers |

# Checklist: Administering a DHCP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Administering a DHCP server**
The following table lists optional actions you can take to administer your DHCP server, depending on your network configuration.

| Step | Reference |
|---|---|
| If clients share a commonly defined vendor type or a common need for a unique DHCP options configuration, configure and assign option classes. | Using option classes; Manage Options and Classes |
| If predefined options do not meet your needs, define additional option types. | Define a new option; Assigning options |
| If the default lease time of eight days is too long or too short, adjust the length of lease durations.<br><br>For example, you can shorten leases for wireless clients on wireless subnets, or lengthen leases for subnets whose DHCP clients all have reservations. | Change or view scope properties; Managing leases |
| If you need to implement support for more than one logical network for each physical network, create and use superscopes. | Create a superscope; Using superscopes |
| If you have a routed network, determine if you need to implement relay agents or update routers, and if so, review relay agent configurations and related design issues. | DHCP/BOOTP Relay Agents; Configure the DHCP Relay Agent |
| If you want to deploy your DHCP server as a multicast server, create multicast scopes. | Create a multicast scope; Using multicast scopes |
| If you want to support BOOTP clients on your network, configure the BOOTP table. | Add an entry to the BOOTP table; View the BOOTP table |
| If DHCP clients have problems accessing network resources, verify DHCP service to clients. | Verify, release, or renew a client address lease; Ipconfig; Command-line utilities |
| If you want to manually configure IP addresses for clients running Windows 2000, Windows XP, or Windows Server 2003 family, disable automatic address configuration. | Disable automatic address configuration |
| If you want to track DHCP server performance, use System Monitor. | Monitoring DHCP server performance; System Monitor overview |
| If you want to allow users to have read-only or administrative access to the DHCP server, add users to DHCP groups. | DHCP groups; Default groups |
| If you are having DHCP-related problems, troubleshoot and resolve common DHCP-related problems, or review DHCP server audit log files. | DHCP Troubleshooting; Analyzing server log files |

# Checklist: Installing a MADCAP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a MADCAP server**

| Step | Reference |
|---|---|
| Review Multicast Address Dynamic Client Allocation Protocol (MADCAP) concepts such as IP multicasting and multicast address allocation. | Multicast address allocation; Understanding Multicasting |
| Determine the multicast IP address range or ranges for use in your MADCAP server configuration. The ranges you select depend on whether you need to use administrative local scoping or global scoping practices to provide multicast address allocation service on your network. | Using multicast scopes |
| Manually configure TCP/IP properties for the DHCP (MADCAP) server computer. | Configure TCP/IP for static addressing |
| To install a MADCAP server, install the DHCP service.<br><br>**Important**<br><br>• The DHCP service can be used to deploy MADCAP servers independently of how DHCP servers are used on your network. For example, to install Windows Server 2003 DHCP for MADCAP service only:<br><br>1. Create multicast scopes.<br><br>2. Do not create other scopes or superscopes.<br><br>Only where other scopes or superscopes are configured, does the server computer also function as a DHCP server. | Install a DHCP server |
| Open the DHCP console. | Open the DHCP console |
| Authorize the DHCP (MADCAP) server in Active Directory. | Authorize a DHCP server in Active Directory |
| (Optional) If the DHCP (MADCAP) server is multihomed, you can disable service bindings for any connection not to be used to listen and provide service to DHCP (MADCAP) clients. | Selectively set DHCP server bindings for network connections; Multihomed DHCP servers |
| Create a multicast scope. | Create a multicast scope |
| If necessary, you can additionally manage multicast IP address allocation by creating and removing scopes, or by modifying multicast scope properties, such as scope TTL and lifetime settings as needed. | Manage Multicast Scopes |
| (Optional) For additional fault tolerance, you can implement DHCP clustering support using a server cluster.<br><br>**Note**<br><br>• DHCP clustering servers are implemented using the Clustering service, which is provided only with Windows Server 2003, Enterprise Edition. | Cluster support for DHCP servers |

# New ways to do DHCP tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for managing DHCP servers using the DHCP console. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Install a DHCP server on this computer | **Network** in Control Panel | Add or Remove Programs<br><br>For more information, see Install a DHCP server. |
| Start the DHCP console | Click **Start**, point to **Administrative Tools**, and then click **DHCP Manager**. | The DHCP console.<br><br>For more information, see Open the DHCP console. |
| Authorize a DHCP server in Active Directory prior to use | Not available in Windows NT Server 4.0 | Log on first as a user who is a member of the Enterprise Admins group for the enterprise where you need to authorize the server.<br><br>For more information, see Authorize a DHCP server in Active Directory. |
| Start or stop a DHCP server | **Services** in Control Panel | The DHCP console.<br><br>For more information, see Start or stop a DHCP server. |
| Add a server to the DHCP console to manage | **Server** menu in **DHCP Manager** | The DHCP console.<br><br>For more information, see Connect to a DHCP server. |
| Add and configure a scope for a DHCP server | **Create** option on **Scope** menu in **DHCP Manager** | The DHCP console.<br><br>For more information, see Create a new scope. |

# DHCP Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Use the 80/20 design rule for balancing scope distribution of addresses where multiple DHCP servers are deployed to service the same scope.**

  Using more than one DHCP server on the same subnet provides increased fault tolerance for servicing DHCP clients located on it. With two DHCP servers, if one server is unavailable, the other server can take its place and continue to lease new addresses or renew existing clients.

  A common practice when balancing a single network and scope range of addresses between two DHCP servers is to have 80 percent of the addresses distributed by one DHCP server and the remaining 20 percent provided by a second. For more information and an example of this concept, see Configuring scopes.

- **Use superscopes for multiple DHCP servers on each subnet in a LAN environment.**

  When started, each DHCP client broadcasts a DHCP discover message (DHCPDISCOVER) to its local subnet to attempt to find a DHCP server. Because DHCP clients use broadcasts during their initial startup, you cannot predict which server will respond to the DHCP discover request of a client if more than one DHCP server is active on the same subnet.

  For example, if two DHCP servers service the same subnet and its clients, clients can be leased at either server. Actual leases distributed to clients can depend on which server responds first to any given client. Later, the server first selected by the client to obtain its lease might be unavailable when the client attempts to renew.

  If renewal fails, the client then delays trying to renew its lease until it enters the rebinding state. In this state, the client broadcasts to the subnet to locate a valid IP configuration and continue without interruption on the network. At this point, a different DHCP server might respond to the client request. If this occurs, the responding server might send a DHCP negative acknowledgement message (DHCPNAK) in reply. This can occur even if the original server that first leased the client is available on the network.

  To avoid these problems when using more than one DHCP server on the same subnet, use a new superscope configured similarly at all servers. The superscope should include all valid scopes for the subnet as member scopes. For configuring member scopes at each server, addresses must only be made available at one of the DHCP servers used on the subnet. For all other servers in the subnet, use exclusion ranges for the same scope ranges of addresses when configuring the corresponding scopes.

  For more information, see Using superscopes.

- **Deactivate scopes only when removing a scope permanently from service.**

  Once you activate a scope, it should not be deactivated until you are ready to retire the scope and its included range of addresses from use on your network.

  Once a scope is deactivated, the DHCP server no longer accepts those scope addresses as valid addresses. This is only useful when the intention is to permanently retire a scope from use. Otherwise, deactivating a scope causes undesired DHCP negative acknowledgement messages (DHCPNAKs) to be sent to clients.

  If the intent is only to affect temporary deactivation of scope addresses, editing or modifying exclusion ranges in an active scope achieves the intended result without undesired results.

  For more information, see Manage Scopes.

- **Use server-side conflict detection on DHCP servers only when it is needed.**

Conflict detection can be used by either DHCP servers or clients to determine whether an IP address is already in use on the network before leasing or using the address.

DHCP client computers running Windows 2000 or Windows XP that obtain an IP address use a gratuitous ARP request to perform client-based conflict detection before completing configuration and use of a server offered IP address. If the DHCP client detects a conflict, it will send a DHCP decline message (DHCPDECLINE) to the server.

If your network includes legacy DHCP clients (clients running a version of Windows earlier than Windows 2000), you can use server-side conflict detection provided by the DHCP Server service under specific circumstances. For example, this feature might be useful during failure recovery when scopes are deleted and recreated. For more information, see DHCP Troubleshooting.

By default, the DHCP service does not perform any conflict detection. To enable conflict detection, increase the number of ping attempts that the DHCP service performs for each address before leasing that address to a client. Note that for each additional conflict detection attempt that the DHCP service performs, additional seconds are added to the time needed to negotiate leases for DHCP clients.

Typically, if DHCP server-side conflict detection is used, you should set the number of conflict detection attempts made by the server to use one or two pings at most. This provides the intended benefits of this feature without decreasing DHCP server performance.

For more information, see Enable address conflict detection.

- **Reservations should be created on all DHCP servers that can potentially service the reserved client.**

  You can use a client reservation to ensure that a DHCP client computer always receives the same IP address lease at startup. If you have more than one DHCP server reachable by a reserved client, add the reservation at each of your other DHCP servers.

  This allows the other DHCP servers to honor the client IP address reservation made for the reserved client. Although the client reservation is only acted upon by the DHCP server where the reserved address is part of the available address pool, you can create the same reservation on other DHCP servers that exclude this address.

  For more information, see Add a client reservation.

- **For server performance, note that DHCP is disk-intensive and purchase hardware with optimal disk performance characteristics.**

  DHCP causes frequent and intensive activity on server hard disks. To provide the best performance, consider RAID solutions when purchasing hardware for your server computer that improves disk access time.

  When evaluating performance of your DHCP servers, you should evaluate DHCP as part of making a full performance evaluation of the entire server. By monitoring system hardware performance in the most demanding areas of utilization (CPU, memory, disk input/output), you obtain the best assessment of when a DHCP server is overloaded or in need of an upgrade.

  Note that the DHCP service includes several System Monitor counters that can be used to monitor service. For more information, see Monitoring DHCP server performance.

- **Keep audit logging enabled for use in troubleshooting.**

  By default, the DHCP service enables audit logging of service-related events. Audit logging provides a long-term service monitoring tool that makes limited and safe use of server disk resources. For more information, see Audit logging.

  For more information on interpreting server audit log files, see Analyzing server log files.

- **Reduce lease times for DHCP clients that use Routing and Remote Access service for remote access.**

If Routing and Remote Access service is used on your network to support dial-up clients, you can adjust the lease time on scopes that service these clients to less than the default of eight days. One recommended way to support remote access clients in your scopes is to add and configure the built-in Microsoft vendor class provided for the purpose of client identification.

- **Increase the duration of scope leases for large, stable, fixed networks if available address space is plentiful.**

  For small networks (for example, one physical LAN not using routers), the default lease duration of eight days is a typical period. For larger routed networks, consider increasing the length of scope leases to a longer period of time, such as 16-24 days. This can reduce DHCP-related network broadcast traffic, particularly if client computers generally remain in fixed locations and scope addresses are plentiful (at least 20 percent or more of the addresses are still available).

- **Integrate DHCP with other services, such as WINS and DNS.**

  WINS and DNS can both be used for registering dynamic name-to-address mappings on your network. To provide name resolution services, you must plan for interoperability of DHCP with these services. Most network administrators implementing DHCP also plan a strategy for implementing DNS and WINS servers.

- **For routed networks, either use relay agents or set appropriate timers to prevent undesired forwarding and relay of BOOTP and DHCP message traffic.**

  If you have multiple physical networks connected through routers, and you do not have a DHCP server on each network segment, the routers must be capable of relaying BOOTP and DHCP traffic. If you do not have such routers, you can set up the DHCP Relay Agent component on at least one server running Windows Server 2003 in each routed subnet. The relay agent relays DHCP and BOOTP message traffic between the DHCP-enabled clients on a local physical network and a remote DHCP server located on another physical network.

  When using relay agents, be sure to set the initial time delay in seconds that relay agents wait before relaying messages on to remote servers. For more information on DHCP relay agents, see DHCP/BOOTP Relay Agents.

- **Use the appropriate number of DHCP servers for the number of DHCP-enabled clients on your network.**

  In a small LAN (for example, one physical subnet not using routers), a single DHCP server can serve all DHCP-enabled clients. For routed networks, the number of servers needed increases, depending on several factors, including the number of DHCP-enabled clients, the transmission speed between network segments, the speed of network links, whether DHCP service is used throughout your enterprise network or only on selected physical networks, and the IP address class of the network. For more information on determining how many DHCP servers to set up, see Planning DHCP networks.

- **For DNS dynamic updates performed by the DHCP service, use the default client preference settings.**

  The Windows Server 2003 DHCP service can be configured to perform DNS dynamic updates for DHCP clients based on how clients request these updates to be done. This setting provides the best use of the DHCP service to perform dynamic updates on behalf of its clients as follows:

  - DHCP client computers running Windows 2000, Windows XP, or a Windows Server 2003 operating system explicitly request that the DHCP server only update pointer (PTR) resource records used in DNS for the reverse lookup and resolution of the client's IP address to its name. These clients update their address (A) resource records for themselves.

  - Clients running earlier versions of Windows cannot make an explicit request for DNS dynamic update protocol preference. For these clients, the DHCP service updates both the PTR and the A resource records when the service is configured to do so.

  For more information, see Using DNS servers with DHCP, Enable DNS dynamic updates for clients, and Configure DNS dynamic update credentials.

- **Use the manual backup and restore methods in the DHCP server console.**

  Use the **Backup** command on the **Action** menu of the DHCP console to perform full backup of the DHCP service at an interval that protects you from significant data loss. When you use the manual backup method, all DHCP server data is included in the backup, including all scope information, log files, registry keys, and DHCP server configuration information (except DNS dynamic update credentials). Do not store these backups on the same hard drive upon which the DHCP service is installed, and make sure that the access control list (ACL) for the backup folder only contains the Administrators group and DHCP Administrator groups as members. In addition to performing manual backups, backup to other locations, such as a tape drive, and make sure unauthorized persons do not have access to your backup copies. You can use Windows Backup for this purpose. For more information, see Best practices for Backup.

  When restoring the DHCP service, use a backup created with the manual **Backup** command or a copy of the database created with synchronous backup by the DHCP service. In addition, use the **Restore** command on the **Action** menu in the DHCP console to restore a DHCP server.

  For more information, see Backing up the DHCP database and Restoring server data.

- **Follow the recommended process for moving a DHCP server database from old server computer hardware to new hardware.**

  Moving a DHCP server database can be problematic. To manage moving the server database more easily, choose and follow a process tried and used by Microsoft Product Support Services such as the following:

  - For restoring server data at the same server computer, see Restoring server data.

  - For moving DHCP server data to another server computer, such as in the case of hardware failure or data recovery, see Move a DHCP database to another server.

- **Before you install a DHCP server, identify the following:**

  - The hardware and storage requirements for the DHCP server.

    For more information, see Planning DHCP networks.

  - Which computers you can immediately configure as DHCP clients for dynamic TCP/IP configuration and which computers you should manually configure with static TCP/IP configuration parameters, including static IP addresses.

    For more information, see Checklist: Configuring TCP/IP.

  - The DHCP option types and their values to be predefined for DHCP clients.

    For more information, see Assigning options.

# DHCP How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Use DHCP client alternate configuration
- Install and Administer Servers
- Manage Server Access
- Monitor and Configure Servers
- Manage Scopes
- Manage Superscopes
- Manage Multicast Scopes
- Manage Options and Classes
- Manage Clients and Leases
- Manage Reservations and Reserved Clients
- Manage Databases
- Manage BOOTP and Remote Access Clients
- Use DHCP Command-line Tools

# Use DHCP client alternate configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use DHCP client alternate configuration**

1. Open Network Connections.
2. Right-click the network connection that you want to configure, and then click **Properties**. The **Local Area Connection Properties** dialog box appears.
3. On the **General** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**. The **Internet Protocol (TCP/IP) Properties** dialog box appears.
4. On the **General** tab, click **Obtain an IP address automatically**.
5. Click the **Alternate Configuration** tab, and then click **User configured**.
6. In **IP address** and **Subnet mask**, type an IP address and subnet mask.
7. You can perform any of the following optional tasks:
   - In **Default gateway**, type the default gateway address.
   - In **Preferred DNS server**, type the primary Domain Name Service (DNS) server address.
   - In **Alternate DNS server**, type the secondary DNS server address.
   - In **Preferred WINS server**, type the primary Windows Internet Name Service (WINS) server address.
   - In **Alternate WINS server**, type the secondary WINS server address.

**Note**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DHCP client alternate configuration

# Install and Administer Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and administer servers**

- Install a DHCP server
- Open the DHCP console
- Connect to a DHCP server
- Disconnect from a DHCP server
- Start or stop a DHCP server
- Reconcile the DHCP database
- Remove or uninstall DHCP

# Install a DHCP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a DHCP server**

1. Open Windows Components Wizard.
2. Under **Components**, scroll to and click **Networking Services**.
3. Click **Details**.
4. Under **Subcomponents of Networking Services**, click **Dynamic Host Configuration Protocol (DHCP)**, and then click **OK**.
5. Click **Next**. If prompted, type the full path to the Windows Server 2003 distribution files, and then click **Next**.

   Required files are copied to your hard disk.

## Notes

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- DHCP servers must be configured with a static IP address.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open the DHCP console
Authorize a DHCP server in Active Directory

# Open the DHCP console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open the DHCP console**

- Open DHCP.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- The DHCP console is an administrative tool for managing DHCP servers. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Install and Administer Servers
DHCP servers

# Connect to a DHCP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a DHCP server**

1. Open DHCP.
2. In the console tree, click **DHCP**.
3. On the **Action** menu, click **Add Server**.
4. In the **Add Server** dialog box, do one of the following:
   - Click **This server**, enter the name of the DHCP server that you want to connect to, and then click **OK**.
   - Click **This authorized DHCP server**, click the DHCP server that you want to connect to, and then click **OK**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- If you are operating DHCP in an Active Directory environment and are connecting to a server for the first time, you need to first authorize the new DHCP server in the directory service. For more information, see Related Topics.
- If DHCP is installed and running on this computer, you do not need to make a manual connection using the wizard. In most cases, the local DHCP server automatically appears in the list of servers when the DHCP console starts, and you connect to the server by clicking it in the console tree.
- Only users belonging to the following groups can connect to a DHCP server:
  - DHCP Users
  - DHCP Administrators
  - Domain Admins
  - Enterprise Admins

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Disconnect from a DHCP server
Authorize a DHCP server in Active Directory
Authorizing DHCP servers

# Disconnect from a DHCP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect from a DHCP server**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Delete**.
4. When prompted to confirm removing the server, click **Yes**.

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Connect to a DHCP server

# Start or stop a DHCP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start or stop a DHCP server**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, point to **All Tasks** and then click one of the following:

   - To start the service, click **Start**.
   - To stop the service, click **Stop**.
   - To interrupt the service, click **Pause**.
   - To stop and then automatically restart the service, click **Restart**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- After you pause or stop the server, the **Resume** option appears and can be clicked to immediately resume service.
- You can also perform these tasks at a command prompt by using the following commands:
  - **net start dhcpserver**
  - **net stop dhcpserver**
  - **net pause dhcpserver**
  - **net continue dhcpserver**
- You can also perform these tasks at the **netsh>** command prompt or in a script using the Netsh commands for DHCP. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install a DHCP server

Open the DHCP console
Net services commands
Netsh commands for DHCP
Netsh DHCP example
Enter a netsh context

# Reconcile the DHCP database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reconcile the DHCP database**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Reconcile All Scopes**.
4. In the **Reconcile All Scopes** dialog box, click **Verify**.

   Inconsistencies found are reported in the status window.

5. If the database is found to be consistent, click **OK**.

   If the database is not consistent, click the displayed addresses that need to be reconciled, and click **Reconcile** to repair inconsistencies.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Reconciling scopes
Reconcile a scope
Restoring server data

# Remove or uninstall DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove or uninstall DHCP**

1. If necessary, stop the DHCP server on this computer.

   For more information, see Related Topics.

2. Open Windows Components Wizard.

3. Under **Components**, scroll to and click **Networking Services**.

4. Click **Details**.

5. In **Subcomponents of Networking Services**, scroll to **Dynamic Host Configuration Protocol (DHCP)** and clear the check box.

## Notes

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

- When you remove DHCP, all DHCP files are deleted from the server except program files that are in use. When the system is restarted, these files are removed as well.

- If you do not stop the DHCP Server service in step 1, this procedure only deletes the DHCP server registry entries.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start or stop a DHCP server
The DHCP database

# Manage Server Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage server access**

- Authorize a DHCP server in Active Directory
- Unauthorize a DHCP server in Active Directory
- Delegate ability to authorize DHCP servers to a non-enterprise administrator
- Selectively set DHCP server bindings for network connections
- Add a user or group as a DHCP user at a member server
- Add a user or group as a DHCP user at a domain controller
- Add a user or group as a DHCP administrator at a member server
- Add a user or group as a DHCP administrator at a domain controller

# Authorize a DHCP server in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To authorize a DHCP server in Active Directory**

1. Open DHCP.
2. In the console tree, click **DHCP**.
3. On the **Action** menu, click **Manage authorized servers**.

   The **Manage Authorized Servers** dialog box appears.
4. Click **Authorize**.
5. When prompted, type the name or IP address of the DHCP server to be authorized, and then click **OK**.

## Notes

- For a DHCP server to be authorized in an Active Directory domain environment, you must first be logged on as a member of the Enterprise Admins group for the enterprise where the server is being added.
- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- This procedure is usually only needed if you are running a DHCP server on a member server. In most cases, if you are installing a DHCP server on a computer also running as a domain controller, the server is automatically authorized the first time you add the server to the DHCP console.
- The fully qualified domain name (FQDN) of the DHCP server cannot exceed 64 characters. If the FQDN of the DHCP server exceeds 64 characters, the attempt to authorize the server fails with the error message, "A constraint violation has occurred." If your DHCP server FQDN exceeds 64 characters, authorize the server using the IP address of the server instead of its FQDN.
- DHCP servers running Windows 2000 cannot be authorized using the Windows Server 2003 administrative tools (the DHCP console and the Netsh commands for DHCP) unless Windows 2000 Server Service Pack 2 is installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Unauthorize a DHCP server in Active Directory
Delegating administration
Authorizing DHCP servers
Delegate ability to authorize DHCP servers to a non-enterprise administrator

# Unauthorize a DHCP server in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To unauthorize a DHCP server in Active Directory**

1. Open DHCP.
2. In the console tree, click **DHCP**.
3. On the **Action** menu, click **Manage authorized servers**.

   The **Manage Authorized Servers** dialog box appears.
4. In **Authorized DHCP servers**, select the server you want to unauthorize.
5. Click **Unauthorize**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Removing authorization for a DHCP server from the directory causes that server to stop responding to client requests.
- If the server you unauthorize is the only DHCP server, clients will be unable to obtain IP address leases through DHCP.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delegating administration
Authorizing DHCP servers
Delegate ability to authorize DHCP servers to a non-enterprise administrator

# Delegate ability to authorize DHCP servers to a non-enterprise administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delegate ability to authorize DHCP servers to anon-enterpriseadministrator**

1. Open Active Directory Sites and Services while logged on as an enterprise administrator.

   For more information, see Notes.

2. On the **View** menu, click **Show Services Node**.

3. In the console tree, click **NetServices**.

   **Where?**

   - Active Directory Sites and Services/Services/NetServices

4. On the **Action** menu, click **Delegate Control**.

5. In the Delegation of Control Wizard, click **Next**.

6. On the **Users or Groups** page, click **Add**.

7. In the **Select Users, Computers, or Groups** dialog box, under **Enter the object names to select**, type the name of the user you want to add to the DHCP Administrators group, and then click **OK**. To add multiple users in one operation, type user names separated by semicolons, and then click **OK**. When you finish adding users, click **Next**.

8. On the **Tasks to Delegate** page, click **Create a custom task to delegate**, and then click **Next**.

9. On the **Active Directory Object Type** page, under **Delegate control of**, click **This folder, existing objects in this folder, and creation of new objects in this folder**, and then click **Next**.

10. On the **Permissions** page, under **Permissions**, select the **Full Control** check box, and then click **Next**.

    This enables access control rights to the user or group account that you delegated in this procedure so that the user or group can add, modify, or delete new objects of DHCP class type in the NetServices folder.

## Notes

- To be logged on as an enterprise administrator, you must log on using a member account in the Enterprise Admins group. By default, this can typically be done by logging on as local administrator at the first domain controller created in your enterprise.

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

- For more information on setting permissions, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authorizing DHCP servers
**Other Resources**
Delegate control

# Selectively set DHCP server bindings for network connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To selectively set DHCP server bindings for network connections**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. Click **Bindings**.
6. In the **Bindings** dialog box, for the list of **Connections and server bindings**, select or clear check boxes for any statically configured network connections that should be enabled or disabled for use at the applicable DHCP server. Click **OK**.
7. On the **Advanced** tab, click **Apply**.

**Notes**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- If a network connection has more than one IP address configured, only the first configured IP address is used or affected by the DHCP server bindings. Any additional configured IP addresses are not used or configurable for bindings and are never used to service clients.
- The DHCP Server service never binds to any wide area network (WAN) or DHCP-enabled network connections on a server computer. For this reason, these connections do not appear in the **Bindings** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Multihomed DHCP servers
Cluster support for DHCP servers
Troubleshooting DHCP servers

# Add a user or group as a DHCP user at a member server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a user or group as a DHCP user at a member server**

1. Open Computer Management if the DHCP server is installed as a domain member server (or a stand-alone workgroup server).

2. In the console tree, click **Local Users and Groups**.

   **Where?**

   - Computer Management (Local)/System Tools/Local Users and Groups

3. In the details pane, double-click **Groups**, and then click **DHCP Users**.

4. On the **Action** menu, click **Properties**.

5. On the **General** tab, click **Add**.

   The **Select Users** dialog box opens.

6. In **Enter the object names to select**, type the name of the user you want to add as a DHCP User, and then click **OK**. To add multiple users in one operation, type user names separated by semicolons, and then click **OK**.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- The DHCP Users group provides a way to grant read-only console access to the DHCP server. Members of DHCP Users group are granted the right to view, but not to modify, data for the applicable server in the DHCP console. For more information about DHCP groups, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a user or group as a DHCP administrator at a member server
Add a user or group as a DHCP user at a domain controller
Delegate ability to authorize DHCP servers to a non-enterprise administrator
DHCP groups
New features for DHCP

# Add a user or group as a DHCP user at a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a user or group as a DHCP user at a domain controller**

1. Open Active Directory Users and Computers if the DHCP server is installed at a domain controller.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*applicable domain*/Users

3. In the details pane, click **DHCP Users**.
4. On the **Action** menu, click **Properties**.
5. Click the **Members** tab, and then click **Add**.

   The **Select Users** dialog box opens.

6. In **Enter the object names to select**, type the name of the user you want to add as a DHCP User, and then click **OK**. To add multiple users in one operation, type user names separated by semicolons, and then click **OK**.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- The DHCP Users group provides a way to grant read-only console access to the DHCP server. Members of DHCP Users group are granted the right to view, but not to modify, data for the applicable server in the DHCP console. For more information about DHCP groups, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a user or group as a DHCP administrator at a member server
Delegate ability to authorize DHCP servers to a non-enterprise administrator
DHCP groups
New features for DHCP

# Add a user or group as a DHCP administrator at a member server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a user or group as a DHCP administrator at a member server**

1. Open Computer Management if the DHCP server is installed as a domain member server (or a stand-alone workgroup server).

2. In the console tree, click **Local Users and Groups**.

   **Where?**

   - Computer Management (Local)/System Tools/Local Users and Groups

3. In the details pane, double-click **Groups**, and then click **DHCP Administrators**.

4. On the **Action** menu, click **Properties**, and then click **Add**.

5. In **Enter the object names to select**, type the name of the user you want to add as a DHCP Administrator, and then click **OK**. To add multiple users in one operation, type user names separated by semicolons, and then click **OK**.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- The DHCP Administrators group provides a way to grant limited administrative access to the DHCP server. Members of the DHCP Administrators group are granted the right to fully administer the applicable server using DHCP tools (the DHCP console and Netsh commands for DHCP); they do not have access to other services running on the server and are not able to perform other administrative actions on the server computer. For more information about DHCP groups, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
New features for DHCP
DHCP tools
DHCP groups
Understanding Groups

# Add a user or group as a DHCP administrator at a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a user or group as a DHCP administrator at a domain controller**

1. Open Active Directory Users and Computers if the DHCP server is installed at a domain controller.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*applicable domain*/Users

3. In the details pane, click **DHCP Administrators**.
4. On the **Action** menu, click **Properties**.
5. Click the **Members** tab, and then click **Add**.

   The **Select Users** dialog box opens.

6. In **Enter the object names to select**, type the name of the user you want to add as a DHCP Administrator, and then click **OK**. To add multiple users in one operation, type user names separated by semicolons, and then click **OK**.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- The DHCP Administrators group provides a way to grant limited administrative access to the DHCP server. Members of the DHCP Administrators group are granted the right to fully administer the applicable server using DHCP tools (the DHCP console and Netsh commands for DHCP); they do not have access to other services running on the server and are not able to perform other administrative actions on the server computer. For more information about DHCP groups, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a user or group as a DHCP administrator at a member server
Add a user or group as a DHCP user at a domain controller
Add a user or group as a DHCP user at a member server
DHCP groups
Delegate ability to authorize DHCP servers to a non-enterprise administrator
New features for DHCP

# Monitor and Configure Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor and configure servers**

- Configure DNS dynamic update credentials
- Enable server statistics to automatically refresh
- Enable DHCP server logging
- Enable address conflict detection
- Enable DNS dynamic updates for clients
- Manually refresh server statistics
- View DHCP server statistics
- View DHCP server properties

# Configure DNS dynamic update credentials

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure DNS dynamic update credentials**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. In ***Server* Properties**, click the **Advanced** tab, and then click **Credentials**.

   The **DNS dynamic update credentials** dialog box appears.

5. In **DNS dynamic update credentials**, type the information required to provide credentials that will determine DNS record ownership, and then click **OK**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- If the supplied credentials belong to an object (such as a computer) that is a member of the **DnsUpdateProxy** security group, the next object to register the same name record in DNS will become the record owner.
- If you specified the credentials (user name, domain, and password) that the DHCP server uses when registering DHCP client computers in DNS, these credentials are not backed up with either synchronous or asynchronous backup. After a DHCP database is restored, new credentials must be configured.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using DNS servers with DHCP
Enable DNS dynamic updates for clients
Dynamic update
Allow dynamic updates

# Enable server statistics to automatically refresh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable server statistics to automatically refresh**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, for **Automatically update statistics every**, specify the number of hours and minutes to use as a refresh interval, and then click **OK**.

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manually refresh server statistics
View DHCP server statistics

# Enable DHCP server logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable DHCP server logging**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, select **Enable DHCP audit logging**, and then click **OK**.

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Audit logging
Analyzing server log files

# Enable address conflict detection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable address conflict detection**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. For **Conflict detection attempts**, type a number greater than 0 (zero) and less than six, and then click **OK**.

   The number you type determines how many times the DHCP server tests an IP address before leasing it to a client.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- When conflict detection attempts are set, the DHCP server uses the Packet Internet Groper (ping) process to test available scope IP addresses before including these addresses in DHCP lease offers to clients.
- A successful ping means the IP address is in use on the network. Therefore, the DHCP server does not offer to lease the address to a client. If the ping request fails and times out, the IP address is not in use on the network. In this case, the DHCP server offers to lease the address to a client.
- Each additional conflict detection attempt delays the DHCP server response by a second while waiting for the ping request to time out. This increases the load on the server. A value of no greater than two (2) for ping attempts is recommended.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DHCP Best Practices
Monitor and Configure Servers
Manage Reservations and Reserved Clients

# Enable DNS dynamic updates for clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable DNS dynamic updates for clients**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. Click the **DNS** tab.
5. Select the **Enable DNS dynamic updates according to the settings below** check box.
6. If needed, modify other DNS update settings.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- DHCP clients running Windows 2000, Windows XP, or a Windows Server 2003 operating system explicitly request that the DHCP server update only the pointer (PTR) resource records that are used in DNS for the reverse lookup and resolution of the IP address of the client to the name of the client. These clients update their address (A) resource records for themselves.
- Clients running earlier versions of Windows cannot make an explicit request for DNS dynamic update protocol preference. For these clients, the DHCP service updates both the PTR and the A resource records.
- For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using DNS servers with DHCP
Configure DNS dynamic update credentials
Enable DNS for DHCP-enabled clients

# Manually refresh server statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manually refresh server statistics**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Display Statistics**.
4. In the ***Server* Statistics** dialog box, click **Refresh**.

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View DHCP server statistics
Enable server statistics to automatically refresh

# View DHCP server statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view DHCP server statistics**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Display Statistics**.

### Note

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manually refresh server statistics
Enable server statistics to automatically refresh

# View DHCP server properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view DHCP server properties**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. View or modify server properties as needed.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor and Configure Servers
Manage Reservations and Reserved Clients
DHCP servers

# Manage Scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage scopes**

- Create a new scope
- Remove a scope
- Exclude an address from a scope
- Activate a scope
- Deactivate a scope
- Change or view scope properties
- Reconcile a scope

# Create a new scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new scope**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **New Scope**.
4. Follow the instructions in the New Scope Wizard.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Before you create the scope, determine starting and ending IP addresses to be used within it.
- Depending on the starting and ending IP addresses for your scope, the DHCP console suggests a default subnet mask useful for most networks. If you know a different subnet mask is required for your network, you can modify the value as needed.
- When you finish creating a new scope, you might need to complete additional tasks, such as activating the scope for use or assigning scope options.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Activate a scope
Assign a scope-based option
Configuring scopes

# Remove a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a scope**

1. Open DHCP.
2. In the console tree, click the applicable scope.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*

3. On the **Action** menu, click **Delete**.
4. When prompted, click **Yes** to delete the scope.

### Important

- Before removing a scope, be sure to have deactivated the scope long enough to migrate clients to a different scope. To migrate clients, either wait for 50% of the configured scope lease time to elapse or manually renew clients in another active scope. Once all clients have been moved or forced to seek leases in another scope, you can safely remove the inactive scope.

### Note

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Removing scopes
Activate a scope
Assign a scope-based option
Configuring scopes

# Exclude an address from a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To exclude an address from a scope**

1. Open DHCP.
2. In the console tree, click **Address Pool**.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Address Pool

3. On the **Action** menu, click **New Exclusion Range**.
4. In the **Add Exclusion** dialog box, type the **Start IP address** that you want to exclude from this scope.

   To exclude a range of more than one IP address, type an **End IP address**.

5. Click **Add**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Addresses that are excluded from the scope are not offered by the DHCP server to clients obtaining leases from this scope.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring scopes

# Activate a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate a scope**

1. Open DHCP.
2. In the console tree, click the applicable scope.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*

3. On the **Action** menu, click **Activate**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Activating a scope is needed only for new scopes to start lease distribution. You must activate a scope to make it available for DHCP clients.
- Do not activate a scope until you specify the options you want.
- The **Action** menu command changes to **Deactivate** when the selected scope is currently activated. Do not deactivate a scope unless you are retiring it permanently from use on your network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Deactivate a scope
Setting Up Scopes

# Deactivate a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To deactivate a scope**

1. Open DHCP.
2. In the console tree, click the applicable scope.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*

3. On the **Action** menu, click **Deactivate**.

**Important**

- Do not use **Deactivate** to pause a scope. Instead, adjust the available scope address range by setting exclusions for scope addresses.
- The **Action** menu command name changes to **Activate** when the selected scope is currently deactivated.

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Activate a scope
Setting Up Scopes

# Change or view scope properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change or view scope properties**

1. Open DHCP.
2. In the console tree, click the applicable scope.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*

3. On the **Action** menu, click **Properties**.
4. View or modify scope properties as needed.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring scopes
Setting Up Options

# Reconcile a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reconcile a scope**

1. Open DHCP.
2. In the console tree, click the applicable scope.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*

3. On the **Action** menu, click **Reconcile**.
4. In the **Reconcile** dialog box, click **Verify**.

   Inconsistencies found are reported in the status window.

5. If the scope is found to be consistent, click **OK**.

   If the scope is not consistent, select the displayed addresses that need to be reconciled, and click **Reconcile** to repair inconsistencies.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Reconciling scopes
Reconcile the DHCP database
Restoring server data

# Manage Superscopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage superscopes**

- Create a superscope
- Remove a superscope
- Add a scope to a superscope
- Activate a superscope
- Deactivate a superscope

# Create a superscope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a superscope**

1. Open DHCP.

2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **New Superscope**.

   This menu option only appears if at least one scope, that is not currently part of a superscope, has been created at the server.

4. Follow the instructions in the New Superscope Wizard.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

- A superscope can have scopes added to it, either during or after its creation.

- Scopes that are contained within superscopes are sometimes called *child scopes* or *member scopes*.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using superscopes
Remove a superscope

# Remove a superscope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a superscope**

1. Open DHCP.
2. In the console tree, click the applicable superscope.

   **Where?**

   - DHCP/*applicable DHCP server/applicable superscope*

3. On the **Action** menu, click **Delete**.
4. When prompted, click **Yes** to delete the superscope.

**Notes**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Removing a superscope does not remove any of the scopes included in it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using superscopes
Remove a scope

# Add a scope to a superscope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a scope to a superscope**

1. Open DHCP.
2. In the console tree, click the applicable scope.

   **Where?**

   - DHCP/*applicable DHCP server/applicable scope*

3. On the **Action** menu, click **Add To Superscope**.
4. In the **Available superscopes** list, click the superscope to which you want to add this scope.

**Notes**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- In order to complete this procedure, at least one superscope must already exist at the applicable DHCP server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using superscopes
Create a superscope

# Activate a superscope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate a superscope**

1. Open DHCP.
2. In the console tree, locate the applicable superscope.

   **Where?**

   - DHCP/*applicable DHCP server/applicable superscope*

3. On the **Action** menu, click **Activate**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Activating a superscope is needed only for new superscopes. When you activate a superscope, it activates all member scopes in the superscope, allowing the DHCP server to lease IP addresses from the superscope to clients on your network. You must activate a superscope to make IP addresses for all of its member scopes available for use by DHCP clients.
- Do not activate a superscope until you specify the options you want for each of its member scopes.
- The **Action** menu command changes to **Deactivate** when the selected scope is currently activated. Do not deactivate a superscope unless you intend to permanently retire all of its member scopes from use on your network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Deactivate a scope
Setting Up Options
Setting Up Scopes

# Deactivate a superscope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To deactivate a superscope**

1. Open DHCP.
2. In the console tree, click the applicable superscope.

   **Where?**

   - DHCP/*applicable DHCP server/applicable superscope*

3. On the **Action** menu, click **Deactivate**.

**Important**

- The **Action** menu command name changes to **Activate** when the selected superscope is currently deactivated.

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Deactivate a scope
Using superscopes

# Manage Multicast Scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage multicast scopes**

- Create a multicast scope
- Remove a multicast scope
- Set the multicast scope Time-To-Live (TTL) value
- Specify a finite lifetime for a multicast scope

# Create a multicast scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a multicast scope**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **New Multicast Scope**.
4. Follow the instructions in the New Multicast Scope Wizard.

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using multicast scopes
Remove a multicast scope
Set the multicast scope Time-To-Live (TTL) value
Specify a finite lifetime for a multicast scope

# Remove a multicast scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a multicast scope**

1. Open DHCP.
2. In the console tree, click the applicable multicast scope.

   **Where?**

   - DHCP/*applicable DHCP server/applicable multicast scope*

3. On the **Action** menu, click **Delete**.
4. When prompted, click **Yes** to remove the multicast scope.

**Notes**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Removing a multicast scope does not necessarily stop the use of multicast groups associated with scope addresses on your network.
- If computers on your network are configured by means other than DHCP to use these same addresses, the multicast group might still be used even after this procedure is complete.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using multicast scopes
Create a multicast scope
Set the multicast scope Time-To-Live (TTL) value
Specify a finite lifetime for a multicast scope

# Set the multicast scope Time-To-Live (TTL) value

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the multicast scope Time-To-Live (TTL) value**

1. Open DHCP.
2. In the console tree, click the applicable multicast scope.

   **Where?**

   - DHCP/*applicable DHCP server/applicable multicast scope*

3. On the **Action** menu, click **Properties**.
4. In **Time to live**, specify a new value.

**Notes**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- The multicast Time-To-Live (TTL) value specifies the number of routers (hops) that multicast traffic is permitted to pass through before expiring on the network. For each router (hop), the original specified TTL is decremented by one (1). When its TTL reaches a value of zero (0), each multicast datagram expires and is no longer forwarded through the network to other subnets.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using multicast scopes
Create a multicast scope
Remove a multicast scope
Specify a finite lifetime for a multicast scope

# Specify a finite lifetime for a multicast scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify a finite lifetime for a multicast scope**

1. Open DHCP.
2. In the console tree, click the applicable multicast scope.

   **Where?**

   - DHCP/*applicable DHCP server/applicable multicast scope*

3. On the **Action** menu, click **Properties**.
4. Click the **Lifetime** tab.
5. For **Multicast scope lifetime**, click **Multicast scope expires on**.
6. Using the controls, specify a date and time when the multicast scope is to be expired and removed from the applicable DHCP server.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- By default, multicast scopes do not expire. When created, they exist until they are manually removed. This procedure, however, can be used to provide a means of having the DHCP server automatically remove the multicast scope at an appointed future date and time.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using multicast scopes
Create a multicast scope
Remove a multicast scope
Set the multicast scope Time-To-Live (TTL) value

# Manage Options and Classes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage options and classes**

- Assign a server-based option
- Assign a scope-based option
- Assign an advanced (class-based) option
- Create a new user or vendor class
- Define a new option
- Modify an option
- Remove an option

# Assign a server-based option

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign a server-based option**

1. Open DHCP.
2. In the console tree, click **Server Options**.

   **Where?**

   - DHCP/*applicable DHCP server*/Server Options

3. On the **Action** menu, click **Configure Options**.
4. In list of **Available Options**, select the check box for the first option you want to configure.
5. Under **Data entry**, type the information required for this option.
6. Repeat the two previous steps for any other server options you want to specify, and then click **OK**.
7. In addition, you can click the **Advanced** tab and specify additional server options to be applied only for members of select user or vendor classes.

   Options assigned at this level are only provided to clients identified as members of the classes specified at this tab.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Server options are inherited as default values for all scopes created at the applicable DHCP server. Option values you configure here are overridden by any scope-based or client-based values assigned for the same options.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign a scope-based option
Assign an option to a reserved client
Manage Options and Classes
DHCP options
Assigning options

# Assign a scope-based option

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign a scope-based option**

1. Open DHCP.
2. In the console tree, click **Scope Options**.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Scope Options

3. On the **Action** menu, click **Configure Options**.
4. In list of **Available Options**, select the check box for the first option you want to configure.
5. Under **Data entry**, type the information required for this option.
6. Repeat the two previous steps for any other server options you want to specify, and then click **OK**.
7. In addition, you can click the **Advanced** tab and specify additional scope options to be applied only for members of select user or vendor classes.

   Options assigned at this level are only provided to clients identified as members of the classes specified at this tab.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Scope options are inherited as default values for all clients of the applicable scope. Option values you configure here are overridden by any reserved client-based values assigned for the same options.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign a server-based option
Assign an option to a reserved client
Assigning options
DHCP options

# Assign an advanced (class-based) option

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign an advanced (class-based) option**

1. Open DHCP.
2. In the console tree, choose one:
   - Click **Server Options**, if you want to assign advanced options for use by all new scopes at the DHCP server.

     **Where?**

     - DHCP/*applicable DHCP server*/Server Options
   - Click **Scope Options**, if you want to assign advanced options for use by a specific scope at the DHCP server.

     **Where?**

     - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Scope Options
3. Right-click the folder, and click **Configure Options**.
4. Click the **Advanced** tab.
5. Select from the following:
   - For **Vendor class**, scroll to select the defined vendor class at the DHCP server (or use the default, **DHCP Standard Options**) for assigning options for the server or scope.
   - For **User class**, scroll to select the defined user class at the DHCP server (or use the default, **Default User Class**) for assigning options for the server or scope.
6. In the list of **Available Options**, select the check box for the first option you want to configure.
7. Under **Data entry**, type the information required for this option.
8. Repeat the two previous steps for any other options you want to specify, and then click **OK**.

**Notes**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- If a class-based option is assigned, it only applies to DHCP clients that are identified as members of the specified user or vendor class.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Manage Options and Classes
Assigning options
Using option classes
Understanding user and vendor classes

# Create a new user or vendor class

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new user or vendor class**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, choose one of the following:

   - To create a new user class, click **Define User Classes**.
   - To create a new vendor class, click **Define Vendor Classes**.

4. Click **Add**.
5. In **New Class**, type the required information.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- In the **DHCP User Classes** and **DHCP Vendor Classes** dialog boxes, you can also click either **Edit** or **Remove** to modify or delete option classes.
- In the **New Class** dialog box, for **ID**, set the same binary or ASCII string that matches the DHCP class ID string which is either implemented or set at member clients.
- You can use the **ipconfig /showclassid** command to display all the DHCP class IDs allowed for any network adapter installed on your computer.
- For DHCP client computers running Windows 2000 or Windows XP, the **ipconfig /setclassid** command can be used to set the specified DHCP class ID string. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set DHCP class ID information at a client computer
Assigning options

# Define a new option

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define a new option**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Set Predefined Options**.
4. In **Predefined Options and Values**, click **Add**.
5. In **Option Type**, type the required information to define the new option, and then click **OK**.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- By default, standard DHCP options and BOOTP extensions as described in RFC 2132 are predefined for use at DHCP servers running Windows Server 2003 . For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Remove an option
Modify an option
Setting Up Options

# Modify an option

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify an option**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Set Predefined Options**.
4. In **Predefined Options and Values**, select the following:

   - For **Option class**, select the class where the option to be modified is used.
   - For **Option name**, scroll to select the option you want to modify.

5. Click **Edit**.
6. Modify the **Name** and **Description** fields for the option as needed.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- For options already defined, only the **Name** and **Description** fields can be modified using this procedure. To change other fields (such as **Data type**, **Code**, or the applicable class where the option is used), the existing option must be removed and a new option defined.
- By default, standard DHCP options and BOOTP extensions as described in RFC 2132 are predefined for use at DHCP servers running Windows Server 2003 . For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Define a new option
Remove an option
Setting Up Options

# Remove an option

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove an option**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP
   - *applicable DHCP server*

3. On the **Action** menu, click **Set Predefined Options**.
4. In **Predefined Options and Values**, select the following:

   - For **Option class**, select the class where the option to be removed is used.
   - For **Option name**, scroll to select the user-defined option you want to remove.

5. Click **Delete**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- By default, standard DHCP options and BOOTP extensions as described in RFC 2132 are predefined at DHCP servers running Windows Server 2003 for use and cannot be removed using this procedure. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Define a new option
Modify an option
Setting Up Options

# Manage Clients and Leases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage clients and leases**

- View client lease information
- Delete a client lease
- Verify, release, or renew a client address lease
- Set DHCP class ID information at a client computer
- Show DHCP class ID information at a client computer

# View client lease information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view client lease information**

1. Open DHCP.
2. In the console tree, click **Address Leases**.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Address Leases

3. In the details pane, locate the client lease for which you want to view information.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- For most clients displayed in the DHCP console, you cannot edit the displayed lease information. Only reserved scope clients permit some of this information to be edited.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a client reservation
Delete a client lease

# Delete a client lease

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a client lease**

1. Open DHCP.
2. In the console tree, click **Address Leases**.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Address Leases

3. In the details pane, click the IP address of the client you want to delete.
4. On the **Action** menu, click **Delete**.
5. Make a reservation with the IP address, or exclude it from the range.
6. Force the client with the existing lease to give it up. At the command prompt on the client computer, type **ipconfig /release**.
7. If needed, give the client a new IP address. At the command prompt on the client computer, type **ipconfig /renew**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- If you want, you can cancel leases for all clients in the scope. To cancel all current leases, select all of the clients in **Address Leases**, right-click, and then click **Delete**.
- Deleting a client lease does not prevent the server from offering that IP address again in a future client lease. It also does not prevent the client from requesting the address again, unless steps 5 to 7 are followed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Verify, release, or renew a client address lease
Add a client reservation
Exclude an address from a scope

# Verify, release, or renew a client address lease

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify, release, or renew a client address lease**

1. At a DHCP-enabled client computer running Windows XP or a member of Windows Server 2003 family, open a Command Prompt.

2. Use the Ipconfig command-line utility to verify, release, or renew the lease of the client with a DHCP server, as follows:

   - To verify a DHCP client lease, type **ipconfig** to view lease-status information.

     Or, type **ipconfig /all**.

   - To release a DHCP client lease, type **ipconfig /release**.

   - To renew a DHCP client lease, type **ipconfig /renew**.

**Note**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DHCP Troubleshooting
Ipconfig
Troubleshooting TCP/IP

# Set DHCP class ID information at a client computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set DHCP class ID information at a client computer**

1. At a DHCP-enabled client computer running Windows XP, open a Command Prompt.

2. Use the Ipconfig command-line utility to set the DHCP class ID the client uses when obtaining its lease from the DHCP server.

   You can type the **ipconfig /setclassid** command as demonstrated in the following example command, which sets an ASCII string ("MyNewClassId") as the DHCP class ID string for the local area network connection in use at the client computer:

```
C:\>ipconfig /setclassid "Local Area Connection" MyNewClassId

Windows XP IP Configuration

        DHCP ClassId successfully modified for adapter"Local Area Connection"
```

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.

- You can use the **ipconfig /showclassid** command to display all the DHCP class IDs allowed for any network adapter installed on your computer.

- For a DHCP options class to be in effect between the DHCP server and its clients, both must be configured with a matching DHCP class ID string that identifies the client as a member of a specific user or vendor options class when it obtains its address lease from the server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assigning options
Manage Options and Classes
Show DHCP class ID information at a client computer
Using option classes
Ipconfig

# Show DHCP class ID information at a client computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To show DHCP class ID information at a client computer**

1. At a DHCP-enabled client computer running Windows XP or Windows 2000, open a Command Prompt.

2. Use the Ipconfig command-line utility to show the DHCP class ID that the client uses when obtaining its lease from the DHCP server.

   You can type the **ipconfig /showclassid** command as demonstrated in the following example command. This example shows an ASCII string ("MyNewClassId") that is currently set as the DHCP class ID for the local area network connection at the client computer:

```
C:\>ipconfig /showclassid "Local Area Connection"

Windows XP IP Configuration

DHCP Class ID for Adapter "Local Area Connection":
        DHCP ClassID Name . . . . . . . . . : Default BOOTP Class
        DHCP ClassID Description  . . . . . : User class for BOOTP clients
        DHCP ClassID Name . . . . . . . . . : Default Remote Access Class
        DHCP ClassID Description  . . . . . : User class for remote access clients

        Host Name . . . . . . . . . . . . . : HOSTA
        Primary DNS Suffix  . . . . . . . . : example.microsoft.com
        Node Type . . . . . . . . . . . . . : Hybrid
        IP Routing Enabled. . . . . . . . . : No
        WINS Proxy Enabled. . . . . . . . . : No

Ethernet adapter Local Area Connection:

        Connection-specific DNS Suffix  . :
        Description . . . . . . . . . . . . : Combo PCMCIA EthernetCard
        Physical Address. . . . . . . . . . : 00-00-00-00-7C-DC
        DHCP Enabled. . . . . . . . . . . . : Yes
        Autoconfiguration Enabled . . . . . : Yes
        IP Address. . . . . . . . . . . . . : 10.0.0.51
        Subnet Mask . . . . . . . . . . . . : 255.255.255.0
        Default Gateway . . . . . . . . . . : 10.0.0.1
        DHCP Class ID . . . . . . . . . . . : MyNewClassId
        DNS servers . . . . . . . . . . . . : 10.0.0.3
        Primary WINS server . . . . . . . . : 10.0.0.5
```

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.

- For a DHCP options class to be in effect between the DHCP server and its clients, both must be configured with a matching DHCP class ID string that identifies the client as a member of a specific user or vendor options class when it obtains its address lease from the server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Ipconfig
Manage Options and Classes
Set DHCP class ID information at a client computer
Setting Up Options
Understanding user and vendor classes

# Manage Reservations and Reserved Clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage reservations and reserved clients**

- Add a client reservation
- Change information for a reserved client
- Assign an option to a reserved client
- Configure advanced (class-based) options for a reserved client
- Change the reserved IP address for a client

# Add a client reservation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a client reservation**

1. Open DHCP.
2. In the console tree, click **Reservations**.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Reservations

3. On the **Action** menu, click **New Reservation**.
4. In **New Reservation**, type the information required to complete the client reservation.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**.

5. To add the client reservation to the scope, click **Add**.
6. Repeat the two previous steps for any other client reservations you want to add, and then click **Close**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- When adding a reservation, the DHCP server checks for the address you enter in the database to make sure it is not currently in use by another client.
- Reservations are applied on the basis of the DHCP client media access control (MAC) address for the applicable network connection. To locate this address, at the client computer, type **ipconfig /all** at a command prompt, and view the Physical Address for the connection. Alternately, you can right-click the connection in the Network Connections folder, and select **Properties** to view this information.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View the status of a local area connection
Using client reservations
Change information for a reserved client
Configure advanced (class-based) options for a reserved client
Change the reserved IP address for a client

# Change information for a reserved client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change information for a reserved client**

1. Open DHCP.
2. In the console tree, click **Reservations**.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Reservations

3. In the details pane, click the reserved client for which you want to change information.
4. On the **Action** menu, click **Properties**.
5. Modify the reservation properties you want to change, and then click **OK**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Reservation IP addresses cannot be modified. If you want to change the IP address of a reservation, delete the reservation, and then recreate it with the new IP address.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a client reservation
Using client reservations

# Assign an option to a reserved client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign an option to a reserved client**

1. Open DHCP.
2. In the console tree, click **Reservations**.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Reservations

3. In the details pane, click the reserved client for which you want to assign options.
4. On the **Action** menu, click **Configure Options**.
5. In the list of **Available Options**, select the check box for the first option you want to configure.
6. Under **Data entry**, type the information required for this option.
7. Repeat the two previous steps for any other reservation options you want to specify, and then click **OK**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Reservation options configured here override option values assigned at either the server or scope levels of configuration.
- Locally defined options configured manually at the client override reservation options.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign a server-based option
Assign a scope-based option
Using option classes
Manage Options and Classes
DHCP options
Assigning options

# Configure advanced (class-based) options for a reserved client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure advanced (class-based) options for a reserved client**

1. Open DHCP.
2. In the console tree, click **Reservations**.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Reservations

3. In the details pane, click the reserved client for which you want to assign options.
4. On the **Action** menu, click **Configure Options**.
5. Click the **Advanced** tab.
6. Select from the following:

   - For **Vendor class**, scroll to select the defined vendor class at the DHCP server (or use the default, **DHCP Standard Options**) for assigning options to the reserved client.
   - For **User class**, scroll to select the defined user class at the DHCP server (or use the default, **Default User Class**) for assigning options to the reserved client.

7. In the list of **Available Options**, select the check box for the first option you want to configure.
8. Under **Data entry**, type the information required for this option.
9. Repeat the two previous steps for any other options you want to specify, and then click **OK**.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- If a class-based option is assigned, it only applies if the reserved client is identified as a member of the specified user or vendor classes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign an advanced (class-based) option
Using option classes
Understanding user and vendor classes

Assigning options
Using client reservations

# Change the reserved IP address for a client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the reserved IP address for a client**

1. Open DHCP.
2. In the console tree, click **Reservations**.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*/Reservations

3. In the details pane, select the reserved client for which you want to change the IP address.
4. To delete the current reservation, on the **Action** menu, click **Delete**.
5. Add a new reservation in DHCP that uses the new IP address for the client.

   For more information, see Related Topics.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- The IP address in an existing client reservation cannot be modified directly. To achieve similar results, the existing reservation must be deleted and a new reservation added.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add a client reservation
Change information for a reserved client
Configure advanced (class-based) options for a reserved client

# Manage Databases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage databases**

- Back up the DHCP database
- Move a DHCP database to another server
- Restore the DHCP database
- Set the DHCP database backup path
- Set the DHCP database path

# Back up the DHCP database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up the DHCP database**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server/*

3. On the **Action** menu, click **Backup**.
4. In the **Browse For Folder** dialog box, select the folder where you want to store the backup DHCP database, and then click **OK**.

## Notes

- To perform this procedure you must be a member of the Administrators group or the DHCP Administrators group on the DHCP server.
- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- You must choose a local drive for the DHCP database backup folder.
- The DHCP service creates an automatic backup of the DHCP database during normal operation. The default storage location for this copy of the database backup is *systemroot*\System32\Dhcp\Backup.
- It is recommended that you back up your DHCP database to a location other than the local drive using Windows Backup (ntbackup.exe) or non-Microsoft backup software. For more information, see Related Topics.
- When you store a manually created backup of the DHCP database in the same location as the synchronous backup that the DHCP server creates every 60 minutes, the manual backup is overwritten when automatic backup occurs.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Backing up the DHCP database
Restoring server data
Restore the DHCP database
Using Network Monitor
Backup
Checklists for Backing Up and Restoring Data
Best practices for Backup

# Move a DHCP database to another server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move a DHCP database to another server**
This topic provides details on how to move a DHCP database from one server computer (the source server) to another server computer (the destination server).

**To back up the DHCP database (at the source server)**

1. Open DHCP.

2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Backup**.

4. In the **Browse For Folder** dialog box, select the folder that will contain the backup DHCP database, and then click **OK**.

5. Stop the DHCP server.

   This prevents the server from assigning new address leases to clients after the database has been backed up.

   For more information, see Start or stop a DHCP server.

6. Disable the DHCP Server service in the list of services.

   This prevents the DHCP server from starting after the database has been transferred. For more information, see Enable or disable a service for a hardware profile.

7. Copy the folder that contains the backup DHCP database to the destination DHCP server computer.

### Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

- You must choose a local drive for the DHCP database backup folder.

- When you stop the DHCP service, DHCP clients cannot obtain address leases or renewals from the server.

- The DHCP service creates an automatic backup of the DHCP database during normal operation. The default storage location for this copy of the database backup is *systemroot*\System32\Dhcp\Backup.

**To restore the DHCP database (at the destination server)**

1. If you have not already installed the DHCP server software, install it.

   For more information, see Install a DHCP server.

2. Open DHCP.

3. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

4. On the **Action** menu, click **Restore**.

5. In the **Browse For Folder** dialog box, select the folder that contains the backup DHCP database, and then click **OK**.

## Notes

- To perform this procedure you must be a member of the Administrators group or the DHCP Administrators group on the source and destination DHCP servers.

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

- Only DHCP databases from the same language version can be restored. For example, a DHCP database from DHCP server running an English language version of the operating system cannot be restored to a DHCP server running a Japanese language version of the operating system.

- To restore the DHCP database, the DHCP service is temporarily stopped. When it is stopped, DHCP clients are unable to contact the DHCP server and obtain IP addresses.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Restore the DHCP database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore the DHCP database**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP*applicable DHCP server*

3. On the **Action** menu, click **Restore**.
4. In the **Browse For Folder** dialog box, select the folder that contains the backup DHCP database, and then click **OK**.

## Notes

- To perform this procedure you must be a member of the Administrators group or the DHCP Administrators group on the DHCP server.
- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Only a DHCP database backed up from the same Windows Server version can be restored. Restoring a DHCP database from a different Windows server version is not supported. To migrate a DHCP database from one Windows Server version to another, use the netsh export/import command.
- Only DHCP databases from the same language version can be restored. For example, a DHCP database from DHCP server running an English language version of the operating system cannot be restored to a DHCP server running a Japanese language version of the operating system.
- To restore the DHCP database, the DHCP service is temporarily stopped. When it is stopped, DHCP clients are unable to contact the DHCP server and obtain IP addresses.
- For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Backing up the DHCP database
Back up the DHCP database
Restoring server data

# Set the DHCP database backup path

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the DHCP database backup path**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. In the ***Server* Properties** dialog box, click the **Advanced** tab, click **Backup path**, and delete the default backup path (C:\WINDOWS\System32\dhcp\backup).
5. Type the backup path you prefer, or click **Browse** to select a folder location.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- This procedure describes how to specify the path DHCP uses when backing up the database. It does not set the path used by DHCP to store the primary database on the applicable server.
- By default, the DHCP database is automatically backed up every hour. However, when you change the backup folder location, an automatic backup of the database does not occur. You can back up the database immediately after changing the backup folder location by using **Backup** on the DHCP console **Action** menu. Automatic backups will then occur every hour.
- For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the DHCP database path
Backing up the DHCP database
Restoring server data
Restore the DHCP database

# Set the DHCP database path

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the DHCP database path**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. In ***Server* Properties**, click the **Advanced** tab, select **Database path**, and then delete the default database path (C:\WINDOWS\System32\Dhcp).
5. Type the path you prefer, or click **Browse** to select a folder location.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- This procedure describes how to specify the path DHCP uses when storing the database on the applicable DHCP server. It does not set the path used by DHCP when you back up a database with the **Backup** command.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the DHCP database backup path
Backing up the DHCP database
Restoring server data
Restore the DHCP database

# Manage BOOTP and Remote Access Clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage BOOTP and remote access clients**

- View the BOOTP table
- Add an entry to the BOOTP table
- Add dynamic BOOTP client support to a scope
- Change the lease time for remote access clients

# View the BOOTP table

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the BOOTP table**

1. Open DHCP.
2. In the console tree, click the applicable DHCP server.

   **Where?**

   - DHCP/*applicable DHCP server*

3. On the **Action** menu, click **Properties**.
4. Verify that the **Show the BOOTP table folder** check box is selected.

   If the check box is cleared, select it and then click **OK**.

5. In the console tree, click **BOOTP Table**.

   **Where?**

   - DHCP/*applicable DHCP server*/BOOTP Table

   BOOTP entries, if any have been added to the BOOTP table, appear in the details pane.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- By default, the BOOTP table is empty. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add an entry to the BOOTP table
Add dynamic BOOTP client support to a scope
BOOTP and DHCP

# Add an entry to the BOOTP table

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an entry to the BOOTP table**

1. Open DHCP.
2. In the console tree, click **BOOTP Table**.

   **Where?**

   - DHCP/*applicable DHCP server*/BOOTP Table

   If the BOOTP table does not appear, see Related Topics.
3. On the **Action** menu, click **New Boot Image**.
4. In the **Add BOOTP Entry** dialog box, enter the required information.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**.
5. Click **Add**.
6. Repeat the previous two steps for each additional BOOTP entry you want to add, and then click **Close**.

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add dynamic BOOTP client support to a scope
View the BOOTP table
BOOTP and DHCP

# Add dynamic BOOTP client support to a scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add dynamic BOOTP client support to a scope**

1. Open DHCP.
2. In the console tree, click the applicable scope.

   **Where?**

   - DHCP/*applicable DHCP server*/*applicable superscope* (if used)/*applicable scope*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. For **Assign IP addresses dynamically to clients of**, click either **Both** or **BOOTP only**.
6. If needed, in **Lease duration for BOOTP clients** adjust the default lease length.

   The lease length is used by the DHCP server to determine how long to preserve leases made dynamically for BOOTP clients of the applicable scope.

## Notes

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Unlike DHCP, where clients can initiate a protocol message to release or renew their address lease, BOOTP clients can only obtain or negotiate their IP address lease at boot time. For this reason, it is recommended that you select a lease length for BOOTP clients that exceeds the time between when BOOTP clients are restarted on the network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add an entry to the BOOTP table
BOOTP and DHCP
View the BOOTP table

# Change the lease time for remote access clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the lease time for remote access clients**

1. Open DHCP.
2. In the console tree, click **Scope Options**.

   **Where?**

   - DHCP/*applicable DHCP server/applicable scope*

3. Right-click **Scope Options** and click **Configure Options**.
4. Click **Advanced**.
5. In **User class**, select **Default Routing and Remote Access Class**.
6. In **Available Options**, click the check box for **051 Lease**.
7. In **Long**, type the lease value in seconds.

   For example, if you want the lease time to be one hour, type: 3600

**Note**

- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Setting Up Options
Using Routing and Remote Access servers with DHCP

# Use DHCP Command-line Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use DHCP command-line tools**

- Use DHCP commands interactively at the command prompt
- Use DHCP commands non-interactively or in scripts

# Use DHCP commands interactively at the command prompt

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use DHCP commands interactively at the command prompt**

1. Open Command Prompt.
2. Type **netsh**.
3. At the **netsh>** command prompt, type **dhcp**.
4. At the **dhcp>** command prompt, type:

   **server***ServerName*

   Or, type:

   **server***IPAddress*

   Where *ServerName* and *IPAddress* are the server that you want to manage.
5. Once connected, you can use any supported Netsh DHCP context command. Type **/?** or **help** to display the immediate context command menu, or type **list** to list all of the Netsh DHCP context commands available for use.

## Notes

- You must be a member of the Administrators group or the DHCP Administrators group to administer DHCP servers using the DHCP console or the Netsh commands for DHCP. You should restrict the membership of these groups to the minimum number of users necessary to administer the server. In addition, only members of the Domain Admins group can authorize or unauthorize a DHCP server in Active Directory. If there are users who need read-only access to the DHCP console, add them to the DHCP Users group instead of the DHCP Administrators group. For more information about using Netsh commands for DHCP, see Related Topics.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- The *****ServerName* parameter is optional and can be left unspecified if you are performing commands locally at the DHCP server. If the server being managed is remote, this parameter is required, and the name of the server must be specified.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing DHCP from the command line
Use DHCP commands non-interactively or in scripts
Configure the command prompt

# Use DHCP commands non-interactively or in scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use DHCP commands non-interactively or in scripts**

1. Open Command Prompt.

2. At the command prompt or in scripted usage, you can use Netsh commands for DHCP non-interactively by using the following fully qualified command syntax:

   **netsh dhcp server** [****_ServerName_] _SubCommand_

## Notes

- You must be a member of the Administrators group or the DHCP Administrators group to administer DHCP servers using the DHCP console or the Netsh commands for DHCP. You should restrict the membership of these groups to the minimum number of users necessary to administer the server. In addition, only members of the Domain Admins group can authorize or unauthorize a DHCP server in Active Directory. If there are users who need read-only access to the DHCP console, add them to the DHCP Users group instead of the DHCP Administrators group. For more information about using Netsh commands for DHCP, see Related Topics.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The ****_ServerName_ parameter is optional and can be left unspecified if you are performing commands locally at the DHCP server. If the server being managed is remote, this parameter is required, and the name of the server must be specified.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing DHCP from the command line
Use DHCP commands interactively at the command prompt
Configure the command prompt
Netsh commands for DHCP

# DHCP Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**
This section provides general background information about Dynamic Host Configuration Protocol (DHCP) and server and client software provided by Microsoft and the Windows Server 2003 family.

- DHCP Overview
- Understanding DHCP
- Administering DHCP
- DHCP Resources

# DHCP Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP overview**

This section covers:

- Security information for DHCP
- DHCP defined
- Benefits of using DHCP
- How DHCP works
- DHCP terminology
- DHCP tools
- Multicast address allocation
- New features for DHCP
- Planning DHCP networks

# Security information for DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for DHCP**
It is important to follow best practices for security when using DHCP servers on your network. For more information, see Best practices for security.

The following are known security issues for DHCP and related protocols:

- **DHCP is an unauthenticated protocol.**

  When a user connects to the network, the user is not required to provide credentials in order to obtain a lease. An unauthenticated user can, therefore, obtain a lease for any DHCP client whenever a DHCP server is available to provide a lease. Any option values that the DHCP server provides with the lease, such as WINS server or DNS server IP addresses, are available to the unauthenticated user. If the DHCP client is identified as a member of a user class or vendor class, the options that are associated with the class are also available.

  Malicious users with physical access to the DHCP-enabled network can instigate a denial-of-service attack on DHCP servers by requesting many leases from the server, thereby depleting the number of leases that are available to other DHCP clients.

  **Recommendations:**

  - Ensure that unauthorized persons do not have physical or wireless access to your network.

  - Enable audit logging for every DHCP server on your network. Regularly check audit log files, and monitor them when the DHCP server receives an unusually high number of lease requests from clients. Audit log files provide the information that you need to track the source of any attacks made against the DHCP server. The default location of audit log files is %windir%\System32\Dhcp. For more information, see Enable DHCP server logging, Audit logging, and Analyzing server log files. You can also check the system event log for explanatory information about the DHCP Server service.

    **Note**

When clients running Microsoft® Windows® XP use 802.1X-enabled local area network (LAN) switches or wireless access points to access the network , authentication occurs before the DHCP server assigns a lease, thereby providing greater security for DHCP.

- **Denial-of-service attacks against the DNS server can be made through the DHCP server.**

  When the DHCP server is configured to act as a DNS proxy server for DHCP clients and to perform DNS dynamic updates, it is possible for a malicious user to perform a denial-of-service attack against both the DHCP server and the DNS server by flooding the DHCP server with requests for leases.

  **Recommendations:**

  - Ensure that unauthorized persons do not have physical or wireless access to your network.

  - Use the DHCP audit logs, located by default at %windir%\System32\Dhcp, to monitor DNS dynamic updates by the DHCP server. The following event IDs are used for DNS dynamic update events:

| Event ID | DHCP Event |
|---|---|
| 30 | DNS dynamic update request to the DNS server |

| | |
|---|---|
| 31 | DNS dynamic update failed |
| 32 | DNS dynamic update successful |

The IP address of the DHCP client is included in the DHCP audit log, providing the ability to track the source of the denial-of-service attack. For more information, see Analyzing server log files and Audit logging.

- **Unauthorized, non-Microsoft DHCP servers can lease IP addresses to DHCP clients.**

  Only DHCP servers running Windows 2000 or Windows Server 2003 can be authorized in Active Directory®. If a DHCP server running Windows 2000 or Windows Server 2003 discovers it is not authorized in Active Directory, the DHCP server stops servicing DHCP clients. Because of this authorization feature, if a malicious or incompetent user installs an unauthorized DHCP server running Windows 2000 or Windows Server 2003 on the organization network, the server cannot assign incorrect or conflicting leases, configure DHCP clients with inaccurate options, or disrupt network services.

  Non-Microsoft DHCP server software does not include the authorization feature that is included in Windows 2000 and Windows Server 2003 DHCP. Because DHCP clients broadcast DHCP discover messages to the nearest DHCP server, if a malicious user installs a non-Microsoft DHCP server on the organization network, nearby DHCP clients will receive incorrect leases that might conflict with the IP addresses assigned to other DHCP clients on the network. In addition, DHCP clients that obtain a lease from the non-Microsoft DHCP server can be configured by the server with option information that is inaccurate. This might reroute network traffic, causing the network to function improperly.

  **Recommendation:**

  - Ensure that unauthorized persons do not have physical or wireless access to your network.

**Additional recommendations**

Before you install and configure DHCP for your network, consider:

- **Restricting who can administer the DHCP service.**

  You must be a member of the Administrators group or the DHCP Administrators group to administer DHCP servers using the DHCP console or the Netsh commands for DHCP. In addition, only members of the Domain Admins group can authorize or unauthorize a DHCP server in Active Directory. Restrict the membership of these groups to the minimum number of users necessary to administer the server.

  If there are users who need read-only access to the DHCP console, add them to the DHCP Users group instead of to the DHCP Administrators group. For more information, see DHCP groups.

# DHCP defined

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP defined**

Dynamic Host Configuration Protocol (DHCP) is an IP standard for simplifying management of host IP configuration. The DHCP standard provides for the use of DHCP servers as a way to manage dynamic allocation of IP addresses and other related configuration details for DHCP-enabled clients on your network.

Every computer on a TCP/IP network must have a unique IP address. The IP address (together with its related subnet mask) identifies both the host computer and the subnet to which it is attached. When you move a computer to a different subnet, the IP address must be changed. DHCP allows you to dynamically assign an IP address to a client from a DHCP server IP address database on your local network:




For TCP/IP-based networks, DHCP reduces the complexity and amount of administrative work involved in reconfiguring computers.

The Microsoft® Windows Server 2003 family provides an RFC-compliant DHCP service you can use to manage IP client configuration and automate IP address assignment on your network.

# Benefits of using DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Benefits of using DHCP**

DHCP provides the following benefits for administering your TCP/IP-based network:

- **Safe and reliable configuration**

  DHCP avoids configuration errors caused by the need to manually type in values at each computer. Also, DHCP helps prevent address conflicts caused by a previously assigned IP address being reused to configure a new computer on the network.

- **Reduces configuration management**

  Using DHCP servers can greatly decrease time spent configuring and reconfiguring computers on your network. Servers can be configured to supply a full range of additional configuration values when assigning address leases. These values are assigned using DHCP options.

  Also, the DHCP lease renewal process helps assure that where client configurations need to be updated often (such as users with mobile or portable computers who change locations frequently), these changes can be made efficiently and automatically by clients communicating directly with DHCP servers.

# How DHCP works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How DHCP works**

DHCP uses a client-server model. The network administrator establishes one or more DHCP servers that maintain TCP/IP configuration information and provide it to clients. The server database includes the following:

- Valid configuration parameters for all clients on the network.
- Valid IP addresses maintained in a pool for assignment to clients, plus reserved addresses for manual assignment.
- Duration of a lease offered by the server. The lease defines the length of time for which the assigned IP address can be used.

With a DHCP server installed and configured on your network, DHCP-enabled clients can obtain their IP address and related configuration parameters dynamically each time they start and join your network. DHCP servers provide this configuration in the form of an address-lease offer to requesting clients.

# DHCP terminology

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP terminology**

| Term | Description |
|---|---|
| scope | A *scope* is the full consecutive range of possible IP addresses for a network. Scopes typically define a single physical subnet on your network to which DHCP services are offered. Scopes also provide the primary way for the server to manage distribution and assignment of IP addresses and any related configuration parameters to clients on the network. |
| superscope | A *superscope* is an administrative grouping of scopes that can be used to support multiple logical IP subnets on the same physical subnet. Superscopes only contain a list of *member scopes* or *child scopes* that can be activated together. Superscopes are not used to configure other details about scope usage. For configuring most properties used within a superscope, you need to configure member scope properties individually. |
| exclusion range | An *exclusion range* is a limited sequence of IP addresses within a scope, excluded from DHCP service offerings. Exclusion ranges assure that any addresses in these ranges are not offered by the server to DHCP clients on your network. |
| address pool | After you define a DHCP scope and apply exclusion ranges, the remaining addresses form the available *address pool* within the scope. Pooled addresses are eligible for dynamic assignment by the server to DHCP clients on your network. |
| lease | A *lease* is a length of time that a DHCP server specifies, during which a client computer can use an assigned IP address. When a lease is made to a client, the lease is *active*. Before the lease expires, the client typically needs to renew its address lease assignment with the server. A lease becomes *inactive* when it expires or is deleted at the server. The duration for a lease determines when it will expire and how often the client needs to renew it with the server. |
| reservation | You use a *reservation* to create a permanent address lease assignment by the DHCP server. Reservations assure that a specified hardware device on the subnet can always use the same IP address. |
| option types | *Option types* are other client configuration parameters a DHCP server can assign when serving leases to DHCP clients. For example, some commonly used options include IP addresses for default gateways (routers), WINS servers, and DNS servers. Typically, these option types are enabled and configured for each scope. The DHCP console also permits you to configure default option types that are used by all scopes added and configured at the server. Most options are predefined through RFC 2132, but you can use the DHCP console to define and add custom option types if needed. |
| options class | An *options class* is a way for the server to further manage option types provided to clients. When an options class is added to the server, clients of that class can be provided class-specific option types for their configuration. For Microsoft® Windows® 2000 and Windows XP, client computers can also specify a class ID when communicating with the server. For earlier DHCP clients that do not support the class ID process, the server can be configured with default classes to use instead when placing clients in a class. Options classes can be of two types: vendor classes and user classes. |

# DHCP tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP tools**
The primary tools that you use to manage DHCP servers are the DHCP console and the Netsh commands for DHCP.

**The Netsh commands for DHCP**
You can type the Netsh commands for DHCP in a command prompt window at the **netsh dhcp>** command prompt, or you can execute the Netsh commands for DHCP in batch files and other scripts. All of the functionality that is available in the DHCP console is available at the **netsh dhcp>** command prompt.

For more information about Netsh commands for DHCP (including syntax, parameters, and examples), see Netsh commands for DHCP. For an example of how to use Netsh commands for DHCP in batch files and scripts, see Netsh DHCP example.

**The DHCP console**
The DHCP console is added to the Administrative Tools folder in Control Panel when you install a DHCP server running Windows Server 2003. The DHCP console appears as a Microsoft Management Console (MMC) snap-in, to further integrate DHCP administration into your total network management.

After you install a DHCP server, you can use the DHCP console or the Netsh commands for DHCP to perform these basic administrative server tasks:

1. Create scopes.
2. Add and configure superscopes and multicast scopes.
3. View and modify scope properties for scopes, such as setting additional exclusion ranges.
4. Activate scopes, multicast scopes, or superscopes.
5. Monitor scope leasing activity by reviewing the active leases for each scope.
6. Create reservations in scopes as needed for DHCP clients that require a permanent IP address for leased use.

In addition, you can also use the DHCP console or the Netsh commands for DHCP to perform the following optional or advanced setup tasks:

- Add new custom default option types.
- Add and configure any user or vendor-defined option classes.
- Further configure other server properties, such as audit logging or BOOTP tables.

The DHCP console also contains enhancements that were suggested by network managers. These include enhanced server performance monitoring, more predefined DHCP option types, dynamic update support for clients running earlier versions of Windows, and detection of unauthorized DHCP servers on your network.

**Notes**

- The Microsoft® Windows Server 2003 family includes new features beyond those provided in earlier versions. For more information, see New features for DHCP.
- The DHCP console provides new ways to perform familiar DHCP administrative tasks previously performed in Windows NT® 4.0 using DHCP Manager. For more information, see New ways to do DHCP tasks.
- For a simple list of Netsh commands for DHCP (not including syntax, parameters, and examples), see Managing DHCP

from the command line. To learn more about using the Netsh command-line tool, see Use DHCP Command-line Tools.

- DHCP servers running Windows 2000 cannot be authorized in Active Directory® using the Windows Server 2003 administrative tools (the DHCP console and the Netsh commands for DHCP) unless Windows 2000 Server Service Pack 2 is installed.

# Multicast address allocation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multicast address allocation**
DHCP servers running Microsoft® Windows Server 2003 allow the assignment of multicast addresses, in addition to the assignment of unicast addresses.

A proposed Internet Engineering Task Force (IETF) standard defines a client/server protocol, the Multicast Address Dynamic Client Allocation Protocol (MADCAP), for defining multicast address allocation between supporting servers and clients. The primary benefit of the proposed MADCAP standard is that you can use it to leverage your existing network infrastructure for assigning multicast addresses in the same way as other IP addresses are currently assigned using DHCP.

Typical applications for multicast are conferencing and audio, which usually require users to specially configure multicast addresses. Unlike IP broadcasts, which are received by all computers or other hosts on the network, a multicast address is a group of computers, using the concept of a group membership to identify the computers to which the message is to be sent.

The multicast address allocation feature has two parts:

1. A MADCAP server, which distributes multicast addresses.

   To use this feature, the administrator first configures multicast scopes and the corresponding multicast IP ranges on the server, through the DHCP console. Once these scopes are configured and activated, the DHCP Server service can provide multicast addresses in the same way it normally provides unicast IP addresses where DHCP scopes are used.

2. Client-side application programming interfaces (APIs), which MADCAP clients can use to request, renew, or release multicast addresses.

   Applications that implement MADCAP client support can call these APIs to request a multicast address from a given multicast scope at the MADCAP server. The underlying MADCAP implementation uses protocol messages similar to the DHCP messaging format to negotiate the multicast address lease between the client and server.

For more information, see Using multicast scopes.

**Notes**

- The standards for multicast address allocation are managed by the Multicast address-allocation (MALLOC) working group of the IETF.

- You can obtain RFCs from the Request for Comments Web site.

  This Web site is currently maintained by members of the Information Sciences Institute (ISI), who publish a classified listing of all RFCs. RFCs are classified as one of the following: approved Internet standard, proposed Internet standard (circulated in draft form for review), Internet best practices, or For Your Information (FYI) documents.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# New features for DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features for DHCP**

Microsoft® Windows Server 2003 Dynamic Host Configuration Protocol (DHCP) provides the following new features:

- **DHCP client alternate configuration**

  With DHCP client alternate configuration, you can easily move a computer between two or more networks, one configured with static IP addresses and one or more configured with DHCP, without needing to reconfigure network adapter parameters, such as IP address, subnet mask, default gateway, preferred and alternate Domain Name Service (DNS) servers, and Windows Internet Name Service (WINS) servers. For more information, see DHCP client alternate configuration and Use DHCP client alternate configuration.

- **DHCP database backup and restore**

  You can back up and restore the DHCP database using the **Backup** and **Restore** commands in the DHCP console. For more information about the **Backup** command, see Backing up the DHCP database and Back up the DHCP database. For more information about the **Restore** command, see Restoring server data and Restore the DHCP database.

The following features were new for Windows 2000 Server and are also provided with the Windows Server 2003 family:

- **Local groups for providing limited server and console access to DHCP administrators and users**

  Two local administrative groups are created when a DHCP server is installed on a server computer: the DHCP Users and DHCP Administrators groups. The DHCP Users group can be used to provide read-only console access to the server. Where used, this group grants members the ability to view, but not modify, server data.

  The DHCP Administrators group provides a higher level of administrative access. Members added to this group can fully manage and modify any data at the DHCP server, but are not provided full unlimited administrative access, such as if the local Administrators group were used instead to provide administrative access to the applicable DHCP server. For more information, see DHCP groups.

- **Automatic assignment of IP addresses**

  For Windows 98, Windows 2000, Windows XP, and the Windows Server 2003 family, DHCP-enabled clients are automatically assigned a temporary IP configuration if a DHCP server is unavailable to provide a lease when the client starts on a network. For DHCP clients running these versions of Windows, the client tries to contact a DHCP server in the background every five minutes to eventually obtain its valid IP address lease.

  Automatic assignment is always transparent to users, who are not prompted if the client fails to obtain a lease from a DHCP server. Addresses are self-assigned from a network address range that is reserved for private TCP/IP usage and not used on the Internet.

- **Enhanced performance monitoring and server reporting capabilities**

  DHCP is critical to the success of your network infrastructure. Without working DHCP servers, IP clients can lose some or all of their ability to access or use your network. Because many experienced network administrators realize the importance of closely monitoring DHCP servers, several changes have been made to assist in this area.

  Performance monitoring counters are available on computers running Windows 2000 or Windows Server 2003 operating systems to help you monitor DHCP server performance on your network. For more information, see Monitoring DHCP server performance.

- **Expanded scope support for multicast scopes and superscopes**

Microsoft DHCP server supports additional scopes that you can use to streamline your ongoing administration of IP address configurations.

New multicast scopes make it possible for DHCP-enabled clients to lease Class D-type IP addresses (224.0.0.0 to 239.255.255.255) for participation in multicast groups. For more information, see Using multicast scopes.

Superscopes (a late addition to Windows NT® Server 4.0) are useful for creating an administrative grouping of member scopes. Superscopes can be helpful when you want to renumber or expand the IP address space on your network without disturbing currently active scopes. For more information, see Using superscopes.

- **Support for user-specified and vendor-specified option classes**

  You can use this feature to separate and distribute appropriate options for clients with similar or special configuration needs.

  For example, you might assign all DHCP-enabled clients on the same floor of your building to the same option class. You could use this class (configured with the same DHCP Class ID value) to distribute other options data during the lease process, overriding any scope or global default options. In this way, options appropriate to a set of class-member clients in the same network location (such as a specific default gateway or parent domain name) are applied as class-specific options. For more information, see Using option classes.

- **Integration of DHCP with DNS**

  A DHCP server running Windows 2000 or Windows Server 2003 can enable dynamic updates in the DNS namespace for any of its clients that support these updates. Scope clients can then use DNS dynamic update protocol to update their host name-to-address mapping information (which is stored in *zones* on the DNS server) whenever changes occur to their DHCP-assigned address. For more information, see Using DNS servers with DHCP.

- **Detection of unauthorized DHCP servers through the use of Active Directory® integration**

  When a DHCP server is unintentionally started on a network, it can cause a variety of network problems. To prevent this scenario from occurring, Windows 2000 Server and the Windows Server 2003 family provide a method for authorizing a DHCP server as well as a means to detect and shut down unauthorized servers.

  When a domain member DHCP server running Windows 2000 or Windows Server 2003 attempts to start on the network, Active Directory is queried, and the IP address of the server is compared to the list of authorized DHCP servers. If a match is found, the server computer is authorized as a DHCP server. If a match is not found, the following occurs:

  - The server is not authorized in Active Directory.
  - The server is identified as an unauthorized server.
  - The server stops leasing IP addresses to DHCP clients.

  In order to use DHCP server authorization, you must deploy Active Directory and the DHCP service on a server running Windows 2000 or Windows Server 2003 . Other DHCP servers do not support this feature.

  For more information, see Authorizing DHCP servers.

- **Dynamic support for BOOTP clients.**

  The DHCP service provides additional support for BOOTP clients in large enterprise networks through the addition of *dynamic BOOTP*.

  Dynamic BOOTP is an extension of the BOOTP protocol that permits the DHCP server to configure BOOTP clients without having to use explicit fixed address configuration. This feature provides easier administration of large BOOTP networks by allowing automatic distribution of IP addresses (much the same way that DHCP does) without having to change client

side behavior.

# Planning DHCP networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning DHCP networks**

This topic covers the following DHCP network planning issues.

- How to determine the number of DHCP servers to use
- How to support DHCP clients on additional subnets
- Planning for routed DHCP networks
- Other planning considerations for enterprise networks

**How to determine the number of DHCP servers to use**

Because there is no fixed limit to the maximum number of clients a DHCP server can service or to the number of scopes you can create on a DHCP server, the primary factors to consider when you determine the number of DHCP servers to use are network architecture and server hardware. For example, in a single subnet environment, only one DHCP server is necessary, although you may want to use two servers or deploy a DHCP server cluster for increased fault tolerance. In multiple subnet environments, routers must forward DHCP messages between subnets, so router performance can affect your DHCP service. In both cases, DHCP server hardware affects service to clients.

For more information on deploying a DHCP server cluster, see Cluster support for DHCP servers.

When you determine the number of DHCP servers to use, consider the following:

- The location of routers on the network and whether you want a DHCP server on each subnet.

  When you extend the use of a DHCP server across more than one network, you often need to configure additional DHCP relay agents and, in some cases, use superscopes as well.

- The transmission speed between the segments for which DHCP service is provided.

  If you have slower WAN links or dial-up links, you may need a DHCP server on both sides of these links to service clients locally.

- The speed of the server disk drives and the amount of random access memory (RAM) installed in the DHCP server computer.

  In order to maximize DHCP server performance, use the fastest disk drives and the most RAM possible. Carefully evaluate disk access time and average times for disk read and write operations when you plan for your DHCP server hardware needs.

- Practical constraints based on the IP address class selected for use and other server configuration details.

You can test your DHCP servers before deployment on the organization network in order to determine the limitations and abilities of your hardware and to see whether network architecture, traffic, and other factors affect DHCP server performance. Hardware and configuration tests also allow you to determine the number of scopes to configure at each server.

To provide a general idea of DHCP server performance, a DHCP server running Microsoft® Windows Server 2003 was run in a test lab environment, and a custom stress application was used against the server. The details of this lab test can help you configure your tests and determine how many DHCP servers to use on your network:

**Server and network configuration**

- Processors: Two x86 Family 6 Model 7 Stepping 3 GenuineIntel ~498 megahertz (MHz)

- Total physical memory: 256.00 megabytes (MBs)
- Network adapters: Three Ethernet 802.3 100 megabits per second (mbps)
- Subnets serviced: Six, four of which are separated from the test server by routers running the DHCP Relay Agent service.
- Operating system: Windows Server 2003, Enterprise Edition
- Number of scopes: 5,155
- DHCP database size at maximum load: 2 gigabytes (GBs)
- Additional factors: Several thousand exclusion ranges, option values, and reservations are configured in the scopes on the server.

**Test details**
The DHCP server was hit with both valid and invalid DHCP client lease and renewal requests from a client/attack simulation stress application on six subnets for 1,152 hours (48 days).

**Test results**
The DHCP server provided clients with the following service during the test:

| DHCP server function | Message and lease volume handled (1,152 hours) |
|---|---|
| Lease assignments | 68,412,059 |
| DHCP discover messages | 20,039,592 |
| DHCP offer messages | 20,039,253 |
| DHCP request messages | 57,559,426 |
| DHCP acknowledgement messages | 57,470,934 |
| DHCP negative acknowledgement messages sent by the server | 484,012 |
| DHCP decline messages | 190,901 |
| Lease releases | 0 |

**Notes**

- These test results are intended to provide a general idea of DHCP server performance capacity. Tests were performed on average hardware and do not imply any limitation of the DHCP service. In addition, the large size of the test DHCP server database (2 GBs) is extremely unusual and was produced due to the high volume of DHCP traffic that the test application generated. In real world environments, typical database sizes are in the tens of MBs or less.
- When you add a large number of scopes to the server, be aware that each scope creates a corresponding need for additional, incremental increases to the amount of disk space used for the DHCP server registry and for the server paging file. For more information, see Change the size of the virtual memory paging file.
- DHCP servers running Windows Server 2003 provide performance monitoring tools that you can use to test and monitor your servers. For more information, see Monitoring DHCP server performance.

**A DHCP planning alternative: configuring standby servers**
Most networks need one primary online DHCP server and one other DHCP server acting as a secondary or backup server. If you choose not to implement two DHCP servers using the 80/20 rule for balancing scopes but want to continue to provide a measure of potential fault tolerance, you might consider implementing a backup or hot standby DHCP server as an alternative.

For more information, see the section "Using the 80/20 rule for scopes" in Configuring scopes.

In the hot standby configuration, the standby DHCP server is another server computer that is installed and configured identically to your primary DHCP server. In the case of the standby, the only difference is that the standby server and its scopes are not activated for use under normal conditions. Duplicate scopes are configured but not activated except when critically needed, such as to substitute for the primary DHCP server because it has stopped or been taken offline for an extended period of time.

Because the hot standby solution requires special attention to its configuration and also manual administration for ensuring fall-over transition for DHCP clients to use it, it is less recommended as a planning alternative than the use of two to three DHCP servers that balance active scope use.

**Supporting additional subnets**
For DHCP service to support additional subnets on your network, you must first determine whether the routers used to connect adjoining subnets can support relaying of BOOTP and DHCP messages. If routers cannot be used for DHCP and BOOTP relay, you can set up either of the following for each subnet:

- A computer running Windows NT® Server 4.0, Windows 2000 Server, or a Windows Server 2003 operating system configured to use the DHCP Relay Agent component.

  This computer simply forwards messages back and forth between clients on the local subnet and a remote DHCP server, using the IP address of the remote server. The DHCP Relay Agent service is available only on computers running Windows NT Server 4.0, Windows 2000 Server, or Windows Server 2003 operating systems.

- A computer running a Windows Server 2003 operating system configured as a DHCP server for the local subnet.

  This server computer must contain and manage scope and other address-configurable information for the local subnet it serves.

For more information, see Relay agent configurations.

**Planning for routed DHCP networks**
In routed networks that use subnets to divide network segments, planning options for DHCP services must observe some specific requirements for a full implementation of DHCP services to function. These requirements include:

- One DHCP server must be located on at least one subnet in the routed network.

- For a DHCP server to support clients on other remote subnets separated by routers, a router or remote computer must be used as a DHCP and BOOTP relay agent to support forwarding of DHCP traffic between subnets.

The following illustration shows a simple routed network with DHCP implemented.




**Enterprise planning considerations**
For an enterprise DHCP network, you should:

- Plan the physical subnets of the network and relative placement of DHCP servers.

- Specify the DHCP option types and their values to be predefined per scope for the DHCP clients.

  This can include planning for scopes based on the needs of particular groups of users. For example, a unit that frequently moves computers to different locations, like a marketing group that uses portable computers docked at different stations, shorter lease durations can be defined for the related scopes. This approach collects IP addresses that are changed

frequently and disposed of and returns them to the pool of available addresses that can be used for new lease offerings.

- Recognize the impact that slower links have on your WAN environment.

  Place DHCP, WINS, and DNS servers to maximize response time and minimize low-speed traffic.

As one example of planning for a large enterprise network, the segmenting of the WAN into logical subnets can match the physical structure of the internetwork. Then one IP subnet can serve as the backbone and, this backbone in relation to each physical subnet maintains a separate IP subnet address.

# Understanding DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding DHCP**
This section covers:

- DHCP servers
- DHCP options
- DHCP groups
- Understanding user and vendor classes
- BOOTP and DHCP
- The DHCP database
- Audit logging
- DHCP client alternate configuration
- DHCP/BOOTP Relay Agents

# DHCP servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP servers**

Configuring DHCP servers for a network provides the following benefits:

- The administrator can assign and specify global and subnet-specific TCP/IP parameters centrally for use throughout the entire network.
- Client computers do not require manual TCP/IP configuration.
- When a client computer moves between subnets, its old IP address is freed for reuse. The client reconfigures its TCP/IP settings automatically when the computer is restarted in its new location.
- Most routers can forward DHCP and BOOTP configuration requests, so DHCP servers are not required on every subnet in the network.

**How clients use servers**

A computer running Windows XP becomes a DHCP client if **Obtain an IP address automatically** is selected in its TCP/IP properties. When a client computer is set to use DHCP, it accepts a lease offer and can receive from the server the following:

- Temporary use of an IP address known to be valid for the network it is joining.
- Additional TCP/IP configuration parameters for the client to use in the form of *options* data.

Also, if conflict detection is configured, the DHCP server attempts to ping each available address it intends to offer prior to presenting the address in a lease offer to a client. This ensures that each IP address offered to clients is not already in use by another non-DHCP computer that uses manual TCP/IP configuration.

**How servers provide optional data**

In addition to an IP address, DHCP servers can be configured to provide optional data to fully configure TCP/IP for clients. Some of the most common DHCP option types configured and distributed by the DHCP server during leases include:

- Default gateways (routers), which are used to connect a network segment to other network segments.
- Other optional configuration parameters to assign to DHCP clients, such as IP addresses for the DNS servers or WINS servers that the client can use in resolving network host names.

For more information, see DHCP options.

# DHCP options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP options**

DHCP provides an internal framework for passing configuration information on to clients on your network. Configuration parameters and other control information are carried in tagged data items stored within protocol messages exchanged between the DHCP server and its clients. These data items are called *options*.

Most standard DHCP options are currently defined in Request for Comments (RFCs) published by the Internet Engineering Task Force (IETF). The full set of standard DHCP options are described specifically in RFC 2132, "DHCP Options and BOOTP Vendor Extensions."

All DHCP options mentioned in RFC 2132 are predefined for you to configure and use at any DHCP server running Windows Server 2003 . If needed, you can also use the DHCP console to define new DHCP options at each server.

Even though most DHCP servers can assign many options, most DHCP clients are typically designed to request or support only a subset of the full RFC-specified standard options set.

- For more information about DHCP options, see "DHCP Options" at the Microsoft Windows Resource Kits Web site.

**How options are applied**

Options can be managed using different levels assigned for each managed DHCP server, including:

- **Server options** These options are applied for all scopes defined at a DHCP server.
- **Scope options** These options are applied specifically to all clients that obtain a lease within a particular scope.
- **Class options** These options are applied only to clients that are identified as members of a specified user or vendor class when obtaining a lease.
- **Reservation options** These options apply only for a single reserved client computer and require a reservation to be used in an active scope.

For more information, see Setting Up Options.

# DHCP groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP groups**
When you install the DHCP Server service, two domain local groups are created: DHCP Users and DHCP Administrators.

**DHCP Users**
Members of the DHCP Users group have read-only DHCP console access to the server, which allows DHCP Users to view, but not to modify, server data, including DHCP server configuration, registry keys, DHCP log files, and the DHCP database. DHCP Users cannot create scopes, modify option values, create reservations or exclusion ranges, or modify the DHCP server configuration in any other way.

**DHCP Administrators**
Members of the DHCP Administrators group can view and modify any data at the DHCP server. DHCP Administrators can create and delete scopes, add reservations, change option values, create superscopes, or perform any other activity needed to administer the DHCP server, including export or import of the DHCP server configuration and database. DHCP Administrators perform these tasks using the Netsh commands for DHCP or the DHCP console. For more information, see DHCP tools.

Members of the DHCP Administrators group do not have unlimited administrative rights. For example, if a DHCP server is also configured as a DNS server, a member of the DHCP Administrators group can view and modify the DHCP configuration but cannot modify DNS server configuration on the same computer.

Because members of the DHCP Administrators group have rights on the local computer only, DHCP Administrators cannot authorize or unauthorize DHCP servers in Active Directory. Only members of the Domain Admins group can perform this task. If you want to authorize or unauthorize a DHCP server in a child domain, you must have enterprise administrator credentials for the parent domain. For more information about authorizing DHCP servers in Active Directory, see Authorizing DHCP servers and Authorize a DHCP server in Active Directory.

**Note**

- To be logged on as an enterprise administrator, you must log on using a member account in the Enterprise Admins group. By default, this is done by logging on as local administrator at the first domain controller created in your enterprise.

**Using groups to administer DHCP servers in a domain**
When you add a user or group to a DHCP Users or DHCP Administrators group on a DHCP server, the rights of the DHCP group member do not apply to all of the DHCP servers in the domain. The rights apply only to the DHCP service on the local computer. In order to assign rights to a user or group that apply to all of the DHCP servers in the domain, you can add the user or group to the DHCP Administrators group on each DHCP server in the domain. This approach is practical only on a small network with one or two DHCP servers.

An alternative approach is to add the user or group to the global group called Domain Admins. Members of the Domain Admins group can administer all of the DHCP servers in the domain, however the administrative rights of Domain Admins are not limited to the DHCP service on each DHCP server. Domain Admins have complete administrative rights on all of the computers in the domain. For this reason, it is recommended that you do not add users or groups to the Domain Admins group solely because you want to assign administrative rights for the DHCP service on all of the DHCP servers in the domain.

If you have more than two DHCP servers on your network, you can use the following steps to most effectively add members to a DHCP group:

1. In the Active Directory Users and Computers console, create a group with universal or global scope. For more information, see Create a new group.

2. Add users or groups to the new group you created. For more information, see Add a member to a group.

3. At each DHCP server in the domain, add the new group (that you created in the first step) to either the DHCP Users group or the DHCP Administrators group, depending on the rights and permissions that you want group members to have.

For more information about how to add members to groups and other related tasks, see Manage Server Access.

For more information about groups, see Group scope and Default groups.

For more information about delegating administrative credentials, see Delegate ability to authorize DHCP servers to a non-enterprise administrator.

# Understanding user and vendor classes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding user and vendor classes**

There are significant differences in the way in which vendor and user classes are available for use by DHCP clients and servers. The following table compares and contrasts the features of these classes.

| User classes | Vendor classes |
|---|---|
| Part of proposed DHCP standard drafts.<br><br>For updated information, obtain the latest draft from the Request for Comments Web site. | Part of accepted DHCP standard drafts (RFC 2131, 2132). |
| Clients identify user class membership with User class information. | Clients identify vendor class membership by using the Vendor class identifier option. |
| User class IDs can be set and viewed at DHCP clients running Windows 2000, Windows XP, or Windows Server 2003 operating systems using the **ipconfig** command.<br><br>For more information, see:<br><br>• Set DHCP class ID information at a client computer<br><br>• Show DHCP class ID information at a client computer | Vendor class IDs are configured internally by DHCP client system vendors, such as Microsoft and other vendors.<br><br>When adding a vendor class at the server, you need to specify class data at the server that matches identically the class ID used by same-vendor clients. For more information, contact the applicable vendor. |
| With user classes, you can take advantage of predefined classes that can be used to support specialized groups of clients, such as BOOTP or Routing and Remote Access clients. | With vendor classes, Microsoft has implemented vendor classes to the provided classification of its DHCP clients by Windows 98, Windows 2000, or Microsoft DHCP clients which share the use of the Microsoft-reserved vendor prefix ("MSFT") in any class-related data used. |
| User classes are typically created for administrative reasons, such as to identify all DHCP clients in a specific site or location.<br><br>For example, all computers and printers located on the same floor in an office building might be assigned or configured as part of the same user class. | Vendor classes are typically created for the enhancement of vendor-specific DHCP functionality.<br><br>For example, Microsoft provides several additional vendor-specific options, such as the ability to disable NetBIOS over TCP/IP at DHCP clients running Windows 2000 or Windows XP. |
| User classes are most helpful to users for managing DHCP option assignments based on their needs.<br><br>If your DHCP network is large enough to benefit from user classes, you can freely customize your DHCP environments with them. In most cases, the user class provides a way to assign and override standard DHCP option values set at either the server, scope, or reserved client level. | Vendor classes are most helpful to vendors for managing DHCP option assignments based on vendor-specific needs without disturbing other non-vendor DHCP clients.<br><br>With a vendor class, options are passed to clients using the Vendor specific information option, which encapsulates the options it sends. This helps non-vendor clients to recognize the information as vendor-specific and ignore and discard it as needed. |

**Notes**

- For more information about DHCP options, see "DHCP Options" at the Microsoft Windows Resource Kits Web site.
- For more information, see Using option classes or Create a new user or vendor class.
- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# BOOTP and DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**BOOTP and DHCP**

The bootstrap protocol (BOOTP) is a host configuration protocol developed before DHCP. DHCP improves on BOOTP and resolves specific limitations BOOTP had as a host configuration service. RFC 951 defines BOOTP.

**BOOTP/DHCP similarities**

Because of the relationship between BOOTP and DHCP, both protocols share some defining characteristics. The common elements include:

- **The format structure each uses to exchange messages between the server and clients**

  BOOTP and DHCP use nearly identical request messages (sent by clients) and reply messages (sent by servers). Messages in either of these protocols use a single User Datagram Protocol (UDP) datagram of 576 bytes to enclose each protocol message. Message headers are the same for both BOOTP and DHCP with one exception: the final message header field used to carry optional data. For BOOTP, this optional field is called the *vendor-specific area* and is limited to 64 octets. For DHCP, this area is called the *options* field and can carry up to 312 octets of DHCP options information.

- **Use of well-known UDP ports for client/server communication**

  Both BOOTP and DHCP use the same reserved protocol ports for sending and receiving messages between servers and clients. Both BOOTP and DHCP servers use UDP port 67 to listen for and receive client request messages. BOOTP and DHCP clients typically reserve UDP port 68 for accepting message replies from either a BOOTP server or DHCP server.

  Because DHCP and BOOTP messages use nearly identical format types and packet structures, and typically use the same well-known service ports, BOOTP or DHCP relay agent programs usually treat BOOTP and DHCP messages as essentially the same message type, without differentiating between them.

- **IP address distribution as an integral part of configuration service**

  Although both BOOTP and DHCP allocate IP addresses to clients during startup, they use different methods of allocation. BOOTP typically provides fixed allocation of a single IP address for each client, permanently reserving this address in the BOOTP server database. DHCP typically provides dynamic, leased allocation of available IP addresses, reserving each DHCP client address temporarily in the DHCP server database.

**BOOTP/DHCP differences**

There are significant differences in the way in which BOOTP and DHCP perform host configuration. The following table compares and contrasts the features of the two protocols that vary.

| BOOTP | DHCP |
|---|---|
| Designed prior to DHCP. | Designed after BOOTP. |
| Intended to configure diskless workstations with limited boot capabilities. | Intended to configure frequently relocated networked computers (such as portables) that have local hard drives and full boot capabilities. |
| Dynamic BOOTP has a default 30-day expiration on IP address leases. | DHCP has a default eight-day expiration on IP address leases. |
| Supports a limited number of client configuration parameters called *vendor extensions*. | Supports a larger and extensible set of client configuration parameters called *options*. |

| | |
|---|---|
| Describes a two-phase bootstrap configuration process, as follows:<br>● Clients contact BOOTP servers to perform address determination and boot file name selection.<br>● Clients contact Trivial File Transfer Protocol (TFTP) servers to perform file transfer of their boot image. | Describes a single-phase boot configuration process whereby a DHCP client negotiates with a DHCP server to determine its IP address and obtain any other initial configuration details it needs for network operation. |
| BOOTP clients do not rebind or renew configuration with the BOOTP server except when the system restarts. | DHCP clients do not require a system restart to rebind or renew configuration with the DHCP server. Instead, clients automatically enter a rebinding state at set timed intervals to renew their leased address allocation with the DHCP server. This process occurs in the background and is transparent to the user. |

# The DHCP database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The DHCP database**

There is no set limit to the number of records that a DHCP server can store. The size of the database is dependent upon the number of DHCP clients on the network. The DHCP database grows over time as a result of clients starting and stopping on the network.

The size of the DHCP database is not directly proportional to the number of active client lease entries. Over time, as some DHCP client entries become obsolete and are deleted, there remains some unused space.

To recover the unused space, the DHCP database is compacted. Starting with Windows NT Server 4.0, dynamic database compaction occurs on DHCP servers as an automatic background process during idle time or after a database update.

**Note**

- When using or upgrading DHCP databases used with earlier versions of Windows NT Server, the DHCP database requires conversion. In the Windows Server 2003 family, if a previous copy of the DHCP server database is detected during an upgrade, it is automatically converted to the current database version and format. For larger databases, this process can increase the time needed to complete the installation.

**DHCP database files**

The DHCP server database in Windows Server 2003 family uses the Exchange Server JET storage engine. When you install the DHCP service, the files shown in the following table are automatically created in the *Systemroot*\System32\Dhcp directory.

**DHCP service database files**

| File | Description |
|---|---|
| Dhcp.mdb | The DHCP server database file. |
| Dhcp.tmp | A temporary file used by the DHCP database as a swap file during database index maintenance operations. This file sometimes remains in the *Systemroot*\System32\Dhcp directory after a system failure. |
| J50.log and J50#####.log | A log of all database transactions. This file is used by the DHCP database to recover data when necessary. |
| J50.chk | A checkpoint file. |

**Important**

- The J50.log file, J50#####.log file, Dhcp.mdb file, and Dhcp.tmp file should not be removed or altered.

The DHCP server database is a dynamic database that is updated as DHCP clients are assigned or as they release their TCP/IP configuration parameters. Because the DHCP database is not a distributed database like the WINS server database, maintaining the DHCP server database is less complex.

The DHCP database and related registry entries are automatically backed up at a specific interval (60 minutes by installation default). You can change this installation default by changing the value of **BackupInterval** in the following registry key:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DHCPServer\Parameters**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

# Audit logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Audit logging**
DHCP servers running Windows Server 2003 include several logging features and server parameters that provide enhanced auditing capabilities. You can specify the following features:

- The directory path in which the DHCP server stores audit log files. DHCP audit logs are located by default at %windir%\System32\Dhcp.
- A maximum size restriction (in megabytes) for the total amount of disk space available for all audit log files created and stored by the DHCP service.
- An interval for disk checking that is used to determine how many times the DHCP server writes audit log events to the log file before checking for available disk space on the server.
- A minimum size requirement (in megabytes) for server disk space that is used during disk checking to determine if sufficient space exists for the server to continue audit logging.

### Notes

- You can selectively enable or disable the audit logging feature at each DHCP server. For more information, see Enable DHCP server logging.
- Only the directory path in which the DHCP server stores audit log files can be modified using the DHCP console. To do so, first select the applicable DHCP server in the console tree. On the **Action** menu, click **Properties**. Next, click the **Advanced** tab and edit **Audit log file path** as necessary.
- Other audit logging parameters described above are adjusted through registry-based configuration changes. For more information about the Windows Server 2003 DHCP server registry, see "Dynamic Host Configuration Protocol" at the Microsoft Windows Resource Kits Web site.

**How audit logging works**
The audit logging behavior discussed in this section applies only to Windows 2000 DHCP and Windows Server 2003 DHCP. It replaces the previous DHCP logging behavior used in earlier versions of Windows NT Server, which do not perform audit checks and use only a single log file named *Dhcpsrv.log* for logging service events.

The formatted structure of DHCP server logs and the level of reporting maintained for audited logging are the same as in earlier DHCP server versions provided with Windows NT Server. For more information, see Analyzing server log files.

**Naming audit log files**
The DHCP Server service bases the name of the audit log file on the current day of the week, as determined by checking the current date and time at the server.

For example, when the DHCP server starts, if the current date and time are:

Monday, April 7, 2003, 04:56:42 P.M.

then the server audit log file is named:

DhcpSrvLog-Mon

**Starting a daily audit log**
When a DHCP server starts or a new day begins (when the local time on the computer is 12:00 A.M.), the server writes a header message in the audit log file, indicating that logging has started. Then, depending on whether the audit log file is a new or existing file, the following actions occur:

- If the file already existed without modification for more than a day, it is overwritten.
- If the file already existed but was modified within the previous 24 hours, the file is not overwritten. Instead, new logging activity is appended to the end of the existing file.

**Disk checks**

After audit logging starts, the DHCP server performs disk checks at regular intervals, to ensure both the ongoing availability of server disk space and that the current audit log file does not become too large or that log-file growth is not occurring too rapidly.

The DHCP server performs a full disk check whenever either of the following conditions occurs:

- A set number of server events are logged.
- When the date changes on the server computer.

By default, the DHCP server performs a periodic disk space check for every 50 events it writes to the audit log. A date change can also be detected by the DHCP server when the server computer reaches 12:00 A.M. on its locally set clock.

Each time a disk check is completed, the server determines whether disk space is filled. The disk is considered full if either of the following conditions is true:

- **Disk space on the server computer is lower than the required minimum amount for DHCP audit logging.**

  By default, if the amount of disk space remaining on the server disk reaches less than 20 megabytes, audit logging is halted.

- **The current audit log file is larger than one-seventh (1/7) of the maximum allotted space or size for the combined total of all audit logs currently stored on the server.**

  At the time of the disk check, the DHCP server compares the exact size (in megabytes) of the current audit log file with a value obtained by dividing the current value for the maximum number of log files the server permits to be stored simultaneously before overwriting and discarding of older log files. By default, seven is the maximum number of log files the server permits to be stored, one for each day of the week. Assuming the default is set, the largest size that the current audit log file can reach is one megabyte.

In either case, if the disk is full, the DHCP server closes the current file and ignores further requests to log audit events until either 12:00 A.M. or until disk status is improved and the disk is no longer full.

Even if audit logged events are ignored because of a disk full condition, the DHCP server continues disk checking every 50 events (or the currently set interval) to determine whether disk conditions have improved. If subsequent disk checks determine that the required amount of server disk space is available, the DHCP server reopens the current log file and resumes logging.

**Ending a daily audit log**

At 12:00 A.M. local time on the server computer, the DHCP server closes the existing log and moves to the log file for the next day of the week. For example, if the day of the week changes at 12:00 A.M. from Wednesday to Thursday, the log file named DhcpSrvLog-Wed is closed, and the file named DhcpSrvLog-Thu is opened and used for logging events.

# DHCP client alternate configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP client alternate configuration**

With DHCP client alternate configuration, you can easily move a computer between two or more networks, one configured with static IP addresses and one or more configured with DHCP. Alternate configuration provides simplified computer migration (for example, a laptop) between networks without requiring that you reconfigure network adapter parameters such as IP address, subnet mask, default gateway, preferred and alternate Domain Name Service (DNS) servers, and Windows Internet Name Service (WINS) servers.

**Configuration options**

When you configure TCP/IP properties for a local area network connection, you have the following options:

- **Static IP address configuration**

  When you click**Use the following IP address**, you can provide values for static IP address settings, such as IP address, subnet mask, default gateway, preferred and alternate DNS servers, and WINS servers. However, if you click **Obtain an IP address automatically** to change the configuration of your network adapter to a DHCP client configuration, all static IP address settings are lost. Additionally, if you move the computer and configure it for another network, when you return to the original network you will need to reconfigure the computer with the original static IP address settings.

- **Dynamic IP address configuration without alternate configuration**

  When you click **Obtain an IP address automatically**, your computer acts as a DHCP client and obtains IP address, subnet mask, and other configuration parameters from your network DHCP server. If the DHCP server is unavailable, the network adapter is configured using IP autoconfiguration.

- **Dynamic IP address configuration with alternate configuration**

  When you click **Obtain an IP address automatically**, click the **Alternate Configuration** tab, and type an alternate configuration, you can move your computer between one statically configured network (such as a home network) and one or more dynamically configured networks (such as a corporate network) without changing any settings. If the DHCP server is unavailable (for example, when your computer is connected to your home network), the network adapter is automatically configured with your alternate configuration, and the computer functions correctly on the network. When you move the computer back to the dynamically configured network and the DHCP server is available, the network adapter is automatically configured with the dynamic configuration assigned by the DHCP server. The alternate configuration is used only when the DHCP client cannot locate a DHCP server.

**DHCP server discover attempts**

If you use DHCP without an alternate configuration and the DHCP client cannot locate a DHCP server, IP autoconfiguration is used to configure the network adapter. The DHCP client continues to attempt to discover a DHCP server on the network every five minutes. If a DHCP server is found, the network adapter is assigned a valid DHCP IP address lease.

If you use DHCP with an alternate configuration, and the DHCP client cannot locate a DHCP server, the alternate configuration is used to configure the network adapter. Normally, there are no additional discovery attempts. However, a DHCP server discovery attempt will occur if:

- The network adapter is disabled and then enabled again.
- Media (such as network cabling) is disconnected and then reconnected.
- The TCP/IP settings for the adapter are changed, and DHCP remains enabled after these changes.

If a DHCP server is found, the network adapter is assigned a valid DHCP IP address lease.

**Notes**

- It is possible that a static alternate configuration could conflict with the configuration of another computer on a network. For example, a client using alternate configuration might have the same IP address as another computer on the network. If this occurs, Address Resolution Protocol (ARP) detects the conflict, and the network adapter of the computer configured with the alternate configuration is automatically set to 0.0.0.0. No additional attempts are made to discover a DHCP server, obtain a lease, or use the static alternate configuration.

- When switching from an alternate configuration to a configuration using IP autoconfiguration, the network adapter is configured with IP autoconfiguration settings (169.254.*x.x*), and a DHCP server discovery attempt is initiated at the same time. If the discovery is successful, the adapter is assigned a valid DHCP lease.

- For more information, see Use DHCP client alternate configuration.

# DHCP/BOOTP Relay Agents

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP/BOOTP relay agents**
This section covers:

- Understanding relay agents
- Relay agent configurations
- Relay agent design issues

# Understanding relay agents

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding relay agents**
A relay agent is a small program that relays DHCP/BOOTP messages between clients and servers on different subnets. DHCP/BOOTP relay agents are part of the DHCP and BOOTP standards and function according to the Request for Comments (RFCs), standard documents that describe protocol design and related behavior.

**DHCP/BOOTP router support**
In TCP/IP networking, routers are used to interconnect hardware and software used on different physical network segments called *subnets* and forward IP packets between each of the subnets. To support and use DHCP service across multiple subnets, routers connecting each subnet should comply with DHCP/BOOTP relay agent capabilities described in RFC 1542.

To be compliant with RFC 1542 and provide relay agent support, each router must be able to recognize BOOTP and DHCP protocol messages and process (relay) them appropriately. Because routers interpret DHCP messages as BOOTP messages (such as a UDP message sent through the same UDP port number and containing shared message structure), a router with BOOTP-relay agent capability typically relays DHCP packets and any BOOTP packets sent on the network.

If a router cannot function as a DHCP/BOOTP relay agent, each subnet must have either its own DHCP server or another computer that can function as a relay agent on that subnet. In cases where it is impractical or impossible to configure routers to support DHCP/BOOTP relay, you can configure a computer running Windows NT Server 4.0 or later to act as a relay agent by installing the DHCP Relay Agent service.

In most cases, routers support DHCP/BOOTP relay. If your routers do not, contact your router manufacturer or supplier to find out if a software or firmware upgrade is available to support this feature.

**How relay agents work**
A relay agent relays DHCP/BOOTP messages that are broadcast on one of its connected physical interfaces, such as a network adapter, to other remote subnets to which it is connected by other physical interfaces. The following illustration shows how client C on Subnet 2 obtains a DHCP address lease from DHCP server 1 on Subnet 1.




1. DHCP client C broadcasts a DHCP/BOOTP discover message (DHCPDISCOVER) on Subnet 2, as a User Datagram Protocol (UDP) datagram using the well-known UDP server port of 67 (the port number reserved and shared for BOOTP and DHCP server communication).

2. The relay agent, in this case a DHCP/BOOTP relay-enabled router, examines the gateway IP address field in the DHCP/BOOTP message header. If the field has an IP address of 0.0.0.0, the agent fills it with the relay agent or router's IP address and forwards the message to the remote Subnet 1 where the DHCP server is located.

3. When DHCP server 1 on remote Subnet 1 receives the message, it examines the gateway IP address field for a DHCP scope that can be used by the DHCP server to supply an IP address lease.

4. If DHCP server 1 has multiple DHCP scopes, the address in the gateway IP address field (GIADDR) identifies the DHCP scope from which to offer an IP address lease.

   For example, if the gateway IP address (GIADDR) field has an IP address of 10.0.0.2, the DHCP server checks its available set of address scopes for a scope range of addresses that matches the class A IP network that includes the gateway

address as a host. In this case, the DHCP server would make a check for a scope of addresses between 10.0.0.1 and 10.0.0.254. If a matching scope exists, the DHCP server selects an available address from the matched scope to use in an IP address lease offer response to the client.

5. When DHCP server 1 receives the DHCPDISCOVER message, it processes and sends an IP address lease offer (DHCPOFFER) directly to the relay agent identified in the gateway IP address (GIADDR) field.

6. The router then relays the address lease offer (DHCPOFFER) to the DHCP client.

   The client IP address is still unknown, so it has to be broadcast on the local subnet. Similarly, a DHCPREQUEST message is relayed from client to server, and a DHCPACK message is relayed from server to client, according to RFC 1542.

**Note**

- For more information, see Relay agent design issues.

# Relay agent configurations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Relay agent configurations**

There are several relay agent configuration options available if you plan to incorporate a relay agent into your DHCP/BOOTP-enabled network. These options include using non-Microsoft routers, Windows Server 2003 Routing and Remote Access service, and the DHCP Relay Agent component provided in Windows NT Server 4.0.

**Non-Microsoft router**

This example shows a non-Microsoft router configuration. The router is running a relay agent forwarding DHCP requests between Subnet A and Subnet B. Typically, the relay agent on the router must be configured with the IP address of the DHCP server.




**The Routing and Remote Access service**

This example shows the Windows Server 2003 Routing and Remote Access configuration. The DHCP Relay Agent in the Windows Server 2003 family must be configured with the IP address of the DHCP server in order to relay DHCP requests between Subnet A and Subnet B. For more information about setting up the DHCP Relay Agent, see Configure the DHCP Relay Agent.




**The Windows Server 2003 family**

This example shows how a standard IP router can be used with a Windows Server 2003 operating system to relay DHCP requests between Subnet A and Subnet B.



# Relay agent design issues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Relay agent design issues**

There are important design issues to consider before you set up relay agents on your network.

When you have multiple DHCP servers, Microsoft recommends that you place your DHCP servers on different subnets, to achieve a degree of fault tolerance, rather than having all of them on one subnet. The servers should not have common IP addresses in their scopes (each server should have a unique pool of addresses).

If the DHCP server on the local subnet shuts down, requests are relayed to a remote subnet. The DHCP server at that location can respond to DHCP requests if it maintains a scope of IP addresses for the requesting subnet. If the remote server has no scope defined for the requesting subnet, it cannot provide IP addresses even if it has available addresses for other scopes. If each DHCP server has a pool of addresses for each subnet, then it can provide IP addresses for remote clients whose own DHCP server is offline.

**Recommended**

The following relay agent implementation is recommended for best routed DHCP network performance. It uses two DHCP servers that are attached to two different subnets. The relay agents operating on each of the routers used to connect the subnets have varied delay intervals (one is set to 4 seconds, the other uses no delay). This eliminates the risk of undesirable floods of DHCP packets through randomly selected network paths.




**Not recommended**

The following relay agent implementation is not recommended. It has the same number of DHCP servers and subnets as in the previous example, but less caution has been used. There is a higher risk of DHCP discovery messages (DHCPDISCOVERs), which are broadcast by the DHCP clients at startup, being randomly forwarded or duplicated when destined for the other subnets. Moreover, the relay agents do not use any delay interval. This increases the likelihood that both relay agents (routers) duplicate any client broadcasted messages, which could potentially flood the remote subnets.


Ethernet LAN
10.10.10.0
network
IP router
with relay agent
Set agent to use:
10.20.20.1, 10.20.20.2
10.30.30.1, 10.30.30.2
[Set delay to 0]
Ethernet LAN
10.20.20.0
network
IP router
with relay agent
Set agent to use:
10.20.20.1, 10.20.20.2
10.30.30.1, 10.30.30.2
[Set delay to 0]
10.20.20.1
10.20.20.2
DHCP server
DHCP server
10.30.30.1
10.30.30.2
Ethernet LAN
10.30.30.0
network

# Administering DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering DHCP**

This section provides information about administering DHCP:

- Managing DHCP from the command line
- DHCP performance monitoring reference
- Managing DHCP Servers
- Setting Up Scopes
- Setting Up Options
- DHCP Interoperability Issues

# Managing DHCP from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing DHCP from the command line**

The Netsh commands for DHCP offer a command-line tool that helps with the administration of DHCP servers and provides an equivalent alternative to console-based management. This can be useful in the following situations:

- When managing DHCP servers in wide area networks (WANs), commands can be used in interactive mode at the **netsh>** command prompt to better manage across slow-speed network links.
- When managing a large number of DHCP servers, commands can be used in batch mode to help script and automate recurring administrative tasks that need to be performed for all DHCP servers.

For a comprehensive **netsh dhcp** command reference (including syntax, parameters, and examples), see Netsh commands for DHCP.

**Notes**

- For information about how to use the Netsh commands provided for DHCP, see Use DHCP Command-line Tools.
- For an example of how to use Netsh commands for DHCP in batch files and scripts, see Netsh DHCP example.
- For more information about Netsh, see The Netsh Command-Line Utility.

# DHCP performance monitoring reference

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP performance monitoring reference**

The Windows Server 2003 DHCP service includes a set of performance counters that can be used to monitor various types of server activity. To access these counters, you must use System Monitor. For more information, see System Monitor overview.

By default, performance monitoring of DHCP servers is available for use after the DHCP service component is installed. DHCP performance is measured using the following metrics and counters:

| Name | Description |
|---|---|
| Packets Received/sec | The number of message packets received per second by the DHCP server. A large number indicates heavy DHCP-related message traffic to the server. |
| Duplicates Dropped/sec | The number of duplicated packets per second dropped by the DHCP server. This number can be affected by multiple DHCP relay agents or network interfaces forwarding the same packet to the server. A large number here indicates that either clients are probably timing out too fast or the server is not responding fast enough. |
| Packets Expired/sec | The number of packets per second that expire and are dropped by the DHCP server. When a DHCP-related message packet is internally queued for 30 seconds or more, it is determined to be stale and expired by the server. A large number here indicates that the server is either taking too long to process some packets while other packets are queued and becoming stale, or traffic on the network is too high for the server to manage. |
| Milliseconds per packet (Avg.) | The average time in milliseconds used by the DHCP server to process each packet it receives. This number can vary depending on the server hardware and its I/O subsystem. A sudden or unusual increase might indicate a problem, either with the I/O subsystem becoming slower or because of an intrinsic processing overhead on the server computer. |
| Active Queue Length | The current length of the internal message queue of the DHCP server. This number equals the number of unprocessed messages received by the server. A large number might indicate heavy server traffic. |
| Conflict Check Queue Length | The current length of the conflict check queue for the DHCP server. This queue holds messages without responses while the DHCP server performs address conflict detection. A large value here might indicate that **Conflict Detection Attempts** has been set too high or that there is unusually heavy lease traffic at the server. |
| Discovers/sec | The number of DHCP discover messages (DHCPDISCOVERs) received per second by the server. These messages are sent by clients when they start on the network and obtain a new address lease. A sudden or unusual increase indicates a large number of clients are attempting to initialize and obtain an IP address lease from the server, such as when a number of client computers are started at any given time. |
| Offers/sec | The number of DHCP offer messages (DHCPOFFERs) sent per second by the DHCP server to clients. A sudden or unusual increase in this number indicates heavy traffic on the server. |
| Requests/sec | The number of DHCP request messages (DHCPREQUESTs) received per second by the DHCP server from clients. A sudden or unusual increase in this number indicates a large number of clients trying to renew their leases with the DHCP server. This might indicate that scope lease durations are too short. |
| Informs/sec | The number of DHCP information messages (DHCPINFORMs) received per second by the DHCP server. DHCP inform messages are used when the DHCP server queries for the directory service for the enterprise root and when dynamic updates are being done on behalf of clients by the server. |
| Acks/sec | The number of DHCP acknowledgment messages (DHCPACKs) sent per second by the DHCP server to clients. A sudden or unusual increase in this number indicates that a large number of clients are being renewed by the DHCP server. This might indicate that scope lease durations are too short. |

| | |
|---|---|
| Nacks/sec | The number of DHCP negative acknowledgment messages (DHCPNAKs) sent per second by the DHCP server to clients. A very high value might indicate potential network trouble in the form of misconfiguration of either the server or clients. When servers are misconfigured, one possible cause is a deactivated scope. For clients, a very high value could be caused by computers moving between subnets, such as laptop portables or other mobile devices. |
| Declines/sec | The number of DHCP decline messages (DHCPDECLINEs) received per second by the DHCP server from clients. A high value indicates that several clients have found their address to be in conflict, possibly indicating network trouble. In this situation, it can help to enable conflict detection on the DHCP server. This should only be used temporarily. Once the situation returns to normal, it should be turned off. |
| Releases/sec | The number of DHCP release messages (DHCPRELEASEs) received per second by the DHCP server from clients. This number only exists if a DHCP client sends a release to the server. This can occur manually, such as when the **ipconfig** command is used at the client computer. Release messages can also be sent by the client if it is configured with the "Release DHCP lease on shutdown" option. Because clients rarely release their address, the value of this counter remains low for many DHCP network configurations. For more information about DHCP options, see "DHCP Options" at the Microsoft Windows Resource Kits Web site. |

**Note**

- Values for the following registry key can be used by a member of the local Administrators group:

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\DHCPServer\Performance**

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

  By changing values stored at this registry path location, you can deny or modify permissions that determine who can view System Monitor counters used for the DHCP server on the server computer. Values under this key are used by the System Monitor DLL (Perfctrs.dll) to identify each of the counters described. For more information, see System Monitor overview.

# Managing DHCP Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing DHCP servers**

This section covers:

- Authorizing DHCP servers
- Monitoring DHCP server performance
- Restoring server data
- Backing up the DHCP database
- Analyzing server log files
- Cluster support for DHCP servers

# Authorizing DHCP servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authorizing DHCP servers**

The Windows Server 2003 family provides integrated security support for networks that use Active Directory. This support adds and uses a class of objects that is part of the base directory schema, providing the following enhancements:

- A list of IP addresses available for the computers that you authorize to operate as DHCP servers on your network.
- Detection of unauthorized DHCP servers and prevention of their starting or running on your network.

The following sections discuss:

- Background information about the detection of unauthorized DHCP servers.
- How computers are authorized in Active Directory to provide DHCP service.
- How an unauthorized server is detected and prevented from providing DHCP service.
- Notes and limitations for implementing DHCP service, depending on whether the Active Directory directory service is available.

**Background on unauthorized DHCP servers**

When configured correctly and authorized for use on a network, DHCP servers provide a useful and intended administrative service. However, when a misconfigured or unauthorized DHCP server is introduced into a network, it can cause problems. For example, if an unauthorized DHCP server starts, it might begin either leasing incorrect IP addresses to clients or negatively acknowledging DHCP clients attempting to renew current address leases.

Either of these configurations can produce further problems for DHCP-enabled clients. For example, clients that obtain a configuration lease from the unauthorized server can fail to locate valid domain controllers, preventing clients from successfully logging on to the network.

To resolve these issues, DHCP servers running Windows Server 2003 are verified as authorized in Active Directory before they can service clients. This avoids most of the accidental damage caused by running DHCP servers with incorrect configurations or correct configurations on the wrong network.

**How DHCP servers are authorized**

The authorization process for DHCP server computers depends on the installed role of the server on your network. In the Windows Server 2003 family there are three roles or server types for which each server computer can be installed:

- **Domain controller.** The computer keeps and maintains a copy of the Active Directory database and provides secure account management for domain member users and computers.
- **Member server.** The computer is not operating as a domain controller but has joined a domain in which it has a membership account in the Active Directory database.
- **Stand-alone server.** The computer is not operating as a domain controller or a member server in a domain. Instead, the server computer is made known to the network through a specified workgroup name, which can be shared by other computers, but is used only for browsing purposes and not to provide secured logon access to shared domain resources.

If you deploy Active Directory, all computers operating as DHCP servers must be either domain controllers or domain member servers before they can be authorized and provide DHCP service to clients.

Although it is not recommended, you can use a stand-alone server as a DHCP server as long as it is not on a subnet with any authorized DHCP servers. When a stand-alone DHCP server detects an authorized server on the same subnet, it automatically stops leasing IP addresses to DHCP clients.

For more information about DHCP server installation, see Install a DHCP server and Open the DHCP console.

For more information about authorizing a DHCP server in Active Directory, see Authorize a DHCP server in Active Directory.

For more information about delegating administrative credentials, see Delegate ability to authorize DHCP servers to a non-enterprise administrator.

**How unauthorized servers are detected**
DHCP servers running Windows Server 2003 provide detection of both authorized and unauthorized servers using the following specific enhancements for the DHCP standard:

- The use of information messaging between DHCP servers using the DHCP information message (DHCPINFORM).
- The addition of several new vendor-specific option types, for communicating information about the root domain.

A DHCP server running Windows Server 2003 uses the following process to determine whether Active Directory is available. If found, the server ensures that it is authorized by adhering to the following procedure, depending on whether it is a member server or a stand-alone server:

- For **member servers** (a server joined to a domain that is part of the enterprise), the DHCP server queries Active Directory for the list of authorized DHCP server IP addresses.

  If the server finds its IP address in the authorized list, it initializes and starts providing DHCP service to clients. If it does not find itself in the authorized list, it does not initialize and stops providing DHCP services.

  When installed in a multiple forest environment, DHCP servers seek authorization from within their forest only. Once authorized, DHCP servers in a multiple forest environment lease IP addresses to all reachable clients. Therefore, if clients from another forest are reached using routers with DHCP/BOOTP forwarding enabled, the DHCP server leases IP addresses to them.

  If Active Directory is not available, the DHCP server continues to operate in its last known state.

- For **stand-alone servers** (a server not joined to any domain or part of an existing enterprise). When the DHCP service starts, it sends a DHCP information message (DHCPINFORM) request to the reachable network, using the local limited broadcast address (255.255.255.255) to locate the root domain on which other DHCP servers are installed and configured.

  This message includes several vendor-specific option types that are known and supported by other DHCP servers running Windows Server 2003 . When received by other DHCP servers, these option types enable the query and retrieval of information about the root domain. When queried, the other DHCP servers reply with DHCP acknowledgement messages (DHCPACK) to both acknowledge and answer with Active Directory root domain information.

  If the stand-alone server receives no reply, it initializes and starts providing DHCP services to clients. If the stand-alone server receives a reply from a DHCP server that is authorized in Active Directory, the stand-alone server does not initialize and does not provide DHCP services to clients.

Authorized servers repeat the detection process at a default interval of 60 minutes. Unauthorized servers repeat the detection process at a default interval of 10 minutes.

Efforts to detect unauthorized servers are noted as "Restarting rogue detection" entries in the audit log.

**Notes**

- The process of authorizing DHCP servers is useful only for DHCP servers running Windows 2000 or Windows Server 2003 .
- When a DHCP server running Windows Server 2003 is installed in a Windows NT 4.0 domain, the server initializes and begins serving DHCP clients in the absence of directory services. However, if there is a Windows Server 2003 domain on the same subnet or on a connected network with routers configured for DHCP or BOOTP forwarding, the DHCP server in the Windows NT 4.0 domain detects its own unauthorized status and ceases to provide IP address leases to clients. If you

authorize the DHCP server in Active Directory, it can provide DHCP service to the clients in the Windows NT 4.0 domain.

- For the directory authorization process to work properly, it is necessary that the first DHCP server introduced on your network participate in Active Directory. This requires that the server be installed as either a domain controller or a member server. When planning or deploying Active Directory with Windows Server 2003 DHCP, it is important that you do not install your first DHCP server as a stand-alone server.

- Most commonly, there is one enterprise root and, therefore, only a single point for directory authorization of the DHCP servers. However, there is no restriction on authorizing DHCP servers for more than one enterprise root.

- The fully qualified domain name (FQDN) of the DHCP server cannot exceed 64 characters. If the FQDN of the DHCP server exceeds 64 characters, the attempt to authorize the server fails with the error message, "A constraint violation has occurred." If your DHCP server FQDN exceeds 64 characters, authorize the server using the server's IP address instead of its FQDN.

- DHCP servers running Windows 2000 cannot be authorized using the Windows Server 2003 administrative tools (the DHCP console and the Netsh commands for DHCP) unless Windows 2000 Server Service Pack 2 is installed.

# Monitoring DHCP server performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring DHCP server performance**
Because DHCP servers are of critical importance in most environments, monitoring their performance can help in troubleshooting cases where server performance has been degraded.

The Windows Server 2003 family provides a set of DHCP server performance counters that can be used to measure and monitor various aspects of server activity, such as the following:

- All types of DHCP messages sent and received by the DHCP service.
- The average amount of processing time spent by the DHCP server per message packet sent and received.
- The number of message packets dropped because of internal delays at a DHCP server computer.

For more information, see DHCP performance monitoring reference.

# Restoring server data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restoring server data**

Restoring the DHCP server database is useful in situations where the database has either become corrupted or lost. In order to successfully restore the database, you need to back up the database on a regular basis. By default, DHCP performs synchronous backups that occur automatically at a 60 minute interval to the folder location *systemroot*\System32\Dhcp\Backup, but you can also perform manual backups, or use backup software to copy the database to other locations. For more information, see Backing up the DHCP database and Backup.

You can restore the DHCP database from a backup copy of the database file, Dhcp.mdb, using the **Restore** command on the **Action** menu in the DHCP console. When you restore the DHCP database from a backup copy of Dhcp.mdb, the following information is configured on the server:

- All scopes, including superscopes and multicast scopes
- Reservations
- Leases
- All options, including server options, scope options, reservation options, and class options
- All registry keys and other configuration settings (for example, audit log settings and folder location settings) set in DHCP server properties.

**Note**

- The DNS dynamic update credentials (user name, domain, and password) that the DHCP server uses when registering DHCP client computers in DNS are not backed up with any backup method, therefore they are not restored during the restore procedure. After restoring your DHCP database, configure credentials for your server. For more information, see Configure DNS dynamic update credentials.

For more information, see Restore the DHCP database and Move a DHCP database to another server.

# Backing up the DHCP database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backing up the DHCP database**

Maintaining a backup of the DHCP database protects you from data loss if the DHCP database is lost (for example, due to hard disk failure) or becomes corrupted. There are three backup methods supported by the DHCP Server service:

- Synchronous backups that occur automatically. The default backup interval is 60 minutes.
- Asynchronous (manual) backups, performed by using the **Backup** command on the DHCP console. For more information about asynchronous backups, see Back up the DHCP database.
- Backups using Windows Backup (ntbackup.exe) or non-Microsoft backup software. For more information about Windows Backup, see Backup.

When a synchronous or asynchronous backup occurs, the entire DHCP database is saved, including the following:

- All scopes, including superscopes and multicast scopes
- Reservations
- Leases
- All options, including server options, scope options, reservation options, and class options
- All registry keys and other configuration settings (for example, audit log settings and folder location settings) set in DHCP server properties. These settings are stored in the following registry subkey:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DHCPServer\Parameters**

  To back up this subkey, open Registry Editor, and save the specified key to a text file. For more information, see Open Registry Editor and Import or Export Registry Keys.

### Note

- The DNS dynamic update credentials (user name, domain, and password) that the DHCP server uses when registering DHCP client computers in DNS are not backed up with any backup method.

**Backup locations**

The default DHCP database backup path is *systemroot*\System32\Dhcp\Backup. You can change the database backup folder by selecting a different local folder during a manual backup or by changing the backup folder location in DHCP server properties. For more information, see Set the DHCP database backup path.

When planning your backup strategy, consider the following:

- Synchronous backups are performed while the DHCP service is running. Similarly, when you perform asynchronous (manual) backups, you do not need to stop the DHCP service unless you are moving your database to a new server. For more information, see Move a DHCP database to another server.
- The backup folder location must be a local path.

### Notes

- This topic assumes that the DHCP database and its backup directories are located on the default drive and path locations.

- The default automatic database backup interval is 60 minutes. You can change the automatic backup interval through the following registry key:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DHCPServer\Parameters\BackupInterval**

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

- Do not store DHCP server backups on the same hard drive on which the DHCP service is installed, and make sure that the access control list (ACL) for the backup folder contains only the Administrators group and the DHCP Administrator group as members. In addition to performing manual backups, back up to other locations, such as a tape drive, and make sure that unauthorized persons do not have access to your backup copies. You can use Windows Backup for this purpose. For more information, see Best practices for Backup.

# Analyzing server log files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Analyzing server log files**

In Windows Server 2003 family, DHCP server log files are designed to use Audit logging to permit log files to remain enabled for use without requiring added monitoring or administering to manage log file growth or conserve disk resources. DHCP audit logs are located by default at %windir%\System32\Dhcp.

The following section outlines the format of these log files and how they can be used to gather more information about DHCP Server service operations on the network.

**DHCP server log file format**

DHCP server logs are comma-delimited text files with each log entry representing a single line of text. Following are the fields (and the order in which they appear) in a log file entry:

*ID, Date, Time, Description, IP Address, Host Name, MAC Address*

Each of these fields is described in detail in the following table:

| Field | Description |
|---|---|
| ID | A DHCP server event ID code. |
| Date | The date on which this entry was logged on the DHCP server. |
| Time | The time at which this entry was logged on the DHCP server. |
| Description | A description of this DHCP server event. |
| IP Address | The IP address of the DHCP client. |
| Host Name | The host name of the DHCP client. |
| MAC Address | The media access control address used by the network adapter hardware of the client. |

**DHCP server log: Common event codes**

DHCP server audit log files use reserved event ID codes to provide information about the type of server event or activity logged. The following table describes these event ID codes in more detail.

| Event ID | Description |
|---|---|
| 00 | The log was started. |
| 01 | The log was stopped. |
| 02 | The log was temporarily paused due to low disk space. |
| 10 | A new IP address was leased to a client. |
| 11 | A lease was renewed by a client. |
| 12 | A lease was released by a client. |
| 13 | An IP address was found in use on the network. |
| 14 | A lease request could not be satisfied because the address pool of the scope was exhausted. |
| 15 | A lease was denied. |
| 20 | A BOOTP address was leased to a client. |

**DNS dynamic update events**

When the DHCP server is configured to perform DNS dynamic updates on behalf of DHCP clients, you can use the DHCP audit logs to monitor update requests by the DHCP server to the DNS server, DNS record update successes, and DNS record update failures. The following event ID's are used for DNS dynamic update events:

| ID number | DHCP Event |
|---|---|
| 30 | DNS dynamic update request |
| 31 | DNS dynamic update failed |
| 32 | DNS dynamic update successful |

The IP address of the DHCP client computer is included in the DHCP audit log, providing the ability to track the source in the event of a denial of service attack.

**DHCP server logs: Server authorization events**

The following are additional server log event ID codes and descriptions. These events can appear in logs made by DHCP servers running Windows Server 2003 . They pertain to the applicable DHCP server and its authorization status when deployed in Active Directory environments.

| Event ID | Description |
|---|---|
| 50 | **Unreachable domain**<br><br>The DHCP server could not locate the applicable domain for its configured Active Directory installation. |
| 51 | **Authorization succeeded**<br><br>The DHCP server was authorized to start on the network. |
| 52 | **Upgraded to a** Windows Server 2003 **operating system**<br><br>The DHCP server was recently upgraded to a Windows Server 2003 operating system, and, therefore, the unauthorized DHCP server detection feature (used to determine whether the server has been authorized in Active Directory) was disabled. |
| 53 | **Cached Authorization**<br><br>The DHCP server was authorized to start using previously cached information. Active Directory was not currently visible at the time the server was started on the network. |
| 54 | **Authorization failed**<br><br>The DHCP server was not authorized to start on the network. When this event occurs, it is likely followed by the server being stopped. |
| 55 | **Authorization (servicing)**<br><br>The DHCP server was successfully authorized to start on the network. |
| 56 | **Authorization failure, stopped servicing**<br><br>The DHCP server was not authorized to start on the network and was shut down by the operating system. You must first authorize the server in the directory before starting it again. For more information, see Authorize a DHCP server in Active Directory. |
| 57 | **Server found in domain**<br><br>Another DHCP server exists and is authorized for service in the same domain. |
| 58 | **Server could not find domain**<br><br>The DHCP server could not locate the specified domain. |

| | |
|---|---|
| 59 | **Network failure**<br><br>A network-related failure prevented the server from determining if it is authorized. |
| 60 | **No DC is DS Enabled**<br><br>No Windows Server 2003 domain controller (DC) was located. For detecting whether the server is authorized, a DC that is enabled for Active Directory is needed. |
| 61 | **Server found that belongs to DS domain**<br><br>Another DHCP server was found on the network that belongs to the Active Directory domain. |
| 62 | **Another server found**<br><br>Another DHCP server was found on the network. |
| 63 | **Restarting rogue detection**<br><br>The DHCP server is trying once more to determine whether it is authorized to start and provide service on the network. |
| 64 | **No DHCP enabled interfaces**<br><br>The DHCP server has its service bindings or network connections configured so that it is not enabled to provide service. This usually means one of the following:<br><br>1. The network connections of the server are either not installed or not actively connected to a network.<br><br>2. The server has not been configured with at least one static IP address for one of its installed and active network connections. For more information, see Configure TCP/IP for static addressing.<br><br>3. All of the statically configured network connections for the server are disabled. For more information, see Multihomed DHCP servers. |

**Example: Excerpt from a sample DHCP server audit log**
The following is a brief excerpt of sample log activity from an audit log generated by the DHCP Server service:

```
ID Date,Time,Description,IP Address,Host Name,MAC Address
00,04/19/99,12:43:06,Started,,,
60,04/19/99,12:43:21,No DC is DS Enabled,,MYDOMAIN,
63,04/19/99,12:43:28,Restarting rogue detection,,,
01,04/19/99,13:11:13,Stopped,,,
00,04/19/99,12:43:06,Started,,,
55,04/19/99,12:43:54,Authorized(servicing),,MYDOMAIN,
```

In this sample, the DHCP server could not be authorized when initially started and is subsequently stopped. Once authorized, the server is then able to start and service clients.

### Note

- For more information, see Audit logging.

# Cluster support for DHCP servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster support for DHCP servers**
The Windows Server 2003 DHCP Server service is a cluster-aware application. You can implement additional DHCP (or MADCAP) server reliability by deploying a DHCP server cluster using the Cluster service provided with Windows Server 2003, Enterprise Edition.

By using clustering support for DHCP, you can implement a local method of DHCP server failover, achieving greater fault tolerance. You can also enhance fault tolerance by combining DHCP server clustering with a remote failover configuration, such as by using a split scope configuration.

**Requirements for setting up a DHCP clustered resource**
The DHCP Server service, when used in a server cluster, requires the following resource types: a disk resource, an IP address resource, and a name resource.

In this type of configuration, a virtual IP address is defined for the IP address resource using the Cluster Administrator tool. This address must be a static IP address, not one obtained from another DHCP server.

The DHCP Server service then binds to this virtual IP address. This IP address must then be used to authorize the DHCP Server service in Active Directory so it can function correctly and service clients on the network.

Once the virtual IP is set using cluster configuration, each network adapter installed on each node (that is, a server in the cluster) needs an IP address configuration. Because the DHCP Server service does not bind to adapter IP addresses, you can either provide them through DHCP or configure them statically. If you use static configuration, the IP addresses for each linked pair of network adapters (linked node-to-node) should be from the same subnet.

Addresses that are statically configured using TCP/IP properties from the Network Connections folder can also be viewed through server bindings properties in the DHCP console. However, for a DHCP server cluster, the bindings information that appears in these views does not correspond to the virtual IP address used for the server cluster. Therefore, these settings do not apply.

When a scope is created for a DHCP clustered server, the virtual IP addresses used in the cluster context must be excluded so that they are not distributed to clients. Additionally. you should configure the database path, audit log file path, and the database backup path using the Cluster Administrator tool on the shared disk.

**Other options for DHCP failover**
Another way to implement DHCP remote failover is to deploy two DHCP servers in the same network that share a split scope configuration based on the 80/20 rule. For more information, see the section " Using the 80/20 rule for scopes" in Configuring scopes.

For more information about installing a DHCP resource type on a Windows Server 2003, Enterprise Edition, cluster server, see Checklist: Installing a DHCP Service resource.

### Notes

- The Jet database upon which the DHCP Server service relies is not cluster aware. If the clustered DHCP database becomes inaccessible to the DHCP Server service, the service will report thousands of ESENT application log errors, and DHCP service by the cluster will be interrupted.

- If you need only to remotely manage DHCP server clusters that are deployed on remote computers running Windows Server 2003, Enterprise Edition, you can use Cluster Administrator to do this.

- For more information about configuring DHCP clustering, see "Dynamic Host Configuration Protocol" at the Microsoft Windows Resource Kits Web site.

# Setting Up Scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up scopes**

This section covers:

- Configuring scopes
- Using superscopes
- Using multicast scopes
- Removing scopes
- Managing leases
- Using client reservations
- Reconciling scopes

# Configuring scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring scopes**

A scope is an administrative grouping of IP addresses for computers on a subnet that use the DHCP service. The administrator first creates a scope for each physical subnet and then uses the scope to define the parameters used by clients. A scope has the following properties:

- A range of IP addresses from which to include or exclude addresses used for DHCP service lease offerings.
- A subnet mask, which determines the subnet for a given IP address.
- A scope name assigned when it is created.
- Lease duration values, which are assigned to DHCP clients that receive dynamically allocated IP addresses.
- Any DHCP scope options configured for assignment to DHCP clients, such as DNS server, router IP address, and WINS server address.
- Reservations, optionally used to ensure that a DHCP client always receives the same IP address.

**Before adding scopes**

A DHCP scope consists of a pool of IP addresses on a given subnet, such as 192.168.0.1 to 192.168.0.254, that the DHCP server can lease to clients.

Each subnet can have only a single DHCP scope with a single continuous range of IP addresses. To use several address ranges within a single scope or subnet for DHCP service, you must first define the scope and then set any needed exclusion ranges.

- **Defining the scope**

  Use the entire range of consecutive IP addresses that make up the local IP subnet for which you are enabling DHCP service.

- **Setting exclusion ranges**

  You should set exclusion ranges for any IP addresses within the scope that you do not want the DHCP server to offer or use for DHCP assignment. For example, you can exclude the first 10 addresses in the previous example scope by creating an exclusion for 192.168.0.1 to 192.168.0.10.

  By setting an exclusion for these addresses, you specify that DHCP clients are never offered these addresses when they request leased configuration from the server. Excluded IP addresses can be active on your network, but only by manually configuring these addresses at hosts that do not use DHCP to obtain an address.

**Creating scopes**

When you create a DHCP scope, you use the DHCP console to enter the following required information:

- A scope name, assigned by you or the administrator who created the scope.
- A subnet mask used to identify the subnet to which an IP address belongs.
- A range of IP addresses contained within the scope.
- A time interval (known as a *lease duration*) that specifies how long a DHCP client can use an assigned IP address before it must renew its configuration with the DHCP server.

**Using the 80/20 rule for scopes**

For balancing DHCP server usage, a good practice is to use the "80/20" rule to divide the scope addresses between the two DHCP servers. If Server 1 is configured to make available most (approximately 80%) of the addresses, then Server 2 can be configured to make the other addresses (approximately 20%) available to clients. The following illustration is an example of the 80/20 rule:




## Notes

- For detailed information on applying the 80/20 rule, see DHCP Best Practices.
- When creating a new scope, the IP address used to create it should not include addresses of existing statically configured computers, such as the DHCP server. Either these static addresses should be outside the range for the scope, or they should be excluded from the scope address pool.

**After scopes are added**

After you define a scope, you can additionally configure the scope by performing the following tasks:

- **Set additional exclusion ranges.**

  You can exclude any other IP addresses that must not be leased to DHCP clients. You should use exclusions for all devices that must be statically configured. The excluded ranges should include all IP addresses that you assigned manually to other DHCP servers, non-DHCP clients, diskless workstations, or Routing and Remote Access and PPP clients.

- **Create reservations.**

  You can choose to reserve some IP addresses for permanent lease assignment to specified computers or devices on your network. You should make reservations only for devices that are DHCP-enabled and that must be reserved for specific purposes on your network (such as print servers).

- **Adjust the length of lease durations.**

  You can modify the lease duration to be used for assigning IP address leases. The default lease duration is eight days.

  For most LANs, the default value is acceptable but can be further increased if computers seldom move or change locations. Also, infinite lease times can be set but should be used with caution. For information about circumstances under which modifying this setting is most useful, see Managing leases.

- **Configure options and classes to be used with the scope.**

  To provide full configuration for clients, DHCP options need to be configured and enabled for the scope. For more information, see Setting Up Options.

  For more advanced discrete management of scope clients, you can add or enable user or vendor-defined option classes. For more information, see Using option classes.

**Notes**

- After you define and configure a scope, the scope must be *activated* before the DHCP server begins providing service to clients. However, you should not activate a new scope until you have specified the DHCP options for it. For information on how to activate a scope, see Activate a scope.

- After you activate a scope, you should not change the range of scope addresses.

# Using superscopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using superscopes**

A superscope is an administrative feature of DHCP servers running Windows Server 2003 that you can create and manage through the DHCP console. Using a superscope, you can group multiple scopes as a single administrative entity. With this feature, a DHCP server can:

- Support DHCP clients on a single physical network segment (such as a single Ethernet LAN segment) where multiple logical IP networks are used. When more than one logical IP network is used on each physical subnet or network, such configurations are often called *multinets*.

- Support remote DHCP clients located on the far side of DHCP and BOOTP relay agents (where the network on the far side of the relay agent uses multinets).

In multinet configurations, you can use DHCP superscopes to group and activate individual scope ranges of IP addresses used on your network. In this way, the DHCP server computer can activate and provide leases from more than one scope to clients on a single physical network.

Superscopes can resolve certain types of DHCP deployment issues for multinets, including situations in which:

- The available address pool for a currently active scope is nearly depleted, and more computers need to be added to the network. The original scope includes the full addressable range for a single IP network of a specified address class. You need to use another IP network range of addresses to extend the address space for the same physical network segment.

- Clients must be migrated over time to a new scope (such as to renumber the current IP network from an address range used in an existing active scope to a new scope that contains another IP network range of addresses).

- You want to use two DHCP servers on the same physical network segment to manage separate logical IP networks.

**Superscope configurations for multinets**

The following section shows how a simple DHCP network consisting originally of one physical network segment and one DHCP server, can be extended to use superscopes for support of multinet configurations.

**Example 1: Non-routed DHCP server (before superscope)**

In this preliminary instance, a small local area network (LAN) with one DHCP server supports a single physical subnet, Subnet A. The DHCP server in this configuration is limited to leasing addresses to clients on this same physical subnet.

The following illustration shows this example network in its original state. At this point, no superscopes have been added and a single scope, Scope 1, is used to service all DHCP clients on Subnet A.




Single scope here configured as follows:
**Scope1:** 192.168.1.1 - 192.168.1.254
**Subnet mask:** 255.255.255.0
**Excluded addresses:** 192.168.1.1 - 192.168.1.10

**Example 2: Superscope for non-routed DHCP server supporting local multinets**

To include multinets implemented for client computers on Subnet A, the same network segment where the DHCP server is

located, you can configure a superscope which includes as members: the original scope (Scope 1) and additional scopes for the logical multinets for which you need to add support (Scope 2, Scope 3).

This illustration shows the scope and superscope configuration to support the multinets on the same physical network (Subnet A) as the DHCP server.




Superscope here with member scopes:
**Scope 1:** 192.168.1.1 - 192.168.1.254
**Scope 2:** 192.168.2.1 - 192.168.2.254
**Scope 3:** 192.168.3.1 - 192.168.3.254

**Subnet mask for all scopes:** 255.255.255.0
**Excluded addresses for member scopes:**
**Scope 1:** 192.168.1.1 - 192.168.1.10
**Scope 2:** 192.168.2.1 - 192.168.2.10
**Scope 3:** 192.168.3.1 - 192.168.3.10

**Example 3: Superscope for routed DHCP server with relay agent supporting remote multinets**
To include multinets implemented for client computers on Subnet B, the remote network segment located across a router from the DHCP server on Subnet A, you can configure a superscope which includes as members: the additional scopes for the logical multinets for which you need to add remote support (Scope 2, Scope 3).

Note that since the multinets are for the remote network (Subnet B), the original scope (Scope 1) does not need to be part of the added superscope.

This illustration shows the scope and superscope configuration to support the multinets on the remote physical network (Subnet B) away from the DHCP server.




Scope for local Subnet A:
**Scope 1:** 192.168.1.1 - 192.168.1.254
**Subnet mask:** 255.255.255.0
**Excluded addresses:** 192.168.1.1 - 192.168.1.10

Superscope added here with
member scopes for Subnet B:
**Scope 2:** 192.168.2.1 - 192.168.2.254
**Scope 3:** 192.168.3.1 - 192.168.3.254
**Subnet mask:** 255.255.255.0
**Excluded addresses:**
**(Scope 2)** 192.168.2.1 - 192.168.2.10
**(Scope 3)** 192.168.3.1 - 192.168.3.10

**Note**

- A DHCP relay agent is used for DHCP servers to support clients on remote subnets, as shown in Example 3. For more information, see Understanding relay agents.

# Using multicast scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using multicast scopes**

Multicast scopes are supported through the use of Multicast Address Dynamic Client Allocation Protocol (MADCAP), a proposed standard protocol for performing multicast address allocation. The MADCAP protocol describes how multicast address allocation (or MADCAP) servers can dynamically provide IP addresses to other computers (MADCAP clients) on your network.

Typically, a MADCAP client might also be a multicast server (MCS) used to support IP multicasting. An MCS manages the shared or group use of the allocated multicast IP address and streams data traffic to members that share the use of the specified group address.

Once an MCS is configured and allocated a group IP address to use, any multicast clients that have registered their membership with the MCS can receive streams sent to this address. By registering with the MCS, clients can participate efficiently in the stream process, such as for real-time video or audio network transmissions. The MCS also manages the multicast group list, updating its membership and status so that multicast traffic is received by all current members.

**Background on multicast addresses**

Ordinarily, you use DHCP scopes to provide client configurations by allocating ranges of IP addresses from the standard address classes, Class A, B, or C. By using DHCP scopes you can assign IP addresses from the ranges provided by these classes for your DHCP clients to be configured to use unicast (or point-to-point) directed communication between other TCP/IP networked computers.

The multicast address range uses an additional address class, Class D, that includes IP addresses that range from 224.0.0.0 to 239.255.255.255 for use in IP multicasting. Addresses in this class are used for multicasting only and not for regular DHCP scopes.

In all TCP/IP networks, each host is required to first be configured with its own IP address, taken from one of the standard address classes. You must assign this required unicast IP address before you can configure a host to support and use secondary IP addresses, such as a multicast IP address.

Another difference between unicast and multicast addresses, is that a group of TCP/IP host computers are intended to be able to share the use of a multicast IP address. This is not normally the case with unicast IP addresses, which are intended to be assigned individually to only one configured host and not shared with other hosts.

When the destination address for an IP datagram is an IP multicast address, the datagram is forwarded to all members of a *multicast group*, which is a set of zero or more hosts identified by the address. The membership of a multicast group is dynamic in that individual hosts can join or leave the group at any time.

Membership and use of multicast groups is unrestricted and can be compared to membership and use of a group e-mail address: Group membership can be any size, and hosts can be members of many multicast groups.

You can permanently reserve multicast group addresses or temporarily assign and use them as needed on your network. For a permanent group IP address to be reserved for Internet use, you must register it with the Internet Assigned Numbers Authority (IANA).

For multicast IP addresses not permanently reserved with the IANA, all Class D addresses that remain unreserved can then be used dynamically to assign and form temporary multicast groups. These temporary groups can exist as long as one or more hosts on the network are configured with the group address and actively share its use.

For more information, see Understanding Multicasting or Internet Group Management Protocol (IGMP).

**Determining the ranges to use for multicast scopes**

When deciding the IP address ranges to use for multicast scopes on your MADCAP server, there are two overall best practices recommended by the Multicast Allocation (MALLOC) working group, an IETF team of industry volunteers who help establish multicast address allocation standards. These practices include:

- **Administrative scoping** This is most useful when you are using multicast IP addresses privately in your own network. It is similar to private IP addressing, as used currently in the unicast IP address spaces (such as the use of the 10.0.0.0 network address space). It is fully discussed in RFC 2365, "Administratively Scoped IP Multicast."

With administratively scoped multicast IP, the range most recommended that you begin with is the 239.192.0.0 range. This range is known as the IPv4 Organization Local Scope and has a subnet mask of 255.252.0.0 (14-bits in length). It is intended for use by an entire organization setting multicast scopes privately for its own internal or organizational use. Starting with this address, you can create a considerable number of addresses--up to $2^{18}$, or 262,144 group addresses--for use in all subnets within your organization's network.

For more information, refer directly to the RFC, which can be obtained at the Request for Comments Web site.

- **Global scoping** This practice is most useful when you are using multicast group IP addresses in a public network address space, particularly the Internet.

  Because most organizations requiring one or more public addresses likely have already been assigned some public unicast IP addresses, a subnet allocation scheme has been proposed. For MADCAP, the 233.0.0.0 range of the Class D address space is recommended for use as a global scope range.

  When the 233.0.0.0 range is used, the allocation of global addresses works in the following way:

  1. A network registry, such as the Internet Assigned Numbers Authority (IANA), allocates and reserves the first 8 bits of the range (for example, the "233" portion of this range).

  2. The next 16 bits (the two middle octet numbers of addresses in this range) are based on a previously assigned Autonomous System (AS) number. This number is recorded with the applicable Internet Assigned Numbers Authority (IANA) registry for your region.

     If you already register your IP addresses with a regional network registry, you might know your AS number. You can also look it up easily using the Whois database system on the Internet. If you are obtaining IP addresses through an Internet service provider (ISP), you might contact them to find out what your AS number is.

     AS numbers are allocated to the regional registries by the IANA. If you or your ISP are located in the United States, you can apply directly to the American Registry for Internet Numbers (ARIN) to obtain an AS number. For more information, see the ARIN Web site.

     For other regions, AS numbers can also be obtained from appropriate regional registries. Other regional registries include the following:

ARIN or <hostmaster@arin.net> for the Americas, Caribbean, and Africa.

RIPE-NCC <ncc@ripe.net> for Europe.

AP-NIC <admin@apnic.net> for the Asia-Pacific region.

1. The last 8 bits in the address are local use bits.

   These provide the IP address range from which to configure any multicast scopes for group addresses you want to use publicly on the Internet. With this global scoping system in use, a subnet mask of 255.255.255.0 should be applied. This provides each organization with an assigned AS number of up to 255 multicast group addresses for use on the Internet.

   For more information on either the practice of global scoping for multicast address allocation or the use and allocation of AS numbers, use the IETF Web site. For global scoping, refer directly to the current available draft entitled "Static Allocation in 233/8," and for more information on AS numbering, refer to RFC 1930, "Guidelines for creation, selection, and registration of an Autonomous System (AS)."

**MADCAP and DHCP**

The Windows Server 2003 DHCP service supports both the DHCP and MADCAP protocols. These protocols function separately and are not dependent on each other. For example, a DHCP client might or might not be a MADCAP client and a MADCAP client might or might not be a DHCP client.

It is also worth mentioning that the DHCP Server service can be used to deploy MADCAP servers independent of how DHCP servers are used on your network. For example, to install Windows Server 2003 DHCP for MADCAP service only:

1. Create multicast scopes.

2. Do not create other scopes or superscopes.

Only where other scopes or superscopes are configured does the server computer also function as a DHCP server.

For more information about MADCAP, see Multicast address allocation.

**Important**

- Multicast scopes and MADCAP only provide a mechanism for dynamically allocating IP address configuration for multicast-ranged IP addresses. Other network configuration details are normally required to enable multicasting for your deployment needs. For more information, see Checklist: Installing IP multicast video conferencing.

**Notes**

- Multicast scopes do not require or support the use of DHCP options, but can be configured with a finite lifetime, enabling the multicast scope to expire and be removed by the server.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

- For more information, see Manage Multicast Scopes.

# Removing scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Removing scopes**

When a subnet is no longer in use, or any other time you want to remove an existing scope, you can remove it at the DHCP console. In most cases, scopes are removed only when you need to renumber your network to use a different IP address range.

To retire a current scope and renumber your network to use a different scope, be sure to do the following first:

1. Create the new scope using a different range of addresses.
2. Configure options for the new scope.
3. Activate the new scope and deactivate the old scope.
4. When all clients are no longer using the old scope, remove the old scope.

**Scope migration strategies**

To remove old scopes without disrupting DHCP service to clients on the subnet, use the following suggestions to design your scope migration strategy:

- Activate a new replacement scope for the subnet before the old scope is deactivated. You can create a replacement scope on a different DHCP server or on the same server.
- Deactivate a scope before removing it. After a scope is deactivated, it does not acknowledge lease or renewal requests. DHCP clients will be unable to renew their leases from the deactivated scope and the original server, so the clients will seek new leases from an available DHCP server.
- Superscopes, reservations, and exclusion ranges can be used individually or together to provide a smooth migration of clients from one scope to another. For example, you can create a superscope that contains both a new scope and the original scope that has been targeted for deactivation. If you create an exclusion range on the original scope, clients with leases that fall within the exclusion range are unable to renew their leases and are assigned leases from the new scope instead. Or, if you create reservations for clients in the new scope, when a client attempts to renew its lease from the original scope, the client is instead assigned the address for which a reservation has been configured.

**Note**

- When manually reconfiguring Windows XP-based client computers, you can use the commands **ipconfig /release** and **ipconfig /renew** at the command prompt to release an old IP address and acquire a new lease from a replacement scope.

# Managing leases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing leases**

IP addresses are leased by the DHCP server to its clients. Each lease has an expiration date, which the client must renew if it is going to continue to use that address.

Leases are retained in the DHCP server database four hours after expiration. This grace period protects a client lease in case the client and server are in different time zones, the individual computer clocks are not synchronized, or the client computer is off the network when the lease expires. Expired leases are included in the list of active leases and are distinguished by a different icon.

You can delete the lease of any DHCP client in the scope. The main reason for doing so is to remove a lease that conflicts with an IP address exclusion or a client reservation that you want to add. Deleting a lease has the same effect as if the client lease expired: the next time that client computer starts, it must enter the initialization state and obtain new TCP/IP configuration information from a DHCP server. There is nothing, however, to prevent the client from obtaining a new lease for the same IP address. You must make the address unavailable before the client requests another lease.

Delete only entries for clients that are no longer using the assigned DHCP lease or that are to be moved immediately to a new address. Deleting an active client could result in duplicate IP addresses on the network, because deleted addresses will be assigned to new active clients. After you delete the lease of the client and set a reservation or exclusion, you should always use **ipconfig /release** on the client computer to force the client to free its IP address.

**Note**

- You can view lease information for individual clients by first selecting the Active Leases folder for a scope in the DHCP console. When Active Leases for a scope is selected as the console view, the list of currently active client leases for the scope appears in the details pane.

**Modifying default lease durations for a scope**

When a scope is created, the default lease duration is set to eight days. In most cases, this value is sufficient. However, because lease renewal is an ongoing process that can affect the performance of DHCP clients and your network, it is sometimes useful to change the lease duration. Use the following guidelines to decide how to best modify lease duration settings for improving DHCP performance on your network:

- If you have a large number of IP addresses available and configurations that rarely change on your network, increase the lease duration to reduce the frequency of lease renewal queries between clients and the DHCP server. This reduces some network traffic caused by clients renewing their leases.

- If you have a limited number of IP addresses available and if client configurations change or clients move often on your network, reduce the lease duration to promote scavenging of old IP addresses by the DHCP server. This increases the rate at which addresses are returned to the available address pool for reassignment to new clients.

**Note**

- You can also set a separate lease time, different from the default scope lease time used by your DHCP clients, such as alternate lease times for supporting other types of clients such as BOOTP or Routing and Remote Access clients. For more information, see Manage BOOTP and Remote Access Clients.

# Using client reservations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using client reservations**

With client reservations, you can reserve a specific IP address for permanent use by a DHCP client. Typically, you will need to do this if the client uses an IP address that was assigned using another method for TCP/IP configuration.

If multiple DHCP servers are configured with a scope that covers the range of the reserved IP address, the client reservation must be made and duplicated at each of these DHCP servers. Otherwise, the reserved client computer can receive a different IP address, depending on the responding DHCP server.

If you want to change a reserved IP address for a current client, you have to remove the existing address reservation of the client, and then add a new reservation. You can change any other information about a reserved client while keeping the reserved IP address.

If you are reserving an IP address for a new client, or an address that is different from its current one, you should verify that the address has not already been leased by the DHCP server. Reserving an IP address in a scope does not automatically force a client currently using that address to stop using it.

If the address is already in use, the client using the address must first release is by issuing a DHCP release message (DHCPRELEASE). To make this happen, at the command prompt of a client computer running Windows 2000 or Windows XP, type **ipconfig /release**.

Reserving an IP address at the DHCP server also does not force the new client for which the reservation is made to immediately move to that address. In this case, too, the client must first issue a DHCP request message (DHCPREQUEST). To make this happen, at the command prompt of a client computer running Windows 2000 or Windows XP, type **ipconfig /renew**.

For clients using Windows 95 or Windows 98, you can use the Winipcfg.exe program to cause release or renewal of the reserved IP address in DHCP. For clients using MS-DOS, and some clients using other operating systems, a computer must be restarted for the change to take effect.

Once these changes are made, the reserved client is leased the IP address now reserved for its permanent use each time it renews its lease with the DHCP server.

**Notes**

- Reserved clients can have DHCP options configured specifically for their use. When options are configured for a reserved client, these values override any option type parameters distributed via server-based, scope-based, or class-based options assignment.

- You can also use the **ipconfig** command on computers running earlier versions of Windows NT to release and renew DHCP leases. However, not all other **ipconfig** command options are available when executed at clients running under these earlier versions.

- You can create reservations using any IP address in the scope's address range, even if the IP address is also within an exclusion range. Because of this design, when the 80/20 rule is implemented and all addresses in the scope are excluded (80% at one server, 20% at the other), reservations still function properly.

# Reconciling scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Reconciling scopes**

Reconciling scopes can fix inconsistencies, such as incorrect or missing information for client IP addresses that are stored in scope lease information. Scope IP address lease information is stored in two forms by the DHCP Server service:

- Detailed IP address lease information, stored in the DHCP database.
- Summary IP address lease information, stored in the DHCP database.

When reconciling scopes, the detail and summary entries are compared to find inconsistencies.

To correct and repair these inconsistencies, you need to reconcile any scope inconsistencies found when performing this operation. Once you select and reconcile scope inconsistencies, the DHCP service either restores those IP addresses to the original owner or creates a temporary reservation for those addresses. These reservations are valid for the lease time assigned to the scope. When the lease time expires, the addresses are recovered for future use.

For more information, see Reconcile the DHCP database or Reconcile a scope.

# Setting Up Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up options**

This section covers:

- Assigning options
- Using option classes

# Assigning options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

### Assigning options

You can manage DHCP options at several distinct levels:

- **Predefined options.** At this level, you control which types of options are predefined for the DHCP server to expose as available options for assigning from any of the option configuration dialog boxes (such as **Server Options**, **Scope Options**, or **Reservation Options**) available through the DHCP console. You can add or remove options to and from the predefined list of standard options as needed. Although options are made available in this way, they are not assigned values until administratively configured at either the server, scope, or reservation.

- **Server options.** Assign values here (using the **General** tab) for options that should apply to or be inherited by all scopes and clients of the DHCP server as defaults. Options configured here can have their values overridden by different values if those values are set at either a scope, options class, or reserved client level.

- **Scope options.** Assign values here (using the **General** tab) for options that should apply only to clients of an applicable scope selected in the DHCP console tree. Options configured here can have their values overridden by different values if those values are set at either an options class or reserved client level.

- **Reservation options.** Assign values for options that should apply only to a specific reserved DHCP client. To use this level of assignment, you must first add a reservation for the applicable client to the applicable DHCP server and scope where the client is to obtain its IP address. These options are set for an individual DHCP client configured with an address reservation in a scope. Only properties manually configured at the client computer can override options assigned at this level.

- **Class options.** When using any of the option configuration dialog boxes (**Server Options**, **Scope Options** or **Reservation Options**), you can click the **Advanced** tab to configure and enable options for assignment to identifying member clients of a specified user or vendor class.

  Depending on the context, only those DHCP clients that identify themselves according to the selected class are distributed options data you have configured specifically for that class. For example, if a class-assigned option is set at a scope, only clients of that scope that indicate class membership during leasing activity are configured with class-assigned option values. Other non-member clients are configured using scope option values set from the **General** tab.

  Options configured here can override values assigned and set at the same context (either server, scope, or reservation options) or values inherited from options configured at a higher context. However, the ability of the client to indicate membership in a specific options class is typically the decisive criteria for using this level of options assignment. For more information, see Using option classes.

### Guidelines for assigning options

The following guidelines can help you determine what level to assign the options you use for clients on your network.

- Add or define new custom option types only if you have new software or applications that require a nonstandard DHCP option.

- If your DHCP server manages many scopes for a large network, be selective when assigning **Server Options**. These options apply by default to *all* clients of a DHCP server computer, unless otherwise overridden.

- Use **Scope Options** for assigning most options that clients use. In most networks, this level is typically preferred for assigning and enabling the use of DHCP options.

- Use **Class Options** if you have a mixture of DHCP clients with diverse needs that are able to identify a specific class on the DHCP server when obtaining a lease. For example, if you have a limited number of DHCP client computers running

Windows 2000, these clients can be configured to receive vendor-specific options that other clients do not use.

- Use **Reservation Options** for individual DHCP clients in your network that have special configuration requirements.

  For any hosts (that is, computers or other networked devices) that do not support DHCP or are not recommended to use it, you can also consider excluding IP addresses for those computers and devices and manually setting the IP address configuration directly at the applicable host. For example, you often need to statically configure the IP address for routers.

**Commonly used options**

After you set basic TCP/IP configuration settings (such as IP address, subnet mask, and default gateway) for clients, most clients also need the DHCP server to provide other information through DHCP options. The most common of these include the following:

- **Routers.** A preferred list of IP addresses for routers on the same subnet as DHCP clients. The client can then contact these routers as needed to forward IP packets destined for remote hosts.

- **DNS servers.** IP addresses for DNS name servers that DHCP clients can contact and use to resolve a domain host name query.

- **DNS domain.** Specifies the domain name that DHCP clients should use when resolving unqualified names during DNS domain name resolution.

- **WINS node type.** A preferred NetBIOS name resolution method for the DHCP client to use (such as b-node for broadcast only or h-node for a hybrid of point-to-point and broadcast methods).

- **WINS server.** IP addresses of primary and secondary WINS servers for the DHCP client to use.

**Note**

- When the DHCP service is installed, no DHCP option definitions are created. Option definitions are created only when the DHCP console is opened for the first time. If you want to configure the value of an option on a DHCP server on which the console has never been opened, you must first use the Netsh DHCP context command **server add optiondef** (at a **netsh dhcp>** command prompt or in a batch file or script) to create the option definition. To use Netsh to assign a value to the option definition you have created, use the Netsh DHCP context command **server set optionvalue**. For more information, see Netsh commands for DHCP.

# Using option classes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using option classes**

Option classes offer an additional method for grouping DHCP-provided configuration details for clients within a scope. For computers running Windows Server 2003 operating systems, there are two types of option classes that can be used for submanaging options:

- User classes, for assigning options to clients identified as sharing a common need for similar DHCP options configuration.
- Vendor classes, for assigning vendor-specific options to clients identified as sharing a commonly defined vendor type.

For more information, see Understanding user and vendor classes.

**How user classes work**

User classes allow DHCP clients to differentiate themselves by specifying a User Class option. When available for client use, this option includes a user-determined class ID that can help to group clients of similar configuration needs within a scope. For example, you might support users and computers with mobile computing needs by configuring a user class at the DHCP server and setting the related class ID at the client computers.

A user class is useful when you need to keep separate options that cover the special needs of identifying client computers, such as providing a shorter lease time for portable computers that move frequently or use remote access often. In this example, you could configure the DHCP server to distribute different options that are specific to the needs of clients.

For instance, shorter leases could be assigned to mobile clients. For more information, see Change the lease time for remote access clients.

The user class feature gives you greater flexibility in configuring DHCP clients on your network, but is not required for standard DHCP use. If user-defined option classes are not configured, options are provided through the applicable server, scope, or client option settings instead. For more information, see Assigning options.

DHCP servers running Windows Server 2003 , which support recognition of user class IDs, perform the following added steps to lease clients identifying themselves according to this process:

1. The server determines whether the user class identified by the client in its lease request is a recognized class, previously defined on the server.

   If a predefined user classes exists at the server and is configured, class-based assignment is enabled. For other user classes, you must first add and configure them at the server before they are available for use.

2. If the user class is recognized, the server determines whether any additional DHCP options are configured for this class in the active lease context (either the scope or a client reservation) where the server is leasing the client.

3. If the scope or reservation options are configured to provide options for the user-defined class of the client, the server returns those options to the client as part of its DHCP acknowledgment message (DHCPACK), which is sent to confirm the lease.

**Additional user class examples**

You can define specific user class identifiers to convey information about client software configuration, its physical location in a building, or about its user preferences. For example, an identifier can specify that DHCP clients are members of a user-defined class called "2nd floor, West," which has need for a special set of router, DNS, and WINS server settings.

You can also use the Microsoft predefined user classes for isolating configuration details specific for clients with special needs, such as BOOTP or Routing and Remote Access service. For more information, see Assign an advanced (class-based) option.

## Notes

- In addition to configuring or adding user classes at the DHCP server, at computers running Windows 2000, or Windows XP, you can set a DHCP class identifier with the **ipconfig /setclassid** command. For more information, see Set DHCP class ID information at a client computer.

- When you configure user classes at the DHCP server, be certain that the class identifying data you set at both the DHCP server and member client computers is identically matching binary or ASCII data. For more information, see Create a new user or vendor class.

- The user class option only permits one class ID to be used for identifying clients. Each client computer can only be identified as a member of a single user class at the DHCP server. You can use additional user classes and make new hybrids from your other user classes, as needed.

**How vendor classes work**

Vendor-defined option classes can be used by DHCP clients that are configured to optionally identify themselves by their vendor type to the DHCP server when obtaining a lease. For a client to identify its membership in a vendor class, the client provides a value in the **Vendor class identifier** option when it requests or selects a lease from the server.

The vendor class identifier information is a string of character data interpreted by a DHCP server. Microsoft supports vendor class identification for its DHCP clients running Windows 98 or Windows XP.

Most vendor types are derived from standard reserved hardware and operating system type abbreviation codes listed in Request for Comments (RFC) 1700. For more information, see TCP/IP RFCs.

Other vendors might choose to define their own specific vendor class IDs for either conveying vendor-specific information or handling special needs for their DHCP clients.

DHCP servers running Windows Server 2003 , which support recognition of vendor class IDs, perform the following additional steps to lease the clients that identify themselves according to this process:

1. The server determines whether the vendor class identified by the client in its lease request is a recognized class, previously defined on the server.

   Only Microsoft vendor classes are predefined at the server. For other vendor classes, you must manually add and configure these options at your DHCP servers running Windows Server 2003 before they are available for use.

2. If the vendor class is a recognized one, the server determines whether any additional DHCP options are configured for this class in the active lease context (either the scope or a client reservation) where the server is leasing the client.

3. If the scope or reservation options are configured to provide options for the vendor-defined class of the client, the server returns them using the **Vendor specific information** option as part of its DHCP acknowledgment message (DHCPACK), which is sent to confirm the lease.

Vendor classes permit other system vendors to support custom applications for DHCP in mixed vendor environments. Vendor-specific options, when provided, are used in addition to any of the standard DHCP options assigned or required for DHCP.

For more information about DHCP, see "Dynamic Host Configuration Protocol" at the Microsoft Windows Resource Kits Web site or refer directly to RFC 2132.

**Notes**

- For more information about DHCP options, see "DHCP Options" at the Microsoft Windows Resource Kits Web site.

- In most cases, the default (or unspecified) vendor class, DHCP standard options, can provide a useful default for grouping any DHCP clients that do not specify or recognize vendor class IDs.

- When you add new vendor classes at the DHCP server, be sure that the vendor class ID data you set at the server matches the actual vendor class ID used by clients for your vendor. For more information, contact the appropriate vendor.

# DHCP Interoperability Issues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interoperability issues**

The following section covers issues that affect the use of the DHCP Server service with other services or network configurations.

- Using DNS servers with DHCP
- Using Routing and Remote Access servers with DHCP
- Multihomed DHCP servers

# Using DNS servers with DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using DNS servers with DHCP**
When installing the Windows Server 2003 DHCP service, you can configure the server to perform updates on behalf of its DHCP clients to any Domain Name System (DNS) servers that support dynamic updates.

**How DHCP/DNS update interaction works**
The DHCP server can be used to register and update the pointer (PTR) and host (A) resource records on behalf of its DHCP-enabled clients.

This process requires the use of an additional DHCP option, the Client FQDN option (option 81). This option permits the client to provide its fully qualified domain name (FQDN) as well as instructions to the DHCP server on how it would like the server to process DNS dynamic updates (if any) on its behalf.

When this option is issued by a qualified DHCP client, such as a DHCP-enabled computer running Windows 2000, Windows XP, or a Windows Server 2003 operating system, option 81 is processed and interpreted by a DHCP server running Windows Server 2003 to determine how the server initiates updates on behalf of the client. If the server is configured to perform DNS dynamic updates, it takes one of the following actions:

- The DHCP server updates both DNS A and PTR records if requested by clients using option 81.
- The DHCP server updates DNS A and PTR records regardless of whether the client requests this action or not.

In addition, the DHCP server can dynamically update DNS A and PTR records on behalf of legacy clients that are not capable of sending option 81 to the server. You can also configure the DHCP server to discard client A and PTR records when the client lease is deleted.

The DHCP server might be configured in one of the following ways:

- The DHCP server registers and updates client information with the authoritative DNS server of the zone in which the DHCP server is located *according to the DHCP client request*.

  This is the default configuration for DHCP servers running Windows Server 2003 and DHCP clients running Windows 2000, Windows XP, or a Windows Server 2003 operating system. In this mode, the DHCP client can request the way in which the DHCP server performs updates of its host (A) and pointer (PTR) resource records. If possible, the DHCP server accommodates the client request for handling updates to its name and IP address information in DNS.

  To modify this setting, select the **Dynamically update DNS A and PTR records only if requested by the DHCP clients** check box, which is located in **Properties** on the **DNS** tab on the applicable DHCP server or on one of its scopes.

- The DHCP server *always* registers and updates client information in DNS.

  This is a modified configuration supported for DHCP servers running Windows Server 2003 and DHCP clients running Windows 2000, Windows XP, or a Windows Server 2003 operating system. In this mode, the DHCP server always performs updates of the client's FQDN, leased IP address information, and both its host (A) and pointer (PTR) resource records, regardless of whether the client has requested to perform its own updates.

  To modify this setting, select the **Enable DNS dynamic updates according to the settings below** check box and click **Always dynamically update DNS A and PTR records**, which is located in **Properties** on the **DNS** tab on the applicable DHCP server or on one of its scopes.

- The DHCP server *never* registers and updates client information in DNS.

  To set this behavior, the DHCP server must be configured to disable performance of DHCP/DNS proxied updates. By disabling this feature, no client host (A) or pointer (PTR) resource records are updated in DNS for DHCP clients.

If necessary, this change in setting can be made at DHCP servers running Windows Server 2003 by clearing the **Enable DNS dynamic updates according to the settings below** check box, which is located in **Properties** on the **DNS** tab on the applicable DHCP server or one of its scopes. By default, updates are always performed for newly installed DHCP servers running Windows Server 2003 and any new scopes created for them.

**Advanced DHCP/DNS server configuration options**

In addition to these standard DHCP/DNS interactions, the DHCP server can be configured to perform these optional update tasks as follows:

1. The server can selectively be configured to *not* send updates for discarding a client host (A) resource record when the client lease expires.

   When the DHCP server is enabled to perform DNS updates, it always sends updates to discard the client pointer (PTR) resource records when the lease expires. Whether the server also does this with client host (A) resource records when the lease of a client expires (by default, the server discards these) is a configurable option.

   To modify this at the applicable DHCP server, clear the **Discard forward (name-to-address) lookups when leases expires** check box in **Properties** on the **DNS** tab.

2. The server can be selectively configured to *not* send updates for clients unable to use the Client FQDN option (option 81), to request the way that updates are handled.

   By default, the DHCP server does not send updates for clients that do not support option 81.

   To modify this setting, select the **Dynamically update DNS A and PTR records for DHCP clients that do not request updates (for example, clients running Windows NT 4.0)** check box, which is located in **Properties** on the **DNS** tab on the applicable DHCP server or one of its scopes.

**Windows DHCP clients and DNS dynamic update protocol**

DHCP clients running Windows 2000, Windows XP, or a Windows Server 2003 operating system interact differently than earlier versions of Windows when performing the DHCP/DNS interactions previously described. The following examples and graphics show how this process varies in different cases.

**Example 1: DHCP/DNS update interaction for DHCP clients running Windows 2000, Windows XP , or a Windows Server 2003 operating system**

DHCP clients running Windows 2000, Windows XP, or a Windows Server 2003 operating system interact with DNS dynamic update protocol as follows:

1. The client initiates a DHCP request message (DHCPREQUEST) to the server and includes DHCP option 81. By default, the client requests that the DHCP server register the DNS PTR record, while the client registers its own DNS A record.

2. The server returns a DHCP acknowledgment message (DHCPACK) to the client, granting an IP address lease and including DHCP option 81. If the DHCP server is configured with the default settings (dynamically update DNS A and PTR records only if requested by the DHCP clients), then option 81 instructs the client that the DHCP server will register the DNS PTR record and the client will register the DNS A record.

3. Asynchronously, the client registers its DNS A record, and the DHCP server registers the DNS PTR record of the client.




**Example 2: DHCP/DNS update interaction for earlier Windows DHCP clients (prior to Windows 2000)**

Earlier versions of Windows DHCP clients do not support the DNS dynamic update process directly, and therefore, cannot directly interact with the DNS server. For these DHCP clients, updates are typically handled as follows:

1. The client initiates a DHCP request message (DHCPREQUEST) to the server. The request does not include DHCP option 81.

2. The server returns a DHCP acknowledgment message (DHCPACK) to the client, granting an IP address lease, without DHCP option 81.

3. The server then sends updates to the DNS server for the forward lookup record of the client, which is a host (A) resource record. The server also sends updates for the reverse lookup record of the client, which is a pointer (PTR) resource record.




**DNS record ownership and the DnsUpdateProxy group**
As previously described, you can configure a DHCP server so that it dynamically registers host (A) and pointer (PTR) resource records on behalf of DHCP clients. In this configuration, the use of secure dynamic update with DNS servers might cause stale resource records.

For example, suppose the following sequence of events occurs:

1. A DHCP server running Windows Server 2003 (DHCP1) performs a secure dynamic update on behalf of one of its clients for a specific DNS domain name.

2. Because the DHCP server successfully created the name, it becomes the owner of the name.

3. Once the DHCP server becomes the owner of the name, only that DHCP server can update the DNS records for that name.

In some circumstances, this can cause problems. For example, if DHCP1 fails and a second backup DHCP server comes online, the second server cannot update the client name because it is not the owner of the name.

In another example, if the DHCP server performs DNS dynamic updates for legacy DHCP clients (clients running a version of Windows earlier than Windows 2000), and those clients are later upgraded to Windows 2000, Windows XP, or a Windows Server 2003 operating system, the upgraded client cannot take ownership of or update its own DNS records.

To solve this problem, the built-in security group called DnsUpdateProxy is provided. If all DHCP servers are added as members of the DnsUpdateProxy group, then the records of one server can be updated by another server if the first server fails. Also, because all of the objects that are created by the members of the DnsUpdateProxy group are not secured, the first user (that is not a member of the DnsUpdateProxy group) to modify the set of records that is associated with a DNS name becomes its owner. When legacy clients are upgraded, they can therefore take ownership of their name records at the DNS server. If every DHCP server registering resource records for legacy clients is a member of the DnsUpdateProxy group, the problems discussed earlier are eliminated.

You can configure the DnsUpdateProxy security group through Active Directory Users and Computers. For more information, see Add a member to a group.

**Securing records when using the DnsUpdateProxy group**
DNS domain names that are registered by the DHCP server are not secure when the DHCP server is a member of the DnsUpdateProxy group. The host (A) resource record for the DHCP server itself is an example of such a record. Also, because objects created by the members of the DnsUpdateProxy group are not secured, it is impossible to use this group effectively in an Active Directory integrated zone that allows only secure dynamic updates unless you take additional steps to allow records created by members of the group to be secured.

To protect against unsecured records, or to allow members of the DnsUpdateProxy group to register records in zones that allow only secured dynamic updates, you can create a dedicated user account and configure DHCP servers to perform DNS dynamic updates with the user account credentials (user name, password, and domain). The credentials of one dedicated user

account can be used by multiple DHCP servers.

A dedicated user account is a user account whose sole purpose is supplying DHCP servers with credentials for DNS dynamic update registrations. When you create a dedicated user account and configure DHCP servers with the account credentials, each DHCP server supplies these credentials when registering names on behalf of DHCP clients using DNS dynamic update. The dedicated user account should be created in the forest where the primary DNS server for the zone to be updated resides. The dedicated user account can also be located in another forest as long as the forest it resides in has a forest trust established with the forest containing the primary DNS server for the zone to be updated. For more information about establishing forest trusts, see Forest trusts.

When the DHCP Server service is installed on a domain controller, configuring the DHCP server with the credentials of the dedicated user account will prevent the server from inheriting, and possibly misusing, the power of the domain controller. When installed on a domain controller, the DHCP Server service inherits the security permissions of the domain controller and has the authority to update or delete any DNS record that is registered in a secure Active Directory-integrated zone (this includes records that were securely registered by other computers running Windows 2000 or a Windows Server 2003 operating system, including domain controllers).

It is necessary to configure a dedicated user account and configure the DHCP server with the account credentials under the following circumstances:

- A domain controller is configured to function as a DHCP server.
- The DHCP server is configured to perform DNS dynamic updates on behalf of DHCP clients.
- The DNS zones to be updated by the DHCP server are configured to allow only secure dynamic updates.

Once you have created a dedicated user account, you can configure DHCP servers with the user account credentials by using the DHCP console or by using the Netsh DHCP context command **server set dnscredentials**. For more information on configuring credentials using the DHCP console, see Configure DNS dynamic update credentials. For a comprehensive reference about Netsh commands for DHCP (including syntax, parameters, and examples for the **set dnscredentials** command at the **server** context of **netsh dhcp**), see Netsh commands for DHCP.

For more information on groups and accounts, see Create a new user account and Default groups.

**Additional resources on DHCP and DNS interaction**
Because of the close integration of DHCP and DNS described in this section, you might want to investigate or review these additional topics as resources for your deployment issues related to this feature.

For more information:

- On the DNS dynamic update protocol and considerations for securing dynamic updates when using Active Directory, see Dynamic update.
- On how to configure DNS servers to enable some or all updates for specific zones, see Allow dynamic updates and Allow only secure dynamic updates.
- On how to enable DHCP servers to initiate updates on behalf of its clients for specific scopes (or all scopes), see Enable DNS dynamic updates for clients.
- On how you can modify or set computers to initiate dynamic updates on their own behalf, see Configure TCP/IP to use DNS.
- On DNS resource records, see DNS or Resource records reference.

### Notes

- Microsoft supports the DHCP/DNS update interaction currently under final review as a proposed Internet standard Request for Comment (RFC) document by the Internet Engineering Task Force (IETF).

  This interactive process is fully described in the applicable draft, "Interaction between DHCP and DNS." To obtain the current or final version of this draft, see the Request for Comments Web site.

- If you are using legacy DNS servers that do not support dynamic updates, you can implement the following changes for dynamically updating and registering names and addresses for your DHCP clients:

    1. If you are using WINS and legacy Windows DNS servers (that is, DNS servers running Windows NT Server 4.0), you might be able to use WINS lookup integration to support resolution of client names and addresses on your network that are not updated dynamically in DNS.

    2. Upgrade or replace older DNS servers with DNS servers running Windows Server 2003 .

- For more information about DHCP, see "Dynamic Host Configuration Protocol" at the Microsoft Windows Resource Kits Web site.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

For more information about the DNS registration changes related to DHCP Servers running Windows Server 2003, see DNS registration changes for Windows Server 2003 based DHCP Servers

# DNS registration changes for Windows Server 2003 based DHCP Servers

Updated: August 19, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This article describes the DNS registration changes for DHCP Servers that are running Windows Server® 2003.

**Overview**

When the DHCP Server role is installed on a domain controller, the DHCP Server inherits the security permissions of the domain controller. To prevent possible misuse of the domain controller's elevated permissions, DHCP Servers that are installed on Windows Server 2003 do not register DNS records on behalf of the clients that are associated with the DHCP Server unless the DHCP Server is explicitly configured with DNS credentials. This behavior prevents a potential escalation of privilege to clients of the DHCP Server.

**Configuring DNS credentials on a DHCP Server that is running Windows Server 2003**

Follow these steps to configure the DNS credentials of a DHCP Server that is installed on a domain controller running the Windows Server 2003 operating system:

1. Start the DHCP management console. Click **Start**, then click **Run** and type **dhcpmgmt.msc**.

2. In the DHCP navigation pane, expand the server node and right-click **IPv4**, and then click **Properties**.

3. In the resulting dialog box, click the **Advanced** tab, and then click **Credentials**.

4. In the **DNS dynamic update credentials** dialog box, enter the DNS credentials (User name, Domain, and Password) for the account that is used by the DHCP Server for DNS registrations.

| **Note** |
| --- |
| These credentials should belong to a member of the DnsAdmins or the DnsUpdateProxy group. |

**Related information**

For more information about creating and adding a user account as a member of a group by using the Active Directory® Users and Groups console, see Changing group memberships.

# Using Routing and Remote Access servers with DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Routing and Remote Access servers with DHCP**

The DHCP Server service can be deployed along with the Routing and Remote Access service to provide remote access clients with a dynamically assigned IP address during connection. When these services are used together on the same server computer, the information provided during dynamic configuration is provided differently than in the case of typical DHCP configuration for LAN-based clients.

In LAN environments, DHCP clients negotiate and receive the following configuration information, based entirely on settings configured in the DHCP console for the DHCP server:

- A leased IP address provided from the available address pool of an active scope on the DHCP server. The DHCP server directly manages and distributes the address to the LAN-based DHCP client.
- Additional parameters and other configuration information provided through assigned DHCP options in the address lease. The values and list of options used correspond to option types configured and assigned on the DHCP server.

When a Routing and Remote Access server provides dynamic configuration for dial-up clients, it first performs the following steps:

- When the Routing and Remote Access server starts with the **Use DHCP to assign remote TCP/IP addresses** option, it instructs the DHCP client to obtain 10 IP addresses from a DHCP server.
- The remote access server utilizes the first of these 10 IP addresses obtained from the DHCP server for the remote access server interface.
- The remaining nine addresses are allocated to TCP/IP-based clients as they dial in to establish a session with the remote access server.

IP addresses that are freed when remote access clients disconnect are reused. When all 10 IP addresses are used, the remote access server obtains 10 more from a DHCP server. When the Routing and Remote Access service is stopped, all IP addresses obtained through DHCP are released.

When the Routing and Remote Access server uses this type of proactive caching of DHCP address leases for dial-up clients, it records the following information for each lease response it obtains from the DHCP server:

- The IP address of the DHCP server
- The client leased IP address (for later distribution to the Routing and Remote Access client)
- The time at which the lease was obtained
- The time at which the lease expires
- The duration of the lease

All other DHCP option information returned by the DHCP server (such as server, scope, or reservation options) is discarded. When the client dials in to the server and requests an IP address (that is, when **Server Assigned IP Address** is selected), it uses a cached DHCP lease to provide the dial-up client with dynamic IP address configuration.

When the IP address is provided to the dial-up client, the client is unaware that the IP address has been obtained through this intermediate process between the DHCP server and the Routing and Remote Access server. The Routing and Remote Access server maintains the lease on behalf of the client. Therefore, the only information that the client receives from the DHCP server is the IP address lease.

In dial-up environments, DHCP clients negotiate and receive dynamic configuration using the following modified behavior:

- A leased IP address from the Routing and Remote Access server cache of DHCP scope addresses. The Routing and Remote Access server obtains and renews its cached address pool with the DHCP server.

- If additional parameters and other configuration information provided through assigned DHCP options in the address lease is normally provided by the DHCP server, this information is returned to the Routing and Remote Access client based on TCP/IP properties configured on the Routing and Remote Access server.

**Note**

- DHCP servers running Windows Server 2003 provide a predefined user class, the **Default Routing and Remote Access Class**, for assigning options that are specific and provided only to Routing and Remote Access clients.

# Multihomed DHCP servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multihomed DHCP servers**

A multihomed DHCP server is a computer running a Windows Server 2003 operating system that uses the DHCP service for more than a single network connection. For a server computer to be multihomed, each network connection must attach the computer to more than a single physical network. This requires that additional hardware (in the form of multiple installed network adapters) be used on the computer.

A computer running a Windows Server 2003 operating system can perform as a multihomed DHCP server. For multihomed servers, the DHCP service binds to the first IP address statically configured for each network connection in use.

By default, the service bindings depend on whether the first network connection is configured dynamically or statically for TCP/IP. Based on the method of configuration it uses, reflected by its current settings in **Internet Protocol (TCP/IP)** properties, the DHCP Server service performs default service bindings as follows:

- If the first network connection uses a manually specified IP address, the connection is enabled in server bindings. For this to occur, a value for **IP address** must be configured and the **Use the following IP address** option selected in Internet Protocol (TCP/IP) properties. In this mode, the DHCP server listens for and provides service to DHCP clients.

- If the first network connection uses an IP address configured dynamically, the connection is disabled in server bindings. This occurs when the **Obtain an IP address automatically** option is selected in Internet Protocol (TCP/IP) properties. For computers running Windows Server 2003 operating systems, this is the default setting. In this mode, the DHCP server does not listen for and provide service to DHCP clients until a static IP address is configured.

- The DHCP server will bind to the first static IP address configured on each adapter.

### Notes

- By design, DHCP server bindings are enabled and disabled on a per-connection, not per-address basis. All bindings are based on the first configured IP address for each connection appearing in the Network Connections folder. If additional static IP addresses (for example, as set in Advanced TCP/IP properties) are configured for the applicable connection, these addresses are never used by DHCP servers running Windows Server 2003 and are inconsequential for server bindings.

- For more information, see About Network Connections.

- DHCP servers running Windows Server 2003 never bind to any of the NDISWAN or DHCP-enabled interfaces used on the server. These interfaces are not displayed in the DHCP console under the current server bindings list because they are never used for DHCP service. Only additional network connections that have a primary static IP address configured can appear in the server bindings list (or be selectively enabled or disabled there).

- For more information, see Selectively set DHCP server bindings for network connections.

**Example: Configuring a multihomed DHCP server**

The following illustration is an example of a multihomed DHCP server with three network adapters installed. Each adapter is configured to lease addresses on separate physical subnets.




Note that in the graphic, each network adapter is configured statically with a single IP address. In this network example, three Class C ranges of IP addresses are used:

| Subnet | IP Address Range |
|---|---|
| Subnet A | 192.168.0.1 - 192.168.0.254 |
| Subnet B | 192.168.1.1 - 192.168.1.254 |
| Subnet C | 192.168.2.1 - 192.168.2.254 |

Each of the three adapters in this example connects the server to three different physical subnets (Subnets A, B, and C). To achieve the intended result of having DHCP provide leased configuration service to all clients in each of the respective subnets, two configuration details are essential and must be verified during deployment plans:

1. The server must use a statically configured IP address on the same subnet as the IP address range it services.

2. The server must have each of its valid subnet IP addresses excluded from the scope used to offer leases to clients.

The default subnet mask value (255.255.255.0) is used for all three subnets, providing 254 IP addresses available per subnet for lease to DHCP clients.

To provide a small pool of ten IP addresses available for statically configured devices such as the DHCP server, printers, and other devices, an exclusion range is created in the scope for each subnet:

| Subnet | Exclusion Range |
|---|---|
| Subnet A | 192.168.0.1 - 192.168.0.10 |
| Subnet B | 192.168.1.1 - 192.168.1.10 |
| Subnet C | 192.168.2.1 - 192.168.2.10 |

Exclusion ranges prevent the DHCP server from assigning IP addresses to DHCP clients, so all of the IP addresses contained in the exclusion ranges will remain available for static configurations even after the DHCP server is started and begins servicing clients.

Addresses that are not excluded remain available for assignment to DHCP clients. These ranges of IP addresses form each subnet's address pool:

| Subnet | Address Pool |
|---|---|
| Subnet A | 192.168.0.11 - 192.168.0.254 |
| Subnet B | 192.168.1.11 - 192.168.1.254 |

| | |
|---|---|
| Subnet C | 192.168.2.11 - 192.168.2.254 |

**Note**

- For more information on IP subnetting, see Subnet masks or the Microsoft Windows Server 2003 *Resource Kit*.

# DHCP Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This section covers:

- DHCP updated technical information
- DHCP console icons reference

# DHCP updated technical information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP updated technical information**

Updated technical information for DHCP can be obtained by searching technical articles available from the Personal Support Center at the Microsoft Web site. For more information, see Updated technical information.

You can also obtain additional in-depth information on DHCP in the Microsoft Windows Server 2003 *Resource Kit*.

For more information about DHCP, see "Dynamic Host Configuration Protocol" at the Microsoft Windows Resource Kits Web site.

For more information about DHCP options, see "DHCP Options" at the Microsoft Windows Resource Kits Web site.

# DHCP console icons reference

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP console icons reference**
DHCP provides context-specific icons to dynamically represent the state of console objects. Objects include the DHCP servers, clients, scopes, and option types that are used and managed for each server.

**General usage of console icons**
Some of the icons used in DHCP have specific meanings, but most are used to communicate general status information about the current state of the console. For example, when an icon is used with an hourglass superimposed, such as:

it means that the console is attempting to connect to the server.

Likewise, when an icon is used with a red circle and a superimposed 'X', such as:

it often means that DHCP has lost its connection to the server.

**Glossary of specific icons**
The following is a partial list of console icons used in DHCP.

- Server-related icons
- Scope-related icons
- Option-related icons
- Other console icons

**Server-related icons**

| Icon | Description |
|---|---|
| | DHCP server added to console. |
| | DHCP server connected and active in console. |
| | DHCP server connected but not authorized in Active Directory for use on your network. |
| | DHCP server connected but current user does not have the administrative credentials to manage the server. |
| | DHCP server warning. Available addresses for server scopes are 90 percent or more leased and in use. This means that the server is nearly depleted of available addresses to lease to clients. |
| | DHCP server alert. No addresses are available from server scopes because the maximum (100 percent) of the addresses allocated for use are currently leased. This represents a failure of the DHCP server on the network because it is not able to lease or service clients. |

**Scope-related icons**

| Icon | Description |
|---|---|
| | Scope or superscope is active. |
| | Scope or superscope is inactive. |

| | |
|---|---|
| | Scope or superscope warning. Scope warning: 90 percent or more of the scope's IP addresses are in use. Superscope warning: If any scope within the superscope has a warning, the superscope has a warning. |
| | Scope or superscope alert. Scope alert: All IP addresses have been allocated by the DHCP server and are in use. No more clients can obtain IP addresses from the DHCP server because it has no more IP addresses to allocate. Superscope alert: At least one scope contained in the superscope has all IP addresses allocated by the DHCP server. No clients can obtain an IP address from the range defined in the scope that is 100 percent allocated. If other scopes within the superscope contain available addresses, the DHCP server can allocate addresses from these scopes. |

**Option-related icons**

| Icon | Description |
|---|---|
| | Server options folder. |
| | Server option. |
| | Scope options folder. |
| | Scope option. |
| | Reservation option. |

**Other console icons**

| Icon | Description |
|---|---|
| | Root of the DHCP console. |
| | Address pool folder. |
| | Scope allocation range. Addresses in this range are allocated to the available address pool used to offer leases to clients. |
| | Scope exclusion range. Addresses in this range are excluded from the available address pool used to offer leases to clients. |
| | Active leases folder. |
| | Active lease: this address is not available for lease by the DHCP server. |
| | Expired lease: this address is available for lease by the DHCP server. |
| | Active lease, DNS dynamic update pending. This address is not available for lease by the DHCP server. |
| | Client is using dial-up network connection through a remote access server. |
| | An individual reservation and the reservations folder. |
| | Bootstrap Protocol (BOOTP) table. |
| | BOOTP entry in the table. |

# DHCP Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

- Troubleshooting DHCP clients
- Troubleshooting DHCP servers

# Troubleshooting DHCP clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting DHCP clients**
What problem are you having?

- The DHCP client does not have an IP address configured or indicates that its IP address is 0.0.0.0.
- The DHCP client appears to have automatically assigned itself an IP address that is incorrect for the current network.
- The DHCP client appears to be missing some network configuration details or is unable to perform related tasks, such as resolving names.
- The DHCP client appears to have incorrect or incomplete options, such as an incorrect or missing router (default gateway) configured for the subnet on which it is located.
- Many DHCP clients are unable to get IP addresses from the DHCP server.
- The DHCP client appears to be affected by another problem not described above.

**The DHCP client does not have an IP address configured or indicates that its IP address is 0.0.0.0.**

- **Cause:** The client was not able to contact a DHCP server and obtain an IP address lease, either because of a network hardware failure or because the DHCP server is unavailable.
- **Solution:** Verify that the client computer has a valid functioning network connection. First, check that related client hardware (cables and network adapters) are working properly at the client using basic network and hardware troubleshooting steps.

  If the client hardware appears to be prepared and functioning properly, check that the DHCP server is available on the network by pinging it from another computer on the same network as the affected DHCP client.

**The DHCP client appears to have automatically assigned itself an IP address that is incorrect for the current network.**

- **Cause:** The client running Windows 98, Windows Millennium Edition, or Windows XP could not find a DHCP server and has used IP autoconfiguration to configure its IP address.

  In some larger networks, disabling IP autoconfiguration might be desirable for network administration.
- **Solution:** First, use the ping command to test connectivity from the client to the server. Your next step would be to either verify or manually attempt to renew the client lease. Depending on your network requirements, it might be necessary to disable IP autoconfiguration at the client. You can learn more about IP autoconfiguration and how it works prior to making this decision.
- **See also:** Test a TCP/IP configuration by using the ping command; Verify, release, or renew a client address lease; Configure TCP/IP for automatic addressing; Disable automatic address configuration

**The DHCP client appears to be missing some network configuration details or is unable to perform related tasks, such as resolving names.**

- **Cause:** The client might be missing DHCP options in its leased configuration, either because the DHCP server is not configured to distribute them or the client does not support the options distributed by the server.
- **Solution:** For Microsoft DHCP clients, verify that the most commonly used and supported options have been configured

at the server, scope, client, or class level of options assignment.

- **See also:** Manage Options and Classes; Assigning options

**The DHCP client appears to have incorrect or incomplete options, such as an incorrect or missing router (default gateway) configured for the subnet on which it is located.**

- **Cause:** The client has the full and correct set of DHCP options assigned but its network configuration does not appear to be working correctly.

  If the DHCP server is configured with an incorrect DHCP router option (option code 3) for the default gateway address of the client, clients running Windows NT, Windows 2000, or Windows XP use the correct address. However, DHCP clients running Windows 95 use the incorrect address.

- **Solution:** Change the IP address list for the router (default gateway) option at the applicable DHCP scope and server. If you are configuring the router option as a **Server Option** at the affected DHCP server, remove it there and set the correct value in the **Scope Options** node for the applicable DHCP scope that services the client.

  In rare instances, you might have to configure the DHCP client to use a specialized list of routers different from other scope clients. In such cases, you can add a reservation and configure the router option list specifically for the reserved client.

- **See also:** Assign a server-based option; Assign a scope-based option; Assign an option to a reserved client; Add a client reservation; Assigning options

  For more information about DHCP options, see "DHCP Options" at the Microsoft Windows Resource Kits Web site.

**Many DHCP clients are unable to get IP addresses from the DHCP server.**

- **Cause:** The IP address of the DHCP server was changed and now DHCP clients cannot get IP addresses.

- **Solution:** A DHCP server can only service requests for a scope that has a network ID that is the same as the network ID of its IP address.

  Make sure that the DHCP server IP address falls in the same network range as the scope it is servicing. For example, a server with an IP address in the 192.168.0.0 network cannot assign addresses from scope 10.0.0.0 unless superscopes are used.

- **See also:** DHCP Best Practices; Using superscopes; Configuring scopes

- **Cause:** The DHCP clients are located across a router from the subnet where the DHCP server resides and are unable to receive an address from the server.

- **Solution:** A DHCP server can provide IP addresses to client computers on remote multiple subnets only if the router that separates them can act as a DHCP relay agent.

  Completing the following steps might correct this problem:

1. Configure a BOOTP/DHCP Relay Agent on the client subnet (that is, the same physical network segment). The relay agent can be located on the router itself, on a computer running Windows NT Server and the DHCP Relay Agent component, on a computer running Windows 2000 Server with the Routing and Remote Access service enabled and configured as a DHCP Relay Agent, or on a computer running a Windows Server 2003 operating system with the Routing and Remote Access service enabled and configured as a DHCP Relay Agent.

2. At the DHCP server, do the following:

   a. Configure a scope to match the network address on the other side of the router where the affected clients are

located.

b. In the scope, make sure that the subnet mask is correct for the remote subnet.

c. Use a default gateway on the network connection of the DHCP server in such a way that it is not using the same IP address as the router that supports the remote subnet where the clients are located.

d. Do not include this scope (that is, the one for the remote subnet) in superscopes configured for use on the same local subnet or segment where the DHCP server resides.

3. Make sure there is only one logical route between the DHCP server and the remote subnet clients.

- **See also:** DHCP Best Practices; DHCP/BOOTP Relay Agents; BOOTP and DHCP

- **Cause:** Multiple DHCP servers exist on the same local area network (LAN).

- **Solution:** Make sure that you do not configure multiple DHCP servers on the same LAN with overlapping scopes.

  You might want to rule out the possibility that one of the DHCP servers in question is a computer running Small Business Server. On a computer running Small Business Server, the DHCP Server service automatically stops when it detects another DHCP server on the LAN.

- **See also:** DHCP Best Practices; Configuring scopes

**The DHCP client appears to be affected by another problem not described above.**

- **Cause:** My problem is not described above.

- **Solution:** Search the Microsoft Web site for updated technical information that might relate to the problem you have observed. If necessary, you can obtain information and instructions that pertain to your current problem or issue.

  If you are connected to the Internet, the latest updates for the Windows Server 2003 family are available at the Microsoft Web site.

- **See also:** DHCP updated technical information; DHCP; Using the Windows Deployment and Resource Kits; Microsoft Web site

# Troubleshooting DHCP servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting DHCP servers**
What problem are you having?

- The DHCP server is stopped.
- The DHCP server is unable to provide service to clients.
- One of two DHCP servers on the same subnet is not servicing clients.
- The DHCP server appears to have suffered some data corruption or loss.
- The server appears to be affected by another problem not described above.

**The DHCP server is stopped.**

- **Cause:** The DHCP server has not been authorized to operate on the network.
- **Solution:** Authorize the DHCP server in the enterprise where it is being used.
- **See also:** Authorizing DHCP servers; Authorize a DHCP server in Active Directory; Delegate ability to authorize DHCP servers to a non-enterprise administrator
- **Cause:** Configuration details might be incorrect or missing at the server.
- **Solution:** If you have just completed setting up or administering the DHCP server, you might want to review the provided checklist to see if you have missed a crucial step in the installation process.

  To help prevent the most common types of problems, review DHCP best practices for tips on deploying and managing your servers.

  Because many DHCP failures are first detected as client-side errors, you might want to start by investigating the problem there.
- **See also:** Checklist: Installing a DHCP server; DHCP Best Practices; Troubleshooting DHCP clients
- **Cause:** The DHCP server has been stopped.
- **Solution:** Check the system event log and DHCP server audit log files for details.

  When the DHCP Server service either stops or cannot start, useful explanatory information about the source of the service failure or shutdown can generally be found in these logs.
- **See also:** Audit logging; Analyzing server log files; Using the Event Viewer snap-in; Search for specific events

**The DHCP server is unable to provide service to clients.**

- **Cause:** The server is a multihomed computer and is not providing service on one or more of its network connections.
- **Solution:** Review Windows Server 2003 DHCP binding defaults for network connections based on whether you have elected to either statically or dynamically configure TCP/IP for any or all installed connections on the server computer.

Also, review an example of multihomed DHCP server configuration to see if you have missed any critical details.

- **See also:** Multihomed DHCP servers; Selectively set DHCP server bindings for network connections

- **Cause:** Scopes or superscopes on the DHCP server have not been either configured or activated for use.

- **Solution:** Add scopes and make sure that they are correctly configured along with any DHCP scope options that need to be assigned for client use.

- **See also:** Configuring scopes; Create a new scope; Assign a scope-based option

- **Cause:** The server is located on a different subnet as some of its clients and is not providing service to clients on remote subnets.

- **Solution:** If you are using a DHCP server in a routed network, you might want to review issues related to DHCP relay agents and the appropriate use of superscopes.

- **See also:** DHCP/BOOTP Relay Agents; Using superscopes; Create a superscope

- **Cause:** The scope in use is full and can no longer lease addresses to requesting clients.

- **Solution:** If the DHCP server does not have IP addresses available to provide to its clients, it returns DHCP negative acknowledgment messages (DHCPNAKs) to them. When this occurs, consider the following possible solutions:

1. Expand the address range by increasing the **End IP address** for the current scope.

2. Create a new additional scope and a superscope, then add the current scope and the new scope to the superscope.

3. Create a new scope or extend the range. Optimally, you could renumber your current IP network. Deactivate the old scope as needed, and then configure and activate the new one.

4. Reduce the lease duration. This can help to expedite the reclaiming of lapsed scope addresses.

   Other DHCP-related procedures and techniques might also help to accelerate or ease the transition from an existing scope being retired to a new scope created to take its place at the server. These include deleting client leases from the scope being retired, excluding addresses from that scope, and then deactivating it once the new scope has been activated. This ensures that the DHCP client obtains leasing from the new scope.

- **See also:** Managing leases; Removing scopes; Delete a client lease; Exclude an address from a scope; Activate a scope; Deactivate a scope

- **Cause:** The range of IP addresses being offered by the DHCP server is in conflict with the range of addresses being offered by another DHCP server on the network.

  You can confirm that this is the likely cause if DHCP server logs still indicate that DHCP negative acknowledgment messages (DHCPNAKs) are being returned to requesting clients or if you have tried unsuccessfully at the client to renew its lease manually.

- **Solution:** Modify the scope address pool for the scopes at each DHCP server to ensure that scope IP addresses do not overlap. You can add exclusions to the scopes as necessary, delete client leases, and temporarily enable server-side conflict detection to assist in solving the problem.

- **See also:** Delete a client lease; Exclude an address from a scope; Enable address conflict detection; DHCP Best Practices

**One of two DHCP servers on the same subnet is not servicing clients.**

- **Cause:** The DHCP server is not authorized in Active Directory.

- **Solution:** If the DHCP server is a domain member, authorize the server in Active Directory.

  In some circumstances you might accidentally have a standalone server and a domain member server on the same subnet. When the standalone server detects the domain member server, it attempts to verify that it is authorized in Active Directory. Even if a domain controller resides on the same subnet as the standalone DHCP server, the DHCP server cannot verify its status with the domain controller because the DHCP server is not a domain member. When the standalone server is unable to access a domain controller to discover whether it is authorized, it stops servicing clients and displays the red icon in the DHCP console that indicates the server is unauthorized. If you want the standalone server to service clients on the subnet, remove the authorized DHCP server from the subnet.

- **See also:** Authorizing DHCP servers; Authorize a DHCP server in Active Directory; DHCP console icons reference

**The DHCP server appears to have suffered some data corruption or loss.**

- **Cause:** The DHCP server database has become corrupted or is missing server data, possibly reporting JET database errors.

- **Solution:** Use DHCP server data recovery options to restore the database and correct any of the reported errors. You can also use the Reconcile feature in the DHCP console to verify and reconcile any database inconsistencies that the server is able to find.

- **See also:** Restoring server data; The DHCP database; Reconciling scopes; Reconcile the DHCP database; Reconcile a scope

**The server appears to be affected by another problem not described above.**

- **Cause:** My problem is not described above.

- **Solution:** Search the Microsoft Web site for updated technical information that might relate to the problem you have observed. You can obtain information and instructions that pertain to your current problem or issue.

  If you have connection to the Internet, the latest updates for members of the Windows Server 2003 family are available at the Microsoft Web site.

- **See also:** DHCP updated technical information; DHCP; Using the Windows Deployment and Resource Kits; Microsoft Web site

# WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS**
Windows Internet Name Service (WINS) provides a dynamic replicated database service that can register and resolve NetBIOS names to IP addresses used on your network. The Microsoft® Windows Server 2003 family provides WINS, which enables the server computer to act as a NetBIOS name server and register and resolve names for WINS-enabled client computers on your network as described in the NetBIOS over TCP/IP standards.

- Before you begin using this computer as a WINS server, see Checklist: Installing a WINS server.
- To find features that have been moved in the Windows Server 2003 family, see New ways to do WINS tasks.
- For tips about using WINS, see WINS Best Practices.
- For help with specific tasks, see WINS How To....
- For general background information, see WINS Concepts.
- For problem-solving instructions, see WINS Troubleshooting.

# Checklist: Installing a WINS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Checklist: Installing a WINS server**

| Step | Reference |
|---|---|
| Review WINS concepts such as NetBIOS names, WINS servers, WINS clients, and replication partners. | Understanding WINS |
| Review best practices for implementing WINS on your network. | WINS Best Practices |
| Determine the number of WINS servers you need to install and where to locate each server on your network. | Planning WINS networks |
| Manually configure TCP/IP properties for the WINS server computer. | Configure TCP/IP for static addressing |
| Install a primary WINS server. | Install a WINS server |
| Open the WINS console. | Open the WINS console |
| Modify WINS server defaults, if needed. | Modifying WINS server defaults |
| Install a secondary WINS server. | Install a WINS server |
| Configure replication settings on both primary and secondary WINS servers. Choose among these replication partnership methods:<br>• Push partner only<br>• Pull partner only<br>• Push/pull partner (recommended) | Configuring WINS replication; Configure Replication |
| View and filter the WINS database view of active registrations as needed. | View WINS Records |
| Add and remove mappings in the WINS database as needed. | Manage Static Mappings; Delete or tombstone an entry in the WINS database |
| Back up, scavenge, and compact the WINS database regularly. | Manage the WINS Database |
| View WINS server statistics periodically to monitor performance. | View WINS server statistics |
| Update client configurations on your network to correctly specify primary/secondary WINS servers based on any WINS servers you have added.<br>Depending on whether you configure clients statically or dynamically for TCP/IP, either:<br>• Configure WINS-related settings in advanced TCP/IP properties.<br>• Configure and enable WINS-related options in DHCP scopes so that IP addresses of WINS servers and NetBIOS node types are provided to clients. | Assign a scope-based option; Configure TCP/IP to use WINS |
| Verify registration of client NetBIOS names in WINS as needed. | Verify WINS registration of client NetBIOS names |

| | |
|---|---|
| Troubleshoot common WINS-related problems. | WINS Troubleshooting |

# New ways to do WINS tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**New ways to do familiar tasks**

The following table lists common tasks for managing WINS servers. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Install a WINS server on this computer | **Network** in Control Panel | Windows Components Wizard.<br><br>For more information, see Install a WINS server. |
| Start WINS management tool | **Start** menu shortcut to WINS Manager listed in **Administrative Tools** | WINS.<br><br>For more information, see Open the WINS console. |
| Start or stop a WINS server | **Services** in Control Panel | Computer Management.<br><br>For more information, see Start or stop a WINS server. |
| Add a server to WINS management tool and manage it | **Add WINS Server** on the **Server** menu in WINS Manager | See Add a WINS server to the console. |
| Add and configure a replication partner for a WINS server | **Replication Partners** on the **Server** menu in WINS Manager | See Add a replication partner. |

# WINS Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Best practices**

- **Use default settings to configure WINS servers.**

  The preconfigured WINS settings provide the optimal configuration for most conditions and should be used in most WINS network installations. If you modify the default settings, be sure that the need to do so is clear and necessary, and that you understand all of the implications.

- **Avoid using static WINS entries for anything other than mission-critical servers.**

  Static WINS entries require further administrative action to ensure their successful and intended use. However, static entries can be useful for specific purposes, such as to safeguard against WINS registration of the names used by critically important servers.

  For example, a static entry in WINS can be used to prevent other computers from taking over the use of the name of a critical server while that server is unavailable.

  The biggest disadvantage of using static WINS entries is that it complicates administration of name and address changes in your network. For example, if either the IP address or the computer name of a static WINS entry changes, you might have to manually update other configurations also, such as DHCP servers, DNS servers, end systems, and Lmhosts files.

  If you decide to use static WINS entries, the following are suggestions for configuring other WINS and DHCP service properties and avoiding common problems:

  - For each IP address used in static WINS mapping, consider using a corresponding client address reservation to reserve the IP address at the DHCP server. For more information, see Using client reservations.

  - If you use static WINS entries only to support temporary changes on your network, keep the default **Overwrite unique static mappings at this server (migrate on)** setting selected in the WINS console.

    When **Overwrite unique static mappings at this server (migrate on)** is selected, any temporary unique or multihomed static entries that you add can be challenged and dynamically updated by clients. An attempt later by any client to dynamically register a unique or multihomed name entry over an existing static entry of the same name results in a challenge process.

    In the challenge, the IP address in the static mapping is compared to any IP address that the named client attempts to dynamically register in WINS. If the two addresses are different, and the static IP address is determined to no longer be active or in use, the IP mapping can be migrated (from a static to dynamic mapping), and the IP address is updated in WINS.

  - If you use static WINS entries on a permanent basis, you can disable **Overwrite unique static mappings at this server (migrate on)**. This prevents a dynamic WINS entry from overriding a static WINS entry that maps to the name and address of a critical server on your network.

- **Schedule consistency checking for an off-peak time.**

  For the Windows Server 2003 family, WINS consistency checking is available through the WINS console. Periodically, you should use this feature to check the WINS database for consistency. However, consistency checking is network and resource intensive for the WINS server computer. For this reason, run WINS consistency checks during times of low traffic, such as at night or on weekends.

- **Select Push/Pull when configuring replication partners.**

  In general, Push/Pull is a simple and effective way to ensure full WINS replication between partners. For most WINS installations, avoid the use of limited replication partnerships (Push Only or Pull Only) between WINS servers. In some large enterprise WINS networks, limited replication partnering can effectively support replication over slow wide area network (WAN) links. However, when you plan limited WINS replication, note the design and configuration when choosing to use limited replication partnerships.

- **For best results in WINS replication and convergence time, use a hub-and-spoke design model.**

  Convergence time is a critical part of WINS planning. Convergence time is the time needed to replicate a new entry in a WINS database from the WINS server that owns the entry to all other WINS servers on the network. To review the convergence time for your WINS network design, ask the following question for each of your WINS servers:

  *How long does it take for a change in WINS data at one WINS server to replicate and appear at other WINS servers in the network?*

  In most cases, the hub-and-spoke model provides a simple and effective planning method for organizations that require full and speedy convergence with minimal administrative intervention. For example, you can use this model for organizations with centralized headquarters or a corporate data center (the *hub*) and several branch offices (the *spokes*). A second, redundant hub (that is, a second WINS server in the central location) can increase the fault tolerance for WINS.

  For more information on how to configure a simple hub-and-spoke model, view the following graphic.




- **Use only as many WINS servers as you need.**

  Too many WINS servers on a network can complicate any problems that arise. Be conservative when adding WINS servers to your network, using the least number of WINS servers to support all clients while maintaining acceptable performance from each server. If you are currently designing a WINS installation that includes more than 20 WINS servers, seek assistance from Microsoft Product Support Services before finalizing your plans.

  By default, most WINS clients first try to communicate using directed point-to-point datagrams with their configured primary WINS servers. One WINS server adequately supports a small routed internetwork. However, at least two WINS servers are recommended, to provide a fault-tolerant WINS installation.

  For more information on WINS server planning, see Planning WINS networks.

- **To maximize server performance, purchase hardware with optimal disk performance characteristics to handle WINS.**

WINS causes frequent and intensive activity on server hard disks. To provide for the best performance, consider RAID solutions that improve disk-access time when you purchase new hardware to use as a WINS server. You should view WINS as part of a full performance evaluation of the server. By monitoring system hardware performance in the most demanding areas of utilization (that is, CPU, memory, and disk input/output) you make the best assessments of when a WINS server is overloaded and should be upgraded.

- **When using multihomed servers, follow well-known considerations for configuring WINS and TCP/IP properties.**

  When using multihomed WINS servers, it is best to configure all server IP addresses as replication partners with other WINS servers. Using a multihomed WINS server as a router or gateway between disjointed networks is not recommended.

- **Monitor and perform regular, offline compaction.**

  Although manual compaction of the WINS server database is not as important for the Windows Server 2003 family as it was for earlier versions, it is still useful. For disk defragmentation and improved server disk performance, use offline WINS compaction regularly (such as monthly or weekly). Then, monitor any changes to the size of the server database file, Wins.mdb, which is located by default in the *systemroot*\System32\WINS\ folder.

  At each offline compaction, it is useful to check the file size of Wins.mdb both before and after compaction, to measure growth and reduction. This information can be helpful in determining the actual benefits to using offline compaction. Based on this information, you can then gauge how often to repeat offline compaction for measurable gains.

- **Perform regular backups of the WINS database.**

  Besides using server tape backup of the WINS server computer, the WINS console offers a backup option which you can use to restore a server database if corruption occurs. For more information, see Troubleshooting WINS databases.

- **Configure clients with more WINS server IP addresses.**

  In earlier versions of Windows NT, you could manually configure clients to use only a primary and secondary WINS server. For the Windows Server 2003 family, you can configure WINS clients (either manually or through DHCP) to use up to 12 WINS servers, two of which can be used to register clients. You can either statically configure this list (using the **Internet Protocol (TCP/IP) Properties** dialog box) or distribute it to DHCP clients (using option 44). A larger list of WINS servers gives clients additional fault tolerance for WINS when their configured primary WINS server is not available.

- **Use NBTSTAT -RR to register and troubleshoot client connectivity.**

  For the Windows Server 2003 family, the NBTSTAT -RR command-line option forces immediate release and renewal of local client names. You can use this command-line option to refresh the client entries in WINS and replicate them to other replication partners without requiring a client reboot.

- **Configure each WINS server computer to point to itself.**

  Each WINS server that you install on your network must register its own set of NetBIOS unique and group names in WINS. To prevent WINS service problems that can occur when a WINS record becomes split (that is, when names registered for a particular WINS server are owned by different WINS servers), each WINS server computer should point only to its own IP address when configuring its TCP/IP properties.

# WINS How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**How to...**

- Install and Manage WINS Servers
- Modify Server Properties
- Manage WINS Console Data
- Manage WINS Clients
- View WINS Records
- Manage the WINS Database
- Configure Replication
- Manage Replication
- Use Persistent Connections
- Manage Static Mappings
- Use WINS Command-line Tools

# Install and Manage WINS Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Install and manage WINS servers**


- Install a WINS server
- Open the WINS console
- Add a WINS server to the console
- Remove a WINS server from the console
- Start or stop a WINS server
- Decommission a WINS server
- Remove or uninstall WINS

# Install a WINS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To install a WINS server**

1. Open Windows Components Wizard.
2. Under **Components**, scroll to and click **Networking Services**.
3. Click **Details**.
4. Under **Subcomponents of Networking Services**, click **Windows Internet Name Service (WINS)**, click **OK**, and then click **Next**.
5. If prompted, type the full path to the Windows Server 2003 distribution files and click **Continue**.

   Required files are copied to your hard disk, and server software can be used after restarting the system.

## Notes

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- It is strongly recommended that you manually configure the computer to use a static IP address.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
WINS How To...

# Open the WINS console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To open the WINS console**

- Open WINS.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- The WINS console is an administrative tool for managing WINS servers. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
WINS How To...

# Add a WINS server to the console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To add a WINS server to the console**

1. Open WINS.
2. In the console tree, click **WINS**.
3. On the **Action** menu, click **Add Server**.
4. In **Add Server**, under **WINS Server**, type the appropriate server information.

   Optionally, click **Browse** to locate the WINS server computer by name on your network.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- If WINS is locally installed on a server, the local WINS server is added to the console automatically. For remote WINS servers, you must add them to the console, and then save the file.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Install and Manage WINS Servers

# Remove a WINS server from the console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To remove a WINS server from the console**

1. Open WINS.
2. Click the WINS server that you want to remove.
3. On the **Action** menu, click **Delete**.
4. When prompted "Are you sure you want to remove WINS server *servername* from this list?," click **Yes**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- This procedure affects only the management console at which you are working, not the WINS server itself which continues to provide service and registration on the network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add a WINS server to the console

# Start or stop a WINS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To start or stop a WINS server**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, point to **All Tasks**, and then click one of the following:

   - To start the service, click **Start**.
   - To stop the service, click **Stop**.
   - To interrupt the service, click **Pause**.
   - To stop and then automatically restart the service, click **Restart**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- After you pause or stop the service, the **Resume** option appears and can be clicked to immediately resume WINS service.
- You can also perform most of these tasks at a command prompt by using the following commands:
  - **net start wins**
  - **net stop wins**
  - **net pause wins**
  - **net continue wins**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
WINS How To...

# Decommission a WINS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To decommission a WINS server**

1. Before decommissioning a WINS server, make sure that any computers previously configured to be WINS clients of this server are reconfigured to point to other servers as their primary or secondary WINS servers. Reconfiguration is necessary only if these clients are to continue using WINS to register and resolve network names.

2. Open WINS.

3. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

4. On the **Action** menu, click **Delete Owner**.

5. In **Delete Owner**, in **Delete this owner**, click the IP address for the WINS server you want to decommission.

   If the WINS server is not running locally on this computer, it might take awhile to load the records for the selected server.

6. In **Use this operation to delete the selected owner and its records**, click **Replicate deletion to other WINS servers (tombstone)**, and then click **OK**.

   When prompted to confirm tombstoning, click **Yes**.

7. In the console tree, click **Replication Partners**.

8. On the **Action** menu, click **Replicate Now**.

9. Once you have verified that records tombstoned in step 6 have been replicated to other partner servers, stop and remove WINS on the decommissioned server.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

- If the WINS server you want to decommission does not appear in the console tree, you can add it. For more information, see Related Topics.

- *Tombstoning* ensures that the WINS servers that are replication partners with the decommissioned WINS server are updated properly so that they delete the records. If tombstone status is not properly replicated, you can manually delete the records at each WINS server on which tombstone replication failed.

- WINS supports remote record deletion for servers running Windows NT Server 4.0, if updated to Service Pack 4 (SP4) or later.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
View WINS server statistics
Add a WINS server to the console
Remove a WINS server from the console
Remove or uninstall WINS

# Remove or uninstall WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To remove or uninstall WINS**

1. Stop the WINS service on this computer. For more information, see Related Topics.
2. Open **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. In **Windows Components**, select the **Networking Services** check box in the list of components, and then click **Details**.
5. In **Subcomponents of Networking Services**, scroll to **Windows Internet Name Service (WINS)**, and clear the check box.

## Notes

- To open Add/Remove Windows Components, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- When you remove WINS, all WINS files are deleted from the server except program files that are in use. When the system is restarted, these files are removed as well.
- If you do not stop the WINS service in step 1, this procedure only deletes the WINS server registry entries.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Start or stop a WINS server

# Modify Server Properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Modify server properties**

- Modify backup settings
- Modify name record settings
- Modify WINS logging properties
- Modify burst handling properties
- Set the WINS database backup path
- Set the WINS database path
- Increase the starting version count
- Disable LAN Manager compatibility

# Modify backup settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To modify backup settings**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, under **Database backup**, modify default backup settings as needed.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

### Note

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the WINS database path
Back up the WINS database
Start or stop a WINS server
Restore the WINS database

# Modify name record settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To modify name record settings**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Intervals** tab.
5. Modify default settings as needed.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- For most networks the default interval settings, **Renewinterval**, **Extinction interval**, **Extinction timeout**, and **Verificationinterval**, are acceptable and do not require modification.
- As an option, you can click **Restore Default** to restore WINS default intervals used for registering, renewing, and verifying name records in the database.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modifying WINS server defaults

# Modify WINS logging properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To modify WINS logging properties**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. Select the **Log detailed events to Windows event log** check box.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- When **Log detailed events to Windows event log** is checked, WINS logs additional detailed information to the system event log, which can be viewed by using Event Viewer. Because this feature can degrade system performance when enabled, it is turned off by default.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modifying WINS server defaults

# Modify burst handling properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To modify burst handling properties**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Enable burst handling**, modify default settings as needed.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, burst handling is enabled and the burst queue sized to **Medium**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Burst handling

# Set the WINS database backup path

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To set the WINS database backup path**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. On the **General** tab, in **Database backup**, select **Default backup path**, and type the path to the folder where you want to back up the WINS database. Or, use **Browse** to select a location.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- You can also set the WINS database backup path while performing a backup. For more information, see Related Topics.
- There is not a default database backup path, and database backups do not occur by default. After a database backup path is set, backups occur every twenty four hours.
- When you choose a backup folder location by typing a path or using the **Browse** button, WINS automatically adds two subfolders to the path, Wins_bak and New. For example, if you type C:\Backups, the folder where WINS stores the backup database is C:\Backups\Wins_bak\New.
- This procedure describes how to specify the path that WINS uses when backing up the database. It does not set the path used by WINS to store the primary database on the applicable server. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Back up the WINS database
Restore the WINS database
Modify backup settings
The WINS database

# Set the WINS database path

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To set the WINS database path**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. In **Database path**, modify the database path as needed.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**.

## Caution

- This setting affects WINS back up and restore operations. If you change the path used here, be sure to run WINS backup again to archive the database with the updated path information. Otherwise, WINS cannot restore previously archived copies of the database where another path was specified, for either the WINS database path or the WINS database backup path.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- This procedure describes how to set the path used by WINS to store the database on the applicable server. It does not specify the path that WINS uses when backing up the database. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Back up the WINS database
Restore the WINS database
Modify backup settings
Set the WINS database backup path
The WINS database

# Increase the starting version count

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To increase the starting version count**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. For **Starting version ID (hexadecimal)**, enter a new number to be used by WINS as the new starting version ID when WINS is next restarted.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Because a higher version ID is used by WINS during replication to determine between two servers storing the same record which version is more recent, increasing the starting version ID count can be useful in some cases where you are attempting to affect a forced update and replication of names between servers. For more information, see Related Topics.
- If you enter a number that is less than the actual value of the highest version ID currently in use by the selected server, WINS might increase the actual number used for its start version count when service is restarted.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage the WINS Database
Troubleshooting WINS databases

# Disable LAN Manager compatibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To disable LAN Manager compatibility**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. Click **Advanced**.
5. Clear the **Use computer names that are compatible with LAN Manager** check box.

**Note**

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modifying WINS server defaults
NetBIOS names reference
Network security: LAN Manager authentication level
Security information for IAS

# Manage WINS Console Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Manage WINS console data**

- View WINS server statistics
- Reset WINS server statistics
- Modify auto-refresh of WINS statistics
- Export WINS console data to a text file
- Modify the WINS console view

# View WINS server statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To view WINS server statistics**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Display Server Statistics**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- To update the display while viewing WINS statistics, click **Refresh**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Reset WINS server statistics
Modify auto-refresh of WINS statistics

# Reset WINS server statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To reset WINS server statistics**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Display Server Statistics**.
4. Click **Reset**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- To update the display while viewing WINS statistics, click **Refresh**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
View WINS server statistics

# Modify auto-refresh of WINS statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To modify auto-refresh of WINS statistics**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. Under the **Automatically update statistics every** check box, modify the refresh interval in hours, minutes, or seconds, as needed.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, statistics automatically refresh every 10 minutes. As an option, you can also disable this feature by clearing the **Automatically update statistics every** check box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View WINS server statistics
Reset WINS server statistics
Modifying WINS server defaults

# Export WINS console data to a text file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To export WINS console data to a text file**

1. Open WINS.
2. In the console tree, click the node item for which you want to export all currently displayed console data.

   For example, if you want to export all console data about a specific WINS server, click this server in the console tree.
3. On the **Action** menu, click **Export List**.
4. Enter a name for the text (*.txt) file to be used to export the output of this operation.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- For the **Active Registrations** node, you need to use **Display Records** to query for and display records for the applicable server prior to performing this procedure. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modifying WINS server defaults
Search for names that begin with a specified set of characters
Filter the WINS database view
Show all records in the WINS database

# Modify the WINS console view

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To modify the WINS console view**

1. Open WINS.
2. In the console tree, click **WINS**.
3. On the **Action** menu, click **Properties**.
4. Modify console defaults to specify how WINS servers added to the console are displayed.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, the WINS console orders servers by name and uses only short (host-only) DNS names when displaying WINS servers that are added to the console. In addition, WINS does not validate its cache of WINS servers, which can prolong the length of time needed to load the WINS console when started.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modifying WINS server defaults

# Manage WINS Clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Manage WINS clients**

- Configure WINS for non-DHCP clients
- Configure WINS for DHCP-enabled clients
- Disable WINS/NetBT name resolution
- Verify WINS registration of client NetBIOS names

# Configure WINS for non-DHCP clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To configure WINS for non-DHCP clients**

- WINS servers must be manually added for each network connection that does not use DHCP. For more information, see Related Topics.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you manually configure a list of WINS servers for the connection that later leases and obtains its IP address using DHCP, the list you configure supercedes any WINS servers returned in an options list provided by the DHCP server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configure TCP/IP to use WINS

# Configure WINS for DHCP-enabled clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To configure WINS for DHCP-enabled clients**

- Clients running Windows 2000, Windows XP, or the Windows Server 2003 family are configured by default to use a DHCP server to obtain an IP address. No further client-side configuration is needed.

  Instead, configure your DHCP server to assign WINS-related options needed to enable DHCP-enabled clients on your network to use WINS. For more information, see Related Topics.

**Note**

- In most cases, WINS clients only need to be provided the following NetBIOS over TCP/IP (NetBT) option types for basic WINS configuration: 44 (WINS servers) and 46 (WINS node type). These options must be configured at the DHCP server providing each client its dynamically leased IP address.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage Options and Classes
DHCP options

# Disable WINS/NetBT name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To disable WINS/NetBT name resolution**

1. Open Network Connections.
2. Click the Local Area Connection to be statically configured and then, from the **File** menu, click **Properties**.
3. In the list of components, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Advanced**.
5. Click the **WINS** tab.
6. Click **Disable NetBIOS over TCP/IP**.

   Optionally, you can select **Use NetBIOS setting from the DHCP server** if you are using a DHCP server that can selectively enable and disable NetBIOS configuration through DHCP option types.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- NetBIOS over TCP/IP can also be disabled for computers running Windows XP using DHCP option types supported by the Windows Server 2003 DHCP Server service. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DHCP options

# Verify WINS registration of client NetBIOS names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To verify WINS registration of client NetBIOS names**

1. On the WINS client computer, at the command prompt, type **nbtstat -n**.

   The local NetBIOS names table for the WINS client computer is listed. Verify that each name indicates **Registered** in the **Status** column.

2. The **Status** column should indicate a status of either Registered, Registering, or Conflict.

   For cases where the name is shown to be either Registered or Registering, you can type **nbtstat -RR** to release and refresh registration of the local NetBIOS names of the client with WINS.

   For cases where the Status column indicates a conflict for the name, try either of the following to resolve the conflicted name state at the client computer:

   - For WINS clients running Windows 2000, Windows XP, or the Windows Server 2003 family, you can selectively disconnect and reconnect the network connection in the Network Connections folder to force release and refresh of all client NetBIOS names in WINS.

     For more information, see Related Topics.

   - For earlier versions of Windows, shut down and restart the system to achieve similar results.

3. After the release/refresh process is completed, repeat **nbtstat -n** to verify that all local names have a status of Registered in WINS.

### Notes

- For name conflicts, you may only need to perform step 2 if the name in conflict is critically necessary for the WINS client to maintain visibility or locate other computers on the network. Names that fit this description might include the [00h] and [20h] type names, if they match the client NetBIOS computer name. The need to resolve a name conflict for other types of names might be less critical to support needs.

- The **nbtstat** command can also be used to perform other useful procedures related to WINS client support. For more information, see Related Topics.

- When you type **nbtstat -n** at a command prompt and the local NetBIOS name table is displayed, the status column should indicate that each name is registered, such as shown in the sample command output below:

```
Node IpAddress: [10.0.0.1] Scope Id: []

NetBIOS Local Name Table
Name                Type         Status
---------------------------------------------------
MY_PC           <20>  UNIQUE      Registered
MY_PC           <00>  UNIQUE      Registered
WORKGROUP       <00>  GROUP       Registered
MY_PC           <03>  UNIQUE      Registered
ADMINISTRATOR   <03>  UNIQUE      Registered
WORKGROUP       <1E>  GROUP       Registered
WORKGROUP       <1D>  UNIQUE      Registered
```

```
..__MYBROWSE__. <01>  GROUP       Registered
```

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a remote network
Disconnect from a remote network
Command-line utilities
WINS Resolution

# View WINS Records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**View WINS records**

- Start a record search
- Stop a record search
- Search for names that begin with a specified set of characters
- Search for records by IP address and subnet mask
- Show all records in the WINS database
- View properties for a selected WINS record
- View or download the WINS database by owner
- Filter the WINS database view
- Clear all WINS data filters
- Add, modify, or remove a filter type

# Start a record search

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To start a record search**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display Records**.
4. Select the appropriate filters and type in search criteria, if necessary.
5. Click **Find Now**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, in the **Display Records** dialog box, all NetBIOS name record owners are selected on the **Record Owners** tab, and all NetBIOS name suffix record types are selected on the **Record Types** tab. If you have a large network or a large number of WINS servers, running a query without refining these selections could cause degradation of network performance due to the large data load transfer.
- You must select at least one record owner and at least one record type in order for any search operation to successfully return one or more name records.
- If you select the **Enable result caching** check box, the search results will be downloaded from the WINS server and stored in local random access memory (RAM). Because search results are stored in RAM, subsequent record searches on this name pattern will be faster. If a large number of records are returned, however, caching of records might take a long time and could adversely affect overall network performance.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Finding and viewing WINS records
Stop a record search

# Stop a record search

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To stop a record search**

1. On the **Action** menu, click **Stop Loading Records**.

   A message box will ask, "Are you sure you want to stop downloading the records from the database?"

2. Click **Yes**.

## Notes

- This procedure assumes that both the WINS console and the **Display Records** dialog box are open.
- A record search must be currently operating in order for this procedure to work.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start a record search
Finding and viewing WINS records

# Search for names that begin with a specified set of characters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To search for names that begin with a specified set of characters**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display Records**.
4. Select the **Filter records matching this Name pattern** check box.
5. In **Filter records matching this Name pattern**, type one or more characters to be used in the search.
6. Click **Find Now** to begin the search.

   WINS searches the server database and displays any names that begin with the specified characters in uppercase.

**Notes**

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Any lowercase text entered in **Filter records matching this Name pattern** is converted to uppercase when WINS performs the record search. WINS does uppercase searches only and does not display any records that contain lowercase characters.
- By default, in the **Display Records** dialog box, all NetBIOS name record owners are selected on the **Record Owners** tab, and all NetBIOS name suffix record types are selected on the **Record Types** tab. If you have a large network or a large number of WINS servers, running a query without refining these selections could cause degradation of network performance due to the large data load transfer.
- You must select at least one record owner and at least one record type in order for the search operation to successfully return one or more name records.
- If you select the **Enable result caching** check box, the search results will be downloaded from the WINS server and stored in local random access memory (RAM). Because search results are stored in RAM, subsequent record searches on this name pattern will be faster. If a large number of records are returned, however, caching of records might take a long time and could adversely affect overall network performance.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Show all records in the WINS database
Filter the WINS database view
Finding and viewing WINS records

# Search for records by IP address and subnet mask

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

#### To search for records by IP address and subnet mask

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display Records**.
4. Select the **Filter records matching this IP address** check box.
5. In **Filter records matching this IP address**, type a valid IP address.
6. Select the **Match the IP address based on this subnet mask** check box.
7. In **Match the IP address based on this subnet mask**, type a valid subnet mask.
8. Click **Find Now** to begin the search.

   WINS searches the server database and displays any records in which the network ID portion of the IP address matches the parameters entered into the **Display Records** dialog box.

### Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, in the **Display Records** dialog box, all NetBIOS name record owners are selected on the **Record Owners** tab, and all NetBIOS name suffix record types are selected on the **Record Types** tab. If you have a large network or a large number of WINS servers, running a query without refining these selections could cause degradation of network performance due to the large data load transfer.
- You must select at least one record owner and at least one record type in order for the search operation to successfully return one or more name records.
- If you select the **Enable result caching** check box, the search results will be downloaded from the WINS server and stored in local random access memory (RAM). Because search results are stored in RAM, subsequent record searches on this name pattern will be faster. If a large number of records are returned, however, caching of records might take a long time and could adversely affect overall network performance.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

#### See Also

**Concepts**

Finding and viewing WINS records
Stop a record search
View or download the WINS database by owner

# Show all records in the WINS database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To show all records in the WINS database**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display Records**.
4. Click **Find Now**.

   The entire WINS database appears in the details pane.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, in the **Display Records** dialog box, all NetBIOS name record owners are selected on the **Record Owners** tab, and all NetBIOS name suffix record types are selected on the **Record Types** tab. If you have a large network or a large number of WINS servers, running a query without refining these selections could cause degradation of network performance due to the large data load transfer.
- You must select at least one record owner and at least one record type in order for the search operation to successfully return one or more name records.
- If you select the **Enable result caching** check box, the search results will be downloaded from the WINS server and stored in local random access memory (RAM). Because search results are stored in RAM, subsequent record searches on this name pattern will be faster. If a large number of records are returned, however, caching of records might take a long time and could adversely affect overall network performance.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Search for names that begin with a specified set of characters
Show all records in the WINS database
Filter the WINS database view
Finding and viewing WINS records

# View properties for a selected WINS record

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To view properties for a selected WINS record**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display records**.
4. To display records in the details pane, perform one of the following actions:

   - To find a record by NetBIOS name, select the **Filter records matching this Name pattern** check box, type the leading characters for the name or names that you want to view, and then click **Find Now**.
   - To find a record by IP address, select the **Filter records matching this IP address** check box, type the IP address, and then click **Find Now**.
   - To find a record by owner, click the **Record Owners** tab. In the **Display records for these owners** list, select the check box for any of the name record owners that you want included in filtered viewing of the WINS database, clear the check box for any name record owners you want excluded from filtered viewing of the WINS database, and then click **Find Now**.

5. In the details pane, click the record that you want to view.
6. On the **Action** menu, click **Properties**.

**Notes**

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, in the **Display Records** dialog box, all NetBIOS name record owners are selected on the **Record Owners** tab, and all NetBIOS name suffix record types are selected on the **Record Types** tab. If you have a large network or a large number of WINS servers, running a query without refining these selections could cause degradation of network performance due to the large data load transfer.
- You must select at least one record owner and at least one record type in order for the search operation to successfully return one or more name records.
- If you select the **Enable result caching** check box, the search results will be downloaded from the WINS server and stored in local random access memory (RAM). Because search results are stored in RAM, subsequent record searches on this name pattern will be faster. If a large number of records are returned, however, caching of records might take a long time and could adversely affect overall network performance.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Search for names that begin with a specified set of characters
Show all records in the WINS database
Filter the WINS database view
Clear all WINS data filters
Add, modify, or remove a filter type
Finding and viewing WINS records

# View or download the WINS database by owner

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To view or download the WINS database by owner**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display Records**.
4. Click the **Record Owners** tab.
5. In the **Display records for these owners** list, perform one or more of the following actions:

   - Select the check box for any name record owners that you want to include in the filtered view of the WINS database.
   - Clear the check box for any name record owners that you want to exclude from the filtered view of the WINS database.
   - Click **Select Local** to query records only for local owners.
   - Click **Select All** to query records for all owners.

6. Click **Find Now**.

   After you click **Find Now**, the filtered WINS database appears in the details pane and shows only the records that are registered to the selected owners.

**Notes**

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, in the **Display Records** dialog box, all NetBIOS name record owners are selected on the **Record Owners** tab, and all NetBIOS name suffix record types are selected on the **Record Types** tab. If you have a large network or a large number of WINS servers, running a query without refining these selections could cause degradation of network performance due to the large data load transfer.
- You must select at least one record owner and at least one record type in order for the search operation to successfully return one or more name records.
- If you select the **Enable result caching** check box, the search results will be downloaded from the WINS server and stored in local random access memory (RAM). Because search results are stored in RAM, subsequent record searches on this name pattern will be faster. If a large number of records are returned, however, caching of records might take a long time and could adversely affect overall network performance.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Filter the WINS database view
Show all records in the WINS database
Finding and viewing WINS records

# Filter the WINS database view

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To filter the WINS database view**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display records**.
4. Click the **Record Types** tab.
5. Clear the check box for any name record types you want excluded from filtered viewing of the WINS database.

   When you click **Find Now**, the filtered WINS database appears in the details pane.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, in the **Display Records** dialog box, all NetBIOS name record owners are selected on the **Record Owners** tab, and all NetBIOS name suffix record types are selected on the **Record Types** tab. If you have a large network or a large number of WINS servers, running a query without refining these selections could cause degradation of network performance due to the large data load transfer.
- You must select at least one record owner and at least one record type in order for the search operation to successfully return one or more name records.
- If you select the **Enable result caching** check box, the search results will be downloaded from the WINS server and stored in local random access memory (RAM). Because search results are stored in RAM, subsequent record searches on this name pattern will be faster. If a large number of records are returned, however, caching of records might take a long time and could adversely affect overall network performance.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Clear all WINS data filters
Add, modify, or remove a filter type
Finding and viewing WINS records

# Clear all WINS data filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To clear all WINS data filters**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display Records**.
4. Click the **Record Types** tab.
5. Click **Clear All**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, all predefined NetBIOS name type filters are included in the filter mask. To restore the default settings, click **Select all**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Filter the WINS database view
Add, modify, or remove a filter type
Finding and viewing WINS records

# Add, modify, or remove a filter type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To add, modify, or remove a filter type**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display Records**.
4. Click the **Record Types** tab.
5. Click either **Add**, **Edit**, or **Delete** to add, modify, or remove name type filters.
6. When adding a new filter type, click **Add**, and complete the following for **Add Record Type**:

   - For **ID (hexadecimal)**, specify a unique 2-digit hexadecimal number that can be used to identify a distinct set of records in the WINS database.
   - For **Description**, enter descriptive text that indicates how the new filter type is to be used.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Only user-defined filters can be modified or deleted. Predefined filters provided for default use by WINS cannot be modified or deleted.
- For any search operation to successfully return one or more name records, at least one record type must be selected.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Filter the WINS database view
View or download the WINS database by owner
Finding and viewing WINS records

# Manage the WINS Database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Manage the WINS database**

- Scavenge the WINS database
- Back up the WINS database
- Restore the WINS database
- Delete or tombstone an entry in the WINS database
- Delete an owner from the WINS database
- Enable periodic consistency checking of the WINS database
- Manually check database consistency
- Check version number consistency between servers

# Scavenge the WINS database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To scavenge the WINS database**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Scavenge Database**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Like other databases, you must periodically maintain the WINS database to remove outdated information. This process is called *scavenging*, which removes old and released entries that were registered at another WINS server and replicated to the server being maintained.
- WINS automatically scavenges the server database at intervals you specify in server properties.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
WINS Database Management
Maintaining the WINS database

# Back up the WINS database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To back up the WINS database**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Back Up Database**. The **Browse for Folder** dialog box opens.
4. Select the location for the WINS database backup. Click **OK**.
5. After backup is completed, click **OK**.

### Important

- Do not specify a network drive as the backup location. In addition, if you change the WINS backup or database path in server properties, perform new backups to ensure successful future restorations of the WINS database.

### Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- The WINS console provides backup tools so that you can back up and restore the WINS database. After you specify a backup folder for the database, WINS performs complete database backups every twenty four hours, using the specified folder.
- WINS can also be configured to back up the database automatically when the service is stopped or the server computer is shut down properly. For more information, see Related Topics.
- When WINS backs up the server database, it creates a **\wins_bak\new** folder under the backup folder you have specified for **Default backup path** in ***Server* Properties**. Actual backups of the WINS database (Wins.mdb) are stored in this folder.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modify backup settings
Set the WINS database path
Set the WINS database backup path
Restore the WINS database
The WINS database

# Restore the WINS database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

#### To restore the WINS database

1. Stop WINS.

   For more information, see Related Topics.

2. Delete all files in the folder path on the WINS server computer where you are restoring the database.

   This path is determined by the current setting of **Database path**, which is located on the **Advanced** tab in ***Server* Properties**.

3. Open WINS.

4. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

5. On the **Action** menu, click **Restore Database**.

6. In the **Browse for Folder** dialog box, click the same folder (such as *C:*\) that was used previously to back up the local WINS database, and then click **OK**.

### Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- In step 1, wait for WINS to completely stop before proceeding. A busy server might require several minutes before it can stop WINS, and you can proceed with restoring the database.
- When restoring the database, WINS requires that the backup path used for restoration match the path that was originally specified in ***Server* Properties** as the **Backup path**.
- WINS does not permit you to restore the server databases remotely from another computer.
- The restore option is only available when viewing a computer on which WINS is stopped. In some cases, you might need to refresh the WINS console to activate the restore option. If necessary, click **Refresh** from the **Action** menu to refresh the server node so that the WINS console detects that WINS is stopped on the server.
- In case of database corruption, the WINS database can be restored from the backup. You can sometimes repair WINS database corruption by incrementing the starting version count for the WINS server and restarting the server.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Set the WINS database path
Set the WINS database backup path
Back up the WINS database
Increase the starting version count
Start or stop a WINS server
The WINS database

# Delete or tombstone an entry in the WINS database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To delete or tombstone an entry in the WINS database**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Display Records** to query for and display the names you want to delete from the database.
4. In the details pane, click the entry you want to delete.

   You can hold down the CTRL or SHIFT keys to select multiple entries.

5. On the **Action** menu, click **Delete**.
6. In the **Delete Record** dialog box, click one of the following options for deleting the selected records:

   - **Delete the record only from this server**.
   - **Replicate deletion of the record to other servers (tombstone)**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- *Tombstoning* marks records as inactive for a period of time. When the WINS database on this server replicates these records to other WINS servers, the other servers preserve the tombstone state and can eventually remove these entries from a replicated WINS network of server databases.
- In step 6, if you click **Delete the record only from this server** and another WINS server owns the records, then deleted records can reappear after WINS replication occurs with other partner servers. To avoid this, click **Replicate deletion of the record to other servers (tombstone)**.
- Once an entry has been tombstoned and its status is replicated to other servers, the tombstoned entries can be removed from WINS, in either of two ways:

  - Automatically, using the **Extinction Interval**, which is set from the **Intervals** tab in WINS server properties.
  - Manually, by performing a scavenging operation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Deleting and tombstoning records

Scavenge the WINS database

# Delete an owner from the WINS database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To delete an owner from the WINS database**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Delete Owner**.
4. In **Delete Owner**, click the owner you want to delete.
5. Select from the following options for deleting records in the WINS database associated with the owner you are deleting:

   - To delete the owner and all of its records from this WINS server only, click **Delete from this server only**.
   - To replicate deletion of the selected owner and all its records at this and other WINS servers, click **Replicate deletion to other servers (tombstone)**.

**Notes**

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- *Tombstoning* marks records as inactive for a period of time. When the WINS database on this server replicates these records to other WINS servers, the other servers preserve the tombstone state and can eventually remove these entries from a replicated WINS network of server databases.
- In step 5, if you click **Delete from this server only**, records first introduced by the deleted owner can reappear after WINS replication with other partner servers that have replicated these records. To avoid this, click **Replicate deletion to other servers (tombstone)**.
- Once an entry has been tombstoned and its status is replicated to other servers, the tombstoned entries can be removed from WINS, in either of two ways:

  - Automatically, using the **Extinction Interval**, which is set from the **Intervals** tab in WINS server properties.
  - Manually, by performing a scavenging operation.

  For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Deleting and tombstoning records

Scavenge the WINS database

Modify name record settings

# Enable periodic consistency checking of the WINS database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable periodic consistency checking of the WINS database**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Properties**.
4. Click the **Database Verification** tab.
5. In **Database Verification settings**, select the **Verify database consistency every** *number of hours* **hours** check box.
6. If necessary, modify default settings to be used when consistency checks are performed.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Because consistency checking might affect WINS server performance, this feature is not enabled by default.
- When enabled, the default settings for this feature provide for up to 30,000 records in the WINS server database to be checked on a daily basis, beginning at 2:00AM each night. During each of these scheduled consistency checks, the selected WINS server compares local copies of records it stores with the same record as stored at the WINS server listed as the owner for each record.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checking the WINS database for consistency
Scavenge the WINS database

# Manually check database consistency

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To manually check database consistency**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Verify Database Consistency**.
4. When prompted if you want to continue, click **Yes** to start consistency checking.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Checking consistency can be resource-intensive when there is a large number of names to be checked or a large number of WINS servers involved. You might want to perform this only during off-peak hours to minimize the impact on your WINS servers and network.
- Consistency checking can also be scheduled through server properties. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable periodic consistency checking of the WINS database
Checking the WINS database for consistency

# Check version number consistency between servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To check version number consistency between servers**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click **Verify Version ID Consistency**.
4. When prompted if you want to continue, click **Yes** to start this operation.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- This operation gets the Owner-to-Highest-Version-ID mappings from different WINS servers. It checks the consistency of the WINS databases at all tested servers by ensuring that each server always has the highest version ID among the network of WINS servers for any records that it owns.
- Checking version ID number consistency can be resource-intensive when there are a large number of WINS servers tested.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Enable periodic consistency checking of the WINS database
Checking the WINS database for consistency

# Configure Replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Configure replication**

- Add a replication partner
- Delete a replication partner
- Enable replication with non-partner WINS servers
- Modify push partner properties
- Modify pull partner properties
- Enable automatic partner configuration

# Add a replication partner

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To add a replication partner**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **New Replication Partner**.
4. In **New Replication Partner**, type the name or IP address of the WINS server to add as a replication partner.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, replication partners are configured as push/pull type partners. Once you have added a partner, however, you can change it to be a Push only or Pull only type partner. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change a replication partner type
Delete a replication partner
Modify push partner properties
Modify pull partner properties
WINS replication overview

# Delete a replication partner

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To delete a replication partner**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. In the details pane, click the replication partner you want to remove.
4. On the **Action** menu, click **Delete**.
5. When prompted, click **Yes** to confirm that you want to delete the partner.

### Note

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a replication partner
Change a replication partner type
WINS replication overview

# Enable replication with non-partner WINS servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable replication with non-partner WINS servers**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Properties**.
4. Clear the **Replicate only with partners** check box, and then click **Apply**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, WINS only replicates with other WINS servers configured as replication partners.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
WINS replication overview
Configuring WINS replication

# Modify push partner properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To modify push partner properties**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Properties**.
4. Click the **Push Replication** tab.
5. As needed, reconfigure push partner defaults to be used by this server when it performs replication with any of its newly added push partners.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

### Note

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Push partners
Configuring WINS replication

# Modify pull partner properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To modify pull partner properties**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Properties**.
4. Click the **Pull Replication** tab.
5. As needed, reconfigure pull partner defaults to be used by this server when it performs replication with any of its newly added pull partners.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

### Note

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Pull partners
Blocking replication partners
Configuring WINS replication

# Enable automatic partner configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable automatic partner configuration**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Properties**.
4. Click the **Advanced** tab.
5. Select the **Enable automatic partner configuration** check box.
6. If necessary, modify the multicast interval and multicast Time-to-Live (TTL) parameters to be used by this server when automatic partner configuration is enabled for use by WINS.

   To view a description of a dialog box item, right-click the item, and then click **What's This?**

## Note

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configuring WINS replication
WINS replication overview

# Manage Replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Manage replication**

- Start replication immediately
- Send a replication trigger to a selected partner
- Send a replication trigger to any specified WINS server
- Accept replication of records owned by a specified server
- Block replication of records owned by a specified server
- Change a replication partner type

# Start replication immediately

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To start replication immediately**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Replicate Now**.
4. When prompted, click **Yes** to start replication immediately.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- This operation queues replication requests for triggering replication with any of its push and pull partner servers. For this operation to have any affect, replication partners must already be added at this server.
- Using Event Viewer, you can monitor the system event log at this server for more information about the status of when actual replication is completed between this server and its partner servers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add a replication partner
Delete a replication partner
Send a replication trigger to any specified WINS server

# Send a replication trigger to a selected partner

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To send a replication trigger to a selected partner**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. In the details pane, click the partner that you want to start replication with.
4. On the **Action** menu, click either **Start Push Replication** or **Start Pull Replication** depending on whether you want to start push or pull replication to the selected partner.
5. Confirm starting of either push or pull replication to the selected partner.
6. If using pull replication, in **Start Pull Replication**, click **Yes** to confirm.
7. If using push replication, in **Start Push Replication**, choose a push replication method, and then click **OK**:

   a. Click **Start for this partner only** to start push replication for the selected partner only.

   b. Click **Propagate to all partners** to request that the selected partner in turn propagate a push replication request to its configured partners.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- When this procedure is completed, replication with the partner server might not start immediately. Actual replication depends on server activity and the priority of other requests queued at the selected partner.
- By default, WINS only enables replication between configured replication partners. To replicate with non-partners, see Related Topics.
- When a push replication is started using this procedure, actual replication requires the selected partner to complete a subsequent pull replication. This occurs because WINS replication is essentially a pull operation. For more examples and information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start replication immediately

Send a replication trigger to any specified WINS server
Enable replication with non-partner WINS servers
WINS replication overview
Pull partners
Blocking replication partners
Configuring WINS replication

# Send a replication trigger to any specified WINS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To send a replication trigger to any specified WINS server**

1. Open WINS.
2. In the console tree, click the applicable WINS server.

   **Where?**

   - WINS/*applicable WINS server*

3. On the **Action** menu, click either **Start Push Replication** or **Start Pull Replication**.
4. In Start Replication, for **WINS server**, type or browse for the specified WINS server with which you want to start replication, and then click **OK**.
5. Confirm starting of either push or pull replication.
6. If using pull replication, in **Start Pull Replication**, click **Yes** to confirm.
7. If using push replication, in **Start Push Replication**, choose a push replication method, and then click **OK**:

   a. Click **Start for this partner only** to start push replication for the selected server only.

   b. Click **Propagate to all partners** to request that the selected server in turn propagate a push replication request to its configured partners.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- By default, WINS only enables replication between configured replication partners. To use this procedure, first configure the server to enable it to replicate with non-partners. For more information, see Related Topics.
- When this procedure is completed, replication with the specified server might not start immediately. Actual replication depends on server activity and the priority of other requests queued at the server.
- When a push replication request is started using this procedure, actual replication requires the specified WINS server to complete a subsequent pull replication. This occurs because WINS replication is essentially a pull operation. For more examples and information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable replication with non-partner WINS servers

Start replication immediately
Send a replication trigger to a selected partner
WINS replication overview
Pull partners
Blocking replication partners
Configuring WINS replication

# Accept replication of records owned by a specified server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To accept replication of records owned by a specified server**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Properties**.
4. On the **Advanced** tab, click **Accept records only for these owners**, and then click **Add**.
5. In **Add Server**, type the NetBIOS name or IP address of the owner (a WINS server) from which you want to replicate records to this server.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Once a server has been added to the list of accepted partners, any records owned by the listed servers are entered into the local database during subsequent replication cycles with other partner servers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Blocking replication partners
WINS replication overview
Deleting and tombstoning records
Accepting replication partners

# Block replication of records owned by a specified server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To block replication of records owned by a specified server**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Properties**.
4. On the **Advanced** tab, click **Block records for these owners**, and then click **Add**.
5. In **Add Server**, type the NetBIOS name or IP address of the owner (a WINS server) from which you want to block records from being replicated to this server.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Once a server has been added to the list of blocked partners, any records owned by the listed servers are not entered into the local database during subsequent replication cycles with other partner servers.
- This feature was previously implemented as a registry-only configuration in Windows NT Server, using the PersonaNonGrata registry parameter.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Blocking replication partners
WINS replication overview
Deleting and tombstoning records
Accepting replication partners
Accept replication of records owned by a specified server

# Change a replication partner type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To change a replication partner type**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. In the details pane, click the partner for which you want to change the type of replication it uses with the applicable WINS server.
4. On the **Action** menu, click **Properties**.
5. Click the **Advanced** tab.
6. For **Replication partner type**, select a new type such as **Push**, **Pull**, or **Push/Pull**.

### Note

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Pull partners
Blocking replication partners
Push partners
WINS replication overview
WINS Best Practices
Configuring WINS replication

# Use Persistent Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Use persistent connections**

- Enable the use of persistent connections with a partner server
- Enable the use of persistent connections with pull partners
- Enable the use of persistent connections with push partners

# Enable the use of persistent connections with a partner server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable the use of persistent connections with a partner server**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. In the details pane, click the partner server for which you want to enable the use of persistent connections for replication purposes.
4. On the **Action** menu, click **Properties**.
5. Click the **Advanced** tab.
6. Select the appropriate **Use persistent connection for replication** check box, depending on whether the partner server is a **Push**, **Pull**, or **Push/Pull** partner of this server.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- In step 6, if the partner is configured as a **Push/Pull** partner of this server, you can select both **Use persistent connection for replication** check boxes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Persistent connections

# Enable the use of persistent connections with pull partners

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable the use of persistent connections with pull partners**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Properties**.
4. Click the **Pull Replication** tab.
5. Select the **Use persistent connections for pull replication partners** check box.

### Note

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable the use of persistent connections with a partner server
Enable the use of persistent connections with push partners
Persistent connections
Pull partners
Blocking replication partners

# Enable the use of persistent connections with push partners

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable the use of persistent connections with push partners**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Properties**.
4. Click the **Push Replication** tab.
5. Select the **Use persistent connections for push replication partners** check box.

### Note

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable the use of persistent connections with a partner server
Enable the use of persistent connections with pull partners
Persistent connections
Push partners

# Manage Static Mappings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Manage static mappings**

- Add a static mapping entry
- Edit a static mapping entry
- Import a file containing static mapping entries
- Enable static mappings to be overwritten during replication

# Add a static mapping entry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To add a static mapping entry**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **New Static Mapping**.

   The **New Static Mapping** dialog box appears.

4. Enter the following information to complete the static mapping:

   - In **Computer name**, type the NetBIOS name of the computer.
   - In **NetBIOS scope (optional)**, you can type a NetBIOS scope identifier, if one is used, for the computer. Otherwise, leave this field blank.
   - In **Type**, click one of the supported types to indicate whether this entry is a **Unique**, **Group**, **Domain Name**, **Internet**, or **Multihomed** type entry.
   - In **IP address**, type the address for the computer.

5. Click **Apply** to add the static mapping entry to the database.
6. To add more static mapping entries, repeat steps 4 and 5. When you are done adding static mapping entries, click **OK**.

### Important

- Because each static mapping is added to the database when you click **Apply**, you cannot cancel work in the dialog box. If you make a mistake when typing a name or address for a mapping, delete the incorrect mapping, and add a new one using the correct information.

### Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- In general, only use static WINS mappings for computer names (not group names) that do not directly use WINS to register their names. Static mappings can overwrite dynamic mappings at other WINS servers and can be increasingly difficult to remove from all servers once they are introduced into a replicated WINS environment.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using static mappings
Delete or tombstone an entry in the WINS database
Deleting and tombstoning records
Blocking replication partners

# Edit a static mapping entry

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To edit a static mapping entry**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. In the details pane, double-click the static mapping you want to edit.
4. Edit the following information for the existing static mapping, as needed:

   - For **IP Address**, you can change the address or, for supported types, you can add or remove addresses from the address list and change the order of addresses.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Only the IP address for an entry can be edited. To change other properties for an entry, you must delete the existing entry and add a new entry.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a static mapping entry
Using static mappings

# Import a file containing static mapping entries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To import a file containing static mapping entries**

1. Open WINS.
2. In the console tree, click **Active Registrations**.

   **Where?**

   - WINS/*applicable WINS server*/Active Registrations

3. On the **Action** menu, click **Import LMHOSTS File**.
4. In the **Open** dialog box, locate and click the Lmhosts-formatted file that you want to use to import static mappings, and then click **Open**.

   WINS immediately begins importing the file into the server database. Wait for WINS to finish this before you click **OK**.

## Notes

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.
- Typically, the Lmhosts file is stored in the *systemroot*\System32\Drivers\Etc folder.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using static mappings

# Enable static mappings to be overwritten during replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable static mappings to be overwritten during replication**

1. Open WINS.
2. In the console tree, click **Replication Partners**.

   **Where?**

   - WINS/*applicable WINS server*/Replication Partners

3. On the **Action** menu, click **Properties**.
4. If needed, click the **General** tab.
5. Select the **Overwrite unique static mappings at this server (migrate on)** check box, and then click **Apply**.

### Note

- To open WINS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **WINS**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using static mappings

# Use WINS Command-line Tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Use WINS command-line tools**

- Use WINS commands interactively at the command prompt
- Use WINS commands non-interactively or in scripts

# Use WINS commands interactively at the command prompt

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To use WINS commands interactively at the command prompt**

1. Open **Command Prompt**.

2. Type **netsh**.

3. At the **netsh>** command prompt, type **wins**.

4. At the **wins>** command prompt, type:

   **server ** *ServerName*

   Or, type:

   **server** *ip_address*

   Where *ServerName* and *ip_address* are the server that you want to manage.

   Once a server that you have rights to manage has been successfully contacted, you should see "You have Read and Write access to the server *ServerName*."

5. Once you are connected to this WINS server, you can use any supported command.

   Type **/?** or **help** to display the immediate WINS subcommand menu. Or, type **list** to display all of the Netsh subcommands that are available for use with WINS.

## Notes

- You must be a member of the Administrators group to change configuration information on WINS servers using the WINS console or the **netsh wins** commands. Administrator group membership should be restricted to the minimum number of users necessary to administer the server.

- If there are users who need read-only access to the WINS console, add them to the WINS Users group instead of the Administrators group. WINS Users can search for WINS records and view replication partners and other configuration information, but they cannot change settings on the WINS server.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing WINS from the command line
Use WINS commands non-interactively or in scripts
Configure the command prompt

# Use WINS commands non-interactively or in scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To use WINS commands non-interactively or in scripts**

1. Open **Command Prompt**.

2. At the command prompt or in scripted form, you can use Netsh commands for WINS non-interactively by using the following fully qualified command syntax:

   **netsh wins server** [****_ServerName_] _subcommand_

## Notes

- You must be a member of the Administrators group to change configuration information on WINS servers using the WINS console or the **netsh wins** commands. Administrator group membership should be restricted to the minimum number of users necessary to administer the server.

- If there are users who need read-only access to the WINS console, add them to the WINS Users group instead of the Administrators group. WINS Users can search for WINS records and view replication partners and other configuration information, but they cannot change settings on the WINS server.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The ****_ServerName_ parameter is optional when performing WINS commands locally on the same computer where the command is executed. If the WINS server being managed is on a remote computer, this parameter is required and the remote server name must be used.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing WINS from the command line
Use WINS commands interactively at the command prompt
Configure the command prompt

# WINS Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Concepts**
This section provides general background information about Windows Internet Name Service (WINS) and NetBIOS over TCP/IP (NetBT) support provided by Microsoft Windows operating systems.

- WINS Overview
- Understanding WINS
- Administering WINS
- WINS Resources

# WINS Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS overview**


- Security information for WINS
- WINS defined
- Benefits of using WINS
- How WINS works
- WINS tools
- New features for WINS
- Planning WINS networks

# Security information for WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Security information for WINS**

It is important to follow best practices for security when using WINS servers on your network. For more information, see Best practices for security.

The following are known security issues for WINS and related protocols:

- **The NetBIOS naming protocol is an unauthenticated protocol.**

  Although the NetBIOS naming protocol can be used with network protocols other than TCP/IP, WINS was designed specifically to support NetBIOS over TCP/IP (NetBT).

  When a user connects to the network, the user is not required to provide credentials before using NetBT or requesting a name service from WINS servers. Malicious users with physical access to the WINS-enabled network can instigate a denial-of-service attack on WINS servers by requesting many name resolution requests or name registration requests from the server, thereby depleting the ability of the WINS server to service requests from other, non-malicious clients.

  **Recommendations:**

  - Ensure that unauthorized persons do not have physical or wireless access to your network.
  - Enable WINS logging in the event log. When **Log detailed events to Windows event log** is selected, WINS logs additional detailed information to the system event log, which allows you to monitor events on the WINS server. For more information, see Modify WINS logging properties and Event Viewer overview.

- Monitor the system event log for burst handling events. Burst mode is enabled by default. If you have disabled burst handling, enable it. The following event messages are logged for name refresh requests by WINS clients:

  - WINS 4338: The WINS server started the burst handling of incoming requests. WINS does this to handle a sudden increase of incoming requests. WINS will also log an event indicating when the burst handling stops. If you see this event frequently, you may want to upgrade the WINS server to a faster computer. You can also try to adjust the burst handling count using the WINS Manager tool.
  - WINS 4339: The WINS server completed the burst handling of incoming requests because the queue length became a quarter of what it was when burst handling began.

    You can also use Event Viewer when you suspect that server performance might be degraded due to an attack. For more information, see Modify burst handling properties.

  - Use Network Monitor if you suspect an attack on your WINS server. You can use Network Monitor to detect networking problems by capturing and viewing the frames (or packets) that a computer receives or sends on a local area network (LAN). When you view incoming packets, you can discover essential information about an attacker, including the IP address of their computer. For more information, see Network Monitor.

- **WINS database and log files are not protected if you move them from the default location.**

  The default location of the WINS database and log files is the *systemroot*\System32\Wins folder. By default, the access control list on the Wins folder provides the best security for the WINS database and log files. The access control list is a list of users and groups that can access the folder. In addition, each object (user or group) is assigned specific permissions that dictate what actions the object can perform. For more information, see Access Control.

  If you change the location of the WINS database files, the new folder location is not protected by the proper access

control list by default, thereby creating a security risk. The same is true when you configure a WINS database backup path. The folder you create when you designate the backup path is not protected by the proper access control list by default, and a security risk is created.

**Recommendations:**

- If you move the WINS files to another folder or create a new folder for backup files, create an access control list for the new folders that contains only the Local System account and the Administrators group. In addition, ensure that the new folder does not inherit unwanted permissions from any folder above it in Windows Explorer.
- Do not store WINS files in a personal folders (My Documents or My Pictures) or in a subfolder of a personal folder.
- Store WINS database and backup files on a computer that is physically secured and protected from unauthorized access.
- Turn on file auditing for individual WINS files, to track the users and groups that access the files. For more information, see Apply or modify auditing policy settings for a local file or folder.
- Do not move the WINS files to a remote storage device or map the WINS files to an external drive or storage device for which no access control list can be created. Also, do not move or copy WINS files to an anonymous File Transfer Protocol (FTP) site or any other unprotected location.

For more information, see Set the WINS database path, Set the WINS database backup path, Modify WINS logging properties, and The WINS database.

- **Name records for mission-critical servers can be overwritten, causing network traffic to be redirected.**

  If a computer connected to the network has the same NetBIOS name as a mission-critical server, such as a domain controller, Web server, or email server, the name records that map the mission-critical server name to a different and incorrect IP address overwrite the WINS name records of the mission-critical server. This can occur when:

  - The original owner of the record is unresponsive, such as when the server is rebooting, the server is disconnected from the network, or when a denial of service attack is being performed against the server.
  - The original record was registered dynamically.

  When this occurs, the mission-critical server name maps to a different IP address, and all of the traffic that was intended for that server is redirected to the new computer, which might be operated by a malicious user.

  **Recommendation:**

  - Use static WINS entries for mission-critical servers to prevent other computers from taking over the use of the name of a critical server. Static WINS entries also prevent other computers from appending additional IP addresses to the name record of a mission-critical server or domain. For more information, see WINS Best Practices, Using static mappings, and Add a static mapping entry.

**Additional recommendations**

Before you install and configure WINS for your network, consider:

- **Restricting who can enable, configure, and disable WINS.**

  You must be a member of the Administrators group to change configuration information on WINS servers using the WINS console or the **netsh wins** commands. Restrict the membership of the Administrator group to the minimum number of users necessary to administer the server.

  If there are users who need read-only access to the WINS console, add them to the WINS Users group instead of to the

Administrators group. WINS Users can search for WINS records and view replication partners and other configuration information, but they cannot change settings on the WINS server. For more information, see WINS tools, Finding and viewing WINS records, and View WINS Records.

# WINS defined

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS defined**

Windows Internet Name Service (WINS) provides a distributed database for registering and querying dynamic mappings of NetBIOS names for computers and groups used on your network. WINS maps NetBIOS names to IP addresses and was designed to solve the problems arising from NetBIOS name resolution in routed environments. WINS is the best choice for NetBIOS name resolution in routed networks that use NetBIOS over TCP/IP.

NetBIOS names are used by earlier versions of Microsoft® Windows® operating systems to identify and locate computers and other shared or grouped resources needed to register or resolve names for use on the network.

NetBIOS names are a requirement for establishing networking services in earlier versions of Microsoft operating systems. Although the NetBIOS naming protocol can be used with network protocols other than TCP/IP, WINS was designed specifically to support NetBIOS over TCP/IP (NetBT).

WINS simplifies managing the NetBIOS namespace in TCP/IP-based networks. The following illustration shows a typical series of events involving WINS clients and servers.




In this example, the following occurs:

1. A WINS client, HOST-A, registers any of its local NetBIOS names with WINS-A, its configured WINS server.

2. Another WINS client, HOST-B, queries WINS-A to locate the IP address for HOST-A on the network.

3. WINS-A replies with the IP address for HOST-A, 192.168.1.20.

WINS reduces the use of local IP broadcasts for NetBIOS name resolution and enables users to locate systems on remote networks easily. Because WINS registrations are done automatically each time clients start and join the network, the WINS database is automatically updated when dynamic address configuration changes are made. For example, when a DHCP server issues a new or changed IP address to a WINS-enabled client computer, WINS information for the client is updated. This requires no manual changes to be made by either a user or network administrator.

**Notes**

- The WINS protocol is based on and is compatible with the protocols defined for NetBIOS name service specified in RFCs 1001 and 1002, so it is interoperable with any other implementations of these RFCs.

- Replication of NetBIOS names data in WINS is Microsoft proprietary technology and is not interoperable with other NetBIOS name servers.

# Benefits of using WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Benefits of using WINS**

WINS provides the following benefits for administering your TCP/IP-based network:

- A dynamic name-to-address database that maintains support for computer name registration and resolution.
- Centralized management of the name-to-address database, alleviating the need for managing Lmhosts files.
- Reduction of NetBIOS-based broadcast traffic on subnets, by permitting clients to query WINS servers to directly locate remote systems.
- Support for earlier Microsoft® Windows® and NetBIOS-based clients on your network, permitting these types of clients to browse lists for remote Windows domains without requiring a local domain controller to be present on each subnet.
- Support for DNS-based clients by enabling those clients to locate NetBIOS resources when WINS lookup integration is implemented.

# How WINS works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**How WINS works**

By default, when a computer running Microsoft® Windows® 2000, Windows XP, or a Windows Server 2003 operating system is configured with WINS server addresses (either manually or through DHCP) for its name resolution, it uses hybrid node (h-node) as its node type for NetBIOS name registration unless another NetBIOS node type is configured. For NetBIOS name query and resolution, it also uses h-node behavior, but with a few differences.

For NetBIOS name resolution, a WINS client typically performs the following general sequence of steps to resolve a name:

1. Client checks to see if the name queried is its local NetBIOS computer name, which it owns.

2. Client checks its local NetBIOS name cache of remote names. Any name resolved for a remote client is placed in this cache where it remains for 10 minutes.

3. Client forwards the NetBIOS query to its configured primary WINS server. If the primary WINS server fails to answer the query--either because it is not available or because it does not have an entry for the name--the client will try to contact other configured WINS servers in the order they are listed and configured for its use.

4. Client broadcasts the NetBIOS query to the local subnet.

5. Client checks the Lmhosts file for a match to the query, if it is configured to use the Lmhosts file.

6. Client tries the Hosts file and then a DNS server, if it is configured for one.

**Notes**

- For a more detailed description of the default NetBIOS name resolution process as it is performed by computers running Windows XP or a Windows Server 2003 operating system, see Resolving names.

- In general, Windows operating systems support two major methods of performing network name resolution. These are:

    - **Host name resolution.** This is Windows sockets-based name resolution which implements the gethostbyname () API to search for a host IP address based on a queried host name. This method relies on either a Hosts file or querying DNS to perform resolution of the name.

    - **NetBIOS name resolution.** This is name resolution that uses the NetBIOS redirector to search for an address based on a queried NetBIOS name. This method relies on either an Lmhosts file or querying WINS to perform resolution of the name.

  WINS clients running Windows 2000, Windows XP, or a Windows Server 2003 operating system are configured by default to first use DNS to resolve names that are longer than 15 characters or that make use of periods (".") within the name. For names that are less than 15 characters and do not use periods, DNS is also used again as a final option after a WINS query fails if the client is configured to use a DNS server.

# WINS tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS tools**
The primary tools that you use to manage WINS servers are the WINS console and the Netsh WINS commands.

**The Netsh WINS commands**
You can type the Netsh WINS commands in a command prompt window at the **netsh wins>** command prompt, or you can execute the Netsh WINS commands in batch files and other scripts. All of the functionality that is available in the WINS console is available at the **netsh wins>** command prompt.

For more information, see Use WINS Command-line Tools. For a simple list of Netsh WINS commands, see Managing WINS from the command line.

For a comprehensive Netsh WINS command reference (including syntax, parameters and examples), see Netsh commands for WINS.

For an example of how to use Netsh WINS commands in batch files and scripts, see Netsh WINS example.

**The WINS console**
The WINS console is added to the Administrative Tools folder in Control Panel when you install a WINS server running a Microsoft® Windows Server 2003 server operating system. The WINS console appears as a Microsoft Management Console (MMC) snap-in, to further integrate WINS administration into your total network management.

After you install a WINS server, you can use the WINS console or the Netsh WINS commands to perform these basic administrative server tasks:

1. View and filter NetBIOS name registrations stored at the WINS server for client names used on your network.
2. Add and configure replication partners for the WINS server.
3. Perform maintenance-related tasks, including back up, restoration, compaction, and scavenging, on the WINS server database.

In addition, you can also use the WINS console to perform the following optional or advanced setup tasks:

- View and modify WINS properties, such as the **Renew Interval** and other intervals that are used when registering, renewing, and verifying name records stored in the server database.
- Add and configure static WINS mappings, if they are needed for use on your network.
- Delete or tombstone WINS records that appear in WINS server data used throughout your network.

## Notes

- The Windows Server 2003 family includes new features beyond those provided in earlier versions. For more information, see New features for WINS.
- The WINS console provides new ways to perform familiar WINS administrative tasks previously performed in Windows NT® 4.0 using WINS Manager. For more information, see New ways to do WINS tasks.

# New features for WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**New features for WINS**

Microsoft® Windows Server 2003 Windows Internet Name Service (WINS) provides the following new features:

- **Advanced WINS database search and filtering capabilities**

  Improved filtering and new search functions help you locate records by showing only those records that fit the criteria you specify. These functions are particularly useful in analyzing very large WINS databases. You can use multiple criteria to perform advanced searches for WINS database records. This improved filtering capability allows you to combine filters for customized and precise query results. Available filters include: record owner, record type, NetBIOS name, and IP address with or without subnet mask.

  Because you can now store query results in the cache of the random access memory (RAM) on your local computer, the performance of subsequent queries is increased, and network traffic is reduced. For more information, see Finding and viewing WINS records and View WINS Records.

- **Accepting replication partners**

  When determining a replication strategy for your organization, you can define a list that controls the source of incoming name records during pull replication between Windows Internet Name Service (WINS) servers.

  In addition to blocking name records from specific replication partners, you can also choose to accept only name records owned by specific WINS servers during replication, excluding the name records of all servers that are not on the list. For more information, see Accepting replication partners and Accept replication of records owned by a specified server.

The following features were new for Windows 2000 Server and are also provided with the Windows Server 2003 family:

- **Command-line administration of WINS servers using Netsh**

  In addition to a full graphical user interface for managing WINS servers, the Windows Server 2003 family provides a fully equivalent WINS command-line based tool called Netsh. For more information, see Managing WINS from the command line. For a comprehensive netsh WINS command reference (including syntax, examples, and parameters) see Netsh commands for WINS. For an example of how to use Netsh WINS commands in batch files and scripts, see Netsh WINS example.

- **Persistent connections**

  You can now configure each WINS server to maintain a persistent connection with one or more replication partners. This increases the speed of replication and eliminates the overhead of opening and terminating connections. For more information, see Persistent connections.

- **Manual tombstoning**

  You can manually mark a record for eventual deletion (*tombstoning*). The tombstone state of the record replicates to other WINS servers, preventing any replicated copies of the deleted records from reappearing at the same server where they were originally deleted. For more information, see Deleting and tombstoning records.

- **Improved management utility**

  The WINS console is fully integrated with the Microsoft Management Console (MMC), a powerful and user-friendly environment you can customize for your efficiency. Because all server administrative utilities included for use in the Windows Server 2003 family are part of MMC, new MMC-based utilities are easier to use as they operate more

predictably and follow a common design.

- **Improved ease of use for both using and managing advanced WINS features**

  WINS now includes the ability to block records by a specific owner or WINS replication partner (formerly known as Persona Non Grata), and to use the feature **Overwrite unique static mappings at this server** (formerly known as **migrate on**). For more information, see either Blocking replication partners or Using static mappings.

- **Dynamic record deletion and multi-select**

  These features help you more easily manage the WINS database. With the WINS snap-in, you can simply point, click, and delete either one or more WINS static or dynamic-type entries. Previously, this functionality was only available when using command-based utilities, such as Winscl.exe, designed for earlier versions of WINS. It is also possible now to delete records that use names based on non-alphanumeric characters.

- **Record verification and version number validation**

  These features quickly check the consistency of names stored and replicated on your WINS servers. Record verification compares the IP addresses returned by a NetBIOS name query of different WINS servers. Version number validation examines the owner address-to-version number mapping tables. For more information, see Checking the WINS database for consistency.

- **Export function**

  When you *export*, you place WINS data in a comma-delimited text file. You can export this file to Microsoft Excel, reporting tools, scripting programs, or similar programs for analysis and reporting. For more information, see Export WINS console data to a text file.

- **Increased fault tolerance for clients**

  Clients running Windows 98, Windows 2000, Windows Millennium Edition, Windows XP, or a Windows Server 2003 operating system can specify more than two WINS servers (up to a maximum of 12 addresses) per interface, using the first two servers in the list for NetBIOS name registration. The extra WINS server addresses are used only if the primary and secondary WINS servers fail to respond. WINS clients also benefit from the use of burst handling, which is enabled by default. For more information, see Burst handling and "Primary/Secondary WINS servers" in WINS servers.

- **Dynamic re-insertion of client names**

  WINS clients no longer have to be restarted after they use WINS to force re-insertion and update of local NetBIOS names. The Nbstat command includes a new option, **-RR**, which provides this functionality. If WINS clients are running Windows NT® 4.0, updated to Service Pack 4 or later, the **-RR** option can also be used with those clients. For more information, see Verify WINS registration of client NetBIOS names.

- **Read-only console access to the WINS console.**

  This feature provides a special purpose local users' group, the WINS Users group, which is automatically added when WINS is installed. By adding members to this group, you can provide read-only access in the WINS console to WINS-related information on this server computer for users that are not administrators. You can allow a user in this group to view, but not modify, information and properties stored at a specific WINS server. For more information, see Default groups.

- **Improved database engine.**

  WINS uses the same performance-enhanced database engine technology that is also used in Active Directory®.

# Planning WINS networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Planning WINS networks**

Before you install WINS servers on your network, plan the following tasks:

1. *Determine the number of WINS servers that you need.*

   One WINS server can handle NetBIOS name resolution requests for a large number of computers. However, consider the location of routers on your network and the distribution of clients in each subnet when you decide how many WINS servers are actually required.

2. *Plan the replication partnerships.*

   Determine whether to configure WINS servers as either pull or push partners, and set partner preferences for each server. For more information, see Configuring WINS replication.

3. *Assess the impact of WINS communication traffic on slower links.*

   Although WINS is designed to help reduce broadcast traffic between local subnets, it does create some traffic between servers and clients. This can be particularly important if you use WINS on routed TCP/IP networks.

   Also consider the effects of slower speed links (such as those typically used for wide area networks) on both replication traffic between WINS servers and NetBIOS registration and renewal traffic required for WINS clients.

4. *Assess the level of fault tolerance within your network for WINS.*

   Consider how shutting down a WINS server (even temporarily) affects your network. Use additional WINS servers for failure recovery, backup, and redundancy.

5. *Test and revise your planned WINS installation.*

   By testing the performance of your network installation of WINS, you can better identify potential problems before they occur. Use WINS server performance counters, which are available through the use of System Monitor.

**Planning how many WINS servers to use**

On a smaller network, a single WINS server can adequately service up to 10,000 clients for NetBIOS name resolution requests. To provide additional fault tolerance, you can configure a second computer running a Microsoft® Windows Server 2003 server operating system as a secondary (or backup) WINS server for clients. If you use only two WINS servers, you can easily set them up as replication partners of each other. For simple replication between two servers, one server should be set as a pull partner and the other as a push partner. Replication can be either manual or automatic, which you can configure by selecting the **Enable automatic partner configuration** check box on the **Advanced** tab of the **Replication Partner Properties** dialog box.

A larger network sometimes requires more WINS servers for several reasons including, most importantly, the number of client connections per server. The number of users that each WINS server can support varies with usage patterns, data storage, and the processing capabilities of the WINS server computer. Some enterprise network environments require more robust hardware to handle WINS activity, so you might benefit from upgrading the server computer.

When planning your servers, remember that each WINS server can simultaneously handle hundreds of registrations and queries per second. The WINS development and testing team compiled the following statistics as a profile of WINS performance during preliminary testing of Windows Server 2003 operating systems:

| | |
|---|---|
| Average registrations/second | 300 |

| | |
|---|---|
| Average queries/second | 350 |

These measurements were based on a server computer running WINS with no other additional services running and installed on currently available hardware. The specific hardware used includes the following: a single Intel Pentium II 350 MHz processor, 128 MB of RAM, and a standard IDE hard drive, which was used to store the WINS database. As a point of reference, WINS records data exchanged between servers and clients is typically small. On average, WINS records are about 40 bytes in size.

A conservative recommendation for large networks is to install one WINS server and a backup server for every 10,000 computers on the network. This reflects the previously suggested baseline performance and provides for large-scale power outages that could cause many WINS computers to restart simultaneously. Capacity is dependent upon server and network resources. Processor speed, available RAM, database storage space, and network resources can have a significant impact.

Two factors can enhance WINS server performance:

- Using two installed processors on the computer running WINS. This can increase performance by almost 25 percent.
- Using a dedicated high-performance disk drive, separate from the system drive, for the WINS database.

### Notes

- After you establish WINS servers on your intranet, you can adjust the *renew interval*, which is the time between a WINS client name registration and name renewal. By using this setting to reduce the number of registrations, you can tune server response time.

  For more information, see Modify name record settings and Modifying WINS server defaults.

- The default WINS database location is *systemroot*\System32\Wins. If you choose to move the WINS database to another location, you can change the database path using the WINS console.

  For more information, see Set the WINS database path.

**Planning for fault tolerance**

To ensure you are planning a fault-tolerant WINS design, ask the following question for each server on your network: *What happens to WINS if this server shuts down or if clients cannot reach it?*

To help answer the question, consider the two common scenarios in which a WINS server could fail to perform in its role on a network:

1. Hardware or power failures that require downtime for server repair or maintenance.
2. A network link or router failure that isolates a WINS server from clients.

To plan for these, determine the maximum length of time any given WINS server could be out of service on your network, factoring in the length of both planned and unplanned outages.

Also consider what happens to your WINS clients if their primary WINS server is shut down. By maintaining and assigning secondary WINS servers for clients, you can reduce the effects of a single WINS server being offline.

**Planning for network performance**

Begin planning by estimating the amount of network traffic between WINS clients and WINS servers under normal conditions. You should predict and monitor the following:

- NetBIOS names commonly registered by WINS clients.
- WINS registration and renewal caused by daily startup of clients.
- Roving users and their effect when moving within a routed network.

- WAN links and their effect on replication performance and convergence.

When planning for WINS network performance, it is important to understand the types of traffic that WINS clients typically generate on a network, as listed below.

- **Name registration**

  For more information, see Registering names.

- **Name refresh**

  For more information, see Renewing names.

- **Name release**

  For more information, see Releasing names.

- **Name query**

  For more information, see Resolving names.

**Daily start up of WINS clients**
When a WINS-enabled client starts on the network, it sends a name registration request for the computer name, user name, domain name, and any additional Windows NT® network client services running on the computer.

A WINS client running Windows XP usually registers more NetBIOS names than other WINS-enabled clients. The name registration requests generated by a computer running under these versions of Windows can include the following: File Server, Domain, the Replicator service, the Messenger service, the Computer Browser service, and additional network service names.

An active WINS client name registration in a WINS server database replicates to all pull partners configured on that WINS server. Eventually, the active name registration replicates to all WINS servers on the network.

When a WINS client is shut down at the end of a workday, it releases the name. When the computer is restarted at the beginning of next workday, the WINS client registers the name again with the WINS server and receives a new version ID. This new, active name registration entry replicates to the pull partners of the WINS server.

Assuming that this cycle occurs each day for each computer, the number of name registration entries replicated each day is approximately the number of computers that start each day multiplied by the number of NetBIOS names registered at each computer.

On large networks (50,000 or more computers), the biggest traffic load represents the name registration requests generated when WINS clients start on the network.

**Planning considerations for routed WINS networks**
When planning for WINS client traffic on large routed networks, consider the effect of name query, registration, and response traffic routed between subnets. Name requests and responses that occur at the daily startup of computers must pass through the traffic queues on the routers, and can cause delays at peak times.

Roving or mobile users also present challenges for planning WINS in routed networks. When a user stops a computer and then moves it to another location--starting the computer on a different subnet with a different primary WINS server--the resulting changes can generate name-challenge traffic for WINS.

Typically, the name registration request is answered with a brief "Wait for Acknowledgment" message. Assuming the name being registered was previously replicated to the new server, it challenges the IP address currently in its database for this name with a brief query. When, as expected, there is no reply, the WINS server retries the challenge two more times for a total of three attempts. If all three attempts fail, the challenge is successful and the new WINS server updates the name registration entry in its database with the new IP address, incrementing the version ID. The new version ID indicates that the entry must be replicated from its new owner--the new WINS server--to other WINS servers on the network.

You can estimate WINS client traffic based on the behavior of the WINS clients, as described in the preceding sections. However, when estimating WINS client traffic, also consider the network topology and the design or configuration of the

network routers. In some cases, it is not possible to predict the traffic load on a specific network router because the routers are configured to autonomously route traffic based on factors other than traffic load.

# Understanding WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Understanding WINS**


- WINS Components
- WINS Replication
- WINS Resolution
- WINS Features

# WINS Components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS components**

- WINS servers
- WINS clients
- WINS proxies
- The WINS database

# WINS servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS servers**

WINS consists of two main components, the WINS server and WINS clients.

The WINS server handles name registration requests from WINS clients, register their names and IP addresses, and responds to NetBIOS name queries submitted by clients, returning the IP address of a queried name if it is listed in the server database.

Also, as the following graphic shows, WINS servers can replicate the contents of their databases (which contain NetBIOS computer name mappings to IP addresses) to other WINS servers. When a WINS-enabled client computer (such as a workstation computer on either Subnet 1 or Subnet 2) starts on the network, its computer name and IP address are sent in a registration request directly to its configured primary WINS server, WINS-A. Because WINS-A is the server that registers these clients, it is said to be the *owner* for the records of the clients in WINS.




In this example, server WINS-A has clients registering with it that are both local (that is, clients on Subnet 2 where it resides) and remote (clients located across a router on Subnet 1). A second WINS server, WINS-B, is located on Subnet 3 and only owns mappings for local clients registering from its same subnet. WINS-A and WINS-B can later complete replication of their databases so that records for clients on all three subnets are in the WINS database at both servers. For more information, see WINS Replication.

**Primary/Secondary WINS servers**

WINS servers are used by clients in one of two ways: either as a primary or secondary WINS server.

The difference between primary and secondary WINS servers is not based in any way on the servers (that for all functional purposes are the same in WINS). The difference occurs at the client which differentiates and orders the list of WINS servers when provided more than one WINS server to use.

For most cases, the client contacts the primary WINS server for all of its NetBIOS name service functions (name registration, name renewal, name release, and name query and resolution). The only case where secondary WINS servers are ever used is when the primary WINS server is either:

1. Unavailable on the network when the service request is made, or

2. Unable to resolve a name for the client (in the case of a name query).

In the case of a failure by the primary WINS server, the client requests the same service function from its secondary WINS servers. If more than two WINS servers are configured at the client, the additional WINS servers are tried until the list is exhausted or one of the secondary WINS server succeeds in processing and responding to the request. After a secondary WINS server is used, a client periodically tries to switch back to its primary WINS server for future service requests.

For most recent WINS clients (Windows XP and Windows 2000), a list of up to 12 secondary WINS servers can be configured (either manually through TCP/IP properties or dynamically by a DHCP server providing a list using DHCP option type 44). This feature is useful in an environment where there is a large number of mobile clients and NetBIOS-based resources, and services are used often. Because in these types of environments, the WINS database may not be consistent throughout the network of

WINS servers because of convergence issues, it can be helpful for clients to be able to query more than two WINS servers.

However, this option should not be overused needlessly as it produces a tradeoff in terms of the actual benefit for adding fault tolerance by listing additional WINS servers. The benefit of this feature must be weighed against the fact that for each additional WINS server listed, the time to process a query request fully in WINS takes incrementally longer to achieve. For example, if a WINS client tries three or more WINS servers before failing, it can substantially delay the processing of the name query before alternative methods for resolution are tried, such as searching of a local Hosts file or querying a DNS server.

# WINS clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS clients**

WINS clients attempt to register their names with a WINS server when they start or join the network. Clients thereafter query the WINS server to resolve remote names as needed.

WINS-enabled clients are computers that are able to be configured to make direct use of a WINS server. Most WINS clients typically have more than one NetBIOS name that they must register for use with the network. These names are used to publish various types of network services, such as the Messenger or Workstation service, that each computer can use in various ways to communicate with other computers on the network.

Microsoft supports WINS clients running on the following platforms:

- Windows Server 2003 operating systems
- Windows XP operating systems
- Windows Millennium Edition
- Windows 2000 operating systems
- Windows NT operating systems
- Windows 98
- Windows 95
- Windows for Workgroups
- Microsoft LAN Manager
- MS-DOS clients
- OS/2 clients
- Linux and UNIX clients (with Samba installed)

WINS-enabled clients communicate with the WINS server to:

- Register client names in the WINS database.
- Renew client names with the WINS database.
- Release client names from the WINS database.
- Resolve names by obtaining mappings from the WINS database for user names, NetBIOS names, DNS names, and IP addresses.

Clients that are not WINS-enabled can participate in these processes in a limited way by using WINS proxies. For more information, see WINS proxies.

# WINS proxies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS proxies**

A WINS proxy is a WINS client computer configured to act on behalf of other host computers that cannot directly use WINS. WINS proxies help resolve NetBIOS name queries for computers located on routed TCP/IP networks.

By default, most computers not able to use WINS use broadcasts to resolve NetBIOS name queries and register their NetBIOS names on the network. You can configure a WINS proxy to listen on behalf of these computers and to query WINS for names not resolved by broadcast.

WINS proxies are only useful or necessary on networks that include NetBIOS broadcast-only (or b-node) clients. For most networks, WINS-enabled clients are common and WINS proxies are typically not needed.

WINS proxies are WINS-enabled computers that listen for b-node NetBIOS name service functions (name registration, name release, and name query) and can respond for those names that are remote and not used on the local network. Proxies communicate directly with a WINS server to retrieve the information necessary to respond to these local broadcasts.

WINS proxies are used in the following ways:

- When a b-node client registers its name, the proxy checks the name against the WINS server database. If the name exists in the WINS database, the proxy might send a negative registration response back to the b-node client attempting to register the name.
- When a b-node client releases its name, the proxy deletes the client name from its remote name cache.
- When a b-node client sends a name query, the proxy attempts to resolve the name using information either locally contained in its cache of remote names or through the use of information it obtains from the WINS server.

**How WINS proxies resolve names**

The following graphic illustrates the use of a WINS proxy, HOST-B, for a subnet containing a b-node client, HOST-A.




In this example, the following steps are used by the WINS proxy to resolve names for the b-node computer:

1. HOST-A broadcasts a NetBIOS name query to the local subnet.
2. HOST-B accepts the broadcast and checks its cache for the appropriate NetBIOS computer name-to-IP address mapping.
3. HOST-B processes the request.

   If HOST-B has a name-to-IP address mapping cached that matches the one HOST-A requested, it returns this information to HOST-A. If not, HOST-B queries a WINS server for the mapping requested by HOST-A.

4. When HOST-B receives the requested name-to-IP address mapping from its configured WINS server (in this example, WINS-A), it briefly caches this information.

   By default, a WINS proxy caches the remote name mappings that it queries in WINS for 6 minutes, but this value is configurable with a minimum value of 1 minute.

5. HOST-B can then use this mapping information to respond to subsequent NetBIOS name queries broadcast from either HOST-A or other b-node computers on the subnet.

**Notes**

- Because WINS servers do not respond to broadcasts, you should configure a computer as a WINS proxy on each subnet that contains non-WINS computers that must use broadcasts to resolve NetBIOS names.

- When a WINS proxy is used to answer a query for a multihomed client or a group record containing a list of IP addresses, only the first listed address is returned to the b-node client.

- For more in-depth information on configuring or using WINS proxy, see Updated technical information or "Windows Internet Name Service" at the Microsoft Windows Resource Kits Web site.

# The WINS database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**The WINS database**

The WINS database stores and replicates the NetBIOS name-to-IP address mappings for your network. In the Windows Server 2003 family, the WINS database uses the Extensible Storage Engine (ESE).

**Compacting the database**

There is no built-in limit to the number of records that a WINS server can replicate or store. The size of the database depends on the number of WINS clients on the network. The WINS database changes over time as clients log on and log off the network.

However, the size of the WINS database is not directly proportional to the number of active client entries. Over time, as some WINS client entries become obsolete and are deleted, the size of the WINS database grows larger than the actual space currently in use by the database. This is because the space used to store obsolete records is not automatically reclaimed by the server once the space is freed and no longer in use.

Compacting the WINS database recovers the unused space. Dynamic database compaction occurs on WINS servers as an automatic background process during idle time after a database update. Compaction can also be done manually offline. Windows NT Server 4.0, Windows 2000, and the Windows Server 2003 family support both dynamic and manual compaction. Windows NT Server 3.51 (or earlier) supports only manual compacting of the WINS server database.

Although dynamic compacting greatly reduces the need for offline compaction, offline compaction reclaims the space better and should be done periodically. How frequently you manually compact the WINS database depends on your network. For large, busy networks with 1,000 or more WINS clients, you should compact offline each month. Smaller networks usually require less frequent manual compaction.

Because the dynamic database compaction occurs while the database is in use, you do not need to stop the WINS server during this process. However, for manual compacting, you must stop the WINS server and take it offline.

**Backing up the WINS database**

The WINS console provides the tools you need to maintain, view, back up, and restore the WINS server database. You should back up the database whenever you back up other files on the WINS server.

**WINS database files**

WINS uses the Jet database format for storing its data. Jet produces the J<*n*>.log and other files in the *systemroot*\System32\Wins folder to increase the speed and efficiency of data storage.

The following table discusses the files that are created and used by the Jet database in each WINS server.

| File | Description |
|---|---|

| | |
|---|---|
| J50.log and J50#####.log | A log of all transactions done with the WINS database. This file is used by WINS to recover data if necessary.<br><br>To increase speed and efficiency of data storage, the Jet database writes current transactions to log files rather than directly to the database. Therefore, the most current view of the data includes both the database and any transactions in the log files. Both of these files are used for recovery if the WINS service abruptly or unexpectedly stops. If the service stops in an unexpected manner, the log files are automatically used to re-create the correct state of the WINS database.<br><br>Log files maintain a specific size; however, they can grow quickly on a busy WINS server. Inevitably, WINS writes more transactions to a log than the log size can accommodate. When a log file is filled, it is renamed, indicating that it is an older log and not in use. A new transaction log is created with the J<*n*>.log file name, where <*n*> is a decimal number, such as J50.log. The naming format of the previous log file is Jet*XXXXX*.log, where each *X* denotes a hexadecimal number from 0 to F. Previous log files are maintained in the same folder as the current log files.<br><br>The log files are processed (all entries written to the database) and deleted every three hours. Processing and deletion also occur upon successful WINS database backup or when the WINS server is shut down properly. If many J<*n*>.log files accumulate, you should schedule frequent backups to maintain the logs.<br><br>After the entries are processed, you can manually delete the log files; however, this prevents a successful recovery of the database if it should be needed. For this important reason, do not manually delete or remove the log files from the system until a backup has been performed. |
| J50.chk | A checkpoint file that indicates the location of the last information successfully written from the transaction logs to the database. In a data recovery scenario, the checkpoint file indicates where the recovery or replaying of data should begin. This checkpoint file is updated every time data is written to the database file (Wins.mdb). |
| Wins.mdb | The WINS server database file, which contains two tables: the IP address-to-Owner ID mapping table and the Name-to-IP address mapping table. |
| Winstmp.mdb | A temporary file that is created by the WINS server service. This file functions as a swap file during index maintenance operations and can remain in the *systemroot*\System32\Wins directory after a system failure. |
| Res#.log | These are reserved log files, which function in emergencies when the server runs out of disk space. If a server attempts to create another transaction log file, and there is insufficient disk space, the server flushes any outstanding transactions into these reserved log files. The service then shuts down and logs an event to Event Viewer. |

**Important**

- The J50.log, J50#####.log, Wins.mdb, Winstmp.mdb, and Res#.log files should not be removed or altered.

# WINS Replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS replication**

- WINS replication overview
- Pull partners
- Push partners

# WINS replication overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Replication overview**

Where multiple WINS servers are used on your network, they can be configured to replicate records in their databases to other servers. This process is shown in the following figure. Two WINS servers, WINS-A and WINS-B, are both configured to fully replicate their records with each other.




By using replication between these WINS servers, a consistent set of WINS information is maintained and distributed throughout the network. For example, in the previous figure a WINS client (HOST-1) on Subnet 1 registers its name with its primary WINS server, WINS-A. Another WINS client (HOST-2) on Subnet 3 registers its name with its primary WINS server, WINS-B. If either of these hosts later attempts to locate the other host using WINS--for example, HOST-1 queries to find an IP address for HOST-2--replication between the WINS servers (WINS-A and WINS-B) makes resolving this query possible.

For replication to work, each WINS server must be configured with at least one other WINS server as its replication partner. This ensures that a name registered with one WINS server is eventually replicated to all other WINS servers in the network. A replication partner can be added and configured as either a pull partner, a push partner or to use both types of replication as a push/pull partner. The push/pull partner type is the default configuration and is the type recommended for use in most cases.

When WINS servers replicate, a latency period exists before a client name-to-address mapping from any given server is propagated to all other WINS servers in the network. This latency is known as the *convergence time* for the entire WINS system. For example, a name release request by a client will not propagate as fast as a name registration request. This is intended by design because it is common for client names to be released and then reused with the same mapping as computers are restarted or periodically shut down. Replicating each of these name releases would unnecessarily increase the network load of replication.

Also, when a WINS client computer is shut down improperly, as in the case of an unexpected power outage, the registered names of the computer are not released as they normally would be by sending a name release request to the WINS server. Therefore, the presence of a name-to-address record in the WINS database does not necessarily mean that a client computer is still using the name or its associated IP address. It only means that at some recent time in the past, a computer by the registered name claimed the mapped IP address for its use.

**Notes**

- The primary and secondary WINS servers of any client should typically have a push and pull relationship with each other. However, a direct push/pull configuration is not required as long as similar results can be indirectly achieved during a normal replication cycle.

- WINS replication is always incremental, meaning that only changes in the database are replicated each time replication occurs, not the entire database.

# Pull partners

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Pull partners**

A pull partner is a WINS server that pulls or requests replication of updated WINS database entries from other WINS servers (those configured to use it as a push partner) at a configured interval. This is done by requesting entries with a higher version ID than the last entry received from its configured partner.

You can enable WINS to pull (request) a replication with its configured partners when one of the following events occurs:

- When the WINS server starts.
- When a set amount of time has elapsed.

The amount of time between pull replications is indicated in the **Replication interval**, which can be specified either globally for all pull partners or individually for specific partners. When the time has elapsed for each interval, WINS sends a pull replication trigger to its pull partners. The partners can then respond with incremental changes, if any have been made since the last time replication occurred.

To configure which of these options WINS uses, you need to configure pull parameters, either the default parameters that the server uses for all pull partners or the parameters in effect for a specific partner. For more information, see Modify pull partner properties.

Additionally, you have the administrative option of manually sending a pull replication trigger to another pull partner of the selected WINS server. For more information, see Send a replication trigger to a selected partner.

**How WINS determines which entries to pull replicate**

Because push-based replication uses notification, it does not transfer any records until a subsequent pull trigger is returned by the notified partner. For this reason, pull replication can be viewed as the single process used to replicate WINS data between servers.

All WINS servers maintain an internal mapping table--an owner-version mapping table--which is used to complete the replication process. This table is built dynamically at startup by the server and is held in memory.

The owner-version mapping table contains the following:

1. An owner ID that maps to the IP address for each WINS server that appears in the database as an owner for records currently stored there.

   For practical purposes, WINS automatically assigns and increments an integer value for the owner ID and then maps that value to a corresponding server IP address. However, when you examine this information in the WINS console, only the owner IP address is actually displayed.

2. When replication occurs with another WINS server, the version ID field for all records currently in the database is scanned for the highest version ID associated with each identified distinct owner.

A sample Owner/Version ID table for WINS-A when it is started might look like the following, if viewed using **View Records** in the WINS console:

| Owner ID (not displayed) | Owner | Name | Highest ID |
|---|---|---|---|
| (0) | 192.168.1.5 | WINS-A | 1F |
| (1) | 192.168.2.5 | WINS-H | 125 |
| (2) | 192.168.3.5 | WINS-B | 31 |

| | | | |
|---|---|---|---|
| (3) | 192.168.4.5 | WINS-C | 100 |
| (4) | 192.168.5.5 | WINS-D | 200 |

The owner-version mappings are used by WINS to determine the state of the database at any point in time before or after replication. The following example provides more detailed information about how WINS manages replication.

**Example: How pull replication works**

The following graphic demonstrates how four WINS servers--WINS-A, WINS-B, WINS-C, and WINS-D--are all configured to be spokes (push/pull partners) of a central hub server, WINS-H.




Before replication with WINS-H, the owner-version mapping table at WINS-A contains a set of highest version IDs for all owners in its database. For the purpose of this example, assume these values are the same as those shown in the sample table provided in the previous section.

In reality, pull replication might be occurring dynamically between all four spokes at various times with WINS-H, so this table might not remain static for long. For a basic understanding of the details of an individual pull replication cycle, however, it is easier to focus on what might happen when just one of the spokes--for example, WINS-A--replicates its database with WINS-H.

In the graphic shown, the following changes in the replicated WINS network occur:

1. The three other spokes (WINS-B, WINS-C and WINS-D) each replicate with WINS-H, based on the highest version ID.

    a. WINS-B has added or updated one new record since its last replication cycle, which WINS-H replicates and adds to its database.

    b. WINS-C has the same version ID and does not have any new records to be replicated at WINS-H.

    c. WINS-D has updated a total of 20 records since its last replication cycle, which WINS-H replicates and adds to its database.

2. WINS-H sends a trigger to WINS-A initiating replication.

    The updated owner-version mapping table information is provided as part of the trigger.

3. WINS-A responds by pulling entries from WINS-H.

    WINS-A determines which owner servers to pull from based on the difference between its locally held highest version ID for each owner and the corresponding updated highest version IDs provided by WINS-H. WINS-A specifies the range of records to be pulled and replicated by comparing versions.

**Notes**

- In the case where another pull replication to the same partner is already queued, WINS sees the subsequent request as further down in the queue and discards it in favor of keeping the one with a higher placement in the service process queue.

- In addition to modifying settings for pull replication properties, WINS also provides the option to enable the use of persistent connections with either a specified pull partner or all pull partners of a WINS server.

  For more information, see either Persistent connections or Enable the use of persistent connections with a partner server.

# Push partners

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Push partners**

A push partner is a WINS server that pushes or notifies other WINS servers (those configured to use it as a pull partner) of the need to replicate their database entries at a configured interval. To configure a WINS server to use push replication, you can choose from several WINS console configurable options. For example, you can enable WINS to push notify configured partners when one of the following events occurs:

- When the WINS server starts.
- Whenever an IP address change occurs for a name-to-address mapping in the WINS server database.
- When a value for **Number of changes in version ID before replication** is specified--either globally for all push partners or individually for specific push partners--WINS can use this value as a threshold for determining how many incremental changes must be made to the database before it starts push replication to its partners.

To configure any of these options for WINS use, you need to configure push parameters, either the default parameters that the server uses for all push partners or the parameters in effect for a specific partner. For more information, see Modify push partner properties.

Additionally, you can manually start push replication to other WINS servers. For more information, see Send a replication trigger to a selected partner or Send a replication trigger to any specified WINS server.

**How updates and WINS settings affect push replication**

By default, WINS does not perform push replication unless other related settings are also configured on the **Push Replication** tab in **Replication Partners Properties**. For example, WINS uses an *update count* value of 0. This is set using **Number of changes in version ID before replication**.

Push replication only occurs at configured replication intervals, unless other push replication settings are configured as well, such as:

- For **Start push replication**, selected events such as the **At service startup** or the **When address changes** check boxes are selected.
- The **Use persistent connections for push replication partners** check box is selected, which can affect how the value of **Number of changes in version ID before replication** is handled by WINS.

Persistent connections are enabled by default, and replication with partners occurs dynamically as records are updated in the local WINS database. A persistent connection alleviates the need for connections to be opened and closed with each update of the WINS database between WINS server replication partners. Replication with partners occurs at a rate specified by the **Number of changes in version ID before replication** parameter. With this value set to the default of zero, replication takes place as soon as a change occurs. This default can be changed so that replication occurs at longer intervals.

With persistent connections enabled for push replication, the default value of 0 causes the local WINS server to send a push trigger and notify its push partners each time an update occurs. An update is defined as an incremental increase to the highest version ID in the local WINS database for records the server owns. This can occur when a new name record is registered and added locally or if an IP address change occurs for an existing record.

With persistent connections, you can reduce the update frequency for push notifications, by specifying a non-zero value instead. If a value greater than zero is used, WINS only starts push replication when the highest ID for records it owns has increased an equal number of times.

For example, the current highest version ID for an example server, WINS-A, is currently 1C4 (which is 452 in converted decimal notation). WINS-A also has its **Number of changes in version ID before replication** set to 2 for either all or some of its push partners.

Suppose that the following four distinct sets of database changes are then made to the WINS server database at WINS-A:

1. HOST-A refreshes its previously registered name record, though it does not change its IP address.

2. HOST-B registers a name record, creating a new dynamic name-to-address mapping to the database.

3. HOST-C registers a previously released name record and, in the process, reports an updated change of IP address for this record.

4. A pull replication occurs at a specified interval with another WINS server.

Once these changes are entered into the database, the newly revised highest version ID for WINS-A is 1C6 (454 decimal). This is because the changes made to WINS server data that were caused by items 2 and 3 (either a client registering a new mapping or updating the IP address used with a previous mapping) each caused the version ID at that server to increment.

For items 1 and 4 in the previous list of changes, the highest version ID at the WINS server does not increase.

After these changes occur, any or all partners of WINS-A that are configured with a **Number of changes in version ID before replication** of 2 would then be push notified by WINS sending a replication trigger to those servers.

**Example: How push with propagation works**
The following figure shows an example of the use of push partners in a replicated WINS network of four servers. In this example, two of the servers--WINS-A and WINS-B--are located at regional data center locations.




These servers are both central hub servers in a multiple hub-and-spoke enterprise WINS replication design. In order to ensure frequent convergence of names data between the separate regional hubs, both WINS-A and WINS-B maintain a full push/pull relationship with each other.

WINS-C and WINS-D are site or branch-level spoke servers used on the same replicated WINS network. Respectively, they are both configured as push-only partners of WINS-A and WINS-B.

If updates occur at one hub WINS-A, it is unlikely a remote spoke operating as a push-only partner (WINS-D) would receive those updates in a timely or predictable manner. Although replication between the hubs might occur frequently enough to pose no significant issues, setting the update count to an arbitrarily high value at WINS-B might delay convergence at WINS-D until more than one replication cycle has been completed.

To administratively correct this delay, you can initiate a push replication trigger at WINS-A to be sent to WINS-B, which would queue a replication immediately between the hubs. In doing so, if you also chose the **Propagate to all partners** when manually starting push replication, then WINS-B would propagate the push replication trigger to WINS-D. This would result in an early notification to WINS-D of the need to use a pull replication to WINS-B, obtaining any interim updates from WINS-A.

**Note**

- In addition to modifying settings for push replication properties, WINS also provides you the option to enable the use of persistent connections with either a specified push partner or all push partners of a WINS server.

  For more information, see either Persistent connections or Enable the use of persistent connections with a partner server.

# WINS Resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS resolution**


- Resolution overview
- Registering names
- Releasing names
- Renewing names
- Resolving names

# Resolution overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Resolution overview**

WINS is used to resolve NetBIOS names, but in order for name resolution to be effective, clients need to be able to dynamically add, remove, or update their names in WINS. The following sections provide a functional description of each of these processes, specifically how names are registered, renewed, released, and resolved for clients on a WINS-based network.

This process of WINS client/server communications is outlined in the following graphic:




In a WINS system, all names are registered with a WINS server. The names are stored in a database on the WINS server which answers requests for name-to-IP address resolution based on the entries in this database.

Redundancy and load balancing are maintained with more than one WINS server in the network. Periodically, servers replicate their database entries among one another in order to maintain a consistent view of the NetBIOS namespace.

Each name has an entry in the database. It is owned by the WINS server it registered with and is a replica on all other WINS servers. Each entry has a state associated with it--the entry can be in the active, released, or extinct (also known as tombstone) state. Entries are also assigned a version ID. This number is used in the replication process and is discussed in more detail in **WINS replication**. For more information, see Push partners.

WINS also allows the registration of static names. This enables the administrator to register names for servers running operating systems that are not capable of dynamic name registration. WINS distinguishes between dynamic and static entries. Static names are treated somewhat differently than dynamic names. For more information, see Using static mappings.

# Registering names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Registering names**

Name registration is a WINS client requesting the use of a NetBIOS name on the network. The request may be for a unique (exclusive) or a group (shared) name. NetBIOS applications can also register one or more names.

As illustrated in the following figure, a WINS client (HOST-C) sends a Name Registration Request directly to its configured WINS server, WINS-A.




WINS-A can accept or reject the name registration request by issuing either a positive or negative name registration response to HOST-C. The action taken by WINS-A depends on several factors:

- Whether the name already exists in the server database at WINS-A
- If a record of the name exists, the state in the server database of the record at WINS-A might make a difference. Also, if the recorded IP address for the name is the same or different from the IP address of the requesting client, HOST-C.
- Whether the request is for a unique or group name entry.

If the name does not exist in the database, it is accepted as a new registration and the following steps occur:

1. The name for HOST-C is entered with a new version ID, given a time stamp, and marked with the owner ID of the WINS server.

   The time stamp is calculated based on adding the value of the **Renew interval** (which by default, is 6 days) set on the WINS server to the current date and time of the server.

2. A positive registration response is sent back to HOST-C with a Time-to-Live (TTL) value equal to the time stamp recorded for the name at WINS-A.

**Name registered for same IP address**

If the name HOST-C is already entered in the database with the same IP address as that being requested, the action taken depends on the state and ownership of the existing name.

1. If the entry is marked as *active* and the entry is owned by the server (WINS-A), the server updates the time stamp for the record and returns a positive Response back to the client.

2. If the entry is marked as being either *released* or *tombstoned*, or if the entry is owned by another WINS server, the registration is treated as a new registration. Time stamp, version ID, and ownership are all updated and a positive response is returned.

**Name registered for different IP address**

In the case where the name exists in the WINS database but with a different IP address than that being requested, the WINS

server is expected to avoid duplicate names. If the database entry is in the released or tombstone state, the WINS server is free to assign that name.

If, however, the entry is in the active state, the node holding the name is challenged to determine if it still exists on the network. In this case, the WINS server (WINS-A) could perform a challenge of the name and do the following:

1. WINS-A sends a Wait for Acknowledgment (WACK) response to the requesting client (HOST-C), specifying a time in the TTL field that the client should be prepared to wait for a response.

2. WINS-A then issues a name query request to the node currently registered for this name in the server database.

3. If the node still exists, it sends a positive response back to WINS-A.

4. WINS-A, in turn, sends a negative name registration response back to the requesting client (HOST-C), rejecting the name registration.

5. If a positive response is not received back from the first challenge query made by WINS-A, two subsequent name query requests are made. After three attempts with no response, the challenge process is complete and a positive registration response is returned to the requesting client (HOST-C) and the name updated in the server database for the new client registration.

**Notes**

- Unlike WINS-enabled clients, which can directly contact a WINS server, non-WINS clients (such as NetBT b-node clients) must register and then continuously defend their names by sending and replying to broadcast name queries within the local network.

- NetBIOS names are registered with the Windows Internet Name Service (WINS) and are normally released when the computer is properly shutdown. If the computer was not properly shutdown, or the computer was not able to contact a WINS server during shutdown, the **nbtstat** command can be used to refresh in WINS the local names for this computer. This can be useful for mobile or portable computers that are moved between different locations on a network.

  For more information, see Verify WINS registration of client NetBIOS names.

# Releasing names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Releasing names**

Name release occurs when a WINS client computer finishes using a particular name, when a proper shutdown occurs. In releasing its name, a WINS client notifies its WINS server (or potentially other computers on the network) that it is no longer using its registered name.

As shown in the following figure, when a WINS-enabled client (HOST-C) releases its name, the following steps occur:




1. The computer named HOST-C either shuts down properly or a user enters the **nbtstat -RR** command, initiating a name release request to be sent to the WINS server, WINS-A.

2. WINS-A marks the related database entry for HOST-C as *released*.

   If the entry remains released for a period of time, WINS-A marks the entry as *tombstoned*, updates the version ID for the entry, and notifies other WINS servers of the change.

3. WINS-A returns a release confirmation message to the WINS client, HOST-C.

If a name entry is marked as released, the WINS server can immediately update or revise a marked name entry when a new registration request arrives from a WINS client with the same name but with a different IP address. This is possible because the WINS database shows that the WINS client at the old IP address is no longer using that name. For example, this can happen when a DHCP-enabled laptop changes subnets.

Name release is most often used to simplify WINS registration for clients that shut down and restart on the network. If a computer released its name during a normal shutdown, the WINS server does not challenge the name when the computer reconnects. If a proper shutdown did not occur, the name registration with a new IP address causes the WINS server to challenge the previous registration. When the challenge fails (because the client computer is no longer using the old IP address), the registration succeeds.

In some cases, a client is not able to release its name by contacting the WINS server, so broadcasts must be used to release a name. This can happen when a WINS-enabled client shuts down without receiving confirmation of its name release by the WINS server.

# Renewing names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Renewing names**

Periodic name renewal is required for WINS client computers to renew their NetBIOS name registrations with the WINS server. The WINS server treats name renewal requests similarly to new name registrations.

When a client computer first registers with a WINS server, the WINS server returns a message with a Time-To-Live (TTL) value that indicates when the client registration expires or needs to be renewed. If renewal does not occur by that time, the name registration expires on the WINS server and the name entry is eventually removed from the WINS database. However, static WINS name entries do not expire, and therefore, do not need to be renewed in the WINS server database. For more information, see Using static mappings.

The default **Renew interval** for entries in the WINS database is six days. Renewal occurs every three days for most WINS clients because WINS clients attempt to renew their registrations when 50 percent of the TTL value has elapsed.

A name must be refreshed before this interval ends or it will be released. Names are refreshed by sending a name refresh request to the WINS server, as shown in the following figure.




It is the responsibility of the client (HOST-C) to refresh its name before the **Renew interval** expires. If the WINS server, WINS-A, does not respond to the refresh request, the client (HOST-C) can increase the rate at which it attempts to refresh its name.

**Important**

- In most cases, the default value is the appropriate **Renew interval**. You should always set the same **Renew interval** for all server replication partners whenever multiple WINS servers are used.

- Incorrectly adjusting the **Renew interval** can adversely affect system and network performance.

# Resolving names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Resolving names**

Name resolution for WINS clients is an extension of the same name resolution process used by all Microsoft NetBIOS over TCP/IP (NetBT) clients to resolve NetBIOS name queries on your network. The actual method of name resolution is transparent to the user.

For Windows XP and Windows 2000, WINS clients use the following sequential flow of options for resolving a name once the query has been made by using **net use** or a similar NetBIOS-based application.

1. Determine if the name is longer than 15 characters or if it contains periods ("."). If so, query DNS for the name.
2. Determine if the name is stored in the remote name cache at the client.
3. Contact and try configured WINS servers to attempt to resolve the name using WINS.
4. Use local IP broadcasts to the subnet.
5. Check an Lmhosts file if **Enable LMHOSTS Lookup** is enabled in the **Internet Protocol (TCP/IP)** properties for the connection.
6. Check a Hosts file.
7. Query a DNS server.

For a more detailed explanation and flowchart of this process, see the *Microsoft Windows Server 2003 Resource Kit*.

# WINS Features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS features**

- Persistent connections
- Burst handling
- Deleting and tombstoning records
- Accepting replication partners
- Blocking replication partners

# Persistent connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Persistent connections**

WINS includes persistent connections between WINS server replication partners. This feature is helpful because WINS servers typically disconnect from their replication partners each time replication is completed. In many cases, where WINS servers are interconnected through high speed local area network (LAN) links it is preferable to keep connections open, instead of closing them after replication is completed.

For earlier versions of WINS, servers had to establish new WINS-server connections each time replication was performed. Each connection required a modest number of processor cycles. For efficiency, network managers often configured systems to accumulate a larger number of changes before connecting for replication. These reconnections sometimes caused delays (perhaps as long as several minutes) in the replication process, which meant that inconsistencies in server data commonly existed when WINS was updated.

Persistent connections increase the speed of replication because a server can immediately send records to its partners, without incurring the cost of establishing temporary connections each time. Every record is immediately updated across the network upon registration in WINS, making replication more consistent. The bandwidth used by persistent connections is minimal because the connection is usually idle.

# Burst handling

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Burst handling**

WINS servers can now support handling of high-volume (*burst*) server loads. Bursts occur when a large number of WINS clients actively and simultaneously try to register their local names in WINS, such as when a power failure occurs. When power is later restored, many users start and register their names simultaneously on the network, which creates high levels of WINS traffic. With burst-mode support, a WINS server can respond positively to these client requests, even before it processes and physically enters those updates in the WINS server database.

**How burst handling works**

Burst mode uses a burst-queue size as a threshold value. This value determines how many name registration and name refresh requests sent by WINS clients are processed normally before burst-mode handling is started. By default, the value is 500. A WINS server initiates burst handling whenever the number of WINS client registration requests exceeds the burst queue size.

In burst handling, additional client requests are immediately answered with a positive response by the WINS server. The response also includes a varied Time-To-Live (TTL) to clients, which helps regulate the client registration load and distribute processing of the requests over time. This slows the refresh and retry rate for new WINS clients and regulates the burst of WINS client traffic.

For example, if the burst queue size is 500 entries and more than 500 requests are active, the WINS server replies immediately to the next 100 WINS client registration requests by sending early success responses with a starting TTL value of 5 minutes. For each additional round of 100 client requests, 5 minutes more are added to the TTL up to a maximum of 50 minutes. If WINS client traffic still arrives at burst levels, the next round of 100 client requests are answered with the initial TTL value of 5 minutes, and the entire process for incrementing the response TTL is repeated.

This behavior continues until the WINS server reaches its maximum intake level of 25,000 name registration and refresh queries. At this point, the WINS server begins dropping queries.

Using the WINS console, you can choose to configure the level of burst handling used by the server which modifies the size of the burst queue to accommodate either a low, medium or large burst situation.

### Note

- Any WINS server running Windows NT Server 4.0 (with Service Pack 3 or later installed), Windows 2000, or Windows Server 2003 can use burst handling.

# Deleting and tombstoning records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Deleting and tombstoning records**
WINS now provides improved database management through support for the following deletion operations:

- Simple deletion, for deleting WINS database records stored on a single server database.
- Tombstoned deletion, for deleting WINS database records replicated to databases on other WINS servers.
- Multiple-group selection of displayed database records, for either simple or tombstoned deletion.

The WINS console also provides a simple and convenient utility for administratively removing records of any type, regardless of whether they are statically or dynamically added. In previous releases of Windows NT Server, other available WINS management tools could only administratively delete entries (such as static mappings) that were added in the same way.

**How simple deletion works**
Simple deletion removes the records that are selected in the WINS console only from the local WINS server you are currently managing. If the WINS records deleted in this way exists in WINS data replicated to other WINS servers on your network, these additional records are not fully removed. Also, records that are simply deleted on only one server can reappear after replication between the WINS server where simple deletion was used and any of its replication partners.

**How tombstoning works**
Tombstoning marks the selected records as *tombstoned*, that is, marked locally as *extinct* and immediately released from active use by the local WINS server. This method allows the tombstoned records to remain present in the server database for purposes of subsequent replication of these records to other servers.

When the tombstoned records are replicated, the tombstone status is updated and applied by other WINS servers that store replicated copies of these records. Each replicating WINS server then updates and tombstones these records. When all WINS servers complete replication of these records and a specified length of time (determined by the **Verification Interval** for each WINS server) elapses, then the records are automatically removed from WINS.

In most cases, WINS records should be tombstoned at the WINS server that originally owned them, to prevent deleted records from reappearing in WINS after subsequent replication with other servers.

The owner of a given WINS server record is typically the first server contacted by the WINS client during the registration process and the server first actually used to register the local names of the client in WINS. In most cases, the WINS server that owns a client name record in WINS corresponds to the primary WINS server configured on the WINS client computer. If the configured primary WINS server is not available during client registration, a configured secondary WINS server can perform the actual registration of the client name and become the owner.

**Example: Tombstoned deletion in replicated WINS**
Where tombstoned deletion is used at an owner WINS server in a replicated WINS network, the following sequence of events occurs to effect full removal of the selected records from all WINS servers that replicate the selected records.

1. The owner WINS server marks and changes the status of the selected WINS records in its local WINS server database from Active to Tombstoned.

   WINS then treats the records as inactive and released from use. When these records are tombstoned locally, the owner WINS server does not respond to or resolve NetBIOS name queries for these names unless the records are registered again by the WINS client.

2. The owner WINS server replicates the selected records as tombstoned at subsequent replication cycles.

   The records are not forcibly and immediately removed from WINS, but are flagged or marked for eventual deletion. The exact replication cycle interval is configured by the server **Intervals** properties, which you set in the WINS console. Records are not removed from WINS data until the time specified in the **Extinction Interval** has elapsed. This allows

other WINS servers to be notified that these records are no longer in use, update their replicated mappings for these records, and further replicate this updated WINS data to other servers.

3. Records become extinct on all replicated WINS servers and are eventually removed physically from all WINS servers.

   After all WINS servers that participate in replication complete a full replication cycle and arrive at a consistent state, the tombstoned records expire and are removed from the WINS database of each server at the next database scavenging operation. After scavenging occurs on all servers, the records no longer appear in the WINS console and are no longer physically stored in the WINS database.

**Note**

- Even if records are manually tombstoned (or otherwise marked as released by WINS), released records remain in the WINS database for some time before being completely removed during subsequent scavenging operations. Typically, the period before records extinction is the sum of the following settings:

  - Extinction interval
  - Extinction time-out
  - Verification interval

  You can view and configure these settings on the **Intervals** tab, located in the **WINS <*Server name*> Properties** dialog box. By default, the sum of these settings is 36 days.

# Accepting replication partners

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Accepting replication partners**

When determining a replication strategy for your organization, you can define a list that controls the source of incoming name records during pull replication between Windows Internet Name Service (WINS) servers.

In addition to blocking name records from specific replication partners, you can also accept name records owned by specific WINS servers during replication, thereby excluding the name records of all servers that are not on the list. Although both options exist, only one can be used. You can either block replication of records from specified servers or accept replication of records from specified servers.

Pull replication based on a list of accepted WINS servers can occur either during normal pull replication or in response to push replication.

To use the **Accept records only for these owners** feature, see Accept replication of records owned by a specified server.

**Deciding whether to block or accept name records**

When using multiple WINS servers, you can either block or accept records owned by specific servers. For example, when you remove a WINS server from service, you can block the server records from pull replication. For more information, see Blocking replication partners.

If your WINS server topology is designed around the recommended hub-and-spoke model, you can use **Accept records only for these owners**. In this way, you can maintain push-pull replication between key WINS servers (hubs) without replicating the name records for other replication partners (spokes) of each individual server.

### Note

- When choosing between blocking or accepting WINS server name records, it is useful to use the shorter of two lists. For example, if you have twenty WINS servers and you want to replicate records owned by only five of them, you can choose **Accept records only for these owners** instead of **Block records for these owners**. In doing so, you need only to create a list of five servers instead of 15.

# Blocking replication partners

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Blocking replication partners**

After a WINS server is removed from a network, records it owns often continue to be distributed to other servers in a replicated WINS environment. In some cases, WINS either does not delete entries owned by servers no longer active on the network or, by default, preserves stale mappings for an extended period of time. The reasons for this can include:

- Static mappings originally made at the inactive WINS server replicate indefinitely among the other active WINS servers in the network, unless manually deleted or tombstoned.

  By default, WINS servers preserve static name-to-address mappings whenever such mappings are challenged. This can occur when a WINS client attempts to register or update a static mapping with a dynamic-type mapping of the same name. This behavior, however, can be modified. For more information, see Using static mappings.

- Dynamic mappings originally registered at the inactive WINS server are not immediately removed from the WINS database.

  In the case of dynamic entries, WINS must check with an owner server before entries are deleted. If a WINS server is no longer in use, other WINS servers conservatively retain records they cannot verify as expired on an owner server. For example, even if an owner is deleted in the WINS console for a particular server, these unwanted or stale names can continue to be replicated throughout the WINS network from other servers storing replica entries of these names.

  WINS processes removal of dynamic-type records according to the **Intervals** tab settings for **WINS Server Properties**, which can be configured at each server. In most default cases, a typical extinction period of 36 days passes before a record is fully removed from the database through the scavenging process.

For the Windows Server 2003 family, the **Block records for these owners** feature (formerly known as PersonaNonGrata) provides a way to block further replication of entries owned by inactive servers. It allows you to specify the IP addresses of inactive WINS servers that own records you want to assure are removed from WINS in a timely manner.

Because it can prevent any further replication of unwanted or stale mappings between active WINS servers remaining on the network, this feature is most useful when you need to remove WINS servers from your network.

**Example: Blocking records owned by a WINS server**

The following graphic shows a typical WINS replication design based on the recommended hub-and-spoke WINS replication design model.




In the example replicated WINS network, WINS-D, is an inactive spoke server in the process of being removed from the network. Because WINS-A is the hub in this design, it must be configured to block records for WINS-D in order to prevent

replica entries for WINS-D that are present at other spokes from replicating back to WINS-A and reappearing throughout the network.

When WINS-D server is given blocked partner status, it can be deleted as an owner in the WINS console at WINS-A and its entries will not be replicated back from other spoke servers, such as WINS-B, WINS-C or WINS-E. For more information, see Delete an owner from the WINS database.

**Notes**

- You can also choose to accept name records only from specific WINS servers during pull replication. For more information, see Accepting replication partners.

- If you are not using a hub-and-spoke WINS replication model, it might be necessary to enforce blocking at other servers in your network. If this is the case, you might need to list the blocked partner by blocking records owned by servers or for WINS running on earlier versions, add and configure the PersonaNonGrata registry parameter at each server. These steps can help achieve the intended results of full replication for records owned by the partner to other WINS servers in the network.

- After adding IP addresses to the list of blocked partners at a WINS server, it is recommended to stop and restart WINS to ensure that the change is in effect the next time replication at that WINS server occurs. For more information, see Start or stop a WINS server.

# Administering WINS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Administering WINS**

This section provides information about administering WINS:

- Managing WINS from the command line
- Configuring WINS replication
- Using static mappings
- Modifying WINS server defaults
- WINS interoperability issues
- WINS Database Management

# Managing WINS from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Managing WINS from the command line**
The Netsh commands for WINS offer a fully equivalent command-line tool for administering WINS servers that can be used as an alternative to console-based management. This is useful in the following situations:

- When managing WINS servers in wide area networks (WANs), commands can be used in interactive mode at the netsh command prompt to better manage over slow-speed network links.

- When managing a large number of WINS servers, commands can be used in batch mode to script and automate recurring administrative tasks that need to be performed for all WINS servers.

For more information about how to use netsh commands provided for WINS, see Use WINS Command-line Tools.

**Notes**

- For a comprehensive netsh WINS command reference (including syntax and parameters), see Netsh commands for WINS.

- For an example of how to use netsh WINS commands in batch files and scripts, see Netsh WINS example.

- The terms Persona Non Grata and Persona Grata refer to the Windows WINS server features **Block records for these owners** and **Accept records only for these owners**, respectively.

- For more information about netsh, see The Netsh Command-Line Utility.

# Configuring WINS replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Configuring WINS replication**

Before configuring replication, you should first carefully design and review your WINS replication topology. For wide area networks (WANs), this planning can be critical to the success of your deployment and use of WINS. In general, WINS offers you the following options and possibilities to choose from when you are configuring replication:

- You can manually configure WINS replication for a WAN environment.
- For some larger campus networks, you might also configure WINS to replicate within a LAN environment.
- In smaller or bounded LAN installations, you might consider enabling and using WINS automatic partner configuration for simplified setup of WINS replication.
- In some larger or more global installations, you might have to configure WINS across untrusted Windows NT domains.

**Replication across wide area networks**

When configuring WINS replication, the two most important planning issues are:

- Whether your WINS replication occurs over slower WAN links.
- The length of time required for all replicated changes in the WINS database to converge and achieve consistency on the network.

The frequency of WINS database replication between WINS servers is a major planning issue. The WINS server database should be replicated frequently enough to prevent the downtime of a single WINS server from affecting the reliability of the mapping information in other WINS servers. However, the time interval between replications should not be so small that it interferes with network throughput.

Network topology can influence your decision on replication frequency. For example, if your network has multiple hubs connected by relatively slow WAN links, you can configure WINS database replication between WINS servers on the slow links to occur less frequently than replication on the local area network or on fast WAN links. This reduces traffic across the slow link and reduces contention between replication traffic and WINS client name queries.

For example, as shown in the following figure, the intervals used for configuring pull replication for WINS servers could be set according to proximity between servers and the likelihood of a high-speed link connecting servers configured as replication partners of each other.




In this example, two WINS servers--SEA-WINS and ATL-WINS--are located in Seattle and Atlanta. Another two servers--MEX-WINS and SYD-WINS--are located outside the United States in Mexico City and Sydney. Because SEA-WINS and ATL-WINS are located across a higher-speed WAN link, they are configured to replicate once every hour. For the servers outside the United States, replication intervals might be set higher, such as either 3 or 12 hours, because lower-speed links are associated with these connections.

**Planning the convergence time**

*Convergence time* is the time needed to replicate a new entry in a WINS database from the WINS server that owns the entry to all other WINS servers on the network. When planning for WINS servers, you must decide what is acceptable as the convergence time for your network.

Using the previous example, if a WINS client in Seattle registers its name with the WINS server in the figure labeled SEA-WINS, other WINS clients in remote locations might query for the client name and not find it in WINS. In some cases, if the address of the client has been changed, the remote locations would only have a mapping for the previous IP address of the client in their local server databases.

WINS clients that query any of the other WINS servers in the figure (that is, ATL-WINS, MEX-WINS, SYD-WINS) only receive a positive and correct response when the new or updated mapping for the client has been fully replicated from the owning WINS server, SEA-WINS, to all of the other servers. Replications from SEA-WINS could be triggered in one of two ways:

- A push replication trigger is sent based on the frequency of updates made at SEA-WINS.
- A pull replication trigger is sent to other WINS servers based on the configured **Replication interval** set in **Replication Partners Properties** at SEA-WINS.

Given this configuration, a new or updated entry at SEA-WINS is only replicated when the pull replication interval expires, but queries at remote sites for the new name mapping might not succeed until later. This is because the time interval for replication to ATL-WINS is one hour; it is three hours for MEX-WINS; and it is 12 hours for SYD-WINS. In this case, WINS convergence time is calculated as:

```
12 hours + 2 * (1 hour) = 14 hours
```

In reality, name query requests can succeed before the convergence time elapses. This happens when the entries replicate over a shorter path than the worst-case path. It also happens when the **Number of changes in version ID before replication** threshold in push replication settings passes before the **Replication Interval** expires in pull replication settings. This results in earlier replication of the new entry. The longer the replication path, the longer the convergence time.

**Notes**

- After each administrator configures the respective WINS server, you can use **Replicate Now** (in the WINS console) to start immediate replication between the configured WINS servers.
- When replication occurs, you can use **Display Records** to view records for the remote WINS server database. However, to make changes to the remote WINS database, you must be logged on under an account with required user rights in the respective domains for both servers involved.

- Replication can also occur between two separate stand-alone servers that participate in a workgroup.

**Replication across local area networks**
When configuring WINS replication for local area networks (LANs), the issues are similar to those that occur in WAN environments, although less critical.

Because the speed of the underlying network links for LANs are much greater than for WANs, it might be acceptable to increase the frequency of WINS database replication, by optimizing push and pull parameters for LAN-based replication partners. For push/pull partners, you can do this by decreasing the **Number of changes in version ID before replication** and **Replication interval** settings accordingly, from what you would use for WAN-based partners on slower links.

For example, between LAN-based replication partners it often works to enable WINS to use a persistent connection between the servers. If a persistent connection was not used, the normal update count threshold would default to a minimum of 20. With persistence, a smaller update count value can be specified.

Next, you could specify a much smaller number, such as a value of one to three in **Number of changes in version ID before replication** before WINS sends a push replication trigger to the other partner. For pull partners, you might also consider setting the **Replication interval** to a value in minutes, instead of hours.

As in WAN replication planning, the WINS server database should be replicated frequently enough to prevent the downtime of a single WINS server from affecting the reliability of the mapping information in other WINS servers. However, the time interval between replications should not be so small that it interferes with network throughput.

For network-intensive environments, it is a good idea to use a network monitoring tool such as Network Monitor to help measure and determine how to optimize your WINS replication strategy.

**Automatic partner configuration**
You can configure a WINS server to automatically configure other WINS server computers as its replication partners. With this automatic partner configuration, other WINS servers are discovered when they join the network and are added as replication partners.

Automatic configuration is possible because each WINS server announces its presence on the network through periodic multicasts. These announcements are sent as IGMP messages for the multicast group address of 224.0.1.24 (the well-known multicast IP address reserved for WINS server use).

Automatic replication configuration monitors multicast announcements from other WINS servers and automatically performs the following configuration steps:

- Adds the IP addresses for the discovered servers to its list of replication partner servers.
- Configures the discovered servers as both push and pull partners.
- Configures pull replication at two-hour intervals with the discovered servers.

If a remote server is discovered and added as a partner through multicasting, it is removed as a replication partner when WINS shuts down properly. To have automatic partner information persist when WINS restarts, you must manually configure the partners.

To manually configure replication with other WINS servers, use the WINS console to specify roles for each partner and any related information.

### Notes

- Automatic partner configuration is typically useful in small networks, such as single subnet LAN environments. It can, however, be used in routed networks. For WINS multicast support in routed networks, the forwarding of multicast traffic is made possible by configuring routers for each subnet to forward multicast traffic for the reserved WINS server multicast group. The destination IP address for this group is 224.0.1.24 between routed subnets.
- Because periodic multicast announcements between WINS servers can add traffic to your network, automatic partner configuration is recommended only if you have a small number of installed WINS servers (typically, three or fewer) on the reachable network.

**Replication between untrusted domains**

It is possible to set up WINS replication between one or more WINS servers in domains that do not have a trust relationship. You can do this without a valid user account in the untrusting domain. To configure replication, an administrator for each WINS server must use the WINS console to manually configure that server to permit this replication.

# Using static mappings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Using static mappings**

Mapped name-to-address entries can be added to WINS in either of two ways:

- Dynamically, by WINS-enabled clients directly contacting a WINS server to register, release, or renew their NetBIOS names in the server database.
- Manually, by an administrator using the WINS console or command-line tools to add or delete statically mapped entries in the server database.

Static entries are only useful when you need to add a name-to-address mapping to the server database for a computer that does not directly use WINS. For example, in some networks, servers running other operating systems cannot register a NetBIOS name directly with a WINS server. Although these names might be added to and resolved from an Lmhosts file or by querying a DNS server, you might consider using a static WINS mapping instead.

**Migrating static mappings to dynamic mappings**

Unlike dynamic mappings, which can age and be automatically removed from WINS over time, static mappings can remain in WINS indefinitely unless administrative action is taken.

By default, if during an update process WINS is presented with both a static and dynamic-type entry for the same name, it preserves the static-type entry. You can, however, use the **Overwrite unique static mappings at this server (migrate on)** feature provided with WINS to change this behavior.

For example, the following graphic shows the default migration behavior for a WINS server, WINS-A, when presented with a registration request by a WINS client, Host C, that challenges a static mapping already in the server database.




In this example, WINS-A prevents the client from registering and updating the statically mapped information for Host C. In the next graphic, the same situation is repeated but in this example, the Migrate feature has been turned on at WINS-A.




As shown in the revised graphic, enabling the Migrate feature at WINS-A permits previously entered static mappings for Host B and C to be overwritten by WINS client computers that attempt to register dynamically and update entries for these same

names. This feature is useful in simplifying the update of legacy NetBIOS b-node (broadcast-only) type clients that are later upgraded to support WINS.

**Important**

- Never add static mappings for computers that can directly use WINS. Doing so can cause undesirable problems, such as demanding additional work to migrate, delete, or tombstone these mappings in a replicated WINS environment.

**Notes**

- You can add static entries either interactively using the WINS console or command-based tools, or by importing an Lmhosts file.

- The **Overwrite unique static mappings at this server (migrate on)** feature was designed primarily to migrate records for older non-WINS clients that do not dynamically register themselves in WINS. Some special group names, such as domain [1Ch] names are not affected by this setting.

# Modifying WINS server defaults

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Modifying server defaults**

WINS uses several default server configuration parameters that determine how NetBIOS name records are managed in the WINS server database. These parameters affect the default settings applied by each WINS server for any records stored locally.

The following sections describe the default, minimum, and maximum settings that can be applied for database records. Included is the examination of behavior differences in servers running under each of the released versions of WINS. In most cases, default values are acceptable. In some cases, you might choose to modify the defaults to accommodate special circumstances, such as when a host name change must be made or when you renumber your network so that hosts use different IP addresses.

All of these settings are configured on the **Intervals** tab of the **WINS <*Server name*> Properties** dialog box.

**Renew interval**

The WINS renew interval setting specifies how often a client renews registration of its names. The following table shows the supported range of values for **Renew interval**:

| | |
|---|---|
| **Default:** | 5 hours (Windows NT Server 3.5)<br>4 days/96 hours (Windows NT Server 3.51)<br>6 days/144 hours (Windows NT Server 4.0, Windows 2000 Server, and the Windows Server 2003 family) |
| **Maximum:** | 999 hours (Windows NT Server 3.5, 3.51, and 4.0)<br>365 days (Windows 2000 Server and the Windows Server 2003 family) |
| **Minimum:** | 40 minutes (all versions of Windows NT Server, Windows 2000 Server, and the Windows Server 2003 family) |

When adjusting this value from its service default, consider the following guidelines:

| **Increasing the renew interval ...** | **Decreasing the renew interval ...** |
|---|---|
| Causes client names to be renewed with WINS less often. | Causes client names to be renewed with WINS more often. |
| Reduces client registration activity and performance loads on the WINS server. | Increases client registration activity and performance loads on the WINS server. |
| Allows WINS data to become less consistent with the current state of the network (when WINS names are changed or modified, changes are reflected less quickly in the WINS database). | Allows WINS data to become more consistent with the current state of the network (when WINS names are changed or modified, changes are reflected more quickly in the WINS database). |

**Important**

- Any changes made to the WINS renew interval also directly affect the extinction interval and extinction time-out values, which are dependent on this setting.

**Extinction interval**

The WINS extinction interval specifies the time interval between when an entry is marked as Released and when it is marked as Extinct in the WINS database. The default value is dependent on the renew interval and, if the WINS server has replication partners, on the maximum replication time interval. The following table shows the supported range of values for **Extinction interval**:

| | |
|---|---|
| **Default:** | (Dependent on **Renew interval**) |

| | |
|---|---|
| **Maximum:** | 999 hours (Windows NT Server 3.5, 3.51, and 4.0)<br><br>365 days (Windows 2000 Server and the Windows Server 2003 family) |
| **Minimum:** | 40 minutes (Windows NT Server 3.5, 3.51, Windows 2000 Server, and the Windows Server 2003 family)<br><br>1 hour (Windows NT Server 4.0) |

**Extinction timeout**

The WINS extinction timeout specifies the interval between when an entry is marked extinct and when the entry is scavenged from the database. The default value is dependent on the renew interval and, if the WINS server has replication partners, on the maximum replication time interval. The following table shows the supported range of values for **Extinction time-out**:

| | |
|---|---|
| **Default:** | (Dependent on **Renew interval**) |
| **Maximum:** | 999 hours (Windows NT Server 3.5, 3.51, and 4.0)<br><br>365 days (Windows 2000 Server and the Windows Server 2003 family) |
| **Minimum:** | 24 hours (all versions of Windows NT Server, Windows 2000 Server, and the Windows Server 2003 family) |

**Verification interval**

The WINS verification interval specifies the time after which the WINS server must verify that old names that it does not own are still active. The default is dependent on the extinction interval. The following table shows the supported range of values for **Verification interval**:

| | |
|---|---|
| **Default:** | (Dependent on **Renew interval**) |
| **Maximum:** | 9999 hours (Windows NT Server 3.5, 3.51, and 4.0)<br><br>365 days (Windows 2000 Server and the Windows Server 2003 family) |
| **Minimum:** | 24 days (all versions of Windows NT Server, Windows 2000 Server, and the Windows Server 2003 family) |

# WINS interoperability issues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Interoperability issues**

**Using DHCP with WINS**
When using DHCP and WINS together on your network, consider the following options for interoperation:

- **Use additional DHCP scope options.**

  Specifically, use DHCP options to assign WINS node types (option type 46) and to identify WINS servers for use by DHCP clients (option type 44). In some cases, this can involve adjusting these option types for each physical subnet in which DHCP and WINS are implemented.

- **Assign a length of time for DHCP lease durations comparable to the time WINS uses for renew intervals, when both services are deployed.**

  By default, DHCP leases are eight days in length and the WINS renew interval is six days. If lease lengths for DHCP widely differ from WINS renew intervals, the effect on your network can be an increase in lease-management traffic and might cause a WINS registration for both services. If you shorten or lengthen the DHCP lease time for clients, then modify the WINS renew interval accordingly.

**Using DNS with WINS**
If a large percentage of clients use NetBIOS and you are using DNS, consider using **WINS lookup on your DNS servers.** If WINS lookup is enabled on Microsoft DNS Service, WINS is used for final resolution of any names that are not found using DNS resolution. The WINS forward lookup and WINS-R reverse lookup records are supported only by DNS. If you use servers on your network that do not support DNS, use the DNS console to ensure that these WINS records are not propagated to DNS servers that do not support WINS lookup.

**Multihomed WINS servers**
For multihomed WINS servers, where more than one local network connection is present and active, WINS does not guarantee the binding order for NetBIOS. Therefore, for a multihomed WINS server to function properly, all installed connections should be configured as routable interfaces.

All multihomed WINS servers should have their primary IP addresses assigned to each network connection, configured as push and pull replication partners at other servers that can be reached through each of the respective connections. When configuring replication partners, you can ensure that a specific network connection is used by specifying an IP address for the remote multihomed server you are adding at a WINS server.

Otherwise, when replication partners are specified and resolved by a name entered in the WINS console, it is possible that a packet generated by WINS could use any of the installed interfaces and their respective IP addresses. This seemingly random behavior results from WINS referring to its local IP routing table, which contains all of the installed IP interfaces, before it sends packets to the remote server.

# WINS Database Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS database management**

- Finding and viewing WINS records
- Maintaining the WINS database
- Checking the WINS database for consistency

# Finding and viewing WINS records

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Finding and viewing WINS records**

The WINS console provides many methods for filtering and displaying WINS database records. Record searches are initiated in the WINS console **Display Records** dialog box. When WINS database records match search parameters, one or more records are displayed in the details pane of the WINS console. While a search is in progress, search statistics are dynamically displayed in the status bar above the details pane. Search statistics include the number of records scanned and filtered. When no WINS database records match search parameters, no records are displayed in the details pane. The methods of filtering WINS database records are detailed in the paragraphs that follow.

For information on how to access the **Display Records** dialog box, see Start a record search.

**Filtering records for display**

When you create a search filter, you choose from three categories of filters:

- Record owners - A record query can be based on the name records of one or more name record owners.
- Record types - A record query can be based on one or more NetBIOS name suffix record types.
- NetBIOS name and IP address - A record query can be based on all or part of a NetBIOS name, on IP address, or it can be based on NetBIOS name and IP address, including or excluding subnet mask as a search parameter.

You can shape a record filter from all three categories above. For example, in one filter you can specify one or more record types, one or more record owners, a partial or complete NetBIOS name, and IP address.

The properties of each returned record are displayed in the details pane. In addition to displaying the Record Name, Type, IP Address and Owner, the details pane also displays State (Active, Released or Tombstoned), Static (x for static, null or blank for not static), Expiration (date and time of record expiration), and Version information for each record.

Because most names in the database are computer names and computer names are always registered as uppercase, WINS only does uppercase searches. If you type lowercase letters in **Filter records matching this Name pattern**, WINS converts lowercase letters to uppercase letters before it performs the search.

You can alter this behavior, however, by using the **Case sensitive match** check box. When checked, the record search is performed matching the supplied criteria, including capitalization.

**Record owners**

For a search operation to successfully return one or more name records, you must select at least one record owner. You can select any combination of record owners or all record owners.

This filter is ideal for doing full display searches and obtaining either a full or filtered view of the WINS database for one or more servers that own records registered in the database.

Because full database views potentially require WINS to download large amounts of data from one or more servers, its use should be handled carefully, especially on WINS networks where a large number of WINS servers are deployed. However, for in-depth troubleshooting and close viewing of local WINS databases (either local viewing on the same WINS server or viewing a WINS server that is on the same local area network), this feature should not present an overly demanding load for either the WINS server or the local area network involved.

However, if multiple servers are selected in **Display records for these owners**, or if a remote owner is selected (such as another WINS server that is located on a wide area network), a substantial or undesirable increase in network traffic can occur. For this reason, you should test the impact of a multiple-server search on network traffic before you use it during peak periods, or you should use a multiple-server search during off-peak hours only.

If a search generates considerable undesired network traffic, the search can be stopped before it completes. For more information, see Stop a record search.

When viewing records by owner, you can additionally filter WINS data to be included in the view by selecting and using

NetBIOS name suffix types as criteria.

**Record types**
For a search operation to successfully return one or more name records, you must select at least one NetBIOS name suffix type. You can select any combination of record types or all record types.

The WINS console provides you with the ability to add, edit, and remove your own NetBIOS name suffix filter types. It also provides a predefined list that is sufficient for typical use.

Filtering records by suffix type allows for more precise network troubleshooting. For example, to isolate browsing problems associated with the Computer Browser service when WINS clients are using browsing on a WAN, you can filter WINS data by including only those records that use the [1Bh] suffix type, which maps the domain master browser record for each primary domain controller (PDC) in a Windows NT domain environment.

For more information, see Add, modify, or remove a filter type.

**NetBIOS name and IP address filters**
Following are some of the methods used for filtering and displaying records based on NetBIOS name and IP address. These filters can be used in different combinations to achieve varying results; however, you must use at least one record type and one record owner for the NetBIOS name and IP address filters to function properly.

**Filter records matching this Name pattern** is used to query the WINS database and filter displayed results to a limited number of NetBIOS name records that start with a specified set of one or more characters. This filter is ideal for targeting either a known name or similar names that are located on a specific WINS server. For example, you can use this filter to troubleshoot WINS name resolution problems in a wide area network (WAN) by verifying that a known NetBIOS name exists in the data that is stored at any one of your remote WINS servers. This filter typically generates minimal network traffic and can be used routinely for locating the source of "name not found" errors that can result from NetBIOS name resolution.

**Filter records matching this Name pattern** locates records through either a full NetBIOS name or a specified set of matching characters. For WINS to include a targeted name in the results of a find operation, the name must begin with the exact characters that are specified in **Filter records matching this Name pattern**. However, names that match the specified search criteria can contain additional characters after the specified string. For example, if "A," "AB," and "AXB" are all NetBIOS name records that begin with the letter "A" in the WINS database, then entering "A" as the search string finds all three records. If you enter "AB" as the search string, only the record "AB" is found. If you enter "ABC" as the search string, none of the three example records are found or returned in your search results.

**Filter records matching IP address** searches for and displays records that have a specific IP address.

**Filter records matching this Name pattern** and **Filter records matching IP address** can be used together to search for and display only records in which the name and IP address match the search criteria.

**Filter records matching IP address** and **Match the IP address based on subnet mask** can be used together to search based on an IP address and subnet mask pair. The subnet mask indicates which sections of the IP address are interpreted as a network ID and which sections of the IP address are interpreted as a host ID. The search is performed for all records with a network ID that matches the IP address search criteria. For example, if the IP address search criteria is 192.168.0.1, and the subnet mask search criteria is 255.255.0.0, the first two sections of the IP address are interpreted as network ID, and all of the name records that have an IP address of 192.168.*x.y* -- where *x* is any number between 1 and 254 and *y* is any number between 1 and 254 -- are returned to the details pane. In another example, if the IP address search criteria is 192.168.0.1, and the subnet mask search criteria is 255.255.255.0, the first three sections of the IP address are interpreted as network ID, and all name records that have an IP address of 192.168.0.*x* -- where *x* is any number between 1 and 254 -- are returned to the details pane.

For more information about IP addressing and subnet masks, see IP addressing and Subnet masks.

For more information about using NetBIOS name filtering, see Search for names that begin with a specified set of characters.

**Result caching**
Result caching allows you to download the records of WINS name record owners to the local computer. This provides faster performance during subsequent query operations on the record set. The record set that is returned to the local computer is stored in local random access memory (RAM). If a large record set is stored in RAM, performance on the local computer might suffer, depending on the amount of RAM and the size of the record set.

If your network has large WINS databases and many WINS servers, it is recommended that you test the performance of **Enable result caching** during off-hours, before you use it during peak network hours.

When **Enable result caching** is checked, caching behavior varies depending on the selected search criteria.

**Queries for records from one WINS server**
If you query for the records of one name record owner, with or without additional query parameters, all records for that owner will be returned and cached in RAM. The additional parameters, if any -- such as name or IP address -- will cause the records to be filtered so that only records matching the detailed query parameters are displayed for view in the details pane. New record searches on the same owner's records with the same or different parameters will cause the locally cached record set to be queried, not the WINS server. The performance of these subsequent queries should be high, depending on the size of the record set and the amount, type and speed of RAM.

**Queries for records from more than one WINS server**
If the query includes requests for records from more than one name record owner:

- The entire WINS database is downloaded to the local RAM cache if no other parameters (such as NetBIOS name or IP address) are specified. All subsequent queries run on the locally cached record set, not on the WINS server.

- Only records containing the additional query parameters are downloaded to the local computer's RAM cache if the query contains other parameters, such as NetBIOS name or IP address.

When you close the WINS console, all records in the RAM cache are cleared. If you reopen the console and initiate a new search, unless you select **Enable result caching**, all queries will be made to the selected WINS server, and only the records that match the query parameters are returned to the local computer and displayed in the details pane. No records will be cached in RAM, and each subsequent query will go to the server, regardless of the query parameters that you have selected.

# Maintaining the WINS database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Maintaining the WINS database**

The WINS console provides tools for maintaining the WINS server database. For example, you can use the WINS console to back up and restore WINS server database files. This should be done periodically in a way that is similar to other backup operations that are planned and used for saving files on a server computer.

Scavenging and compaction are additional maintenance operations that can be used selectively to help reduce database size and reclaim disk space that is no longer used by the database. Dynamic compaction is performed automatically by WINS at periodic intervals; manual compaction is not generally necessary.

**Backing up the database**

The WINS console provides its own backup feature for backing up the WINS server database during normal operation. After you specify a backup directory for the database, WINS performs complete database backups every twenty four hours, by default.

You should also periodically back up the registry entries for the WINS server.

**Scavenging the database**

*Scavenging* the WINS server database is the process of deleting and removing obsolete information that remains after changes occur. Like any database, the WINS server database of address mappings needs to be periodically scavenged and backed up.

The local WINS server database sometimes retains both released and old entries that were registered at another WINS server (and replicated to the local WINS server). Automatic scavenging takes place at intervals defined by the relationship between the Renewal and Extinct intervals you define. You can also manually clean the database.

The following table describes the results of scavenging the WINS database for a WINS server. The first column shows the record criteria and state of any database entries before scavenging is started. The second column indicates the state of the same entries after the scavenging process.

| Entry state before scavenging | Entry state after scavenging |
|---|---|
| Active names owned by this WINS server for which the **Renew Interval** has expired | Marked as Released. |
| Released names owned by this WINS server for which the **Extinction interval** has expired | Marked as Tombstoned. |
| Tombstoned names for which the **Extinction timeout** has expired | Deleted and removed from the database. |
| Tombstoned names, replicated from other servers for which the **Extinction timeout** has expired | Deleted and removed from the database. |
| Active names replicated from other servers for which the **Verification interval** has expired | Revalidated. |
| Tombstoned or deleted names replicated from other servers | Removed from the database. |

# Checking the WINS database for consistency

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Checking the WINS database for consistency**

Consistency checking helps maintain database integrity among WINS servers in a large network. When you initiate consistency checking at the WINS console, all records are pulled from each owner listed in the current server database, including other WINS servers that are indirect (not directly configured) replication partners.

All records pulled from remote databases are compared to records in the local database, using the following checks for consistency:

- If the record in the local database is identical to the record pulled from the owner database, its time stamp is updated.
- If the record in the local database has a lower version ID than the record pulled from the owner database, the pulled record is added to the local database, and the original local record is marked for deletion.

**Verifying registered name records on WINS servers**

The WINS console provides a **Check Registered Names** command that enables you to perform the following verification process among WINS servers on your networks:

- A NetBIOS name records list, that can be specified either manually or within a stored file called Names.txt, that contains names to be used for testing and verification.
- A WINS server list, that can be specified either manually or within a stored file called Servers.txt, that contains a list of WINS servers to be used in verifying NetBIOS names you specify for testing and verification checking.

WINS uses a NetBIOS names list (such as Names.txt) to build a query and test for those records with any matched records currently registered and stored in local databases at other WINS servers. The list of WINS servers used for checking the specified names is determined by the IP addresses of WINS servers included in a server list (such as Servers.txt) which is the list of WINS servers to query.

Once initiated, WINS completes the name verification check as follows:

1. The names list is used to build a name query list.
2. Names are verified by query for the first listed name, starting with the WINS server specified in the server list.
3. For each name in the names list, WINS queries each server to verify that the name is registered.

   If a failure is encountered, WINS can retry the query up to three times at the same server.
4. Once the first listed name has been queried and checked at all servers, WINS proceeds to the next name in the list.

The process is repeated until all listed names have been checked at all listed servers. The results and outcome of each name/server query are displayed in **Checking Name Registrations**.

Typical results include the following:

| Status | Description |
|---|---|
| ok | Indicates both that the query was successfully completed and that the specified name was verified at the current WINS server. |
| no response | Indicates that the current WINS server did not respond to the NetBIOS name query. |

| | |
|---|---|
| name not found | Indicates that the current WINS server responded to the NetBIOS name query but could not resolve the specified name. |

When all specified names have been queried at all specified servers, a summary is provided, which includes the following information:

- Any specified NetBIOS names that were verified to their mapped IP address.
- Specified NetBIOS names for which not all WINS servers resolved to the same IP address, indicating that the name is stored and used inconsistently in WINS.
- WINS servers that generated an incomplete response, indicating that they did not have all of the names specified for verification.
- WINS servers that did not respond.
- Specified names for which an IP address could not be verified.

**Notes**

- If a WINS database is extremely large, the consistency checking process can be highly resource intensive, affecting network, disk, and processor usage. For this reason, consider doing this during off-peak hours when other network traffic loads and usage are minimal.
- If the **Verify name records with all partners of each selected server** check box is cleared, listed names are only verified at each of the specified servers. When this option is selected, replication partners of each listed server are additionally queried for any specified names.
- For Windows 2000 and the Windows Server 2003 family, use the WINS console to perform consistency checking. For earlier versions, consistency checking must be enabled, either through changes to the Windows NT registry or triggered through use of the WINSCL utility, provided as part of the *Windows NT 4.0 Server Resource Kit*.

# WINS Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Resources**

- NetBIOS names reference
- WINS updated technical information

# NetBIOS names reference

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**NetBIOS names reference**

For all Microsoft operating systems that support and use NetBIOS names, the first 15 characters of a name can be specified by a user. However, the 16th character of the name (00-FF hex) always indicates a resource type.

This section contains additional information and examples of NetBIOS names used by Microsoft networking components.

**NetBIOS name types: an overview**

The following are the most common types of NetBIOS names that can be used in the WINS database. When creating static mappings, these general types are available for selected use.

| Type | Description |
|---|---|
| Unique | Used to associate the computer specified by name in Computer name and a single IP address in IP address for this static mapped entry. When this type is selected, three types of records are statically added to the WINS database for the specified computer name. These are the [00h]WorkStation, [03h]Messenger and [20h] File Server types. |
| Group | Also referred to as a normal group. This type is used to add a static entry for the computer, specified by name in a static mapping, to a workgroup used on your network. If this type is used, the IP address for the computer is not stored in WINS, but is resolved through local subnet broadcasts. |
| Domain Name | Indicates a *domain name* [1C] mapped entry for locating Windows NT domain controllers. |
| Internet Group | Used for special user-defined administrative groups. You can use this to group resources. For example, you can indicate a group of file or print servers for organizing shared resources that are visible when browsing your network places. Each Internet group is represented by a shared *group name* of [20h] type in the WINS database. |
| Multihomed | Used to register a unique name for a computer that has more than one IP address (either multiple adapters each using a single address or one network adapter configured with multiple IP addresses). |

**NetBIOS Unique Names**

| Format | Description |
|---|---|
| *computer_name*[00h] | Registered by the Workstation Service on the WINS client. In general, this name is called the *NetBIOS computer name*. |
| *computer_name*[03h] | Registered by the Messenger service on the WINS client. This service is used by the client when sending and receiving messages. This name is usually appended to both the NetBIOS computer name for the WINS client computer and the name of the user currently logged on to that computer when sending messages on the network. |
| *computer_name*[06h] | Registered by the Routing and Remote Access service on the WINS client (when the service is started). |
| *domain_name*[1Bh] | Registered by each Windows NT Server 4.0 domain controller running as the domain master browser. This name record is used to allow remote browsing of domains. When a WINS server is queried for this name, a WINS server returns the IP address of the computer that registered this name. |
| *computer_name*[1Fh] | Registered by the Network Dynamic Data Exchange (NetDDE) services. This appears only if the NetDDE services are started on the computer. |
| *computer_name*[20h] | Registered by the Server service on the WINS client. This service is used to provide points of service to the WINS client to provide sharing of its files on the network. |
| *computer_name*[21h] | Registered by the RAS Client service on the WINS client (when the RAS Client is started). |

| | |
|---|---|
| *computer_name*[BEh] | Registered by the Network Monitoring Agent Service and appearing only if the service is started on the WINS client computer. If the computer name has fewer than 15 characters, the remaining character spaces are padded with plus (**+**) symbols. |
| *computer_name*[BFh] | Registered by the Network Monitoring Utility (included with Microsoft Systems Management Server). If the computer name has fewer than 15 characters, the remaining character spaces are padded with plus (**+**) symbols. |
| *username*[03h] | User names for currently logged-on users are registered in the WINS database. Each user name is registered by the Server service component so that users can receive any **net send** commands sent to the user name. If more than one user logs on with the same user name, only the first computer logged on with that name registers the name. |

**NetBIOS Group Names**

| Format | Description |
|---|---|
| *domain_name*[00h] | Registered by the Workstation Service so that it can receive browser broadcasts from LAN Manager-based computers. |
| *domain_name*[1Ch] | Registered for use by the domain controllers within the domain. These contain up to 25 IP addresses. |
| *domain_name*[1Dh] | Registered for use by master browsers, of which there is only one per subnet. Backup browsers use this name to communicate with the master browser, retrieving the list of available servers from the master browser. WINS servers always return a positive registration response for *domain_name*[1D], even though the WINS server does not *register* this name in its database. Therefore, when a WINS server is queried for the *domain_name*[1D], the WINS server returns a negative response, which forces the client to broadcast for name resolution. |
| *group_name*[1Eh] | A normal group name. Any computers configured to be network browsers can broadcast to this name, and listen for broadcasts to this name, to elect a master browser. A statically mapped group name uses this name to register itself on the network. When a WINS server receives a name query for a name ending with [1E], the WINS server always returns the network broadcast address for the local network of the requesting client. The client can then use this address to broadcast to the group members. These broadcasts are for the local subnet and should not cross routers. |
| *group_name*[20h] | A special group name called the *Internet Group* is registered with WINS servers to identify groups of computers for administrative purposes. For example, "printersg" could be a registered group name used to identify an administrative group of print servers. |
| --__MSBROWSE__[01h] | Registered by the master browser for each subnet. When a WINS server receives a name query for this name, the WINS server always returns the network broadcast address for local network of the requesting client. |

# WINS updated technical information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**WINS updated technical information**

Updated technical information for WINS can be obtained by searching technical articles available from the Personal Support Center at the Microsoft Web site. For more information, see Updated technical information.

You can also obtain additional in-depth information on WINS in the Microsoft Windows Server 2003 Resource Kit:

- For more information about WINS, see "Windows Internet Name Service" at the Microsoft Windows Resource Kits Web site.
- For more information about LMHOSTS files, see "LMHOSTS File" at the Microsoft Windows Resource Kits Web site.

# WINS Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Troubleshooting**

- Troubleshooting WINS clients
- Troubleshooting WINS servers
- Troubleshooting WINS databases
- Troubleshooting WINS replication

# Troubleshooting WINS clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Troubleshooting WINS clients**
What problem are you having?

- Name resolution failed at client computer.
- Clients do not use load balancing when connecting to a multihomed server.

**Name resolution failed at client computer.**
**Cause:** The type of name resolution used is not appropriate. Determine whether WINS or DNS is being used for name resolution and if the name failure is occurring for a NetBIOS or fully qualified domain name (FQDN).

**Solution:** NetBIOS names are 15 characters or less and are typically used to provide named access to NetBIOS services based on a computer or group name, such as Messenger, Server, or Workstation. For example, the short NetBIOS name "PRINT-SRV1" might refer to a computer used for print and file sharing. It might also be known by a full DNS computer name "print-srv1.example.microsoft.com" when referring to this same computer as a Windows Server 2003 family shared or published printer resource.

For the DNS example name used above, Winsock-based host name resolution is used to associate an IP address to a named computer and further identify it in a system of domains, either on an intranet or the Internet. This is needed for an entirely different set of applications such as to access Web, SMTP mail, FTP, or LDAP servers.

The key point is that each of these name schemes, WINS or DNS, require name resolution for different names used in different applications, and for locating different types of network services. NetBIOS has its own set of applications that demand NetBIOS name resolution, the absence of which can cause network failures. These can be repaired or detected through an Lmhosts file used temporarily to troubleshoot or correct the name resolution failure.

DNS and FQDNs use Winsock-style host name resolution and are primarily used for Active Directory and Internet use. For this type of name resolution, you can also repair or detect most problems through a Hosts file used temporarily to troubleshoot or correct the name resolution failure.

Verify the exact form of the name that has failed during resolution and determine whether the correct type of name resolution was involved.

**See also:** NetBIOS name resolution; Host name resolution

**Cause:** The client could be using an application or version of Windows that requires the use of WINS to resolve names.

**Solution:** Not all Windows computers or applications require WINS or NetBIOS over TCP/IP (NetBT). For example, if the failed name resolution occurred as part of a URL entered in a Web browser or FTP program or as part of an address in an Internet e-mail program, the problem is more likely related to DNS.

Name resolution could fail when any clients need access to shared resources not published through Active Directory. Resources such as these could include legacy file servers, print servers, domain controllers or member servers used to complete logon or browsing of Windows NT domains. Some examples of applications that a client might use and need WINS to assist in name resolution include My Network Places, the **Map Network Drive** feature in Windows Explorer, or the **net** command (Net.exe) and any of its supported options, such as **net use** or **net view**.

**Cause:** The client computer may not be correctly configured, or may not be able to use WINS.

**Solution:** First, determine whether the client is configured or enabled to use both TCP/IP and WINS. Client configuration of WINS-related settings can be done in one of the two following ways:

1. Manually, by an administrator setting the client TCP/IP configuration.
2. Dynamically, by a DHCP server providing the client with TCP/IP configuration.

In most cases, computers running earlier versions of Microsoft operating systems are already able to use WINS once TCP/IP is installed and configured for the client. For Windows XP, administrators can optionally disable NetBIOS over TCP/IP (NetBT) for each client. By disabling NetBT, WINS cannot be used at the client.

In addition, determine whether the client has a valid IP configuration. To check a client IP configuration, use the **ipconfig /all** command. You can use **ipconfig /all|more** as a way to slow or pause viewing of the output of this command for screen-by-screen review.

In the command output, verify that the client computer has the following configured for the network where it is connected:

- a valid IP address
- a valid subnet mask
- a default gateway
- Primary and secondary WINS servers

If the client does not have a valid configuration, you can either use the **ipconfig /renew** command to force the client to obtain a renewal of its IP configuration with the DHCP server or you can update the TCP/IP configuration for the client manually.

**See also:** Configure WINS for non-DHCP clients; Configure WINS for DHCP-enabled clients

**Cause:** The client may not have basic connectivity with its configured WINS servers.

**Solution:** To verify whether a client has basic TCP/IP access to the WINS server, first try pinging the WINS server by its IP address.

For example, if the client uses a primary WINS server of 10.0.0.1, type **ping 10.0.0.1** at the command prompt on the client computer. If you are not sure what the IP address is for the WINS server, you can usually learn it by repeating the previous command, using **ipconfig /all|more** if necessary to pause the displayed command output so you can read the WINS server information.

If the WINS server responds to a direct ping of its IP address, try using the **nbtstat -RR** command at both the client and also at the resource server that the client is attempting to locate by name. This command forces the WINS client services on each computer to send name release and refresh requests to the WINS server and re-register their names.

If the WINS server does not respond to direct pinging of its IP address, this indicates that the source of the problem more likely stems from network connectivity issues between the client and WINS server. If this is the case, follow basic TCP/IP network troubleshooting steps to fix the problem.

**See also:** Verify WINS registration of client NetBIOS names; Troubleshooting WINS replication

**Cause:** The primary or secondary WINS server may not be able to service the client.

**Solution:** At the primary or secondary WINS server for the client, use either Event Viewer or the WINS console to see if WINS is started and currently running. If WINS is running on the server, search for the name previously requested by the client to see if it is in the WINS server database. If the name does not appear in the server database, check whether replication is configured correctly and operational between your WINS servers.

**See also:** View Event Logs; Search for names that begin with a specified set of characters; Troubleshooting WINS replication

**Clients do not use load balancing when connecting to a multihomed server.**
**Cause:** The WINS server providing name resolution for the multihomed server name to network clients is running an earlier version of Windows NT Server that has not been updated.

For all versions of Windows that support NetBIOS over TCP/IP (NetBT), network clients configured to use WINS attempt to resolve any NetBIOS-based name for a computer to its IP address or addresses. To perform this resolution, clients query the WINS server for the address of the multihomed server and receive a list of addresses that are always in the same order. NetBT then takes and sorts the list, arranging multiple IP addresses according to the following three categories:

- Any IP addresses located on the same subnet as the client are given first priority in the list.
- Next, IP addresses located within the same network class are listed.

- Finally, any remaining addresses are listed.

Then, NetBT starts at the top of this list and pings the first address to make sure it is valid. If it gets a reply, it uses the address. If there is no reply, it goes on to the next address.

The problem is that for earlier versions of WINS, within these three categories, the order of addresses is left unchanged from the list provided by the WINS server. This means that, as long as the first address in the list is online, it will be the one that is always used by every WINS client. This does not provide load balancing or 'round-robin' rotation, as DNS servers can perform when more than one name-to-address mapping is available.

**Solution:** Upgrade WINS servers to a Windows Server 2003 operating system; or update servers running earlier versions of Windows 2000 and Windows NT Server to the most current Service Pack

# Troubleshooting WINS servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Troubleshooting WINS servers**
What problem are you having?

- Unable to send and receive data.
- The server intermittently loses its ability to resolve names.

**Unable to send and receive data.**
**Cause:** Configuration details might be incorrect or missing.

**Solution:** To help prevent the most common types of problems, review WINS best practices for deploying and managing your WINS servers. Most WINS-related problems start as failed queries at a client, so it is best to start there.

**See also:** WINS Best Practices; Troubleshooting WINS clients

**Cause:** The WINS server might not be able to service the client.

**Solution:** At the primary or secondary WINS server for the client that cannot locate a name, use **Event Viewer** or the WINS console to see if WINS is started and currently running. If WINS is running on the server, search for the name previously requested by the client to see if it is in the WINS server database.

**See also:** Search for names that begin with a specified set of characters

**Cause:** WINS server is failing or registering database corruption errors.

**Solution:** Use WINS database recovery techniques to help restore WINS operations.

**See also:** Troubleshooting WINS databases

**Cause:** The name does not appear in the server database.

**Solution:** Check whether replication is configured correctly and is operational between your WINS servers.

**See also:** Troubleshooting WINS replication

**Cause:** The WINS server answering clients might be returning stale or incorrect information to clients. The WINS server might be affected by a static mapping blocking the dynamic registration of a valid name-to-address mapping for the computer hosting the service or resource in question.

**Solution:** In general, static mappings are not recommended for clients that can dynamically use WINS to update their name and address information. If the information returned to a client during name resolution is incorrect or stale, check to see if the name entry in the WINS server database was entered statically. If the name was entered using a static type entry, you can choose from the following options to update WINS:

- Enable **Migrate (Overwrite unique static record with dynamic record)** in **Replication Partners Properties**.
- Edit the static mapping to update the mapped address information.
- Delete the static entry from WINS.

**See also:** Enable static mappings to be overwritten during replication; Edit a static mapping entry; Delete or tombstone an entry in the WINS database

**Cause:** Replication might not be occurring between all WINS servers.

**Solution:** In some WINS designs, the use of one-way replication partnerships, such as push-only or pull-only partners, can

create situations where WINS names data is not regularly replicated to all network servers.

**See also:** Troubleshooting WINS replication

**The server intermittently loses its ability to resolve names.**

**Cause:** There might be a split registration problem, where the WINS server is registering its names in WINS at two servers on the network. This is possible when the WINS server settings configured in TCP/IP properties at the server computer are pointing to IP address of remote WINS servers and are not configured to use the IP address of the local WINS server.

**Solution:** Re-configure client TCP/IP properties at the WINS server to have its primary and secondary WINS servers point to the IP address of the local server computer.

**See also:** Configure TCP/IP to use WINS; WINS Best Practices

# Troubleshooting WINS databases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Troubleshooting WINS databases**
What problem are you having?

- WINS backups are failing consistently.
- My WINS database might be corrupt.

**WINS backups are failing consistently.**
**Cause:** The specified WINS backup directory might be a remote drive.

**Solution:** Make sure that the path for the WINS backup directory is on a local disk on the WINS server. WINS cannot back up its database files to a remote drive.

**See also:** Modify backup settings

**My WINS database might be corrupt.**
**Cause:** The WINS database might need to be deleted, either fully or partially, so its contents can be replaced from another WINS server to restore consistency and data integrity between WINS servers on your network.

**Solution:** Depending upon the extent of the corruption, and whether the corrupted records are at an owner or replica WINS server, you can rebuild the local WINS database through replication with another WINS server.

To remove corrupted data at a replica (non-owner) server, use the **Delete Owner** command option in the WINS console. When using this feature, you can delete the owner and any of its associated records in question from local WINS data.

For an owner WINS server, you can restore data integrity by starting WINS with a fresh database and then initiating replication with a configured replication partner. At the local WINS server that has ownership of corrupted records, perform the following steps:

1. Stop WINS.
2. Delete all current WINS database files.

   These files are located at the path currently set in **Database path** on the **Advanced** tab in *WINS server* **Properties** for the local WINS server.
3. Restart WINS.

When WINS is started with a blank database, it will replicate back all records (including those for which it is the registered owner) that are contained in data stored at other servers. This is because it starts with a fresh version ID that is lower than the highest version ID in use for data at any of the other WINS servers on your network.

WINS servers use record version IDs to keep track of which records are new and which need to be replicated to other WINS servers. Every record in a WINS database has an owner associated with it. This owner is generally the server with which the WINS client originally registered.

Normally, the owner (in this case, the local WINS server) does not receive its own records from its replication partners. For this reason, if records are lost from the local WINS database, it can typically be difficult to replace these records, even if they exist at other partners. In this solution, this difficulty is solved because records that exist on other WINS servers are certain to have higher version IDs than the local server after you delete the local database and restart WINS.

On a WINS network with many records, two registry entries affect the ability of WINS to perform local database replacement as described in earlier steps. The first value, InitTimeReplication, keeps the server from accepting queries or registrations until the database is reconstructed. The second, InitTimePause, causes WINS to pause while the replication takes place. The location

and configurable settings for these keys are discussed more in the following tables.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**InitTimeReplication**

**Location:**

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\WINS\Partners\Push (or Pull)

**Type:**

REG_DWORD

**Default:**

1

**Values:**

1 = The WINS server pulls replicas of new WINS database entries from its partners when the system is initialized or when a replication-related parameter changes.

0 = Replication occurs only as often as specified in the value of **TimeInterval**.

**InitTimePause**

**Location:**

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\WINS\Parameters

**Type:**

REG_DWORD

**Default:**

0

**Values:**

0 = WINS does not start in paused state.

1 = WINS starts in paused state.

The value set determines whether WINS starts in a paused state and remains in that state until its first replication is completed. In the paused state, WINS does not accept any name registrations, releases, or queries. WINS remains in the paused state until it has replicated with its partners or until its first replication attempt has failed.

If the value of **InitTimePause** is set to **1**, then **InitTimeReplication** (in the Pull partners subkey) should either be set to **1** or deleted from the registry.

# Troubleshooting WINS replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Troubleshooting WINS replication**
What problem are you having?

- I can't locate the source of "duplicate name" error messages.
- I need to locate the source of "network path not found" error messages on a WINS client.
- The WINS server cannot pull or push replications to another WINS server.
- I am having intermittent problems with WINS not replicating between servers on my network.

**I can't locate the source of "duplicate name" error messages.**
**Cause:** You might need to manually remove static records, or enable static mappings to be overwritten during replication.

**Solution:** If the duplicate name exists already as a static mapping, you can tombstone or delete the entry. If possible, this should be done at the WINS server that owns the static mapping for the duplicate name record. As a preventive measure, in cases where a static mapping needs to be replaced and updated by a dynamic mapping, you can also enable **Overwrite unique static mappings at this server (migrate on)** in **Replication Partners** Properties.

**See also:** Delete or tombstone an entry in the WINS database; Enable static mappings to be overwritten during replication

**I need to locate the source of "network path not found" error messages on a WINS client.**
**Cause:** The network path might contain the name of a server computer configured as a p-node, m-node, or h-node and its IP address is different from the one in the WINS database. In this case, the IP address for this computer might have changed recently and the new address has not yet replicated to the local server.

**Solution:** Check the WINS database for the name and IP address mapping information. If they are not current, you can start replication at the WINS server that owns the name record requiring an update at other servers.

**See also:** Manage Replication

**Cause:** If the computer name is configured as b-node, its name might be missing in the WINS database.

**Solution:** Check the WINS database for the name. If the name is not present in WINS and clients on other subnets need to locate the b-node computer by name, you can add a static mapping in WINS for the computer name and IP address.

**See also:** Manage Static Mappings

**The WINS server cannot pull or push replications to another WINS server.**
**Cause:** If the other WINS server is located across a router, there might be a failed router or other network problem.

**Solution:** Check for network connectivity to the server. Ping the other WINS server and troubleshoot any network or TCP/IP-related problems.

**See also:** Troubleshooting TCP/IP

**Cause:** If the other WINS server is located on the same network or can be successfully pinged, the server might not be configured correctly as either a push or pull partner of the other server.

For example, suppose the two WINS servers are called WINS-A and WINS-B. If WINS-A needs to perform pull replications with WINS-B, make sure it is a push partner of WINS-B. Likewise, if WINS-A should push replications to WINS-B, it should be a pull partner of WINS-B.

**Solution:** Determine the current configuration of replication partners between the servers. If partnerships are missing at either server, add the missing partner IP addresses as needed for the appropriate server. If needed, you can also change the replication partner type, such as from either **Pull only** or **Push only**, to **Push/Pull**.

**See also:** Configure Replication; Change a replication partner type

**I am having intermittent problems with WINS not replicating between servers on my network.**
**Cause:** The replication pattern in use on your network might not be correct or appropriate to use.

**Solution:** In general, deployments of more than 20 WINS servers are strongly discouraged. Also, for best results and simpler administration, follow a hub-and-spoke replication topology when designing your replicated WINS network that utilizes push/pull partnerships between each WINS hub server and its member spoke servers.

If a single hub-and-spoke exceeds the recommended total of 20 WINS servers for a single enterprise, you should contact Microsoft Consulting Services or Microsoft Product Support Services about how to better revise or reduce your current WINS installation. For larger or enterprise installations, multiple hub-and-spoke designs are acceptable to use.

In rare cases, the limited use of push-only and pull-only partner relationships might be needed. You should, however, carefully review added WINS administration issues when these configurations are deployed. At a minimum, establish reliable support procedures for occasions when you might need to manually trigger replication between WINS servers configured to operate using these types of limited replication partnering.

**Cause:** The version ID might not be incrementing for WINS entries when replicating on all servers.

**Solution:** The version ID is incremented in the WINS database by each server that owns and registers a name record. This identifier is a hexadecimal value stored with each name record in the database and is used for version tracking when the record is replicated to other servers.

Version IDs are only incremented for certain types of record changes. For example, a typical refresh of a name in WINS does not increment the version ID, but other changes, such as an updated IP address, can cause the version ID to increment in most cases.

When the version ID is not consistently incrementing for a name record at all servers in the replicated WINS network, you can use either the WINS console or command-line options to administratively increase the starting version count for the server and correct the problem.

**See also:** Increase the starting version count; WINS Best Practices

# Help: Windows Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Firewall**
Windows Firewall is a host firewall technology that replaces Internet Connection Firewall in Windows Server 2003 with Service Pack 1 (SP1) and Windows XP with Service Pack 2 (SP2). Windows Firewall provides stateful inspection of incoming Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) traffic and is designed to protect you from network attacks that pass through your perimeter network or originate inside your organization, such as Trojan horse attacks, port scanning attacks, and worms. You should run Windows Firewall on each of your clients and servers, thereby extending your defense-in-depth strategy to the innermost layer of your security architecture.

- To learn about differences between Windows Firewall and its predecessor, Internet Connection Firewall, see Help: Windows Firewall features.
- For help with specific Windows Firewall tasks, see Help: Windows Firewall How To....
- For general background information, see Help: Windows Firewall Concepts.

# Help: Windows Firewall Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Firewall Concepts**

- Help: Windows Firewall overview
- Help: Windows Firewall features
- Help: Understanding Windows Firewall
- Help: Administering Windows Firewall

# Help: Windows Firewall overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Firewall overview**

To protect your network from the growing number of malicious network attacks, you need to adopt a defense-in-depth strategy. A defense-in-depth strategy incorporates a variety of network security technologies into your security architecture and implements those technologies in a layered structure that extends from your perimeter network (outer layer) to every computer in your organization (inner layer). By providing host firewall protection on the innermost layer of your network, Windows Firewall, a new security component in Windows XP with Service Pack 2 and Windows Server 2003 with Service Pack 1, can be an effective part of your defense-in-depth security strategy.

**What is Windows Firewall?**

Windows Firewall is a host firewall technology, so it runs on each of your clients and servers, providing protection from network attacks that pass through your perimeter network or originate inside your organization, such as Trojan horse attacks, port scanning attacks, and worms. Like many firewall technologies, Windows Firewall is a stateful firewall, so it inspects and filters all TCP/IP version 4 (IPv4) and TCP/IP version 6 (IPv6) traffic. Unsolicited incoming traffic is dropped unless it is a response to a request by the host (solicited traffic) or it is specifically allowed (that is, has been added to the exceptions list). You can specify traffic to be added to the exceptions list according to port number, application name, or service name by configuring Windows Firewall settings. With the exception of some Internet Control Message Protocol (ICMP) messages, Windows Firewall allows all outgoing traffic.

You cannot use Windows Firewall as a perimeter firewall solution and you should not use it as a substitute for other network security technologies. Windows Firewall is designed to be a supplemental security solution; it should be part of a security architecture that implements a variety of security technologies, such as border routers, perimeter firewalls, virtual private networking (VPN), and Internet Protocol security (IPSec).

**How does Windows Firewall work?**

When an incoming packet reaches your computer, Windows Firewall inspects the packet and determines whether it meets criteria specified on the exceptions list. If the packet matches an entry in the list, Windows Firewall passes the packet to the TCP/IP protocol for further processing. If the packet does not match an entry in the list, Windows Firewall silently discards the packet and creates an entry in the Windows Firewall logging file (if the logging file is enabled). The entries in the exceptions list can consist of program names, system service names, TCP ports, and UDP ports. There is no way to create an entry in the exceptions list based on the IP Protocol field in the IP header.

There are only two conditions under which traffic can pass through Windows Firewall:

- When the incoming traffic corresponds to a recent outgoing request, in which case the response traffic is then considered incoming solicited traffic.

  An example of this occurs when a Domain Name System (DNS) Name Query Request message is sent to a DNS server. Windows Firewall adds an entry to the state filter table so that the corresponding DNS Name Query Response message sent by the DNS server can be passed to the TCP/IP protocol for further processing.

- When you configure Windows Firewall for exceptions, in which case you allow a computer using Windows Firewall to accept unsolicited incoming traffic when acting as a server, a listener, or a peer.

  An example of this occurs when your computer is acting as a Web server. In this case, you must configure Windows Firewall to allow Web traffic so that the computer can respond to requests from Web clients. You can configure exceptions based on programs, in which case the ports opened by the program are added to the exceptions list automatically, or on TCP or UDP ports, in which case the ports are opened whether the application or services using them is active or not.

**Common Windows Firewall scenarios**

With a few exceptions, Windows Firewall can be enabled on all configurations of Windows XP with SP2 or Windows Server 2003 with SP1. Therefore, it is recommended that you enable Windows Firewall on every client and server in your organization, including bastion hosts and other servers in your perimeter network, mobile and remote clients that connect to your network, and all clients and servers in your internal network. In addition, Windows Firewall can be enabled for any network architecture. Therefore, it is also recommended that you enable Windows Firewall regardless of how you have

designed and implemented your perimeter network or your internal network.

There are a few circumstances that might require you to disable Windows Firewall.

**Servers that are running Routing and Remote Access**
You should not enable Windows Firewall on a server that is running Routing and Remote Access. Routing and Remote Access has its own firewall, which makes Windows Firewall redundant and unnecessary, and Routing and Remote Access might not run if you have Windows Firewall enabled.

**Servers that are running perimeter firewalls**
You should not run Windows Firewall on a server that is running a perimeter firewall, such as Microsoft Internet Security and Acceleration (ISA) Server 2004. The protection that Windows Firewall provides in this situation is redundant and unnecessary. In addition, Windows Firewall can cause perimeter firewalls such as ISA Server to function improperly.

**Computers that are running non-Microsoft host firewalls**
You should disable Windows Firewall on a computer that is running a non-Microsoft host firewall. Although it is possible to run two host firewalls on a single computer, it is not recommended that you do so. Host firewall implementations vary widely and there is no guarantee that a non-Microsoft host firewall and Windows Firewall will work well together. Also, using multiple host firewalls on a single computer does not necessarily decrease your attack surface because the configuration settings (traffic rules) for each firewall would likely be the same and therefore redundant. Also, the use of multiple host firewalls on a single computer increases your operational overhead without any substantial gain in security.

In addition, you might want to disable Windows Firewall if a client or server requires you to open numerous ports or allow a large number of applications and services to receive unsolicited traffic. Because a significant volume of network traffic will be allowed to pass through Windows Firewall anyway, by disabling Windows Firewall, you eliminate the operational overhead associated with Windows Firewall configuration and maintenance. You also avoid any performance impact related to Windows Firewall. However, you should closely evaluate the design of any client or server that requires you to open numerous ports. Clients and servers that are configured for numerous roles or are configured to provide numerous services can be a critical point of failure in your organization and usually indicate poor infrastructure design.

**Windows Firewall impacts**
When it is enabled in its default configuration, Windows Firewall blocks all unsolicited incoming network traffic on all network connections. While blocking unsolicited incoming traffic reduces your attack surface and increases your level of security, it can cause some applications and services to stop working properly. For this reason, you might need to configure Windows Firewall so that unsolicited incoming traffic is allowed on certain ports or for certain applications and services. Configuration changes are usually not required on clients that are running Windows XP because the applications and services that run on a client do not often need to receive unsolicited incoming traffic. However, configuration changes are usually required on servers that are running Windows Server 2003 because many of the applications and services that run on servers, by their nature, need to receive and process unsolicited incoming traffic.

**Impacts on Windows XP**
Windows Firewall affects only the applications and services that listen for and respond to unsolicited incoming network traffic. There are not many applications and services that do this on clients because clients are typically configured as service consumers, not service providers. In other words, most of the network traffic that passes through a client computer is initiated by the client in the form of an outgoing request to another computer. Therefore, on clients, you usually do not need to configure Windows Firewall settings. There are, however, some scenarios for which you might have to create Windows Firewall exceptions:

**File and Printer Sharing**
If you want to share files, folders, or printers so that other computers can access these resources on a client computer, you need to configure Windows Firewall on the client computer so that unsolicited incoming traffic is allowed on TCP ports 139 and 445, and UDP ports 137 and 138.

**Remote Assistance**
If you want administrators and help desk technicians to use Remote Assistance to administer clients, you need to configure Windows Firewall on your clients so that the Remote Assistance application is allowed to process unsolicited incoming traffic.

**Remote Desktop**
If you want administrators and help desk technicians to use Remote Desktop to administer clients, you need to configure Windows Firewall on your clients so that unsolicited incoming traffic is allowed on TCP port 3389.

**UPnP**
If your clients need to process UPnP requests from other computers, you need to configure Windows Firewall on your clients so that unsolicited incoming traffic is allowed on TCP port 2869 and UDP port 1900.

In addition, you might have to configure Windows Firewall settings on clients that run certain types of e-mail and antivirus

programs. E-mail programs that connect to an Exchange server typically receive unsolicited incoming traffic when the Exchange server sends a message to the client informing the client that new mail has arrived. Centrally-managed antivirus programs typically receive unsolicited traffic when new virus signatures are pushed down to clients from a server.

**Impacts on Windows Server 2003**

Unlike clients, servers typically run numerous applications and services that listen for and respond to unsolicited incoming network traffic. Therefore, on servers, you usually need to configure some Windows Firewall settings so that unsolicited traffic is not blocked by Windows Firewall. The configuration changes that you need to make to Windows Firewall vary according to the server role and the applications and services that are running on the server. The following are some common server roles that require you to configure Windows Firewall settings:

| Server role | Configuration required |
|---|---|
| DHCP server | Allow unsolicited traffic through TCP port 67. |
| DNS server | Allow unsolicited traffic through TCP port 53. |
| POP3 server | Allow unsolicited traffic through TCP port 110. If you are securing POP3 traffic with Secure Sockets Layer (SSL), allow unsolicited traffic through TCP port 995. |
| Telnet server | Allow unsolicited traffic through TCP port 23. |

In addition to configuring Windows Firewall for various server roles, applications, and services, you need to configure Windows Firewall settings on every server that you intend to manage remotely. If you are managing the server by using the Administration Tools Pack on a client running Windows XP, then you need to configure Windows Firewall settings on the server to ensure that Windows Firewall allows the unsolicited traffic from the remote computer to reach the applications and services that are running on the server.

## Notes

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**See Also**

**Concepts**

Help: Windows Firewall How To...
Help: Understanding Windows Firewall
Help: Administering Windows Firewall

# Help: Windows Firewall features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Firewall features**
Windows Firewall has several features and functional differences that make it different from its predecessor, Internet Connection Firewall (ICF). Some of the features and functional differences between Windows Firewall and Internet Connection Firewall include:

- Global configuration options that apply to all connections
- New Windows Firewall component in Control Panel
- New Don't Allow Exceptions operating mode
- Startup security
- Incoming traffic scoping for Internet Protocol version 4 (IPv4)
- Exceptions that can be specified by program or system service name
- Built-in support for Internet Protocol version 6 (IPv6)
- New configuration options

**Global configuration options that apply to all connections**
Windows Firewall allows you to configure settings that apply to all the connections of the computer. In Windows Server 2003 with no service packs installed, ICF settings are configured per connection, which means that if you want to enable ICF on multiple connections and configure exceptions, you must configure each connection separately. When you change a global Windows Firewall setting, the change is applied to all the connections on which Windows Firewall is enabled. Windows Firewall does allow per-connection configuration; however, connection-specific configuration overrides global configuration.

**New Windows Firewall component of Control Panel**
To enable or disable ICF, you select or clear the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box, which appears on the **Advanced** tab of a connection's properties dialog box. You can also click the **Settings** button on a connection's properties dialog box and configure general exceptions, logging settings, and ICMP exceptions.

In Windows Firewall, the check box on the **Advanced** tab of a connection's properties dialog box has been replaced with a **Settings** button, which launches the new Windows Firewall component in Control Panel. You can also start Windows Firewall by clicking **Windows Firewall** in Control Panel. The Windows Firewall dialog box allows you to configure general settings, exceptions for programs (applications and services), connection-specific settings, log settings, and ICMP exceptions. You can also reset Windows Firewall to its default configuration.

**New operating mode**
ICF is either enabled (that is, allows only solicited traffic and traffic that is specified in the exceptions list) or disabled (allows all traffic). Windows Firewall has these same operating modes, but allows you to select a new operating mode that enables Windows Firewall and ignores all exceptions. When Windows Firewall is running in this new mode, all unsolicited incoming traffic is dropped, including traffic that is specified in the exceptions list. This new mode can be used when connecting to the Internet from a public location, such as a hotel or airport, or inside an organization's network to temporarily lock down computers during a network attack or when a malicious program is spreading. After the network attack is over and appropriate updates have been installed to prevent future attacks, Windows Firewall can be returned to the normal operating mode that allows exceptions. All of the original settings for exceptions are maintained.

**Startup security**
ICF is active only on the connections on which it is enabled when the Internet Connection Firewall/Internet Connection Sharing (ICS) service starts successfully. Therefore, there is a delay between when the computer is active on the network and when the connections are protected with ICF. This delay makes the computer accessible to attacks during startup.

In Windows Firewall, there is a startup policy that performs stateful packet filtering, thereby allowing the computer to perform basic networking startup tasks using Dynamic Host Configuration Protocol (DHCP) and the Domain Name System (DNS) protocol to configure the computer and communicate with a domain controller to obtain Group Policy updates. The startup policy does not include any settings that you have configured for Windows Firewall, including settings that you have configured in Group Policy. After the Windows Firewall/ICS service is started, Windows Firewall uses the settings that you configured for Windows Firewall and removes the startup policy. The startup policy settings cannot be configured.

**Incoming traffic scoping**
In ICF, allowed traffic can originate from any IPv4 address. In Windows Firewall, you can specify where allowed traffic originates. This configuration setting is known as the scope option and it allows you to limit the scope of an exception to traffic that originates from any IPv4 address; originates from an IPv4 address that can be reached directly (based on IPv4 routing table entries); originates from one or more specific IPv4 addresses or from one or more ranges of IPv4 addresses.

For incoming IPv6 traffic, Windows Firewall allows you to specify whether exceptions allow traffic that originates from any IPv6 address or from an IPv6 address that can be reached directly (based on IPv6 routing table entries).

**Exceptions can be specified by program name**
In ICF, you manually configure exceptions by specifying the set of TCP and UDP ports that correspond to the traffic of a specific program, which can be either an application or service. This can make configuration difficult for users who do not know the set of TCP and UDP ports for the application or service. Also, this configuration does not work for applications that do not listen on a specific set of UDP or TCP ports.

To make it easier to configure exceptions, Windows Firewall allows you to specify the set of TCP or UDP ports or the file name of the program. When the program runs, Windows Firewall monitors the ports on which the program listens and opens the ports.

To allow you to quickly enable exceptions for commonly allowed, incoming unsolicited traffic, Windows Firewall has predefined exceptions, such as File and Printer Sharing and Remote Desktop. In addition, the notification mechanism in Windows Firewall prompts local administrators with a Windows Security Alert message to allow them to add new programs to the program exceptions list automatically.

**Built-in support for IPv6**
IPv6 support is included with Windows Firewall and is automatically enabled on all IPv6 connections. Both IPv4 and IPv6 share the same settings for traffic that is on the exceptions list. For example, if you allow file and print sharing traffic, then both IPv4- and IPv6-based unsolicited incoming file and printer sharing traffic is allowed.

### Note

Windows Firewall does not support the same scope options for IPv6 traffic as IPv4 traffic.

**New configuration options**
The only way to enable or disable ICF is through the Network Connections folder, the Network Setup Wizard, or the Internet Connection Wizard. To configure exceptions, you must either use the Network Connections folder or your application must be ICF-aware, in which case it adds exceptions to the exceptions list automatically.

In Windows Firewall, there are several ways to configure settings and options, including the following:

**Netsh commands**
Netsh is a command-line tool that you can use to configure settings for network components. However, to configure a component, the component must support a set of commands through a **netsh** context. Windows Firewall provides the **firewall** context, which can be used to configure Windows Firewall settings through a suite of commands. Using **netsh firewall**, you can create scripts to automatically configure a set of Windows Firewall settings for both IPv4 and IPv6 traffic. You can also use **netsh firewall** commands to display the configuration and status of the Windows Firewall.

**Group Policy settings**
Windows Firewall provides several new Group Policy settings so you can centrally configure and manage large numbers of computers in an organization that uses the Active Directory directory service. These new Group Policy settings allow you to configure Windows Firewall operational modes, exceptions, and other settings.

**Profile-based settings**
Windows Firewall allows you to configure settings in two different profiles: a domain profile and a standard profile. The domain profile is the set of Windows Firewall settings that are needed when the computer is connected to the network that contains the domain controllers of the organization. For example, the domain profile might contain exceptions for the applications needed by a managed computer in an enterprise network. The standard profile is the set of Windows Firewall settings that are needed when the computer is not connected to the network that contains the domain controllers of the organization (for example, when an employee uses the organization's laptop on the road and connects to the Internet using a

public broadband or wireless Internet service provider). Because the organization's laptop is directly connected to the Internet, the standard profile should contain more restrictive settings than the domain profile.

### Notes

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Help: Understanding Windows Firewall
Help: Administering Windows Firewall
Help: Windows Firewall How To...

# Help: Understanding Windows Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Windows Firewall**

- Help: Understanding Windows Firewall notifications
- Help: Understanding Windows Firewall exceptions
- Help: Understanding Windows Firewall profiles
- Help: Understanding Windows Firewall scope options

# Help: Understanding Windows Firewall notifications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Windows Firewall notifications**

By default, Windows Firewall displays a **Windows Security Alert** dialog box (referred to as a notification) when a program attempts to listen for unsolicited incoming traffic. If you are a member of the Administrators group on the computer, the notification gives you the ability to add the program to the exceptions list. If you are not a member of the Administrators group on the computer, the notification informs you that a program attempted to listen for incoming traffic but was blocked. You are not given the ability to add the program to the exceptions list if you are not a member of the Administrators group on the computer.

Windows Firewall does not display a notification in the user interface (UI) when a system service attempts to listen for incoming traffic on a port. This is also true for any program that runs like a system service (that is, a program that runs under an account that has higher privilege than a user account, for example, the Local System account, or a program that runs while there is no user logged onto the computer). However, you can use the security event log to determine whether a system service, or a program that runs like a system service, attempts to listen for incoming traffic. To do this you must enable the **Audit process tracking** policy and the **Audit policy change** policy. When you do this, Windows Firewall will write a Failure Audit with Event ID 861 to the security event log any time a program or system service attempts to listen for incoming traffic.

You can disable the notification feature in Windows Firewall by using the **Exceptions** tab in Windows Firewall in Control Panel, the **netsh firewall set notifications** command, or the **Windows Firewall: Prohibit notifications** policy setting. This might be necessary on servers that are remotely managed.

**See Also**

**Concepts**

Help: Windows Firewall overview
Help: Administering Windows Firewall
Help: Enable or disable Windows Firewall notifications

# Help: Understanding Windows Firewall exceptions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Windows Firewall exceptions**

When you first turn on Windows Firewall or you restore Windows Firewall default settings, all unsolicited incoming traffic is blocked on all network connections. This means that any program or system service that attempts to listen for traffic on a TCP or UDP port will be unable to receive network traffic. To allow programs and system services to receive unsolicited traffic through these ports, you need to add the program or system service to the Windows Firewall exceptions list, or you need to determine which port or ports the program or system service uses and add the port or ports to the Windows Firewall exceptions list. Adding programs, system services, and ports to the exceptions list is the most commonly used way to control which traffic is allowed to pass through Windows Firewall.

**Note**

You can also control which traffic is allowed to pass through Windows Firewall by configuring the **Windows Firewall: Allow authenticated IPSec bypass** Group Policy setting. However, this is not a commonly used method of controlling the traffic that passes through Windows Firewall.

**Configuring exceptions**

You can add exceptions to the Windows Firewall exceptions list on a global basis or on a per-connection basis. Global exceptions are applied to all of the network connections on a computer, including new connections that you create. The recommended method for configuring global Windows Firewall exceptions is to use the Security Configuration Wizard (SCW).You can also add global exceptions for programs, system services, and ports to the Windows Firewall exceptions list by using the **Exceptions** tab in Windows Firewall in Control Panel. You can also add global exceptions for programs, system services, and ports by using the **netsh firewall** commands and the Windows Firewall Group Policy settings. Per-connection exceptions are applied to specific network connections. You can add per-connection exceptions for system services and ports by configuring **Network Connections Settings** on the **Advanced tab** in Windows Firewall in Control Panel and by using the **netsh firewall set portopening** and **netsh firewall set icmpsettings** commands. You cannot configure per-connection exceptions by using the Windows Firewall Group Policy settings.

You can also configure exceptions by enabling or disabling the preconfigured Windows Firewall exceptions by using the **Exceptions** tab in Windows Firewall in Control Panel, the **netsh firewall set** commands, or the Windows Firewall Group Policy settings. The following table lists the preconfigured exceptions, all of which are disabled by default:

| Exception | Description |
|---|---|
| File and Printer Sharing | Opens TCP ports 139 and 445 and UDP ports 137 and 138. Allows a computer to receive unsolicited traffic to shared files, folders, and printers. |
| Remote Desktop | Opens TCP port 3389. Allows a computer to be managed remotely with the Remote Desktop Connection feature. |
| UPnP Framework | Opens TCP port 2869 and UDP port 1900. Allows a computer to receive UPnP discovery requests from other computers and devices. |
| Remote administration | Opens TCP ports 135 and 445. Allows Svchost.exe and Lsass.exe to receive unsolicited incoming traffic and allows hosted services to open additional dynamically-assigned ports, typically in the range of 1024 to 1034. Allows a computer to be remotely managed with administrative tools such as the Microsoft Management Console (MMC) and Windows Management Instrumentation (WMI). |
| ICMP | Includes a suite of Internet Control Messages Protocol (ICMP) exceptions. ICMP exceptions allow you to control how a computer responds to and sends ICMP messages, including messages that are used by the **ping** command. |

**Note**

You can programmatically configure Windows Firewall exceptions by using the Windows Firewall application programming interface (API).

**Mitigating the risks associated with exceptions**

Each time you add a program, system service, or port to the exceptions list, you make your computer more accessible to attack. A common network attack uses port scanning software to identify computers that have open and unprotected ports. By adding numerous programs, system services, and ports to the exceptions list, you defeat the purpose of a firewall and increase the attack surface of your computer. This problem typically occurs when you configure a server for several different roles, and you need to open numerous ports to accommodate each of the server roles. You should closely evaluate the design of any server that requires you to open numerous ports. Servers that are configured for numerous roles or configured to provide numerous services can be a critical point of failure in your organization and usually indicate poor infrastructure design.

To help decrease your security risk, follow these guidelines:

**Create an exception only when you need it**
If you think a program or system service might require a port for unsolicited incoming traffic, do not add the program or system service to the exceptions list until you verify that the program or system service attempted to listen for unsolicited traffic. By default, Windows Firewall displays a notification when a program attempts to listen for unsolicited traffic. You can also use the security event log to determine whether a system service has attempted to listen for unsolicited traffic.

**Never allow an exception for a program that you do not recognize**
If Windows Firewall notifies you that a program has attempted to listen for unsolicited traffic, verify the name of the program and the executable (.exe) file before you add the program to the exceptions list. Likewise, if you use the security event log to identify system services that have attempted to listen for unsolicited traffic, verify that the service is a legitimate system service before you add a port to the exceptions list for the system service.

**Remove an exception when you no longer need it**
If you add a program, system service, or port to the exceptions list on a server, and then change the server's role or reconfigure the services and applications on the server, be sure to update the exceptions list and remove all of the exceptions that are no longer necessary.

**Best practices for creating exceptions**
In addition to the general guidelines for managing exceptions, use the following best practices when adding a program, system service, or port to the exceptions list.

**Adding programs**
Always try to add a program (.exe file) or a system service that runs within an .exe file to the exceptions list before you try to add a port. When you add a program to the exceptions list, Windows Firewall dynamically opens the required ports for the program. When the program is running, Windows Firewall allows incoming traffic through the required ports; when the program is not running, Windows Firewall blocks any incoming traffic that is sent to the ports.

**Adding system services**
Do not add system services to the exceptions list if the system service runs within Svchost.exe. Adding Svchost.exe to the exceptions list allows any system service running inside every instance of Svchost.exe to receive unsolicited incoming traffic. You should add system services to the exceptions list only if the system service runs within an .exe file or you can enable a preconfigured Windows Firewall system service exception, such as the UPnP Framework exception or the File and Printer Sharing exception.

**Adding ports**
You should add a port to the exceptions list as a last resort. When you add a port to the exceptions list, Windows Firewall allows incoming traffic through the port, regardless whether there is a program or system service listening on the port for incoming traffic.

**See Also**
**Concepts**
Help: Windows Firewall overview
Help: Administering Windows Firewall
Help: Configure Exceptions and Notifications
Help: Configure Windows Firewall Auditing and Logging
Security Configuration Wizard

# Help: Understanding Windows Firewall profiles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Windows Firewall profiles**

You can configure Windows Firewall settings in two different profiles: a domain profile and a standard profile. The domain profile settings are used when a computer is connected to a network in which the computer's domain account resides. The standard profile settings are used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network like the Internet. Both the domain profile and standard profile contain the same set of Windows Firewall settings. However, the standard profile settings are typically more restrictive than the domain profile because the standard profile settings do not need to include applications and services that are used in a managed domain environment only.

When you configure Windows Firewall in Control Panel, you are configuring settings in the profile that is currently in use. You cannot use Windows Firewall in Control Panel to configure settings in the profile that is not currently in use. You can tell which profile is currently in use by looking at the bottom of the **General** tab. You can also use the **netsh firewall show currentprofile** command to determine which profile is currently in use.

When you configure Windows Firewall Group Policy settings, you can configure individual policy settings in either the domain profile or the standard profile. Likewise, when you use the **netsh firewall** commands to configure Windows Firewall settings, you can choose to configure settings in either the domain profile or the standard profile.

If you do not configure standard profile settings, their default values are applied whenever Windows Firewall uses the standard profile. Therefore, it is highly recommended that you configure both domain and standard profile settings and that you enable the Windows Firewall for both profiles.

**See Also**

**Concepts**

Help: Windows Firewall overview
Help: Administering Windows Firewall
Help: Determine which profile Windows Firewall is using

# Help: Understanding Windows Firewall scope options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Windows Firewall scope options**
When you configure a program or port exception, you can also configure the scope options for the exception. Scope options are used to define which computers are allowed to send traffic through an exception. You can access the scope options by pressing the **Change Scope** button in the **Add a Port**, **Add a Program**, **Edit a Port**, or **Edit a Program** dialog boxes. There are three scope options:

**Any computer (including those on the Internet)**
This setting allows traffic from any Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) address through the exception. This setting might make your computer more accessible to malicious users or programs on the Internet.

**My network (subnet) only**
This setting allows traffic only from IPv4 or IPv6 addresses that can be reached directly by your computer. Windows Firewall determines whether the source IPv4 or IPv6 address of the incoming packet can be reached directly by querying the IPv4 and IPv6 routing tables. You can see all destinations that are considered directly reachable by typing the **route print** command at a command prompt. For the IPv4 routing table, all IPv4 addresses that match the routes in which the IPv4 address of the **Gateway** column equals the IPv4 address of the **Interface** column are considered directly reachable. For the IPv6 routing table, all IPv6 addresses that match routes in which the **Gateway** column is set to **On-link** are considered directly reachable. Therefore, the set of directly reachable addresses depends on your networking configuration, as specified by the IPv4 and IPv6 configuration of LAN-based connections (such as Ethernet and 802.11 wireless), dial-up connections, and broadband Internet connections. In some Internet configurations, all destinations are considered directly reachable.

For example, for a computer that is only directly connected to a private home network, the set of directly reachable unicast addresses is confined to those that match the IPv4 network ID of the private subnet. If the network connection is configured with an IPv4 address of 192.168.0.99 with a subnet mask of 255.255.0.0, only traffic from IPv4 addresses in the range 192.168.0.0 to 192.168.255.255 is allowed.

In another example, for a computer that is directly connected to both a private home network and the Internet through a cable modem, the set of directly reachable unicast addresses are those that match either the network ID of the private subnet or the cable modem provider subnet. For example, if the private network connection is configured with an IPv4 address of 192.168.0.1 and a subnet mask of 255.255.0.0 and the cable modem connection is configured with an IPv4 address of 131.107.17.211 and a subnet mask of 255.255.255.0, traffic received by either network connection is allowed from IPv4 addresses in the ranges from 192.168.0.0 to 192.168.255.255 and from 131.107.17.0 to 131.107.17.255.

**Custom list**
This setting allows you to specify one or more IPv4 addresses or IPv4 address ranges separated by commas. IPv4 address ranges typically correspond to subnets. For IPv4 addresses, type the IPv4 address in dotted decimal notation. For IPv4 address ranges, you can specify the range using a dotted decimal subnet mask or a prefix length. When you use a dotted decimal subnet mask, you can specify the range as an IPv4 network ID (such as 10.47.81.0/255.255.255.0) or by using an IPv4 address within the range (such as 10.47.81.231/255.255.255.0). When you use a network prefix length, you can specify the range as an IPv4 network ID (such as 10.47.81.0/24) or by using an IPv4 address within the range (such as 10.47.81.231/24). The following is an example custom list: 10.91.12.56,10.7.14.9/255.255.255.0,10.116.45.0/255.255.255.0,172.16.31.11/24,172.16.111.0/24

You cannot specify a custom list for IPv6 traffic.

### Note

When you configure and enable an exception, you are instructing Windows Firewall to allow specific unsolicited incoming traffic sent from the specified scope (from any address, from an address that can be reached directly, or from a custom list). For any scope, enabling an exception makes the computer accessible to attacks based on incoming unsolicited traffic from computers that are assigned the allowed addresses and from malicious computers that spoof traffic. There is no way to prevent spoofed attacks from the Internet on connections assigned public IPv4 addresses except by disabling the exception. Therefore, you should try to configure scope options so that the number of computers that are allowed to send unsolicited traffic through an exception is kept to a minimum. This will reduce, but not eliminate, the likelihood of a spoof attack.

**See Also**
**Concepts**
Help: Windows Firewall overview
Help: Administering Windows Firewall

# Help: Administering Windows Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Windows Firewall**

- Help: Administering Windows Firewall with Group Policy
- Help: Administering Windows Firewall with Netsh
- Help: Administering Windows Firewall through Control Panel

# Help: Administering Windows Firewall with Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Windows Firewall with Group Policy**
If your organization uses the Active Directory directory service, you can configure and manage Windows Firewall with the new Windows Firewall Group Policy settings. When you use Group Policy to configure Windows Firewall, local administrators and users will be unable to locally configure some Windows Firewall configuration settings by using Windows Firewall in Control Panel. Specifically, when a Windows Firewall setting is configured through Group Policy, the equivalent setting appears dimmed in Windows Firewall in Control Panel.

**Windows Firewall Group Policy Settings**
There are two sets of Windows Firewall policy settings:

- The domain profile settings, which are used when a computer is connected to a network that contains the organization's domain controllers.
- The standard profile settings, which are used when a computer is connected to a network that does not contain the organization's domain controllers.

If you do not configure profile settings, their default values are applied. Therefore, it is highly recommended that you configure both domain and standard profile settings so that you maintain your desired level of security in the event your computer changes from one profile to another. Also, the standard profile settings are typically more restrictive that the domain profile settings because the standard profile settings do not need to include applications and services that are used only in a managed domain environment. Both the domain profile and standard profile contain the same set of Windows Firewall configuration settings.

The following Group Policy settings are available for managing Windows Firewall. These settings can be found at the following location in the Group Policy Editor snap-in:

- Computer Configuration\Administrative Templates\Network\Network Connections\Windows Firewall\Domain Profile or Standard Profile

| Group Policy setting | Description |
|---|---|
| **Windows Firewall: Protect all network connections** | Used to specify that all network connections have Windows Firewall enabled. |
| **Windows Firewall: Do not allow exceptions** | Used to specify that all unsolicited incoming traffic is dropped, including traffic that has been added to the exceptions list. |
| **Windows Firewall: Define program exceptions** | Used to define by application file names traffic that has been added to the exceptions list. |
| **Windows Firewall: Allow local program exceptions** | Used to enable local configuration of program exceptions. |
| **Windows Firewall: Allow remote administration exception** | Used to enable remote procedure calls (RPC) and Distributed Component Object Model (DCOM), which are necessary for many remote administration using tools such as Microsoft Management Console (MMC) and Windows Management Instrumentation (WMI). |
| **Windows Firewall: Allow file and printer sharing exception** | Used to specify whether file and printer sharing traffic is allowed. |
| **Windows Firewall: Allow ICMP exceptions** | Used to specify the types of unsolicited Internet Control Message Protocol (ICMP) traffic allowed. |

| | |
|---|---|
| **Windows Firewall: Allow Remote Desktop exception** | Used to specify whether the computer can accept a Remote Desktop-based connection request. |
| **Windows Firewall: Allow UPnP framework exception** | Used to specify whether the computer can participate in UPnP discovery. |
| **Windows Firewall: Prohibit notifications** | Used to disable notifications when an application uses new Windows Firewall application programming interfaces (APIs) to request traffic that has been added to the exceptions list. |
| **Windows Firewall: Allow logging** | Used to enable logging of discarded traffic, successful connections, and configure log file settings. |
| **Windows Firewall: Prohibit unicast response to multicast or broadcast requests** | Used to discard the unicast packets sent in response to a multicast or broadcast request. |
| **Windows Firewall: Define port exceptions** | Used to specify by TCP and UDP ports traffic that has been added to the exceptions list. |
| **Windows Firewall: Allow local port exceptions** | Used to enable local configuration of port exceptions. |

You can also configure the **Windows Firewall: Allow authenticated IPSec bypass police setting**, which can be found at the following location in the Group Policy Editor snap-in:

Computer Configuration\Administrative Templates\Network\Network Connections\Windows Firewall

This policy setting allows unsolicited incoming messages from specified systems that authenticate using IPSec.

**Notes**

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- You cannot use Group Policy to configure Windows Firewall per-connection settings.
- Group Policy settings must be refreshed before they take effect.

**See Also**
**Concepts**
Help: Understanding Windows Firewall
Help: Administering Windows Firewall
Help: Windows Firewall How To...

# Help: Administering Windows Firewall with Netsh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Windows Firewall with Netsh**

Windows Firewall provides a **firewall** context, which you can use to configure Windows Firewall settings. To run Windows Firewall **netsh** commands, you must start **netsh** from the command prompt and then change to the **firewall** context by typing **firewall**. You can also run **netsh firewall** commands all on one line at the command prompt.

You cannot use the **netsh firewall** commands to remotely configure Windows Firewall settings, and you cannot use the **dump** command to create a script based on the current Windows Firewall configuration.

The following **netsh firewall** commands can be used in scripts or at the command line to configure Windows Firewall.

| Command | Description |
|---|---|
| **add allowedprogram** | Used to add a program-based exception. |
| **set allowedprogram** | Used to modify the settings of an existing program-based exception. |
| **delete allowedprogram** | Used to delete an existing program-based exception. |
| **set icmpsetting** | Used to specify Internet Control Message Protocol (ICMP) traffic that has been added to the exceptions list. |
| **set multicastbroadcastresponse** | Used to specify the unicast response to a multicast or broadcast request behavior. |
| **set notifications** | Used to specify the notification behavior. |
| **set logging** | Used to specify logging options. |
| **set opmode** | Used to specify the operating mode of Windows Firewall either globally or for a specific connection (interface). |
| **add portopening** | Used to create a port-based exception. |
| **set portopening** | Used to modify the settings of an existing port-based exception. |
| **delete portopening** | Used to delete an existing port-based exception. |
| **set service** | Used to enable or disable the predefined file and printer sharing, remote administration, Remote Desktop, and UPnP exceptions. |
| **show allowedprogram** | Displays the programs that have been added to the exceptions list. |
| **show config** | Displays the local configuration information. |
| **show currentprofile** | Displays the current profile. |
| **show icmpsetting** | Displays the ICMP settings. |
| **show logging** | Displays the logging settings. |
| **show multicastbroadcastresponse** | Displays multicast or broadcast response settings. |
| **show notifications** | Displays the current settings for notifications. |
| **show opmode** | Displays the operational mode. |
| **show portopening** | Displays the ports that have been added to the exceptions list. |

| | |
|---|---|
| **show service** | Displays the services. |
| **show state** | Displays the current state information. |

**Note**

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Netsh overview
Help: Understanding Windows Firewall
Help: Administering Windows Firewall through Control Panel
Help: Administering Windows Firewall with Group Policy
Help: Windows Firewall How To...

# Help: Administering Windows Firewall through Control Panel

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Windows Firewall through Control Panel**

The most common way to configure Windows Firewall settings on a single computer is to use Windows Firewall in Control Panel. You must be a member of the Administrators group on the local computer to configure settings in Windows Firewall in Control Panel. If you are not a member of the Administrators group, all of the settings in Windows Firewall will appear dimmed.

The Windows Firewall user interface consists of three tabs: the **General** tab, the **Exceptions** tab, and the **Advanced** tab.

**General tab**

You can disable and enable Windows Firewall for all connections on the **General** tab. The settings can be configured as follows:

| Setting | Description |
|---|---|
| **On** | Enables Windows Firewall for all of the network connections that are selected on the **Advanced** tab. Windows Firewall is enabled to allow only solicited traffic and incoming traffic that has been added to the exceptions list. |
| **Don't allow exceptions** | Allows solicited incoming traffic only. Incoming traffic that has been added to the exceptions list is not allowed. The settings on the **Exceptions** tab are ignored and all of the network connections are protected, regardless of the settings on the **Advanced** tab. In addition, Windows Firewall does not display a notification when a program attempts to listen for unsolicited incoming traffic. |
| **Off** | Disables Windows Firewall. This is not recommended, especially for network connections that are directly accessible from the Internet, unless you are already using a non-Microsoft host firewall. |

On Windows Server 2003, the default setting for Windows Firewall is **Off** for all connections and all newly created connections. If you turn Windows Firewall on, Windows Firewall can affect the communications of programs or services that rely on unsolicited incoming traffic. In this case, you must identify those programs that are no longer working and add them to the exceptions list or determine which ports the programs use and add the ports to the exceptions list.

**Exceptions tab**

You can configure the following settings on the **Exceptions** tab:

| Setting | UI element | Description |
|---|---|---|
| Program and port exceptions | **Add Program** button | Used to display the **Add a Program** dialog box, which allows you to select a program from a list or browse for a program's executable (.exe) file. |
| | **Add Port** button | Used to display the **Add a Port** dialog box, which allows you to specify a name for the port exception, a port number, and a port type (TCP or UDP). |
| | **Edit** button | Used to display the **Edit a Port** or **Edit a Program** dialog box, which allows you to modify the exception settings for the port or program.<br><br>All of the programs or services enabled from the **Exceptions** tab are enabled for all of the connections that are selected on the **Advanced** tab.<br><br>For more information about program exceptions, see Help: Understanding Windows Firewall exceptions. |

| | | |
|---|---|---|
| Program and port scope options | **Change Scope** button | Available in the **Add a Port**, **Add a Program**, **Edit a Port**, or **Edit a Program** dialog boxes. There are three scope options:<br>• Any computer (including those on the Internet)<br>• My network (subnet) only<br>• Custom list<br><br>For more information about scope options, see Help: Understanding Windows Firewall scope options. |
| Windows Firewall notifications | **Display a notification when Windows Firewall blocks a program** check box | Select or clear to configure the way Windows Firewall handles notifications.<br><br>If selected, Windows Firewall displays a notification when a program that is not in the exceptions list attempts to listen for incoming traffic.<br><br>If cleared, notifications do not appear.<br><br>For more information about notifications, see Help: Understanding Windows Firewall notifications. |

**Advanced tab**

You can configure the following settings on the **Advanced** tab:

| Setting | Description |
|---|---|
| Network Connection Settings | Used to configure the following settings:<br>• Specify individual connection on which Windows Firewall is enabled. To enable Windows Firewall on a connection, select the check box next to the connection name. To disable Windows Firewall on a connection, clear the check box. By default, all of the network connections have Windows Firewall enabled. If a network connection does not appear in **Network Connection Settings**, then it is a non-standard networking connection, such as a custom dialer from an Internet service provider (ISP).<br><br>If you clear all of the check boxes in **Network Connection Settings**, then Windows Firewall is not protecting your computer even if, on the **General** tab, you have selected **On**. Also, the settings in **Network Connection Settings** are ignored if you have selected **Don't allow exceptions** on the **General** tab, in which case all interfaces are protected.<br><br>• Configure advanced settings for an individual network connection. Click the network connection name, and then clicking **Settings**. The **Advanced Settings** dialog box is displayed.<br><br>On the **Advanced Settings** dialog box, from the **Services** tab, you can configure specific services (by TCP or UDP port only). From the **ICMP** tab, you can enable specific types of ICMP traffic. |
| Security Logging | Configure Windows Firewall logging options by clicking **Settings** in **Security Logging**. When you do this, the **Log Settings** dialog box is displayed, which allows you to configure whether to log discarded (dropped) packets or successful connections and specify the name, location, and maximum size of the log file. By default, the log file is named pfirewall.log and it is saved in the systemroot folder. The default maximum size is 4096 kilobytes (KB). |
| ICMP Settings | Configure Internet Control Message Protocol (ICMP) exceptions by clicking **Settings** in **ICMP**. When you do this, the **ICMP** dialog box is displayed, which allows you to enable and disable the types of incoming ICMP messages that Windows Firewall allows for all the connections selected on the **Advanced** tab. ICMP messages are used for diagnostics, reporting error conditions, and configuration. By default, no ICMP messages in the list are allowed. |

| | |
|---|---|
| Restore Defaults | Restore Windows Firewall default settings. All of the entries in the exceptions list are deleted and all settings and options are restored to their original state. Clicking **Restore Defaults** also enables Windows Firewall on all connections, which can cause your programs and system services to behave improperly because this is not the standard. This might also cause Internet Connection Sharing (ICS) and Network Bridge to fail. |

**Note**

- If a setting is managed by Group Policy, or you do not have the administrative rights to configure a setting, the setting will appear dimmed when you open Windows Firewall in Control Panel.

**See Also**
**Concepts**
Help: Understanding Windows Firewall
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Windows Firewall How To...
Help: Understanding Windows Firewall scope options

# Help: Windows Firewall How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Firewall How To...**

- Help: Start, Stop, and Reset Windows Firewall
- Help: Configure Exceptions and Notifications
- Help: Configure Windows Firewall Auditing and Logging
- Help: Configure Windows Firewall for IPSec Bypass and Multicast Traffic
- Help: Secure Windows Firewall

# Help: Start, Stop, and Reset Windows Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Start, Stop, and Reset Windows Firewall**

- Help: Use SCW to start and configure Windows Firewall
- Help: Turn Windows Firewall on or off
- Help: Turn Windows Firewall on or off for a specific connection
- Help: Turn Windows Firewall on with no exceptions
- Help: Restore Windows Firewall default settings

# Help: Use SCW to start and configure Windows Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use SCW to start and configure Windows Firewall**

1. Open the Security Configuration Wizard (SCW), and follow the steps in the wizard until you reach the **Network Security** page.

2. Clear the **Skip this section** check box, and click **Next**.

3. Under **Select the ports to open**, review the ports that SCW will add to the Windows Firewall exceptions list. To prevent SCW from adding a port to the exceptions list, clear the check box that is next to the port.

   If you want to add additional ports or applications to the Windows Firewall exceptions list, click **Add**, and do the following:

   - To add a port to the exceptions list, in **Port number**, enter the port number, select **TCP** or **UDP** or both **TCP** and **UDP**, and click **OK**.

   - To add an application to the exceptions list, click the **Approve Application** tab, and in **Application path**, enter the application path and name of the executable (.exe) file, and then click **OK**. Try to use environment variables for paths so that the administrative templates can be applied to any computer.

4. Click **Next**.

5. On the **Confirm Port Configuration** page, verify the port settings that SCW will use to configure Windows Firewall, and click **Next**.

6. On the **Registry Settings** page, select the **Skip this section** check box, and click **Next**.

7. On the **Audit Policy** page, select the **Skip this section** check box, and click **Next**.

8. On the **Internet Information Services** page, select the **Skip this section** check box, and click **Next**.

9. On the **Save Security Policy** page, click **Next**.

10. Enter a name and description for the security policy, and click **Next**.

11. Click **Apply now** to apply the security policy and configure Windows Firewall, and click **Next**.

12. When SCW has applied the security policy, click **Next**, and then click **Finish**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- SCW is an optional component in Windows Server 2003 with Service Pack 1 (SP1). SCW is not installed by default. You must install it before you can use it.

- SCW displays different pages depending on the configuration settings you choose.

- If you choose **Remote access/VPN server**, **Internet Connection Sharing server**, or **Microsoft Internet Security and Acceleration Server 2004** during the Role-Based Service Configuration portion of SCW, you will not be able to select Network Security portion of SCW, which is where you configure Windows Firewall settings.

**See Also**
**Concepts**
Help: Understanding Windows Firewall exceptions
Security Configuration Wizard

# Help: Turn Windows Firewall on or off

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To turn Windows Firewall on or off**

- Open Windows Firewall.

  To turn Windows Firewall on, click **On**, and then click **OK**.

  To turn Windows Firewall off, click **Off**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.
- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.
- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- If you start Windows Firewall after you start applications that are listed in the Windows Firewall exceptions list, you might have to restart your computer. Windows Firewall cannot track the state of a program's traffic if the program was started before Windows Firewall/Internet Connection Sharing service was started. This problem typically occurs when you start a program, and then try to start Windows Firewall for the first time.

**See Also**
**Concepts**
Help: Windows Firewall Concepts
Help: Turn Windows Firewall on with no exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using

# Help: Turn Windows Firewall on or off for a specific connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To turn Windows Firewall on or off for a specific connection**

1. Open Windows Firewall.

2. Click the **Advanced** tab.

   To turn Windows Firewall off for a connection, in **Network Connection Settings**, clear the check box next to the connection that you do not want to secure with Windows Firewall, and then click **OK**.

   To turn Windows Firewall on for a connection, in **Network Connection Settings**, select the check box next to the connection that you want to secure with Windows Firewall, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Help: Understanding Windows Firewall scope options
Help: Administering Windows Firewall with Netsh
Help: Determine which profile Windows Firewall is using

# Help: Turn Windows Firewall on with no exceptions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To turn Windows Firewall on or off with no exceptions**

1. Open Windows Firewall.
2. Click **On**, select the **Don't allow exceptions** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.
- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.
- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using

# Help: Restore Windows Firewall default settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore Windows Firewall default settings**

1. Open Windows Firewall.
2. Click the **Advanced** tab, click **Restore Defaults**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.
- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Restoring the default settings will undo all of the changes you have made to Windows Firewall. Restoring the default settings will also turn Windows Firewall on for all network connections, which might cause some programs to stop working. In addition, restoring the default settings might cause Internet Connection Sharing and Network Bridge to not work correctly.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Help: Windows Firewall Concepts
Help: Administering Windows Firewall with Netsh

# Help: Configure Exceptions and Notifications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Exceptions and Notifications**

- Help: Add a port to the Windows Firewall exceptions list
- Help: Edit or delete a port exception
- Help: Add a program to the Windows Firewall exceptions list
- Help: Edit or delete a program exception
- Help: Add a system service to the Windows Firewall exceptions list
- Help: Edit or delete a system service exception
- Help: Enable or disable the Remote Desktop exception
- Help: Enable or disable the File and Printer Sharing exception
- Help: Enable or disable the UPnP Framework exception
- Help: Enable or disable the remote administration exception
- Help: Change the scope of a Windows Firewall exception
- Help: Configure ICMP exceptions
- Help: Enable or disable Windows Firewall notifications
- Help: Determine which profile Windows Firewall is using

# Help: Add a port to the Windows Firewall exceptions list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a port to the Windows Firewall exceptions list**

1. Open Windows Firewall.
2. Click the **Exceptions** tab, and then click **Add Port**.
3. In **Name**, type a name for the port exception. This name will appear in the exceptions list.
4. In **Port number**, type the number of the port that you want to add to the exceptions list.
5. Click either **TCP** or **UDP**, depending on the type of traffic the port handles, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.
- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.
- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Adding a port to the exceptions list can lessen the security of your computer because the port will be open any time the computer is running. You should add ports to the exceptions list only when it is not possible to add a program to the exceptions list.
- You should configure scope options for any exceptions that you enable.

**See Also**
**Concepts**
Help: Understanding Windows Firewall exceptions

Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Edit or delete a port exception

Updated: November 6, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit or delete a port exception**

1. Open Windows Firewall.

2. Click the **Exceptions** tab, and then click the port exception that you want to edit or delete.

   If you want to delete the port exception, click **Delete**, and then click **OK**.

   If you want to edit the port exception, click **Edit**, make changes to the port exception settings, and then click **OK**.

If you want to remove all your program and port exceptions, you can restore your firewall to the default settings. Restoring the default settings deletes all the Windows Firewall settings that you made since Windows was installed. This may cause some programs to stop working. For more information, see Help: Restore Windows Firewall default settings.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- Adding a port to the exceptions list can lessen the security of your computer because the port will be open any time the computer is running. You should add ports to the exceptions list only when it is not possible to add a program to the exceptions list.

- You should configure scope options for any exceptions that you enable.

**See Also**
**Concepts**

Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Add a program to the Windows Firewall exceptions list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a program to the Windows Firewall exceptions list**

1. Open Windows Firewall.

2. Click the **Exceptions** tab, and then click **Add program**.

   If the program that you want to add to the exceptions list appears in **Programs**, click the program, and then click **OK**.

   If the program that you want to add to the exceptions list does not appear in **Programs**, click **Browse**, and then do the following:

   - In the **Browse** dialog box, find the executable (.exe) file or command (.com) file for the program that you want to add to the exceptions list, click **Open**, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- You should configure scope options for any exceptions that you enable.

**See Also**
**Concepts**
Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy

Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Edit or delete a program exception

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit or delete a program exception**

1. Open Windows Firewall.

2. Click the **Exceptions** tab, and then click the program exception that you want to edit or delete.

   If you want to delete the program exception, click **Delete**, and then click **Yes**.

   If you want to edit the program exception, click **Edit**, make the changes you want to the program exception settings, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- You should configure scope options for any exceptions that you enable.

**See Also**
**Concepts**
Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Add a system service to the Windows Firewall exceptions list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a system service to the Windows Firewall exceptions list**

1. Open Windows Firewall, and then click the **Advanced** tab.

2. In **Network Connection Settings**, click the connection that the system service uses to communicate, and click **Settings**.

   If the system service that you want to add to the exceptions list appears in **Services**, select the check box for the system service, and click **OK**.

   If the system service that you want to add to the exceptions list does not appear in **Services**, click **Add**, and in the **Service Settings** dialog box, do the following:

   a. In **Description of service**, type a display name for the system service. The display name appears in the exceptions list.

   b. In **Name or IP address of the computer hosting this service on your network**, type the name or IP address of the local computer.

   c. In **External port number for this service**, type the port number that receives unsolicited traffic for this system service. Usually, this is the same as the IP address of the local computer.

   d. In **Internal port number for this service**, type the port number to which the local computer should forward (or translate) incoming traffic.

   e. Click either **TCP** or **UDP**, depending on the type of traffic your internal and external ports handle, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- You should configure scope options for any exceptions that you enable.

**See Also**
**Concepts**
Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Edit or delete a system service exception

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit or delete a system service exception**

1. Open Windows Firewall, and click the **Advanced** tab.

2. In **Network Connection Settings**, click the connection that the system service uses to communicate, and then click **Settings**.

3. Click the system service that you want to edit or delete.

   If you want to delete the system service exception, click **Delete**, and then click **OK**.

   If you want to edit the system service exception, click **Edit**, make changes to the system service exception settings, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- You should configure scope options for any exceptions that you enable.

- You cannot change the name of a system service exception by clicking **Edit**. To rename a system service exception, delete the exception and create a new exception with the new name.

**See Also**
**Concepts**

Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Enable or disable the Remote Desktop exception

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable the Remote Desktop exception**

1. Open Windows Firewall.

2. Click the **Exceptions** tab.

   To enable the Remote Desktop exception, in **Programs and Services**, select the **Remote Desktop** check box, and click **OK**.

   To disable the Remote Desktop exception, in **Programs and Services**, clear the **Remote Desktop** check box, and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- The Remote Desktop exception allows unsolicited incoming traffic through TCP port 3389.

- Disabling the Remote Desktop exception can prevent other computers from using Remote Desktop Connection to connect to your computer.

- You should configure scope options for any exceptions that you enable.

**See Also**
**Concepts**

Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Enable or disable the File and Printer Sharing exception

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable the File and Printer Sharing exception**

1. Open Windows Firewall.

2. Click the **Exceptions** tab.

   To enable the File and Printer Sharing exception, in **Programs and Services**, select the **File and Printer Sharing** check box, and click **OK**.

   To disable the File and Printer Sharing exception, in **Programs and Services**, clear the **File and Printer Sharing** check box, and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- The File and Printer Sharing exception allows unsolicited incoming traffic through TCP ports 139 and 445, and UDP ports 137 and 138.

- Enabling the File and Printer Sharing exception allows incoming Internet Control Message Protocol (ICMP) echo traffic, which enables computers to ping your computer. To block incoming ICMP echo traffic, you must edit the File and Printer Sharing exception so that TCP port 445 is disabled.

- Editing or disabling the File and Printer Sharing exception can prevent other computers from accessing your shared folders, files, and printers. Be sure to test for the impact these changes might have on file and printer sharing functionality.

- You should configure scope options for any exceptions that you enable.

- Enabling the File and Printer Sharing exception will allow other computers to reach your computer with the **ping** command.

**See Also**
**Concepts**
Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Enable or disable the UPnP Framework exception

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable the UPnP Framework exception**

1. Open Windows Firewall.

2. Click the **Exceptions** tab.

   To enable the UPnP Framework exception, in **Programs and Services**, select the **UPnP Framework** check box, and click **OK**.

   To disable the UPnP Framework exception, in **Programs and Services**, clear the **UPnP Framework** check box, and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- The UPnP Framework exception allows unsolicited incoming traffic through TCP port 2869 and UDP port 1900.

- You might not want to disable or edit the UPnP Framework exception if your computer uses Internet Connection Sharing to provide other computers access to the Internet. Editing or disabling the UPnP Framework exception can prevent computers and devices from detecting and using the shared Internet connection. Be sure to test for the impact these changes might have on Internet Connection Sharing functionality.

- You should configure scope options for any exceptions that you enable.

**See Also**

**Concepts**

Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Enable or disable the remote administration exception

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable the remote administration exception**

- To enable the remote administration exception, type the following at the command prompt, and press ENTER:

  **netsh firewall set service remoteadmin enable**

- To disable the remote administration exception, type the following at the command prompt, and press ENTER:

  **netsh firewall set service remoteadmin disable**

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- You cannot use Windows Firewall in Control Panel to configure the remote administration exception.
- The remote administration exception allows unsolicited incoming traffic through TCP ports 135 and 445 and through ports that are dynamically assigned by the RPC endpoint mapper. The remote administration exception also allows Svchost.exe and Lsass.exe to receive unsolicited incoming traffic.
- You should enable the remote administration exception only if your remote administrative tools require remote procedure calls (RPC) and Distributed Component Object Model (DCOM). Malicious users often attempt to attack networks and computers using RPC and DCOM. It is recommended that you contact the manufacturer of your remote administration tool to determine if it requires RPC and DCOM communication. If it does not, do not enable the remote administration exception.
- You should configure scope options for any exceptions that you enable.
- Enabling the remote administration exception will allow other computers to reach your computer with the **ping** command.

**See Also**
**Concepts**
Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Change the scope of a Windows Firewall exception

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the scope of a Windows Firewall exception**

1. Open Windows Firewall, and click the **Exceptions** tab.

2. In **Programs and Services**, click the exception that you want to configure, and click **Edit**.

3. Click **Change scope**, and do one of the following:

   If you want to enable the exception for all computers, click **Any computer (including those on the Internet)**, and then click **OK**.

   If you want to enable the exception for computers that can be reached directly from your local subnet, click **My network (subnet) only**, and then click **OK**.

   If you want to enable the exception for specific IP addresses or IP address ranges, click **Custom list**, enter the specific IP addresses or IP address ranges, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- The **My network (subnet) only** and **Custom list** scope options do not necessarily increase your security. Malicious users can circumvent these scope restrictions by spoofing an IP address that appears to be directly reachable or is part of the custom list.

- When you use the **My network (subnet) only** scope option, Windows Firewall does not perform any IP configuration

testing or subnet mask interpretation to determine whether the traffic is coming from a computer that is a neighbor on a locally attached subnet. Instead, Windows Firewall determines whether the traffic is originating from an address that can be reached directly based on routes in the IP version 4 (IPv4) routing table. If it is, this traffic is considered to match the **My network (subnet) only** scope.

**See Also**

**Concepts**

Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using

# Help: Configure ICMP exceptions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure ICMP exceptions**

1. Open Windows Firewall, and click the **Advanced** tab.
2. In **ICMP**, click **Settings**.
3. In the **ICMP Settings** dialog box, select the check box for each ICMP setting that the computer should respond to, or clear the check box for each ICMP setting that the computer should not respond to, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.
- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.
- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- You can also configure ICMP exceptions on a per-connection basis.
- If you clear the **Allow incoming echo request** check box to prevent your computer from responding to ping traffic, you might also have to configure the remote administration exception and the File and Printer Sharing exception so that incoming traffic is not allowed on TCP port 445. These exceptions open TCP port 445, which allows incoming echo request traffic.
- You can enable other computers to reach your computer with the **ping** command by selecting the **Allow incoming echo request** check box.
- You should configure scope options for any exceptions that you enable.

**See Also**

**Concepts**

Help: Understanding Windows Firewall exceptions
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using
Help: Understanding Windows Firewall scope options

# Help: Enable or disable Windows Firewall notifications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable Windows Firewall notifications**

1. Open Windows Firewall.

2. Click the **Exceptions** tab.

   To enable notifications, select the **Display a notification when Windows Firewall blocks a program** check box, and click **OK**.

   To disable notifications, clear the **Display a notification when Windows Firewall blocks a program** check box, and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- Windows Firewall cannot display notifications if you select the **Don't allow exceptions** check box in Windows Firewall in Control Panel or if you enable the **Do not allow exceptions** policy setting in the Group Policy Object Editor. There are no Windows Firewall settings that allow you to override this behavior.

- Windows Firewall does not display notifications for programs that rely on the Winsock driver to dynamically bind to a UDP port. If a program uses this method (sometimes referred to as wildcard binds) to bind to a UDP port, you might be able to use the **netstat** command and other troubleshooting tools to determine which UDP port is being used.

- Windows Firewall does not display notifications for system services or programs that run when there is no user logged on to the computer. Most system services run this way, as do many server programs. However, Windows Firewall does

write an event (Event ID 861) to the security event log for any system service or program that attempts to listen on a port, no matter if a notification is displayed.

**See Also**
**Concepts**
Help: Understanding Windows Firewall notifications
Help: Enable auditing of Windows Firewall events
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using

# Help: Determine which profile Windows Firewall is using

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To determine which profile Windows Firewall is using**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall show currentprofile**

## Notes

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Windows Firewall settings can be configured in two profiles: the domain profile, which is used when a computer is connected to a network in which the computer's domain account resides; and the standard profile, which is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network.

**See Also**
**Concepts**
Help: Administering Windows Firewall with Netsh

# Help: Configure Windows Firewall Auditing and Logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Windows Firewall Auditing and Logging**

- Help: Enable auditing of Windows Firewall events
- Help: Change the name and location of the Windows Firewall log file
- Help: Change the size of the Windows Firewall log file
- Help: Enable or disable the Windows Firewall log
- Help: View the Windows Firewall log
- Help: View Windows Firewall events in Event Viewer

# Help: Enable auditing of Windows Firewall events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable auditing of Windows Firewall events**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Group Policy settings in your organization.

2. Open **Windows Settings**, open **Security Settings**, open **Local Policies**, and then click **Audit Policy**.

3. Double-click **Audit process tracking**, select the **Success** and **Failure** check boxes, and then click **OK**.

4. Double-click **Audit policy change**, select the **Success** and **Failure** check boxes, and then click **OK**.

## Notes

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- You can determine when a system service listens on a port by enabling auditing and viewing security event log entries.

- You must enable auditing for Windows Firewall events to appear in the security event log.

**See Also**
**Concepts**
Help: Windows Firewall overview
Help: View Windows Firewall events in Event Viewer

# Help: Change the name and location of the Windows Firewall log file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the name and location of the Windows Firewall log file**

1. Open Windows Firewall, and click the **Advanced** tab.
2. In **Security Logging**, click **Settings**.
3. Click **Save As**, and navigate to the location where you want to save the log file.
4. In **File name**, type the new log file name, or leave it blank to accept the default name of pfirewall.log.
5. Click **Save**, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.
- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Help: Windows Firewall Concepts
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy

# Help: Change the size of the Windows Firewall log file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the size of the Windows Firewall log file**

1. Open Windows Firewall, and click the **Advanced** tab.
2. In **Security Logging**, click **Settings**.
3. In **Size limit (KB)**, enter the maximum size of the log file, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.
- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- The default size of the Windows Firewall log file is 4,096 kilobytes (KB) and the maximum size is 32,767 KB.
- When the firewall log reaches the size that you set in **Size limit (KB)**, the log file is renamed and a new log file is created.

**See Also**
**Concepts**
Help: Windows Firewall Concepts
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy

# Help: Enable or disable the Windows Firewall log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable the Windows Firewall log file**

1. Open Windows Firewall, and click the **Advanced** tab.

2. In **Security Logging**, click **Settings**.

   To enable the log file, select the **Log dropped packets** check box or the **Log successful connections** check box or both check boxes, and click **OK**.

   To disable the log file, clear the **Log dropped packets** check box and the **Log successful connections** check box, and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.

- You can also use the **netsh** command with the **firewall** context to perform this procedure and configure other Windows Firewall settings.

- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.

- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.

- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Help: Windows Firewall Concepts
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy

# Help: View the Windows Firewall log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the Windows Firewall log**

1. Open Windows Firewall, and click the **Advanced** tab.
2. In **Security Logging**, click **Settings**.
3. In the **Log Settings** dialog box, click **Save As**.
4. Right-click the log file name, and then click **Open**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To start Windows Firewall, click **Start**, point to **Control Panel**, and then click **Windows Firewall**.
- If a Windows Firewall setting appears dimmed in the graphical user interface, and on the **General** tab, you see **For your security, some settings are controlled by Group Policy**, the setting might be managed by Group Policy. If all Windows Firewall settings appear dimmed, and on the **General** tab, you see **You must be a computer administrator to change these settings**, you do not have administrative rights to configure Windows Firewall.
- On Windows Server 2003, Windows Firewall is turned off by default and the Windows Firewall/Internet Connection Sharing service is disabled by default. You might have to start the Windows Firewall/Internet Connection Sharing service if you try to perform this procedure and you have never started Windows Firewall.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- By default, the Windows Firewall log file is named pfirewall.log and is saved in the systemroot folder.
- You must enable the Windows Firewall log file for it to appear in the systemroot folder.

**See Also**
**Concepts**
Help: Windows Firewall Concepts

# Help: View Windows Firewall events in Event Viewer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view Windows Firewall events in Event Viewer**

1. Open Event Viewer.
2. Click **Security**, and then look for event IDs 848 through 861.

## Notes

- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- You must enable the **Audit process tracking** and **Audit policy change** policy settings to view Windows Firewall events in the security event log.

**See Also**
**Concepts**
Help: Windows Firewall Concepts
Help: Enable auditing of Windows Firewall events

# Help: Configure Windows Firewall for IPSec Bypass and Multicast Traffic

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Windows Firewall for IPSec and Multicast Traffic**

- Help: Allow IPSec traffic to bypass Windows Firewall
- Help: Block unicast responses to a computer's outgoing multicast traffic

# Help: Allow IPSec traffic to bypass Windows Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow IPSec traffic to bypass Windows Firewall**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.

2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, and then click **Windows Firewall**.

3. In the details pane, double-click **Windows Firewall: Allow authenticated IPSec bypass**.

4. In the **Windows Firewall: Allow authenticated IPSec bypass properties** dialog box, on the **Settings** tab, click **Enabled**.

5. In **Define IPSec peers to be exempted from firewall policy**, type the Security Descriptor Definition Language (SDDL) string that corresponds to the group accounts for the computers to which this policy applies, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- This procedure allows unsolicited incoming messages from specified systems that authenticate using IPSec.

- You cannot use Windows Firewall in Control Panel or the **netsh firewall** command to perform this procedure.

- The format of the SDDL string for a single group is:

  O:DAG:DAD:(A;;RCGW;;;SID)

  Where SID is the Security Identifier (SID) of a group account.

- You can use the Getsid.exe tool to obtain the SID of a group account. Getsid.exe is typically used to compare the SIDs of two accounts on different domain controllers, but you can also use it to obtain the SID of a specified user or group account. For more information about Getsid.exe, see the Windows Server 2003 Resource Kit Tools Web site.

- If you have more than one group, then the syntax for the SDDL string is:

  O:DAG:DAD:(A;;RCGW;;;SID1) (A;;RCGW;;;SID2) (A;;RCGW;;;SID3)...

- If you enable the **Windows Firewall: Allow authenticated IPSec bypass** policy setting, and then later disable the policy setting or change it to **Not Configured**, Windows Firewall deletes the SDDL string.

**See Also**
**Concepts**
Help: Administering Windows Firewall with Group Policy

# Help: Block unicast responses to a computer's outgoing multicast traffic

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To block unicast responses to a computer's outgoing multicast traffic**

- Type the following at the command prompt, and press ENTER:

  **netsh firewall set multicastbroadcastresponse disable**

## Notes

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- You can also use Group Policy settings to perform this procedure and configure other Windows Firewall settings.
- You can configure Windows Firewall settings in the standard profile or the domain profile. The domain profile is used when a computer is connected to a network in which the computer's domain account resides. The standard profile is used when a computer is connected to a network in which the computer's domain account does not reside, such as a public network or the Internet. Make sure Windows Firewall is using the correct profile when you perform this procedure.
- You cannot use Windows Firewall in Control Panel to perform this procedure.
- If you perform this procedure, and this computer sends multicast or broadcast messages to other computers, Windows Firewall blocks the unicast responses sent by those other computers.
- By default, when this computer sends a multicast or broadcast message to other computers, Windows Firewall waits as long as three seconds for unicast responses from the other computers and then blocks all later responses.

**See Also**
**Concepts**
Help: Administering Windows Firewall with Netsh
Help: Administering Windows Firewall with Group Policy
Help: Determine which profile Windows Firewall is using

# Help: Secure Windows Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Secure Windows Firewall**

- Help: Prevent administrators from configuring local port exceptions
- Help: Prevent administrators from configuring local program exceptions
- Help: Prevent administrators from running Windows Firewall

# Help: Prevent administrators from configuring local port exceptions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prevent administrators or users from configuring local port exceptions**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.

2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then open either **Domain Profile** or **Standard Profile**, depending on which profile you want to configure.

3. In the details pane, double-click **Windows Firewall: Allow local port exceptions**.

4. In the **Windows Firewall: Allow local port exceptions properties** dialog box, on the **Settings** tab, click **Disabled**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- You cannot use Windows Firewall in Control Panel or the **netsh firewall** command to perform this procedure.

- Group Policy settings must be refreshed before they take effect.

**See Also**
**Concepts**
Help: Administering Windows Firewall with Group Policy

# Help: Prevent administrators from configuring local program exceptions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prevent administrators or users from configuring local program exceptions**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.

2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then open either **Domain Profile** or **Standard Profile**, depending on which profile you want to configure.

3. In the details pane, double-click **Windows Firewall: Allow local program exceptions**.

4. In the **Windows Firewall: Allow local program exceptions properties** dialog box, on the **Settings** tab, click **Disabled**, and then click **OK**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- You cannot use Windows Firewall in Control Panel or the **netsh firewall** command to perform this procedure.

- Group Policy settings must be refreshed before they take effect.

**See Also**
**Concepts**
Help: Administering Windows Firewall with Group Policy

# Help: Prevent administrators from running Windows Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prevent administrators or users from running Windows Firewall**

1. Open the Group Policy Object Editor snap-in to edit the Group Policy object (GPO) that is used to manage Windows Firewall settings in your organization.

2. Open **Computer Configuration**, open **Administrative Templates**, open **Network**, open **Network Connections**, open **Windows Firewall**, and then open either **Domain Profile** or **Standard Profile**, depending on which profile you want to configure.

3. In the details pane, double-click **Windows Firewall: Protect all network connections**.

4. In the **Windows Firewall: Protect all network connections properties** dialog box, on the **Settings** tab, click **Disabled**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- You cannot use Windows Firewall in Control Panel or the **netsh firewall** command to perform this procedure.

- Group Policy settings must be refreshed before they take effect.

**See Also**
**Concepts**
Help: Administering Windows Firewall with Group Policy

# Managing Remote Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing remote connections**
This section covers:

- Routing and Remote Access
- Connection Manager Administration Kit
- Connection Point Services
- Virtual Private Networks
- Internet Authentication Service

# Routing and Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing and Remote Access**

The Routing and Remote Access service in the Microsoft® Windows Server 2003 family provides:

- Multiprotocol LAN-to-LAN, LAN-to-WAN, virtual private network (VPN), and network address translation (NAT) routing services. For more information, see Routing.

- Dial-up and VPN remote access services. For more information, see Remote Access.

# Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing**

Routing and Remote Access provides multiprotocol LAN-to-LAN, LAN-to-WAN, virtual private network (VPN), and network address translation (NAT) routing services. Routing and Remote Access is intended for use by system administrators who are already familiar with routing protocols and services, and routable protocols such as TCP/IP and AppleTalk.

- Before configuring the server running Routing and Remote Access to act as a router, see Checklist: Installing and configuring the router.
- To find features that have been moved or changed in the Microsoft® Windows Server 2003 family, see New ways to do familiar Routing tasks.
- For tips about using routing, see Routing best practices.
- For help with specific tasks, see Routing How To....
- For general background information, see Routing Concepts.
- For problem-solving instructions, see Troubleshooting routing.

# Checklist: Installing and configuring the router

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring the router**

| Step | Reference |
|---|---|
| Review key concepts. | Routing Overview |
| Verify the compatibility of all hardware to be installed in the server running Routing and Remote Access. | Support resources |
| Install the hardware. | Manufacturer's documentation |
| For a new network, decide which protocols and routing features of Routing and Remote Access you want to use. For an existing network, determine which protocols and routing features of Routing and Remote Access you need to enable. | Routing Overview and documentation about your network |
| Install and configure the protocols. | Network communications |
| Enable and configure the Routing and Remote Access service. | Enable the Routing and Remote Access service |
| (Optional) Design and deploy static IP routing. | Setting Up a Static Routed IP Internetwork |
| (Optional) Design and deploy the RIP-for-IP routing protocol. | Setting Up a RIP-for-IP Routed Internetwork |
| (Optional) Design and deploy the OSPF routing protocol. This feature is not available on the Itanium-based versions of the Windows operating systems. | Setting Up an OSPF Routed Internetwork |
| (Optional) Add and configure the DHCP Relay Agent. | Configure the DHCP Relay Agent |
| (Optional) Configure IP multicast support. | Configure IP Multicast Support |
| (Optional) Design and deploy network address translation. | Setting Up Network Address Translation |
| (Optional) Configure the appropriate IP packet filters. | Manage Packet Filters |
| (Optional) Design and deploy demand-dial routing. | Setting Up Demand-Dial Routing |

# New ways to do familiar Routing tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring the routing features of Routing and Remote Access. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0 or Windows NT version 4.0 with the Routing and Remote Access Service (RRAS). The user interface for the Windows 2000 version of Routing and Remote Access differs slightly from that for the Windows Server 2003 family, but the steps for most common tasks have not changed.

Please choose your upgrade path:

- Windows NT 4.0 to a product in the Windows Server 2003 family
- Windows NT 4.0 with RRAS to a product in the Windows Server 2003 family

**Windows NT 4.0 to a product in the Windows Server 2003 family**

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Install, configure, and remove routing protocols | **Services** tab of **Network** in Control Panel | Open Routing and Remote Access. For more information, see Manage Routing Protocols and Routing How To.... |
| View the IP routing table | The **route print** command at a Windows NT command prompt. This displays the IP Forwarding table, the actual table that is used to forward IP traffic. | Open Routing and Remote Access. The routing table displayed in Routing and Remote Access displays the table based on updates from routing protocols, dial-on demand connections, and static routes. The table maintained in Routing and Remote Access is used to update the IP Forwarding table. You can view the IP Forwarding table directly by using the **route print** command at the command prompt. For more information, see View routing tables. |

**Windows NT 4.0 with RRAS to a product in the Windows Server 2003 family**

| If you want to | In Windows NT 4.0 with RRAS use | In this version of Windows use |
|---|---|---|
| Install, configure, and remove routing protocols | Routing and RAS Admin | Open Routing and Remote Access. For more information, see Manage Routing Protocols and Routing How To.... |
| View the IP routing table | Routing and RAS Admin | Open Routing and Remote Access. For more information, see View routing tables. |

# Routing best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**
For best practice information about design, security, and deployment for the server running Routing and Remote Access, see the following:

- Security information for routing
- Setting Up a Static Routed IP Internetwork
- Setting Up a RIP-for-IP Routed Internetwork
- Setting Up an OSPF Routed Internetwork
- Setting Up Network Address Translation
- Setting Up Demand-Dial Routing
- Deploying Router-to-Router VPNs

**Note**

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Routing How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

This section provides instructions for installing and configuring the routing components of Routing and Remote Access.

- Enable the Routing and Remote Access service
- Enable LAN and WAN routing
- Manage Routing Interfaces
- Manage Devices and Ports
- Manage Routing Protocols
- Manage Static Routes
- Manage Routers
- Manage Packet Filters
- Configure RIP for IP
- Configure OSPF
- Configure NAT/Basic Firewall
- Enable ICMP router discovery
- Configure the DHCP Relay Agent
- Configure IP Multicast Support
- Enable AppleTalk routing
- Connect Your Internal Network to the Internet
- Gather Troubleshooting Information

**Note**

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Enable the Routing and Remote Access service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable the Routing and Remote Access service**

1. If this server is a member of an Active Directory domain and you are not a domain administrator, instruct your domain administrator to add the computer account of this server to the **RAS and IAS Servers** security group in the domain of which this server is a member. The domain administrator can add the computer account to the **RAS and IAS Servers** security group by using Active Directory Users and Computers or with the **netsh ras add registeredserver** command. If this server is using local authentication or is authenticating against a RADIUS server, skip this step.

2. Open Routing and Remote Access.

3. By default, the local computer is listed as a server.

   To add another server, in the console tree, right-click **Server Status**, and then click **Add Server**.

   In the **Add Server** dialog box, click the applicable option, and then click **OK**.

4. In the console tree, right-click the server you want to enable, and then click **Configure and Enable Routing and Remote Access**.

5. Follow the instructions in the Routing and Remote Access Server Setup Wizard.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Installing the router
Installing the remote access server
Netsh commands for remote access

# Enable LAN and WAN routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable LAN and WAN routing**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to enable routing, and then click **Properties**.
3. On the **General** tab, select the **Router** check box, and then do one of the following:
   - For LAN routing, click **Local area network (LAN) routing only**.
   - For LAN and WAN routing, click **LAN and demand-dial routing**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Common routing problems
Routing Overview

# Manage Routing Interfaces

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage routing interfaces**


- Add a routing interface
- Add a demand-dial interface
- Configure demand-dial filters
- Configure dial-out hours
- Configure the calling router for certificate-based EAP
- Configure the answering router for certificate-based EAP
- Configure a pre-shared key for a demand-dial routing interface
- Delete an Interface

# Add a routing interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a routing interface**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. Right-click **General**, and then click **New Interface**.
4. In **Interfaces**, click the interface you want to add, and then click **OK**.
5. If applicable, complete any configuration dialog boxes for the interface.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Routing interfaces, devices, and ports

# Add a demand-dial interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a demand-dial interface**

1. Open Routing and Remote Access.
2. In the console tree, click **Network Interfaces**.

   **Where?**

   - Routing and Remote Access/*server name*/Network Interfaces

3. Right-click **Network Interfaces**, and then click **New Demand-dial Interface**.
4. Follow the instructions in the Demand-Dial Interface wizard.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Demand-dial routing
Understanding Demand-Dial Routing
Setting Up Demand-Dial Routing

# Configure demand-dial filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure demand-dial filters**

1. Open Routing and Remote Access.
2. In the console tree, click **Network Interfaces**.

   **Where?**

   - Routing and Remote Access/*server name*/Network Interfaces

3. In the details pane, right-click the demand-dial interface you want to configure, and then click **Set IP Demand-dial filters**.
4. In the **Set Demand-Dial Filters** dialog box, do one or more of the following:

   - To add a demand-dial filter, click **New**, fill in the required information for the filter, and then click **OK**.
   - To modify the settings for a demand-dial filter, click the filter, click **Edit**, edit the filter, and then click OK.
   - To delete a configured demand-dial filter, click the filter, and then click **Delete**.

5. After adding or editing a filter, in the **Set Demand-dial Filters** dialog box, click the filter, and in **Initiate connection**, click the appropriate action for the filter.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Demand-dial routing
Demand-dial routing design considerations

# Configure dial-out hours

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure dial-out hours**

1. Open Routing and Remote Access.
2. In the console tree, click **Network Interfaces**.

   **Where?**

   - Routing and Remote Access/*server name*/Network Interfaces

3. In the details pane, right-click the demand-dial interface you want to configure, and then click **Dial-out Hours**.
4. In the **Dial-out Hours** dialog box, select the appropriate day and time boxes to either permit or deny dial-out access, click either **Permitted** or **Denied**, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Demand-dial routing
Demand-dial routing design considerations

# Configure the calling router for certificate-based EAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the calling router for certificate-based EAP**

1. Open Routing and Remote Access.
2. In the console tree, click **Network Interfaces**.

   **Where?**

   - Routing and Remote Access/*server name*/Network Interfaces

3. In the details pane, right-click the appropriate demand-dial interface, and then click **Properties**.
4. On the **Security** tab, click **Advanced (custom settings)**, and then click **Settings**.
5. Under **Logon security**, click **Use Extensible Authentication Protocol (EAP)**, click **Smart card or other certificate (TLS) (encryption enabled)**, and then click **Properties**.
6. In the **Smart Card or Other Certificate (TLS) Properties** dialog box, click **Use a certificate on this computer**.
7. To enable validation of the server certificate, select the **Validate server certificate** check box, select the **Connect to these servers** check box, and then type the DNS domain name of the answering router preceded by a period.
8. In **Trusted Root Certification Authorities**, click the root certification authority of the answering router, and then click **OK**.
9. Click **OK** to save changes to the security configuration, and then click **OK** again to save changes to the demand-dial interface.
10. In the details pane, right-click the demand-dial interface, and then click **Set credentials**.
11. In **User name on certificate**, click the user certificate for this demand-dial connection, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- If the root certification authority for the answering router does not appear, the root certification authority certificate for the answering router may be in the personal store rather than trusted root certification authorities store.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files
EAP
Deploying certificate-based authentication for demand-dial routing
Business partner demand-dial connection
Branch office demand-dial connection
Configure the answering router for certificate-based EAP
Move a certificate

# Configure the answering router for certificate-based EAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the answering router for certificate-based EAP**

1. Open Routing and Remote Access.
2. Right-click the name of the remote access router, and then click **Properties**.
3. On the **Security** tab, click **Authentication Methods**.
4. In the **Authentication Methods** dialog box, select the **Extensible authentication protocol (EAP)** check box, click **OK**, and then click **OK** again.
5. In the console tree, double-click the router name, and then click **Remote Access Policies**.
6. In the details pane, right-click the remote access policy that will be used by your certificate-based routers, click **Properties**, and then click **Edit Profile**.
7. On the **Authentication** tab, select the **Extensible Authentication Protocol** check box, click **Smart card or other certificate (TLS)**, and then click **Configure**.
8. In the **Smart Card or Other Certificate (TLS) Properties** dialog box, select the machine certificate you want to use, and then click **OK**.
9. Click **OK** to save the settings of the profile, and then click **OK** again to save the settings of the policy.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
EAP
Deploying certificate-based authentication for demand-dial routing
Business partner demand-dial connection
Branch office demand-dial connection
Configure the calling router for certificate-based EAP

# Configure a pre-shared key for a demand-dial routing interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a pre-shared key for a demand-dial routing interface**

1. Open Routing and Remote Access.
2. Click **Network Interfaces**.

   **Where?**

   - Routing and Remote Access/*server name*/Network Interfaces

3. Right-click the interface you want to configure, and click **Properties**.
4. On the **Security** tab, click **IPSec Settings**.
5. Select the **Use pre-shared key for authentication** check box, and type the pre-shared key in **Key**.
6. Click **OK** twice.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- A pre-shared key can be any non-null string of any combination of up to 255 characters. Use care when typing the pre-shared key because any differences between the pre-shared key configured on the client and the pre-shared key configured on the server will prevent connections from being made.
- If the VPN type of the demand-dial routing interface you want to configure is not set to **L2TP IPSec VPN** or **Automatic**, the **IPSec Settings** button is not available.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Pre-shared keys
Configure a pre-shared key on a remote access server
Add a demand-dial interface

# Delete an Interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Delete an interface**

- Delete an interface from the router
- Delete an interface from a routable protocol
- Delete an interface from a routing protocol

# Delete an interface from the router

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an interface from the router**

1. Open Routing and Remote Access.
2. In the console tree, click **Network Interfaces**.

   **Where?**

   - Routing and Remote Access/*server name*/Network Interfaces

3. In the details pane, right-click the interface you want to delete, and then click **Delete**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Delete an interface from a routable protocol
Delete an interface from a routing protocol

# Delete an interface from a routable protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an interface from a routable protocol**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface you want to delete, and then click **Delete**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Delete an interface from the router
Delete an interface from a routing protocol

# Delete an interface from a routing protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an interface from a routing protocol**

1. Open Routing and Remote Access.
2. In the console tree, click a routing protocol.

   **Where?**

   - Routing and Remote Access*server name*\IP Routing*routing protocol*

3. In the details pane, right-click the interface you want to delete, and then click **Delete**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Delete an interface from the router
Delete an interface from a routable protocol

# Manage Devices and Ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage devices and ports**

- Enable routing on ports
- Add PPTP or L2TP ports

# Enable routing on ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable routing on ports**

1. Open Routing and Remote Access.
2. In the console tree, click **Ports**.

   **Where?**

   - Routing and Remote Access/*server name*/Ports

3. Right-click **Ports**, and then click **Properties**.
4. On the **Devices** tab, click the device you want to configure, and then click **Configure**.
5. In the **Configure Device** dialog box, select the connection options you want to enable for the device, and then click **OK**.
6. Click **Apply**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Routing interfaces, devices, and ports
Add PPTP or L2TP ports

# Add PPTP or L2TP ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add PPTP or L2TP ports**

1. Open Routing and Remote Access.
2. In the console tree, click **Ports**.

   **Where?**

   - Routing and Remote Access/*server name*/Ports

3. Right-click **Ports**, and then click **Properties**.
4. In the **Ports Properties** dialog box, click either **WAN Miniport (PPTP)** or **WAN Miniport (L2TP)**, and then click **Configure**.
5. In **Maximum ports**, type the number of ports, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.
- You cannot set the number of PPTP ports to 0. To disable PPTP, clear the **Remote access connections (inbound only)** and **Demand-dial routing connections (inbound and outbound)** check boxes.
- When you increase the number of ports for **WAN Miniport (PPTP)** or **WAN Miniport (L2TP)** from less than 1000 ports to greater than 1000 ports, you must restart the computer to apply the change.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Point-to-Point Tunneling Protocol
Layer Two Tunneling Protocol
Devices and ports
Enable routing on ports

# Manage Routing Protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage routing protocols**

- Add an IP routing protocol
- Delete an IP routing protocol
- Change preference levels

# Add an IP routing protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an IP routing protocol**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. Right-click **General**, and then click **New Routing Protocol**.
4. In **Routing protocols**, click the protocol you want to add, and then click **OK**.
5. If applicable, complete any configuration dialog boxes for the protocol.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Delete an IP routing protocol

# Delete an IP routing protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an IP routing protocol**

1. Open Routing and Remote Access.
2. In the console tree, click an **IP routing protocol**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/*IP routing protocol*

3. Right-click the IP routing protocol you want to delete, and then click **Delete**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add an IP routing protocol

# Change preference levels

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change preference levels**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. Right-click **General**, and then click **Properties**.
4. On the **Preference Levels** tab, click a route source, and then click either **Move Up** or **Move Down** to change the preference level of the protocol.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
IP Routing Protocols

# Manage Static Routes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage static routes**

- Add a static route
- Delete a static route
- Add a default static IP route
- Perform Auto-Static Updates

# Add a static route

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a static route**

1. Open Routing and Remote Access.
2. In the console tree, right-click **Static Routes**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/Static Routes

3. Click **New Static Route**.
4. In the **Static Route** dialog box, enter the interface, destination, network mask, gateway, and metric. If this is a demand-dial interface, **Gateway** is unavailable. You can also select the **Use this route to initiate demand-dial connections** check box to initiate a demand-dial connection for traffic that matches the route.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Delete a static route
Add a default static IP route
Perform Auto-Static Updates

# Delete a static route

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a static route**

1. Open Routing and Remote Access.
2. In the console tree, click **Static Routes**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/Static Routes

3. In the details pane, right-click the static route you want to delete, and then click **Delete**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add a static route
Add a default static IP route
Perform Auto-Static Updates

# Add a default static IP route

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a default static IP route**

1. Open Routing and Remote Access.
2. In the console tree, right-click **Static Routes**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/Static Routes

3. Click **New Static Route**.
4. In **Interface**, click the interface you want to use for the default route.
5. In **Destination**, type **0.0.0.0**.
6. In **Network mask**, type **0.0.0.0**.
7. In **Gateway**, do one of the following:

   - If the interface is a demand-dial interface, **Gateway** is unavailable. Select the **Use this route to initiate demand-dial connections** check box to initiate a demand-dial connection for traffic matching the route.
   - If the interface for this route is a LAN connection, such as Ethernet or token ring, type the IP address of the router interface that is on the same network segment as the LAN interface.

8. In **Metric**, type **1**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Add a static route
Delete a static route
Perform Auto-Static Updates

# Perform Auto-Static Updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Perform auto-static updates**

- Perform manual auto-static updates
- Perform scheduled auto-static updates

# Perform manual auto-static updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To perform manual auto-static updates**

1. Open Routing and Remote Access.

2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the demand-dial interface you want to update, and then click **Update Routes**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- When an auto-static update is performed, the existing auto-static routes for the interface are deleted before the update is requested from the other router. If there is no response to the request, then the router cannot replace the routes it has deleted. This may lead to a loss of connectivity to remote networks.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Demand-dial routing updates
Perform scheduled auto-static updates

# Perform scheduled auto-static updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To perform scheduled auto-static updates**

1. Create a batch file that contains the **netsh** commands you need to perform an auto-static update.

2. Open Task Scheduler, double-click **Add Scheduled Task**, and then follow the instructions in the Scheduled Task wizard to schedule the batch file to run.

## Notes

- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.

- When an auto-static update is performed, the existing auto-static routes for the interface are deleted before the update is requested from the other router. If there is no response to the request, then the router cannot replace the routes it has deleted. This may lead to a loss of connectivity to remote networks.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Demand-dial routing updates
Perform manual auto-static updates
Netsh overview

# Manage Routers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage routers**

- Administer a remote router
- View routing tables
- Reset the Routing and Remote Access service

# Administer a remote router

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To administer a remote router**

1. Open Routing and Remote Access.
2. In the console tree, right-click **Server Status**, and then click **Add Server**.
3. In the **Add Server** dialog box, do one of the following:
   - Click **The following computer**, and type the computer name or IP address of the server.
   - Click **All routing and remote access servers in the domain**, and then type the domain containing the server you want to administer. Click **OK**, and then select the server.
   - Click **Browse Active Directory**, click **Next**, and in the **Find Routers or Remote Access Servers** dialog box, select the check boxes next to the types of servers that you want to search for. Click **OK**, and then select the server.
4. Once the remote server appears as an item in the console tree, you can administer it.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- You can administer a server running Routing and Remote Access by using Terminal Services.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Enable the Routing and Remote Access service

# View routing tables

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view routing tables**

1. Open Routing and Remote Access.
2. In the console tree, right-click **Static Routes**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/Static Routes

3. Click **Show IP Routing Table**.

- The following table shows the detail views that are available in Routing and Remote Access by right-clicking the component.

| Component | Views |
|---|---|
| IP Routing/General | TCP/IP information |
| | Multicast forwarding table |
| | Multicast statistics |
| IP Routing/General/Interface | TCP/IP information |
| | Address translations |
| | IP addresses |
| | IP routing table |
| | TCP connections |
| | UDP listener ports |
| IP Routing/Static Routes | IP routing table |
| IP Routing/OSPF | Areas |
| | Link state database |
| | Neighbors |
| | Virtual interfaces |
| IP Routing/RIP | Neighbors |
| IP Routing/NAT/Basic Firewall | DHCP Allocator information |
| | DNS Proxy information |
| IP Routing/NAT/Basic Firewall/Interface | Mappings |
| IP Routing/IGMP | Group table |

| | |
|---|---|
| IP Routing/IGMP/Interface | Interface group table |

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Understanding the IP routing table

# Reset the Routing and Remote Access service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reset the Routing and Remote Access service**

1. Open Routing and Remote Access.
2. Right-click the computer name for which you want to reset the Routing and Remote Access service, and then click **Disable Routing and Remote Access**.
3. Right-click the computer name for which you want to reset the Routing and Remote Access service, and then click **Configure and Enable Routing and Remote Access**.
4. Follow the instructions in the Routing and Remote Access Server Setup wizard.
5. When prompted to start the Routing and Remote Access service, click **Yes**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This procedure resets the current configuration back to the default values based on your selection of options in the Routing and Remote Access wizard. This includes removing all IP routing protocols and their configurations.
- Resetting the Routing and Remote Access service does not reset your remote access policies.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Enable the Routing and Remote Access service
Administer a remote router
Re-create the default remote access policy

# Manage Packet Filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage packet filters**

- Add a packet filter
- Modify a packet filter
- Delete a packet filter
- Add local host filters
- Add PPTP Filters
- Add L2TP over IPSec Filters

# Add a packet filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a packet filter**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to add a filter, and then click **Properties**.
4. On the **General** tab, click either **Inbound Filters** or **Outbound Filters**.
5. In the **Inbound Filters** or **Outbound Filters** dialog box, click **New**.
6. In the **Add IP Filter** dialog box, type the settings for the filter, and then click **OK**.
7. In **Filter action**, select the appropriate filter action, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Modify a packet filter
Delete a packet filter

# Modify a packet filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify a packet filter**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to modify a filter, and then click **Properties**.
4. On the **General** tab, click either **Inbound Filters** or **Outbound Filters**.
5. In the **Inbound Filters** or **Outbound Filters** dialog box, click the filter you want to modify, and then click **Edit**.
6. In the **Edit IP Filter** dialog box, change the appropriate settings for the filter, and then click **OK**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Add a packet filter
Delete a packet filter

# Delete a packet filter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a packet filter**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to delete a filter, and then click **Properties**.
4. On the **General** tab, click either **Inbound Filters** or **Outbound Filters**.
5. In the **Inbound Filters** or **Outbound Filters** dialog box, click the filter you want to delete, and then click **Delete**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Add a packet filter
Modify a packet filter

# Add local host filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add local host filters**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to set the filter, and then click **Properties**.
4. On the **General** tab, click **Inbound Filters**.
5. In the **Inbound Filters** dialog box, click **New**, and create a set of five input filters.
6. In the **Inbound Filters** dialog box, click **Drop all packets except those that meet the criteria below**, and then click **OK**.

The five filters are described in the following table.

| Filter | Example |
|---|---|
| Accept packets sent directly to your computer. | Your router is configured with an IP address of 10.1.1.1 and a subnet mask of 255.255.0.0. To allow packets with a destination of your router, add a filter with a **Destination IP address** of 10.1.1.1, a **Destination Subnet mask** of 255.255.255.255, and select **Any** as the type of protocol. |
| Accept packets broadcast to the local subnet. | The second filter enables you to receive packets that are going to the 10.1.0.0 subnetted network. Add a filter with a **Destination IP address** of 10.1.255.255, a **Destination Subnet mask** of 255.255.255.255, and select **Any** as the type of protocol. |
| Accept packets broadcast to all subnets of a class-based network. | Set this filter to allow packets going to all subnets of the network 10.0.0.0. Add a filter with a **Destination IP address** of 10.255.255.255, a **Destination Subnet mask** of 255.255.255.255, and select **Any** as the type of protocol. |
| Accept packets sent to the limited broadcast (the all-1s address) | Add a filter with a **Destination IP address** of 255.255.255.255, a **Destination Subnet mask** of 255.255.255.255, and select **Any** as the type of protocol. |
| Accept all IP multicast packets | Add a filter with a **Destination IP address** of 224.0.0.0, a **Destination Subnet mask** of 240.0.0.0, and select **Any** as the type of protocol. |

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Add a packet filter
Modify a packet filter
Delete a packet filter

# Add PPTP Filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Add PPTP filters**

- Select the PPTP interface
- Set PPTP input filters
- Set PPTP output filters

# Select the PPTP interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select the PPTP interface**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, click the interface on which you want to enable PPTP filtering, scroll to the **IP Address** column, and write down the IP address assigned to the interface.
4. Right-click the interface, and then click **Properties**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- All six filters work together to complete PPTP filtering. The PPTP filtering is not secure unless all six filters are set correctly.
- If the six filters are the only filters configured, then the only traffic that is allowed in and out of the interface is PPTP traffic to and from the PPTP server and PPTP client on the computer running Routing and Remote Access.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Add a packet filter
Set PPTP input filters
Set PPTP output filters

# Set PPTP input filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set PPTP input filters**

To set PPTP input filters, you must configure up to three input filters and select the appropriate filter action.

**To add the first input filter**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to add a filter, and then click **Properties**.
4. On the **General** tab, click **Inbound Filters**.
5. In the **Inbound Filters** dialog box, click **New**.
6. In the **Add IP Filter** dialog box, select the **Destination network** check box.
7. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.
8. In **Protocol**, click **Other**.
9. In **Protocol number**, type **47**, and then click **OK**.

**To add the second input filter**

1. In the **Inbound Filters** dialog box, click **New**.
2. In the **Add IP Filter** dialog box, select the **Destination network** check box.
3. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.
4. In **Protocol**, click **TCP**.
5. In **Source port**, type **0**.
6. In **Destination port**, type **1723**, and then click **OK**.

**To add the third input filter (optional)**

If the PPTP server computer is also used as a PPTP client, you need to configure an additional filter.

1. In the **Inbound Filters** dialog box, click **New**.
2. In the **Add IP Filter** dialog box, select the **Destination network** check box.
3. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.

4. In **Protocol**, click **TCP [established]**.

5. In **Source port**, type **1723**.

6. In **Destination port**, type **0**, and then click **OK**.

**To select the filter action for the input filters**

- In the **Inbound Filters** dialog box, click **Drop all packets except those that meet the criteria below**, and then click **OK**.

## Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Add a packet filter
Select the PPTP interface
Set PPTP output filters

# Set PPTP output filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set PPTP output filters**

To set PPTP output filters, you must configure up to three output filters and select the appropriate filter action.

**To add the first output filter**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to add a filter, and then click **Properties**.
4. On the **General** tab, click **Outbound Filters**.
5. In the **Outbound Filters** dialog box, click **New**.
6. In the **Add IP Filter** dialog box, select the **Source network** check box.
7. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.
8. In **Protocol**, click **Other**.
9. In the **Protocol** box, type **47**, and then click **OK**.

**To add the second output filter**

1. In the **Outbound Filters** dialog box, click **New**.
2. In the **Add IP Filter** dialog box, select the **Source network** check box.
3. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.
4. In **Protocol**, click **TCP**.
5. In **Source port**, type **1723**.
6. In **Destination port**, type **0**, and then click **OK**.

**To add the third output filter (optional)**

If the PPTP server computer is also used as a PPTP client, you need to configure an additional filter.

1. In the **Outbound Filters** dialog box, click **New**.
2. In the **Add IP Filter** dialog box, select the **Source network** check box.
3. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.

4. In **Protocol**, click **TCP**.

5. In **Source port**, type **0**.

6. In **Destination port**, type **1723**, and then click **OK**.

**To select the filter action for the output filters**

- In the **Outbound Filters** dialog box, click **Drop all packets except those that meet the criteria below**, and then click **OK**.

## Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Add a packet filter
Select the PPTP interface
Set PPTP input filters

# Add L2TP over IPSec Filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Add L2TP over IPSec filters**

- Select the L2TP over IPSec interface
- Set L2TP over IPSec input filters
- Set L2TP over IPSec output filters

# Select the L2TP over IPSec interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select the L2TP over IPSec interface**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, click the interface on which you want to enable L2TP over IPSec filtering, scroll to the **IP Address** column, and write down the IP address assigned to the interface.
4. Right-click the interface, and then click **Properties**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- All four L2TP over IPSec input and output filters work together to complete L2TP over IPSec filtering. The L2TP over IPSec filtering is not secure unless all four filters are set correctly.
- If the four L2TP over IPSec filters are the only filters that are configured, then the only traffic that is allowed in and out of the interface is L2TP over IPSec traffic to and from the L2TP server and L2TP client on the computer running Routing and Remote Access.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Add a packet filter
Set L2TP over IPSec input filters
Set L2TP over IPSec output filters

# Set L2TP over IPSec input filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set L2TP over IPSec input filters**

To set L2TP over IPSec input filters, you must configure the filters and select the appropriate filter action.

**To add the first L2TP over IPSec input filter**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to add a filter, and then click **Properties**.
4. On the **General** tab, click **Inbound Filters**.
5. In the **Inbound Filters** dialog box, click **New**.
6. In the **Add IP Filter** dialog box, select the **Destination network** check box.
7. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.
8. In **Protocol**, click **UDP**.
9. In **Source port**, type **0**.
10. In **Destination port**, type **500**, and then click **OK**.

**To add the second L2TP over IPSec input filter**

1. On the **General** tab, click **Inbound Filters**.
2. In the **Inbound Filters** dialog box, click **New**.
3. In the **Add IP Filter** dialog box, select the **Destination network** check box.
4. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.
5. In **Protocol**, click **UDP**.
6. In **Source port**, type **0**.
7. In **Destination port**, type **1701**, and then click **OK**.

**To add the third L2TP over IPSec input filter**

1. On the **General** tab, click **Inbound Filters**.

2. In the **Inbound Filters** dialog box, click **New**.

3. In the **Add IP Filter** dialog box, select the **Destination network** check box.

4. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.

5. In **Protocol**, click **UDP**.

6. In **Source port**, type **0**.

7. In **Destination port**, type **4500**, and then click **OK**.

#### To select the filter action for the input filter

- In the **Inbound Filters** dialog box, click **Drop all packets except those that meet the criteria below**, and then click **OK**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Add a packet filter
Select the L2TP over IPSec interface
Set L2TP over IPSec output filters

# Set L2TP over IPSec output filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set L2TP over IPSec output filters**

To set L2TP over IPSec output filters, you must configure the filters and select the appropriate filter action.

**To add the first L2TP over IPSec output filter**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to add a filter, and then click **Properties**.
4. On the **General** tab, click **Outbound Filters**.
5. In the **Outbound Filters** dialog box, click **New**.
6. In the **Add IP Filter** dialog box, select the **Source network** check box.
7. In **IP Address**, type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.
8. In **Protocol**, click **UDP**.
9. In **Source port**, type **500**.
10. In **Destination port**, type **0**, and then click **OK**.

**To add the second L2TP over IPSec output filter**

1. On the **General** tab, click **Outbound Filters**.
2. In the **Outbound Filters** dialog box, click **New**.
3. In the **Add IP Filter** dialog box, select the **Source network** check box.
4. In **IP Address** type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.
5. In **Protocol**, click **UDP**.
6. In **Source port**, type **1701**.
7. In **Destination port**, type **0**, and then click **OK**.

**To add the third L2TP over IPSec output filter**

1. On the **General** tab, click **Outbound Filters**.

2. In the **Outbound Filters** dialog box, click **New**.

3. In the **Add IP Filter** dialog box, select the **Source network** check box.

4. In **IP Address** type the IP address of the interface, and in **Subnet mask**, type **255.255.255.255**.

5. In **Protocol**, click **UDP**.

6. In **Source port**, type **4500**.

7. In **Destination port**, type **0**, and then click **OK**.

#### To select the filter action for the output filter

- In the **Outbound Filters** dialog box, click **Drop all packets except those that meet the criteria below**, and then click **OK**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Packet filtering
Add a packet filter
Select the L2TP over IPSec interface
Set L2TP over IPSec input filters

# Configure RIP for IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure RIP for IP**


- Configure RIP version 2
- Enable authentication
- Enable auto-static update mode
- Add route filters
- Add peer filters
- Add a unicast neighbor
- Enable silent RIP
- Configure Advanced RIP Settings

# Configure RIP version 2

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure RIP version 2**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. In the details pane, right-click the interface you want to configure for RIP version 2, and then click **Properties**.
4. On the **General** tab, in **Outgoing packet protocol**, do one of the following:

   - If there are RIP version 1 routers on the same network as this interface, click **RIP version 2 broadcast**.
   - If there are only RIP version 2 routers on the same network as this interface, or if the interface is a demand-dial interface, click **RIP version 2 multicast**.

5. On the **General** tab, in **Incoming packet protocol**, do one of the following:

   - If there are RIP version 1 routers on the same network as this interface, click **RIP version 1 and 2**.
   - If there are only RIP version 2 routers on the same network as this interface, or if the interface is a demand-dial interface, click **RIP version 2 only**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP
RIP-for-IP design considerations
RIP-for-IP security

# Enable authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable authentication**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. In the details pane, right-click the interface on which you want to set authentication, and then click **Properties**.
4. On the **General** tab, select the **Activate authentication** check box, and in **Password**, type a password.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- Authentication is used only for RIP version 2.
- All RIP version 2 routers on a network must use matching passwords or the routers cannot process each other's route announcements.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP
RIP-for-IP design considerations
RIP-for-IP security
Configure RIP version 2

# Enable auto-static update mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable auto-static update mode**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. In the details pane, right-click the interface on which you want to set auto-static update mode, and then click **Properties**.
4. On the **General** tab, in **Operation mode**, click **Auto-static update mode**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP
RIP-for-IP design considerations
Configure RIP version 2
Demand-dial routing updates

# Add route filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add route filters**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. In the details pane, right-click the interface on which you want to set route filters, and then click **Properties**.
4. On the **Security** tab, in **Action**, do one of the following:

   - If you want filters to be applied to incoming RIP-for-IP announcement packets, click **For incoming routes**.
   - If you want filters to be applied to outgoing RIP-for-IP announcement packets, click **For outgoing routes**.

5. Do one of the following:

   - To include all routes, click **Accept all routes**.
   - To specify a range of addresses to accept or exclude, click **Accept all routes in the ranges listed** or **Ignore all routes in the ranges listed**, type in the range, and click **Add**.

6. After you set route filters, click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP
RIP-for-IP security

# Add peer filters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add peer filters**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. Right-click **RIP**, and then click **Properties**.
4. On the **Security** tab, do one of the following:

   - To process all announcements from all routers, click **Accept announcements from all routers**.
   - To process announcements only from specific routers, click **Accept announcements from listed routers only**.
   - To discard all announcements from specific routers, click **Ignore announcements from listed routers**.

5. If you have selected **Accept announcements from listed routers only** or **Ignore announcements from listed routers**, in **Router IP address**, type the IP address of the router, and then click **Add**. You must add each router individually.
6. Click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- The action you select applies to all the peer routers; you cannot set the action per router.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP
RIP-for-IP security

# Add a unicast neighbor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a unicast neighbor**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. In the details pane, right-click the interface on which you want to set unicast neighbors, and then click **Properties**.
4. On the **Neighbors** tab, do one of the following:

   - To disable the use of neighboring routers, click **Use broadcast or multicast only**.
   - To use neighboring routers in addition to broadcast or multicast, click **Use neighbors in addition to broadcast or multicast**.
   - To use neighboring routers instead of broadcast or multicast, click **Use neighbors instead of broadcast or multicast**.

5. If you chose to use neighbors, in the **IP address** box, type the IP address of the unicast neighbor, and then click **Add**. You must add each neighbor separately.
6. Click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- The option you select applies to all of the unicast neighbors.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP
RIP-for-IP design considerations
RIP-for-IP security

# Enable silent RIP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable silent RIP**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. In the details pane, right-click the interface on which you want to set silent RIP mode, and then click **Properties**.
4. On the **General** tab, in **Outgoing packet protocol**, click **Silent RIP**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP
RIP-for-IP design considerations
RIP-for-IP security

# Configure Advanced RIP Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure advanced RIP settings**

- Set timers for periodic update mode
- Set split-horizon processing
- Set poison-reverse processing

# Set timers for periodic update mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set timers for periodic update mode**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. In the details pane, right-click the interface on which you want to set timer values, and then click **Properties**.
4. On the **Advanced** tab, in **Periodic announcement interval (seconds)**, **Time before routes expire (seconds)**, and **Time before route is removed (seconds)**, type the appropriate time values in seconds.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP

# Set split-horizon processing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set split-horizon processing**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. In the details pane, right-click the interface on which you want to configure split-horizon processing, and then click **Properties**.
4. On the **Advanced** tab, either select the **Enable split-horizon processing** check box to enable split-horizon processing or clear the **Enable split-horizon processing** check box to disable split-horizon processing, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP

# Set poison-reverse processing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set poison-reverse processing**

1. Open Routing and Remote Access.
2. In the console tree, click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. In the details pane, right-click the interface on which you want to enable poison-reverse processing, and then click **Properties**.
4. On the **Advanced** tab, select the **Enable poison-reverse processing** check box, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
RIP for IP

# Configure OSPF

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure OSPF**

- Add the OSPF routing protocol
- Configure OSPF Global Settings
- Configure OSPF Interface Settings

## Notes

- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

# Add the OSPF routing protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add the OSPF routing protocol**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. Right-click **General**, and then click **New Routing Protocol**.
4. In the **New Routing Protocol** dialog box, click **Open Shortest Path First (OSPF)**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
OSPF
Add an interface to OSPF

# Configure OSPF Global Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure OSPF global settings**

- Create an OSPF area
- Configure ranges for an OSPF area
- Configure an ASBR
- Add a virtual interface

## Notes

- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

# Create an OSPF area

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an OSPF area**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. Right-click **OSPF**, and then click **Properties**.
4. On the **Areas** tab, click **Add**.
5. On the **General** tab, in **Area ID**, type the dotted decimal number that identifies the area.
6. To use a plaintext password, verify that the **Enable plaintext password** check box is selected.
7. To mark the area as a stub, select the **Stub area** check box.
8. In **Stub metric**, click the arrows to set the stub metric.
9. To import routes of other areas into the stub area, select the **Import summary advertisements** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.
- You cannot configure the backbone area (area ID 0.0.0.0) as a stub area, and you cannot configure virtual links through stub areas.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
OSPF
Configure ranges for an OSPF area

# Configure ranges for an OSPF area

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure ranges for an OSPF area**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. Right-click **OSPF**, and then click **Properties**.
4. On the **Areas** tab, click **Add**.
5. On the **Ranges** tab, in **Destination**, type the IP network ID for the range.
6. In **Network mask**, type the associated mask for the range, and then click **Add**.
7. To delete a range, click the range you want to delete, and then click **Remove**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
OSPF
Create an OSPF area

# Configure an ASBR

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure an ASBR**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. Right-click **OSPF**, and then click **Properties**.
4. On the **General** tab, click **Enable autonomous system boundary router**.
5. To configure external route sources, do the following:

   a. On the **External Routing** tab, click either **Accept routes from all route sources except those selected** or **Ignore routes from all route sources except those selected**.

   b. Select or clear the appropriate check boxes next to the route sources.

6. To configure external route filters, do the following:

   a. On the **External Routing** tab, click **Route Filters**.

   b. In **Destination** and in **Network mask**, type the route you want to filter, and then click **Add**.

   c. Repeat the preceding step for all the routes you want to filter.

   d. Click either **Ignore listed routes** or **Accept listed routes** depending on the appropriate filtering action, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files
OSPF
OSPF design considerations
Troubleshooting OSPF

# Add a virtual interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a virtual interface**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. Right-click **OSPF**, and then click **Properties**.
4. On the **Virtual Interfaces** tab, click **Add**.
5. In **Transit area ID**, click the transit area over which you are connecting the virtual link.
6. In **Virtual neighbor router ID**, type the OSPF router ID of the router at the other endpoint of the virtual link.
7. In **Transit delay (seconds)**, click the arrows to set the transit delay in seconds.
8. In **Retransmit interval (seconds)**, click the arrows to set the retransmit interval in seconds.
9. In **Hello interval (seconds)**, click the arrows to set the hello interval in seconds.
10. In **Dead interval (seconds)**, click the arrows to set the dead interval in seconds.
11. If the backbone area is configured to have a password, in **Plaintext password**, type a password.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.
- To configure a virtual link, each virtual link neighbor must add and configure a virtual interface.
- The hello interval, dead interval, and password must match between the virtual link neighbors.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files
OSPF
OSPF design considerations
Troubleshooting OSPF

# Configure OSPF Interface Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure OSPF interface settings**

- Add an interface to OSPF
- Configure an OSPF interface
- Set a password on an OSPF interface
- Configure the hello and dead intervals
- Add an OSPF neighbor

## Notes

- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

# Add an interface to OSPF

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an interface to OSPF**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access\server name\IP Routing\OSPF

3. Right-click **OSPF**, and then click **New Interface**.
4. In **Interfaces**, click the interface you want to add, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
OSPF
Configure an OSPF interface

# Configure an OSPF interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure an OSPF interface**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. In the details pane, right-click the interface you want to configure, and then click **Properties**.
4. On the **General** tab, select the **Enable OSPF for this address** check box.
5. In **Area ID**, click the ID of the area to which the interface belongs.
6. In **Router priority**, click the arrows to set the priority of the router over the interface.
7. In **Cost**, click the scroll arrows to set the cost of sending a packet over the interface.
8. If the area to which the interface belongs is enabled for passwords, in **Password**, type a password.
9. Under **Network type**, click the type of OSPF interface.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.
- If the interface has multiple IP addresses configured on it, then the **General** tab displays an **IP Address** box. You need to configure OSPF settings for each IP address on the interface.
- A router priority of **0** means that the router cannot become an OSPF-designated router.
- A larger cost means the interface is not used as much as other interfaces with a lower cost.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files

OSPF

Set a password on an OSPF interface

Configure the hello and dead intervals

# Set a password on an OSPF interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set a password on an OSPF interface**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. In the details pane, right-click the interface on which you want to set a password, and then click **Properties**.
4. On the **General** tab, in **Password**, type a password.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.
- All interfaces in the same area that are on the same network must use identical passwords.
- Interfaces in the same area that are on different networks can have different passwords.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
OSPF
OSPF security
Create an OSPF area
Configure an OSPF interface
Configure the hello and dead intervals

# Configure the hello and dead intervals

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the hello and dead intervals**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. In the details pane, right-click the interface you want to configure, and then click **Properties**.
4. On the **Advanced** tab, in **Hello interval (seconds)**, type the number of seconds between transmissions of hello packets by the router on the interface.
5. In **Dead interval (seconds)**, type the number of seconds before a silent remote router is declared down.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.
- The dead interval should be a multiple of the hello interval (commonly 4).
- The hello and dead intervals must be the same for all routers attached to a common network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
OSPF
Configure an OSPF interface
Set a password on an OSPF interface

# Add an OSPF neighbor

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an OSPF neighbor**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. In the details pane, right-click the interface you want to configure, and then click **Properties**.
4. On the **NBMA Neighbors** tab, in **Neighbor IP address**, type the IP interface address (not the router ID) of a router that is attached to the nonbroadcast network.
5. In **Router priority**, click the arrows to set the priority for the neighbor, and then click **Add**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.
- You need to repeat the preceding steps for all the routers that are attached to the nonbroadcast network.
- The priority determines the neighbor's eligibility to become a designated router. When an interface to a nonbroadcast network comes up, the router sends hello packets only to those neighbors eligible to become a designated router, until the identity of the designated router is discovered.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
OSPF
OSPF design considerations
Configure an OSPF interface
Set a password on an OSPF interface
Configure the hello and dead intervals

# Configure NAT/Basic Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure NAT/Basic Firewall**


- Add network address translation
- Add and configure an interface for network address translation
- Enable network address translation addressing
- Enable network address translation name resolution
- Configure interface IP address ranges
- Configure Basic Firewall for a public interface
- Configure services and ports
- Allow or deny ICMP messages

# Add network address translation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add network address translation**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. Right-click **General**, and then click **New Routing Protocol**.
4. In the **Select Routing Protocol** dialog box, click **NAT/Basic Firewall**, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Understanding Network Address Translation
Add and configure an interface for network address translation
Deploying network address translation

# Add and configure an interface for network address translation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add and configure an interface for network address translation**

1. Open Routing and Remote Access.
2. In the console tree, click **NAT/Basic Firewall**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/NAT/Basic Firewall

3. Right-click **NAT/Basic Firewall**, and then click **New Interface**.
4. In **Interfaces**, click the interface you want to add, and then click **OK**.
5. Do one of the following:

   - If this interface connects to the Internet, on the **NAT/Basic Firewall** tab, click **Public interface connected to the Internet** and select the **Enable NAT on this interface** check box. If you want the public interface to be protected by dynamic packet filters, select the **Enable a basic firewall on this interface** check box.
   - If this interface connects to the small office or home network, on the **NAT/Basic Firewall** tab, click **Private interface connected to private network**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- You can enable Basic Firewall without NAT for a public interface by clicking **Basic Firewall only**.
- For a dial-up connection to the Internet, select the demand-dial interface that is configured to connect to your Internet service provider (ISP).
- For a permanent connection to the Internet, select the permanent interface that is connected to your ISP.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Understanding Network Address Translation
Network address translation design considerations

Deploying network address translation
Add network address translation
Configure Basic Firewall for a public interface

# Enable network address translation addressing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable network address translation addressing**

1. Open Routing and Remote Access.
2. In the console tree, click **NAT/Basic Firewall**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/NAT/Basic Firewall

3. Right-click **NAT/Basic Firewall**, and then click **Properties**.
4. On the **Address Assignment** tab, select the **Automatically assign IP addresses by using the DHCP allocator** check box.
5. (Optional) To allocate to DHCP clients on the private network, in **IP address** and **Mask**, configure the range of IP addresses.
6. (Optional) To exclude addresses from allocation to DHCP clients on the private network, click **Exclude**, click **Add**, and then configure the addresses.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Understanding Network Address Translation
Network address translation design considerations
Deploying network address translation

# Enable network address translation name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable network address translation name resolution**

1. Open Routing and Remote Access.
2. In the console tree, click **NAT/Basic Firewall**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/NAT/Basic Firewall

3. Right-click **NAT/Basic Firewall**, and then click **Properties**.
4. For host name resolution to a DNS server, on the **Name Resolution** tab, select the **Clients using Domain Name System (DNS)** check box.
5. If you want the connection to the Internet to be initiated when a host on the private network sends a DNS name query to the Network Address Translation (NAT) computer, select the **Connect to the public network when a name needs to be resolved** check box, and then click the name of the appropriate demand-dial interface in **Demand-dial interface**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Understanding Network Address Translation
Network address translation design considerations
Deploying network address translation

# Configure interface IP address ranges

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure interface IP address ranges**

1. Open Routing and Remote Access.
2. In the console tree, click **NAT/Basic Firewall**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/NAT/Basic Firewall

3. In the details pane, right-click the interface you want to configure, and then click **Properties**.
4. On the **Address Pool** tab, click **Add**, and do one of the following:

   - If you are using a range of IP addresses that can be expressed with an IP address and a subnet mask, in **Start address**, type the starting IP address, and in **Mask**, type the subnet mask.
   - If you are using a range of IP addresses that cannot be expressed with an IP address and a subnet mask, in **Start address**, type the starting IP address, and in **End address**, type the ending IP address.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- If you have multiple address ranges, you can use **Add** to add each one separately.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Understanding Network Address Translation
Network address translation design considerations
Deploying network address translation

# Configure Basic Firewall for a public interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Basic Firewall**

1. Open Routing and Remote Access.
2. In the console tree, click **NAT/Basic Firewall**.

   **Where?**

   - Routing and Remote Access/*ServerName*/IP Routing/NAT/Basic Firewall

3. In the details pane, right-click the interface you want to configure, and then click **Properties**.
4. On the **NAT/Basic Firewall** tab, do one of the following:

   - Click **Public interface connected to the Internet**, and select the **Enable a basic firewall on this interface** check box.
   - Click **Basic firewall only**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- You cannot enable Internet Connection Firewall (ICF) or Windows Firewall on a server running Routing and Remote Access.
- You can allow or prevent certain types of network traffic from reaching your network by configuring additional static packet filters. You can also accept or deny certain types of Internet Control Message Protocol (ICMP) messages, traffic required for particular services, or traffic that travels through particular ports.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Basic Firewall
Configure services and ports
Allow or deny ICMP messages
Internet Control Message Protocol (ICMP)

# Configure services and ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure services and ports**

1. Open Routing and Remote Access.
2. In the console tree, click **NAT/Basic Firewall**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/NAT/Basic Firewall

3. In the details pane, right-click the interface you want to configure, and then click **Properties**.
4. On the **Services and Ports** tab, do one of the following:

   - If the port is associated with a common network service, review the list in **Services**, select the service, review the settings displayed, and click **OK**.
   - If the port is not associated with a common network service, or if that service is not listed under **Services**, click **Add**, provide a name for the service and the settings required for the incoming and outgoing ports, and click **OK**.
   - If you want to disable a port that is associated with a common network service, review the list in **Services**, and clear the appropriate check box.
   - If you want to disable a previously added port, review the list in **Services**, and the clear the appropriate check box for the service that you added.
   - If you want to remove a previously added port, review the list in **Services**, select the service that you added, and click **Delete**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Understanding Network Address Translation
Network address translation design considerations
Deploying network address translation

Allow or deny ICMP messages
Packet filtering

# Allow or deny ICMP messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow or deny ICMP messages**

1. Open Routing and Remote Access.
2. In the console tree, click **NAT/Basic Firewall**.

   **Where?**

   - Routing and Remote Access/*ServerName*/IP Routing/NAT/Basic Firewall

3. In the details pane, right-click the interface you want to configure, and then click **Properties**.
4. Click the ICMP tab, and in **Allow the following**, do the following:

   - To allow a type of ICMP message to reach your network, select the appropriate check box.
   - To prevent a type of ICMP message from reaching your network, clear the appropriate check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- To display a description of a message type, click the message type. The description appears under Description.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Understanding Network Address Translation
Network address translation design considerations
Deploying network address translation
Configure services and ports
Internet Control Message Protocol (ICMP)

# Enable ICMP router discovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable ICMP router discovery**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface you want to enable, and then click **Properties**.
4. On the **General** tab, select the **Enable router discovery advertisements** check box.
5. In **Advertisement lifetime (minutes)**, type or select the time after which a router is considered down after hearing its last router advertisement.
6. In **Minimum time (minutes)**, type or select the minimum rate at which the router periodically sends ICMP router advertisements.
7. In **Maximum time (minutes)**, type or select the maximum rate at which the router periodically sends ICMP router advertisements.

   Based on the values of minimum time and maximum time, the router sends periodic ICMP router advertisements between the minimum and maximum times.

8. In **Level of preference**, type or select the level of preference for this router to be a default gateway for hosts.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
ICMP router discovery

# Configure the DHCP Relay Agent

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure the DHCP Relay Agent**


- Add the DHCP Relay Agent
- Configure global DHCP Relay Agent properties
- Enable the DHCP Relay Agent on a router interface

# Add the DHCP Relay Agent

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add the DHCP Relay Agent**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. Right-click **General**, and then click **New Routing Protocol**.
4. In the **Select Routing Protocol** dialog box, click **DHCP Relay Agent**, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
DHCP Relay Agent
Configure global DHCP Relay Agent properties
Enable the DHCP Relay Agent on a router interface
Troubleshooting the DHCP Relay Agent

# Configure global DHCP Relay Agent properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure global DHCP Relay Agent properties**

1. Open Routing and Remote Access.
2. In the console tree, click **DHCP Relay Agent**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/DHCP Relay Agent

3. Right-click **DHCP Relay Agent**, and then click **Properties**.
4. On the **General** tab, in **Server address**, type the IP address of your DHCP server, and then click **Add**.
5. Repeat step 4 for each DHCP server you need to add, and then click **OK**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
DHCP Relay Agent
Add the DHCP Relay Agent
Enable the DHCP Relay Agent on a router interface
Troubleshooting the DHCP Relay Agent

# Enable the DHCP Relay Agent on a router interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable the DHCP Relay Agent on a router interface**

1. Open Routing and Remote Access.
2. In the console tree, click **DHCP Relay Agent**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/DHCP Relay Agent

3. Right-click **DHCP Relay Agent**, and then click **New Interface**.
4. Click the interface you want to add, and then click **OK**.
5. In the **DHCP Relay Properties** dialog box, on the **General** tab, verify that the **Relay DHCP packets** check box is selected.
6. If needed, in **Hop-count threshold** and **Boot threshold (seconds)**, click the arrows to modify the thresholds.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
DHCP Relay Agent
Add the DHCP Relay Agent
Configure global DHCP Relay Agent properties
Troubleshooting the DHCP Relay Agent

# Configure IP Multicast Support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure IP multicast support**

- Configure the IGMP Routing Protocol
- Configure multicast scopes
- Configure multicast boundaries
- Configure multicast heartbeat

# Configure the IGMP Routing Protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure the IGMP routing protocol**

- Add the IGMP routing protocol
- Enable IGMP router and IGMP proxy mode
- Configure IGMP router settings

# Add the IGMP routing protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add the IGMP routing protocol**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. Right-click **General**, and then click **New Routing Protocol**.
4. In the **Select Routing Protocol** dialog box, click **IGMP Router and Proxy**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- By default, IGMP protocol version 3 is selected. To change the IGMP protocol version, select a different version when configuring a new interface for IGMP, or change the version an interface uses by right-clicking the interface and selecting **Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Multicast forwarding overview
Multicast routing
Enable IGMP router and IGMP proxy mode
Configure IGMP router settings

# Enable IGMP router and IGMP proxy mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable IGMP router and IGMP proxy mode**

1. Open Routing and Remote Access.
2. In the console tree, click **IGMP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/IGMP

3. Right-click **IGMP**, and then click **New Interface**.
4. In **Interfaces**, click the interface you want to enable, and then click **OK**.
5. In the **General** dialog box, verify that the **Enable IGMP** check box is selected.
6. Under **Mode**, do one of the following:

   - To enable IGMP router mode, click **IGMP router**.
   - To enable IGMP proxy mode, click **IGMP proxy**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- You can only enable IGMP proxy mode on a single interface.
- By default, IGMP protocol version 3 is selected. To change the IGMP protocol version, select a different version when configuring a new interface for IGMP, or change the version an interface uses by right-clicking the interface and selecting **Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Multicast forwarding overview
Multicast routing
Add the IGMP routing protocol
Configure IGMP router settings

# Configure IGMP router settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure IGMP router settings**

1. Open Routing and Remote Access.
2. In the console tree, click **IGMP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/IGMP

3. In the details pane, right-click the interface you want to configure, and then click **Properties**.
4. Click the **Router** tab, and then configure the appropriate settings.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- Typically, you do not need to modify the default IGMP router settings. You should only modify these settings after you have made a careful analysis of your network.
- Items on the **Router** tab are available only if IGMP router mode is enabled on the selected interface.
- By default, IGMP protocol version 3 is selected. To change the IGMP protocol version, select a different version when configuring a new interface for IGMP, or change the version an interface uses by right-clicking the interface and selecting **Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Multicast forwarding overview
Add the IGMP routing protocol
Enable IGMP router and IGMP proxy mode

# Configure multicast scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure multicast scopes**

1. Open Routing and Remote Access.
2. In the console tree, right-click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. Click **Properties**.
4. On the **Multicast Scopes** tab, click **Add**.
5. In the **Add Scoped Boundary** dialog box, do the following:

   - In **Scope name**, type the name of the multicast scope.
   - In **IP address**, type the base IP address of the range.
   - In **Mask**, type the mask of the range.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- You can also edit or remove an existing multicast scope by clicking **Edit** or **Remove** on the **Multicast Scopes** tab.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Multicast boundaries and heartbeat
Configure multicast boundaries
Expressing an IP address range with a mask

# Configure multicast boundaries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure multicast boundaries**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to configure multicast boundaries, and then click **Properties**.
4. On the **Multicast Boundaries** tab, do the following:

   - To add a scope to the list of multicast scope boundaries, in **Scope**, click a configured scope, and then click **Add**.
   - To remove a scope from the list of multicast scope boundaries, in **Scoped boundaries**, click the scope, and then click **Remove**.
   - To enable TTL scoping, select the **Activate TTL boundary** check box, and in **TTL value**, click the arrows to set a TTL value.
   - To set a rate limit, in **Rate limit (KBps)**, type the rate in kilobytes per second.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- Only scopes configured as multicast scopes are listed in **Scope**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Multicast boundaries and heartbeat
Configure multicast scopes

# Configure multicast heartbeat

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure multicast heartbeat**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, right-click the interface on which you want to configure multicast heartbeat, and then click **Properties**.
4. On the **Multicast Heartbeat** tab, select the **Enable multicast heartbeat detection** check box, and do the following:

   - To add a multicast heartbeat group, in **Multicast heartbeat group**, type the IP address of the multicast heartbeat group.
   - To configure a quiet time before alerting, in **Quiet time before alerting (minutes)**, click the arrows to set a quiet time in minutes.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Multicast boundaries and heartbeat

# Enable AppleTalk routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable AppleTalk routing**

1. Open Routing and Remote Access.
2. In the console tree, right-click **AppleTalk Routing**.

   **Where?**

   - Routing and Remote Access*server name*\AppleTalk Routing

3. Click **Enable AppleTalk Routing**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Unicast routing overview

# Connect Your Internal Network to the Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connect your internal network to the Internet**

- Connect to the Internet by using an ISDN device
- Connect to the Internet by using a modem
- Connect to the Internet by using a T1 line
- Connect to the Internet by using Frame Relay
- Connect to the Internet by using PPPoE

# Connect to the Internet by using an ISDN device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to the Internet by using an ISDN device**

1. Install the ISDN hardware in your computer.
2. On your desktop, right-click **My Computer**, click **Properties**, and on the **Hardware** tab, click **Add Hardware Wizard**.
3. After the ISDN adapter is added and the server is restarted, verify that the ISDN adapter appears as a port under **Ports** in Routing and Remote Access.
4. To configure the ISDN port, right-click **Ports**, and then click **Properties**.
5. In the **Port Properties** dialog box, select the ISDN device you want to configure, click **Configure**, select the **Demand-dial routing connections (inbound and outbound)** check box, and then click **OK**.
6. To create a demand-dial interface and configure it to use the ISDN device to dial in to the ISP, right-click **Network Interfaces**, and then click **New Demand-dial interface**.
7. Add a default route that uses the newly created demand-dial interface.

### Note

- A default route has a **Destination** of 0.0.0.0, a **Network mask** of 0.0.0.0, and a **Metric** of 1. Because the dial-up connection to the ISP is a point-to-point link, the **Gateway** IP address is not configurable.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Routed Connection to the Internet
Add a demand-dial interface
Add a default static IP route

# Connect to the Internet by using a modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to the Internet by using a modem**

1. Install the modem hardware on your computer.
2. Add the modem to your computer. In Control Panel, double-click **Phone and Modem Options**.
3. If you are prompted to provide location information, do so and then click **OK**.
4. On the **Modems** tab, click **Add**. Follow the steps in the Add Hardware Wizard.
5. After the modem is added, verify that the modem appears as a port under **Ports** in Routing and Remote Access.
6. To configure the modem port, right-click **Ports**, and then click **Properties**.
7. In the **Port Properties** dialog box, select the modem you want to configure, click **Configure**, select the **Demand-dial routing connections (inbound and outbound)** check box, and then click **OK**.
8. To create a demand-dial interface and configure it to use the modem to dial in to the ISP, right-click **Network Interfaces**, and then click **New Demand-dial interface**.
9. Add a default network route that uses the newly created demand-dial interface.

### Note

- A default network route has a **Destination** of 0.0.0.0, a **Network mask** of 0.0.0.0, and a **Metric** of 1. Because the dial-up connection to the ISP is a point-to-point link, the **Gateway** IP address is not configurable.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Routed Connection to the Internet
Add a demand-dial interface
Add a default static IP route

# Connect to the Internet by using a T1 line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to the Internet by using a T1 line**

1. Install the T1 adapter in your computer.
2. On your desktop, right-click **My Computer**, click **Properties**, and on the **Hardware** tab, click **Add Hardware Wizard**.
3. Configure the T1 adapter with the IP address allocated by your ISP.
4. After the T1 adapter is added and the server is restarted, verify that the T1 adapter appears as a network adapter under **Network Interfaces** in Routing and Remote Access.
5. Add a default network route that uses the T1 interface.

### Note

- A default network route has a **Destination** of 0.0.0.0, a **Network mask** of 0.0.0.0, and a **Metric** of 1. Because the T1 connection to the ISP is a point-to-point link, the **Gateway** IP address is not configurable.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Routed Connection to the Internet
Add a default static IP route

# Connect to the Internet by using Frame Relay

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to the Internet by using Frame Relay**

1. Install the Frame Relay adapter in your computer.
2. On your desktop, right-click **My Computer**, click **Properties**, and on the **Hardware** tab, click **Add Hardware Wizard**. Follow the prompts to configure the Frame Relay adapter.
3. After the Frame Relay adapter is added and the server restarted, verify that the Frame Relay adapter appears as a routing interface under **Network Interfaces** in Routing and Remote Access.
4. Add a default network route that uses the Frame Relay adapter.

### Note

- A default network route has a **Destination** of 0.0.0.0, **Network mask** of 0.0.0.0, and **Metric** of 1. Because the Frame Relay connection to the ISP is a point-to-point link, the **Gateway** IP address is not configurable.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Routed Connection to the Internet
Add a default static IP route

# Connect to the Internet by using PPPoE

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to the Internet by using PPPoE**

1. Open Routing and Remote Access.
2. In the console tree, click **Ports**.

   **Where?**

   - Routing and Remote Access/*ServerName*/Ports

3. Verify that the PPPoE adapter appears as a port.
4. In the console tree, right-click Network Interfaces, click **New Demand-dial Interface**, and follow the steps of the Demand Dial Interface wizard.
5. In **Connection Type**, click **Connect using PPP over Ethernet (PPPoE)**.
6. Complete the wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- The PPPoE adapter should appear as a port by default.
- Any interface you create using the Demand-Dial Interface Wizard appears under **Network Interfaces**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Internet Connection Sharing and network address translation
Routed Connection to the Internet
Add a demand-dial interface
Add a default static IP route
Point-to-Point Protocol over Ethernet (PPPoE) Protocols
Make a broadband connection using Point-to-Point Protocol over Ethernet (PPPoE)

# Gather Troubleshooting Information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Gather troubleshooting information**

- Determine the status of the Routing and Remote Access service
- Check the state of an interface
- Verify that OSPF is enabled
- Determine whether the router is receiving routes
- View RIP neighbors
- View OSPF neighbors
- Log details for a routing protocol
- Gather Demand-Dial Troubleshooting Information

**Note**

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Determine the status of the Routing and Remote Access service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To determine the status of the Routing and Remote Access service**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. In the details pane, for **Routing and Remote Access**, verify that the **Status** column is **Started**.
4. If the service is not started, right-click **Routing and Remote Access**, and then click **Start**.
5. If the router does not start, check the system event log for error messages.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Common routing problems

# Check the state of an interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check the state of an interface**

1. Open Routing and Remote Access.
2. In the console tree, click **General**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/General

3. In the details pane, click the interface you want to check, and then verify the values under the **Administrative Status** and **Operational Status** columns.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Common routing problems

# Verify that OSPF is enabled

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify that OSPF is enabled**

1. Open Routing and Remote Access.
2. In the console tree, click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. In the details pane, right-click the OSPF interface you want to verify, and then click **Properties**.
4. On the **General** tab, verify that the **Enable OSPF for this address** check box is selected.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Troubleshooting OSPF

# Determine whether the router is receiving routes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To determine whether the router is receiving routes**

1. Open Routing and Remote Access.
2. In the console tree, right-click **Static Routes**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/Static Routes

3. Click **Show IP routing table**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Common routing problems

# View RIP neighbors

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view RIP neighbors**

1. Open Routing and Remote Access.
2. In the console tree, right-click **RIP**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/RIP

3. Click **Show Neighbors**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Troubleshooting RIP for IP

# View OSPF neighbors

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view OSPF neighbors**

1. Open Routing and Remote Access.
2. In the console tree, right-click **OSPF**.

   **Where?**

   - Routing and Remote Access/*server name*/IP Routing/OSPF

3. Click **Show Neighbors**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Troubleshooting OSPF

# Log details for a routing protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To log details for a routing protocol**

1. Open Routing and Remote Access.
2. In the console tree, right-click the routing protocol for which you want to log details, and click **Properties**.

   **Where?**

   - Routing and Remote Access/*ServerName*/IP Routing/*RoutingProtocol*

3. Click the level of detail you want to log for this routing protocol, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- You can choose what level of detail to log for any protocol that appears under the IP Routing node in Routing and Remote Access.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Configure logging
Using event logging
Logging

# Gather Demand-Dial Troubleshooting Information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Gather demand-dial troubleshooting information**

- Check the unreachability reason
- Check the credentials being used by the calling router
- Check the authentication provider being used by the answering router
- Check the status of the port on the answering router
- Check the status of the demand-dial interface
- Check the static route for a demand-dial interface

# Check the unreachability reason

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check the unreachability reason**

1. Open Routing and Remote Access.
2. In the console tree, click **Network Interfaces**.

   **Where?**

   - Routing and Remote Access/*server name*/Network Interfaces

3. In the details pane, right-click the demand-dial interface you want to check, and then click **Unreachability reason**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Troubleshooting demand-dial routing

# Check the credentials being used by the calling router

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check the credentials being used by the calling router**

1. Open Routing and Remote Access.
2. In the console tree, click **Network Interfaces**.

   **Where?**

   - Routing and Remote Access\server name\Network Interfaces

3. In the details pane, right-click the demand-dial interface you want to check, and then click **Set Credentials**.
4. Verify the user name and domain name. If necessary, enter the credentials again.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Troubleshooting demand-dial routing

# Check the authentication provider being used by the answering router

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check the authentication provider being used by the answering router**

1. Open Routing and Remote Access.
2. In the console tree, right-click the router on which you want to check authentication, and then click **Properties**.
3. On the **Security** tab, verify that the router is configured with the correct security options.
4. If you are using Internet Protocol Security (IPSec), click **IPSec Settings** and verify that the IPSec settings are correct.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Troubleshooting demand-dial routing

# Check the status of the port on the answering router

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check the status of the port on the answering router**

1. Open Routing and Remote Access.
2. In the console tree, click **Ports**.

   **Where?**

   - Routing and Remote Access/*server name*/Ports

3. In the details pane, click the port used to answer the calling router and verify that the **Status** column is **Active**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Troubleshooting demand-dial routing

# Check the status of the demand-dial interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check the status of the demand-dial interface**

1. Open Routing and Remote Access.
2. In the console tree, click **Network Interfaces**.

   **Where?**

   - Routing and Remote Access/*server name*/Network Interfaces

3. In the details pane, in the **Connection State** column, verify that the demand-dial interface that matches the user name for the credentials of the calling router is **Connected**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Troubleshooting demand-dial routing

# Check the static route for a demand-dial interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check the static route for a demand-dial interface**

1. Open Routing and Remote Access.
2. In the console tree, click **Static Routes**.

   **Where?**

   - Routing and Remote Access/*ServerName*/IP Routing/Static Routes

3. In the **Interface** column, find the name of the demand-dial interface for which you want to check the static route.
4. Verify that the destination IP address, the network mask, and the gateway IP address (if any) are correct for the router to which you want to connect.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Troubleshooting demand-dial routing
Understanding Demand-Dial Routing
Setting Up a Static Routed IP Internetwork
Demand-dial routing

# Routing Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Routing Overview
- Understanding Routing
- Using Routing
- Routing Tools and Utilities

# Routing Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing overview**
Routing and Remote Access is a full-featured software router and an open platform for routing and internetworking. It offers routing services to businesses in local area network (LAN) and wide area network (WAN) environments or over the Internet by using secure virtual private network (VPN) connections.

An advantage of the Routing and Remote Access service is integration with the Microsoft® Windows Server 2003 family. The Routing and Remote Access service delivers many cost-saving features and works with a wide variety of hardware platforms and hundreds of network adapters. The Routing and Remote Access service is extensible with application programming interfaces (APIs) that developers can use to create custom networking solutions and that new vendors can use to participate in the growing business of open internetworking.

**Note**

- A computer running a member of the Windows Server 2003 family and the Routing and Remote Access service that provides LAN and WAN routing services is hereafter referred to as a server running Routing and Remote Access.

The server running Routing and Remote Access is intended for use by system administrators who are already familiar with routing protocols and routing services. Through Routing and Remote Access, administrators can view and manage both routers and remote access servers on their networks. For security information about enabling the routing features of Routing and Remote Access, see Security information for routing.

Routing and Remote Access includes the following features:

- Multiprotocol unicast routing for Internet Protocol (IP) and AppleTalk. For more information, see Unicast routing overview.
- Industry-standard unicast IP routing protocols:
  - Open Shortest Path First (OSPF) This feature is not available on the Itanium-based versions of the Windows operating systems. This content is not available in this preliminary release.
  - Routing Information Protocol (RIP) versions 1 and 2

  For more information, see Unicast routing overview.
- IP multicast services (Internet Group Management Protocol (IGMP) router mode and IGMP proxy mode) to enable the forwarding of IP multicast traffic. For more information, see Multicast Forwarding and Routing Overview.
- The Routing And Remote Access Server Setup Wizard, which includes a set of common server configurations to help meet your network needs. For more information, see Common server configurations for remote access servers.
- Support for multiple network interfaces. For more information, see Network interfaces.
- IP network address translation (NAT) services to simplify the connection of small office or home office (SOHO) networks to the Internet. For more information, see Internet Connection Sharing and network address translation.
- A simple packet filtering service that can be enabled for any public interface, even one that is also configured for network address translation. For more information, see Basic Firewall.
- Static IP packet filtering for security and performance. For more information, see Packet filtering.
- Demand-dial routing over dial-up WAN links. For more information, see Demand-dial routing.

- Virtual private network (VPN) support with the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol (L2TP) over Internet Protocol security (IPSec), referred to as L2TP/IPSec. For more information, see Routing over VPN connections.

- Industry standard support for a Dynamic Host Configuration Protocol (DHCP) Relay Agent for IP. For more information, see DHCP Relay Agent.

- Industry standard support for router advertisement by using Internet Control Message Protocol (ICMP) router discovery. For more information, see ICMP router discovery.

- A graphical user interface for remote monitoring and configuration. For more information, see Administration and management tools.

- A command-line interface for running scripts and automating configuration and remote monitoring. For more information, see Administration and management tools.

- Support for Windows power management capabilities. For more information, see Power management support.

- Simple Network Management Protocol (SNMP) management capabilities with support for popular management information bases (MIBs). For more information, see SNMP support.

- Extensive support for media, including Ethernet, token ring, Fiber Distributed Data Interface (FDDI), asynchronous transfer mode (ATM), Integrated Services Digital Network (ISDN), T-Carrier, Frame Relay, xDSL, cable modems, X.25, and analog modems.

- APIs for routing protocols, administration, and the user interface to enable value-added development.

# Unicast routing overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Unicast routing overview**

Unicast routing is the forwarding of traffic destined to a single location on an internetwork from a source host to a destination host by using routers. An internetwork is at least two networks that are connected by routers. A router is a network layer intermediate system that is used to connect networks together based on a common network layer protocol such as TCP/IP. A network is a portion of the networking infrastructure (encompassing repeaters, hubs, and bridges/Layer 2 switches) that is bounded by routers and is associated with the same network layer address known as a network address or network ID.

A typical router is connected to two or more networks over LAN or WAN media. Computers on a network can send packets to computers on other networks by forwarding the packets to the router. The router examines the packet and uses the destination network address in the packet header to decide which interface to use to forward the packet. Through routing protocols (OSPF, RIP, and others), a router learns network information (such as network addresses) from neighboring routers, and then propagates this information to routers on other networks to enable connectivity between all computers on all networks.

The server running Routing and Remote Access can route IP and AppleTalk traffic.

**IP routing**

The IP network protocol is part of a suite of Internet protocols known as Transmission Control Protocol/Internet Protocol (TCP/IP). IP is used to communicate across any set of interconnected IP networks. IP routers are either static routers (routes are established by an administrator and do not change until the administrator changes them) or dynamic routers (routes are updated dynamically by using routing protocols).

IP routing is the forwarding of IP traffic from a source host to a destination host through IP routers. At each router, the next hop is determined by matching the destination IP address within the packet with the best route in the routing table.

Routing and Remote Access includes support for two IP unicast routing protocols:

- Routing Information Protocol (RIP) for IP
- Open Shortest Path First (OSPF) This feature is not available on the Itanium-based versions of the Windows operating systems. This content is not available in this preliminary release.

Routing and Remote Access is not limited to supporting just RIP for IP and OSPF. The server running Routing and Remote Access is an extensible platform; other vendors can create additional IP routing protocols, such as Interior Gateway Routing Protocol (IGRP) and Border Gateway Protocol (BGP).

For more information about the IP unicast routing protocols supported by Routing and Remote Access, see RIP for IP and OSPF.

**AppleTalk routing**

AppleTalk is used primarily in Apple Macintosh environments.

Routing and Remote Access includes support for the Routing Table Maintenance Protocol (RTMP) routing protocol.

# Multicast Forwarding and Routing Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multicast forwarding and routing overview**
Unicasting is the sending of network traffic to a specific endpoint. Multicasting is the sending of network traffic to a group of endpoints. Only those members in the group of endpoints that are listening for the multicast traffic (the multicast group) process the multicast traffic. All other nodes ignore the multicast traffic.

You can use multicast traffic to:

- Discover resources on the internetwork.
- Support datacasting applications such as file distribution or database synchronization.
- Support multicasting multimedia applications such as digitized audio and video.

Multicast forwarding is the intelligent forwarding of multicast traffic. Multicast routing is the propagation of multicast group listening information. Routing and Remote Access supports multicast forwarding and limited multicast routing.

This section covers:

- Multicast forwarding overview
- Multicast routing overview

For more detailed information about multicasting, see Understanding Multicasting.

# Multicast forwarding overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multicast forwarding overview**
With multicast forwarding, a router forwards multicast traffic to networks where other multicast devices are listening. Multicast forwarding prevents the forwarding of multicast traffic to networks where there are no nodes listening.

For multicast forwarding to work across an internetwork, nodes and routers must be multicast-capable.

**Multicast-capable nodes**
A multicast-capable node must be able to:

- Send and receive multicast packets.
- Register the multicast addresses being listened to by the node with local routers, so that multicast packets can be forwarded to the network of the node.

All computers running a member of the Microsoft® Windows Server 2003 family are IP multicast-capable and can both send and receive IP multicast traffic. IP multicasting applications that send multicast traffic must construct IP packets with the appropriate IP multicast address as the destination IP address. IP multicasting applications that receive multicast traffic must inform the TCP/IP protocol that they are listening for all traffic to a specified IP multicast address.

IP nodes use the Internet Group Management Protocol (IGMP) to register their interest in receiving IP multicast traffic from IP routers. IP nodes that use IGMP send out an IGMP Membership Report message to notify their local routers that they are listening on a specific IP multicast address.

**Multicast-capable routers**
A multicast-capable router must be able to:

- Listen for all multicast traffic on all attached networks. Upon receiving multicast traffic, forward the multicast packet to attached networks where nodes are listening or where downstream routers have nodes that are listening.

  In the Windows Server 2003 family, the ability to listen for all multicast traffic and forward multicast packets is provided by the TCP/IP protocol. The TCP/IP protocol uses a multicast forwarding table to make decisions on where to forward incoming multicast traffic.

- Listen for IGMP Membership Report messages and update the TCP/IP multicast forwarding table.

  In the Windows Server 2003 family, the ability to listen for IGMP Membership Report messages and update the TCP/IP multicast forwarding table is provided by the IGMP routing protocol on an interface operating in IGMP router mode.

- Use a multicast routing protocol to propagate multicast group listening information to other multicast-capable routers.

  Routing and Remote Access does not provide any multicast routing protocols. However, the Routing and Remote Access service is an extensible platform and can support multicast routing protocols.

**Note**

- The IGMP routing protocol component provided with Routing and Remote Access is not a multicast routing protocol. The maintenance of entries in the IP multicast forwarding table is a function of IGMP.

**The IGMP routing protocol component**
Maintaining entries in the TCP/IP multicast forwarding table is implemented through the IGMP routing protocol component, a component that is added as an IP routing protocol by using Routing and Remote Access. Once the IGMP routing protocol component is added, router interfaces are added to IGMP. You can configure each interface added to the IGMP routing protocol component in one of two operating modes:

1. IGMP router mode

2. IGMP proxy mode

**IGMP router mode**
In the Windows Server 2003 family, the ability to listen for IGMP Membership Report packets and track group membership is provided by an interface running in IGMP router mode. You must enable IGMP router mode on the interfaces where listening multicast hosts are located.

For more information about the IGMP router mode, see Multicast forwarding.

**IGMP proxy mode**
An interface running in IGMP proxy mode acts as a proxy multicast host that sends IGMP Membership Report messages on one interface for IGMP Membership Report messages received on all other interfaces running in IGMP router mode. The upstream router attached to the network of the IGMP proxy mode interface receives the IGMP proxy mode Membership Report packets and adds them to its own multicast tables. In this way, the upstream router knows to forward multicast packets to the network segment of the IGMP proxy mode interface for multicast groups that were registered by hosts attached to network segments of the IGMP proxy mode router.

When the upstream router forwards multicast traffic to the IGMP proxy mode interface network, it is forwarded to the appropriate hosts on the IGMP router mode interface networks by the TCP/IP protocol.

All nonlocal multicast traffic received on all interfaces running IGMP router mode is forwarded using the IGMP proxy mode interface. The upstream router that receives the forwarded multicast traffic can elect to forward the multicast traffic or discard it. Using IGMP proxy mode, multicast sources on networks attached to the server running Routing and Remote Access can send multicast traffic to multicast hosts attached to upstream multicast routers.

IGMP proxy mode is designed to pass IGMP Membership Report messages from a single router intranet to a multicast-capable portion of the Internet. The multicast-capable portion of the Internet is known as the Internet multicast backbone, or MBone. With IGMP proxy mode enabled on the Internet interface, hosts on the single router intranet can receive multicast traffic from multicast sources on the Internet and send multicast traffic to hosts on the Internet.

For more information about using IGMP proxy mode interfaces, see Multicast routing.

# Multicast routing overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multicast routing overview**

With multicast routing, multicast-capable routers communicate multicast group membership information to each other so that intelligent multicast forwarding decisions are made across the internetwork. Multicast routers use multicast routing protocols to communicate multicast group information with each other.

Examples of multicast routing protocols include Distance Vector Multicast Routing Protocol (DVMRP), Multicast Extensions to OSPF (MOSPF), Protocol Independent Multicast-Sparse Mode (PIM-SM), and Protocol Independent Multicast-Dense Mode (PIM-DM). The Windows Server 2003 family does not include any multicast routing protocols. However, Routing and Remote Access is an extensible platform; other vendors can create multicast routing protocols.

**Note**

- Open Shortest Path First (OSPF) is not available on Microsoft® Windows® XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Security information for routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for routing**

Before you enable the routing functions of Routing and Remote Access, you should carefully review your network infrastructure so that you can configure Routing and Remote Access to best meet your needs for security and functionality. Security for Routing and Remote Access can be considered in three parts: securing the server running Routing and Remote Access, securing the network traffic between the server and its clients, and using secure authentication methods.

Unless your server running Routing and Remote Access is functioning only as a router, you should also consider security for your remote access solution. For more information, see Security information for remote access.

In addition to understanding the security implications of enabling Routing and Remote Access on your network, you should also practice general server security, as described in Security. For example, you should use only servers that are running Microsoft® Windows® 2000 or a member of the Microsoft® Windows Server 2003 family, if possible. You will decrease your network security if you use servers that are running Routing and Remote Access on the same network as servers that are running Microsoft Windows NT® 4.0 and either Remote Access Service (RAS) or Routing and Remote Access Service (RRAS). For more information, see Domain and forest functionality and Member server in a domain.

Before you configure and enable the routing features of Routing and Remote Access, you should consider the following:

- **Who can enable, configure, and disable Routing and Remote Access**. You must be a member of the Administrators group to use the Routing and Remote Access snap-in, either on its own or within the Microsoft Management Console (MMC). Similarly, you must be a member of the Administrators group to run most of the **netsh routing** commands from the command line. You should restrict the membership of the Administrators group to the minimum number of users necessary to administer the server. For more information, see Default local groups and Default groups.

- **How to open Routing and Remote Access**. As a security best practice, you should consider using the **runas** command to open Routing and Remote Access rather than logging on with administrative credentials. For more information about why you might not want to log on with administrative credentials, see Why you should not run your computer as an administrator. To open Routing and Remote Access when you are not logged on as a member of the Administrators group, type the following at the command prompt:

  **runas /user:**[*Domain/*]*UserName***"mmc %windir%\system32\rrasmgmt.msc"**

  The user name that you provide must correspond to the Administrator account or an account that is a member of the Administrators group, and you must provide the password for the account when prompted. For more information, see Runas.

- **Which routing components your network requires**. You should document how your network is designed, what routing components it requires, and how those components are configured. This information can help you maintain your network and identify potential areas requiring special concern or improvement.

- **Whether you can simplify your network topology**. Simpler networks are easier to maintain, and they provide fewer attack points. Simplifying your network might include:

  - **Minimizing the number of network interfaces**. Although most servers running Routing and Remote Access are multihomed computers, you should configure as few network interfaces as possible. If possible, you should not configure the server running Routing and Remote Access with more than one public interface. For more information, see Network interfaces.

  - **Minimizing the number of routes**. You should use as few demand-dial and static routes as possible. For more information, see Demand-dial routing and Static routing design considerations.

  - **Minimizing the number of routing protocols in use**. You should track what routing protocols are enabled and how they are configured, and then you should configure and enable as few routing protocols as possible. For

example, unless your network requires AppleTalk, you should not enable AppleTalk routing on the server running Routing and Remote Access.

- **Which tunneling protocols to use**. If you are configuring a router-to-router virtual private network (VPN), you should consider using Layer Two Tunneling Protocol (L2TP) instead of Point-to-Point Tunneling Protocol (PPTP). For more information, see VPN Router-to-Router tunneling protocols and Security information for remote access.

- **Whether to use packet filters, firewalls, and other network traffic protection**. You should restrict routing network traffic to the minimum level that your network requires. You should also consider using IP filters in your default remote access policy to ensure that inbound network traffic comes from authorized remote access clients.

    - In Routing and Remote Access, you can configure static packet filters on network interfaces and enable Basic Firewall on public interfaces. You can use these security features alone or in conjunction with other software, such as Internet Authentication Service (IAS). For more information, see Basic Firewall, Packet filtering, and Internet Authentication Service.

    - You cannot configure Routing and Remote Access if Internet Connection Firewall (ICF) or Windows Firewall is enabled. Before you configure Routing and Remote Access, you must disable Internet Connection Firewall or Windows Firewall. For more information about Windows Firewall, see Help: Windows Firewall.

- **What level of logging to use**. You should consider how much information you want to log for each routing protocol. Routing and Remote Access has several logging options. Additionally, you can use features in IAS for more detail and centralized accounting and auditing. For more information, see Logging and Internet Authentication Service.

# Common server configurations for remote access servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Common configurations for remote access servers**

When you run the Routing and Remote Access Server Setup Wizard, the wizard prompts you to choose the configuration path that most closely resembles the remote access solution that you want to deploy. If none of the wizard configuration paths meets your needs exactly, you can further configure your server after the wizard finishes, or you can choose the custom configuration path. However, if you choose the custom configuration path, you must manually configure all elements of Routing and Remote Access. The most common remote access solutions include dial-up connections, virtual private network (VPN) connections, and secure connections between two private networks.

**Remote access (dial-up)**




If you choose this path, the server running Routing and Remote Access is configured to allow remote access clients to connect to the private network by dialing into a modem bank or other dial-up equipment. To configure this type of server in the wizard, click **Remote Access**, select the **Dial-up** check box, and follow the steps. After the wizard finishes, you can configure additional options. For example, you can configure how the server answers the call, how the server verifies which remote access clients have permission to connect to the private network, and whether the server routes network traffic between remote access clients and the private network. For more information about dial-up connections, see Dial-Up Networking and Setting Up Dial-Up Remote Access.

**Remote access (VPN)**




If you choose this path, the server running Routing and Remote Access is configured to allow remote access clients to connect to the private network across the Internet. To configure this type of server in the wizard, click **Remote Access**, select the **VPN** check box, and follow the steps. After the wizard finishes, you can configure additional options. For example, you can configure how the server verifies which VPN clients have permission to connect to the private network and whether the server routes network traffic between VPN clients and the private network. For more information about VPNs, see Deploying VPNs for Remote Access and Common configuration for the VPN server.

**Network address translation (NAT)**




If you choose this path, the server running Routing and Remote Access is configured to share an Internet connection with

computers on the private network and to translate traffic between its public address and the private network. Computers on the Internet will not be able to determine the IP addresses of computers on the private network. To configure this type of server in the wizard, click **Network address translation (NAT)**, and follow the steps. After the wizard finishes, you can configure additional options. For example, you can configure packet filters and choose what services to allow on the public interface. For more information about NAT, see Setting Up Network Address Translation and Translated Connection to the Internet.

**VPN and NAT**




If you choose this path, the server running Routing and Remote Access is configured to provide NAT for the private network and to accept VPN connections. Computers on the Internet will not be able to determine the IP addresses of computers on the private network. However, VPN clients will be able to connect to computers on the private network as if they were physically attached to the same network. To configure this type of server in the wizard, click **Virtual Private Network (VPN) access and NAT**, and follow the steps. For more information about NAT and VPN, see New features for virtual private networks and Network address translation design considerations.

**Secure connection between two private networks**




If you choose this path, two servers running Routing and Remote Access are configured to send private data securely across the Internet. You must choose this path when you run the Routing and Remote Access Server Setup Wizard on each server. The connection between the two servers can be persistent (always on) or on demand (demand-dial). To configure this type of server in the wizard, click **Secure connection between two private networks**, and follow the steps. After the wizard finishes, you can configure each server with additional options. For example, you can configure what routing protocols each server accepts and how each server routes traffic between the two networks. For more information about connections between two networks, see Deploying Router-to-Router VPNs, Branch Office over the Internet, and Dial-Up Branch Office Network.

# Network interfaces

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network interfaces**
Network interfaces enable any server running Routing and Remote Access to communicate with other computers over private or public networks. Network interfaces have two aspects that relate to Routing and Remote Access: the physical hardware, such as a network adapter, and the network interface configuration.

**Physical hardware**
Network interface hardware, commonly called network adapters, can be any adapter that plugs into the system bus of a computer and allows that computer to connect to a network. Most servers running Routing and Remote Access have at least two network adapters. This is required if the server running Routing and Remote Access is acting as a router between two network segments. A server that has two or more network adapters connecting to separate networks is called multihomed.

**Important**

- Because multihomed servers are accessible from more than one network, they have more security requirements than a server that is connected to a single network. Each network interface must be configured correctly in order to protect the server and the private networks to which the server connects. For more information, see Security information for routing and Security information for remote access.

The server running Routing and Remote Access usually detects all network adapters automatically when the Routing and Remote Access Server Setup Wizard is run. If you install an additional network adapter after the wizard has been run, you can add and configure the interface using Routing and Remote Access. For more information, see Add a routing interface.

**Network interface configuration**
In Routing and Remote Access, network interfaces fall into the following categories:

- **Private interface**. A private interface is a network adapter that is physically connected to a private network. Most private networks are configured with a private network IP address range, and the private interface is also configured with a private address. Because a private network is, in theory, composed of known users and computers, you generally have fewer security considerations for a private interface than for a public interface. For more information, see Internet private addresses.

- **Public interface**. A public interface is a network adapter that is physically connected to a public network, such as the Internet. The public interface is configured with a public IP address. You can configure a public interface to perform network address translation (NAT). Because a public interface is theoretically accessible by anyone on the public network, security considerations are generally higher for a public interface than for a private interface.

- **Demand-dial interface**. Demand-dial interfaces connect specific routers on either public or private networks. A demand-dial interface can be either on-demand (activated only when needed) or persistent (always connected). For more information, see Demand-dial routing.

In addition to configuring each network interface as a public, private, or demand-dial interface, you can configure packet filters, addresses, and other options for network interfaces. Some options for public interfaces, such as Basic Firewall, are not available for private interfaces. For more information about the kind of configuration you might want to consider and example implementations, see Routing Scenarios and Remote Access Scenarios.

# Internet Connection Sharing and network address translation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Connection Sharing and network address translation**
To connect a small office or home office (SOHO) network to the Internet, you can use one of two methods:

1. Routed connection

   For a routed connection, the server running Routing and Remote Access acts as an IP router that forwards packets between SOHO hosts and Internet hosts. While conceptually simple, a routed connection requires knowledge of IP address and routing configuration for SOHO hosts and the server running Routing and Remote Access. However, routed connections allow all IP traffic between SOHO hosts and Internet hosts. For more information, see SOHO Network to the Internet.

2. Translated connection

   For a translated connection, the server running Routing and Remote Access acts as an network address translator: an IP router that translates addresses for packets being forwarded between SOHO hosts and Internet hosts. This is called network address translation, or NAT. Translated connections that use a server running Routing and Remote Access require less knowledge of IP addressing and routing and provide a simplified configuration for SOHO hosts. However, translated connections might not allow all IP traffic between SOHO hosts and Internet hosts.

You can configure a translated connection to the Internet by using one of the following methods.

- You can use the Internet Connection Sharing feature of Network Connections. Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- You can use the Network Address Translation (NAT) routing protocol provided with the Routing and Remote Access snap-in on servers running Windows Server 2003 and Routing and Remote Access.

Both Internet Connection Sharing and NAT provide translation, addressing, and name resolution services to SOHO hosts.

Internet Connection Sharing is designed to provide a single step of configuration (a single check box) in order to provide a translated connection to the Internet for all of the computers on the SOHO network. However, once enabled, Internet Connection Sharing does not allow further configuration beyond the configuration of services and ports on the SOHO network. For example, Internet Connection Sharing is designed for a single IP address obtained from an Internet service provider (ISP) and does not allow you to change the range of IP addresses allocated to SOHO hosts. For more information, see Connecting to the Internet in a home or small office network.

The NAT routing protocol component is designed to provide maximum flexibility in the configuration of the server running Routing and Remote Access to provide a translated connection to Internet. The NAT routing protocol component requires more configuration steps; however, each step of the configuration is customizable. Most of the configuration can be accomplished using the Routing and Remote Access server Setup Wizard. The NAT routing protocol component allows for ranges of IP addresses from an ISP and the configuration of the range of IP addresses allocated to SOHO hosts. For more information, see Understanding Network Address Translation.

The following table summarizes the features and capabilities of Internet Connection Sharing and the NAT routing protocol component.

| **Internet Connection Sharing** | **Network address translation** |
|---|---|
| Single check box configuration | Manual configuration |
| Single public IP address | Multiple public IP addresses |

| | |
|---|---|
| Fixed address range for SOHO hosts | Configurable address range for SOHO hosts |
| Single SOHO interface | Multiple SOHO interfaces |

**Notes**

- Internet Connection Sharing is a feature designed to connect SOHO networks to the Internet. The NAT routing protocol component is designed to connect small to medium business networks to the Internet. Neither Internet Connection Sharing nor the NAT routing protocol component is designed to:
    - Directly connect SOHO networks.
    - Connect networks within an intranet.
    - Directly connect branch office networks to a corporate network.
    - Connect branch office networks to a corporate network over the Internet.

# Basic Firewall

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing your network with Basic Firewall**
You can use Basic Firewall to help secure your network from unsolicited public network traffic, such as traffic sent from the Internet. People who send such traffic might be trying to access your network without your permission. You can enable Basic Firewall for any public interface, including one that also provides network address translation (NAT) for your network.

**How Basic Firewall works**
Basic Firewall is a stateful firewall, which combines dynamic packet filtering of network traffic with a set of static packet filters. Basic Firewall monitors traffic that travels through the interface for which Basic Firewall is enabled.

If the interface is configured for private network traffic only, Basic Firewall will route traffic among the computers on the private network only. Basic Firewall will route traffic between a private network and virtual private network (VPN) client computers. However, computers that are part of a private network will not be able to detect computers outside of the private network, and computers that are not part of the private network will not be able to detect computers that belong to the private network.

If the interface is configured for private network traffic and to provide NAT, each packet's source and destination addresses are recorded in a table. All traffic from the public network is compared to the entries in the table. Traffic from the public network can reach the private network only if the table contains an entry that shows that the communication exchange originated from within the private network. In this way, Basic Firewall prevents unsolicited traffic from a public network (such as the Internet) from reaching a private network.

For more information about packet filters and how they work, see Packet filtering. For information on configuring Basic Firewall, see Configure Basic Firewall for a public interface.

**Considerations when using Basic Firewall**
You do not need to use Basic Firewall if your network has other firewall software installed. Basic Firewall can provide protection only for public interfaces. You cannot enable Basic Firewall on a private interface. You must configure each instance of Basic Firewall individually to work correctly with all services and ports on that interface. For more information, see Network interfaces.

You can configure exceptions for certain types of traffic if Basic Firewall blocks traffic that your network should receive. For example, some programs, such as e-mail programs, might behave differently if Basic Firewall is enabled. For information on configuring exceptions, see Configure services and ports, Packet filtering, and Allow or deny ICMP messages.

You can further protect your private network by configuring static packet filters on private interfaces. Unlike Basic Firewall, static packet filters can be configured on private interfaces.

You can configure Basic Firewall when you configure and enable Routing and Remote Access or after Routing and Remote Access has been enabled. You cannot enable Internet Connection Firewall or Windows Firewall if Routing and Remote Access is configured and enabled. You cannot configure and enable Routing and Remote Access if Internet Connection Firewall or Windows Firewall is enabled. For more information about Windows Firewall, see Help: Windows Firewall.

**Basic Firewall logging and settings**
By default, Basic Firewall does not log an event when it discards unsolicited network traffic. Notifications of discarded traffic can be sent frequently enough to be a distraction to the network administrator. You can enable logging if you want to be notified when Basic Firewall discards traffic. For example, the default configuration of Basic Firewall does not allow echo requests sent from the Internet to reach a private network. If you enable logging and you do not allow echo requests to reach your network, the echo request fails, and a log entry that notes the failed attempt is generated. For more information, see Configure logging and Log details for a routing protocol.

# Packet filtering

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Packet filtering**
Routing and Remote Access supports IP packet filtering, which specifies what type of traffic is allowed into and out of the router. The packet filtering feature is based on exceptions. You can set packet filters per interface and configure them to do one of the following:

- Pass through all traffic except packets prohibited by filters.
- Discard all traffic except packets allowed by filters.

For more information about packet filtering, see Manage Packet Filters.

For detailed information about IP packet filtering, including examples of filtering configurations and fragmentation filtering, see "Part One: Routing" at the Microsoft® Windows® Resource Kits Web site.

**Protocols and Ports for Packet Filtering**
The following table details some of the common ports and protocols that you might want to allow, depending on your remote access configuration. Not all ports listed here might be required for your remote access server. For example, if you are allowing only Layer Two Tunneling Protocol (L2TP), you would not configure a filter for Point-to-Point Tunneling Protocol (PPTP). Similarly, this table might not contain all of the ports that your specific network needs.

**Ports Used for Protocols**

| Protocol | Port | Used For |
|---|---|---|
| TCP | 25 | Simple Mail Transfer Protocol (SMTP) |
| TCP | 67 | Dynamic Host Control Protocol (DHCP) (if the remote access server uses an external DHCP server) |
| TCP | 80 | World Wide Web (HyperText Transfer Protocol (HTTP)) |
| TCP | 110 | Post Office Protocol, version 3 (POP3) |
| TCP | 1701 | L2TP |
| TCP | 1723 | PPTP |
| TCP | 7250 | Network Access Quarantine Control (Remote Access Quarantine Client (RQC) messages from client computers) |
| UDP | 53 | Domain Name Service (DNS) (for name resolution of external Web sites) |
| UDP | 67 | DHCP (if the remote access server uses an external DHCP server) |
| UDP | 500 | Internet Protocol Security (IPSec) |
| UDP | 1701 | L2TP |
| UDP | 1723 | PPTP |
| UDP | 4500 | IPSec with network address translation (NAT) |
| | 47 | Generic Routing Encapsulation (GRE) |
| | 50 | Encapsulating Security Payload (ESP) (for firewalls that use NAT traversal [NAT-T]) |

**Note**

- To support Windows Update, you must allow TCP traffic to travel inbound and outbound on port 80 and UDP traffic to

travel inbound and outbound on port 53. Depending on your network configuration, you might have to configure these filters on your remote access server, on your firewall, or both.

**See Also**
**Concepts**
Basic Firewall
Help: Windows Firewall
Set L2TP over IPSec input filters
Set PPTP input filters

# Demand-dial routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Demand-dial routing**

Routing and Remote Access also includes support for demand-dial routing (also known as dial-on-demand routing). By using a demand-dial interface, the router can initiate a connection to a remote site when the packet to be routed is received by the router. The connection becomes active only when data is sent to the remote site. When no data has been sent over the link for a specified amount of time, the link is disconnected. By making a demand-dial connection, you can use existing dial-up telephone lines instead of leased lines for low-traffic situations. Demand-dial routing can significantly reduce your connection costs.

Routing and Remote Access also includes support for demand-dial filters and dial-out hours.

You can use demand-dial filters to specify what types of traffic are allowed to create the connection. Demand-dial filters are separate from IP packet filters, which you configure to specify what traffic is allowed into and out of an interface once the connection is made. For more information about configuring demand-dial filters, see Configure demand-dial filters.

You can set dial-out hours to specify the hours that a router is allowed to dial out to make demand-dial connections. For more information about configuring dial-out hours, see Configure dial-out hours.

You can configure when the router accepts incoming connections through remote access policies. For more information, see Introduction to remote access policies.

For more information about demand-dial routing and an example, see Understanding Demand-Dial Routing.

# Routing over VPN connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing over VPN connections**

Conventional routing occurs between routers over either LAN-based shared access technologies, such as Ethernet or token ring, or WAN-based point-to-point technologies, such as T1 or Frame Relay. With conventional WAN technologies, IP packets are forwarded between two routers over a physical or logical point-to-point connection. This connection is dedicated to the customer across a private data network that is provided by the WAN service provider.

With the advent of the Internet, you can now route packets between routers that are connected to the Internet across a virtual connection that emulates the properties of a dedicated, private, point-to-point connection. This type of connection is known as a *router-to-router virtual private network (VPN)* connection. With router-to-router VPN connections, you can replace expensive long-haul WAN links with short-haul WAN links to your local Internet service provider (ISP).

To emulate a private, point-to-point connection, a packet that is forwarded between routers is encapsulated, or wrapped, with an additional header that provides routing information that is needed to reach the endpoint. The endpoints of the connection are the routers. The portion of the virtual private networking connection in which your data is encapsulated is called the tunnel.

For secure VPN connections, the data portion of your packets is encrypted. Intercepted packets are undecipherable without the encryption keys. The portion of the virtual private networking connection in which your data is encrypted is called the *virtual private network (VPN) connection*. In router-to-router VPN connections, the tunnel and the VPN connection are the same.

For general information about VPN technology, see Virtual Private Networks.

For more information about router-to-router VPNs, see Understanding Router-to-Router VPNs.

For information about designing and deploying a router-to-router VPN connection, see Deploying Router-to-Router VPNs.

For an example of a router-to-router VPN connection, see Branch Office over the Internet.

**Note**

- On Microsoft® Windows Server™ 2003, Web Edition, and Microsoft Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# DHCP Relay Agent

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP Relay Agent**
The DHCP Relay Agent component is a Bootstrap Protocol (BOOTP) relay agent that relays Dynamic Host Configuration Protocol (DHCP) messages between DHCP clients and DHCP servers on different IP networks. The DHCP Relay Agent is compliant with RFC 1542, "Clarifications and Extensions for the Bootstrap Protocol." For each IP network segment that contains DHCP clients, either a DHCP server or a computer acting as a DHCP Relay Agent is required.

For information about installing and configuring the DHCP Relay Agent component, see Configure the DHCP Relay Agent.

### Notes

- You cannot use the DHCP Relay Agent component on a computer running the DHCP service, the network address translation (NAT) routing protocol component with automatic addressing enabled, or Internet Connection Sharing (ICS).

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# ICMP router discovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ICMP router discovery**

To configure IP hosts with the IP addresses of local routers and provide a way for hosts to sense routers that are down, you can use Internet Control Message Protocol (ICMP) messages for router solicitation and advertisement (as described in RFC 1256, "ICMP Router Discovery Messages"):

- Router solicitations are sent by hosts to discover routers on their networks.
- Router advertisements are sent by routers in response to a router solicitation and periodically to notify hosts on the network that the router is still available.

The Microsoft® Windows Server 2003 family implementation of the TCP/IP protocol supports ICMP router solicitations and the receipt of ICMP router advertisements, but they are disabled by default. Routing and Remote Access supports ICMP router advertisements. For more information about enabling ICMP router solicitation, see Enable ICMP router discovery.

# Administration and management tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administration and management tools**

Routing and Remote Access includes the following routing administration and management tools:

- Routing and Remote Access

  The primary administration and management tool for configuring routing is Routing and Remote Access (available in the Administrative Tools folder in Control Panel).

  With Routing and Remote Access, you can:

  - Configure and enable, disable, start, and stop the Routing and Remote Access service.
  - Manage local and remote routers.
  - Configure routing behavior and routing protocols.
  - Configure demand-dial routing behavior.
  - Configure remote access policies.
  - Monitor active connections.
  - Configure local logging behavior.

  Members of the Administrators group can use Routing and Remote Access. If the computer is a member of a domain, members of the Domain Administrators group can use Routing and Remote Access.

- The Netsh command-line utility

  You can use the Netsh command-line utility to locally configure routing from the command prompt. The Netsh utility can use script files to automate configuration tasks.

  For more information about the Netsh utility, see The Netsh Command-Line Utility.

**Notes**

- You cannot use a server running a member of the Microsoft® Windows Server 2003 family to administer servers running Microsoft Windows NT® 4.0 and Remote Access Service (RAS) or to administer servers running Windows NT 4.0 and Routing and Remote Access Service (RRAS). For more information, see Administration tools in mixed environments.

- You cannot remotely configure AppleTalk routing properties with Routing and Remote Access.

# Power management support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Power management support**

Power management support for Routing and Remote Access is listed in the following table. The first column indicates the configuration of the server running Routing and Remote Access, the second column indicates whether the computer goes into sleep mode based on an Advanced Configuration and Power Interface (ACPI) or Advanced Power Management (APM)-initiated sleep mode request, and the last column indicates whether the computer goes into sleep mode based on a user-initiated sleep mode request.

| Configuration of the server running Routing and Remote Access | ACPI/APM-initiated sleep mode | User-initiated sleep mode |
|---|---|---|
| LAN and WAN router | No | Yes |
| LAN only router | No | No |
| Router with an active VPN connection | No | Yes |
| Router with active network address translation | No | Yes |
| Remote access server | No | Yes |
| Remote access server with an active dial-up connection | No | Yes |
| Remote access server with an active VPN connection | No | Yes |

**Note**

- If there are active connections when a user-initiated sleep mode request is made, you are not informed. The active connections are terminated.

# SNMP support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SNMP support**

If you want the Routing and Remote Access service to participate in a Simple Network Management Protocol (SNMP) environment as an SNMP agent, you need to install the SNMP service. For more information, see Install the SNMP service.

Routing and Remote Access supports the following management information bases (MIBs):

- Internet MIB II

  Objects in the Internet MIB II are documented in RFC 1213, "Management Information Base for Network Management of TCP/IP-based internets: MIB-II."

- IP Forwarding Table MIB

  Objects in the IP Forwarding Table MIB are documented in RFC 1354, "IP Forwarding Table MIB."

- Microsoft RIP Version 2 for Internet Protocol MIB

- Wellfleet-Series7-MIB for OSPF. Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

- Microsoft BOOTP for Internet Protocol MIB

- Internet Group Management Protocol MIB

  Objects in the Internet Group Management Protocol MIB are documented in RFC 2933, "Internet Group Management Protocol MIB."

For more information about the SNMP service, see SNMP.

# Understanding Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding routing**

The routing services included with Windows NT Server 4.0, MultiProtocol Routing (MPR) version 1.0, provided a limited set of LAN-based routing services suitable for smaller organizations and branch offices. In June 1996, Microsoft offered the Routing and Remote Access Service as a free download component from the Microsoft Web site. The Routing and Remote Access Service for Windows NT 4.0, also known as RRAS, integrated routing and remote access services, extended the internetworking and routing capabilities available in MPR 1.0, and enabled routing over WANs and dial-up links. Windows 2000 offered Routing and Remote Access, a new service distinct from RRAS, that included expanded dial-in and VPN networking capabilities. The version of Routing and Remote Access included with the Windows Server 2003 family extends the capabilities of the Windows 2000 Routing and Remote Access service with simplified configuration for demand-dial routing, dynamic as well as static packet filtering, IPSec translation over network address translation (NAT) interfaces, and many other features.

This section covers:

- Understanding Unicast Routing
- Understanding Multicasting
- Understanding Network Address Translation
- Understanding Demand-Dial Routing
- Understanding Router-to-Router VPNs

# Understanding Unicast Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding unicast routing**
Unicast routing, the forwarding of traffic to a specific destination address, is facilitated by knowing where traffic can be delivered in the internetwork. At each hop in the path from the source to the destination, a routing decision is made.

This section covers:

- Routing tables
- Routing Configurations
- IP Routing Protocols
- Routing interfaces, devices, and ports

# Routing tables

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing tables**

The routing decision is aided by knowing which network addresses (or network IDs) are available in the internetwork. This knowledge is obtained from a database called the routing table. The routing table is a series of entries called routes that contain information on where the network IDs of the internetwork are located. The routing table is not exclusive to a router. Hosts (non-routers) may also have a routing table that is used to determine the optimal route.

**Types of routing table entries**

Each entry in the routing table is considered a route and is one of the following types:

- Network route

  A network route provides a route to a specific network ID in the internetwork.

- Host route

  A host route provides a route to an internetwork address (network ID and node ID). Host routes are typically used to create custom routes to specific hosts to control or optimize network traffic.

- Default route

  A default route is used when no other routes in the routing table are found. For example, if a router or host cannot find a network route or host route for the destination, the default route is used. The default route simplifies the configuration of hosts. Rather than configuring hosts with routes for all the network IDs in the internetwork, a single default route is used to forward all packets with a destination network or internetwork address that was not found in the routing table.

**Routing table structure**

The following illustration shows the structure of the routing table.




Each entry in the routing table consists of the following information fields:

- Network ID

  The network ID or an internetwork address for a host route. On IP routers, there is an additional subnet mask field that determines the IP network ID from a destination IP address.

- Forwarding address

  The address to which the packet is forwarded. The forwarding address is a hardware address or an internetwork address. For networks to which the host or router is directly attached, the forwarding address field may be the address of the interface that is attached to the network.

- Interface

  The network interface that is used when packets are forwarded to the network ID. This is a port number or other type of

logical identifier.

- Metric

  A measurement of the preference of a route. Typically, the lowest metric is the most preferred route. If multiple routes exist to a given destination network, the route with the lowest metric is used. Some routing algorithms only store a single route to any network ID in the routing table, even when multiple routes exist. In this case, the metric is used by the router to determine which route to store in the routing table.

**Note**

- The preceding list is intended to be a representative list of fields in the routing tables used by routers. Actual fields in the routing tables for different routable protocols may vary.

# Routing Configurations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing configurations**
You can use routers in many different topologies and network configurations. When you add a server running Routing and Remote Access configured as a router to your network, you must select:

- The protocols to be routed (IP or AppleTalk) by the router.
- Routing protocols (RIP or OSPF) for each protocol to be routed.
- LAN or WAN media (network adapters, modems, or other dial-up equipment).

This section describes three typical routing scenarios:

- A simple routing scenario
- A multiple-router scenario
- A demand-dial routing scenario

For more detailed routing configurations, see Routing Scenarios.

**Note**

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# A simple routing scenario

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**A simple routing scenario**

The following illustration shows a simple network configuration with a server running Routing and Remote Access that connects two LAN segments (Networks A and B). In this configuration, routing protocols are not necessary because the router is connected to all the networks to which it needs to route packets.



# A multiple-router scenario

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**A multiple-router scenario**

The following illustration shows a more complex router configuration.




In this configuration, there are three networks (Networks A, B, and C) and two routers (Routers 1 and 2). Router 1 is on Networks A and B, and Router 2 is on Networks B and C. Router 1 must notify Router 2 that Network A can be reached through Router 1, and Router 2 must notify Router 1 that Network C can be reached through Router 2. This information is automatically communicated through the use of routing protocols, such as RIP or OSPF. When a user on Network A wants to communicate with a user on Network C, the user's computer on Network A forwards the packet to Router 1. Router 1 then forwards the packet to Router 2. Router 2 then forwards the packet to the user's computer on Network C.

Without the use of routing protocols, a network administrator has to enter static routes into the routing tables of Router 1 and Router 2. While static routes work, they do not scale well to larger internetworks or recover from changes in the internetwork topology.

**Note**

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# A demand-dial routing scenario

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**A demand-dial routing scenario**

The following illustration shows a router configuration that uses demand-dialing.




Networks A and B are geographically separated and, for the amount of traffic that is transferred between the networks, a leased WAN link is not economical. Router 1 and Router 2 can connect over an analog phone line by using modems on both ends (or another type of connectivity, such as ISDN). When a computer on Network A initiates communication with a computer on Network B, Router 1 establishes a phone connection with Router 2. The modem connection is maintained as long as there are packets going back and forth. When the connection is idle, Router 1 hangs up to reduce connection costs.

For a detailed explanation of what happens during a demand-dial connection, see Understanding Demand-Dial Routing.

# IP Routing Protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP routing protocols**
In dynamic IP routing environments, IP routing information is propagated by using IP routing protocols. The two most common IP routing protocols that are used on intranets are Routing Information Protocol (RIP) and Open Shortest Path First (OSPF).

You can run multiple routing protocols on the same intranet. In this case, you must configure which routing protocol is the most preferred source of learned routes by configuring preference levels. The preferred routing protocol is the source of the route that is added to the routing table, regardless of the metric of the learned route. For example, if the metric of an OSPF learned route is 5 and the metric of the corresponding RIP learned route is 3 and OSPF is the preferred routing protocol, then the OSPF route is added to the IP routing table and the RIP route is ignored. For more information, see Change preference levels.

This section provides information about:

- RIP for IP
- OSPF

**Notes**

- If you are using multiple IP routing protocols, configure only a single routing protocol per interface.
- OSPF is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# RIP for IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RIP for IP**

The Routing Information Protocol (RIP) is designed for exchanging routing information within a small to medium-size internetwork.

The biggest advantage of RIP is that it is extremely simple to configure and deploy. The biggest disadvantage of RIP is its inability to scale to large or very large internetworks. The maximum hop count used by RIP routers is 15. Networks that are 16 hops or more away are considered unreachable. As internetworks grow larger in size, the periodic announcements by each RIP router can cause excessive traffic. Another disadvantage of RIP is its high recovery time. When the internetwork topology changes, it may take several minutes before the RIP routers reconfigure themselves to the new internetwork topology. While the internetwork reconfigures itself, routing loops may form that result in lost or undeliverable data.

Initially, the routing table for each router includes only the networks that are physically connected. A RIP router periodically sends announcements that contain its routing table entries to inform other local RIP routers of the networks it can reach. RIP version 1 uses IP broadcast packets for its announcements. RIP version 2 uses multicast or broadcast packets for its announcements.

RIP routers can also communicate routing information through triggered updates. Triggered updates occur when the network topology changes and updated routing information is sent that reflects those changes. With triggered updates, the update is sent immediately rather than waiting for the next periodic announcement. For example, when a router detects a link or router failure, it updates its own routing table and sends updated routes. Each router that receives the triggered update modifies its own routing table and propagates the change.

Routing and Remote Access supports RIP versions 1 and 2. RIP version 2 supports multicast announcements, simple password authentication, and more flexibility in subnetted and Classless InterDomain Routing (CIDR) environments.

The Windows Server 2003 family implementation of RIP has the following features:

- Selection of which RIP version to run on each interface for incoming and outgoing packets.
- Split-horizon, poison-reverse, and triggered-update algorithms that are used to avoid routing loops and speed recovery of the internetwork when topology changes occur.
- Route filters for choosing which networks to announce or accept.
- Peer filters for choosing which router's announcements are accepted.
- Configurable announcement and route aging timers.
- Simple password authentication support.
- The ability to disable subnet summarization.

For information about designing and deploying a RIP-for-IP internetwork, see Setting Up a RIP-for-IP Routed Internetwork.

For detailed information about the operation of RIP for IP, see "Unicast IP Routing" at the Microsoft Windows Resource Kits Web site.

**Note**

- If you are using multiple IP routing protocols, configure only a single routing protocol per interface.

# OSPF

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**OSPF**

Open Shortest Path First (OSPF) is designed for exchanging routing information within a large or very large internetwork. This feature is not available on the Itanium-based versions of the Windows operating systems. This content is not available in this preliminary release.

The biggest advantage of OSPF is that it is efficient; OSPF requires very little network overhead even in very large internetworks. The biggest disadvantage of OSPF is its complexity; OSPF requires proper planning and is more difficult to configure and administer.

OSPF uses a Shortest Path First (SPF) algorithm to compute routes in the routing table. The SPF algorithm computes the shortest (least cost) path between the router and all the networks of the internetwork. SPF-calculated routes are always loop-free.

Instead of exchanging routing table entries like RIP routers, OSPF routers maintain a map of the internetwork that is updated after any change to the network topology. This map, called the link state database, is synchronized between all the OSPF routers and is used to compute the routes in the routing table. Neighboring OSPF routers form an adjacency, which is a logical relationship between routers to synchronize the link state database.

Changes to internetwork topology are efficiently flooded across the entire internetwork to ensure that the link state database on each router is synchronized and accurate at all times. Upon receiving changes to the link state database, the routing table is recalculated.

As the size of the link state database increases, memory requirements and route computation times increase. To address this scaling problem, OSPF divides the internetwork into areas (collections of contiguous networks) that are connected to each other through a backbone area. Each router only keeps a link state database for those areas that are connected to the router. Area border routers (ABRs) connect the backbone area to other areas.

To further reduce the amount of routing information flooded into areas, OSPF allows the use of stub areas. A stub area can contain a single entry and exit point (a single ABR), or multiple ABRs when any of the ABRS can be used to reach external route destinations.

The following illustration shows a diagram of an OSPF internetwork.




OSPF has the following advantages over RIP:

- OSPF-calculated routes are always loop-free.
- OSPF can scale to large or very large internetworks.
- Reconfiguration for network topology changes is faster.

The implementation of OSPF in Routing and Remote Access has the following features:

- Route filters for controlling interaction with other routing protocols.
- Dynamic reconfiguration of all OSPF settings.

- Coexistence with RIP.
- Dynamic addition and deletion of interfaces.

For information about designing and deploying an OSPF internetwork, see Setting Up an OSPF Routed Internetwork.

For detailed information about the operation of OSPF, see "Unicast IP Routing" at the Microsoft Windows Resource Kits Web site.

**Notes**

- Routing and Remote Access does not support the use of OSPF in a demand-dial configuration that uses nonpermanent, dial-up links.
- If you are using multiple IP routing protocols, configure only a single routing protocol per interface.

# Routing interfaces, devices, and ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing interfaces, devices, and ports**
Routing and Remote Access views the installed networking equipment as a series of routing interfaces, devices, and ports.

- A routing interface is a physical or logical interface over which unicast or multicast packets are forwarded.
- A device represents hardware or software that creates physical or logical point-to-point connections.
- A port is a communication channel that supports a single point-to-point connection.

**Routing interface**
The server running Routing and Remote Access uses a routing interface to forward unicast IP or AppleTalk packets and multicast IP packets. There are two types of routing interfaces:

- LAN interfaces

  A LAN interface is a physical interface that typically represents a local area connection that uses local area networking technology such as Ethernet or token ring. A LAN interface reflects an installed network adapter. An installed WAN adapter is sometimes represented as a LAN interface. For example, some Frame Relay adapters create a separate logical LAN interface for each configured virtual circuit. LAN interfaces are always active and typically do not require an authentication process to become active.

- Demand-dial interfaces

  A demand-dial interface is a logical interface that represents a point-to-point connection. The point-to-point connection is based on either a physical connection, such as two routers connected over an analog phone line that uses modems, or a logical connection, such as two routers connected over a virtual private network connection that uses the Internet. Demand-dial connections are either on-demand (the point-to-point connection is only established when needed) or persistent (the point-to-point connection is established and then remains in a connected state). Demand-dial interfaces typically require an authentication process to become connected. The equipment needed by a demand-dial interface is a port on a device.

You can view the installed and configured routing interfaces by clicking **Network Interfaces** in Routing and Remote Access.

**Device**
A device is the hardware or software that provide ports that demand-dial and remote access connections use to establish point-to-point connections. Devices can be physical, such as a modem, or virtual, such as virtual private network (VPN) protocols. Devices can support a single port, such as a modem, or multiple ports, such as modem bank hardware that can terminate 64 different incoming analog phone calls. An example of a virtual multiport device is the Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP). Each of these tunneling protocols supports multiple VPN connections.

You can view the installed devices by viewing the properties of **Ports** in Routing and Remote Access.

**Port**
A port is a channel of a device that supports a single point-to-point connection. For single-port devices such as modems, the device and the port are indistinguishable. For multiport devices, the port is the subdivision of the device over which a separate point-to-point communication is possible. For example, Primary Rate Interface (PRI) ISDN adapters support two separate channels called B channels. The ISDN adapter is a device. Each B channel is a port because a separate point-to-point connection occurs over each B channel.

You can view the installed ports by clicking **Ports** in Routing and Remote Access.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider

using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

# Understanding Multicasting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding multicasting**
Multicasting is useful for the point-to-multipoint delivery of information on an internetwork. You can use three mechanisms for point-to-multipoint delivery:

1. Send the information to each endpoint individually by using unicast addresses. The disadvantages of this method are the duplication of network traffic and the overhead of maintaining a list of unicast endpoints.

2. Send the information in a single packet by using a broadcast address. The advantages of this method are the use of a single packet and no overhead for keeping lists of recipients. The disadvantages are the use of broadcast packets (which disturbs all nodes on the network) and the fact that broadcasts are not forwarded by routers. A broadcast packet reaches everyone on a network, but not everyone on the internetwork.

3. Send the information in a single packet by using a multicast address. The advantages of this method are the use of a single packet and no overhead for keeping lists of recipients. Unlike broadcast packets, multicast traffic does not disturb those nodes who are not listening for it. Routers can be multicast-capable and forward the multicast packet to all networks where there is at least one node listening.

Multicasting is the best choice for point-to-multipoint delivery.

This section covers:

- Multicast forwarding
- Multicast routing
- Multicast boundaries and heartbeat

# Multicast forwarding

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multicast forwarding**
Multicast forwarding, the intelligent forwarding of multicast traffic, is provided by the TCP/IP protocol and the IGMP routing protocol component for interfaces running in IGMP router mode, on a server running Routing and Remote Access.

The TCP/IP protocol performs the following multicast forwarding functions:

- Receives and processes all multicast traffic

  The Windows Server 2003 family implementation of the TCP/IP protocol receives and processes all multicast traffic on interfaces that are configured for IGMP router mode. All IP multicast traffic is either forwarded to a process running on the router or intelligently forwarded to network segments that contain either multicast group members or routers connected to downstream network segments that contain multicast group members.

- Forwards multicast packets to appropriate interfaces

  Upon receipt of a multicast packet, Routing and Remote Access consults a multicast forwarding table to decide whether to forward the packet to any of its other attached interfaces.

Interfaces running in IGMP router mode perform the following multicast forwarding functions:

- Places the network adapter in multicast-promiscuous mode

  IGMP router mode sets the network adapter to multicast-promiscuous mode. In multicast-promiscuous mode, all multicast packets received by the network adapter are passed to the networking layers for processing. Not all network adapters are capable of multicast-promiscuous mode.

- Tracks multicast group membership

  The IGMP router mode interface listens for IGMP Membership Report messages on locally attached networks and compiles a list of multicast accounting information as a series of {receiver network, multicast group} pairs. The receiver network is the IP network ID of a network segment containing a listening node. The multicast group is the specific multicast address that is registered by a listening node. The multicast-capable router does not track the IP addresses of the hosts that are listening, only the IP network ID where at least one host is listening.

  To ensure that hosts are still listening on their registered multicast addresses, the IGMP router periodically queries each network. The response to the query is an IGMP Membership Report message. If there are multiple IGMP routers on a single network, one IGMP router is elected the querier and performs all periodic queries.

- Updates to the TCP/IP multicast forwarding table

  Based on the current state of listening hosts on IGMP router mode interfaces, the TCP/IP multicast forwarding table is updated with the appropriate entries.

# Multicast routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multicast routing**
Multicast routing, the propagation of multicast listening information, is provided by multicast routing protocols such as Distance Vector Multicast Routing Protocol (DVMRP). The Windows Server 2003 family does not provide any multicast routing protocols. However, you can use the IGMP routing protocol and IGMP router mode and IGMP proxy mode to provide multicast forwarding in a single router intranet or when connecting a single router intranet to the Internet.

**Single-router intranets**
For an intranet that connects multiple networks by using a single router, you can enable IGMP router mode on all router interfaces to provide multicast forwarding support between multicast sources and multicast listening hosts on any network.

For more information, see Add the IGMP routing protocol and Enable IGMP router and IGMP proxy mode.

**Single-router intranets and the Internet**
If the server running Routing and Remote Access is attached to the MBone (the multicast-capable portion of the Internet) through an Internet service provider (ISP), you can use IGMP proxy mode to send and receive multicast traffic to and from the Internet.

In the following illustration, IGMP proxy mode is enabled on the Internet interface and IGMP router mode is enabled on the intranet interface. Multicast hosts register themselves locally and the IGMP proxy mode interface registers their memberships to the multicast-capable router at the ISP. Multicast traffic from the Internet is forwarded to the ISP router. The ISP router forwards the multicast traffic to the server running Routing and Remote Access, which then forwards the traffic to listening hosts on your intranet.




When multicast traffic is sent by an intranet host, it is forwarded over the IGMP proxy mode interface to the ISP router. The ISP router then forwards it to the appropriate downstream router. In this way, Internet hosts can receive multicast traffic sent by intranet hosts.

For more information, see Add the IGMP routing protocol and Enable IGMP router and IGMP proxy mode.

### Note

- The IGMP routing protocol that uses IGMP router mode and IGMP proxy mode is not the equivalent of a multicast routing protocol. A multicast routing protocol is required for multicast forwarding and routing support in a multiple-router intranet. You may configure the IGMP router mode and IGMP proxy mode interfaces to provide multicast forwarding support in multiple-router intranets. However, it is not recommended or supported.

# Multicast boundaries and heartbeat

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multicast boundaries and heartbeat**
The Windows Server 2003 family also provides multicast boundary and multicast heartbeat support.

**Multicast boundaries**
Multicast boundaries are administrative barriers to the forwarding of IP multicast traffic. Without boundaries, an IP multicast router forwards all appropriate IP multicast traffic. You can create multicast boundaries by specifying a range of IP addresses, known as a multicast scope, or by the value of the Time to Live (TTL) field in the IP header.

**Scope-based boundaries**
Scope-based boundaries prevent the forwarding of IP multicast traffic with a specified group IP address or range of IP addresses.

**TTL-based boundaries**
TTL-based boundaries prevent the forwarding of IP multicast traffic with a TTL less than a specified value. TTL-based boundaries apply to all multicast packets regardless of the multicast group. TTL-based boundaries are less effective than scope-based boundaries.

For more information, see Configure multicast boundaries.

**Multicast heartbeat**
Multicast heartbeat is the ability of the server running Routing and Remote Access to listen for a regular multicast notification to a specified group address to verify that IP multicast connectivity is available on the network. If the heartbeat is not received within a configured amount of time, the server running Routing and Remote Access sets a flag on the configured interface. The status of this flag can be polled by a program to provide notification that the multicast heartbeat is no longer present. For more information, see "IP Multicast Support" at the Microsoft Windows Resource Kits Web site.

For more information, see Configure multicast heartbeat.

# Understanding Network Address Translation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding network address translation**
With network address translation (NAT), you can configure your home network or small office network to share a single connection to the Internet. NAT consists of the following components:

- **Translation component**

  The server running Routing and Remote Access on which NAT is enabled translates the IP addresses and TCP/UDP port numbers of packets that are forwarded between the private network and the Internet.

- **Addressing component**

  The network address translation computer provides IP address configuration information to the other computers on the home network. The addressing component is a simplified DHCP server that allocates an IP address, a subnet mask, a default gateway, and the IP address of a DNS server. You must configure computers on the home network as DHCP clients in order to receive the IP configuration automatically. The default TCP/IP configuration for computers running any member of the Windows Server 2003 family, Windows XP, Windows 2000, Windows NT, Windows Millennium Edition, Windows 98, or Windows 95 is as a DHCP client.

- **Name-resolution component**

  The network address translation computer becomes the DNS server for the other computers on the home network. When name resolution requests are received by the network address translation computer, it forwards the name-resolution requests to the Internet-based DNS server for which it is configured and returns the responses to the home network computer.

For more information about network address translation, see:

- Internet private addresses

- A NAT example

For more information about configuring NAT, see Setting Up Network Address Translation.

For an example of using NAT, see SOHO Network to the Internet.

**Note**

- Because NAT includes addressing and name resolution components that provide DHCP and DNS services for hosts on the private network, you cannot run:

  - The DHCP service or the DHCP Relay Agent if NAT addressing is enabled.

  - The DNS service if NAT TCP/IP networking name resolution is enabled.

# Internet private addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet private addresses**

To communicate on the Internet, you must use addresses that have been allocated by the Internet Assigned Numbers Authority (IANA). Addresses allocated by IANA can receive traffic from Internet locations and are known as public addresses. A typical small business or home office is allocated a public address (or addresses) from its Internet service provider (ISP), who has received a range of public addresses.

To allow multiple computers in the small office or home office to communicate on the Internet, each computer must have its own public address. This requirement places great stress on the available pool of public addresses.

To relieve this stress, IANA has provided for an address reuse scheme by reserving network IDs for private internetworks. The private network IDs include:

- 10.0.0.0 with the subnet mask 255.0.0.0
- 172.16.0.0 with the subnet mask 255.240.0.0
- 192.168.0.0 with the subnet mask 255.255.0.0

For more information about portions of the IP address space that are reserved for private intranets, see RFC 1918, "Address Allocation for Private Internets." All addresses in these ranges are known as private addresses.

Private addresses cannot receive traffic from Internet locations. Therefore, if an intranet is using private addresses and communicating with Internet locations, the private address must be translated to a public address. A network address translator is placed between an intranet that uses private addresses and the Internet, which uses public addresses. Outgoing packets from the intranet have their private addresses translated by network address translation (NAT) into public addresses. Incoming packets from the Internet have their public addresses translated by NAT into private addresses. For more information, see A NAT example.

For more information about configuring NAT for the home office and small business, see SOHO Network to the Internet.

# A NAT example

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**A NAT example**

If a small business is using the 192.168.0.0 network ID for its intranet and has been granted the public address of *w1.x1.y1.z1* by its Internet service provider (ISP), then network address translation (NAT) maps all private addresses on 192.168.0.0 to the IP address of *w1.x1.y1.z1*. If multiple private addresses are mapped to a single public address, NAT uses dynamically chosen TCP and UDP ports to distinguish one intranet location from another.

**Note**

- The use of *w1.x1.y1.z1* and *w2.x2.y2.z2* is intended to represent valid public IP addresses as allocated by the Internet Assigned Numbers Authority (IANA) or an ISP.

The following illustration shows an example of using NAT to transparently connect an intranet to the Internet.




If a private user at 192.168.0.10 uses a Web browser to connect to the Web server at *w2.x2.y2.z2*, the user's computer creates an IP packet with the following information:

- Destination IP address: *w2.x2.y2.z2*
- Source IP address: 192.168.0.10
- Destination port: TCP port 80
- Source port: TCP port 5000

This IP packet is then forwarded to the NAT protocol, which translates the addresses of the outgoing packet to the following:

- Destination IP address: *w2.x2.y2.z2*
- Source IP address: *w1.x1.y1.z1*
- Destination port: TCP port 80
- Source port: TCP port 1025

The NAT protocol keeps the mapping of {192.168.0.10, TCP 1025} to {*w1.x1.y1.z1*, TCP 5000} in a table.

The translated IP packet is sent over the Internet. The response is sent back and received by the NAT protocol. When received, the packet contains the following public address information:

- Destination IP address: *w1.x1.y1.z1*

- Source IP address: *w2.x2.y2.z2*
- Destination port: TCP port 1025
- Source port: TCP port 80

The NAT protocol checks its translation table and maps the public addresses to private addresses and forwards the packet to the computer at 192.168.0.10. The forwarded packet contains the following address information:

- Destination IP address: 192.168.0.10
- Source IP address: *w2.x2.y2.z2*
- Destination port: TCP port 5000
- Source port: TCP port 80

For outgoing packets from the NAT protocol, the source IP address (a private address) is mapped to the ISP allocated address (a public address), and the TCP/UDP port numbers are mapped to a different TCP/UDP port number.

For incoming packets to the NAT protocol, the destination IP address (a public address) is mapped to the original intranet address (a private address), and the TCP/UDP port numbers are mapped back to their original TCP/UDP port numbers.

**Note**

- Packets that contain the IP address only in the IP header are properly translated by NAT. Packets that contain the IP address within the IP payload may not be properly translated by NAT.

# Understanding Demand-Dial Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding demand-dial routing**

While the concept of demand-dial routing is fairly simple, configuration of demand-dial routing is relatively complex. This complexity is due to the following factors:

- **Connection endpoint addressing**

  The connection must be made over public data networks, such as the analog phone system. The endpoint of the connection must be identified by a phone number or other endpoint identifier.

- **Authentication and authorization of the caller**

  Anyone calling the server running Routing and Remote Access must be authenticated and authorized. Authentication is based on the caller's set of credentials that are passed during the connection establishment process. The credentials that are passed must correspond to a user account. Authorization is granted based on the dial-in permission of the user account and remote access policies. For more information, see Introduction to remote access policies.

- **Differentiation between remote access clients and routers**

  Both routing and remote access services coexist on the same server running Routing and Remote Access. Both remote access clients and routers can call the same phone number. The server running Routing and Remote Access that answers the call must be able to distinguish a remote access client from a router that is calling to create a demand-dial connection. To differentiate a remote access client from a demand-dial router, the user name in the authentication credentials sent by the calling router must exactly match the name of a demand-dial interface on the answering router. Otherwise, the incoming connection is assumed to be a remote access connection.

- **Configuration of both ends of the connection**

  Both ends of the connection must be configured even if only one end of the connection is initiating a demand-dial connection. Configuring only one side of the connection means that packets are successfully routed only in one direction. Normal communication requires that information travel in both directions.

- **Configuration of static routes**

  You should not use dynamic routing protocols over temporary dial-up demand-dial connections. Therefore, routes for network IDs that are available across the demand-dial interface must be added to the routing table as static routes. You can accomplish this manually or by using auto-static updates.

For more information, see:

- Demand-dial routing example
- The demand-dial connection process
- Demand-dial routing updates
- Unnumbered connections
- One-way initiated demand-dial connections

For information about designing and deploying demand-dial routing, see Setting Up Demand-Dial Routing.

# Demand-dial routing example

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Demand-dial routing example**

This section gives an example of the complexity (and elegance) of demand-dial routing. The following illustration shows the configuration between two offices that want to use demand-dial IP routing.




The Seattle office has a server running Routing and Remote Access that acts as a remote access server and demand-dial router. All computers in the Seattle office are connected to the 172.16.1.0 network (subnet mask 255.255.255.0). The Seattle router (hereafter referred to as Router 1) has a modem connected to its COM1 port and the phone number of the modem is 555-0111.

The New York office has a server running Routing and Remote Access that acts as a remote access server and demand-dial router. All computers in the New York office are connected to the 172.16.2.0 network (subnet mask 255.255.255.0). The New York router (hereafter referred to as Router 2) has a modem connected to its COM2 port and the phone number of the modem is 555-0122.

The user on the computer with the IP address of 172.16.1.10 needs to be able to connect to the user on the computer with the IP address of 172.16.2.20 and vice versa.

### Note

- To configure demand-dial routing, use the Demand-dial Interface wizard, and do not skip the **Add a Static Route** dialog box. For more information about adding a demand-dial interface, see Add a demand-dial interface.

**Configuring Router 1**

The configuration for demand-dial routing on Router 1 consists of the following three steps:

1. Create a demand-dial interface.
2. Create a static route.
3. Create an account that Router 2 uses when calling Router 1.

**Creating a demand-dial interface**

By using Routing and Remote Access with the Demand-dial Interface wizard, the administrator at Router 1 creates a demand-dial interface called DD_NewYork with the following configuration:

- Equipment: Modem on COM1
- Phone number: 555-0122
- Protocols: TCP/IP
- Authentication credentials: DD_Seattle with a password

**Creating a static route**

By using Routing and Remote Access, the Demand-dial Interface wizard will already have the name and destination of the static route with the following configuration:

- Interface: DD_NewYork
- Destination: 172.16.2.0
- Network mask: 255.255.255.0
- Metric: 1

**Note**

- Because the demand-dial connection is a point-to-point connection, the gateway IP address is not configurable.

**Creating an account with dial-in permissions**

By using Active Directory Users and Computers or Local Users and Groups, the administrator at Router 1 creates a user account with the following settings:

- Account name: DD_NewYork with a password.
- Account settings: Clear the **User must change password at next logon** check box and select the **Password never expires** check box.

The DD_NewYork account is granted dial-in permissions either through the dial-in properties of the user account or through remote access policies. For more information, see Introduction to remote access policies.

**Configuring Router 2**

The configuration for demand-dial routing on Router 2 consists of the following three steps:

1. Create a demand-dial interface.
2. Create a static route.
3. Create an account that Router 1 uses when calling Router 2.

**Creating a demand-dial interface**

By using Routing and Remote Access, the administrator at Router 2 creates a demand-dial interface called DD_Seattle with the following configuration:

- Equipment: Modem on COM2
- Phone number: 555-0111
- Protocols: TCP/IP
- Authentication credentials: DD_NewYork with a password

**Creating a static route**

By using Routing and Remote Access, the Demand-dial Interface wizard will already have the name and destination of the static route with the following configuration:

- Interface: DD_Seattle
- Destination: 172.16.1.0
- Network mask: 255.255.255.0
- Metric: 1

### Note

- Because the demand-dial connection is a point-to-point connection, the gateway IP address is not configurable.

**Creating an account with dial-in permissions**

By using Active Directory Users and Computers or Local Users and Groups, the administrator at Router 2 creates a user account with the following settings:

- Account name: DD_Seattle with a password.
- Account settings: Clear the **User must change password at next logon** check box and select the **Password never expires** check box.

The DD_Seattle account is granted dial-in permissions either through the dial-in properties of the user account or through remote access policies. For more information, see Introduction to remote access policies.

**The resulting configuration**

The following illustration shows the demand-dial routing configuration in terms of the demand-dial interfaces, static routes, and user accounts for the routers for the Seattle and New York offices.




### Note

- In order for two-way initiated demand-dial routing to work properly, authentication credentials (such as user account name) for the demand-dial interface must be created on both servers running Routing and Remote Access. This example shows a proper configuration and is summarized in the following table.

| Router | Demand-dial interface name | User account name |
|---|---|---|
| Router 1 | DD_NewYork | DD_Seattle |
| Router 2 | DD_Seattle | DD_NewYork |

For more information about how the demand-dial connection process works for this example configuration, see The demand-dial connection process.

# The demand-dial connection process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The demand-dial connection process**

When the user at 172.16.1.10 tries to connect to a resource at 172.16.2.20, the following events occur:

1. Packets from 172.16.1.10 destined for 172.16.2.20 are forwarded to Router 1.
2. Router 1 receives the packet from 172.16.1.10 and checks its routing table. A route to 172.16.2.20 is found that uses the DD_NewYork interface.
3. Router 1 checks the state of the DD_NewYork interface and finds it is in a disconnected state.
4. Router 1 retrieves the configuration of the DD_NewYork demand-dial interface.
5. Based on the DD_NewYork configuration, Router 1 uses the modem on COM1 to dial the number 555-0122.
6. Router 2 answers the incoming call.
7. Router 2 requests authentication credentials from the incoming caller.
8. Router 1 sends the user name DD_Seattle with its associated password.
9. Upon receipt of the authentication credentials, Router 2 checks the user name and password against the security features of the Windows Server 2003 family and verifies that Router 1 has dial-in permission through the dial-in properties of the DD_Seattle user account and the configured remote access policies.
10. Router 2 must now determine whether the incoming caller is a dial-up networking client or a router that is creating a demand-dial connection. Router 2 looks in its list of demand-dial interfaces to find one that matches the user name sent by Router 1 as part of the authentication credentials. Router 2 finds a demand-dial interface DD_Seattle that matches the user name credential.
11. Router 2 changes the state of the DD_Seattle demand-dial interface to a connected state.
12. Router 1 forwards the packet from the computer at 172.16.1.10 across the demand-dial connection to Router 2.
13. Router 2 receives the packet and forwards it to the computer at 172.16.2.20.
14. The response to the connection request by the computer at 172.16.1.10 is forwarded to Router 2 by the computer at 172.16.2.20.
15. Router 2 receives the packet destined for 172.16.1.10 and checks its routing table. A route to 172.16.1.10 is found by using the DD_Seattle interface.
16. Router 2 checks the state of the DD_Seattle interface and finds it is in a connected state.
17. Router 2 forwards the packet to Router 1.
18. Router 1 forwards the packet to the computer at 172.16.1.10.

If the user name credential does not match the name of an appropriate demand-dial interface, the calling entity is identified as a remote access client, which can result in routing problems.

For example, if Router 1 uses DialUpRouter1 as its user name credential, then Router 1 is identified as a remote access client rather than a router (assuming DialUpRouter1 is a valid account with dial-in permissions). Packets are routed from the user at 172.16.1.10 to the user at 172.16.2.20 as described in the preceding process.

However, response packets from 172.16.2.20 to 172.16.1.10 are forwarded to Router 2 which, upon inspecting its routing table, determines that the interface to use is DD_Seattle. DD_Seattle is in a disconnected state. Based on the configuration for DD_Seattle, COM2 is to be used. However, COM2 is currently being used for a remote access client (Router 2). Router 2 then tries to locate another modem that is not being used. If found, Router 2 dials Router 1 and forwards the packet after the connection has been established. If another modem cannot be found, the response packets from 172.16.2.20 to 172.16.1.10 are dropped.

# Demand-dial routing updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Demand-dial routing updates**

While demand-dial routing can save connection costs, typical routing protocols rely on a periodic advertising process to communicate routing information. For example, RIP for IP advertises the contents of its routing table every 30 seconds on all interfaces. This behavior is not a problem for permanently connected LAN or WAN lines. For usage-sensitive dial-up WAN lines, this type of periodic behavior could cause the router to call another router every 30 seconds, which may result in an undesirable phone bill. Therefore, you should not run routing protocols across temporary dial-up WAN lines.

If you do not use routing protocols to update the routing tables, then you must enter the routes as static routes. The static routes that correspond to the network IDs available across the interface are entered manually or automatically. The automatic entering of static routes for demand-dial interfaces is known as auto-static updates and is supported by the server running Routing and Remote Access. Auto-static updates are supported when you use RIP for IP, but not OSPF.

When instructed, a demand-dial interface that is configured for auto-static updates sends a request across an active connection to request all of the routes of the router on the other side of the connection. In response to the request, all of the routes of the requested router are automatically entered as static routes in the routing table of the requesting router. The static routes are persistent; they are kept in the routing table even if the interface becomes disconnected or the router is restarted. An auto-static update is a one-time, one-way exchange of routing information.

You can automate and schedule auto-static updates by executing the update as a scheduled task. For more information, see Scheduling auto-static updates.

**Notes**

- The "auto" in auto-static refers to the automatic adding of the requested routes as static routes in the routing table. The sending of the request for routes is performed through an explicit action: either through Routing and Remote Access or the Netsh utility while the demand-dial interface is in a connected state. Auto-static updates are not automatically performed every time a demand-dial connection is made.

- When an auto-static update is requested, the existing auto-static routes are deleted before the update is requested from other routers. If there is no response to the request, then the router cannot replace the routes it has deleted. This may lead to a loss of connectivity to remote networks.

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Unnumbered connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Unnumbered connections**

Demand-dial connections use the Point-to-Point Protocol (PPP) during the connection establishment process. For a TCP/IP-based connection, TCP/IP connection settings are negotiated. An important element of a TCP/IP-based PPP connection is the allocation of an IP address.

When a call is initiated between demand-dial routers (either servers running Routing and Remote Access or servers running Windows NT 4.0 Routing and Remote Access Service (RRAS)), the calling router requests an IP address from the answering router. Additionally, the answering router requests an IP address from the calling router. If both routers have an IP address to give each other, the logical interface on the PPP connection for each router is assigned an IP address from the other router. This is known as a numbered connection.

Windows NT 4.0 RRAS demand-dial connections require a numbered connection. If either the calling or the answering router does not have an IP address to assign to the other router, the connection establishment process fails. In order for both routers to assign IP addresses to each other, you must configure both routers for the proper IP address assignment behavior through the properties of RRAS. A common configuration problem with RRAS is that the default configuration for IP address assignment is to obtain an IP address through DHCP, and no DHCP server is available on the network. Because no IP address is allocated to the other router, the connection establishment process fails.

This configuration issue is no longer a problem with the Routing and Remote Access service for two reasons:

1. Use of Automatic Private IP Addressing (APIPA) addresses

   In Routing and Remote Access, the default IP address assignment configuration is still to obtain an IP address from a DHCP server. However, if a DHCP server is not available, APIPA addresses in the range from 169.254.0.1 through 169.254.255.254 are used. Therefore, the default configuration does not prevent two routers from establishing a numbered connection. For more information about APIPA, see New features for TCP/IP.

2. Support for unnumbered connections

   By default, demand-dial connections do not require a numbered connection. The calling and answering routers still request an IP address from each other during the connection establishment process. But if one of the routers rejects the request, both routers continue with the connection establishment process. The logical interface on the point-to-point connection does not have an assigned IP address. This is known as an unnumbered connection.

When you use an unnumbered connection, on the router that requests an IP address but is rejected, a warning is logged in the system event log with the following characteristics:

- Source: RemoteAccess
- Event ID: 20165
- Message text: A connection has been established on port [number] using interface [name of demand-dial interface], but no IP address was obtained.

On the router from which an IP address is requested but is not given, a warning is logged in the system event log with the following characteristics:

- Source: RemoteAccess
- Event ID: 20166
- Message text: A connection has been established on port [number] using interface [name of demand-dial interface], but the remote side got no IP address.

With unnumbered connections, servers running Routing and Remote Access can connect to other routers that do not assign an IP address during the connection establishment process. For example, some Internet service providers (ISPs) prefer to use an unnumbered connection to conserve IP addresses.

**Important**

- The routing protocols provided with Routing and Remote Access cannot operate over unnumbered connections. If you use unnumbered connections, you must employ static routing. This limitation is not an issue when you are using a server running Routing and Remote Access to connect to the Internet and you configure a default static IP route, rather than run a routing protocol.

# One-way initiated demand-dial connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**One-way initiated demand-dial connections**
If a demand-dial connection is always initiated by one router, a one-way initiated demand-dial connection, then you can simplify the configuration of the answering router by configuring static routes on the user account of the calling router.

Using the example scenario described in Demand-dial routing example, the Seattle router (Router 1) always calls the New York router (Router 2).

**Configuring Router 1**
The configuration of Router 1 is the same as described in Demand-dial routing example.

**Configuring Router 2**
The configuration for demand-dial routing on Router 2 consists of just two steps:

1. Create a user account that Router 1 uses when calling Router 2.

2. Configure static routes on the user account used by Router 1 for the network IDs of the Seattle office.

**Creating a user account with dial-in permissions**
By using Active Directory Users and Computers or Local Users and Groups, the administrator at Router 2 creates a user account with the following settings:

- Account name: DD_Seattle with a password.

- Account settings: Clear the **User must change password at next logon** check box and select the **Password never expires** check box.

The DD_Seattle account is granted dial-in permissions either through the dial-in properties of the user account or through remote access policies. For more information, see Introduction to remote access policies.

**Configuring the user account with static routes**
By using Active Directory Users and Computers or Local Users and Groups, the administrator at Router 2 obtains dial-in properties on the DD_Seattle user account and then adds the static routes of the Seattle office.

For more information, see Configure static routes for a dial-in user.

For this example, the network administrator adds the following route:

- Destination: 172.16.1.0

- Network mask: 255.255.255.0

- Hops: 1

**Note**

- You can only add static routes to the dial-in properties of a user account on a server running Routing and Remote Access acting as stand-alone router (not a member of a domain) or a server running Routing and Remote Access that is a member of a Windows 2000 native or a Windows Server 2003 domain.

**Connection process for a one-way initiated demand-dial connection**
When the user at 172.16.1.10 tries to connect to a resource at 172.16.2.20, the following events occur:

1. Packets from 172.16.1.10 destined for 172.16.2.20 are forwarded to Router 1.

2. Router 1 receives the packet from 172.16.1.10 and checks its routing table. A route to 172.16.2.20 is found by using the DD_NewYork interface.

3. Router 1 checks the state of the DD_NewYork interface and finds it is in a disconnected state.

4. Router 1 retrieves the configuration of the DD_NewYork demand-dial interface.

5. Based on the DD_NewYork configuration, Router 1 uses the modem on COM1 to dial the number 555-0122.

6. Router 2 answers the incoming call.

7. Router 2 requests authentication credentials from the incoming caller.

8. Router 1 sends the user name DD_Seattle with its associated password.

9. Upon receipt of the authentication credentials, Router 2 checks the user name and password against the security features of the Windows Server 2003 family and verifies that Router 1 has dial-in permission through the dial-in properties of the DD_Seattle user account and the configured remote access policies.

10. Router 2 retrieves the static route (172.16.1.0 with the subnet mask of 255.255.255.0) that is configured on the DD_Seattle user account and creates a corresponding static route in its routing table. If Router 2 is configured with routing protocols, Router 2 uses routing protocols to communicate with neighboring routers so that the route to the Seattle network is propagated to all of the routers in the New York office.

11. Router 2 must now determine whether the incoming caller is a dial-up networking client or a router creating a demand-dial connection. Router 2 looks in its list of demand-dial interfaces and does not find one called DD_Seattle. Therefore, Router 2 considers the connection to the Seattle office to be a remote access connection.

12. Router 1 forwards the packet from the computer at 172.16.1.10 across the demand-dial connection to Router 2.

13. Router 2 receives the packet and forwards it to the computer at 172.16.2.20.

14. The response to the connection request by the computer at 172.16.1.10 is forwarded to Router 2 by the computer at 172.16.2.20.

15. Router 2 receives the packet destined for 172.16.1.10 and checks its routing table. A route to 172.16.1.10 is found by using the connection to Router 1.

16. Router 2 forwards the packet to Router 1.

17. Router 1 forwards the packet to the computer at 172.16.1.10.

**Note**

- When the connection is made, the static routes on the user account of the calling router are added to the routing table of the answering router. If routing protocols are used to propagate the new static route, then there is a delay between the time the connection is made and the time when all of the routers on the intranet of the answering router are aware of the new route. Therefore, hosts on the intranet of the calling router may experience a delay between the time that the connection is made and the time when they begin to receive traffic back from hosts on the intranet of the answering router.

# Understanding Router-to-Router VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding router-to-router VPNs**

This section covers:

- Components of a router-to-router VPN connection
- VPN Router-to-Router tunneling protocols
- A dial-up router-to-router VPN connection

# Components of a router-to-router VPN connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Components of a router-to-router VPN connection**

A router-to-router VPN connection includes the following components:

- **Virtual private network (VPN) clients**

  The VPN client is the router that initiates the VPN connection, also known as the calling router. For router-to-router VPN connections, you can configure servers running Routing and Remote Access (available in both Windows 2000 and the Windows Server 2003 family) and computers running Windows NT Server 4.0 and the Routing and Remote Access Service (RRAS) as VPN clients.

- **VPN servers**

  The VPN server is the router that accepts the connection from the calling router, also known as the answering router. For router-to-router VPN connections, you can configure both servers running Routing and Remote Access and computers running Windows NT Server 4.0 and RRAS as VPN servers. On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

- **LAN and remote access protocols**

  LAN protocols are used to transport information. Remote access protocols are used to negotiate connections and provide framing for LAN protocol data. Routing and Remote Access supports the routing of TCP/IP LAN protocol packets by using the PPP remote access protocol across a router-to router VPN connection.

- **Tunneling protocols**

  Tunneling protocols are used by VPN clients and VPN servers to manage tunnels and send tunneled data. Routing and Remote Access includes support for the Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) tunneling protocols. Windows NT Server 4.0 with RRAS only includes the PPTP tunneling protocol.

- **WAN options**

  VPN servers are typically connected to the Internet by using permanent WAN connections such as T1 or Frame Relay. VPN clients are connected to the Internet by using permanent WAN connections or by dialing in to a local Internet service provider (ISP) that uses standard telephone lines or ISDN. Once connected to the Internet, the VPN client can connect to the VPN server.

- **Demand-dial interfaces**

  The VPN client (the calling router) must have a demand-dial interface configured for:

  - The host name or IP address of the interface of the VPN server on the Internet.
  - A PPTP port (for a PPTP-based VPN connection) or an L2TP port (for an L2TP-based connection).
  - The user account credentials (user name, domain, password) for a user account that can be validated by the VPN server.
  - For an L2TP/IPSec-based connection, a valid certificate that can be validated by the VPN server.

The VPN server (the answering router) must have a demand-dial interface with the same name as the user account that is used by the VPN client (the calling router) in order for the incoming connection to be recognized and accepted as a demand-dial connection.

- **User accounts**

  A user account must be created for the calling router. This account can be created automatically when you run the Routing and Remote Access with the Routing and Remote Access Server Setup Wizard. This account must have dial-in permissions either through the dial-in properties of the user account or through remote access policies. For more information, see Introduction to remote access policies.

- **Static routes or routing protocols**

  In order for each router to forward packets across the router-to-router VPN connection, each router must contain the appropriate routes in the routing tables. Routes are added to the routing tables of both routers either as static routes or by enabling a routing protocol to operate across a persistent router-to-router VPN connection.

- **Security options**

  Because the router-to-router VPN connection is validated by the server running Routing and Remote Access, you can utilize all of the security features of the Windows Server 2003 family, including domain security, data encryption, Remote Authentication Dial-In User Service (RADIUS), smart cards, and callback. For more information, see Remote Access Security.

The following illustration shows the components of a router-to-router VPN connection.



# VPN Router-to-Router tunneling protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tunneling protocols**

Routing and Remote Access provides two tunneling protocols for creating router-to-router VPN connections:

- Point-to-Point Tunneling Protocol (PPTP)
- Layer Two Tunneling Protocol (L2TP)

### Note

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**PPTP**

PPTP is a VPN tunneling protocol. PPTP is an extension of the Point-to-Point Protocol (PPP) and leverages the authentication, compression, and encryption mechanisms of PPP.

PPTP is installed with TCP/IP. By default, Routing and Remote Access is configured for five PPTP ports. To increase the number of PPTP ports, see Add PPTP or L2TP ports. You can enable PPTP ports for inbound remote access and demand-dial routing connections by using the Routing and Remote Access Server Setup Wizard. To enable PPTP ports for routing after the wizard is run, see Enable routing on ports.

Two primary services of virtual private networking are encapsulation and encryption.

**Encapsulation**

A PPP frame containing an IP datagram is wrapped with a Generic Routing Encapsulation (GRE) header and an IP header. In the IP header is the source and destination IP address that correspond to the VPN client and VPN server.

The following illustration shows PPTP encapsulation for a PPP frame.




**Encryption**

The PPP frame is encrypted with Microsoft Point-to-Point Encryption (MPPE) by using encryption keys generated from the MS-CHAP or EAP-TLS authentication process. Virtual private networking clients must use either the MS-CHAP or EAP-TLS authentication protocol in order to encrypt PPP payloads. PPTP does not provide encryption services. PPTP encapsulates a previously encrypted PPP frame.

### Note

- It is possible to have a nonencrypted PPTP connection where the PPP payload is sent in plaintext. However, a nonencrypted PPTP connection is not recommended for virtual private network connections over the Internet because communications of this type are not secure.

**L2TP**

L2TP is an industry-standard Internet tunnelling protocol that provides encapsulation for sending PPP frames across packet-oriented media. Like PPTP, L2TP leverages the authentication and compression mechanisms of PPP. Unlike PPTP, L2TP does not utilize Microsoft Point-to-Point Encryption (MPPE) to encrypt PPP frames. L2TP relies on Internet Protocol security (IPSec)

for encryption services.

The result is that L2TP-based virtual private networking connections are a combination of L2TP and IPSec. Both L2TP and IPSec must be supported by both routers. For more information about IPSec, see Internet Protocol Security Overview. L2TP is described in RFC 2661.

L2TP is installed with TCP/IP. By default, Routing and Remote Access is configured for five L2TP ports. To increase the number of L2TP ports, see Add PPTP or L2TP ports. You can enable L2TP ports for inbound remote access and demand-dial routing connections by using the Routing and Remote Access Server Setup Wizard. To enable L2TP ports for routing after the wizard is run, see Enable routing on ports.

Two primary services of virtual private networking are encapsulation and encryption.

**Encapsulation**
Encapsulation for L2TP over IPSec packets consists of two layers of encapsulation.

1. L2TP encapsulation

   A PPP frame containing an IP datagram is wrapped with an L2TP header and a UDP header.

2. IPSec encapsulation

   The resulting L2TP message is then wrapped with an IPSec Encapsulating Security Payload (ESP) header and trailer, an IPSec Authentication trailer that provides message integrity and authentication, and a final IP header. In the IP header is the source and destination IP address that correspond to the VPN client and VPN server.

The following illustration shows L2TP and IPSec encapsulation for a PPP frame.




**Encryption**
The L2TP message is encrypted with DES and 3DES encryption algorithms by using encryption keys generated from the IPSec authentication process.

### Note

- It is possible to have a non-IPSec-based (nonencrypted) L2TP connection where the PPP payload is sent in plaintext. However, a nonencrypted L2TP connection is not recommended for virtual private network connections over the Internet because communications of this type are not secure.

# A dial-up router-to-router VPN connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**A dial-up router-to-router VPN connection**

A router-to-router VPN connection is typically used to connect remote offices together when both routers are connected to the Internet through permanent WAN links, such as T1 or Frame Relay. In this configuration, the VPN connection is always available. However, when a permanent WAN link is not possible or practical, you can configure a dial-up router-to-router VPN connection.

The following illustration shows a dial-up router-to-router VPN connection.




A dial-up router-to-router VPN connection consists of two demand-dial interfaces and two static routes that are configured on the VPN client (the calling router):

1. A demand-dial interface to dial in to a local Internet service provider (ISP).
2. A demand-dial interface for the router-to-router VPN connection.
3. A static host route to dynamically connect to the Internet.
4. A static route to reach the locations of the corporate office.

A dial-up router-to-router VPN connection is automatically established when you route traffic to a specific location. For example, in a branch office configuration, when a packet is received to be routed to the corporate office, the branch office router uses a dial-up link to connect to a local Internet service provider (ISP) and then creates a router-to-router VPN connection with the corporate office router located on the Internet.

**Note**

- This discussion assumes that the corporate office router (the answering router) is connected to the Internet by using a permanent WAN link. It is possible to have both routers connected to the Internet by using a dial-up WAN link. However, this is only feasible if the ISP supports demand-dialing routing to customers; the ISP calls the customer router when an IP datagram is to be delivered to the customer. Demand-dial routing to customers is not widely supported by ISPs.

To configure a dial-up router-to-router VPN connection at the branch office router, do the following:

1. Create a demand-dial interface for the Internet connection that is configured for the appropriate equipment (a modem or ISDN device), the phone number of the local ISP, and the user name and password to gain Internet access.
2. Create a demand-dial interface for the VPN connection with the corporate office router that is configured for a VPN port (a PPTP or L2TP port), the IP address of the interface on the Internet for the corporate office router, and a user name and password that can be verified by the VPN server. The user name must match the name of a demand-dial interface on the corporate office router.
3. Create a static host route to the corporate office router that uses the ISP demand-dial interface.
4. Create a static route (or routes) for the IP network IDs of the corporate intranet that uses the VPN demand-dial interface.

To configure the corporate office router, do the following:

1. Create a demand-dial interface for the VPN connection with the branch office that is configured for a VPN device (a PPTP or L2TP port). The demand-dial interface must have the same name as the user name in the authentication credential that is used by the branch office router to create the VPN connection.

2. Create a static route (or routes) for the IP network IDs of the branch office that uses the VPN demand-dial interface.

The router-to-router VPN connection is automatically initiated by the branch office router through the following process:

1. Packets sent to a corporate hub network location from a computer in the branch office are forwarded to the branch office router.

2. The branch office router checks its routing table and finds a route to the corporate intranet network ID, which uses the VPN demand-dial interface.

3. The branch office router checks the state of the VPN demand-dial interface and finds it is in a disconnected state.

4. The branch office router retrieves the configuration of the VPN demand-dial interface.

5. Based on the VPN demand-dial interface configuration, the branch office router attempts to initialize a router-to-router VPN connection at the IP address of the corporate office router on the Internet.

6. To establish a VPN, either a TCP (by using PPTP) or UDP (by using L2TP/IPSec) packet must be sent to corporate office router that acts as the VPN server. The VPN establishment packet is created.

7. To forward the VPN establishment packet to the corporate office router, the branch office router checks its routing table and finds the host route that is using the ISP demand-dial interface.

8. The branch office router checks the state of the ISP demand-dial interface and finds it is in a disconnected state.

9. The branch office router retrieves the configuration of the ISP demand-dial interface.

10. Based on the ISP demand-dial interface configuration, the branch office router uses its modem to dial and establish a connection with its local ISP.

11. Once the ISP connection is made, the VPN establishment packet is sent to the corporate office router.

12. A router-to-router VPN connection is negotiated between the branch office router and the corporate office router. As part of the negotiation, the branch office router sends authentication credentials that are verified by the corporate office router.

13. The corporate office router checks its demand-dial interfaces and finds one that matches the user name sent during authentication and changes the interface to a connected state.

14. The branch office router forwards the routed packet across the VPN and the corporate office router forwards the packet to the appropriate intranet location.

15. When the intranet location responds to the packet sent to it by the user in the branch office, the packet is forwarded to the corporate office router.

16. The corporate office router checks its routing table and finds a route to the branch office network that uses the VPN demand-dial interface.

17. The corporate office router checks the state of the VPN demand-dial interface and finds it is in a connected state.

18. The response packet is forwarded across the Internet by using the VPN connection.

19. The response packet is received by the branch office router and is forwarded to the original user.

# Using Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using routing**

This section covers:

- Installing the router
- Managing routing from the command line
- Deploying Routing
- Routing Scenarios

# Installing the router

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing the router**
To install and configure the server running Routing and Remote Access as a router, you must be logged on as a member of the local Administrators group.

**Hardware requirements**
Before you configure Routing and Remote Access as a router, all hardware needs to be installed and working. Depending on your network and requirements, you might need the following hardware:

- A LAN or WAN adapter with a WHQL signed driver.
- One or more compatible modems and an available COM port.
- A multiport adapter for acceptable performance with multiple remote connections.
- An X.25 smart card (if you are using an X.25 network).
- An ISDN adapter (if you are using an ISDN line).

To verify the compatibility of all hardware in a computer running Routing and Remote Access, see Support resources.

**Installation**
When you install the Windows Server 2003 family, the routing component is automatically installed. However, the Routing and Remote Access service is installed in a disabled state. For more information, see Enable the Routing and Remote Access service.

# Managing routing from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing routing from the command line**

In addition to using Routing and Remote Access, you can use command-line utilities to manage routing for users, computers, or services. By combining command-line utilities within a batch file, you can simplify and automate complicated daily, weekly, or monthly tasks on one or more servers.

For example, you can use the Pathping command to diagnose network latency or loss problems. The **pathping** command sends Internet Control Message Protocol (ICMP) Echo Request messages to each router between a source and a destination over a period of time. Results based on the packets returned from each router are then computed and displayed. The Tracert command can be used to determine and display a path to a destination IP address that you want to reach. Used for diagnostic purposes only, the Ping command verifies IP-level connectivity to another TCP/IP computer. You can use Netsh overview to perform more advanced management tasks, such as scripting.

For more information about managing routing with command-line utilities, see Routing Tools and Utilities. For information about other command-line utilities, see Command-line reference A-Z.

# Deploying Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying routing**
This section covers:

- Setting Up a Static Routed IP Internetwork
- Setting Up a RIP-for-IP Routed Internetwork
- Setting Up an OSPF Routed Internetwork
- Setting Up Network Address Translation
- Setting Up Demand-Dial Routing
- Deploying Router-to-Router VPNs

**Note**

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Setting Up a Static Routed IP Internetwork

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up a static routed IP internetwork**
A static routed IP internetwork does not use routing protocols such as RIP for IP or OSPF to communicate routing information between routers. All of the routing information is stored in a static routing table on each router. You need to ensure that each router has the appropriate routes in its routing table so that traffic can be exchanged between any two endpoints on the IP internetwork.

This section covers:

- The static routed environment
- Static routing design considerations
- Static routing security
- Deploying static routing

**Note**

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# The static routed environment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The static routed environment**

A static routed IP environment is best suited to a small, single-path, static IP internetwork:

- A small internetwork is defined as 2 to 10 networks.
- Single-path means that there is only a single path for packets to travel between any two endpoints on the internetwork.
- Static means that the topology of the internetwork does not change over time.

Candidates for a static routed environment include:

- A small business.
- A home office IP internetwork.
- A branch office with a single network.

  Rather than running a routing protocol across a typically low-bandwidth WAN link, a single default route at the branch office router ensures that all traffic not destined for a computer on the branch office network is routed to the main office.

The disadvantages of static routing are:

- No fault tolerance

  If a router or link goes down, static routers do not sense the fault and inform other routers of the fault. While this is a concern on large, corporate internetworks, a small office (with two routers and three networks based on LAN links) does not go down often enough to justify deploying a multipath topology and a routing protocol.

- Administrative overhead

  If a new network is added or removed from the internetwork, routes to the new network must be manually added or removed. If a new router is added, it must be properly configured for the routes of the internetwork.

# Static routing design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Static routing design considerations**
To prevent problems, you should consider the following design issues before you implement static routing.

**Peripheral router configuration**
To simplify configuration, you can configure peripheral routers with a default route that points to the neighboring router. A peripheral router is a router attached to multiple networks, only one of which has a neighboring router.

**Default routes and routing loops**
It is recommended that you do not configure two neighboring routers with default routes that are pointing to each other. A default route passes all traffic that is not on a directly connected network to the configured router. Two routers that have default routes pointing to each other may produce routing loops for traffic with an unreachable destination.

**Demand-dial environments**
You can implement static routing across demand-dial links in one of two ways:

- Default route

  You can configure a default route on the branch office router that uses the demand-dial interface. The advantage of a default route is that a single route only needs to be added once. The disadvantage of a default route is that any traffic, including traffic for unreachable destinations, that is not on the branch office network causes the branch office router to call the main office.

- Auto-static routes

  Auto-static routes are static routes that are automatically added to the routing table for a router after routes are requested across a demand-dial connection by using the RIP-for-IP routing protocol. The advantage of auto-static routes is that unreachable destinations do not cause the router to call the main office. The disadvantage of auto-static routes is that they must be periodically updated to reflect the networks that are reachable at the main office. If a new network is added to the main office and the branch office has not performed an auto-static update, all destinations on the new main office network are unreachable from the branch office.

# Static routing security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Static routing security**

To prevent the intentional or unintentional modification of static routes on the routers, you should:

- Employ physical security so that the routers cannot be accessed by users.
- Assign administrator rights only to those users who will be running Routing and Remote Access.

# Deploying static routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying static routing**

If static routing is appropriate for your IP internetwork, you can perform the following steps to deploy static routing:

1. Draw a map of the topology of your IP internetwork that shows the separate networks and the placement of routers and hosts (nonrouter computers that run TCP/IP).

2. For each IP network (a cabling system bounded by one or more routers), assign a unique IP network ID (also known as an IP network address).

3. Assign IP addresses to each router interface. It is a common industry practice to assign the first IP addresses of a given IP network to router interfaces. For example, for an IP network ID of 192.168.100.0 with a subnet mask of 255.255.255.0, the router interface is assigned the IP address of 192.168.100.1.

4. For peripheral routers, configure a default route on the interface that has a neighboring router. The use of default routes on peripheral routers is optional.

5. For each nonperipheral router, compile a list of routes that need to be added as static routes to the routing table for that router. Each route consists of a destination network ID, a subnet mask, a gateway (or forwarding) IP address, a metric (number of router hops to reach the network), and the interface to be used to reach the network.

6. For nonperipheral routers, add the static routes compiled in step 5 to each router. You can add static routes by using Routing and Remote Access.

7. When your configuration is complete, use the **ping** and **tracert** commands to test connectivity between host computers so that all routing paths are checked. For more information about the **ping** and **tracert** commands, see Using the ping command and Using the tracert command.

For information about troubleshooting static routing, see Common routing problems and Understanding the IP routing table.

# Setting Up a RIP-for-IP Routed Internetwork

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up a RIP-for-IP routed internetwork**

A RIP-for-IP routed internetwork uses the RIP-for-IP routing protocol to dynamically communicate routing information between routers. Deployed properly, a RIP-for-IP environment automatically adds and removes routes as networks are added and removed from the internetwork. You need to ensure that each router is properly configured so that the RIP-based route announcements are sent and received by all RIP routers on the internetwork.

This section covers:

- The RIP-for-IP environment
- RIP-for-IP design considerations
- RIP-for-IP security
- Deploying RIP for IP

# The RIP-for-IP environment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The RIP-for-IP environment**

A RIP-for-IP routed environment is best suited to a small-to-medium-sized, multipath, dynamic IP internetwork:

- A small-to-medium-sized internetwork is defined as 10 to 50 networks.
- Multipath means that there are multiple paths for packets to travel between any two endpoints on the internetwork.
- Dynamic means that the topology of the internetwork changes over time because networks are added and removed and links occasionally go down and come back up.

Candidates for a RIP-for-IP routed environment include:

- A medium-sized business.
- A large branch office or satellite office with multiple networks.

# RIP-for-IP design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RIP-for-IP design considerations**
To prevent problems, you should consider the following design issues before you implement RIP for IP.

**Decreased diameter of 14 routers**
The maximum diameter of RIP internetworks is 15 routers. The diameter is a measure of the size of an internetwork in terms of hops or other metrics. However, the server running Routing and Remote Access considers all non-RIP learned routes to be at a fixed hop count of 2. Static routes, even static routes for directly connected networks, are considered non-RIP learned routes. When a server running Routing and Remote Access acting as a RIP router advertises its directly connected networks, it advertises them at a hop count of 2 even though there is only one physical router to cross. Therefore, a RIP-based internetwork that uses servers running Routing and Remote Access have a maximum physical diameter of 14 routers.

**RIP costs**
RIP uses the hop count as a metric to determine the best route. Using the number of routers to cross as the basis for selecting the best route may lead to undesired routing behavior. For example, if two sites were connected together by using a T1 link and a lower-speed satellite link as a backup, both links are considered the same metric. When the router is given a choice between two routes of the same lowest metric (hop count), the router is free to choose between them.

If the router chooses the satellite link, then the slower backup link is used rather than the higher bandwidth link. To prevent the satellite link from being chosen, you can assign a custom cost to the satellite interface. For example, if you assign a cost of 2 to the satellite interface (rather than the default of 1), then the best route is always the T1 link. If the T1 link goes down, the satellite link is chosen as the next best route.

If you are using custom costs to indicate link speed, delay, or reliability factors, ensure that the accumulated costs (hop counts) between any two endpoints on the internetwork do not exceed 15.

**Mixed RIP version 1 and RIP version 2 environments**
For maximum flexibility, you should use RIP version 2 in your RIP-for-IP internetwork. If there are routers in your internetwork that do not support RIP version 2, you can use a mixed environment of RIP v1 and RIP v2. However, RIP v1 does not support classless interdomain routing (CIDR) or variable-length subnet masks (VLSM) implementations. If you have support for CIDR and VLSM in one part of your internetwork but not another, you may experience routing problems.

If a network is using a mixture of RIP v1 and RIP v2 routers, then you must configure the interfaces of the server running Routing and Remote Access to advertise by using either RIP v1 broadcasts or RIP v2 broadcasts and accept either RIP v1 or RIP v2 announcements.

**RIP version 2 authentication**
If you use RIP version 2 simple password authentication, then you must configure all of the RIP v2 interfaces on the same network with the same case-sensitive password. You can use the same password for all the networks of your internetwork or you can vary the password for each network.

**RIP version 2 across demand-dial links**
If you use RIP to perform auto-static updates across demand-dial links, then you must configure each demand-dial interface to use RIP v2 multicast announcements and to accept RIP v2 announcements. Otherwise, the RIP request for routes sent by the requesting router is not responded to by the router on the other side of the demand-dial link.

**RIP over Frame Relay**
Since RIP is primarily a broadcast and multicast-based protocol, you need a special configuration for proper RIP operation over a nonbroadcast technology such as Frame Relay. How RIP is configured for Frame Relay depends on how the Frame Relay virtual circuits appear as network interfaces on servers running Routing and Remote Access. Either the Frame Relay adapter appears as a single adapter for all the virtual circuits (single adapter model), or each virtual circuit appears as a separate adapter (multiple adapter model).

**Single-adapter model**
With the single adapter model, also known as the nonbroadcast multiple access (NBMA) model, the network of the Frame Relay service provider (also known as the Frame Relay cloud) is treated as an IP network and the endpoints are assigned IP addresses from a designated IP network ID. To ensure that RIP traffic is received by all of the appropriate endpoints on the cloud, you must configure the Frame Relay interface to unicast its RIP announcements to all of the appropriate endpoints. You

do this by configuring RIP neighbors. For more information about configuring RIP neighbors, see Add a unicast neighbor.

Also, in a spoke and hub Frame Relay topology, the Frame Relay interface for the hub router must have split-horizon processing disabled. Otherwise, the spoke routers never receive each other's routes. For more information, see Set split-horizon processing.

**Multiple-adapter model**

With the multiple adapter model, each Frame Relay virtual circuit appears as a point-to-point link with its own network ID, and the endpoints are assigned IP addresses from a designated IP network ID. Because each virtual circuit is its own point-to-point connection, you can either broadcast (assuming both endpoints are on the same IP network ID) or multicast RIP announcements.

### Note

- The preceding discussion also applies to other nonbroadcast technologies such as X.25 and ATM.

**Silent RIP hosts**

A silent RIP host (a nonrouter) processes received RIP announcements but does not make RIP announcements. The processed RIP announcements are used to build the routing table for the host. You do not need to configure silent RIP hosts with a default gateway. Silent RIP is commonly used in UNIX environments. If there are silent RIP hosts on a network, you must determine which version of RIP they support. If the silent RIP hosts only support RIP v1, then you must use RIP v1 on the network for that host.

# RIP-for-IP security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RIP-for-IP security**

In addition to the security steps listed in Static routing security, you can enhance RIP-for-IP security through:

- RIP version 2 authentication
- Peer security
- Route filters
- Neighbors

**RIP version 2 authentication**

To prevent the corruption of RIP routes by an unauthorized RIP router in a RIP version 2 environment, you can configure RIP v2 router interfaces to use simple password authentication. Received RIP announcements that do not match the configured password are discarded. Note that the password is sent in plaintext. Any user with a network sniffer, such as Microsoft Network Monitor, can capture the RIP v2 announcements and view the password. For more information, see Enable authentication.

**Peer security**

You can configure each RIP router with a list of routers (by IP address) from which RIP announcements are accepted. By default, RIP announcements from all sources are accepted. By configuring a list of RIP peers, RIP announcements from unauthorized RIP routers are discarded. For more information about configuring peer security, see Add peer filters.

**Route filters**

You can configure route filters on each RIP interface so that the only routes considered for addition to the routing table are those that reflect reachable network IDs within the internetwork. For example, if an organization is using subnets of the private network ID 10.0.0.0, route filtering can be used so that the RIP routers discard all routes except those within the 10.0.0.0 network ID. For more information about configuring RIP route filters, see Add route filters.

**Neighbors**

By default, RIP either broadcasts (RIP version 1 or RIP version 2) or multicasts (RIP v2 only) announcements. To prevent RIP traffic from being received by any node except neighboring RIP routers, the server running Routing and Remote Access can unicast RIP announcements. While originally intended for use by nonbroadcast multiaccess (NBMA) network technologies such as Frame Relay, the configuration of RIP neighbors ensures that RIP announcements are directed to neighboring RIP routers. For more information about configuring RIP neighbors, see Add a unicast neighbor.

### Note

- When you configure RIP routers to unicast to neighboring RIP routers, the ability of silent RIP hosts to receive RIP traffic is impaired. Therefore, you need to either add the silent RIP hosts as neighbors or configure the server running Routing and Remote Access to broadcast or multicast (in addition to unicasting to RIP neighbors).

# Deploying RIP for IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying RIP for IP**

While basic RIP version 1 functionality is easy to configure and deploy, RIP version 2 and advanced RIP capabilities, such as peer security and route filtering, require additional configuration and testing. To make troubleshooting and problem isolation easier, it is recommended that you deploy your RIP-based internetwork in the following stages:

1. Set up basic RIP and make sure it is working.

2. Add advanced capabilities one at a time, testing after each feature is added.

**Deploying RIP**

To deploy RIP, you can perform the following steps:

1. Draw a map of the topology of your IP internetwork that shows the separate networks and the placement of routers and hosts (nonrouter computers running TCP/IP).

2. For each IP network (a cabling system bounded by one or more routers), assign a unique IP network ID (also known as an IP network address).

3. Assign IP addresses to each router interface. It is a common industry practice to assign the first IP addresses of an IP network to router interfaces. For example, for an IP network ID of 192.168.100.0 with a subnet mask of 255.255.255.0, the router interface is assigned an IP address of 192.168.100.1.

4. For each router interface, designate whether that interface will be configured for RIP v1 or RIP v2. If an interface is configured for RIP v2, designate whether the RIP v2 announcements will be broadcast or multicast.

5. By using Routing and Remote Access, add the RIP routing protocol and configure the appropriate interfaces for RIP v1 or RIP v2 for each server running Routing and Remote Access.

6. When your configuration is complete, allow a few minutes for the routers to update each other's routing tables and then test the internetwork.

For more information, see Configure RIP for IP.

**Testing a RIP internetwork**

To test your RIP internetwork, you can perform the following steps:

1. To verify that a server running Routing and Remote Access is receiving RIP announcements from all of its neighboring RIP routers, view the RIP neighbors for the router. For more information about viewing RIP neighbors, see View RIP neighbors.

2. For each RIP router, view the IP routing table and verify that all of the routes that should be learned from RIP are present. For more information about viewing the IP routing table, see View routing tables.

3. Use the **ping** and **tracert** commands to test connectivity between host computers so that all routing paths are checked. For more information about the **ping** and **tracert** commands, see Using the ping command and Using the tracert command.

For information about troubleshooting RIP for IP, see Troubleshooting RIP for IP.

# Setting Up an OSPF Routed Internetwork

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up an OSPF routed internetwork**

An Open Shortest Path First (OSPF) routed internetwork uses the OSPF routing protocol to dynamically communicate routing information between routers. Deployed properly, an OSPF environment automatically adds and removes routes as networks are added and removed from the internetwork. You need to ensure that each router is properly configured so that the OSPF-based route advertisements are propagated to OSPF routers on the internetwork. This feature is not available on the Itanium-based versions of the Windows operating systems. This content is not available in this preliminary release.

This section covers:

- The OSPF environment
- OSPF design considerations
- OSPF security
- Deploying OSPF
- Deploying a single-area OSPF internetwork

# The OSPF environment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The OSPF environment**

An Open Shortest Path First (OSPF) routed environment is best suited to a large-to-very-large, multipath, dynamic IP internetwork.

- A large-to-very-large internetwork contains more than 50 networks.
- Multipath means that there are multiple paths for packets to travel between any two endpoints on the internetwork.
- Dynamic means that the topology of the internetwork changes over time because networks are added and removed and links occasionally go down and come back up.

Candidates for an OSPF routed environment include:

- A corporate or institutional campus.
- A worldwide corporate or institutional internetwork.

**Notes**

- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

# OSPF design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**OSPF design considerations**
To prevent problems, you should consider the following design issues before you implement Open Shortest Path First (OSPF). This feature is not available on the Itanium-based versions of the Windows operating systems. This content is not available in this preliminary release.

**OSPF design**
There are three levels of OSPF design:

- Autonomous system design
- Area design
- Network design

To aid in configuration and to prevent problems, you should consider the following elements for each level of OSPF design.

**Autonomous system design**
The following guidelines are recommended when designing an OSPF autonomous system:

- Subdivide the OSPF autonomous system into areas that can be summarized.
- If possible, subdivide your IP address space into a network/area/subnet/host hierarchy.
- Make the backbone area a single high-bandwidth network.
- Create stub areas whenever possible.
- Avoid virtual links whenever possible.

**Area design**
The following guidelines are recommended when designing each OSPF area:

- Ensure that all areas are assigned network IDs that can be expressed as a small number of summary routes.
- If an area can be summarized with a single route, make the area ID the single route being advertised.
- Ensure that multiple area border routers (ABRs) for the same area are summarizing the same routes.
- Ensure that there are no back doors between areas and that all inter-area traffic crosses the backbone area.
- Keep areas under 100 networks.

**Network design**
The following guidelines are recommended when designing each network:

- Assign router priorities so that the least busy routers are the designated router and backup designated router.
- Designate link costs to reflect bit rate, delay, or reliability characteristics.
- Assign a password.

**OSPF over Frame Relay**
Although OSPF packets are typically multicast, OSPF was designed to operate over a nonbroadcast technology such as Frame Relay. How OSPF is configured for Frame Relay depends on how the Frame Relay virtual circuits appear as network interfaces on the server running Routing and Remote Access. Either the Frame Relay adapter appears as a single adapter for all the virtual circuits (single adapter model) or each virtual circuit appears as a separate adapter (multiple adapter model).

**Single-adapter model**
With the single adapter model, also known as the nonbroadcast multiple access (NBMA) model, the network for the Frame Relay service provider (also known as the Frame Relay cloud) is treated as an IP network and the endpoints on the cloud are assigned IP addresses from a designated IP network ID. To ensure that OSPF traffic is received by all of the appropriate endpoints on the cloud, you must configure the Frame Relay interface to unicast its OSPF announcements to all of the appropriate endpoints. For the server running Routing and Remote Access, this is done by designating the interface as a nonbroadcast multiple access (NBMA) network and adding OSPF neighbors. To add OSPF neighbors, see Add an OSPF neighbor.

Also, in a spoke and hub Frame Relay topology, the Frame Relay interface for the hub router must have a router priority set to 1 or greater and the Frame Relay interfaces for the spoke routers must have a router priority set to 0. Otherwise, the hub router, which is the only router that can communicate with all of the spoke routers, cannot become the designated router and adjacencies cannot form across the Frame Relay network.

**Multiple-adapter model**
With the multiple adapter model, each Frame Relay virtual circuit appears as a point-to-point link with its own network ID, and the endpoints are assigned IP addresses from a designated IP network ID. Since each virtual circuit is its own point-to-point connection, you can configure the interface for the point-to-point network type.

### Note

- The preceding discussion also applies to other NBMA technologies such as X.25 and ATM.

**Using virtual links**
An OSPF routed internetwork can be subdivided into areas, which are collections of contiguous networks. All areas are connected together through a common area called the backbone area. A router that connects an area to the backbone area is called an area border router (ABR). Normally, ABRs have a physical connection to the backbone area. When it is not possible or practical to have an ABR of an area physically connected to the backbone area, you can use a virtual link to connect the ABR to the backbone.

A virtual link is a logical point-to-point connection between an ABR of an area and an ABR that is physically connected to the backbone area. For example, a virtual link is configured between the ABR of Area 2 and the ABR of Area 1. The ABR of Area 1 is physically connected to the backbone area. Area 1 is known as the transit area, the area across which the virtual link is created in order to logically connect Area 2 to the backbone.

To create a virtual link, both routers, called virtual link neighbors, are configured with the transit area, the router ID of the virtual link neighbor, matching hello and dead intervals, and a matching password.

For more information, see Add a virtual interface.

To troubleshoot a virtual link, see Troubleshooting OSPF.

**External routes and ASBRs**
The set of OSPF routers in an organization defines an OSPF autonomous system (AS). By default, only OSPF routes corresponding to directly connected network segments are propagated within the AS. An external route is any route that is not within the OSPF AS. External routes can come from many sources:

- Other routing protocols such as RIP for IP (v1 and v2).
- Static routes.
- Routes set on the router through SNMP.

External routes are propagated throughout the OSPF AS through one or more autonomous system boundary routers (ASBRs). An ASBR advertises external routes within the OSPF AS. For example, if you need to advertise the static routes of a server running Routing and Remote Access, you need to enable that router as an ASBR.

For more information, see Configure an ASBR.

**External route filters**
By default, OSPF routers acting as ASBRs import and advertise all external routes. You may want to filter out external routes to keep the ASBR from advertising improper routes. External routes can be filtered on the ASBR by:

- External route source

  You can configure the ASBR to accept or ignore the routes of certain external sources such as routing protocols (RIP v2) or other sources (static routes or SNMP).

- Individual route

  You can configure the ASBR to accept or discard specific routes by configuring one or multiple {Destination, Network mask} pairs.

**OSPF on a remote access server**
If you configure a server running Routing and Remote Access using OSPF to also act as a remote access server and the static IP address pool address ranges are for a separate subnet, in order to properly advertise the route that represents all remote access clients:

1. Enable the server running Routing and Remote Access as an autonomous system boundary router (ASBR).

2. Configure OSPF route filters to **Accept listed routes** and then add the routes that correspond to the address ranges of your static IP address pool.

For more information, see Configure an ASBR.

**OSPF on an answering demand-dial router**
If you configure a server running Routing and Remote Access by using OSPF to act as the answering router in a one-way initiated connection, in order to properly advertise the routes that represent the network segments of your calling routers:

1. Enable the answering router as an autonomous system boundary router (ASBR).

2. If the answering router is also configured with static IP address pool address ranges that are for a separate subnet, add the routes that correspond to the static routes on the user accounts being used by the calling routers.

For more information, see Configure an ASBR.

For more information about using static routes for one-way initiated demand-dial connections, see One-way initiated demand-dial connections.

# OSPF security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**OSPF security**

In addition to the security steps listed in Static routing security, you can enhance Open Shortest Path First (OSPF) security through:

- Authentication
- External route filters on ASBRs

**Authentication**

By default, OSPF interfaces on the server running Routing and Remote Access are configured to send the simple password of "12345678" in their OSPF Hello messages. The simple password helps prevent the corruption of OSPF data from an unauthorized OSPF router on a network. The password is sent in plaintext. Any user with a network sniffer, such as Microsoft Network Monitor, can capture the OSPF Hello messages and view the password.

**External route filters on ASBRs**

To prevent the propagation of invalid routes into the OSPF autonomous system (AS) from external sources such as RIP or static routes, you can configure autonomous system boundary routers (ASBRs) with route filters. You can configure ASBR route filters so that any route that matches a configured list is discarded, or any route that does not match a configured list is discarded.

For more information, see Configure an ASBR.

### Notes

- You can only use external route filters to filter routes from non-OSPF sources. There is no capability to filter OSPF routes within the OSPF autonomous system.
- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.

# Deploying OSPF

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying OSPF**

Deploying Open Shortest Path First (OSPF) requires careful planning and configuration at three levels:

- The autonomous system
- The area
- The network

**Planning the autonomous system**

For the OSPF autonomous system (AS), you need to:

1. Subdivide the OSPF AS into areas that can be easily summarized by using summary routes.
2. Designate the backbone area.
3. Assign area IDs.
4. Identify virtual links.
5. Identify area border routers (ABRs).
6. Identify stub areas.
7. Identify autonomous system boundary routers (ASBRs).

**Planning each area**

For each router, you need to:

1. Add the areas to which the router is connected.
2. If the area is a stub area, enable the area as a stub area.
3. If the router is an ABR, optionally configure the ranges that summarize the IP networks within the area.
4. If the router is an ABR that uses a virtual link, add the virtual interface.
5. If the router is an ASBR, enable ASBR and configure optional external route filters.

**Planning each network**

For each IP address for each router interface that uses OSPF, you need to:

1. Add the interface to the OSPF routing protocol.
2. Enable OSPF on the interface.
3. Configure the interface for the appropriate area ID.
4. Configure the interface for the appropriate router priority.

5. Configure the interface for the appropriate link cost.

6. Configure the interface for the appropriate password.

7. Configure the interface for the appropriate network type.

8. If the interface is a single-adapter Frame Relay (or X.25 or ATM) interface, configure the nonbroadcast multiple access (NBMA) neighbors.

For more information, see Configure OSPF.

**Testing OSPF**

To test your OSPF internetwork, you can perform the following steps:

1. To verify that a server running Routing and Remote Access is receiving OSPF announcements from all of its adjacent OSPF routers, view the OSPF neighbors for the router. For more information, see View OSPF neighbors.

2. For each router, view the IP routing table and verify that all of the routes that should be learned from OSPF are present. For more information, see View routing tables.

3. Use the **ping** and **tracert** commands to test connectivity between host computers so that all routing paths are checked. For more information about the **ping** and **tracert** commands, see Using the ping command and Using the tracert command.

For information about troubleshooting OSPF, see Troubleshooting OSPF.

### Notes

- This feature is not available on the Itanium-based versions of the Windows operating systems.

- This content is not available in this preliminary release.

# Deploying a single-area OSPF internetwork

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying a single-area OSPF internetwork**
While Open Shortest Path First (OSPF) is a routing protocol designed to scale to very large internetworks, the planning and implementation of a large scale OSPF internetwork is complex and time-consuming. However, you do not need a large or very large internetwork to take advantage of the advanced features of OSPF.

In the Windows Server 2003 family implementation of OSPF, the default values of global and interface settings make it very easy to create a single-area OSPF internetwork with minimal configuration. The single area is the backbone area (0.0.0.0).

**OSPF default global settings**
The OSPF default global settings are as follows:

- The router identification is set to the IP address of the first IP binding at the time of the installation of the OSPF routing protocol.
- The router is not configured as an OSPF autonomous system boundary router.
- A single area, the backbone area, is configured and enabled for a plaintext password. It is not configured as a stub area and there are no address ranges.
- There are no configured virtual interfaces, and external routes are not filtered.

For a single-area OSPF internetwork, no changes to the OSPF default global settings are required.

**OSPF default interface settings**
The OSPF default interface settings for LAN interfaces are as follows:

- By default, OSPF is *not* enabled to run over the interface.
- The area ID is set to the backbone area (0.0.0.0).
- The router priority is set to 1. With multiple OSPF routers on the same network with the same router priority, the OSPF designated router and backup designated router are elected based on the router with the highest router ID.
- The cost is set to 2.
- The password is set to 12345678.
- The network type is set to broadcast for LAN interfaces.
- There are no configured neighbors.
- The hello interval is set to 10 seconds.
- The dead interval is set to 40 seconds.

For a single-area OSPF internetwork, the only required change to the OSPF default interface settings is to enable OSPF to run over the interface.

**Deploying OSPF**
To deploy an OSPF single-area internetwork that consists of LAN interfaces, perform the following steps on each server running Routing and Remote Access:

1. Enable the Routing and Remote Access service. For more information, see Enable the Routing and Remote Access service.

2. Add the OSPF routing protocol. For more information, see Add the OSPF routing protocol.

3. Add the routing interfaces of the router to the OSPF routing protocol, enabling OSPF for each interface. For more information, see Add an interface to OSPF.

**Testing OSPF**

To test your single-area OSPF internetwork, you can perform the following steps:

1. To verify that a server running Routing and Remote Access is receiving OSPF announcements from all of its adjacent OSPF routers, view the OSPF neighbors for the router. For more information, see View OSPF neighbors.

2. For each router, view the IP routing table and verify that all of the routes that should be learned from OSPF are present. For more information, see View routing tables.

3. Use the **ping** and **tracert** commands to test connectivity between host computers so that all routing paths are checked. For more information about the **ping** and **tracert** commands, see Using the ping command and Using the tracert command.

For information about troubleshooting OSPF, see Troubleshooting OSPF.

### Notes

- This feature is not available on the Itanium-based versions of the Windows operating systems.

- This content is not available in this preliminary release.

# Setting Up Network Address Translation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up network address translation**

This section covers:

- Network address translation design considerations
- Deploying network address translation

# Network address translation design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network address translation design considerations**
To prevent problems, you should consider the following design issues before you implement network address translation.

**Private network addressing**
You should use the following IP addresses from the Internet Assigned Numbers Authority (IANA) private IP network IDs: 10.0.0.0 with a subnet mask of 255.0.0.0, 172.16.0.0 with a subnet mask of 255.240.0.0, and 192.168.0.0 with a subnet mask of 255.255.0.0. By default, network address translation uses the private network ID 192.168.0.0 with the subnet mask of 255.255.255.0 for the private network.

For more information, see Enable network address translation addressing.

If you are using public IP addresses that have not been allocated by IANA or your ISP, then you may be using the IP network ID of another organization on the Internet. This is known as illegal or overlapping IP addressing. If you are using overlapping public addresses, then you cannot reach the Internet resources of the overlapping addresses. For example, if you use 1.0.0.0 with the subnet mask of 255.0.0.0, then you cannot reach any Internet resources of the organization that is using the 1.0.0.0 network.

You can also exclude specific IP addresses from the configured range. Excluded addresses are not allocated to private network hosts.

**Single or multiple public addresses**
If you are using a single public IP address allocated by your ISP, no other IP address configuration is necessary. If you are using multiple IP addresses allocated by your ISP, then you must configure the network address translation (NAT) interface with your range of public IP addresses. For the range of IP addresses given to you by your ISP, you must determine whether the range of public IP addresses can be expressed by using an IP address and a mask.

If you are allocated a number of addresses that is a power of 2 (2, 4, 8, 16, and so on), you might be able to express the range by using a single IP address and mask. For example, if you are given the four public IP addresses 206.73.118.212, 206.73.118.213, 206.73.118.214, and 206.73.118.215 by your ISP, then you can express these four addresses as 206.73.118.212 with a mask of 255.255.255.252.

If your IP addresses are not expressible as an IP address and a subnet mask, you can enter them as a range or series of ranges by indicating the starting and ending IP addresses.

For more information, see Configure interface IP address ranges.

**Allowing inbound connections**
Normal NAT usage from a home or small business allows outbound connections from the private network to the public network. Some programs that run from the private network create connections to Internet resources. The return traffic from the Internet can be translated because the connection was initiated from the private network.

To allow Internet users to access resources on your private network, you must do the following:

- Configure a static IP address configuration on the resource server including IP address (from the range of IP addresses allocated by the NAT computer), subnet mask (from the range of IP addresses allocated by the NAT computer), default gateway (the private IP address of the NAT computer), and DNS server (the private IP address of the NAT computer).
- Exclude the IP address being used by the resource computer from the range of IP addresses being allocated by the NAT computer.
- Configure a special port. A special port is a static mapping of a public address and port number to a private address and port number. A special port maps an inbound connection from an Internet user to a specific address on your private network. By using a special port, you can create a Web server on your private network that is accessible from the Internet.

For more information, see Configure services and ports.

**Configuring applications and services**

You may need to configure applications and services to work properly across the Internet. For example, if users on your small office or home office (SOHO) network want to play the game Asheron's Call with other users on the Internet, network address translation must be configured for Asheron's Call.

For more information, see Configure services and ports.

**VPN connections from a translated SOHO network**

To access a private intranet using a virtual private network (VPN) connection from a translated SOHO network, you can use the Point-to-Point Tunneling Protocol (PPTP) and create a VPN connection from a host on the SOHO network to the VPN server of the private intranet on the Internet.

The Windows Server 2003 family allows L2TP/IPSec VPN connections to work with NAT. For more information, see Virtual Private Networks.

# Deploying network address translation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying network address translation**

To deploy network address translation for a small office or home office network, you need to configure:

- The network address translation computer.
- Other computers on the small office or home network.

**Configuring the network address translation computer**

To configure the network address translation (NAT) computer, you can complete the following steps:

1. Install and enable the Routing and Remote Access service.

   In the Routing and Remote Access Server Setup Wizard, choose **Network address translation (NAT)**. After the wizard is finished, all of the configuration for NAT is complete. You do not need to complete steps 2 through 8.

   If you have already enabled the Routing and Remote Access service, then complete steps 2 through 8 as needed.

   For information about installing and enabling the Routing and Remote Access service, see Enable the Routing and Remote Access service.

2. Configure the IP address of the private network interface.

   For the IP address of the LAN adapter that connects to the home network, you need to configure the following:

   - IP address: 192.168.0.1
   - Subnet mask: 255.255.255.0
   - No default gateway

   **Note**

   - The IP address in the preceding configuration for the home network interface is based on the default address range of 192.168.0.0 with a subnet mask of 255.255.255.0, which is configured for the addressing component of network address translation. If you change this default address range, you should change the IP address of the private interface for the network address translation computer to be the first IP address in the configured range. Using the first IP address in the range is a recommended practice, not a requirement of the network address translation components.

3. Enable routing on your dial-up port.

   If your connection to the Internet is a permanent connection that appears as a LAN interface (such as DDS, T-Carrier, Frame Relay, permanent ISDN, xDSL, or cable modem) or if you are connecting your server running Routing and Remote Access to another router before the connection to the Internet, and the LAN interface is configured with an IP address, subnet mask, and default gateway either statically or through DHCP, skip to step 6.

   For information about enabling routing on your dial-up port, see Enable routing on ports.

4. Create a demand-dial interface to connect to your Internet service provider.

You need to create a demand-dial interface that is enabled for IP routing and uses your dial-up equipment and the credentials that you use to dial your Internet service provider (ISP). For more information about creating demand-dial interfaces, see Add a demand-dial interface.

5. Create a default static route that uses the Internet interface.

   For a default static route, you need to select the demand-dial interface (for dial-up connections) or LAN interface (for permanent or intermediate router connections) that is used to connect to the Internet. The destination is 0.0.0.0 and the network mask is 0.0.0.0. For a demand-dial interface, the gateway IP address is not configurable.

   For more information about configuring a default static route, see Add a default static IP route.

6. Add the network address translation (NAT) routing protocol.

   For information about adding the network address translation (NAT) IP routing protocol, see Add network address translation.

7. Add your Internet and home network interfaces to the NAT routing protocol.

   For information about adding interfaces to the NAT IP routing protocol, see Add and configure an interface for network address translation.

8. Enable NAT addressing and name resolution.

   For information about enabling network address translation addressing, see Enable network address translation addressing.

   For information about enabling network address translation name resolution, see Enable network address translation name resolution.

   **Note**

   - The NAT addressing feature only assigns addresses from a single range that corresponds to a single subnet. If multiple home network LAN interfaces are added to the NAT routing protocol, a single subnet configuration (where all LAN interfaces are connected to the same network) is assumed. If the LAN interfaces correspond to different networks, connectivity between clients on different networks who receive addresses from the NAT computer may not be possible.

**Configuring other computers on the small office or home network**

You need to configure the TCP/IP protocol on the other computers on the small office or home network to obtain an IP address automatically, and then restart them. When the computers on the home network receive their IP address configuration from the network address translation computer, they are configured with:

- IP address (from the address range of 192.168.0.0 with a subnet mask of 255.255.255.0).
- Subnet mask (255.255.255.0).
- Default gateway (the IP address of a directly reachable IP router on the small office or home network).
- DNS server (the IP address of the interface for the NAT computer on the small office or home network).

**Advanced NAT settings**

To configure advanced NAT settings, you can do the following:

- If you have been given a range of IP addresses from your ISP, configure the range of IP addresses on your Internet interface.

- If there are services running on the private network that need to be accessed by users from the Internet, add a special port that maps the public IP address and port number to a private IP address and port number.

For more information, see Configure NAT/Basic Firewall.

For information about troubleshooting network address translation, see Troubleshooting network address translation.

# Setting Up Demand-Dial Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up demand-dial routing**

This section covers:

- Demand-dial routing design considerations
- Demand-dial routing security
- Deploying demand-dial routing
- Deploying certificate-based authentication for demand-dial routing

# Demand-dial routing design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Demand-dial routing design considerations**
To prevent problems, you should consider the following design issues before you implement demand-dial routing.

**On-demand or persistent connections**
You must decide whether your demand-dial connections will be on-demand or persistent:

- On-demand demand-dial connections are used when the cost of using the communications link is time-sensitive. For example, long distance analog phone calls are charged on a per-minute basis. With on-demand connections, the connection is made when traffic is forwarded, and the connection is terminated when the link is not used. You can configure idle disconnect behavior for the answering router by setting an idle disconnect through the profile properties of the remote access policy that is used for the demand-dial connection. You can configure idle disconnect behavior for the calling router on the **Options** tab on the properties of the demand-dial interface.

- Persistent demand-dial connections use a dial-up link but can be left in a connected state 24 hours a day without incurring additional usage charges. Examples of persistent demand-dial connections include local calls that use analog phone lines and flat-rate ISDN.

**One-way or two-way initiated connections**
You must decide whether your demand-dial connections will be initiated by one router or by both routers:

- With one-way initiated connections, one router is always the answering router and one router is always the calling router. The answering router accepts the connection and the calling router initiates the connection. One-way initiated connections are well suited to an on-demand spoke and hub topology where the branch office router is the only router that initiates the connection. One-way initiated connections require the following:

    - The answering router is configured as a LAN and demand-dial router.

    - A user account is added for the authentication credentials of the calling router that is accessed and validated by the answering router.

    - A demand-dial interface is configured at the answering router with the same name as the user account that is used by the calling router. The demand-dial interface is not used to dial out; therefore it is not configured with the phone number of the calling router or with valid user credentials.

  For an alternate configuration that does not require a demand-dial interface on the answering router, see One-way initiated demand-dial connections.

- With two-way initiated connections, either router can be the answering router or the calling router depending on who is initiating the connection. Both routers must be configured to initiate and accept a demand-dial connection. You can use two-way initiated connections when traffic from either router can create the demand-dial connection. Two-way initiated demand-dial connections require the following:

    - Both routers are configured as LAN and demand-dial routers.

    - User accounts are added for both routers so that the authentication credentials of the calling router are accessed and validated by the answering router.

    - Demand-dial interfaces, with the same name as the user account that is used by the calling router, are fully configured at both routers, including the phone number of the answering router and user account credentials to authenticate the calling router.

**Restricting the initiation of an on-demand connection**
To prevent the calling router from making unnecessary on-demand dial-up connections, which may result in excessive phone charges, you can restrict the calling router from making connections in two ways:

- Demand-dial filtering

  You can use demand-dial filtering to configure either the types of IP traffic that do not cause a connection to be made or the types of IP traffic that cause a connection to be made. For more information, see Configure demand-dial filters.

- Dial-out hours

  You can use dial-out hours to configure the hours that a calling router is either permitted or denied to make a demand-dial connection. For more information, see Configure dial-out hours.

**IP packet filters and demand-dial filters**
Demand-dial filters are applied before the connection is made. IP packet filters are applied after the connection is made. To prevent the demand-dial connection from being established for traffic that is discarded by the IP packet filters:

- If you have configured a set of output IP packet filters with the **Receive all packets except those that meet the criteria listed below** option, then configure the same set of filters as demand-dial filters with **Initiate connection** set to **For all traffic except**.

- If you have configured a set of output IP packet filters with the **Drop all packets except those that meet the criteria listed below** option, then configure the same set of filters as demand-dial filters with **Initiate connection** set to **Only for the following traffic**.

**Routing**
Both routers on a demand-dial connection must have the appropriate routes in their routing tables to forward traffic across the connection. Additionally, routes in the routers of the intranets of both demand-dial routers must be present that support the two-way forwarding of traffic between any two destinations. Routes can be static or dynamic.

Static routing is recommended for on-demand connections. You can add static routes to the routing table either manually or through an auto-static update. For more information, see Add a static route or Demand-dial routing updates.

Dynamic routing is recommended for persistent connections by adding the demand-dial interface to an IP routing protocol.

**Creating a remote access policy for demand-dial connections**
By using remote access policies, you can create a policy that requires demand-dial connections to use a specific authentication method and encryption strength.

For example, you can create a user group account called Demand-Dial Routers whose members are the user accounts that are used by calling routers when a demand-dial connection is created. Then, you create a policy with one condition: the **Windows-Group** is set to **Demand-Dial Routers**. Finally, you configure the profile for the policy to select a specific authentication method and encryption strength.

For more information, see Introduction to remote access policies.

**User accounts for demand-dial connections**
When you create a user account by using the Demand-Dial Interface wizard, the remote access permission is set to **Allow access** even though for a new account in a Windows 2000 or Windows Server 2003 domain or a stand-alone router, the default remote access permission for newly created accounts is set to **Control access through Remote Access Policy**. This behavior may cause some confusion if you are using the access-by-policy administrative model. In the access-by-policy administrative model, the remote access permission of all user accounts is set to **Control access through Remote Access Policy** and the remote access permission of individual policies are set to either **Grant remote access permission** or **Deny remote access permission**. For more information, see Remote Access Policies Examples.

When the user account is created, it is created with the current default password settings and policies set for your domain. Verify that the user accounts that are used by calling routers have the following settings on the **General** tab on the properties of the user account:

- The **User must change password at next logon** check box is cleared.

If the check box is selected, then you must manually clear this setting. Accounts created with the Demand-Dial Interface wizard automatically have this check box cleared. If you do not clear this setting, then a demand-dial router cannot connect by using this account. When the calling router sends its credentials, the calling router is prompted to change the password. Because the initiation of a demand-dial connection is not an interactive process involving a user, the calling router is unable to change the password and aborts the connection attempt.

- The **Password never expires** check box is selected.

  Because the demand-dial connection process is not interactive, if the password expires, the calling router is prompted to change the expired password and the connection attempt is ended.

# Demand-dial routing security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Demand-dial routing security**
Security is less of a concern for demand-dial connections than for router-to-router VPN connections because data is not traveling across a public network like the Internet. However, data may be intercepted as it travels through the infrastructure of your telecommunications provider.

In addition to the security steps listed in Static routing security, you can enhance demand-dial routing security through:

- Strong authentication
- Data encryption

**Strong authentication**
For authentication, use the strongest authentication scheme that is possible for your demand-dial configuration. The strongest authentication scheme is the use of EAP-TLS with certificates. For more information, see Deploying certificate-based authentication for demand-dial routing.

Otherwise, use MS-CHAP v2 authentication and enforce the use of strong passwords on your network. For more information, see Enable authentication protocols.

**Data encryption**
For encryption, you can use either link encryption or end-to-end encryption:

- Link encryption encrypts the data only on the link between the two routers. You can use 128-bit Microsoft Point-to-Point Encryption (MPPE). 40-bit MPPE can be used for compatibility with older versions of Microsoft operating systems. You must use MPPE in conjunction with either MS-CHAP or EAP-TLS authentication.
- End-to-end encryption encrypts the data between the source host and its final destination. You can use Internet Protocol security (IPSec) to encrypt data from the source host to the destination host across the demand-dial link.

  For more information, see Internet Protocol Security (IPSec).

To require encryption, clear the **No encryption** option and select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile that is used by your calling routers. For more information, see Configure encryption.

# Deploying demand-dial routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying demand-dial routing**

In order to create a two-way, initiated demand-dial routing connection from a branch office router to a corporate office router, you must perform the following:

1. Configure the corporate office router to initiate and receive demand-dial connections with the branch office router.
2. Configure the branch office router to initiate and receive demand-dial connections with the corporate office router.
3. Initiate the demand-dial connection from either the branch office router or the corporate office router.

### Notes

- This deployment assumes a demand-dial connection between a branch office router and a corporate office router. You can also apply this deployment to the demand-dial connections between two corporate offices.
- For maximum flexibility, the demand-dial connection is a two-way initiated connection that is initiated by either the branch office router or the corporate office router.

The following illustration shows elements of Routing and Remote Access that provide demand-dial routed connections.




**Configuring the corporate office router**

If you want your corporate office router to support two-way initiated demand-dial connections, complete the following steps:

- Configure the connection to the corporate office intranet.
- Configure the LAN and WAN router.
- Configure ports to allow demand-dial connections.
- Configure demand-dial interfaces.
- Configure static routes.
- Configure remote access policies.

**Configuring the connection to the intranet**

The connection to the intranet is a LAN adapter that is installed in the computer. You need to verify that the LAN adapter is compatible with your operating system. For more information, see Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.
- IP addresses of the corporate DNS and WINS servers.

Because the corporate office router will act as a router between the corporate office and the branch office, it must be configured with either static routes or with routing protocols so that all of the destinations on the corporate network are reachable from the corporate office router.

**Configuring the LAN and demand-dial router**

You can enable LAN and demand-dial routing by installing the Routing and Remote Access service and using the Routing and Remote Access Server Setup Wizard. For more information, see Enable the Routing and Remote Access service.

If you have already run the Routing and Remote Access Server Setup Wizard but did not choose the routing options, you can configure LAN and demand-dial routing through the properties of the router in Routing and Remote Access. For more information, see View properties of the remote access server.

To allow demand-dial connections, you need to either run the Routing and Remote Access Server Setup Wizard and choose the routing options, or configure the following settings manually:

- **General**

  Verify that the **Router** check box and **LAN and demand-dial routing** are selected for the server running Routing and Remote Access.

- **Security**

  - **Authentication Methods**

    Select the authentication methods that are supported by the router to authenticate the credentials of demand-dial routers. For servers running Routing and Remote Access configured as demand-dial routers, select either MS-CHAP v2 or EAP (if smart cards or machine certificates are available) authentication.

  - **Authentication Provider**

    You can verify the credentials of demand-dial routers by using the security features of the Windows Server 2003 family or a RADIUS server. If RADIUS is selected, you need to configure RADIUS server settings for your RADIUS server or RADIUS proxy.

  - **Accounting Provider**

    You can record demand-dial router activity for analysis or accounting purposes by selecting and configuring an accounting provider.

- **IP**

  Click **Static address pool** and configure the ranges of IP addresses that are dynamically allocated to demand-dial routers.

- **PPP**

  Select the **Link control protocol (LCP) extensions** check box.

  Select the **Software compression** check box.

**Configuring demand-dial interfaces**

For each branch office router, you can create a demand-dial interface by using the Demand-Dial Interface wizard. In the wizard, enter the following:

- **Interface name**

The name of the interface that represents the connection to the branch office and that has the same name as the user account used by the branch office router. For example, for a router in the New York branch office, enter **NewYorkRouter**.

- **Connection type**

  Click **Connect using a modem, ISDN adapter, or other physical device**.

- **Select a device**

  Click the device being used to create the connection.

- **Phone number or address**

  Because the corporate router can also initiate a demand-dial connection, the phone number of the branch office router is required.

- **Protocols and security**

  Select the protocols to route and the **Add a user account so a remote router can dial in** check boxes.

- **Dial-in credentials**

  Type the domain and password for the account that will be used to authenticate the branch office router. The Demand-Dial Interface wizard automatically creates the account and sets its remote access permission to **Allow access**. The name of the account is the same as the name of the demand-dial interface. For example, for the New York branch office router, the name of the account is **NewYorkRouter**.

- **Dial-out credentials**

  Because the corporate router can also initiate a demand-dial connection, the name, domain, and password are required. The user credentials that you enter are used to authenticate the corporate office router when it initiates a connection with the branch office router. For example, for the corporate office router, the name of the account is **CorpHub**.

**Configuring static routes**

You need to add static routes so that traffic to the branch office is forwarded by using the appropriate demand-dial interface. For each route of each branch office, configure the interface, destination, network mask, and metric. For interface, you need to select the demand-dial interface that corresponds to the branch office.

For example, the route that corresponds to the New York branch office is 192.168.25.0 with a subnet mask of 255.255.255.0. This route becomes the static route with the following configuration:

- Interface: NewYorkRouter
- Destination: 192.168.25.0
- Network mask: 255.255.255.0
- Metric: 1

**Note**

- Because the demand-dial connection is a point-to-point connection, the gateway IP address is not configurable.

For more information, see Add a static route.

**Configuring remote access policies**
By using the Demand-Dial Interface wizard, the dial-in properties of user accounts that are used by branch office routers are already configured to allow remote access.

If you want to grant remote access to the demand-dial-based branch office routers based on group membership, do the following:

1. For a stand-alone router, use Local Users and Groups and set dial-in properties to **Allow access** for all users.

2. For a router in a domain, use Active Directory Users and Computers and set dial-in properties to **Control access through Remote Access Policy** for all users.

3. Create a user group whose members can create demand-dial connections with the corporate office router. For example, create a user group named BranchOfficeRouters.

4. Add the appropriate user accounts that correspond to the accounts that are used by the branch office routers to the group.

5. Delete the default remote access policies.

6. Create a new remote access policy with the following properties:

    - Set **Policy name** to **Demand-dial connection if member of BranchOfficeRouters** (example).

    - Set the **Windows-Groups** attribute to **BranchOfficeRouters** (example).

    - Set the **NAS-Port-Type** condition to all except **Virtual (VPN)**.

The default settings for encryption on the **Encryption** tab on the properties of a remote access policy profile are to allow no encryption and all levels of encryption strength. To require encryption for demand-dial connections, clear the **No Encryption** option and select the encryption strength you want to use.

For more information, see Configure encryption.

**Configuring the branch office router**
If you want your branch office router to support two-way initiated demand-dial connections, complete the following steps:

- Configure the connection to the branch office intranet.

- Configure the LAN and WAN router.

- Configure ports to allow demand-dial connections.

- Configure a demand-dial interface.

- Configure static routes.

- Configure remote access policies.

**Configuring the connection to the branch office intranet**
The connection to the branch office network is a LAN adapter installed in the computer. To verify whether the LAN adapter is compatible with your operating system, see Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.

- IP addresses of the branch office DNS and WINS servers.

Because the branch office router will act as a router between the corporate office and the branch office, it must be configured with either static routes or with routing protocols so that all of the destinations on the branch office network are reachable from the branch office router.

**Configuring the LAN and demand-dial router**

You can enable the LAN and demand-dial router by installing the Routing and Remote Access service and using the Routing and Remote Access Server Setup Wizard. For more information, see Enable the Routing and Remote Access service.

If you have already run the Routing and Remote Access Server Setup Wizard but did not choose the routing options, you can configure LAN and demand-dial routing through the properties on the router in Routing and Remote Access. For more information, see View properties of the remote access server.

To allow demand-dial connections, you need to either run the Routing and Remote Access Server Setup Wizard and choose the routing options, or configure the following settings manually:

- **General**

  Verify that the **Router** check box and **LAN and demand-dial routing** are selected.

- **Security**

  - **Authentication Methods**

    Select the authentication methods that are supported by the router to authenticate the credentials of demand-dial routers. For servers running Routing and Remote Access configured to act as demand-dial routers, select either MS-CHAP v2 or EAP (if smart cards or machine certificates are available) authentication.

  - **Authentication Provider**

    You can verify the credentials of demand-dial routers by using the security features of the Windows Server 2003 family or a RADIUS server. If RADIUS is selected, you need to configure RADIUS server settings for your RADIUS server or RADIUS proxy.

  - **Accounting Provider**

    You can record demand-dial router activity for analysis or accounting purposes by selecting and configuring an accounting provider.

- **IP**

  Click **Static address pool** and configure the ranges of IP addresses that are dynamically allocated to demand-dial routers.

- **PPP**

  Select the **Link control protocol (LCP) extensions** check box.

  Select the **Software compression** check box.

**Configuring a demand-dial interface**

You can create a demand-dial interface by using the Demand-Dial Interface wizard. In the wizard, enter the following:

- **Interface name**

  The name of the interface is the name of the user account configured on the corporate server running Routing and Remote Access and used by the branch office router to connect to the corporate network. For example, enter **CorpHub**.

- **Connection type**

Click **Connect using a modem, ISDN adapter, or other physical device**.

- **Select a device**

  Click the device being used to create the connection.

- **Phone number or address**

  Because the branch office router can also initiate a demand-dial connection, the phone number of the corporate office router is required.

- **Protocols and security**

  Select the protocols to route and the **Add a user account so a remote router can dial in** check boxes.

- **Dial-in credentials**

  Type the domain and password for the account that will be used to authenticate the corporate office router. The Demand-Dial Interface wizard automatically creates the account and sets its remote access permission to **Allow access**. The name of the account is the same as the name of the demand-dial interface. For example, for the corporate office router, the name of the account is **CorpHub**.

- **Dial-out credentials**

  Because the branch office router can also initiate a demand-dial connection, the name, domain, and password are required. The user credentials that you enter are used to authenticate the branch office router when it initiates a connection with the corporate office router. For example, for the branch office router in the New York branch office, the name of the account is **NewYorkRouter**.

**Note**

- In order for two-way initiated demand-dial routing to work properly, the user name of the calling router must match the name of a demand-dial interface on both sides of the connection. The following table shows an example.

| Router | Demand-dial interface name | User account name |
|---|---|---|
| Corporate office router | NewYorkRouter | CorpHub |
| Branch office router | CorpHub | NewYorkRouter |

**Configuring static routes**

You need to add static routes so that traffic to the corporate office is forwarded by using the appropriate demand-dial interface. For each route of the corporate office, configure the interface, destination, network mask, and metric. For interface, you need to select the demand-dial interface that corresponds to the corporate office previously created.

For example, the route that corresponds to the corporate office is 10.0.0.0 with a subnet mask of 255.0.0.0. To configure this route as a static route, set the following:

- Interface: CorpHub
- Destination: 10.0.0.0
- Network mask: 255.0.0.0
- Metric: 1

**Note**

- Because the demand-dial connection is a point-to-point connection, the gateway IP address is not configurable.

For more information, see Add a static route.

**Configuring remote access policies**
By using the Demand-Dial Interface wizard, the dial-in properties of user accounts that are used by corporate office routers are already configured to allow remote access.

If you want to grant remote access to the demand-dial-based corporate office routers based on group membership, do the following:

1. For a stand-alone router, use Local Users and Groups and set dial-in properties to **Allow access** for all users.

2. For a router in a domain, use Active Directory Users and Computers and set dial-in properties to **Control access through Remote Access Policy** for all users.

3. Create a user group whose members can create demand-dial connections with the branch office router. For example, create a user group named DemandDialRouters.

4. Add the appropriate user accounts that correspond to the accounts that are used by the corporate office router or other branch office routers to the group.

5. Delete the default remote access policies.

6. Create a new remote access policy with the following properties:

    - Set **Policy name** to **Demand-dial connection if member of DemandDialRouters** (example).

    - Set the **Windows-Groups** attribute to **DemandDialRouters** (example).

    - Set the **NAS-Port-Type** condition to all except **Virtual (VPN)**.

The default settings for encryption on the **Encryption** tab on the properties of a remote access policy profile are to allow no encryption and all levels of encryption strength. To require encryption for demand-dial connections, clear the **No Encryption** option and select the encryption strength you want to use.

For more information, see Configure encryption.

**Initiating the demand-dial connection**
To connect the branch office router to the corporate office router, do one of the following:

- From the branch office router

  In Routing and Remote Access, right-click the demand-dial interface that connects to the corporate office (in this example, the CorpHub demand-dial interface), and then click **Connect**.

- From the corporate office router

  In Routing and Remote Access, right-click the demand-dial interface that connects to the branch office (in this example, the NewYorkRouter demand-dial interface), and then click **Connect**.

For information about troubleshooting demand-dial routing, see Troubleshooting demand-dial routing.

# Deploying certificate-based authentication for demand-dial routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying certificate-based authentication for demand-dial routing**

The use of certificates for authentication of calling routers is the strongest form of authentication in the Windows Server 2003 family. For certificate-based authentication of demand-dial connections, you must use the Extensible Authentication Protocol (EAP) with the **Smart card or other certificate (TLS)** EAP type, also known as EAP-Transport Level Security (EAP-TLS). EAP-TLS requires the use of user certificates for the calling router and machine certificates (also known as computer certificates) for the answering router.

The deployment of certificate-based authentication for demand-dial routing typically occurs in the following situations:

- Business partner demand-dial connection
- Branch office demand-dial connection

# Business partner demand-dial connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Business partner demand-dial connection**
To use certificates for a two-way initiated, mutually authenticated, demand-dial configuration between two business partners (in this example, Company A and Company B), you must perform the following:

- Configure the calling and answering routers for demand-dial routing.
- Install computer certificates on the calling router and answering router computers.
- Configure the domain for Web-based certificate enrollment.
- At Company A, create a user account for the Company B router and export a certificate for the user account.
- At Company B, create a user account for the Company A router and export a certificate for the user account.
- At Company A, import the certificate from Company B.
- Configure the Company A router to support certificate-based authentication as a calling router and as an answering router.
- At Company B, import the certificate from Company A.
- Configure the Company B router to support certificate-based authentication as a calling router and as an answering router.

**Configuring the calling and answering routers for demand-dial routing**
Configure the Routing and Remote Access calling and answering routers as described in Deploying demand-dial routing for dial-in demand-dial routing or Router-to-router VPN Deployment for VPN demand-dial routing.

**Installing computer certificates on the calling router and answering router computers**
To use EAP-TLS, a computer certificate (also known as a machine certificate) must be installed on the authenticating server and the remote access client. In order to install a computer certificate, a certification authority must be present to issue certificates. Once the certification authority is configured, you can install a certificate in three different ways:

- By configuring the automatic enrollment, or autoenrollment, of computer certificates to computers in a Windows Server 2003 domain.
- By using the Certificates snap-in to obtain a computer certificate.
- By using your browser to connect to the CA Web enrollment pages to install a certificate on the local computer or to a floppy disk for installation on another computer, such as non-domain member computers that cannot obtain a certificate through autoenrollment.

Based on the certificate policies in your organization, you only need to perform one of these allocations.

For more information, see Network access authentication and certificates.

To configure a certification authority and install the computer certificate, perform the following steps:

1. Install the Certificate Services component as an enterprise root certification authority (CA). This step is only necessary if you do not already have an enterprise root CA.

    a. If necessary, promote the computer that will be a CA to a domain controller (DC).

b. Install the Certificate Services component as an enterprise root CA. For more information, see Install an enterprise root certification authority.

2. Configure the CA to issue certificates that permit exportable keys. To do so, you must clone and rename the **router (offline request)** template, choose to make the keys exportable, and add the new template to the templates the CA can use to issue certificates. For more information, see To establish the certificate types that an enterprise certification authority can issue.

3. Do one of the following:

   - To auto-enroll computer certificates, configure the domain. For more information, see Configure automatic certificate allocation from an enterprise CA.

     To create a computer certificate for the calling or answering router that is a member of the domain for which auto-enrollment is configured (as well as other computers that are members of the domain), restart the computer or type **gpupdate /Target:Computer /Force** from the command prompt.

   - To manually enroll computer certificates, use the Certificates snap-in or the CA Web enrollment pages to install the CA root certificate. For more information, see Manage certificates for a computer and Request a certificate.

In order for the CA to issue certificates for the calling router, you must configure the domain for Web-based enrollment. For more information, see Set up certification authority Web enrollment support.

**Creating a user account and exporting its certificate for the Company B router**

To create a dial-in user account for the Company B router and export the user certificate of the user account, do the following:

1. Log on as a domain administrator.

2. Create a user account that the Company B router will use when it dials the Company A router. For more information, see Create a new user account.

3. Obtain a certificate that has an exportable key from the certification authority through Web-based enrollment. This certificate might be called **router (offline request)**, or have another name. For more information, see Install a router (offline request) certificate.

4. Export the exportable key certificate to a .cer file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, do not export the private key.

5. Map the newly created certificate (the .cer file) to the user account that was created for the Company B router. For more information, see Map a certificate to a user account.

6. Export the certificate to a .pfx file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, export the private key, select the **Delete the private key if the import is successful** check box, and click **Include all certificates in the certification path if possible**. Save this file to a floppy disk to send to the network administrator at Company B.

7. Send the floppy disk that contains the Company B dial-in account user certificate file to the network administrator at Company B.

**Creating a user account and exporting its certificate for the Company A router**

To create a dial-in user account for the Company A router and export the user certificate of the user account, do the following:

1. Log on as a domain administrator.

2. Create a user account that the Company A router will use when it dials the Company B router. For more information, see

Create a new user account.

3. Obtain a certificate that has exportable keys from the certification authority through Web-based enrollment. This certificate might be called **router (offline request)**, or it might have another name. For more information, see Install a router (offline request) certificate.

4. Export the certificate to a .cer file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, do not export the private key.

5. Map the newly created certificate (the .cer file) to the user account created for the Company A router. For more information, see Map a certificate to a user account.

6. Export the certificate to a .pfx file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, export the private key, select the **Delete the private key if the import is successful** check box, and click **Include all certificates in the certification path if possible**. Save this file to a floppy disk to send to the network administrator at Company A.

7. Send the floppy disk that contains the Company A dial-in account user certificate file to the network administrator at Company A.

**Importing the certificates from Company B**
Upon receipt at Company A of the floppy disk that contains the certificate file from Company B, on the Company A router, import the user certificate. For more information, see Import a certificate.

**Configuring the Company A router to support certificate-based authentication**
To configure the Company A router for certificate-based authentication as an answering router, see Configure the answering router for certificate-based EAP.

To configure the Company A router for certificate-based authentication as a calling router, see Configure the calling router for certificate-based EAP.

**Importing the certificates from Company A**
Upon receipt at Company B of the floppy disk that contains the certificate files from Company A, on the Company B router, import the user certificate. For more information, see Import a certificate.

**Configuring the Company B router to support certificate-based authentication**
To configure the Company B router for certificate-based authentication as an answering router, see Configure the answering router for certificate-based EAP.

To configure the Company B router for certificate-based authentication as a calling router, see Configure the calling router for certificate-based EAP.

# Branch office demand-dial connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Branch office demand-dial connection**
To use certificates for a two-way initiated, mutually authenticated, demand-dial configuration between two routers in the same organization (in this example, a branch office router and a corporate office router), you must perform the following:

- Configure the calling and answering routers for demand-dial routing.
- Install a computer certificate on the corporate office router.
- Configure the domain for Web-based certificate enrollment.
- Create user accounts and export certificates.
- Import the dial-out user certificate on the corporate office router.
- Configure the corporate office router to support certificate-based authentication as a calling router and as an answering router.
- Import the dial-in certificate on the branch office router.
- Configure the branch office router to support certificate-based authentication as a calling router.
- Connect to the corporate office and join the organization domain.

**Configuring the calling and answering routers for demand-dial routing**
Configure the calling and answering routers as described in Deploying demand-dial routing for dial-up demand-dial routing or Router-to-router VPN Deployment for VPN demand-dial routing.

**Installing a computer certificate on the corporate office router**
In order to configure EAP-TLS on the corporate office router, you must install a computer certificate (also known as a machine certificate). In order to install a computer certificate, a certification authority must be present to issue certificates. Once the certification authority is configured, you can install a certificate in three different ways:

- By configuring the automatic enrollment, or autoenrollment, of computer certificates to computers in a Windows Server 2003 domain.
- By using the Certificates snap-in to obtain a computer certificate.
- By using your browser to connect to the CA Web enrollment pages to install a certificate on the local computer or to a floppy disk for installation on another computer, such as non-domain member computers that cannot obtain a certificate through autoenrollment.

Based on the certificate policies in your organization, you only need to perform one of these allocations.

For more information, see Network access authentication and certificates.

To configure a certification authority and install the computer certificate, perform the following steps:

1. Install the Certificate Services component as an enterprise root certification authority. This step is only necessary if you do not already have an enterprise root certification authority (CA).
    a. If necessary, promote the computer that will be a CA to a domain controller (DC).
    b. Install the Certificate Services component as an enterprise root CA. For more information, see Install an enterprise root

certification authority.

2. Configure the CA to issue certificates with exportable keys. For more information, see To establish the certificate types that an enterprise certification authority can issue.

3. Do one of the following:

   - To auto-enroll computer certificates, configure the domain. For more information, see Configure automatic certificate allocation from an enterprise CA.

     To create a computer certificate for the calling or answering router that is a member of the domain for which autoenrollment has been configured (as well as other computers that are members of the domain), restart the computer or type **gpupdate /Target:***Computer***/Force** from the command prompt.

   - To manually enroll computer certificates, use the Certificates snap-in or the CA Web enrollment pages to install the CA root certificate. For more information, see Manage certificates for a computer and Request a certificate.

In order for the CA to issue certificates for the calling router, you must configure the domain for Web-based enrollment. For more information, see Set up certification authority Web enrollment support.

**Creating user accounts and exporting certificates**

To create dial-in and dial-out user accounts and export certificates, do the following:

1. Log on as a domain administrator.

2. Create a user account that the corporate office router will use when it dials the branch office router (the dial-out account). For more information, see Create a new user account.

3. Obtain a certificate that has exportable keys for the dial-out account from the certification authority through Web-based enrollment. This certificate might be called **router (offline request)**, or it might have another name. For more information, see Install a router (offline request) certificate.

4. Export the certificate for the dial-out account to a .cer file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, do not export the private key.

5. Map the newly created certificate (the .cer file) to the dial-out user account. For more information, see Map a certificate to a user account.

6. Export the certificate of the dial-out account to a .pfx file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, export the private key and click **Delete the private key if the import is successful** and select the option to **Include all certificates in the certification path if possible**.

7. Create a user account that the branch office router will use when it dials the corporate office router (the dial-in account). For more information, see Create a new user account.

8. Obtain a certificate that has exportable keys for the dial-in account from the certification authority through Web-based enrollment. This certificate might be called **router (offline request)**, or it might have another name. For more information, see Install a router (offline request) certificate.

9. Export the certificate for the dial-in account to a .cer file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, do not export the private key.

10. Map the newly created certificate (the .cer file) to the dial-in user account. For more information, see Map a certificate to a user account.

11. Export the certificate of the dial-in account to a .pfx file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, export the private key and click **Delete the private key if the import is successful**. Save this file to a floppy disk to send to the network administrator at the branch office.

12. Send the floppy disk that contains the dial-in account user certificate file to the network administrator at the branch office.

**Importing the dial-out certificate on the corporate office router**
On the corporate office router, import the user certificate for the dial-out account. For more information, see Import a certificate.

**Configuring the corporate office router to support certificate-based authentication**
To configure the corporate office router for certificate-based authentication as an answering router, see Configure the answering router for certificate-based EAP.

To configure the corporate office router for certificate-based authentication as a calling router, see Configure the calling router for certificate-based EAP.

**Importing the certificate on the branch office router**
Upon receipt at the branch office of the floppy disk that contains the certificate file from the corporate office, import the user certificate for the dial-in account. For more information, see Import a certificate.

**Configuring the branch office router to support certificate-based authentication**
To configure the branch office router for certificate-based authentication as a calling router, see Configure the calling router for certificate-based EAP.

**Connecting to the corporate office and joining the organization domain**
To connect to the corporate office and join the organization domain, do the following:

1. From the branch office, connect to the corporate office by right-clicking the demand-dial interface, and then clicking **Connect**.

2. Once connected, the branch office router joins the domain through the **Computer Name** tab (in the properties of My Computer).

3. After joining the domain, restart the branch office router.

4. After restarting the branch office router, connect to the corporate office router again.

5. Once connected, the branch office router receives domain policy and a computer certificate (if auto-enrollment of computer certificates is configured). If auto-enrollment of computer certificates is not configured, obtain a computer certificate through the Certificates snap-in. For more information, see Manage certificates for a computer and Request a certificate.

6. Once a computer certificate is obtained, configure the branch office router for certificate-based authentication as an answering router. For more information, see Configure the answering router for certificate-based EAP.

At this point, you can install a domain controller in the branch office by using the demand-dial connection to the corporate office. For more information on installing a domain controller, see Checklist: Installing a domain controller.

### Notes

- The ability of the branch office router to join the domain and the installation of a domain controller depends on DNS name resolution. Ensure that both the router and the domain controller computer are configured with the proper DNS server IP addresses.

- By default, an answering router checks the certificate revocation list when authenticating the calling router. Because the root CA computer is always reachable by the corporate office router, the certificate revocation list can always be checked. However, the root CA computer is not reachable by the branch office router until after the connection is made. If the root CA computer cannot be reached, then Active Directory is checked. In this case, the branch office router accesses its local domain controller for the revocation list. If the certificate revocation list is not published in Active Directory, then the branch office router acting as the answering router rejects the connection attempt. To prevent this problem, do one of the following:

    - Publish the certificate revocation list in Active Directory. For more information, see Schedule the publication of the certificate revocation list or Manually publish the certificate revocation list.

    - On the branch office router, set the following registry value to 1:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RasMan\PPP\EAP\13\IgnoreRevocationOffline**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

# Deploying Router-to-Router VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying router-to-router VPNs**

This section covers:

- Planning for router-to-router VPNs
- Router-to-router VPN Deployment

# Planning for router-to-router VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for router-to-router VPNs**

This section covers:

- Router-to-router VPN design considerations
- Router-to-router VPN security considerations

# Router-to-router VPN design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Router-to-router VPN design considerations**
To prevent problems, you should consider the following design issues before you implement router-to-router VPN connections:

- PPTP-based or L2TP/IPSec connections
- Installing certificates
- On-demand or persistent connections
- Restricting the initiation of an on-demand connection
- One-way or two-way initiated connections
- Number of PPTP or L2TP ports needed
- Configuring firewall packet filters
- Routing
- Single hop across VPN connection
- Creating a remote access policy for router-to-router VPN connections

For more information, see Router-to-router VPN connection.

**PPTP-based or L2TP/IPSec connections**
A VPN server running Windows Server 2003 provides support for both Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP). When choosing between PPTP and L2TP/IPSec router-to-router VPN solutions, consider the following:

- PPTP can be used for router-to-router VPN connections for routers running Windows NT Server 4.0 with the Routing and Remote Access Service (RRAS), Windows 2000 Server, or Windows Server 2003 operating systems. PPTP does not require a public key infrastructure (PKI) to issue computer certificates. By using encryption, PPTP-based VPN connections provide data confidentiality--captured packets cannot be interpreted without the encryption key. PPTP-based VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data origin authentication (proof that the data was sent by the authorized user).

- L2TP can be used only with routers running Windows 2000 Server or Windows Server 2003 operating systems. When both types of routers are used, a public key infrastructure (PKI) is required to issue computer certificates to all routers. Routers running Windows Server 2003 operating systems additionally support a single preshared key configured on the answering router and all calling routers. By using Internet Protocol Security (IPSec), L2TP/IPSec VPN connections provide data confidentiality, data integrity, and data origin authentication.

**Installing certificates**
To use certificate authentication for L2TP/IPSec router-to-router VPN connections, you must install computer certificates, also known as machine certificates, on both routers. For more information, see Computer certificates for L2TP/IPSec VPN connections and Network access authentication and certificates.

**On-demand or persistent connections**
You must decide whether your router-to-router VPN connections will be on-demand or persistent:

- On-demand demand-dial connections require that the answering router is permanently connected to the Internet. The calling router connects to the Internet by using a dial-up link such as an analog phone line or ISDN. You need to configure a single demand-dial interface at the answering router. You need to configure two demand-dial interfaces at the calling router: one to connect to a local Internet service provider (ISP) and one for the router-to-router VPN connection. Demand-dial router-to-router VPN connections also require an additional host route in the IP routing table of the calling router. For more information, see A dial-up router-to-router VPN connection.

- Persistent connections require that both routers are connected to the Internet by using permanent WAN connections. You only need to configure a single demand-dial interface at each router. Permanent connections can be initiated and left in a connected state 24 hours a day.

**Restricting the initiation of an on-demand connection**

To prevent the calling router from making unnecessary connections, you can restrict the calling router from making on-demand router-to-router VPN connections in two ways:

- Demand-dial filtering

  You can use demand-dial filtering to configure either the types of IP traffic that do not cause a demand-dial connection to be made or the types of IP traffic that cause a connection to be made. For more information, see Configure demand-dial filters.

- Dial-out hours

  You can use dial-out hours to configure the hours that a calling router is either permitted or denied to make a router-to-router VPN connection. For more information, see Configure dial-out hours.

**One-way or two-way initiated connections**

You must decide whether your router-to-router VPN connections will be initiated by one router or by both routers:

- With one-way initiated connections, one router is the VPN server and one router is the VPN client. The VPN server accepts the connection and the VPN client initiates the connection. One-way initiated connections are well suited to a permanent connection spoke-and-hub topology where the branch office router is the only router that initiates the connection. One-way initiated connections require that:

  - The VPN server (the answering router) is configured as a LAN and demand-dial router.

  - A user account is added for the authentication credentials of the calling router that is accessed and validated by the answering router.

  - A demand-dial interface is configured at the answering router with the same name as the user account that is used by the calling router. This demand-dial interface is not used to dial out, therefore it is not configured with the host name or IP address of the calling router or with valid user credentials.

  For more information, see One-way initiated demand-dial connections.

- With two-way initiated connections, either router can be the VPN server or the VPN client depending on who is initiating the connection. Both routers must be configured to initiate and accept a VPN connection. You can use two-way initiated connections when the router-to-router VPN connection is not up 24 hours a day and traffic from either router is used to create the on-demand connection. Two-way initiated router-to-router VPN connections require that:

  - Both routers are connected to the Internet by using a permanent WAN link.

  - Both routers are configured as LAN and demand-dial routers.

  - User accounts are added for both routers so that the authentication credentials for the calling router are accessed and validated by the answering router.

- Demand-dial interfaces, with the same name as the user account that is used by the calling router, must be fully configured at both routers, including settings for the host name or IP address of the answering router and user account credentials to authenticate the calling router.

**Configuring firewall packet filters**
If you have a firewall, you must configure packet filters on the firewall to allow traffic between the VPN router and the routers on the Internet. For more information, see VPN servers and firewall configuration.

**Routing**
Both routers on a router-to-router VPN connection must have the appropriate routes in their routing tables to forward traffic across the connection. Routes can be static or dynamic. You can add static routes to the routing table either manually or through an auto-static update. You can add dynamic routes to the routing table by adding the VPN connection demand-dial interface to a routing protocol. However, enabling a routing protocol on the VPN connection demand-dial interface is only recommended when the demand-dial interface is permanently connected.

### Note

- Unlike demand-dial routing by using direct links, you cannot use a default IP route for the VPN demand-dial interface to summarize all the routes of the corporate office. Because the branch office router is connected to the Internet, the default route must be used to summarize all the routes of the Internet and configured to use the interface that connects the router to the Internet.

**Single hop across VPN connection**
For the purposes of designing a routed infrastructure, you can consider the router-to-router VPN connection as a single hop regardless of how many routers are crossed when the encapsulated data is sent across the Internet.

**Creating a remote access policy for router-to-router VPN connections**
By using remote access policies, you can create a policy that requires router-to-router VPN connections to use a specific authentication method and encryption strength.

For example, you can create an Active Directory group called VPN Routers whose members are the user accounts that are used by calling routers when a router-to-router VPN connection is created. Then, you create a policy with two conditions on the policy: the **NAS-Port-Type** is set to **Virtual (VPN)** and the **Windows-Group** is set to **VPN Routers**. Finally, you configure the profile for the policy to select a specific authentication method and encryption strength.

You can also use the **Tunnel-Type** condition to create separate remote access policies for PPTP and L2TP connections. For example, to require a specific authentication method and encryption strength for PPTP connections, set the **Tunnel-Type** condition to **Point-to-Point Tunneling Protocol**.

# Router-to-router VPN security considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Router-to-router VPN security considerations**
In addition to the security steps listed in Static routing security, you can enhance router-to-router VPN security through:

- Strong authentication
- Data encryption
- PPTP or L2TP/IPSec packet filtering
- Firewall packet filtering

For more information, see Router-to-router VPN connection.

**Strong authentication**
For authentication, use the strongest authentication scheme that is possible for your router-to-router VPN configuration. The strongest authentication scheme is the use of EAP-TLS with certificates. For more information, see Deploying certificate-based authentication for demand-dial routing.

Otherwise, use MS-CHAP version 2 authentication and enforce the use of strong passwords on your network. For more information, see MS-CHAP version 2.

**Data encryption**
For encryption, you can use either link encryption or end-to-end encryption:

- Link encryption encrypts the data only on the link between the two routers. You can use 128-bit Microsoft Point-to-Point Encryption (MPPE) for high security with PPTP connections. Otherwise, you can use either 56-bit or 40-bit MPPE. 40-bit MPPE is used with older versions of Microsoft operating systems. MPPE is only used in conjunction with either MS-CHAP, MS-CHAP v2, or EAP-TLS authentication. L2TP/IPSec connections use DES or 3DES encryption.
- End-to-end encryption encrypts the data between the source host and its final destination. You can use IPSec to encrypt data from the source host to the destination host across the demand-dial link.

To specify encryption strength, select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile that is used by your calling routers. The **No Encryption** option is cleared by default; do not select this option or your network security is in jeopardy. For more information, see Security issues for VPN and Configure encryption.

**Packet filtering**
To secure the calling or answering corporate router from sending or receiving any traffic on its Internet interface except router-to-router VPN traffic, you need to ensure that PPTP or L2TP/IPSec input and output filters are configured on the router interface that corresponds to the connection to the Internet.

When you configure your VPN server using the Routing and Remote Access Server Setup Wizard, PPTP or L2TP/IPSec input and output filters are configured automatically. For more information about PPTP filters, see Add PPTP Filters. For more information about L2TP/IPSec filters, see Add L2TP over IPSec Filters.

Because IP routing is enabled on the Internet interface, if PPTP or L2TP/IPSec filters are not configured on the Internet interface, then any traffic received on the Internet interface is routed, which may forward unwanted Internet traffic to your intranet.

**Firewall packet filtering**
It is a common practice to provide protection to intranet hosts, such as a VPN server, from Internet hosts using a firewall. If you have a firewall, you must configure packet filters on the firewall to allow traffic between the VPN router and the routers on the Internet. For more information, see VPN servers and firewall configuration.

# Router-to-router VPN Deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Router-to-router VPN deployment**
This section covers:

- PPTP-based router-to-router VPN deployment
- L2TP-based router-to-router VPN deployment

# PPTP-based router-to-router VPN deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPTP-based router-to-router VPN deployment**

To create a PPTP-based router-to-router VPN connection to send private data across the Internet, you must perform the following:

1. Configure the router running a Windows Server 2003 operating system at the corporate office to receive PPTP connections from a branch office router.
2. Configure the router running a Windows Server 2003 operating system at the branch office to initiate a PPTP connection with the corporate office router.
3. Initiate the PPTP connection from the branch office router.

**Note**

- These steps assume that the PPTP-based router-to-router VPN connection is between a corporate office and a branch office. However, you can also apply these steps to the VPN connection between two corporate offices.

**Configuring the corporate office router**

If you want your router running a Windows Server 2003 operating system in the corporate to support multiple branch office PPTP connections, complete the following steps:

- Configure the connection to the Internet.
- Configure the connection to the intranet.
- Configure the corporate router.
- Configure demand-dial interfaces.
- Configure firewall packet filters.
- Configure remote access policies.

The following illustration shows the elements of a PPTP-based router-to-router VPN connection on a computer running a Windows Server 2003 operating system.




For more information, see Router-to-router VPN connection and Point-to-Point Tunneling Protocol.

**Note**

- To simplify configuration, the branch office router always initiates the PPTP connection.

**Configuring the connection to the Internet**
The connection to the Internet is a dedicated connection--a WAN adapter that is installed in the computer. The WAN adapter is typically a DDS, T1, Fractional T1, or Frame Relay adapter. You must contract with a local telephone company to run the appropriate physical wiring to your premises. You need to verify that the WAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

The WAN adapter includes drivers that are installed in Windows Server 2003 operating systems so that the adapter appears as a network adapter.

You need to configure the following TCP/IP settings on the WAN adapter:

- IP address and subnet mask assigned from the InterNIC or an Internet service provider (ISP).
- Default gateway of the ISP router.

**Configuring the connection to the intranet**
The connection to the intranet is a LAN adapter that is installed in the computer. You need to verify that the LAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.
- DNS and WINS name servers of corporate intranet name servers.

Because the corporate router will route traffic between the corporate office and the branch office, you must configure the corporate router with either static routes or with routing protocols so that all of the destinations on the corporate network are reachable from the corporate router.

**Configuring the corporate router**
You need to enable the corporate router by installing the Routing and Remote Access service. For more information, see Enable the Routing and Remote Access service.

**Configuring demand-dial interfaces**
For each branch office router, you can create a demand-dial interface by using the Demand-Dial Interface Wizard. In the wizard, configure the following:

- **Interface Name**

  The name of the interface that represents the connection to the branch office. For example, for a router in the New York branch office, type **NewYorkRouter**.

- **Connection Type**

  Click **Connect using virtual private networking (VPN)**.

- **VPN Type**

  Click **Point-to-Point Tunneling Protocol (PPTP)**.

- **Destination Address**

  Because the corporate router will not initiate the VPN connection, no phone number or address is required.

- **Protocols and Security**

  Select the **Add a user account so a remote router can dial in** check box.

- **Configuring static routes**

  You need to add static routes so that traffic to the branch office is forwarded by using the appropriate demand-dial interface. For each route of each branch office, configure the interface, destination, network mask, and metric. For the interface, you need to select the demand-dial interface that corresponds to the branch office.

  For example, the route that corresponds to the New York branch office is 192.168.25.0 with a subnet mask of 255.255.255.0. This route becomes the static route with the following configuration:

  - Interface: NewYorkRouter
  - Destination: 192.168.25.0
  - Network mask: 255.255.255.0
  - Metric: 1

- **Dial-out Credentials**

  Because the corporate router will not initiate the VPN connection, type in any name, domain, and password.

- **Dial-in Credentials**

  Type the domain and password for the account that will be used to authenticate the branch office router. The Demand-Dial Interface Wizard automatically creates the account and sets its remote access permission to **Allow access**. The name of the account is the same as the name of the demand-dial interface. For example, for the New York branch office router, the name of the account is **NewYorkRouter**.

**Note**

- Because the PPTP connection is a point-to-point connection, the Gateway IP address is not configurable.

For more information, see Add a static route.

**Configuring firewall packet filters**
If you are using a firewall in the corporate office, you need to configure PPTP packet filters on your firewall to allow PPTP traffic between the branch office routers and the corporate office router. For more information, see VPN servers and firewall configuration.

**Configuring remote access policies**
By using the Demand-Dial Interface Wizard, the dial-in properties of user accounts that are used by branch office routers are already configured to allow remote access.

If you want to grant remote access to the PPTP-based branch office routers based on group membership, do the following:

1. For a stand-alone router that is not a member of a domain, use Local Users and Groups and set dial-in properties to **Allow access** for all users.

2. For a directory services-based router, use Active Directory Users and Computers and set dial-in properties to **Control access through Remote Access Policy** for all users.

3. Create an Active Directory group whose members can create virtual private networking connections with the VPN server. For example, BranchOfficeRouters.

4. Add the appropriate user accounts that correspond to the accounts that are used by the branch office routers to the Active Directory group.

5. Create a new remote access policy with the following properties:

    - Set **Policy name** to **VPN Access if member of BranchOfficeRouters** (example).
    - Set the **Windows-Groups** condition to **BranchOfficeRouters** (example).
    - Set the **NAS-Port-Type** condition to **Virtual (VPN)**.
    - Set the **Tunnel-Type** condition to **Point-to-Point Tunneling Protocol**.
    - Select the **Grant remote access permission** option.

6. If this computer is only used to provide router-to-router VPN connections, you need to delete the default remote access policies. Otherwise, move the default remote access policy so that it is evaluated last.

For encryption, the default setting allows no encryption and all levels of encryption strength. To require encryption, clear the **No Encryption** option and select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile that is used by your calling routers.

For more information, see Configure encryption.

**Configuring the branch office router**
If you want your router running a Windows Server 2003 operating system in the branch office to initiate a PPTP connection with the corporate office router, complete the following steps:

- Configure the connection to the Internet.
- Configure the connection to the branch office network.
- Configure a demand-dial interface.
- Configure static routes.
- Configure firewall packet filters.

### Note

- To simplify this configuration, the branch office router always initiates the PPTP connection.

**Configuring the connection to the Internet**
The connection to the Internet is a dedicated connection--a WAN adapter that is installed in the computer. The WAN adapter is typically a DDS, T1, Fractional T1, or Frame Relay adapter. You must contract with a local telephone company to run the appropriate physical wiring to your premises. You need to verify that the WAN adapter is compatible with the server operating systems. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

The WAN adapter includes drivers that are installed in Windows Server 2003 operating systems so that the adapter appears as a network adapter.

You need to configure the following TCP/IP settings on the WAN adapter:

- IP address and subnet mask assigned from the InterNIC or an Internet service provider (ISP).
- Default gateway of the ISP router.

**Configuring the connection to the branch office network**
The connection to the branch office network is a LAN adapter that is installed in the computer. You need to verify that the WAN

adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.
- DNS and WINS name servers of branch office name servers.

**Configuring a demand-dial interface**

You can create a demand-dial interface by using the Demand-Dial Interface Wizard. In the wizard, configure the following:

- **Interface Name**

  Type the name of the interface that represents the connection to the corporate office. For example, type **CorpOffice**.

- **Connection Type**

  Click **Connect using virtual private networking (VPN)**.

- **VPN Type**

  Click **Point-to-Point Tunneling Protocol (PPTP)**.

- **Destination Address**

  Type the IP address or host name that is assigned to the Internet interface of the router at the corporate office. If you enter a host name, verify that the host name resolves to the proper IP address.

- **Dial-out Credentials**

  Type the name, domain name, and password of the user account that corresponds to this branch office router. The credentials are the same as those entered in the **Dial-Out Credentials** page of the Demand-Dial Interface Wizard when the demand-dial interface for this branch office was created on the corporate router.

**Configuring static routes**

You need to add static routes so that traffic to the corporate office is forwarded by using the appropriate demand-dial interface. For each route of the corporate office, configure the interface, destination, network mask, and metric. For the interface, select the demand-dial interface that corresponds to the corporate office previously created.

For example, the route that corresponds to the corporate office is 10.0.00 with a subnet mask of 255.0.0.0. This route becomes a static route with the following configuration:

- Interface: CorpOffice
- Destination: 10.0.0.0
- Network mask: 255.0.0.0
- Metric: 1

**Note**

- Because the PPTP connection is a point-to-point connection, the Gateway IP address is not configurable.

For more information, see Add a static route.

**Configuring firewall packet filters**

If you are using a firewall in the branch office, you need to configure PPTP packet filters on your firewall to allow PPTP traffic between the corporate office router and the branch office router. For more information, see VPN servers and firewall configuration.

**Initiating the PPTP router-to-router VPN connection**

To test the branch office router to the corporate router, in Routing and Remote Access, right-click the demand-dial interface that connects to the corporate office, and then click **Connect**.

For information about troubleshooting a router-to-router VPN, see Troubleshooting router-to-router VPNs.

# L2TP-based router-to-router VPN deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**L2TP-based router-to-router VPN deployment**

To create an L2TP/IPSec router-to-router VPN connection to send private data across the Internet, you must perform the following:

1. Configure the router running a Windows Server 2003 operating system at the corporate office to receive L2TP connections from a branch office router.
2. Configure the router running a Windows Server 2003 operating system at the branch office to initiate an L2TP connection with the corporate office router.
3. Initiate the L2TP connection from the branch office router.

**Note**

- These steps assume that the L2TP/IPSec router-to-router VPN connection is between a corporate office and a branch office. However, you can also apply these steps to a VPN connection between two corporate offices.

**Configuring the corporate office router**

If you want your router running a Windows Server 2003 operating system in the corporate office to support multiple branch office L2TP connections, complete the following steps:

- Configure the connection to the Internet.
- Configure the connection to the intranet.
- Install a computer certificate.
- Configure the corporate router.
- Configure demand-dial interfaces.
- Configure firewall packet filters.
- Configure remote access policies.

The following illustration shows the elements of a L2TP/IPSec router-to-router VPN connection on a computer running a Windows Server 2003 operating system.




For more information, see Router-to-router VPN connection and Layer Two Tunneling Protocol.

**Note**

- To simplify this configuration, the branch office router always initiates the L2TP connection.

**Configuring the connection to the Internet**

The connection to the Internet is a dedicated connection--a WAN adapter that is installed in the computer. The WAN adapter is typically a DDS, T1, Fractional T1, or Frame Relay adapter. You must contract with a local telephone company to run the appropriate physical wiring to your premises. You need to verify that the WAN adapter is compatible with the Windows Server 2003 operating systems. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

The WAN adapter includes drivers that are installed in Windows Server 2003 operating systems so that the adapter appears as a network adapter.

You need to configure the following TCP/IP settings on the WAN adapter:

- IP address and subnet mask assigned from the InterNIC or an Internet service provider (ISP).
- Default gateway of the ISP router.

**Configuring the connection to the intranet**

The connection to the intranet is a LAN adapter that is installed in the computer. You need to verify that the LAN adapter is compatible with the Windows Server 2003 operating systems. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

You need to configure the following TCP/IP setting on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.

If the corporate router is running applications, you need to configure the following TCP/IP setting on the LAN adapter:

- DNS and WINS name servers of corporate intranet name servers.

Because the corporate router will route traffic between the corporate office and the branch office, you must configure the corporate router with either static routes or with routing protocols so that all of the destinations on the corporate network are reachable from the corporate router.

**Installing a computer certificate**

Although you can use either computer certificates or a preshared key to provide authentication for IPSec security associations for an L2TP/IPSec connection, computer certificates are the recommended method. Therefore, you must install a computer certificate on the corporate router.

For more information about installing a computer certificate on the corporate router, see Computer certificates for L2TP/IPSec VPN connections and Network access authentication and certificates.

**Configuring the corporate router**

You need to enable the corporate router by installing the Routing and Remote Access service. For more information, see Enable the Routing and Remote Access service.

**Configuring demand-dial interfaces**

For each branch office router, you can create a demand-dial interface by using the Demand-Dial Interface Wizard. In the wizard, configure the following:

- **Interface Name**

  Type the name of the interface that represents the connection to the branch office. For example, for a router in the New York branch office, type **NewYorkRouter**.

- **Connection Type**

  Click **Connect using virtual private networking (VPN)**.

- **VPN Type**

  Click **Layer 2 Tunneling Protocol (L2TP)**.

- **Destination Address**

  Because the corporate router will not initiate the VPN connection, no address is required.

- **Protocols and Security**

  Select the protocols you want to route, and then select the **Add a user account so a remote router can dial in** check box.

- **Configuring static routes**

  You need to add static routes so that traffic to the branch office is forwarded by using the appropriate demand-dial interface. For each route of each branch office, configure the interface, destination, network mask, and metric. For the interface, you need to select the demand-dial interface that corresponds to the branch office.

  For example, the route that corresponds to the New York branch office is 192.168.25.0 with a subnet mask of 255.255.255.0. This route becomes the static route with the following configuration:

  - Interface: NewYorkRouter
  - Destination: 192.168.25.0
  - Network mask: 255.255.255.0
  - Metric: 1

- **Dial-out Credentials**

  Because the corporate router will not initiate the VPN connection, type in any name, domain, and password.

- **Dial-in Credentials**

  Type the domain and password for the account that will be used to authenticate the branch office router. The Demand-Dial Interface Wizard automatically creates the account and sets its remote access permission to **Allow access**. The name of the account is the same as the name of the demand-dial interface. For example, for the New York branch office router, the name of the account is **NewYorkRouter**.

**Note**

- Because the L2TP connection is a point-to-point connection, the Gateway IP address is not configurable.

For more information, see Add a static route.

**Configuring firewall packet filters**
If you are using a firewall in the corporate office, you need to configure L2TP/IPSec packet filters on your firewall to allow L2TP/IPSec traffic between branch office routers and the corporate office router. For more information, see VPN servers and firewall configuration.

**Configuring remote access policies**
By using the Demand-Dial Interface Wizard, the dial-in properties of user accounts that are used by branch office routers are already configured to allow remote access.

If you want to grant remote access to the L2TP/IPSec branch office routers based on group membership, do the following:

1. For a stand-alone router that is not a member of a domain, use Local Users and Groups and set dial-in properties to **Allow access** for all users.

2. For a directory services-based router, use Active Directory Users and Computers and set dial-in properties to **Control access through Remote Access Policy** for all users.

3. Create a Windows Server 2003 Active Directory group whose members can create virtual private networking connections with the VPN server. For example, BranchOfficeRouters.

4. Add the appropriate user accounts that corresponds to the accounts that are used by the branch office routers to the Active Directory group.

5. Create a new remote access policy with the following properties:

   - Set **Policy name** to **VPN Access if member of BranchOfficeRouters** (example).
   - Set the **Windows-Groups** condition to **BranchOfficeRouters** (example).
   - Set the **NAS-Port-Type** condition to **Virtual (VPN)**.
   - Set the **Tunnel-Type** condition to **Layer Two Tunneling Protocol**.
   - Select the **Grant remote access permission** option.

6. If this computer is only used to provide router-to-router VPN connections, you need to delete the default remote access policies. Otherwise, move the default remote access policy so that it is evaluated last.

For encryption, the default setting allows no encryption and all levels of encryption strength. To require encryption, clear the **No Encryption** option and select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile that is used by your calling routers.

For more information, see Configure encryption.

**Configuring the branch office router**

If you want your router running a Windows Server 2003 operating system in the branch office to initiate an L2TP connection with the corporate office router, complete the following steps:

- Configure the connection to the Internet.
- Configure the connection to the branch office network.
- Install a computer certificate.
- Configure a demand-dial interface.
- Configure static routes.
- Configure firewall packet filters.

**Note**

- To simplify configuration, the branch office router always initiates the L2TP connection.

**Configuring the connection to the Internet**

The connection to the Internet is a dedicated connection--a WAN adapter that is installed in the computer. The WAN adapter is

typically a DDS, T1, Fractional T1, or Frame Relay adapter. You must contract with a local telephone company to run the appropriate physical wiring to your premises. You need to verify that the WAN adapter is compatible with the Windows Server 2003 operating systems. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

The WAN adapter includes drivers that are installed in Windows Server 2003 operating systems so that the adapter appears as a network adapter.

You need to configure the following TCP/IP settings on the WAN adapter:

- IP address and subnet mask assigned from the InterNIC or an Internet service provider (ISP).
- Default gateway of the ISP router.

**Configuring the connection to the branch office network**

The connection to the branch office network is a LAN adapter that is installed in the computer. You need to verify that the LAN adapter is compatible with the Windows Server 2003 operating systems. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.
- DNS and WINS name servers of branch office name servers.

Because the branch office router will act as a router between the corporate office and the branch office, you must configure the branch office router with either static routes or with routing protocols so that all of the destinations on the branch office network are reachable from the branch office router.

**Installing a computer certificate**

Because the corporate office router is using computer certificates to authenticate the IPSec security association for the L2TP/IPSec connection, you must install a computer certificate on the branch office router in order for an L2TP/IPSec connection to be successfully established.

For more information about installing a computer certificate on the branch office router, see Computer certificates for L2TP/IPSec VPN connections and Network access authentication and certificates.

**Configuring a demand-dial interface**

You can create a demand-dial interface by using the Demand-Dial Interface Wizard. In the wizard, configure the following:

- **Interface Name**

  Type the name of the interface that represents the connection to the corporate office. For example, type **CorpOffice**.

- **Connection Type**

  Click **Connect using virtual private networking (VPN)**.

- **VPN Type**

  Click **Layer 2 Tunneling Protocol (L2TP)**.

- **Destination Address**

  Type the IP address or host name that is assigned to the Internet interface of the router at the corporate office. If you enter a host name, verify that the host name resolves to the proper IP address.

- **Protocols and Security**

  Select the protocols you want to route.

- **Dial-out Credentials**

  Type the name, domain name, and password of the user account that corresponds to this branch office router. The credentials are the same as those entered in the **Dial-Out Credentials** page of the Demand-Dial Interface Wizard when the demand-dial interface for this branch office was created on the corporate router.

**Configuring static routes**
You need to add static routes so that traffic to the corporate office is forwarded by using the appropriate demand-dial interface. For each route of the corporate office, configure the interface, destination, network mask, and metric. For the interface, select the demand-dial interface that corresponds to the corporate office previously created.

For example, the route that corresponds to the corporate office is 10.0.00 with a subnet mask of 255.0.0.0. This route becomes a static route with the following configuration:

- Interface: CorpOffice
- Destination: 10.0.0.0
- Network mask: 255.0.0.0
- Metric: 1

### Note

- Because the L2TP connection is a point-to-point connection, the Gateway IP address is not configurable.

For more information, see Add a static route.

**Configuring firewall packet filters**
If you are using a firewall in the branch office, you need to configure L2TP/IPSec packet filters on your firewall to allow L2TP/IPSec traffic between the corporate office router and the branch office router. For more information, see VPN servers and firewall configuration.

**Initiating the L2TP router-to-router VPN connection**
To connect the branch office router to the corporate router, in Routing and Remote Access, right-click the demand-dial interface that connects to the corporate office, and then click **Connect**.

For information about troubleshooting a router-to-router VPN, see Troubleshooting router-to-router VPNs.

# Routing Scenarios

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing scenarios**
This section describes typical network configurations that use Routing and Remote Access. Each scenario reviews the configuration of network media, addresses, routing protocols, and other services. While your network configuration may be different than those described here, the basic concepts apply.

The network addresses in these scenarios use private address ranges as specified by RFC 1597, "Address Allocation for Private Internets." If your network is connected to the Internet, you can contact an Internet service provider (ISP) to receive network addresses acquired from the Internet Assigned Numbers Authority (IANA).

**Notes**

- A computer running a member of the Windows Server 2003 family and the Routing and Remote Access service that provides LAN and WAN routing services is hereafter referred to as a server running Routing and Remote Access, or in specific contexts, as the Windows Server 2003 family routing service.

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

The following routing scenarios are described:

- SOHO Network to the Internet
- Small Office Network
- Medium-Size Office Network
- Corporate Network
- Dial-Up Branch Office Network
- Branch Office over the Internet

# SOHO Network to the Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SOHO network to the Internet**

To connect a small office or home office (SOHO) network to the Internet, you can use one of two methods:

1. Routed connection

   For a routed connection, the server running Routing and Remote Access acts as an IP router that forwards packets between SOHO hosts and Internet hosts. For more information, see Routed Connection to the Internet.

2. Translated connection

   For a translated connection, server running Routing and Remote Access and NAT acts as a network address translator, an IP router that translates addresses for packets that are forwarded between SOHO hosts and Internet hosts. For more information, see Translated Connection to the Internet.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Routed Connection to the Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routed connection to the Internet**
This scenario describes a small office or home office (SOHO) network that connects to the Internet by using a routed connection.

A SOHO network has the following characteristics:

- One network segment.
- A single protocol: TCP/IP.
- Demand-dial or dedicated-link connections to the Internet service provider (ISP).

The following illustration shows an example of a SOHO network.




The server running Routing and Remote Access is configured with a network adapter for the media that is used in the home network (for example, Ethernet) and an ISDN adapter or an analog modem. You can use a leased line or other permanent connection technologies, such as xDSL and cable modems, but this scenario describes the more typical configuration that uses a dial-up link to a local ISP.

This section covers:

- Planning for a routed connection
- Configuring a routed connection
- Testing a routed connection

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Planning for a routed connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for a routed connection**

While offering maximum flexibility in the types of traffic that can be sent and received, a routed connection requires more configuration. The server running Routing and Remote Access must be manually configured with an IP address configuration and, if DHCP is not being used, all the SOHO hosts must also be manually configured with an IP address configuration. If DHCP is being used, a DHCP server on the SOHO network must be configured with the proper scope and scope options. For more information about DHCP, see DHCP

For this scenario, an Internet service provider (ISP) supplies a range of IP addresses to use for the SOHO hosts and the IP address of a DNS server for host name resolution.

Routing protocols are not needed to propagate IP routing information on the SOHO network. The computers on the SOHO network are either manually configured with the default gateway IP address of the server running Routing and Remote Access or receive a default gateway IP address through DHCP. The server running Routing and Remote Access is not configured with a default gateway but with a default route. For an example of adding the default route on the server running Routing and Remote Access, see Add a default static IP route.

If the ISP is on the MBone, the multicast backbone of the Internet, it can forward IP multicast traffic from the Internet. To receive the multicast traffic, the interfaces on the server running Routing and Remote Access are configured for the IGMP router mode and IGMP proxy mode. For more information, see Configuring a routed connection.

A routed connection to the Internet means that communication can occur with any host on the Internet. It also means that the SOHO hosts are exposed to attackers on the Internet. To provide security, packet filters are configured on the server running Routing and Remote Access to prevent unwanted Internet traffic on the SOHO network. For more information, see Packet filtering.

Computers on the home network are configured to use any services provided by an ISP, such as Network News Transfer Protocol (NNTP) servers and mail servers.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Configuring a routed connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring a routed connection**
To configure a routed connection to the Internet for a small office or home office (SOHO) network, the following are configured:

- The server running Routing and Remote Access.
- Other computers on the SOHO network.

**Configuring the server running Routing and Remote Access to act as a router**
To configure the server running Routing and Remote Access to act as a router, the following steps are completed:

1. The TCP/IP protocol on the server running Routing and Remote Access for the SOHO network interface is configured with:

   - IP address (from the address range obtained from the ISP).
   - Subnet mask (from the address range obtained from the ISP).
   - DNS server (from the IP address received from the ISP).

   TCP/IP is configured through the properties of the TCP/IP protocol for the local area connection in Network Connections.

   **Note**

   - Do not configure a default gateway.

2. The Routing and Remote Access service is installed and enabled.

   For information on installing and enabling the Routing and Remote Access service, see Enable the Routing and Remote Access service.

3. Routing on the dial-up port is enabled.

   If the connection to the Internet is a permanent connection that appears as a LAN interface (such as DDS, T-Carrier, Frame Relay, permanent ISDN, xDSL, or cable modem), or if you are connecting your server running Routing and Remote Access to another router before the connection to the Internet, skip to step 5.

   For information about enabling routing on the dial-up port, see Enable routing on ports.

4. A demand-dial interface is created to connect to the ISP.

   A demand-dial interface is created that is enabled for IP routing and uses the dial-up equipment and credentials that are used to dial the ISP. For more information about creating demand-dial interfaces, see Add a demand-dial interface.

5. A default static route is created that uses the Internet interface.

   For a default static route, the demand-dial interface (for dial-up connections) or LAN interface (for permanent or intermediate router connections) that is used to connect to the Internet is selected. The destination is 0.0.0.0 and the network mask is 0.0.0.0. For a demand-dial interface, the gateway IP address is not configurable. For a LAN interface that

is a point-to-point connection to your ISP, the gateway address is 0.0.0.0.

For more information about configuring a default static route, see Add a default static IP route.

6. Multicast support is configured (optional).

   To add multicast support to the SOHO network:

   - The IGMP routing protocol is added. For more information, see Add the IGMP routing protocol.
   - IGMP router mode is enabled on the interface that is connected to the home network. For more information, see Enable IGMP router and IGMP proxy mode.
   - IGMP proxy mode is enabled on the interface that is connected to the ISP. For more information, see Enable IGMP router and IGMP proxy mode.

**Configuring other computers on the SOHO network**

The TCP/IP protocol on the SOHO hosts is configured with:

- IP address (from the address range obtained from the ISP).
- Subnet mask (from the address range obtained from the ISP).
- Default gateway (the IP address assigned to the SOHO server running Routing and Remote Access network adapter).
- DNS server (from the IP address received from the ISP).

TCP/IP is configured through the properties of the TCP/IP protocol for the local area connection in Network Connections.

### Notes

- The previous configuration of SOHO hosts assumes that TCP/IP is configured manually. To automatically configure TCP/IP for SOHO hosts, you must install and configure a DHCP server. For more information, see DHCP Overview.
- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Testing a routed connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Testing a routed connection**

From a computer in the home network, the **ping** command is used to ping a computer on the Internet and activate the demand-dial connection to the ISP. Once connected, the **ping** command is used again. If a reply is received, the packets are being routed correctly. For more information, see Using the ping command.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Translated Connection to the Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Translated connection to the Internet**
This scenario describes a small office or home office (SOHO) network that connects to the Internet by using a translated connection.

The network configuration is simplified through the use of the network address translation (NAT) routing protocol, which provides network address translation, addressing, and name resolution services for SOHO network computers.

A SOHO network has the following characteristics:

- One network segment.
- A single protocol: TCP/IP.
- Demand-dial or dedicated-link connections to the Internet service provider (ISP).

The following illustration shows an example of a SOHO network.




The server running Routing and Remote Access is configured with a network adapter for the media that is used in the home network (for example, Ethernet) and an ISDN adapter or an analog modem. You can use a leased line or other permanent connection technologies, such as xDSL and cable modems, but this scenario describes the more typical configuration that uses a dial-up link to a local ISP.

This section covers:

- Planning for a translated connection
- Configuring a translated connection
- Testing a translated connection

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Planning for a translated connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for a translated connection**

For this scenario, the network address translation addressing component on the network address translation computer automatically assigns unique addresses from the Internet Assigned Numbers Authority (IANA) private network IP of 192.168.0.0 with a subnet mask of 255.255.255.0 to all the other computers on the home network.

Traditional routing protocols such as RIP and OSPF are not needed to propagate IP routing information on the home network. The computers on the home network receive the default gateway IP address of the network address translation computer. The network address translation computer must be configured with a default route. For an example of adding the default route on the server running Routing and Remote Access, see Add a default static IP route.

If the Internet service provider (ISP) is on the MBone, the multicast backbone of the Internet, it can forward IP multicast traffic from the Internet. To receive the multicast traffic, the interfaces on the server running Routing and Remote Access are configured for IGMP router mode or IGMP proxy mode. For more information, see Configuring a translated connection.

Computers on the home network are configured to use any services provided by an ISP, such as Network News Transfer Protocol (NNTP) servers and mail servers.

**Notes**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Configuring a translated connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring a translated connection**
To configure the server running Routing and Remote Access to connect a small office or home office (SOHO) network to the Internet with a translated connection by using the network address translation (NAT) routing protocol and a dial-up connection to an Internet service provider (ISP), the steps in Deploying network address translation are completed.

**Notes**

- This scenario assumes the use of the NAT routing protocol for the translated connection. You can simplify the configuration of a translated connection by using the Internet Connection Sharing feature of Network Connections. For more information about the differences between Internet Connection Sharing and NAT, see Internet Connection Sharing and network address translation. For information about Internet Connection Sharing, see Connecting to the Internet in a home or small office network.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

For more information about setting up network address translation, see Setting Up Network Address Translation.

**Configuring multicast support (optional)**
To add multicast support to the home network:

- The IGMP routing protocol is added. For more information, see Add the IGMP routing protocol.

- IGMP router mode is enabled on the interface that is connected to the home network. For more information, see Enable IGMP router and IGMP proxy mode.

- IGMP proxy mode is enabled on the interface that is connected to the ISP. For more information, see Enable IGMP router and IGMP proxy mode.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Testing a translated connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Testing a translated connection**

From a computer in the small office or home office network, the **ping** command is used to ping a computer on the Internet and activate the demand-dial connection to the ISP. Once connected, the **ping** command is used again. If a reply is received, the packets are being routed correctly. For more information, see Using the ping command.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Small Office Network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Small office network**
A small office network has the following characteristics:

- Few network segments (for example, one segment on each floor or wing of a building).
- A closed network with no connections to or from another network such as the Internet.
- Support for the IP protocol.

The following illustration shows an example of a small office network.




In this small office network scenario, Routers 1 and 2 are configured with a network adapter for each medium that is used on Networks A, B, and C.

This section covers:

- Planning addresses for the small office network
- Routing protocols for the small office network
- Other services for the small office network
- Configuring the small office network
- Testing the small office network

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Planning addresses for the small office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning addresses for the small office network**
The following sections describe how IP addresses are assigned for this small office network scenario.

**IP addressing**
IP addresses are assigned based on the private network ID of 192.168.0.0 for each network segment that uses a subnet mask of 255.255.255.0. This provides for growth of up to 254 computers on each network segment.

The following table shows the assignment of IP addresses for this scenario.

**IP addresses for the small office network**

| Segment | IP address with a subnet mask | Range of host IDs |
|---|---|---|
| Network A | 192.168.1.0, 255.255.255.0 | 192.168.1.1-192.168.1.254 |
| Network B | 192.168.2.0, 255.255.255.0 | 192.168.2.1-192.168.2.254 |
| Network C | 192.168.3.0, 255.255.255.0 | 192.168.3.1-192.168.3.254 |

After planning the network for this small office, addresses are assigned (either manually or by using DHCP) in the ranges described in the preceding table for all other computers on Networks A, B, and C.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Routing protocols for the small office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing protocols for the small office network**
For the small office network, IP routing is implemented through the Routing Information Protocol (RIP) or through static routes.

**IP routing protocols**
In this scenario, configuring RIP for IP is easier than configuring static routes. For more information, see Configuring the small office network.

### Note

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Other services for the small office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Other services for the small office network**

For this scenario, the DHCP service is used to automatically configure IP addresses and other information on client computers. Because DHCP is used, the DHCP Relay Agent must be configured on the interface for Router 1 on Network A and on the interface for Router 2 on Network C. When the DHCP Relay Agent is used, DHCP client computers on Networks A and C can acquire addresses from the DHCP server on Network B. For information about setting up a DHCP server, see DHCP Overview.

### Notes

- You do not need to use the DHCP Relay Agent on a router if there is a DHCP server on each network segment.
- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Configuring the small office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the small office network**
To configure the small office network in this scenario, the following steps are completed:

1. Network adapters are installed and configured.
2. The Routing and Remote Access service is installed.
3. Routing protocols are added to each router.
4. The DHCP relay agent is installed and configured.
5. A WINS or DNS name server is installed.

These steps are outlined in the following sections and are intended as general guidelines for setting up a small office routed IP network.

**Installing and configuring network adapters**
To install and configure network adapters, the following steps are performed:

1. Two network adapters are installed in every router.
2. The drivers for the network adapters are installed.
3. The TCP/IP protocol is installed.
4. IP addresses are configured on the network adapters.

The following table shows the IP addresses for this small office network scenario.

| Router | Network adapter connected to | IP address |
|---|---|---|
| Router 1 | Network A | 192.168.1.1 |
| | Network B | 192.168.2.1 |
| Router 2 | Network B | 192.168.2.2 |
| | Network C | 192.168.3.1 |

**Installing the Routing and Remote Access service**
For this small office network scenario, the Routing and Remote Access service is installed. For more information about installing the Routing and Remote Access service, see Enable the Routing and Remote Access service.

**Configuring RIP for dynamic IP routing**
To implement dynamic IP routing for this small office network, the RIP routing protocol is added and then both LAN interfaces are added to RIP. For more information, see Add an IP routing protocol and Add a routing interface.

**Installing and configuring the DHCP Relay Agent**
In this small office network scenario, for Router 1, the DHCP Relay Agent is enabled for the interface that is on Network A. For Router 2, the DHCP Relay Agent is enabled for the interface that is on Network C. The DHCP server is also configured so that the clients on all the networks can use DHCP.

For information about configuring the DHCP Relay Agent, see Configure the DHCP Relay Agent.

**Installing a WINS or DNS name server**

To access network resources by using NetBIOS or domain names, a WINS or DNS name server is installed. For more information, see WINS Overview and DNS Overview.

### Note

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Testing the small office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Testing the small office network**

For this scenario, the ability of the small office network to route packets is tested by answering the following questions:

1. Are the packets being routed correctly?

   From a computer on each network, the **ping** command is used to ping a computer on the other networks. If a reply is not received from a network, the packets are not being routed correctly. For more information, see Using the ping command.

2. Are the RIP packets being sent correctly?

   To ensure that RIP packets are sent correctly, the neighboring RIP routers are viewed from one RIP router. Neighboring RIP routers are routers connected to a common network. For information about viewing RIP neighbors, see View RIP neighbors.

3. Are all the networks being routed to?

   To ensure that all the networks are being routed to, the IP routing table is viewed from a router. For more information, see Understanding the IP routing table.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Medium-Size Office Network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Medium-size office network**

A medium-size office network has the following characteristics:

- Several LAN segments with a backbone (for example, one segment on each floor or wing of a building).
- Dial-up connections for users who connect from home or while traveling.
- Internet connections.

The following illustration shows an example of a medium-size office network.




A medium-size office network typically uses a few different types of network media. The different office segments can use 10-megabit-per-second (Mbps) Ethernet or token ring networks, but the backbone network that is used to connect to the different networks and host servers can use 100-Mbps Ethernet, Fiber Distributed Data Interface (FDDI), or other types of networks.

The following table shows the network media that is used in this medium-size office network scenario.

| Router | Role | Medium |
|---|---|---|
| Router 1 | Connects to Network A | One Ethernet (10 Mbps) or token ring adapter |
| | Connects to the backbone | One Ethernet (100 Mbps) or FDDI adapter |
| Router 2 | Connects to Network B | One Ethernet (10 Mbps) or token ring adapter |
| | Connects to the backbone | One Ethernet (100 Mbps) or FDDI adapter |
| Router 3 | Connects to Network C | One Ethernet (10 Mbps) or token ring adapter |
| | Connects to the backbone | One Ethernet (100 Mbps) or FDDI adapter |
| Router 4 | Connects to the backbone | One Ethernet (100 Mbps) or FDDI adapter |
| | Connects to dial-up networking clients | Modems and/or ISDN adapters |

This section covers:

- Planning addresses for the medium-size office network

- Routing protocols for the medium-size office network
- Other services for the medium-size office network
- Configuring the medium-size office network
- Testing the medium-size office network

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Planning addresses for the medium-size office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning addresses for the medium-size office network**
The following sections describe how IP addresses are assigned for this medium-size office network scenario.

**IP addressing**
Network IDs are assigned based on the private network ID of 192.168.0.0 for each network segment that uses a subnet mask of 255.255.255.0. This provides for growth of up to 254 computers on each network segment.

The following table shows the assignment of IP addresses for this scenario.

**IP addresses for the medium-size office network**

| Segment | IP network ID with mask | Range of host IDs |
|---|---|---|
| Backbone | 192.168.1.0, 255.255.255.0 | 192.168.1.1-192.168.1.254 |
| Network A | 192.168.2.0, 255.255.255.0 | 192.168.2.1-192.168.2.254 |
| Network B | 192.168.3.0, 255.255.255.0 | 192.168.3.1-192.168.3.254 |
| Network C | 192.168.4.0, 255.255.255.0 | 192.168.4.1-192.168.4.254 |

After planning the network for this medium-size office, addresses are assigned (either manually or by using DHCP) in the ranges described in the preceding table for all other computers on Networks A, B, C, and D.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Routing protocols for the medium-size office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing protocols for the medium-size office network**
The IP routing protocol for medium-size networks is Routing Information Protocol (RIP) version 2. In this scenario, RIP v2 is configured on Routers 1, 2, 3, and 4.

### Note

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Other services for the medium-size office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Other services for the medium-size office network**

For this scenario, the DHCP service is used to automatically configure IP addresses and other information on client computers. Because DHCP is used, the DHCP Relay Agent must be configured on Routers 1, 2, and 3. When the DHCP Relay Agent is used, client computers on Networks A, B, and C can acquire addresses from the DHCP server on the backbone. For information about setting up a DHCP server, see DHCP Overview.

### Notes

- You do not need to use the DHCP Relay Agent on a router if you have a DHCP server on each network segment.
- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Configuring the medium-size office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the medium-size office network**

To configure the medium-size office network described in this scenario, the following steps are completed:

1. Network adapters are installed and configured.
2. The Routing and Remote Access service is installed.
3. RIP is configured.
4. Remote access devices are configured.
5. The DHCP Relay Agent is installed and configured.
6. A WINS or DNS name server is installed.

These steps are outlined in the following sections and are intended as general guidelines for setting up and testing a medium-size office routed IP network.

**Installing and configuring network adapters**

To install and configure network adapters, the following steps are performed:

1. Two network adapters are installed in every router.
2. The drivers are installed for the network adapters.
3. The TCP/IP protocol is installed.
4. IP addresses are configured on the network adapters.

The following table shows the IP addresses for this medium-size office network scenario.

| Router | Network adapter connected to | IP address |
|---|---|---|
| Router 1 | Backbone | 192.168.1.1 |
| | Network A | 192.168.2.1 |
| Router 2 | Backbone | 192.168.1.2 |
| | Network B | 192.168.3.1 |
| Router 3 | Backbone | 192.168.1.3 |
| | Network C | 192.168.4.1 |
| Router 4 | Backbone | 192.168.1.4 |

**Installing the Routing and Remote Access service**

For this medium-size office network scenario, the Routing and Remote Access service is installed and enabled as a LAN router on every router. For more information about installing the Routing and Remote Access service, see Enable the Routing and Remote Access service.

On Router 4, the Routing and Remote Access service is installed and enabled as a LAN router and as a remote access server.

An IP address pool is created with a starting address of 192.168.1.225 and an ending IP address of 192.168.1.254. This pool allows up to 29 simultaneously connected dial-up networking clients. IP addresses for interfaces on the backbone network must now be selected so that they are within the range 192.168.1.1 to 192.168.1.224.

For more information about creating IP address pools, see Create a static IP address pool.

**Configuring RIP**
On each router, the RIP protocol is configured. To configure RIP, the RIP routing protocol is added to IP, each interface is added to RIP, and then RIP version 2 is enabled on the RIP interface that is connected to the backbone. For more information, see Configure RIP version 2.

This scenario assumes that all the routers on the medium-size office network are servers running Routing and Remote Access that are configured for RIP version 2. If there are existing routers on the network that only support RIP version 1, then you must install and configure RIP version 1.

**Configuring remote access devices**
In this routing scenario, Router 4 is a remote access server for dial-up networking clients. On Router 4, remote access devices such as modems and ISDN adapters are installed. For more information, see Installing Dial-Up Equipment.

**Installing and configuring the DHCP Relay Agent**
To use DHCP on the network, the DHCP Relay Agent is installed on Routers 1, 2, and 3, the interfaces that are not on the backbone are configured to use the DHCP Relay Agent, and the DHCP server is configured for use by clients.

For information about configuring the DHCP Relay Agent, see Configure the DHCP Relay Agent.

**Installing a WINS or DNS name server**
To access network resources by using NetBIOS or domain names, a WINS or DNS name server is installed. For more information, see WINS Overview and DNS Overview.

## Note

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Testing the medium-size office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Testing the medium-size office network**

For this scenario, the ability of the medium-size office network to route packets is tested by answering the following questions:

1. Are the packets being routed correctly?

   From a computer on each network, the **ping** command is used to ping a computer on the other networks. If a reply is not received from a network, the packets are not being routed correctly. For more information, see Using the ping command.

2. Are the RIP packets being sent correctly?

   To ensure that RIP packets are sent correctly, the neighboring RIP routers are viewed from one RIP router. Neighboring RIP routers are routers connected to a common network. For information about viewing RIP neighbors, see View RIP neighbors.

3. Are all the networks being routed to?

   To ensure that all the networks are being routed to, the IP routing table is viewed from a router. For more information, see Understanding the IP routing table.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Corporate Network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Corporate network**

A typical corporate network has the following characteristics:

- Many LAN segments with a backbone (for example, one segment on each floor or wing of several buildings).
- More than one network protocol.
- Areas configured with Open Shortest Path First (OSPF).
- Dial-up connections for users who connect from home or while traveling.
- Leased-line connections to branch offices.
- Demand-dial connections to branch offices.
- Internet connections.

The following illustration shows an example of a corporate network.




A corporate network typically uses different types of network media. The different office segments can use 10-megabit-per-second (Mbps) Ethernet or token ring networks, but the backbone network that is used to connect to the different networks and host servers is usually made up of 100-Mbps Ethernet or Fiber Distributed Data Interface (FDDI). Connections to external networks (the Internet) are over leased lines or packet-switched services such as Frame Relay. Connections to branch offices are over either switched media (ISDN or analog modems), dedicated media (leased lines or Frame Relay), or the Internet.

You can also use a tunneling protocol to connect the branch offices to the corporate network over the Internet, as shown by Router 10 and Network G. For more information on this type of configuration, see Branch Office over the Internet.

This corporate network scenario depicts branch office connections that use switched, demand-dial links and dedicated links, as shown in the following table.

**Network media for the corporate network**

| Router | Role | Medium |
|---|---|---|
| Router 1 | Connects to the backbone | One Ethernet (100 Mbps) or FDDI adapter |
| | Connects to Network C | One Ethernet (10 Mbps) or token ring adapter |
| Router 2 | Connects to Network A | One Ethernet (10 Mbps) or token ring adapter |
| | Connects to Network C | One Ethernet (10 Mbps) or token ring adapter |
| Router 3 | Connects to Network B | One Ethernet (10 Mbps) or token ring adapter |
| | Connects Network C | One Ethernet (10 Mbps) or token ring adapter |
| Router 4 | Connects to the backbone | One Ethernet (100 Mbps) or FDDI adapter |
| | Connects to Network D | One Ethernet (10 Mbps) or token ring adapter |
| Router 5 | Connects to the backbone | One Ethernet (100 Mbps) or FDDI adapter |
| | Connects to Router 6 | One or more ISDN adapters and/or analog modems |
| | Connects to Router 7 | One Frame Relay adapter |
| Router 6 | Connects to Router 5 | One or more ISDN adapters and/or analog modems |
| | Connects to Network F | One Ethernet (10 Mbps) or token ring adapter |
| Router 7 | Connects to Router 5 | One Frame Relay adapter |
| | Connects to Network E | One Ethernet (10 Mbps) or token ring adapter |
| Router 8 | Connects to the backbone | One Ethernet (100 Mbps) or FDDI adapter |
| | Connects to dial-up networking clients | Analog modems and/or ISDN adapters |
| Router 9 | Connects to the backbone | One Ethernet (100 Mbps) or FDDI adapter |
| | Connects to the Internet | One leased line (56K/T1) or Frame Relay adapter |
| Router 10 | Connects to Network G | One Ethernet (100 Mbps) or FDDI adapter |
| | Connects to the Internet | One leased line (56K/T1) or Frame Relay adapter |

This section covers:

- Planning addresses for the corporate network
- Routing protocols for the corporate network
- Other services for the corporate network
- Configuring the corporate network
- Testing the corporate network

**Notes**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.
- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Planning addresses for the corporate network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning addresses for the corporate network**
The following sections describe how IP addresses are assigned for this corporate network scenario.

**IP addressing**
Network IDs are assigned based on the private network ID of 172.16.0.0 for each network segment that uses a subnet mask of 255.255.255.0. This provides for growth of up to 254 computers on each network segment.

The following table shows the assignment of IP addresses for this scenario.

**IP addresses for the corporate network**

| Segment | IP network ID with mask | Range of host IDs |
|---|---|---|
| Backbone | 172.16.1.0, 255.255.255.0 | 172.16.1.1-172.16.1.254 |
| Network A | 172.16.2.0, 255.255.255.0 | 172.16.2.1-172.16.2.254 |
| Network B | 172.16.3.0, 255.255.255.0 | 172.16.3.1-172.16.3.254 |
| Network C | 172.16.4.0, 255.255.255.0 | 172.16.4.1-172.16.4.254 |
| Network D | 172.16.5.0, 255.255.255.0 | 172.16.5.1-172.16.5.254 |
| Network E | 172.16.6.0, 255.255.255.0 | 172.16.6.1-172.16.6.254 |
| Network F | 172.16.128.0, 255.255.255.0 | 172.16.128.1-172.16.128.254 |
| Network G | 172.16.129.0, 255.255.255.0 | 172.16.129.1-172.16.129.254 |

After planning the network, addresses are assigned (either manually or by using DHCP) in the ranges described in the preceding table for all computers on the backbone and Networks A, B, C, D, E, F, and G.

If more than 254 addresses are needed on a segment, a new segment is configured rather than increasing the network address range. To provide for growth in the network, branch offices are assigned a 254-address range, even if there are not that many computers.

However, if a branch office has 10 to 15 people, the extra addresses in a 254-address range should not be wasted. Instead, variable-length subnets are used to divide the 254 addresses into multiple subnets of different lengths. Both the Routing Information Protocol (RIP) version 2 and the Open Shortest Path First (OSPF) protocols support variable-length subnets.

### Notes

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Routing protocols for the corporate network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing protocols for the corporate network**
The following section describes the use of OSPF as the IP routing protocol for this corporate network scenario.

**IP routing protocols**
Corporate networks typically use the OSPF protocol because it is efficient in large networks. For this corporate network scenario, OSPF is configured on all the servers running Routing and Remote Access within the corporate site. Different types of routes are used depending on the type of WAN connection. For a demand-dial link, you must use either static routes or auto-static routes that use RIP. For dedicated WAN links, you can use static routes, RIP, or OSPF.

For this scenario, Routers 5, 6, 7, 9, and 10 are also configured with either static routes or RIP version 2 on the interfaces that are not configured for OSPF. Static routes are preferable to using RIP version 2 when the link bandwidth is not adequate or when routing update traffic is considered an unnecessary overhead cost.

### Notes

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Other services for the corporate network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Other services for the corporate network**

For this scenario, the DHCP service is used for automatic configuration of IP addresses and other information on client computers. Because DHCP is used, the DHCP Relay Agent must be configured on the servers running Routing and Remote Access acting as Routers 1, 2, 3, 4, and 7. When the DHCP Relay Agent is used, client computers on Networks A, B, D, and G can acquire addresses from the DHCP server on the backbone. For information about setting up a DHCP server, see DHCP Overview.

**Notes**

- You do not need to use the DHCP Relay Agent on your router if you have a DHCP server on each network.
- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Configuring the corporate network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the corporate network**

To configure the corporate network described in this scenario, the following steps are completed at each router, as shown in the following table.

These steps are intended as general guidelines for setting up and testing a corporate routed IP network.

**Individual configurations for each router in the corporate network**

| Router | Configuration step |
|---|---|
| Router 1 | 1. Two network adapters are installed and then the drivers for the network adapters are installed.<br><br>2. The address 172.16.1.1 is assigned to the network adapter that is connected to the backbone, and 172.16.4.1 is assigned to the network adapter that is connected to Network C.<br><br>3. The Routing and Remote Access service is installed and LAN routing is enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access. For more information, see Enable the Routing and Remote Access service.<br><br>4. The OSPF routing protocol is configured. For more information, see Configure OSPF.<br><br>The area 0.0.0.1 is added to the router. By default, the backbone area 0.0.0.0 is added to the router.<br><br>On the interface connected to the backbone, the area is configured as 0.0.0.0.<br><br>On the interface connected to Network C, the area is configured as 0.0.0.1 for Area 1.<br><br>Passwords are configured for each interface.<br><br>5. The DHCP Relay Agent is installed and configured on the interface on Network C. For more information, see Configure the DHCP Relay Agent. |

| | |
|---|---|
| Router 2 | 1. Two network adapters are installed and then the drivers are installed for the network adapters.<br><br>2. The address 172.16.4.3 is assigned to the network adapter that is connected to Network C, and 172.16.2.1 is assigned to the network adapter that is connected to Network A.<br><br>3. The Routing and Remote Access service is installed and LAN routing is enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.<br><br>4. The OSPF routing protocol is configured.<br><br>The area 0.0.0.1 is added to the router. The backbone area 0.0.0.0 is deleted, which was added to the router by default.<br><br>On the interface connected to Network A, the area is configured as 0.0.0.1 for Area 1.<br><br>On the interface connected to Network C, the area is configured as 0.0.0.1 for Area 1.<br><br>Passwords are configured for each interface.<br><br>5. The DHCP Relay Agent is installed and configured on the interface connected to Network A. |
| Router 3 | 1. Two network adapters are installed and then the drivers for the network adapters are installed.<br><br>2. The address 172.16.4.2 is assigned to the network adapter that is connected to Network C, and 172.16.3.1 is assigned to the network adapter that is connected to Network B.<br><br>3. The Routing and Remote Access service is installed and LAN routing is enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.<br><br>4. The OSPF routing protocol is configured.<br><br>The area 0.0.0.1 is added to the router. The backbone area 0.0.0.0 is deleted, which was added to the router by default.<br><br>On the interface connected to Network B, the area is configured as 0.0.0.1 for Area 1.<br><br>On the interface connected to Network C, the area is configured as 0.0.0.1 for Area 1.<br><br>Passwords are configured for each interface.<br><br>5. The DHCP Relay Agent is installed and configured on the interface connected to Network B. |

| | |
|---|---|
| Router 4 | 1. Two network adapters are installed and then the drivers for the network adapters are installed.<br><br>2. The address 172.16.1.2 is assigned to the network adapter that is connected to the backbone, and 172.16.5.1 is assigned to the network adapter that is connected to Network D.<br><br>3. The Routing and Remote Access service is installed and LAN routing is enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.<br><br>4. The OSPF routing protocol is configured.<br><br>The area 0.0.0.2 is added to the router. By default, the backbone area 0.0.0.0 is added to the router.<br><br>On the interface connected to the backbone, the area is configured as 0.0.0.0.<br><br>On the interface connected to Network D, the area is configured as 0.0.0.2 for Area 2.<br><br>Passwords are configured for each interface.<br><br>5. The DHCP Relay Agent is installed and configured on the interface connected to Network D. |
| Router 5 | 1. One network adapter, one Frame Relay card, and a modem, ISDN, or remote access device are installed and then the drivers are installed.<br><br>2. The address 172.16.1.3 is assigned to the network adapter that is connected to the backbone, and the address 172.16.7.1 is assigned to the Frame Relay adapter.<br><br>3. The Routing and Remote Access service is installed and LAN and demand-dial routing are enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.<br><br>4. A pool of addresses is assigned to the router with a starting address of 172.16.1.193 and an ending IP address of 172.16.1.194. For more information, see Create a static IP address pool.<br><br>5. The OSPF routing protocol is configured.<br><br>The router is configured as an autonomous system boundary router.<br><br>On the interface connected to the backbone, the area is configured as 0.0.0.0.<br><br>6. RIP for IP is enabled on the Frame Relay interface that connects to Router 7.<br><br>7. Static routes are configured on the interface that connects to Router 6. For more information, see Add a static route. |

| | |
|---|---|
| Router 6 | 1. One network adapter and one modem, ISDN, or remote access device are installed and then the drivers for the adapters are installed.<br><br>For more information about the configuration of Router 5 and Router 6 to support a demand-dial connection, see Dial-Up Branch Office Network.<br><br>2. The address 172.16.128.1 is assigned to the network adapter that is connected to Network F.<br><br>3. The Routing and Remote Access service is installed and LAN and WAN routing are enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.<br><br>4. A pool of addresses is assigned to the router with a starting address of 172.16.128.253 and an ending IP address of 172.16.128.254. For more information, see Create a static IP address pool.<br><br>5. Static routes are configured on the interface that connects to Router 5. For more information, see Add a static route. |
| Router 7 | 1. One network adapter and one Frame Relay adapter are installed and then the drivers for the adapters are installed.<br><br>2. The address 172.16.7.2 is assigned to the Frame Relay adapter that is connected to Router 5, and 172.16.6.1 is assigned to the network adapter that is connected to Network G.<br><br>3. The Routing and Remote Access service is installed and LAN routing is enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.<br><br>4. Static routes, or any routing protocol that is used on the branch office LAN, are configured.<br><br>5. RIP for IP is configured on the interface that connects to Router 5.<br><br>6. The DHCP Relay Agent is installed and configured on the interface connected to Network E. |
| Router 8 | 1. One network adapter is installed and then the drivers for the network adapter are installed.<br><br>2. The address 172.16.1.4 is assigned to the network adapter that is connected to the backbone.<br><br>3. The Routing and Remote Access service is installed and LAN routing and remote access are enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.<br><br>4. Remote access devices are installed.<br><br>5. A pool of addresses is assigned with a starting address of 172.16.1.192 and an ending IP address of 172.16.1.254. This pool defines a range of 62 IP addresses for dial-up networking clients.<br><br>6. The OSPF routing protocol is configured.<br><br>On the interface that is connected to the backbone, the area is configured as 0.0.0.0.<br><br>Passwords are configured for each interface. |

| | |
|---|---|
| Router 9 | 1. Two network adapters are installed and then the drivers for the network adapters are installed.<br><br>2. The address 172.16.1.5 is assigned to the network adapter that is connected to the backbone. The IP address (as allocated by the Internet service provider) is assigned to the interface that is connected to the Internet.<br><br>3. The Routing and Remote Access service is installed and LAN and WAN routing and remote access are enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.<br><br>4. A pool of addresses is assigned with a starting address of 172.16.1.65 and an ending IP address of 172.16.1.66.<br><br>5. The OSPF routing protocol is configured.<br><br>On the interface that is connected to the backbone, the area is configured as 0.0.0.0.<br><br>Passwords are configured for each interface. |
| Router 10 | 1. One network adapter and one modem, ISDN, or remote access device are installed and then the drivers for the adapters are installed.<br><br>2. The address 172.16.129.1 is assigned to the network adapter that is connected to Network F.<br><br>3. The Routing and Remote Access service is installed and LAN and WAN routing are enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.<br><br>4. A pool of addresses is assigned to the router with a starting address of 172.16.129.253 and an ending IP address of 172.16.129.254.<br><br>5. Static routes are configured on the interface that connects to Router 9.<br><br>For more information about the configuration of Router 9 and Router 10 to support a demand-dial PPTP connection over the Internet, see Branch Office over the Internet. |

**Installing a WINS or DNS name server**
To access network resources by using NetBIOS or domain names, a WINS or DNS name server is installed. For more information, see WINS Overview and DNS Overview.

### Notes

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Testing the corporate network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Testing the corporate network**

For this scenario, the ability of the corporate network to route packets is tested by answering the following questions:

1. Are the packets being routed correctly?

   From a computer on each network, the **ping** command is used to ping a computer on the other networks. If a reply is not received from a network, the packets are not being routed correctly. For more information, see Using the ping command.

2. Is OSPF working correctly?

   The status of OSPF is checked by viewing OSPF neighbors. For more information, see View OSPF neighbors. This feature is not available on the Itanium-based versions of the Windows operating systems. This content is not available in this preliminary release.

   For more information about troubleshooting OSPF, see Troubleshooting OSPF.

3. Are all the networks being routed to?

   To ensure that all the networks are being routed to, the IP routing table is viewed from a router. For more information, see Understanding the IP routing table.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Dial-Up Branch Office Network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up branch office network**

The Corporate Network scenario depicted a network that included branch offices. This scenario describes the configuration of a network for a branch office that is connected through a dial-up link.

A typical branch office network has the following characteristics:

- One LAN segment.
- Demand-dial connections to the corporate office.

The following illustration shows an example of a dial-up branch office network.




In this scenario, the server running Routing and Remote Access must be configured with a network adapter for the media that is used in the branch office (for example, Ethernet) and an ISDN device or analog modem for connection to the corporate office.

For more information, see:

- Planning addresses for the branch office network
- Routing protocols for the branch office network
- Other services for the branch office network
- Configuring the branch office network
- Testing the branch office network

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Planning addresses for the branch office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning addresses for the branch office network**
For IP addressing, the branch office network segment is assigned a network range of 254 addresses to enable potential growth. In this scenario, the branch office network (Network F) uses the address of 172.16.128.0 for the network and 255.255.255.0 for the mask.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Routing protocols for the branch office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing protocols for the branch office network**
For this branch office scenario, static routes are used for IP routing.

**IP routing protocols**
It is not necessary to use IP routing protocols to propagate IP routing information to the branch office. Instead, static routes are configured on the server running Routing and Remote Access. For an example of adding a static route, see Configuring the branch office network.

If there are many static routes or static services to add, RIP is configured to use auto-static updates. Auto-static updates enable the router to automatically send a request for a list of routes from other routers. The list is then converted to static routes. For more information, see Perform Auto-Static Updates.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Other services for the branch office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Other services for the branch office network**

Using a long-haul analog phone line or an ISDN line unnecessarily can be expensive. To prevent a demand-dial connection from initiating except when information resources such as a Web site or file shares are accessed, DHCP and name servers are installed on the branch office network.

In this scenario, the DHCP service is used for automatic configuration of IP addresses and other information on client computers. Installing and configuring the DHCP Relay Agent on Router 6 may cause Router 6 to dial the corporate hub router every time a DHCP packet is sent by a computer in the branch office network. Therefore, a DHCP server is installed on the branch office network. For information about setting up a DHCP server, see DHCP Overview.

Name servers (WINS and DNS) are installed on the branch office network to enable local name resolution. The branch office name servers must be configured to replicate with the corporate name servers. For more information, see WINS Overview and DNS Overview.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Configuring the branch office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the branch office network**
To configure the dial-up branch office network described in this scenario, the following steps are completed:

1. A network adapter is installed and configured.
2. The Routing and Remote Access service is installed.
3. An IP address pool is configured.
4. Remote access devices are installed and configured.
5. Demand-dial interfaces are created.
6. Dial-in accounts are created.
7. IP static routes are added.
8. A DHCP, WINS, or DNS name server is installed.

These steps are outlined in the following sections and are intended as general guidelines for setting up and testing a branch office network.

**Installing and configuring the network adapter**
To configure Router 5 and Router 6, the following steps are performed:

1. The network adapter is installed.
2. The driver for the network adapter is installed.
3. The IP address on the network adapter is configured through the properties of the TCP/IP protocol.

In this branch office network scenario, the IP addresses are assigned as follows:

- For Router 5, the IP address 172.16.1.3 is assigned to the network adapter that connects to the backbone.
- For Router 6, the IP address 172.16.128.1 is assigned to the network adapter that connects to the branch office network.

**Installing the Routing and Remote Access service**
For this branch office network scenario, the Routing and Remote Access service is installed and LAN and demand-dial routing are enabled. For more information, see Enable the Routing and Remote Access service. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access.

**Configuring an IP address pool**
In this scenario, the branch office router (Router 6) is configured with an IP address pool with a single range, as shown in the following table.

| Starting address of the range | Ending address of the range | Number of addresses in the range |
|---|---|---|
| 172.16.128.253 | 172.16.128.254 | 2 |

The corporate router (Router 5) is also configured with an IP address pool with a single range, as shown in the following table.

| Starting address of the range | Ending address of the range | Number of addresses in the range |
|---|---|---|
| 172.16.1.133 | 172.16.1.134 | 2 |

For information about configuring an IP address pool, see Create a static IP address pool.

**Installing and configuring dial-up devices**

For both Router 5 and Router 6, modem or ISDN devices are installed and configured. For more information, see Installing Dial-Up Equipment.

**Creating demand-dial interfaces**

On Router 6, a demand-dial interface is created by using the Demand-Dial Interface wizard with the following configuration:

- Name: CorpOffice
- Protocols to route: IP
- Modem or adapter: Installed modem or ISDN device
- Phone number: Phone number of Router 5
- Dial-out credentials: BranchOffice

On Router 5, a demand-dial interface is created by using the Demand-Dial Interface wizard with the following configuration:

- Name: BranchOffice
- Protocols to route: IP
- Modem or adapter: Installed modem or ISDN device
- Phone Number: Phone number of Router 6
- Dial-out credentials: CorpOffice

For more information about the Demand-Dial Interface wizard, see Add a demand-dial interface.

**Creating dial-in accounts**

For this branch office network scenario, the dial-in accounts are created as follows:

- On Router 6, the account CorpOffice is added and dial-in permissions are enabled.
- On Router 5, the account BranchOffice is added and dial-in permissions are enabled.

For information about creating dial-in accounts, see Configure dial-in user properties.

**Adding IP static routes**

For this branch office scenario, on the branch office router (Router 6), a static default route is added. The demand-dial interface that connects to the corporate network is used.

The following route specifies that if a packet is sent to a computer that is not on the branch office network, a demand-dial connection is made with the corporate office. Because the connection between the branch office and the corporate office is a point-to-point connection, the gateway IP address is not configurable.

| Interface | Destination | Network mask | Metric |
|---|---|---|---|
| CorpOffice | 0.0.0.0 | 0.0.0.0 | 2 |

For more information about configuring a static default route, see Add a default static IP route.

On the corporate router (Router 5), a static route is added in order to connect to the branch office. The following route specifies that all packets destined for the branch office network are sent over the demand-dial interface.

| Interface | Destination | Network mask | Metric |
|---|---|---|---|
| BranchOffice | 172.16.128.0 | 255.255.255.0 | 2 |

For information about adding static routes, see Add a static route.

The corporate router (Router 5) is configured as an autonomous system boundary router so that the static routes are advertised to other OSPF routers.

**Installing a DHCP, WINS, or DNS name server**

A DHCP server is installed on the branch office network so that the clients can receive IP addresses. For more information about DHCP, see DHCP Overview. To access network resources by using NetBIOS or domain names, a WINS or DNS name server is installed. For more information, see WINS Overview and DNS Overview.

### Notes

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Testing the branch office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Testing the branch office network**

From a computer in the branch office, the **ping** command is used to ping a computer on the corporate network and activate the demand-dial connection with the corporate network router. Once connected, the **ping** command is used again. If a reply is received from that computer, the packets are being routed correctly.

For more information, see Using the ping command.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Branch Office over the Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Branch office over the Internet**
Branch offices can be connected through leased or dial-up lines to a corporate network. Using leased or dial-up lines across long distances can be expensive. Routing and Remote Access enables branch offices to connect to the corporate networks by using the Internet. In this scenario, a branch office has a dial-up connection to a local Internet service provider (ISP). The branch office router then makes a secure, encrypted tunnel by using the Point-to-Point Tunneling Protocol (PPTP) across the Internet to the corporate network. This network configuration can result in cost savings, because the dial-up line is local instead of long distance.

**Note**

- To establish the tunnel, the branch office router must know the IP address of the corporate network router.

The following illustration shows an example of a branch office that connects to the corporate network by using the Internet.




In this scenario, the server running Routing and Remote Access acting as a branch office router makes a demand-dial PPTP connection to the server running Routing and Remote Access on the corporate network.

The branch office router must be configured with a network adapter for the medium that is used in the branch office (for example, Ethernet) and an ISDN adapter or analog modem for connection to the ISP. You can use a leased line to connect to the ISP, but this scenario discusses only demand-dial connections. The corporate office router must be connected to the Internet by using a leased line.

This section covers:

- Planning addresses for the branch office over the Internet
- Routing protocols for the branch office over the Internet
- Other services for the branch office over the Internet
- Configuring the PPTP connection at the branch office
- Configuring the PPTP connection at the corporate office
- Testing the PPTP connection

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Planning addresses for the branch office over the Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning addresses for the branch office over the Internet**
To enable potential growth, a network range of 254 IP addresses is assigned to the branch office network segment. In this scenario, the branch office network segment (Network G) uses the address of 172.16.129.0 for the network and 255.255.255.0 for the mask.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Routing protocols for the branch office over the Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing protocols for the branch office over the Internet**
For this branch office scenario, because the connection between the branch office and the corporate office is not permanent, static routes are used for IP routing.

**IP routing protocols**
It is not necessary to use IP routing protocols to propagate IP routing information to the branch office. Instead, static routes are configured on the server running Routing and Remote Access. For an example of adding a static route, see Configuring the PPTP connection at the branch office.

If there are many static routes or static services to add, RIP is configured to use auto-static updates. Auto-static updates enable the router to automatically send a request for a list of routes from other routers. This list is then converted to static routes. Auto-static updates must be performed on both sides of the connection. For more information, see Perform Auto-Static Updates.

### Note

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Other services for the branch office over the Internet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Other services for the branch office over the Internet**
To maximize performance and minimize the amount of data that is sent over the connection, DHCP and name servers are installed on the branch office network. The branch office name servers are configured to replicate with the corporate name servers.

In this scenario, the DHCP service is used for automatic configuration of IP addresses and other information on client computers. A DHCP server is installed on the branch office network. For information about setting up a DHCP server, see DHCP Overview.

To access network resources by using NetBIOS or domain names, a WINS or DNS name server is installed. For more information, see WINS Overview and DNS Overview.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Configuring the PPTP connection at the branch office

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the PPTP connection at the branch office**
To configure a PPTP connection to a corporate network at a branch office router, you must create a demand-dial interface for the ISP connection, a demand-dial interface for the PPTP connection to the corporate router, and a static route for each interface. For this scenario, the following steps are completed:

1. Network adapters are installed and configured.
2. An IP address pool is configured.
3. The Routing and Remote Access service is installed.
4. Remote access devices are installed and configured.
5. Demand-dial interfaces are created.
6. IP static routes are added.
7. PPTP filters are added.
8. Remote access policies are configured.

These steps are outlined in the following sections and are intended as general guidelines for setting up a branch office PPTP connection to a corporate network.

**Installing and configuring network adapters**
To install and configure network adapters, the following steps are performed:

1. On the server running Routing and Remote Access, one network adapter is installed.
2. The driver for the network adapter is installed.
3. The IP address on the network adapter is configured through the properties of the TCP/IP protocol.

For this scenario, the address 172.16.129.1 is assigned to the network adapter that is connected to the branch office network.

**Installing the Routing and Remote Access service**
For this branch office network scenario, the Routing and Remote Access service is installed and LAN and demand-dial routing are enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access. For more information, see Enable the Routing and Remote Access service.

**Configuring an IP address pool**
In this scenario, the branch office router (Router 10) is configured with an IP address pool with a single range, as shown in the following table.

| Starting address of the range | Ending address of the range | Number of addresses in the range |
|---|---|---|
| 172.16.129.253 | 172.16.129.254 | 2 |

For information about configuring an IP address pool, see Create a static IP address pool.

**Creating demand-dial interfaces**
Two demand-dial interfaces are created: one for the modem or remote access device and one for the PPTP VPN device.

A demand-dial interface for the ISP connection is created by using the Demand-Dial Interface wizard with the following configuration:

- Name: ISP
- Protocols to route: IP
- Modem or adapter: Installed modem or ISDN device
- Phone number: Phone number of local ISP
- Dial-out credentials: ISP user account

A demand-dial interface for the PPTP connection is created by using the Demand-Dial Interface wizard with the following configuration:

- Name: CorpPPTP
- Protocols to route: IP
- Modem or adapter: VPN1 device
- Phone number: *w.x.y.z* (IP address of the Internet interface for Router 9)
- Dial-out credentials: BranchOfficePPTP

For more information about the Demand-Dial Interface wizard, see Add a demand-dial interface.

**Note**

- The use of *w.x.y.z* is intended to represent a valid public IP address as allocated by the Internet Assigned Numbers Authority (IANA) or an ISP. For your configuration, substitute your allocated public address for *w.x.y.z*.

Routing and Remote Access now displays three interfaces that correspond to:

1. The network adapter that connects to the branch office.
2. The ISP connection.
3. The PPTP connection.

**Adding IP static routes**

The PPTP connection to the corporate network is made by adding two IP static routes: a route to the ISP and a route to the PPTP server.

To call the PPTP server, a static route is added with the characteristics shown in the following table. This route specifies that if a packet's destination is on the corporate network, the PPTP demand-dial interface is used.

| Interface | Destination | Network mask | Metric |
|---|---|---|---|
| CorpPPTP | 172.16.0.0 | 255.255.0.0 | 1 |

To call the ISP, another static route is added with the characteristics shown in the following table. This route specifies that if a PPTP connection is made to the corporate PPTP server, the ISP demand-dial interface is used.

| Interface | Destination | Network mask | Metric |
|---|---|---|---|

| | | | |
|---|---|---|---|
| ISP | *w.x.y.z* | 255.255.255.255 | 1 |

For more information about adding static routes, see Add a static route.

**Note**

- The use of *w.x.y.z* is intended to represent a valid public IP address as allocated by the Internet Assigned Numbers Authority (IANA) or an ISP. For your configuration, substitute your allocated public address for *w.x.y.z*.

This configuration enables the router that is sending packets to 172.16.0.0 (the corporate network) to first connect to the ISP and then connect to the corporate office PPTP server (Router 9). After the interfaces are added and the static routes or routing protocols are set up, a manual connection to the corporate network is not needed except during the testing process.

**Adding PPTP filters**

To prevent the branch office router from sending or receiving any traffic except the packets sent over the PPTP connection, PPTP filters are added to the ISP demand-dial interface. For information about adding PPTP filters, see Add PPTP Filters.

**Configuring remote access policies**

A new remote access policy is created with the following properties:

- **Policy name** is set to **PPTP Access for Branch Office Routers** (example).
- The **NAS-Port-Type** condition is set to **Virtual (VPN)**.
- The **Tunnel-Type** condition is set to **Point-to-Point Tunneling Protocol**.
- The **Grant remote access permission** option is selected.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Configuring the PPTP connection at the corporate office

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the PPTP connection at the corporate office**
To route PPTP tunneled data to a branch office router, you must create a demand-dial interface for the PPTP connection to the branch office router and a static route. For this scenario, the following steps are completed:

1. Network adapters are installed and configured.
2. The Routing and Remote Access service is installed.
3. An IP address pool is configured.
4. A demand-dial interface is created.
5. A dial-in account is created.
6. An IP static route is added.
7. PPTP filters are added for security.
8. Remote access policies are configured.

These steps are outlined in the following sections and are intended as general guidelines for setting up a branch office PPTP connection to a corporate network.

**Installing and configuring network adapters**
To install and configure network adapters, the following steps are performed:

1. One network adapter is installed to connect to the backbone and one network adapter is installed to permanently connect to the Internet.
2. The drivers for the network adapters are installed.
3. The IP addresses on the network adapters are configured through the properties of the TCP/IP protocol.

In this network scenario, the IP addresses are configured as follows:

1. The IP address 172.16.1.5 is assigned to the network adapter that connects to the backbone.
2. The IP address *w.x.y.z* is assigned to the network adapter that connects to the Internet.

**Note**

- The use of *w.x.y.z* is intended to represent a valid public IP address as allocated by the Internet Assigned Numbers Authority (IANA) or an ISP. For your configuration, substitute your allocated public address for *w.x.y.z*.

**Installing the Routing and Remote Access service**
For this branch office network scenario, the Routing and Remote Access service is installed and LAN and WAN routing are enabled. To allow remote access clients to create remote access VPN connections to the corporate network, remote access is also enabled. The network adapters that are installed automatically appear as interfaces in Routing and Remote Access. For more information, see Enable the Routing and Remote Access service.

**Configuring an IP address pool**

In this scenario, the corporate router (Router 9) is configured with an IP address pool with a single range, as shown in the following table.

| Starting address of the range | Ending address of the range | Number of addresses in the range |
|---|---|---|
| 172.16.1.137 | 172.16.1.138 | 2 |

For information about configuring an IP address pool, see Create a static IP address pool.

**Creating a demand-dial interface**
A demand-dial interface is created by using the Demand-Dial Interface wizard with the following configuration:

- Name: BranchOfficePPTP
- Protocols to route: IP
- Modem or adapter: VPN1 device
- Phone number: <none>
- Dial-out credentials: <none>

For more information about the Demand-Dial Interface wizard, see Add a demand-dial interface.

Routing and Remote Access now displays three interfaces that correspond to:

1. The network adapter that is connected to the backbone.
2. The network adapter that is connected to the Internet.
3. The PPTP connection.

**Creating a dial-in account**
The account BranchOfficePPTP is added and dial-in permissions are enabled.

For more information, see Configure dial-in user properties.

**Adding an IP static route**
At the corporate router, a single IP static route is added so that traffic can be routed to the branch office network.

To route packets to the branch office, a static route is added with the characteristics shown in the following table. This route specifies that if a packet is to be delivered to the branch office network, the PPTP demand-dial interface is used.

| Interface | Destination | Network mask | Metric |
|---|---|---|---|
| BranchOfficePPTP | 172.16.129.0 | 255.255.255.0 | 1 |

For more information about adding static routes, see Add a static route.

**Adding PPTP filters**
In this scenario, Router 9 is also configured as a Microsoft Proxy Server. The Microsoft Proxy Server configuration settings are used to add PPTP filters to the Internet interface. For information about adding filters by using Microsoft Proxy Server, see the Microsoft Proxy Server documentation.

**Configuring remote access policies**
A new remote access policy is created with the following properties:

- **Policy name** is set to **PPTP Access for Corporate Office Router** (example).
- The **NAS-Port-Type** condition is set to **Virtual (VPN)**.

- The **Tunnel-Type** condition is set to **Point-to-Point Tunneling Protocol**.
- The **Grant remote access permission** option is selected.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Testing the PPTP connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Testing the PPTP connection**

From a computer on the branch office network, the **ping** command is used to ping a computer on the corporate network and activate the demand-dial PPTP connection. Once connected, the **ping** command is used again. If a reply is received from that computer, the packets are being routed correctly. For more information, see Using the ping command.

For this scenario, the status of both the ISP and PPTP connections is determined by checking each of the following conditions:

1. Is there a connection to the ISP?

   From the branch office server running Routing and Remote Access, the connection state of the ISP demand-dial interface is checked.

2. Is there a connection by PPTP to the corporate network?

   From the branch office server running Routing and Remote Access, a manual connection to the corporate office is made by using the PPTP demand-dial interface. Once connected, a computer is pinged on the corporate network.

For more information, see Troubleshooting demand-dial routing.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Routing Tools and Utilities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing tools and utilities**
This section covers:

- Administration tools in mixed environments
- Understanding the IP routing table
- Using the ping command
- Using the tracert command
- Using the pathping command
- Using the IP multicasting utilities
- Using event logging
- Using tracing
- The Netsh Command-Line Utility

# Administration tools in mixed environments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administration tools in mixed environments**
A mixed environment contains a combination of at least two of the following: servers running Windows NT 4.0 with Remote Access Service (RAS), servers running Windows NT 4.0 with Routing and Remote Access Service (RRAS), servers running Windows 2000 and Routing and Remote Access, and servers running members of the Windows Server 2003 family and Routing and Remote Access.

The platforms use different administration tools, as follows:

- With Windows 2000 and the Windows Server 2003 family, you use Routing and Remote Access or the Netsh command-line utility.
- With servers running Windows NT 4.0 with RRAS, you use the Routing and RAS Admin administrative tool.
- With servers running Windows NT 4.0 with RAS, you use the Remote Access Admin administrative tool.

The following sections describe the capabilities and behavior of each of these administrative tools when they are used to administer a platform for which they were not designed.

**Using Routing and Remote Access to administer a server running Windows NT 4.0 with RRAS**
You can use Routing and Remote Access to administer a server acting as a router and running Windows NT 4.0 with RRAS. However, due to the differences in configuration and features, you cannot do the following:

- Configure the remote access properties. On Windows NT 4.0, you must configure these properties on the server running Windows NT 4.0 with RRAS through Network in Control Panel.
- Install or remove the Routing and Remote Access service.
- Run Telnet.
- Configure remote access policies.
- Configure logging.
- Configure dial-in hours.
- Configure demand-dial filters.
- Install IGMP or network address translation (NAT) routing protocols.
- Obtain multicast statistics.
- Configure TCP/IP or IPX settings.

### Notes

- IPX routing is not supported on servers running the Windows Server 2003 family version of Routing and Remote Access.
- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

**Using Routing and Remote Access to administer a server running Windows NT 4.0 with RAS**

You cannot use Routing and Remote Access to manage a server running Windows NT 4.0 with RAS. When you select a server running Windows NT 4.0 with RAS, the Remote Access Admin administrative tool is started.

**Using Netsh to administer a server running Windows NT 4.0 with RAS or a server running Windows NT 4.0 with RRAS**
You cannot use the Netsh command-line utility to manage a server running Windows NT 4.0 with RAS or a server acting as a router and running Windows NT 4.0 with RRAS.

**Using Routing and RAS Admin to administer a server running Routing and Remote Access**
You cannot use Routing and RAS Admin to manage the routing components of Routing and Remote Access. When you select a server running Routing and Remote Access configured as a router, the router is displayed. However, there are no routing nodes and you cannot configure the routing components.

**Using Routing and RAS Admin to administer a server running Windows NT 4.0 with RAS**
You cannot use Routing and RAS Admin to manage a server running Windows NT 4.0 with RAS. When you select a server running Windows NT 4.0 with RAS, the Remote Access Admin administrative tool is started.

**Using Remote Access Admin to administer a server running Windows NT 4.0 with RRAS or a server running Routing and Remote Access**
You cannot use Remote Access Admin to manage a server acting as a router and running Windows NT 4.0 with RRAS or the routing components of Routing and Remote Access. When you select a server acting as a router and running Windows NT 4.0 with RRAS or a server running Routing and Remote Access, Remote Access Admin displays the "Remote Access Service is not started on the selected server" message.

# Understanding the IP routing table

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding the IP routing table**

To troubleshoot routing problems, you must understand the routing table. Every computer that runs TCP/IP makes routing decisions that are determined by the IP routing table.

For information about displaying the IP routing table, see View routing tables.

The following illustration shows an example of a routing table.

RRAS-ROUTER1 - IP Routing Table

| Destination | Network mask | Gateway | Interface | Metric | Protocol |
|---|---|---|---|---|---|
| 10.57.76.0 | 255.255.255.0 | 10.57.76.1 | Local Area C... | 1 | Local |
| 10.57.76.1 | 255.255.255.255 | 127.0.0.1 | Loopback | 1 | Local |
| 10.255.255.255 | 255.255.255.255 | 10.57.76.1 | Local Area C... | 1 | Local |
| 127.0.0.0 | 255.0.0.0 | 127.0.0.1 | Loopback | 1 | Local |
| 127.0.0.1 | 255.255.255.255 | 127.0.0.1 | Loopback | 1 | Local |
| 192.168.45.0 | 255.255.255.0 | 192.168.45.1 | Local Area C... | 1 | Local |
| 192.168.45.1 | 255.255.255.255 | 127.0.0.1 | Loopback | 1 | Local |
| 224.0.0.0 | 224.0.0.0 | 192.168.45.1 | Local Area C... | 1 | Local |
| 224.0.0.0 | 224.0.0.0 | 10.57.76.1 | Local Area C... | 1 | Local |
| 255.255.255.255 | 255.255.255.255 | 192.168.45.1 | Local Area C... | 1 | Local |
| 255.255.255.255 | 255.255.255.255 | 10.57.76.1 | Local Area C... | 1 | Local |

The IP routing table contains information in the following columns:

- **Destination**

  The destination is the destination host, subnet address, network address, or default route. The destination for a default route is 0.0.0.0.

- **Network mask**

  The network mask is used in conjunction with the destination to determine when a route is used. For example, a host route has a mask of 255.255.255.255, a default route has a mask of 0.0.0.0, and a subnet or network route has a mask between these two extremes.

  A mask of 255.255.255.255 means that only an exact match of the destination uses this route. A mask of 0.0.0.0 means that any destination can use this route. When a mask is written in binary, a **1** is significant (must match) and a **0** is insignificant (does not need to match).

  For example, a destination of 172.16.8.0 has a network mask of 255.255.248.0. This network mask means that the first two octets must match exactly, the first five bits of the third octet must match (248=11111000) and the last octet does not matter. The third octet of 172.16.8.0 (that is, 8) equals **00001**000 in binary. Without changing the first 5 bits (the masked-off portion shown in **bold**), you can go up to 15, or **00001**111 in binary. So a route with a destination of 172.16.8.0 and a mask of 255.255.248.0 applies to all packets destined for 172.16.**8**.0 through 172.16.**15**.255.

- **Gateway**

  The gateway is the IP address of the next router where a packet needs to be sent. On a LAN link (such as Ethernet or token ring), the gateway must be directly reachable by this router by using the interface indicated in the **Interface** column. On a LAN link, both the gateway and interface determine how the traffic is being forwarded by the router. For a demand-dial interface, the gateway address is not configurable. On a point-to-point link, the interface determines how the traffic is being forwarded by the router.

- **Interface**

  The interface indicates the LAN or demand-dial interface that is to be used to reach the next router.

- **Metric**

The metric indicates the relative cost of using the route to reach the destination. A typical metric is hops, or the number of routers to cross to reach the destination. If there are multiple routes with the same destination, the route with the lowest metric is the best route.

- **Protocol**

  The protocol shows how the route was learned. If the **Protocol** column lists RIP, OSPF, or anything other than Local, then the router is receiving routes. Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Using the ping command

Updated: August 22, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the ping command**

If you are having connectivity problems, you can use the **ping** command to check the destination IP address you want to reach and record the results. The **ping** command displays whether the destination responded and how long it took to receive a reply. If there is an error in the delivery to the destination, the **ping** command displays an error message.

You can use the **ping** command to:

- Ping your computer (by address, not host name) to determine that TCP/IP is functioning. (Pinging your computer does not verify that your network adapter is functioning.)
- Ping the local router to determine whether the router is running.
- Ping beyond your local router.

The following table shows some useful **ping** command options.

| Option | Use |
|---|---|
| **-n** *Count* | Determines the number of echo requests to send. The default is 4 requests. |
| **-w** *Timeout* | Enables you to adjust the time-out (in milliseconds). The default is 1,000 (a 1-second time-out). |
| **-l** *Size* | Enables you to adjust the size of the ping packet. The default size is 32 bytes. |
| **-f** | Sets the Do Not Fragment bit on the ping packet. By default, the ping packet allows fragmentation. |

The following example illustrates how to send two pings, each 1,450 bytes in size, to IP address 131.107.8.1:

```
C:\>ping -n 2 -l 1450 131.107.8.1
Pinging 131.107.8.1 with 1450 bytes of data:

Reply from 131.107.8.1: bytes=1450 time<10ms TTL=32
Reply from 131.107.8.1: bytes=1450 time<10ms TTL=32

Ping statistics for 131.107.8.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate roundtrip times in milliseconds:
    Minimum = 0ms, Maximum =  10ms, Average =  2ms
```

By default, **ping** waits 4,000 milliseconds (4 seconds) for each response to be returned before displaying the "Request Timed Out" message. If the remote system being pinged is across a high-delay link, such as a satellite link, responses may take longer to be returned. You can use the **-w** (wait) option to specify a longer time-out.

For more information about other **ping** options, see Command-line utilities.

To check connectivity by using the **ping** command, at the command prompt, type **ping** and the IP address you want to reach.

A response of "Destination net unreachable" means there was no route to the destination. You need to check the routing table on the router listed in the "Reply from" address in the "Destination net unreachable" message. For more information about the routing table, see Understanding the IP routing table.

A response of "Request timed out" means that there was no response to the ping in the default time period (1 second). You can check for the following:

- A router is down.

  To check the routers in the path between the source and the destination, use the **tracert** command. For more information, see Using the tracert command.

- The destination host is down.

  Physically verify that the host is running or check connectivity through another protocol.

- There is no route back to your computer.

  If the host is running, you can check for a return route by viewing the default gateway and local routing table on the destination host.

- The latency of the response is more than one second.

  Use the **-w** option on the **ping** command to increase the time-out. For example, to allow responses within 5 seconds, use **ping -w 5000**.

# Using the tracert command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the tracert command**

If you are having connectivity problems, you can use the **tracert** command to check the path to the destination IP address that you want to reach and record the results. The **tracert** command displays the series of IP routers that are used in delivering packets from your computer to the destination and how long it took on each hop. If the packets are unable to be delivered to the destination, the **tracert** command displays the last router that successfully forwarded your packets.

For more information about the **tracert** command, you can type **tracert -?** at a command prompt.

The most common use of **tracert** is as follows:

**tracert** *IP address* [**-d**]

This returns a list of the routers that are crossed to get to *IP address*. By using the **-d** option, the router path is displayed faster because **tracert** does not try to resolve the names of the routers in the path.

# Using the pathping command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the pathping command**

The **pathping** command is a route tracing tool that combines features of the **ping** and **tracert** commands with additional information that neither of those tools provides. The **pathping** command sends packets to each router on the way to a final destination over a period of time, and then computes results based on the packets returned from each hop. Since the command shows the degree of packet loss at any given router or link, it is easy to determine which routers or links might be causing network problems. A number of switches are available, as shown in the following table.

| Switch | Name | Function |
|---|---|---|
| **-n** | Hostnames | Does not resolve addresses to host names. |
| **-h** | Maximum hops | Maximum number of hops to search for target. |
| **-g** | Host-list | Loose source route along host list. |
| **-p** | Period | Number of milliseconds to wait between pings. |
| **-q** | Num_queries | Number of queries per hop. |
| **-w** | Timeout | Waits this many milliseconds for each reply. |
| **-T** | Layer Two tag | Attaches a Layer Two priority tag (for example, for IEEE 802.1p) to the packets and sends it to each of the network devices in the path. This helps in identifying the network devices that do not have Layer Two priority configured properly. The -T switch is used to test for Quality of Service (QoS) connectivity. |
| **-R** | RSVP test | Checks to determine whether each router in the path supports the Resource Reservation Protocol (RSVP), which allows the host computer to reserve a certain amount of bandwidth for a data stream. |

The default number of hops is 30, and the default wait time before a time-out is 3 seconds. The default period is 250 milliseconds, and the default number of queries to each router along the path is 100.

The following is a typical **pathping** report. The compiled statistics that follow the hop list indicate packet loss at each individual router.

```
D:\>pathping -n microsoft

Tracing route to microsoft [157.54.1.196]
over a maximum of 30 hops:
  0  172.16.87.35
  1  172.16.87.218
  2  192.168.52.1
  3  192.168.80.1
  4  157.54.247.14
  5  157.54.1.196

Computing statistics for 125 seconds...
            Source to Here   This Node/Link
Hop  RTT    Lost/Sent = Pct  Lost/Sent = Pct  Address
  0                                           172.16.87.35
                                 0/ 100 =  0%   |
  1   41ms     0/ 100 =  0%     0/ 100 =  0%  172.16.87.218
                                13/ 100 = 13%   |
  2   22ms    16/ 100 = 16%     3/ 100 =  3%  192.168.52.1
                                 0/ 100 =  0%   |
  3   24ms    13/ 100 = 13%     0/ 100 =  0%  192.168.80.1
                                 0/ 100 =  0%   |
  4   21ms    14/ 100 = 14%     1/ 100 =  1%  157.54.247.14
                                 0/ 100 =  0%   |
```

```
  5   24ms    13/ 100 = 13%     0/ 100 =  0%  157.54.1.196

Trace complete.
```

When **pathping** is run, you first see the results for the route as it is tested for problems. This is the same path that is shown by the **tracert** command. The **pathping** command then displays a busy message for the next 125 seconds (this time varies by the hop count). During this time, **pathping** gathers information from all the routers previously listed and from the links between them. At the end of this period, it displays the test results.

The two rightmost columns--**This Node/Link Lost/Sent=Pct** and **Address**--contain the most useful information. The link between 172.16.87.218 (hop 1), and 192.168.52.1 (hop 2) is dropping 13 percent of the packets. All other links are working normally. The routers at hops 2 and 4 also drop packets addressed to them (as shown in the **This Node/Link** column), but this loss does not affect their forwarding path.

The loss rates displayed for the links (marked as a **|** in the rightmost column) indicate losses of packets being forwarded along the path. This loss indicates link congestion. The loss rates displayed for routers (indicated by their IP addresses in the rightmost column) indicate that those routers' CPUs might be overloaded. These congested routers might also be a factor in end-to-end problems, especially if packets are forwarded by software routers.

# Using the IP multicasting utilities

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the IP multicasting utilities**

IP multicasting utilities in the Windows Server 2003 family consist of the following:

- The **mrinfo** command.
- Netsh support for multicast troubleshooting.
- Support for the **mtrace** command.

**The mrinfo command**

The Windows Server 2003 family includes the **mrinfo** command that displays the configuration of a multicast router. You can use the configuration information to aid in the troubleshooting of multicast forwarding and routing problems.

The **mrinfo** command queries a specified multicast router with an Internet Group Management Protocol (IGMP) message. The response to the query contains a version number, the list of interfaces and the neighbors on each interface, metrics, Time to Live (TTL) thresholds, and flags. The syntax of the **mrinfo** command is:

**mrinfo** [**-n**] [ **-i***address* ] [ **-r***retry_count* ] [ **-t***timeout_count* ] *multicast_router*

- The **-n** option displays IP addresses in numeric format.
- The **-i** option specifies the IP address of the interface from which you want to send the **mrinfo** query. By default, the interface from which to send the **mrinfo** query is determined by the IP routing table.
- The **-r** option specifies the neighbor query retry limit. The default value is 3.
- The **-t** option specifies how long in seconds **mrinfo** waits for a neighbor query reply. The default value is 4.

The following is an example of the **mrinfo** command:

```
C:\>mrinfo 10.1.0.1
10.1.0.1(test1.microsoft.com) [version 18.55,mtrace,snmp]:
10.1.0.1 -> 0.0.0.0 (local) [1/0/querier/leaf]
10.2.0.1 -> 10.2.0.2 (test2.microsoft.com) [1/0]
10.2.0.1 -> 10.2.0.3 (test3.microsoft.com) [1/0]
10.3.0.1 -> 0.0.0.0 (local) [1/0/querier/leaf]
```

In the above example, **mrinfo** is run against the multicast router at 10.1.0.1. The first line shows the multicast router configuration: version number (for servers running Routing and Remote Access, the version number reflects the build number of the operating system) and flags (mtrace and snmp supported).

Each additional line displays the interfaces on the multicast router and the neighbors on each interface. Interfaces 10.1.0.1 and 10.3.0.1 have no neighbors. Interface 10.2.0.1 has two neighbors, 10.2.0.2 and 10.2.0.3. For each line, **mrinfo** displays the interface and neighbor, the domain name for the neighbor, the multicast routing metric, the TTL threshold, and flags indicating its role on the network such as the IGMP querier of the network (querier) or whether it has no neighbors (leaf).

**Netsh support for multicast troubleshooting**

To view multicast tables and gather information to aid in the troubleshooting of multicast routing and forwarding problems, you can use the following **netsh** commands:

- **netsh routing ip show mfe**

  Displays the entries in the multicast forwarding table. This is equivalent to the Multicast Forwarding Table available in

Routing and Remote Access. To view the Multicast Forwarding Table from within Routing and Remote Access, under **IP Routing**, right-click **General**, and then click **Show Multicast forwarding table**.

- **netsh routing ip show mfestats**

  Displays packet statistics and input and output interface information for multicast forwarding entries in the multicast forwarding table. This is equivalent to the Multicast Statistics table available in Routing and Remote Access. To view the Multicast Statistics table from within Routing and Remote Access, under **IP Routing**, right-click **General**, and then click **Show Multicast statistics**.

- **netsh interface ip show joins**

  Displays the multicast groups locally joined on each interface.

**Support for the mtrace command**

Although the Windows Server 2003 family does not provide a version of the Mtrace multicasting utility, the server running Routing and Remote Access does respond to **mtrace** command queries from other Mtrace utilities.

# Using event logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using event logging**

Routing and Remote Access logs routing events extensively in the system event log. You can use information in the event logs to troubleshoot routing processes.

Four levels of logging are available:

- Log errors only
- Log errors and warnings (the default)
- Log the maximum amount of information
- Disable event logging

For example, if an OSPF router is unable to establish an adjacency on an interface, you can:

1. Disable OSPF on the interface.
2. Change the level of logging for OSPF to log the maximum amount of information.
3. Enable OSPF on the interface.
4. Examine the system event log for information on the OSPF adjacency process.
5. Change the level of logging for OSPF to log error information only.

You can further troubleshoot the adjacency problem by analyzing the OSPF entries in the system event log.

Logging is available in the following dialog boxes:

- **General** tab of the properties of IP Routing\General
- **General** tab of the properties of IP Routing\Network Address Translation (NAT)
- **General** tab of the properties of IP Routing\RIP
- **General** tab of the properties of IP Routing\OSPF
- **General** tab of the properties of IP Routing\IGMP

**Notes**

- Logging consumes system resources and should be used sparingly to help identify network problems. After the event has been logged or the problem is identified, you should immediately reset logging to its default value (log errors and warnings).
- When logging the maximum amount of information, the logging information can be complex and very detailed. Some of this information is useful only to Microsoft Product Support Services engineers or to network administrators who are very experienced with Routing and Remote Access.
- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the

Windows Server 2003 family.

# Using tracing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using tracing**

The server running Routing and Remote Access has an extensive tracing capability that you can use to troubleshoot complex network problems. You can enable the components in the Windows Server 2003 family to log tracing information to files. You must enable the tracing function by changing settings in the registry under:

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

You enable tracing for each routing protocol by setting the registry values described later in this section. You can enable and disable tracing for routing protocols while the router is running. Each routing protocol is capable of tracing and appears as a subkey (such as **OSPF** and **RIPV2**) under the preceding registry key.

To enable tracing for each routing protocol, you can configure the following registry entries for each protocol key:

- **EnableFileTracing REG_DWORD 1**

  You can enable logging tracing information to a file by setting **EnableFileTracing** to **1**. The default value is **0**.

- **FileDirectory REG_EXPAND_SZ** *Path*

  You can change the default location of the tracing files by setting **FileDirectory** to the path you want. The file name for the log file is the name of the component for which tracing is enabled. By default, log files are placed in the *systemroot*\Tracing folder.

- **FileTracingMask REG_DWORD** *LevelOfTracingInformationLogged*

  **FileTracingMask** determines how much tracing information is logged to the file. The default value is **FFFF0000**.

- **MaxFileSize REG_DWORD** *SizeOfLogFile*

  You can change the size of the log file by setting different values for **MaxFileSize**. The default value is **10000** (64K).

**Notes**

- Tracing consumes system resources and should be used sparingly to help identify network problems. After the trace is captured or the problem is identified, you should immediately disable tracing. Do not leave tracing enabled on multiprocessor computers.

- The tracing information can be complex and very detailed. Most of the time this information is useful only to Microsoft Product Support Services engineers or to network administrators who are very experienced with Routing and Remote Access.

# The Netsh Command-Line Utility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The Netsh command-line utility**
Netsh is a command-line and scripting utility for networking components for local or remote computers. The Netsh utility can also save a configuration script in a text file for archival purposes or for configuring other servers.

The Netsh utility is a shell that can support multiple components through the addition of Netsh helper DLLs. A Netsh helper DLL extends Netsh functionality by providing additional commands to monitor or configure a specific networking component. Each Netsh helper DLL provides a context, a group of commands for a specific networking component. Within each context, subcontexts can exist. For example, within the routing context, the subcontext **ip** exists to group IP routing commands together.

Netsh command-line options include the following:

- **-a** *AliasFile*

  Specifies that an alias file is used. An alias file contains a list of **netsh** commands and an aliased version so that you can use the aliased command line in place of the **netsh** command. You can use alias files to map commands that may be more familiar in other platforms to the appropriate **netsh** command.

- **-c** *Context*

  Specifies the context of the command that corresponds to an installed helper DLL.

- *Command*

  Specifies the **netsh** command to execute.

- **-f** *ScriptFile*

  Specifies that all of the **netsh** commands in the *ScriptFile* file are run.

- **-r** *RemoteMachine*

  Specifies that the **netsh** commands are run on a remote computer specified by either its name or IP address.

You can abbreviate commands to the shortest unambiguous string. For example, issuing the command **sh ip int** is equivalent to issuing **show ip interface**. **Netsh** commands can be either global or context-specific. Global commands can be issued in any context and are used for general Netsh utility functions. Context-specific commands vary according to the context. You can log commands issued to a log file to create an audit trail of a **netsh** command session.

The following table lists the **netsh** global commands.

| Command | Description |
|---|---|
| **..** | Moves up one context level. |
| **?** or **help** | Displays command-line Help. |
| **show version** | Displays the current version of Windows and the Netsh utility. |
| **show netdlls** | Displays the current version of installed Netsh helper DLLs. |
| **add helper** | Add a Netsh helper DLL. |
| **delete helper** | Removes a Netsh helper DLL. |

| | |
|---|---|
| **show helper** | Displays the installed Netsh helper DLLs. |
| **cmd** | Creates a command window. |
| **online** | Sets the current mode to online. |
| **offline** | Sets the current mode to offline. |
| **set mode** | Sets the current mode to online or offline. |
| **show mode** | Displays the current mode. |
| **flush** | Discards any changes in offline mode. |
| **commit** | Commits changes made in offline mode. |
| **set audit-logging** | Turns on or off the logging facility. |
| **show audit-logging** | Displays current audit logging settings. |
| **set loglevel** | Sets level of logging information. |
| **show loglevel** | Displays the level of logging information. |
| **set machine** | Configures the computer on which the **netsh** commands are executed. |
| **show machine** | Displays the computer on which the **netsh** commands are executed. |
| **exec** | Executes a script file containing **netsh** commands. |
| **quit** or **bye** or **exit** | Exits the Netsh utility. |
| **add alias** | Adds an alias to an existing command. |
| **delete alias** | Deletes an alias to an existing command. |
| **show alias** | Displays all defined aliases. |
| **dump** | Writes configuration to a text file. |
| **popd** | A scripting command that pops a context from the stack. |
| **pushd** | A scripting command that pushes the current context on the stack. |

The Netsh utility has the following command modes:

- Online

  In online mode, commands issued at a Netsh command prompt are executed immediately.

- Offline

  In offline mode, commands issued at a Netsh command prompt are accumulated and executed as a batch by issuing the **commit** global command. You can discard accumulated commands by issuing the **flush** global command.

- Script

  With either the **-f** command-line option or by issuing the **exec** global command at a Netsh command prompt, all the **netsh** commands in the specified file are executed.

To create a script of the current configuration, use the global **dump** command. The **dump** command outputs the current running configuration in terms of **netsh** commands. You can use the script created by this command to configure a new server or to reconfigure the existing server. If you are making extensive changes to the configuration of a component, it is

recommended that you begin the configuration session with the **dump** command, in case you need to restore the configuration prior to changes being made.

For more information on **netsh** commands for the Routing and Remote Access service, see:

- Interface commands
- IP routing commands
- Scheduling auto-static updates
- Netsh commands for remote access

# Interface commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interface commands**

The following table lists the **netsh** commands that you can type at a command prompt to administer interface settings on a server running Routing and Remote Access. If there are multiple commands for a particular function, they are indicated by separating the individual commands with a slash (/).

When you type the command at the command prompt, precede each command with **netsh**. For the exact syntax of each command, type the command with the **?** option. For example, to receive command-line Help on the **netsh interface** command, type **netsh interface ?** at the command prompt.

| Command | Description |
|---|---|
| **interface set/show interface** | Enables, disables, connects, disconnects, and displays the configuration of demand-dial interfaces. |
| **interface set/show credentials** | Configures or displays the user name, password, and domain name on a demand-dial interface. |

# IP routing commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP routing commands**

The following table lists the **netsh** commands that you can type at a command prompt to administer IP settings on a server running Routing and Remote Access. If there are multiple commands for a particular function, they are indicated by separating the individual commands with a slash (/). For example, the command **routing ip set/show loglevel** is actually two separate commands: **routing ip set loglevel** and **routing ip show loglevel**.

When you type the command at the command prompt, precede each command with **netsh**. For the exact syntax of each command, type the command with the **?** option. For example, to receive command-line Help on the **netsh routing ip set interface** command, type **netsh routing ip set interface ?** at the command prompt.

| Command | Description |
|---|---|
| **routing ip add/delete/set/show interface** | Adds, deletes, configures, or displays general IP routing settings on a specified interface. |
| **routing ip add/delete/set/show filter** | Adds, deletes, configures, or displays IP packet filters on a specified interface. |
| **routing ip add/delete/show boundary** | Adds, deletes, or displays multicast boundary settings on a specified interface. |
| **routing ip add/delete/set/show rtmroute** | Adds, deletes, configures, or displays a non-persistent Route Table Manager route. |
| **routing ip add/delete/set/show persistentroute** | Adds, deletes, configures, or displays persistent routes. |
| **routing ip add/delete/set/show preferenceforprotocol** | Adds, deletes, configures, or displays the preference level for a routing protocol. |
| **routing ip add/delete/set/show scope** | Adds, deletes, or displays a multicast scope. |
| **routing ip set/show loglevel** | Configures or displays the global IP logging level. |
| **routing ip show helper** | Displays all Netsh utility subcontexts of IP. |
| **routing ip show protocol** | Displays all running IP routing protocols. |
| **routing ip show mfe** | Displays multicast forwarding entries. |
| **routing ip show mfestats** | Displays multicast forwarding entry statistics. |
| **routing ip show boundarystats** | Displays IP multicast boundaries. |
| **routing ip show rtmdestinations** | Displays destinations in the Route Table Manager routing table. |
| **routing ip show rtmroutes** | Displays routes in the Route Table Manager routing table. |

| | |
|---|---|
| **routing ip nat set/show global** | Configures or displays global network address translation (NAT) settings. |
| **routing ip nat add/delete/set/show interface** | Adds, deletes, configures, or displays network address translation (NAT) settings for a specified interface. |
| **routing ip nat add/delete addressrange** | Adds or deletes an address range to the NAT interface public address pool. |
| **routing ip nat add/delete addressmapping** | Adds or deletes a NAT address mapping. |
| **routing ip nat add/delete portmapping** | Adds or deletes a NAT port mapping. |
| **routing ip autodhcp set/show global** | Configures or displays global DHCP allocator parameters. |
| **routing ip autodhcp set/show interface** | Configures or displays DHCP allocator settings for a specified interface. |
| **routing ip autodhcp add/delete exclusion** | Adds or deletes an exclusion from the DHCP allocator range of addresses. |
| **routing ip dnsproxy set/show global** | Configures or displays global DNS proxy parameters. |
| **routing ip dnsproxy set/show interface** | Configures or displays DNS proxy parameters for a specified interface. |
| **routing ip igmp set/show global** | Configures or displays IGMP global settings. |
| **routing ip igmp add/delete/set/show interface** | Adds, deletes, configures, or displays IGMP on the specified interface. |
| **routing ip igmp show grouptable** | Displays the IGMP host groups table. |
| **routing ip igmp show ifstats** | Displays the IGMP statistics for each interface. |
| **routing ip igmp show iftable** | Displays the IGMP host groups for each interface. |
| **routing ip igmp show proxygrouptable** | Displays the IGMP group table for the IGMP proxy interface. |
| **routing ip igmp show rasgrouptable** | Displays the group table for internal interface used by the remote access server. |
| **routing ip ospf set/show global** | Configures or displays global OSPF settings. This feature is not available on the Itanium-based versions of the Windows operating systems. This content is not available in this preliminary release. |
| **routing ip ospf add/delete/set/show interface** | Adds, removes, configures, or displays OSPF on a specified interface. |
| **routing ip ospf add/delete/set/show area** | Adds, removes, configures, or displays an OSPF area. |

| | |
|---|---|
| **routing ip ospf add/delete/set/show virtif** | Adds, removes, configures, or displays an OSPF virtual interface. |
| **routing ip ospf add/delete/show neighbor** | Adds, removes, configures, or displays an OSPF neighbor. |
| **routing ip ospf add/delete/show protofilter** | Adds, removes, configures, or displays routing information sources for OSPF external routes. |
| **routing ip ospf add/delete/show routefilter** | Adds, removes, configures, or displays route filtering for OSPF external routes. |
| **routing ip ospf show areastats** | Displays OSPF area statistics. |
| **routing ip ospf show lsdb** | Displays the OSPF link state database. |
| **routing ip ospf show virtifstats** | Displays OSPF virtual link statistics. |
| **routing ip relay set global** | Configures DHCP Relay Agent global settings. |
| **routing ip relay add/delete/set interface** | Adds, removes, or configures DHCP Relay Agent settings on a specified interface. |
| **routing ip relay add/delete dhcpserver** | Adds or removes a DHCP server IP address to the list of DHCP server addresses. |
| **routing ip relay show ifbinding** | Displays IP address bindings for interfaces. |
| **routing ip relay show ifconfig** | Displays DHCP Relay Agent configuration for each interface. |
| **routing ip relay show ifstats** | Displays DHCP statistics for each interface. |
| **routing ip rip set/show global** | Configures or displays global RIP-for-IP settings. |
| **routing ip rip add/delete/set/show interface** | Adds, deletes, configures, or displays RIP-for-IP settings on a specified interface. |
| **routing ip rip add/delete peerfilter** | Adds or removes a RIP peer filter. |
| **routing ip rip add/delete acceptfilter** | Adds or removes a RIP route filter to the list of routes being accepted. |
| **routing ip rip add/delete announcefilter** | Adds or removes a RIP route filter to the list of routes being announced. |
| **routing ip rip add/delete/show neighbor** | Adds, removes, or displays a RIP neighbor. |
| **routing ip rip set/show flags** | Configures and displays advanced settings of RIP for IP on a specified interface. |

| | |
|---|---|
| **routing ip rip show globalstats** | Displays global RIP parameters. |
| **routing ip rip show ifbinding** | Displays IP address bindings for interfaces. |
| **routing ip rip show ifstats** | Displays RIP statistics for each interface. |

# Scheduling auto-static updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scheduling auto-static updates**

You can automate scheduled updates by using a combination of Netsh utility scripts and Task Scheduler. To perform an automated auto-static update by using Routing Information Protocol (RIP) for IP, use the following **netsh** commands:

**netsh interface set interface name=** *DemandDialInterfaceName* **connect=CONNECTED**

**netsh routing ip rip update** *DemandDialInterfaceName*

**netsh interface set interface name=** *DemandDialInterfaceName* **connect=DISCONNECTED**

For example, to automatically update the RIP-for-IP routes by using a demand-dial connection called CorpHub, you type the following **netsh** commands:

**netsh interface set interface name=CorpHub connect=CONNECTED**

**netsh routing ip rip update CorpHub**

**netsh interface set interface name=CorpHub connect=DISCONNECTED**

You can run these commands from a batch file or you can place them in a Netsh script file. For example, the script file Corphub.scp runs the following commands for CorpHub:

**interface set interface name=CorpHub connect=CONNECTED**

**routing ip rip update CorpHub**

**interface set interface name=CorpHub connect=DISCONNECTED**

To run the Corphub.scp script, type the following at a command prompt:

**netsh -f corphub.scp**

After the batch file or Netsh script file is created, you can execute the batch file or Netsh script on a scheduled basis by using Task Scheduler.

# Netsh commands for remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for remote access**

The following table lists the Netsh commands that you can type at a command prompt to administer a server running Routing and Remote Access. If there are multiple commands for a particular function, they are indicated by separating the individual commands with a slash (/). For example, the command **ras add/delete/show registeredserver** is actually three separate commands: **ras add registeredserver**, **ras delete registeredserver**, and **ras show registeredserver**.

When you type the command at the command prompt, precede each command with **netsh**. For the exact syntax of each command, type the command with the **?** option. For example, to receive command-line Help on the **netsh ras add registeredserver** command, type **netsh ras add registeredserver ?** at the command prompt.

| Command | Description |
|---|---|
| **ras add/delete/show registeredserver** | Configures or displays whether the specified remote access server computer is a member of the **RAS and IAS Servers** security group in the Active Directory directory service of the specified domain. |
| **ras show activeservers** | Displays current servers running Routing and Remote Access on your network. |
| **ras set/show authmode** | Configures or displays whether and when dial-in connections are authenticated. |
| **ras add/delete/show authtype** | Configures or displays the permitted authentication types. |
| **ras add/delete/show client** | Configures or displays currently connected remote access clients. |
| **ras add/delete/show link** | Configures or displays the configuration of software compression and link control protocol (LCP) extensions. |
| **ras add/delete/show multilink** | Configures or displays Multilink and Bandwidth Allocation Protocol (BAP) settings. |
| **ras set/show tracing** | Configures or displays tracing settings. |
| **ras set/show user** | Configures or displays remote access settings for user accounts. |
| **ras ip set access** | Configures whether IP traffic from remote access clients is forwarded to the networks to which the remote access server is connected. |
| **ras ip set addrassign** | Configures the method by which the remote access server assigns IP addresses to incoming connections. |
| **ras ip set addrreq** | Configures whether remote access clients or demand-dial routers can request their own IP addresses. |
| **ras ip show config** | Displays IP remote access configuration. |
| **ras ip set negotiation** | Configures whether IP is negotiated for remote access connections. |
| **ras ip delete pool** | Deletes the static IP address pool. |
| **ras ip add/delete range** | Adds or removes a range of addresses from the static IP address pool. |
| **ras appletalk set access** | Configures whether AppleTalk traffic from remote access clients is forwarded to the networks to which the remote access server is connected. |

| | |
|---|---|
| **ras appletalk show config** | Displays AppleTalk remote access configuration. |
| **ras appletalk set negotiation** | Configures whether AppleTalk is negotiated for remote access connections. |
| **ras aaaa set/show accounting** | Configures or displays the accounting provider. |
| **ras aaaa add/delete/set/show acctserver** | Configures or displays RADIUS accounting servers. |
| **ras aaaa set/show authentication** | Configures or displays the authentication provider. |
| **ras aaaa add/delete/set/show authserver** | Configures or displays RADIUS authentication servers. |

For more information about Netsh, see The Netsh Command-Line Utility.

# Troubleshooting routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section covers:

- Common routing problems
- Troubleshooting RIP for IP
- Troubleshooting OSPF
- Troubleshooting network address translation
- Troubleshooting the DHCP Relay Agent
- Troubleshooting IP multicast
- Troubleshooting demand-dial routing

**Note**

- Open Shortest Path First (OSPF) is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Common routing problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Common routing problems**

What problem are you having?

**The server running Routing and Remote Access is not properly forwarding traffic.**

**Cause:** The Routing and Remote Access service is not running.

**Solution:** Verify that the Routing and Remote Access service is running.

**See also:** Determine the status of the Routing and Remote Access service

**Cause:** Routing is not enabled.

**Solution:** Verify that routing is enabled.

**See also:** Enable LAN and WAN routing

**Cause:** IP routing is not enabled.

**Solution:** Verify that IP routing is enabled on the **IP** tab for the properties of the server.

**See also:** View properties of the remote access server

**Cause:** The router is not receiving routes from routing protocols.

**Solution:** Verify whether the router is receiving routes from routing protocols.

For example, in the IP routing table, the **Protocol** column shows how the route was learned. If the **Protocol** column lists RIP or OSPF (or anything other than Local), then the router is receiving routes.

**See also:** Determine whether the router is receiving routes; Understanding the IP routing table

**Cause:** The proper routing protocols are not enabled on the proper interfaces.

**Solution:** Verify that the proper routing protocols are enabled on the proper interfaces.

In Routing and Remote Access, you can click **IP Routing** to see what routing protocols are installed. You can then click a routing protocol to see what interfaces are using that routing protocol.

For an interface running OSPF, you must verify that OSPF is enabled on that interface. This feature is not available on the Itanium-based versions of the Windows operating systems. This content is not available in this preliminary release.

**See also:** Verify that OSPF is enabled

**Cause:** Your TCP/IP configuration is incorrect.

**Solution:** Verify your TCP/IP configuration by using the **ipconfig** command.

Unlike hosts, routers are typically configured with a static TCP/IP configuration rather than by using DHCP. Verify that you are using the correct subnet mask for your environment. To verify that you are using the correct subnet mask, you can use the **ping** command to try to ping hosts on your network segment and your neighboring routers.

**See also:** Using the ping command

**Cause:** Your default route configuration is incorrect.

**Solution:** Verify your default route configuration.

Unlike hosts, routers are not normally configured with a default gateway. Instead, a default route is used to forward all traffic with an unknown destination. If default routing is to be used on the router, a default route is either learned through routing protocols or configured on the router. If a default route is not being learned through routing protocols (for example, when you connect a network to the Internet), you should configure a static default route on the router.

For a statically configured default route, you need to check the configuration to make sure the route is using the correct interface. If the interface being used is a LAN interface, such as Ethernet or token ring, the **Gateway** IP address for the route must be a directly reachable IP address on the same network as the selected interface.

**See also:** Add a default static IP route

To reset Routing and Remote Access back to default values, see Reset the Routing and Remote Access service.

# Troubleshooting RIP for IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting RIP for IP**
What problem are you having?

- There are improper routes in a mixed RIP v1 and RIP v2 environment.
- Silent RIP hosts are not receiving routes.
- RIP routers are not receiving expected routes.
- Auto-static RIP updates are not working.
- Host or default routes are not being propagated.

**There are improper routes in a mixed RIP v1 and RIP v2 environment.**
**Cause:** RIP v2 routers are configured to multicast announcements. Multicast announcements are never received by RIP v1 routers.

**Solution:** On networks containing RIP v1 routers, verify that RIP v2 is configured to broadcast its announcements on networks that contain RIP v1 routers and verify that the RIP v2 router interfaces are configured to accept both RIP v1 and RIP v2 announcements.

**See also:** Configure RIP version 2

**Silent RIP hosts are not receiving routes.**
**Cause:** RIP v2 routers are configured to multicast announcements. Multicast announcements are never received by silent RIP hosts.

**Solution:** If there are silent RIP hosts on a network that are not receiving routes from the local RIP router, verify the version of RIP supported by the silent RIP hosts. For example, if the silent RIP hosts only support listening for broadcasted, RIP v1 announcements, you cannot use RIP v2 multicasting.

If you are using the RIP listener component, available on Microsoft Windows NT Workstation version 4.0, Service Pack 4, or the Windows 2000 RIP Listening Service, you must configure your RIP routers for RIP v1 or RIP v2 broadcasting.

**See also:** Configure RIP version 2

**RIP routers are not receiving expected routes.**
**Cause:** You are deploying variable length subnetting, disjointed subnets, or supernetting in a RIP v1 or mixed RIP v1 and RIP v2 environment.

**Solution:** Do not deploy variable length subnetting, disjointed subnets, or supernetting in a RIP v1 or mixed RIP v1 and RIP v2 environment.

**See also:** "Unicast IP Routing" at the Microsoft Windows Resource Kits Web site.

**Cause:** Your password is not matched for all the RIP v2 interfaces on a network segment.

**Solution:** If authentication is enabled, verify that all interfaces on the same network are using the same case-sensitive password.

**See also:** Enable authentication

**Cause:** RIP peer filtering is not configured correctly.

**Solution:** If RIP peer filtering is being used, verify that the correct IP addresses for the neighboring peer RIP routers are configured.

**See also:** Add peer filters

**Cause:** RIP route filtering is not configured correctly.

**Solution:** If RIP route filtering is being used, verify that the ranges of network IDs for your internetwork are included or are not being excluded.

**See also:** Add route filters

**Cause:** RIP neighbors are not configured correctly.

**Solution:** If RIP neighbors are configured, verify that the correct IP addresses are configured for the unicasted RIP announcements.

**See also:** Add a unicast neighbor

**Cause:** IP packet filtering is preventing the receiving (through input filters) or sending (through output filters) of RIP traffic.

**Solution:** Verify that IP packet filtering on the router interfaces is not preventing the receiving (through input filters) or sending (through output filters) of RIP traffic. RIP traffic uses User Datagram Protocol (UDP) port 520.

**See also:** Manage Packet Filters

**Cause:** TCP/IP filtering is preventing the receiving of RIP traffic.

**Solution:** Verify that TCP/IP filtering on the router interfaces is not preventing the receiving of RIP traffic. RIP traffic uses UDP port 520.

**See also:** Configure TCP/IP to use TCP/IP filtering

**Cause:** You are using auto-static RIP and you did not do an initial manual update.

**Solution:** When you use auto-static RIP on a demand-dial interface, the first time you make a connection, you must manually update routes.

You must also update routes manually on the router for the corresponding interface. The routes then appear in the IP routing table.

**See also:** Perform Auto-Static Updates; View routing tables

**Auto-static RIP updates are not working.**
**Cause:** The demand-dial interfaces are configured to broadcast announcements.

**Solution:** For dial-up demand-dial interfaces that use auto-static updates, configure the demand-dial interfaces to use RIP v2 multicast as the outgoing packet protocol. When a router calls another router, each router receives an IP address from the IP address pool of the other router, which is on a different subnet. Because broadcasted RIP messages are addressed to the subnet broadcast address, each router does not process the other router's broadcasted request for routes. Using multicasting, RIP requests and announcements are processed regardless of the subnet for the router interfaces.

**See also:** Configure RIP version 2

**Host or default routes are not being propagated.**
**Cause:** RIP by default is not configured to propagate host or default routes.

**Solution:** If host routes or default routes need to be propagated, change the default settings on the **Advanced** tab of the properties of a RIP interface.

# Troubleshooting OSPF

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting OSPF**
What problem are you having?

- OSPF adjacency is not forming.
- A virtual link is not forming.
- Lack of OSPF routes or existence of improper OSPF routes.

**OSPF adjacency is not forming.**
**Cause:** OSPF is not enabled on the interface.

**Solution:** Verify that OSPF is enabled on the interface on the network segment where an adjacency should form. By default, when you add an interface to the OSPF routing protocol, OSPF is disabled for the interface and must be manually enabled.

**See also:** Verify that OSPF is enabled

**Cause:** Adjacency should not form.

**Solution:** Verify that the two neighboring routers should form an adjacency. If the two routers are the only routers on the network, an adjacency should form. If there are more than two routers on the network, adjacencies only form with the designated router (DR) and backup designated router (BDR). If the two routers have already formed adjacencies with the DR and the BDR, they cannot form adjacencies with each other. In this case, the neighbor should appear as **2-way** under **Neighbor State** (in the **Neighbors** dialog box from IP Routing\General).

**See also:** View routing tables

**Cause:** No IP connectivity to a neighboring router.

**Solution:** Ping the neighboring router to ensure basic IP and network connectivity. You can use the **tracert** command to trace the route to the neighboring router. There should not be any routers between the neighboring routers.

**See also:** Using the ping command; Using the tracert command

**Cause:** Adjacency cannot form because of mismatched OSPF configuration.

**Solution:** Use OSPF logging to log errors and warnings to record information on why the adjacency is not forming.

**See also:** Using event logging

**Cause:** Mismatched authentication setting or mismatched passwords.

**Solution:** Verify that the areas are enabled for authentication and the OSPF interfaces are using the same password. Servers running Routing and Remote Access configured to act as OSPF routers have authentication enabled by default and the default password is 12345678. You need to change the authentication to match all neighboring OSPF routers on the same network. The password can vary per network.

**See also:** Create an OSPF area; Set a password on an OSPF interface

**Cause:** The neighboring routers are not configured for the same hello or dead interval.

**Solution:** Verify that the neighboring routers are configured for the same hello and dead intervals. By default the hello interval is 10 seconds and the dead interval is 40 seconds.

**See also:** Configure the hello and dead intervals

**Cause:** Mismatched stub area configuration on neighboring routers.

**Solution:** Verify that the routers agree as to whether the area to which the common network segment belongs is a stub area

or not.

**See also:** Create an OSPF area

**Cause:** Neighboring routers have mismatched Area IDs.

**Solution:** Verify that the interfaces of the neighboring routers are configured with the same Area ID.

**See also:** Create an OSPF area

**Cause:** Non-Broadcast Multiple Access (NBMA) interfaces are not configured with OSPF neighbors.

**Solution:** If the routers are on a NBMA network such as X.25 or Frame Relay that operate in single interface mode, their neighbors must be manually configured by using the unicast IP address of the neighbor or neighbors to which the OSPF packets need to be sent.

**See also:** Add an OSPF neighbor; OSPF design considerations

**Cause:** All router interfaces are set to a router priority of 0.

**Solution:** On broadcast networks (Ethernet, token ring, FDDI) or NBMA networks (X.25 and Frame Relay operating in single interface mode), verify that all routers do not have a router priority of 0. At least one router must have a router priority of 1 or greater so that it can become the DR for the network.

**See also:** Configure an OSPF interface

**Cause:** IP packet filtering is preventing the receiving (through input filters) or sending (through output filters) of OSPF traffic.

**Solution:** Verify that IP packet filtering on the router interfaces is not preventing the receiving (through input filters) or sending (through output filters) of OSPF traffic. OSPF traffic uses IP protocol number 89.

**See also:** Manage Packet Filters

**Cause:** TCP/IP filtering is preventing the receiving of OSPF traffic.

**Solution:** Verify that TCP/IP filtering on the router interfaces is not preventing the receiving of OSPF traffic. OSPF traffic uses IP protocol number 89.

**See also:** Configure TCP/IP to use TCP/IP filtering

**A virtual link is not forming.**
**Cause:** Mismatched configuration of password, hello interval, or dead interval.

**Solution:** Verify that the virtual link neighbor routers are configured for the same password, hello interval, and dead interval.

**See also:** Add a virtual interface

**Cause:** The router ID of the virtual link neighbor is configured incorrectly.

**Solution:** For each router, verify that the router ID of the virtual link neighbor is correctly configured. The router ID is found on the **General** tab on the properties of the OSPF routing protocol.

**See also:** Add a virtual interface

**Cause:** Virtual link neighbors are configured for the incorrect transit area ID.

**Solution:** Verify that both virtual link neighbors are configured for the same correct transit area ID.

**See also:** Add a virtual interface

**Cause:** The retransmit interval is not long enough.

**Solution:** For large internetworks with substantial round-trip delays across the transit area, verify that the retransmit interval is long enough.

**See also:** Add a virtual interface

**Lack of OSPF routes or existence of improper OSPF routes.**
**Cause:** Not receiving summarized routes.

**Solution:** If you are not receiving summarized OSPF routes for an area, verify that the area border routers (ABRs) for the area

are configured with the proper {Destination, Network mask} pairs summarizing that area's routes.

**See also:** Configure ranges for an OSPF area

**Cause:** You are receiving both individual and summarized OSPF routes for an area.

**Solution:** Verify that all the ABRs for the area are configured with the proper {Destination, Network mask} pairs summarizing that area's routes.

**See also:** Configure ranges for an OSPF area

**Cause:** You are not receiving external routes from the autonomous system boundary router (ASBR).

**Solution:** Verify that the source and route filtering that is configured on the ASBR is not too restrictive, preventing proper routes from being propagated to the OSPF autonomous system. External source and route filtering is configured on the **External Routing** tab on the properties of the OSPF routing protocol.

**See also:** Configure an ASBR

**Cause:** All ABRs are not connected to the backbone.

**Solution:** Verify that all ABRs are either physically connected to the backbone or logically connected to the backbone by using a virtual link. There should not be backdoor routers, which are routers that connect two areas without going through the backbone.

**Notes**

- This feature is not available on the Itanium-based versions of the Windows operating systems.
- This content is not available in this preliminary release.
- For more information, see "Unicast IP Routing" at the Microsoft Windows Resource Kits Web site.

# Troubleshooting network address translation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting network address translation**
What problem are you having?

- The network address translation (NAT) computer is not properly translating packets.
- Private network hosts are not receiving IP address configuration.
- Name resolution for private network hosts is not working.

**The network address translation (NAT) computer is not properly translating packets.**
**Cause:** NAT interfaces are not properly configured.

**Solution:** You must add both public (Internet) and private (small office or home office) interfaces to the NAT routing protocol.

You need to verify that the interface on the server running Routing and Remote Access that connects to the Internet is configured for translation. The **Public interface connected to the Internet** option on the **General** tab of the properties of the Internet interface should be selected.

You need to verify that the interfaces on the server running Routing and Remote Access that connects to the small office or home office are properly configured. The **Private interface connected to private network** option on the **General** tab of the properties of the home network interface should be selected.

**See also:** Add and configure an interface for network address translation

**Cause:** TCP/UDP port translation is not enabled.

**Solution:** If you have more private hosts than public IP addresses, you need to verify that the **Translate TCP/UDP headers (recommended)** check box on the **General** tab of the properties of the public interface is selected.

**Cause:** Your range of public addresses is not configured correctly.

**Solution:** If you have multiple public IP addresses, you need to verify that they are properly entered on the **Address Pool** tab of the properties of the Internet interface. If your address pool includes an IP address that was not allocated to you by your ISP, then inbound Internet traffic that is mapped to that IP address may be routed by the ISP to another location.

**See also:** Configure interface IP address ranges

**Cause:** The traffic being forwarded by the network address translation computer is not translatable.

**Solution:** If you have some programs that do not seem to work through the NAT computer, you can try running them from the NAT computer. If they work from the NAT computer and not from a computer on the private network, then the payload of the program may not be translatable.

**Cause:** Your range of private addresses is configured incorrectly.

**Solution:** You need to verify that the range of addresses configured on the **Addressing** tab of the properties of the NAT routing protocol corresponds to a private network (10.0.0.0 with a subnet mask of 255.0.0.0, 172.16.0.0 with a subnet mask of 255.240.0.0, 192.168.0.0 with a subnet mask of 255.255.0.0) or a subnet of a private network. The default range of addresses is 192.168.0.0 with a subnet mask of 255.255.255.0.

**See also:** Enable network address translation addressing

**Cause:** IP packet filtering is preventing the receiving or sending of IP traffic.

**Solution:** You need to verify that IP packet filtering on the private network and Internet interfaces is not preventing the receiving (input filters) or sending (output filters) of Internet-based traffic.

**See also:** Manage Packet Filters

**Private network hosts are not receiving IP address configuration.**
**Cause:** Network address translation addressing is not enabled on the private interface.

**Solution:** You need to verify that network address translation addressing is enabled on the interface that corresponds to your small office or home office network segment.

**See also:** Enable network address translation addressing

**Name resolution for private network hosts is not working.**
**Cause:** Network address translation name resolution is not enabled on the private interface.

**Solution:** You need to verify that network address translation name resolution is enabled on the interface that corresponds to your small office or home office network segment.

**See also:** Enable network address translation name resolution

**Cause:** The network address translation computer is not properly configured for name resolution.

**Solution:** If computers on the small office or home office network are not able to resolve names to IP addresses, then you can check the name resolution configuration of the network address translation computer by using the **ipconfig** command. There are two ways that you can configure name resolution when dialing an ISP:

- Statically assigned name servers

  You must manually configure the TCP/IP protocol with the IP address (or addresses) of the name servers provided by the ISP. If you have statically assigned name servers, you can use the **ipconfig** command at any time to get the IP addresses of your configured name servers.

- Dynamically assigned name servers

  Manual configuration is not required. The IP addresses of the name servers provided by the ISP are dynamically assigned whenever you dial the ISP. If you have dynamically assigned name servers, you must run the **ipconfig** command *after* a connection to the ISP has been made.

# Troubleshooting the DHCP Relay Agent

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting the DHCP Relay Agent**
What problem are you having?

- The DHCP Relay Agent is not providing relay services for DHCP clients on a network segment.

**The DHCP Relay Agent is not providing relay services for DHCP clients on a network segment.**
**Cause:** The interface on the server running Routing and Remote Access that connects to the network segment where the DHCP clients are located is not added to the DHCP Relay Agent IP routing protocol.

**Solution:** Verify that the interface on the server running Routing and Remote Access that connects to the network segment where the DHCP clients are located is added to the DHCP Relay Agent IP routing protocol.

**See also:** Enable the DHCP Relay Agent on a router interface

**Cause:** The **Relay DHCP packets** check box is not selected for the DHCP Relay Agent interface that is connected to the network segment where the DHCP clients are located.

**Solution:** Verify that the **Relay DHCP packets** check box is selected for the DHCP Relay Agent interface that is connected to the network segment where the DHCP clients are located.

**Cause:** The IP addresses of DHCP servers configured on the global properties of the DHCP Relay Agent are incorrect.

**Solution:** Verify that the IP addresses of DHCP servers configured on the global properties of the DHCP Relay Agent are the correct IP addresses for DHCP servers on your internetwork.

**See also:** Configure global DHCP Relay Agent properties

**Cause:** The correctly configured DHCP servers are not reachable.

**Solution:** From the router with the DHCP Relay Agent enabled, use the **ping** command to ping each of the DHCP servers that are configured in the global DHCP Relay Agent dialog box. If you cannot ping the DHCP servers from the DHCP Relay Agent router, troubleshoot the lack of connectivity between the DHCP Relay Agent router and the DHCP server or servers.

**See also:** Troubleshooting TCP/IP

**Cause:** IP packet filtering is preventing the receiving (through input filters) or sending (through output filters) of DHCP traffic.

**Solution:** Verify that IP packet filtering on the router interfaces is not preventing the receiving (through input filters) or sending (through output filters) of DHCP traffic. DHCP traffic uses the User Datagram Protocol (UDP) ports of 67 and 68.

**See also:** Manage Packet Filters

**Cause:** TCP/IP filtering is preventing the receiving of DHCP traffic.

**Solution:** Verify that TCP/IP filtering on the router interfaces is not preventing the receiving of DHCP traffic. DHCP traffic uses the UDP ports of 67 and 68.

**See also:** Configure TCP/IP to use TCP/IP filtering

# Troubleshooting IP multicast

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting IP multicast**
What problem are you having?

**The IGMP router does not work correctly.**

**Cause:** The network adapters that are installed may not support multicast promiscuous mode.

**Solution:** Multicast promiscuous mode is a special mode of the network adapter in which all multicasted traffic is passed up to the network operating system. Not all network adapters support multicast promiscuous mode. If a server running Windows Server 2003 fails to put the network adapter into multicast promiscuous mode during the initialization of the IGMP router component, then an event is logged in the System Event Log. The event has the following attributes:

- Source: IP Router Manager
- Event ID: 20157
- Message text: The interface [name of the interface] could not be enabled for multicast. IGMP router will not be activated over this interface.

**See also:** Using the IP multicasting utilities

# Troubleshooting demand-dial routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting demand-dial routing**
What problem are you having?

- On-demand connection is not made automatically.
- Unable to make a demand-dial connection.
- Unable to reach locations beyond the calling router or answering router.
- Auto-static updates are not working.

**On-demand connection is not made automatically.**
**Cause:** IP routing is not enabled.

**Solution:** Verify that IP routing is enabled on the **IP** tab on the properties of the server.

**See also:** View properties of the remote access server

**Cause:** Static routes are configured incorrectly.

**Solution:** Verify that the correct static routes exist and are configured with the appropriate demand-dial interface.

**See also:** Add a static route

**Cause:** A static route cannot initiate a demand-dial connection.

**Solution:** For the static routes that use a demand-dial interface, verify that the **Use this route to initiate demand-dial connections** check box is selected.

**See also:** Add a static route

**Cause:** The demand-dial interface is disabled.

**Solution:** Verify that the demand-dial interface is not in a disabled state. To enable, under **Network Interfaces**, right-click the demand-dial interface, and then click **Enable**.

**Cause:** Dial-out hours are preventing the connection from initiating.

**Solution:** Verify that the dial-out hours for the demand-dial interface on the calling router are not preventing the connection attempt.

**See also:** Configure dial-out hours

**Cause:** Demand-dial filters are preventing the connection from initiating.

**Solution:** Verify that the demand-dial filters for the demand-dial interface on the calling router are not preventing the connection attempt.

**See also:** Configure demand-dial filters

**Unable to make a demand-dial connection.**
**Cause:** The Routing and Remote Access service is not started on the calling router and the answering router.

**Solution:** Verify the state of the Routing and Remote Access service on the calling router and the answering router.

**See also:** Monitor the Routing and Remote Access service

**Cause:** The demand-dial interface is in an unreachable state.

**Solution:** Determine the unreachability reason. If a demand-dial connection fails, the demand-dial interface is in an

unreachable state. The Routing and Remote Access service records the reason why the connection failed through the unreachability reason. Based on the information in the unreachability reason, you can perform further troubleshooting.

**See also:** Check the unreachability reason

**Cause:** LAN and WAN routing is not enabled on the calling router and the answering router.

**Solution:** Enable **LAN and demand-dial routing** on the calling router and the answering router.

**See also:** Enable LAN and WAN routing

**Cause:** Demand-dial connections for the protocols being routed are not allowed.

**Solution:** Verify that IP-based remote access and demand-dial connections are allowed (on the **IP** tab on the properties of the server).

**See also:** View properties of the remote access server

**Cause:** Dial-up, PPTP, or L2TP ports are not enabled for inbound and outbound demand-dial routing connections.

**Solution:** Enable dial-up, PPTP, or L2TP ports for inbound and outbound demand-dial routing connections.

**See also:** Enable routing on ports

**Cause:** All of the ports on the calling or answering router are already being used by currently connected remote access clients or demand-dial routers.

**Solution:** Verify that not all of the ports are already being used by clicking **Ports** in Routing and Remote Access. For VPN connections, if necessary, change the number of PPTP or L2TP ports to allow more concurrent connections.

**See also:** Add PPTP or L2TP ports

**Cause:** For VPN connections, the tunneling protocol of the calling router is not supported by the answering router.

By default, VPN-based demand-dial interfaces use the **Automatic** server type option, which means that they try to establish a PPTP-based VPN connection first, and then they try an L2TP/IPSec VPN connection. If calling routers use either the **Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP)** server type option, verify that the selected tunneling protocol is supported by the answering router.

By default, a server running Routing and Remote Access is a PPTP and L2TP-capable demand-dial router. To create a PPTP-only router, set the number of L2TP ports to zero. To create an L2TP-only router, set the number of PPTP ports to one.

**Solution:** Verify that the appropriate number of PPTP or L2TP ports is configured on the calling router and answering router.

**See also:** Add PPTP or L2TP ports

**Cause:** The calling router and the answering router in conjunction with a remote access policy are not configured to use at least one common authentication method.

**Solution:** Configure the calling router and the answering router in conjunction with a remote access policy to use at least one common authentication method.

**Cause:** The calling router and the answering router in conjunction with a remote access policy are not configured to use at least one common encryption strength.

**Solution:** Configure the calling router and the answering router in conjunction with a remote access policy to use at least one common encryption strength.

**See also:** Configure encryption

**Cause:** The account option on the **Account** tab on the properties of the user account of the calling router is configured to change the password at the next logon attempt.

**Solution:** Clear the **User must change password at next logon** check box on the **Account** tab on the properties of the user account of the calling router.

**Cause:** The password of the user account of the calling router has expired.

**Solution:** Reset the password and select the **Password never expires** check box on the **Account** tab on the properties of the user account of the calling router.

**Cause:** The demand-dial connection does not have the appropriate permissions through dial-in properties of the user account and remote access policies.

**Solution:** Verify that the demand-dial connection has the appropriate permissions through dial-in properties of the user account and remote access policies. If the answering router is configured to use Windows Authentication, the remote access policies stored on the answering router are used. If the answering router is configured to use RADIUS authentication and the RADIUS server being used is an Internet Authentication Service (IAS) server, then the remote access policies of the IAS server are used.

In order for the connection to be established, the settings of the connection attempt must:

- Match all of the conditions of at least one remote access policy.
- Be granted remote access permission through the local user account (set to **Allow access**) or through the domain user account (set to **Control access through Remote Access Policy**) and the remote access permission of the matching remote access policy (set to **Grant remote access permission**).
- Match all the settings of the profile.
- Match all the settings of the dial-in properties of the user account.

**See also:** Introduction to remote access policies; Accepting a connection attempt

**Cause:** The settings of the remote access policy profile are in conflict with properties of the answering router.

**Solution:** Verify that the settings of the remote access policy profile are not in conflict with properties of the answering router.

The properties of the remote access policy profile and the properties of the answering router both contain settings for:

- Multilink
- Bandwidth allocation protocol
- Authentication protocols

If the settings of the profile of the matching remote access policy are in conflict with the settings of the answering router, the connection attempt is rejected. For example, if the matching remote access policy profile specifies that the EAP-TLS authentication protocol must be used and EAP is not enabled on the answering router, the connection attempt is rejected.

**See also:** Enable authentication protocols; Configure authentication

**Cause:** The credentials of the calling router (user name, password, and domain name) are incorrect and cannot be validated by the answering router.

**Solution:** Verify that the credentials of the calling router (user name, password, and domain name) are correct and can be validated by the answering router.

**Cause:** There are not enough addresses in the static IP address pool.

**Solution:** If the answering router is configured with a static IP address pool, verify that there are enough addresses in the pool. If all of the addresses in the static pool have been allocated to connected remote access clients or demand-dial routers, the answering router is unable to allocate an IP address, and the connection attempt is rejected. Modify the static IP address pool if needed.

**See also:** Create a static IP address pool

**Cause:** The authentication provider of the answering router is improperly configured.

**Solution:** Verify the configuration of the authentication provider. You can configure the answering router to use either the security features of the Windows Server 2003 family or RADIUS to authenticate the credentials of the VPN client.

**See also:** Use RADIUS authentication

**Cause:** The answering router cannot access Active Directory.

**Solution:** For an answering router that is a member server of a domain that is configured for Windows Authentication, verify that:

- The **RAS and IAS Servers** security group exists. If not, then create the group and set the group type to **Security** and the group scope to **Domain local**.

- The **RAS and IAS Servers** security group has Read permission to the **RAS and IAS Servers Access Check** object.

- The computer account of the answering router computer is a member of the **RAS and IAS Servers** security group. You can use the **netsh ras show registeredserver** command to view the current registration. You can use the **netsh ras add registeredserver** command to register the server in a specified domain.

  If you add or remove the answering router computer to the **RAS and IAS Servers** security group, the change will not take effect immediately (due to the way that the Windows Server 2003 family caches Active Directory information). For the change to take effect immediately, you need to restart the answering router computer.

- The answering router has joined the domain.

**See also:** Create a new group; Verify permissions for the RAS and IAS security group; Netsh commands for remote access

**Cause:** An answering router running Windows NT 4.0 with the Routing and Remote Access Service (RRAS) cannot validate connection requests.

**Solution:** If calling routers are dialing in to an answering router running Windows NT 4.0 with RRAS that is a member of a domain, verify that the Everyone group is added to the Pre-Windows 2000 Compatible Access group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, type **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** at the command prompt on a domain controller computer and then restart the domain controller computer.

**Cause:** The Fax Service is enabled and your modem does not support adaptive answer.

**Solution:** Verify that if the Fax Service and the Routing and Remote Access service are sharing the same modem, that the modem supports adaptive answer. If the modem does not support adaptive answer, you must disable the Fax Service on the modem to receive incoming demand-dial or remote access connections.

**Cause:** The calling router is not correctly configured for certificate-based demand-dial routing.

**Solution:** Verify that EAP is enabled as an authentication protocol. Verify that the correct certificate for the root certification authority certificate of the answering router is selected. Verify that the correct router (offline request) certificate is selected when configuring the credentials of the demand-dial interface.

**See also:** Configure the calling router for certificate-based EAP; Deploying certificate-based authentication for demand-dial routing

**Cause:** The answering router is not correctly configured for certificate-based demand-dial routing.

**Solution:** Verify that EAP is enabled as an authentication protocol. Verify that the correct computer certificate for the root certification authority of the answering router is selected.

**See also:** Configure the answering router for certificate-based EAP; Deploying certificate-based authentication for demand-dial routing

**Cause:** You are using MS-CHAP v1 and a user password over 14 characters long.

**Solution:** Either use a different authentication protocol such as MS-CHAP v2 or change your password so that it is 14 characters or less.

**See also:** MS-CHAP; Enable authentication protocols

**Unable to reach locations beyond the calling router or answering router.**
**Cause:** IP routing is not enabled.

**Solution:** Verify that IP routing is enabled (on the **IP** tab on the properties of the server).

**See also:** View properties of the remote access server

**Cause:** The appropriate demand-dial interface has not been added to the protocol being routed.

**Solution:** Add the appropriate demand-dial interface to the protocol being routed.

**See also:** Add a routing interface

**Cause:** There are no routes on both sides of the demand-dial connection that support the two-way exchange of traffic.

**Solution:** Unlike a remote access connection, a demand-dial connection does not automatically create a default route. You need to create routes on both sides of the demand-dial connection so that traffic can be routed to and from the other side of the demand-dial connection.

You can manually add static routes to the routing table, or you can add static routes through routing protocols. For persistent demand-dial connections, you can enable Open Shortest Path First (OSPF) or Routing Information Protocol (RIP) across the demand-dial connection. OSPF is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family. For on-demand demand-dial connections, you can automatically update routes through an auto-static RIP update.

**See also:** Add an IP routing protocol; Add a static route; Perform Auto-Static Updates

**Cause:** There are no routes on the calling router and answering router's intranet that support the forwarding of packets between hosts on the intranets.

**Solution:** Verify that there are routes in the intranet routers on the calling router and answering router's intranet so that all locations on both networks are reachable. You can add routes to the routers of each intranet through static routes or by enabling a routing protocol on the intranet interface of the calling and answering routers.

**See also:** Unicast routing overview; Deploying Routing

**Cause:** A two-way initiated, demand-dial connection is being interpreted by the answering router as a remote access connection.

**Solution:** In order for the answering router to determine that the incoming call is a router rather than remote access client, the user name of the credentials for the calling router must match the name of a demand-dial interface that is configured on the answering router.

If the incoming caller is a router, the port on which the call was received shows a status of **Active** and the corresponding demand-dial interface is in a **Connected** state. If the name of the user name credential for the calling router appears under **Remote Access Clients** in Routing and Remote Access, then the calling router has been interpreted by the answering router as a remote access client.

For two-way initiated connections, either router can be the calling router or the answering router. The user names and demand-dial interface names must be properly matched. For example, two-way initiated connections would work under the following configuration:

- Router 1 has a demand-dial interface called NEW-YORK which is configured to use SEATTLE as the user name when sending authentication credentials.

- Router 2 has a demand-dial interface called SEATTLE which is configured to use NEW-YORK as the user name when sending authentication credentials.

This example assumes that the SEATTLE user name can be validated by Router 2 and the NEW-YORK user name can be validated by Router 1.

**See also:** Demand-dial routing example; Check the status of the port on the answering router; Check the status of the demand-dial interface

**Cause:** Static routes on the user account of the calling router for a one-way initiated demand-dial connection are not configured correctly.

**Solution:** For a one-way initiated demand-dial connection, verify that the appropriate static routes are enabled on the user account of the calling router and that the answering router is configured with a routing protocol so that when a connection is made, the static routes of the user account of the calling router are advertised to neighboring routers.

**See also:** One-way initiated demand-dial connections; Configure static routes for a dial-in user

**Cause:** Packet filters on the demand-dial interfaces of the calling router or answering router are preventing the flow of traffic.

**Solution:** Verify that there are no IP packet filters on the demand-dial interfaces of the calling router or the answering router that prevent the sending or receiving of TCP/IP traffic.

You can configure each demand-dial interface with IP input and output filters to control the exact nature of TCP/IP traffic that is allowed into and out of the demand-dial interface.

**See also:** Manage Packet Filters

**Cause:** Packet filters on the remote access policy profile are preventing the flow of IP traffic.

**Solution:** Verify that there are no configured TCP/IP packet filters on the profile properties of the remote access policy being used by the demand-dial connection (or the RADIUS server if Internet Authentication Service is used) that are preventing the sending or receiving of TCP/IP traffic.

You can use remote access policies to configure TCP/IP input and output packet filters that control the exact nature of TCP/IP traffic allowed on the demand-dial connection.

**See also:** Configure IP options

**Auto-static updates are not working.**
**Cause:** For IP auto-static updates, each demand-dial interface is not configured for RIP v2 multicast.

**Solution:** Verify on both routers that the demand-dial interface is added to the RIP routing protocol, that its operation mode is set to **Auto-static update** mode, the outgoing packet protocol is **RIP version 2 multicast**, and that the incoming packet protocol is set to **RIP version 1 and 2**.

**See also:** Configure RIP version 2; Enable auto-static update mode

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access**
The remote access features of the Microsoft® Windows Server 2003 family enable remote or mobile workers who use dial-up communication links to access corporate networks as if they were directly connected. Remote access also provides virtual private network (VPN) services so that users can access corporate networks over the Internet.

- Before installing Routing and Remote Access, see Checklist: Installing and configuring the remote access server.
- To find features that have been moved in the Windows Server 2003 family, see New ways to do familiar Remote Access tasks.
- For tips about using remote access, see Best Remote Access practices.
- For help with specific tasks, see Remote Access How To....
- For general background information, see Remote Access Concepts.
- For problem-solving instructions, see Remote Access Troubleshooting.

# Checklist: Installing and configuring the remote access server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring the remote access server**

| Step | Reference |
|---|---|
| Review key concepts. | Remote Access Overview |
| Verify the compatibility of all hardware to be installed in the server running Routing and Remote Access. | Compatible Hardware and Software at Support resources |
| Install the hardware. | Manufacturer's documentation |
| Verify that the hardware is successfully installed. | Device Manager |
| Install and configure modems. | Checklist: Installing and configuring modems |
| Install and configure the protocols. | Network communications |
| Enable and configure the Routing and Remote Access service. | Enable the Routing and Remote Access service |
| Configure dial-in properties and remote access policies for dial-in permission, authentication, and encryption settings. | Introduction to remote access policies |
| (Optional) Configure the server running Routing and Remote Access as a RADIUS client to an IAS server for centralized remote access policy management. | Using RADIUS |

# New ways to do familiar Remote Access tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring the remote access features of Routing and Remote Access. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0 or Windows NT version 4.0 with the Routing and Remote Access Service (RRAS). The user interface for the Windows 2000 version of Routing and Remote Access differs slightly from that for the Windows Server 2003 family, but the steps for most common tasks have not changed.

Please choose your upgrade path:

- Windows NT 4.0 to a product in the Windows Server 2003 family
- Windows NT 4.0 with RRAS to a product in the Windows Server 2003 family

**Windows NT 4.0 to a product in the Windows Server 2003 family**

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Enable and configure remote access | **Services** tab of **Network** in Control Panel | Open Routing and Remote Access. For more information, see Enable the Routing and Remote Access service and Configure the Remote Access Server. |
| Configure authentication and encryption options | **Services** tab of **Network** in Control Panel | Open Routing and Remote Access. For more information, see Configure Remote Access Security. |
| Manage remote access servers and clients | Remote Access Admin | Open Routing and Remote Access. For more information, see Administer the Remote Access Server and Manage Remote Access Clients. |
| Grant remote access permission to user accounts | Remote Access Admin or User Manager for Domains | Open Computer Management or Active Directory Users and Computers. For more information, see Configure remote access permission for a user. |

**Windows NT 4.0 with RRAS to a product in the Windows Server 2003 family**

| If you want to | In Windows NT 4.0 with RRAS use | In this version of Windows use |
|---|---|---|
| Enable and configure remote access | **Services** tab of **Network** in Control Panel | Open Routing and Remote Access. For more information, see Enable the Routing and Remote Access service and Configure the Remote Access Server. |
| Configure authentication and encryption options | **Services** tab of **Network** in Control Panel | Open Routing and Remote Access. For more information, see Configure Remote Access Security. |
| Manage remote access servers and clients | Routing and RAS Admin | Open Routing and Remote Access. For more information, see Administer the Remote Access Server and Manage Remote Access Clients. |
| Grant remote access permission to user accounts | Remote Access Admin or User Manager for Domains | Open Computer Management or Active Directory Users and Computers. For more information, see Configure remote access permission for a user. |

# Best Remote Access practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

The following list provides best practices for implementing and configuring the remote access server and is based on recommendations from Microsoft Product Support Services:

- **Plan for security**

  When designing and implementing your remote access solution, be sure to plan for the most secure deployment possible. For more information, see Security information for remote access.

- **Use DHCP to obtain IP addresses**

  If you installed a DHCP server, configure the remote access server to use DHCP to obtain IP addresses for remote access clients. If you did not install a DHCP server, configure the remote access server with a static IP address pool, which is a subset of addresses from the subnet to which the remote access server is attached.

- **Use strong authentication**

  Use strong passwords more than 8 characters long that contain a mixture of uppercase and lowercase letters, numbers, and permitted punctuation. Do not use passwords based on names or words. Strong passwords are more resistant to a dictionary attack, where an attacker attempts to determine a password by sending a series of commonly used names and words.

  Although EAP-TLS works with registry-based certificates, for security reasons you should use EAP-TLS with smart cards only.

  If you are using MS-CHAP, use MS-CHAP version 2. You can obtain the latest MS-CHAP updates for Windows NT version 4.0, Windows 98, and Windows 95. For more information, see MS-CHAP version 2.

- **Avoid configuring different remote access policies for the same user**

  If a user dials in by using a multilink connection, all connections beyond the first connection are connected by using the remote access policy that matched the first connection.

# Remote Access How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

This section provides instructions for installing and configuring Routing and Remote Access to act as remote access server:

- Enable the Routing and Remote Access service
- Configure the Remote Access Server
- Configure Dial-up Equipment
- Configure Remote Access Security
- Configure dial-in user properties
- Administer the Remote Access Server

# Enable the Routing and Remote Access service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable the Routing and Remote Access service**

1. If this server is a member of an Active Directory domain and you are not a domain administrator, instruct your domain administrator to add the computer account of this server to the **RAS and IAS Servers** security group in the domain of which this server is a member. The domain administrator can add the computer account to the **RAS and IAS Servers** security group by using Active Directory Users and Computers or with the **netsh ras add registeredserver** command. If this server is using local authentication or is authenticating against a RADIUS server, skip this step.

2. Open Routing and Remote Access.

3. By default, the local computer is listed as a server.

   To add another server, in the console tree, right-click **Server Status**, and then click **Add Server**.

   In the **Add Server** dialog box, click the applicable option, and then click **OK**.

4. In the console tree, right-click the server you want to enable, and then click **Configure and Enable Routing and Remote Access**.

5. Follow the instructions in the Routing and Remote Access Server Setup Wizard.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Installing the router
Installing the remote access server
Netsh commands for remote access

# Configure the Remote Access Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure the remote access server**

- View properties of the remote access server
- Enable the remote access server
- Create a static IP address pool
- Configure AppleTalk remote access
- Enable Multilink
- Enable BAP and BACP
- Enable or disable broadcast name resolution
- Re-create the default remote access policy
- Configure Ports

# View properties of the remote access server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view properties of the remote access server**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to view properties, and then click **Properties**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Enable authentication protocols
Enable the remote access server

# Enable the remote access server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable the remote access server**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to enable remote access, and then click **Properties**.
3. On the **General** tab, select the **Remote access server** check box.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Remote Access Troubleshooting
View properties of the remote access server

# Create a static IP address pool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a static IP address pool**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to create a static IP address pool, and then click **Properties**.
3. On the **IP** tab, click **Static address pool**, and then click **Add**.
4. In **Start IP address**, type a starting IP address, and then either type an ending IP address for the range in **End IP address** or type the number of IP addresses in the range in **Number of addresses**.
5. Click **OK**, and then repeat steps 3 and 4 for as many ranges as you need to add.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- If the static IP address pool consists of ranges of IP addresses that are for a separate subnet, then you need to either enable an IP routing protocol on the remote access server computer or add static IP routes consisting of the {IP Address, Mask} of each range to the routers of the intranet. If the routes are not added, then remote access clients cannot receive traffic from resources on the intranet.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Expressing an IP address range with a mask
Remote Access Troubleshooting

# Configure AppleTalk remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure AppleTalk remote access**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to configure AppleTalk remote access, and then click **Properties**.
3. On the **AppleTalk** tab, configure the AppleTalk remote access settings.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Enable AppleTalk routing
Unicast routing overview

# Enable Multilink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Multilink**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to enable Multilink, and then click **Properties**.
3. On the **PPP** tab, select the **Multilink connections** check box.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Multilink and BAP
Enable BAP and BACP

# Enable BAP and BACP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable BAP and BACP**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to enable BAP and BACP, and then click **Properties**.
3. On the **PPP** tab, select the **Dynamic bandwidth control using BAP and BACP** check box.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Multilink and BAP
New features of remote access
Enable Multilink

# Enable or disable broadcast name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable broadcast name resolution**

1. Open Routing and Remote Access.
2. Right-click the name of the server on which you want to enable or disable broadcast name resolution, and then click **Properties**.
3. On the **IP** tab, select or clear the **Enable broadcast name resolution** check box, and click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- Broadcast name resolution is enabled by default.
- If you enable or disable broadcast name resolution, you must restart the Routing and Remote Access service on the server for which you enabled or disabled broadcast name resolution.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
New features of remote access
Broadcast name resolution

# Re-create the default remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To re-create the default remote access policy**

1. Open Routing and Remote Access.
2. In the console tree, right-click **Remote Access Policies**.

   **Where?**

   - Routing and Remote Access/*server name*/Remote Access Policies

3. Click **New Remote Access Policy**.
4. In the New Remote Access Policy wizard, do the following:

   - Choose **Set up a custom policy**.
   - In **Policy name**, type a default name, and then click **Next**.

     When the Routing and Remote Access service is installed, the default remote access policies are **Connections to Microsoft Routing and Remote Access server** and **Connections to other access servers**. However, you are not required to use these names.

   - On the Policy Conditions page, click **Add**, click **Day-and-Time-Restrictions**, click **Add**, and configure the attribute to permit all times on all days. Click **OK**, and then click **Next**.
   - On the Permissions page, click **Deny remote access permission**, and then click **Next**.
   - On the Profile page, click **Next**. Do not make changes to the profile for this policy.
   - Click **Finish**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- By following the steps in this procedure, you will re-create the policy conditions of the **Connections to other access servers** policy.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Working with MMC console files
Introduction to remote access policies
Add a remote access policy

# Configure Ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure ports**

- Configure ports for remote access
- Set the phone number on a port
- Add PPTP or L2TP ports

# Configure ports for remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure ports for remote access**

1. Open Routing and Remote Access.
2. In the console tree, right-click **Ports**.

   **Where?**

   - Routing and Remote Access/*server name*/Ports

3. Click **Properties**.
4. In the **Ports Properties** dialog box, click a device, and then click **Configure**.
5. In the **Configure Device** dialog box, do one or more of the following:

   - To enable remote access, select the **Remote access connections (inbound only)** check box.
   - To enable demand-dial routing, select the **Demand-dial routing connections (inbound and outbound)** check box.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Devices and ports
Set the phone number on a port
Multilink and BAP

# Set the phone number on a port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the phone number on a port**

1. Open Routing and Remote Access.
2. In the console tree, click **Ports**.

   **Where?**

   - Routing and Remote Access/*server name*/Ports

3. Right-click **Ports**, and then click **Properties**.
4. In the **Ports Properties** dialog box, click the port that corresponds to the dial-up equipment, and then click **Configure**.
5. In **Phone number for this device**, type the phone number of the port, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- You also set the phone number on a port when you are using the **Restrict Dial-in to this number only** dial-in constraint of a remote access policy profile and your phone line and installed phone equipment do not support Called Line Identification (CLID).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Devices and ports
Multilink and BAP

# Add PPTP or L2TP ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add PPTP or L2TP ports**

1. Open Routing and Remote Access.
2. In the console tree, click **Ports**.

   **Where?**

   - Routing and Remote Access/*server name*/Ports

3. Right-click **Ports**, and then click **Properties**.
4. In the **Ports Properties** dialog box, click either **WAN Miniport (PPTP)** or **WAN Miniport (L2TP)**, and then click **Configure**.
5. In **Maximum ports**, type the number of ports, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.
- You cannot set the number of PPTP ports to 0. To disable PPTP, clear the **Remote access connections (inbound only)** and **Demand-dial routing connections (inbound and outbound)** check boxes.
- When you increase the number of ports for **WAN Miniport (PPTP)** or **WAN Miniport (L2TP)** from less than 1000 ports to greater than 1000 ports, you must restart the computer to apply the change.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Point-to-Point Tunneling Protocol
Layer Two Tunneling Protocol
Devices and ports
Enable routing on ports

# Configure Dial-up Equipment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure dial-up equipment**

- Configure a direct serial connection
- Configure modem-pooling equipment
- Configure an X.25 smart card
- Configure other security devices

# Configure a direct serial connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a direct serial connection**

1. Open Control Panel.
2. Double-click **Phone and Modem Options**, click the Modems tab, and then click **Add**.
3. In the Add Hardware wizard, select the **Don't detect my modem; I will select it from a list** check box, and then click **Next**.
4. In **Manufacturer**, click **Standard Modem Types**.
5. In **Models**, click **Communications cable between two computers**, and then click **Next**.
6. Follow the remaining instructions in the Install New Modem wizard.

## Notes

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- A null modem must be configured on both the client and the server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Data compression
Modems

# Configure modem-pooling equipment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure modem-pooling equipment**

1. Open Control Panel.
2. Configure your modem-pooling equipment to behave like one of the modem types listed in the Add Hardware wizard. (In other words, the modem-pooling equipment must generate and accept command strings as if it were a modem of the chosen type.) The switching equipment must also have the same RS-232 signal behavior as the specified modem.
3. Connect COM ports to the equipment and configure the ports for remote access in Routing and Remote Access.

## Notes

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- It is recommended that you configure the modem-pooling equipment as a Hayes-compatible modem, a widely known standard.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Modems
Configure ports for remote access

# Configure an X.25 smart card

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure an X.25 smart card**

1. Install the X.25 smart card (according to the manufacturer's instructions).
2. Make sure your X.25 smart card is configured with the X.3-parameter values as shown in the following table.
3. Configure the X.25 ports for remote access in Routing and Remote Access.

| Parameter number | X.3 parameter | Value |
|---|---|---|
| 1 | PAD Recall | 0 |
| 2 | Echo | 0 |
| 3 | Data Fwd Char | 0 |
| 4 | Idle Timer | 1 |
| 5 | Device Control | 0 |
| 6 | PAD Service Signals | 1 |
| 7 | Break Signal | 0 |
| 8 | Discard Output | 0 |
| 9 | Padding after CR | 0 |
| 10 | Line Folding | 0 |
| 11 | *Not Set* | |
| 12 | Flow Control | 0 |
| 13 | Linefeed Insertion | 0 |
| 14 | Padding after LF | 0 |
| 15 | Editing | 0 |
| 16 | Character Delete | 0 |
| 17 | Line Delete | 0 |
| 18 | Line Display | 0 |
| 19 | Editing PAD Srv Signals | 0 |
| 20 | Echo Mask | 0 |
| 21 | Parity Treatment | 0 |
| 22 | Page Wait | 0 |

**Important**

- Failure to set these values as shown prevents remote access that uses an X.25 smart card from functioning properly. For

more information on setting these values, see the instructions provided with your X.25 smart card.

### Notes

- A communication driver for the X.25 smart card that emulates communication ports is supplied by the hardware manufacturer or by another company.

- Make sure that the speed of the leased line can support all the serial communication (COM) ports at all speeds at which clients will dial in. For example, four clients connecting at 9,600 bps (through dial-up PADs) will require a 38,400-bps (4 times 9,600) leased line on the server end. If the leased line does not have adequate bandwidth, it can cause time-outs and can cause the performance for connected clients to degrade.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
X.25
Configure ports for remote access

# Configure other security devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure other security devices**

- If the modem on the remote access server is different from the modem listed in the security host section of the Modem.inf file, you need to customize the Modem.inf file on the remote access server to link the security host to the modem on the server.

- The remote user must activate the Terminal utility to interact with the security host.

**Note**

- To use a Security Dynamics security host, you must order two connectors through your Security Dynamics provider to permit initialization of the remote access server modem. When you order, specify that you want the dial-out option. The provider will then send you an AND gate and a jumper box. For the ACM/400 security host, you will also receive different software.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Remote Access Security hosts

# Configure Remote Access Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure remote access security**

- Use Windows Authentication
- Use Windows Accounting
- Use RADIUS authentication
- Use RADIUS accounting
- Enable authentication protocols
- Enable EAP
- Configure EAP-RADIUS
- Enable reversibly encrypted passwords in a domain
- Configure logging
- Configure smart card remote access
- Configure a pre-shared key on a remote access server
- Configure a pre-shared key on a remote access client
- Verify permissions for the RAS and IAS security group
- Map a certificate to a user account
- Manage Machine Certificates

# Use Windows Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use Windows Authentication**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to configure authentication, and then click **Properties**.
3. On the **Security** tab, in **Authentication provider**, click **Windows Authentication**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Remote Access Authentication

# Use Windows Accounting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use Windows Accounting**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to configure Windows Accounting, and then click **Properties**.
3. On the **Security** tab, in **Accounting provider**, click **Windows Accounting**, and then click **OK**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Logging

# Use RADIUS authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use RADIUS authentication**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to configure RADIUS authentication, and then click **Properties**.
3. On the **Security** tab, in **Authentication provider**, click **RADIUS authentication**, and then click **Configure**.
4. In the **RADIUS Authentication** dialog box, click **Add**.
5. In the **Add RADIUS Server** dialog box, configure the settings for your RADIUS authentication server, and then click **OK**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Introduction to remote access policies
Using RADIUS

# Use RADIUS accounting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use RADIUS accounting**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to configure RADIUS accounting, and then click **Properties**.
3. On the **Security** tab, in **Accounting provider**, click **RADIUS accounting**, and then click **Configure**.
4. In the **RADIUS Accounting** dialog box, click **Add**.
5. In the **Add RADIUS Server** dialog box, configure the settings for your RADIUS accounting server, and then click **OK**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Remote Access Authentication
Introduction to remote access policies
Using RADIUS

# Enable authentication protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable authentication protocols**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to enable authentication protocols, and then click **Properties**.
3. On the **Security** tab, click **Authentication Methods**.
4. In the **Authentication Methods** dialog box, select the appropriate check boxes for the authentication protocols that the remote access server will use to authenticate remote clients, and then click **OK**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Remote Access Authentication Methods

# Enable EAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable EAP**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to configure EAP, and then click **Properties**.
3. On the **Security** tab, click **Authentication Methods**.
4. In the **Authentication Methods** dialog box, select the **Extensible authentication protocol (EAP)** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- When you enable EAP, all the installed EAP types are enabled. By default, EAP-MD5 CHAP and EAP-TLS are installed if the server is a member of a domain. If the server is not a member of a domain, only EAP-MD5 is installed. To see the installed EAP types, click **EAP Methods**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
EAP
Configure EAP-RADIUS
Configure authentication
Network access authentication and certificates

# Configure EAP-RADIUS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure EAP-RADIUS**

1. Open Routing and Remote Access.
2. Right-click the server name for which you want to configure EAP-RADIUS, and then click **Properties**.
3. On the **Security** tab, click **Authentication Methods**.
4. In the **Authentication Methods** dialog box, select the **Extensible authentication protocol (EAP)** check box, and then click **OK**.
5. In **Authentication provider**, click **RADIUS Authentication**, and then click **Configure**.
6. In the **RADIUS Authentication** dialog box, click **Add**.
7. In the **Add RADIUS Server** dialog box, configure the settings for your RADIUS authentication server, and then click **OK**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
EAP
Enable EAP
Configure authentication

# Enable reversibly encrypted passwords in a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable reversibly encrypted passwords in a domain**

1. Open Active Directory Users and Computers.
2. In the console tree, double-click **Active Directory Users and Computers**, right-click the domain name, and then click **Properties**.
3. On the **Group Policy** tab, click **Default Domain Policy**, and then click **Edit**.
4. In the console tree, click **Password Policy**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Account Policies/Password Policy

5. In the details pane, double-click **Store passwords using reversible encryption for all users in the domain**.
6. Click **Enabled**, and then click **OK**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- To open Group Policy on a stand-alone computer that is not a member of a domain, click **Start**, click **Run**, type **gpedit.msc**, and then click **OK**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
CHAP

# Configure logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure logging**

1. Do one of the following:

   - Open Routing and Remote Access.

     Double-click **Routing and Remote Access**, and then double-click the server name on which you want to configure logging.

   - Open Internet Authentication Service.

     Double-click **Internet Authentication Service**.

2. In the console tree, click **Remote Access Logging**.

3. In the details pane, right-click any log file, and then click **Properties**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- To open Internet Authentication Service, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Internet Authentication Service**.

- If the Logging folder does not appear in Routing and Remote Access, then Windows Accounting is not enabled.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Logging
Use Windows Authentication
Use Windows Accounting

# Configure smart card remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure smart card remote access**

1. Open Routing and Remote Access.
2. Right-click the name of the remote access server, and then click **Properties**.
3. On the **Security** tab, click **Authentication Methods**.
4. In the **Authentication Methods** dialog box, confirm that the **Extensible authentication protocol (EAP)** check box is selected, click **OK**, and then click **OK** again.
5. In the console tree, click **Remote Access Policies**.

   **Where?**

   - Routing and Remote Access/*ServerName*/Remote Access Policies

6. In the details pane, right-click the remote access policy that your smart card remote access clients will use, click **Properties**, and then click **Edit Profile**.
7. On the **Authentication** tab, click **EAP Methods**.
8. In **Select EAP Providers**, click **Add**.
9. Click the smart card EAP method you want to use, and then click **OK**.
10. If the smart card EAP method you added requires a certificate, click the method, click **Edit**, and choose a certificate.
11. Click **OK** to save the settings of the profile, and then click **OK** again to save the settings of the policy.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Using smart cards for remote access
Enable the IAS server to read user accounts in Active Directory

# Configure a pre-shared key on a remote access server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a pre-shared key on a remote access server**

1. Open Routing and Remote Access.
2. Right-click the name of the server on which you want to configure a pre-shared key, and then click **Properties**.
3. On the **Security** tab, select the **Allow custom IPSec policy for L2TP connection** check box.
4. Type the pre-shared key in **Pre-shared Key**.
5. Click **Apply**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- A pre-shared key can be any non-null string of any combination of up to 255 characters. Use care when typing the pre-shared key because any differences between the pre-shared key configured on the client and the pre-shared key configured on the server will prevent connections from being made.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Pre-shared keys
Configure a pre-shared key on a remote access client

# Configure a pre-shared key on a remote access client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a pre-shared key on a remote access client**

1. Open Network Connections.
2. Right-click the virtual private network connection for which you want to configure a pre-shared key, and click **Properties**.
3. On the **Security** tab, click **IPSec Settings**.
4. Select the **Use pre-shared key for authentication** check box, and type the pre-shared key in **Key**.
5. Click **OK** twice.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- A pre-shared key can be any non-null string of any combination of up to 255 characters. Use care when typing the pre-shared key because any differences between the pre-shared key configured on the client and the pre-shared key configured on the server will prevent connections from being made.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Pre-shared keys
Configure a pre-shared key on a remote access server
Make a virtual private network (VPN) connection

# Verify permissions for the RAS and IAS security group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify permissions for the RAS and IAS security group**

1. Open Active Directory Users and Computers.
2. On the **View** menu, confirm that **Advanced Features** is selected.
3. In the console tree, right-click **RAS and IAS Servers Access Check**.

   **Where?**

   - Active Directory Users and Computers/*domain name*/System/RAS and IAS Servers Access Check

4. Click **Properties**.
5. On the **Security** tab, do one of the following:

   - In **Group or User names**, if **RAS and IAS Servers** is listed, click **RAS and IAS Servers (***DomainName***\RAS and IAS Servers)**, and under **Permissions for RAS and IAS Servers**, verify that the **Read**, **Write**, **Create All Child Objects**, and **Delete All Child Objects** check boxes are selected.
   - In **Group or User names**, if **RAS and IAS Servers** is not listed, add **RAS and IAS Servers**. Then, in **Group or user names**, click **RAS and IAS Servers**, and in **Permissions for RAS and IAS Servers (***DomainName***\RAS and IAS Servers)**, select the **Read**, **Write**, **Create All Child Objects**, and **Delete All Child Objects** check boxes.

**Note**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Remote Access Troubleshooting
Troubleshooting demand-dial routing

# Map a certificate to a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To map a certificate to a user account**

1. Open Active Directory Users and Computers.
2. On the **View** menu, select **Advanced Features**.
3. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers\domain node\Users

   Or, click the folder that contains the user account.
4. In the details pane, click the user account to which you want to map a certificate.
5. On the **Action** menu, click **Name Mappings**.
6. In the **Security Identity Mapping** dialog box, on the **X.509 Certificates** tab, click **Add**.
7. Type the name and path of the .cer file that contains the certificate you want to map to this user account, and then click **Open**.
8. Do one of the following:

| To | Do this |
|---|---|
| Map the certificate to one account (one-to-one mapping) | Confirm that both the **Use Issuer for alternate security identity** and the **Use Subject for alternate security identity** check boxes are selected. |
| Map any certificate that has the same subject to the user account, regardless of the issuer of the certificate (many-to-one mapping) | Clear the **Use Issuer for alternate security identity** check box, and confirm that the **Use Subject for alternate security identity** check box is selected. |
| Map any certificate that has the same issuer to the user account, regardless of the subject of the certificate (many-to-one mapping) | Clear the **Use Subject for alternate security identity** check box, and confirm that the **Use Issuer for alternate security identity** check box is selected. |

**Notes**

- To perform this procedure, you must be a member of the Account Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- The certificate you are mapping to a user account must be in Distinguished Encoding Rules (DER) or Base64 encoded binary format. For instructions on exporting an existing certificate to a .cer file, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Mapping certificates to user accounts
Export a certificate
Importing and exporting certificates
Map an account from a trusted non-Windows kerberos realm to a user account

# Manage Machine Certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage machine certificates**

- Install the certificate of the enterprise root CA as a trusted root CA certificate
- Configure automatic certificate allocation from an enterprise CA
- Submit a user certificate request via the Web
- Install a router (offline request) certificate

# Install the certificate of the enterprise root CA as a trusted root CA certificate

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install the certificate of the enterprise root CA as a trusted root CA certificate**

1. Open Active Directory Users and Computers.

2. In the console tree, double-click **Active Directory Users and Computers**, right-click the domain name in which your CA lives, and then click **Properties**.

3. On the **Group Policy** tab, click **Default Domain Policy**, and then click **Edit**.

4. In the console tree, right-click **Trusted Root Certification Authorities**, point to **All Tasks**, and then click **Import**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Trusted Root Certification Authorities

   The Certificate Manager Import wizard appears.

5. Click **Next**, and follow the instructions in the wizard to import the .crt file of the CA (located in the *systemroot*\System32\CertSrv\CertEnroll folder) into the Trusted Root Certification Authorities store.

6. Type the following at the command prompt:

   **gpupdate /Target:**{*Computer* | *User*} **/Force**

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
EAP
Using smart cards for remote access
Deploying certificate-based authentication for demand-dial routing
Business partner demand-dial connection
Branch office demand-dial connection

# Configure automatic certificate allocation from an enterprise CA

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure automatic certificate allocation from an enterprise CA**

1. Open Active Directory Users and Computers.
2. In the console tree, double-click **Active Directory Users and Computers**, right-click the domain name in which your CA lives, and then click **Properties**.
3. On the **Group Policy** tab, click **Default Domain Policy**, and then click **Edit**.
4. In the console tree, right-click **Automatic Certificate Request Settings**, point to **New**, and then click **Automatic Certificate Request**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Public Key Policies/Automatic Certificate Request Settings

5. When the Automatic Certificate Request wizard appears, click **Next**.
6. In **Certificate templates**, click **Computer**, and then click **Next**.

   Your enterprise root CA appears on the list.

7. Click the CA, click **Next**, and then click **Finish**.
8. To create a computer certificate for the CA computer, type the following at the command prompt:

   **gpupdate /target:***Computer*

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
EAP
Using smart cards for remote access

Deploying certificate-based authentication for demand-dial routing
Business partner demand-dial connection
Branch office demand-dial connection

# Submit a user certificate request via the Web to a Windows Server 2003 CA

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To submit a user certificate request via the Web to a Windows Server 2003 CA**

1. Open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the Web server running Windows Server 2003 where the certification authority (CA) you want to access is located.
3. Click **Request a certificate**.
4. On **Request a Certificate**, click **User Certificate**.
5. On the User Certificate - Identifying Information page, do one of the following:
   - Read the message "No further identifying information is required. To complete your certificate, press submit".
   - Enter your identifying information for the certificate request.
6. (Optional) Click **More Options** to specify the cryptographic service provider (CSP) and whether you want to enable strong private key protection. (This means that you will receive a prompt every time that the private key associated with the certificate is used.)
7. Click **Submit**.
8. Do one of the following:
   - If you see the **Certificate Pending** Web page, see Related Topics below for the procedure to check on a pending certificate.
   - If you see the **Certificate Issued** Web page, click **Install this certificate**.
9. If you are finished using the Certificate Services Web pages, close Internet Explorer.

### Note

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Windows Server 2003 Certificate Services Web pages
Check on a pending certificate request to a Windows Server 2003 CA
Submit an advanced certificate request via the Web to a Windows Server 2003 CA

Request a certificate from a Windows Server 2003 CA using a PKCS #10 or PKCS #7 file
Use Windows 2000 Certificate Services Web Pages

# Install a router (offline request) certificate

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a router (offline request) certificate**

1. Open Internet Explorer.
2. In Internet Explorer, in **Address**, type the address of the certification authority (CA) that issues computer certificates. The address is the name of the server followed by **/certsrv**.
3. On the **Welcome** page, click **Request a certificate**, click **Advanced certificate request**, and then click **Create and submit a request to this CA**.
4. In **Certificate Template**, select **Router (Offline request)**, or the name of the template that the CA administrator directed you to choose.
5. In **Name**, type the user account name that is used by the calling router.
6. Under **Key Options**, select the **Mark keys as exportable** and **Store certificate in the local computer certificate store** check boxes.
7. Confirm the other options that you want, and then click **Submit**.
8. A message appears that asks you to confirm that you trust this Web site and that you want to request a certificate. Click **Yes**, and then, on the **Certificate Issued** page, click **Install this certificate**.

   A message informs you that a new certificate has been successfully installed.

## Notes

- To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.
- Windows 2000 CAs included a template, called **Router (Offline request)**, which allowed exportable keys. In the Windows Server 2003 family, CAs require you to create a special template to allow keys to be exported.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
EAP
Using smart cards for remote access
Deploying certificate-based authentication for demand-dial routing
Business partner demand-dial connection
Branch office demand-dial connection

# Configure dial-in user properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure dial-in user properties**

- Configure remote access permission for a user
- Configure caller ID and callback
- Configure static IP address assignment
- Configure static routes for a dial-in user

# Configure remote access permission for a user

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure remote access permission for a user**

1. Do one of the following:

   - If the remote access server is part of a Windows Server 2003 domain:

     | |
     |---|
     | Open Active Directory Users and Computers. |
     | In the console tree, click **Users** (console tree location is: Active Directory Users and Computers/*domain name*/Users). |
     | In the details pane, right-click a user name, and then click **Properties**. |

   - If the remote access server is a stand-alone server (not part of a domain):

     | |
     |---|
     | Open Computer Management. |
     | In the console tree, click **Users** (console tree location is: Computer Management/System Tools/Local Users and Groups/Users). |
     | In the details pane, right-click a user name, and then click **Properties**. |

2. On the **Dial-in** tab, under **Remote Access Permission (Dial-in or VPN)**, click either **Allow access**, **Control access through Remote Access Policy**, or **Deny access**, and then click **OK**.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Enable the IAS server to read user accounts in Active Directory
Guest authentication
Configure caller ID and callback

# Configure caller ID and callback

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure caller ID and callback**

1. Do one of the following:

| | |
|---|---|
| If the remote access server is part of a Windows Server 2003 domain: | a. Open Active Directory Users and Computers.<br>b. In the console tree, click **Users**.<br>• ***Where?***<br>Active Directory Users and Computers/*domain name*/Users<br>c. In the details pane, right-click a user name, and then click **Properties**. |
| If the remote access server is a stand-alone server (not a part of a domain): | a. Open Computer Management.<br>b. In the console tree, click **Users**.<br>• ***Where?***<br>Computer Management/System Tools/Local Users and Groups/Users<br>c. In the details pane, right-click a user name, and then click **Properties**. |

2. On the **Dial-in** tab, do the following:

   - For caller ID, select the **Verify Caller-ID** check box, and then type the number that the user is calling from.
   - For callback, under **Callback Options**, click the callback option you want to set for this user.

**Notes**

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- If the remote access server is part of a Windows NT 4.0 domain, you can use the Windows NT 4.0 User Manager for Domains administrative tool to set callback options for user accounts.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Enable the IAS server to read user accounts in Active Directory
Caller ID and callback
Remote Access Troubleshooting

# Configure static IP address assignment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a static IP address assignment**
**If the remote access server is part of a Windows Server 2003 domain, do the following:**

1. Open Active Directory Users and Computers.
2. In the console tree, click **Users**.

   **Where?**

   - Active Directory Users and Computers/*domain name*/Users

3. In the details pane, right-click a user name, and then click **Properties**.
4. On the **Dial-in** tab, select the **Assign a Static IP Address** check box, and then type the static IP address for this user.

**If the remote access server is a stand-alone server (not part of a domain), do the following:**

1. Open Computer Management.
2. In the console tree, click **Users**.

   **Where?**

   - Computer Management/System Tools/Local Users and Groups/Users

3. In the details pane, right-click a user name, and then click **Properties**.
4. On the **Dial-in** tab, select the **Assign a Static IP Address** check box, and then type the static IP address for this user.

### Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- The server running Routing and Remote Access will provide a dynamic IP address to a user if the static address you specify is not valid or is already in use by another user.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files

Enable the IAS server to read user accounts in Active Directory
Configure remote access permission for a user

# Configure static routes for a dial-in user

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure static routes for a dial-in user**

1. Do one of the following:

   - If the remote access server is part of a Windows Server 2003 domain:

     | |
     |---|
     | Open Active Directory Users and Computers. |
     | In the console tree, click **Users** (console tree location is: Active Directory Users and Computers/*domain name*/Users). |
     | In the details pane, right-click a user name, and then click **Properties**. |

   - If the remote access server is a stand-alone server (not a part of a domain):

     | |
     |---|
     | Open Computer Management. |
     | In the console tree, click **Users** (console tree location is: Computer Management/System Tools/Local Users and Groups/Users). |
     | In the details pane, right-click a user name, and then click **Properties**. |

2. On the **Dial-in** tab, select the **Apply Static Routes** check box.

3. In the **Static Routes** dialog box, click **Add Route**.

4. In **Destination**, **Network Mask**, and **Metric**, type the destination network address, the network mask, and the number of hops to the destination network for this static route, and then click **OK**.

5. Repeat steps 3 and 4 for each static route you want to configure for this user.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- Metrics greater than 99 might not function.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files

Enable the IAS server to read user accounts in Active Directory
One-way initiated demand-dial connections
Troubleshooting demand-dial routing

# Administer the Remote Access Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer the remote access server**

- Add a remote access server
- Monitor the Routing and Remote Access service
- Start and stop the Routing and Remote Access service
- Administer remote access servers running Windows NT 4.0
- Manage Remote Access Clients

# Add a remote access server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a remote access server**

1. Open Routing and Remote Access.
2. In the console tree, double-click **Routing and Remote Access**, right-click **Server Status**, and then click **Add Server**.
3. In the **Add Server** dialog box, do one of the following:
   - To add the computer on which Routing and Remote Access is currently running, click **This computer**.
   - To add a computer on another network, click **The following computer**, and then type the computer name or IP address of the server.
   - To add a server from a particular domain, click **All routing and remote access servers in the domain**, type the domain containing the server you want to administer, and then click **OK**.
   - To add a server, click **Browse the Active Directory**, and then click **Next**. In the **Find Routers or Remote Access Servers** dialog box, click the types of servers you want to search for, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Enable the Routing and Remote Access service

# Monitor the Routing and Remote Access service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To monitor the Routing and Remote Access service**

1. Open Routing and Remote Access.

2. In the console tree, double-click **Routing and Remote Access**, and then click **Server Status**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- The Routing and Remote Access window shows the server name, the state of the Routing and Remote Access service (started or stopped), the total ports on the server, the ports in use, and the amount of time that the server has been up since the Routing and Remote Access service was last started.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Start and stop the Routing and Remote Access service
Remote Access Troubleshooting

# Start and stop the Routing and Remote Access service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start and stop the Routing and Remote Access service**

1. Open Routing and Remote Access.
2. In the console tree, click **Server Status**.
3. In the details pane, right-click a server name, point to **All Tasks**, and do one of the following:
   - To start the service, click **Start**.
   - To stop the service, click **Stop**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Monitor the Routing and Remote Access service
Remote Access Troubleshooting

# Administer remote access servers running Windows NT 4.0

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To administer remote access servers running Windows NT 4.0**

1. On your server running Windows NT 4.0, click **Start**, click **Run**, type **rasadmin**, and then click **OK**.
2. On the **Server** menu, click **Select Domain or Server**.
3. In **Domain**, type a domain name or server name.

   Or, in **Select Domain**, click a domain or server.

## Notes

- Remote Access Admin is a tool specific to Windows NT 4.0. It is not available on computers running a member of the Windows Server 2003 family.
- If you type a server name in **Domain**, you must precede the name with two backslashes (****).
- Domains in **Select Domain** are all those trusted by the domain to which you belong. To view the servers in a domain, double-click the domain name.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Administration tools in mixed environments

# Manage Remote Access Clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage remote access clients**

- View connected remote access clients
- Disconnect a remote access client
- Send a message to a single remote access client
- Send a message to all remote access clients

# View connected remote access clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view connected remote access clients**

1. Open Routing and Remote Access.
2. In the console tree, click **Remote Access Clients**.

   **Where?**

   - Routing and Remote Access/*server name*/Remote Access Clients

3. In the details pane, right-click a user name, and then click **Status**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Disconnect a remote access client
Send a message to a single remote access client
Send a message to all remote access clients

# Disconnect a remote access client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect a remote access client**

1. Open Routing and Remote Access.
2. In the console tree, click **Remote Access Clients**.

   **Where?**

   - Routing and Remote Access/*server name*/Remote Access Clients

3. In the details pane, right-click a user name, and then click **Disconnect**.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Send a message to a single remote access client
Send a message to all remote access clients
View connected remote access clients

# Send a message to a single remote access client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send a message to a single remote access client**

1. Open Routing and Remote Access.
2. In the console tree, click **Remote Access Clients**.

   **Where?**

   - Routing and Remote Access/*server name*/Remote Access Clients

3. In the details pane, right-click a user name, and then click **Send message**.
4. In **Send Message**, type your message, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Disconnect a remote access client
Send a message to all remote access clients
View connected remote access clients

# Send a message to all remote access clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send a message to all remote access clients**

1. Open Routing and Remote Access.
2. In the console tree, right-click **Remote Access Clients**.

   **Where?**

   - Routing and Remote Access/*server name*/Remote Access Clients

3. Click **Send to all**.
4. In **Send Message**, type your message, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Disconnect a remote access client
Send a message to a single remote access client
View connected remote access clients

# Remote Access Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Remote Access Overview
- Understanding Remote Access
- Using Remote Access
- Resources

# Remote Access Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access overview**
This section covers:

- Overview of remote access
- Security information for remote access
- The remote access server as a dial-up networking server
- The remote access server as a virtual private networking server
- New features of remote access
- Common server configurations for remote access servers

# Overview of remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Overview of remote access**

By configuring Routing and Remote Access to act as a remote access server, you can connect remote or mobile workers to organization networks. Remote users can work as if their computers are physically connected to the network.

Users run remote access software and initiate a connection to the remote access server. The remote access server, which is a server running Routing and Remote Access, authenticates users and services sessions until terminated by the user or network administrator. All services typically available to a LAN-connected user (including file and print sharing, Web server access, and messaging) are enabled by means of the remote access connection.

Remote access clients use standard tools to access resources. For example, on a server running Routing and Remote Access, clients can use Windows Explorer to make drive connections and to connect to printers. Connections are persistent: Users do not need to reconnect to network resources during their remote sessions. Because drive letters and universal naming convention (UNC) names are fully supported by remote access, most commercial and custom applications work without modification.

A server running Routing and Remote Access provides two different types of remote access connectivity:

1. Dial-up networking

   Dial-up networking is when a remote access client makes a nonpermanent, dial-up connection to a physical port on a remote access server by using the service of a telecommunications provider such as analog phone, ISDN, or X.25. The best example of dial-up networking is that of a dial-up networking client who dials the phone number of one of the ports of a remote access server.

   Dial-up networking over an analog phone or ISDN is a direct physical connection between the dial-up networking client and the dial-up networking server. You can encrypt data sent over the connection, but it is not required.

   For more information, see The remote access server as a dial-up networking server.

2. Virtual private networking

   Virtual private networking is the creation of secured, point-to-point connections across a private network or a public network such as the Internet. A virtual private networking client uses special TCP/IP-based protocols called tunneling protocols to make a virtual call to a virtual port on a virtual private networking server. The best example of virtual private networking is that of a virtual private networking client who makes a virtual private network connection to a remote access server that is connected to the Internet. The remote access server answers the virtual call, authenticates the caller, and transfers data between the virtual private networking client and the corporate network.

   In contrast to dial-up networking, virtual private networking is always a logical, indirect connection between the virtual private networking client and the virtual private networking server over a public network such as the Internet. To ensure privacy, you must encrypt data sent over the connection.

   For more information, see The remote access server as a virtual private networking server.

# Security information for remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for remote access**

Before you enable the remote access functions of Routing and Remote Access, you should carefully review your network infrastructure so that you can configure Routing and Remote Access to best meet your needs for security and functionality. Security for Routing and Remote Access can be considered in three parts: securing the server running Routing and Remote Access, securing the network traffic between the server and its clients, and using secure authentication methods.

Unless your server running Routing and Remote Access is functioning only as a remote access server, you should also consider security for your routing implementation. For more information, see Security information for routing.

In addition to understanding the security implications of enabling Routing and Remote Access on your network, you should also practice general server security, as described in Security. For example, you should use only servers that are running Microsoft® Windows® 2000 or a member of the Microsoft Windows Server 2003 family, if possible. You will decrease your network security if you use servers that are running Routing and Remote Access on the same network as servers that are running Microsoft Windows NT 4.0 and either Remote Access Service (RAS) or Routing and Remote Access Service (RRAS). For more information, see Domain and forest functionality and Member server in a domain.

Before you configure and enable the remote access features of Routing and Remote Access, you should consider the following:

- **Who can enable, configure, and disable Routing and Remote Access**. You must be a member of the Administrators group to use the Routing and Remote Access snap-in, either on its own or as part of the Microsoft Management Console (MMC). Similarly, you must be a member of the Administrators group to run most of the **netsh ras** commands from the command line. You should restrict the membership of the Administrators group to the minimum number of users necessary to administer the server. For more information, see Default local groups and Default groups.

- **How to open Routing and Remote Access**. As a security best practice, consider using the **runas** command to open Routing and Remote Access rather than logging on with administrative credentials. For more information about why you might not want to log on with administrative credentials, see Why you should not run your computer as an administrator. To open Routing and Remote Access when you are not logged on as a member of the Administrators group, type the following at the command prompt:

  **runas /user:**[*Domain/*]*UserName***"mmc %windir%\system32\rrasmgmt.msc"**

  The user name that you provide must correspond to the Administrator account or an account that is a member of the Administrators group, and you must provide the password for the account when prompted. For more information, see Runas.

- **What client operating systems to support for remote access**. You should consider requiring your remote users to run particular operating systems (such as Microsoft Windows XP, Windows 2000, and members of the Windows Server 2003 family). There are two reasons to consider requiring your users to run particular operating systems. First, not all operating systems have equal levels of security in their file systems and in their user accounting. Second, not all remote access features are available on all operating systems. For more information, see Dial-up networking clients and Virtual private networking clients. Based on what client operating systems you support, you should also consider:

  - **What authentication methods to support**. You should consider requiring your remote users to be authenticated with more secure authentication protocols, such as Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) or Extensible Authentication Protocol (EAP), rather than allowing your remote users to use protocols such as Password Authentication Protocol (PAP), Shiva Password Authentication Protocol (SPAP), and Challenge Handshake Authentication Protocol (CHAP). For more information, see Remote Access Authentication Methods.

  - **What level of data encryption to support**. You should consider requiring your remote users to encrypt data with the highest level of data encryption possible. Windows XP and the Windows Server 2003 family support 128-bit encryption keys. For more information, see Remote Access Data encryption.

- **Whether to require strong passwords**. You should consider requiring your remote users to provide passwords that contain a certain number of characters and that contain a combination of uppercase letters, lowercase letters, numerals, and special characters. One of the worst vulnerabilities for a remote access solution is the use of weak or easily guessable passwords by remote users. For more information, see Strong passwords and Passwords must meet complexity requirements.

- **Whether to provide users with a managed client**. You should consider increasing the security of your remote access solution by using the Connection Manager Administration Kit (CMAK) wizard to create a customized connection for your users. By customizing the connection, you can control how your users connect to your network, and you can also simplify troubleshooting. For more information, see Connection Manager Administration Kit.

- **How to document your remote access solution**. You should consider documenting your remote access needs and how your network is configured to meet them. This can help you maintain your network and identify potential areas for concern or improvement. For more information, see Dial-up remote access design considerations and Remote access VPN design considerations.

- **Whether to use Windows Authentication or RADIUS**. If you have more than one server running Routing and Remote Access, you should consider using Remote Authentication Dial-In User Service (RADIUS) instead of Windows Authentication. By using RADIUS, you can centralize authentication, authorization, and auditing of remote access connections, and you can more easily manage your remote access policies. If you use RADIUS, you should consider encrypting your RADIUS traffic with Internet Protocol security (IPSec). For more information, see Remote Access Authentication, Using RADIUS, and Internet Authentication Service.

- **Whether to restrict remote access**. You should consider restricting remote access to your network to certain users, user groups, times, or client configurations by configuring remote access policies. You can implement remote access policies that restrict remote access to only those users or groups who require it and who meet the security standard for your network. For more information, see Remote Access Policies.

# The remote access server as a dial-up networking server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The remote access server as a dial-up networking server**

Routing and Remote Access provides traditional dial-up remote access to support mobile users or home users who are dialing in to organization intranets. Dial-up equipment that is installed on the server running Routing and Remote Access answers incoming connection requests from dial-up networking clients. The remote access server answers the call, authenticates and authorizes the caller, and transfers data between the dial-up networking client and the organization intranet.

The following illustration shows dial-up networking functionality.




For more information, see Dial-Up Networking.

# The remote access server as a virtual private networking server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The remote access server as a virtual private networking server**

A virtual private networking connection emulates a point-to-point connection. To emulate a point-to-point connection, data is encapsulated, or wrapped, with an additional IP header that provides routing information to reach a virtual private networking server. The portion of the virtual private networking connection in which the data is encapsulated is called the tunnel.

For secure virtual private networking, data is encrypted before it is encapsulated. Intercepted packets are unintelligible without the encryption keys. The portion of the virtual private networking connection in which your data is encrypted is called the virtual private network (VPN) connection.

VPN connections are created, managed, and terminated by using special protocols called tunneling protocols. Both the virtual private networking client and the virtual private networking server must support the same tunneling protocol to create a virtual private networking connection. A server running Routing and Remote Access is a virtual private networking server for both the Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) tunneling protocols. For more information, see VPN Tunneling Protocols.

The following illustration shows virtual private networking functionality.




For more information, see Virtual Private Networking.

**Note**

- On Microsoft® Windows Server™ 2003, Web Edition, and Microsoft Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# New features of remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features of remote access**

Routing and Remote Access provides the following new features:

- Support for pre-shared keys in L2TP/IPSec authentication

  You can configure a server running Routing and Remote Access to use a pre-shared key instead of a certificate for L2TP/IPSec VPN connections. For more information, see Pre-shared keys.

- Integration of network address translation (NAT) and static and dynamic packet filtering

  NAT interfaces can be configured to work with Basic Firewall, a simple dynamic packet filter, or with incoming or outgoing static packet filters. For more information, see Basic Firewall.

- Support for L2TP/IPSec connections over NAT

  Routing and Remote Access seamlessly translates L2TP/IPSec VPN connections over a NAT interface. For more information, see Virtual Private Networks.

- Broadcast name resolution for small networks without a local WINS or DNS server

  Small networks or branch offices without local WINS or DNS servers can resolve corporate resource network names. For more information, see Broadcast name resolution.

# Common server configurations for remote access servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Common configurations for remote access servers**

When you run the Routing and Remote Access Server Setup Wizard, the wizard prompts you to choose the configuration path that most closely resembles the remote access solution that you want to deploy. If none of the wizard configuration paths meets your needs exactly, you can further configure your server after the wizard finishes, or you can choose the custom configuration path. However, if you choose the custom configuration path, you must manually configure all elements of Routing and Remote Access. The most common remote access solutions include dial-up connections, virtual private network (VPN) connections, and secure connections between two private networks.

**Remote access (dial-up)**




If you choose this path, the server running Routing and Remote Access is configured to allow remote access clients to connect to the private network by dialing into a modem bank or other dial-up equipment. To configure this type of server in the wizard, click **Remote Access**, select the **Dial-up** check box, and follow the steps. After the wizard finishes, you can configure additional options. For example, you can configure how the server answers the call, how the server verifies which remote access clients have permission to connect to the private network, and whether the server routes network traffic between remote access clients and the private network. For more information about dial-up connections, see Dial-Up Networking and Setting Up Dial-Up Remote Access.

**Remote access (VPN)**




If you choose this path, the server running Routing and Remote Access is configured to allow remote access clients to connect to the private network across the Internet. To configure this type of server in the wizard, click **Remote Access**, select the **VPN** check box, and follow the steps. After the wizard finishes, you can configure additional options. For example, you can configure how the server verifies which VPN clients have permission to connect to the private network and whether the server routes network traffic between VPN clients and the private network. For more information about VPNs, see Deploying VPNs for Remote Access and Common configuration for the VPN server.

**Network address translation (NAT)**




If you choose this path, the server running Routing and Remote Access is configured to share an Internet connection with

computers on the private network and to translate traffic between its public address and the private network. Computers on the Internet will not be able to determine the IP addresses of computers on the private network. To configure this type of server in the wizard, click **Network address translation (NAT)**, and follow the steps. After the wizard finishes, you can configure additional options. For example, you can configure packet filters and choose what services to allow on the public interface. For more information about NAT, see Setting Up Network Address Translation and Translated Connection to the Internet.

**VPN and NAT**




If you choose this path, the server running Routing and Remote Access is configured to provide NAT for the private network and to accept VPN connections. Computers on the Internet will not be able to determine the IP addresses of computers on the private network. However, VPN clients will be able to connect to computers on the private network as if they were physically attached to the same network. To configure this type of server in the wizard, click **Virtual Private Network (VPN) access and NAT**, and follow the steps. For more information about NAT and VPN, see New features for virtual private networks and Network address translation design considerations.

**Secure connection between two private networks**




If you choose this path, two servers running Routing and Remote Access are configured to send private data securely across the Internet. You must choose this path when you run the Routing and Remote Access Server Setup Wizard on each server. The connection between the two servers can be persistent (always on) or on demand (demand-dial). To configure this type of server in the wizard, click **Secure connection between two private networks**, and follow the steps. After the wizard finishes, you can configure each server with additional options. For example, you can configure what routing protocols each server accepts and how each server routes traffic between the two networks. For more information about connections between two networks, see Deploying Router-to-Router VPNs, Branch Office over the Internet, and Dial-Up Branch Office Network.

# Understanding Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding remote access**
This section covers:

- Dial-Up Networking
- Virtual Private Networking
- Devices and ports
- VPN Tunneling Protocols
- Remote Access Security
- Remote Access Policies
- Multilink and BAP
- Logging
- The remote access server and network management
- Broadcast name resolution

# Dial-Up Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up networking**
This section covers:

- Components of dial-up networking
- Dial-up networking clients
- Dial-up networking servers

# Components of dial-up networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Components of dial-up networking**

Dial-up networking includes the following components:

- **Dial-up networking servers**

  You can configure a server running Routing and Remote Access to provide dial-up networking access to an entire network or restrict access to the shared resources of the remote access server only.

- **Dial-up networking clients**

  Remote access clients running a member of the Windows Server 2003 family, Windows XP, Windows 2000, Windows NT running Remote Access Service (RAS) or Routing and Remote Access Service (RRAS), Windows Millennium Edition, Windows 98, Windows 95, or Mac OS can all connect to a server running Routing and Remote Access.

- **LAN and remote access protocols**

  Application programs use LAN protocols to transport information. Remote access protocols are used to negotiate connections and provide framing for LAN protocol data that is sent over wide area network (WAN) links. Routing and Remote Access supports LAN protocols such as TCP/IP and AppleTalk, which enable access to Internet, UNIX, Mac OS, and Novell NetWare resources. Routing and Remote Access supports remote access protocols such as PPP.

- **WAN options**

  Clients can dial in by using standard telephone lines and a modem or modem pool. Faster links are possible by using ISDN. You can also connect remote access clients to remote access servers by using X.25 or ATM. Direct connections are also supported through an RS-232C null modem cable, a parallel port connection, or an infrared connection.

- **Security options**

  The Windows Server 2003 family provides logon and domain security, support for security hosts, data encryption, Remote Authentication Dial-In User Service (RADIUS), smart cards, remote access account lockout, remote access policies, and callback for secure network access for dial-up clients.

The following illustration shows the dial-up networking components. Your actual implementation and configuration of dial-up networking may vary.



# Dial-up networking clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up networking clients**

A dial-up client that connects to a server running Routing and Remote Access can be any PPP client. The client must have installed a modem, an analog telephone line or other WAN connection, and remote access software.

You can automate the connection process for clients by using a simple batch file and the **rasdial** command. You can also schedule automatic backups to or from remote computers by using dial-up networking and Task Scheduler.

**Windows-based PPP clients**

Windows-based PPP clients that use TCP/IP can access a server running Routing and Remote Access. Windows-based PPP clients cannot use the AppleTalk protocol. The remote access server automatically negotiates authentication with PPP clients.

Remote access support for Windows-based PPP dial-up networking clients is outlined in the following table.

| Dial-up networking client | Supported remote access PPP features | Unsupported remote access PPP features |
|---|---|---|
| Microsoft® Windows Server 2003 family, Windows XP, Windows 2000 | Multilink, Bandwidth Allocation Protocol (BAP), Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Challenge Handshake Authentication Protocol (CHAP), Shiva Password Authentication Protocol (SPAP), Password Authentication Protocol (PAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2), and Extensible Authentication Protocol (EAP) | |
| Windows NT® version 4.0 | Multilink, MS-CHAP, CHAP, SPAP, PAP, and MS-CHAP v2 (with Windows NT 4.0 Service Pack 4 and later) | BAP and EAP |
| Windows NT version 3.5*x* | MS-CHAP, CHAP, SPAP, and PAP | Multilink, BAP, MS-CHAP v2, and EAP |
| Windows Millennium Edition, Windows 98 | Multilink, MS-CHAP, CHAP, SPAP, PAP, and MS-CHAP v2 (with the Windows 98 Service Pack 1 and later) | BAP and EAP |
| Windows 95 | MS-CHAP, CHAP, SPAP, and PAP (with the Windows Dial-Up Networking 1.3 Performance & Security Upgrade for Windows 95 and later) | MS-CHAP v2, Multilink, BAP, and EAP |

**Note**

- Windows 95 with the Windows Dial-Up Networking 1.3 Performance & Security Upgrade for Windows 95 supports MS-CHAP v2 over virtual private network (VPN) connections but not over dial-up connections. For more information about virtual private networks and networking upgrades, see the Microsoft Web site.

**Other PPP clients**

Other PPP client software that uses TCP/IP or AppleTalk can access a server running Routing and Remote Access. The remote access server automatically negotiates authentication with PPP clients. No special configuration of the server running Routing and Remote Access is required for other PPP clients except to ensure that both the remote access server and the other PPP client are configured for the same LAN and authentication protocols.

For more information, see the documentation provided with your PPP client.

# Dial-up networking servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up networking servers**

Administrators can use Routing and Remote Access to configure a server to provide remote access, view connected users, and monitor remote access traffic.

For dial-up networking access, the server must have at least one modem or a multiport adapter, and analog telephone lines or other WAN connections. If the server provides access to a network, you must install and connect a separate network adapter to the network to which the server provides access.

You can configure a server running Routing and Remote Access both during and after you run the Routing and Remote Access Server Setup wizard.

For more information about installing the Routing and Remote Access service, see Enable the Routing and Remote Access service.

For more information about deploying a dial-up networking remote access server, see Setting Up Dial-Up Remote Access.

# Virtual Private Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private networking**
This section covers:

- Components of virtual private networking
- Virtual private networking clients
- Virtual private networking servers

# Components of virtual private networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Components of virtual private networking**

A virtual private network can include the following components:

- **Virtual private network (VPN) servers**

  You can configure your computer to provide access to an entire network or to restrict access to just the resources of the computer acting as the VPN server.

- **VPN clients**

  VPN clients are either individual users who obtain a remote access VPN connection or routers that obtain a router-to-router VPN connection. VPN clients running a member of the Windows server 2003 family, Windows XP, Windows 2000, Windows NT 4.0, Windows 95, Windows 98, or Windows Millennium Edition can create remote access VPN connections to the VPN server. Computers running either a member of the Windows server 2003 family or Windows 2000 and Routing and Remote Access, or Windows NT server 4.0 and the Routing and Remote Access Service (RRAS), can create router-to-router VPN connections. VPN clients can also be any Point-to-Point Tunneling Protocol (PPTP) client or Layer Two Tunneling Protocol (L2TP) client that uses Internet Protocol security (IPSec).

- **LAN and remote access protocols**

  LAN protocols are used by application programs to transport information. Remote access protocols are used to negotiate connections and provide framing for LAN protocol data that is sent over wide area network (WAN) links. Routing and Remote Access supports the PPP remote access protocol. Windows Server 2003, Datacenter Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Web Edition; and Windows Server 2003, Standard Edition all support LAN protocols such as TCP/IP and AppleTalk, which enable access to Internet, UNIX, Apple Macintosh, and Novell NetWare resources.

- **Tunneling protocols**

  VPN clients use tunneling protocols to create secured connections to a VPN server by using the PPTP or L2TP tunneling protocols.

- **WAN options**

  VPN servers are connected to the Internet by using permanent WAN connections such as T1 or Frame Relay. VPN clients are connected to the Internet by using permanent WAN connections or by dialing in (by using standard analog telephone lines or ISDN) to a local Internet service provider (ISP).

- **Security options**

  Routing and Remote Access supports logon and domain security, support for security hosts, data encryption, smart cards, IP packet filtering, and caller ID to provide secure network access for VPN clients.

The following illustration shows all the virtual private networking components and possible configurations. Your actual implementation and configuration of virtual private networking may vary.




**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Virtual private networking clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private networking clients**

Virtual private networking clients that connect to Routing and Remote Access can be computers running a member of the Windows Server 2003 family, Windows XP, Windows 2000, Windows NT 4.0, Windows 95, Windows 98, or Windows Millennium Edition. The client must be able to send TCP/IP packets to the remote access server over the Internet. Therefore, either a network adapter or modem with an analog telephone line or other WAN connection to the Internet is required.

**Tunneling protocols for Microsoft virtual private network clients**

Support for tunneling protocols by Microsoft virtual private networking clients is outlined in the following table.

| Virtual private networking client | Supported tunneling protocols | Unsupported tunneling protocols |
|---|---|---|
| Windows Server 2003 family, Windows 2000, and Windows XP | Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) | |
| Windows NT version 4.0 | PPTP | L2TP |
| Windows Millennium Edition and Windows 98 | PPTP | L2TP |
| Windows 95 with the Windows Dial-Up Networking 1.3 Performance & Security Upgrade for Windows 95 | PPTP | L2TP |
| Windows Millennium Edition, Windows 98, Windows NT 4.0 Workstation, and Windows 95 with Microsoft L2TP/IPSec VPN Client installed | L2TP | |

**Notes**

- Windows NT version 3.5*x* does not support either PPTP or L2TP.
- For more information about virtual private networks, Microsoft L2TP/IPSec VPN Client, and networking upgrades, see the Microsoft Web site.

**Authentication for Microsoft virtual private network clients**

Support for authentication protocols by Microsoft virtual private networking clients is outlined in the following table.

| Virtual private networking client | Supported remote access authentication protocols | Unsupported remote access authentication protocols |
|---|---|---|
| Windows Server 2003 family, Windows XP, and Windows 2000 | Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Challenge Handshake Authentication Protocol (CHAP), Shiva Password Authentication Protocol (SPAP), Password Authentication Protocol (PAP), MS-CHAP version 2 (MS-CHAP v2), and Extensible Authentication Protocol (EAP) | |
| Windows NT version 4.0 | MS-CHAP, CHAP, SPAP, PAP, and MS-CHAP v2 (with the Windows NT 4.0 Service Pack 4 and later) | EAP |
| Windows 98 | MS-CHAP, CHAP, SPAP, PAP, and MS-CHAP v2 (with the Windows 98 Service Pack 1 and later) | EAP |

| | | |
|---|---|---|
| Windows 95 | MS-CHAP, CHAP, SPAP, PAP, and MS-CHAP v2 (with the Dial-Up Networking 1.4 Upgrade) | EAP |

**Note**

- Windows 95 with the Dial-Up Networking 1.4 Upgrade supports MS-CHAP v2 over virtual private network (VPN) connections but not over dial-up connections. For more information about virtual private networks and networking upgrades, see the Microsoft Web site .

**Other virtual private network (VPN) clients**

Other virtual private network clients that use PPTP or L2TP with Internet Protocol security (IPSec) can access a server running Routing and Remote Access. However, if you want secure VPN connections, you must make sure that clients support the proper encryption. For PPTP, Microsoft Point-to-Point Encryption (MPPE) must be supported. For L2TP, IPSec encryption must be supported.

For more information, see the documentation provided with the virtual private network client.

# Virtual private networking servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private networking servers**

You can use Routing and Remote Access to configure a VPN server, view connected users, and monitor remote access traffic.

For virtual private network access from the Internet, the server typically has a permanent connection to the Internet. A nonpermanent connection to the Internet is possible if the Internet service provider (ISP) supports demand-dial connections; the connection is created when traffic is delivered to the VPN server. However, this is not a common configuration. If the VPN server provides access to a network, you must install and connect a separate network adapter to the network to which the VPN server provides access.

On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

You can configure remote access servers after the Routing and Remote Access Server Setup Wizard is run.

For more information about installing the Routing and Remote Access service, see Enable the Routing and Remote Access service.

For more information about deploying a remote access VPN server, see Setting Up Remote Access VPNs.

# Devices and ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Devices and ports**
A server running Routing and Remote Access views the installed networking equipment as a series of devices and ports.

- A device represents hardware or software that can create physical or logical point-to-point connections.
- A port is a communication channel that can support a single point-to-point connection.

**Device**
A device is the hardware or software that provides ports that remote access connections can use to establish point-to-point connections. Devices are physical, such as a modem, or virtual, such as virtual private network (VPN) protocols. Devices can support a single port, such as a modem, or multiple ports, such as modem bank hardware that can terminate 64 different incoming analog phone calls. Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) are examples of virtual multiport devices. Each of these tunneling protocols supports multiple VPN connections.

To see the installed devices, you can view the properties of **Ports** in Routing and Remote Access.

**Port**
A port is a channel of a device that can support a single point-to-point connection. For single-port devices such as modems, the device and the port are indistinguishable. For multiport devices, the port is the subdivision of the device over which a separate point-to-point connection is possible. For example, Primary Rate Interface (PRI) ISDN adapters support two separate channels called B channels. The ISDN adapter is a device. Each B channel is a port because a separate point-to-point connection can occur over each B channel.

You can view the dial-up ports by clicking **Ports** in Routing and Remote Access.

### Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# VPN Tunneling Protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN tunneling protocols**
This section covers:

- Point-to-Point Tunneling Protocol
- Layer Two Tunneling Protocol

# Point-to-Point Tunneling Protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Point-to-Point Tunneling Protocol**
Point-to-Point Tunneling Protocol (PPTP) is a tunneling protocol first supported in Windows NT 4.0 and Windows 98. PPTP is an extension of Point-to-Point Protocol (PPP) and leverages the authentication, compression, and encryption mechanisms of PPP. Client support for PPTP is built-in to the Windows XP remote access client.

VPN server support for PPTP is built in to members of the Windows Server 2003 family. PPTP is installed with the TCP/IP protocol. Depending on your choices when running the Routing and Remote Access Server Setup Wizard, PPTP is configured for five or 128 PPTP ports. For more information, see Checklist: Installing and configuring a PPTP server.

PPTP and Microsoft Point-to-Point Encryption (MPPE) provide the primary VPN services of encapsulation and encryption of private data.

**Encapsulation**
A PPP frame (an IP, IPX, or Appletalk datagram) is wrapped with a Generic Routing Encapsulation (GRE) header and an IP header. In the IP header is the source and destination IP address that correspond to the VPN client and VPN server.

The following illustration shows PPTP encapsulation for a PPP frame.




**Encryption**
The PPP frame is encrypted with Microsoft Point-to-Point Encryption (MPPE) by using encryption keys generated from the MS-CHAP, MS-CHAP v2, or EAP-TLS authentication process. Virtual private networking clients must use the MS-CHAP, MS-CHAP v2, or EAP-TLS authentication protocol in order for the payloads of PPP frames to be encrypted. PPTP is taking advantage of the underlying PPP encryption and encapsulating a previously encrypted PPP frame.

For more information on deploying PPTP-based VPN connections, see Deploying Virtual Private Networks. For a PPTP-based VPN server checklist, see Checklist: Installing and configuring a PPTP server. For an example implementation of PPTP-based VPN connections, see Virtual Private Network Implementation Examples. For information on configuring PPTP-based VPN connections in a test lab, see Virtual Private Network Test Lab.

## Notes

- It is possible to have a nonencrypted PPTP connection where the PPP frame is sent in plaintext. However, a nonencrypted PPTP connection is not recommended for VPN connections over the Internet because communications of this type are not secure.

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Layer Two Tunneling Protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Layer Two Tunneling Protocol**
Layer Two Tunneling Protocol (L2TP) is an RFC-based tunneling protocol that is an industry standard and was first supported in the Windows 2000 client and server operating systems. Unlike PPTP, L2TP in servers running Windows Server 2003 does not utilize Microsoft Point-to-Point Encryption (MPPE) to encrypt Point-to-Point Protocol (PPP) datagrams. L2TP relies on Internet Protocol security (IPSec) for encryption services. The combination of L2TP and IPSec is known as L2TP/IPSec. L2TP/IPSec provides the primary virtual private network (VPN) services of encapsulation and encryption of private data.

Both L2TP and IPSec must be supported by both the VPN client and the VPN server. Client support for L2TP is built-in to the Windows XP remote access client, and VPN server support for L2TP is built in to members of the Windows Server 2003 family.

L2TP is installed with the TCP/IP protocol. Depending on your choices when running the Routing and Remote Access Server Setup Wizard, L2TP is configured for five or 128 L2TP ports. For more information, see Checklist: Installing and configuring an L2TP server.

For more information about IPSec, see Internet Protocol Security Overview.

**Encapsulation**
Encapsulation for L2TP/IPSec packets consists of two layers:

1. L2TP encapsulation

   A PPP frame (an IP datagram or an IPX datagram) is wrapped with an L2TP header and a UDP header.

2. IPSec encapsulation

   The resulting L2TP message is then wrapped with an IPSec Encapsulating Security Payload (ESP) header and trailer, an IPSec Authentication trailer that provides message integrity and authentication, and a final IP header. In the IP header is the source and destination IP address that corresponds to the VPN client and VPN server.

The following illustration shows L2TP and IPSec encapsulation for a PPP datagram.




**Encryption**
The L2TP message is encrypted with either Data Encryption Standard (DES) or Triple DES (3DES) by using encryption keys generated from the Internet Key Exchange (IKE) negotiation process.

For more information on deploying L2TP/IPSec VPN connections, see Deploying Virtual Private Networks. For configuration checklists, see VPN Checklists. For an example implementation of L2TP/IPSec VPN connections, see Virtual Private Network Implementation Examples. For information on configuring L2TP/IPSec VPN connections in a test lab, see Virtual Private Network Test Lab.

### Notes

- It is possible to have a nonencrypted L2TP connection where the PPP frame is sent in plaintext. However, a nonencrypted L2TP connection is not recommended for virtual private network connections over the Internet because communications of this type do not provide data confidentiality.

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Remote Access Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access security**
This section covers:

- Remote Access Authentication Methods
- Authentication vs. authorization
- Pre-shared keys
- Remote Access Data encryption
- Remote access dial-in permissions
- Caller ID and callback
- Remote Access Security hosts
- Remote access account lockout

# Remote Access Authentication Methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication methods**
The authentication of remote access clients is an important security concern. Authentication methods typically use an authentication protocol that is negotiated during the connection establishment process. The Windows Server 2003 family also supports unauthenticated access.

This section covers:

- EAP
- MS-CHAP
- MS-CHAP version 2
- CHAP
- SPAP
- PAP
- Unauthenticated access

# EAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**EAP**

With the Extensible Authentication Protocol (EAP), an arbitrary authentication mechanism authenticates a remote access connection. The exact authentication scheme to be used is negotiated by the remote access client and the authenticator (either the remote access server or the Remote Authentication Dial-In User Service [RADIUS] server). Routing and Remote Access includes support for EAP-TLS and MD5-Challenge by default. You can plug in other EAP modules to the server running Routing and Remote Access to provide other EAP methods.

EAP allows for an open-ended conversation between the remote access client and the authenticator. The conversation consists of authenticator requests for authentication information and the responses by the remote access client. For example, when EAP is used with security token cards, the authenticator can separately query the remote access client for a name, PIN, and card token value. As each query is asked and answered, the remote access client passes through another level of authentication. When all questions have been answered satisfactorily, the remote access client is authenticated.

A specific EAP authentication scheme is known as an EAP type. Both the remote access client and the authenticator must support the same EAP type for successful authentication to occur.

The Windows Server 2003 family includes an EAP infrastructure, two EAP types, and the ability to pass EAP messages to a RADIUS server (EAP-RADIUS).

**EAP infrastructure**

EAP is a set of internal components that provide architectural support for any EAP type in the form of a plug-in module. For successful authentication, both the remote access client and authenticator must have the same EAP authentication module installed. The Windows Server 2003 family provides two EAP types: MD5-Challenge and EAP-TLS. EAP-TLS is only available for members of a domain. You can also install additional EAP types. The components for an EAP type must be installed on every remote access client and every authenticator.

**MD5-Challenge**

Message Digest 5 Challenge (MD5-Challenge) is a required EAP type that uses the same challenge handshake protocol as PPP-based CHAP, but the challenges and responses are sent as EAP messages.

A typical use for MD5-Challenge is to authenticate the credentials of remote access clients by using user name and password security systems. You can also use MD5-Challenge to test EAP interoperability.

**EAP-TLS**

EAP-Transport Level Security (EAP-TLS) is an EAP type that is used in certificate-based security environments. If you are using smart cards for remote access authentication, you must use the EAP-TLS authentication method. The EAP-TLS exchange of messages provides mutual authentication, negotiation of the encryption method, and encrypted key determination between the remote access client and the authenticator. EAP-TLS provides the strongest authentication and key determination method.

EAP-TLS is supported only on servers that are running Routing and Remote Access, that are configured to use Windows Authentication or RADIUS, and that are members of a domain. A remote access server running as a stand-alone server or a member of a workgroup does not support EAP-TLS.

For information about configuring smart cards for remote access clients, see Using smart cards for remote access.

**EAP-RADIUS**

EAP-RADIUS is not an EAP type, but the passing of EAP messages of *any* EAP type by an authenticator to a RADIUS server for authentication. For example, for a remote access server that is configured for RADIUS authentication, the EAP messages sent between the remote access client and remote access server are encapsulated and formatted as RADIUS messages between the remote access server and the RADIUS server.

EAP-RADIUS is used in environments where RADIUS is used as the authentication provider. An advantage of using EAP-RADIUS is that EAP types do not need to be installed at each remote access server, only at the RADIUS server. In the case of an IAS server, you only need to install EAP types on the IAS server.

In a typical use of EAP-RADIUS, a server running Routing and Remote Access is configured to use EAP and to use an IAS server for authentication. When a connection is made, the remote access client negotiates the use of EAP with the remote access server. When the client sends an EAP message to the remote access server, the remote access server encapsulates the EAP

message as a RADIUS message and sends it to its configured IAS server. The IAS server processes the EAP message and sends a RADIUS-encapsulated EAP message back to the remote access server. The remote access server then forwards the EAP message to the remote access client. In this configuration, the remote access server is only a pass-through device. All processing of EAP messages occurs at the remote access client and the IAS server.

Routing and Remote Access can be configured to authenticate locally, or to a RADIUS server. If Routing and Remote Access is configured to authenticate locally, all EAP methods will be authenticated locally. If Routing and Remote Access is configured to authenticate to a RADIUS server, all EAP messages will be forwarded to the RADIUS server with EAP-RADIUS.

For more information about configuring a server running Routing and Remote Access for EAP-RADIUS, see Configure EAP-RADIUS.

**Enabling EAP**

To enable EAP-based authentication, you must do the following:

1. Enable EAP as an authentication protocol on the remote access server. For more information, see Enable EAP.

2. Enable EAP and, if needed, configure the EAP type on the appropriate remote access policy. For more information, see Introduction to remote access policies and Configure authentication.

3. Enable and configure EAP on the remote access client. For more information, see Extensible Authentication Protocol (EAP).

**Note**

- Make sure your network access server (NAS) supports EAP before you enable it on a remote access policy on an IAS server. For more information, see your NAS documentation.

# MS-CHAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MS-CHAP**
The Windows Server 2003 family includes support for the Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), also known as MS-CHAP version 1. MS-CHAP is a nonreversible, encrypted password authentication protocol. The challenge handshake process works as follows:

1. The authenticator (the remote access server or the IAS server) sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string.

2. The remote access client sends a response that contains the user name and a nonreversible encryption of the challenge string, the session identifier, and the password.

3. The authenticator checks the response and, if valid, the user's credentials are authenticated.

If you use MS-CHAP as the authentication protocol, then you can use Microsoft Point-to-Point Encryption (MPPE) to encrypt the data sent on the PPP or PPTP connection.

The Windows Server 2003 family also includes support for MS-CHAP version 2. MS-CHAP version 2 provides stronger security for remote access connections than MS-CHAP. You should consider using MS-CHAP version 2 instead of MS-CHAP. For more information, see MS-CHAP version 2.

**Enabling MS-CHAP**
To enable MS-CHAP-based authentication, you must do the following:

1. Enable MS-CHAP as an authentication protocol on the remote access server. For more information, see Enable authentication protocols. MS-CHAP is enabled by default.

2. Enable MS-CHAP on the appropriate remote access policy. For more information, see Introduction to remote access policies and Configure authentication. MS-CHAP is enabled by default.

3. Enable MS-CHAP on the remote access client. For more information, see Microsoft Challenge Handshake Authentication Protocol (MS-CHAP).

### Notes

- MS-CHAP (version 1 and version 2) is the only authentication protocol provided with the Windows Server 2003 family that supports password change during the authentication process.

- By default, the Windows Server 2003 family implementation of MS-CHAP v1 does not support LAN Manager authentication. If you want to allow the use of LAN Manager authentication with MS-CHAP v1 for older operating systems such as Windows NT 3.5x and Windows 95, you must set the following registry value to 1 on the authenticating server:

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy\Allow LM Authentication**

  Windows 2000 Server supports LAN Manager authentication by default. Upgrading a computer running Windows 2000 Server to a member of the Windows Server 2003 family preserves the existing **Allow LM Authentication** setting.

- If MS-CHAP v1 is used as the authentication protocol, a 40-bit encrypted connection cannot be established if the user's password is larger than 14 characters. This behavior affects both dial-up and virtual private network-based remote access and demand dial connections.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

# MS-CHAP version 2

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MS-CHAP version 2**

The Windows Server 2003 family includes support for version 2 of the Microsoft Challenge Handshake Authentication Protocol (MS-CHAP v2), which provides stronger security for remote access connections. MS-CHAP v2 solves some issues of MS-CHAP version 1, as shown in the following table.

| MS-CHAP version 1 issue | MS-CHAP version 2 solution |
|---|---|
| LAN Manager encoding of the response used for backward compatibility with older Microsoft remote access clients is cryptographically weak. | MS-CHAP v2 no longer allows LAN Manager encoded responses. |
| LAN Manager encoding of password changes is cryptographically weak. | MS-CHAP v2 no longer allows LAN Manager encoded password changes. |
| Only one-way authentication is possible. The remote access client cannot verify that it is dialing in to its organization's remote access server or a masquerading remote access server. | MS-CHAP v2 provides two-way authentication, also known as mutual authentication. The remote access client receives verification that the remote access server that it is dialing in to has access to the user's password. |
| With 40-bit encryption, the cryptographic key is based on the user's password. Each time the user connects with the same password, the same cryptographic key is generated. | With MS-CHAP v2, the cryptographic key is always based on the user's password and an arbitrary challenge string. Each time the user connects with the same password, a different cryptographic key is used. |
| A single cryptographic key is used for data sent in both directions on the connection. | With MS-CHAP v2, separate cryptographic keys are generated for transmitted and received data. |

MS-CHAP v2 is a one-way encrypted password, mutual authentication process that works as follows:

1. The authenticator (the remote access server or the IAS server) sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string.

2. The remote access client sends a response that contains:

   - The user name.

   - An arbitrary peer challenge string.

   - A one-way encryption of the received challenge string, the peer challenge string, the session identifier, and the user's password.

3. The authenticator checks the response from the client and sends back a response containing:

   - An indication of the success or failure of the connection attempt.

   - An authenticated response based on the sent challenge string, the peer challenge string, the encrypted response of the client, and the user's password.

4. The remote access client verifies the authentication response and, if correct, uses the connection. If the authentication response is not correct, the remote access client terminates the connection.

**Enabling MS-CHAP v2**

To enable MS-CHAP v2-based authentication, you must do the following:

1. Enable MS-CHAP v2 as an authentication protocol on the remote access server. For more information, see Enable authentication protocols. MS-CHAP v2 is enabled by default.

2. Enable MS-CHAP v2 on the appropriate remote access policy. For more information, see Introduction to remote access policies and Configure authentication. MS-CHAP v2 is enabled by default.

3. Enable MS-CHAP v2 on the remote access client. For more information, see Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2).

## Notes

- Windows 95 with the Windows Dial-Up Networking 1.3 Performance & Security Upgrade for Windows 95 supports MS-CHAP v2 for virtual private network (VPN) connections but not for dial-up connections.

- MS-CHAP (version 1 and version 2) is the only authentication protocol provided with the Windows Server 2003 family that supports password change during the authentication process.

- Make sure your network access server (NAS) supports MS-CHAP v2 before you enable it on a remote access policy on an IAS server. For more information, see your NAS documentation.

# CHAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CHAP**

The Challenge Handshake Authentication Protocol (CHAP) is a challenge-response authentication protocol that uses the industry-standard Message Digest 5 (MD5) hashing scheme to encrypt the response. CHAP is used by various vendors of network access servers and clients. A server running Routing and Remote Access supports CHAP so that remote access clients that require CHAP are authenticated. Because CHAP requires the use of a reversibly encrypted password, you should consider using another authentication protocol such as MS-CHAP version 2.

To enable CHAP-based authentication, you must do the following:

1. Enable CHAP as an authentication protocol on the remote access server. For more information, see Enable authentication protocols. CHAP is disabled by default.

2. Enable CHAP on the appropriate remote access policy. For more information, see Introduction to remote access policies and Configure authentication.

3. Enable storage of a reversibly encrypted form of the user's password.

   You can enable storage of a reversibly encrypted form of the user's password per user account or enable storage for all accounts in a domain. For more information, see Enable reversibly encrypted passwords in a domain.

4. Force a reset of the user's password so that the new password is in a reversibly encrypted form.

   When you enable passwords to be stored in a reversibly encrypted form, the current passwords are not in a reversibly encrypted form and are not automatically changed. You must either reset user passwords or set user passwords to be changed the next time each user logs on. For more information, see Reset a user password. Once the password is changed, it is stored in a reversibly encrypted form.

   If you set user passwords to be changed the next time a user logs on, the user must log on by using a LAN connection and change the password before they attempt to log on with a remote access connection by using CHAP. You cannot change passwords during the authentication process by using CHAP--the logon attempt fails. One workaround for the remote access user is to temporarily log on by using MS-CHAP to change the password.

5. Enable CHAP on the remote access client. For more information, see Challenge Handshake Authentication Protocol (CHAP).

**Notes**

- If your password expires, CHAP cannot change passwords during the authentication process.

- Make sure your network access server (NAS) supports CHAP before you enable it on a remote access policy on an IAS server. For more information, see your NAS documentation.

- You cannot use Microsoft Point-to-Point Encryption (MPPE) with CHAP.

# SPAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SPAP**
The Shiva Password Authentication protocol (SPAP) is a reversible encryption mechanism employed by Shiva. A computer running Windows XP Professional, when connecting to a Shiva LAN Rover, uses SPAP, as does a Shiva client that connects to a server running Routing and Remote Access. This form of authentication is more secure than plaintext but less secure than CHAP or MS-CHAP.

To enable SPAP-based authentication, you must do the following:

1. Enable SPAP as an authentication protocol on the remote access server. For more information, see Enable authentication protocols. SPAP is disabled by default.

2. Enable SPAP on the appropriate remote access policy. For more information, see Configure authentication. SPAP is disabled by default.

3. Enable SPAP on the remote access client. For more information, see Shiva Password Authentication Protocol (SPAP).

**Important**

- When you enable SPAP as an authentication protocol, the same user password is always sent in the same reversibly-encrypted form. This makes SPAP authentication susceptible to replay attacks, where an attacker captures the packets of the authentication process and replays the responses to gain authenticated access to your intranet. The use of SPAP is discouraged, especially for virtual private network connections.

**Notes**

- If your password expires, SPAP cannot change passwords during the authentication process.

- Make sure your network access server (NAS) supports SPAP before you enable it on a remote access policy on an IAS server. For more information, see your NAS documentation.

- You cannot use Microsoft Point-to-Point Encryption (MPPE) with SPAP.

# PAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PAP**
Password Authentication Protocol (PAP) uses plaintext passwords and is the least secure authentication protocol. It is typically negotiated if the remote access client and remote access server cannot negotiate a more secure form of validation.

To enable PAP-based authentication, you must do the following:

1. Enable PAP as an authentication protocol on the remote access server. For more information, see Enable authentication protocols. PAP is disabled by default.

2. Enable PAP on the appropriate remote access policy. For more information, see Configure authentication. PAP is disabled by default.

3. Enable PAP on the remote access client. For more information, see Password Authentication Protocol (PAP).

### Important

- When you enable PAP as an authentication protocol, user passwords are sent in plaintext form. Anyone capturing the packets of the authentication process can easily read the password and use it to gain unauthorized access to your intranet. The use of PAP is highly discouraged, especially for virtual private network connections.

### Notes

- By disabling the support for PAP on the remote access server, plaintext passwords are never sent by the dial-up client. Disabling support for PAP increases authentication security, but remote access clients who only support PAP cannot connect.

- If your password expires, PAP cannot change passwords during the authentication process.

- Make sure your network access server (NAS) supports PAP before you enable it on a remote access policy on an IAS server. For more information, see your NAS documentation.

- You cannot use Microsoft Point-to-Point Encryption (MPPE) with PAP.

# Unauthenticated access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Unauthenticated access**
The Windows Server 2003 family supports unauthenticated access, which means that user credentials (a user name and password) are not required. There are some situations where unauthenticated access is useful. This section covers:

- DNIS authorization
- ANI/CLI authentication
- Guest authentication

**Important**

- When you enable unauthenticated access, remote access users are connected without sending user credentials. An unauthenticated remote access client does not negotiate the use of a common authentication protocol during the connection establishment process and does not send a user name or password.

  Unauthenticated access with remote access clients could occur when the authentication protocols configured by the remote access client do not match those configured on the remote access server. In this case, the use of a common authentication protocol is not negotiated and the remote access client does not send a user name and password.

# DNIS authorization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNIS authorization**

Dialed Number Identification Service (DNIS) enables authorization of a connection attempt based on the number called. DNIS identifies the number that was called to the receiver of the call and is provided by most standard telephone companies. For example, you could allow all connections that connect through a specific toll-free number. To identify DNIS-based connections and apply the appropriate connection settings, you must do the following:

1. Enable unauthenticated access on the remote access server. For more information, see Enable authentication protocols.

2. Create a remote access policy on the authenticating server (remote access server or IAS server) for DNIS-based authorization with the Called-Station-Id condition set to the phone number. For more information, see Introduction to remote access policies and Add a remote access policy.

3. Enable unauthenticated access on the remote access policy for DNIS-based authorization. For more information, see Configure authentication.

### Notes

- If your phone service or hardware does not support DNIS, the passing of the number that was called, you can manually set the phone number of the port. For more information, see Set the phone number on a port.

- IAS does not support proxied DNIS access requests.

# ANI/CLI authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ANI/CLI authentication**
Automatic Number Identification/Calling Line Identification (ANI/CLI) authentication is the authentication of a connection attempt based on the phone number of the caller. ANI/CLI service returns the number of the caller to the receiver of the call and is provided by most standard telephone companies.

ANI/CLI authentication is different from caller ID authorization. In caller ID authorization, the caller sends a valid user name and password. The caller ID that is configured for the dial-in property on the user account must match the connection attempt; otherwise, the connection attempt is rejected. In ANI/CLI authentication, a user name and password are not sent.

To identify ANI/CLI-based connections and apply the appropriate connection settings, you must do the following:

1. Enable unauthenticated access on the remote access server. For more information, see Enable authentication protocols.

2. Enable unauthenticated access on the appropriate remote access policy for ANI/CLI-based authentication. For more information, see Introduction to remote access policies and Configure authentication.

3. Create a user account for each number that will be calling for which you want to provide ANI/CLI authentication. The name of the user account must match the number that the user is dialing from. For example, if a user is dialing in from 555-0100, create a "5550100" user account.

4. Set the following registry value to 31 on the authenticating server (either the remote access server or the IAS server):

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy\User Identity Attribute**

   This registry setting tells the authenticating server to use the calling number (RADIUS attribute 31, Calling-Station-ID) as the identity of the calling user. The user identity is set to the calling number only when there is no user name being supplied in the connection attempt.

   To always use the calling number as the user identity, set the following registry value to 1 on the authenticating server:

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy\Override User-Name**

   However, if you set **Override User-Name** to 1 and the **User Identity Attribute** to 31, the authenticating server can only perform ANI/CLI-based authentication. Normal authentication by using authentication protocols such as MS-CHAP, CHAP, and EAP is disabled.

### Note

- Changes to the registry settings will not take effect until the Routing and Remote Access service or the Internet Authentication Service are restarted.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

# Guest authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Guest authentication**
By default, the Windows Server 2003 family includes a user account called Guest. With guest authentication, during the authentication process, the caller does not send a user name or password. If unauthenticated access is enabled, by default the Guest account is used as the identity of the caller.

To enable Guest account access, you must do the following:

1. Enable unauthenticated access on the remote access server. For more information, see Enable authentication protocols.

2. Enable unauthenticated access on the appropriate remote access policy. For more information, see Configure authentication.

3. Enable the Guest account. For more information, see To enable a disabled user account.

4. Set the remote access permission on the Guest account to either **Allow access** or **Control access through Remote Access Policy** depending on your remote access policy administrative model. For more information, see Remote Access Policies Examples and Configure remote access permission for a user.

If you want to enable a guest account that is not named **Guest**, create a user account and set the remote access permission to either **Allow access** or **Control access through Remote Access Policy**. Then, set the following registry value on the authenticating server (either the remote access server or the IAS server) to the name of the account:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy\Default User Identity**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Notes**

- Changes to the registry setting will not take effect until the Routing and Remote Access service or the Internet Authentication Service are restarted.

- For more information about how to modify registry keys and set values, see "Remote Access Server" at the Microsoft Windows Resource Kits Web site.

# Authentication vs. authorization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication vs. authorization**
The distinction between authentication and authorization is important in understanding why connection attempts are either accepted or denied:

- Authentication is the verification of the credentials of the connection attempt. This process consists of sending the credentials from the remote access client to the remote access server in an either plaintext or encrypted form by using an authentication protocol.

- Authorization is the verification that the connection attempt is allowed. Authorization occurs after successful authentication.

For a connection attempt to be accepted, the connection attempt must be both authenticated and authorized. It is possible for the connection attempt to be authenticated by using valid credentials, but not authorized. In this case, the connection attempt is denied.

If a remote access server is configured for **Windows Authentication**, the security features of the Windows Server 2003 family are used to verify the credentials for authentication, and the dial-in properties of the user account and locally stored remote access policies are used to authorize the connection. If the connection attempt is both authenticated and authorized, the connection attempt is accepted.

For more information, see Introduction to remote access policies.

If the remote access server is configured for RADIUS authentication, the credentials of the connection attempt are passed to the RADIUS server for authentication and authorization. If the connection attempt is both authenticated and authorized, the RADIUS server sends an accept message back to the remote access server and the connection attempt is accepted. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends a reject message back to the remote access server and the connection attempt is rejected.

If the RADIUS server is a computer running the Internet Authentication Service (IAS), the IAS server performs authentication through selected authentication features and authorization through the dial-in properties of the user account and remote access policies stored on the IAS server.

# Pre-shared keys

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Pre-shared keys**

A pre-shared key is a string of Unicode characters used to authenticate L2TP/IPSec connections. You can configure Routing and Remote Access to authenticate VPN connections that support a pre-shared key. The use of pre-shared keys is supported on many operating systems, including the Windows Server 2003 family and Windows XP. You can also configure a server running the Windows Server 2003 family version of Routing and Remote Access to use a pre-shared key to authenticate connections from other routers.

**Advantages and disadvantages of pre-shared keys**

Pre-shared key authentication does not require the hardware and configuration investment of a public key infrastructure (PKI), which is necessary for using computer certificates for L2TP/IPSec authentication. Pre-shared keys are simple to configure on a remote access server, and they are relatively simple to configure on a remote access client. They can be made transparent to the user if they are issued within Connection Manager profiles. If you are in the process of establishing a PKI or you manage an Active Directory domain, you can configure Routing and Remote Access to accept an L2TP/IPSec connection using a computer certificate or a pre-shared key.

However, a single remote access server can utilize only one pre-shared key for all L2TP/IPSec connections that require a pre-shared key for authentication. Therefore, you must issue the same pre-shared key to all L2TP/IPSec VPN clients that connect to the remote access server using a pre-shared key. Unless you distribute the pre-shared key within a Connection Manager profile, each user must manually type the pre-shared key. This limitation further reduces the security of the deployment and increases the probability of error. Moreover, if the pre-shared key on a remote access server is changed, a client with a manually configured pre-shared key will be unable to connect to that server until the pre-shared key on the client is changed. If the pre-shared key was distributed to the client within a Connection Manager profile, that profile must be reissued with the new pre-shared key and reinstalled on the client computer. Unlike certificates, the origin and the history of a pre-shared key cannot be determined. For these reasons, the use of pre-shared keys to authenticate L2TP/IPSec connections is considered a relatively weak authentication method. If you want a long term, strong authentication method, you should consider using a PKI.

**Considerations when choosing a pre-shared key**

A pre-shared key is a sequence of characters that is configured on both the remote access server and the L2TP/IPSec client. The pre-shared key can be any non-null string of any combination of up to 256 Unicode characters. When you choose a pre-shared key, consider the fact that users who utilize the New Connection wizard to create the VPN client connection must type the pre-shared key manually. A key that is long and complex enough to provide adequate security might be difficult for the majority of your users to type accurately. If the pre-shared key presented by the VPN client deviates in any way from the pre-shared key configured on the remote access server, client authentication will fail.

When the pre-shared key is first stored, the remote access server and the VPN client attempt to convert the Unicode string into ASCII. If the attempt is successful, the ASCII version of the string will be used for authentication. This strategy ensures that the pre-shared key is not corrupted in transmission by any devices that do not comply with the Unicode standard, such as routers from other companies. If the pre-shared key cannot be stored as ASCII, the Unicode string will be used. If the Unicode pre-shared key must be processed by any device that does not comply with the Unicode standard, the connection attempt will almost certainly fail.

**Pre-shared keys and the Connection Manager Administration Kit**

By using the Connection Manager Administration Kit (CMAK) wizard, you can create a customized connection experience for your users. You can use the CMAK wizard to create a VPN connection profile that includes a pre-shared key. Because profiles are self-extracting, the user does not have to type in the pre-shared key or even to know that a pre-shared key exists. You can further increase the security of your Connection Manager profile distribution by encrypting the pre-shared key with a personal identification number (PIN). In this way, not only does the user never see or have to type the pre-shared key, you can distribute the profile and the PIN separately, reducing the chance that an unauthorized user might obtain access to your network.

For more information about Connection Manager profiles, see Connection Manager Administration Kit. For more information about VPNs, see Virtual Private Networks and Virtual private networking with IPSec. For more information about PKIs, see Deploying a Public Key Infrastructure.

# Remote Access Data encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data encryption**

You can use data encryption to protect the data that is sent between the remote access client and the remote access server. Data encryption is important for anyone who wants or needs secure data transfer, including financial institutions, law-enforcement and government agencies, and corporations. For installations where data confidentiality is required, the network administrator can set the remote access server to require encrypted communications. Users who connect to that server must encrypt their data or the connection attempt is denied.

For dial-up networking connections, you can protect your data by encrypting it on the communications link between the remote access client and the remote access server. You should use data encryption when there is a risk of unauthorized interception of transmissions on the communications link between the remote access client and the remote access server. For dial-up networking connections, Routing and Remote Access uses Microsoft Point-to-Point Encryption (MPPE).

For virtual private networking connections, you can protect your data by encrypting it between the ends of the virtual private network (VPN). You should always use data encryption for VPN connections when private data is sent across a public network such as the Internet, where there is always a risk of unauthorized interception. For VPN connections, Routing and Remote Access uses MPPE with the Point-to-Point Tunneling Protocol (PPTP) and IP Security (IPSec) encryption with the Layer Two Tunneling Protocol (L2TP).

Because data encryption is performed between the VPN client and VPN server, it is not necessary to use data encryption on the communication link between a dial-up client and its Internet service provider (ISP). For example, a mobile user uses a dial-up networking connection to dial in to a local ISP. Once the Internet connection is made, the user creates a VPN connection with the corporate VPN server. If the VPN connection is encrypted, there is no need to use encryption on the dial-up networking connection between the user and the ISP.

MPPE and IPSec are configured on the **Encryption** tab on the properties of a remote access policy to use 40-bit (the **Basic** setting), 56-bit (the **Strong** setting), or 128-bit (the **Strongest** setting) encryption keys. You should use 40-bit encryption keys to connect with older operating systems that do not support 56-bit or 128-bit encryption keys. Otherwise, use 56-bit encryption keys. Encryption keys are determined at the time of the connection. MPPE requires the use of the MS-CHAP (v1 or v2) or EAP-TLS authentication protocols.

**Notes**

- Data encryption for PPP or PPTP connections is available only if MS-CHAP (v1 or v2) or EAP-TLS is used as the authentication protocol. Data encryption for L2TP connections relies on IPSec, which does not require any specific authentication protocol.

- Remote access data encryption does not provide end-to-end data encryption. End-to-end encryption is data encryption between the client application and the server that hosts the resource or service being accessed by the client application. To get end-to-end data encryption, use IPSec to create a secure connection *after* the remote access connection has been made.

# Remote access dial-in permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access dial-in permissions**

After a remote access server is installed, you must specify from which users the remote access server can accept a connection. For a server running Routing and Remote Access, authorization is determined by the dial-in properties on the user account, the remote access policies, or both. For more information, see Configure Remote Access Policies.

You do not need to create user accounts just for remote access users. Remote access servers use the user accounts specified in the available user accounts databases.

**How security works at connection**

The following steps describe what happens during a call from a remote access client to a server running Routing and Remote Access that is configured to use Windows Authentication:

1. A remote access client dials a remote access server.
2. The server sends a challenge to the client.
3. The client sends an encrypted response to the server that consists of a user name, a domain name, and a password.
4. The server checks the response against the appropriate user accounts database.
5. If the account is valid and the authentication credentials are correct, the server uses the dial-in properties of the user account and remote access policies to authorize the connection.

If callback is enabled, the server hangs up the connection, calls the client back, and continues the connection negotiation process.

**Notes**

- Steps 2 and 3 assume that the remote access client and the remote access server use the MS-CHAP v1 or CHAP authentication protocols. The sending of client credentials may vary for other authentication protocols.
- If the remote access server is a member of domain and the user response does not contain a domain name, then the domain name of the remote access server is used. If you want to use a different domain name than that of the remote access server, set the following registry value on the remote access client to the name of the domain that you want to use:

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn\DefaultDomain**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Security after the connection is made**

Credentials used for remote access only provide a communication channel to the target network. The client does not log on to the network as a result of a remote access connection. Each time the client attempts to access a network resource, it will be challenged for credentials. If it does not respond to the challenge with acceptable credentials, the access attempt will fail. Windows XP and the Windows Server 2003 family adds a feature to simplify remote access. After a successful connection, Windows XP and Windows Server 2003 family remote access clients will cache these credentials as default credentials for the duration of the remote access connection. When a network resource challenges the remote access client, the client provides the cached credentials without requiring the user to enter them again.

# Caller ID and callback

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Caller ID and callback**
As an additional measure of security, Routing and Remote Access offers caller ID and callback features, which ensure that only users from specific locations can access the remote access server. These features also save telephone charges for the user.

**Caller ID**
When you set dial-in security by using the caller ID feature, you specify the phone number from which the user must call in. If the user does not call in from that specific phone number, the connection attempt is rejected by the remote access server.

Caller ID must be supported by the caller, the phone system between the caller and the remote access server, and the remote access server. Caller ID on the remote access server consists of call answering equipment that supports the passing of caller ID information and the appropriate drivers that support the passing of caller ID information to the Routing and Remote Access service.

If you configure a caller ID phone number for a user and you do not have support for the passing of caller ID information all the way from the caller to the Routing and Remote Access service, the connection attempt is denied.

The caller ID feature is designed to provide a higher degree of security for telecommuters. The disadvantage of configuring caller ID is that the user can only dial in from a specific phone line.

For virtual private network (VPN) connections, the caller ID is the IP address of the VPN client. If the VPN client has a fixed IP address, this can be a useful security feature.

**Callback**
When you use the callback feature, the user initiates a call and connects with the remote access server. After authentication and authorization, the remote access server then drops the call and calls back a moment later to a negotiated or preassigned callback number.

You configure each user's callback options when you grant remote access permission. For more information, see Configure dial-in user properties.

There are three callback options to choose from:

- **No callback** (the default)
- **Set by caller**
- **Always call back to**

**Note**

- Until the user has been authenticated, authorized, and called back (if callback is set), no data from the dial-up networking client or the remote access server is transferred.

**No callback**
If the user account is not configured for callback, the remote access server establishes a connection as soon as the connection attempt has been accepted. The **No callback** option does not provide any additional security.

**Set by caller**
Although the **Set by caller** option is not really a security feature, it is useful for clients who call from various locations and phone numbers. It also minimizes telephone charges for these users. When the user's call reaches the remote access server, the following events occur:

1. After authentication and authorization of the connection attempt, the **Callback** dialog box appears on the user's computer.
2. The user types the current callback number in the dialog box.

3. The callback number is sent to the server.

4. The call is terminated.

5. The server calls the client back at the callback number.

6. Once reconnected, the client and server continue the connection negotiation.

**Always call back to**

For additional security, select the **Always call back to** option and type the number of the phone to which the user's dial-up equipment is connected. When the user's call reaches the remote access server, the following events occur:

1. After authentication and authorization of the connection attempt, the server sends a message announcing that the user will be called back.

2. The server disconnects and calls the user back at the preset number.

3. Once reconnected, the client and server continue the connection negotiation.

You should set this option for stationary remote computers, such as those used by telecommuters in home offices.

The disadvantage of configuring callback to always call a specific number is that the user can only dial in from a specific location.

### Notes

- You can also configure the **Set by caller** callback option for groups by setting the **Service-Type** condition of a remote access policy to **Callback-Framed-User**. For more information, see Elements of a remote access policy.

- Because of the way that callback connections are processed, you cannot configure both a caller ID and callback that is set to either **Set by caller** or **Always call back to**.

- Callback over a primary rate interface (PRI) ISDN channel may not work properly if a service is listening on the other ISDN channel. When the remote access server calls back, an ISDN channel is picked to receive the call. If the ISDN channel is not the same one used to make the initial call, the remote access client or demand-dial router does not recognize the incoming call as the remote access server callback and drops the call.

# Remote Access Security hosts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security hosts**

A security host is an authentication device that verifies whether a caller from a remote access client is authorized to connect to the remote access server. This verification supplements security already supplied by the server running Routing and Remote Access. This verification is only applicable to dial-up or modem connections.

The security host sits between the remote access client and the remote access server. The security host generally provides an extra layer of security by requiring a hardware key of some sort in order to provide authentication. Verification that the remote access client is in physical possession of the key takes place before access to the remote access server is granted. This open architecture allows customers to choose from a variety of security hosts to augment the security in Routing and Remote Access.

For example, one kind of security system consists of two hardware devices: the security host and the security card. The security host is installed between the remote access server and its modem. The security card is a small unit the size of a credit card, resembling a pocket calculator without keys. The security card displays a different access number every minute. This number is synchronized with the same number calculated in the security host every minute. When connecting, the remote user sends the number on the security card to the host. If the number is correct, the security host connects the remote access client with the remote access server.

Another kind of security host prompts the remote access client to type in a user name (which may or may not be the same as the remote access client name) and a password (which differs from the remote access client password).

You must configure the security host to allow the remote access server to initialize the modem before the security functions take effect. The remote access server must also be able to directly initialize the modem that is connected to the security host without security checks from the security host. The security host might interpret the attempt of the remote access server to initialize the modem as an attempt to dial out.

For more information, see Configure other security devices.

# Remote access account lockout

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access account lockout**

You can use the remote access account lockout feature to specify how many times a remote access authentication fails against a valid user account before the user is denied access. Remote access account lockout is especially important for remote access virtual private network (VPN) connections over the Internet. An attacker on the Internet can attempt to access an organization intranet by sending credentials (valid user name, guessed password) during the VPN connection authentication process. During a dictionary attack, the attacker sends hundreds or thousands of credentials by using a list of passwords based on common words or phrases.

With remote access account lockout enabled, a dictionary attack is thwarted after a specified number of failed attempts. As the network administrator, you must decide on two remote access account lockout variables:

1. The number of failed attempts before future attempts are denied.

   After each failed attempt, a failed attempts counter for the user account is incremented. If the user account's failed attempts counter reaches the configured maximum, future attempts to connect are denied.

   A successful authentication resets the failed attempts counter when its value is less than the configured maximum. In other words, the failed attempts counter does not accumulate beyond a successful authentication.

2. How often the failed attempts counter is reset.

   The failed attempts counter is periodically reset to 0. If an account is locked out after the maximum number of failed attempts, the failed attempts counter is automatically reset to 0 after the reset time.

You enable the remote access account lockout feature by changing settings in the registry on the computer that provides the authentication. If the remote access server is configured for Windows Authentication, modify the registry on the remote access server computer. If the remote access server is configured for RADIUS authentication and Internet Authentication Service (IAS) is being used, modify the registry on the IAS server.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

To enable remote access account lockout, you must set the **MaxDenials** entry in the registry to **1** or greater. **MaxDenials** is the maximum number of failed attempts before the account is locked out. You set the **MaxDenials** entry in the following registry subkey:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout**

By default, **MaxDenials** is set to **0**, which means that remote access account lockout is disabled.

To modify the amount of time before the failed attempts counter is reset, you must set the **ResetTime (mins)** entry in the registry to the required number of minutes. You set the **ResetTime (mins)** entry in the following registry subkey:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout**

By default, **ResetTime (mins)** is set to **0xb40**, or 2,880 minutes (48 hours).

**Manually resetting an account that is locked out**

To manually reset a user account that has been locked out before the failed attempts counter is automatically reset, delete the following registry subkey that corresponds to the user's account name:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout\ *domain name:user name***

When the lockout count for a user account is reset to 0 due to either a successful authentication or an automatic reset, the registry subkey for the user account is deleted.

**Notes**

- Remote access account lockout is not related to the **Account locked out** setting on the **Account** tab on the properties of a user account.

- The remote access account lockout feature does not distinguish between malicious users who attempt to access your intranet and authentic users who attempt remote access but have forgotten their current passwords. Users who have forgotten their current password typically try the passwords that they remember and, depending on the number of attempts and the **MaxDenials** setting, may have their accounts locked out.

- If you enable the remote access account lockout feature, a malicious user can deliberately force an account to be locked out by attempting multiple authentications with the user account until the account is locked out, thereby preventing the authentic user from being able to log on.

# Remote Access Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access policies**
This section covers:

- Introduction to remote access policies
- Dial-in properties of a user account
- Elements of a remote access policy
- Accepting a connection attempt
- Remote Access Policies Examples

# Introduction to remote access policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introduction to remote access policies**

In Windows NT versions 3.5, 3.51, and 4.0, authorization was based on a **Grant dial-in permission to user** option in either User Manager or the Remote Access Admin utility. An individual user configured callback options. For the Routing and Remote Access service and Internet Authentication Service (IAS) in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000, network access authorization is granted on the basis of user account dial-in properties and remote access policies.

Remote access policies are an ordered set of rules that define how connections are either authorized or rejected. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting.

If a connection is authorized, the remote access policy profile specifies a set of connection restrictions. The dial-in properties of the user account also provide a set of restrictions. Where applicable, user account connection restrictions override the remote access policy profile connection restrictions.

For information about connection restrictions that can be configured on a user account, see Dial-in properties of a user account. For information about connection restrictions that can be configured through a remote access policy profile, see Elements of a remote access policy.

Remote access policies validate a number of connection settings before authorizing the connection, including the following:

- Remote access permission
- Group membership
- Type of connection
- Time of day
- Authentication methods
- Advanced conditions:
    - Access server identity
    - Access client phone number or MAC address
    - Whether user account dial-in properties are ignored
    - Whether unauthenticated access is allowed

After the connection is authorized, remote access policies can also be used to specify connection restrictions, including the following:

- Idle timeout time
- Maximum session time
- Encryption strength
- IP packet filters
- Advanced restrictions:

- IP address for PPP connections
- Static routes

Additionally, you can vary connection restrictions based on the following settings:

- Group membership
- Type of connection
- Time of day
- Authentication methods
- Identity of the access server
- Access client phone number or MAC address
- Whether unauthenticated access is allowed

For example, you can have policies that specify different maximum session times for different types of connections or groups. Additionally, you can specify restricted access for business partners or unauthenticated connections.

**Authorizing access**
There are two ways to use remote access policies to grant authorization:

1. By user
2. By group

**Authorization by user**
If you are managing authorization by user, set the remote access permission on the user or computer account to either **Grant access** or **Deny access** and, optionally, create different remote access policies based on different types of connections. For example, you might want to have one remote access policy that is used for dial-up connections and a different remote access policy that is used for wireless connections. Managing authorization by user is recommended only when you have a small number of user or computer accounts to manage.

If you are managing authorization by user, the basic process for authorizing a connection attempt occurs as follows:

- If the connection attempt matches all policy conditions, check the remote access permission setting of the account.
    - If the remote access permission is set to **Grant access**, apply the connection settings of the policy profile and user account.
    - If the remote access permission is set to **Deny access**, reject the connection attempt.
- If the connection attempt does not match all policy conditions, process the next remote access policy.

If the connection attempt does not match all conditions of any remote access policy, reject the connection attempt.

**Authorization by group**
If you are managing authorization by group, set the remote access permission on the user account to **Control access through Remote Access Policy** and create remote access policies that are based on different types of connections and group membership. For example, you might want to have one remote access policy for dial-up connections for employees (members of the Employees group) and a different remote access policy for dial-up connections for contractors (members of the Contractors group).

If you are managing authorization by group, the basic process for authorizing a connection attempt occurs as follows:

- If the connection attempt matches all policy conditions, check the remote access permission of the remote access policy.
    - If the remote access permission is set to **Grant remote access permission**, apply the connection settings of the policy profile and user account.
    - If the remote access permission is set to **Deny remote access permission**, reject the connection attempt.
- If the connection attempt does not match all policy conditions, process the next remote access policy.

If the connection attempt does not match all conditions of any remote access policy, reject the connection attempt.

**Notes**

- Setting the remote access permission on user accounts to **Control access through Remote Access Policy** and not using groups to manage network access is not recommended.
- For more information about how connection attempts are processed, see Remote Access Policies Examples

**Remote access policies and the Routing and Remote Access service**

For servers running the Routing and Remote Access service that are configured for the Windows authentication provider, remote access policies are administered from Routing and Remote Access and apply only to the connections of the Routing and Remote Access service server. For servers running the Routing and Remote Access service that are configured for the RADIUS authentication provider, you can no longer configure remote access policies using Routing and Remote Access. If the RADIUS server is an IAS server, remote access policies are administered from Internet Authentication Service.

To centrally manage a single set of remote access policies for multiple remote access or VPN servers using an IAS server, you must do the following:

1. Install the Internet Authentication Service (IAS) as a Remote Authentication Dial-In User Service (RADIUS) server. For more information, see Checklist: Configuring IAS for dial-up and VPN access.
2. Configure IAS with RADIUS clients that correspond to each of the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition remote access or VPN servers. For more information, see Add RADIUS clients.
3. On the IAS server, create the central set of policies that will be used by all servers running the Routing and Remote Access service. For more information, see Add a remote access policy.
4. Configure each of the servers running the Routing and Remote Access service as RADIUS clients to the IAS server. For more information, see Use RADIUS authentication.

For more information about deploying IAS for centralized remote access policy management, see Using RADIUS for multiple remote access servers.

Centralized management of remote access policies is also used when you have remote access servers that are running Windows NT 4.0 and the Routing and Remote Access Service (RRAS). You can configure this server as a RADIUS client to an IAS server. However, a remote access server running Windows NT 4.0 without RRAS cannot be configured to use centralized remote access policies.

# Dial-in properties of a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-in properties of a user account**

In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; the user account for a stand-alone server or server running Active Directory contains a set of dial-in properties that are used when allowing or denying a connection attempt made by a user. On a stand-alone server, you can set the dial-in properties on the **Dial-in** tab in the user account in Local Users and Groups. On a server running Active Directory, you can set the dial-in properties on the **Dial-in** tab in the user account in Active Directory Users and Computers. On these servers, you cannot use the Windows NT 4.0 User Manager for Domains administrative tool.

The dial-in properties for a user account are:

- **Remote Access Permission (Dial-in or VPN)**

  You can use this property to set remote access permission to be explicitly allowed, denied, or determined through remote access policies. In all cases, remote access policies are used to authorize the connection attempt. If access is explicitly allowed, remote access policy conditions, user account properties, or profile properties can still deny the connection attempt. The **Control access through Remote Access Policy** option is only available on user accounts in a Windows 2000 native domain, a Windows server 2003 domain, or for local accounts on stand-alone servers running Windows 2000, Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition.

  By default, the Administrator and Guest accounts on a stand-alone server or in a Windows 2000 native domain are set to **Control access through Remote Access Policy**. In a Windows 2000 mixed domain, they are set to **Deny access**. New accounts that are created on a stand-alone server or in a Windows 2000 native domain are set to **Control access through Remote Access Policy**. New accounts that are created in a Windows 2000 mixed domain are set to **Deny access**.

  For more information, see Domain and forest functionality.

- **Verify Caller ID**

  If this property is enabled, the server verifies the caller's phone number. If the caller's phone number does not match the configured phone number, the connection attempt is denied.

  Caller ID must be supported by the caller, the phone system between the caller and the remote access server, and the remote access server. On a computer running the Routing and Remote Access service, caller ID support consists of call answering equipment that provides caller ID information and the appropriate Windows Server 2003, Standard Edition ; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition driver to pass the information to the Routing and Remote Access service.

  If you configure a caller ID phone number for a user, and you do not have support for the passing of caller ID information from the caller to the Routing and Remote Access service, the connection attempt is denied.

- **Callback Options**

  If this property is enabled, the server calls the caller back during the connection process. The phone number that is used by the server is set by either the caller or the network administrator.

- **Assign a Static IP Address**

  You can use this property to assign a specific IP address to a user when a connection is made.

- **Apply Static Routes**

You can use this property to define a series of static IP routes that are added to the routing table of the server running the Routing and Remote Access service when a connection is made. This setting is designed for user accounts that a router running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition uses for demand-dial routing. For more information, see One-way initiated demand-dial connections.

For information about configuring dial-in properties for a user account, see Configure dial-in user properties.

**Support for ignoring the dial-in properties of user accounts**
In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; you can configure a RADIUS attribute to ignore the dial-in properties of user and computer accounts in the profile properties of a remote access policy. To support multiple types of connections for which IAS provides authentication and authorization, it might be necessary to disable the processing of user account dial-in properties. This can be done to support scenarios in which specific dial-in properties are not required.

For example, the caller-ID, callback, static IP address, and static routes properties are designed for a client that is dialing into a network access server (NAS). These settings are not designed for wireless access points. A wireless access point that receives these settings in the RADIUS message from the IAS server might be unable to process them, which could cause the wireless client to become disconnected. When IAS provides authentication and authorization for users who are both dialing in and accessing the organization network through wireless technology, the dial-in properties must be configured to support either dial-in connections (by setting dial-in properties) or wireless connections (by not setting dial-in properties).

You can use IAS to enable the processing of dial-in properties for user and computer accounts in some scenarios (such as dial-in), and to disable the processing of dial-in properties for user and computer accounts in other scenarios (such as wireless and authenticating switch). This is done by configuring the **Ignore-User-Dialin-Properties** attribute on the **Advanced** tab of the profile settings for a remote access policy. For more information, see Add RADIUS attributes to a remote access policy.

The **Ignore-User-Dialin-Properties** attribute is set to the following:

- To enable the processing of dial-in properties for a user account, delete the **Ignore-User-Dialin-Properties** attribute or set it to **False**. For example, for a remote access policy that is designed for dial-in connections, no additional configuration is required.

- To disable the processing of dial-in properties for a user account, set the **Ignore-User-Dialin-Properties** attribute to the value of **True**. For example, this is set for the remote access policy that is used for wireless or authenticating switch connections. When the dial-in properties of the user account are ignored, remote access permission is determined by the setting in the remote access policy.

### Notes

- For a user account in a Windows NT 4.0 domain or a Windows 2000 mixed domain:

    - Only the dial-in properties named **Remote Access Permission (Dial-in or VPN)** (**Allow access** and **Deny access** options) and **Callback Options** are available.

    - You can use the Windows NT 4.0 User Manager for Domains administrative tool to grant or deny dial-in access and set callback options.

- If a user account is in a Windows 2000 native domain, the callback number can be up to 128 characters in length. If a user account is on a stand-alone remote access server, in a Windows NT 4.0 domain, or in a Windows 2000 mixed domain, the callback number can contain 24 through 48 characters. This is due to the compressed format for storing callback numbers.

- When a remote access server running Windows NT 4.0 uses a Windows 2000 native domain to obtain the dial-in properties of a user account, the **Control access through Remote Access Policy** remote access permission setting is retrieved as **Deny access**. Callback settings are retrieved correctly.

  For more information, see Domain and forest functionality.

- A remote access server running Windows NT 4.0 cannot use remote access policies. If you upgrade this remote access server with the Routing and Remote Access Service (RRAS) and configure it for RADIUS authentication, you can use the remote access policies of an IAS server. Alternately, you can upgrade this remote access server with a member of the Windows 2000 server family; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition.

- User accounts upgraded to Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition that were configured with dial-in permission enabled are set to **Allow access**. Those that were configured with dial-in permission disabled are set to **Deny access**.

- It is possible to ignore all of the user or computer account properties. For more information, see the section "Ignoring the dial-in properties of user accounts" in New features for IAS.

- In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; computer accounts also have dial-in properties and are similar to user accounts. This new feature allows computers to be authenticated as if they were users, and is used when an IEEE 802.1x Ethernet client uses EAP-TLS and an installed computer certificate to authenticate itself to an Ethernet switch.

- The **Ignore-User-Dialin-Properties** attribute disables the use of all dial-in properties for the user account. Specific dial-in properties cannot be selectively disabled.

# Elements of a remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Elements of a remote access policy**

A remote access policy is a named rule that consists of the following elements:

- Conditions
- Remote access permission
- Profile

**Conditions**

Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt. If there are multiple conditions, then all of the conditions must match the settings of the connection attempt in order for it to match the policy.

The following table shows the condition attributes that you can set for a remote access policy.

| Attribute name | Description |
|---|---|
| Authentication Type | The type of authentication that is being used by the access client. Authentication types include CHAP, EAP, MS-CHAP, and MS-CHAP v2. |
| Called Station ID | The phone number of the network access server (NAS). This attribute is a character string. You can use pattern matching syntax to specify area codes. For more information, see Pattern matching syntax. In order to receive called station ID information during a call, the phone line, the hardware, and the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition driver for the hardware must support passing the called ID. Otherwise, the called station ID is manually set for each port. To manually set phone numbers on ports, see Set the phone number on a port. |
| Calling Station ID | The phone number used by the caller. This attribute is a character string. You can use pattern matching syntax to specify area codes. For more information, see Pattern matching syntax. For more information about obtaining the calling station ID, see Caller ID and callback. |
| Client Friendly Name | The name of the RADIUS client that is requesting authentication. This name is configured in **Friendly name** on the **Settings** tab in the properties of a RADIUS client in Internet Authentication Service. This attribute is a character string. You can use pattern matching syntax to specify client names. For more information, see Pattern matching syntax. This attribute is used by the IAS server. |
| Client IP Address | The IP address of the RADIUS client. This can be the IP address of either the originating RADIUS client or an intermediate RADIUS proxy. This attribute is a character string. You can use pattern matching syntax to specify IP networks. For more information, see Pattern matching syntax. This attribute is used by the IAS server. |
| Client Vendor | The vendor of the network access server (NAS) that is requesting authentication. The Routing and Remote Access service is the **Microsoft** NAS manufacturer. You can use this attribute to configure separate policies for different NAS manufacturers who are RADIUS clients to an IAS server. This attribute is used by the IAS server. Ensure that you configure the NAS as a RADIUS client on the IAS server. For more information, see Add RADIUS clients. |
| Day and Time Restrictions | The day of the week and the time of day of the connection attempt. The day and time is relative to the day and time of the server providing the authorization. |
| Framed Protocol | The type of framing for incoming packets. Examples are PPP, SLIP, Frame Relay, and X.25. This attribute is used by the IAS server. |

| | |
|---|---|
| NAS Identifier | The name of the network access server (NAS). This attribute is a character string. You can use pattern matching syntax to specify multiple NASs. For more information, see Pattern matching syntax. This attribute is used by the IAS server. |
| NAS IP Address | The IP address of the NAS (the RADIUS client) that sent the message. This attribute is a character string. You can use pattern matching syntax to specify IP networks. For more information, see Pattern matching syntax. This attribute is used by the IAS server. |
| NAS Port Type | The type of media that is used by the access client. Examples include analog phone lines (known as async), ISDN, tunnels or virtual private networks (known as virtual), IEEE 802.11 wireless, and Ethernet switches. |
| Service Type | The type of service that is being requested. Examples include framed (such as PPP connections) and login (such as Telnet connections). For more information about RADIUS service types, see RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)." This attribute is used by the IAS server. |
| Tunnel Type | The type of tunnel that is being created by the requesting client. Tunnel types include the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol (L2TP), which are used by Windows XP; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000, remote access clients and demand-dial routers. You can use this condition to specify profile settings, such as authentication methods or encryption strengths for a specific type of tunneling technology. |
| Windows Groups | The names of the groups to which the user or computer account that is attempting the connection belongs. There is no condition attribute for a specific user or computer name. It is not necessary to have a separate remote access policy for each group. Instead, you can use multiple groups or nested groups to consolidate and delegate the administration of group membership. For a remote access or IAS server in a Windows 2000 native domain, you can use universal groups. For more information, see Domain and forest functionality. |

**Notes**

- If conditions that use an IAS server attribute are evaluated against a Routing and Remote Access service server that is configured for Windows authentication, they do not match and the policy is not applied.
- Not all access servers send all of the IAS server-specific attributes.
- You cannot use the built-in local or domain groups for the Windows Groups attribute.

**Remote access permission**

If all conditions of a remote access policy are met, remote access permission is either granted or denied. You can use either the **Grant remote access permission** option or the **Deny remote access permission** option to set remote access permission for a policy.

Remote access permission is also granted or denied for each user account. The user remote access permission overrides the policy remote access permission. When remote access permission on a user account is set to the **Control access through Remote Access Policy** option, the policy remote access permission determines whether the user is granted access.

Granting access through either the user account permission setting or the policy permission setting is only the first step in accepting a connection. The connection attempt is subject to the settings of both the user account dial-in properties and policy profile properties. If the connection attempt does not match the settings of the user account or policy profile properties, the connection attempt is rejected.

**Profile**

A remote access policy profile is a set of properties that are applied to a connection when it is authorized--either through the user account or policy permission settings. A profile consists of the following groups of properties:

- Dial-in constraints
- IP
- Multilink

- Authentication
- Encryption
- Advanced

**Dial-in constraints**

You can set the following dial-in constraints:

- **Minutes server can remain idle before it is disconnected**

  The time after which a connection is disconnected when there is no activity. By default, this property is not set. Additionally, the Routing and Remote Access service does not disconnect an idle connection.

- **Minutes client can be connected**

  The maximum amount of time that a connection is connected. The connection is disconnected by the access server after the maximum session length. By default, this property is not set and there is no maximum session limit.

- **Allow access only on these days and at these times**

  The days of the week and hours of each day that a connection is allowed. If the day and time of the connection attempt do not match the configured day and time limits, the connection attempt is rejected. By default, this property is not set and there are no day or time limits. The Routing and Remote Access service does not disconnect active connections that were previously connected at a time when connection attempts were allowed.

- **Allow access to this number only**

  The specific phone number that a caller must call in order for a connection to be allowed. If the dial-in number of the connection attempt does not match the configured dial-in number, the connection attempt is rejected. By default, this property is not set and all dial-in numbers are allowed.

- **Allow access only through these media**

  The specific types of media, such as modem (also known as async), ISDN, virtual private network (known as virtual), or 802.11 wireless that a caller must use in order for a connection to be allowed. If the dial-in medium of the connection attempt does not match the configured dial-in media, the connection attempt is rejected. By default, this property is not set and all media types are allowed.

For information about setting dial-in constraints on a profile, see Configure dial-in constraints.

**IP**

You can set IP properties that specify IP address assignment behavior. You have the following options:

- The access server must supply an IP address.
- The access client can request an IP address.
- IP address assignment is determined by the access server (this is the default setting).
- A static IP address is assigned. A static IP address assigned to the user account overrides this setting. The IP address assigned is typically used to accommodate vendor-specific attributes for IP addresses.

You can also use the **IP** tab to define IP packet filters that apply to remote access connection traffic. This is intended for use with the Routing and Remote Access service. You can configure IP traffic that is allowed to remote access clients (output filters) or from remote access clients (input filters) on an exception basis. Either all traffic is allowed, except traffic specified by filters; or all traffic is blocked, except traffic specified by filters . Remote access policy profile filtering applies to all remote access

connections that match the remote access policy.

Filters configured on the **IP** tab are applied to client connections on a Windows Server 2003 remote access server, but are not applied to dial-on-demand connections.

For information about setting IP options on a profile, see Configure IP options.

**Multilink**
You can set Multilink properties that both enable Multilink and determine the maximum number of ports that a Multilink connection can use. Additionally, you can set Bandwidth Allocation Protocol (BAP) policies that both determine BAP usage and specify when extra BAP lines are dropped. The Multilink and BAP properties are specific to the Routing and Remote Access service. By default, Multilink and BAP are disabled.

The Routing and Remote Access service must have Multilink and BAP enabled in order for the Multilink properties of the profile to be enforced. For more information about enabling Multilink and BAP for the Routing and Remote Access service, see Enable Multilink and Enable BAP and BACP.

For information about setting Multilink options on a profile, see Configure multilink options.

**Authentication**
You can set authentication properties to both enable the authentication types that are allowed for a connection and specify the EAP type that must be used. Additionally, you can configure the EAP type. By default, Microsoft Encrypted Authentication (MS-CHAP) and Microsoft Encrypted Authentication version 2 (MS-CHAP v2) are enabled. With Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; you can specify whether users can change their expired passwords using MS-CHAP and MS-CHAP v2 (enabled by default). For more information, see Authentication methods.

The Routing and Remote Access service must have the corresponding authentication types enabled in order for the authentication properties of the profile to be enforced. For more information about enabling authentication methods for the Routing and Remote Access service, see Enable authentication protocols.

For information about setting authentication options on a profile, see Configure authentication.

**Encryption**
You can set encryption properties for the following encryption strengths:

- **No Encryption**

  This option allows a nonencrypted connection. To require encryption, clear this option.

- **Basic**

  For dial-up and PPTP-based VPN connections, Microsoft Point-to-Point Encryption (MPPE) is used with a 40-bit key. For L2TP/IPSec VPN connections, 56-bit Data Encryption Standard (DES) encryption is used.

- **Strong**

  For dial-up and PPTP VPN connections, MPPE is used with a 56-bit key. For L2TP/IPSec VPN connections, 56-bit DES encryption is used.

- **Strongest**

  For dial-up and PPTP VPN connections, MPPE is used with a 128-bit key. For L2TP/IPSec VPN connections, triple DES (3DES) encryption is used.

These settings are intended to be used with Routing and Remote Access. When other access servers are used, select **No Encryption** and clear all encryption strengths.

For information about setting encryption options on a profile, see Configure encryption.

**Advanced**
You can set advanced properties to specify the series of RADIUS attributes that are sent back by the IAS server to be evaluated by the RADIUS client. RADIUS attributes are specific to performing RADIUS authentication, and are ignored by a server that is

running the Routing and Remote Access service and is configured for Windows authentication. By default, **Framed-Protocol** is set to **PPP** and **Service-Type** is set to **Framed**.

The only advanced attributes that are used by the Routing and Remote Access service are **Account-Interim-Interval**, **Framed-Protocol**, **Framed-MTU**, **Reply-Message**, and **Service-Type**.

For information about setting advanced options on a profile, see Add RADIUS attributes to a remote access policy.

**Notes**

- Some elements of a remote access policy correspond to RADIUS attributes that are used during RADIUS-based authentication. For remote access policies on an Internet Authentication Service (IAS) server, verify that the access servers used are sending RADIUS attributes that correspond to the configured remote access policy conditions and profile settings. If a NAS does not send a RADIUS attribute that corresponds to a remote access policy condition or profile setting, then all RADIUS authentications from that NAS are denied.

- You can only use the Generate-Session-Timeout attribute if your user account database is a Security Accounts Manager (SAM) database or is the user account database for an Active Directory domain. When the value of **Generate-Session-Timeout** is set to **True**, the **ForceLogoff** value for a SAM database should be set to **0**. In the Local Security Settings console, ForceLogoff is changed to zero when **Network security: Force logoff when logon hours expire** is enabled. For information about how to change the ForceLogoff value for the SAM database at the command prompt, see Net accounts.

- You can configure wireless connection policy so that wireless clients periodically reauthenticate. This ensures that the client Wired Equivalent Privacy (WEP) encryption keys are changed often enough to provide adequate security for the wireless connection. To configure reauthentication, set the session timeout in your remote access policy or connection request policy for wireless connections (using the Session-Timeout attribute) to the required interval (for example, 10 minutes). Additionally, configure the Termination-Action attribute with the **Attribute value** set to **RADIUS-Request**. If the Termination-Action attribute is not set to **RADIUS-Request**, wireless access points might end the connection during reauthentication. For more information, see your hardware documentation.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Accepting a connection attempt

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accepting a connection attempt**

When a user attempts a connection, the connection attempt is accepted or rejected, based on the following logic:

1. The first policy in the ordered list of remote access policies is checked. If there are no policies, reject the connection attempt.

2. If all conditions of the policy do not match the connection attempt, go to the next policy. If there are no more policies, reject the connection attempt.

3. If all conditions of the policy match the connection attempt, check the value of the **Ignore-User-Dialin-Properties** attribute.

4. If the **Ignore-User-Dialin-Properties** attribute is set to **False**, check the remote access permission setting for the user attempting the connection.

   - If **Deny access** is selected, reject the connection attempt.

   - If **Allow access** is selected, apply the user account and profile properties. If the connection attempt does not match the settings of the user account and profile properties, reject the connection attempt. If the connection attempt matches the settings of the user account and profile properties, accept the connection attempt.

   - If the remote access permission is not set to **Allow access** or **Deny access**, the remote access permission must be set to **Control access through Remote Access Policy**. Check the remote access permission setting of the policy.

     If **Deny remote access permission** is selected, reject the connection attempt.

     If **Grant remote access permission** is selected, apply the user account and profile properties. If the connection attempt does not match the settings of the user account and profile properties, reject the connection attempt. If the connection attempt matches the settings of the user account properties and profile, accept the connection attempt.

5. If the **Ignore-User-Dialin-Properties** attribute is set to **True**, check the remote access permission setting of the policy.

   - If **Deny remote access permission** is selected, reject the connection attempt.

   - If **Grant remote access permission** is selected, apply the profile properties. If the connection attempt does not match the settings of the profile properties, reject the connection attempt. If the connection attempt matches the settings of the profile properties, accept the connection attempt. The following illustration shows the logic of remote access policies.




**Notes**

- The profile and user account settings for the first matching policy are applied to the connection. If a connection does not match the profile or user account settings of the remote access policy, the additional remote access policies are not tried.

- A connection attempt might not match any of the remote access policies. If this is the case, the connection attempt is rejected regardless of the remote access permission setting on the user account.

- The remote access polices are tried in order. The more specific remote access policies are typically placed in order ahead of the more general remote access policies.

- The **Ignore-User-Dialin-Properties** attribute is a new feature for Windows Server 2003, Standard Edition ; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition that allows you to ignore all of the dial-in properties of a user account. For more information, see New features for IAS.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

For examples of how different connection attempts are processed, see Remote Access Policies Examples.

# Remote Access Policies Examples

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access policies examples**

Remote access policies provide a tremendous amount of flexibility in accepting connections based on policy evaluation order, policy conditions, user and policy dial-in permissions, and both policy profile and user account properties.

This section describes some examples in which remote access policies are used:

- Grant remote access per user
- Allow dial-up connection using group membership
- Restrict connections to a maximum session time
- Wireless access example
- Wireless access with secure password authentication
- Ethernet switch access
- Force VPN clients to use strongest encryption
- Apply packet filters for business partner extranet
- Send NAS-specific RADIUS attributes to a RADIUS client

# Grant remote access per user

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Grant remote access per user**

In this example, the network administrator wants to maintain the same administrative model as used in Windows NT 4.0. Dial-in access is granted on a per-user basis by modifying the remote access permission setting of the user account dial-in properties. The remote access permission is set to either **Allow access** or **Deny access**. This is the default behavior for remote access policies.

You can create a policy named **Allow access if dial-in permission is enabled**. The settings on this policy are configured as follows:

- A single condition that consists of the **Day-And-Time-Restrictions** attribute is set for all times on all days.
- The **Deny remote access permission** option is selected.
- The profile is set to the default settings.

When a user attempts a remote access connection, the following logic is used to accept or reject the attempt (assuming that the default settings are set for all the other user account dial-in properties):

1. The policy **Allow access if dial-in permission is enabled** is evaluated.
2. The settings of the connection attempt match the conditions of the policy (the connection attempt occurs on any day or at any time).
3. On the user account, the setting in **Remote Access Permission (Dial-in or VPN)** is evaluated.
    - If **Allow access** is selected, the connection is accepted, subject to the settings of the user account dial-in and profile properties.
    - If **Deny access** is selected, the connection is rejected.

# Allow dial-up connection using group membership

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Allow dial-up connection using group membership**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants to allow dial-up connections for only those user accounts that belong to a specific set of groups. After remote access permission is set for all user accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common dial-up policy with the following settings:

   - Policy name: Dial-up access for allowed groups
   - Access Method: Dial-up access
   - User or Group: Select **Group**, and then specify the set of groups for which dial-up access is allowed.
   - Authentication methods: Select **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** and **Microsoft Encrypted Authentication**.
   - Policy Encryption Level: Select all check boxes.

     For more information, see Add a remote access policy.

2. Delete the default policies.

   For more information, see Delete a remote access policy.

**Notes**

- It is not necessary to have a separate remote access policy for each group. You can specify multiple groups or you can use nested groups to both delegate administration of and consolidate group membership. For a remote access or IAS server in a Windows 2000 native domain or a Windows Server 2003 domain, you can use universal groups. For more information, see Domain and forest functionality.
- You cannot use built-in or local domain groups when specifying group names.

# Restrict connections to a maximum session time

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restrict connections to a maximum session time**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

To conserve limited remote access server connection resources, the network administrator wants to limit the maximum session time for all dial-up users (members of the DialUpUsers group) to 20 minutes.

To implement this restriction through a remote access policy, the administrator completes the following steps:

**Use the New Remote Access Policy Wizard to create a custom policy with the following settings:**

- Policy name: Maximum session time of 20 minutes
- Conditions: NAS-Port-Type matches Async (Modem); Windows-Groups matches DialUpUsers
- Permission: Grant remote access permission
- Profile settings, **Dial-in Constraints** tab: Select the **Minutes client can be connected** check box, and then type **20**.

For more information, see Add a remote access policy.

**Delete the default policies.**
For more information, see Delete a remote access policy.

# Wireless access example

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless access example**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator requires that all IEEE 802.11 wireless clients (members of the WirelessClients group) use certificate-based authentication through the EAP-TLS authentication protocol and 128-bit encryption. Wireless clients that do not have a user certificate installed are placed on a specific virtual local area network (VLAN) from which they can obtain a user certificate. This unauthenticated connection lasts for 10 minutes.

In this case, two policies are needed:

1. A policy that requires wireless connections to use EAP-TLS.

2. A policy that accepts unauthenticated wireless connections and configures the wireless access point for a VLAN with an ID of 3. VLAN 3 contains a certificate server from which the wireless client can obtain a user certificate.

After remote access permission is set for all user accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common wireless policy with the following settings:

   - Policy name: Wireless access

   - Access Method: Wireless access

   - User or Group: Select **Group**, and then specify the WirelessClients group (example).

   - Authentication methods: Select **Smart Card or other Certificate**.

   - Policy Encryption Level: Check the **Strongest encryption (MPPE 128-bit)** check box, and then clear all other check boxes.

2. Use the New Remote Access Policy Wizard to create a custom wireless policy with the following settings:

   - Policy name: First-time wireless access

   - Conditions: NAS-Port-Type matches Wireless-Other or Wireless-IEEE 802.11; Windows-Groups matches WirelessUsers; Authentication-Type matches Unauthenticated

   - Permission: Grant remote access permission

   - Profile settings, **Dial-in Constraints** tab: Select the **Minutes client can be connected** check box, and then type **10**.

   - Profile settings, **Advanced** tab: Add the **Tunnel-Type** attribute with the value of **Virtual LANs (VLAN)**; Add the **Tunnel-Pvt-Group-ID** attribute with the value of **3**.

3. Delete the default policies.

   For more information, see Delete a remote access policy.

**Notes**

- For this example, a public key infrastructure (PKI) that has allocated the required computer certificate to the IAS server is being used. Additionally, the wireless access point is configured to use both an IAS server as its RADIUS server and EAP. For more information about wireless access deployment, see Checklist: Configuring the IAS server and wireless access points for wireless access.

- You can configure wireless connection policy so that wireless clients periodically reauthenticate. This ensures that the client Wired Equivalent Privacy (WEP) encryption keys are changed often enough to provide adequate security for the wireless connection. To configure reauthentication, set the session timeout in your remote access policy or connection request policy for wireless connections (using the Session-Timeout attribute) to the required interval (for example, 10 minutes). Additionally, configure the Termination-Action attribute with the **Attribute value** set to **RADIUS-Request**. If the Termination-Action attribute is not set to **RADIUS-Request**, wireless access points might end the connection during reauthentication. For more information, see your hardware documentation.

# Wireless access with secure password authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless access with secure password authentication**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator requires that all IEEE 802.11 wireless clients (members of the *WirelessClientsGroupName*) use Protected Extensible Authentication Protocol, or PEAP, for secure password authentication on the network. PEAP uses Transport Level Security (TLS) to create an end-to-end (EAP client to EAP authenticator) encrypted connection after verifying the identity of the authenticator (typically, an IAS server). The IAS server authenticates itself to the client with a certificate that includes the Server Authentication purpose in Enhanced Key Usage (EKU) extensions.

Because PEAP is used with EAP-MS-CHAPv2 as the authentication type, user authentication is accomplished with password-based credentials that are typed by the user. PEAP with EAP-MS-CHAPv2, also called PEAP-EAP-MS-CHAPv2, eliminates the need to deploy certificates to wireless clients for the purpose of client computer and user authentication.

For more information, see Network access authentication and certificates.

**Note**

- You can also use PEAP with EAP-TLS (also known as PEAP-EAP-TLS), which provides strong security. PEAP-EAP-TLS uses a public key infrastructure with certificates, which are used for server authentication, and either smart cards or certificates for client computer and user authentication. When you use PEAP-EAP-TLS, client certificate information is encrypted, providing improved security over EAP-TLS without PEAP.

After remote access permission is set for all user accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common wireless policy with the following settings:

   - Policy name: Wireless access
   - Access Method: Wireless access
   - User or Group: Select **Group**, and then specify the WirelessClients group (example).
   - Authentication methods: Select **Protected Extensible Authentication Protocol (PEAP)**.
   - Policy Encryption Level: Check the **Strongest encryption (MPPE 128-bit)** check box, and then clear all other check boxes.

2. Delete the default policies.

   For more information, see Delete a remote access policy.

When you select PEAP as the authentication method, the default authentication type is EAP-MS-CHAPv2. The following defaults are used:

- **Number of authentication retries:** 0
- **Allow client to change password** (not selected)

To change these settings, complete the steps in Configure PEAP and EAP methods.

**Notes**

- Although the deployment of certificates for user and client computer authentication is more difficult than using PEAP with EAP-MS-CHAPv2, certificates are more secure and PEAP with EAP-TLS with smart cards or certificates is the recommended authentication method for wireless clients.

- You can configure wireless connection policy so that wireless clients periodically reauthenticate. This ensures that the client Wired Equivalent Privacy (WEP) encryption keys are changed often enough to provide adequate security for the wireless connection. To configure reauthentication, set the session timeout in your remote access policy or connection request policy for wireless connections (using the Session-Timeout attribute) to the required interval (for example, 10 minutes). Additionally, configure the Termination-Action attribute with the **Attribute value** set to **RADIUS-Request**. If the Termination-Action attribute is not set to **RADIUS-Request**, wireless access points might end the connection during reauthentication. For more information, see your hardware documentation.

# Ethernet switch access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ethernet switch access example**

In this example, the network administrator is managing authorization by using groups. All computer accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants all Ethernet switch clients (members of the SwitchClients group) to use password-based authentication through the Extensible Authentication Protocol (EAP) and the Message Digest 5 (MD5) Challenge EAP type. After remote access permission is set for all computer accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common dial-up policy with the following settings:

   - Policy name: Ethernet access for switch clients
   - Access Method: Ethernet
   - User or Group: Select **Group**, and then specify the SwitchClients group (example).
   - Authentication methods: Select the **MD5-Challenge** EAP type.

     For more information, see Add a remote access policy.

2. Delete the default policies.

   For more information, see Delete a remote access policy.

**Note**

- This example assumes that the Ethernet switch has been configured to use both EAP and an IAS server as its RADIUS server. For more information about Ethernet switch deployment, see Ethernet switch access with IAS

# Force VPN clients to use strongest encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Force VPN clients to use strongest encryption**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants all virtual private network (VPN) clients (members of the group VPNClients) to use strong encryption. For PPTP connections, 128-bit Microsoft Point-to-Point Encryption (MPPE) must be used. For L2TP over IPSec connections, Triple DES (3DES)must be used. After remote access permission is set for all user accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common VPN policy with the following settings:

   - Policy name: VPN access with strongest encryption
   - Access Method: VPN access
   - User or Group: Select **Group**, and then specify the VPNClients group (example).
   - Authentication methods: Select **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** and **Microsoft Encrypted Authentication**.
   - Policy Encryption Level: Select the **Strongest encryption (IPSec triple DES or MPPE 128-bit)** check box, and then clear all other check boxes.

     For more information, see Add a remote access policy.

2. Delete the default policies.

   For more information, see Delete a remote access policy.

# Apply packet filters for business partner extranet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Apply packet filters for business partner extranet**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants to restrict access for all business partner virtual private network (VPN) connections to the resources of an extranet--a specific network segment that contains Web and file servers. In this example, the extranet is the subnet of 192.168.47.0/24. All business partner user accounts are members of the Partners group.

After remote access permission is set for all user accounts, the administrator completes the following steps.

**Use the New Remote Access Policy Wizard to create a custom policy with the following settings:**

- Policy name: Business partner connections to extranet
- Conditions: Windows-Groups matches Partners
- Permission: Grant remote access permission
- Profile settings, **Dial-in Constraints** tab: Select **Allow access only through these media** and the **Virtual (VPN)** media type.
- Profile settings, **IP** tab, Input packet filter:
    - Deny all traffic except those listed below
    - Destination network, IP address: 192.168.47.0
    - Destination network, Subnet mask: 255.255.255.0
    - Protocol: Any
- Profile settings, **IP** tab, Output packet filter:
    - Deny all traffic except those listed below
    - Source network, IP address: 192.168.47.0
    - Source network, Subnet mask: 255.255.255.0
    - Protocol: Any
- Profile settings, **Authentication** tab: Select both **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** and **Microsoft Encrypted Authentication**. For both of these authentication types, select **User can change password after it has expired**.
- Clear all other check boxes.
- Profile settings, **Encryption** tab: Select the **Strongest encryption** check box, and then clear all other check boxes

  For more information, see Add a remote access policy.

**Delete the default policies.**

For more information, see Delete a remote access policy.

# Send NAS-specific RADIUS attributes to a RADIUS client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Send NAS-specific RADIUS attributes to a RADIUS client**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants to allow dial-up connections for only those user accounts in the DialUpUsers group and send a series of network access server (NAS) attributes that are specific to the Ascend NAS. After remote access permission is set for all user accounts, the administrator completes the following steps:

1. When the NAS is added as a RADIUS client, configure its manufacturer.

   For more information about configuring a RADIUS client for an IAS server, see Add RADIUS clients.

2. Use the New Remote Access Policy Wizard to create a custom policy with the following settings:

   - Policy name: Allow dial-up connections and send attributes to Ascend NAS
   - Conditions: Client-Vendor matches Ascend Communications, Inc.; NAS-Port-Type matches Async (Modem) ; Windows-Groups matches DialUpUsers
   - Permission: Grant remote access permission
   - Profile settings, **Advanced** tab: Add the appropriate Ascend RADIUS attributes with their correct values.

   For more information, see Add a remote access policy.

3. Delete the default policies.

   For more information, see Delete a remote access policy.

# Multilink and BAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multilink and BAP**

A server running Routing and Remote Access supports Multilink and the Bandwidth Allocation Protocol (BAP). With Multilink, multiple physical links appear as a single logical link over which data is sent and received. A good example is the aggregation of both B channels of an ISDN Basic Rate Interface (BRI) connection. Multilink is the recommended method of combining multiple B channels of a BRI connection because the support for bonding, the combining of ISDN B channels through hardware support, is specific to the ISDN adapter. You can use Multilink for any ISDN adapter. Multilink must be supported on both sides of the connection.

While Multilink allows for multiple physical links to be aggregated, Multilink does not provide a mechanism to adapt to changing bandwidth conditions by adding extra links when needed or terminating extra links when unneeded. This additional capability is provided by the Bandwidth Allocation Protocol (BAP).

For example, a Multilink and BAP-enabled remote access client and remote access server create a Multilink connection that consists of a single physical link. As the utilization of the single link rises to a configured level, the remote access client uses a BAP request message to request an additional link. The BAP request message specifies the type of link desired, such as analog phone, ISDN, or X.25. The remote access server then sends a BAP response message that contains the phone number of an available port on the remote access server of the same type as specified by the remote access client in the BAP request.

**Enabling Multilink and BAP**

To enable Multilink and BAP, you must do the following:

1. Enable dynamic bandwidth control on the remote access server by using Routing and Remote Access. For more information, see Enable BAP and BACP. Multilink and BAP are enabled by default.

2. Enable Multilink on the appropriate remote access policy. For more information, see Configure multilink options. By default, the policy profile is configured to use the settings of the remote access server.

3. Configure Multilink and BAP behavior on the remote access client. For more information, see Configuring multiple device dialing.

To set the phone number on a port that is sent back to a BAP-enabled client, see Set the phone number on a port.

### Note

- When Multilink and BAP are used in combination with the callback feature set to always call back to the same number, a concentrator must exist on the caller side that can distribute incoming calls to the same number on various ports.

# Logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logging**

A server running Routing and Remote Access supports three types of logging:

1. Event logging

   Event logging is the recording of events in the system event log. Event logging is typically used for troubleshooting or for notifying network administrators of unusual events. For more information, see Remote Access Troubleshooting tools.

2. Local authentication and accounting logging

   A server running Routing and Remote Access supports the logging of authentication and accounting information for remote access connections in local logging files when Windows Authentication or Windows Accounting is enabled. This logging is separate from the events recorded in the system event log. You can use the information that is logged to track remote access usage and authentication attempts. Authentication and accounting logging is especially useful for troubleshooting remote access policy issues. For each authentication attempt, the name of the remote access policy that either accepted or rejected the connection attempt is recorded.

   The authentication and accounting information is stored in a configurable log file or files stored in the *systemroot*\System32\LogFiles folder. The log files are saved in Internet Authentication Service (IAS) or database-compatible format, meaning that any database program can read the log file directly for analysis.

   To configure authentication and accounting logging, you must first enable either Windows Authentication or Windows Accounting. For more information, see Use Windows Accounting. Then, you can configure the type of activity to log (accounting or authentication activity) and log file settings such as log file format. For more information, see Configure logging.

   For information about the database-compatible format of the log file, see Interpreting database-import log files.

   For information about the IAS format of the log file, see IAS-Formatted Log Files.

3. RADIUS-based authentication and accounting logging

   A server running Routing and Remote Access supports the logging of authentication and accounting information for remote access connections at a Remote Authentication Dial-In User Service (RADIUS) server when RADIUS authentication and accounting are enabled. This logging is separate from the events recorded in the system event log. You can use the information that is logged on your RADIUS server to track remote access usage and authentication attempts. For more information, see Use RADIUS authentication and Use RADIUS accounting.

   If your RADIUS server is running IAS, then authentication and accounting information is logged in log files stored on the IAS server. For more information, see Logging user authentication and accounting requests.

You can also use the Netsh command-line utility available in the Windows Server 2003 family to collect detailed logs and information about a remote access connection. For instance, you can use the **netsh ras set tracing** command to enable or disable tracing for a specific component. For more information, see Netsh commands for remote access (ras).

# The remote access server and network management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The remote access server and network management**

A server running Routing and Remote Access can participate in a Simple Network Management Protocol (SNMP) environment as an SNMP agent. For this configuration, you need to install the SNMP service. For more information, see Install the SNMP service.

The remote access server records management information in various object identifiers of the Internet MIB II that is installed with the SNMP service. Objects in the Internet MIB II are documented in RFC 1213, "Management Information Base for Network Management of TCP/IP-based internets: MIB-II."

For more information about the SNMP service, see SNMP.

# Broadcast name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Broadcast name resolution**
If you enable broadcast name resolution on a server running Routing and Remote Access, a remote access client can resolve full computer names and NetBIOS names of computers and other resources on a remote network that does not have either a DNS or a WINS server configured. This feature requires no configuration of the remote access clients, and it is enabled by default on the remote access server.

When broadcast name resolution is enabled, the remote access server uses a NetBIOS over TCP/IP (NetBT) proxy to resolve names as follows:

1. A remote access client that must resolve a name to an IP address broadcasts a NetBIOS Name Query packet across all interfaces.

2. The remote access server receives the NetBIOS Name Query packet and checks its cache for the appropriate mapping of name to IP address.

   - If the remote access server has the appropriate mapping, it sends a NetBIOS Name Query Response packet to the remote access client.

   - If the remote access server does not have the appropriate mapping, it broadcasts the NetBIOS Name Query packet on all its LAN interfaces. The resource whose name is being resolved sends a NetBIOS Name Query Response packet to the remote access client.

   After the remote access server routes the NetBIOS Name Query Response packet from the resource to the remote access client, the mappings remain in the cache of the server for a limited period of time. The default period of time is 10 minutes.

3. After this point, the remote access client uses the IP address of the resource to send packets to that resource.

As a result, a remote access client can resolve names of resources on a remote network as if the remote access client were on the same LAN as the resources.

By default, broadcast name resolution allows remote access clients to resolve names on a remote network but not for resources on a LAN to resolve the names of remote access clients. Remote access servers with broadcast name resolution enabled do not, by default, broadcast NetBIOS Name Query packets to interfaces that connect to the Internet. Unless you change this behavior, a remote access client cannot resolve the name of another remote access client. If you have a network that averages 10 connections or fewer, you might want to change this behavior.

**A broadcast name resolution example**
Remote Access Client A has established a connection to the remote access server for a small business network. The remote access server has three LAN interfaces (LAN1, LAN2, and LAN3) and an external interface. One of the resources on the small business network is Network Printer B, which is accessed through LAN3. Remote Access Client A has the NetBIOS name but not the IP address of Network Printer B. Remote Access Client A broadcasts a NetBIOS Name Query packet for Network Printer B. The remote access server broadcasts the NetBIOS Name Query packet to LAN1, LAN2, and LAN3 but not to the external interface. Network Printer B responds with a NetBIOS Name Query Response packet that is addressed to Remote Access Client A. The remote access server routes the NetBIOS Name Query Response packet to Remote Access Client A.

# Using Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using remote access**
This section covers:

- Installing the remote access server
- Managing remote access from the command line
- Installing Dial-Up Equipment
- Deploying Remote Access
- Remote Access Scenarios

# Installing the remote access server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing the remote access server**
To enable and configure the remote access server, you must be logged on as a member of the Administrators group.

**Hardware requirements**
Before you install the Routing and Remote Access service, all hardware should be installed and working. Depending on your network and your requirements, you might need the following hardware:

- Network adapter with a WHQL signed driver
- One or more compatible modems and an available COM port
- Multiport adapter for acceptable performance with multiple remote connections
- X.25 smart card (if you are using an X.25 network)
- ISDN adapter (if you are using an ISDN line)

To verify the compatibility of all hardware in a computer running Routing and Remote Access, see Support resources.

**Installing the software**
When you install a member of the Windows Server 2003 family, the remote access component is automatically installed. However, the Routing and Remote Access service is installed in a disabled state.

For more information, see Enable the Routing and Remote Access service.

# Managing remote access from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing remote access from the command line**

In addition to using the Routing and Remote Access snap-in, you can use command-line utilities to manage and troubleshoot remote access for users, computers, or services. For example, you can use Tracert to check the path to a remote access client and record the results. Used for diagnostic purposes only, Ping verifies IP-level connectivity to another TCP/IP computer. You can use Netsh overview to perform more advanced management tasks, such as server configuration and scripting.

For information about other command-line utilities, see Command-line reference A-Z. For more information about managing remote access with Netsh, see Netsh commands for remote access.

# Installing Dial-Up Equipment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing dial-up equipment**

This section provides information about the types of dial-up and remote access equipment used by the remote access server:

- Modems
- ISDN
- X.25
- ATM

# Modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modems**
Modems are the most commonly used dial-up equipment and offer up to 56,000 bps over an analog phone line. To ensure that your modems work with the Windows Server 2003 family, see Support resources. Microsoft has tested and verified the modems on this list with a server running Routing and Remote Access.

Modems from different manufacturers, and even different models from one manufacturer, may be incompatible in some settings and circumstances. Even modems that claim to follow the Hayes AT command set may, at times, be unable to communicate with other Hayes-compatible modems.

For general information about modems, see Installing and configuring Modems.

For information about troubleshooting modems, see Troubleshooting modems.

**Connecting without a modem**
You can connect two computers without a modem through a direct serial, direct parallel, or infrared connection. Although a direct connection eliminates the need for a network adapter, it is a slow link.

For more information, see:

- Using direct connections
- Establishing infrared network connections between computers

**Modem-pooling equipment**
The Windows Server 2003 family works with a variety of modem-pooling equipment.

For more information, see Configure modem-pooling equipment.

# ISDN

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ISDN**

Integrated Services Digital Network (ISDN) offers a much faster communication speed than an analog telephone line. The analog telephone line can communicate at speeds up to 56 kilobits per second (Kbps), whereas ISDN communicates at speeds of 64 or 128 Kbps. Businesses that need this kind of speed usually have a large telecommuting workforce or need to do extensive administrative tasks remotely, such as install software on off-site workstations.

**Installing an ISDN adapter**

A primary rate interface (PRI) ISDN line comes with two B channels and one D channel. Each B channel transmits data at 64 Kbps. The D channel is for signaling and transmits data at 16 Kbps.

You need to install ISDN adapters on the server and on each client. Then you can configure each B channel to act as a port, or you can configure all B channels to act as a single port. For instructions on installing an ISDN adapter, see the adapter's documentation. If you cannot configure all of the B channels on your ISDN adapter to act as a single port, then use Multilink to aggregate the multiple B channels into a single logical port.

If your business has a large number of people calling in to the remote access server, you can configure each channel to operate as a separate port. This configuration allows the greatest number of people to call in.

However, if your business has few people calling in, but needs more data-transmission speed, you can configure both channels to act as a single port. With this configuration, line speed increases to 128 Kbps. If you have installed more than one adapter, you can combine the channels on each adapter and get even faster transmission speed.

For information about installing an ISDN adapter, see To install an Install an ISDN adapter.

# X.25

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**X.25**

X.25 is a packet-switching technology that transmits data reliably across a packet-switching network. You can reach the X.25 network through a direct connection or through an asynchronous connection that dials in to a packet assembler/disassembler (PAD). A server running Routing and Remote Access only supports direct connections to X.25 networks by using an X.25 smart card.

After you install and enable the Routing and Remote Access service, you can set up the remote access server for an X.25 network. For more information, see Configure an X.25 smart card.

**Notes**

- The remote access server does not support callback on X.25 networks.
- An X.25 smart card is not related to a smart card that is used for authentication.

# ATM

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ATM**

Asynchronous transfer mode (ATM) is a fixed-length, cell-based, packet-switching technology that transmits data across an ATM network. You can reach the ATM network through a direct connection such as digital subscriber line (DSL) or cable modem, or through an ATM on-demand connection. The Routing and Remote Access service supports Point-to-Point Protocol (PPP) connections over switched virtual circuit (SVC) or permanent virtual circuit (PVC) ATM connections.

PPP connections over ATM require an ATM adapter with the appropriate support for the physical connection to the ATM network. Once installed, the ATM adapter appears as a dial-up device. For more information, see Devices and ports.

For more information on ATM and ATM support in the Windows Server 2003 family, see Windows ATM Services.

# Deploying Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying remote access**
This section covers:

- Using smart cards for remote access
- Setting Up Dial-Up Remote Access
- Setting Up Remote Access VPNs
- Remote Access Authentication
- Using RADIUS

# Using smart cards for remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using smart cards for remote access**

The use of smart cards for user authentication is the strongest form of authentication in the Windows Server 2003 family. For remote access connections, you must use the Extensible Authentication Protocol (EAP) with the **Smart card or other certificate (TLS)** EAP type, also known as EAP-Transport Level Security (EAP-TLS). To use smart cards for remote access authentication, you must do the following:

- Configure remote access on the remote access server.
- Install a computer certificate on the remote access server computer.
- Configure the **Smart card or other certificate (TLS)** EAP type in remote access policies.
- Enable smart card authentication on the dial-up or VPN connection on the remote access client.

**Configuring the remote access server to provide remote access**

You can configure the server running Routing and Remote Access as described in Using the remote access server as a corporate remote access server for dial-in remote access or Deploying Remote Access VPNs for VPN remote access.

**Installing a computer certificate on the remote access server**

In order to configure EAP-TLS on the remote access server, you must install a computer certificate, also known as a machine certificate. To install a computer certificate on the remote access server, a certification authority must be present to issue certificates. Once the certification authority is configured, you can install a certificate on the remote access server in two different ways:

1. By configuring the automatic allocation of computer certificates to computers in a domain.
2. By using Certificate Manager to obtain a computer certificate.

Based on the certificate policies in your organization, you only need to perform one of these two allocations.

To configure a certification authority and install the computer certificate, perform the following steps:

1. Install the Certificate Services component as an enterprise root certification authority. This step is only necessary if you do not already have an enterprise root certification authority (CA). For more information, see Install an enterprise root certification authority.
2. Do one of the following:
   - For auto-enrollment of computer certificates, configure the domain. For more information, see Configure automatic certificate allocation from an enterprise CA.

     To create a computer certificate for the remote access server computer that is a member of the domain for which auto-enrollment is configured (as well as other computers that are members of the domain), restart the remote access server or type **secedit /refreshpolicy machine_policy** from the command prompt.
   - To manually enroll computer certificates, use Certificate Manager to install the CA root certificate. For more information, see Manage certificates for a computer and Request a certificate.

**Important**

- EAP-TLS authentication requires that the remote access client computer certificate is configured with the Client

Authentication purpose and the remote access server computer certificate is configured with the Server Authentication purpose. A certificate purpose is identified using an object identifier (also known as OID). EAP-TLS checks the EKU extension for the client computer certificate to determine whether the Client Authentication purpose is present. When the EKU extension includes the Client Authentication purpose, EAP-TLS can use the certificate for authentication. For more information, see Network access authentication and certificates.

**Enabling a smart card logon process for the domain**

To enable a smart card logon process for the domain, you can perform the following procedures:

1. Prepare a certification authority to issue smart card certificates

2. Prepare a smart card certificate enrollment station

**Configuring the remote access server for smart card remote access**

To configure the server running Routing and Remote Access for smart card remote access, see Configure smart card remote access.

**Configuring the remote access client for smart card remote access**

You need to install a smart card reader on the remote access client computer. For more information, see Install a smart card reader on a computer.

Once a smart card reader is installed on the computer, you are prompted whether you want to use the smart card for authentication when you create dial-up or VPN connections.

For existing dial-up or VPN connections, you can enable smart card authentication on the properties of the dial-up or VPN connection. For more information, see Enable smart card or other certificate authentication.

# Setting Up Dial-Up Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up dial-up remote access**

This section covers:

- Dial-up remote access design considerations
- Dial-up remote access security
- Deploying Dial-Up Remote Access

# Dial-up remote access design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up remote access design considerations**
To prevent problems, you should consider the following design issues before you implement dial-up remote access connections.

**IP address allocation**
Determine whether the remote access server will use DHCP or a static IP address pool to obtain addresses for dial-up clients. If you use a static IP address pool, determine whether the pool will be ranges of addresses that are a subset of addresses from the IP network to which the server is attached or a separate subnet. If the static IP address pool address ranges represent a different subnet, ensure that routes to the address ranges exist in the routers of your intranet so that traffic to connected remote access clients is forwarded to the remote access server.

**Number of incoming ports needed**
Determine the maximum number of dial-up remote access clients that dial in at one time. Based on the number, you need to obtain modem bank equipment and phone lines that meet that need. Once the driver for the modem bank adapter is installed, verify that all of the ports of the modem bank device are configured to allow remote access. For more information, see Configure ports for remote access.

**Deciding on a remote access policy administrative model**
Before setting the dial-in permission on user accounts and creating remote access policies, you need to decide on a remote access policy administrative model. In the Windows Server 2003 family, there are two primary models for administering remote access permissions and connection settings:

1. Access by user.

2. Access by policy in an Active Directory domain.

Controlling access by policy centralizes many basic administrative tasks. For more information, see Introduction to remote access policies and Remote Access Policies Examples.

**Creating a remote access policy for dial-up remote access connections**
By using remote access policies, you can create a policy that requires dial-up connections to use a specific authentication method and encryption strength.

For example, you can create a group called Dial-up Users whose members are the user accounts of the users who are creating dial-up remote access connections. Then, you create a policy with two conditions on the policy: the **NAS-Port-Type** is set to all types except **Virtual (VPN)** and the **Windows-Groups** is set to *Dial-up Users* (the name of the group you created). Finally, you configure the profile for the policy to select a specific authentication method and encryption strength.

For more information, see Introduction to remote access policies.

**Using an IAS server for centralized authentication, authorization, and accounting**
If you have multiple remote access servers, you need to configure remote access policies and logging for each remote access server. If you want to take advantage of centralized remote access policies, accounting, and logging, configure the remote access servers as Remote Authentication Dial-In User Service (RADIUS) clients to a single server running the Internet Authentication Service (IAS) as a RADIUS server.

You should also use an IAS server if you have servers running Windows NT 4.0 and the Routing and Remote Access Service (RRAS) and you want to take advantage of remote access policies. You cannot configure remote access servers running Windows NT 4.0 as RADIUS clients. You must upgrade a remote access server running Windows NT 4.0 to a server running Windows NT 4.0 and RRAS.

For more information, see Using RADIUS for multiple remote access servers.

**Using Connection Manager**
For a large remote access deployment, you can use Connection Manager and the Connection Manager Administration Kit to provide a custom dialer with preconfigured connections to all remote access clients across your organization.

For more information about Connection Manager, see Connection Manager Administration Kit.

# Dial-up remote access security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up remote access security**
You can enhance dial-up remote access security through:

- Strong authentication
- Data encryption

**Strong authentication**
For authentication, use the strongest authentication scheme that is possible for your dial-up remote access configuration. The strongest authentication is the use of EAP-TLS with smart cards. For more information, see Using smart cards for remote access.

Otherwise, use MS-CHAP v2 authentication and enforce the use of strong passwords on your network. For more information, see MS-CHAP version 2.

**Data encryption**
For encryption, you can use either link encryption or end-to-end encryption:

- Link encryption encrypts the data only on the link between the remote access client and the remote access server, or the two servers running Routing and Remote Access. For dial-up remote access connections, you must use Microsoft Point-to-Point Encryption (MPPE) in conjunction with either MS-CHAP or EAP-TLS authentication.
- End-to-end encryption encrypts the data between the source host and its final destination. You can use IPSec to encrypt data from the remote access client to the destination host after the remote access connection is made.

To require encryption, clear the **No Encryption** option and select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile that is used by your dial-up remote access clients. For more information, see Configure encryption.

# Deploying Dial-Up Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying dial-up remote access**

This section covers:

- Using the remote access server as a corporate remote access server
- Using the remote access server for Internet access

# Using the remote access server as a corporate remote access server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the remote access server as a corporate remote access server**

You can use a server running Routing and Remote Access to provide dial-up access to your corporate intranet. If you want the remote access server to support multiple, dial-up networking TCP/IP-based connections, complete the following steps:

- Configure the connection to the intranet.
- Configure the connection to the dial-up networking clients.
- Configure the dial-in ports.
- Configure the remote access server.
- Configure multicast support.
- Configure remote access policies.

The following illustration shows the elements of a server running Routing and Remote Access that provides dial-up access to a corporate intranet.




**Configuring the connection to the intranet**

The connection to the intranet from Routing and Remote Access is a LAN adapter installed in the computer. To verify that the LAN adapter is compatible with Windows, see Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.
- Default gateway of a local router.
- IP address of DNS and WINS servers of corporate intranet servers.

**Configuring the connection to dial-up networking clients**

To allow the connection of multiple, simultaneous dial-up clients, you must have modem pooling equipment (hereafter known as the modem bank) with the appropriate connections to the local telecommunications provider.

Remote access hardware devices include an adapter that installs on the server running Routing and Remote Access. To verify that the hardware is compatible with Windows, see Support resources.

The modem bank adapter includes drivers that are installed on the server running Routing and Remote Access so that the modem bank appears as a device with multiple modem ports.

**Configuring the dial-in ports**

All the modem bank ports are listed as separate ports under **Ports** in Routing and Remote Access. You should configure all the modem bank ports for remote access.

For more information about configuring ports for remote access, see Configure ports for remote access.

**Configuring the remote access server**

You can configure the properties of the remote access server in Routing and Remote Access. For more information, see View properties of the remote access server.

To allow multiple dial-up networking clients to access the corporate intranet, you need to configure the following settings:

- **General**

  Verify that the **Remote access server** check box is selected.

- **Security**

  - **Authentication Methods**

    The server is configured by default to accept certain authentication methods. You can control which authentication methods you accept through remote access policies, except for CHAP.

  - **Authentication Provider**

    You can verify the credentials of dial-up networking clients by using Windows Authentication or a RADIUS server. If RADIUS is selected, you need to configure RADIUS server settings for your RADIUS server or RADIUS proxy.

  - **Accounting Provider**

    You can record dial-up client networking activity for analysis or accounting purposes by selecting and configuring an accounting provider.

- **IP**

  Verify that the **Enable IP routing** and **Allow IP-based remote access and demand-dial connections** check boxes are selected.

  If you are using a DHCP server to allocate IP addresses for remote access clients, click **Dynamic Host Configuration Protocol (DHCP)**. If not, click **Static address pool** and configure the ranges of IP addresses that are dynamically allocated to dial-up networking clients.

  If the static IP address pool consists of ranges of IP addresses that are for a separate subnet, then you need to either enable an IP routing protocol on the remote access server computer or add static IP routes for each range to all the routers on the other intranet. If the routes are not added, then remote access clients cannot receive traffic from resources on the intranet.

  For more information about configuring IP address pools, see Create a static IP address pool.

**Configuring multicast support**

Depending on the options selected when the Routing and Remote Access Server Setup Wizard was run, multicast support may already be enabled. To configure multicast support, you need to complete the following steps:

1. Add the IGMP Router and Proxy routing protocol. For more information, see Add the IGMP routing protocol.

2. Add the **Internal** interface to the IGMP routing protocol and configure it in IGMP router mode. For more information, see Enable IGMP router and IGMP proxy mode.

3. Add the interface that represents the permanent connection to the intranet to the IGMP routing protocol and configure the interface in IGMP proxy mode. For more information, see Enable IGMP router and IGMP proxy mode.

**Configuring remote access policies**

If you want to authorize remote access to the dial-up networking clients based on the access-by-user administrative model, do the following:

1. For a stand-alone remote access server, use Local Users and Groups and set dial-in properties to **Allow access** for those users who will be making remote access connections.

2. For a remote access server in an Active Directory domain, use Active Directory Users and Computers and set dial-in properties to **Allow access** for those users who will be making remote access connections.

If you want to grant remote access to the dial-up networking clients based on group membership and an access-by-policy administrative model, do the following:

1. For a stand-alone remote access server, use Local Users and Groups and set dial-in properties to **Control access through Remote Access Policy** for all users.

2. Do one of the following:

   - For a remote access server that is a member of a Windows 2000 mixed domain, use Active Directory Users and Computers and set dial-in properties to **Allow access** for all users.

   - For a remote access server that is a member of a Windows 2000 native or Windows Server 2003 domain, use Active Directory Users and Computers and set dial-in properties to **Control access through Remote Access Policy** for all users.

3. Create a group whose members will be able to create dial-up networking connections with the remote access server. For example, create a group named Remote_Users.

4. Add the appropriate user accounts to the new group.

5. Delete the default remote access policies.

6. Create a new remote access policy with the following properties:

   - Set **Policy name** to **Remote Access if member of RAS_Users** (example).

   - Set the **Windows-Groups** attribute to **RAS_Users** (example).

   - Select the **Grant remote access permission** option.

For more information, see Remote Access Policies Examples.

For encryption, the default setting allows Microsoft Point-to-Point Encryption (MPPE) when requested by the remote access client. To force encryption for dial-up networking connections, you need to modify the encryption settings on the policy profile to require encryption. For dial-up networking connections, clear the **No encryption** option and select the following options on the **Encryption** tab on the properties of the remote access policy profile:

- **Basic**

  You should use this option when communicating with Microsoft dial-up networking clients who are using lower encryption. This option uses Microsoft Point-to-Point Encryption (MPPE) and a 40-bit encryption key.

- **Strong**

  You should use this option when communicating with Windows Server 2003 family, Windows XP, Windows 2000, Windows Millennium Edition, and Windows 98 dial-up networking clients who are using moderate encryption. This option uses MPPE and a 56-bit encryption key.

- **Strongest**

You should use this option when communicating with dial-up networking clients who are using the highest level of encryption. This option uses MPPE and a 128-bit encryption key.

For more information, see Configure encryption.

# Using the remote access server for Internet access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the remote access server for Internet access**

You can use a server running Routing and Remote Access to provide both access to the Internet and traditional dial-up Internet service provider (ISP) services. If you want the remote access server to support multiple dial-up TCP/IP-based connections, complete the following steps:

- Configure the connection to the Internet.
- Configure the connection to the dial-up clients.
- Configure the dial-in ports.
- Configure the remote access server.
- Configure multicast support.
- Configure remote access policies.

The following illustration shows the elements of an ISP based on the Windows Server 2003 family.




**Configuring the connection to the Internet**

The connection to the Internet from Routing and Remote Access is a dedicated connection--a WAN adapter installed in the computer. The WAN adapter is typically a DDS, T1, Fractional T1, or Frame Relay adapter. As the ISP, you must contract with a local telephone company to run the appropriate physical wiring to your premises. You need to verify that the WAN adapter is compatible with Windows. (See Support resources.)

The WAN adapter includes drivers that are installed on the server running Routing and Remote Access so that the WAN adapter appears as a network adapter.

You need to configure the following TCP/IP settings on the WAN adapter:

- IP address and subnet mask assigned from the Internet Assigned Numbers Authority (IANA) or a downstream ISP.
- Default gateway of the downstream ISP or network access point (NAP) router.

**Configuring the connection to the dial-up clients**

To allow the connection of multiple, simultaneous dial-up clients, you must have modem pooling equipment (hereafter known as the modem bank) with the appropriate connections to the local telecommunications provider.

Typical modem banks include an adapter that installs on the server running Routing and Remote Access. You need to verify that the modem bank adapter is compatible with Windows. (See Support resources.)

The modem bank adapter includes drivers that are installed on the server running Routing and Remote Access so that the modem bank appears as a series of modem ports.

**Configuring the remote access server**

You can configure the properties of the remote access server in Routing and Remote Access. For more information, see View properties of the remote access server.

You can configure the remote access server to allow multiple dial-up clients access to the Internet during the Routing and Remote Access Server Setup wizard. If you did not configure this during setup, you need to manually configure the following settings:

- **General**

  Verify that the **Remote access server** check box is selected.

- **Security**

  - **Authentication Methods**

    Select the authentication methods that are supported by the remote access server to authenticate the credentials of dial-up clients. Microsoft dial-up networking clients typically use MS-CHAP authentication. Other dial-up networking clients might use CHAP, SPAP, and PAP authentication.

  - **Authentication Provider**

    You can verify the credentials of dial-up clients by using Windows Server 2003 family security features or a RADIUS server. If RADIUS is selected, you need to configure RADIUS server settings for your RADIUS server or RADIUS proxy.

  - **Accounting Provider**

    You can record dial-up client activity for analysis or accounting purposes by selecting and configuring an accounting provider.

- **IP**

  Verify that the **Enable IP routing** and **Allow IP-based remote access and demand-dial connections** check boxes are selected.

  Click **Static address pool** and configure the ranges of IP addresses that are dynamically allocated to dial-up networking clients. Or, if a DHCP server is available, click **Dynamic Host Allocation Protocol (DHCP)**.

  For more information about configuring IP address pools, see Create a static IP address pool.

**Configuring multicast support**

Depending on the options selected when the Routing and Remote Access Server Setup Wizard was run, multicast support might already be enabled. To configure multicast support manually, you need to complete the following steps:

1. Add the IGMP Router and Proxy routing protocol. For more information, see Add the IGMP routing protocol.

2. Add the **Internal** interface to the IGMP routing protocol and configure it in IGMP router mode. For more information, see Enable IGMP router and IGMP proxy mode.

3. Add the interface that represents the permanent connection to the Internet to the IGMP routing protocol and configure the interface in IGMP proxy mode. For more information, see Enable IGMP router and IGMP proxy mode.

**Configuring remote access policies**

If you want to authorize remote access to the dial-up networking clients by user, do the following:

1. For a stand-alone remote access server, use Local Users and Groups and set dial-in properties to **Allow access** for those users who will be making remote access connections.

2. For a remote access server in an Active Directory domain, use Active Directory Users and Computers and set dial-in properties to **Allow access** for those users who will be making remote access connections.

If you want to grant remote access to the dial-up networking clients based on group membership and an access-by-policy administrative model, do the following:

1. For a stand-alone remote access server, use Local Users and Groups and set dial-in properties to **Control access through Remote Access Policy** for all users.

2. Do one of the following:

    - For a remote access server that is a member of a Windows 2000 mixed domain, use Active Directory Users and Computers and set dial-in properties to **Allow access** for all users.

    - For a remote access server that is a member of a Windows 2000 native or Windows Server 2003 domain, use Active Directory Users and Computers and set dial-in properties to **Control access through Remote Access Policy** for all users.

3. Create a user group whose members will be able to create dial-up networking connections with the remote access server. For example, create a group named Clients.

4. Add the appropriate user accounts to the new group.

5. Delete the default remote access policies.

6. Create a new remote access policy with the following properties:

    - Set **Policy name** to **Remote Access if member of Clients** (example).

    - Set the: **Windows-Groups** attribute to **Clients** (example).

    - Select the **Grant remote access permission** option.

# Setting Up Remote Access VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up remote access VPNs**

This section covers:

- Remote access VPN design considerations
- Remote access VPN security considerations
- Deploying Remote Access VPNs

# Remote access VPN design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access VPN design considerations**

To prevent problems, you should consider the following design issues before you implement remote access VPN connections:

- Choosing between PPTP-based connections or L2TP/IPSec connections
- Installing certificates
- Configuring firewall packet filters
- Creating a remote access policy for remote access VPN connections
- Using an IAS server
- Using Connection Manager

For more information, see Remote access VPN connection.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Choosing between PPTP-based connections or L2TP/IPSec connections**

A VPN server running Windows Server 2003 provides support for both Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP). When choosing between PPTP-based and L2TP/IPSec remote access VPN solutions, consider the following:

- PPTP can be used with a variety of Microsoft clients including Windows 95 (with the Dial-up Networking Upgrade 1.3 and later), Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000, and Windows XP. PPTP does not require a public key infrastructure (PKI) to issue computer certificates. By using encryption, PPTP-based VPN connections provide data confidentiality (captured packets cannot be interpreted without the encryption key). PPTP-based VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data origin authentication (proof that the data was sent by the authorized user).
- L2TP can only be used with client computers running Windows 2000 or Windows XP. L2TP supports either computer certificates or a preshared key as the authentication method for Internet Protocol security (IPSec). Computer certificate authentication, the recommended authentication method, requires a PKI to issue computer certificates to the VPN server computer and all VPN client computers. By using IPSec, L2TP/IPSec VPN connections provide data confidentiality, data integrity, and data authentication.

**Installing certificates**

In order to create L2TP/IPSec remote access VPN connections using computer certificate authentication for IPSec, you must install computer certificates, also known as machine certificates, on the VPN client and the VPN server. For more information, see Network access authentication and certificates and Computer certificates for L2TP/IPSec VPN connections.

**Configuring firewall packet filters**

If you have a firewall, you must configure packet filters on the firewall to allow traffic between VPN clients on the Internet and the VPN server. For more information, see VPN servers and firewall configuration.

**Creating a remote access policy for remote access VPN connections**
By using remote access policies, you can create a policy that requires remote access VPN connections to use a specific authentication method and encryption strength.

For example, you can create an Active Directory group called VPN Users whose members are the user accounts of the users creating remote access VPN connections across the Internet. Then, you create a policy with two conditions on the policy: **NAS-Port-Type** is set to **Virtual (VPN)** and **Windows-Group** is set to **VPN Users**. Finally, you configure the profile for the policy to select a specific authentication method and encryption strength.

For more information, see Introduction to remote access policies.

**Using an IAS server**
If you have multiple VPN servers running Windows Server 2003 , you need to configure remote access policies and logging for each VPN server. If you want to take advantage of centralized remote access policies and logging, you can configure the VPN servers as Remote Authentication Dial-In User Service (RADIUS) clients to a single computer (a RADIUS server) running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; and Internet Authentication Service (IAS).

You should also use an IAS server if you have VPN servers running Windows NT 4.0 with the Routing and Remote Access Service (RRAS), and you want to take advantage of remote access policies available in products in the Windows Server 2003 family. You cannot configure VPN servers running Windows NT 4.0 as RADIUS clients. You must upgrade a VPN server running Windows NT 4.0 to a VPN server running Windows NT 4.0 and RRAS.

For more information, see Using RADIUS for multiple remote access servers.

**Using Connection Manager**
For a large remote access VPN deployment, you can use Connection Manager and the Connection Manager Administration Kit to provide a custom dialer with preconfigured VPN connections to all remote access clients across your organization.

For more information about Connection Manager, see Connection Manager Administration Kit. For information about implementing VPN support with the Connection Manager Administration Kit, see Implementing VPN support.

# Remote access VPN security considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access VPN security considerations**
You can enhance remote access VPN security through:

- L2TP connections
- Strong authentication
- Data encryption
- PPTP or L2TP/IPSec packet filtering
- Firewall packet filtering
- Multi-use VPN servers
- Preventing traffic routed from remote access clients
- Remote access policy profile packet filtering

For more information, see Remote access VPN connection.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**L2TP connections**
A VPN server running Windows Server 2003 provides support for both Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP). By using encryption, PPTP-based VPN connections provide data confidentiality (captured packets cannot be interpreted without the encryption key). PPTP-based VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data origin authentication (proof that the data was sent by the authorized user). By using Internet Protocol Security (IPSec), L2TP/IPSec VPN connections provide data confidentiality, data integrity, and data origin authentication. However, L2TP can only be used with Windows 2000 and Windows XP clients.

L2TP/IPSec connections can use the following methods to authenticate the IPSec security association:

- A computer certificate

  Computer certificates are the recommended method of authentication as they provide strong authentication and are very difficult to impersonate or compromise. Computer certificate authentication requires a public key infrastructure (PKI) to issue computer certificates to the VPN server computer and all VPN client computers.

- A preshared key

  A preshared key is a string of text that is configured on both the VPN client and VPN server. Preshared key is a relatively weak authentication method, therefore it is only recommended that you use preshared key authentication while your PKI is being deployed or when preshared key is required by the VPN server to which your VPN clients are connecting.

**Strong user authentication**
For user authentication, use the strongest authentication scheme that is possible for your remote access VPN configuration. The strongest authentication scheme is the use of EAP-TLS with certificates. For more information, see Deploying Certificate-based Authentication for VPN Connections and Network access authentication and certificates.

Otherwise, use MS-CHAP version 2 authentication and enforce the use of strong passwords on your network. For more information, see MS-CHAP version 2.

**Data encryption**
For encryption, you can use either link encryption or end-to-end encryption with link encryption:

- Link encryption encrypts the data only on the link between the VPN client and the VPN server. For PPTP connections, you must use Microsoft Point-to-Point Encryption (MPPE) in conjunction with either MS-CHAP, MS-CHAP v2, or EAP-TLS authentication. For L2TP/IPSec connections, IPSec provides encryption on the link between the VPN client and the VPN server.

- End-to-end encryption encrypts the data between the source host and its final destination. You can use IPSec *after* the VPN connection is made to encrypt data from the source host to the destination host.

To require encryption, select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile that is used by your remote access VPN clients. For more information, see Configure encryption.

Encryption on client computers is configured in the properties of the network connection.

**PPTP or L2TP/IPSec filtering**
To secure the VPN server from sending or receiving any traffic on its Internet interface except VPN traffic, you need to ensure that PPTP or L2TP/IPSec input and output filters are configured on the interface that corresponds to the connection to the Internet. For more information about PPTP filters, see Add PPTP Filters. For more information about L2TP/IPSec filters, see Add L2TP over IPSec Filters.

Because IP routing is enabled on the Internet interface, if PPTP or L2TP/IPSec filters are not configured on the Internet interface, then any traffic received on the Internet interface is routed, which may forward unwanted Internet traffic to your intranet.

**Firewall packet filtering**
It is a common practice to provide protection to intranet hosts, such as a VPN server, from Internet hosts using a firewall. If you have a firewall, you must configure packet filters on the firewall to allow traffic to and from VPN clients on the Internet and the VPN server. For more information, see VPN servers and firewall configuration.

**Multi-use VPN servers**
Due to the way in which routes are created on VPN remote access clients, it is possible for the VPN client to send some unencrypted traffic to the VPN server rather than through the encrypted tunnel of the VPN connection. For example, the VPN client might connect to other services running on the VPN server without routing traffic across the VPN connection. Only traffic that is destined to the public IP address of the VPN server is sent in the clear. If traffic from the client uses the internal or private interface IP address of the VPN server, however, it will be encrypted.

When a remote access VPN client creates a VPN connection with a VPN server, it creates a series of routes in the IP routing table on the VPN client:

- Default route that uses the VPN connection

  The new default route for the VPN connection effectively replaces the existing default route for the duration of the connection. After the connection is made, all traffic that does not match an address on the directly connected network or the address of the VPN server is sent in an encrypted form over the VPN connection.

- Host route to the VPN server that uses the local network connection

  The host route for the address of the VPN server is created so that the VPN server is reachable using the locally attached network. If the host route is not present, VPN traffic to the VPN server cannot be sent.

The result of having the host route for the VPN server is that all traffic that is sent to applications running on the VPN server is not sent across the VPN connection, but is instead sent in an unencrypted form across the network between the VPN client and VPN server.

For example, when a remote access VPN client creates a VPN connection with a VPN server and then accesses a shared file on the VPN server computer, that traffic is not sent using the VPN connection. The file sharing traffic is sent in plaintext over the network between the VPN client and VPN server.

Additionally, if packet filters are configured on the VPN server that only allow VPN connection traffic, all other traffic sent to the VPN server is discarded. In this situation, all attempts to connect to services running on the VPN server will fail because traffic attempting to connect to those services are not sent over the VPN connection.

The key to which address is used by the VPN client to access services running on the VPN server lies in the way that the name of the VPN server is resolved. Typical users and applications refer to network resources using names, rather than IP addresses. The name must be resolved to an IP address using either DNS or WINS. If the intranet DNS and WINS infrastructures never contain a record mapping the name of the VPN server to the public IP address of the VPN server, traffic to services running on the VPN server is always sent across the VPN connection.

To prevent the VPN server from dynamically registering the public IP address of its Internet interface in the intranet DNS, obtain properties of the **Internet Protocol (TCP/IP)** component of the Internet connection in the **Network Connections** folder. Click **Advanced**. In the **Advanced TCP/IP Settings** dialog box, click the **DNS** tab, and then clear the **Register this connection's addresses in DNS** check box.

To prevent the VPN server from registering the public IP address of its Internet interface with intranet WINS servers, obtain properties of the **Internet Protocol (TCP/IP)** component of the Internet connection in the **Network Connections** folder. Click **Advanced**. In the **Advanced TCP/IP Settings** dialog box, click the **WINS** tab, and then click **Disable NetBIOS over TCP/IP**.

Before the VPN connection is made, the VPN client uses the Internet DNS infrastructure to resolve the name of the VPN server computer to its public IP addresses. After the VPN connection is made, assuming that intranet DNS and WINS servers are configured either during the PPP connection process or through the relaying of the DHCPInform message, the VPN client uses the intranet DNS and WINS infrastructures to resolve the name of the VPN server computer to its intranet IP addresses.

**Preventing traffic routed through remote access clients**

Once a remote access client successfully establishes a PPTP or L2TP connection, by default any packet sent over the connection is received by the VPN server and forwarded. Packets sent over the connection can include:

- Packets originated from the remote access client computer.
- Packets sent to the remote access client computer by other computers.

When the remote access client computer makes the VPN connection, by default it creates a default route so that all traffic that matches the default route is sent over the VPN connection. If other computers are forwarding traffic to the remote access VPN client, treating the remote access client computer as a router, then that traffic may also be forwarded across the VPN connection. This is a security problem because the computer that is sending traffic to the remote access VPN client has not been authenticated by the VPN server. The computer sending traffic to the remote access VPN client computer has physical network access to the same resources as the authenticated remote access VPN client computer.

To prevent the VPN server from sending traffic across the VPN connection for computers other than authenticated remote access VPN client computers, configure the following remote access policy packet filters on the remote access policy that is used for your VPN connections:

- From client filter:

  Filter action set to **Deny all traffic except those listed below**.

| IP Packet Filter Field | Setting |
|---|---|
| Source Address | User's address |
| Source Network Mask | User's mask |
| Destination Address | Any |
| Destination Network Mask | Any |
| Protocol | Any |

- To client filter:

  Filter action set to **Deny all traffic except those listed below**.

| IP Packet Filter Field | Setting |
|---|---|
| Source Address | Any |
| Source Network Mask | Any |
| Destination Address | User's address |
| Destination Network Mask | User's mask |
| Protocol | Any |

This set of IP packet filters will cause the VPN server to discard all traffic sent across the VPN connection except those that either originated from or is sent to authenticated remote access VPN clients.

**Remote access policy profile packet filtering**

To define the specific types of IP traffic that are allowed for remote access VPN connections, you need to configure IP packet filters on the profile for the remote access policy that is used for your remote access VPN connections. With profile packet filters, you can configure the IP traffic that is allowed from remote access clients (**From client** filters) or to remote access clients (**To client** filters) on an exception basis: either all traffic except traffic specified by filters or no traffic except traffic specified by filters. Remote access policy profile filtering applies to all remote access connections that match the remote access policy. For more information, see Elements of a remote access policy and Configure IP options.

# Deploying Remote Access VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying remote access VPNs**

This section covers:

- PPTP-based remote access VPN deployment
- L2TP-based remote access VPN deployment

# PPTP-based remote access VPN deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPTP-based remote access VPN deployment**

You can use remote access with products in the Windows Server 2003 family to provide access to a corporate intranet for remote access clients who are making PPTP connections across the Internet. If you want your remote access server to support multiple PPTP connections, complete the following steps:

- Configure the connection to the Internet.
- Configure the connection to the intranet.
- Configure the remote access server as a corporate intranet router.
- Configure the VPN server.
- Configure firewall packet filters.
- Configure remote access policies.

The following illustration shows the elements of a remote access server running Windows Server 2003 that provides PPTP-based remote access to a corporate intranet.




For more information, see Remote access VPN connection and Point-to-Point Tunneling Protocol.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Configuring the connection to the Internet**

The connection to the Internet from a computer running a Windows Server 2003 operating system is a dedicated connection--a WAN adapter installed in the computer. The WAN adapter is typically a DDS, T1, Fractional T1, or Frame Relay adapter. You must contract with a local telephone company to run the appropriate physical wiring to your premises. You need to verify that the WAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

The WAN adapter includes drivers that are installed in memobers of the Windows Server 2003 family so that the WAN adapter appears as a network adapter.

You need to configure the following TCP/IP settings on the WAN adapter:

- IP address and subnet mask assigned from the InterNIC or an Internet service provider (ISP).
- Default gateway of the ISP router.

So that VPN clients can connect to your VPN server by name (rather than IP address), you can request DNS registration of your VPN server by your ISP.

**Configuring the connection to the intranet**
The connection to the intranet from a computer running a Windows Server 2003 operating system is a LAN adapter that is installed in the computer. You need to verify that the LAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.
- DNS and WINS name servers of corporate intranet name servers.

**Configuring the remote access server as a corporate intranet router**
In order for the remote access server to properly forward traffic on the corporate intranet, you must configure it as a router with either static routes or routing protocols so that all of the locations of the intranet are reachable from the remote access server. For information about routing concepts, see Routing Overview. For information about setting up the remote access server as a router, see Deploying Routing.

**Configuring the VPN server**
You can configure your VPN server by running the Routing and Remote Access Server Setup Wizard. You can configure the following settings using the wizard:

- A basic firewall on the public interface.
- The method by which the VPN server assigns addresses to remote access clients (either using addresses that the VPN server obtains from a DHCP server or using addresses from a specified range of addresses that you have configured on the VPN server).
- Forwarding of authorization and authentication messages to a Remote Access Dial-In User Service (RADIUS) server (configuration of the VPN server as a RADIUS client).

Once the wizard is run, these Routing and Remote Access settings are automatically configured:

- Network interfaces
- PPTP and L2TP ports (5 or 128 of each, depending on your choices when running the wizard)
- Multicast support using Internet Group Messaging Protocol (IGMP)
- IP routing
- Installation of the DHCP Relay Agent component.

For more information on enabling Routing and Remote Access and running the wizard, see Enable the Routing and Remote Access service.

**Configuring firewall packet filters**
If you are using a firewall, you need to configure PPTP packet filters on your firewall to allow PPTP traffic between Internet-based VPN clients and the VPN server computer. For more information, see VPN servers and firewall configuration.

**Configuring remote access policies**
For an access-by-user administrative model, you need to set the remote access permission to **Allow access** on the user accounts for those users who will be making VPN connections. For an access-by-policy model, make the appropriate changes to the remote access permission of the user accounts. For more information, see Introduction to remote access policies.

To configure a remote access policy to control the authentication and encryption options for VPN connections, you can create a remote access policy with the following settings:

- Set **Policy name** to **VPN Access** (example).
- For conditions, set the **NAS-Port-Type** condition to **Virtual (VPN)** and the **Tunnel-Type** condition to **Point-to-Point Tunneling Protocol**.
- Select the **Grant remote access permission** option.
- For profile settings, select the appropriate authentication and encryption options.

Then, either delete the default remote access policies or position them after the new policy. This remote access policy allows all users with remote access permission to create a VPN connection.

If you want to distinguish dial-up remote access users from VPN remote access users, do the following:

1. Create an Active Directory group whose members can create virtual private networking connections with the VPN server. For example, VPN_Users.
2. Add the appropriate user accounts to the new Active Directory group.
3. Create a new remote access policy with the following properties:
    - Set **Policy name** to **VPN Access if member of VPN_Users** (example).
    - For conditions, set the **Windows-Groups** condition to **VPN_Users** (example), set the **NAS-Port-Type** condition to **Virtual (VPN)**, and set the **Tunnel-Type** condition to **Point-to-Point Tunneling Protocol**.
    - Select the **Grant remote access permission** option.
4. Move the default remote access policies after the new policy.

The default encryption settings allow no encryption and all levels of encryption strength. To require encryption, clear the **No Encryption** option and select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile. The encryption strengths are:

- **Basic**

  You can use this option when communicating with older Microsoft dial-up networking clients. This option uses Microsoft Point-to-Point Encryption (MPPE) and a 40-bit encryption key.

- **Strong**

  You can use this option when communicating with dial-up networking clients running Windows 98, Windows 2000, or Windows XP. This option uses MPPE and a 56-bit encryption key.

- **Strongest**

  You can use this option when communicating with dial-up networking clients running Windows 98, Windows 2000, or Windows XP. This option uses MPPE and a 128-bit encryption key.

For more information, see Configure encryption.

# L2TP-based remote access VPN deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**L2TP-based remote access VPN deployment**

You can use remote access with products in the Windows Server 2003 family to provide access to a corporate intranet for remote access clients who are making Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec) connections across the Internet. If you want your remote access server to support multiple L2TP/IPSec connections, complete the following steps:

- Configure the connection to the Internet.
- Configure the connection to the intranet.
- Configure the remote access server as a corporate intranet router.
- Configure the VPN server.
- Install certificates.
- Configure firewall packet filters.
- Configure remote access policies.

The following illustration shows the elements of a remote access server running Windows Server 2003 that provides L2TP/IPSec remote access to a corporate intranet.




For more information, see Remote access VPN connection and Layer Two Tunneling Protocol.

### Notes

- The following configuration assumes that computer certificates, also known as machine certificates, are already installed on the VPN server and remote access client computers. For more information, see Computer certificates for L2TP/IPSec VPN connections and Network access authentication and certificates.
- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Configuring the connection to the Internet**

The connection to the Internet from a computer running a Windows Server 2003 family operating system is a dedicated connection--a WAN adapter installed in the computer. The WAN adapter is typically a DDS, T1, Fractional T1, or Frame Relay adapter. You must contract with a local telephone company to run the appropriate physical wiring to your premises. You need to verify that the WAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

The WAN adapter includes drivers that are installed in members of the Windows Server 2003 family so that the WAN adapter appears as a network adapter.

You need to configure the following TCP/IP settings on the WAN adapter:

- IP address and subnet mask assigned from the InterNIC or an Internet service provider (ISP).
- Default gateway of the ISP router.

To enable VPN clients to connect to your VPN server by name rather than by IP address, request that your ISP register your VPN server in DNS.

**Configuring the connection to the intranet**
The connection to the intranet from a computer running a Windows Server 2003 operating system is a LAN adapter that is installed in the computer. You need to verify that the LAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.
- DNS and WINS name servers of corporate intranet name servers.

**Configuring the remote access server as a corporate intranet router**
In order for the remote access server to properly forward traffic on the corporate intranet, you must configure it as a router with either static routes or routing protocols so that all of the locations of the intranet are reachable from the remote access server. For information about routing concepts, see Routing Overview. For information about setting up the remote access server as a router, see Deploying Routing.

**Configuring the VPN server**
You can configure your VPN server by running the Routing and Remote Access Server Setup Wizard. You can use the wizard to configure the following settings:

- A basic firewall on the public interface.
- The method by which the VPN server assigns addresses to remote access clients (either using addresses that the VPN server obtains from a DHCP server or using addresses from a specified range of addresses that you configure).
- Forwarding of authorization and authentication messages to a Remote Authentication Dial-In User Service (RADIUS) server (configuration of the VPN server as a RADIUS client).

Once the wizard is run, the following Routing and Remote Access settings are automatically configured:

- Network interfaces
- PPTP and L2TP ports (five or 128 of each, depending on the choices you made when you ran the wizard)
- Multicast support using Internet Group Management Protocol (IGMP)
- IP routing
- Installation of the DHCP Relay Agent component.

For more information about enabling Routing and Remote Access and running the wizard, see Enable the Routing and Remote Access service.

**Installing certificates**
In order to create L2TP/IPSec remote access VPN connections using computer certificate authentication for IPSec, you must install computer certificates, also known as machine certificates, on the VPN client and the VPN server. For more information, see Network access authentication and certificates and Computer certificates for L2TP/IPSec VPN connections.

**Configuring firewall packet filters**

If you are using a firewall, you need to configure L2TP/IPSec packet filters on your firewall to allow L2TP/IPSec traffic between Internet-based VPN clients and the VPN server computer. For more information, see VPN servers and firewall configuration.

**Configuring remote access policies**

For an access-by-user administrative model, you need to set the remote access permission to **Allow access** on the user accounts for those users who will be making VPN connections. For an access-by-policy model, make the appropriate changes to the remote access permission of the user accounts. For more information, see Introduction to remote access policies.

To configure a remote access policy to control the authentication and encryption options for VPN connections, create a remote access policy with the following settings:

- Set **Policy name** to **VPN Access** (example).
- For conditions, set the **NAS-Port-Type** condition to **Virtual (VPN)** and the **Tunnel-Type** condition to **Layer Two Tunneling Protocol**.
- Select the **Grant remote access permission** option.
- For profile settings, select the appropriate authentication and encryption options.

Then, either delete the default remote access policies or position them after the new policy. This remote access policy allows all users with remote access permission to create a VPN connection.

If you want to distinguish dial-up remote access users from VPN remote access users, do the following:

1. Create an Active Directory group whose members can create virtual private networking connections with the VPN server. For example, VPN_Users.
2. Add the appropriate user accounts to the new Active Directory group.
3. Create a new remote access policy with the following properties:
    - Set **Policy name** to **VPN Access if member of VPN_Users** (example).
    - For conditions, set the **Windows-Groups** condition to **VPN_Users** (example), set the **NAS-Port-Type** condition to **Virtual (VPN)**, and set the **Tunnel-Type** condition to **Layer Two Tunneling Protocol**.
    - Select the **Grant remote access permission** option.
4. Move the default remote access policies after the new policy.

The default encryption settings allow no encryption and all levels of encryption strength. To require encryption, clear the **No Encryption** option and select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile. The encryption strengths are:

- **Basic**

  This option uses 56-bit DES encryption.

- **Strong**

  This option uses 56-bit DES encryption.

- **Strongest**

  This option uses triple DES (3DES) encryption.

For more information, see Configure encryption.

# Remote Access Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication**

This section provides information about administering user dial-in properties on a remote access server that is configured as a:

- Stand-alone server
- Member server in a domain

# Stand-alone server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Stand-alone server**

For a server running Routing and Remote Access that is not a member of a domain, only the local accounts database is available for setting remote access permissions. You cannot use domain-based account databases.

You can administer user account dial-in properties through the properties of a dial-up connection in the Network Connections folder or through Local Users and Groups. The list of users is based on the local user accounts.

# Member server in a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Member server in a domain**

For a server running Routing and Remote Access that is a member server in a domain, four accounts databases are available:

1. Local.
2. Windows NT 4.0 domain.
3. Windows 2000 mixed domain.
4. Windows 2000 native or Windows Server 2003 domain.

The following table describes the management of user dial-in properties for each accounts database.

| Accounts database | User dial-in properties |
|---|---|
| Local | Administered by using Local Users and Groups. |
| Windows NT 4.0 domain | Administered by using the Windows NT 4.0 User Manager for Domains administrative tool. |
| Windows 2000 mixed domain | Administered by using either Active Directory Users and Computers or the Windows NT 4.0 User Manager for Domains administrative tool. |
| Windows 2000 native domain | Administered by using Active Directory Users and Computers. |

**Note**

- In order for the remote access server to access user account dial-in properties stored in Active Directory, the Routing and Remote Access service must run in the security context of a computer account that is a member of the **RAS and IAS Servers** security group. When you run the Routing and Remote Access Server Setup Wizard to configure Routing and Remote Access in a domain, the computer account of the remote access server is automatically added to the **RAS and IAS Servers** security group. If the remote access server needs to access user dial-in properties in other domains, you must manually add the computer account to the **RAS and IAS Servers** security group of the other domains either through Active Directory Users and Computers or by typing **netsh ras add registeredserver** at a command prompt. For more information, see Netsh commands for remote access and Remote Access Troubleshooting.

# Using RADIUS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using RADIUS**

This section covers:

- Using RADIUS for multiple remote access servers
- Using RADIUS in a heterogeneous remote access infrastructure

# Using RADIUS for multiple remote access servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using RADIUS for multiple remote access servers**

If you have more than one remote access server, rather than administer the remote access policies of all the remote access servers separately, you can configure a single server with the Internet Authentication Service (IAS) as a Remote Authentication Dial-In User Service (RADIUS) server and configure the remote access servers as RADIUS clients. The IAS server provides centralized remote access authentication, authorization, accounting, and auditing.

For more information about IAS, see Features of IAS.

If you want to configure the remote access servers and the IAS server, complete the following steps:

- Configure the remote access server.
- Configure the remote access server for RADIUS authentication.
- Configure the remote access server for RADIUS accounting.
- Configure the IAS server.

**Configuring the remote access server**

You must configure the server running Routing and Remote Access to provide remote access to either dial-up networking clients or virtual private networking clients. For more information, see the following:

- Using the remote access server as a corporate remote access server
- Using the remote access server for Internet access
- Deploying Remote Access VPNs
- Router-to-router VPN Deployment

**Configuring the remote access server for RADIUS authentication**

When you configure the properties of the server running Routing and Remote Access, select **RADIUS authentication** as the authentication provider. For more information, see Use RADIUS authentication.

When you add a RADIUS server, you must configure the following:

- **server name**

  The host name or IP address of the IAS server.

- **Secret**

  The server running Routing and Remote Access and the IAS server share a secret that is used to encrypt messages sent between them. You must configure both the remote access server and the IAS server to use the same shared secret.

- **Port**

  The remote access server must send its authentication requests to the UDP port on which the IAS server is listening. The default value of 1812 is based on RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)" and does not need to be changed when you are using an IAS server.

**Configuring the remote access server for RADIUS accounting**

When you configure the properties of the server running Routing and Remote Access, select **RADIUS accounting** as the accounting provider. For more information, see Use RADIUS accounting.

When you add a RADIUS server, you must configure the following:

- **server name**

  The host name or IP address of the IAS server.

- **Secret**

  The server running Routing and Remote Access and the IAS server share a secret that is used to encrypt messages sent between them. You must configure both the remote access server and the IAS server to use the same shared secret.

- **Port**

  The remote access server must send its accounting requests to the UDP port on which the IAS server is listening. The default value of 1813 is based on RFC 2866, "RADIUS Accounting" and does not need to be changed when using an IAS server.

**Configuring the IAS server**

You need to configure the IAS server for domains and port numbers and register each of the remote access servers as clients. For more information, see Checklist: Configuring IAS for dial-up and VPN access.

Once the remote access servers are configured to use RADIUS authentication, the remote access policies stored on the remote access servers are no longer used. Instead, the remote access policies stored on the IAS server are used. Therefore, if one of the remote access servers contains the current set of remote access policies that are applied to all of the remote access servers, you can copy the remote access policies to the IAS server. For more information, see Copy the IAS configuration to another server.

### Notes

- If you have servers running Windows NT 4.0 and the Routing and Remote Access Service (RRAS) and you want to use remote access policies to authenticate incoming remote access connection attempts, you must configure the server running RRAS as a RADIUS client to an IAS server. You cannot configure remote access servers running Windows NT 4.0 as RADIUS clients. You must upgrade a remote access server running Windows NT 4.0 to a server running Windows NT 4.0 and RRAS.

- To provide redundancy and fault tolerance, configure two IAS servers, a primary and a backup, and copy the remote access policies from the primary to the backup. Then configure each remote access server with two RADIUS servers that correspond to the primary and backup IAS servers. If the primary IAS server becomes unavailable, then the remote access servers automatically begin to use the secondary IAS server.

- IAS is included as a part of Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

# Using RADIUS in a heterogeneous remote access infrastructure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using RADIUS in a heterogeneous remote access infrastructure**

For some Internet service providers and corporate remote access environments, the remote access equipment consists of multiple remote access devices of different types from different manufacturers. In a heterogeneous remote access environment, a single standard must exist for providing authentication of remote access user credentials and accounting of remote access activity.

In many of these environments, the Remote Authentication Dial-In User Service (RADIUS) is used. RADIUS is a client/server protocol where RADIUS clients send authentication and accounting requests to a RADIUS server. The RADIUS server checks the remote access authentication credentials on the user accounts and logs remote access accounting events.

You can configure a server running Routing and Remote Access as a RADIUS client. In this scenario, the network also includes a RADIUS server called Internet Authentication Service (IAS) that can be used by Routing and Remote Access. For more information about IAS, see Features of IAS.

If you want to configure a server running Routing and Remote Access as a RADIUS client, complete the following steps:

- Configure the remote access server.
- Configure the remote access server for RADIUS authentication.
- Configure the remote access server for RADIUS accounting.

The following illustration shows the elements of a server running Routing and Remote Access that uses RADIUS in a heterogeneous remote access infrastructure.




**Configuring the remote access server**

You must configure the server running Routing and Remote Access to provide remote access to either dial-up networking clients or virtual private networking clients. For more information, see the following:

- Using the remote access server as a corporate remote access server
- Using the remote access server for Internet access
- Deploying Remote Access VPNs
- Router-to-router VPN Deployment

**Configuring the remote access server for RADIUS authentication**

When you configure the properties of the server running Routing and Remote Access, select **RADIUS authentication** as the authentication provider. For more information, see Use RADIUS authentication.

When you add a RADIUS server, you must configure the following:

- **Server Name**

The host name or IP address of the computer that is running the RADIUS server process.

- **Secret**

  The RADIUS client (the server running Routing and Remote Access) and the RADIUS server share a secret that is used to encrypt messages sent between them. You must configure both the RADIUS client and the RADIUS server to use the same shared secret.

- **Port**

  The RADIUS client must send its authentication requests to the UDP port on which the RADIUS server is listening. The default value of 1812 is based on RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)." For some older RADIUS servers based on the earliest implementations of RADIUS, the port should be set to 1645.

**Configuring the remote access server for RADIUS accounting**

When you configure the properties of the server running Routing and Remote Access, select **RADIUS accounting** as the accounting provider. For more information, see Use RADIUS accounting.

When you add a RADIUS server, you must configure the following:

- **Server Name**

  The host name or IP address of the computer that is running the RADIUS server process.

- **Secret**

  The RADIUS client (the server running Routing and Remote Access) and the RADIUS server share a secret that is used to encrypt messages sent between them. You must configure both the RADIUS client and the RADIUS server to use the same shared secret.

- **Port**

  The RADIUS client must send its accounting requests to the UDP port on which the RADIUS server is listening. The default value of 1813 is based on RFC 2866, "RADIUS Accounting." For some older RADIUS servers, the port should be set to 1646.

# Remote Access Scenarios

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access scenarios**
This section covers:

- Internet access for A. Datum Corporation
- Remote Access for Electronic, Inc.

# Internet access for A. Datum Corporation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet access for A. Datum Corporation**
This section describes how remote access is configured for a fictional company using the features available in the Windows Server 2003 family. While your network configuration may be different than described here, you can apply the basic concepts.

A. Datum Corporation is an Internet service provider that provides dial-up Internet access and Internet services for customers in the greater Los Angeles metropolitan area. A. Datum has implemented a remote access solution that connects dial-up customers to the Internet.

**The network configuration**
The A. Datum network is configured as follows:

- A. Datum has obtained a class B network ID of 131.107.0.0 with a subnet mask of 255.255.0.0 from the Internet Assigned Numbers Authority (IANA). A. Datum connects to the Internet backbone by using a T3 WAN link. The IP address of the WAN adapter is 131.107.0.1.

- To receive up to 512 simultaneous incoming calls, A. Datum is using a modem bank switch that is connected to the local telephone service provider. The modem bank switch connects to the remote access server by using a modem bank adapter. Dial-up clients can dial in to A. Datum at the phone number 555-0122.

The following illustration shows the network configuration of the A. Datum remote access server.




The remote access server computer is configured as follows:

**1. Install the WAN adapter in the remote access server**
The network adapter that is used to connect to the Internet is installed according to the adapter manufacturer's instructions. Once the driver is installed and functioning, the adapter appears as a local area connection in the Network Connections folder.

**2. Configure TCP/IP on the WAN adapter**
The IP address of 131.107.0.1 with the subnet mask 255.255.255.255 is configured. DNS server addresses are configured. WINS server addresses and a default gateway are not configured.

**3. Install the Routing and Remote Access service**
The Routing and Remote Access Server Setup Wizard is run. Within the wizard, both remote access and routing for LAN and demand-dial routing are enabled, and all ports are enabled for both routing and remote access. For more information, see Enable the Routing and Remote Access service.

**4. Install the modem bank adapter in the remote access server**
The modem bank adapter that is used to connect to the modem bank switch is installed according to the adapter manufacturer's instructions. Once the driver is installed and functioning, the device and its ports appear under **Ports** in Routing and Remote Access.

**5. Configure the ports of the modem bank device for remote access**
All of the ports of the modem bank device are enabled for remote access. For more information, see Configure ports for remote access.

**6. Configure a static IP address pool**
A static IP address pool with a starting IP address of 131.107.192.1 and an ending IP address of 131.107.255.254 is configured. This creates a static address pool for up to 16,381 remote access clients.

For more information, see Create a static IP address pool.

**7. Configure a static route to reach Internet locations**

To reach Internet locations, a static route is configured with the following settings:

- Interface: The WAN adapter attached to the Internet
- Destination: 0.0.0.0
- Network mask: 0.0.0.0
- Gateway: 0.0.0.0
- Metric: 1

### Note

- Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway address. The gateway address of 0.0.0.0 is an example. 0.0.0.0 is reserved as the unspecified IP address.

**User accounts and remote access policy configuration**

The A. Datum remote access server is a stand-alone server that uses the access-by-user administrative model. Access is controlled by setting the dial-in permission of individual user accounts to either **Allow access** or **Deny access**. The default remote access policies are not deleted.

### Note

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Remote Access for Electronic, Inc.

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access for Electronic, Inc.**
This section describes how remote access is configured for a fictional company by using the features available in the Windows Server 2003 family. While your network configuration may be different than described here, you can apply the basic concepts.

Electronic, Inc. is an electronics design and manufacturing company that has implemented a remote access solution that connects dial-up and VPN remote access users to the Electronic, Inc. intranet.

To deploy a remote access solution for Electronic, Inc., the network administrator performs an analysis and makes design decisions regarding:

- The network configuration
- The remote access policy configuration
- The domain configuration

**The network configuration**
The key elements of the network configuration are:

- The Electronic, Inc. corporate intranet uses the private network address block of 172.16.0.0 with a subnet mask of 255.240.0.0.
- The remote access server computer is connected to an intranet network that contains a router that connects to the rest of the Electronic, Inc. corporate intranet. The network ID of the intranet network segment attached to the remote access server computer is 172.31.248.0 with a subnet mask of 255.255.255.0.
- The remote access server computer is configured with a static pool of IP addresses that represents the separate subnet of 172.31.252.0 with a subnet mask of 255.255.252.0.

The following illustration shows the network configuration of the Electronic, Inc. remote access server.




Based on the network configuration of the Electronic, Inc. corporate campus intranet, the remote access server computer is configured as follows.

**1. Install hardware in the remote access server**
The network adapter that is used to connect to the intranet segment is installed according to the adapter manufacturer's instructions. Once the driver is installed and functioning, the adapter appears as a local area connection in the Network Connections folder.

**2. Configure TCP/IP on the LAN adapter**
The IP address of 172.31.248.1 with the subnet mask 255.255.255.0 is configured. DNS and WINS server addresses are also configured. A default gateway is not configured.

**3. Install the Routing and Remote Access service**
The Routing and Remote Access Server Setup Wizard is run. Within the wizard, both remote access and LAN and demand-dial routing are enabled, and all ports are enabled for both routing and remote access. For more information, see Enable the Routing and Remote Access service.

**4. Configure a static IP address pool**

A static IP address pool with a starting IP address of 172.31.252.1 and an ending IP address of 172.31.255.254 is configured. This creates a static address pool for up to 1,021 remote access clients.

For more information, see Create a static IP address pool.

**5. Configure a static route on the remote access server computer to reach intranet locations**
To reach intranet locations, a static route is configured on the remote access server computer with the following settings:

- Interface: The LAN adapter attached to the intranet
- Destination: 172.16.0.0
- Network mask: 255.240.0.0
- Gateway: 172.31.248.2
- Metric: 1

For more information, see Add a static route.

**6. Configure a static route on the router to reach remote access clients**
To reach remote access clients, a static route is configured on the router with the following settings:

- Interface: The LAN adapter attached to the network segment of the remote access server
- Destination: 172.31.252.0
- Network mask: 255.255.252.0
- Gateway: 172.31.248.1
- Metric: 1

This static route is propagated to the other routers in the Electronic, Inc. intranet through the use of routing protocols.

For more information, see Add a static route.

**The remote access policy configuration**
To ease the transition from a Windows NT 4.0 environment, the network administrator for Electronic, Inc. decides on an access-by-user administrative model. Remote access is controlled by setting the dial-in permission of individual user accounts to either **Allow access** or **Deny access**. Remote access policies are used to apply different connection settings based on group membership. The default remote access policies are deleted.

**The domain configuration**
To take advantage of the ability to apply different connection settings to different groups, the following groups are created:

- DialUp_Users

  Used for dial-up remote access connections.

- VPN_Users

  Used for remote access VPN connections.

Based on this configuration, the following remote access scenarios are described:

- Dial-up remote access for employees
- VPN remote access for employees

- Centralized authentication by using IAS

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Dial-up remote access for employees

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up remote access for employees**
Employees of Electronic, Inc. use MS-CHAP or MS-CHAP v2 authentication for dial-up remote access and encryption is allowed but not required. To deploy remote access for Electronic, Inc. employees by using dial-up remote access connections, the following configuration is implemented.

**Network configuration**
For dial-up remote access, the Electronic, Inc. network is also configured as follows:

- To receive up to 48 simultaneous incoming calls, Electronic, Inc. uses a modem bank switch that is connected to the local telephone service provider. The modem bank switch connects to the remote access server by using a modem bank adapter. Dial-up clients can dial in to the Electronic, Inc. intranet at the phone number 555-0111.

The following illustration shows the configuration of the Electronic, Inc. remote access server for dial-up remote access connections.




The remote access server computer is configured as follows:

**1. Install modem bank adapter in the remote access server**
The modem bank adapter that is used to connect to the modem bank switch is installed according to the adapter manufacturer's instructions. Once the driver is installed and functioning, the device and its ports appear under **Ports** in Routing and Remote Access.

**2. Configure the ports of the modem bank device for remote access**
All of the ports of the modem bank device are enabled for remote access. For more information, see Configure ports for remote access.

**Domain configuration**
For each employee that is allowed dial-up remote access:

- The remote access permission on the dial-in properties of the user account is set to **Allow access**.
- The user account is added to the DialUp_Users group.

**Remote access policy configuration**
To define the authentication and encryption settings for remote access dial-up clients, the following remote access policy is created:

- Policy name: Dial-up Remote Access Clients
- Conditions:
    - **NAS-Port-Type** is set to all types *except* Virtual (dial-up)
    - **Windows-Groups** is set to DialUp_Users

- Permission is set to **Grant remote access permission**
- Profile settings:
    - **Authentication** tab: **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** and **Microsoft Encrypted Authentication (MS-CHAP)** are enabled.
    - **Encryption** tab: Select all the options.

**Note**

- In the access-by-user administrative model, the remote access permission on the remote access policy has no effect on granting remote access permission. However, the network administrator for Electronic, Inc. set the remote access permission on the policy to **Grant remote access permission** so that an eventual transition to an access-by-policy administrative model does not require changing all the remote access permission settings on all of the configured remote access policies.

**Dial-up remote access client configuration**

The New Connection wizard is used to create a dial-up connection with the following setting:

- Phone number: 555-0111

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# VPN remote access for employees

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN remote access for employees**
VPN remote access for employees of Electronic, Inc. is PPTP-based, uses MS-CHAP v2 authentication, and requires strong encryption. To deploy remote access for Electronic, Inc. employees by using remote access VPN connections across the Internet, the following configuration is implemented.

**Network configuration**
For VPN remote access, the Electronic, Inc. network is configured as follows:

- The remote access server computer is directly attached to the Internet by using a T3 (also known as a DS-3) dedicated WAN link.
- The IP address of the T3 WAN adapter on the Internet is 157.54.130.1 as allocated by Electronic, Inc.'s Internet service provider (ISP). The IP address of the WAN adapter is referred to on the Internet by the name vpn.electronic.microsoft.com.

The following illustration shows the configuration of the Electronic, Inc. remote access server for remote access VPN connections.




The remote access server computer is configured as follows:

**1. Install T3 WAN adapter in the remote access server**
The T3 WAN adapter that is used to connect to the Internet is installed according to the adapter manufacturer's instructions. Once the driver is installed and functioning, the adapter appears as a local area connection in Device Manager or the properties of Ports.

**2. ConfigureTCP/IPon the WAN adapter**
For the WAN adapter, the IP address of 157.54.130.1 with a subnet mask 255.255.255.255 is configured. A default gateway is not configured.

**3. Configure a static route to reach Internet locations**
To reach Internet locations, a static route is configured with the following settings:

- Interface: The WAN adapter attached to the Internet
- Destination: 0.0.0.0
- Network mask: 0.0.0.0
- Gateway: 157.54.130.1
- Metric: 1

**Note**

- Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway

address. The gateway address of 0.0.0.0 is an example. 0.0.0.0 is reserved as the unspecified IP address.

**4. Increase the number of PPTP ports**
By default, only five PPTP ports are enabled for VPN connections. The number of PPTP ports is increased to 1,000. For more information, see Add PPTP or L2TP ports.

**5. Configure PPTP packet filters**
PPTP packet filters are configured on the WAN adapter that connects to the Internet. For more information, see Add PPTP Filters.

**Domain configuration**
For each employee that is allowed VPN remote access:

- The remote access permission on the dial-in properties of the user account is set to **Allow access**.
- The user account is added to the VPN_Users group.

**Remote access policy configuration**
To define the authentication and encryption settings for remote access VPN clients, the following remote access policy is created:

- Policy name: Remote Access VPN Clients
- Conditions:
    - **NAS-Port-Type** is set to **Virtual (VPN)**
    - **Tunnel-Type** is set to **Point-to-Point Tunneling Protocol**
    - **Windows-Groups** is set to *VPN_Users* (example)
    - **Called-Station-ID** is set to *157.54.130.1* (example)
- Permission is set to **Grant remote access permission**
- Profile settings:
    - **Authentication** tab: **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** is enabled.
    - **Encryption** tab: Clear all the options except **Strongest**.

### Notes

- The **Called-Station-ID** is set to the IP address of the Internet interface for the remote access server. Only tunnels initiated from the Internet are allowed. Tunnels initiated from the Electronic, Inc. intranet are not permitted. Electronic, Inc. users that require Internet access from the Electronic, Inc. intranet must go through the Electronic, Inc. proxy server (not shown), where Internet access is controlled and monitored.
- In the access-by-user administrative model, the remote access permission on the remote access policy has no effect on granting remote access permission. However, the network administrator for Electronic, Inc. set the remote access permission on the policy to **Grant remote access permission** so that an eventual transition to an access-by-policy administrative model does not require changing all the remote access permission settings on all of the configured remote access policies.

**PPTP-based remote access client configuration**
The New Connection wizard is used to create a VPN connection with the following settings:

- VPN connection type: PPTP
- Host name or IP address: vpn.electronic.microsoft.com

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Centralized authentication by using IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Centralized authentication by using IAS**

To spread the load between dial-up and VPN-based remote access users, the network administrator of Electronic, Inc. decided to configure a separate VPN server. To centralize the authentication and accounting functions for both the remote access server and the VPN server, a computer running Internet Authentication Service (IAS) with the IP address of 172.31.248.9 functions as a RADIUS server.

The following illustration shows the configuration of the Electronic, Inc. network for IAS-based centralized authentication and accounting.




The IAS server has its own set of remote access polices. Rather than maintain two different sets of remote access policies, one set for the remote access server and one for the VPN server, both the remote access server and the VPN server are configured to use RADIUS authentication. Therefore, both the remote access server and the VPN server are configured as RADIUS clients to the IAS server. The IAS server provides remote access authentication, authorization, and accounting for both the remote access server and the VPN server.

**Remote access policy configuration**

Once the server running Routing and Remote Access is configured to use RADIUS authentication, the remote access policies stored on the remote access server are no longer used. Instead, the remote access policies stored on the IAS server are used. Therefore, the current set of remote access policies is copied to the IAS server.

For more information, see Copy the IAS configuration to another server.

**RADIUS configuration**

To configure RADIUS authentication and accounting, the network administrator for Electronic, Inc. configures the following:

- The IAS server is configured for two RADIUS clients: the dial-up remote access server and the VPN server. For more information, see Internet Authentication Service and Add RADIUS clients.
- The dial-up server running Routing and Remote Access is configured to use RADIUS authentication and accounting (at the IP address of 172.31.248.9) and a shared secret. For more information, see Use RADIUS authentication and Use RADIUS accounting.
- The VPN server running Routing and Remote Access is configured to use RADIUS authentication and accounting (at the IP address of 172.31.248.9) and a shared secret. For more information, see Use RADIUS authentication and Use RADIUS accounting.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This section covers:

- Remote access RADIUS attributes
- Netsh commands for remote access
- Expressing an IP address range with a mask
- Pattern matching syntax
- Remote Access Troubleshooting tools

# Remote access RADIUS attributes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access RADIUS attributes**

The following lists the RADIUS attributes that are supported by Routing and Remote Access for the various RADIUS packet types. For more information about these attributes, see RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)" and RFC 2548, "Microsoft Vendor-specific RADIUS Attributes."

**Access-Request**

- User-Name (RADIUS attribute 1)

  A string value that contains the user principal name or the Windows NT name without the domain.

- User-Password (RADIUS attribute 2)

  A string value that contains the user password and is only sent if Password Authentication Protocol (PAP) is negotiated as the authentication protocol.

- CHAP-Password (RADIUS attribute 3)

  Contains the response value provided by the Challenge Handshake Authentication Protocol (CHAP) in response to the challenge.

- NAS_IP_Address (RADIUS attribute 4)

  Contains the IP address of the remote access server.

- NAS-Port (RADIUS attribute 5)

  Indicates the number of the port on the remote access server on which the incoming call was received.

- Service-Type (RADIUS attribute 6)

  The only value that is sent is "Framed" (2).

- Framed-Protocol (RADIUS attribute 7)

  The only value that is sent is "PPP" (1).

- Framed-MTU (RADIUS attribute 12)

  Used in conjunction with EAP authentication to notify the RADIUS server of the maximum transmission unit (MTU) negotiated with the client, so that the RADIUS server does not send EAP messages that cannot be delivered over the link.

- State (RADIUS attribute 24)

  The remote access server never sends this attribute in the initial Access-Request. If EAP is used as the authentication protocol and a State attribute is received in an Access-Challenge packet, that State attribute is returned unmodified in the next Access-Request packet sent.

- Called-Station-Id (RADIUS attribute 30)

  Telephone number on which the call was received. For virtual private network (VPN) connections, the IP address of the VPN server.

- Calling-Station-Id (RADIUS attribute 31)

  Telephone number on which the call was made. For virtual private network (VPN) connections, the IP address of the VPN client.

- NAS-Identifier (RADIUS attribute 32)

  The fully qualified domain name of the remote access server.

- CHAP-Challenge (RADIUS attribute 60)

  The challenge sent by the remote access server during CHAP authentication.

- NAS-Port-Type (RADIUS attribute 61)

  The only values sent are **Async (Modem)** (0), **ISDN Sync** (2), **ISDN Async V.120** (3), **ISDN Async V.110** (4), and **Virtual (VPN)** (5). Virtual port is used to indicate VPN connections.

- Tunnel-Type (RADIUS attribute 64)

  The only values sent are PPTP (1) and L2TP (3).

- Tunnel-Medium-Type (RADIUS attribute 65)

  The only value sent is IP (1).

- Tunnel-Client-EndPoint (RADIUS attribute 66) (PPTP only)

- Tunnel-Server-EndPoint (RADIUS attribute 67) (PPTP only)

- Connect-Info (RADIUS attribute 77)

  Contains whatever data TAPI returns about the call.

- EAP-Message (RADIUS attribute 79)

- Message Authenticator (RADIUS attribute 80)

  Always sent if EAP is used for authentication. Otherwise, it is configurable on the properties of the RADIUS server.

- MS-CHAP-Response (vendor type 1)

- MS-CHAP-CPW-2 (vendor type 4)

- MS-Chap-LM-Enc-Pw (vendor type 5)

- MS-CHAP-NT-Enc-PW (vendor type 6)

- MS-Ras-Vendor (vendor type 9)

- MS-Chap-Challenge (vendor type 11)

- MS-Ras-Version (vendor type 18)

**Access-Accept**

- Service-Type (RADIUS attribute 6)

  Only Framed (2) and Callback Framed (4) are accepted. If any other Service-Type is received, the call is dropped.

- Framed-Protocol (RADIUS attribute 7)

  Only PPP (1) is accepted. If any other Framed-Protocol is received, the call is dropped.

- Framed-IP-Address (RADIUS attribute 8)

  The only acceptable values are 0xFFFFFFFF (user selects address) and 0xFFFFFFFE (remote access server selects address). If any other Framed-IP-Address is received, the call is dropped.

- Framed-MTU (RADIUS attribute 12)

  Not used.

- Framed-Compression (RADIUS attribute 13)

  The only recognized values are None (0) and VJ TCP/IP header compression (1). All other values are ignored.

- Reply-Message (RADIUS attribute 18)

  Returned in CHAP, PAP, and MS-CHAP (both v1 and v2) Success packets and in the EAP Notification message.

- Framed-Route (RADIUS attribute 22)

- State (RADIUS attribute 24)

  Returned to the remote access server unchanged.

- Class (RADIUS attribute 25)

  Sent unchanged to accounting server in Accounting Start message.

- Session-Timeout (RADIUS attribute 27)

- Idle-Timeout (RADIUS attribute 28)

- Termination-Action (RADIUS attribute 29)

  Not used.

- Port-Limit (RADIUS attribute 62)

- EAP-Message (RADIUS attribute 79)

- Message Authenticator (RADIUS attribute 80)

  Only accepted if EAP is used for authentication.

- Acct-Interim-Interval (RADIUS attribute 85)

- MS-MPPE-Encryption-Policy (vendor type 7)

- MS-MPPE-Encryption-Types (vendor type 8)

- MS-CHAP-Domain (vendor type 10)
- MS-BAP-Usage (vendor type 13)
- MS-Link-Utilization-Threshold (vendor type 14)
- MS-Link-Drop-Time-Limit (vendor type 15)
- MS-Filter (vendor type 22)

**Access-Challenge**
The Access-Challenge is only used with EAP. Otherwise, the receipt of an Access-Challenge is treated as Access-Reject.

- State (RADIUS attribute 24)
- Session-Timeout (RADIUS attribute 27)
- EAP-Message (RADIUS attribute 79)
- Message Authenticator (RADIUS attribute 80)

**Accounting-Request**

- User-Name (RADIUS attribute 1)
- NAS_IP_Addresss (RADIUS attribute 4)
- NAS-Port (RADIUS attribute 5)
- Service-Type (RADIUS attribute 6)
- Framed-Protocol (RADIUS attribute 7)
- Framed-IP-Address (RADIUS attribute 8)
- Framed-MTU (RADIUS attribute 12)
- Framed-Compression (RADIUS attribute 13)
- Framed-Route (RADIUS attribute 22)
- Class (RADIUS attribute 25)
- Session-Timeout (RADIUS attribute 27)
- Idle-Timeout (RADIUS attribute 28)
- Termination-Action (RADIUS attribute 29)
- Called-Station-Id (RADIUS attribute 30)
- Calling-Station-Id (RADIUS attribute 31)
- NAS-Identifier (RADIUS attribute 32)

- NAS-Port-Type (RADIUS attribute 61)
- Port-Limit (RADIUS attribute 62)
- Tunnel-Type (RADIUS attribute 64)
- Tunnel-Medium-Type (RADIUS attribute 65)
- Tunnel-Client-EndPoint (RADIUS attribute 66)
- Tunnel-Server-EndPoint (RADIUS attribute 67)
- Connect-Info (RADIUS attribute 77)
- Acct-Status-Type (RADIUS attribute 40)

  "On" is sent when the Routing and Remote Access service is started. "Off" is sent if the Routing and Remote Access service is gracefully stopped. "Start" and "Stop" are sent at the beginning and end of a user connection. "Interim-Update" is sent at approximately the interval specified in the Acct-Interim-Interval attribute (some random jitter is applied) and only if the Acct-Interim-Interval attribute was returned in the Access-Accept message.

- Acct-Delay-Time (RADIUS attribute 41)

  Five seconds are added on every retransmission (regardless of the actual time between retransmissions).

- Acct-Input-Octets (RADIUS attribute 42)
- Acct-Output-Octets (RADIUS attribute 43)
- Acct-Session-Id (RADIUS attribute 44)
- Acct-Authentic (RADIUS attribute 45)
- Acct-Session-Time (RADIUS attribute 46)
- Acct-Input-Packets (RADIUS attribute 47)
- Acct-Output-Packets (RADIUS attribute 48)
- Acct-Termination-Cause (RADIUS attribute 49)

  The only values sent are 1 (User Request), 4 (Idle Timeout), 5 (Session Timeout), 6 (Admin Reset), and 8 (Port Error).

- Acct-Multi-Session-Id (RADIUS attribute 50)
- Acct-Link-Count (RADIUS attribute 51)
- MS-Ras-Vendor (vendor type 9)
- MS-CHAP-Domain (vendor type 10)
- MS-Ras-Version (vendor type 18)
- MS-Filter (vendor type 22)

# Netsh commands for remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Netsh commands for remote access**

The following table lists the Netsh commands that you can type at a command prompt to administer a server running Routing and Remote Access. If there are multiple commands for a particular function, they are indicated by separating the individual commands with a slash (/). For example, the command **ras add/delete/show registeredserver** is actually three separate commands: **ras add registeredserver**, **ras delete registeredserver**, and **ras show registeredserver**.

When you type the command at the command prompt, precede each command with **netsh**. For the exact syntax of each command, type the command with the **?** option. For example, to receive command-line Help on the **netsh ras add registeredserver** command, type **netsh ras add registeredserver ?** at the command prompt.

| Command | Description |
|---|---|
| **ras add/delete/show registeredserver** | Configures or displays whether the specified remote access server computer is a member of the **RAS and IAS Servers** security group in the Active Directory directory service of the specified domain. |
| **ras show activeservers** | Displays current servers running Routing and Remote Access on your network. |
| **ras set/show authmode** | Configures or displays whether and when dial-in connections are authenticated. |
| **ras add/delete/show authtype** | Configures or displays the permitted authentication types. |
| **ras add/delete/show client** | Configures or displays currently connected remote access clients. |
| **ras add/delete/show link** | Configures or displays the configuration of software compression and link control protocol (LCP) extensions. |
| **ras add/delete/show multilink** | Configures or displays Multilink and Bandwidth Allocation Protocol (BAP) settings. |
| **ras set/show tracing** | Configures or displays tracing settings. |
| **ras set/show user** | Configures or displays remote access settings for user accounts. |
| **ras ip set access** | Configures whether IP traffic from remote access clients is forwarded to the networks to which the remote access server is connected. |
| **ras ip set addrassign** | Configures the method by which the remote access server assigns IP addresses to incoming connections. |
| **ras ip set addrreq** | Configures whether remote access clients or demand-dial routers can request their own IP addresses. |
| **ras ip show config** | Displays IP remote access configuration. |
| **ras ip set negotiation** | Configures whether IP is negotiated for remote access connections. |
| **ras ip delete pool** | Deletes the static IP address pool. |
| **ras ip add/delete range** | Adds or removes a range of addresses from the static IP address pool. |
| **ras appletalk set access** | Configures whether AppleTalk traffic from remote access clients is forwarded to the networks to which the remote access server is connected. |

| | |
|---|---|
| **ras appletalk show config** | Displays AppleTalk remote access configuration. |
| **ras appletalk set negotiation** | Configures whether AppleTalk is negotiated for remote access connections. |
| **ras aaaa set/show accounting** | Configures or displays the accounting provider. |
| **ras aaaa add/delete/set/show acctserver** | Configures or displays RADIUS accounting servers. |
| **ras aaaa set/show authentication** | Configures or displays the authentication provider. |
| **ras aaaa add/delete/set/show authserver** | Configures or displays RADIUS authentication servers. |

For more information about Netsh, see The Netsh Command-Line Utility.

# Expressing an IP address range with a mask

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Expressing an IP address range with a mask**

In some cases, Routing and Remote Access allows you to enter a range of IP addresses as a starting and ending range. For example, when creating a DHCP scope, you enter the IP address range as 10.0.0.1 through 10.0.0.199. In this case, a suitable range for an arbitrary number of required IP addresses is specified.

In other cases, Routing and Remote Access requires you to enter a range of IP addresses as an IP address and mask. In these cases, a range for an arbitrary number of required IP addresses cannot be specified due to the mathematical relationship between the IP address and the mask. When expressing an IP address range in this way, you must choose a base IP address and a mask.

- An IP address is a 32-bit number that is expressed as a series of 4 octets converted to base 10 and separated by the period character. This is commonly known as dotted decimal notation. The base IP address contains a series of bits that are fixed and a series of bits that are variable. The bits that are fixed must be a set of contiguous bits starting from the high order bit. The range of IP addresses formed from the base IP address is the set of addresses defined by all of the possible values of the variable bits.

- The mask is a 32-bit number (also expressed in dotted decimal notation) that is used to explicitly define which bits in the base IP address are fixed and which bits are variable. In the mask, a 1 bit is used to indicate a bit in the IP address that is fixed. A fixed bit cannot be used to express the range. A 0 bit in the mask is used to indicate a bit in the IP address that is variable.

Because the bits that are fixed in the base IP address must be contiguous from the high order bit, a mask must consist of a series of contiguous 1 bits that represent the fixed portion of the address followed by a series of contiguous 0 bits that represent the variable portion of the address.

The following table shows examples of valid masks.

| Subnet mask (binary) | Subnet mask (decimal) |
|---|---|
| 11111111 11111111 11110000 00000000 | 255.255.240.0 |
| 11111111 11111111 11111111 00000000 | 255.255.255.0 |
| 11111111 11111111 11111111 11100000 | 255.255.255.224 |

The following table shows examples of invalid masks.

| Subnet mask (binary) | Subnet mask (decimal) |
|---|---|
| 11111111 11111111 10110000 00000000 | 255.255.176.0 |
| 11111111 11111111 11110000 11110000 | 255.255.240.240 |
| 11111111 11111111 11111111 00000111 | 255.255.255.7 |

The base IP address, the mask, and the range of addresses are related to each other by a mathematical operation called a bit-wise logical AND. In a bit-wise logical AND, the 32 bits of an IP address are lined up with the 32 bits of the mask. For each bit, a logical AND is performed. In a logical AND, the result is a 1 if both bits being compared are a 1 and a 0 otherwise.

Because the variable bits in the mask are always 0, the bit-wise logical AND of any IP address in the range (including the base IP address) and the mask always produces the same result: the base IP address.

The base IP address cannot have a 1 where there is a 0 in the mask. A bit-wise logical AND of a 1 with a 0 always produces a 0. Therefore, a base IP address with a 1 where there is a 0 in the mask will, when a bit-wise logical AND is used with the mask, not produce the base IP address. This is the reason for the error message: "The network mask entered is not valid. The destination address cannot be more specific than the network mask."

The base IP address that contains a 1 where there is a 0 is more specific than the mask. A base IP address can be less specific than the mask, but not more specific.

The following table shows an example of a valid address range.

| | Binary | Decimal |
|---|---|---|
| **Base IP address** | 11000000 10101000 10011100 00000000 | 192.168.156.0 |
| **Mask** | 11111111 11111111 11111100 00000000 | 255.255.252.0 |

In this example, the first 22 bits are fixed and the last 10 bits are variable. Notice how the bit-wise logical AND of the base IP address and the mask produces the base IP address (192.168.156.0 AND 255.255.252.0 = 192.168.156.0).

The following table shows an example of a invalid address range.

| | Binary | Decimal |
|---|---|---|
| **Base IP address** | 11000000 10101000 00111001 00000000 | 192.168.57.0 |
| **Mask** | 11111111 11111111 11110000 00000000 | 255.255.240.0 |

In this example, there are two bits in the base IP address set to 1 that are not included in the mask. Notice how the bit-wise logical AND of the base IP address and the mask does not produce the base IP address (192.168.57.0 AND 255.255.240.0 = 192.168.48.0).

**Multicast scopes**

The base IP address and mask that you choose for multicast scopes depends on two factors:

1. The originating address space

   The originating address space is the original set of addresses from which a range is being defined. For example, when defining multicast scopes, the originating address space is from 239.0.0.0 through 239.254.255.255.

2. The number of addresses needed

   The following table lists the mask needed based on a required number of addresses.

| Number of addresses required | Number of fixed bits | Number of variable bits | Mask |
|---|---|---|---|
| 1-2 | 31 | 1 | 255.255.255.254 |
| 3-4 | 30 | 2 | 255.255.255.252 |
| 5-8 | 29 | 3 | 255.255.255.248 |
| 9-16 | 28 | 4 | 255.255.255.240 |
| 17-32 | 27 | 5 | 255.255.255.224 |
| 33-64 | 26 | 6 | 255.255.255.192 |
| 65-128 | 25 | 7 | 255.255.255.128 |
| 129-256 | 24 | 8 | 255.255.255.0 |
| 257-512 | 23 | 9 | 255.255.254.0 |
| 513-1,024 | 22 | 10 | 255.255.252.0 |
| 1,025-2,048 | 21 | 11 | 255.255.248.0 |
| 2,049-4,096 | 20 | 12 | 255.255.240.0 |

| | | | |
|---|---|---|---|
| 4,097-8,192 | 19 | 13 | 255.255.224.0 |
| 8,193-16,384 | 18 | 14 | 255.255.192.0 |
| 16,385-32,768 | 17 | 15 | 255.255.128.0 |
| 32,769-65,536 | 16 | 16 | 255.255.0.0 |
| 65,537-131,072 | 15 | 17 | 255.254.0.0 |
| 131,073-262,144 | 14 | 18 | 255.252.0.0 |
| 262,145-524,288 | 13 | 19 | 255.248.0.0 |
| 524,289-1,048,576 | 12 | 20 | 255.240.0.0 |
| 1,048,577-2,097,152 | 11 | 21 | 255.224.0.0 |
| 2,097,153-4,194,304 | 10 | 22 | 255.192.0.0 |
| 4,194,305-8,388,608 | 9 | 23 | 255.128.0.0 |
| 8,388,609-16,777,216 | 8 | 24 | 255.0.0.0 |

Based on the base IP address, the range is defined as follows:

- The first IP address in the range is the base IP address.
- The last IP address in the range is the address obtained by setting all the variable bits to a 1 and then expressing the result in dotted decimal notation.

The following table shows an example for a multicast scope.

| | **Decimal** | **Binary** |
|---|---|---|
| **IP address** | 239.0.64.0 | 11101111 00000000 01000000 00000000 |
| **Mask** | 255.255.192.0 | 11111111 11111111 11000000 00000000 |
| **First IP address** | 239.0.64.0 | 11101111 00000000 01000000 00000000 |
| **Last IP address** | 239.0.128.255 | 11101111 00000000 01111111 11111111 |

# Pattern matching syntax

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Pattern matching syntax**

Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition support regular expressions for pattern matching, which is widely used in UNIX environments. You can use this syntax to specify the conditions of remote access policy attributes and RADIUS realms, as specified in the following table.

| Character | Description | Example |
|---|---|---|
| \ | Marks the next character as a character to match. | /n/ matches the character "n". The sequence /\n/ matches a line feed or newline character. |
| ^ | Matches the beginning of the input or line. | |
| $ | Matches the end of the input or line. | |
| * | Matches the preceding character zero or more times. | /zo*/ matches either "z" or "zoo." |
| + | Matches the preceding character one or more times. | /zo+/ matches "zoo" but not "z." |
| ? | Matches the preceding character zero or one times. | /a?ve?/ matches the "ve" in "never." |
| . | Matches any single character except a newline character. | |
| **(** *pattern* **)** | Matches *pattern* and remembers the match. To match ( ) (parentheses), use "\(" or "\)". | |
| *x* **\|** *y* | Matches either *x* or *y*. | /z\|food?/ matches "zoo" or "food." |
| **{** *n* **}** | Matches exactly *n* times (*n* is a nonnegative integer). | /o{2}/ does not match the "o" in "Bob," but matches the first two instances of the letter o in "foooood." |
| **{** *n* **,}** | Matches at least *n* times (*n* is a nonnegative integer). | /o{2,}/ does not match the "o" in "Bob" but matches all of the instances of the letter o in "foooood." /o{1,}/ is equivalent to /o+/. |
| **{** *n* **,** *m* **}** | Matches at least *n* and at most *m* times (*m* and *n* are nonnegative integers). | /o{1,3}/ matches the first three instances of the letter o in "fooooood." |
| **[** *xyz* **]** | Matches any one of the enclosed characters (a character set). | /[abc]/ matches the "a" in "plain." |
| **[^** *xyz* **]** | Matches any characters that are not enclosed (a negative character set). | /[^abc]/ matches the "p" in "plain." |
| **\b** | Matches a word boundary (for example, a space). | /ea*r\b/ matches the "er" in "never early." |
| **\B** | Matches a nonword boundary. | /ea*r\B/ matches the "ear" in "never early." |
| **\d** | Matches a digit character (equivalent to [0-9]). | |
| **\D** | Matches a nondigit character (equivalent to [^0-9]). | |
| **\f** | Matches a form feed character. | |
| **\n** | Matches a line feed character. | |

| | | |
|---|---|---|
| **\r** | Matches a carriage return character. | |
| **\s** | Matches any white space character including space, tab, and form feed (equivalent to [ \f\n\r\t\v]). | |
| **\S** | Matches any non-white space character (equivalent to [^ \f\n\r\t\v]). | |
| **\t** | Matches a tab character. | |
| **\v** | Matches a vertical tab character. | |
| **\w** | Matches any word character, including underscore (equivalent to [A-Za-z0-9_]). | |
| **\W** | Matches any nonword character, excluding underscore (equivalent to [^A-Za-z0-9_]). | |
| **** *num* | Refers to remembered matches (?*num*, where *num* is a positive integer). For example, \1 replaces what is stored in the first remembered match. This option can be used only in the **Replace** text box when configuring attribute manipulation. | |
| **/** *n* **/** | Allows the insertion of ASCII codes into regular expressions (?*n*, where *n* is an octal, hexadecimal, or decimal escape value). | |

**Examples for remote access policy attributes**

The following examples describe the use of the pattern matching syntax to specify remote access policy attributes:

- To specify all phone numbers within the 899 area code, the syntax is:

  899.*

- To specify a range of all IP addresses that begin with 192.168.1, the syntax is:

  192\.168\.1\..+

**Examples for manipulation of the realm name in the User-Name attribute**

The following examples describe the use of the pattern matching syntax to manipulate realm names for the User-Name attribute, which is located on the **Attribute** tab in the properties of a connection request policy. For more information, see Configure attribute manipulation.

- To remove the realm portion of the User-Name attribute

  In the outsourced dial scenario, the Internet service provider (ISP) might require a realm name to route the authentication request. However, the IAS server might not recognize the realm name portion of the user name. Therefore, the realm name must be removed before it is forwarded to the IAS server.

  - **Find:** @microsoft\.com
  - **Replace:**

- To replace *user@example.microsoft.com* with *example.microsoft.com\user*

  - **Find:** (.*)@(.*)
  - **Replace:** $2\$1

- To replace *domain\user* with *specific_domain\user*

  - **Find:** (.*)\\(.*)

- **Replace:** *specific_domain*\$2

- To replace *user* with *user@specific_domain*

  - **Find:** $

  - **Replace:** *@specific_domain*

**Example for RADIUS message forwarding by a proxy server**
You can create routing rules that forward RADIUS messages with a specific realm name to a specific set of RADIUS servers when IAS is used as a RADIUS proxy. Following is a recommended syntax for routing requests based on realm name.

| NetBIOS name | Pattern |
|---|---|
| WCOAST | ^wcoast\\ |

In the following example, wcoast.microsoft.com is a unique user principal name (UPN) suffix for the DNS or Active Directory domain wcoast.microsoft.com. Using the supplied pattern, the IAS proxy can route messages based on domain NetBIOS name or UPN suffix.

| NetBIOS name | UPN suffix | Pattern |
|---|---|---|
| WCOAST | wcoast.microsoft.com | ^wcoast\\\|@wcoast\.microsoft\.com$ |

**Ping user-name**
Some RADIUS proxy servers and network access servers periodically send authentication and accounting requests (known as ping requests) to verify that the IAS server is present on the network. These ping requests include fictional user names. When IAS processes these requests, the event and accounting logs become filled with access reject records, making it more difficult to keep track of valid records.

When you configure a registry entry for **ping user-name**, IAS matches the registry entry value against the user name value in ping requests by other servers. A **ping user-name** registry entry specifies the fictional user name (or a user name pattern, with variables, that matches the fictional user name) sent by RADIUS proxy servers and network access servers. When IAS receives ping requests that match the **ping user-name** registry entry value, IAS rejects the authentication requests without processing the request. IAS does not record transactions involving the fictional user name in any log files, which makes the event log easier to interpret.

### Note

- **Ping user-name** is not installed by default. You must add **ping user-name** to the registry. You can add an entry to the registry using Registry Editor. For more information, see Registry Editor.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Adding ping user-name to the registry**
**Ping user-name** can be added to the following registry key as a string value by a member of the local Administrators group:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IAS\Parameters**

| Name | Type | Data |
|---|---|---|
| ping user-name | REG_SZ | *User name* |

### Note

- To indicate more than one user name for a **ping user-name** value, enter a name pattern, such as a Domain Name

System (DNS) name including wildcard characters, in **Data**.

# Remote Access Troubleshooting tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting tools**
To help you gather information to troubleshoot problems with the remote access server, the following troubleshooting tools are available:

- Authentication and accounting logging. For more information see Logging.
- Event logging records errors, warnings, and events in the system event log.
- PPP logging creates a log file of the PPP control messages that are sent and received during a PPP connection.
- Tracing records the sequence of programming functions called during a process to a file. For more information, see Using tracing.
- Network Monitor captures traffic sent between a remote access server and dial-up client during the PPP connection process and during data transfer.
- Netsh commands for remote access configure and display remote access settings for a server running Routing and Remote Access. For more information, see Netsh commands for remote access (ras).

**Remote Access Diagnostics**
You can use the Remote Access diagnostic functions available in the Windows Server 2003 family to collect detailed logs and information about a remote access connection. Diagnostics can be set from the **Diagnostics** tab of the properties of the connection in the Network Connection folder, or you can run specific diagnostics by typing the **netsh ras diagnostics** command at the command line.

**Audit logging and Oakley logging for L2TP/IPSec connections**
You can use the Audit Logging feature to monitor IPSec-related events. This is the fastest and simplest method to troubleshoot unsuccessful L2TP/IPSec connections. You can also enable the Oakley log to record all ISAKMP main mode or quick mode negotiations. For more information, see IPSec troubleshooting tools.

**Event logging**
You can use event logging to record remote access server errors, warnings, and other detailed information in the system event log. You can enable event logging on the **Event Logging** tab on the properties of a remote access server. For more information, see View properties of the remote access server.

**PPP logging**
PPP logging records the series of programming functions and PPP control messages during a PPP connection and is a valuable source of troubleshooting information when you are troubleshooting the failure of a PPP connection. To enable PPP logging, select the **Enable Point-to-Point Protocol (PPP) logging** option on the **PPP** tab on the properties of a remote access server. For more information, see View properties of the remote access server.

The PPP log in Windows NT 4.0 has been replaced by the tracing function. To duplicate the PPP log, you need to enable file tracing for the PPP key. By default, the PPP log is stored as the ppp.log file in the *systemroot*\Tracing folder. For more information on tracing options, see Using tracing.

**Network Monitor**
Network Monitor can capture the network traffic between a dial-up networking client and the remote access server. By analyzing remote access traffic, you can find answers to remote access problems and possible solutions or workarounds. For more information, see Install and Configure Network Monitor.

### Note

- The proper interpretation of remote access traffic with Network Monitor requires an in-depth understanding of PPP and other protocols. You can save Network Monitor traces as files and send them to Microsoft Product Support Services for analysis.

# Remote Access Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Unable to establish a remote access connection.
- Unable to access resources beyond the remote access server.
- Callback is not working.
- Not sure how to configure the remote access server to make as well as receive remote access connections (dial out and dial in).

**Unable to establish a remote access connection.**
**Cause:** Your modem is not working properly.

**Solution:** Verify that your modem is working properly.

**See also:** Troubleshooting modems

**Cause:** The Routing and Remote Access service is not started on the remote access server.

**Solution:** Verify the state of the Routing and Remote Access service on the remote access server.

**See also:** Monitor the Routing and Remote Access service

**Cause:** Remote access is not enabled on the remote access server.

**Solution:** Enable remote access on the remote access server.

**See also:** Enable the remote access server

**Cause:** Dial-in, PPTP, or L2TP ports are not enabled for inbound remote access connections.

**Solution:** Enable PPTP, L2TP, or dial-in ports for inbound remote access connections, as needed.

**See also:** Configure ports for remote access

**Cause:** The LAN protocols being used by the remote access clients are not configured on the remote access server to allow remote access connections.

**Solution:** Verify that remote access connections are allowed on the **IP** or **AppleTalk** tabs on the properties of a server. If you do not have one or more of these tabs, then the corresponding network protocol is not installed on the server.

**See also:** View properties of the remote access server

**Cause:** All of the remote access ports on the remote access server are already being used by currently connected remote access clients or demand-dial routers.

**Solution:** You can verify whether all of the remote access ports on the remote access server are not already being used by clicking **Ports** in Routing and Remote Access. If all ports are busy, disconnect any inactive connections, or consider adding more ports.

**Cause:** The remote access client and the remote access server in conjunction with a remote access policy are not configured to use at least one common authentication method.

**Solution:** Configure the remote access client and the remote access server in conjunction with a remote access policy to use at least one common authentication method.

If a remote access client running Windows 95 is attempting a dial-up connection, verify that the remote access server is not

requiring only MS-CHAP v2 authentication. A remote access client running Windows 95 supports the use of MS-CHAP v2 over virtual private network (VPN) connections, not the use of MS-CHAP v2 over dial-up connections.

**Cause:** The remote access client and the remote access server in conjunction with a remote access policy are not configured to use at least one common encryption strength.

**Solution:** Configure the remote access client and the remote access server in conjunction with a remote access policy to use at least one common encryption strength.

**See also:** Configure encryption

**Cause:** The remote access connection does not have the appropriate permissions through dial-in properties of the user account and remote access policies.

**Solution:** Verify that the remote access connection has the appropriate permissions through dial-in properties of the user account and remote access policies.

In order for the connection to be established, the settings of the connection attempt must:

- Match all of the conditions of at least one remote access policy.
- Be granted remote access permission through the local user account (set to **Allow access**), or through the domain user account (set to **Control access through Remote Access Policy**) and the remote access permission of the matching remote access policy (set to **Grant remote access permission**).
- Match all the settings of the profile.
- Match all the settings of the dial-in properties of the user account.

**See also:** Introduction to remote access policies; Accepting a connection attempt

**Cause:** The settings of the remote access policy profile are in conflict with properties of the remote access server.

**Solution:** Verify that the settings of the remote access policy profile are not in conflict with properties of the remote access server.

The properties of the remote access policy profile and the properties of the remote access server both contain settings for:

- Multilink
- Bandwidth allocation protocol
- Authentication protocols

If the settings of the profile of the matching remote access policy are in conflict with the settings of the remote access server, the connection attempt is rejected. For example, if the matching remote access policy profile specifies that the EAP-TLS authentication protocol must be used and EAP is not enabled on the remote access server, the connection attempt is rejected.

**See also:** Enable authentication protocols; Configure authentication

**Cause:** The credentials of the remote access client (user name, password, and domain name) are incorrect and cannot be validated by the remote access server.

**Solution:** Verify that the credentials of the remote access client (user name, password, and domain name) are correct and can be validated by the remote access server.

**Cause:** There are not enough addresses in the static IP address pool.

**Solution:** If the remote access server is configured with a static IP address pool, verify that there are enough addresses in the pool. If all of the addresses in the static pool have been allocated to connected remote access clients, the remote access server is unable to allocate an IP address, and the connection attempt is rejected. Modify the static IP address pool if needed.

**See also:** Create a static IP address pool

**Cause:** The authentication provider of the remote access server is improperly configured.

**Solution:** Verify the configuration of the authentication provider. You can configure the remote access server to use either Windows Authentication or Remote Authentication Dial-In User Service (RADIUS) to authenticate the credentials of the remote access client.

**See also:** Use RADIUS authentication

**Cause:** The remote access server cannot access Active Directory.

**Solution:** For a remote access server that is a member server of a domain that is configured for Windows Authentication, verify that:

- The **RAS and IAS Servers** security group exists. If not, then create the group and set the group type to **Security** and the group scope to **Domain local**.

- The **RAS and IAS Servers** security group has Read permission to the **RAS and IAS Servers Access Check** object.

- The computer account of the remote access server computer is a member of the **RAS and IAS Servers** security group. You can use the **netsh ras show registeredserver** command to view the current registration. You can use the **netsh ras add registeredserver** command to register the server in a specified domain.

  If you add or remove the remote access server to the **RAS and IAS Servers** security group, the change does not take effect immediately (due to the way that Active Directory information is cached). For the change to take effect immediately, you need to restart the remote access server computer.

- The remote access server has joined the domain.

**See also:** Create a new group; Verify permissions for the RAS and IAS security group; Netsh commands for remote access

**Cause:** A remote access server running Windows NT 4.0 cannot validate connection requests.

**Solution:** If remote access clients are dialing in to a remote access server running Windows NT 4.0 that is a member of a Windows 2000 mixed domain, verify that the Everyone group is added to the Pre-Windows 2000 Compatible Access group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, type **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** at the command prompt on a domain controller computer and then restart the domain controller computer.

**See also:** Member server in a domain

**Cause:** The Windows Fax service is enabled and your modem does not support adaptive answer.

**Solution:** If the Windows Fax service and the Routing and Remote Access service are sharing the same modem, verify that the modem supports adaptive answer. If the modem does not support adaptive answer, you must disable fax on the modem to receive incoming remote access connections.

**Cause:** You are using MS-CHAP v1 and a user password over 14 characters long.

**Solution:** Either use a different authentication protocol such as MS-CHAP v2 or change your password so that it is 14 characters or less.

**See also:** MS-CHAP; Enable authentication protocols

**Unable to access resources beyond the remote access server.**
**Cause:** For IP-based remote access clients, IP routing is not enabled.

**Solution:** Verify that IP routing is enabled on the **IP** tab on the properties of the server.

**See also:** View properties of the remote access server

**Cause:** For AppleTalk-based remote access clients, network access is not allowed.

**Solution:** Verify that network access is allowed on the **AppleTalk** tab on the properties of a server. If you do not have the AppleTalk tab, then AppleTalk is not installed on the server.

**See also:** View properties of the remote access server

**Cause:** A static IP address pool is configured but there are no routes back to the remote access clients.

**Solution:** If the remote access server is configured to use a static IP address pool, verify that the routes to the ranges of addresses of the static IP address pool are reachable by the hosts and routers of the intranet. If not, then IP routes consisting of the address ranges of the static IP address pool, as defined by the IP address and mask of each range, must be added to the routers of the intranet or the routing protocol of your routed infrastructure on the remote access server must be enabled. If the routes to the remote access client subnets are not present, remote access clients cannot receive traffic from locations on the intranet. A route for the network is implemented either through static routing entries or through a routing protocol, such as Open Shortest Path First (OSPF) or Routing Information Protocol (RIP). OSPF is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

If the remote access server is configured to use DHCP for IP address allocation and no DHCP server is available, the remote access server allocates addresses from the Automatic Private IP Addressing (APIPA) address range from 169.254.0.1 through 169.254.255.254. Allocating APIPA addresses for remote access clients works only if the network to which the remote access server is attached is also using APIPA addresses.

If the remote access server is using APIPA addresses when a DHCP server is available, verify that the proper adapter is selected from which to obtain DHCP-allocated IP addresses. By default, the remote access server randomly chooses the adapter to use to obtain IP addresses through DHCP. If there is more than one LAN adapter, then the Routing and Remote Access service may choose a LAN adapter for which there is no DHCP server available.

If the static IP address pool consists of ranges of IP addresses that are a subset of the range of IP addresses for the network to which the remote access server is attached, verify that the ranges of IP addresses of the static IP address pool are not assigned to other TCP/IP nodes, either through static configuration or through DHCP.

**See also:** Create a static IP address pool

**Cause:** Packet filters on the remote access policy profile are preventing the flow of IP traffic.

**Solution:** Verify that there are no configured TCP/IP packet filters on the profile properties of the remote access policies on the remote access server (or the RADIUS server if Internet Authentication Service is used) that are preventing the sending or receiving of TCP/IP traffic.

You can use remote access policies to configure TCP/IP input and output packet filters that control the exact nature of TCP/IP traffic allowed on the remote access connection. Verify that the profile TCP/IP packet filters are not preventing the flow of needed traffic.

**See also:** Configure IP options

**Callback is not working.**
**Cause:** Callback is improperly configured on the dial-in properties of the user account.

**Solution:** Verify the callback configuration on the dial-in properties of the user account.

**See also:** Configure caller ID and callback

**Cause:** **Link control protocol (LCP) extensions** is disabled on the **PPP** tab for the properties of the remote access server.

**Solution:** Enable **Link control protocol (LCP) extensions** on the **PPP** tab for the properties of the remote access server.

**See also:** View properties of the remote access server

**Cause:** Callback numbers may be truncated when a remote access server running Windows NT 4.0 requests dial-in properties of a user account in a Windows 2000 native or Windows Server 2003 domain.

**Solution:** Reconfigure callback numbers on the dial-in properties of the user account.

**See also:** Configure caller ID and callback

For more information on troubleshooting remote access VPN connections, see Troubleshooting remote access VPNs.

To reset Routing and Remote Access back to default values, see Reset the Routing and Remote Access service.

**Not sure how to configure the remote access server to make as well as receive remote access connections (dial out and dial in).**
**Cause:** The remote access server is either not configured as a remote access server (dial in) or not configured for demand-dial routing (dial out).

**Solution:** In the console tree, right-click the remote access server, and click **Properties**. On the **General** tab, make sure that both the **Router** and the **Remote access server** check boxes are selected, and make sure that the **LAN and demand-dial routing** option is selected under **Router**. If these options are not configured, configure them, and click **OK**. In the console tree, click **Network Interfaces** and make sure that you have configured demand-dial interfaces for dialing out to the servers or services to which you want to connect. If no demand-dial interfaces exist, right-click **Network Interfaces**, and click **New Demand-dial Interface**.

**See also:** Demand-dial routing; Routing over VPN connections; Overview of remote access

# Connection Manager Administration Kit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Manager Administration Kit**
Connection Manager is a versatile client dialer and connection software that you can customize by using the Connection Manager Administration Kit (CMAK) wizard.

- For an overview of CMAK and its audience, see Connection Manager Administration Kit Welcome.
- Before installing CMAK, see Before You Start: Understanding Connection Manager and the Administration Kit.
- To understand your customization options and how to select the options most appropriate to your needs and those of your users, see Phase One: Planning.
- For information on how to design and develop custom graphics, text, programs, and other elements to support your needs, see Phase Two: Developing Custom Elements.
- For information on how to create your service profiles using the CMAK wizard and advanced customization techniques, see Phase Three: Running the CMAK Wizard and Creating a Service Profile.
- For information on delivery options and methods, see Phase Four: Preparing for Integration, Delivery, and Installation.
- For information on testing and verifying your customizations, see Phase Five: Testing Your Deliverables.
- For problem-solving instructions and other support information, including answers to Frequently Asked Questions about Connection Manager and CMAK, see Phase Six: Supporting Your Customers.

# Connection Manager Administration Kit Welcome

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Welcome**
Before you run the Connection Manager Administration Kit (CMAK) wizard, you should plan and develop the Connection Manager custom elements that you want in your client dialer. Then, simply run the CMAK wizard, answering the questions and providing the information about your custom elements. The CMAK wizard then builds a service profile, which is a set of files which you distribute to your users so that they can easily install and run your custom version of Connection Manager.

**Audience**
The CMAK wizard and this documentation is designed for the network administrators, information officers, and other team members responsible for the design, development, testing, distribution, and support of connection software for the customers who connect to your Internet or network service.

**About this documentation**
The CMAK wizard provides defaults that support quick and easy creation of a basic Connection Manager service profile. If you want to use all of the defaults and do not want to take advantage of the many customizable elements, you probably do not need this documentation.

However, if you want a custom dialer that promotes your brand and supports more advanced functions, such as automatic phone book updates and support for virtual private networking, you need this documentation to help you effectively plan and implement your custom elements.

### Notes

- The information in this documentation is based on the assumption that you have already installed the CMAK software. For more information on evaluating Connection Manager, see the Microsoft Web site.

- You can display this documentation directly from the CMAK wizard by clicking **Help**. When you click **Help** in a pane of the CMAK wizard, the Help topic associated with that specific pane appears.

# Before You Start: Understanding Connection Manager and the Administration Kit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Before you start: Understanding Connection Manager and the Administration Kit**
Connection Manager provides a customized remote access client so your customers can connect to your service using connection features that you define. The Connection Manager Administration Kit (CMAK) wizard simplifies the customization process by providing a graphical interface that you use to specify custom elements for your service. It then builds your customized installation package for you. This section covers the information you need to know before starting the customization process:

- Connection Manager components and system requirements
- Connection Manager features
- Connection Manager user interface
- The Connection Manager Administration Kit Wizard
- The Connection Manager Administration Kit and the customization process
- Connection Point Services and Connection Manager

# Connection Manager components and system requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Manager components and system requirements**
This section covers the system requirements for Connection Manager and its components.

Connection Manager requires at least 2 megabytes (MB) of free disk space and one of the following:

- Windows 95
- Windows 98
- Windows NT Workstation 4.0
- Windows 2000
- Windows Millennium Edition
- Windows XP
- Windows Server 2003 family
- Internet Explorer 4.01, 5.0, 5.5, or 6.0. Users do not have to have Internet Explorer set as their default browser or use the software.

Disk-space requirements for each system vary, based on the elements included in the service profile. For information on determining the size of the complete deliverable, see Planning for effective implementation.

For dial-up connections, Connection Manager also requires a 28.8 modem or faster connection. Connection Manager can automatically configure the modem. If you include Connection Manager in the service profile, the user's modem is configured as needed to support Connection Manager.

Connection Manager requires Transmission Control Protocol/Internet Protocol (TCP/IP) (for peer connections in a network), Telephony Application Programming Interface (TAPI) (for dialing support), and Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP) (to establish a VPN connection), but when Connection Manager installs, it automatically installs the appropriate level of software on computers with operating systems that do not have the appropriate level. There are two exceptions. On Windows NT 4.0, the user must manually set up PPTP (in **Remote Access Setup**). For information that you can provide to your users on how to set up PPTP on Windows NT 4.0, see Providing user documentation. On Windows 95, Windows 98, Windows Millennium Edition, and Windows NT 4.0, users must install Microsoft L2TP/IPSec VPN Client in order to set up L2TP. For more information about Microsoft L2TP/IPSec VPN Client, see the Microsoft Web site .

If your service profile supports PPTP VPN connections, any users who are running Windows 95 on their computers must have the latest connection software installed on their computers to establish a VPN connection. For more information on this requirement, see Implementing VPN support. For information that you can provide to your users to update their connection software, see Providing user documentation.

**Notes**

- It is recommended that your clients install the latest service pack for their operating systems before installing a Connection Manager service profile.
- If your Connection Manager profile includes phone book downloads, your users must have a version of Internet Explorer installed on their computers, as noted above. Users do not have to have Internet Explorer set as their default browser or use the software.

# Connection Manager features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Manager features**

Connection Manager 1.3 provides support for local and remote connections to your service using a network of access points, such as those available worldwide through Internet service providers (ISPs). If your service requires secure connections over the Internet, you can use Connection Manager to establish virtual private network (VPN) connections to your service. Following are some of the features supported by Connection Manager 1.3.

**New features of Connection Manager 1.3**

Connection Manager provides the following new features:

- **VPN server selection** If you include a VPN file and configure VPN entries, your users will be able to choose what VPN server to use when making a connection.
- **Log files** Log files for a connection can be created, cleared, and viewed by the user.
- **Favorites** Users can define the properties of a connection, save these settings for future use, and pick the favorite to use from the Connection Manager user interface. For example, a user can configure location properties and dialing rules for connecting to the office from home and from a frequent business travel location, and then the user can choose between these two sets of settings when connecting.
- **Provide user name and password for the service profile** An administrator can pre-populate the User Name and Password fields of a Connection Manager profile for first-time or permanent use.
- **Terminal windows** Connection Manager 1.3 supports the use of a terminal window on all supported platforms except Windows NT 4.0.
- **Network Access Quarantine Control** Formerly a feature installed through the Windows Server 2003 Resource Kit Tools, the Remote Access Quarantine Server (RQS) and Remote Access Quarantine Client (RQC) components are available as part of Windows Server 2003 operating systems with Service Pack 1 installed. RQC is available as a pre-defined custom action in the Connection Manager Administration Kit wizard, although you will have to customize it to meet the needs of your individual network. For more information about Network Access Quarantine Control, see Related Topics.

**General connection support**

Connection Manager supports a variety of features that both simplify and enhance implementation of connection support for you and your users. Most of these features can be incorporated using the Connection Manager Administration Kit (CMAK) wizard:

| Feature | Support |
|---|---|
| Routing table updates | You can reconfigure user routing tables in order to better manage your network traffic and security. |
| Automatic proxy configuration | You can configure user proxy settings to ensure that the user has appropriate access to internal and external resources while connected to your service. |
| Branding | You can customize the graphics, icons, messages, Help, and phone book support in Connection Manager to provide an identity and support that are unique to your service or corporation. You can include custom logos, customer support, and phone book information to identify and represent your company. |

| | |
|---|---|
| Custom actions and monitored applications | You can incorporate custom functionality, including your own programs to enhance the connection experience for your users. These programs can be automatically run at various points during the connection process, such as when users log on or disconnect. Connection Manager also supports pre-connect and pre-tunnel actions. And you can set up monitored applications to automatically disconnect after the program ends. |
| Multiple user support for each service profile | You can provide support for users who share computers. User profiles allow two or more people to use the same computer and the same service profile. Credentials are maintained, based on the logon ID of the user, so users do not have to re-enter them for each connection. |
| Simplified distribution | Using the CMAK wizard to automatically build your service profile (the customized software required for your users to run Connection Manager), you create a self-installing executable file that can be distributed on compact disc or downloaded by your users. |
| Custom phone books | You can specify the phone books to be provided to your users. You can download your phone book to users and provide automatic phone book updates when your users log on. And to simplify maintenance, you can combine existing phone books by merging existing service profiles. |

**VPN connection support**

In addition to the support for basic dial-up connections, you can also use the CMAK wizard to incorporate support for VPN connections using a tunneling protocol to tunnel through a public network (such as when dialing into an ISP to access a corporate server). Connection Manager can create a VPN connection through a pre-existing dial-up session, local area network (LAN), or digital subscriber line (DSL).

| Feature | Support |
|---|---|
| VPN server selection | You can allow your users to choose a VPN server to use when connecting to your service. |
| Direct connections | If you specify support for VPN, you can provide support for direct connections (sometimes referred to as *always on* or *persistent* connections). Support for direct connections includes support for cable, asymmetric digital subscriber lines (ADSLs), and other types of direct connections. |
| Protocols | You can specify whether or not to support VPN, which enables your users to connect to your service using Point-to-Point Tunneling Protocol (PPTP) or, if running Windows 2000, Windows XP, or a member of the Windows Server 2003 family, Layer Two Tunneling Protocol (L2TP) with IPSec encryption as the tunneling protocol. These protocols enable improved security for direct and dial-up connections. |

**Note**

- You must build profiles on a server operating system that is designed for the same processor set as the client computers on which users will install the profile. For example, if you want to build a profile for users running Windows XP Professional x64 Edition, you must build that profile on a server that is running an x64-based version of Windows Server 2003.

**See Also**

**Concepts**

IAS Network Access Quarantine Control

**Other Resources**

Network Access Quarantine Control

Step-by-Step Guide for Creating and Testing Connection Manager Profiles in a Test Lab

Step-by-Step Guide for Setting Up Network Quarantine and Remote Access Certificate Provisioning in a Test Lab

# Connection Manager user interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Manager user interface**
The Connection Manager user interface consists of a logon dialog box and a set of property dialog boxes with which users can customize their connection settings. What users see depends largely on how you customize your service profile.

**Logon dialog box**
The logon dialog box is often the only Connection Manager screen your users see when they access your service. In this dialog box, users specify their credentials (such as user name and user password) so they can connect to your network servers. You can customize the service name, the graphic, support information, and many other elements of this dialog box. Users can customize this logon box if they configure favorite connection settings. If two or more favorites have been configured, the user can select between them from the logon dialog box.

**Properties dialog box**
The *ServiceProfileServiceName* **Properties** dialog box appears when a user clicks **Properties** in the logon dialog box.

### Note

- The title bar of this dialog box contains the service name of the service profile in front of the word **Properties**, but it is referred to as simply the **Properties** dialog box throughout this documentation.

**Properties, General tab**
On the **General** tab, users can specify the way their system connects to an Internet service provider to access your service:

- If the service profile that you provide supports virtual private network (VPN) connections, users can choose to use a direct connection that has already been established through the Internet or use a dial-up connection, and they can specify phone numbers and other information needed to establish the dial-up connection.

- If the service profile that you provide does not support VPN connections, users do not see the direct connections option, but they are still able to specify phone numbers and other information for their dial-up connections.

Using this tab, users specify the following information for dial-up connections:

- **The phone number to be dialed to establish a connection.** Users can type a number or click the **Phone Book** button to select a phone number.

  When users click the **Phone Book** button, they can select a Point of Presence (POP) from the list in the **Phone Book** dialog box. Only those access numbers that are supported by the selected service type, country (or region), and state (or province) appear in the list.

  Additional numbers appear in the **More access numbers** box only if you configure the phone book to support this second set of numbers (such as the phone numbers you specify as surcharge numbers when you use Connection Point Services to create your phone book). If the phone book has not been set up with additional access numbers for the selected geographic location, the box will be empty. If you specify in the Connection Manager Administration Kit (CMAK) wizard that custom text is to be displayed next to the **More access numbers** box, the text appears only when numbers are available for the selected geographic location.

- **Dialing rules for the dialing location.** If users click **Dialing Rules** on the **General** tab, they can specify information about the location from which they are currently dialing, such as calling card numbers and codes for accessing outside and long-distance lines, or information about other locations for future use.

- **The type of device to be used to connect to your service.** Only installed and available device types appear in the **Connect using** list.

- **Whether to save the settings on the General tab as a favorite for later use.** Users can name and save all the

current settings on the General tab and then select one of these favorite sets of settings from the Connection Manager logon dialog box.

**Properties, Internet Logon tab**

On the **Internet Logon** tab, the user types a user name and password to connect to the Internet for a VPN connection. This tab appears only if, in the CMAK wizard, you:

- Specify support for VPN connections, and
- Do not select the **Use the same user name and password for VPN and dial-up connections** check box.

**Properties, Options tab**

On the **Options** tab, users can change the connection defaults for establishing and keeping active connections. Users can change both the number of redial attempts that occur and the amount of idle time that elapses before Connection Manager ends the connection. Users can also choose whether to enable logging for this connection. If a log file has been created, users can choose to view or clear it.

**Properties, VPN tab**

On the **VPN** tab, users can choose a VPN server for the connection. This tab appears only if, in the CMAK wizard, you do each of the following three things:

- Specify support for VPN connections.
- Select the **Allow the user to choose a VPN server before connecting** check box.
- Include a VPN file with your service profile.

For more information, see Implementing VPN support and Incorporating VPN entries.

**Properties, Advanced tab**

On the **Advanced** tab, users can enable Internet Connection Sharing, Internet Connection Firewall, Windows Firewall, or any combination of these services for the connection. This tab only appears on computers running Windows XP or a member of the Windows Server 2003 family. Users who are not members of the Administrators group will not see this tab.

For more information, see Internet Connection Firewall and Connecting to the Internet in a home or small office network.

### Notes

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

**Additional screens**

Depending on how you customize the service profile, other screens appear during the connection process. For example, you can specify whether a license agreement appears when users install Connection Manager. You can also incorporate Help screens and other messages into the service profile.

# The Connection Manager Administration Kit Wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The Connection Manager Administration Kit Wizard**
The Connection Manager Administration Kit (CMAK) wizard guides you through the process of building service profiles that exactly meet your business needs. Before you can build these profiles, you must install the CMAK wizard.

**New features in the CMAK wizard**
The CMAK wizard offers many new features, including some that help you do the following:

- provide routing table updates that apply only while clients are connected to your server ("split tunneling")
- Automatically configure Internet Explorer proxy settings for a client computer
- Enable clients to choose which virtual private network (VPN) server to use when they connect
- Include a pre-shared key for L2TP/IPSec authentication
- Encrypt a pre-shared key profile with a personal identification number (PIN) that the user must type in before the profile can be installed

Additionally, the CMAK wizard gives you greater customization and control when you build the service profile, including:

- More types of custom actions and full configuration of all custom actions from within the wizard
- An advanced customization pane in which you can edit .cms and .cmp files as the profile is built
- Improved interface and text throughout the wizard

**Differences from earlier versions of the CMAK wizard**
If you have used an earlier version of the CMAK wizard, you should be aware of the following issues:

- If you have upgraded your operating system, you might have to manually migrate some profile files.
- If you have manually edited the .inf file of a profile created with an earlier version of the CMAK wizard, you must reapply the edits to the .inf file if you change the profile with the current version of the CMAK wizard.
- You can no longer use the CMAK wizard to incorporate an animation in your service profile.
- If you add a program under the **[RunPostSetupCommandsSection]** of the .inf file of your profile, you must add the file as an additional file when you edit the profile with the CMAK wizard.

**Installing the CMAK wizard**
To install Connection Manager components, including the CMAK wizard, on a computer running a member of the Windows Server 2003 family:

1. Open the Windows Components wizard.
2. In **Components**, click **Management and Monitoring Tools** (but do not select or clear the check box), and then click **Details**.
3. Select the **Connection Manager Administration Kit** check box.
4. Click **OK**, click **Next**, and then click **Finish**.

**Notes**

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.

To install the CMAK wizard on a computer running Windows XP Professional, you must install the Windows Server 2003 Administration Tools Pack. To install the Administration Tools Pack:

1. Insert the installation CD for a member of the Windows Server 2003 family into the CD-ROM drive of the computer running Windows XP Professional.

2. Click **Perform additional tasks**.

3. Click **Browse this CD**.

4. Double-click the I386 folder.

5. Double-click the Adminpk.msi file and follow the instructions on your screen.

**Starting the CMAK wizard**

To start the CMAK Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Connection Manager Administration Kit**.

# The Connection Manager Administration Kit and the customization process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The Connection Manager Administration Kit and the customization process**

Connection Manager is *brandable*, which means that you can customize the installation package that you deliver to your customers so that Connection Manager reflects the identity of your organization. You determine which functions and features you want to include, and how Connection Manager will appear to your customers. You do this by using the Connection Manager Administration Kit (CMAK) wizard to build custom service profiles.

A service profile consists of all of the files needed by Connection Manager to run the installation file and then enable users to establish connections with your service. You can maximize or minimize the identification of your service or organization, depending on what you decide to include in a service profile. For example, you can include custom corporate logos or other graphics, custom icons, and your own online Help. You can also add monitored applications, specify actions to run before, during, and after a connection, and customize other features available in Connection Manager.

To customize the features and support provided in Connection Manager for your service, you can customize your service profile in two ways:

- **Run the CMAK wizard (required to build a service profile).**

  You can use the CMAK wizard to customize most features of Connection Manager at the same time that you build a service profile. For example, you can use the wizard to specify the name and graphics that your users see when running Connection Manager, whether virtual private network (VPN) connections to your service are available to your users, and whether your custom connect actions should run when a connection is established. How much you customize depends on your needs and those of your users. When you finish running the wizard, it automatically builds your customized service profile for you.

- **Use advanced customization techniques (optional).**

  In some instances, there may be features or functions that you want to customize that are not available through the wizard. To add these features or functions, you can use the **Advanced customization** pane to edit the .cms and .cmp files when you build a service profile. After the service profile has been built, you can manually edit any service profile file, but you must run the CMAK wizard again after editing the files in this way in order to incorporate your changes. For example, you can change case sensitivity for passwords by editing the value for the **PasswordHandling** key in the **Advanced Customization** pane.

  For more information, see Methods of editing service profile files.

Each service profile that you create includes the ready-to-use custom features that you specify in a complete, self-installing package so users can seamlessly connect to your service.

Although using the CMAK wizard to build your custom service profiles is simple, creating the custom elements requires careful planning and development. This guide covers the six phases of planning, development, and implementation:

| Phase | Tasks | Persons responsible |
|---|---|---|
| 1 | Decide and document which features to customize for Connection Manager. The information you need to complete these tasks is covered in Phase One: Planning. | Network administrator or technical lead |
| 2 | Design and develop a phone book and any other custom elements you want, such as graphics, text, and related programs, as needed. The information you need to complete these tasks is covered in Phase Two: Developing Custom Elements. | Technical support team |

| | | |
|---|---|---|
| 3 | Use the CMAK wizard to build each custom service profile, and then use advanced customization techniques, as appropriate, to complete the customization of each, as covered in Phase Three: Running the CMAK Wizard and Creating a Service Profile. | Technical lead or other member of technical support team |
| 4 | Verify distribution readiness, including addressing any integration requirements, as covered in Phase Four: Preparing for Integration, Delivery, and Installation. | Technical lead or other member of technical support team |
| 5 | Test your Connection Manager deliverables, including all custom elements, on all target platforms and operating systems, using the process and criteria covered in Phase Five: Testing Your Deliverables. | Testers or other members of technical support team |
| 6 | Provide effective troubleshooting and other support to your customers. Phase Six: Supporting Your Customers. | Customer support team |

# Connection Point Services and Connection Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Point Services and Connection Manager**
Connection Point Services (CPS) enables you to automatically distribute and update custom phone books. These phone books contain one or more Point of Presence (POP) entries, with each POP supplying a telephone number that provides dial-up access to an Internet access point. The phone books give users complete POP information, so when they travel they can connect to different Internet access points rather than being restricted to a single POP.

Without the ability to update phone books (a task CPS handles automatically), users would have to contact technical support to be informed of changes in POP information and to reconfigure their client dialer software.

CPS has two components:

- **Phone Book Administrator**, a tool used to create and maintain the phone book database and to publish new phone book information to the Phone Book Service.

- **Phone Book Service**, an Internet Information Services (IIS) extension. Phone Book Service automatically checks subscribers' or corporate employees' current phone books and, if necessary, downloads a phone book update.

CPS is available with the Windows Server 2003 family.

For information on setting up a phone book, see Providing phone book support. For more information on CPS, see Connection Point Services.

# Phase One: Planning

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phase one: Planning**
The information in this section supports the planning that you need to do before you run the Connection Manager Administration Kit (CMAK) wizard to create the customized service profiles that you need:

- Security information for the Connection Manager Administration Kit. This section provides an overview of basic security considerations when you design and deploy your customized service profiles.

- Deciding what to customize. This section provides an overview of the Connection Manager features and functions that can be customized and included in your service profile.

- Planning for effective implementation. This information covers the factors that you should consider before starting the planning for an individual service profile, so you will understand how your decisions will impact the service-profile deliverables. It includes information on the impacts of integrating your service profile with another installation package or otherwise tailoring your delivery and installation methods.

- Creating profiles in multiple languages. This information covers the factors you should consider when designing a multinational or multilingual deployment of a service profile.

# Security information for the Connection Manager Administration Kit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for the Connection Manager Administration Kit**

You can increase the security of your remote access solution by using the Connection Manager Administration Kit (CMAK) wizard to customize and control the connection experience for your users. By customizing the connection, you can control how your users connect to your network, and you can simplify troubleshooting.

Before you create a profile, you should understand your remote access solution, including its security requirements and limitations. You can consider remote access security in three parts: securing servers, including remote access servers and the computers on which you store profiles; understanding and securing the computers on which users will install profiles; and designing, creating, and distributing the profile itself. For more information about securing servers, see Security information for remote access, Security issues for VPN, and Security.

When you design your service profile, you should consider the following:

- **Who can create or edit service profiles using the CMAK wizard**. A user must be a member of the Administrators group to install the CMAK wizard. A user must be a member of either the Administrators or the Power Users group to run the CMAK wizard.

- **Who can edit service profile files using a plain text editor**. Service profile files are plain text files that users can edit using a plain text editor, such as Notepad, instead of using the CMAK wizard. By using a plain text editor instead of the CMAK wizard, users do not need to be members of the Administrators or the Power Users group to edit or to delete these files. You can help prevent users from changing these files on the server by limiting access to the directories that contain service profile files. Recognize that, after you distribute the service profile, users can edit the service profile files on their own computers. For more information, see Methods of editing service profile files and Including Connection Manager in custom applications.

- **Which operating systems your users will use**. You should consider two things regarding the operating systems on your users' computers: the level of security inherent in the operating system itself, as well as the security features that you can configure for the operating system. Not all operating systems support all Connection Manager features. For example, Windows 98 does not support user certificates. Therefore, if you want to deploy a highly secure VPN solution, you might need to limit your deployment to users running Windows XP on hard disks formatted with NTFS.

  Depending on how your users have configured their computers, you might have these security considerations:

  - Whether the user's hard disk is formatted with FAT or NTFS. FAT systems are not as secure as NTFS systems.

  - What authentication and tunneling protocols are supported. Not all operating systems support all of the same protocols as the Windows Server 2003 family. For more information, see Dial-up networking clients and Virtual private networking clients.

  - Whether your users will install the profile for individual use only or for all users of that computer (Windows 2000, Windows XP, and the Windows Server 2003 family only). You should encourage your users to install the profile for individual use only so that only the user who installed the profile and members of the Administrators or System Operators groups can modify the service profile files. If a user installs a profile for all users, all users of that computer can modify or delete the service profile files.

  - Whether the profile allows users to save the password for the profile. You can hide the check box that allows users to save passwords by configuring the **HideRememberPassword** key. For Windows XP and the Windows Server 2003 family, you can configure the **GlobalCredentials** key to prevent users from saving a user name and password for any other users of that computer. This key prevents other users from using another user's credentials. The connection will close automatically if Fast User Switching is used.

- If necessary, you can create different service profiles for different operating systems, restricting access for users with less secure systems.

- **Whether your service profile will include a phone book or multiple phone books**. If your service profile will include one or more phone books, you will need to create the phone book with Phone Book Administrator (PBA), and you will need to include a Phone Book Service (PBS) server in your deployment. For more information, see Security information for Connection Point Services.

- **How you will distribute your service profile**. You can distribute your service profile in several ways, including on a Web site and on portable media such as floppy disks. Depending on what kind of service profile you create, you might want to consider how to control the distribution of your profile. For example, if you create a profile with a pre-shared key, you will want to limit distribution of the profile to authorized users, particularly if you do not encrypt the profile with a personal identification number (PIN). For more information about distributing service profiles, see Planning for effective implementation.

  For detailed information about deploying Connection Manager, see "Deploying Remote Access Clients Using Connection Manager" at the Microsoft Windows Resource Kits Web site.

# Deciding what to customize

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deciding what to customize**

You use the Connection Manager Administration Kit (CMAK) wizard to do most of the customization for a service profile. Before running the CMAK wizard, you need to decide what to customize and then provide the information and files needed to run the wizard and implement your decisions. To help prepare for creating a service profile, you can complete the planning worksheet that is available in the Microsoft Windows Server 2003 *Resource Kit*. For more information about Connection Manager, see "Deploying Remote Access Clients Using Connection Manager" at the Microsoft Windows Resource Kits Web site.

Although you can run the wizard and build your service profile without providing any custom information except for the name of your service and a file naming convention to be used for your service profile files, you can customize many other things to meet the needs of your service. For example, you can:

- Reconfigure user routing tables with your service profile in order to better manage your network traffic and security.
- Reconfigure user proxy settings to ensure that the user has appropriate access to internal and external resources while connected to your service.
- Merge previously created service profiles into the one you are building, in order to incorporate phone books and other features from those service profiles.
- Provide a customer-support phone number or other customer-support information in the Connection Manager logon dialog box.
- Provide a realm name (prefix or suffix) that is appended to the user name (to specify network routing and authentication information needed by your Internet service provider, without requiring your users to understand and incorporate the information themselves).
- Specify dial-up entries to be associated with phone numbers in the phone book you provide to your users.
- Support virtual private network (VPN) connections over the Internet.
- Incorporate custom actions that run at specific points during the connection process.
- Incorporate your own graphics in place of the default graphics provided in Connection Manager logon and phone book dialog boxes.
- Provide a phone book and phone-book update capabilities to your users.
- Incorporate your own custom icons in place of the default Connection Manager icons.
- Customize the menu associated with the notification area shortcut when your users are connected.
- Customize the default Help file provided with Connection Manager.
- Include Connection Manager 1.3 in your service profile.
- Incorporate an end-user license agreement (EULA) in your service profile that users must accept in order to install your service profile.
- Include additional support files in your service profile, such as user documentation files and custom programs and files needed to support custom actions that you have defined.

- Enhance security by enabling server callback to users.

You should consider each of these decisions for each service profile you create. Depending on the target audience and other factors, your decisions will vary for each service profile. Before you implement your decisions, you should also consider the effects of your implementation and any potential integration issues. For information on planning for integration, distribution, installation, and other implementation considerations, see Planning for effective implementation. This section provides the information you need to support effective and efficient planning and development efforts.

**Note**

- You must build profiles on a server operating system that is designed for the same processor set as the client computers on which users will install the profile. For example, if you want to build a profile for users running Windows XP Professional x64 Edition, you must build that profile on a server that is running an x64-based version of Windows Server 2003.

# Planning for effective implementation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for effective implementation**
When you make decisions about customization, you should be aware of how your decisions affect Connection Manager implementation methods. When developing plans for a service profile, consider implementation impacts on:

- The number of service profiles to be developed.
- The distribution methods to be used.
- Streamlining installation with additional components.
- Integrating your Connection Manager service profile with an Internet Explorer installation package.
- The size of the service profile.
- Phone book support to be implemented.
- Maintenance and update requirements for the service profile, including all files included in the profile.
- The operating systems and processors of the client computers for which you want to create service profiles.

The following sections discuss how customizing your service profile can affect how it is implemented. For additional information on setting up multiple profiles and how each service profile is set up and maintained, see Setting up your service profiles.

**How many service profiles do you need?**
You can create additional service profiles to support specialized needs of specific groups or personnel. For example, if you support multiple geographic locations, you might want to create separate profiles for the specific locations, incorporating a unique local phone book in each. Or, you might want to create separate profiles for each operating system supported by your service.

If a user installs more than one service profile, each profile is installed in a separate folder. All profiles share the same copy of Connection Manager software.

To provide support for another language, you must run the Connection Manager Administration Kit (CMAK) wizard on a computer with an international version of a Microsoft operating system that uses the appropriate language.

**Note**

- You must build profiles on a server operating system that is designed for the same processor set as the client computers on which users will install the profile. For example, if you want to build a profile for users running Windows XP Professional x64 Edition, you must build that profile on a server that is running an x64-based version of Windows Server 2003.

**What distribution methods are appropriate?**
You can distribute your service profile by compact disc or downloadable file (over a private network or the Internet). You should consider the size of your customized profile, including all files that you add to it, when you determine which distribution method to use. Remember that download time can significantly affect users, especially those running modems at relatively slow speeds.

If you are using the Internet Explorer Administration Kit 5.5 (IEAK 5.5) to distribute Internet Explorer 5.5 to your users, you can simplify installation for them by including a Connection Manager profile in an Internet Explorer installation package. Files included in the Internet Explorer installation package should not be duplicated in the Connection Manager profile.

**What is the size of your service profile?**

It is important that the size of your service profile be compatible with your distribution method. Consider the features you want and how you plan to distribute the profile: by compact disc or download.

For example, if you distribute by download, consider how the download time affects users. In some cases, you might decide to use or offer users an alternate distribution method. The CMAK wizard does not support profiles that span multiple disks and does not prompt you if the profile does not fit on one disk.

A service profile that includes Connection Manager uses about 800 kilobytes (KB) of disk space. A profile without Connection Manager can use as little about 200 KB. The total size of your service profile depends on your customization options, such as the method you use to provide phone books, whether you include custom actions or other optional programs, and which graphics you incorporate.

**Note**

- If disk space allows, include Connection Manager in the service profile to ensure that your users always have the latest version.

**How can you streamline installation for your users?**
You can simplify installation and updates for your users by incorporating programs and files that can be automatically installed and configured as part of the Connection Manager installation package. You can tailor these features to the individual needs of specific user groups to control the size of the service profile that you distribute and to promote the effective use of your service. For example, if you support a corporate account with marketing and accounting departments, you might provide separate profiles for each department, so that you could provide the marketing department with a phone book covering the entire sales area, and you could provide the accounting department with a phone book containing only the numbers of the local access points.

**How can you integrate your service profile with an Internet Explorer installation package?**
Some of your planning decisions will depend on whether you are integrating Connection Manager and your custom service profile with an Internet Explorer installation package or other implementation package. You should determine your delivery method before creating a service profile, to determine any impacts of integration.

For example, if you decide to integrate your Connection Manager service profile with an installation package created with the Internet Explorer Administration Kit (IEAK), you might want to coordinate the graphics included in your profile with those you include in the installation package.

Also, if you decide to deliver your Connection Manager service profile with an Internet Explorer installation package, you must create the Connection Manager service profile first. Unlike the IEAK 5 wizard, which had a button to call CMAK directly, the IEAK 5.5 wizard has no direct link to CMAK.

After building your service profile, check the information in Phase Four: Preparing for Integration, Delivery, and Installation to determine additional customization that may be needed to integrate your service profile with other packages.

**What phone book support is required?**
You can provide a single phone number for user connection, or you can provide additional support for roaming access by providing a phone book of access numbers. Using Connection Point Services, you can provide phone book support with automatic updates, but you need to create the initial phone books and decide how the phone book support is to be implemented before running the CMAK wizard. For more information, see Connection Point Services and Connection Manager.

**What is required to maintain and use existing service profiles?**
You can incorporate many types of files in a service profile. You can edit or merge other profiles, include program files to support custom actions and monitored applications, and provide user documentation to customize Connection Manager appropriately for your users.

In order to maintain these service profiles so that you always use the appropriate version of an incorporated file, you should understand how the CMAK wizard incorporates and uses such files. For information about how files and folders for your profile are set up, see Setting up your service profiles.

When you edit a service profile that contains a merged service profile, the CMAK wizard attempts to use the latest version of the merged profile. If the CMAK wizard cannot find the merged profile in the appropriate \Program Files\CMAK\Profiles*ServiceProfileFileName* folder, it defaults to the current version of the merged profile that was originally stored with the profile being edited.

Plan your profiles, especially those with incorporated files, so that the information and elements are easily maintained and updated as needed.

# Creating profiles in multiple languages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating profiles in multiple languages**
You can use language-specific versions of the Connection Manager Administration Kit (CMAK) wizard to create service profiles in different languages. If you create your profiles with the same wizard options in different languages, you can provide your users with the same connection experience independent of language. For example, you can allow your users who speak French and your users who speak Japanese to connect to your service or private network in the same way.

**Creating profiles for users of operating systems other than Windows XP or the Windows Server 2003 family**
You can purchase installation CDs for the Windows Server 2003 family in a variety of languages. However, each CD is specific to only one language, known here as the system language. To build a profile for users who are not running Windows XP or a member of the Windows Server 2003 family, you will need to purchase a separate CD for each language in which you intend to create a profile. For example, if you want to create a German service profile for users who are not running Windows XP or a member of the Windows Server 2003 family, you must purchase and install the German version of a Windows Server 2003 operating system.

The reason that you must purchase a separate CD is that profiles built with the Windows Server 2003 family require Connection Manager 1.3 to function. Users who are not running Windows XP or a member of the Windows Server 2003 family will not have Connection Manager 1.3, so you must include it when you build a profile for them. The version of Connection Manager 1.3 you include will appear in your system language, regardless of whether it matches the system language of your users. The only way to change your system language to match the system language of your users is to install a version of the operating system that is specific to the system language of your users. Even if you create a profile on a Multi-Language User Interface (MUI) version of the Windows Server 2003 family, Connection Manager still appears in the system language that is specific to that version.

**Creating profiles for users of Windows XP or the Windows Server 2003 family**
For users who are running Windows XP or a member of the Windows Server 2003 family, you do not need to include Connection Manager 1.3 when you create a profile. However, you must change your input language and your **Language for non-Unicode programs** settings so that they match the system language of your users. Your input language determines which characters you can type in the CMAK wizard and as you create phone books. The CMAK wizard itself and Phone Book Administrator will still appear in your system language. If your input language differs from your system language, a warning message appears when you start the CMAK wizard. The message should remind you that, if you include Connection Manager 1.3, it will appear on your users' computers in your system language, which might not match that of your users. If you do not plan to include Connection Manager 1.3, you can close the warning message and proceed with the wizard.

For more information about changing your input language, see Specify regional and language options and user's name during Setup and Regional and Language Options overview.

# Phase Two: Developing Custom Elements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phase two: Developing custom elements**

Depending on what you customize, you might need additional resources to prepare your elements. For example, if you want to include a phone book in your service profile, you must get copies of the phone book and region files. Or, if you want to include custom icons and bitmaps, you might need additional programs or graphics specialists to prepare them.

This section provides guidelines for the technical team members who are responsible for creating the graphics, information, programs, and other special elements to customize Connection Manager for your needs:

- Setting up your service profiles
- Specifying service and file names
- Providing a realm name
- Merging phone books and other features from existing service profiles
- Implementing VPN support
- Incorporating VPN entries
- Including a pre-shared key
- Providing phone book support
- Incorporating custom dial-up entries
- Including routing table updates
- Using automatic proxy configuration
- Incorporating custom actions
- Incorporating custom graphics
- Incorporating custom icons
- Customizing the notification area shortcut menu
- Providing custom Windows Help
- Providing support information in the logon dialog box
- Adding Connection Manager 1.3 to the service profile
- Incorporating a license agreement
- Including additional files
- Providing user documentation

- Advanced Customization

These guidelines can be reproduced and distributed to the individuals responsible for doing the design and development work.

# Setting up your service profiles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up your service profiles**

Each time you run the Connection Manager Administration Kit (CMAK) wizard, you create a service profile. You can create as many profiles as needed to support each specific target audience. After creating an initial profile, you can use the CMAK wizard to edit it. You can save the edited profile under the same name to overwrite the existing profile, or you can specify a different service name and file name to save it to a separate profile (preserving the existing profile as it was before you started the edit).

**Notes**

- To create a service profile for another language, use the version of the operating system and the CMAK wizard that are appropriate for the language.
- When you edit a service profile created using a previous version of Connection Manager, the CMAK wizard automatically upgrades the profiles to Connection Manager 1.3. Any changes you have made manually to the .inf file (using advanced customization techniques) will not be retained in the upgraded file. However, the CMAK wizard creates a backup of the old .inf file before creating the updated file.
- If you change the file name, be sure to change the service name also. If a user installs two service profiles with the same service name but different file names, the service profiles do not work correctly.

Each service profile should be specific to a target audience. For example, you might set up separate profiles if:

- You want to make each service profile specific to a single operating system.
- You specify phone books to support each of your distinct geographic areas or corporate organizations.
- Different users require different connection methods (dial-up versus direct).
- You need to support virtual private network (VPN) connections only for specific users.
- The individual corporations or groups you support each want to use their own logos or other proprietary symbols or text.
- The help desk or other support is not the same for everyone and you want to provide specific information to each user group.
- You want to establish different administrative authorities for the groups you support.

**Note**

- To simplify administration of multiple service profiles, you can merge existing profiles into a parent service profile. See Merging phone books and other features from existing service profiles.

You must run the CMAK wizard again for each service profile you want to create, including each service name or file name you want to present to your users. For more information on service and file names, see Specifying service and file names.

It is easier to create additional service profiles if you edit an existing profile and then change the service name and file name.

When you create a service profile, the CMAK wizard copies all files that are incorporated in the profile into a \Program Files\CMAK\Profiles*ServiceProfileFileName* folder. When you edit existing profiles, the CMAK wizard always pulls the incorporated files from the ...*ServiceProfileFileName* folder. To make changes to the incorporated files, edit the incorporated file in the ...*ServiceProfileFileName* folder and then run the CMAK wizard again to update the service profile. Or, if you edit the file that is not in the ...*ServiceProfileFileName* folder, start the CMAK wizard, add the updated file (by browsing to find it), and then finish running the wizard.

**Note**

- Connection Manager supports user installation of multiple service profiles, running multiple instances of Connection Manager, and multiple user profiles for each service profile. If a user installs more than one service profile, each is installed in a separate folder. You can provide as many profiles to your users as necessary to support your service, but each must have a different service name and file name.

# Specifying service and file names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Specifying service and file names**
You must specify a *service name* and a *file name* in the CMAK wizard to identify your service profile.

**Service name**
The service name identifies your service to the users. This name is limited to 35 characters. This name appears at several points when the user runs your service:

- In the title bar text for the logon dialog box for your service.
- In dialog boxes that appear during installation of Connection Manager and the service profile.
- As the desktop icon name for your service.
- On the program button on the taskbar that appears when the program is running.
- In the ToolTip that appears when the mouse points to the notification area shortcut, which is located at the far right of the taskbar.
- Within the service-profile files, for entries such as the default dial-up entry.
- In the **Network Connections** (Windows XP and the Windows Server 2003 family), **Network and Dial-up Connections** (Windows 2000), or **Dial-up Networking** (Windows 95, Windows 98, and Windows Millennium Edition) folder, as the name of the remote access connection.

This name is often a company or other business name. For example, the fictitious company Awesome Computers, which is used for many of the examples in this documentation, might specify "Awesome Computers" as the service name. Or, if Awesome Computers wanted to set up separate service profiles for each department, the service name for one of them might be "Awesome Computers - Marketing."

### Important

- Do not use any of the following characters in your service profile name: *** \ / : ? " < > | [ ] & ´ % ' `**

**File name**
The file name is the name given to the files created by the CMAK wizard when building the service profile. The file name can be up to 8 alphanumeric characters, with the following limitations:

- Do not use any of the following: **! , ; <space> * = \ / : ? ' " < > | . % + [ ] $ # & @ ( ) { } ´ ` ~**
- Do not use extended character sets in a file name. A file name created with extended character sets (such as those containing ü, é, and other extended characters) might not be properly installed.

### Notes

- Make sure that the service and file names are different from all other service and file names that you provide to your customers. If two service profiles on the same computer have the same service or file name, the associated connection icons do not work correctly.
- When editing an existing profile, if you do not want to write over the existing service profile, you must change the file name.

The file name is used for the following files created by the CMAK wizard, including:

- The folder containing the files needed to install Connection Manager and the service profile, located at \Program Files\CMAK\Profiles*ServiceProfileFileName*.
- The self-extracting executable (.exe) file.
- The .cms file, which contains the configuration data for Connection Manager features, such as the names of the phone book and connection methods.
- The .cmp file, which stores user-specified information, such as the name of the .cms file to be used.
- The .inf file, which is used for the setup and installation of Connection Manager on the user's computer.
- The .sed file, which specifies the contents of the .cab file and how to compress the files when the service profile is built.

The file name for the service profile is also used for the folder created on the user's computer when the .exe file is run.

For more information on what files are created by the CMAK wizard, see Using the CMAK wizard to build a service profile.

# Providing a realm name

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providing a realm name**
Realm names are used for network routing and authentication. They provide the identification necessary to forward authentication requests to the server that holds the user's credentials. Realm information is only used for dial-up connections.

Using the CMAK wizard, you can specify a realm user-name prefix (such as **Awesome/**) or suffix (such as **@Awesome**) to be added to the user name that a user enters when connecting to your service. The realm name includes any separator characters, such as **@** or **/**. If you specify the realm, users do not have to provide it themselves because Connection Manager automatically places the realm name before or after the user's name so that the user does not have to understand network routing in order to log on.

If your service uses multiple networks or ISPs for access, you need to provide different realm information for each. In this case, create service profiles for each of the networks and then merge them together in a *top-level* profile.

**Notes**

- If the top-level profile contains realm information, it will take precedence over the information in the component profiles.
- For more information on merging profiles to incorporate multiple realms, see Merging phone books and other features from existing service profiles.

Internet Authentication Service (IAS) supports the use of realm names for remote authentication. You can use the realm name as authentication routing information for remote roaming users (users who travel and connect through different Internet access points). When the user connects through a network access server (NAS), the NAS sends the realm name, the user name, and the user password to the authentication server. For more information on IAS, see Internet Authentication Service.

# Merging phone books and other features from existing service profiles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Merging phone books and other features from existing service profiles**

If you have information in existing service profiles that you need in the service profile you are building, you can use the Connection Manager Administration Kit (CMAK) wizard to *merge* much of the information from existing profiles into the profile you are building. A profile that contains information from other service profiles is called the *top-level* profile. A service profile that has its information merged into another service profile is called a *component* profile.

**Why merge service profiles?**

Merging service profiles is especially useful for incorporating information from multiple phone books, including service types, Point of Presence (POP) names, and access numbers, and to consolidate different Internet access points that are covered by more than one of your phone books. By merging existing service profiles into the service profile you are building, you can make several dissimilar networks appear to be a single, cohesive service.

For example, you can merge service profiles from different geographical locations to create a worldwide profile. Or you can create profiles that have specific administration requirements and network configurations, including authentication and TCP/IP characteristics and, after they are created, you can merge them together within a top-level service profile. Merging service profiles is especially useful where multiple entities, whether corporations or Internet service providers (ISPs), share in providing a service. The following examples show how you can use merged service profiles to simplify network access through diverse access points:

- **Example 1: Creating a service profile to support access to the Internet from remote service areas.** If a local ISP contracts with a roaming consortium to provide dial-up services outside the ISP's primary service area, the local ISP could create a top-level service profile containing the local ISP's network information and merging the service profiles of the other members of the roaming consortium. Each component service profile could contain the individual phone-book and network attributes, including a unique realm name, of an individual ISP. Merging all of the component service profiles would support the need for users of the local ISP service to roam outside of the ISP's primary service area by enabling them to connect through the cooperating ISPs. All ISPs can maintain their own service profiles, simply redistributing profiles to other members of the consortium whenever they are updated. The consortium itself might also create a top-level service profile aggregating the services of all member ISPs.

- **Example 2: Implementing a corporate virtual private network (VPN).** Corporations using VPN connections for a session by establishing a tunnel through the Internet to the corporation often use ISP-provided dial-up access to support corporate access from remote locations. By merging service profiles, the corporation could combine its own phone book of internal dial-up numbers with those of remote ISPs contracted to provide access points for the remote corporate users, enabling VPN connections to be established even for dial-up connections through the public network access point of the ISP. The corporation and each ISP could maintain its own service profiles, but the corporation could merge all of the ISP service profiles into the corporate profile. The distribution of the top-level service profile by the corporation to all corporate users would support access from multiple locations. This access would appear to users as if it were made using the corporation's own access points.

**What information is merged**

The following table shows which features of a component service profile are merged into the top-level service profile and how the CMAK wizard incorporates the merged information. Only the items from the component service profile that are listed in the **Merged features** column are recognized and used by the top-level service profile.

| Merged features | Implementation notes |
|---|---|
| File name | Files contained in the component service profiles retain the names used by the component service profiles. |
| Merged profiles | You can merge only one level of service profiles. For example, if you merge service profile A, which contains component profiles B and C, these second-level component profiles (B and C) are ignored. |

| | |
|---|---|
| Realm name | If a realm name is specified in the top-level service profile, it will be used in all instances, and all realm names specified in the component service profiles are ignored. If no realm name is specified in the top-level service profile, the realm names specified in the component service profiles are used for connections that are made using entries from the component service profiles. |
| Dial-up entries | The script specified for each dial-up entry must have a unique name than all other scripts contained in the top-level service profile and any component profiles. |
| Phone books | If the same service types exist in more than one profile, the information is merged. If duplicate states (or provinces), or POPs exist among top-level and component service profiles, Connection Manager removes the duplicates when it runs the top-level service profile. Connection Manager treats two phone book entries as duplicates if they display the same information in the user interface (except that Connection Manager is not case sensitive). If two phone book entries appear as duplicates but have separate configuration data in the POP, Connection Manager still removes one of the phone book entries. |
| Download phone book | Each time the user establishes a connection, Connection Manager automatically checks for the phone books with each server that is running Phone Book Service (PBS). PBS checks the phone book version and, if necessary, downloads phone book updates to the user. Phone books specified in the top-level and component phone books are downloaded, as appropriate to the respective service profiles. |
| URL for downloads | The URLs specified in the top-level and component service profiles are used for downloads to phone books for the respective service profiles. |

The top-level service profile does not merge information related to the following features of the component profile. To affect these features, you must specify them in the top-level service profile:

- Service name
- Support information specified for the logon dialog box
- VPN support
- Custom actions
- Monitored applications
- Custom graphics for the logon dialog box and **Phone Book** dialog box
- Message to display next to the **More access numbers** box when alternate access numbers are available for the selected geographic location
- Custom icons
- Customization options for the notification area shortcut menu
- Custom Help files
- Option to include Connection Manager software in the service profile
- License agreement
- Additional files

**Note**

- To support VPN connections for component service profiles, in the **VPN Support** pane of the CMAK wizard, under **Establish a virtual private network (VPN) connection when the user dials a phone number from these phone books**, select the **Phone books from the merged profiles** check box.

**Where the files are located**
When you create a service profile, the CMAK wizard copies all files referenced in the service profile, including the component service-profile files, into the \Program Files\CMAK\Profiles*ServiceProfileFileName* folder for that service profile.

When you use the CMAK wizard to merge features, the top-level service profile includes references to the component (merged) service profiles. For example, for phone books in merged service profiles, the resulting service profile installs multiple files in the folder specified for the profile rather than combining the phone books into a single file. However, when users select a POP involving phone books merged from two other service profiles, the users see the information from all the phone books as a single phone book.

When you edit a top-level service profile, the CMAK wizard attempts to use the latest version of the component service profile. If the CMAK wizard cannot find the merged service profile in the \Program Files\CMAK\Profiles*MergedServiceProfileFileName* folder (for the top-level service profile), it uses the current version of the merged service profile that was originally stored with the service profile being edited.

**Special considerations for building merged profiles**
When building merged service profiles:

- If you merge a Connection Manager 1.0 service profile into a Connection Manager 1.3 service profile, the CMAK wizard automatically updates the merged information to Connection Manager 1.3. Any changes you have made manually to the .inf file (using advanced customization techniques) will not be retained in the upgraded file. However, the CMAK wizard creates a backup of the old .inf file before creating the updated file.

- You should provide separate service profiles for each network that has unique characteristics or requirements.

- You must ensure that all scripts specified for dial-up entries in the merged service profile are available when you build the top-level service profile.

# Implementing VPN support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Implementing VPN support**

A virtual private network (VPN) connection provides remote access to private networks over the Internet (or other network). A VPN connection is established using a tunneling protocol to establish a tunnel through the Internet to the private network. To establish a VPN connection, users can connect directly to the private network through the tunnel by either dialing into a local Internet service provider (ISP) or by using a pre-existing connection to the Internet (direct connection).

This means that remote users can connect using worldwide Internet access points to access private resources as easily as local users. For corporations, it means that employees have remote access over the Internet to their corporate private networks, usually through a local phone number. By having users create VPN connections rather than dial-up connections, the company does not incur the considerable expenses associated with long distance telephone service without compromising security.

You can use the Connection Manager Administration Kit (CMAK) wizard to set up VPN connections, including:

- Make a direct connection to a private network.
- Tunnel to a private network over a public network, such as the Internet.
- Select a VPN server for the connection.

Using the CMAK wizard, you can set up support for VPN connections using Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP). To use these VPN connections, you must have a remote access server on the private network configured to support the appropriate protocol.

**Important**

- A VPN connection requires the computer on which Connection Manager is run to be configured to use the same network protocols that the private network uses.
- VPN connections on computers running Windows Millennium Edition, Windows 98, or Windows 95 require that users have the latest **Dial-up Networking** software installed on their computers. Users can obtain the current **Dial-up Networking** software by downloading it as a Windows Update from the Microsoft Web site. You should document this requirement for your users. For more information on verifying and installing **Dial-up Networking** software, see Providing user documentation.

You can create a service profile with which users can establish direct VPN connections using a tunnel over pre-existing connections. For example, users who are already connected to the Internet through DSL or a cable modem can then connect to your private network using a VPN-only profile you created. To create a VPN-only profile:

1. Start the CMAK wizard.
2. On the **VPN Support** pane, select the **Phone book from this profile** check box, and specify a VPN server (by DNS name or IP address) or specify a VPN file (by typing the full path to the file).
3. Configure the VPN entry or entries with the correct security and addressing information for your network.
4. Do not specify a phone book file on the **Phone Book** pane, and clear the **Automatically download phone book updates** check box. You do not need to configure the default dial-up entry for the profile.
5. On the last pane of the wizard, select the **Advanced Customization** check box.
6. On the **Advanced Customization** pane, click the profile .cms file, click **[Connection Manager]**, and set the value of the **Dialup** key to zero.

7. Click **Apply**, and finish the wizard.

This process will create a VPN-only profile, without a **General** tab in the **Properties** dialog box for the profile. Users of your profile will not see any phone or dialing information.

Setting up a dial-up profile to support VPN connections might add an **Internet Logon** tab to the **Properties** dialog box. On this tab, users type the user name and password for the Internet service provider. In the logon dialog box, users type a private network user name, password, and logon domain.

To implement VPN connections in your service profiles, you must specify the address of at least one server to be used and how to handle authentication. As you run the CMAK wizard:

- You must have the VPN server address, specified either as a Domain Name System (DNS) name or as an IP address.
- If you do not want to let the server assign addresses for DNS and Windows Internet Name Service (WINS) servers, you can specify the addresses to be used.
- You must specify whether users enter the same user name and password for both logging on to the Internet using a dial-up connection and logging on to the private network server using a VPN connection. If you specify that the same credentials are to be used for both, the **Internet Logon** tab does not appear and the user only has to enter credentials once to connect.
- You must specify the VPN security settings to be used for connections made with this profile. By default, the security setting is selected to **Configure both basic and advanced settings**, which requires a configuration for profiles installed on Windows 95, Windows 98, Windows NT 4.0, and Windows Millinnium Edition, and a different configuration for profiles installed on Windows 2000, Windows XP, and the Windows Server 2003 family. Windows 2000, Windows XP, and Windows Server 2003 operating systems have more advanced security features and can support a higher level of security for your connection. You can choose to require the use of advanced security settings only, although this setting might prevent some clients on other operating systems from connecting with your profile. You can also choose to only use basic security settings that all operating systems that support Connection Manager can utilize. However, this means that all profiles will use the basic security settings, regardless of whether the operating system can support a higher level of security.

For more information on keys to be used with direct connections, see Advanced Customization.

**Enabling user choice of VPN server**

You can allow your users to choose from multiple VPN servers when they connect to your service. For example, a user could select a VPN server that is close to his or her location, or a user could choose a VPN server that has higher security settings. For service profiles that support this option, a **VPN** tab appears in the properties dialog box for the service profile. This tab contains a customized text message and a list of VPN servers.

To enable your users to choose a VPN server from a list when they connect to your service, you must create a VPN file before you start the CMAK wizard. A VPN file is a text file that you can create with any plain-text editor, such as Notepad. The table below identifies sections, keys, and values that compose a VPN file.

| **[Section] or key name** | **Value** |
|---|---|
| [Settings] | Section header for the keys that contain VPN settings. |
| Default | The friendly name of the default VPN server for this service profile. If you create this key but do not give it a value or if you do not create this key, the user must select a VPN server the first time the user connects with the service profile. |
| UpdateURL | The URL of the Web server that contains updates for this VPN file. If this key is added, a post-connect action is added automatically to the service profile. This post-connect action updates the VPN file. |
| Message | The text message that appears on the **VPN** tab in the properties dialog box for the service profile. This message cannot exceed 256 characters, and it must be a single paragraph. If this key is created but left blank, default text is used. |

| | |
|---|---|
| [VPN Servers] | Section header for the keys that identify VPN servers. |
| *FriendlyName* | The name of this key is the friendly name of one of the VPN servers available to your users. The value of this key is the DNS or IP address of the VPN server for which you named the key. If you want this VPN server to use a specific set of network and security settings, you must follow the DNS address with a comma and the name of a VPN entry that you will edit in the **VPN Entries** pane of the CMAK wizard. If you do not specify a VPN entry, the VPN server will use the settings for the default VPN entry. |

If the fictional company Awesome Computers wanted to provide its users with a choice of VPN servers, an administrator might create a VPN file similar to the example below.

```
[Settings]
default=Awesome Computers HQ
UpdateURL=http://awesomecomputers.microsoft.com/VPNfile.txt
Message=Please select a server from the following list. You might want to choose a
server closest to your location or to your data. Windows 98 users should choose Awesome
Computers Windows 98.

[VPN Servers]
Awesome Computers HQ=awesomecomputers.microsoft.com
Awesome Computers New York=ny.awesomecomputers.microsoft.com
Awesome Computers Spain=es.awesomecomputers.microsoft.com,Awesome International VPN
Settings
Awesome Computers Madagascar=ma.awesomecomputers.microsoft.com,Awesome International
VPN Settings
Awesome Computers Windows 98=awesomecomputers.microsoft.com,Awesome 98 VPN Settings
```

In the above example, the VPN servers Awesome Computers HQ and Awesome Computers New York use the settings specified in the default VPN entry. Awesome Computers Spain and Awesome Computers Madagascar will use the settings specified in the VPN entry named Awesome International VPN Settings. A VPN server has been created for clients running Windows 98, and it uses the settings specified in the VPN entry named Awesome 98 VPN Settings. Users will see the friendly names of the VPN servers you have defined.

**See Also**
**Concepts**
Incorporating VPN entries
Advanced Customization

# Incorporating VPN entries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incorporating VPN entries**
You can configure connection and security settings for the VPN server or servers that support your service profile. If you customize these settings, you can better handle unique network authentication or routing requirements of particular service profiles.

Before you can configure a VPN entry to support a VPN server or servers, you must specify whether users of your service profile must use a particular server or they can choose between servers. You specify this in the **VPN Support** pane of the Connection Manager Administration Kit (CMAK) wizard. If you allow your users to choose between servers, each server you specify appears in a list on the **VPN** tab of the properties dialog box for the service profile. For more information, see Implementing VPN support.

You can specify the following settings for each VPN entry:

- Whether to enable file and printer sharing
- Whether to enable clients to log on to a network
- Whether to assign static or dynamic IP addresses for Domain Name System (DNS) and Windows Internet Name Service (WINS) servers
- Whether to make this connection the default gateway for the client
- What security settings to apply, based on the client operating system
    - What authentication methods to use
    - What encryption methods to use

You must edit each VPN entry to provide the network and security configuration necessary for a client to connect to any of your VPN servers that use that entry. For example, if you want to configure a profile to require smart cards, on the **VPN Entries** pane, click the VPN entry that you want to configure, click **Edit**, click the **Security** tab, click **Configure** for advanced security settings, click **Use Extensible Authentication Protocol (EAP)**, and then click **Smart Card or other certificate (encryption enabled)**.

All VPN servers use the default VPN entry unless you specify a different entry in the VPN file. If you do not edit the default VPN entry, it will use default settings provided by the CMAK wizard, which might not be appropriate for your network.

### Notes

- The VPN entry that you add or edit in **VPN Entries** must have exactly the same name as one you specified in the VPN file. If you add a VPN entry (by clicking **New**), you must add a matching entry manually to the VPN file to make the VPN server available to your users.
- On the **General** tab, the option **Disable file and printer sharing** affects only computers running Windows NT 4.0, Windows 2000, Windows XP, or a member of the Windows Server 2003 family.
- On the **General** tab, the option **Enable clients to log on to a network** affects only computers running Windows 95, Windows 98, or Windows Millennium Edition.

**See Also**
**Concepts**
Implementing VPN support
Advanced Customization

Remote Access Overview
VPN Tunneling Protocols
Remote Access Authentication Methods

# Including a pre-shared key

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Including a pre-shared key**
You can use a pre-shared key instead of a certificate for L2TP/IPSec authentication of your VPN clients that are running Windows XP or a member of the Windows Server 2003 family. Pre-shared keys do not require a public key infrastructure (PKI) for deployment, but they are a relatively weak authentication method. You can increase the security of your pre-shared key deployment by encrypting the pre-shared key with a personal identification number (PIN), which your users must enter before the profile will install.

**Advantages and disadvantages of pre-shared keys**
Using pre-shared key authentication does not require the infrastructure investment of a PKI. Pre-shared keys and certificates can be used on the same remote access server. If the pre-shared key on the remote access server is changed, no pre-shared key VPN clients will be able to access the server until the pre-shared key on the clients are changed by reissuing and reinstalling the service profile. Unlike certificates, the origin and history of a pre-shared key cannot be determined. Administrators who want a long-term, strong authentication method should consider using certificates.

**Considerations when creating pre-shared keys**
A pre-shared key is a string of text that is configured on both the remote access server and the client. A remote access server running a member of the Windows Server 2003 family can be configured with only a single pre-shared key. All L2TP/IPSec VPN clients that connect to the remote access server using a pre-shared key must use the same pre-shared key. The pre-shared key can be any non-null string of any combination of between 8 and 256 characters. If there is any deviation between the pre-shared key on the remote access server and the pre-shared key presented by the VPN client, client authentication will fail. For more information about routing and remote access on computers running any member of the Windows Server 2003 family, see Routing and Remote Access.

**Increasing the security of a pre-shared key by encrypting it with a PIN**
You can increase the security of the distribution of your pre-shared key profile by encrypting the pre-shared key with a PIN. By distributing the service profile and the PIN separately, you will reduce the chances that unauthorized users can install the profile and gain access to your network. A PIN must contain no fewer than 4 but no more than 15 characters. Be sure to distribute the PIN to your users in a secure manner. This PIN will be requested only when the user installs the profile.

**Configuring a service profile to use a pre-shared key**
To configure a service profile to use a pre-shared key:

1. On the **VPN Entries** pane, click the VPN entry that you want to configure to use a pre-shared key, and click **Edit**.
2. On the **Security** tab, under **Common security settings**, click **Configure**.
3. Under **VPN strategy**, click **Use L2TP/IPSec if available**, and select the **Use a pre-shared key when using L2TP/IPSec** check box. Finish configuring any other aspects of security that your profile requires, and click **OK**.
4. Edit each VPN entry that requires a pre-shared key, and click **Next**.
5. On the **Pre-shared Key** pane, type a pre-shared key.
6. Do one of the following:
   - Type a PIN to encrypt the pre-shared key.
   - Clear the **Encrypt the pre-shared key using a PIN** check box.
7. Click **Next**, and finish building the profile.

# Providing phone book support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providing phone book support**

Connection Point Services (CPS) provides support for creating, maintaining, and updating phone books. You can also create your own phone book (without CPS), but the complexity of the format and requirements makes this difficult if you have more than a single Point of Presence (POP). For more information, see Connection Point Services and Connection Manager.

**Specifying how the phone book is to be provided**

Using the Connection Manager Administration Kit (CMAK) wizard, you can either:

- Include a phone book in your service profile, without support for automatic upgrades. To do this, in the CMAK wizard:

  1. In the **Phone Book** pane, browse to the phone book (.pbk) file that you want to include in your service profile. Make sure that you have a region file with the same name located in the same folder as the .pbk file. For example, if your phone book file name is AwesomeE.pbk, the region file should be AwesomeE.pbr.

  2. Make sure that the **Automatically download phone book updates** check box is **not** selected.

- Include a phone book in your service profile, with support for automatic upgrades. To do this, in the CMAK wizard:

  1. In the **Phone Book** pane, browse to the phone book (.pbk) file that you want to include in your service profile. Make sure that you have a region file with the same name located in the same folder as the .pbk file. For example, if your phone book file name is AwesomeE.pbk, the region file should be AwesomeE.pbr. Make sure the **Automatically download phone book updates** check box is selected.

  2. In the **Phone Book Updates** pane, type the URL for the Phone Book Service (PBS) server from which upgrades are to be downloaded.

- Do not include a phone book in your service profile, but download the phone book (and upgrades) to your users later. To do this, in the CMAK wizard:

  1. In the **Phone Book** pane, select **Automatically download phone book updates**, but leave the **Phone book file** box blank.

  2. In the **Phone Book Updates** dialog box, type the name of the phone book (.pbk) file to be downloaded (but do not include the file name extension) and the URL for the PBS server from which it (and the upgrades) can be downloaded.

  If you use this method of providing the phone book, you will also be able to automatically download updates to your users, using the same phone book name and URL.

The phone book name that you provide must match the name of the phone book file created with CPS. This file name consists of a maximum of 8 alphanumeric characters and cannot include spaces or a file name extension. If you specify support for automatic updates, each time the user connects to your service, the PBS server automatically downloads any updates to the phone book.

**Using more access numbers**

Using Phone Book Administrator in CPS, you can set up phone books to list certain POPs under **Access numbers** and other POPs under **More access numbers**. If you want to separate phone numbers into two sets, you must configure the phone book for this purpose before you use it in a service profile. You can separate the numbers in different ways, such as by separating standard access numbers from 800 and 888 numbers or by separating no-surcharge access numbers from access numbers that incur a surcharge when a user uses one of them to connect. For more information, see Connection Point Services.

If your phone book specifies a second set of access numbers, those numbers appear in the **More access numbers** box in the **Phone Book** dialog box.

To customize information about these numbers:

- Provide a line of text to appear in the dialog box, next to the **More access numbers** box. This line of text should be no more than about 100 characters. (The maximum size varies based on line breaks and other formatting considerations.) This line of text appears only when alternate access numbers are available for the selected geographic area.

- Edit the definition Help for the list of alternate phone numbers. The default Help states that the user can be charged extra for using alternate phone numbers. If this is not true for your alternate phone numbers, edit these Help topics in the Cmtrb32.rtf file: IDH_PHONEBOOK_ACCESSNUMS and IDH_PHONEBOOK_MORE. For more information on customizing Help, see Providing custom Windows Help.

**Advanced customization for the phone book**
Through advanced customization, you can also:

- Control which service types (such as modem or ISDN) appear in the **Phone Book** dialog box.

- Change the message displayed during phone book updates (such as "Downloading new Awesome Computers backup company phone book.").

- Control the delay before downloading phone book updates.

For more information on advanced customization techniques, see Advanced Customization.

**Creating a single-POP phone book**
You can include a default phone book that contains a single phone number. This is useful if you want your customers to access your service through one POP. When you want to create a phone book for multiple POPs, you should use CPS.

To create a single-POP phone book, you must create both a phone book (.pbk) file and an region (.pbr) file. In the CMAK wizard, you specify only the name of the .pbk file, but the wizard requires that a .pbr file with the same name be located in the same folder as the .pbk file.

**Creating a .pbr file for a single-POP phone book**
The region file you create for your single-POP phone book must be in plain-text (.txt) and must have two lines:

- The first line contains a number specifying how many regions are in the region file.

- The second line contains the name of the state or province (or other geographical area). Do not use commas or semi-colons in this line. The file name (excluding the extension) cannot exceed eight characters.

For example, if the POP were located in Washington, the contents of the single-POP .pbr file would be:

```
1
Washington
```

**To create a .pbr file for a single-POP phone book**

1. Start a text editor such as Notepad.

2. Type a **1** and press **Enter**.

3. Type the name of the region and press **Enter**.

4. Save the file in plain-text (text-only) format, naming it in the following format:

   *PBname*.pbr

**Creating a .pbk file for a single-POP phone book**
The file you create for your single-POP phone book:

- Must be located in the same folder as the .pbr file. The file name (excluding the extension) cannot exceed eight characters and cannot contain spaces.

- Must be in plain-text (.txt), comma-delimited format, and must include fields in the following sequence and without any spaces (if necessary, please scroll this page to read all 11 fields):

  *Index,Country/dependencyID,RegionID,POPName,AreaCode,PhoneNumber,MinBaud,MaxBaud,Reserved,ServiceType,DUNEntry*

  For example, a single-POP phone book for our fictitious company, Awesome Computers, might have the following specification:

```
1,44,1,AwesomeE - Basecamp,999,555-0134,9600,56000,0,0,Awesome Dun
```

**Notes**

- The TAPI codes are updated as needed. The TAPI codes available to a user depends on the operating system they are running and which service updates have been installed. Ensure that the TAPI code you provide is correct for the level of operating system that your users are running.

- For more information about what to specify for each field, see Connection Point Services.

# Incorporating custom dial-up entries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incorporating custom dial-up entries**

The dial-up entries referenced in the **Dial-up Networking Entries** pane of the Connection Manager Administration Kit (CMAK) wizard correspond to dial-up entries specified in the Phone Book Administrator (PBA) of Connection Point Services. The CMAK wizard provides a method of customizing how users access specific Points of Presence (POPs). For more information, see Connection Point Services and Connection Manager.

Customizing the way in which specific entries are handled can provide more effective handling of unique network authentication or routing requirements for POPs in your phone book. For each entry, you can specify the Domain Name System (DNS) or Windows Internet Name Service (WINS) addresses to be used for POPs associated with that entry and, if desired, a script to be associated with it.

Before you run the CMAK wizard, you must decide which dial-up entries you want to customize, and the options to use for each.

For each dial-up entry, you must specify:

- The name of the entry. If you have already specified a phone book to use with this service profile, any dial-up entry names used by the POPs (as configured in PBA) will automatically appear in the **Dial-up networking entries** box on this pane. If you have not yet specified a dial-up entry to use with a POP but want to add an entry, you must type the entry name exactly as it will appear in the **Dial-Up Networking entry** box in PBA. If you want to set up a default for your service profile, the name must be specified as your *ServiceProfileServiceName* (conforming to the file naming rules for a service profile).

- Whether to let the server assign the addresses or to use specific DNS addresses, WINS addresses, or both. If you want to assign specific addresses, you can specify both primary and secondary DNS and WINS IP addresses.

- If your service requires it, a script (.scp) file can be associated with this entry.

- You must specify the security settings to be used for connections made with this profile. By default, the security setting is selected to **Configure both basic and advanced settings**, which requires a configuration for profiles installed on Windows 95, Windows 98, Windows NT 4.0, and Windows Millinnium Edition, and a different configuration for profiles installed on Windows 2000, Windows XP, and the Windows Server 2003 family. Windows 2000, Windows XP, and Windows Server 2003 operating systems have more advanced security features and can support a higher level of security for your connection. You can choose to require the use of advanced security settings only, although this setting might prevent some clients on other operating systems from connecting with your profile. You can also choose to only use basic security settings that all operating systems that support Connection Manager can utilize. However, this means that all profiles will use the basic security settings, regardless of whether the operating system can support a higher level of security.

For example, for Awesome Computers (the service name of our fictitious company):

- To set up custom properties for a specific entry, you could specify **Awesome Dun** in the **Dial-up networking entries** box in the CMAK wizard. In this case, all POPs that have **Awesome Dun** specified as the phone book entry would use the addresses, scripts, or both as specified in the CMAK wizard for the Awesome Dun entry.

- To set up default properties for all other POPs in the phone book of the Awesome Computers service profile (any without a specific entry), you could specify **Awesome1** in the **Dial-up networking entries** field. In this case, all phone book entries without explicit dial-up entries would use the addresses and scripts specified in the CMAK wizard for this default entry.

**To set up custom dial-up entries**

- Add phone book entries to the appropriate POPs:

1. Start **Phone Book Administrator** and click the phone book you want.

2. Add or edit a POP to include the dial-up entry on the **Settings** tab.

For more information, see Connection Point Services.

- Using the CMAK wizard, build a service profile and set up dial-up entries as follows:

  - If the entry you want to configure already exists, select it and click **Edit**.

  - If the entry does not exist, in the **Dial-up Networking Entries** pane, click **New**. In **New Dial-up Networking Entry**, type the same dial-up entry as the one you entered in the phone book file. If you want to assign specific DNS or WINS addresses or specify scripts for the entry, specify the appropriate information.

**Notes**

- Remember, the dial-up entry that you set up in your service profile must either exactly match the specific entries you set up in the phone book or, if defining the default, it must be the default service name.

- Under the **General** tab, the option "Disable file and printer sharing" only affects computers running Windows NT 4.0, Windows 2000, Windows XP, and members of the Windows Server 2003 family.

- Under the **General** tab, the option "Enable clients to log on to a network" only affects computers running Windows 95 and Windows 98, but it has security implications for networks that support Windows NT, Windows 2000, Windows XP, and the Windows Server 2003 family.

- For additional customization options for these dial-up entries, see Advanced Customization.

**Using phone-book Dial-up Networking Entries**

Connection Manager reads the entry specified in **Dial-up Networking Entries** associated with the phone number, and uses the custom properties (such as a DNS server and script) that you specified in the CMAK wizard for that entry to establish a connection.

If Connection Manager does not detect a specific dial-up entry for the phone number in the service profile, it then tries to find the default entry (the **Dial-up Networking** entry named *ServiceProfileFileName* in the .cms file).

# Including routing table updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Including routing table updates**
You can alter client routing tables with your service profile in order to better manage your network traffic and security. You can include a routing table update file in the service profile, provide a URL to an update file hosted on a server, or both. You can choose to undo any routing changes you made when the connection ends.

**Advantages of including routing table update files**
Routing table update files are plain-text files that you can create and edit in any plain-text editor, such as Notepad. The files you create or edit modify the routes on client computers. You can manage significant network security requirements by including update files in your service profile. The benefits include the ability to:

- Control user access to a corporate intranet
- Route all users through the same remote access server
- Redirect Internet traffic for VPN connections without using a proxy server

**Distribution channels**
You can provide update files for use in a service profile in one of three ways:

- Include a file with the service profile package (static)
- Include a URL to a file with the service profile and host the file on a Web server (dynamic)
- Include a static file with the service profile and include a URL to an additional file

You can distribute an update file by including it with a service profile. When the user installs the profile, the file is installed in the same directory as the profile. You should consider two factors when designing a profile using this method of distribution. The routes in the file on the client computer are not secure, and you must create and distribute another profile if you want to update the file.

**Important**

- If the user has not logged on with administrative credentials for the client computer, routing table updates will not be applied. The user will still be able to connect using your service profile, but the routes on the client computer will not be updated.
- The user will not be able to connect if the client computer cannot find a routing table update file you included in a service profile.
- You should not change the value of the **Gateway_On_Remote** key when using routing table updates. By default, the dial-up entry and the VPN entry will have **Make this connection the default gateway** selected. You should leave this default in place, and remove any gateways by using the **REMOVE_GATEWAY** command in the routing table update file itself.

You can provide a URL to an update file hosted on a Web server. The file is downloaded every time the user completes the connection. The user cannot easily review any file you distribute in this way. You can also update that file at any time.

If you provide a URL to an update file hosted on a Web server, you must make two additional decisions. The first is whether to terminate the connection if the server is unreachable. The second is whether to delete the default gateway on the client computer.

You can provide an update file with the service profile and a URL to an additional update file. This approach allows more flexibility in your design. You can add or delete routes as necessary to the hosted file, but you cannot change the file included

with the service profile without recreating and redistributing the profile. You still need to decide whether to terminate the connection if the Web server is unreachable and whether to delete the default gateway on the client computer.

**Syntax of routing table update files**

Routing table update files are constructed as plain-text files. Each line in those files adds or deletes a route. The syntax of each line is as follows:

*Command Destination* **MASK** *Netmask Gateway* **METRIC** *Metric* **IF** *Interface*

Certain parameters can contain the value of **default**. In those cases, the appropriate information from the client computer is used. The table below provides more information on each parameter:

| Parameter | Description |
|---|---|
| *Command* | Commands are one of the following:<br><br>**ADD** adds a route.<br><br>**DELETE** deletes a route.<br><br>**REMOVE_GATEWAY** removes the default gateway. |
| *Destination* | Specifies the destination IP address on the target network. |
| **MASK** | Specifies that the next parameter is the netmask value. |
| *Gateway* | Specifies the gateway. Should contain the value of **default**. |
| **METRIC** | Specifies that the next parameter is the cost for the destination. Should contain the value of **default**. The METRIC parameter is optional, and its default value is 1. |
| **IF** | Specifies that the next parameter is the interface number. Should almost always contain the value of **default**. |

Therefore, if you want the connection to follow a different route, you could include a line in your update file similar to the following example:

**ADD** 192.168.123.231 **MASK** 255.255.255.0 default **METRIC** default **IF** default

To delete a route, you could include a line in your update file similar to the following example:

**DELETE** 192.168.123.232 **MASK** 255.255.255.0 default **METRIC** default **IF** default

**Important**

- If your update file includes a delete command for a non-existent route, the connection will fail.

To remove the default gateway, you can add the following line to the update file:

**REMOVE_GATEWAY**

Unlike adding or removing a route, no further syntax is required after the **REMOVE_GATEWAY** command.

**Routing table precedence**

During the connection process, update files are applied in a specific sequence. You should keep this sequence in mind when you design your profile.

1. All update files are retrieved.

2. If you provide an update file with the service profile (static file), routing information from that file is applied to the client computer first.

3. If you provide a URL to an update file, routing information from that file is applied to the client computer after the static file. The downloaded file is then immediately deleted.

4. If you included the **REMOVE_GATEWAY** command in either update file, this command is executed after all other routes

have been applied.

Because of this order of precedence, a dynamic file could add routes, replace routes specified in the static file included with the service profile, or both.

If you use the default interface value, the routing changes you specify are reversed when the user disconnects.

# Using automatic proxy configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using automatic proxy configuration**
Conflicting or improperly configured proxy settings on user computers can restrict user access to external sites while the user is connected to a corporate network. By reconfiguring user proxy settings, you can ensure that the user has appropriate access to internal and external resources.

**Requirements for proxy configuration**
In order for you to use this feature:

- Your users must have Internet Explorer 4.01, Internet Explorer 5.01, Internet Explorer 5.5, or Internet Explorer 6 installed on the computer.
- Your network must include a proxy server.

Internet Explorer 4.0 requires a slightly different approach than Internet Explorer 4.01, Internet Explorer 5.01, Internet Explorer 5.5, and Internet Explorer 6. With Internet Explorer 4.0, proxy settings apply to all connections on the computer. With Internet Explorer 4.01, Internet Explorer 5.01, Internet Explorer 5.5, and Internet Explorer 6, different settings can apply to different connections on the same computer. If any of your users run Internet Explorer 4.0, you should choose to restore the computer's previous proxy settings when the user disconnects.

**Note**

- You can choose to restore the previous proxy settings when users disconnect without adversely affecting users running later versions of Internet Explorer. Therefore, if you are building a profile for use with multiple versions of Internet Explorer, you should choose to restore the client settings.

**Syntax of proxy configuration files**
In order to use automatic proxy configuration in a service profile, you must create a plain text file that specifies the following values:

| Key | Value |
|---|---|
| [Automatic Proxy] | No value. Section header for automatic proxy settings. |
| [Manual Proxy] | No value. Section header for manual proxy settings. |
| AutoProxyEnable | 0 = Automatic proxy detection is disabled.1 = Automatic proxy detection is enabled.<br><br>Note: This feature is not supported for Internet Explorer 4.0 and earlier. |
| AutoConfigScriptEnable | 0 = Automatic proxy script will not be executed. 1 = Automatic proxy script will be executed. |
| AutoConfigScript | The URL that points to the automatic proxy script. |
| ProxyEnable | 0 = Manual proxy detection is disabled. 1 = Manual proxy detection is enabled. |
| ProxyServer | The name and the port of the proxy server. |
| ProxyOverride | The Domain Name System (DNS) addresses, separated by semi-colons, that should not be passed on to the proxy server. You can use the value <local> to bypass the proxy server for local addresses. If you include this value, it should be the final value in the string. |

The following is an example of a simple proxy text file for Awesome Computers, simpleproxy.txt:

```
[Automatic Proxy]
AutoProxyEnable=1
```

```
AutoConfigScriptEnable=1
AutoConfigScript=www.awesome.com/proxysettings.txt

[Manual Proxy]
ProxyEnable=0
```

You can configure far more complex proxy settings. In the following example, complexproxy.txt, automatic proxy detection is used where possible, and a manual proxy supports Internet Explorer 4.0 users, who cannot use automatic proxy detection. Different protocols use different proxy servers, and several addresses have been excluded from the proxy.

```
[Automatic Proxy]
AutoProxyEnable=1

[Manual Proxy]
ProxyEnable=1
ProxyServer=ftp=FTPProxyServer:2010;gopher=GopherProxyServer:1066;http=HTTPProxyServer:
42;https=SecureProxyServer:1984;socks=SocksProxyServer:5150
ProxyOverride=DNSaddress1;DNSaddress2;<local>
```

For more information about proxy settings in Internet Explorer, refer to the Internet Explorer Administration Kit.

For more information about Connection Manager keys and values, see Advanced Customization.

**Note**

- Users who select the **Log on using dial-up connection check box** gain access through an account designed for this method of logging on. The Log on to Windows account has only minimal permissions, and any customized settings for Connection Manager, including proxy configuration, are applied to the Log on to Windows account, rather than to the account of the user who logged on. For more information, see Incorporating Connection Manager with logon security.

# Incorporating custom actions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incorporating custom actions**
You can enhance the connection experience for your users by providing additional programs that start seamlessly during the connection to your service. You can use the Connection Manager Administration Kit (CMAK) wizard to include custom actions in your service profiles to automatically start programs when users connect to your service. These custom actions can use programs that users have already installed, or you can include the programs with your service profile.

**What custom actions are supported and when do they run?**
Using the CMAK wizard, you can specify custom actions to run at nine points during the connection process. You can run:

- Pre-init actions
- Pre-connect actions
- Pre-dial actions (for dial-up connections only)
- Pre-tunnel actions (for virtual private network (VPN) connections only)
- Post-connect actions
- Monitored actions
- Disconnect actions
- On cancel actions
- On error actions

**Pre-init actions**
As soon as users start Connection Manager, it runs the pre-init actions you have specified in your service profile. These actions are run before the Connection Manager logon screen appears.

### Note

- Pre-init actions will run when the **Properties** dialog box is opened for the service profile.

**Pre-connect actions**
As soon as users click **Connect**, Connection Manager runs the pre-connect actions you have specified in your service profile. These are run before Connection Manager establishes a connection to your service. For actions that relate specifically to dialing, use pre-dial actions. For actions that relate specifically to tunneling, use pre-tunnel actions, not pre-dial actions.

You cannot set up a pre-connect action that requires an interaction with your service. If interaction is required for part of the action, split the action into a pre-connect action and either a pre-tunnel or post-connect action, depending on the interaction requirements.

For example, recording information to a local log at the start of the logon process can be a pre-connect action, but if you want the log to post to your service, you might set up a separate upload program to be run as a post-connect, or disconnect action.

**Pre-dial actions**
Connection Manager runs pre-dial actions after users click **Connect** but before the computer starts to dial the connection to your service. Pre-connect actions run before pre-dial actions.

For example, pre-dial actions might include a program to change the telephone number for the service profile, or they might include a program that affects how the profile dials the connection.

**Pre-tunnel actions (for VPN connections only)**
Connection Manager runs pre-tunnel actions after establishing a connection with the Internet server but before establishing a tunnel to the VPN server. This type of action is available only if you set up VPN in your service profile, and it will run only when users are using the VPN connection option.

For VPN connections, set up all custom actions that require interaction with an Internet server as pre-tunnel actions, because Connection Manager has limited access to an Internet server after a tunnel is established.

For example, this might be a program that verifies user credentials for accessing a corporate server behind a firewall or that checks the expiration dates of client certificates.

For direct VPN connections, Connection Manager runs both pre-connect and pre-tunnel actions before connecting.

**Post-connect actions**
Connection Manager runs post-connect actions after establishing a connection or, if using a VPN connection, after establishing the tunnel. Each post-connect action that you specify in the CMAK wizard runs every time the user connects to your service, whether the connection is a dial-up connection or a direct connection.

In the CMAK wizard **Custom Actions** pane, you might have the following built-in post-connect action:

- **Download phone-book updates**. If you set up the service profile to automatically download phone book updates, Connection Manager includes this post-connect action which sends a phone book request to the server running Phone Book Service (PBS). The server compares the client phone-book version with the most recent files available at the server and then downloads any available phone-book updates. The next time the user starts Connection Manager, the updated access numbers are available. You must provide the name of the phone book and the URL of the PBS server in order to set up this custom action. For more information on phone book updates, see Connection Point Services and Connection Manager.

**Note**

- If you have Connection Manager 1.0 service profiles that include connect actions, you should know that custom actions in Connection Manager 1.3 are the same as connect actions in Connection Manager 1.0. In fact, in Connection Manager 1.3, post-connect actions are specified in the [Connect Actions] section of the .cms file, as there is no section with the title of [Post-Connect actions]. This terminology has been maintained to support migration from the previous version.

**Monitored actions**
Connection Manager runs monitored actions after establishing a connection and, for a VPN connection, after establishing the tunnel. Each monitored action that you specify in the CMAK wizard runs every time the user connects to your service, whether the connection is a dial-up connection or a direct connection. Monitored actions cause the connection to end after specified programs have ended. For example, if you want to automatically run an e-mail program, and then automatically end the connection when the user quits the e-mail program, you must set up e-mail as a monitored action.

Monitored actions are similar to post-connect actions, except that:

- All monitored actions must be .exe files, because monitored actions are run asynchronously. A dynamic-link library (DLL) runs synchronously, so it cannot be run as a monitored action.

- Connection Manager automatically monitors the status of all monitored actions and starts the disconnect sequence when the last monitored action closes.

**Note**

- When the last monitored action closes, Connection Manager provides a prompt to allow the user to override the disconnect sequence. The user can click **Stay Online** to prevent disconnection, click **Disconnect Now** to disconnect immediately, or do nothing to let the connection close automatically after 30 seconds have elapsed.

As with post-connect actions, you must provide the appropriate information for each monitored action specified in the service profile, including the application to run, whether to include the application in the service profile, and any command-line parameters that you want to run with the application. Unlike post-connect actions, you do not specify a descriptive name for a monitored action.

**Important**

- Specify only applications that are complete program files. You can include only one file for each monitored action. If you specify an application that requires other files or is a *self-extracting* executable (.exe) file as a monitored action, these files might cause the program or connection to be interrupted. If a monitored action requires any files in addition to the executable file, you should specify them as Additional Files in the CMAK wizard. These files will be installed with the service profile but not run. For more information on how to include these support files in the service profile, see Including additional files.

You can specify a variety of applications as monitored actions. For example, you might want to include a commercially available e-mail application, a customized stock-ticker application, and an internal accounting application. The CMAK wizard can specify all three types of applications. You can also select whether the application for each monitored action is to be included in your service profile.

**Notes**

- If you have Connection Manager 1.0 service profiles that include connect actions, you should know that Monitored Actions in Connection Manager 1.3 are the same as Auto Applications in Connection Manager 1.0. In fact, in Connection Manager 1.3, you specify monitored actions in the [Auto-Applications] section of the .cms file, because there is no section with the title of [Monitored Applications] in that file. This terminology has been maintained to support migration from the previous version.

- If the application is one that users already have installed on their computers, do not include it in the service profile. Doing so will increase the size of your service profile and might cause problems on the users' computers.

- You should not specify Internet Explorer 4.0 or 4.01 as a monitored action if your users have installed an active desktop. Upgrading your users to the most recent version of Internet Explorer is advisable if you want to specify Internet Explorer as a monitored action.

**Disconnect actions**

Connection Manager runs disconnect actions immediately before disconnecting from your service. You can use a disconnect action for routine administration. For example, you might set up a custom action to collect status information from your service, such as total minutes online (if this information is tracked by your service). This information can then be displayed for the user.

**Important**

- Disconnect actions will run even if the disconnect was not initialized by Connection Manager. For example, if the user's connection is terminated by a disruption in telephone service, Connection Manager will attempt to run the disconnect actions specified in the service profile after the unexpected termination.

**On cancel actions**

As soon as users click **Cancel** during a connection attempt, Connection Manager runs the on cancel actions you have specified in the service profile. On cancel actions do not run if the user clicks the **Cancel** button to close the Connection Manager program. You might use on cancel actions to reverse any pre-init actions you specified in the service profile.

**On error actions**

Connection Manager runs on error actions whenever an error occurs during the connection. You might use on error actions to collect information about the error for later support.

**What types of programs can be used as custom actions?**

A custom action can be:

- A dynamic-link library (DLL)

- An executable file such as a .bat, .exe, or .cmd file

- A shell-executable file, such as a .txt or .doc file

**Including Dynamic Link Libraries (DLLs)**
Only DLLs run synchronously, meaning that Connection Manager starts the action and then waits for the function to return before continuing. The first argument, specified as the **Parameters** in the **Add/Edit Custom Actions**pane of the CMAK wizard, is the function name within the DLL to call. The syntax for each of these custom actions is:

```
HRESULT WINAPI function
    [IN] HWND hWndParent
    [IN] HINSTANCE hinstDll
    [IN] LPCSTR pszCommandLine
    [IN] DWORD dwReserved
```

Supported DLL parameters are:

| Parameter | Result |
|---|---|
| hWndParent | Handle to the Connection Manager logon dialog box or NULL; used as the parent window for any user interface displayed by the custom action. |
| hinstDll | Handle to the instance of this DLL. |
| pszCommandLine | String pointer to the command-line arguments. |
| dwReserved | Reserved for future use. |

DLLs often have a comment (entered in the description field of the CMAK wizard) that appears while it is being run. This comment has the form **Running Custom Action** *Description*. The Connection Manager window is frozen and does not accept input while this type of action runs.

Custom action DLLs require a return value. If the return value is less than 0, the DLL call failed (SUCCEEDED macro returns False). In the case of a failure, an error message (including the return value) is displayed in the form "**Custom Action** ***Description*** failed *return value*." If a pre-init, pre-dial, pre-connect, pre-tunnel, or post-connect DLL action fails, the connection attempt is ended.

The DLL is unloaded after the function call.

### Notes

- For .dll files that require extended times for the program to run, you should implement WM_PAINT messages in the .dll file to ensure that graphics are not disrupted by user actions.

- You must provide the exact name of exported functions when building custom action DLLs. If you are using a language that changes the exported name (such as C++), you can use a .def file to preserve the exported function names.

**Including executable files**
This type of custom action is run asynchronously, which means that Connection Manager starts the executable, including any command-line parameters, and continues without waiting for the process to exit. No comment appears in the Connection Manager status pane while these programs are run.

**Including shell-executable files**
These actions run asynchronously. Each shell-executable file needs to be opened by a particular program.

If this program is not included in Microsoft operating systems, you might need to include it with your service profile.

**How to specify a custom action**
In the CMAK wizard **Custom Actions** pane, click **New** to add a custom action or **Edit** to edit an existing custom action. A dialog box will appear, where you must provide the following information for each custom action:

- **Description:** Provide a short description of the program to run. If you do not provide a description, the CMAK wizard uses the program name (specified in the **Program to run**) as the description.

- **Program to run:** Specify the file name of the program or file, including the file extension (such as .dll, .exe, .bat, or .cmd).

- **Parameters:** Specify any command-line parameters that you want to run with the program. The following special command-line parameter macros are supported:

| Macro | Replaced with |
|---|---|
| %ServiceName% | The service name of the profile. |
| %UserPrefix% | The user-name prefix used for this connection. |
| %UserSuffix% | The user-name suffix used for this connection. |
| %UserName% | The user name, without any specified realm user-name prefix or suffix. |
| %Profile% | The location and file name of the active connection profile (.cmp) file. |
| %ServiceDir% | The path to the profile directory. |
| %Domain% | The domain for the connection. |
| %InetUserName% | The user name for the Internet connection. |
| %ConnectionType% | A value that identifies the connection type, 0, 1 or 2.<br>0 = dial-up<br>1 = direct<br>2 = double-dial (also known as VPN) |
| %DialRasPhoneBook% | The full path to the phone book file. |
| %TunnelRasPhoneBook% | The full path to the phone book file used for the VPN portion of this connection. |
| %DialRasEntry% | The service name or remote access entry name for the dial-up connection. |
| %TunnelRasEntry% | The service name or remote access entry name for the tunnel connection. |
| %AutoRedial% | A Boolean value that is 1 if this dial attempt is an automatic redial, or 0 if not an automatic redial. |
| %PopName% | The description for the phone number in the CM phone book. |
| %ErrorCode% | The Win32 error code. |
| %LastErrorSource% | The origin of the last error. |
| %CurrentFavorite% | The location settings selected in the **Use settings for** box. |
| %TunnelServerAddress% | The IP or DNS address of the VPN server, if any. |
| %ClientIPAddress% | The IP address of the computer on which the Connection Manager profile is installed. |
| %ServerIPAddress% | The IP address of the remote access server, if any. |
| %Interactive% | A Boolean value, used in a custom action to determine whether or not to display a user interface. You can incorporate this macro in a custom action to display an interactive user interface (such as an error message) only when Connection Manager is running in an interactive state, and to take other action (such as failing gracefully without an error message) when Connection Manager is running in a non-interactive state. This macro was designed to be used in conjunction with the **Program interacts with the user** check box. For example, if you deselect **Program interacts with the user**, you should add the **%interactive%** macro to the custom action parameters so that the custom action can behave according to the state in which Connection Manager is run. This might require modification to the custom action itself. |

When parameters are specified, Connection Manager reads the parameter string and passes it to the **pszCommandLine** parameter of the .dll or sends it as part of the command line to an executable (available through the Windows API **GetCommandLine**).

If the custom action is a DLL, then the first token of the parameter string is the function entry point to call. This token is removed from the parameters before the string is passed to the DLL (so the **pszCommandLine** string does not start with the .dll entry point name as the first parameter). When creating the DLL, ensure that the exported name of the DLL matches the name specified as the first parameter, as name decoration can result in different names. (For example, in C++, name decoration can cause the exported name to be different from the actual function name.)

If you use the macros shown in the previous table, Connection Manager replaces them with the actual run-time information for the parameter. For instance, if you wanted a custom action to collect information about which users are using which POPs, you might create a .dll (or other executable) for this purpose, and Connection Manager could pass it information such as user name (using the **%UserName%**macro), realm information (using the **%Prefix%** or **%Suffix%** macro), and POP name (using the **%PopName%** macro). To do this, in **Parameters**, type the following: **LogUserAndPopToServer %UserPrefix% %UserName% %UserSuffix% %PopName% http://logserver.awesome.microsoft.com %AutoRedial%**

Depending on the user connection method, the information passed to the DLL might be:

```
NULL jobrown @awesome.microsoft.com "Seattle 28,800"
http://logserver.awesome.microsoft.com 0
```

Connection Manager cannot denote a null pointer in a string, so it passes the string NULL (which it passes for all undefined parameters). Only use quotes when a parameter contains spaces, such as "Seattle 28,800". The number variables have to be parsed from the string. If you are using an executable other than a DLL, the function name (in the above example, **LogUserAndPopToServer**) is not required.

- **Action Type:** Specify the time at which the action is performed by selecting **Pre-init**, **Pre-connect**, **Pre-dial**, **Pre-tunnel**, **Post-connect**, **Monitored**, **Disconnect**, **On cancel**, or **On error**.

- **Run this custom action for:** Specify what kinds of connections for which this custom action should be run. You can choose to run this action for **All connections**, **All connections that involve dial-up**, **All connections that involve a tunnel**, **Connections that use only a tunnel**, or **Connections that use only dial-up**.

- **Include the custom action program with this service profile**. Select this check box to include the program with the service profile. If the program you want to run is not part of a standard Windows operating system or is not otherwise available to your users, you should include the program in your service profile. The custom action program is installed in the service profile folder.

- **Program interacts with the user**. This check box is selected by default. If this check box is selected, Connection Manager will not run the custom action if it is in a non-interactive state. If this check box is not selected, a program called by the custom action that interacts with the user will halt the connection process indefinitely, waiting for a response from the user. For more deterministic behavior, you can deselect this check box and use the *%interactive%* macro.

Specify only programs that are complete program files. You can include one program file per custom action. If any non-executable files are required to support a custom action, you should include them as additional files, not as custom actions. These additional files should be files that are installed with the service profile but do not have to be run. For more information on how to include these support files in the service profile, see Including additional files.

**Important**

- If a custom action runs when users log on to Windows, the custom action runs with system permissions. You should ensure that any custom actions that you enable do not pose unintended security risks.

**Notes**

- Custom actions are run in the order shown in the CMAK wizard.

- Not all system macros are supported on all Windows operating systems.

- By default, custom actions are not run when a profile is used to log on to Windows. For more information about whether to enable custom actions for use when logging on to Windows, see Incorporating Connection Manager with logon security.

- If you want settings not found in the Connection Manager macros, other settings can be accessed programmatically

through remote access Application Programming Interfaces (APIs), which are documented in the Microsoft Web site.

- Sample code for a custom action that dynamically updates the proxy server setting for direct Internet connection users based on their selection of a VPN server is available in the Remote Access Service sample section of the Microsoft Platform Software Development Kit. For more information, see the Microsoft Web site.

# Incorporating custom graphics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incorporating custom graphics**
You must specify in the CMAK wizard whether to use the Connection Manager default bitmaps or your own graphics. You can customize the user interface to include your company logo or another image that identifies your organization. You can replace the graphics in both the logon dialog box and the **Phone Book** dialog box. The replacement graphic must be a bitmap (.bmp) file.

**Logon graphic**
You can create a custom bitmap or use an existing bitmap to replace the default Connection Manager logon bitmap that appears at the top of the logon dialog box. The default bitmap is 330-by-140 pixels.

**Phone book graphic**
You can customize the bitmap in the **Phone Book** dialog box that users see when they select a phone number for dial-up connections. This bitmap appears on the left side of the **Phone Book** dialog box. The default bitmap is 114-by-309 pixels.

**Creating bitmaps**
You should consider the following when creating bitmaps for either of the two replaceable graphics:

- Each bitmap looks best if it has the same proportions as the default bitmap.

  Connection Manager automatically scales each bitmap to appear at the appropriate size, but the image may be distorted if the proportions are not correct.

- Connection Manager uses the same set of bitmaps, regardless of the display mode.

- You should use the Internet Explorer 256-color (halftone) palette when creating bitmaps.

Be sure to test each bitmap on all target operating systems and platforms.

You can also customize the three icons your users see. For more information, see Incorporating custom icons.

# Incorporating custom icons

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incorporating custom icons**

You can provide custom icons for use with your service profile. You can provide one file that contains icons in multiple sizes or you can provide one file for each icon in each size. You use the CMAK wizard to specify the icons your users see.

- The large program icon is used as the desktop icon. It also appears in the ALT+TAB switch window, the Network Connections folder (Windows XP and the Windows Server 2003 family only), and in several other pop-up dialog boxes. This icon is 32 × 32 pixels.

- The small title-bar icon appears at the left edge of the title bar in the logon dialog box, in various pop-up dialog boxes and, when the program is running, on its program button on the taskbar. This icon is 16 × 16 pixels.

- After a connection is established, the small notification area shortcut appears in the notification area at the far right of the taskbar. This shortcut icon is 16 × 16 pixels.

Connection Manager contains default images for these three icons. If you prefer to use your own icons in place of the default icons, you can use an icon-editing program to create them. Consider creating icons that are similar to one another to unify the appearance of the service profile. To ensure that your icons appear correctly, consider creating icons that match the size and overall shape of the default icons.

A single icon file (.ico) can have multiple images, each with a different size, color depth, or both. If you create an icon that you want to use in place of all three default icons, you should build the icon to include both 16 × 16 pixel and 32 × 32 pixel icons in the same file. If you include an icon in only one size, the icon is scaled to fit the assigned area when it appears, but the image might be dithered or distorted. To specify a single custom icon to be used for all service-profile icons, you must specify the name of the custom .ico file in all three boxes in the **Icons** pane of the CMAK wizard. If a box is left blank, the appropriate default is used.

**Note**

- You can also customize the menu that appears when a user right-clicks the notification area shortcut. For more information on customizing this shortcut menu, see Customizing the notification area shortcut menu.

# Customizing the notification area shortcut menu

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Customizing the notification area shortcut menu**

The notification area shortcut appears in the notification area of the taskbar when the user is connected to your service. The notification area shortcut menu appears when a user right-clicks this notification area shortcut. This menu provides users with a shortcut to specific programs. On the Windows XP and Windows Server 2003 family platforms, the default commands on the menu are **Status**, **Disconnect**, and **Open Network Connections**. You can use the CMAK wizard to customize the shortcut menu by specifying which, if any, additional commands appear on the menu.

For each command that you want to add to the menu (other than the default commands), in the **Notification Area Shortcut Menu** pane, click **Add** and then specify:

- The name of the command as you want it to appear on the shortcut menu. If you do not specify a name for the command, the menu displays the file name and extension of the program to be run.
- The file name (including the extension) of the program that runs when the user selects the associated command from the menu.
- Any command-line parameters that you want to run with the program.
- Whether to include the program file in the service profile. If the program is not part of the standard operating systems of the target platforms or is otherwise not available on your users' computers, you can include it in the service profile. At installation, the service profile copies the program to the \Cm*ServiceProfileFileName* folder.

You can add keyboard access to a shortcut menu command by designating one of the letters in the command name as the access key. To specify an access key, type **&** immediately before the letter you want to use. For example, to add "Notepad" as a command name with the letter *N* specified as the access key, type **&Notepad**. When the user opens the shortcut menu, the access-key letter (*N*, in this example) is underlined in the command name (N**otepad**). The user would press **N** to carry out the **Notepad** command.

If you want to include the ampersand (**&**) character in the command name, type **&&**. For example, to use the command name **Update & Refresh**, type **Update && Refresh**. One ampersand appears in the shortcut menu, not as an access key. (**Update & Refresh**).

**Notes**

- You can control whether a notification area shortcut appears in the notification area of the taskbar by editing the **NoTrayIcon** key in the **[Connection Manager]** section of the .cms file.
- The default commands are not shown in the CMAK wizard in **Shortcut menu commands** or included in the service-profile files. They are part of the Connection Manager software and cannot be deleted.
- Only .exe files and other shell executables (such as URLs and other files with registered extensions) are supported as menu commands. DLLs are not supported.
- System macros are not supported for notification area shortcut commands.

For more information on using advanced customization techniques to edit the .cms file, see Advanced Customization.

# Providing custom Windows Help

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providing custom Windows Help**

Using the CMAK wizard, you can customize the Connection Manager Help, which consists of definition Help and troubleshooting Help:

- Definition Help provides a brief description of each interface item. This type of Help topic appears in a small pop-up window. Users display topics by right-clicking an interface item and then clicking **What's This?** in the pop-up menu that appears. In all dialog boxes except the logon dialog box, users can also display definition Help by clicking the question-mark icon in the title bar and then clicking an interface item.

- Troubleshooting Help provides a structured approach to diagnosing and resolving many common connection problems. Users start the **Troubleshooting Help for Connection Manager** by clicking **Help** in the logon dialog box. The question-and-answer format guides users through the process of determining the most likely causes of the problem. The troubleshooter then provides possible solutions, based on the user's feedback.

Both the definition and troubleshooting Help files are compiled into a single .hlp file. This compiled file is incorporated into the service profile for distribution to users. The default Help is Cmmgr32.hlp. If you do not want to use the default Help, you can use the CMAK wizard to specify custom Help to replace the default in your service profile.

### Note

- There are two areas of Help information that you cannot customize in the definition Help. The first non-customizable Help is for the **Dialing Properties** dialog box, which users see when they click the **Dialing Rules** button on the **General** tab of the **Properties** dialog box. This dialog box, along with its associated Help, is a component of the Telephony Application Programming Interface (TAPI) and, although it is incorporated as a part of Connection Manager, it cannot be modified. The second area of non-customizable Help is for the **Advanced** tab. The contents of this tab is a component of the operating system, and, although it is incorporated as part of Connection Manager, it cannot be modified.

**How to customize the default Help files**

If you want to include custom Windows Help, you must prepare your files before you run the CMAK wizard. Your CMAK folder contains all the source files for the default Help. To create custom Help, you need to copy the default Help source files, including the project [.hpj] file, rename them, edit the renamed files, and recompile them. Putting them in a separate folder helps prevent the original files from being accidentally altered. Following is information for creating custom Help, including a description of all of the source files and the tools required to modify these source files. (A basic understanding of Windows Help systems and user-education concepts is assumed.)

### Notes

- Be careful when modifying the Help files. Troubleshooting any changes you make to these files may be difficult. Avoid changing Windows Help file footnotes. Do not change any of the topic IDs (the number sign [**#**] footnotes).

- When specifying your custom Help file in the CMAK wizard, browse to it. If you type the name, no validation is done on the file name (including the file extension).

- Be sure to test the Help files on all target platforms before delivering your service profile to your users. Test Help by starting Connection Manager and using both the definition Help and troubleshooting Help. Testing Help by opening the .hlp file outside of Connection Manager does not allow verification of functionality.

**Source files**

The following files are included with the Connection Manager Administration Kit. By copying and editing these source files, compiling your customized project (.hpj) file, and then specifying the name of your compiled custom Help file in the CMAK wizard, you can provide custom Help to your users.

- Cmmgr32.hlp (the default compiled Help file) is installed on your hard disk in the %windir%\system32 folder.
- The following online Help components are installed on your hard disk in the \Program Files\CMAK\Support\CMHelp folder.
    - Cmmgr32.hpj (the Help project file)
    - Cmctx32.rtf (the text document source for definition Help)
    - Cmtrb32.rtf (the text document source for troubleshooting Help)
    - Cmmgr32.h (the header file containing definition Help IDs)

**Note**

- Graphics (such as bullets) and the template file (What7.dot) are not included, so you should create these files as appropriate to the look you want to provide in your Help.

You must have all of these files if you want to customize and rebuild the online Help system. Suggestions for customizing entries in each file are discussed later in this section.

**Important**

- To customize the Help files for a service profile, copy the default Help files (except for the cmmgr32.hlp file), rename them, and then edit only the renamed files. Do not edit the original default files. When you rename the files, be sure to update the names referenced in the .hpj file, including the name of the .hlp file. If you customize Help but do not rename the .hlp file, the CMAK wizard could accidentally incorporate both the default and custom files when building the service profile.

**Note**

- In addition to Help files, you can provide custom readme and procedural documentation to your users. For more information on the additional documentation that you can customize for your users, see Providing user documentation.

**Required tools**

To customize Help files, you need both of the following:

- A word-processing program, such as Microsoft Word, that reads and writes .rtf files.
- A 32-bit Help compiler, such as Help Workshop version 4.01.

**Note**

- You can obtain the Help compiler (Hcw.exe) and associated files by downloading both the Microsoft Platform SDK build environment and the additional SDK components (tools) from the Microsoft SDKs at the Microsoft Developer Network (MSDN) on the Microsoft Web site. However, only limited documentation on how to create and edit Help is available in the SDK.

**What should be customized?**

The most commonly customized Help information is that contained in the .rtf files, including the text for both definition and troubleshooting Help.

**Customization suggestions for definition Help**

The text for definition Help is in the Cmctx32.rtf file. In the following list of customization suggestions, each topic is referred to by its topic ID (the **#** footnote).

| Help topic ID | Customization information |
|---|---|
| IDH_LOGON_SVCMSG | You can use this topic to provide definition Help for the support text in the logon dialog box. This topic is available to your users only if you specify support information when you run the CMAK wizard. The default text for this message is "No Help topic is associated with this item." You can modify this text to provide assistance for the support information specified for your users. For example, if you provide a phone number for customer support in the logon dialog box, you might want to provide a description of your customer support service in this topic. |
| IDH_PHONEBOOK_SERVICETYP | In the **Phone Book** dialog box, this topic provides Help for the **Service Type** list. You can customize this topic to reflect the service types that you support and provide a brief description of each. |
| IDH_PHONEBOOK_ACCESSNUMS and IDH_PHONEBOOK_MORE | In the **Phone Book** dialog box, these topics provide Help for the **Access Numbers** and **More Access Numbers** lists. You can edit the text to change it or add information about the available numbers, such as which numbers incur user fees and which do not. The default text specifies that additional charges are incurred for using these numbers. |
| IDH_PHONEBOOK_SVCMSG | In the **Phone Book** dialog box, this topic provides Help for the phone book customer service message area (next to the **More access numbers** box). This message appears only if alternate access numbers are available for the selected geographic area. You can use this topic to provide additional customer service information about the numbers that you make available in the **More access numbers** box. |

Throughout the Cmctx32.rtf file, you might want to convert the term "your service" to the name of your service. For example, IDH_LOGON_NAME contains: "The name you use to log on to your service." Changing "your service" to "Awesome Computers" (the name of the fictitious company used in this documentation) would change this information to: "The name you use to log on to Awesome Computers."

#### Customization suggestions for troubleshooting Help

Text for troubleshooting Help is in the Cmtrb32.rtf file. The following list of customization suggestions includes reference topics by title and by topic ID (the **#** footnote).

| Troubleshooter topic ID | Customization information |
|---|---|
| ICM_TRB (Logon troubleshooter) | This topic is the starting point for users to identify possible problems. If you have specific problems that you want to address, you can add your own paths to the troubleshooter, starting at this point or any other point. You can add topics to cover as many situations as you want. |
| ICM_TRB_CUSTSUP_NUMS (Contact customer support) | This topic provides users with customer support information. Add your specific customer support phone numbers, hours of operation, and any other useful information. Users display this topic from other topics that recommend contacting customer support. |
| ICM_TRB_GIVEUP and ICM_TRB_GIVEUP_PHONE (Next steps) | One of these two topics appears at the end of all troubleshooting paths. You can provide additional instructions for users who try all troubleshooting possibilities and still are not able to connect to your service. |

As with the definition Help, you might want to change the term "your service" to the name of your service, wherever it appears in this file.

#### Customizing and compiling the project file

If you customize the .rtf files, you must customize the project (.hpj) file (using Help Workshop 4.01 or another appropriate Help editor) to include the names of these custom .rtf files. You can also change other project information in the .hpj file. For example, if you do not want to use "Troubleshooting Help for Connection Manager" in the title bar of the Help screens, you can change the title in the .hpj file by changing "TITLE" under **[Options]** and "proc4" under **[Windows]**.

After you have finished editing the .hpj and .rtf files, you must compile the .hpj file to create the .hlp file. When you run the CMAK wizard to create your service profile, specify the name of your custom .hlp file in the wizard.

### Note

- The .h (mapping) file contains the topic IDs for the definition Help. Editing this file may prevent Connection Manager from displaying the appropriate definition Help.

# Providing support information in the logon dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providing support information in the logon dialog box**

Using the CMAK wizard, you can specify a line of information to be displayed in the logon dialog box, directly above **Connection status**. This line of text can be a maximum of approximately 50 characters and usually gives users a number to call for support (for example, "Call 1-800-###-#### for customer service."). No default text exists for this option. If you do not specify a line of text, this area of the logon dialog box will be left blank.

It is a good idea to repeat this contact information in your readme files or other installation information in case a user is unable to display the logon dialog box.

If the contact information is dependent on the location from which the user is dialing, you may want to customize Help to list all customer support numbers (to support remote users roaming outside of their home service area). In the line of text in the logon dialog box, you can then include a reference to Help. (For example, "For customer support phone numbers, click Help.") If you do this, add a link to the initial Help screen that jumps to the topic ID containing your support numbers. For more information on customizing Help, see Providing custom Windows Help.

# Adding Connection Manager 1.3 to the service profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding Connection Manager 1.3 to the service profile**

Users must install Connection Manager 1.3 to run a Connection Manager 1.3 service profile. If your users do not already have it installed, you should include the software in your service profile. In the CMAK wizard, you can specify that this software be included in the service profile being built, so the software is installed when the user installs the service profile.

If you include Connection Manager 1.3 in the service profile and a user already has a previous version of Connection Manager installed, the software is automatically upgraded at installation. Connection Manager cannot be installed if the current or a newer version is already installed on the user's computer. For example, Connection Manager 1.3 is a component of the Windows XP operating system, so it does not get installed on computers running Windows XP. However, the service profile installs, even if the Connection Manager software included in the profile does not install.

### Notes

- You must include the Connection Manager 1.3 software for profiles intended for installation on computers running Windows 2000.

- Be sure to include the size of the Connection Manager software when calculating the size of your installation package. For more information on the size of the Connection Manager software, see Planning for effective implementation.

# Incorporating a license agreement

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incorporating a license agreement**

Using the CMAK wizard, you can include your own license agreement with your service profile. This license agreement appears when the user clicks **OK** to begin the initial installation of your service profile. If the user accepts the license agreement, installation continues. If not, installation is cancelled.

Use a text-editor program such as Notepad to create your license agreement as a text (.txt) file. To avoid formatting problems, do not use hard returns (forced line breaks) at the end of lines. Line-wrapping for the license agreement is done automatically when the service profile is built.

# Including additional files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Including additional files**

Using the CMAK wizard, you can include additional files with your service profile to help your users implement your service. All files that are needed for your service that are not specified elsewhere in the CMAK wizard should be specified in the **Additional Files** pane. This section contains information on files that you might want to include as **Additional Files.**

**Custom action and monitored application files**

If you create a custom action, you include the program for the custom action by specifying it in the **Add/Edit Custom Actions** pane. However, you can specify only one program there for each custom action. You must specify any other files required to implement the custom action as **Additional Files**. The same is true for monitored applications. For more information on the files that should be included with custom actions and monitored applications, see Incorporating custom actions.

**User documentation files**

If you create user documentation for your service, you can include this information in your service profile by specifying the user documentation files as **Additional Files**.

**Other files**

If, after running the CMAK wizard, you do any advanced customization by editing any of the service-profile files, you may need to include additional files to support some of your advanced customization. You specify any files needed to implement advanced customization as **Additional Files**. For more information on this type of advanced customization, see Advanced Customization.

# Providing user documentation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Providing user documentation**

In addition to Providing custom Windows Help with your documentation, you can incorporate other types of user documentation in your service profile. You can also develop and incorporate your own user documentation. Start by reviewing the information provided:

- Sample procedures, included later in this topic:
    - "Setting up Connection Manager," including instructions for installing Point-to-Point Tunneling Protocol (PPTP) on Windows NT 4.0
    - "Establishing a connection"
    - "Removing Connection Manager"

To reduce your customer-support effort, make sure your setup and usage instructions clearly and completely guide users through both the setup of Connection Manager and connection to your service.

To include user documentation in your service profile, you can specify the files as additional files (in the **Additional Files** pane of the CMAK wizard). For more information, see Including additional files.

### Note

- Information included as additional documentation is not available to users until after the installation of your service profile. You may want to include some of the information (such as installation procedures) in a readme file to accompany your service profile, in a form that can be accessed by users before installation.

**User procedures**

You should provide user documentation to help your users install, run, and remove (if necessary) Connection Manager. You can customize the following procedures and put them in a readme file or other user documentation.

Throughout these user procedures, you should substitute the name of your service profile for each occurrence of *ServiceProfileFileName*, regardless of the context in which it is used. (The name you substitute should not be in italics.)

The following procedures contain glossary entries and links to a Web site. Depending on the format and support files you provide with your documentation, these may not work as implemented here. Change these as appropriate for your deliverables.

When you run the CMAK wizard to build the service profile, specify the file names of all user documentation that you want incorporated in the service profile (to be available after the service profile is installed) as additional files.

**Setting up Connection Manager**

Instruct your users to perform the following procedure for setting up Connection Manager.

**To set up Connection Manager**

1. Use the appropriate method to find the *ServiceProfileFileName* file:
    - If you received the program for your service on a floppy disk or compact disc:

    a. Insert it in the appropriate drive.

    b. Use Windows Explorer to find the file.

- If you downloaded Connection Manager from the Internet or a private network, browse to your download folder.

2. Double-click *ServiceProfileFileName*.exe.

3. After installation has finished, double-click the icon on your desktop to start the service. If prompted to install additional components, click **OK**.

**Verifying and installing the latest Dial-up Networking software for Windows 95**
If your service profile supports virtual private network (VPN) connections, your users who are running Windows 95 on their computers must have the latest version of **Dial-up Networking** installed on their computers. Provide the following instructions to your VPN users to ensure installation of the correct level:

**To verify and install Dial-up Networking software for Windows 95**

- If you are running Windows 95 on your computer, you must obtain and install the current **Dial-up Networking** software by downloading it as a Windows Update from the Microsoft Web site.

**Installing PPTP on Windows NT 4.0**
Users must manually install PPTP on computers running Windows NT 4.0. Provide the following instructions for users who need to install PPTP.

**To install PPTP on a computer running Windows NT 4.0**

1. Make sure your Windows NT 4.0 distribution disks are available during installation to provide access to system files. It is recommended that the latest Windows NT service pack be installed. Windows NT Service Pack 3 provides support for multiple modems or dial-up scripting. Windows NT Service Pack 4 provides additional support for virtual private network (VPN) connections. Users can obtain the current **Dial-up Networking** software by downloading it as a Windows Update from the Microsoft Web site.

2. Click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Network**.

3. On the **Protocols** tab, click **Add**.

4. In **Network Protocol**, click **Point To Point Tunneling Protocol**, and click **OK**.

   If Setup needs to copy files, click **Continue** when prompted.

5. When prompted, click **OK** to set the number of virtual private networks to one (1).

6. In the **Remote Access Setup** dialog box, click **Add**, click **VPN1 - RASPPTM**, and then click **OK**.

7. Click **Configure**, and make sure the **Dial out only** check box is selected.

8. In the **Remote Access Setup** dialog box, click **Network**, make sure **TCP/IP** is selected, and then click **Continue**.

   You can now close the Network utility and restart your computer, and then reinstall any service packs required.

**Establishing a connection**
Have users follow these steps to connect to your service:

**To establish a connection on a computer running Windows XP or a member of the Windows Server 2003 family**

1. On the desktop, double-click the *ServiceProfileFileName* icon.

2. Type your *user name* and password.

   If you are using a VPN connection to connect to a private network, you might be prompted for a second user name and password. If so, type the user name and password that you use for the Internet service provider (ISP).

3. If you are using Connection Manager to connect for the first time, click **Properties**, click **Phone Book**, and then select a phone number from the list of access numbers. (Or you can type the phone number that you want to use to access the service.)

4. Click **OK** twice.

5. Click **Connect** to establish a connection.

   When you connect, the service might update your phone book if new or changed phone book information is available.

### Notes

- You can also start Connection Manager using the program icon in the **Network Connections** folder.

- Connection Manager supports user installation of multiple service profiles, running multiple instances of Connection Manager, and multiple user profiles for each service profile. If a user installs more than one service profile, each is installed in a separate folder. You can provide as many profiles to your users as necessary to support your service, but each must have a different service name and file name.

- For help with establishing a connection or with troubleshooting, use Help:

  - For troubleshooting Help, click **Help** in the Connection Manager logon dialog box.

  - For more information about an item in the dialog boxes, click the question mark [**?**] button in the title bar, and then click the item.

    Or right-click the item, and then click **What's This?**.

#### Removing your service

Have users follow these steps if they want to remove your service.

### Note

- This procedure removes only the service profile. The Connection Manager software remains on the user's computer, as users may require it for other profiles on the computer.

#### To remove your service

1. Right-click the desktop icon for your service.

2. Click **Delete**.

3. When prompted to continue with the deletion, click **OK**.

### Note

- Removing the software for your service does not remove the Connection Manager software required to run your service. You must remove the Connection Manager software separately.

# Advanced Customization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Advanced customization**
The CMAK wizard supports most of the customization features that administrators need to build a custom Connection Manager service profile. However, you can customize additional features by editing the service profile files, thereby changing the manner in which Connection Manager handles certain functions.

This section covers:

- Uses for advanced customization
- Methods of editing service profile files
- Templates and service profile files
- Information files
- Tables of Key Names and Values

# Uses for advanced customization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Uses for advanced customization**

By using advanced customization techniques, you can change the implementation of many Connection Manager functions. For example, you can:

- Create a custom button that appears on the Connection Manager user interface to start applications or run processes.
- Prepopulate user name, password, phone number, and dialing properties for use in your service profile.
- Show a particular set of phone numbers based on user selection of connection device.
- Remove the **General** tab from the **Properties** dialog boxes for VPN-only service profiles.
- Control whether the **Save password** check box appears on the Connection Manager user interface.
- Control whether the **Logon domain** field appears.
- Specify your multilink strategy.
- Specify logging behavior.
- Turn off balloon-tip help for Connection Manager.

**Note**

- Refer to Tables of Key Names and Values for complete lists of keys and values you can modify.

# Methods of editing service profile files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### Methods of editing service profile files

The CMAK wizard provides the ability to customize the most frequently used elements of a service profile. However, you might want to customize elements that you cannot configure in the wizard. By adding or editing keys and values to the individual files that make up a service profile, you can provide even higher levels of customization to meet the needs of the most advanced environments.

#### Service profile file types

Service profiles include files of the following types:

- **Service provider (.cms) files.** The .cms files specify the configuration of the phone book and most of the other functions of your service profiles. Most advanced customization for a service profile is done by editing the .cms file for a particular service profile.

- **Connection profile (.cmp) files.** The .cmp files contain information specified by the user. Because users can overwrite any changes you make to this file, any changes you make to this file will be available only the first time the service profile is used. Editing this file allows you to provide a first-time-only population of user information. For example, you can fill in the user name, the password, and the phone number fields for the first user to install and open your service profile.

- **Information (.inf) files.** The .inf files specify installation information for your service profiles. Although you can edit some setup and uninstall information in an .inf file, you should test the installation after you make any changes. You cannot edit information files from the **Advanced Customization** pane of the CMAK wizard. To edit .inf files, you must use a plain-text editor.

- **Connection profile (.sed) files.** The .sed files contain the instructions for building a self-extracting executable (.exe) file for your service profiles. You should never edit any .sed file.

**Important**

- You should not edit these files without a basic understanding of system configuration files. Troubleshooting changes to service-profile files is difficult, so you should be careful if you modify them. Read section and key descriptions carefully. It is advisable to make a backup of these files before you edit them.

#### Editing files

You can edit specific files used to build the service profile in one of two ways:

- Select the **Advanced customization** check box in the **Ready to Build the Service Profile** pane of the CMAK wizard. If you do this and then click Next, the **Advanced Customization** pane appears.

- Edit files after profile creation with any plain-text editor, such as Notepad.

You can use the **Advanced customization** pane to edit the .cms and .cmp files when you build a service profile. You can select or type in section names, keys, and values. You must click **Apply** after you specify each change. Some advantages of this method are that the changes you make are built directly into the service profile when you click **Next** and that you can make many changes quickly. However, if you mistype a section name or a key name and click **Apply**, you cannot edit or delete that section or key name directly. You can delete a key only by setting the value to blank, and you can delete a section only by setting one of the keys in it to blank.

**Important**

- If you set the value of a single key to blank in the **Advanced Customization** pane, you delete that key and all other keys in that section, as well as the section itself.

- Any user can modify the .cms and .cmp files of a service profile that is installed on a computer running Windows 98 or Windows Millennium Edition. Users who install profiles for individual use on computers that are running Windows Server 2003, Windows XP, or Windows 2000 can modify .cms and .cmp files of the profiles they install. Only members of the Administrators or System Operators groups can modify .cms and .cmp files of profiles installed for All Users on computers that are running Windows Server 2003, Windows XP, or Windows 2000. All members of the Administrators group on a local computer can modify the .cms and .cmp files of any service profile that is installed on that computer.

To customize a service profile after you have built it, you can edit any of the service profile files with any plain-text editor, such as Notepad. However, you must run the CMAK wizard again after editing the files in this way in order to incorporate your changes.

For more information, see Deciding what to customize, Uses for advanced customization, and ISP and Service Type keys.

# Templates and service profile files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Templates and service profile files**

You can edit service profile files to customize functions and features that are not controlled through the CMAK wizard. To do this advanced customization, you can edit two sets of service profile files:

- The template files that are used by the CMAK wizard to build all service profiles
- The files that are built when you run the CMAK wizard to create a specific service profile

Service profile files consist of the following file types:

- **Service provider (.cms) files.** The .cms files specify the configuration of the phone book and most of the other functions of your service profiles. Most advanced customization for a service profile is done by editing the .cms file for that service profile.
- **Connection profile (.cmp) files.** The .cmp files contain information specified by the user. Because users can overwrite this information, any changes you make to this file will be available only the first time the service profile is used. Editing this file allows you to provide a first-time-only population of user information. For example, you can fill in the user name, the password, and the phone number fields for the first user to install and open your service profile.
- **Information (.inf) files.** The .inf files specify the installation information for your service profiles. Although you can edit some setup and uninstall information in an .inf file, you should test the installation after you make the changes. Information files cannot be edited from the CMAK wizard **Advanced Customization** pane. To edit them, you must use a text editor.
- **Connection profile (.sed) files.** The .sed files contain the instructions for building a self-extracting executable (.exe) file for your service profiles. You should never edit any .sed file.

**Important**

- You should not edit these files without a basic understanding of system configuration files. Troubleshooting changes to service-profile files is difficult, so you should be careful if you modify them. Read section and key descriptions carefully. It is advisable to make a backup of these files before you edit them.
- Do not save template files in Unicode.

**Editing template files**

You can edit template files, specific service profile files, or both, depending on whether you want to apply your customization options to all service profiles or to a specific service profile only.

**Template files**

When you install CMAK, the following template files are installed in the \Program Files\CMAK\Support folder:

- Template.cms
- Template.cmp
- Template.inf
- Template.sed

These files contain the default settings that the CMAK wizard uses to build service profiles. If you want to incorporate the same custom functions in all service profiles that you build, edit the template files before building the service profiles. You must edit the template files with a plain-text editor, such as Notepad. You cannot edit template files with the CMAK wizard.

Generally, you should edit only the .cms template file.

**Note**

- Changes you make to the template files apply only to subsequent new service profiles you build using the CMAK wizard.

**Built files**

When you use the CMAK wizard to build a service profile, the wizard creates a self-extracting executable (.exe) file and several other files containing the settings that you specify. The service profile files are saved in the \Program Files\CMAK\Profiles*ServiceProfileFileName* folder. All service profile files have the name *ServiceProfileFileName*, which you specify in the CMAK wizard. In addition to the .exe file, the wizard creates the following files for each service profile:

- *ServiceProfileFileName*.cms
- *ServiceProfileFileName*.cmp
- *ServiceProfileFileName*.inf
- *ServiceProfileFileName*.sed

Changes to template files are reflected only in profiles built after you make the change. If you want profiles built before you make the change to match the template, you must update the profile. To incorporate template changes into a built file, you must run the CMAK wizard again.

For more information, see Deciding what to customize, Uses for advanced customization, and Finishing customization after running the CMAK wizard.

# Information files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Information files**
Only customize entries under the section headers in the information files as specified in this topic.

**Information (.inf) file entries**
By editing a profile's .inf file, you can specify commands and programs that run when the user installs or uninstalls your service profile. For example, you can specify:

- Setup programs that integrate any custom software that you include.
- Programs that automatically display a readme file.
- Configuration utilities or cleanup programs.

To add custom setup and uninstall programs to your service profile, you must first create the service profile by running the CMAK wizard and then you must specify any files needed as additional files. After you have built the service profile, specify the necessary commands or programs in the appropriate sections of the .inf file, and then run the CMAK wizard again to incorporate your changes.

The .inf file contains four section headers under which you can include commands and programs to be run at specific points. Use file names of eight characters or less. Do not use spaces in the file names.

| This .inf file section header | Contains commands or programs that run |
|---|---|
| RunPreSetupCommandsSection | Before Setup begins |
| RunPostSetupCommandsSection | After Setup finishes installing the service profile |
| RunPreUnInstCommandsSection | Before the uninstall program begins |
| RunPostUnInstCommandsSection | After the uninstall program runs |

Each entry beneath the header should include the name of the command or program to be run and any command line parameters. Each command specified in the .inf file is run only once during an installation or uninstallation process.

**Important**

- Incorrectly specified commands or programs in the .inf file can cause significant installation problems. Do not add entries to any of the sections in the .inf file except those previously specified. Thoroughly test all service profiles containing custom .inf entries for version, file, or other conflicts. Be sure that the file location is specified correctly.
- If you edit Connection Manager 1.0 service profiles in Connection Manager Administration Kit 1.3, the profile is upgraded to Connection Manager 1.3 and the new .inf file replaces the previous .inf file. After the upgrade, you must add your entries to the new .inf file if you want to continue using your custom setup and uninstall commands in the service profiles that you build.

For more information, see Phase Four: Preparing for Integration, Delivery, and Installation, Incorporating Connection Manager with logon security, and Finishing customization after running the CMAK wizard.

# Tables of Key Names and Values

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tables of key names and values**

The tables in this section discuss the valid keys for each section of the .cms and .cmp files, the function and use of each key, and where each key can be set (in the CMAK wizard, the .cms file, or the .cmp file).

- ISP and Service Type keys
- Connection Manager keys
- Custom actions keys
- Networking keys
- Other sections, keys, and values

# ISP and Service Type keys

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ISP and Service Type keys**

You can change how Connection Manager defines ISPs and service types. To do this, you need to edit the service (.cms) file. The table in this topic covers the key names and values that are contained in the [ISP] and [Service Type] sections of the .cms file. This topic also describes the formats and conventions you should use in the .cms file.

**Formats and conventions**

Connection Manager creates service (.cms) files, which use the same format as Windows initialization (.ini) files and information (.inf) files. These files have several things in common:

- Each file is divided into sections, with specific entries, called keys, assigned to the sections. A section heading is enclosed within brackets (**[ ]**). For example, **[Strings]** is a section containing text strings for messages. The keys for a section are not enclosed within brackets.
- Each value specified for a numeric key must be an integer. Any non-integer value is truncated to an integer or ignored as invalid data. For example, 1.1 is truncated to 1 or ignored.
- The order in which sections appear is not important; a particular section can be placed anywhere in the file.
- Each comment in the file begins with a semicolon (**;**).

For more information on these formats and conventions, see the *Microsoft Platform Software Development Kit* (SDK) at the Microsoft Web site.

**Table of ISP and Service Type keys**

The following table shows the valid keys for the [ISP] and [Service Type] sections of the .cms file, the function and use of each key, and where you must set each key. You can set different keys in the Connection Manager Administration Kit (CMAK) wizard or the .cms file. If you change the .cms file with a text editor, rather than by using the **Advanced Customization** pane of the CMAK wizard, you must rebuild the profile afterward with the CMAK wizard.

### Notes

- In the following table, keys controlled by specific panes of the CMAK wizard are identified in the **Set in** column as **CMAK**. Although these keys appear in the .cms file, you should use the CMAK wizard to update these keys. If you change these keys by editing the .cms file, the CMAK wizard might overwrite them if you run it again on the same profile. Keys set in the **Advanced Customization** pane of the CMAK wizard are identified as .cms or .cmp, as appropriate.
- Specify file locations either with the full path or the relative path from the .cmp file directory, as appropriate to the key.

| [Section] or key | Definition and values | Set in |
|---|---|---|
| **[ISP]** | Connection types (such as modem or ISDN) that you want to list in the **Phone Book** dialog box and how you want phone numbers to appear for each connection type. The Mask and Match keys specify whether the number is listed under **Access numbers** or under **More access numbers**. Each Mask and Match key has a corresponding option in **Phone Book Administrator**. | CMAK<br>.cms |
| PBFile | Location of the phone book (.pbk) file used by the service profile. | CMAK |
| RegionFile | Location of the region (.pbr) file used by the service profile. | CMAK |
| PBURL | Uniform Resource Locator (URL) for the phone book update. | CMAK |
| Mask&*FilterName* | Mask value for *FilterName*. | .cms |

| | | |
|---|---|---|
| Match&*FilterName* | Match value for *FilterName*. | .cms |
| FilterA&*ServiceProfile* | List of filter names to be applied to the access numbers phone-book settings field, for example:<br><br>FilterA& = NosurchargeSignon<br><br>The selected phone numbers are listed in the Connection Manager **Phone Book** dialog box, in **Access numbers**.<br><br>Default = NosurchargeSignon | .cms |
| FilterB&*ServiceProfile* | List of filter names to be applied to the referenced phone book options field, for example:<br><br>FilterB& = SurchargeSignon<br><br>The selected phone numbers are listed in the Connection Manager **Phone Book** dialog box, in **More access numbers**.<br><br>Default = SurchargeSignon. | .cms |
| References | Service profiles merged through the CMAK wizard. Profile names are separated by commas, tabs, or spaces. Each name is used as a reference to that profile. | CMAK |
| CMSFile&*MergedProfile* | Location of a component service profile .cms file for a merged service profile. The format is CMSFile&*MergedProfile=Location*, where *MergedProfile* is the name of the top-level service profile and *Location* is the location of the .cms file for the component service profile, for example:<br><br>CMSFile&Awesome1 = Awesome1\AwesomeE.cms | CMAK |
| FilterA&*MergedProfile* | List of filter names to be applied to the component profile's phone book options field. The format is FilterA&*MergedProfile = FilterName*, where *MergedProfile* is the name of the merged service profile, as specified in the CMSFile*MergedProfile* key, and *FilterName* is the name of the filter to be applied, for example:<br><br>FilterA&AwesomeE = NoSurchargeSignon<br><br>The selected phone numbers are merged into the list in the Connection Manager **Phone Book** dialog box, in **Access numbers**. | CMAK |
| FilterB&*MergedProfile* | List of filter names to be applied to the component profile's phone book options field. The format is FilterB&*MergedProfile = FilterName*, where *MergedProfile* is the name of the merged service profile, as specified in the .cms file **References** key, and *FilterName* is the name of the filter to be applied, for example:<br><br>FilterB&AwesomeE = SurchargeSignon<br><br>The selected phone numbers are merged into the list in the Connection Manager **Phone Book** dialog box, in **More access numbers**. | CMAK |
| **[Service Types]** | Items that are listed in the **Phone Book** dialog box, under **Service type**. By default, this section contains two service types. At least two service types must be defined in order for phone numbers to appear in the Properties dialog box of the Connection Manager client. | CMAK<br><br>.cms |
| *ServiceType* | A service type as a filter name to be applied to the referenced phone book options field. The format of the entry is *ServiceType=FilterName*, where *ServiceType* is the name as it appears in the **Phone Book** dialog box, in the **Service type** list, and *FilterName* is the name of the filter to be associated with the entry, for example:<br><br>ISDN Multicast = MulticastISDN<br><br>A service type name should include no more than 25 characters. You can concatenate service type names by inserting a space between each name. At least two service types must be defined in order for phone numbers to appear in the Properties dialog box of the Connection Manager client. | .cms |

# Connection Manager keys

Updated: December 12, 2008

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Manager keys**
You can change how Connection Manager handles aspects of its own operation. To do this, you need to edit the service (.cms) file. The table in this topic covers the key names and values that relate to the operation of Connection Manager. It also describes the formats and conventions you should use in the .cms file.

**Formats and conventions**
Connection Manager creates service (.cms) files and connection profile (.cmp) files, which use the same format as Windows initialization (.ini) files and information (.inf) files. These files have several things in common:

- Each file is divided into sections, with specific entries, called keys, assigned to the sections. A section heading is enclosed within brackets (**[ ]**). For example, **[Strings]** is a section containing text strings for messages. The keys for a section are not enclosed within brackets.
- Each value specified for a numeric key must be an integer. Any noninteger value is truncated to an integer or ignored as invalid data. For example, 1.1 is truncated to 1 or ignored.
- The order in which sections appear is not important; a particular section can be placed anywhere in the file.
- Each comment in the file begins with a semicolon (**;**).

For more information on these formats and conventions, see the *Microsoft Platform Software Development Kit* (SDK) at the Microsoft Web site.

**Table of Connection Manager keys**
The following table shows the valid keys for the [Connection Manager] section of the .cms and .cmp files, the function and use of each key, and where each key can be set. You can set different keys in the Connection Manager Administration Kit (CMAK) wizard, the .cms file, or the .cmp file. If you change the .cms file with a text editor, rather than by using the **Advanced Customization** pane of the CMAK wizard, you must rebuild the profile afterward with the CMAK wizard.

### Notes

- In the following table, keys controlled by specific panes of the CMAK wizard are identified in the **Set in** column as **CMAK**. Although these keys appear in the .cms file, you should use the CMAK wizard to update these keys. If you change these keys by editing the .cms file, the CMAK wizard might overwrite them if you run it again on the same profile. Keys set in the Advanced Customization portion of the CMAK wizard are identified as .cms or .cmp, as appropriate.
- Any changes you make to the .cmp file will be available only the first time the service profile is used. Editing this file allows you to provide a first-time-only population of user information. For example, you can fill in the user name, the password and the phone number fields for the first user to install and open your service profile.
- Make sure that locations are specified either as the full path or in reference to the location of the .cmp file.

| [Section] or key | Description and values | Set in |
|---|---|---|
| **[Connection Manager]** | Miscellaneous service profile keys and values. | CMAK<br>.cms<br>.cmp |

| | | |
|---|---|---|
| Dialup | Whether dial-up connections are supported. If this key is set to 0, the **General** tab of the Properties dialog box will not be displayed in Connection Manager.<br><br>0 = No.<br><br>1 (default) = Yes. | .cms |
| Direct | Whether direct connections are supported.<br><br>0 = No.<br><br>1 (default) = Yes.<br><br>This key is valid only for virtual private network (VPN) service profiles. | .cms |
| ConnectionType | The connection method to be used.<br><br>0 (default) = Connect the VPN by using a dial-up connection.<br><br>1 = Connect the VPN by using a direct connection (for example, DSL).<br><br>The direct connection (for example, DSL) must already exist and cannot be set up by using CMAK. This key is valid only if both the **Dialup** key and the **Direct** key values are set to **1**. The user can change this option in the **Properties** dialog box, on the **General** tab. | .cmp |
| CMSFile | Name of the service profile file: (*Path**ServiceProfileFileName*.cms), as entered in **File name** on the **Service and File Names** pane. Do not use any of the following in this name: **! @ ^ & ( ) { } `´ ~ , ; <space> * = \ / : ? ' " < > \| . % + [ ] $ #**. Do not alter this value in the .cmp file. If a different name is required, use the CMAK wizard to edit the existing profile, specify new names for both the **Service name** and the **File name**, and then finish running the wizard. | CMAK |
| ServiceName | The primary identifier for the profile. This text appears in the Connection Manager logon dialog box as the service name and as the name of the connection.<br><br>Maximum = approximately 35 characters. Do not use any of the following in this name: *** \ / : ? " < > \| [ ]** | CMAK |
| ServiceMessage | Connection Manager logon dialog box service message, usually a customer-support number for the service.<br><br>Maximum = Approximately 50 characters. | CMAK |
| UserName | Name used to authenticate the user, for first-time use. | .cmp |
| InternetUserName | The user name provided for the dial-up portion of a double-dial VPN connection, for first-time use. | .cmp |
| UseSameUserName | Whether to use a separate user name and password when connecting to the Internet service provider (for a VPN connection).<br><br>0 (default) = Use the separate user name and password as specified by the user in the **Properties** dialog box, on the **Internet Logon** tab. Service profiles built to support VPN connections contain this additional tab.1 = Use the same user name and password when connecting to the Internet service provider as is specified in the Connection Manager logon dialog box for logging on to the private service. | CMAK |
| UseWinLogonCredentials | Whether or not to incorporate the user name, password, and domain specified when the user logged on most recently to Windows (Windows 2000, Windows XP, and members of the Windows Server 2003 family only).<br><br>0 = Do not use credentials.<br><br>1 (default) = Use credentials. | .cms |

| | | |
|---|---|---|
| GlobalCredentials | Windows XP and Windows Server 2003 family only. Whether to save the user name and password supplied for first-time use of the connection for all users of the connection.<br><br>0 (default) = Save the user name and password for a single user.<br><br>1 = Save the user name and password for all users of the connection. | .cmp |
| KeepDefaultCredentials | Whether to keep first-time credentials such as user name and password in the .cmp file of the profile, for persistent default display and usage for all first-time users of the profile. For single-user profiles, this information will only be displayed once.<br><br>0 (default) = Do not keep first-time credentials in the .cmp file.<br><br>1 = Keep first-time credentials in the .cmp file. | .cmp |
| UserNamePrefix | Prefix added to the user name. Enables access to a service that requires a realm name as a user name prefix. This key will not be used in a VPN connection. | CMAK |
| UserNameSuffix | Suffix appended to the user name. Enables access to a service that requires a realm name as a user name suffix. This key will not be used in a VPN connection. | CMAK |
| MaxUserName | Maximum length of the user name.<br><br>Default = 256 bytes. | .cms |
| UserNameOptional | Whether the user must type a user name in the Connection Manager logon and Internet logon dialog boxes.<br><br>0 (default) = User name required.<br><br>1 = User name not required. | .cms |
| HideUserName | Whether the **User name** box appears in the Connection Manager logon dialog box.<br><br>0 (default) = Display the box.<br><br>1 = Do not display the box.<br><br>If the box is hidden, the data is considered optional. | .cms |
| HideInternetUsername | Whether the **User name** box appears in the **Properties** dialog box, **Internet Logon** tab.<br><br>0 (default) = Display the box.<br><br>1 = Do not display the box.<br><br>If the box is hidden, the data is considered optional. If **HideInternetPassword** and **HideInternetUsername** are both 1, the **Internet Logon** tab is not displayed. | .cms |
| Password | The password that appears automatically the first time the profile is used. | .cmp |
| InternetPassword | The password provided for the dial-up portion of a double-dial VPN connection. The password appears automatically the first time the user tries to connect using a double-dial VPN connection. | .cmp |
| PasswordHandling | Method of handling the password during the logon process.<br><br>0 (default) = Normal (case sensitive).<br><br>1 = Lowercase password.<br><br>2 = Uppercase password. | .cms |

| | | |
|---|---|---|
| MaxPassword | Maximum length of the password that can be typed by the user.<br><br>Default = 127 bytes. | .cms |
| PasswordOptional | Whether the user must specify a password in the Connection Manager logon and Internet logon dialog boxes.<br><br>0 (default) = Password required.<br><br>1 = Password not required. | .cms |
| RememberPassword | Whether to have the **Save password** check box in the Connection Manager logon dialog box selected by default for a populated password. This key is ignored if the **Password** key is not set in the .cmp file.<br><br>0 (default) = The check box is not selected by default.<br><br>1 = The check box is selected by default, and the password is saved. | .cmp |
| RememberInternetPassword | Whether to have the **Save password** check box on the **Internet Logon** tab selected by default for a populated Internet password. This key is ignored if the **InternetPassword** key is not set in the .cmp file.<br><br>0 (default) = The check box is not selected by default.<br><br>1 = The check box is selected by default, and the password is saved. | .cmp |
| HideRememberInternetPassword | Whether to show the **Save Password** check box on the **Internet Logon** tab. If this key is not specified, the value of **HideRememberPassword** is used by default.<br><br>0 = Display check box.<br><br>1 = Do not display check box. | .cms |
| ResetPassword | Whether the **New** button appears in the Connection Manager logon dialog box to enable a user to change the password. You can use this button to set up a custom password option. Type the location (including the name) of the executable (.exe) program you use to support this process. You must provide an executable program for this option and include it as an additional file when you run the CMAK wizard. | .cms |
| HidePassword | Whether the **Password** box appears in the Connection Manager logon dialog box when using dial-up connections.<br><br>0 (default) = Display the box.<br><br>1 = Do not display the box.<br><br>If the box is hidden, the data is considered optional. | .cms |
| HideInternetPassword | Whether the **Password** box appears in the **Properties** dialog box, on the **Internet Logon** tab.<br><br>0 (default) = Display the box.<br><br>1 = Do not display the box.<br><br>If the box is hidden, the data is considered optional. If **HideInternetPassword** and **HideInternetUsername** are both 1, the **Internet Logon** tab is not displayed. | .cms |
| DomainOptional | Whether to require the user to enter a domain name when making a connection.<br><br>0 = Require a domain.<br><br>1 (default) = Do not require a domain. | .cms |

| | | |
|---|---|---|
| HideRememberPassword | Whether to display the **Save password** check box in the Connection Manager logon dialog box and on the **Internet Logon** tab.<br><br>0 (default) = Display the check box.<br><br>1 = Do not display the check box.<br><br>If **HideRememberPassword** is set to **1**, it overrides the **HideDialAutomatically** setting. | .cms |
| Domain | Logon domain name that appears automatically the first time the profile is used. | .cmp |
| MaxDomain | Maximum length of the domain name that can be typed by the user.<br><br>Default = 15 characters. | .cms |
| HideDomain | Whether the **Logon domain** box appears in the Connection Manager logon dialog box when using dial-up connections.<br><br>0 = Display the box.<br><br>1 = Do not display the box.<br><br>The default is 1. If the profile supports VPN connections, CMAK will set this value to 0 if you have not previously set a value. If the box is hidden, the data is considered optional. | .cms |
| DialAutomatically | Whether to have the **Connect automatically** check box in the Connection Manager logon dialog box selected by default and to connect automatically when the service profile is used for the first time.<br><br>0 (default) = Do not select the check box.<br><br>1 = Select the check box and connect automatically. | .cmp |
| DialAutoMessage | Text displayed when the **Dial Automatically** check box has been selected. | .cms |
| HideDialAutomatically | Whether the **Connect automatically** check box appears in the Connection Manager logon dialog box.<br><br>0 (default) = Display the check box.<br><br>1 = Do not display the check box.<br><br>If **HideRememberPassword** is set to **1**, that key value overrides the **HideDialAutomatically** setting. | .cms |
| Phone0 | Primary phone number populated for the user the first time the service profile is used, and used as the default phone number any time another phone number is not selected. | .cmp |
| Phone1 | Backup phone number populated for the user the first time the service profile is used. | .cmp |
| PhoneCanonical0 | The canonical format of the primary phone number the first time the service profile is used. | .cmp |
| PhoneCanonical1 | The canonical format of the backup phone number the first time the service profile is used. | .cmp |
| MaxPhoneNumber | Specifies a character limit for the phone number fields. The default for Windows 95, Windows 98, and Windows Millennium Edition is 36; the default for Windows NT 4.0, Windows 2000, Windows XP, and members of the Windows Server 2003 family is 80. | .cms |
| PhoneCountry0 | Telephony Application Programming Interface (TAPI) country/region code for the primary phone number set as default for the user the first time the service profile is used. | .cmp |

| | | |
|---|---|---|
| PhoneCountry1 | Telephony Application Programming Interface (TAPI) country/region code for the backup phone number set as default for the user the first time the service profile is used. | .cmp |
| UseDialingRules0 | Whether dialing rules are applied to the primary phone number the first time the service profile is used.<br><br>0 = Use the form of the number exactly as it appears in the phone book or as it is typed.<br><br>1 (default) = Apply dialing rules to the number when displaying and dialing. | .cmp |
| UseDialingRules1 | Whether dialing rules are applied to the backup number the first time the service profile is used.<br><br>0 = Use the form of the number exactly as it appears in the phone book or as it is typed.<br><br>1 (default) = Apply dialing rules to the number when displaying and dialing. | .cmp |
| NoDialingRules | Whether the user can use dialing rules on phone numbers selected from a phone book. Dialing rules check boxes are grayed out when this key value is set to 1.<br><br>0 (default) = Allow dialing rules.<br><br>1 = Do not allow dialing rules. | .cms |
| PhoneSource0 | Service profile for the primary phone number to be used the first time the user connects. | .cmp |
| PhoneSource1 | Service profile for the backup phone number to be used the first time the user connects. | .cmp |
| Description0 | Description of the primary phone number provided the first time the service profile is opened. Appears in the **Phone Book** dialog box as a Point of Presence (POP) name. | .cmp |
| Description1 | Description of the backup phone number provided the first time the service profile is opened. Appears in the **Phone Book** dialog box as a POP name. | .cmp |
| DUN0 | Dial-up entry used for the primary phone number the first time the service profile is opened. Value must exactly match the entry name specified in **Phone Book Administrator**. | .cmp |
| DUN1 | Dial-up entry used for the backup phone number the first time the service profile is opened. Value must exactly match the entry name specified in **Phone Book Administrator**. | .cmp |
| ServiceType0 | Default service type from the phone book for the primary phone number the first time that the service profile is opened. | .cmp |
| ServiceType1 | Default service type from the phone book for the backup phone number the first time that the service profile is opened. | .cmp |
| Region0 | Default state or province from the phone book for the primary phone number the first time that the service profile is opened. | .cmp |
| Region1 | Default state or province from the phone book for the backup phone number the first time that the service profile is opened. | .cmp |

| | | |
|---|---|---|
| Version | Version of the phone book. If you use Connection Point Services (CPS) to provide phone books, the CPS phone book server will use this value to determine how to handle phone book updates. If the version on the client is more than five versions out of date, the CPS phone book server downloads the complete phone book. Otherwise, the CPS phone book server downloads a delta file to the client. | .cms |
| DUN | Specifies the default network entry used when the **Dial-up Networking entry** field for a POP is blank.<br><br>Default = *ServiceProfileServiceName*. | .cms |
| RedialCount | Number of redials to attempt the first time the service profile is opened.<br><br>Default = 3. | .cmp |
| RedialDelay | Delay (in seconds) between dials. The delay is the elapsed time between the point of failure of the first call and the dialing of the first digit of the next phone number.<br><br>Default = 5 seconds. | .cms |
| IdleTimeout | Time (in minutes) that the connection must be idle before disconnecting the first time the service profile is used to connect.<br><br>Default = 10 minutes.<br><br>If set to 0, no time-out occurs. | .cmp |
| IdleThreshold | Maximum bytes per minute that can be received if an idle state is to be detected for the connection.<br><br>Default = 0.<br><br>This key is not available for use on Windows NT 4.0 systems. | .cms |
| AutoReconnect | Whether to automatically attempt to reconnect when a connection is terminated unexpectedly.<br><br>0 (default) = Do not automatically attempt to reconnect.<br><br>1 = Attempt to reconnect when the connection is terminated unexpectedly. | .cms |
| NoPromptReconnect | Whether the user is prompted for reconnection when the connection is inadvertently dropped.<br><br>0 (default) = Prompt the user.<br><br>1 = Do not prompt the user. | .cms |
| Icon | Location and name of the icon (.ico) file for the large program icon.<br><br>Default = The default value for this key is blank. When nothing is specified, the standard Connection Manager program icon is used. | CMAK |
| SmallIcon | Location and name of the icon (.ico) file for the small title-bar icon.<br><br>Default = The default value for this key is blank. When nothing is specified, the standard Connection Manager title-bar icon is used. | CMAK |
| TrayIcon | Location and name of the icon (.ico) file for the notification area shortcut.<br><br>Default = The default value for this key is blank. When nothing is specified, the Connection Manager notification area shortcut is used. | CMAK |

| | | |
|---|---|---|
| HideTrayIcon | Whether to display notification area shortcuts.<br><br>*NULL* = If **TrayIcon** is set, display the notification area shortcut specified. Otherwise, display the connection-status notification area shortcut ("blinking lights"). The *NULL* value is available for Windows 2000, Windows XP, and members of the Windows Server 2003 family only.<br><br>0 = If **TrayIcon** is set, display the notification area shortcut specified and the connection-status notification area shortcut ("blinking lights"). Otherwise, display the Connection Manager notification area shortcut and the connection-status notification area shortcut ("blinking lights").<br><br>1 = Do not display any notification area shortcuts.<br><br>If you choose not to display any notification area shortcuts, edit the Help file to remove references to it. | .cms |
| Logo | Location and name of the .bmp file displayed in the Connection Manager logon dialog box.<br><br>Default = The default value for this key is blank. When nothing is specified, the standard Connection Manager logon bitmap is used. | CMAK |
| PBLogo | Location and name of the .bmp file for the **Phone Book** dialog box.<br><br>Default = The default value for this key is blank. When nothing is specified, the standard Phone Book bitmap is used. | CMAK |
| PBMessage | Text in the **Phone Book** dialog box next to **More access numbers** when alternate access numbers are available for the selected geographic region.<br><br>Maximum = Approximately 100 characters. | CMAK |
| PBUpdateMessage | Text that appears in the dialog box during a phone book update.<br><br>Default = "Downloading new phone book." | .cms |
| DownloadDelay | Time (in seconds) to delay before downloading phone book updates. The delay time starts when the connection is established.<br><br>Default = 15 seconds.<br><br>This entry is valid only when, in the CMAK wizard, in the **Phone Book** dialog box, **Automatically download phone book updates** is selected. | .cms |
| HideDelay | Time (in seconds) to delay before displaying the **Phone Book Updates** dialog box (with progress indicators and message).<br><br>Default = Do not display the dialog box. | .cms |
| HelpFile | Location (including name) of the custom Help file.<br><br>Default Help file = The default value for this key is blank. When nothing is specified, the standard Help file for service profiles is used. | CMAK |
| Tunnel | Whether VPN connections are enabled for the phone numbers in the top-level phone book.<br><br>0 (default) = Not enabled.<br><br>1 = Enabled. | CMAK |
| TunnelDUN | Specifies the default networking entry used for VPN connections.<br><br>Default = *ServiceProfileServiceName* tunnel. | .cms |

| | | |
|---|---|---|
| TunnelFile | The relative path and file name to the VPN file. If you use macros in the relative path, make sure the macros are recognized on all operating systems to which you want to distribute the profile. | CMAK |
| TunnelAddress | IP address or domain name for the VPN server. | CMAK |
| TunnelReferences | Whether VPN connections are enabled for phone numbers taken from a component profile phone book in a top-level service profile.<br><br>0 (default) = Do not use a VPN connection when using any of these phone numbers.<br><br>1 = Use a VPN connection when using a phone number from a component profile. | CMAK |
| AnimatedLogo | Whether an animation is used instead of a static bitmap in the Connection Manager logon dialog box.<br><br>0 (default) = Do not use animated graphic.<br><br>1 = Use animated graphic.<br><br>The **[Animated Logo]** section specifies the location of the animation (.spl or .swf) file. This function is provided for compatibility with Connection Manager 1.0, which supports Macromedia Flash animation files. | CMAK<br><br>.cms |
| CheckOSComponents | Whether Connection Manager checks the operating system components.<br><br>0 (default) = Check the operating system components only when an error occurs.<br><br>1 = Always check the operating system components on service profile initialization. | .cms |
| DoNotCheckBindings | Whether to include the TCP/IP bindings when checking operating system components.<br><br>0 (default) = Include the TCP/IP bindings in the TCP/IP check and bind all that are not bound.<br><br>1 = Do not include the TCP/IP bindings when checking operating system components.<br><br>TCP/IP services that sit between the TCP/IP protocol and the dial-up adapter are protected from checking only if this option is set to 1 in all installed service profiles. This option does not prevent bindings from being checked by other service profiles or programs. | .cms |
| IsdnDialMode | Whether Connection Manager binds two channels on the same ISDN device.<br><br>0 (default) = Do not bind (use single channel only).<br><br>1 = Always bind the two channels (connect both channels automatically). If a dual channel connection cannot be made, the connection attempt will fail. This entry is valid only for Windows NT 4.0, Windows 2000, Windows Millennium Edition, Windows XP, and members of the Windows Server 2003 family.<br><br>2 = Attempt to bind the two channels (connect with both channels). If the two channels cannot be bonded, use single channel. This entry is valid only for Windows 2000, Windows XP, and members of the Windows Server 2003 family. | .cms |

| | | |
|---|---|---|
| EnableICF | Windows XP and Windows Server 2003 family only. Whether to turn on Internet Connection Firewall (ICF) or Windows Firewall for this connection with default settings.<br><br>0 = Do not adjust ICF or Windows Firewall settings. Default for the original releases of the Windows XP and Windows Server 2003 operating systems.<br><br>1 = Enable ICF or Windows Firewall for the connection with default settings. Default for Windows XP with Service Pack 2 and for Windows Server 2003 with Service Pack 1.<br><br>Note: Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition. | .cms |
| DisableICS | Windows XP and Windows Server 2003 family only. Whether to suppress the user's ability to share the connection with Internet Connection Sharing (ICS).<br><br>0 (default) = Do not suppress ICS settings.<br><br>1 = Disable ICS settings for this connection.<br><br>Note: Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems. | .cms |
| HideAdvancedTab | Windows XP and Windows Server 2003 family only. Whether to suppress the display of the Advanced tab in the Connection Manager user interface for this connection.<br><br>0 (default) = Do not suppress display of the Advanced tab.<br><br>1 = Suppress display of the Advanced tab. | .cms |
| CustomButtonText | Text displayed on a custom button on the main connect screen. This feature is not supported when logging on to Windows. | .cms |
| CustomButtonToolTip | Text displayed on a tooltip for a custom button on the main connect screen. This feature is not supported when logging on to Windows. | .cms |
| MonitorCallingProgram | When Connection Manager is called programmatically, whether to handle Connection Manager as a monitored custom action when the calling process terminates.<br><br>0 = Do not monitor.<br><br>1 (default) = Monitor. | .cms |
| EnableLogging | Whether a log file is created for the connection the first time the service profile is opened. The log file is named *ServiceProfileServiceName*.log for an all user profile, and *ServiceProfileServiceName* (Single User).log for a single user profile.<br><br>0 = Do not create a log file.<br><br>1 (default) = Create a log file. | .cmp |
| HideBalloonTips | Whether to suppress display of balloon-tip help for this connection.<br><br>0 (default) = Do not suppress balloon-tip help.<br><br>1 = Suppress balloon-tip help. | .cms |
| InternetConnection | Windows XP and Windows Server 2003 family only. Whether to inform the operating system that this connection is a connection to the Internet.<br><br>0 (default) = Do not inform the operating system that this connection connects to the Internet for the purposes of automatically connecting when no connection is present.<br><br>1 = Allow the system to register this connection as a connection to the Internet. | .cms |

# Custom actions keys

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Custom actions keys**

You can change how Connection Manager handles custom actions. To do this, you need to edit the service (.cms) file. The table in this topic covers the key names and values that pertain to the custom actions sections of the .cms file. This topic also describes the formats and conventions you should use in the .cms file.

**Formats and conventions**

Connection Manager creates service (.cms) files, which use the same format as Windows initialization (.ini) files and information (.inf) files. These files have several things in common:

- Each file is divided into sections, with specific entries, called keys, assigned to the sections. A section heading is enclosed within brackets (**[ ]**). For example, **[Strings]** is a section containing text strings for messages. The keys for a section are not enclosed within brackets.
- Each value specified for a numeric key must be an integer. Any noninteger value is truncated to an integer or ignored as invalid data. For example, 1.1 is truncated to 1 or ignored.
- The order in which sections appear is not important; a particular section can be placed anywhere in the file.
- Each comment in the file begins with a semicolon (**;**).

For more information on these formats and conventions, see the *Microsoft Platform Software Development Kit* (SDK) at the Microsoft Web site.

**Table of custom actions keys**

The following table shows the valid keys for custom actions in the .cms file, the function and use of each key, and where each key can be set. You can set different keys in the Connection Manager Administration Kit (CMAK) wizard or in the .cms file. If you change the .cms file with a text editor, rather than by using the **Advanced Customization** pane of the CMAK wizard, you must rebuild the profile afterward with the CMAK wizard.

### Notes

- In the following table, keys controlled by specific panes of the CMAK wizard are identified in the **Set in** column as **CMAK**. Although these keys appear in the .cms file, you should use the CMAK wizard to update these keys. If you change these keys by editing the .cms file, the CMAK wizard might overwrite them if you run it again on the same profile. Keys set in the **Advanced Customization** pane of the CMAK wizard are identified as .cms or .cmp, as appropriate.
- Make sure that locations are specified either as the full path or in reference to the location of the .cmp file.

| [Section] or key name | Value | Set in |
|---|---|---|
| **[Pre-Init Actions]** | Pre-init actions that run as soon as the profile is accessed. This section uses the following keys: **#**, **#&Flags**, and **#&Description**. | CMAK .cms |
| **[Pre-Connect Actions]** | Pre-connect actions that run before a connection is established. This section uses the following keys: **#**, **#&Flags**, and **#&Description**. | CMAK .cms |
| **[Pre-Dial Actions]** | Pre-dial actions that run immediately before dialing. This section uses the following keys: **#**, **#&Flags**, and **#&Description**. | CMAK .cms |

| | | |
|---|---|---|
| **[Pre-Tunnel Actions]** | Pre-tunnel actions that run after a connection is established but before a VPN connection is established. This section uses the following keys: **#**, **#&Flags**, and **#&Description**. | CMAK .cms |
| **[Connect Actions]** | Post-connect actions that run after a connection is established and, if using a VPN connection, after the tunnel is established. This section uses the following keys: **#**, **#&Flags**, and **#&Description**. | CMAK .cms |
| **[Disconnect Actions]** | Disconnect actions that run immediately before disconnecting. If a connection is prematurely dropped, these actions are run immediately after the disconnect. This section uses the following keys: **#**, **#&Flags**, and **#&Description**. | CMAK .cms |
| **[On-Cancel Actions]** | On cancel actions that run when the user cancels a session after the connection attempt has been started (any time that the user clicks **Connect** and then clicks **Cancel**). This section uses the following keys: **#**, **#&Flags**, and **#&Description**. | CMAK .cms |
| **[On-Error Actions]** | On error actions that run when an error occurs. This section uses the following keys: **#**, **#&Flags**, and **#&Description**. | CMAK .cms |
| **[Auto Applications]** | Monitored actions that run during a connection. Monitored actions are similar to custom actions, but they cause the disconnect sequence to start as soon as the last monitored application closes. This section uses the following keys: **#**, **#&Flags**, and **#&Description**. Unlike other custom actions, monitored actions do not support DLLs. | CMAK .cms |
| # | A single custom action. The format of the value is # = *Program*, where # is an ID number from 0 to *n* and *Program* is the complete name of the program, including the extension and the path to the program. Programs start in sequence: 0 first, 1 next, and so on. Environmental variables, such as **%UserName%** and **%ErrorCode%**, are set in Connection Manager and can be used as parameters for custom action programs. | CMAK |
| #&Flags | Whether a custom action runs when a specific type of connection is established. The format of the value is #&Flags = *n*, where # is the ID number of a defined custom action and *n* is a value indicating how dial-up and direct connections for that custom action are handled:<br>• If *n* = 0 (default), all connections trigger this custom action. (The default key does not appear in the .cms file.)<br>• If *n* = 1, all direct tunnel connections trigger this custom action.<br>• If *n* = 2 , all connections that involve dialing (direct or double) trigger this custom action.<br>• If *n* = 4, only connections that use direct dial-up trigger this custom action.<br>• If *n* = 8, only connections that involve a tunnel trigger this custom action. | CMAK |
| #&Description | A short description of a custom action. The format of the value is #&Description = *text*, where # is the number of a defined custom action and *text* is the descriptive information for that custom action. This option can provide useful information when a synchronous custom action does not run properly.<br>Default = the name of the custom action program. | CMAK |

# Networking keys

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Networking keys**
You can change how Connection Manager handles networking. To do this, you need to edit the service (.cms) file. The table in this topic covers the key names and values that pertain to the networking sections of the .cms file. This topic also describes the formats and conventions you should use in the .cms file.

**Formats and conventions**
Connection Manager creates service (.cms) files, which use the same format as Windows initialization (.ini) files and information (.inf) files. These files have several things in common:

- Each file is divided into sections, with specific entries, called keys, assigned to the sections. A section heading is enclosed within brackets (**[ ]**). For example, **[Strings]** is a section containing text strings for messages. The keys for a section are not enclosed within brackets.

- Each value specified for a numeric key must be an integer. Any noninteger value is truncated to an integer or ignored as invalid data. For example, 1.1 is truncated to 1 or ignored.

- The order in which sections appear is not important; a particular section can be placed anywhere in the file.

- Each comment in the file begins with a semicolon (**;**).

For more information on these formats and conventions, see the *Microsoft Platform Software Development Kit* (SDK) at the Microsoft Web site.

**Table of networking keys**
The following table shows the valid keys for networking in the .cms file, the function and use of each key, and where each key can be set. You can set different keys in the Connection Manager Administration Kit (CMAK) wizard or the .cms file. If you change the .cms file with a text editor, rather than by using the **Advanced Customization** pane of the CMAK wizard, you must rebuild the profile afterward with the CMAK wizard.

### Notes

- In the following table, keys controlled by specific panes of the CMAK wizard are identified in the **Set in** column as **CMAK**. Although these keys appear in the .cms file, you should use the CMAK wizard to update these keys. If you change these keys by editing the .cms file, the CMAK wizard might overwrite them if you run it again on the same profile. Keys set in the **Advanced Customization** pane of the CMAK wizard are identified as .cms or .cmp, as appropriate.

- Make sure that locations are specified either as the full path or in reference to the location of the .cmp file.

| [Section] or key | Value | Set in |
|---|---|---|
| **[Server& *TunnelDUN* ]** | Configuration to be used for all VPN connections. In this section title, ***TunnelDUN*** is the value of the **TunnelDUN** key in the **[Connection Manager]** section, for example, **[Server&Awesome1 Tunnel]**. This section uses the same format as the **[Server&*EntryName*]** section, including all keys valid under that section. | CMAK<br><br>.cms |
| **[Server& *EntryName* ]** | Configuration where *EntryName* is the name of the dial-up entry, for example, **[Server&Awesome1]**. | CMAK |

| | | |
|---|---|---|
| SecureLocalFiles | For Windows NT, Windows 2000, Windows XP, and the Windows Server 2003 family only. Whether to share local files and allow printer sharing. If the value of the key is set to 1, it will disable NetBT on Windows 2000, Windows XP, and the Windows Server 2003 family. This key value will be ignored on Windows XP and the Windows Server 2003 family if a value is assigned to the SecureClientForMSNet key, the SecureFileAndPrint key, or both.<br><br>0 (default) = Share local files.<br><br>1 = Do not share local files. | CMAK |
| SecureClientForMSNet | For Windows XP and Windows Server 2003 family only. Whether to allow the use of NetBT. If this key is set, the value of the SecureLocalFiles key will be ignored on Windows XP and the Windows Server 2003 family.<br><br>0 (default) = Allow the use of NetBT.<br><br>1 = Do not allow the use of NetBT. | .cms |
| SecureFileAndPrint | For Windows XP and Windows Server 2003 family only. Whether to enable file and print sharing. If this key is set, the value of the SecureLocalFiles key will be ignored on Windows XP and the Windows Server 2003 family.<br><br>0 (default) = Enable.<br><br>1 = Disable. Files and print resources will not be shared. | .cms |
| DataEncrypt | Whether to force data encryption. On Windows 2000, Windows XP, and Windows Server 2003 family, if the EncryptionType key is set, the value of the DataEncrypt key is ignored.<br><br>0 (default for dial-up entries) = Do not encrypt data.<br><br>1 (default for VPN entries) = Encrypt data. | CMAK |
| EncryptionType | Which encryption level is to be used.<br><br>0 = None.<br><br>1 (default for VPN) = Require.<br><br>3 (default for dial-up) = Optional.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |
| DontNegotiateMultilink | For Windows XP and Windows Server 2003 family only. Whether to allow multilinking.<br><br>0 (default) = Allow multilink.<br><br>1 = Do not allow multilink. | .cms |
| DontUseRasCredentials | For Windows XP and Windows Server 2003 family only. Whether to cache the user's credentials in **Stored User Names and Passwords** for use when accessing resources after logging onto the network.<br><br>0 (default) = Cache and use credentials.<br><br>1 = Do not cache and use credentials. | .cms |
| NetworkLogon | Whether to log on to the network.<br><br>0 = Do not log on. This is the default for dial-up connections.<br><br>1 = Log on. This is the default for VPN connections. | CMAK |
| SW_Compress | Whether to negotiate PPP software compression of data.<br><br>0 = Do not negotiate.<br><br>1 (default) = Negotiate. | .cms |

| | | |
|---|---|---|
| Disable_LCP | Whether to enable Link Control Protocol (LCP) extensions.<br><br>0 (default) = Enable.<br><br>1 = Disable. | .cms |
| PW_Encrypt | Whether to encrypt a password when authenticating.<br><br>0 (default) = Do not encrypt.<br><br>1 = Encrypt. | CMAK |
| PW_EncryptMS | Whether to require MS-CHAP version 1 for this connection.<br><br>0 = Do not require MS-CHAP v1.<br><br>1 = Require MS-CHAP v1. | CMAK |
| Negotiate_TCP/IP | Whether to negotiate TCP/IP.<br><br>0 = Do not negotiate.<br><br>1 (default) = Negotiate. | .cms |
| Negotiate_IPX | Whether to negotiate the IPX protocol.<br><br>0 (default) = Do not negotiate.<br><br>1 = Negotiate.<br><br>The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family. | .cms |
| Negotiate_NetBEUI | Whether to negotiate the NetBIOS extended user interface (NetBEUI) protocol.<br><br>0 (default) = Do not negotiate.<br><br>1 = Negotiate | .cms |
| Require_PAP | Whether Password Authentication Protocol (PAP) is allowed.<br><br>0 (default) = Not allowed.<br><br>1 = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |
| Require_SPAP | Whether SPAP is allowed.<br><br>0 (default) = Not allowed.<br><br>1 = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |
| Require_EAP | Whether Extensible Authentication Protocol (EAP) is allowed. Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family.<br><br>0 (default) = Not allowed.<br><br>1 = allowed. | CMAK |
| Require_CHAP | Whether Challenge Handshake Authentication Protocol (CHAP) is allowed.<br><br>0 = Not allowed.<br><br>1 (default) = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |

| | | |
|---|---|---|
| Require_MSCHAP | Whether MSCHAP is allowed.<br><br>0 = Not allowed.<br><br>1 (default) = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |
| Require_MSCHAP2 | Whether MSCHAP2 is allowed.<br><br>0 = Not allowed.<br><br>1 (default) = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |
| Require_W95MSCHAP | Whether W95MSCHAP is allowed. If the value for this key is set to 1, the value for **Require_MS-CHAP** must also be set to 1.<br><br>0 (default) = Not allowed.<br><br>1 = Allowed.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | .cms |
| Custom_Security | Applies specific security settings on computers that are running Windows 2000, Windows XP, or a member of the Windows Server 2003 family.<br><br>0 = Use standard configuration.<br><br>1 (default) = Use custom configuration.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |
| CustomAuthKey | The ID of the Extensible Authentication Protocol (EAP) to be used with this connection. The format is CustAuthKey=*ID*, where *ID* is the Internet Assigned Numbers Authority (IANA) ID for the EAP. If no ID is specified, no attempt will be made to read the CustomAuthData value.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |
| CustomAuthData# | The configuration data for the EAP. The format is CustomAuthData#=*EAPConfigurationData*, where # is the number of the specific line of data, for example:<br><br>CustomAuthData0=1AFGKLMOFGH23K<br><br>Each line can contain a maximum of 128 characters. If the ID has more than 128 characters, it must be specified as multiple lines, incrementing the # for each line of data. For example:<br><br>CustomAutData0=*First128Characters*<br><br>CustomAuthData1=*Next128Characters*<br><br>CustomAuthData2=*FinalCharacters*<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |
| EnforceCustomSecurity | Whether to enforce the user to select a phone book entry that includes support for advanced security settings.<br><br>0 (default) = Do not enforce phone book entry selection (Windows 95, Windows 98, Windows NT 4.0, and Windows Millennium Edition only).<br><br>1 = Enforce selection of phone book entries that support advanced security settings. | CMAK |

| | | |
|---|---|---|
| DisableNbtOverIP | Whether to disable the NetBT routing protocol for this connection.<br><br>0 (default) = Do not disable NetBT.<br><br>1 = Leave NetBT enabled. | .cms |
| **[TCP/IP& *TunnelDUN* ]** | TCP/IP configuration to be used for all VPN connections. In this section title, ***TunnelDUN*** is the value of the **TunnelDUN** key in the **[Connection Manager]** section, for example, **[TCP/IP&Awesome1 Tunnel]**. This section uses the same format as the **[TCP/IP&*EntryName*]** section, including all keys valid under that section. This section is created by the CMAK wizard, but the keys can be edited in the .cms file. | CMAK<br><br>.cms |
| **[TCP/IP& *EntryName* ]** | Configuration where *EntryName* is the name of the dial-up entry, for example: **[TCP/IP&Awesome1]** | CMAK<br><br>.cms |
| Specify_IP_Address | Whether to specify a static IP address.<br><br>0 (default) = Do not specify.<br><br>1 = Specify. | .cms |
| IP_Address | Static IP addresses. | .cms |
| Specify_Server_Address | Whether to specify static DNS or Windows Internet Name Service (WINS) addresses.<br><br>0 (default) = Do not specify.<br><br>1 = Specify. | CMAK |
| DNS_Address | DNS address. | CMAK |
| DNS_Alt_Address | Alternate DNS address. | CMAK |
| WINS_Address | WINS address. | CMAK |
| WINS_Alt_Address | Alternate WINS address. | CMAK |
| IP_Header_Compress | Whether to enable IP header compression.<br><br>0 = Not enabled.<br><br>1 (default) = Enabled. | CMAK |
| Gateway_On_Remote | Whether to use the default gateway on the remote private network.<br><br>0 = Do not use.<br><br>1 (default) = Use. | CMAK |
| DNSSuffix | DNS suffix to use for the connection. Valid only for Windows XP and Windows Server 2003 family. | .cms |
| TcpWindowSize | For Windows 2000, Windows XP, and the Windows Server 2003 family only. Sets the size (expressed in bytes) of the window offered by the connection. Number of bytes of information to send before expecting a return packet. Minimum value is 4,096; maximum value is 65,500. If this key is not set, the default window size for the operating system and the connection type will be used. | .cms |
| **[Networking& *TunnelDUN* ]** | The single configuration to be used for all VPN connections. In this section title, ***TunnelDUN*** is the value of the **TunnelDUN** key in the **[Connection Manager]** section, for example, **[Networking&Awesome1 Tunnel]**. | CMAK |

| | | |
|---|---|---|
| VpnStrategy | Which tunneling protocol to use when establishing a VPN connection.<br><br>1 = Use PPTP only.<br><br>2 = Try PPTP first.<br><br>3 = Use L2TP only.<br><br>4 = Try L2TP first.<br><br>Valid only for Windows 2000, Windows XP, and the Windows Server 2003 family. | CMAK |
| UsePreSharedKey | Windows 2000, Windows XP, and the Windows Server 2003 family only. Whether to use a pre-shared key for L2TP/IPSec authentication.<br><br>0 (default) = Do not use a pre-shared key.<br><br>1 = Use a pre-shared key. | CMAK |
| VPNEntry | Specifies to CMAK whether the entry is a VPN entry and not a dial-up entry.<br><br>0 = Dial-up entry.<br><br>1 = VPN entry. | CMAK |
| **[Scripting&** *EntryName* **]** | Location (path and file name) of the script (.scp) file, where *EntryName* is the name of the dial-up entry, for example: **[Scripting&Awesome1]**). | CMAK |
| Name | Location (including path and file name) of the script (.scp) file for the dial-up entry. | CMAK |
| UseTerminalWindow | For Windows 2000, Windows XP, and the Windows Server 2003 family only. Whether to use a terminal window when connecting to a service.<br><br>0 (default) = Do not use a terminal window.<br><br>1 = Use a terminal window. | .cms |

# Other sections, keys, and values

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Other sections, keys, and values**
This topic covers sections, keys, and values that might appear in service (.cms) files, particularly those created with older versions of the Connection Manager Administration Kit (CMAK) wizard.

**Formats and conventions**
Connection Manager creates service (.cms) files, which use the same format as Windows initialization (.ini) files and information (.inf) files. These files have several things in common:

- Each file is divided into sections, with specific entries, called keys, assigned to the sections. A section heading is enclosed within brackets (**[ ]**). For example, **[Strings]** is a section containing text strings for messages. The keys for a section are not enclosed within brackets.
- Each value specified for a numeric key must be an integer. Any noninteger value is truncated to an integer or ignored as invalid data. For example, 1.1 is truncated to 1 or ignored.
- The order in which sections appear is not important; a particular section can be placed anywhere in the file.
- Each comment in the file begins with a semicolon (**;**).

For more information on these formats and conventions, see the *Microsoft Platform Software Development Kit* (SDK) at the Microsoft Web site.

**Table of other sections, keys, and values**
The following table shows valid sections, keys, and values in the .cms file, the function and use of each key, and where each key can be set. You can set different keys in the CMAK wizard or the .cms file. If you change the .cms file with a text editor, rather than by using the **Advanced Customization** pane of the CMAK wizard, you must rebuild the profile afterward with the CMAK wizard.

### Notes

- In the following table, keys controlled by specific panes of the CMAK wizard are identified in the **Set in** column as **CMAK**. Although these keys appear in the .cms file, you should use the CMAK wizard to update these keys. If you change these keys by editing the .cms file, the CMAK wizard might overwrite them if you run it again on the same profile. Keys set in the **Advanced Customization** pane of the CMAK wizard are identified as .cms or .cmp, as appropriate.
- Make sure that locations are specified either as the full path or in reference to the location of the .cmp file.

| [Section] or key | Value | Set in |
|---|---|---|
| **[Logging]** | Section header for logging keys. | .cms |
| FileDirectory | Location of the log file. If the value is blank, the log file is created in the temp folder. | .cms |
| MaxFileSize | The maximum file size (in kilobytes) before the log is cleared. The file size is checked at initialization to prevent data loss during a connection. The default is 100. | .cms |
| **[Menu Options]** | Shortcut menu that appears when a user right-clicks the notification area shortcut for your service. | CMAK<br>.cms |

| | | |
|---|---|---|
| *MenuItemName* | Command to be run from the shortcut menu. Shortcut menu commands use the format *MenuItemName* = *Program*, where *MenuItemName* is the descriptive name that appears on the shortcut menu and *Program* is the name of the executable file (.exe or other shell executable), including any parameters required for the program. For example, if you wanted to add Notepad to the menu with the command name "Editor," the line would be **Editor = Notepad.exe**.<br><br>Default commands = **Status** and **Disconnect**.<br><br>These programs are available in a shortcut menu that appears when a user right-clicks the notification area shortcut. Applications are listed in the menu in the same sequence as they are listed in the .cms file. | CMAK |
| **[Animated Logo]** | Animation for the Connection Manager logon dialog box.<br><br>The **[Animated Logo]** section specifies the location of the animation (.spl or .swf) file. This function is provided for compatibility with Connection Manager 1.0, which supports Macromedia Flash animation files.<br><br>**AnimatedLogo** in the **[Connection Manager]** section determines whether an animation is used. | .cms |
| Movie | Location of the animation file displayed in place of the static bitmap in the logon dialog box. Settings in the **[Animated Actions]** section control how the animation plays based on the connection state. | .cms |
| *Parameter* | Other animation parameters passed from software used to create the animation, such as Quality = 3. | .cms |
| **[Animated Actions]** | When to play segments of the animation. | CMAK<br><br>.cms |
| Initial | Starting frame number; plays when Connection Manager starts. | .cms |
| Dialing0 | Starting frame number; plays while the primary phone number is dialed. | .cms |
| Pausing | Starting frame number; plays during pauses between dials. | .cms |
| Dialing1 | Starting frame number; plays while the backup phone number is dialed. | .cms |
| Authenticating | Starting frame number; plays during user authentication. | .cms |
| Tunneling | Starting frame number; plays while establishing a VPN connection. | .cms |
| Connected | Starting frame number; plays when a connection is established. | .cms |
| Error | Starting frame number; plays at any connection error. | .cms |
| **[CustomButton Actions]** | Controls what programs start when the custom button is clicked. The format of the value is # = *Program*, where # is an ID number from 0 to *n* and *Program* is the complete name of the program, including the extension and the path to the program. Programs start in sequence: 0 first, 1 next, and so on. Custom button programs cannot be DLL files. | CMAK<br><br>.cms |
| **[Profile Format]** | Service profile version. | CMAK<br><br>.cms<br><br>.cmp |
| Version | Format version of the profile or service file. | .cms<br><br>.cmp |

# Phase Three: Running the CMAK Wizard and Creating a Service Profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phase three: Running the CMAK wizard and creating a service profile**
Using the Connection Manager Administration Kit (CMAK) wizard, you can quickly and easily build the required service profiles to configure Connection Manager.

All custom elements must be ready before you run the CMAK wizard to build a service profile. Get organized before you start the wizard, and ensure that you have everything you need, including all external files required to implement the custom elements.

This section covers information that you need to know to run the CMAK wizard and create a service profile:

- Preparing to run the CMAK wizard
- Using the CMAK wizard to build a service profile
- Finishing customization after running the CMAK wizard

# Preparing to run the CMAK wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Preparing to run the CMAK wizard**
You run the Connection Manager Administration Kit (CMAK) wizard to create a service profile. As you step through the wizard, you are prompted to select customizable options and type the location of customized elements that you have prepared. Before you can run the wizard, you must set it up on your computer and ensure that all custom elements are available. To help prepare for creating a service profile, you can complete the planning worksheet that is available in the Microsoft Windows Server 2003 *Resource Kit*. For more information about Connection Manager, see "Deploying Remote Access Clients Using Connection Manager" at the Microsoft Windows Resource Kits Web site.

**Setting up the CMAK wizard**
No set up of the CMAK wizard is required once it is installed, but before running the CMAK wizard, make sure that you have read the information in the readme file (CMAKRead.htm) to ensure that you understand any restrictions on how to use the CMAK wizard to set up your service profile. For more information, see The Connection Manager Administration Kit Wizard.

**Editing the templates**
Before you use the CMAK wizard to build a service profile, locate the advanced customization information provided by the network administrator or other technical lead. Advanced customization of the templates must be done before you run the wizard. For more information on editing the template files, see Advanced Customization.

**Setting up a folder for custom elements**
All custom elements prepared by you and other members of your team must be accessible as you run the CMAK wizard. Create a folder in which to place all custom files to be incorporated into the service profile. Using a single location for custom elements makes it easier to build profiles when you run the CMAK wizard. You might want to set up a folder at a network location so your team can share a common location to save and access files as they develop the custom elements.

**Creating multiple customized versions**
To build multiple service profiles, you should create a separate worksheet for each profile to be built. To simplify building a service profile similar to one you already have, you can duplicate the existing service profile and edit it as appropriate.

**To build a new service profile from an existing profile**

1. In the **Service Profile Selection** pane, click **Existing profile**, and then select the appropriate existing profile from the drop-down list.

2. In the **Service and File Names** pane, specify a new service name and file name for the new profile.

3. Proceed through the wizard, changing any specific options as appropriate, to build the new service profile.

### Important

- The existing service profile is not affected, as long as you specify a different service name and file name for the service profile you are building. If you do not specify a new file name, the CMAK wizard will prompt you that you want to overwrite the existing service profile. If you specify a new file name but do not specify a new service name, users who have the existing service profile installed and then install the new service profile will experience problems in connecting using the service profile.

**Updating and editing existing service profiles**
When you create a service profile, the CMAK wizard copies all incorporated files into the CMAK folder for that service profile (Program Files\CMAK\Profiles*ServiceProfileFileName*). When you edit an existing profile, the wizard always pulls the incorporated files from this CMAK folder. To make changes to the incorporated files, you can either edit the incorporated file in this CMAK folder and then run the wizard again to update the service profile, or you can start the wizard, delete the referenced files, add the updated file by browsing to it, and then finish running the wizard.

### Notes

- When you edit an existing service profile created with a previous version of the Connection Manager Administration Kit,

the CMAK wizard automatically updates the profile. Any changes you have made manually to the .inf file (using advanced customization techniques) will not be retained in the upgraded file. However, the CMAK wizard creates a backup of the old .inf file before creating the updated file.

- When you use the CMAK wizard to edit a top-level service profile, the wizard attempts to use the latest version of the merged service profile. If the CMAK wizard cannot find the merged service profile in the \Program Files\CMAK\Profiles*ServiceProfileFileName* folder, it defaults to the current version of the merged service profile, originally stored with the service profile that you are editing.

# Using the CMAK wizard to build a service profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the CMAK wizard to build a service profile**
You create a service profile in three overall steps:

1. Start the Connection Manager Administration Kit (CMAK) wizard.
2. Run the CMAK wizard and enter all the necessary information.
3. Let the CMAK wizard build your customized service profile.

**Starting the CMAK wizard**
Do not start the CMAK wizard until you have collected all the information you need and have prepared all the custom bitmaps, Help files, and other items, as described in earlier sections.

**To start the CMAK wizard**

- To start the CMAK Wizard, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and double-click **Connection Manager Administration Kit**.

**Running the CMAK wizard**
Other sections in this documentation provide details on how to build a service profile, describing them in the order in which panes appear when you run the CMAK wizard. To help prepare for creating a service profile, you can complete the planning worksheet that is available in the Microsoft Windows Server 2003 *Resource Kit*. For more information about Connection Manager, see "Deploying Remote Access Clients Using Connection Manager" at the Microsoft Windows Resource Kits Web site. If you completed the planning worksheet, based on the options you select in the wizard, some panes might be skipped. These should correspond to nonselected items on the worksheet.

When specifying the file location of any custom element (such as the phone book, bitmaps, icons, and so on), type the complete path or browse to find the appropriate file.

After you enter information and select options on each pane, click **Next** to advance to the next pane. If you want to change information that you entered in previous panes, you can click **Back** as many times as necessary to find the settings you want to change. Make your changes, and then click **Next** to continue moving forward through the CMAK wizard.

If at any time you want more information about the information to be included in a particular pane of the CMAK wizard, click **Help** in the pane to open a topic that contains Help specific to the content of that pane.

**Finishing the CMAK wizard and building a service profile**
When you finish entering information and are ready for the CMAK wizard to build a service profile, click **Next**. After the CMAK wizard builds the service profile, a dialog box appears, indicating the folder and file name that were used to save the service profile.

### Caution

- After you click **Next** in **Ready to Build the Service Profile**, the **Back** button will be grayed out. To change the profile, you will have to run the CMAK wizard again.
- To perform advanced customization using the CMAK wizard, you must select the **Advanced Customization** check box before clicking **Next**. For more information, see Advanced Customization.

### Note

- After you finish, you can continue with manual advanced customization, including editing the .cms files and other service-profile files. For more information, see Advanced Customization. Be sure to rerun the CMAK wizard after completing advanced customization.

# Finishing customization after running the CMAK wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Finishing customization after running the CMAK wizard**

You can make most advanced customization changes within the **Advanced Customization** pane in the CMAK wizard. After you use the CMAK wizard to build a service profile, you can make further advanced customization changes to the service profile files by editing them with any plain-text editor, such as Notepad. You must run the CMAK wizard again after making these changes.

### Note

- The changes made to the .cms file for a service profile are specific to that service profile. To make a change that you want applied to all future service profiles, edit the template file. For more information on editing template files, see Templates and service profile files.

**To edit a .cms file for the service profile**

1. Locate the *ServiceProfileFileName*.cms file in the \Program Files\CMAK\Profiles*ServiceProfileFileName* folder and make backup copies of the file.

2. Use a text editor such as Notepad to edit the file. Be sure to save your changes before you close the file.

3. Restart the CMAK wizard, and then click **Next** to proceed to the **Service Profile Selection** pane.

4. Click **Existing profile** and then select the name of the service profile being edited.

5. Proceed through all other panes without making any changes, and then click **Finish** to rebuild your service profile. During the rebuilding process, the advanced customization that you did in the .cms file is applied to the service profile.

**To edit other files for the service profile**

Most advanced customization is done in the .cms file, but you can also edit the .inf and .cmp files. For more information, see Advanced Customization.

### Important

- If you do edit the .inf file, be extremely cautious.

If the service profile specifies it, edit these other files using the same procedures as specified above for the .cms files. These files are located in the same folder as the .cms file. To save time, edit all required files before rerunning the CMAK wizard to incorporate the changes.

# Phase Four: Preparing for Integration, Delivery, and Installation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phase four: Preparing for integration, delivery, and installation**
Connection Manager and your service profile can be delivered and installed on its own or integrated with another installation package, such as Internet Explorer. How you package and deliver your customized version will affect the way it is provided to your users. This section covers those options and considerations:

- Including Connection Manager in custom applications
- Incorporating Connection Manager with logon security
- Signing your code

# Including Connection Manager in custom applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Including Connection Manager in custom applications**
You can include Connection Manager and a service profile with your own custom installation package. The Connection Manager installation program, cmstp.exe, has several parameters that control how a service profile package is installed. You can also call Connection Manager programmatically from within a custom application.

### Note

- If you use the Internet Explorer Administration Kit (IEAK) to integrate a Connection Manager service profile with an Internet Explorer installation package, the Internet Explorer Customization wizard incorporates the command-line parameters needed to integrate the two installation programs.

**Delivering a custom installation package**
Using command-line parameters, you can integrate Connection Manager in an installation package and deliver it to users in a variety of ways. For example, you might want to:

- Distribute and install the Connection Manager service profile over a corporate network using Microsoft Systems Management Server (SMS) to automatically handle the process without user intervention.
- Integrate the Connection Manager service profile with another product's installation package and install them both in a single process.
- Post the Connection Manager service profile to a Web site to enable users to access and install the service profile in a standard way (specifying certain installation options for the user).

You can set up these distribution and installation options programmatically only after you have used the Connection Manager Administration Kit (CMAK) wizard to build your service profile.

**Syntax**
The command-line syntax for cmstp.exe follows one of two formats. One is for use with a service profile self-extracting executable (.exe) file, and the other is for use with a service profile information (.inf) file.

If you are building a custom installation or uninstallation package with the service profile executable file, you will want to use the following syntax:

*ServiceProfileFileName* **.exe** [**/q:a**] **/c:"cmstp.exe***ServiceProfileFileName***.inf** [*Parameters*]**"**

If you are building a custom installation or uninstallation package without the service profile executable file, you must use the following format in your custom package, and you must know the full path to the .inf file:

**cmstp.exe** [*Parameters*] **"***[drive][full path]ServiceProfileFileName***.inf"**

### Notes

- *ServiceProfileFileName* is the file name of the service profile you build with the CMAK wizard.
- The optional parameter, **/q:a**, prevents the profile from opening automatically after it has been installed. The verification message that the installation has succeeded will still appear.
- The command, **/c:"cmstp.exe**, tells the extraction program to override the installation defaults.
- Do not use brackets [] when specifying the full path to the .inf file.
- *Parameters* are the command-line parameters supported by the cmstp.exe installer.

- You can copy your service profile executable (.exe) file during installation and store it in a known directory. This action helps ensure that you will be able to refer to the file in the future.

**Parameters**

The following parameters are supported for installation:

- **/ni** - Do not create a desktop icon (valid only on Windows 95, Windows 98, Windows NT 4.0, and Windows Millennium Edition). On these platforms, the desktop icon is the only method provided by Connection Manager to enable users to open or delete a profile. If you use this parameter, you must provide your users with a different way to open or delete the service profile.

- **/ns** - Do not create a desktop shortcut to the connection in the **Network and Dial-up Connections** folder (Windows 2000 only) or the **Network Connections** folder (Windows XP and the Windows Server 2003 family only). Unless the silent install parameter is also specified, this only overrides the default user preference, so the user will still see the check box in the user interface. The first time a user installs a profile, the default is not to have a desktop shortcut. The user's preference is stored in the registry and is used for each subsequent installation.

- **/s** - Install silently. This option allows automatic installation of the service profile without requiring user action to complete the installation. The profile will not automatically start after it is installed.

- **/sd** - Set the installed profile as the default connection for Internet Explorer.

- **/su** - Install for a single user (valid only on Windows 2000, Windows XP, and the Windows Server 2003 family). If this parameter is used, the user does not have the option to install the software for use by all users.

- **/au** - Specifies that the service profile should be installed for all users. This parameter is only valid for computers running a member of the Windows Server 2003 family, Windows 2000, or Windows XP.

The parameters supported for uninstalling are as follows:

- **/u** - Uninstall the service profile. The only other parameter that can be used with this parameter is the **/s** parameter.

**Important**

- Any user can modify the .cms and .cmp file of a service profile that is installed on a computer running Windows 98 or Windows Millennium Edition. Users who install profiles for individual use on computers that are running Windows Server 2003, Windows XP, or Windows 2000 can modify .cms and .cmp files of the profiles they install. Only members of the Administrators or System Operators groups can modify .cms and .cmp files of profiles installed for All Users on computers that are running Windows Server 2003, Windows XP, or Windows 2000. Any member of the Administrators group on a local computer can modify the .cms and .cmp files of any service profile that is installed on that computer.

**Notes**

- To display Help for the command-line installation parameters, type **cmstp.exe /?**.

- To find the full path to the service profile information (.inf) file, you can use information extracted from registry keys, as shown in the examples below.

If you do not specify parameters for the cmstp.exe file, the following defaults are used:

- On a computer running Windows 95, Windows 98, Windows NT 4.0, or Windows Millennium Edition, a desktop icon is created, the profile is installed for all users, and the profile automatically starts after it is installed.

- On a computer running Windows 2000, the profile is installed for use by all users. If the Connection Manager software is included in the profile, that software is also installed.

- On a computer running Windows XP or a member of the Windows Server 2003 family, the profile is installed for use by all users, but Connection Manager is not installed.

- If a user is not logged on as an Administrator on a computer running Windows 2000, Windows XP, or a member of the Windows Server 2003 family, the software is installed for use by a single user.

#### Integrating the Connection Manager software within a custom application

You can start Connection Manager from within a custom application. A Connection Manager service profile presents a customized connection experience for your users. By integrating Connection Manager with a custom application, you can ensure that all your users connect in exactly the same way.

If you are writing an application for computers running Windows 2000, Windows XP, or a member of the Windows Server 2003 family and you want to open the Connection Manager logon screen, you can use the **RasDialDlg** function. If you want to open the **Properties** dialog box (for instance, if you want the users to select a phone number before dialing), you can use the **RasEntryDlg** function. For more information about these functions and how to use them, search for these function names on the Microsoft Web site or consult the Microsoft Platform Software Development Kit.

The **RasDialDlg** and **RasEntryDlg** functions are not supported for Windows 95, Windows 98, or Windows Millennium Edition. Although these functions are supported on Windows NT 4.0, using them in conjunction with a Connection Manager service profile is not supported. If you are writing an application for computers running one of these operating systems, you can instead use a function such as **ShellExecute** or **CreateProcess** to start cmmgr32.exe. To open the Connection Manager logon screen, use cmmgr32.exe, followed by the full path to the .cmp file of the service profile you want to open. For example, to open the Awesome Computers service profile on a computer running Windows 98, use:

**cmmgr32.exe "c:\Windows\All Users\Application Data\Microsoft\Network\Connections\Cm\Awesome1.cmp"**

If you want to open the **Properties** dialog box, use **cmmgr32.exe** followed by **/settings** followed by the full path to the .cmp file. For example, to open the **Properties** dialog box for the Awesome Computers service profile on a computer running Windows 98, use:

**cmmgr32.exe /settings "c:\Windows\All Users\Application Data\Microsoft\Network\Connections\Cm\Awesome1.cmp"**

### Notes

- Cmmgr32 runs as an asynchronous action. No response is returned to your program unless you add a custom action to the service profile to provide one. You can also wait for cmmgr32 to finish.

- If the path to the service profile includes spaces, use double quotes (") around the path.

#### Examples

The following examples describe a possible installation and a possible uninstallation for the Awesome Computers service profile.

#### Installation example for Awesome Computers

Awesome Computers wants to build a custom installation package and include a service profile, Awesome1.exe, in the package. The company wants a silent install. The company does not want to create a desktop icon or a desktop shortcut to the profile. The company would like the Awesome1 service profile to become the default connection for Internet Explorer.

In its custom installation package, Awesome Computers includes this command:

**Awesome1.exe /q:a c:"cmstp.exe Awesome1.inf /s /ni /ns /sd"**

#### Uninstallation example for Awesome Computers

Awesome Computers wants to build a custom uninstallation package. The company does not know whether its users still have the Awesome1.exe file on their computers, so the company must get the .cmp path from the registry and extrapolate the .inf path from that information. The company wants a silent uninstall.

1. The custom uninstallation application queries one of two registry keys.

   All users profile:

   **HKLM\Software\Microsoft\Connection Manager\Mappings**

Single user profile:

**HKCU\Software\Microsoft\Connection Manager\Mappings**

The call provides a name value of **Awesome Computers** (the full service profile name) and returns a data value of the full path to the Awesome1.cmp file.

2. The custom uninstall application changes the Awesome1.cmp path into the Awesome1.inf path by replacing Awesome1.cmp with Awesome1\Awesome1.inf. For example, a .cmp path of C:\Documents and Settings\All Users\Application Data\Microsoft\Network\Connections\Cm\Awesome1.cmp must be converted to C:\Documents and Settings\All Users\Application Data\Microsoft\Network\Connections\Cm\Awesome\Awesome1.inf.

3. The custom uninstallation application includes this command:

   **cmstp.exe /u /s "c:\Documents and Settings\All Users\Application Data\Microsoft\Network\Connections\Cm\Awesome\Awesome1.inf"**

### Notes

- The credentials of the user installing or uninstalling the package will affect how cmstp.exe operates. If the user does not have administrative credentials, he or she will not be able to install or uninstall a profile for all users. The profile will be installed for his or her own use.

- If you use the following key, the initial portion of the data value might contain a system macro.
  **HKCU\Software\Microsoft\Connection Manager\Mappings**

For more information, see Verifying customization and Troubleshooting process.

# Incorporating Connection Manager with logon security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Incorporating Connection Manager with logon security**
On computers running Windows 2000, Windows XP, or a member of the Windows Server 2003 family, a user can choose to log on to Windows using a Connection Manager profile.

**Logging on using a Connection Manager profile**
A user can log on using a dial-up connection (including a Connection Manager connection) only if the computer is a member of a domain. If the computer is not a member of a domain, the **Log on using dial-up connection** check box does not appear.

In the **Log On to Windows** dialog box, the user can select the **Log on using dial-up connection** check box. After clicking **OK**, the user is prompted to choose a network connection.

After the user chooses a connection, the Connection Manager logon screen appears. Connection Manager starts the connection process when the user clicks **Connect**.

**Administrative issues when using Connection Manager profiles to log on**
If you intend to build and distribute a service profile that will be regularly used to log on, you might want to consider the following:

- By default, Connection Manager uses the credentials the user types at the Windows logon screen. If you want Connection Manager to ignore those credentials, add the key **UseWinLogonCredentials=0** to the **[Connection Manager]** section of the .cms file for your service profile.

- Users who select the **Log on using dial-up connection** check box gain access to Windows through an account designed for this method of logging on. Any customized settings for Connection Manager are applied to the Log on to Windows account, rather than to the account of the user who logged on. The user can customize the Log on to Windows account in the same ways as any other account (for example, save settings for phone numbers). However, permissions granted to specific users, such as access to Help, are not granted to users logging on to Windows, even after the connection completes.

For security reasons, the following functionality is not available when a user logs on to Windows using a dial-up connection:

- If the .cms file for the profile contains the **ResetPassword** key, any value for that key is ignored.

- Custom buttons do not appear or function.

- The **View Log** button is disabled.

- The **Advanced** tab does not appear in the properties dialog box for the service profile.

- Keys in the .cmp file that provide user credentials for first-time use (such as **UserName**) are not available. Internet connection credential keys such as **InternetUserName** are available.

- Macromedia animation does not display. If you are building a profile that uses Macromedia animation, you should include a default bitmap to display in this circumstance.

- Custom actions do not run unless those actions have been enabled by editing the registry. Even after custom actions have been enabled, the custom action will not automatically work with registered shell extensions. You must indicate the program to use with the file, as well as the path to the file.

- Connection Manager Help (including context-sensitive help) is disabled during the connection process.

- If a system component is missing when the user attempts to log on using a dial-up connection, the connection attempt fails. The user will not be able to install the system component until he or she logs on.

**Enabling custom actions that run when a user logs on by using a dial-up connection**

For a custom action to run during the logon process, you must enable that action by specifying a value for the following registry key:

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Connection Manager** *ProfileName* **\WinLogon Actions**

The format for the value of the registry key is:

| Field | Description |
|---|---|
| Name | The name of the executable file that will run. |
| Type | REG_DWORD |
| Data | A value indicating the location of the executable. |

Supported values for the Data field are:

| Value | Location of executable |
|---|---|
| 0x00000000(0) | *%windir%*\system32 |
| 0x00000001(1) | Profile directory |

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

The following example enables automatic phone book downloads at connection time. If you set the value of:

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Connection Manager** *Profile Name* **\WinLogon Actions**

as:

**cmdl32.exe REG_DWORD 0x00000000(0)**

then you must also include:

**[Xnstall.AddReg.AllUsers] "HKLM", "%AppAct%\%ServiceName%\WinLogon Actions","cmdl32.exe",0x00010001,0**

in the .inf file of the profile.

**Important**

- Modifying .inf files can cause significant installation problems. Thoroughly test all service profiles containing custom .inf entries for version, file, or other conflicts. Be sure that the file location is specified correctly.

- If a custom action runs when users log on to Windows, the custom action runs with system permissions. You should ensure that any custom actions that you enable do not pose unintended security risks.

- Any user can modify the .cms and .cmp files of a service profile that is installed on a computer running Windows 98 or Windows Millennium Edition. Users who install profiles for individual use on computers that are running Windows Server 2003, Windows XP, or Windows 2000 can modify .cms and .cmp files of the profiles they install. Only members of the Administrators or System Operators groups can modify .cms and .cmp files of profiles installed for All Users on computers that are running Windows Server 2003, Windows XP, or Windows 2000. All members of the Administrators group on a local computer can modify the .cms and .cmp files of any service profile that is installed on that computer.

# Signing your code

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Signing your code**

After you have completed all customization, it is recommended that you sign your code before distributing it. Signing your code with a digital signature assures your users that the programs are from a known source and have not been altered.

You should digitally sign your customized service profile and any custom programs to ensure that users do not receive warnings when they install Connection Manager. Users may be prevented from installing ActiveX controls and other packages that are not signed.

Further information on code signing can be found in the Internet Explorer Administration Kit (IEAK). If you include a Connection Manager service profile in an Internet Explorer installation package created using Internet Explorer Connection wizard in the IEAK and you specify the digital signature option in the wizard, the digital signature support provided for the installation package will include the Connection Manager service profile.

For information on code signing, go to the Microsoft Developer Network (MSDN) on the Microsoft Web site, click **Search** on the MSDN screen, and type **digital signature** or **code signing** as search words for a Boolean search on the entire Microsoft site.

# Phase Five: Testing Your Deliverables

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phase five: Testing your deliverables**
Before you distribute the service profile that you created with the CMAK wizard, you must test it. This section describes how to install and test service profiles. It also describes things you should consider if you want to provide users with your own installation program to install service profiles.

This section covers:

- Test process
- Verifying customization

# Test process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Test process**

Test your service profiles to ensure that all planned functionality is implemented and working as planned.

Install and use each of your service profiles on each target operating system. If you distribute the service profile on disk and by e-mail, test each method. If you are providing multiple language versions, test each one on the operating system with the appropriate national language.

**Note**

- If you have users with previous versions of your service profile, test Connection Manager on computers on which your previous service profiles are installed, as well as on computers that have not had previous versions installed.

Following are basic testing steps:

For each service profile and scenario, use Connection Manager to do all of the following:

**To test your service profile**

1. Install the service profile, as described in a readme or other documentation that you provide to the user.
2. Provide the required user data (user name, user password, and so on) needed to connect.
3. Select a Point of Presence (POP) from the phone book and establish a connection.
4. View network resources while connected to your service (check connection while viewing URL).
5. Disconnect from your service.
6. Uninstall the service profile.

**Test criteria**

As you install and use Connection Manager, ensure that each service profile has the appearance, features, functionality, and usability that you want:

- Do the screens contain the appropriate titles, text, and graphics, and does each screen appear completely and correctly?
- Are all of the required programs included in the service profile? If not, are all required programs available on all target operating systems?
- Do all menu items for the notification area shortcut function appropriately?
- Do you have appropriate Help topics, including troubleshooting Help, for all functions and features? Do these Help files match the service profile? If you customized any Help topics, do they appear as expected?
- Is everything formatted and spelled correctly, including the license agreement, the screen text, the text in the Connection Manager dialog box, and other user information?
- Is all the information needed to install and use Connection Manager and the service profile available when and where it is needed?
- If you provided custom dial-up entries (especially scripts), do they work as intended?

- If you support such connections, can you establish both direct and VPN connections?
- If you included any pre-connect, pre-tunnel, post-connect, pre-dial, on-error, on-cancel, or disconnect custom actions or monitored applications, do they work as expected?
- If you edited your .cms and other service-profile files, do your custom entries perform as expected?
- If you use code signing, does it display the authentication notice as expected?
- Can you connect using POPs in the phone book for each of the service types you support?
- If you use a firewall, are the firewall proxy settings correct?
- Do the components of Connection Point Services provide the required functionality, as described in the Connection Point Services documentation?

**Note**

- You must build profiles on a server operating system that is designed for the same processor set as the client computers on which users will install the profile. For example, if you want to build a profile for users running Windows XP Professional x64 Edition, you must build that profile on a server that is running an x64-based version of Windows Server 2003.

# Verifying customization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Verifying customization**
There are several instances in which customization requirements are very stringent. This section repeats the information that is most commonly overlooked.

**Verify that the service name and file name are valid**
If you are updating an existing service profile, the service and file names should be identical to the previous ones. If you are creating a new profile (from scratch or by editing an existing service profile and renaming it), make sure that the service name and file name are not identical to any other service profiles that you provide to your customers. If you do not specify a file name that is different from the one users already have installed on their computers, the original service profile will be overwritten. If you specify a new file name but do not specify a new service name, users who have the existing service profile installed and then install the new service profile with the different service name will experience problems in connecting. The problems can include not being able to connect using either or both service profiles.

If you need to make changes to an existing profile, use the Connection Manager Administration Kit (CMAK) wizard to edit the profile instead of creating a new file with the same name.

**Verify that the dial-up entry is correct in CMAK and the phone book**
The phone book dial-up entry that you set up in the phone book must exactly match the entry you set up in your service profile. Use a text editor such as Notepad to open the Phone Book Administrator file and verify that the entries you specified in the phone book (at the end of each line) match the entries you set up in the CMAK wizard in **Dial-up Networking Entry**.

**Verify that all required protocols are specified correctly**
Verify that the required network protocols, including TCP/IP, IPX, and NetBEUI (as appropriate), are specified in the **[Server&*EntryName*]** section of the .cms file for the specified dial-up entry. The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

**Verify that proxy settings are set as required to support successful connection to your service**
If your service requires special proxy settings, verify that any required custom actions have been designed and implemented to adjust the settings on the user's computer as required.

**Verify that resident programs on users' computers can be opened**
If your custom actions or other custom settings depend on the program being available on your users' computers (that is, you are not providing it in your service profile), verify that the program is in the default path and available to Connection Manager. Not all programs that appear on the **Start** menu are necessarily set up to use the default path. Connection Manager may not be able to find the installed program if it is not in the default path.

**Verify that your service-profile updates are complete**
When you edit a service profile created using a previous version of Connection Manager, the CMAK wizard automatically upgrades the profiles to Connection Manager 1.3. Any changes you have made manually to the .inf file (using advanced customization techniques) will not be retained in the upgraded file. However, the CMAK wizard creates a backup copy of the old .inf file before creating the updated file.

**Verify that your time-outs are not preventing access to your domain**
For VPN connections, users sometimes get a message that the domain controller couldn't be found, even though the routing information is correct. In this case, setting the LanaTimeout entry in the .cms file may correct the problem. This increases the time (in seconds) to wait before attempting to connect to a VPN server. Increase this value if the connection consistently fails to find a domain controller.

# Phase Six: Supporting Your Customers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phase six: Supporting your customers**

Effective customer support depends on your ability to resolve customer questions, problems and needs quickly and effectively. This section contains information intended to help you and your users address the more common situations:

- Troubleshooting process
- Frequently Asked Questions

# Troubleshooting process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting process**

The first level of troubleshooting is provided in the Help files that are included with your service profile. This troubleshooting Help information makes it possible for many problems to be resolved directly for the user. Each problem that cannot be resolved leads to a Help topic that instructs the user to call customer support. You should customize this Help topic with a phone number for your customer support group, additional information specific to your service, and other information that may help your users.

To resolve problems most effectively:

1. Make sure that all system requirements have been met.

2. Make sure that the right access numbers are used and that they are entered appropriately. Verify that the users understand that dialing rules are applied only to numbers selected from a phone book. If a number is typed or edited by the user, the user must then type all dialing information, including information such as long-distance and external access numbers. If the user chooses to use dialing rules, the user must configure those dialing rules appropriately for the location.

3. Make sure that users know how to access the Connection Manager troubleshooter Help, and verify the path that they have taken through the troubleshooter Help.

4. Check the appropriate readme files (including CMAKRead.htm) to determine if the problem is a known issue.

5. Look for common mistakes. Make sure that:

    - The modem is on.

    - The appropriate version of Connection Manager is installed. To determine this, instruct your users to check the About tab of the **Properties** dialog box (Windows NT 4.0, Windows 95, Windows 98, and Windows Millennium Edition only) or the log file.

    - The latest version of the appropriate connection software is installed.

    - The latest service pack is installed.

    - The Internet Explorer proxy settings are correct.

    - All required network protocols are installed and configured correctly on the user's computer. (If the user can get to the network but cannot reach any services, a protocol on the user's computer may be misconfigured.)

    - If your service profile supports VPN connections, Point-to-Point Tunneling Protocol (PPTP) is installed and configured correctly and completely. (This is usually a problem only on Windows NT 4.0, where Connection Manager does not automatically install PPTP.)

    - The dial-up entry used is appropriate for the user’s operating system.

    - If a user has installed remote access or PPTP on a computer running Windows NT 4.0, the user reinstalled any service pack previously installed on their computer.

    - If a user is using Internet Connection Sharing, Internet Connection Firewall, or Windows Firewall, the connection is shared correctly, and the firewall is not blocking the connection. For more information, see Connecting to the Internet in a home or small office network, Internet Connection Firewall, and Help: Windows Firewall.

**Notes**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

6. Users have received all workarounds and solutions, as appropriate. For information on networking problems, including those involving remote access, search the Microsoft Developer Network (MSDN) on the Microsoft Web site.

7. Ask users to generate and send a log file for the connection. The log file contains valuable troubleshooting information. Logging is typically enabled by default. Users can enable or disable logging after the profile is installed.

   - If logging was enabled, instruct the user to send the log file. Make sure that the user did not turn off logging. Depending on the size of the log file, decide whether you want the user to clear the log file and try the connection again before sending the log file. To direct the user to the log file, instruct the user to click **Properties** on the Connection Manager logon screen. Instruct the user to click the **Options** tab and then to click the **View Log** button. After the user has opened the log file, instruct the user to save the file and send it to you.
   - If logging was not enabled, instruct the user to turn on logging, try the connection, and then send the log file. To enable logging, instruct the user to click **Properties** on the Connection Manager logon screen. Instruct the user to click the **Options** tab and to select the **Enable Logging** check box.
   - If the user is attempting to log on to Windows using a Connection Manager profile, the **View Log** button will be unavailable. To find the log file, instruct the user to log on to the computer without using a Connection Manager profile. After the user has logged on, instruct the user to click **Start**, point to **Search**, and click **For Files or Folders**. After the **Search Results** dialog box appears, instruct the user to click **Search Options**, select the **Advanced Options** check box and to select the **Search system folders** check box. Instruct the user to search for the log file by typing in the first three letters of the service profile name, followed by an asterisk (*) and the log file extension (.log). For example, if the service profile is named **Awesome**, the user should search for **Awe*.log**. The file is stored in %windir%\temp by default.

**Note**

- On computers running Windows NT 4.0, Windows 2000, Windows XP, or a member of the Windows Server 2003 family, the log file will be in Unicode format. To view the log file on an operating system other than these, open the log file with a Unicode text viewer, such as Microsoft Word.

For more information about Connection Manager, see "Connection Manager" at the Microsoft Windows Resource Kits Web site.

# Frequently Asked Questions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Frequently Asked Questions**
The following are some of the most frequently asked questions regarding Connection Manager and the Connection Manager Administration Kit (CMAK). These questions are based on actual calls to Microsoft product support services and might save you time and effort.


Unexpected profile actions

- My users can connect to the server when using the service profile, but authentication fails.
- My VPN service profile takes too long to connect.
- The connection failed to complete after the user was authenticated.
- My ISDN users connect at 64 kilobits per second.
- My merged service profile does not work the way I expect.
- The Connection Manager interface appears unexpectedly.
- Nothing happens when my Windows 2000 users double-click the connection icon in the Network and Dial-up Connections folder.
- My service profile does not work properly when users log on to Windows with it.

Profile administration

- The phone book for the service profile is not being updated.
- I cannot populate a previously installed service profile with user credentials and phone numbers.
- I want to dynamically change the VPN address for a service profile.

Profile creation and development

- I do not understand the difference between networking settings and dial-up entry names.
- I want to create a profile in which some numbers dial directly to my corporate network and other numbers connect through a virtual private network (VPN) tunnel.
- I want to create a VPN-only profile.
- I want to provide user credentials and phone numbers for my service profile for first-time users.
- I want to create Connection Manager profiles in more than one language.


**Unexpected profile actions**
**My users can connect to the server when using the service profile, but authentication fails.**
**Cause:** The server running Routing and Remote Access cannot connect to the authentication server.

**Solution:** Verify connectivity between the server running Routing and Remote Access and the authentication server.

**Cause:** The authentication server has been improperly specified or configured for dial-up access.

**Solution:** Verify that the authentication server specified in the service profile is correct. Check the user account and make sure that dial-up access has been enabled. Look up any Group Policy rules that might restrict access by time or date.

**Cause:** Settings for one or more dial-up entries have been improperly specified.

**Solution:** Check the security settings for each dial-up entry, including component service profiles. The simplest way to do this is to run the CMAK wizard and choose to edit the service profile. When you reach **Dial-up Networking Entries**, click each entry, and choose **Edit**. On the **Security** tab, make sure that all of the security settings are properly specified. One common error is to select **Use advanced security settings** for a service profile issued to users using Windows 98.

**Solution:** Make sure that the phone book files (.pbk) have the correct settings for each dial-up entry. For each phone book you include in the service profile, open the .pbk file, and check the **Dial-up Networking entry** field, which is the eleventh field, in each POP. Make sure that this field is specified appropriately for each POP in each phone book.

**See also:** Advanced Customization; Add, edit, or delete POPs by command line

**Cause:** Prefix and suffix information might have been specified differently in one or more of the component profiles merged into the issued service profile.

**Solution:** Check the **UserNamePrefix** and **UserNameSuffix** keys in every .cms file in the service profile. A value for this key specified in the component profile will override the top-level service profile settings.

**See also:** Merging phone books and other features from existing service profiles; Providing a realm name; Advanced Customization

**My VPN service profile takes too long to connect.**
**Cause:** Certificates are invalid or missing.

**Solution:** Delays in excess of one minute often indicate a failure in the L2TP/IPSec protocol negotiation. Make sure all service profile users have the newest certificates installed on their computers.

**See also:** Certificate Services

**Cause:** A round robin server solution is causing delays in authentication.

**Solution:** If you are using a round robin server solution for L2TP/IPSec VPN servers, a failure to authenticate on one or more servers can result in variable delays. Depending on the profile settings, Connection Manager might redial each VPN server multiple times. Make sure that each VPN server is configured to accept the profile and that each server is responding.

**See also:** Configuring round robin

**The connection failed to complete after the user was authenticated.**
**Cause:** The service profile contains improper post-connect actions.

**Solution:** Check the **[Connect Actions]** section in the service profile's .cms files for proper behavior on target platforms.

**See also:** Incorporating custom actions

**Cause:** The routing table update file is attempting to delete a route that does not exist.

**Solution:** Check the routing table update file for outdated or inaccurate route changes.

**See also:** Including routing table updates

**My ISDN users connect at 64 kilobits per second.**
**Cause:** The **ISDNDialMode** key value is not set to bind dual channels.

**Solution:** Use advanced customization techniques to change the value of the **ISDNDialMode** key to the appropriate value and reissue the profile.

**See also:** Advanced Customization; Connection Manager keys

**My merged service profile does not work the way I expect.**
**Cause:** The settings in component service profiles conflict.

**Solution:** Check phone book filtering. Your region settings might not be the same in all phone book region (.pbr) files. If you

are using custom service types in any or all phone books, you need to apply these to each .cms file.

**Solution:** Check the realm name information in each .cms file included in the merged service profile. Any realm information in a component profile will be overridden by realm information specified in the .cms file of the top-level profile. You should leave the realm name field in the top-level service profile blank if you need to merge profiles with different realm information.

**Solution:** Check the tunnel addressing information. Tunnel addressing information specified for the top-level service profile is used for all phone books merged into it.

**Solution:** Check the settings for the **DUN** and **TUNNELDUN** keys in all the .cms files included in the merged service profile. The default dial-up settings originate in the .cms file associated with the user-selected phone book, while VPN settings originate from the top-level .cms file.

**See also:** Advanced Customization; Providing a realm name; Implementing VPN support; Merging phone books and other features from existing service profiles

**The Connection Manager interface appears unexpectedly.**
**Cause:** Internet Explorer is improperly configured for your service.

**Solution:** Assist your users in changing their Internet Explorer options. In Internet Explorer, click the **Tools** menu. Click **Internet Options**, click the **Connections** tab, and choose the appropriate dial-up option. The default option is to always dial the default connection. However, your profile might work best if a connection is never dialed or if a connection is dialed only when an Internet-based program is started and a network connection is not present.

**Cause:** Custom actions requiring Internet access continue to run after the connection has been terminated.

**Solution:** Test the service profile to make sure that all custom actions terminate properly. If necessary, reissue the service profile after you add a disconnect action that ensures that programs terminate properly.

**See also:** Incorporating custom actions

**Nothing happens when my Windows 2000 users double-click the connection icon in the Network and Dial-up Connections folder.**
**Cause:** The installation of the Connection Manager 1.3 software did not complete correctly. Mismatched Connection Manager software binaries are causing the connection software to fail.

**Solution:** Instruct the user to uninstall the Connection Manager 1.3 software and then to reinstall the profile. To uninstall Connection Manager 1.3, click **Start**, point to **Settings**, click **Control Panel**, double-click **Add or Remove Programs**, click **Microsoft Connection Manager 1.3**, and click **Remove**. This process restores the Windows 2000 Connection Manager binaries on the user's computer. The user should then reinstall the Connection Manager service profile that included the Connection Manager 1.3 software.

**See also:** Troubleshooting process; Adding Connection Manager 1.3 to the service profile

**My service profile does not work properly when users log on to Windows with it.**
**Cause:** Custom actions have not been enabled to run when users log on to Windows with your service profile.

**Solution:** Check whether you need to change user registry settings in order to run custom actions when users log on to Windows with your profile. Consider issuing a profile without custom actions for use in logging on to Windows.

**See also:** Incorporating Connection Manager with logon security

**Cause:** Connection Manager treats "Log on using dial-up connection" as a separate user with limited permissions. Settings applied to "Log on using dial-up connection," such as proxy configuration, are not applied to individual user accounts after logging on.

**Solution:** Consider advising your users to log on to their local computers before using a Connection Manager profile to log on to your domain.

**See also:** Incorporating Connection Manager with logon security

**Profile administration**
**The phone book for the service profile is not being updated.**
**Cause:** Improperly set Internet proxy settings is preventing phone book downloads.

**Solution:** Check the proxy settings for the service profile. If you are using automatic proxy configuration in your service profile, check those settings and the proxy server.

**See also:** Common problems and their solutions; Using automatic proxy configuration

**Cause:** One or more Phone Book Service (PBS) servers is out of service.

**Solution:** Check the log files on your PBS servers for error messages.

**See also:** Phone Book Service error messages

**Cause:** A URL in the service profile points to a server that cannot be found.

**Solution:** Check the URLs specified in the service profile. You can do this by editing the profile in the CMAK wizard or by checking the settings directly in the **[ISP]** section of the .cms file. If the profile in question was created by merging other service profiles, check the URLs in all the .cms files used by the profile.

**See also:** Advanced Customization; Providing phone book support; Merging phone books and other features from existing service profiles

**I cannot populate a previously installed service profile with user credentials and phone numbers.**
**Cause:** Updated profiles cannot overwrite user settings.

**Solution:** Populate the credentials in the .cmp file of the service profile when you create it. Run the CMAK wizard again. Issue the new service profile to your users. Include a custom uninstallation and reinstallation package or instructions on how to uninstall the old service profile and install the new version.

**I want to dynamically change the VPN address for a service profile.**
**Solution:** Write a custom pre-tunnel action that adjusts the VPN address in the .cms file as needed. This feature might not be available in future releases. Consider the security implications of dynamically changing a VPN address before you implement this solution.

**See also:** Incorporating custom actions

**Profile creation and development**
**I do not understand the difference between networking settings and dial-up entry names.**
**Solution:** Dial-up entry names identify combinations of networking settings. You can specify different networking settings for each dial-up entry. You usually name dial-up entries on the **Settings** tab of the **Add POP** pane of Phone Book Administrator. Those names appear in the **Dial-up Networking Entries** pane of the CMAK wizard. Connection Manager uses a dial-up entry to determine which combination of networking settings are required for a service profile.

**See also:** Incorporating custom dial-up entries

**I want to create a profile in which some numbers dial directly to my corporate network and other numbers connect through a virtual private network (VPN) tunnel.**
**Solution:** Create two phone books, one for a direct dial connection and the other for a VPN tunnel to your network. Create a profile with your direct dial phone book. Create a top-level profile that merges the direct dial profile and incorporates the VPN phone book. On the **VPN Support** pane, select the **Phone book from this profile** check box.

**See also:** Merging phone books and other features from existing service profiles; Implementing VPN support

**I want to create a VPN-only profile.**
**Solution:** To create a VPN-only profile, start the CMAK wizard. Select the **Phone book from this profile** check box on the **VPN Support** pane, and specify a VPN server or a VPN file. Configure the VPN entry or entries with the correct security and addressing information for your network. Do not specify a phone book file on the **Phone Book** pane, and clear the **Automatically download phone book updates**check box. You do not need to configure the default dial-up entry for the profile. On the last pane of the wizard, select the **Advanced Customization** check box. On the **Advanced Customization** pane, click the profile .cms file, click the **[Connection Manager]** section, and set the value of the **Dialup** key to zero. Click **Apply**, and finish the wizard. This will create a VPN-only profile, without a **General** tab in the **Properties** dialog box for the profile. Users of your profile will not see any phone or dialing information.

**See also:** Implementing VPN support; Incorporating VPN entries

**I want to provide user credentials and phone numbers for my service profile for first-time users.**
**Solution:** Use advanced customization techniques to assign values to the keys for which you want to provide first-time use data, such as **Username**. These keys must be assigned in the .cmp file, and they will only be available the first time the profile is used. For passwords, you must edit three keys: **PCS**, **RememberPassword**, and **Password**. The **PCS** key should always be set to 0; the values of the other keys will vary according to your profile needs. If you want this password to be available to all users who connect from the same computer, you must edit an additional key, **KeepDefaultCredentials**. If this key is not set to the appropriate value, only the person who installs this profile will have access to the credentials you provide for first-time use. Consider the security implications of providing a password for the first user who installs and uses the profile.

**See also:** Advanced Customization

**I want to create Connection Manager profiles in more than one language.**
**Solution:** You can use the Multi-User Language Interface (MUI) with the CMAK wizard to build a Connection Manager profile in a language other than the one installed with your operating system. You can include the Connection Manager software with that profile, but Connection Manager will appear in the language installed with your operating system, not in the MUI language. You can choose not to include the Connection Manager software with the profile, but your users might not have access to all the current features of Connection Manager.

To include a version of Connection Manager in the same language as the profile, you must build the profile using the CMAK wizard on an operating system that was installed with the appropriate language. The CMAK wizard comes with Windows Server 2003 operating systems. You can also install the CMAK wizard on Windows XP Professional with the Windows Server 2003 Administration Tools Pack.

**See also:** Creating profiles in multiple languages

# Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Point Services**
Connection Point Services (CPS) works in conjunction with Connection Manager to automate the process of updating your users' computers with new Points of Presence (POP). CPS consists of Phone Book Service (a utility for distributing your phone books) and Phone Book Administrator (a tool to create and maintain your phone book files). For complete information about Connection Manager, see Connection Manager Administration Kit.

- Before installing, configuring, or administering CPS, see Connection Point Services Checklists.
- For helpful tips, see Connection Point Services Best Practices.
- For help with specific tasks, see Connection Point Services How To....
- For general information, see Connection Point Services Concepts.
- For problem-solving strategies and error messages, see Connection Point Services Troubleshooting.

# Connection Point Services Checklists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

- Checklist: Installing Connection Point Services
- Checklist: Configuring Connection Point Services
- Checklist: Administering phone books with Phone Book Administrator
- Checklist: Administering phone books by command line

# Checklist: Installing Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing Connection Point Services**

While installing a Windows Server 2003 family operating system, you might have chosen to add Connection Manager components, which include Connection Point Services (CPS). If you did not install CPS at that time, you can install it now.

| Step | Reference |
| --- | --- |
| Connection Point Services (CPS): what it is and how it works. | Understanding Connection Point Services |
| Important information to prepare you for installing CPS. | Install Connection Point Services |
| Install CPS | Install and Configure Connection Point Services |

# Checklist: Configuring Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring Connection Point Services**

When you configure Connection Point Services (CPS), you provide security by setting permissions for directory access.

| Step | Reference |
|---|---|
| Connection Point Services (CPS): what it is and how it works. | Understanding Connection Point Services |
| Set up an account with administrative credentials on your Phone Book Service (PBS) host. | Manage Users; Active Directory |
| Add a user account on your PBS server for posting phone books, and make that account a member of the Guests group. Disable this user account from logging on locally and from logging on through Terminal Services. | Set permissions early; Create a local user account; Add a member to a local group; Deny log on locally; Deny log on through Terminal Services |
| Configure the FTP service for specific users for use when posting phone book data. | Set up FTP accounts for known users |
| Set permissions for the PBS folder (Windows NTFS file system only). | Set permissions for the Phone Book Service folder (NTFS only) |
| Set Write permissions for FTP. | Set the Write permission for the FTP virtual directory |

# Checklist: Administering phone books with Phone Book Administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Administering phone books with Phone Book Administrator**

Connection Point Services (CPS) provides two ways to administer phone books: by using the Phone Book Administrator (PBA) interface and by using the command line. This checklist can help guide you through administration with Phone Book Administrator. You are advised to create and publish your first phone book using the PBA interface, even if you plan to use the command line for your usual publishing procedures.

| Step | Reference |
|---|---|
| Connection Point Services (CPS): what it is and how it works. | Understanding Connection Point Services |
| Create a phone book. | Add and configure a new phone book |
| To add, edit, delete, or import regions. | Add, edit, or delete a region<br><br>Import a region list |
| Add POPs to your phone book. | Add or edit a POP |
| Publish the phone book. | Publish a phone book |

# Checklist: Administering phone books by command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Administering phone books by command line**

Connection Point Services (CPS) provides two ways to administer phone books: by using the Phone Book Administrator (PBA) interface and by using the command line. This checklist can help guide you through administration by command line. However, you are advised to create and publish your first phone book by using the PBA interface, even if you plan to use the command line for your usual publishing procedures.

| Step | Reference |
|---|---|
| Connection Point Services (CPS): what it is and how it works | Understanding Connection Point Services |
| Create a phone book and then set or edit the phone book options. | Create and Manage Phone Books by Command Line |
| Import a region file | Add, edit, or delete regions by command line<br><br>Import a region list by command line |
| Import a POP file | Add, edit, or delete POPs by command line<br><br>Import an existing POP list by command line |
| Publish the phone book | Publish a phone book by command line |

# Connection Point Services Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

The following list provides recommended practices for configuring and administering Connection Point Services.

- Specify POP settings (sometimes called *service types*) so users can search for the POPs that suit their needs and filter out the ones they do not want. For more information, see Make the most of POP settings.
- Set up permissions before attempting to publish phone books to Phone Book Service (PBS). For more information, see Set permissions early.
- Secure the Phone Book Administrator (PBA) and PBS folders to avoid synchronization problems and to prevent unauthorized personnel from posting to the host server. For more information, see Safeguard your files.
- Back up phone book databases and other important files to preserve POP and region information. For more information, see Back up your data.

# Make the most of POP settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Make the most of POP settings**

POP settings (sometimes called *service types*) provide you with the ability to specify the way users are presented with POP choices. For example, you can set some POPs to appear only when a user is dialing with a modem. Users can then use the differences between POP connections to search for the POPs that suit their needs and filter out the ones they do not want. The more you can differentiate between POPs, the less a user will need to search.

The following table shows available POP settings and what they do. The first column lists options for each POP setting, with default options in bold type. You can set these options when you add or edit POPs, whether you are using the Phone Book Administrator (PBA) graphical interface or command-line interface.

| POP setting | Action determined for the POP |
|---|---|
| **Sign on** | Users can connect to the POP to sign on to your service. |
| Sign up | Users can connect to the POP to sign up for your service. |
| **Modem** | Users can connect to the POP by modem. |
| **ISDN** | Users can connect to the POP using an ISDN connection device. |
| Custom 1 | The POP is assigned a custom service setting. You can specify custom service settings in the Connection Manager service profile. |
| **Multicast** | The POP supports IP multicast protocols. |
| Surcharge | The access number for the POP appears in the **More access numbers** box in the Connection Manager **Phone Book** dialog box. You can charge your users an additional fee when they use this POP. |
| Custom 2 | The POP is assigned a custom service setting. You can specify custom service settings in the Connection Manager service profile. |

# Set permissions early

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set permissions early**

Before you can publish phone books to a Phone Book Service (PBS) server, you must create a specific user account that has the necessary permissions to post your phone books. The NTFS file system should be running on the server host.

Failure to set permissions is likely to prevent you from posting phone books to the server.

- **Set up a user account on your PBS server for posting phone books.** Set permissions for the user account so that users have the minimum necessary permissions to post to your server. This user account should be a local user account on the PBS server, not on a domain, even if the PBS server is a member of a domain. For additional security, you should make this account a member of the Guests group, and you should disable this user account from logging on locally and from logging on through Terminal Services. To post phone books, the user account must have a password, and the password cannot be null. The user account must also have Write permission to the PBSData folder in the FTP virtual directory. For more information, see Create a local user account, Set the Write permission for the FTP virtual directory, Add a member to a local group, Deny log on locally, and Deny log on through Terminal Services.

- **Configure FTP service for specific users when posting phone book data**. When you install PBS, FTP is also installed with anonymous FTP access disabled. Because passwords for FTP sessions are sent as plaintext, you should not use the user account that you use for posting phone book data for any other purpose. Do not activate this account until you are ready to post data, and disable this account as soon as you finish posting data. For more information, see Set up FTP accounts for known users and Disable or activate a local user account.

- **Set permissions for the Phone Book Service folder** (optional but recommended; NTFS file system only). For added security, you should change access to the Phone Book Service folder from all authenticated users to specific user accounts or groups. For more information, see Set permissions for the Phone Book Service folder (NTFS only) and Security information for Connection Point Services.

- **Set the Write permission for the PBSData folder in the FTP virtual directory.** Before you can post, you must set this permission on the FTP virtual root. For enhanced security, set this permission only when you are about to post to the server, and then clear it immediately after posting. For more information, see Set the Write permission for the FTP virtual directory.

# Safeguard your files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Safeguard your files**
To avoid phone book synchronization problems and to prevent unauthorized personnel from changing files and posting to the host server, you should secure the Phone Book Administrator (PBA) and Phone Book Service (PBS) folders.

**Avoid synchronization and version problems during remote administration**
For each phone book server, you must have only one PBA folder, which must reside on your primary PBA computer. One administration tool can post multiple phone books, each to a different server, but each phone book should be administered with only one PBA computer.

You can administer phone books from the primary PBA computer or from a remote computer. A remote computer could be a portable computer or other computer that you plan to use when you cannot administer from your primary PBA computer. For example, you might be traveling and need to administer from a laptop computer, or you may want to administer from one of several workstations located in your office. You can install PBA on any remote computer, but after installing PBA you must delete the PBA folder from the remote computer to avoid synchronization and version problems.

Deleting only the PBA folder prevents an accidental start of PBA from the remote computer, but it leaves the necessary system files intact.

For more information on administering from a remote computer, see Administer from a remote computer or Administer from a remote computer by command line.

**Limit access to the PBS and PBA directories**
Connection Point Services setup configures the PBS folder with default permissions. For added security, you should limit access to the PBS folder to specific users.

If you want only certain people to post phone books to the server, make sure that the primary PBA computer is secure by protecting the PBA folder.

For more information on securing your PBS and PBA directories, see Set Permissions.

**Important**

- If your primary PBA computer is used by several people, users who do not have access to the PBA folder still have a link to Phone Book Administrator on the **Start** menu. If these users click that link, a message box gives them a choice between deleting the link, fixing it, or canceling out of the message box. They must click **Cancel**. Do not allow them to delete or fix it. Deleting this link could delete it for everyone.

# Back up your data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Back up your data**

The process of backing up and restoring data can vary greatly, depending on your network configuration. Backing up and restoring is critical to preserving POP and region information. Backing up the database and .cab files on the server is not as critical as backing up the phone book databases and other important files located on the primary computer where Phone Book Administrator (PBA) resides. Be sure to back up the following three important data files:

- PBServer.mdb
- *phone_book*.mdb
- PBAdmin.ini

# Connection Point Services How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install and Configure Connection Point Services
- Monitor the Activity of Connection Point Services
- Administer Phone Books with Phone Book Administrator
- Administer Phone Books by Command Line
- Stop Phone Book Service
- Remove Connection Point Services

# Install and Configure Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and configure Connection Point Services**

- Install Phone Book Service on a server
- Install Phone Book Administrator
- Test for production readiness
- Set Permissions

# Install Phone Book Service on a server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Phone Book Service on a server**

1. Open the Windows Components wizard.
2. In **Components**, click **Management and Monitoring Tools** (but do not select or clear the check box), and then click **Details**.
3. In **Subcomponents of Management and Monitoring Tools**, select the **Connection Point Services** check box.
4. Click **OK**, and then click **Next**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.
- If you are installing your server for the first time, skip step 1 of the installation procedure.
- The FTP service must reside on your server. If you do not want FTP to run while PBA is not in use, stop the service through the Internet Service Manager and then enable it when you are ready to post your phone book. Do not uninstall FTP, because doing so will automatically uninstall PBS.
- During installation, you will be advised that PBS requires Internet Server Application Programming Interface (ISAPI), part of Internet Information Services (IIS). If you do not choose to enable IIS, PBS will not function properly. You will have to enable PBS at a later time by using the IIS Security Wizard.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Connection Point Services
Major components and their roles
How the components work together
Set permissions early
Safeguard your files

# Install Phone Book Administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Phone Book Administrator**

1. Insert your Windows Server 2003 family CD into your CD-ROM drive.
2. Click **Perform additional tasks**.
3. Click **Browse this CD**.
4. Browse to SUPPORT\PBA.
5. Double-click **PBAInst.exe**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- You can install Phone Book Administrator (PBA) from VALUEADD\MSFT\MGMT\PBA on the product CD for the Windows XP Professional operating system or the original release of a Windows Server 2003 operating system.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Connection Point Services
Major components and their roles
How the components work together
Set permissions early
Safeguard your files

# Test for production readiness

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test for production readiness**

1. Create a test phone book, add one or more POPs, and then publish the phone book.

2. On the Phone Book Service (PBS) host, browse to \Program Files\Phone Book Service\Data and verify the presence of a new data folder named for the phone book and containing a .cab file.

3. Using the Connection Manager Administration Kit, create a service profile that is set to download this phone book from the correct PBS host. Do not include the phone book files when you create this profile.

4. Install the new service profile on your test client computer.

5. Edit the Connection Manager service provider (.cms) file to prevent automatic deletion of the .tmp file created during a phone book download:

   a. Using Notepad on the test client computer, open the .cms file.

   b. Edit the key 0=cmdl32.exe %profile% located in the [Connect Actions] section by typing **/no_delete**. After editing, the section should appear as follows:

   ```
   [Connect Actions]
   0=cmdl32.exe %profile% /no_delete
   ```

   c. Save the edited .cms file, and use the profile to connect Connection Manager to Connection Point Services.

   d. Once a connection is established and information is downloaded, disconnect from the server.

6. Using Notepad on the client computer, open the .tmp file and verify that PBS is working properly. The name of the .tmp file varies, but it is located in the same directory as the .cmp file of the service profile. The directory will likely contain more than one .tmp file. The .tmp file created by a successful phone book download should contain control characters and file names such as *[phone_book]*.pbk. If PBS is not working properly, all the .tmp files will contain error messages.

7. Open Connection Manager and verify whether the phone book appears when you click **Phone Book**.

**Notes**

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- Do not open your .tmp file before disconnecting from the server. Doing so could prevent a proper phone book update.

- You can simulate a client request by typing the following in the browser **Address** box on the test client computer and then pressing ENTER:

  **http://localhost/pbserver/pbserver.dll?osarch=0&ostype=0&osver=1&cmver=1&lcid=1033&pbver=0&pb=***[phone_book_name]*

  For example, if the phone book name is Awesome, you would type:

  **http://localhost/pbserver/pbserver.dll?**

**osarch=0&ostype=0&osver=1&cmver=1&lcid=1033&pbver=0&pb=Awesome**

If you want to see Phone Book Service error messages while you simulate client requests, you must turn off the display of friendly HTTP error messages in Internet Explorer. On the **Tools** menu, click **Internet Options**, and click the **Advanced** tab. Clear the **Show friendly HTTP error messages** check box.

- If your .tmp file does not contain a phone book, the following list might help you identify areas for further investigation.
    - An empty .tmp file indicates that the phone book request did not generate a response from the PBS host.
    - An error number of 5*xx* typically indicates that PBS is not installed or is not running.
    - An error number of 4*xx* typically indicates that you requested a phone book that did not exist in the PBS database.
    - An error number of 204 indicates that the client phone book version number is high enough that you do not need a phone book update.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add or edit a POP
Add a new phone book by command line
Set or edit phone book options by command line
Add, edit, or delete POPs by command line
Publish a phone book
Publish a phone book by command line
Set permissions early
Safeguard your files
Major components and their roles
Administer phone books
Connection Point Services Troubleshooting
Connection Manager Administration Kit

# Set Permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set permissions**

- Set permissions for the Phone Book Service folder (NTFS only)
- Set up FTP accounts for known users
- Secure your Phone Book Administrator folder (NTFS only)
- Set the Write permission for the FTP virtual directory

# Set permissions for the Phone Book Service folder (NTFS only)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions for the Phone Book Service folder (NTFS only)**

1. Open Windows Explorer.
2. Double-click the Program Files folder, right-click the Phone Book Service folder, and then click **Properties**.
3. At the bottom of the **Security** tab, click **Advanced**. Clear the **Allow inheritable permissions from the parent to propagate to this object and all child objects.** check box.
4. In the **Security** dialog box, click **Remove**.
5. Click **Add**, type the name of the user or group to which you are granting posting permissions, and then click **OK**.
6. In the **Permissions** list, select the appropriate check boxes.
7. Click **OK** for the server to register the changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- Before anyone attempts to post to the server, verify that you have set the Write permission for the FTP virtual directory.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the Write permission for the FTP virtual directory
Set up FTP accounts for known users
Set permissions early
Administer phone books
How the components work together

# Set up FTP accounts for known users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up FTP accounts for known users**

1. Open Internet Information Services (IIS) Manager.
2. In the console tree, double-click the server name, double-click **FTP Sites**, and then click **Default FTP Site**.
3. On the **Action** menu, click **Properties**.
4. On the **Security Accounts** tab, make sure that the **Allow Anonymous Connections** check box is cleared. If the check box is selected, clear it.
5. If a warning appears indicating that unencrypted passwords can be transmitted over the network, click **Yes** to close the warning.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Internet Services Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Services Manager**.
- Do not allow Write permissions to the PBSData folder until you are ready post data, and remove the Write permission as soon as you finish posting data. For more information, see Related Topics.
- You should create a user account that is used solely for posting phone books. This user account should be a local user account on the PBS server, not on a domain, and it should be a member of the Guests group. Do not activate this account until you are ready to post data, and disable this account as soon as you finish posting data. You should also disable this user account from logging on locally and from logging on through Terminal Services. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the Write permission for the FTP virtual directory
Set permissions for the Phone Book Service folder (NTFS only)
Set permissions early
Administer phone books
How the components work together
Deny log on locally
Deny log on through Terminal Services

# Secure your Phone Book Administrator folder (NTFS only)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To secure your Phone Book Administrator folder (NTFS only)**

1. Open Windows Explorer.
2. Double-click the Program Files folder, right-click the PBA folder, and then click **Properties**.
3. At the bottom of the **Security** tab, click **Advanced**. Clear the **Inherit from parent the permission entries that apply to child objects** check box.
4. In the **Security** dialog box, click **Remove**.
5. Click **Add**, type the name of the user or group to which you are granting posting permissions, and then click **OK**.
6. In the **Permissions** list, select the appropriate check boxes, and then click **OK** for the server to register the changes.

### Important

- If your primary PBA computer is used by several people, users who do not have access to the PBA folder still have a link to Phone Book Administrator on the **Start** menu. If these users click that link, a message box gives them a choice between deleting the link, fixing it, or canceling out of the message box. They must click **Cancel**. Do not allow them to delete or fix it. Deleting this link could delete it for everyone.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- Before anyone attempts to post to the server, verify that you have set the Write permission for the FTP virtual directory. For information about the Write permission, see Related Topics.
- You must have an account on the server to which you are posting.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the Write permission for the FTP virtual directory
Set permissions for the Phone Book Service folder (NTFS only)
Set up FTP accounts for known users
Set permissions early
Administer phone books
How the components work together

# Set the Write permission for the FTP virtual directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the Write permission for the FTP virtual directory**

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS)**.
2. In the console tree, double-click the server name.
3. Double-click **FTP Sites**, double-click **Default FTP Site**, and then click **PBSData**.
4. On the **Action** menu, click **Properties**.
5. On the **Virtual Directory** tab, select the **Write** check box, and then click **OK** for the server to register the changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- For enhanced security, set the Write permission only when you are about to post to the server, and then clear it immediately after posting.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set permissions for the Phone Book Service folder (NTFS only)
Set up FTP accounts for known users
Set permissions early
Safeguard your files
Major components and their roles

# Monitor the Activity of Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor the activity of Connection Point Services**

- Monitor the activity of Phone Book Service
- Monitor Phone Book Activity

# Monitor the activity of Phone Book Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To monitor the activity of Phone Book Service**

1. Open Performance.
2. In the console tree, click System Monitor.
3. In the details pane, right-click the empty graph, and then click **Add Counters**.
4. Click **Select counters from computer** and enter a computer name.
5. In **Performance object**, click **PBServer Monitor**.
6. Do one of the following:
   - To select all counters, click **All counters**, and then click **Add**.
   - To select specific counters, click **Select counters from list**, click the counter you want to use, and then click **Add**.

     Repeat this step for each counter you want to add.
7. Click **Close**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- PBServer Monitor counters are not available unless PBS has been started. PBS starts automatically when a client computer requests a phone book. To simulate a client request, type the following in the **Address** box of your browser, and then press ENTER:

  `http://localhost/pbserver/pbserver.dll?osarch=0&ostype=0&osver=1&cmver=1&lcid=1033&pbver=0&pb=`*[phone_book_name]*

  For example, if the phone book name is Awesome, you would type:

  `http://localhost/pbserver/pbserver.dll?osarch=0&ostype=0&osver=1&cmver=1&lcid=1033&pbver=0&pb=Awesome`

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

# Monitor Phone Book Activity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor phone book activity**

- View phone book files
- View the history of a phone book

# View phone book files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view phone book files**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book.
3. On the **Tools** menu, click **View Phone Book Files**.
4. In **File**, click the phone book you want to view.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.
- A phone book file contains a list of all of the available (in-service) POPs that will be downloaded to your users when you publish the phone book. If you have previously published the phone book, this list also contains new updates for POPs that you've previously published. POP fields appear in the same order as you would create them in the command-line interface.
- You should inspect your phone book files before posting them:
    - After you create new POPs.
    - After you edit existing POPs.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administer phone books
Phone book history

# View the history of a phone book

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the history of a phone book**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book.
3. On the **File** menu, click **View Log File**.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.
- The log file shows the date and time you created or changed the POP or region and the action you took.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Example phone book history
Administer phone books
Phone book history
Install Phone Book Administrator
Connection Point Services Troubleshooting

# Administer Phone Books with Phone Book Administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer phone books with Phone Book Administrator**

- Administer from a remote computer
- Create and Manage Phone Books
- Create and Manage Regions
- Create and Manage POPs

# Administer from a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To administer from a remote computer**

1. Install Phone Book Administrator (PBA) on your remote computer.
2. Delete *only* the PBA *folder* from the remote computer where you just installed PBA.
3. Establish a network connection with your primary PBA computer.
4. Establish a network connection with your primary Phone Book Service (PBS) computer.
5. On your remote computer, map a network drive to the PBA directory on your primary PBA computer.
6. Start PBA on your primary PBA computer from your remote computer.
7. Administer phone books as if you were at your primary PBA computer.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- The default location of the PBA folder is [*drive*:]\Program Files\PBA.
- A remote computer could be a portable computer or other computer that you plan to use when you cannot administer from your primary PBA computer. For example, you might be traveling and unable to administer from the primary PBA computer in your office.
- Deleting the PBA folder from the remote computer prevents an accidental start of PBA from the remote computer (which can result in data files being out of sync), but it leaves the necessary system files intact.
- For each server, you must have only one PBA folder (the one on your primary PBA computer). One administration tool can post multiple phone books, each to a different server, but each phone book should be administered with only one PBA computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Safeguard your files
Install Phone Book Administrator
Connection Point Services Troubleshooting

# Create and Manage Phone Books

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and manage phone books**

- Add and configure a new phone book
- Edit phone book options
- Copy a phone book
- Delete a phone book
- Print a phone book
- Publish a phone book

# Add and configure a new phone book

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add and configure a new phone book**

1. Open Phone Book Administrator.

2. On the **File** menu, click **New Phone Book**.

3. In **New phone book name**, type the phone book name, and then click **OK**.

   A phone book name must contain no more than eight characters, cannot consist of all digits, and must not contain a space or any symbols, including but not limited to: **! , ; * = / \ : ? ' " < > | . & % {} [] @ ( ) ´ ` ~**

   If you are creating a customized phone book for use with a Connection Manager profile, the phone book name must match the one you typed in the **Phone Book Updates** pane in the Connection Manager Administration Kit Wizard.

4. In **Phone books**, click the new phone book.

5. On the **Tools** menu, click **Options**.

6. In the **Options** dialog box, type the server name or address, user name, and password used to post phone books to the Phone Book Service (PBS) server.

   Do not use spaces in these entries.

**Notes**

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.

- The account used for posting the phone book must have a password. The password cannot be null or contain spaces.

- You can change options for this phone book at any time by repeating steps 5 and 6.

- When you publish the phone book, the user name and password will be used to send it using FTP to the server specified in the **Options** dialog box.

- For better security, do not provide a user name and password in the **Options** dialog box until you are ready to post a phone book. After you have posted a phone book, clear the user name and password entries before you close Phone Book Administrator (PBA).

- Phone Book Administrator is available only in English, French, German, Japanese, and Spanish. However, you can create a phone book in another language by running PBA on an operating system optimized for that language.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Install Phone Book Administrator

Administer phone books

Connection Point Services Troubleshooting

Connection Manager Administration Kit

# Edit phone book options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit phone book options**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book name.
3. On the **Tools** menu, click **Options**.
4. In the **Options** dialog box, edit the information as necessary.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.
- You can change options for this phone book at any time by repeating steps 3 and 4.
- When you publish the phone book, the user name and password will be used to send it through FTP to the server specified in the **Options** dialog box.
- For better security, do not provide a user name and password in the **Options** dialog box until you are ready to post a phone book. After you have posted a phone book, clear the user name and password entries before you close Phone Book Administrator.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Phone Book Administrator
Administer phone books
Connection Point Services Troubleshooting

# Copy a phone book

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy a phone book**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book to be copied.
3. On the **File** menu, click **Copy Phone Book**.
4. In **New phone book name**, type a new name for the copy.

   A phone book name must contain no more than eight characters, cannot consist of all digits, and must not contain a space or any of the following symbols: **! , ; * = / \ : ? ' " < > | . & % {} [] @ ( ) ´ ` ~**

   If you are creating a customized phone book for use with a Connection Manager profile, the phone book name must match the one you typed in the **Phone Book Updates** pane in the Connection Manager Administration Kit Wizard.

   If you intend to administer a phone book from the command line, the path to the phone book should not contain any spaces.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.
- You cannot copy a phone book to a pre-existing phone book name.
- You can change options for this phone book at any time by clicking the phone book in **Phone books** and then, on the **Tools** menu, clicking **Options**.
- All phone book data in the original phone book file is maintained in the copy except for the version number and the history. The default release directory is changed to the new name and the version number is reset to 1.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Edit phone book options
Install Phone Book Administrator
Administer phone books
Connection Point Services Troubleshooting
Connection Manager Administration Kit

# Delete a phone book

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a phone book**

1. Open Phone Book Administrator.
2. In the **Phone books** list, click the phone book to be deleted.
3. On the **File** menu, click **Delete Phone Book**.

   A dialog box appears asking whether you want to delete the phone book.
4. Click **Yes**.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.
- Deleting a phone book removes the indexed references. However, released .cab files remain on the Phone Book Administrator (PBA) and Phone Book Service (PBS) file systems. If you want to remove the .cab files, you must manually delete the phone book directory and all its contents. You should delete the directory and files from the PBS server any time you delete a phone book.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View phone book files
Delete phone book data
Install Phone Book Administrator
Administer phone books
Connection Point Services Troubleshooting

# Print a phone book

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print a phone book**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book that you want to print.
3. In **Filter by**, click a filter name.

   If you want to print all of the POPs in the phone book, click **All POPs**, and then go on to step 5.
4. In **Containing**, type the attribute, such as a numerical area code, that must be shared by all POPs on the list you want to print.
5. Click **Apply**.
6. On the **File** menu, click **Print POP List**.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Find POPs
Install Phone Book Administrator
Administer phone books
Connection Point Services Troubleshooting

# Publish a phone book

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To publish a phone book**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book.
3. If you have not yet specified a server address for this phone book, on the **Tools** menu, click **Options**, and then type a server address, user name, and password for your Phone Book Service (PBS) host.
4. On the **Tools** menu, click **Publish Phone Book**.
5. In **Release directory**, browse to the directory in which you want to keep your files.
6. Click **Create**, and then click **Post**.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.
- Before attempting to post, make sure that you have set the Write permission for the Phone Book Service (PBS) FTP virtual directory. If the Write permission is not set, the phone book will not post.
- If you specify a release directory on a remote computer, you may have to reset the directory every time you publish a phone book.
- Every time you propagate a phone book by posting, it increments the version number of the phone book. If you post to multiple servers, the version numbers will be inaccurate. To propagate a phone book to multiple servers, you can use the replication capability of the Windows Server 2003 family.
- You cannot create or post a phone book if no changes have been made to that phone book.
- Phone Book Administrator is available only in English, French, German, Japanese, and Spanish. However, you can create a phone book in another language by running PBA on an operating system optimized for that language.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the Write permission for the FTP virtual directory
Install Phone Book Administrator
Phone Book Service activity
How the components work together
Connection Point Services Troubleshooting

# Create and Manage Regions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and manage regions**

- Import a region list
- Add, edit, or delete a region

# Import a region list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import a region list**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book.
3. On the **Tools** menu, click **Regions Editor**.
4. In the **Regions Editor** dialog box, click **Import**.
5. Browse to the region file, and then click **Open**.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.
- Duplicate regions will be ignored.
- The Regions Editor is not case-sensitive.
- The Regions Editor does not accept more than 30 characters for a region name.
- A region name must not contain a double quotation mark (") or a comma (,).
- To import correctly into Phone Book Administrator (PBA), a region file must be in .pbr format. A number on the first line corresponds with the number of regions the file contains. For example:

```
3
Alaska
Arizona
California
```

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Phone Book Administrator
Administer phone books
Important phone book files
Example region files containing U.S. states
Connection Point Services Troubleshooting

# Add, edit, or delete a region

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or delete a region**

1. Open Phone Book Administrator.
2. On the **Tools** menu, click **Regions Editor**.
3. In the **Regions Editor** dialog box, do one of the following, and then click **OK**:
   - To add a region, click **Add**, and then type a region name of no more than 30 characters.

     A region name must not contain a double quotation mark (") or a comma (,).
   - To edit a region, click the region, click **Edit**, and then type the change.
   - To delete a region, click the region, and then click **Delete**.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Phone Book Administrator
Administer phone books
Important phone book files
Example region files containing U.S. states
Connection Point Services Troubleshooting

# Create and Manage POPs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and manage POPs**

- Add or edit a POP
- Delete a POP
- Assign a region to a POP
- Find POPs
- Print a POP list

# Add or edit a POP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add or edit a POP**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book.
3. Do one of the following:
   - To add a POP, on the **Edit** menu, click **Add POP**.
   - To edit a POP, click a POP entry in the **POP Name** list, and then on the **Edit** menu, click **Edit POP**.
4. On the **Access Information** tab, enter or edit the information.
   - For countries or dependencies without area codes, type a space in **Area code**.
   - **In Service** status includes the POP in future updates of the phone book file, and **Not In Service** status excludes them.
5. On the **Settings** tab, do the following (optional):
   - In **POP settings**, select the options that apply to this POP. Entries in **bold text** are automatically preselected.

| **Select** | **If** |
|---|---|
| **Sign on** | Subscribers must sign on each time they dial in. |
| Sign up | The POP allows people to sign up with your service. |
| Surcharge | You can charge subscribers a surcharge for this POP. |
| **Modem** | Subscribers can dial in using a modem. |
| **ISDN** | Subscribers can dial in using an ISDN line. |
| **Multicast** | The POP supports IP multicast. |
| Custom 1 and Custom 2 | Your service has assigned a value to these service settings in the Connection Manager file. |

   - In **Dial-Up Networking entry**, type a name that matches the name in the **Dial-up Networking Entries** pane in the Connection Manager Administration Kit Wizard.
   - In **Analog speed**, enter the minimum and maximum rate supported.
6. On the **Comments** tab, type any comments (for administrative reference only).
7. Click **Save**.

**Notes**

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.
- The country.txt file included in the PBA directory will help you determine the required country or dependency code.
- When you set the surcharge key, Connection Manager treats the POP as an alternate access number. For more information, see Related Topics.
- Sign up is not a default option in Connection Manager. You need to specify it, just as you do in Custom 1 and Custom 2.
- Do not clear the **Sign on** check box. This option must remain selected so that Connection Manager can display the access number.
- Although Phone Book Administrator is only available in English, French, German, Japanese, and Spanish, and the country/dependency list is only available in English, country/dependency names chosen from the English list appear in the same language as the Connection Manager profile. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign a region to a POP
Install Phone Book Administrator
Administer phone books
Service types (POP settings)
Connection Point Services Troubleshooting
Connection Manager Administration Kit
Providing phone book support
Incorporating custom dial-up entries
Creating profiles in multiple languages

# Delete a POP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a POP**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book.
3. In the **POP Name** list, click a POP.
4. On the **Edit** menu, click **Delete POP**.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Phone Book Administrator
Administer phone books
Service types (POP settings)
Connection Point Services Troubleshooting

# Assign a region to a POP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign a region to a POP**

1. Open Phone Book Administrator.
2. In **Phone books**, click a phone book.

   If you have not already added the region that you want to assign to this POP, add it to the phone book now.

   a. On the **Tools** menu, click **Regions Editor**.

   b. In the **Regions Editor** dialog box, click **Add**, and then type a region name of not more than 30 characters.

      A region name must not contain a double quotation mark (") or a comma (,).
3. In the **POP Name** list, click a POP.
4. On the **Edit** menu, click **Edit POP**.
5. In **Region**, click a region, and then click **OK**.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add, edit, or delete a region
Install Phone Book Administrator
Administer phone books
Service types (POP settings)
Example region files containing U.S. states
Connection Point Services Troubleshooting

# Find POPs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find POPs**

1. Open Phone Book Administrator.
2. In **Phone books**, click a phone book.
3. To find POPs by access number, area code, country/dependency, name, region, or status, click an option in **Filter by**.
4. In **Containing**, type a whole or partial word or number.
5. Click **Apply**.

   POPs that match your filter criteria appear in the list of POPs.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.
- You can sort your POP list by POP name, area code, access number, country/dependency, region, or status by clicking column titles.
- To turn filtering off and find all POPs in the phone book, in **Filter by**, click **All POPs**, and then click **Apply**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Phone Book Administrator
Administer phone books
Service types (POP settings)
Connection Point Services Troubleshooting

# Print a POP list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print a POP list**

1. Open Phone Book Administrator.
2. In **Phone books**, click the phone book whose POP list you want to print.
3. In **Filter by**, click a filter name.

   If you want to print all of the POPs in the phone book, click **All POPs**, and then go on to step 5.
4. In **Containing**, type the attribute, such as a numerical area code, that must be shared by all POPs on the list you want to print.
5. Click **Apply**.
6. On the **File** menu, click **Print POP List**.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Phone Book Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Phone Book Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Find POPs
Install Phone Book Administrator
Administer phone books
Service types (POP settings)
Connection Point Services Troubleshooting

# Administer Phone Books by Command Line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer phone books by command line**

- Administer from a remote computer by command line
- Create and Manage Phone Books by Command Line
- Create and Manage Regions by Command Line
- Create and Manage POPs by Command Line

# Administer from a remote computer by command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To administer from a remote computer by command line**

1. Install Phone Book Administrator (PBA) on your remote computer.
2. Delete *only* the PBA *directory* from the remote computer where you just installed PBA.
3. Establish a network connection with your primary PBA computer.
4. Establish a network connection with your primary Phone Book Service (PBS) computer.
5. On your remote computer, map a network drive to the PBA directory on your primary PBA computer.
6. Change your directory on the remote computer to the mapped network drive.
7. Administer phone books as if you were at your primary PBA computer.

### Important

- If you intend to administer PBA from the command line, do not use spaces in the path to pbadmin.exe or in the path to the phone book files.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- The default location of the PBA directory is [*drive*:]\Program Files\PBA.
- A remote computer could be a portable or other computer that you plan to use when you cannot administer from your primary PBA computer. For example, you might be traveling and be unable to administer from the PBA computer in your office.
- Deleting the PBA directory prevents an accidental start of PBA from the remote computer, but it leaves the necessary system files intact.
- You must administer each phone book using only one installation of PBA. Administering the same phone book using multiple installations of PBA causes version conflicts that will prevent your users from obtaining the most recent version of the phone book.
- You can publish multiple phone books to the same phone book server, even if you administer each phone book with a different installation of PBA.
- You must not publish the same phone book to more than one phone book server. Publishing the same phone book to more than one phone book server causes version conflicts that will prevent your users from obtaining the most recent version of the phone book. However, you can replicate the phone book to multiple servers from the server to which you published the phone book. You can use the Distributed File System (DFS) console to automate replication, or you can replicate the phone book yourself by copying the folder to the other servers. For more information about DFS, see Distributed File System in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Safeguard your files
Connection Point Services Troubleshooting
Install Phone Book Administrator
Distributed File System (DFS)

# Create and Manage Phone Books by Command Line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and manage phone books by command line**

- Add a new phone book by command line
- Set or edit phone book options by command line
- Publish a phone book by command line

# Add a new phone book by command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a new phone book by command line**

1. Open Command prompt.

2. At the command prompt, type the following:

   **cd \Program Files\PBA**

3. At the command prompt, type the following, using the appropriate file name for the new phone book:

   **Pbadmin.exe /N***[phone_book_name]*

   A phone book name must contain no more than eight characters, cannot consist of all digits, and must not contain a space or any of the following symbols: **! , ; * = / \ : ? ' " < > | . & % {} [] @ ( ) ´ ` ~**

   For example, if you want to create a phone book with the name of Awesome, you type:

   **Pbadmin.exe /N Awesome**

   Command-line parameters are case-sensitive.

4. Press ENTER to add your new phone book to the database.

### Important

- If you intend to administer PBA from the command line, do not use spaces in the path to pbadmin.exe or in the path to the phone book files.

### Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- To open the command prompt window, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are creating a customized phone book for use with a Connection Manager profile, the phone book name must match the one you typed in the **Phone Book** pane in the Connection Manager Administration Kit Wizard.

- For more information about the **pbadmin** command, type **pbadmin /?** at the command prompt, or see **pbadmin** in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set or edit phone book options by command line

Administer phone books
Connection Point Services Troubleshooting
Pbadmin

# Set or edit phone book options by command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set or edit phone book options by command line**

1. Open Command prompt.

2. At the command prompt, type the following:

   **cd \Program Files\PBA**

3. At the command prompt, type the following, using the appropriate phone book file name, and then a server name or address, user name, and password to set or edit those options for the phone book:

   **Pbadmin.exe /O***[phone_book_name][server_name][user_name][password]*

   For example, if you want to set options for a phone book with the name of Awesome, using a server name of Rainier, a user name of jobrown, and a password of Tr6f23nN, you type:

   **Pbadmin.exe /O Awesome Rainier jobrown Tr6f23nN**

   Command-line parameters are case-sensitive.

4. Press ENTER.

### Important

- If you intend to administer PBA from the command line, do not use spaces in the path to pbadmin.exe or in the path to the phone book files.

### Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- To open the command prompt window, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- For more information about the **pbadmin** command, type **pbadmin /?** at the command prompt, or see **pbadmin** in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a new phone book by command line
Administer phone books
Connection Point Services Troubleshooting
Pbadmin

# Publish a phone book by command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To publish a phone book by command line**

1. Open Command prompt.
2. At the command prompt, type the following:

   **cd \Program Files\PBA**

3. At the command prompt, type the following:

   **Pbadmin.exe /B** *[phone_book_name]*

   Command-line parameters are case-sensitive.

4. Press ENTER.

### Important

- If you intend to administer PBA from the command line, do not use spaces in the path to pbadmin.exe or in the path to the phone book files.

### Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open the command prompt window, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Before you post, make sure that:
  - The host server is running.
  - You are connected to the network.
  - You have FTP, File, and Write permissions to post to the host server.
  - The account used for posting the phone book has a password. The password cannot be null.
  - You have set the host server name, user name, and password in the command line (by using the /O parameter) or in the PBA graphical interface. For more information on how to set a host server name, a user name, and a password, see Related Topics.
- For more information about the **pbadmin** command, type **pbadmin /?** at the command prompt, or see **pbadmin** in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Administer phone books
Major components and their roles
How the components work together
Connection Point Services Troubleshooting
Set Permissions
Set or edit phone book options by command line
Pbadmin

# Create and Manage Regions by Command Line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and manage regions by command line**

- Add, edit, or delete regions by command line
- Import a region list by command line

# Add, edit, or delete regions by command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or delete regions by command line**

1. Do one of the following:

    - If no region input file exists, create one by using a plain-text editor such as Notepad.
    - Export data from your existing database into a plain-text input file. Make sure that the data is formatted according to the steps in this procedure.

2. Open the plain-text input file, and do one of the following:

    - To add a region, on a new line, type the following:

A unique ID number for the region you are adding.

A comma delimiter.

A new region name of no more than 30 characters and with no double quotation marks (") or commas (,).

For example, with an ID number of 13 and a region name of Alaska, you type:

**13,Alaska**

- To edit a region, on a new line, type the following:

The ID number of the region.

A comma delimiter.

A new region name of no more than 30 characters and with no double quotation marks (") or commas (,).

For example, for an ID number of 12, and a new name of California, you type:

**12,California**

- To delete a region, on a new line, type the following:

The ID number of the record to be deleted, followed by a comma.

Leave the name field blank.

For example, if the ID number of the field were 12, you type:

**12,**

1. Press ENTER.

2. Repeat steps 2 and 3 for each region you want to add, edit, or delete, and then save the file with a .txt extension.

3. Import the region file from the command line.

**Important**

- If you intend to administer PBA from the command line, do not use spaces in the path to pbadmin.exe or in the path to the phone book files.

### Notes

- Do not leave blank lines in the file.
- POPs and regions must be in separate files.
- A command line region record consists of two fields: an ID number (integer) and a region name, delimited by a comma. The ID number must be unique. When it matches the ID number of an existing region, it overwrites that region.
- Phone book files saved in formats other than plain text (for example, Microsoft Word, WordPad, or Rich-Text Format) do not work properly.
- Before you can import phone book or region files into a particular phone book, the phone book name must be registered in the Phone Book Administrator (PBA) phone books list. If the phone book name does not appear in that list, you must use PBA to create the phone book.
- Command-line region files do not match the region-file format imported or generated by the PBA user interface.
- For a list of U.S. states that you can copy, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Import a region list by command line
Add and configure a new phone book
Example region files containing U.S. states
Administer phone books
Example region files containing U.S. states
Connection Point Services Troubleshooting

# Import a region list by command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import a region list by command line**

1. Open Command prompt.
2. At the command prompt, type the following:

   **cd \Program Files\PBA**

3. At the command prompt, type the following, using the appropriate paths and file names:

   **Pbadmin.exe /I***[phone_book_name]***/R***[region_file_path]******[region_file_name]*

   Command-line parameters are case-sensitive.

4. Press ENTER.

### Important

- If you intend to administer PBA from the command line, do not use spaces in the path to pbadmin.exe or in the path to the phone book files.

### Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open the command prompt window, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- For more information about the **pbadmin** command, type **pbadmin /?** at the command prompt, or see **pbadmin** in Related Topics.
- Before you can import phone book or region files into a particular phone book, the phone book name must be registered in the Phone Book Administrator (PBA) phone books list. If the phone book name does not appear in that list, you must use PBA to create the phone book.
- Phone book files saved in formats other than plain text (for example, Microsoft Word, WordPad, or Rich-Text Format) do not work properly.
- Command-line region files do not match the region-file format imported or generated by the PBA user interface.
- POPs and regions must be in separate files.
- If you have not yet posted your phone book, you can import up to 32,000 POPs. If, however, you have already posted, the limit is 6,000 POPs (combined adds, edits, and deletes) for each subsequent post.
- If you are importing POP files and region files by using separate commands, and if your POPs require region information, import your region file first.
- Phone Book Administrator should not be open while you are importing files.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add, edit, or delete regions by command line
Add and configure a new phone book
Administer phone books
Example region files containing U.S. states
Connection Point Services Troubleshooting
Pbadmin

# Create and Manage POPs by Command Line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and manage POPs by command line**

- Add, edit, or delete POPs by command line
- Import an existing POP list by command line

# Add, edit, or delete POPs by command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or delete POPs by command line**

1. Do one of the following:

   - If no POP input file exists, create one by using a plain-text editor such as Notepad.
   - Export data from your existing database into a plain-text input file. Make sure that the data is formatted according to the steps in this procedure.

2. Open the plain-text input file and do one of the following:

   - To add or edit a POP: On a new line and using the following syntax for each POP you want to add or edit, type the data for each field in the order given. Follow each data item except the last with a comma and a space. These items are described in the following table.

     *[POP index]*, *[Country/dependency ID]*, *[Region ID]*, *[POP name]*, *[Area code]*, *[Access number]*, *[Minimum baud]*, *[Maximum baud]*, *[Reserved flag]*, *[Service flag]*, *[Dial-up entry]*

     For example:

     1, 1, 0, Redmond, 425, 555-0134, 9600, 56000, 0, 120, ISP

| Field | Description | Data type | Characters allowed | Limit |
|---|---|---|---|---|
| POP index | Numerical POP identifier | Integer | All numbers | 32,000 rows |
| Country/dependency ID | TAPI code | Integer | All numbers | No limit |
| Region ID | Numerical region identifier | Integer | All numbers (A zero (0) identifies all regions.) | No limit |
| POP name | Name given to the POP (maximum of 30 characters) | Text | All except comma (,) and double quotation mark (") | 30 characters |
| Area code | Area code for the access number | Text | All numbers, hyphen (-), space ( ) | 10 characters |
| Access number | Phone number that subscribers can dial to reach your service | Text | All numbers, number sign (#), asterisk (*), hyphen (-), space ( ) | 35 characters |
| Minimum baud | Minimum baud rate of the POP | Integer | Typical modem speeds | No limit |
| Maximum baud | Maximum baud rate of the POP | Integer | Typical modem speeds | No limit |
| Reserved flag | Unassigned flag that must always be set to zero (0) | Integer | Zero (0) | No limit |

| | | | | |
|---|---|---|---|---|
| Service flag | Decimal value representing a combination of POP settings | Integer | See Related Topics | No limit |
| Dial-up entry | Name of the entry created in the **Dial-up Networking Entries** pane in the Connection Manager Administration Kit, used in the Connection Manager profile (.cms file) | Text | All except comma (,), ampersand (&), percent (%), apostrophe ('), acute (´), agave (`), and double quotation mark (") | 50 characters |

- To delete a POP:On a new line, first type the index number of the POP you want to delete. In each of the next nine fields, type a zero, followed by a comma. Leave the 11th field blank.

  In the following example, the POP index number is 12:

  12,0,0,0,0,0,0,0,0,0,

3. Press ENTER.

4. Repeat steps 2 and 3 for each POP you want to add, edit, or delete, and then save the file with a .txt extension.

5. Import the file.

**Important**

- If you intend to administer PBA from the command line, do not use spaces in the path to pbadmin.exe or in the path to the phone book files.

**Notes**

- POPs and regions must be in separate files.

- The ID number must be unique. If it matches the ID number of an existing POP, it overwrites that POP.

- Phone book files saved in formats other than plain text (for example, Microsoft Word, WordPad, or Rich-Text Format) do not work properly.

- Before you can import phone book files into a particular phone book, the phone book name must be registered in the **Phone books** list in Phone Book Administrator (PBA). If the phone book name does not appear in that list, you must create the phone book.

- If you are importing POP files and region files by using separate commands, and if your POPs require region information, import your region file first.

- The country.txt file included in the PBA directory will help you determine the required country or dependency code you will need for your POP files.

- Phone Book Administrator (PBA) is available only in English, French, German, Japanese, and Spanish. However, you can create a phone book in another language by running PBA on an operating system optimized for that language.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Import an existing POP list by command line
Service types (POP settings)
Add a new phone book by command line
Administer phone books
Use your existing database with Phone Book Administrator
Connection Point Services Troubleshooting

# Import an existing POP list by command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import an existing POP list by command line**

1. Open Command prompt.
2. At the command prompt, type the following:

   **cd \Program Files\PBA**

3. At the command prompt, type the following, using the appropriate paths and file names:

   **Pbadmin.exe /l***[phone_book_name]***/P***[POP_file_path]******[POP_file_name]

   Command-line parameters are case-sensitive.

4. Press ENTER.

### Important

- If you intend to administer PBA from the command line, do not use spaces in the path to pbadmin.exe or in the path to the phone book files.

### Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open the command prompt window, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- For more information about the **pbadmin** command, type **pbadmin /?** at the command prompt, or see **pbadmin** in Related Topics.
- Before you can import phone book or region files into a particular phone book, the phone book name must be registered in the Phone Book Administrator (PBA) phone books list. If the phone book name does not appear in that list, you must use PBA to create the phone book.
- Phone book files saved in formats other than plain text (for example, Microsoft Word, WordPad, or Rich-Text Format) do not work properly.
- Command-line region files do not match the region-file format imported or generated by the PBA user interface.
- POPs and regions must be in separate files.
- If you have already posted, your import file must contain no more than 6,000 POPs (combined adds, edits, and deletes) for each subsequent post.
- If you are importing POP files and region files by using separate commands, and if your POPs require region information, import your region file first.
- Phone Book Administrator should not be open while you are importing files.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add, edit, or delete regions by command line
Import a region list by command line
Add, edit, or delete POPs by command line
Administer phone books
Use your existing database with Phone Book Administrator
Connection Point Services Troubleshooting
Pbadmin

# Stop Phone Book Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To stop Phone Book Service

1. Click **Start**, point to **Administrative Tools**, and then click **Internet Information Services (IIS)**.
2. Double-click the name of your server.
3. If Phone Book Service is running in its own process, double-click the default Web site for your server. Right-click **PBServer**, and click **Properties**.
4. If Phone Book Service is not running in its own process, right-click the default Web site for your server, and click **Properties**. Click the **Home Directory** tab.
5. In the **Application Settings** area, click **Unload**, and click **OK**.

### Caution

- If you unload applications from the **Home Directory** tab of **Default Web Site Properties**, you will unload all Web applications for the default Web server that are not running in their own processes. You might need to restart some or all of your Web applications manually.

### Note

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Remove Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remove Connection Point Services**

- Change directory access (NTFS partitions only)
- Remove Phone Book Service from a server
- Remove Phone Book Administrator
- Delete phone book data

# Change directory access (NTFS partitions only)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change directory access (NTFS partitions only)**

1. Open Windows Explorer.
2. Do one of the following:
   - If you are changing access for Phone Book Service, browse to the Program Files folder, right-click the Phone Book Service folder, and then click **Properties**.
   - If you are changing access for Phone Book Administrator, browse to the Program Files folder, right-click the PBA folder, and then click **Properties**.
3. On the **Security** tab, in **Name**, click the group to which you belong.
4. In **Permissions**, select the **Full Control** check box under **Allow**, and then click **OK** to register the changes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Remove Phone Book Service from a server
Remove Phone Book Administrator
Delete phone book data
Set permissions early
Safeguard your files
Connection Point Services Troubleshooting

# Remove Phone Book Service from a server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove Phone Book Service from a server**

1. Open the Windows Components wizard.
2. In **Components**, click **Management and Monitoring Tools** (but do not select or clear its check box), and then click **Details**.
3. In **Subcomponents of Management and Monitoring Tools**, clear the **Connection Point Services** check box.
4. Click **OK**, and then click **Next**.

### Important

- If you have not already installed Internet Information Services (IIS), IIS components will be installed when you install Connection Point Services components. If you remove Connection Point Services components from your server, you do not automatically remove IIS components. You must use the Windows Components Wizard to remove IIS components. For more information, see the Microsoft Web site.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.
- When you remove Phone Book Service (PBS), not all data files are removed. You must delete those files manually.
- Before deleting data, you may first need to change security settings in order to obtain access.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete phone book data
Change directory access (NTFS partitions only)
Remove Phone Book Administrator
Major components and their roles
How the components work together
Administer phone books
Connection Point Services Troubleshooting

# Remove Phone Book Administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove Phone Book Administrator**

1. Open Add or Remove Programs.
2. Select **Phone Book Administrator**.
3. Click **Change/Remove**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**.
- When you remove Phone Book Administrator (PBA), not all data files are removed. These include phone book databases and .pbk, .pbr, and .cab files, both on the PBA computer and on the computer running Phone Book Service. You must delete those files manually.
- Before deleting data, you may first need to change security settings in order to obtain access.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Remove Phone Book Service from a server
Change directory access (NTFS partitions only)
Delete phone book data
Major components and their roles
How the components work together
Administer phone books
Connection Point Services Troubleshooting

# Delete phone book data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete phone book data**

1. Open Windows Explorer.
2. Double-click the Program Files folder, and then do one of the following:
   - If you are deleting Phone Book Service data, delete the Phone Book Service folder.
   - If you are deleting Phone Book Administrator data, delete the PBA folder.

## Notes

- You must be logged on as a member of the Administrators or the Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change directory access (NTFS partitions only)
Remove Phone Book Administrator
Remove Phone Book Service from a server
Major components and their roles
How the components work together
Administer phone books
Connection Point Services Troubleshooting

# Connection Point Services Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

- Understanding Connection Point Services
- Using Connection Point Services
- Connection Point Services Resources

# Understanding Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Connection Point Services**


- Connection Point Services overview
- Major components and their roles
- How the components work together
- Security information for Connection Point Services
- Example system configurations

# Connection Point Services overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection Point Services overview**

Connection Point Services (CPS) is a component that provides Internet service providers (ISPs) and corporations with the ability to automatically update phone book files residing on the client computers of employees or Internet users. These phone book files work with Connection Manager (CM), configurable connection software that also resides on client computers. For complete information about Connection Manager, see the Connection Manager Administration Kit.

As a CPS administrator, you can:

- Make it easy for your new users to get online by providing updated access numbers.
- Create and manage phone book databases containing customized information about local access numbers that provide connectivity to points of presence (POPs) in a wide area network.
- Automatically download new access numbers to your users when POPs change.
- Host multiple phone books on a single server.

CPS consists of two major components, Phone Book Administrator (PBA) and Phone Book Service (PBS). For more information about PBA and PBS, see Major components and their roles.

PBS must run on your server, but administrators often use PBA on workstations. PBS is a component in Management and Monitoring Tools on your Microsoft® Windows Server 2003 family CD. PBA is available from the Valueadd directory on your product CD for Microsoft Windows XP Professional or the original release of Windows Server 2003. You can also install PBA from the Support folder on your product CD for Windows Server 2003 with Service Pack 1.

# Major components and their roles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Major components and their roles**
Connection Point Services (CPS) consists of two major components: Phone Book Service (PBS) and Phone Book Administrator (PBA).

**Phone Book Service**
Phone Book Service (PBS) is an Internet Information Services (IIS) extension. When users query your phone book server, PBS automatically compares their phone book versions with the most recent files in your database and then downloads the appropriate updates to the Connection Manager profiles residing on their computers.

To run PBS, your server must be using IIS running both File Transfer Protocol (FTP) and the World Wide Web (WWW) services. The FTP service accepts updated phone book information posted from Phone Book Administrator (PBA), and the WWW service handles Connection Manager queries. For more information about PBS and security, see Security information for Connection Point Services.

**Phone Book Administrator (PBA)**
PBA enables you to do the following:

- Create and edit up to one hundred unique phone books. Each phone book is a collection of POPs, each one represented by a local access number for a specific region within a country or dependency.
- Choose service types (connection settings, such as ISDN access or multicast connection) for each POP. Service types, also called *POP settings*, help users to easily find POPs that suit their needs.
- Add and delete new regions and POPs.
- Associate POPs with network configurations defined in Connection Manager.
- Publish new phone book information to PBS.
- Using the command-line interface, import phone book information from other applications for distribution to end users.

Before publishing, or *posting*, to the server, PBA uses a utility to compress the phone book files, increasing the speed at which you can download the new information to your users.

**Connection Manager**
Connection Manager is not an actual component of CPS; it is a companion software package that you must have in order to use CPS.

Connection Manager is a client dialer and connection software tool that you can customize by using the Connection Manager Administration Kit (CMAK), which is a wizard that you run to build a customized service profile (a set of files that you distribute to your users so they can easily install and run your custom version of Connection Manager). Connection Manager supports a variety of features:

- Branding: You can customize the graphics, icons, logon messages, and Help.
- Custom actions: You can customize functionality, including your own programs to enhance the connection experience for your users.
- Simplified distribution: You can create a self-installing executable file that can be distributed on compact disc or downloaded by your users.
- Custom phone books: Using the CMAK wizard, you can specify the phone books to be provided to your users. Using PBS, you can provide automatic phone book updates when your users log on. And to simplify maintenance, you can combine existing phone books by merging existing service profiles.

**Notes**

- PBA and PBS need a TCP/IP connection in order to post phone book information through FTP.

- FTP service must reside on your server. If you do not want FTP to run while PBA is not in use, stop the service through the Internet Services Manager and then enable it when you are ready to post. Do not uninstall FTP, because doing so will automatically uninstall PBS.

- You can use the current version of PBS with earlier versions of PBA.

# How the components work together

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How the components work together**

When you provide Phone Book Administrator (PBA) with new phone book information, PBA posts the updates to the Phone Book Service (PBS). The next time a user connects to your network using a Connection Manager profile, Connection Manager queries PBS for updated phone book information. PBS compares the phone book name and version sent from Connection Manager with the updated phone book files and then downloads any new phone book information to the client computer.

PBA relies on the FTP service to post the phone book files to PBS. PBS relies on the WWW service, provided through Internet Information Services (IIS), to receive phone book queries from Connection Manager and download the appropriate updates.

PBA and PBS need a TCP/IP connection in order to transfer phone book information through FTP.

The following diagram and accompanying table illustrate the publishing process.




| Step 1: Establish database | Step 2: Post to server | Step 3: Update client computer |
|---|---|---|
| You can use PBA to create new phone book files. Or, if you already have a database in another application, you can import that data file using the PBA command-line interface. | Imported files or data you entered using PBA are stored in the phone book database. When you publish phone books, PBA publishes changes to PBS through FTP. | Using Connection Manager, clients query the phone book server. If the versions of their phone books are outdated, PBS automatically sends the appropriate update. |

**Notes**

- You can use the current version of PBS with earlier versions of PBA.
- You must administer each phone book using only one installation of PBA. Administering the same phone book using multiple installations of PBA causes version conflicts that will prevent your users from obtaining the most recent version of the phone book.
- You can publish multiple phone books to the same server running a member of the Windows Server 2003 family, even if you administer each phone book with a different installation of PBA.
- You must not publish the same phone book to more than one phone book server. Publishing the same phone book to more than one phone book server causes version conflicts that will prevent your users from obtaining the most recent version of the phone book. However, you can replicate the phone book to multiple servers from the server to which you published the phone book. You can use the Distributed File System (DFS) console to automate replication, or you can replicate the phone book yourself by copying the folder to the other servers. For more information about DFS, see Distributed File System (DFS).
- The FTP service must reside on your server. If you do not want FTP to run while PBA is not in use, stop the service through the Internet Services Manager and then enable it when you are ready to post. Do not uninstall FTP, because doing so will automatically uninstall PBS.

# Security information for Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for Connection Point Services**
Before you deploy Connection Point Services (CPS), you should recognize and plan for its specific security requirements. CPS comprises Phone Book Administrator (PBA) and Phone Book Service (PBS). Each component has its own implications, and there are additional security considerations for the interaction between the two programs.

To deploy CPS securely, you must also consider the general security of all of your servers, including your PBS server. For more information about server security, see Security.

When you plan your CPS deployment, you should consider the following:

- **Using PBA or PBS**. A user must be a member of either the Power Users or the Administrators group in order to use PBA or to start or stop PBS. A user must be a member of the Administrators group to install PBA or PBS. You should limit the membership of the Power Users and Administrators groups to the minimum number of users required to administer the server.

- **Managing the File Transfer Protocol (FTP) service**. PBA requires FTP to publish and update phone books to the PBS server. The PBS server requires FTP to receive phone books. To increase security, you might want to:

    - Start the FTP service on the computer running PBA immediately before you publish or update a phone book, and stop the service immediately after you finish. Make sure that no processes are using the FTP service before you stop it. For more information about how to start and stop services such as FTP, see Start, stop, pause, resume, or restart a service.

    - Do not enable anonymous access for FTP on your PBS server. Anonymous access is disabled by default when you install PBS, even if it was previously enabled. You should not install PBS on a server that allows anonymous FTP access. You can also check a log file to see who has posted files to that server. For more information, see Set up FTP accounts for known users.

    - Restrict the account that you create on the PBS server for posting phone book data to minimal permissions, and use that account only for posting phone book data. The user posting phone book data to the PBS server must have a user account on that server, the account must have a password, and Write permission to the PBSData virtual directory must be enabled. This user account should be a local user account on the PBS server, not on a domain, even if the PBS server is a member of a domain. For additional security, you should make this user account a member of the Guests group, and you should disable this user account from logging on locally and from logging on through Terminal Services. For more information, see Add a member to a local group, Deny log on locally, and Deny log on through Terminal Services.

    - FTP automatically sends passwords in plaintext. Do not activate the user account that you created for posting phone book data until you are ready to post data, and disable this account as soon as you finish posting data. Consider changing the password for this user account often, possibly after each posting. If one or more PBS servers are directly connected to the Internet without a firewall, you should consider alternatives to using FTP across the Internet. For example, you can host a PBS server inside your network and use file replication over a virtual private network (VPN) connection to update the PBS servers on the Internet. For more information, see Administer phone books and Disable or activate a local user account.

- **Protecting phone book files**. Phone book files are plain text files that any user can edit with any plain text editor, such as Notepad. Users do not need to be members of the Power Users or the Administrators group to read, change, or delete these files. To minimize the possibility that users might change these files, you might want to:

    - Limit access to the directories that contain phone book files so that only specific users or user groups can change or delete those files. For example, you can change the permissions on the PBA directories to allow only users with administrative credentials to access those directories. By default, any member of the Power Users or the Administrators groups can access the PBS and PBA directories. For more information on securing your PBS and

PBA directories, see Set Permissions.

- If you are issuing a phone book as part of a Connection Manager profile, encourage your users to install the profile for their individual use only (Windows 2000, Windows XP, and the Windows Server 2003 family only). If a user installs a profile for all users, any user can modify or delete phone book information. For more information, see Security information for the Connection Manager Administration Kit.

- **Securing Internet Information Services (IIS)**. PBS requires IIS to function. You should take steps to secure your PBS server, including but not limited to running the IIS Security Wizard. For more information, see Installing IIS.

# Example system configurations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Example system configurations**
Connection Point Services (CPS) lends itself to various configurations, according to your needs. A few examples follow.

**Dedicated single server**
**Example uses**:

- Testing
- Small company
- Small Internet service provider

You can maintain both Phone Book Service (PBS) and Phone Book Administrator (PBA) on a single computer running an operating system in the Windows Server 2003 family.

Even though PBA posts to the same server on which it resides, you must use the same procedures for setting permissions and posting phone books as you would with any other configuration.

**Dedicated Phone Book Service server with a Phone Book Administrator client**
**Example uses**:

- Medium to large corporations
- When ownership and responsibilities for phone book administration and server maintenance are split between groups

In this configuration, PBS and PBA are installed on separate computers. PBA could be installed on a server or on a workstation running Windows XP Professional. The following illustration shows this configuration.




**Dedicated Phone Book Service server with remote administration of Phone Book Administrator client**
**Example use**: When the primary computer running Phone Book Administrator is not physically accessible to the administrator, you can use this dual-mode system.

You can configure PBA to run on a primary (dedicated) computer and on a remote workstation at the same time.

The following illustration shows this configuration.

All data files reside on the primary computer, never on the remote workstation. The remote workstation accesses the data files on the primary computer.

**Multiple servers with firewall**
**Example uses**:

- Highly secured environment
- Very large Internet service providers

- Phone book replication among multiple Internet service providers

You can install PBA on a primary computer and on multiple remote workstations. PBS is installed on a staging server and on multiple host servers residing in a less secure environment outside a firewall. The following illustration shows this configuration.




The remote workstations access phone book data on the primary PBA computer. Phone book updates are posted to the staging server. Using a content replication method, phone book updates are then copied from the staging server through the firewall to the host servers.

# Using Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Connection Point Services**


- Determine how many servers you need
- Install Connection Point Services
- Administer phone books
- Use your existing database with Phone Book Administrator
- Example Implementations
- Monitor Activity

# Determine how many servers you need

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Determine how many servers you need**

Every time a user logs on, the Phone Book Service (PBS) server is "hit" with an HTTP query, initiating the phone-book update process. Along with hardware throughput and update (.cab) file size, the number of hits and the processor speed are important factors in determining how many servers you may need.

The following table details how many hits the listed processors are capable of handling on a dedicated server with 128 megabytes (MB) of RAM.

| Processor | Hits/sec | Hits/hour | Hits/day |
|---|---|---|---|
| Pentium III 600 MHz | 275 | 990,000 | 23,760,000 |
| Pentium III 500 MHz (dual processors) | 250 | 900,000 | 21,600,000 |
| Celeron 400 MHz | 125 | 450,000 | 10,800,000 |
| Pentium II 300 MHz | 125 | 450,000 | 10,800,000 |

To get an accurate data pool to determine if an additional server is needed, take an average of the number of logons your network receives during peak hours for a week.

**Notes**

- The maximum number of hits per second depends on the size of the update (.cab) file. Larger update files can slow performance by as much as 25 percent. The numbers in the table assume small update files of about 5 kilobytes (KB).
- Your server performance might increase if you run PBS in the Internet Information Services (IIS) process.

# Install Connection Point Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install Connection Point Services**

Before you install Connection Point Services (CPS), you must prepare for the installation:

- Determine and verify the computer names and IP addresses of all the computers on which you will be installing CPS. You might want to have a dedicated server so phone book requests do not interfere with the response time of your other services.
- Be aware that the Phone Book Service (PBS) server will be using Internet Information Services (IIS) to run the FTP and WWW services. This is necessary for PBS to receive posts from Phone Book Administrator (PBA) and to download updated files to client computers. For more information about PBS and security, see Security information for Connection Point Services.
- If you have multiple servers running PBS, consider how you want to replicate data. You can use PBA to publish a phone book to only one server. You will need to replicate the phone book data to the other servers.
- Document all of your processes, as well as the choices and tradeoffs that you made. This information will be valuable for troubleshooting your system and for future reference.

For information on CPS installation and configuration, see Install and Configure Connection Point Services.

After installing Connection Point Services, you must set directory access permissions and test your installation.

To test the installation, create a test phone book and post it to the server. However, to fully test server functionality, you must connect to the server with Connection Manager, using a service profile built by the Connection Manager Administration Kit wizard.

The FTP service already resides on your server. If you do not want FTP to run while PBA is not in use, stop the service and then enable it when you are ready to post.

**Important**

- Do not uninstall FTP, because doing so will automatically uninstall PBS.
- You must administer each phone book using only one installation of PBA. Administering the same phone book using multiple installations of PBA causes version conflicts that will prevent your users from obtaining the most recent version of the phone book.
- You can publish multiple phone books to the same phone book server, even if you administer each phone book with a different installation of PBA.
- You must not publish the same phone book to more than one phone book server. Publishing the same phone book to more than one phone book server causes version conflicts that will prevent your users from obtaining the most recent version of the phone book. However, you can replicate the phone book to multiple servers from the server to which you published the phone book. You can use the Distributed File System (DFS) console to automate replication, or you can replicate the phone book yourself by copying the folder to the other servers. For more information about DFS, see Distributed File System (DFS).

# Administer phone books

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer phone books**
Connection Point Services (CPS) provides two ways to administer phone books: by using the Phone Book Administrator (PBA) interface and by using the command line. You are advised to create and publish your first phone book by using the PBA interface, even if you plan to use the command line for your usual publishing procedures.

**Phone books**
Establishing a new phone book is as easy as giving it a name and configuring it with a server name or address, user name, and password. Because Connection Manager (CM) must be configured to query the Phone Book Service database for phone books, the phone book names you choose must match the ones in the service profile created by the Connection Manager Administration Kit (CMAK) wizard. If your Connection Manager profile will use multiple phone books, name them intuitively, according to their use.

For a comprehensive understanding of how CPS works with Connection Manager, run the CMAK wizard and explore CM concepts.

Make sure that you plan a method for maintaining phone books. Ask yourself:

- Who will be entering data?
- If you will be using the command line for administration, who will create and maintain data files? Will you be exporting data from an existing database or typing new files by hand?
- Do you want to use the same phone book for all locales, or should you translate and maintain multiple language versions: one for each country or dependency? Phone Book Administrator is available only in English, French, German, Japanese, and Spanish. However, you can create a phone book in another language by running PBA on an operating system optimized for that language.

After you create a phone book, you must add POPs, and you can create regions. You can update either of these as often as necessary.

For information on how to create a phone book, see Add and configure a new phone book or Add a new phone book by command line.

**Regions**
To reduce the number of POPs that Connection Manager displays at any one time, you may want to designate areas within a country or dependency as regions. Such a filtered list makes the task of choosing the closest POP easier for your users.

For information on how to add regions, see Add, edit, or delete a region or Add, edit, or delete regions by command line.

**Points of Presence (POPs)**
The primary purpose of adding and editing POPs is to make updated access information available to your users.

POP information specifies how Connection Manager establishes a connection and how access numbers are displayed to users. For example, when you select the **Surcharge** POP setting, Connection Manager displays the POP in the **More access numbers** box instead of the **Access numbers** box. This feature can be used to differentiate automatically between remote-access numbers for corporate private networks and secure connections.

In preparing to create a POP, ask yourself the following questions:

- In what country or dependency does the POP exist? The list of countries and dependencies provided by CPS associates the Telephony Applications Program Interface (TAPI) country/dependency code with the POP. Although Phone Book Administrator is only available in English, French, German, Japanese, and Spanish, and the country/dependency list only appears in English, country/dependency names chosen from the English list appear in the same language as the Connection Manager profile.
- Does your POP belong to a region?

- What name will you give the POP?
- What is the area code for the access number?
- What is the access number?
- What POP settings will you need? For information about POP settings, see Service types (POP settings).
- What special dial-up settings will you need? Connection Manager and CPS must use the same dial-up entry name.
- What is the minimum baud for the POP?
- What is the maximum baud for the POP?
- What is the POP status: In Service or Not In Service? (When you publish to Phone Book Server, only in-service POPs are used to create or update the phone book files).

For information on how to create POPs, see Add or edit a POP or Add, edit, or delete POPs by command line.

**Publishing**

Before attempting to publish a phone book, make sure that you have set the Write permission for the FTP virtual directory on the computer running Phone Book Service. If the Write permission is not set, the phone book will not post.

You can publish multiple phone books to a server. However, every time you publish a phone book, CPS increments the version number of the phone book. If you publish the same phone book to multiple servers, the version numbers will be inaccurate. To publish a phone book to multiple servers, you can use file replication.

For information on how to publish a phone book, see Publish a phone book or Publish a phone book by command line.

# Use your existing database with Phone Book Administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use your existing database with Phone Book Administrator**

If you already maintain your phone book data in other databases, you can import both entire files and incremental updates to your phone books by using the command line interface to Phone Book Administrator (PBA).

You can write a query that exports the changes from your database into two text files: one for regions and one for POPs. Once you have these files, you can use the command-line interface to import changes.

You should keep a record in your database of the last time that you created import files for PBA. You can use this time stamp to create just the change (delta) files to be imported into PBA.

For information on creating new POP and region files, and exporting data, see Add, edit, or delete POPs by command line and Add, edit, or delete regions by command line.

**Notes**

- POPs and regions must be in separate files.
- Command-line region files do not match the region-file format imported or generated by the PBA user interface.
- You might want to create and publish your first phone book with the user interface to ensure that all authentication and configuration issues are worked out.
- If you have not yet posted your phone book, you can import up to 32,000 POPs. If, however, you have already posted, the limit is 6,000 POPs (combined adds, edits, and deletes) for each subsequent post.

# Example Implementations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Example implementations**

- ISP model
- Corporation model

# ISP model

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ISP model**

The following hypothetical situation illustrates how Connection Point Services (CPS) can support an expanding Internet service provider (ISP).

Awesome Computers, an ISP, is growing rapidly and must accommodate increasing numbers of subscribers. Many of the subscribers access Awesome Computers locally, but the ISP also has corporate customers whose employees want increased access capabilities while traveling. At the same time, Awesome Computers wants to replace its expensive, 800-number access charges with a remote-access solution that costs less.

To accommodate its growing business needs, the company plans to expand its existing remote access infrastructure and also offer a "Business Class" set of higher-quality Points of Presence (POPs) for demanding corporate customers. Awesome Computers joins an alliance of ISPs in order to provide coverage outside its service area.

Because many of the new subscribers are inexperienced computer users, Awesome Computers realizes that increasing the number of regional access numbers and providing remote POPs will generate a higher volume of calls to technical support. Faced with an immediate need to automate the process of updating subscribers with access numbers, the company decides to use CPS.

Awesome Computers installed Phone Book Service (PBS) and Phone Book Administrator (PBA) on a computer running an operating system in the Windows Server 2003 family. Personnel authorized by Awesome Computers use the administration tool to create a phone book that contains all of the trunk-line and connect information that the ISP offers for Internet access. When access numbers change, the personnel can quickly make the changes and publish them.

Publishing creates the updated phone book files and copies them to PBS. Awesome Computers keeps all the phone books for each corporate customer on the server running PBS. In addition, Awesome Computers has chosen to deploy a number of PBS servers within the infrastructure, thereby ensuring accessibility for users who connect to the Internet from various geographical regions.

With the file replication capability of the Windows Server 2003 family, the phone book can easily be automatically distributed from a master PBS server to a number of subordinate PBS servers.

Awesome Computers uses the Connection Manager Administration Kit wizard to create a customized Connection Manager profile that includes the company logo and tailored online Help. Awesome Computers then packages the profile and Connection Manager for distribution on compact disc or the World Wide Web.

After Connection Manager and the service profile have been installed on individual computers, CPS automatically downloads the most current list of local access phone numbers every time the subscriber connects to Awesome Computers.

# Corporation model

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Corporation model**

The following hypothetical situation illustrates how Connection Point Services (CPS) can help a corporation offer a cost-effective means of telecommuting.

Lucerne Publishing is a multinational company that spends a significant amount of money on maintaining toll-free numbers and long-distance dial-up fees so its employees can telecommute. The toll-free numbers are used by home-office employees who live outside the calling area and by employees on business trips. Home-office employees living inside the calling area use a local access number.

The company administrators decide to use CPS. They realize that by merging access numbers in their existing phone book with the access numbers of a large ISP they can save money on long distance charges while maintaining a high level of security.

Employees making phone calls from within the calling area of Lucerne Publishing connect to the company remote-access server. Employees residing or traveling outside the calling area of the corporation are able to call local ISP phone numbers, thereby connecting to the ISP remote-access server and avoiding long distance charges previously incurred by the maintenance of the toll-free numbers.

Lucerne Publishing uses the Connection Manager Administration Kit and Connection Manager to automatically differentiate between its local access numbers and the Points of Presence (POPs) provided by the ISP outside the local calling area. As a result, the difference between local access to the corporate private network and the extended remote access over the Internet using a secure connection is transparent to their employees. Connection Manager connects to the corporate private network when an employee uses a local access number and establishes a secure connection to the corporate private network when an employee uses an ISP phone number.

Lucerne Publishing and the ISP each maintain Phone Book Service (PBS) servers for their respective phone books. Each PBS server automatically downloads the most current list of local access phone numbers every time the employee connects to Lucerne Publishing.

# Monitor Activity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor activity**

- Phone Book Service activity
- Phone book history

# Phone Book Service activity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phone Book Service activity**

Counters in System Monitor produce statistical information on the processor, memory, cache, network, and other object types.

The following table lists PBServer Monitor counters. (Delta files contain only the current, incremental phone book updates, not the entire phone book file.)

| Name | Description |
|---|---|
| Delta Upgrade Hits | The total number of delta phone book files (incremental phone book changes only) that have been successfully downloaded. |
| Delta Upgrade Hits/Sec | The number of delta phone book files that are successfully downloaded per second. |
| Error Hits | The total number of requests that could not be satisfied by PBS. |
| Error Hits/Sec | The number of requests per second that could not be satisfied by PBS. |
| Full Upgrade Hits | The total number of full phone book files that have been successfully downloaded. |
| Full Upgrade Hits/Sec | The number of full phone book files that are successfully downloaded per second. |
| No Upgrade Hits | The total number of requests for phone book files that did not require an update to the phone book file on the client computer. |
| No Upgrade Hits/Sec | The number of requests for phone book files per second that do not require an update to the phone book file on the client computer. |
| Total Hits | The total number of requests for phone book files. |
| Total Hits/Sec | The number of requests per second for phone book files. |

**Notes**

- For information on how to monitor phone book activity, see Monitor the activity of Phone Book Service

- PBServer Monitor counters are not available unless PBS has been started. PBS starts automatically when a client computer requests a phone book. To simulate a client request, you can type:

  **http://localhost/pbserver/pbserver.dll?osarch=0&ostype=0&osver=1&cmver=1&lcid=1033&pbver=0&pb=***[phone_book_name]*

  in the **Address** box of your browser, and then press ENTER. For example, if the phone book name is Awesome, you would type:

  **http://localhost/pbserver/pbserver.dll?osarch=0&ostype=0&osver=1&cmver=1&lcid=1033&pbver=0&pb=Awesome**

# Phone book history

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phone book history**

Phone Book Administrator (PBA) maintains a log for each phone book that records its activity, including creating, deleting, or editing a POP, as well as posting the phone book to the server. If you move a phone book log file off the system or change its name, PBA creates a new file the next time activity occurs for that phone book.

Because these log files can become large enough over time to take up valuable hard disk space and to inhibit searching for useful information, you should save and archive them periodically.

For information on monitoring phone book history, see View the history of a phone book.

# Connection Point Services Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

- Important phone book files
- Example region files containing U.S. states
- Service types (POP settings)
- Example phone book history

# Important phone book files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Important phone book files**

Connection Point Services (CPS) needs services, software, and data files to update the database, post the database and data files to the server, query the server database, and download the appropriate file to Connection Manager.

Phone Book Administrator (PBA) creates the database and files for the server and then transfers the data using the File Transfer Protocol (FTP). Phone Book Service (PBS) compares phone book versions based on the query variables sent by Connection Manager and responds with the appropriate compressed phone book (.cab) file.

Phone book data elements are described in the following table.

| Element | Description |
|---|---|
| Cabinet (.cab) files | Compressed data files that contain phone book information for Connection Manager. There are two types of .cab files, Full.cab and Delta.cab files. |
| Phone book (.pbk) files | Text files that contain lists of POP information. Phone book files are compressed before they are transferred to the server. |
| Region (.pbr) file | Text file that helps categorize POPs. A POP can reference one of many geographical regions listed in the Region file. |
| Version (.ver) file | Text file that keeps track of the latest phone-book release. |
| PBServer.mdb | A database that contains updated query tables for your phone books. PBA transfers the database to the server as Newpb.mdb every time a phone book is published. |
| *Phone_book*.mdb | A database that contains POP information. |

# Example region files containing U.S. states

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Example region files containing U.S. states**

The format you use for the list of U.S. states depends on whether you are using Phone Book Administrator or the command line. You will need to use .pbr format (the first list) for Phone Book Administrator. If you are importing by command line, use the second list.

The following list shows the names of all fifty U.S. states, plus the District of Columbia, in .pbr format:

```
51
Alabama
Alaska
Arizona
Arkansas
California
Colorado
Connecticut
Delaware
District of Columbia
Florida
Georgia
Hawaii
Idaho
Illinois
Indiana
Iowa
Kansas
Kentucky
Louisiana
Maine
Maryland
Massachusetts
Michigan
Minnesota
Mississippi
Missouri
Montana
Nebraska
Nevada
New Hampshire
New Jersey
New Mexico
New York
North Carolina
North Dakota
Ohio
Oklahoma
Oregon
Pennsylvania
Rhode Island
South Carolina
South Dakota
Tennessee
Texas
Utah
Vermont
Virginia
Washington
West Virginia
Wisconsin
Wyoming
```

The following list shows the names of all fifty U.S. states, plus the District of Columbia, in command-line format.

```
1,Alabama
```

2,Alaska
3,Arizona
4,Arkansas
5,California
6,Colorado
7,Connecticut
8,Delaware
9,District of Columbia
10,Florida
11,Georgia
12,Hawaii
13,Idaho
14,Illinois
15,Indiana
16,Iowa
17,Kansas
18,Kentucky
19,Louisiana
20,Maine
21,Maryland
22,Massachusetts
23,Michigan
24,Minnesota
25,Mississippi
26,Missouri
27,Montana
28,Nebraska
29,Nevada
30,New Hampshire
31,New Jersey
32,New Mexico
33,New York
34,North Carolina
35,North Dakota
36,Ohio
37,Oklahoma
38,Oregon
39,Pennsylvania
40,Rhode Island
41,South Carolina
42,South Dakota
43,Tennessee
44,Texas
45,Utah
46,Vermont
47,Virginia
48,Washington
49,West Virginia
50,Wisconsin
51,Wyoming

# Service types (POP settings)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Service types (POP settings)**

Connection Point Services (CPS) uses a combination of eight service *types* (sometimes called POP settings) to determine how a connection is established.

Service types enable your users to filter for POPs specific to their needs.

Each service type has two possible options with values of 0 (zero) or 1 (one), and you convert these eight values into a decimal number to use in the **Service Flag** field of the POP file you plan to import.

If one of the flags for a POP is not set in the **Service Types** section of the .cms file in the Connection Manager profile, then even if you set the flag for this POP in Phone Book Administrator (PBA), a user may not be able to filter for the POP in **Access numbers** in the Connection Manager interface. For example, service types Custom 1 and Custom 2 are advanced and require manual editing of the Connection Manager service profile.

For each service type, the bitflag table (later in this topic) contains four fields. The first field is the bit number representing each option of any given service type. The second field contains integers representing the availability of each of the service type options. The third and fourth fields contain text that describes the service type option.

Options shown in bold are the default values for POPs in Connection Manager.

Immediately after the following table is an example of its use.

**Bitflag table for service types**

| Bit number | Option values | Option descriptions | Notes |
|---|---|---|---|
| 0 | **0**<br>1 | **Sign on**<br>Not Sign on | Users can connect to the POP to sign on to your service. If you set this option to 1, you must also set the same option to 1 in Connection Manager. |
| 1 | **0**<br>1 | **Not Sign up**<br>Sign up | Users can connect to the POP to sign up for your service. If you set this option to 1, you must also set the same option to 1 in Connection Manager. |
| 2 | **0**<br>1 | **Modem**<br>Not Modem | Users can connect to the POP by modem. |
| 3 | **0**<br>1 | **ISDN**<br>Not ISDN | Users can connect to the POP using an ISDN connection device. |
| 4 | **0**<br>1 | Not Custom 1<br>Custom 1 | The POP is assigned a custom service setting. You can specify custom service settings in the Connection Manager service profile. |
| 5 | **0**<br>1 | **Multicast**<br>Not Multicast | The POP supports IP multicast protocols. |
| 6 | **0**<br>1 | **Not Surcharge**<br>Surcharge | The access number for the POP appears in the **More access numbers** box in the Connection Manager **Phone Book** dialog box. You can charge your users an additional fee when they use this POP. |

| | | | |
|---|---|---|---|
| 7 | **0**<br><br>1 | Not Custom 2<br><br>Custom 2 | The POP is assigned a custom service setting. You can specify custom service settings in the Connection Manager service profile. |

To determine the correct decimal value to use for the **Service Flag** field in your command-line import file, create an eight-bit binary number where each bit has the value given in the table. Then, convert this binary number into a decimal number.

**Example**

The following list begins with the service type with the least mathematical significance. If you created a POP with these service types, the binary value would be 01111000, which converts to 120 in decimal notation:

Sign on (0)

Not Sign up (0)

Modem (0)

Not ISDN (1)

Custom 1 (1)

Not Multicast (1)

Surcharge (1)

Not Custom 2 (0)

A POP using this decimal number (shown in *italics*) as its service flag might look like this:

1, 1, 0, Redmond, 425, 555-0135, 9600, 56000, 0, *120*, ISP

# Example phone book history

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Example phone book history**

The example later in this topic shows the POP activity of one phone book. Each change is on a new line that begins with a date, and fields are separated by commas. Each line shows the following:

- Date and time of the change
- Action taken (for example, POP Add or POP Edit)
- Key field (POP name)
- Eleven fields representing the POP data:
    1. POP index
    2. Country/dependency
    3. Region identifier
    4. POP name
    5. Area code
    6. Access number
    7. Minimum baud rate
    8. Maximum baud rate
    9. Reserved flag
    10. Service flag
    11. Dial-up entry
- POP status (in or out of service)
- Date of change (repeated)
- Two reserved fields

The following example shows phone book history as it appears in the phone-book history text file. Note that the five field names in the first line of the text file (Date, Action, and so on) do not list every data field.

```
Date, Action, Key Field, Record ( ; delimited )

1/27/01 3:40:35 PM, POP Add, Albany, 3;1;0;Albany;912;555-0100;0;56000;0;48;ISP;1;1/27/01;;0

1/27/01 3:40:38 PM, POP Edit, Albuquerque, 5;1;0;Albuquerque;505;555-0121;0;56000;0;48;;1;1/27/01;;0

1/27/01 10:07:18 AM, POP Delete, Portland, 305;1;0;Portland;1;1;0;0;0;48;;1;1/27/01;;0
```

The following table supplies field names and their corresponding POP data for all of the fields in the first data line of the

example.

| Name of field | Data |
|---|---|
| Date and time | 1/27/01 3:40:35 PM |
| Action | POP Add |
| Key field | Albany |
| POP index | 3 |
| Country/dependency | 1 |
| Region ID | 0 |
| POP name | Albany |
| Area code | 912 |
| Access number | 555-0100 |
| Minimum baud rate | 0 |
| Maximum baud rate | 56000 |
| Reserved flag | 0 |
| Service flag | 48 |
| Dial-up entry | ISP |
| POP status | 1 |
| Date of change | 1/27/01 |
| Reserved | |
| Reserved | 0 |

**Notes**

- Phone book history files can also contain information on region and publishing activities.
- The dial-up entry name corresponds exactly with the dial-up entry created in the **Dial-up Networking Entries** pane in the Connection Manager Administration Kit Wizard.

# Connection Point Services Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

- Phone Book Service error messages
- Phone Book Administrator command-line error messages
- Common problems and their solutions

# Phone Book Service error messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phone Book Service error messages**

The following table describes Phone Book Service (PBS) events and the actions you can take to correct errors.

| MessageID | Type | Message | Action to perform |
|---|---|---|---|
| 30002 | Error | Unable to get query parameters | No action is required. If the problem persists, restart IIS Admin Service. |
| 30003 | Error | Error in sending HTML header | Make sure that IIS is running properly. If the problem persists, restart IIS Admin Service. |
| 30004 | Error | Error in sending HTML content | Make sure that IIS is running properly. If the problem persists, restart IIS Admin Service. |
| 30005 | Error | Cannot open file | Make sure the file exists and the file permissions are set properly. |
| 30010 | Error | Internal Error | Free up the memory currently in use by closing any nonessential applications. You might also want to increase the physical or virtual memory of your system. Also, make sure that IIS is running properly. |
| 30015 | Error | Failure when waiting for child thread to terminate | Restart IIS Admin service. |

# Phone Book Administrator command-line error messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phone Book Administrator command-line error messages**

The following table describes Phone Book Administrator error messages that you may encounter while importing files by command line and the actions you can take to correct errors. After you correct the problem, remember to re-import your files.

| Message | Action to perform |
|---|---|
| Error during file creation | Make sure that you have set the appropriate file permissions for the PBA and release directories. See Secure your Phone Book Administrator folder (NTFS only). |
| Invalid path for input file | Make sure you have typed the full path. |
| Problem while reading import file | Make sure the file is in plain text format. |
| Phone book passed as a command-line parameter does not exist | Make sure you have typed the correct phone book name as it appears in the Phone Book Administrator interface. |
| Error during command line processing | See Import.log in the application directory for details. |
| Wrong number of fields | Make sure that all of your POPs have 11 fields and all regions have two fields. |
| Length of input field exceeds allowed value | Edit the input field. For field restrictions, see Add, edit, or delete POPs by command line. |
| Numeric input field contains non-numeric characters | Replace with numeric characters. |
| Incomplete configuration data for publishing | Set the options for this phone book. See Set or edit phone book options by command line. |
| Duplicate region name found | Delete duplicate region name from the input file. |
| Region specified for phone book import does not exist | Add or import the missing region. See Add, edit, or delete regions by command line or Import a region list by command line. |
| Country or dependency specified for phone book import does not exist | Specify the correct TAPI code for country or dependency as defined in the phone book. You can find the default list in the Country.txt file. |
| Total length of the Area Code and Access Number fields exceeds allowed value | Edit area code and access number to contain a total of no more than 36 characters. |
| Wrong or missing command line option | At the command prompt, type **pbadmin /?** for a list of valid command-line parameters. |

# Common problems and their solutions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Common problems and their solutions**
What problem are you having?

- I need more help using Phone Book Administrator.
- I am unable to post to Phone Book Service (PBS).
- I am unable to import a phone book file or region file from the command line.
- No POPs appear in Connection Manager.
- When a client computer sent a query to Phone Book Service, the phone book did not download.
- Phone Book Administrator will not display particular characters.
- The application log available through Event Viewer does not contain messages about phone book activity.

**I need more help using Phone Book Administrator.**
**Solution:** In the **Phone Book Administrator** dialog box, open the **Help** menu, click **What's This?**, and then click the component you do not understand. A message will be displayed with information about the component. In other dialog boxes, click **?** and then click the component you do not understand.

**I am unable to post to Phone Book Service (PBS).**
**Possible solutions:**

Make sure that:

- The server is turned on.
- You are connected to the network.
- The PBS server has FTP installed and running.
- FTP is not set to **Allow only anonymous connections**.
- Write permission is turned on for the FTP virtual directory.
- The server name or address, user name, and password are typed correctly in the **Options** dialog box. Null passwords and passwords that contain spaces are not supported.
- You have logged on with an account that has permission to write to the hard drive on the server. These include FTP service, share, and file permissions (NTFS only).
- You have connectivity to the PBS host. Try PING and TRACERT to verify that the TCP/IP connection is correctly configured.
- The proxy settings in Internet Explorer are set correctly. The proxy might have changed, or someone might have reset it.
- Firewall issues are not interfering.

**See also:** Set permissions early; Set the Write permission for the FTP virtual directory; Set up FTP accounts for known users

**I am unable to import a phone book file or region file from the command line.**
**Possible solutions:**

Make sure that:

- You have the appropriate permissions to alter the phone book database.
- You typed the full path to the phone book file or region file. The file path cannot contain spaces.
- Your file is in plain-text format.
- You typed the correct phone book name as it appears in the Phone Book Administrator interface.
- All of the POPs have 11 fields. All regions have two fields.
- You did not exceed the limit for any of the input fields.
- None of the numeric input fields contain nonnumeric characters.
- The country, region, or dependency that you specified exists in the phone book.
- You used the correct import command. At the command prompt, type **pbadmin /?** for a list of valid command-line parameters.
- You used the correct import format for your region file. Command-line region files do not match the region-file format imported or generated by the Phone Book Administrator user interface.
- You did not try to import more than 6,000 POPs. If you have not yet posted your phone book, you can import up to 32,000 POPs. If, however, you have already posted, the limit is 6,000 POPs (combined adds, edits, and deletes) for each subsequent post.

**See also:** Set permissions early; Use your existing database with Phone Book Administrator; Import a region list by command line; Phone Book Administrator command-line error messages

**No POPs appear in Connection Manager.**
**Possible solutions:**

Make sure that:

- The profile created with the Connection Manager Administration Kit wizard is set up correctly to use your phone books.
- No files are missing. If they are, use your backup copies to restore them.
- The name of the phone book in Phone Book Administrator matches the name typed in when the profile was built using the Connection Manager Administration Kit wizard.
- You did not exceed the limit for any of the input fields. (For a table that describes limits, see Add, edit, or delete POPs by command line.)
- None of the numeric input fields contain nonnumeric characters.
- You specified the correct country or dependency for the POP.
- You and your users have the appropriate version of Internet Explorer installed. Phone book downloads require that Internet Explorer 5.01, Internet Explorer 5.5, or Internet Explorer 6 be installed on the computer. Internet Explorer does not have to be set as the default browser.

**See also:** Back up your data; Use your existing database with Phone Book Administrator

**When a client computer sent a query to Phone Book Service, the phone book did not download.**
**Possible solutions:**

Make sure that:

- The server is turned on.
- The client computer is connected to the network.
- The name of the phone book in Phone Book Administrator (PBA) matches the name in Connection Manager.
- The user gave the phone book enough time to download.
- WWW is currently running on the server. Even though WWW is installed, it may not be running.
- The proxy settings in the client computer are set correctly.
- The current network configuration still passes your previous tests for production readiness.
- The security attribute is set properly on the phone book update files. It should not be set to **Read Only**.
- Firewall issues are not interfering with phone book downloads.
- Your users have the appropriate version of Internet Explorer installed. Phone book downloads require that Internet Explorer 5.01, Internet Explorer 5.5, or Internet Explorer 6 be installed on the computer. Internet Explorer does not have to be set as the default browser.
- Anonymous access is enabled for the Phone Book Service virtual root directory. Connection Manager requests phone books through the anonymous account.
- You have looked at the application log in Event Viewer on the PBS host server for error messages. If Event Viewer messages have been enabled for PBS and the phone book file is missing, you will see error 30005: "Cannot open file."
- The phone book file is present on the host server, and that the host server is accessible. The phone book file might be missing from the host server, or access to the host server might be denied.
- The phone book is posted to the phone book server. A profile might have been configured to ask for a phone book, but the administrator might not have posted that phone book to the server.
- The phone book has been posted to the production environment. The phone book might have been posted to a staging server but not replicated to the production environment. Try editing a POP in the phone book and then posting the file again. (You can edit the POP to repeat the existing data.) When you edit a POP, PBA changes the version number of the phone book. PBA does not post to the host server unless the phone book version number changes.

If your Connection Manager profile uses Connection Manager 1.3, have the user enable logging and check the Connection Manager log file for more information about the problem.

**See also:** Set permissions early; Phone Book Service error messages; Test for production readiness; Connection Manager Administration Kit

**Phone Book Administrator will not display particular characters.**
**Cause:** Phone Book Administrator does not support Unicode.

**Solution:** Set **Language for non-Unicode programs** to match the language that you want to display:

1. Open Regional and Language Options. To open Regional and Language Options, click **Start**, click **Control Panel**, and then double-click **Regional and Language Options**.
2. Click the **Advanced** tab. In the **Language for non-Unicode programs** section, click the appropriate language version.

3. Click **Apply**.

   This may cause your system to restart.

**The application log available through Event Viewer does not contain messages about phone book activity.**
**Cause:** PBS cannot, by default, write to the application log.

**Solution:** On computers that are running a Windows Server 2003 family operating system, applications that are members of Guests cannot, by default, write to the application log. You can change this behavior, but you should first consider whether the security implications of doing so outweigh the benefits. To change this behavior, you can set the following registry key to a value of zero:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\EventLog\Application\RestrictGuestAccess**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

# Virtual Private Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private networks**
The Microsoft® Windows Server 2003 Routing and Remote Access service provides virtual private network (VPN) services for remote access and router-to-router VPN connections by using either Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP) with Internet Protocol security (IPSec).

- Before installing a VPN server, see VPN Checklists.
- To find features that have been moved in the Windows Server 2003 family, see New ways to do VPN tasks.
- For tips about using VPNs, see VPN best practices.
- For general background information, see VPN Concepts.
- For problem-solving instructions, see VPN Troubleshooting.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# VPN Checklists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

This section covers:

- Checklist: Installing and configuring a VPN server
- Checklist: Installing and configuring a PPTP server
- Checklist: Installing and configuring an L2TP server

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Checklist: Installing and configuring a VPN server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring a VPN server**

| Step | Reference |
|---|---|
| Review key concepts. | Virtual Private Networks Overview |
| Install the hardware. | Manufacturer's documentation |
| Verify the compatibility of all hardware to be installed in the computer running a Windows Server 2003 operating system. | The Compatible Hardware and Software section at Support resources |
| Verify that the hardware is successfully installed. | Device Manager |
| Install the Routing and Remote Access service, and configure the server using the Routing and Remote Access Server Setup Wizard. | Enable the Routing and Remote Access service |
| Configure dial-in properties and remote access policies for dial-in permission, authentication, and encryption settings. | Introduction to remote access policies |
| Configure packet filters on your firewall. | VPN servers and firewall configuration |
| (Optional) Configure the virtual private network (VPN) server as a RADIUS client. | Use RADIUS authentication |

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Checklist: Installing and configuring a PPTP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring a PPTP server**

| Step | Reference |
|---|---|
| Review key concepts. | Virtual Private Networks Overview |
| Install the hardware. | Manufacturer's documentation |
| Verify the compatibility of all hardware to be installed in the computer running a Windows Server 2003 operating system. | The Compatible Hardware and Software section at Support resources |
| Verify that the hardware is successfully installed. | Device Manager |
| Install the Routing and Remote Access service, and configure the server using the Routing and Remote Access Server Setup Wizard. | Enable the Routing and Remote Access service |
| Configure packet filters on your firewall. | VPN servers and firewall configuration |
| Configure dial-in properties and remote access policies for dial-in permission, authentication, and encryption settings. | Introduction to remote access policies |

**Notes**

- Point-to-Point Tunneling Protocol (PPTP) is automatically installed.
- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Checklist: Installing and configuring an L2TP server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring an L2TP server**

| Step | Reference |
|---|---|
| Review key concepts. | Virtual Private Networks Overview |
| Install the hardware. | Manufacturer's documentation |
| Verify the compatibility of all hardware to be installed in the computer running a Windows Server 2003 operating system. | The Compatible Hardware and Software section at Support resources |
| Verify that the hardware is successfully installed. | Device Manager |
| Install the Routing and Remote Access service, and configure the server using the Routing and Remote Access Server Setup Wizard. | Enable the Routing and Remote Access service |
| Configure packet filters on your firewall. | VPN servers and firewall configuration |
| Install computer certificates on the Layer Two Tunneling Protocol (L2TP) server computer and L2TP client computers. | Computer certificates for L2TP/IPSec VPN connections;Network access authentication and certificates |
| Configure dial-in properties and remote access policies for dial-in permission, authentication, and encryption settings. | Introduction to remote access policies |

**Notes**

- L2TP is automatically installed.
- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# New ways to do VPN tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring virtual private networks (VPNs). The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT 4.0 or Windows NT 4.0 with the Routing and Remote Access Service (RRAS).

**Notes**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.
- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

Please choose your upgrade path:

- Windows NT 4.0 to a product in the Windows Server 2003 family
- Windows NT 4.0 with RRAS to a product in the Windows Server 2003 family

**Windows NT 4.0 to a product in the Windows Server 2003 family**

| If you want to | In Windows NT 4.0 use | In this version of Windows |
|---|---|---|
| Install the Point-to-Point Tunneling Protocol (PPTP) | **Protocols** tab of **Network** in Control Panel | PPTP is automatically installed. |
| Set the maximum number of PPTP connections | **Protocols** tab of **Network** in Control Panel | Use Properties of **Ports** in Routing and Remote Access. |
| Set PPTP packet filtering | The **Advanced IP Addressing** dialog box from the properties of the TCP/IP protocol (**Protocols** tab of **Network** in Control Panel) | PPTP packet filtering is set by default when you run the Routing and Remote Access Setup Wizard or set manually when you use Properties of the IP interface in **General** under **IP Routing** in Routing and Remote Access. |

**Windows NT 4.0 with RRAS to a product in the Windows Server 2003 family**

| If you want to | Windows NT 4.0 with RRAS use | In this version of Windows |
|---|---|---|
| Install the Point-to-Point Tunneling Protocol (PPTP) | **Protocols** tab of **Network** in Control Panel | PPTP is automatically installed. |
| Set the maximum number of PPTP connections | **Protocols** tab of **Network** in Control Panel | Use Properties of **Ports** in Routing and Remote Access. |

| | | |
|---|---|---|
| Set PPTP packet filtering | Properties of the IP interface in **Summary** under **IP Routing** in Routing and RAS Admin | PPTP packet filtering is set by default when you run the Routing and Remote Access Setup Wizard or set manually when you use Properties of IP interface in **General** under **IP Routing** in Routing and Remote Access. |

# VPN best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Use the virtual private network (VPN) server to allocate IP address leases to remote access clients.**

  If your VPN server handles more than 20 concurrent remote access connections, then use the VPN server to allocate IP address leases to remote access clients. When you create the IP address pool on the VPN server, be certain not to allocate addresses already in use by DHCP servers on your network.

  If you did not install a DHCP server, and you have a single subnet with computers configured with static IP addresses, configure the VPN server with an IP address pool that is a subset of addresses for the subnet to which the VPN server is attached. For more information, see Create a static IP address pool.

- **Use strong authentication.**

  - Use strong passwords that are more than 8 characters long and that contain a mixture of uppercase and lowercase letters, numbers, and permitted punctuation. Do not use passwords based on names or words. Strong passwords are more resistant to a dictionary attack, where an unauthorized user attempts to determine a password by sending a series of commonly used names and words.

  - Although EAP-TLS works with registry-based user certificates, it is highly recommended that you use only EAP-TLS with smart cards for remote access VPN connections. Because smart cards are distributed to trusted users and require a personal identification number (PIN), they are more secure than registry-based user certificates.

  - If you are using password-based user authentication, then use MS-CHAP version 2. You can obtain the latest MS-CHAP updates for VPN clients running Windows NT 4.0, Windows 95, or Windows 98 from Microsoft. For more information, see MS-CHAP version 2.

- **Use strong encryption.**

  Use the strongest level of encryption that your situation allows. The levels of encryption that are available when configuring the profile properties of a remote access policy are: Basic, Strong, and Strongest. For more information, see Configure encryption and Elements of a remote access policy.

# VPN Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**
This section provides general background information about the virtual private network (VPN) services included with the Windows Server 2003 family:

- Virtual Private Networks Overview
- Understanding Virtual Private Networks
- Deploying Virtual Private Networks
- VPN Resources

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Virtual Private Networks Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private networks overview**
This section covers:

- Security issues for VPN
- Introduction to virtual private networks
- New features for virtual private networks
- VPN Connections
- Types of Virtual Private Networks

# Security issues for VPN

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for VPN**

It is important to follow best practices for security when using VPN servers on your network. For more information, see Best practices for security.

If your VPN servers are configured as Remote Authentication Dial-In User Service (RADIUS) clients, see Security information for IAS.

The following are known security issues for VPN connections:

- **Many authentication methods are too weak to provide adequate security for most organizations.**

  Many organizations use authentication methods that expose their network to a variety of security attacks. The most secure method of authentication is Extensible Authentication Protocol-Transport Level Security (EAP-TLS) when used in conjunction with smart cards. However, EAP-TLS and smart cards require a public key infrastructure (PKI), which can be complicated to deploy.

**Recommendations:**

- When you use password-based authentication, enforce strong password policies on your network to make dictionary attacks more difficult. For more information, see Strong passwords.

- Make sure that Password Authentication Protocol (PAP), Shiva Password Authentication Protocol (SPAP), and Challenge Handshake Authentication Protocol (CHAP) are disabled. PAP, SPAP, and CHAP are disabled by default on the profile of a remote access policy.

- Disable Microsoft Challenge Handshake Authentication Protocol (MS-CHAP). MS-CHAP is enabled by default on the profile of a remote access policy. Before you disable MS-CHAP, ensure that all of the clients of your access server are capable of using MS-CHAP version 2 (MS-CHAP v2). If your server is not running Microsoft® Windows Server 2003 , you will need the latest service packs and dial-up networking updates in order to support MS-CHAP v2.

- If you do not disable MS-CHAP, disable LAN Manager authentication for MS-CHAP. LAN Manager authentication for MS-CHAP is enabled by default on clients running earlier versions of Windows.

- To disable LAN Manager authentication, set the registry value **Allow LM Authentication** to **0** at the following registry key:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Note**

- You can change an entry in the registry using Registry Editor. For more information, see Registry Editor and Change a value.

- Make sure that MS-CHAP v2 is enabled. MS-CHAP v2 is enabled by default on the profile of a remote access policy.

- For wireless clients running Windows XP, enable Protected Extensible Authentication Protocol (PEAP) with Extensible Authentication Protocol (EAP)-MS-CHAP v2. PEAP with EAP-MS-CHAP v2 provides mutual authentication: the IAS server is authenticated with a certificate, while user authentication is accomplished with password-based credentials. For more information, see New features for IAS. For an example of a remote access policy that uses PEAP, see Wireless access with secure password authentication.

- Enable EAP-TLS. EAP-TLS is disabled by default on the profile of a remote access policy. When you use the EAP-TLS authentication protocol, you must install a computer certificate on the IAS server. For client and user authentication, you can install a certificate on the client computer, or you can use smart cards. Before you deploy certificates, you must design the certificate with the correct requirements. For more information, see Network access authentication and certificates and Computer certificates for certificate-based authentication.

- **Remote access account lockout can deny network access to authorized users.**

  When remote access account lockout is enabled, if a malicious user attempts a dictionary attack with the logon name of an authorized user, both the malicious user and the authorized user are locked out of the account until the account lockout threshold is reached.

**Recommendations:**

- Be cautious when defining account lockout policy. For more information, see Password Best practices.

- When the lockout count for a user account is reset to 0, due to either a successful authentication or an automatic reset, the registry subkey for the user account is deleted. To manually unlock a user account (before the account lockout threshold is automatically reset), delete the following registry subkey that corresponds to the account name of the user:

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout***DomainName:UserName*

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Notes**

- You can change an entry in the registry using Registry Editor. For more information, see Registry Editor and Change a value.

- Remote access account lockout is not related to the **Account locked out** setting on the **Account** tab on the properties of a user account.

  Remote access account lockout does not distinguish between malicious users who attempt to access your intranet and authentic users who attempt remote access but have forgotten their current passwords. Users who have forgotten their current password typically try the passwords that they remember and, depending on the number of attempts and the MaxDenials setting, might be locked out of their accounts.

- **The default domain functional level setting (Windows 2000 mixed) can create security holes.**

  You should use only servers that are running Windows 2000 or Windows Server 2003 , if possible. Using servers running Routing and Remote Access in conjunction with servers that are running Windows NT® 4.0 and either Remote Access Service (RAS) or Routing and Remote Access Service (RRAS) decreases your network security. For more information, see Domain and forest functionality and Member server in a domain.

**Additional recommendations:**

- **Allow the VPN server or a DHCP server to assign IP address leases to VPN clients.**

  Do not allow VPN clients to specify their own IP address. The method by which the VPN server assigns addresses to remote access clients (either using addresses that the VPN server obtains from a DHCP server or using addresses from a specified range of addresses that you configure) is set when you run the Routing and Remote Access Server Setup Wizard. For more information, see L2TP-based remote access VPN deployment and DHCP.

- **Always use Layer Two Tunneling Protocol/Internet Protocol security (L2TP/IPSec) connections for the strongest encryption.**

  Use smart cards for the most secure authentication and certificates when smart cards are not practical.

- **Use certificates for Internet Key Exchange (IKE) negotiations.**

  Do not use preshared keys, as they are susceptible to dictionary attacks.

For more information about security considerations and best practices, see:

- Network access authentication and certificates
- VPN servers and firewall configuration
- VPN best practices
- VPN Security
- Remote access VPN security considerations
- Router-to-router VPN security considerations
- Deploying Certificate-based Authentication for VPN Connections
- Event Viewer

- Security information for IAS
- Security information for remote access
- Security information for remote access

# Introduction to virtual private networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private networks**

A virtual private network (VPN) is the extension of a private network that encompasses links across shared or public networks like the Internet. With a VPN, you can send data between two computers across a shared or public network in a manner that emulates a point-to-point private link. Virtual private networking is the act of creating and configuring a virtual private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information, which allows the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The link in which the private data is encapsulated and encrypted is a virtual private network (VPN) connection.

The following illustration shows the logical equivalent of a VPN connection.




Users working at home or on the road can use VPN connections to establish a remote access connection to an organization server by using the infrastructure provided by a public network such as the Internet. From the user's perspective, the VPN is a point-to-point connection between the computer (the VPN client) and an organization server (the VPN server). The exact infrastructure of the shared or public network is irrelevant because it appears logically as if the data is sent over a dedicated private link.

Organizations can also use VPN connections to establish routed connections with geographically separate offices or with other organizations over a public network such as the Internet while maintaining secure communications. A routed VPN connection across the Internet logically operates as a dedicated WAN link.

With both remote access and routed connections, an organization can use VPN connections to trade long-distance dial-up or leased lines for local dial-up or leased lines to an Internet service provider (ISP).

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. For more information about feature availability on Windows Server 2003, Web Edition, see Overview of Windows Server 2003, Web Edition. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

There are two types of Point-to-Point Protocol (PPP)-based VPN technology in the Microsoft Windows 2003 family:

1. Point-to-Point Tunneling Protocol (PPTP)

   PPTP uses user-level PPP authentication methods and Microsoft Point-to-Point Encryption (MPPE) for data encryption.

2. Layer Two Tunneling Protocol (L2TP) with Internet Protocol security (IPSec)

   L2TP uses user-level PPP authentication methods and computer-level certificates with IPSec for data encryption, or IPsec in tunnel mode, in which IPsec itself provides encapsulation (for IP traffic only).

For more information on the types of VPN connections, see VPN Connections. For information on how to design and deploy VPN connections, see Deploying Virtual Private Networks.

# New features for virtual private networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features for virtual private networks**

The Microsoft® Windows® Server 2003 family provides the following new features for virtual private networks (VPNs):

- Network address translation (NAT) transparency

  VPN servers running Windows Server 2003 support Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec) traffic that originates from VPN clients behind NATs. For this feature to function properly, the client computer must support the following IPSec Protocol Working Group Internet drafts:

  - Negotiation of NAT-Traversal in the IKE (draft-ietf-ipsec-nat-t-ike-02.txt).
  - UDP Encapsulation of IPsec Packets (draft-ietf-ipsec-udp-encaps-02.txt).

**Important**

- An IPSec NAT-T deployment for Windows that includes VPN servers that are located behind network address translators is not recommended. When a server is behind a network address translator, and the server uses IPSec NAT-T, unintended behavior might occur because of the way network address translators translate network traffic.

- VPN deployment with Network Load Balancing

  VPN servers running Windows Server 2003 support VPN deployment in conjunction with Network Load Balancing to create high-availability VPN solutions. Support is included for Point-to-Point Tunneling Protocol (PPTP) and L2TP/IPSec VPN solutions with Network Load Balancing.

- NetBIOS over TCP/IP (NetBT) name resolution proxy

  When a remote access VPN client connects to a VPN server, it relies on DNS or WINS servers on the target network for name resolution.

  DNS and WINS servers are common in organizations that use DNS for host name resolution or WINS for NetBIOS name resolution. However, in a small office or home office environment, these servers might not be present. In this case, the names of the computers to which communications are attempted need to be resolved by another means for successful communication to occur.

  NetBIOS names can be resolved, without the use of a WINS or DNS server, by a VPN server running Windows Server 2003 because Windows Server 2003 operating systems include the NetBT proxy. With NetBT proxy, connected VPN clients and the nodes on network segments to which the VPN server is attached can resolve each other's NetBIOS names, as follows:

  1. When a VPN client needs to resolve a name to an IP address without a configured DNS or WINS server, it sends a NetBIOS Name Query packet to the VPN server. When a node on a network segment that is attached to the VPN server needs to resolve a name to an IP address without a configured DNS or WINS server, it sends a NetBIOS Name Query packet as a broadcast on the network segment.

  2. The VPN server receives the NetBIOS Name Query packet, checks its local cache of resolved NetBIOS names and, if the name is not found, forwards the NetBIOS Name Query as a NetBIOS broadcast on all attached interfaces, except the logical interface that is connected to all VPN clients.

  3. The node that has registered the NetBIOS name sends a positive NetBIOS Name Query Response to the VPN

server.

4. The VPN server receives the positive NetBIOS Name Query Response and then forwards it on the interface from which the NetBIOS Name Query packet originated to the sending node.

The result is that network nodes on network segments that are attached to the VPN server (and all connected VPN clients) can automatically resolve each other's names without a DNS or WINS server.

- Preshared key configuration for L2TP connections

  The Windows Server 2003 family supports both computer certificates and a preshared key as authentication methods to establish an IPSec security association for L2TP connections. A preshared key is a string of text that is configured on both the VPN client and VPN server. Preshared key is a relatively weak authentication method, therefore it is recommended that you use preshared key authentication only while your public key infrastructure (PKI) is being deployed to obtain computer certificates or when a preshared key is required by VPN clients. You can enable the use of a preshared key for L2TP connections and specify the preshared key from the **Security** tab on the properties of a server running Routing and Remote Access. For more information, see View properties of the remote access server.

  Preshared key authentication is also supported by remote access VPN clients running Windows XP. You can enable preshared key authentication and configure a preshared key from **IPSec settings** on the **Security** tab on the properties of a VPN connection in Network Connections. For more information, see Configure a pre-shared key on a remote access client.

  Preshared key authentication is also supported for site-to-site VPN connections between computers running a Windows Server 2003 operating system. You can enable preshared key authentication and configure a preshared key for demand-dial interfaces from **IPSec settings** on the **Security** tab on the properties of a demand-dial interface in Routing and Remote Access. For more information, see Configure a pre-shared key for a demand-dial routing interface and Configure a pre-shared key on a remote access server.

The following features for virtual private networks were new in Windows 2000:

- L2TP

  In addition to PPTP, the the Windows Server 2003 family includes the industry standard L2TP, which is used in conjunction with IPSec to create secure virtual private network connections. The use of L2TP and IPSec is referred to as L2TP/IPSec.

- Remote access policies

  Remote access policies are a set of conditions and connection settings that allow network administrators more flexibility in setting remote access permissions and connection restraints. With remote access policies, you can force the use of strong authentication and encryption for VPN users and use a different set of authentication and encryption constraints for dial-up users.

  For more information, see Remote Access Policies.

- MS-CHAP version 2

  Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) version 2 significantly strengthens the security for the passing of security credentials and the generation of encryption keys during the negotiation of a remote access connection. MS-CHAP version 2 was specifically designed for authenticating VPN connections.

  For more information, see MS-CHAP version 2.

- Extensible Authentication Protocol

  Extensible Authentication Protocol (EAP) allows new authentication methods to be used, a feature that is especially

important for the deployment of security based on smart cards. EAP is an architecture that allows other authentication modules to plug into PPP in Windows 2000, Windows XP, and products in the Windows Server 2003 family. Both EAP-MD5 CHAP and EAP-TLS (used for smart card and certificate-based authentication) are supported. You can configure your VPN server to handle authentication, or the VPN server can be configured to forward authentication requests to RADIUS servers.

For more information, see EAP and Internet Authentication Service.

- Remote access account lockout

  Remote access account lockout is a security feature that denies access after a configured number of failed authentication attempts. Remote access account lockout is designed to help prevent dictionary attacks. A dictionary attack occurs when an unauthorized user attempts to log on by using a known user name and a list of common words as the password. Remote access account lockout is disabled by default.

  For more information, see Remote access account lockout.

# VPN Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN connections**
Creating a VPN connection is very similar to establishing a point-to-point connection by using dial-up connections and demand-dial routing. There are two types of VPN connections:

- Remote access VPN connection
- Router-to-router VPN connection

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Remote access VPN connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access VPN connection**
A remote access client (a single user computer) makes a remote access VPN connection that connects to a private network. The VPN server provides access to the entire network to which the VPN server is attached. The packets sent from the remote client across the VPN connection originate at the remote access client computer.

The remote access client (the VPN client) authenticates itself to the remote access server (the VPN server) and, for mutual authentication, the server authenticates itself to the client.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

Computers running Microsoft Windows 95, Windows 98, Windows NT® 4.0, Windows 2000, Windows Millennium Edition, Windows XP, or a Windows Server 2003 operating system can create remote access VPN connections to a VPN server running Windows Server 2003. VPN clients may also be any non-Microsoft Point-to-Point Tunneling Protocol (PPTP) client or Layer Two Tunneling Protocol (L2TP) client with Internet Protocol security (IPSec).

For more information on PPTP and L2TP, see Understanding VPN Tunneling Protocols. For more information on deploying remote access VPN connections, see Deploying VPNs for Remote Access. For an example implementation of a remote access VPN connection, see Virtual Private Network Implementation Examples.

# Router-to-router VPN connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Router-to-router VPN connection**

A router makes a router-to-router VPN connection that connects two portions of a private network. The VPN server provides a routed connection to the network to which the VPN server is attached. On a router-to-router VPN connection, the packets sent from either router across the VPN connection typically do not originate at the routers.

The calling router (the VPN client) authenticates itself to the answering router (the VPN server) and, for mutual authentication, the answering router authenticates itself to the calling router.

Computers running Microsoft® Windows NT® 4.0 with the Routing and Remote Access Service (RRAS), Windows® 2000 Server, or a Windows 200 Server or a Windows Server 2003 operating system can create router-to-router VPN connections. VPN clients may also be any non-Microsoft Point-to-Point Tunneling Protocol (PPTP)-capable router or Layer Two Tunneling Protocol (L2TP)-capable router with Internet Protocol security (IPSec).

For more information on PPTP and L2TP, see Understanding VPN Tunneling Protocols. For more information on deploying router-to-router VPN connections, see Deploying Router-to-Router VPNs. For an example implementation of router-to-router VPN connections, see Virtual Private Network Implementation Examples. For information about creating router-to-router connections by using IPSec Encapsulating Security Payload (ESP) tunnel mode, see Tunnel mode.

# Types of Virtual Private Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Types of virtual private networks**

You can use VPN connections whenever you need a secure point-to-point connection to connect users or networks. Typical VPN connections are either Internet-based or intranet-based. This section covers:

- Internet-based VPNs
- Intranet-based VPNs

# Internet-based VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet-based VPNs**
By using an Internet-based VPN connection, you can avoid long-distance and 1-800 telephone charges while taking advantage of the global availability of the Internet.

**Remote access over the Internet**
Rather than making a long distance or 1-800 call to a corporate or outsourced network access server (NAS), a remote access client can call a local ISP. By using the established physical connection to the local ISP, the remote access client initiates a VPN connection across the Internet to the organization's VPN server. Once the VPN connection is created, the remote access client can access the resources of the private intranet.

The following illustration shows remote access over the Internet.




For more information on deploying remote access VPN connections across the Internet, see Deploying VPNs for Remote Access. For an example implementation of a remote access VPN connection across the Internet, see Virtual Private Network Implementation Examples.

**Connecting networks over the Internet**
When networks are connected over the Internet, a router forwards packets to another router across a VPN connection. This is known as a router-to-router VPN connection. To the routers, the VPN operates as a data-link layer link.

The following illustration shows connecting networks over the Internet.




**Using dedicated WAN links**
Rather than using an expensive long-distance dedicated WAN link between branch offices, the branch office routers are connected to the Internet by using local dedicated WAN links to a local ISP. A router-to-router VPN connection is then initiated by either router across the Internet. Once connected, routers can forward directed or routing protocol traffic to each other by using the VPN connection.

**Using dial-up WAN links**
Rather than making a long distance or 1-800 call to a corporate or outsourced NAS, a branch office router can call a local ISP. By using the established connection to the local ISP, a router-to-router VPN connection is initiated by the branch office router to the corporate office router across the Internet. The corporate office router acts as a VPN server and must be connected to a local ISP by using a dedicated WAN link.

It is possible to have both the corporate office and the branch office connected to the Internet by using a dial-up WAN link. However, this is only feasible if the ISP supports demand-dialing routing to customers--the ISP calls the customer router when an IP datagram is to be delivered to the customer. Demand-dial routing to customers is not widely supported by ISPs.

For more information on deploying router-to-router VPN connections across the Internet, see Deploying Router-to-Router VPNs. For an example implementation of router-to-router VPN connections across the Internet, see Virtual Private Network Implementation Examples.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Intranet-based VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Intranet-based VPNs**
The intranet-based VPN connection takes advantage of IP connectivity on an organization intranet.

**Remote access over an intranet**
On some organization intranets, the data of a department, such as a human resources department, is so sensitive that the department's network is physically disconnected from the rest of the organization's intranet. While this protects the department's data, it creates information accessibility problems for those users who are not physically connected to the separate network.

With a VPN connection, the department's network is physically connected to the organization intranet but separated by a VPN server. The VPN server does not provide a direct routed connection between the organization intranet and the department's network. Users on the organization intranet with the appropriate user rights can establish a remote access VPN connection with the VPN server and access the protected resources of the sensitive department's network. Additionally, all communication across the VPN connection is encrypted for data confidentiality. For those users who do not have the user rights to establish a VPN connection, the department's network is hidden from view.

The following illustration shows remote access over an intranet.




For more information on deploying remote access VPN connections across an intranet, see Deploying VPNs for Remote Access.

**Connecting networks over an intranet**
You can also connect two networks over an intranet by using a router-to-router VPN connection. Organizations with departments in separate locations, whose data is highly sensitive, may use a router-to-router VPN connection to communicate with each other.

For example, the finance department might need to communicate with the human resources department to exchange payroll information. The finance department and the human resources department are connected to the common intranet with computers that can act as VPN routers. Once the VPN connection is established, users on computers on either network can exchange sensitive data across the corporate intranet.

The following illustration shows connecting networks over an intranet.




For more information on deploying router-to-router VPN connections across an intranet, see Deploying Router-to-Router VPNs.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Understanding Virtual Private Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding virtual private networks**
This section covers:

- Components of virtual private networks
- Properties of VPN connections
- A dial-up router-to-router VPN connection
- VPN servers and firewall configuration
- VPN servers and Microsoft ISA Server
- Network access authentication and certificates
- VPN Tunneling Protocols
- VPN Security

# Components of virtual private networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Components of virtual private networks**

A VPN connection includes the following components:

- VPN server

  A computer that accepts VPN connections from VPN clients.

- VPN client

  A computer that initiates a VPN connection to a VPN server. A VPN client can be an individual computer or a router.

- Tunnel

  The portion of the connection in which your data is encapsulated.

- VPN connection

  The portion of the connection in which your data is encrypted. For typical secure VPN connections, the data is encrypted and encapsulated along the same portion of the connection.

  **Note**

  - It is possible to create a tunnel and send the data through the tunnel without encryption. This is not a VPN connection because the private data is sent across a shared or public network in an unencrypted and easily readable form.

- Tunneling protocols

  Protocols that are used to manage tunnels and encapsulate private data. Data that is tunneled must also be encrypted to be a VPN connection. The Windows Server 2003 family includes the PPTP and L2TP tunneling protocols. For more information, see Point-to-Point Tunneling Protocol and Layer Two Tunneling Protocol.

- Tunneled data

  Data that is usually sent across a private point-to-point link.

- Transit internetwork

  The shared or public network crossed by the encapsulated data. For the Windows Server 2003 family, the transit internetwork is always an IP internetwork. The transit internetwork can be the Internet or a private IP-based intranet.

The following illustration shows the components of a virtual private network.




**Notes**

- Typically, the tunnel and the VPN connection are along the same portion of the connection. However, in a compulsory

tunneling configuration, the tunnel (the encapsulation) and the VPN connection (the encryption) are not defined along the same portion of the connection.

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Properties of VPN connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Properties of VPN connections**
VPN connections that use PPTP and L2TP/IPSec have the following properties:

- Encapsulation
- Authentication
- Data encryption

**Encapsulation**
With VPN technology, private data is encapsulated with a header that provides routing information, which allows the data to traverse the transit internetwork. For examples of encapsulation, see Understanding VPN Tunneling Protocols.

**Authentication**
Authentication for VPN connections takes three different forms:

1. User-level authentication by using PPP authentication

   To establish the VPN connection, the VPN server authenticates the VPN client that is attempting the connection by using a Point-to-Point Protocol (PPP) user-level authentication method and verifies that the VPN client has the appropriate authorization. If mutual authentication is used, the VPN client also authenticates the VPN server, which provides protection against computers that are masquerading as VPN servers.

2. Computer-level authentication by using IKE

   To establish an IPSec security association, the VPN client and the VPN server use the Internet Key Exchange (IKE) protocol to exchange either computer certificates or a preshared key. In either case, the VPN client and server authenticate each other at the computer level. Computer certificate authentication is highly recommended as it is a much stronger authentication method. For more information, see Internet Key Exchange. Computer-level authentication is only done for L2TP/IPSec connections.

3. Data origin authentication and data integrity

   To verify that the data sent on the VPN connection originated at the other end of the connection and was not modified in transit, the data contains a cryptographic checksum based on an encryption key known only to the sender and the receiver. Data origin authentication and data integrity are only available for L2TP/IPSec connections.

**Data encryption**
To ensure confidentiality of the data as it traverses the shared or public transit internetwork, the data is encrypted by the sender and decrypted by the receiver. The encryption and decryption processes depend on both the sender and the receiver using a common encryption key.

Intercepted packets sent along the VPN connection in the transit internetwork are unintelligible to anyone who does not have the common encryption key. The length of the encryption key is an important security parameter. You can use computational techniques to determine the encryption key. However, such techniques require more computing power and computational time as the encryption keys get larger. Therefore, it is important to use the largest possible key size to ensure data confidentiality.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time.

Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# A dial-up router-to-router VPN connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**A dial-up router-to-router VPN connection**

A router-to-router VPN connection is typically used to connect remote offices together when both routers are connected to the Internet through permanent WAN links, such as T1 or Frame Relay. In this configuration, the VPN connection is always available. However, when a permanent WAN link is not possible or practical, you can configure a dial-up router-to-router VPN connection.

The following illustration shows a dial-up router-to-router VPN connection.




A dial-up router-to-router VPN connection consists of two demand-dial interfaces and two static routes that are configured on the VPN client (the calling router):

1. A demand-dial interface to dial in to a local Internet service provider (ISP).
2. A demand-dial interface for the router-to-router VPN connection.
3. A static host route to dynamically connect to the Internet.
4. A static route to reach the locations of the corporate office.

A dial-up router-to-router VPN connection is automatically established when you route traffic to a specific location. For example, in a branch office configuration, when a packet is received to be routed to the corporate office, the branch office router uses a dial-up link to connect to a local Internet service provider (ISP) and then creates a router-to-router VPN connection with the corporate office router located on the Internet.

**Note**

- This discussion assumes that the corporate office router (the answering router) is connected to the Internet by using a permanent WAN link. It is possible to have both routers connected to the Internet by using a dial-up WAN link. However, this is only feasible if the ISP supports demand-dialing routing to customers; the ISP calls the customer router when an IP datagram is to be delivered to the customer. Demand-dial routing to customers is not widely supported by ISPs.

To configure a dial-up router-to-router VPN connection at the branch office router, do the following:

1. Create a demand-dial interface for the Internet connection that is configured for the appropriate equipment (a modem or ISDN device), the phone number of the local ISP, and the user name and password to gain Internet access.
2. Create a demand-dial interface for the VPN connection with the corporate office router that is configured for a VPN port (a PPTP or L2TP port), the IP address of the interface on the Internet for the corporate office router, and a user name and password that can be verified by the VPN server. The user name must match the name of a demand-dial interface on the corporate office router.
3. Create a static host route to the corporate office router that uses the ISP demand-dial interface.
4. Create a static route (or routes) for the IP network IDs of the corporate intranet that uses the VPN demand-dial interface.

To configure the corporate office router, do the following:

1. Create a demand-dial interface for the VPN connection with the branch office that is configured for a VPN device (a PPTP or L2TP port). The demand-dial interface must have the same name as the user name in the authentication credential that is used by the branch office router to create the VPN connection.

2. Create a static route (or routes) for the IP network IDs of the branch office that uses the VPN demand-dial interface.

The router-to-router VPN connection is automatically initiated by the branch office router through the following process:

1. Packets sent to a corporate hub network location from a computer in the branch office are forwarded to the branch office router.

2. The branch office router checks its routing table and finds a route to the corporate intranet network ID, which uses the VPN demand-dial interface.

3. The branch office router checks the state of the VPN demand-dial interface and finds it is in a disconnected state.

4. The branch office router retrieves the configuration of the VPN demand-dial interface.

5. Based on the VPN demand-dial interface configuration, the branch office router attempts to initialize a router-to-router VPN connection at the IP address of the corporate office router on the Internet.

6. To establish a VPN, either a TCP (by using PPTP) or UDP (by using L2TP/IPSec) packet must be sent to corporate office router that acts as the VPN server. The VPN establishment packet is created.

7. To forward the VPN establishment packet to the corporate office router, the branch office router checks its routing table and finds the host route that is using the ISP demand-dial interface.

8. The branch office router checks the state of the ISP demand-dial interface and finds it is in a disconnected state.

9. The branch office router retrieves the configuration of the ISP demand-dial interface.

10. Based on the ISP demand-dial interface configuration, the branch office router uses its modem to dial and establish a connection with its local ISP.

11. Once the ISP connection is made, the VPN establishment packet is sent to the corporate office router.

12. A router-to-router VPN connection is negotiated between the branch office router and the corporate office router. As part of the negotiation, the branch office router sends authentication credentials that are verified by the corporate office router.

13. The corporate office router checks its demand-dial interfaces and finds one that matches the user name sent during authentication and changes the interface to a connected state.

14. The branch office router forwards the routed packet across the VPN and the corporate office router forwards the packet to the appropriate intranet location.

15. When the intranet location responds to the packet sent to it by the user in the branch office, the packet is forwarded to the corporate office router.

16. The corporate office router checks its routing table and finds a route to the branch office network that uses the VPN demand-dial interface.

17. The corporate office router checks the state of the VPN demand-dial interface and finds it is in a connected state.

18. The response packet is forwarded across the Internet by using the VPN connection.

19. The response packet is received by the branch office router and is forwarded to the original user.

# VPN servers and firewall configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

There are two approaches to using a firewall with a VPN server:

- **VPN server in front of the firewall**. The VPN server is attached to the Internet, and the firewall is between the VPN server and the intranet.

- **VPN server behind the firewall**. The firewall is attached to the Internet, and the VPN server is between the firewall and the intranet.

**VPN server in front of the firewall**

When the VPN server is in front of the firewall and attached to the Internet, you need to add packet filters to the Internet interface that allow only VPN traffic to and from the IP address of the VPN server's Internet interface.

For inbound traffic, when the tunneled data is decrypted by the VPN server, it is forwarded to the firewall. Through the use of its filters, the firewall allows the traffic to be forwarded to intranet resources. Because the only traffic that crosses the VPN server is generated by authenticated VPN clients, in this scenario, firewall filtering can be used to prevent VPN users from accessing specific intranet resources. Because the only Internet traffic allowed on the intranet must pass through the VPN server, this approach also prevents the sharing of File Transfer Protocol (FTP) or Web intranet resources with non-VPN Internet users.

The following illustration shows the VPN server in front of the firewall.




For the Internet interface on the VPN server, configure the following input and output filters using Routing and Remote Access:

**Packet filters for Point-to-Point Tunneling Protocol (PPTP)**

Configure the following input filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP destination port of 1723.

  This filter allows PPTP tunnel maintenance traffic from the PPTP client to the PPTP server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.

  This filter allows PPTP tunneled data from the PPTP client to the PPTP server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP [established] source port of 1723.

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. TCP [established] traffic is accepted only when the VPN server initiated the TCP connection.

Configure the following output filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP source port of 1723.

  This filter allows PPTP tunnel maintenance traffic from the VPN server to the VPN client.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.

  This filter allows PPTP tunneled data from the VPN server to the VPN client.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP destination port of 1723.

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. TCP [established] traffic is sent only when the VPN server initiated the TCP connection.

For more information, see Add PPTP Filters.

**Packet filters for Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec)**
Configure the following input filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 500.

  This filter allows Internet Key Exchange (IKE) traffic to the VPN server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 1701.

  This filter allows L2TP traffic from the VPN client to the VPN server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 4500.

  This filter allows IPSec network address translator traversal (NAT-T) traffic.

Configure the following output filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 500.

  This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 1701.

  This filter allows L2TP traffic from the VPN server to the VPN client.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 4500.

  This filter allows IPSec NAT-T traffic.

For more information, see Add Add L2TP over IPSec Filters.

There are no filters required for Encapsulating Security Payload (ESP) traffic at the IP protocol of 50. The ESP header is removed by the IPSec components before the L2TP packet is passed to Routing and Remote Access.

**Important**

- An IPSec NAT-T deployment for Windows that includes VPN servers that are located behind network address translators is not recommended. When a server is behind a network address translator, and the server uses IPSec NAT-T, unintended behavior might occur because of the way network address translators translate network traffic.

**VPN server behind the firewall**
In a more common configuration, the firewall is attached to the Internet, and the VPN server is an intranet resource that is attached to the Perimeter Network. The VPN server has an interface on both the perimeter network and the intranet. In this scenario, the firewall must be configured with input and output filters on its Internet interface that allow tunnel maintenance traffic and tunneled data to pass to the VPN server. Additional filters can allow traffic to pass to Web, FTP, and other types of servers on the perimeter network. For an additional layer of security, the VPN server can also be configured with PPTP or L2TP/IPSec packet filters on its perimeter network interface.

Because the firewall does not have the encryption keys for each VPN connection, it can filter only on the plaintext headers of the tunneled data. In other words, all tunneled data passes through the firewall. This is not a security concern, however, because the VPN connection requires an authentication process that prevents unauthorized access beyond the VPN server.

The following illustration shows the VPN server behind the firewall on the perimeter network.




For both the Internet and perimeter network interfaces on the firewall, configure the following input and output filters using the firewall's configuration software:

**Packet filters for PPTP**
Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

**Filters on the Internet interface**
Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter allows PPTP tunnel maintenance traffic from the PPTP client to the PPTP server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data from the PPTP client to the PPTP server.

- Destination IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should be used only in conjunction with the PPTP packet filters described in VPN server in front of the firewall and configured on the VPN server's perimeter network interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet that use this port.

Configure the following output filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter allows PPTP tunnel maintenance traffic from the VPN server to the VPN client.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data from the VPN server to the VPN client.

- Source IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should be used only in conjunction with the PPTP packet filters described in VPN server in front of the firewall and configured on the VPN server's perimeter network interface. By allowing all traffic from the VPN server

to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

**Filters on the perimeter network interface**

Configure the following input filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter allows PPTP tunnel maintenance traffic from the VPN server to the VPN client.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data from the VPN server to the VPN client.

- Source IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should be used only in conjunction with the PPTP packet filters described in VPN server in front of the firewall and configured on the VPN server's perimeter network interface. By allowing all traffic from the VPN server to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

  This filter allows PPTP tunnel maintenance traffic from the PPTP client to the PPTP server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

  This filter allows PPTP tunneled data from the PPTP client to the PPTP server.

- Destination IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should be used only in conjunction with the PPTP packet filters described in VPN server in front of the firewall and configured on the VPN server's perimeter network interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

**Packet filters for L2TP/IPSec**

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

**Filters on the Internet interface**

Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4).

  This filter allows IKE traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194).

  This filter allows IPSec NAT-T traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic from the VPN client to the VPN server.

Configure the following output packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4).

  This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and UDP source port of 4500.

  This filter allows IPSec NAT-T traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic from the VPN server to the VPN client.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

**Important**

- An IPSec NAT-T deployment for Windows that includes VPN servers that are located behind network address translators is not recommended. When a server is behind a network address translator, and the server uses IPSec NAT-T, unintended behavior might occur because of the way network address translators translate network traffic.

**Filters on the perimeter network interface**

Configure the following input packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4).

  This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and UDP source port of 4500.

  This filter allows IPSec NAT-T traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic from the VPN server to the VPN client.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4).

  This filter allows IKE traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194).

  This filter allows IPSec NAT-T traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

  This filter allows IPSec ESP traffic from the VPN client to the VPN server.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

**Important**

- An IPSec NAT-T deployment for Windows that includes VPN servers that are located behind network address translators

is not recommended. When a server is behind a network address translator, and the server uses IPSec NAT-T, unintended behavior might occur because of the way network address translators translate network traffic.

# VPN servers and Microsoft ISA Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN servers and Microsoft ISA Server**
Microsoft Internet Security and Acceleration (ISA) Server can be installed on a virtual private network (VPN) server running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition, and provides the capabilities of both the Windows Server 2003 Routing and Remote Access service and the ISA Server.

When ISA Server is installed, all of the packet filters that are configured on the VPN server are ignored by ISA. ISA creates its own packet filters when it is configured to allow VPN connections. After ISA is installed, use ISA Administrator, instead of Routing and Remote Access, to configure additional filters.

ISA Server includes two wizards that you can use to connect two ISA server computers with a router-to-router VPN connection:

- Local ISA VPN Wizard

  This wizard is used to create a .vpc (VPN configuration) file.

- Remote ISA VPN Wizard

  This wizard is used to read the .vpc file. After the servers have been configured, both ISA servers can initiate and accept VPN connections.

For more information about ISA Server packet filtering capabilities and support for VPNs, see Microsoft Internet Security and Acceleration (ISA) Server online Help.

# Network access authentication and certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access authentication and certificates**
This topic contains the following sections:


- Overview
- Certificate requirements for EAP
- Computer authentication by IPSec
- Certificate-based authentication and wireless clients
- Certificate enrollment methods and domain membership
- Choosing a certificate enrollment method
- CA Web enrollment services


**Overview**
Certificates are used for network access authentication because they provide strong security for authenticating users and computers and eliminate the need for less secure password-based authentication methods. This topic describes how Internet Authentication Service (IAS) and virtual private network (VPN) servers use Extensible Authentication protocol-Transport Level Security (EAP-TLS), protected Extensible Authentication protocol (PEAP), or Internet protocol security (IPSec) to perform certificate-based authentication for many types of network access, including VPN and wireless connections. Additionally, this topic describes certificate enrollment methods to help you determine the best certificate enrollment method for your use.

When discussing authentication, a server is defined as a VPN or IAS server that is a TLS end point. You can configure VPN servers to perform network access authentication without IAS, or you can use IAS for authentication when you have multiple Remote Access Dial-In User Service (RADIUS) clients (such as VPN servers and wireless access points) on your network.

Two authentication methods use certificates: Extensible Authentication protocol-Transport Level Security (EAP-TLS) and protected Extensible Authentication protocol (PEAP). Both methods always use certificates for server authentication. Depending on the authentication type configured with the authentication method, certificates might be used for user authentication and client authentication. For more information, see EAP and PEAP.

The use of certificates for authentication of VPN connections is the strongest form of authentication available with the Windows Server 2003 family. You must use certificate-based authentication for VPN connections based on Layer Two Tunneling protocol over Internet protocol security (L2TP/IPSec). Point-To-Point Tunneling protocol (PPTP) connections do not require certificates, although you can configure PPTP connections to use certificates for computer authentication when you use EAP-TLS as the authentication method. For wireless clients (computing devices with wireless network adapters, such as your portable computer or personal digital assistant), PEAP with EAP-TLS and smart cards or certificates is the recommended authentication method. For more information, see Layer Two Tunneling Protocol, Point-to-Point Tunneling Protocol, and Wireless Networking.

**Notes**

- During authentication attempts for PPTP connections, computer authentication does not occur. When EAP-TLS is used as the authentication method with PPTP connections, user authentication is performed with a certificate from the certificate store on the local computer or from a smart card.

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

In order to deploy certificates to users and computers, you must first do the following:

- Design a public key infrastructure (PKI).
- Deploy a certification authority (CA) using Certificate Services.
- Design and publish one or more certificates using Certificate Templates, which is installed with Certificate Services.
- Select one or more distribution methods (also known as certificate enrollment methods) to install the certificates on computers and distribute them to users.

### Notes

- This topic assumes that you have designed and deployed a PKI according to the guidelines provided in Certificate Services Best practices. For more information, see Public Key Infrastructure.
- PEAP is not supported for VPN or dial-up connections.

#### Example certificate deployments

- Remote access VPN connections

  L2TP/IPSec or PPTP connections between VPN server and VPN client with EAP-TLS used as the authentication method. IPSec uses computer certificates between the client and the server for authentication, and EAP-TLS uses a certificate (from a smart card or the user's local certificate store) for user authentication. The IAS or VPN server certificate must contain the server Authentication purpose, and the client computer or user certificate must contain the Client Authentication purpose in the Enhanced Key Usage (EKU) extensions of the certificate.

  For more information, see L2TP-based remote access VPN deployment and PPTP-based remote access VPN deployment.

- Router-to-router VPN connections

  L2TP/IPSec connections for dedicated or dial-on-demand connections between servers with EAP-TLS as the authentication method. Both servers must have certificates that contain the server Authentication and Client Authentication purposes in EKU extensions.

  For more information, see Deploying Router-to-Router VPNs.

- IEEE 802.1X wireless or switch clients

  PEAP-EAP-MS-CHAPv2 is configured as the authentication method, and the **Validate server certificate** option is enabled on the client computers. The IAS server certificate EKU extensions contain the server Authentication purpose, and the certificate is used to identify the server to the client. User authentication is accomplished with user name and password.

  For more information, see Enable smart card or other certificate authentication and Configuring wireless network settings on client computers.

- IEEE 802.1X wireless or switch clients

  L2TP/IPSec is used, and EAP-TLS with certificates is configured as the authentication method. The IAS server certificate contains the server Authentication purpose in EKU extensions to identify itself to the client, and the client uses a certificate (from a smart card or the user's local certificate store) to identify itself to the IAS server. (The wireless access points are configured as RADIUS clients on the IAS server, which is the EAP authenticator.)

**Note**

- User authentication and network access authorization are performed by either the VPN server or a RADIUS server, such as Internet Authentication Service (IAS), as configured on the VPN gateway. For more information, see Use RADIUS authentication.

**Certificate requirements for EAP**

When you use EAP with a strong EAP type (such as TLS with smart cards or certificates), both the client and the server use certificates to verify their identities to each other. Certificates must meet specific requirements in order to allow the server and the client to use them for successful authentication.

One such requirement is that the certificate is configured with one or more purposes in EKU extensions that correlate to the certificate use. For example, a certificate used for the authentication of a client to a server must be configured with the Client Authentication purpose. Similarly, a certificate used for the authentication of a server must be configured with the server Authentication purpose. When certificates are used for authentication, the authenticator examines the client certificate, seeking the correct purpose object identifier in EKU extensions. For example, the object identifier for the Client Authentication purpose is 1.3.6.1.5.5.7.3.2.

Certificate Templates allow customization of the certificates issued by Certificate Services, including both how certificates are issued and what they contain, including their purposes. In Certificate Templates, you can use a default template, such as the Computer template, to define the template that the CA uses to assign certificates to computers. You can also create a certificate template and assign purposes in EKU extensions to the certificate. By default, the Computer template includes the Client Authentication purpose and the server Authentication purpose in EKU extensions.

The certificate template that you create can include any purpose for which the certificate will be used. For example, if you use smart cards for authentication, you can include the Smart Card Logon purpose in addition to the Client Authentication purpose.

When using IAS, you can configure IAS to check certificate purposes before granting network authorization. IAS can check additional EKUs and Issuance Policy purposes (also known as Certificate Policies).

Some non-Microsoft CA software might contain a purpose named All, which represents all possible purposes. This is indicated by a blank (or null) EKU extension. Although All means all possible purposes, the All purpose cannot be substituted for the Client Authentication purpose, the server authentication purpose, or any other purpose related to network access authentication.

For more information, see Certificate Templates and Manage Certificate Templates for an Enterprise Certification Authority.

You can view certificates and their purposes by using the Certificates console to open the certificate store. Certificate stores can be viewed in Logical Store mode or in Purpose mode. For information about how to view certificate purposes, see Display certificate stores in Purpose mode.

**Minimum certificate requirements**

All certificates that are used for network access authentication must meet the requirements for X.509 certificates and work for connections that use Secure Sockets Layer-Transport Level Security (SSL/TLS). After this minimum requirement is met, both client and server certificates have additional requirements.

**Client certificate requirements**

With EAP-TLS or PEAP-EAP-TLS, the server accepts the client authentication attempt when the certificate meets the following requirements:

- The client certificate is issued by an enterprise CA or mapped to a user or computer account in Active Directory.

- The user or computer certificate on the client chains to a trusted root CA, includes the Client Authentication purpose in EKU extensions (the object identifier for Client Authentication is 1.3.6.1.5.5.7.3.2), and fails neither the checks that are performed by CryptoAPI and specified in the remote access policy nor the Certificate object identifier checks that are

specified in IAS remote access policy.

- The 802.1X client does not use registry-based certificates that are either smart card-logon or password-protected certificates.

- For user certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate contains the user principal name (UPN). To configure the UPN name in a certificate template:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.

4. In **Include this information in alternate subject name**, select **User principal name (UPN)**.

- For computer certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate must contain the client's fully qualified domain name (FQDN), which is also called the DNS name. To configure this name in the certificate template:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.

4. In **Include this information in alternate subject name**, select **DNS name**.

With PEAP-EAP-TLS and EAP-TLS, clients display a list of all installed certificates in the Certificates snap-in, with the following exceptions:

- Wireless clients do not display registry-based and smart card-logon certificates.

- Wireless clients and VPN clients do not display password-protected certificates.

- Certificates that do not contain the Client Authentication purpose in EKU extensions are not displayed.

**Server certificate requirements**

Clients can be configured to validate server certificates by using the **Validate server certificate** option. With PEAP-EAP-MS-CHAPv2, PEAP-EAP-TLS, or EAP-TLS as the authentication method, the client accepts the server's authentication attempt when the certificate meets the following requirements:

- The Subject name contains a value. If you issue a certificate to your IAS server that has a blank Subject, the certificate is not available to authenticate your IAS server. To configure the certificate template with a Subject name:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.

4. In **Subject name format**, select a value other than **None**.

- The computer certificate on the server chains to a trusted root CA and does not fail any of the checks that are performed by CryptoAPI and specified in the remote access policy.

- The IAS or VPN server computer certificate is configured with the server Authentication purpose in EKU extensions (the object identifier for server Authentication is 1.3.6.1.5.5.7.3.1).

- The server certificate is configured with a required cryptographic service provider (CSP) value of **Microsoft RSA SChannel Cryptographic provider**. To configure the required CSP:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Request Handling** tab, and then click **CSPs**.

4. In **CSP Selection**, select **Requests must use one of the following CSPs**.

5. In **CSPs**, select the **Microsoft RSA SChannel Cryptographic provider** checkbox. Clear all other checkboxes in **CSPs**.

- The Subject Alternative Name (SubjectAltName) extension, if used, must contain the DNS name of the server. To configure the certificate template with the DNS name of the enrolling server:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.

4. In **Include this information in alternate subject name**, select **DNS name**.

With PEAP and EAP-TLS, servers display a list of all installed certificates in the computer's certificate store, with the following exceptions:

- Certificates that do not contain the server Authentication purpose in EKU extensions are not displayed.

- Certificates that do not contain a Subject name are not displayed.

- Servers do not display registry-based and smart card-logon certificates.

**Note**

- You can designate which certificate is used by the IAS or VPN server in the remote access policy profile. When you configure EAP and PEAP authentication methods that require a certificate for server authentication and you do not select a specific certificate in the **Smart Card or other Certificate properties** dialog box, the IAS or VPN server automatically selects a certificate from the computer certificate store. If the server obtains a newer certificate, it uses the new certificate. This might cause the IAS or VPN server to use a certificate that is not correctly configured for authentication, causing authentication to fail as the result. To prevent this, always select a server certificate when configuring PEAP and EAP authentication methods that require one. For more information, see Configure PEAP and EAP methods.

**Computer authentication by IPSec**

Computer authentication is first performed during L2TP/IPSec connection attempts between remote access client and server. When a secure channel is established between the client and the server, the user authentication and authorization attempt proceeds.

Computer authentication by IPSec is performed using preshared keys or computer certificates. The recommended method of authentication is the use of a PKI and certificates. If certificates are used, a computer certificate is required to establish IPSec trust during the Internet Key Exchange (IKE) negotiation in L2TP/IPSec-based VPN connections.

In order for a VPN server running Windows Server 2003 and a VPN client running Windows 2000 or Windows XP to establish trust for an L2TP/IPSec VPN connection, both computers must have a computer certificate that was issued by the same trusted enterprise root CA. When both computers have and exchange certificates that were issued by the trusted root CA during IPSec negotiation, they extend trust to each other, and the security association is made.

**How VPN servers use certificates**

When an L2TP/IPSec VPN connection is attempted between VPN client and server, computer authentication fails if the VPN client certificate (from a smart card or the certificate store on the local computer) is not configured with the Client Authentication purpose in EKU extensions, and the VPN server certificate is not configured with the server Authentication purpose in EKU extensions. IPSec checks the EKU extensions for the client certificate to determine whether the Client Authentication purpose object identifier is present. When the EKU extension includes the Client Authentication purpose object identifier, IPSec can use the certificate for authentication. Similarly, the client checks the VPN server certificate for the server Authentication purpose object identifier in EKU extensions.

Although VPN servers that end connections for remote users need a certificate configured with the server Authentication purpose in EKU extensions only, a VPN servers that is used as an end point for a VPN connection with another VPN server originates (as a client) and ends (as a server) VPN connections. For this reason, the certificate on these servers must contain both the server Authentication purpose and the Client Authentication purpose in EKU extensions. In addition, due to the way in which automatic certificate selection functions, both purposes (server Authentication and Client Authentication) must be contained in the same certificate.

Automatic certificate selection can provide IPSec with any certificate in the certificate store that chains to the trusted enterprise root CA, regardless of the purposes contained in the certificate. If two certificates are installed on the server--one that contains the Client Authentication purpose and one that contains the server Authentication purpose--it is possible that automatic certificate selection could use the wrong certificate for authentication. For example, the certificate with the Client Authentication purpose might be required, but the certificate provided through automatic certificate selection contains the server Authentication purpose. In this and similar instances, computer authentication fails.

**Certificate-based authentication and wireless clients**

IEEE 802.1X authentication provides authenticated access to 802.11 wireless networks and to wired Ethernet networks. 802.1X provides support for secure EAP types, such as TLS with smart cards or certificates. You can configure 802.1X with EAP-TLS in a variety of ways. If the **Validate server certificate** option is configured on the Windows XP Professional client, the client authenticates the server by using its certificate. Client computer and user authentication can be accomplished using certificates from the client certificate store or a smart card, providing mutual authentication.

With wireless clients, PEAP-EAP-MS-CHAPv2 can be used as the authentication method. PEAP-EAP-MS-CHAPv2 is a password-based user authentication method that uses TLS with server certificates. During PEAP-EAP-MS-CHAPv2 authentication, the IAS or RADIUS server supplies a certificate to validate its identity to the client (if the **Validate server certificate** option is configured on the Windows XP Professional client). Client computer and user authentication is accomplished with passwords, which eliminates some of the difficulty of deploying certificates to wireless client computers.

802.1X wireless and switch clients can also use PEAP-EAP-TLS, which provides strong security. PEAP-EAP-TLS uses a public key infrastructure (PKI) with certificates used for server authentication and either smart cards or certificates for client computer and user authentication. For more information, see PEAP.

For more information about 802.1X authentication, see Understanding 802.1X authentication for wireless networks.

**Certificate enrollment methods and domain membership**

The domain membership of computers for which you want to enroll certificates affects the certificate enrollment method that you can choose. Certificates for domain member computers can be enrolled automatically (also known as auto-enrollment), while an administrator must enroll certificates for non-domain member computers using the Web or a floppy disk. The certificate enrollment method for non-domain member computers is known as a trust bootstrap process, through which certificates are created and then manually requested or distributed securely by administrators, to build common trust.

**Domain member certificate enrollment**

If your VPN server, IAS server, or client running Windows 2000 or Windows XP is a member of a domain running Windows Server 2003 and Active Directory, you can configure the auto-enrollment of computer and user certificates. After auto-enrollment is configured and enabled, all domain member computers receive computer certificates when Group Policy is refreshed next, whether the refresh is triggered manually with the **gpupdate** command, or by logging on to the domain.

If your computer is a member of a domain where Active Directory is not installed, you can install computer certificates

manually by requesting them through Certificates.

For more information, see Computer certificates for L2TP/IPSec VPN connections.

**Note**

- Computers running Windows 2000 can auto-enroll computer certificates only.

**Non-domain member certificate enrollment**

Certificate enrollment for computers that are not domain members cannot be done with auto-enrollment. When a computer is joined to a domain, a trust is established that allows auto-enrollment to occur without administrator intervention. When a computer is not joined to a domain, trust is not established and a certificate is not issued. Trust must be established using one of the following methods:

- An administrator (who is, by definition, trusted) must request a computer or user certificate using the CA Web enrollment tool.

- An administrator must save a computer or user certificate to a floppy disk and install it on the non-domain member computer. Or, when the computer is not accessible to the administrator (for example, a home computer connecting to an organization network with an L2TP/IPSec VPN connection), a domain user whom the administrator trusts can install the certificate.

- An administrator can distribute a user certificate on a smart card (computer certificates are not distributed on smart cards).

Many network infrastructures contain VPN and IAS servers that are not domain members. For example, a VPN server in a perimeter network might not be a domain member for security purposes. In this case, a computer certificate with the server Authentication purpose contained in the EKU extensions must be installed on the non-domain member VPN server before it can successfully negotiate L2TP/IPSec-based VPN connections with clients. Note that if the non-domain member VPN server is used as an end point for a VPN connection with another VPN server, EKU extensions must contain both the server Authentication and Client Authentication purposes.

**Choosing a certificate enrollment method**

If you are running an enterprise certification authority (CA) on a computer running Windows Server 2003, Standard Edition, you can use the following table to determine the best certificate enrollment method for your requirements:

| Object and domain membership | Certificate template | Certificate purposes | Preferred certificate enrollment method | Alternate certificate enrollment method |
|---|---|---|---|---|
| VPN or IAS server, domain member | Computer | Server Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| VPN server with site-to-site connection, domain member | Computer | Server Authentication and Client Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| Windows XP client, domain member | Computer | Client Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| VPN or IAS server, non-domain member | Computer | Server Authentication | CA Web enrollment tool | Install from a floppy disk |
| VPN server with site-to-site connection, non-domain member | Computer | Server Authentication and Client Authentication | CA Web enrollment tool | Install from a floppy disk |
| Windows XP client, non-domain member | Computer | Client Authentication | CA Web enrollment tool | Install from a floppy disk |
| User, domain user | User | Client Authentication | Auto-enrollment | Use a smart card or the CA Web enrollment tool |

If your certification authority is on a computer running one of the following operating systems, the RAS and IAS servers and Workstation Authentication templates are available for use:

- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition
- Windows Server 2003, Enterprise Edition for Itanium-based Systems
- Windows Server 2003, Datacenter Edition for Itanium-based Systems
- Windows Server 2003, Enterprise x64 Edition
- Windows Server 2003, Datacenter x64 Edition

Use the following table to determine when to use these templates.

| Object and domain membership | Certificate template | Certificate purpose | Preferred certificate enrollment method | Alternate certificate enrollment method |
|---|---|---|---|---|
| VPN or IAS server, domain member | RAS and IAS server | Server Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| Windows XP client, domain member | Workstation Authentication | Client Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| VPN or IAS server, non-domain member | RAS and IAS server | Server Authentication | CA Web enrollment tool | Install from a floppy disk |
| Windows XP client, non-domain member | Workstation Authentication | Client Authentication | CA Web enrollment tool | Install from a floppy disk |

**Important**

If your server running IAS is not a domain controller but is a member of a domain with a Windows 2000 mixed functional level, you must add the server to the access control list (ACL) of the **RAS and IAS server** certificate template. You must also configure the correct permissions for autoenrollment. There are different procedures for adding single servers and groups of servers to the ACL.

To add an individual server to the ACL for the **RAS and IAS server** certificate template:

In Certificate Templates, select the template **RAS and IAS server**, and then add the IAS server to the template **Security** properties. For more information, see Allow subjects to request a certificate that is based on the template. After you have added your IAS server to the ACL, grant **Read**, **Enroll**, and **Autoenroll** permissions.

To manage a group of servers, add the servers to a new global or universal group, and then add the group to the ACL of the certificate template:

1. In Active Directory Users and Computers, create a new global or universal group for IAS servers. For more information, see Create a new group. Next, add to the group all computers that are IAS servers, are not domain controllers, and are members of a domain with a Windows 2000 mixed functional level. For more information, see Add a member to a group.

2. In Certificate Templates, select the **RAS and IAS server** template, and then add the group you created to the template **Security** properties by completing the steps in Allow subjects to request a certificate that is based on the template. After you have added your new group, grant **Read**, **Enroll**, and **Autoenroll** permissions.

For more information, see Domain and forest functionality.

**CA Web enrollment services**

A set of CA Web pages is provided with Certificate Services in the Windows Server 2003 family. These Web enrollment pages allow you to connect to the CA through a Web browser and perform common tasks, such as requesting certificates from a CA.

For more information, see Certification authority Web enrollment services and Set up certification authority Web enrollment support.

After your PKI has been designed and deployed and you are assigned the required permissions on the certificate templates, you can request certificates for computers by using the CA Web enrollment pages. You can request a certificate for a non-domain member computer using the CA Web enrollment tool by completing the following steps. You must be an administrator on the local computer to complete these steps.

**Install a computer certificate on the local computer**

1. Use your Web browser to log onto the Web enrollment tool at **http://***serverName***/certsrv**, where *serverName* is the name of the server hosting the CA, the CA Web pages, or both. Click **Request a certificate**.

2. In **Request a certificate**, click **Advanced Certificate Request**.

3. In **Advanced Certificate Request**, click **Create and submit a request to this CA**. A Web form that contains the sections **Certificate Template**, **Key Options**, and **Additional Options** becomes available. Select the following:

   - **Certificate template:** Administrator
   - **Key options:** Create new key set
   - **CSP:** Microsoft Enhanced Cryptographic provider v1.0
   - **Key usage:** Exchange
   - **Key size:** 1024
   - **Automatic key container name:** Selected
   - **Mark keys as exportable:** Selected
   - **Use local machine store:** Selected
   - **Request format:** PKCS 10
   - **Hash:** SHA-1
   - **Friendly name:** VPN server name

4. Click **Submit**.

5. Click **Install this certificate**

   The certificate you requested is installed on your local computer.

**Save a computer certificate to floppy disk**

1. Use your Web browser to log onto the Web enrollment tool at **http://***serverName***/certsrv**, where *serverName* is the name of the server hosting the CA, the CA Web pages, or both. Click **Request a certificate**.

2. In **Request a certificate**, click **Advanced Certificate Request**.

3. In **Advanced Certificate Request**, click **Create and submit a request to this CA**. A Web form that contains the sections **Certificate Template**, **Key Options**, and **Additional Options** becomes available. Select the following:

   - **Certificate template:** Administrator

- **Key options:** Create new key set
- **CSP:** Microsoft Enhanced Cryptographic provider v1.0
- **Key usage:** Exchange
- **Key size:** 1024
- **Automatic key container name:** Selected
- **Mark keys as exportable:** Selected
- **Export keys to file:** Selected
- **Full path name:** The path location where you want to save the file and the file name (for example, *\\serverName\ShareName\FileName.pvk*.).
- **Request format:** PKCS 10
- **Hash:** SHA-1
- **Friendly name:** VPN server name

4. Click **Submit**.
5. Create and confirm your private key password.
6. In **Certificate Issued**, select **DER encoded** and **Download certificate chain**.
7. Insert a floppy disk in your disk drive.
8. When prompted by your browser, select **Save this file to disk**, and then click **OK**.
9. Browse to your floppy drive, and then click **OK**.

   The certificate you requested is downloaded and saved to floppy disk.

For more information, see Submit an advanced certificate request via the Web.

For information about installing a certificate from floppy to the certificate store of a local computer, see Import a certificate.

**Note**

- If you have issued a certificate to your IAS server that has a blank Subject, the certificate is not available to authenticate your IAS server. To change this, you can use Certificate Templates to create a new certificate for enrollment on your IAS server. In the certificate properties, on the **Subject Name** tab, in **Subject name format**, select a value other than **None**.

# Understanding VPN Tunneling Protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN tunneling protocols**
This section covers:

- Point-to-Point Tunneling Protocol
- Layer Two Tunneling Protocol

# Point-to-Point Tunneling Protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Point-to-Point Tunneling Protocol**
Point-to-Point Tunneling Protocol (PPTP) is a tunneling protocol first supported in Windows NT 4.0 and Windows 98. PPTP is an extension of Point-to-Point Protocol (PPP) and leverages the authentication, compression, and encryption mechanisms of PPP. Client support for PPTP is built-in to the Windows XP remote access client.

VPN server support for PPTP is built in to members of the Windows Server 2003 family. PPTP is installed with the TCP/IP protocol. Depending on your choices when running the Routing and Remote Access Server Setup Wizard, PPTP is configured for five or 128 PPTP ports. For more information, see Checklist: Installing and configuring a PPTP server.

PPTP and Microsoft Point-to-Point Encryption (MPPE) provide the primary VPN services of encapsulation and encryption of private data.

**Encapsulation**
A PPP frame (an IP, IPX, or Appletalk datagram) is wrapped with a Generic Routing Encapsulation (GRE) header and an IP header. In the IP header is the source and destination IP address that correspond to the VPN client and VPN server.

The following illustration shows PPTP encapsulation for a PPP frame.




**Encryption**
The PPP frame is encrypted with Microsoft Point-to-Point Encryption (MPPE) by using encryption keys generated from the MS-CHAP, MS-CHAP v2, or EAP-TLS authentication process. Virtual private networking clients must use the MS-CHAP, MS-CHAP v2, or EAP-TLS authentication protocol in order for the payloads of PPP frames to be encrypted. PPTP is taking advantage of the underlying PPP encryption and encapsulating a previously encrypted PPP frame.

For more information on deploying PPTP-based VPN connections, see Deploying Virtual Private Networks. For a PPTP-based VPN server checklist, see Checklist: Installing and configuring a PPTP server. For an example implementation of PPTP-based VPN connections, see Virtual Private Network Implementation Examples. For information on configuring PPTP-based VPN connections in a test lab, see Virtual Private Network Test Lab.

## Notes

- It is possible to have a nonencrypted PPTP connection where the PPP frame is sent in plaintext. However, a nonencrypted PPTP connection is not recommended for VPN connections over the Internet because communications of this type are not secure.

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Layer Two Tunneling Protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Layer Two Tunneling Protocol**
Layer Two Tunneling Protocol (L2TP) is an RFC-based tunneling protocol that is an industry standard and was first supported in the Windows 2000 client and server operating systems. Unlike PPTP, L2TP in servers running Windows Server 2003 does not utilize Microsoft Point-to-Point Encryption (MPPE) to encrypt Point-to-Point Protocol (PPP) datagrams. L2TP relies on Internet Protocol security (IPSec) for encryption services. The combination of L2TP and IPSec is known as L2TP/IPSec. L2TP/IPSec provides the primary virtual private network (VPN) services of encapsulation and encryption of private data.

Both L2TP and IPSec must be supported by both the VPN client and the VPN server. Client support for L2TP is built-in to the Windows XP remote access client, and VPN server support for L2TP is built in to members of the Windows Server 2003 family.

L2TP is installed with the TCP/IP protocol. Depending on your choices when running the Routing and Remote Access Server Setup Wizard, L2TP is configured for five or 128 L2TP ports. For more information, see Checklist: Installing and configuring an L2TP server.

For more information about IPSec, see Internet Protocol Security Overview.

**Encapsulation**
Encapsulation for L2TP/IPSec packets consists of two layers:

1. L2TP encapsulation

   A PPP frame (an IP datagram or an IPX datagram) is wrapped with an L2TP header and a UDP header.

2. IPSec encapsulation

   The resulting L2TP message is then wrapped with an IPSec Encapsulating Security Payload (ESP) header and trailer, an IPSec Authentication trailer that provides message integrity and authentication, and a final IP header. In the IP header is the source and destination IP address that corresponds to the VPN client and VPN server.

The following illustration shows L2TP and IPSec encapsulation for a PPP datagram.




**Encryption**
The L2TP message is encrypted with either Data Encryption Standard (DES) or Triple DES (3DES) by using encryption keys generated from the Internet Key Exchange (IKE) negotiation process.

For more information on deploying L2TP/IPSec VPN connections, see Deploying Virtual Private Networks. For configuration checklists, see VPN Checklists. For an example implementation of L2TP/IPSec VPN connections, see Virtual Private Network Implementation Examples. For information on configuring L2TP/IPSec VPN connections in a test lab, see Virtual Private Network Test Lab.

### Notes

- It is possible to have a nonencrypted L2TP connection where the PPP frame is sent in plaintext. However, a nonencrypted L2TP connection is not recommended for virtual private network connections over the Internet because communications of this type do not provide data confidentiality.

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# VPN Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**VPN security**
This section covers:

- Authorization
- Authentication of VPN clients
- Data encryption between VPN server and client
- Using packet filtering

# Authorization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authorization**
Virtual private network (VPN) connections are only accepted for those users and routers that have been authorized. In the Windows Server 2003 family, authorization of VPN connections is determined by the dial-in properties on the user account and remote access policies. For more information, see Remote Access Policies.

You do not need to create additional user accounts just for VPN connections. VPN servers use the user accounts specified in the available user accounts databases according to security in Windows Server 2003 operating systems.

**How security works at connection**
The following steps describe what happens during a connection attempt from a Point-to-Point Tunneling Protocol (PPTP)-based VPN client to a PPTP-based VPN server running Windows Server 2003:

1. The VPN client creates a PPTP tunnel with the VPN server.
2. The server sends a challenge to the client.
3. The client sends an encrypted response to the server.
4. The server checks the response against the user accounts database.
5. If the account is valid, and the connection is authorized, the server accepts the connection subject to the remote access policies and user account properties for the VPN client.

**Note**

- Steps 2 through 4 assume that the VPN client and the VPN server use the MS-CHAP or CHAP authentication protocols. The sending of client credentials may vary for other authentication protocols.

The following steps describe what happens during a connection attempt from a Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec)-based VPN client to an L2TP/IPSec-based VPN server running Windows Server 2003:

1. The IPSec security association is created by using computer certificates and the Internet Key Exchange (IKE) negotiation process.
2. The VPN client creates an L2TP tunnel with the VPN server.
3. The server sends a challenge to the client.
4. The client sends an encrypted response to the server.
5. The server checks the response against the user accounts database.
6. If the account is valid, and the connection is authorized, the server accepts the connection subject to the remote access policies and user account properties for the VPN client.

**Note**

- Steps 3 through 5 assume that the VPN client and the VPN server use the MS-CHAP v1 or CHAP authentication protocols. The sending of client credentials may vary for other authentication protocols.

**Security after the connection is made**

After passing authorization and authentication and connecting to the LAN, VPN clients for remote access VPN connections can only access network resources for which they have permission. Remote access VPN clients are subject to security in Windows Server 2003 operating systems--just as they are at the office. In other words, remote access VPN clients cannot do anything for which they lack sufficient rights, nor can they access resources for which they do not have permission.

The VPN server must authenticate the remote access VPN clients before they can access or generate traffic on the network. This authentication is a separate step from logging on to a Windows Server 2003 domain.

You use packet filters based on a remote access policy profile to restrict access for remote access VPN connections for IP traffic. With profile packet filters, you can configure the IP traffic that is allowed out of the connection (output filters) or into the connection (input filters) on an exception basis: either all remote access traffic except traffic specified by the filters or no traffic except traffic specified by the filters. Remote access policy profile filtering applies to all remote access connections that match the remote access policy.

For more information, see Elements of a remote access policy.

**Credential caching**

When remote access servers are configured to accept a user's domain credentials (user name and password) as the means for authenticating access requests, the credentials used for remote access are the same as those needed for access to network resources (such as file servers, application servers, and e-mail servers). Each time the remote access client tries to access a server on the network, the server challenges the client for its credentials to ensure that only authorized users can utilize server resources.

If the client computer is a member of the server's domain, then the client responds automatically with the user's logon credentials and gains access to resources for which the user has security permissions. If the client computer is not a member of the server's domain (for example, when the client is a non-domain member home computer), then the client does not have the appropriate logon credentials to meet the server's challenge; for each network resource the user attempts to access, the user is either prompted to enter credentials or is denied access.

When the client is running Windows XP or a Windows Server 2003 operating system, the credentials that are used for authentication to the remote access server are cached, or saved, by the client. If the non-domain member client attempts to access a domain server for which it has not previously established its credentials, the client responds to the server's challenge with the cached remote access credentials and gains access to resources for which the user has security permissions. Because the credentials are cached and sent automatically upon server challenge, users on non-domain member computers are not prompted for credentials every time they attempt to access network resources.

# Authentication of VPN clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication**

The authentication of virtual private network (VPN) clients by the VPN server is a vital security concern. Authentication takes place at two levels:

1. Computer-level authentication

   When Internet Protocol security (IPSec) is used for a Layer Two Tunneling Protocol (L2TP) over IPSec (L2TP/IPSec) VPN connection, computer-level authentication is performed through the exchange of computer certificates or a preshared key during the establishment of the IPSec security association. For more information, see Internet Key Exchange.

2. User-level authentication

   Before data can be sent over the Point-to-Point Tunneling Protocol (PPTP) or L2TP tunnel, the remote access client or demand-dial router that requests the VPN connection must be authenticated. User-level authentication occurs through the use of a Point-to-Point Protocol (PPP) authentication method. For more information, see Remote Access Authentication Methods.

For more information, see Network access authentication and certificates.

# Data encryption between VPN server and client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data encryption**

You must use data encryption to provide data confidentiality for the data that is sent between the VPN client and the VPN server across a shared or public network, where there is always a risk of unauthorized interception. You can configure the VPN server to force encrypted communications. Users who connect to that server must encrypt their data or a connection is not allowed. For VPN connections, the Windows Server 2003 family uses Microsoft Point-to-Point Encryption (MPPE) with the Point-to-Point Tunneling Protocol (PPTP) and Internet Protocol security (IPSec) encryption with the Layer Two Tunneling Protocol (L2TP).

Because data encryption is performed between the VPN client and VPN server, data encryption is not necessary on the communication link between a dial-up client and its Internet service provider (ISP). For example, a mobile user uses a dial-up connection to dial in to a local ISP. Once the Internet connection is made, the user creates a VPN connection with the corporate VPN server. If the VPN connection is encrypted, encryption is not needed on the dial-up connection between the user and the ISP.

### Notes

- Data encryption for Point-to-Point Protocol (PPP) or PPTP connections is available only if you use MS-CHAP, MS-CHAP v2, or EAP-TLS as the user-level authentication method. Data encryption for L2TP connections relies on IPSec computer-level authentication, which does not require any specific user-level authentication method.

- VPN data encryption does not provide end-to-end data encryption. End-to-end encryption is data encryption between the client application and the server hosting the resource or service that is accessed by the client application. To get end-to-end data encryption, you can use IPSec to create a secure connection *after* the VPN connection is made.

# Using packet filtering

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Packet filtering**

To secure the VPN server from sending or receiving any traffic on its Internet interface except VPN traffic, you need to use PPTP or L2TP/IPSec input and output filters on the interface that corresponds to the connection to the Internet. For router-to-router VPN connections, you must also configure the calling router (the VPN client) with PPTP or L2TP/IPSec packet filters, a step that is done automatically when you configure your remote access server using the Routing and Remote Access Setup Wizard.

Because, by default, IP routing is enabled on intranet interfaces and the interface that corresponds to the connection to the Internet, the computer running a Windows Server 2003 operating system forwards IP packets between the Internet and your intranet. This provides a direct, routed connection between your intranet and possible attackers on the Internet. To protect your intranet so that the only traffic that is forwarded to the intranet is the traffic that is sent and received over secure VPN connections, you must use PPTP or L2TP/IPSec filters on the Internet interface.

If you have a firewall, you must configure packet filters on the firewall to allow traffic between the VPN router and the routers on the Internet. For more information, see VPN servers and firewall configuration.

For more information about PPTP filters, see Add PPTP Filters. For more information about L2TP/IPSec filters, see Add L2TP over IPSec Filters.

# Deploying Virtual Private Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying virtual private networks**

This section covers:

- Deploying VPNs for Remote Access
- Deploying Router-to-Router VPNs
- Deploying Certificate-based Authentication for VPN Connections
- Virtual Private Network Implementation Examples
- Virtual Private Network Test Lab

# Deploying VPNs for Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying remote access VPNs**
This section covers:

- Planning for Remote Access VPNs
- Deploying Remote Access VPNs

# Planning for Remote Access VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for remote access VPNs**

This section covers:

- Remote access VPN design considerations
- Remote access VPN security considerations

# Remote access VPN design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access VPN design considerations**

To prevent problems, you should consider the following design issues before you implement remote access VPN connections:

- Choosing between PPTP-based connections or L2TP/IPSec connections
- Installing certificates
- Configuring firewall packet filters
- Creating a remote access policy for remote access VPN connections
- Using an IAS server
- Using Connection Manager

For more information, see Remote access VPN connection.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Choosing between PPTP-based connections or L2TP/IPSec connections**

A VPN server running Windows Server 2003 provides support for both Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP). When choosing between PPTP-based and L2TP/IPSec remote access VPN solutions, consider the following:

- PPTP can be used with a variety of Microsoft clients including Windows 95 (with the Dial-up Networking Upgrade 1.3 and later), Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000, and Windows XP. PPTP does not require a public key infrastructure (PKI) to issue computer certificates. By using encryption, PPTP-based VPN connections provide data confidentiality (captured packets cannot be interpreted without the encryption key). PPTP-based VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data origin authentication (proof that the data was sent by the authorized user).
- L2TP can only be used with client computers running Windows 2000 or Windows XP. L2TP supports either computer certificates or a preshared key as the authentication method for Internet Protocol security (IPSec). Computer certificate authentication, the recommended authentication method, requires a PKI to issue computer certificates to the VPN server computer and all VPN client computers. By using IPSec, L2TP/IPSec VPN connections provide data confidentiality, data integrity, and data authentication.

**Installing certificates**

In order to create L2TP/IPSec remote access VPN connections using computer certificate authentication for IPSec, you must install computer certificates, also known as machine certificates, on the VPN client and the VPN server. For more information, see Network access authentication and certificates and Computer certificates for L2TP/IPSec VPN connections.

**Configuring firewall packet filters**

If you have a firewall, you must configure packet filters on the firewall to allow traffic between VPN clients on the Internet and the VPN server. For more information, see VPN servers and firewall configuration.

**Creating a remote access policy for remote access VPN connections**
By using remote access policies, you can create a policy that requires remote access VPN connections to use a specific authentication method and encryption strength.

For example, you can create an Active Directory group called VPN Users whose members are the user accounts of the users creating remote access VPN connections across the Internet. Then, you create a policy with two conditions on the policy: **NAS-Port-Type** is set to **Virtual (VPN)** and **Windows-Group** is set to **VPN Users**. Finally, you configure the profile for the policy to select a specific authentication method and encryption strength.

For more information, see Introduction to remote access policies.

**Using an IAS server**
If you have multiple VPN servers running Windows Server 2003 , you need to configure remote access policies and logging for each VPN server. If you want to take advantage of centralized remote access policies and logging, you can configure the VPN servers as Remote Authentication Dial-In User Service (RADIUS) clients to a single computer (a RADIUS server) running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; and Internet Authentication Service (IAS).

You should also use an IAS server if you have VPN servers running Windows NT 4.0 with the Routing and Remote Access Service (RRAS), and you want to take advantage of remote access policies available in products in the Windows Server 2003 family. You cannot configure VPN servers running Windows NT 4.0 as RADIUS clients. You must upgrade a VPN server running Windows NT 4.0 to a VPN server running Windows NT 4.0 and RRAS.

For more information, see Using RADIUS for multiple remote access servers.

**Using Connection Manager**
For a large remote access VPN deployment, you can use Connection Manager and the Connection Manager Administration Kit to provide a custom dialer with preconfigured VPN connections to all remote access clients across your organization.

For more information about Connection Manager, see Connection Manager Administration Kit. For information about implementing VPN support with the Connection Manager Administration Kit, see Implementing VPN support.

# Remote access VPN security considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access VPN security considerations**
You can enhance remote access VPN security through:

- L2TP connections
- Strong authentication
- Data encryption
- PPTP or L2TP/IPSec packet filtering
- Firewall packet filtering
- Multi-use VPN servers
- Preventing traffic routed from remote access clients
- Remote access policy profile packet filtering

For more information, see Remote access VPN connection.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**L2TP connections**
A VPN server running Windows Server 2003 provides support for both Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP). By using encryption, PPTP-based VPN connections provide data confidentiality (captured packets cannot be interpreted without the encryption key). PPTP-based VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data origin authentication (proof that the data was sent by the authorized user). By using Internet Protocol Security (IPSec), L2TP/IPSec VPN connections provide data confidentiality, data integrity, and data origin authentication. However, L2TP can only be used with Windows 2000 and Windows XP clients.

L2TP/IPSec connections can use the following methods to authenticate the IPSec security association:

- A computer certificate

  Computer certificates are the recommended method of authentication as they provide strong authentication and are very difficult to impersonate or compromise. Computer certificate authentication requires a public key infrastructure (PKI) to issue computer certificates to the VPN server computer and all VPN client computers.

- A preshared key

  A preshared key is a string of text that is configured on both the VPN client and VPN server. Preshared key is a relatively weak authentication method, therefore it is only recommended that you use preshared key authentication while your PKI is being deployed or when preshared key is required by the VPN server to which your VPN clients are connecting.

**Strong user authentication**
For user authentication, use the strongest authentication scheme that is possible for your remote access VPN configuration. The strongest authentication scheme is the use of EAP-TLS with certificates. For more information, see Deploying Certificate-based Authentication for VPN Connections and Network access authentication and certificates.

Otherwise, use MS-CHAP version 2 authentication and enforce the use of strong passwords on your network. For more information, see MS-CHAP version 2.

**Data encryption**
For encryption, you can use either link encryption or end-to-end encryption with link encryption:

- Link encryption encrypts the data only on the link between the VPN client and the VPN server. For PPTP connections, you must use Microsoft Point-to-Point Encryption (MPPE) in conjunction with either MS-CHAP, MS-CHAP v2, or EAP-TLS authentication. For L2TP/IPSec connections, IPSec provides encryption on the link between the VPN client and the VPN server.

- End-to-end encryption encrypts the data between the source host and its final destination. You can use IPSec *after* the VPN connection is made to encrypt data from the source host to the destination host.

To require encryption, select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile that is used by your remote access VPN clients. For more information, see Configure encryption.

Encryption on client computers is configured in the properties of the network connection.

**PPTP or L2TP/IPSec filtering**
To secure the VPN server from sending or receiving any traffic on its Internet interface except VPN traffic, you need to ensure that PPTP or L2TP/IPSec input and output filters are configured on the interface that corresponds to the connection to the Internet. For more information about PPTP filters, see Add PPTP Filters. For more information about L2TP/IPSec filters, see Add L2TP over IPSec Filters.

Because IP routing is enabled on the Internet interface, if PPTP or L2TP/IPSec filters are not configured on the Internet interface, then any traffic received on the Internet interface is routed, which may forward unwanted Internet traffic to your intranet.

**Firewall packet filtering**
It is a common practice to provide protection to intranet hosts, such as a VPN server, from Internet hosts using a firewall. If you have a firewall, you must configure packet filters on the firewall to allow traffic to and from VPN clients on the Internet and the VPN server. For more information, see VPN servers and firewall configuration.

**Multi-use VPN servers**
Due to the way in which routes are created on VPN remote access clients, it is possible for the VPN client to send some unencrypted traffic to the VPN server rather than through the encrypted tunnel of the VPN connection. For example, the VPN client might connect to other services running on the VPN server without routing traffic across the VPN connection. Only traffic that is destined to the public IP address of the VPN server is sent in the clear. If traffic from the client uses the internal or private interface IP address of the VPN server, however, it will be encrypted.

When a remote access VPN client creates a VPN connection with a VPN server, it creates a series of routes in the IP routing table on the VPN client:

- Default route that uses the VPN connection

  The new default route for the VPN connection effectively replaces the existing default route for the duration of the connection. After the connection is made, all traffic that does not match an address on the directly connected network or the address of the VPN server is sent in an encrypted form over the VPN connection.

- Host route to the VPN server that uses the local network connection

  The host route for the address of the VPN server is created so that the VPN server is reachable using the locally attached network. If the host route is not present, VPN traffic to the VPN server cannot be sent.

The result of having the host route for the VPN server is that all traffic that is sent to applications running on the VPN server is not sent across the VPN connection, but is instead sent in an unencrypted form across the network between the VPN client and VPN server.

For example, when a remote access VPN client creates a VPN connection with a VPN server and then accesses a shared file on the VPN server computer, that traffic is not sent using the VPN connection. The file sharing traffic is sent in plaintext over the network between the VPN client and VPN server.

Additionally, if packet filters are configured on the VPN server that only allow VPN connection traffic, all other traffic sent to the VPN server is discarded. In this situation, all attempts to connect to services running on the VPN server will fail because traffic attempting to connect to those services are not sent over the VPN connection.

The key to which address is used by the VPN client to access services running on the VPN server lies in the way that the name of the VPN server is resolved. Typical users and applications refer to network resources using names, rather than IP addresses. The name must be resolved to an IP address using either DNS or WINS. If the intranet DNS and WINS infrastructures never contain a record mapping the name of the VPN server to the public IP address of the VPN server, traffic to services running on the VPN server is always sent across the VPN connection.

To prevent the VPN server from dynamically registering the public IP address of its Internet interface in the intranet DNS, obtain properties of the **Internet Protocol (TCP/IP)** component of the Internet connection in the **Network Connections** folder. Click **Advanced**. In the **Advanced TCP/IP Settings** dialog box, click the **DNS** tab, and then clear the **Register this connection's addresses in DNS** check box.

To prevent the VPN server from registering the public IP address of its Internet interface with intranet WINS servers, obtain properties of the **Internet Protocol (TCP/IP)** component of the Internet connection in the **Network Connections** folder. Click **Advanced**. In the **Advanced TCP/IP Settings** dialog box, click the **WINS** tab, and then click **Disable NetBIOS over TCP/IP**.

Before the VPN connection is made, the VPN client uses the Internet DNS infrastructure to resolve the name of the VPN server computer to its public IP addresses. After the VPN connection is made, assuming that intranet DNS and WINS servers are configured either during the PPP connection process or through the relaying of the DHCPInform message, the VPN client uses the intranet DNS and WINS infrastructures to resolve the name of the VPN server computer to its intranet IP addresses.

**Preventing traffic routed through remote access clients**

Once a remote access client successfully establishes a PPTP or L2TP connection, by default any packet sent over the connection is received by the VPN server and forwarded. Packets sent over the connection can include:

- Packets originated from the remote access client computer.
- Packets sent to the remote access client computer by other computers.

When the remote access client computer makes the VPN connection, by default it creates a default route so that all traffic that matches the default route is sent over the VPN connection. If other computers are forwarding traffic to the remote access VPN client, treating the remote access client computer as a router, then that traffic may also be forwarded across the VPN connection. This is a security problem because the computer that is sending traffic to the remote access VPN client has not been authenticated by the VPN server. The computer sending traffic to the remote access VPN client computer has physical network access to the same resources as the authenticated remote access VPN client computer.

To prevent the VPN server from sending traffic across the VPN connection for computers other than authenticated remote access VPN client computers, configure the following remote access policy packet filters on the remote access policy that is used for your VPN connections:

- From client filter:

  Filter action set to **Deny all traffic except those listed below**.

| IP Packet Filter Field | Setting |
|---|---|
| Source Address | User's address |
| Source Network Mask | User's mask |
| Destination Address | Any |
| Destination Network Mask | Any |
| Protocol | Any |

- To client filter:

  Filter action set to **Deny all traffic except those listed below**.

| IP Packet Filter Field | Setting |
|---|---|
| Source Address | Any |
| Source Network Mask | Any |
| Destination Address | User's address |
| Destination Network Mask | User's mask |
| Protocol | Any |

This set of IP packet filters will cause the VPN server to discard all traffic sent across the VPN connection except those that either originated from or is sent to authenticated remote access VPN clients.

**Remote access policy profile packet filtering**

To define the specific types of IP traffic that are allowed for remote access VPN connections, you need to configure IP packet filters on the profile for the remote access policy that is used for your remote access VPN connections. With profile packet filters, you can configure the IP traffic that is allowed from remote access clients (**From client** filters) or to remote access clients (**To client** filters) on an exception basis: either all traffic except traffic specified by filters or no traffic except traffic specified by filters. Remote access policy profile filtering applies to all remote access connections that match the remote access policy. For more information, see Elements of a remote access policy and Configure IP options.

# Deploying Remote Access VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying remote access VPNs**

This section covers:

- PPTP-based remote access VPN deployment
- L2TP-based remote access VPN deployment

# PPTP-based remote access VPN deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPTP-based remote access VPN deployment**

You can use remote access with products in the Windows Server 2003 family to provide access to a corporate intranet for remote access clients who are making PPTP connections across the Internet. If you want your remote access server to support multiple PPTP connections, complete the following steps:

- Configure the connection to the Internet.
- Configure the connection to the intranet.
- Configure the remote access server as a corporate intranet router.
- Configure the VPN server.
- Configure firewall packet filters.
- Configure remote access policies.

The following illustration shows the elements of a remote access server running Windows Server 2003 that provides PPTP-based remote access to a corporate intranet.




For more information, see Remote access VPN connection and Point-to-Point Tunneling Protocol.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Configuring the connection to the Internet**

The connection to the Internet from a computer running a Windows Server 2003 operating system is a dedicated connection--a WAN adapter installed in the computer. The WAN adapter is typically a DDS, T1, Fractional T1, or Frame Relay adapter. You must contract with a local telephone company to run the appropriate physical wiring to your premises. You need to verify that the WAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

The WAN adapter includes drivers that are installed in memobers of the Windows Server 2003 family so that the WAN adapter appears as a network adapter.

You need to configure the following TCP/IP settings on the WAN adapter:

- IP address and subnet mask assigned from the InterNIC or an Internet service provider (ISP).
- Default gateway of the ISP router.

So that VPN clients can connect to your VPN server by name (rather than IP address), you can request DNS registration of your VPN server by your ISP.

**Configuring the connection to the intranet**
The connection to the intranet from a computer running a Windows Server 2003 operating system is a LAN adapter that is installed in the computer. You need to verify that the LAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.
- DNS and WINS name servers of corporate intranet name servers.

**Configuring the remote access server as a corporate intranet router**
In order for the remote access server to properly forward traffic on the corporate intranet, you must configure it as a router with either static routes or routing protocols so that all of the locations of the intranet are reachable from the remote access server. For information about routing concepts, see Routing Overview. For information about setting up the remote access server as a router, see Deploying Routing.

**Configuring the VPN server**
You can configure your VPN server by running the Routing and Remote Access Server Setup Wizard. You can configure the following settings using the wizard:

- A basic firewall on the public interface.
- The method by which the VPN server assigns addresses to remote access clients (either using addresses that the VPN server obtains from a DHCP server or using addresses from a specified range of addresses that you have configured on the VPN server).
- Forwarding of authorization and authentication messages to a Remote Access Dial-In User Service (RADIUS) server (configuration of the VPN server as a RADIUS client).

Once the wizard is run, these Routing and Remote Access settings are automatically configured:

- Network interfaces
- PPTP and L2TP ports (5 or 128 of each, depending on your choices when running the wizard)
- Multicast support using Internet Group Messaging Protocol (IGMP)
- IP routing
- Installation of the DHCP Relay Agent component.

For more information on enabling Routing and Remote Access and running the wizard, see Enable the Routing and Remote Access service.

**Configuring firewall packet filters**
If you are using a firewall, you need to configure PPTP packet filters on your firewall to allow PPTP traffic between Internet-based VPN clients and the VPN server computer. For more information, see VPN servers and firewall configuration.

**Configuring remote access policies**
For an access-by-user administrative model, you need to set the remote access permission to **Allow access** on the user accounts for those users who will be making VPN connections. For an access-by-policy model, make the appropriate changes to the remote access permission of the user accounts. For more information, see Introduction to remote access policies.

To configure a remote access policy to control the authentication and encryption options for VPN connections, you can create a remote access policy with the following settings:

- Set **Policy name** to **VPN Access** (example).
- For conditions, set the **NAS-Port-Type** condition to **Virtual (VPN)** and the **Tunnel-Type** condition to **Point-to-Point Tunneling Protocol**.
- Select the **Grant remote access permission** option.
- For profile settings, select the appropriate authentication and encryption options.

Then, either delete the default remote access policies or position them after the new policy. This remote access policy allows all users with remote access permission to create a VPN connection.

If you want to distinguish dial-up remote access users from VPN remote access users, do the following:

1. Create an Active Directory group whose members can create virtual private networking connections with the VPN server. For example, VPN_Users.
2. Add the appropriate user accounts to the new Active Directory group.
3. Create a new remote access policy with the following properties:
    - Set **Policy name** to **VPN Access if member of VPN_Users** (example).
    - For conditions, set the **Windows-Groups** condition to **VPN_Users** (example), set the **NAS-Port-Type** condition to **Virtual (VPN)**, and set the **Tunnel-Type** condition to **Point-to-Point Tunneling Protocol**.
    - Select the **Grant remote access permission** option.
4. Move the default remote access policies after the new policy.

The default encryption settings allow no encryption and all levels of encryption strength. To require encryption, clear the **No Encryption** option and select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile. The encryption strengths are:

- **Basic**

  You can use this option when communicating with older Microsoft dial-up networking clients. This option uses Microsoft Point-to-Point Encryption (MPPE) and a 40-bit encryption key.

- **Strong**

  You can use this option when communicating with dial-up networking clients running Windows 98, Windows 2000, or Windows XP. This option uses MPPE and a 56-bit encryption key.

- **Strongest**

  You can use this option when communicating with dial-up networking clients running Windows 98, Windows 2000, or Windows XP. This option uses MPPE and a 128-bit encryption key.

For more information, see Configure encryption.

# L2TP-based remote access VPN deployment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**L2TP-based remote access VPN deployment**

You can use remote access with products in the Windows Server 2003 family to provide access to a corporate intranet for remote access clients who are making Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec) connections across the Internet. If you want your remote access server to support multiple L2TP/IPSec connections, complete the following steps:

- Configure the connection to the Internet.
- Configure the connection to the intranet.
- Configure the remote access server as a corporate intranet router.
- Configure the VPN server.
- Install certificates.
- Configure firewall packet filters.
- Configure remote access policies.

The following illustration shows the elements of a remote access server running Windows Server 2003 that provides L2TP/IPSec remote access to a corporate intranet.




For more information, see Remote access VPN connection and Layer Two Tunneling Protocol.

### Notes

- The following configuration assumes that computer certificates, also known as machine certificates, are already installed on the VPN server and remote access client computers. For more information, see Computer certificates for L2TP/IPSec VPN connections and Network access authentication and certificates.
- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

**Configuring the connection to the Internet**

The connection to the Internet from a computer running a Windows Server 2003 family operating system is a dedicated connection--a WAN adapter installed in the computer. The WAN adapter is typically a DDS, T1, Fractional T1, or Frame Relay adapter. You must contract with a local telephone company to run the appropriate physical wiring to your premises. You need to verify that the WAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

The WAN adapter includes drivers that are installed in members of the Windows Server 2003 family so that the WAN adapter appears as a network adapter.

You need to configure the following TCP/IP settings on the WAN adapter:

- IP address and subnet mask assigned from the InterNIC or an Internet service provider (ISP).
- Default gateway of the ISP router.

To enable VPN clients to connect to your VPN server by name rather than by IP address, request that your ISP register your VPN server in DNS.

**Configuring the connection to the intranet**
The connection to the intranet from a computer running a Windows Server 2003 operating system is a LAN adapter that is installed in the computer. You need to verify that the LAN adapter is compatible with products in the Windows Server 2003 family. To verify compatibility, see the Compatible Hardware and Software section at Support resources.

You need to configure the following TCP/IP settings on the LAN adapter:

- IP address and subnet mask assigned from the network administrator.
- DNS and WINS name servers of corporate intranet name servers.

**Configuring the remote access server as a corporate intranet router**
In order for the remote access server to properly forward traffic on the corporate intranet, you must configure it as a router with either static routes or routing protocols so that all of the locations of the intranet are reachable from the remote access server. For information about routing concepts, see Routing Overview. For information about setting up the remote access server as a router, see Deploying Routing.

**Configuring the VPN server**
You can configure your VPN server by running the Routing and Remote Access Server Setup Wizard. You can use the wizard to configure the following settings:

- A basic firewall on the public interface.
- The method by which the VPN server assigns addresses to remote access clients (either using addresses that the VPN server obtains from a DHCP server or using addresses from a specified range of addresses that you configure).
- Forwarding of authorization and authentication messages to a Remote Authentication Dial-In User Service (RADIUS) server (configuration of the VPN server as a RADIUS client).

Once the wizard is run, the following Routing and Remote Access settings are automatically configured:

- Network interfaces
- PPTP and L2TP ports (five or 128 of each, depending on the choices you made when you ran the wizard)
- Multicast support using Internet Group Management Protocol (IGMP)
- IP routing
- Installation of the DHCP Relay Agent component.

For more information about enabling Routing and Remote Access and running the wizard, see Enable the Routing and Remote Access service.

**Installing certificates**
In order to create L2TP/IPSec remote access VPN connections using computer certificate authentication for IPSec, you must install computer certificates, also known as machine certificates, on the VPN client and the VPN server. For more information, see Network access authentication and certificates and Computer certificates for L2TP/IPSec VPN connections.

**Configuring firewall packet filters**

If you are using a firewall, you need to configure L2TP/IPSec packet filters on your firewall to allow L2TP/IPSec traffic between Internet-based VPN clients and the VPN server computer. For more information, see VPN servers and firewall configuration.

**Configuring remote access policies**

For an access-by-user administrative model, you need to set the remote access permission to **Allow access** on the user accounts for those users who will be making VPN connections. For an access-by-policy model, make the appropriate changes to the remote access permission of the user accounts. For more information, see Introduction to remote access policies.

To configure a remote access policy to control the authentication and encryption options for VPN connections, create a remote access policy with the following settings:

- Set **Policy name** to **VPN Access** (example).
- For conditions, set the **NAS-Port-Type** condition to **Virtual (VPN)** and the **Tunnel-Type** condition to **Layer Two Tunneling Protocol**.
- Select the **Grant remote access permission** option.
- For profile settings, select the appropriate authentication and encryption options.

Then, either delete the default remote access policies or position them after the new policy. This remote access policy allows all users with remote access permission to create a VPN connection.

If you want to distinguish dial-up remote access users from VPN remote access users, do the following:

1. Create an Active Directory group whose members can create virtual private networking connections with the VPN server. For example, VPN_Users.
2. Add the appropriate user accounts to the new Active Directory group.
3. Create a new remote access policy with the following properties:
    - Set **Policy name** to **VPN Access if member of VPN_Users** (example).
    - For conditions, set the **Windows-Groups** condition to **VPN_Users** (example), set the **NAS-Port-Type** condition to **Virtual (VPN)**, and set the **Tunnel-Type** condition to **Layer Two Tunneling Protocol**.
    - Select the **Grant remote access permission** option.
4. Move the default remote access policies after the new policy.

The default encryption settings allow no encryption and all levels of encryption strength. To require encryption, clear the **No Encryption** option and select the appropriate encryption strengths on the **Encryption** tab of the remote access policy profile. The encryption strengths are:

- **Basic**

  This option uses 56-bit DES encryption.

- **Strong**

  This option uses 56-bit DES encryption.

- **Strongest**

  This option uses triple DES (3DES) encryption.

For more information, see Configure encryption.

# Deploying Certificate-based Authentication for VPN Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying certificate-based authentication for VPN connections**
The use of certificates for authentication of VPN connections is the strongest form of authentication in the Windows Server 2003 family. You must use certificate-based authentication for Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec)-based VPN connections and with smart cards. With smart cards, you must use the Extensible Authentication Protocol (EAP) with the **Smartcard or other certificate (TLS)** EAP type, also known as EAP-Transport Level Security (EAP-TLS).

This section covers:

- Computer certificates for L2TP/IPSec VPN connections
- Smart cards and remote access VPN connections
- Deploying certificate-based authentication for demand-dial routing

# Computer certificates for L2TP/IPSec VPN connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computer certificates for L2TP/IPSec VPN connections**
The Windows Server 2003 family supports two authentication methods for Layer Two Tunneling Protocol over Internet Protocol security (L2TP/IPSec)-based VPN connections: computer certificates, also known as machine certificates, and preshared keys.

In order to create an L2TP/IPSec connection using the computer certificate authentication method, you must install a certificate in the local computer certificate store on the VPN client and VPN server computer. To install a computer certificate, a certification authority must be present to issue certificates. Once the certification authority is configured, you can install a certificate in three different ways:

- By configuring the automatic enrollment, or auto-enrollment, of computer certificates to computers in a Windows Server 2003 domain.
- By using the Certificates snap-in to obtain a computer certificate.
- By using your browser to connect to the CA Web enrollment pages to install a certificate on the local computer or to a floppy disk for installation on another computer, such as a user's home computer.

Based on the certificate policies in your organization, you need to perform only one of these allocations.

Non-domain member computers cannot obtain certificates through auto-enrollment. For more information, see Network access authentication and certificates.

To configure a CA and install the computer certificate, perform the following steps:

1. If you do not already have an enterprise root CA:
   a. Promote the computer that will be a CA to a domain controller (DC), if necessary.
   b. Install the Certificate Services component as an enterprise root CA on a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. For more information, see Install an enterprise root certification authority.
2. To auto-enroll computer certificates, configure Windows Server 2003 domain. For more information, see Configure automatic certificate allocation from an enterprise CA.

   To create a computer certificate for the VPN server that is a member of the domain for which auto-enrollment is configured (as well as other computers that are members of the domain), restart the computer or type **gpupdate /target:computer** from a Windows Server 2003 command prompt.
3. To manually enroll computer certificates, use the Certificates snap-in to install the CA root certificate. For more information, see Use Windows Server 2003 Certificate Services Web Pages, Manage certificates for a computer and Request a certificate.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Smart cards and remote access VPN connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Smart cards and remote access VPN connections**
The use of smart cards for user authentication is the strongest form of authentication in the Windows Server 2003 family. For remote access VPN connections, you must use Extensible Authentication Protocol (EAP) with the **Smart card or other certificate (TLS)** EAP type, also known as EAP-Transport Level Security (EAP-TLS).

To use smart cards for remote access VPN authentication, you must do the following:

- Configure the VPN server to provide remote access VPN services.
- Install a computer certificate on the VPN server.
- Enable a smart card logon process for the domain.
- Enable EAP and configure the **Smartcard or other certificate (TLS)** EAP type on the VPN server computer.
- Enable smart card authentication on the VPN connection on the remote access client computer.

**Configuring the VPN server to provide remote access VPN services**
Configure the VPN server running Windows Server 2003 as described in Deploying Remote Access VPNs.

**Installing a computer certificate on the VPN server**
In order to configure EAP-TLS on the VPN server computer, you must install a computer certificate, also known as a machine certificate. To install a computer certificate on the VPN server computer, a certification authority (CA) must be present to issue certificates. Once the CA is configured, you can install a certificate on the VPN server computer in one of three different ways:

- By configuring the automatic enrollment, or auto-enrollment, of computer certificates to computers in a Windows Server 2003 domain.
- By using the Certificates snap-in to request and install a computer certificate.
- By using your browser to connect to the CA Web enrollment pages to request and install a certificate on the local computer or to a floppy disk for installation on another computer, such as a user's home computer.

Based on the certificate policies in your organization, you need to perform only one of these enrollment methods.

For more information, see Network access authentication and certificates.

To configure a CA and install the computer certificate, perform the following steps:

1. If you do not already have an enterprise root CA:

   a. Promote the computer that will be a CA to a domain controller (DC), if necessary.

   b. Install the Certificate Services component on a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; as an enterprise root CA. For more information, see Install an enterprise root certification authority.

2. To auto-enroll computer certificates, configure the Windows Server 2003 domain. For more information, see Configure automatic certificate allocation from an enterprise CA.

   To create a computer certificate for the VPN server that is a member of the domain for which auto-enrollment is configured (as well as other computers that are members of the domain), restart the computer or type **gpupdate**

**/target:computer** from a Windows Server 2003 command prompt.

3. To manually enroll computer certificates, use the Certificates snap-in to install the CA root certificate, or use your browser to connect to the CA Web enrollment pages. For more information, see Use Windows Server 2003 Certificate Services Web Pages, Manage certificates for a computer, and Request a certificate.

### Note

- VPN servers that terminate connections for remote users must have a certificate that contains the Server Authentication purpose in the Enhanced Key Usage extensions. VPN servers that are used for site-to-site VPN connections both originate and terminate connections. For this reason the certificate on these servers must contain both the Server Authentication purpose and the Client Authentication purpose in the Enhanced Key Usage extensions. For more information, see Network access authentication and certificates.

#### Enabling a smart card logon process for the domain

To enable a smart card logon process for the domain, complete the following procedures:

1. Prepare a certification authority to issue smart card certificates
2. Prepare a smart card certificate enrollment station
3. Set up a smart card for user logon

#### Configuring the VPN server for smart card remote access

To configure the VPN server running Windows Server 2003 for smart card remote access, see Configure smart card remote access.

#### Configuring the remote access client for smart card remote access

You need to install a smart card reader on the remote access client computer. For more information, see Install a smart card reader on a computer.

Once a smart card reader is installed on the computer running Windows XP, then you are prompted about whether you want to use the smart card for authentication when you create dial-up or VPN connections.

For existing dial-up or VPN connections, you can enable smart card authentication from the properties of the dial-up or VPN connection. For more information, see Enable smart card or other certificate authentication.

### Note

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Deploying certificate-based authentication for demand-dial routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying certificate-based authentication for demand-dial routing**

The use of certificates for authentication of calling routers is the strongest form of authentication in the Windows Server 2003 family. For certificate-based authentication of demand-dial connections, you must use the Extensible Authentication Protocol (EAP) with the **Smart card or other certificate (TLS)** EAP type, also known as EAP-Transport Level Security (EAP-TLS). EAP-TLS requires the use of user certificates for the calling router and machine certificates (also known as computer certificates) for the answering router.

The deployment of certificate-based authentication for demand-dial routing typically occurs in the following situations:

- Business partner demand-dial connection
- Branch office demand-dial connection

# Business partner demand-dial connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Business partner demand-dial connection**
To use certificates for a two-way initiated, mutually authenticated, demand-dial configuration between two business partners (in this example, Company A and Company B), you must perform the following:

- Configure the calling and answering routers for demand-dial routing.
- Install computer certificates on the calling router and answering router computers.
- Configure the domain for Web-based certificate enrollment.
- At Company A, create a user account for the Company B router and export a certificate for the user account.
- At Company B, create a user account for the Company A router and export a certificate for the user account.
- At Company A, import the certificate from Company B.
- Configure the Company A router to support certificate-based authentication as a calling router and as an answering router.
- At Company B, import the certificate from Company A.
- Configure the Company B router to support certificate-based authentication as a calling router and as an answering router.

**Configuring the calling and answering routers for demand-dial routing**
Configure the Routing and Remote Access calling and answering routers as described in Deploying demand-dial routing for dial-in demand-dial routing or Router-to-router VPN Deployment for VPN demand-dial routing.

**Installing computer certificates on the calling router and answering router computers**
To use EAP-TLS, a computer certificate (also known as a machine certificate) must be installed on the authenticating server and the remote access client. In order to install a computer certificate, a certification authority must be present to issue certificates. Once the certification authority is configured, you can install a certificate in three different ways:

- By configuring the automatic enrollment, or autoenrollment, of computer certificates to computers in a Windows Server 2003 domain.
- By using the Certificates snap-in to obtain a computer certificate.
- By using your browser to connect to the CA Web enrollment pages to install a certificate on the local computer or to a floppy disk for installation on another computer, such as non-domain member computers that cannot obtain a certificate through autoenrollment.

Based on the certificate policies in your organization, you only need to perform one of these allocations.

For more information, see Network access authentication and certificates.

To configure a certification authority and install the computer certificate, perform the following steps:

1. Install the Certificate Services component as an enterprise root certification authority (CA). This step is only necessary if you do not already have an enterprise root CA.

    a. If necessary, promote the computer that will be a CA to a domain controller (DC).

b. Install the Certificate Services component as an enterprise root CA. For more information, see Install an enterprise root certification authority.

2. Configure the CA to issue certificates that permit exportable keys. To do so, you must clone and rename the **router (offline request)** template, choose to make the keys exportable, and add the new template to the templates the CA can use to issue certificates. For more information, see To establish the certificate types that an enterprise certification authority can issue.

3. Do one of the following:

   - To auto-enroll computer certificates, configure the domain. For more information, see Configure automatic certificate allocation from an enterprise CA.

     To create a computer certificate for the calling or answering router that is a member of the domain for which auto-enrollment is configured (as well as other computers that are members of the domain), restart the computer or type **gpupdate /Target:Computer /Force** from the command prompt.

   - To manually enroll computer certificates, use the Certificates snap-in or the CA Web enrollment pages to install the CA root certificate. For more information, see Manage certificates for a computer and Request a certificate.

In order for the CA to issue certificates for the calling router, you must configure the domain for Web-based enrollment. For more information, see Set up certification authority Web enrollment support.

**Creating a user account and exporting its certificate for the Company B router**

To create a dial-in user account for the Company B router and export the user certificate of the user account, do the following:

1. Log on as a domain administrator.

2. Create a user account that the Company B router will use when it dials the Company A router. For more information, see Create a new user account.

3. Obtain a certificate that has an exportable key from the certification authority through Web-based enrollment. This certificate might be called **router (offline request)**, or have another name. For more information, see Install a router (offline request) certificate.

4. Export the exportable key certificate to a .cer file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, do not export the private key.

5. Map the newly created certificate (the .cer file) to the user account that was created for the Company B router. For more information, see Map a certificate to a user account.

6. Export the certificate to a .pfx file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, export the private key, select the **Delete the private key if the import is successful** check box, and click **Include all certificates in the certification path if possible**. Save this file to a floppy disk to send to the network administrator at Company B.

7. Send the floppy disk that contains the Company B dial-in account user certificate file to the network administrator at Company B.

**Creating a user account and exporting its certificate for the Company A router**

To create a dial-in user account for the Company A router and export the user certificate of the user account, do the following:

1. Log on as a domain administrator.

2. Create a user account that the Company A router will use when it dials the Company B router. For more information, see

Create a new user account.

3. Obtain a certificate that has exportable keys from the certification authority through Web-based enrollment. This certificate might be called **router (offline request)**, or it might have another name. For more information, see Install a router (offline request) certificate.

4. Export the certificate to a .cer file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, do not export the private key.

5. Map the newly created certificate (the .cer file) to the user account created for the Company A router. For more information, see Map a certificate to a user account.

6. Export the certificate to a .pfx file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, export the private key, select the **Delete the private key if the import is successful** check box, and click **Include all certificates in the certification path if possible**. Save this file to a floppy disk to send to the network administrator at Company A.

7. Send the floppy disk that contains the Company A dial-in account user certificate file to the network administrator at Company A.

**Importing the certificates from Company B**
Upon receipt at Company A of the floppy disk that contains the certificate file from Company B, on the Company A router, import the user certificate. For more information, see Import a certificate.

**Configuring the Company A router to support certificate-based authentication**
To configure the Company A router for certificate-based authentication as an answering router, see Configure the answering router for certificate-based EAP.

To configure the Company A router for certificate-based authentication as a calling router, see Configure the calling router for certificate-based EAP.

**Importing the certificates from Company A**
Upon receipt at Company B of the floppy disk that contains the certificate files from Company A, on the Company B router, import the user certificate. For more information, see Import a certificate.

**Configuring the Company B router to support certificate-based authentication**
To configure the Company B router for certificate-based authentication as an answering router, see Configure the answering router for certificate-based EAP.

To configure the Company B router for certificate-based authentication as a calling router, see Configure the calling router for certificate-based EAP.

# Branch office demand-dial connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Branch office demand-dial connection**
To use certificates for a two-way initiated, mutually authenticated, demand-dial configuration between two routers in the same organization (in this example, a branch office router and a corporate office router), you must perform the following:

- Configure the calling and answering routers for demand-dial routing.
- Install a computer certificate on the corporate office router.
- Configure the domain for Web-based certificate enrollment.
- Create user accounts and export certificates.
- Import the dial-out user certificate on the corporate office router.
- Configure the corporate office router to support certificate-based authentication as a calling router and as an answering router.
- Import the dial-in certificate on the branch office router.
- Configure the branch office router to support certificate-based authentication as a calling router.
- Connect to the corporate office and join the organization domain.

**Configuring the calling and answering routers for demand-dial routing**
Configure the calling and answering routers as described in Deploying demand-dial routing for dial-up demand-dial routing or Router-to-router VPN Deployment for VPN demand-dial routing.

**Installing a computer certificate on the corporate office router**
In order to configure EAP-TLS on the corporate office router, you must install a computer certificate (also known as a machine certificate). In order to install a computer certificate, a certification authority must be present to issue certificates. Once the certification authority is configured, you can install a certificate in three different ways:

- By configuring the automatic enrollment, or autoenrollment, of computer certificates to computers in a Windows Server 2003 domain.
- By using the Certificates snap-in to obtain a computer certificate.
- By using your browser to connect to the CA Web enrollment pages to install a certificate on the local computer or to a floppy disk for installation on another computer, such as non-domain member computers that cannot obtain a certificate through autoenrollment.

Based on the certificate policies in your organization, you only need to perform one of these allocations.

For more information, see Network access authentication and certificates.

To configure a certification authority and install the computer certificate, perform the following steps:

1. Install the Certificate Services component as an enterprise root certification authority. This step is only necessary if you do not already have an enterprise root certification authority (CA).
    a. If necessary, promote the computer that will be a CA to a domain controller (DC).
    b. Install the Certificate Services component as an enterprise root CA. For more information, see Install an enterprise root

certification authority.

2. Configure the CA to issue certificates with exportable keys. For more information, see To establish the certificate types that an enterprise certification authority can issue.

3. Do one of the following:

   - To auto-enroll computer certificates, configure the domain. For more information, see Configure automatic certificate allocation from an enterprise CA.

     To create a computer certificate for the calling or answering router that is a member of the domain for which autoenrollment has been configured (as well as other computers that are members of the domain), restart the computer or type **gpupdate /Target:***Computer***/Force** from the command prompt.

   - To manually enroll computer certificates, use the Certificates snap-in or the CA Web enrollment pages to install the CA root certificate. For more information, see Manage certificates for a computer and Request a certificate.

In order for the CA to issue certificates for the calling router, you must configure the domain for Web-based enrollment. For more information, see Set up certification authority Web enrollment support.

**Creating user accounts and exporting certificates**

To create dial-in and dial-out user accounts and export certificates, do the following:

1. Log on as a domain administrator.

2. Create a user account that the corporate office router will use when it dials the branch office router (the dial-out account). For more information, see Create a new user account.

3. Obtain a certificate that has exportable keys for the dial-out account from the certification authority through Web-based enrollment. This certificate might be called **router (offline request)**, or it might have another name. For more information, see Install a router (offline request) certificate.

4. Export the certificate for the dial-out account to a .cer file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, do not export the private key.

5. Map the newly created certificate (the .cer file) to the dial-out user account. For more information, see Map a certificate to a user account.

6. Export the certificate of the dial-out account to a .pfx file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, export the private key and click **Delete the private key if the import is successful** and select the option to **Include all certificates in the certification path if possible**.

7. Create a user account that the branch office router will use when it dials the corporate office router (the dial-in account). For more information, see Create a new user account.

8. Obtain a certificate that has exportable keys for the dial-in account from the certification authority through Web-based enrollment. This certificate might be called **router (offline request)**, or it might have another name. For more information, see Install a router (offline request) certificate.

9. Export the certificate for the dial-in account to a .cer file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, do not export the private key.

10. Map the newly created certificate (the .cer file) to the dial-in user account. For more information, see Map a certificate to a user account.

11. Export the certificate of the dial-in account to a .pfx file. For more information, see Export a certificate. Within the Certificates snap-in Export wizard, export the private key and click **Delete the private key if the import is successful**. Save this file to a floppy disk to send to the network administrator at the branch office.

12. Send the floppy disk that contains the dial-in account user certificate file to the network administrator at the branch office.

**Importing the dial-out certificate on the corporate office router**
On the corporate office router, import the user certificate for the dial-out account. For more information, see Import a certificate.

**Configuring the corporate office router to support certificate-based authentication**
To configure the corporate office router for certificate-based authentication as an answering router, see Configure the answering router for certificate-based EAP.

To configure the corporate office router for certificate-based authentication as a calling router, see Configure the calling router for certificate-based EAP.

**Importing the certificate on the branch office router**
Upon receipt at the branch office of the floppy disk that contains the certificate file from the corporate office, import the user certificate for the dial-in account. For more information, see Import a certificate.

**Configuring the branch office router to support certificate-based authentication**
To configure the branch office router for certificate-based authentication as a calling router, see Configure the calling router for certificate-based EAP.

**Connecting to the corporate office and joining the organization domain**
To connect to the corporate office and join the organization domain, do the following:

1. From the branch office, connect to the corporate office by right-clicking the demand-dial interface, and then clicking **Connect**.

2. Once connected, the branch office router joins the domain through the **Computer Name** tab (in the properties of My Computer).

3. After joining the domain, restart the branch office router.

4. After restarting the branch office router, connect to the corporate office router again.

5. Once connected, the branch office router receives domain policy and a computer certificate (if auto-enrollment of computer certificates is configured). If auto-enrollment of computer certificates is not configured, obtain a computer certificate through the Certificates snap-in. For more information, see Manage certificates for a computer and Request a certificate.

6. Once a computer certificate is obtained, configure the branch office router for certificate-based authentication as an answering router. For more information, see Configure the answering router for certificate-based EAP.

At this point, you can install a domain controller in the branch office by using the demand-dial connection to the corporate office. For more information on installing a domain controller, see Checklist: Installing a domain controller.

### Notes

- The ability of the branch office router to join the domain and the installation of a domain controller depends on DNS name resolution. Ensure that both the router and the domain controller computer are configured with the proper DNS server IP addresses.

- By default, an answering router checks the certificate revocation list when authenticating the calling router. Because the root CA computer is always reachable by the corporate office router, the certificate revocation list can always be checked. However, the root CA computer is not reachable by the branch office router until after the connection is made. If the root CA computer cannot be reached, then Active Directory is checked. In this case, the branch office router accesses its local domain controller for the revocation list. If the certificate revocation list is not published in Active Directory, then the branch office router acting as the answering router rejects the connection attempt. To prevent this problem, do one of the following:

    - Publish the certificate revocation list in Active Directory. For more information, see Schedule the publication of the certificate revocation list or Manually publish the certificate revocation list.

    - On the branch office router, set the following registry value to 1:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RasMan\PPP\EAP\13\IgnoreRevocationOffline**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

# Virtual Private Network Implementation Examples

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private network implementation examples**
This section describes how common VPN scenarios are configured for a fictional company by using computers running Windows Server 2003 operating systems. While your network configuration may be different than those described here, you can apply the basic concepts.

Electronic, Inc. is an electronics design and manufacturing company with a main corporate campus in New York and branch offices and distribution business partners throughout the United States. Electronic, Inc. has implemented a VPN solution by using computers running Windows Server 2003, Standard Edition, that connect remote access users, branch offices, and business partners.

The VPN server at the corporate office provides both remote access and router-to-router PPTP and L2TP/IPSec VPN connections and the routing of packets to intranet and Internet locations. For more information, see Common configuration for the VPN server.

Based on the common configuration of the VPN server, the following virtual private network scenarios are described:

- Remote access for employees
- On-Demand Branch Office
- Persistent Branch Office
- Extranet for Business Partners
- Dial-up and VPNs with RADIUS

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Common configuration for the VPN server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Common configuration for the VPN server**
To deploy a VPN solution for Electronic, Inc., the network administrator performs an analysis and makes design decisions regarding:

- The network configuration
- The remote access policy configuration
- The domain configuration
- The security configuration

**The network configuration**
The key elements of the network configuration are:

- The Electronic, Inc. corporate intranet uses the private networks of 172.16.0.0 with a subnet mask of 255.240.0.0 and 192.168.0.0 with a subnet mask of 255.255.0.0. The corporate campus network segments use subnets of 172.16.0.0 and the branch offices use subnets of 192.168.0.0.
- The VPN server computer is directly attached to the Internet by using a T3 (also known as a DS-3) dedicated WAN link.
- The IP address of the WAN adapter on the Internet is 207.209.68.1 as allocated by the Internet service provider (ISP) for Electronic, Inc. The IP address of the WAN adapter is referred to on the Internet by the domain name vpn.electronic.microsoft.com.
- The VPN server computer is directly attached to an intranet network segment that contains a RADIUS server, a file and Web server for business partner access, and a router that connects to the rest of the Electronic, Inc. corporate campus intranet. The intranet network segment has the IP network ID of 172.31.0.0 with the subnet mask of 255.255.0.0.
- The VPN server computer is configured with a static pool of IP addresses to allocate to remote access clients and calling routers. The static pool of IP addresses is a subset of the intranet network segment (an on-subnet address pool).

The following illustration shows the network configuration of the Electronic, Inc. VPN server.




**Note**

- While the illustration shows the RADIUS server and file and Web server, the RADIUS server is not discussed until Dial-up and VPNs with RADIUS, and the file and Web servers are not discussed until Extranet for Business Partners.

Based on the network configuration of the Electronic, Inc. corporate campus intranet, the VPN server computer is configured as follows.

**1. Install hardware in the VPN server**
The network adapter that is used to connect to the intranet segment and the WAN adapter that is used to connect to the Internet are installed according to the adapter manufacturer's instructions. Once drivers are installed and functioning, both adapters appear as local area connections in the Network Connections folder.

**2. Configure TCP/IP on the LAN and WAN adapter**
For the LAN adapter, an IP address of 172.31.0.1 with a subnet mask 255.255.0.0 is configured. For the WAN adapter, an IP address of 207.209.68.1 with a subnet mask 255.255.255.255 is configured. A default gateway is not configured for either adapter. DNS and WINS server addresses are also configured.

**3. Install the Routing and Remote Access service**
The Routing and Remote Access Server Setup Wizard is run. Within the wizard, the **Remote Access (dial-up or VPN)** option is selected. For more information, see Enable the Routing and Remote Access service.

While running the wizard, a static IP address pool with a starting IP address of 172.31.255.1 and an ending IP address of 172.31.255.254 is configured. This creates a static address pool for up to 253 VPN clients.

For more information, see Create a static IP address pool.

The default method of authenticating remote access and demand-dial connections is to use Windows authentication, which is appropriate in this configuration containing only one VPN server. For information on the use of RADIUS authentication for Electronic, Inc., see Dial-up and VPNs with RADIUS. For more information on the use of Windows and RADIUS authentication, see Authentication vs. authorization.

**4. Enable the EAP authentication method**
To enable the use of smart card-based remote access VPN clients and certificate-based calling routers, the network administrator enables Extensible Authentication Protocol (EAP) on the VPN server.

For more information, see Enable EAP.

**5. Configure static routes on the VPN server to reach intranet and Internet locations**
To reach intranet locations, a static route is configured with the following settings:

- Interface: The LAN adapter attached to the intranet
- Destination: 172.16.0.0
- Network mask: 255.240.0.0
- Gateway: 172.31.0.2
- Metric: 1

This static route simplifies routing by summarizing all destinations on the Electronic, Inc. intranet. This static route is used so that the VPN server does not need to be configured with a routing protocol.

To reach Internet locations, a static route is configured with the following settings:

- Interface: The WAN adapter attached to the Internet
- Destination: 0.0.0.0
- Network mask: 0.0.0.0
- Gateway: 0.0.0.0
- Metric: 1

This static route summarizes all destinations on the Internet. This route allows the VPN server to respond to a remote access client or demand-dial router VPN connection from anywhere on the Internet.

**Note**

- Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway. The gateway address of 0.0.0.0 is an example. 0.0.0.0 is the unspecified IP address.

**Setting the phone number for the PPTP and L2TP devices**

To aid in the configuration of remote access policies to confine VPN connections from Internet users, the port properties for the **WAN Miniport (PPTP)** and **WAN Miniport (L2TP)** devices are modified with the IP address of the VPN server's Internet interface in the **Phone number for this device** field. For more information, see Set the phone number on a port.

**Configure a static route on the intranet router to reach all branch offices**

To reach branch office locations from the intranet router, a static route is configured with the following settings:

- Interface: The LAN adapter attached to the intranet
- Destination: 192.168.0.0
- Network mask: 255.255.0.0
- Gateway: 172.31.0.1
- Metric: 1

This static route simplifies routing by summarizing all destinations at branch offices of Electronic, Inc.

**The remote access policy configuration**

Electronic, Inc. has migrated to a Windows 2000 native domain and the network administrator for Electronic, Inc. decides on an access-by-policy administrative model. The remote access permission on all user accounts is set to **Control access through Remote Access Policy**. The granting of remote access permission to connection attempts is controlled by the remote access permission setting on the first matching remote access policy. Remote access policies are used to apply different VPN connection settings based on group membership, and the default remote access policies are deleted.

For more information, see Introduction to remote access policies.

**The domain configuration**

To take advantage of the ability to apply different connection settings to different types of VPN connections, the following Active Directory groups are created:

- VPN_Users

  Used for remote access VPN connections

- VPN_Routers

  Used for router-to-router VPN connections from Electronic, Inc. branch offices

- VPN_Partners

  Used for router-to-router VPN connections from Electronic, Inc. business partners

**Note**

- All users and groups in this implementation example are created in the electronic.microsoft.com Active Directory domain.

**The security configuration**

To enable L2TP/IPSec connections, the use of smart cards by remote access clients, and the use of EAP-TLS by routers, the

Electronic, Inc. domain is configured to auto-enroll computer certificates to all domain members.

For more information, see Deploying Certificate-based Authentication for VPN Connections.

**Notes**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Remote access for employees

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access for employees**
Remote access for Electronic, Inc. employees is deployed by using remote access VPN connections across the Internet based on the settings configured in Common configuration for the VPN server and the following additional settings.

The following illustration shows the Electronic, Inc. VPN server that provides remote access VPN connections.




**Domain configuration**
For each employee that is allowed VPN access:

- The remote access permission on the dial-in properties of the user account is set to **Control access through Remote Access Policy**.
- The user account is added to the VPN_Users Active Directory group.

**Remote access policy configuration**
To define the authentication and encryption settings for remote access VPN clients, the following remote access policy is created:

- Policy name: Remote Access VPN Clients
- Conditions:
    - **NAS-Port-Type** is set to Virtual (VPN)
    - **Windows-Groups** is set to VPN_Users
    - **Called-Station-ID** is set to 207.209.68.1
- Permission is set to **Grant remote access permission**
- Profile settings:
    - **Authentication** tab: **Extensible Authentication Protocol** is enabled and **Smartcard or other certificate (TLS)** is configured to use the installed computer certificate (also known as the machine certificate). **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** and **Microsoft Encrypted Authentication (MS-CHAP)** are also enabled.
    - **Encryption** tab: **Strong** and **Strongest** are the only options that are selected.

### Note

- The **Called-Station-ID** condition is set to the IP address of the Internet interface for the VPN server. Only tunnels initiated from the Internet are allowed. Tunnels initiated from the Electronic, Inc. intranet are not permitted. Electronic, Inc. users that require Internet access from the Electronic, Inc. intranet must go through the Electronic, Inc. proxy server (not shown), where Internet access is controlled and monitored.

#### PPTP-based remote access client configuration

The New Connection Wizard is used on client computers to create a VPN connection with the following setting:

- Host name or IP address: vpn.electronic.microsoft.com

- On the **Networking** tab, **Type of dial-up server I am calling** is set to **Point-to-Point Tunneling Protocol (PPTP)**. This is done to provide better performance when connecting. When **Type of dial-up server I am calling** is set to **Automatic**, an Internet Protocol security (IPSec) security association (SA) for an Layer Two Tunneling Protocol (L2TP) connection is attempted first. By configuring the connection for PPTP, the IPSec SA for an L2TP connection is not attempted.

#### L2TP/IPSec remote access client configuration

The remote access computer logs on to the Electronic, Inc. domain using a LAN connection to the Electronic, Inc. intranet and receives a certificate through auto-enrollment. Then, the New Connection Wizard is used to create a VPN connection with the following setting:

- Host name or IP address: vpn.electronic.microsoft.com

The VPN connection settings are modified as follows:

- On the **Networking** tab, **Type of dial-up server I am calling** is set to **Layer-2 Tunneling Protocol (L2TP)**. When **Type of dial-up server I am calling** is set to **Automatic**, an IPSec SA for an L2TP connection is attempted first. If the IPSec SA is not successful, then a PPTP connection is attempted. In this case, the network administrator for Electronic, Inc. does not want remote access clients that are capable of establishing an L2TP connection to fall back to the PPTP connection.

### Notes

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# On-Demand Branch Office

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**On-demand branch office**

The Portland and Dallas branch offices of Electronic, Inc. are connected to the corporate office by using on-demand router-to-router VPN connections. Both the Portland and Dallas offices contain a small number of employees who only need occasional connectivity with the corporate office. The routers running Windows Server 2003, Standard Edition, in the Portland and Dallas offices are equipped with an ISDN adapter that dials a local Internet service provider to gain access to the Internet, and then a router-to-router VPN connection is made across the Internet. When the VPN connection is not used for five minutes, the routers at the branch offices terminate the VPN connection.

The Dallas branch office uses the IP network ID of 192.168.28.0 with a subnet mask of 255.255.255.0. The Portland branch office uses the IP network ID of 192.168.4.0 with a subnet mask of 255.255.255.0.

To simplify the configuration, the VPN connection is a one-way initiated connection that is always initiated by the branch office router. For more information, see One-way initiated demand-dial connections.

The following illustration shows the Electronic, Inc. VPN server that provides on-demand branch office connections.




To deploy on-demand router-to-router VPN connections to connect the Portland and Dallas branch offices to the corporate office based on the settings configured in Common configuration for the VPN server, the following additional settings are configured.

**Domain configuration**

For the VPN connection to the Dallas office, the user account VPN_Dallas is created with the following settings:

- Password of nY7W{q8~=z3.
- For the dial-in properties on the VPN_Dallas account, the remote access permission is set to **Control access through Remote Access Policy**, and the static route is 192.168.28.0 with a subnet mask of 255.255.255.0 is added.
- For the account properties on the VPN_Dallas account, the **Password never expires** account option is enabled.
- The VPN_Dallas account is added to the VPN_Routers group.

For the VPN connection to the Portland office, the user account VPN_Portland is created with the following settings:

- Password of P*4s=wq!Gx1.
- For the dial-in properties on the VPN_Portland account, the remote access permission is set to **Control access through Remote Access Policy**, and the static route is 192.168.4.0 with a subnet mask of 255.255.255.0 is added.
- For the account properties on the VPN_Portland account, the **Password never expires** account option is enabled.
- The VPN_Portland account is added to the VPN_Routers group.

**Remote access policy configuration**

To define the authentication and encryption settings for the VPN routers, the following remote access policy is created:

- Policy name: VPN Routers
- Conditions:
  - **NAS-Port-Type** is set to Virtual (VPN)
  - **Windows-Groups** is set to VPN_Routers
  - **Called-Station-ID** is set to 207.209.68.1
- Permission is set to **Grant remote access permission**
- Profile settings:
  - **Authentication** tab: **Extensible Authentication Protocol** is enabled and **Smartcard or other certificate (TLS)** is configured to use the installed computer certificate (also known as the machine certificate). **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** is also enabled.
  - **Encryption** tab: **Strong** and **Strongest** are the only options that are selected.

**Note**

- The **Called-Station-ID** is set to the IP address of the Internet interface for the VPN server. Only tunnels initiated from the Internet are allowed. Tunnels initiated from the Electronic, Inc. intranet are not permitted. Electronic, Inc. users that require Internet access from the Electronic, Inc. intranet must go through the Electronic, Inc. proxy server (not shown), where Internet access is controlled and monitored.

For more information about the branch office router configuration, see:

- The Dallas office: PPTP-based on-demand branch office
- The Portland office: L2TP-based on-demand branch office

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# PPTP-based on-demand branch office

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPTP-based on-demand branch office**

The Dallas branch office is a PPTP-based branch office that uses a router running Windows Server 2003, Standard Edition, to create an on-demand, router-to-router VPN connection with the corporate office router in New York as needed. When the connection is made and is idle for five minutes, the connection is terminated.

To deploy a PPTP, one-way initiated, on-demand, router-to-router VPN connection to the corporate office based on the settings configured in Common configuration for the VPN server and On-Demand Branch Office, the following settings are configured on the Dallas router.

**Demand-dial interface for the connection to the ISP**

To connect the Dallas office router to the Internet by using a local ISP, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  ISP

- **Connection type**

  **Connect using a modem, ISDN adapter, or other physical device** is selected.

- **Select a device**

  The appropriate ISDN device is selected.

- **Phone number or address**

  Phone number of the ISP for the Dallas office.

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  Static route for Electronic, Inc. VPN server: To create the connection to the Dallas ISP when the router-to-router VPN connection needs to be made, the following static route is configured:

  - Interface: ISP
  - Destination: 207.209.68.1
  - Network mask: 255.255.255.255
  - Metric: 1

- **Dial-out credentials**

  - User name: Dallas office ISP account name.
  - Password: Dallas office ISP account password.

- Confirm password: Dallas office ISP account password.

For more information, see Add a demand-dial interface.

**Demand-dial interface for router-to-router VPN connection**
To connect the Dallas office router to the VPN server by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  CorpHQ

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Point to Point Tunneling Protocol (PPTP)** is selected.

- **Destination address**

  207.209.68.1

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  To make all locations on the corporate intranet reachable, the following static route is configured:

  - Interface: CorpHQ
  - Destination: 172.16.0.0
  - Network mask: 255.240.0.0
  - Metric: 1

  To make all locations on Electronic, Inc. branch offices reachable, the following static route is configured:

  - Interface: CorpHQ
  - Destination: 192.168.0.0
  - Network mask: 255.255.0.0
  - Metric: 1

- **Dial-out credentials**

  - User name: VPN_Dallas
  - Domain: electronic.microsoft.com
  - Password: nY7W{q8~=z3

- Confirm password: nY7W{q8~=z3

**PPTP packet filters on the demand-dial interface that connects to the ISP**
To ensure that only PPTP-based traffic is allowed on the connection to the Internet, PPTP packet filters are configured on the ISP demand-dial interface. For more information, see Add PPTP Filters.

### Note

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# L2TP-based on-demand branch office

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**L2TP-based on-demand branch office**

The Portland branch office is an L2TP/IPSec branch office that uses a router running Windows Server 2003, Standard Edition, to create an on-demand, router-to-router VPN connection with the corporate office router in New York as needed. When the connection is made and is idle for five minutes, the connection is terminated.

To deploy an L2TP, one-way initiated, on-demand, router-to-router VPN connection to the corporate office based on the settings configured in Common configuration for the VPN server and On-Demand Branch Office, the following settings are configured on the Portland router.

**Certificate configuration**

The Portland router was configured by the Electronic, Inc. network administrator while it was physically connected to the Electronic, Inc. intranet and then shipped to the Portland site. While the Portland router was connected to the Electronic, Inc. intranet, a computer certificate was installed through auto-enrollment.

**Demand-dial interface for the connection to the ISP**

To connect the Portland office router to the Internet by using a local ISP, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  ISP

- **Connection type**

  **Connect using a modem, ISDN adapter, or other physical device** is selected.

- **Select a device**

  The appropriate ISDN device is selected.

- **Phone number or address**

  Phone number of the ISP for the Portland office.

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  The following static route for the Electronic, Inc. VPN server is added automatically when the wizard is run to create the connection to the Portland ISP:

  - Interface: ISP
  - Destination: 207.209.68.1
  - Network mask: 255.255.255.255
  - Metric: 1

- **Dial-out credentials**

- User name: Portland office ISP account name.
- Password: Portland office ISP account password.
- Confirm password: Portland office ISP account password.

**Demand-dial interface for router-to-router VPN connection**

To connect the Portland office router to the VPN server by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  CorpHQ

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Layer-2 Tunneling Protocol (L2TP)** is selected.

- **Destination address**

  207.209.68.1

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  Static route for corporate headquarters and branch offices: To make all locations on the corporate intranet reachable, the following static route is configured:

  - Interface: CorpHQ
  - Destination: 172.16.0.0
  - Network mask: 255.240.0.0
  - Metric: 1

  To make all locations on Electronic, Inc. branch offices reachable, the following static route is configured:

  - Interface: CorpHQ
  - Destination: 192.168.0.0
  - Network mask: 255.255.0.0
  - Metric: 1

- **Dial-out credentials**

  - User name: VPN_Portland

- Domain: electronic.microsoft.com
- Password: P*4s=wq!Gx1
- Confirm password: P*4s=wq!Gx1

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Persistent Branch Office

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Persistent branch office**

The Chicago and Phoenix branch offices of Electronic, Inc. are connected to the corporate office by using persistent router-to-router VPN connections that stay connected 24 hours a day. The routers running Windows Server 2003, Standard Edition, in the Chicago and Phoenix offices are equipped with T1 WAN adapters that have a permanent connection to a local Internet service provider to gain access to the Internet.

The Chicago branch office uses the IP network ID of 192.168.9.0 with a subnet mask of 255.255.255.0. The Chicago branch office router uses the public IP address of 131.107.0.1 for its Internet interface. The Phoenix branch office uses the IP network ID of 192.168.14.0 with a subnet mask of 255.255.255.0. The Phoenix branch office router uses the public IP address of 131.107.128.1 for its Internet interface.

The VPN connection is a two-way initiated connection. The connection is initiated from either the branch office router or the corporate office router. Two-way initiated connections require the creation of demand-dial interfaces, remote access policies, IP address pools, and packet filters on the routers on both sides of the connection.

The following illustration shows the Electronic, Inc. VPN server that provides persistent branch office connections.




To deploy persistent router-to-router VPN connections to connect the Chicago and Phoenix branch offices to the corporate office based on the settings configured in Common configuration for the VPN server, the following additional settings are configured.

**Domain configuration**

For the Chicago office VPN connection that is initiated by the Chicago router, the user account VPN_Chicago is created with the following settings:

- Password of U9!j5dP(%q1.
- For the dial-in properties on the VPN_Chicago account, the remote access permission is set to **Control access through Remote Access Policy**.
- For the account properties on the VPN_Chicago account, the **Password never expires** account option is enabled.
- The VPN_Chicago account is added to the VPN_Routers group.

For the Phoenix office VPN connection that is initiated by the Phoenix router, the user account VPN_Phoenix is created with the following settings:

- Password of z2F%s)bW$4f.
- For the dial-in properties on the VPN_Phoenix account, the remote access permission is set to **Control access through Remote Access Policy**.

- For the account properties on the VPN_Phoenix account, the **Password never expires** account option is enabled.
- The VPN_Phoenix account is added to the VPN_Routers group.

For the Chicago office VPN connection and the Phoenix office VPN connection that are initiated by the corporate headquarters router, the user account VPN_CorpHQ is created with the following settings:

- Password of o3\Dn6@`-J4.
- For the dial-in properties on the VPN_CorpHQ account, the remote access permission is set to **Control access through Remote Access Policy**.
- For the account properties on the VPN_CorpHQ account, the **Password never expires** account option is enabled.
- The VPN_CorpHQ account is added to the VPN_Routers group.

**Remote access policy configuration**
Remote access policies must be configured at the VPN server, the Chicago router, and the Phoenix router.

**Remote access policy configuration at the VPN server**
The remote access policy configuration for the VPN server is the same as described in On-Demand Branch Office.

**Remote access policy configuration at the Chicago router**
To define the authentication and encryption settings for the VPN connections, the default policies are deleted, and the following remote access policy is created:

- Policy name: VPN Routers
- Conditions:
    - **NAS-Port-Type** is set to Virtual (VPN)
    - **Windows-Groups** is set to VPN_Routers
    - **Called-Station-ID** is set to 131.107.0.1
- Permission is set to **Grant remote access permission**
- Profile settings:
    - **Authentication** tab: **Extensible Authentication Protocol** is enabled and **Smartcard or other certificate (TLS)** is configured to use the installed computer certificate (also known as the machine certificate). **Microsoft Encrypted Authentication version 2 (MS-CHAP v2) is also enabled.**
    - **Encryption** tab: **Strong** and **Strongest** are the only options that are selected.

**Note**

- The **Called-Station-ID** is set to the IP address of the Internet interface for the branch office router. Only tunnels initiated from the Internet are allowed. Tunnels initiated from the Electronic, Inc. branch office network are not permitted.

**Remote access policy configuration at the Phoenix router**
To define the authentication and encryption settings for the VPN connections, the default policies are deleted, and the following remote access policy is created:

- Policy name: VPN Routers

- Conditions:
    - **NAS-Port-Type** is set to Virtual (VPN)
    - **Windows-Groups** is set to VPN_Routers
    - **Called-Station-ID** is set to 131.107.128.1
- Permission is set to **Grant remote access permission**
- Profile settings:
    - **Authentication** tab: **Extensible Authentication Protocol** is enabled and **Smartcard or other certificate (TLS)** is configured to use the installed computer certificate (also known as the machine certificate). **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** is also enabled.
    - **Encryption** tab: **Strong** and **Strongest** are the only options that are selected.

### Note

- The **Called-Station-ID** is set to the IP address of the Internet interface for the branch office router. Only tunnels initiated from the Internet are allowed. Tunnels initiated from the Electronic, Inc. branch office network are not permitted.

#### IP address pool configuration

IP address pools must be configured at the VPN server, the Chicago router, and the Phoenix router.

#### IP address pool configuration at the VPN server

The IP address pool configuration for the VPN server is the same as described in Common configuration for the VPN server.

#### IP address pool configuration at the Chicago router

A static IP address pool with a starting IP address of 192.168.9.248 and an ending IP address of 192.168.9.253 is configured. This creates a static address pool for up to five VPN clients.

For more information, see Create a static IP address pool.

#### IP address pool configuration at the Phoenix router

A static IP address pool with a starting IP address of 192.168.14.248 and an ending IP address of 192.168.14.253 is configured. This creates a static address pool for up to five VPN clients.

For more information, see Create a static IP address pool.

For more information about the corporate router and branch office router configuration, see:

- The Chicago office: PPTP-based persistent branch office
- The Phoenix office: L2TP-based persistent branch office

### Note

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# PPTP-based persistent branch office

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPTP-based persistent branch office**
The Chicago branch office is a PPTP-based branch office that uses a router running Windows Server 2003, Standard Edition, to create a persistent, router-to-router VPN connection with the corporate office router in New York. The connection is never terminated even when idle.

To deploy a PPTP, two-way initiated, persistent, router-to-router VPN connection to the corporate office based on the settings configured in Common configuration for the VPN server and Persistent Branch Office, the following settings are configured on the corporate router and Chicago router.

**VPN server configuration**
The VPN server is configured with a demand-dial interface and a static route.

**Demand-dial interface for router-to-router VPN connection**
To connect the corporate office router to the Chicago router by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  VPN_Chicago

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Point to Point Tunneling Protocol (PPTP)** is selected.

- **Destination address**

  131.107.0.1

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  Static route for Chicago office network: To make all locations on the Chicago network reachable, the following static route is configured:

  - Interface: VPN_Chicago
  - Destination: 192.168.9.0
  - Network mask: 255.255.255.0
  - Metric: 1

- **Dial-out credentials**

  - User name: VPN_CorpHQ

- Domain: electronic.microsoft.com
- Password: o3\Dn6@`-J4
- Confirm password: o3\Dn6@`-J4

Once the demand-dial interface is created, the following change is made:

- For the properties of the demand-dial interface, on the **Options** tab, under **Connection type**, the **Persistent connection** option is selected. To obtain properties of a demand-dial interface, click **Routing Interfaces**, right-click the desired demand-dial interface, and then click **Properties**.

**Chicago router configuration**
The Chicago router is configured with a demand-dial interface and a static route.

**Demand-dial interface for router-to-router VPN connection**
To connect the Chicago office router to the corporate office router by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  VPN_CorpHQ

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Point to Point Tunneling Protocol (PPTP)** is selected.

- **Destination address**

  207.209.68.1

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  Static route for the Electronic, Inc. VPN server: To make the Electronic, Inc. VPN server on the Internet reachable, the following static route is configured:

  - Interface: The WAN adapter attached to the Internet
  - Destination: 207.209.68.1
  - Network mask: 255.255.255.255
  - Gateway: 0.0.0.0
  - Metric: 1

  Static routes for corporate intranet and branch offices: To make all locations on the corporate intranet reachable, the following static route is configured:

    - Interface: VPN_CorpHQ
    - Destination: 172.16.0.0
    - Network mask: 255.240.0.0
    - Metric: 1

  To make all locations on Electronic, Inc. branch offices reachable, the following static route is configured:

    - Interface: VPN_CorpHQ
    - Destination: 192.168.0.0
    - Network mask: 255.255.0.0
    - Metric: 1

- **Dial-out credentials**
    - User name: VPN_Chicago
    - Domain: electronic.microsoft.com
    - Password: U9!j5dP(%q1
    - Confirm password: U9!j5dP(%q1

Once the demand-dial interface is created, the following change is made:

- For the properties of the demand-dial interface, on the **Options** tab, under **Connection type**, the **Persistent connection** option is selected.

**Notes**

- Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway. The gateway address of 0.0.0.0 is an example. 0.0.0.0 is the unspecified IP address.

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# L2TP-based persistent branch office

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**L2TP-based persistent branch office**

The Phoenix branch office is an L2TP/IPSec branch office that uses a router running Windows Server 2003, Standard Edition, to create a persistent, router-to-router VPN connection with the corporate office router in New York. The connection is never terminated even when idle.

To deploy an L2TP, two-way initiated, persistent, router-to-router VPN connection to the corporate office based on the settings configured in Common configuration for the VPN server and Persistent Branch Office, the following settings are configured on the corporate router and Phoenix router.

**VPN server configuration**

The VPN server is configured with a demand-dial interface and a static route.

**Demand-dial interface for router-to-router VPN connection**

To connect the corporate office router to the Phoenix router by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  VPN_Phoenix

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Layer-2 Tunneling Protocol (L2TP)** is selected.

- **Destination address**

  131.107.128.1

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  Static route for Phoenix office network: To make all locations on the Phoenix network reachable, the following static route is configured:

  - Interface: VPN_Phoenix
  - Destination: 192.168.14.0
  - Network mask: 255.255.255.0
  - Metric: 1

- **Dial-out credentials**

  - User name: VPN_CorpHQ

- Domain: electronic.microsoft.com
- Password: o3\Dn6@`-J4
- Confirm password: o3\Dn6@`-J4

Once the demand-dial interface is created, the following change is made:

- For the properties of the demand-dial interface, on the **Options** tab, under **Connection type**, the **Persistent connection** option is selected.

**Phoenix router configuration**

The Phoenix router was configured by the Electronic, Inc. network administrator while connected to the Electronic, Inc. intranet and then shipped to the Phoenix site. While the Phoenix router was connected to the Electronic, Inc. intranet, a computer certificate was installed through auto-enrollment. Additionally, the Phoenix router computer was configured with a demand-dial interface and a static route.

**Demand-dial interface for router-to-router VPN connection**

To connect the Phoenix office router to the corporate office router by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  VPN_CorpHQ

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Layer-2 Tunneling Protocol (L2TP)** is selected.

- **Destination address**

  207.209.68.1

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  Static route for the Electronic, Inc. VPN server: To make the Electronic, Inc. VPN server on the Internet reachable, the following static route is configured:

  - Interface: The WAN adapter attached to the Internet
  - Destination: 207.209.68.1
  - Network mask: 255.255.255.255
  - Gateway: 0.0.0.0
  - Metric: 1

  **Note**

- Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway. The gateway address of 0.0.0.0 is an example. 0.0.0.0 is the unspecified IP address.

Static route for corporate intranet and branch offices: To make all locations on the corporate intranet reachable, the following static route is configured:

- Interface: VPN_CorpHQ
- Destination: 172.16.0.0
- Network mask: 255.240.0.0
- Metric: 1

To make all locations on Electronic, Inc. branch offices reachable, the following static route is configured:

- Interface: VPN_CorpHQ
- Destination: 192.168.0.0
- Network mask: 255.255.0.0
- Metric: 1

- **Dial-out credentials**
  - User name: VPN_Phoenix
  - Domain: electronic.microsoft.com
  - Password: z2F%s)bW$4f
  - Confirm password: z2F%s)bW$4f

Once the demand-dial interface is created, the following change is made:

- For the properties of the demand-dial interface, on the **Options** tab, under **Connection type**, the **Persistent connection** option is selected.

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Extranet for Business Partners

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Extranet for business partners**

The network administrator for Electronic, Inc. has created an extranet, a portion of the Electronic, Inc. private network that is available to business partners through secured VPN connections. The Electronic, Inc. extranet is the network attached to the Electronic, Inc. VPN server and contains a file server and a Web server. Parts distributors Tasmanian Traders and Parnell Aerospace are Electronic, Inc. business partners and connect to the Electronic, Inc. extranet by using on-demand, router-to-router VPN connections. An additional access policy is used to ensure that the business partners can only access the extranet file server and Web server.

The file server on the Electronic, Inc. extranet is configured with an IP address of 172.31.0.10, and the Web server is configured with an IP address of 172.31.0.11. Tasmanian Traders uses the public network ID of 131.107.254.0 with a subnet mask of 255.255.255.0. Parnell Aerospace uses the public network ID of 131.107.250.0 with a subnet mask of 255.255.255.0. To ensure that the extranet Web server and file server can reach the business partners, static routes are configured on the file server and Web server for each of the business partner networks that use the gateway address of 172.31.0.1

To simplify configuration, the VPN connection is a one-way initiated connection. The connection is always initiated by the business partner router. For more information, see One-way initiated demand-dial connections.

The following illustration shows the Electronic, Inc. VPN server that provides extranet connections for business partners.




To deploy business partner, on-demand, one-way initiated, router-to-router VPN connections to connect Tasmanian Traders and Parnell Aerospace to the Electronic, Inc. extranet based on the settings configured in Common configuration for the VPN server, the following additional settings are configured.

**Domain configuration**

For the VPN connection to Tasmanian Traders, the user account PTR_Tasmanian is created with the following settings:

- Password of Y8#-vR7?]fl.
- For the dial-in properties on the PTR_Tasmanian account, the remote access permission is set to **Control access through Remote Access Policy**.
- For the account properties on the PTR_Tasmanian account, the **Password never expires** account option is enabled.
- The PTR_Tasmanian account is added to the VPN_Partners group.

For the VPN connection to Parnell Aerospace, the user account PTR_Parnell is created with the following settings:

- Password of W@8c^4r-;2\.
- For the dial-in properties on the PTR_Parnell account, the remote access permission is set to **Control access through Remote Access Policy**.

- For the account properties on the PTR_Parnell account, the **Password never expires** account option is enabled.
- The PTR_Parnell account is added to the VPN_Partners group.

**Routing configuration**

In order for traffic to be sent back to the Tasmanian Traders and Parnell Aerospace business partners, additional routes are added to the routing tables on the file server and Web server computers. On both computers, the following commands were issued at the Windows Server 2003 command prompt.

- route -p add 131.107.254.0 mask 255.255.255.0 172.31.0.1
- route -p add 131.107.250.0 mask 255.255.255.0 172.31.0.1

**Remote access policy configuration**

To define the authentication and encryption settings for business partner VPN connections, the following remote access policy is created:

- Policy name: VPN Partners
- Conditions:
    - **NAS-Port-Type** is set to Virtual (VPN)
    - **Windows-Groups** is set to VPN_Partners
    - **Called-Station-ID** is set to 207.209.68.1
- Permission is set to **Grant remote access permission**
- Profile settings:
    - **Authentication** tab: **Extensible Authentication Protocol** is enabled and **Smartcard or other certificate (TLS)** is configured to use the installed computer certificate (also known as the machine certificate). **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** is also enabled.
    - **Encryption** tab: **Strong** and **Strongest** are the only options that are selected.

**Demand-dial interface configuration**

To connect the Tasmanian Traders router to the Electronic, Inc. VPN server by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  PTR_Tasmanian

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Point to Point Tunneling Protocol (PPTP)** is selected.

- **Destination address**

  Leave blank.

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Dial-out credentials**

  - User name: none
  - Domain: none
  - Password: none
  - Confirm password: none

**Note**

- The destination address is left blank, and the user name is set to "none" because this demand-dial interface is never used to initiate the connection. The business partner always initiates the connection to the Electronic, Inc. VPN server.

To connect the Parnell Aerospace router to the Electronic, Inc. VPN server by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  PTR_Parnell

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Point to Point Tunneling Protocol (PPTP)** is selected.

- **Destination address**

  Leave blank.

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Dial-out credentials**

  - User name: none
  - Domain: none Password: none
  - Confirm password: none

**Note**

- The destination address is left blank and the user name is set to "none" because this demand-dial interface is never used to initiate the connection. The business partner always initiates the connection to the Electronic, Inc. VPN server.

**Static route for Tasmanian Traders**
To make all locations at Tasmanian Traders reachable, the following static route is configured:

- Interface: PTR_Tasmanian
- Destination: 131.107.254.0
- Network mask: 255.255.255.0
- Metric: 1

**Static route for Parnell Aerospace**
To make all locations at Parnell Aerospace reachable, the following static route is configured:

- Interface: PTR_Parnell
- Destination: 131.107.250.0
- Network mask: 255.255.255.0
- Metric: 1

**IP packet filter configuration**
To confine the traffic from the Tasmanian Traders business partner to the file server and Web server, the following IP packet filters are configured on the PTR_Tasmanian demand-dial interface:

- Input filters
    - Filter action: **Deny all traffic except those listed below**
    - Filter 1: Destination network IP address of 172.31.0.10 and subnet mask of 255.255.255.255
    - Filter 2: Destination network IP address of 172.31.0.11 and subnet mask of 255.255.255.255
- Output filters
    - Filter action: **Deny all traffic except those listed below**
    - Filter 1: Source network IP address of 172.31.0.10 and subnet mask of 255.255.255.255
    - Filter 2: Source network IP address of 172.31.0.11 and subnet mask of 255.255.255.255

To confine the traffic from the Parnell Aerospace business partner to the file server and Web server, the following IP packet filters are configured on the PTR_Parnell demand-dial interface:

- Input filters
    - Filter action: **Deny all traffic except those listed below**
    - Filter 1: Destination network IP address of 172.31.0.10 and subnet mask of 255.255.255.255
    - Filter 2: Destination network IP address of 172.31.0.11 and subnet mask of 255.255.255.255
- Output filters
    - Filter action: **Deny all traffic except those listed below**

- Filter 1: Source network IP address of 172.31.0.10 and subnet mask of 255.255.255.255
- Filter 2: Source network IP address of 172.31.0.11 and subnet mask of 255.255.255.255

**Note**

- The **Called-Station-ID** is set to the IP address of the Internet interface for the VPN server. Only tunnels initiated from the Internet are allowed. Tunnels initiated from the Electronic, Inc. intranet are not permitted. Electronic, Inc. users that require Internet access from the Electronic, Inc. intranet must go through the Electronic, Inc. proxy server (not shown), where Internet access is controlled and monitored.

For more information about the business partner router configuration, see:

- Tasmanian Traders: PPTP-based extranet for business partners
- Parnell Aerospace: L2TP-based extranet for business partners

**Note**

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# PPTP-based extranet for business partners

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPTP-based extranet for business partners**

Tasmanian Traders is a business partner that uses a router running a Windows Server 2003 operating system to create an on-demand, PPTP-based, router-to-router VPN connection with the Electronic, Inc. corporate office router in New York as needed. When the connection is created and is idle for five minutes, the connection is terminated. The Tasmanian Traders router is connected to the Internet by using a permanent WAN connection.

To deploy a PPTP, one-way initiated, on-demand, router-to-router VPN connection to the corporate office based on the settings configured in Common configuration for the VPN server and Extranet for Business Partners, the following settings are configured on the Tasmanian Traders router:

**Demand-dial interface for router-to-router VPN connection**

To connect the Tasmanian Traders router to the Electronic, Inc. VPN server by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  Electronic

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Point to Point Tunneling Protocol (PPTP)** is selected.

- **Destination address**

  207.209.68.1

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  Static route for Electronic, Inc. extranet: To make all locations on the Electronic, Inc. extranet reachable, the following static route is configured:

  - Interface: Electronic
  - Destination: 172.31.0.0
  - Network mask: 255.255.0.0
  - Metric: 1

- **Dial-out credentials**

  - User name: PTR_Tasmanian
  - Domain: electronic.microsoft.com

- Password: Y8#-vR7?]fI
- Confirm password: Y8#-vR7?]fI

**Notes**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# L2TP-based extranet for business partners

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**L2TP-based extranet for business partners**

Parnell Aerospace is a business partner that uses a router running Windows Server 2003, Standard Edition, to create an on-demand, L2TP/IPSec, router-to-router VPN connection with the Electronic, Inc. corporate office router in New York as needed. When the connection is created and is idle for five minutes, the connection is terminated. The Parnell Aerospace router is connected to the Internet by using a permanent WAN connection.

To deploy an L2TP, one-way initiated, on-demand, router-to-router VPN connection to the corporate office based on the settings configured in Common configuration for the VPN server and Extranet for Business Partners, the following settings are configured on the Parnell Aerospace router.

**Certificate configuration**

The Parnell Aerospace router was configured by the Electronic, Inc. network administrator while physically connected to the Electronic, Inc. intranet and then shipped to the network administrator at Parnell Aerospace. While the Parnell Aerospace router was connected to the Electronic, Inc. intranet, a computer certificate was installed through auto-enrollment.

**Demand-dial interface for router-to-router VPN connection**

To connect the Parnell Aerospace router to the Electronic, Inc. VPN server by using a router-to-router VPN connection over the Internet, a demand-dial interface is created by using the Demand-Dial Interface Wizard with the following settings:

- **Interface name**

  Electronic

- **Connection type**

  **Connect using virtual private networking (VPN)** is selected.

- **VPN type**

  **Layer-2 Tunneling Protocol (L2TP)** is selected.

- **Destination address**

  207.209.68.1 (the IP address of the Electronic, Inc. VPN server's interface on the Internet)

- **Protocols and security**

  The **Route IP packets on this interface** check box is selected.

- **Static Routes for Remote Networks**

  Static route for Electronic, Inc. extranet: To make all locations on the Electronic, Inc. extranet reachable, the following static route is configured:

  - Interface: Electronic
  - Destination: 172.31.0.0
  - Network mask: 255.255.0.0
  - Metric: 1

- **Dial-out credentials**

- User name: PTR_Parnell
- Domain: electronic.microsoft.com
- Password: W@8c^4r-;2\
- Confirm password: W@8c^4r-;2\

**Notes**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Dial-up and VPNs with RADIUS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up and VPNs with RADIUS**
In addition to VPN-based remote access, the network administrator for Electronic, Inc. wants to provide modem-based dial-up remote access for employees of the New York office. All employees of the New York office belong to a Windows Server 2003 operating system group called NY_Employees. A separate remote access server running Windows Server 2003, Standard Edition, provides dial-up remote access at the phone number 555-0111. Rather than administer the remote access policies of both the VPN server and the remote access server separately, the network administrator is using a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; with Internet Authentication Service (IAS) as a RADIUS server. The IAS server has an IP address of 172.31.0.9 on the Electronic, Inc. extranet and provides centralized remote access authentication, authorization, and accounting for both the remote access server and the VPN server.

The following illustration shows the Electronic, Inc. RADIUS server that provides authentication and accounting for the VPN server and the remote access server.




**Domain configuration**
For each New York office employee that is allowed dial-up access, the remote access permission for the dial-in properties of the user account is set to **Control access through Remote Access Policy**.

**Remote access policy configuration**
Remote access policies must be modified in two ways:

1. The existing remote access policies that are configured on the VPN server running Windows Server 2003 must be copied to the IAS server.

2. A new remote access policy is added for dial-up remote access clients on the IAS server.

**Copying the remote access policies**
Once the VPN server running Windows Server 2003 is configured to use RADIUS authentication, the remote access policies stored on the VPN server are no longer used. Instead, the remote access policies stored on the IAS server running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. Therefore, the current set of remote access policies is copied to the IAS server.

For more information, see Copy the IAS configuration to another server.

**Creating a new remote access policy for dial-up remote access clients**
To define the authentication and encryption settings for dial-up connections by employees of the New York office, the following remote access policy is created on the RADIUS server computer:

- Policy name: Dial-Up for New York Employees

- Conditions:

    - **NAS-Port-Type** is set to all types *except* Virtual (VPN)

    - **Windows-Groups** is set to NY_Employees

- Permission is set to **Grant remote access permission**

- Profile settings:

    - **Authentication** tab: **Extensible Authentication Protocol** is enabled and **Smartcard or other certificate (TLS)** is configured to use the installed computer certificate (also known as the machine certificate). **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** and **Microsoft Encrypted Authentication (MS-CHAP)** are also enabled.

    - **Encryption** tab: All options are selected.

**RADIUS configuration**

To configure RADIUS authentication and accounting, the network administrator for Electronic, Inc. configures the following:

- The RADIUS server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; and with IAS installed. IAS is configured for two RADIUS clients; the remote access server and the VPN server. For more information, see Internet Authentication Service and Add RADIUS clients.

- The remote access server running a member of the Windows Server 2003 family is configured to use RADIUS authentication and accounting at the IP address of 172.31.0.9 and a shared secret. For more information, see Use RADIUS authentication and Use RADIUS accounting.

- The VPN server running Windows Server 2003 is configured to use RADIUS authentication and accounting at the IP address of 172.31.0.9 and to use a shared secret. For more information, see Use RADIUS authentication and Use RADIUS accounting.

**Dial-up remote access client configuration**

The New Connection Wizard is used to create a dial-up connection with the following setting:

- Phone number: 555-0111

### Note

- The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

# Virtual Private Network Test Lab

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private network test lab**
This section provides detailed information about how you can use five computers to create a test lab with which to configure and test the virtual private network (VPN) features in the Windows Server 2003 family. These instructions are designed to take you through a set of tasks that expose you to VPN connections and their functionality. Beyond the set of tasks, you can use these instructions to create a functioning VPN configuration. You can then use this configuration to experiment with VPN features and functionality in the Windows Server 2003 family prior to deployment on your production network.

This section covers:

- Setting up the VPN test network infrastructure
- Virtual private network test lab tasks

**Note**

- The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

For more information about deploying VPN connections, see Deploying Virtual Private Networks. For configuration checklists, see VPN Checklists. For an example implementation of VPN connections, see Virtual Private Network Implementation Examples.

# Setting up the VPN test network infrastructure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up the infrastructure**

The infrastructure for the VPN test lab network consists of five computers performing the following services:

- A computer running Windows Server 2003, Standard Edition, named DC1 that is acting as a domain controller, a Domain Name System (DNS) server, and a certification authority (CA).
- A computer running Windows Server 2003, Standard Edition, named IAS1 that is acting as a Remote Authentication Dial-in User Service (RADIUS) server.
- A computer running Windows Server 2003, Standard Edition, named IIS1 that is acting as a Web server and file sharing server.
- A computer running Windows Server 2003, Standard Edition, named VPN1 that is acting as a VPN server. VPN1 has two network adapters installed.
- A computer running Windows XP Professional named CLIENT1 that is acting as a VPN client.

The following illustration shows the configuration of the VPN test lab.




There is a network segment representing a corporate intranet and a network segment representing the Internet. All computers on the corporate intranet are connected to a common hub or Layer 2 switch. All computers on the Internet are connected to a separate common hub or Layer 2 switch. Private addresses are used throughout the test lab configuration. The private network of 172.16.0.0/24 is used for the intranet. The private network of 10.0.0.0/24 is used for the simulated Internet.

Each computer is manually configured with the appropriate IP address, subnet mask, and DNS server IP address. There are no Dynamic Host Configuration protocol (DHCP) or Windows Internet Name Service (WINS) servers present.

The following sections describe the configuration for each of the computers in the test lab. To reconstruct this test lab, configure the computers in the order presented.

### Note

- The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

**DC1**

DC1 is a computer running Windows Server 2003, Standard Edition, that is providing the following services:

- A domain controller for the testlab.microsoft.com domain.
- A DNS server for the testlab.microsoft.com DNS domain.
- The enterprise root certification authority (CA) for the testlab.microsoft.com domain.

To configure DC1 for these services, perform the following steps.

1. Install Windows Server 2003, Standard Edition, as a stand-alone server.
2. Configure the TCP/IP protocol with the IP address of 172.16.0.1 and the subnet mask of 255.255.255.0. For more information, see Configure TCP/IP for static addressing.
3. Run Active Directory Installation Wizard for a new domain called testlab.microsoft.com in a new forest. Install the DNS service when prompted.
4. Install the Certificate Services component as an enterprise root CA. For more information, see Install an enterprise root certification authority.
5. Configure the testlab.microsoft.com domain for automatic enrollment of computer certificates. For more information, see Configure automatic certificate allocation from an enterprise CA.

**IAS1**

IAS1 is a computer running Windows Server 2003, Standard Edition, that is providing RADIUS authentication, authorization, and accounting for VPN1 (the VPN server computer).

To configure IAS1 as a RADIUS server, perform the following steps:

1. On DC1, add a computer account for the IAS1 computer. For more information, see Create a new computer account.
2. Install Windows Server 2003, Standard Edition, as a stand-alone server.
3. Configure the TCP/IP protocol with the IP address of 172.16.0.2, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.
4. Join IAS1 to the testlab.microsoft.com domain. For more information, see Join a domain.
5. Install Internet Authentication Service (IAS) and register it in the testlab.microsoft.com domain. For more information, see Install IAS and Enable the IAS server to read user accounts in Active Directory.

**IIS1**

IIS1 is a computer running Windows Server 2003, Standard Edition, and Internet Information Services (IIS). It is providing Web server services for intranet clients. To configure IIS1 as a Web server, perform the following steps:

1. On DC1, add a computer account for the IIS1 computer. For more information, see Create a new computer account.
2. Install Windows Server 2003, Standard Edition, as a stand-alone server, then install IIS.

   For more information, see Installing IIS.
3. Configure the TCP/IP protocol with the IP address of 172.16.0.3, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

4. Join IIS1 to the testlab.microsoft.com domain. For more information, see Join a domain.

5. Install IIS.

6. To determine whether the Web server is working correctly, run your Web browser on IAS1. When prompted by the Internet Connection Wizard, configure the wizard for a LAN connection. In your Web browser, in **Address**, type **http://IIS1.testlab.microsoft.com/wnetStndS_v_s_rgb.gif**. You should see a Windows Server 2003, Standard Edition, graphic.

7. On IIS1, in Windows Explorer, share the root directory of **Local Disk (C:)** using the share name **ROOT** to the group **Everyone** with full access. For more information, see Share a folder or drive.

**VPN1**

VPN1 is a computer running Windows Server 2003, Standard Edition, that is providing VPN server services for Internet-based VPN clients. To configure VPN1 as a VPN server, perform the following steps:

1. On DC1, add a computer account for VPN1. For more information, see Create a new computer account.

2. Install Windows Server 2003, Standard Edition, as a stand-alone server.

3. For the intranet local area connection, configure the TCP/IP protocol with the IP address of 172.16.0.4, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

4. For the Internet local area connection, configure the TCP/IP protocol with the IP address of 10.0.0.2 and the subnet mask of 255.255.255.0. For more information, see Configure TCP/IP for static addressing.

5. Join VPN1 to the testlab.microsoft.com domain. For more information, see Join a domain.

6. Configure and enable the Routing and Remote Access service. For more information, see Enable the Routing and Remote Access service. In the Routing and Remote Access server Setup Wizard, select **Virtual private network (VPN) server** from the list of common configurations. When prompted for IP address assignment, select **From a specified range of addresses** and configure the range 172.16.0.248 to 172.16.0.255. Do not configure RADIUS authentication.

**CLIENT1**

CLIENT1 is a computer running Windows XP Professional that is acting as a VPN client and gaining remote access to intranet resources across the simulated Internet. To configure CLIENT1 as a VPN client, perform the following steps:

1. On DC1, add a computer account for CLIENT1. For more information, see Create a new computer account.

2. Connect CLIENT1 to the intranet network segment.

3. On CLIENT1, install Windows XP Professional as a workgroup computer.

4. Configure the TCP/IP protocol with the IP address of 172.16.0.5, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

5. Join CLIENT1 to the testlab.microsoft.com domain. For more information, see Join a domain.

6. Configure the TCP/IP protocol with the IP address of 10.0.0.1, the subnet mask of 255.255.255.0, and no DNS server IP address. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

7. Shut down the CLIENT1 computer.

8. Disconnect the CLIENT1 computer from the intranet network segment, and connect it to the simulated Internet network

segment.

9. Restart the CLIENT computer and log on using the cached credentials of the testlab.microsoft.com administrator account.

# Virtual private network test lab tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual private network test lab tasks**
The following tasks are designed to take you through the most common elements of setting up VPN support with the Windows Server 2003 family:

- PPTP-based remote access
- L2TP/IPSec remote access
- RADIUS authentication and accounting
- Remote access policies for both PPTP and L2TP connections

### Note

- The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

**PPTP-based remote access**
To create a PPTP-based remote access VPN connection between CLIENT1 and VPN1 and test whether intranet resources are available, perform the following steps:

**1. Create a user account.**

1. On DC1, use Active Directory Users and Computers to create a user account called **PPTPUser** and set the password. For more information, see Create a new computer account.
2. Set the remote access permission on the **Dial-in** tab to **Allow access**. For more information, see Configure remote access permission for a user.

**2. Create the PPTP connection.**

1. On CLIENT1, use the New Connection Wizard to create a new PPTP-based connection named **PPTPtoCorpnet**, using the IP address of 10.0.0.2. For more information, see Make a virtual private network (VPN) connection.
2. Obtain properties on the newly created **PPTPtoCorpnet** connection, and click the **Networking** tab.
3. Under **Type of VPN server I am calling**, select **PPTP VPN**, and click **OK**.

**3. Make the PPTP connection.**

1. On CLIENT1, double-click the **PPTPtoCorpnet** connection.
2. When prompted for credentials, type **PPTPUser** as the user name, **testlab.microsoft.com** as the domain name, and the password.

**4. Access Web server and file share on the intranet.**

1. On CLIENT1, run your Web browser.

2. When prompted by the Internet Connection Wizard, configure it for a LAN connection. In your Web browser, type **http://IIS1.testlab.microsoft.com/wnetStndS_v_s_rgb.gif**.

   You should see a Windows Server 2003, Standard Edition, graphic.

3. On CLIENT1, click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**.

   You should see the contents of the Local Drive (C:) on IIS1.

**5. Disconnect the PPTP connection.**

- On CLIENT1, right-click the **PPTPtoCorpnet** connection, and then click **Disconnect**.

**L2TP/IPSec remote access**

To create an L2TP/IPSec remote access VPN connection between CLIENT1 and VPN1 and test whether intranet resources are available, perform the following steps:

**1. Create a user account.**

1. On DC1, use Active Directory Users and Computers to create a user account called **L2TPUser**, and set the password. For more information, see Create a new computer account.

2. Set the remote access permission on the **Dial-in** tab to **Allow access**. For more information, see Configure remote access permission for a user.

**2. Create the L2TP connection.**

1. On CLIENT1, use the New Connection Wizard to create a new L2TP/IPSec connection named **L2TPtoCorpnet**, using the IP address of 10.0.0.2. For more information, see Make a virtual private network (VPN) connection.

2. Obtain properties on the newly created **L2TPtoCorpnet** connection, and click the **Networking** tab.

3. Under **Type of VPN server I am calling**, select **L2TP IPSec VPN**, and click **OK**.

**3. Make the L2TP connection.**

1. On CLIENT1, double-click the **L2TPtoCorpnet** connection.

2. When prompted for credentials, type **L2TPUser** as the user name, **testlab.microsoft.com** as the domain name, and the password.

**4. Access Web server and file share on the intranet.**

1. On CLIENT1, run your Web browser.

2. When prompted by the Internet Connection Wizard, configure it for a LAN connection. In your Web browser, type **http://IIS1.testlab.microsoft.com/wnetStndS_v_s_rgb.gif**.

   You should see a Windows Server 2003, Standard Edition, graphic.

3. On CLIENT1, click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**.

   You should see the contents of the Local Drive (C:) on IIS1.

**5. Disconnect the L2TP connection.**

- On CLIENT1, right-click the **L2TPtoCorpnet** connection, and then click **Disconnect**.

**RADIUS authentication and accounting**

To configure RADIUS authentication and accounting for VPN connections, perform the following steps:

**1. Configure IAS1 for VPN1 as a RADIUS client.**

- On IAS1, add VPN1 as a RADIUS client with the Microsoft Client-Vendor, at the IP address of 172.16.0.4, and set the shared secret. For more information, see Add RADIUS clients.

**2. Configure IAS1 to log authentication events.**

- On IAS1, enable the logging of accounting and authentication requests in the Internet Authentication Service console, in the **Remote Access Logging** folder, in the properties of the **Local File** object, on the **Settings** tab.

**3. Configure VPN1 for IAS1 as a RADIUS server.**

- On VPN1, add IAS1 as a RADIUS server for both the authentication and accounting provider, at the IP address of 172.16.0.2, and set the shared secret. For more information, see Use RADIUS authentication and Use RADIUS accounting.

**4. Make a PPTP connection.**

1. On CLIENT1, create a PPTP connection with VPN1 using the **PPTPtoCorpnet** connection.
2. Disconnect the connection. For more information, see Disconnect from a remote network.

**5. Make an L2TP connection.**

1. On CLIENT1, create an L2TP connection with VPN1 using the **L2TPtoCorpnet** connection.
2. Disconnect the connection. For more information, see Disconnect from a remote network.

**6. Check the system event log for RADIUS events.**

- On IAS1, use Event Viewer to view IAS events in the system event log for the PPTP and L2TP connections that were recently created using CLIENT1.

**7. Check RADIUS authentication and accounting logs.**

- On IAS1, use Windows Explorer to open the *systemroot*\System32\Logfiles\Iaslog.log file. Note the authentication and accounting entries for the PPTP and L2TP connection. For more information on the IAS log file format, see Interpreting IAS-formatted log files.

**Remote access policies for both PPTP and L2TP connections**

To create separate remote access policies for PPTP and L2TP connections, perform the following steps:

**1. On IAS1, create a new remote access policy with the following settings:**

- **Policy name**: PPTP connections
- **Policy Conditions**:
    - **NAS-Port-Type** matches **Virtual (VPN)**

- **Tunnel-Type** matches **Point-to-Point Tunneling Protocol (PPTP)**

- **If a connection request matches the specified conditions**: Grant remote access permission

- Profile settings, **IP** tab, output packet filter:

  Filter action: **Deny all traffic except those listed below**

  Destination network, IP address: 172.16.0.1

  Destination network, Subnet mask: 255.255.255.255

  Protocol: Any

- Profile settings, **IP** tab, input packet filter:

  Filter action: **Deny all traffic except those listed below**

  Source network, IP address: 172.16.0.1

  Destination network, Subnet mask: 255.255.255.255

  Protocol: Any

**2. On IAS1, create a new custom remote access policy with the following settings:**

- **Policy name**: L2TP connections

- **Policy Conditions**:

  **NAS-Port-Type** matches **Virtual (VPN)**

  **Tunnel-Type** matches **Layer Two Tunneling Protocol (L2TP)**

- **If a connection request matches the specified conditions**: Grant remote access permission

- Profile settings, **IP** tab, output packet filter:

  Filter action: **Deny all traffic except those listed below**

  Destination network, IP address: 172.16.0.2

  Destination network, Subnet mask: 255.255.255.255

  Protocol: Any

- Profile settings, **IP** tab, input packet filter:

  Filter action: **Deny all traffic except those listed below**

  Source network, IP address: 172.16.0.2

  Destination network, Subnet mask: 255.255.255.255

  Protocol: Any

**3. Make a PPTP connection and test connectivity.**

1. On CLIENT1, make a VPN connection with VPN1 using the **PPTPtoCorpnet** connection.

2. Use the **ping** command to ping DC1 at its IP address of 172.16.0.1.

3. Use the **ping** command to ping IAS1 at its IP address of 172.16.0.2. This command fails because packet filtering for all connections that match the **PPTP connections** policy allows only traffic sent to and from the IP address of 172.16.0.1.

4. Disconnect the **PPTPtoCorpnet** connection.

**4. Make an L2TP connection and test connectivity.**

1. On CLIENT1, make a VPN connection with VPN1 using the **L2TPtoCorpnet** connection.

2. Use the **ping** command to ping IAS1 at its IP address of 172.16.0.2.

3. Use the **ping** command to ping DC1 at its IP address of 172.16.0.1. This command fails because packet filtering for all connections that match the L2TP connections policy allows only traffic sent to and from the IP address of 172.16.0.2.

4. Disconnect the **L2TPtoCorpnet** connection.

**5. Check the system event log for IAS events.**

- On IAS1, use Event Viewer to view the IAS events in the system event log for the PPTP and L2TP connections that were recently created by CLIENT1. Note that the authentication event message text contains the name of the remote access policy that accepted the connection.

# VPN Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

The Microsoft Windows Server 2003 Deployment Kit contains advanced information about virtual private networks (VPNs).

- For more information about configuring, administering, and optimizing VPNs, see "Internetworking Guide" at the Microsoft Windows Resource Kits Web site.
- For more information about designing remote access solutions, see "Designing Network Services" at the Microsoft Windows Resource Kits Web site.
- For additional information, see VPN Technical White Papers at the Microsoft Web site.

For additional information about troubleshooting with the remote access server, see Remote Access Troubleshooting tools.

# Remote Access Troubleshooting tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting tools**
To help you gather information to troubleshoot problems with the remote access server, the following troubleshooting tools are available:

- Authentication and accounting logging. For more information see Logging.
- Event logging records errors, warnings, and events in the system event log.
- PPP logging creates a log file of the PPP control messages that are sent and received during a PPP connection.
- Tracing records the sequence of programming functions called during a process to a file. For more information, see Using tracing.
- Network Monitor captures traffic sent between a remote access server and dial-up client during the PPP connection process and during data transfer.
- Netsh commands for remote access configure and display remote access settings for a server running Routing and Remote Access. For more information, see Netsh commands for remote access (ras).

**Remote Access Diagnostics**
You can use the Remote Access diagnostic functions available in the Windows Server 2003 family to collect detailed logs and information about a remote access connection. Diagnostics can be set from the **Diagnostics** tab of the properties of the connection in the Network Connection folder, or you can run specific diagnostics by typing the **netsh ras diagnostics** command at the command line.

**Audit logging and Oakley logging for L2TP/IPSec connections**
You can use the Audit Logging feature to monitor IPSec-related events. This is the fastest and simplest method to troubleshoot unsuccessful L2TP/IPSec connections. You can also enable the Oakley log to record all ISAKMP main mode or quick mode negotiations. For more information, see IPSec troubleshooting tools.

**Event logging**
You can use event logging to record remote access server errors, warnings, and other detailed information in the system event log. You can enable event logging on the **Event Logging** tab on the properties of a remote access server. For more information, see View properties of the remote access server.

**PPP logging**
PPP logging records the series of programming functions and PPP control messages during a PPP connection and is a valuable source of troubleshooting information when you are troubleshooting the failure of a PPP connection. To enable PPP logging, select the **Enable Point-to-Point Protocol (PPP) logging** option on the **PPP** tab on the properties of a remote access server. For more information, see View properties of the remote access server.

The PPP log in Windows NT 4.0 has been replaced by the tracing function. To duplicate the PPP log, you need to enable file tracing for the PPP key. By default, the PPP log is stored as the ppp.log file in the *systemroot*\Tracing folder. For more information on tracing options, see Using tracing.

**Network Monitor**
Network Monitor can capture the network traffic between a dial-up networking client and the remote access server. By analyzing remote access traffic, you can find answers to remote access problems and possible solutions or workarounds. For more information, see Install and Configure Network Monitor.

### Note

- The proper interpretation of remote access traffic with Network Monitor requires an in-depth understanding of PPP and other protocols. You can save Network Monitor traces as files and send them to Microsoft Product Support Services for analysis.

# VPN Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section covers:

- Troubleshooting remote access VPNs
- Troubleshooting router-to-router VPNs

# Troubleshooting remote access VPNs

Updated: August 28, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting remote access VPNs**
What problem are you having?

- I received error 800, which says the VPN server is unreachable
- I received error 721, which says the remote computer is not responding
- I received error 741/742, which says there is an encryption mismatch error
- I am unable to establish a remote access VPN connection
- VPN clients are unable to access resources beyond the VPN server

**I received error 800, which says the VPN server is unreachable.**

- **Cause:** PPTP/L2TP packets from the VPN client cannot reach the VPN server.
- **Solution:**
  - Assuming ping is allowed (that is, not blocked) between the VPN client and the VPN server, ping the VPN server. Confirm that you have network connectivity and that packet filters configured on the VPN server are not preventing communication.
  - By default, when you configure the Routing and Remote Access service for the VPN server role, only PPTP and L2TP/IPsec packets are allowed on the Internet interface of the VPN server. Internet Control Message Protocol (ICMP) packets used by the ping command are filtered out. To enable ping on the VPN server, add an input filter and an output filter that allow ICMP traffic.
  - To allow PPTP traffic, configure the network firewall to open TCP port 1723 and to forward IP protocol 47 for Generic Routing Encapsulation (GRE) traffic to the VPN server. Some firewalls refer to IP protocol 47 as VPN or PPTP pass-through. Not all firewalls support IP protocol 47. You may need to upgrade the firmware.
  - To allow L2TP traffic, configure the network firewall to open UDP port 1701 and to allow IPsec ESP formatted packets (IP protocol 50).
  - Check that the VPN server is listening for PPTP traffic on TCP port 1723 and is listening for L2TP traffic on UDP port 1701. On the VPN server, from the command line, type:

    **netstat -aon**

    You should see the following:

    | | | | |
    |---|---|---|---|
    | TCP | 0.0.0.0:1723 | 0.0.0.0:0 | LISTENING |
    | UDP | 0.0.0.0:1701 | *:* | |

  - You can use the Portqry command-line tool to determine if PPTP and L2TP packets are being dropped between the VPN client and VPN server. For more information, see Description of the Portqry.exe command-line utility on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=71609).

To query the PPTP port, use the following command: `portqry.exe -n <server_IP_address> -e 1723`

To query the L2TP port, use the following command: `portqry.exe -n <server_IP_address> -e 1701 -p UDP`

If the query output includes "NOT LISTENING/FILTERED", check the port configuration on the RRAS server.

**I received error 721, which says the remote computer is not responding.**

- **Cause:** This issue can occur if the network firewall does not permit GRE traffic (IP protocol 47). PPTP uses GRE for tunneled data.

- **Solution:** Configure the network firewall between the VPN client and the server to permit GRE. In addition, make sure that the network firewall permits TCP traffic on port 1723. Both of these conditions must be met to establish VPN connectivity by using PPTP. For more information, see article 888201, "You receive an "Error 721" error message when you try to establish a VPN connection through your Windows Server-based remote access server" in the Microsoft Knowledge Base.

| **Note** |
| --- |
| The firewall might be on or in front of the VPN client or in front of the VPN server. |

**I received error 741/742, which says there is an encryption mismatch error.**

- **Cause:** These errors occur if the VPN client requests an encryption level that is not valid or if the VPN server does not support an encryption type requested by the client.

- **Solution:**

  - Check the properties (**Security** tab) of the VPN connection on the VPN client. If **Require data encryption (disconnect if none)** is selected, clear the selection and retry the connection.

  - If you are using Internet Authentication Service (IAS), check the encryption level in the remote access policy in the IAS console or policies on other RADIUS servers. Ensure that the encryption level requested by the VPN client is selected on the VPN server.

**I am unable to establish a remote access VPN connection.**

- **Cause:** The user account used by the remote access client is locked out, expired, disabled, or the time the connection is being made does not correspond to the configured logon hours.

- **Solution:** Modify the account properties of the user account.

**See also:** Manage Users, Groups, and Computers

**Cause:** The Routing and Remote Access service is not started on the VPN server.

**Solution:** Verify the state of the Routing and Remote Access service on the VPN server.

**See also:** Monitor the Routing and Remote Access service; Start and stop the Routing and Remote Access service

**Cause:** Remote access is not enabled on the VPN server.

**Solution:** Enable remote access on the VPN server.

**See also:** Enable the remote access server

**Cause:** PPTP or L2TP ports are not enabled for inbound remote access requests.

**Solution:** Enable PPTP ports, L2TP ports, or both, for inbound remote access requests.

**See also:** Configure ports for remote access

**Cause:** The LAN protocols being used by the VPN clients are not enabled for remote access on the VPN server.

**Solution:** Enable the LAN protocols being used by the VPN clients for remote access on the VPN server.

**See also:** View properties of the remote access server

**Cause:** All of the PPTP or L2TP ports on the VPN server are already being used by currently connected remote access clients or demand-dial routers.

**Solution:** Verify that all of the PPTP or L2TP ports on the VPN server are not already being used by clicking **Ports** in Routing and Remote Access. If necessary, change the number of PPTP or L2TP ports to allow more concurrent connections.

**See also:** Add PPTP or L2TP ports

**Cause:** The tunneling protocol of the VPN client is not supported by the VPN server.

By default, remote access VPN clients running Windows XP or a Windows Server 2003 operating system use the **Automatic** server type option, which means that they try to establish a PPTP-based VPN connection first, and then they try an L2TP/IPSec-based VPN connection. If VPN clients use either the **Point-to-Point Tunneling Protocol (PPTP)** or **Layer Two Tunneling Protocol (L2TP)** server type option, verify that the selected tunneling protocol is supported by the VPN server.

By default, a computer running a Windows Server 2003 operating system and the Routing and Remote Access service is a PPTP and L2TP server with five L2TP ports and five PPTP ports. To create a PPTP-only server, set the number of L2TP ports to 0. To create an L2TP-only server, set the number of PPTP ports to 1 and disable remote access inbound connections and demand-dial connections.

**Solution:** Verify that the appropriate number of PPTP or L2TP ports is configured.

**See also:** Add PPTP or L2TP ports

**Cause:** The VPN client and the VPN server in conjunction with a remote access policy are not configured to use at least one common authentication method.

**Solution:** Configure the VPN client and the VPN server in conjunction with a remote access policy to use at least one common authentication method.

**See also:** Configure authentication

**Cause:** The VPN client and the VPN server in conjunction with a remote access policy are not configured to use at least one common encryption method.

**Solution:** Configure the VPN client and the VPN server in conjunction with a remote access policy to use at least one common encryption method.

**See also:** Configure encryption

**Cause:** The VPN connection does not have the appropriate permissions through dial-in properties of the user account and remote access policies.

**Solution:** Verify that the VPN connection has the appropriate permissions through dial-in properties of the user account and remote access policies.

In order for the connection to be established, the settings of the connection attempt must:

- Match all of the conditions of at least one remote access policy.
- Be granted remote access permission through the user account (set to **Allow access**) or through the user account (set to **Control access through Remote Access Policy**) and the remote access permission of the matching remote access policy (set to **Grant remote access permission**).
- Match all the settings of the profile.
- Match all the settings of the dial-in properties of the user account.

**See also:** Introduction to remote access policies; Accepting a connection attempt

**Cause:** The settings of the remote access policy profile are in conflict with properties of the VPN server.

The properties of the remote access policy profile and the properties of the VPN server both contain settings for:

- Multilink
- Bandwidth allocation protocol
- Authentication protocols

If the settings of the profile of the matching remote access policy are in conflict with the settings of the VPN server, the connection attempt is rejected. For example, if the matching remote access policy profile specifies that the EAP-TLS authentication protocol must be used, and EAP is not enabled on the VPN server, the connection attempt is rejected.

**Solution:** Verify that the settings of the remote access policy profile are not in conflict with properties of the VPN server.

**See also:** Enable authentication protocols; Configure authentication

**Cause:** The credentials of the VPN client (user name, password, and domain name) are incorrect and cannot be validated by the VPN server.

**Solution:** Verify that the credentials of the VPN client (user name, password, and domain name) are correct and can be validated by the VPN server.

**Cause:** There are not enough addresses in the static IP address pool.

**Solution:** If the VPN server is configured with a static IP address pool, verify that there are enough addresses in the pool. If all of the addresses in the static pool have been allocated to connected VPN clients, the VPN server is unable to allocate an IP address, and the connection attempt is rejected. Modify the static IP address pool if needed.

**See also:** Create a static IP address pool

**Cause:** The VPN client is configured to request its own IPX node number, and the VPN server is not configured to allow IPX clients to request their own IPX node number.

**Solution:** Configure the VPN server to allow IPX clients to request their own IPX node number.

**Cause:** The VPN server is configured with a range of IPX network numbers that are being used elsewhere on your IPX internetwork.

**Solution:** Configure the VPN server with a range of IPX network numbers that is unique to your IPX internetwork.

**Cause:** The authentication provider of the VPN server is improperly configured.

**Solution:** Verify the configuration of the authentication provider. You can configure the VPN server to use either Windows or RADIUS to authenticate the credentials of the VPN client.

**See also:** Use RADIUS authentication

**Cause:** The VPN server cannot access Active Directory.

**Solution:** For a VPN server that is a member server in a Windows 2000 mixed domain or a Windows 2000 native domain that is configured for Windows authentication, verify that:

- The **RAS and IAS Servers** security group exists. If not, then create the group and set the group type to **Security** and the group scope to **Domain local**.
- The **RAS and IAS Servers** security group has Read permission to the **RAS and IAS Servers Access Check** object.
- The computer account of the VPN server computer is a member of the **RAS and IAS Servers** security group. You can use the **netsh ras show registeredserver** command to view the current registration. You can use the **netsh ras add registeredserver** command to register the server in a specified domain.

If you add the VPN server computer to (or remove the VPN server computer from) the **RAS and IAS Servers** security group, the change does not take effect immediately (due to the way that Active Directory information is cached). For the change to take effect immediately, you need to restart the VPN server computer.

- For a Windows 2000 native domain, the VPN server has joined the domain.

**See also:** Create a new group; Verify permissions for the RAS and IAS security group; Netsh commands for remote access; Domain and forest functionality

**Cause:** A Windows NT 4.0 VPN server cannot validate connection requests.

**Solution:** If VPN clients are dialing in to a VPN server running Windows NT 4.0 that is a member of a Windows 2000 mixed domain, verify that the Everyone group is added to the Pre-Windows 2000 Compatible Access group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, type **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** at a command prompt on a domain controller computer and then restart the domain controller computer.

**See also:** Domain and forest functionality

**Cause:** The VPN server is unable to communicate with the configured RADIUS server.

**Solution:** If your RADIUS server is only reachable through your Internet interface, add an input filter and output filter to the Internet interface for UDP port 1812 (based on RFC 2138, "Remote Authentication Dial-In User Service (RADIUS)") or UDP port 1645 (for older RADIUS servers) for RADIUS authentication and UDP port 1813 (based on RFC 2139, "RADIUS Accounting") or UDP port 1646 (for older RADIUS servers) for RADIUS accounting.

**See also:** Add a packet filter

**Cause:** Cannot ping the VPN server from the Internet.

**Solution:** Due to the PPTP and L2TP/IPSec packet filtering that is configured on the Internet interface of the VPN server, Internet Control Message Protocol (ICMP) packets used by the **ping** command are filtered out. To enable ping capability on the VPN server, you need to add an input filter and an output filter that allow traffic for IP protocol 1 (ICMP traffic).

**See also:** Add a packet filter

**VPN clients are unable to access resources beyond the VPN server.**
**Cause:** The LAN protocols used by remote access VPN clients are not enabled to allow access to the network to which the VPN server is attached.

**Solution:** Configure the LAN protocols used by the remote access VPN clients to allow access to the network to which the VPN server is attached.

**See also:** View properties of the remote access server

**Cause:** A static IP address pool is configured but there are no routes back to the remote access VPN clients.

**Solution:** If the VPN server is configured to use a static IP address pool, verify that the routes to the ranges of addresses defined by the static IP address pool are reachable by the hosts and routers of the intranet. If not, then IP routes consisting of the address ranges of the static IP address pool, as defined by the IP address and mask of each range, must be added to the routers of the intranet, or the routing protocol of your routed infrastructure on the VPN server must be enabled. If the routes to the remote access VPN client subnets are not present, remote access VPN clients cannot receive traffic from locations on the intranet. A route for the network is implemented either through static routing entries or through a routing protocol, such as Open Shortest Path First (OSPF) or Routing Information Protocol (RIP).

If the VPN server is configured to use DHCP for IP address allocation, and no DHCP server is available, the VPN server allocates addresses from the Automatic Private IP Addressing (APIPA) address range from 169.254.0.1 through 169.254.255.254. Allocating APIPA addresses for remote access clients works only if the network to which the VPN server is attached is also using APIPA addresses.

If the VPN server is using APIPA addresses when a DHCP server is available, verify that the proper adapter is selected from which to obtain DHCP-allocated IP addresses. By default, the VPN server randomly chooses the adapter to use to obtain IP addresses through DHCP. If there is more than one LAN adapter, then the Routing and Remote Access service may choose a LAN adapter for which there is no DHCP server available.

If the static IP address pool consists of ranges of IP addresses that are a subset of the range of IP addresses for the network to

which the VPN server is attached, verify that the ranges of IP addresses of the static IP address pool are not assigned to other TCP/IP nodes, either through static configuration or through DHCP.

**See also:** Create a static IP address pool

**Cause:** Packet filters on the remote access policy profile are preventing the flow of IP traffic.

**Solution:** Verify that there are no configured TCP/IP packet filters on the profile properties of the remote access policies on the VPN server (or the RADIUS server if Internet Authentication Service is used) that are preventing the sending or receiving of TCP/IP traffic.

You can use remote access policies to configure TCP/IP input and output packet filters that control the exact nature of TCP/IP traffic allowed on the VPN connection. Verify that the profile TCP/IP packet filters are not preventing the flow of needed traffic.

**See also:** Configure IP options

For more information about troubleshooting remote access connections, see Remote Access Troubleshooting.

**Note**

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling Protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling Protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

# Troubleshooting router-to-router VPNs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting router-to-router VPNs**
What problem are you having?

- Unable to establish a router-to-router VPN connection.
- Unable to send and receive data.

**Unable to establish a router-to-router VPN connection.**
**Cause:** The user account used by the calling router is locked out, expired, disabled, or the time the connection is being made does not correspond to the configured logon hours.

**Solution:** Modify the account properties of the user account.

**See also:** Manage Users, Groups, and Computers

**Cause:** The Routing and Remote Access service is not started on the VPN client (the calling router) and the VPN server (the answering router).

**Solution:** Verify the state of the Routing and Remote Access service on the VPN client and the VPN server.

**See also:** Monitor the Routing and Remote Access service; Start and stop the Routing and Remote Access service

**Cause:** LAN and WAN routing is not enabled on the calling router and the answering router.

**Solution:** Enable **Local and remote routing (LAN and WAN router)** on the calling router and the answering router.

**See also:** Enable LAN and WAN routing

**Cause:** PPTP or L2TP ports are not enabled for inbound and outbound demand-dial routing connections.

**Solution:** Enable PPTP ports, L2TP ports, or both, for inbound and outbound demand-dial routing connections.

**See also:** Enable routing on ports

**Cause:** All of the PPTP or L2TP ports on the calling or answering router are already being used by currently connected remote access clients or demand-dial routers.

**Solution:** Verify that all of the PPTP or L2TP ports on the VPN server are not already being used by clicking **Ports** in Routing and Remote Access. If necessary, change the number of PPTP or L2TP ports to allow more concurrent connections.

**See also:** Add PPTP or L2TP ports

**Cause:** The tunneling protocol of the calling router is not supported by the answering router.

By default, demand-dial interfaces on a computer running a Windows Server 2003 operating system use the **Automatic** server type option, which means that they try to establish an L2TP/IPSec-based VPN connection first, and then they try a PPTP-based VPN connection. If calling routers use either the **Point-to-Point Tunneling Protocol (PPTP)** or **Layer-2 Tunneling Protocol (L2TP)** server type option, verify that the selected tunneling protocol is supported by the answering router.

By default, a computer running a Windows Server 2003 operating system and the Routing and Remote Access service is a PPTP and L2TP-capable demand dial router with five L2TP ports and five PPTP ports. To create a PPTP-only router, set the number of L2TP ports to 0. To create an L2TP-only router, set the number of PPTP ports to 1 and disable remote access inbound connections and demand-dial connections.

**Solution:** Verify that the appropriate number of PPTP or L2TP ports is configured on the calling router and answering router.

**See also:** Add PPTP or L2TP ports

**Cause:** The calling router and the answering router in conjunction with a remote access policy are not configured to use at least one common authentication method.

**Solution:** Configure the calling router and the answering router in conjunction with a remote access policy to use at least one common authentication method.

**See also:** Configure authentication

**Cause:** The calling router and the answering router in conjunction with a remote access policy are not configured to use at least one common encryption method.

**Solution:** Configure the calling router and the answering router in conjunction with a remote access policy to use at least one common encryption method.

**See also:** Configure encryption

**Cause:** The VPN connection does not have the appropriate permissions through dial-in properties of the user account and remote access policies.

**Solution:** Verify that the VPN connection has the appropriate permissions through dial-in properties of the user account and remote access policies.

In order for the connection to be established, the settings of the connection attempt must:

- Match all of the conditions of at least one remote access policy.
- Be granted remote access permission through the user account (set to **Allow access**) or through the user account (set to **Control access through Remote Access Policy**) and the remote access permission of the matching remote access policy (set to **Grant remote access permission**).
- Match all the settings of the profile.
- Match all the settings of the dial-in properties of the user account.

**See also:** Introduction to remote access policies; Accepting a connection attempt

**Cause:** The settings of the remote access policy profile are in conflict with properties of the answering router.

The properties of the remote access policy profile and the properties of the answering router both contain settings for:

- Multilink
- Bandwidth allocation protocol
- Authentication protocols

If the settings of the profile of the matching remote access policy are in conflict with the settings of the answering router, the connection attempt is rejected. For example, if the matching remote access policy profile specifies that the EAP-TLS authentication protocol must be used, and EAP is not enabled on the answering router, the connection attempt is rejected.

**Solution:** Verify that the settings of the remote access policy profile are not in conflict with properties of the VPN server.

**See also:** Enable authentication protocols; Configure authentication

**Cause:** The credentials of the calling router (user name, password, and domain name) are incorrect and cannot be validated by the answering router.

**Solution:** Verify that the credentials of the calling router (user name, password, and domain name) are correct and can be validated by the answering router.

**Cause:** There are not enough addresses in the static IP address pool.

**Solution:** If the answering router is configured with a static IP address pool, verify that there are enough addresses in the pool. If all of the addresses in the static pool have been allocated to connected remote access clients or demand-dial routers, the answering router is unable to allocate an IP address, and the connection attempt is rejected. Modify the static IP address pool if needed.

**See also:** Create a static IP address pool

**Cause:** The authentication provider of the answering router is improperly configured.

**Solution:** Verify the configuration of the authentication provider. You can configure the answering router to use either Windows or RADIUS to authenticate the credentials of the VPN client.

**See also:** Use RADIUS authentication

**Cause:** The answering router cannot access Active Directory.

**Solution:** For an answering router that is a member server in a Windows 2000 mixed domain or a Windows 2000 native domain that is configured for the Windows authentication provider, verify that:

- The **RAS and IAS Servers** security group exists. If not, then create the group and set the group type to **Security** and the group scope to **Domain local**.

- The **RAS and IAS Servers** security group has Read permission to the **RAS and IAS Servers Access Check** object.

- The computer account of the answering router computer is a member of the **RAS and IAS Servers** security group. You can use the **netsh ras show registeredserver** command to view the current registration. You can use the **netsh ras add registeredserver** command to register the server in a specified domain.

  If you add the answering router computer to (or remove the answering router from) the **RAS and IAS Servers** security group, the change does not take effect immediately (due to the way that Active Directory information is cached). For the change to take effect immediately, you need to restart the answering router computer.

- For a Windows 2000 native domain, the answering router has joined the domain.

**See also:** Create a new group; Verify permissions for the RAS and IAS security group; Netsh commands for remote access; Domain and forest functionality

**Cause:** An answering router running Windows NT 4.0 with the Routing and Remote Access Service (RRAS) cannot validate connection requests.

**Solution:** If calling routers are dialing in to a answering router running Windows NT 4.0 with RRAS that is a member of a Windows 2000 mixed domain, verify that the Everyone group is added to the Pre-Windows 2000 Compatible Access group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, type **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** at a command prompt on a domain controller computer and then restart the domain controller computer.

**See also:** Domain and forest functionality

**Cause:** The answering router is unable to communicate with the configured RADIUS server.

**Solution:** If your RADIUS server is only reachable through your Internet interface, add an input filter and output filter to the Internet interface for UDP port 1812 (based on RFC 2138, "Remote Authentication Dial-In User Service (RADIUS)") or UDP port 1645 (for older RADIUS servers) for RADIUS authentication and UDP port 1813 (based on RFC 2139, "RADIUS Accounting") or UDP port 1646 (for older RADIUS servers) for RADIUS accounting.

**See also:** Add a packet filter

**Cause:** Cannot ping the answering router from the Internet.

**Solution:** Due to the PPTP and L2TP/IPSec packet filtering that is configured on the Internet interface of the answering router, Internet Control Message Protocol (ICMP) packets used by the **ping** command are filtered out. To enable ping capability on the answering router, you need to add an input filter and an output filter that allow traffic for IP protocol 1 (ICMP traffic).

**See also:** Add a packet filter

**Unable to send and receive data.**
**Cause:** The appropriate demand-dial interface has not been added to the protocol being routed.

**Solution:** Add the appropriate demand-dial interface to the protocol being routed.

**See also:** Add a routing interface

**Cause:** There are no routes on both sides of the router-to-router VPN connection that support the two-way exchange of traffic.

**Solution:** Unlike a remote access VPN connection, a router-to-router VPN connection does not automatically create a default route. You need to create routes on both sides of the router-to-router VPN connection so that traffic can be routed to and from the other side of the router-to-router VPN connection.

You can manually add static routes to the routing table, or you can add static routes through routing protocols. For persistent VPN connections, you can enable Open Shortest Path First (OSPF) or Routing Information Protocol (RIP) across the VPN connection. For on-demand VPN connections, you can automatically update routes through an auto-static RIP update.

**See also:** Add an IP routing protocol; Add a static route; Perform Auto-Static Updates

**Cause:** A two-way initiated, router-to-router VPN connection is being interpreted by the answering router as a remote access connection.

**Solution:** If the user name in the credentials of the calling router appears under **Dial-In Clients** in Routing and Remote Access, the answering router interpreted the calling router as a remote access client. Verify that the user name in the credentials of the calling router matches the name of a demand-dial interface on the answering router.

If the incoming caller is a router, the port on which the call was received shows a status of **Active** and the corresponding demand-dial interface is in a **Connected** state.

**See also:** Check the status of the port on the answering router; Check the status of the demand-dial interface

**Cause:** Packet filters on the demand-dial interfaces of the calling router and answering router are preventing the flow of traffic.

**Solution:** Verify that there are no packet filters on the demand-dial interfaces of the calling router and answering router that prevent the sending or receiving of traffic.

You can configure each demand-dial interface with IP input and output filters to control the exact nature of TCP/IP traffic that is allowed into and out of the demand-dial interface.

**See also:** Manage Packet Filters

**Cause:** Packet filters on the remote access policy profile are preventing the flow of IP traffic.

**Solution:** Verify that there are no configured TCP/IP packet filters on the profile properties of the remote access policies on the VPN server (or the RADIUS server if Internet Authentication Service is used) that are preventing the sending or receiving of TCP/IP traffic.

You can use remote access policies to configure TCP/IP input and output packet filters that control the exact nature of TCP/IP traffic allowed on the VPN connection. Verify that the profile TCP/IP packet filters are not preventing the flow of needed traffic.

**See also:** Configure IP options

For more information on troubleshooting demand-dial connections, see Troubleshooting demand-dial routing.

# Internet Authentication Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Authentication Service**
Internet Authentication Service (IAS) in Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy. As a RADIUS server, IAS performs centralized connection authentication, authorization, and accounting for many types of network access including wireless, authenticating switch, and remote access dial-up and virtual private network (VPN) connections. As a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers. RADIUS is an Internet Engineering Task Force (IETF) standard. For more detailed information, see Features of IAS.

- Before installing an IAS server, see IAS checklists.
- For tips about using IAS, see IAS Best Practices.
- For help with specific tasks, see IAS How To....
- For general background information, see IAS Concepts.
- For problem-solving instructions, see IAS Troubleshooting.

**Notes**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# IAS checklists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

- Checklist: Configuring IAS for dial-up and VPN access
- Checklist: Configuring IAS to outsource dial-up access
- Checklist: Configuring IAS to forward requests
- Checklists: Configuring IAS for Wireless Access
- Checklists: Configuring IAS for Authenticated Switch Access

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Configuring IAS for dial-up and VPN access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring IAS for dial-up and VPN access**

| Step | Reference |
|---|---|
| Review RADIUS and IAS concepts. | IAS Overview; Understanding IAS |
| Review IAS implementation best practices. | IAS Best Practices |
| Review IAS security issues. | Security information for IAS |
| Install IAS on the servers to be used as IAS servers. (One server is configured first, and then the configuration is copied to the other server or servers.) | Install IAS |
| Configure the properties of one IAS server, including the ports used and event log settings. | Configure IAS Properties |
| Add the access servers (the dial-up network access servers and VPN servers) as clients on the same IAS server. | Configure RADIUS Clients |
| Set up the remote access policies that reflect your dial-up and VPN connection requirements on the same IAS server. | Introduction to remote access policies; Configure Remote Access Policies |
| Configure logging methods for user authentication and accounting requests. | Configure Logging for User Authentication and Accounting |
| Copy the IAS configuration from the first IAS server to additional IAS servers. | Copy the IAS configuration to another server |
| Register the IAS servers in the appropriate Active Directory domains. | Enable the IAS server to read user accounts in Active Directory |
| Verify the configuration of RADIUS accounting and authentication on the access servers, ensuring that IAS is configured on all access servers as the authentication and accounting provider. | IAS as a RADIUS server design considerations; Use RADIUS authentication; Use RADIUS accounting;Remote Access; Manufacturer's documentation |
| Verify that the access servers are properly configured for dial-up and VPN connections. | Routing and Remote Access; Manufacturer's documentation |

**Notes**

- For more information about how to deploy IAS with Routing and Remote Access, see Dial-up and VPN remote access.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Configuring IAS to outsource dial-up access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring IAS to outsource dial-up access**

| Step | Reference |
|---|---|
| Review RADIUS and IAS concepts. | IAS Overview; Understanding IAS |
| Review IAS implementation best practices. | IAS Best Practices |
| Review IAS security issues. | Security information for IAS |
| Configure the firewall to support the IAS servers in the perimeter network. | IAS and firewalls |
| Install IAS on the servers to be used as primary and backup IAS servers. | Install IAS |
| Add the Internet service provider (ISP) RADIUS proxies and the VPN servers as clients on the primary IAS server. | Configure RADIUS Clients; Contact the Internet service provider (ISP) |
| Configure the remote access policies that reflect your dial-up and VPN connection requirements on the primary IAS server. | Introduction to remote access policies; Configure Remote Access Policies |
| Configure logging methods for user authentication and accounting requests. | Configure Logging for User Authentication and Accounting |
| Copy the IAS configuration from the primary IAS server to the backup and remote IAS servers. | Copy the IAS configuration to another server |
| Register the primary and backup IAS servers in the appropriate Active Directory domains. | Enable the IAS server to read user accounts in Active Directory |
| Verify the configuration of RADIUS accounting and authentication on the VPN servers. Ensure that IAS is configured on all VPN servers as the authentication and accounting provider. | IAS as a RADIUS server design considerations; Use RADIUS authentication; Use RADIUS accounting; Remote Access; Manufacturer's documentation |
| Verify that the VPN servers are correctly configured for VPN connections. | Virtual Private Networks;Routing and Remote Access; Manufacturer's documentation |
| Verify the configuration of the RADIUS proxies at the ISP. | Contact the ISP |
| Verify that Connection Manager service profiles have been configured to support access to the corporate network, and that they work correctly (with Connection Point Services support, as required). | Connection Manager Administration Kit |

**Notes**

- For information about how to deploy IAS with outsourced dial-up access, see Outsourced VPN remote access.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Configuring IAS to forward requests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring IAS to forward requests**

| Step | Reference |
|---|---|
| Review RADIUS and IAS concepts. | IAS Overview; Understanding IAS |
| Review IAS implementation best practices. | IAS Best Practices |
| Review IAS security issues. | Security information for IAS |
| Install IAS on the computer that will be used as the IAS proxy server. | Install IAS |
| Configure the properties of the IAS proxy server, including the ports used and event log settings. | Configure IAS Properties |
| Add the RADIUS clients on the IAS proxy server. Configure the shared secret that is common to both the IAS proxy server and the RADIUS client (from which requests are received). | Add RADIUS clients |
| On the IAS proxy server, use the New Remote RADIUS Server Group Wizard to create a remote server group with one or more RADIUS servers to which RADIUS messages are forwarded. Configure RADIUS ports and shared secrets that are common to both the IAS proxy server and the RADIUS servers (to which requests are forwarded). | Add a remote RADIUS server group |
| On the IAS proxy server, use the New Connection Request Policy Wizard to create a connection request policy to forward connection requests and accounting information to the remote RADIUS server group. | Add a connection request policy |
| Configure logging methods for user authentication and accounting requests. | Configure Logging for User Authentication and Accounting |
| Optional. Copy the IAS proxy configuration from the IAS proxy server to additional IAS proxy servers. | Copy the IAS configuration to another server |
| Register the IAS proxy servers in the appropriate Active Directory domains. | Enable the IAS server to read user accounts in Active Directory |
| Verify the configuration of RADIUS accounting and authentication on the RADIUS clients. | IAS as a RADIUS server design considerations; Use RADIUS authentication; Use RADIUS accounting;Remote Access; Manufacturer's documentation |
| Configure the IAS servers as RADIUS clients on the RADIUS servers (to which requests are forwarded). | Add RADIUS clients; Manufacturer's documentation |

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklists: Configuring IAS for Wireless Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists: Configuring IAS for wireless access**
You can use the following checklists to configure both the IAS server and wireless access points, and to install user certificates on wireless clients.

- Checklist: Configuring the IAS server and wireless access points for wireless access
- Checklist: Installing a user certificate on a wireless client over a wireless connection
- Checklist: Installing a user certificate on a wireless client over an unauthenticated Ethernet connection
- Checklist: Installing a user certificate from floppy disk on a wireless client

### Note

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Configuring the IAS server and wireless access points for wireless access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring the IAS server and wireless access points for wireless access**

| Step | Reference |
|---|---|
| Review RADIUS and IAS concepts. | IAS Overview; Understanding IAS |
| Review IAS implementation best practices. | IAS Best Practices |
| Review IAS security issues. | Security information for IAS |
| If you are using certificates to authenticate wireless clients, install a computer certificate on the IAS server computers. | Computer certificates for certificate-based authentication |
| Install IAS on the servers to be used as primary and backup IAS servers. | Install IAS |
| Configure the properties of the primary IAS server, including the ports used and event log settings. | Configure IAS Properties |
| Configure logging methods for user authentication and accounting requests. | Configure Logging for User Authentication and Accounting |
| Add the wireless access points as RADIUS clients on the primary IAS server. | Add RADIUS clients |
| Use the New Remote Access Policy Wizard to create a common policy for wireless access. | Add a remote access policy |
| If you are using secure password authentication through Protected Extensible Authentication Protocol (PEAP), also called PEAP-EAP-MS-CHAP v2, configure authentication methods for the remote access policy. | Configure PEAP and EAP methods; PEAP |
| If you want client and server certificate authentication using Protected Extensible Authentication Protocol (PEAP), also called PEAP-EAP-TLS, configure authentication methods for the remote access policy. | Configure PEAP and EAP methods; PEAP |
| If you want client and server certificate authentication using PEAP-EAP-TLS, install a certificate on the wireless client from floppy disk, or deploy smart cards. | Checklist: Installing a user certificate from floppy disk on a wireless client; Checklist: Deploying smart cards for logging on to Windows; PEAP; Network access authentication and certificates |
| If you are using certificate authentication with EAP-TLS and initially installing a user certificate on your wireless clients over a wireless connection, enable guest authentication. | Guest authentication |
| If you are using certificate authentication with EAP-TLS and initially installing a user certificate on your wireless clients over a wireless connection, create a group named Guests and add the Guest account as a member. | Create a new group; Add a member to a group |

| | |
|---|---|
| If you are using certificate authentication with EAP-TLS and initially installing a user certificate on your wireless clients over a wireless connection, use the New Remote Access Policy Wizard to create a custom policy for new wireless clients (clients that do not have user certificates). Set the **NAS-Port-Type** condition to **Wireless-IEEE 802.11** and **Wireless-Other**, and the **Windows-Groups** condition to **Guests**. On the **Dial-in Constraints** tab of the profile, restrict the maximum session time to 10 minutes. On the **Advanced** tab of the profile, add the **Tunnel-Type** attribute with the value of **Virtual LANs (VLAN)**, and then add the **Tunnel-Pvt-Group-ID** attribute with the VLAN ID value that corresponds to guest wireless clients. | Add a remote access policy |
| Copy the IAS configuration from the primary IAS server to the backup IAS server. | Copy the IAS configuration to another server |
| Register the primary and backup IAS servers in the appropriate Active Directory domains. | Enable the IAS server to read user accounts in Active Directory |
| Verify the configuration of the wireless access points. Ensure that the RADIUS servers used for authentication and accounting for the wireless access point are the IAS server computers. | Manufacturer's documentation |
| Optional. Install a user certificate on wireless clients over a wireless connection. | Checklist: Installing a user certificate on a wireless client over a wireless connection |
| Optional. Install user certificates on wireless clients over an unauthenticated Ethernet connection. | Checklist: Installing a user certificate on a wireless client over an unauthenticated Ethernet connection |
| Optional. Install user certificates from floppy disk on wireless clients. | Checklist: Installing a user certificate from floppy disk on a wireless client |

For more information, see Wireless access and Outsourced wireless access.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Installing a user certificate on a wireless client over a wireless connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a user certificate on a wireless client over a wireless connection**

| Step | Reference |
|---|---|
| Configure the IAS server and wireless access points for wireless access. | Checklist: Configuring the IAS server and wireless access points for wireless access |
| Configure the network connection that corresponds to your wireless network adapter for EAP-TLS authentication. | Right-click the network connection, and then click **Properties**. Click the **Authentication** tab. Select **Smart Card or other Certificate** as the Extensible Authentication Protocol (EAP) type. Click **Properties**, and then click **Use a certificate on this computer**. Click **OK** three times to save changes to the network connection. |
| Move your wireless client computer within range of the configured wireless access point. | |
| After your wireless client computer has guest access, join the domain. | Join a domain |
| Restart the computer when prompted and log on to the domain. You will still have a guest connection. | |
| Install a user certificate on the wireless client computer. | Request a certificate;<br><br>Submit a user certificate request via the Web |
| After 10 minutes, your wireless client computer is disconnected because of the dial-in constraint in the new wireless client remote access policy that has been configured on the IAS server. When the wireless client computer reconnects, it will authenticate with the new user certificate and receive an authenticated connection to the wireless access point. | |

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Installing a user certificate on a wireless client over an unauthenticated Ethernet connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a user certificate on a wireless client over an unauthenticated Ethernet connection**

| Step | Reference |
|---|---|
| Configure the IAS server and wireless access points for wireless access. | Checklist: Configuring the IAS server and wireless access points for wireless access |
| Connect your wireless client computer to the organization network using an unauthenticated Ethernet connection. | |
| Join the domain. | Join a domain |
| Restart the computer when prompted and log on to the domain. | |
| Install a user certificate on the wireless client computer. | Request a certificate;<br><br>Submit a user certificate request via the Web |
| Configure the network connection that corresponds to your wireless network adapter for EAP-TLS authentication. | Right-click the network connection, and then click **Properties**. Click the **Authentication** tab. Select **Smart Card or other Certificate** as the Extensible Authentication Protocol (EAP) type. Click **Properties**, and then click **Use a certificate on this computer**. Click **OK** three times to save changes to the network connection. |
| Move your wireless client computer within range of the configured wireless access point. | |
| When the wireless client computer connects, it will authenticate with the new user certificate and receive an authenticated connection to the wireless access point. | |

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Installing a user certificate from floppy disk on a wireless client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a user certificate from floppy disk on a wireless client**

| Step | Reference |
|---|---|
| Configure the IAS server and wireless access points for wireless access. | Checklist: Configuring the IAS server and wireless access points for wireless access |
| Submit a user certificate request using the CA Web enrollment tool and save the certificate to floppy disk. | To save a certificate to floppy disk, submit an advanced certificate request by using the CA Web enrollment tool. You can connect to the CA Web enrollment tool at http://*ServerName*/certsrv, where *ServerName* is the name of the Web server at which the certification authority you want to access is located. Click **Advanced certificate request**, and then click **Create and submit a request to this CA**.<br><br>On the **Advanced Certificate Request** page, select **Export keys to file**, and then type a path statement (including file name) in **Full path statement**. In **Additional Options**, select the PKCS 10 format, and then click **Submit**. When prompted, type a password and password confirmation in **Create Private Key Password**, and then click **OK**. In **Certificate Issued**, select **DER encoded** and **Download certificate chain**. When prompted, save the file in a selected location. For more information, see Submit an advanced certificate request via the Web to a Windows Server 2003 CA. |
| Import the user certificate into the certificate store of the computer for which you want to use the certificate. | Insert the floppy disk that contains the user certificate into the floppy drive and run the Certificate Import Wizard. For more information, see Import a certificate. |

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklists: Configuring IAS for Authenticated Switch Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists: Configuring IAS for authenticated switch access**
You can use the following checklists to configure both the IAS server and Ethernet switch clients, and to install user certificates on Ethernet clients.

- Checklist: Configuring the IAS server for authenticated switch access
- Checklist: Configuring a first-time Ethernet switch client for authenticated access
- Checklist: Installing a computer certificate on an Ethernet client over an unauthenticated Ethernet connection
- Checklist: Installing a user certificate from floppy disk on an Ethernet client

### Note

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Configuring the IAS server for authenticated switch access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring the IAS server for authenticated switch access**

| Step | Reference |
|---|---|
| Review RADIUS and IAS concepts. | IAS Overview; Understanding IAS |
| Review IAS implementation best practices. | IAS Best Practices |
| Review IAS security issues. | Security information for IAS |
| If you are using certificates for authentication, install a computer certificate on IAS server computers. | Computer certificates for certificate-based authentication |
| Install IAS on the servers to be used as primary and backup IAS servers. | Install IAS |
| Configure the properties of the primary IAS server, including the ports used and event log settings. | Configure IAS Properties |
| Configure logging methods for user authentication and accounting requests. | Configure Logging for User Authentication and Accounting |
| Add the Ethernet switches as RADIUS clients on the primary IAS server. | Add RADIUS clients |
| Use the New Remote Access Policy Wizard to create a common policy for Ethernet access. | Add a remote access policy |
| If you are using certificate authentication and initially installing computer certificates on your Ethernet clients over an authenticated Ethernet connection, enable guest authentication. | Guest authentication |
| If you are using certificate authentication and initially installing a computer certificate on your Ethernet clients over an authenticated Ethernet connection, create a group named Guests and add the Guest account as a member. | Create a new group; Add a member to a group |
| If you are using certificate authentication and initially installing computer certificates on your Ethernet clients over an authenticated Ethernet connection, use the New Remote Access Policy Wizard to create a custom policy for new Ethernet switch clients (clients that do not have computer certificates that are used in the authentication process). Set the **NAS-Port-Type** condition to **Ethernet** and the **Windows-Groups** condition to Guests. On the **Dial-in Constraints** tab of the profile, restrict the maximum session time to 10 minutes. On the **Advanced** tab of the profile, add the **Tunnel-Type** attribute with the value of **Virtual LANs (VLAN)**, and then add the **Tunnel-Pvt-Group-ID** attribute with the VLAN ID value that corresponds to guest Ethernet clients. | Add a remote access policy |

| | |
|---|---|
| Copy the IAS configuration from the primary IAS server to the backup IAS server. | Copy the IAS configuration to another server |
| Register the primary and backup IAS servers in the appropriate Active Directory domains. | Enable the IAS server to read user accounts in Active Directory |
| Verify the configuration of the Ethernet switches. Ensure that the RADIUS servers used for the Ethernet switch authentication and accounting are the IAS server computers. | Manufacturer's documentation |
| Optional. Configure first-time Ethernet switch clients for authenticated access. | Checklist: Configuring a first-time Ethernet switch client for authenticated access |
| Optional. Install a computer certificate on Ethernet clients over an unauthenticated Ethernet connection. | Checklist: Installing a computer certificate on an Ethernet client over an unauthenticated Ethernet connection |
| Optional. Install user certificates from floppy disk on Ethernet clients. | Checklist: Installing a user certificate from floppy disk on an Ethernet client |

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Configuring a first-time Ethernet switch client for authenticated access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring a first-time Ethernet switch client for authenticated access**
This checklist applies only to Ethernet client computers that use EAP-TLS authentication.

| Step | Reference |
|---|---|
| Configure the IAS server for authenticated switch access. | Checklist: Configuring the IAS server for authenticated switch access |
| Configure the network connection that corresponds to your Ethernet network adapter for EAP-TLS authentication. | Right-click the network connection, and then click **Properties**. Click the **Authentication** tab, and then select **Smart Card or other Certificate** as the Extensible Authentication Protocol (EAP) type. Click **Properties**, and then click **Use a certificate on this computer**. Click **OK** three times to save changes to the network connection. |
| Connect your Ethernet client computer to the Ethernet switch. | |
| After your Ethernet client computer has unauthenticated access, join the domain. | Join a domain |
| Restart the computer when prompted and log on to the domain. Although you will still have a guest connection, the Ethernet client computer is issued a computer certificate during the startup and domain logon process. | |
| After 10 minutes, your Ethernet client computer is disconnected because of the dial-in constraint in the first-time Ethernet client remote access policy that has been configured on the IAS server. When the Ethernet client computer reconnects, it will authenticate with the new computer certificate and receive an authenticated connection to the Ethernet switch. | |

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Installing a computer certificate on an Ethernet client over an unauthenticated Ethernet connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a computer certificate on an Ethernet client over an unauthenticated Ethernet connection**
This checklist applies only to Ethernet client computers that use EAP-TLS authentication.

| Step | Reference |
|---|---|
| Configure the IAS server for authenticated switch access. | Checklist: Configuring the IAS server for authenticated switch access |
| Connect your Ethernet client computer to the organization network with an unauthenticated Ethernet connection. | |
| Join the domain. | Join a domain |
| Restart the computer when prompted and log on to the domain. | |
| Install a user certificate on the Ethernet client computer. | Request a certificate;<br><br>Submit a user certificate request via the Web |
| Configure the network connection that corresponds to your Ethernet network adapter for EAP-TLS authentication. | Right-click the network connection, and then click **Properties**. Click the **Authentication** tab. Select **Smart Card or other Certificate** as the Extensible Authentication Protocol (EAP) type. Click **Properties**, and then click **Use a certificate on this computer**. Click **OK** three times to save changes to the network connection. |
| Connect your Ethernet client computer to the organization network by using an authenticated Ethernet connection. | |
| When the Ethernet client computer connects, it will authenticate with the new user certificate and receive an authenticated connection to the network. | |

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Checklist: Installing a user certificate from floppy disk on an Ethernet client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a user certificate from floppy disk on an Ethernet client**
This checklist applies only to Ethernet client computers that use EAP-TLS authentication.

| Step | Reference |
|---|---|
| Configure the IAS server for authenticated switch access. | Checklist: Configuring the IAS server for authenticated switch access |
| Submit a user certificate request using the CA Web enrollment tool and save the certificate to floppy disk. | To save a certificate to floppy disk, submit an advanced certificate request by using the CA Web enrollment tool. You can connect to the CA Web enrollment tool at http://*ServerName*/certsrv, where *ServerName* is the name of the Web server at which the certification authority you want to access is located. Click **Advanced certificate request**, and then click **Create and submit a request to this CA**.<br><br>On the **Advanced Certificate Request** page, select **Export keys to file**, and then type a path statement (including file name) in **Full path statement**. In **Additional Options**, select the PKCS 10 format, and then click **Submit**. When prompted, type a password and password confirmation in **Create Private Key Password**, and then click **OK**. In **Certificate Issued**, select **DER encoded** and **Download certificate chain**. When prompted, choose to save the file, and then select a location for it. For more information, see Submit an advanced certificate request via the Web to a Windows Server 2003 CA. |
| Import the user certificate into the certificate store of the computer for which you want to use the certificate. | Insert the floppy disk that contains the user certificate into the floppy drive and run the Certificate Import Wizard. For more information, see Import a certificate. |

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# IAS Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS Best Practices**
This topic provides best practices for implementing and configuring IAS and is based on recommendations from Microsoft Product Support Services.

**Installation suggestions**
Before installing IAS, do the following:

- Install and test each of your access servers using local authentication methods before making them RADIUS clients.

- After you install and configure IAS, save the configuration by using the **netsh aaaa show config** > *path**file.txt* command. For more information, see Netsh commands for AAAA. Save the IAS configuration with the **netsh aaaa show config** > *path**file.txt* command each time a change is made.

- Do not install Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition on the same partition as Windows 2000. These operating systems use common files in the *systemroot*\Program Files folder to access the IAS database. If you decide to install Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition on the same partition as Windows 2000, IAS in Windows 2000 can no longer access remote access policies or remote access logging.

- Do not configure a server running IAS or Routing and Remote Access and Windows Server 2003 as a member of a Windows NT Server 4.0 domain if your user accounts database is stored on a Windows Server 2003 domain controller in another domain. Doing this will cause Lightweight Directory Access Protocol (LDAP) queries from the IAS server to the Windows Server 2003 domain controller to fail.

- Instead, configure your server running IAS or Routing and Remote Access and Windows Server 2003 as a member of a Windows Server 2003 domain. Alternately, you can configure a server running IAS and Windows Server 2003 as a proxy server that forwards authentication and accounting requests to another server running IAS and Windows Server 2003 that can access the user accounts database on the Windows Server 2003 domain controller. For more information, see Deploying IAS as a RADIUS Proxy.

**Security issues**
When you are administering an IAS server remotely, do not send sensitive or confidential data (for example shared secrets or passwords) over the network in plaintext. There are two recommended methods for remote administration of IAS servers:

- Use Terminal Services to access the IAS server.

  When you use Terminal Services, data is not sent between client and server. Only the user interface of the server (for example, the operating system desktop and IAS console image) is sent to the Terminal Services client, which is named Remote Desktop Connection in Windows XP. The client sends keyboard and mouse input, which is processed locally by the server that has Terminal Services enabled. When Terminal Services users log on, they can view only their individual client sessions, which are managed by the server and are independent of each other. In addition, Remote Desktop Connection provides 128-bit encryption between client and server. For more information, see Terminal Services.

- Use IPSec to encrypt confidential data.

  You can use IPSec to encrypt communication between the IAS server and the remote client computer that is being used to administer it. In order to administer the server remotely, the Windows Server 2003 Administration Tools Pack must be installed on the client computer, and the IAS snap-in must be added to the Microsoft Management Console (MMC). For more information, see IPSec Policy Rules.

Your IAS server provides authentication, authorization, and accounting for connection attempts to your organization network. You can protect your IAS server and RADIUS messages from unwanted internal and external intrusion. For more information about how to secure your IAS server, see Securing IAS.

For additional information about securing RADIUS traffic when IAS is used as a RADIUS server, see IAS as a RADIUS server security considerations. For additional information about securing RADIUS traffic when IAS is used as a RADIUS proxy, see IAS as a RADIUS proxy security considerations.

**Use the Runas command to administer local IAS servers**

You can use the **Runas** command to perform administrative tasks when you are logged on as a member of a group that does not have the required administrative credentials (such as the Users group or the Power Users group). Logging onto your server without administrative credentials is recommended because it protects the computer from a variety of possible security attacks, such as the accidental installation of a computer virus.

**Logging**

There are two types of logging in IAS:

1. Event logging for IAS You can use event logging to record IAS events in the system event log. This is used primarily for auditing and troubleshooting connection attempts.

2. Logging user authentication and accounting requests You can log user authentication and accounting requests to log files in text format or database format, or you can log to a stored procedure in a SQL Server 2000 database. Request logging is used primarily for connection analysis and billing purposes, and is also useful as a security investigation tool, providing you with a method of tracking down the activity an attacker.

To make the most effective use of IAS logging:

- Turn on logging (initially) for both authentication and accounting records. Modify these selections after you have determined what is appropriate for your environment.

- Ensure that event logging is configured with a capacity that is sufficient to maintain your logs.

- Back up all log files on a regular basis, since they cannot be recreated when they are damaged or deleted.

- Use the RADIUS Class attribute to both track usage and simplify the identification of which department or user to charge for usage. Although the automatically generated Class attribute is unique for each request, duplicate records might exist in cases where the reply to the access server is lost and the request is resent. You might need to delete duplicate requests from your logs to accurately track usage.

- To provide failover and redundancy with SQL Server logging, place two computers running SQL Server on different subnets. Use the SQL Server Create Publication Wizard to set up database replication between the two servers. For more information, see SQL Server 2000 documentation.

You can use the Iasparse.exe tool in the \Support\Tools folder on the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition compact disc to view IAS logs.

For more information, see Remote Access Logging.

**Performance-tuning IAS**

- To optimize IAS authentication and authorization response times and minimize network traffic, install IAS on a domain controller.

- When universal principal names (UPNs) or Windows Server 2003 domains are used, IAS uses the global catalog to authenticate users. To minimize the time it takes to do this, install IAS on either a global catalog server or a server that is on the same subnet. For more information, see The role of the global catalog. For more information about domain functionality, see Domain and forest functionality.

- When you have remote RADIUS server groups configured and, in IAS **Connection Request Policies**, you clear the **Record accounting information on the servers in the following remote RADIUS server group** check box, these

groups are still sent network access server (NAS) start and stop notification messages. This creates unnecessary network traffic. To eliminate this traffic, disable NAS notification forwarding for individual servers in each remote RADIUS server group by clearing the **Forward network start and stop notifications to this server** check box. For more information, see Configure the authentication and accounting settings of a group member and Configure accounting.

#### Using IAS in large organizations

- If you are using remote access policies to restrict access for all but certain groups, create a universal group for all of the users for whom you want to allow access, and then create a remote access policy that grants access for this universal group. Do not put all of your users directly into the universal group, especially if you have a large number of them on your network. Instead, create separate groups that are members of the universal group, and add users to those groups. For more information about universal groups, see Group scope. For more information about restricting or granting access to a group, see Allow dial-up connection using group membership.

- Use a user principal name to refer to users whenever possible. A user can have the same user principal name regardless of domain membership. This practice provides scalability that might be required in organizations with a large number of domains.

- If the IAS server is on a computer other than a domain controller and it is receiving a very large number of authentication requests per second, you can improve performance by increasing the number of concurrent authentications between the IAS server and the domain controller.

  To do this, edit the following registry key: **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters**. Add a new value named **MaxConcurrentApi** and assign to it a value from 2 through 5.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

### Notes

- If you assign a value to **MaxConcurrentApi** that is too high, your IAS server might place an excessive load on your domain controller.

- To effectively balance the load of either a large number of authorizations or a large volume of RADIUS authentication traffic (such as a large wireless implementation using certificate-based authentication), install IAS as a RADIUS server on all of your domain controllers. Next, configure two or more IAS proxies to forward the authentication requests between the access servers and the RADIUS servers. Next, configure your access servers to use the IAS proxies as RADIUS servers. For more information, see Using IAS proxy for load balancing.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# IAS How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install IAS
- Manage a remote IAS server
- Copy the IAS configuration to another server
- Enable the IAS server to read user accounts in Active Directory
- Start or stop IAS
- Configure IAS Properties
- Configure RADIUS Clients
- Configure Remote Access Policies
- Configure Connection Request Processing
- Configure Logging for User Authentication and Accounting

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Install IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To install IAS

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. In the **Windows Components Wizard** dialog box, click **Networking Services**, and then click **Details**.
4. In the **Networking Services** dialog box, select **Internet Authentication Service**, click **OK**, and then click **Next**.
5. If prompted, insert your Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition compact disc.
6. After IAS is installed, click **Finish**, and then click **Close**.

### Notes

- You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might prevent you from completing this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Manage a remote IAS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage a remote IAS server**

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. In the MMC console, click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Internet Authentication Service (IAS)**, and then click **Add**.
4. Select the computer for which you want to manage IAS:

| To | Do this |
|---|---|
| Manage IAS on the computer on which this console is running | Click **Local computer (the one this console is running on)**. |
| Manage IAS on a remote computer | Click **Another computer**, and then type the name or IP address of the remote IAS server. |

5. Click **Finish**, click **Close**, and then click **OK**.

### Note

- You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable the IAS server to read user accounts in Active Directory
Configure IAS Properties
Start or stop IAS

# Copy the IAS configuration to another server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy the IAS configuration to another server**

1. Open **Command Prompt**.
2. At the command prompt, type **netsh aaaa show config >***path\file.txt*.

   This stores configuration settings (including registry settings) in a text file. The path can be relative or absolute, or it can be a UNC path.
3. Copy the file you created to the destination computer.
4. At a command prompt on the destination computer, type **netsh exec** *path\file.txt*.

   A message appears indicating whether the update was successful.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- You do not need to stop IAS on the destination computer to run the **netsh exec** command. When the command is run, IAS is automatically refreshed with the updated configuration settings.
- Because IAS configurations are not encrypted in the text file, sending it over a network might pose a security risk.
- Do not use this procedure if the source IAS database is a higher version number than the version number of the destination IAS database. You can view the version number of the IAS database from the display of the **netsh aaaa show config** command.
- This procedure replicates the entire IAS, remote access policy, connection request policy, registry, and logging configuration.
- You can copy a remote access configuration from a server running Routing and Remote Access and Windows 2000 or Windows Server 2003 to an IAS server. Use the **netsh aaaa dump** command at the remote access server to create a file that contains the configuration, copy it to the IAS server, and then use the **netsh exec** command to import the configuration to the IAS server. Next, configure RADIUS clients on the IAS server.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

The Netsh Command-Line Utility

Netsh commands for remote access

Managing multiple IAS servers

# Enable the IAS server to read user accounts in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable the IAS server to read user accounts in Active Directory**
**To register the IAS server in the default domain using Internet Authentication Service:**

1. Log on to the IAS server with an account that has administrative credentials for the domain.
2. Open Internet Authentication Service.
3. Right-click **Internet Authentication Service**, and then click **Register Server in Active Directory**. When the **Register Internet Authentication Service in Active Directory** dialog box appears, click **OK**.

**To register the IAS server in the default domain using the netsh command:**

1. Log on to the IAS server with an account that has administrative credentials for the domain.
2. Open Command Prompt.
3. At the command prompt, type: **netsh ras add registeredserver**

**To register the IAS server in the default domain using Active Directory Users and Computers:**

1. Log on to the IAS server with an account that has administrative credentials for the domain.
2. Open Active Directory Users and Computers.
3. In the console tree, click the **Users** folder in the appropriate domain.
4. In the details pane, right-click **RAS and IAS Servers**, and then click **Properties**.
5. In the **RAS and IAS Servers Properties** dialog box, on the **Members** tab, add each of the IAS servers.

**To register the IAS server in another domain using Active Directory Users and Computers:**

1. Log on to the IAS server with an account that has administrative credentials for the domain.
2. Open Active Directory Users and Computers.
3. In the console tree, click the **Users** folder in the appropriate domain.
4. In the details pane, right-click **RAS and IAS Servers**, and then click **Properties**.
5. In the **RAS and IAS Servers Properties** dialog box, on the **Members** tab, add each of the IAS servers.

**To register the IAS server in another domain using the netsh command:**

1. Log on to the IAS server with an account that has administrative credentials for the domain.
2. Open Command Prompt.
3. At the command prompt, type **netsh ras add registeredserver** *Domain IASServer*, where *Domain* is the DNS domain name of the domain and *IASServer* is the name of the IAS server computer.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- After you register the service in Active Directory, you can verify the security settings. For more information, see Related Topics.

- You can add the IAS server to the **RAS and IAS Servers** group with the Dsmod tool.

- For more information about how to use IAS or Routing and Remote Access servers as stand-alone servers or member servers in Windows NT 4.0 domains, Windows 2000 mixed domains, Windows 2000 native domains, and Windows Server 2003 domains, see the links below.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Verify permissions for the RAS and IAS security group
Stand-alone server
Member server in a domain
Dsmod
Domain and forest functionality

# Start or stop IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To start or stop IAS

1. Open Internet Authentication Service.
2. In the console tree, right-click **Internet Authentication Service**, and then click **Start Service**.

   Or, to stop the service, click **Stop Service**.

### Note

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure IAS Properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure IAS properties**

- Configure IAS port information
- Configure event logging for IAS

# Configure IAS port information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure IAS port information**

1. Open Internet Authentication Service.

2. Right-click **Internet Authentication Service**, and then click **Properties**.

3. Click the **Ports** tab, and then examine the settings for ports. If your RADIUS authentication and RADIUS accounting UDP ports vary from the default values provided (1812 and 1645 for authentication, and 1813 and 1646 for accounting), type your port settings in **Authentication** and **Accounting**.

   To use multiple port settings for authentication or accounting requests, separate the port numbers with commas.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

- To configure a remote IAS server, you must have administrative credentials on the remote server.

- The values of 1812 for authentication and 1813 for accounting are RADIUS standard ports defined in RFCs 2865 and 2866. However, many access servers use ports 1645 for authentication requests and 1646 for accounting requests by default. Whatever port numbers you decide to use, make sure that IAS and your access server are configured to use the same ones.

- You can specify an IP address to which the RADIUS messages must be sent with the following syntax: *IPAddress***:***UDPPort*. For example, if you have multiple network adapters and you only want to receive RADIUS authentication messages sent to the IP address of 10.0.0.99 and UDP port 1812, you would type **10.0.0.99:1812** in **Authentication**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
RADIUS protocol
Working with MMC console files

# Configure event logging for IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure event logging for IAS**

1. Open Internet Authentication Service.
2. Right-click **Internet Authentication Service (IAS)**, and then click **Properties**.
3. On the **General** tab, select each required option, and then click **OK**.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- Rejected, discarded, and successful authentication requests are logged by default.
- Selecting **Successful authentication requests** can result in logging extremely large volumes of data. Before selecting this option, verify that the Event Viewer is configured with a maximum log size that will accommodate this type of event logging.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Event Viewer overview
Event Viewer

# Configure RADIUS Clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure RADIUS clients**

- Add RADIUS clients
- Delete a RADIUS client
- Edit RADIUS client configuration
- Configure the Message Authenticator attribute and shared secret

### Note

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Add RADIUS clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add RADIUS clients**

1. Open Internet Authentication Service.
2. Right-click **RADIUS Clients**, and then click **New RADIUS Client**.
3. Use the New RADIUS Client Wizard to both add and configure a client.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- If the client is a network access server (NAS) and you are planning to use NAS-specific remote access policies for configuration (for example, a remote access policy that contains vendor-specific attributes), in the New RADIUS Client Wizard, click **Client Vendor**, and then select the name of the manufacturer. If you do not know the name of the manufacturer or it is not in the list, click **RADIUS Standard**. If the RADIUS client is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition and the Routing and Remote Access service, click **Microsoft**.
- When you have a number of RADIUS clients on the same subnet (such as, numerous wireless access points), you can simplify RADIUS client administration by specifying an address range while running the wizard, instead of specifying the IP address or DNS name of a single RADIUS client. All of the RADIUS clients in the range must use the same configuration and shared secret. The address range for RADIUS clients is expressed in the network prefix length notation *w.x.y.z/p*, where *w.x.y.z* is the dotted decimal notation of the address prefix and *p* is the prefix length (the number of high order bits that define the network prefix). This is also known as Classless Inter-Domain Routing (CIDR) notation. An example is 192.168.21.0/24. To convert from subnet mask notation to network prefix length notation, p is the number of high order bits set to one in the subnet mask.
- If your access server supports use of the Message Authenticator attribute (also known as the signature attribute), in the New RADIUS Client Wizard, click **Request must contain the Message Authenticator attribute**. If the access server does not support the Message Authenticator attribute, do not select this option.
- If IAS receives an access request from a RADIUS proxy, it cannot detect the manufacturer of the NAS that originated the request. This can cause problems if you plan to use remote access policy conditions that are based on the client vendor and have at least one client that is defined as a RADIUS proxy. Requests from the proxy might not match any of the remote access policies and be denied.
- Shared secrets are case-sensitive. Verify that the client's shared secret and the shared secret you type in **Shared secret** are identical. For more information, see Related Topics.
- If the client address cannot be resolved when you click **Verify**, make sure that the DNS name that you typed is correct.
- You can use the friendly name that you provide for your RADIUS clients with the **Client-Friendly-Name** condition in remote access policies. For more information, see Related Topics.
- Enabling the use of the Message Authenticator attribute provides additional security when PAP, CHAP, MS-CHAP, and MS-CHAP v2 are used for authentication. EAP uses the Message Authenticator attribute by default and does not require that you enable it. For more information, see Related Topics.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum

of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared secrets
Message Authenticator attribute
IAS as a RADIUS server design considerations
Elements of a remote access policy
Delete a RADIUS client

# Delete a RADIUS client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a RADIUS client**

1. Open Internet Authentication Service.
2. In the console tree, click **RADIUS Clients**.
3. In the details pane, right-click the client that you want to delete, and then click **Delete**.
4. In the **Delete Client** dialog box, click **Yes**.

### Note

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
IAS as a RADIUS server design considerations
Edit RADIUS client configuration
Add RADIUS clients

# Edit RADIUS client configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit RADIUS client configuration**

1. Open Internet Authentication Service.
2. In the console tree, click **RADIUS Clients**.
3. In the details pane, double-click the client for which you want to edit the configuration.
4. Make changes to the RADIUS client configuration.
5. Click **OK** to save your changes.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- If IAS receives an access request from a RADIUS proxy server, IAS cannot detect the manufacturer of the NAS that originated the request. This can cause problems if you plan to use remote access policy conditions that are based on the client vendor and have at least one client that is defined as a RADIUS proxy server. Requests from the proxy might not match any of the remote access policies and be denied.
- Shared secrets are case-sensitive. Verify that the client's shared secret and the shared secret that you type in **Shared secret** are identical. For more information, see Related Topics.
- If the client address cannot be resolved when you click **Verify**, make sure the DNS name that you typed is correct.
- You can use the friendly name that you provide for your RADIUS clients with the **Client-Friendly-Name** condition in remote access policies. For more information, see Related Topics.
- Enabling the use of the Message Authenticator attribute provides additional security when PAP, CHAP, MS-CHAP, and MS-CHAP v2 are used for authentication. EAP uses the Message Authenticator attribute by default and does not require that you enable it. For more information, see Related Topics.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared secrets

Message Authenticator attribute
IAS as a RADIUS server design considerations
Elements of a remote access policy
Configure the Message Authenticator attribute and shared secret

# Configure the Message Authenticator attribute and shared secret

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the Message Authenticator attribute and shared secret**

1. Open Internet Authentication Service.
2. In the console tree, click **RADIUS Clients**.
3. In the details pane, double-click the client for which you want to edit the configuration.
4. To enable or disable use of the Message Authenticator attribute, click **Request must contain the Message Authenticator attribute**.
5. In **Shared secret**, type a shared secret. In **Confirm shared secret**, retype the shared secret.
6. Click **OK** to save your changes.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- When you enable the Message Authenticator attribute, the entire RADIUS message is encrypted and the shared secret is used as the key. Shared secrets are case-sensitive. Verify that the client's shared secret and the shared secret that you type in **Shared secret** are identical. For more information, see Related Topics.
- Enabling the use of the Message Authenticator attribute provides additional security when PAP, CHAP, MS-CHAP, and MS-CHAP v2 are used for authentication. EAP uses the Message Authenticator attribute by default and does not require that you enable it. For more information, see the topics listed below.
- Shared secrets should be comprised of a 22-character or longer random sequence of letters, numbers, and punctuation and should be changed often to protect your IAS server and your RADIUS clients from online dictionary attacks. For more information, see the topics listed below.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared secrets
Message Authenticator attribute
IAS as a RADIUS server design considerations
Elements of a remote access policy

# Configure Remote Access Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure remote access policies**

- Add a remote access policy
- Configure a remote access policy to grant or deny access
- Configure a condition of a remote access policy
- Change the policy evaluation order
- Rename a remote access policy
- Delete a remote access policy
- Recreate the default remote access policy
- Configure a Profile for a Remote Access Policy

# Add a remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a remote access policy**

1. Do one of the following:

    - Open Routing and Remote Access and, if necessary, double-click **Routing and Remote Access** and the server name.
    - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, right-click **Remote Access Policies**, and then click **New Remote Access Policy**.
3. Use the New Remote Access Policy Wizard to create a policy.

For a common policy, you must choose the following:

- An access method:
    - VPN access
    - Dial-up access
    - Wireless access
    - Ethernet
- Whether to grant access permissions by user or by group.
- Authentication methods.
- Levels of allowed encryption (depending on the access method selected).

  For a custom policy, you must configure the following:

- A set of policy conditions.
- Whether remote access permission for the policy is granted or denied.
- Remote access policy profile settings.

**Notes**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

- You cannot use this procedure to configure a remote access policy unless a remote access server is already installed on your network. After installing one or more remote access servers, use this procedure and the New Remote Access Policy Wizard to configure a new remote access policy.

- When you use the New Remote Access Policy Wizard to create a common policy:

    - The access method is used to automatically configure the NAS Port Type condition.

    - If you choose to grant access by group, the Windows Groups condition is automatically set to the selected groups.

    - The authentication settings are used to set the profile settings on the **Authentication** tab.

    - The encryption levels are used to set the profile settings on the **Encryption** tab.

    - The remote access permission is always set to **Grant remote access permission**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure a condition of a remote access policy
Configure a remote access policy to grant or deny access
Configure a Profile for a Remote Access Policy
Delete a remote access policy
Rename a remote access policy

# Configure a remote access policy to grant or deny access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a remote access policy to grant or deny access**

1. Do one of the following:

   - Open Routing and Remote Access and, if necessary, double-click **Routing and Remote Access** and the server name.
   - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, click **Remote Access Policies**.
3. In the details pane, double-click the policy that you want to configure.
4. Under **If a connection request matches the specified conditions**:

   - To grant dial-up permission for these users, click **Grant remote access permission**.
   - To deny dial-up permission for these users, click **Deny remote access permission**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- The **Remote Access Permission** setting on the dial-in properties of the user or computer account overrides this setting if it is set to either **Allow access** or **Deny access**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure a condition of a remote access policy
Change the policy evaluation order
Configure a Profile for a Remote Access Policy
Delete a remote access policy
Rename a remote access policy

# Configure a condition of a remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a condition of a remote access policy**

1. Do one of the following:
   - Open Routing and Remote Access, double-click **Routing and Remote Access**, and then double-click the server name to expand the console tree.
   - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.
2. In the console tree, click **Remote Access Policies**.
3. In the details pane, double-click the policy that you want to configure.
4. To configure a new condition, click **Add**, and then do the following:
   - In the **Select Attribute** dialog box, click the attribute to add, and then click **Add**.
   - In the attribute dialog box, configure the attribute, and then click **OK**.
5. To change the configuration of an existing condition:
   - Click the condition, and then click **Edit**.
   - In the attribute dialog box, specify the required settings, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the policy evaluation order
Configure a remote access policy to grant or deny access
Configure a Profile for a Remote Access Policy
Delete a remote access policy

Rename a remote access policy

# Change the policy evaluation order

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the policy evaluation order**

1. Do one of the following:

    - Open Routing and Remote Access and, if necessary, double-click **Routing and Remote Access** and the server name.

    - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, click **Remote Access Policies**.

3. In the details pane, right-click the policy that you want to move, and then click either **Move Up** (to move the policy up one level) or **Move Down** (to move the policy down one level).

**Notes**

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure a condition of a remote access policy
Configure a remote access policy to grant or deny access
Configure a Profile for a Remote Access Policy
Delete a remote access policy
Rename a remote access policy

# Rename a remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a remote access policy**

1. Do one of the following:

    - Open Routing and Remote Access and, if necessary, expand **Routing and Remote Access** and the server name.
    - Open Internet Authentication Service and, if necessary, expand **Internet Authentication Service**.

2. In the console tree, click **Remote Access Policies**.
3. In the details pane, right-click the policy that you want to rename, and then click **Rename**.
4. Type a new name for the policy, and then press ENTER.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure a condition of a remote access policy
Configure a remote access policy to grant or deny access
Configure a Profile for a Remote Access Policy
Delete a remote access policy
Change the policy evaluation order

# Delete a remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a remote access policy**

1. Do one of the following:

    - Open Routing and Remote Access and, if necessary, double-click **Routing and Remote Access** and the server name.

    - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, click **Remote Access Policies**.

3. In the details pane, right-click the policy that you want to delete, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure a condition of a remote access policy
Configure a remote access policy to grant or deny access
Configure a Profile for a Remote Access Policy
Change the policy evaluation order
Rename a remote access policy

# Recreate the default remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To recreate the default remote access policy**

1. Do one of the following:

   - Open Routing and Remote Access and, if necessary, double-click both **Routing and Remote Access** and the server name.
   - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. Right-click **Remote Access Policies** and then click **New Remote Access Policy**.

3. In the Add Remote Access Policy Wizard, create a custom remote access policy with the following settings:

   - In **Policy name**, type a default name, and then click **Next**.
   - In **Conditions**, click **Add**, click **Day-and-Time-Restrictions**, click **Add**, and then configure the attribute to permit all times on all days.
   - Click **OK**, and then click **Next**.
   - In **Permissions**, click **Deny remote access permission**, and then click **Next**.
   - In **Profile**, click **Next**. Do not make any changes to the profile.
   - Click **Finish**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Introduction to remote access policies
Add a remote access policy

# Configure a Profile for a Remote Access Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure a profile for a remote access policy**

- Configure dial-in constraints
- Configure IP options
- Configure multilink options
- Configure authentication
- Configure PEAP and EAP methods
- Configure encryption
- Configure Attributes for a Profile

# Configure dial-in constraints

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To configure dial-in constraints

1. Do one of the following:

    - Open Routing and Remote Access, double-click **Routing and Remote Access**, and then double-click the server name to expand the console tree.

    - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, click **Remote Access Policies**.

3. In the details pane, double-click the remote access policy that you want to configure.

4. Click **Edit Profile**.

5. On the **Dial-in Constraints** tab, specify any required settings.

### Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

- Dial-in constraints should work with most network access servers, however, some network access servers might not support the RADIUS attributes that are required to enable dial-in constraints. For more information, see your NAS documentation.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure IP options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To configure IP options

1. Do one of the following:

    - Open Routing and Remote Access and, if necessary, double-click **Routing and Remote Access** and the server name.

    - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, click **Remote Access Policies**.

3. In the details pane, double-click the remote access policy that you want to configure.

4. Click **Edit Profile**.

5. On the **IP** tab, specify any required settings.

### Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure multilink options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To configure multilink options

1. Do one of the following:

    - Open Routing and Remote Access and, if necessary, double-click **Routing and Remote Access** and the server name.

    - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, click **Remote Access Policies**.

3. In the details pane, double-click the remote access policy that you want to configure.

4. Click **Edit Profile**.

5. On the **Multilink** tab, specify any required settings.

### Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure authentication**

1. Do one of the following:

   - Open Routing and Remote Access and, if necessary, double-click **Routing and Remote Access** and the server name.

   - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, click **Remote Access Policies**.

3. In the details pane, double-click the policy that you want to configure.

4. Click **Edit Profile**.

5. On the **Authentication** tab, specify any required settings.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

- When using the EAP-TLS authentication method with certificates, TLS uses cached certificate properties instead of reading the certificate from the certificate store. If a certificate is either changed or deleted and replaced by a new certificate, TLS continues using outdated cached certificate information until the cache expires or is refreshed. If you change or replace a certificate, you can refresh the TLS cache by restarting the server computer.

- For information about configuring specific authentication protocols, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
EAP
MS-CHAP
CHAP
PAP
Unauthenticated access

# Configure PEAP and EAP methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure PEAP and EAP methods**

1. Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.
2. In the console tree, click **Remote Access Policies**.
3. In the details pane, double-click the policy that you want to configure.
4. Click **Edit Profile**.
5. On the **Authentication** tab, click **EAP Methods**.
6. In **Select EAP providers**, click **Add**. Select the authentication methods that you want to use, and then click **OK**.
7. In **Select EAP providers**, click the EAP type that you want to configure, and then click **Edit**. Depending on the EAP type selected, one of the following dialog boxes is displayed:
   - If **Protected EAP (PEAP)** is selected, the **Protected EAP Properties** dialog box opens. In **Certificate Issued**, select the certificate that the server uses to identify itself to client computers. To enable PEAP fast reconnect for 802.11 wireless client computers, click **Enable Fast Reconnect**. Secure password user authentication with EAP-MSCHAPv2 is the default in **EAP Types**. To configure EAP-MSCHAPv2 properties, click **Edit**. To configure certificate or smart card user authentication click **Add**. In **Authentication methods,** click **Smart Card or other certificate**, and then click **OK**.
   - If **Smart Card or other Certificate Properties** is selected, the **Smart Card or other Certificate Properties** dialog box opens. In **Certificate issued to**, select the certificate that the server uses to authenticate to client computers.
8. In **Select EAP providers**, click **Move Up** or **Move Down** to specify the negotiation order of EAP methods. The server starts negotiation with the client according to the order specified in **EAP types**.

**Notes**

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- If the EAP types displayed in the **Select EAP Providers** dialog box are selected as **(Server-Configured)**, then the configuration of the EAP type is common to the server and applies to all remote access policies that are configured on the server. Otherwise, the EAP type configuration is specific to the current remote access policy, in which case the EAP method must be configured for each policy to which you want it to apply.
- PEAP is used as an authentication method for 802.11 wireless client computers, but is not supported for virtual private network (VPN) or other remote access clients. Because of this, you can only configure PEAP as the authentication method for a remote access policy when you are using IAS.
- When using the PEAP-EAP-TLS and EAP-TLS authentication methods with certificates, TLS uses cached certificate properties instead of reading the certificate from the certificate store. If a certificate is either changed or deleted and

replaced by a new certificate, TLS continues using outdated cached certificate information until the cache expires or is refreshed. If you change or replace a certificate, you can refresh the TLS cache by restarting the server computer.

- EAP-TLS is not supported if the IAS or remote access server is not a member of the domain.

- If you have issued a certificate to your IAS server that has a blank Subject, the certificate is not available to authenticate your IAS server. To change this, you can use Certificate Templates to create a new certificate for enrollment on your IAS server. In the certificate properties, on the **Subject Name** tab, in **Subject name format**, select a value other than **None**.

- For information about configuring specific authentication protocols, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
EAP
MS-CHAP
CHAP
PAP
Unauthenticated access
PEAP
Network access authentication and certificates

# Configure encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To configure encryption

1. Do one of the following:

    - Open Routing and Remote Access and, if necessary, double-click **Routing and Remote Access** and the server name.

    - Open Internet Authentication Service and, if necessary, double-click **Internet Authentication Service**.

2. In the console tree, click **Remote Access Policies**.

3. In the details pane, double-click the remote access policy that you want to configure.

4. Click **Edit Profile**.

5. On the **Encryption** tab, specify any required settings.

### Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure Attributes for a Profile

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure attributes for a profile**

- Add RADIUS attributes to a remote access policy
- Configure vendor-specific attributes for a remote access policy

# Add RADIUS attributes to a remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add RADIUS attributes to a remote access policy**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote Access Policies**.
3. Right-click the policy for which you want to configure a RADIUS attribute, and then click **Properties**.
4. Click **Edit Profile**, click the **Advanced** tab, and then click **Add**.
5. In the list of available RADIUS attributes, double-click the attribute that you want to add to the profile. If the **Multivalued Attribute Information** dialog box appears, click **Add**.
6. In **Attribute**, supply a value from the following table, and then click **Add**. If you want to add another value to the attribute, click **Add** again, and then configure the attribute.

| If the attribute format is | In Attribute value | Attributes in this format include |
|---|---|---|
| Integer | Type the value specified in your access server documentation. | Acct-Interim-Interval |
| Enumerator | Select the value from the list. | Framed-Protocol |
| InetAddr | Type the IP address to be specified in the attribute. | Framed-IP-Netmask |
| Octet String | Type the string or the hexadecimal value that you provide. | Tunnel-Password |
| String | Type the string that you provide. | Reply-Message |
| Boolean | Select either **True** or **False**. | Ignore-User-Dialin-Properties |

**Notes**

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- You can only use the Generate-Session-Timeout attribute if your user account database is either a Security Accounts Manager (SAM) database or the user account database for an Active Directory domain. When the value of **Generate-Session-Timeout** is set to **True**, the ForceLogoff value for a SAM database should be set to **0** (zero). In Local Security Settings, ForceLogoff is set to **0** when **Network security: Force logoff when logon hours expire** is enabled. For more information, see Edit security settings on a Group Policy object. For information about changing the ForceLogoff value for the SAM database at the command prompt, see Net accounts.
- You can configure wireless connection policy so that wireless clients periodically reauthenticate. This ensures that the client Wired Equivalent Privacy (WEP) encryption keys are changed often enough to provide adequate security for the wireless connection. To configure reauthentication, set the session timeout in your remote access policy or connection request policy for wireless connections (using the Session-Timeout attribute) to the interval you prefer (for example, 10 minutes). Additionally, configure the Termination-Action attribute with **Attribute value** set to **RADIUS-Request**. If the Termination-Action attribute is not set to **RADIUS-Request**, wireless access points might end the connection during reauthentication. For more information, see your hardware documentation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Elements of a remote access policy
Vendor-specific attribute overview
Configure vendor-specific attributes for a remote access policy

# Configure vendor-specific attributes for a remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure vendor-specific attributes for a remote access policy**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote Access Policies**.
3. Right-click the policy for which you want to configure a vendor-specific attribute (VSA), and then click **Properties**.
4. Click **Edit Profile**, click the **Advanced** tab, and then click **Add**.
5. Look at the list of available RADIUS attributes to determine whether your vendor-specific attribute is already in it. If it is, double-click it, and then configure it as specified in your NAS documentation.

If the vendor-specific attribute is not in the list of available RADIUS attributes:

1. Double-click **Vendor-Specific**, and then click **Add**.
2. Specify the vendor for your network access server.

   To specify the NAS vendor by selecting the name from the list, click **Select from list**, and then select the vendor of the NAS for which you are configuring the VSA. If you are using Routing and Remote Access as your NAS, click **Microsoft**.

   If the vendor is not listed, you can specify the vendor by typing the vendor code. To do this, click **Enter Vendor Code**, and then type the vendor code in the space provided. If you do not know the vendor code for your NAS, see IETF RFC 1007 for a list of SMI Network Management Private Enterprise Codes.
3. Specify whether the attribute conforms to the RADIUS RFC VSA specification. If you are not sure, see your NAS documentation.
4. If your attribute conforms, click **Yes. It conforms**, and then click **Configure Attribute**.
   - In **Vendor-assigned attribute number**, type the number that is assigned to the attribute (the numbers available are 0 through 255).
   - In **Attribute format**, specify the format for the attribute, and then in **Attribute value**, type the value that you are assigning to the attribute.
5. If the attribute does not conform, click **No. It does not conform**, and then click **Configure Attribute**. In **Hexadecimal attribute value**, type the value for the attribute.

**Notes**

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- To see configuration examples for vendor-specific attributes that both conform and do not conform to the RADIUS RFC, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Vendor-specific attribute overview
Sample VSA for a Cisco NAS
Sample VSA for a US Robotics NAS

# Configure Connection Request Processing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure connection request processing**

- Configure Connection Request Policies
- Configure Remote RADIUS Server Groups

# Configure Connection Request Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure connection request policies**

- Add a connection request policy
- Configure a condition of a connection request policy
- Change the policy evaluation order
- Delete a connection request policy
- Configure a Profile for a Connection Request Policy

# Add a connection request policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a connection request policy**

1. Open Internet Authentication Service.
2. In the console tree, right-click **Connection Request Policies**, and then click **New Connection Request Policy**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Connection Request Policies

3. Use the New Connection Request Policy Wizard to configure your connection request policy.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- You are not required to configure a remote RADIUS server group before you create a connection request policy that forwards authentication or accounting messages. You can create a remote RADIUS server group in the New Connection Request Policy Wizard.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Connection request policies

# Configure a condition of a connection request policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a condition of a connection request policy**

1. Open Internet Authentication Service.
2. In the console tree, click **Connection Request Policies**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Connection Request Policies

3. In the details pane, double-click the policy that you want to configure.
4. To configure a new condition, click **Add**, and then do the following:

   - In the **Select Attribute** dialog box, click the attribute to add, and then click **Add**.
   - In the attribute dialog box, configure the attribute, and then click **OK**.

5. To change the configuration of an existing condition:

   - Click the condition, and then click **Edit**.
   - In the attribute dialog box, specify your required settings, and then click **OK**.

**Note**

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Connection request policies

# Change the connection request policy evaluation order

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the policy evaluation order**

1. Open Internet Authentication Service.
2. In the console tree, click **Connection Request Policies**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Connection Request Policies

3. In the details pane, right-click the policy that you want to move, and then click either **Move Up** (to move the policy up one level) or **Move Down** (to move the policy down one level).

**Note**

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Connection request policies

# Delete a connection request policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a connection request policy**

1. Open Internet Authentication Service.
2. In the console tree, click **Connection Request Policies**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Connection Request Policies

3. In the details pane, right-click the policy that you want to delete, and then click **Delete**.

### Note

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Connection request policies

# Configure a Profile for a Connection Request Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure a profile for a connection request policy**

- Configure authentication
- Configure accounting
- Configure attribute manipulation
- Configure advanced attributes

# Configure authentication in a connection request policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure authentication**

1. Open Internet Authentication Service.
2. In the console tree, click **Connection Request Policies**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Connection Request Policies

3. In the details pane, right-click the policy that you want to configure, and then click **Properties**.
4. Click **Edit Profile**.
5. On the **Authentication** tab, specify authentication settings.

**Note**

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Connection request policies

# Configure accounting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure accounting**

1. Open Internet Authentication Service.
2. In the console tree, click **Connection Request Policies**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Connection Request Policies

3. In the details pane, right-click the policy that you want to configure, and then click **Properties**.
4. Click **Edit Profile**.
5. On the **Accounting** tab, specify accounting settings.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- When you have remote RADIUS server groups configured and, in RADIUS proxy server **Connection Request Policies**, you clear the **Record accounting information on the servers in the following remote RADIUS server group** check box, these groups are still sent network access server (NAS) start and stop notification messages. This creates unnecessary network traffic. To eliminate this traffic, disable NAS notification forwarding for individual servers in each remote RADIUS server group by clearing the **Forward network start and stop notifications to this server** check box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Connection request policies
Configure the authentication and accounting settings of a group member

# Configure attribute manipulation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure attribute manipulation**

1. Open Internet Authentication Service.
2. In the console tree, click **Connection Request Policies**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Connection Request Policies

3. In the details pane, right-click the policy that you want to configure, and then click **Properties**.
4. Click **Edit Profile**.
5. On the **Attribute** tab, specify find and replace rules for RADIUS attributes.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

### Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- When the user name does not contain a domain name, IAS supplies one. By default, the domain name supplied is that of which the IAS server is a member. You can specify this domain name through the following registry setting:

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn\DefaultDomain**

- Some third party network access servers delete or modify the domain name as specified by the user. As the result, the network access request is authenticated against the default domain, which might not be the domain for the user's account. To resolve this problem, configure your RADIUS servers to change the user name into the correct format with the accurate domain name.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Connection request policies

# Configure advanced attributes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure advanced attributes**

1. Open Internet Authentication Service.
2. In the console tree, click **Connection Request Policies**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Connection Request Policies

3. In the details pane, right-click the policy that you want to configure, and then click **Properties**.
4. Click **Edit Profile**.
5. On the **Advanced** tab, specify connection attributes.

**Note**

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Connection request policies

# Configure Remote RADIUS Server Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure remote RADIUS server groups**

- Add a remote RADIUS server group
- Delete a remote RADIUS server group
- Configure the Members of a Remote RADIUS Server Group

# Add a remote RADIUS server group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a remote RADIUS server group**

1. Open Internet Authentication Service.
2. In the console tree, right-click **Remote RADIUS Server Groups**, and then click **New Remote RADIUS Server Group**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Remote RADIUS Server Groups

3. Use the New Remote RADIUS Server Group Wizard to configure your remote RADIUS server group.

### Note

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Remote RADIUS server groups

# Delete a remote RADIUS server group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a remote RADIUS server group**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote RADIUS Server Groups**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Remote RADIUS Server Groups

3. In the details pane, right-click the remote RADIUS server group that you want to delete, and then click **Delete**.

### Note

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Remote RADIUS server groups

# Configure the Members of a Remote RADIUS Server Group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure the members of a remote RADIUS server group**

- Add a member to a remote RADIUS server group
- Remove a member from a remote RADIUS server group
- Configure the name of a group member
- Configure the authentication and accounting settings of a group member
- Configure the load balancing properties of a group member

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Add a member to a remote RADIUS server group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a member to a remote RADIUS server group**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote RADIUS Server Groups**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Remote RADIUS Server Groups

3. In the details pane, right-click the remote RADIUS server group that you want to configure, and then click **Properties**.
4. Click **Add**, and then configure address, authentication, accounting, and load balancing settings.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Remote RADIUS server groups

# Remove a member from a remote RADIUS server group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a member from a remote RADIUS server group**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote RADIUS Server Groups**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Remote RADIUS Server Groups

3. In the details pane, right-click the remote RADIUS server group that you want to configure, and then click **Properties**.
4. Click the server to remove from the group and then click **Remove**.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Remote RADIUS server groups

# Configure the name of a group member

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the name of a group member**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote RADIUS Server Groups**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Remote RADIUS Server Groups

3. In the details pane, right-click the remote RADIUS server group that you want to configure, and then click **Properties**.
4. On the **General** tab, click the server that you want to configure, and then click **Properties**.
5. On the **Address** tab, type the name or IP address of the server.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Remote RADIUS server groups

# Configure the authentication and accounting settings of a group member

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the authentication and accounting settings of a group member**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote RADIUS Server Groups**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Remote RADIUS Server Groups

3. In the details pane, right-click the remote RADIUS server group that you want to configure, and then click **Properties**.
4. On the **General** tab, click the server that you want to configure, and then click **Properties**.
5. On the **Authentication/Accounting** tab, specify authentication and accounting settings.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- When you have remote RADIUS server groups configured and, in RADIUS proxy server **Connection Request Policies**, you clear the **Record accounting information on the servers in the following remote RADIUS server group** check box, these groups are still sent network access server (NAS) start and stop notification messages. This creates unnecessary network traffic. To eliminate this traffic, disable NAS notification forwarding for individual servers in each remote RADIUS server group by clearing the **Forward network start and stop notifications to this server** check box. For more information, see Configure the authentication and accounting settings of a group member and Configure accounting.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Remote RADIUS server groups
Configure accounting
IAS Best Practices

# Configure the load balancing properties of a group member

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the load balancing properties of a group member**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote RADIUS Server Groups**.

   **Where?**

   - Internet Authentication Service (Local)/Connection Request Processing/Remote RADIUS Server Groups

3. In the details pane, right-click the remote RADIUS server group that you want to configure, and then click **Properties**.
4. On the **General** tab, click the server that you want to configure, and then click **Properties**.
5. On the **Load balancing** tab, specify load balancing settings.

## Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Introduction to connection request processing
Remote RADIUS server groups

# Configure Logging for User Authentication and Accounting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure logging for user authentication and accounting**

- Select requests to be logged
- Configure log file properties

# Select requests to be logged

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select requests to be logged**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote access logging**.
3. In the details pane, right-click **Local File** or **SQL Server**, and then click **Properties**.
4. On the **Settings** tab, select one or more check boxes for recording authentication and accounting requests in the IAS log files:
   - To capture accounting requests and responses, select **Accounting requests**.
   - To capture authentication requests, Access-Accept messages, and Access-Reject messages, select **Authentication requests**.
   - To capture periodic status updates, such as interim accounting packets, select **Periodic status**.

**Notes**

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
- It is recommended that you initially select **Accounting requests** and **Authentication requests**. You can change your selections as your requirements change.
- To use SQL Server logging, you must have a computer running an XML-compliant version of SQL Server.
- The **Authentication requests** option alerts you to both problems with transaction volume and unauthorized attempts to access resources.
- Interim accounting messages, recorded when you select **Periodic status**, are sent by the access server only when:
  1. You have added and configured the **Acct-Interim-Interval** RADIUS accounting attribute on the **Advanced** tab of the appropriate remote access policy.
  2. The access server supports the Acct-Interim-Interval attribute and the sending of interim accounting messages.
- Logging of client certificate validation failures is a secure channel (Schannel) event, and is not enabled on the IAS server by default. You can enable additional Schannel events (including logging of the client certificate validation failure) by changing the following registry key value from **1** (**REG_DWORD** type, data **0x00000001**) to **3** (**REG_DWORD** type, data **0x00000003**):

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\EventLogging**

  **Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add RADIUS attributes to a remote access policy
Vendor-specific attribute overview
Event logging for IAS
SQL Server database logging

# Configure log file properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure log file properties**

1. Open Internet Authentication Service.
2. In the console tree, click **Remote Access Logging**.
3. In the details pane, choose **Local File** or **SQL Server**, and then follow the appropriate directions below:
   - Local File
   - SQL Server

**Local File**

1. In the details pane, right-click **Local File**, and then click **Properties**.
2. On the **Log File** tab, in **Directory**, type the location where log files are to be stored. The default location is the *systemroot*\System32\LogFiles folder.
3. In **Format**, click **Database-compatible**. Or, to keep your log files in IAS format, click **IAS**.
4. To open new log files at specific intervals, click the interval that you want to use:
   - For heavy transaction volume and logging activity, click **Daily**.
   - For lesser transaction volumes and logging activity, click **Weekly** or **Monthly**.
   - To store all transactions in one log file, click **Never (unlimited file size)**.
   - To limit the size of each log file, click **When log file reaches this size**, and then type a file size--after which a new log is created. The default size is 10 MB.
5. To automatically delete log files when the disk is full, click **When disk is full delete older log files.** If the oldest log file is the current log file, it is not deleted.

**SQL Server**

1. In the details pane, right-click **SQL Server**, and then click **Properties**.
2. In **Log the following information**, select the information that you want to log:
   - To log all accounting requests, click **Accounting requests**.
   - To log authentication requests, click **Authentication requests**.
   - To log periodic status, such as interim accounting requests, click **Periodic status**.

3. To configure the number of concurrent sessions allowed between the IAS server and the SQL Server database, type a number in **Maximum number of concurrent sessions**.

4. To configure the SQL Server data source, click **Configure**, and then, on the **Connection** tab, specify the following:

   - To specify the name of the server on which the database is stored, type or select a name in **Select or enter a server name**.

   - To specify the authentication method with which to log on to the server, click **Use Windows NT integrated security**. Or, click **Use a specific user name and password**, and then type credentials in **User name** and **Password**.

   - To allow a blank password, click **Blank password**.

   - To store the password, click **Allow saving password**.

   - To specify which database to connect to on the computer running SQL Server, click **Select the database on the server**, and then select a database name from the list.

5. To test the connection between the IAS server and the computer running SQL Server, click **Test Connection**.

### Notes

- To open Internet Authentication Service, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.

- To prevent the log files from filling the hard drive on which you are storing them, it is strongly recommended that you keep them on a partition that is separate from the system partition.

- To send the log file data for collection by another process, you can configure IAS to write to a named pipe. To use named pipes, set the log file folder to **\\.\pipe** or *****ComputerName***\pipe**. The named pipe server program creates a named pipe called **\\.\pipe\iaslog.log** to accept the data. In the **Local file properties** dialog box, in **Create a new log file**, select **Never (unlimited file size)** when you use named pipes.

- The log file directory can be created by using system environment variables (instead of user variables), such as %systemdrive%, %systemroot%, and %windir%. For example, the following path, using the environment variable %windir%, locates the log file at the system directory in the subfolder \System32\Logs (that is, %windir%\System32\Logs\).

- Switching log file formats does not cause a new log to be created. If you change log file formats, the file that is active at the time of the change will contain a mixture of the two formats (records at the start of the log will have the previous format, and records at the end of the log will have the new format).

- If you are administering an IAS server remotely, you cannot browse the directory structure. If you need to log accounting information to a remote server, specify the log file name by typing a Universal Naming Convention (UNC) name, such as *MyLogServer**LogShare*.

- If you do not supply a full path statement in **Log File Directory**, the default path is used. For example, if you type **IASLogFile** in **Log File Directory**, the file is located at *systemroot*\System32\IASLogFile.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Disk quotas overview
Remote Access Logging
Logging user authentication and accounting requests
SQL Server database logging
Command shell overview

# IAS Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- IAS Overview
- Understanding IAS
- Deploying IAS
- IAS Resources

# IAS Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS overview**

This section covers:

- Security information for IAS
- Introducing IAS
- Features of IAS
- RADIUS protocol
- New features for IAS
- Differences in IAS implementations
- How IAS Works

# Security information for IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for IAS**

It is important to follow best practices for security when using IAS servers on your network. For more information, see Best practices for security.

Following is security information for IAS and related protocols:

- **Deployments of PEAP and EAP unprotected by PEAP should not use the same authentication type**

  When you deploy both PEAP and EAP unprotected by PEAP, do not use the same EAP authentication type with and without PEAP. For example, if you deploy PEAP with EAP-TLS (PEAP-EAP-TLS), do not also deploy EAP-TLS without PEAP. Deploying authentication methods with the same type -- one with and the other without the protection of PEAP -- creates a security vulnerability.

- **The RADIUS User-Password hiding mechanism might not provide sufficient security for passwords.**

  The RADIUS hiding mechanism uses the RADIUS shared secret, the Request Authenticator, and the use of the MD5 hashing algorithm to encrypt the User-Password and other attributes, such as Tunnel-Password and MS-CHAP-MPPE-Keys. RFC 2865 notes the potential need for evaluating the threat environment and determining whether additional security should be used.

  You can provide additional protection for hidden attributes by using Internet Protocol Security (IPSec) with Encapsulating Security Payload (ESP) and an encryption algorithm, such as Triple DES (3DES), to provide data confidentiality for the entire RADIUS message.

  **Recommendations:**

  - Use IPSec to provide additional security for RADIUS clients and servers. For more information, see Securing RADIUS traffic with IPSec.
  - If you are using IPSec with ESP and an encryption algorithm is not available, do the following:

1. Require the use of the Message Authenticator attribute for all Access-Request messages. For more information, see Message Authenticator attribute and Configure the Message Authenticator attribute and shared secret.
2. Use cryptographically strong Request Authenticators.
3. Require the use of strong user passwords.
4. Use authentication counting and account lockout to help prevent a dictionary attack against a user password.
5. Use a long shared secret with a random sequence of letters, numbers, and punctuation. Change it often to help protect your IAS server and your RADIUS clients from dictionary attacks. For more information, see Shared secrets.

- **Many authentication methods are too weak to provide adequate security for most organizations.**

  Many organizations use authentication methods that expose their network to a variety of security attacks. The most secure method of authentication is Extensible Authentication Protocol-Transport Level Security (EAP-TLS) when used in conjunction with smart cards. However, EAP-TLS and smart cards require a public key infrastructure (PKI), which can be complicated to deploy.

  **Recommendations:**

  - When you use password-based authentication, enforce strong password policies on your network to make dictionary attacks more difficult.
  - Make sure that Password Authentication Protocol (PAP), Shiva Password Authentication Protocol (SPAP), and Challenge Handshake Authentication Protocol (CHAP) are disabled. PAP, SPAP, and CHAP are disabled by default on the profile of a remote access policy.
  - Disable Microsoft Challenge Handshake Authentication Protocol (MS-CHAP). MS-CHAP is enabled by default on the profile of a remote access policy. Before you disable MS-CHAP, ensure that all of the clients of your access server are capable of using MS-CHAP version 2 (MS-CHAP v2). If your server is not running Windows Server 2003 , you will need the latest service packs and dial-up networking updates in order to support MS-CHAP v2.
  - If you do not disable MS-CHAP, disable Microsoft LAN Manager authentication for MS-CHAP. LAN Manager authentication for MS-CHAP is enabled by default on clients running earlier versions of Windows.

    To disable LAN Manager authentication, set the registry value **Allow LM Authentication** to **0** at the following registry key:

    **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Policy**

**Caution**

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

- **Remote access account lockout can cause authorized users to be denied network access.**

  If a malicious user attempts a dictionary attack with an authorized user logon name when remote access account lockout is enabled, both the malicious user and the authorized user are locked out of the account until the configured **Account lockout threshold** is reached.

  **Recommendations:**

  - When the lockout count for a user account is reset to zero, due to either a successful authentication or an automatic reset, the registry subkey for the user account is deleted. To manually unlock a user account (before the **Account lockout threshold** is automatically reset), you can delete the following registry subkey that corresponds to the user's account name:

    **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout***DomainName***:***UserName*

    **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the **Last Known Good Configuration** startup option if problems are encountered after manual changes have been applied.

  **Notes**

  - Remote access account lockout is not related to the **Account locked out** setting on the **Account** tab in the properties of a user account.

  - With remote access account lockout, there is no distinction between malicious users who attempt to access your intranet and authentic users who attempt remote access but have forgotten or repeatedly mistyped their current passwords. Depending on the number of attempts and the MaxDenials setting in the registry, these users might have their accounts locked out.

- **Moving IAS log files from the default location results in unprotected log files.**

  When IAS logging is enabled, log files are located by default in the *systemroot*\System32\LogFiles folder. By default, the access control list on the LogFiles folder provides the best security for the IAS log files. The access control list is a list of users and groups that can access the folder. In addition, each user or group is assigned specific permissions that determine what actions the user or group can take with the folder. For more information, see Access Control.

  If you change the location of the IAS log files, however, the new folder location is not protected with the usual access control list by default. This creates a security risk.

  **Recommendations:**

  - If you move the IAS log files to another folder, create an access control list that contains only the Local System account and the Administrators group. In addition, ensure that the new log file folder does not inherit permissions from any other folder.

  - Do not store IAS log files in any personal folders, in My Documents, or in any subfolder of My Documents.

  - Store IAS log files on a computer that is physically secured and protected from unauthorized access.

  - Use file-level auditing for individual IAS log files to track the users and groups that access them. For more information, see Apply or modify auditing policy settings for a local file or folder.

  - Do not move or map the IAS log files to a remote storage device or an external drive for which no access control list can be created. Also do not move or copy IAS log files to an anonymous FTP site or any other unprotected location.

- **IAS server log data is sent in plaintext to remote computers running SQL server**

  If your IAS server is configured to log user authentication and accounting requests to a remote computer running SQL Server, the log data is sent unencrypted over the network, posing a high security risk.

  **Recommendations:**

  - Establish an encrypted connection at the computer running SQL Server for communication with the IAS server. For more information, see the SQL Server documentation.

  For more information about logging IAS data to a SQL Server database, see SQL Server database logging and Configure log file properties.

For additional information about security considerations and best practices, consult the following resources:

- Securing IAS
- IAS as a RADIUS server security considerations
- IAS Best Practices
- IAS as a RADIUS proxy security considerations
- Remote Access Logging
- Event Viewer
- Security issues for VPN
- Security information for remote access
- Security information for routing

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Introducing IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introducing IAS**

Internet Authentication Service (IAS) in Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition is the Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy. As a RADIUS server, IAS performs centralized connection authentication, authorization, and accounting for many types of network access, including wireless, authenticating switch, dial-up and virtual private network (VPN) remote access, and router-to-router connections. As a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers. IAS supports the Internet Engineering Task Force (IETF) standards for RADIUS described in RFCs 2865 and 2866.

IAS enables the use of a heterogeneous set of wireless, switch, remote access, or VPN equipment. You can use IAS with the Routing and Remote Access service, which is available in Microsoft Windows 2000, Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

When an IAS server is a member of an Active Directory® domain, IAS uses the directory service as its user account database and is part of a single sign-on solution. The same set of credentials is used for network access control (authenticating and authorizing access to a network) and to log on to an Active Directory domain.

Internet service providers (ISPs) and organizations that maintain network access have the increased challenge of managing all types of network access from a single point of administration, regardless of the type of network access equipment used. The RADIUS standard supports this functionality in both homogeneous and heterogeneous environments. RADIUS is a client-server protocol that enables network access equipment (used as RADIUS clients) to submit authentication and accounting requests to a RADIUS server.

A RADIUS server has access to user account information and can check network access authentication credentials. If the user's credentials are authentic and the connection attempt is authorized, the RADIUS server authorizes the user's access on the basis of specified conditions and logs the network access connection in an accounting log. The use of RADIUS allows the network access user authentication, authorization, and accounting data to be collected and maintained in a central location, rather than on each access server.

With IAS, organizations can also outsource remote access infrastructure to a service provider while retaining control over user authentication, authorization, and accounting.

Different IAS configurations can be created for the following solutions:

- Wireless access
- Organization dial-up or virtual private network (VPN) remote access
- Outsourced dial-up or wireless access
- Internet access
- Authenticated access to extranet resources for business partners

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Features of IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Features of IAS**

IAS includes support for the following features:

- A variety of authentication methods

  IAS supports a number of authentication protocols and allows you to add custom methods that meet your authentication requirements. The supported authentication methods are:

  - Password-based Point-to-Point Protocol (PPP) authentication protocols

    Password-based PPP authentication protocols, such as Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), and MS-CHAP version 2 (MS-CHAP v2), are supported. For more information, see Authentication methods.

  - Extensible Authentication Protocol (EAP)

    An Internet standards-based infrastructure that allows the addition of arbitrary authentication methods, such as smart cards, certificates, one-time passwords, and token cards. A specific authentication method that uses the EAP infrastructure is an EAP type. IAS includes support for EAP-Message Digest 5 (MD5) and EAP-Transport Level Security (EAP-TLS). For more information, see EAP.

- A variety of authorization methods

  IAS supports a number of authorization methods and allows you to add custom methods that meet your authorization requirements. The supported authorization methods are:

  - Dialed Number Identification Service (DNIS)

    The authorization of a connection attempt that is based on the number called. DNIS supplies the number that was called to the call receiver and is provided by most standard telephone companies.

  - Automatic Number Identification/Calling Line Identification (ANI/CLI)

    The authorization of a connection attempt that is based on the phone number of the caller. ANI/CLI service supplies the number of the caller to the call receiver and is provided by most standard telephone companies.

  - Guest authorization

    The use of the Guest account as the identity of the user when the connection is made without user credentials (user name and password).

    For more information, see Unauthenticated access.

- Heterogeneous access servers

  IAS supports access servers that support RADIUS RFCs 2865 and 2866. In addition to dial-up access servers (also known as network access servers), IAS also supports the following:

  - Wireless access points

    By using remote access policies and the Wireless-IEEE 802.11 port type condition, IAS can be used as a RADIUS server for wireless access points that use RADIUS for authentication and authorization of wireless nodes. For

more information, see New features for IAS.

- Authenticating switches

  By using remote access policies and the Ethernet port type condition, IAS can be used as a RADIUS server for Ethernet network switches that use RADIUS for authentication and authorization of switch nodes. For more information, see New features for IAS.

- Integration with the Routing and Remote Access service

  Both IAS and the Routing and Remote Access service share remote access policies and log file capabilities. This integration provides a consistent implementation for both IAS and the Routing and Remote Access service. It enables you to deploy Routing and Remote Access in small sites without the requirement of a separate, centralized IAS server. It also provides the ability to scale to a centralized remote access management model, whenever you have multiple Routing and Remote Access servers in your organization. IAS, in conjunction with Routing and Remote Access servers, implements a single point of administration for remote access to your network for outsourced-dial, demand-dial, and VPN access. The policies within IAS at a central and large site can be exported to the independent Routing and Remote Access server in a small site.

- RADIUS proxy

  IAS allows an incoming RADIUS request to be forwarded to another RADIUS server for the processing of authentication and authorization, or accounting. As a RADIUS proxy, IAS can be used whenever the RADIUS request must be routed to another RADIUS server. IAS can forward requests based on user name, access server IP address, access server identifier, and other conditions. For more information, see New features for IAS.

- Outsourced dial-up and wireless network access

  Outsourced dialing (also known as wholesale dialing) provides a contract between an organization (the customer) and an Internet service provider (ISP). The ISP allows organization employees to connect to its network before the VPN tunnel to the organization’s private network is established. When an organization employee connects to the ISP's NAS, the authentication and usage records are forwarded to the IAS server at the organization. The IAS server enables the organization to control user authentication, track usage, and determine which employees are allowed to access the ISP's network.

  The advantage of outsourced dialing is the potential savings. By using an ISP's routers, network access servers, and wide area network (WAN) links instead of purchasing your own, you might save a great deal on hardware (infrastructure) costs. By dialing into the ISP with worldwide connections, you might significantly decrease your long-distance phone bill. And, by moving support requirements to the provider, you might eliminate administrative costs.

  You can also outsource wireless access. A vendor can provide wireless access in a remote location and use your user name to forward the connection request to a RADIUS server that is under your control for authentication and authorization. A good example is wireless Internet access in an airport.

- Centralized user authentication and authorization

  To authenticate a connection request, IAS validates the connection credentials against user accounts in the local Security Accounts Manager (SAM), a Microsoft® Windows NT® Server 4.0 domain, or an Active Directory® domain. For an Active Directory domain, IAS supports the use of Active Directory user principal names (UPNs) and universal groups.

  To authorize a connection request, IAS uses the dial-in properties of the user account that correspond to both the connection credentials and remote access policies. Although it is relatively easy to manage remote access permission for each user account, this approach does not scale well as an organization grows. Remote access policies provide a more powerful and flexible way to manage remote access permission. You can authorize network access based on various conditions, including:

  - User account membership in a group.

- The time of day or day of the week.
- The type of media through which the user is connecting (for example, wireless, Ethernet switch, modem, or VPN).
- The phone number that the user calls.
- The access server from which the request arrives.

  By configuring profile settings on remote access policies, you can control many connection parameters, including:

- The use of specific authentication methods.
- The idle timeout.
- The maximum time of a single session.
- The number of links in a multilink session.
- The use of encryption and its strength.
- The use of packet filters to control what the remote access user can access when connected to the network. For example, you can use filters to control which IP addresses, hosts, and ports the user is allowed to use in sending or receiving packets.
- The creation of a compulsory tunnel that forces all packets from that connection to be securely tunneled through the Internet and terminated in a private network.
- The virtual local area network identifier (VLAN ID) for wireless or Ethernet connections.

- Centralized administration for all of your access servers

  Support for the RADIUS standard allows IAS to control connection parameters for any access server that implements RADIUS. The RADIUS standard also allows individual remote access vendors to create proprietary extensions called vendor-specific attributes (VSAs). IAS has incorporated the extensions from a number of vendors in its dictionary. Additional VSAs can be added to the profile of individual remote access policies. For more information, see Vendor-specific attribute overview.

- Centralized auditing and usage accounting

  Support for the RADIUS standard allows IAS to collect, at a central location, the usage (accounting) records sent by all access servers. IAS stores audit information (for example, authentication accepts and rejects) and usage information (for example, connect and disconnect records) in log files. IAS supports a log file format that can be directly imported into a database. The data can then be analyzed with any standard data analysis application. For more information, see Logging user authentication and accounting requests.

- Snap-in-based management tool

  IAS provides an administration tool named the Internet Authentication Service snap-in. You can run Internet Authentication Service from Administrative Tools to administer IAS on a local computer. Or, you can add the Internet Authentication Service snap-in to the Microsoft Management Console (MMC) to administer IAS running on either a local computer or a remote computer.

- Local or remote monitoring with tools provided in Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition

  You can monitor IAS locally or on a remote computer with tools provided in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, including Event Viewer, System

Monitor, and Simple Network Management Protocol (SNMP). You can also use Network Monitor to capture RADIUS messages for detailed traffic analysis and troubleshooting.

- Scalability

  You can use IAS in a variety of network configurations of varying size, from stand-alone servers for small networks to large organization and ISP networks.

- Support for multiple IAS servers

  The synchronization of the configuration of multiple IAS servers can be performed with the Netsh command-line tool. For more information see Copy the IAS configuration to another server.

- Extensibility

  The Windows Server 2003 Platform Software Development Kit (SDK) contains two smaller networking SDKs--the IAS SDK and the EAP SDK.

  You can use the IAS SDK to:

  - Return custom attributes to the access server in addition to those returned by IAS. For example, you can create a customized module to assign IP addresses.
  - Control the number of user network sessions.
  - Import usage and audit data directly into an Open Database Connectivity (ODBC)-compliant database.
  - Create customized authorization modules.
  - Create customized authentication modules (non-EAP).

You can use the EAP SDK to create EAP types. For more information, see EAP.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# RADIUS protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**RADIUS protocol**

Remote Authentication Dial-In User Service (RADIUS) is an industry standard protocol described in RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)," and RFC 2866, "RADIUS Accounting." RADIUS is used to provide authentication, authorization, and accounting services. A RADIUS client (typically a dial-up server, VPN server, or wireless access point) sends user credentials and connection parameter information in the form of a RADIUS message to a RADIUS server. The RADIUS server authenticates and authorizes the RADIUS client request, and sends back a RADIUS message response. RADIUS clients also send RADIUS accounting messages to RADIUS servers. Additionally, the RADIUS standards support the use of RADIUS proxies. A RADIUS proxy is a computer that forwards RADIUS messages between RADIUS-enabled computers.

RADIUS messages are sent as User Datagram Protocol (UDP) messages. UDP port 1812 is used for RADIUS authentication messages and UDP port 1813 is used for RADIUS accounting messages. Some network access servers might use UDP port 1645 for RADIUS authentication messages and UDP port 1646 for RADIUS accounting messages. By default, IAS supports receiving RADIUS messages destined to both sets of UDP ports. For information about changing the UDP ports that are used by IAS, see Configure IAS port information. Only one RADIUS message is included in the UDP payload of a RADIUS packet.

RFCs 2865 and 2866 define the following RADIUS message types:

- Access-Request

  Sent by a RADIUS client to request authentication and authorization for a connection attempt.

- Access-Accept

  Sent by a RADIUS server in response to an Access-Request message. This message informs the RADIUS client that the connection attempt is authenticated and authorized.

- Access-Reject

  Sent by a RADIUS server in response to an Access-Request message. This message informs the RADIUS client that the connection attempt is rejected. A RADIUS server sends this message if either the credentials are not authentic or the connection attempt is not authorized.

- Access-Challenge

  Sent by a RADIUS server in response to an Access-Request message. This message is a challenge to the RADIUS client that requires a response.

- Accounting-Request

  Sent by a RADIUS client to specify accounting information for a connection that was accepted.

- Accounting-Response

  Sent by the RADIUS server in response to the Accounting-Request message. This message acknowledges the successful receipt and processing of the Accounting-Request message.

A RADIUS message consists of a RADIUS header and zero or more RADIUS attributes. Each RADIUS attribute specifies a piece of information about the connection attempt. For example, there are RADIUS attributes for the user name, the user password, the type of service requested by the user, and the IP address of the access server. RADIUS attributes are used to convey information between RADIUS clients, RADIUS proxies, and RADIUS servers. For example, the list of attributes in the Access-Request message includes information about the user credentials and the parameters of the connection attempt. In contrast, the Access-Accept message includes information about the type of connection that can be made, connection constraints, and any vendor-specific attributes (VSAs).

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# New features for IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features for IAS**

The following are new features for IAS in Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition:

- Support for RADIUS proxy
- Support for mapping network authentication and authorization for IAS proxy
- Support for IEEE 802.1x wireless and authenticating switches
- Support for Protected Extensible Authentication Protocol (PEAP) for 802.11 wireless clients
- Enhanced EAP configuration for remote access policies
- Support for IAS Network Access Quarantine Control
- Support for logging to an XML-compliant SQL server database
- Ignoring the dial-in properties of user accounts
- Support for configuring RADIUS clients by IP address range
- Support for computer authentication
- Support for checking user certificate purposes
- Improved attribute manipulation
- Support for the Authentication Type remote access policy condition
- Improved support for the Class attribute

Each of these new features is described below.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

**Support for RADIUS proxy**

Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition include RADIUS proxy support in addition to RADIUS server support. With RADIUS proxy support, you can use IAS as a RADIUS message router or forwarder. Based on attributes in the incoming RADIUS message, the RADIUS proxy forwards the message to a specific RADIUS server or client.

IAS implements RADIUS proxy support through connection request processing and remote RADIUS server groups. Connection

request processing is a set of rules that determine how an incoming connection request is processed: It is either authenticated and authorized by the IAS server (when the IAS server is being used as a RADIUS server), or it is forwarded to another RADIUS server for authentication and authorization (when the IAS server is being used as a RADIUS proxy). When IAS is being used as a proxy, it forwards the connection request to a member of a remote RADIUS server group.

By correctly configuring connection request processing rules and remote RADIUS server groups, an IAS server can be used as a RADIUS server for some requests and a RADIUS proxy for others. The IAS proxy also provides the ability to balance RADIUS requests across multiple RADIUS servers, to detect both which RADIUS servers are no longer available and when they become available again. For more information about how RADIUS proxies are used and when to use IAS as a RADIUS proxy, see IAS as a RADIUS proxy.

For more information about configuring IAS as a RADIUS proxy, see Deploying IAS as a RADIUS Proxy and Checklist: Configuring IAS to forward requests.

**Support for mapping network authentication and authorization for IAS proxy**
The proxy component of IAS supports the ability to separate the authentication and authorization of connection requests from access servers. The IAS proxy can forward password-based user credentials to an external RADIUS server for authentication, and perform authorization against a user account in an Active Directory® domain and a locally configured remote access policy. Alternate user authentication databases can be used but connection authorization and restrictions are still determined through local administration.

For example, a visitor from a partner company can be authenticated by using the partner company RADIUS server and user database. Authorization to access your network is accomplished based on connection request policy settings on your IAS server and a Windows account user database in an Active Directory domain that is established for visitor accounts.

You can configure the proxy component with the **Remote-RADIUS-to-Windows-User-Mapping** attribute in the advanced properties of a connection request policy.

For more information, see Mapping network authentication and authorization and Connection request policies.

**Support for IEEE 802.1x wireless and authenticating switches**
IAS provides authentication, authorization, and accounting for connections that use the link-layer standard IEEE 802.1x for wireless and switch access.

**Wireless access point support**
Wireless networking technology has become more widely used since the adoption of wireless industry standards such as IEEE 802.11 and 802.11b. You can use wireless networking to roam within a building or area and connect automatically to a network when you are in the vicinity of a wireless access point.

While providing convenience, wireless networking technologies and wireless access points have the following security risks:

- Anyone who has a compatible wireless network adapter can gain access to the network.
- Wireless networking signals use radio waves to send and receive information. Anyone within an appropriate distance to a wireless access point can detect and receive all data sent to and from the wireless access point.

To counter the first security risk, wireless access points must require authentication and authorization of the wireless node before data can be sent to and received from the network that is attached to the wireless access point. To provide their own authentication and authorization, each wireless access point should require a user account database with each user's authentication credentials and a set of rules by which authorization is granted. Because this is difficult to manage, some wireless access points are RADIUS clients that use the industry standard RADIUS protocol to send connection request and accounting messages to a central RADIUS server. The RADIUS server has access to a user account database and a set of rules for granting authorization. The RADIUS server processes the wireless access point's connection request and either accepts or rejects the connection request.

To counter the second security risk, the data sent between the wireless nodes and the wireless access points must be encrypted. Because of this, the authentication method that is used by the wireless node must allow for the determination of encryption keys that are used to encrypt data.

IAS supports wireless access authentication through the following:

- The use of Extensible Authentication Protocol-Transport Level Security (EAP-TLS) to provide certificate-based authentication and the management of encryption keys that are used to encrypt data sent between wireless access points and wireless nodes.

- The use of the Wireless-IEEE 802.11 port types when configuring the NAS-Port-Type condition of remote access policies.

  By using these port types, you can create a separate remote access policy that contains connection parameters and encryption settings that are specifically designed for wireless nodes.

- By providing a means to grant guest access for a wireless client that does not have a certificate installed.

For more information about configuring IAS for wireless access authentication, see Checklist: Configuring the IAS server and wireless access points for wireless access and Wireless access.

For an example of a remote access policy for wireless clients, see Wireless access example.

**Authenticating switch support**
Network switches are a commonly used technology. A network switch provides selective filtering of the traffic received on each port of the switch and allows better traffic management by segmenting a network. A typical switch, however, sends and receives packets to any node that is connected to the switch. This can be a problem for switch ports in a conference room within an organization. A person from another organization can simply connect their network adapter to the switch port and gain instant access to the network.

To prevent physical layer access to the network, newer switches can be configured to require switch node authentication and authorization before data can be sent to and received from the network that is attached to the switch. To provide their own authentication and authorization, each switch is required to have a user account database with each user's authentication credentials and a set of rules by which authorization is granted. Because this is difficult to manage, newer switches can be RADIUS clients, using the industry standard RADIUS protocol to send connection request and accounting messages to a central RADIUS server. The RADIUS server has access to a user account database and a set of rules for granting authorization. The RADIUS server processes the switch's connection request and either grants the connection request or rejects it.

As the data sent to and from the switch node travels on the wire between the switch node and the switch, encryption of the sent data is neither required nor typically implemented.

IAS supports switch access authentication through the following:

- EAP-TLS is used to provide computer or user certificate-based authentication.

- The Ethernet port type is used when configuring the NAS-Port-Type condition of remote access policies. By using this port type, you can create a separate remote access policy containing connection parameters that are specifically designed for switch nodes.

- Virtual local area network (VLAN) switching is used on the basis of whether or not the switch access client is authenticated.

For more information about configuring IAS for switch access authentication, see Checklist: Configuring the IAS server for authenticated switch access and Ethernet switch access.

**Support for Protected Extensible Authentication Protocol (PEAP) for 802.11 wireless clients**
Protected Extensible Authentication Protocol (PEAP) provides protection for clients and authenticators (IAS or RADIUS servers) that are using EAP. The next generation of EAP, PEAP uses Transport Level Security (TLS) to create end-to-end communication between client and authenticator after the identity of the authenticator is verified.

PEAP creates a session key that is used to derive keys for the encryption of data traffic between the wireless client and the wireless access point. It also provides smooth and efficient roaming between wireless access points that are configured as RADIUS clients on the same IAS server. You can use PEAP with two different authentication types: EAP-MS-CHAPv2 or EAP-TLS. PEAP with EAP-MS-CHAPv2 (also known as PEAP-EAP-MS-CHAPv2), provides secure password authentication for wireless and authenticated switch clients. PEAP-EAP-MS-CHAPv2 can be used to complement the security of MS-CHAPv2 and does not require deployment of user or computer certificates. You can also use PEAP with EAP-TLS (also known as PEAP-EAP-TLS), which provides strong security through a public key infrastructure (PKI) with certificates, which is used for server authentication, and either smart cards or certificates for client and user authentication. When you use PEAP-EAP-TLS, client certificate information is encrypted, providing improved security over the use of EAP-TLS without PEAP.

The following table shows the strengths of PEAP-EAP-MS-CHAPv2 and how it compares to MS-CHAPv2:

| Feature/function | MS-CHAPv2 | PEAP-EAP-MS-CHAPv2 |
|---|---|---|
| Provides client authentication using passwords | Yes | Yes |
| Ensures that the server has access to credentials | Yes | Yes |
| Authenticates the server | No | Yes |
| Prevents wireless access point spoofing | No | Yes |
| Prevents an unauthorized server from negotiating the least secure authentication method | No | Yes |
| Uses TLS keys generated with a public key | No | Yes |
| Provides end-to-end encryption | No | Yes |
| Prevents dictionary or brute force attacks | No | Yes |
| Prevents replay | No | Yes |
| Allows chaining of authentication methods | No | Yes |
| Requires client trust of certificates provided by the server | No | Yes |

For more information, see PEAP.

**Enhanced EAP configuration for remote access policies**

In Microsoft Windows 2000 Server, you can select only a single EAP type for a remote access policy. All connections matching the conditions of the policy must use this single EAP type, which is selected in the policy profile settings. Additionally, the configuration of an EAP type is global for all of the remote access policies configured on the IAS server. This can cause problems when you want to either individually configure the properties for EAP types for each policy or select multiple EAP types by group or for a specific kind of network connection.

In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; these limitations are removed for IAS. For example, you can select different computer certificates for VPN connections and EAP-TLS authentication for wireless connections, or you can select multiple EAP types for wireless connections in the circumstance where some of your wireless clients use EAP-TLS authentication and some of them use PEAP with EAP-MS-CHAPv2.

For more information, see Configure PEAP and EAP methods, PEAP, and EAP.

**Support for IAS Network Access Quarantine Control**

IAS Network Access Quarantine Control provides phased network access, which restricts the access of remote clients to quarantine mode until each client is either verified as meeting or configured according to organization network access policy. After the client computer configuration is verified as meeting organization network policy, the quarantine restrictions, which consist of Quarantine IP-Filters and Session Timers, are removed and standard remote access policy is applied to the connection.

For more information, see IAS Network Access Quarantine Control.

**Support for logging to an XML-compliant SQL server database**

You can use an XML-compliant database, such as Microsoft SQL server™ 2000, to log user authentication and accounting requests, which are received from one or more access servers, to a central data source. Log data is passed from IAS to a procedure stored in a database that supports both structured query language (SQL) and extensible markup language (XML).

For more information, see SQL Server database logging.

**Ignoring the dial-in properties of user accounts**

You can configure a RADIUS attribute on the profile properties of a remote access policy to ignore the dial-in properties of user accounts. The dial-in properties of the user account contain the following:

- Remote access permission

- Caller-ID
- Callback options
- Static IP address
- Static routes

To support multiple types of connections for which IAS provides authentication and authorization, it might be necessary to disable the processing of user account dial-in properties. This can be done to support scenarios in which specific dial-in properties are not required.

For example, the caller-ID, callback, static IP address, and static routes properties are designed for a client that is dialing into a network access server (NAS), not for wireless access points. A wireless access point that receives these settings in a RADIUS message might not be able to process them, which could cause the wireless client to be disconnected. When IAS provides authentication and authorization for users who are both dialing in and accessing the organization network through wireless technology, the dial-in properties must be configured to support either dial-in connections (by setting dial-in properties) or wireless connections (by not setting dial-in properties).

You can use IAS to enable dial-in properties processing for the user account in some scenarios (such as dial-in) and to disable dial-in properties processing for user account dial-in properties in other scenarios (such as wireless and authenticating switch). This is accomplished by configuring the **Ignore-User-Dialin-Properties** attribute on the **Advanced** tab of the profile settings for a remote access policy. For more information, see Add RADIUS attributes to a remote access policy.

The **Ignore-User-Dialin-Properties** attribute is set to the following:

- To enable user account dial-in properties processing, delete the **Ignore-User-Dialin-Properties** attribute or set it to **False**. For example, for a remote access policy that is designed for dial-in connections, no additional configuration is required.
- To disable user account dial-in properties processing, set the **Ignore-User-Dialin-Properties** attribute to the value of **True**. For example, this is set for the remote access policy that is designed for wireless or authenticating switch connections. When the dial-in properties of the user account are ignored, remote access permission is determined by the remote access permission setting in the remote access policy.

You can also use this attribute to manage network access control through groups and the remote access permission on the remote access policy. By setting the **Ignore-User-Dialin-Properties** attribute to the value of **True**, the remote access permission on the user account is ignored. The disadvantage to using the **Ignore-User-Dialin-Properties** attribute in this way is that you cannot use the additional dial-in properties of caller-ID, callback, static IP address, and static routes.

For more information about setting advanced attributes, see Add RADIUS attributes to a remote access policy.

**Note**

- The **Ignore-User-Dialin-Properties** attribute disables the use of all dial-in properties for the user account. Specific dial-in properties cannot be selectively disabled.

**Support for configuring RADIUS clients by IP address range**

For IAS in Windows 2000, you must specify a RADIUS client by IP address or by Domain Name System (DNS) name. In addition, you must configure each RADIUS client separately, even if you have a number of RADIUS clients on the same subnet. While this is not an issue for typical dial-in or VPN access server configurations, numerous wireless access points can be placed on the same subnet. To simplify administration of RADIUS clients for this type of network configuration, IAS in Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition allows you to specify a RADIUS client by using an address range. All of the RADIUS clients in the range must use the same configuration and shared secret.

The address range for RADIUS clients is expressed in the network prefix length notation *w.x.y.z/p*, where *w.x.y.z* is the dotted decimal notation of the address prefix and *p* is the prefix length (the number of high order bits that define the network prefix). This is also known as Classless Inter-Domain Routing (CIDR) notation. An example is 192.168.21.0/24. To convert from subnet mask notation to network prefix length notation, *p* is the number of high order bits set to one in the subnet mask.

For more information, see Configure RADIUS Clients.

**Support for computer authentication**
Active Directory and IAS in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition support the authentication of computer accounts by using standard user authentication methods such as Point-to-Point Protocol (PPP). This allows a computer and its computer certificate to be authenticated for wireless or authenticating switch access clients.

**Support for checking user certificate purposes**
To enforce the use of specific types of user certificates for specific connection types, you can configure IAS to check the purposes (also known as object identifiers or OIDs) of certificates in their Enhanced Key Usage (EKU) extensions. When EAP-TLS is used, the access client can send a smart card certificate or a user certificate from the client computer certificate store during the authentication process. Because smart card certificates are more secure than certificates from the client computer certificate store, you might want to ensure that remote access VPN connections use a smart card certificate for user authentication instead of a locally-installed user certificate. In this case, you can configure a remote access policy to require that the user certificate EKU extensions contain the Smart Card Logon purpose (the object identifier is 1.3.6.1.4.1.311.20.2.2). When the VPN client offers the smart card certificate for user authentication, the IAS server checks the EKU extensions according to remote access policy. Because the Smart Card Logon purpose is present in the VPN client certificate, authentication succeeds. For more information, see Network access authentication and certificates.

You can configure a list of object identifiers that are required to be present in the user certificate (which is offered by the access client). To do this, use the **Allow certificates with these OIDs** attribute on the **Advanced** tab in the properties of a remote access policy profile. No required object identifiers are specified by default. For more information, see Add RADIUS attributes to a remote access policy.

**Improved attribute manipulation**
In Windows 2000, you can use IAS to manipulate the contents of the User-Name RADIUS attribute. Using connection request policies in IAS for Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; you can manipulate the User-Name, Called-Station-ID, and Calling-Station-ID RADIUS attributes. For more information, see Connection request policies.

**Support for the Authentication Type remote access policy condition**
In IAS for Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; you can create remote access policies by using the Authentication Type condition. You can use it to specify connection constraints that are based on the authentication protocol or method that is used by the access client. For more information, see Elements of a remote access policy.

**Improved support for the Class attribute**
In Windows 2000, IAS automatically generates a value for the Class attribute and appends it to the existing value of the Class attribute received in the RADIUS request message. The result is the Class attribute in the RADIUS response message. In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition:

1. You can disable the automatic generation of a value for the Class attribute by using **Generate-Class-Attribute** on the **Advanced** tab in the properties of a remote access policy profile. Automatic generation of a value for the Class attribute is disabled by default. For more information, see Add RADIUS attributes to a remote access policy.

2. Instead of appending the generated value of the Class attribute to the existing Class attribute, IAS creates a separate Class attribute. The RADIUS response message contains both the original Class attribute and the second Class attribute that is generated by IAS.

# Differences in IAS implementations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Differences in IAS implementations**

A previous version of IAS was released with the Microsoft® Windows NT® 4.0 Option Pack. The following table describes the differences between this version and the one available in Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

| **Windows NT 4.0 IAS** | **IAS in Microsoft® Windows Server™ 2003 operating systems** |
|---|---|
| If a domain name is not specified during authentication, the IAS server authenticates the user against only the local SAM database. | IAS resolves a user name without a specified domain name by using the following sequence:<br>1. IAS determines a default domain from the registry, if one is specified.<br>2. If the IAS server is a member of a domain, IAS authenticates the user against that domain.<br>3. If the IAS server is not a member of a domain, IAS authenticates the user against the local SAM database. |
| IAS does not use callback permissions for all user objects. | IAS uses callback permissions for all user objects. |
| IAS log files are written in ASCII. | IAS log files are multilanguage and are written in UTF-8. |

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# How IAS Works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How IAS works**
This section covers:

- IAS as a RADIUS server
- IAS as a RADIUS proxy

**Note**

- You can configure IAS in Microsoft® Windows Server™ 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Microsoft Windows Server 2003, Enterprise Edition, and Microsoft Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# IAS as a RADIUS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS as a RADIUS server**

Internet Authentication Service (IAS) can be used as a RADIUS server to perform authentication, authorization, and accounting for RADIUS clients. A RADIUS client can be either an access server or a RADIUS proxy. When Internet Authentication Service (IAS) is used as a RADIUS server, it provides the following:

- A central authentication and authorization service for all access requests that are sent by RADIUS clients.

  IAS uses either a Microsoft® Windows NT® Server 4.0 domain, an Active Directory® domain, or the local Security Accounts Manager (SAM) to authenticate user credentials for a connection attempt. IAS uses the dial-in properties of the user account and remote access policies to authorize a connection.

- A central accounting recording service for all accounting requests that are sent by RADIUS clients.

  Accounting requests are stored in a local log file for analysis.

The following illustration shows IAS as a RADIUS server for a variety of access clients and a RADIUS proxy. IAS uses an Active Directory domain for user credential authentication of incoming RADIUS Access-Request messages.




When IAS is used as a RADIUS server, RADIUS messages provide authentication, authorization, and accounting for network access connections in the following way:

1. Access servers, such as dial-up network access servers, VPN servers, and wireless access points, receive connection requests from access clients.

2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the IAS server.

3. The IAS server evaluates the Access-Request message.

4. If required, the IAS server sends an Access-Challenge message to the access server. The access server processes the challenge and sends an updated Access-Request to the IAS server.

5. The user credentials are checked and the dial-in properties of the user account are obtained by using a secure connection to a domain controller.

6. The connection attempt is authorized with both the dial-in properties of the user account and remote access policies.

7. If the connection attempt is both authenticated and authorized, the IAS server sends an Access-Accept message to the access server.

   If the connection attempt is either not authenticated or not authorized, the IAS server sends an Access-Reject message to the access server.

8. The access server completes the connection process with the access client and sends an Accounting-Request message to the IAS server, where the message is logged.

9. The IAS server sends an Accounting-Response to the access server.

**Note**

- The access server also sends Accounting-Request messages for the following:

  - During the time in which the connection is established.

  - When the access client connection is closed.

  - When the access server is started and stopped.

You can use IAS as a RADIUS server when:

- You are using a Windows NT Server 4.0 domain, an Active Directory domain, or the local Security Accounts Manager (SAM) as your user account database for access clients.

- You are using the Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 Routing and Remote Access service on multiple dial-up servers, VPN servers, or demand-dial routers and you want to centralize both the configuration of remote access policies and connection logging for accounting.

- You are outsourcing your dial-in, VPN, or wireless access to a service provider. The access servers use RADIUS to authenticate and authorize connections that are made by members of your organization.

- You want to centralize authentication, authorization, and accounting for a heterogeneous set of access servers.

For more information, see Deploying IAS as a RADIUS Server.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# IAS as a RADIUS proxy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS as a RADIUS proxy**

Internet Authentication Service (IAS) can be used as a RADIUS proxy to provide the routing of RADIUS messages between RADIUS clients (access servers) and RADIUS servers that perform user authentication, authorization, and accounting for the connection attempt. When used as a RADIUS proxy, IAS is a central switching or routing point through which RADIUS access and accounting messages flow. IAS records information in an accounting log about the messages that are forwarded.

The following illustration shows IAS as a RADIUS proxy between RADIUS clients (access servers) and either RADIUS servers or another RADIUS proxy.




When IAS is used as a RADIUS proxy between a RADIUS client and a RADIUS server, RADIUS messages for network access connection attempts are forwarded in the following way:

1. Access servers, such as dial-up network access servers, VPN servers, and wireless access points, receive connection requests from access clients.

2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the IAS server that is being used as the IAS RADIUS proxy.

3. The IAS RADIUS proxy receives the Access-Request message and, based on the locally configured connection request policies, determines where to forward the Access-Request message.

4. The IAS RADIUS proxy forwards the Access-Request message to the appropriate RADIUS server.

5. The RADIUS server evaluates the Access-Request message.

6. If required, the RADIUS server sends an Access-Challenge message to the IAS RADIUS proxy, where it is forwarded to the access server. The access server processes the challenge with the access client and sends an updated Access-Request to the IAS RADIUS proxy, where it is forwarded to the RADIUS server.

7. The RADIUS server authenticates and authorizes the connection attempt.

8. If the connection attempt is both authenticated and authorized, the RADIUS server sends an Access-Accept message to the IAS RADIUS proxy, where it is forwarded to the access server.

   If the connection attempt is either not authenticated or not authorized, the RADIUS server sends an Access-Reject message to the IAS RADIUS proxy, where it is forwarded to the access server.

9. The access server completes the connection process with the access client and sends an Accounting-Request message to the IAS RADIUS proxy. The IAS RADIUS proxy logs the accounting data and forwards the message to the RADIUS server.

10. The RADIUS server sends an Accounting-Response to the IAS RADIUS proxy, where it is forwarded to the access server.

You can use IAS as a RADIUS proxy when:

- You are a service provider that offers outsourced dial, virtual private network (VPN), or wireless network access services to multiple customers. Your network access servers send connection requests to the IAS RADIUS proxy. Based on the realm portion of the user name in the connection request, the IAS RADIUS proxy forwards the connection request to a RADIUS server that is maintained by the customer and can authenticate and authorize the connection attempt. For more information, see Realm names.

- You want to provide authentication and authorization for user accounts that are not members of either the domain in which the IAS server is a member or another domain that has a two-way trust with the domain in which the IAS server is a member. This includes accounts in untrusted domains, one-way trusted domains, and other forests. Instead of configuring your access servers to send their connection requests to an IAS RADIUS server, you can configure them to send their connection requests to an IAS RADIUS proxy. The IAS RADIUS proxy uses the realm name portion of the user name and forwards the request to an IAS server in the correct domain or forest. Connection attempts for user accounts in one domain or forest can be authenticated for network access servers in another domain or forest.

  IAS supports authentication across forests without a RADIUS proxy when the two forests contain only domains that consist of domain controllers running Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. The forest functional level must be Windows Server 2003 , and there must be a two-way trust relationship between forests. If you use EAP-TLS with certificates as your authentication method, however, you must use a RADIUS proxy for authentication across forests that consist of Windows Server 2003 domains.

- You want to perform authentication and authorization by using a database that is not a Windows account database. In this case, connection requests that match a specified realm name are forwarded to a RADIUS server, which has access to a different database of user accounts and authorization data. Examples of other user databases include Novell Directory Services (NDS) and Structured Query Language (SQL) databases.

- You want to process a large number of connection requests. In this case, instead of configuring your RADIUS clients to attempt to balance their connection and accounting requests across multiple RADIUS servers, you can configure them to send their connection and accounting requests to an IAS RADIUS proxy. The IAS RADIUS proxy dynamically balances the load of connection and accounting requests across multiple RADIUS servers and increases the processing of large numbers of RADIUS clients and authentications per second.

- You want to provide RADIUS authentication and authorization for outsourced service providers and minimize intranet firewall configuration. An intranet firewall is between your perimeter network (the network between your intranet and the Internet) and intranet. By placing an IAS server on your perimeter network, the firewall between your perimeter network and intranet must allow traffic to flow between the IAS server and multiple domain controllers. When replacing the IAS server with an IAS proxy, the firewall must allow only RADIUS traffic to flow between the IAS proxy and one or multiple IAS servers within your intranet.

For more information, see Deploying IAS as a RADIUS Proxy.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Understanding IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding IAS**
This section covers:

- Components of a RADIUS infrastructure
- Realm names
- IAS and tunnels
- Managing multiple IAS servers
- Network access authentication and certificates
- Securing IAS
- Vendor-Specific Attributes
- Authentication methods
- Remote Access Policies
- Connection Request Processing
- Remote Access Logging

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Components of a RADIUS infrastructure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Components of a RADIUS infrastructure**
A RADIUS authentication, authorization, and accounting infrastructure consists of the following components:

- Access clients
- Access servers (RADIUS clients)
- RADIUS proxies
- RADIUS servers
- User account databases

These components are shown in the following illustration.




**Access clients**
An access client is a device that requires some level of access to a larger network. Examples of access clients are dial-up or virtual private network (VPN) clients, wireless clients, or LAN clients connected to a switch.

**Access servers (RADIUS clients)**
An access server is a device that provides some level of access to a larger network. An access server using a RADIUS infrastructure is also a RADIUS client, sending connection requests and accounting messages to a RADIUS server. Examples of access servers are:

- Network access servers (NASs) that provide remote access connectivity to an organization network or the Internet. An example is a Windows 2000 computer running the Routing and Remote Access service and providing either traditional dial-up or virtual private network (VPN) remote access services to an organization's intranet.
- Wireless access points that provide physical layer access to an organization's network, using wireless-based transmission and reception technologies.
- Switches that provide physical layer access to an organization's network, using traditional LAN technologies such as Ethernet.

**RADIUS proxies**
A RADIUS proxy is a device that forwards or routes RADIUS connection requests and accounting messages between RADIUS clients (and RADIUS proxies) and RADIUS servers (or RADIUS proxies). The RADIUS proxy uses information within the RADIUS message, such as the User-Name or Called-Station-ID RADIUS attributes, to route the RADIUS message to the appropriate

RADIUS server.

A RADIUS proxy can be used as a forwarding point for RADIUS messages when the authentication, authorization, and accounting must occur at multiple RADIUS servers in different organizations.

**RADIUS servers**
A RADIUS server is a device that receives and processes connection requests or accounting messages sent by RADIUS clients or RADIUS proxies. In the case of connection requests, the RADIUS server processes the list of RADIUS attributes in the connection request. Based on a set of rules and the information in the user account database, the RADIUS server either authenticates and authorizes the connection and sends back an Access-Accept message or sends back an Access-Reject message. The Access-Accept message can contain connection restrictions that are implemented by the access server for the duration of the connection.

**User account databases**
The user account database is the list of user accounts and their properties that can be checked by a RADIUS server to verify authentication credentials and user account properties containing authorization and connection parameter information.

The user account databases that IAS can use are the local Security Accounts Manager (SAM), a Microsoft Windows NT 4.0 domain, or the Active Directory directory service. For Active Directory, IAS can provide authentication and authorization for user or computer accounts in the domain in which the IAS server is a member, two-way trusted domains, and trusted forests with domain controllers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition.

If the user accounts for authentication reside in a different type of database, IAS can be configured as a RADIUS proxy to forward the authentication request to a RADIUS server that does have access to the user account database. Different databases for Active Directory include untrusted forests, untrusted domains, or one-way trusted domains.

### Note

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Realm names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Realm names**

When an access client sends user credentials, a user name is often included. Within the user name are two elements:

1. Identification of the user account name
2. Identification of the user account location

For example, for the user name user1@microsoft.com, user1 is the user account name and microsoft.com is the location of the user account. The identification of the location of the user account is known as a realm. There are different forms of realm names:

- The realm name can be a prefix.

  For example, Microsoft\user1, where Microsoft is the name of a Windows NT 4.0 domain.

- The realm name can be a suffix.

  For example, user1@microsoft.com, where microsoft.com is either a DNS domain name or the name of an Active Directory domain.

You can require users to specify the realm name by using the correct syntax when they connect. For example, you can require that the user type the user name (including the user account name and the realm) in **User name** in the **Connect** dialog box when making a dial-up or VPN connection in Network Connections. You can also add the realm name automatically to the user account name. For example, you can specify a realm name and syntax when you create a custom dialing package with the Connection Manager Administration Kit, so that the user has to specify only the user account name. For more information, see Connection Manager Administration Kit.

The user name is passed from the access client to the access server during the authentication phase of the connection attempt. This user name becomes the User-Name RADIUS attribute in the Access-Request message that is sent by the access server to its RADIUS server.

If the RADIUS server is an IAS server, the Access-Request message is evaluated against the set of configured connection request policies. Conditions on the connection request policy can include the specification of the contents of User-Name attribute. You can configure a set of connection request policies that are specific to the realm name within the User-Name attribute of incoming messages. This allows to you create routing rules that forward RADIUS messages with a specific realm name to a specific set of RADIUS servers when IAS is used as a RADIUS proxy.

Before the RADIUS message is either processed locally (when IAS is being used as a RADIUS server) or forwarded to another RADIUS server (when IAS is being used as a RADIUS proxy), the User-Name attribute in the message might be modified by attribute manipulation rules that are configured on the profile of the first matching connection request policy. Attribute manipulation rules for the User-Name attribute might change the following:

- Remove the realm name from the user name (also known as realm stripping).

  For example, the user name user1@microsoft.com is changed to user1.

- Change the realm name but not its syntax.

  For example, the user name user1@microsoft.com is changed to user1@wcoast.microsoft.com.

- Change the syntax of the realm name.

  For example, the user name microsoft\user1 is changed to user1@microsoft.com.

Attribute manipulation rules are configured in the **Attribute** tab on the profile of a connection request policy. For more information, see Configure attribute manipulation. IAS attribute manipulation rules use regular expression syntax. For more

information about this syntax and examples, see Pattern matching syntax.

After the User-Name attribute is modified according to attribute manipulation rules, additional settings on the profile of the first matching connection request policy are used to determine whether:

- The IAS server processes the Access-Request message locally (when IAS is being used as a RADIUS server).
- The IAS server forwards the message to another RADIUS server (when IAS is being used as a RADIUS proxy).

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Notes**

- When the user name does not contain a domain name, IAS supplies one. By default, the IAS-supplied domain name is the domain of which the IAS server is a member. You can specify the IAS-supplied domain name through the following registry setting:

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RasMan\PPP\ControlProtocols\BuiltIn\DefaultDomain**

- Some third party network access servers delete or modify the domain name as specified by the user. As the result, the network access request is authenticated against the default domain, which might not be the domain for the user's account. To resolve this problem, configure your RADIUS servers to change the user name into the correct format with the accurate domain name.

# IAS and tunnels

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS and tunnels**

Tunneling, also known as encapsulation, is a method of using an internetwork infrastructure of one protocol to transfer a payload. Sometimes, the payload consists of the frames (or packets) of another protocol. Instead of being sent as the originating host produces it, the frame is encapsulated with an additional header. The additional header provides routing information so that the encapsulated payload can traverse an intermediate internetwork (also known as a transit internetwork). The encapsulated packets are then routed between tunnel endpoints over the transit internetwork. After the encapsulated payload packets reach their destination on the transit internetwork, the frame is de-encapsulated and forwarded to its final destination.

The entire process of the encapsulation, transmission, and de-encapsulation of packets is called tunneling. The logical path in the transit internetwork through which the encapsulated packets travel is called a tunnel.

With remote access virtual private network (VPN) connections, there are two types of tunneling:

1. Voluntary tunneling

2. Compulsory tunneling

**Voluntary tunneling**

A user or client computer can issue a VPN request to configure and create a voluntary tunnel. In this case, the user's computer is a tunnel endpoint that acts as the tunnel client. Voluntary tunneling occurs when a workstation or router uses tunneling client software to create a VPN connection to the target tunnel server. In order to accomplish this, the appropriate tunneling protocol must be installed on the client computer. In a dial-up situation, which is the most common use, the client must establish a dial-up connection to the internetwork before the client can set up a tunnel. A good example of this is the dial-up Internet user, who must dial an ISP and obtain an Internet connection before a tunnel over the Internet is created.

Voluntary tunneling is not different from other types of network access, and IAS can be used for authentication, authorization, and accounting.

**Compulsory tunneling**

Compulsory tunneling is the creation of a secure tunnel by another computer or network device on the client computer's behalf. Compulsory tunnels are configured and created automatically for users without their knowledge or intervention. With a compulsory tunnel, the user's computer is not a tunnel endpoint. Another device between the user's computer and the tunnel server is the tunnel endpoint, which acts as the tunnel client.

Some vendors that sell dial-up access servers provide for the creation of a tunnel on behalf of a dial-up client. The computer or network device that provides the tunnel for the client computer is known as a front-end processor (FEP) in PPTP, an L2TP access concentrator (LAC) in L2TP, or an IP Security gateway in IPSec. The term FEP is used to describe tunnel creation functionality, regardless of the protocol used. To perform its function, the FEP must have the appropriate tunneling protocol installed and must be capable of establishing the tunnel when the client computer attempts a connection. The Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000 Routing and Remote Access service cannot be used as a FEP.

An organization can contract with an ISP to deploy a nationwide set of FEPs. These FEPs can establish tunnels across the Internet to a VPN server that is connected to the organization private network, thereby consolidating calls from geographically diverse locations into a single Internet connection at the organization network.

There are two types of compulsory tunneling. In the first type, the tunnel is created before the access client is authenticated. Based on the realm name or the caller ID of the access client, the FEP sends an Access-Request to an IAS server. The IAS server sends back an immediate Access-Accept with RADIUS attributes for the tunnel creation without performing authentication and authorization. After the tunnel is created, the access client authenticates against the tunnel server.

In the second type of compulsory tunneling, the tunnel is created after the access client is authenticated by the FEP. In this case, the FEP sends the Access-Request message with the client credentials to an IAS server. The IAS server authenticates and authorizes the connection attempt and returns RADIUS attributes in the Access-Accept message, which specify to the NAS how to initiate a tunnel to a VPN server. The tunnel-endpoint (the VPN server at which the tunnel is terminated), can be changed on

the basis of conditions in a remote access policy. For example, the tunnel-endpoint can be changed on the basis of the user name or the user account group membership. Controlling compulsory tunnels with remote access policies provides more flexibility than static tunneling (which requires a dedicated network access server) or realm-based tunneling (which requires that all users in a given realm use the same tunnel settings).

**RADIUS attributes used with compulsory tunneling**

The following is a list of RADIUS attributes that are used to carry compulsory tunneling information from the IAS server to the FEP. The same set of attributes is used regardless of the type of compulsory tunnel.

- Attributes used only for authorization:
    - Tunnel-Preference
    - Tunnel-Password

      This attribute should not be used with RADIUS proxies.
- Attributes used in authorization and accounting:
    - Tunnel-Type (for example, PPTP and L2TP)
    - Tunnel-Medium-Type (for example, X.25, ATM, Frame Relay, and IP)
    - Tunnel-Client-Endpoint
    - Tunnel-Server-Endpoint
    - Tunnel-Private-Group-ID
    - Tunnel-Assignment-ID
    - Tunnel-Client-Auth-ID
    - Tunnel-Server-Auth-ID
- Attributes used only for accounting:
    - Acct-Tunnel-Connection-ID

**Compulsory tunneling example**

The following table lists the set of compulsory tunneling RADIUS attributes that specify a compulsory tunnel. This tunnel uses the Layer Two Tunneling Protocol (L2TP) with a VPN server at the IP address of 131.107.9.41, and a tunnel password of og*37y#cW@95?4xT.

| Attribute | Value |
|---|---|
| Framed-Protocol | PPP |
| Service-Type | Outbound-User |
| Tunnel-Medium-Type | IP (IP version 4) |
| Tunnel-Password | og*37y#cW@95?4xT |
| Tunnel-Server-Endpt | 131.107.9.41 |
| Tunnel-Type | Layer Two Tunneling Protocol (L2TP) |

To configure IAS, add your versions of all of these attributes and their appropriate values on the **Advanced** tab in the profile

properties of the applicable remote access policy. For more information, see Add RADIUS attributes to a remote access policy.

**Notes**

- You can use a RADIUS Access-Accept packet to send tunnel attributes that describe multiple tunnels. However, IAS only returns information for a single tunnel.

- IAS only supports one compulsory tunnel for each connection request and remote access policy.

# Managing multiple IAS servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing multiple IAS servers**

To prevent a single IAS server from becoming a single point of failure for authentication, authorization, and accounting, IAS servers are often deployed either in pairs (a primary and a backup) to provide fault tolerance, or in multiples to balance the load of large numbers of authentication and accounting requests.

When you deploy more than one IAS server computer to provide the same authentication, authorization, and accounting service to RADIUS clients and proxies, you must synchronize the configuration of the IAS server computers. You can do this through the following:

- Internet Authentication Service snap-in

  You can use the Internet Authentication Service snap-in to manage both local and remote IAS server computers. If you make a configuration change to one IAS server, you must make the same configuration change to all of the IAS servers that provide the same service. For information about managing a remote IAS server computer, see Manage a remote IAS server. This method might be acceptable for a small number of servers and minor configuration changes.

- Copying the configuration of one IAS server computer to another

  With this method, you can use the **netsh aaaa show config** to export the configuration of one IAS server to a Netsh script file. Next, you can use the **netsh exec** command to import the IAS server configuration on another IAS server. For more information, see Copy the IAS configuration to another server. This method is recommended for a large number of configuration changes or a larger number of servers. You can use this method to centrally manage RADIUS and remote access policy configuration in a large enterprise network.

When you type **netsh aaaa show config**, it creates a Netsh script file that is designed to be used with the **netsh exec** command. The contents of the display of the **netsh aaaa show config** command includes:

- IAS settings
- RADIUS clients
- Remote access policies
- Connection request policies
- Remote access logging settings

For more information, see Netsh commands for AAAA.

The Routing and Remote Access service uses the same database as IAS. There are no settings, however, for IAS, RADIUS clients, or connection request policies. If settings are present, they are ignored by the Routing and Remote Access service. Because of this, it is possible to copy the configuration of an IAS server computer to a Routing and Remote Access server. For example, you can update the remote access policies of a branch office demand-dial router computer.

**Notes**

- The **netsh aaaa show config** and **netsh exec** commands can be run without stopping the Internet Authentication Service or the Routing and Remote Access service.
- The version number of the IAS database from which the configuration was exported, should be less than or equal to the version of the IAS database to which the configuration is being imported. To view the version of the IAS database from which the configuration was exported, view the Netsh command script file, which is created with the **netsh aaaa show config** command. To view the version of the IAS database on the IAS server computer on which the Netsh script is being

run, use the **netsh aaaa show version** command.

- The **netsh aaaa show config** and **netsh exec** commands are used to export and import the entire database. You cannot specify which elements of the IAS database to export or import.

- Do not copy the configuration from a Routing and Remote Access server to an IAS server. By default, the Routing and Remote Access service version of the IAS database does not contain all of the configuration items contained in the full IAS database.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Network access authentication and certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network access authentication and certificates**
This topic contains the following sections:


- Overview
- Certificate requirements for EAP
- Computer authentication by IPSec
- Certificate-based authentication and wireless clients
- Certificate enrollment methods and domain membership
- Choosing a certificate enrollment method
- CA Web enrollment services


**Overview**
Certificates are used for network access authentication because they provide strong security for authenticating users and computers and eliminate the need for less secure password-based authentication methods. This topic describes how Internet Authentication Service (IAS) and virtual private network (VPN) servers use Extensible Authentication protocol-Transport Level Security (EAP-TLS), protected Extensible Authentication protocol (PEAP), or Internet protocol security (IPSec) to perform certificate-based authentication for many types of network access, including VPN and wireless connections. Additionally, this topic describes certificate enrollment methods to help you determine the best certificate enrollment method for your use.

When discussing authentication, a server is defined as a VPN or IAS server that is a TLS end point. You can configure VPN servers to perform network access authentication without IAS, or you can use IAS for authentication when you have multiple Remote Access Dial-In User Service (RADIUS) clients (such as VPN servers and wireless access points) on your network.

Two authentication methods use certificates: Extensible Authentication protocol-Transport Level Security (EAP-TLS) and protected Extensible Authentication protocol (PEAP). Both methods always use certificates for server authentication. Depending on the authentication type configured with the authentication method, certificates might be used for user authentication and client authentication. For more information, see EAP and PEAP.

The use of certificates for authentication of VPN connections is the strongest form of authentication available with the Windows Server 2003 family. You must use certificate-based authentication for VPN connections based on Layer Two Tunneling protocol over Internet protocol security (L2TP/IPSec). Point-To-Point Tunneling protocol (PPTP) connections do not require certificates, although you can configure PPTP connections to use certificates for computer authentication when you use EAP-TLS as the authentication method. For wireless clients (computing devices with wireless network adapters, such as your portable computer or personal digital assistant), PEAP with EAP-TLS and smart cards or certificates is the recommended authentication method. For more information, see Layer Two Tunneling Protocol, Point-to-Point Tunneling Protocol, and Wireless Networking.

**Notes**

- During authentication attempts for PPTP connections, computer authentication does not occur. When EAP-TLS is used as the authentication method with PPTP connections, user authentication is performed with a certificate from the certificate store on the local computer or from a smart card.

- On Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 Point-to-Point Tunneling protocol (PPTP) ports, and you can create up to 1,000 Layer Two Tunneling protocol (L2TP) ports. However, Windows Server 2003, Web Edition, can accept only one virtual private network (VPN) connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

In order to deploy certificates to users and computers, you must first do the following:

- Design a public key infrastructure (PKI).
- Deploy a certification authority (CA) using Certificate Services.
- Design and publish one or more certificates using Certificate Templates, which is installed with Certificate Services.
- Select one or more distribution methods (also known as certificate enrollment methods) to install the certificates on computers and distribute them to users.

### Notes

- This topic assumes that you have designed and deployed a PKI according to the guidelines provided in Certificate Services Best practices. For more information, see Public Key Infrastructure.
- PEAP is not supported for VPN or dial-up connections.

#### Example certificate deployments

- Remote access VPN connections

  L2TP/IPSec or PPTP connections between VPN server and VPN client with EAP-TLS used as the authentication method. IPSec uses computer certificates between the client and the server for authentication, and EAP-TLS uses a certificate (from a smart card or the user's local certificate store) for user authentication. The IAS or VPN server certificate must contain the server Authentication purpose, and the client computer or user certificate must contain the Client Authentication purpose in the Enhanced Key Usage (EKU) extensions of the certificate.

  For more information, see L2TP-based remote access VPN deployment and PPTP-based remote access VPN deployment.

- Router-to-router VPN connections

  L2TP/IPSec connections for dedicated or dial-on-demand connections between servers with EAP-TLS as the authentication method. Both servers must have certificates that contain the server Authentication and Client Authentication purposes in EKU extensions.

  For more information, see Deploying Router-to-Router VPNs.

- IEEE 802.1X wireless or switch clients

  PEAP-EAP-MS-CHAPv2 is configured as the authentication method, and the **Validate server certificate** option is enabled on the client computers. The IAS server certificate EKU extensions contain the server Authentication purpose, and the certificate is used to identify the server to the client. User authentication is accomplished with user name and password.

  For more information, see Enable smart card or other certificate authentication and Configuring wireless network settings on client computers.

- IEEE 802.1X wireless or switch clients

  L2TP/IPSec is used, and EAP-TLS with certificates is configured as the authentication method. The IAS server certificate contains the server Authentication purpose in EKU extensions to identify itself to the client, and the client uses a certificate (from a smart card or the user's local certificate store) to identify itself to the IAS server. (The wireless access points are configured as RADIUS clients on the IAS server, which is the EAP authenticator.)

**Note**

- User authentication and network access authorization are performed by either the VPN server or a RADIUS server, such as Internet Authentication Service (IAS), as configured on the VPN gateway. For more information, see Use RADIUS authentication.

**Certificate requirements for EAP**

When you use EAP with a strong EAP type (such as TLS with smart cards or certificates), both the client and the server use certificates to verify their identities to each other. Certificates must meet specific requirements in order to allow the server and the client to use them for successful authentication.

One such requirement is that the certificate is configured with one or more purposes in EKU extensions that correlate to the certificate use. For example, a certificate used for the authentication of a client to a server must be configured with the Client Authentication purpose. Similarly, a certificate used for the authentication of a server must be configured with the server Authentication purpose. When certificates are used for authentication, the authenticator examines the client certificate, seeking the correct purpose object identifier in EKU extensions. For example, the object identifier for the Client Authentication purpose is 1.3.6.1.5.5.7.3.2.

Certificate Templates allow customization of the certificates issued by Certificate Services, including both how certificates are issued and what they contain, including their purposes. In Certificate Templates, you can use a default template, such as the Computer template, to define the template that the CA uses to assign certificates to computers. You can also create a certificate template and assign purposes in EKU extensions to the certificate. By default, the Computer template includes the Client Authentication purpose and the server Authentication purpose in EKU extensions.

The certificate template that you create can include any purpose for which the certificate will be used. For example, if you use smart cards for authentication, you can include the Smart Card Logon purpose in addition to the Client Authentication purpose.

When using IAS, you can configure IAS to check certificate purposes before granting network authorization. IAS can check additional EKUs and Issuance Policy purposes (also known as Certificate Policies).

Some non-Microsoft CA software might contain a purpose named All, which represents all possible purposes. This is indicated by a blank (or null) EKU extension. Although All means all possible purposes, the All purpose cannot be substituted for the Client Authentication purpose, the server authentication purpose, or any other purpose related to network access authentication.

For more information, see Certificate Templates and Manage Certificate Templates for an Enterprise Certification Authority.

You can view certificates and their purposes by using the Certificates console to open the certificate store. Certificate stores can be viewed in Logical Store mode or in Purpose mode. For information about how to view certificate purposes, see Display certificate stores in Purpose mode.

**Minimum certificate requirements**

All certificates that are used for network access authentication must meet the requirements for X.509 certificates and work for connections that use Secure Sockets Layer-Transport Level Security (SSL/TLS). After this minimum requirement is met, both client and server certificates have additional requirements.

**Client certificate requirements**

With EAP-TLS or PEAP-EAP-TLS, the server accepts the client authentication attempt when the certificate meets the following requirements:

- The client certificate is issued by an enterprise CA or mapped to a user or computer account in Active Directory.

- The user or computer certificate on the client chains to a trusted root CA, includes the Client Authentication purpose in EKU extensions (the object identifier for Client Authentication is 1.3.6.1.5.5.7.3.2), and fails neither the checks that are performed by CryptoAPI and specified in the remote access policy nor the Certificate object identifier checks that are

specified in IAS remote access policy.

- The 802.1X client does not use registry-based certificates that are either smart card-logon or password-protected certificates.

- For user certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate contains the user principal name (UPN). To configure the UPN name in a certificate template:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.

4. In **Include this information in alternate subject name**, select **User principal name (UPN)**.

- For computer certificates, the Subject Alternative Name (SubjectAltName) extension in the certificate must contain the client's fully qualified domain name (FQDN), which is also called the DNS name. To configure this name in the certificate template:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.

4. In **Include this information in alternate subject name**, select **DNS name**.

With PEAP-EAP-TLS and EAP-TLS, clients display a list of all installed certificates in the Certificates snap-in, with the following exceptions:

- Wireless clients do not display registry-based and smart card-logon certificates.

- Wireless clients and VPN clients do not display password-protected certificates.

- Certificates that do not contain the Client Authentication purpose in EKU extensions are not displayed.

**Server certificate requirements**

Clients can be configured to validate server certificates by using the **Validate server certificate** option. With PEAP-EAP-MS-CHAPv2, PEAP-EAP-TLS, or EAP-TLS as the authentication method, the client accepts the server's authentication attempt when the certificate meets the following requirements:

- The Subject name contains a value. If you issue a certificate to your IAS server that has a blank Subject, the certificate is not available to authenticate your IAS server. To configure the certificate template with a Subject name:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.

4. In **Subject name format**, select a value other than **None**.

- The computer certificate on the server chains to a trusted root CA and does not fail any of the checks that are performed by CryptoAPI and specified in the remote access policy.

- The IAS or VPN server computer certificate is configured with the server Authentication purpose in EKU extensions (the object identifier for server Authentication is 1.3.6.1.5.5.7.3.1).

- The server certificate is configured with a required cryptographic service provider (CSP) value of **Microsoft RSA SChannel Cryptographic provider**. To configure the required CSP:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Request Handling** tab, and then click **CSPs**.

4. In **CSP Selection**, select **Requests must use one of the following CSPs**.

5. In **CSPs**, select the **Microsoft RSA SChannel Cryptographic provider** checkbox. Clear all other checkboxes in **CSPs**.

- The Subject Alternative Name (SubjectAltName) extension, if used, must contain the DNS name of the server. To configure the certificate template with the DNS name of the enrolling server:

1. Open Certificate Templates.

2. In the details pane, right-click the certificate template that you want to change, and then click **properties**.

3. Click the **Subject Name** tab, and then click **Build from this Active Directory information**.

4. In **Include this information in alternate subject name**, select **DNS name**.

With PEAP and EAP-TLS, servers display a list of all installed certificates in the computer's certificate store, with the following exceptions:

- Certificates that do not contain the server Authentication purpose in EKU extensions are not displayed.

- Certificates that do not contain a Subject name are not displayed.

- Servers do not display registry-based and smart card-logon certificates.

**Note**

- You can designate which certificate is used by the IAS or VPN server in the remote access policy profile. When you configure EAP and PEAP authentication methods that require a certificate for server authentication and you do not select a specific certificate in the **Smart Card or other Certificate properties** dialog box, the IAS or VPN server automatically selects a certificate from the computer certificate store. If the server obtains a newer certificate, it uses the new certificate. This might cause the IAS or VPN server to use a certificate that is not correctly configured for authentication, causing authentication to fail as the result. To prevent this, always select a server certificate when configuring PEAP and EAP authentication methods that require one. For more information, see Configure PEAP and EAP methods.

**Computer authentication by IPSec**

Computer authentication is first performed during L2TP/IPSec connection attempts between remote access client and server. When a secure channel is established between the client and the server, the user authentication and authorization attempt proceeds.

Computer authentication by IPSec is performed using preshared keys or computer certificates. The recommended method of authentication is the use of a PKI and certificates. If certificates are used, a computer certificate is required to establish IPSec trust during the Internet Key Exchange (IKE) negotiation in L2TP/IPSec-based VPN connections.

In order for a VPN server running Windows Server 2003 and a VPN client running Windows 2000 or Windows XP to establish trust for an L2TP/IPSec VPN connection, both computers must have a computer certificate that was issued by the same trusted enterprise root CA. When both computers have and exchange certificates that were issued by the trusted root CA during IPSec negotiation, they extend trust to each other, and the security association is made.

**How VPN servers use certificates**

When an L2TP/IPSec VPN connection is attempted between VPN client and server, computer authentication fails if the VPN client certificate (from a smart card or the certificate store on the local computer) is not configured with the Client Authentication purpose in EKU extensions, and the VPN server certificate is not configured with the server Authentication purpose in EKU extensions. IPSec checks the EKU extensions for the client certificate to determine whether the Client Authentication purpose object identifier is present. When the EKU extension includes the Client Authentication purpose object identifier, IPSec can use the certificate for authentication. Similarly, the client checks the VPN server certificate for the server Authentication purpose object identifier in EKU extensions.

Although VPN servers that end connections for remote users need a certificate configured with the server Authentication purpose in EKU extensions only, a VPN servers that is used as an end point for a VPN connection with another VPN server originates (as a client) and ends (as a server) VPN connections. For this reason, the certificate on these servers must contain both the server Authentication purpose and the Client Authentication purpose in EKU extensions. In addition, due to the way in which automatic certificate selection functions, both purposes (server Authentication and Client Authentication) must be contained in the same certificate.

Automatic certificate selection can provide IPSec with any certificate in the certificate store that chains to the trusted enterprise root CA, regardless of the purposes contained in the certificate. If two certificates are installed on the server--one that contains the Client Authentication purpose and one that contains the server Authentication purpose--it is possible that automatic certificate selection could use the wrong certificate for authentication. For example, the certificate with the Client Authentication purpose might be required, but the certificate provided through automatic certificate selection contains the server Authentication purpose. In this and similar instances, computer authentication fails.

**Certificate-based authentication and wireless clients**

IEEE 802.1X authentication provides authenticated access to 802.11 wireless networks and to wired Ethernet networks. 802.1X provides support for secure EAP types, such as TLS with smart cards or certificates. You can configure 802.1X with EAP-TLS in a variety of ways. If the **Validate server certificate** option is configured on the Windows XP Professional client, the client authenticates the server by using its certificate. Client computer and user authentication can be accomplished using certificates from the client certificate store or a smart card, providing mutual authentication.

With wireless clients, PEAP-EAP-MS-CHAPv2 can be used as the authentication method. PEAP-EAP-MS-CHAPv2 is a password-based user authentication method that uses TLS with server certificates. During PEAP-EAP-MS-CHAPv2 authentication, the IAS or RADIUS server supplies a certificate to validate its identity to the client (if the **Validate server certificate** option is configured on the Windows XP Professional client). Client computer and user authentication is accomplished with passwords, which eliminates some of the difficulty of deploying certificates to wireless client computers.

802.1X wireless and switch clients can also use PEAP-EAP-TLS, which provides strong security. PEAP-EAP-TLS uses a public key infrastructure (PKI) with certificates used for server authentication and either smart cards or certificates for client computer and user authentication. For more information, see PEAP.

For more information about 802.1X authentication, see Understanding 802.1X authentication for wireless networks.

**Certificate enrollment methods and domain membership**

The domain membership of computers for which you want to enroll certificates affects the certificate enrollment method that you can choose. Certificates for domain member computers can be enrolled automatically (also known as auto-enrollment), while an administrator must enroll certificates for non-domain member computers using the Web or a floppy disk. The certificate enrollment method for non-domain member computers is known as a trust bootstrap process, through which certificates are created and then manually requested or distributed securely by administrators, to build common trust.

**Domain member certificate enrollment**

If your VPN server, IAS server, or client running Windows 2000 or Windows XP is a member of a domain running Windows Server 2003 and Active Directory, you can configure the auto-enrollment of computer and user certificates. After auto-enrollment is configured and enabled, all domain member computers receive computer certificates when Group Policy is refreshed next, whether the refresh is triggered manually with the **gpupdate** command, or by logging on to the domain.

If your computer is a member of a domain where Active Directory is not installed, you can install computer certificates

manually by requesting them through Certificates.

For more information, see Computer certificates for L2TP/IPSec VPN connections.

**Note**

- Computers running Windows 2000 can auto-enroll computer certificates only.

**Non-domain member certificate enrollment**

Certificate enrollment for computers that are not domain members cannot be done with auto-enrollment. When a computer is joined to a domain, a trust is established that allows auto-enrollment to occur without administrator intervention. When a computer is not joined to a domain, trust is not established and a certificate is not issued. Trust must be established using one of the following methods:

- An administrator (who is, by definition, trusted) must request a computer or user certificate using the CA Web enrollment tool.

- An administrator must save a computer or user certificate to a floppy disk and install it on the non-domain member computer. Or, when the computer is not accessible to the administrator (for example, a home computer connecting to an organization network with an L2TP/IPSec VPN connection), a domain user whom the administrator trusts can install the certificate.

- An administrator can distribute a user certificate on a smart card (computer certificates are not distributed on smart cards).

Many network infrastructures contain VPN and IAS servers that are not domain members. For example, a VPN server in a perimeter network might not be a domain member for security purposes. In this case, a computer certificate with the server Authentication purpose contained in the EKU extensions must be installed on the non-domain member VPN server before it can successfully negotiate L2TP/IPSec-based VPN connections with clients. Note that if the non-domain member VPN server is used as an end point for a VPN connection with another VPN server, EKU extensions must contain both the server Authentication and Client Authentication purposes.

**Choosing a certificate enrollment method**

If you are running an enterprise certification authority (CA) on a computer running Windows Server 2003, Standard Edition, you can use the following table to determine the best certificate enrollment method for your requirements:

| Object and domain membership | Certificate template | Certificate purposes | Preferred certificate enrollment method | Alternate certificate enrollment method |
|---|---|---|---|---|
| VPN or IAS server, domain member | Computer | Server Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| VPN server with site-to-site connection, domain member | Computer | Server Authentication and Client Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| Windows XP client, domain member | Computer | Client Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| VPN or IAS server, non-domain member | Computer | Server Authentication | CA Web enrollment tool | Install from a floppy disk |
| VPN server with site-to-site connection, non-domain member | Computer | Server Authentication and Client Authentication | CA Web enrollment tool | Install from a floppy disk |
| Windows XP client, non-domain member | Computer | Client Authentication | CA Web enrollment tool | Install from a floppy disk |
| User, domain user | User | Client Authentication | Auto-enrollment | Use a smart card or the CA Web enrollment tool |

If your certification authority is on a computer running one of the following operating systems, the RAS and IAS servers and Workstation Authentication templates are available for use:

- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition
- Windows Server 2003, Enterprise Edition for Itanium-based Systems
- Windows Server 2003, Datacenter Edition for Itanium-based Systems
- Windows Server 2003, Enterprise x64 Edition
- Windows Server 2003, Datacenter x64 Edition

Use the following table to determine when to use these templates.

| Object and domain membership | Certificate template | Certificate purpose | Preferred certificate enrollment method | Alternate certificate enrollment method |
|---|---|---|---|---|
| VPN or IAS server, domain member | RAS and IAS server | Server Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| Windows XP client, domain member | Workstation Authentication | Client Authentication | Auto-enrollment | Request a certificate with the Certificates snap-in |
| VPN or IAS server, non-domain member | RAS and IAS server | Server Authentication | CA Web enrollment tool | Install from a floppy disk |
| Windows XP client, non-domain member | Workstation Authentication | Client Authentication | CA Web enrollment tool | Install from a floppy disk |

**Important**

If your server running IAS is not a domain controller but is a member of a domain with a Windows 2000 mixed functional level, you must add the server to the access control list (ACL) of the **RAS and IAS server** certificate template. You must also configure the correct permissions for autoenrollment. There are different procedures for adding single servers and groups of servers to the ACL.

To add an individual server to the ACL for the **RAS and IAS server** certificate template:

In Certificate Templates, select the template **RAS and IAS server**, and then add the IAS server to the template **Security** properties. For more information, see Allow subjects to request a certificate that is based on the template. After you have added your IAS server to the ACL, grant **Read**, **Enroll**, and **Autoenroll** permissions.

To manage a group of servers, add the servers to a new global or universal group, and then add the group to the ACL of the certificate template:

1. In Active Directory Users and Computers, create a new global or universal group for IAS servers. For more information, see Create a new group. Next, add to the group all computers that are IAS servers, are not domain controllers, and are members of a domain with a Windows 2000 mixed functional level. For more information, see Add a member to a group.

2. In Certificate Templates, select the **RAS and IAS server** template, and then add the group you created to the template **Security** properties by completing the steps in Allow subjects to request a certificate that is based on the template. After you have added your new group, grant **Read**, **Enroll**, and **Autoenroll** permissions.

For more information, see Domain and forest functionality.

**CA Web enrollment services**

A set of CA Web pages is provided with Certificate Services in the Windows Server 2003 family. These Web enrollment pages allow you to connect to the CA through a Web browser and perform common tasks, such as requesting certificates from a CA.

For more information, see Certification authority Web enrollment services and Set up certification authority Web enrollment support.

After your PKI has been designed and deployed and you are assigned the required permissions on the certificate templates, you can request certificates for computers by using the CA Web enrollment pages. You can request a certificate for a non-domain member computer using the CA Web enrollment tool by completing the following steps. You must be an administrator on the local computer to complete these steps.

**Install a computer certificate on the local computer**

1. Use your Web browser to log onto the Web enrollment tool at **http://***serverName***/certsrv**, where *serverName* is the name of the server hosting the CA, the CA Web pages, or both. Click **Request a certificate**.

2. In **Request a certificate**, click **Advanced Certificate Request**.

3. In **Advanced Certificate Request**, click **Create and submit a request to this CA**. A Web form that contains the sections **Certificate Template**, **Key Options**, and **Additional Options** becomes available. Select the following:

   - **Certificate template:** Administrator
   - **Key options:** Create new key set
   - **CSP:** Microsoft Enhanced Cryptographic provider v1.0
   - **Key usage:** Exchange
   - **Key size:** 1024
   - **Automatic key container name:** Selected
   - **Mark keys as exportable:** Selected
   - **Use local machine store:** Selected
   - **Request format:** PKCS 10
   - **Hash:** SHA-1
   - **Friendly name:** VPN server name

4. Click **Submit**.

5. Click **Install this certificate**

   The certificate you requested is installed on your local computer.

**Save a computer certificate to floppy disk**

1. Use your Web browser to log onto the Web enrollment tool at **http://***serverName***/certsrv**, where *serverName* is the name of the server hosting the CA, the CA Web pages, or both. Click **Request a certificate**.

2. In **Request a certificate**, click **Advanced Certificate Request**.

3. In **Advanced Certificate Request**, click **Create and submit a request to this CA**. A Web form that contains the sections **Certificate Template**, **Key Options**, and **Additional Options** becomes available. Select the following:

   - **Certificate template:** Administrator

- **Key options:** Create new key set
- **CSP:** Microsoft Enhanced Cryptographic provider v1.0
- **Key usage:** Exchange
- **Key size:** 1024
- **Automatic key container name:** Selected
- **Mark keys as exportable:** Selected
- **Export keys to file:** Selected
- **Full path name:** The path location where you want to save the file and the file name (for example, *\\serverName\ShareName\FileName.pvk*.).
- **Request format:** PKCS 10
- **Hash:** SHA-1
- **Friendly name:** VPN server name

4. Click **Submit**.
5. Create and confirm your private key password.
6. In **Certificate Issued**, select **DER encoded** and **Download certificate chain**.
7. Insert a floppy disk in your disk drive.
8. When prompted by your browser, select **Save this file to disk**, and then click **OK**.
9. Browse to your floppy drive, and then click **OK**.

   The certificate you requested is downloaded and saved to floppy disk.

For more information, see Submit an advanced certificate request via the Web.

For information about installing a certificate from floppy to the certificate store of a local computer, see Import a certificate.

**Note**

- If you have issued a certificate to your IAS server that has a blank Subject, the certificate is not available to authenticate your IAS server. To change this, you can use Certificate Templates to create a new certificate for enrollment on your IAS server. In the certificate properties, on the **Subject Name** tab, in **Subject name format**, select a value other than **None**.

# Securing IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing IAS**

This section covers:

- IAS and firewalls
- Shared secrets
- Message Authenticator attribute
- Remote access account lockout
- IAS Network Access Quarantine Control
- Securing RADIUS traffic with IPSec

# IAS and firewalls

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS and firewalls**
In the most common configuration, the firewall is connected to the Internet and the IAS server is another intranet resource that is connected to the perimeter network.

To reach the domain controller within the intranet, the IAS server might have:

- An interface on the perimeter network and an interface on the intranet (IP routing is not enabled).

- A single interface on the perimeter network. In this configuration, IAS communicates with intranet domain controllers through another firewall that connects the perimeter network to the intranet. This configuration is shown in the following drawing.




**Configuring the Internet firewall**
The firewall that is connected to the Internet must be configured with input and output filters on its Internet interface (and, optionally, its network perimeter interface), to allow the forwarding of RADIUS messages between the IAS server and RADIUS clients or proxies on the Internet. Additional filters can be used to allow the passing of traffic to Web servers, VPN servers, and other types of servers on the perimeter network.

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

**Filters on the Internet interface**
Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1812 (0x714).

  This filter allows RADIUS authentication traffic from Internet-based RADIUS clients to the IAS server. This is the default UDP port that is used by IAS, as defined in RFC 2865. If you are using a different port, substitute that port number for 1812, which is used in this example.

- Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1813 (0x715).

  This filter allows RADIUS accounting traffic from Internet-based RADIUS clients to the IAS server. This is the default UDP port that is used by IAS, as defined in RFC 2866. If you are using a different port, substitute that port number for 1813, which is used in this example.

- (Optional) Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1645 (0x66D).

  This filter allows RADIUS authentication traffic from Internet-based RADIUS clients to the IAS server. This is the UDP port that is used by older RADIUS clients.

- (Optional) Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1646 (0x66E).

  This filter allows RADIUS accounting traffic from Internet-based RADIUS clients to the IAS server. This is the UDP port that is used by older RADIUS clients.

Configure the following output filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the IAS server's perimeter network interface and UDP source port of 1812 (0x714).

  This filter allows RADIUS authentication traffic from the IAS server to Internet-based RADIUS clients. This is the default UDP port that is used by IAS, as defined in RFC 2865. If you are using a different port, substitute that port number for 1812, which is used in this example.

- Source IP address of the IAS server's perimeter network interface and UDP source port of 1813 (0x715).

  This filter allows RADIUS accounting traffic from the IAS server to Internet-based RADIUS clients. This is the default UDP port that is used by IAS, as defined in RFC 2866. If you are using a different port, substitute that port number for 1813, which is used in this example.

- (Optional) Source IP address of the IAS server's perimeter network interface and UDP source port of 1645 (0x66D).

  This filter allows RADIUS authentication traffic from the IAS server to Internet-based RADIUS clients. This is the UDP port that is used by older RADIUS clients.

- (Optional) Source IP address of the IAS server's perimeter network interface and UDP source port of 1646 (0x66E).

  This filter allows RADIUS accounting traffic from the IAS server to Internet-based RADIUS clients. This is the UDP port that is used by older RADIUS clients.

**Filters on the perimeter network interface**

Configure the following input filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the IAS server's perimeter network interface and UDP source port of 1812 (0x714).

  This filter allows RADIUS authentication traffic from the IAS server to Internet-based RADIUS clients. This is the default UDP port that is used by IAS, as defined in RFC 2865. If you are using a different port, substitute that port number for 1812, which is used in this example.

- Source IP address of the IAS server's perimeter network interface and UDP source port of 1813 (0x715).

  This filter allows RADIUS accounting traffic from the IAS server to Internet-based RADIUS clients. This is the default UDP port that is used by IAS, as defined in RFC 2866. If you are using a different port, substitute that port number for 1813, which is used in this example.

- (Optional) Source IP address of the IAS server's perimeter network interface and UDP source port of 1645 (0x66D).

  This filter allows RADIUS authentication traffic from the IAS server to Internet-based RADIUS clients. This is the UDP port that is used by older RADIUS clients.

- (Optional) Source IP address of the IAS server's perimeter network interface and UDP source port of 1646 (0x66E).

  This filter allows RADIUS accounting traffic from the IAS server to Internet-based RADIUS clients. This is the UDP port that is used by older RADIUS clients.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1812 (0x714).

  This filter allows RADIUS authentication traffic from Internet-based RADIUS clients to the IAS server. This is the default UDP port that is used by IAS, as defined in RFC 2865. If you are using a different port, substitute that port number for 1812, which is used in this example.

- Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1813 (0x715).

  This filter allows RADIUS accounting traffic from Internet-based RADIUS clients to the IAS server. This is the default UDP port that is used by IAS, as defined in RFC 2866. If you are using a different port, substitute that port number for 1813, which is used in this example.

- (Optional) Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1645 (0x66D).

  This filter allows RADIUS authentication traffic from Internet-based RADIUS clients to the IAS server. This is the UDP port that is used by older RADIUS clients.

- (Optional) Destination IP address of the IAS server's perimeter network interface and UDP destination port of 1646 (0x66E).

  This filter allows RADIUS accounting traffic from Internet-based RADIUS clients to the IAS server. This is the UDP port that is used by older RADIUS clients.

For added security, you can use the IP addresses of each RADIUS client that sends the packets through the firewall to define specific filters for traffic between the client and the IP address of the IAS server on the perimeter network.

**Configuring the intranet firewall**
The firewall that is connected to the intranet must be configured with input and output filters on its perimeter network interface (and, optionally, its intranet interface), to allow the forwarding of RADIUS messages between the IAS server on the perimeter network and domain controllers in the intranet. Additional filters can allow the passing of traffic to Web, VPN, and other types of servers on the perimeter network.

Separate input and output packet filters can be configured on the perimeter network interface and the intranet interface.

**Filters on the perimeter network interface**
Configure the following input packet filters on the perimeter network interface of the intranet firewall to allow the following types of traffic:

- Source IP address of the IAS server's perimeter network interface.

  This filter allows traffic from the IAS server on the perimeter network.

Configure the following output filters on the perimeter network interface of the intranet firewall to allow the following types of traffic:

- Destination IP address of the IAS server's perimeter network interface.

  This filter allows traffic to the IAS server on the perimeter network.

**Filters on the intranet interface**
Configure the following input filters on the intranet interface of the firewall to allow the following types of traffic:

- Destination IP address of the IAS server's perimeter network interface.

  This filter allows traffic to the IAS server on the perimeter network.

Configure the following output packet filters on the intranet interface of the firewall to allow the following types of traffic:

- Source IP address of the IAS server's perimeter network interface.

  This filter allows traffic from the IAS server on the perimeter network.

# Shared secrets

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shared secrets**

A shared secret is a text string that serves as a password between:

- A RADIUS client and RADIUS server.
- A RADIUS client and a RADIUS proxy.
- A RADIUS proxy and a RADIUS server.

For a configuration that uses a RADIUS client, a RADIUS proxy, and a RADIUS server, the shared secret that is used between the RADIUS client and the RADIUS proxy can be different than the shared secret used between the RADIUS proxy and the RADIUS server.

Shared secrets are used to verify that RADIUS messages, with the exception of the Access-Request message, are sent by a RADIUS-enabled device that is configured with the same shared secret. Shared secrets also verify that the RADIUS message has not been modified in transit (message integrity). The shared secret is also used to encrypt some RADIUS attributes, such as User-Password and Tunnel-Password. To provide verification for Access-Request messages, you can enable use of the RADIUS Message Authenticator attribute for both the RADIUS client configured on the IAS server and the access server. For more information, see Message Authenticator attribute and Configure the Message Authenticator attribute and shared secret.

When creating and using a shared secret:

- You must use the same case-sensitive shared secret on both RADIUS devices.
- Use a different shared secret for each RADIUS server-RADIUS client pair.
- To ensure a random shared secret, generate a random sequence at least 22 characters long.
- You can use any standard alphanumeric and special characters.
- You can use a shared secret of up to 128 characters in length. To protect your IAS server and your RADIUS clients from brute force attacks, use long shared secrets (more than 22 characters).
- Make the shared secret a random sequence of letters, numbers, and punctuation and change it often to protect your IAS server and your RADIUS clients from dictionary attacks. Shared secrets should contain characters from each of the following three groups:

| Group | Examples |
|---|---|
| Letters (uppercase and lowercase) | A, B, C and a, b, c |
| Numerals | 0, 1, 2, 3 |
| Symbols (all characters not defined as letters or numerals) | Exclamation point (!), asterisk (*), colon (:) |

The stronger your shared secret, the more secure are the attributes (for example, those used for passwords and encryption keys) that are encrypted with it. An example of a strong shared secret is 8d#>9fq4bV)H7%a3-zE13sW.

**Notes**

- When Password Authentication Protocol (PAP) is used between an access client and an access server (a RADIUS client), the access server encrypts the PAP password by using the shared secret and sends it in an Access-Request packet. If the access server sends the Access-Request message to a RADIUS proxy, the RADIUS proxy must first decrypt the PAP

password with the shared secret that was used between the RADIUS proxy and the access server. Next, it encrypts the PAP password with the shared secret that was used between the RADIUS proxy and the RADIUS server before forwarding the Access-Request message. Because a malicious user or process at a RADIUS proxy can record user names and passwords for PAP connections after they are decrypted but before they are encrypted, the use of PAP is highly discouraged.

- If you specify RADIUS clients by using an IP address range, all RADIUS clients within the address range must use the same shared secret.

- If you are using a password-based authentication method, it is strongly recommended that you use MS-CHAP v2, MS-CHAP, or CHAP with strong passwords to provide password protection from dictionary attacks.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Message Authenticator attribute

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Authenticator attribute**

When you configure IAS for a RADIUS client, you configure the IP address of the client. If an incoming RADIUS Access-Request message does not originate from at least one of the IP addresses of configured clients, IAS automatically discards the message, providing protection for an IAS server. However, source IP addresses can be spoofed (substituted with other IP addresses).

To provide protection from spoofed Access-Request messages and RADIUS message tampering, each RADIUS message can be additionally protected with the RADIUS Message Authenticator attribute, which is described in RFC 2869, "RADIUS Extensions."

The RADIUS Message Authenticator attribute is a Message Digest 5 (MD5) hash of the entire RADIUS message. The shared secret is used as the key. If the RADIUS Message Authenticator attribute is present, it is verified. If it fails verification, the RADIUS message is discarded. If the client settings require the Message Authenticator attribute and it is not present, the RADIUS message is discarded.

For information about how to configure the use of the Message Authenticator attribute for RADIUS clients of an IAS server, see Configure the Message Authenticator attribute and shared secret. For information about how to configure the use of the Message Authenticator attribute for the Routing and Remote Access service, see Use RADIUS authentication.

# Remote access account lockout

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access account lockout**

You can use the remote access account lockout feature to specify how many times a remote access authentication fails against a valid user account before the user is denied access. Remote access account lockout is especially important for remote access virtual private network (VPN) connections over the Internet. An attacker on the Internet can attempt to access an organization intranet by sending credentials (valid user name, guessed password) during the VPN connection authentication process. During a dictionary attack, the attacker sends hundreds or thousands of credentials by using a list of passwords based on common words or phrases.

With remote access account lockout enabled, a dictionary attack is thwarted after a specified number of failed attempts. As the network administrator, you must decide on two remote access account lockout variables:

1. The number of failed attempts before future attempts are denied.

   After each failed attempt, a failed attempts counter for the user account is incremented. If the user account's failed attempts counter reaches the configured maximum, future attempts to connect are denied.

   A successful authentication resets the failed attempts counter when its value is less than the configured maximum. In other words, the failed attempts counter does not accumulate beyond a successful authentication.

2. How often the failed attempts counter is reset.

   The failed attempts counter is periodically reset to 0. If an account is locked out after the maximum number of failed attempts, the failed attempts counter is automatically reset to 0 after the reset time.

You enable the remote access account lockout feature by changing settings in the registry on the computer that provides the authentication. If the remote access server is configured for Windows Authentication, modify the registry on the remote access server computer. If the remote access server is configured for RADIUS authentication and Internet Authentication Service (IAS) is being used, modify the registry on the IAS server.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

To enable remote access account lockout, you must set the **MaxDenials** entry in the registry to **1** or greater. **MaxDenials** is the maximum number of failed attempts before the account is locked out. You set the **MaxDenials** entry in the following registry subkey:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout**

By default, **MaxDenials** is set to **0**, which means that remote access account lockout is disabled.

To modify the amount of time before the failed attempts counter is reset, you must set the **ResetTime (mins)** entry in the registry to the required number of minutes. You set the **ResetTime (mins)** entry in the following registry subkey:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout**

By default, **ResetTime (mins)** is set to **0xb40**, or 2,880 minutes (48 hours).

**Manually resetting an account that is locked out**

To manually reset a user account that has been locked out before the failed attempts counter is automatically reset, delete the following registry subkey that corresponds to the user's account name:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout\ *domain name:user name***

When the lockout count for a user account is reset to 0 due to either a successful authentication or an automatic reset, the registry subkey for the user account is deleted.

**Notes**

- Remote access account lockout is not related to the **Account locked out** setting on the **Account** tab on the properties of a user account.

- The remote access account lockout feature does not distinguish between malicious users who attempt to access your intranet and authentic users who attempt remote access but have forgotten their current passwords. Users who have forgotten their current password typically try the passwords that they remember and, depending on the number of attempts and the **MaxDenials** setting, may have their accounts locked out.

- If you enable the remote access account lockout feature, a malicious user can deliberately force an account to be locked out by attempting multiple authentications with the user account until the account is locked out, thereby preventing the authentic user from being able to log on.

# IAS Network Access Quarantine Control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS Network Access Quarantine Control**

IAS Network Access Quarantine Control provides phased network access for remote client computers by restricting them to a quarantine mode. After the client computer configuration is either brought into or determined to be in accordance with your organization's network policy, quarantine restrictions, which consist of Quarantine IP-Filters and Session Timers, are removed and standard remote access policy is applied to the connection.

Network Access Quarantine Control provides protection when users in your organization accidentally reconfigure key settings and do not restore them before connecting to your network. For example, a user might disable antivirus software that is required while connected to your network. Although Network Access Quarantine Control does not protect against attackers, computer configurations for authorized users can be verified and, if necessary, corrected before they can access the network. A timer setting is also available, which you can use to specify an interval at which the connection is dropped if the client fails to meet configuration requirements.

You can use the Routing and Remote Access service to process the specific RADIUS options sent by IAS, complete any required client configuration work, and remove the quarantine condition (or drop the connection) based on success or failure.

**Quarantine mode**

Quarantine mode is a set of network restrictions that are configured in IAS and implemented by the remote access server for each connection. You can use a Quarantine IP Filter to restrict access to a specific set of servers (for example, servers on a specified virtual Local Area Network) and a Quarantine Session Timer to restrict the amount of time the client can remain connected in quarantine mode. You can set these filters in the IAS console.

**Configuring Network Access Quarantine Control**

You can implement Network Access Quarantine Control with one or more servers running Windows Server 2003 and remote access, one or more servers running Windows Server 2003 and IAS, a Connection Manager (CM) profile created with Connection Manager Administration Kit (CMAK), an administrator-provided script, and two additional components: the notifier component and the listener component.

The notifier component is included in the CM profile and installed on the client computer. The notifier component sends notification to the remote access server when the administrator-provided script has run successfully on the client.

The listener component is installed on the remote access server, and receives notification from the notifier component that the script on the client has successfully performed all configuration checks. After the listener component receives notification, it removes the client from quarantine mode, and the remote access server applies standard remote access policy to the client.

You can create your own notifier and listener components, or you can use Rqs.exe (a listener component also called the Remote Access Quarantine Service) and Rqc.exe (a notifier component), which are available in the Windows Deployment and Resource Kits. For more information, see Windows Resource Kit Tools Help.

If you are running Windows Server 2003 with Service Pack 1, you can install the Remote Access Quarantine Service (Rqs.exe) using the Windows Components Wizard. In the wizard, the Remote Access Quarantine Service is available in Networking Services. In addition, you can use the Windows Components Wizard to install CMAK. After you install CMAK, Rqc.exe is located on your hard drive at *%systemroot%\Program Files\cmak\support*.

When you create the CM profile, you can include the administrator-provided script and Rqc.exe, which are distributed to and installed on remote access client computers. This profile can be installed on the following client operating systems:

- Windows XP Professional
- Windows XP Home Edition
- Windows 2000 Professional
- Windows Millennium Edition
- Windows 98 Second Edition

For more information about CMAK, see Connection Manager Administration Kit.

**Caution**

- Placing all remote access clients in quarantine mode without a way to remove quarantine policy and apply full access policy might prevent all remote access clients from establishing network connections.

Before you implement a remote access policy with Quarantine IP-Filters and Session Timers on your network, you must complete the following steps:

1. Create a client-side script that validates client configuration information.

2. Create a notification component that provides verification to the remote access server that the script has successfully run. If you do not want to create a notification component, you can use Rqc.exe.

3. Create a listener component to install on remote access servers (that can receive information from the notification component), and then remove the client from quarantine mode, applying the full access policy. If you do not want to create a listener component, you can use Rqs.exe.

4. Create a CM profile with CMAK. Include the client-side script and the notification component in the profile.

5. Distribute the CM profile for installation on remote access client computers.

**Configure two remote access policies on remote access or IAS servers**

On your remote access servers or IAS servers, use the **New Remote Access Policy Wizard** to create these remote access policies:

- A full access policy

- A quarantine policy

The notification and listener components of Network Access Quarantine Control use port 7250 by default. When you create the quarantine policy, you must configure **Inbound Filters** (also called **Input Filters**) to allow network traffic on port 7250, or the notification component (Rqc.exe), which runs on client computers, cannot notify the remote access server listener component (Rqs.exe) that the script has run successfully. You can specify another port, but you must also configure the listener and notification components to use this new port.

When you use the **New Remote Access Policy Wizard** to create quarantine policy, do the following to add quarantine IP filters and session timers:

1. On the **profile** page, click **Edit profile**.

2. In **Edit Dial-In profile**, click the **Advanced** tab, and then click **Add**.

**To add a quarantine IP filter**

1. In **Attribute**, select **MS-Quarantine-IPFilter**, and then click **Add**.

2. In **IP Filter Attribute Information**, click **Input Filters**.

3. In **Inbound Filters**, click **New**.

4. In **Add IP Filter**, click **Destination network**, and then type values in **IP address** and **Subnet mask** that match your network configuration.

5. In **protocol**, select **TCP**.

6. In **Source port**, type **7250**.

7. In **Destination port**, type **7250**, and then click **OK**.

8. In **Inbound Filters**, for **Filter action**, click **Permit only the packets listed below**, and then click **OK**.

**To add a quarantine session timer**

- In **Attribute**, select **MS-Quarantine-Session-Timeout**, and then click **Add**.

- In **Attribute value**, type the number of seconds that represents the maximum amount of time (in seconds) that you want client computers to remain connected in quarantine mode, and then click **OK**.

After you have configured both the full access and quarantine policies, make sure that the quarantine policy is first in the list of remote access policies.

For more information about the **New Remote Access Policy Wizard**, see Add a remote access policy.

### Notes

- IAS Network Access Quarantine Control must be used in combination with remote access policy (user-based authentication and access control) to provide security for your network. The quarantine filters and timers can be defined in a remote access policy on either the remote access server or the IAS server. When the connection is made, the access server implements the restrictions (Quarantine IP Filters and Quarantine Session Timers) that are returned by IAS.

- Notification by the client (to the server) that the administrator-provided script has successfully run is not secure and can be spoofed by a malicious user or attacker. Network Access Quarantine Control is only intended to help verify a managed client computer configuration.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Securing RADIUS traffic with IPSec

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing RADIUS traffic with IPSec**

Internet Protocol security (IPSec) provides you with the ability to secure RADIUS servers against unwanted traffic by filtering on specific network adapters (allowing or blocking specific protocols) and enabling you to choose source IP addresses from which traffic is allowed. For organizational units, you can create IPSec policies, which are stored in Active Directory. Or, you can create local policies on RADIUS servers, and apply these policies to specific computers. If you create IPSec policies for an organizational unit, the policy is applied through Group Policy.

For more information about organizational units and Group Policy, see Organizational units, Group Policy integration, and Create a new organizational unit.

Before you create IPSec filters, determine the type of traffic that you want to allow for each RADIUS server. Filters that are too strict might block acceptable network traffic. For example, if IAS is installed on a domain controller and all IP traffic, except RADIUS traffic, is blocked on all ports, user queries for Active Directory objects (such as printers) on default global catalog port 3268 will fail. Conversely, IP filters that are too general expose the RADIUS server to unwanted traffic. For more information, see Special IPSec considerations and Filter list.

RADIUS messages are sent with the User Datagram Protocol (UDP). UDP port 1812 is used for RADIUS authentication messages and UDP port 1813 is used for RADIUS accounting messages. When you create inbound and outbound filters with IPSec, UDP traffic must be allowed on these ports. However, some network access servers might use UDP port 1645 for RADIUS authentication messages and UDP port 1646 for RADIUS accounting messages. By default, IAS supports both sets of ports. If your network access servers use UDP ports 1645 and 1646, you can create IPSec filters that allow traffic on these ports. For information about changing the UDP ports that IAS uses, see Configure IAS port information. For more information about defining IP filter lists, see Internet Protocol Security (IPSec) and Define IP Filter Lists.

You can use the IP Security Policies snap-in and Group Policy to configure IPSec policy for organizational units. For more information, see Group Policy.

You can use the Local Security Policy console to configure IPSec policy for individual RADIUS servers. For more information, see Local Security Policy.

For information about how to create IP filters for remote access clients, see IAS Network Access Quarantine Control.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Vendor-Specific Attributes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Vendor-specific attributes**
This section covers:

- Vendor-specific attribute overview
- Sample VSA for a Cisco NAS
- Sample VSA for a US Robotics NAS

# Vendor-specific attribute overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Vendor-specific attribute overview**

IAS provides the ability to specify RADIUS attributes that are returned with a RADIUS response message. These RADIUS attributes can be specified for each remote access policy and are configured on the **Advanced** tab in the properties of a profile for a remote access policy. For more information, see Configure Attributes for a Profile.

In addition to the standard RADIUS attributes, which are described in RFCs 2865 and 2866, you might also need to configure vendor-specific attributes (VSAs). VSAs allow vendors to support their own proprietary RADIUS attributes that are not included in RFCs 2865 and 2866. IAS includes VSAs from a number of vendors in its dictionary; however, the VSAs provided with IAS in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition do not include VSAs for all vendors. There is no process for centrally updating and distributing new versions of the IAS dictionary. Instead, you can manually create new VSAs for each remote access policy profile.

Some network access server (NAS) manufacturers use vendor-specific attributes (VSAs) to provide functionality that is not supported in standard attributes. IAS enables you to create or edit VSAs to take advantage of proprietary functionality supported by some NAS vendors. For more information about vendor-specific functionality and the VSAs that you can configure, see your access server documentation.

Before adding a VSA, check the list of attributes in the IAS dictionary. If the required VSA is present, use it. If not, you must add a VSA to the profile. VSAs are added by specifying the **Vendor-Specific** attribute on the **Advanced** tab in the properties of a profile for a remote access policy. For more information, see Configure vendor-specific attributes for a remote access policy.

If you want to add VSAs to the remote access policy profile to support proprietary functionality for your access point, you must determine whether the VSA conforms to the format recommended in RFC 2865.

- If it does, you must specify:
    - A network access vendor by either name or vendor code.
    - A vendor-assigned attribute number.
    - The attribute format (that is, the type of data, such as string or hexadecimal).
    - The attribute value.
- If it does not, you must specify:
    - A network access vendor by either name or vendor code.
    - A hexadecimal attribute value that represents the attribute data.

**Ordering of multiple VSAs**

If you need to configure more than one VSA, you should arrange them in the appropriate order in the **Multivalued Attribute Information** dialog box. For example, if you are using a filtering VSA that automatically disconnects users who do not meet specific criteria, you should ensure that the attribute appears at the top of the list. Use the **Move Up** and **Move Down** buttons in the **Multivalued Attribute Information** dialog box to arrange the attributes in the correct order.

# Sample VSA for a Cisco NAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sample VSA for a Cisco NAS**

This sample Cisco vendor-specific attribute conforms to the format for vendor-specific attributes (type 26) that is specified in RFC 2865.

The following table lists the information used to configure a Cisco VSA to specify a primary DNS server with an IP address of 10.10.10.10.

| Information | Description |
|---|---|
| Vendor ID | This is the unique ID for Cisco. When you select Cisco as the network access server vendor, this ID (9) is automatically supplied. |
| Cisco-assigned attribute number | This is the vendor-type number (1) for vendor-specific attributes that take the attribute-value pair form, specified in Cisco documentation as cisco-avpair. |
| Syntax | If the attribute is mandatory, the syntax is *Protocol***:***Attribute* **=** *Value*. If the attribute is optional, the attribute-value pair is separated by an asterisk (*****) instead of an equal sign (**=**). In this example, *Protocol* is a value of the Cisco protocol attribute for a specific type of authorization. *Attribute* and *Value* represent an appropriate attribute/value (AV) pair defined in the Cisco TACACS+ specification. This allows the full set of features available for TACACS+ authorization to be used for RADIUS. The Cisco format, used to specify a primary DNS server, is **ip:dns-servers=10.10.10.10**. |

To specify the primary DNS server of 10.10.10.10, configure an RFC-compliant VSA:

1. Open the IAS console. To open the IAS console, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Authentication Service**.
2. In the console tree, click **Remote Access Policies**.
3. Right-click the policy for which you want to configure a vendor-specific attribute (VSA), and then click **Properties**.
4. Click **Edit Profile**, click the **Advanced** tab, and then click **Add**.
5. In the **Add Attribute** dialog box, scroll the **Attribute** list and double-click **Vendor-Specific**. The **Multivalued Attribute Information** dialog box opens. Click **Add**.
6. The **Vendor-Specific Attribute Information** dialog box opens. In **Select from list**, select **Cisco** as the network access server vendor, and then click **Yes, it conforms**.
7. Click **Configure Attribute**. The **Configure VSA (RFC-compliant)** dialog box opens.
8. In **Vendor-assigned attribute number**, type **1**.
9. In **Attribute format**, select **String**.
10. In **Attribute Value**, type **ip:dns-servers=10.10.10.10**.

For more information, see Configure vendor-specific attributes for a remote access policy. For more information about Cisco VSAs, see your Cisco documentation.

# Sample VSA for a US Robotics NAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sample VSA for a US Robotics NAS**

This sample US Robotics vendor-specific attribute does not conform to the format for vendor-specific attributes (VSAs) that is specified in RFC 2865. Therefore, this US Robotics VSA must be entered in hexadecimal format.

The following table lists the information that is used to configure a US Robotics VSA to specify a primary DNS/NBNS server with an IP address of 10.10.10.10.

| Information | Description |
|---|---|
| Vendor ID | This is the unique ID for US Robotics. When you select US Robotics as the network access server vendor, this ID (429) is automatically supplied. |
| US Robotics VSA Indicator | The VSA indicator for the primary DNS/NBNS server is 0x900F. |
| String | The hexadecimal encoding of the string 10.10.10.10. (the IP address of a primary DNS/NBNS server) is 0x31302E31302E31302E31302E. |

To specify an IP address of 10.10.10.10 for a primary DNS/NBNS server:

- Configure a non-RFC-compliant VSA.

  For more information, see Configure vendor-specific attributes for a remote access policy.

  - Select **US Robotics** as the network access server vendor.
  - In **Hexadecimal attribute value**, in the **Configure VSA (non-RFC-compliant)** dialog box, type **0x900f31302E31302E31302E31302E**.

For more information about US Robotics VSAs, see your US Robotics documentation.

# Authentication Methods for use with IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication methods**
The authentication of access clients is an important security concern. Authentication methods typically use an authentication protocol that is negotiated during the connection establishment process. IAS also supports unauthenticated connections.

For information about using authentication methods, see the following:

- PEAP
- EAP
- MS-CHAP
- MS-CHAP version 2
- CHAP
- PAP
- Unauthenticated access

Each authentication method has advantages and disadvantages in terms of security, usability, and breadth of support. The authentication method used is determined by the configuration of both the access server and client. Consult your access server documentation to determine which authentication methods are supported.

You can configure IAS to accept multiple authentication methods. You can also configure your access servers to attempt to negotiate a connection by using the most secure protocol first, and then the next most secure, and so on down to the least secure. For example, the Routing and Remote Access service tries to negotiate a connection using EAP first, then MS-CHAP v2, then MS-CHAP, then CHAP, then SPAP, and then PAP. When EAP is chosen as the authentication method, the negotiation of the EAP type occurs between the access client and the IAS server.

In addition, you can use remote access policies for various authentication methods, depending on the type of client being authenticated. For example, you can create two remote access policies, one for virtual private network (VPN) clients and one for wireless clients--each of which uses a different authentication method. The remote access policy for VPN clients can be configured to use EAP-TLS with smart cards or certificates as the authentication method and authentication type, while the remote access policy for wireless clients can be configured to use PEAP-EAP-MS-CHAP v2, which provides secure password authentication.

**Certificate deployments and Active Directory replication**
Some authentication methods, such as PEAP and EAP, can use certificates for authentication of computers and users. Latency in Active Directory replication might temporarily affect the ability of a client or server to obtain a certificate from a certification authority (CA). If a computer configured to use certificates for authentication cannot enroll a certificate, authentication fails.

For example, if a new computer running IAS is joined to the domain by a local domain controller, the IAS server is rebooted and Group Policy is applied. After Group Policy is applied and certificate autoenrollment is configured, the IAS server attempts to enroll a certificate with a CA. The CA queries Active Directory to determine whether the IAS server has the security permissions required to enroll a certificate. If the domain controller that joined the IAS server to the domain has replicated IAS server domain membership information across the domain, the CA can verify that the IAS server should receive a certificate. The CA then enrolls a certificate to the IAS server. You can manually refresh Group Policy by logging onto the domain or by running the **gpupdate** command.

When information about the IAS server domain membership has not replicated to other domain controllers and the CA queries a domain controller that does not yet have information about the IAS server, the CA cannot verify that the IAS server is running on a computer that is a domain member. Because the CA cannot verify that the IAS server has the security permissions to enroll a certificate, the CA does not enroll a certificate for the IAS server. As the result, authentication fails and the server cannot log on to the network.

For more information, see Network access authentication and certificates and Remote Access Policies.

**Notes**

- IEEE 802.1x authentications for wireless and switch clients use only the PEAP and EAP authentication methods.
- IAS also supports authentication for computer accounts.
- When using the EAP-TLS authentication method with certificates, TLS uses cached certificate properties instead of reading the certificate from the certificate store. If a certificate is either changed or deleted and replaced by a new certificate, TLS continues using outdated cached certificate information until the cache expires or is refreshed. If you change or replace a certificate, you can refresh the TLS cache by restarting the server computer.

# PEAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PEAP**

Protected Extensible Authentication Protocol (PEAP) is a new member of the family of Extensible Authentication Protocol (EAP) protocols. PEAP uses Transport Level Security (TLS) to create an encrypted channel between an authenticating PEAP client, such as a wireless computer, and a PEAP authenticator, such as an Internet Authentication Service (IAS) or Remote Authentication Dial-In User Service (RADIUS) server. PEAP does not specify an authentication method, but provides additional security for other EAP authentication protocols, such as EAP-MSCHAPv2, that can operate through the TLS encrypted channel provided by PEAP. PEAP is used as an authentication method for 802.11 wireless client computers, but is not supported for virtual private network (VPN) or other remote access clients.

To enhance both the EAP protocols and network security, PEAP provides:

- Protection for the EAP method negotiation that occurs between client and server through a TLS channel. This helps prevent an attacker from injecting packets between the client and the network access server (NAS) to cause the negotiation of a less secure EAP method. The encrypted TLS channel also helps prevent denial of service attacks against the IAS server.
- Support for the fragmentation and reassembly of messages, allowing the use of EAP types that do not provide this.
- Wireless clients with the ability to authenticate the IAS or RADIUS server. Because the server also authenticates the client, mutual authentication occurs.
- Protection against the deployment of an unauthorized wireless access point (WAP) when the EAP client authenticates the certificate provided by the IAS server. In addition, the TLS master secret created by the PEAP authenticator and client is not shared with the access point. Because of this, the access point cannot decrypt the messages protected by PEAP.
- PEAP fast reconnect, which reduces the delay in time between an authentication request by a client and the response by the IAS or RADIUS server, and allows wireless clients to move between access points without repeated requests for authentication. This reduces resource requirements for both client and server.

**PEAP authentication process**

There are two stages in the PEAP authentication process between PEAP client and authenticator. The first stage sets up a secure channel between the PEAP client and the authenticating server. The second stage provides EAP authentication between the EAP client and authenticator.

**TLS encrypted channel**

The wireless client associates with a wireless access point. An IEEE 802.11-based association provides an Open System or Shared Key Authentication before a secure association is created between the client and access point. After the IEEE 802.11-based association is successfully established between the client and access point, the TLS session is negotiated with the access point. After authentication is successfully completed between the wireless client and the server (for example, an IAS server), the TLS session is negotiated between them. The key that is derived during this negotiation is used to encrypt all subsequent communication.

**EAP-authenticated communication**

Complete EAP communication, including EAP negotiation, occurs through the TLS channel. The IAS server authenticates the user and client computer with the method that is determined by the EAP type and selected for use within PEAP (either EAP-TLS or EAP-MS-CHAPv2). The access point only forwards messages between wireless client and RADIUS server--the access point (or a person monitoring it) cannot decrypt these messages because it is not the TLS end point.

**802.11 wireless deployments using PEAP**

You can choose between two EAP types for use with PEAP: EAP-MS-CHAPv2 or EAP-TLS. EAP-MS-CHAPv2 uses credentials (user name and password) for user authentication, and a certificate in the server computer certificate store for server authentication. EAP-TLS uses either certificates installed in the client computer certificate store or a smart card for user and client computer authentication, and a certificate in the server computer certificate store for server authentication.

**PEAP with EAP-MS-CHAPv2**

PEAP with EAP-MS-CHAPv2 (PEAP-EAP-MS-CHAPv2) is easier to deploy than EAP-TLS because user authentication is accomplished with password-based credentials (user name and password) instead of certificates or smart cards--only the IAS or RADIUS server is required to have a certificate. Additionally, the server certificate can be issued by a public certification authority (CA) that is trusted by the client computer (that is, the public CA certificate already exists in the Trusted Root Certification Authority folder on the client computer certificate store). In this case, the server certificate is not downloaded and added to the client trusted root certificate store, and the user is not prompted to make a decision about whether to trust the server.

PEAP-EAP-MS-CHAPv2 provides improved security over MS-CHAPv2 by using mutual authentication, preventing an unauthorized server from negotiating the least secure authentication method, and providing key generation with TLS. PEAP-EAP-MS-CHAPv2 requires that the client trust certificates provided by the server.

For server and client computer certificate requirements, see Network access authentication and certificates. For an example wireless access policy using PEAP-EAP-MS-CHAPv2, see Wireless access with secure password authentication.

**PEAP with EAP-TLS**

Public Key certificates provide a much stronger authentication method than those that use password-based credentials. PEAP with EAP-TLS (PEAP-EAP-TLS) uses certificates for server authentication and either certificates or smart cards for user and client computer authentication. To use PEAP-EAP-TLS, you must deploy a public key infrastructure (PKI).

For more information, see Network access authentication and certificates.

**PEAP fast reconnect**

PEAP fast reconnect enables wireless clients to move between wireless access points on the same network without being reauthenticated each time they associate with a new access point.

Wireless access points are configured as RADIUS clients to RADIUS servers. If a wireless client roams between access points that are configured as clients to the same RADIUS server, the client is not required to be authenticated with each new association. When a client moves to an access point that is configured as a RADIUS client to a different RADIUS server, although the client is reauthenticated, this process occurs much more efficiently.

PEAP fast reconnect reduces the response time for authentication between client and authenticator because the authentication request is forwarded from the new server to the original server. Because both the PEAP client and authenticator use previously cached TLS connection properties (the collection of which is named the TLS handle), the authenticator can quickly determine that the client connection is a reconnect.

If the original PEAP authenticator is unavailable, full authentication must occur between the client and the new authenticator. The new PEAP authenticator's TLS handle is cached by the client. The client can cache TLS handles for multiple PEAP authenticators. For smart carts or PEAP-EAP-MSCHAPv2 authentication, the user is asked to supply the PIN or credentials, respectively.

With PEAP-EAP-MS-CHAPv2 authentication:

| **When the new access point is a client to the same RADIUS server** | **When the new access point is a client to a new RADIUS server** |
|---|---|
| The user is not prompted for credentials each time the client computer associates with a new access point. | The user is prompted for credentials on this initial association. The next time the client computer associates with an access point that is a client to this server, user credentials are not required. |
| The RADIUS server is not required to provide a certificate. | The RADIUS server provides a certificate on this initial association so that the wireless client can authenticate to the RADIUS server. The next time the client computer associates with an access point that is a client to this server, the server is not required to be reauthenticated. |

With PEAP-EAP-TLS authentication:

| **When the new access point is a client to the same RADIUS server** | **When the new access point is a client to a new RADIUS server** |
|---|---|

| | |
|---|---|
| The client and server are not required to exchange certificates. | The client and server exchange certificates on this initial association. The next time the client computer associates with an access point that is a client to this server, certificates are not exchanged. |
| The user is not prompted for a smart card personal identification number (PIN) each time the client computer associates with a new access point. | The user is prompted for a smart card PIN on this initial association. The next time the client computer associates with an access point that is a client to this server, the user is not prompted for the PIN. |

For more information about RADIUS clients, see Components of a RADIUS infrastructure.

For more information about RADIUS proxy servers, see IAS as a RADIUS proxy.

To enable PEAP fast reconnect:

- Both the PEAP client (802.11 wireless client) and PEAP authenticator (RADIUS server) must have fast reconnect enabled.
- All access points to which the PEAP client roams must be configured as RADIUS clients to a RADIUS server (the PEAP authenticator) for which PEAP is configured as the authentication method for wireless connections.
- All access points to which the PEAP client associates must be configured to prefer the same RADIUS server (PEAP authenticator) in order to avoid being prompted for credentials from every RADIUS server. If the access point cannot be configured to prefer a RADIUS server, you can configure an IAS RADIUS proxy with a preferred RADIUS server.

For more information, see Configure PEAP and EAP methods.

**Notes**

- When using the PEAP-EAP-TLS and EAP-TLS authentication methods with certificates, TLS uses cached certificate properties instead of reading the certificate from the certificate store. If a certificate is either changed or deleted and replaced by a new one, TLS continues using outdated cached certificate information until the cache expires or is refreshed. If you change or replace a certificate, you can refresh the TLS cache by restarting the server computer.
- PEAP does not support guest authentication, which has a blank user name and password.
- When you deploy both PEAP and EAP unprotected by PEAP, do not use the same EAP authentication type with and without PEAP. For example, if you deploy PEAP with EAP-TLS (PEAP-EAP-TLS), do not also deploy EAP-TLS without PEAP. Deploying authentication methods with the same type -- one with and the other without the protection of PEAP -- creates a security vulnerability.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# EAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**EAP**

With the Extensible Authentication Protocol (EAP), an arbitrary authentication mechanism authenticates a remote access connection. The exact authentication scheme to be used is negotiated by the remote access client and the authenticator (either the remote access server or the Remote Authentication Dial-In User Service [RADIUS] server). Routing and Remote Access includes support for EAP-TLS and MD5-Challenge by default. You can plug in other EAP modules to the server running Routing and Remote Access to provide other EAP methods.

EAP allows for an open-ended conversation between the remote access client and the authenticator. The conversation consists of authenticator requests for authentication information and the responses by the remote access client. For example, when EAP is used with security token cards, the authenticator can separately query the remote access client for a name, PIN, and card token value. As each query is asked and answered, the remote access client passes through another level of authentication. When all questions have been answered satisfactorily, the remote access client is authenticated.

A specific EAP authentication scheme is known as an EAP type. Both the remote access client and the authenticator must support the same EAP type for successful authentication to occur.

The Windows Server 2003 family includes an EAP infrastructure, two EAP types, and the ability to pass EAP messages to a RADIUS server (EAP-RADIUS).

**EAP infrastructure**

EAP is a set of internal components that provide architectural support for any EAP type in the form of a plug-in module. For successful authentication, both the remote access client and authenticator must have the same EAP authentication module installed. The Windows Server 2003 family provides two EAP types: MD5-Challenge and EAP-TLS. EAP-TLS is only available for members of a domain. You can also install additional EAP types. The components for an EAP type must be installed on every remote access client and every authenticator.

**MD5-Challenge**

Message Digest 5 Challenge (MD5-Challenge) is a required EAP type that uses the same challenge handshake protocol as PPP-based CHAP, but the challenges and responses are sent as EAP messages.

A typical use for MD5-Challenge is to authenticate the credentials of remote access clients by using user name and password security systems. You can also use MD5-Challenge to test EAP interoperability.

**EAP-TLS**

EAP-Transport Level Security (EAP-TLS) is an EAP type that is used in certificate-based security environments. If you are using smart cards for remote access authentication, you must use the EAP-TLS authentication method. The EAP-TLS exchange of messages provides mutual authentication, negotiation of the encryption method, and encrypted key determination between the remote access client and the authenticator. EAP-TLS provides the strongest authentication and key determination method.

EAP-TLS is supported only on servers that are running Routing and Remote Access, that are configured to use Windows Authentication or RADIUS, and that are members of a domain. A remote access server running as a stand-alone server or a member of a workgroup does not support EAP-TLS.

For information about configuring smart cards for remote access clients, see Using smart cards for remote access.

**EAP-RADIUS**

EAP-RADIUS is not an EAP type, but the passing of EAP messages of *any* EAP type by an authenticator to a RADIUS server for authentication. For example, for a remote access server that is configured for RADIUS authentication, the EAP messages sent between the remote access client and remote access server are encapsulated and formatted as RADIUS messages between the remote access server and the RADIUS server.

EAP-RADIUS is used in environments where RADIUS is used as the authentication provider. An advantage of using EAP-RADIUS is that EAP types do not need to be installed at each remote access server, only at the RADIUS server. In the case of an IAS server, you only need to install EAP types on the IAS server.

In a typical use of EAP-RADIUS, a server running Routing and Remote Access is configured to use EAP and to use an IAS server for authentication. When a connection is made, the remote access client negotiates the use of EAP with the remote access server. When the client sends an EAP message to the remote access server, the remote access server encapsulates the EAP

message as a RADIUS message and sends it to its configured IAS server. The IAS server processes the EAP message and sends a RADIUS-encapsulated EAP message back to the remote access server. The remote access server then forwards the EAP message to the remote access client. In this configuration, the remote access server is only a pass-through device. All processing of EAP messages occurs at the remote access client and the IAS server.

Routing and Remote Access can be configured to authenticate locally, or to a RADIUS server. If Routing and Remote Access is configured to authenticate locally, all EAP methods will be authenticated locally. If Routing and Remote Access is configured to authenticate to a RADIUS server, all EAP messages will be forwarded to the RADIUS server with EAP-RADIUS.

For more information about configuring a server running Routing and Remote Access for EAP-RADIUS, see Configure EAP-RADIUS.

**Enabling EAP**

To enable EAP-based authentication, you must do the following:

1. Enable EAP as an authentication protocol on the remote access server. For more information, see Enable EAP.

2. Enable EAP and, if needed, configure the EAP type on the appropriate remote access policy. For more information, see Introduction to remote access policies and Configure authentication.

3. Enable and configure EAP on the remote access client. For more information, see Extensible Authentication Protocol (EAP).

**Note**

- Make sure your network access server (NAS) supports EAP before you enable it on a remote access policy on an IAS server. For more information, see your NAS documentation.

# MS-CHAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MS-CHAP**
The Windows Server 2003 family includes support for the Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), also known as MS-CHAP version 1. MS-CHAP is a nonreversible, encrypted password authentication protocol. The challenge handshake process works as follows:

1. The authenticator (the remote access server or the IAS server) sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string.

2. The remote access client sends a response that contains the user name and a nonreversible encryption of the challenge string, the session identifier, and the password.

3. The authenticator checks the response and, if valid, the user's credentials are authenticated.

If you use MS-CHAP as the authentication protocol, then you can use Microsoft Point-to-Point Encryption (MPPE) to encrypt the data sent on the PPP or PPTP connection.

The Windows Server 2003 family also includes support for MS-CHAP version 2. MS-CHAP version 2 provides stronger security for remote access connections than MS-CHAP. You should consider using MS-CHAP version 2 instead of MS-CHAP. For more information, see MS-CHAP version 2.

**Enabling MS-CHAP**
To enable MS-CHAP-based authentication, you must do the following:

1. Enable MS-CHAP as an authentication protocol on the remote access server. For more information, see Enable authentication protocols. MS-CHAP is enabled by default.

2. Enable MS-CHAP on the appropriate remote access policy. For more information, see Introduction to remote access policies and Configure authentication. MS-CHAP is enabled by default.

3. Enable MS-CHAP on the remote access client. For more information, see Microsoft Challenge Handshake Authentication Protocol (MS-CHAP).

### Notes

- MS-CHAP (version 1 and version 2) is the only authentication protocol provided with the Windows Server 2003 family that supports password change during the authentication process.

- By default, the Windows Server 2003 family implementation of MS-CHAP v1 does not support LAN Manager authentication. If you want to allow the use of LAN Manager authentication with MS-CHAP v1 for older operating systems such as Windows NT 3.5x and Windows 95, you must set the following registry value to 1 on the authenticating server:

  **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy\Allow LM Authentication**

  Windows 2000 Server supports LAN Manager authentication by default. Upgrading a computer running Windows 2000 Server to a member of the Windows Server 2003 family preserves the existing **Allow LM Authentication** setting.

- If MS-CHAP v1 is used as the authentication protocol, a 40-bit encrypted connection cannot be established if the user's password is larger than 14 characters. This behavior affects both dial-up and virtual private network-based remote access and demand dial connections.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

# MS-CHAP version 2

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MS-CHAP version 2**

The Windows Server 2003 family includes support for version 2 of the Microsoft Challenge Handshake Authentication Protocol (MS-CHAP v2), which provides stronger security for remote access connections. MS-CHAP v2 solves some issues of MS-CHAP version 1, as shown in the following table.

| MS-CHAP version 1 issue | MS-CHAP version 2 solution |
|---|---|
| LAN Manager encoding of the response used for backward compatibility with older Microsoft remote access clients is cryptographically weak. | MS-CHAP v2 no longer allows LAN Manager encoded responses. |
| LAN Manager encoding of password changes is cryptographically weak. | MS-CHAP v2 no longer allows LAN Manager encoded password changes. |
| Only one-way authentication is possible. The remote access client cannot verify that it is dialing in to its organization's remote access server or a masquerading remote access server. | MS-CHAP v2 provides two-way authentication, also known as mutual authentication. The remote access client receives verification that the remote access server that it is dialing in to has access to the user's password. |
| With 40-bit encryption, the cryptographic key is based on the user's password. Each time the user connects with the same password, the same cryptographic key is generated. | With MS-CHAP v2, the cryptographic key is always based on the user's password and an arbitrary challenge string. Each time the user connects with the same password, a different cryptographic key is used. |
| A single cryptographic key is used for data sent in both directions on the connection. | With MS-CHAP v2, separate cryptographic keys are generated for transmitted and received data. |

MS-CHAP v2 is a one-way encrypted password, mutual authentication process that works as follows:

1. The authenticator (the remote access server or the IAS server) sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string.

2. The remote access client sends a response that contains:

   - The user name.

   - An arbitrary peer challenge string.

   - A one-way encryption of the received challenge string, the peer challenge string, the session identifier, and the user's password.

3. The authenticator checks the response from the client and sends back a response containing:

   - An indication of the success or failure of the connection attempt.

   - An authenticated response based on the sent challenge string, the peer challenge string, the encrypted response of the client, and the user's password.

4. The remote access client verifies the authentication response and, if correct, uses the connection. If the authentication response is not correct, the remote access client terminates the connection.

**Enabling MS-CHAP v2**

To enable MS-CHAP v2-based authentication, you must do the following:

1. Enable MS-CHAP v2 as an authentication protocol on the remote access server. For more information, see Enable authentication protocols. MS-CHAP v2 is enabled by default.

2. Enable MS-CHAP v2 on the appropriate remote access policy. For more information, see Introduction to remote access policies and Configure authentication. MS-CHAP v2 is enabled by default.

3. Enable MS-CHAP v2 on the remote access client. For more information, see Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2).

## Notes

- Windows 95 with the Windows Dial-Up Networking 1.3 Performance & Security Upgrade for Windows 95 supports MS-CHAP v2 for virtual private network (VPN) connections but not for dial-up connections.

- MS-CHAP (version 1 and version 2) is the only authentication protocol provided with the Windows Server 2003 family that supports password change during the authentication process.

- Make sure your network access server (NAS) supports MS-CHAP v2 before you enable it on a remote access policy on an IAS server. For more information, see your NAS documentation.

# CHAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**CHAP**

The Challenge Handshake Authentication Protocol (CHAP) is a challenge-response authentication protocol that uses the industry-standard Message Digest 5 (MD5) hashing scheme to encrypt the response. CHAP is used by various vendors of network access servers and clients. A server running Routing and Remote Access supports CHAP so that remote access clients that require CHAP are authenticated. Because CHAP requires the use of a reversibly encrypted password, you should consider using another authentication protocol such as MS-CHAP version 2.

To enable CHAP-based authentication, you must do the following:

1. Enable CHAP as an authentication protocol on the remote access server. For more information, see Enable authentication protocols. CHAP is disabled by default.

2. Enable CHAP on the appropriate remote access policy. For more information, see Introduction to remote access policies and Configure authentication.

3. Enable storage of a reversibly encrypted form of the user's password.

   You can enable storage of a reversibly encrypted form of the user's password per user account or enable storage for all accounts in a domain. For more information, see Enable reversibly encrypted passwords in a domain.

4. Force a reset of the user's password so that the new password is in a reversibly encrypted form.

   When you enable passwords to be stored in a reversibly encrypted form, the current passwords are not in a reversibly encrypted form and are not automatically changed. You must either reset user passwords or set user passwords to be changed the next time each user logs on. For more information, see Reset a user password. Once the password is changed, it is stored in a reversibly encrypted form.

   If you set user passwords to be changed the next time a user logs on, the user must log on by using a LAN connection and change the password before they attempt to log on with a remote access connection by using CHAP. You cannot change passwords during the authentication process by using CHAP--the logon attempt fails. One workaround for the remote access user is to temporarily log on by using MS-CHAP to change the password.

5. Enable CHAP on the remote access client. For more information, see Challenge Handshake Authentication Protocol (CHAP).

**Notes**

- If your password expires, CHAP cannot change passwords during the authentication process.

- Make sure your network access server (NAS) supports CHAP before you enable it on a remote access policy on an IAS server. For more information, see your NAS documentation.

- You cannot use Microsoft Point-to-Point Encryption (MPPE) with CHAP.

# PAP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PAP**
Password Authentication Protocol (PAP) uses plaintext passwords and is the least secure authentication protocol. It is typically negotiated if the remote access client and remote access server cannot negotiate a more secure form of validation.

To enable PAP-based authentication, you must do the following:

1. Enable PAP as an authentication protocol on the remote access server. For more information, see Enable authentication protocols. PAP is disabled by default.

2. Enable PAP on the appropriate remote access policy. For more information, see Configure authentication. PAP is disabled by default.

3. Enable PAP on the remote access client. For more information, see Password Authentication Protocol (PAP).

### Important

- When you enable PAP as an authentication protocol, user passwords are sent in plaintext form. Anyone capturing the packets of the authentication process can easily read the password and use it to gain unauthorized access to your intranet. The use of PAP is highly discouraged, especially for virtual private network connections.

### Notes

- By disabling the support for PAP on the remote access server, plaintext passwords are never sent by the dial-up client. Disabling support for PAP increases authentication security, but remote access clients who only support PAP cannot connect.

- If your password expires, PAP cannot change passwords during the authentication process.

- Make sure your network access server (NAS) supports PAP before you enable it on a remote access policy on an IAS server. For more information, see your NAS documentation.

- You cannot use Microsoft Point-to-Point Encryption (MPPE) with PAP.

# Unauthenticated access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Unauthenticated access**
The Windows Server 2003 family supports unauthenticated access, which means that user credentials (a user name and password) are not required. There are some situations where unauthenticated access is useful. This section covers:

- DNIS authorization
- ANI/CLI authentication
- Guest authentication

**Important**

- When you enable unauthenticated access, remote access users are connected without sending user credentials. An unauthenticated remote access client does not negotiate the use of a common authentication protocol during the connection establishment process and does not send a user name or password.

  Unauthenticated access with remote access clients could occur when the authentication protocols configured by the remote access client do not match those configured on the remote access server. In this case, the use of a common authentication protocol is not negotiated and the remote access client does not send a user name and password.

# DNIS authorization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DNIS authorization**

Dialed Number Identification Service (DNIS) enables authorization of a connection attempt based on the number called. DNIS identifies the number that was called to the receiver of the call and is provided by most standard telephone companies. For example, you could allow all connections that connect through a specific toll-free number. To identify DNIS-based connections and apply the appropriate connection settings, you must do the following:

1. Enable unauthenticated access on the remote access server. For more information, see Enable authentication protocols.

2. Create a remote access policy on the authenticating server (remote access server or IAS server) for DNIS-based authorization with the Called-Station-Id condition set to the phone number. For more information, see Introduction to remote access policies and Add a remote access policy.

3. Enable unauthenticated access on the remote access policy for DNIS-based authorization. For more information, see Configure authentication.

### Notes

- If your phone service or hardware does not support DNIS, the passing of the number that was called, you can manually set the phone number of the port. For more information, see Set the phone number on a port.

- IAS does not support proxied DNIS access requests.

# ANI/CLI authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ANI/CLI authentication**
Automatic Number Identification/Calling Line Identification (ANI/CLI) authentication is the authentication of a connection attempt based on the phone number of the caller. ANI/CLI service returns the number of the caller to the receiver of the call and is provided by most standard telephone companies.

ANI/CLI authentication is different from caller ID authorization. In caller ID authorization, the caller sends a valid user name and password. The caller ID that is configured for the dial-in property on the user account must match the connection attempt; otherwise, the connection attempt is rejected. In ANI/CLI authentication, a user name and password are not sent.

To identify ANI/CLI-based connections and apply the appropriate connection settings, you must do the following:

1. Enable unauthenticated access on the remote access server. For more information, see Enable authentication protocols.

2. Enable unauthenticated access on the appropriate remote access policy for ANI/CLI-based authentication. For more information, see Introduction to remote access policies and Configure authentication.

3. Create a user account for each number that will be calling for which you want to provide ANI/CLI authentication. The name of the user account must match the number that the user is dialing from. For example, if a user is dialing in from 555-0100, create a "5550100" user account.

4. Set the following registry value to 31 on the authenticating server (either the remote access server or the IAS server):

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy\User Identity Attribute**

   This registry setting tells the authenticating server to use the calling number (RADIUS attribute 31, Calling-Station-ID) as the identity of the calling user. The user identity is set to the calling number only when there is no user name being supplied in the connection attempt.

   To always use the calling number as the user identity, set the following registry value to 1 on the authenticating server:

   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy\Override User-Name**

   However, if you set **Override User-Name** to 1 and the **User Identity Attribute** to 31, the authenticating server can only perform ANI/CLI-based authentication. Normal authentication by using authentication protocols such as MS-CHAP, CHAP, and EAP is disabled.

### Note

- Changes to the registry settings will not take effect until the Routing and Remote Access service or the Internet Authentication Service are restarted.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

# Guest authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Guest authentication**
By default, the Windows Server 2003 family includes a user account called Guest. With guest authentication, during the authentication process, the caller does not send a user name or password. If unauthenticated access is enabled, by default the Guest account is used as the identity of the caller.

To enable Guest account access, you must do the following:

1. Enable unauthenticated access on the remote access server. For more information, see Enable authentication protocols.

2. Enable unauthenticated access on the appropriate remote access policy. For more information, see Configure authentication.

3. Enable the Guest account. For more information, see To enable a disabled user account.

4. Set the remote access permission on the Guest account to either **Allow access** or **Control access through Remote Access Policy** depending on your remote access policy administrative model. For more information, see Remote Access Policies Examples and Configure remote access permission for a user.

If you want to enable a guest account that is not named **Guest**, create a user account and set the remote access permission to either **Allow access** or **Control access through Remote Access Policy**. Then, set the following registry value on the authenticating server (either the remote access server or the IAS server) to the name of the account:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy\Default User Identity**

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Notes**

- Changes to the registry setting will not take effect until the Routing and Remote Access service or the Internet Authentication Service are restarted.

- For more information about how to modify registry keys and set values, see "Remote Access Server" at the Microsoft Windows Resource Kits Web site.

# Remote Access Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access policies**
This section covers:

- Introduction to remote access policies
- Dial-in properties of a user account
- Elements of a remote access policy
- Accepting a connection attempt
- Remote Access Policies Examples

# Introduction to remote access policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introduction to remote access policies**

In Windows NT versions 3.5, 3.51, and 4.0, authorization was based on a **Grant dial-in permission to user** option in either User Manager or the Remote Access Admin utility. An individual user configured callback options. For the Routing and Remote Access service and Internet Authentication Service (IAS) in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000, network access authorization is granted on the basis of user account dial-in properties and remote access policies.

Remote access policies are an ordered set of rules that define how connections are either authorized or rejected. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting.

If a connection is authorized, the remote access policy profile specifies a set of connection restrictions. The dial-in properties of the user account also provide a set of restrictions. Where applicable, user account connection restrictions override the remote access policy profile connection restrictions.

For information about connection restrictions that can be configured on a user account, see Dial-in properties of a user account. For information about connection restrictions that can be configured through a remote access policy profile, see Elements of a remote access policy.

Remote access policies validate a number of connection settings before authorizing the connection, including the following:

- Remote access permission
- Group membership
- Type of connection
- Time of day
- Authentication methods
- Advanced conditions:
    - Access server identity
    - Access client phone number or MAC address
    - Whether user account dial-in properties are ignored
    - Whether unauthenticated access is allowed

After the connection is authorized, remote access policies can also be used to specify connection restrictions, including the following:

- Idle timeout time
- Maximum session time
- Encryption strength
- IP packet filters
- Advanced restrictions:

- IP address for PPP connections
- Static routes

Additionally, you can vary connection restrictions based on the following settings:

- Group membership
- Type of connection
- Time of day
- Authentication methods
- Identity of the access server
- Access client phone number or MAC address
- Whether unauthenticated access is allowed

For example, you can have policies that specify different maximum session times for different types of connections or groups. Additionally, you can specify restricted access for business partners or unauthenticated connections.

**Authorizing access**
There are two ways to use remote access policies to grant authorization:

1. By user
2. By group

**Authorization by user**
If you are managing authorization by user, set the remote access permission on the user or computer account to either **Grant access** or **Deny access** and, optionally, create different remote access policies based on different types of connections. For example, you might want to have one remote access policy that is used for dial-up connections and a different remote access policy that is used for wireless connections. Managing authorization by user is recommended only when you have a small number of user or computer accounts to manage.

If you are managing authorization by user, the basic process for authorizing a connection attempt occurs as follows:

- If the connection attempt matches all policy conditions, check the remote access permission setting of the account.
    - If the remote access permission is set to **Grant access**, apply the connection settings of the policy profile and user account.
    - If the remote access permission is set to **Deny access**, reject the connection attempt.
- If the connection attempt does not match all policy conditions, process the next remote access policy.

If the connection attempt does not match all conditions of any remote access policy, reject the connection attempt.

**Authorization by group**
If you are managing authorization by group, set the remote access permission on the user account to **Control access through Remote Access Policy** and create remote access policies that are based on different types of connections and group membership. For example, you might want to have one remote access policy for dial-up connections for employees (members of the Employees group) and a different remote access policy for dial-up connections for contractors (members of the Contractors group).

If you are managing authorization by group, the basic process for authorizing a connection attempt occurs as follows:

- If the connection attempt matches all policy conditions, check the remote access permission of the remote access policy.
    - If the remote access permission is set to **Grant remote access permission**, apply the connection settings of the policy profile and user account.
    - If the remote access permission is set to **Deny remote access permission**, reject the connection attempt.
- If the connection attempt does not match all policy conditions, process the next remote access policy.

If the connection attempt does not match all conditions of any remote access policy, reject the connection attempt.

**Notes**

- Setting the remote access permission on user accounts to **Control access through Remote Access Policy** and not using groups to manage network access is not recommended.
- For more information about how connection attempts are processed, see Remote Access Policies Examples

**Remote access policies and the Routing and Remote Access service**

For servers running the Routing and Remote Access service that are configured for the Windows authentication provider, remote access policies are administered from Routing and Remote Access and apply only to the connections of the Routing and Remote Access service server. For servers running the Routing and Remote Access service that are configured for the RADIUS authentication provider, you can no longer configure remote access policies using Routing and Remote Access. If the RADIUS server is an IAS server, remote access policies are administered from Internet Authentication Service.

To centrally manage a single set of remote access policies for multiple remote access or VPN servers using an IAS server, you must do the following:

1. Install the Internet Authentication Service (IAS) as a Remote Authentication Dial-In User Service (RADIUS) server. For more information, see Checklist: Configuring IAS for dial-up and VPN access.
2. Configure IAS with RADIUS clients that correspond to each of the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition remote access or VPN servers. For more information, see Add RADIUS clients.
3. On the IAS server, create the central set of policies that will be used by all servers running the Routing and Remote Access service. For more information, see Add a remote access policy.
4. Configure each of the servers running the Routing and Remote Access service as RADIUS clients to the IAS server. For more information, see Use RADIUS authentication.

For more information about deploying IAS for centralized remote access policy management, see Using RADIUS for multiple remote access servers.

Centralized management of remote access policies is also used when you have remote access servers that are running Windows NT 4.0 and the Routing and Remote Access Service (RRAS). You can configure this server as a RADIUS client to an IAS server. However, a remote access server running Windows NT 4.0 without RRAS cannot be configured to use centralized remote access policies.

# Dial-in properties of a user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-in properties of a user account**

In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; the user account for a stand-alone server or server running Active Directory contains a set of dial-in properties that are used when allowing or denying a connection attempt made by a user. On a stand-alone server, you can set the dial-in properties on the **Dial-in** tab in the user account in Local Users and Groups. On a server running Active Directory, you can set the dial-in properties on the **Dial-in** tab in the user account in Active Directory Users and Computers. On these servers, you cannot use the Windows NT 4.0 User Manager for Domains administrative tool.

The dial-in properties for a user account are:

- **Remote Access Permission (Dial-in or VPN)**

  You can use this property to set remote access permission to be explicitly allowed, denied, or determined through remote access policies. In all cases, remote access policies are used to authorize the connection attempt. If access is explicitly allowed, remote access policy conditions, user account properties, or profile properties can still deny the connection attempt. The **Control access through Remote Access Policy** option is only available on user accounts in a Windows 2000 native domain, a Windows server 2003 domain, or for local accounts on stand-alone servers running Windows 2000, Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition.

  By default, the Administrator and Guest accounts on a stand-alone server or in a Windows 2000 native domain are set to **Control access through Remote Access Policy**. In a Windows 2000 mixed domain, they are set to **Deny access**. New accounts that are created on a stand-alone server or in a Windows 2000 native domain are set to **Control access through Remote Access Policy**. New accounts that are created in a Windows 2000 mixed domain are set to **Deny access**.

  For more information, see Domain and forest functionality.

- **Verify Caller ID**

  If this property is enabled, the server verifies the caller's phone number. If the caller's phone number does not match the configured phone number, the connection attempt is denied.

  Caller ID must be supported by the caller, the phone system between the caller and the remote access server, and the remote access server. On a computer running the Routing and Remote Access service, caller ID support consists of call answering equipment that provides caller ID information and the appropriate Windows Server 2003, Standard Edition ; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition driver to pass the information to the Routing and Remote Access service.

  If you configure a caller ID phone number for a user, and you do not have support for the passing of caller ID information from the caller to the Routing and Remote Access service, the connection attempt is denied.

- **Callback Options**

  If this property is enabled, the server calls the caller back during the connection process. The phone number that is used by the server is set by either the caller or the network administrator.

- **Assign a Static IP Address**

  You can use this property to assign a specific IP address to a user when a connection is made.

- **Apply Static Routes**

You can use this property to define a series of static IP routes that are added to the routing table of the server running the Routing and Remote Access service when a connection is made. This setting is designed for user accounts that a router running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition uses for demand-dial routing. For more information, see One-way initiated demand-dial connections.

For information about configuring dial-in properties for a user account, see Configure dial-in user properties.

**Support for ignoring the dial-in properties of user accounts**
In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; you can configure a RADIUS attribute to ignore the dial-in properties of user and computer accounts in the profile properties of a remote access policy. To support multiple types of connections for which IAS provides authentication and authorization, it might be necessary to disable the processing of user account dial-in properties. This can be done to support scenarios in which specific dial-in properties are not required.

For example, the caller-ID, callback, static IP address, and static routes properties are designed for a client that is dialing into a network access server (NAS). These settings are not designed for wireless access points. A wireless access point that receives these settings in the RADIUS message from the IAS server might be unable to process them, which could cause the wireless client to become disconnected. When IAS provides authentication and authorization for users who are both dialing in and accessing the organization network through wireless technology, the dial-in properties must be configured to support either dial-in connections (by setting dial-in properties) or wireless connections (by not setting dial-in properties).

You can use IAS to enable the processing of dial-in properties for user and computer accounts in some scenarios (such as dial-in), and to disable the processing of dial-in properties for user and computer accounts in other scenarios (such as wireless and authenticating switch). This is done by configuring the **Ignore-User-Dialin-Properties** attribute on the **Advanced** tab of the profile settings for a remote access policy. For more information, see Add RADIUS attributes to a remote access policy.

The **Ignore-User-Dialin-Properties** attribute is set to the following:

- To enable the processing of dial-in properties for a user account, delete the **Ignore-User-Dialin-Properties** attribute or set it to **False**. For example, for a remote access policy that is designed for dial-in connections, no additional configuration is required.

- To disable the processing of dial-in properties for a user account, set the **Ignore-User-Dialin-Properties** attribute to the value of **True**. For example, this is set for the remote access policy that is used for wireless or authenticating switch connections. When the dial-in properties of the user account are ignored, remote access permission is determined by the setting in the remote access policy.

### Notes

- For a user account in a Windows NT 4.0 domain or a Windows 2000 mixed domain:

    - Only the dial-in properties named **Remote Access Permission (Dial-in or VPN)** (**Allow access** and **Deny access** options) and **Callback Options** are available.

    - You can use the Windows NT 4.0 User Manager for Domains administrative tool to grant or deny dial-in access and set callback options.

- If a user account is in a Windows 2000 native domain, the callback number can be up to 128 characters in length. If a user account is on a stand-alone remote access server, in a Windows NT 4.0 domain, or in a Windows 2000 mixed domain, the callback number can contain 24 through 48 characters. This is due to the compressed format for storing callback numbers.

- When a remote access server running Windows NT 4.0 uses a Windows 2000 native domain to obtain the dial-in properties of a user account, the **Control access through Remote Access Policy** remote access permission setting is retrieved as **Deny access**. Callback settings are retrieved correctly.

  For more information, see Domain and forest functionality.

- A remote access server running Windows NT 4.0 cannot use remote access policies. If you upgrade this remote access server with the Routing and Remote Access Service (RRAS) and configure it for RADIUS authentication, you can use the remote access policies of an IAS server. Alternately, you can upgrade this remote access server with a member of the Windows 2000 server family; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition.

- User accounts upgraded to Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition that were configured with dial-in permission enabled are set to **Allow access**. Those that were configured with dial-in permission disabled are set to **Deny access**.

- It is possible to ignore all of the user or computer account properties. For more information, see the section "Ignoring the dial-in properties of user accounts" in New features for IAS.

- In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; computer accounts also have dial-in properties and are similar to user accounts. This new feature allows computers to be authenticated as if they were users, and is used when an IEEE 802.1x Ethernet client uses EAP-TLS and an installed computer certificate to authenticate itself to an Ethernet switch.

- The **Ignore-User-Dialin-Properties** attribute disables the use of all dial-in properties for the user account. Specific dial-in properties cannot be selectively disabled.

# Elements of a remote access policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Elements of a remote access policy**

A remote access policy is a named rule that consists of the following elements:

- Conditions
- Remote access permission
- Profile

**Conditions**

Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt. If there are multiple conditions, then all of the conditions must match the settings of the connection attempt in order for it to match the policy.

The following table shows the condition attributes that you can set for a remote access policy.

| Attribute name | Description |
|---|---|
| Authentication Type | The type of authentication that is being used by the access client. Authentication types include CHAP, EAP, MS-CHAP, and MS-CHAP v2. |
| Called Station ID | The phone number of the network access server (NAS). This attribute is a character string. You can use pattern matching syntax to specify area codes. For more information, see Pattern matching syntax. In order to receive called station ID information during a call, the phone line, the hardware, and the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition driver for the hardware must support passing the called ID. Otherwise, the called station ID is manually set for each port. To manually set phone numbers on ports, see Set the phone number on a port. |
| Calling Station ID | The phone number used by the caller. This attribute is a character string. You can use pattern matching syntax to specify area codes. For more information, see Pattern matching syntax. For more information about obtaining the calling station ID, see Caller ID and callback. |
| Client Friendly Name | The name of the RADIUS client that is requesting authentication. This name is configured in **Friendly name** on the **Settings** tab in the properties of a RADIUS client in Internet Authentication Service. This attribute is a character string. You can use pattern matching syntax to specify client names. For more information, see Pattern matching syntax. This attribute is used by the IAS server. |
| Client IP Address | The IP address of the RADIUS client. This can be the IP address of either the originating RADIUS client or an intermediate RADIUS proxy. This attribute is a character string. You can use pattern matching syntax to specify IP networks. For more information, see Pattern matching syntax. This attribute is used by the IAS server. |
| Client Vendor | The vendor of the network access server (NAS) that is requesting authentication. The Routing and Remote Access service is the **Microsoft** NAS manufacturer. You can use this attribute to configure separate policies for different NAS manufacturers who are RADIUS clients to an IAS server. This attribute is used by the IAS server. Ensure that you configure the NAS as a RADIUS client on the IAS server. For more information, see Add RADIUS clients. |
| Day and Time Restrictions | The day of the week and the time of day of the connection attempt. The day and time is relative to the day and time of the server providing the authorization. |
| Framed Protocol | The type of framing for incoming packets. Examples are PPP, SLIP, Frame Relay, and X.25. This attribute is used by the IAS server. |

| | |
|---|---|
| NAS Identifier | The name of the network access server (NAS). This attribute is a character string. You can use pattern matching syntax to specify multiple NASs. For more information, see Pattern matching syntax. This attribute is used by the IAS server. |
| NAS IP Address | The IP address of the NAS (the RADIUS client) that sent the message. This attribute is a character string. You can use pattern matching syntax to specify IP networks. For more information, see Pattern matching syntax. This attribute is used by the IAS server. |
| NAS Port Type | The type of media that is used by the access client. Examples include analog phone lines (known as async), ISDN, tunnels or virtual private networks (known as virtual), IEEE 802.11 wireless, and Ethernet switches. |
| Service Type | The type of service that is being requested. Examples include framed (such as PPP connections) and login (such as Telnet connections). For more information about RADIUS service types, see RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)." This attribute is used by the IAS server. |
| Tunnel Type | The type of tunnel that is being created by the requesting client. Tunnel types include the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol (L2TP), which are used by Windows XP; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000, remote access clients and demand-dial routers. You can use this condition to specify profile settings, such as authentication methods or encryption strengths for a specific type of tunneling technology. |
| Windows Groups | The names of the groups to which the user or computer account that is attempting the connection belongs. There is no condition attribute for a specific user or computer name. It is not necessary to have a separate remote access policy for each group. Instead, you can use multiple groups or nested groups to consolidate and delegate the administration of group membership. For a remote access or IAS server in a Windows 2000 native domain, you can use universal groups. For more information, see Domain and forest functionality. |

**Notes**

- If conditions that use an IAS server attribute are evaluated against a Routing and Remote Access service server that is configured for Windows authentication, they do not match and the policy is not applied.
- Not all access servers send all of the IAS server-specific attributes.
- You cannot use the built-in local or domain groups for the Windows Groups attribute.

**Remote access permission**

If all conditions of a remote access policy are met, remote access permission is either granted or denied. You can use either the **Grant remote access permission** option or the **Deny remote access permission** option to set remote access permission for a policy.

Remote access permission is also granted or denied for each user account. The user remote access permission overrides the policy remote access permission. When remote access permission on a user account is set to the **Control access through Remote Access Policy** option, the policy remote access permission determines whether the user is granted access.

Granting access through either the user account permission setting or the policy permission setting is only the first step in accepting a connection. The connection attempt is subject to the settings of both the user account dial-in properties and policy profile properties. If the connection attempt does not match the settings of the user account or policy profile properties, the connection attempt is rejected.

**Profile**

A remote access policy profile is a set of properties that are applied to a connection when it is authorized--either through the user account or policy permission settings. A profile consists of the following groups of properties:

- Dial-in constraints
- IP
- Multilink

- Authentication
- Encryption
- Advanced

**Dial-in constraints**

You can set the following dial-in constraints:

- **Minutes server can remain idle before it is disconnected**

  The time after which a connection is disconnected when there is no activity. By default, this property is not set. Additionally, the Routing and Remote Access service does not disconnect an idle connection.

- **Minutes client can be connected**

  The maximum amount of time that a connection is connected. The connection is disconnected by the access server after the maximum session length. By default, this property is not set and there is no maximum session limit.

- **Allow access only on these days and at these times**

  The days of the week and hours of each day that a connection is allowed. If the day and time of the connection attempt do not match the configured day and time limits, the connection attempt is rejected. By default, this property is not set and there are no day or time limits. The Routing and Remote Access service does not disconnect active connections that were previously connected at a time when connection attempts were allowed.

- **Allow access to this number only**

  The specific phone number that a caller must call in order for a connection to be allowed. If the dial-in number of the connection attempt does not match the configured dial-in number, the connection attempt is rejected. By default, this property is not set and all dial-in numbers are allowed.

- **Allow access only through these media**

  The specific types of media, such as modem (also known as async), ISDN, virtual private network (known as virtual), or 802.11 wireless that a caller must use in order for a connection to be allowed. If the dial-in medium of the connection attempt does not match the configured dial-in media, the connection attempt is rejected. By default, this property is not set and all media types are allowed.

For information about setting dial-in constraints on a profile, see Configure dial-in constraints.

**IP**

You can set IP properties that specify IP address assignment behavior. You have the following options:

- The access server must supply an IP address.
- The access client can request an IP address.
- IP address assignment is determined by the access server (this is the default setting).
- A static IP address is assigned. A static IP address assigned to the user account overrides this setting. The IP address assigned is typically used to accommodate vendor-specific attributes for IP addresses.

You can also use the **IP** tab to define IP packet filters that apply to remote access connection traffic. This is intended for use with the Routing and Remote Access service. You can configure IP traffic that is allowed to remote access clients (output filters) or from remote access clients (input filters) on an exception basis. Either all traffic is allowed, except traffic specified by filters; or all traffic is blocked, except traffic specified by filters . Remote access policy profile filtering applies to all remote access

connections that match the remote access policy.

Filters configured on the **IP** tab are applied to client connections on a Windows Server 2003 remote access server, but are not applied to dial-on-demand connections.

For information about setting IP options on a profile, see Configure IP options.

**Multilink**
You can set Multilink properties that both enable Multilink and determine the maximum number of ports that a Multilink connection can use. Additionally, you can set Bandwidth Allocation Protocol (BAP) policies that both determine BAP usage and specify when extra BAP lines are dropped. The Multilink and BAP properties are specific to the Routing and Remote Access service. By default, Multilink and BAP are disabled.

The Routing and Remote Access service must have Multilink and BAP enabled in order for the Multilink properties of the profile to be enforced. For more information about enabling Multilink and BAP for the Routing and Remote Access service, see Enable Multilink and Enable BAP and BACP.

For information about setting Multilink options on a profile, see Configure multilink options.

**Authentication**
You can set authentication properties to both enable the authentication types that are allowed for a connection and specify the EAP type that must be used. Additionally, you can configure the EAP type. By default, Microsoft Encrypted Authentication (MS-CHAP) and Microsoft Encrypted Authentication version 2 (MS-CHAP v2) are enabled. With Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; you can specify whether users can change their expired passwords using MS-CHAP and MS-CHAP v2 (enabled by default). For more information, see Authentication methods.

The Routing and Remote Access service must have the corresponding authentication types enabled in order for the authentication properties of the profile to be enforced. For more information about enabling authentication methods for the Routing and Remote Access service, see Enable authentication protocols.

For information about setting authentication options on a profile, see Configure authentication.

**Encryption**
You can set encryption properties for the following encryption strengths:

- **No Encryption**

  This option allows a nonencrypted connection. To require encryption, clear this option.

- **Basic**

  For dial-up and PPTP-based VPN connections, Microsoft Point-to-Point Encryption (MPPE) is used with a 40-bit key. For L2TP/IPSec VPN connections, 56-bit Data Encryption Standard (DES) encryption is used.

- **Strong**

  For dial-up and PPTP VPN connections, MPPE is used with a 56-bit key. For L2TP/IPSec VPN connections, 56-bit DES encryption is used.

- **Strongest**

  For dial-up and PPTP VPN connections, MPPE is used with a 128-bit key. For L2TP/IPSec VPN connections, triple DES (3DES) encryption is used.

These settings are intended to be used with Routing and Remote Access. When other access servers are used, select **No Encryption** and clear all encryption strengths.

For information about setting encryption options on a profile, see Configure encryption.

**Advanced**
You can set advanced properties to specify the series of RADIUS attributes that are sent back by the IAS server to be evaluated by the RADIUS client. RADIUS attributes are specific to performing RADIUS authentication, and are ignored by a server that is

running the Routing and Remote Access service and is configured for Windows authentication. By default, **Framed-Protocol** is set to **PPP** and **Service-Type** is set to **Framed**.

The only advanced attributes that are used by the Routing and Remote Access service are **Account-Interim-Interval**, **Framed-Protocol**, **Framed-MTU**, **Reply-Message**, and **Service-Type**.

For information about setting advanced options on a profile, see Add RADIUS attributes to a remote access policy.

**Notes**

- Some elements of a remote access policy correspond to RADIUS attributes that are used during RADIUS-based authentication. For remote access policies on an Internet Authentication Service (IAS) server, verify that the access servers used are sending RADIUS attributes that correspond to the configured remote access policy conditions and profile settings. If a NAS does not send a RADIUS attribute that corresponds to a remote access policy condition or profile setting, then all RADIUS authentications from that NAS are denied.

- You can only use the Generate-Session-Timeout attribute if your user account database is a Security Accounts Manager (SAM) database or is the user account database for an Active Directory domain. When the value of **Generate-Session-Timeout** is set to **True**, the **ForceLogoff** value for a SAM database should be set to **0**. In the Local Security Settings console, ForceLogoff is changed to zero when **Network security: Force logoff when logon hours expire** is enabled. For information about how to change the ForceLogoff value for the SAM database at the command prompt, see Net accounts.

- You can configure wireless connection policy so that wireless clients periodically reauthenticate. This ensures that the client Wired Equivalent Privacy (WEP) encryption keys are changed often enough to provide adequate security for the wireless connection. To configure reauthentication, set the session timeout in your remote access policy or connection request policy for wireless connections (using the Session-Timeout attribute) to the required interval (for example, 10 minutes). Additionally, configure the Termination-Action attribute with the **Attribute value** set to **RADIUS-Request**. If the Termination-Action attribute is not set to **RADIUS-Request**, wireless access points might end the connection during reauthentication. For more information, see your hardware documentation.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Accepting a connection attempt

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accepting a connection attempt**

When a user attempts a connection, the connection attempt is accepted or rejected, based on the following logic:

1. The first policy in the ordered list of remote access policies is checked. If there are no policies, reject the connection attempt.

2. If all conditions of the policy do not match the connection attempt, go to the next policy. If there are no more policies, reject the connection attempt.

3. If all conditions of the policy match the connection attempt, check the value of the **Ignore-User-Dialin-Properties** attribute.

4. If the **Ignore-User-Dialin-Properties** attribute is set to **False**, check the remote access permission setting for the user attempting the connection.

   - If **Deny access** is selected, reject the connection attempt.

   - If **Allow access** is selected, apply the user account and profile properties. If the connection attempt does not match the settings of the user account and profile properties, reject the connection attempt. If the connection attempt matches the settings of the user account and profile properties, accept the connection attempt.

   - If the remote access permission is not set to **Allow access** or **Deny access**, the remote access permission must be set to **Control access through Remote Access Policy**. Check the remote access permission setting of the policy.

     If **Deny remote access permission** is selected, reject the connection attempt.

     If **Grant remote access permission** is selected, apply the user account and profile properties. If the connection attempt does not match the settings of the user account and profile properties, reject the connection attempt. If the connection attempt matches the settings of the user account properties and profile, accept the connection attempt.

5. If the **Ignore-User-Dialin-Properties** attribute is set to **True**, check the remote access permission setting of the policy.

   - If **Deny remote access permission** is selected, reject the connection attempt.

   - If **Grant remote access permission** is selected, apply the profile properties. If the connection attempt does not match the settings of the profile properties, reject the connection attempt. If the connection attempt matches the settings of the profile properties, accept the connection attempt. The following illustration shows the logic of remote access policies.




**Notes**

- The profile and user account settings for the first matching policy are applied to the connection. If a connection does not match the profile or user account settings of the remote access policy, the additional remote access policies are not tried.

- A connection attempt might not match any of the remote access policies. If this is the case, the connection attempt is rejected regardless of the remote access permission setting on the user account.

- The remote access polices are tried in order. The more specific remote access policies are typically placed in order ahead of the more general remote access policies.

- The **Ignore-User-Dialin-Properties** attribute is a new feature for Windows Server 2003, Standard Edition ; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition that allows you to ignore all of the dial-in properties of a user account. For more information, see New features for IAS.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

For examples of how different connection attempts are processed, see Remote Access Policies Examples.

# Remote Access Policies Examples

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access policies examples**

Remote access policies provide a tremendous amount of flexibility in accepting connections based on policy evaluation order, policy conditions, user and policy dial-in permissions, and both policy profile and user account properties.

This section describes some examples in which remote access policies are used:

- Grant remote access per user
- Allow dial-up connection using group membership
- Restrict connections to a maximum session time
- Wireless access example
- Wireless access with secure password authentication
- Ethernet switch access
- Force VPN clients to use strongest encryption
- Apply packet filters for business partner extranet
- Send NAS-specific RADIUS attributes to a RADIUS client

# Grant remote access per user

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Grant remote access per user**

In this example, the network administrator wants to maintain the same administrative model as used in Windows NT 4.0. Dial-in access is granted on a per-user basis by modifying the remote access permission setting of the user account dial-in properties. The remote access permission is set to either **Allow access** or **Deny access**. This is the default behavior for remote access policies.

You can create a policy named **Allow access if dial-in permission is enabled**. The settings on this policy are configured as follows:

- A single condition that consists of the **Day-And-Time-Restrictions** attribute is set for all times on all days.
- The **Deny remote access permission** option is selected.
- The profile is set to the default settings.

When a user attempts a remote access connection, the following logic is used to accept or reject the attempt (assuming that the default settings are set for all the other user account dial-in properties):

1. The policy **Allow access if dial-in permission is enabled** is evaluated.
2. The settings of the connection attempt match the conditions of the policy (the connection attempt occurs on any day or at any time).
3. On the user account, the setting in **Remote Access Permission (Dial-in or VPN)** is evaluated.
    - If **Allow access** is selected, the connection is accepted, subject to the settings of the user account dial-in and profile properties.
    - If **Deny access** is selected, the connection is rejected.

# Allow dial-up connection using group membership

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Allow dial-up connection using group membership**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants to allow dial-up connections for only those user accounts that belong to a specific set of groups. After remote access permission is set for all user accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common dial-up policy with the following settings:

   - Policy name: Dial-up access for allowed groups
   - Access Method: Dial-up access
   - User or Group: Select **Group**, and then specify the set of groups for which dial-up access is allowed.
   - Authentication methods: Select **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** and **Microsoft Encrypted Authentication**.
   - Policy Encryption Level: Select all check boxes.

     For more information, see Add a remote access policy.

2. Delete the default policies.

   For more information, see Delete a remote access policy.

**Notes**

- It is not necessary to have a separate remote access policy for each group. You can specify multiple groups or you can use nested groups to both delegate administration of and consolidate group membership. For a remote access or IAS server in a Windows 2000 native domain or a Windows Server 2003 domain, you can use universal groups. For more information, see Domain and forest functionality.
- You cannot use built-in or local domain groups when specifying group names.

# Restrict connections to a maximum session time

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restrict connections to a maximum session time**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

To conserve limited remote access server connection resources, the network administrator wants to limit the maximum session time for all dial-up users (members of the DialUpUsers group) to 20 minutes.

To implement this restriction through a remote access policy, the administrator completes the following steps:

**Use the New Remote Access Policy Wizard to create a custom policy with the following settings:**

- Policy name: Maximum session time of 20 minutes
- Conditions: NAS-Port-Type matches Async (Modem); Windows-Groups matches DialUpUsers
- Permission: Grant remote access permission
- Profile settings, **Dial-in Constraints** tab: Select the **Minutes client can be connected** check box, and then type **20**.

For more information, see Add a remote access policy.

**Delete the default policies.**
For more information, see Delete a remote access policy.

# Wireless access example

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless access example**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator requires that all IEEE 802.11 wireless clients (members of the WirelessClients group) use certificate-based authentication through the EAP-TLS authentication protocol and 128-bit encryption. Wireless clients that do not have a user certificate installed are placed on a specific virtual local area network (VLAN) from which they can obtain a user certificate. This unauthenticated connection lasts for 10 minutes.

In this case, two policies are needed:

1. A policy that requires wireless connections to use EAP-TLS.

2. A policy that accepts unauthenticated wireless connections and configures the wireless access point for a VLAN with an ID of 3. VLAN 3 contains a certificate server from which the wireless client can obtain a user certificate.

After remote access permission is set for all user accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common wireless policy with the following settings:

   - Policy name: Wireless access

   - Access Method: Wireless access

   - User or Group: Select **Group**, and then specify the WirelessClients group (example).

   - Authentication methods: Select **Smart Card or other Certificate**.

   - Policy Encryption Level: Check the **Strongest encryption (MPPE 128-bit)** check box, and then clear all other check boxes.

2. Use the New Remote Access Policy Wizard to create a custom wireless policy with the following settings:

   - Policy name: First-time wireless access

   - Conditions: NAS-Port-Type matches Wireless-Other or Wireless-IEEE 802.11; Windows-Groups matches WirelessUsers; Authentication-Type matches Unauthenticated

   - Permission: Grant remote access permission

   - Profile settings, **Dial-in Constraints** tab: Select the **Minutes client can be connected** check box, and then type **10**.

   - Profile settings, **Advanced** tab: Add the **Tunnel-Type** attribute with the value of **Virtual LANs (VLAN)**; Add the **Tunnel-Pvt-Group-ID** attribute with the value of **3**.

3. Delete the default policies.

   For more information, see Delete a remote access policy.

**Notes**

- For this example, a public key infrastructure (PKI) that has allocated the required computer certificate to the IAS server is being used. Additionally, the wireless access point is configured to use both an IAS server as its RADIUS server and EAP. For more information about wireless access deployment, see Checklist: Configuring the IAS server and wireless access points for wireless access.

- You can configure wireless connection policy so that wireless clients periodically reauthenticate. This ensures that the client Wired Equivalent Privacy (WEP) encryption keys are changed often enough to provide adequate security for the wireless connection. To configure reauthentication, set the session timeout in your remote access policy or connection request policy for wireless connections (using the Session-Timeout attribute) to the required interval (for example, 10 minutes). Additionally, configure the Termination-Action attribute with the **Attribute value** set to **RADIUS-Request**. If the Termination-Action attribute is not set to **RADIUS-Request**, wireless access points might end the connection during reauthentication. For more information, see your hardware documentation.

# Wireless access with secure password authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless access with secure password authentication**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator requires that all IEEE 802.11 wireless clients (members of the *WirelessClientsGroupName*) use Protected Extensible Authentication Protocol, or PEAP, for secure password authentication on the network. PEAP uses Transport Level Security (TLS) to create an end-to-end (EAP client to EAP authenticator) encrypted connection after verifying the identity of the authenticator (typically, an IAS server). The IAS server authenticates itself to the client with a certificate that includes the Server Authentication purpose in Enhanced Key Usage (EKU) extensions.

Because PEAP is used with EAP-MS-CHAPv2 as the authentication type, user authentication is accomplished with password-based credentials that are typed by the user. PEAP with EAP-MS-CHAPv2, also called PEAP-EAP-MS-CHAPv2, eliminates the need to deploy certificates to wireless clients for the purpose of client computer and user authentication.

For more information, see Network access authentication and certificates.

**Note**

- You can also use PEAP with EAP-TLS (also known as PEAP-EAP-TLS), which provides strong security. PEAP-EAP-TLS uses a public key infrastructure with certificates, which are used for server authentication, and either smart cards or certificates for client computer and user authentication. When you use PEAP-EAP-TLS, client certificate information is encrypted, providing improved security over EAP-TLS without PEAP.

After remote access permission is set for all user accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common wireless policy with the following settings:

   - Policy name: Wireless access
   - Access Method: Wireless access
   - User or Group: Select **Group**, and then specify the WirelessClients group (example).
   - Authentication methods: Select **Protected Extensible Authentication Protocol (PEAP)**.
   - Policy Encryption Level: Check the **Strongest encryption (MPPE 128-bit)** check box, and then clear all other check boxes.

2. Delete the default policies.

   For more information, see Delete a remote access policy.

When you select PEAP as the authentication method, the default authentication type is EAP-MS-CHAPv2. The following defaults are used:

- **Number of authentication retries:** 0
- **Allow client to change password** (not selected)

To change these settings, complete the steps in Configure PEAP and EAP methods.

**Notes**

- Although the deployment of certificates for user and client computer authentication is more difficult than using PEAP with EAP-MS-CHAPv2, certificates are more secure and PEAP with EAP-TLS with smart cards or certificates is the recommended authentication method for wireless clients.

- You can configure wireless connection policy so that wireless clients periodically reauthenticate. This ensures that the client Wired Equivalent Privacy (WEP) encryption keys are changed often enough to provide adequate security for the wireless connection. To configure reauthentication, set the session timeout in your remote access policy or connection request policy for wireless connections (using the Session-Timeout attribute) to the required interval (for example, 10 minutes). Additionally, configure the Termination-Action attribute with the **Attribute value** set to **RADIUS-Request**. If the Termination-Action attribute is not set to **RADIUS-Request**, wireless access points might end the connection during reauthentication. For more information, see your hardware documentation.

# Ethernet switch access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ethernet switch access example**

In this example, the network administrator is managing authorization by using groups. All computer accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants all Ethernet switch clients (members of the SwitchClients group) to use password-based authentication through the Extensible Authentication Protocol (EAP) and the Message Digest 5 (MD5) Challenge EAP type. After remote access permission is set for all computer accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common dial-up policy with the following settings:

   - Policy name: Ethernet access for switch clients
   - Access Method: Ethernet
   - User or Group: Select **Group**, and then specify the SwitchClients group (example).
   - Authentication methods: Select the **MD5-Challenge** EAP type.

     For more information, see Add a remote access policy.

2. Delete the default policies.

   For more information, see Delete a remote access policy.

**Note**

- This example assumes that the Ethernet switch has been configured to use both EAP and an IAS server as its RADIUS server. For more information about Ethernet switch deployment, see Ethernet switch access with IAS

# Force VPN clients to use strongest encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Force VPN clients to use strongest encryption**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants all virtual private network (VPN) clients (members of the group VPNClients) to use strong encryption. For PPTP connections, 128-bit Microsoft Point-to-Point Encryption (MPPE) must be used. For L2TP over IPSec connections, Triple DES (3DES)must be used. After remote access permission is set for all user accounts, the administrator completes the following steps:

1. Use the New Remote Access Policy Wizard to create a common VPN policy with the following settings:

   - Policy name: VPN access with strongest encryption
   - Access Method: VPN access
   - User or Group: Select **Group**, and then specify the VPNClients group (example).
   - Authentication methods: Select **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** and **Microsoft Encrypted Authentication**.
   - Policy Encryption Level: Select the **Strongest encryption (IPSec triple DES or MPPE 128-bit)** check box, and then clear all other check boxes.

     For more information, see Add a remote access policy.

2. Delete the default policies.

   For more information, see Delete a remote access policy.

# Apply packet filters for business partner extranet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Apply packet filters for business partner extranet**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants to restrict access for all business partner virtual private network (VPN) connections to the resources of an extranet--a specific network segment that contains Web and file servers. In this example, the extranet is the subnet of 192.168.47.0/24. All business partner user accounts are members of the Partners group.

After remote access permission is set for all user accounts, the administrator completes the following steps.

**Use the New Remote Access Policy Wizard to create a custom policy with the following settings:**

- Policy name: Business partner connections to extranet
- Conditions: Windows-Groups matches Partners
- Permission: Grant remote access permission
- Profile settings, **Dial-in Constraints** tab: Select **Allow access only through these media** and the **Virtual (VPN)** media type.
- Profile settings, **IP** tab, Input packet filter:
    - Deny all traffic except those listed below
    - Destination network, IP address: 192.168.47.0
    - Destination network, Subnet mask: 255.255.255.0
    - Protocol: Any
- Profile settings, **IP** tab, Output packet filter:
    - Deny all traffic except those listed below
    - Source network, IP address: 192.168.47.0
    - Source network, Subnet mask: 255.255.255.0
    - Protocol: Any
- Profile settings, **Authentication** tab: Select both **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** and **Microsoft Encrypted Authentication**. For both of these authentication types, select **User can change password after it has expired**.
- Clear all other check boxes.
- Profile settings, **Encryption** tab: Select the **Strongest encryption** check box, and then clear all other check boxes

  For more information, see Add a remote access policy.

**Delete the default policies.**

For more information, see Delete a remote access policy.

# Send NAS-specific RADIUS attributes to a RADIUS client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Send NAS-specific RADIUS attributes to a RADIUS client**
In this example, the network administrator is managing authorization by using groups. All user accounts have the **Remote Access Permission (Dial-in or VPN)** option set to **Control access through Remote Access Policy**.

The network administrator wants to allow dial-up connections for only those user accounts in the DialUpUsers group and send a series of network access server (NAS) attributes that are specific to the Ascend NAS. After remote access permission is set for all user accounts, the administrator completes the following steps:

1. When the NAS is added as a RADIUS client, configure its manufacturer.

   For more information about configuring a RADIUS client for an IAS server, see Add RADIUS clients.

2. Use the New Remote Access Policy Wizard to create a custom policy with the following settings:

   - Policy name: Allow dial-up connections and send attributes to Ascend NAS
   - Conditions: Client-Vendor matches Ascend Communications, Inc.; NAS-Port-Type matches Async (Modem) ; Windows-Groups matches DialUpUsers
   - Permission: Grant remote access permission
   - Profile settings, **Advanced** tab: Add the appropriate Ascend RADIUS attributes with their correct values.

   For more information, see Add a remote access policy.

3. Delete the default policies.

   For more information, see Delete a remote access policy.

# Connection Request Processing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection request processing**
This section covers:

- Introduction to connection request processing
- Processing a connection request
- Elements of Connection Request Processing
- Connection Request Processing Examples

# Introduction to connection request processing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introduction to connection request processing**
In Windows 2000, Internet Authentication Service (IAS) can only be used as a RADIUS server. In Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, IAS can be used as either a RADIUS server or a RADIUS proxy.

When IAS is used as a RADIUS server:

- Access-Request messages are authenticated through Active Directory or a Windows NT server 4.0 domain, or through the Security Account Manager (SAM) in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. They are authorized with the user or computer account dial-in properties and remote access policies. Additionally, IAS can return an Access-Accept message without authenticating and authorizing the Access-Request message.

- Accounting-Request messages are logged in a local log file based on remote access logging settings.

**Note**

- IAS can be configured for separate use as an authentication RADIUS server or an accounting RADIUS server.

When IAS is used as a RADIUS proxy:

- Access-Request messages are forwarded to another RADIUS server for authentication and authorization.

- Accounting-Request messages are logged in a local log file (based on remote access logging settings) and forwarded to another RADIUS server for accounting.

**Note**

- IAS can be configured for separate use as an authentication RADIUS proxy or an accounting RADIUS proxy.

To determine whether a specific connection attempt request or an accounting message received from a RADIUS client should be processed locally or forwarded to another RADIUS server, the IAS server uses connection request processing. Connection request processing is a combination of:

- Connection request policies

  Connection request policies determine, for any incoming RADIUS request message, whether the message is processed locally or forwarded to another RADIUS server. For more information, see Connection request policies.

- Remote RADIUS server groups

  Remote RADIUS server groups specify multiple RADIUS servers when an incoming RADIUS request message is forwarded. For more information, see Remote RADIUS server groups.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an

unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Processing a connection request

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Processing a connection request**

When an IAS server receives a valid RADIUS request message from a RADIUS client, it is processed as follows:

1. The first policy in the ordered list of connection request policies is checked. If there are no policies and the RADIUS message is an Access-Request message, an Access-Reject message is immediately sent. If there are no policies and the RADIUS message is an Accounting-Request message, the message is discarded.

2. If all conditions of the policy do not match the RADIUS request message, then go to the next policy. If there are no additional policies and the RADIUS message is an Access-Request message, an Access-Reject message is immediately sent. If there are no additional policies and the RADIUS message is an Accounting-Request message, the message is discarded.

3. If all conditions of the policy match the RADIUS request message, then apply the attribute manipulation rules that are configured in the policy.

4. If the RADIUS message is an Access-Request message, check the authentication settings on the profile of the matching policy.

   - If the authentication setting requires using this server for authentication, use a Windows NT 4.0 domain or the Active Directory directory service, or the local Security Account Manager (SAM) on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; for authentication. Use the matching remote access policy and user or computer account dial-in properties for authorization. For more information, see Accepting a connection attempt.

   - If the authentication setting requires forwarding the Access-Request message to a remote RADIUS server group, forward the Access-Request message to a member of the remote RADIUS server on the basis of priority and weight settings.

   - If the authentication setting requires accepting the connection attempt without performing authentication or authorization, an Access-Accept message is immediately sent.

5. If the RADIUS message is an Accounting-Request message, check the accounting settings on the profile of the matching policy.

   - If the accounting settings specify that accounting information is to be sent to a member of a remote RADIUS server group, forward the accounting message to the appropriate RADIUS server in the group on the basis of priority and weight settings.

   - Based on the settings for remote access logging, record the information in the local accounting log file.

The following illustration shows the logic of connection request policies.




For examples of how different connection requests are processed, see Connection Request Processing Examples.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Elements of Connection Request Processing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Elements of connection request processing**

This section covers:

- Connection request policies
- Remote RADIUS server groups

# Connection request policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection request policies**

Connection request policies are sets of conditions and profile settings that give network administrators flexibility in configuring how incoming authentication and accounting request messages are handled by the IAS server. With connection request policies, you can create a series of policies so that some RADIUS request messages sent from RADIUS clients are processed locally (IAS is being used as a RADIUS server) and other types of messages are forwarded to another RADIUS server (IAS is being used as a RADIUS proxy). This capability allows IAS to be deployed in many new RADIUS scenarios.

With connection request policies, you can use IAS as a RADIUS server or as a RADIUS proxy, based on the time of day and day of the week, by the realm name in the request, by the type of connection being requested, by the IP address of the RADIUS client, and so on.

It is important to remember that with connection request policies, a RADIUS request message is processed only if the settings of the incoming RADIUS request message match at least one of the connection request policies. For example, if the settings of an incoming RADIUS Access-Request message do not match at least one of the connection request policies, an Access-Reject message is sent.

For more information about how incoming RADIUS request messages from RADIUS clients are processed, see Processing a connection request.

To create a connection request policy, see Configure Connection Request Policies.

A connection request policy is a named rule that consists of the following elements:

- Conditions
- Profile

Connection request policies, like remote access policies, operate on the following principle:

- Based on the RADIUS request message matching all of the conditions of a connection request policy, the profile settings of the policy are applied to determine the processing of the message.

**Conditions**

Connection request policy conditions are one or more RADIUS attributes that are compared to the attributes of the incoming RADIUS request message. If there are multiple conditions, then all of the conditions must match the attributes of the incoming RADIUS message in order for the RADIUS request message to match the policy.

The following table shows the condition attributes that you can set for a connection request policy.

| Attribute name | Description |
|---|---|
| Called Station ID | The phone number of the network access server (NAS). This attribute is a character string. You can use pattern matching syntax to specify area codes. For more information, see Pattern matching syntax. |
| Calling Station ID | The phone number used by the caller (the access client). This attribute is a character string. You can use pattern matching syntax to specify area codes. For more information, see Pattern matching syntax. |
| Client Friendly Name | The name of the RADIUS client computer that is requesting authentication. This attribute is a character string. You can use pattern matching syntax to specify client names. For more information, see Pattern matching syntax. |
| Client IP Address | The IP address of the network access server (the RADIUS client). This attribute is a character string. You can use pattern matching syntax to specify IP networks. For more information, see Pattern matching syntax. |

| | |
|---|---|
| Client Vendor | The vendor of the network access server (NAS) that is requesting authentication. A computer running the Routing and Remote Access service is the Microsoft NAS manufacturer. You can use this attribute to configure separate policies for different NAS manufacturers. This attribute is a character string. You can use pattern matching syntax. For more information, see Pattern matching syntax. |
| Day and Time Restrictions | The day of the week and the time of day of the connection attempt. The day and time is relative to the day and time of the IAS server. |
| Framed Protocol | The type of framing for incoming packets. Examples are PPP, SLIP, Frame Relay, and X.25. |
| NAS Identifier | The name of the network access server (NAS). This attribute is a character string. You can use pattern matching syntax to specify NAS identifiers. For more information, see Pattern matching syntax. |
| NAS IP Address | The IP address of the network access server (the RADIUS client). This attribute is a character string. You can use pattern matching syntax to specify IP networks. For more information, see Pattern matching syntax. |
| NAS Port Type | The type of media used by the access client. Examples are analog phone lines (known as async), ISDN, tunnels or virtual private networks (known as virtual), IEEE 802.11 wireless, and Ethernet switches. |
| Remote RADIUS to Windows User Mapping | Specifies that Windows authorization occurs for users authenticated by a remote RADIUS server. For example, visitors to your network from partner organizations can be authenticated by their organization's RADIUS server, and can use a Windows user account at your organization to access a guest local area network (LAN) on your network. For more information, see Mapping network authentication and authorization |
| Service Type | The type of service being requested. Examples include framed (for example, PPP connections) and login (for example, Telnet connections). For more information about RADIUS service types, see RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)." |
| Tunnel Type | The type of tunnel that is being created by the requesting client. Tunnel types include the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol (L2TP), used by Windows XP; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000 remote access clients and demand-dial routers. |
| User Name | The user name that is used by the access client in the RADIUS message. This attribute is a character string that typically contains a realm name and a user account name. For more information about realm names, see Realm names. You can use pattern matching syntax to specify user names. For more information, see Pattern matching syntax. |

**Profile**

A connection request policy profile is a set of properties that are applied to an incoming RADIUS message. A connection request policy profile consists of the following groups of properties:

- Authentication
- Accounting
- Attribute manipulation
- Advanced

**Authentication**

You can set the following authentication options that are used for RADIUS Access-Request messages:

- Authenticate requests on this server.

  Use a Windows NT 4.0 domain or the Active Directory directory service, or the local Security Account Manager (SAM) on Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; for both authentication and the matching remote access policy and user account dial-in properties for

authorization. In this case, the IAS server is being used as a RADIUS server.

- Forward requests to another RADIUS server in a remote RADIUS server group.

  Forward the Access-Request message to another RADIUS server in a specified remote RADIUS server group. If the IAS server receives a valid Access-Accept message that corresponds to the Access-Request message, the connection attempt is considered authenticated and authorized. In this case, the IAS server is being used as a RADIUS proxy.

- Accept the connection attempt without performing authentication or authorization.

  Do not check authentication of the user credentials and authorization of the connection attempt. An Access-Accept message is immediately sent to the RADIUS client. This setting is used for some types of compulsory tunneling where the access client is tunneled before the user's credentials are authenticated. For more information, see IAS and tunnels.

  This authentication option cannot be used when the access client's authentication protocol is MS-CHAP v2 or EAP-TLS, both of which provide mutual authentication. In mutual authentication, the access client proves that it is a valid access client to the authenticating server (the IAS server), and the authenticating server proves that it is a valid authenticating server to the access client. When this authentication option is used, the Access-Accept message is returned. However, the authenticating server does not provide validation to the access client and mutual authentication fails.

For information about configuring authentication settings on a connection request profile, see Configure authentication.

**Accounting**
You can set the following accounting options that are used for RADIUS Accounting-Request messages:

- Forward accounting information in a specific remote RADIUS server group.

  Pass the accounting request to another RADIUS server in a specified remote RADIUS server group. In this case, the IAS server is acting as a RADIUS proxy.

### Note

- IAS always records the accounting information for Accounting-Request messages in local log files on the basis of remote access logging settings. For more information, see Select requests to be logged.

For information about setting accounting settings on a connection request profile, see Configure accounting.

**Attribute manipulation**
You can configure a set of find and replace rules that manipulate the text strings of one of the following attributes:

- User Name
- Called Station ID
- Calling Station ID

Find and replace rule processing occurs for one of the above attributes before the RADIUS message is subject to authentication and accounting settings.

Configuring attribute manipulation for the User Name attribute is equivalent to configuring realm replacement rules for IAS in Windows 2000.

If you are using the MS-CHAP v2 authentication protocol, you cannot manipulate the User Name attribute if the connection request policy is used to forward the RADIUS message. The only exception occurs when a backslash (\) character is used and the manipulation only affects the information to the left of it. A backslash character is typically used to indicate a domain name (the information to the left of the backslash character) and a user account name within the domain (the information to the right of the backslash character). In this case, only attribute manipulation rules that modify or replace the domain name are allowed.

For information about configuring attribute manipulation settings on a connection request profile, see Configure attribute

manipulation. For information about the syntax of find and replace rules, see Pattern matching syntax.

### Notes

- Find and replace rules apply only to a single attribute. You cannot configure find and replace rules for each attribute.
- You cannot add to the list of attributes available for manipulation.

#### Advanced

You can set advanced properties to specify the series of RADIUS attributes that are:

- Added to the RADIUS response message when the IAS server is being used as a RADIUS authentication or accounting server.

  When there are attributes specified on both a remote access policy and the connection request policy, the attributes that are sent in the RADIUS response message are the combination of the two sets of attributes.

- Added to the RADIUS message when the IAS server is being used as a RADIUS authentication or accounting proxy. If the attribute already exists in the message that is forwarded, it is replaced with the value of the attribute specified in the connection request policy.

For information about setting advanced options on a profile, see Configure advanced attributes.

#### The default connection request policy

A default connection request policy named **Use Windows authentication for all users** is created when you install IAS. This policy has the following configuration:

- The **Day-and-Time-Restrictions** condition is set to all times and all days.
- Authentication is set to authenticate requests on this server.
- Accounting is not set to forward accounting information to a remote RADIUS server group.
- Attribute manipulation rules are not configured.
- Advanced attributes are not configured.

The default connection request policy uses IAS as a RADIUS server. To configure an IAS server to act as a RADIUS proxy, you must also configure a remote RADIUS server group. You can create a new remote RADIUS server group while you are creating a new connection request policy with the New Connection Request Policy Wizard. You can either delete the default connection request policy or verify that the default connection request policy is the last policy processed. For an example of a connection request policy that uses IAS as a RADIUS proxy, see Remote RADIUS server.

### Note

- If IAS and the Routing and Remote Access service are installed on the same computer, and the Routing and Remote Access service is configured for Windows authentication and accounting, it is possible for Routing and Remote Access service authentication and accounting requests to be forwarded to a RADIUS server. This can occur when Routing and Remote Access service authentication and accounting requests match a connection request policy that is configured to forward them to a remote RADIUS server group.

# Remote RADIUS server groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote RADIUS server groups**

A remote RADIUS server group is a named group that contains one or more RADIUS servers. When IAS is being used as a RADIUS proxy for RADIUS request messages, a remote RADIUS server group must be specified. This group is used to facilitate the common configuration of both a primary and at least one backup RADIUS server. You can specify various settings to either determine the order in which the servers are used or distribute the RADIUS messages across all servers in the group.

Each server in the group has the following settings:

- Name or address

  Each group member must have a unique name within the group. The name can be an IP address or a name that can be resolved to its IP address. For more information, see Configure the name of a group member.

- Authentication and accounting

  When IAS is used as a RADIUS proxy, it is acting as a RADIUS client to another RADIUS server. Therefore, each group member must be configured to send RADIUS messages to the correct User Datagram Protocol (UDP) port that is used by the RADIUS server for RADIUS traffic. Additionally, each group member must be configured for the correct shared secret. The default authentication port is 1812. The default accounting port is 1813. For more information, see Configure the authentication and accounting settings of a group member.

- Load balancing

  A priority setting is used to indicate which member of the group is the primary server (the priority is set to 1). For group members that have the same priority, a weight setting is used to calculate how often RADIUS messages are sent to each of them. You can use additional settings to configure the way in which the IAS server detects when a group member first becomes unavailable and when it becomes available after it has been determined to be unavailable. For more information, see Configure the load balancing properties of a group member.

After a remote RADIUS server group is configured, it can be specified in the authentication and accounting settings of a connection request policy. Because of this, you can configure a remote RADIUS server group first. Next, you can configure the connection request policy to use the newly configured remote RADIUS server group. Alternately, you can use the New Connection Request Policy Wizard to create a new remote RADIUS server group while you are creating the connection request policy.

For information about creating a remote RADIUS server group, see Add a remote RADIUS server group.

**Note**

- Remote RADIUS server groups are separate from Windows groups.

# Connection Request Processing Examples

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connection request processing examples**
This section covers:

- Remote RADIUS server
- Disabling caller-set callback
- Compulsory tunnels

# Remote RADIUS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote RADIUS server**

To provide RADIUS authentication, authorization, and accounting for an outsourced dial-up and Internet vendor, an IAS server computer that is used as a RADIUS proxy is placed in the organization's perimeter network. To forward RADIUS request messages, where the User-Name attribute contains the realm name example.microsoft.com, to a remote RADIUS server group that corresponds to primary and backup IAS servers inside an organization intranet, the following connection request policy is configured:

1. Use the New Remote RADIUS Server Group Wizard to create a new remote RADIUS server group named **Intranet IAS servers**. Configure the group with primary and backup servers that correspond to the two IAS server computers in the organization intranet.

   For more information, see Add a remote RADIUS server group.

2. Use the New Connection Request Policy Wizard to create a new connection request policy named **Forward to intranet IAS servers**. Configure the policy to forward RADIUS requests, configure the realm name of example.microsoft.com (but do not remove the realm name before authentication), and select the remote RADIUS server group named **Intranet IAS servers**.

   For more information, see Add a connection request policy.

Based on this connection request policy, all RADIUS request messages that contain the realm name example.microsoft.com in the User-Name attribute are forwarded to an IAS server in the organization intranet.

- Delete the default policy named **Use Windows authentication for all users**.

For more information, see Delete a connection request policy.

# Disabling caller-set callback

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disabling caller-set callback**

In this example, a network administrator wants to disable the use of callback when callback settings are determined by the dial-up access client (also known as caller-set callback). If caller-set callback is not being used, do not disable or modify the use of callback.

To implement this connection request policy example, complete the following steps:

1. Use the New Connection Request Policy Wizard to create a new common connection request policy named **Forward to intranet IAS servers**. Configure the policy to forward RADIUS requests, configure the realm name of example.microsoft.com, and select the remote RADIUS server group named **Intranet IAS servers**. Do not remove the realm name before authentication.

   For more information, see Add a connection request policy.

2. Use the New Remote RADIUS Server Wizard to create a new remote RADIUS server group named **Intranet IAS servers**. Configure the group with primary and backup servers that correspond to the two IAS server computers in the organization intranet.

   For more information, see Add a remote RADIUS server group.

3. Create a new custom connection request policy with the following settings:

   - Policy name: Forward to intranet IAS servers but remove caller-set callback
   - Conditions: User-Name matches *example.microsoft.com*; Service-Type matches Callback Login
   - Profile settings, **Authentication** tab: Select **Forward requests to the following remote RADIUS server group for authentication**, and then select the **Intranet IAS servers** group.
   - Profile settings, **Accounting** tab: Select **Record accounting information on the servers in the following remote RADIUS server group**, and then select the **Intranet IAS servers** group.
   - Profile settings, **Attribute** tab: Select the **User-Name** attribute and add a find and replace rule that finds **example.microsoft.com** and replaces it with nothing.
   - Profile settings, **Advanced** tab: Add the **Service-Type** attribute with the value **Framed**.

   For more information, see Add a connection request policy.

4. Delete the default policy named **Use Windows authentication for all users**.

   For more information, see Delete a connection request policy.

Based on this connection request policy, all RADIUS request messages that contain the realm name example.microsoft.com in the User-Name attribute are forwarded to an IAS server in the organization intranet. When the Access-Accept message is sent to the IAS server computer in the perimeter network, the Callback-ID and Callback-Number RADIUS attributes in the message are set to blank values--effectively disabling callback.

# Compulsory tunnels

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Compulsory tunnels**
In this example, a network administrator wants to enable the use of compulsory tunnels for organization employees when they dial an outsourced Internet vendor.

To provide RADIUS authentication, authorization, and accounting for the outsourced dial-up and Internet vendor, an IAS server computer, used as a RADIUS proxy, is placed on the organization's perimeter network. A connection request policy is configured on the IAS server computer to forward RADIUS request messages, where the User-Name attribute contains the realm name example.microsoft.com, to a remote RADIUS server group that corresponds to primary and backup IAS servers inside an organization intranet. The policy is also configured to add the set of attributes required in creating a compulsory tunnel to a VPN server with the Layer Two Tunneling Protocol (L2TP). The VPN server has the IP address of 131.107.9.41 and uses the example tunnel password of og*37y#cW@95?4xT.

To implement this connection request policy example, complete the following steps:

1. Use the New Remote RADIUS Server Wizard to create a new remote RADIUS server group named **Intranet IAS servers**. Configure the group with primary and backup servers that correspond to the two IAS server computers in the organization intranet.

   For more information, see Add a remote RADIUS server group.

2. Use the New Connection Request Policy Wizard to create a new connection request policy named **Forward to intranet IAS servers**. Configure the policy to forward RADIUS requests and the realm name of example.microsoft.com, and then select the remote RADIUS server group named **Intranet IAS servers**. Do not remove the realm name before authentication.

   For more information, see Add a connection request policy.

3. Configure the advanced properties for the profile of the connection request policy named **Forward to intranet IAS servers** with the following:

   - Add the **Framed-Protocol** attribute with the value set to **PPP**.
   - Add the **Service-Type** attribute with the value set to **Outbound-User**.
   - Add the **Tunnel-Medium-Type** attribute with the value set to **IP**.
   - Add the **Tunnel-Password** attribute with the value set to **og*37y#cW@95?4xT**.
   - Add the **Tunnel-Server-Endpt** attribute with the value set to **131.107.9.41**.
   - Add the **Tunnel-Type** attribute with the value set to **Layer Two Tunneling Protocol (L2TP)**.

     For more information, see Configure advanced attributes.

4. Delete the default policy named **Use Windows authentication for all users**.

   For more information, see Delete a connection request policy.

Based on this connection request policy, all RADIUS request messages that contain the realm name example.microsoft.com in the User-Name attribute are forwarded to an IAS server in the organization intranet. When the Access-Accept message is sent back to the IAS server computer in the perimeter network, the set of attributes that are used to create a compulsory tunnel are

added to the message before they are returned to the outsourced vendor's Front End Processor (FEP).

For more information about compulsory tunneling, see IAS and tunnels.

# Remote Access Logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote access logging**

This section contains information about remote access logging in IAS. There are three types of logging:

1. Event logging for IAS

   Event logging is the recording of IAS events in the system event log. It is used primarily for auditing and troubleshooting connection attempts.

2. Logging user authentication and accounting requests

   Recording user authentication and accounting requests in log files is used primarily for connection analysis and billing purposes, and is also useful as a security investigation tool, providing you with a method of tracking down the activity of a malicious user after an attack.

3. SQL Server database logging

   Recording user authentication and accounting requests in a SQL Server XML-compliant database enables multiple IAS servers to share one data source, and provides the advantages of using a relational database.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Event logging for IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Event logging for IAS**

IAS errors are automatically recorded in the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition system event logs. IAS authentication events are recorded in the system event log on the basis of event logging settings. You can individually enable the logging of:

- Rejected or discarded authentication events

  These events are useful for understanding why a particular request was rejected or ignored. Requests can be rejected or ignored for many reasons, including the following:

  - The RADIUS message is not formatted according to RFCs 2865 or 2866.
  - The RADIUS client is unknown.
  - The RADIUS client has multiple IP addresses and sent the request on an address other than the one defined in IAS.
  - The shared secret is invalid.
  - The message authenticator (also known as a digital signature) sent by the client is invalid.
  - IAS was unable to locate the user name's domain.
  - IAS was unable to connect to the user name's domain.
  - IAS was unable to access the user account in the domain.

  If the request was rejected, the information in the event text includes the user name, access server identifiers, the authentication type, the name of the first matching remote access policy, the reason for the rejection, and other information.

  The logging of rejected or discarded authentication events is enabled by default.

  Secure channel (Schannel) is a security support provider (SSP) that supports a set of Internet security protocols, such as Secure Sockets Layer (SSL) and Transport Level Security (TLS). These protocols provide identity authentication and secure, private communication through encryption.

  Logging of client certificate validation failures is a secure channel event, and is not enabled on the IAS server by default. You can enable additional secure channel events by changing the following registry key value from **1** (**REG_DWORD** type, data **0x00000001**) to **3** (**REG_DWORD** type, data **0x00000003**):

  **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\EventLogging**

  **Caution**

  - Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

  For more information, see The Registry.

- Successful authentication requests

If the request was accepted, the information in the event text includes the user name, access server identifiers, the authentication type, and the name of the first matching remote access policy.

The logging of successful authentication events is enabled by default.

For information about configuring event logging options for IAS, see Configure event logging for IAS.

For information about how to view event logs in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; see Event Viewer.

For more information about troubleshooting IAS, see IAS Troubleshooting.

# Logging user authentication and accounting requests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logging user authentication and accounting requests**
You can use IAS to create log files based on the authentication and accounting requests received from access servers, and collect this information in a central location. By setting up and using log files to track authentication information, such as each connection accept and connection reject, you can simplify administration of your service. You can set up and use logs to track accounting information (such as logon and logoff records) to maintain records for billing purposes.

When you set up logging, you can specify:

- Which requests are logged.
- The log-file format.
- How often new logs are started.
- Automatic deletion of the oldest log file when the disk is full.
- Where log files are recorded.
- What the log file records contain.

You can use the IAS console to specify which requests are logged. For more information, see Select requests to be logged.

You can use the IAS console to specify the log-file format, how often new logs are started, and where log files are recorded. For more information, see Configure log file properties.

**Which requests are logged**
You can select the following types of requests to log:

- Accounting requests, including the following:
    - Accounting-on requests, which are sent by the access server to indicate that it is online and ready to accept connections.
    - Accounting-off requests, which are sent by the access server to indicate that it is going offline.
    - Accounting-start requests are sent by the access server (after the user is accepted by the IAS server) to indicate the start of a user session.
    - Accounting-stop requests are sent by the access server to indicate the end of a user session.
- Authentication requests, including the following:
    - Authentication requests are sent by the access server on behalf of the connecting user. These entries in the log contain only incoming attributes.
    - Authentication accepts and rejects, which are sent by IAS to the access server, indicating whether the user should be accepted or rejected. These entries contain only outgoing attributes.
- Periodic status, which is used to obtain interim accounting requests sent by some access servers during sessions.

  Accounting-interim requests, which are sent periodically by the access server during a user session. This type of request can be used when the Acct-Interim-Interval RADIUS attribute is configured to support periodic requests in the remote

access profile on the IAS server.

### Notes

- All types of request logging are disabled by default.
- Initially, it is recommended that you enable the logging of accounting and authentication requests. You can refine your logging methods after you determine which data best matches your needs.
- You can also log accounting, authentication, and periodic status to a SQL Server database. For more information, see SQL Server database logging.

#### The log-file format

You can configure your IAS servers to log data with either the IAS format or the database-import log format.

- If you select the IAS format, attributes are logged in the form of attribute-value pairs. This format has the following characteristics:
    - The sequence of the attributes is dependent on the access server that is sending a request.
    - The logged attributes include RADIUS standard, IAS-specific, and vendor-specific.
    - All attributes that contain unprintable characters or any delimiters are printed in hexadecimal format (for example, 0x026).
- If you select the database-import log format, attributes are logged in a format that supports importing the log into databases. This format has the following characteristics:
    - The attributes of all records are recorded in the same sequence (predefined by IAS), regardless of which access server sent the request.
    - If the attribute is not present in the request or reply (for example, if an attribute is not received from the access server), then the field in the log is empty.
    - The specific set of attributes that are logged is in a sequence that is predefined in IAS. Although the attributes included in this set are limited in number, they are those that are generally most useful in request tracking and analysis.

### Notes

- IAS format is enabled by default.
- Switching log file formats does not cause a new log to be created. If you change log file formats, the file that is active at the time of the change will contain a mixture of the two formats (that is, records at the start of the log have the previous format and records at the end of the log have the new format).

#### How often new logs are started

When you configure your servers, you can specify whether new logs are started daily, weekly, monthly, or after the log reaches a specific size. You can also specify that a single log is maintained continually (regardless of file size), but this is not recommended. The file naming convention for logs is determined by the amount of time that each log is used, which you specify by setting the change frequency. If the log file change frequency is **Daily**, the log file name format is **in***yymmdd* (where *y* is year, *m* is month, and *d* is day).**log**. If the log file change frequency is **Weekly**, the log file name format is **in***yymmww*.**log** (where *w* is week). If the log file change frequency is **Monthly**, the log file name format is **in***yymm*.**log**. Following are some example log file names:

| Log file name | Change frequency | File creation date |
|---|---|---|
| in020528.log | Daily | May 28, 2002 |
| in020304.log | Weekly | The fourth week of March, 2002 |
| in0304.log | Monthly | April, 2003 |

**Notes**

- By default, the new log file change frequency is set to **Monthly**.
- When the log file change frequency is **Never (unlimited file size)**, the IAS log file is named **iaslog.log**.
- When the log file change frequency is **When log file reaches this size**, the file name is **iaslog***n***.log**, where *n* represents the maximum size of the file.

**Automatic deletion of the oldest log file when the disk is full**

You can specify whether IAS deletes the oldest log file when the disk is full. IAS determines which file is oldest by checking the file name rather than the time stamp. The currently configured log file change frequency (where new files are created at either a specified time or when the current file reaches a specified size) determines which log file IAS deletes.

IAS only deletes the oldest file from the same change frequency--it does not delete a file with a different file name format. For example, if the currently configured log file change frequency is set to **Monthly** (which has a log file name format of **in***yymm***.log**), IAS will not delete a log file named in021231.log even if it is the oldest log file, because the log file name and name format indicate that it was created with the change frequency set to **Daily**. The only time that IAS might delete a file from another change frequency occurs when the two change frequency settings are **Daily** and **Weekly**, and the file name can be interpreted both as one of the first four days of the month and as one of the four weeks of the month, as demonstrated in the following examples:

| The file name | Is interpreted as both |
|---|---|
| in020501.log | A **Daily** file created on May 1, 2002 or a **Weekly** file created in the first week of May, 2002. |
| in020502.log | A **Daily** file created on May 2, 2002 or a **Weekly** file created in the second week of May, 2002. |
| in020503.log | A **Daily** file created on May 3, 2002 or a **Weekly** file created in the third week of May, 2002. |
| in020504.log | A **Daily** file created on May 4, 2002 or a **Weekly** file created in the fourth week of May, 2002. |

**Notes**

- If the oldest log file is the log file that is currently in use, the log file is not deleted.
- Because changing the log file change frequency might result in the overwriting of existing logs, you should copy logs to a separate file before making the change.
- For more information about how to effectively configure and use log files, see IAS Best Practices and Configure log file properties.

**Where log files are recorded**

By default, log files are located in the *systemroot*\System32\LogFiles folder. You can specify a different location.

The log file directory can be created by using system environment variables (instead of user variables) such as, %systemdrive%, %systemroot%, and %windir%. For example, the following path, using the environment variable %windir%, locates the log file at the system directory in the subfolder \System32\Logs (that is, %windir%\System32\Logs\).

For information about how to set up the authentication and accounting logging service on an IAS server, see Configure Logging for User Authentication and Accounting.

For information about how to import a log file into a database, see Importing IAS log files into a database.

**Log capture for other processes**
To send log data directly to another process, you can configure IAS to write to a named pipe, instead of a file. To use named pipes, set the log file folder to **\\.\pipe** or *****ComputerName***\pipe**. The named pipe server program should create a named pipe called **\\.\pipe\iaslog.log** to accept the data.

**What the log file records contain**
Attributes are recorded in UTF-8 encoding in a comma-delimited format. The format of the records in a log file depends on the file format.

In IAS-formatted log files, each record starts with a fixed-format header, which consists of the access server IP address, user name, record date, record time, service name, and computer name, followed by attribute-value pairs.

For more information about the attributes and other data logged in the IAS format, see IAS-Formatted Log Files.

In database-import log files, each record contains attribute values in a consistent sequence, including the computer name, service name, record date, and record time. An access server might not use all of the attributes specified in the database-import log format, but the comma-delimited location for each of these pre-defined attributes is maintained, even for attributes that have no value specified in a record.

For more information about the attributes and other data logged in the database-import log format, see Database-Import Log Files.

### Note

- This documentation contains logging information that is specific to IAS. The types of data that your access server sends might vary, depending upon the manufacturer. For information about the specific data that your access server sends and how to configure accounting packets, see your access server documentation.

# SQL Server database logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SQL server database logging**

You can use SQL server database logging to log user authentication and accounting requests, which are received from one or more network access servers, to a centralized data source. Log data is passed from IAS to a stored procedure in a database that supports both structured query language (SQL) and extensible markup language (XML), such as one created with Microsoft SQL Server 2000. A stored procedure is a custom program, created by a SQL server database programmer, that runs within the SQL server database environment.

Logging to a relational database instead of a standard text file (that is, either IAS format or database-compatible format) has the advantage that relationships between data tables enable the flexible creation of dynamic data views by using queries and reports.

**Which requests are logged**

By default, no data is logged. When you configure logging, all required attributes, accounting, and authentication data that is normally logged in either IAS or database-compatible format is logged to the SQL server database. You can log the following information in a SQL server database:

- Accounting requests, including the following:
    - Accounting-on requests, which are sent by the access server to indicate that it is online and ready to accept connections.
    - Accounting-off requests, which are sent by the access server to indicate that it is going offline.
    - Accounting-start requests, which are sent by the access server (after the user is accepted by the IAS server) to indicate the start of a user session.
    - Accounting-stop requests, which are sent by the access server to indicate the end of a user session.
- Authentication requests, including the following:
    - Authentication requests, which are sent by the access server on behalf of the connecting user. These entries in the log contain only incoming attributes.
    - Authentication accepts and rejects, which are sent by IAS to the access server, indicating whether the user should be accepted or rejected. These entries contain only outgoing attributes.
- Periodic status, which is used to obtain interim accounting requests sent by some access servers during sessions.

Accounting-interim requests, which are sent periodically by the access server during a user session, can also be logged to a SQL server database. This type of request can be used when the Acct-Interim-Interval RADIUS attribute is configured to support periodic requests in the remote access profile on the IAS server.

**Attributes that are recorded by SQL server**

Unlike database-import log files, which use a fixed sequence of attributes, the sequence of the attributes in a SQL server log file depends upon the sequence used by the access server. For more information about the sequence of these records, see the documentation for the access server. The table in Interpreting IAS-formatted log files lists, in numerical order, RADIUS attributes that can be found in a SQL server database and configured for IAS logging. This table does not cover vendor-specific attributes (VSAs). For information about VSAs, see Interpreting IAS IDs for vendor-specific attributes. For more information about VSAs that are supported by your access server, see your access server documentation.

**Attributes that are not recorded by SQL server**

Although most of the attributes that are sent by an access server and supported in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition are logged when you specify the use

of SQL server logging, some attributes are not logged because they contain sensitive information (for example, passwords) that should not be included. The table in Interpreting IAS-formatted log files lists some of the attributes that are not logged.

You can use the IAS console to specify which requests are logged. For more information, see Select requests to be logged.

You can use the IAS console to specify additional SQL server logging properties. For more information, see Configure log file properties.

**Notes**

- All types of request logging are disabled by default.
- Initially, it is recommended that you enable the logging of accounting and authentication requests. You can refine your logging methods after you determine your required data.
- You can also log accounting, authentication, and periodic status in IAS or database-compatible format. For more information, see Logging user authentication and accounting requests.

**How IAS and SQL server logging works**

When IAS is configured to log to a SQL server database, the following occurs:

1. The IAS server receives or produces the data that it is configured to log to a SQL server database. For example, a network access server configured as a RADIUS client to the IAS server sends the IAS server incoming attributes for an authentication request by a user or computer.
2. The IAS server packages the data as an XML document.
3. The IAS server initiates a connection with the server running SQL server. If the IAS server is unable to establish a connection with SQL server, the IAS server stops processing authentication requests and users cannot log on to the network.

   IAS connects to SQL server using password-based credentials that you can configure on the **Connection** tab in the **Data Link Properties** dialog box.

   If you select **Use Windows NT integrated security**, the IAS server attempts to authenticate the connection with the local computer account. If you select **Use a specific name and password**, you must specify the user name and password for a valid user account that you have already configured on your SQL server.

   For more information about configuring SQL server data link properties, see Configure log file properties.
4. If a connection with the computer running SQL server is successful, the IAS server signals SQL server to run a stored procedure named **report_event**. IAS then sends the accounting data (an XML document) to the computer running SQL server. The stored procedure that is called by IAS must be named **report_event**, or logging will fail.

   The **report_event** stored procedure uses only one parameter. This parameter is a Unicode string of type **ntext**. The Unicode string contains an XML document which in turn contains the accounting data to be logged.
5. The **report_event** stored procedure in the SQL server database processes the XML document and stores the processed data in SQL server database tables.
6. SQL server notifies the IAS server that the transaction is successful.
7. The IAS server is free to process additional data, repeating this process as required.

**How IAS creates an XML document from accounting and authentication data**

If you select the SQL server log format, attribute-value pairs are converted to XML format.

When IAS creates an XML document from attributes, accounting and authentication data, the XML document includes the attribute ID or name, the attribute value, and the data type of the attribute value. With SQL server logging, there are five date

types for attribute values: Non-negative integers, strings (text), hexadecimal numbers, IPv4 addresses, and date and time. The following example XML document, which is created by IAS, includes the attribute name (User-Name), the attribute value (DOMAIN\username), and the data type of (1), indicating a value data type of string:

```
<Event>
    <User-Name data_type="1">DOMAIN\username</User-Name>
</Event>
```

The next example XML document includes the attribute name (NAS-IP-Address), the attribute value (192.168.0.1), and the data type (3), indicating a value data type of IP address:

```
<Event>
    <NAS-IP-Address data_type="3">192.168.0.1</NAS-IP-Address>
</Event>
```

The next example XML document includes the attribute name (Provider-Type), the attribute value (1), and the data type of (0), indicating a value data type of non-negative integers:

```
<Event>
    <Provider-Type data_type="0">1</Provider-Type>
</Event>
```

The last example XML document, which is created by IAS, includes the three previous examples combined into one typical XML document:

```
<Event>
    <User-Name data_type="1">DOMAIN\username</User-Name>
    <NAS-IP-Address data_type="3">192.168.0.1</NAS-IP-Address>
    <Provider-Type data_type="0">1</Provider-Type>
</Event>
```

**Configuring SQL server database logging**

To log information to a SQL server database, you must configure the SQL server data link in **Remote Access Logging** in the IAS console. For more information, see Configure log file properties.

To provide centralized logging, failover, and redundancy with SQL server logging, you can place two SQL server databases on different subnets. Use the SQL server Create Publication Wizard to set up database replication between the two servers.

For larger IAS deployments, you should use more than two computers running SQL server and you should not configure all IAS servers to log to the same one. When you have a primary and backup IAS server configuration, the primary IAS server should log to a computer running SQL server that is other than the one to which the backup IAS server logs. With this configuration, if the primary computer running SQL server fails, preventing the primary IAS server from servicing clients, the backup IAS server, which is configured to log to a different computer running SQL server, can service clients until the primary server is back on line.

If the computer running SQL server runs out of disk space and can no longer perform IAS logging, IAS discards all accounting and authentication requests, preventing client authentication through the IAS server. Make sure you maintain adequate hard disk space on your computer running SQL server.

The performance of the SQL server database (throughput and latency) affects the response time of the IAS server. During an authentication attempt, IAS does not respond to the RADIUS client (network access server) until after SQL server notifies the IAS server that the last data write operation is successful. Because of this, you should ensure that the computer running SQL server is designed to handle the IAS logging load and is able to respond with minimum delay.

In addition, some RADIUS clients might retry authentication attempts if a response is not received from the IAS server within a specific time period (for example, one second). This could add to the load on the computers running IAS and SQL server during periods of slow response due to network conditions. You should configure the network access servers with an appropriate retry interval to prevent overloading the computers running IAS and SQL server with retry attempts.

For more information about SQL server configuration and best practices, see the SQL server 2000 documentation.

**Requirements**

Logging to a SQL server database requires the use of one or more servers running Windows Server 2003 and IAS and an SQLXML-compliant database, such as Microsoft SQL server 2000. Because this feature requires SQLXML, Microsoft Data Engine (MSDE) cannot be used instead of SQL server.

# Deploying IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying IAS**
This section covers:

- Deploying IAS as a RADIUS Server
- Deploying IAS as a RADIUS Proxy
- Setting up IAS in a Test Lab

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Deploying IAS as a RADIUS Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying IAS as a RADIUS server**
This section covers:

- Planning for IAS as a RADIUS Server
- IAS as a RADIUS server

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Planning for IAS as a RADIUS Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for IAS as a RADIUS server**
This section covers:

- IAS as a RADIUS server design considerations
- IAS as a RADIUS server security considerations
- Computer certificates for certificate-based authentication

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# IAS as a RADIUS server design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS as a RADIUS server design considerations**

Consider the following design issues when deploying IAS as a RADIUS server:

- The common uses of IAS as a RADIUS server

  For more information about common ways to use IAS as a RADIUS server, see IAS as a RADIUS server.

- IAS server domain membership

  You must decide in which domain the IAS server computer is a member. For multiple domain environments, an IAS server can authenticate credentials for user accounts in the domain of which it is a member and all domains that trust this domain. To read the dial-in properties for user accounts, however, you must add the computer account of the IAS server to the RAS and IAS servers group for each domain. For more information, see Enable the IAS server to read user accounts in Active Directory. Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition also provide authentication across forests. For more information, see Accessing resources across forests.

- RADIUS ports

  You can configure the IAS server to receive RADIUS messages that are sent to UDP ports other than the default ports of 1812 and 1645 (for RADIUS authentication) and ports 1813 and 1646 (for RADIUS accounting). For more information, see Configure IAS port information.

- RADIUS clients

  A RADIUS client can be an access server (for example, a dial-up or VPN server, a wireless access point, or an Ethernet switch) or a RADIUS proxy. IAS supports all access servers and RADIUS proxies that comply with RFC 2865, "Remote Authentication Dial-in User Service (RADIUS)."

  Configure each access server or RADIUS proxy that sends RADIUS request messages to the IAS server as a RADIUS client on the IAS server. For each RADIUS client, you can configure a friendly name, an IP address or DNS name, the client vendor, the shared secret, and whether to use the RADIUS Message Authenticator attribute. For more information, see Configure RADIUS Clients.

  You can specify IP addresses or DNS names for RADIUS clients. In most cases, it is better to specify RADIUS clients with IP addresses. When you use IP addresses, IAS is not required to resolve host names at startup and will start much more quickly as the result. This is especially beneficial if your network contains a large number of RADIUS clients. You can use DNS names to specify RADIUS clients when you require something other than administrative flexibility (for example, the ability to map multiple RADIUS client IP addresses to a single DNS name).

  IAS in Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition allows you to specify a RADIUS client by using an address range. The address range for RADIUS clients is expressed in the network prefix length notation *w.x.y.z/p*, where *w.x.y.z* is the dotted decimal notation of the address prefix and *p* is the prefix length (the number of high order bits that define the network prefix). This is also known as Classless Inter-Domain Routing (CIDR) notation. An example is 192.168.21.0/24. To convert from subnet mask notation to network prefix length notation, *p* is the number of high order bits in the subnet mask that are set to 1.

- Third-party access points

  To determine whether a third-party access point is interoperable with IAS as a RADIUS server, refer to the third-party access point documentation for its use of RADIUS attributes and vendor-specific attributes.

To test basic interoperability for PPP connections, configure the access point and the access client to use Password Authentication Protocol (PAP). Use additional PPP authentication protocols until you have tested the ones that you are using for network access.

- Connection request policy configuration

  The default connection request policy named **Use Windows authentication for all users** is configured for IAS when it is used as a RADIUS server. For more information about the settings of the default connection request policy, see Connection request policies. Additional connection request policies can be used to define more specific conditions, manipulate attributes, and specify advanced attributes. Use the New Connection Request Policy Wizard to create either common or custom connection request policies. For more information about connection request policies and how to configure them when IAS is being used as a RADIUS server, see Introduction to connection request processing.

- Realm replacement

  To correctly replace or convert realm names within the user name of a connection request, you must configure attribute manipulation rules for the User-Name attribute on the appropriate connection request policy. For more information, see Connection request policies and Configure attribute manipulation.

- Remote access policy configuration

  Remote access policies can be used to define specific conditions, set dial-in constraints, set allowed authentication protocols and encryption strength, and specify advanced attributes. Use the New Remote Access Policy Wizard to create either common or custom policies. For more information, see Introduction to remote access policies.

- Remote access policies and authorization by user or group

  In small organizations, you can manage authorization by setting the remote access permission on each user account. For a large organization, set the remote access permission on each user account to be controlled through remote access policy. Next, configure remote access policies to grant access by using group membership. For more information, see Introduction to remote access policies.

- The use of additional RADIUS attributes and vendor-specific attributes

  If you plan to return additional RADIUS attributes or vendor-specific attributes (VSAs) with the responses to RADIUS requests, you must add the RADIUS attributes or VSAs to the appropriate remote access policy. For more information, see Vendor-specific attribute overview and Configure Attributes for a Profile.

- The use of logging

  You can enable event logging for authentication events to assist with troubleshooting and debugging connection attempts. For more information, see Event logging for IAS. For information about logging user authentication and accounting requests for analysis and connection billing, see Logging user authentication and accounting requests.

- Interim accounting

  To use interim accounting, first verify that your access server supports sending interim accounting messages. Next, add the Acct-Interim-Interval RADIUS attribute from the **Advanced** tab on the profile settings of the appropriate remote access policy. Configure the Acct-Interim-Interval attribute with the interval (in minutes) to send periodic interim accounting messages. For more information, see Add RADIUS attributes to a remote access policy. Next, enable logging of periodic status. For more information, see Select requests to be logged.

  The Routing and Remote Access service supports sending interim accounting messages and is enabled from the properties of a RADIUS server when configuring the RADIUS accounting provider. For more information, see Use RADIUS accounting.

- Performance and capacity planning

For information about tuning performance and using IAS in large organizations, see IAS Best Practices.

- Multiple IAS servers

  To provide fault tolerance for RADIUS-based authentication and accounting, you should always use at least two IAS servers. One IAS server is used as the primary RADIUS server and the other is used as a backup. Access servers are configured for both IAS servers. They switch to using the backup IAS server when the primary IAS server becomes unavailable. For information about how to synchronize the configuration of multiple IAS servers, see Managing multiple IAS servers.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# IAS as a RADIUS server security considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS as a RADIUS server security considerations**

Consider the following security issues when deploying IAS as a RADIUS server:

- Shared secrets

  Configure strong shared secrets and change them frequently to prevent dictionary attacks. Strong shared secrets are a long (more than 22 characters) sequence of random letters, numbers, and punctuation. For more information, see Shared secrets.

- Message Authenticator attribute

  To ensure that an incoming RADIUS Access-Request message, for connection requests that use the PAP, CHAP, MS-CHAP, and MS-CHAP v2 authentication protocols, was sent from a RADIUS client configured with the correct shared secret, you can use the RADIUS Message Authenticator attribute (also known as a digital signature or the signature attribute). You must enable the use of the Message Authenticator attribute on both the IAS server (as part of the configuration of the RADIUS client in Internet Authentication Service) and the RADIUS client (the access server or RADIUS proxy). Ensure that the RADIUS client supports the Message Authenticator attribute before enabling it. The Message Authenticator attribute is always used with EAP without having to enable it on the IAS server and access server. For more information, see Edit RADIUS client configuration.

  For information about enabling the RADIUS Message Authenticator attribute for your access server, see the appropriate access server documentation. For the Routing and Remote Access service, the use of the RADIUS Message Authenticator attribute is enabled from the properties of a RADIUS server when you configure the RADIUS authentication provider. For more information, see Use RADIUS authentication.

- Firewall configuration

  If your IAS server is on a perimeter network (also known as a demilitarized zone or DMZ), configure your Internet firewall (between your perimeter network and the Internet) to allow RADIUS messages to pass between your IAS server and RADIUS clients on the Internet. You might need to configure an additional firewall that is placed between your perimeter network and your intranet, to allow traffic to flow between the IAS server on the perimeter network and domain controllers on the intranet. For more information, see IAS and firewalls.

- Authentication protocols

  IAS supports several different authentication protocols. The order of authentication protocols, from the most secure to the least secure, is: PEAP-EAP-TLS (for wireless clients and authenticated switch clients only), EAP-TLS, PEAP-EAP-MS-CHAPv2 (for wireless clients and authenticated switch clients only), MS-CHAP v2, MS-CHAP, EAP-MD5, CHAP, and PAP. Microsoft recommends using only the strongest authentication protocols required for your configuration. For password-based authentication protocols, strong password policies must be enforced to protect from dictionary attacks. The use of PAP is not recommended unless it is required. For more information, see Authentication methods.

- Remote access account lockout

  To provide protection for online dictionary attacks launched against access servers by using known user names, you can enable remote access account lockout. Remote access account lockout disables remote access for user accounts after a configured number of failed connection attempts has been reached. For more information, see Remote access account lockout.

  Remote access account lockout can also be used to prevent a malicious user from intentionally locking out a domain account by attempting multiple dial-up or VPN connections with the wrong password. You can set the number of failed attempts for remote access account lockout to a number that is lower than the logon retries for domain account lockout.

By doing this, remote access account lockout occurs before domain account lockout, which prevents the domain account from being intentionally locked out.

- Certificate-based authentication

  When you use the EAP-TLS authentication protocol, you must install a computer certificate on the IAS server. For client and user authentication, you can either install a certificate on the client computer or use smart cards. Before certificate enrollment occurs, the certificate must be designed with the correct requirements and purposes. For more information, see Network access authentication and certificates and Computer certificates for certificate-based authentication.

- Authentication of wireless clients using Protected Extensible Authentication Protocol (PEAP)

  You can use PEAP with EAP-TLS (also known as PEAP-EAP-TLS) to deploy a certificates with PEAP. PEAP with EAP-MS-CHAPv2 (also known as PEAP-EAP-MS-CHAPv2) provides secure password authentication. PEAP-EAP-TLS uses a public key infrastructure (PKI) with certificates for server authentication and either smart cards or certificates for client and user authentication. When you use PEAP-EAP-TLS, client certificate information is encrypted.

  Although using PEAP-EAP-TLS with smart cards for certificate deployment is the most secure authentication method, the complexity and cost of deploying certificates to wireless and authenticated switch clients might not be practical for your organization. PEAP-EAP-MS-CHAPv2 balances security with deployment cost and complexity. Unlike MS-CHAPv2, PEAP-EAP-MS-CHAPv2 is a mutual authentication method that uses Transport Level Security (TLS) to create an end-to-end encrypted connection between client and authenticator. The client authenticates the server by using the server's certificate and the server authenticates the client with password-based credentials. For an example of a remote access policy using PEAP, see Wireless access with secure password authentication.

- Registration of the IAS server in Active Directory

  Before the IAS server can access Active Directory domains to authenticate user credentials and user access account properties, the IAS server must be registered in those domains. For more information, see Enable the IAS server to read user accounts in Active Directory.

- Using IPSec filters to lock down IAS servers

  You can configure IPSec filters on a granular level to allow specific traffic in and out of network interfaces on RADIUS servers. These filters can be applied to organizational units and stored in Active Directory, or they can be created and applied to individual servers. For more information, see Securing RADIUS traffic with IPSec.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Computer certificates for certificate-based authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Computer certificates for certificate-based authentication**
Before using Extensible Authentication Protocol-Transport Level Security (EAP-TLS), you must install a computer certificate (also known as a machine certificate) on the Internet Authentication Service (IAS) server computer. The installed computer certificate must be issued from a certification authority (CA) that can follow a certificate chain to a root CA that is trusted by the access clients. Additionally, in order for the IAS server to validate the user or computer certificate of the access client, the IAS server must install the certificate of the root CA that issued the user or computer certificate to the access clients.

For more information, see Network access authentication and certificates.

To install a computer certificate, a CA must be available to issue certificates. After the CA is configured, you can install a computer certificate on the IAS server:

- By configuring the automatic allocation of computer certificates to computers in an Active Directory domain.
- By using Certificate Manager to obtain a computer certificate.
- By using Microsoft Internet Explorer and Web-based enrollment.

To configure a CA and install the computer certificate, complete the following steps:

1. Install the Certificate Services component as an enterprise root CA. For more information, see Install an enterprise root certification authority. This step is only required if you do not already have an enterprise root CA.

   If necessary, configure the computer that is the CA to be a domain controller.

2. To install the computer certificates through auto-enrollment, configure Group Policy on the Active Directory domain for automatic allocation of computer certificates. For more information, see Configure automatic certificate allocation from an enterprise CA.

   To create a computer certificate for the IAS server that is a member of the domain for which auto-enrollment is configured (and other computers that are members of the domain), restart the computer or type **gpupdate /target:computer** from a Command Prompt.

3. To manually enroll computer certificates, use Certificate Manager to install both the CA root certificate of the access client and the computer certificate for the IAS server. For more information, see Manage certificates for a computer and Request a certificate.

4. To manually enroll computer certificates with Internet Explorer, see Submit an advanced certificate request via the Web.

**Notes**

- While Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition allow you to install multiple computer certificates, only a single certificate can be selected for all remote access policies that specify authentication through EAP-TLS.
- An IAS server does not use a new certificate revocation list (CRL), which is published by the certification authority (CA), until the old CRL has expired. If a certificate has expired and the IAS server is checking against an old but unexpired CRL, the IAS server will allow clients with revoked certificates to connect to the network. To prevent this from happening, you can publish CRLs with short expiration times (one hour or longer).
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address,

but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Configuring IAS as a RADIUS Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS as a RADIUS server**
This section covers:

- Wireless access
- Dial-up and VPN remote access
- Extranet access
- Internet access
- Outsourced VPN remote access
- Outsourced wireless access
- Ethernet switch access with IAS

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Wireless access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless access**

This topic describes how IAS can be used to support authentication, authorization, and accounting for wireless connections to an organization. This topic describes a typical configuration for an organization that uses:

- Two IAS servers.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication. If only one RADIUS server is configured and it becomes unavailable, wireless access clients cannot connect. By using two IAS servers and configuring all wireless access points (RADIUS clients) for both the primary and secondary IAS servers, the RADIUS clients can detect when the primary RADIUS server is unavailable and automatically fail over to the secondary IAS server.

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints. To both optimize IAS authentication and authorization response times and minimize network traffic, IAS is installed on domain controllers.

- A certificate infrastructure.

  The Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication protocol is used with locally installed user certificates to authenticate wireless clients.

- Wireless remote access policies.

  Remote access policies are configured for wireless connections so that authenticated wireless users obtain unrestricted access to the organization intranet and guest wireless users obtain restricted access to a virtual LAN (VLAN) that contains a certificate server.

- Multiple wireless access points.

  Multiple third-party wireless access points provide wireless access in different buildings of an organization.

- A certificate server on a separate VLAN for new wireless clients

  To accommodate new wireless clients, a certificate server is placed on a separate virtual local area network (VLAN). The certificate server is used by a new wireless client to obtain a user certificate for subsequent wireless access. After 10 minutes of receiving a guest connection, the wireless client is disconnected. When the wireless client computer reconnects, it authenticates by using the newly installed user certificate and obtains authenticated wireless access to the network.

The following illustration shows a wireless configuration.




When you configure IAS so that wireless clients can access the certificate server on the VLAN, it is important to note that the term guest refers to user identity and the term unauthenticated refers to user credentials. The following rules illustrate the differences between the Guest account and remote access policy unauthenticated access:

**User identity rules**

| Condition | IAS action |
|---|---|
| If the wireless client has a UserID | IAS maps to the corresponding Windows account. |
| If the wireless client does not have a UserID or uses a null UserID | IAS attempts to default to the Guest account (in the domain in which the IAS server is a member). |

**Credential rules**

| Condition | IAS action |
|---|---|
| If the wireless client provides credentials | The authentication method in use specifies whether to accept the connection attempt. EAP-MD5 guest authentication is standard user authentication with a UserID and password (which can be null). However, secure channel (Schannel) does not support authentication to a Windows account that is based on a NULL certificate. |
| If the wireless client does not provide credentials | IAS checks the remote access policy to determine whether unauthenticated access is allowed. EAP-TLS supports one-way authorization, such as unauthenticated access where the wireless client does not send credentials. However, EAP-MD5 does not support one-way authorization. |

**Notes**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory domains, the certificate infrastructure, or the wireless access points. For more information about how to deploy these components, see the appropriate Help topics.

- This topic describes the bootstrapping of wireless clients by using a VLAN that contains a certificate server. You can also bootstrap wireless clients by using an unauthenticated Ethernet connection or a certificate that has been extracted to a floppy disk.

To configure IAS for this example, complete the following steps:

- Configure Active Directory for user accounts and groups.

- Configure the primary IAS server on a domain controller.

- Configure the secondary IAS server on a different domain controller.

- Configure RADIUS accounting and authentication on wireless access points.

**Configuring user accounts and groups**

To configure user accounts and groups, do the following:

1. Ensure that all users that are making wireless connections have a corresponding user account.

2. Manage your wireless access by user by setting the remote access permission on user accounts to **Allow access** or **Deny access**. To manage your wireless access by group, set the remote access permission on user accounts to **Control access through Remote Access Policy**. For more information, see Configure remote access permission for a user.

3. Organize your wireless access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For example, create a universal group named WirelessUsers that contains global groups of wireless user accounts. For more information, see Group scope.

4. Configure the Guest account to allow guest access for new wireless clients. For more information, see Guest authentication. Enable reversibly encrypted password storage on the Guest account. For more information, see User and computer accounts.

5. Create a group named Guests, and add the Guest account as a member. For more information, see Create a new group.

6. Configure the domain in which the IAS server computers will be members for auto-enrollment of computer certificates. For more information, see Configure automatic certificate allocation from an enterprise CA.

**Configuring the primary IAS server on a domain controller**

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, see Authentication across forests.

4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.

5. If needed, configure additional UDP ports for authentication and accounting messages that are sent by RADIUS clients. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

6. Add the wireless access points as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets. Enable the use of the Message Authenticator attribute only when it is supported by the wireless access point.

7. Use the New Remote Access Policy Wizard to create a common wireless policy with the following settings:

   - Policy name: Wireless access
   - Access Method: Wireless access
   - User or Group: Select **Group**, and then specify the WirelessUsers group (example).
   - Authentication methods: Select **Smart Card or other Certificate**. If you have multiple computer certificates, click **Configure**, and then select the appropriate computer certificate.

- Policy Encryption Level: Select the **Strongest encryption** check box, and then clear all other check boxes.

8. Use the New Remote Access Policy Wizard to create a custom wireless policy with the following settings:

    - Policy name: New wireless access
    - Conditions:
    - NAS-Port-Type matches Wireless-Other or Wireless-IEEE 802.11
    - Windows-Groups matches Guests
    - Permission: Grant remote access permission
    - Profile settings, **Dial-in Constraints** tab:
    - Select the **Minutes client can be connected** check box, and then type **10**.
    - Profile settings, **Advanced** tab:
    - Add the **Tunnel-Type** attribute with the value of **Virtual LANs (VLAN)**.
    - Add the **Tunnel-Pvt-Group-ID** attribute with the value of the VLAN ID of the VLAN that contains the certificate server for new wireless clients.

For additional examples of remote access policies, see Remote Access Policies Examples.

1. Delete the default remote access policies. For more information, see Delete a remote access policy.

**Configuring the secondary IAS server on a different domain controller**

To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, see Authentication across forests.

4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS authentication and accounting on wireless access points**

Configure your third-party wireless access point as a RADIUS client with two RADIUS servers (the primary and secondary IAS servers). For more information, see the documentation for the wireless access point.

# Dial-up and VPN remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up and VPN remote access**

This topic describes how IAS is used to support authentication, authorization, and accounting for dial-up and virtual private network (VPN) connections to an organization. This topic describes a typical configuration for an organization that uses:

- Two IAS servers.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication. If only one RADIUS server is configured and it becomes unavailable, dial-up and VPN access clients cannot connect. By using two IAS servers and configuring all dial-up and VPN servers (RADIUS clients) for both the primary and secondary IAS servers, RADIUS clients can detect when the primary RADIUS server is unavailable and automatically fail over to the secondary IAS server.

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-up properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints. To optimize IAS authentication and authorization response times and minimize network traffic, IAS is installed on domain controllers.

- A smart card certificate infrastructure.

  The Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication protocol is used with smart cards to authenticate employee connections and provide the highest level of security for authentication.

- Custom remote access policies.

  Remote access policies are configured to specify, based on group membership, the different requirements for various types of connections and users.

- Multiple dial-up and VPN servers.

  Dial-up and VPN servers consist of computers running different versions of Windows and the Routing and Remote Access service, and third-party network access server (NAS) devices.

The following illustration shows the common dial-up and VPN remote access configuration.




**Note**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory, the smart card certificate infrastructure, or the dial-up and VPN servers. For more information about how to deploy these components, see the appropriate Help topics.

To configure IAS for this example, complete the following steps:

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.
- Configure RADIUS accounting and authentication on dial-up and VPN servers.

**Configuring user accounts and groups**

To configure user accounts and groups, do the following:

1. Ensure that all users that are making remote access connections have a corresponding user account. This includes employees, contractors, vendors, and business partners.
2. Set the remote access permission on user accounts to **Allow access** or **Deny access** to manage network access by user. Or, to manage network access by group, set the remote access permission on user accounts to **Control access through Remote Access Policy**. For more information, see Configure remote access permission for a user.
3. Organize remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For more information, see Group scope.
4. If you are using the Challenge-Handshake Authentication Protocol (CHAP), enable support for reversibly encrypted passwords in the appropriate domains. For more information, see Enable reversibly encrypted passwords in a domain.
5. If you are using certificate-based authentication, configure the domain in which IAS server computers will be members for the auto-enrollment of computer certificates. For more information, see Configure automatic certificate allocation from an enterprise CA.

**Configuring the primary IAS server on a domain controller**

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see Install IAS.
2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.
3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that these domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, see Authentication across forests.
4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.
5. If needed, configure additional UDP ports for authentication and accounting messages that are sent by RADIUS clients. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.
6. Add the dial-up and VPN servers as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets. Enable the use of the Message Authenticator attribute, but only when it is also supported by the RADIUS client.
7. Create remote access policies that reflect your remote access usage scenarios.

For example, to configure a remote access policy to permit members of the Contractors group dial-up connections from 8:00 AM to 5:00 PM, Monday through Friday, create a new custom remote access policy with the following settings:

- Policy name: Contractor connections
- Conditions: Windows-Groups matches Contractors
- Permission: Grant remote access permission
- Profile settings, **Dial-in constraints** tab: **Allow access only on these days and at these times** is set to 8:00 AM to 5:00 PM, Monday through Friday; **llow access only through these media** is set to Async (Modem).

All other settings are at their default values.

8. As another example, to configure a remote access policy that requires VPN connections for members of the Employees group to use smart card authentication and 128-bit encryption, use the New Remote Access Policy Wizard to create a common VPN policy with the following settings:

   - Policy name: VPN connections
   - Access Method: VPN access
   - User or Group: Select **Group**, and then specify the Employees group (example).
   - Authentication methods: Select **Smart Card or other Certificate**, and then clear all other check boxes.
   - Policy Encryption Level: Select the **Strongest encryption** check box, and then clear all other check boxes.

     For additional examples of remote access policies, see Remote Access Policies Examples.

     If you have created new remote access policies, either delete the default remote access policy named **Allow access if dial-up permission is enabled**, or move it so that it is the last policy to be evaluated. For more information, see Delete a remote access policy and Change the policy evaluation order.

**Configuring the secondary IAS server on a different domain controller**

To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, see Authentication across forests.

   For additional examples of remote access policies, see Remote Access Policies Examples.

If you have created new remote access policies, either delete the default remote access policy named **Allow access if dial-up permission is enabled**, or move it so that it is the last policy to be evaluated. For more information, see Delete a remote access policy and Change the policy evaluation order.

4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS authentication and accounting on dial-up and VPN servers**

To configure each dial-up or VPN server to use the primary and secondary IAS servers for authentication, authorization, and accounting of remote access connections, do the following:

1. If the dial-up or VPN server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service; configure the primary and secondary IAS servers as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting.

2. If the dial-up or VPN server is a computer running Windows NT Server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT Server 4.0 online Help for information about how to configure the primary and secondary IAS servers as RADIUS servers for RADIUS authentication.

3. If the dial-up or VPN server is a third-party network access server (NAS), see the documentation for the NAS to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS servers).

# Extranet access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Extranet access**

This topic describes how IAS and the Routing and Remote Access service are used to support an extranet. An extranet is an extension of an organization intranet that is used to facilitate communication with business partners. It provides business partners with limited access to resources on an organization's intranet through a virtual private network (VPN) connection over the Internet. This enhances the speed and efficiency of the business relationship. The organization has multiple business partners, each of which need access to only a specific subnet or set of addresses within the organization's network. Network access control is defined at the IP level using IP packet filters for specific connection categories. The filters that are applied to VPN connections can be specified by using remote access policies.

This topic describes a typical configuration for an organization extranet that uses:

- Two IAS servers.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication. If only one RADIUS server is configured and it becomes unavailable, VPN users cannot connect. By using two IAS servers and configuring all VPN servers (RADIUS clients) for both the primary and secondary IAS servers, the RADIUS clients can detect when the primary RADIUS server is unavailable and automatically fail over to the secondary IAS server.

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints. To optimize IAS authentication and authorization response times and minimize network traffic, IAS is installed on domain controllers.

- Custom remote access policies.

  Remote access policies are configured to specify, based on group membership, how different business partners can access different extranet resources.

- VPN servers.

  VPN servers consist of computers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000 and the Routing and Remote Access service.

The following illustration shows the configuration for VPN-based remote access to the extranet by business partners.




**Note**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory, the VPN servers, or the resources on the extranet. For more information about how to deploy these components, see the appropriate Help topics.

To configure IAS for this example, complete the following steps:

- Configure Active Directory for user accounts, groups, and certificate auto-enrollment.
- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.
- Configure RADIUS authentication and accounting on VPN servers.

**Configuring user accounts, groups, and certificate auto-enrollment**
To configure user accounts, groups, and certificate auto-enrollment, do the following:

1. Ensure that all business partners that are making remote access connections have a corresponding user account. Configure the user principal names (UPNs) for business partner accounts by using the organization name of the business partner. For example, for a user account named Someone for the business partner Microsoft, configure the UPN as someone@example.microsoft.com.
2. Manage your business partner's extranet access by user by setting the remote access permission on user accounts to **Allow access** or **Deny access**. To manage your business partner's extranet by group, set the remote access permission on user accounts to **Control access through Remote Access Policy**. For more information, see Configure remote access permission for a user.
3. Organize your business partner user accounts into the appropriate universal and nested global groups in order to take advantage of group-based remote access policies. Create global groups for each business partner and then add the appropriate user accounts. Next, create universal groups for business partner categories that contain the global groups as members. For example, for the business partner Microsoft, create a global group named Microsoft and add all Microsoft user accounts to the group. Then, because Microsoft is a supplier business partner, create a universal group named Suppliers and add the Microsoft global group as a member. For more information, see Group scope.

**Configuring the primary IAS server on a domain controller**
To configure the primary IAS server on a domain controller, do the following:

**On the domain controller, install IAS as an optional networking component.**
For more information, see Install IAS.

**Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain.**
For more information, see Enable the IAS server to read user accounts in Active Directory.

If the IAS server is authenticating connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

**Enable file logging of accounting, authentication, and periodic status events.**
For more information, see Configure log file properties.

**Add the VPN servers as RADIUS clients of the IAS server.**
For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

**Create remote access policies that reflect your required business partner extranet usage.**
For example, to create a remote access policy that restricts all VPN remote access traffic for members of the Suppliers group to the subnet of 192.168.47.0/24, use the New Remote Access Policy Wizard to create a new custom remote access policy with the following settings:

- Policy name: Supplier business partner connections
- Conditions: Windows-Groups matches Suppliers
- Permission: Grant remote access permission

- Profile settings, **Dial-in Constraints** tab: **Allow access only through these media** and the **Virtual (VPN)** media type are selected.
- Profile settings, **IP** tab, Input packet filter:
    - Deny all traffic except those listed below
    - Destination network, IP address: 192.168.47.0
    - Destination network, Subnet mask: 255.255.255.0
    - Protocol: Any
- Profile settings, **IP** tab, Output packet filter:
    - Deny all traffic except those listed below
    - Source network, IP address: 192.168.47.0
    - Source network, Subnet mask: 255.255.255.0
    - Protocol: Any
- Profile settings, **Authentication** tab: Select the **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** check box, and then clear all other check boxes.
- Profile settings, **Encryption** tab: Select the **Strongest** check box, and then clear all other check boxes.

For more information, see Add a remote access policy.

As another example, to create a remote access policy that restricts all VPN remote access traffic for members of the business partners Investors group to the IP address of 10.59.192.17, create a new remote access policy with the following settings:

- Policy name: Investor business partner connections
- Conditions: Windows-Groups matches Investors
- Permission: Grant remote access permission
- Profile settings, **Dial-in Constraints** tab: **Allow access only through these media** and the **Virtual (VPN)** media type are selected.
- Profile settings, **IP** tab, Input packet filter:
    - Deny all traffic except those listed below
    - Destination network, IP address: 10.59.192.17
    - Destination network, Subnet mask: 255.255.255.255
    - Protocol: Any
- Profile settings, **IP** tab, Output packet filter:
    - Deny all traffic except those listed below

    - Source network, IP address: 10.59.192.17
    - Source network, Subnet mask: 255.255.255.255
    - Protocol: Any

- Profile settings, **Authentication** tab: Select the **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)** check box, and then clear all other check boxes.
- Profile settings, **Encryption** tab: Select the **Strongest** check box, and then clear all other check boxes.

For more information, see Add a remote access policy.

For additional examples of remote access policies, see Remote Access Policies Examples.

Delete the default remote access policies, or move them so that they are the last policies to be evaluated. For more information, see Delete a remote access policy and Change the policy evaluation order.

**Ensure that a computer certificate that can be validated by computers using smart cards is installed on the IAS server computer.**
For more information, see Computer certificates for certificate-based authentication.

**Configuring the secondary IAS server on a different domain controller**
To configure the secondary IAS server on a different domain controller, do the following:

**On the other domain controller, install IAS as an optional networking component.**
For more information, see Install IAS.

**Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain.**
For more information, see Enable the IAS server to read user accounts in Active Directory.

If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

**Ensure that a computer certificate that can be validated by computers using smart cards is installed on the IAS server computer.**
For more information, see Computer certificates for certificate-based authentication.

**Copy the configuration of the primary IAS server to the secondary IAS server.**
For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS authentication and accounting on the VPN servers**
To configure each VPN server to use the primary and secondary IAS servers for authentication, authorization, and accounting of remote access connections, do the following:

1. If the VPN server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, configure the primary and secondary IAS servers as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting.

2. If the VPN server is a computer running Windows NT server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT server 4.0 online Help for information about how to configure the primary and secondary IAS servers as RADIUS servers for RADIUS authentication.

# Internet access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet access**

This topic describes how IAS is used to authenticate and authorize Internet connections for the customers of an Internet service provider (ISP). In this example, the ISP has a single data center that supports a large number of customers. Dial-up servers are distributed with multiple Points of Presence (POPs). The ISP offers two service plans: A basic plan for users with dial-up modems and a premium plan that provides support for ISDN connections. Connection authorization is provided to customers based on the subscription plan.

For each basic plan customer, a Connection Manager service profile is used to enable access to the Internet through the ISP's POPs. For each premium plan user, the same Connection Manager service profile is used to enable ISDN access. The service profile contains both dial-up and ISDN POPs. Connection Manager makes the appropriate POPs available based on the type of connection device that the user selects.

This topic describes a configuration for an ISP that uses:

- Two IAS servers.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication. If only one RADIUS server is configured and it becomes unavailable, ISP customers cannot connect to the Internet. By using two IAS servers and configuring all dial-up servers (RADIUS clients) for both the primary and secondary IAS servers, RADIUS clients detect when the primary RADIUS server is unavailable and automatically fail over to the secondary IAS server.

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints. To optimize IAS authentication and authorization response times and minimize network traffic, IAS is installed on domain controllers.

- Custom remote access policies.

  Remote access policies are configured to specify, based on group membership, the different types of connection constraints for the two service plans (basic and premium).

- Multiple dial-up servers.

  Dial-up servers consist of third-party network access server (NAS) devices, and computers running different versions of Windows and the Routing and Remote Access service. The dial-up servers are geographically distributed at various POPs so that ISP customers can access the Internet with a local call.

The following illustration shows the configuration of an ISP providing Internet access.




**Note**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory, the dial-up servers, or the Connection Manager profile. For more information about how to deploy these components, see the

appropriate Help topics.

To configure IAS to support Internet access for an ISP, complete the following steps:

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.
- Configure RADIUS accounting and authentication on dial-up servers.

**Configuring user accounts and groups**

To configure user accounts and groups, do the following:

1. Ensure that all customers that are making Internet connections have a corresponding user account. The user principal names that are created should contain the @ (at) sign, followed by the ISP name (for example, *UserName***@***ISPName*).
2. If you want to manage access by group, ensure that all user accounts are configured with the **Control access through Remote Access Policy** remote access permission. For more information, see Configure remote access permission for a user.
3. Create two universal groups: one for basic plan users (named BasicPlanUsers) and one for premium plan users (named PremiumPlanUsers). Because of the large number of potential users, create global groups and make user accounts members of global groups. Next, add the global groups as members of the universal groups. For more information, see Group scope.
4. If some customers are using the Challenge-Handshake Authentication Protocol (CHAP), you must enable support for reversibly encrypted passwords in the appropriate domains. For more information, see Enable reversibly encrypted passwords in a domain.

**Configuring the primary IAS server on a domain controller**

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see Install IAS.
2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.
3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, seeEnable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.
4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.
5. If needed, configure additional UDP ports for authentication and accounting messages that are sent by RADIUS clients. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.
6. Add the dial-up servers as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets. Enable the use of the Message Authenticator attribute only when it is supported by the RADIUS client.
7. Create remote access policies for both basic and premium plan customers.

Create two remote access policies. One policy is used to authorize basic plan connections (by using the BasicPlanUsers group). The other policy is used to authorize premium plan connections (by using the PremiumPlanUsers group).

8. Authorize connections for basic plan users by using the New Remote Access Policy Wizard to create a common VPN policy with the following settings:

   - Policy name: Basic plan connections
   - Conditions: Windows-Groups matches BasicPlanUsers
   - Permission: Grant remote access permission
   - Profile settings, **Dial-in Constraints** tab: **Allow access only through these media** and the **Async (Modem)** media type are selected.
   - Profile settings, **Authentication** tab: The **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)**, **Microsoft Encrypted Authentication (MS-CHAP)**, and **Encrypted Authentication (CHAP)** protocols are selected.

     For more information, see Add a remote access policy.

9. Authorize connections for premium plan users by creating a new remote access policy with the following settings:

   - Policy name: Premium plan connections
   - Conditions: Windows-Groups matches PremiumPlanUsers
   - Permission: Grant remote access permission
   - Profile settings, **Dial-in Constraints** tab: **Allow access only through these media** is selected, and then the **ISDN Async V.110**, **ISDN Async V.120**, and **ISDN Async** media type are selected.
   - Profile settings, **Authentication** tab: The **Microsoft Encrypted Authentication version 2 (MS-CHAP v2)**, **Microsoft Encrypted Authentication (MS-CHAP)**, and **Encrypted Authentication (CHAP)** protocols are selected.

10. Because all of the ISP customers are using either the basic or premium plan, delete the default remote access policies so that no other types of connections are authorized. For more information, see Delete a remote access policy.

    For additional examples of remote access policies, see Remote Access Policies Examples.

**Configuring the secondary IAS server on a different domain controller**

To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the secondary IAS server is authenticating connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, seeEnable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS accounting and authentication on dial-up servers**

To configure each dial-up server to use the primary and secondary IAS servers for authentication, authorization, and accounting of remote access connections, do the following:

1. If the dial-up server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, configure the primary and secondary IAS servers as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting

2. If the dial-up server is a computer running Windows NT server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT server 4.0 online Help for information about how to configure the primary and secondary IAS servers as RADIUS servers for RADIUS authentication.

3. If the dial-up server is a third-party network access server (NAS), see the documentation for the NAS to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS servers).

# Outsourced VPN remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Outsourced VPN remote access**

This topic describes how IAS is used by both an outsourced service provider and an organization to outsource VPN-based remote access. The outsourced service provider offers worldwide Points of Presence (POPs), which organization employees can call. After the call is made, the employee's computer makes a Point-to-Point Tunneling Protocol (PPTP) VPN connection to one of the organization's VPN servers.

The organization in this example configuration has a large central location with remote users. Each user requires secure access to the organization network. The organization has determined that it is more cost-effective to provide VPN remote access through outsourcing than to implement and maintain its own dial-up remote access infrastructure.

The encryption and authentication requirements for the Internet connection are less stringent than those required for the VPN connection. For example, the Internet connection uses CHAP without encryption, while the VPN connection uses smart cards and Extensible Authentication Protocol-Transport Level Security (EAP-TLS) with 128-bit encryption.

This topic describes a configuration for an organization that uses:

- Two IAS servers.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication, authorization, and accounting. If only one RADIUS server is configured and it becomes unavailable, dial-up and VPN users cannot connect. By using two IAS servers, the IAS proxies in the ISP network can detect when the primary RADIUS server becomes unavailable and automatically fail over to the secondary IAS server. The IAS servers are placed in the perimeter network.

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints.

- Smart card certificate infrastructure.

  The Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication protocol is used with smart cards to authenticate VPN remote access users and provide the highest level of security for VPN authentication.

- Custom remote access policies.

  Remote access policies are configured to specify, based on group membership, the different types of connection constraints for users.

- Connection Manager.

  Connection Manager is used to create a profile with all of the phone numbers of the ISP's POPs. This automates both the dialing of the ISP and the creation of the VPN connection. Remote access users select the location from which they are dialing, and then select either a local or other appropriate ISP phone number (POP) from the phone book in the Connection Manager service profile. Remote access users connect to the ISP's dial-up access server through the Challenge Handshake Password Authentication (CHAP). The realm name portion of the user name is used by the ISP's proxies to forward authentication requests to an IAS proxy in the organization’s perimeter network. After the ISP connection is made, the Connection Manager profile automatically initiates a VPN connection to one of the organization's VPN servers.

- VPN servers.

  VPN servers consist of computers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000 and the Routing and Remote Access service.

This topic also describes a configuration for an ISP that uses:

- Two IAS proxies in the ISP's network.

  The ISP uses RADIUS proxies on its network to forward RADIUS request messages between the ISP dial-up servers and the RADIUS servers of multiple customers. Two IAS proxies are used to provide fault tolerance for RADIUS authentication.

- Dial-up access servers.

  Dial-up servers consist of computers running:

  - Windows Server 2003, Standard Edition
  - Windows Server 2003, Enterprise Edition
  - Windows Server 2003, Datacenter Edition
  - Windows 2000 and the Routing and Remote Access service
  - Third-party network access server (NAS) devices

The following illustration shows the outsourced VPN configuration.




**Note**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory domains, the smart card certificate infrastructure, Connection Manager, or the VPN servers. For more information about how to deploy these components, see the appropriate Help topics.

To configure IAS for this example, complete the following steps:

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server in the perimeter network.
- Configure the secondary IAS server in the perimeter network.
- Configure RADIUS authentication and accounting on VPN servers.
- Configure the Internet firewall to support RADIUS traffic.
- Configure the primary IAS proxy at the service provider.
- Configure the secondary IAS proxy at the service provider.

- Configure RADIUS accounting and authentication on the dial-up access servers at the service provider.

**Configuring user accounts and groups**

To configure user accounts and groups, do the following:

1. Ensure that all users who are making remote access connections have a corresponding user account.

2. If you want to manage access by group, ensure that all user accounts are configured with the **Control access through Remote Access Policy** remote access permission. For more information, see Configure remote access permission for a user.

3. Create two universal groups: one for employees (named Employees) and one for contractors (named Contractors). Because of the large number of potential users, create global groups and add the user accounts as members to the appropriate group. Next, add the global groups as members to the universal groups. For more information, see Group scope.

4. Because the ISP requires the use of Challenge-Handshake Authentication Protocol (CHAP) authentication for the dial-up connection, you must enable support for reversibly encrypted passwords for the appropriate domains. For more information, see Enable reversibly encrypted passwords in a domain.

**Configuring the primary IAS server in the perimeter network**

To configure the primary IAS server in the perimeter network, do the following:

1. On a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the perimeter network; install IAS as an optional networking component. For more information, see Install IAS. The computer on which IAS is installed is not required to be dedicated to forwarding RADIUS messages. You can install IAS on a Web server, file server, or DNS server.

2. Configure the IAS server computer to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the IAS server is authenticating connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the IAS server is authenticating connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, see Authentication across forests.

4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.

5. If needed, configure additional UDP ports for RADIUS messages that are sent by the ISP's RADIUS proxies. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

6. Add the service provider's RADIUS proxies as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

7. Create remote access policies that allow VPN remote access for employees and deny it for contractors.

   Use the New Remote Access Policy Wizard to create a common remote access policy that both allows VPN remote access for members of the Employees group and requires the use of EAP-TLS and 128-bit encryption with the following settings:

- Policy name: VPN connections for employees
- Access Method: VPN access
- User or Group: Select **Group**, and then specify the Employees group (example).
- Authentication methods: Select the **Extensible Authentication Protocol** check box, ensure that the **Smart Card or other Certificate** EAP type is selected, and then clear all other check boxes.
- Policy Encryption Level: Select the **Strongest encryption** check box, and then clear all other check boxes.

  Create a custom remote access policy that denies VPN remote access for members of the Contractors group with the following settings:

- Policy name: **Deny VPN connections for contractors**
- Conditions:
- Windows-Groups matches Contractors
- NAS-port-type matches Virtual (VPN)
- Permission: Deny remote access permission

  For additional examples of remote access policies, see Remote Access Policies Examples.

8. Delete the default remote access policies. For more information, see Delete a remote access policy.

**Configuring the secondary IAS server in the perimeter network**

To configure the secondary IAS server on another computer in the perimeter network, do the following:

1. On another computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows server 2003, Datacenter Edition in the perimeter network; install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the secondary IAS server is authenticating connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the secondary IAS server is authenticating connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, see Authentication across forests.

4. Copy the configuration from the primary IAS proxy to the secondary IAS proxy in the perimeter network. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS authentication and accounting on the VPN servers**

To configure each VPN server to use the primary and secondary IAS servers for authentication, authorization, and accounting of remote access connections, do the following:

1. If the VPN server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise

Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, configure the primary and secondary IAS servers in the perimeter network as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting.

2. If the VPN server is a computer running Windows NT Server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT Server 4.0 online Help for information about how to configure the primary and secondary IAS servers as RADIUS servers for RADIUS authentication.

3. If the VPN server is a third-party network access server (NAS), see the documentation for the NAS to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS servers).

**Configuring the Internet firewall to support RADIUS traffic**

1. Configure the Internet firewall to allow RADIUS traffic between the IAS servers in the perimeter network and the IAS proxies in the ISP's network.

   For more information, see IAS and firewalls

**Configuring the primary IAS proxy at the service provider**

1. On a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the ISP's network; install IAS as an optional networking component. For more information, see Install IAS. The computer on which IAS is installed is not required to be dedicated to forwarding RADIUS messages. You can install IAS on a Web server, file server, or DNS server.

2. If needed, configure additional UDP ports for authentication and accounting messages that are sent by the ISP's dial-up access servers. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and 1813 and 1646 for accounting.

3. Add the ISP's dial-up access servers as RADIUS clients of the IAS proxy. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

4. Create a connection request policy that forwards RADIUS request messages on the basis of the ISP customer’s realm name.

   Use the New Connection Request Policy Wizard to create a connection request policy with the condition that the user name matches the realm name for the customer's organization and forwards RADIUS request messages to a remote RADIUS server for authentication. In the New Connection Request Policy Wizard, use the New Remote RADIUS Server Group Wizard to create a remote RADIUS server group with members that include the two IAS servers in the customer's perimeter network.

   For more information, see Add a connection request policy.

5. Delete the default connection request policy named **Use Windows authentication for all users**. For more information, see Delete a connection request policy.

**Configuring the secondary IAS proxy at the service provider**

To configure the secondary IAS proxy on another computer in the perimeter network, do the following:

1. On another computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the ISP's network; install IAS as an optional networking component. For more information, see Install IAS.

2. Copy the configuration of the primary IAS proxy to the secondary IAS proxy in the ISP's network. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS accounting and authentication on the dial-up access servers at the service provider**

To configure each dial-up or VPN server to use the primary and secondary IAS servers for authentication, authorization, and accounting of remote access connections, do the following:

1. If the dial-up or VPN server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, configure the primary and secondary IAS proxies in the ISP's network as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting.

2. If the dial-up or VPN server is a computer running Windows NT Server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT Server 4.0 online Help for information about how to configure the primary and secondary IAS proxies in the ISP's network as RADIUS servers for RADIUS authentication.

3. If the dial-up or VPN server is a third-party network access server (NAS), see the documentation for the NAS to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS proxies) in the ISP's network.

# Outsourced wireless access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Outsourced wireless access**

This topic describes how IAS is used by both an outsourced wireless provider and an organization to outsource wireless Internet access. The wireless service provider offers remote wireless access points in locations such as airports. Employees of an organization can use these points to access the Internet. The wireless service provider forwards the access request of a client to an IAS server in the organization. The IAS server provides authentication and authorization. The wireless service provider determines the employee's organization from the realm name portion of the employee's user name during the authentication process. This enables the service provider to bill the organization for the employee's use of the wireless connection.

This topic describes a configuration for an organization that uses:

- Two IAS servers.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication, authorization, and accounting. If only one RADIUS server is configured and it becomes unavailable, outsourced wireless clients cannot connect. By using two IAS servers, the IAS proxies in the wireless service provider's network can detect when the primary RADIUS server is unavailable and automatically fail over to the secondary IAS server. The IAS servers are placed in the perimeter network.

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints.

- A certificate infrastructure.

  The Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication protocol is used to authenticate wireless clients with locally installed user certificates.

- A wireless remote access policy.

  A common remote access policy is configured so that authenticated wireless users can access the Internet.

This topic also describes a configuration for a wireless service provider that uses:

- Two IAS proxies in the wireless service provider's network.

  The wireless service provider uses RADIUS proxies in its network to forward RADIUS request messages between the wireless access points and the RADIUS servers. Two IAS proxies are used to provide fault tolerance for RADIUS authentication.

- Multiple wireless access points.

  Third-party wireless access points provide wireless access in different airports.

The following illustration shows the outsourced wireless configuration.




**Note**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory domains, the certificate infrastructure, or the wireless access points. For more information about how to deploy these components, see the appropriate Help topics.

To configure IAS for this example, complete the following steps:

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server in the perimeter network.
- Configure the secondary IAS server in the perimeter network.
- Configure the Internet firewall to support RADIUS traffic.
- Configure the primary IAS proxy at the service provider.
- Configure the secondary IAS proxy at the service provider.
- Configure RADIUS accounting and authentication on the wireless access points at the service provider.

**Configuring user accounts and groups**

To configure user accounts and groups, do the following:

1. Ensure that all users that are making remote access connections have a corresponding user account.
2. If you want to manage access by group, ensure that all user accounts are configured with the **Control access through Remote Access Policy** remote access permission. For more information, see Configure remote access permission for a user.
3. Organize your wireless access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For example, create a universal group named WirelessUsers that contains global groups of wireless user accounts. For more information, see Group scope.

**Configuring the primary IAS server in the perimeter network**

To configure the primary IAS server in the perimeter network, do the following:

1. On a computer runningWindows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the perimeter network; install IAS as an optional networking component. For more information, see Install IAS. The computer on which IAS is installed is not required to be dedicated to forwarding RADIUS messages. You can install IAS on a Web server, file server, or DNS server.
2. Configure the IAS server computer to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the IAS server is authenticating connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the IAS server is authenticating connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, see Authentication across forests.

4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.

5. If needed, configure additional UDP ports for RADIUS messages that are sent by the wireless service provider's RADIUS proxies. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

6. Add the service provider's RADIUS proxies as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

7. Use the New Remote Access Policy Wizard to create a common wireless policy with the following settings:

   - Policy name: Wireless access
   - Access Method: Wireless access
   - User or Group: Select **Group**, and then specify the WirelessUsers group (example).
   - Authentication methods: Select **Smart Card or other Certificate**. If you have multiple computer certificates, click **Configure**, and then select the appropriate computer certificate.
   - Policy encryption level: Clear the **Basic encryption (MPPE 40-bit)** and **Basic encryption (MPPE 56-bit)** check boxes.

     For additional examples of remote access policies, see Remote Access Policies Examples.

8. Delete the default remote access policies, or ensure that they are the last policies to be evaluated. For more information, see Delete a remote access policy and Change the policy evaluation order.

**Configuring the secondary IAS server in the perimeter network**

To configure the secondary IAS server on another computer in the perimeter network, do the following:

1. On another computer runningWindows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the perimeter network; install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the secondary IAS server is authenticating connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the secondary IAS server is authenticating connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, see

Authentication across forests.

4. Copy the configuration of the primary IAS proxy to the secondary IAS proxy in the perimeter network. For more information, see Copy the IAS configuration to another server.

**Configuring the Internet firewall to support RADIUS traffic**

To configure the Internet firewall to support RADIUS traffic, do the following:

1. Configure the Internet firewall to allow RADIUS traffic between the IAS servers in the perimeter network and the IAS proxies in the wireless service provider's network.

   For more information, see IAS and firewalls.

**Configuring the primary IAS proxy at the service provider**

To configure the primary IAS proxy at the service provider, do the following:

1. On a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the wireless service provider's network; install IAS as an optional networking component. For more information, see Install IAS. The computer on which IAS is installed is not required to be dedicated to forwarding RADIUS messages. You can install IAS on a Web server, file server, or DNS server.

2. If needed, configure additional UDP ports for authentication and accounting messages that are sent by the wireless service provider's access points. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

3. Add the wireless service provider's access points as RADIUS clients of the IAS proxy. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

4. Create a connection request policy that forwards RADIUS request messages based on the realm name of the wireless service provider's customer.

   Use the New Connection Request Policy Wizard to create a connection request policy that forwards connection requests to a remote RADIUS server group for authentication and where the user name matches the realm name for the customer's organization. In the New Connection Request Policy Wizard, use the New Remote RADIUS server Group Wizard to create a remote RADIUS server group with members that include the two IAS servers in the customer's perimeter network.

   For more information, see Add a connection request policy.

5. Delete the default connection request policy named **Use Windows authentication for all users**. For more information, see Delete a connection request policy.

**Configuring the secondary IAS proxy at the service provider**

To configure the secondary IAS proxy on another computer in the perimeter network, do the following:

1. On another computer runningWindows Server 2003, Standard Edition; Windows Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the wireless service provider network; install IAS as an optional networking component. For more information, see Install IAS.

2. Copy the configuration of the primary IAS proxy to the secondary IAS proxy in the wireless service provider's network. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS accounting and authentication on the wireless access points at the service provider**

Configure each wireless access point as a RADIUS client with two RADIUS servers (the primary and secondary IAS servers). For more information, see the documentation for the wireless access point.

# Ethernet switch access with IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Ethernet switch access**

This topic describes how IAS can be used to support authentication, authorization, and accounting for authenticated Ethernet switch connections to an organization. In contrast to user-authenticated connections, such as dial-up and VPN connections, Ethernet connections use computer certificates and computer authentication to validate the computer that is connecting to the Ethernet switch. This topic describes a typical configuration for an organization that uses:

- Two IAS servers.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication. If only one RADIUS server is configured and it becomes unavailable, authenticating switch clients cannot connect. By using two IAS servers and configuring all authenticating Ethernet switches (RADIUS clients) for both the primary and secondary IAS servers, the RADIUS clients can detect when the primary RADIUS server is unavailable and automatically fail over to the secondary IAS server.

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints. To optimize IAS authentication and authorization response times and minimize network traffic, IAS is installed on domain controllers.

- A computer certificate infrastructure.

  The Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication protocol is used with computer certificates to authenticate Ethernet connections and provide the highest level of security for authentication.

- Ethernet switch remote access policies.

  Remote access policies are configured for authenticated switch connections so that guest switch users obtain restricted access to a virtual LAN (VLAN) that contains a certificate server and authenticated switch users obtain unrestricted access to the organization intranet.

- Multiple authenticating switches.

  Multiple third-party Ethernet authenticating switches provide access in different buildings of an organization.

- A certificate server on a separate VLAN for new Ethernet switch clients.

  A certificate server is placed on a separate virtual local area network (VLAN) and used by a new switch client to obtain a computer certificate for subsequent switch access. After 10 minutes of receiving a guest connection, the switch client is disconnected. When the switch client computer reconnects, it authenticates by using the newly installed computer certificate and obtains authenticated switch access to the network.

The following illustration shows an Ethernet switch configuration.




**Notes**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory domains, the certificate infrastructure, or Ethernet switches. For more information about how to deploy these components, see the appropriate Help topics.

- This topic describes the bootstrapping of new Ethernet clients by using a VLAN that contains a certificate server. You can also bootstrap new Ethernet clients by using an unauthenticated Ethernet connection or a certificate that has been extracted to a floppy disk.

To configure IAS for this example, complete the following steps:

- Configure Active Directory for computer accounts and groups.

- Configure the primary IAS server on a domain controller.

- Configure the secondary IAS server on a different domain controller.

- Configure RADIUS authentication and accounting on Ethernet switches.

**Configuring computer accounts and groups**

To configure computer accounts and groups, do the following:

1. Ensure that all computers that are making authenticating switch connections have a corresponding computer account.

2. If you want to manage switch access by group, set the remote access permission on computer accounts to **Control access through Remote Access Policy**. For more information, see Configure remote access permission for a user.

3. Organize switch access computers into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For example, create a universal group named SwitchComputers that contains global groups of switch computer accounts. For more information, see Group scope.

4. Configure the domain in which the IAS server computers are members for the auto-enrollment of computer certificates. For more information, see Configure automatic certificate allocation from an enterprise CA.

5. Configure the Guest account to allow guest access for new Ethernet clients. For more information, see Guest authentication. Enable reversibly encrypted password storage on the Guest account. For more information, see User and computer accounts.

6. Create a group named Guests, and add the Guest account as a member. For more information, see Create a new group.

**Configuring the primary IAS server on a domain controller**

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the IAS server computer (the domain controller) to read the properties of user and computer accounts in the

domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the IAS server authenticates connection attempts for user and computer accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user and computer accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the IAS server authenticates connection attempts for user and computer accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, see Authentication across forests.

4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.

5. If needed, configure additional UDP ports for authentication and accounting messages that are sent by RADIUS clients. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

6. Add the authenticating Ethernet switches as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets. Enable the use of the Message Authenticator attribute only when it is supported by the Ethernet switch.

7. Use the New Remote Access Policy Wizard to create a common Ethernet policy with the following settings:

   - Policy name: Ethernet switch access
   - Access Method: Ethernet
   - User or Group: Select **Group**, and then specify the SwitchComputers group (example).
   - Authentication methods: Select **Smart Card or other Certificate**. If you have multiple computer certificates, click **Configure**, and then select the appropriate computer certificate.

8. If you are using certificate-based authentication, you can use the New Remote Access Policy Wizard to create a custom Ethernet switch policy with the following settings:

   - Policy name: New switch access
   - Conditions: NAS-Port-Type matches Ethernet and Windows-Groups matches Guests
   - Permission: Grant remote access permission
   - Profile settings, **Dial-in Constraints** tab: Select the **Minutes client can be connected** check box, and then type **10**.
   - Profile settings, **Authentication** tab: Select **Extensible Authentication Protocol** and the **MD-5 Challenge** EAP type, and then clear all other check boxes.
   - Profile settings, **Advanced** tab: Add the **Tunnel-Type** attribute with the value of **Virtual LANs (VLAN)**; Add the **Tunnel-Pvt-Group-ID** attribute with the value of the VLAN ID that corresponds to the VLAN that contains the certificate server for new switch clients.

   For additional examples of remote access policies, see Remote Access Policies Examples.

9. Delete the default remote access policies, or ensure that they are the last policies to be evaluated. For more information, see Delete a remote access policy and Change the policy evaluation order.

**Configuring the secondary IAS server on a different domain controller**

To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, see Authentication across forests.

4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS authentication and accounting on Ethernet switches**

To configure your third-party Ethernet switch RADIUS client with two RADIUS servers (the primary and secondary IAS servers), see the Ethernet switch documentation.

# Deploying IAS as a RADIUS Proxy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying IAS as a RADIUS proxy**
This section covers:

- Planning for IAS as a RADIUS Proxy
- IAS as a RADIUS proxy

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Planning for IAS as a RADIUS Proxy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for IAS as a RADIUS proxy**
This section covers:

- IAS as a RADIUS proxy design considerations
- IAS as a RADIUS proxy security considerations

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# IAS as a RADIUS proxy design considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS as a RADIUS proxy design considerations**

Consider the following design issues when deploying IAS as a RADIUS proxy:

- The common uses of IAS as a RADIUS proxy

  For more information about common ways to use IAS as a RADIUS proxy, see IAS as a RADIUS proxy.

- Connection request policy configuration

  The default connection request policy named **Use Windows authentication for all users** is configured for IAS when it is used as a RADIUS server. To create a connection request policy to use IAS as a RADIUS proxy, you must first create a remote RADIUS server group. Next, create a connection request policy that forwards authentication requests to a remote RADIUS server group. Finally, either delete the default connection request policy or move the new connection request policy so that it is evaluated first. For more information, see Introduction to connection request processing. For an example of a connection request policy configured when IAS is being used as a RADIUS proxy, see Remote RADIUS server.

- Realm replacement and attribute manipulation

  To convert realm names and configure RADIUS message forwarding based on the realm name, you must configure attribute manipulation for the User-Name attribute on the appropriate connection request policy. For more information, see Connection request policies and Configure attribute manipulation. For an example of a connection request policy that forwards authentication requests based on a specific realm name, see Remote RADIUS server.

  If you are using the MS-CHAP v2 authentication protocol, you cannot manipulate the User Name attribute if the connection request policy is used to forward the RADIUS message. The only exception occurs when a backslash character (\) is used and the manipulation only affects the information to the left of it. A backslash character is typically used to indicate a domain name (the information to the left of it) and a user account name within the domain (the information to the right of it). In this case, only attribute manipulation rules that modify or replace the domain name are allowed.

- The use of additional RADIUS attributes and vendor-specific attributes

  If you plan to return additional RADIUS attributes and vendor-specific attributes (VSAs) with RADIUS requests, you must add the RADIUS attributes and VSAs to the appropriate connection request policy. For more information, see Vendor-specific attribute overview and Configure advanced attributes.

- Remote RADIUS server group configuration

  To forward RADIUS requests, the connection request policy must be configured to use a remote RADIUS server group (a list of one or more RADIUS servers to which RADIUS messages are forwarded). For more information, see Remote RADIUS server groups and Configure Remote RADIUS Server Groups.

- Copy logging information at the IAS proxy

  The IAS proxy can record all RADIUS accounting information in the local log file. This creates a central location for all authentication and accounting information for all of the access servers of the IAS proxy. For more information, see Connection request policies.

- Authentication and accounting ports

  When you configure a server in a remote RADIUS server group, you can configure custom UDP ports to which RADIUS authentication and accounting messages are sent. The default UDP port for authentication requests is 1812. The default

UDP port for accounting requests is 1813. For more information, see Remote RADIUS server groups and Configure the authentication and accounting settings of a group member.

- Load balancing and failure detection

  When you configure multiple servers in a remote RADIUS server group, you can configure settings that determine how the IAS server balances the load of authentication and accounting requests over the RADIUS servers in the group. You can use additional settings to configure IAS to detect and recover from the failure of a remote RADIUS server group member. For more information, see Configure the load balancing properties of a group member.

- Multiple IAS servers

  To provide fault tolerance for RADIUS-based authentication and accounting, you should always use at least two IAS servers. One IAS server is used as the primary RADIUS proxy and the other is used as a backup. Access servers or other RADIUS proxies are configured for both IAS servers. They switch to using the backup IAS proxy when the primary IAS proxy becomes unavailable. For information about how to synchronize the configuration of multiple IAS servers, see Managing multiple IAS servers.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# IAS as a RADIUS proxy security considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS as a RADIUS proxy security considerations**

Consider the following security issues when deploying IAS as a RADIUS proxy:

- Shared secrets

  Configure strong shared secrets to prevent dictionary attacks and change them frequently. Strong shared secrets are a long (more than 22 characters) sequence of random letters, numbers, and punctuation. For more information, see Shared secrets.

- Firewall configuration

  If your IAS proxy is on a perimeter network, configure your Internet firewall (between your perimeter network and the Internet) to allow RADIUS messages to pass between your IAS proxy and RADIUS clients on the Internet. You might need to configure an additional firewall that is placed between your perimeter network and your intranet, to allow RADIUS traffic to flow between the IAS proxy on the perimeter network and an IAS server on the intranet. For more information, see IAS and firewalls.

- Message Authenticator attribute

  You can use the RADIUS Message Authenticator attribute (also known as a digital signature or the signature attribute) to ensure that RADIUS Access-Request messages for connection requests were sent from a RADIUS client configured with the correct shared secret. The Message Authenticator attribute is always used with EAP--without having to enable it on the IAS server and access server. For the PAP, CHAP, MS-CHAP, and MS-CHAP v2 authentication protocols, you must enable the use of the Message Authenticator attribute on both the IAS server (as part of the configuration of the RADIUS client in Internet Authentication Service) and the RADIUS client (the access server or RADIUS proxy). Ensure that the RADIUS client supports the Message Authenticator attribute before enabling it. For more information, see Edit RADIUS client configuration.

  For information about enabling the RADIUS Message Authenticator attribute for your access server, see the appropriate access server documentation. For the Routing and Remote Access service, the use of the RADIUS Message Authenticator attribute is enabled from the properties of a RADIUS server when you configure RADIUS authentication. For more information, see Use RADIUS authentication.

- Using IPSec filters to lock down IAS proxy servers

  You can configure IPSec filters to allow specific network traffic to pass through network interfaces on RADIUS servers. These filters can be applied to organizational units and stored in Active Directory, or they can be created and applied to individual servers. For more information, see Securing RADIUS traffic with IPSec.

- Password Authentication Protocol (PAP)

  The use of Password Authentication Protocol (PAP) is strongly discouraged when using RADIUS proxies. For more information, see IAS and firewalls.

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by

specifying an IP address range.

# Configuring IAS as a RADIUS Proxy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS as a RADIUS proxy**
This section covers:

- Authentication across forests
- Outsourced dial and a proxy in the perimeter network
- Using a different authentication database
- Using IAS proxy for load balancing
- Mapping network authentication and authorization

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Authentication across forests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication across forests**
This topic describes how IAS can be used as both a RADIUS server and proxy to provide authentication between two forests. This configuration can also be used to provide authentication between both untrusted domains and one-way trusted domains.

### Notes

- When using EAP-TLS with certificates as the authentication method, you need to use one or more RADIUS proxy servers that forward authentication requests to the appropriate forest, even when the forests have a two-way, transitive trust relationship.

- When using authentication methods other than EAP-TLS with certificates, IAS supports authentication without a RADIUS proxy across two forests when they consist of only domain controllers running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition; and contain only Windows Server 2003 domains. For more information, see Accessing resources across forests and Domain and forest functionality.

The following configuration is for an organization that uses:

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints. To optimize IAS authentication and authorization response times and minimize network traffic, IAS is installed on domain controllers.

- Two IAS servers in each forest.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication, authorization, and accounting in each forest. If only one RADIUS server is configured and it becomes unavailable, access clients for that forest cannot connect. By using two IAS servers and configuring the IAS proxies for both primary and secondary IAS servers, the IAS proxies can detect when the primary RADIUS server is unavailable and automatically fail over to the secondary IAS server.

- Remote access policies.

  Remote access policies are configured to specify, based on group membership, the different types of connection constraints for users.

- Two IAS proxies.

  Two IAS proxies are used to provide fault tolerance for RADIUS requests that are sent from access servers.

The following illustration shows the authentication across forests configuration that is using IAS proxies.




**Note**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory domains or access servers. For more information about how to deploy these components, see the appropriate Help topics.

To configure IAS for this example, complete the following steps:

- Configure the Active Directory forests for user accounts and groups.
- Configure the primary IAS server on a domain controller in the first forest.
- Configure the secondary IAS server on a different domain controller in the first forest.
- Configure the primary IAS server on a domain controller in the second forest.
- Configure the secondary IAS server on a different domain controller in the second forest.
- Configure the primary IAS proxy.
- Configure the secondary IAS proxy.
- Configure RADIUS authentication and accounting on access servers.

**Configuring forests for user accounts and groups**

To configure forests for user accounts and groups, do the following:

1. Ensure that all users who are making network access connections have a corresponding user account.
2. Manage your network access by group by setting the remote access permission on user accounts to **Control access through Remote Access Policy**. For more information, see Configure remote access permission for a user.
3. Organize your remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For more information, see Group scope.
4. If you are using the Challenge-Handshake Authentication Protocol (CHAP), enable support for reversibly encrypted passwords for the appropriate domains. For more information, see Enable reversibly encrypted passwords in a domain.

**Configuring the primary IAS server on a domain controller in the first forest**

To configure the primary IAS server on a domain controller in the first forest, do the following:

1. On the domain controller in the first forest, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the IAS server is authenticating connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.

5. Add the IAS proxies as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

6. Create the appropriate remote access policies for access clients in the first forest.

   For examples of remote access policies, see Remote Access Policies Examples.

**Configuring the secondary IAS server on a different domain controller in the first forest**
To configure the secondary IAS server on a different domain controller in the first forest, do the following:

1. On the other domain controller in the first forest, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see Copy the IAS configuration to another server.

**Configuring the primary IAS server on a domain controller in the second forest**
To configure the primary IAS server on a domain controller in the second forest, do the following:

1. On the domain controller in the second forest, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.

5. Add the IAS proxies as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

6. Create the appropriate remote access policies for access clients in the second forest.

   For examples of remote access policies, see Remote Access Policies Examples.

**Configuring the secondary IAS server on a different domain controller in the second forest**

To configure the secondary IAS server on a different domain controller in the second forest, do the following:

1. On the other domain controller in the second forest, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see Copy the IAS configuration to another server.

**Configuring the primary IAS proxy**

To configure the primary IAS proxy, do the following:

1. On a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; install IAS as an optional networking component. For more information, see Install IAS. The computer on which IAS is installed is not required to be dedicated to forwarding RADIUS messages. For example, You can install IAS on a Web server, file server, or DNS server.

2. If needed, configure additional UDP ports for RADIUS messages that are sent by the access servers. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

3. Add the access servers as RADIUS clients of the IAS proxy. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

4. Create a connection request policy that forwards RADIUS request messages (that are based on the realm name of accounts in the first forest) to IAS servers in the first forest.

   Use the New Connection Request Policy Wizard to create a connection request policy that forwards connection requests to a remote RADIUS server group and where the realm name matches the realm name of the user accounts in the first forest. Clear the check box that removes the realm name for authentication. In the New Connection Request Policy Wizard, use the New Remote RADIUS server Group Wizard to create a remote RADIUS server group with members that include the two IAS servers in the first forest.

   For more information, see Add a connection request policy.

5. Create a connection request policy that forwards RADIUS request messages (that are based on the realm name of accounts in the second forest) to IAS servers in the second forest.

   Use the New Connection Request Policy Wizard to create a connection request policy that forwards connection requests to a remote RADIUS server group and where the realm name matches the realm name of the user accounts in the second forest. Clear the check box that removes the realm name for authentication. In the New Connection Request Policy Wizard, use the New Remote RADIUS server Group Wizard to create a remote RADIUS server group with members that include the two IAS servers in the second forest.

For more information, see Add a connection request policy.

6. Delete the default connection request policy named **Use Windows authentication for all users**. For more information, see Delete a connection request policy.

**Configuring the secondary IAS proxy**

To configure the secondary IAS proxy on another computer, do the following:

1. On another computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; install IAS as an optional networking component. For more information, see Install IAS.

2. Copy the configuration of the primary IAS proxy to the secondary IAS proxy. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS authentication and accounting on the access servers**

To configure each access server to use the primary and secondary IAS proxies for the authentication, authorization, and accounting of network connections, do the following:

1. If the dial-up or VPN server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, configure the primary and secondary IAS proxies as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting.

2. If the dial-up or VPN server is a computer running Windows NT Server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT Server 4.0 online Help for information about how to configure the primary and secondary IAS proxies as RADIUS servers for RADIUS authentication.

3. If the dial-up server, VPN server, wireless access point, or authenticating switch is a third-party access server, see the documentation for the access server to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS proxies).

For more information, see Forest trusts, Accessing resources across forests, Upgrading from a Windows NT domain, and Best practices for Active Directory Domains and Trusts.

### Note

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Outsourced dial and a proxy in the perimeter network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Outsourced dial and a proxy in the perimeter network**

This topic describes how IAS can be used as a RADIUS proxy to forward messages between the RADIUS proxies of an outsourced dial service provider and RADIUS servers within an organization intranet. The outsourced service provider offers worldwide Points of Presence (POPs), which the employees of an organization can call. After the call is completed, the employee's computer is connected to the organization's intranet.

This topic describes a configuration for an organization that uses:

- Two IAS servers.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication, authorization, and accounting. If only one RADIUS server is configured and it becomes unavailable, dial-in and VPN users cannot connect. By using two IAS servers and configuring both primary and secondary IAS proxies in the perimeter network for both the primary and secondary IAS servers, the IAS proxies can detect when the primary RADIUS server is unavailable and automatically fail over to the secondary IAS server.

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints. To optimize IAS authentication and authorization response times and minimize network traffic, IAS is installed on domain controllers.

- Custom remote access policies. Remote access policies are configured to specify, based on group membership, the different types of connection constraints for users.

- Two IAS proxies in the perimeter network.

  To simplify the configuration of the intranet firewall, IAS proxies are used in the perimeter network instead of IAS servers. When IAS servers are used in the perimeter network, you must either configure each IAS server computer with two network adapters (one connected to the perimeter network and one connected to the intranet), or configure the intranet firewall to allow Active Directory traffic between the IAS servers on the perimeter network and all of the domain controllers in the intranet. By replacing the IAS server computers with IAS proxies, the IAS proxy computers do not require two network adapters. Additionally, the intranet firewall is only required to be configured to allow RADIUS traffic between the IAS proxies in the perimeter network and the IAS servers in the intranet.

  Two IAS proxies are used in the perimeter network to provide fault tolerance for RADIUS requests that are sent from the service provider's proxies.

- Connection Manager.

  To automate service provider dialing, Connection Manager is used to create a profile with all of the phone numbers of the service provider POPs. Remote access users select the location from which they are dialing, and then select the appropriate phone number (POP) from the phone book in the Connection Manager service profile for the service provider. The remote access user connects to the service provider's dial-in access server through Challenge Handshake Password Authentication (CHAP). The realm name portion of the user name is used by the service provider's proxies to forward the authentication requests to an IAS proxy in the perimeter network of the organization.

This topic also describes a configuration for a service provider that uses:

- Two IAS proxies in the service provider's network.

  The service provider uses RADIUS proxies in its network to forward RADIUS request messages between the dial-in

servers of the service provider and the RADIUS servers or proxies of multiple customers across the Internet. Two IAS proxies are used to provide fault tolerance for RADIUS authentication.

- Dial-up access servers.

  Dial-up servers consist of computers running:

  - Windows Server 2003, Standard Edition
  - Windows Server 2003, Enterprise Edition
  - Windows Server 2003, Datacenter Edition
  - Windows 2000 and the Routing and Remote Access service
  - Third-party network access server (NAS) devices

The following illustration shows the outsourced dial configuration using IAS proxies in the organization's intranet.




**Note**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory domains or Connection Manager. For more information about how to deploy these components, see the appropriate Help topics.

To configure IAS for this example, complete the following steps:

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.
- Configure RADIUS authentication and accounting on VPN servers.
- Configure the primary IAS proxy in the perimeter network.
- Configure the secondary IAS proxy in the perimeter network.
- Configure the intranet and Internet firewalls to support RADIUS traffic.

- Configure the primary IAS proxy at the service provider.
- Configure the secondary IAS proxy at the service provider.
- Configure RADIUS accounting and authentication on the dial-up access servers at the service provider.

**Configuring user accounts and groups**

To configure user accounts and groups, do the following:

1. Ensure that all users who are making remote access connections have a corresponding user account.
2. Manage your network access by group by setting the remote access permission user accounts to **Control access through Remote Access Policy**. For more information, see Configure remote access permission for a user.
3. Organize your remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For more information, see Group scope.
4. Because the outsourced service provider requires the use of Challenge-Handshake Authentication Protocol (CHAP) authentication for the dial-in connection, you must enable support for reversibly encrypted passwords for the appropriate domains. For more information, see Enable reversibly encrypted passwords in a domain.

**Configuring the primary IAS server on a domain controller**

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see Install IAS.
2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.
3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the IAS server is authenticating connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, see Authentication across forests.
4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.
5. Add the IAS proxies in the perimeter network as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.
6. Create remote access policies that accurately reflect your remote access usage.

   For example, to configure a custom remote access policy to permit members of the Contractors group to connect from 8:00 AM to 5:00 PM, Monday through Friday, use the New Remote Access Policy Wizard to create a new custom remote access policy with the following settings:

   - Policy name: **Contractor connections**
   - Conditions: Windows-Groups matches Contractors
   - Permission: **Grant remote access permission**
   - Profile settings, **Dial-in Constraints** tab: **Allow access only on these days and at these times** is set to 8:00 AM

to 5:00 PM, Monday through Friday.

To configure a remote access policy that requires members of the Executives group to use EAP-TLS authentication and 128-bit encryption, use the New Remote Access Policy Wizard to create a common remote access policy with the following settings:

- Policy name: High security connections for executives
- Access Method: Dial-up access
- User or Group: Select **Group**, and then specify the Executives group (example).
- Authentication methods: Select **Extensible Authentication Protocol (EAP)** and, in **Type**, select **Smart Card or other Certificate**. If you have multiple computer certificates, click **Configure**, select the appropriate computer certificate, and then clear all other check boxes.
- Policy Encryption Level: Select the **Strongest encryption (MPPE 128-bit)** check box, and then clear all other check boxes.

For additional examples of remote access policies, see Remote Access Policies Examples.

If you have created new remote access policies, either delete the default remote access policies, or ensure that they are the last policies to be evaluated. For more information, see Delete a remote access policy and Change the policy evaluation order.

**Configuring the secondary IAS server on a different domain controller**

To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see Install IAS.
2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.
3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, see Authentication across forests.
4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS authentication and accounting on the VPN servers**

To configure each VPN server to use the primary and secondary IAS servers for authentication, authorization, and accounting of remote access connections, do the following:

1. If the VPN server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, configure the primary and secondary IAS servers in the perimeter network as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting.

2. If the VPN server is a computer running Windows NT server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT server 4.0 online Help for instructions about how to configure the primary and secondary IAS servers as RADIUS servers for RADIUS authentication.

3. If the VPN server is a third-party network access server (NAS), see the documentation for the NAS to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS servers).

**Configuring the primary IAS proxy in the perimeter network**

To configure the primary IAS proxy in the perimeter network, do the following:

1. On a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the perimeter network; install IAS as an optional networking component. For more information, see Install IAS. The computer on which IAS is installed is not required to be dedicated to forwarding RADIUS messages. You can install IAS on a Web server, file server, or DNS server.

2. If needed, configure additional UDP ports for RADIUS messages that are sent by the service provider's RADIUS proxies. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

3. Add the service provider's RADIUS proxies as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

4. Create a connection request policy that forwards RADIUS request messages based on the realm name of your organization.

   Use the New Connection Request Policy Wizard to create a connection request policy that forwards connection requests to a remote RADIUS server group and where the realm name matches the realm name of the user accounts in your organization. Clear the check box that removes the realm name for authentication. In the New Connection Request Policy Wizard, use the New Remote RADIUS server Group Wizard to create a remote RADIUS server group with members that include the two IAS servers within your intranet.

   For more information, see Add a connection request policy.

5. Delete the default connection request policy named **Use Windows authentication for all users**. For more information, see Delete a connection request policy.

**Configuring the secondary IAS proxy in the perimeter network**

To configure the secondary IAS proxy on another computer in the perimeter network, do the following:

1. On another computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the perimeter network; install IAS as an optional networking component. For more information, see Install IAS.

2. Copy the configuration of the primary IAS proxy to the secondary IAS proxy in the perimeter network. For more information, see Copy the IAS configuration to another server.

**Configuring the intranet and Internet firewalls to support RADIUS traffic**

To configure the intranet and Internet firewalls to support RADIUS traffic, do the following:

1. Configure the intranet firewall to allow RADIUS traffic between the IAS proxies in the perimeter network and the IAS servers in the intranet.

2. Configure the Internet firewall to allow RADIUS traffic between the IAS proxies in the perimeter network and the IAS proxies in the service provider's network.

   For more information, see IAS and firewalls

**Configuring the primary IAS proxy at the service provider**
To configure the primary IAS proxy at the service provider, do the following:

1. On a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the service provider's network; install IAS as an optional networking component. For more information, see Install IAS. The computer on which IAS is installed is not required to be dedicated to forwarding RADIUS messages. You can install IAS on a Web server, file server, or DNS server.

2. If needed, configure additional UDP ports for authentication and accounting messages that are sent by the service provider's dial-up access servers. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

3. Add the service provider's dial-up access servers as RADIUS clients of the IAS proxy. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

4. Create a connection request policy that forwards RADIUS request messages based on the realm name of the service provider's customer.

   Use the New Connection Request Policy Wizard to create a connection request policy that forwards connection requests to a remote RADIUS server group and where the realm name matches the realm name of the customer's organization. Clear the check box that removes the realm name for authentication. In the New Connection Request Policy Wizard, use the New Remote RADIUS server Group Wizard to create a remote RADIUS server group with members that include the two IAS proxies in the customer's perimeter network.

   For more information, see Add a connection request policy.

5. Delete the default connection request policy named **Use Windows authentication for all users**. For more information, see Delete a connection request policy.

**Configuring the secondary IAS proxy at the service provider**
To configure the secondary IAS proxy on another computer in the perimeter network, do the following:

1. On another computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition in the service provider's network; install IAS as an optional networking component. For more information, see Install IAS.

2. Copy the configuration of the primary IAS proxy to the secondary IAS proxy in the service provider's network. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS accounting and authentication on the dial-up servers at the service provider**
To configure each dial-up server to use the primary and secondary IAS proxies in the service provider's network for the authentication, authorization, and accounting of dial-up connections, do the following:

1. If the dial-up or VPN server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, configure the primary and secondary IAS proxies in the service provider's network as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting.

2. If the dial-up or VPN server is a computer running Windows NT server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT server 4.0 online Help for information about how to configure the primary and secondary IAS proxies in the service provider's network as RADIUS servers for RADIUS authentication.

3. If the dial-up or VPN server is a third-party network access server (NAS), see the documentation for the NAS to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS proxies in the

service provider's network).

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Using a different authentication database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using a different authentication database**
This topic describes how IAS can be used as a RADIUS proxy to forward RADIUS messages between access servers and two different user account databases. An Active Directory forest and another user account database are used as examples.

The following configuration is for an organization that uses:

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints. To optimize IAS authentication and authorization response times and minimize network traffic, IAS is installed on domain controllers.

- Another user account database.

  There is another user account database that is used to authenticate and authorize connection attempts. Examples of other user account databases are Novell Directory Services (NDS) and Structured Query Language (SQL). Because IAS can only access Windows-based user account databases, a third-party RADIUS server must be used to access another user account database.

- Two IAS servers for each authentication database.

  Two IAS servers (one primary and one secondary) are used to provide fault tolerance for RADIUS-based authentication, authorization, and accounting for each user account database. If only one RADIUS server is configured and it becomes unavailable, access clients for that user account database cannot connect. By using two IAS servers and configuring the IAS proxies for both primary and secondary IAS servers for each account database, the IAS proxies can detect when the primary RADIUS server is unavailable and automatically fail over to the secondary IAS server.

- Remote access policies.

  Remote access policies are configured to specify, based on group membership, the different types of connection constraints for users.

- Two IAS proxies.

  Two IAS proxies are used to provide fault tolerance for RADIUS requests that are sent from access servers.

The following illustration shows the use of IAS proxies to forward RADIUS messages between the access servers and the RADIUS servers that are accessing two different user account databases.



**Note**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory domains, access servers, the third-party RADIUS servers, or the other user account database. For more information about how to deploy these components, see the appropriate Help topics.

To configure IAS for this example, complete the following steps:

- Configure the Active Directory forest for user accounts and groups.
- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.
- Configure the primary and secondary RADIUS servers for the other user account database.
- Configure the primary IAS proxy.
- Configure the secondary IAS proxy.
- Configure RADIUS authentication and accounting on the access servers.

**Configuring the forests for user accounts and groups**

To configure the forest for user accounts and groups, do the following:

1. Ensure that all users that are making network access connections have a corresponding user account.
2. Manage your network access by group by setting the remote access permission on user accounts to **Control access through Remote Access Policy**. For more information, see Configure remote access permission for a user.
3. Organize your remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For more information, see Group scope.
4. If you are using the Challenge-Handshake Authentication Protocol (CHAP), enable support for reversibly encrypted passwords for the appropriate domains. For more information, see Enable reversibly encrypted passwords in a domain.

**Configuring the primary IAS server on a domain controller**

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see Install IAS.
2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.
3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, see Authentication across forests.
4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.
5. Add the IAS proxies as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

6. Create the appropriate remote access policies for access clients.

   For examples of remote access policies, see Remote Access Policies Examples.

**Configuring the secondary IAS server on a different domain controller**

To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see Install IAS.

2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.

   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, see Authentication across forests.

4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see Copy the IAS configuration to another server.

**Configuring the primary and secondary RADIUS servers for the other user account database**

To configure the primary and secondary RADIUS servers to access the additional user account database and to add the two IAS proxies as RADIUS clients, see the documentation for the RADIUS servers.

**Configuring the primary IAS proxy**

To configure the primary IAS proxy, do the following:

1. On a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; install IAS as an optional networking component. For more information, see Install IAS. The computer on which IAS is installed is not required to be dedicated to forwarding RADIUS messages. For example, you can install IAS on a Web server, file server, or DNS server.

2. If needed, configure additional UDP ports for RADIUS messages that are sent by the access servers. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

3. Add the access servers as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

4. Create a connection request policy that forwards RADIUS request messages to the IAS servers where the realm name matches the user accounts in the forest.

   Use the New Connection Request Policy Wizard to create a connection request policy that forwards connection requests to a remote RADIUS server group and where the realm name matches the realm name for the user accounts in the forest. Clear the check box that removes the realm name for authentication. In the New Connection Request Policy Wizard, use the New Remote RADIUS server Group Wizard to create a remote RADIUS server group with members that include the two IAS servers.

   For more information, see Add a connection request policy.

5. Create a connection request policy that forwards RADIUS request messages based on the realm name of accounts in the additional user account database to the RADIUS servers.

Use the New Connection Request Policy Wizard to create a connection request policy that forwards connection requests to a remote RADIUS server and where the realm name matches the realm name for the user accounts in the additional user account database. Clear the check box that removes the realm name for authentication. In the New Connection Request Policy Wizard, use the New Remote RADIUS server Group Wizard to create a remote RADIUS server group with members that include the two RADIUS servers.

For more information, see Add a connection request policy.

6. Delete the default connection request policy named **Use Windows authentication for all users**. For more information, see Delete a connection request policy.

**Configuring the secondary IAS proxy**

To configure the secondary IAS proxy on another computer, do the following:

1. On another computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; install IAS as an optional networking component. For more information, see Install IAS.

2. Copy the configuration of the primary IAS proxy to the secondary IAS proxy. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS authentication and accounting on the access servers**

To configure each access server to use the primary and secondary IAS proxies for the authentication, authorization, and accounting of network connections, do the following:

1. If the dial-up or VPN server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, configure the primary and secondary IAS proxies as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting.

2. If the dial-up or VPN server is a computer running Windows NT server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT server 4.0 online Help for information about how to configure the primary and secondary IAS proxies as RADIUS servers for RADIUS authentication.

3. If the dial-up server, VPN server, wireless access point, or authenticating switch is a third-party access server, see the documentation for the access server to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS proxies).

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Using IAS proxy for load balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using IAS proxy for load balancing**

This topic describes how IAS proxies, placed between access servers and IAS servers, can be used to balance the load of a large volume of authentication traffic. Without IAS proxies, each access server balances its RADIUS requests across multiple RADIUS servers and detects unavailable RADIUS servers. By using IAS proxy, consistent load balancing is used to spread the load of authentication, authorization, and accounting traffic across all the IAS servers in the organization. Additionally, there is a consistent scheme for failure detection and RADIUS server failover.

The following configuration is for an organization that uses:

- Active Directory domains.

  Active Directory domains contain the user accounts, passwords, and dial-in properties that each IAS server requires to authenticate user credentials and evaluate both authorization and connection constraints. To optimize IAS authentication and authorization response times and minimize network traffic, IAS is installed on domain controllers.

- Multiple IAS servers.

  To balance the load of RADIUS authentication, authorization, and accounting across multiple domain controllers, IAS is installed as a RADIUS server on them.

- Remote access policies.

  Remote access policies are configured to specify, based on group membership, the different types of connection constraints for users.

- Two IAS proxies.

  Two IAS proxies are used to provide fault tolerance for RADIUS requests that are sent from access servers.

The following illustration shows IAS proxies used to balance the load of RADIUS traffic across multiple IAS servers.




**Note**

- This topic only describes how to configure IAS. It does not describe the configuration of Active Directory domains or access servers. For more information about how to deploy these components, see the appropriate Help topics.

To configure IAS for this example, complete the following steps:

- Configure the Active Directory forests for user accounts and groups.

- Configure the IAS server on domain controllers.
- Configure the primary IAS proxy.
- Configure the secondary IAS proxy.
- Configure RADIUS authentication and accounting on access servers.

**Configuring forests for user accounts and groups**

To configure Active Directory for user accounts and groups, do the following:

1. Ensure that all users that are making network access connections have a corresponding user account.
2. Manage your network access by group by setting the remote access permission on user accounts to **Control access through Remote Access Policy**. For more information, see Configure remote access permission for a user.
3. Organize your remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For more information, see Group scope.
4. If you are using the Challenge-Handshake Authentication Protocol (CHAP), enable support for reversibly encrypted passwords for the appropriate domains. For more information, see Enable reversibly encrypted passwords in a domain.

**Configuring the IAS server on domain controllers**

To configure IAS as a RADIUS server on each domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see Install IAS.
2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see Enable the IAS server to read user accounts in Active Directory.
3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains. For more information, see Enable the IAS server to read user accounts in Active Directory. For more information about trust relationships, see Trust direction.
4. Enable file logging for accounting and authentication events. For more information, see Configure log file properties.
5. Add the IAS proxies as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.
6. Create the appropriate remote access policies for access clients.

   For examples of remote access policies, see Remote Access Policies Examples.

**Configuring the primary IAS proxy**

To configure the primary IAS proxy, do the following:

1. On a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; install IAS as an optional networking component. For more information, see Install IAS. The computer on which IAS is installed is not required to be dedicated to forwarding RADIUS messages. For example, You can install IAS on a Web server, file server, or DNS server.
2. If needed, configure additional UDP ports for RADIUS messages that are sent by the access servers. For more information, see Configure IAS port information. By default, IAS uses UDP ports 1812 and 1645 for authentication and ports 1813 and 1646 for accounting.

3. Add the access servers as RADIUS clients of the IAS server. For more information, see Add RADIUS clients. Verify that you are configuring the correct name or IP address and shared secrets. For more information, see Shared secrets.

4. Use the New Remote RADIUS server Group Wizard to create a custom remote RADIUS server group. Add each IAS server as a member of the remote RADIUS server group and configure each group member with the priority of 1 and a weight of 50.

5. Create a connection request policy that forwards RADIUS request messages to the IAS servers where the realm name matches the accounts in the domain.

   Use the New Connection Request Policy Wizard to create a connection request policy that forwards connection requests to a remote RADIUS server and where the realm name matches the realm name for the user accounts in the forest. Clear the check box that removes the realm name for authentication. Select the previously created remote RADIUS server group as the group to forward connection requests.

   For more information, see Add a connection request policy.

6. Delete the default connection request policy named **Use Windows authentication for all users**. For more information, see Delete a connection request policy.

**Configuring the secondary IAS proxy**

To configure the secondary IAS proxy on another computer, do the following:

1. On another computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition; install IAS as an optional networking component. For more information, see Install IAS.

2. Copy the configuration of the primary IAS proxy to the secondary IAS proxy. For more information, see Copy the IAS configuration to another server.

**Configuring RADIUS authentication and accounting on the access servers**

To configure each access server to use the primary and secondary IAS proxies for the authentication, authorization, and accounting of network connections, do the following:

1. If the dial-up or VPN server is a computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows 2000 and the Routing and Remote Access service, configure the primary and secondary IAS proxies as RADIUS servers for both RADIUS authentication and accounting. For more information, see Use RADIUS authentication and Use RADIUS accounting.

2. If the dial-up or VPN server is a computer running Windows NT server 4.0 and the Routing and Remote Access Service (RRAS), see the Windows NT server 4.0 online Help for information about how to configure the primary and secondary IAS proxies as RADIUS servers for RADIUS authentication.

3. If the dial-up or VPN server is a third-party access server, see the documentation for the access server to determine how to configure it as a RADIUS client with two RADIUS servers (the primary and secondary IAS proxies).

### Note

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Mapping network authentication and authorization

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Mapping network authentication and authorization**

You can use IAS proxy to split network access authorization and authentication between two user account databases. You can authenticate visitors (such as partner organization employees) with their password-based credentials stored in an external partner organization user database, while performing authorization with your local Active Directory user account database. To split authentication and authorization, you must create a local user account for each visitor. The visitor account user name must map to an account in the partner organization user database with the same user name. When the visitor logs on to your network, the external user database provides authentication, while the local user database and remote access policies are used to process authorization.

The forwarding of authentication requests to the partner organization RADIUS server is necessary only when visitor authentication is performed with password-based credentials. Visitors can also be authenticated by mapping certificates to local user accounts. If a partner organization uses certificate-based authentication, you can authenticate visitors by trusting the partner organization certification authority (CA) and mapping visitor certificates to local user accounts. Because the visitor is authenticated locally with a certificate, it is not necessary to forward authentication requests to the remote RADIUS server at the partner organization.

For more information about certificate mapping, see Mapping certificates to user accounts, Map a certificate to a user account, and Map an account from a trusted non-Windows kerberos realm to a user account.

**Mapping with password-based authentication**

When partner organizations use password-based authentication, you must configure local user accounts for visitors and configure the IAS proxy server to forward authentication requests to a remote RADIUS server group that contains the partner company RADIUS server.

**Configuring visitor user accounts**

If you are using password-based authentication, you must create a user account for each visitor. The visitor user account must be created so that the user name in the local database and the user name in the partner organization user database match exactly. To allow visitors access to your organization network while authenticating them at their partner organization, perform the following actions to configure visitor accounts:

- **Add the** user principal name suffix **to the domain.**

  For more information, see Add user principal name suffixes.

- **Create an** organizational unit **that will contain the user accounts for visitors.**

  For more information, see Create a new organizational unit.

- **Create user accounts in the organizational unit for visitors.**

  When you create the user account, the user logon name must exactly match the visitor's user logon name as configured in the user account database at the partner organization. For example, if the user's user logon name at Microsoft Corporation is user@microsoft.com, designate "user" as the **User logon name** while creating the new user account with **Active Directory Users and Computers**. Also, select the correct UPN suffix from the drop-down list of suffixes. For example, if in step one you added the UPN suffix "microsoft.com", select "@microsoft.com" from the list. For more information, see Create a new user account

- **Configure user account properties**

  Make sure that the user account is configured in a way that will allow the visitor network access, but will also contain any necessary security restrictions. For example, in user account properties, on the **Dial-in** tab, you can set **Remote Access Permission (Dial-in or VPN)** to **Allow access**.

**Configuring the IAS proxy server**

To allow visitors access to your organization network, you need to perform the following actions on your IAS proxy server:

- **Configure the partner organization RADIUS server as a member of a Remote RADIUS server group.**

  Run the New Remote RADIUS Server Group Wizard, create and name a custom group, and then add the partner organization RADIUS server (or servers) to the group. While adding each server, configure the shared secret (obtained from the partner organization) on the **Authentication/Accounting** tab of the **Add RADIUS server** dialog box.

  For more information, see Add a remote RADIUS server group and Configure the Members of a Remote RADIUS Server Group.

- **Configure a connection request policy for visitor access.**

  To determine whether a specific connection attempt request or an accounting message received from a RADIUS client should be processed locally or forwarded to another RADIUS server, the IAS server uses connection request processing. When you allow visitor access with external authentication, you must configure a minimum of two connection request policies on your IAS proxy server. One policy allows your organization employees access, while denying access to other network access attempts. Another policy allows visitor access. When you configure your policy evaluation order, place the policy or policies for your organization users last, with the visitor policy processed first. For more information, see Change the policy evaluation order.

  To configure a connection request policy, run the New Connection Request Policy Wizard, create and name a custom policy, and then add **Policy Conditions**. You can configure as many conditions as you want, but you must include the **User-Name** attribute. Use wildcard and other characters described in Pattern matching syntax to configure a value that must be matched by the user principal name (UPN) suffix your visitors will use when logging on. For example, if you want to authenticate all visitors from Microsoft Corporation with this connection request policy, you can type "^.@microsoft\.com$" in **User-Name**. In **Request Processing Method**, click **Edit Profile**, and then click the **Advanced** tab. Click **Add**, and then add the attribute **Remote-RADIUS-to-Windows-User-Mapping**, with the attribute value set to **True**. This attribute tells the IAS proxy server to authenticate users (whose properties match other policy conditions) against the RADIUS servers configured in the Remote RADIUS Server Group. On the **Attribute** tab, you can configure replacement text for the **User-Name** attribute, as well as configure any other attributes that you require. On the **Authentication** tab, click **Forward requests to the following remote RADIUS server group for authentication**, and then choose the Remote RADIUS server group that can authenticate your visitors.

  For more information, see Add a connection request policy, Realm names, Configure Connection Request Processing, and Connection Request Processing.

- **Configure a remote access policy for visitor access.**

  Because you can create remote access policies that are based on different types of connections and group membership, you can create a remote access policy used specifically to process network access attempts by visitors. At minimum, you need two remote access policies: one for your employees and a different one for visitors. For more information, see Introduction to remote access policies and Remote Access Policies Examples.

- **Ensure that network access servers are configured as RADIUS clients to this IAS proxy.**

  When visitors log on to your network, they do so through network access servers, such as wireless access points. These access points must be configured as RADIUS clients to your IAS proxy server.

  For more information, see Configure RADIUS Clients.

After the IAS proxy is configured to forward authentication requests and user accounts for visitors are created in Active Directory, the password-based authentication request is forwarded to the partner organization RADIUS server when the visitor logs onto your organization network. If the user is authenticated, authorization is performed against your organization Active Directory user database. Active Directory returns user attributes (values configured for the user account, such as **Allow access** on the **Dial-in** tab of user properties), and IAS performs policy evaluation against these local user account attributes.

**Important**

- By configuring a local user account with the UPN of an account in the remote partner organization user account database, you map one account to the other. If the partner organization changes the user name of an account you have mapped, you must change the user name in your user account database, too.

**Mapping with certificate-based authentication**

Certificate-based authentication is more secure than password-based authentication. In addition, when you map certificates to user accounts instead of using password-based authentication methods, an authentication request is not forwarded to the partner organization RADIUS server and user account database. For these reasons, certificate mapping enhances security and can significantly reduce logon time for users.

A certificate is mapped to a user account in one of two ways: a single certificate is mapped to a single user account (one-to-one mapping) or multiple certificates are mapped to one user account (many-to-one mapping).

**One-to-one mapping**

For one-to-one mapping of certificates to visitor user accounts, you must perform the following steps:

- **Create a user account for each visitor according to the steps in "Configuring visitor user accounts."**

  Because the user account itself is not mapped to the user account at the partner organization, the user account names do not have to exactly match the user account names in the partner organization user accounts database. Configuring the local account with the same user name, however, can make it easier to configure realm manipulation rules.

- **Map the certificate to a user account.**

  For more information, see Mapping certificates to user accounts and Map a certificate to a user account

- **Perform cross-certification or authorize the partner organization CA as a qualified subordinate certification authority using a computer running** Windows Server 2003, Enterprise Edition **and Certificate Services.**

  Performing cross-certification or authorizing the partner organization CA as a qualified subordinate CA is recommended for domains with a Windows Server 2003 domain functional level and clients running Windows XP. If your domains are Windows 2000 native, it is recommended that you use a certificate trust list (CTL) instead.

  For more information, see Domain and forest functionality, Qualified subordination, Qualified subordination overview, and Perform qualified subordination.

- **Configure your IAS server or IAS proxy server**

  Because authentication requests are not forwarded to external RADIUS servers when authentication is performed with certificates, you are not required to use IAS as a proxy server.

  - Configure a connection request policy for visitor access with the New Connection Request Policy Wizard.

    For more information, see Add a connection request policy, Pattern matching syntax, Realm names, Configure Connection Request Processing, and Connection Request Processing.

  - Create a realm manipulation rule to map the partner organization realm name, which is contained in the user's certificate, to the user's local user account.

    For example, if user@tailspintoys.com is mapped to user@microsoft.com, the realm manipulation rule must replace "@tailspintoys.com" with "@microsoft.com" using pattern-matching syntax.

    For more information, see Realm names and Pattern matching syntax.

  - Configure a remote access policy for visitor access. For more information, see Introduction to remote access policies and Remote Access Policies Examples.

  - Ensure that network access servers are configured as RADIUS clients to this IAS server.

When visitors log on to your network, they do so through network access servers, such as wireless access points. These access points must be configured as RADIUS clients to your IAS proxy server or IAS server. For more information, see Configure RADIUS Clients.

**Many-to-one mapping**

Many-to-one certificate mapping allows you to map many certificates to one user account. After you have trusted the enterprise root CA of a partner organization, you can map all certificates issued by the partner organization CA to one account that you create in your local domain. This solution provides ease of management for many users, and eliminates the need to create an individual user account for every visitor to your organization.

For many-to-one mapping of certificates, you must perform the following steps:

- **Perform cross-certification or authorize the partner organization CA as a qualified subordinate certification authority using a computer running** Windows Server 2003, Enterprise Edition **and Certificate Services.**

  Performing cross-certification or authorizing the partner organization CA as a qualified subordinate CA is recommended for domains with a Windows Server 2003 domain functional level and clients running Windows XP. If your domains are Windows 2000 native, it is recommended that you use a certificate trust list (CTL) instead.

  For more information, see Domain and forest functionality, Qualified subordination, Qualified subordination overview, and Perform qualified subordination.

- **Create one user account and map the partner organization certificate to the account.**

  For more information, see Mapping certificates to user accounts and Map a certificate to a user account.

- **Configure your IAS server or IAS proxy server**

  Because authentication requests are not forwarded to external RADIUS servers when authentication is performed with certificates, you are not required to use IAS as a proxy server.

  - Configure a connection request policy for visitor access with the New Connection Request Policy Wizard.

    For more information, see Add a connection request policy, Pattern matching syntax, Realm names, Configure Connection Request Processing, and Connection Request Processing.

  - Create a realm manipulation rule to map the partner organization realm name, which is contained in the user's certificate, to the user's local user account.

    For example, if user@tailspintoys.com is mapped to user@microsoft.com, the realm manipulation rule must replace "@tailspintoys.com" with "@microsoft.com" using pattern-matching syntax.

    For more information, see Realm names and Pattern matching syntax.

  - Configure a remote access policy for visitor access. For more information, see Introduction to remote access policies and Remote Access Policies Examples.

  - Ensure that network access servers are configured as RADIUS clients to this IAS server.

    When visitors log on to your network, they do so through network access servers, such as wireless access points. These access points must be configured as RADIUS clients to your IAS proxy server or IAS server. For more information, see Configure RADIUS Clients.

For more information about designing your IAS proxy server deployment, see IAS as a RADIUS proxy design considerations.

**Notes**

- After you run the New Remote RADIUS Server Group Wizard, the New Connection Request Policy Wizard is opened by

default. If you do not want the New Connection Request Policy Wizard to open, click **Start the New Connection Request Policy Wizard when this wizard closes** in the New Remote RADIUS Server Group Wizard.

- For strong security between your IAS proxy server and your partner organization RADIUS servers, you can use Internet Protocol security (IPSec). For more information, see Securing RADIUS traffic with IPSec. For additional security recommendations and concerns, see IAS as a RADIUS proxy security considerations and Security information for IAS.

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Setting up IAS in a Test Lab

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up IAS in a test lab**

This section describes how you can use five computers to create a test lab to configure and test the Internet Authentication Service (IAS) features in Windows Server 2003, Standard Edition. These instructions are designed to walk you through a set of tasks that expose you to IAS and its associated functionality. Beyond the set of tasks, you can use these instructions to create a functioning IAS configuration and experiment with IAS features and functionality before deploying these technologies on your production network.

This section covers:

- Setting up the IAS test lab infrastructure
- IAS test lab tasks

**Notes**

- The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.
- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Setting up the IAS test lab infrastructure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up the infrastructure**

The infrastructure for the IAS test lab network consists of five computers performing the following services:

- A computer running Windows Server 2003, Standard Edition that is used as both a domain controller and a Domain Name System (DNS) server. This computer is named DC1.
- A computer running Windows Server 2003, Standard Edition that is used as a Remote Authentication Dial-in User Service (RADIUS) server. This computer is named IAS1.
- A computer running Windows Server 2003, Standard Edition that is used as a RADIUS proxy. This computer is named IAS2.
- A computer running Windows Server 2003, Standard Edition that is used as a VPN server. This computer is named VPN1.
- A computer running Windows XP Professional that is used as a VPN client. This computer is named CLIENT1.

The following illustration shows the configuration of the IAS test lab.




There are separate network segments representing a corporate intranet and the Internet. All computers on the corporate intranet are connected to a common hub or Layer 2 switch. All computers on the Internet are connected to a separate common hub or Layer 2 switch. Private addresses are used throughout the test lab configuration. The private network ID 172.16.0.0/24 is used for the intranet, and the private network ID 10.0.0.0/24 is used for the simulated Internet.

Each computer is manually configured with the appropriate IP address, subnet mask, and DNS server. There are no Dynamic Host Configuration protocol (DHCP) or Windows Internet Name Service (WINS) servers present.

The following sections describe how each of the computers in the test lab are configured. To reconstruct this test lab, please configure the computers in the order presented.

**Note**

- The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

**DC1**

DC1 is a computer running Windows Server 2003, Standard Edition. It is providing the following services:

- A domain controller for the testlab.microsoft.com domain.
- A DNS server for the testlab.microsoft.com DNS domain.
- A certification authority for the testlab.microsoft.com domain.

**Note**

- The domain name testlab.microsoft.com is used here for example purposes only. You can use any domain name in your test lab configuration.

To configure DC1 to perform these services, perform the following steps:

1. Install Windows Server 2003, Standard Edition as a standalone server.
2. After restarting, log on using the local Administrator account.
3. Configure the TCP/IP protocol with the IP address of 172.16.0.1 and the subnet mask of 255.255.255.0. For more information, see Configure TCP/IP for static addressing.
4. At the command prompt, run **dcpromo** for a new domain named testlab.microsoft.com in a new forest. Install the DNS service when prompted.
5. After restarting, log on using the Administrator account to the testlab.microsoft.com domain.
6. Install Certificate Services as an enterprise root certification authority. For more information, see Install an enterprise root certification authority.
7. Configure the testlab.microsoft.com domain for automatic enrollment of computer certificates. For more information, see Configure automatic certificate allocation from an enterprise CA.

**IAS1**

IAS1 is a computer running Windows Server 2003, Standard Edition. It is providing RADIUS authentication, authorization, and accounting for VPN1, the VPN server computer.

To configure IAS1 as a RADIUS server, perform the following steps:

1. On DC1, add a computer account for IAS1. For more information, see Create a new computer account.
2. Install Windows Server 2003, Standard Edition as a standalone server.
3. After restarting, log on using the local Administrator account.
4. Configure the TCP/IP protocol with the IP address of 172.16.0.2, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.
5. Join IAS1 to the testlab.microsoft.com domain. For more information, see Join a domain.
6. After restarting, log on using the Administrator account to the testlab.microsoft.com domain.
7. Install the Internet Authentication Service. For more information, see Install IAS.
8. Configure the IAS server to access user account properties in the testlab.microsoft.com domain. For more information, see Enable the IAS server to read user accounts in Active Directory.

9. Install Network Monitor. For more information, see Install Network Monitor.

**IAS2**

IAS2 is a computer running Windows Server 2003, Standard Edition. IAS2 is a RADIUS proxy for VPN1, the VPN server computer.

To configure IAS2 as a RADIUS proxy, perform the following steps:

1. Install Windows Server 2003, Standard Edition as a standalone server.

2. After restarting, log on using the local Administrator account.

3. Configure the TCP/IP protocol with the IP address of 172.16.0.3, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

4. Install the Internet Authentication Service. For more information, see Install IAS.

**VPN1**

VPN1 is a computer running Windows Server 2003, Standard Edition. It is providing VPN server services for VPN clients. To configure VPN1 as a VPN server, perform the following steps:

1. From DC1, add a computer account for the VPN1 computer. For more information, see Create a new computer account.

2. Install Windows Server 2003, Standard Edition as a standalone server.

3. After restarting, log on using the local Administrator account.

4. For the intranet local area connection, configure the TCP/IP protocol with the IP address of 172.16.0.4, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

5. For the Internet local area connection, configure the TCP/IP protocol with the IP address of 10.0.0.2 and the subnet mask of 255.255.255.0. For more information, see Configure TCP/IP for static addressing.

6. Join the VPN1 computer to the testlab.microsoft.com domain. For more information, see Join a domain.

7. After restarting, log on using the Administrator account in the testlab.microsoft.com domain.

8. Configure and enable the Routing and Remote Access service. For more information, see Enable the Routing and Remote Access service. In the Routing and Remote Access server Setup Wizard, select **Virtual private network (VPN) server** from the **Common Configurations** dialog box. When prompted for IP address assignment, select **From a specified range of addresses**, and then configure the range 172.16.0.248 through 172.16.0.255. Do not configure RADIUS authentication.

**CLIENT1**

CLIENT1 is a computer running Windows XP Professional. It is used as a VPN client to gain remote access to intranet resources across the simulated Internet. To configure CLIENT1 as a VPN client, perform the following steps:

1. On DC1, add a computer account for CLIENT1. For more information, see Create a new computer account.

2. Connect CLIENT1 to the intranet network segment.

3. On CLIENT1, install Windows XP Professional as a workgroup computer.

4. After restarting, log on using the local Administrator account.

5. Type **gpupdate /target:computer** at a command prompt to ensure that a computer certificate is installed on CLIENT1.

6. Configure the TCP/IP protocol with the IP address of 172.16.0.5, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

7. Join CLIENT1 to the testlab.microsoft.com domain. For more information, see Join a domain.

8. After restarting, log on using the Administrator account in the testlab.microsoft.com domain.

9. Configure the TCP/IP protocol with the IP address of 10.0.0.1 and the subnet mask of 255.255.255.0. Do not configure a DNS server IP address. For more information, see Configure TCP/IP for static addressing and Configure TCP/IP to use DNS.

10. Shut down CLIENT1, disconnect it from the intranet network segment, and then connect it to the simulated Internet network segment.

11. Restart CLIENT1.

12. After restarting, log on using the local Administrator account.

# IAS test lab tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS test lab tasks**

The following tasks are designed to take you through the common elements of setting up RADIUS server and proxy support with Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition:

- PPTP and L2TP remote access VPN connections
- IAS as a RADIUS server for the VPN server
- Remote access policies for different types of VPN connections
- IAS as a RADIUS proxy for the VPN server

### Note

- The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

**PPTP and L2TP remote access VPN connections**

To create a PPTP-based remote access VPN connection between CLIENT1 and VPN1, do the following:

**1. Create a user account**

1. On DC1, use Active Directory Users and Computers to create a user account named **PPTPUser** with a password. For more information, see Create a new user account.
2. Set the remote access permission on the **Dial-in** tab to **Allow access**. For more information, see Configure remote access permission for a user.

**2. Create the PPTP connection**

1. On CLIENT1, use the Make New Connection Wizard to create a new virtual private network connection named **PPTP**, with the IP address of 10.0.0.2. For more information, see Make a virtual private network (VPN) connection.
2. On CLIENT1, right-click the new **PPTP** connection, and then click **Properties**.
3. Click the **Networking** tab, and then in **Type of VPN**, click **PPTP VPN**.
4. Click **OK** to save changes to the PPTP connection.

**3. Make the PPTP connection**

1. On CLIENT1, double-click the **PPTP** connection.
2. In the **PPTP** dialog box, type **PPTPUser@testlab.microsoft.com** as the user name, type the password, and then select the **Save this user name and password to use when** check box.
3. Click **Connect**.

**4. Disconnect the PPTP connection**

- On CLIENT1, right-click the **PPTP** connection, and then click **Disconnect**.

To create an L2TP/IPSec remote access VPN connection between CLIENT1 and VPN1, do the following:

**1. Create a user account**

1. On DC1, use Active Directory Users and Computers to create a user account named **L2TPUser** with a password. For more information, see Create a new user account.

2. Set the remote access permission on the **Dial-in** tab to **Allow access**. For more information, see Configure remote access permission for a user.

**2. Create the L2TP connection**

1. On CLIENT1, use the Make New Connection Wizard to create a new L2TP/IPSec connection named **L2TP**, with the IP address of **10.0.0.2**. For more information, see Make a virtual private network (VPN) connection.

2. Right-click the new **L2TP** connection, and then click **Properties**.

3. Click the **Networking** tab, and then in **Type of VPN**, click **L2TP IPSec VPN**.

4. Click **OK** to save changes to the **L2TP** connection.

**3. Make the L2TP connection**

1. On CLIENT1, double-click the **L2TP** connection.

2. In the **L2TP** dialog box, type **L2TPUser@testlab.microsoft.com** as the user name, type the password, and then select the **Save this user name and password to use when** check box.

3. Click **Connect**.

**IAS as a RADIUS server for the VPN server**

To configure IAS1 as the RADIUS authentication and accounting server for the VPN server, do the following:

**1. Add VPN1 as a RADIUS client to IAS1**

- On IAS1, add VPN1 as a RADIUS client with the client address of **172.16.0.4**, the client vendor of **Microsoft**, and a shared secret. For more information, see Add RADIUS clients.

**2. Configure IAS1 to log authentication events**

- On IAS1, configure remote access logging of authentication events. For more information, see Select requests to be logged.

**3. Configure VPN1 for RADIUS and IAS1 as the RADIUS server**

- Configure VPN1 to use RADIUS authentication. Configure IAS1 as the RADIUS server with the IP address of 172.16.0.1 and the shared secret. For more information, see Use RADIUS authentication.

- Configure VPN1 to use RADIUS accounting. Configure IAS1 as the RADIUS server with the IP address of 172.16.0.1 and the shared secret. For more information, see Use RADIUS accounting.

**4. Make a PPTP connection**

1. On CLIENT1, use the **PPTP** connection to create a PPTP-based VPN connection with VPN1. For more information, see Connect to a remote network.

2. At a Command Prompt, use the ping command to ping DC1 at its IP address of 172.16.0.1.

3. Use the ping command to ping IAS1 at its IP address of 172.16.0.2.

4. Disconnect the **PPTP** connection. For more information, see Disconnect from a remote network.

**5. Make an L2TP connection**

1. On CLIENT1, use the **L2TP** connection to create an L2TP/IPSec VPN connection with VPN1.

2. Use the ping command to ping DC1 at its IP address of 172.16.0.1.

3. Use the ping command to ping IAS1 at its IP address of 172.16.0.2.

4. Disconnect the **L2TP** connection.

**6. Check the system event log for RADIUS events**

- On IAS1, use Event Viewer to view IAS events in the system event log for the PPTP and L2TP connections that were recently created using CLIENT1.

**7. Check RADIUS authentication and accounting logs**

- On IAS1, use Windows Explorer to open the systemroot\System32\Logfiles\Iaslog.log file. View the authentication and accounting entries for the connections that were recently created using CLIENT1. For more information about the IAS log file format, see Interpreting IAS-formatted log files.

**8. Capture RADIUS traffic with Network Monitor**

1. On IAS1, run Network Monitor from the Administrative Tools folder.

2. Use Network Monitor to capture network traffic. For more information, see Capture network frames.

3. On CLIENT1, make a VPN connection with VPN1 with the **PPTP** connection.

4. On CLIENT1, disconnect the **PPTP** connection.

5. On IAS1, stop Network Monitor and view the captured frames with the protocol of **RADIUS**. For more information, see Expand and collapse frame details.

**Remote access policies for different types of VPN connections**
To create remote access policies for different types of VPN connections, do the following:

**1. Create different remote access policies for different types of connections**

1. On IAS1, use the New Remote Access Policy Wizard to create a new custom remote access policy with the following settings:

- Policy name: PPTP connections

- Conditions: NAS-Port-Type matches Virtual (VPN) and Tunnel-Type matches Point-to-Point Tunneling Protocol (PPTP)

- Permission: Grant remote access permission
- Profile settings, IP tab, Input packet filter
    - Deny all traffic except those listed below
    - Destination network, IP address: 172.16.0.1
    - Destination network, Subnet mask: 255.255.255.255
    - Protocol: Any
- Profile settings, IP tab, Output packet filter:
    - Deny all traffic except those listed below
    - Source network, IP address: 172.16.0.1
    - Destination network, Subnet mask: 255.255.255.255
    - Protocol: Any

For more information, see Add a remote access policy.

1. Use the New Remote Access Policy Wizard to create a new custom remote access policy with the following settings:

- Policy name: L2TP connections
- Conditions: NAS-Port-Type matches Virtual (VPN) and Tunnel-Type matches Layer Two Tunneling Protocol (L2TP)
- Permission: Grant remote access permission
- Profile settings, IP tab, Input packet filter:
    - Deny all traffic except those listed below
    - Destination network, IP address: 172.16.0.2
    - Destination network, Subnet mask: 255.255.255.255
    - Protocol: Any
- Profile settings, IP tab, Output packet filter:
    - Deny all traffic except those listed below
    - Source network, IP address: 172.16.0.2
    - Destination network, Subnet mask: 255.255.255.255
    - Protocol: Any

**2. Make a PPTP connection and test connectivity**

1. On CLIENT1, make a VPN connection with VPN1 using the **PPTP** connection.

2. Use the ping command to ping DC1 at its IP address of 172.16.0.1.

3. Use the ping command to ping IAS1 at its IP address of 172.16.0.2.

   This command fails because packet filtering for all connections that match the **PPTP connections** policy allows only traffic sent to and from the IP address of 172.16.0.1.

4. Disconnect the **PPTP** connection.

**3. Make an L2TP connection and test connectivity**

1. On CLIENT1, make a VPN connection with VPN1 using the **L2TP** connection.

2. Use the ping command to ping IAS1 at its IP address of 172.16.0.2.

3. Use the ping command to ping DC1 at its IP address of 172.16.0.1.

   This command fails because packet filtering for all connections that match the **L2TP connections** policy allows only traffic sent to and from the IP address of 172.16.0.2.

4. Disconnect the **L2TP** connection.

**4. Check the system event log for RADIUS events**

- On IAS1, use Event Viewer to view the IAS events in the system event log for the connections that were recently created. Note that the authentication event message text contains the name of the remote access policy that accepted the connection.

**IAS as a RADIUS proxy for the VPN server**

To configure IAS2 as the RADIUS proxy for the VPN server, do the following:

**1. Reconfigure VPN1 with IAS2 as its RADIUS server**

1. On VPN1, remove IAS1 as the RADIUS server for the RADIUS authentication provider. Next, configure IAS2 as the RADIUS server for the RADIUS authentication provider with the IP address of 172.16.0.3 and the shared secret. For more information, see Use RADIUS authentication.

2. On VPN1, remove IAS1 as the RADIUS server for the RADIUS accounting provider. Next, configure IAS2 as the RADIUS server for the RADIUS accounting provider with the IP address of 172.16.0.3 and the shared secret. For more information, see Use RADIUS accounting.

**2. Reconfigure IAS1 with IAS2 as a RADIUS client**

1. On IAS1, remove VPN1 as a RADIUS client. For more information, see Delete a RADIUS client.

2. On IAS1, add IAS2 as a RADIUS client with the client address of 172.16.0.4, the client vendor of Microsoft, and a shared secret. For more information, see Add RADIUS clients.

**3. Configure IAS2 as a RADIUS proxy between VPN1 and IAS1**

1. On IAS2, add VPN1 as a RADIUS client with the client address of **172.16.0.4**, the client vendor of **Microsoft**, and a shared secret. For more information, see Add RADIUS clients.

2. On IAS2, create a remote RADIUS server group that contains IAS1. For more information, see Add a remote RADIUS server group. In the New Remote RADIUS server Group Wizard, configure the following:

- Group name: Testlab
- Primary server: 172.16.0.2
- Shared secret: SharedSecret
  Do not configure a backup server.

3. Create a new connection request policy. For more information, see Add a connection request policy. In the New Connection Request Policy Wizard, configure the following:

   - Policy name: Proxy
   - Select the option to **Forward connection requests to a remote RADIUS server for authentication**
   - Realm name: testlab.microsoft.com
   - Clear the **Before authentication, remove the realm name from the user name: check box**.
   - server group: Testlab

**Note**

- IAS1 and IAS2 must use the same shared secret. Additionally, IAS2 and VPN1 must use the same shared secret. The shared secret that is used between IAS1 and IAS2 can be different than the shared secret that is used between IAS2 and VPN1.

#### 4. Make a PPTP connection

1. On CLIENT1, make a VPN connection with VPN1 using the PPTP connection.
2. Use ping to ping DC1 at its IP address of 172.16.0.1.
3. Use ping to ping IAS1 at its IP address of 172.16.0.2.

   This command fails because packet filtering for all connections that match the PPTP connections policy allows only traffic sent to and from the IP address of 172.16.0.1.
4. Disconnect the PPTP connection.

In this configuration, the authentication and accounting messages that are sent by the VPN server are being forwarded from IAS2 to IAS1.

# IAS Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This section provides information about:

- Troubleshooting tools to use with IAS
- Pattern matching syntax
- IAS Resources: Windows Deployment and Resource Kits
- Database-Import Log Files
- IAS-Formatted Log Files

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Troubleshooting tools to use with IAS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting tools to use with IAS**
To help you gather information to troubleshoot problems with IAS, the following troubleshooting tools are available:

- IAS event logging and Event Viewer
- Network Monitor
- System Monitor counters
- SNMP Service

**IAS event logging and Event Viewer**
Use Event Viewer to obtain information about hardware and software problems and to monitor all security events--including informational, warning, and error events.

To troubleshoot IAS authentication attempts in the system event log, enable event logging for IAS. Viewing the authentication attempts in this log is useful in troubleshooting remote access policies. When you have multiple remote access policies configured, you can use the system event log to determine the name of the remote access policy that either accepted or rejected the connection attempt. Enabling IAS event logging and reading the text of IAS authentication events in the system event log is the most useful tool for troubleshooting failed IAS authentications. For more information, see Event logging for IAS.

**Network Monitor**
You can use Network Monitor, available in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition; or a commercial packet analyzer (also known as a network sniffer), to capture and view RADIUS authentication and accounting messages that are sent to and from the IAS server. Network Monitor includes a RADIUS parser, which you can use to view the attributes of a RADIUS message and troubleshoot connection issues.

For more information, see Network Monitor.

**System Monitor counters**
You can use Performance console overview to monitor the resource use of specific components and program processes. With System Monitor, you can use charts and reports to determine how efficiently your server uses IAS and both identify and troubleshoot potential problems.

You can use System Monitor to monitor the following IAS-related performance objects:

- IAS Accounting Client
- IAS Accounting server
- IAS Authentication Client
- IAS Authentication server

**SNMP Service**
You can use SNMP to monitor status information for your IAS server. IAS supports the RADIUS Authentication server MIB (as specified in RFC 2619) and the RADIUS Accounting server MIB (as specified in RFC 2621).

# Pattern matching syntax

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Pattern matching syntax**

Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition support regular expressions for pattern matching, which is widely used in UNIX environments. You can use this syntax to specify the conditions of remote access policy attributes and RADIUS realms, as specified in the following table.

| Character | Description | Example |
|---|---|---|
| \ | Marks the next character as a character to match. | /n/ matches the character "n". The sequence /\n/ matches a line feed or newline character. |
| ^ | Matches the beginning of the input or line. | |
| $ | Matches the end of the input or line. | |
| * | Matches the preceding character zero or more times. | /zo*/ matches either "z" or "zoo." |
| + | Matches the preceding character one or more times. | /zo+/ matches "zoo" but not "z." |
| ? | Matches the preceding character zero or one times. | /a?ve?/ matches the "ve" in "never." |
| . | Matches any single character except a newline character. | |
| **(** *pattern* **)** | Matches *pattern* and remembers the match. To match ( ) (parentheses), use "\(" or "\)". | |
| *x* **\|** *y* | Matches either *x* or *y*. | /z\|food?/ matches "zoo" or "food." |
| **{** *n* **}** | Matches exactly *n* times (*n* is a nonnegative integer). | /o{2}/ does not match the "o" in "Bob," but matches the first two instances of the letter o in "foooood." |
| **{** *n* **,}** | Matches at least *n* times (*n* is a nonnegative integer). | /o{2,}/ does not match the "o" in "Bob" but matches all of the instances of the letter o in "foooood." /o{1,}/ is equivalent to /o+/. |
| **{** *n* **,** *m* **}** | Matches at least *n* and at most *m* times (*m* and *n* are nonnegative integers). | /o{1,3}/ matches the first three instances of the letter o in "fooooood." |
| **[** *xyz* **]** | Matches any one of the enclosed characters (a character set). | /[abc]/ matches the "a" in "plain." |
| **[^** *xyz* **]** | Matches any characters that are not enclosed (a negative character set). | /[^abc]/ matches the "p" in "plain." |
| **\b** | Matches a word boundary (for example, a space). | /ea*r\b/ matches the "er" in "never early." |
| **\B** | Matches a nonword boundary. | /ea*r\B/ matches the "ear" in "never early." |
| **\d** | Matches a digit character (equivalent to [0-9]). | |
| **\D** | Matches a nondigit character (equivalent to [^0-9]). | |
| **\f** | Matches a form feed character. | |
| **\n** | Matches a line feed character. | |

| | | |
|---|---|---|
| **\r** | Matches a carriage return character. | |
| **\s** | Matches any white space character including space, tab, and form feed (equivalent to [ \f\n\r\t\v]). | |
| **\S** | Matches any non-white space character (equivalent to [^ \f\n\r\t\v]). | |
| **\t** | Matches a tab character. | |
| **\v** | Matches a vertical tab character. | |
| **\w** | Matches any word character, including underscore (equivalent to [A-Za-z0-9_]). | |
| **\W** | Matches any nonword character, excluding underscore (equivalent to [^A-Za-z0-9_]). | |
| **** *num* | Refers to remembered matches (?*num*, where *num* is a positive integer). For example, \1 replaces what is stored in the first remembered match. This option can be used only in the **Replace** text box when configuring attribute manipulation. | |
| **/** *n* **/** | Allows the insertion of ASCII codes into regular expressions (?*n*, where *n* is an octal, hexadecimal, or decimal escape value). | |

**Examples for remote access policy attributes**

The following examples describe the use of the pattern matching syntax to specify remote access policy attributes:

- To specify all phone numbers within the 899 area code, the syntax is:

  899.*

- To specify a range of all IP addresses that begin with 192.168.1, the syntax is:

  192\.168\.1\..+

**Examples for manipulation of the realm name in the User-Name attribute**

The following examples describe the use of the pattern matching syntax to manipulate realm names for the User-Name attribute, which is located on the **Attribute** tab in the properties of a connection request policy. For more information, see Configure attribute manipulation.

- To remove the realm portion of the User-Name attribute

  In the outsourced dial scenario, the Internet service provider (ISP) might require a realm name to route the authentication request. However, the IAS server might not recognize the realm name portion of the user name. Therefore, the realm name must be removed before it is forwarded to the IAS server.

  - **Find:** @microsoft\.com
  - **Replace:**

- To replace *user@example.microsoft.com* with *example.microsoft.com\user*

  - **Find:** (.*)@(.*)
  - **Replace:** $2\$1

- To replace *domain\user* with *specific_domain\user*

  - **Find:** (.*)\\(.*)

- **Replace:** *specific_domain*\$2

- To replace *user* with *user@specific_domain*

  - **Find:** $

  - **Replace:** *@specific_domain*

**Example for RADIUS message forwarding by a proxy server**
You can create routing rules that forward RADIUS messages with a specific realm name to a specific set of RADIUS servers when IAS is used as a RADIUS proxy. Following is a recommended syntax for routing requests based on realm name.

| NetBIOS name | Pattern |
|---|---|
| WCOAST | ^wcoast\\ |

In the following example, wcoast.microsoft.com is a unique user principal name (UPN) suffix for the DNS or Active Directory domain wcoast.microsoft.com. Using the supplied pattern, the IAS proxy can route messages based on domain NetBIOS name or UPN suffix.

| NetBIOS name | UPN suffix | Pattern |
|---|---|---|
| WCOAST | wcoast.microsoft.com | ^wcoast\\\|@wcoast\.microsoft\.com$ |

**Ping user-name**
Some RADIUS proxy servers and network access servers periodically send authentication and accounting requests (known as ping requests) to verify that the IAS server is present on the network. These ping requests include fictional user names. When IAS processes these requests, the event and accounting logs become filled with access reject records, making it more difficult to keep track of valid records.

When you configure a registry entry for **ping user-name**, IAS matches the registry entry value against the user name value in ping requests by other servers. A **ping user-name** registry entry specifies the fictional user name (or a user name pattern, with variables, that matches the fictional user name) sent by RADIUS proxy servers and network access servers. When IAS receives ping requests that match the **ping user-name** registry entry value, IAS rejects the authentication requests without processing the request. IAS does not record transactions involving the fictional user name in any log files, which makes the event log easier to interpret.

### Note

- **Ping user-name** is not installed by default. You must add **ping user-name** to the registry. You can add an entry to the registry using Registry Editor. For more information, see Registry Editor.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Adding ping user-name to the registry**
**Ping user-name** can be added to the following registry key as a string value by a member of the local Administrators group:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IAS\Parameters**

| Name | Type | Data |
|---|---|---|
| ping user-name | REG_SZ | *User name* |

### Note

- To indicate more than one user name for a **ping user-name** value, enter a name pattern, such as a Domain Name

System (DNS) name including wildcard characters, in **Data**.

# IAS Resources: Windows Deployment and Resource Kits

Published: January 21, 2005

Updated: November 25, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more information about IAS, see Internet Authentication Service on TechNet at http://www.microsoft.com/ias.

# Database-Import Log Files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Database-import log files**

This section provides information about using database-import log files:

- Importing IAS log files into a database
- Interpreting database-import log files

# Importing IAS log files into a database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Importing IAS log files into a database**

When you select **Database compatible file format** for IAS accounting and authentication logs, the information is recorded in a comma-delimited format. The consistent structure of this format makes it possible to use a database program, such as Microsoft Access, to import a log into a table, which presents data in a way that is easily viewed and interpreted. And, after importing a log, you can filter and analyze the data in a way that is useful for your accounting and auditing processes.

**Importing IAS log files into a table with Microsoft Access 2002 or Microsoft Access 97**

The table that you create with Access is formatted according to the type of data that is contained in each field. Although this example shows you how to create a new table, you can also import a log into an existing table.

Before creating the table, add the two lines in the next paragraph to the beginning of the log file created by IAS. The first line contains the attribute names of each field in the log and Access uses them to create the field names in its table. Using attribute names in the table makes it easier to interpret the log entries. The second line is used by Access to set up the appropriate data type for each column in its table. After you import the log file, you should delete these entries from the Access table.

```
ComputerName,ServiceName,Record-Date,Record-Time,Packet-Type,User-Name,Fully-Qualified-
User-Name,Called-Station-ID,Calling-Station-ID,Callback-Number,Framed-IP-Address,NAS-
Identifier,NAS-IP-Address,NAS-Port,Client-Vendor,Client-IP-Address,Client-Friendly-
Name,Event-Timestamp,Port-Limit,NAS-Port-Type,Connect-Info,Framed-Protocol,Service-
Type,Authentication-Type,NP-Policy-Name,Reason-Code,Class,Session-Timeout,Idle-
Timeout,Termination-Action,EAP-Friendly-Name,Acct-Status-Type,Acct-Delay-Time,Acct-
Input-Octets,Acct-Output-Octets,Acct-Session-ID,Acct-Authentic,Acct-Session-Time,Acct-
Input-Packets,Acct-Output-Packets,Acct-Terminate-Cause,Acct-Multi-Ssn-ID,Acct-Link-
Count,Acct-Interim-Interval,Tunnel-Type,Tunnel-Medium-Type,Tunnel-Client-Endpt,Tunnel-
Server-Endpt,Acct-Tunnel-Conn,Tunnel-Pvt-Group-ID,Tunnel-Assignment-ID,Tunnel-
Preference,MS-Acct-Auth-Type,MS-Acct-EAP-Type,MS-RAS-Version,MS-RAS-Vendor,MS-CHAP-
Error,MS-CHAP-Domain,MS-MPPE-Encryption-Types,MS-MPPE-Encryption-Policy
x,x,01/01/1999,12:00:00,0,x,x,x,x,x,x,x,x,0,0,x,x,01/01/1999
12:00:00,0,0,x,0,0,0,x,0,x,0,0,0,x,0,0,0,0,x,0,0,0,0,0,x,0,0,0,0,x,x,x,x,x,0,0,0,x,0,x,
x,0,0
```

- To create the table using Microsoft Access 2002
- To create the table using Microsoft Access 97

**To create the table using Microsoft Access 2002**

1. Rename log files and give them a file name extension of **.txt**.
2. In Access 2002, click **Blank Database**.
3. In **File New Database**, specify a *FileName*, and then click **Create**.
4. Click **File**, click **Get External Data**, and then click **Import**.
5. In **Import**, in **Files of type**, select **Text Files**, locate the IAS log file, select it, and then click **Import**.
6. In the Import Text Wizard, click **Advanced**.
7. In **Import Specification**:
   - In **File Format**, select **Delimited**.
   - In **Field Delimiter**, select the **,** (comma).

- In the **Text Qualifier**, select **"** (quotation mark).
- In the **Dates, Times, and Numbers**, select **Four Digit Years** and **Leading Zeros in Dates**, and then type the appropriate **Date Order** (such as **MDY**), **Date Delimiter** (such as **/** or forward slash), **Time Delimiter** (such as **:** or colon), and **Decimal Symbol** (such as **.** or period).

8. In the **Import Text Wizard** dialog box, click **Next**, select **First Row Contains Field Names**, and then click **Next**.
9. Select **In a New Table**, and then click **Next**.
10. Leave the defaults in **Field Options**, and then click **Next**.
11. Select **Let Access add primary key**, and then click **Next**.
12. In **Import to Table**, type the name of the new table, click **Finish**, and then click **OK**.
13. In the *FileName***:Database** dialog box, select the name of your database, and then click **Open** to view your table.

**To create the table using Microsoft Access 97**

1. Rename log files and give them a file name extension of **.txt**.
2. In Access 97, click **File**, and then click **New Database**.
3. In **New**, click **Blank Database**, and then click **OK**.
4. In **File New Database**, specify a *FileName*, and then click **Create**.
5. In the *FileName***:Database** dialog box, click **New**.
6. In **New Table**, select **Import Table**, and then click **OK**.
7. In **Import**, in **Files of type**, select **Text Files**, locate the IAS log file, select it, and then click **Import**.
8. In the Import Text Wizard, select **Advanced**.
9. In **Advanced**:

    - In **File Format**, select **Delimited**.
    - In **Field Delimiter**, select the **,** (comma).
    - In the **Text Qualifier**, select **"** (quotation mark).
    - In **File Origin**, select **Windows (ANSI)**.
    - In the **Dates, Times, and Numbers**, select **Four Digit Years** and **Leading Zeros in Dates**, and then type the appropriate **Date Order** (such as **MDY**), **Date Delimiter** (such as **/** or forward slash), **Time Delimiter** (such as **:** or colon), and **Decimal Symbol** (such as **.** or period).

10. In **Import Text Wizard**, click **Next**, select **First Row Contains Field Names**, and then click **Next**.
11. Select **In a New Table**, and then click **Next**.

12. Leave the defaults in **Field Options**, and then click **Next**

13. Select **Let Access add primary key**, and then click **Next**.

14. In **Import to Table**, type the name of the new table, click **Finish**, and then click **OK**.

15. In the *FileName***:Database** dialog box, select the name of your database, and then click **Open** to view your table.

# Interpreting database-import log files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interpreting database-import log files**

Unlike IAS-formatted log files, database-import log files present the data in a standard sequence and use a structure that is identical, regardless of the network access server (NAS) that sends the data. This consistent sequence and structure helps simplify accounting and authentication records. Data can be easily exported to a database.

### Note

- Although IAS supports both IAS-formatted and database-import log files, the database-import log format should be used in most instances because it supports the use of ODBC-compliant tools.

**Entries recorded in database-import log files**

The following are example entries (Access-Request and Access-Accept) from a database-import log file.

```
"CLIENTCOMP","IAS",06/17/1999,13:04:33,1,"client",,,,,,,,,9,"10.10.10.10","iasclient",,,,,,,1,,0,,,,,,,,
```

| | | | |
|---|---|---|---|
| | Called-Station-ID | Text | The phone number dialed by the user. |
| | Calling-Station-ID | Text | The phone number from which the call originated. |
| | Callback-Number | Text | The callback phone number. |
| | Framed-IP-Address | Text | The framed address to be configured for the user. |
| | NAS-Identifier | Text | The text that identifies the NAS originating the request. |
| | NAS-IP-Address | Text | The IP address of the NAS originating the request. |
| | NAS-Port | Number | The physical port number of the NAS originating the request. |
| 9 | Client-Vendor | Number | The manufacturer of the NAS (this is an IAS-internal attribute). |
| "10.10.10.10" | Client-IP-Address | Text | The IP address of the RADIUS client (this is an IAS-internal attribute). |
| "iasclient" | Client-Friendly-Name | Text | The friendly name for the RADIUS client (this is an IAS-internal attribute). |
| | Event-Timestamp | Time | The date and time that this event occurred on the NAS. |
| | Port-Limit | Number | The maximum number of ports that the NAS provides to the user. |
| | NAS-Port-Type | Number | The type of physical port that is used by the NAS originating the request. |
| | Connect-Info | Text | Information that is used by the NAS to specify the type of connection made. Typical information includes connection speed and data encoding protocols. |
| | Framed-Protocol | Number | The protocol to be used. |
| | Service-Type | Number | The type of service that the user has requested. |
| 1 | Authentication-Type | Number | The authentication scheme, which is used to verify the user and can be:<br>• 1 = PAP<br>• 2 = CHAP<br>• 3 = MS-CHAP<br>• 4 = MS-CHAP v2<br>• 5 = EAP<br>• 7 = None<br>• 8 =Custom<br><br>This is an IAS-internal attribute. |
| | NP-Policy-Name | Text | The friendly name of the remote access policy that either granted or denied access. This attribute is logged in Access-Accept and Access-Reject messages. If a user is rejected because none of the remote access policies matched, then this attribute is blank. |
| 0 | Reason-Code | Number | The reason for rejecting a user, which can be:<br>• 0 = IAS_SUCCESS<br>• 1 = IAS_INTERNAL_ERROR<br>• 2 = IAS_ACCESS_DENIED<br>• 3 = IAS_MALFORMED_REQUEST |

| | | | |
|---|---|---|---|
| | | | • 4 = IAS_GLOBAL_CATALOG_UNAVAILABLE<br>• 5 = IAS_DOMAIN_UNAVAILABLE<br>• 6 = IAS_SERVER_UNAVAILABLE<br>• 7 = IAS_NO_SUCH_DOMAIN<br>• 8 = IAS_NO_SUCH_USER<br>• 16 = IAS_AUTH_FAILURE<br>• 17 = IAS_CHANGE_PASSWORD_FAILURE<br>• 18 = IAS_UNSUPPORTED_AUTH_TYPE<br>• 32 = IAS_LOCAL_USERS_ONLY<br>• 33 = IAS_PASSWORD_MUST_CHANGE<br>• 34 = IAS_ACCOUNT_DISABLED<br>• 35 = IAS_ACCOUNT_EXPIRED<br>• 36 = IAS_ACCOUNT_LOCKED_OUT<br>• 37 = IAS_INVALID_LOGON_HOURS<br>• 38 = IAS_ACCOUNT_RESTRICTION<br>• 48 = IAS_NO_POLICY_MATCH<br>• 64 = IAS_DIALIN_LOCKED_OUT<br>• 65 = IAS_DIALIN_DISABLED<br>• 66 = IAS_INVALID_AUTH_TYPE<br>• 67 = IAS_INVALID_CALLING_STATION<br>• 68 = IAS_INVALID_DIALIN_HOURS<br>• 69 = IAS_INVALID_CALLED_STATION<br>• 70 = IAS_INVALID_PORT_TYPE<br>• 71 = IAS_INVALID_RESTRICTION<br>• 80 = IAS_NO_RECORD<br>• 96 = IAS_SESSION_TIMEOUT<br>• 97 = IAS_UNEXPECTED_REQUEST<br><br>This is an IAS-internal attribute. |
| | Class | Text | The attribute that is sent to client in an Access-Accept packet. |
| | Session-Timeout | Number | The length of time (in seconds) before the session is terminated. |
| | Idle-Timeout | Number | The length of idle time (in seconds) before the session is terminated. |

| | | | |
|---|---|---|---|
| | Termination-Action | Number | The action that the NAS should take when service is completed. |
| | EAP-Friendly-Name | Text | The friendly name that is used with Extensible Authentication Protocol (EAP). |
| | Acct-Status-Type | Number | The number that specifies whether an accounting packet starts or stops a bridging, routing, or Terminal Server session. |
| | Acct-Delay-Time | Number | The length of time (in seconds) for which the NAS has been sending the same accounting packet. |
| | Acct-Input-Octets | Number | The number of octets received during the session. |
| | Acct-Output-Octets | Number | The number of octets sent during the session. |
| | Acct-Session-Id | Text | The unique numeric string that identifies the server session. |
| | Acct-Authentic | Number | The number that specifies which server authenticated an incoming call. |
| | Acct-Session-Time | Number | The length of time (in seconds) for which the session has been active. |
| | Acct-Input-Packets | Number | The number of packets received during the session. |
| | Acct-Output-Packets | Number | The number of packets sent during the session. |
| | Acct-Terminate-Cause | Number | The reason that a connection was terminated. |
| | Acct-Multi-Ssn-ID | Text | The unique numeric string that identifies the multilink session. |
| | Acct-Link-Count | Number | The number of links in a multilink session. |
| | Acct-Interim-Interval | Number | The length of interval (in seconds) between each interim update that the NAS sends. |
| | Tunnel-Type | Number | The tunneling protocol to be used. |
| | Tunnel-Medium-Type | Number | The medium to use when creating a tunnel for protocols. For example, L2TP packets can be sent over multiple link layers. |
| | Tunnel-Client-Endpt | Text | The IP address of the tunnel client. |
| | Tunnel-Server-Endpt | Text | The IP address of the tunnel server. |
| | Acct-Tunnel-Connection | Text | An identifier assigned to the tunnel. |
| | Tunnel-Pvt-Group-ID | Text | The group ID for a specific tunneled session. |
| | Tunnel-Assignment-ID | Text | The tunnel to which a session is assigned. |
| | Tunnel-Preference | Number | The preference of the tunnel type, as indicated with the Tunnel-Type attribute when multiple tunnel types are supported by the access server. |
| | MS-Acct-Auth-Type | Number | A Routing and Remote Access service attribute. For more information, see RFC 2548. |
| | MS-Acct-EAP-Type | Number | A Routing and Remote Access service attribute. For more information, see RFC 2548. |
| | MS-RAS-Version | Text | A Routing and Remote Access service attribute. For more information, see RFC 2548. |

| | | | |
|---|---|---|---|
| | MS-RAS-Vendor | Number | A Routing and Remote Access service attribute. For more information, see RFC 2548. |
| | MS-CHAP-Error | Text | A Routing and Remote Access service attribute. For more information, see RFC 2548. |
| | MS-CHAP-Domain | Text | A Routing and Remote Access service attribute. For more information, see RFC 2548. |
| | MS-MPPE-Encryption-Types | Number | A Routing and Remote Access service attribute. For more information, see RFC 2548. |
| | MS-MPPE-Encryption-Policy | Number | A Routing and Remote Access service attribute. For more information, see RFC 2548. |

# IAS-Formatted Log Files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IAS-formatted log files**

This section provides information about using IAS-formatted log files:

- Interpreting IAS-formatted log files
- Interpreting IAS IDs for vendor-specific attributes

# Interpreting IAS-formatted log files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interpreting IAS-formatted log files**

In the Windows NT 4 version of IAS, log files are formatted by using a method in which attributes are logged as attribute-value pairs. This formatting is supported in IAS in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000. The logs that use this format are referred to as IAS-formatted log files. However, in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; and Windows 2000, this format supports the inclusion of additional information in the log file:

- In addition to accounting messages (Accounting-On, Accounting-Off, Accounting-Start, Accounting-Stop, and Accounting-Interim), the IAS server also logs authentication messages (Access-Request, Access-Accept, and Access-Reject).
- All string attributes that contain either unprintable characters or delimiters are printed in hexadecimal format (for example, 0x026).
- If IAS receives an attribute (RADIUS-standard or vendor-specific) that is not defined in the IAS dictionary, it is logged as a string.

### Notes

- Unless you have migration, compatibility, or other issues that require you to use the IAS format, you should use the database-import log format. Although a database-import log file contains a smaller subset of attributes, it contains the attributes required to support most tracking and accounting activities.
- To help read and interpret the log files in IAS format, the Iasparse tool is included in the IAS Resources: Windows Deployment and Resource Kits. This tool can be used to parse log files and provide output in a readable format.

**Entries recorded in IAS-formatted log files**

The following is an example entry (Access-Request) from an IAS-formatted log file.

```
10.10.10.10,client,06/04/1999,14:42:19,IAS,CLIENTCOMP,6,2,7,1,5,9,61,5,64,1,65,1,31,1
```

The format of this record, which is the same for all records in your log file, includes a header, followed by the attribute-value pairs for all attributes that are contained in the packet.

The first six record fields make up the header and include:

| Value shown in example | Attribute | ID | Data type | Represents |
|---|---|---|---|---|
| 10.10.10.10 | NAS-IP-Address | IAS Header | Text | The IP address of the NAS that is sending the request. |
| client | User-Name | IAS Header | Text | The user name that is requesting access. |
| 06/04/1999 | Record-Date | IAS Header | Time | The date that the log is written. |
| 14:42:19 | Record-Time | IAS Header | Time | The time that the log is written. |

| | | | | |
|---|---|---|---|---|
| IAS | Service-Name | IAS Header | Text | The name of the service that is running on the RADIUS server. |
| CLIENTCOMP | Computer-Name | IAS Header | Text | The name of the RADIUS server. |

Beyond the header, RADIUS attributes and values are listed in pairs in the following format:

*<AttributeNumber1>,<ValueForAttributeNumber1>,<AttributeNumber2>,<ValueForAttributeNumber2>,*

For example, the two fields after the header contain a 6 and a 2, which can be interpreted as follows:

- The number 6 represents the RADIUS ID for the Service-Type.
- The number 2 represents the attribute value for the Service-Type. The RADIUS protocol specifies the following values for the Service-Type attribute:
    - 1 = Login
    - 2 = Framed
    - 3 = Callback Login
    - 4 = Callback Framed
    - 5 = Outbound
    - 6 = Administrative
    - 7 = NAS Prompt
    - 8 = Authenticate Only
    - 9 = Callback NAS Prompt

The value of this attribute is 2 (Framed).

This attribute-value pair is interpreted as Service-Type = Framed, which indicates to the IAS server to provide a framed protocol for the user–for example, Point-to-Point Protocol (PPP) or SLIP.

The following table describes the RADIUS attributes, listed in numerical order, which can be found in an IAS-formatted log file record. Unlike database-import log files, which use a fixed sequence of attributes, the sequence of the attributes in IAS-formatted log files depends upon the sequence used by the access server. For additional information about the sequence of these records, see the documentation for the access server.

**Notes**

- This table does not cover vendor-specific attributes (VSAs). For information about interpreting VSAs, see Interpreting IAS IDs for vendor-specific attributes. For more information about VSAs that are supported by your access server, see your access server documentation.
- The entries in the **ID** column, which begin with "IAS," are IAS-specific attributes. They are not found in the RADIUS protocol.

| **Attribute** | **ID** | **Data type** | **Represents** |
|---|---|---|---|
| User-Name | 1 | Text | The user identity, as specified by the user. |

| | | | |
|---|---|---|---|
| NAS-IP-Address | 4 | Text | The IP address of the NAS originating the request. |
| NAS-Port | 5 | Number | The physical port number of the NAS originating the request. |
| Service-Type | 6 | Number | The type of service that the user has requested. |
| Framed-Protocol | 7 | Number | The protocol to be used. |
| Framed-IP-Address | 8 | Text | The framed address to be configured for the user. |
| Framed-IP-Netmask | 9 | Text | The IP netmask to be configured for the user. |
| Framed-Routing | 10 | Number | The Routing method to be used by the user. |
| Filter-ID | 11 | Text | The name of the filter list for the user requesting authentication. |
| Framed-MTU | 12 | Number | The maximum transmission unit to be configured for the user. |
| Framed-Compression | 13 | Number | The compression protocol to be used. |
| Login-IP-Host | 14 | Number | The IP address of the host to which the user should be connected. |
| Login-Service | 15 | Number | The service that connects the user to the login host. |
| Login-TCP-Port | 16 | Number | The TCP port to which the user should be connected. |
| Reply-Message | 18 | Text | The message displayed to the user when an authentication request is accepted. |
| Callback-Number | 19 | Text | The callback phone number. |
| Callback-ID | 20 | Text | The name of a location to be called by the access server when performing callback. |
| Framed-Route | 22 | Text | The routing information that is configured on the access client. |
| Framed-IPX-Network | 23 | Number | The IPX network number to be configured on the NAS for the user. |

| | | | |
|---|---|---|---|
| Class | 25 | Text | The attribute sent to the client in an Access-Accept packet, which is useful for correlating Accounting-Request packets with authentication sessions. The format is:<br>• **Type** Contains the value 25 (1 octet).<br>• **Length** Contains a value of 20 or greater (1 octet).<br>• **Checksum** Contains an Adler-32 checksum that is computed over the remainder of the Class attribute (4 octets).<br>• **Vendor-ID** Contains the ID of the access server vendor (4 octets). The high-order octet is 0 and the low-order 3 octets are the SMI Network Management Private Enterprise Code of the vendor in network byte order, as defined in RFC 1007 "Vendor SMI Network Management Private Enterprise Codes".<br>• **Version** Contains the value of 1 (2 octets).<br>• **server-Address** Contains the IP address of the RADIUS server that issued the Access-Challenge. For multihomed servers, this is the address of the network interface that received the original Access-Request (2 octets).<br>• **Service-Reboot-Time** Specifies the time at which the first serial number was returned (8 octets).<br>• **Unique-Serial-Number** Contains a unique number to distinguish an individual connection attempt (8 octets).<br>• **String** Contains information that is used to classify accounting records for additional analysis (0 or more octets). In IAS, the Class attribute from the profile is copied into the String field.<br>The Class attribute is used to match the accounting and authentication records if the Class attribute is sent by the network access server in the accounting request packets. The combination of Serial-Number, Service-Reboot-Time, and server-Address must be a unique identification for each authentication that the server accepts. |
| Vendor-Specific | 26 | Text | The attribute that is used to support proprietary NAS features. |
| Session-Timeout | 27 | Number | The length of time (in seconds) before a session is terminated. |
| Idle-Timeout | 28 | Number | The length of idle time (in seconds) before a session is terminated. |
| Termination-Action | 29 | Number | The action that the NAS should take when service is completed. |
| Called-Station-ID | 30 | Text | The phone number that is dialed by the user. |
| Calling-Station-ID | 31 | Text | The phone number from which the call originated. |
| NAS-Identifier | 32 | Text | The string that identifies the NAS originating the request. |
| Login-LAT-Service | 34 | Text | The host with which the user is to be connected by LAT. |

| | | | |
|---|---|---|---|
| Login-LAT-Node | 35 | Text | The node with which the user is to be connected by LAT. |
| Login-LAT-Group | 36 | Text | The LAT group codes for which the user is authorized. |
| Framed-AppleTalk-Link | 37 | Number | The AppleTalk network number for the serial link to the user (this is used only when the user is a router). |
| Framed-AppleTalk-Network | 38 | Number | The AppleTalk network number that the NAS must query for existence in order to allocate the user's AppleTalk node. |
| Framed-AppleTalk-Zone | 39 | Text | The AppleTalk default zone for the user. |
| Acct-Status-Type | 40 | Number | The number that specifies whether an accounting packet starts or stops a bridging, routing, or Terminal server session. |
| Acct-Delay-Time | 41 | Number | The length of time (in seconds) for which the NAS has been sending the same accounting packet. |
| Acct-Input-Octets | 42 | Number | The number of octets received during the session. |
| Acct-Output-Octets | 43 | Number | The number of octets sent during the session. |
| Acct-Session-ID | 44 | Text | The unique numeric string that identifies the server session. |
| Acct-Authentic | 45 | Number | The number that specifies which server has authenticated an incoming call. |
| Acct-Session-Time | 46 | Number | The length of time (in seconds) for which the session has been active. |
| Acct-Input-Packets | 47 | Number | The number of packets received during the session. |
| Acct-Output-Packets | 48 | Number | The number of packets sent during the session. |
| Acct-Terminate-Cause | 49 | Number | The reason that a connection was terminated. |
| Acct-Multi-SSN-ID | 50 | Text | The unique numeric string that identifies the multilink session. |
| Acct-Link-Count | 51 | Number | The number of links in a multilink session. |
| Event-Timestamp | 55 | Time | The date and time that this event occurred on the NAS. |
| NAS-Port-Type | 61 | Number | The type of physical port that is used by the NAS originating the request. |
| Port-Limit | 62 | Number | The maximum number of ports that the NAS provides to the user. |
| Login-LAT-Port | 63 | Number | The port with which the user is connected by Local Area Transport (LAT). |
| Tunnel-Type | 64 | Number | The tunneling protocols to be used. |

| | | | |
|---|---|---|---|
| Tunnel-Medium-Type | 65 | Number | The transport medium to use when creating a tunnel for protocols. For example, L2TP packets can be sent over multiple link layers. |
| Tunnel-Client-Endpt | 66 | Text | The IP address of the tunnel client. |
| Tunnel-server-Endpt | 67 | Text | The IP address of the tunnel server. |
| Acct-Tunnel-Connection | 68 | Text | An identifier assigned to the tunnel. |
| Password-Retry | 75 | Number | The number of times a user can try to be authenticated before the NAS terminates the connection. |
| Prompt | 76 | Number | A number that indicates to the NAS whether or not it should (Prompt=1) or should not (Prompt=0) echo the user's response as it is typed. |
| Connect-Info | 77 | Text | Information that is used by the NAS to specify the type of connection made. Typical information includes connection speed and data encoding protocols. |
| Configuration-Token | 78 | Text | The type of user profile to be used (sent from a RADIUS proxy server to a RADIUS proxy client) in an Access-Accept packet. |
| Tunnel-Pvt-Group-ID | 81 | Text | The group ID for a particular tunneled session. |
| Tunnel-Assignment-ID | 82 | Text | The tunnel to which a session is to be assigned. |
| Tunnel-Preference | 83 | Number | A number that indicates the preference of the tunnel type, as indicated with the Tunnel-Type attribute when multiple tunnel types are supported by the access server. |
| Acct-Interim-Interval | 85 | Number | The length of interval (in seconds) between each interim update sent by the NAS. |
| Ascend | 107 to 255 | Text | The vendor-specific attributes for Ascend. For more information, see the Ascend documentation. |
| Client-IP-Address | IAS 4108 | Text | The IP address of the RADIUS client. |
| NAS-Manufacturer | IAS 4116 | Number | The manufacturer of the NAS. |
| MS-CHAP-Error | IAS 4121 | Number | The error data that describes an MS-CHAP transaction. |
| Authentication-Type | IAS 4127 | Number | The authentication scheme that is used to verify the user. |
| Client-Friendly-Name | IAS 4128 | Text | The friendly name for the RADIUS client. |
| SAM-Account-Name | IAS 4129 | Text | The user account name in the Security Accounts Manager (SAM) database. |
| Fully-Qualified-User-Name | IAS 4130 | Text | The user name in canonical format. |
| EAP-Friendly-Name | IAS 4132 | Text | The friendly name that is used with Extensible Authentication Protocol (EAP). |

| | | | |
|---|---|---|---|
| Packet-Type | IAS 4136 | Number | The type of packet, which can be:<br>• 1=Accept-Request<br>• 2=Access-Accept<br>• 3=Access-Reject<br>• 4=Accounting-Request |
| Reason Code | IAS 4142 | Number | The reason for rejecting a user, which can be:<br>• 00 = Success<br>• 01 = Internal error<br>• 02 = Access denied<br>• 03 = Malformed request<br>• 04 = Global catalog unavailable<br>• 05 = Domain unavailable<br>• 06 = server unavailable<br>• 07 = No such domain<br>• 08 = No such user<br>• 16 = Authentication failure<br>• 17 = Password change failure<br>• 18 = Unsupported authentication type<br>• 19 = No reversibly encrypted password is stored for the user account<br>• 32 = Local users only<br>• 33 = Password must be changed<br>• 34 = Account disabled<br>• 35 = Account expired<br>• 36 = Account locked out<br>• 37 = Invalid logon hours<br>• 38 = Account restriction<br>• 48 = Did not match remote access policy<br>• 49 = Did not match connection request policy |

| | | | |
|---|---|---|---|
| | | | • 64 = Dial-in locked out<br>• 65 = Dial-in disabled<br>• 66 = Invalid authentication type<br>• 67 = Invalid calling station<br>• 68 = Invalid dial-in hours<br>• 69 = Invalid called station<br>• 70 = Invalid port type<br>• 71 = Invalid restriction<br>• 80 = No record<br>• 96 = Session timed out<br>• 97 = Unexpected request |
| NP-Policy-Name | IAS 4149 | Text | The friendly name of a remote access policy. |

**Attributes that are not recorded in IAS-formatted log files**

Although most of the attributes that are sent by an access server and supported in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition are logged when you specify the use of IAS-formatted log files, some attributes are not logged because they contain sensitive information (for example, passwords) that should not be included. The following table lists some of the attributes that are not logged.

| **Attribute name** | **ID/Description** |
|---|---|
| User-Password | 2 |
| CHAP-Password | 3 |
| State | 24 |
| Proxy-State | 33 |
| CHAP-Challenge | 60 |
| Tunnel-Password | 69 |
| EAP-Message | 79 |
| Signature | 80 |
| MS-CHAP-Challenge | Microsoft vendor-specific attribute |
| MS-CHAP-Response | Microsoft vendor-specific attribute |
| MS-CHAP-CPW-1 | Microsoft vendor-specific attribute |
| MS-CHAP-CPW-2 | Microsoft vendor-specific attribute |
| MS-CHAP-LM-Enc-PW | Microsoft vendor-specific attribute |
| MS-CHAP-NT-Enc-PW | Microsoft vendor-specific attribute |

| | |
|---|---|
| MS-CHAP-MPPE-Keys | Microsoft vendor-specific attribute |
| MS-MPPE-Send-Key | Microsoft vendor-specific attribute |
| MS-MPPE-Recv-Key | Microsoft vendor-specific attribute |
| MS-Filter | Microsoft vendor-specific attribute |
| MS-CHAP2-Response | Microsoft vendor-specific attribute |
| MS-CHAP2-Success | Microsoft vendor-specific attribute |
| MS-CHAP2-CPW | Microsoft vendor-specific attribute |

# Interpreting IAS IDs for vendor-specific attributes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Interpreting IAS IDs for vendor-specific attributes**
When IAS logs any of the following vendor-specific attributes, it assigns a unique ID to the attribute in order to facilitate interpretation of the log files.

**Note**

- The ID (number) that is assigned by IAS to a vendor-specific attribute might not be the same as the ID that is assigned by the vendor. If you are using the IAS-formatted log file format and want to determine the ID used by the vendor, see your network access server (NAS) documentation. If you are using database-import log files, the ID (number) that is assigned to the attribute does not appear in the log. Use the following table to determine the data type for the attributes listed.

| ID | Attribute name | Vendor | Data type |
|---|---|---|---|
| 107 | Ascend-Calling-Subaddress | Ascend | Text |
| 108 | Ascend-Callback-Delay | Ascend | Number |
| 109 | Ascend-Endpoint-Disc | Ascend | Text |
| 110 | Ascend-Remote-FW | Ascend | Text |
| 111 | Ascend-Multicast-G-Leave-Delay | Ascend | Number |
| 112 | Ascend-CBCP-Enable | Ascend | Number |
| 113 | Ascend-CBCP-Mode | Ascend | Number |
| 114 | Ascend-CBCP-Delay | Ascend | Number |
| 115 | Ascend-CBCP-Trunk-Group | Ascend | Number |
| 116 | Ascend-Appletalk-Route | Ascend | Text |
| 117 | Ascend-Appletalk-Peer-Mode | Ascend | Number |
| 118 | Ascend-Route-Appletalk | Ascend | Number |
| 119 | Ascend-FCP-Parameter | Ascend | Text |
| 120 | Ascend-Modem-Port-No | Ascend | Number |
| 121 | Ascend-Modem-Slot-No | Ascend | Number |
| 122 | Ascend-Modem-Shelf-No | Ascend | Number |
| 123 | Ascend-CallAttempt-Limit | Ascend | Number |
| 124 | Ascend-CallBlock-Duration | Ascend | Number |
| 125 | Ascend-Maximum-Call-Duration | Ascend | Number |
| 126 | Ascend-Route-Preference | Ascend | Number |
| 127 | Ascend-Tunneling-Protocol | Ascend | Number |
| 128 | Ascend-Shared-Profile-Enable | Ascend | Number |

| | | | |
|---|---|---|---|
| 129 | Ascend-Primary-Home-Agent | Ascend | Text |
| 130 | Ascend-Secondary-Home-Agent | Ascend | Text |
| 131 | Ascend-Dialout-Allowed | Ascend | Number |
| 132 | Ascend-Client-Gateway | Ascend | Number |
| 133 | Ascend-BACP-Enable | Ascend | Number |
| 134 | Ascend-DHCP-Maximum-Leases | Ascend | Number |
| 135 | Ascend-Client-Primary-DNS | Ascend | String |
| 136 | Ascend-Client-Secondary-DNS | Ascend | String |
| 137 | Ascend-Client-Assign-DNS | Ascend | Number |
| 138 | Ascend-User-Acct-Type | Ascend | Number |
| 139 | Ascend-User-Acct-Host | Ascend | Number |
| 140 | Ascend-User-Acct-Port | Ascend | Number |
| 141 | Ascend-User-Acct-Key | Ascend | Text |
| 142 | Ascend-User-Acct-Base | Ascend | Number |
| 143 | Ascend-User-Acct-Time | Ascend | Number |
| 144 | Ascend-Assign-IP-Client | Ascend | Number |
| 145 | Ascend-Assign-IP-Server | Ascend | Number |
| 146 | Ascend-Assign-IP-Global-Pool | Ascend | Text |
| 147 | Ascend-DHCP-Reply | Ascend | Number |
| 148 | Ascend-DHCP-Pool-Number | Ascend | Number |
| 149 | Ascend-Expect-Callback | Ascend | Number |
| 150 | Ascend-Event-Type | Ascend | Number |
| 151 | Ascend-Session-Svr-Key | Ascend | Text |
| 152 | Ascend-Multicast-Rate-Limit | Ascend | Number |
| 153 | Ascend-IF-Netmask | Ascend | Number |
| 154 | Ascend-Remote-Addr | Ascend | Number |
| 155 | Ascend-Multicast-Client | Ascend | Number |
| 156 | Ascend-FR-Circuit-Name | Ascend | Text |
| 157 | Ascend-FR-Link-Up | Ascend | Number |
| 158 | Ascend-FR-Nailed-Grp | Ascend | Number |
| 159 | Ascend-FR-Type | Ascend | Number |
| 160 | Ascend-FR-Link-Mgt | Ascend | Number |

| | | | |
|---|---|---|---|
| 161 | Ascend-FR-N391 | Ascend | Number |
| 162 | Ascend-FR-DCE-N392 | Ascend | Number |
| 163 | Ascend-FR-DTE-N392 | Ascend | Number |
| 164 | Ascend-FR-DCE-N393 | Ascend | Number |
| 165 | Ascend-FR-DTE-N393 | Ascend | Number |
| 166 | Ascend-FR-T391 | Ascend | Number |
| 167 | Ascend-FR-T392 | Ascend | Number |
| 168 | Ascend-Bridge-Address | Ascend | Text |
| 169 | Ascend-TS-Idle-Limit | Ascend | Number |
| 170 | Ascend-TS-Idle-Mode | Ascend | Number |
| 171 | Ascend-DBA-Monitor | Ascend | Number |
| 172 | Ascend-Base-Channel-Count | Ascend | Number |
| 173 | Ascend-Minimum-Channels | Ascend | Number |
| 174 | Ascend-IPX-Route | Ascend | Text |
| 175 | Ascend-FT1-Caller | Ascend | Number |
| 176 | Ascend-Backup | Ascend | Text |
| 177 | Ascend-Call-Type | Ascend | Number |
| 178 | Ascend-Group | Ascend | Text |
| 179 | Ascend-FR-DLCI | Ascend | Number |
| 180 | Ascend-FR-Profile-Name | Ascend | Text |
| 181 | Ascend-Ara-PW | Ascend | Text |
| 182 | Ascend-IPX-Node-Addr | Ascend | Text |
| 183 | Ascend-Home-Agent-IP-Addr | Ascend | Text |
| 184 | Ascend-Home-Agent-Password | Ascend | Text |
| 185 | Ascend-Home-Network-Name | Ascend | Text |
| 186 | Ascend-Home-Agent-UDP-Port | Ascend | Number |
| 187 | Ascend-Multilink-ID | Ascend | Number |
| 188 | Ascend-Num-In-Multilink | Ascend | Number |
| 189 | Ascend-First-Dest | Ascend | Number |
| 190 | Ascend-Pre-Input-Octets | Ascend | Number |
| 191 | Ascend-Pre-Output-Octets | Ascend | Number |
| 192 | Ascend-Pre-Input-Packets | Ascend | Number |
| 193 | Ascend-Pre-Output-Packets | Ascend | Number |

| | | | |
|---|---|---|---|
| 194 | Ascend-Maximum-Time | Ascend | Number |
| 195 | Ascend-Disconnect-Cause | Ascend | Number |
| 196 | Ascend-Connect-Progress | Ascend | Number |
| 197 | Ascend-Data-Rate | Ascend | Number |
| 198 | Ascend-Pre-Session-Time | Ascend | Number |
| 199 | Ascend-Token-Idle | Ascend | Number |
| 200 | Ascend-Token-Immediate | Ascend | Number |
| 201 | Ascend-Require-Auth | Ascend | Number |
| 202 | Ascend-Number-Sessions | Ascend | Text |
| 203 | Ascend-Authen-Alias | Ascend | Text |
| 204 | Ascend-Token-Expiry | Ascend | Number |
| 205 | Ascend-Menu-Selector | Ascend | Text |
| 206 | Ascend-Menu-Item | Ascend | Text |
| 207 | Ascend-PW-Warntime | Ascend | Number |
| 208 | Ascend-PW-Lifetime | Ascend | Number |
| 209 | Ascend-IP-Direct | Ascend | Text |
| 210 | Ascend-PPP-VJ-Slot-Comp | Ascend | Number |
| 211 | Ascend-PPP-VJ-1172 | Ascend | Number |
| 212 | Ascend-PPP-Async-Map | Ascend | Number |
| 213 | Ascend-Third-Prompt | Ascend | Text |
| 214 | Ascend-Send-Secret | Ascend | Text |
| 215 | Ascend-Receive-Secret | Ascend | Text |
| 216 | Ascend-IPX-PeerMode | Ascend | Number |
| 217 | Ascend-IP-Pool-Definition | Ascend | Text |
| 218 | Ascend-Assign-IP-Pool | Ascend | Number |
| 219 | Ascend-FR-Direct | Ascend | Number |
| 220 | Ascend-FR-Direct-Profile | Ascend | Text |
| 221 | Ascend-FR-Direct-DLCI | Ascend | Number |
| 222 | Ascend-Handle-IPX | Ascend | Number |
| 223 | Ascend-Netware-Timeout | Ascend | Number |
| 224 | Ascend-IPX-Alias | Ascend | Number |
| 225 | Ascend-Metric | Ascend | Number |
| 226 | Ascend-PRI-Number-Type | Ascend | Number |

| | | | |
|---|---|---|---|
| 227 | Ascend-Dial-Number | Ascend | Text |
| 228 | Ascend-Route-IP | Ascend | Number |
| 229 | Ascend-Route-IPX | Ascend | Number |
| 230 | Ascend-Bridge | Ascend | Number |
| 231 | Ascend-Send-Auth | Ascend | Number |
| 232 | Ascend-Send-Passwd | Ascend | Text |
| 233 | Ascend-Link-Compression | Ascend | Number |
| 234 | Ascend-Target-Util | Ascend | Number |
| 235 | Ascend-Maximum-Channels | Ascend | Number |
| 236 | Ascend-Inc-Channel-Count | Ascend | Number |
| 237 | Ascend-Dec-Channel-Count | Ascend | Number |
| 238 | Ascend-Seconds-Of-History | Ascend | Number |
| 239 | Ascend-History-Weigh-Type | Ascend | Number |
| 240 | Ascend-Add-Seconds | Ascend | Number |
| 241 | Ascend-Remove-Seconds | Ascend | Number |
| 242 | Ascend-Data-Filter | Ascend | String |
| 243 | Ascend-Call-Filter | Ascend | String |
| 244 | Ascend-Idle-Limit | Ascend | Number |
| 245 | Ascend-Preempt-Limit | Ascend | Number |
| 246 | Ascend-Callback | Ascend | Number |
| 247 | Ascend-Data-Svc | Ascend | Number |
| 248 | Ascend-Force56 | Ascend | Number |
| 249 | Ascend-Billing-Number | Ascend | Text |
| 250 | Ascend-Call-By-Call | Ascend | Number |
| 251 | Ascend-Transit-Number | Ascend | Text |
| 252 | Ascend-Host-Info | Ascend | Text |
| 253 | Ascend-PPP-Address | Ascend | Number |
| 254 | Ascend-MPP-Idle-Percent | Ascend | Number |
| 255 | Ascend-Xmit-Rate | Ascend | Number |
| 4120 | MS-CHAP-Domain | Microsoft | String |
| 4121 | MS-CHAP-Error | Microsoft | String |
| 4134 | MS-Acct-Auth-Type | Microsoft | Number |
| 4135 | MS-Acct-EAP-Type | Microsoft | Number |

| | | | |
|---|---|---|---|
| 4143 | MS-Filter | Microsoft | String |
| 4147 | MS-RAS-Vendor | Microsoft | Number |
| 4148 | MS-RAS-Version | Microsoft | String |
| 4150 | MS-Primary-DNS-Server | Microsoft | String |
| 4151 | MS-Secondary-DNS-Server | Microsoft | String |
| 4152 | MS-Primary-NBNS-Server | Microsoft | String |
| 4153 | MS-Secondary-NBNS-Server | Microsoft | String |
| 5000 | Cisco-AV-Pair | Cisco | Text |
| 11000 | USR-Last-Number-Dialed-Out | US Robotics | Text |
| 11001 | USR-Last-Number-Dialed-In-DNIS | US Robotics | Text |
| 11002 | USR-Last-Callers-Number-ANI | US Robotics | Text |
| 11003 | USR-Channel | US Robotics | Number |
| 11004 | USR-Event-ID | US Robotics | Number |
| 11005 | USR-Event-Date-Time | US Robotics | Number |
| 11006 | USR-Call-Start-Date-Time | US Robotics | Number |
| 11007 | USR-Call-End-Date-Time | US Robotics | Number |
| 11008 | USR-Default-DTE-Data-Rate | US Robotics | Number |
| 11009 | USR-Initial-Rx-Link-Data-Rate | US Robotics | Number |
| 11010 | USR-Final-Rx-Link-Data-Rate | US Robotics | Number |
| 11011 | USR-Initial-Tx-Link-Data-Rate | US Robotics | Number |
| 11012 | USR-Final-Tx-Link-Data-Rate | US Robotics | Number |
| 11013 | USR-Chassis-Temperature | US Robotics | Number |
| 11014 | USR-Chassis-Temp-Threshold | US Robotics | Number |
| 11015 | USR-Actual-Voltage | US Robotics | Number |
| 11016 | USR-Expected-Voltage | US Robotics | Number |
| 11017 | USR-Power-Supply-Number | US Robotics | Number |
| 11018 | USR-Card-Type | US Robotics | Number |
| 11019 | USR-Chassis-Slot | US Robotics | Number |
| 11020 | USR-Sync-Async-Mode | US Robotics | Number |
| 11021 | USR-Originate-Answer-Mode | US Robotics | Number |
| 11022 | USR-Modulation-Type | US Robotics | Number |
| 11023 | USR-Initial-Modulation-Type | US Robotics | Number |
| 11024 | USR-Connect-Term-Reason | US Robotics | Number |

| | | | |
|---|---|---|---|
| 11025 | USR-Failure-to-Connect-Reason | US Robotics | Number |
| 11026 | USR-Equalization-Type | US Robotics | Number |
| 11027 | USR-Fallback-Enabled | US Robotics | Number |
| 11028 | USR-Connect-Time-Limit | US Robotics | Number |
| 11029 | USR-Number-of-Rings-Limit | US Robotics | Number |
| 11030 | USR-DTE-Data-Idle-Timout | US Robotics | Number |
| 11031 | USR-Characters-Sent | US Robotics | Number |
| 11032 | USR-Characters-Received | US Robotics | Number |
| 11033 | USR-Blocks-Sent | US Robotics | Number |
| 11034 | USR-Blocks-Received | US Robotics | Number |
| 11035 | USR-Blocks-Resent | US Robotics | Number |
| 11036 | USR-Retrains-Requested | US Robotics | Number |
| 11037 | USR-Retrains-Granted | US Robotics | Number |
| 11038 | USR-Line-Reversals | US Robotics | Number |
| 11039 | USR-Number-Of-Characters-Lost | US Robotics | Number |
| 11040 | USR-Number-of-Blers | US Robotics | Number |
| 11041 | USR-Number-of-Link-Timeouts | US Robotics | Number |
| 11042 | USR-Number-of-Fallbacks | US Robotics | Number |
| 11043 | USR-Number-of-Upshifts | US Robotics | Number |
| 11044 | USR-Number-of-Link-NAKs | US Robotics | Number |
| 11045 | USR-DTR-False-Timeout | US Robotics | Number |
| 11046 | USR-Fallback-Limit | US Robotics | Number |
| 11047 | USR-Block-Error-Count-Limit | US Robotics | Number |
| 11048 | USR-DTR-True-Timeout | US Robotics | Number |
| 11049 | USR-Security-Login-Limit | US Robotics | Number |
| 11050 | USR-Security-Resp-Limit | US Robotics | Number |
| 11051 | USR-DTE-Ring-No-Answer-Limit | US Robotics | Number |
| 11052 | USR-Back-Channel-Data-Rate | US Robotics | Number |
| 11053 | USR-Simplified-MNP-Levels | US Robotics | Number |
| 11054 | USR-Simplified-V42bis-Usage | US Robotics | Number |
| 11055 | USR-Mbi-Ct-PRI-Card-Slot | US Robotics | Number |
| 11056 | USR-Mbi-Ct-TDM-Time-Slot | US Robotics | Number |
| 11057 | USR-Mbi-Ct-PRI-Card-Span-Line | US Robotics | Number |

| | | | |
|---|---|---|---|
| 11058 | USR-Mbi-Ct-BChannel-Used | US Robotics | Number |
| 11059 | USR-Physical-State | US Robotics | Number |
| 11060 | USR-Packet-Bus-Session | US Robotics | Number |
| 11061 | USR-Server-Time | US Robotics | Number |
| 11062 | USR-Channel-Connected-To | US Robotics | Number |
| 11063 | USR-Slot-Connected-To | US Robotics | Number |
| 11064 | USR-Device-Connected-To | US Robotics | Number |
| 11065 | USR-NFAS-ID | US Robotics | Number |
| 11066 | USR-Q931-Call-Reference-Value | US Robotics | Number |
| 11067 | USR-Call-Event-Code | US Robotics | Number |
| 11068 | USR-DS0 | US Robotics | Number |
| 11069 | USR-DS0s | US Robotics | Text |
| 11070 | USR-Gateway-IP-Address | US Robotics | String |
| 11071 | USR-Call-Arrival-in-GMT | US Robotics | Number |
| 11072 | USR-Call-Connect-in-GMT | US Robotics | Number |
| 11073 | USR-Call-Terminate-in-GMT | US Robotics | Number |
| 11074 | USR-IDS0-Call-Type | US Robotics | Number |
| 11075 | USR-Call-Reference-Number | US Robotics | Number |
| 11076 | USR-CDMA-Call-Reference-Number | US Robotics | Number |
| 11077 | USR-Mobile-IP-Address | US Robotics | String |
| 11078 | USR-IWF-IP-Address | US Robotics | String |
| 11079 | USR-Called-Party-Number | US Robotics | Text |
| 11080 | USR-Calling-Party-Number | US Robotics | Text |
| 11081 | USR-Call-Type | US Robotics | Number |
| 11082 | USR-ESN | US Robotics | Text |
| 11083 | USR-IWF-Call-Identifier | US Robotics | Number |
| 11084 | USR-IMSI | US Robotics | Text |
| 11085 | USR-Service-Option | US Robotics | Number |
| 11086 | USR-Disconnect-Cause-Indicator | US Robotics | Number |
| 11087 | USR-Mobile-NumBytes-Txed | US Robotics | Number |
| 11088 | USR-Mobile-NumBytes-Rxed | US Robotics | Number |
| 11089 | USR-Num-Fax-Pages-Processed | US Robotics | Number |
| 11090 | USR-Compression-Type | US Robotics | Number |

| | | | |
|---|---|---|---|
| 11091 | USR-Call-Error-Code | US Robotics | Number |
| 11092 | USR-Modem-Setup-Time | US Robotics | Number |
| 11093 | USR-Call-Connecting-Time | US Robotics | Number |
| 11094 | USR-Connect-Time | US Robotics | Number |
| 11095 | USR-RMMIE-Manufacutere-ID | US Robotics | Number |
| 11096 | USR-RMMIE-Product-Code | US Robotics | Text |
| 11097 | USR-RMMIE-Serial-Number | US Robotics | Text |
| 11098 | USR-RMMIE-Firmware-Version | US Robotics | Text |
| 11099 | USR-RMMIE-Firmware-Build-Date | US Robotics | Text |
| 11100 | USR-RMMIE-Status | US Robotics | Number |
| 11101 | USR-RMMIE-Num-Of-Updates | US Robotics | Number |
| 11102 | USR-RMMIE-x2-Status | US Robotics | Number |
| 11103 | USR-RMMIE-Planned-Disconnect | US Robotics | Number |
| 11104 | USR-RMMIE-Last-Update-Time | US Robotics | Number |
| 11105 | USR-RMMIE-Last-Update-Event | US Robotics | Number |
| 11106 | USR-RMMIE-Rcv-Tot-PwrLvl | US Robotics | Number |
| 11107 | USR-RMMIE-Rcv-PwrLvl-3300Hz | US Robotics | Number |
| 11108 | USR-RMMIE-Rcv-PwrLvl-3750Hz | US Robotics | Number |
| 11109 | USR-RMMIE-PwrLvl-NearEcho-Canc | US Robotics | Number |
| 11110 | USR-RMMIE-PwrLvl-FarEcho-Canc | US Robotics | Number |
| 11111 | USR-RMMIE-PwrLvl-Noise-Lvl | US Robotics | Number |
| 11112 | USR-RMMIE-PwrLvl-Xmit-Lvl | US Robotics | Number |
| 11113 | USR-PW-IFilter-IP | US Robotics | Text |
| 11114 | USR-PW-IFilter-IPX | US Robotics | Text |
| 11115 | USR-PW-IFilter-SAP | US Robotics | Text |
| 11116 | USR-PW-OFilter-IP | US Robotics | Text |
| 11117 | USR-PW-OFilter-IPX | US Robotics | Text |
| 11118 | USR-PW-OFilter-SAP | US Robotics | Text |
| 11119 | USR-PW-VPN-ID | US Robotics | Number |
| 11120 | USR-PW-VPN-Name | US Robotics | Text |
| 11121 | USR-PW-VPN-Neighbor | US Robotics | String |
| 11122 | USR-PW-Framed-Routing-V2 | US Robotics | Number |
| 11123 | USR-PW-VPN-Gateway | US Robotics | Text |

| | | | |
|---|---|---|---|
| 11124 | USR-PW-Tunnel-Authentication | US Robotics | Text |
| 11125 | USR-PW-Index | US Robotics | Number |
| 11126 | USR-PW-Cutoff | US Robotics | Text |
| 11127 | USR-PW-Packet | US Robotics | Text |
| 11128 | USR-Primary-DNS-Server | US Robotics | String |
| 11129 | USR-Secondary-DNS-Server | US Robotics | String |
| 11130 | USR-Primary-NBNS-Server | US Robotics | String |
| 11131 | USR-Secondary-NBNS-Server | US Robotics | String |
| 11132 | USR-Syslog-Tap | US Robotics | Number |
| 11133 | USR-Log-Filter-Packet | US Robotics | Number |
| 11134 | USR-Chassis-Call-Slot | US Robotics | Number |
| 11135 | USR-Chassis-Call-Span | US Robotics | Number |
| 11136 | USR-Chassis-Call-Channel | US Robotics | Number |
| 11137 | USR-Keypress-Timeout | US Robotics | Number |
| 11138 | USR-Unauthenticated-Time | US Robotics | Number |
| 11139 | USR-VPN-Encryptor | US Robotics | Text |
| 11140 | USR-VPN-GW-Location-ID | US Robotics | Text |
| 11141 | USR-Re-Chap-Timeout | US Robotics | Number |
| 11142 | USR-CCP-Algorithm | US Robotics | Number |
| 11143 | USR-ACCM-Type | US Robotics | Number |
| 11144 | USR-Connect-Speed | US Robotics | Number |
| 11145 | USR-Framed-IP-Address-Pool-Name | US Robotics | Text |
| 11146 | USR-MP-EDO | US Robotics | String |
| 11147 | USR-Local-Framed-IP-Addr | US Robotics | String |
| 11148 | USR-Framed-IPX-Route | US Robotics | String |
| 11149 | USR-MPIP-Tunnel-Originator | US Robotics | String |
| 11150 | USR-Bearer-Capabilities | US Robotics | Number |
| 11151 | USR-Speed-Of-Connection | US Robotics | Number |
| 11152 | USR-Max-Channels | US Robotics | Number |
| 11153 | USR-Channel-Expansion | US Robotics | Number |
| 11154 | USR-Channel-Decrement | US Robotics | Number |
| 11155 | USR-Expansion-Algorithm | US Robotics | Number |
| 11156 | USR-Compression-Algorithm | US Robotics | Number |

| | | | |
|---|---|---|---|
| 11157 | USR-Receive-Acc-Map | US Robotics | Number |
| 11158 | USR-Transmit-Acc-Map | US Robotics | Number |
| 11159 | USR-Compression-Reset-Mode | US Robotics | Number |
| 11160 | USR-Min-Compression-Size | US Robotics | Number |
| 11161 | USR-IP | US Robotics | Number |
| 11162 | USR-IPX | US Robotics | Number |
| 11163 | USR-Filter-Zones | US Robotics | Number |
| 11164 | USR-Appletalk | US Robotics | Number |
| 11165 | USR-Bridging | US Robotics | Number |
| 11166 | USR-Spoofing | US Robotics | Number |
| 11167 | USR-Host-Type | US Robotics | Number |
| 11168 | USR-Send-Name | US Robotics | Text |
| 11169 | USR-Send-Password | US Robotics | Text |
| 11170 | USR-Start-Time | US Robotics | Number |
| 11171 | USR-End-Time | US Robotics | Number |
| 11172 | USR-Send-Script1 | US Robotics | Text |
| 11173 | USR-Reply-Script1 | US Robotics | Text |
| 11174 | USR-Send-Script2 | US Robotics | Text |
| 11175 | USR-Reply-Script2 | US Robotics | Text |
| 11176 | USR-Send-Script3 | US Robotics | Text |
| 11177 | USR-Reply-Script3 | US Robotics | Text |
| 11178 | USR-Send-Script4 | US Robotics | Text |
| 11179 | USR-Reply-Script4 | US Robotics | Text |
| 11180 | USR-Send-Script5 | US Robotics | Text |
| 11181 | USR-Reply-Script5 | US Robotics | Text |
| 11182 | USR-Send-Script6 | US Robotics | Text |
| 11183 | USR-Reply-Script6 | US Robotics | Text |
| 11184 | USR-Terminal-Type | US Robotics | Text |
| 11185 | USR-Appletalk-Network-Range | US Robotics | Number |
| 11186 | USR-Local-IP-Address | US Robotics | Text |
| 11187 | USR-Routing-Protocol | US Robotics | Number |
| 11188 | USR-Modem-Group | US Robotics | Text |
| 11189 | USR-IPX-Routing | US Robotics | Number |

| | | | |
|---|---|---|---|
| 11190 | USR-IPX-WAN | US Robotics | Number |
| 11191 | USR-IP-RIP-Policies | US Robotics | Number |
| 11192 | USR-IP-RIP-Simple-Auth-Password | US Robotics | Text |
| 11193 | USR-IP-RIP-Input-Filter | US Robotics | Text |
| 11194 | USR-IP-Call-Input-Filter | US Robotics | Text |
| 11195 | USR-IPX-RIP-Input-Filter | US Robotics | Text |
| 11196 | USR-MP-MRRU | US Robotics | Number |
| 11197 | USR-IPX-Call-Input-Filter | US Robotics | Text |
| 11198 | USR-AT-Input-Filter | US Robotics | Text |
| 11199 | USR-AT-RTMP-Input-Filter | US Robotics | Text |
| 11200 | USR-AT-Zip-Input-Filter | US Robotics | Text |
| 11201 | USR-AT-Call-Input-Filter | US Robotics | Text |
| 11202 | USR-ET-Bridge-Input-Filter | US Robotics | Text |
| 11203 | USR-IP-RIP-Output-Filter | US Robotics | Text |
| 11204 | USR-IP-Call-Output-Filter | US Robotics | Text |
| 11205 | USR-IPX-RIP-Output-Filter | US Robotics | Text |
| 11206 | USR-IPX-Call-Output-Filter | US Robotics | Text |
| 11207 | USR-AT-Output-Filter | US Robotics | Text |
| 11208 | USR-AT-RTMP-Output-Filter | US Robotics | Text |
| 11209 | USR-AT-Zip-Output-Filter | US Robotics | Text |
| 11210 | USR-AT-Call-Output-Filter | US Robotics | Text |
| 11211 | USR-ET-Bridge-Output-Filter | US Robotics | Text |
| 11212 | USR-ET-Bridge-Call-Output-Filter | US Robotics | Text |
| 11213 | USR-IP-Default-Route-Option | US Robotics | Number |
| 11214 | USR-MP-EDO-HIPER | US Robotics | String |
| 11215 | USR-Modem-Training-Time | US Robotics | Number |
| 11216 | USR-Interface-Index | US Robotics | Number |
| 11217 | USR-Tunnel-Security | US Robotics | Number |
| 11218 | USR-Port-Tap | US Robotics | Number |
| 11219 | USR-Port-Tap-Format | US Robotics | Number |
| 11220 | USR-Port-Tap-Output | US Robotics | Number |
| 11221 | USR-Port-Tap-Facility | US Robotics | Number |
| 11222 | USR-Port-Tap-Priority | US Robotics | Number |

| | | | |
|---|---|---|---|
| 11223 | USR-Port-Tap-Address | US Robotics | String |
| 11224 | USR-MobileIP-Home-Agent-Address | US Robotics | String |
| 11225 | USR-Tunneled-MLPP | US Robotics | Number |
| 11226 | USR-Multicast-Proxy | US Robotics | Number |
| 11227 | USR-Multicast-Receive | US Robotics | Number |
| 11228 | USR-Multicast-Forwarding | US Robotics | Number |
| 11229 | USR-IGMP-Query-Interval | US Robotics | Number |
| 11230 | USR-IGMP-Maximum-Response-Time | US Robotics | Number |
| 11231 | USR-IGMP-Robustness | US Robotics | Number |
| 11232 | USR-IGMP-Version | US Robotics | Number |
| 11233 | USR-IGMP-Routing | US Robotics | Number |
| 11234 | USR-VTS-Session-Key | US Robotics | Text |
| 11235 | USR-Orig-NAS-Type | US Robotics | Text |
| 11236 | USR-Call-Arrival-Time | US Robotics | Number |
| 11237 | USR-Call-End-Time | US Robotics | Number |
| 11238 | USR-Rad-Multicast-Routing-Ttl | US Robotics | Number |
| 11239 | USR-Rad-Multicast-Routing-Rate-Limit | US Robotics | Number |
| 11240 | USR-Rad-Multicast-Routing-Protocol | US Robotics | Number |
| 11241 | USR-Rad-Multicast-Routing-Boundary | US Robotics | Text |
| 11242 | USR-Rad-Dvmrp-Metric | US Robotics | Number |
| 11243 | USR-Chat-Script-Name | US Robotics | Text |
| 11244 | USR-Chat-Script-Rules | US Robotics | Text |
| 11245 | USR-Rad-Location-Type | US Robotics | Number |
| 11246 | USR-Tunnel-Switch-Endpoint | US Robotics | Text |
| 11247 | USR-OSPF-Addressless-Index | US Robotics | Number |
| 11248 | USR-Callback-Type | US Robotics | Number |
| 11249 | USR-Tunnel-Auth-Hostname | US Robotics | Text |
| 11250 | USR-Acct-Reason-Code | US Robotics | Number |
| 11251 | USR-DNIS-ReAuthentication | US Robotics | Number |
| 11252 | USR-PPP-Source-IP-Filter | US Robotics | Text |
| 11253 | USR-Auth-Mode | US Robotics | Number |
| 11254 | USR-NAS-Type | US Robotics | Number |
| 11255 | USR-Request-Type | US Robotics | Number |

# IAS Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section covers:

- Troubleshooting IAS as a RADIUS server
- Troubleshooting IAS as a RADIUS proxy

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Troubleshooting IAS as a RADIUS server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting IAS as a RADIUS server**
What problem are you having?

- A connection attempt fails.
- Connection attempts using Protected Extensible Authentication Protocol (PEAP) fail.
- Computers newly joined to the domain cannot enroll certificates.
- Connection attempts using Password Authentication Protocol (PAP) fail.
- Connection attempts using Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP) fail.
- Users with 802.1x authenticated switch client connections with EAP or PEAP are not prompted to input credentials and cannot log on.
- Users who should not have remote access permission can log on
- Users can log on successfully, but their logon attempts do not appear in the log files.
- You have just changed your Active Directory domain from a Windows 2000 mixed domain to a Windows 2000 native domain or a Windows Server 2003 domain, and IAS servers can no longer authenticate valid connection requests.
- Interim accounting messages are not being recorded in the remote access log file
- Third-party network access servers (NASs) drop connections and IAS sends a Microsoft vendor-specific attribute (MS-RAS-MPPE-Encryption-Policy) to the network access server instead of the NAS's vendor ID.
- MS-CHAP v2 or EAP-TLS authentications are being rejected for valid connection attempts.
- You can no longer see the Remote Access Policies node in Routing and Remote Access.
- When you try to open the IAS console, you receive an error message and the console does not open.

**A connection attempt fails.**
**Cause**: IAS has either rejected or discarded the connection attempt for a variety of reasons.

**Solution**: Verify that IAS is configured for the logging of rejected authentication attempts to the event log. Try the connection again, and then check the system event log for an IAS event for the connection attempt. Use the information in the log to determine the reason the connection attempt was either rejected or discarded and the way to fix the problem.

**See also**: Event logging for IAS; Configure event logging for IAS

**Connection attempts using Protected Extensible Authentication Protocol (PEAP) fail.**
**Cause**: The wireless remote access policy on the IAS server is not configured to use PEAP as the authentication method.

**Solution**: Create a wireless remote access policy on the IAS server that is configured with PEAP-EAP-MS-CHAPv2 or PEAP-EAP-TLS.

**See also**: Configure Remote Access Policies; Configure PEAP and EAP methods; Wireless access with secure password authentication

**Cause**: Remote access policy on the IAS server is configured to reject the user or computer.

**Solution**: Make sure that your remote access policy and user account settings are not configured to reject user authentication and authorization attempts. For example, if the user belongs to a group that is restricted in remote access policies to accessing the network between 9 A.M. and 5 P.M., any attempt by the user to log on between 5:01 P.M. and 8:59 A.M. will fail. To resolve the issue, change the day and time restrictions in remote access policies, or move the user to another group with different or no day and time restrictions. In addition, user account settings in Active Directory might prevent the user from accessing the network.

**See also**: Introduction to remote access policies; Elements of a remote access policy; Dial-in properties of a user account

**Cause**: You are using PEAP-EAP-MS-CHAPv2 as the authentication method and the IAS server does not have a certificate to use for server authentication to the client.

**Solution**: If the IAS server is a domain member computer, auto-enroll a certificate that contains the Server Authentication purpose in Enhanced Key Usage extensions and is configured according to minimum server certificate requirements. If the IAS server is not a domain member computer, enroll a certificate that contains the Server Authentication purpose in Enhanced Key Usage extensions (and which is configured according to minimum server certificate requirements) using the CA Web enrollment tool, or install a certificate on the IAS server using a floppy disk.

**See also**: The section "Certificate requirements for EAP" in Network access authentication and certificates.

**Cause**: You are using PEAP-EAP-MS-CHAPv2 as the authentication method and when prompted to make a decision whether or not to trust the server certificate, the user has chosen not to trust the server certificate.

**Solution**: Instruct users to trust the server certificate at next log on.

**Cause**: You are using PEAP-EAP-TLS as the authentication method and the client computer does not have a certificate that contains the Client Authentication purpose in Enhanced Key Usage extensions and is configured according to minimum client certificate requirements.

**Solution**: If the client is a domain member computer, auto-enroll a certificate that contains the Client Authentication purpose in Enhanced Key Usage extensions and is configured according to minimum client certificate requirements. If the client is not a domain member computer, enroll a certificate that contains the Client Authentication purpose in Enhanced Key Usage extensions (and which is configured according to minimum client certificate requirements) using the CA Web enrollment tool, or install a certificate on the IAS server using a floppy disk.

**See also**: The section "Certificate requirements for EAP" in Network access authentication and certificates.

**Cause**: The client computer does not trust the CA that issued the server certificate.

**Solution**: Add your organization's root certification authority certificate to the Trusted Root Certification Authorities folder in the computer certificate store.

**Cause**: You are using PEAP-EAP-TLS as the authentication method and the IAS server does not have a certificate to use for server authentication to the client.

**Solution**: If the IAS server is a domain member computer, auto-enroll a certificate that contains the Server Authentication purpose in Enhanced Key Usage extensions and is configured according to minimum server certificate requirements. If the IAS server is not a domain member computer, enroll a certificate that contains the Server Authentication purpose in Enhanced Key Usage extensions (and which is configured according to minimum server certificate requirements) using the CA Web enrollment tool, or install a certificate on the IAS server using a floppy disk.

**See also**: The section "Certificate requirements for EAP" in Network access authentication and certificates.

**Cause**: Group Policy on domain member computers has not been refreshed after the configuration of certificate auto-enrollment.

**Solution**: Refresh Group Policy on domain member computers. After auto-enrollment is configured and enabled, all domain member computers will receive computer certificates at the next refresh of Group Policy, whether the refresh is triggered manually with the gpupdate command or when the computer next logs on to the domain.

**See also**: Network access authentication and certificates; Configure automatic certificate allocation from an enterprise CA; Use Windows Server 2003 Certificate Services Web Pages; Manage certificates for a computer; and Request a certificate

**Computers newly joined to the domain cannot enroll certificates.**
**Cause**: When a Public Key Infrastructure (PKI) is deployed and autoenrollment is configured for the domain, Active Directory replication latency can temporarily affect the ability of a client or server to obtain a certificate from a certification authority.

**Solution**: Verify that Active Directory has replicated to all domain controllers, and then refresh Group Policy to enroll a certificate on the computer. You can refresh Group Policy by running the gpupdate command or by logging on to the domain.

**See also**: Checklist: Configuring certificate autoenrollment; Authentication methods; Troubleshooting replication

**Connection attempts using Password Authentication Protocol (PAP) fail.**
**Cause**: The password is incorrect.

**Solution**: Try the connection again and verify that the password is correct.

**Cause**: Mismatched shared secrets.

**Solution**: Verify that the same shared secret is configured on the access server and for the RADIUS client in Internet Authentication Service.

**See also**: Edit RADIUS client configuration; Shared secrets

**Connection attempts using Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP) fail.**
**Cause**: The password is incorrect.

**Solution**: Try the connection again and verify that the password is correct.

**Cause**: The access client is using LAN Manager authentication.

**Solution**: By default, MS-CHAP for Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition does not support LAN Manager authentication. If you want to allow the use of LAN Manager authentication with MS-CHAP for older Microsoft operating systems such as Windows NT 3.5x and Windows 95, you must set the following registry value to 1 on the IAS server:
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RemoteAccess\Policy\Allow LM Authentication

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**See also**: MS-CHAP; Add a value to a registry key entry

**Users with 802.1x authenticated switch client connections with EAP or PEAP are not prompted to input credentials and cannot log on.**
**Cause**: The icon for the network connection must be displayed in the notification area on the client.

**Solution**: Configure the network connection properties on the client so that the network connection icon is displayed in the notification area when connected.

**See also**: Show a network connection icon in the notification area

**Cause**: The switch is not configured to initiate authentication with an EAP-Request-Identity packet when it detects that a client is active on the port.

**Solution**: Configure the switch to initiate authentication.

**Users who should not have remote access permission can log on.**
**Cause**: A remote access policy might be granting access to these users.

**Solution**: Check the policy list to ensure that you have not included groups in which the user is a member that should be denied access.

**Cause**: The remote access permission on the user account dial-in properties is set to Grant access, which overrides the remote access permission of the remote access policy.

**Solution**: Check the dial-in properties for the user object.

**See also**: Configure remote access permission for a user; Dial-in properties of a user account

**Cause**: Remote access policies might be evaluated in the wrong order.

**Solution**: Authorization is granted or denied by the first policy with conditions that match the connection attempt. Use Move Up to move the policy that denies access to these users, so that it is higher in the list.

**See also**: Configure remote access permission for a user; Dial-in properties of a user account

**Users can log on successfully, but their logon attempts do not appear in the log files.**
**Cause**: IAS is not configured to log successful authentication attempts.

**Solution**: Configure IAS to log successful authentication attempts in the event log.

**See also**: Event logging for IAS; Configure event logging for IAS

**You have just changed your Active Directory domain from a Windows 2000 mixed domain to a Windows 2000 native domain or a Windows Server 2003 domain, and IAS servers can no longer authenticate valid connection requests.**
**Cause**: After you change a domain from a Windows 2000 mixed domain to a Windows 2000 native domain or a Windows Server 2003 domain, every domain controller in the domain must be restarted in order for the change to replicate.

**Solution**: Restart the domain controllers so that the IAS servers can authenticate valid connection requests.

**Interim accounting messages are not being recorded in the remote access log file.**
**Cause**: Remote access logging is not configured to record periodic status.

**Solution**: Configure remote access logging to record periodic status.

**Cause**: You have not added and configured the Acct-Interim-Interval RADIUS accounting attribute on the Advanced tab of the appropriate remote access policy.

**Solution**: Add and configure the Acct-Interim-Interval RADIUS accounting attribute on the Advanced tab of the appropriate remote access policy. The value of the Acct-Interim-Interval attribute is the number of seconds between interim accounting updates.

**Cause**: The access server does not support interim accounting or is not configured to send interim accounting messages.

**Solution**: Verify that the access server supports interim accounting and the Acct-Interim-Interval attribute. If so, consult the access server documentation for instructions about how to enable interim accounting. The Routing and Remote Access service supports sending interim accounting messages and is enabled from the properties of a RADIUS server when configuring the RADIUS accounting provider.

**See also**: Logging user authentication and accounting requests; Add RADIUS attributes to a remote access policy; Use RADIUS accounting

**Third-party network access servers (NASs) drop connections and IAS sends a Microsoft vendor-specific attribute (MS-RAS-MPPE-Encryption-Policy) to the network access server instead of the NAS's vendor ID.**
**Cause**: The encryption policy as configured in remote access policies is not set to the default.

**Solution**: The Microsoft VSA MS-RAS-MPPE-Encryption-Policy is sent to third-party network access servers if the encryption policy as configured in remote access policies is not set to the default. The default encryption policy allows all encryption levels and does not contain any attributes specific to the Routing and Remote Access service, so applying default encryption settings will prevent IAS from sending any Microsoft VSAs to third-party NASs. Take the following steps in the profile of the remote access policy:

- On the Encryption tab, return the encryption policy to the default. Select all encryption levels except No encryption.
- On the Advanced tab, remove all attributes specific to the Routing and Remote Access service from the list of attributes for the remote access policy.
- On the IP tab, remove all IP filters.
- On the Multilink tab, select Server settings determine Multilink usage. Return all other settings to the default (not selected).

**See also**: Configure encryption; Configure vendor-specific attributes for a remote access policy; Vendor-specific attribute overview; Recreate the default remote access policy; Configure IP options

**MS-CHAP v2 or EAP-TLS authentications are being rejected for valid connection attempts.**
**Cause**: The authentication method (in the profile settings of the connection request policy) that matched the connection request is configured for the Accept the connection attempt without performing authentication or authorization setting. MS-CHAP v2 and EAP-TLS are mutual authentication protocols. In mutual authentication, the access client proves that it is a valid

access client to the authenticating server and the authenticating server proves that it is a valid authenticating server to the access client. When this authentication option is used, although the connect request is granted, the authenticating server does not provide validation to the access client. As the result, mutual authentication fails.

**Solution**: Reconfigure connection request policies so that MS-CHAP v2 and EAP-TLS connections are not being accepted without authentication and authorization.

**See also**: Connection request policies

**You can no longer see the Remote Access Policies node in Routing and Remote Access.**
**Cause**: When a Routing and Remote Access server is configured to use RADIUS authentication, the remote access policies on the Routing and Remote Access server are never evaluated for authorization.

**Solution**: This is intended behavior. If the RADIUS server of the server running the Routing and Remote Access service is an IAS server, remote access policies are configured through Internet Authentication Service.

**When you try to open the IAS console, you receive an error message and the console does not open.**
**Cause**: To administer IAS, you must have Administrative credentials.

**Solution**: You can use the Runas command to perform tasks (for example, opening the IAS console), when you are logged on as a member of a group that does not have administrative credentials (such as Users or Power Users). Alternately, you can log on as Administrator to open the console.

**See also**: Run a program with administrative credentials; Create a shortcut using the runas command; Runas

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Troubleshooting IAS as a RADIUS proxy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting IAS as a RADIUS proxy**
What problem are you having?

- No responses to Access-Request messages are sent through the IAS proxy for valid connection attempts.
- MS-CHAP v2 authentications are being rejected for valid connection attempts.
- When you try to open the IAS console, you receive an error message and the console does not open.

**No responses to Access-Request messages are sent through the IAS proxy for valid connection attempts.**
**Cause:** A mismatched shared secret is configured between the IAS proxy and the RADIUS server or proxy to which the RADIUS messages are being forwarded.

**Solution:** Verify that the shared secret that is configured on the IAS proxy for the appropriate remote RADIUS server group member matches the shared secret of the RADIUS client of the RADIUS server or proxy to which the RADIUS messages are being sent.

**See also:** Shared secrets

**MS-CHAP v2 authentications are being rejected for valid connection attempts.**
**Cause:** Attribute manipulation rules on the profile setting for the connection request policy that is forwarding the requests are modifying the User-Name attribute. This causes the validation of the MS-CHAP v2 hash on the authenticating server to fail. The only exception is when a backslash (\) character is used and the manipulation only affects the information to the left of it. A backslash character is typically used to indicate a domain name (the information to the left of the backslash character) and a user account name within the domain (the information to the right of the backslash character). In this case, attribute manipulation rules that only modify or replace the domain name are allowed.

**Solution:** Reconfigure attribute manipulation rules so that they are not modifying the User-Name attribute.

**See also:** Connection request policies

**When you try to open the IAS console, you receive an error message and the console does not open.**
**Cause:** To administer IAS, you must have administrative credentials.

**Solution:** You can use the Runas command to perform tasks (for example, opening the IAS console), when you are logged on as a member of a group that does not have administrative credentials (such as Users or Power Users). Alternately, you can log on as Administrator to open the console.

**See also:** Run a program with administrative credentials, Create a shortcut using the runas command, Runas

**Note**

- You can configure IAS in Windows Server 2003, Standard Edition, with a maximum of 50 RADIUS clients and a maximum of 2 remote RADIUS server groups. You can define a RADIUS client using a fully qualified domain name or an IP address, but you cannot define groups of RADIUS clients by specifying an IP address range. If the fully qualified domain name of a RADIUS client resolves to multiple IP addresses, the IAS server uses the first IP address returned in the DNS query. With IAS in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, you can configure an unlimited number of RADIUS clients and remote RADIUS server groups. In addition, you can configure RADIUS clients by specifying an IP address range.

# Configuring Communication Links

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring communication links**
This section covers:

- Network Connections
- Installing and configuring Modems
- Wireless Networking
- HyperTerminal
- Telnet

# Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Connections**
Network Connections provides connectivity between your computer and the Internet, a network, or another computer. With Network Connections, you can configure settings to reach local or remote network resources or functions.

Network Connections combines Microsoft® Windows® NT 4.0 Dial-Up Networking with features that were formerly located in the Network Control Panel, such as network protocol and service configuration. Each connection in the Network Connections folder contains a set of features that creates a link between your computer and another computer or network. By using Network Connections, performing a task, such as modifying a network protocol, is as easy as right-clicking a connection and then clicking **Properties**.

- Before using Network Connections, see Checklist: Configuring network connections.
- Before setting up a home or small office network, see Checklist: Setting up a home or small office network.
- To find features that have been moved, see New ways to do Network Connections tasks.
- For tips about using Network Connections, see Network Connections Best Practices.
- For help with specific tasks, see Network Connections How To....
- For general background information, see Network Connections Concepts.
- For problem-solving instructions, see Network Connections Troubleshooting.

# Checklist: Configuring network connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configuring network connections**

| | Step | Reference |
|---|---|---|
| | Review key concepts. | Network Connections Overview |
| | Connect to a local area network. | Make a local area connection |
| | Connect to a network by using a phone line. | Make a dial-up connection by using a phone line;<br>Connect to a remote network |
| | Connect to an Internet service provider (ISP). | Make an Internet connection;<br>Connect to a remote network |
| | Connect to a private network through the Internet (virtual private network). | Make a virtual private network (VPN) connection;<br>Connect to a remote network |
| | Create a direct connection to another computer. | Make a direct network connection;<br>Connect to a remote network |
| | Let other computers connect to this computer. | Accept incoming network connections |
| | Monitor a connection. | Monitor a connection |

# Checklist: Setting up a home or small office network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Setting up a home or small office network**

| Step | Reference |
|---|---|
| Decide whether to use home and small office networking or Routing and Remote Access for your home or small office network. | Home and small office networking compared to Routing and Remote Access |
| If you decide to use home and small office networking, review home and small office networking concepts. | Home and Small Office Networking; Using Home and Small Office Networking Features |
| Decide how you want to save your Internet connection credentials (you will be prompted to select credential settings when you fist attempt to activate your Internet connection or when you run the New Connection Wizard). | Dial-up credentials |
| Familiarize yourself with the icons that indicate that Internet Connection Firewall (ICF), Windows Firewall, Network Bridge, or Internet Connection Sharing (ICS) have been enabled on a network connection. | Network connection icons |
| Enable ICF or Windows Firewall to protect your home or small office network. | Enable or disable ICF |
| Set up Network Bridge to connect network segments. | Set up Network Bridge |
| Enable ICS to provide a shared Internet connection. | Enable ICS |

**Notes**

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

# New ways to do Network Connections tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

This table lists common tasks for Network Connections. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Connect to a dial-up server | **Dial-Up Networking** in My Computer | Network Connections. For more information, see Connect to a remote network. |
| Connect to a network by using a modem or ISDN line | **Dial-Up Networking** in My Computer (click **New**) | Network Connections. For more information, see Make a dial-up connection by using a phone line. |
| Connect to a network by using a virtual private network (VPN) | **Dial-Up Networking** in My Computer (click **New**) | Network Connections. For more information, see Make a virtual private network (VPN) connection. |
| Configure a computer for incoming connections | **Network** in Control Panel | Network Connections. For more information, see Accept incoming network connections. |
| Configure the TCP/IP protocol | **Network** in Control Panel | Network Connections. For more information, see Configure TCP/IP settings. |
| Configure a connection | **Dial-Up Networking** in My Computer (click **More**) | Network Connections. For more information, see Configure a connection to a remote network. |
| Add a client or service | **Network** in Control Panel | Network Connections. For more information, see Add a network component. |
| Add a protocol | **Network** in Control Panel | Network Connections. For more information, see Add a network component. |
| Add optional networking components | **Network** in Control Panel | Network Connections. For more information, see Add optional networking components. |
| Monitor incoming or outgoing connections | **Dial-Up Monitor** in Control Panel | Network Connections. For more information, see Monitor a connection. |

# Network Connections Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

The following list provides best practices for implementing Network Connections:

- **When using multiple network adapters, rename each local area connection.**

  Your operating system detects network adapters and automatically creates a local area connection in the Network Connections folder for each network adapter. If more than one network adapter is installed, you can eliminate possible confusion by immediately renaming each local area connection to reflect the network that it connects to.

  You need to add or enable the network clients, services, and protocols that are required for each connection. When you do so, the client, service, or protocol is added or enabled in all other network and dial-up connections.

- **Verify required connection settings for your network adapter.**

  If your network administrator or Internet service provider (ISP) requires static settings, you may need one or more of the following:

  - A specific IP address.
  - DNS addresses.
  - DNS domain name.
  - Default gateway address.
  - WINS addresses.

  DHCP is enabled by default. This gets your IP address and other configuration information automatically. Automated IP settings are used for all connections, and they eliminate the need to configure settings such as DNS, WINS, and so on.

  For more information about enabling DHCP, see Configure TCP/IP settings.

- **Create dial-up, VPN, or direct connections by using the New Connection wizard.**

  After you create a connection, you can copy the connections, rename them and modify the connection settings. By doing so, you can easily create different connections to accommodate multiple modems, ISPs, dialing profiles, and so on.

  For more information, see Copy a network or dial-up connection.

- **Specify the order in which network providers and protocols are accessed.**

  By changing the order of protocols bound to those providers, you can improve performance. For example, if your LAN connection is enabled to access NetWare and Microsoft Windows networks, which use IPX and TCP/IP, but your primary connection is to a Microsoft Windows network that uses TCP/IP, you can move **Microsoft Windows Network** to the top of the **Network providers** list on the **Provider Order** tab, and move **Internet Protocol (TCP/IP)** to the top of the **File and Printer Sharing for Microsoft Networks** binding on the **Adapters and Bindings** tab.

  The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

  For more information about modifying the order of network providers, see Modify the network provider order. For more information about modifying the order of protocol bindings, see Modify the protocol bindings order.

- **Only install and enable the network protocols that you need.**

  Limiting the number of protocols on your computer enhances its performance and reduces network traffic.

  If your computer encounters a problem with a network or dial-up connection, it attempts to establish connectivity by using every network protocol that is installed and enabled. By only installing and enabling the protocols that your system can use, the operating system does not attempt to connect with protocols it cannot use, and returns status information to you more efficiently.

# Network Connections How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Connect to a remote network
- Disconnect from a remote network
- Add optional networking components
- Make Network and Dial-up Connections
- Configure Network and Dial-up Connections
- Monitor Network Connections
- Set up and Manage Home and Small Office Networking

# Connect to a remote network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a remote network**

1. Open Network Connections.

2. Double-click the connection you want to connect to the network.

3. If you are prompted, in the **Connect***connection type* dialog box, type your user name, password, and logon domain.

   If **Logon domain** does not appear in this dialog box, and you want to log on to a domain, type your user name and the domain name in one of two ways:

   - Your user principal name prefix (your user name) and your user principal name suffix (your domain name), joined by the "at" sign (@). For example, user@sales.westcoast.microsoft.com.

   - Your domain name and your user name, separated by the backslash (\) character. For example, sales\user.

     Note that the suffix in the first example is a fully-qualified DNS domain name. Your administrator might have created an alternative suffix to simplify the logon process. For example, creating a user principal name suffix of "microsoft" allows the same user to log on by using the much simpler user@microsoft.com. For more information, see your system administrator.

   Instead of having to type your domain name, you can also enable Windows to include your logon domain to the **Connect***connection type* dialog box.

   a. Click **Properties** if the button is enabled.

   b. On the **Options** tab, select the **Include Windows logon domain** check box.

   Once you are connected to the network, you can minimize your connection window and use e-mail, Windows Explorer, and so on.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- To connect to a remote network you can also right-click the connection that you want to connect to the network and click **Connect**.

- Authentication happens at two levels: logging on to the local computer domain and logging on as a valid user.

- If you are on a local area network, your computer logs on to the local computer domain automatically even before you see the logon prompt. When you type your user credentials at the logon prompt, the computer checks your user credentials, opens a user context, and makes your personal data available to you. You are now logged on at both levels.

- If you are not on a local area network, the attempt to log on to the local computer domain fails silently. In this case, user logon can proceed in two ways.

  - When you log on at the logon prompt, the computer checks your cached credentials from the last time you were actually online with the domain controller. If they match, a user context associated with your name is opened. You

are now logged on to the computer as a valid user but you are not logged on to the network. To access your personal files you can now create a remote access connection by using your remote access credentials or by using your cached credentials if your computer is a member of the domain.

- Select the **Log on using dial-up networking** check box if your computer is a member of the domain. In this case, the computer will attempt to make a remote access connection before processing your credentials. Once connected it uses the user logon credentials as though you were on a local area network.

- The logon domain name you type should name the domain that the server is in. This is not the DNS domain name given by some PPP/SLIP providers.

- If you are using IP connectivity on your local network connection and on your remote connection, you may not be able to see all computers on your local network. This is because after you connect, your remote connection becomes your default path for network routing. Consequently, you see computers on the remote network, and you see other computers on the same LAN segment to which your computer is connected. But you cannot communicate with computers on networks that were previously reached through a router on your local LAN.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Automating a Network or Dial-up Connection
Authentication process
Make a virtual private network (VPN) connection

# Disconnect from a remote network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect from a remote network**

1. Open Network Connections.
2. Right-click the connection you want to disconnect, and then click **Disconnect**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you disable a local area connection, it disables the network adapter driver until the connection is reactivated.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a remote network
Make a virtual private network (VPN) connection

# Add optional networking components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add optional networking components**

1. Open Network Connections.
2. On the **Advanced** menu, click **Optional Networking Components**.
3. Do one of the following:
   - To install all optional networking components, select the **Management and Monitoring Tools**, **Networking Services**, and **Other Network File and Print Services** check boxes, and then click **Next**.
   - To install specific optional networking components, select its corresponding check box, and then click **Details**.

     In **Subcomponents of***component*, select the check boxes next to the components you want to install, and then click **OK**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Examples of optional networking components include Simple Network Management Protocol (SNMP), Simple TCP/IP Services, and Print Services for UNIX.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Optional networking components
New ways to do Network Connections tasks

# Make Network and Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Make network and dial-up connections**

- Make a virtual private network (VPN) connection
- Make a local area connection
- Manage Network Connections
- Set up ISDN Connections
- Set up Dial-up Connections
- Configure Dial-up Options
- Make Internet Connections
- Make Direct Network Connections
- Set up Incoming Connections

# Make a virtual private network (VPN) connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a virtual private network (VPN) connection**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Connect to the network at my workplace**, and then click **Next**.
4. Click **Virtual Private Network connection**, and click **Next**.
5. Type a name for this connection in **Company Name**, and click **Next**.
6. If you have already established a dial-up connection, do one of the following:
   - If you need to establish a connection with your ISP or some other network before tunneling to your destination computer or network, click **Automatically dial this initial connection**, click a connection in the list, and then click **Next**.
   - If you do not want to automatically dial an initial connection, click **Do not dial the initial connection**, and then click **Next**.
7. Type the host name or IP address of the VPN server to which you are connecting, and then click **Next**.
8. Do one of the following:
   - If you want this connection to be made available to all users on this computer, click **Anyone's use**, and then click **Next**.
   - If you want to reserve the connection for your own use, click **My use only**, and then click **Next**.
9. If you want to create a shortcut for this connection on your desktop, select the **Add a shortcut to this connection to my desktop** check box.
10. Click **Finish**.

### Notes

- To make the connection available to all users, you must be logged on as Administrator or as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You can create multiple VPN connections by copying them in the Network Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple hosts, security options, and so on.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Virtual Private Network (VPN) Connections
Troubleshooting network and dial-up connections
PPP operation and protocols
Outsourcing dial-up networks
Point-to-Point Protocol (PPP)

# Make a local area connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a local area connection**

- Typically, most computers running a member of the Windows Server 2003 family are connected to a local area network. When you start your computer, your network adapter is detected and the local area connection automatically starts. Unlike other types of connections, the local area connection is created automatically, and you do not have to click the local area connection in order to start it.

Once the local area connection is running, you can choose to disable it. To do so, do the following.

1. Open Network Connections.
2. Right-click the **LAN or High Speed Internet** connection that you want to disable.
3. Click **Disable**.

If you need to enable a disabled **LAN or High Speed Internet** connection, do the following.

1. Right-click the connection that you want to enable.
2. Click **Enable**.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- A local area connection is automatically created for each network adapter that is detected.
- If more than one network adapter is installed, you can eliminate possible confusion by immediately renaming each local area connection to reflect the network that it connects to.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using local area connections
View the status of a local area connection
Set up Network Components
Network connection icons

# Manage Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage network connections**

- Copy a network or dial-up connection
- Show a network connection icon in the notification area
- Rename a network connection icon
- Repair a LAN or high-speed Internet connection

# Copy a network or dial-up connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy a network or dial-up connection**

1. Open Network Connections.
2. Right-click your dial-up, VPN, or direct connection, and then click **Create Copy**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You cannot copy local area and incoming connections.
- You can create multiple dial-up, virtual private network, or direct connections by copying them in the Network Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple modems, ISPs, dialing profiles, and so on.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Network connection icons

# Show a network connection icon in the notification area

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To show a network connection icon in the notification area**

1. Open Network Connections.
2. Right-click the icon for the network connection that you want to display in the notification area, and then click **Properties**.
3. Select the **Show icon in notification area when connected** check box, and then click **OK**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If a network connection icon is displayed in the notification area, you can access basic connection and statistical information by pointing to the icon.
- If the **Show icon in notification area when connected** check box is selected for the shared Internet connection, an icon for the shared Internet connection appears in the notification area each time that connection is enabled. When the shared Internet connection is disabled, the icon disappears from the notification area.
- It is a good practice to select this check box when you are using wireless connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Home and small office networking features
Rename a network connection icon

# Rename a network connection icon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a network connection icon**

1. Open Network Connections.
2. Right-click the icon for the network connection that you want to rename, and click **Rename**.
3. Type a new name for the selected network connection icon.

## Notes

- You must be logged on as a member of the Administrators group to rename a network connection icon. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You should rename each network connection icon with a name that is easy to remember.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a network connection icon in the notification area
Home and small office networking features

# Repair a LAN or high-speed Internet connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To repair a LAN or high-speed Internet connection**

1. Open Network Connections.
2. Before using the **Repair** option, try the following steps:
   - Check the status indicator lights on the DSL or cable modem to verify that the connection to your Internet service provider (ISP) is working.
   - Turn off the DSL or cable modem, and then turn it on again.
   - If your Internet connection is not working after these steps, contact your ISP.
3. If media is connected properly, under **LAN or High-Speed Internet**, right-click the connection that you want to repair, and then click **Repair**.

## Notes

- You must be a member of the Administrators group, the Power Users group, or the Network Configuration Operators group to complete this procedure. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you have more than one network adapter installed in your computer, it is recommended that you rename your Internet connection to easily identify it. For example, you could call your Internet connection "Internet connection" or use the name of your ISP.
- The **Repair** option is not available for media disconnected connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet service providers
Copy a network or dial-up connection
Network connection icons

# Set up ISDN Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up ISDN connections**

- Make a dial-up connection by using an ISDN line
- Install an ISDN adapter
- Configure ISDN settings

# Make a dial-up connection by using an ISDN line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a dial-up connection by using an ISDN line**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. In **Network Connection Type**, do one of the following:
   - Click **Connect to the Internet**. In **Internet Connection**, click **Connect using a dial-up modem**, and then click **Next**.
   - Click **Connect to the network at my workplace**. In **Network Connection**, click **Dial-up connection**, and then click **Next**.
4. Follow the instructions in the New Connection wizard.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To enable automatic dialing for all users, create a connection, save your user name and password for all users, and then enable Internet Connection Sharing for the new connection.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using ISDN lines
Configuring multiple device dialing
Install an ISDN adapter
Configure ISDN settings
Troubleshooting network and dial-up connections

# Install an ISDN adapter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install an ISDN adapter**

1. Install an Integrated Services Digital Network (ISDN) adapter and start your computer. If you install more than one ISDN adapter, you must restart your computer after you have installed the ISDN drivers. Otherwise, all available ISDN ports may not show up when you configure your network connection for ISDN.

2. You can use Device Manager to configure your system so that your ISDN adapter knows what type of telephone switch it is connected to. The switch type simply refers to the brand of equipment and software revision level that the telephone company uses to provide you with ISDN service. There are only a few types of switches in the world and usually just one in countries/regions other than the United States.

## Notes

- To use Device Manager, you must be logged on as an administrator or a member of the Administrators group.

- To open Device Manager, click **Start**, right-click **My Computer**, click **Manage**, and then click **Device Manager**.

- If your ISDN adapter is internal to your computer, it appears in **Network adapters**. If the ISDN adapter is external to your computer, it appears in **Modems**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure ISDN settings
ISDN access
Using ISDN lines
Set up ISDN Connections

# Configure ISDN settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure ISDN settings**

1. Open Network Connections.
2. Right-click the dial-up connection that uses ISDN, and then click **Properties**.
3. On the **General** tab, in **Connect using**, click the ISDN device, and then click **Configure**.
4. In the **ISDN Configure** dialog box, do one or both of the following:
   - In **Line type**, click the line type you want to use. The line types are listed from highest to lowest quality.
   - If you want to start negotiation with the line type you selected, and then negotiate with a lower-quality line type depending on the line's condition, select the **Negotiate line type** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Depending on the type of ISDN adapter that you are using, when you click **Configure** in step 3, you may be presented with **Modem Configuration** settings.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using ISDN lines
Troubleshooting network and dial-up connections

# Set up Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up dial-up connections**

- Make a dial-up connection by using a phone line
- Configure a modem for a dial-up connection
- Make a dial-up connection by using X.25
- Start Remote Access Auto Connection Manager

# Make a dial-up connection by using a phone line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a dial-up connection by using a phone line**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. In **Network Connection Type**, do one of the following:
   - Click **Connect to the Internet**. In **Internet Connection**, click **Connect using a dial-up modem**, and then click **Next**.
   - Click **Connect to the network at my workplace**. In **Network Connection**, click **Dial-up connection**, and then click **Next**.
4. Follow the instructions in the New Connection wizard.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To enable automatic dialing for all users, create a connection, save your user name and password for all users, and then enable Internet Connection Sharing for the new connection.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using phone lines and modems
Using an intermediary device
Install a modem
Configure a modem for a dial-up connection
Install a modem

# Configure a modem for a dial-up connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a modem for a dial-up connection**

1. Open Network Connections.
2. Right-click the dial-up connection you want to configure, and then click **Properties**.
3. On the **General** tab, under **Connect using**, click the modem you want to configure, and then click **Configure**.
4. Under **Hardware Features**, select the check box options you want to enable.
5. If you want to enable the modem speaker, select the **Enable modem speaker** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To ensure compatibility, you need to use the same kind of modem as the one connected to the remote access server, select the same initial speed, and enable the same features. If you do not select the same model, at least select a modem with the same ITU-T standard as the modem on the server.
- Selecting a feature that is not supported by your modem has no effect on its performance.
- Network Connections automatically configures connections according to the devices that are available. For example, you might use a laptop and docking station when you are in the office, with your connection configured to use the docking station modem. When you are on the road, and the docking station modem is not available, the connection is automatically configured to use the laptop's PCMCIA modem. When you return to the office and re-dock your laptop into its docking station, Network Connections detects that the docking station modem is available again, and automatically reconfigures the connection to use the docking station modem.
- Modem changes made as described in this procedure are unique to this connection. To make global changes, you will need to change each modem.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install a modem
Information about specific modems

# Make a dial-up connection by using X.25

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a dial-up connection by using X.25**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. In **Network Connection Type**, click **Connect to the Internet** or **Connect to the network at my workplace**, and then click **Next**.
4. In **Network Connection**, click **Connect using a dial-up modem** or **Dial-up connection**, and then click **Next**.
5. Follow the instructions in the New Connection wizard.
6. When the wizard is finished, right-click the new connection, and then click **Properties**.
7. On the **Options** tab, click **X.25**.
8. In **Network**, click your X.25 network provider.
9. In **X.121 address**, type the X.121 address (the X.25 equivalent of a phone number) for the server you want to call.
10. If required by the X.25 host computer, in **User Data**, type additional connection information.
11. In **Facilities**, type any additional facility options you want to request from your X.25 provider. (For example, some providers support the **/r** option to specify reverse charging.)

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To enable automatic dialing for all users, create a connection, save your user name and password for all users, and then enable Internet Connection Sharing for the new connection.
- For details on facility options, consult your X.25 documentation or provider.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**See Also**
**Concepts**
Using X.25
Troubleshooting network and dial-up connections

# Start Remote Access Auto Connection Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Remote Access Auto Connection Manager**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. In the details pane, right-click **Remote Access Auto Connection Manager**, and then click **Start**.

   Under the **Status** column, **Started** appears.

### Note

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure redial options
Configuring redial options
Working with MMC console files

# Configure Dial-up Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure dial-up options**

- Configure client callback options
- Configure autodial
- Enable logging
- Enable multiple device dialing
- Configure phone number modifiers
- Configure redial options
- Assign multiple phone numbers to a connection
- Enable operator-assisted calls or manual dialing

# Configure client callback options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure client callback options**

1. Open Network Connections.
2. On the **Advanced** menu, click **Remote Access Preferences**.
3. On the **Callback** tab, do one of the following:
   - If you do not want to use callback, click **No callback**.
   - If you want to decide whether to use callback at the time you connect, click **Ask me during dialing when the server offers**.
   - If you want to consistently use callback, click **Always call me back at the number(s) below**, and select the modem or device you want to be called back.

     If **Phone number** is blank for the device you have selected, click **Edit**, and then type the number.

     If you want to remove a modem or device from the list of possible callback devices, click the modem or device, and then click **Delete**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Your system administrator can override any callback settings that you configure. Through server settings, the administrator can deny the use of callback, allow you to set callback options, or require callback to a particular number. Contact your system administrator to determine your callback options.
- To use callback, Link Control Protocol (LCP) extensions must be enabled.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Callback
Start requesting LCP extensions in PPP
PPP operation and protocols

# Configure autodial

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure autodial**

1. Open Network Connections.
2. On the **Advanced** menu, click **Remote Access Preferences**.
3. On the **Autodial** tab, select the check box next to the location where you want to enable autodial, and then click **OK**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Autodial is enabled by default in Windows XP. In support of this feature, the Remote Access Auto Connection Manager is also started by default.
- Autodial maps the resources that are required to initially make a connection, but it does not map resources after the connection is made. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Autodial
Automating a Network or Dial-up Connection
Start Remote Access Auto Connection Manager

# Enable logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable logging**

1. Open Network Connections.
2. On the **Advanced** menu, click **Remote Access Preferences**.
3. On the **Diagnostics** tab, do one or more of the following:
   - To enable logging for all remote access connections, select the **Enable logging** check box.
   - To clear logs created during a previous connection, click **Clear Log**.
   - To generate a report, click **Generate Report**.
4. If you clicked **Generate Report** in the previous step, do one of the following:
   - To view the diagnostic report, click **View report**.
   - To save the report to a file, click **Save report to this file**, and specify the location and the name of the file you want to create. Specify the location by typing it or by clicking **Browse** and browsing to the location you want.
   - To send the report to an e-mail address, click **Send report to this e-mail address**, and type the e-mail address to which you want to send the report.
   - To quickly generate a more concise report, select the **Generate summary report** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Creating a diagnostic report in HTML can take several minutes. If you need only configuration and installation information, you can create a summary report.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure autodial
Configure callback for incoming connections
Logging remote access connections

# Enable multiple device dialing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable multiple device dialing**

1. Open Network Connections.
2. Right-click the connection on which you want to enable the dialing of multiple devices, and then click **Properties**.
3. On the **Options** tab, in **Multiple devices**, do one of the following:
   - If you want to dial only the first available device, click **Dial only first available device**.
   - If you want to use all of your devices, click **Dial all devices**.
   - If you want to dynamically dial and hang up devices as needed, click **Dial devices only as needed**, and then click **Configure**.

     In **Automatic dialing**, click the **Activity at least** percentage and **Duration at least** time you want to set. Another line is dialed when connection activity reaches this level for the amount of time that you specify.

     In **Automatic hangup**, click the **Activity no more than** percentage and **Duration at least** time you want to set. A device is disconnected when connection activity decreases to this level for at least the amount of time that you specify.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you selected **Dial devices only as needed**, the last multilinked device ignores the **Automatic hangup** setting, and a twenty-minute time-out is used for the last device.
- If you use multiple devices to dial a server that requires callback, then only one of your multilinked devices is called back. This is because only one number is stored in a user account. Therefore, only one device connects and all other devices fail to complete the connection, and your connection loses Multilink functionality.

  You can avoid this problem if the multilinked phonebook entry is an ISDN with two channels that have the same phone number.
- Multiple device dialing is available only if multiple adapters are available on the computer.
- If you select **Dial all devices**, dropped links in the multilinked bundle are not automatically reinitialized. You can force links to reinitialize by selecting **Dial devices only as needed**, then **Configure**, and then setting easily achieved **Automatic dialing** conditions which cause another line to be dialed. For example, set **Activity at least** to 1% and **Duration at least** to 3 seconds.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Configuring multiple device dialing
ISDN access
Incoming connections
Configuring dial-up options
Using ISDN lines

# Configure phone number modifiers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure phone number modifiers**

1. Open Network Connections.
2. Right-click the dial-up connection you want to configure, and then click **Properties**.
3. On the **General** tab, modify the phone number by using one or more characters from the following list.

| Character | Function |
|---|---|
| , (comma) | Pauses briefly (two seconds for most modems) before continuing. |
| P | Switches from touch-tone to pulse/rotary dialing. |
| T | Switches from pulse/rotary to touch-tone dialing. |
| $ | Waits for a calling card prompt tone. |

For example, when you dial in from a hotel room, you may have to add 9 to the number you are calling to get an outside line. Your number may look similar to the following:

9,555-0100

The comma after the 9 gives a pause long enough to get an outside line before continuing to dial the full number.

**Notes**

- To configure phone number modifiers, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- By selecting the **Use dialing rules** check box, and then clicking **Dialing Rules**, you can create a **Location** for which specific dialing rules automatically apply. For example, if you frequently connect from home and need to disable call waiting, you can create a location for use from home that automatically dials the code to disable call waiting before the connection call is placed.
- If you see unfamiliar characters being dialed before your number, check to make sure the **Use dialing rules** check box is not selected. If it is selected, a location may be enforcing a dialing rule that you do not need. For example, a location may be selected that always dials 1 to accommodate long distance dialing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Specify area code rules
Use a calling card for long distance dialing
Change dialing properties for a location
Create a new location
Specify numbers to dial an outside line

Modify an existing calling card
Create a new calling card

# Configure redial options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure redial options**

1. Open Network Connections.
2. Right-click the dial-up connection you want to configure, and then click **Properties**.
3. On the **Options** tab, in **Redialing options**, click the arrows to set the number of times the dial-up connection is automatically redialed if the first attempt to connect fails.
4. In **Time between redial attempts**, click the down arrow to set the amount of time to pause between the end of one attempt and the beginning of another.
5. In **Idle time before hanging up**, click the down arrow to set the amount of idle time before a client hangs up a connection.
6. If you want the connection to automatically redial if the line is dropped, select the **Redial if line is dropped** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The **Time between redial attempts** pause lets the device reset itself before redialing. The default is one minute. If that is not enough time, increase this setting. You can also experiment with shorter times, but if you make the pause too short, the device does not have time to reset itself.
- Keep in mind that the remote access server also has a disconnect timer. If **Idle time before hanging up** is set to a high time-out number, there is no guarantee that the connection will not disconnect before the time passes. This is because the server may choose to disconnect you before the time set here.
- If you are using Internet Explorer, the browser's own settings for dial-up connections may conflict with these redial options.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring redial options
Start Remote Access Auto Connection Manager

# Assign multiple phone numbers to a connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign multiple phone numbers to a connection**

1. Open Network Connections.
2. Right-click the dial-up connection to which you want to assign multiple phone numbers, and then click **Properties**.
3. On the **General** tab, click **Alternates**.
4. If you want the connection to attempt a different phone number if the first number in the list is not successful, select the **If number fails, try next number** check box.
5. If you want the connection to move the first successful phone number to the top of the list under **Phone numbers** and use it on the next connection attempt, click **Move successful number to top of list**.
6. To add a new phone number to the list, click **Add**, and in **Phone number**, type the phone number.
7. If you want the connection to use established area code and dialing rules, select the **Use dialing rules** check box.
8. When you are finished adding the new phone number, click **OK**.
9. If you want to change the order of phone numbers under **Phone numbers**, click a phone number, and then click the **Up** button or **Down** button until the phone number is in the correct order.
10. If you want to remove a phone number under **Phone numbers**, click the phone number, and then click **Delete**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you have alternate phone numbers defined for a connection, you can specify a particular phone number at the time you initiate your connection.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Configure Dial-up Options
Configuring dial-up options

# Enable operator-assisted calls or manual dialing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable operator-assisted calls or manual dialing**

1. Open Network Connections.
2. On the **Advanced** menu, click **Operator-Assisted Dialing**.
3. Double-click the connection you want to dial.
4. Pick up the telephone handset, and then dial the number or ask the operator to dial it for you.

   The number assigned to the entry is displayed in the dialog box for easy reference.
5. Immediately after you have finished dialing, click **Dial**.
6. Hang up the handset only after the modem takes control of the line, which is typically signaled with a click followed by silence.

   It is always safe to replace the handset once Network Connections begins verifying your user name and password. The status message will remind you of this.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If operator-assisted dialing is enabled, a check mark appears next to **Operator-Assisted Dialing** on the **Advanced** menu.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Make Internet Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Make Internet connections**

- Make an Internet connection
- Make a broadband connection using Point-to-Point Protocol over Ethernet (PPPoE)

# Make an Internet connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make an Internet connection**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Connect to the Internet**, and then click **Next**.
4. Click **Connect using a dial-up modem**, and click **Next**.
5. If prompted, select the check box for the modem you want to use for this connection, and click **Next**.
6. Type the name of your ISP in **ISP Name**, and click **Next**.
7. Type the phone number of your ISP, including the area code, in **Phone number**, and click **Next**.
8. Make the connection available to any user or only to yourself by selecting **Anyone's use** or **My use only**.
9. In **User name**, type the account name your ISP assigned to you. In **Password** and **Confirm password**, type the password for your account name, and click **Next**.
10. If you want to create a shortcut for this connection on your desktop, select the **Add a shortcut to this connection to my desktop** check box.
11. Click **Finish.**

## Notes

- To make a connection available to any user, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To enable automatic dialing for all users, create a connection, save your user name and password for all users, and then enable Internet Connection Sharing for the new connection.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Internet Connections
Troubleshooting network and dial-up connections
Make a virtual private network (VPN) connection

# Make a broadband connection using Point-to-Point Protocol over Ethernet (PPPoE)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a broadband connection using Point-to-Point Protocol over Ethernet (PPPoE)**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Connect to the Internet**, and then click **Next**.
4. To create a high-speed connection that requires you to sign in, click **Connect using a broadband connection that requires a user name and password**, and click **Next**.
5. Type the name of your ISP in **ISP Name**, and click **Next**.
6. Do one of the following:
   - If you want this connection to be made available to all users on your network, click **Anyone's use**, and then click **Next**.
   - If you want to reserve the connection for your own use, click **My use only**, and then click **Next**.
7. In **User name**, type the account name your ISP assigned to you. In **Password** and **Confirm password**, type the password for your account name, and click **Next**.
8. To create a shortcut for this connection on your desktop, select the **Add a shortcut to this connection to my desktop** check box.
9. Click **Finish**.

## Notes

- To make a connection available to any user, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Point-to-Point Protocol over Ethernet (PPPoE) Protocols
Point-to-Point Protocol (PPP)
Make Internet Connections

# Make Direct Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Make direct network connections**

- Make a direct network connection
- Make a direct network connection by using a DirectParallel cable

# Make a direct network connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a direct network connection**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Set up an advanced connection**, and then click **Next**.
4. Click **Connect directly to another computer**, click **Next**, and then follow the instructions in the New Connection Wizard.

## Notes

- To create a direct network connection that acts as a host, you must be logged on as Administrator or be a member of the Administrators group. Guest direct network connections do not require administrator credentials.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you specify your connection as a host when you create it, the connection appears as **Incoming Connections** in the Network Connections folder.
- You can create multiple direct connections by copying them in the Network Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple ports, host computers, and so on.
- Direct connections can bypass authentication requirements. This is useful for devices such as handheld PCs. You must configure this setting in the host incoming connection. For more information, see Related Topics.
- If you create a direct connection by using a serial (RS-232C) cable, the port that you select in the New Connection Wizard is enabled for connections that use a null modem.
- If you are logged on to your computer as Administrator or a member of the Administrators group when you create a direct connection, you are presented with a list of connection devices to choose from that includes all of the parallel ports for the computer, infrared ports that are installed and enabled, and COM ports. If you need to use a COM port for a direct connection, ask your system administrator to configure one of the COM ports on your computer with a null modem by using Phone and Modem Options in Control Panel.
- Users do not need to use direct connections to allow access to shared resources, such as files and printers, over a local area network. In order to enable shared access to resources on the local computer, you must enable file and print sharing, share the resources, and then set up the appropriate permissions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using direct connections
Grant incoming connection access rights to your computer

Share a folder or drive

# Make a direct network connection by using a DirectParallel cable

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a direct network connection by using a DirectParallel cable**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Set up an advanced connection**, and then click **Next**.
4. Click **Connect directly to another computer**, click **Next**, and then do one of the following:
   - If you want your computer to act as the host, click **Host**, and click **Next**. Click **Direct Parallel (LPT1)** in the **Device for this connection**, and click **Next**.
   - If you want your computer to act as the guest, click **Guest**, and click **Next**. Type the name of the other computer in **Computer Name**, and click **Next**.
5. Follow the remaining instructions in the New Connection Wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- For information about contacting Parallel Technologies or their local dealers to order DirectParallel cables, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DirectParallel cable access
To contact Parallel Technologies

# Set up Incoming Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up incoming connections**

- Accept incoming network connections
- Grant incoming connection access rights to your computer
- Configure an incoming connection to use TCP/IP
- Configure callback for incoming connections

# Accept incoming network connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To accept incoming network connections**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. Click **Set up an advanced connection**, and then click **Next**.
4. Click **Accept incoming connections**, click **Next**, and then follow the instructions in the New Connection Wizard.

## Notes

- To create an incoming network connection, you must be a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you make another incoming network connection, and you use the New Connection Wizard again, the existing incoming network connection is reconfigured.
- For large numbers of incoming connections on a server that operates as part of a distributed network or as a domain controller, use Routing and Remote Access to create a remote access server.
- If your incoming connection and Fax Service have problems working together (for example, you cannot receive incoming connection calls on a device enabled for Fax receive), the modem may not support adaptive answer. Check your modem documentation to verify that you need to disable Fax receive for that device to accept incoming connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Incoming connections
Incoming connection clients
Configure an incoming connection to use TCP/IP
Grant incoming connection access rights to your computer

# Grant incoming connection access rights to your computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To grant incoming connection access rights to your computer**

1. Open Network Connections.
2. Right-click **Incoming Connections**, and then click **Properties**.
3. On the **General** tab, do one or more of the following:
   - In **Devices**, click the check boxes next to the devices through which incoming connections can connect.
   - If you have more than one device selected, and you want to enable Multilink dialing, select the **Enable multilink** check box.
   - If you want to enable others to make private connections to your computer by tunneling through the Internet or other network, select the **Allow others to make private connections to my computer by tunneling through the Internet or other network** check box.
   - If you want to show icons on the taskbar when incoming connections are connected, select the **Show icon in notification area when connected** check box.
4. On the **Users** tab, do one or more of the following:
   - If you want to allow a user to connect, select the check box next to the user's name.
   - If you want to revoke a user's right to connect, clear the check box next to a user's name.
   - To create a new user and grant them access to connect, click **New**.
   - To remove a user, click **Delete**.
   - To modify a user's name, password, and callback options, click **Properties**.
   - If you want directly connected devices to connect without providing a password, select the **Always allow directly connected devices such as palmtop computers to connect without providing a password** check box.
5. On the **Networking** tab, do one or more of the following:
   - If you want to enable a network component, in **Network components**, select the check box next to the component name.
   - If you want to disable a network component, in **Network components**, clear the check box next to the component name.
   - To add a network component, click **Install**.
   - To remove a network component, click **Uninstall**.

- To configure a network component, click **Properties**.

**Notes**

- To modify an incoming connection, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you enable or disable TCP/IP for incoming connections, the settings affect only incoming connections. Other protocols, such as AppleTalk, are enabled or disabled for all connection types.

Related Topics

Configure callback for incoming connections

Incoming connections

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure an incoming connection to use TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure an incoming connection to use TCP/IP**

1. Open Network Connections.
2. Right-click **Incoming Connections**, and then click **Properties**.
3. On the **Networking** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. If you want incoming users to be able to access the local area network on which the dial-up computer resides, select the **Allow callers to access my local area network** check box.
5. Do one of the following:

   - If you want to automatically assign TCP/IP addresses, select the **Assign TCP/IP addresses automatically using DHCP** check box.
   - If you want to specify TCP/IP addresses, click **Specify TCP/IP addresses**, and do the following:

     In **From**, type the starting IP address.

     In **To**, type the ending IP address.

     Based on the addresses you type in **From** and **To**, the number of addresses that are allocated is displayed in **Total**. The starting and ending IP addresses that you choose depend on the originating address space and the number of addresses needed. For more information on defining IP addresses for incoming connections, see Related Topics.

6. If you want incoming connections to be able to use a specific TCP/IP address, select the **Allow calling computer to specify its own IP address** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Most TCP/IP networks use subnets in order to effectively manage routed IP addresses. The closest matching subnet is calculated for the range that you specified in **From** and **To**. The range of addresses in the closest matching subnet may exceed the range that you specified. Unless the addresses specified in **From** and **To** are subnet boundaries, the range based on the calculated subnet will be larger than the range that you specified. To avoid this, specify a range that falls on subnet boundaries. For example, if you are using the 10.0.0.0 private network ID for your intranet, a range that falls on subnet boundaries is 10.0.1.168 to 10.0.1.175. Or, if you are using the 192.168.0.0 private network ID for your intranet, a range that falls on subnet boundaries is 192.168.1.0 to 192.168.1.255.
- If you select **Allow callers to access my local area network** or **Allow calling computer to specify its own IP address**, these changes take effect immediately.
- If you select **Assign TCP/IP addresses automatically using DHCP**, but there is no DHCP server available, random addresses from the range 169.254.0.1 to 169.254.255.254 are assigned.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Incoming connections and IP addresses
To configure a default gateway

# Configure callback for incoming connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure callback for incoming connections**

1. Open Network Connections.
2. Right-click **Incoming Connections**, and then click **Properties**.
3. On the **Users** tab, click the user for which you want to configure callback, and then click **Properties**.
4. On the **Callback** tab, do one of the following:
   - If you do not want to use callback for incoming connections, click **Do not allow callback**.
   - If you want to give incoming connections the option of setting a specific callback number, click **Allow the caller to set the callback number**.
   - If you want to consistently use the same callback number, click **Always use the following callback number**, and type the number.

**Note**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Grant incoming connection access rights to your computer
Callback

# Configure Network and Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure network and dial-up connections**

- Configure a connection to a remote network
- Set up Network Components
- Secure Network or Dial-up Connections

# Configure a connection to a remote network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a connection to a remote network**

1. Open Network Connections.
2. Right-click the dial-up connection that you want to configure, and then click **Properties**.
3. Do one or more of the following:
   - To configure dialing devices, phone numbers, host address, country/region codes, or dialing rules, click the **General** tab.
   - To configure dialing and redialing options, multilink configuration, or X.25 parameters, click the **Options** tab.
   - To configure identity authentication, data encryption, or terminal window and scripting options, click the **Security** tab.
   - To configure the remote access server and protocols used for this connection, click the **Networking** tab.
   - To enable or disable Internet Connection Firewall, Windows Firewall, Internet Connection Sharing, and on-demand dialing, click the **Advanced** tab. For more information, see Related Topics.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- Depending on the type of connection you are configuring, different options appear in the properties dialog box.
- For more information about a specific item on a tab, right-click the item, and then click **What's This?**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring a network or dial-up connection
Set up Network Components
Secure Network or Dial-up Connections
Configuring Authentication and Data Encryption

Enable ICS
Enable or disable ICF
Enable ICS on-demand dialing

# Set up Network Components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up network components**

- Configure TCP/IP settings
- Add a network component
- Remove a network component
- Enable a network component
- Disable a network component
- Modify the protocol bindings order
- Modify the network provider order
- Configure Client for Microsoft Networks
- Configure File and Printer Sharing for Microsoft Networks
- PPP Settings

# Configure TCP/IP settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure TCP/IP settings**

1. Open Network Connections.
2. Right-click the connection you want to configure, and then click **Properties**.
3. Do one of the following:
   - If the connection is a local area connection, on the **General** tab, in **This connection uses the following items**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
   - If the connection is a dial-up or VPN connection, on the **Networking** tab, in **This connection uses the following items**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
   - If the connection is an incoming connection, on the **Networking** tab, in **Network components**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Do one of the following:
   - If you want IP settings to be assigned automatically, click **Obtain an IP address automatically**, and then click **OK**.
   - If you want to specify an IP address or a DNS server address, do the following:

     Click **Use the following IP address**, and in **IP address**, type the IP address.

     Click **Use the following DNS server addresses**, and in **Preferred DNS server** and **Alternate DNS server**, type the addresses of the primary and secondary DNS servers.
5. To configure DNS, WINS, and IP settings, click **Advanced**.
6. In a local area connection, selecting the **Obtain an IP address automatically** option enables the **Alternate Configuration** tab. Use this to enter alternate IP settings if your computer is used on more than one network. To configure DNS, WINS, and IP settings, click **User configured** on the **Alternate Configuration** tab.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You should use automated IP settings (DHCP) whenever possible for all connections because they eliminate the need to configure settings such as IP address, DNS server address, and WINS server address.
- The **Alternate Configuration** settings specify a second set of IP settings that are used when a DHCP server is not available. This is very useful for laptop users who often switch between two different network environments such as DHCP and static IP network environments.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
TCP/IP
New ways to do Network Connections tasks

# Add a network component

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a network component**

1. Open Network Connections.
2. Right-click the connection to which you want to add a network component, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, click **Install**.
   - If this is a dial-up, VPN, or incoming connection, on the **Networking** tab, click **Install**.
4. In the **Select Network Component Type** dialog box, click either **Client**, **Service**, or **Protocol**, and then click **Add**.
5. Do one of the following:
   - If you do not have an installation disk for the component, click the appropriate client, service, or protocol, and then click **OK**.
   - If you have an installation disk for the component, click the appropriate client, service, or protocol, click **Have Disk**, insert the installation disk into the selected drive, and then click **OK**.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You should only install the network components that you need, for the following reasons:
  - Network performance is enhanced and network traffic is reduced when only the required protocols and clients are installed.
  - If a problem with a network or dial-up connection occurs, your computer attempts to establish connectivity by using every network protocol that is installed and enabled. By only installing and enabling the protocols that your system can use, the operating system does not attempt to connect with protocols it cannot use, and returns status information to you more efficiently.
  - Excessive services can hinder performance on your local computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Remove a network component
Disable a network component
Enable a network component

## Add optional networking components

# Remove a network component

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a network component**

1. Open Network Connections.
2. Right-click the connection on which you want to remove a network component, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, click the client, service, or protocol you want to remove, and then click **Uninstall**.
   - If this is a dial-up or VPN connection, on the **Networking** tab, in **This connection uses the following items**, click the client, service, or protocol you want to remove, and then click **Uninstall**.
   - If this is an incoming connection, on the **Networking** tab, in **Network components**, click the client, service, or protocol you want to remove, and then click **Uninstall**.
4. In the **Uninstall***component name* dialog box, click **Yes**.

**Note**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Add a network component
Enable a network component
Disable a network component

# Enable a network component

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable a network component**

1. Open Network Connections.
2. Right-click a connection, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, select the check box next to the client, service, or protocol you want to enable.
   - If this is a dial-up or VPN connection, on the **Networking** tab, in **This connection uses the following items**, select the check box next to the client, service, or protocol you want to enable.
   - If this is an incoming connection, on the **Networking** tab, in **Network components**, select the check box next to the client, service, or protocol you want to enable.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You should only enable the network components that you need, for the following reasons:
  - Network performance is enhanced and network traffic is reduced when only the required protocols and clients are enabled.
  - If your computer encounters a problem with a network or dial-up connection, it attempts to establish connectivity by using every network protocol that is installed and enabled. By only enabling the protocols that your system can use, the operating system does not attempt to connect with protocols it cannot use, and returns status information to you more efficiently.
  - Excessive services can hinder performance on your local computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Disable a network component

# Disable a network component

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable a network component**

1. Open Network Connections.
2. Right-click the connection on which you want to disable a network component, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, clear the check box next to the client, service, or protocol you want to disable.
   - If this is a dial-up, VPN, or incoming connection, on the **Networking** tab, in **Use these components with this connection**, clear the check box next to the client, service, or protocol you want to disable.

**Note**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Enable a network component

# Modify the protocol bindings order

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify the protocol bindings order**

1. Open Network Connections.
2. On the **Advanced** menu, click **Advanced Settings**.
3. In **Connections**, click the connection that you want to modify.
4. Click the protocol that you want to move up or down in the list, and then click the up or down arrow button.

## Notes

- To modify the protocol bindings order, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Configuring a network or dial-up connection
Set up Network Components

# Modify the network provider order

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify the network provider order**

1. Open Network Connections.
2. Click the connection you want to modify, and on the **Advanced** menu, click **Advanced Settings**.
3. On the **Provider Order** tab, in **Network Providers**, click the provider you want to move up or down in the list, and then click the **Up** or **Down** button.

## Notes

- To modify the network provider order, you must be logged on as a member of the Administrators group.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Network providers are accessed in the order they appear in this list.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Connections Best Practices
Set up Network Components
Configuring a network or dial-up connection

# Configure Client for Microsoft Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Client for Microsoft Networks**

1. Open Network Connections.
2. Right-click a connection, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, click **Client for Microsoft Networks**, and then click **Properties**.
   - If this is a dial-up, VPN, or incoming connection, on the **Networking** tab, in **Use these components with this connection**, click **Client for Microsoft Networks**, and then click **Properties**.
4. In **Name service provider**, click an RPC name service provider.
5. If you selected **DCE Cell Directory Service**, in **Network address** type the network address of your provider.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Windows Locator is the default name service provider.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Client for Microsoft Networks
Set up Network Components

# Configure File and Printer Sharing for Microsoft Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure File and Printer Sharing for Microsoft Networks**

1. Open Network Connections.
2. Right-click a connection, and then click **Properties**.
3. Do one of the following:
   - If this is a local area connection, on the **General** tab, in **This connection uses the following items**, click **File and Printer Sharing for Microsoft Networks**, and then click **Properties**.
   - If the connection is a dial-up or VPN connection, on the **Networking** tab, in **This connection uses the following items**, click **File and Printer Sharing for Microsoft Networks**, and then click **Properties**.
   - If the connection is a incoming connection, on the **Networking** tab, in **Network components**, click **File and Printer Sharing for Microsoft Networks**, and then click **Properties**.
4. Do one of the following:
   - To optimize the server for a small number of clients, click **Minimize memory used**.
   - To optimize the server for a mixed usage of file and printer sharing in addition to other services, such as serving as an interactive workstation, click **Balance**.
   - To dedicate as many resources as possible to file and print server services, click **Maximize data throughput for file sharing**.
   - To optimize server memory for distributed applications that perform their own memory caching, such as Microsoft SQL Server, click **Maximize data throughput for network applications**.
5. To specify that LAN Manager 2.*x* clients on your network can browse for shared resources on this computer, select the **Make browser broadcasts to LAN Manager 2.x clients** check box.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You can only configure File and Printer Sharing for Microsoft Networks on a server.
- To share local folders, you must enable File and Printer Sharing for Microsoft Networks. For more information, see Related Topics.
- The File and Printer Sharing for Microsoft Networks component is the equivalent of the Server service in Windows NT 4.0.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
File and Printer Sharing for Microsoft Networks
Share a folder or drive

# PPP Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPP settings**

- Stop requesting LCP extensions in PPP
- Start requesting LCP extensions in PPP
- Turn off IP header compression in PPP
- Disable software compression

# Stop requesting LCP extensions in PPP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop requesting LCP extensions in PPP**

1. Disconnect from the remote server.
2. Open Network Connections.
3. Right-click the dial-up connection you want to configure, and then click **Properties**.
4. On the **Networking** tab, click **Settings**.
5. In the **PPP Settings** box, clear the **Enable LCP extensions** check box.

### Note

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting network and dial-up connections
Start requesting LCP extensions in PPP
PPP operation and protocols

# Start requesting LCP extensions in PPP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start requesting LCP extensions in PPP**

1. Disconnect from the remote server.
2. Open Network Connections.
3. Right-click the dial-up connection you want to configure, and then click **Properties**.
4. On the **Networking** tab, click **Settings**, and then select the **Enable LCP extensions** check box.

### Note

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting network and dial-up connections
Stop requesting LCP extensions in PPP
PPP operation and protocols

# Turn off IP header compression in PPP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To turn off IP header compression in PPP**

1. Disconnect from the remote server.
2. Open Network Connections.
3. Right-click the dial-up connection you want to configure, and then click **Properties**.
4. On the **Networking** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
5. Click **Advanced**, and then clear the **Use IP header compression** check box.

## Notes

- To turn off IP header compression, you must be logged on as an administrator or a member of the Administrators group. Some connections do not display the **IP header compression** check box.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting network and dial-up connections
ISP access methods
PPP Settings

# Disable software compression

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable software compression**

1. Open Network Connections.
2. Right-click the dial-up, VPN, or incoming connection you want to configure, and then click **Properties**.
3. On the **Networking** tab, click **Settings**, and clear the **Enable software compression** check box.

### Note

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting network and dial-up connections
Start requesting LCP extensions in PPP
Stop requesting LCP extensions in PPP
Using phone lines and modems

# Secure Network or Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Secure network or dial-up connections**

- Configure identity authentication and data encryption settings
- Enable smart card or other certificate authentication
- Terminal and Scripting Options

For information on security considerations for routing, remote access and VPN, see Security information for routing, Security information for remote access and Security issues for VPN.

# Configure identity authentication and data encryption settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure identity authentication and data encryption settings**

1. Open Network Connections.
2. Right-click the dial-up connection you want to configure, and then click **Properties**.
3. On the **Security** tab, do one of the following:

   - To select preconfigured combinations of identity authentication methods and data encryption requirements, click **Typical (recommended settings)**, and in **Validate my identity as follows**, click a method to use for validation.

     Depending on your **Validate my identity as follows** selection, you can select (enable) or clear (disable) the **Automatically use my Windows logon name and password (and domain if any)** or **Require data encryption (disconnect if none)** check boxes. The following table shows which selections are available.

     If **Require data encryption (disconnect if none)** is not enabled, encryption is optional. To prevent encryption, click **Advanced (custom settings)**, and then click **Settings**. In **Data encryption**, click **No encryption allowed (server will disconnect if it requires encryption)**.

     For dial-up connections:

     | Validate my identity as follows | Automatically use my Windows logon name and password (and domain if any) | Require data encryption (disconnect if none) |
     |---|---|---|
     | Allow unsecured password | Unavailable | Unavailable |
     | Require secured password | Available | Available |
     | Use smart card | Unavailable | Available |

     For VPN connections:

     | Validate my identity as follows | Automatically use my Windows logon name and password (and domain if any) | Require data encryption (disconnect if none) |
     |---|---|---|
     | Require secured password | Available | Available (enabled by default) |
     | Use smart card | Unavailable | Available (enabled by default) |

   - To individually enable, configure, and disable authentication methods and encryption requirements, click **Advanced (custom settings)**, and then click **Settings**.

**Important**

- Modifying **Advanced (custom settings)** requires a knowledge of security protocols.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- Among the advanced custom settings is the **Allow older MS-CHAP version for Windows 95 servers** check box. You must select this check box to configure a connection to a server running Windows 95.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dial-up connection authentication and data encryption
VPN connection authentication and data encryption

# Enable smart card or other certificate authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable smart card or other certificate authentication**

1. Open Network Connections.
2. Right-click the dial-up, VPN, or incoming connection on which you want to use smart card or other certificate authentication, and then click **Properties**.
3. If you are using typical settings for your smart card, on the **Security** tab, click **Typical (recommended settings)**, and in the **Validate my identity as follows** list, click **Use smart card**, and then click **OK**.
4. If you are individually enabling, configuring, and disabling authentication methods and encryption requirements, on the **Security** tab, click **Advanced (custom settings)**, and click **Settings**. In **Logon security**, click **Use Extensible Authentication Protocol (EAP)**, click **Smart card or other certificate (encryption enabled)**, click **Properties**, and then do the following:
   - If you want to use the certificate that resides on your smart card, click **Use my smart card**.
   - If you want to use the certificate that resides in the certificate store on your computer, click **Use a certificate on this computer**.
   - If you want to verify that the server certificate presented to your computer has not expired, has the correct signature, and has a trusted root certification authority, select the **Validate server certificate** check box.
   - If you only want to connect to servers within a particular domain, select the **Connect only if server name ends with** check box, and then type the name of the domain.
   - To specify that the root certification authority for your server certificate must be in a particular root certification authority, in **Trusted root certification authority**, click the appropriate certification authority.
   - To use a different user name when the user name in the smart card or certificate is not the same as the user name in the domain that you are logging on to, select the **Use a different user name for the connection** check box.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If, for example, you only want to connect to servers that reside in Microsoft.com, you type **Microsoft.com** in **Connect only if server name ends in**.
- If, for example, you are working for a consulting company where you need to log on to the domain of the company to which you are assigned, but your smart card contains a user name specific to your home company, select the **Use a different user name for the connection** check box.
- If you select the **Use a different user name for the connection** check box, your certificate is exported without private keys and submitted to your system administrator to be explicitly mapped to your domain user account. For more information, see Related Topics.
- If you select the **Connect only if server name ends with** check box, and do not type a domain name, at the time you connect you are prompted to use the domain name in the server certificate.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Smart card and other certificate authentication
Mobile users and certificates
Map a certificate to a user account

# Terminal and Scripting Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Terminal and scripting options**

- Use the Terminal feature to log on to a remote computer
- Activate a logon script
- Log and view modem commands

# Use the Terminal feature to log on to a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use the Terminal feature to log on to a remote computer**

1. Open Network Connections.
2. Right-click the dial-up connection on which you want to use Terminal, and then click **Properties**.
3. On the **Security** tab, select the **Show terminal window** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- After you connect, the After Dial Terminal window appears and prompts you from the remote computer display. You then use this window to log on to the remote computer. After you have completed all interactions with the remote computer, click **Done**. At this point, authentication on the remote access server begins.
- If the logon sequence does not vary, you can write a script that automatically passes information to the remote computer during the logon sequence, which enables completely automatic connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using the Terminal feature for the remote logon process
Automating the Dial-up Logon Process by Using Switch.inf Scripts
Activate a logon script
Log and view modem commands

# Activate a logon script

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To activate a logon script**

1. Open Network Connections.
2. Right-click the dial-up connection on which you want to run a script, and then click **Properties**.
3. On the **Security** tab, select the **Run script** check box.
4. To use an existing script file, click **Browse**.

   Or, to use Switch.inf, click **Edit**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- In addition to Switch.inf, the Windows 95 scripting language is supported for computers running the Windows Server 2003 family that connect to computers running operating systems other than Windows. You may find the Windows 95 scripting language easier to use than Switch.inf. For more information, see Related Topics.
- If you click **Edit** and load Switch.inf into Notepad, immediately resave the file with a different name to preserve the default Switch.inf file.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Creating Switch.inf Scripts for Network and Dial-up Connections
Log and view modem commands
Troubleshooting logon scripts

# Log and view modem commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To log and view modem commands**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem whose commands you want to log.
3. Click **Properties**, then click the **Diagnostics** tab.
4. In **Logging**, select the **Append to Log** check box if you are working on a computer running Windows XP or **Record a Log** check box if you are working on a computer running a member of the Windows Server 2003 family.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Click **View Log** to display the log file. This file identifies the modem and displays information about how it is functioning.
- Commands sent to the modem are captured in the file:

  *systemroot*\System32\ModemLog_*Model*.txt.

  The path is usually C:\Windows\System32\. *Model* is the name of the modem as it appears in the list of installed modems on the **Modems** tab of **Phone and Modem Options**.
- In Windows XP, logging is always turned on, and the log is overwritten at the beginning of every session unless you select the **Append to Log** check box. In the Windows Server 2003 family, logging is turned off unless you select the **Record a Log** check box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modem diagnostics
Test a modem
Log and view modem commands
Troubleshooting modems

# Monitor Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor network connections**

- Monitor a connection
- View the status of a local area connection

# Monitor a connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To monitor a connection**

1. Open Network Connections.
2. Do one of the following:
   - To monitor activity for the current session, right-click the connection you want to monitor, and then click **Status**.
   - To automatically enable the status monitor each time the connection is active, right-click the connection, click **Properties**, and then select the **Show icon in taskbar when connected** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If your computer is set up to accept incoming connections, a connection with a user name assigned to it appears in the Network Connections folder as each user connects. You can view the progress of incoming connections by right-clicking a named connection, and then clicking **Status**.
- You can also monitor a connection by clicking **Status** on the **File** menu or by clicking the status icon in the taskbar. You can also right-click the taskbar icon and click **Status**.
- If you want to view the general status (**Disconnected**, **Connected**) of all of your connections, on the **View** menu, click **Details**.
- You also can disconnect an active connection by clicking **Disconnect** in the **Status** dialog box.
- By default, the status monitor is disabled for local area connections, but enabled for all other types of connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitoring network and dial-up connections
Troubleshooting network and dial-up connections
Log and view modem commands
Configure a connection to a remote network

# View the status of a local area connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the status of a local area connection**

1. Open Network Connections.
2. Do one of the following:
   - To monitor activity for the current session, right-click the local area connection, and then click **Status**.
   - To automatically enable the status monitor each time the connection is active, right-click the local area connection, click **Properties**, and then select the **Show icon in taskbar when connected** check box.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You can view and modify the network components of a local area connection, such as network protocol settings, by clicking **Properties**.
- By default, the status monitor is disabled for local area connections, but enabled for all other types of connections.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using local area connections
Make a local area connection

# Set up and Manage Home and Small Office Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up and manage home and small office networking**

- Set up and Manage Internet Connection Firewall (ICF)
- Set up and Manage Network Bridge
- Set up and Manage Internet Connection Sharing (ICS)
- Manage service definitions for ICF or ICS

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Set up and Manage Internet Connection Firewall (ICF)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up and manage Internet Connection Firewall (ICF)**

- Enable or disable ICF
- Enable or disable Internet Control Message Protocol requests for ICF
- Enable or disable ICF security log options
- View the ICF security log
- Change the file size of the ICF security log
- Change the path and file name of the ICF security log
- Restore the default settings to the ICF security log

# Enable or disable ICF

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable ICF**

1. Open Network Connections.
2. Under **Dial-up** or **LAN or High-Speed Internet**, right-click the connection that you want to protect, and then click **Properties**.
3. On the **Advanced** tab, under **Internet Connection Firewall**, do one of the following:
   - To enable Internet Connection Firewall (ICF), select the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box.
   - To disable ICF, clear the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box.

**Caution**

- If you clear the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box, the firewall on this connection is disabled. When the firewall is disabled, your computer is exposed to possible external attacks through the unprotected connection.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You need to enable ICF only on connections that connect to the Internet. You should not enable ICF on virtual private network (VPN) connections or on network client computers because ICF interferes with network activities such as file and printer sharing.
- If ICF is enabled on two or more connections on a single computer, the ICF settings are global. If you select a setting or change a setting in **Logging Options** or **Services** on any one of the connections on which ICF is enabled, that setting is applied to all of the ICF firewalls on that computer.
- ICF cannot be enabled on the private connections of the Internet Connection Sharing host computer, on Network Bridge, or on the adapters that are included in Network Bridge.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- This topic applies only to product features available in the original release of Windows Server 2003.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Home and small office networking features
Internet Connection Firewall
Internet Connection Firewall security log
Service definitions for Internet Connection Firewall and Internet Connection Sharing

# Enable or disable Internet Control Message Protocol requests for ICF

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable Internet Control Message Protocol requests for ICF**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **ICMP** tab, do one of the following:
   - To enable Internet Control Message Protocol (ICMP) options, select the check box next to each type of request for information to which you would like your computer to respond.
   - To disable ICMP options, clear any or all of the ICMP check boxes.

### Caution

- All ICMP options are disabled by default. If you enable ICMP options on a computer running ICF, that computer may become visible to the Internet and vulnerable to attacks.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable ICF
Home and small office networking features
Internet Connection Firewall

# Enable or disable ICF security log options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable ICF security log options**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, under **Logging Options**, select from the following:

   To enable ICF security log options:

   - To enable the logging of all dropped packets that originate from either the home or small office network or the Internet, select the **Log dropped packets** check box.
   - To enable the logging of all successful connections that originate from either the home or small office network or the Internet, select the **Log successful connections** check box.

   Or, to disable ICF security log options:

   - To disable the logging of all dropped packets that originate from either the home or small office network or the Internet, clear the **Log dropped packets** check box.
   - To disable the logging of all successful connections that originate from either the home or small office network or the Internet, clear the **Log successful connections** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Security logging is not enabled by default when ICF is enabled.
- ICF functions whether security logging is enabled or disabled.
- If ICF is enabled on two or more connections on a single computer, the ICF settings are global. If you select a setting or change a setting in **Logging Options** or **Services** on any one of the connections on which ICF is enabled, that setting is applied to all of the ICF firewalls on that computer.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall security log
Internet Connection Firewall
Enable or disable ICF

# View the ICF security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the ICF security log**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, under **Log file options**, in **Name**, click **Browse**.
5. Scroll to the ICF security log, right-click the file, and then click **Open**.

## Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The default name for the ICF security log is pfirewall.log, and the default location is in the Windows directory.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall
Internet Connection Firewall security log
Enable or disable ICF
Enable or disable ICF security log options

# Change the file size of the ICF security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the file size of the ICF security log**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, under **Log file options**, in **Size limit**, use the arrow buttons to adjust the size limit.

### Caution

- If the maximum allowable size for pfirewall.log is exceeded, the information that pfirewall.log contains is saved as pfirewall.log.old, and new information is saved as pfirewall.log. If the maximum allowable size for pfirewall.log is exceeded again, the information that pfirewall.log contains is saved as pfirewall.log.old, and the information that had been in pfirewall.log.old is overwritten. New information is always saved in pfirewall.log.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If ICF is enabled on two or more connections on a single computer, the ICF settings are global. If you select a setting or change a setting in **Logging Options** or **Services** on any one of the connections on which ICF is enabled, that setting is applied to all of the ICF firewalls on that computer.
- The default size limit of the ICF security log is 4,096 kilobytes.
- You must enable ICF before you can adjust the size limit of the ICF security log.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall security log
Enable or disable ICF
Internet Connection Firewall

# Change the path and file name of the ICF security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the path and file name of the ICF security log**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, under **Log file options**, click **Browse**, and then do the following:
   - To change the path of the ICF security log, navigate to the location where you want the ICF security log to be saved. (If you do not want to change the path, do not navigate to a new location.)
   - To change the file name of the ICF security log, in **File name**, type a new name, and then click **Open**. (If you do not want to change the default file name, in **File name**, type **pfirewall**, and then click **Open**.)

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The default name for the ICF security log is pfirewall.log, and the default location is in the Windows directory.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall security log
Enable or disable ICF
Enable or disable ICF security log options

# Restore the default settings to the ICF security log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore the default settings to the ICF security log**

1. Open Network Connections.
2. Right-click the connection on which Internet Connection Firewall (ICF) has been enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Security Logging** tab, click **Restore Defaults**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The default size limit of the log is 4,096 kilobytes.
- The default name for the ICF security log is pfirewall.log, and the default location is in the Windows directory.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Internet Connection Firewall security log
Home and small office networking features
Internet Connection Firewall
Enable or disable ICF

# Set up and Manage Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up and manage Network Bridge**

- Set up Network Bridge
- Add a connection to or remove a connection from Network Bridge
- Enable or disable Network Bridge
- Remove Network Bridge

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Set up Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up Network Bridge**

1. Open Network Connections.
2. Under **LAN or High-Speed Internet**, hold down the CTRL key, and then click each of the adapters that you want to include in the bridge.
3. Right-click one of the connections that you have selected, and then click **Bridge Connections**.

**Caution**

- If Internet Connection Firewall (ICF), Windows Firewall, or Internet Connection Sharing (ICS) are not enabled on your network, do not set up Network Bridge between the public Internet connection and the private network connection. Setting up Network Bridge between the public Internet connection and the private network connection creates an unprotected link between your network and the Internet, leaving your network vulnerable to external attacks. When ICF, Windows Firewall, or ICS are enabled, this risk is mitigated.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Connections that run ICF, Windows Firewall, or adapters that provide shared Internet connections cannot be a part of the Network Bridge and do not appear under **LAN or High-Speed Internet**. Similarly, the **Bridge Connections** menu command is available only for adapters that you can add as connections to Network Bridge.
- You must select at least two private network connections in order for **Bridge Connections** to create a bridge.
- You can create only one Network Bridge on a computer, but the bridge can accommodate numerous network connections.
- Network adapters that previously appeared under **LAN or High-Speed Internet**, appear under **Network Bridge** when you add them as connections to the bridge. Network adapters that provide Internet connectivity, such as DSL and cable modems, cannot be bridged and always appear under **LAN or High-Speed Internet**.
- Only Ethernet, IEEE-1394 adapters, or Ethernet-compatible adapters, such as wireless and Home Phoneline Network adapters, can be a part of Network Bridge.
- Network Bridge with a wireless or IEEE-1394 connection supports traffic using Internet Protocol version 4 (IPv4) only.
- You cannot create Network Bridge on a computer running Windows 2000 or earlier versions of Windows.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a connection to or remove a connection from Network Bridge
Enable or disable Network Bridge
Remove Network Bridge
Home and small office networking features
Network Bridge

# Add a connection to or remove a connection from Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a connection to or remove a connection from Network Bridge**

1. Open Network Connections.
2. Under **Network Bridge**, right-click **Network Bridge**, and then click **Properties**.
3. On the **General** tab, under **Adapters**, do one of the following.
   - To add a connection to Network Bridge, select the check box next to each adapter you want to add to the bridge, and then click **OK**.
   - To remove a connection from Network Bridge, clear the check box next to each adapter you want to remove from the bridge, and then click **OK**.

**Caution**

- If Internet Connection Firewall (ICF), Windows Firewall, or Internet Connection Sharing (ICS) are not enabled on your network, do not set up Network Bridge between the public Internet connection and the private network connection. Setting up Network Bridge between the public Internet connection and the private network connection creates an unprotected link between your network and the Internet, leaving your network vulnerable to external attacks. When ICF, Windows Firewall, or ICS are enabled, this risk is mitigated.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you remove network adapters from Network Bridge and less than two adapters remain, Network Bridge does not perform its intended function but continues to use system resources.
- You must set up Network Bridge before you can add connections to it. For more information, see Related Topics.
- Connections that run ICF, Windows Firewall, or adapters that provide shared Internet connections cannot be a part of Network Bridge and do not appear under **Adapters**. Similarly, the **Add to Bridge** menu command is available only for adapters that you can add as connections to Network Bridge.
- Network adapters that previously appeared under **LAN or High-Speed Internet**, appear under **Network Bridge** when you add them as connections to the bridge. Network adapters that provide Internet connectivity, such as DSL and cable modems, cannot be bridged and always appear under **LAN or High-Speed Internet**.
- Only Ethernet, IEEE-1394 adapters, or Ethernet-compatible adapters, such as wireless and Home Phoneline Network adapters, can be a part of Network Bridge.
- Network Bridge with a wireless or IEEE-1394 connection supports traffic using Internet Protocol version 4 (IPv4) only.

- You can also add a connection to or remove a connection from Network Bridge by opening Network Connections, right-clicking the adapter that you want to add, and then clicking **Add to Bridge**.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up Network Bridge
Enable or disable Network Bridge
Remove Network Bridge
Home and small office networking features
Network Bridge

# Enable or disable Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable Network Bridge**

1. Open Network Connections.
2. Under **Network Bridge**, right-click **Network Bridge**, and do one of the following:
   - To enable Network Bridge, click **Enable**.
   - To disable Network Bridge, click **Disable**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you remove or disable Network Bridge, computers on your network are unable to communicate with computers that are on different network segments. Some of the computers on your network may lose Internet connectivity.
- You must set up Network Bridge before you can enable or disable it. For more information, see Related Topics.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up Network Bridge
Add a connection to or remove a connection from Network Bridge
Remove Network Bridge
Home and small office networking features
Network Bridge

# Remove Network Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove Network Bridge**

1. Open Network Connections.
2. Under **Network Bridge**, right-click **Network Bridge**, and then click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- If you remove or disable Network Bridge, computers on your network are unable to communicate with computers that are on different network segments. Some of the computers on your network may lose Internet connectivity.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set up Network Bridge
Add a connection to or remove a connection from Network Bridge
Enable or disable Network Bridge
Home and small office networking features
Network Bridge

# Set up and Manage Internet Connection Sharing (ICS)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up and manage Internet Connection Sharing (ICS)**

- Enable ICS
- Enable ICS on-demand dialing
- Configure client computers for ICS

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Enable ICS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable ICS**

1. Open Network Connections.
2. Under **Dial-up** or **LAN or High-Speed Internet**, right-click the connection on which you want to enable (ICS), and then click **Properties**.
3. On the **Advanced** tab, under **Internet Connection Sharing**, select the **Allow other network users to connect through this computer's Internet connection** check box.
4. If present, under **Home networking connection**, in the **Select a private network connection** drop-down box, select the connection that connects directly to your network. For more information, see Notes.

**Important**

- You can enable ICS only when two or more network adapters are installed.
- If your computer has multiple private connections that connect to private LAN segments, you should bridge the private connections before you enable ICS. For information about how to set up Network Bridge, see Related Topics.
- When ICS is enabled, it automatically provides an IP address for the private connection of the ICS host.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- **Home networking connection** and the **Select a private network connection** drop-down box are present only when two or more adapters that can connect to the private network are installed.
- For information about the protocols, services, interfaces, and routes that are automatically configured by ICS when you connect to the Internet in a home or small office network, see Related Topics.
- If you are having difficulty connecting the ICS host to the Internet, you can use **Repair** to obtain refreshed IP address configurations from your Internet service provider (ISP).
- If a computer on your network is unable to communicate with other network computers, you can use **Repair** to force the connection to refresh its IP address configuration from the ICS host.
- Before using **Repair**, ensure that the client computers are plugged into the network and that the ICS host is running, configured, and connected to the network and the Internet.
- For information about how to configure client computers for ICS, see Related Topics.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connecting to the Internet in a home or small office network
Manage service definitions for ICF or ICS
Network Bridge
Configure client computers for ICS
Connecting to the Internet in a home or small office network

# Enable ICS on-demand dialing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable ICS on-demand dialing**

1. Open Network Connections.
2. Right-click the dial-up connection on which you want to enable ICS on-demand dialing, and then click **Properties**.
3. On the **Advanced** tab, under **Internet Connection Sharing**, do one of the following:
   - To enable ICS on-demand dialing, select the **Establish a dial-up connection whenever a computer on my network attempts to access the Internet** check box.
   - To disable ICS on-demand dialing, clear the **Establish a dial-up connection whenever a computer on my network attempts to access the Internet** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- In order for ICS on-demand dialing to work correctly, the shared Internet connection must be available for all ICS clients, and the dial-up credentials must be saved for all users. For more information, see Related Topics.
- **Establish a dial-up connection whenever a computer on my network attempts to access the Internet** is selected by default. Unless you clear the check box, clients can activate the shared Internet connection from their computers by starting their Internet browsers or mail programs.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Troubleshooting demand-dial routing
Dial-up credentials

# Configure client computers for ICS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure client computers for ICS**

1. Insert the Windows Server 2003, Standard Edition or Windows Server 2003, Enterprise Edition CD into the CD-ROM drive of your computer running Windows 98, Windows 98 Second Edition, or Windows Millennium Edition.

2. When the Welcome screen appears, click **Perform additional tasks**, and then click **Browse this CD**.

3. Double-click **Support**, and then double-click **Netsetup**.

4. Follow the instructions that appear on your screen.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- Enable Internet Connection Sharing (ICS) on your ICS host computer before you attempt to configure client computers for your home or small office network. For more information, see Related Topics.

- You can run Netsetup.exe only on computers running Windows 98, Windows 98 Second Edition, or Windows Millennium Edition.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connecting to the Internet in a home or small office network
Enable ICS
Home and Small Office Networking Overview

# Manage service definitions for ICF or ICS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage service definitions for ICF or ICS**

1. Open Network Connections.
2. Right-click the shared Internet connection or the Internet connection on which Internet Connection Firewall (ICF) or Internet Connection Sharing (ICS) is enabled, and then click **Properties**.
3. On the **Advanced** tab, click **Settings**.
4. On the **Services** tab, do one of the following:
   - If you want to enable a service, under **Services** select the check box next to service that you want to enable, and then enter the required information.
   - If you want to add a service definition, click **Add**, and then enter the required information.
   - If you want to edit a service definition, click the service that you want to edit, click **Edit**, and then change the appropriate information.
   - If you want to delete a service definition, click the service that you want to delete, and then click **Delete**.

**Caution**

- This feature is intended for advanced users. When you add or edit a service definition, you create a port mapping that can expose your network to the Internet. For more information about service definitions, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Before you can add, edit, or delete service definitions, you must enable ICF, ICS, or both.
- You can obtain the correct configuration values for the TCP or UDP port number from the service's documentation or Web site.
- No two service definitions can use the same port numbers simultaneously.
- For service definitions that are preconfigured, you can edit only the field that is labeled **Name or IP address of the computer hosting this service on your network**.
- For service definitions that you have added, you can edit only the fields that are labeled **Name or IP address of the computer hosting this service on your network**, **External Port number for this service**, and **Internal Port number for this service**.
- You can delete only the service definitions that you have added. You cannot delete service definitions that are preconfigured.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

- This topic applies only to product features available in the original release of Windows Server 2003.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Service definitions for Internet Connection Firewall and Internet Connection Sharing
Internet Connection Firewall
Enable or disable ICF

# Network Connections Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Network Connections:

- Network Connections Overview
- Understanding Network Connections
- Using Network Connections
- Administering Network Connections

# Network Connections Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Connections overview**
This section covers:

- About Network Connections
- Hardware requirements for network and dial-up connections
- Network connection types
- Dial-up connections
- Network connection icons
- Protocols and Access Methods
- Home and Small Office Networking Overview

# About Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**About Network Connections**

Network Connections provides connectivity between your computer and the Internet, a network, or another computer. With Network Connections, you can gain access to network resources and functionality, whether you are physically located at the location of the network or in a remote location. Connections are created, configured, stored, and monitored from within the Network Connections folder. For information about creating connections, see Make Network and Dial-up Connections. For a list of network connection icons, see Network connection icons.

Each connection in the Network Connections folder contains a set of features that you can use to create a link between your computer and another computer or network. Outgoing connections contact a remote access server by using a configured access method (LAN, modem, ISDN line, cable modem, DSL line, X.25 line, and so on) to establish a connection with a remote network. Conversely, an incoming connection enables a computer running Microsoft® Windows® XP or a stand-alone computer running a member of the Windows Server 2003 family to be contacted by other computers. This means your computer can operate as a remote access server. Whether you are connected locally (LAN), remotely (dial-up, VPN, and so on), or both, you can configure any connection so that it can perform any needed network function. For example, you can print to network printers, access network drives and files, browse other networks, or access the Internet.

Because all services and communication methods are configured within the connection, you do not need to use external management tools to configure connection settings. For example, the settings for a dial-up connection include features to be used before, during, and after connecting. These include the modem you use to dial, the type of password encryption you want to use upon connecting, and the network protocols you use after you connect. Connection status, which includes the duration and speed of a connection, is viewed from the connection itself; you do not need to use an external status tool. For information about configuring a connection, see Configure a connection to a remote network.

Logon and domain security, support for security hosts, data encryption, authentication, and callback provide secure network access for network and dial-up connections. For more information about security features, see Security Features of Network Connections.

# Hardware requirements for network and dial-up connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware requirements for network and dial-up connections**
Depending on your configuration, you may need some or all of the following hardware:

- One or more network adapters with a Network Driver Interface Specification (NDIS) driver for LAN connectivity
- One or more compatible modems and an available COM port
- ISDN adapter (if you are using an ISDN line)
- DSL adapter
- X.25 adapter or PAD (if you are using X.25)
- Analog telephone line, ISDN line, X.25 line, or DSL line
- Smart card reader
- Wireless adapter

To verify the compatibility of all hardware, see the compatibility information at Support resources.

For information about configuring a connection, see Configure a connection to a remote network. For a list of network connection icons, see Network connection icons.

# Network connection types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network connection types**

You can use Network Connections to create several types of connections between computers. The following table shows each type of connection you can create, the communication methods you can use to create each type of connection, and an example of each connection type. For more information, click the appropriate connection type in the table.

| Connection type | Communication method | Example |
|---|---|---|
| Dial-up connections | Modem, ISDN, X.25 | Connect to a private network or the Internet by using remote access |
| Virtual Private Network (VPN) Connections | VPN connections to organization networks over the Internet using PPTP or L2TP | Connect securely to a private network over the Internet |
| Connection Manager Administration Kit | See dial-up or VPN connections | Connect to a private network using a collection of dial-up or VPN settings provided by your network administrator in a self-installing profile |
| Using local area connections | Ethernet, token ring, cable modem, DSL, FDDI, IP over ATM, IrDA, wireless, WAN technologies (T1, Frame Relay), PPPoE | Connect directly to a local area network, to your cable modem, or to your DSL modem through an Ethernet adapter or a similar device |
| Using direct connections | Serial cabling, infrared link, DirectParallel cable | Connect a handheld computer running Microsoft® Windows® CE to a desktop computer so you can synchronize information |
| Incoming connections | See dial-up, VPN, or direct connections | Accept dial-up, VPN, or direct connections from other computers |

For a list of icons that appear in the Network Connections folder, see Network connection icons.

# Dial-up connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up connections**

A dial-up connection connects you to a network or the Internet by using a device that utilizes the telephone network. This device can be a modem that uses a standard phone line, an ISDN adapter with a high-speed ISDN line, DSL, broadband, a cable modem, or an X.25 network.

If you are a typical user, you may have one or two dial-up connections, perhaps to the Internet and to your corporate network.

You can create multiple dial-up connections by copying them in the Network Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple modems, dialing profiles, and different destinations.

For more information about creating dial-up connections, see Set up Dial-up Connections.

# Network connection icons

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network connection icons**

The following table displays the different icons that appear in the Network Connections folder.

| Icon | Description | Location |
|---|---|---|
| | The local area connection is active | Network Connections folder |
| | The local area connection is disabled | Network Connections folder |
| | The default remote access connection | Network Connections folder |
| | An incoming connection | Network Connections folder |
| | The media is disconnected | Network Connections folder |
| | The media is disconnected | Taskbar |
| | The driver is disabled | Network Connections folder |
| | The adapter in this connection has an invalid IP address | Network Connections folder |
| | The adapter in this connection has an invalid IP address | Taskbar |
| | A connection on which Internet Connection Sharing (ICS) has been enabled. | Network Connections folder |
| | A connection on which Internet Connection Firewall (ICF) has been enabled. | Network Connections folder |
| | A connection on which Network Bridge has been enabled. | Network Connections folder |

**Notes**

- This topic applies only to product features available in the original release of Windows Server 2003.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Protocols and Access Methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Protocols and access methods**
This section covers:

- Network communications
- TCP/IP
- IPX
- AppleTalk
- ISDN access
- Serial (RS-232) cable access
- DirectParallel cable access
- SLIP
- Point-to-Point Protocol (PPP)

**Note**

- The IPX/SPX protocol is not available on the Itanium-based versions of the Windows operating systems.

# Network communications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network communications**
Network protocols, access methods, and server protocols help create a communications link between your computer and a network.

Whether information is transported from your computer to a server by using a direct serial cable, or across a secure VPN from your Internet service provider (ISP) to the corporate network, different combinations of access methods and protocols are used to send information. These access methods and protocols are listed as follows:

**Network protocols**

- TCP/IP
- IPX
- AppleTalk

**Access methods**

- Using phone lines and modems
- ISDN access
- Serial (RS-232) cable access
- DirectParallel cable access

**Server protocols**

- Point-to-Point Protocol (PPP)
- SLIP
- Layer Two Tunneling Protocol (L2TP)
- Point-to-Point Tunneling Protocol (PPTP)

The IPX/SPX protocol is not available on the Itanium-based versions of the Windows operating systems.

# TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**TCP/IP**
Transmission Control Protocol/Internet Protocol (TCP/IP) is the most popular protocol, and the basis for the Internet. Its routing capabilities provide maximum flexibility in an enterprise-wide network.

On a TCP/IP network, you must provide IP addresses to clients. Clients may also require a naming service or a method for name resolution. This section explains IP addressing and name resolution for Network Connections on TCP/IP networks. It also describes the FTP and Telnet tools that are provided by the Microsoft® Windows Server 2003 family.

**Assigning IP addresses to remote access connections**
Each remote computer that connects to a remote access server on a Windows TCP/IP network is provided an IP address by the server. This IP address is either automatically provided through the server by DHCP on the target network or selected from the static range that is assigned to the remote access server. For more information see Incoming connections and IP addresses.

If you use a specific IP address, the remote access server must be configured to permit users to request a specific address.

**Name resolution for network and dial-up connections**
In addition to requiring an IP address, network and dial-up connections for a NetBT gateway on a TCP/IP network may require a mechanism to resolve computer names to IP addresses. You can use any one of the following mechanisms for name resolution: DNS, WINS, broadcast name resolution, and the Hosts and Lmhosts files.

Incoming connections clients are assigned the same WINS and DNS servers that are assigned to the incoming connections computer. The TCP/IP settings of a network or dial-up connection can override these automatic assignments.

Incoming connections that do not have a DNS or WINS server rely completely on broadcast name resolution, which occurs by using broadcasts to locate network resources that match a computer name.

**Internet tools**
The Windows Server 2003 family provides File Transfer Protocol (FTP) and Telnet. FTP is a character-based utility that permits you to connect to FTP servers and transfer files. Telnet is a command-line application that lets you log on to remote computers and issue commands as if you were at the keyboard of the remote computer. Multiple variations of FTP, Telnet, and other services based on earlier Internet standards are also available on the Internet or commercially.

For information about configuring a connection, see Configure a connection to a remote network. For information about configuring TCP/IP, see Configure TCP/IP settings.

### Note

- FTP and Telnet tools run on top of TCP/IP networks.

# IPX

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IPX**
Internetwork Packet Exchange (IPX) is the protocol that is used on many Novell NetWare networks. This section explains how to integrate Network Connections into an IPX network.

**Support for NetWare and IPX**
Network Connections supports NetWare and IPX networks by including:

- NWLink IPX/SPX/NetBIOS Compatible Transport Protocol, also known as NWLink, which is the Microsoft® Windows Server 2003 family implementation of the IPX protocol.
- Client Service for NetWare client component, which provides access to NetWare shared file and printer resources.
- Point-to-Point Protocol (PPP) IPX Configuration Protocol (IPXCP), which supports IPX over PPP for dial-up, virtual private network (VPN), and incoming connections. PPP connections use IPXCP, as defined in RFC 1552, "The PPP Internetwork Packet Exchange Control Protocol (IPXCP)," to configure PPP connections for IPX.

For information about configuring a connection, see Configure a connection to a remote network. For information about configuring NWLink IPX/SPX/NetBIOS Compatible Transport Protocol, see Configure NWLink.

**Notes**

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.
- IPX routing is not supported in the Windows Server 2003 family.
- IPX cannot be configured for incoming or dial-up connections in the Windows Server 2003 family.

# AppleTalk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AppleTalk**

Networking on an Apple Macintosh computer is implemented through the AppleTalk protocol. Applications and processes can communicate across a single AppleTalk network or an AppleTalk internet, which is a number of interconnected AppleTalk networks. By using AppleTalk, applications and processes can transfer and exchange data and share resources, such as printers and file servers. AppleTalk remote access is supported by the AppleTalk Control Protocol (ATCP).

With ATCP, Macintosh clients can run the AppleTalk network protocol over the Point-to-Point Protocol (PPP). With ATCP support, a remote user can access a Web server over TCP/IP, print a document to an AppleTalk printer, and connect to a Macintosh file server (over TCP/IP or AppleTalk), all with the same dial-up connection over PPP.

For more information, see AppleTalk over PPP.

For information about configuring a connection, see Configure a connection to a remote network.

# ISDN access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ISDN access**

To enhance your network connection speed, you can use an Integrated Services Digital Network (ISDN) line. Whereas standard phone lines typically transmit at up to 56 kilobits per second (Kbps), ISDN lines can transmit at speeds of 64 or 128 Kbps.

An ISDN line must be installed by the phone company at both the server and at the remote site. ISDN also requires that you install an ISDN adapter on the server and your computer. Costs for ISDN equipment and lines may be higher than standard modems and phone lines. However, the speed of communication reduces the duration of the connection, possibly saving toll charges.

An ISDN line comes with two B channels that transmit data at 64 Kbps, and one D channel for signaling that transmits data at 16 Kbps. You can configure each B channel to operate as a port. With some ISDN drivers, you can aggregate the channels. This means you can statically assign a higher bandwidth by configuring both B channels to act as a single port. With this configuration, line speed increases to 128 Kbps. For information about configuring a connection to use ISDN, see Make a dial-up connection by using an ISDN line or Configure ISDN settings.

The Multilink feature performs channel aggregation for ISDN. Multilink combines multiple physical links into a logical bundle. This aggregate link increases the bandwidth of a connection. In addition, you can allocate multiple links dynamically, which means ISDN lines are used only as they are required. This eliminates excess bandwidth, representing a significant efficiency advantage to users. You can configure the conditions under which extra lines are dialed, and underused lines are disconnected, through Network Connections settings. For more information, see Configuring multiple device dialing.

# Serial (RS-232) cable access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Serial (RS-232) cable access**

Suppose you have two networks in the same location that are not physically connected. To use resources on both networks from one computer, you can use RS-232 cables and COM ports to create null modems. You connect an RS-232 cable from a COM port to a COM port on each remote access server. To the servers, the COM ports then appear to be modems. You can use server remote access, a part of the integrated Routing and Remote Access service, to create network access.

For information about creating a direct connection by using serial access, see Make a direct network connection.

# DirectParallel cable access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DirectParallel cable access**

The DirectParallel feature supports computer-to-computer connections by using standard and enhanced (ECP) parallel ports, which use Parallel Technologies' Basic or Fast parallel cables. For information about configuring a direct network connection to use a DirectParallel cable, see Make a direct network connection by using a DirectParallel cable. Direct cable connection cables are available from several sources. You can obtain DirectParallel cables from Parallel Technologies or from their local dealers. For more information, see To contact Parallel Technologies.

# SLIP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SLIP**
Serial Line Internet Protocol (SLIP) is an older remote access standard typically used by UNIX remote access servers. Network Connections supports SLIP for dial-up connections.

When you connect to a SLIP server, a **Windows Terminal** dialog box appears for an interactive logon session with the UNIX SLIP server. The UNIX logon overrides and prevents the remote access logon session from appearing. After a connection is established, remote network access becomes transparent to the user.

The RFCs supported in remote access are:

- RFC 1055, "A Nonstandard for Transmission of IP Datagrams Over Serial Lines: SLIP"
- RFC 1144, "Compressing TCP/IP Headers for Low-Speed Serial Links"

For information about configuring a network or dial-up connection, see Configure a connection to a remote network.

**Notes**

- Microsoft® Windows® XP does not support SLIP for incoming connections.
- You must use the TCP/IP protocol and a serial COM port to connect to a SLIP server.

# Point-to-Point Protocol (PPP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Point-to-Point Protocol (PPP)**
The Point-to-Point Protocol (PPP) is a set of standard protocols that allow remote access software from different vendors to interoperate. A PPP-enabled connection can dial into remote networks through any industry-standard PPP server. PPP also permits a remote access server to receive calls from, and provide network access to, other vendors' remote access software that complies with the PPP standards.

The PPP standards also permit advanced features that are not available with older standards such as SLIP. PPP supports several authentication methods, as well as data compression and encryption. With most PPP implementations, you can automate the entire logon sequence.

PPP also supports multiple LAN protocols. You can use TCP/IP or IPX as your network protocol.

The IPX/SPX protocol is not available on the Itanium-based versions of the Windows operating systems.

PPP is the basis for the Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) protocols, which are used in secure virtual private network (VPN) connections.

PPP is the primary standard for most remote access computing.

This section covers:

- PPP operation and protocols
- Point-to-Point Protocol over Ethernet (PPPoE) Protocols

# PPP operation and protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**PPP operation and protocols**

Connections made using Point-to-Point Protocol (PPP) must adhere to standards established in PPP RFCs. For information on PPP RFCs, see "Introduction to TCP/IP" at the Microsoft Windows Resource Kits Web site. This topic gives an overview of PPP operations and the protocols used in a PPP connection. For information about how to configure PPP settings, see PPP Settings.

**PPP connection process**

After a physical or logical connection is made to a PPP-based remote access server, the following negotiations establish a PPP connection:

1. Negotiate the use of the link

   PPP uses Link Control Protocol (LCP) to negotiate link parameters such as the maximum PPP frame size, the use of Multilink, and the use of a specific PPP authentication protocol.

2. Authenticate the remote access computer

   The remote access client and the remote access server exchange messages according to the negotiated authentication protocol. If EAP is being used, the client and the server negotiate a specific EAP method, known as an EAP type, and then exchange the messages of that EAP type.

3. Use callback

   If callback for the dial-up connection is configured, the physical connection is terminated, and the remote access server calls the remote access client back.

4. Negotiate the use of network protocols

   This process involves using a series of network control protocols (NCPs) to configure the network protocols the remote access client uses. For example, the NCP for TCP/IP is Internet Protocol Control Protocol (IPCP). If the remote access client is configured to use TCP/IP, IPCP is used to assign an IP address to the client and to configure the client with the IP addresses of DNS and WINS servers.

The resulting connection remains active until the line is disconnected for any of the following reasons:

- The user or the administrator explicitly hangs up the line.
- The line is disconnected due to idle time-out.
- An unrecoverable link error occurs.

**Link Control Protocols**

Link Control Protocol (LCP) establishes and configures PPP framing. PPP framing defines how data is encapsulated before it is transmitted on the wide area network. The PPP standard framing format ensures that remote access software from any vendor can communicate and recognize data packets from any remote access software that adheres to the PPP standards.

PPP and the Microsoft® Windows Server 2003 family use variants of High-level Data Link Control (HDLC) framing for serial or ISDN connections.

**Authentication protocols**

Authentication protocols are negotiated immediately after link quality is determined and before the network layer is negotiated. For more information about available authentication protocols, see Authentication Protocols and Methods.

**Network control protocols**

Network control protocols establish and configure different network protocol settings for TCP/IP, IPX, and AppleTalk. The following table describes network control protocols used with a PPP connection.

The IPX/SPX protocol is not available on the Itanium-based versions of the Windows operating systems.

| Network control protocol | Description |
|---|---|
| Internet Protocol Control Protocol (IPCP) | IPCP is used to configure TCP/IP on the remote access client. Configuration parameters include an IP address and the IP addresses of DNS and WINS servers. |
| Internetwork Packet Exchange Control Protocol (IPXCP) | IPXCP is used to configure IPX on the remote access client. IPXCP is widely implemented by PPP vendors. Configuration parameters include the IPX network and node numbers. |
| AppleTalk Control Protocol (ATCP) | ATCP is used to configure AppleTalk on the remote access client. Configuration parameters include an AppleTalk address. |

**Note**

- The **PPP settings** dialog box includes the option **Negotiate multi-link for single-link connections**. This option enables separation of high and low-priority channels over a single link connection. If your remote access server supports this feature, you might notice improved audio quality. However, because this feature is incompatible with many remote access servers, you should leave it disabled unless instructed otherwise. For more information on PPP settings, see PPP Settings.

# Point-to-Point Protocol over Ethernet (PPPoE) Protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Point-to-Point Protocol over Ethernet (PPPoE) Protocols**
Using Point-to-Point Protocol over Ethernet (PPPoE) and a broadband modem, LAN clients can gain individual authenticated access to high-speed data networks. By combining Ethernet and Point-to-Point Protocol (PPP), PPPoE provides an efficient way to create a separate connection to a remote server for each user. Access, billing, and choice of service are managed on a per-user basis, rather than a per-site basis. Clients can create both dial-up and dial-on-demand connections using PPPoE.

**PPPoE RFCs**
Requests for Comments (RFCs) are an evolving series of technical reports, proposals for protocols, and protocol standards used by the Internet community. PPPoE is defined in RFCs published by the Internet Engineering Task Force (IETF) and other working groups. The following is an example of a PPPoE RFC:

**RFC number** - 2516A

**Title** - Method for Transmitting PPP Over Ethernet (PPPoE)

**Obtaining RFCs**
You can obtain RFCs from the RFC Editor Web site. This Web site is currently maintained by members of the Information Sciences Institute (ISI) who publish a classified listing of all RFCs. RFCs are classified as one of the following: approved Internet standards, proposed Internet standards (circulated in draft form for review), Internet best practices, or For Your Information (FYI) documents.

### Notes

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.
- PPPoE is supported on clients running Microsoft® Windows® XP. All members of the Windows Server 2003 family can connect to and route PPPoE traffic. However, they cannot accept an incoming PPPoE connection.

For more information, see Make a broadband connection using Point-to-Point Protocol over Ethernet (PPPoE).

# Home and Small Office Networking Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Home and small office networking overview**
With home and small office networking, even if you do not have the resources, time, or experience of a network administrator, you can build a small network in your home or small office. Home and small office networking is a group of features that limits hardware requirements and automates configurations for routing and private IP addressing. This reduces the work that is typically required to connect and configure a network. It also reduces the associated costs.

Home and small office networking provides many of the benefits of traditional networks and offers benefits that are uniquely suited for a home or small office network. With home and small office networking you can:

- Share one Internet connection with all of the computers on your network.
- Provide local private network services, such as name resolution and addressing, automatically.
- Share printers or other hardware devices, such as cameras and scanners, from any computer on your network.
- Share folders and files with other computers on your network.
- Use one computer to secure your entire network and protect your Internet connection.
- Easily establish network connections across two or more types of network media, such as IEEE 1394, Ethernet, wireless, and Home Phoneline Network adapters.

This section covers:

- Home and small office networking features
- Home and small office networking compared to Routing and Remote Access

# Home and small office networking features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Home and small office networking features**
Home and small office networking uses the following features to simplify the implementation and administration of small networks.

**Internet Connection Firewall (ICF) and Windows Firewall**
Home and small office networking offers computers that are directly connected to the Internet the added security of firewall protection. The firewall checks all communications that cross the connection between your network and the Internet and is selective about which responses from the Internet it allows. ICF and Windows Firewall protect only the computer on which they are enabled. However, if ICF or Windows Firewall are enabled on the Internet Connection Sharing (ICS) host computer, ICS clients that use the shared Internet connection for Internet connectivity are protected because they cannot be seen from outside your network. For this reason, you should always enable ICF or Windows Firewall on the ICS host computer. In addition, if there are clients on your network with direct Internet connections, or if you have a stand-alone computer that is connected to the Internet, then it is recommended that you enable ICF or Windows Firewall on those Internet connections as well.

For more information about ICF, see Internet Connection Firewall. For more information about Windows Firewall, see Help: Windows Firewall.

**Network Bridge**
Network Bridge removes the need for routing and bridging hardware in a home or small office network that consists of multiple LAN segments. With Network Bridge, multiple LAN segments become a single IP subnet, even if the LAN segments are of mixed network media types. Network Bridge automates the configuration and management of the address allocation, routing, and name resolution that is typically required in a network that consists of multiple LAN segments.

For more information see Network Bridge.

**Internet Connection Sharing (ICS)**
ICS allows you to provide Internet access for a home or small office network by using one common connection as the Internet gateway. The ICS host is the only computer that is directly attached to the Internet. Multiple ICS clients simultaneously use the common Internet connection and benefit from Internet services as if the clients were directly connected to the Internet service provider (ISP) (for more information, see Dial-up credentials). Security is enhanced when ICS is enabled because only the ICS host computer is visible to the Internet. The addresses of ICS clients are hidden from the Internet rendering ICS clients invisible to the Internet. In addition, ICS simplifies the configuration of small networks by providing local private network services, such as name resolution and addressing.

For more information about ICS, see Connecting to the Internet in a home or small office network.

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

# Home and small office networking compared to Routing and Remote Access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Home and small office networking compared to Routing and Remote Access**
Home and small office networking and the Routing and Remote Access service are both included in Microsoft® Windows® XP and the Windows Server 2003 family to provide the following networking capabilities:

- Network address translation (NAT)
- Packet filtering
- Dial-up access
- Virtual private network (VPN) access
- Automated IP addressing for network clients
- Name resolution
- Internet connections

Because home and small office networking features and Routing and Remote Access share common drivers, they cannot coexist, and you need to decide which to implement in order to meet the needs of your networking environment.

**Home and small office networking**
You can use home and small office networking features--Internet Connection Firewall (ICF), Windows Firewall, Network Bridge, and Internet Connection Sharing (ICS)--to simplify the configuration of your home or small office network. These features are intended for networks that consist of two to 10 computers when you want to:

- Protect computers that are connected to the Internet using an easily configured firewall. For more information, see Internet Connection Firewall.
- Connect LAN segments simply, by using the **Bridge Connections** menu command. For more information, see Network Bridge.
- Provide simultaneous Internet access for up to 10 computers using a single dial-up or high-speed Internet connection. For more information, see Connecting to the Internet in a home or small office network.
- Support up to 10 simultaneous dial-up or VPN connections to provide remote computers with access to network resources. For more information, see Incoming connections.

**Routing and Remote Access**
You can use Routing and Remote Access in combination with Internet Security and Acceleration (ISA) server to meet the needs of your small business with high security needs, your medium-sized private business, or your enterprise network that spans multiple subnets and supports up to 1,000 computers running Windows Server 2003, Standard Edition or up to 5,000 computers running Windows Server 2003, Enterprise Edition. Internet connectivity for branch offices is routed through the corporate routing and firewall infrastructure. Use Routing and Remote Access when you want to:

- Provide local and branch office computers with high-security Internet access.
- Connect branch offices with corporate intranets, and share resources as if all computers are connected to the same LAN.

- Protect network interfaces with static packet filters or dynamic packet filters.
- Support up to 1,000 simultaneous dial-up or VPN connections to provide remote computers with access to corporate network resources.

**Comparisons**

The following table summarizes the differences between the way that home and small office networking features (and incoming connections) and Routing and Remote Access (with ISA server) implement basic networking services.

| Service | Description | Home and small office networking features (and incoming connections) | Routing and Remote Access (with ISA server) |
|---|---|---|---|
| **NAT** | Hides internally managed IP addresses from external networks by translating private internal addresses to public external addresses. This reduces IP address registration costs by letting you use unregistered IP addresses internally, with translation to a small number of registered IP addresses externally. It also hides the internal network structure, reducing the risk of attacks against internal systems. | ICS | Routing and Remote Access NAT |
| **Dynamic packet filtering** | Provides protection from unsolicited traffic for the private network. Permits only traffic that is sent in response to an internal request. | ICF, Windows Firewall | Basic Firewall |
| **VPN and dial-up access** | Allows clients on a remote computer to connect to a private network and to access network resources as if the computer was physically attached to the network. | Incoming connections | Routing and Remote Access |
| **Address assignment** | Automates the assignment of client IP addresses on the private network in order to configure clients to allow client-access to network resources. | ICS DHCP allocator | Routing and Remote Access NAT, with a DHCP allocator or a DHCP server |
| **DNS name resolution** | Converts the names of computers and other network devices, such as printers, to IP addresses. | ICS DNS Proxy | NAT, with a DNS proxy or a DNS server |
| **Internet connections** | Provides a high-speed or dial-up connection for the computer that is connected to the Internet to use to publish its services to the private network. | Configured through Network Connections | Configured through Routing and Remote Access |

The following table summarizes which components, features, and connection types are best suited to provide networking services in a given networking scenario.

| Service | Scenario: Small office, non-domain network | Scenario: Small office, domain network | Scenario: Medium office network with NAT traversal | Scenario: Medium office network without NAT traversal | Scenario: Enterprise network | Scenario: Branch office network |
|---|---|---|---|---|---|---|
| **NAT** | ICS | ICS | Routing and Remote Access | Routing and Remote Access | Routing and Remote Access | Routing and Remote Access with ISA (on the corporate network) |

| | | | | | | |
|---|---|---|---|---|---|---|
| **Packet filtering** | ICF, Windows Firewall | ICF, Windows Firewall | Basic Firewall or ISA | Basic Firewall or ISA | ISA | ISA (on the corporate network) |
| **VPN/remote access** | Incoming connections | Incoming connections | Routing and Remote Access | Routing and Remote Access | Routing and Remote Access | Routing and Remote Access with ISA (on the corporate network) |
| **Address assignment** | ICS DHCP allocator | DHCP server | DHCP server | DHCP server | DHCP server | DHCP server |
| **Internet connection** | Dial-up, ISDN, broadband, DSL, or LAN | Dial-up, ISDN, broadband, DSL, or LAN (with ISA) | Any combination of T1 or T3, dial-up, ISDN, broadband, DSL, and LAN | Any combination of dial on demand (DoD) routing, T1 or T3, dial-up, ISDN, broadband, DSL, or LAN | Any combination of dial on demand (DoD) routing, T1 or T3, dial-up, ISDN, broadband, DSL, or LAN | Any combination of dial on demand (DoD) routing, T1 or T3, dial-up, ISDN, broadband, DSL, or LAN |
| **NAT traversal** | Not available | Not available | ISA | Does not apply | ISA | ISA (on the corporate network) |
| **DNS** | DNS proxy | DNS server | DNS server | DNS server | DNS server | DNS server |

**Notes**

- In the case of ICS with a DNS or DHCP server, ICS discovers the DHCP service or DNS service and, if present, disables the DNS proxy, DHCP allocator, or both.
- Because ICS, ICF, and Windows Firewall are not compatible with Routing and Remote Access, you cannot enable ICS or ICF if Routing and Remote Access is configured. Likewise, if ICS or ICF are enabled, they must first be disabled in order to configure Routing and Remote Access.
- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Understanding Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Network Connections**
This section covers:

- Security for Network Connections
- Tunneling Protocols
- Home and Small Office Networking

# Security for Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security for network connections**

This section covers:

- Security Features of Network Connections
- Intermediary Devices

# Security Features of Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security features of Network Connections**

You can configure your dial-up, virtual private network (VPN), broadband, and direct connections to enforce various levels of password authentication and data encryption. Authentication methods range from unencrypted to custom, using Extensible Authentication Protocol (EAP). EAP provides an infrastructure to support a wide range of authentication methods, including smart cards, certificates, one-time passwords, and public keys. You can also specify the type of data encryption, depending on the level of password authentication (MS-CHAP, MS-CHAP v2, or EAP-TLS) that you select. Finally, you can configure callback options to increase dial-up security.

User accounts and domains provide security with encrypted authentication and encrypted data. For more information about security, see Security.

This section covers:

- Callback
- How security works at connection
- Group Policy for Network Connections
- Group Policy settings that prohibit home and small office networking on your domain
- Dial-up credentials
- Data Encryption and Encryption Methods
- Network Authentication

# Callback

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Callback**
The callback feature instructs the remote access server to disconnect after authentication and then to call you back. Callback provides cost advantages to you and security advantages to your network. By immediately hanging up and then dialing you back, callback reduces phone charges. Required callback enhances network security by ensuring that only users from specific locations can access the server. By dropping the call, and then calling back a moment later to the preassigned callback number, most impersonators can be thwarted.

If callback has been required by your administrator, the following steps occur when your call reaches the remote access server:

1. The server determines whether your user name and password are correct.

2. If they are correct, the server disconnects and then returns the call.

If your administrator has configured callback for your connection through the **Remote Access Preferences** menu option, the following steps occur when your call reaches the remote access server:

1. The server determines whether your user name and password are correct.

2. If they are correct, the **Callback** dialog box appears on your computer.

3. You then type the current callback number in the dialog box and wait for the server to disconnect and return the call.

   Optionally, you can press ESC to cancel the callback process and remain connected.

For more information about configuring your callback options, see Configure client callback options.

**Notes**

- Your callback options are determined by your system administrator. The administrator may deny the use of callback, allow user-defined callback, or require callback to a particular number. You can contact your system administrator to determine your callback options.

- If your computer is configured to accept incoming connections, you can enforce callback options on that computer. For more information, see Configure callback for incoming connections.

# How security works at connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How security works at connection**

The following steps describe what happens during a call to a remote access server:

1. The remote access client computer calls a remote access server.

2. Depending on the authentication methods that have been negotiated, one of the following happens:

   - If you are using PAP or SPAP:

     The client sends its user name and password to the server.

     The server checks the account credentials against the user database.

   - If you are using CHAP or MS-CHAP:

     The server sends a challenge to the client.

     The client sends an encrypted response to the server.

     The server checks the response against the user database.

   - If you are using MS-CHAP v2:

     The server sends a challenge to the client.

     The client sends an encrypted response and a challenge to the server.

     The server checks the encrypted response of the client against the user database and sends the client its own encrypted response.

     The client verifies the encrypted response of the server.

   - If you are using certificate-based authentication:

     The server sends its computer certificate to the client.

     If the client is configured for certificate-based authentication, the client validates the server certificate.

     The client sends its user certificate to the server.

     The server verifies that the user certificate is valid and has not been revoked.

3. If the connection attempt is both authenticated (the user credentials are valid, the user account is enabled and not locked out, and the connection is occurring within the allowed logon hours) and authorized (for incoming connections, the user account has remote access permission), the server accepts the remote access connection.

   If callback is enabled, the server calls your computer back and repeats steps 2 through 3.

**Notes**

- For the Routing and Remote Access service, authorization of the connection is determined by the dial-in properties of the

user account and remote access policies. Remote access policies are a set of conditions and connection settings that give network administrators more flexibility in granting remote access permissions and usage. If the settings of the connection do not match at least one of the remote access policies that apply to your connection, the connection attempt is rejected, regardless of the dial-in properties of the user account.

- Network Connections authentication methods control access to your network, not to resources on the network. Once you are connected to a network, resource access control is managed through access control lists (ACLs) in various ways: Local Users and Groups, Active Directory, Group Policy, file and printer sharing, and so on.

- Authentication methods are designed for remote computing environments where the computer that is dialing in does not have access to the network until credentials, such as user name and password, are validated.

- For information about configuring security options for your connections, such as identity authentication protocols and data encryption settings, see Secure Network or Dial-up Connections.

# Group Policy for Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy for Network Connections**
By using the Group Policy snap-in, an administrator can configure the desktop settings of many computers at once. Some Group Policy settings apply specifically to Network Connections, such as user access to connections and the ability to change the properties of connections. These settings can be found in the Administrative Templates of the Group Policy snap-in.

### Where?

- *Group_Policy_object_name*/Computer Configuration (or User Configuration)/Administrative Templates/Network/Network Connections

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- If you do not have write access to a Group Policy object, you cannot see the settings contained in that object.

For an overview of Group Policy, Group Policy objects, and administrative templates, see Group Policy overview. For more information on how to use Group Policy to define policy settings that are applied to computers or users, see Group Policy (pre-GPMC).

# Group Policy settings that prohibit home and small office networking on your domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy settings that prohibit home and small office networking on your domain**

Group Policy settings define the various components of the user's environment that a system administrator can manage. The following Group Policy settings are available to prohibit home and small office networking on your domain:

- **prohibit use of Internet Connection Sharing on your DNS domain network** determines whether computer users with administrator accounts can enable and configure Internet Connection Sharing (ICS) on network connections on your domain.

- **prohibit use of Internet Connection Firewall on your DNS domain network** determines whether computer users with administrator accounts can enable and configure Internet Connection Firewall (ICF) on network connections on your domain.

- **prohibit installation and configuration of Network Bridge on your DNS domain network** determines whether computer users with administrator accounts can enable Network Bridge on your domain.

**Group Policy for ICS**

ICS allows a user with an administrator account to provide Internet access for a home or small office network, using one common connection as the Internet gateway, and to provide local private network services, such as name resolution and addressing. The ICS host is the only computer that is directly connected to the Internet. Multiple ICS clients simultaneously use the common (shared) Internet connection. ICS is available only on computers that have two or more network connections. Users with administrator accounts enable ICS by accessing the **Advanced** tab on network connections. For more information about ICS, see Connecting to the Internet in a home or small office network.

If you do not want users on your domain to be able to enable or configure ICS, apply the **prohibit use of Internet Connection Sharing on your DNS domain network** Group Policy setting. When you apply the **prohibit use of Internet Connection Sharing on your DNS domain network** Group Policy setting, users with administrator accounts cannot enable or configure ICS because the **Advanced** tab on network connections becomes unavailable, and the related wizard options on computers running Windows XP Professional are blocked.

Computer users with user rights are prohibited from configuring ICS, regardless of this Group Policy setting.

**Group Policy for ICF**

ICF allows a user with an administrator account to activate a firewall to protect the public connection of a small network or a single computer that is connected to the Internet. ICF is enabled by default only on dial-up connections. Users with administrator accounts can enable ICF on other types of connections, such as LAN and wireless connections, by accessing the **Advanced** tab on network connections. For more information about ICF, see Internet Connection Firewall.

If you do not want users to be able to enable or configure ICF, apply the **prohibit use of Internet Connection Firewall on your DNS domain network** Group Policy setting. When you apply the **prohibit use of Internet Connection Firewall on your DNS domain network** Group Policy setting, users with administrator accounts cannot enable or configure ICF because the **Advanced** tab becomes unavailable, and the related wizard options on computers running Windows XP Professional are blocked.

Computer users with user rights are prohibited from configuring ICF, regardless of this Group Policy setting.

**Group Policy for Network Bridge**

Network Bridge automates the configuration that is required to route traffic between multi-segment networks comprised of a single type of media or mixed media. With Network Bridge, no configuration is required, and you do not need purchase additional hardware such as routers or bridges. The Network Bridge menu command **Bridge Connections** is available only when two or more network adapters are present. By default, Network Bridge is disabled, but users with administrator accounts can use **Bridge Connections** to enable Network Bridge. For more information about Network Bridge, see Network Bridge.

If you do not want users on your domain to be able to enable or configure Network Bridge, apply the **prohibit installation and configuration of Network Bridge on your DNS domain network** Group Policy setting. When you apply the **prohibit**

**installation and configuration of Network Bridge on your DNS domain network** Group Policy setting, users with administrator accounts cannot enable or configure Network Bridge because the **Bridge Connections** command is removed from the menu on network connection icons, and the related wizard options on computers running Windows XP Professional are blocked.

Computer users with user rights are prohibited from configuring Network Bridge, regardless of this Group Policy setting.

**Important**

The Group Policy settings for ICS and ICF are location-aware. A location-aware Group Policy setting applies only when a computer is connected to the domain that it was connected to when the setting was last refreshed. If a computer, such as a laptop computer, is connected to your domain, but the setting on the computer was last refreshed on another domain, then the Group Policy setting on your network do not apply to that computer.

For example, if ICS is enabled on a computer that is on your network, and then you apply the Group Policy setting for ICS, ICS continues to run as normal until you restart the computer or log off of the network and then log on again. The same is true if the Group Policy setting for ICS is set on your network, and then a computer on which ICS is enabled joins your network. In this case also, ICS continues to run as normal until you restart the new computer or log off of the network and then log on again. This example applies to ICF as well.

The Group Policy setting for Network Bridge works somewhat differently. If Network Bridge is enabled on a computer that is on your network, and then you apply the Group Policy setting for Network Bridge, or if the Group Policy setting for Network Bridge is set on your network, and then a computer on which Network Bridge is enabled joins your network, Network Bridge continues to exist in the Network Connections folder and to use system resources, but the functionality of Network Bridge changes. In these cases, Network Bridge continues to receive and send data over the network connections that are included in the bridge but, for security reasons, no longer forwards traffic from one network connection to another. This change is immediate; it does not require that you restart the computer or log off of the network and then log on again.

For an overview of Group Policy, see Group Policy (pre-GPMC).

**Note**

- This topic applies only to product features available in the original release of Windows Server 2003.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.
- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Dial-up credentials

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dial-up credentials**

Credentials are required to successfully connect to the Internet using dial-up connections. Your credentials are the combination of your ISP account name and password. Your computer user profile identifies your account as either an administrator account or a user account. If you log on using user credentials, you will have two options for how your credentials are saved when you create a dial-up connection. If you log on with administrative credentials, you will have three options for how your credentials are saved when you create a dial-up connection. If you have account profiles with both administrator and user credentials, only the credentials for the account that you have logged on with will apply.

If you log on with user account credentials, you can create only dial-up connections for **My use only**. These connections will be successful only when you are logged on to the computer, and they will not be accessible to any other users. When creating user account dial-up connections, you can either save your credentials as **Me only** (for your use only) or not save your credentials. If you save your credentials as **Me only**, and you are logged on to the computer using the account you used when you created the connection, your credentials will be automatically supplied when using the connection. If you do not save your credentials, a dialog box will be provided for you to enter your credentials each time you are logged on to the computer using the account you used when creating the connection and you attempt to connect to the Internet.

If you log on with administrative credentials, you can create connections for either **Anyone's use** or **My use only**. If you create a connection for **My use only**, the connection options are the same as those created by anyone with a user account. Connections created for **Anyone's use** can be seen by any user who logs on to that computer, but connection success to the Internet is affected by how credentials were saved. When creating a connection for **Anyone's use**, you can save your credentials so that they will be supplied automatically for **Anyone who logs on to this computer** or **Me only**. You can also choose not to save your credentials.

If you save your credentials for **Anyone who logs on to this computer**, your saved credentials will automatically be supplied for any user who has logged on to the computer and attempts to connect to the Internet using that connection or on-demand dialing. If you save your credentials but do not make them available for **Anyone who logs on to this computer**, then your saved credentials will be automatically supplied only if you are logged on to the computer. Other users will see a dialog box that prompts them for your user name and password; they will be unable to connect to the Internet using this connection unless they can provide your credentials. If you do not save your credentials, a dialog box will appear and prompt you or any other user to enter your credentials before the connection to the Internet can be completed.

The following summarizes the experience of a user logging on with user credentials and experience of a user logging on with administrative credentials:

**Logging on with user credentials**

| Connection type created | Credential save options | Local computer access |
|---|---|---|
| • **My use only** | • **Me only**<br>• Not saved | This connection will work only when you are logged on with the account that you used when you created the connection. If you did not save your credentials, you will be prompted to supply them every time you use the connection. You cannot configure Internet Connection Sharing (ICS), Internet Connection Firewall (ICF), or Windows Firewall on this connection. If you want to enable ICS or configure ICF or Windows Firewall, you should log on with administrative credentials before you create the dial-up connection because only administrators can access the **Advanced** tab. If you want to enable ICS, ICF, or Windows Firewall on a dial-up connection that was created with user credentials, the user who created the connection must be promoted to an administrator. |

| ICS client Internet access through My use only ICS connections |
|---|

| |
|---|
| Users of ICS client computers can connect to the Internet through connections on the ICS host computer that are saved for **My use only** if:<br>• The user who created the connection saved credentials and is logged on to the ICS host computer.<br>• The user who created the connection did not save credentials but that user or another user can supply those credentials when prompted by the ICS host computer to do so. |

**Logging on with administrative credentials**

| Connection type created | Credential save options | Local computer access |
|---|---|---|
| • **My use only** | • **Me only**<br>• Not saved | This connection will work only when you are logged on with the account that you used when you created the connection. If you did not save your credentials, you will be prompted to supply them every time you use the connection. You cannot configure Internet Connection Sharing (ICS), Internet Connection Firewall (ICF), or Windows Firewall on this connection. If you want to enable ICS or configure ICF or Windows Firewall, you should log on with administrative credentials before you create the dial-up connection because only administrators can access the **Advanced** tab. If you want to enable ICS, ICF, or Windows Firewall on a dial-up connection that was created with user credentials, the user who created the connection must be promoted to an administrator. |
| • **Anyone's use** | • **Anyone who uses this computer**<br>• **Me only**<br>• Not saved | All users can see the Internet connection. If you saved your credentials for all users, any user who is logged on to the computer can connect to the Internet. If you saved your credentials as **Me only**, the connection is automatically completed only if you log on to the computer as the administrator who saved the credentials when you created the connection. If you did not save your credentials, when you or another user attempts to connect to the Internet, your credentials must be supplied when prompted in order to complete the connection. You can access the **Advanced** tab to enable ICS or configure ICF or Windows Firewall because you logged on as an administrator when you created the dial-up connection. |

| **ICS client Internet access through Anyone's use connections** |
|---|
| Users of ICS client computers can connect to the Internet through connections on the ICS host computer that are saved for **Anyone's use** only if:<br>• The administrator who created the connection saved credentials for use by anyone who uses the computer. (This attempt will succeed whether a user is logged on to the ICS host computer or the ICS client computer.)<br>• The administrator who created the connection saved credentials as **User only** and is now logged on to the ICS host computer with the same credentials.<br>• The administrator who created the connection did not save credentials, but a user can supply those credentials when prompted by the ICS host computer. |

You must enable on-demand dialing in order for programs to automatically initiate dial-up connections to your Internet service provider. For more information, see Enable ICS on-demand dialing.

You must enable ICS in order for computers on your network to share a common Internet connection. For more information, see Connecting to the Internet in a home or small office network.

**Notes**

- As a security measure, you should log on using user credentials rather than administrative credentials.
- Internet Connection Firewall is included only in the original releases of Windows Server 2003, Standard Edition, and Windows Server 2003, Enterprise Edition.

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

- Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

# Data Encryption and Encryption Methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data encryption and encryption methods**

This section covers:

- Data encryption
- Microsoft Point-to-Point Encryption (MPPE)
- Internet Protocol Security (IPSec) Encryption

# Data encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data encryption**

You can think of encryption as locking something valuable into a strong box with a key. Sensitive data is encrypted by using an encryption algorithm and a key, which renders it unreadable without the knowledge of the key. Data encryption keys are determined at connect time between the connecting computers. The use of data encryption can be initiated by your computer or by the server you are connecting to.

Network Connections supports two types of encryption:

- Microsoft MPPE, which uses RSA RC4 encryption.
- An implementation of Internet Protocol security (IPSec) that uses Data Encryption Standard (DES) encryption.

Both MPPE and IPSec support multiple levels of encryption, as shown in the following table.

| Encryption type | Level of encryption supported |
|---|---|
| MPPE Standard | 40-bit, 56-bit |
| MPPE Strong | 128-bit |
| IPSec DES | 56-bit |
| IPSec Triple DES | 3DES |

For more information, see the following:

- Microsoft Point-to-Point Encryption (MPPE)
- Internet Protocol Security (IPSec) Encryption

**Notes**

- Server controls are flexible and can be set to deny the use of encryption, require a specific encryption strength, or allow your computer to select an encryption strength. By default, most servers allow encryption and allow clients to select their encryption methods. This works for most computers. A system administrator sets encryption requirements. To determine your options, contact your system administrator.
- To enable MPPE-based data encryption for dial-up or virtual private network (VPN) connections, you must select the MS-CHAP, MS-CHAP v2, or EAP-TLS authentication methods. These authentication methods generate the keys used in the encryption process.
- Virtual Private Networks (VPNs) use encryption depending on the type of server they are connecting to. If the VPN connection is configured to connect to a PPTP server, then MPPE encryption is used. If the VPN is configured to connect to an L2TP server, then IPSec encryption methods are used. If the VPN is configured for an Automatic server type, which is the default selection, then L2TP, and its associated IPSec encryption, will be attempted first, then PPTP, and its associated MPPE encryption.

# Microsoft Point-to-Point Encryption (MPPE)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Point-to-Point Encryption (MPPE)**
Microsoft Point-to-Point Encryption (MPPE) encrypts data in Point-to-Point Protocol (PPP)-based dial-up connections or Point-to-Point Tunneling Protocol (PPTP) virtual private network (VPN) connections. 128-bit key (strong), 56-bit key, and 40-bit key (standard) MPPE encryption schemes are supported. MPPE provides data security for the PPTP connection that is between the VPN client and the VPN server.

For more information about configuring connections, see Configure a connection to a remote network.

**Note**

- MPPE requires encryption keys as generated by Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2), or Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication process.

# Internet Protocol Security (IPSec) Encryption

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol security (IPSec)**
The long-term direction for secure networking, IPSec is a suite of cryptography-based protection services and security protocols. Because it requires no changes to applications or protocols, you can easily deploy IPSec for existing networks.

IPSec provides computer-level authentication, as well as data encryption, for VPN connections that use the L2TP protocol. IPSec is negotiated between your computer and an L2TP-based VPN server before an L2TP connection is established. This negotiation secures both passwords and data.

L2TP uses standard PPP-based authentication protocols, such as EAP, MS-CHAP, MS-CHAP v2, CHAP, SPAP, and PAP with IPSec.

Encryption is determined by the IPSec Security Association, or SA. A security association is a combination of a destination address, a security protocol, and a unique identification value, called a Security Parameters Index (SPI). The available encryptions include:

- Data Encryption Standard (DES), which uses a 56-bit key.
- Triple DES (3DES), which uses three 56-bit keys and is designed for high-security environments.

For more information about IPSec, see Internet Protocol Security (IPSec).

For more information about configuring connections, see Configure a connection to a remote network.

# Network Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication**
This section covers:

- Authentication process
- Authentication Protocols and Methods

# Intermediary Devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Intermediary devices**
This section covers:

- Using an intermediary device
- Security hosts
- Writing scripts for intermediary devices

# Tunneling Protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tunneling protocols**
This section covers:

- Point-to-Point Tunneling Protocol (PPTP)
- Layer Two Tunneling Protocol (L2TP)

# Home and Small Office Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Home and small office networking**
This section covers:

- Internet Connection Firewall
- Internet Connection Firewall security log
- Network Bridge
- Service definitions for Internet Connection Firewall and Internet Connection Sharing

**Note**

- Internet Connection Sharing and Network Bridge are not included in Windows Server 2003, Web Edition; Windows Server 2003, Datacenter Edition; and the Itanium-based versions of the original release of the Windows Server 2003 operating systems.

# Using Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Network Connections**

Network Connections is automatically installed. This section covers:

- Monitoring network and dial-up connections
- Making Network Connections
- Using Home and Small Office Networking Features

# Administering Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Network Connections**

This section covers:

- Managing network connections from the command line
- Logging remote access connections
- Configuring Network and Dial-up Connections

# Network Connections Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section covers:

- Troubleshooting network and dial-up connections
- Error messages in Network Connections (Windows Server 2003)

# Installing and configuring Modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modems**
Modem support is provided by the TAPI Unimodem 5 Service Provider. It enables you to install and use almost any modem available today. Modems are used by programs such as HyperTerminal, Network Connections, and Fax, as well as by many products, including software required to gain access to Internet service providers.

- Before you begin using modems with this computer, see Checklist: Installing and configuring modems.
- To find features that have been moved in the Microsoft® Windows Server 2003 family, see New ways to do familiar modem tasks.
- For tips about using modems, see Modem best practices.
- For help with specific tasks, see Modems How To ....
- For general background information, see Modems Concepts.
- For problem-solving instructions, see Modems Troubleshooting.

# Checklist: Installing and configuring modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring modems**

| Step | Reference |
|---|---|
| Review key concepts. | Modem overview; Understanding Modems; Dialing rules overview |
| Install modem hardware. | Installing a modem |
| Install the modem. | Install a modem; Install an unsupported modem |
| Configure the installed modem. | Configure an installed modem; Change How Modems Initialize and Connect; Change Modem Data and Hardware Settings |
| Set dialing properties for your location. | Dialing rules overview; Change dialing properties for a location |
| Connect using your modem. | Connecting using your modem |

# New ways to do familiar modem tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring modems. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Install a modem | **Modems** in Control Panel. | Open Phone and Modem Options. For more information, see Install a modem. |
| Configure a modem | **Modems** in Control Panel. | Open Phone and Modem Options. For more information, see Configure an installed modem. |
| Test or troubleshoot a modem | Terminal program, such as **Hyperterminal**. | Open Phone and Modem Options. For more information, see Test a modem. |

# Modem best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

The following list provides best practices for installing and using modems. Review this list before you use modems.

- **Turn on any external modems before you start Windows.**

  If an external Plug and Play modem is not turned on before you start Windows, Windows may not detect it correctly.

  For more information, see Troubleshooting modems.

- **If Windows does not support your modem make and model, check the manufacturer's Web site for the latest installation or** .inf **file.**

  The installation files, particularly for new modems, are often added or updated by manufacturers. If you cannot find your modem listed in Windows, look in the manufacturer's documentation to find if the modem has the same characteristics as another supported modem.

  Some legacy modems are not supported. For a list of supported modems, see Support resources.

  For more information, see Installing a modem.

- **Use the Phone and Modem Options diagnostics utility to test and troubleshoot a modem.**

  To see if the modem is responding to commands, in Control Panel, double-click **Phone and Modem Options**. In **Phone and Modem Options** click the **Modems** tab, right-click the modem that you are interested in, click **Properties**, and then click the **Diagnostics** tab. You can also configure logging and view logs at this location.

  For more information, see Modem diagnostics, Test a modem, Log and view modem commands.

- **Before you connect to another computer or online service provider, check the hardware settings for your modem.**

  The data connection parameters for two modems need to be identical for them to connect successfully. Refer to the documentation provided by your service provider for the correct settings. Typical settings are:

  - 8 data bits
  - None for parity
  - One stop bit

  Bulletin boards and most service providers use these settings. If these do not work, try seven data bits, even parity, and one stop bit. A few online service providers use these settings. Other settings are extremely rare.

  For more information, see Hardware settings.

- **When you install the same make and model of modem on multiple ports, first install it on one port, and then duplicate the installation for the remaining ports.**

  For more information, see Duplicate a modem installation on multiple ports.

- **To change the same setting for modems installed on multiple ports, first change the setting on one modem and then copy the change to the remaining modems.**

For more information, see Copy modem properties to modems on multiple ports.

- **Rarely, you may have to configure the COM port before installing a modem.**

  The COM port is automatically configured for most modems. The following exceptions may occur:

  - If you install an internal modem card that is not compatible with Plug and Play specifications, you may have to configure its built-in COM port by using **Add Hardware** in Control Panel before you can install the modem.

  - If you are using PCMCIA drivers, your computer will detect and configure the PCMCIA modem cards automatically when you first insert them. If the card is not detected, you might need to run **Add Hardware** in Control Panel to configure the card's built-in COM port. You then install the PCMCIA modem by using the **Phone and Modem Options** in Control Panel.

  For more information, see Install a modem, Troubleshooting network and dial-up connections.

- **If you create a direct connection to a computer, Windows CE device, or other device, make sure that the maximum port speed in the connection properties is set as high as supported by the device.**

  The initial maximum port speed for a direct connection created in the **Phone and Modem Options** in Control Panel is 19.2 Kbps. If the device on a connection supports higher data transmission speeds, you can increase the maximum port speed to that supported by the device. To find the correct speed, refer to the manufacturer's documentation provided with the device.

  For more information, see Change the maximum modem port speed, Optimizing data transfer speed.

# Modems How To ...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to ...**

- Install or Remove Modems
- Install and Configure Modems on Multiple Ports
- Test Modems
- Change How Modems Initialize and Connect
- Change Modem Data and Hardware Settings

# Install or Remove Modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install or remove modems**

- Install a modem
- Install an unsupported modem
- Configure an installed modem
- Change the COM port for a modem
- Remove a modem

# Install a modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install a modem**

1. Open Phone and Modem Options in Control Panel.
2. (Optional) If you are prompted for location information, enter dialing information for your location and click **OK**.
3. On the **Modems** tab, click **Add**.
4. Follow the instructions in the Install New Modem wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- If the Install New Modem wizard doesn't detect your modem or you cannot find it listed, see Related Topics for instructions on installing an unsupported modem.
- If you are installing an external modem, if a cable is not provided, refer to the manufacturer's instructions for cabling requirements. Most common cables will work, but some cables do not have all the pins connected. Do not use the 9-to-25 pin converters that come with most mouse hardware, because some of them do not carry modem signals.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing a modem
Remove a modem
Install an unsupported modem
Configure an installed modem
Change the COM port for a modem

# Install an unsupported modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install an unsupported modem**

1. Open Phone and Modem Options in Control Panel.

   If you are prompted for location information, enter dialing information for your location and click **OK**.

2. On the **Modems** tab, click **Add**.

3. In the Install New Modem wizard, select the **Don't detect my modem; I will select it from a list** check box, and then click **Next**.

4. Do one of the following:

| If | Do this |
|---|---|
| A modem is listed that is a good match for your modem | On the list of modems, click the manufacturer and model that is the best match for your modem, and then click **Next**. |
| Your modem is not listed, and you have an installation disk or .inf file provided by the manufacturer | Click **Have Disk** and specify the location of the disk or .inf file. |

5. Follow the remaining instructions provided by the Install New Modem wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

- If possible, get the latest installation disk or modem .inf file designed for your operating system from the manufacturer. This can usually be downloaded from the manufacturer's Web pages. If the installation files for your operating system are unavailable, you can try using an installation disk or .inf file for Windows 95, Windows 98, Windows NT, Windows 2000, Windows Millennium Edition, or Windows XP. It may work correctly in the Windows Server 2003 family as well.

- If no installation file is available for the modem, the manufacturer's instructions that came with the modem may specify a compatible modem. Otherwise, click **(Standard Modem Types)** in **Manufacturers**, and try choosing the standard modem with the speed matching your modem.

- If a cable is not provided with an external modem, refer to the manufacturer's instructions for cabling requirements. Most but not all off-the-shelf cables will work. Some cables do not have all the pins connected. Do not use the 9-to-25 pin converters that come with most mouse hardware, because some of them do not carry modem signals.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Installing a modem
Remove a modem
Install a modem
Configure an installed modem
Change the COM port for a modem

# Configure an installed modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure an installed modem**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem you want to configure, and then click **Properties**.
3. Enter the changes you want:

| To | Do this |
|---|---|
| Change the speaker volume, maximum port speed, or wait for a dial tone before dialing. | Click the **Modem** tab. |
| View diagnostic information, including commands supported by your modem, or to view a log of modem activity. | Click the **Diagnostics** tab. |
| Add extra initialization commands. | Click the **Advanced** tab. |
| Change call preferences (manual dialing or disconnecting) or data connection preferences (port speed, data protocol, compression, or flow control). | Click the **Advanced** tab, and then click **Change Default Preferences**. |
| Change hardware settings (data bits, parity, stop bits, or modulation). | Click the **Advanced** tab, click **Change Default Preferences**, and then click the **Advanced** tab. |

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Not all default configurations in the connection interface are secure. Ensure that you change settings appropriately.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing a modem
Remove a modem
Install a modem
Install an unsupported modem
Change the COM port for a modem

# Change the COM port for a modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the COM port for a modem**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem you want to configure, and then click **Properties**.
3. On the **Advanced** tab, click **Advanced Port Settings**.

   If this option is not displayed, the modem does not support changing the assigned COM port
4. In **Com Port Number**, click the port number you want to use.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Do not change this setting unless it is required by your system configuration. The communications port is normally assigned correctly when you install a modem.
- The advanced COM port settings are available only for certain brands of modems. If you do not have one of these modems, to change the assigned COM port, you must reinstall the modem and choose the desired COM port.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing a modem
Remove a modem
Install a modem
Install an unsupported modem
Configure an installed modem

# Remove a modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a modem**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to delete.
3. Click **Remove**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing a modem
Install a modem
Install an unsupported modem
Configure an installed modem
Change the COM port for a modem

# Install and Configure Modems on Multiple Ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and configure modems on multiple ports**

- Duplicate a modem installation on multiple ports
- Copy modem properties to modems on multiple ports

# Duplicate a modem installation on multiple ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To duplicate a modem installation on multiple ports**

1. Install the modem you want to duplicate on at least one port.

   See Related Topics if you have not done this already.

2. Open Phone and Modem Options in Control Panel.

3. On the **Modems** tab, right-click the modem you want to duplicate to other ports.

4. Click **Duplicate**.

5. Choose to duplicate the modem installation on all ports in the list or only on selected ports.

   - Click **All ports** to duplicate the modem installation on all ports shown on the list.

   - Click **Selected ports** and then select the ports to duplicate the modem installation only on selected ports.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

- Use this procedure to quickly duplicate a modem installation when you are installing the same make and model of modem on more than one port.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install a modem
Install an unsupported modem
Copy modem properties to modems on multiple ports

# Copy modem properties to modems on multiple ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy modem properties to modems on multiple ports**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, right-click the modem with the properties you want to copy.
3. Click **Copy properties**.
4. Right-click the modem to which you want to copy the properties, and then click **Apply properties**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Use this procedure to quickly duplicate changes to modem properties from one modem to modems on other ports.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure an installed modem
Duplicate a modem installation on multiple ports

# Test Modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Test modems**

- Test a modem
- Log and view modem commands

# Test a modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test a modem**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to test.
3. Click **Properties**, then click **Diagnostics**, then click **Query Modem**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- This procedure sends standard AT commands to your modem and displays the responses to indicate features supported by your modem.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modem diagnostics
Log and view modem commands
Troubleshooting modems

# Log and view modem commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To log and view modem commands**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem whose commands you want to log.
3. Click **Properties**, then click the **Diagnostics** tab.
4. In **Logging**, select the **Append to Log** check box if you are working on a computer running Windows XP or **Record a Log** check box if you are working on a computer running a member of the Windows Server 2003 family.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Click **View Log** to display the log file. This file identifies the modem and displays information about how it is functioning.
- Commands sent to the modem are captured in the file:

  *systemroot*\System32\ModemLog_*Model*.txt.

  The path is usually C:\Windows\System32\. *Model* is the name of the modem as it appears in the list of installed modems on the **Modems** tab of **Phone and Modem Options**.
- In Windows XP, logging is always turned on, and the log is overwritten at the beginning of every session unless you select the **Append to Log** check box. In the Windows Server 2003 family, logging is turned off unless you select the **Record a Log** check box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Modem diagnostics
Test a modem
Log and view modem commands
Troubleshooting modems

# Change How Modems Initialize and Connect

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Change how modems initialize and connect**

- Specify initialization commands
- Select a country/region
- Wait for dial tone before dialing
- Adjust modem speaker volume
- Specify disconnect times

# Specify initialization commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify initialization commands**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to configure.
3. Click **Properties**, then click the **Advanced** tab.
4. In **Extra initialization commands**, type the commands that should be sent to the modem when it is initialized.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Type the commands exactly as they should be sent, with no extra spaces or characters. The AT prefix is unnecessary.
- The commands you enter here are sent after the standard initialization commands for the selected modem. This enables you to change the modem behavior or configuration as required by your communication software or to communicate with a service provider. For more information about the commands recognized by your modem, refer to the manufacturer's documentation or Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Wait for dial tone before dialing
Specify disconnect times
Adjust modem speaker volume
Select a country/region

# Select a country/region

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To select a country/region**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to configure.
3. Click **Properties**, and then click the **Advanced** tab.
4. In **Country/Region Select**, click the country/region from which you are dialing.

## Notes

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- You might not be able to configure your modem to accommodate the telephone networks of different countries/regions. If you should be able to configure your modem for different countries/regions but **Country/Region Select** does not appear, contact your modem manufacturer for an update.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Wait for dial tone before dialing
Specify initialization commands
Specify disconnect times
Adjust modem speaker volume
Dialing rules overview

# Wait for dial tone before dialing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To wait for dial tone before dialing**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to configure, then click **Properties**.
3. On the **Modem** tab, select the **Wait for dial tone before dialing** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- If you clear the **Wait for dial tone before dialing** check box, your modem will dial without waiting for a dial tone. This may be necessary if a phone line does not provide a standard dial tone.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify initialization commands
Specify disconnect times
Adjust modem speaker volume
Select a country/region
Dialing rules overview

# Adjust modem speaker volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To adjust modem speaker volume**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to configure, then click **Properties**.
3. On the **Modem** tab, under **Speaker volume**, move the slider to the left or right to adjust the volume.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Some modems do not support this functionality.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Wait for dial tone before dialing
Specify initialization commands
Specify disconnect times
Select a country/region

# Specify disconnect times

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify disconnect times**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to configure, then click **Properties**.
3. Click the **Advanced** tab, then click **Change Default Preferences**.
4. On the **General** tab, under **Call preferences**:

| To | Do this |
|---|---|
| Automatically disconnect when a connection is not used for a specified time. | Select the **Disconnect a call if idle for more than** check box, and then type the number of minutes the modem should wait before disconnecting. |
| Change the time allowed to make a new connection. | Type the number of seconds in **Cancel the call if not connected within**. |

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Note that these are the default settings. If these parameters are specified by a program, the program's settings take precedence.
- If you are using a network connection, the disconnect time is determined by the redialing options for the network connection rather than by the setting in the modem's properties. See Related Topics for the procedure to configure a connection.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Wait for dial tone before dialing
Specify initialization commands
Adjust modem speaker volume
Select a country/region
Dialing rules overview
Configure a connection to a remote network

# Change Modem Data and Hardware Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Change modem data and hardware settings**

- Change the maximum modem port speed
- Change data connection preferences
- Change hardware settings
- Designate hardware or software compression

# Change the maximum modem port speed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the maximum modem port speed**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to configure, then click **Properties**.
3. On the **Modem** tab, in the **Maximum Port Speed** list, click the speed for the modem.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- This value is usually set correctly when the modem is installed.
- This speed represents the maximum speed programs are allowed to transmit data to the modem and is normally faster than the modem speed. For example, this is usually set to 57,600 bps for a 33.6 Kbps (V.34) modem.
- Changing this setting also sets the current speed for the modem. To set a different value for the current speed, see Related Topics for the procedure to change the data connection preferences.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change hardware settings
Change data connection preferences
Designate hardware or software compression
Improving modem throughput speeds
Attaining fast speeds with a 56 Kbps modem
Optimizing data transfer speed
Data transfer speed

# Change data connection preferences

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change data connection preferences**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to configure.
3. Click **Properties**, then click the **Advanced** tab.
4. Click **Change Default Preferences**.
5. On the **General** tab, under **Data Connection Preferences**, change the setting for **Port speed**, **Data protocol**, **Compression**, or **Flow control**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- These values are usually set correctly when the modem is installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the maximum modem port speed
Change hardware settings
Designate hardware or software compression
Data compression
Data transfer speed
Optimizing data transfer speed

# Change hardware settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change hardware settings**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to configure.
3. Click **Properties**, then click the **Advanced** tab.
4. Click **Change Default Preferences**, then click the **Advanced** tab.
5. Under **Hardware Settings**, change the setting for **Data bits**, **Parity**, **Stop bits**, or **Modulation**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- These values are usually set correctly when the modem is installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the maximum modem port speed
Change data connection preferences
Designate hardware or software compression

# Designate hardware or software compression

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To designate hardware or software compression**

1. Open Phone and Modem Options in Control Panel.
2. On the **Modems** tab, click the modem that you want to configure.
3. Click **Properties**, then click the **Advanced** tab.
4. Click **Change Default Preferences**.
5. In **Data Connection Preferences**, in the **Compression** list, click **Enabled** or **Disabled**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Some communications programs provide efficient software compression. If you are using a program that provides software compression, turning off hardware compression often provides the highest data throughput speed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the maximum modem port speed
Change hardware settings
Change data connection preferences
Data compression

# Modems Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

- Modem overview
- Understanding Modems
- Using Modems

# Modem overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modem overview**

Telephone lines that are used in a conventional phone system are designed to transmit human speech with analog signals. Analog signals vary continuously, like waves, along their length. Computers store and process data in digital format and communicate both internally and with each other in binary digits.

When two computers communicate with each other over a conventional phone line, a modem translates the binary information from the computer at the sending end to an analog signal that can pass over the phone line. At the receiving end, another modem translates the analog signal back to binary information that can be used by the computer. The conversion from binary to analog information is called *modulation* and the conversion back from analog to binary is called *demodulation*. The word *modem*, short for modulator/demodulator, is the name for the device that performs these conversions.

Both standard modems and fax machines perform this conversion. When you communicate over analog telephone lines, you use standard modems. There are also ISDN and cable modems that provide digital communication. They enable computers to transmit digital information directly without the conversion over special lines. The most recent developments with the fastest transmissions are broadband technologies such as cable or DSL (Digital Subscriber Line) that transmit over cable or phone lines as waves similar to radio or television.

The speed at which modems transmit data is called the *throughput*. This is measured in bits per second (bps).

The conversion between digital and analog data is governed by proprietary and international standards, called *modem protocols*. Improving protocols have allowed faster data throughput rates. Besides using the available bandwidth more effectively, protocols have incorporated data compression to speed the throughput and error control for greater reliability. The protocols for modems operating at 56 kilobits per second (Kbps) are now designed to support data that is received from an Internet service provider that is connected to a digital line but that sends data with the standard digital-to-analog conversion.

The Microsoft® Windows Server 2003 family supports many different modems. If you are looking for a new modem that is compatible with Windows, see the compatibility information in Support resources. Microsoft has tested these modems with Network Connections. If you have a modem that is not on this list, it may still work perfectly well with Windows, especially if the manufacturer provides an installation disk or .inf file for Windows. Some modems are compatible with a supported modem and may be installed by choosing the supported information. Check the documentation provided with the modem for details. You may also find installation files or other useful information at the manufacturer's Web site.

# Understanding Modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding modems**

This section presents concepts related to modem use:

- Data transfer speed
- Data compression
- Flow control
- Hardware settings

# Data transfer speed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data transfer speed**

When you initiate a dial-up connection, there are three separate components that determine the speed of the connection:

- The *local interface speed* is the speed of the connection between your computer and your modem, also called the DTE rate.
- The *connection speed* is the speed at which data is transmitted between modems over the communication lines, also called the DCE rate.
- The *remote interface speed* is the speed at which data is transmitted between the receiving modem and computer.

The actual speed at which data is transmitted between computers depends on all three rates. When you buy a modem, the manufacturer normally specifies the maximum connection speed. The maximum port speed that can be set in **Phone and Modem Options** represents the maximum local interface speed that communications programs can use to send data to the modem. The remote interface speed may also limit transmission speeds. For example, an online service provider must support the same 56 Kbps protocol used by your modem for you to receive high speed transmissions with a 56 Kbps modem. (See Attaining fast speeds with a 56 Kbps modem.)

Most high-speed modems today fix their local interface speed at a given value instead of changing it according to the connection speed. The local interface speed is then normally higher than the actual transmission of data over the phone line. For example, a V.34 modem that supports 33.6 Kbps may have its maximum port speed set for 115.2 Kbps.

This difference, called speed buffering, allows a modem to seem faster than the actual connection speed. Using this, a modem can compress the data and send it in fewer bits. Since there are fewer bits, it takes less time to transmit than the uncompressed data. The V.34 modem may then obtain throughput rates of 33.6 Kbps even though the line is carrying fewer actual bits than that.

# Data compression

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Data compression**
Data compression enables information to be transmitted beyond the actual connection speed. Data, particularly text and graphics, usually contain repeated sequences of identical information. Data compression works by replacing many characters of repeated information with a few characters and transmitting only one copy of repeated sequences of data.

Common methods for compressing data include MNP 3, MNP 4, MNP 5, and V.42bis.

The effectiveness of compression varies, depending on the data. Text and graphics may often be compressed considerably. Compression for data that is already compressed may be less efficient because the modem is trying to compress data that is already compressed. V.42bis handles compressed data better than the other protocols.

For any error correction or data compression protocol to work, modems on both ends of the connection must support the protocol. Once two modems are connected, they automatically negotiate to determine the best mutual protocol. For example, if V.42bis fails, the backup is MNP 5, and if it fails, the backup is "no compression."

To achieve maximum rates when transferring data, hardware flow control is required, and the maximum port speed must be set for two to four times the rated speed of the modem. This is usually configured correctly when you install a modem. If you know which type of compression is supported by your modem, you can calculate this value by multiplying the modem speed by the ratio shown above. For example, for a 14.4 kilobit per second (Kbps) modem with V.42bis compression, you would perform the following calculation.

14.4 Kbps x 4 = 57.6 Kbps

This value is the maximum rate that data may be transferred by the modem. The actual throughput depends on the data.

**Software data compression**
Communication software, such as Network Connections, may support data compression. For example, using a 56 Kbps V.90 modem, you can enable software compression, and experience an average throughput of 44 Kbps. Tests show that software compression can result in higher data transfer rates than hardware compression.

If you use software compression, you may want to check the settings. For best results with software compression, you can check your modem's properties for the following:

- Enable modem error correction by setting **Data Protocol** to **Standard EC** or **Forced EC**.
- Disable hardware compression by setting **Compression** to **Disabled**.
- Enable the software compression provided by your communication software.

See Change data connection preferences to change this setting for a modem.

# Flow control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Flow control**
There are six different links involved in a connection: transmitting computer to transmitting modem, transmitting modem to receiving modem, receiving modem to computer, and the reverse of all three links. They can all have different data transmission rates. When the receiving modem is temporarily unable to accept the data, it needs a way to tell the transmitting modem to slow down or wait for it. Flow control is the method by which a modem controls the rate at which other modems send data to it.

If flow control is improperly configured, it may be impossible to connect to a remote system, the transfer rate may be considerably slowed, or the connection may be disrupted. If you see many errors and retransmittal of data when downloading files, check both the modem and the communication program settings for flow control. Flow control settings for both the communication program and the modem must be the same. Many communication programs set these automatically based on the modem, but some must be configured separately.

**Hardware flow control**
Hardware flow (RTS/CTS) control depends on the modem to control the flow of data. This should be used with all high-speed modems or modems that compress data.

**Software flow control**
Software flow (also called XON/XOFF or CTRL+S/CTRL+Q) control uses data characters to indicate that the data flow should start or stop. This enables a modem to send a control character to signal another modem to stop transmitting while it catches up.

Software flow control is slower and usually less desirable than hardware flow control. Software flow control is used only for transmitting text. It cannot be used for binary file transfer because binary data may contain the special flow control characters.

See Change data connection preferences to set hardware or software flow control for a modem.

# Hardware settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware settings**
The connection preferences for a modem provide hardware settings that determine how data is sent and reassembled by modems. To change hardware settings, see Change hardware settings.

**Data bits**
Data bits are the number of bits in a word. Most systems now use eight bits to represent a single data character (extended ASCII). In rare instances, some older systems still use seven bits.

**Parity**
Parity controls how modems check for errors. Line noise can insert extra bits into data being transmitted over a phone line. In older modulation types, parity checked for these errors. With parity checking, the transmitting modem adds a *parity bit* to the data packet in order to make the number of 1 bits in the packet odd or even. The receiving modem adds up the number of 1 bits that is received and accepts or rejects the packet depending on whether the sum agrees with the parity bit.

Several settings for parity are available:

| Set this | To do this |
|---|---|
| Even | Set the parity bit to 0 or 1 to make the number of 1 bits even. |
| Odd | Set the parity bit to 0 or 1 to make the number of 1 bits odd. |
| None | Send no parity bit. |
| Mark | Set the parity bit always to 1. |
| Space | Set the parity bit always to 0. |

Most modem connections now use more reliable and sophisticated methods of error checking, so this is usually set to None.

**Stop bits**
Stop bits frame data packets in asynchronous communication. These tell the receiving modem that a byte has been sent. Modern asynchronous protocols never require more than one stop bit.

**Modulation**
In the hardware settings, modulation is the protocol that a modem uses to translate digital information to the audio tones transmitted over a standard telephone line. Both computers must be using the same type of modulation to exchange information successfully.

You can switch the modulation type to be compatible with the modem signals for the computer to which you are trying to connect. Most modems fall into the standard category, but if you are having trouble connecting, try switching to a nonstandard modulation type.

# Using Modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using modems**

This section contains information you can use to install and configure modems:

- Installing a modem
- Setting hardware connections
- Optimizing data transfer speed
- Attaining fast speeds with a 56 Kbps modem
- Modem diagnostics
- Connecting using your modem

# Installing a modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing a modem**
To install the modem hardware on your computer, refer to the manufacturer's documentation. Generally, the following instructions apply:

- If you have an internal modem with jumpers, set the jumpers for Plug and Play for the Windows Server 2003 family or, if that option is unavailable, for Windows 95, Windows 98, Windows NT, Windows 2000, or Windows Millennium Edition.
- Install an internal modem in an empty slot. Attach an external serial modem to an unused COM port.
- If you have an external modem, ensure that it is attached to the computer, plugged into the power source, and turned on before you turn on the computer and start Windows.
- Ensure that your modem is properly connected to the phone line and to your computer.

Most modems manufactured now are compatible with Plug and Play and installed automatically after they are connected to the computer. However, older modems and external serial modems may not be detected. If your modem is not installed automatically, then use **Phone and Modem Options** in **Control Panel** to install it manually. To install a modem, see Install a modem. To choose the modem type manually, see Install an unsupported modem.

# Setting hardware connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting hardware connections**
A modem has to be connected to the telephone system and to a power source.

**Telephone connections**
Both internal and external modems use a modular phone cord, called an RJ-11 coupler. Usually, a new modem includes the coupler. If you have an older, four-pronged jack, you may also need an adapter to connect to your phone line.

If your office has digital, PBX or multiline phones, you cannot connect a modem directly to the phone system. You need to either install a standard line or install a device that connects your modem to one of these systems. See your computer or telephone supplier for specialized solutions.

**Power connections**
Internal modems are powered directly by the computer. External modems usually come with a power adapter, but are sometimes powered by the computer. Most portable modems for portable computers use batteries, although they often come with an adapter too. Some are powered directly from the computer. Modems use different sizes and styles of adapters, so if you lose the adapter, you need to get a new one from the modem or computer manufacturer. Refer to the manufacturer's documentation for additional information about the power requirements of your computer.

**Cabling Connections**
External serial modems connect to the computer with a serial cable, also called an RS-232 cable. If a cable is not provided with your modem, make sure you get one that fits both your computer's serial port and the modem's port.

# Optimizing data transfer speed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Optimizing data transfer speed**

The speed settings for a modem are usually set correctly during installation. There are two different speeds set in **Phone and Modem Options**: **Maximum Port Speed** in the general modem properties and **Port Speed** in the default preferences. Both of these specify the speed at which communication programs send data to the modem, not the speed at which data is transferred between your modem and the remote modem.

The **Maximum Port Speed** specifies the maximum speed programs can transmit data to a modem. Most programs set their own speed, so this limits the speed at which they can send data. The **Port Speed** is the current port speed for the modem. These two settings are usually the same, although there are some instances in which one should be changed. Do not change these speeds just to match the rated speed for your modem.

**Measuring your modem speed**

There are several separate components that determine the speed of the connection. When a program or hardware reports transmission speeds, note carefully what is being reported. For example, if a V.34 modem rated at 33.6 kilobits per second (Kbps) reports a connection speed of 38.4 Kbps or higher, it is reporting the port speed that data is transmitted from the computer to the modem rather than the rate that data is transmitted between modems.

For practical purposes, you want to know the *throughput speed*, the speed that your data is transferred between your communications program and the program on another computer. Communications programs usually report the throughput speed and normally provide the speed in either bits per second (bps) or characters per second (cps). For serial communications, speed in bits per second can be divided by 10 to approximate the characters transmitted per second.

Check the throughput speed after you start communicating with the remote computer. The initial throughput speed is not necessarily a good estimate of the transmission speed. Modems often shift speed up or slow down after the initial connection. Modems can fall forward or fall back depending on the quality of the connection and other factors.

**Transmission speeds for direct serial connections**

When you create a direct serial connection to another computer or device, you may need to adjust the maximum port speed for the connection. The initial setting for a direct serial connection is 19.2 Kbps. For a Windows CE device or other device that supports higher data transmission speeds, change the maximum port speed to the correct speed. Refer to the manufacturer's instructions for the transmission speed supported by the device.

For additional information, see Attaining fast speeds with a 56 Kbps modem, Data transfer speed, Improving modem throughput speeds, Change data connection preferences, or Change the maximum modem port speed.

# Attaining fast speeds with a 56 Kbps modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Attaining fast speeds with a 56 Kbps modem**

A modem connection must fulfill three requirements to support a 56 kilobits per second (Kbps) (also called V.90) connection.

- The host server must use a digital connection to the network. Your Internet service provider can tell you if they support 56 Kbps service.
- Both ends of the connection must support the same protocol, the V.90 standard or either of its predecessors, K56flex, or 3COM/USR X2. For example, if your Internet service provider has a V.90 device, your modem must support the V.90 protocol.
- There can only be one analog connection between your modem and the host computer. The phone line in most homes is an analog line.

If a connection does not meet these requirements, a modem falls back to the fastest protocol that works for the connection. For example, a 56 Kbps V.90 modem falls back to the 33.6 Kbps V.34 protocol if it cannot make a V.90 connection. Even if your connection fulfills these requirements, other factors may reduce either the transmission speed or the number of times that you successfully obtain the highest speed connection. For example, old lines or lines that are subject to interference may reduce transmission speeds. Maximum throughput speeds of 26 Kbps are not unusual in these cases.

Devices to improve the quality of your telephone service may also hamper 56 Kbps V.90 modem connections. Load coils found on long wire lengths to improve voice quality do not usually prevent V.90 connections, but can reduce the speed. Digital pads, which balance the volume of voice calls, usually do not prevent V.90 connections, but they can reduce the speed. Analog pads prevent V.90 connections, because they convert the digital data to analog to balance the volume and then back to digital. This inserts an additional analog section in the line.

In practice, the 56 Kbps speed supported by the V.90 and other protocols is unattainable. U.S. government regulations to safeguard public phone systems right now limit transmission speeds to 53 Kbps. Phone-line noise and other limitations of phone systems usually keep average transmissions in the 40 to 50 Kbps range.

For more information, see Data transfer speed, Improving modem throughput speeds, and Optimizing data transfer speed.

# Modem diagnostics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modem diagnostics**

You can verify that your modem is working properly by using the diagnostics that are available through the **Phone and Modem Options** icon in Control Panel. (For more information on this procedure, see Test a modem.)

When you query a modem, Windows tries the following commands and displays the results.

| Command | Response |
|---|---|
| ATQ0V1E0 | Initializes the query. |
| AT+GMM | Model identification (ITU V.250 recommendation is not supported by all modems). |
| AT+FCLASS=? | Fax classes supported by the modem, if any. |
| AT#CLS=? | Shows whether the modem supports the Rockwell voice command set. |
| ATI*n* | Displays manufacturer's information for *n* = 1 through 7. This provides information such as the port speed, the result of a checksum test, and the model information. Check the manufacturer's documentation for the expected results. |
| AT+GCI? | Displays the current selected country/region. |
| AT+GCI=? | Displays all countries/regions supported by this modem. |

You can also record a log of commands as they are sent to the modem by communication programs or the operating system. For more information, see Log and view modem commands.

# Connecting using your modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connecting using your modem**

The following uses of a modem are supported:

| To | See |
|---|---|
| Make a phone connection to another computer. | HyperTerminal |
| Make a direct connection another computer. | Using direct connections |
| Log onto a network or connect to another computer. | About Network Connections |

**Notes**

- You can also install communications programs provided by other manufacturers. They can provide added or different functionality or serve a specific purpose, such as connecting to an online service.
- Because connection settings are not cleared after a call is completed, connection details (including the phone number) can be misused by malicious users.

# Modems Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

- Troubleshooting modems
- Improving modem throughput speeds
- Information about specific modems

# Troubleshooting modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting modems**
What problem are you having?

- The modem diagnostics indicate that an external serial modem is not receiving commands.
- The modem cable is good, but the modem still does not receive commands.
- The modem is installed correctly but the diagnostics indicate that it is not responding correctly.
- According to the modem diagnostics, the modem works but you still cannot make a connection.
- You cannot connect or are having problems connecting to another computer or an online service provider.
- Your PCMCIA modem card was not detected automatically when you inserted it.
- An RS-232 serial device is functioning poorly or not at all.
- A connection in Network Connections reports that a port is in use or not configured for remote access.

**The modem diagnostics indicate that an external serial modem is not receiving commands.**
**Cause:** If the modem diagnostics indicate that the modem is not receiving commands, the modem cabling may be faulty.

**Solution:** Try connecting the modem with a new cable.

**The modem cable is good, but the modem still does not receive commands.**
**Cause:** The modem is installed incorrectly.

**Solution:** Check the modem's documentation to make sure that you installed it correctly.

**The modem is installed correctly but the diagnostics indicate that it is not responding correctly.**
**Cause:** The incorrect model and make were specified during installation in Windows or an obsolete installation (.inf) file was used to install the modem.

**Solution:** Check the modem documentation for the correct model and make or a compatible model and make, and reinstall the modem. You can check with the manufacturer to see if there are more recent installation files available for installing the modem.

**According to the modem diagnostics, the modem works but you still cannot make a connection.**
**Cause:** The modem is connected incorrectly to the phone line or there is a problem with the phone line.

**Solution:** Check the connection to the phone line. If the modem is connected correctly, have a telephone professional check your phone line.

**You cannot connect or are having problems connecting to another computer or an online service provider.**
**Cause:** The data connection parameters (see Hardware settings) for the two modems are not the same. Data connection parameters for two modems need to be identical for them to connect successfully.

**Solution:** Check the hardware settings for both computers. Typical settings are eight data bits, None for parity, one stop bit.

If you are connecting to an online service provider, refer to the documentation provided by your service provider for the correct settings. These are typically eight data bits, None for parity, one stop bit. Bulletin boards and most service providers use these settings. If these do not work, try seven data bits, even parity, and one stop bit. A few online service providers use these settings. Other settings are extremely rare.

**Your PCMCIA modem card was not detected automatically when you inserted it.**
**Cause:** The card's built-in COM port is not configured.

**Solution:** Use Add Hardware in Control Panel to configure the card's built-in COM port. You can then install the PCMCIA modem by using the Phone and Modem Options in Control Panel. (See Install a modem.)

**An RS-232 serial device is functioning poorly or not at all.**
**Cause:** Some RS-232 serial devices now support speeds faster than the 115.2 Kbps currently supported by the standard serial ports installed on most computers. (For example, some serial devices support 230 Kbps.)

**Solution:** If the **Port speed** in the connection preferences is faster than 115.2 Kbps, then try reducing the speed to 115.2 Kbps. You can leave the **Maximum Port Speed** in the modem properties at the higher setting. If the serial port on your computer supports speeds faster than 115.2 Kbps, then set the **Port speed** to whichever is slower, the serial port speed or the device speed.

**A connection in Network Connections reports that a port is in use or not configured for remote access.**
**Cause:** There are several possible causes:

- When you started Windows, a Plug and Play modem was not connected or it was turned off, so it was not detected correctly.
- The port is in use by another program, such as Fax.
- The connection may not have been updated to reflect changes in communications equipment, such as new serial ports or modems.

**Solution:** Use one of the following solutions:

- If the modem was disconnected or turned off, reconnect or turn on your modem, and then either restart Windows or use Device Manager to scan for new hardware.
- If the communications port is already in use by another connection, hang up the connection, or modify the connection to use a different port.
- If you have changed your hardware (such as a new serial port or modem), you must reconfigure the connection.

**See also:** Configure a connection to a remote network; Reinstall a Plug and Play device

**Notes**

- To diagnose and correct modem problems with the interactive modem troubleshooter, see List of troubleshooters and click **Modem**. For modem diagnostics, see Test a modem. Refer to the modem manufacturer's documentation or customer service representative for the correct modem commands and responses.
- It is usually a good idea to note changes to modem settings and switches as you make them. This enables you to return the modem to its previous state, if necessary.
- If you have modem problems, you may want to report them to the modem manufacturer or your supplier. In your report, include your modem log, which records information about your attempts to connect to other systems. For information about creating a modem log, see Log and view modem commands.

# Improving modem throughput speeds

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Improving modem throughput speeds**

If you have a 33.6 kilobits per second (Kbps) (V.34 protocol) modem and consistently connect at 26 Kbps or higher, or you have a 56 Kbps (V.90, K56flex, or X2 protocols) modem and consistently receive data from your Internet service provider or online provider at 40,000 bps or higher, there is not much you can do to go faster. Higher speeds are not possible with most phone circuits. The rated speeds are only possible with a near-perfect connection.

If you consistently connect at slower speeds (24 Kbps or less for a 33.6 Kbps modem or less than 40 Kbps for a 56 Kbps modem), you can try the following:

- Check that the correct manufacturer and model are listed for your modem in Phone and Modem Options. If your modem is not listed correctly, try following the steps in Install an unsupported modem to install the correct modem.
- Check if the manufacturer has a new installation disk or .inf file available. Files can often be downloaded from a manufacturer's Web site or bulletin board.
- For a 33.6 Kbps modem, make sure the **Maximum Port Speed** is set at 38,400 bps or higher (57,600 is recommended). For a 56 Kbps modem make sure it is set at 115,200 bps. See Change the maximum modem port speed for instructions.
- Make sure the modem is operating at the correct temperature. Most modems operate better if they do not get too hot or too cold. Do not block air flow by placing things on top, under, or around your modem or computer.
- Try connecting to numbers that you know support a given speed. This lets you know whether the problem is at your end or the other end.
- Make sure a local number is actually a local call. Many online service providers forward local calls to a distant location to provide local calling service. This works well at 14,400 bps, but can degrade the signal and limit higher speeds.
- If your phone line has a surge protector, you can try using your modem without it. Note, however, that doing so eliminates the protection provided to your equipment by the surge protector.
- Have an expert check the telephone wiring between the telephone company connection and your computer. Faulty wiring, soldered connections, long connecting cables, inferior equipment (phones, fax machines, or even other modems) can all degrade the signal. To distinguish between wiring and line problems, check the voice quality, both while the local wiring is connected and on the telephone company line while it is disconnected from the local wiring. Many phone companies now connect their service with a modular plug, which consumers can disconnect to perform this check.

Telephone companies normally do not guarantee data throughput speeds. When you ask your telephone company to fix a noisy line, you may be told you need a data line or *conditioned circuit*. This is usually an expensive solution. Modems can perform well on a good quality voice circuit. Ideally you want your system tested with a sophisticated line tester (a "bit error rate tester" or "data test set") rather than the regular line tester. You can also ask the company to switch you to a different cable pair from the central office. Sometimes, people solve the problem by ordering a different phone number.

For additional information, see Attaining fast speeds with a 56 Kbps modem, Data transfer speed, and Optimizing data transfer speed

# Information about specific modems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Information about specific modems**
This section gives additional information about specific modems that are supported by Network Connections. For a list of supported modems, see Support resources.

**Characteristics of specific modems**
Modems supported by Network Connections do not necessarily work in all modes with other modems on the list of compatible modems.

For example, some V.32bis modems do not work with port speeds of 14,400 bps although they communicate at a modem-to-modem speed of 14,400 bps. If your V.32bis modem does not work at an initial speed of 14,400 bps, you should change the speed to 19,200 bps. For information about changing initial modem speed, see Configure a modem for a dial-up connection for a dial-up connection.

**Internal modems**
Certain internal modems cannot be automatically detected. You must install the modem manually through Control Panel.

**Callback**
Modems supported as server modems may not support callback when used as client modems.

# Wireless Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless networking**
Wireless technology makes it possible for you to use a wide range of devices to access data from anywhere in the world. Wireless networks reduce or eliminate the high cost of laying expensive fiber and cabling and provide backup functionality for wired networks. To ensure that wireless networks and devices are compatible, cost-effective, and secure, organizations and special interest groups are working to develop standards for wireless communications.

The Microsoft® Windows Server 2003 family provides support for infrared communication and for 802.11 wireless networks. In Windows Server 2003, Standard Edition, support is provided for infrared communication. 802.11 wireless networking is supported in all versions of the Windows Server 2003 family.

- For help with specific tasks, see Wireless Networking How To....
- For general background information, see Wireless Networking Concepts.
- For problem-solving instructions, see Windows Server 2003 Wireless Troubleshooting.

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition supports infrared networking.

# Wireless Networking How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Access Active Directory-based wireless network policies
- Define Active Directory-based Wireless Network Policies
- Define Wireless Network Policies on a Client Computer
- Use Resultant Set of Policy (RSoP) to View Wireless Network Policy Assignments
- Monitor Wireless Network Activity
- Manage Infrared Communication

# Access Active Directory-based wireless network policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To access Active Directory-based wireless network policies**

1. Open Active Directory Users and Computers.

2. In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

   **Where?**

   - Active Directory Users and Computers [*DomainControllerName.DomainName*]/*Domain*/*OrganizationalUnit*/*ChildOrganizationalUnit*...

3. Click **Properties**, and then click the **Group Policy** tab.

4. Click **Edit** to open the Group Policy object that you want to edit. Or, click **New** to create a new Group Policy object, and then click **Edit**.

5. In the Group Policy console tree, click **Wireless Network (IEEE 802.11) Policies**.

   **Where?**

   - *PolicyName* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/Wireless Network (IEEE 802.11) Policies

## Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have permission to edit Group Policy objects (for information, see Related Topics). As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.

- In **Wireless Network (IEEE 802.11) Policies**, you can add, edit, or remove wireless network policies, define preferred wireless networks, and define 802.1X authentication. For more information, see Related Topics.

- You can define only one wireless network policy for a domain or an organizational unit. If a wireless network policy already exists, the option to create a new wireless network policy will not be available.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add, edit, or remove Active Directory-based wireless network policies
Define 802.1X authentication for wireless networks in Group Policy
Group Policy (pre-GPMC)
Define preferred wireless networks in Group Policy

Security information for wireless networks
Understanding 802.1X authentication for wireless networks

# Define Active Directory-based Wireless Network Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define Active Directory-based wireless network policies**

- Add, edit, or remove Active Directory-based wireless network policies
- Define preferred wireless networks in Group Policy
- Define 802.1X authentication for wireless networks in Group Policy

# Add, edit, or remove Active Directory-based wireless network policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove Active Directory-based wireless network policies**

1. In **Wireless Network (IEEE 802.11) Policies**, choose whether to add a new policy or edit or remove an existing one:

   - To add a new policy, click **Wireless Network (IEEE 802.11) Policies**. Next, on the **Action** menu, click **Create Wireless Network Policy**. Follow the instructions in the Wireless Network Policy Wizard until the **Properties** dialog box for your new policy appears.
   - To edit an existing policy, right-click the policy that you want to edit, and then click **Properties**.
   - To remove an existing policy, right-click the policy that you want to delete, and then click **Delete**.

2. If you are adding or editing a policy, in **Properties**, click the **General** tab, and then do any of the following:

   - To name the policy or change the existing name of the policy, type a unique name in **Name**.
   - To provide a description of the wireless network policy, such as which groups or domains it affects, type a description in **Description**.
   - To specify how often Active Directory is polled for updates, type a value in **Check for policy changes every***number***minutes**.
   - To specify the type of wireless network that clients can access, click a network type in the **Networks to access** list.
   - To allow clients to use Windows to configure their wireless network settings, select the **Use Windows to configure wireless network settings for clients** check box. This check box is selected by default. For information about this option, see Notes.

     Or, to prevent clients from using Windows to configure their wireless network settings, clear the **Use Windows to configure wireless network settings for clients** check box.
   - To allow clients to automatically connect to available networks that do not appear on the **Preferred Networks** tab, select the **Automatically connect to non-preferred networks** check box.

     Or, to ensure that clients connect only to networks that appear on the **Preferred Networks** tab, clear the **Automatically connect to non-preferred networks** check box.

3. If you are adding or editing a policy, click the **Preferred Networks** tab, and do any of the following:

   - To change the order in which connection attempts to preferred wireless networks are made, click the wireless network that you want to move to a new position on the list, and then click **Move Up** or **Move Down**.
   - To specify a new preferred wireless network, click **Add**, and then configure the settings as needed. For more information, see Related Topics.
   - To change the settings for a preferred wireless network, click the wireless network for which you want to change

settings, click **Edit**, and then change the settings as needed. For more information, see Related Topics.

- To remove a wireless network from the list of preferred networks, under **Networks**, click the wireless network that you want to remove, and then click **Remove**.

### Important

- It is highly recommended that you use 802.1X authentication whenever you connect to an 802.11 wireless network. 802.1X is an IEEE standard that enhances security and deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting. For information about how to define 802.1X authentication settings in Group Policy, see Related Topics.

- For enhanced security, in Windows XP Service Pack 1 and in the Windows Server 2003 family, 802.1X authentication is available only for access point (infrastructure) networks that require the use of a network key (WEP). WEP provides data confidentiality by encrypting the data that is sent between wireless clients and wireless access points. For additional information about security for wireless networks, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have permission to edit Group Policy objects (for more information, see Related Topics). As a security best practice, consider using Run asto perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open **Wireless Network (IEEE 802.11) Policies**, you must access Active Directory-based wireless network policies. For more information, see Related Topics.

- When you allow clients to use Windows to configure their wireless network settings, they can connect to the wireless networks that you specify, and they are notified when new wireless networks are available. After the client selects a wireless network, the Windows Configuration service automatically configures the wireless network adapter to match the settings of that network, and a network connection attempt is made. For more information, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Define preferred wireless networks in Group Policy
Access Active Directory-based wireless network policies
Define 802.1X authentication for wireless networks in Group Policy
Configuring wireless network settings on client computers
Security information for wireless networks
Understanding 802.1X authentication for wireless networks

# Define preferred wireless networks in Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define preferred wireless networks in Group Policy**

1. In **Wireless Network (IEEE 802.11) Policies**, double-click the wireless network policy that you want to configure.
2. On the **Preferred Networks** tab, choose whether to add a new preferred wireless network or edit or remove an existing one:
   - To add a new preferred wireless network, click **Add**.
   - To edit an existing preferred wireless network, click **Edit**.
   - To remove an existing preferred wireless network, click **Remove**.
3. If you are adding or editing a preferred wireless network, in **Properties**, click the **Network Properties** tab, and then do any of the following:
   - To name the preferred wireless network or to change the existing name, type a unique name in Name.
   - To provide a description of the preferred wireless network, such as the type of network, whether a network key is required, and whether IEEE 802.1X authentication is enabled, type a description in **Description**.
   - If a network key is required to connect to the preferred wireless network, do the following:

     To specify that a network key is used for authentication to the wireless network, from the drop-down list in **Network authentication**, select one of the following:

     - **Shared**
     - **WPA**
     - **WPA-PSK**

     If you do not select this check box, open system authentication is used. For more information, see Notes.

| |
|---|
| To specify that a network key is used to encrypt the data that is sent over the network, from the drop-down list in **Data encryption**, select one of the following:<br><br>For Shared, select one of the following:<br><br>• **Disabled**<br><br>• **WEP**<br><br>For WPA, select one of the following:<br><br>• **WEP**<br><br>• **TKIP**<br><br>• **AES**<br><br>For WPA-PSK, select one of the following:<br><br>• **WEP**<br><br>• **TKIP**<br><br>• **AES**<br><br>**Note**<br><br>By default, **Data encryption, WEP** is selected. |
| To specify that a network key is automatically provided for clients (for example, if you are using 802.1X for dynamic key distribution), select the **The key is provided automatically** check box. |

- If the preferred wireless network connection that you are configuring is to a computer-to-computer (ad hoc) network, select the **This is a computer-to-computer (ad hoc) network; wireless access points are not used** check box.

4. To specify that 802.1X authentication is used for the preferred wireless network connection, click the **IEEE 802.1x** tab, and then configure settings as needed. For more information, see Related Topics.

**Important**

- It is highly recommended that you use 802.1X authentication whenever you connect to an 802.11 wireless network. 802.1X is an IEEE standard that enhances security and deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting. For information about how to define 802.1X authentication settings in Group Policy, see Related Topics.

**Important**

- For enhanced security, in Windows XP Service Pack 1 and in the Windows Server 2003 family, 802.1X authentication is available only for access point (infrastructure) networks that require the use of a network key (WEP). WEP provides data confidentiality by encrypting the data that is sent between wireless clients and wireless access points. For additional information about security for wireless networks, see Related Topics.

**Important**

- For enhanced security and connectivity, do not select the **Network authentication (Shared mode)** check box. When this check box is selected, shared key authentication is used, rather than open system authentication. Shared key authentication is less secure than open system authentication because it requires the exchange of a secret key that is

shared by all wireless access points and clients and therefore is more vulnerable to known-text attacks. In addition, if you select this check box for a wireless network that has multiple wireless access points, clients will lose network connectivity when they travel from one wireless access point to a new wireless access point.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have permission to edit Group Policy objects (for more information, see Related Topics). As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open **Wireless Network (IEEE 802.11) Policies**, you must access Active Directory-based wireless network policies. For more information, see Related Topics.

- WPA includes two modes: one using 802.1X and RADIUS authentication, which is simply known as WPA, and another, simpler scheme for SOHO environments using a pre-shared key (known as WPA-PSK). For environments with a RADIUS infrastructure, WPA supports EAP and RADIUS.

- Although WPA support is now mandatory for all Wi-Fi certified hardware, existing network equipment that is not WPA-compliant must be upgraded to support WPA.

- Support for Wireless Policy Group Policy controlled settings is not available in Windows Server operating systems before Windows Server 2003 with Service Pack 1 (SP1). There is no way to script WLAN settings for Windows XP without Windows Server 2003 SP1.

- If the **Network authentication (Shared mode)** check box is not selected (that is, if open system authentication is used), no authentication is provided. Open system authentication only performs identity verification.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add, edit, or remove Active Directory-based wireless network policies
Access Active Directory-based wireless network policies
Define 802.1X authentication for wireless networks in Group Policy
Configuring wireless network settings on client computers
Security information for wireless networks
Understanding 802.1X authentication for wireless networks
Group Policy (pre-GPMC)

# Define 802.1X authentication for wireless networks in Group Policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define 802.1X authentication for wireless networks in Group Policy**

1. In **Wireless Network (IEEE 802.11) Policies**, double-click the wireless network policy for which you want to configure 802.1X authentication.

2. On the **Preferred Networks** tab, under **Networks**, choose whether to configure 802.1X authentication for an existing wireless network or for a new wireless network:

   - To configure 802.1X authentication for an existing wireless network, click the wireless network for which you want to configure 802.1X authentication, and then click **Edit**.

   - To configure 802.1X authentication for a new wireless network, click **Add**.

3. On the **IEEE 802.1x** tab, do one of the following:

   - To enable IEEE 802.1X authentication for this wireless network, select the **Enable network access control using IEEE 802.1x** check box. This check box is selected by default.

   - To disable IEEE 802.1X authentication for this wireless network, clear the **Enable network access control using IEEE 802.1x** check box.

4. In **EAPOL-Start message**, specify whether to transmit Extensible Authentication Protocol over LAN (EAPOL)-Start message packets and, if so, how to transmit them.

5. In **Parameters (seconds)**, specify EAPOL-Start message packet parameters.

6. In **EAP type**, click the EAP type to be used with this wireless network.

7. If you select **Smart Card or other certificate** in **EAP type**, click **Settings** and, in **Smart Card or other Certificate Properties**, do the following:

   - To allow wireless clients to use the certificate that resides on their smart card for authentication, click **Use my smart card**.

   - To allow wireless clients to use the certificate that resides in the certificate store on their computer for authentication, click **Use a certificate on this computer**, and then specify whether to use simple certificate selection.

   - To verify that the server certificate presented to client computers is still valid, select the **Validate server certificate** check box, click the **Connect to these servers** check box, specify the server or servers to which client computers will automatically connect, and then specify the trusted root certification authorities.

   - To view detailed information about the selected root certification authority, click **View Certificate**.

   - To allow users to specify a different user name when the user name in their smart card or certificate is not the same as the user name in the domain to which they are logging on, select the **Use a different user name for the**

**connection** check box.

8. If you select **Protected EAP (PEAP)** in **EAP type**, click **Settings**, and then do the following:

   To verify that the server certificate presented to client computers is still valid, select the **Validate server certificate** check box, click the **Connect to these servers** check box, specify the server or servers to which client computers will automatically connect, and then specify the trusted root certification authorities.

   In **Select Authentication Method**, click the authentication method that clients are to use within PEAP, and then click **Configure**.

   - If you selected **Secured password (EAP-MSCHAP v2)**, in **EAP MSCHAPv2 Properties**, specify whether to use the user name and password (and domain, if applicable) that users on client computers type in the Windows logon screen for authentication, and then click **OK**.

   - If you selected **Smart Card or other certificate**, in **Smart Card or other Certificate Properties**, configure the settings as needed by following the instructions in step 7, and then click **OK**.

   To enable fast reconnect for wireless clients, select the **Enable Fast Reconnect** check box. For more information about PEAP fast reconnect, see Notes.

9. On the **IEEE 802.1x** tab, do the following:

   To specify that client computers attempt authentication to the network if user information or computer information is not available, select the **Authenticate as guest when user or computer information is unavailable** check box.

   To specify that client computers attempt authentication to the network if a user is not logged on, select the **Authenticate as computer when computer information is available** check box, and in **Computer authentication**, click an option to specify how the computer should attempt authentication. For information about each of the options that you can select for **Computer authentication**, see Notes.

**Important**

- It is highly recommended that you use 802.1X authentication whenever you connect to an 802.11 wireless network. 802.1X is an IEEE standard that enhances security and deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting.

- For enhanced security, in Windows XP Service Pack 1 and in the Windows Server 2003 family, 802.1X authentication is available only for access point (infrastructure) networks that require the use of a network key (WEP). WEP provides data confidentiality by encrypting the data that is sent between wireless clients and wireless access points. For additional information about security for wireless networks, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have permission to edit Group Policy objects (for more information, see Related Topics). As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open **Wireless Network (IEEE 802.11) Policies**, you must access Active Directory-based wireless network policies. For more information, see Related Topics.

- To define 802.1X authentication, you must select an existing preferred wireless network, or you must define a new preferred wireless network. For information about how to define preferred wireless networks, see Related Topics.

- PEAP fast reconnect allows roaming users to maintain continuous wireless network connectivity when traveling between different wireless access points on the same network, as long as each wireless access point is configured as a client of the same IAS (RADIUS) server. In addition, both the wireless client and the RADIUS server must have fast reconnect enabled.

- If you select the **Authenticate as computer when computer information is available** check box, you can select one of the following options:

    - **With user authentication**. When this option is selected, when users are not logged on to the computer, authentication is performed using the computer credentials. After a user logs on to the computer, authentication is maintained using the computer credentials. If a user travels to a new wireless access point, authentication is performed using the user credentials.

    - **With user re-authentication (recommended)**. When this option is selected, when users are not logged on to the computer, authentication is performed using the computer credentials. After a user logs on to the computer, authentication is performed using the user credentials. When a user logs off of the computer, authentication is performed using the computer credentials.

    - **Computer only.** When this option is selected, authentication is always performed using the computer credentials. User authentication is never performed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy (pre-GPMC)
Add, edit, or remove Active Directory-based wireless network policies
Access Active Directory-based wireless network policies
Define preferred wireless networks in Group Policy
Security information for wireless networks
Understanding 802.1X authentication for wireless networks

# Define Wireless Network Policies on a Client Computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Define wireless network policies on a client computer**

- Use Windows to configure wireless network settings on a client computer
- Add, edit, or remove wireless network connections on a client computer
- Define a wireless network connection on a client computer
- Define 802.1X authentication for wireless networks on a client computer
- Connect to an available wireless network

# Use Windows to configure wireless network settings on a client computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To use Windows to configure wireless network settings on a client computer

1. Right-click the wireless network connection icon in the notification area, and then click **View Available Wireless Networks**. For more information, see Notes.

2. Under **Related Tasks**, click **Change the order of preferred networks**.

3. On the **Wireless Networks** tab, do one of the following:

    - To use Windows to configure wireless network settings on your computer, select the **Use Windows to configure my wireless network settings** check box. This check box is selected by default. For information about this option, see Notes.

    - If you do not want to use Windows to configure wireless network settings on your computer, clear the **Use Windows to configure my wireless network settings** check box.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- tok:wirelessicon

- When an error is detected that might limit or prevent a connection to a wireless network, the wireless warning icon is displayed in the notification area.

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- When you select the **Use Windows to configure my wireless network settings** check box, you enable the Wireless Configuration service. When the Wireless Configuration service is enabled, you can connect to an existing wireless network, change wireless network connection settings, configure a new wireless network connection, and manage preferred wireless networks. You will be notified when new wireless networks are available. After you select a wireless network, your wireless network adapter will be automatically configured to match the settings of that network, and a network connection attempt will be made.

- The **Wireless Networks** tab appears only if you are using a wireless network adapter that supports the Wireless Configuration service. If you are unsure whether your wireless network adapter supports the Wireless Configuration service, contact the wireless network adapter manufacturer.

- For more information about configuring wireless network settings, including information about the Wireless Configuration service, see Related Topics.

- If you are using non-Microsoft wireless networking software to configure wireless network settings, clear the **Use Windows to configure my wireless network settings** check box.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Show a network connection icon in the notification area
Define 802.1X authentication for wireless networks on a client computer
Security information for wireless networks
Configuring wireless network settings on client computers
Wireless networking overview
Understanding 802.1X authentication for wireless networks
Connect to an available wireless network

# Add, edit, or remove wireless network connections on a client computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add, edit, or remove wireless network connections on a client computer**

1. Right-click the wireless network connection icon in the notification area, and then click **View Available Wireless Networks**. For more information, see Notes.

2. To add, modify, change the preferred order of, or remove wireless networks, under **Related Tasks**, click **Change the order of preferred networks**, and on the **Wireless Networks** tab, do any of the following:

| To | Do this |
|---|---|
| Remove a preferred wireless network | Under **Preferred networks**, click the wireless network that you want to remove from the list, and then click **Remove**. |
| Add a new wireless network | Select the **Use Windows to configure my network settings** check box. Under **Preferred networks**, click **Add**. |
| Automatically connect to available networks that do not appear in the **Preferred networks** list | Under **Preferred networks**, click **Advanced**, and then select the **Automatically connect to non-preferred networks** check box. |
| Modify settings for an existing wireless network | Select the **Use Windows to configure my network settings** check box. Under **Preferred networks**, click the wireless network connection that you want to modify, and then click **Properties**. |
| Change the order in which connection attempts to preferred networks are made | Under **Preferred networks**, click the wireless network that you want to move to a new position in the list, and then click **Move up** or **Move down** until the wireless network is at the desired position. |

3. If you are adding or modifying a wireless network, click the **Association** tab, configure wireless network settings as needed, and then click **OK**. For more information, see Related Topics.

4. To define 802.1X authentication for the wireless network connection, click the **Authentication** tab, configure the settings as needed, and then click **OK**. For more information, see Related Topics.

5. On the **Wireless networks** tab, under **Available networks**, click on **View Wireless Networks**, and then do any of the following:

| To | Do this |
|---|---|
| Update the list of available wireless networks that are within range of your computer | Under **Network Tasks**, click **Refresh network list**. |

| | |
|---|---|
| Connect to a wireless network after configuring network settings on the **Wireless Networks** tab | Under **Choose a wireless network**, click the network to which you want to connect, and then click **Connect**.<br><br>**Important**<br><br>If a network does not broadcast its network name, it does not appear under **Choose a wireless network**. To connect to an access point (infrastructure) network that you know is available but that does not appear under **Available network**, follow the instructions in the previous table for adding a new wireless network. |

**Important**

- It is highly recommended that you use 802.1X authentication whenever you connect to an 802.11 wireless network. 802.1X is an IEEE standard that enhances security and deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting. For more information, see Related Topics.

- For enhanced security, in Windows XP Service Pack 1 and in the Windows Server 2003 family, 802.1X authentication is available only for access point (infrastructure) networks that require the use of a network key (WEP). WEP provides data confidentiality by encrypting the data that is sent between wireless clients and wireless access points. For additional information about security for wireless networks, see Related Topics.

- If you attempt to connect to a computer-to-computer network or an access point network that does not require the use of a network key, the settings in the **Authentication** tab are not available, and you cannot configure 802.1X authentication for the connection.

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- tok:wirelessicon

- When an error is detected that might limit or prevent a connection to a wireless network, the wireless warning icon is displayed in the notification area.

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- The **Wireless Networks** tab appears only if you are using a wireless network adapter that supports the Wireless Configuration service. If you are unsure whether your wireless network adapter supports the Wireless Configuration service, contact the wireless network adapter manufacturer. For more information about automatic wireless network configuration, see Related Topics.

- If you are unable to connect to an existing wireless network, and the name of the network that you want to connect to appears under the **Preferred networks** list, click the network, and then click **Properties**. In **Wireless Network Properties**, review the settings to ensure that they are correct.

- If the **Preferred networks** list includes both access point (infrastructure) and computer-to-computer (ad hoc) networks, you cannot move a computer-to-computer network to a position in the list that is higher than the position of an access point network. To connect to a computer-to-computer (ad hoc) network when both computer-to-computer and access point (infrastructure) networks are within range of your computer, click **Advanced**, and then click **Computer-to-computer (ad hoc) networks only**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Show a network connection icon in the notification area
Define 802.1X authentication for wireless networks on a client computer
Security information for wireless networks
Configuring wireless network settings on client computers
Wireless networking overview
Understanding 802.1X authentication for wireless networks
Connect to an available wireless network

# Define a wireless network connection on a client computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define a wireless network connection on a client computer**

1. Right-click the wireless network connection icon in the notification area, and then click **View Available Wireless Networks**. For more information, see Notes.

2. Under **Related Tasks**, click **Change the order of preferred networks**, and then on the **Wireless Networks** tab do one of the following:

| To | Do this |
|---|---|
| Add and configure a new wireless network | Click **Add**, and then on the **Association** tab, in **Network name (SSID)**, type a unique name (also known as a service set identifier). |
| Reconfigure the settings for an existing wireless network | Click on the wireless network that you want to reconfigure, and then click **Properties**. |

3. If a network key is required to connect to the wireless network, under **Wireless network key**, do the following:

   - To specify that a network key for Wired Equivalent Privacy (WEP) is used to encrypt the data that is sent over the network, in **Data encryption**, select **WEP**. WEP is selected by default.

   - To specify that a network key is used for authentication to the wireless network, in **Network Authentication**, select **Shared**, **WPA**, or **WPA-PSK**. For more information, consult your network administrator, or see the product documentation for your access point.

   - If the network key is automatically provided for you (for example, if 802.1X is being used for dynamic key distribution), select the **The key is provided for me automatically** check box, and leave **Network key** and **Confirm network key** blank.

   - If the network key is not automatically provided for you, clear the **The key is provided for me automatically** check box, and in **Network key**, type the key, and then type it again in **Confirm network key**. For information about how to specify a network key, see Notes.

   - To specify the location where the network key is stored, in **Key index**, type a number, or click the arrows to select a number.

4. If you are configuring settings to a wireless network that is for a computer-to-computer (ad hoc) network, select the **This is a computer-to-computer (ad hoc) network; wireless access points are not used** check box.

5. To specify that 802.1X authentication is used for the wireless network connection, click the **Authentication** tab, and then configure settings as needed. For more information, see Related Topics.

**Important**

- It is highly recommended that you use 802.1X authentication whenever you connect to an 802.11 wireless network. 802.1X is an IEEE standard that enhances security and deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting. For more information, see Related Topics.

- In Windows XP Service Pack 1 and in the Windows Server 2003 family, 802.1X authentication is available only for access point (infrastructure) networks that require the use of a network key (WEP). WEP provides data confidentiality by

encrypting the data that is sent between wireless clients and wireless access points. For additional information about security for wireless networks, see Related Topics.

- For enhanced security, in Windows XP Service Pack 2 and in the Windows Server 2003 family with Service Pack 1, Wi-Fi Protected Access (WPA) has been included to strengthen WEP. For more information, see Security information for wireless networks.

- If you attempt to connect to a computer-to-computer network or an access point network that does not require the use of a network key, the settings in the **Authentication** tab are not available, and you cannot configure 802.1X authentication for the connection.

- To enhance security and connectivity, do not select the **Network Authentication (Shared mode)** check box. When this check box is selected, shared key authentication is used, rather than open system authentication. Shared key authentication is less secure than open system authentication because it requires the exchange of a secret key that is shared by all wireless access points and clients and therefore is more vulnerable to known-text attacks. In addition, if you select this check box for a wireless network that has multiple wireless access points, you will lose network connectivity when you travel from one wireless access point to a new wireless access point. To determine whether the wireless network you are connecting to has multiple wireless access points, use Wireless Monitor. For information about how to use Wireless Monitor to view details about wireless access points, see Related Topics.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- tok:wirelessicon

- When an error is detected that might limit or prevent a connection to a wireless network, the wireless warning icon is displayed in the notification area.

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- The **Wireless Networks** tab appears only if you are using a wireless network adapter that supports the Wireless Configuration service. If you are unsure whether your wireless network adapter supports the Wireless Configuration service, contact the wireless network adapter manufacturer. For more information about using Windows to configure wireless network settings on your computer, see Related Topics.

- You must type a network key in hexadecimal or ASCII format. If you use hexadecimal format, specify a key that is 10 or 26 characters in length. If you use ASCII format, specify a key that is 5 or 13 characters in length.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a network connection icon in the notification area
Define 802.1X authentication for wireless networks on a client computer
Configuring wireless network settings on client computers
Understanding 802.1X authentication for wireless networks
Security information for wireless networks
View details about wireless network access points
Connect to an available wireless network
Use Windows to configure wireless network settings on a client computer

# Define 802.1X authentication for wireless networks on a client computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To define 802.1X authentication for wireless networks on a client computer**

1. Right-click the wireless network connection icon in the notification area, and then click **View Available Wireless Networks**. For more information, see Notes.

2. Under **Related Tasks**, click **Change the order of preferred networks**.

3. On the **Wireless Networks** tab, under **Preferred networks**, click the wireless network connection for which you want to configure 802.1X authentication, and then click **Properties**.

4. On the **Authentication** tab, do one of the following:

   - To enable IEEE 802.1X authentication for this connection, select the **Enable IEEE 802.1x authentication for this network** check box. This check box is selected by default.

   - To disable IEEE 802.1X authentication for this connection, clear the **Enable IEEE 802.1x authentication for this network** check box.

5. In **EAP type**, click the Extensible Authentication Protocol (EAP) type to be used with this connection.

6. If you select **Smart Card or other Certificate** in **EAP type**, click **Properties** and, in **Smart Card or other Certificate Properties**, do the following:

   - To use the certificate that resides on your smart card for authentication, click **Use my smart card**.

   - To use the certificate that resides in the certificate store on your computer for authentication, click **Use a certificate on this computer**, and then specify whether to use simple certificate selection.

   - To verify that the server certificate presented to your computer is still valid, select the **Validate server certificate** check box, specify the server or servers to which your computer will automatically connect, and then specify the trusted root certification authorities.

   - To view detailed information about the selected root certification authority, click **View Certificate**.

   - To use a different user name when the user name in the smart card or certificate is not the same as the user name in the domain to which you are logging on, select the **Use a different user name for the connection** check box.

7. If you select **Protected EAP (PEAP)** in **EAP type**, click **Properties**, and then do the following:

   - To verify that the server certificate presented to your computer is still valid, select the **Validate server certificate** check box, specify the server or servers to which your computer will automatically connect, and then specify the trusted root certification authorities.

   - In **Select Authentication Method**, click the authentication method that you want to use within PEAP, and then click **Configure**.

If you select **Secured password (EAP-MSCHAP v2)**, then, in **EAP MSCHAP v2 Properties**, specify whether to use the user name and password (and domain, if applicable) that you type in the Windows logon screen for authentication, click **OK**, and then click **OK** again.

If you select **Smart Card or other certificate**, then, in **Smart Card or other Certificate Properties**, configure the settings as needed by following the instructions in step 6, click **OK**, and then click **OK** again.

8. On the **Authentication** tab, do the following:

   - To specify that the computer attempt authentication to the network if a user is not logged on, select the **Authenticate as computer when computer information is available** check box. This check box is selected by default.

   - To specify that the computer attempt authentication to the network if user information or computer information is not available, select the **Authenticate as guest when user or computer information is unavailable** check box.

**Important**

- It is highly recommended that you use 802.1X authentication whenever you connect to an 802.11 wireless network. 802.1X is an IEEE standard that enhances security and deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting. For more information, see Related Topics.

- For enhanced security, in Windows XP Service Pack 1 and in the Windows Server 2003 family, 802.1X authentication is available only for access point (infrastructure) networks that require the use of a network key (WEP). WEP provides data confidentiality by encrypting the data that is sent between wireless clients and wireless access points. For additional information about security for wireless networks, see Related Topics.

- If you attempt to connect to a computer-to-computer network or an access point network that does not require the use of a network key, the settings in the **Authentication** tab are not available, and you cannot configure 802.1X authentication for the connection.

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- tok:wirelessicon

- When an error is detected that might limit or prevent a connection to a wireless network, the wireless warning icon is displayed in the notification area.

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- To define 802.1X authentication, you must select an existing wireless network connection, or you must add a new wireless network connection. For information about how to add a new wireless network connection, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a network connection icon in the notification area
Add, edit, or remove wireless network connections on a client computer

Configuring wireless network settings on client computers
Connect to an available wireless network
Understanding 802.1X authentication for wireless networks
Security information for wireless networks
Wireless networking overview

# Connect to an available wireless network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to an available wireless network**

1. Right-click the wireless network connection icon in the notification area, and then click **View Available Wireless Networks**. For more information, see Notes.

2. Click on the wireless network that you want to connect to, and then click **Connect**.

3. To configure additional wireless network settings, or if you are having difficulty making a connection to the wireless network that you selected, under **Related Tasks**, click **Change the order of preferred networks**, and then configure the settings in the **Wireless Networks** tab. For more information, see Notes.

## Important

- It is highly recommended that you use 802.1X authentication whenever you connect to an 802.11 wireless network. 802.1X is an IEEE standard that enhances security and deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting. For more information, see Related Topics.

- For enhanced security, in Windows XP Service Pack 1 and in the Windows Server 2003 family, 802.1X authentication is available only for access point (infrastructure) networks that require the use of a network key (WEP). WEP provides data confidentiality by encrypting the data that is sent between wireless clients and wireless access points. For additional information about security for wireless networks, see Related Topics.

- If you attempt to connect to a computer-to-computer network or an access point network that does not require the use of a network key, the settings in the **Authentication** tab are not available, and you cannot configure 802.1X authentication for the connection.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- tok:wirelessicon

- When an error is detected that might limit or prevent a connection to a wireless network, the wireless warning icon is displayed in the notification area.

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- On the **Wireless Networks** tab, you can view and modify additional wireless configuration settings. However, this tab appears only if you are using a wireless network adapter that supports the Wireless Configuration service. For more information, see Related Topics.

- You must type a network key in hexadecimal or ASCII format. If you use hexadecimal format, specify a key that is 10 or 26 characters in length. If you use ASCII format, specify a key that is 5 or 13 characters in length.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Show a network connection icon in the notification area
Add, edit, or remove wireless network connections on a client computer
Security information for wireless networks
Configuring wireless network settings on client computers
Define 802.1X authentication for wireless networks on a client computer
Understanding 802.1X authentication for wireless networks

# Use Resultant Set of Policy (RSoP) to View Wireless Network Policy Assignments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Resultant Set of Policy (RSoP) to view wireless network policy assignments**

- Use RSoP to view wireless network policy assignments for a computer
- Use RSoP to view wireless network policy assignments for members of a Group Policy container

# Use RSoP to view wireless network policy assignments for a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use RSoP to view wireless network policy assignments for a computer**

1. Open Resultant Set of Policy.

   An RSoP query runs in logging mode. For more information, see Related Topics.

2. In the console tree, click **Wireless Network (IEEE 802.11) Policies**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Wireless Network (IEEE 802.11) Policies

   The list of all wireless network policies is displayed in the details pane.

3. To view detailed policy settings for the wireless network policy that is being applied, in the details pane, right-click the policy with the highest precedence, and then click **Properties**.

## Notes

- To open RSoP, click **Start**, click **Run**, type **rsop.msc**, and then click **OK**.
- To run an RSoP query in logging mode, see Related Topics.
- The RSoP console displays detailed wireless network policy settings for only the policy that is being applied. The detailed policy settings include the list of preferred wireless networks that clients can connect to and, for each preferred wireless network, the wireless network key settings, whether 802.1X authentication is enabled, and, if so, the authentication settings.
- The details pane in RSoP provides two ways to indicate which wireless network policy is being applied:
  - The lower the number in the **Precedence** column, the higher the precedence of the wireless network policy. The wireless network policy with a precedence number of 1 is the policy that is being applied.
  - In the **Name** column, the folder icon that precedes the name of the wireless network policy that is being applied contains a green dot.
- To hide or display columns in the RSoP console, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run RSoP
Access RSoP data for an existing computer and user (logging mode)

RSoP logging mode
Hide or display columns in an MMC console
Using Resultant Set of Policy to view wireless network policy assignments

# Use RSoP to view wireless network policy assignments for members of a Group Policy container

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use RSoP to view wireless network policy assignments for members of a Group Policy container**

1. On a domain controller, open Resultant Set of Policy (RSoP) as a Microsoft Management Console (MMC). For more information, see Related Topics.

2. On the **Action** menu, click **Generate RSoP Data**.

3. On the **Welcome to the Resultant Set of Policy Wizard** page, click **Next**.

4. On the **Mode Selection** page, click **Planning mode**, and then follow the instructions in the wizard to run a planning mode query. For more information, see Related Topics.

5. In the console tree, click **Wireless Network (IEEE 802.11) Policies**.

   **Where?**

   - Computer Configuration/Windows Settings/Security Settings/Wireless Network (IEEE 802.11) Policies

   The list of all wireless network policies is displayed in the details pane.

6. To view detailed policy settings for the wireless network policy that is being applied, in the details pane, right-click the policy with the highest precedence, and then click **Properties**.

**Notes**

- The RSoP console only displays detailed wireless network policy settings for only the policy that is being applied. The detailed policy settings include the list of preferred wireless networks that clients can connect to, and for each preferred wireless network, the wireless network key settings, whether 802.1X authentication is enabled, and if so, the authentication settings.

- The details pane in RSoP provides two ways to indicate which wireless network policy is being applied:

  - The lower the number in the **Precedence** column, the higher the precedence of the wireless network policy. The wireless network policy with a precedence number of 1 is the policy that is being applied.

  - In the **Name** column, the folder icon that precedes the name of the wireless network policy that is being applied contains a green dot.

- You can run an RSoP planning mode query only on a domain controller (when you run a planning mode query, you must specify the name of the domain controller). However, you can specify any wireless network client as the target for the query, provided you have the appropriate permissions to do so.

- To hide or display columns in the RSoP console, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open RSoP as an MMC snap-in
Simulate RSoP data for a computer and user (planning mode)
Hide or display columns in an MMC console
Using Resultant Set of Policy to view wireless network policy assignments

# Monitor Wireless Network Activity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor wireless network activity**

- Add the Wireless Monitor snap-in
- Enable or disable Wireless Client Information logging
- Delete Wireless Client Information logs
- View details about wireless network access points
- View details about wireless network clients

# Add the Wireless Monitor snap-in

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add the Wireless Monitor snap-in**

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.
3. In the **Add Standalone Snap-in** dialog box, click **Wireless Monitor**, and then click **Add**.
4. Click **Close**, and then click **OK**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To save the console settings, on the **File** menu, click **Save**.
- To open a saved console file, open MMC. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open a saved MMC console for local and remote computers
Open MMC

# Enable or disable Wireless Client Information logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable Wireless Client Information logging**

1. Create a console containing Wireless Monitor. Or, open a saved console file containing Wireless Monitor.
2. Double-click **Wireless Client Information**.

   **Where?**

   - Wireless Monitor/*ServerName*/Wireless Client Information

3. Right-click **Wireless Client Information**, and then do one of the following:

- To disable logging for wireless network clients, click **Disable Logging**.
- To enable logging for wireless network clients, click **Enable Logging**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- For information about adding the Wireless Monitor snap-in, see Related Topics.
- For information about the items displayed in Wireless Monitor, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Open a saved MMC console for local and remote computers
Save an MMC console file
Logging and viewing wireless network activity
Add the Wireless Monitor snap-in

# Delete Wireless Client Information logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete Wireless Client Information logs**

1. Create a console containing Wireless Monitor. Or, open a saved console file containing Wireless Monitor.
2. Double-click **Wireless Client Information**.

   **Where?**

   - Wireless Monitor/*ServerName*/Wireless Client Information

3. Right-click **Wireless Client Information**, and then click **Delete Logs**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- For information about adding the Wireless Monitor snap-in, see Related Topics.
- For information about the items displayed in Wireless Monitor, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Open a saved MMC console for local and remote computers
Save an MMC console file
Logging and viewing wireless network activity
Add the Wireless Monitor snap-in

# View details about wireless network access points

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view details about wireless network access points**

1. Create a console containing Wireless Monitor. Or, open a saved console file containing Wireless Monitor.
2. Double-click **Access Point Information**.

   **Where?**

   - Wireless Monitor/*ServerName*/Access Point Information

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- For information about adding the Wireless Monitor snap-in, see Related Topics.
- For information about the items displayed in Wireless Monitor, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Open a saved MMC console for local and remote computers
Save an MMC console file
Logging and viewing wireless network activity
Add the Wireless Monitor snap-in

# View details about wireless network clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view details about wireless network clients**

1. Create a console containing Wireless Monitor. Or, open a saved console file containing Wireless Monitor.
2. Double-click **Wireless Client Information**.

   **Where?**

   - Wireless Monitor/*ServerName*/Wireless Client Information

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- For information about adding the Wireless Monitor snap-in, see Related Topics.
- For information about the items displayed in Wireless Monitor, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Open a saved MMC console for local and remote computers
Save an MMC console file
Logging and viewing wireless network activity
Add the Wireless Monitor snap-in

# Manage Infrared Communication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage infrared communication**

- Set up the Computer for Infrared Communication
- Configure Infrared Communication
- Use Infrared for Dial-up or Direct Network Connections
- Send Files over an Infrared Link
- Print to an infrared printer
- Transfer images from a digital camera to your computer

# Set up the Computer for Infrared Communication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up the computer for infrared communication**

- Verify infrared support on a computer
- Add a new infrared device to a serial port
- View infrared device properties
- Establish an infrared link

# Verify infrared support on a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To verify infrared support on a computer**

1. Verify that your computer has an infrared transceiver (the transceiver will appear as a small, dark red window).

   If your computer does not have an infrared transceiver, see your manufacturer's documentation to verify whether IrDA functionality is supported.

2. Open Device Manager.

3. Double-click **Infrared Devices**.

   If you have an infrared transceiver, but **Infrared Devices** does not appear in Device Manager, you do not have an infrared device installed.

   If no infrared devices are listed, do one or both of the following:

   - Enter BIOS setup to verify whether the infrared device is enabled in BIOS, in IrDA, or Fast IrDA (FIR) mode. For information about how to enter BIOS setup, see your computer manufacturer's documentation. If the infrared device is disabled in BIOS, you might be able to use your computer's BIOS setup to enable it.

     **Caution** Using a BIOS setup utility incorrectly can cause your computer to operate incorrectly. Microsoft cannot guarantee that problems resulting from the incorrect use of the BIOS setup can be solved. Use this utility at your own risk. Before you make any changes, make a note of the original settings so that you can reinstate them if necessary.

   - See your computer manufacturer's documentation to determine whether the BIOS needs to be updated.

4. Verify that at least one infrared device is listed and that it is enabled.

   If infrared devices are listed but they are not enabled, verify that the infrared devices are installed correctly on your computer. For more information see Related Topics.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open Device Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Computer Management**, and then click **Device Manager**.

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

- Infrared transceivers are now installed in nearly all new portable computers. If your computer does not have an infrared transceiver, but it supports IrDA, and IrDA is enabled in BIOS, you can install an external infrared transceiver.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Add a new infrared device to a serial port

Enable a device

# Add a new infrared device to a serial port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a new infrared device to a serial port**

1. Connect the infrared device to a serial port on your computer, and note the serial port number (typically COM1 or COM2).
2. In Control Panel, open Add Hardware.
3. In **Welcome to the Add Hardware Wizard**, click **Next**.
4. In **Is the hardware connected?**, click **Yes, I have already connected the hardware**, and then click **Next**.
5. In **Installed hardware**, click **Add a new hardware device**, and then click **Next**.
6. In **The wizard can help you install other hardware**, click **Install the hardware that I manually select from a list**, and then click **Next**.
7. In **Common hardware types**, click **Infrared devices**, and then click **Next**.
8. In **Select Infrared Device**, do one of the following:
   - If you do not have an installation disk for the infrared device, under **Manufacturer**, click the manufacturer, under **Infrared Device**, click the infrared device, and then click **Next**.
   - If you have an installation disk for a compatible driver, click **Have Disk**, insert the disk, click **Next**, and then follow any additional instructions to install the device.
9. In **The wizard is ready to install your hardware**, click **Next**.
10. In **Infrared Device Installation**, select the serial port to which your infrared device is attached, and then click **Next**.

**Important**

- For enhanced security, performance, and stability, it is recommended that you use only signed device drivers. For more information, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- This procedure applies only if an infrared device was not installed automatically when you turned on your computer. Typically, USB, PCMCIA, and other Plug and Play compliant infrared devices are automatically detected and installed. These devices do not need additional configuration to be installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing infrared devices
Set up the Computer for Infrared Communication
Signed Drivers

# View infrared device properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view infrared device properties**

1. Open Wireless Link in Control Panel.
2. On the **Hardware** tab, click the infrared device for which you want to view properties, and then click **Properties**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- The Wireless Link icon appears in Control Panel only if an infrared device is already installed on your computer.
- If a device that is connected to your computer is not listed, you need to add it. For more information, see Related Topics.
- The infrared device **Properties** dialog box shows the device status, driver files, and power management settings for an infrared device.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a new infrared device to a serial port
Verify infrared support on a computer

# Establish an infrared link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To establish an infrared link**

1. Verify that the devices that you want to communicate over an infrared link have infrared functionality enabled and work correctly.

   For information about verifying infrared functionality on your computer, see Related Topics. For information about verifying infrared functionality on other devices, see the device manufacturer's documentation.

   Align your devices so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other.

2. When the devices are correctly aligned, the icon appears on the taskbar.

### Important

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Related Topics.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

- The infrared transceiver is the small, dark red window on your portable computer, printer, digital camera, adapter, or other device.

- You can choose whether to hide or display the Wireless Link taskbar icon when two infrared devices are aligned. By default, the Wireless Link taskbar icon is displayed. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Verify infrared support on a computer
Security information for infrared communication
Transferring files
Hide or display the Wireless Link taskbar icon
Add a new infrared device to a serial port

# Configure Infrared Communication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure infrared communication**

- Enable or prevent receipt of files sent over an infrared link
- Change the default location for receipt of files sent over an infrared link
- Hide or display the Wireless Link taskbar icon
- Enable or disable notification of files received over an infrared link
- Enable or disable sound notification when an infrared device is nearby
- Change settings to transfer images from a digital camera

# Enable or prevent receipt of files sent over an infrared link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or prevent receipt of files sent over an infrared link**

1. Open Wireless Link in Control Panel.
2. On the **Infrared** tab, do one of the following:
   - To enable your computer to receive files from others, select the **Allow others to send files to your computer using infrared communications** check box. This check box is selected by default.
   - To prevent your computer from receiving files from others, clear the **Allow others to send files to your computer using infrared communications** check box.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- The Wireless Link icon appears in Control Panel only if an infrared device is already installed on your computer.
- If you clear the **Allow others to send files to your computer using infrared communications** check box, you can still send files to other computers by using Wireless Link.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable notification of files received over an infrared link
Enable or disable sound notification when an infrared device is nearby

# Change the default location for receipt of files sent over an infrared link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the default location for receipt of files sent over an infrared link**

1. Open Wireless Link in Control Panel.
2. On the **Infrared** tab, under **Default location for received files**, click **Browse**.

   By default, files that are sent over an infrared link are copied to the desktop of the recipient computer.
3. Enter a new location, and then click **OK**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- The Wireless Link icon appears in Control Panel only if an infrared device is already installed on your computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or prevent receipt of files sent over an infrared link
Enable or disable notification of files received over an infrared link
Change the default location for receipt of files sent over an infrared link

# Hide or display the Wireless Link taskbar icon

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To hide or display the Wireless Link taskbar icon**

1. Open Wireless Link in Control Panel.
2. On the **Infrared** tab, do one of the following:
   - To hide the **Wireless Link** taskbar icon, clear the **Display an icon on the taskbar indicating infrared activity** check box. This check box is selected by default.
   - To display the **Wireless Link** taskbar icon, select the **Display an icon on the taskbar indicating infrared activity** check box.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- The Wireless Link icon appears in Control Panel only if an infrared device is already installed on your computer.

  When another infrared transceiver is in range, this option causes the Infrared status icon to appear on the taskbar. The icon changes to whenever the infrared connection is active. To display the current status of the connection, move the mouse pointer over the icon.

  To open the properties for Wireless Link, right-click the icon, and then click **Properties**.

  To transfer files, right-click the icon, and then click **Transfer Files** to open the **Wireless Link** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure Infrared Communication

# Enable or disable notification of files received over an infrared link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable notification of files received over an infrared link**

1. Open Wireless Link in Control Panel.
2. On the **Infrared** tab, do one of the following:
   - If you do not want to be notified when you receive files over an infrared link, clear the **Notify me when receiving files** check box. This check box is selected by default.
   - To be notified when you receive files sent over an infrared link, select the **Notify me when receiving files** check box.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- The Wireless Link icon appears in Control Panel only if an infrared device is already installed on your computer.
- If the **Notify me when receiving files** check box is selected, the Wireless Link status dialog box appears when files are being received. The status dialog box indicates that files are being received and shows the name of the computer that is sending the files.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or prevent receipt of files sent over an infrared link
Enable or disable notification of files received over an infrared link
Change the default location for receipt of files sent over an infrared link

# Enable or disable sound notification when an infrared device is nearby

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable sound notification when an infrared device is nearby**

1. Open Wireless Link in Control Panel.
2. On the **Infrared** tab, do one of the following:
   - To enable sound notification when an infrared device is nearby, select the **Play sound when Infrared device is nearby** check box. This check box is selected by default.
   - To disable sound notification when an infrared device is nearby, clear the **Play sound when Infrared device is nearby** check box.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- The Wireless Link icon appears in Control Panel only if an infrared device is already installed on your computer.
- If the **Play sound when infrared device is nearby** check box is selected, but you do not hear the notification sound, make sure that your infrared devices are correctly aligned and that an infrared link is established. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable notification of files received over an infrared link
Enable or prevent receipt of files sent over an infrared link
Hide or display the Wireless Link taskbar icon
Establish an infrared link

# Change settings to transfer images from a digital camera

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change settings to transfer images from a digital camera**

1. Open Wireless Link in Control Panel.
2. On the **Image Transfer** tab, select the **Use Wireless Link to transfer images from a digital camera to your computer** check box.
3. In **Transfer images to a subfolder of**, specify the default location for images transferred by Wireless Link.
4. To automatically open Windows Explorer to the default location after receiving files, select the **Explore location after receiving pictures** check box.

## Caution

- To enhance security, do not enable infrared image transfer (IrTran-P), that is, do not select the **Use Wireless Link to transfer images from a digital camera to your computer** check box. If you select this check box, all incoming files sent over IrTran-P are automatically accepted. Because incoming files might contain harmful programs, ensure that the files originate from a trustworthy source. Do not open files if you cannot verify the source, do not recognize the file format, or are unsure of the content. Instead, delete the files immediately. For more information about infrared security considerations, see Related Topics.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- The Wireless Link icon appears in Control Panel only if an infrared device is already installed on your computer.
- If you are using synchronization software, note that certain software will automatically disable infrared image transfer (IrTran-P). As a result, you will not be able to use Wireless Link to transfer images from a digital camera to your computer. In such cases, you might have to remove the synchronization software before you can use Wireless Link to transfer digital camera images. To determine whether your synchronization software disables infrared image transfer, see your software manufacturer's product documentation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Transfer images from a digital camera to your computer
Transferring digital camera images to the computer
Security information for infrared communication

# Use Infrared for Dial-up or Direct Network Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use infrared for dial-up or direct network connections**

- Make a dial-up connection using an infrared-enabled cellular phone
- Establish a direct infrared network connection between computers
- Use an infrared network connection

# Make a dial-up connection using an infrared-enabled cellular phone

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a dial-up connection using an infrared-enabled cellular phone**

1. Make sure that infrared is activated on the cellular phone. To check compatibility, see your cellular phone manufacturer's documentation.

2. Establish an infrared link between the cellular phone and the computer by aligning them so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other. For more information, see Related Topics.

3. If you are establishing an infrared link between the cellular phone and the computer for the first time, the Found New Hardware balloon appears when the cellular phone is detected. Do any of the following before proceeding to step 4:

   - If the device driver for the cellular phone already exists on your computer, the driver is installed automatically, and you are notified that the cellular phone is ready to use.

   - If the device driver for the cellular phone does not already exist on your computer, the Found New Hardware Wizard appears. Do one of the following:

     | |
     |---|
     | If you have a setup or installation disk for the driver files, select **Install from a list or specific location (advanced)**, and then follow the instructions in the wizard. |
     | If you do not have a setup or installation disk, select **Install software automatically (recommended)**, and then follow the instructions in the wizard. To complete this task, you need administrative credentials. For more information, see Notes. |

   **Important**

   - If you are installing an unsigned device driver for the cellular phone, a message appears to notify you that the driver is not signed. For enhanced security, performance, and stability, it is recommended that you use only signed device drivers. For more information, see Related Topics.

4. Open Network Connections.

5. Double-click **New Connection Wizard**, and then click **Next**.

6. In **Network Connection Type**, do one of the following:

   - To connect to the Internet, click **Connect to the Internet**.

   - To connect to a business network, click **Connect to the network at my workplace**.

7. Click **Next**, and then do one of the following:

   - If you are connecting to the Internet, in **Internet Connection**, click **Connect using a dial-up modem**, click **Next**, and then follow the instructions in the remaining wizard pages.

- If you are connecting to a business network, in **Network Connection**, click **Dial-up connection**, click **Next**, and then follow the instructions in the remaining wizard pages.

8. To complete the connection, in the **Connect *connection type*** dialog box, type your user name and password, and then click **Dial**.

### Important

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Related Topics.

### Notes

- To install new driver files on a computer, you must have be a member of the Administrators group on the local computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Establish an infrared link
Security information for infrared communication
Infrared implementation
Signed Drivers

# Establish a direct infrared network connection between computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To establish a direct infrared network connection between computers**

1. Open Network Connections.
2. Double-click **New Connection Wizard**, and then click **Next**.
3. In **Network Connection Type**, click **Advanced connection**, and then click **Next**.
4. In **Advanced Connection Options**, click **Connect directly to another computer**, and then click **Next.**
5. In **Host or Guest**, specify whether this computer is the host or the guest by doing one of the following:
   - If this computer has information that will be accessed by another computer, click **Host**.
   - If this computer is being used to access information on another computer, click **Guest**.
6. Click **Next**, and then do one of the following:
   - If you specified this computer as a host, in **Devices for this connection**, select **Infrared Port**, click **Next**, and then follow the instructions in the remaining wizard pages.
   - If you specified this computer as a guest, in **Connection Name**, type the name of the computer that you want to access, and then click **Next**. In **Select a device**, select **Infrared Port**, click **Next**, and then follow the instructions in the remaining wizard pages.

     To complete the connection, in the **Connect *connection type*** dialog box, type your user name and password, and then click **Connect**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- To examine or change properties for this connection, right-click its icon in Network Connections.
- After you add an infrared device, you may have to restart the computer before you can select **Infrared Port** in step six.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Establishing infrared network connections between computers

Use an infrared network connection

Add a new infrared device to a serial port

# Use an infrared network connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use an infrared network connection**

1. Establish an infrared link by aligning your computers so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other. For more information, see Related Topics.
2. Open Network Connections.
3. Double-click the connection that you want to use to access the network.
4. Verify that the connection information is correct, and then click **OK**.
5. On the **File** menu, click **Connect**.

### Important

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Related Topics.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- When you make a network connection over infrared, the infrared link takes the place of a cable. You do not need to use a serial or parallel cable.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Establish an infrared link
Security information for infrared communication
Establish a direct infrared network connection between computers
Make a dial-up connection using an infrared-enabled cellular phone
Network Connections

# Send Files over an Infrared Link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Send files over an infrared link**

- Send files with the Wireless Link dialog box
- Send files with Send To
- Send files using the command line
- Drag files to an infrared link

# Send files with the Wireless Link dialog box

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send files with the Wireless Link dialog box**

1. Establish an infrared link by aligning your computers so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other. For more information, see Related Topics.

2. Click the IrDA status icon on the taskbar.

3. In the **Wireless Link** dialog box, specify the files you want to send, and then click **Send**.

### Important

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Related Topics.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

  While another infrared transceiver is in range, you can also double-click the Wireless Link icon on the desktop to open the **Wireless Link** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Establish an infrared link
Security information for infrared communication
Transferring files

# Send files with Send To

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send files with Send To**

1. Establish an infrared link by aligning your computers so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other. For more information, see Related Topics.

2. On the desktop, in Windows Explorer or in **My Computer**, select the files you want to send.

3. Right-click the selected files, point to **Send To**, and then click **Nearby Computer**.

### Important

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Related Topics.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

- If **Nearby Computer** does not appear on the **Send To** menu (as described in step 3 of this procedure), make sure that your infrared devices are correctly aligned and that an infrared link is established. The **Nearby Computer** menu command only appears when an infrared link is active.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Establish an infrared link
Security information for infrared communication
Transferring files

# Send files using the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send files using the command line**

1. Establish an infrared link by aligning your computers so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other. For more information, see Related Topics.

2. Open Command Prompt.

3. At the command prompt, type:

   **irftp**[*Drive:*] [[*Path*] *FileName*] [**/h**]

| Value | Description |
|---|---|
| *Drive:* | Specifies the drive that contains the files that you want to send over an infrared link. |
| [*Path*] *FileName* | Specifies the location and name of the file or set of files that you want to send over an infrared link. |
| **/h** | Specifies hidden mode. When hidden mode is used, the **Wireless Link** dialog box does not display when the files are sent. |

### Important

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Related Topics.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

- If you do not specify a drive, a path, and a name for the file or set of files that you want to send, the **Wireless Link** dialog box opens, so that you can select the files without using the command line.

- For information about other command-line utilities or manageability, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Establish an infrared link

Security information for infrared communication

Transferring files

Command-line reference A-Z

**Other Resources**

Irftp

# Drag files to an infrared link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To drag files to an infrared link**

1. Establish an infrared link by aligning your computers so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other. For more information, see Related Topics.

2. In Windows Explorer, in **My Computer**, or on the desktop, select one or more files.

   Drag the selected files onto the desktop icon, which represents the wireless connection.

### Important

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Related Topics.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Establish an infrared link
Security information for infrared communication
Transferring files

# Print to an infrared printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print to an infrared printer**

1. Verify that the infrared printer you want to use is correctly installed.

2. Establish an infrared link between the computer and the printer by aligning them so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other. For more information, see Related Topics.

3. Open the document that you want to print, and then proceed with printing.

### Important

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Related Topics.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a printer attached to your computer
Security information for infrared communication
Establish an infrared link
Print a document
Printing to an infrared printer

# Transfer images from a digital camera to your computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To transfer images from a digital camera to your computer**

1. Open Wireless Link in Control Panel.
2. On the **Image Transfer** tab, select the **Use Wireless Link to transfer images from a digital camera to your computer** check box.
3. Establish an infrared link by aligning your computers so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other. For more information about establishing infrared links, see Related Topics.
4. Refer to the camera manufacturer's instructions to select and transfer images.

**Caution**

- To enhance security, do not enable infrared image transfer (IrTran-P), that is, do not select the **Use Wireless Link to transfer images from a digital camera to your computer** check box. If you select this check box, all incoming files sent over IrTran-P are automatically accepted. Because incoming files might contain harmful programs, ensure that the files originate from a trustworthy source. Do not open files if you cannot verify the source, do not recognize the file format, or are unsure of the content. Instead, delete the files immediately. For more information about infrared security considerations, see Related Topics.

**Important**

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Related Topics.

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.
- The Wireless Link icon appears in Control Panel only if an infrared device is already installed on your computer.
- The Wireless Link icon will automatically create a folder for the transferred images, accept the images, and open the folder when the image transfer is completed.
- Wireless Link cannot initiate the transfer of images. The camera must initiate the transfer of images.
- To use Wireless Link to transfer images, the camera must support infrared image transfer (IrTran-P). If it does not support this protocol, you need to use a program provided by the camera manufacturer or by another party. If you use

another program, make sure you clear the **Use Wireless Link to transfer images from a digital camera to your computer** check box.

- If you are using synchronization software, note that certain software will automatically disable IrTran-P. As a result, you will not be able to use Wireless Link to transfer images from a digital camera to your computer. In such cases, you might have to remove the synchronization software before you can use Wireless Link to transfer digital camera images. To determine whether your synchronization software disables infrared image transfer, see your software manufacturer's product documentation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Establish an infrared link
Change settings to transfer images from a digital camera
Security information for infrared communication

# Wireless Networking Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about wireless networking:

- Wireless networking overview
- Wireless Link overview
- Understanding Wireless Networking
- Using Wireless Networking
- Wireless Networking Resources

# Wireless networking overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless networking overview**
Wireless networking technologies range from global voice and data networks, which allow users to establish wireless connections across long distances, to infrared light and radio frequency technologies that are optimized for short-range wireless connections. Devices commonly used for wireless networking include portable computers, desktop computers, hand-held computers, personal digital assistants (PDAs), cellular phones, pen-based computers, and pagers. Wireless technologies serve many practical purposes. For example, mobile users can use their cellular phone to access e-mail. Travelers with portable computers can connect to the Internet through base stations installed in airports, railway stations, and other public locations. At home, users can connect devices on their desktop to synchronize data and transfer files.

**Defining standards**
To lower costs, ensure interoperability, and promote the widespread adoption of wireless technologies, organizations such as the Institute of Electrical and Electronics Engineers (IEEE), Internet Engineering Task Force (IETF), Wireless Ethernet Compatibility Alliance (WECA), and the International Telecommunication Union (ITU) are participating in several major standardization efforts. For example, IEEE working groups are defining how information is transferred from one device to another (whether radio waves or infrared light is used, for example) and how and when a transmission medium should be used for communications. In developing wireless networking standards, organizations such as the IEEE address power management, bandwidth, security, and issues that are unique to wireless networking.

**Wireless network types**
As with wired networks, wireless networks can be classified into different types based on the distances over which data can be transmitted.

**Wireless wide area networks (WWANs)**
WWAN technologies enable users to establish wireless connections over remote public or private networks. These connections can be maintained over large geographical areas, such as cities or countries, through the use of multiple antenna sites or satellite systems maintained by wireless service providers. Current WWAN technologies are known as second-generation (2G) systems. Key 2G systems include Global System for Mobile Communications (GSM), Cellular Digital Packet Data (CDPD), and Code Division Multiple Access (CDMA). Efforts are under way to transition from 2G networks, some of which have limited roaming capabilities and are incompatible with each other, to third-generation (3G) technologies that would follow a global standard and provide worldwide roaming capabilities. The ITU is actively promoting the development of a global standard for 3G.

**Wireless metropolitan area networks (WMANs)**
WMAN technologies enable users to establish wireless connections between multiple locations within a metropolitan area (for example, between multiple office buildings in a city or on a university campus), without the high cost of laying fiber or copper cabling and leasing lines. In addition, WMANs can serve as backups for wired networks, should the primary leased lines for wired networks become unavailable. WMANs use either radio waves or infrared light to transmit data. Broadband wireless access networks, which provide users with high-speed access to the Internet, are in increasing demand. Although different technologies, such as the multichannel multipoint distribution service (MMDS) and the local multipoint distribution services (LMDS), are being used, the IEEE 802.16 working group for broadband wireless access standards is still developing specifications to standardize development of these technologies.

**Wireless local area networks (WLANs)**
WLAN technologies enable users to establish wireless connections within a local area (for example, within a corporate or campus building, or in a public space, such as an airport). WLANs can be used in temporary offices or other spaces where the installation of extensive cabling would be prohibitive, or to supplement an existing LAN so that users can work at different locations within a building at different times. WLANs can operate in two different ways. In infrastructure WLANs, wireless stations (devices with radio network cards or external modems) connect to wireless access points that function as bridges between the stations and the existing network backbone. In peer-to-peer (ad hoc) WLANS, several users within a limited area, such as a conference room, can form a temporary network without using access points, if they do not require access to network resources.

In 1997, IEEE approved the 802.11 standard for WLANs, which specifies a data transfer rate of 1 to 2 megabits per second (Mbps). Under 802.11b, which is emerging as the new dominant standard, data is transferred at a maximum rate of 11 Mbps over a 2.4 gigahertz (GHz) frequency band. Another newer standard is 802.11a, which specifies data transfer at a maximum rate of 54 Mbps over a 5 GHz frequency band.

**Wireless personal area networks (WPANs)**

WPAN technologies enable users to establish ad hoc, wireless communications for devices (such as PDAs, cellular phones, or laptops) that are used within a personal operating space (POS). A POS is the space surrounding a person, up to a distance of 10 meters. Currently, the two key WPAN technologies are Bluetooth and infrared light. Bluetooth is a cable replacement technology that uses radio waves to transmit data to a distance of up to 30 feet. Bluetooth data can be transferred through walls, pockets, and briefcases. Technology development for Bluetooth is driven by the Bluetooth Special Interest Group (SIG), which published the Bluetooth version 1.0 specification in 1999. Alternatively, to connect devices at a very close range (1 meter or less), users can create infrared links.

To standardize the development of WPAN technologies, IEEE has established the 802.15 working group for WPANs. This working group is developing a WPAN standard, based on the Bluetooth version 1.0 specification. Key goals for this draft standard are low complexity, low power consumption, interoperability, and coexistence with 802.11 networks.

# Wireless Link overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Wireless Link overview**
After an infrared device is installed, the Wireless Link icon appears in Control Panel. You can use the Wireless Link icon to configure file transfer and image transfer options and then to send files or transfer images quickly and easily over an infrared link.

When another infrared transceiver comes in range, the Wireless Link icon appears on the desktop and on the taskbar. You can then:

- Send files by specifying a location and one or more files in the **Wireless Link** dialog box. For more information, see Send files with the Wireless Link dialog box.
- Send files by dragging file icons to the Wireless Link icon on the desktop and dropping them there. For more information, see Drag files to an infrared link.
- Transfer images from a digital camera that supports the infrared image transfer (IrTran-P) protocol to your computer. For more information, see Transfer images from a digital camera to your computer.

**Note**

- In the Microsoft® Windows Server 2003 family, only Windows Server 2003, Standard Edition supports infrared networking.

# Understanding Wireless Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding wireless networking**
This section covers:

- Understanding 802.1X authentication for wireless networks
- Understanding wireless provisioning services
- Understanding Infrared Networking

# Understanding 802.1X authentication for wireless networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding 802.1X authentication for wireless networks**
802.1X is an IEEE standard for authenticated network access to wired Ethernet networks and wireless 802.11 networks. IEEE 802.1X enhances security and deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting.

**EAP, EAP-TLS, EAP-MS-CHAP v2, and PEAP authentication**
The support that 802.1X provides for Extensible Authentication Protocol (EAP) types allows you to choose from several different authentication methods for wireless clients and servers.

**EAP**
802.1X uses EAP for message exchange during the authentication process. With EAP, an arbitrary authentication method, such as certificates, smart cards, or credentials, is used. EAP allows for an open-ended conversation between an EAP client (such as a wireless computer) and an EAP server (such as an Internet Authentication Service (IAS) server). The conversation consists of requests for authentication information by the server and responses by the client. In order for authentication to be successful, the client and the server must use the same authentication method.

**EAP-TLS**
EAP-Transport Layer Security (TLS) is an EAP type that is used in certificate-based security environments, and it provides the strongest authentication and key determination method. EAP-TLS provides mutual authentication, negotiation of the encryption method, and encrypted key determination between the client and the authenticating server. If you want to use certificates or smart cards for user and client computer authentication, you must use EAP-TLS or, for enhanced security, Protected EAP (PEAP) with EAP-TLS.

**EAP-MS-CHAP v2**
EAP-Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) is a mutual authentication method that supports password-based user or computer authentication. During the EAP-MS-CHAP v2 authentication process, both the server and client must prove that they have knowledge of the user's password in order for authentication to succeed. With EAP-MS-CHAP v2, after successful authentication, users can change their passwords, and they are notified when their passwords expire.

### Note

- EAP-MS-CHAP v2 is available only with PEAP.

**PEAP**
PEAP is an authentication method that uses TLS to enhance the security of other EAP authentication protocols. PEAP provides the following benefits: an encryption channel to protect EAP methods running within PEAP, dynamic keying material generated from TLS, fast reconnect (the ability to reconnect to a wireless access point by using cached session keys, which allows for quick roaming between wireless access points), and server authentication that can be used to protect against the deployment of unauthorized wireless access points.

**PEAP authentication process**
The PEAP authentication process consists of two main phases:

1. **Server authentication and the creation of a TLS encryption channel**. The server identifies itself to a client by providing certificate information to the client. After the client verifies the identity of the server, a master secret is generated. The session keys that are derived from the master secret are then used to create a TLS encryption channel that encrypts all subsequent communication between the server and the wireless client.

2. **EAP conversation and user and client computer authentication**. A complete EAP conversation between the client and the server is encapsulated within the TLS encryption channel. With PEAP, you can use any one of several EAP authentication methods, such as passwords, smart cards, and certificates, to authenticate the user and client computer.

The session keys that are generated during the PEAP authentication process provide keying material for the Wired Equivalent Privacy (WEP) encryption keys that encrypt the data that is sent between wireless clients and wireless access points.

You can use PEAP with any of the following authentication methods for wireless authentication:

- EAP-TLS, which uses certificates for server authentication and either certificates or smart cards for user and client computer authentication.
- EAP-MS-CHAP v2, which uses certificates for server authentication and credentials for user authentication.
- Non-Microsoft EAP authentication methods.

### Notes

- PEAP is not supported for use with EAP-MD5.
- PEAP is available as an authentication method for 802.11 wireless clients, but it is not supported for virtual private network (VPN) clients or other remote access clients. Therefore, you can configure PEAP as the authentication method for a remote access policy only when you are using Internet Authentication Service (IAS).

**PEAP support for fast reconnect**

When clients connect to an 802.11 wireless network, the authenticated session has a set expiration interval that the network administrator determines to limit the duration of authenticated sessions. To avoid the notices that require authenticated clients to provide their credentials to reauthenticate and resume a session, you can enable the fast reconnect option.

PEAP supports fast reconnect, which allows roaming users to maintain continuous wireless network connectivity when traveling between different wireless access points on the same network, as long as each wireless access point is configured as a client of the same IAS (RADIUS) server. In addition, fast reconnect must be enabled on both the wireless client and the RADIUS server.

After the initial PEAP authentication succeeds, users are not prompted to supply their credentials (if PEAP with EAP-MS-CHAP v2 is used for authentication) or their personal identification number (PIN) (if PEAP with smart cards is used for authentication) each time they travel to a new wireless access point. Instead, when PEAP fast reconnect is enabled, after the initial PEAP authentication succeeds, the client and the server cache TLS session keys. When users associate with a new wireless access point, the client and the server use the cached keys to reauthenticate each other until the cache has expired. Because the keys are cached, the RADIUS server can quickly determine that the client connection is a reconnect. This reduces the delay in time between an authentication request by a client and the response by the RADIUS server. It also resource requirements for the client and the server.

If the original RADIUS server is not used, full authentication must occur, and the user is again prompted for credentials or a PIN. This can occur in the following situations:

- The user associates with a new wireless access point that is configured as a client of a new RADIUS server.
- The user associates with the same wireless access point, but the wireless access point forwards the authentication request to a new RADIUS server.

In both situations, after the initial authentication with the new RADIUS server succeeds, the client caches the new TLS session keys. Clients can cache TLS session keys for multiple RADIUS servers.

For more information, see EAP, MS-CHAP version 2, and PEAP.

For more information about configuring the session time-out for wireless remote access policies, see Introduction to remote access policies, Configure Remote Access Policies, and Elements of a remote access policy.

**Security and deployment considerations**

When choosing an authentication method, balance the level of security that you require with the effort that you want to devote to deployment. For the highest level of security, choose PEAP with certificates (EAP-TLS). For the greatest ease of deployment, choose PEAP with passwords (EAP-MS-CHAP v2).

Although both PEAP with EAP-TLS and EAP-TLS alone provide strong security through the use of certificates for server authentication and either certificates or smart cards for client computer and user authentication, when PEAP with EAP-TLS is used, client certificate information is encrypted. PEAP with EAP-MS-CHAP v2 requires the least effort to deploy because client authentication is password-based, so certificates or smart cards do not need to be installed on clients. Because PEAP creates an

end-to-end encrypted channel before EAP-MS-CHAP v2 authentication occurs, the authentication exchange reduces the risk of offline dictionary attacks.

For information about certificate requirements for 802.1X authentication methods, see Network access authentication and certificates. For information about deploying smart cards, see Checklist: Deploying smart cards for logging on to Windows.

**Important**

- When you deploy both PEAP and EAP unprotected by PEAP, do not use the same EAP authentication type with and without PEAP. For example, if you deploy PEAP with EAP-TLS (PEAP-EAP-TLS), do not also deploy EAP-TLS without PEAP. Deploying authentication methods with the same type--one with and the other without the protection of PEAP--creates a security vulnerability.

- Certificate Templates provide a new version 2 certificate template, named Workstation Authentication. This certificate template, which can be configured for auto-enrollment, is available for enterprise certification authorities (CAs) on computers running Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or the 64-bit versions of the Windows Server 2003 family. Certificates based on this template allow client computers to authenticate their identity to servers. You should use this certificate template if you are both using certificate-based authentication for wireless clients running Windows XP and providing an enterprise CA on a computer running Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or the 64-bit versions of Windows Server 2003 family. If you use the Computer certificate template in this case, client authentication will fail. For more information, see Certificate Templates and Manage Certificate Templates for an Enterprise Certification Authority.

**How 802.1X works on 802.11 wireless networks**

802.1X implements port-based network access control. Port-based network access control uses the physical characteristics of a switched local area network (LAN) infrastructure to authenticate computing devices that are attached to a LAN port and to prevent access to that port when the authentication process fails.

During a port-based network access control interaction, a LAN port adopts one of two roles: authenticator or supplicant. In the role of authenticator, a LAN port enforces authentication before it allows user access to the services that can be accessed through that port. In the role of supplicant, a LAN port requests access to the services that can be accessed through the authenticator's port. An authentication server, which can either be a separate entity or co-located with the authenticator, checks the supplicant's credentials on behalf of the authenticator. The authentication server then responds to the authenticator, indicating whether the supplicant is authorized to access the authenticator's services.

The authenticator’s port-based network access control defines two logical data paths to the LAN through one physical LAN port. The first data path, the uncontrolled port, allows data exchange between the authenticator and a computing device on the LAN, regardless of the authentication state of that device. This is the path that EAPOL (EAP over LAN) messages will take. The second data path, the controlled port, allows data exchange between an authenticated LAN user and the authenticator. This is the path that all other network traffic will take, after the computing device is authenticated.

**802.1X and IAS**

To support authentication, authorization, and accounting for wireless network connections, you can use 802.1X with IAS. IAS is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy server. When RADIUS is implemented, a wireless access point prevents data traffic from being forwarded to a wired network or to another wireless client without a valid authentication key. The process of obtaining a valid authentication key is as follows:

1. When a wireless client comes within range of a wireless access point, the wireless access point challenges the client.

2. The wireless client sends its identity to the wireless access point, which forwards this information to a RADIUS server.

3. The RADIUS server requests the credentials of the wireless client to verify the identity of the client. As part of this request, the RADIUS server specifies the type of credentials that are required.

4. The wireless client sends its credentials to the RADIUS server.

5. The RADIUS server verifies the credentials of the wireless client. If the credentials are valid, the RADIUS server sends an encrypted authentication key to the wireless access point.

6. The wireless access point uses this authentication key to securely transmit per-station unicast session and multicast authentication keys to the wireless client.

For information about configuring IAS for wireless access authentication, see Checklist: Configuring the IAS server and wireless access points for wireless access and Wireless access.

For information about how to configure 802.1X settings for clients, see Define 802.1X authentication for wireless networks on a client computer.

For general information about security for wireless networks, see Security information for wireless networks.

# Understanding wireless provisioning services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding wireless provisioning services**

Wireless Provisioning Services (WPS) are enhancements to wireless networking in Windows XP and the Windows Server 2003 family. WPS enables wireless networks to provide a consistent sign-up and configuration experience while providing more secure access to network resources.

When you connect to a wireless network that supports WPS, you are prompted to download additional files that are used by WPS to enable a more secure, consistent, and automatic setup. These files are generally safe to download, but the information contained within them might not be accurate. You should verify that the Web site from which the files will be downloaded is one that you expected based on your location, or one that you trust to provide you with the information.

If you choose to download the files, XML files are downloaded from the provider and stored on your computer for as long as you wish to use the wireless network. The files contain no personal information about you or your computer.

If you do not choose to download the files, you cannot connect to the wireless network.

When the files are downloaded, the Wireless Network Registration Wizard requests information from you. Some networks might require payment. Other networks might simply require acknowledgement of a usage policy. Follow the instructions in the Wireless Network Registration Wizard to provide the appropriate information to the network provider.

After you have signed up with a wireless network that supports WPS, the files stored on your computer can be updated automatically. If you choose this option, the wireless provisioning services in your computer will connect to the wireless network provider and update the information stored on your computer. This occurs on a schedule determined by the wireless network provider (typically once a week or once a month). This update process will only occur while you are already connected to the Internet and will not interfere with your use of your Internet connection. If you choose not to update the files automatically, the files for a specific network will be updated the next time you connect to that wireless network. (This is required to ensure that your computer has the right information for that network.)

For more information about connecting to wireless networks, see Connect to an available wireless network.

For conceptual information about configuring wireless client computer settings, see Configuring wireless network settings on client computers.

# Understanding Infrared Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding infrared networking**
This section covers:

- Infrared communication
- Infrared implementation
- Links between infrared devices

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Infrared communication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Infrared communication**
Infrared communication uses infrared light to transfer data. Infrared light is also used almost universally by TV and VCR remote controls. In computers, infrared communication offers an alternative to cables and floppy disks. Infrared communication provides a point-to-point, low-cost way to connect computers with each other or with devices and appliances. Many cellular phones are equipped with infrared ports that allow them to be connected to a computer for dial-up networking connections.

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Infrared implementation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Infrared implementation**
Infrared data transfer is implemented according to the Infrared Data Association (IrDA) standards and protocols. These standards are designed to allow low-cost components and low power requirements, and to enable connections by simply pointing infrared devices at each other. Infrared transceivers are now installed in nearly all new portable computers. For a computer that does not have a built-in infrared transceiver, you can install an external infrared transceiver.

IrDA is a half-duplex, short-range data transfer technology. The IrDA protocols specify the procedures that support link initialization, device address discovery, connection startup and data rate negotiation, information exchange, disconnection, link shutdown, and device address conflict resolution.

**Infrared device support**
Infrared functionality is provided through the Wireless Link file transfer feature, infrared printing (IrLPT), infrared image transfer (IrTran-P), and infrared networking (IrNET and IrComm) capabilities. Additionally, the IrDA Winsock API supports programs created by other software and equipment manufacturers. These manufacturers sell programs that use the Winsock API (or proprietary interfaces) to provide infrared connections to printers, modems, digital pagers, personal digital assistants, electronic cameras, organizers, cellular phones, and hand-held computers.

### Important

- For enhanced security, performance, and stability, it is recommended that you use only signed device drivers. For more information, see Signed Drivers.

**Infrared transmission speeds**
Many laptops, notebook computers, and hand-held devices now include infrared transceiver ports that provide infrared asynchronous serial transmission with maximum data transfer speeds of either 115.2 kilobytes per second (Kbps) or 4 megabits per second (Mbps), and sometimes 16 Mbps.

**Serial IrDA 115.2 Kbps support (SIR)**
The Serial IrDA (SIR) implementation of IrDA provides a maximum data transfer speed of 115.2 Kbps. The primary benefit of this implementation is that existing serial hardware can be used without additional cost.

**Fast IrDA 4.0 Mbps support (FIR) and Very Fast IrDA 16.0 Mbps support (VFIR)**
The Fast IrDA (FIR) implementation of IrDA provides a maximum data transfer speed (of 4 Mbps) that easily adjusts to slower devices. Very Fast IrDA (VFIR), which provides 16 Mbps half-duplex data transmission, is also supported. FIR and VFIR devices can communicate with SIR devices.

**IrComm and IrNET support**
With IrComm, you can use your computer and infrared-equipped cellular phone to connect to your Internet service provider or to send fax transmissions without additional equipment. With IrNET, you can establish a point-to-point connection between your computer and another computer or another infrared-equipped device. You can also establish a connection between your computer and a network access point.

**IrLPT printer support**
If your computer came with an internal infrared device, or if you install an infrared transceiver, an infrared port appears as a local port in the **Add Printer** dialog box. When you associate a printer with this port and then print to that printer, the IrLPT protocol is used to transmit the data.

**IrTran-P image transfer support**
The IrTran-P image transfer protocol is used in digital cameras and other digital image capture devices. You can use this feature to receive digital images from a camera or other device that supports IrTran-P to your computer through an IrDA connection. The IrTran-P service is implemented as a listen-only service; it never initiates the IrTran-P connection.

### Important

- To enhance security, do not enable IrTran-P. When you enable IrTran-P, all incoming files sent over IrTran-P are automatically accepted. Because incoming files might contain harmful programs, ensure that the files originate from a

trustworthy source. Do not open files if you cannot verify the source, do not recognize the file format, or are unsure of the content. Instead, delete the files immediately. For more information about infrared security considerations, see Security information for infrared communication.

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Links between infrared devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Links between infrared devices**
Infrared data-link transmission occurs between two infrared devices. All transmissions over the link go from the primary (commanding) device to the secondary (receiving) device. The primary role is determined dynamically when a link is established and continues until the connection is closed. Any station that is capable can assume the primary role. When two computers are in range, either one can assume the primary role, so a user can initiate a data transfer on either computer. Some devices may only be capable of assuming the secondary role.

**How an infrared connection is made**
A device creates a link by automatically detecting the other or by a user request. The commanding station sends a connection request at 9,600 bits per second (bps) to the other device (including information such as an address, data rate, and other capabilities). The responding device assumes the secondary role and returns information that contains its address and capabilities. The primary and secondary stations then change the data rate and link parameters to the common set defined by the initial information transfer. Finally, the primary station sends data to the secondary station confirming the connection. The devices are then connected and begin data transfer under control of the primary device.

**Multiple infrared connections**
The Winsock API supports multiple simultaneous connections over a single IrDA link. This feature allows several different programs to use a single infrared device simultaneously. For example, you might place a portable computer within range of a desktop computer to simultaneously send and receive mail, update your calendar and contacts, and print deferred documents. Each activity can be controlled by a separate program on the laptop computer that locates and connects to its corresponding program on the desktop computer.

**Multiple infrared devices**
A single infrared device cannot link simultaneously with more than one other infrared device. However, you can install multiple infrared devices on a computer to provide simultaneous links to multiple infrared devices. For example, using separate infrared devices, a desktop computer can simultaneously print to an infrared port, communicate with a portable computer, and dial a network connection.

For additional information, see Installing infrared devices.

### Important

- For enhanced security, performance, and stability, it is recommended that you use only signed device drivers. For more information, see Signed Drivers.

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Security information for infrared communication.

### Note

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Using Wireless Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using wireless networking**
This section covers:

- Security information for wireless networks
- Creating, modifying, and assigning Active Directory-based wireless network policies
- Using Resultant Set of Policy to view wireless network policy assignments
- Configuring wireless network settings on client computers
- Logging and viewing wireless network activity
- Using Infrared Communication

# Security information for wireless networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for wireless networks**
Wireless networking technologies provide convenience and mobility, but they also introduce security risks on your network. For example, unless authentication and authorization mechanisms are implemented, anyone who has a compatible wireless network adapter can access the network. Without encryption, wireless data is sent in plaintext, so anyone within sufficient distance of a wireless access point can detect and receive all data sent to and from a wireless access point.

The following security mechanisms enhance security over wireless networks:

- Windows Firewall
- 802.11 identity verification and authentication
- 802.11 Wired Equivalent Privacy (WEP) encryption
- Wi-Fi Protected Access (WPA)
- 802.1X authentication
- IAS support for 802.1X authentication

**Windows Firewall**
Windows Firewall runs on each of your clients and servers. It provides protection from network attacks that pass through your perimeter network or originate inside your organization, such as Trojan horse attacks, port-scanning attacks, and worms. Like many firewall technologies, Windows Firewall is a stateful firewall. It drops all incoming traffic unless it has been sent in response to a request from the computer or has been specified as allowed. With the exception of some Internet Control Message Protocol (ICMP) messages, Windows Firewall allows all outgoing traffic.

Windows Firewall is not included in the original release of the Windows Server 2003 operating systems.

In Windows Server 2003 with Service Pack 1 (SP1), Windows Firewall is disabled by default for all connections, including LAN (wired and wireless), dial-up, and virtual private network (VPN) connections. Windows Firewall is also disabled by default for new connections.

If your wireless computers do not already have a firewall installed and running, it is recommended that you turn on Windows Firewall. Right-click the wireless icon in the notification area, click **Change Windows Firewall settings**, and then select **On**.

For information about Windows Firewall, see Help: Turn Windows Firewall on or off for a specific connection, Help: Add a program to the Windows Firewall exceptions list, Help: Add a port to the Windows Firewall exceptions list, and Windows Firewall and Message Queuing.

**802.11 identity verification and authentication**
For identity verification and authentication, IEEE 802.11 defines the open system and shared key authentication subtypes:

- Open system authentication does not actually provide authentication; it performs only identity verification through the exchange of messages between the initiator (wireless client) and the receiver (wireless access point).
- Shared key authentication does provide authentication by verifying that the initiator has knowledge of a shared secret. Under the 802.11 standard, it is assumed that the shared secret is sent to the wireless access point over a secure channel that is independent of 802.11.

For information about how to specify the authentication subtype that you want to use, see Define preferred wireless networks in Group Policy or Define a wireless network connection on a client computer.

**Important**

- To enhance security and connectivity, do not use shared key authentication. Shared key authentication is less secure than open system authentication because it requires the exchange of a secret key that is shared by all wireless access points and clients and therefore is more vulnerable to known-text attacks. In addition, if you use shared key authentication for a wireless network that has multiple wireless access points, you will lose network connectivity when you travel from one wireless access point to a new wireless access point. In this case, you will lose connectivity because your network key will no longer match the shared key used by all wireless access points. To determine whether a wireless network you are connecting to has multiple wireless access points, use Wireless Monitor. For information about how to use Wireless Monitor to view details about wireless access points, see View details about wireless network access points.

**802.11 WEP encryption**
For encryption, 802.11 defines the WEP algorithm. WEP provides data confidentiality by encrypting the data that is sent between wireless clients and wireless access points.

To encrypt the data that is sent over wireless networks, WEP uses the RC4 stream cipher with either a standard 40-bit encryption key or, in some implementations, a 104-bit encryption key. A stream cipher is a method of encrypting text (to produce ciphertext) in which a cryptographic key and algorithm are applied to each binary digit in a data stream, one bit at a time. The RC4 stream cipher, designed by RSA Data Security, Inc can accept keys of arbitrary length. Data integrity is provided through an integrity check value (ICV) in the encrypted portion of the wireless frame.

**WPA**
WPA is a new wireless network security technology developed by the Wi-Fi Alliance. Wi-Fi Protected Access strengthens existing cryptographic weaknesses of WEP, and introduces a method to generate and distribute encryption keys automatically. The solution also introduces an integrity check on the data so an attacker cannot modify packets of information being communicated. To improve enterprise-level user authentication, Wi-Fi Protected Access authenticates every user on the network while keeping those users from joining rogue networks. WPA presents a practical answer to dealing with the weaknesses of WEP, based on available technologies and offering forward compatibility with 802.11i and backward compatibility with existing 802.11 solutions.

**802.1X authentication**
802.1X is an IEEE standard for authenticated network access to wired Ethernet networks and wireless 802.11 networks. IEEE 802.1X provides support for centralized user identification, authentication, dynamic key management, and accounting. The 802.1X standard enhances security by allowing a computer and a network to authenticate each other, generating a per-user/per-session key to encrypt data over wireless connections and providing the ability to dynamically change keys.

### Important

- For enhanced security in Windows XP Service Pack 1 and the Windows Server 2003 family, 802.1X authentication is available only for access point (infrastructure) networks that require the use of a network key (WEP).

- It is highly recommended that you use 802.1X authentication whenever you connect to an 802.11 wireless network. If you connect to an 802.11 wireless network, and you do not enable 802.1X, the data that you send is more vulnerable to attacks, such as offline traffic analysis, bit flipping, and malicious packet injection.

**EAP authentication methods**
802.1X uses Extensible Authentication Protocol (EAP) for message exchange during the authentication process. With EAP, an arbitrary authentication method, such as passwords, smart cards, or certificates, is used to authenticate the wireless connection. The support that 802.1X provides for EAP types allows you to use any of the following authentication methods:

- EAP-Transport Level Security (EAP-TLS), which uses certificates for server authentication and either smart cards or certificates for user and client computer authentication.

- Protected EAP with EAP-Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP with EAP-MS-CHAPv2), which uses certificates for server authentication and credentials (user name and password) for user authentication.

- PEAP with EAP-TLS, which uses certificates for server authentication and either smart cards or certificates for user and client computer authentication.

For more information, see EAP, MS-CHAP version 2, Understanding 802.1X authentication for wireless networks, and PEAP.

**Security and ease of deployment**

When choosing an authentication method, balance the level of security that you require with the effort that you want to devote to deployment. For the highest level of security, choose PEAP with certificates (EAP-TLS). PEAP uses TLS to enhance the security of other EAP authentication protocols. With PEAP, TLS is used to create an end-to-end encrypted channel between an EAP client (such as a wireless computing device) and an EAP server (such as an Internet Authentication Service (IAS) server). Although both PEAP with EAP-TLS and EAP-TLS alone provide strong security through the use of certificates for server authentication and either certificates or smart cards for client computer and user authentication, when PEAP with EAP-TLS is used, client certificate information is encrypted.

For the greatest ease of deployment, choose PEAP with passwords (EAP-MS-CHAP v2). PEAP with EAP-MS-CHAP v2 requires the least effort to deploy, because client authentication is password-based, so certificates or smart cards do not need to be installed on clients. Because PEAP creates an end-to-end encrypted channel before EAP-MS-CHAP v2 authentication occurs, the authentication exchange reduces the risk of offline dictionary attacks.

The session keys that are generated during the PEAP authentication process provide keying material for the WEP encryption keys that encrypt the data that is sent between wireless clients and wireless access points. In addition, PEAP supports fast reconnect. PEAP fast reconnect allows roaming users to maintain continuous wireless network connectivity when traveling between different wireless access points on the same network, as long as each wireless access point is configured as a client of the same IAS (RADIUS) server.

For information about certificate requirements for 802.1X authentication methods, see Network access authentication and certificates. For information about deploying smart cards, see Checklist: Deploying smart cards for logging on to Windows.

**Important**

- When you deploy both PEAP and EAP unprotected by PEAP, do not use the same EAP authentication type with and without PEAP. For example, if you deploy PEAP with EAP-TLS (PEAP-EAP-TLS), do not also deploy EAP-TLS without PEAP. Deploying authentication methods with the same type--one with and the other without the protection of PEAP--creates a security vulnerability.

**IAS support for 802.1X authentication**

To enhance the security and deployment of wireless networks, you can use 802.1X with IAS, the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy server. When RADIUS is implemented, wireless access points that are configured as RADIUS clients use the RADIUS protocol to send connection requests and accounting messages to a central RADIUS server. The RADIUS server accesses a user account database and set of rules for granting authorization, processes the wireless access point's connection request, and then either accepts or rejects the connection request.

For more information about 802.1X authentication, see Understanding 802.1X authentication for wireless networks. For information about wireless network client configuration, see Configuring wireless network settings on client computers. For information about configuring IAS for wireless access authentication, see Checklist: Configuring the IAS server and wireless access points for wireless access and Wireless access.

# Creating, modifying, and assigning Active Directory-based wireless network policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating, modifying, and assigning Active Directory-based wireless network policies**
Wireless network settings can be configured locally, by users on client computers, or centrally. To enhance the deployment and administration of wireless networks, you can use Group Policy to create, modify, and assign wireless network policies centrally, for Active Directory clients. When you use Group Policy to define wireless network policies, you can specify whether clients can use Windows to configure wireless network connection settings, whether to enable 802.1X authentication for wireless network connections, and the preferred wireless networks that clients can connect to.

When you assign an Active Directory-based wireless network policy, the policy is stored in Active Directory and it is cached in the registry on the local computer. If a computer is temporarily disconnected from a Windows 2000 or a Windows Server 2003 domain, the Active Directory policy that is cached in the local registry is applied instead, if such a policy has been cached.

A Group Policy object defines access, configuration, and usage settings for accounts and resources. Wireless network policies can be assigned to the Group Policy object of a site, domain, or organizational unit. When the wireless network policy is applied to one of the Group Policy objects for the Active Directory object, the wireless network policy is propagated to any computer accounts that are affected by that Group Policy object. For more information, see Group Policy overview.

When assigning a wireless network policy in Active Directory, consider the following:

- Only a single wireless network policy can be assigned at a specific level in Active Directory.
- A wireless network policy that is assigned for the domain takes precedence over local wireless network settings on a client computer, when a client computer is a member of the domain.
- A wireless network policy that is assigned to an organizational unit in Active Directory takes precedence over a domain-level policy for members of that organizational unit.
- An organizational unit inherits the policy of its parent organizational unit unless either policy inheritance is explicitly blocked or policy is explicitly assigned.
- The highest possible level of the Active Directory hierarchy should be used to assign policies to reduce the amount of configuration and administration required.

For information about how to define Active Directory-based wireless network policies, see Define Active Directory-based Wireless Network Policies.

**Note**

- The integration of IPSec policies with Active Directory described here does not apply to computers running Windows XP Home Edition. You cannot administer Active Directory-based IPSec policy from a computer running Windows XP Home Edition because these computers cannot join Active Directory domains.

# Using Resultant Set of Policy to view wireless network policy assignments

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Resultant Set of Policy to view wireless network policy assignments**

Resultant Set of Policy (RSoP) is an addition to Group Policy that you can use to view wireless network policy assignments for a computer or for members of a Group Policy container. This information can help you troubleshoot policy precedence issues and plan your deployment.

To view wireless network policy assignments in RSoP, you must first open the RSoP MMC console, and then run a query. RSoP provides two types of queries: logging mode queries (for viewing wireless network policy assignments for a computer) and planning mode queries (for viewing wireless network policy assignments for members of a Group Policy container).

**Logging mode queries**

You can run an RSoP logging mode query to view all of the wireless network policies that are assigned to an wireless network client. The query results display the precedence of each wireless network policy assignment, so that you can quickly determine which wireless network policies are assigned but are not being applied and which wireless network policy is being applied. The RSoP console also displays detailed settings (that is, whether 802.1X authentication is enabled, the list of preferred wireless networks that clients can connect to, and wireless network key settings) for the wireless network policy that is being applied.

When you run a logging mode query, RSoP retrieves policy information from the Windows Management Instrumentation (WMI) repository on the target computer, and then displays this information in the RSoP console. In this way, RSoP provides a view of the policy settings that are being applied to a computer at a given time.

For information about how to run RSoP logging mode queries for wireless network policies, see Use RSoP to view wireless network policy assignments for a computer. For general information about RSoP logging mode queries, see RSoP logging mode. For general information about RSoP, see RSoP overview.

**Planning mode queries**

You can run an RSoP planning mode query to view all of the wireless network policies that are assigned to members of a Group Policy container. For example, a planning mode query can be useful if you are planning a company reorganization and you want to move computers from one organizational unit to a new organizational unit. By supplying the appropriate information and then running a planning mode query, you can determine which wireless network policies that are assigned but are not being applied to the new organizational unit and which wireless network policy is being applied. In this way, you can identify which policy would be applied if you were to move the computers to the new organizational unit. As with logging mode queries, when you run a planning mode query, the RSoP console displays detailed policy settings for the wireless network policy that is being applied.

When you run a planning mode query, RSoP retrieves the names of the target user, computer, and domain controller from the WMI repository on the domain controller. WMI then uses the Group Policy Data Access Service (GPDAS) to create the policy settings that would be applied to the target computer, based on the RSoP query settings that you entered. RSoP reads the policy settings from the WMI repository on the domain controller, and then displays this information in the RSoP console user interface.

For information about how to run RSoP planning mode queries for wireless network policies, see Use RSoP to view wireless network policy assignments for members of a Group Policy container. For general information about RSoP planning mode queries, see RSoP planning mode.

**Note**

- You can run an RSoP planning mode query only on a domain controller (when you run a planning mode query, you must explicitly specify the domain controller name). However, you can specify any wireless network client as the target for the query, provided you have the appropriate permissions to do so.

**Assigning and processing wireless network policies in Group Policy**

Wireless network policies can be assigned from and stored in Active Directory, as part of Group Policy, or they can be assigned and stored locally, on a computer. When a computer is joined to an Active Directory domain, the domain-level wireless network policy applies. If a computer is not joined to an Active Directory domain, the local settings apply.

Group Policy settings are contained in Group Policy objects, which are linked with specific Active Directory objects (sites, domains, and organizational units). When a wireless network policy is assigned to a Group Policy object for an Active Directory object (such as an organizational unit), that wireless network policy is propagated to any computer accounts that are affected by the Group Policy object.

Multiple Group Policy objects, each of which can contain an wireless network policy, can be assigned to a computer account. When multiple wireless network policies are assigned, the last policy that is processed is the policy that is applied (that is, the last policy takes the highest precedence and overrides the settings of any wireless network policy assignments that were processed earlier).

Policy precedence is based on the Group Policy inheritance model. The policy used is the policy assigned at the lowest level of the domain hierarchy for the domain container of which the computer is a member. For example, if there are wireless network policies that are configured for both the domain and for an organizational unit within the domain, the computers that are members of the domain use the domain wireless network policies. The computers that are members of the organizational unit within the domain use the organizational unit wireless network policies. If there are multiple organizational units, members of each organizational unit use the wireless network policy assigned to the organizational unit that is closest in level to their container in the Active Directory hierarchy. If no wireless network policies are configured for Active Directory, or if a computer is not connected to an Active Directory domain, the local wireless settings are used.

**Wireless network policy information displayed in the RSoP console**

The RSoP snap-in simplifies the task of determining which wireless network policy is being applied by displaying the following information for each Group Policy object that contains a wireless network policy assignment: the name of the wireless network policy, the name of the Group Policy object that the wireless network policy is assigned to, the wireless network policy precedence (the lower the number, the higher the precedence), and the name of the site, domain, and organizational unit to which the Group Policy object containing the wireless network policy applies (that is, the scope of management for the Group Policy object).

For general information about Group Policy, see Group Policy overview.

# Configuring wireless network settings on client computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring wireless network settings on client computers**

The Wireless Configuration service in Windows XP and the Windows Server 2003 family supports the IEEE 802.11 standard for wireless networks and minimizes the configuration that is required to access wireless networks. This service is enabled by default and allows you to travel to different wireless networks without reconfiguring all of the network connection settings on your computer at each location. As you move from one location to another, the Wireless Configuration service searches and sends notifications when there are wireless networks available. When connecting to a new wireless network, you will need only to configure the settings specific to that network. The Wireless Configuration service updates your wireless network adapter to match those settings and attempts to connect to the wireless network.

The Windows Configuration service allows you to use Windows to configure your wireless network settings. When you use Windows to configure your wireless network settings, you can create a list of preferred wireless networks and you can specify the order in which to attempt connections to these networks.

For more information, see Define Wireless Network Policies on a Client Computer.

**Viewing available 802.11 wireless networks**

In addition to the wireless network notifications that are provided by the Wireless Configuration service, in Windows XP with Service Pack 2 and Windows Server 2003 with Service Pack 1, computers equipped with wireless network adapters have an enhanced interface called **Wireless Network Connection**.

The **Wireless Network Connection** interface was designed with several user benefits in mind. The interface makes it easier to view information about available 802.11 wireless networks, including network type (infrastructure or ad hoc), wireless signal strength, state (connected or disconnected), and whether networks are security enabled. The **Wireless Network Connection** interface also makes it easier to configure and control connections to 802.11 wireless networks. Direct access to 802.11 wireless properties, where both general 802.11 wireless settings and advanced configurations are made, can be accessed through the following interface links: **Related Tasks**, **Change the order of preferred networks**, and **Change advanced settings**.

To control connections to wireless networks, under **Choose a wireless network**, click **Connect** or **Disconnect**. To refresh the list of available wireless networks, use **Network Tasks**.

As an example, while attempting to connect to a wireless network, under **Choose a wireless network**, you might see status that reads **Limited or no connectivity**. Click the **Limited or no connectivity** link for more information, including possible remedies for connectivity problems.

**Wireless icons**

You can access the **Wireless Network Connection** interface by using new wireless icons that are provided in both the notification area and in Network Connections. To open the **Wireless Network Connection** interface, right-click the wireless icon, and then select **View Available Wireless Networks**. For more information, see Define Wireless Network Policies on a Client Computer.

**Notification area wireless icons**

In Windows XP with Service Pack 2 and Windows Server 2003 with Service Pack 1, wireless icons are displayed to indicate the state of the wireless connection:

| Notification area icon | Wireless connection state |
|---|---|
| | Connected |
| | Connecting |
| | Not connected |
| | Warning |

The warning icon is displayed when an IP address error is detected while attempting to connect to a wireless network. In the event of an error, the warning icon will always be displayed in the notification area, provided the device has not been disabled or the **Show icon in notification area when connected** check box on the properties page of the device has not been

cleared.

**Note**

- When the wireless connection is in any of these states, you can right-click the wireless connection icon in the notification area and then select **View Available Wireless Networks**. Wireless adapters that have been disabled by right-clicking the wireless connection icon, and then clicking **Disable**, will not appear in the notification area. Disabled wireless adapters can be enabled only by right-clicking the wireless network connection icon in Network Connections, and then clicking **Enable**.

**Network Connections wireless icons**

There are five views that can be applied to icons in Network Connections: Thumbnails, Tiles, Icons, List, and Details. In each of these views, the various states of wireless connections are represented in a similar way to the wireless icons in the notification area. For example, a wireless icon with a red X indicates there is no connection to a wireless network; a wireless icon with a yellow triangle indicates an error state.

Wireless icons in Network Connections do not indicate connecting or transitioning states. However, if you select the wireless icon and then expand **Details** in the left-hand pane, you have access to the following information:

| **Information type** | **Example or explanation:** |
|---|---|
| Connection status | • Acquiring Network address<br>• Not connected, Firewalled<br>• Connected, Firewalled |
| Assigned IP Address | Displays the IP address that is assigned to the wireless adapter |
| Subnet Mask | Displays the associated subnet mask. |
| Method of IP address and subnet mask assignment | • Assigned by DHCP<br>• Manually Configured |
| Network type | • Access Point<br>• Peer-to-Peer (same as computer to computer) |
| Network | The name of the network. This is also known as the SSID. |
| Signal Strength: | Excellent |

**Selecting a wireless network type**

When you configure new or existing wireless network connections or connect to an available wireless network, you can choose from the following wireless network types:

- **Access point (infrastructure)**

  In access point wireless networks, wireless clients (computing devices with wireless network adapters, such as your portable computer or personal digital assistant) connect to wireless access points. The access points function as bridges between wireless clients and the existing network backbone. As you move from one location to another, and the signal for one wireless access point weakens, or the access point becomes congested with traffic, you can connect to a new access point. For example, if you work in a large corporation, you might connect to several different wireless access points as you move between different floors of a building or different buildings in a campus, while still maintaining uninterrupted access to network resources.

- **Computer-to-computer (ad hoc)**

In computer-to-computer wireless networks, wireless clients connect to each other directly, rather than through wireless access points. For example, if you are in a meeting with co-workers, and you do not need to gain access to network resources, your wireless device can make direct connections to the wireless devices of your co-workers, and you can form a temporary network.

- **Any available network (access point preferred)**

  In access point preferred wireless networks, a connection to an access point wireless network is always attempted first, if there are any available. If an access point network is not available, a connection to a computer-to-computer wireless network is attempted. For example, if you use your laptop at work in an access point wireless network, and then you take your laptop home to use in your computer-to-computer home network, the Windows Configuration service will change your wireless network settings as needed so that you can connect to your home network.

#### Configuring wireless network security settings

802.1X is an IEEE standard for authenticated network access to wired Ethernet networks and wireless 802.11 networks. IEEE 802.1X enhances security and deployment by providing support for centralized user identification, authentication, dynamic key management, and accounting. For enhanced security in Windows XP Service Pack 1 and the Windows Server 2003 family, 802.1X authentication is available only for access point (infrastructure) networks that require the use of a network key (WEP). Wired Equivalent Privacy (WEP) provides data confidentiality by encrypting the data that is sent between wireless clients and wireless access points.

**Important**

- It is highly recommended that you use 802.1X authentication whenever you connect to an 802.11 wireless network. If you connect to an 802.11 wireless network without using 802.1X authentication, the data that you send is more vulnerable to attacks, such as offline traffic analysis, bit flipping, and malicious packet injection. For more information about security considerations for wireless networks, see Security information for wireless networks.

#### Configuring wireless network key settings

When Wired Equivalent Privacy (WEP) is enabled, a network key can be used to encrypt data that is sent over wireless networks. The network key can be provided for you automatically (for example, if 802.1X is being used for dynamic key distribution), or you can specify the key by typing it yourself. If you specify the key yourself, you can also specify the key index (the location where a specific key is stored).

A wireless client can be configured with up to four keys (the key index values are 0, 1, 2, and 3). When a wireless access point or a wireless client transmits an encrypted message using a key that is stored in a specific key index, the transmitted message indicates the key index that was used to encrypt the message body. The receiving wireless access point or wireless client can retrieve the key that is stored at the key index and use it to decode the encrypted message body.

**Important**

- To enhance security and connectivity, do not use shared key authentication. Shared key authentication is less secure than open system authentication because it requires the exchange of a secret key that is shared by all wireless access points and clients and therefore is more vulnerable to known-text attacks. In addition, if you use shared key authentication for a wireless network that has multiple wireless access points, you will lose network connectivity when you travel from one wireless access point to a new wireless access point. In this case, you will lose connectivity because your network key will no longer match the shared key used by all wireless access points. To determine whether a wireless network you are connecting to has multiple wireless access points, use Wireless Monitor. For information about how to use Wireless Monitor to view details about wireless access points, see View details about wireless network access points.

#### Configuring 802.1X authentication settings for wireless network connections

The support that 802.1X provides for Extensible Authentication Protocol (EAP) types allows you to select from several authentication methods for wireless network connections. For more information about choosing an authentication method, see Understanding 802.1X authentication for wireless networks.

# Logging and viewing wireless network activity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Logging and viewing wireless network activity**

Wireless Monitor allows you to view details about access points and wireless clients. You can use this information to troubleshoot your wireless service. The Wireless Configuration service logs information in Wireless Monitor that allows you to:

- Identify service configuration changes.
- Check the events logged in the Wireless Configuration service log that are generated from outside of your network, such as media event notifications, 802.1X events, and timer expiration events.
- Check how the Wireless Configuration service reacts to external events by following transitions, as they are reflected in the log.

For information about how to use Wireless Monitor, see Monitor Wireless Network Activity.

**Wireless Monitor access point information**

The following table describes the items that are displayed when you use Wireless Monitor to view details about **Wireless Access Point Information**.

| Access point statistics | Description |
|---|---|
| Network Name | Displays the names of the networks that are within the reception range of the wireless adapter on your computer. |
| Network Type | Displays the type of network listed in the logged event. Network Types include Access Point and Peer to Peer. |
| MAC Address | Displays the media access control (MAC) address of the networks that are within the reception range of the wireless adapter on your computer. |
| Privacy | Displays whether privacy is enabled or disabled for any network within the reception range of the wireless adapter on your computer. |
| Signal Strength | Displays the strength of the signals that are broadcast from the networks that are within the reception range of the wireless adapter on your computer. IEEE specifies that 802.11 wireless devices receive at a signal strength range between -76 dBmW (decibel milliwatts) and -10 dBmW, with -10 dBmW indicating the strongest signal; some receivers that are more sensitive may be able to accept weaker signals, possibly as weak as -85 dBmW to -90 dBmW. |
| Radio Channel | Displays the radio channels on which the networks that are within the reception range of the wireless adapter on your computer are broadcasting. |
| Access Point Rate | Displays the data rate that the wireless network will support. |
| Network Adapter GUID | Displays a 16-byte value generated by the operating system in order to identify, uniquely, each wireless adapter on your computer. A globally unique identifier (GUID) is used to identify a particular device or component. |

**Wireless Monitor wireless client information**

The following table describes the items that are displayed when you use Wireless Monitor to view details about **Wireless Client Information**.

| Wireless client statistics | Description |
|---|---|
| Source | Identifies the software that generated the event. Events displayed in Wireless Monitor are generated either by the Wireless Configuration service or EAP over LAN (EAPOL). |
| Type | Displays the type of event. The event type is the classification of the event, such as Error, Warning, Information, or Packet. |
| Time | Displays the time at which the event was logged. |
| Local MAC Address | Displays the media access control (MAC) address of the local network adapter. |
| Remote MAC Address | Displays the MAC address of the remote network interface. This is frequently a wireless access point, or, in wireless computer-to-computer networks, this is the MAC address of the network adapter in the remote computer. |
| Network Name | Displays the name of the wireless network for which the event was generated. Under the 802.11 standard for wireless networks, the network name is also known as a service set identifier (SSID). |
| Description | Provides a brief summary of the logged event. (For a list of the events that are logged in the **Description** field and for brief explanations of these events, see the following table. |

**Wireless Client Information events**

The following table describes the events that are logged in the **Description** column of Wireless Monitor for **Wireless Client Information**.

| Description | Explanation |
|---|---|
| Adding interface "x." | Indicates that a new network adapter has been inserted into the computer running the Wireless Configuration service. This does not necessarily mean that the adapter is a wireless adapter. The Wireless Configuration service determines whether this is a wireless adapter and, if so, ceases to control it (the Wireless Configuration service removes the internal context of the adapter from its hashes). The description includes the NDIS adapter as parameter "x." |
| Configuration "x" has a default random WEP key. Authentication is disabled. Assuming invalid configuration. | Indicates that the Wireless Configuration service associated successfully to a network after providing a random Wired Equivalent Privacy (WEP) key. If 802.1X is not taking control over the connection, then this is considered to be an invalid configuration. If 802.1X is disabled, the Wireless Configuration service deletes this configuration. If 802.1X is taking control, the Wireless Configuration service does not delete the configuration, and the Wireless Configuration service retains the configuration as "Successful." The log contains the SSID of the network to which it is associated as parameter "x." |
| Deleting configuration "x" and moving on (If no better match is found, the configuration will be revived). | Indicates that the Wireless Configuration service is failing to detect the presence of an ad hoc network. This means that the adapter might be the only 802.11 wireless device in the ad hoc network, in which case there is no "Media connect" notification and media status shows "Disconnected" after the two seconds allowed to associate. If this happens, the Wireless Configuration service deletes the network but tags the deleted network such that association will be reattempted if all of the other networks fail. If all other attempts to associate with a network fail, then the tagged network is "Revived," and the Wireless Configuration service will mark it "Successful," regardless of the media indications. This is how the system becomes the first node in the ad hoc network. The description includes the SSID of the ad hoc network that failed as parameter "x." |

| | |
|---|---|
| Deleting configuration "x" and moving on. | Indicates that the Wireless Configuration service attempted to associate to an infrastructure network and failed. There are two possible reasons for the failure:<br>• The Wireless Configuration service provided the 802.11 settings, but two seconds elapsed and there was no "Media connect" notification received or, at the end of the two seconds, media status indicated "Disconnected."<br>• The association succeeded, but 802.1X requested that the Wireless Configuration service disassociate from this network explicitly (as a result of an authentication failure). The description includes the SSID of the network that failed as parameter "x." |
| Driver failed scanning, rescheduling another scan in 5sec. | Indicates that the Wireless Configuration service queried the driver, and the driver reported an error. This can happen if the driver cannot serve the request immediately. In such instances, the Wireless Configuration service schedules another attempt in five seconds. This event does not occur regularly (ideally, it does not occur at all). If this event is logged often, it could indicate that the wireless adapter drivers have been corrupted. |
| Failed to associate to any wireless network. | Indicates that the Wireless Configuration service is in a "Failed" state. In most cases, this event is logged every minute, when:<br>• The Wireless Configuration service is disabled.<br>• The wireless adapter is outside of the broadcast range of any preferred networks.<br>• There is no available wireless network.<br>This description indicates that no networks are detected to which the wireless adapter can associate. In the "Failed" state, the Wireless Configuration service provides a random SSID for this purpose. |
| Hard resetting interface. | Indicates that the wireless adapter has been hard-reset. This usually means that any existing association is broken before the Wireless Configuration service reanalyzes the available networks and the user preferences and configures the appropriate 802.11 settings on the adapter. This event is usually logged when you roam from one network to another and is associated with a "Media disconnect" notification from the previous network. This event is also logged once every minute if the Windows Configuration service is looping through a failure state, indicating that the Wireless Configuration service could not associate to any network. This event does not occur in the case of a steady successful, association. |
| Initiating scanning for wireless networks. | Indicates that the Wireless Configuration service is initiating a scanning cycle. Subsequently, the Wireless Configuration service causes the driver to scan for networks, query the list of available networks, and either associate to a more appropriate network or, if it detects there is no change, retain the current successful association (if any). This event is logged once every minute. |
| No configuration change. Still associated to "x." | Indicates that the Wireless Configuration service processed the list of available networks and the list of user preferred networks and determined there is no change that would allow it to break the existing successful association. Subsequently, the Wireless Configuration service will transition into "Association successful" state. This event is logged once every minute in the case of a steady successful association. The description includes the SSID of the associated network as parameter "x." |
| No configurations left in the selection list. | Indicates that the Wireless Configuration service went through all of the selected wireless networks and failed to associate to any of them. This could happen either as a result of an NDIS media status indicating "Disconnected" or as a result of the 802.1X forcing the networks to be deleted because of a failed authentication. This description indicates that the Wireless Configuration service is transitioning into a "Failed" state. This is not necessarily an error; it could mean the service is not configured, that no preferred network is available, or that no wireless networks are available at all. |

| | |
|---|---|
| Plumbing configuration SSID: "x," Network Type: "y." | Indicates that the Wireless Configuration service processed the list of available networks and the list of user preferred networks and determined something has changed such that the adapter needs to be reconfigured. This description indicates that a current, successful association is broken. This event occurs when the computer is disconnected from one network or when a new, preferred network became available. This event does not occur in the case of a successful, steady association. The description includes the new SSID that has been configured on the adapter ("x") and the infrastructure parameter ("y;" where 0 = an ad hoc network, and 1 = an infrastructure network). |
| Processing command Next Configuration. | Indicates that 802.1X directed the Wireless Configuration service to break the association to the current network but retain it for subsequent attempts. This event indicates that 802.1X attempted one type of authentication and failed. 802.1X can try authentication if the wireless adapter reassociates with the network. |
| Processing command Remove Configuration. | Indicates that the Wireless Configuration service is directed by 802.1X to break the existing association with the current network and delete it from its list such that it does not refresh that association again. This event indicates that 802.1X failed all attempts for all types of authentication. |
| Processing command Update data. | Indicates that 802.1X is providing new data to the Wireless Configuration service in order to have it associated with the network to which it is currently associated. |
| Processing user command Refresh. | Indicates that the Wireless Configuration service received a command to refresh. Most commonly, this command is a result of the user clicking the **Refresh** in the user interface. This action interrupts the one minute scanning cycle by instantly initiating scanning. |
| Processing user command Reset. | Indicates that the Wireless Configuration service received a command to reset. Most commonly this command is issued by 802.1X whenever something changes in the configuration. This command breaks any existing association while the Wireless Configuration service provides new 802.11 settings. |
| Received Device Arrival notification for "x." | Indicates that the Wireless Configuration service received a "Device arrival" notification. The description includes the NDIS of the device being added as parameter "x." |
| Received Device Removal notification for "x." | Indicates that the Wireless Configuration service received a "Device removal" notification. The description includes the NDIS of the device being removed as parameter "x." |
| Received Media Connect notification. | Indicates that the Wireless Configuration service received a "Media connect" notification. This event can occur repeatedly while the Wireless Configuration service configures the initial association. This description can also indicate that the adapter is randomly switching from one access point to another, which highlights problems with the wireless adapter driver. |
| Received Media Disconnect notification. | Indicates that the Wireless Configuration service received a "Media disconnect" notification from the driver. This event occurs when the Wireless Configuration service provides the adapter with the SSID of a wireless network has been detected, but the adapter fails to associate with that network. This event also occurs whenever the user roams out of the area covered by the network. |

| | |
|---|---|
| Received Timeout notification. | Indicates that an internal Wireless Configuration service timer elapsed. This event description should occur at the following times:<br>• Three seconds after the Wireless Configuration service initiates a scanning cycle.<br>• Two seconds after the Wireless Configuration service provides any network configuration, if there is no "Media connect" notification in between.<br>• Every minute while in "Failed" or "Association successful" states. |
| Scan completed. | Indicates that the Wireless Configuration service has transitioned into scanning cycle after allowing three seconds for the driver of the wireless adapter to scan for wireless networks. Once in this state, the Wireless Configuration service will query the driver in order to get all 802.11 settings, the most important one being the list of available networks. This event occurs every minute. |
| Skipping configuration "x" for now, attempt authentication later. | Indicates that 802.X directed the Wireless Configuration service to skip the current successful association and move to another association. The network configuration is retained for another attempt at a subsequent loop (802.1X will attempt computer authentication instead of user authentication). The description includes the SSID of the network that is being disassociated as parameter "x." |
| Wireless configuration has been changed via an administrative call. | Indicates that a configuration change has been received through an administrative call. This is either the result of the user changing the wireless configuration or a policy object being applied. |
| Wireless Configuration service failed to start. | Indicates that the service failed to initialize. |
| Wireless Configuration service was started successfully. | Indicates that the service has started successfully. This is an indication that the service:<br>• Initialized all of its global objects successfully.<br>• Initialized the logging database framework successfully.<br>• Initialized 802.1X authentication successfully.<br>• Initialized the policy module successfully.<br>• Registered with the Service Control Manager (SCM) successfully, including registering for device notifications.<br>• Registered with Windows Management Instrumentation (WMI) for NDIS media event notifications successfully.<br>• Registered as an RPC server successfully. |
| Wireless interface successfully associated to "x" [MAC "y"]. | Indicates that the Wireless Configuration service associated to a network successfully and that network is sending a "Success" notification to 802.1X. The description includes the SSID of the network to which it is associated as parameter "x." The MAC address of the remote endpoint is also included as parameter "y." |

# Using Infrared Communication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using infrared communication**
You can transfer data over infrared connections with computers, printers, cameras, and other devices that use IrDA protocols. In addition, you can use an infrared-enabled cellular phone to make a dial-up connection to the Internet or to your corporate network. This section covers:

- Security information for infrared communication
- Installing infrared devices
- Establishing infrared network connections between computers
- Transferring files
- Printing to an infrared printer
- Transferring digital camera images to the computer

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Security information for infrared communication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security information for infrared communication**

Infrared provides a secure and simple way to transfer and synchronize data. In order for infrared communication to work properly, infrared transceivers must operate on a line-of-sight basis--they must be placed at a 30-degree angle and no farther than 1 meter (approximately 40 inches) apart. Because infrared operates at such a short distance and a narrow angle, it is relatively difficult for an attacker to intercept data that is sent over infrared. However, infrared does not provide data encryption. Because data is sent in plaintext, it is vulnerable to packet sniffing attacks. Therefore, take precautions to ensure that data that is sent over infrared is not intercepted.

To enhance security when using infrared communication, do the following:

- **Do not enable infrared image transfer (IrTran-P).**

  IrTran-P is disabled by default (that is, the option to use Wireless Link to transfer images from a digital camera to your computer is disabled). If you enable this option, all of the incoming files that are sent over IrTran-P are automatically accepted. Because incoming files might contain harmful programs, ensure that the files originate from a trustworthy source. Do not open files if you cannot verify the source, do not recognize the file format, or are unsure of the content. Instead, delete the files immediately.

  For more information, see Transferring digital camera images to the computer and Infrared implementation.

- **Align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices.**

  Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an outside infrared device. To further enhance security, ensure that all infrared devices and data sources are trustworthy. Finally, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible.

  For more information, see Links between infrared devices and Establish an infrared link.

- **Use only signed device drivers.**

  For security, performance, and stability, it is recommended that you use only signed device drivers.

  For more information, see Signed Drivers.

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Installing infrared devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing infrared devices**
Most internal infrared devices are installed by Setup or when you start your computer after adding one of these devices. However, when you attach a serial infrared transceiver to a serial (COM) port or to a USB port, you do need to install it as described in this topic. This topic also describes how to install an internal infrared device that is not automatically detected.

**Installing an infrared device connected to a serial port**
If you have a desktop computer or a portable computer without an internal infrared device, you can connect a serial infrared transceiver to a serial (COM) port or a USB port. To do so, attach the infrared transceiver to the port, note which port you are using, and then follow the steps in Add a new infrared device to a serial port.

**Installing an undetected internal infrared device**
If you add an internal infrared device to a computer with Windows XP Home Edition, Windows XP Professional, or Windows Server 2003, Standard Edition already installed, the device is typically detected and installed the next time you start the computer. If this does not occur, you might have to install the device manually. To do so, follow the instructions in Add a new infrared device to a serial port.

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Establishing infrared network connections between computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing infrared network connections between computers**
An infrared network connection allows you to establish a direct connection between two infrared-enabled devices without using modems, cables, or network hardware. Instead, you align the two devices to establish an infrared link, and then use Network Connections to create an infrared network connection.

**How you can use infrared network connections**
You can use infrared network connections to perform a variety of tasks. For example, you can:

**Access the Internet from a public location**
While at an airport, a hotel, or any other location that provides public Internet access, you can sit at a kiosk and use your laptop, personal digital assistant (PDA), or other infrared-enabled device to connect to the Internet. Kiosks at such locations provide a built-in infrared port, so that you can align your infrared-enabled device with the built-in port and establish an infrared network connection.

When making this type of infrared network connection, you configure the infrared device as the guest. The server that your device connects to is the host.

**Access shared information on another computer**
You can establish an infrared network connection between two computers to access shared resources on one of the computers. For example, if you have a portable computer that you want to use to access files in a shared folder on your desktop computer, you can establish an infrared network connection between the two computers.

When making this type of infrared network connection, you configure the portable computer as the guest and the desktop computer as the host.

**Creating network connections on both computers**
To create an infrared network connection, follow the steps in Establish a direct infrared network connection between computers on both computers. When you create the network connections, specify the computer that has the information you want to access as the host computer, and the computer you are going to use to access the information as the guest computer.

For both computers, specify the infrared port as the device that will be used to make the connection.

**Configuring permissions on the host computer**
If you are connecting two Windows XP Home Edition, Windows XP Professional, or Windows Server 2003, Standard Edition computers, you can access shared folders on the host. However, before you can make the connection, make sure that the user name for the connection is given membership in a group on the host computer. For information on how to add users to your computer, see Create a local user account. The new user must have sufficient permissions to perform any expected tasks. If you are adding a user so that someone else can access your computer, you might want to add the new user to the Guest, User, or Power Users groups.

**Accessing the local computer**
After you set up network connections on both the host and guest computers, follow the steps in Use an infrared network connection. Working on the guest computer, you can then map shared drives on the host computer and use **My Computer** or Windows Explorer to open folders and access files.

## Notes

- If you only want to send files from one infrared-enabled device to another, you do not need to establish an infrared network connection. Instead, you only need to establish an infrared link, as described in Establish an infrared link.
- Infrared network connections use the Point-to-Point protocol (PPP). Typically, this protocol does not require additional configuration.
- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Transferring files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Transferring files**
You can use Wireless Link to send selected files or an entire folder to another computer running Windows 98, Windows 2000, Windows XP Home Edition, Windows XP Professional, or Windows Server 2003, Standard Edition.

**Establishing an infrared connection**
To send files, first establish an infrared link with another computer or other infrared device, as described in the procedure Establish an infrared link. When an infrared device is in range, the infrared connection status icon appears on the taskbar, and the Wireless Link item appears on the desktop. These indicate only that an infrared device is in range. A data connection is made when Wireless Link or another program communicates over the link. While a program is connected over a link, the taskbar icon changes to the icon. As long as an infrared transceiver is in range, you can transfer data over the link.

The IrDA Winsock API enables programs to communicate with any device that implements one of the supported protocols, described in Infrared implementation. Simultaneous connections over an infrared link are supported (unless the link is being used for a network connection), so more than one program can use the same link. If you use another program, refer to the documentation provided with it for procedures to transfer data.

**Important**

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Security information for infrared communication.

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Printing to an infrared printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing to an infrared printer**

Printing to a printer that is connected by an infrared link is very much the same as printing to any other printer. After you establish an infrared connection between your computer and a printer, the printer is typically installed automatically by Plug and Play. If your computer is not infrared-capable, you can attach an infrared transceiver to your serial (COM) port. For more information, see Add a new infrared device to a serial port.

After you add an infrared printer, the only difference between printing to it and printing to any other printer is that you must first establish an infrared link between the computer and the printer by aligning them so that the infrared transceivers are within 1 meter (approximately 40 inches) of each other, and the transceivers are pointing at each other. For more information, see Establish an infrared link.

**Important**

- For enhanced security, performance, and stability, it is recommended that you use only signed device drivers. For more information, see Signed Drivers.

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Security information for infrared communication.

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Transferring digital camera images to the computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Transferring digital camera images to the computer**

The **Image Transfer** tab of the **Wireless Link** dialog box includes a check box in which you can specify whether you want to transfer images from a digital camera to your computer using Wireless Link.

The **Image Transfer** tab also specifies the default location for image files received from a digital camera. This location is used by Wireless Link. Other programs may use other locations or require you to specify a location when you begin the transfer. You can also elect to have Windows Explorer automatically open to this location after images have been transferred.

The transfer is initiated by the digital camera rather than the computer. To transfer images, point the infrared window on your camera at the infrared transceiver on the computer. Then follow the camera manufacturer's instructions to select the picture and send it to the computer. Some cameras can transfer more than one picture at a time. The exact steps to transfer pictures depends on the make and model of the camera, and should be described in the instructions for operating the camera.

To configure Wireless Link to transfer camera images, see Change settings to transfer images from a digital camera. To transfer camera images to your computer using Wireless Link, see Transfer images from a digital camera to your computer.

**Caution**

- To enhance security, do not enable infrared image transfer (IrTran-P), that is, do not select the **Use Wireless Link to transfer images from a digital camera to your computer** check box. If you select this check box, all incoming files sent over IrTran-P are automatically accepted. Because incoming files might contain harmful programs, ensure that the files originate from a trustworthy source. Do not open files if you cannot verify the source, do not recognize the file format, or are unsure of the content. Instead, delete the files immediately. For more information about infrared security considerations, see Security information for infrared communication.

**Important**

- To enhance security, align infrared devices so that the infrared transceivers are between 0.1 meter (approximately 4 inches) and 0.5 meter (approximately 20 inches) apart whenever you establish an infrared link between two devices. Although there can be up to 1 meter (approximately 40 inches) between infrared transceivers without a loss in data transfer capability, if you place the transceivers closer to each other, you minimize the risk of interference from an attacker. To further enhance security, ensure that all infrared devices and data sources are trustworthy, and, if you are transferring infrared data with another person, conduct the transfer in a private location whenever possible. For more information about infrared security considerations, see Security information for infrared communication.

**Notes**

- If you are using synchronization software, note that certain software will automatically disable infrared image transfer (IrTran-P). As a result, you will not be able to use Wireless Link to transfer images from a digital camera to your computer. In such cases, you might have to remove the synchronization software before you can use Wireless Link to transfer digital camera images. To determine whether your synchronization software disables infrared image transfer, see your software manufacturer's product documentation.

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

# Wireless Networking Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**
This page provides sources of additional information about wireless technologies.

**Microsoft product support**
For Microsoft product support, see Microsoft product Support Services at the Microsoft Web site.

**Windows Resource Kits**
The Windows Resource Kits contain extensive technical information about wireless networking protocols and deploying wireless technologies.

- For information about wireless networking protocols, see "Wireless Networking protocols" at the Microsoft Windows Resource Kits Web site.

- For information about deploying wireless technologies, see "Deploying a Wireless LAN" at the Microsoft Windows Resource Kits Web site.

**The Infrared Data Association (IrDA)**
The Infrared Data Association (IrDA) is an international organization that sets standards for infrared data connections. The infrared support provided by Windows XP Home Edition, Windows XP Professional, and Windows Server 2003, Standard Edition is designed to meet these standards.

For more information about IrDA, see the Infrared Data Association Web site.

### Note

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Windows Server 2003 Wireless Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This topic provides basic troubleshooting information for wireless computers connecting to Windows Server 2003 with Service Pack 1 (SP1) domain wireless networks. If you are looking for wireless troubleshooting for a Windows XP home office environment, see the following:

- Troubleshooting Microsoft Windows XP-based Wireless Networks in the Small Office or Home Office at http://go.microsoft.com/fwlink/?LinkID=55033
- How to troubleshoot wireless network connections in Windows XP at http://go.microsoft.com/fwlink/?LinkId=83226
- Troubleshooting IEEE 802.11 Wireless Access with Microsoft Windows at http://go.microsoft.com/fwlink/?LinkId=83227
- Configuring Windows XP IEEE 802.11 Wireless Networks for the Home and Small Business at http://go.microsoft.com/fwlink/?LinkID=30045

**Quick lists for common connectivity problems**

This section provides a series of tables and lists that can help you to quickly identify conditions that can cause connectivity problems. The quick lists are presented in two categories: by symptom and by network type.

Quick lists by symptom

- Symptom: Inability to connect
- Symptom: Intermittent connectivity
- Symptom: Incorrect, missing, or stale visible networks
- Symptom: Wireless client has associated, but there is no valid IP address configuration or no network connectivity
- Symptom: There are no visible wireless networks or the Wireless Networks tab is not present in the Network Connections folder

Quick lists by network type

- General network connectivity problems
- Domain network connectivity problems
- 802.1X-authenticated network connectivity problems

**Quick lists by symptom**

The following series of tables present common symptoms, their causes, and likely solutions.

**Symptom: Inability to connect**

| Possible Causes | Corrective Measures |
|---|---|

| | |
|---|---|
| • Incorrect or incompatible wireless network configuration. For example, shared key authentication is configured on the wireless AP, and the wireless client is attempting open system authentication.<br>• Inadvertent media access control (MAC) address filtering.<br>• The wireless network name is not visible.<br>• The wireless AP and wireless network adapter are not using the same 802.11 standard (for example, you are using an 802.11b network adapter and an 802.11a wireless AP).<br>• | • Verify that the wireless network configurations between the wireless client and wireless AP are compatible.<br>• Double-check the steps you followed during configuration. User error is a common source of incorrect configuration.<br>• |
| • Radio frequency (RF) interference from nearby devices, such as cordless phone and Bluetooth devices.<br>• Wireless client is at the periphery of the RF range of the wireless AP. | • Review the wireless network environment and network topology.<br>• Double-check the steps you followed during configuration. User error is a common source of incorrect configuration. |
| • Improperly functioning or outdated wireless network adapter driver. | • Obtain and install the most recent version of the wireless network adapter driver. |

**Symptom: Intermittent connectivity**

| Possible Causes | Corrective Measures |
|---|---|
| • Improperly functioning or outdated wireless network adapter driver.<br>• Improperly functioning wireless AP. | • Obtain and install the most recent version of the wireless network adapter driver. |

**Symptom: Incorrect, missing, or stale visible networks**

| Possible Causes | Corrective Measures |
|---|---|
| • Improperly functioning or outdated wireless network adapter driver.<br>• Improperly functioning radio equipment on wireless AP or wireless network adapter.<br>• Malfunctioning wireless network adapter drivers cannot detect and register visible networks. | • Obtain and install the most recent version of the wireless network adapter driver.<br>• Run the Repair tool on the wireless network adapter.<br>**Note**<br>To run Repair, right-click the wireless connection icon in the notification area or in Network Connections, and the click **Repair**. |

**Symptom: Wireless client has associated, but no there is no valid IP address configuration or no network connectivity**

| Possible Causes | Corrective Measures |
|---|---|

| | |
|---|---|
| • Authentication problem.<br>• Incorrect encryption key.<br>• Corrupt, expired, or missing certificates.<br>• Improperly functioning wireless AP. | • Verify that the wireless network configurations between the wireless client and wireless AP are compatible.<br>• If you are using a static wired equivalent privacy (WEP) key, verify that it has been correctly configured.<br>**Note**<br>Due to known security issues with WEP encryption, it is recommended that you use only Wi-Fi Protected Access version 2 (WPA2) (preferred) or WPA.<br>• Verify whether other computers that connect to the wireless AP have the same problem. If all wireless clients of the same wireless AP have the same problem, make sure the wireless AP settings match the **IAS RADIUS Clients** settings configured for the specific wireless AP. Also make sure that the **Wireless Network (IEEE 802.11) Policy** configuration settings for client computers match the wireless AP settings. |

**Symptom: There are no visible wireless networks or the Wireless Networks tab is not present in the Network Connections folder**

| Possible Causes | Corrective Measures |
|---|---|
| • The Wireless ZeroConfig Service is not running.<br>• Improperly functioning or outdated wireless network adapter driver.<br>• On a portable computer, the external switch for the wireless antenna might be turned off. | • Check to see if the Wireless ZeroConfig Service is running by checking the state is **Started** in the Services Microsoft Management Console (MMC) snap-in.<br>• Use the Services console to configure the Wireless ZeroConfig Service to start automatically.<br>• A wireless network adapter driver that fails in the early stages of service startup can cause the Wireless ZeroConfig Service not to initialize on that interface. If this is determined to be the problem, install the most recent version of the driver for that adapter. |

**Quick lists by network type**

The following quick lists are not exhaustive catalogs of connectivity problems. They provide information about the types of conditions that can cause connectivity problems.

For the purposes of this document, network connectivity problems fall into three groups:

- General network connectivity problems
- Domain network connectivity problems
- 802.1X-authenticated network connectivity problems

**General network connectivity problems**

These types of problems can occur on networks ranging from small office/home office (SOHO) workgroup-based networks to enterprise networks.

| Possible Causes | Corrective Measures |
|---|---|

| | |
|---|---|
| A wireless setting mismatch exists between the wireless AP and the wireless client.<br><br>For example, the wireless AP is configured to use WPA2 and the client is configured with WPA-PSK. | For a single computer, manually configure the settings to match.<br><br>For multiple computers, use the **Wireless Network (IEEE 802.11) Policy** to configure settings to match the settings configured on the wireless AP. |
| The wireless adapter is disabled in Network Connections. | Right-click the disabled connection, and then click **Enable**. |
| The external switch that controls the wireless antenna is turned off. | Set the external switch to the On position. |
| The wireless network adapter is malfunctioning. | Install the most recent driver for the adapter. If that fails to fix the problem, replace the wireless adapter, if possible.<br><br>**Note**<br>In cases where the adapter is built into the computer, you can try to install a second wireless adapter, such as a Universal Serial Bus (USB) adapter to test connectivity. If you install a second wireless adapter, it is recommended that you first disable the original wireless connection. In Network Connections, right-click the wireless connection, and then click **Disable**. Install the second wireless adapter, and then attempt to connect to the network. |
| Network clients configured with static IP addresses are configured by using an IP address that is not in the correct IP address range or by using a different subnet mask. | Configure the IP settings with a unique IP address in the correct IP address range, using the correct subnet mask. |
| The DHCP service is enabled on the wireless router to allocate IP addresses to network clients, but one or more network clients are configured with a static IP address. | Configure the **Internet Protocol (TCP/IP)** properties of the wireless adapter to **Obtain an IP address automatically**. |
| The DHCP server is disconnected from the network, powered off, or the service is not running. | In a SOHO network, the DHCP service is typically provided by the wireless router or by Internet Connection Sharing (ICS). Restore the DHCP service. |
| In a SOHO network:<br><br>• In a new wireless network or when replacing your modem or wireless AP, you have not registered your modem with your Internet service provider (ISP), or your router media access control (MAC) address. Modem or router registration varies by ISP.<br><br>• You have not configured the public connection on the router to accept DHCP leases from the ISP network. For example, you have configured the public connection on the wireless router with a static IP address. | • Contact your ISP to register your device.<br><br>• Configure the public interface to accept DHCP leases from the ISP network. |

**Domain network connectivity problems**

In addition to the general network connectivity problems, these types of problems commonly occur on domain networks, ranging from small organizations to enterprise networks.

**Active Directory**

| Possible Causes | Corrective Measures |
|---|---|
| The user does not have an account in the **Active Directory Users and Computers** snap-in. | Create an account for the user. |
| The **Dial-in** properties of the user account or computer account in **Active Directory Users and Computers** is set to **Deny access**. | Set the user and computer account **Dial-in** properties to **Allow access**.<br>**Note**<br>In 802.1X networks that use IAS to authenticate connections, set the user account **Dial-in** properties to **Control access through Remote Access Policy**. |
| The user account is disabled. | In **Active Directory Users and Computers**, in **Users**, right-click the account, and then click **Enable**. |
| The user account has expired. | In **Active Directory Users and Computers**, right-click the account, click **Properties**, and then on the **Account** tab, in **Account expires**, select **Never**, or in **End of** set a new expiration date. |
| The user is attempting a connection at a prohibited time, as specified in the logon hours of the user account (the default setting is **Logon Permitted** for all hours). | In **Active Directory Users and Computers**, in the user account properties, on the **Account** tab, click **Logon Hours**, and then configure the settings to specify the hours that the client is allowed to connect to the network. |
| The user is attempting a prohibited connection by using a computer that is not specified in the **Log On To** setting of the user account properties, or the default setting **All computers** is not selected. | In **Active Directory Users and Computers**, in the user account properties, on the **Account** tab, click **Log On To**, and then either select **All computers**, or select **The following computers**. In **Computer name**, specify the computers that the user is allowed to use to connect to the network. |
| The Domain Name System (DNS) service is stopped or is not configured. | On your DNS server, in the **Services** snap-in, right click **DNS Server**, and then click **Start**. |

**Users and Computers**

| Possible Causes | Corrective Measures |
|---|---|
| The client computer is not a member of the domain. | Join the computer to the domain. |
| The client is attempting to log on to the domain by using non-domain credentials. A common error with new computers or user accounts is that users log on to the computer by using their computer account. | At the log on window, in **Log on to**, select the domain if it is available, and then use the domain user account. For newly joined computers, logon using the domain name and user account in the format *DomainName\UserName*. |

**DHCP**

| Possible Causes | Corrective Measures |
|---|---|
| The DHCP scope is full, and therefore the DHCP server cannot lease addresses to requesting DHCP clients. | If the DHCP scope does not use the full address range, edit the scope to expand the address range. |
| The IP address of the DHCP server was changed and now DHCP clients cannot get IP addresses. | Make sure that the static IP address and subnet mask of the DHCP server are within the addressing scheme of the subnet. |
| The DHCP service is stopped. | On your DHCP server, in the **Services** snap-in, right-click **DHCP Server**, and then click **Start**. |

| | |
|---|---|
| On a newly configured DHCP server:<br>• The DHCP server is not authorized in Active Directory.<br>• The IP address range of the DHCP scope is incorrectly specified.<br>• The DHCP scope is not activated.<br>• The DHCP server is not on the same subnet as the clients.<br>• The DHCP server is offline. | In the **DHCP** snap-in, right-click the domain container, and then click **Authorize**.<br>Set the IP address range and subnet mask in the scope to match the addressing scheme of your subnet.<br>In the **DHCP** snap-in, right-click the domain container, right-click the scope for the subnet that is presenting connectivity problems, and then click **Activate**.<br>Physically connect the DHCP server to the correct subnet.<br>Restart the DHCP server. |

**802.1X-authenticated network connectivity problems**

This section provides examples of configuration problems that are specific to networks that deploy 802.1X-authenticating wireless APs and IAS for 802.1X-authenticated connections. In an 802.1X network, consider the following examples in addition to the examples listed in the previous two sections.

**Active Directory problems**

| Possible Causes | Corrective Measures |
|---|---|
| The Active Directory domain functional level is not raised to **Windows Server 2003**. IAS Remote Authentication Dial-In User Service (RADIUS) settings require the Windows Server 2003 domain functional level. | In the **Active Directory Domains and Trusts** snap-in, click **Action**, and then click **Raise Domain Functional Level**.<br>**Important**<br>If domain controllers on your network are running Windows NT 4.0 and earlier, then do not raise the domain functional level to Windows 2000 native. After the domain functional level is set to Windows 2000 native, it cannot be changed back to Windows 2000 mixed.<br>Likewise, if domain controllers on your network are running Windows 2000 or Windows NT 4.0 and earlier, then do not raise the domain functional level to Windows Server 2003. After the domain functional level is set to Windows Server 2003, it cannot be changed back to Windows 2000 mixed or Windows 2000 native. |
| In **Active Directory Users and Computers**, the dial-in properties of the user account are not configured to **Control access through Remote Access Policy**. | In the **Active Directory Users and Computers** snap-in, in the domain container, open **Users**, right-click the user account, click **Properties**, and then on the **Dial-in tab**, select **Control access through Remote Access Policy**. |
| The IAS remote access policy grants access for members of an Active Directory security group. However, the user is not a member of the security group that is specified in the remote access policy. | In **Active Directory Users and Computers**, in the domain container, open **Users**, right-click the security group that is specified in the applicable IAS remote access policy, click **Properties**. On the **Members** tab, click **Add**, and then in **Enter the object names to select**, type the user account to add the user to the security group. |

| | |
|---|---|
| The authentication method specified in **Wireless Network (IEEE 802.11) Policies** does not match the authentication method specified in the IAS remote access policy.<br><br>For example, if network clients running Windows Vista are configured by **Wireless Network (IEEE 802.11) Policies** to use Protected Extensible Authenticated Protocol Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP-MS-CHAPv2) authentication, but no IAS remote access policy specifies PEAP-MS-CHAPv2 authentication, the mismatch prevents client authentication. | Change the Extensible Authenticated Protocol (EAP) authentication method specified in either the IAS remote access policy or in **Wireless Network (IEEE 802.11) Policies** to match the method deployed for your network. |

**Client**

| Possible Causes | Corrective Measures |
|---|---|
| The Wireless ZeroConfig Service is not running. By default, the Wireless ZeroConfig Service starts automatically. | For an individual computer, manually start the service in the **Services** snap-in.<br><br>For multiple computers, in Group Policy Object Editor, open **Wireless Network (IEEE 802.11) Policies**. In the details pane, right-click the applicable wireless policy, click **Properties**, and then on the **General** tab, select **Use Windows to configure wireless network settings for clients**. |
| When using Extensible Authenticated Protocol-Transport Layer Security (EAP-TLS) authentication, the client does not have a certificate that contains the Client Authentication purpose in the Enhanced Key Usage extension. Instead, the client is configured according to minimum client certificate requirements. | Instruct the user to log on to the network by using a wired Ethernet connection and domain credentials; the enhanced certificate is automatically installed. |

**Certificate Services**

| Possible Causes | Corrective Measures |
|---|---|
| When using EAP-TLS, the user does not have a client certificate. | Instruct the user to log on to the network by using a wired Ethernet connection and domain credentials; the certificate is automatically installed. |
| The client does not have a corresponding root CA certificate that matches the issuing CA of the IAS server certificate. | 1. In Group Policy Object Editor, open **Wireless Network (IEEE 802.11) Policies**.<br><br>2. In the details pane, right-click the applicable wireless policy, and then click **Properties**.<br><br>3. On the **Preferred Networks** tab, in **Networks**, select the corresponding wireless network, and then click **Edit**.<br><br>4. On the **802.1x** tab, in **EAP** type, click **Settings**, and then in **Trusted Root Certification Authorities**, select the certificate that matches the IAS server certificate. |

**IAS (RADIUS)**

| Possible Causes | Corrective Measures |
|---|---|
| The RADIUS shared secret on the wireless AP does not match the shared secret configured for RADIUS clients in IAS. | Configure the wireless AP to use the same shared secret. The shared secret is specified in the **RADIUS Clients** node of the **IAS** snap-in. |

| | |
|---|---|
| The IAS remote access policy properties are configured to reject the user or computer requests. For example:<br>• On the **Settings** tab, the properties of the policy are set to **Deny remote access permission**.<br>• On the **Dial-in Constraints** tab of the remote access policy, time restrictions prohibiting the connection are configured.<br>• On the **Dial-in Constraints** tab, an incorrect media type is specified. | • In the **IAS** snap-in, right-click the applicable remote access policy, click **Properties**, and then click **Grant remote access permission**.<br>• In the **IAS** snap-in, right-click the applicable remote access policy, click **Properties**, click **Edit Profile**. On the **Dial-in Constraints** tab, select **Allow access only on these days and at these times**, and then click **Edit** to specify allowed access times.<br>• In the **IAS** snap-in, right-click the applicable remote access policy, click **Properties**, click **Edit Profile**, and then on the **Dial-in Constraints** tab, do one of the following:<br>  • Clear **Allow access only through these media (NAS-Port-Type)**<br>  • Select **Allow access only through these media (NAS-Port-Type)**, and then select **Wireless - IEEE 802.11**. |
| A mismatch exists between the trusted root CA that issued the RADIUS server certificate specified in the IAS remote access policy, and the trusted root CA specified in the selected EAP type in **Wireless Network (IEEE 802.11) Policies**. | 1. In Group Policy Object Editor, open **Wireless Network (IEEE 802.11) Policies**.<br>2. In the details pane, right-click the applicable wireless policy, and then click **Properties**.<br>3. On the **Preferred Networks** tab, in **Networks**, select the corresponding wireless network, and then click **Edit**.<br>4. On the **802.1x** tab, in **EAP type**, click **Settings**, and then in **Trusted Root Certification Authorities**, select the certificate that matches the IAS server certificate. |
| The vendor-specific attributes (VSAs) for the wireless AP are configured incorrectly. | Check the wireless AP product documentation for VSA usage, and then specify the **IAS Vendor specific attributes** in **IAS**. If you are unsure about the correct VSA setting, select **RADIUS Standard**. |
| The IP address of the wireless AP (RADIUS client) specified in IAS is incorrect. | 1. In the **IAS** snap-in, select **RADIUS Clients**.<br>2. In the details pane, right-click the RADIUS client that corresponds to the wireless AP, and then click **Properties**.<br>3. On the **settings** tab, in **Address (IP or DNS)** type the correct IP address, and then click **Verify**. |
| The IAS server certificate has expired. | For information about requesting an IAS Server certificate, see the Windows Server 2003 Help topic Computer certificates for certificate-based authentication. |
| The IAS service is stopped. | In the **Services** snap-in, right-click **Internet Authentication Service**, and then click **Start**. |

| | |
|---|---|
| EAP is configured differently in the applicable remote access policy than it is in **Wired Network (IEEE 802.11) Policy** in Active Directory. | Configure both the IAS remote access policy and **Wired Network (IEEE 802.11) Policy** to use the EAP method that corresponds with your network deployment. |
| On a newly configured IAS server:<br>• IAS is not registered in Active Directory.<br>• The IAS server does not have a server certificate. | • In the **IAS** snap-in, right-click **Internet Authentication Service**, and then click **Register Server in Active Directory**.<br>• See the Windows Server 2003 Help topic Computer certificates for certificate-based authentication for information about requesting an IAS Server certificate. |

**Wireless AP**

| Possible Causes | Corrective Measures |
|---|---|
| The wireless AP does not have the correct or latest firmware. | Contact the wireless AP manufacturer for the latest firmware. |
| The IP address of the wireless AP is incorrectly configured for the subnet. | Configure the wireless AP with a static IP address and subnet mask according to your network TCP/IP addressing scheme. |
| The wireless AP does not specify the correct IP address of the IAS RADIUS server. | On the wireless AP, configure the RADIUS server IP address to match the IP address of your IAS server. |

**Troubleshooting Windows XP wireless connections: quick list**

Wireless connection problems frequently occur for the following reasons:

- Disabled wireless network adapters.
- Incorrectly configured wireless network settings.
- Insufficient credentials for authentication or missing permissions for authorization.
- Distance, interference, or obstructions between wireless devices.

Because successful troubleshooting depends on your ability to identify the source of the problem, Windows Server 2003 with Service Pack 1 (SP1) includes status and warning reporting to help you isolate and resolve wireless network connection problems.

Follow these steps to solve several common problems associated with wireless connections:

- Many portable computers have a switch that can be used to turn 802.11 wireless network adapters on and off. On a laptop computer, the switch might be on the left side of the computer case. Be sure that the switch is turned on. For more information, see the product documentation for your portable computing device.
- Make sure that the wireless adapter has not been disabled. You can disable a wireless adapter through the user interface (UI) by right-clicking on a wireless adapter icon, and then selecting **Disable**. Wireless adapters that have been disabled in this way will not appear in the notification area and can only be enabled in Network Connections. For more information, see Configuring wireless network settings on client computers.
- Use Wireless Auto Configuration to configure wireless network settings. When enabled, Wireless Auto Configuration allows you to connect to an existing wireless network, change wireless network connection settings, configure a connection to a new wireless network, and specify preferred wireless networks. It also notifies you when new wireless networks are available. When you switch wireless networks, your wireless network adapter settings will be dynamically updated to match the settings of that new network and a network connection attempt will be made. For more information, see Use Windows to configure wireless network settings on a client computer.
- If you are connecting to a wireless network for the first time, Wireless Auto Configuration will configure basic network

settings, if the service is enabled. However, you may need to configure additional settings, such as the data encryption type or Wired Equivalent Privacy (WEP) key, if they are not automatically configured for your account in the Active Directory directory service. For more information, see Define Wireless Network Policies on a Client Computer. You might also need to request account permissions from your network administrator.

- Check to see if the desired wireless network appears in the network list. Right-click the wireless icon, and then click **View Available Wireless Networks**. If the desired wireless network does not appear under **Choose a wireless network**, you might be outside of the broadcast range of that network or the network might be suppressing the beaconing signal. First, try to relocate the wireless device to a location that receives a stronger signal. To refresh the network list and get the most current list of wireless networks that are advertising within reception range of your computer, right-click the wireless icon, click **View available wireless networks**, and then, under **Network Tasks**, click **Refresh network list**.

  **Note**

  - Some infrastructure networks suppress the beaconing signal because they do not want to advertise the availability of their wireless network. In these cases, the network will not appear under **Choose a wireless network**. However, you can connect to these networks if you enter all of the correct configuration information, which you can obtain from the network administrator.

    For more information about configuring wireless network settings, see Use Windows to configure wireless network settings on a client computer, Add, edit, or remove wireless network connections on a client computer, and Define a wireless network connection on a client computer.

- Check to see if there is a wireless warning icon in the notification area. You can click the warning icon to get information about the error as well as possible remedies. If you used **View Available Wireless Network** to open the list of available wireless networks, check for a warning where the wireless network is displayed under **Choose a wireless network**. You can click the warning link text to get information about the warning and possible remedies.

- If you have previously connected successfully to a specific network, but the connection to that network failed, right-click the wireless icon, and then click **Repair**. This will disable, and then enable the wireless adapter.

- Balloon notifications appear in the notification area when the status of a wireless connection changes (for example, when status changes from **Connected** to **Not Connected** because the wireless device has moved out of reception range). If a balloon notification indicates an error, you can click the notification for more information and remedies to correct the problem.

**Status and warning reporting**

Windows Server 2003 SP1 monitors wireless connection status to detect connection errors. Status and connection error information is reported in real time in the following ways:

**Notification area icons**

The wireless icon in the notification area changes appearance to indicate the state of the wireless connection:

| Icon | Wireless connection state |
|---|---|
| | Connected |
| | Connecting |
| | Not connected |
| | Warning |

If the warning icon is displayed, click the icon, and then, on the **General** tab of the **Wireless Network Connection Status** dialog box, under **Connection**, click **More information**.

**Hovering pointer display**

You can hover the mouse pointer over the wireless icon in the notification area and instantly see the connection status and other information:

| Status type | Reported status |
| --- | --- |
| Wireless Network Connection | The name of the network to which the computer is currently connected. This is the same as the service set identifier (SSID). |
| Speed | Connection speed (for example, 11.0 Mbps) |
| Signal Strength | Excellent<br><br>Very good<br><br>Poor |
| Status | Connected<br><br>Limited or no connection<br><br>Not connected<br><br>**Note**<br><br>In the event of a service interruption, the status will display the name of the most recent wireless network connection, but it will indicate that the connection has limited or no connectivity. |

**Balloon notifications**

Balloon notifications will appear in the notification area when the status of a wireless connection changes, regardless of whether the wireless icon has been configured to display in the notification area.

**Example notifications**

When Windows has detected an error, you might receive the following notification:

**The connection has limited or no connectivity. You may not be able to access the Internet or some network resources. For more information, click this message.**

To open the **Wireless Network Connection Status** dialog box, click the balloon. On the **Support** tab, under **Connection Status**, a message similar to the following will be displayed:

**Limited or no connectivity. You might not be able to access the Internet or some network resources. This problem occurred because the network did not assign a network address to the computer.**

**To restart the connection and try to establish full connectivity, click Repair.**

After you click **Repair**, the wireless adapter is restarted, and shortly afterward, you will receive a notification similar to the following:

**Wireless Network Connection is now connected**

**Connected to: Gulliver**

**Signal Strength: Very Good**

The balloon notifications are displayed for a short time only. For this reason, there are alternative ways to access the **Support** tab in the **Wireless Connection Status** dialog box.

**Choose a wireless connection**

Wireless network connection status is also reported in **Choose a wireless connection** in the **Wireless Network Connection** dialog box. To control the connection to wireless networks that appear under **Choose a wireless network**, click **Connect** or **Disconnect**. If you have attempted to connect to a wireless network and the connection attempt failed, a status message such as **Limited or no connectivity** might be displayed for that item. To open the **Wireless Network Connection Status** dialog box, click the message text. On the **Support** tab, under **Connection Status**, a message similar to the following will be displayed:

**Limited or no connectivity. You might not be able to access the Internet or some network resources. This problem occurred because the network did not assign a network address to the computer.**

**To restart the connection and try to establish full connectivity, click Repair.**

**Additional troubleshooting tools**

You can use the following tools to troubleshoot 802.11 wireless network connections and infrared devices:

- **802.11 wireless networks**. To troubleshoot problems related to 802.11 wireless network connections, you can view details about wireless access points and wireless clients in Wireless Monitor. For detailed information about the statistics displayed in Wireless Monitor, see Logging and viewing wireless network activity.

- **Infrared devices**. To troubleshoot problems related to infrared devices, you can use the Input Devices Troubleshooter.

**Note**

- In the Windows Server 2003 family, only Windows Server 2003, Standard Edition, supports infrared networking.

**Additional references**

- A Support Guide for Wireless Diagnostics and Troubleshooting at http://go.microsoft.com/fwlink/?LinkId=83229

- Microsoft Wireless Networking home page at http://go.microsoft.com/fwlink/?LinkId=83230

- Windows Vista Wireless Networking Evaluation Guide at http://go.microsoft.com/fwlink/?LinkId=79943

- Step-by-Step Guide for Setting Up Secure Wireless Access in a Test Lab at http://go.microsoft.com/fwlink/?Linkid=28117

# HyperTerminal

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**HyperTerminal**
HyperTerminal is a program that you can use to connect to other computers, Telnet sites, bulletin board systems (BBSs), online services, and host computers, using either a modem, a null modem cable, or a TCP/IP (Winsock) connection:

- For help with specific tasks, see HyperTerminal How To....
- For general background information, see HyperTerminal Concepts.
- For problem-solving instructions, see HyperTerminal Troubleshooting.

# HyperTerminal How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install HyperTerminal
- Set up a new connection
- Change redial settings
- Change the port settings of the modem
- Change the settings for making a call
- Use a saved connection to call a remote computer
- Send a file to a remote computer
- Wait for or stop waiting for a call using HyperTerminal
- Receive a file from a remote computer
- Save incoming data to a text file
- Use HyperTerminal with Minitel
- View East Asian languages with the VT-UTF8 terminal emulator in Hyperterminal
- Change the backscroll buffer size
- Change the number of columns
- Stop using HyperTerminal as your telnet client

# Install HyperTerminal

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install HyperTerminal**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove** Windows **Components**.
3. Click **Accessories and Utilities**, then click **Details**.
4. Click **Communications**, then click **Details**.
5. Select the **HyperTerminal** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the port settings of the modem
Set up a new connection

# Set up a new connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up a new connection**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open HyperTerminal.
2. If the **Location Information** dialog box is displayed, complete the information for your location and click **OK**.
3. On the **File** menu, click **New Connection**.
4. In the **Name** box, type a name that describes the connection.
5. In the **Icon** box, click the appropriate icon, and then click **OK**.
6. In the **Connect To** dialog box, choose which port or modem you want to use in the **Connect using** drop-down box.
7. If you are connecting with a modem, enter the information for the phone number you want to dial. If you are connecting via TCP/IP (Winsock), enter the host address and port number, and then click **OK**.
8. If you are using a modem and want to change your dialing properties, click **Dialing Properties**, choose the appropriate location, and click **Edit**.
9. If you are using a modem and want to change your location, click **Dialing Properties**, and click **New**.
10. If the **Port Settings** dialog box is displayed, complete the information, and click **OK**.
11. If you are using a modem, to dial the call, click **Dial**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- You can only change the dialing properties if you have chosen a modem in the **Connect using** drop-down box.
- If you save this connection, it will be saved using the .ht extension. By default, the connections are saved in the Documents and Settings\user name\Start Menu\Programs\Accessories\Communications\HyperTerminal folder.

**Using a command line**

1. Open Command Prompt.

2. On the volume where the operating system is installed, type: **cd Program Files\Windows NT** and then press ENTER.

3. Type:

   **hypertrm.exe/t***telnet address*[ **:***port*]

| Value | Description |
|---|---|
| telnet address | The address or name of the connection you want to create. This can be in standard Internet dotted notation (for example, w.x.y.z) or can be the site's user-friendly name. |
| port | The number of the port that you want the connection to use. Port 23 is the default. |

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- Using the command line to create a new connection creates a TCP/IP (Winsock) connection.
- If you save this connection, it will be saved using the .ht extension. By default the connections are saved in the Documents and Settings\user name\Start Menu\Programs\Accessories\Communications\HyperTerminal folder.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Change the settings for making a call
Wait for or stop waiting for a call using HyperTerminal

# Change redial settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change redial settings**

1. Open HyperTerminal.
2. Open a saved connection file or create a new connection. However, do not dial the connection. Instead click **Cancel** when the **Connect** dialog box is displayed.
3. On the **File** menu, click **Properties** and click the **Connect To** tab.
4. Select or clear the **Redial on busy** check box depending on what you want to do.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Set up a new connection
Use a saved connection to call a remote computer

# Change the port settings of the modem

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the port settings of the modem**

1. Open HyperTerminal.
2. Open a saved connection file or create a new connection.
3. If the **Connect To** dialog box does not display, on the **File** menu, click **Properties**.
4. On the **Connect To** dialog box (or the **Connect To** tab on the **Properties** dialog box) click **Configure**.
5. Make your changes on the **General** and **Advanced** tabs. Click **OK**, and click **OK** again.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- This procedure applies only to connections using a modem or COM port.
- The changes that you make to the port settings affect only this connection.
- The port settings for your modem are automatically set when you install the modem, and you typically do not need to change these settings. However, you can modify port settings for a connection with special requirements. For example, a server you connect to might require that you disable data compression or reduce the port speed. To get Help for an item in the Connection Preference dialog box, right-click the item, and then click **What's This?**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Change the settings for making a call
Set up a new connection

# Change the settings for making a call

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the settings for making a call**

1. Open HyperTerminal.
2. Create a new connection, or click **Cancel** when the **Connection Description** dialog box is displayed.
3. If you did not create a new connection, open a saved connection.
4. If the **Connect To** dialog box is displayed, make the appropriate changes and click **OK**.
5. To make additional changes, on the **File** menu, click **Properties**.
6. Click the **Connect To** and **Settings** tabs to see the available choices, and make the appropriate changes.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Change the port settings of the modem
Set up a new connection

# Use a saved connection to call a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use a saved connection to call a remote computer**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open HyperTerminal.
2. If the **Connection Description** dialog box is displayed, click **Cancel**.
3. On the **File** menu, click **Open**.
4. In the **File name** box, type or select the name of the connection you want to use and click **Open**.
5. If the **Connect To** dialog box is displayed, type any needed information, such as the Phone number, and click **OK**.
6. Click **Dial**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- HyperTerminal connections are saved using the .ht extension. By default the connections are saved in the Documents and Settings\user name\Start Menu\Programs\Accessories\Communications\HyperTerminal folder.

**Using a command line**

1. Open Command Prompt.
2. On the volume where the operating system is installed, type: **cd Program Files\Windows NT** and then press ENTER.
3. Type:

   **hypertrm.exe** [**/d**] *connection name*

4. If the **Connect To** dialog box is displayed, type any needed information, such as the Phone number, and click **OK**.
5. Click **Dial**.

| Value | Description |
|---|---|

| | |
|---|---|
| connection name | The fully qualified path and name of a previously saved connection. HyperTerminal connections are saved using the .ht extension. By default the connections are saved in the Documents and Settings\user name\Start Menu\Programs\Accessories\Communications\HyperTerminal folder. |

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Send a file to a remote computer
Receive a file from a remote computer
Wait for or stop waiting for a call using HyperTerminal

# Send a file to a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send a file to a remote computer**

1. Open HyperTerminal.
2. Open a saved connection file or create a new connection.
3. Connect to the remote computer.
4. On the **Transfer** menu, click **Send File**.
5. In the **Filename** box, type the path and name of the file you want to send.
6. In the **Protocol** list, click the protocol your computer is using to send the file.
7. Click **Send**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- You must have an active HyperTerminal session connected prior to performing this procedure.
- Both the sending computer and the receiving computer must be using the same file transfer protocol.
- In most cases, you need to configure the file-transfer software on the remote computer so that it receives the file correctly. For more information, contact the administrator of the remote computer.
- If you use the Zmodem protocol to transfer data, the remote computer will receive the file automatically and will not need to perform a manual receive procedure.
- Be aware that the automatic download feature of the Zmodem protocol can pose a security risk by allowing remote users to send files to your computer without your explicit permission. To avoid this risk, you must select a protocol other than Zmodem or clear the **Allow remote host initiated file transfers** check box on the **Settings** tab of **Connection Properties**.
- Incoming files might contain viruses or malicious programs that could compromise or destroy data on your computer. To mitigate this risk, use virus scanning software and ensure that the file is coming from a reliable and trusted source.
- You can save the protocol that is the recommended default by following these same procedures, but click **Close** instead of **Send** in the last step.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Set up a new connection
Use a saved connection to call a remote computer
Receive a file from a remote computer

# Wait for or stop waiting for a call using HyperTerminal

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To wait for or stop waiting for a call using HyperTerminal**

1. Open HyperTerminal.
2. If the **Connection Description** dialog box is displayed, click **Cancel**.
3. On the **Call** menu, click **Wait for a Call** or **Stop Waiting**, depending on what you want to do.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- HyperTerminal does not automatically host calls when unattended, so you will need to manually interact with the user who is calling you.
- Waiting for a call allows others to connect to your computer and upload files. Take the proper precautions to ensure that only authorized individuals connect to your computer and that you approve of any uploads.
- When you disconnect from a call, you must always follow this procedure again to wait for another call.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Receive a file from a remote computer
Use a saved connection to call a remote computer

# Receive a file from a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To receive a file from a remote computer**

1. Open HyperTerminal.
2. Connect to the remote computer using either a new connection or using a saved connection. Alternatively, you can wait for a call and have the remote computer connect to your computer.
3. Use the remote computer's modem file transfer software to send the file to your computer.
4. On the **Transfer** menu on your computer, click **Receive File**.
5. In the **Place received file in the following folder** box, type the path of the folder in which you want to save the downloaded file.
6. In the **Use receiving protocol** list, click the protocol that the remote computer is using to send the file, and then click **Receive**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- You must have an active HyperTerminal session connected prior to performing this procedure.
- If the remote computer is transferring data using the Zmodem protocol and you have allowed remote host initiated file transfers, you do not need to perform this procedure because the Zmodem protocol causes HyperTerminal to receive the file automatically.

  Be aware that the automatic download feature of the Zmodem protocol can pose a security risk by allowing remote users to send files to your computer without your explicit permission. To avoid this risk, you must select a protocol other than Zmodem in the **Receive File** dialog box or clear the **Allow remote host initiated file transfers** check box on the **Settings** tab of **Connection Properties**.
- Incoming files might contain viruses or malicious programs that could compromise or destroy data on your computer. To mitigate this risk, use virus scanning software and ensure that the file is coming from a reliable and trusted source.
- You can save the receiving folder and the protocol that are the recommended defaults by following these same procedures, but click **Close** instead of **Receive** in the last step.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Install HyperTerminal
Send a file to a remote computer
Use a saved connection to call a remote computer
Wait for or stop waiting for a call using HyperTerminal

# Save incoming data to a text file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To save incoming data to a text file**

1. Open HyperTerminal.
2. Connect to the remote computer using either a new connection or using an saved connection. Alternatively, you can wait for a call and have the remote computer connect to your computer.
3. On the **Transfer** menu, click **Capture Text** or, to send the text file directly to a printer, click **Capture to Printer**.
4. In the **File** box, type a descriptive name for the file, and then click **Start**.
5. Any text typed at the remote computer will now be captured in the specified file, or be sent to the printer.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- When saving incoming data to a text file or sending the data to a printer, HyperTerminal will capture all data transmitted during the connection session. You cannot use HyperTerminal to filter the data.
- To stop capturing the text to a file, on the HyperTerminal **Transfer** menu click **Capture Text**, then click **Stop** or **Pause**.
- To stop capturing the text to a printer, on the HyperTerminal Transfer menu click **Capture to Printer**, which will uncheck this selection.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
HyperTerminal overview
Receive a file from a remote computer

# Use HyperTerminal with Minitel

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use HyperTerminal with Minitel**

1. Open HyperTerminal.
2. Open a saved connection file or create a new connection.
3. On the **File** menu, click **Properties**, and then click the **Settings** tab.
4. In the **Emulation** list, click **Minitel**, and then click **OK**.
5. On the **View** menu, click **Font**.
6. In the **Font** list, click **Arial Alternative Symbol** then click **OK**.

**Notes**

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- Arial Alternative is the only font that contains the graphic characters you need for the Minitel service. If this font is not an available option, add the Arial Alternative Symbol and Arial Alternative Regular fonts to your computer.
- The function keys used for Minitel emulation are as follows:

| Function key | Minitel key |
|---|---|
| **HOME** | Sommaire |
| **DELETE** or **F2** | Annulation |
| **PAGE UP** or **F3** | Retour |
| **INSERT** or **F4** | Repetition |
| **END** or **F5** | Guide |
| **BACKSPACE** or **F6** | Correction |
| **PAGE DOWN** or **F7** | Suite |
| **ENTER** or **F8** | Envoi |
| **F9** | Connexion/fin |

- When you emulate a Minitel terminal, the **Port Settings** default values change to:
    - Data bits: 7

- Parity: Even
- Stop bits: 1

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Set up a new connection

# View East Asian languages with the VT-UTF8 terminal emulator in Hyperterminal

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view East Asian languages with the VT-UTF8 terminal emulator in Hyperterminal**

1. Open Regional and Language Options in Control Panel.
2. On the **Languages** tab, click **Install files for East Asian languages** and click **Apply**.
3. After the installation of the files for East Asian languages completes, do not restart your computer. Click the **Advanced** tab in the **Regional and Language Options** dialog box.
4. In the **Language for non-Unicode programs** box, select the language you want to view and click **OK**.
5. When prompted about using your existing files, click **Yes**.
6. Restart your computer.
7. After your computer has been restarted, open Regional and Language Options in Control Panel.
8. On the **Regional Options** tab, verify that the correct language (the East Asian language that you want to view) is selected in the **Standard and formats** box.
9. Close the Regional and Language Options dialog box and open HyperTerminal.
10. If the **Connection Description** dialog box is displayed, click **Cancel**.
11. In the **File** menu, click **Properties**.
12. On the **Connect To** tab, click **Configure**.
13. Change the **Port Speed** to match your port setting (this setting will likely be correct by default).
14. In the **Properties** box, click the **Settings** tab and select **VT-UTF8** in the **Emulation** box and click **OK**.
15. In the HyperTerminal **View** menu, click **Font** and select the correct East Asian language font.
16. Connect to the remote computer. For more information, see Related Topics below.

**Notes**

- You must be logged on as an administrator or a member of the Administrators group (or you must have been delegated the proper authority) to complete the first several steps of this procedure, in which you install the files for East Asian languages. After those files are installed and you restart your computer, the remaining steps do not require you to have administrative credentials. Therefore, as a security best practice, consider performing the final steps as a user without administrative credentials.
- To open Regional and Language Options, click **Start**, click **Control Panel**, and then double-click **Regional and Language Options**.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then

click **HyperTerminal**.

- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.

- The East Asian languages include Chinese, Japanese, and Korean. The files for most other languages are installed on your computer automatically by Windows.

- It is recommended that you install these files only if you are going to be entering or receiving text in one or more of these languages. The Chinese, Japanese, and Korean files require about 230 megabytes (MB) of space on your hard disk. In addition to the language files, Input Method Editors (IMEs) are installed for each of the three languages.

- Similar steps can be used to configure HyperTerminal to view other languages as well.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Set up a new connection

# Change the backscroll buffer size

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the backscroll buffer size**

1. Open HyperTerminal.
2. Open a saved connection file or create a new connection.
3. On the **File** menu, click **Properties**, and click the **Settings** tab.
4. In the **Backscroll buffer lines** box, type the number of communication text lines that you want saved in addition to what is displayed on the screen, and then click **OK**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- The maximum backscroll buffer value is 500 lines.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
HyperTerminal overview
Change the settings for making a call

# Change the number of columns

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the number of columns**

1. Open HyperTerminal.
2. Open a saved connection file or create a new connection.
3. On the **File** menu, click **Properties**.
4. Click the **Settings** tab. In the **Emulation** drop-down box, choose either **VT100**, **VT100J** or **VT-UTF8**.
5. Click the **Terminal Setup** button and as appropriate, either check or clear the **132 column mode** check box. Click **OK**, and click **OK** again.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open HyperTerminal, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, and then click **HyperTerminal**.
- HyperTerminal is not installed by default. For information on how to install HyperTerminal, see Related Topics.
- This procedure is valid only when emulating a VT100, VT100J or VT-UTF8 terminal.
- The default number of columns is 80.
- You can not change the number of rows displayed by HyperTerminal. VT-UTF8 and Minitel terminal emulation will display 25 rows and all other terminal types will display 24 rows.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install HyperTerminal
Change the settings for making a call
Change the port settings of the modem
Set up a new connection

# Stop using HyperTerminal as your telnet client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To stop using HyperTerminal as your telnet client

1. Open My Computer.
2. Click the **Tools** menu, and click **Folder Options**.
3. On the **File Types** tab, scroll down the list of file types to **URL: Telnet Protocol**.
4. Click **URL:Telnet Protocol**, and then click **Advanced**.
5. In the **Actions** box, click **Open** and then click **Edit**.
6. Delete the contents of the **Application used to perform action** field, and type:

   **rundll32.exe url.dll,TelnetProtocolHandler %l**

7. Close these dialog boxes by clicking **OK** to save your settings.

### Important

- Take note of the existing entry in the **Application used to perform action** field before replacing that entry. You can then repeat this procedure using the prior data if you want to again use HyperTerminal as your default telnet client.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open My Computer, click **Start**, and then click **My Computer**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# HyperTerminal Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about HyperTerminal:

- HyperTerminal overview
- Understanding HyperTerminal connections
- Understanding HyperTerminal file transfers
- HyperTerminal Resources

# HyperTerminal overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**HyperTerminal overview**

HyperTerminal is a program that you can use to connect to other computers, Telnet sites, bulletin board systems (BBSs), online services, and host computers, using either your modem, a null modem cable or Ethernet connection.

Though using HyperTerminal with a BBS to access information on remote computers is a practice that has become less common with the availability of the World Wide Web, HyperTerminal is still a useful means of configuring and testing your modem or examining your connection with other sites. For more information, see Set up a new connection.

HyperTerminal records the messages passed to and from the computer or service on the other end of your connection. Therefore, it can serve as a valuable troubleshooting tool when setting up and using your modem. To make sure that your modem is connected properly or to view your modem's settings, you can send commands through HyperTerminal and check the results. HyperTerminal has scroll functionality that allows you to look at received text that has scrolled off the screen.

You can use HyperTerminal to transfer large files from a computer onto your portable computer using a serial port rather than going through the process of setting up your portable computer on a network. For more information, see Send a file to a remote computer.

You can use HyperTerminal to help debug source code from a remote terminal. You can also use HyperTerminal to communicate with older character-based computers.

HyperTerminal is designed to be an easy-to-use tool and is not meant to replace other full-feature tools available on the market. You can use HyperTerminal to perform the specific tasks described above, but do not attempt to use HyperTerminal for more complex communication needs. For more information on what HyperTerminal does and does not support, see the HyperTerminal frequently asked questions list at the Hilgraeve Web Site.

Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Understanding HyperTerminal connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding HyperTerminal connections**

You can create and save multiple connections within HyperTerminal. Each connection can have its own unique configuration, including phone number and port settings such as bits per second(bps), parity checks, and host TCP/IP address.

As part of the connection configuration, HyperTerminal also allows your computer to emulate a variety of terminal types, such as TTY, ANSI and VTx terminals.

The connection that is currently open within HyperTerminal is displayed in the Hyperterminal title bar.

You can only have one connection open per each HyperTerminal session. You can, however, start multiple HyperTerminal sessions, opening a new connection for each session (provided each connection uses a different communication port). To set up a new connection, see Set up a new connection.

When you first open HyperTerminal, you will be prompted to decide if you want HyperTerminal to be your default telnet client. You also will be able to choose not to see this dialog box again. If you choose to not see the dialog box again, but later do want to set Hyperterminal as your default telnet client, see HyperTerminal Troubleshooting.

HyperTerminal cannot automatically provide your login ID and password when you make a connection. If you provide a password when using HyperTerminal for a Telnet session, be aware that this password will be sent to the remote computer using plain text (as with all telnet connections).

# Understanding HyperTerminal file transfers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding HyperTerminal file transfers**

You can use HyperTerminal to send and receive text and data files to a remote computer. The status of the transfer is displayed in the HyperTerminal window during the transfer. You can also capture the status of a transferred file and send the information directly to a printer. For more information, see Save incoming data to a text file.

You can use **Send File** to send data files as described in Send a file to a remote computer. You can send text files as an alternative to keyboard input using **Send Text File**. This is useful, for example, when you need to send a string of keyboard input to a remote computer for troubleshooting purposes, or need to send a lengthy keyboard input string as part of a connection sequence to a remote computer or router.

For information on receiving a file, see Receive a file from a remote computer.

Be aware that you cannot change the size of the HyperTerminal window to view more or less text than what is currently displayed.

# HyperTerminal Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This topic provides sources of additional information about HyperTerminal and modems.

- For a list of HyperTerminal frequently asked questions, see the HyperTerminal frequently asked questions list at the Hilgraeve Web Site.
- For more information about modems in general, see Installing and configuring Modems.
- For more information about supported modems, refer to the hardware compatibility information for products in the Windows Server 2003 family, by clicking the appropriate link in Support resources.
- Most modems support the AT command set. If you are uncertain about the exact command set supported by your modem, check the manufacturer's documentation. The modem documentation may also provide information about special commands supported by the modem. The modem manufacturer's Web site is often a good place to find technical information for your modem and the commands it supports.
- Microsoft Product Support Services provides a searchable database of technical support information for Microsoft products. You can search here for general information about HyperTerminal, modems or for information related to other Microsoft products. You can get to the Microsoft Product Support Services Web site by clicking the appropriate link in Support resources.
- For enhancements and updates to Windows software and other Microsoft software that might affect the use of HyperTerminal or your modem, see Windows Update at the Microsoft Web site.
- Microsoft TechNet is a service that contains resources for Information Technology (IT) professionals. This includes access to white papers and articles written about HyperTerminal by support professionals at Microsoft, career and industry information, discussion groups, technical feedback, and software. For more information about TechNet, see the Microsoft Web site.
- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# HyperTerminal Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- You receive an error when you are performing a file transfer and you are not able to send or receive a file successfully.
- You are unable to connect to a remote computer when using a remote site's user-friendly name.
- You want to use HyperTerminal as your default telnet application, but you are not prompted for that option when you open HyperTerminal.
- The terminal screen size does not change when you maximize HyperTerminal to full screen.
- Information you type is not appearing on the HyperTerminal terminal screen.
- Information you type is appearing double on the HyperTerminal terminal screen.
- ANSI characters are not displaying properly.
- HyperTerminal does not redial when the line is busy.
- After connecting to a remote computer, the terminal screen displays meaningless information.
- You cannot delete characters from the terminal screen.
- The cursor does not blink when you are using HyperTerminal via a Remote Desktop connection.
- You cannot use the traditional Windows control key combinations to perform certain tasks, such as CTRL+C to copy, or CTRL+V to paste.

**You receive an error when you are performing a file transfer and you are not able to send or receive a file successfully.**
Possible Cause

Solution

- The local computer and remote computer are not using the same file transfer protocol.
- Ensure that the file transfer protocol is the same on both the remote and local computer.
- You have excess noise on your phone line.
- Establish a new connection. You might need to use a different phone line.
- The connection to the phone line is interrupted briefly.
- Establish a new connection.
- Your hardware is not working correctly, you are using incorrect drivers, or your modem configuration is incorrect.
- Verify that the hardware is functioning properly and has the latest drivers. Verify that you are using the correct port. Also try your file transfer using different port speeds, baud rates, data bits, parity, stop bits and modulation. These settings

must be the same for both the remote and local computers.

**You are unable to connect to a remote computer when using a remote site's user-friendly name.**
**Cause:** The Domain Name System (DNS) server may not be functioning properly.

**Solution:** Connect to the site using standard Internet dotted notation (for example, w.x.y.z).

**You want to use HyperTerminal as your default telnet application, but you are not prompted for that option when you open HyperTerminal.**
**Cause:** When the dialog box asking you whether to use HyperTerminal as your telnet default client first appeared, you responded **No** and selected **Stop asking me this question**.

**Solution:** Open Registry Editor. In Registry Editor, navigate to the following registry key and change the value to 1:

**HKEY_CURRENT_USER\Software\Microsoft\HyperTerminal\Telnet Check**.

After you change this value, the next time you start HyperTerminal, the dialog box will appear and will ask you whether you want to have HyperTerminal as your default telnet application.

**Notes**

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.
- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

**The terminal screen size does not change when you maximize HyperTerminal to full screen.**
**Cause:** The HyperTerminal terminal screen size is determined by the font size that you are using. It will size itself to display 24 lines of 80 (or 132) characters of whatever font you have selected.

**Solution:** On the HyperTerminal **View** menu, select **Font**. If you want a larger terminal screen select a larger font size. If you want a smaller terminal screen, select a smaller font size.

**Information you type is not appearing on the HyperTerminal terminal screen.**
**Cause:** The terminal screen displays the information sent from the remote computer, not information entered on your local computer. To view what you type, the remote computer must be configured to echo back your input. There also may be a time lag between when you type the information and when it is displayed on your terminal screen.

**Solution:** Ensure that you have a proper connection with the remote computer and ensure that the remote computer is configured to echo user input.

**Cause:** Neither your computer nor the remote computer is configured to echo the information that is typed.

**Solution:** If the remote computer is not configured to echo user input, you can configure your computer to do so with the following steps:

1. On the HyperTerminal **File** menu, click **Properties**.
2. On the **Settings** tab, click **ASCII Setup**.
3. Check the **Echo typed characters locally** box.

**Information you type is appearing double on the HyperTerminal terminal screen.**
**Cause:** Both your computer and the remote computer are configured to echo back your input.

**Solution:** Either ensure that the remote computer is not configured to echo user input or instead configure your computer to not echo user input. To ensure that your computer is not echoing user input:

1. On the HyperTerminal **File** menu, click **Properties**.

2. On the **Settings** tab, click **ASCII Setup**.

3. Clear the **Echo typed characters locally** box.

**ANSI characters are not displaying properly.**
**Cause:** You are not using the Terminal font.

**Solution:** On the HyperTerminal **View** menu, select **Font**. Click **Terminal** and choose **OK**.

**HyperTerminal does not redial when the line is busy.**
**Cause:** The **Redial on busy** option may not be selected.

**Solution:** On the HyperTerminal **File** menu, select **Properties**. Select the **Connect To** tab and ensure that **Redial on busy** is selected.

**After connecting to a remote computer, the terminal screen displays meaningless information.**
**Cause:** You have not selected the correct terminal emulation type.

**Solution:** On the HyperTerminal **File** menu, select **Properties**. Select the **Settings** tab. In the **Emulation** drop-down box, choose the terminal type of the remote computer. If the remote computer's type is not listed in the drop-down box, it is not supported by HyperTerminal.

**You cannot delete characters from the terminal screen.**
**Cause:** The remote computer that you are connected to has control of the characters displayed on your terminal screen. The remote computer is expecting the cursor to be at particular positions on the screen based on the data that it has already sent to the screen. If you alter that screen locally, you potentially disrupt the interaction between you and the remote computer in ways that the host can not predict or control. Therefore, HyperTerminal does not allow you to delete characters from the screen.

**Solution:** Not Applicable

**The cursor does not blink when you are using HyperTerminal via a Remote Desktop connection.**
**Cause:** The cursor blink option is not available when using HyperTerminal via a Remote Desktop connection. This is done to reduce the screen updates in the Remote Desktop session.

**Solution:** Not applicable

**You cannot use the traditional Windows control key combinations to perform certain tasks, such as CTRL+C to copy, or CTRL+V to paste.**
**Cause:** If you are set to **Terminal Keys** in the properties for this connection, pushing control key combinations will send out terminal commands, not the Windows commands. For example, many hosts use CTRL+V to navigate through their systems.

**Solution:** You can change this setting to **Windows Keys** and the control key combinations will work as in Windows. To do so, on the HyperTerminal **File** menu, click **Properties**. Click the **Settings** tab, and click the **Windows Keys** radio button.

### Important

- When using the **Windows Keys** setting, all function, arrow and control keys will be processed locally. Using the **Terminal keys** setting and then using the menus to perform the necessary functions is the recommended option.

### Notes

- For more information on these topics and other frequently asked questions, see the HyperTerminal frequently asked questions list at the Hilgraeve Web site.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials. Exceptions are noted.

# Telnet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Telnet**

By using Telnet Client, you can connect to a remote computer that is running Telnet server software and you can interact with that computer through a console window. For more information about Telnet Client, see Telnet Client.

By using Telnet Server, you can allow users of Telnet client software to log on to your computer and to run character-mode applications on it. For more information about Telnet Server, see Telnet Server.

# Telnet Client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Telnet Client**
With Telnet Client, users can connect to a remote Telnet server and run character-mode applications on the remote computer or administer it.

- For general background information, see Telnet Client Concepts.
- For problem-solving instructions, see Telnet Client Troubleshooting.

# Telnet Client Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

Telnet is a member of the TCP/IP family of protocols, and it allows a user to establish a remote session on a server. The protocol supports only alphanumeric terminals -- that is, it does not support mice and other pointing devices, nor does it support graphical user interfaces. Instead, all commands must be entered at the command line.

The Telnet protocol provides very little security. In a Telnet session that does not use NTLM authentication, all data, including passwords, is transmitted between client and server in plaintext. Because Telnet session traffic is not secure, you need to ensure that no sensitive data is sent during a Telnet session.

This section provides general background information about Telnet Client:

- Telnet Client overview
- Using Telnet Client

# Telnet Client overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Telnet Client overview**

Telnet Client uses the Telnet protocol (part of the TCP/IP protocol suite) to connect to a remote computer over a network. Telnet Client allows a computer to connect to a remote Telnet server and run applications on that server.

All members of the Microsoft® Windows Server 2003 family include an improved, command line-oriented version of Telnet Client. You can use the NTLM authentication process to authenticate clients for Telnet connections between computers running Microsoft Windows NT®, Windows 2000, Windows XP, or the Windows Server 2003 family.

# Using Telnet Client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Telnet Client**

This section contains information on the following:

- Telnet Client authentication
- Telnet commands
- Session mode and command mode

# Telnet Client authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Telnet Client authentication**

You can use your local Windows user name and password or domain account information to access a computer running Telnet Server.

If you do not use the NTLM option for authentication, the user name and password are sent to the Telnet server as plaintext.

If you use NTLM authentication, Telnet Client uses the Windows security context for authentication, and the user is not prompted for a user name and password. The user name and password are encrypted.

Attempts to log on to the Telnet server using NTLM authentication will fail if the password option is set at **User Must Change Password at Next Log on**. To log on successfully, you must log on to the server directly, change your password, and then log on through Telnet Client.

If you connect to a Telnet server using NTLM authentication, you will not be able to access additional network resources because of a limitation of the NTLM authentication. In order to access network resources from a Telnet session, you need to access network drives by providing your user name and password again.

**Note**

- Although authentication might be encrypted, the actual data being transferred can be read by anyone on the network. Because Telnet session traffic is not secure, ensure that no sensitive data is sent or received during a Telnet session.

# Telnet commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Telnet commands**

The Telnet commands allow you to communicate with a remote computer that is using the Telnet protocol. You can run Telnet without parameters in order to enter the Telnet context, indicated by the Telnet prompt (`Microsoft Telnet>`). From the Telnet prompt, use the Telnet commands to manage a computer running Telnet Client.

The Telnet Client command prompt accepts the following commands:

| Command | Description |
|---|---|
| **open** | Use **open** *hostname* to establish a Telnet connection to a host. |
| **close** | Use the **close** command to close an existing Telnet connection. |
| **display** | Use the **display** command to view the current settings for Telnet Client. |
| **send** | Use the **send** command to send commands to the Telnet server. The following commands are supported:<br><br>**ao**<br><br>Abort output command.<br><br>**ayt**<br><br>"Are you there?" command.<br><br>**esc**<br><br>Sends the current escape character.<br><br>**ip**<br><br>Interrupt process command.<br><br>**synch**<br><br>Performs the Telnet synch operation.<br><br>**brk**<br><br>Sends a break signal.<br><br>Anything other than one of the commands listed above will be sent as a string to the Telnet server. For example, **send abcd** will send the string **abcd** to the Telnet server, which will echo the string in the Telnet session window. |
| **quit** | Use the **quit** command to close Telnet Client. |
| **set** | Use the **set** command with one of the following arguments to configure Telnet Client for the current session.<br><br>**bsasdel**<br><br>Backspace will be sent as delete.<br><br>**codeset option**<br><br>Available only when the language is set to Japanese. Set the current code set to option, which can be one of the following:<br><br>• Shift JIS |

| | |
|---|---|
| | • Japanese EUC<br>• JIS Kanji<br>• JIS Kanji (78)<br>• DEC Kanji<br>• NEC Kanji<br><br>The same code set should be set on the remote computer. By default, Telnet Client uses a raster font. Before accessing a remote computer using one of these code sets, you must configure Telnet Client to use a TrueType font to ensure that characters will display properly.<br><br>**crlf**<br><br>New-line mode; causes the RETURN key to send 0x0D, 0x0A.<br><br>**delasbs**<br><br>Delete will be sent as backspace.<br><br>**escape** *character*<br><br>Switches from Telnet session mode to the Telnet command mode. While in Telnet command mode, press ENTER to return to Telnet session mode.<br><br>**localecho**<br><br>Turns on localecho.<br><br>**logfile** *name*<br><br>Specifies the name of the file to which the Telnet log is written for this session. If you do not specify the path of the file, it is created in your current directory. Specifying a log file also turns on logging.<br><br>**logging**<br><br>Turns on logging for this session.<br><br>**mode** {**console** \| **stream**}<br><br>The mode of operation.<br><br>**ntlm**<br><br>Turns on NTLM authentication.<br><br>**term** {**ansi** \| **vt100** \| **vt52** \| **vtnt**}<br><br>The type of terminal you want Telnet Client to emulate.<br><br>**?**<br><br>Displays Help information for **set**. |
| **unset** | Use **unset** to turn off an option previously set using the **set** command. |
| **status** | Use the **status** command to determine whether the computer running Telnet Client is connected. |
| **?/help** | Displays Help information. |

For more information about Telnet Server commands, see Telnet Command-line Reference for Telnet Server.

# Session mode and command mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Session mode and command mode**

There are two modes of operation in Telnet Client: Telnet command mode and Telnet session mode. Telnet command mode allows the Telnet terminal to open a connection to a remote host, close a connection to a remote host, display operating parameters, set terminal options, print the status, and quit the program.

For information about how to set Telnet options and to learn how to send Telnet commands, see Telnet Command-line Reference for Telnet Server.

Once connected to a remote host, Telnet Client is in Telnet session mode. This mode is the most common. After logging on, users are given a Command Prompt window. Users are then able to use character-based applications on any remote Telnet server, as if they were logged on to the system directly.

Once connected to a host, you can return from session mode to command mode to change terminal settings. You can change from a Telnet session mode to a Telnet command mode by pressing **CTRL** + **]**; press ENTER to return to the Telnet session mode. For more information about switching modes, see Telnet commands.

**Important**

- **Use a TrueType font in the Command Prompt window when using Telnet with any of the four supported Asian languages**

  Telnet supports the following four Asian languages: Japanese, Simplified Chinese, Korean, and Traditional Chinese. You must use a TrueType font in the Command Prompt window to properly display the characters used in Asian languages. By default, the Command Prompt window displays characters using raster fonts.

  Because each Telnet session creates a new Command Prompt window, you need to permanently set the properties of the Command Prompt window to use the TrueType fonts if you use Telnet with Asian languages on a regular basis.

  1. To set the Command Prompt window font, right-click in the title bar of the Command Prompt window, and then click **Properties**.
  2. Select a TrueType font, and then click **OK**.
  3. To always use the TrueType font you selected, in the **Apply Properties to Shortcut** dialog box, click **Modify shortcut that started this window**.

- This can also be done using the **Save properties for all future windows with the same title** option. If this option is not set, DBCS (Double Byte Character Set) characters will not be displayed properly in the Telnet session.

**Note**

- Although authentication might be encrypted, the actual data being transferred can be read by anyone on the network. Because Telnet session traffic is not secure, ensure that no sensitive data is sent or received during a Telnet session.

# Telnet Client Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section contains information on troubleshooting Telnet Client. For troubleshooting information on Telnet Server, see Telnet Server Troubleshooting.

What problem are you having?

- Telnet Client is not restored after Services for UNIX is uninstalled.
- Characters and lines on the Telnet Client display are misaligned when connecting to the computer running Telnet Server.

**Telnet Client is not restored after Services for UNIX is uninstalled.**
**Cause:** Setup attempted to restore the original Telnet files but failed.

**Solution:** Restore the Telnet files from the Windows CD. For example, here are the commands you would use to restore the required files if C:\WINDOWS is the system folder, and D is the source CD:

```
C:\>copy D:\I386\telnet.exe C:\WINDOWS\SYSTEM32

C:\>expand -r D:\I386\tlntsvr.ex_ C:\WINDOWS\SYSTEM32

C:\>expand -r D:\I386\tlntsvrp.dl_ C:\WINDOWS\SYSTEM32

C:\>expand -r D:\I386\tlntsess.ex_ C:\WINDOWS\SYSTEM32
```

**Characters and lines on the Telnet Client display are misaligned when connecting to the computer running Telnet Server.**
**Cause:** This problem can be caused by a mismatch between the size of the Telnet Client window and the maximum size of the command window that Telnet Server can support.

**Solution:** Try reducing the size of the Telnet Client window before reconnecting.

# Telnet Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Telnet Server**
Telnet Server provides ASCII terminal sessions to Telnet clients. Telnet Server supports two types of authentication and supports four types of terminals: ANSI, VT-100, VT-52, and VTNT.

- For tips about using Telnet Server, see Telnet Server best practices.
- For general background information, see Telnet Server Concepts.
- For problem-solving instructions, see Telnet Server Troubleshooting.

# Telnet Server best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Provide user-level security**

  To control the access that Telnet users have to files on the server, use only the NTFS file system on the system partition. Create a TelnetClients group, add all Telnet users to that group, and then assign file and directory permissions to control the access that members have to the files and directories.

- **Allow users to disconnect before stopping the service**

  Before you stop Telnet Server or uninstall it, send a message to Telnet client sessions that you are about to stop the service. You can then stop the service after the users have had the opportunity to close their sessions.

- **Make sure that the client code page matches the code page of the UNIX terminal**

  If users will be connecting to Telnet Server from computers running an internationalized version of UNIX (such as a version capable of supporting European languages), you should ensure that the code page used by the command shell can display extended characters properly. To set the code page for all users when they log on, edit *systemroot*\system\login.cmd to add the **chcp** command to set the appropriate code page. For example, to support English and Western European UNIX, add the command **chcp 1252**.

# Telnet Server Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

Telnet is a member of the TCP/IP family of protocols, and it allows a user to establish a remote session on a server. The protocol supports only alphanumeric terminals -- that is, it does not support mice and other pointing devices, nor does it support graphical user interfaces. Instead, all commands must be entered at the command line.

The Telnet protocol provides very little security. In a Telnet session that does not use NTLM authentication, all data, including passwords, is transmitted between client and server in plaintext. Because of this limitation, as well as the general recommendation against allowing untrusted users access to security-critical servers, you should run Telnet Server only on computers that store no sensitive data.

This section contains information about:

- Telnet Server overview
- Telnet Server access
- Telnet Server authentication
- Administering Telnet Server

# Telnet Server overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Telnet Server overview**

Telnet Server is a gateway for Telnet clients. When Telnet Server is running on a computer, users can use Telnet clients to connect to it from remote computers. When a Telnet client connects to a computer running Telnet Server, the remote user is asked to enter a user name and password. By default, only user name and password combinations that are valid on the local server can be used to log on to that server.

Once logged on, a user is given a command prompt that can be used as if it had been opened in a command prompt window locally. By default, however, the user cannot use applications that interact with the desktop.

Telnet uses the Telnet protocol (part of the TCP/IP protocol suite) to connect to a remote computer over the network.

Members of the Administrators group can log on to Telnet Server. Access by other users is controlled by membership in the TelnetClients group. By default, this group contains no entries. If you want to allow users who are not members of the Administrators group to log on to Telnet Server, add the appropriate users or groups to the TelnetClients group.

**Telnet licensing**

With each installation of a member of the Microsoft® Windows Server 2003 family, a license provides one installation of Telnet Server. However, several Telnet sessions can exist simultaneously.

By default, Telnet Server supports a maximum of two Telnet clients at a time. However, the administrator can change this value by using the **tlntadmn** command. For detailed information about using the **tlntadmn** command, see Telnet Command-line Reference for Telnet Server. The number of distinct devices that can connect to a computer running Telnet Server is limited by available resources like RAM, virtual memory, and the number of Client Access Licenses (CALs).

### Note

- After you create a TelnetClients group, you must stop and start the Telnet Server service before users can log on.

# Telnet Server access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Telnet Server access**
By default, only members of the Administrators group can log on to Telnet Server. You can give access to your computer running Telnet Server to other users. To allow other users to log on, create a local user group, on your server, called TelnetClients, and add users to that group. If a group with that name exists on your server, only that group's members can log on to Telnet Server.

- **To create a local user group called TelnetClients:** See Create a local group.

- **To grant access to a computer running Telnet Server**

  1. Open Computer Management.

  2. In the console tree, expand **Local Users and Groups**, and click **Groups**.

  3. Double-click the TelnetClients group.

  4. Click **Add...**

  5. Follow the instructions in the **Select Users** box to add users to the TelnetClients group, and then click **OK**.

**Notes**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- The group name TelnetClients must be spelled exactly as shown.

- After you create a TelnetClients group, you must stop and start the Telnet Server service before users can log on.

# Telnet Server authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Telnet Server authentication**
Authentication is the means by which a user is identified and authorized. Telnet Server supports two methods of authentication: NTLM and plaintext.

If you use NTLM authentication, Windows-based clients use the Windows security context for authentication, and the user is not prompted for a user name and password. The user name and password are encrypted. Sessions that do not use encryption use unencrypted text (also known as plaintext) and are visible across the network.

If you do not use NTLM authentication the user name and password are sent to the computer running Telnet Server as plaintext. Anyone capturing the packets of the authentication process can easily read the password and use it to gain unauthorized access to your intranet. The use of plaintext authentication is therefore highly discouraged.

If **User Must Change Password at Next Logon** is set for a user, an attempt to log on to the computer running Telnet Server will fail. The user must log on to the server directly and change the password, then log on through Telnet.

If you connect to a computer running Telnet Server using NTLM authentication, you will not be able to access additional network resources because of a limitation of the NTLM authentication. In order to access network resources from a Telnet session, you need to access network drives by providing your user name and password again.

NTLM might not be the chosen mode of client authentication. This occurs when your client is:

- A Windows-based client that does not use NTLM by choice
- A UNIX Telnet client

In this scenario, the only other authentication method supported by the computer running Telnet Server is the user name/password method. In this method, the user name and password are sent as plaintext for authentication by the server.

For information about how to set authentication methods, see Telnet Command-line Reference for Telnet Server.

**Note**

- Although authentication data might be encrypted, the actual data being transferred can be read by anyone on the network. Because Telnet session traffic is not secure, ensure that no sensitive data is sent or received during a Telnet session.

# Administering Telnet Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Telnet Server**

This section covers:

- Using the tlntadmn command
- Using login.cmd
- Managing Telnet from the command line

# Using the tlntadmn command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the tlntadmn command**

You can use the **tlntadmn** command to administer a local or a remote computer that is running Telnet Server. This command is new in Windows XP and the Windows Server 2003 family. For detailed information about using the **tlntadmn** command, see Telnet Command-line Reference for Telnet Server.

**System requirements**

Both the computer you are administering and the computer at which you use the **tlntadmn** command must be running Windows NT, Windows 2000, Windows XP, or a member of the Windows Server 2003 family. If either computer is running Windows NT or Windows 2000, it must also be running Windows Services for UNIX 2.0.

**Security considerations**

To administer a local or a remote computer using the **tlntadmn** command, you must log on to the local computer with administrative credentials. To administer a remote computer, you must also do one of the following:

- log on to the local computer with an account that has administrative credentials on the remote computer
- use the **-u** and **-p** parameters to provide administrative credentials for the remote computer

Use of the **-u** and **-p** parameters is a non-NTLM method of authentication. This method is less secure than NTLM because credentials are sent in plaintext. As a security measure, you cannot use the **-s**, **-k**, and **-m** parameters even if you provide credentials using the **-u** and **-p** parameters. The **-s**, **-k**, and **-m** parameters are used to list a session, end a session, and send a message, respectively.

# Using login.cmd

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using login.cmd**

The program login.cmd is run whenever someone logs on to Telnet Server. You can edit the login.cmd file to change the text that appears when someone logs on, and you can use login.cmd to run other programs.

When Telnet Server is installed, the program is copied to %SystemRoot%\system32\login.cmd. This is what it looks like:

```
@echo off
rem
rem  Default global logon script for the Telnet Server
rem
rem  In the default setup, this command script is run
rem  when the initial command shell is invoked. It, in
rem  turn, will try to invoke the individual user's
rem  logon script.
rem
echo
*====================================================
echo
Welcome to Microsoft Telnet Server.
echo
*================================================
cd /d %HOMEDRIVE%\%HOMEPATH%
```

In the following example, the file has been edited to display the text "Welcome to my Telnet Server," and run **net config** to display information about the server configuration:

```
@echo off
rem
rem  Default global logon script for the Telnet Server
rem
rem  In the default setup, this command script is run
rem  when the initial command shell is invoked. It, in
rem  turn, will try to invoke the individual user's
rem  logon script.
rem
echo
*=================================================
echo
Welcome to my Telnet Server
echo
*=================================================
cd /d %HOMEDRIVE%\%HOMEPATH%net config server
```

# Managing Telnet from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Telnet from the command line**

By using the **tlntadmn** command, you can administer a local or a remote computer running Telnet Server. For a list of **tlntadmn** parameters, see Telnet Command-line Reference for Telnet Server.

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

# Telnet Server Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section contains information on troubleshooting Telnet Server. For troubleshooting information on Telnet Client, see Telnet Client Troubleshooting.

What problem are you having?

- Logon failure: The user has not been granted the requested logon type at this computer.
- Access denied: Specified user is not a member of the TelnetClients group.
- Logon failure: The user must change his or her password before logging on the first time.
- Logon failure: Telnet Server allows NTLM authentication only.
- Error while connecting to localhost: Could not open a connection to host on port 23.
- Cannot connect to network drives.
- Users cannot access drive letters mapped by others.
- If you interrupt and then reconnect a Telnet session that is being conducted from a UNIX computer to a Windows-based computer, the cursor might lose alignment with the prompt.

**Logon failure: The user has not been granted the requested logon type at this computer.**
**Cause:** The user does not have permission to log on locally to the computer running Telnet Server.

**Solution:** Ensure that the user, or a group the user belongs to (such as the TelnetClients group), has the right to log on locally. Ensure that the user, or a group that the user belongs to, is not denied the right to log on locally.

**Access denied: Specified user is not a member of the TelnetClients group.**
**Cause:** Users must belong to the TelnetClients group in order to connect to Telnet Server.

**Solution:** Add the user to the TelnetClients group.

**Logon failure: The user must change his or her password before logging on the first time.**
**Cause:** The user's password must be changed.

**Solution:** The user must log on to the computer directly (that is, in person) to change the password.

**Logon failure: Telnet Server allows NTLM authentication only.**
**Cause:** Either the client does not support NTLM authentication and the server requires it, or the client requires NTLM authentication and Telnet Server does not support it.

**Solution:** If the Telnet client does not support NTLM authentication, use a client that does. If that is not possible, enable plaintext authentication on Telnet Server. For maximum security, use only Telnet clients that support NTLM authentication, and enable NTLM authentication and disable plaintext authentication on Telnet Server.

**Error while connecting to localhost: Could not open a connection to host on port 23.**
**Cause:** Telnet Server is not configured to listen on port 23, or Telnet Server is not started.

**Solution:** Ensure that Telnet Server is started and configured to listen on port 23. If that is not possible, specify the configured port after the host name when connecting.

**Cannot connect to network drives.**
**Cause:** NTLM authentication does not allow access to network drives.

**Solution:** Map the network to a drive letter using the **net use /user** command.

**Users cannot access drive letters mapped by others.**

**Cause:** Telnet clients recognize only drive letters mapped during the current Telnet session. Consequently, Telnet users cannot access network drives using drive letters mapped by other Telnet users or by local server users.

**Solution:** Telnet client users must use the **net use /user** command to map network drives to drive letters themselves.

**If you interrupt and then reconnect a Telnet session that is being conducted from a UNIX computer to a Windows-based computer, the cursor might lose alignment with the prompt.**

**Cause:** This problem is caused by the interruption in communication.

**Solution:** If your default shell is the Windows command prompt, type **cls**. If your default shell is the Korn shell or C shell, type **clear**.

# Configuring Communication Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring communication services**

This section covers:

- Telephony

# Telephony

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Telephony**
This implementation of Telephony Application Programming Interface (TAPI) unifies Internet Protocol (IP) and traditional public switched telephone network (PSTN) telephony to support programs that effectively work the same way over intranets, the Internet, and traditional networks. This section provides information for configuring and managing both IP and PSTN telephony servers and clients.

- Before you begin using telephony with this computer, see Telephony Checklists.
- To find features that have been moved in Microsoft® Windows® 2000, Windows XP, and the Windows Server 2003 family, see New ways to do familiar Telephony tasks.
- For tips about using telephony, see Telephony Best practices.
- For help with specific tasks, see Telephony How To....
- For general background information, see Telephony Concepts.
- For problem-solving instructions, see Telephony Troubleshooting.

# Telephony Checklists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Checklists**

- Checklist: Installing IP telephony (H.323)
- Checklist: Installing IP multicast video conferencing
- Checklist: Setting up a telephony server
- Checklist: Setting up telephony clients

# Checklist: Installing IP telephony (H.323)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Checklist: Installing IP telephony (H.323)**

| Step | Reference |
|---|---|
| Review key concepts. | IP Telephony |
| Confirm that TCP/IP is running on your network. | TCP/IP Core Networking Service |
| (Optional) Confirm that you have an Active Directory domain and that a TAPI application directory partition is available on your network.<br><br>TAPI application directory partitions are not required for H.323 IP telephony, but if a TAPI application directory partition is available on your network, users of TAPI-based programs can specify user names, rather than computer names or IP addresses, when making H.323-based IP telephony calls. | TAPI and Active Directory;Tapicfg;Active Directory |
| Configure H.323 gateways and firewall proxies. | Specify an H.323 gateway or proxy |
| Configure H.323 gatekeepers. | Specify an H.323 gatekeeper |
| Make voice or video IP telephony calls. | Manufacturer's documentation |

# Checklist: Installing IP multicast video conferencing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Checklist: Installing IP multicast video conferencing**

| Step | Reference |
|---|---|
| Review key concepts. | IP Multicast Video Conferencing |
| Confirm that TCP/IP is running on your network. | TCP/IP Core Networking Service |
| Confirm that your network is enabled for multicast. | Enabling IP multicast |
| Confirm that you have an Active Directory domain and that a TAPI application directory partition is available on your network. | TAPI and Active Directory;Active Directory |
| Configure MADCAP on your network for IP multicast telephony.<br>• Install the DHCP server if it is not already installed.<br>• For networks using domains, you must authorize the DHCP service in the domain's directory service database. This prevents users from setting up unauthorized DHCP services.<br>You do not need to do this, however, if the server has been authorized already (for example, if it already functioning as a DHCP server).<br>• To have the DHCP server allocate multicast addresses using the MADCAP protocol, create a new multicast scope. | DHCP;Authorize a DHCP server in Active Directory;Create a multicast Telephony scope |
| Set up multimedia devices on client computers. | Setting up clients for IP telephony;Configure multimedia devices;Eliminating acoustic echo |
| Set up conference rooms. | Setting up conference rooms;Eliminating acoustic echo |
| Make multicast video conference calls. | Manufacturer's documentation. |

# Checklist: Setting up a telephony server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Checklist: Setting up a telephony server**

| Step | Reference |
|---|---|
| Review key concepts. | General Telephony Concepts |
| Confirm that you have an Active Directory domain. | Active Directory |
| For the telephony server, set up or select a computer that is running a Windows Server 2003 operating system and that is in the same domain as the telephony clients or that is in a domain that has a two-way trust relationship with the domain containing the telephony clients. | Trusts |
| Install telephony hardware and software on the server. | Manufacturer's documentation |
| Enable the telephony server. | Enable or disable a telephony server |
| Specify a logon account for the telephony server. (You can create a specific account, such as TAPI_ADMIN.) | Change the Telephony service logon account for a telephony server |
| Designate telephony administrators who are not in the administrator group on the telephony server. | Designate telephony administrators |
| Assign telephony users to lines or phones. | Assign a telephony user to a line or phone |

# Checklist: Setting up telephony clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Checklist: Setting up telephony clients**

| Step | Reference |
|---|---|
| Review key concepts. | General Telephony Concepts |
| Confirm that you have an Active Directory domain. | Active Directory |
| (Optional) Upgrade Windows NT 4.0 and Windows 95 clients to TAPI version 2.1 or later, if you have not already done so. | Telephony Resources |
| Install any necessary telephony hardware and software on the clients. | Manufacturer's documentation |
| (Optional) If clients are running Windows 98 or Windows Millennium Edition, enable Windows 98 or Windows Millennium Edition telephony clients. | Enable a Windows 98 or Windows Millennium Edition telephony client |
| On each client, specify the telephony servers to be used by the client. | Specify telephony servers on a client computer; Tcmsetup |
| Configure general dialing rules, area code rules, and calling cards as required for users. | Create a new location |
| Dial and receive phone calls. | Manufacturer's documentation. |

# New ways to do familiar Telephony tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**New ways to do familiar tasks**

The following table lists common tasks for configuring telephony. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 | In this version of Windows use |
|---|---|---|
| **Create a dialing location** | **Telephony** in **Control Panel** | Open Phone and Modem Options. For more information, see Create a new location. |
| **Configure dialing properties for a location** | **Telephony** in **Control Panel** | Open Phone and Modem Options. For more information, see Change dialing properties for a location. |
| **Specify area code rules** | **Telephony** in **Control Panel** | Open Phone and Modem Options. For more information, see Specify area code rules. |
| **Create a calling card** | **Telephony** in **Control Panel** | Open Phone and Modem Options. For more information, see Create a new calling card. |
| **Set up or disable a telephony server** | The **Tcmsetup** command | Open Telephony. For more information, see Enable or disable a telephony server. |
| **Set up or disable a telephony client** | The **Tcmsetup** command | For this task, there is a user interface difference between Windows 2000 and Windows XP and the Windows Server 2003 family.<br><br>If you have telephony servers or clients that run Windows 2000, use the **Tcmsetup** command. For more information, see Specify telephony servers on a client computer.<br><br>If you have telephony servers that run Windows Server 2003 and telephony clients that run Windows XP or a Windows Server 2003 operating system, open Phone and Modem Options. For more information, see Add a telephony service provider. |
| **Assign users to lines or phones** | The **Tcmapp** command | Open Telephony. For more information, see Assign a telephony user to a line or phone. |
| **Add a telephony service provider on a local or remote server** | **Telephony** in **Control Panel** | Open Telephony. For more information, see Add a telephony service provider on a telephony server. |
| **Add a telephony service provider on a client** | **Telephony** in **Control Panel** | Open Phone and Modem Options. For more information, see Add a telephony service provider. |

# Telephony Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Best practices**

- **As a rule, any computer with enough processing power to run a** Windows Server 2003 **operating system can act as a telephony server, linking client computers to hardware, such as a PBX.**

  Other services, especially server-side telephony programs, may have different requirements.

  Remember that, if the telephony server is the PBX, it may make additional processing demands.

- **Obtain service provider components required to communicate with a PBX from a PBX vendor or value-added reseller.**

  TAPI does not supply the service provider components required to communicate with a PBX.

- **Like networking software, a telephony program must be secured to protect privacy and data integrity.**

  The Telephony service must run on a computer that is a member of a Windows Server 2003 domain, and all users must have valid domain accounts. The domain administrator defines which users have permission to communicate with the switch, and the telephony server administrator defines which lines those users can access.

- **Telephony client users must be in the same domain as the telephony server.**

  Make sure that telephony clients are in the same domain as a telephony server or in a domain that has a two-way trust relationship with the domain containing the server. MADCAP servers must also be in the same domain as their client computers.

- **Specify the server on a telephony client or install the Windows Remote Service Provider to make lines or phones available to the client users**

  TAPI is automatically installed on computers running Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003 operating systems. However, unless your servers are running Windows Server 2003 , and your clients are running Windows XP or a Windows Server 2003 operating system, you must specify a server on the telephony clients before the server's lines are available to the client users.

  For more information, see Specify telephony servers on a client computer.

  If your servers are running Windows Server 2003 , and your clients are running Windows XP or a Windows Server 2003 operating system, you only need to add the Windows Remote Service provider.

  For more information, see Add a telephony service provider.

- **Set up a domain and install Active Directory for IP multicast video conferencing (and optionally, for H.323 IP telephony)**.

  In Windows Server 2003 Active Directory, application directory partitions store and replicate static or dynamic program-specific data. TAPI application directory partitions are not required for H.323 IP telephony, but if one is available on your network, users of TAPI-based programs can specify user names, rather than computer names or IP addresses, when making H.323-based IP telephony calls.

  **Note**

- You can still use Internet Locator Service (ILS) servers in your organization, if ILS is needed to support certain programs, because telephony clients running Windows XP or a Windows Server 2003 operating system can query either ILS servers or TAPI application directory partitions.

For more information about TAPI and Active Directory, see TAPI and Active Directory.

- **If you create a multicast scope, be certain to provide a range of addresses that is large enough for your needs.**

  If a range is not large enough, the scope's pool of addresses will be exhausted. The time-out period must also allow addresses to be reused often enough to keep the pool of addresses from being exhausted. A setting of 30 days is normally sufficient.

  For more information, see Using multicast scopes, Create a multicast Telephony scope, and Enabling IP multicast.

# Telephony How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**How to...**

- Create a multicast Telephony scope
- Manage a TAPI Application Directory Partition
- Manage Telephony Servers
- Manage Telephony Service Providers on a Server
- Manage Telephony Clients and Users
- Change How Phone Calls Are Dialed

# Create a multicast Telephony scope

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To create a multicast scope**

1. Open DHCP.
2. In the console tree, right-click the applicable DHCP server, and then click **New Multicast Scope**.

   **Where?**

   - DHCP/*Applicable DHCP server*

3. On the **Welcome** page of the New Multicast Scope Wizard, click **Next**.
4. Type a name and description for the scope, and then click **Next**.
5. In **IP Address range**, type a valid IP address range for the multicast scope.
6. In **TTL**, specify the number of routers that multicast traffic should pass through on your network.
7. In **Add Exclusions**, add any IP addresses or a range of IP addresses to exclude from the distribution, and then click **Next**.
8. In **Lease Duration**, select a DHCP lease duration for the affected IP addresses, and then click **Next**.
9. In **Activate Multicast Scope**, click **Yes** or **No** to specify whether to activate the multicast scope immediately, and then click **Next**.
10. In **Completing the New Multicast Scope Wizard**, click **Finish**.
11. Repeat the procedure to create any additional scopes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open DHCP, click **Start**, click **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **DHCP**.
- Provide a range of addresses that is large enough for your needs, and be sure to set a time-out period that will allow addresses to be reused often enough to keep the scope's pool of addresses from being exhausted. A setting of 30 days is normally sufficient.
- If you are running the MADCAP server inside a domain, it might generate warning messages for granting permissions to the DHCP server. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enabling IP multicast
Using multicast scopes
Connect to a DHCP server
Authorizing DHCP servers

# Manage a TAPI Application Directory Partition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Manage a TAPI application directory partition**

- Create a TAPI application directory partition
- Remove a TAPI application directory partition

# Create a TAPI application directory partition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To create a TAPI application directory partition**

1. Verify that the computer is a domain controller and that the Active Directory directory service is running.

2. Open Command Prompt.

3. At the command prompt, type:

   **tapicfg install /directory:** *PartitionName* [/**server:***DCName*] [**/forcedefault**]

| Value | Description |
|---|---|
| **/directory:** *PartitionName* | Required. Specifies the DNS name of the TAPI application directory partition that is to be created. Note that this name must be a fully qualified domain name. |
| **/server:** *DCName* | Specifies the DNS name of the domain controller on which the TAPI application directory partition is to be created. If the domain controller name is not specified, the name of the local computer is used. |
| **/forcedefault** | Specifies that this directory will be the default TAPI application directory partition for the domain. There can be multiple TAPI application directory partitions in a domain.<br><br>If this application directory partition is the first TAPI application directory partition created on the domain, it is automatically set as the default, regardless of whether you use the **/forcedefault** option. |

**Notes**

- To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory. For more information, see Default local groups and Default groups.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at the command prompt, type:

  **tapicfg help**

- You can also use **tapicfg** to create and remove service connection points that enable telephony clients to efficiently locate TAPI application directory partitions in a domain. For more information, see Related Topics.

- You can also create a TAPI application directory partition by running the Configure Your Server Wizard. If you select **Typical setup for a first server**, a TAPI application directory partition with the name MSTAPI.*DomainName* is automatically created on the first domain controller that is configured in a forest. For more information about the prerequisites for configuring the typical setup for a first server, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Manage Domain Controllers
Remove a TAPI application directory partition
TAPI and Active Directory
Tapicfg
Typical setup for a first server

# Remove a TAPI application directory partition

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To remove a TAPI application directory partition**

1. Open Command Prompt.

2. At the command prompt, type:

   **tapicfg remove / directory:** *PartitionName*

| Value | Description |
|---|---|
| **/directory:** *PartitionName* | Required. Specifies the DNS name of the TAPI application directory partition that is to be removed. Note that this name must be a fully qualified domain name. |

## Notes

- To perform this procedure, you must be a member of the Enterprise Admins group in Active Directory. For more information, see Default local groups and Default groups.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at the command prompt, type:

  **tapicfg help**

- If you do not know the name of the TAPI application directory partition, at the command prompt, type **tapicfg show** to view a list of the TAPI application directory partitions that are in the domain. Type **tapicfg show [/defaultonly]** to view the TAPI application directory partition that is the default for the domain.

- You can also use **tapicfg** to create and remove service connection points that enable telephony clients to efficiently locate TAPI application directory partitions in a domain. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a TAPI application directory partition
TAPI and Active Directory
Tapicfg

# Manage Telephony Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Manage telephony servers**

- Enable or disable a telephony server
- Add a server to the Telephony console
- Delete a server from the Telephony console
- Designate telephony administrators
- Change the Telephony service logon account for a telephony server
- Update information for a telephony server

# Enable or disable a telephony server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable or disable a telephony server**

1. Open Telephony.
2. In the console tree, click the server you want to manage.

   **Where?**

   - Telephony/*server*

3. On the **Action** menu, click **Properties**.
4. On the **Setup** tab, do one of the following:

   - To enable the telephony server, select the **Enable telephony server** check box.
   - To disable the telephony server, clear the **Enable telephony server** check box.

5. If you are enabling the telephony server, under **Account Information**, specify a logon account for the Telephony service. For more information, see Related Topics.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- If the **Enable telephony server** check box is unavailable, you need to start the Telephony service first. For more information, see Related Topics.
- This procedure changes whether clients can use the telephony devices on the server. Changes take effect the next time the system attempts to start the Telephony service.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a server from the Telephony console
Designate telephony administrators
Update information for a telephony server
Start, stop, pause, resume, or restart a service

# Add a server to the Telephony console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To add a server to the Telephony console**

1. Open Telephony.
2. In the console tree, click **Telephony**.

   **Where?**

   - Telephony

3. On the **Action** menu, click **Add Computer**.
4. In the **Select Computer** dialog box, click **Locations**, and then click the domain where the server is located.
5. In **Enter the object name to select**, enter the name of the telephony server to manage, and then click **OK**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- The telephony server must be enabled on the computer you want to administer. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable a telephony server
Delete a server from the Telephony console
Designate telephony administrators
Change the Telephony service logon account for a telephony server
Update information for a telephony server

# Delete a server from the Telephony console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To delete a server from the Telephony console**

1. Open Telephony.
2. In the console tree, click the server you want to remove.

   **Where?**

   - Telephony/*server*

3. On the **Action** menu, click **Delete**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- This procedure removes the server from the Telephony console. It does not disable it or prevent it from being used by telephony clients. For the procedure to disable a telephony server, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable a telephony server
Designate telephony administrators
Change the Telephony service logon account for a telephony server
Update information for a telephony server

# Designate telephony administrators

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To designate telephony administrators**

1. Open Telephony.
2. In the console tree, click the server you want to manage.

   **Where?**

   - Telephony/*server*

3. On the **Action** menu, click **Properties**.
4. On the **Setup** tab, in **Telephony administrators**, click **Add**.
5. In the **Select Users** dialog box, click **Locations**, and then click the domain where the user account is located.
6. In **Enter the object names to select**, enter the name of the user to be designated as a Telephony administrator, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- Use this procedure to enable users who are not administrators for the server to administer the Telephony service.
- If the user's domain is not listed in the **Locations** dialog box, try specifying the domain in **Enter the object names to select** by using the following format:

  *domain\username*

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable a telephony server
Add a server to the Telephony console
Delete a server from the Telephony console
Change the Telephony service logon account for a telephony server
Update information for a telephony server

# Change the Telephony service logon account for a telephony server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To change the Telephony service logon account for a telephony server**

1. Open Telephony.
2. In the console tree, click the server you want to manage.

   **Where?**

   - Telephony/*server*

3. On the **Action** menu, click **Properties**.
4. On the **Setup** tab, in **Account Information**, click **Choose User**.
5. In the **Select User** dialog box, click **Locations**, and then click the domain where the account is located.
6. In **Enter the object name to select**, type the name of the account, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- If the user's domain is not listed in the **Locations** dialog box, try specifying the domain in **Enter the object name to select** by using the following format:

  *domain\username*

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable a telephony server
Delete a server from the Telephony console
Designate telephony administrators
Update information for a telephony server

# Update information for a telephony server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To update information for a telephony server**

1. Open Telephony.
2. In the console tree, click the server you want to manage.

   **Where?**

   - Telephony/*server*

3. In the **Action** menu, click **Refresh**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- This procedure updates the status of lines and phones, displays new lines and phones, and shows changes made by others administering the Telephony service.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable auto updates for telephony devices
Enable or disable a telephony server
Delete a server from the Telephony console
Designate telephony administrators
Change the Telephony service logon account for a telephony server

# Manage Telephony Service Providers on a Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Manage telephony service providers on a server**

- Add a telephony service provider on a telephony server
- View the status of a telephony service provider's lines or phones
- Export telephony service provider information to a text file
- Enable auto updates for telephony devices
- Configure telephony service providers on a telephony server

# Add a telephony service provider on a telephony server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To add a telephony service provider on a telephony server**

1. Open Telephony.
2. In the console tree, click the server you want to manage.

   **Where?**

   - Telephony/*server*

3. In the **Action** menu, click **Manage Providers**.
4. In the **Telephony Providers** dialog box, click **Add**.
5. In the **Add Provider** dialog box, under **Available telephony providers**, select the telephony service provider that you want to add, and then click **Add**.

   If the telephony service provider you want to add does not appear in this list, you must install the appropriate driver files for the telephony service provider. For more information, see Notes.

6. In the **Telephony Providers** dialog box, review the list of installed telephony service providers to confirm that the provider was added to the list, and then click **Close**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must be designated as a telephony administrator (for information about designating telephony administrators, see Related Topics). If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- After you add a telephony service provider, the Telephony console might not display the lines and phones supported by the service provider until after the computer is restarted.
- If the telephony driver is not listed in the **Available telephony drivers** list in the **Add Provider** dialog box, you must install the driver files in the *Systemroot*\system32 folder. Usually you can use the setup program on a telephony device manufacturer's installation disk. For additional information, see the manufacturer's installation instructions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a telephony service provider
Add a telephony service provider on a telephony server
Export telephony service provider information to a text file

View the status of a telephony service provider's lines or phones
Enable auto updates for telephony devices
Designate telephony administrators

# View the status of a telephony service provider's lines or phones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To view the status of a telephony service provider's lines or phones**

1. Open Telephony.
2. In the console tree, click the service provider you want to manage.

**Where?**

- Telephony/*server/service provider*

1. Click **View**.
2. Click **Lines** or **Phones**.

**Notes**

- If the telephony server is not enabled, performing this task does not require administrative credentials. Any user who logs on to the telephony server can view the status of the lines and phones for each telephony service provider on the server. In addition, that user can view the names of all users who are assigned to lines or phones on the server.
- If the telephony server is enabled, to view the status of lines or phones for each telephony service provider on the server you must log on to the telephony server as one of the following:
  - A member of the local Administrators group on the telephony server. (If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.) For more information, see Default local groups and Default groups.
  - A telephony administrator. (For more information about designating telephony administrators, see Related Topics.)
  - A user who is assigned to a line or phone on the telephony server.

  Members of the local Administrators group on the telephony server and telephony administrators can view the names of all users who are assigned to lines or phones on the server. However, a user who is assigned to a line or phone on the server cannot view the names of other users who are assigned to lines or phones on the server. For information about enabling a telephony server, see Related Topics.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Add a telephony service provider on a telephony server
Export telephony service provider information to a text file
Enable auto updates for telephony devices
Update information for a telephony server
Designate telephony administrators

# Export telephony service provider information to a text file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To export telephony service provider information to a text file**

1. Open Telephony.
2. In the console tree, click the service provider you want to manage.

   **Where?**

   - Telephony/*server*/*service provider*

3. On the **Action** menu, click **Export List**.

## Notes

- If the telephony server is not enabled, performing this task does not require administrative credentials. Any user who logs on to the telephony server can export telephony service provider information to a text file. The information that is exported will include a list of telephony lines and phones on the telephony server and the names of users who are assigned to each of the telephony lines or phones.
- If the telephony server is enabled, to view the status of lines or phones for each telephony service provider on the server you must log on to the telephony server as one of the following:
  - A member of the local Administrators group on the telephony server. (If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.)
  - A telephony administrator. (For more information about designating telephony administrators, see Related Topics.)
  - A user who is assigned to a line or phone on the telephony server.

  Members of the local Administrators group on the telephony server and telephony administrators can export the names of users who are assigned to each of the telephony lines or phones on the server. However, a user who is assigned to a line or phone on the server cannot export the names of other users who are assigned to lines or phones on the server. For information about enabling a telephony server, see Related Topics.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- This procedure exports the list of telephony lines or phones to a text file.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Designate telephony administrators
Enable or disable a telephony server

# Enable auto updates for telephony devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable auto updates for telephony devices**

1. Open Telephony.
2. In the console tree, click the server you want to manage.

   **Where?**

   - Telephony/*server*

3. On the **Action** menu, click **Properties**, and then click the **Auto Updates** tab.
4. Select the **Enable device status auto updates** check box.
5. Next to **Interval**, type the **Hours** and **Minutes** between automatic updates.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- The action enabled by this procedure updates only the status information for lines and phones.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a telephony service provider on a telephony server
Export telephony service provider information to a text file
View the status of a telephony service provider's lines or phones
Enable auto updates for telephony devices
Update information for a telephony server

# Configure telephony service providers on a telephony server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To configure telephony service providers on a telephony server**

1. Open Telephony.
2. In the console tree, click the server you want to manage.

   **Where?**

   - Telephony/*server*

3. On the **Action** menu, click **Manage Providers**.
4. Click the telephony service provider you want to configure, and then click **Configure**.

   If **Configure** is not available, then the telephony service provider has no configurable properties.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must be designated as a telephony administrator (for information about designating telephony administrators, see Related Topics). If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.For more information, see Default local groups, Default groups, and Using Run as.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- The settings you can change depend on the telephony service provider.
- You can also configure telephony service providers on the local computer using **Phone and Modem Options** in Control Panel. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure a telephony service provider
Designate telephony administrators

# Manage Telephony Clients and Users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Manage telephony clients and users**

- Assign a telephony user to a line or phone
- Remove users from telephony lines or phones
- Specify telephony servers on a client computer
- Add a telephony service provider
- Configure a telephony service provider
- Specify an H.323 gateway or proxy
- Specify an H.323 gatekeeper
- Specify H.323 call control options
- Enable a Windows 98 or Windows Millennium Edition telephony client

# Assign a telephony user to a line or phone

Updated: May 4, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To assign a telephony user to a line or phone**

1. Open Telephony.
2. In the console tree, click the service provider you want to manage.

   **Where?**

   - Telephony\server\service provider

3. In the details pane, under **Line Name** or **Phone Name**, click the line or phone.
4. On the **Action** menu, click **Edit Users**.
5. In the **Edit Users** dialog box, click **Add**.
6. In the **Select Users** dialog box, click **Locations**, and then click the domain where the user is located.
7. In **Enter the object names to select**, enter the name of the user to be assigned a line or phone, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must be designated as a telephony administrator (for information about designating telephony administrators, see Related Topics). If the computer is joined to a domain, members of the Domain Admins group can perform this procedure only if the Domain Admins group is a member of the Administrators group. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- By default, members of groups with administrative credentials for a telephony server already have access to all telephony lines or phones on that server and therefore do not need to be assigned lines or phones on the server. Members of the local Administrators group can access all lines or phones on that server from any other computer in the domain. By default, members of the Domain Admins group can access all telephony lines and phones on all servers in the domain.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- After you perform this procedure, telephony client programs run by the user will be able to use the line or phone on the telephony server.
- A user added for a phone or line must be in the same domain as the telephony server or in a domain that has a two-way trust relationship with the domain containing the server.
- If the user's domain is not listed in the **Locations** dialog box, try specifying the domain in **Enter the object names to select** by using the following format:

  *domain\username*

- Before a user can use the telephony line or phone, the telephony server must be specified on the client computer. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable a Windows 98 or Windows Millennium Edition telephony client
Remove users from telephony lines or phones
Assign a telephony user to a line or phone
Specify telephony servers on a client computer
View the status of a telephony service provider's lines or phones
Designate telephony administrators

# Remove users from telephony lines or phones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To remove users from telephony lines or phones**

1. Open Telephony.
2. In the console tree, click the service provider you want to manage.

   **Where?**

   - Telephony/*server*/*service provider*

3. In the details pane, under **Line Name** or **Phone Name** click the line or phone.
4. In the **Action** menu, click **Edit Users**.
5. In **Assigned Users**, click a user, and then click **Remove**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must be designated as a telephony administrator (for information about designating telephony administrators, see Related Topics). If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Telephony, click **Start**, click **Run**, and then type **tapimgmt.msc**.
- After you perform this procedure, telephony client programs run by the user will be unable to use the line or phone on the telephony server.
- You can remove multiple users from telephony lines or phones at one time by using the **Tsecimp** command. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View the status of a telephony service provider's lines or phones
Tsecimp
Designate telephony administrators

# Specify telephony servers on a client computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To specify telephony servers on a client computer**

1. Open a Command Prompt window.
2. Type **tcmsetup /c** *server1 server2* **...**

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- The client must be in the same domain as the telephony server or in a fully trusted domain when you run **tcmsetup**.
- In order for TAPI to function correctly, you must add the Windows Remote Service Provider to the client computer or run the **tcmsetup** command. If your telephony clients run Windows XP or a Windows Server 2003 operating system, and your telephony servers run Windows Server 2003, then add the Windows Remote Service Provider to your clients. (The Windows Remote Service Provider enables clients to detect telephony servers.) However, if your telephony clients or servers run Windows NT 4.0 or Windows 2000, you must run **tcmsetup**, as described in this procedure, instead. For more information about adding telephony service providers and using **tcmsetup**, see Related Topics.
- In the **tcmsetup** command-line syntax, *server1 server2* **...** represents the names of one or more telephony servers. Because **tcmsetup** replaces any previous server list, you should include every telephony server that you want to make available to the client when you run **tcmsetup**.
- You may have to restart the Telephony service for this change to take effect. For more information, see Related Topics.
- Before client users can dial out using the phone lines on the telephony server, the telephony server administrator must assign the users to the phone lines. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify telephony servers on a client computer
Remove users from telephony lines or phones
Assign a telephony user to a line or phone
Specify telephony servers on a client computer
Tcmsetup
Add a telephony service provider
Start, stop, pause, resume, or restart a service
Run a program with administrative credentials
Cmd

# Add a telephony service provider

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To add a telephony service provider**

1. Open Phone and Modem Options in Control Panel.

2. On the **Advanced** tab, click **Add**.

3. In the **Add Provider** dialog box, under **Telephony providers**, click the service provider, then click **Add**.

   If the telephony service provider that you want to add does not appear in this list, you must install the appropriate driver files for the telephony service provider. For more information, see Notes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

- In order for TAPI to function correctly, you must add the Windows Remote Service Provider to the client computer or run the **tcmsetup** command. If your telephony clients run Windows XP or a Windows Server 2003 operating system, and your telephony servers run Windows Server 2003 , then add the Windows Remote Service Provider to your clients. (The Windows Remote Service Provider enables clients to detect telephony servers.) However, if your telephony clients or servers run Windows NT 4.0 or Windows 2000, you must run **tcmsetup**, as described in this procedure, instead. For more information about adding telephony service providers and using **tcmsetup**, see Related Topics.

- This procedure only adds telephony service providers that are already installed on your hard drive in the *Systemroot*\system32 folder. Use the setup or installation program provided on the manufacturer's disk to install the driver files for a new telephony service provider.

- A manufacturer's setup program normally adds the new telephony service provider to the list on the **Advanced** tab of **Phone and Modem Options**. If a telephony service provider is not listed on the **Advanced** tab after you run the manufacturer's setup program, use this procedure to add the service provider. If it is not listed in the **Add Provider** dialog box, refer to the manufacturer's documentation for further instructions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a telephony service provider
Configure a telephony service provider
Specify an H.323 gateway or proxy
Enable a Windows 98 or Windows Millennium Edition telephony client
Specify telephony servers on a client computer
Add a telephony service provider on a telephony server
Tcmsetup

# Configure a telephony service provider

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To configure a telephony service provider**

1. Open Phone and Modem Options in Control Panel.
2. On the **Advanced** tab, under **Providers**, click the telephony service provider you want to configure.
3. Click **Configure**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- The settings you can change depend on the service provider. If **Configure** is unavailable, then the service provider does not have any settings that you can change.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a telephony service provider
Configure a telephony service provider
Specify an H.323 gateway or proxy
Enable a Windows 98 or Windows Millennium Edition telephony client
Specify telephony servers on a client computer
Add a telephony service provider on a telephony server
Configure telephony service providers on a telephony server

# Specify an H.323 gateway or proxy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To specify an H.323 gateway or proxy**

1. Open Phone and Modem Options in Control Panel.
2. On the **Advanced** tab, in **Providers**, click **Microsoft H.323 Telephony Service Provider**, and then click **Configure**.
3. In **Configure H.323 Service Provider**, select one of the following:
   - To specify an H.323 gateway, select the **Use H.323 gateway** check box, and then type the computer name or IP address for the gateway.
   - To specify an H.323 proxy, select the **Use H.323 proxy**, and then type the server name or IP address for the proxy.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- An H.323 gateway connects the IP network with a switched service network, such as the public switched telephone network (PSTN). Your telephony administrator can provide the correct name or IP address to enter here.
- A proxy server acts as a firewall or security barrier between your intranet and the Internet, preventing other people on the Internet from gaining access to confidential information on your internal network or computer. Your telephony administrator can provide the correct name or IP address to enter here.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure telephony service providers on a telephony server
Add a telephony service provider
Configure a telephony service provider
Specify an H.323 gateway or proxy
Enable a Windows 98 or Windows Millennium Edition telephony client
Specify telephony servers on a client computer

# Specify an H.323 gatekeeper

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To specify an H.323 gatekeeper**

1. Open Phone and Modem Options in Control Panel.
2. On the **Advanced** tab, in **Providers**, click **Microsoft H.323 Telephony Service Provider**, and then click **Configure**.
3. Under **Calling through a gatekeeper using TAPI 3.1**, select the **Use H.323 gatekeeper** check box, and then type the computer name or IP address (for example, 10.255.255.255) of the gatekeeper that you want to use.
4. Under **Calling through a gatekeeper using TAPI 3.1**, register with the gatekeeper by specifying one or both of the following:
   - To use a telephone number, select the **Log on using my phone number** check box, and then type the telephone number that you want to use.
   - To use an account name, select the **Log on using my account name** check box, and then type the name of the account you want to use.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- An H.323 gatekeeper serves as the central point for all calls within its zone of control. Gatekeepers perform address translation, admissions control, bandwidth control, and zone management functions. In addition, gatekeepers can perform call control, call authorization, bandwidth management, and call management functions.
- If you select the **Use H.323 gatekeeper** check box, any gateway and proxy settings that are specified in the **Configure H.323 Service Provider** dialog box will not be used.
- To use an H.323 gatekeeper, you must register with the gatekeeper by specifying a valid telephone number, account name, or both. Registering with a gatekeeper allows you to receive calls at this computer. If you are not sure which telephone number or account name to use to register with the H.323 gatekeeper, contact your telephony administrator or service provider.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify an H.323 gateway or proxy
Specify H.323 call control options
Configure telephony service providers on a telephony server
Add a telephony service provider

Configure a telephony service provider
Specify an H.323 gateway or proxy
Enable a Windows 98 or Windows Millennium Edition telephony client
Specify telephony servers on a client computer

# Specify H.323 call control options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To specify H.323 call control options**

1. Open Phone and Modem Options in Control Panel.
2. On the **Advanced** tab, in **Providers**, click **Microsoft H.323 Telephony Service Provider**, and then click **Configure**.
3. Under **H.323 call control**, in **H.323 call alerting timeout**, type the amount of time, in milliseconds, an incoming call will ring before it is dropped.
4. Under **H.323 call control**, in **H.323 call listening port**, type the number of the port on which this H.323 client will listen for incoming calls, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify an H.323 gateway or proxy
Specify an H.323 gatekeeper
Configure telephony service providers on a telephony server
Add a telephony service provider
Configure a telephony service provider
Specify an H.323 gateway or proxy
Enable a Windows 98 or Windows Millennium Edition telephony client
Specify telephony servers on a client computer

# Enable a Windows 98 or Windows Millennium Edition telephony client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To enable a Windows 98 or Windows Millennium Edition telephony client**

1. On a Windows 98 or Windows Millennium Edition telephony client computer, click **Start**, point to **Settings**, and then click **Control Panel**.
2. Double-click **Network**.
3. (Optional) If **Client for Microsoft Networks** is not in the list of installed network components, click **Add**, and follow the instructions to add it.
4. Click **File and Print Sharing**, confirm that both check boxes are selected, and then click **OK**.
5. Click the **Access Control** tab, click **User-level access control**, then click **OK**.

### Note

- After you complete this procedure, you need to specify a telephony server on the client. To make the telephony resources on the server available to the client users, the server administrator needs to assign the users to telephones or lines on the server. Users must log on to a trusted domain account. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Telephony Resources
Assign a telephony user to a line or phone
Specify telephony servers on a client computer
Designate telephony administrators

# Change How Phone Calls Are Dialed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Change how phone calls are dialed**

- Create a new location
- Change dialing properties for a location
- Specify area code rules
- Specify numbers to dial an outside line
- Create a new calling card
- Modify an existing calling card
- Use a calling card for long distance dialing
- Specify a long distance carrier

# Create a new location

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To create a new location**

1. Open Phone and Modem Options in Control Panel.
2. On the **Dialing Rules** tab, click **New**.
3. On the **General** tab, in **Location name**, type the name of the new location.
4. Specify the characteristics of your telephone line by providing additional information on the **General** tab, the **Area Code Rules** tab, and the **Calling Card** tab.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- See Related Topics for more information about locations or specifying specific elements of a location.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Create a new calling card
Specify area code rules
Change dialing properties for a location

# Change dialing properties for a location

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To change dialing properties for a location**

1. Open Phone and Modem Options in Control Panel.
2. On the **Dialing Rules** tab, click the location to modify, and then click **Edit**.
3. Click a tab to make changes to the **General** settings, **Area Code Rules**, or **Calling Card** for the location.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- See Related Topics for more information about locations or changing specific elements of a location.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Modify an existing calling card
Create a new calling card
Create a new location
Specify area code rules

# Specify area code rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To specify area code rules**

1. Open Phone and Modem Options in Control Panel.
2. On the **Dialing Rules** tab, click the location to modify, and then click **Edit**.

   To create a new location, click **New**.
3. Click the **Area Code Rules** tab.
4. To add a new area code rule, click **New**.

   To edit an existing rule, click the rule and then click **Edit**.
5. In **Area code**, enter the area or city code.
6. In **Prefixes**, select one of the following:

   - To apply the rule to all prefixes in the area code, select the **Include all the prefixes within this area code** check box.
   - To apply the rule to specific prefixes, select the **Include only the prefixes in the list below** check box, click **Add**, enter one or more prefix, and then click **OK**.
7. In **Rules**, select one or both of the following:

   - To dial a number before the phone numbers that contain these prefixes, select the **Dial** check box, and then type the digits to dial.
   - To dial the area code before the phone number, select the **Include the area code** check box.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Change dialing properties for a location

# Specify numbers to dial an outside line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To specify numbers to dial an outside line**

1. Open Phone and Modem Options in Control Panel.
2. On the **Dialing Rules** tab, click the location to modify, and then click **Edit**.

   To create a new location, click **New**.
3. On the **General** tab, in **Dialing rules**, enter the numbers required to reach an outside line for local and long-distance calls.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Do not use this procedure to specify the numbers required before dialing to a different city, area, or country/region, such as 1, 0, 00, or 011. The numbers specified here are to dial local or long distance calls from within a private branch exchange (PBX), such as used by a company or other organization. The dialing rules for a location automatically insert the number required to call a number with a different city, area, or country/region code.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Create a new location
Change dialing properties for a location

# Create a new calling card

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To create a new calling card**

1. Open Phone and Modem Options in Control Panel.
2. On the **Dialing Rules** tab, click the location to modify, and then click **Edit**.

   To create a new location, click **New**.
3. On the **Calling Card** tab, click **New**.
4. On the **General** tab, enter the following:

| Value | Description |
|---|---|
| **Calling card name** | The name you use to identify the card. |
| **Account number** | The account number provided by the calling card company for the calling card. For some calling cards this is the home or business phone number associated with the account. |
| **Personal Identification Number (PIN)** | Any personal number required to secure your account against unauthorized use. Often a four digit number, this is usually dialed after the account number when you place a call. Some calling cards do not use a PIN. |

5. On the **Long Distance** tab, the **International** tab, and the **Local Calls** tab, type the phone number that is required to access the corresponding service. Include all the digits dialed after you access an outside line.

   - On the **Long Distance** tab, **Access number for long distance calls** is the phone number required to access the calling card service for domestic long distance calls.
   - On the **International** tab, **Access number for international calls** is the phone number required to access the calling card service for international long distance calls.
   - On the **Local Calls** tab, **Access number for local calls** is the phone number required to access the calling card service for local calls. Leave this blank to dial local numbers without using the calling card.
6. On the **Long Distance** tab, the **International** tab, and the **Local Calls** tab, in **Calling card dialing steps**, click the buttons to add the steps that are required to complete a call.

| Click | To dial |
|---|---|
| **Access Number** | The number to access the calling card company. This will dial the access number entered at the top of the tab. |
| **Account Number** | The account number that identifies your account, which was entered on the **General** tab. |
| **PIN** | Your personal identification number, which was entered on the **General** tab, usually dialed immediately after the account number. |

| | |
|---|---|
| **Destination Number** | The country/region code, area code, and phone number. |
| **Wait for Prompt** | Pauses dialing for a dial tone, voice message, or time interval. |
| **Specify Digits** | Any sequence of the digits 0-9, *, and #. |

7. To reposition any step in the order of calling card dialing steps, click the step, and then click **Move Up** or **Move Down**.

   To remove the step, click **Delete**.

### Important

- When you enter a calling card PIN number, this number is displayed in plaintext. Therefore, to enhance security, type this number in a private location whenever possible. If you are not in a private location, take care to ensure that this number is not exposed to another user.

### Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

- If you add **Access Number**, **PIN**, or **Account Number** to the **Calling card dialing steps** and the corresponding field is blank, the program will prompt you to enter the missing value when you click **OK** or **Apply**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Modify an existing calling card
Use a calling card for long distance dialing

# Modify an existing calling card

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To modify an existing calling card**

1. Open Phone and Modem Options in Control Panel.
2. On the **Dialing Rules** tab, click the location to modify, and then click **Edit**.
3. On the **Calling Card** tab, click the card type you want to modify, and then click **Edit**.
4. Click the **General** tab to change the calling card name, number, or your personal identification number (PIN).
5. Click the **Long Distance**, **International**, or **Local Calls** tab to change the access number or dialing steps for the corresponding type of call.

### Important

- When you enter a calling card PIN number, this number is displayed in plaintext. Therefore, to enhance security, type this number in a private location whenever possible. If you are not in a private location, take care to ensure that this number is not exposed to another user.

### Notes

- To perform this procedure, you must be the user who created the calling card. A calling card that is created by one user cannot be modified by a different user.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- If you add **Access Number**, **PIN**, or **Account Number** to the **Calling card dialing steps** and the corresponding field is blank, the program will prompt you to enter the missing value when you click **OK** or **Apply**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Change dialing properties for a location
Create a new calling card
Use a calling card for long distance dialing

# Use a calling card for long distance dialing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To use a calling card for long distance dialing**

1. Open Phone and Modem Options in Control Panel.
2. On the **Dialing Rules** tab, click the location to modify, and then click **Edit**.

   To create a new location, click **New**.
3. On the **Calling Card** tab, click the calling card you want to use.

## Notes

- To perform this procedure, you must be the user who created the calling card.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- If you use more than one calling card, you can create a different location for each one. Telephony programs typically allow you to change the location before dialing.
- If you dial digits to access a long distance carrier (a long distance operator), you can use a calling card to specify the dialing sequence. When you create the calling card for a long distance carrier, create a unique location to go with it.
- See Related Topics for procedures to create calling cards or locations.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Create a new calling card
Modify an existing calling card

# Specify a long distance carrier

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**To specify a long distance carrier**

1. Open Phone and Modem Options in Control Panel.
2. On the **Dialing Rules** tab, click **New**.
3. On the **General** tab, type a **Location name** to identify the long distance carrier, and specify the **Country/region**, **Area code**, and any **Dialing rules**.
4. On the **Calling Card** tab, click **New**, and on the **General** tab, in **Calling card name**, type a name to identify the calling card, such as the name of the long distance carrier.
5. On the **Long Distance** tab and the **International** tab, in the access number box, type the corresponding long distance operator number.
6. On the **Long Distance** tab and the **International** tab, in **Calling card dialing steps**, click **Access Number** and the additional steps for dialing long distance and international calls.
7. On the **Local Calls** tab, leave **Access number for local calls** blank, unless an access number is required, and in **Calling card dialing steps**, enter the steps for dialing local calls.

   To dial only the local phone number, click **Destination Number**, and clear the **Dial the area code** check box.
8. Click **OK** to save the new calling card.
9. On the **Calling Card** tab, in **Card Types**, make sure that the new calling card is selected to be used with the location, and then click **OK**.
10. Under **Locations**, leave this location selected if you want this to be the default long distance carrier for telephony programs, and then click **OK**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- This procedure specifies long distance carrier numbers (also called the long distance operator) required for domestic or international long distance calling. This procedure is not normally required to use the default domestic and international long distance operators for a country/region. The country/region selection for a location normally enables programs to use the default numbers.
- The location selected on the **Dialing Rules** tab of **Phone and Modem Options** is the default location for dialing. If you create more than one location, some telephony programs allow you to choose a different location before dialing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dialing rules overview
Create a new calling card
Use a calling card for long distance dialing
Create a new location

# Telephony Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Concepts**

This section provides general background information about Telephony Application Programming Interface (TAPI):

- Telephony overview
- TAPI and Active Directory
- Telephony service providers
- Understanding Telephony
- Using Telephony
- Telephony Resources

# Telephony overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Telephony overview**

Telephony Application Programming Interface (TAPI) integrates telecommunications with the computer. TAPI supports both traditional (or PSTN) and IP telephony to provide voice, data, and video communication. Supported hardware includes sound and video cards, modems, ISDN lines, ATM networks, and cameras. With this, you can communicate over direct connections to local computers, phone lines, LANs, WANs, and the Internet.

In addition to making and receiving calls, programs can use TAPI to provide enhanced telephony features such as caller ID, call routing, voice mail, and video conferencing. Communication programs may identify the caller, recall and display caller information, and even prioritize or route a call, based on customer information.

With TAPI, there is a common standard so that communication programs can control telephony functions for data, fax, and voice calls.

- TAPI manages all signaling between a computer and a telephone network, including basic functions such as dialing, answering, and hanging up a call.
- TAPI includes supplementary functions that are found in PBX, ISDN, and other phone systems, such as hold, transfer, conference, and call park.
- TAPI provides access to features that are specific to certain service providers, with built-in extensibility to accommodate additional telephony features and networks.
- TAPI allows users to make audio and video telephone calls over IP networks using the H.323 protocol and IP multicast video conferencing.
- TAPI includes Quality of Service (QoS) support to improve conference quality and network manageability.

These capabilities enable TAPI to unify IP and traditional telephony and support programs that effectively work the same way over intranets and the Internet as they do over the traditional telephone network. Furthermore, the current version of TAPI provides compatibility with previous versions to support existing telephony programs.

HyperTerminal and Network Connections are communication programs that use TAPI and that are provided with the Microsoft® Windows Server 2003 family. Fax Services, another communication program that uses TAPI, is available on an optional basis. You can also install other Microsoft or non-Microsoft programs to provide additional telephony support. The common interface that TAPI provides enables different communication programs to work together and share communication devices.

# TAPI and Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**TAPI and Active Directory**

Active Directory is the extensible and scalable directory service for Windows Server 2003 operating systems. TAPI-based programs use Active Directory to facilitate H.323-based IP telephony calls and multicast conferencing.

Active Directory publishes the locations of all TAPI application directory partitions on the network and provides the directory store for TAPI program-specific data. In Active Directory, application directory partitions store and replicate static or dynamic program-specific data. TAPI-based programs can use TAPI application directory partitions to store and retrieve information to facilitate H.323-based IP telephony calls and multicast conferencing.

TAPI provides transparent translation from user or computer names to IP addresses. When you place an H.323-based IP telephony call, you can specify a computer name or a user name, rather than an IP address, because TAPI application directory partitions publish user names, computer names, and IP addresses, and TAPI can translate user or computer names to IP addresses.

For multicast conferencing, TAPI application directory partitions publish information about active users and IP multicast conferences in progress. Information about active users includes the users' names, computer names, and IP addresses. Information about conferences in progress includes the conferences' names, descriptions, times, and originators.

You can create TAPI application directory partitions in one of two ways:

- By using the **Tapicfg** command-line tool. You can also use this command to remove or display TAPI application directory partitions and to set default TAPI application directory partitions. For more information, see Tapicfg.

- By selecting the **Typical setup for a first server** option when running the Configure Your Server Wizard to configure the first domain controller in a domain. When you select this option, the **Tapicfg** command runs transparently to create a TAPI application directory partition called MSTAPI.*DomainName*. This application directory partition is set as the default for the domain. For more information about the prerequisites for using this option, see Typical setup for a first server.

For more information about application directory partitions, see Application directory partitions. For more information about the H.323 IP protocol and the telephony service provider that supports it, see H.323 IP telephony calls.

**Note**

- If Internet Locator Service (ILS) is needed to support certain programs, you can continue to use ILS servers in your organization because telephony clients running Windows Server 2003 operating system can query either ILS servers or TAPI application directory partitions.

# Telephony service providers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Telephony service providers**
The following telephony service providers are supplied with this version of the Windows operating system.

**HID Phone TAPI Service Provider**
The Human Interface Device (HID) Phone TAPI Service Provider allows users of TAPI-based programs to make and receive telephone calls by using universal serial bus (USB) telephone handsets connected to their computers. This service provider is installed with Windows Server 2003 operating systems, and it requires no configuration.

**NDIS Proxy TAPI Service Provider**
The Network Driver Interface Specification (NDIS) Proxy TAPI Service Provider enables TAPI to control NDIS 5 devices. Routing and Remote Access uses this telephony service provider to set up network connections across NDIS WAN devices for My Network Places. This service provider is installed with Windows XP and Windows Server 2003 operating systems. To map addresses for the different types of calls received by the NDIS Proxy TAPI Service Provider on the local computer, see Configure a telephony service provider.

**TAPI Kernel-Mode Service Provider**
The TAPI Kernel-Mode Service Provider allows TAPI to control NDIS 4 devices. Network Connections uses this telephony service provider to set up network connections across NDIS WAN devices for My Network Places. This service provider is installed with Windows XP and Windows Server 2003 operating systems, and it requires no configuration.

**Unimodem 5 Service Provider**
The universal modem driver (Unimodem) provides services for data, fax, and voice modems so that users and program developers will not have to learn or maintain modem AT commands to dial, answer, and configure modems. Unimodem performs these tasks by using files that specify the AT command sets and expected responses for modems. These files usually have an .inf extension. Unimodem supports both data and voice, so a modem can be used for data and fax functions, as a speaker phone, as an answering machine, and so on. This service provider is installed with Windows XP and Windows Server 2003 operating systems.

**Windows Remote Service Provider**
The Windows Remote Service Provider enables a client computer to access and control remote telephony devices connected to a server computer, such as telephones and phone lines. The remote service provider makes the client-to-server link transparent to programs.




The Windows Remote Service Provider resides on the client computer, where it appears as a telephony service provider. When a program makes a TAPI request, the service provider interprets the request and sends it across the computer network to the Telephony service running on the server. The program then uses the telephony service provider for the telephone device installed on the server. By using this solution, clients can share a telephony resource over the network.

The Windows Remote Service Provider is installed with TAPI 2.2 on Windows 2000, Windows Millennium Edition, Windows XP, and Windows Server 2003 operating systems. On Windows 98 and Windows NT Service Pack 4, the Windows Remote Service Provider is included with TAPI 2.1. Windows NT 4.0 and Windows 95 clients need to upgrade to TAPI 2.1.

**H.323 TAPI Service Provider**

The H.323 TAPI Service Provider allows TAPI-supported programs to engage in multimedia calls with any H.323-compliant terminal on the local area network. It is also designed to work well with standard phone or Internet phone gateways. This service provider supports programs that enable computer users to interconnect and communicate with other people over IP-based networks. This service provider is installed with Windows XP and Windows Server 2003 operating systems.

**Multicast Conference TAPI Service Provider**

The Multicast Conference TAPI Service Provider supports IP multicast video conferencing for telephony clients. Using IP multicast, users participate in efficient multicasting, sending data to a single conference IP address instead of copying data to all participants in a conference. This service provider is installed with Windows XP and Windows Server 2003 operating systems.

### Note

- For information about configuring telephony service providers, see Configure telephony service providers on a telephony server.

# Understanding Telephony

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Understanding telephony**

This section covers:

- General Telephony Concepts
- IP Telephony
- IP Multicast Video Conferencing

# General Telephony Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**General telephony concepts**

This section covers:

- Introduction to telephony
- PSTN telephony
- PSTNs, PBXs, and computer networks
- How TAPI supports programs and hardware

# Introduction to telephony

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Introduction to telephony**
Telephony Application Programming Interface (TAPI) provides a uniform set of commands for any supported telephony device that is connected to your computer.

**How TAPI works**
When you use a Windows Server 2003 family program to send faxes, connect to a telephone, make a call using IP, join a conference, or perform other TAPI-supported activities, there are three layers of software that enable you to use a telephony device: a program, TAPI, and a telephony service provider.




- A program enables you to make phone calls, send and receive data or faxes, or join conferences. Microsoft and independent software vendors provide programs that incorporate TAPI functionality. Examples include HyperTerminal, which is included with Windows Server 2003 operating systems. Fax Services, another communication program that uses TAPI, is available on an optional basis.

- TAPI provides telephony functions for programs, such as dialing.

- A telephony service provider translates the commands for a telephony device or telephony protocol. Telephony service providers for modems and several telephony protocols are installed with Windows Server 2003 operating systems, and others are provided by independent hardware vendors.

**How TAPI works with telephony hardware**
In place of proprietary solutions that are difficult and expensive to implement, TAPI provides a single set of commands allowing hardware and software to come from different sources. When a telephony program requests an action, TAPI determines which telephony service provider supports the device, and the service provider sends the appropriate commands to the device.

For example, the Unimodem (Universal Modem) 5 Service Provider is installed with Windows Server 2003 operating systems. When an program initiates the dialing of a phone number, TAPI sends the appropriate commands to the Unimodem 5 Service Provider, and it sends the correct sequence of commands for dialing the number on the particular modem attached to the computer. The program does not need to know anything about the modem.

With TAPI, the Windows Server 2003 operating system provides a telephony platform. The Windows Server 2003 operating system is the operator that interconnects computer and telephony resources to direct and manage calls. TAPI-based software can provide solutions ranging from complete small office systems or home office systems to PBX front ends that integrate the computer and telephone networks of an organization.

**How TAPI makes resources available to users**
Network users make calls using lines. They may also use phones for call management and control. These are supported by the telephony service providers. A telephony administrator makes telephony resources on a server, such as lines and phones, available to users.

**Lines**
In telephony, a line is any media stream (data, voice, or video) supported by a telephony service provider. For a modem, the

line represents an actual telephone line. For other devices, it can represent another media stream, such as a gateway to the Internet. Calls are made using a line provided by a telephony service provider. A single line can support more than one call at a time, such as with call waiting and conferencing. The telephony service provider determines how physical devices are represented by telephony lines.

**Phones**
A phone in TAPI represents terminal equipment that is supported by a telephony service provider. In TAPI, a "phone" closely models a physical telephone, but may not be actual hardware. It includes the concepts of a hookswitch, handset, speaker, microphone, display, lamps, and programmable buttons. As with an actual telephone on your desk, a phone may be used without actually making calls to someone else. For example, you can use a telephone handset to listen to or retrieve voice mail messages.

A phone does not have to represent actual hardware; it can be emulated in software. For instance, this type of phone might be a program that uses TAPI or a telephony service provider. A physical phone is not required to manage a media stream and, although all telephony service providers support lines, not all telephony service providers support phones.

**How TAPI supports telephony**
The Windows Server 2003 family supports computer-telephony integration, from simple small office or home-based solutions to full-scale, enterprise-wide solutions. Low-level telephony solutions enable an individual to:

- Use a computer as an answering machine or voice mailbox.
- Use a computer to send and receive faxes.
- Use a microphone and speaker or a speakerphone connected to a multimedia computer as a phone system.
- Log on to a computer from a remote location.

These solutions are typically supported by a direct connection from a desktop computer to a phone line with a modem. Enterprise-wide solutions rely on a network and enable an organization to:

- Use a computer either as a multiline phone system or as a multifunction PBX controller.
- Provide a voice-mail system across your entire enterprise.
- Provide a fax system across your entire enterprise.
- Provide on-demand audio information services that allows callers to retrieve prerecorded or computer-generated text-to-speech information.
- Fax information to customers on demand, based on touch-tone keys or other input.
- Implement automatic, list-based dialing for telemarketing.
- Deliver database information related to a call (such as a customer profile or account information) at the same time it is switched to a desktop.
- Create a single, easy-to-use client program to manage all communications, such as voice, pager, e-mail, and conferencing.
- Provide dial-up access to the network.

On a network, a hardware interface adapter on a server connects the LAN to a public or private telephone network. The adapter is supported by a telephony service provider. Programs running on the server provide calling features such as conferencing, transferring, forwarding, holding, and automatic redialing. The server controls the calls going through the network and out onto the PBX.

Client programs provide voice, video, fax, and data communications for the user. As this section illustrates, the client programs call functions provided by the Telephony service on the client. The Telephony service on the client uses the Windows Remote

Service Provider to communicate over the network with the telephony service providers on the server, and the actual telephony services are provided by the telephony service providers managing telephony devices on the server.

# PSTN telephony

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**PSTN telephony**

The public switched telephone network (PSTN) is a traditional, circuit-switched network which is optimized for real-time voice communication. When you call someone, you close a switch by dialing and establish a circuit with the other party. The PSTN guarantees Quality of Service (QoS) by dedicating the circuit to your call until you hang up the telephone. Whether you and your party are talking or silent, you continue to use the same circuit until you hang up.

Telephony Application Programming Interface (TAPI) enables programs to communicate easily over the traditional telephone network. TAPI provides support for the direct connection to a PSTN network, automatic phone dialing, and interfaces for conference calling, voice mail, and caller ID.

TAPI helps make the Windows Server 2003 family a powerful and flexible platform for developing and using computer-telephony integration (CTI) programs. Programs can build on TAPI's client-server support to provide a graphical interface for system administration and enhanced services, such as voice mail, call queueing, follow-me call forwarding, telephony-computer integration, and speech recognition. Besides enabling individual programs to provide telephony services, TAPI manages the telephony devices, so more than one program using a line can remain active simultaneously. One program can wait for a call while another is dialing out.

In a client-server environment, telephony can be managed like any other network service. You can specify the lines and phones available to particular users and use domain security to control access to telephony resources. Telephony service providers and all stored parameters can be updated across a LAN, making it easy to set up, use, and manage resources without regard to physical location.

# PSTNs, PBXs, and computer networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**PSTNs, PBXs, and computer networks**

In telephony, desktop computers interact with telephony hardware. Most PBXs implement the computer-telephony integration (CTI) link with hardware that connects the PBX to one or more computers. Communication between a PBX and computers is rarely straightforward, because many vendors implement the CTI link differently. Furthermore, there is no standard for encoding data on the CTI link. As a result, computer-based telephony programs often need to understand the switch's proprietary language, which requires a separate driver. By using TAPI, you can use one API that many programs can use to request telephony services, and supply one translator (called a telephony service provider) for each PBX or other piece of telephony hardware.

Windows supplies several telephony service providers. The telephony service provider that enables telephony client-server functionality is the Microsoft Windows Remote Service Provider. Proprietary hardware may require additional telephony service providers, which are normally packaged with the hardware or available from the equipment manufacturer.

**Microsoft Windows client-server components**

TAPI includes the following components that support client-server CTI.

**Server-side components**

The following components are available on the telephony server.

- **Telephony console.** Enables administrators to administer and monitor a telephony server, lines, phones, and users.
- **Telephony service provider.** Translates between the generalized commands created by the Telephony service and the proprietary commands that the hardware uses. Provided by the hardware vendor.
- **Telephony service.** On the server side, communicates with the remote service provider running on the clients. Also communicates with telephony programs running on the server, such as the client manager program.

**Client-side components**

The following components must run on the client or appear on the client.

- **Telephony service.** On the client side, communicates with client-side telephony programs and acts as a link between those programs and the Windows Remote Service Provider.
- **Windows Remote Service Provider.** Communicates over the network with the Telephony service running on the telephony server computer.

Note that these components do not include client-side TAPI-based programs. Some general purpose programs use TAPI. Specialized programs, such as those that manage call centers, are provided by independent software or hardware vendors.

# How TAPI supports programs and hardware

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**How TAPI supports programs and hardware**

Telephony Application Programming Interface (TAPI) supports programs that share the same communication ports and devices. Communication programs use TAPI to communicate with and control the appropriate hardware. TAPI manages the way programs and their call traffic use ports and devices. Telephony service providers provide support for particular devices or classes of devices.




**Telephony service providers**

Telephony service providers are dynamic link libraries (DLLs) that translate the commands for a specific telephony device, carrying out the low-level tasks required to communicate over telephone and IP networks.

**How TAPI makes services available**

TAPI arbitrates requests from different communication programs, so they can run simultaneously while sharing communication ports and devices. Each service provider controls its hardware. For example, one Windows Server 2003 family program can use TAPI functions to make outgoing calls while other programs are waiting for incoming calls.

Although only one program controls any single call, the devices are available to other programs, and users do not have to close other programs that share a communication port.

**Supplied telephony service providers**

The following table lists the telephony service providers that are supplied with Windows Server 2003 operating systems:

| Telephony service providers | Used for |
|---|---|
| HID Phone | Human Interface Device (HID) telephone communications |
| H.323 | IP telephony multimedia communications |
| Multicast Conference | IP telephony multiparty conferencing |
| Unimodem 5 | Modems |
| NDIS Proxy | Connection-oriented network devices using NDIS 5 |
| TAPI Kernel-Mode | Connects network devices using NDIS 4 |
| Windows Remote | Client-server telephony |

For more information about telephony service providers, see Telephony service providers.

**Additional hardware support**
Other service providers can be installed, in addition to telephony hardware or software.

Hardware manufacturers normally supply device drivers for their hardware. To provide telephony services for the hardware, the manufacturers or independent software vendors provide telephony service providers to control the hardware. Each service provider supports at least one device specific to that category of communication, such as a fax board, an ISDN card, a telephone or a modem. Some service providers support several devices, or even an entire class of devices.

For more detailed information about how telephony works, see Telephony Resources.

# IP Telephony

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**IP telephony**
This section covers:

- IP telephony concepts
- H.323 IP telephony calls

# IP telephony concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**IP telephony concepts**
IP telephony refers to telephone communications over TCP/IP networks. In contrast to the PSTN, which consists of analog and digital signals over a circuit-switched network, IP telephony is packet switched. All information to be transmitted over the network is separated into data packets. Each packet has a header containing its source and destination, a sequence number, a block of data content, and an error-checking code. Routers and servers direct these packets over the network until they arrive at their destination. When the packets arrive, the sequence number is used to reassemble the packets in their original order. Unlike PSTN telephony, which dedicates a circuit for a telephone call, the data packets share a circuit with other transmissions.

**How IP telephony works**
IP telephony blends voice, video, and data by using TCP/IP as a common transport, effectively collapsing these separate channels into one. Clients using IP telephony employ either existing multimedia hardware or a phone plugged into a PSTN adapter.

IP telephony supports voice communication, voice and video mail, and video on demand. IP telephony also enables voice and video conferencing over the Internet and existing IP-based LANs and WANs.

**Benefits of IP telephony**
IP telephony provides the following benefits:

- IP telephony moves multimedia traffic over any network that uses IP. This offers users both flexibility in physical media (PSTN lines, xDSL, ISDN, leased lines, coaxial cable, satellite, and twisted pair) and flexibility of physical location.

- IP telephony lowers costs of existing services, such as voice and broadcast video, and broadens means of communication to include audio and video conferencing.

**TAPI support for IP telephony**
Two telephony service providers provide IP telephony support in the Windows Server 2003 family: the H.323 TAPI Service Provider and the Multicast Conference TAPI Service Provider. These service providers support the H.323 and IP multicast video conferencing standards and enable different IP telephony programs to interoperate with each other. Organizations can now take full advantage of the global reach of the Internet or the immediacy of a corporate intranet for real-time communication.

**Quality of Service for IP telephony**
Packet-switched networks such as the Internet do not dedicate a connection between parties, so they cannot guarantee Quality of Service (QoS) as the PSTN does. Historically, they have been used for programs, such as e-mail and file transfer, where variable QoS is acceptable.

Organizations can now design private networks to provide adequate bandwidth and redundant paths between switching points for high quality IP telephony. New software and hardware have improved QoS, and voice over IP on a WAN can be as good as standard PSTN calls, even if the network is congested. For more information about QoS, see QoS overview.

# H.323 IP telephony calls

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**H.323 IP telephony calls**

The H.323 TAPI Service Provider provides H.323 version 1.0 support for TAPI-based programs. H.323 is a comprehensive (ITU-T) standard. It provides for multimedia communication over computer networks that do not provide a guaranteed quality of service, such as IP-based networks and the Internet. The H.323 standard enables any H.323-compliant terminal to operate with another.

H.323 encompasses call control, multimedia streaming, and bandwidth management for point-to-point calls and multipoint conferences over packet-switched networks. It is used for peer-to-peer, two-way delivery of audio and video phone data. It is designed to interact well with PSTN or Internet phone gateways. For conferencing, it works well for a moderate number of participants.

H.323 support defines the format for audio and video information and how audio and video information is packaged for transmission over the network. Standard audio and video codecs encode and decode input and output from audio and video sources to be communicated between nodes on a network.

TAPI 3.1 in Windows XP and the Windows Server 2003 family support H.323 gateways, gatekeepers, and proxies. To implement H.323 gateways and proxies with the H.323 TAPI Service Provider, see Specify an H.323 gateway or proxy. To implement H.323 gatekeepers, see Specify an H.323 gatekeeper.

# IP Multicast Video Conferencing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**IP multicast video conferencing**

This section covers:

- How TAPI supports IP multicast video conferencing
- Security for multicast video conferencing

# How TAPI supports IP multicast video conferencing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**How TAPI supports IP multicast video conferencing**

The DHCP service provides Multicast Address Dynamic Client Allocation Protocol (MADCAP) support required for IP multicast video conferencing. DHCP provides unicast addresses for point to point communications between two networked computers. These addresses are unique and specific to a single network. In contrast, MADCAP provides multicast IP addresses, which are shared by many computers. These enable data to be forwarded to all members of a multicast group. The membership of a multicast group is dynamic, which means individual hosts can join or leave the group at any time.

MADCAP is part of the DHCP service, which is an optional component of the Windows Server 2003 family. Clients that are manually configured or use DHCP to obtain a unicast IP address lease can also use MADCAP to obtain multicast IP address configuration. Although MADCAP is packaged in the DHCP service, the DHCP and MADCAP services are independent of each other. You can use a MADCAP server to provide multicast addresses and run DHCP on other servers to provide DHCP services on your network.

MADCAP is built on a client/server model, where multicast conference hosts request addresses from a MADCAP server.

TAPI application directory partitions provide Active Directory directory services for IP multicast conferencing and H.323 telephony. They also publish IP multicast conferences and user IP mappings for H.323 IP telephony and enable users to join conferences and make IP telephone calls. For more information about TAPI application directory partitions, see Tapicfg and TAPI and Active Directory.

For information about installing IP multicast and video conferencing on your network, see Enabling IP multicast.

# Security for multicast video conferencing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Security for multicast video conferencing**

The TAPI 3.1 conference security system determines who can create, delete, view, and join conferences. Each published conference can be associated with a security descriptor. This specifies object access rights on a user or group basis.

By associating security descriptors with SDP conference descriptors, conference creators can specify who can enumerate and view conference announcements. User authentication is provided by the Windows security subsystem.

For more information, see Security.

# Using Telephony

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Using telephony**

This section covers:

- Dialing rules overview
- Canonical address format for phone numbers
- Telephony hardware requirements
- Setting up clients for IP telephony
- Setting up conference rooms
- Enabling IP multicast
- Eliminating acoustic echo
- Setting up the PSTN server and client
- Managing telephony from the command line

# Dialing rules overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Dialing rules overview**
Telephony API (TAPI) allows you to configure dialing rules for all telephony programs. If you run a TAPI-supported program before configuring dialing, the program normally prompts you for the minimum information required for dialing. To completely configure dialing rules on a computer, use **Phone and Modem Options** in Control Panel, as described in this topic.

**Dial locations**
A dial location is the set of parameters that is used to analyze telephone numbers and determine the correct sequence of numbers to dial. You can create as many dial locations as you need for different dialing requirements.

A dial location can correspond to an actual geographic location, such as an office or hotel room, where you need to include codes for dialing outside of the company or building. If your work takes you to different cities, countries, or regions, you can create a dial location for each area, city, or country/region code.

A dial location can also be created for dialing rules that are not necessarily associated with a geographic location. For example, you can create different dialing rules for different company accounts, calling cards, or long distance services.

Dial location information includes:

- Name
- Country or region
- Area (or city) code
- Dialing rules
- Area code rules
- Calling card information

**Naming and using dial locations**
Dial locations can be named anything that is helpful for remembering them later. The default dial location is the location selected in the **Phone and Modem Options** dialog box. Many telephony programs allow you to select a different dial location before dialing. When a telephony program dials a long distance number, it uses the dialing information provided by the selected dial location to complete the connection.

The available dial locations are listed on the **Dialing Rules** tab of **Phone and Modem Options** in Control Panel. On this tab, you can select the default dial location, create a new dial location, or edit any existing dial location. To create a new dial location, see Create a new location. If you want to edit an existing dial location, see Change dialing properties for a location.

**General dial location rules**
General dial location rules specify the name of a location, the country or region, and the local area code. They also enable you to access outside lines from a PBX for local and long distance calls, disable call waiting (so incoming calls do not interrupt connections), and specify either tone or pulse dialing.

The default rules for dialing local, long distance, and international calls are determined by the country or region that you have selected. Dialing rules specify the long distance operator and international prefix that is dialed before domestic long distance or international phone numbers.

**Area code rules**
Area code rules enable you to change the dialing patterns for calls to numbers in the same country or region, as described in Specify area code rules. The default area code rule for numbers in another area code often requires the program to dial the long distance operator, the area or city code, and then the phone number. For numbers in the same area code, the program

dials only the subscriber's phone number. The default rules are shown in the following table.

| Format and Example | For |
|---|---|
| *LongDistanceOperator* (*AreaOrCityCode*) *Prefix-SubscriberNumber*<br><br>1 (425) 555-0101 | Numbers in other area or city codes |
| *Prefix-SubscriberNumber*<br><br>555-0101 | Numbers in the same area or city code |

Digits dialed in each of these categories may be different for different countries, regions, or locations.

Use area code rules to specify any of the following exceptions to the default rules:

| Format and Example | For |
|---|---|
| (*AreaOrCityCode*) *Prefix-SubscriberNumber*<br><br>(425) 555-0101 | Numbers in another area or city code. |
| (*AreaOrCityCode*) *Prefix-SubscriberNumber*<br><br>(425) 555-0101 | Numbers in the same area or city code. |
| *LongDistanceOperator*<br>*Prefix-SubscriberNumber*<br><br>1 555-0101 | Numbers in the same area or city code. |
| *LongDistanceOperator* (*AreaOrCityCode*)<br>*Prefix-SubscriberNumber*<br><br>1 (425) 555-0101 | Numbers in the same area or city code. |
| *LongDistanceOperator* (*AreaOrCityCode*)<br>*Prefix-SubscriberNumber*<br><br>01 (425) 555-0101 | Numbers in another city or area code, but the long distance operator is not the default for the country or region. |

If you use a different long distance operator than the country/region default for all area codes, it is much easier to use a calling card for the dial location instead of creating an area code rule for every area or city code that you dial. For more information, see "Calling cards" in this topic.

**Specifying more than one rule for an area code**

When different prefixes in an area code are dialed differently, you can specify more than one rule for the area code. Area code rules are applied from top to bottom in the list. If two rules apply to a particular prefix in an area code, the last rule in the list determines the dialing pattern for that prefix.

When you create a new area code rule, it is added to the bottom of the list. To apply multiple rules correctly then, create area code rules in this order:

1. Any default rule that should apply to all unspecified prefixes. (In this case, use the **Include all the prefixes within this area code** option). This rule should appear at the top of the list.

2. Area code rules that list specific prefixes. (Use the **Include only the prefixes in the list below** option).

Most prefixes in an area code follow the default dialing patterns shown here:

| Dial | For |
|---|---|
| *1* (*AreaOrCityCode*) *Prefix-SubscriberNumber* | Numbers in another area or city code. |
| *Prefix-SubscriberNumber* | Numbers in the same area or city code. |

When most prefixes in an area code do not follow the default dialing patterns, create one rule to specify the new default dialing pattern for the area code and create another rule to specify the few exceptions. This is easier than just creating one area code rule that lists the bulk of prefixes for an area code.

For example, suppose you dial most numbers in an area code in the following format:

(*AreaOrCityCode*) *Prefix-SubscriberNumber*

but three prefixes are dialed in this format:

1 (*AreaOrCityCode*) *Prefix-SubscriberNumber*

First, create an area code rule that applies to all prefixes in the area code and, in that rule under **Rules**, select only the **Include the area code** check box. This creates the default rule for the area code. It will apply to all the prefixes that are not specified by the second rule in the list.

Second, create an area code rule that lists the three prefixes that follow the second dialing pattern. In that rule under **Rules**, select both the **Dial** check box (confirm that the adjacent text box says *1*) and the **Include the area code** check box. This creates a rule that applies only to the three prefixes and should appear second in the list.

**Calling cards**

A dial location can use a calling card to specify the sequence of numbers to be dialed for a particular calling card or long distance service. You may also choose to use a calling card to specify a different long distance operator or international operator than the default operator for a country or region.

A calling card records the calling card number and any personal identification number (PIN). It defines the access numbers and dialing steps for local, long distance, and international calls. After you create a new calling card, you can use it with any dial location. To create or use a calling card, see Create a new calling card or Use a calling card for long distance dialing.

If you use two or more long distance services with different access numbers, you can create a calling card for each service and a different dial location to use each calling card. Telephony programs typically permit you to change the dial location before dialing a number. When you dial a number, select the dial location that corresponds to the long distance service that you want to use.

**Important**

- When you enter a calling card PIN number, this number is displayed in plaintext. Therefore, to enhance security, type this number in a private location whenever possible. If you are not in a private location, take care to ensure that this number is not exposed to another user.

**Specify a long distance carrier**

The long distance carrier number, also called the long distance operator, is dialed before the area code or country/region code to dial domestic or international long distance calls. The country/region selection for a dial location usually enables programs to dial the correct long distance carrier numbers for domestic and international calls. If your long distance carrier is not accessed with the default numbers for the country/region selection, then see Specify a long distance carrier.

# Canonical address format for phone numbers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Canonical address format for phone numbers**

The canonical address format is a universal phone number format recognized by the Windows Server 2003 Telephony API (TAPI). The format explicitly identifies the components of a phone number, which TAPI translates according to a country or region's dialing rules. The canonical address format is:

**+***Country/RegionCode* **(***AreaCode***)** *SubscriberNumber*

For example, this is how you would enter a number for a subscriber in the United States of America in canonical address format:

+1 (425) 555-0100

Some address book programs store numbers in canonical address format. If your computer is not dialing a number correctly, you may try saving or entering the number in this format.

**Components**

- **+**

  Indicates that the number is in canonical address format.

- *Country/RegionCode*

  The standard country/region code that identifies the country or region for a phone number. This contains one or more digits from 0 through 9. The country/region code is delimited by the space that follows it.

- **(***AreaCode***)**

  The area or city code for the phone number. This may contain one or more digits from 0 through 9. The area or city code is delimited by parentheses. This component is omitted for countries/regions that do not use area or city codes.

- *SubscriberNumber*

  The number for a phone subscriber. This contains one or more digits from 0 to 9, formatting characters, or the dialing control characters:

- A a B b C c D d P p T t W w * # ! @ $ ?

  The subscriber number should not contain the following characters:

- ( ) ^

  Formatting characters in the subscriber numbers are normally spaces, periods, and dashes. Use formatting characters to make a phone number easier to read. They do not affect dialing, because TAPI discards them.

# Telephony hardware requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Telephony hardware requirements**

To hook up a regular (PSTN) phone line to your local computer, you must have a voice modem. To send and receive faxes in addition to making voice or data calls, you must have a fax modem.

To use IP telephony service providers, your computer must have a sound card, video camera, speakers, and a microphone (or headset).

To support speech and hearing simultaneously, you must have a full-duplex sound card.

To include video communications, you must have a video capture card.

You may need to install other hardware for special telephony programs. For example, to connect trunk lines for a personal computer-based telephony server or cables to connect to a PBX system, you may need expansion cards. For additional information about your telephony hardware requirements, consult your vendor.

To verify hardware compatibility, see the Compatible Hardware and Software section at Support resources.

# Setting up clients for IP telephony

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Setting up clients for IP telephony**
Each client computer must run Windows XP or a Windows Server 2003 operating system and be equipped with a video camera, sound card, speakers, and microphone (or headset).

**Install Windows XP or a Windows Server 2003 operating system**
All of the Windows XP and Windows Server 2003 components required to support TAPI multicast conferencing on a client computer are included with these operating systems.

**Create computer and user accounts**
The client computer should be added to a domain. When the client computer is added to a domain, a computer account is automatically created for it. The user should also have a user account in a domain, and should log on to the computer with this account.

If computer or user accounts for the domain are not created, users will not be able to access TAPI application directory partitions to participate in multicast conferencing. If Internet Locator Service (ILS) servers are being used to support multicast conferencing, users will have to manually add the ILS servers to multicast conferencing programs. The multicast conferencing capabilities are otherwise still available.

**Install a full duplex audio sound card**
Windows XP and Windows Server 2003 operating systems support a wide range of audio devices. Multicast conferencing requires a sound card capable of full duplex audio functions. Such sound cards allow audio to be played back and recorded simultaneously. The manufacturer's documentation should describe the characteristics of your sound card.

**Verify full duplex audio**
If you are uncertain whether your system supports full duplex audio, install and start Sound Recorder and record a long message (about one minute). (Sound Recorder is an audio program that is available on an optional basis. Rewind the recording and start Sound Recorder again, so it appears in two places on the desktop. In the first instance of Sound Recorder, start playing back the recording, and then use the other instance to simultaneously record another message. If you can record and play back simultaneously, then your system supports full duplex audio.

**Eliminate acoustic echo**
Acoustic echo is a common problem with audio conferencing systems. You can eliminate acoustic echo by using an audio headset or by enabling acoustic echo cancellation (AEC). For more information, see Eliminating acoustic echo.

**Install the video camera**
Video cameras designed to connect to your computer capture and transfer video images to the computer. TAPI multicast conferencing components are able to acquire these images using WDM or Video for Windows interfaces that are built into the operating system.

Many of the older video cameras require the use of a video capture card that fits into an ISA or PCI slot inside your computer. Video capture cards connect to the camera with a cable, using connectors built into the back of the capture card. Windows XP and Windows Server 2003 operating systems provide support for many such cameras.

The newer generation of cameras uses Universal Serial Bus (USB) interfaces to eliminate the need for a video capture card. The camera can be directly plugged into a USB port in the back of your computer without switching off or opening up the computer. Windows XP and Windows Server 2003 operating systems automatically detect these cameras, using Plug and Play, and install all the device drivers necessary to make the camera work.

Ensure that drivers are available for the video camera model that you choose. To verify the compatibility of a driver, click the Compatible Hardware and Software link at Support resources.

**Add USB ports to your computer**
Some older computer models may not contain USB ports. This is easily determined by examining the back of the computer for a rectangular female connector that corresponds to the male USB connector on your camera.

Such computers can be upgraded by using PCI cards that have two to four USB ports on the back. Inexpensive cards can be ordered from online computer retail Web sites.

# Setting up conference rooms

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Setting up conference rooms**
On the computer, conference rooms have similar configurations as workstations, but you must provide audio and visual coverage for a group of people.

**Get a good microphone**
Headsets are not practical for canceling echoes in conference rooms. Special equipment that cancels echoes is required, which contains long range microphones and speakers with automatic gain control. Some of these devices can track a speaker walking around a room and automatically adjust the volume before transmitting it to other participants. Several manufacturers produce this equipment. For more information, see Eliminating acoustic echo.

**Install a large screen television**
Your conference rooms may already be equipped with large screen televisions or projection systems. Connect the video output from a computer to these systems to provide an image that is suitably large for viewing by groups of people.

**Set up logon access**
When a participant joins a multicast conference, TAPI automatically determines the name of the user and displays this to other conference participants. This works fine for client workstations, because users usually log on to Windows using their domain accounts. For a video conferencing computer in a conference room, you may want to create a user account that uses the conference room name or location. This allows participants to associate the name of the conference room with its image.

# Enabling IP multicast

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Enabling IP multicast**

To enable multicast conferencing on your network, you must install Active Directory and configure a server with the MADCAP service. The steps are listed in Checklist: Installing IP multicast video conferencing.

For more information about MADCAP and the Microsoft implementation of multicast conferencing, see How TAPI supports IP multicast video conferencing.

**Checking routers on your network**

The Windows Server 2003 family provides the following support for troubleshooting multicast routing and forwarding problems:

- The **mrinfo** command displays the configuration of a multicast router.
- The following **netsh** commands display multicast group tables and troubleshooting information for interfaces.
    - **netsh routing ip show mfe** displays the entries in the multicast forwarding table.
    - **netsh routing ip show mfestats** displays packet statistics and input and output interface information for multicast forwarding entries in the multicast forwarding table.
    - **netsh interface ip show joins** displays the multicast groups locally joined on each interface.
- Although the Windows Server 2003 family does not provide a version of the Mtrace multicasting utility, a multicast router running Windows Server 2003 does respond to **mtrace** command queries from other Mtrace utilities.

For more information about how to use the **mrinfo** command and **netsh** commands for troubleshooting multicast routing and forwarding problems, see Using the IP multicasting utilities. For general information about how to use Netsh, see Netsh overview.

**Configuring routers on your network**

If there is a router on your network and you want to enable multicast conferencing across the router, you may have to configure your router to handle multicast packets. This can involve several steps:

1. Enable multicast for the network (globally).
2. Determine the interfaces on which to use multicast, and enable multicast on those interfaces.
3. Enable multicast routing protocols on specific devices. For example:
    - PIM Sparse mode for links that have limited bandwidth.
    - PIM Dense mode for links that have broad bandwidth.
    - DVMRP
4. Create access lists specifying the range of multicast group addresses allowed to cross the router.
5. Associate access lists with specific interfaces on different routers.

To configure routing services, see Routing. To configure a particular router, refer to the documentation provided by the manufacturer.

# Eliminating acoustic echo

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Eliminating acoustic echo**

Acoustic echo is a common problem with audio conferencing systems. It originates in the local audio loopback that occurs when your microphone picks up audio signals from your the speaker, and sends it back to the other participant along with your voice. The other participant hears this echo of his or her own voice as he or she speaks. For example:

1. Alice and Bob are participating in a call or conference.
2. Alice speaks into her microphone.
3. Bob hears Alice's voice played by his speaker.
4. Bob's microphone also picks up and transmits Alice's voice.
5. Alice can now hear an echo of her own voice, with a slight delay due to the round-trip transmission time.

Acoustic echo can be caused or exacerbated when very sensitive microphones are used, speaker volume is turned up very high, or the microphone and speaker are very close to one another. Besides being annoying, this can prevent normal conversation among participants in a conference.

**Eliminate echo with audio headsets**

One of the easiest ways to completely eliminate acoustic echo is to use audio headsets. These play audio at low volume very close to a conference participant's ears. This prevents the microphone from picking up the audio signal.

Many different audio headsets are available to use with personal computers. These have connectors designed to fit into the audio jacks on sound cards. Many also have built-in microphones that enhance audio quality by staying a constant distance from a person's mouth, even while one moves or changes position.

**Eliminate echo with acoustic echo cancellation**

In Windows XP and Windows Server 2003 operating systems, TAPI 3.1 provides support for acoustic echo cancellation (AEC). During the process of AEC, a digital signal processor electronically compares the speaker signal to the microphone signal, and then subtracts the speaker signal from the microphone signal. In this way, AEC prevents the speaker sound from being fed back into the microphone and transmitted as an echo.

# Setting up the PSTN server and client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Setting up the PSTN server and client**

Telephony programs can use telephony devices controlled by both local and remote telephony service providers. For remote devices, the telephony server controls user access to lines and phones. The Windows Remote Service Provider on client computers gives users access to telephony devices on remote servers. The following instructions summarize how to set up the telephony server and clients.

**Select a server that is part of a domain**

Make sure that the server that will be used as a telephony server is part of a domain and that the domain is trusted by the telephony users' domains. Add telephony clients to this domain or a fully trusted domain. Clients cannot be logged on to more than one domain at a time. For information about adding local users or groups, see Local Users and Groups overview. For more information about adding users and groups in Active Directory, see Windows interface administrative tool reference A-Z: Active Directory Users and Computers.

**Administer the telephony server**

Log on to the server as an administrator. Administrators of the Telephony service must be in the same domain as the server or in fully trusted domains. Use the Telephony console in the Microsoft Management Console (MMC) to administer the Telephony service on the local computer. To open the Telephony console, see Working with MMC console files.

**Enable the telephony server**

The telephony server is automatically installed with Windows Server 2003 operating systems. To begin using the telephony server, you must first enable it, and then specify the logon account for the Telephony service on the server. You can also specify additional telephony administrators. They do not have to have administrative credentials on the entire server.

Use the Telephony console to enable the telephony server. For more information, see Enable or disable a telephony server.

When you enable the telephony server, enter the user name and password of the domain user account to run the Telephony service on the server. (For more information, see Change the Telephony service logon account for a telephony server.) This account must be in the Administrators group on the server. It must also be in the same domain as the server or in a fully trusted domain. You may want to create a special account in the domain for this purpose. If you do, use an easily recognizable name, such as TAPI_ADMIN.

You should also designate individuals who will administer the telephony server. (For more information, see Designate telephony administrators.) This procedure enables you to allow users who do not have administrative credentials on the telephony server to administer the Telephony service. This procedure does not give users additional administrator rights and permissions on the server. Users who have administrative credentials on the telephony server can administer the Telephony service just as they can administer any other service on the computer. These users must be in the same domain as the telephony server or in fully trusted domains.

**Provide access to Telephony service providers**

Specify which users have access to the telephony service providers on the telephony server. This task is also performed in the Telephony console. You can do this by assigning users to a line or phone belonging to the service providers. (For more information, see Assign a telephony user to a line or phone.) Users need to be in the same domain as the telephony server or in a trusted domain. If you need to delete a user, follow the steps in Remove users from telephony lines or phones.

**Set up the telephony clients**

To use remote PSTN telephony devices on a server, telephony clients must be running TAPI version 2.1 or later. For information about which version of TAPI is included in each Windows operating system, see Telephony Resources.

Computers running Windows 98 or Windows Millennium Edition must have File and Printer Sharing for Microsoft Networks and user-level access control enabled, and they must be running Client for Microsoft Networks. For more information, see Enable a Windows 98 or Windows Millennium Edition telephony client.

For all operating systems, you need to specify the telephony servers to be used by the client. For more information, see Specify telephony servers on a client computer. The clients must be in the same domain as the telephony server or in fully trusted domains.

# Managing telephony from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Managing telephony from the command line**
In addition to the Telephony console, you can use command-line utilities to manage telephony servers and clients. For example, you can use Tapicfg to create, remove, or display a TAPI application directory partition, publish a service connection point for a TAPI application directory partition, or set a default TAPI application directory partition.

You can use Tsecimp to import assignment information from an Extensible Markup (XML) file into a telephony server security file (tsec.ini) and to set up or disable telephony clients.

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

**Note**

- If your telephony clients or servers run Windows NT 4.0 or Windows 2000, you must run **tcmsetup** to set up telephony clients. If your telephony clients run Windows XP or a Windows Server 2003 operating system, and your telephony servers run Windows Server 2003 , you can set up telephony clients by adding the Windows Remote Service Provider to your clients instead. To add the Windows Remote Service Provider, see Add a telephony service provider.

# Telephony Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Resources**
This section includes information about additional telephony resources.

**Service packs and updates**
TAPI 2.2 supports the PSTN client-server telephony functionality provided with Windows NT 4.0 Service Pack 4. TAPI 3.0 and TAPI 3.1 support both PSTN telephony and media stream IP telephony and are compatible with products supported by previous versions of TAPI.

The following table shows the version of TAPI included with each Windows operating system. (Recent releases, service packs, or updates might not appear on this table.)

| Windows release | TAPI version |
|---|---|
| Windows 95 | TAPI 2.0 |
| Windows NT 4.0 | TAPI 2.0 |
| Windows NT 4.0 Service Pack 4 | TAPI 2.1 |
| Windows 98 | TAPI 2.1 |
| Windows Millennium Edition | TAPI 2.2 |
| Windows 2000 | TAPI 2.2, TAPI 3.0 |
| Windows XP | TAPI 2.2, TAPI 3.1 |
| Windows Server 2003 family | TAPI 2.2, TAPI 3.1 |

For the latest Windows updates, including service packs, see the Windows Update section at Support resources.

**Other related Microsoft Web sites**

- To verify hardware compatibility, see Compatible Hardware and Software at Support resources.
- For additional product support, see Microsoft Product Support Services at Support resources.

# Telephony Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Troubleshooting**

This section provides information about troubleshooting telephony in the Windows Server 2003 family.

- Troubleshooting IP telephony
- Troubleshooting PSTN telephony

# Troubleshooting IP telephony

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Troubleshooting IP telephony**
What problem are you having?

- A program fails to get a multicast address from the MADCAP server.
- Programs warn that they are using a local random address allocation when users create conferences.
- A client cannot see all participants in a multicast conference.
- Audio or video does not function correctly in a call or conference.
- There are audio or video device errors during calls or conferences.

**A program fails to get a multicast address from the MADCAP server.**
**Cause:** The server is not reachable from the client, or the server is not set up correctly.

**Solution:** The MADCAP server and client must both be in the same domain. The MADCAP server must be configured correctly.

**See also:** Using multicast scopes; Enabling IP multicast; Checklist: Installing IP multicast video conferencing; Create a multicast Telephony scope.

**Programs warn that they are using a local random address allocation when users create conferences.**
**Cause:** All addresses have been allocated for the scopes assigned for MADCAP.

**Solution:** Expand the size of the MADCAP scope.

**See also:** Using multicast scopes; Enabling IP multicast; Checklist: Installing IP multicast video conferencing; Create a multicast Telephony scope.

**A client cannot see all participants in a multicast conference.**
**Cause:** Servers running Routing and Remote Access may not be configured correctly for multicast telephony.

**Solution:** Check that multicast is enabled in both directions on servers running Routing and Remote Access.

**See also:** Understanding Multicasting; Configure IP Multicast Support; Using the IP multicasting utilities.

**Audio or video does not function correctly in a call or conference.**
**Cause:** Either the audio or video is working incorrectly, or the system is using an older audio or video driver that is incompatible with Windows Server 2003 operating systems. Test audio and video on your computer by doing the following:

- Install and use Sound Recorder to check your microphone and speakers. Sound Recorder is an audio program that is available on an optional basis.
- Use a video capture program to check the video camera and its settings, such as Vidcap32, found in the Platform Software Development Kit (SDK).

**Solution:** Check the audio or video equipment manufacturer's Web site for the latest Windows Server 2003 family-compatible drivers.

**There are audio or video device errors during calls or conferences.**
**Cause:** The audio or video device is being used by another program, or the sound card does not support full-duplex functionality.

**Solution:** If closing another program that is using an audio or video device does not resolve the problem, then check that the

sound card and driver support full-duplex functionality:

- Install and start Sound Recorder, and record a long message (about one minute). (Sound Recorder is an audio program that is available on an optional basis.) Rewind the recording, and start Sound Recorder again, so that it appears in two places on the desktop. In the first instance of Sound Recorder, start playing back the recording you made, and then use the other instance to simultaneously record another message. If you can record and play back simultaneously, then your system supports full-duplex audio.

# Troubleshooting PSTN telephony

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 8 Beta

**Troubleshooting PSTN telephony**
What problem are you having?

- One or more client computers cannot see the telephony server.
- Client computers cannot see lines on the telephony server.
- A client user cannot see lines or phones on the telephony server even though the server is set up correctly and the lines or phones have been assigned to the user.
- A program fails to start after a user canceled out of the telephony location information dialog box.
- A client user cannot see a new line on the telephony server, even though the server administrator has assigned the user to the line.
- A telephony program dials a long distance number incorrectly.

**One or more client computers cannot see the telephony server.**
**Cause:** The server cannot be reached on the network because it is not set up correctly.

**Solution:** Ensure the following:

- The client computers can successfully log on to a domain.
- The client computers can successfully see the server with the **ping** command.
- That the telephony server has been configured correctly.

**See also:** Permit users to log on locally to a domain controller; Diagnose Connections; Routing Tools and Utilities; Manage Telephony Servers.

**Client computers cannot see lines on the telephony server.**
**Cause:** The telephony server has not been set up correctly, or users have not been authorized to access lines on the server.

**Solution:** When a TAPI-based program accesses lines on the telephony server, the user who is running the program process is first authenticated. For the user to see the lines, they must have been assigned to the user. Ensure that the server has been set up correctly and that the user is assigned lines on the server.

**See also:** Manage Telephony Servers; Assign a telephony user to a line or phone.

**A client user cannot see lines or phones on the telephony server even though the server is set up correctly and the lines or phones have been assigned to the user.**
**Cause:** The client computer has not been enabled to use the telephony server.

**Solution:** Use **tcmsetup** on the client computer to specify the telephony server.

**See also:** Specify telephony servers on a client computer; Tcmsetup.

**A program fails to start after a user canceled out of the telephony location information dialog box.**
**Cause:** The program requires TAPI to translate addresses, but address translation requires current telephony location information, which has not been specified.

**Solution:** When the location information dialog box is displayed, the user should enter the country/region, local area code,

external line access settings, and select tone or pulse dialing. The user can instead create a location for telephony.

**See also:** Dialing rules overview; Create a new location.

**A client user cannot see a new line on the telephony server, even though the server administrator has assigned the user to the line.**
**Cause:** When you assign a currently running client user to a line on the telephony server, the new settings will not be available until the user restarts the Telephony service on the client computer.

**Solution:** Stop all client TAPI-based programs on the client so that the Telephony service will shut down. When the client programs restart, they will be able to see the newly allocated lines. If you cannot get TAPI to shut down, restart the client computer.

**See also:** Manage Telephony Clients and Users; Start, stop, pause, resume, or restart a service.

**A telephony program dials a long distance number incorrectly.**
**Cause:** The program is not sending the correct number to TAPI.

**Solution:** Check that the number is entered correctly. If the long distance dialing rules for the country/region code are not being implemented, try entering the phone number in canonical address format:

**+** *Country/RegionCode* **(**AreaCode**)***SubscriberNumber*

This is a universal address format for telephone numbers, and it is recognized by TAPI. The format distinguishes the country/region code, area code, and number, so TAPI will dial according to the dialing rules for the country or region.

**See also:** Canonical address format for phone numbers.

# Managing Bandwidth

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing bandwidth**
This section covers:

- Quality of Service (QoS)
- Windows ATM Services

For information about managing bandwidth with Winsock Direct, see "Windows Server 2003 Network Architecture" at the Microsoft Windows Resource Kits Web site.

# Quality of Service (QoS)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quality of Service (QoS)**
QoS is a set of standards and mechanisms that ensure quality data transmission for QoS-enabled programs.

- For help installing QoS Packet Scheduler, see Install QoS Packet Scheduler.
- For general background information, see QoS Concepts.

# Install QoS Packet Scheduler

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install QoS Packet Scheduler**

1. Open Network Connections.
2. Click any connection, and then, on the **File** menu, click **Properties**.
3. Click **Install**, click **Service**, and then click **Add**.
4. Click **QoS Packet Scheduler**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- When you install QoS Packet Scheduler on any network connection, it is installed on every local network connection.
- Install QoS Packet Scheduler on end-systems on which you want to enable traffic shaping, traffic marking, or both. For computers with dial-up connections, install QoS Packet Scheduler if you want to enable Integrated Services over Slow Links (ISSLOW).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
QoS overview
QoS Resources

# QoS Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Quality of Service (QoS):

- QoS overview
- Differentiated Services Code Point (DSCP) overview
- Planning QoS networks
- QoS Resources

# QoS overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**QoS overview**

Quality of Service (QoS) is a set of service requirements that the network must meet in order to ensure an adequate service level for data transmission. These service requirements are based on industry standards for QoS functionality.

QoS allows real-time programs to make the most efficient use of network bandwidth. Because QoS provides some level of guarantee for sufficient network resources, it gives a shared network a level of service similar to that of a dedicated network.

A QoS guarantee indicates a service level that enables a program to transmit data at a specified rate and to deliver it within a specified time frame.




The goal of QoS is a guaranteed delivery system for network traffic, such as Internet Protocol (IP) packets.

# Differentiated Services Code Point (DSCP) overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Differentiated Services Code Point (DSCP) overview**
Differentiated Services Code Point (DSCP) is a field in an IP packet that enables different levels of service to be assigned to network traffic. This is achieved by marking each packet on the network with a DSCP code and appropriating to it the corresponding level of service.

DSCP is the combination of IP Precedence and Type of Service fields. In order to work with legacy routers that only support IP Precedence, DSCP values are used because they are compatible with IP Precedence fields. For more information, see RFC 2474 at the RFC Editor Web site.

**How DSCP works**
Quality of Service (QoS)-enabled programs request a specific service type for a traffic flow through the generic QoS (GQoS) application programming interface (API). The available service types are:

- Guaranteed service

  Guaranteed service provides high quality, quantifiable guarantees with bounded (guaranteed minimum) latency.

- Controlled load service

  Controlled load service provides high quality, quantifiable guarantees without bounded latency.

DSCP is a six-bit field carrying the default values that are displayed in the following table. The equivalent IP Precedence and IEEE 802.1p values are shown in separate columns.

| Service Type | DSCP | IP Precedence | IEEE 802.1p |
|---|---|---|---|
| Network control | 30 | 6 | 7 |
| Guaranteed | 28 | 5 | 5 |
| Controlled load | 18 | 3 | 3 |
| All other traffic | 0 | 0 | 0 |

**Notes**

- The network control service cannot be requested through the GQoS API. However, it can be requested by network management applications that use the traffic control (TC) API.

- Default DSCP values can be edited remotely using Group Policy.

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Planning QoS networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning QoS networks**
It is recommended that you ensure your network hardware, Windows configuration, programs, and services meet the necessary requirements for Quality of Service (QoS).

**Hardware**
To take advantage of IEEE 802.1p markings, QoS requires that your network adapters and network equipment are 802.1p compliant. Additionally, you may need to configure network hardware for QoS 802.1p markings. For more information about the standards upon which Windows QoS networks are based, see QoS standards (RFCs).

**Windows configuration**
Install QoS Packet Scheduler on end-systems on which you want to enable traffic shaping, traffic marking, or both. For computers with dial-up connections, install QoS Packet scheduler if you want to enable Integrated Services over Slow Links (ISSLOW). For more information, see Install QoS Packet Scheduler.

**Programs and services**
Only programs and services that have been designed to work with QoS can take advantage of QoS functionality. Examples of QoS-aware services and programs are Windows Messenger and NetMeeting.

# QoS Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This topic covers:

- QoS standards (RFCs)
- QoS updated technical information

# QoS standards (RFCs)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**QoS standards (RFCs)**

Request for Comments (RFCs) are an evolving series of reports, proposals for protocols, and protocol standards used by the Internet community. Quality of Service (QoS) standards are defined in RFCs published by the Internet Engineering Task Force (IETF) and other working groups.

The following RFCs cover QoS issues:

| | |
|---|---|
| 2211 | Specification of the Controlled-Load Network Element Service |
| 2212 | Specification of Guaranteed Quality of Service |

For more information about RFCs and how to obtain them, see TCP/IP RFCs.

**Notes**

- For additional information about the most current RFCs related to QoS, see the RFC Editor Web site.
- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# QoS updated technical information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**QoS updated technical information**
The following Web sites post information about Quality of Service (QoS) issues, Internet engineering standards, and other related topics.

- The Internet Engineering Task Force Web site
- The Transport area of the Active working groups page at the Internet Engineering Task Force Web site

For more information about QoS, see "Quality of Service" at the Microsoft Windows Resource Kits Web site.

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Windows ATM Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows ATM services**
Asynchronous transfer mode (ATM) is a standardized set of networking technologies that can be used to build high-speed networks that accommodate Quality of Service (QoS) for connections. Microsoft® Windows® ATM services software allows you to participate in ATM networks if you have ATM adapter hardware installed on your computer.

- Before you begin using ATM with this computer, see Checklist: Installing ATM services.
- For tips about using ATM, see ATM Best practices.
- For help with specific tasks, see ATM How To....
- For general background information, see ATM Concepts.
- For problem-solving instructions, see ATM Troubleshooting.

# Checklist: Installing ATM services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing ATM services**

| | Step | Reference |
|---|---|---|
| | Review ATM concepts such as cells, the ATM model, and ATM QoS. | Understanding ATM |
| | Review LAN emulation (LANE). | Configuring LAN emulation |
| | Review TCP/IP over ATM. | IP over ATM (IP/ATM) |
| | Install an ATM adapter on this computer. | Install an ATM adapter |
| | (Optional) Configure the LAN emulation client. | Configure a LAN emulation client |
| | (Optional) Set up IP/ATM connections and install and configure the ATM ARP/MARS service. | Enable a simple TCP/IP-over-ATM connection; Install and configure the ATM ARP/MARS service |
| | Configure permanent virtual connections (PVCs). | Create a permanent virtual circuit using ATM |
| | Troubleshoot ATM connections. | ATM Troubleshooting |

# ATM Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Use the default emulated LAN (ELAN).**

  On computers running Windows Server 2003 operating systems, asynchronous transfer mode (ATM) services are configured with a default *unspecified_ELAN_name* group name. If you plan to implement LAN emulation (LANE), it is recommended that you use the preconfigured, default unspecified ELAN group name.

  When you purchase an ATM switch, check the product specifications to ensure that it is preconfigured with the default, unspecified ELAN group name.

- **Use supported ATM adapters.**

  Before you buy an ATM adapter, confirm that it is designed for Windows ATM. For more information about compatible hardware, see Support resources.

- **Note configuration parameters before you upgrade from Windows NT 4.0 or Windows 2000 to** Windows Server 2003 **operating systems.**

  Before upgrading from Windows NT 4.0 or Windows 2000, note the following configuration information for each of the LAN emulation clients you plan to upgrade.

  - The ELAN name
  - The media type to be emulated on the LAN
  - ATM addresses for the LAN emulation server (LES) and broadcast and unknown server (BUS) associated with the ELAN
  - The maximum allowable packet size for the ELAN

  After you note these configuration parameters, use the LAN emulation configuration server (LECS) interface on your ATM switch to configure ELANs and their associated parameters, as listed and noted above.

  Next, install the Server software and configure the ELAN name for each LAN emulation client (LEC). For information about configuring the ELAN name, see Configure a LAN emulation client.

- **Use only one ATM ARP/MARS per virtual LAN.**

  If your network uses IP/ATM, it is recommended that you configure only one ATM ARP/MARS for each virtual local area network (VLAN) on your network. If you have multiple ARP servers on the same network segment, and your ARP client is configured with the addresses for these servers, the ARP caches could become out of synch. This can render parts of the network unreachable.

# ATM How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install an ATM adapter
- Configure a LAN emulation client
- Create an encapsulated Ethernet permanent virtual circuit
- Enable a simple TCP/IP-over-ATM connection
- Configure an advanced TCP/IP-over-ATM connection
- Install and configure the ATM ARP/MARS service
- Create a permanent virtual circuit using ATM
- View ATM connection information

# Install an ATM adapter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install an ATM adapter**

- In most cases, if your ATM network adapter hardware supports Plug and Play, your adapter is automatically detected and installed. For more information, see Hardware.

## Notes

- After your ATM adapter is recognized and installed, ATM services are installed and available. By default, the connection is set up as a LAN emulation client (LEC). To view or configure additional services, right-click your ATM connection in Network Connections, and then click **Properties**.

- If your ATM hardware does not support Plug and Play detection, contact your adapter vendor for updated installation software and information about installing your adapter. For more information about installing hardware that does not support Plug and Play, see Hardware.

- To provide end-to-end connectivity, you must configure a permanent virtual circuit (PVC) on the ATM adapter at each end station. You must also configure a PVC on each ATM device (for example, router, ATM switch, and DSLAM) connected between the end stations. For information about configuring a PVC on ATM devices, please see the manufacturers' documentation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure a LAN emulation client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a LAN emulation client**

1. Open Network Connections.
2. Click the ATM connection that corresponds to the ATM network adapter installed on this computer.
3. On the **File** menu, click **Properties**.
4. In **This connection uses the following items**, click **ATM LAN Emulation**, and then click **Properties**.
5. Click **Add**, **Edit**, or **Remove** to add, rename, or delete entries in the list of emulated LANs available for use with this ATM connection.

   The ATM LAN emulation client is configured by default to support a simple connection with another device (typically an ATM switch) that is providing LAN emulation services to your network.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- By default, the LAN emulation client is preconfigured to use the default emulated LAN (ELAN) group name for simple client configuration:

  *<unspecified_ELAN_name>*

  If you do not want to configure a specific ELAN name on your network, use the unspecified ELAN name as the default ELAN group name.

  Where this default ELAN name is not currently configured or supported by other ATM hardware products (such as other switches used on your network), you can do one of the following:

  - Add another supported ELAN name to the list maintained here.
  - Configure a default ELAN on your ATM switches that provide LAN emulation service to your network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install an ATM adapter
Create a permanent virtual circuit using ATM
Configure TCP/IP for static addressing
Configure TCP/IP for dynamic addressing

# Create an encapsulated Ethernet permanent virtual circuit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an encapsulated Ethernet permanent virtual circuit**

1. Open Network Connections.
2. Click the ATM connection that corresponds to the ATM network adapter installed on this computer.
3. On the **File** menu, click **Properties**.
4. On the **General** tab, click **Install**, click **Protocol**, and then click **Add**.
5. Click **Microsoft Ethernet PVC - RFC2684**, click **OK**, and then click **Close**.
6. In the Network Connections folder, right-click the icon for the connection with the device name **Ethernet ATM Miniport**, and then click **Properties**.
7. On the **General** tab, click **Configure**.
8. On the **Advanced** tab, click **Encapsulation Type**.
9. In **Value**, click an encapsulation type.
10. Click **VCI Number**, and then type a number in **Value**.
11. Click **VPI Number**, type a number in **Value**, and then click **OK** twice to apply the configuration you specified.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- Contact your service provider to determine what encapsulation type to choose and what values to type in for virtual circuit identification (VCI) and virtual path identification (VPI).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install an ATM adapter
Example: Permanent virtual circuit (PVC)
Creating permanent virtual circuits
Encapsulated Ethernet over PVC
Create a permanent virtual circuit using ATM

# Enable a simple TCP/IP-over-ATM connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable a simple TCP/IP-over-ATM connection**

1. Open Network Connections.
2. Click the ATM connection that corresponds to the ATM network adapter installed on this computer.
3. Click **File**, and then click **Properties**.
4. In **This connection uses the following items**, select the **Internet Protocol (TCP/IP)** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You can use the default TCP/IP settings for quick ATM ARP client configuration for a single-switch network by installing the ARP/MARS server on one system. If you use another ARP/MARS server, the client must be configured with the ATM address of that server.
- In order for the default setup to work, the ATM ARP/MARS service must be running on a server on the network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install an ATM adapter
Configure an advanced TCP/IP-over-ATM connection
Create a permanent virtual circuit using ATM

# Configure an advanced TCP/IP-over-ATM connection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure an advanced TCP/IP-over-ATM connection**

1. Open Network Connections.
2. Click the ATM connection for the ATM network adapter installed on this computer.
3. On the **File** menu, click **Properties**.
4. In **This connection uses the following items**, select the **Internet Protocol (TCP/IP)** check box, and then click **Properties**.
5. To connect using a permanent virtual circuit (PVC), in **Internet Protocol (TCP/IP) Properties**, specify an IP address for use with this ATM connection. By default, the TCP/IP connection uses a DHCP server to obtain an IP address.
6. Click **Advanced**, and then click the **ATM ARP Client** tab.
7. If you want TCP/IP to be used only over permanent virtual circuits (PVCs) configured using the ATM Call Manager, select **PVC only**. In order for this option to function, you must configure the PVC with an application type of **ATM ARP** in the ATM Call Manager.
8. To connect using a switched virtual circuit (SVC), in **ARP server address list** and **MARS address list**, use **Add**, **Edit**, and **Remove** to add, rename, or delete an address entry or entries for any ATM ARP servers or ATM Multicast Address Resolution Service (MARS) servers on your network. If more than one address is listed, use the **up** and **down** arrow buttons to reorder the list.
9. After you finish configuring these properties, click **OK**.
10. Configure any other advanced TCP/IP properties that this ATM connection will use, such as multiple IP addresses, DNS servers, or WINS servers, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group or the Network Configuration Operators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- For ATM ARP or MARS server lists, you can either use the default address or add entries for other computers and switches on your network. If needed, the entries added to the list can include ATM addresses for other ATM switches that provide comparable ARP/MARS service.
- If you add multiple addresses to the ARP or MARS server address lists, the ATM ARP client tries the first address on the list. If the client fails to connect using that address, it then tries each address in the order that it appears on the list.
- In most cases, default client settings provide rapid configuration, and ATM ARP clients are preconfigured with default ARP/MARS addresses. Advanced client configuration might not be needed.
- Only one ATM ARP/MARS server may use the default address. If clients have difficulty connecting to the server over ATM, you should view the ATM ARP/MARS server Error Log for details.

- By default, the maximum transmission unit is set to 9180. The maximum transmission unit must be the same for all clients using the same server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a permanent virtual circuit using ATM
Enable a simple TCP/IP-over-ATM connection
Configure TCP/IP for static addressing
Configure TCP/IP for dynamic addressing

# Install and configure the ATM ARP/MARS service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install and configure the ATM ARP/MARS service**

1. Open Network Connections.
2. Click the ATM connection that corresponds to the ATM network adapter installed on this computer.
3. On the **File** menu, click **Properties**, and then click **Install**.
4. In **Select Network Component Type**, in the list of network component types, click **Protocol**, and then click **Add**.
5. In **Select Network Protocol**, in the list of network protocols, click **ATM ARP/MARS Service**, and then click **OK**.

   The ATM ARP/MARS service is installed with a default configuration. For more information on modifying these default services, see the notes below or see your ATM networking and product documentation.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- The ATM ARP/MARS service is not available on Windows XP.
- It is recommended that you configure only one ATM ARP/MARS for each virtual local area network (VLAN) on your network. If you have multiple ARP servers on the same network segment, and your ARP client is configured with the addresses for these servers, the ARP caches could become out of synch. This can render parts of the network unreachable.
- The ATM ARP/MARS service supports IP-to-ATM address resolution for ATM network hosts active on the same logical IP subnet as this computer. The ATM ARP service supports classical IP over ATM connections, as described in RFC 2225. The MARS service supports multicast IP over ATM connections, as described in RFC 2022.
- By default, this predetermined address is used for simple configuration: `47007900010200000000000000000A03E00000200`. This number is an NSAP address that Microsoft selected to simplify configuration and enhance interoperability of the ATM ARP/MARS service with other computers running ATM services.
- The ATM ARP/MARS service includes a preconfigured default list of IP address ranges for which broadcast and multicast forwarding are provided by the service. This default is known as hubbed mode operation. If the IP address range list is empty, the ATM ARP/MARS service defaults to nonhubbed mode operation. In this mode, the service will not perform any forwarding for multicast group clients.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install an ATM adapter
Enable a simple TCP/IP-over-ATM connection

Configure an advanced TCP/IP-over-ATM connection

# Create a permanent virtual circuit using ATM

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a permanent virtual circuit using ATM**

1. Open Network Connections.
2. Click the ATM connection that corresponds to the ATM network adapter for which you want to create a permanent virtual circuit (PVC).
3. On the **File** menu, click **Properties**.
4. In **This connection uses the following items**, click **ATM Call Manager**, and then click **Properties**.
5. In **ATM Call Manager** properties, click **Add**.
6. Review and modify PVC settings as needed:
   - For **Name**, you can either use the default unspecified PVC name or type a name. Both are used only for your reference.
   - For **Virtual path ID**, you can either use the default path of **0** or type a number that should be used to identify the virtual path for the connection.
   - For **Virtual circuit ID**, type a number that identifies the virtual circuit within the specified virtual path for the connection.
   - In **Application type**, click the type of application for use with this PVC. If you configured your IP/ATM connection for PVCs only, you must select the application type **ATM ARP** for this PVC.
7. If needed, click **Advanced** to configure any settings that specify a quality of service for use with the PVC.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer. For more information, see Default local groups and Default groups.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- PVCs are not required and are not typically recommended for building most types of ATM network connections.
- To provide end-to-end connectivity, you must configure a PVC on the ATM adapter at each end station. You must also configure a PVC on each ATM device (for example, router, ATM switch, and DSLAM) connected between the end stations. For information about configuring a PVC on ATM devices, please see the manufacturers' documentation.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Install an ATM adapter
Example: Permanent virtual circuit (PVC)
Create an encapsulated Ethernet permanent virtual circuit
Creating permanent virtual circuits

# View ATM connection information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view ATM connection information**

1. Open Command Prompt.
2. At the command prompt, type **atmadm /?** to display syntax and parameters for the **atmadm** command.

**Note**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Atmadm
Managing ATM from the command line

# ATM Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**
This section provides general background information about asynchronous transfer mode (ATM) technologies that are available.

- ATM Overview
- Understanding ATM
- Using ATM
- ATM Resources

# ATM Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ATM overview**

This section covers:

- ATM defined
- Windows ATM services defined
- Benefits of ATM

# ATM defined

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ATM defined**

Asynchronous transfer mode (ATM) refers to a number of related technologies, including software, hardware, and connection media. ATM differs from other existing LAN and WAN technologies and was specifically designed to support high-speed communication. ATM allows networks to use bandwidth resources at maximum efficiency while maintaining Quality of Service (QoS) for users and programs with strict service requirements.

The basic components of ATM are the computers that are connected to the ATM network and the devices that connect these computers and ensure that data is transferred successfully. The computers that are connected to the ATM network are called *end stations*. Routers, DSLAMs, and ATM switches are examples of devices that connect end stations and ensure that data is transferred successfully.

**Asynchronous** means the available network bandwidth is not divided into fixed channels or slots synchronized by a timing mechanism or clock. Devices that communicate using asynchronous communication are not bound by design in terms of their ability to send and receive information at a precise transmission rate. Instead, the sender and receiver negotiate the rate at which they will communicate based on physical hardware limitations and the ability to maintain a reliable flow of information within the network.

**Transfer mode** refers to the way the information is transferred between sender and receiver. In ATM, the concept of small, fixed-length cells is used to structure and parcel data for transfer. By using cells that contrast directly with the variable-length packet mechanism used by most existing network technologies, ATM assures that connections can be negotiated and managed so that no single data type or connection can monopolize the transfer path.

# Windows ATM services defined

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows ATM services defined**

- An ATM UNI Call Manager that conforms to ATM Forum specifications for signaling over ATM and supports the creation of switched virtual circuits (SVCs) and permanent virtual circuits (PVCs).
- Updated NDIS 5.0 ATM miniport drivers that display the **Designed for Microsoft® Windows®** logo and provide limited support for installed ATM adapter hardware. For more information, see Designed for Microsoft Windows logo.
- A LAN emulation client module that conforms to the ATM Forum LANE 1.0 specification, enabling existing programs and protocols designed for use on Ethernet and Token Ring networks to run without modification over an ATM network.
- An ARP/MARS service that enables the Microsoft TCP/IP stack to resolve ATM addresses to hardware addresses for more direct and efficient use of ATM media. The IP-over-ATM module conforms to RFC specifications for Address Resolution Protocol (ARP) over ATM media, as well as support for a multicast address resolution service (MARS).
- Enhanced support for TAPI-based programs and services.
- Support for raw channel access filtering that can be used in program environments that include DirectStreaming.
- Expanded support for PPP dial-up over ATM media, enabling you to create and use a PPP dial-up connection over any supported ATM adapter that is installed in a computer running a Windows Server 2003 operating system. This feature provides support for ATM over DSL, ATM over cable modem, and other ATM on-demand connectivity possibilities.
- Access to native ATM services with Winsock 2.0 over ATM. Available through the Windows Sockets ATM Service Provider, Winsock 2.0 allows direct access to ATM services from user-mode programs. It also allows programs that use TCP/IP as a transport protocol to run over ATM networks and interoperate with standard LAN-based IP clients.
- Encapsulated Ethernet over a permanent virtual connection (PVC), which enables the transmission of Ethernet packets across a PVC on an ATM network adapter. It accomplishes this by acting as an Ethernet adapter to the TCP/IP protocol and using the PVC on the ATM adapter to transfer Ethernet packets.

# Benefits of ATM

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Benefits of ATM**

ATM provides a flexible and scalable solution to the increasing need for quality of service in networks where multiple information types (such as data, voice, and real-time video and audio) are supported. With ATM, each of these information types can pass through a single network connection.

ATM can provide the following benefits:

- High-speed communication
- Connection-oriented service, similar to traditional telephony
- Fast, hardware-based switching
- A single, universal, interoperable network transport
- A single network connection that can reliably mix voice, video, and data
- Flexible and efficient allocation of network bandwidth

# Understanding ATM

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding ATM**

Asynchronous transfer mode (ATM) refers to a number of related and standardized technologies designed to provide high-speed network communication. For more information on these technologies, see the following sections:

- ATM networks
- Example: Switched virtual circuit (SVC)
- Example: Permanent virtual circuit (PVC)
- Encapsulated Ethernet over PVC
- IP over ATM (IP/ATM)

# ATM networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ATM networks**

ATM networks are made up of three distinct elements: users (endpoint devices), switches, and interfaces.

In ATM networking, two types of interfaces exist to describe how these elements communicate: user-to-network interfaces (UNIs) and network-to-network interfaces (NNIs).

The UNI and NNI specifications provide a standard signaling method for end stations and ATM switches to communicate.




**Note**

- The Windows ATM services Call Manager currently supports ATM Forum UNI 3.1. For more information about Windows ATM services, see Windows ATM services defined

# Example: Switched virtual circuit (SVC)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Example: Switched virtual circuit (SVC)**
The following example shows in detail how an SVC connection is established in an ATM network.




1. An ATM user, End station A, sends a SETUP message on the virtual channel connection reserved for signaling (VPI=0, VCI=5) to its directly connected ATM switch, ATM switch 1.

   The SETUP message contains:

   - The 20-byte ATM address of End station B.
   - Quality of Service (QoS) parameters needed for connection.

2. ATM switch 1 analyzes the SETUP message to determine whether it can find a table entry for End station B's switch address and whether it can support the required QoS parameters for the requested connection.

   If ATM switch 1 can find End station B's switch address and support the required QoS parameters for connection, a CALL PROCEEDING message is sent back to End station A.

   If ATM switch 1 either cannot find End station B's switch address or accommodate the connection, a RELEASE message containing standard error codes is sent back to End station A and the call request is rejected.

3. ATM switch 1 sends the SETUP message (described in Step 1) on the virtual channel connection reserved for signaling (VPI=0, VCI=5) to the next switch in the network, ATM switch 2, on its way to the destination, End station B.

   The SETUP message contains the same ATM address of End station B and all required QoS parameter information as included in Step 1. Additionally, the SETUP message now identifies a new VPI/VCI pair for use between ATM switch 1 and ATM switch 2.

4. ATM switch 2 receives the SETUP message. End station B is determined to be a directly connected endpoint registered to ATM switch 2. If the QoS required for the requested connection can be accommodated, ATM switch 2 sends a CALL PROCEEDING message back to ATM switch 1.

5. The SETUP message reaches its final destination, End station B. End station B can respond in either of two ways:

   - If End station B accepts the call, it sends a CALL PROCEEDING message, followed immediately by a CONNECT message back to ATM switch 2.
   - If End station B rejects the call, it sends a RELEASE message back to ATM switch 2. The RELEASE message is passed through the network back to the original source user (End station A) containing standard error codes that describe the reason for rejection.

6. ATM switch 2 sends a CONNECT ACK message back to End station B, then passes the CONNECT message to ATM switch 1.

7. ATM switch 1 sends a CONNECT ACK message back to ATM switch 2, then passes the CONNECT message to End station A.

8. End station A sends a CONNECT ACK message back to ATM switch 1.

The previous example demonstrates a simplified SVC connection using only two intermediate switches, ATM switch 1 and ATM switch 2. In practice, most SVC connections use many intermediate switches. When more intermediate switches are used, each switch forwards SETUP, CONNECT, and RELEASE messages to its downstream neighbor switch, while relaying acknowledgment messages (CONNECT ACK, RELEASE ACK) back to its upstream neighbor switch.

The result of this connection process is a virtual circuit (VC) in both directions between End stations A and B in which:

- All intermediate switches and ATM user hardware at End stations A and B now support and protectively enforce the QoS communication parameters required for this connection.

- All intermediate switches (in this case, ATM switches 1 and 2) have been programmed with switch table entries for mapping the series of VPI/VCI pairs used for routing ATM cells between End stations A and B.

If any intermediate switch cannot accommodate the requested QoS connection parameters, a RELEASE message is sent back in the direction of the source user. The RELEASE message contains the standard error codes and frees the resources being reserved at each device in the connection path and at the source.

# Example: Permanent virtual circuit (PVC)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Example: Permanent virtual circuit (PVC)**
Besides using switching and call message signaling to create circuits on demand between any two possible users, you can also permanently configure an ATM virtual circuit between any two specified service endpoints on an ATM network.

PVCs are used extensively by public ATM service providers to create and establish a complex ATM-based infrastructure for their internal networks. In many cases, the internal ATM infrastructure of the network is built using PVCs with actual end-to-end connections occurring over SVCs.

PVCs can also be useful in some private ATM networks where the following types of situations exist:

- A large campus LAN migrating to a higher-speed ATM backbone. In backbone configurations, connections typically represent a few static configurable switch paths that infrequently change, so permanent configuration of ATM circuits is acceptable.

- A small WAN with a limited number of sites, each requiring a continuous, dedicated high-speed connection that guarantees a fixed quality of service between site locations. With a circuit permanently established, the ATM switches at both WAN sites do not require the added latency and overhead of using call signaling or connection setup and teardown each time ATM cell traffic data is sent over the network. Sent data can be directly forwarded over PVCs established between each site.

# Encapsulated Ethernet over PVC

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encapsulated Ethernet over PVC**

With encapsulated Ethernet over PVC, you can use an ATM network adapter to send and receive Ethernet packets across a permanent virtual circuit (PVC). Before this technology was developed, you had to install a separate driver for each digital subscriber line (DSL) network adapter to allow an Ethernet network to exchange information with an ATM network. With this technology, you need only provide the virtual path identifier (VPI), the virtual channel identifier (VCI), and the encapsulation type required by your service provider. After you configure these settings, the encapsulated Ethernet PVC transmits data through the ATM adapter. Encapsulated Ethernet over PVC performs the following functions:

- It acts as an Ethernet adapter to the TCP/IP protocol.
- It uses a PVC on the ATM adapter to transfer Ethernet packets.
- It provides a mechanism to transport packets across a DSL.
- It eliminates the need to install different drivers for different DSL network adapters.

**Notes**

- The VPI and VCI must match those configured on your service provider's switch or router.
- For information about creating an encapsulated Ethernet PVC, see Create an encapsulated Ethernet permanent virtual circuit
- For more information about encapsulation, see RFC 2684 on the RFC Editor Web site.
- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# IP over ATM (IP/ATM)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP over ATM (IP/ATM)**
This section covers:

- IP/ATM overview
- How IP/ATM works

# IP/ATM overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP/ATM overview**

IP/ATM is a group of services for communicating over an ATM network that can be used as an alternative to LAN emulation. IP/ATM is handled by two main components: the IP/ATM client and the IP/ATM server. The IP/ATM server includes an ATM ARP server and a multicast address resolution service (MARS). IP/ATM network components can reside on a server or an ATM switch.

The main advantage of using IP/ATM is that it is faster than LANE. With IP/ATM, no additional header information is added to packets as they move through the protocol stack. Once an IP/ATM client has established a connection, data can be transferred without modification.

IP/ATM supports the use of a Dynamic Host Configuration Protocol (DHCP) server in the ATM network.

# How IP/ATM works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How IP/ATM works**
The three general stages of IP/ATM operation are:

- Client initialization
- Client registration
- Data transfer

**IP/ATM client startup and registration with a static IP address**
The following example describes each step in establishing an IP/ATM connection for a single IP/ATM client with a static IP address:




1. Client A initializes and gets an ATM address from the ATM switch.
2. Client A connects to the ATM ARP/MARS server and joins the broadcast group. Client A's IP-to-ATM address mapping is added to the ATM ARP server database.
3. Client A contacts Client B, an ATM end station connected to the network, and begins data transfer.

**IP/ATM client startup and registration with DHCP**
The following example describes each step in establishing an IP/ATM connection for a single IP/ATM client obtaining an IP address using Dynamic Host Configuration Protocol (DHCP):

1. The client initializes and gets an ATM address from the ATM switch.
2. The client connects to the ATM ARP/MARS server and joins the broadcast group.
3. The client connects to the multicast server and sends a DHCP request. The multicast server broadcasts the DHCP request to all members of the broadcast group. The DHCP server receives the request.
4. The DHCP server sends a DHCP reply to the multicast server, which then broadcasts the reply to the broadcast group.
5. The client receives the DHCP reply and then registers its IP and ATM addresses with the ATM ARP/MARS server.
6. The client is now ready to contact other hosts and begin data transfer.

**Note**

- For more information about DHCP, see DHCP.

**Broadcasting and multicasting**

A broadcast or multicast is used to establish point-to-multipoint connections between the requesting client and multiple end stations on the network. The process of broadcasting and multicasting on IP/ATM can follow two different methods:

- **Direct point-to-multipoint connection**

  If a client needs to send an IP packet to a broadcast or multicast IP address, it sends a request to the MARS to resolve the IP address to a list of clients. The MARS sends a group of addresses to a client, allowing it to set up a point-to-multipoint connection.

- **Point-to-multipoint connection through a multicast server**

  The MARS can also work with a multicast server. The Windows ATM ARP/MARS has an integrated multicast server, which registers one or more multicast groups with the MARS and receives a list of members in each multicast group from the MARS. The MARS updates the multicast server when clients join or leave a multicast group. When a client makes a multicast or broadcast request to the MARS, the MARS returns only the address of the multicast server. The client contacts the multicast server, which creates a point-to-multipoint connection with the multicast group. The multicast server copies and distributes the packets sent by the client that initiated the point-to-multipoint call to end stations on the multicast list.

# Using ATM

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using ATM**

This section covers:

- ATM addressing
- Configuring LAN emulation
- Creating switched virtual circuits
- Creating permanent virtual circuits
- Modifying ATM defaults
- Managing ATM from the command line
- Well-known addresses

# ATM addressing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ATM addressing**

ATM addresses are needed to support the use of switched virtual connections (SVCs) through an ATM network.

At the simplest level, ATM addresses are 20 bytes long and have three distinct parts:

- **Network prefix**

  The first 13 bytes identify the location of a specific switch in the network. The use of this portion of the address can vary considerably depending on its address format. Each of the three standard ATM addressing schemes provides information about ATM switch locations differently. These schemes include the data country/region code (DCC) format, the international code designator (ICD) format, and the E.164 format proposed by the ITU-T for international telephone numbering use in broadband ISDN networks.

- **Adapter media access control address**

  The next 6 bytes identify a physical endpoint, such as a specific ATM adapter card, using a media access control layer address that is physically assigned to the ATM hardware by its manufacturer. The use and assignment of media access control addresses for ATM hardware are identical to how this addressing works for Ethernet, Token Ring, and other IEEE 802.*x* technologies.

- **Selector (SEL)**

  The last byte is used to select a logical connection endpoint on the physical ATM adapter. Although all ATM addresses fit this basic three-part structure, there are significant differences in the exact format of the first 13 bytes, depending on the addressing format or whether the ATM network is for public or private use.

All of the three ATM address formats that are currently in widespread use (DCC, ICD, and E.164) include the following characteristics:

- They comply with the Network Service Access Point (NSAP) addressing plan as proposed by the Open Standards Interconnection (OSI) protocol suite of the International Standards Organization (ISO).

- Each can be used to establish and interconnect privately built ATM networks that support switched virtual connections (SVCs).

# Configuring LAN emulation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring LAN emulation**

The basic services for an emulated LAN reside on the ATM switch and workstations with ATM hardware installed. An edge device is required if you plan to connect an Ethernet or Token Ring network to your emulated LAN.

| LAN emulation (LANE) component | Where the component resides |
|---|---|
| LANE services | The LES, LECS, and BUS typically reside on the ATM switch. |
| LAN emulation client | The LAN emulation client can reside on workstations, ATM switches, and edge devices. |

**Services**

Most of the configuration for LAN emulation is done on the switch running the LAN emulation services. These services include the LAN emulation server (LES), LAN emulation configuration server (LECS), and the broadcast unknown server (BUS). For more information about LAN emulation, see "Asynchronous Transfer Mode" at the Microsoft Windows Resource Kits Web site. For more information about configuring the LANE services on your ATM switch, see the documentation for your ATM switch.

**Client**

The LAN emulation client (LEC) software resides on workstations, ATM switches, and edge devices. The LEC is integrated in Windows Server 2003 operating systems and is installed by default if you have an ATM network adapter installed on your computer. After installation is complete, you can install and configure any network protocol driver that works over Ethernet or Token Ring, such as TCP/IP or IPX/SPX. For information about installing and configuring network protocols on your ATM ELAN connection, see Configure a connection to a remote network. For more information about configuring the LAN emulation client, see Configure a LAN emulation client.

### Note

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

# Creating switched virtual circuits

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating switched virtual circuits**
Switched virtual circuits (SVCs) are processed by the ATM Call Manager, which handles virtual circuit creation and management. The ATM ARP client and the LAN emulation client support SVCs. Both the LAN emulation client and ATM ARP client are based on the ATM Forum UNI specification.

**LAN emulation with SVCs**
The LAN emulation client running Windows uses SVCs by default. The client must be configured with a valid emulated LAN name to use SVCs. For more information, see Configure a LAN emulation client.

**IP/ATM with SVCs**
The ATM ARP client running Windows uses SVCs by default. The ATM ARP client must be configured with the address of the ATM ARP/MARS on the subnet. For more information, see Configure an advanced TCP/IP-over-ATM connection.

# Creating permanent virtual circuits

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating permanent virtual circuits**

Permanent virtual circuits (PVCs) are typically used by large service providers or in campus backbones to maintain virtual connections in an ATM network.

The ATM Call Manager is used to manually set up PVCs in an emulated LAN or in an IP/ATM network. Use the ATM Call Manager to define the properties for each PVC, including the PVC definition, program usage, and advanced properties.

| PVC property | Description |
|---|---|
| PVC definition | Define the name, virtual path ID, virtual circuit ID, and ATM adaptation layer (AAL). |
| Program usage | Define the program type for use with the PVC and the basic call-matching characteristics of this PVC. |
| Advanced properties | Define the Quality of Service (QoS) transmit and receive parameters and specific layer protocols for call matching and answer matching. |

For more information, see Create a permanent virtual circuit using ATM.

# Modifying ATM defaults

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modifying ATM defaults**

When you install an ATM adapter on computers running Windows Server 2003 operating systems, the default LAN emulation client configuration for ATM services is automatically installed. The LAN emulation client is preconfigured to use the following default emulated LAN (ELAN) group name:

```
<unspecified ELAN name>
```

The *unspecified_ELAN_name* is configured for use with Windows ATM services and by other vendors of ATM switching products. It enables simplified configuration of the ATM LAN emulation client service provided with computers running Windows Server 2003 operating systems, with ATM switches that might be providing LAN emulation services for your network.

Where the unspecified ELAN name is not currently configured or supported by default by other ATM hardware products (such as other switches used on your network), you can do one of the following:

- Add another supported ELAN name to the list maintained here.
- Configure the unspecified ELAN name on your ATM switches that provide LAN emulation service to your network.

For more information about changing the settings of your LAN emulation client, see Configure a LAN emulation client.

# Managing ATM from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing ATM from the command line**
From the command line, you can monitor details of connections and addresses registered by ATM Call Manager on an asynchronous transfer mode (ATM) network. The command line function **atmadm** allows you to view more detailed connection information than the user interface, which only monitors simple ATM connectivity.

To monitor ATM connectivity:

1. Click the icon for your ATM adapter in the Network Connections folder.
2. Click **File**, and then click **Properties**.
3. On the General tab, select the **Show icon in taskbar notification when connected** check box.

### Notes

- For specific information about using the ATM command-line utility, see Atmadm.
- When working with the **atmadm** command, you can list commands currently available for use by typing any of the following: **help/?**, **/?/?**, **-?/?**, **?/?**.

# Well-known addresses

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Well-known addresses**

**LANE well-known addresses**
LAN emulation clients (LEC) running Windows are preconfigured with the well-known ATM address and virtual channel defined in the ATM Forum specification for obtaining the LAN emulation configuration server (LECS) address. The LEC uses these addresses in the following order:

1. It can try the well-known LECS address defined in the ATM protocol.

   ```
   47007900000000000000000000000A03E00000100
   ```

2. It can use the well-known virtual channel defined in the ATM protocol.

3. It can query using the integrated local management interface (ILMI).

**IP/ATM preconfigured address**
Windows ATM ARP clients are preconfigured with the ARP/MARS addresses of Windows ATM services. This address is:

This number is a well-known E.164 address that Microsoft has selected to simplify configuration and enhance interoperability of the ATM ARP/MARS with other computers running ATM services and Windows XP Professional or a Windows Server 2003 operating system.

# ATM Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

This section covers:

- ATM standards
- ATM updated technical information

# ATM standards

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ATM standards**
There are three major groups defining implementation standards for ATM in networks: the ATM Forum, the IETF, and the ITU-T.

**ATM Forum**
The ATM Forum is an international not-for-profit group of ATM hardware manufacturers, network software developers, and network service providers, consisting of workgroups that develop and review ATM specifications.

For more information about the ATM Forum and access to ATM Technical Committee specifications, see the ATM Forum Web site.

**Internet Engineering Task Force (IETF)**
The IETF is the standards organization for the Internet. Within the IETF, the IETF IP-over-ATM working group developed standards for IP traffic over ATM networks. The following chart lists RFCs for implementing IP over ATM.

| RFC | Description |
|---|---|
| 2684 | Multiprotocol encapsulation over ATM adaptation layer 5 |
| 1722 | RIP version 2 protocol applicability statement |
| 1754 | IP over ATM Working Group's Recommendations for the ATM Forum's Multiprotocol BOF - Version 1 |
| 1755 | ATM Signaling Support for IP over ATM |
| 2022 | Support for Multicast over UNI 3.0/3.1 based ATM Networks |
| 2225 | Classical IP and ARP over ATM (makes obsolete RFC 1626 and RFC 1577) |

For more information about the IETF and the RFCs regarding IP over ATM, see the IETF Web site.

**International Telecommunication Union, Telecommunication Standardization Sector (ITU-T)**
The ITU-T defines standards for global telecommunication and services. The ITU-T developed B-ISDN and SONET for broadband transport services.

For additional information about the ITU-T and ITU-T recommendations, see the ITU-T Web site.

### Note

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# ATM updated technical information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ATM updated technical information**

The Personal Support Center at the Microsoft Web site provides additional information about ATM. For more information, see Updated technical information.

For more information about asynchronous transfer mode (ATM), see "Asynchronous Transfer Mode" at the Microsoft Windows Resource Kits Web site.

# ATM Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**ATM Troubleshooting**
What problem are you having?

- Host cannot join an emulated LAN.
- Hosts on my IP/ATM network are unreachable.
- A hardware error occurs when my ATM adapter attempts to initialize.

**Host cannot join an emulated LAN.**
**Cause:** The client and switch must be running a user-to-network interface (UNI) based on parameters outlined in ATM Forum UNI 3.1 or 4.0 specifications. Otherwise, the initial point-to-multipoint (PMP) call setup fails and no connection is made. The LAN emulation client (LEC) provided in the Windows Server 2003 family is based on the ATM Forum UNI 3.1 specification.

**Solution:** Use the Atmadm utility to determine if the setup call is failing. The Atmadm utility is used to monitor connections and addresses registered by the ATM Call Manager on an asynchronous transfer mode (ATM) network. You can use it to display statistics for incoming and outgoing calls on ATM adapters. If Atmadm displays few or no packets received on signaling and ILMI VCs, check your switch connectivity. For example:

```
Signaling and ILMI packets sent     = 90718
Signaling and ILMI packets received = 0
```

**See also:** View ATM connection information; Atmadm

**Hosts on my IP/ATM network are unreachable.**
**Cause:** If you have more than one ATM ARP/MARS for each virtual LAN on your network, the ARP caches can become out of synch. This causes clients to be able to only contact clients registered with the same server. Clients registered with another server are unreachable.

**Solution:** Only use one ATM ARP/MARS for each virtual LAN.

**A hardware error occurs when my ATM adapter attempts to initialize.**
**Cause:** Your ATM adapter is running IP/ATM and LAN emulation (LANE) on the same subnet.

**Solution:** IP/ATM and LAN emulation (LANE) must be configured to run on separate subnets.

# Integrating with Other Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Integrating with other networks**
This section covers:

- Novell NetWare Integration
- Services for Macintosh

# Novell NetWare Integration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Novell NetWare integration**
You can use tools provided with the Microsoft® Windows® Server 2003 family to access resources on Novell NetWare networks:

- To find features that have been moved in the Windows Server 2003 family, see New ways to do familiar Novell NetWare Integration tasks.
- For tips about using NWLink or Client Service for NetWare, see Novell NetWare Integration Best Practices.
- For help with specific tasks, see Novell NetWare Integration How To....
- For general background information, see Novell NetWare Integration Concepts.
- For problem-solving instructions, see Novell NetWare Integration Troubleshooting.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# New ways to do familiar Novell NetWare Integration tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for interoperating with Novell NetWare in Windows. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Install and configure the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) | **Network** in Control Panel | Network Connections. For more information, see Install NWLink and Configure NWLink. |
| Install Client Service for NetWare | **Network** in Control Panel | Network Connections. For more information, see Install Client Service for NetWare. |
| Enable or disable NWLink NetBIOS | **Network** in Control Panel | Network Connections. For more information, see Enable direct hosting of IPX. |

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Novell NetWare Integration Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

This section covers:

- Best practices for NWLink
- Best practices for Client Service for NetWare

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Best practices for NWLink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for NWLink**

**If you do not know what value to use for the internal network number, accept the default value of 00000000.**

- You can determine what internal network number your routers are using by typing **ipxroute config** at a command prompt.
- For more information about setting the internal network number, see Configure NWLink.

**Use automatic frame type detection.**

- You should choose to have the frame type automatically detected rather than setting this value manually. If you must use manual frame type detection, determine what type of frame type you should use before you start configuring the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) by typing **ipxroute config** at a command prompt. In addition, verify that all the server computers on your network that use IPX/SPX are using the same frame type.
- For more information, see Configure NWLink.

**Specify the order that Windows accesses network providers and protocols.**

- By changing your provider order, and by changing the order of protocols bound to those providers, you can improve performance. For example, if your LAN connection is enabled to access NetWare and Microsoft Windows networks and uses TCP/IP and IPX, but your primary connection is to a NetWare network, you can move **NetWare or Compatible Network** to the top of the **Network providers** list on the **Provider Order** tab, and move **NWLink IPX/SPX/NetBIOS Compatible Transport Protocol** to the top of the **File and Printer Sharing for Microsoft Networks** binding on the **Adapters and Bindings** tab.
- For more information about modifying the order of network providers, see Modify the network provider order. For more information about modifying protocol bindings, see Modify the protocol bindings order.

**Only install and enable the network protocols that you need.**

- Limiting the number of protocols on your computer enhances network performance and reduces network traffic.
- If Windows encounters a problem with network connectivity, it attempts to establish connectivity by using every network protocol that is installed. By only installing and enabling the protocols that your system can use, Windows does not attempt to connect with protocols it cannot use, and returns status information to you more efficiently.

**When using direct hosting, if you cannot access a server computer running Windows 95 or Windows 98 from a server running Windows Server 2003 , you must enable NetBIOS over IPX/SPX on the computer running Windows 95 or Windows 98.**

- For more information about enabling NetBIOS, see the Windows 95 or Windows 98 documentation.

**When you use NWLink to connect computers running Windows that use Windows-based networking, ensure that NetBIOS broadcast propagation (type 20 packets) is enabled on all the router interfaces that connect the computers running Windows.**

- You can enable NetBIOS broadcast propagation by using Routing and Remote Access. NetBIOS broadcast propagation is enabled by default.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Best practices for Client Service for NetWare

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for Client Service for NetWare**
**When running Novell Directory Services (NDS) administrative utilities on Windows, you must use NetWare client software.**

- To administer a Novell NDS server from a Windows computer using NDS utilities, you need to run the client software from Novell, not Client Service for NetWare.

**When changing a user's NetWare password, synchronize the password for both NetWare and Windows resources.**

- You can use the **Change Password** dialog box that is available by pressing CTRL+ALT+DELETE. User names and passwords should be the same for all NetWare servers.

**When connecting to NetWare resources on NDS trees, confirm that you are using the proper default tree and context.**

- For more information, see Set a default tree and context.

**Use the 32-bit version of command prompt.**

- If you are running Client Service for NetWare and use a command prompt, you should use the 32-bit version (that is, Cmd.exe), which is available on the **Programs** menu, not the 16-bit version (that is, Command.com).

**Follow the best practices for NWLink.**

- For more information about the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink), see Best practices for NWLink.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Novell NetWare Integration How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install and Configure Client Service for NetWare
- Install and Configure NWLink
- Use NetWare Resources

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Install and Configure Client Service for NetWare

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and configure Client Service for NetWare**

- Install Client Service for NetWare
- Set a preferred server
- Set a default tree and context
- Set print options

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

# Install Client Service for NetWare

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Client Service for NetWare**

1. Open Network Connections.
2. Right-click a local area connection and then click **Properties**.
3. On the **General** tab, click **Install**.
4. In the **Select Network Component Type** dialog box, click **Client**, and then click **Add**.
5. In the **Select Network Client** dialog box, click **Client Service for NetWare**, and then click **OK**.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You must be a member of the Administrators group to install Client Service for NetWare.
- When you install Client Service for NetWare, it is installed for all your connections. For connections that you do not want Client Service for NetWare installed, right-click that connection, click **Properties** and, on either the **General** or **Networking** tab, clear the **Client Services for NetWare** check box.
- To confirm that Client Service for NetWare is working properly, at a command prompt, type **net view /network:nw**. You should see a list of available NetWare servers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install and Configure Client Service for NetWare
Install and Configure NWLink
Use NetWare Resources

# Set a preferred server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set a preferred server**

1. Open Client Service for NetWare.
2. Click **Preferred Server**, and in **Select Preferred Server**, enter the preferred server.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open Client Service for NetWare, click **Start**, click **Control Panel**, and then double-click **CSNW**. To use Client Service for NetWare, you must first install it.
- You should set a default tree and context only in an NDS environment; otherwise, you can set a preferred server.
- If you do not want to set a preferred server, click **None**. You are then attached to the nearest available NetWare server, and your interaction with the NetWare network is through that server. If you do not set a preferred server, and do not select **None**, then you are prompted to set one each time you log on. You can set a preferred server at any time.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install and Configure Client Service for NetWare
Install and Configure NWLink
Use NetWare Resources

# Set a default tree and context

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set a default tree and context**

1. Open Client Service for NetWare.
2. Click **Default Tree and Context**, and in **Tree** and **Context**, type your tree and context.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open Client Service for NetWare, click Start, click Control Panel, and then double-click CSNW. To use Client Service for NetWare, you must first install it.
- You should set a default tree and context only in an NDS environment; otherwise, you can set a preferred server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install and Configure Client Service for NetWare
Install and Configure NWLink
Use NetWare Resources

# Set print options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set print options**

1. Open Client Service for NetWare.
2. Under **Print Options**, do one or more of the following:
   - To eject a page at the end of each document, select the **Add Form Feed** check box.
   - To receive notification when a document prints, select the **Notify When Printed** check box.
   - To print a banner page before each document, select the **Print Banner** check box.
3. Under **Login Script options**, to run a login script from the Netware server, select **Run Login Script**.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open Client Service for NetWare, click **Start**, click **Control Panel**, and then double-click **CSNW**.
- To use Client Service for NetWare, you must first install it.
- You must be a member of the Administrators group to set print options.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install and Configure Client Service for NetWare
Install and Configure NWLink
Use NetWare Resources

# Install and Configure NWLink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and configure NWLink**

- Install NWLink
- Configure NWLink
- Enable direct hosting of IPX

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

# Install NWLink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install NWLink**

1. Open **Network Connections**.
2. Right-click a local area connection and then click **Properties**.
3. On the **General** tab, click **Install**.
4. In the **Select Network Component Type** dialog box, click **Protocol**, and then click **Add**.
5. In the **Select Network Protocol** dialog box, click **NWLink IPX/SPX/NetBIOS Compatible Transport Protocol**, and then click **OK**.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You must be a member of the Administrators group to install NWLink.
- When you install NWLink, it is installed for all your connections. If you do not want NWLink installed for a certain connection, right-click that connection, click **Properties** and, on either the **General** or **Networking** tab, clear the **NWLink IPX/SPX/NetBIOS Compatible Transport Protocol** check box.
- To confirm that NWLink is working properly, at a command prompt, type **ipxroute config**. You should see a table with information about the bindings for which NWLink is configured.

**See Also**
**Concepts**
Use NetWare Resources
Install and Configure Client Service for NetWare
Install and Configure NWLink

# Configure NWLink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure NWLink**

1. Open **Network Connections**.
2. Right-click a local area connection and then click **Properties**.
3. On the **General** tab, click **NWLink IPX/SPX/NetBIOS Compatible Transport Protocol**, and then click **Properties**.
4. On the **General** tab, type a value for **Internal network number**, or leave this setting at the default value of **00000000**.
5. Do one of the following:

- Click **Auto frame type detection**, and then click **OK**.
- Click **Manual frame type detection**, and do the following:
    - Click **Add**.
    - In the **Manual Frame Detection** dialog box, in **Frame type**, click a frame type.
    - In **Network number**, type a network number, and then click **OK**.
    - Repeat these steps for each frame type you want to include, and then click **OK**.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- These instructions are for advanced users. New users should accept default values unless told otherwise by their network administrator.

**Notes**

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To configure NWLink, you must first install the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol and be a member of the Administrators group.
- Additional tunable parameters for NWLink are stored in the registry. In most situations, you should not need to modify the defaults. You should use this procedure only if you want to bind NWLink to a different network adapter or to manually change the frame type.
- By default, NWLink automatically detects the frame type used by the network adapter to which it is bound. If NWLink detects no network traffic or if multiple frame types are detected in addition to the 802.2 frame type, NWLink sets the frame type to 802.2.

  You can determine which external network number, frame type, and internal network number your routers are using by typing **ipxroute config** at a command prompt.

**See Also**

**Concepts**

Install NWLink

Using Ipxroute

Setting the frame type

Best practices for NWLink

# Enable direct hosting of IPX

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable direct hosting of IPX**

1. Open Network Connections.
2. Right-click a local area connection, and then click **Properties**.
3. In the list of installed components, clear the **NWLink NetBIOS** check box.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- To enable direct hosting of IPX, you must first install the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink).
- You must be a member of the Administrators group to enable direct hosting of IPX.
- When you clear the **NWLink NetBIOS** check box, the component is disabled for all IPX interfaces. However, you can clear the **NWLink IPX/SPX/NetBIOS Compatible Transport Protocol** check box on a per-interface basis.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install and Configure Client Service for NetWare
Install and Configure NWLink
Use NetWare Resources

# Use NetWare Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use NetWare resources**

- Connect to a NetWare resource
- Connect to a printer attached to a NetWare server
- View NetWare resources
- View current network connections
- Change your NetWare password

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

# Connect to a NetWare resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a NetWare resource**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open **My Network Places**.
2. Do one of the following:
   - Double-click **NetWare or Compatible Network**.
   - Double-click **Entire Network**, and then double-click **NetWare or Compatible Network**.

     Windows Explorer displays Novell Directory Services (NDS) trees with a tree icon and individual NetWare computers with a computer icon.
3. Double-click a tree or volume to see the contents. You can then double-click those contents to see other computers or volumes.
4. Do one of the following:
   - When you find the volume or folder that you want to access, double-click the volume or folder to expand it.
   - To map a local drive to the volume or folder, click the volume or folder, and then, on the **Tools** menu, click **Map Network Drive**. Follow the instructions in the Map Network Drive wizard.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

**Notes**

- To open My Network Places, click **Start**, and then click **My Computer**. Under **Other Places**, click **My Network Places**.
- When you map a network drive, you are connected by default under the user name and password you used to log on. To connect under a different user name, follow the instructions in the Map Network Drive wizard.
- To view or connect to NetWare resources, you must first install Client Service for NetWare and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.

**Using a command line**

1. Open **Command Prompt**.

2. To connect to an individual NetWare volume, at the command prompt type: **net use** *drive**UNCname*|*NetWareName*

   For example, to use the Universal Naming Convention (UNC) naming syntax to redirect drive G to the \Data\Mydata directory of the Thor volume on a server called Nw4, type**: net use G\\nw4\thor\data\mydata**

   The following message indicates that your user name and password cannot be authenticated:

   ```
   The password is invalid for server name\volume name[\directory name...]
   ```

3. To connect with a valid user name and password, at the command prompt, add your user name and password to the command line by typing: */user: UserName Password*

   For example, to use drive G to connect as Annie with the password Marshmallow to the \Data\Mydata directory within the Thor volume on a server called Nw4, type: **net use G\\nw4\ thor\data\mydata /user:annie marshmallow**

4. To connect to a Novell Directory Services (NDS) tree, at the command prompt, type: **net use***drive:*\\ *TreeName\volume.OrgName.OrgName [/u:UserName.OrgName.OrgName [password]]*

| **Value** | **Description** |
|---|---|
| *TreeName* | Name of the tree volume. |
| *OrgName* | Tree location to which you want to connect. |
| *UserName.OrgName.OrgName* | The user name and context for this tree (unless it is your default tree). |

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

**Notes**

- To open a command prompt, click **Start**, point to **All Programs**, point to **Accessories**, and then click **Command Prompt**.
- To view or connect to NetWare resources, you must first install Client Service for NetWare and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.
- When you connect to NetWare file servers, **net use** is equivalent to the NetWare **map** command on MS-DOS-based NetWare workstations.
- You can connect to Novell Directory Services (NDS) trees and individual NetWare volumes that use bindery-style security. NetWare server volumes, directories, and print queues are represented by Universal Naming Convention (UNC) names and use the same command syntax as Windows-based networks.
- You can type the path to an NDS resource with or without labels (that is, **cn**=, **ou**=, and **o**=).
- When you use a command prompt to connect to NetWare file resources, you can use the next available drive letter by replacing the drive letter with an asterisk (*) in the syntax. For example: **net use ***: *UNCname or NetWarename*
- If you want to be prompted for a password, you can replace the password in the command line with an asterisk (*). When you type your password at the prompt, it does not appear on the screen.
- If you are running Client Service for NetWare and use a command prompt, you should use the 32-bit version (that is, Cmd.exe), which is available on the **Programs** menu, not the 16-bit version (that is, Command.com).
- To view the complete syntax for net use, at a command prompt, type:

**net use /?**

**See Also**
**Concepts**
Install Client Service for NetWare
Install NWLink

# Connect to a printer attached to a NetWare server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To connect to a printer attached to a NetWare server

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open **Printers and Faxes**.
2. Double-click **Add Printer**, and then click **Next**.
3. Follow the remaining instructions in the Add Printer wizard.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- To view or connect to NetWare resources, you must first install Client Service for NetWare and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.
- If a printer driver is not available locally for a NetWare printer, you are prompted to install one.

#### Using a command line

1. Open **Command Prompt**. To redirect output from a port to a print queue, at the command prompt type: **net use***ServerName \\PrintQueue*

   For example, to redirect output from LPT1 to the NetWare print queue called Memos on the server Nw4, type: **net use lpt1 \\nw4\memos**

   This is equivalent to the NetWare capture **q=memos s=nw4 l=1** command line.

2. To send files that do not require formatting to LPT1, redirect output with **net use**, and then, at the command prompt, type: **copy** *FileName ServerName*

   For example: **copy myfile.txt lpt1**

3. To copy a file directly to a print queue, redirect output with **net use**, and then, at the command prompt type: **copy***FileName \\PrintQueue*

   For example: **copy myfile.txt \\nw4\memos**

4. To connect to a printer in a Novell Directory Services (NDS) tree, at the command prompt, type: **net use drive:** *\\TreeName\printer.OrgName.OrgName [/u:UserName.OrgName.OrgName [password]]*

| Value | Description |
| --- | --- |
| *TreeName* | Name of the tree *volume*. |
| *OrgName* | Tree location to which you want to connect. |
| *UserName.OrgName.OrgName* | The user name and context for this tree (unless it is your default tree). |

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

**Notes**

- To open a command prompt, click **Start**, point to **All Programs**, point to **Accessories**, and then click **Command Prompt**.
- To view or connect to NetWare resources, you must first install Client Service for NetWare and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.
- You can type the path to a NDS resource with or without labels (that is, **cn=**, **ou=**, **o=**).
- When you run any application that writes directly to a predefined port, the **net use** command works like the NetWare **capture** utility, associating the NetWare print queue with the port.
- If you are running Client Service for NetWare and use a command prompt, you should use the 32-bit version (that is, Cmd.exe), which is available on the **Programs** menu, not the 16-bit version (that is, Command.com).

**See Also**
**Concepts**
Install Client Service for NetWare
Install NWLink
Set print options

# View NetWare resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view NetWare resources**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open **My Network Places** and do one of the following:

   - Double-click **NetWare or Compatible Network**.
   - Double-click **Entire Network**, and then double-click **NetWare or Compatible Network**.

   Windows Explorer displays Novell Directory Services (NDS) trees with a tree icon and individual NetWare computers with a computer icon.

2. Double-click a tree or volume to see the contents. You can then double-click those contents to see other computers or volumes.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open My Network Places, click **Start**, and then click **My Computer**. Under **Other Places**, click **My Network Places**.
- To view or connect to NetWare resources, you must first install Client Service for NetWare and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.

**Using a command line**

- Open **Command Prompt** and do one of the following:

  - To display a list of NetWare file servers, at the command prompt, type**: net view /network:nw**
  - To display volumes on a specific NetWare file server, at the command prompt, type: **net view ***NWServerName***/network:nw**

    For example, to view the volumes on the NetWare server Nw4, type: **net view \\nw4 /network:nw**

  - To display the contents of a directory on a NetWare file server running bindery security, at the command prompt, type: **dir ***DirectoryPath*
  - To display the contents of a directory on a NetWare file server running Novell Directory Services (NDS), enclose the directory path in quotation marks.

For example, at the command prompt, type: **dir "***NDStree******volume.unit.group***"**

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

**Notes**

- To open a command prompt, click **Start**, point to **All Programs**, point to **Accessories**, and then click **Command Prompt**.
- To view or connect to NetWare resources, you must first install Client Service for NetWare and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.
- **Net view**performs the same function as the NetWare **slist** utility.
- You can type the path to an NDS resource with or without labels (that is, **cn**=, **ou**=, **o**=).
- If you are running Client Service for NetWare and use a command prompt, you should use the 32-bit version (that is, Cmd.exe), which is available on the **Programs** menu, not the 16-bit version (that is, Command.com).

**See Also**
**Concepts**
Connect to a NetWare resource
Connect to a printer attached to a NetWare server
Install Client Service for NetWare
Install NWLink
View current network connections

# View current network connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To view current network connections

- Using the Windows interface to view all network connections
- Using the Windows interface to view NetWare connections only
- Using a command line

#### Using the Windows interface to view all network connections

- Open **My Computer**.

### Notes

- To open My Computer, click **Start**, and then click **My Computer**.
- My Computer displays a red X over any disconnected network drives.

#### Using the Windows interface to view NetWare connections only

- Open Windows Explorer, right-click **My Network Places**, and then click **Who Am I?**

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open Windows Explorer, click **Start**, point to **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
- To show or hide desktop icons, right-click somewhere on the desktop, click **Properties**, click the **Desktop** tab, click **Customize Desktop**, and then, under **Desktop icons**, select the check box next to the icon you want to display, or clear the check box next to the icon you want to hide.

#### Using a command line

1. Open Command Prompt.
2. At the command prompt, type: **net use**

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open a command prompt, click **Start**, point to **All Programs**, point to **Accessories**, and then click **Command Prompt**.

- To view or connect to NetWare resources, you must first install Client Service for NetWare and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.

- If you are running Client Service for NetWare and use a command prompt, you should use the 32-bit version (that is, Cmd.exe), which is available on the **Programs** menu, not the 16-bit version (that is, Command.com).

- To view the complete syntax for **net use**, at a command prompt, type: **net use /?**

**See Also**
**Concepts**
Install Client Service for NetWare
Install NWLink

# Change your NetWare password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change your NetWare password**

- On a NetWare bindery server
- On a NetWare Novell Directory Services (NDS) tree

**On a NetWare bindery server**

1. Open **Command Prompt**.
2. Change to the drive for the NetWare server, and then type:

   **cd \public**

3. Type the following, where *NetWareServer* is the name of the NetWare server for which you want to change your password:

   **setpass***NetWareServer*

4. When prompted for each, type your old password, a new password, and the new password again.

   A message confirms that you have successfully changed your password.

5. If prompted, type **y** to change your password on other NetWare servers that also use your old password.

   Or, to leave your old password unchanged on the other NetWare servers, type **n**.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- To open a command prompt, click **Start**, point to **All Programs**, point to **Accessories**, and then click **Command Prompt**.
- To view or connect to NetWare resources, you must first install Client Service for NetWare and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.
- To change your password on more than one server, connect to all the servers before you type **setpass**.
- If you are running Client Service for NetWare and use a command prompt, you should use the 32-bit version (that is, Cmd.exe), which is available on the **Programs** menu, not the 16-bit version (that is, Command.com).

**On a NetWare Novell Directory Services (NDS) tree**

1. Press CTRL+ALT+DELETE.
2. Click **Change Password**.

3. In **Log on to**, click **NetWare or Compatible Network**.

4. In **Old Password**, type your current password.

5. In **New Password**, type your new password and, in **Confirm New Password**, type your new password again.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

### Notes

- Your password changes on all NDS trees to which you are currently connected. If the old password you specify does not match your current password on any of those trees, you are prompted to type the old password for those trees.

- To view or connect to NetWare resources, you must first install Client Service for NetWare and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.

**See Also**
**Concepts**
Install NWLink
Install Client Service for NetWare
Passwords

# Novell NetWare Integration Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Novell NetWare integration overview
- The NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)
- Install Client Service for NetWare

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Novell NetWare integration overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Novell NetWare integration overview**

The Microsoft® Windows® Server 2003 family and Microsoft® Windows® XP provide several services that enable computers running Windows to coexist and interoperate with Novell NetWare networks and servers. Some of these services are included with the Windows Server 2003 family and Windows XP. Other services are available as separate products.

- **NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)**

  Included with the Windows Server 2003 family operating systems and Windows XP is the Windows implementation of the IPX/SPX protocol. NWLink supports connectivity between computers running Windows and computers running NetWare and compatible systems. You can also use NWLink as a protocol that connects Windows-based computers and computers running MS Client for DOS. For more information about NWLink, see The NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink).

- **Client Service for NetWare**

  Included with the Windows Server 2003 family and Windows XP, Client Service for NetWare makes direct connections from workstations to file and printer resources on NetWare servers. With Client Service for NetWare, you can access servers running either Novell Directory Services (NDS) or bindery security. Logon script support is also included.

- **Microsoft Directory Synchronization Services**

  Included with Services for NetWare (a separate product), Microsoft Directory Synchronization Services (MSDSS) can perform a one-time migration of directory objects or synchronize directory objects from both Novell Directory Services (NDS) and Bindery environments with Microsoft Active Directory.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# The NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)**
This section covers:

- NWLink overview
- Understanding NWLink
- Using NWLink

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# NWLink overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**NWLink overview**
NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) is an implementation of Novell's Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) and NetBIOS protocols. Windows clients can use NWLink to access client and server applications running on Novell NetWare servers. NetWare clients can use NWLink to access client and server applications running on Windows servers. With NWLink, computers running Windows can communicate with other network devices, such as printers, that use IPX/SPX. You can also use NWLink in small networks that only use Windows and other Microsoft client software.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Understanding NWLink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding NWLink**

The NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) is a Network Driver Interface Specification (NDIS)-compliant, native 32-bit implementation of Novell's IPX/SPX protocol. NWLink supports two networking application programming interfaces (APIs): NetBIOS and Windows Sockets. These APIs allow communication among computers running Windows and between computers running Windows and NetWare servers.

The NWLink transport driver is an implementation of the lower-level NetWare protocols, which include IPX, SPX, RIPX (Routing Information Protocol over IPX), and NBIPX (NetBIOS over IPX). IPX controls addressing and routing of packets of data within and between networks. SPX provides reliable delivery through sequencing and acknowledgments. NWLink provides NetBIOS compatibility with a NetBIOS layer over IPX.

**Interoperating with NetWare resources**

Depending on what platform you are using and what resources that you want to access, you may need to use NWLink in conjunction with other tools. The following table lists several interoperability options that use NWLink.

| Platform | Running | Can connect to |
|---|---|---|
| Windows | NWLink | Client/server applications running on a NetWare server |
| Windows | NWLink and Client Service for NetWare | NetWare servers for file and print services |
| NetWare client | IPX with NetBIOS, Named Pipes, or Windows Sockets support | Computers running Windows (with NWLink installed) running IPX applications such as Microsoft SQL Server |
| NetWare client | IPX | Servers running Windows Server 2003 (with NWLink installed) for file and print services |

For a Windows client to access file and print resources on a NetWare server, Client Service for NetWare must be installed on the Windows client in addition to NWLink. For more information, see Novell NetWare integration overview.

**Interoperating with Microsoft resources**

By default, the file and print sharing components of Windows use NetBIOS over IPX to send file and print sharing messages.

Alternatively, you can disable NetBIOS so that the file and print sharing messages are sent directly over IPX. This is known as direct hosting. While direct hosting may be more efficient, it causes an interoperability issue: A direct hosting client can only connect to a direct hosting server. Direct hosting clients include computers running Microsoft Network Client for MS-DOS, Windows for Workgroups, Windows 95, and Windows 98. Direct hosting servers include computers running Microsoft Network Client for MS-DOS, and all Windows-based computers.

The following table lists several interoperability options that use IPX.

| Platform | Running | Can connect to |
|---|---|---|
| Microsoft Network Client for MS-DOS<br><br>Windows for Workgroups<br><br>Windows 95<br><br>Windows 98 | IPX only (direct hosting) | File and print shared resources on computers running:<br><br>Microsoft Network Client for MS-DOS<br><br>Windows for Workgroups<br><br>Windows 95<br><br>Windows 98<br><br>Windows NT<br><br>Windows 2000 |

| | | |
|---|---|---|
| Microsoft Network Client for MS-DOS<br>Windows for Workgroups Windows 95<br>Windows 98<br>Windows NT<br>Windows 2000 | NetBIOS over IPX | File and print shared resources on computers running:<br>Microsoft Network Client for MS-DOS<br>Windows for Workgroups<br>Windows 95<br>Windows 98<br>Windows NT<br>Windows 2000 |

**Important**

- You cannot use direct hosting of IPX to gain access to resources on any computer acting as a direct host server from a computer running Windows. Computers running Windows do not include a direct hosting client for NWlink. For example, while you can connect by using direct hosting of IPX to resources on a computer running Windows from a computer running Windows 95 or Windows 98 (where NetBIOS is disabled by default for performance reasons), you cannot connect by using direct hosting of IPX to the computer running Windows 95 or Windows 98 from the computer running Windows. For detailed instructions about enabling direct hosting, see Enable direct hosting of IPX.

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.

- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.

- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Using NWLink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using NWLink**
This section covers:

- Installing NWLink
- Configuring NWLink
- Routing and NWLink

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Installing NWLink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing NWLink**
You can install the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) when you install a Windows Server 2003family operating system, or you can install NWLink later. You must be logged on as a member of the Administrators group to install and configure NWLink.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Configuring NWLink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring NWLink**
This section covers:

- Setting the frame type
- Setting the external and internal network numbers

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Setting the frame type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting the frame type**

The frame type defines the way in which the network adapter, in a computer running Windows, formats data to be sent over a network. To communicate between a computer running Windows and NetWare servers, you need to configure the IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) on the computer running Windows with the same frame type as the one used by the NetWare servers.

The following table lists the topologies and frame types supported by NWLink.

| Topology | Supported frame type |
|---|---|
| Ethernet | Ethernet II, 802.3, 802.2, and Sub Network Access Protocol (SNAP), which defaults to 802.2 |
| Token ring | 802.5 and SNAP |
| Fiber Distributed Data Interface (FDDI) | 802.2 and 802.3 |

**Note**

- On Ethernet networks, the standard frame type for NetWare 2.2 and NetWare 3.11 is 802.3. Starting with NetWare 3.12, the default frame type was changed to 802.2.

You can choose to automatically detect or manually configure the frame type. However, the frame type is automatically detected when NWLink is loaded. If multiple frame types are detected in addition to the 802.2 frame type, NWLink defaults to the 802.2 frame type.

If the frame type is manually configured, a computer running Windows can use multiple frame types simultaneously.

For more information about configuring the frame type, see Configure NWLink.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Setting the external and internal network numbers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting the external and internal network numbers**

The IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) uses two types of IPX network numbers for routing purposes: external network numbers, referred to as the **Network number** in the **Manual Frame Detection** dialog box, and internal network numbers, referred to as the **Internal network number** in the **NWLink IPX/SPX/NetBIOS Compatible Transport Protocol Properties** dialog box. Both are hexadecimal numbers, with 1 to 8 digits (1 to FFFFFFFE).

The external network number is associated with physical network adapters and networks. To communicate with each other, all computers on the same network that use a given frame type must have the same external network number. If your computer has multiple network adapters that are connected to different networks, you must assign an external network number to each configured frame type and network adapter combination on your computer. If you do not know the appropriate numbers to use, see your NetWare documentation. If you do not set an external network number, the number is automatically detected by the Windows operating system.

The internal network number, also called a virtual network number, identifies a virtual network inside a computer. On a server computer running Windows, programs advertise themselves as being located on the virtual network, not a physical network. When you assign internal network numbers, you gain more efficient routing to the programs in a computer with multiple network adapters that are connected to multiple, interconnected networks. By default, the internal network number is 00000000. However, in each of the following situations, you need to manually assign a unique, nonzero internal network number:

- Your computer is acting as a server (running Windows) for a program that uses the NetWare Service Advertising Protocol (SAP), such as SQL Server or SNA Server.
- Your computer is acting as an IPX router that uses the Windows routing.

For more information about setting the external and internal network numbers, see Configure NWLink.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Routing and NWLink

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing and NWLink**
This section covers:

- Using the Routing Information Protocol
- Using the Service Advertising Protocol
- Using Ipxroute

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Using the Routing Information Protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the Routing Information Protocol**
In networks that use only the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink), routing information is propagated by using the Routing Information Protocol (RIP). NWLink uses the Routing Information Protocol over IPX (RIPX) to determine the best router to use when forwarding IPX traffic to another IPX network.

**RIPX and NWLink**
RIPX is used to determine the frame type and external network number for each interface that is configured for autodetection. When NWLink loads, it sends out a RIPX request to discover the frame type and external network number. IPX routers (and NetWare servers) respond to the request with packets that use the IPX router frame type and contain the external network number of the local network.

**RIPX and routers**
When the IPX host determines the IPX network number of the destination, it compares the destination IPX network number with its own attached external network number. If they are different, the IPX host must forward the IPX packet to an IPX router. To determine the best router to use to forward the IPX traffic, the IPX host sends out a RIP request for the destination IPX network number. All IPX routers on the network of the IPX host that have a route to the destination network number send back a response to the IPX host. The IPX host chooses the best response and forwards the IPX packet to the router that sent the best response.

When the Windows Routing and Remote Access service is installed, the computer running Windows acts as a RIP-based IPX router. RIPX is used to listen for route announcements and periodically send route announcements to maintain an IPX routing table.

For more information about routing with Windows, see Checklist: Installing and configuring the router.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Using the Service Advertising Protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the Service Advertising Protocol**

The NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) uses the Service Advertising Protocol (SAP) to locate the nearest server upon startup and to locate either all services or all services of a specific type. When the Windows Routing and Remote Access service is installed, the computer running Windows also uses SAP to listen for SAP advertisements and periodically make SAP advertisements to maintain a table of available services on the network. If you have installed Routing and Remote Access, you can run just SAP advertising and routing by disabling all other functions of the Routing and Remote Access.

The SAP Agent is a network service that allows services on a server running Windows Server 2003 to advertise themselves by using periodic SAP advertisements. Without the SAP Agent, the services on a computer running Windows are unavailable to computers running NetWare client software as well as computers running Windows that are configured to use just the NWLink protocol.

You can manually install the SAP Agent by using the properties of a connection, but is usually automatically installed when a service that uses the SAP Agent is installed. An example is Microsoft Exchange Server. To make the computer running Microsoft Exchange Server available to computers running NetWare client software and computers running Windows that are configured to use just the NWLink protocol, the Microsoft Exchange Server setup program automatically installs the SAP Agent. Once installed, the Microsoft Exchange Server service uses the SAP Agent to periodically advertise its presence and location so that it can be contacted by Microsoft Exchange clients on an IPX network.

When Routing and Remote Access is running, you can configure SAP settings per network adapter, such as whether to forward broadcasts from one adapter to another or whether to disable SAP over an adapter. These settings are configured through the Routing and Remote Access console. You cannot configure the SAP Agent.

The following table summarizes the relationship between the SAP Agent and Routing and Remote Access.

| Status of Routing and Remote Access/SAP Agent | Routing method |
|---|---|
| Neither SAP Agent nor Routing and Remote Access is running | Services are not advertised/routed. |
| SAP Agent is running (but not Routing and Remote Access) | SAP Agent advertises/routes services. |
| Routing and Remote Access is running (but not SAP Agent) | Routing and Remote Access advertises/routes services. |
| Both SAP Agent and Routing and Remote Access are running | Routing and Remote Access advertises/routes services. |

**Note**

- You do not need to install the SAP Agent service separately when Routing and Remote Access is running, since Routing and Remote Access assumes the responsibilities of the SAP Agent.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Using Ipxroute

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Ipxroute**

You can modify the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) settings that affect routing with the Ipxroute.exe command-line utility (the **ipxroute** command). The Ipxroute utility provides the same functionality as the Route.exe command-line utility that is supplied by Novell for its MS-DOS-based clients.

The Ipxroute utility manages the source routing variables of NWLink. Ipxroute.exe is installed only if the NWLink transport protocol is bound to a network adapter.

The Ipxroute utility changes these settings for the current session only. When you log off, the settings are lost. To permanently change NWLink settings, you must set them in the registry.

For information about the **ipxroute** command, see Ipxroute.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Client Service for NetWare

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Client Service for NetWare**
This section covers:

- Client Service for NetWare overview
- Using Client Service for NetWare

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Client Service for NetWare overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Client Service for NetWare overview**
A user who works locally on a computer running Windows can use Client Service for NetWare to gain direct access to NetWare file and print resources, both on Novell Directory Services (NDS) trees and on servers with bindery security.

Client Service for NetWare depends on and works with the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink). NWLink is an implementation of the Internetwork Packet Exchange (IPX), Sequenced Packet Exchange (SPX), and NetBIOS transport protocols used by the NetWare network. The Microsoft implementations of these protocols can coexist seamlessly with other protocols on the same network adapter.

**Note**

- Client Service for NetWare does not support the IP protocol and therefore cannot be used to interoperate with NetWare 5.*x* in an IP-only environment. To do this, you must load the IPX protocol onto the NetWare 5.*x* server, or use a redirector that is compatible with NetWare core Protocol (NCP) and supports native IP.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Using Client Service for NetWare

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Client Service for NetWare**
This section covers:

- Installing Client Service for NetWare
- Configuring Client Service for NetWare
- Connecting Directly to NetWare Resources

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Installing Client Service for NetWare

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Client Service for NetWare**

You have the option to install client service for NetWare when you install a Windows Server 2003 operating system, or you can install Client Service for NetWare later. You must be logged on as a member of the Administrators group to install and configure Client Service for NetWare. For more information about installing Client Service for NetWare after you have already installed a Windows Server 2003 operating system, see Install Client Service for NetWare.

During installation of Client Service for NetWare, the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) is installed if it is not already on the server. For more information about NWLink, see The NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink).

By default, the NetWare network is placed first in the network search order.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Before you install Client Service for NetWare on a computer, remove any existing client software that is compatible with NCP, including NetWare client software, from the computer.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Configuring Client Service for NetWare

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Client Service for NetWare**
When you first log on after Client Service for NetWare is installed, you are prompted to set your default tree and context or your preferred server. The tree and context define the position of the user object for the user name you use to log on to a Novell Directory Services (NDS) tree. A preferred server is the NetWare server to which you are automatically connected when you log on, if your network does not use NDS.

You can have either a default tree and context or a preferred server, but not both. (In NDS environments, you set a default tree and context.) If you select a default tree and context, you can still access NetWare servers that use bindery security.

For more information, see Set a preferred server and Set a default tree and context. To change your tree and context at a later time, you can use Client Service for NetWare (in Control Panel).

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Connecting Directly to NetWare Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connecting directly to NetWare resources**
Novell Directory Services (NDS) trees (as well as NetWare servers running bindery security) appear under NetWare or Compatible Network in Windows Explorer. You can double-click a tree name to expand it, and then double-click any container object to expand its contents and structure. You can connect to and assign a local drive to any volume, folder, or directory map object anywhere in the tree hierarchy (for which you have credentials). For more information, see Connect to a NetWare resource.

To connect to an NDS printer, you can use the Add Printer wizard, just as you would to connect to any network printer. For more information, see Connect to a printer attached to a NetWare server.

If you have a default tree and context, once you have logged on, you do not need to log on again or supply another password to access any volume in your default tree. If you access another tree, you are prompted to supply a full context (including user name) for that tree.

This section covers:

- Changing the NetWare password
- Using logon scripts
- Managing NetWare file attributes
- Running NetWare utilities and programs that can use NetWare

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Changing the NetWare password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Changing the NetWare password**
Users who use client service for NetWare to directly access NetWare resources can change their passwords on NDS trees on the network. For more information, see Change your NetWare password.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Using logon scripts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using logon scripts**

When a user, running Client Service for NetWare to directly access NetWare resources, first makes a connection to a particular NetWare server, the user's logon script (if any) runs.

### Note

- If you are running Client Service for NetWare and use a command prompt, you should use the 32-bit version (Cmd.exe), which is available on the **Programs** menu, not the 16-bit version (Command.com). While running Client Service for NetWare, if you run a batch file (such as a logon script) that opens an instance of the command prompt, the batch file should begin with the command **#cmd**, not **#command**.

### Important

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Managing NetWare file attributes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing NetWare file attributes**

When you use a computer running Client Service for NetWare to directly access NetWare servers, you can use NetWare utilities (such as the **filer** and **rights** commands), to set attributes that are not supported by the mapping of Windows-to-NetWare file rights. For more information about other supported utilities, see Running NetWare utilities and programs that can use NetWare.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Running NetWare utilities and programs that can use NetWare

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Running NetWare utilities and programs that can use NetWare**

On servers running Windows Server 2003 and Client Service for NetWare, you can run some standard NetWare utilities from a command prompt and many programs that can use NetWare. Client Service for NetWare does not support NetWare 4.*x* or 5.*x* utilities. In addition, for some administrative functions, you must use the management tools included with your Windows Server 2003 family operating system.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Novell NetWare Integration Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Client Service for NetWare does not start.
- Client Service for NetWare starts, but no servers are found.
- Client Service for NetWare starts, but only some resources are found.
- Cannot use the NetWare map command to connect to file or print resources.
- The network does not start when you start your computer.
- A computer running the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) cannot connect to other computers.

**Client Service for NetWare does not start.**
**Cause:** One of the required services or protocols might be unavailable.

**Solution:** Try to start Client Service for NetWare manually. If that fails, use Event Viewer to look at the system log and troubleshoot the problem.

**See also:** Event Viewer

**Client Service for NetWare starts, but no servers are found.**
**Cause:** The frame type or SAP might be configured incorrectly.

**Solution:** Try the following:

- View the network adapter load line in the Autoexec.ncf file of the NetWare server to verify that you are using the correct frame type for the server. For example, suppose a network adapter load line for a server is:

  ```
  load 3C503 FRAME=ETHERNET_802.3 NAME=ETH
  ```

  In this case, the server is bound to a 3Com 503 Ethernet adapter that accepts the raw 802.3 frame format.

  To see the frame type that is set for your network adapter, start Network Connections in Control Panel. If the frame type is detected automatically and the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) detects any frames of type 802.2 or no frames at all, NWLink sets the frame type to 802.2. If the network adapter receives frames of type 802.2, but your NetWare network uses a different frame type, you must set the frame type manually.

  To determine which frame type your routers are using on a server running Windows Server 2003 , at a command prompt, type:

  **ipxroute config**

- If you cannot see NetWare file server resources, ensure that SAP is enabled on all the router interfaces that connect the computers running Windows and the server computers running NetWare software. (A Windows remote access server, which is a server running Windows Server 2003 and the Routing and Remote Access service, is an IPX router.) You can enable SAP on an interface by using Routing and Remote Access.

**See also:** Configure NWLink

**Client Service for NetWare starts, but only some resources are found.**
**Cause:** The tree and context or SAP might be configured incorrectly.

**Solution:** Try the following:

- In a Novell Directory Services (NDS) environment, verify that you are using the correct tree and context.
- For each IPX router, view the SAP service table and verify that all of the services that should be learned from SAP are present. If some of the services are not present, you should verify that SAP filtering is not preventing the propagation of these services. The SAP service table is not the same as running the NetWare **slist** command or running an NDS query for a list of all the servers in the tree.

**See also:** View routing tables

**Cannot use the NetWare map command to connect to file or print resources.**
**Cause:** You might be using the wrong command prompt, or the default environment size of the command prompt window might be too small.

**Solution:** The default environment for 16-bit programs is too small to accommodate the mapping table created by the **map** command. Designate the 32-bit command prompt (that is, Cmd.exe) as the permanent command interpreter, and reset the default environment size allocated to the command prompt window. An environment of 4,096 bytes is large enough to accommodate the **map** command, the mapping table, and the command interpreter.

To make these changes to the environment, type the following line in Config.nt:

**shell=%systemroot%\system32\cmd.exe /e:4096**

This line sets Cmd.exe as the command interpreter for the window as long as it remains open or until you run another **shell** command. The line also permanently allocates 4,096 bytes of environment space to 16-bit programs you run in the window.

**The network does not start when you start your computer.**
**Cause:** You might have duplicate computer names.

**Solution:** You must specify a computer name that is not the same as the name of another computer on the network or the same as the name of a workgroup or a domain on the network.

**A computer running the NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) cannot connect to other computers.**
**Cause:** Routing might be configured incorrectly.

**Solution:** Try the following:

- Verify that you have NetBIOS broadcast propagation (type 20 packets) enabled on all the router interfaces between the computers. (A Windows remote access server, which is a server running Windows Server 2003 and the Routing and Remote Access service, is an IPX router.) You can enable NetBIOS broadcast propagation in Routing and Remote Access.
- Verify that you do not have more than eight routers between the computers running Windows.

**Important**

- Client Service for NetWare is not available on 64-bit versions of Windows Server 2003 and Windows XP.
- Gateway Service for NetWare is included in Windows 2000 Server. It is not included in the Windows Server 2003 family.
- Computers running operating systems in the Windows Server 2003 family cannot be IPX routers.

# Services for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Services for Macintosh**
Services for Macintosh lets x86-based and Apple Macintosh clients share files and printers and remotely connect to a Microsoft® network. The three components of Services for Macintosh are: File Server for Macintosh, Print Server for Macintosh, and the AppleTalk Protocol.

- Before you install Services for Macintosh, see Services for Macintosh Checklists.
- To find features that have been moved in the Microsoft® Windows Server 2003 family, see New ways to do familiar Services for Macintosh tasks.
- For tips about using Services for Macintosh, see Services for Macintosh Best practices.
- For help with specific tasks, see Services for Macintosh How To....
- For general background information, see Services for Macintosh Concepts.
- For problem-solving instructions, see Services for Macintosh Troubleshooting.

**Notes**

- This feature is not included on computers running the Microsoft Windows Server 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Services for Macintosh Checklists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

This section provides checklists to help you do the following tasks:

- Checklist: Installing and configuring File Server for Macintosh
- Checklist: Installing and configuring Print Server for Macintosh
- Checklist: Installing and configuring AppleTalk networking and routing
- Checklist: Installing and configuring remote access for Macintosh

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Checklist: Installing and configuring File Server for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring File Server for Macintosh**

| Step | Resource |
|---|---|
| Prepare for installation.<br>• Ensure that there is an NTFS partition on each computer that will share files with a Macintosh computer.<br>• Ensure that an installation compact disc for your operating system or a network location with installation files for your operating system is available. | Macintosh-accessible volumes |
| Install File Server for Macintosh. | Install File Server for Macintosh |
| Install the client component of Services for Macintosh on the Macintosh client computers. | Install authentication files on the Macintosh client |
| Configure File Server for Macintosh. | Create a new logon message for Macintosh clients \| Set session limits for Macintosh clients \| Make new extension associations for x86-based and Macintosh files |
| Configure Macintosh-accessible volumes. | Create a Macintosh-accessible volume \| Set security options for Macintosh volumes |
| Set security on volumes. | Set security options for Macintosh volumes \| Set permissions on a Macintosh-accessible volume or folder |

# Checklist: Installing and configuring Print Server for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring Print Server for Macintosh**

| Step | Reference |
|---|---|
| Ensure that an installation compact disc for your operating system or a network location with installation files for your operating system is available. | |
| Install Print Server for Macintosh. | Install Print Server for Macintosh |
| Create a printer on the computer running Services for Macintosh. | Create a printer |
| *(optional)* Create a user account for Macintosh print jobs. | Set up a user account for Macintosh print jobs |
| *(optional)* Capture printers. | Capture or release an AppleTalk printing device |
| *(optional)* Create high and low priority printers. | Prioritize print jobs |
| *(optional)* Create printing pools. | Creating printing pools |

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Checklist: Installing and configuring AppleTalk networking and routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring AppleTalk networking and routing**

| Step | Reference |
|---|---|
| Review key concepts. | AppleTalk network media;<br>AppleTalk networks;<br>Phase 2 AppleTalk networks |
| Install and configure the AppleTalk Protocol. | Install the AppleTalk protocol and routing;<br>Configure AppleTalk protocol;<br>Configure AppleTalk routing. |
| Assign network numbers and zones. | Assigning AppleTalk network numbers and ranges;<br>Assigning AppleTalk zones |
| Seed the network. | Working with AppleTalk seed routers |
| Create a router record. | AppleTalk routing information;<br>Creating an AppleTalk router record |

# Checklist: Installing and configuring remote access for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring remote access for Macintosh**

| Step | Reference |
|---|---|
| Review key concepts | AppleTalk over PPP |
| Install and configure the AppleTalk Protocol | Install the AppleTalk protocol and routing;<br><br>Configure AppleTalk protocol |
| Install and configure remote access for Macintosh | Install remote access for Macintosh;<br><br>Configure remote access for Macintosh |
| Allow plaintext password authentication | Configure plaintext password authentication for Macintosh remote access |
| Review key remote access concepts | Remote Access |

# New ways to do familiar Services for Macintosh tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring Services for Macintosh. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0.

**Note**

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Configure File Server for Macintosh | MacFile menu in File Manager | Computer Management, System Tools, Shared Folders. |
| Administer File Server for Macintosh | MacFile in Control Panel or MacFile menu in Server Manager | Computer Management, System Tools, Shared Folders. |
| Create an AppleTalk printer | The Printers folder (add an AppleTalk printer) | Printers, Add Printer. |
| Configure AppleTalk Protocol properties | Network utility in Control Panel | Network Connections, right-click Local Area Connection, Properties, AppleTalk Protocol. |
| Set up user accounts for Macintosh users | User Manager for Domains | Local Accounts or Active Directory Users and Computers |
| Enable AppleTalk users to remotely dial in | Not available in 4.0. | Routing and Remote Access. If you have AppleTalk and remote access, this feature is installed by default. |

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Services for Macintosh Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Avoid sharing the root directory of hard drives**

  While it is easy to share the root directory of a hard drive, it makes troubleshooting the volume much more difficult. When a volume is not working, you can move it to a new directory and delete the old volume and directory. Once the old volume is moved, you can rename the new volume with the old name. This is extremely fast and easy to do, but it will not work on root volumes.

- **Verify routing**

  Verify with your network administrator that your routers are routing AppleTalk if you want communication across the LAN. Also verify, if you will be seeding your network with a computer running Services for Macintosh, all ranges must match for the range. There can be no overlapping between segments and each segment's seed router must contain identical range numbers. Also, it is more efficient for a router to seed an AppleTalk network than it is for a computer running Services for Macintosh.

- **Using POSIX file names**

  Do not use POSIX file names in the same directory tree that Macintosh users can access through Macintosh-accessible volumes. The POSIX subsystem is case-sensitive (that is, you can create one file called accounts, one called ACCOUNTS, one called Accounts, and so on.).

- **Do not share the interrupt setting for the LocalTalk card**

  The interrupt setting used for the LocalTalk card should not be shared with any other device.

# Services for Macintosh How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Install and Configure Services for Macintosh
- Set up Macintosh Client Software
- Manage Macintosh-accessible Volumes
- Associate Windows File Types with Macintosh File Types
- Administer Macintosh Clients
- Manage AppleTalk Printing
- Manage Remote Access for Macintosh

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Install and Configure Services for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install and configure Services for Macintosh**

- Install Print Server for Macintosh
- Install File Server for Macintosh
- Install the AppleTalk protocol and routing
- Install remote access for Macintosh
- Configure AppleTalk protocol
- Configure AppleTalk routing
- Configure remote access for Macintosh

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Install Print Server for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Print Server for Macintosh**

1. Open the Windows Components wizard.
2. In **Components**, click **Other Network File and Print Services** (but do not select or clear its check box), and then click **Details**.
3. Select the **Print Services for Macintosh** check box.

   To install File Server for Macintosh, you can also select the **File Services for Macintosh** check box.
4. Click **OK**, and then click **Next**.

## Notes

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.
- If you have not already installed the AppleTalk protocol, the protocol is installed automatically when you install Print Server for Macintosh.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install File Server for Macintosh
Install the AppleTalk protocol and routing
Checklist: Installing and configuring Print Server for Macintosh

# Install File Server for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install File Server for Macintosh**

1. Open the Windows Components wizard.
2. In **Components**, click **Other Network File and Print Services** (but do not select or clear its check box), and then click **Details**.
3. Select the **File Services for Macintosh** check box.

   To install Print Server for Macintosh, you can also select the **Print Services for Macintosh** check box.
4. Click **OK**, and then click **Next**.

## Notes

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.
- You must have an NTFS partition installed before you install File Server for Macintosh.
- When File Server for Macintosh and TCP/IP are installed, the Apple Filing Protocol (AFP) over TCP/IP is enabled automatically.
- If you have not already installed the AppleTalk protocol, the protocol is installed automatically when you install File Server for Macintosh.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install Print Server for Macintosh
Install the AppleTalk protocol and routing
Checklist: Installing and configuring File Server for Macintosh

# Install the AppleTalk protocol and routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install the AppleTalk protocol and routing**

1. Open Network Connections.
2. Right-click any **Local Area Connection** icon, and then click **Properties**.
3. On the **General** tab, click **Install**.
4. In **Select Network Component Type**, click **Protocol**, and then click **Add**.
5. In **Select Network Protocol**, click **AppleTalk Protocol**, and then click **OK**.

## Notes

- To open Network Connections, click **Start**, point to **Settings**, and click **Network Connections**.
- There are two additional AppleTalk networking components in Services for Macintosh: File Server for Macintosh and Print Server for Macintosh. To fully integrate all three AppleTalk networking components, see Related Topics.
- Installing the AppleTalk protocol on a LAN connection installs and enables it for all LAN connections on your computer. To uninstall AppleTalk, right-click the connection icon, click **Properties**, and clear the **AppleTalk Protocol** check box.
- You can install the AppleTalk protocol without installing File Server for Macintosh or Print Server for Macintosh. For example, you can install the AppleTalk protocol for routing purposes only, or for dial-up (AppleTalk over PPP) purposes only, without needing to install or run File Server for Macintosh or Print Server for Macintosh.
- By installing Print Server for Macintosh, you automatically install the AppleTalk network protocol.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure AppleTalk protocol
Install Print Server for Macintosh
Install File Server for Macintosh
Checklist: Installing and configuring AppleTalk networking and routing

# Install remote access for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install remote access for Macintosh**

1. Open Routing and Remote Access.
2. Right-click your server, and then click **Configure and Enable Routing and Remote Access**.
3. Follow the instructions in the Routing and Remote Access Server Setup Wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- You must install AppleTalk on your computer before you can install remote access for Macintosh.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Installing and configuring remote access for Macintosh
Configure remote access for Macintosh
Remote Access

# Configure AppleTalk protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure AppleTalk protocol**

1. Open Network Connections.
2. Right-click **Local Area Connection**, and then click **Properties**.
3. Select the **AppleTalk Protocol** check box, and then click **Properties**.

   AppleTalk protocol properties, such as a default adapter type and a default zone, are configured per connection.
4. If you want to use this connection as the default adapter, select the **Accept inbound connections on this adapter** check box. This is the default behavior.
5. In **This system will appear in zone**, select a zone for the system.

## Notes

- To open Network Connections, click **Start**, point to **Settings**, and click **Network Connections**.
- You cannot clear the **Accept inbound connections on this adapter** check box by clicking it. If you want to stop an adapter from accepting inbound connections and you have a second adapter installed, you can clear the check box by making the second adapter accept inbound connections. Because only one adapter at a time can be bound to AppleTalk, the first adapter stops accepting inbound connections. If you want to clear the check box but you have no other adapters installed, you must uninstall AppleTalk.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Checklist: Installing and configuring AppleTalk networking and routing
Install the AppleTalk protocol and routing
Remote Access

# Configure AppleTalk routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure AppleTalk routing**

1. Open Routing and Remote Access.
2. Under **Routing and Remote Access**, double-click your server and right-click **AppleTalk Routing**.
3. Click **Enable AppleTalk Routing**.
4. In the **Adapters** list, right-click an adapter, and then click **Properties**.
5. Configure seed routing, network number allocation, and the zone list as appropriate for the computer.

### Note

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the AppleTalk protocol and routing
Configure AppleTalk protocol
Remote Access

# Configure remote access for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure remote access for Macintosh**

1. Open Routing and Remote Access.
2. Right-click your server, and then click **Properties**.
3. On the **AppleTalk** tab, configure remote access options as appropriate for the computer, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- You must install AppleTalk on your computer before you can configure remote access for Macintosh.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the AppleTalk protocol and routing
Configure AppleTalk protocol
Remote Access Overview

# Set up Macintosh Client Software

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up Macintosh client software**

- Gain access to authentication files
- Install authentication files on the Macintosh client

# Gain access to authentication files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To gain access to authentication files**

1. On the **Macintosh Apple** menu, double-click **Chooser.**
2. Double-click the **AppleShare** icon, and then click the AppleTalk zone in which the computer running Services for Macintosh resides.
3. On the list of file servers, click the name of the computer running Services for Macintosh, and then click **OK**.
4. Click **Registered User** or **Guest**, as appropriate, and then click **OK**.
5. Click the Microsoft UAM volume, and then click **OK**.
6. Close the **Chooser** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Services for Macintosh Authentication
Install authentication files on the Macintosh client

# Install authentication files on the Macintosh client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install the authentication files on the Macintosh client**

1. Open Chooser, click **AppleShare**, select the AppleTalk zone in which the computer running Services for Macintosh is located, and then select that file server. Continue to connect and select the Microsoft UAM volume as the item you want to use.

2. On the Macintosh Desktop, double-click the Microsoft UAM volume to open it.

3. Double-click **MS UAM Installer**.

### Note

- When Macintosh users connect to the computer running Services for Macintosh, authentication will be offered.

  How this is done depends on the Macintosh system the user is running. If a Macintosh computer is running System 7.1 or later and the plaintext and guest options are disabled at the server level, users will be given only one choice: Microsoft Authentication. Earlier systems will show both choices: Microsoft Authentication and plaintext password protection in the form of the Apple standard UAMs, even if the plaintext and guest logon options have been disabled and are unavailable to clients.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Services for Macintosh Authentication
Gain access to authentication files

# Manage Macintosh-accessible Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage Macintosh-accessible volumes**

- Create a Macintosh-accessible volume
- View a list of Macintosh-accessible volumes
- Modify Macintosh-accessible volumes
- Set permissions on a Macintosh-accessible volume or folder
- Set security options for Macintosh volumes
- View current users of volumes
- Disconnect Macintosh users and volumes
- Remove a Macintosh-accessible volume

# Create a Macintosh-accessible volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a Macintosh-accessible volume**

1. Open Computer Management.
2. In the console tree, right-click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. Click **New Share**.
4. In Share a Folder Wizard, click **Next**.
5. In **Folder path**, type the drive letter and path to the folder you want to make accessible to Macintosh users and click **Next**.

   Or, click **Browse** to find the folder and click **Next**.

6. In **Share name**, type the name and, optionally, in **Description**, type a description of the Windows share.
7. Select the **Apple Macintosh users** check box and click **Next**.
8. In **Permissions**, click a standard permission set or click **Customize** to set more specific permissions.
9. Click **Finish**.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- A Macintosh-accessible volume is a folder on a computer running Services for Macintosh that is made available to Macintosh clients. Once a Macintosh-accessible volume has been created for Macintosh users and a shared folder created for users of x86-based computers, both types of users can exchange files. Macintosh clients need to run AppleTalk to access AppleTalk printers.
- By default, the folder is shared out to Microsoft Windows users only. If you share out the folder only to Macintosh users, then default permissions for these users are read-only. You can, however, change their permissions using the **Security** tab.
- All Macintosh-accessible volumes must be created on an NTFS partition or on a CDFS volume. If you specify a CDFS volume, the Macintosh-accessible volume will provide read-only access.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Modify Macintosh-accessible volumes
Create a Macintosh-accessible volume
Creating a Macintosh-accessible volume on a CDFS
Working with MMC console files

# View a list of Macintosh-accessible volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view a list of Macintosh-accessible volumes**

1. Open Computer Management.

2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. Look in the **Type** column for Macintosh. This indicates the volumes on this computer that Macintosh users can use over the network.

   If the **Type** column isn't displayed, right-click **Shares**, point to **View**, and then click **Detail**.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Macintosh-accessible volumes
How shared files appear to users

# Modify Macintosh-accessible volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify Macintosh-accessible volumes**

1. Open Computer Management.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. Double-click the Macintosh-accessible volume you want to modify.
4. In the **Properties** dialog box, on the **General** and **Security** tabs, you can change the volume's properties.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can view the list of all shares on the computer running Services for Macintosh as well as modify the properties of each.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a Macintosh-accessible volume
Working with MMC console files

# Set permissions on a Macintosh-accessible volume or folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions on a Macintosh-accessible volume or folder**

1. Open Computer Management.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. Right-click the volume or folder, and then click **Properties**.
4. On the **Security** tab, under **Permissions**, select the appropriate permissions.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Macintosh-accessible volumes

# Set security options for Macintosh volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set security options for Macintosh volumes**

1. Open Computer Management.
2. In the console tree, double-click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. Right-click the volume you want to change, and then click **Properties**.
4. On the **General** tab, select one or more of the **SFM Volume Security** options.

   - **Password.** Enter the password for this volume.
   - **This volume is read-only.** Changes can not be made to files in the volume.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Set permissions on a Macintosh-accessible volume or folder

# View current users of volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view current users of volumes**

1. Open Computer Management.
2. In the console tree, click **Sessions**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Sessions

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- **Sessions** does not dynamically update the list of users. To see the latest list of users, press F5, or right-click **Sessions**, and then click **Refresh**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Macintosh-accessible volumes
File sharing
How shared files appear to users

# Disconnect Macintosh users and volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect Macintosh users and volumes**

1. Open Computer Management.
2. In the console tree, right-click **Sessions**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Sessions

3. Click **Disconnect All Sessions**.

**Caution**

- Send a message to users before disconnecting them or the volumes they are using. Otherwise, they might lose data.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Send a message to all Macintosh clients

# Remove a Macintosh-accessible volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a Macintosh-accessible volume**

1. Open Computer Management.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. Right-click the Macintosh volume you want to remove.
4. Click **Stop Sharing**.

**Notes**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Removing a Macintosh-accessible volume does not delete the corresponding folder on the computer running Services for Macintosh. Removing a volume makes it unavailable only to Macintosh users.
- If Macintosh users are currently connected, use **Send Message** to warn them. Users are likely to lose data if you remove a volume that's being used.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a Macintosh-accessible volume
Modify Macintosh-accessible volumes
Removing a Macintosh-accessible volume
Send a message to all Macintosh clients
Working with MMC console files

# Associate Windows File Types with Macintosh File Types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Associate Windows file types with Macintosh file types**

- Make new extension associations for x86-based and Macintosh files
- Add a file creator and type for Macintosh files
- Edit a description of a Macintosh file type
- Delete a file type and its extension associations

# Make new extension associations for x86-based and Macintosh files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make new extension associations for x86-based and Macintosh files**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click **Configure File Server for Macintosh**.
4. On the **File Association** tab, in **Files with MS-DOS Extension**, type an extension or select one by clicking the arrow.

   If the extension is already associated with a file type and file creator, it will be selected in the association list **With Macintosh Creator and Type**.

5. In the association list **With Macintosh Creator and Type**, click a creator and type to which you want to associate this extension.
6. To associate the extension with the selected creator and type, click **Associate**.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Extension-type associations ensure that users will see the correct icon on their computers for a file stored on the computer running Services for Macintosh. For example, Macintosh users will see a Macintosh-style icon for a Microsoft Excel file, and Windows users will see the Windows-style icon for the same Microsoft Excel file.
- For a list of extension-type associations, see "Services for Macintosh" at the Microsoft Windows Resource Kits Web site. You can create new associations or add, edit, or delete existing file creators and types.
- When you add a new extension-type association, it affects only files that are subsequently created on the server, not currently existing files. You can associate multiple extensions with a Macintosh file type and creator. However, you can associate only one file type and creator with an extension.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
**Other Resources**

Macfile
Services for Macintosh commands

# Add a file creator and type for Macintosh files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a file creator and type for Macintosh files**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click **Configure File Server for Macintosh**.
4. In the **Properties** dialog box, click the **File Association** tab, and then click **Add**.
5. Type the **Creator** and **File Type**, or select them by clicking the arrow, and then, optionally, add a description.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- The new creator and type is added to the file association list.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a file type and its extension associations
Cross-platform applications
Working with MMC console files

# Edit a description of a Macintosh file type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit a description of a Macintosh file type**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click **Configure File Server for Macintosh**.
4. On the **File Association** tab, for files with MS-DOS extensions, type an extension or select one by clicking the arrow.

   If the extension is already associated with a file type and file creator, it will be selected in the association list **With Macintosh Creator and Type**.

5. Click **Edit**, and type the new description.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a file type and its extension associations
Cross-platform applications
Working with MMC console files

# Delete a file type and its extension associations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a file type and its extension associations**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click **Configure File Server for Macintosh**.
4. On the **File Association** tab, click the creator you want to delete.
5. Click **Delete**, and then click **Yes** to confirm that you want to remove the selected file type and associated extensions.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
**Other Resources**
Macfile
Services for Macintosh commands

# Administer Macintosh Clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Macintosh clients**

- Set security options for Macintosh users
- Change the owner of a directory
- View open file forks
- Stop, start, pause, or resume services
- Send a message to all Macintosh clients
- Create a new logon message for Macintosh clients
- Set session limits for Macintosh clients
- Change the name of the computer running File Server for Macintosh

# Set security options for Macintosh users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set security options for Macintosh users**

1. Open Computer Management.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. Double-click the Macintosh-accessible volume you want to modify.
4. Click the **Security** tab, and then modify the security options.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set security options for Macintosh volumes
Working with MMC console files

# Change the owner of a directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the owner of a directory**

1. Open Windows Explorer.
2. Browse to, and then right-click, the folder you want to edit.
3. Click **Properties**, and then, on the **Security** tab, click **Advanced**.
4. On the **Owner** tab, **in Change owner to**, click the new owner.
5. Optionally, select the check box **Replace owner on subcontainers and objects** to assign this ownership to any folders or objects contained in this directory.
6. Click **Apply**.

**Note**

- To open **Windows Explorer**, click **Start**, point to **Programs**, point to **Accessories**, and then click **Windows Explorer**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# View open file forks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view open file forks**

1. Open Computer Management.
2. In the console tree, click **Open Files**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Open Files

3. In the list of open files, scan the **Type** column to find the resource and data files that are open on Macintosh clients connected to the computer running Services for Macintosh.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Macintosh file forks
Working with MMC console files

# Stop, start, pause, or resume services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop, start, pause, or resume services**

1. Open Computer Management.

2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. Double-click the service you want to change, and then click **Start**, **Stop**, **Pause**, or **Resume** as appropriate.

   To change the startup type (for example, to specify manual startup), on the **General** tab, under **Startup**, click the arrow, and then click **Automatic**, **Manual**, or **Disabled**.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Stopping, starting, pausing, and resuming services
Working with MMC console files

# Send a message to all Macintosh clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send a message to all Macintosh clients**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click **Configure File Server for Macintosh**.
4. On the **Sessions** tab, type your message text in **Send message to all Macintosh users of this system**.
5. Click **Send**, and then click **OK**.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can send messages only to all Macintosh users of the server. You can't send messages to individual users or to those users accessing a particular volume.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Disconnect Macintosh users and volumes

# Create a new logon message for Macintosh clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new logon message for Macintosh clients**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click **Configure File Server for Macintosh**.
4. On the **Configuration** tab, in **Logon message**, type the message you want Macintosh users to see when they log on to the computer running Services for Macintosh.

### Note

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a file type and its extension associations
Cross-platform applications
Working with MMC console files

# Set session limits for Macintosh clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set session limits for Macintosh clients**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click **Configure File Server for Macintosh**.
4. On the **Configuration** tab, under **Sessions**, click **Unlimited**.

   Or, click **Limited to**, and then type a number.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Setting session limits determines the number of simultaneous connections Macintosh clients can make to the computer running Services for Macintosh.
- The number of simultaneous connections is limited only by the network media. However, limiting the connections can result in performance increases by the computer running Services for Macintosh.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a file type and its extension associations
Cross-platform applications
Working with MMC console files

# Change the name of the computer running File Server for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the name of the computer running File Server for Macintosh**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Right-click **Configure File Server for Macintosh**.
4. In **Server Name for AppleTalk Workstations**, type a new name for the server, and then click **OK**.
5. In **Services and Applications**, click **Services**, and then right-click **File Server for Macintosh**.
6. Click **Restart**.

### Note

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Administrator and user issues and solutions

# Manage AppleTalk Printing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage AppleTalk printing**

- Share a printer with Macintosh users
- Capture or release an AppleTalk printing device
- Stop and restart Print Server for Macintosh
- Create a printer
- Set up a user account for Macintosh print jobs
- Prioritize print jobs

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Share a printer with Macintosh users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To share a printer with Macintosh users**

1. Open Printers and Faxes.
2. Right-click a printer, and then click **Properties**.
3. On the **Sharing** tab, click **Shared as**, and type a name.
4. On the **General** tab, in **Location**, type the place where users can find their print jobs.

   For example, you can type the room number where the shared printer is located.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- Macintosh users can gain access to any printing device on a computer running Services for Macintosh. Unlike clients running Windows XP and Windows 2000, those running Windows NT 4.0 need to run AppleTalk to access Macintosh volumes and AppleTalk printers.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Capture or release an AppleTalk printing device
Sharing printers among Macintosh and x86-based clients

# Capture or release an AppleTalk printing device

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To capture or release an AppleTalk printing device**

1. Open Printers and Faxes.
2. Right-click a printer, and then click **Properties**.
3. On the **Ports** tab, click the port, and then click **Configure Port**.
4. Either select or clear the **Capture this AppleTalk printing device** check box, and then click **OK**.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- If you want an AppleTalk printing device to be reserved for users of the computer running Services for Macintosh, it must be captured. Capturing gives spooling capabilities to all users who connect to the corresponding printer through the computer running Services for Macintosh.
- Releasing a captured printing device enables any Macintosh client to print to an AppleTalk printer directly. However, the print jobs will not be under the administrator's control.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sharing printers among Macintosh and x86-based clients
Create a printer
Setting up a user account for Macintosh print jobs

# Stop and restart Print Server for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop and restart Print Server for Macintosh**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. Right-click **Print Server For Macintosh**.
4. Click **Stop** or **Start** as appropriate.

**Notes**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You would want to stop and restart Print Server for Macintosh to remove a printing device, for example.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Stopping, starting, pausing, and resuming services
Checklist: Installing and configuring Print Server for Macintosh
Sharing printers among Macintosh and x86-based clients
Print Server for Macintosh
Manage AppleTalk Printing

# Create a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a printer**

1. Open Printers and Faxes.
2. In the **Printers** dialog box, double-click **Add Printer**.
3. Follow the Add Printer wizard instructions and select **Local Printer** (select the **Automatically detect my printer** check box if the printer is directly attached to the computer on which you are creating the printer), and then click **Next**.
4. Click **Create a new port**, click **AppleTalk Printing Devices**, and then click **Next**.
5. Expand the AppleTalk zone where the printer you want to connect to is located, and then select the printer.
6. Complete the Add Printer wizard.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- The printer name can contain up to 32 characters. This name will appear in the title bar of the printer window. By default, it is the name that network users (except MS-DOS users) will see when they share the printer.
- Select the **Share this printer** option during setup. In **Share Name**, specify the printer name that you want MS-DOS clients to see.
- When you are selecting a destination, if the printing device is physically connected to the computer running Services for Macintosh, select the appropriate port.

  If the printing device is on a network, click **Add Port**. In the **Printer Ports** dialog box, click **AppleTalk Printing Devices**, and then click **OK**. In the **Available AppleTalk Printing Devices** dialog box, select a zone and a printer, and then click **OK**.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Creating a printer
Printing device vs. printer

# Set up a user account for Macintosh print jobs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up a user account for Macintosh print jobs**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. Double-click **Print Server For Macintosh**.
4. On the **Log On** tab, click **This Account**, and then type the name of the user account, or click **Browse**.
5. To require a password for Macintosh users of the computer running Services for Macintosh, in **Password** type a password, and then confirm it.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Setting up a user account for Macintosh print jobs
Working with MMC console files

# Prioritize print jobs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To prioritize print jobs**

1. If necessary, create two (or more) printers and share them.
2. Open Printers and Faxes.
3. Right-click a printer, and then click **Properties**.
4. On the **Advanced** tab, in **Priority**, select a priority.

   Higher priority jobs will print before lower priority jobs. For example, a priority-1 job will print before a priority-2 job.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Prioritizing print jobs
Create a printer
Share a printer with Macintosh users

# Manage Remote Access for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage remote access for Macintosh**

- Configure remote access protocols for Macintosh clients
- Configure user authentication for Macintosh remote access
- Configure plaintext password authentication for Macintosh remote access
- View AppleTalk and MacFile events

# Configure remote access protocols for Macintosh clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure remote access protocols for Macintosh clients**

1. Open Routing and Remote Access.
2. Right-click the server name, click **Configure and Enable Routing and Remote Access**, and then follow directions in the Routing and Remote Access Server Setup Wizard.
3. Click **Remote access server**.
4. After the wizard finishes, open Routing and Remote Access again, and then right-click the server name.
5. Click **Properties**, and then click the **PPP** tab to configure PPP settings.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- If you need to add the server, right-click **Server Status**, then click **Add Server**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the AppleTalk protocol and routing
Configure AppleTalk protocol
Remote Access

# Configure user authentication for Macintosh remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure user authentication for Macintosh remote access**

1. Open Routing and Remote Access.
2. Right-click your server, and then click **Properties**.
3. On the **Security** tab, configure remote access options as appropriate for the computer, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group. As a security best practice, consider using the Run As command rather than logging on with administrative credentials. If you have logged on with administrative credentials, you can also open Routing and Remote Access by clicking **Start**, clicking **Control Panel**, double-clicking **Administrative Tools**, and then double-clicking **Routing and Remote Access**. For more information, see Default local groups, Default groups, and Using Run as.
- For Macintosh clients logging in as guest, ensure that the **Unauthenticated access** check box on the Security tab is selected.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the AppleTalk protocol and routing
Configure AppleTalk protocol
Configure plaintext password authentication for Macintosh remote access
Remote Access

# Configure plaintext password authentication for Macintosh remote access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure plaintext password authentication for Macintosh remote access**

1. Open MMC.
2. In the **File** menu, click **Add/Remove Snap-in**.
3. Click **Add**, click **Group Policy Object Editor**, and then click **Add**.
4. In the Select Group Policy Object wizard, click **Finish**, and then click **Close**.
5. In the **Add/Remove Snap-in** dialog box, click **OK**.
6. In the **Console Root**, double-click **Local Computer Policy**, and then double-click **Computer Configuration**.
7. Double-click **Windows Settings**, and then double-click **Security Settings**.
8. Double-click **Account Policies**, and then click **Password Policy**.
9. Right-click **Store password using reversible encryption for all users in the domain**, and then click **Properties**.
10. On the **Security** tab, click **Enabled**. Click **Apply** and then **OK**.

### Note

- To open MMC, click **Start**, click **Run**, and then type **mmc**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Install the AppleTalk protocol and routing
Configure AppleTalk protocol
Remote Access

# View AppleTalk and MacFile events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view AppleTalk and MacFile events**

1. Open Event Viewer.
2. Click **System Log**.
3. Review the System Log and, in the **Source** list, double-click the MacFile, MacSrv, or AppleTalk events you want to view.

## Notes

- To open Event Viewer, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Event Viewer**.
- To see specific events, on the **View** menu, click **Filter**, and then select the filter options you want to apply. (To filter events in Event Viewer, the Computer Browser service must be started.)

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Viewing the event log
View Event Logs

# Services for Macintosh Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

You can use Services for Macintosh to share files and printers between x86-based and Apple Macintosh clients.

After the AppleTalk protocol is set up, a computer running Services for Macintosh can also function as an AppleTalk router. Routing capability is supported for AppleTalk Phase 2.

With Services for Macintosh, Macintosh computers need only the Macintosh operating system software to function as clients; no additional software is required. You can, however, set up and distribute the optional Microsoft® user authentication module (MSUAM), which lets Macintosh users securely log on to the computer running Services for Macintosh using the same logon method as Microsoft Windows® users. For information about MSUAM, see Services for Macintosh Authentication.

Services for Macintosh simplifies administration by maintaining just one set of user accounts instead of separate user accounts (one on the Macintosh server and another on the computer running Services for Macintosh).

This section provides general background information about File Server for Macintosh and Print Server for Macintosh:

- Understanding Services for Macintosh
- Using AppleTalk Networking

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Understanding Services for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Services for Macintosh**

Services for Macintosh is made up of three parts:

- **File Server for Macintosh** (also called MacFile) which lets you designate a folder as a Macintosh-accessible volume, ensures Macintosh file names are legal NTFS names, and handles permissions.
- **Print Server for Macintosh** (also called MacPrint) which lets all network users send print jobs to a spooler on the computer running Services for Macintosh and continue working; they don't have to wait for their print jobs to finish.
- **AppleTalk protocol** is the layer of AppleTalk Phase 2 protocols that delivers data to its destination on the network.

This section contains information on:

- AppleTalk Networking and Routing
- Files
- Printing Benefits
- Security for Services for Macintosh users

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# AppleTalk Networking and Routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AppleTalk networking and routing**
A computer running Services for Macintosh can provide routing and seed routing support. An AppleTalk router broadcasts routing information, such as network addresses, and keeps track of and directs data packets on AppleTalk networks. Seed routers perform these functions and also seed the physical networks on which they reside. Seeding a network means establishing and initializing the network address information for that network.

You can install an unlimited number of network adapters on a computer running Services for Macintosh in order to add to an AppleTalk internet.

For more information on AppleTalk networking and routing, see:

- AppleTalk network media
- AppleTalk networks
- Phase 2 AppleTalk networks
- AppleTalk routing information
- Working with AppleTalk seed routers
- Making a router plan
- Assigning AppleTalk network numbers and ranges
- Assigning AppleTalk zones
- Creating an AppleTalk router record
- AppleTalk over PPP

# AppleTalk network media

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AppleTalk network media**

As you plan how to physically connect your Macintosh and x86-based computers, the first thing to consider is network media. Each network media type has its own method of cabling and network topology, and each requires different network hardware.

You can choose among three types of media:

- Ethernet (Fast, Gigabit)
- Token ring
- LocalTalk

Ethernet is the most common network media in x86-based networking.

In token ring, all computers are connected in a ring or star topology where a token passing scheme is used to send messages. FDDI, although not as common, is based on token ring and is designed to be used with fiber-optic cabling.

LocalTalk, an early media, was used in AppleTalk networking.

# AppleTalk networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AppleTalk networks**

Because AppleTalk networks differ from x86-based networks, you must consider some special concepts and issues when you set up an AppleTalk network.

One important concept is the internet. (Note that this is a different concept than the Transmission Control Protocol/Internet Protocol [TCP/IP] Internet.) Most large AppleTalk networks are not single physical networks in which all computers are attached to the same network cabling system. Instead, they are internets, which are multiple smaller physical networks connected by *routers*.

Routers maintain a map of the physical networks on the internet and forward data received from one physical network to other physical networks. Routers are necessary so that computers on different physical networks can communicate with one another. They also reduce network traffic on the internet by isolating the physical networks. In other words, routers only send data that is usable by a network.

Some routers on the network are seed routers. A seed router initializes and broadcasts routing information about one or more physical networks. This information tells routers where to send each packet of data. Each physical network must have one or more seed routers that broadcast the routing information for that network.

Not all routers are seed routers. Routers that are not seed routers maintain a map of the physical networks on the internet and forward data to the correct physical network. Seed routers perform these functions too, but they also initialize the routing information, such as network numbers and zone lists, for one or more physical networks.

A computer running Services for Macintosh can function as a seed router or as a nonseed router. If it is a seed router, it must be the first server you start so that it can initialize the other routers and nodes with network information. If it is a nonseed router, it cannot be started until a seed router has initialized all ports. You can also use dedicated hardware routers on your network.

# Phase 2 AppleTalk networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Phase 2 AppleTalk networks**
There are two types of AppleTalk networks: Phase 1 and Phase 2.

AppleTalk Phase 1 was the original AppleTalk protocol architecture designed to support networking for small workgroups. Phase 1 could only support a single physical network that had just one network number and one zone.

AppleTalk Phase 2 enhances the routing and naming services of AppleTalk. This means improved network traffic and better router selection. You can now create AppleTalk networks that support more than 254 nodes and have multiple zones.

You must use Phase 2 to run Services for Macintosh.

Phase 2 includes these features:

- enhancements to routing and naming services, such as the ability to support AppleTalk networks with more than 254 nodes
- support for multiple zones per network
- improved network traffic and better router selection
- support for token ring, LocalTalk, Ethernet, and FDDI

LocalTalk networks have a single network number; EtherTalk and TokenTalk networks can be assigned a network range, allowing for more nodes on the network.

**AppleTalk zones**
Each LocalTalk network must be in a single zone; each EtherTalk and TokenTalk network can have multiple zones, and individual nodes on a network can be configured to be in any one of the network's associated zones.

**Number of nodes per network**
A *node* is any type of device on the network. Each client, printer, server, and router is a node on an AppleTalk network.

LocalTalk networks can have as many as 254 nodes (but are actually limited to 32 or fewer nodes because of media capacity); EtherTalk and TokenTalk networks can have as many as 253 nodes for every number in the network range, for a maximum of 16.5 million nodes.

### Caution

- Although EtherTalk and TokenTalk support 16.5 million nodes, do not specify this many nodes; network media cannot physically accommodate this many nodes.

### Notes

- Help and Support uses the terms *Ethernet* and *token ring* in descriptions of network media. For discussions on the implementation of an AppleTalk network on Ethernet or token ring, EtherTalk and TokenTalk are used. For more information, refer to an AppleTalk manual.
- AppleTalk Phase 2 is not available on Macintosh platforms released before Macintosh Plus. It requires the installation of AppleTalk file V53. AppleTalk file V53 is installed automatically when you install either EtherTalk 2.0 or TokenTalk 2.0.

# AppleTalk routing information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AppleTalk routing information**

AppleTalk routing information includes:

- A network number or network range associated with each physical network.
- The zone name or zone list associated with each physical network.
- The default zone for the network (if the network has multiple zones).

The *network number* or *network range* is the address or range of addresses assigned to the network. A network number is unique and identifies a particular AppleTalk physical network. By keeping track of network numbers and network ranges, routers can send incoming data to the correct physical network. A network number can be any number from 1 through 65,535.

LocalTalk networks can have only a single network number; EtherTalk, TokenTalk, and Fiber Distributed Data Interface (FDDI) networks can have network ranges.

A zone is a logical grouping that simplifies browsing the network for resources, such as servers and printers. It is similar to a Windows Server 2003 domain as far as browsing is concerned. In LocalTalk networks, each physical network can be associated with only one zone. However, for EtherTalk, TokenTalk, or FDDI, you have more flexibility in assigning zones. Each EtherTalk, TokenTalk, or FDDI network can have one or more zones associated with it, and each zone can include servers and printers on one or more physical networks. This allows you to group servers and printers logically into zones so that users can easily locate and access the servers and printers, no matter what physical networks they are on.

Each Macintosh client on the network is assigned to a single zone. However, each client can access servers and printers in any zone on the network. Zones make accessing network resources simpler for users. When users use the Chooser to view the network, they see only the resources in a single zone at a time, preventing them from having to navigate through huge numbers of resources on large networks to find the resources that they need. You can put the clients, servers, and printers used by a single group into a single zone so that users will see only the resources they typically use but will still be able to access resources in other zones when required.

A zone list includes all the zones associated with that network. One of these zones is the network's default zone, to which the Macintosh clients on that network are assigned by default. Users can configure the client to be in a different zone, however.

# Working with AppleTalk seed routers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Working with AppleTalk seed routers**

When you install and set up Services for Macintosh, you must specify whether the computer running Services for Macintosh will seed each physical network to which it is attached. For example, a computer running Services for Macintosh attached to three physical AppleTalk networks might serve as a seed router on two of the networks but not on the third.

For networks that the server will seed, specify the routing information. The computer running Services for Macintosh will then function as a seed router, seeding the routing information that you provided. If you specify that a server will not seed a network (that is, if you make it a nonseed router), the port will be seeded by another AppleTalk router attached to it.

**Using multiple seed routers on a network**

To make your network more reliable in case of system failure and power outages, you can install multiple seed routers on the same physical network.

When you install multiple seed routers for a particular network, all the seed routers must seed the same information for that network. When the network starts, the first seed router that starts on the network becomes the actual seed router. Although only one server seeds zone and range information at a time, all other seed routers continue to act as routers.

If the first seed router to start on a network has different routing information than the other seed routers that start later, the information established by the first seed router is used. If a seed router that started subsequently has different information and is running Services for Macintosh, the conflicting information is ignored and an event is written to Event Viewer. If a server that is seed routing stops responding or is shut down, one of the waiting seed routers starts seeding the network zone and range information.

# Making a router plan

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Making a router plan**

After you have diagrammed your network, make a router plan that shows the location and type of each seed router for the network.

To make a router plan, determine:

- The expected number of AppleTalk devices on present and projected Ethernet networks.
- The quantity of network numbers sufficient to satisfy capacity requirements. (Up to *n* * 253 devices can be supported, where *n* is the number of network numbers in the range.)

**Router seeding plan example**

| Cable ID | Network range | Zone list | Seed devices |
|---|---|---|---|
| Ethernet #1 | 16-25 | **Finance** Accounts Payable Accounts Receivable | Server 1 Server 3 |
| Ethernet #2 | 32-37 | **Marketing** Marketing Exec | Server 5 |
| Ethernet #3 | 768 | **Engineering** | Server 6 |
| LocalTalk #1 | 1024 | Finance | Server 3 |
| LocalTalk #2 | 1280 | Marketing | Server 5 |

**Notes**

- In the **Zone list** column, bold formatting indicates the network's default zone.
- Seed device information can indicate the server name, the type of dedicated hardware router, or the location, depending on your needs.

# Assigning AppleTalk network numbers and ranges

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Assigning AppleTalk network numbers and ranges**

Follow these guidelines when you decide how to assign network numbers and network ranges:

- Use network numbers that leave room for expansion.

  For example, if your first AppleTalk zone has a network range of 1024-1099, then make the range of your second AppleTalk zone 1400-1450. This will leave plenty of room for growth of the first AppleTalk network.

- Network numbers are essentially arbitrary. The important thing is for them to be unique and not to overlap (if in a range) with other ranges.

- Base your network ranges on the number of nodes you expect to have in the future on each network.

  Base the extent of a network range on the number of AppleTalk nodes expected on the physical network. The total number of possible AppleTalk nodes is 253 times the number of network numbers in the range.

  For example, a network range of 101 through 103 permits 759 nodes (3 * 253 = 759); a network range of 120 to 129 permits 2530 nodes (10 * 253 = 2530) on a network. Leave room for more nodes than are currently connected.

For more information, see AppleTalk routing information, Configure AppleTalk protocol, Assigning AppleTalk zones, and Working with AppleTalk seed routers.

# Assigning AppleTalk zones

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Assigning AppleTalk zones**

AppleTalk zones are identified by zone names. Follow these guidelines when you decide how to assign zone names:

- You can assign one or more zone names to each Ethernet network. An asterisk (*) cannot be a zone name.
- For each Ethernet network, decide which zone will be the default zone.
- The number of zones your internet has depends on the size of the internet you are planning. If your internet is small, a single zone can be adequate. But if you have a single Phase 2 Ethernet network that spans a large geographic area or contains large numbers of AppleTalk devices (such as printers and servers), use multiple zones to make it easier for users to find the devices they need.

# Creating an AppleTalk router record

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating an AppleTalk router record**

To keep information about your internet for maintenance purposes, create a router record. Include the following information:

- Router location.
- Physical location.
- If the router is a computer running Services for Macintosh, record the server computer name.
- Router type and version.
- The physical networks connected to the router, with the following information for each:
    - Cabling identification.
    - Network media type.
    - Network numbers.
    - Zone names.
    - Default zone.
    - Whether this router is a seed router for the networks attached to it.

# AppleTalk over PPP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**AppleTalk over PPP**

Macintosh users can dial in to a computer running Services for Macintosh by using a PPP client that supports AppleTalk Control Protocol (ATCP). ATCP is installed automatically if remote access and AppleTalk Protocol are installed. The AppleTalk protocol is installed automatically with File Server for Macintosh and Print Server for Macintosh. You can also install AppleTalk protocol separately. The version of ATCP that comes with Services for Macintosh includes the following features:

- Addresses are dynamically allocated.
- Clients are not forced to change their passwords.

ATCP client callback functionality is identical to other dial-up client types. Callback ensures that only users from specific locations can access the dial-up server. This saves toll charges for the user. In addition, ATCP clients can specify their number to be called back at the time they dial in, if their account is callback-enabled. To do this, users specify their user name as *username@phonenumber*. For example, a user named Jeff Smith, whose account is configured for callback, wants to be called back at 5550123. He can enter JeffSmith@5550123 as his user name, and he will be called back at that number.

Guest account authentication and access authorization are the same for ATCP as for any other dial-up method.

ATCP requires that user passwords be saved in reversibly encrypted plaintext format on the dial-up server. (By default, passwords are not stored on the dial-up server in plaintext format.) Because of this requirement, existing users who want to use the ATCP protocol must have their passwords converted to plaintext format. First the administrator must enable the dial-up server to store passwords in plaintext format. For existing accounts, the administrator can then either delete and re-create the accounts or change the password for the accounts. The passwords are then stored in plaintext format, and users can dial in using the ATCP protocol.

All account passwords created after enabling plaintext password storage will be saved in plaintext format.

You can store passwords on a per-user or domain-wide basis. So for ATCP clients, and for Apple Encrypted passwords to work, the administrator must configure encrypted plaintext password storage for each user.

**Note**

- Do not statically configure client addresses, because addresses are dynamically allocated.

# Files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Files**

This section contains information on:

- Macintosh-accessible volumes
- File sharing
- How shared files appear to users
- Cross-platform applications
- Macintosh file forks
- How file names are translated
- Macintosh file administration

# Macintosh-accessible volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Macintosh-accessible volumes**

A computer running Services for Macintosh can store files so that users of both x86-based and Macintosh computers can gain access to them. Users of x86-based computers look for shared files in a shared folder on the computer running Services for Macintosh. Macintosh users look for shared files in the same folder; however, they see the folder as a volume, with familiar folders and files.

A Macintosh user shares a file with users of x86-based computers by storing that file in a Macintosh-accessible volume on the computer running Services for Macintosh. Likewise, a Macintosh user can mount a Macintosh-accessible volume on the desktop to use files stored in shared folders by users of x86-based computers.

All Macintosh-accessible volumes must be created on an NTFS partition or on a Compact Disc File System (CDFS) volume. If you specify a CDFS volume, the Macintosh-accessible volume will provide read-only access.

# File sharing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File sharing**

With Services for Macintosh, both users of x86-based computers and Macintosh users can easily share files stored on a server. On a computer running File Server for Macintosh, files are stored in shared folders or in Macintosh volumes. For a file to be accessible to x86-based clients, it must be in a shared folder (or in a subfolder of a shared folder). Each server can have one or more shared folders. Each shared folder on a server is assigned a unique name.

With File Server for Macintosh, Macintosh users cannot automatically gain access to all shared folders. To make a folder (and consequently its subfolders, which may or may not be shared for users who are not using Macintosh computers) available to Macintosh users, the administrator must designate the folder as a *Macintosh-accessible volume*.

### Notes

- Some users of x86-based computers are familiar with the terms *volume* and *volume labels* as they relate to a hard disk partition. In the context of Services for Macintosh, a volume is either a folder designated as both a shared folder and a Macintosh-accessible volume (meaning that both types of clients can gain access to the files in the volume) or a folder available only to Macintosh users on the network. The only exception to this convention is CD-ROM or Compact Disc File System (CDFS) volumes.

- Macintosh volume names that are created on a computer running a member of the Windows Server 2003 family cannot exceed 12 characters. To create volumes with longer names, use the command-line utility Macfile.exe. For example, to add a volume called Landscape Design on the magnolia server using the Trees folder on drive E:, type:

  macfile volume /add /server:\\magnolia /name:"Landscape Design" /path:e:\trees

# How shared files appear to users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How shared files appear to users**

Shared files appear as expected to users of x86-based computers and to Macintosh computer users. For example, when an MS-DOS user views these shared files, the file names follow the MS-DOS standard naming convention, whether or not they were created that way by a Macintosh user. When a Macintosh user views the files, they appear as Macintosh files on the Macintosh client itself or on Macintosh servers running AppleShare.

Within a folder that is both a shared folder and a Macintosh-accessible volume, users of networked x86-based computers see folders and files. In fact, this is what is actually stored on the server hard disk. To Macintosh users, the volume appears to contain Macintosh files and folders. When Macintosh users browse through the files available on the server, they see icons that represent each file and folder.

Macintosh files and folders can have Macintosh file names, including long names and names containing spaces and other characters. They are not limited to the 8.3 naming convention of the file allocation table (FAT) file system used with the MS-DOS system and some OS/2 computers. File Server for Macintosh and the NTFS file system translate the names so that users can see them.

# Cross-platform applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cross-platform applications**

Many applications have separate (cross-platform) versions for Macintosh and x86-based computers. With Services for Macintosh, Macintosh and x86-based clients can use cross-platform applications to modify the same files.

For example, a person who uses an x86-based version of Microsoft Excel can create a spreadsheet file and store the file on the server in a shared folder that is also configured as a Macintosh-accessible volume. A Macintosh user who opens that shared folder can double-click the file icon, and Microsoft Excel for Macintosh starts and opens the file. The Macintosh user can modify the file and save it. When the user of the x86-based computer opens the file, the modified version of the file appears.

File Server for Macintosh uses extension-type associations to display x86-based files with the correct icon when the Macintosh user is viewing the files in the Finder.

File Server for Macintosh comes with extension-type associations already defined for many applications. These can be extended or modified.

# Macintosh file forks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Macintosh file forks**

Each Macintosh file has two parts, or *forks*: a data fork and a resource fork. The data fork contains the actual data of the file. The resource fork contains Macintosh operating system information about the file, such as code, menu, font, and icon definitions.

When a Macintosh file is shared on a computer running Services for Macintosh, the two forks are saved in a single NTFS file.

# How file names are translated

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How file names are translated**

There are two things to be aware of regarding file-name translation when running Services for Macintosh:

- How Macintosh file names are maintained and presented to various users.
- How the longer NTFS names (more than 32 characters) are presented to Macintosh users.

When a Macintosh user creates a file or a folder on the server and gives it a name, File Server for Macintosh checks it for illegal NTFS characters. If the file name contains illegal NTFS characters, File Server for Macintosh replaces the illegal characters. Otherwise, the original Macintosh name is the same as the NTFS name. Macintosh users see the name as it was created. Windows users see the same name with any illegal characters replaced.

After File Server for Macintosh has replaced illegal NTFS characters, names that are too long for MS-DOS users are shortened to six characters, a tilde (~), and a unique number. Extensions are preserved.

Users creating long NTFS file names (up to 256 characters) should name files with 31 characters (the Macintosh limit) or fewer so that Macintosh users can readily decipher the file names.

Summary:

- A file created using the FAT file naming convention appears as created to NTFS users and Macintosh users.
- A file created using the 31-character limit of the Macintosh system appears as created (except with illegal characters replaced) to NTFS users. MS-DOS users see a short name.
- A file created using the NTFS 256-character limit appears as created to Macintosh users if it has 31 or fewer characters. Otherwise, it appears to both MS-DOS and Macintosh users in the shortened form.

**Note**

- Because MS-DOS users refer to files created by Macintosh users by the translated short names, Macintosh users should give the FAT standard names (eight characters plus an optional period and three-character extension) to files and folders that will also be used by MS-DOS users. This way MS-DOS users won't have to decipher short names. Macintosh users can use long file and folder names if no MS-DOS users need to access those files.

# Macintosh file administration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Macintosh file administration**

In Windows NT 4.0 and earlier, the Macfile program handled Macintosh file administration, including the creation of Macintosh volumes, passwords, security options, user limits, and permissions. You accessed the **Macfile** menu from Control Panel, File Manager, and Server Manager.

On computers that are running Services for Macintosh, Macintosh volumes and files are centrally administered through the Microsoft Management Console (MMC) Shared Folders snap-in. Both Macintosh and non-Macintosh files are administered by using Shared Folders. This improves the integration of Macintosh file management with non-Macintosh files.

**Note**

- You can still use the command-line version of Macfile (Macfile.exe) to administer servers, volumes, files, and folders.

# Printing Benefits

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing benefits**

Users gain three major printing benefits with Print Server for Macintosh:

- Macintosh users can print PostScript jobs to non-PostScript printers directly connected to the computer running Services for Macintosh. To the Macintosh user, these printers appear like the standard LaserWriter.
- Users of x86-based computers can send print jobs to PostScript printers on an AppleTalk network, and they can check print jobs from their clients.
- Macintosh and x86-based print jobs are spooled before they go to the printer. So both Macintosh users and users of x86-based computers can send jobs to the printer and then continue working at their computers. This means that users do not have to wait for their jobs to print before using their computers to do other tasks or wait for a printer to be available.

For more information on printing, see:

- Print Server for Macintosh
- Printing device vs. printer
- Sharing printers among Macintosh and x86-based clients
- Capturing AppleTalk printers
- Avoiding LaserPrep wars

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Print Server for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Print Server for Macintosh**
When Services for Macintosh is set up, several AppleTalk services (such as Print Server for Macintosh) become integrated with the operating system. Print Server for Macintosh makes printers connected to the computer running Services for Macintosh available to Macintosh clients, and it makes AppleTalk PostScript printers (with LaserWriter drivers) available to x86-based clients.

When Print Server for Macintosh receives print jobs, it sends them to a spooler, which is a portion of the print server hard disk. The spooler then sends the print job to the specified printing device; for example, to a printing device on the AppleTalk network. This enables Macintosh users, as well as users of x86-based computers, to submit print jobs and continue working on their computers without waiting for the print job to finish.

Print Server for Macintosh also translates all incoming PostScript files if the print request is to a non-PostScript printer attached to the computer running Services for Macintosh. A Macintosh client (but not a Windows-based client) can send a PostScript job to any printer that is connected to a computer that is running Services for Macintosh.

Whether printing devices are attached to the computer running Services for Macintosh or are located elsewhere on the AppleTalk network, the Printers folder displays a list of print jobs for the respective printers you created to represent the devices. By default, each list presents jobs in first-in, first-out (FIFO) order. You can change the priority of jobs, however, and specify permissions for the printer and times for print jobs to run.

### Notes

- This implementation of PostScript Routing Information Protocol (RIP) supports 300 dpi and Postscript level 1.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Printing device vs. printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing device vs. printer**
Before setting up printers, it's important to understand the distinction between a printing device and a printer that you create using the Add Printer wizard:

- A printing device is the hardware that actually does the printing.

- A printer you create using the Add Printer wizard is a software interface between the document and the printing device. You use the printer you create to send jobs to the printing device, according to the specified priority (for example, on a first-come, first-served basis).

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Sharing printers among Macintosh and x86-based clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sharing printers among Macintosh and x86-based clients**
Without Print Server for Macintosh, Macintosh clients typically send print jobs to AppleTalk printing devices (usually PostScript devices), and x86-based clients send print jobs to non-PostScript devices. Print Server for Macintosh lets both types of clients send print jobs to either type of printing device. In addition, both types of users get the benefit of spooling when they print through the computer running Services for Macintosh.

Users of x86-based computers specify printers on a computer running Services for Macintosh and send print jobs to them as usual, whether the printing device is attached to the server itself or located elsewhere on the network. Similarly, Macintosh users have the familiar Chooser interface to use for connecting to printers that are set up for both AppleTalk printing devices and those attached to a computer running Services for Macintosh.

To get these printing benefits, set up Print Server for Macintosh, and create a user account (such as MACUSERS).

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Capturing AppleTalk printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capturing AppleTalk printers**

When you set up a printer on an AppleTalk network to be used with Services for Macintosh, you can specify whether Print Server for Macintosh will *capture* the printer, that is, prevent the printer from accepting print jobs from any source other than Print Server for Macintosh. Capturing gives you complete control over the printer.

If a printer will be used exclusively by users of x86-based computers, you should capture it. Doing so ensures that users do not accidentally bypass the print server and send print jobs directly to the printer or reset the printer, which might cause spooler problems. You will also avoid LaserPrep Wars.

Note that if a source other than Print Server for Macintosh prints jobs on the printer, you should not capture the printer. For example, don't capture printers if you use Apple LaserShare (which provides spooled printing for Macintosh clients) or if you use a minicomputer that sends print jobs from minicomputer users to the printer.

If a printer is not captured and multiple sources send jobs to that printer, no jobs will be interrupted; however, while the printer is printing a job from one source, it will appear busy to the other sources.

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Avoiding LaserPrep wars

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Avoiding LaserPrep wars**

A condition known as *LaserPrep wars* causes slow printing performance with some AppleTalk networks. Services for Macintosh solves this problem.

LaserPrep wars occur when a network has Macintosh clients that use two or more versions of Chooser Packs, which include a PostScript preparation file (also called a LaserPrep file) and a PostScript driver. A printer can use only one version of the LaserPrep file at a time. When a Macintosh user sends a print job to the printer, the Macintosh computer checks for the printer's version of the LaserPrep file. If the printer currently has a different version than the Macintosh client uses, the Macintosh client sends its version of the LaserPrep file along with the print job and instructs the printer to load that file as the printer's resident LaserPrep file.

Because Macintosh computers with different LaserPrep file versions send print jobs to a printer, different versions of the LaserPrep file are loaded and unloaded on the printer. The printer must load and unload versions of the LaserPrep file and then print a startup page each time a different LaserPrep file becomes resident. This causes performance problems and can reduce the life cycle of the printer.

For example, suppose a Macintosh user whose client uses Chooser Pack version 6.0 sends a document to the printer. The LaserPrep version 6.0 file is made resident on the printer. Then, if the next document sent to the printer comes from a client using Chooser Pack version 7.0, the printer must reset, load LaserPrep 7.0, and print a new startup page before printing the document.

Print Server for Macintosh solves the LaserPrep wars problem by sending the LaserPrep file with each job. This extra effort actually improves overall performance: The printer never has to spend time making a LaserPrep resident or printing a startup page.

Note that for printers on an AppleTalk network, LaserPrep wars are guaranteed to be avoided only if the printer is captured. If the printer is not captured, users who send print jobs directly to the printer, bypassing Print Server for Macintosh, can initiate LaserPrep wars.

LaserPrep wars are always prevented when printers are attached directly to a computer that is set up with Print Server for Macintosh.

### Note

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Security for Services for Macintosh users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security for Services for Macintosh users**
With Services for Macintosh, network security is enforced for Macintosh clients in the same way it is enforced for x86-based clients. File Server for Macintosh translates user identification, authentication (passwords), and permissions so that the integrity of the server is maintained regardless of the type of client used.

For more information on network security, see:

- Accounts for Macintosh clients
- Services for Macintosh Passwords
- Services for Macintosh Authentication
- Services for Macintosh Permissions

# Accounts for Macintosh clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accounts for Macintosh clients**

Services for Macintosh uses the same user accounts database as any member of the Windows Server 2003 family. If you already have Windows Server 2003 family accounts created for the people who will be using Macintosh computers on the network, you don't need to create additional accounts. Create accounts only for users who don't already have accounts on the computer running File Server for Macintosh and Print Server for Macintosh.

One aspect of the Windows Server 2003 family user accounts, the user's *primary group*, applies only to Services for Macintosh. The user's primary group is the group the user works with most, and it should be the group with which the user has the most resource needs in common. When a user creates a folder on a server, the user becomes the owner. The owner's primary group is set as the group associated with the folder. The administrator or owner can change the group associated with the folder.

# Services for Macintosh Passwords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Passwords**
A Macintosh user logs on to a computer running Services for Macintosh in one of three ways:

- As a guest.
- As a user with a plaintext password.
- As a user with an encrypted password.

**Guest logons**
To allow Macintosh clients to connect to a server running Services for Macintosh without providing user credentials, you need to enable guest logons. To do so, you need to enable the guest account in the Computer Management snap-in. The guest account is disabled by default in the Windows Server 2003 family. For more information on enabling guest accounts, see Default local groups, Local user accounts, and Create a local user account.

Once the guest account has been enabled, you can set guest permissions on a volume. To do this, use the options available in the **Permissions for Guests** box in the **Security** tab of the **Microsoft UAM Volume Properties** dialog box.

For more information on how to set guest permissions on Macintosh volumes, see Set permissions on a Macintosh-accessible volume or folder. To understand how permissions work, see Access Control.

**Plaintext passwords**
Plaintext password protection is part of the AppleShare client software on Macintosh computers. It provides less security than encrypted password protection because the passwords are sent over computer networks and can be detected by *sniffers* (network monitors that can look for passwords). Moreover, the AppleShare passwords can contain no more than eight characters, making them relatively easy to guess with conventional methods like dictionary attacks. Plaintext password protection is offered for Macintosh users who use the standard AppleShare client software or System 7 File Sharing.

**Encrypted passwords**
An encrypted password is more secure than a plaintext password. The Windows Server 2003 family encrypts passwords and stores them so that they cannot be directly stolen from the client itself. Encrypted passwords can contain up to 14 characters. File Server for Macintosh offers encrypted passwords to Macintosh clients.

# Services for Macintosh Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication**

With Services for Macintosh, Macintosh computers need only the Macintosh operating system to function as clients; no additional software is required. You can, however, set up and distribute the Microsoft user authentication module (MSUAM), which lets Macintosh clients securely log on to computers running Services for Macintosh using the same method as Windows clients.

The Apple standard UAM provides only minimal encryption and so user passwords can be easily intercepted on the local area network (LAN) or Internet. MSUAM is a stronger password encryption method used to log on to a computer running Services for Macintosh. MSUAM encrypts passwords and stores them on the computer running Services for Macintosh. You can either set up, or instruct Macintosh users to set up, the authentication file for MSUAM on their Macintosh computers over the network. Although MSUAM is not required, it provides secure encrypted authentication to servers running Services for Macintosh.

If MSUAM is set up on a client computer, the **Require strong authentication (NTLMv2)** check box is selected by default. Selecting this check box requires users to authenticate to servers that implement NTLMv2. This setting prevents users from authenticating to servers that are running Windows NT 4.0 without Service Pack 3 and other operating systems that cannot authenticate using NTLMv2. However, users can clear the **Require strong authentication (NTLMv2)** check box to allow themselves to authenticate to servers that do not implement NTLMv2. Using NTLMv2 helps eliminate man-in-the-middle attacks in which an attacker tries to force authentication using the less secure Lan Manager (LM) Authentication protocol.

The default setting for NTLMv2 authentication is restored every time users log off. If users clear the **Require strong authentication (NTLMv2)** check box, a message box appears, warning them of the security risk of doing so. However, users can still clear the check box and authenticate using a weaker protocol, putting your network at risk.

With MSUAM, users can also specify a domain when they log on or change their passwords. This feature ensures that, if they have accounts in multiple domains, the correct one will be used. (To specify a domain, users type *DomainName******UserName*in **Name**.)

Different versions of the Mac OS use different default authentication methods. Ensure that you check the appropriate Mac OS documentation before you configure authentication methods for Macintosh users. For information on authentication and remote access security in the Windows Server 2003 family, see Remote Access Security.

# Services for Macintosh Permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Permissions**

Access to network files and folders is controlled with permissions. Using the Windows Server 2003 family security system, you specify which users can access shares, folders, and files, and how these resources can be used. The Macintosh-style permissions differ in that they can be set for folders only, not files. However, NTFS permissions still apply to both files and folders that are shared as Macintosh accessible volumes.

Services for Macintosh ensures a consistent file-level security for x86-based and Macintosh clients by translating file permissions, which adds a level of security to your network. A Macintosh user sets permissions according to the Macintosh scheme; File Server for Macintosh translates these to the Windows Server 2003 family permissions. The reverse is also true: The Windows Server 2003 family permissions set by x86-based clients are translated to Macintosh-style permissions for Macintosh users.

Both the Administrators and Server Operators groups can administer File Server for Macintosh and Print Server for Macintosh. The Windows Server 2003 family Administrator account always has full permissions on volumes that are running File Server for Macintosh. Server Operators group membership is sufficient for common administrative tasks. You need not log on with administrative credentials to perform these tasks.

For more information on permissions, see:

- Types of permissions
- How file-level permissions are handled
- Translating permissions between x86-based clients and Macintosh clients
- Setting permissions from a Macintosh or x86-based computer
- Volume passwords

# Types of permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Types of permissions**

Users and administrators of x86-based networks use the Windows Server 2003 family permissions. Macintosh users set Macintosh-style permissions on the folders they create.

On a computer running a member of the Windows Server 2003 family, new files and subfolders inherit permissions from the folder in which they are created.

Macintosh files inherit the permissions set on folders. Any Windows Server 2003 family permission specified for a file will be recognized by File Server for Macintosh, even though the Macintosh user won't see any indication in the Finder that these permissions exist. There are four types of Macintosh-style permissions:

- *Cannot Move, Rename, Or Delete* prohibits these actions on a folder.
- *Make Changes* lets a user modify the contents of files in the folder, rename files, move files, create new files, and delete existing files.
- *See Folders* lets a user see what folders are contained in the folder.
- *See Files* lets a user see what files are in the folder and read those files.

A Macintosh user cannot give these permissions to multiple users and groups. Instead, permissions can be assigned to three categories of user.

- **Owner.** The user who created the folder.
- **User/Group.** Similar to the Windows Server 2003 family group associated with the folder. Every folder on a server can have one group associated with it at any one time. The group can be a special group such as *users* or *administrators*, or it can be any other group on the server.
- **Everyone.** All other users of the server, including user accounts with guest access.

The Macintosh security scheme assumes that every folder on a server falls into one of three types: *private information* (accessible only by the owner of the folder); *group information* (accessible by a single workgroup); and *public information* (accessible by everyone).

For example, consider a folder containing information that all members of a certain group should see, but that only one person can change. The person allowed to change the information should be the *Owner* of the folder and should have See Files, See Folders, and Make Changes permissions. The workgroup that uses the folder should be the *Group* associated with the folder and should have only See Files and See Folders permissions. Because no one else needs to see the folder's contents, the *Everyone* category should not be selected.

Although a folder's owner will often be a member of the group associated with the folder, this need not be the case.

With both Macintosh-style and Windows-style permissions, users' access to folders can be defined differently for each folder and subfolder within a tree. For example, you could give a user See Files, See Folders, and Make Changes permissions for one folder, only See Files permission for a subfolder of that folder, and no permissions for another subfolder.

# How file-level permissions are handled

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How file-level permissions are handled**

With a member of the Windows Server 2003 family, users of x86-based computers can assign permissions separately for each file within a folder. The Macintosh computer, however, does not support file-level permissions. When a file has file-level permissions, those permissions apply to Macintosh users only if the permissions are more restrictive than those assigned for the folder that contains the file.

For example, if a Macintosh user has See Files, See Folders, and Make Changes permissions for a folder, the user can read and make changes to files in the folder. However, if that user has only Read permissions for any particular file in that folder, the user can only read that file. Because of the Read file-level permission, the user cannot make changes to the file.

# Translating permissions between x86-based clients and Macintosh clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Translating permissions between x86-based clients and Macintosh clients**

Services for Macintosh translates permissions so that those set by a user of an x86-based computer are translated into the equivalent Macintosh permissions, and vice versa. When a user of an x86-based computer sets permissions for a folder, or when a Macintosh user sets permissions for a folder, permissions are translated according to the following table:

| Windows Server 2003 family permissions | Macintosh permissions |
|---|---|
| Read | See Files, See Folders, or both |
| Write, Delete | Make Changes |

**Note**

- Permissions set within Macintosh networks behave differently from those set from within the Windows Server 2003 family networks, including Macintosh-style permissions. From the Macintosh computer, a right assigned to everyone overrides more restrictive rights set on the owner or a group. From an x86-based computer, permissions assigned to everyone do not override permissions set on the owner or group.

# Setting permissions from a Macintosh or x86-based computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting permissions from a Macintosh or x86-based computer**

A folder's owner can set permissions for the folder. Both the folder's owner and the server administrator can use an x86-based computer to set Windows Server 2003 family permissions for folders on the server. The folder's owner can set permissions for the folder from an x86-based computer because the owner of every folder has the P (Change Permission) permission for that folder.

# Volume passwords

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Volume passwords**
Services for Macintosh provides an extra level of security through Macintosh-accessible *volume passwords*. A volume password is a password you assign to a Macintosh-accessible volume when configuring it. Any Macintosh user who wants to use the volume must type the volume password. Users of x86-based computers do not need to know the volume password to access the folder that corresponds to the Macintosh-accessible volume.

**Notes**

- Volume passwords are optional.
- Volume passwords are case-sensitive.
- When you create a new Macintosh-accessible volume, the default is to have no volume password.
- Because of a constraint with the System 6 and System 7 Finder, you cannot automatically mount a volume with a volume password at startup or by double-clicking an alias. You also cannot automatically mount a volume if the user originally connected to the volume with the Microsoft user authentication module (UAM).

For more information on network security options for users of Services for Macintosh, see Security for Services for Macintosh users.

# Using AppleTalk Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using AppleTalk networking**

This section contains information on:

- Administering Services for Macintosh
- Services for Macintosh Printing
- Services for Macintosh Macintosh-accessible Volumes

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Administering Services for Macintosh

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Services for Macintosh**
This section contains information on:

- Managing Services for Macintosh from the command line
- Stopping, starting, pausing, and resuming services
- Viewing the event log
- Setting up user accounts for Macintosh users

# Managing Services for Macintosh from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Services for Macintosh from the command line**

In addition to using Computer Management, you can also use command-line utilities to manage both File Server for Macintosh and Print Server for Macintosh. For example, use Services for Macintosh commands to start or stop File Server for Macintosh or Print Server for Macintosh. You can also use the Macfile commands to manage File Server for Macintosh servers, volumes, directories, and files.

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools. For more information about Services for Macintosh, see Administer Macintosh Clients.

# Stopping, starting, pausing, and resuming services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Stopping, starting, pausing, and resuming services**

When you set up Services for Macintosh, two services are automatically started: File Server for Macintosh and Print Server for Macintosh. (The AppleTalk Protocol is also started.) At times, you might need to stop these services, as shown in the following table:

| Service | Stop the service to |
|---|---|
| Print Server for Macintosh (MacPrint) | Install another printer driver, to configure a printer, or to immediately see the result of deleting, creating, or changing a printer.<br><br>Stop it to remove the Print Server for Macintosh service. |
| File Server For Macintosh (MacFile) | Change the server name that Macintosh users will see or to remove File Server for Macintosh.<br><br>Pause it when you want to make changes to the server attributes but want to allow current users to continue working. If this service is paused, no new Macintosh users can log on to the computer running Services for Macintosh. |
| AppleTalk Protocol | Change router parameters and the default networking zone and to remove AppleTalk Protocol (which automatically stops File Server for Macintosh and Print Server for Macintosh). |

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Viewing the event log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Viewing the event log**

To check events on the computer running Services for Macintosh, use Event Viewer, which is available in the Administrative Tools folder in Control Panel. If File Server for Macintosh and Print Server for Macintosh are running, you can see events that involved File Server for Macintosh, Print Server for Macintosh, or the AppleTalk Protocol.

# Setting up user accounts for Macintosh users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up user accounts for Macintosh users**

User accounts are created for Macintosh users just as they are created for Windows users. A guest account is automatically created when you install Services for Macintosh. By default, both local guests and guests accessing the server through a client on the network (including a Macintosh client) are allowed. This means that if users log on without a regular user account and password, they will be logged on as a guest.

Guests have some access to shared resources. Guests can do everything that those with a user account can do except keep a local profile on their computers; lock their computers; and create, delete, and modify local groups on their computers.

The guest account cannot be deleted, but it can be disabled. If it is disabled, no network users, including Macintosh users, can log on without a user account and password.

# Services for Macintosh Printing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing**
This section contains information on:

- Creating a printer
- Setting up a user account for Macintosh print jobs
- Using printers on the AppleTalk network
- Prioritizing print jobs
- Creating printing pools

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Creating a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating a printer**

After you have physically attached a printing device to a computer (either directly or on a network), use the Add Printer wizard to create a printer that represents it. You can create more than one printer representing the same printing device.

For example, if you have a printing device in your office but also share it with others over a network, you might want to create two printers representing the printing device. You can create a printer for yourself that is not shared over the network and a second printer that is shared. Then it's easy to control the use of the shared printer. You can set permissions on the shared printer, ensuring that only members of your department can print to it. Or you can set a low priority for it, ensuring that documents you send to the printer will always print before documents sent by those who share it.

Another common example is to create a printer that spools to a printing device at night and another printer that spools to the same printing device during the day.

To create a printer, you must be logged on with sufficient permissions. Members of the Administrators, Server Operators, and Print Operators groups can create printers.

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Setting up a user account for Macintosh print jobs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up a user account for Macintosh print jobs**

After setting up Services for Macintosh, you should create a user account that will be used by all Macintosh clients when printing jobs to captured AppleTalk printing devices or to other devices on the computer running Services for Macintosh. You should also configure Print Server for Macintosh to use this account.

After it is created, the user account (for example, MACUSERS) appears in the list of names that is displayed in the properties dialog box of the print server. You can give specific rights to this user account, just as you would any user account, including Print and No Access.

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Using printers on the AppleTalk network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using printers on the AppleTalk network**

With Services for Macintosh, both x86-based and Macintosh clients can send print jobs to printing devices or spoolers on the AppleTalk network.

The printing device must appear as a LaserWriter in the Chooser, and there must be an appropriate driver for the printing device.

Macintosh clients use printers just as they normally do, through the Chooser. If an AppleTalk printer has been set up, it can be captured so that Macintosh clients cannot access it directly. This causes Macintosh print jobs to go through the computer running Services for Macintosh and to be spooled along with print jobs from x86-based clients.

To allow any Macintosh client to print to an AppleTalk printer directly, disable the capture setting. However, if you do this, the print jobs will not be under the administrator's control.

When an AppleTalk printer is released, any Macintosh user on the AppleTalk network can use the device directly.

A printing device on AppleTalk can be captured when Print Server for Macintosh is set up and a printer is created for it. It must remain captured so that all Macintosh clients send print jobs through the computer running Services for Macintosh. If a printing device has been released, you can recapture it.

You can select another spooler instead of an actual device. However, use this type of configuration with caution. It is possible to create an endless loop of print spooling with this method.

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Prioritizing print jobs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Prioritizing print jobs**

You can create multiple printers that all send print jobs to a single printing device. Each printer has a print-priority level associated with it. If you create two printers and associate them with a single printing device, jobs routed to the printer with the higher priority (lower number) print first.

It's a good idea for users of x86-based computers to create a group that corresponds to each printer. For example, users in Group 1 might have access rights to a priority-1 printer, users in Group 2 might have access rights to a priority-2 printer, and so on. This lets you prioritize print jobs according to the users submitting their jobs.

For Macintosh users, you must create one user account for all incoming print jobs to the computer running Services for Macintosh. Consequently, all Macintosh users sending print jobs through the computer running Services for Macintosh will have the same access rights.

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Creating printing pools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating printing pools**

When you create a printer, you can associate it with more than one printing device in order to form a *printing pool*. To set up a pool, you create a printer and assign it as many output ports as you have identical printing devices. Printing pools have the following characteristics:

- All devices in the pool share the same print property settings and act as a single unit. For example, stopping one device pauses them all.
- Print destinations can be of the same type or mixed (serial, parallel, and network).
- When a job arrives for the printing pool, the spooler on the computer running Services for Macintosh checks the destinations to see which device is idle. The first port selected gets checked first, the second port second, and so on. If your pool consists of a different type of port, make sure you select the fastest port first (network, then parallel, and then serial).
- A printing pool can contain a mixture of printer interface types, but the printing devices must all use the same printer driver.

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Services for Macintosh Macintosh-accessible Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Macintosh-accessible volumes**
This section contains information on:

- Creating a Macintosh-accessible volume
- Creating a Macintosh-accessible volume on a CDFS
- Configuring Macintosh-accessible volumes
- Creating folders in a volume
- Setting permissions for volumes and folders
- Removing a Macintosh-accessible volume

# Creating a Macintosh-accessible volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating a Macintosh-accessible volume**

Just as you can create a shared folder for users of x86-based computers, you can use Shared Folders to designate a folder as a Macintosh-accessible volume. If the folder is to be accessed by x86-based clients as well as Macintosh clients, share the folder by selecting the **Microsoft Windows** check box in the Create Shared Folder wizard. If you do not need to share the files with users of x86-based computers, you can create a volume on a folder. That is, it does not have to be a shared folder.

**Note**

- You cannot designate any folder within a Macintosh-accessible volume as another Macintosh-accessible volume.

# Creating a Macintosh-accessible volume on a CDFS

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating a Macintosh-accessible volume on a CDFS**

To Create a Macintosh-accessible volume on a CDFS volume, follow the same procedure you used to create one on an NTFS-partitioned drive. Note that the CDFS disk is read-only. This means that Services for Macintosh will interpret all security options as See Files and See Folders (read-only).

# Configuring Macintosh-accessible volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Macintosh-accessible volumes**

With the Windows Server 2003 family, you can share folders on the server in any combination. For example, you can share a single folder twice with two different share names, and you can share a folder with one share name and then share a subfolder of that folder with another share name.

However, different rules apply when you use Services for Macintosh to configure Macintosh-accessible volumes. You cannot configure two folders in the same folder tree as volumes.

This means that, when configuring Macintosh-accessible volumes, you cannot configure:

- A single folder twice as two different volumes.
- A folder as a volume if it exists anywhere in the folder tree under another folder configured as a volume.
- A folder as a volume if one of its subfolders, or any subfolder of one of its subfolders, is configured as a volume.

All Macintosh-accessible volumes must be on an NTFS partition or a Compact Disc File System (CDFS) volume. (The CDFS volumes are read-only.) The number of volumes visible to the user is determined by the length of the volume names, which must all fit in a buffer in order to be displayed. (The size of the buffer is determined by an underlying AppleTalk Protocol.) Volume names can have a maximum of 27 characters. Try to strike a balance between clearly naming volumes so that users can identify them easily and keeping the names short so that all of the volume names can be displayed.

**Note**

- Use this formula to determine the number of Macintosh-accessible volume names that can be displayed:

  $N * (M+2) <= 4624$

  where $N$ is the number of volumes and $M$ is the average length of the names in bytes.

# Creating folders in a volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating folders in a volume**

From a computer running Services for Macintosh, you can create folders for a Macintosh-accessible volume or folders for Macintosh clients, just as you would create other folders on Macintosh or Windows-based computers.

On a computer running Services for Macintosh, the folders appear in Windows Explorer as subfolders of the volume. To create another subfolder, you click the **File** menu, point to **New**, and click **Folder**.

On a Macintosh computer, you create folders using the **New Folder** command on the **File** menu. You view and use folders in the Macintosh-accessible volume just as you would any other volume, by using the **View** menu to see the folders organized by Name, Date, Icon, Size, and so forth.

You cannot, however, designate a folder within a Macintosh-accessible volume as another Macintosh-accessible volume.

# Setting permissions for volumes and folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting permissions for volumes and folders**

Just as you set permissions on shared folders to control which users of x86-based computers have access to the folder, you control who can use Macintosh-accessible volumes by setting permissions. Permissions also control the kind of access granted to users. For example, permissions dictate which users can make changes to a folder and which ones can read the content of the folder but not alter it in any way.

**Note**

- Macintosh files inherit the permissions set on folders; you cannot set permissions on files directly.

# Removing a Macintosh-accessible volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Removing a Macintosh-accessible volume**

If you want to make a volume unavailable to Macintosh users, you must remove it. Removing the volume does not delete the files contained in the volume, nor does it delete its status as a shared directory if it has been designated as such for users of x86-based computers. Removing only removes its status as a Macintosh-accessible volume.

# Services for Macintosh Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
Troubleshooting provides solutions for problems that might occur on a computer running Services for Macintosh. This includes problems that both users and administrators might encounter.

For more information on troubleshooting, see:

- Administrator and user issues and solutions
- Printing issues and solutions

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Administrator and user issues and solutions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administrator and user issues and solutions**
What problem are you having?

- A Macintosh-accessible volume is unavailable to a user.
- A Macintosh user's password has expired without notification.
- Users have forgotten their passwords.
- An "incorrect password" message is displayed, even though the password was entered correctly.
- The computer running Services for Macintosh appears in the Chooser on Macintosh clients and then disappears. The appearances are erratic and unpredictable.
- Cannot find a file or folder.
- Cannot save a file with an 8.3 file name from the Macintosh computer.
- Cannot find a server.
- Cannot see any zones within the Chooser on a Macintosh computer.
- The Microsoft user authentication module volume cannot be found.
- The view of a folder is erased or does not match the view selected in the View menu.
- A file is now displayed with the default Windows icon instead of the correct icon.
- A user of an x86-based computer cannot see the contents of a folder.
- A Macintosh user did not receive a server message.
- A user cannot automatically connect to a Macintosh-accessible volume using an alias.

**A Macintosh-accessible volume is unavailable to a user.**
**Cause:** The volume might be configured as a private volume. A private volume is any volume for which the primary and Everyone groups have no access permissions. Only the volume's owner has permissions. In this case, only the owner has access to the volume.

**Solution:** To make the volume accessible to users, the owner should give the primary or Everyone group at least one permission for the volume.

**Cause:** If the Macintosh-accessible volume is on a CDFS volume, and it appears in the Chooser but cannot be selected, the CD-ROM on which it was created might not be in the disk drive.

**Solution:** Be sure that the correct CD-ROM is in the disk drive and that the drive door is closed.

**A Macintosh user's password has expired without notification.**
**Cause:** Users will be notified that their passwords have expired only if the Microsoft user authentication module (MSUAM) files are installed on their clients. If they are using the Apple standard UAM, they will be told only that their logon attempts failed and to try again later. For more information on installing the Macintosh client software, see the Teachtext ReadMe file in the MSUAM volume or see Authentication and Services for Macintosh Concepts.

**Solution:** Install the MSUAM files on the user's client computer.

**Users have forgotten their passwords.**
**Cause:** If users have not logged in for an extended period of time, or the password they used was too complex, they may forget their password.

**Solution:** Assign the users new passwords.

**An "incorrect password" message is displayed, even though the password was entered correctly.**
**Cause:** Users might have two accounts, with different passwords, on separate domains.

**Solution:** Have users enter the domain and then the account name in Name when they log on. For example:

Domain1\user2A

**The computer running Services for Macintosh appears in the Chooser on Macintosh clients and then disappears. The appearances are erratic and unpredictable.**
**Cause:** Two physical AppleTalk networks have been given the same network numbers. The server started first works as expected. When the second server is started, it appears in the Chooser on one Macintosh client, and then disappears and appears in the Chooser on a different client. The order of appearance is unpredictable.

**Solution:** Check the network numbers used for each physical network. When you find the duplicates, change one so that all physical networks use unique network numbers. After you make the change, restart the AppleTalk Protocol on the server on which you made the change. If you find no duplicates, see if your network has a bridge that is filtering packets. It might be filtering out the requests from the second server to find a unique address.

**Cause:** Zones and network numbers are no longer in correspondence. This situation could occur if you recently changed zone names.

**Solution:** When you change the name of a zone, you must shut down the routers directly connected to the networks in question for 10 to 15 minutes before restarting. This allows the other routers to discard old zone information.

**Cannot find a file or folder.**
**Cause:** The user might not have the necessary permissions for the folder that contains the file or folder in question.

**Solution:** The administrator or the owner of the folder can reset permissions to allow the user to access the folder.

**Cannot save a file with an 8.3 file name from the Macintosh computer.**
**Cause:** A file might already exist with that name. However, Macintosh users cannot see it.

**Solution:** Give the file a different name.

**Cannot find a server.**
**Cause:** There may be a problem with the cable system between the client and the server.

**Solution:** Follow these steps:

1. Be sure the cable system between the client and the server is correct. Be sure the network connection, layout, and cable termination conform to the specifications of the cable system you are using.

2. Start with the client that can't find the server. If the cable system is LocalTalk, ensure that the LocalTalk connector box is firmly attached to the printer port on the back of the Macintosh client, not the modem port.

   If the cable system is not LocalTalk, ensure that the network connector is securely connected to its port. In Control Panel, double-click the Network icon to review other network settings.

3. Determine whether other clients are having the same problem.

   If they are, check the cables and connections at the server, and ensure that the server is operating properly. If the server is not the source of the problem, proceed to step 4.

4. Check for breaks in the cable system. If the missing server is on a local network, check each client between the client that can't find the server and the server until you find the server in the Chooser. The break in the cable system is between the client that shows the server in the Chooser and the one that does not.

   If the missing server is on a different physical network on the internet, use your router seeding plan and server

information table to determine which client is the first one beyond the router that links the two networks. Test that client, and then test each client beyond it (in the direction of the server) until the server appears in the Chooser.

If the server was visible at the first client, work backward toward your own network, testing the client adjacent to each router until the server fails to appear in the Chooser. Isolate the break by testing the clients on this network.

5. When you have isolated the network break, check the network cables and connections at that location to make sure all are securely attached, and try again to display the server in the Chooser. If necessary, try replacing cables or connectors.

**Cannot see any zones within the Chooser on a Macintosh computer.**
**Cause:** There might be router problems.

**Solution:** Check for the following:

- The Macintosh computer might be running on an AppleTalk Phase 2 network without the correct Ethernet driver.
- The router might be using Phase 1 and the rest of the internet is using Phase 2.
- The Macintosh computer is configured for one type of network media (such as LocalTalk) but is actually on a network that uses a different media type (such as Ethernet or token ring).

If the problem persists, make sure all routers are configured properly.

**The Microsoft user authentication module volume cannot be found.**
**Cause:** When the computer running Services for Macintosh was installed, there might have been insufficient disk space for the Microsoft user authentication module (UAM) volume. Or the computer running Services for Macintosh might have been installed without an NTFS partition.

**Solution:** You can create the volume by typing and entering the following at the command prompt: **setup /i oemnxpsm.inf /c uaminstall**. This command line copies UAM files to the AppleShare folder in the first NTFS partition and sets up registry values for this volume in the Registry Editor.

**The view of a folder is erased or does not match the view selected in the View menu.**
**Cause:** The Finder occasionally cannot show the correct view of a folder.

**Solution:** The folder owner must log on to the server, connect to the Macintosh-accessible volume, and, on the **View** menu, select a view (such as View By Icon, View By Name, and so on). The selected view then remains in effect.

**A file is now displayed with the default Windows icon instead of the correct icon.**
**Cause:** The application that uses that type of data file might have been removed from the Macintosh computer.

**Solution:** If the file had no resource fork, use the Apple ResEdit utility to reset the file type and file creator of the file. Use this utility only if you are experienced with it.

**A user of an x86-based computer cannot see the contents of a folder.**
**Cause:** The user of the x86-based computer does not have sufficient permissions to view the contents of the folder.

**Solution:** Use the computer running Services for Macintosh to make sure the user has Read permission, or the folder owner can use a Macintosh computer to give the user of the x86-based computer both the See Files and See Folders permissions. (A user of an x86-based computer must have both these permissions to get the Read permission.)

**A Macintosh user did not receive a server message.**
**Cause:** The Macintosh user is running a version of the AppleTalk Filing Protocol that is older than version 2.1.

**Solution:** Make sure that the user upgrades to version 3.0 of AppleShare, which allows users to receive messages.

**A user cannot automatically connect to a Macintosh-accessible volume using an alias.**
**Cause:** Macintosh clients can be configured to automatically connect to volumes when the client is started or when the user double-clicks an alias to an object on a volume. However, automatic connection to volumes is not supported by the Macintosh system software if the volume is configured with a volume password, or if the user originally used the Microsoft user authentication module (MSUAM) to connect to the volume.

**Solution:** If the volume has a password, you can mount it through the Chooser and then use the alias. Or you can specify that it be opened at system startup time when you mount the volume. If you are using MSUAM to log on to the server, you must

mount the volume through the Chooser and then use the alias.

# Printing issues and solutions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing issues and solutions**
What problem are you having?

- AppleTalk printers don't appear in the Printers Folder Available AppleTalk Printer dialog box.
- Messages consistently appear while documents are being printed.
- The PostScript message "Offending command" appears at the end of the printed document or elsewhere.
- Print jobs fail to print.
- Macintosh extended characters (such as bullets, smart quotes, and copyright and trademark symbols) are changed into other characters on the LaserWriter II.

**AppleTalk printers don't appear in the Printers Folder Available AppleTalk Printer dialog box.**
**Cause:** Clicking the AppleTalk zone name does not display the printers in that zone.

**Solution:** You must double-click the zone name in this dialog box.

**Messages consistently appear while documents are being printed.**
**Cause:** The printer needs to be reset.

**Solution:** Reset the printing device by turning it off and then on again.

**The PostScript message "Offending command" appears at the end of the printed document or elsewhere.**
**Cause:**

- A user or administrator might have canceled the print job while it was spooling.
- A user is spooling to a PSTODIB printing device, and it has PostScript level 2 elements or is a PostScript level 2 document.

**Solution:** No action is necessary, and you can reprint the file as desired.

**Print jobs fail to print.**
**Cause:** If one of the printing devices is turned off, all printing devices can stop printing.

**Solution:** Check each printing device that prints jobs for these printers.

**Macintosh extended characters (such as bullets, smart quotes, and copyright and trademark symbols) are changed into other characters on the LaserWriter II.**
**Cause:** If the LaserWriter hasn't been set correctly, printing problems can occur, regardless of how you set the COM port in Control Panel. This problem affects Macintosh computers more frequently than x86-based computers because Macintosh computers use extended characters more often than other clients do.

**Solution:** Set the communications port for the LaserWriter correctly, referring to the owner's manual for the printing device.

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

# Internet and E-mail Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet and E-mail Services**
This section covers:

- Using Internet Explorer
- Using Windows Media Services
- Using the FrontPage Server Extensions 2002
- UDDI Services overview
- E-mail Services
- Background Intelligent Transfer Service

**Note**

- In the Windows Server 2003 family, you can implement enhanced security settings for Internet Explorer for all users and reduce the exposure of your server to Web sites that might pose a security risk. For more information, see Internet Explorer Enhanced Security Configuration.

# Using Internet Explorer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Internet Explorer**

With Internet Explorer, you can search for and view information on your corporate network or the World Wide Web. You can type the address of the Web page you want to visit into the address bar, or click an address from your list of Favorites. Internet Explorer also lets you search the Internet for people, businesses, and information about subjects that interest you.

In the Windows Server 2003 family, you can implement enhanced security settings for Internet Explorer for all users and reduce the exposure of your server to Web sites that might pose a security risk. For more information, see Internet Explorer Enhanced Security Configuration.

To open Internet Explorer, click **Start**, point to **All programs**, and then click **Internet Explorer**.

**Note**

- For information about using Internet Explorer, click the **Help** menu in Internet Explorer.

# Using Windows Media Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Windows Media Services**

With Windows Media Services, you can manage, deliver, and archive Windows Media content, including streaming audio and video, over an intranet or the Internet. Windows Media Services includes server components such as Windows Media Services Administrator for the Web. Windows Media Services is not installed by default.

**To install Windows Media Services**

1. Open the Windows Components Wizard.
2. In the Windows Components Wizard, select the **Windows Media Services** check box, and then click **Details**.
3. In the **Windows Media Services** dialog box, select the **Windows Media Services Administrator for the Web** check box.

## Notes

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- To open Windows Media Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Windows Media Services**.
- Windows Media Services and Windows Media Services Administrator for the Web are not included with Windows Server 2003, Web Edition, and the Itanium-based versions of the Windows operating systems.
- For more information about using Windows Media Services and Windows Media Services Administrator for the Web, in Windows Media Services Administrator for the Web, click **Help**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Using the FrontPage Server Extensions 2002

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the FrontPage Server Extensions 2002**

The Microsoft FrontPage Server Extensions 2002 are a set of programs that work with Microsoft Internet Information Services to support administering, authoring, and browsing a FrontPage-extended Web site. Users running the FrontPage client can then create and manage professional-looking Web sites, whether they're creating a personal Web page or a corporate Internet or intranet site. FrontPage Server Extensions 2002 gives you, as the administrator, tools for managing site security, organizing your content into subwebs, and checking site usage.

Microsoft FrontPage Server Extensions 2002 are not installed by default during Windows Setup.

**To install Microsoft FrontPage Server Extensions 2002**

1. Open **Add or Remove Programs** in Control Panel.
2. Click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, double-click **Application Server**, and then select **Internet Information Services**.

### Notes

- To open FrontPage 2002 Server Extensions, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Information Services**.
- Once installed, the Microsoft FrontPage Server Extensions 2002 will run the next time you restart your server with Internet Information Services enabled.

# UDDI Services overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**UDDI Services overview**

Universal Description, Discovery, and Integration (UDDI) is an industry specification for publishing and locating information about Web services. The Windows Server 2003 family includes UDDI Services, an optional component that provides UDDI capabilities for use within an enterprise or between business partners. UDDI Services is a standards-based XML Web service that enables enterprise developers to efficiently publish, discover, share, and reuse Web services directly through their development tools and business applications.

Built on the Microsoft .NET Framework, UDDI Services is a reliable and scalable solution that provides easy integration with enterprise technologies and tools. Information Technology (IT) managers can leverage native support for standard categorization schemes, Microsoft SQL Server, and Active Directory authentication. As a UDDI version 1.0 and 2.0 API-compliant service, UDDI Services also includes a Web interface that is translated into all of the languages that are supported by the Windows Server 2003 family.

UDDI Services product documentation is available in the UDDI Services Help. To view product documentation, install UDDI Services, and then see Windows Server 2003 Product Help. For information about installing UDDI Services, see Using UDDI Services.

**Notes**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.
- To open UDDI Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **UDDI Services**.
- To run UDDI Services, you must also run Microsoft Internet Information Services (IIS).
- For additional planning, installation, configuration, troubleshooting, development, or other UDDI Services resources, see Enterprise UDDI Services at the Microsoft Web site.

# Using UDDI Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using UDDI Services**
Universal Description, Discovery, and Integration (UDDI) is an industry specification for the description and discovery of Web services. The UDDI specification builds on the Simple Object Access Protocol (SOAP), Extensible Markup Language (XML), and HTTP/S protocol standards developed by the World Wide Web Consortium (W3C) and Internet Engineering Task Force (IETF). UDDI Services is an optional component of the Windows Server 2003 family that provides UDDI capabilities within an enterprise.

UDDI Services is not installed by default during setup of a Windows Server 2003 operating system.

**UDDI Services requirements checklist**
Before you install and deploy UDDI Services, make sure your system meets the following hardware and software requirements.

**Hardware requirements**
In addition to the Windows Server 2003 family minimum requirements, UDDI Services requires:

- 128 megabytes (MB) of RAM minimum
- 80 MB free disk space (plus additional storage for UDDI Services data)

**Software requirements**
You must use one of the following Microsoft operating systems to host UDDI Services:

- Windows Server 2003, Standard Edition
- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition

**Installing UDDI Services**
When you first install UDDI Services, all components are selected by default. By deselecting components, you can decrease or disable UDDI Services functionality. If you are unfamiliar with UDDI Services components or deployment options, complete the default installation, and then view the Deploying and Installing UDDI Services topic in Windows Server 2003 Product Help.

**To install UDDI Services, or add or remove a component**

1. Open the Windows Components Wizard.
2. In the Windows Components Wizard, select the **UDDI Services** check box, and then follow the instructions to install, remove, or add components to UDDI Services.

After you have installed UDDI Services, open UDDI Services.

### Notes

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- To open UDDI Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **UDDI Services**.
- To run UDDI Services, you must also run Microsoft Internet Information Services (IIS).
- You can use Administrative Tools to administer local or remote installations of UDDI Services.

- If you have already installed UDDI Services, see Windows Server 2003 Product Help.

- After installing and configuring UDDI Services, you can view the Web-based user interface on any server to which you have installed the user interface subcomponent. For example, **http://*serverName*/uddi/**, or if you are using a secure sever, **https://*serverName*/uddi/**.

- For more information about UDDI Services, see UDDI Services overview.

- For additional planning, installation, configuration, troubleshooting, development, or other UDDI Services resources, see UDDI Enterprise Services at Microsoft Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# E-mail Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**E-mail Services**
With E-mail Services, you can install components on your computer to configure it as a mail server. E-mail Services includes the Post Office Protocol 3 (POP3) service and the Simple Mail Transfer Protocol (SMTP) service, which retrieve and transfer e-mail, respectively. To provide e-mail services to users, such as sending and receiving e-mail, administrators can create mailboxes on the server.

This section covers:

- Simple Mail Transfer Protocol (SMTP) service
- POP3 Service

# Simple Mail Transfer Protocol (SMTP) service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Simple Mail Transfer Protocol (SMTP) service**
The Simple Mail Transfer Protocol (SMTP) service, which transfers e-mail, is installed as part of E-mail Services, along with the POP3 service.

SMTP controls how e-mail is transported and then delivered across the Internet to the destination server. The SMTP service sends and receives e-mail between servers, whereas the POP3 service retrieves e-mail from the mail server to the user's computer. For more information, see the SMTP documentation Microsoft SMTP Service.

For more information about configuring the SMTP service for use with the POP3 service, see Simple Mail Transfer Protocol (SMTP) and Configure the Simple Mail Transfer Protocol Service.

**Managing the SMTP service**
You can manage the SMTP service from the Internet Information Services (IIS) Microsoft Management Console (MMC) snap-in. To open the IIS snap-in, click **Start**, and then click **Run**. In **Open**, type **inetmgr**, and then click **OK**. If the IIS snap-in does not appear, verify that you have installed IIS.

# POP3 Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**POP3 service**
The POP3 service is an e-mail service that retrieves e-mail. Administrators can use the POP3 service to store and manage e-mail accounts on the mail server.

When the POP3 service is installed on the mail server, users can connect to the mail server and retrieve and download e-mail to their local computer using an e-mail client that supports the POP3 protocol (such as Microsoft Outlook). The POP3 service is used with the Simple Mail Transfer Protocol (SMTP) service, which sends outgoing e-mail.

- Before proceeding, see Checklist: Installing and configuring the POP3 service.
- For tips about using the POP3 service, see Best practices for POP3 service.
- For help with specific tasks, see How To... For E-Mail Services.
- For general background information, see E-mail Services Concepts.
- For problem-solving instructions, see Troubleshooting E-mail services.

# Checklist: Installing and configuring the POP3 service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing and configuring the POP3 service**

| Step | Reference |
|---|---|
| Read recommended topics. | Best practices for POP3 service;<br><br>POP3 service overview;<br><br>E-mail concepts;<br><br>Name resolution for TCP/IP |
| Choose and configure the mail server running the POP3 service with a Domain Name System (DNS) name. It is recommended that you choose a name that reflects the role of the server, such as mailserver1. | Rename a computer |
| Contact your Internet Service Provider (ISP). | Contacting your Internet Service Provider |
| Configure TCP/IP for static addressing on the mail server on which the POP3 service is installed. | Configure TCP/IP for static addressing |
| Install the POP3 service. | Install e-mail services |
| Read "POP3 service authentication methods" and determine which method is best for your deployment. | POP3 Service Authentication Methods |
| Configure the POP3 service to use the authentication method you have chosen. | Set the authentication method |
| Create a domain. | Create a domain |
| Create a mailbox. | Create a mailbox |
| Configure all e-mail clients to use the POP3 service. In addition, if you are using Secure Password Authentication (SPA), you should now configure your e-mail clients to use SPA. | Configure E-mail Clients |
| Manage your mail server. | POP3 service administration tools |

# Best practices for POP3 service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Set the mail store on a volume other than the one on which the operating system is installed.**

- Because the mail store can potentially use large amounts of disk space, you should either set a disk quota limit on the volume of the mail store (to control its disk space usage) or set it to use a volume other than the one where the operating system is installed. This will prevent the possibility of the operating system running out of disk space if the mail store becomes too large. For more information, see Set the mail store. For more information on disk quotas, see Configuring disk quotas for the POP3 service.

**Set the Badmail directory on a volume other than the one on which the operating system is installed.**

- When an e-mail is undeliverable, the Simple Mail Transfer Protocol (SMTP) service returns it to the sender with a non-delivery report (NDR). If the NDR cannot be delivered to the sender, a copy of the message is put in the Badmail directory. To prevent the possibility of the operating system running out of disk space (if the Badmail directory becomes too large), you must move the SMTP Badmail directory to a volume other than the one on which the operating system is installed.

**Use an NTFS file system partition for the mail store and implement disk quotas.**

- NTFS file system partitions provide greater directory and folder security. This enhances the protection of e-mail stored on the local hard disk.

- Disk quotas can prevent the mail store from using an excessive or unanticipated amount of disk space. This is important because excessive or unanticipated use of disk space can adversely affect the performance of the server on which the POP3 service is running.

**Allocate adequate disk space for the mail store.**

- To ensure that you allocate adequate disk space for the mail store, estimate the amount of disk space it will need based on the number of users on the server, the volume of e-mail they will receive, and the average size of the e-mail they will receive.

**Use strong passwords to protect your mail server.**

- Strong passwords protect your mail server from unauthorized access. For more information about passwords, see Strong passwords.

# How To... For E-Mail Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Set Up the POP3 Service
- Change Service Settings
- Manage E-mail Domains
- Manage Mailboxes
- Configure the Simple Mail Transfer Protocol Service
- Configure E-mail Clients

# Set Up the POP3 Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up the POP3 service**

- Install e-mail services
- Set the authentication method
- Set the mail store

# Install e-mail services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install e-mail services**

1. Open Add or Remove Programs.
2. Click **Add/Remove** Windows **Components**.
3. Scroll down and select the **E-mail Services** check box.
4. Click **Next**, and then complete the wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**.
- You can also install e-mail services using the Configure Your Server Wizard. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set Up the POP3 Service
Checklist: Installing and configuring the POP3 service
Mail server role: Configuring a mail server

# Set the authentication method

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the authentication method**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, right-click the *computer_name* node and click **Properties**.

   **Where?**

   - POP3 Service/*computer_name*

3. In **Authentication Method**, click the drop-down menu, and then click the authentication method you want to use.

### Important

- You can change the authentication method only if there are no existing domains defined. If there are existing domains, the option to change the authentication method is disabled.

  If you change this parameter, you must stop and restart the POP3 service and the IIS Admin service. For more information on how to change the POP3 service state and start and stop the IIS Admin service, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- If the computer on which the POP3 service is running is either a member of an Active Directory domain or a domain controller, the default authentication method is Active Directory integrated authentication. If the computer is not in a domain, the default is local Windows accounts authentication.
- If the computer on which the POP3 service is running is a domain controller, the available authentication methods are Active Directory integrated authentication and encrypted password file authentication. If the computer on which the POP3 service is running is a member server in an Active Directory domain, all three authentication methods are available. Otherwise, the available authentication methods are local Windows accounts authentication and encrypted password file authentication.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **winpop set** {**authentication**|**auth**} {**1**|**2**|**3**}

| Value | Description |
|---|---|
| **winpop set** {**authentication**\|**auth**} | Sets the type of authentication to be used by the POP3 service. The **auth** command provides the same functionality as **authentication**. |
| {**1**\|**2**\|**3**} | Specifies the authentication method POP3 uses to verify users who connect to the mail server. The available authentication methods vary, depending on the configuration of your server.<br>● If the computer on which the POP3 service is running is an Active Directory domain controller, the values are as follows:<br>**1** - Active Directory integrated authentication<br>**2** - Encrypted password file authentication<br>● If the computer on which the POP3 service is running is an Active Directory member server, the values are as follows:<br>**1** - Active Directory integrated authentication<br>**2** - Local Windows accounts authentication<br>**3** - Encrypted password file authentication<br>● Otherwise, the following values apply:<br>**1** - Local Windows accounts authentication<br>**2** - Encrypted password file authentication |

**Important**

- You can change the authentication method only if there are no existing domains defined. If there are existing domains, the option to change the authentication method is disabled.

  If you change this parameter, you must stop and restart the POP3 service and the IIS Admin service. For more information on how to change the POP3 service state and start and stop the IIS Admin service, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If the computer on which the POP3 service is running is either a member of an Active Directory domain or a domain controller, the default authentication method is Active Directory integrated authentication. If the computer is not in a domain, the default is local Windows accounts authentication.

- If the computer on which the POP3 service is running is an Active Directory domain controller, the available authentication methods are Active Directory integrated authentication and encrypted password file authentication. If the computer on which the POP3 service is running is a member server in an Active Directory domain, all three authentication methods are available. Otherwise, the available authentication methods are local Windows accounts authentication and encrypted password file authentication.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop set help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
POP3 Service Authentication Methods
Change Service Settings
Change the POP3 service state
Start, stop, pause, resume, or restart a service

# Set the mail store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the mail store**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, right-click the *computer_name* node, and then click **Properties**.

   **Where?**

   - POP3 Service/*computer_name*

3. In **Root Mail Directory**, type the path to the mail store; for example, C:*mailstore*. The maximum value is 260 characters.

**Important**

- If you change this parameter, you must stop and restart the POP3 service and the IIS Admin service. For more information on how to change the POP3 service state and start and stop the IIS Admin service, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- You cannot set the mail store to the root directory of the hard disk, for example C:\, or to a directory in which files are currently in use.
- It is recommended that you locate the mail store on a NTFS file system partition.
- The default mail store is *system_drive*:\Inetpub\mailroot\Mailbox.
- The mail store must be a directory on the local file system or a Universal Naming Convention (UNC) path; mapped drives are not supported.
- If you change the mail store from its default location or from a previous location, and there is stored e-mail, you must manually move the folders in which the e-mail is stored to the new location.
- If you transfer the mail store to a new directory, you must move the mail store directory to ensure the directory retains the correct ownership; copying the mail store will not work.
- If you move the mail store to a new location or restore it from a backup, you must reset the permissions on the directory

using the **winpop set mailroot** command.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- Because the mail store can potentially use large amounts of disk space, you should either set a disk quota limit on the volume of the mail store (to control its disk space usage) or set it to use a volume other than the one where the operating system is installed. This will prevent the possibility of the operating system running out of disk space if the mail store becomes too large. For more information on disk quotas, see Configuring disk quotas for the POP3 service.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **winpop set mailroot***path_to_mail_store*

| Value | Description |
|---|---|
| **winpop set mailroot** | Sets the mail store for the server. |
| *path_to_mail_store* | Specifies a directory on the local file system or a Universal Naming Convention (UNC) path. The maximum allowable value is 260 characters. |

**Important**

- If you change this parameter, you must stop and restart the POP3 service and the IIS Admin service. For more information on how to change the POP3 service state and start and stop the IIS Admin service, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You cannot set the mail store to the root directory of the hard disk (for example C:\), or to a directory in which files are currently in use.

- It is recommended that you locate the mail store on a NTFS file system partition.

- The default mail store is *system_drive*:\Inetpub\mailroot\Mailbox.

- The mail store must be a directory on the local file system or a Universal Naming Convention (UNC) path; mapped drivers are not supported.

- If you change the mail store from its default or from a previous location, and there is stored e-mail, you must manually move the folders in which the e-mail is stored to the new location.

- If you move the mail store to a new location or restore it from a backup, you must reset the permissions on the directory using the **winpop set mailroot** command.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local

computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop set help**

- Because the mail store can potentially use large amounts of disk space, you should either set a disk quota limit on the volume of the mail store (to control its disk space usage) or set it to use a volume other than the one where the operating system is installed. This will prevent the possibility of the operating system running out of disk space if the mail store becomes too large. For more information on disk quotas, see Configuring disk quotas for the POP3 service.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the POP3 service state
POP3 service administration tools
Mail store
Change Service Settings
Start, stop, pause, resume, or restart a service

# Change Service Settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Change service settings**

- Configure the mail server to require Secure Password Authentication
- Change the POP3 service state
- Set the port
- Set the event logging level
- Set the number of sockets
- Set the number of threads
- Connect to a remote server running the POP3 service

# Configure the mail server to require Secure Password Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the mail server to require Secure Password Authentication**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, right-click the *computer_name* node and click **Properties**.

   **Where?**

   - POP3 Service/*computer_name*

3. Select **Require Secure Password Authentication (SPA) for all client connections**.

### Important

- If you change this parameter, you must stop and restart the POP3 service. For more information on changing the POP3 service state, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- SPA supports only Active Directory integrated authentication and local Windows accounts authentication.
- If you enable SPA, users' e-mail clients must also be configured to use SPA. For more information, see Related Topics.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.
- Configuring the mail server to require Secure Password Authentication affects only the POP3 service and not the Simple Mail Transfer Protocol (SMTP) service.

**Using a command line**

1. Open Command Prompt.
2. Type:

**winpop set sparequired** {**0**|**1**}

| Value | Description |
|---|---|
| **winpop set sparequired** | Specifies whether Secure Password Authentication is required for all client connections. |
| {**0**\|**1**} | Specifies whether secure authentication is required from all e-mail clients.<br><br>The default is 0, which specifies that SPA is not required. 1 requires SPA from all e-mail clients and prevents e-mail clients from authenticating by using plaintext authentication. |

**Important**

- If you change this parameter, you must stop and restart the POP3 service. For more information on changing the POP3 service state, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- SPA supports only Active Directory integrated authentication and local Windows accounts authentication.
- If you enable SPA, users' e-mail clients must also be configured to use SPA. For more information, see Related Topics.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.
- To view the complete syntax for this command, at a command prompt, type:

  **winpop set help**

- Configuring the mail server to require Secure Password Authentication affects only the POP3 service and not the Simple Mail Transfer Protocol (SMTP) service.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 Service Authentication Methods
Configure Outlook Express for Secure Password Authentication
Change the POP3 service state
Winpop

# Change the POP3 service state

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the POP3 service state**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, right-click *ComputerName*, point to **All Tasks**, and then click the state you want, as described in the table.

   **Where?**

   - POP3 Service/*computer_name*

| Service state | Description |
|---|---|
| Start | Starts the POP3 service from a stopped or paused state and allows POP3 clients to connect to the mail server and retrieve e-mail. |
| Stop | Stops the POP3 service. All current connections are terminated. |
| Pause | Pauses the POP3 service so that no new connections will be accepted. Existing connections are not affected. |
| Resume | Resumes the POP3 service from a paused state to allow POP3 clients to connect to the server and retrieve e-mail. |
| Restart | Restarts the POP3 service. All current connections are terminated. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **net** {**start**|**stop**} **pop3svc**

| Value | Description |
| --- | --- |
| **net** {**start**\|**stop**} | Starts or stops a service. |
| **pop3svc** | Specifies the POP3 service. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **net start /help**

  **net stop /help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Change Service Settings

# Set the port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To set the port

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open POP3 service.
2. In the console tree, right-click the *computer_name* node and click **Properties**.

   **Where?**

   - POP3 Service/*computer_name*

3. In **Server Port**, type an appropriate port number from 1 to 65535.

### Important

- If you change this parameter, you must stop and restart the POP3 service. For more information on changing the POP3 service state, see Related Topics.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- The port used by the POP3 service is configured by default to 110, the Internet standard for the POP3 protocol. If you change the port from 110, you must notify users to configure their e-mail clients to reflect the change.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **winpop set port***port_number*

| Value | Description |
|---|---|

| | |
|---|---|
| **winpop set port** | Sets the port that POP3 service monitors for clients. |
| *port_number* | Specifies any valid port number from 1 to 65535. |

**Important**

- If you change this parameter, you must stop and restart the POP3 service. For more information on changing the POP3 service state, see Related Topics.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- The port used by the POP3 service is configured by default to 110, the Internet standard for the POP3 protocol. If you change the port from 110, you must notify users to configure their e-mail clients to reflect the change.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop set help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Change Service Settings
Change the POP3 service state

# Set the event logging level

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the event logging level**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, right-click the *computer_name* node, and then click **Properties**.

   **Where?**

   - POP3 Service/*computer_name*

3. In **Logging level**, click the drop-down menu, and then click the logging level you want, as described in the following table.

| POP3 event logging level | Description |
|---|---|
| None | No events are logged. |
| Low | Only critical events are logged. |
| Medium | Critical and warning events are logged. |
| High | Critical, warning, and informational events are logged. |

**Important**

- If you change this parameter, you must stop and restart the POP3 service and the IIS Admin service. For more information on how to change the POP3 service state and manage the IIS Admin service, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **winpop set logging** {**0**|**1**|**2**|**3**}

| Value | Description |
|---|---|
| **winpop set logging** | Sets the event logging level for the POP3 service. For example, to log only critical events type: **winpop set logging 1**. |
| {**0**\|**1**\|**2**\|**3**} | Specifies the level at which events are logged.<br><br>**0** - No events are logged.<br><br>**1** - Only critical events are logged.<br><br>**2** - Critical and warning events are logged.<br><br>**3** - Critical, warning, and informational events are logged. |

**Important**

- If you change this parameter, you must stop and restart the POP3 service and the IIS Admin service. For more information on how to change the POP3 service state and manage the IIS Admin service, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop set help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Change Service Settings
Change the POP3 service state
Event Viewer overview
Event Viewer
Start, stop, pause, resume, or restart a service

# Set the number of sockets

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the number of sockets**

1. Open Command Prompt.
2. Type:

   **winpop set socketmin** *number_of_sockets*

   or

   **winpop set socketmax***number_of_sockets*

   or

   **winpop set socketthreshold***number_of_sockets*

   or

   **winpop set socketbacklog***number_of_sockets*

   or

   **winpop set socket***socketbacklog socketmax socketmin socketthreshold*

| Variable | Description | Numeric value for ***number_of_sockets*** |
|---|---|---|
| **winpop set socketmin** | Sets the minimum number of sockets that the POP3 service can create (*socketmin*). The default setting is 10. | The value of *number_of_sockets* for this parameter must be from 1 to 32,000 and less than *socketmax*. |
| **winpop set socketmax** | Sets the maximum number of sockets the service can create (*socketmax*). Once *socketmax* has been reached, any new connection attempts are held in a queue (determined by *socketbacklog*) until an existing connection is disconnected. The default value is 50. | The value of *number_of_sockets* for this parameter must be from 1 to 32,000 and greater than *socketmin* + *socketthreshold*. |
| **winpop set socketthreshold** | Sets the number of new sockets that the POP3 service can create at a time. Once *socketmin* is reached, *n* number of new sockets are created (if there are still remaining sockets available), where *n* is the value of *socketthreshold*. | The value of *number_of_sockets* for this parameter must be from 1 to 99, or 0 if *socketmin* = *socketmax*. |
| **winpop set socketbacklog** | Sets the number of connections that are held in a queue once *socketmax* has been reached (determined by *socket backlog*). Connection requests that occur after *socketmax* has been reached, and are less than *socketmax*+*socketbacklog*, are accepted, but do not receive any response until a new socket becomes available. Once *socketbacklog* is filled, new connection attempts are denied. | The value of *number_of_sockets* for this parameter must be from 0 to 100. |

| | | |
|---|---|---|
| **winpop set socket** | Sets all the socket values. You must provide a numeric value for each of the four four variables. | Each numeric value should follow the specifications for each variable, as described previously in this table; for example, **winpop set socket 5 40 10 5**. |

**Important**

- Do not change these parameters unless it is necessary to do so. Changing these parameters can adversely affect the performance of the POP3 service.

- If you change these parameters, you must stop and restart the POP3 service. For more information on changing the POP3 service state, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

-

-

- To view the complete syntax for this command, at a command prompt, type:

  **winpop set help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Change Service Settings
Change the POP3 service state

# Set the number of threads

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the number of threads**

1. Open Command Prompt.
2. Type:

   **winpop set threadcount***number_of_threads*

| Value | Description |
|---|---|
| **winpop set threadcount** | Sets the number of threads for the POP3 service. This value is based on the number of CPUs that are installed on the computer on which the POP3 service is running. |
| *number_of_threads* | Specifies the number of CPUs installed on the computer on which the POP3 service is running. This value must be from 1 to 32. |

**Important**

- Do not change this parameter unless it is necessary to do so. Changing this parameter can adversely affect the performance of the POP3 service.
- If you change this parameter, you must stop and restart the POP3 service. For more information on changing the POP3 service state, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.
- To view the complete syntax for this command, at a command prompt, type:

  **winpop set help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Change Service Settings
Change the POP3 service state

# Connect to a remote server running the POP3 service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a remote server running the POP3 service**

1. Open POP3 service.
2. In the console tree, right-click **POP3 Service**, and then click **Connect**.
3. Do one of the following:
   - In **Server Name**, type the name of the computer to which you want to connect, and then click **OK**.
   - Or, click **Browse** to view the list of computers available on your network, select the appropriate computer, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, point to **Administrative Tools**, and then click **POP3 Service**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.
- To connect to a remote mail server running the POP3 service, your computer must also have the POP3 service installed and be on the same local network.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Change Service Settings
Remote administration of POP3 service

# Manage E-mail Domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage domains**

- Create a domain
- Delete a domain
- List domains
- Lock a domain
- Unlock a domain
- View domain statistics

# Create a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a domain**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, right-click the *computer_name* node, point to **New**, and then click **Domain**.

   **Where?**

   - POP3 Service/*computer_name*

3. In **Domain name**, type the name of the domain to be created, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- Contact your Internet Service Provider (ISP) for domain naming requirements and conventions. For information about this, see Related Topics.
- The POP3 service supports top-level and third-level domain names. For example, example.com and mailserver.example.com are both supported.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **winpop add***domain_name*

| Value | Description |
|---|---|
| **winpop add** | Creates a new domain and adds it to the existing POP3 service directory structure. |

| | |
|---|---|
| *domain_name* | Specifies the name of the domain you want to create, for example, *example.com*. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Contact your Internet Service Provider (ISP) for domain naming requirements and conventions. For information about this, see Related Topics.
- The POP3 service supports top-level and third-level domain names. For example, example.com and mailserver.example.com are both supported.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.
- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage E-mail Domains
Contacting your Internet Service Provider

# Delete a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a domain**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, click *ComputerName*, right-click the domain you want to delete, and then click **Delete**.
3. When the confirmation box appears, click **Yes** to delete the domain, all of the mailboxes stored in that domain, and all mail stored in that domain.

### Caution

- Deleting a domain will delete all of the mailboxes in that domain, the corresponding mail store directories, and all e-mail stored in those directories.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- You cannot delete a domain while a user in that domain is connected to the server running the POP3 service.
- If you are using local Windows accounts or Active Directory integrated authentication, and you have created user accounts, the user accounts are not deleted when the corresponding mailboxes are deleted.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **winpop delete***domain_name*

| Value | Description |
|---|---|
| **winpop delete** | Deletes an existing domain. |

| | |
|---|---|
| *domain_name* | Specifies the name of the domain you want to delete, for example, *example.com*. |

**Caution**

- Deleting a domain will delete all of the mailboxes in that domain, the corresponding mail store directories, and all e-mail stored in those directories.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are using local Windows accounts or Active Directory integrated authentication, and you have created user accounts, the user accounts are not deleted when their corresponding mailboxes are deleted.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage E-mail Domains

# List domains

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To list domains**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, click the *computer_name* node. All domains on that server appear in the details pane.

   **Where?**

   - POP3 Service\computer_name

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **POP3 Service**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **winpop list**

| Value | Description |
|---|---|
| **winpop list** | Lists all POP3 service domains. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local

computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage E-mail Domains

# Lock a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To lock a domain

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open POP3 service.
2. In the console tree, click the *computer_name* node.

   **Where?**

   - POP3 Service/*computer_name*

3. In the details pane, right-click the domain you want to lock, and then click **Lock**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- When you lock a domain, you prevent all members of the domain from retrieving their e-mail. Incoming e-mail for the domain is still received and delivered to the appropriate mailbox directory in the mail store. The service will also continue to send outgoing mail.
- When a domain is locked, individual mailboxes cannot be unlocked; the entire domain must be unlocked.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **winpop lock***domain_name*

| Value | Description |
|---|---|
| **winpop lock** | Locks all mailboxes in a domain, which prevents users from downloading mail using the POP3 service. |

| | |
|---|---|
| *domain_name* | Specifies the domain you want to lock; for example, example.com |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- When you lock a domain, you prevent all members of the domain from retrieving their e-mail. Incoming e-mail for the domain is still received and delivered to the appropriate mailbox directory in the mail store. The service will also continue to send outgoing mail.

- When a domain is locked, individual mailboxes cannot be unlocked; the entire domain must be unlocked.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage E-mail Domains

# Unlock a domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To unlock a domain**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, click the *computer_name* node to display the list of domains.

   **Where?**

   - POP3 Service/*computer_name*

3. In the details pane, right-click the locked domain you want to unlock, and then click **Unlock**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- When a domain is locked, individual mailboxes cannot be unlocked; the entire domain must be unlocked.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **winpop unlock***domain_name*

| Value | Description |
|---|---|
| **winpop unlock** | Unlocks a locked domain to permit users to download e-mail again using the POP3 service. |
| *domain_name* | Specifies the name of the domain that you want to unlock, for example, *example.com* |

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group

might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- When a domain is locked, individual mailboxes cannot be unlocked; the entire domain must be unlocked.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage E-mail Domains

# View domain statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view domain statistics**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, click the domain you for which you want to view statistics.

   **Where?**

   - POP3 Service/*computer_name/domain_name*

   The domain statistics appear in the details pane.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- Depending on the number of domains and mailboxes on the server running the POP3 service, this procedure might require considerable CPU resources.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **winpop stat**

| Value | Description |
|---|---|
| **winpop stat** | Displays statistics for the domain, including total number of mailboxes in the domain, lock/unlock status of the domain, total number of messages, and disk usage. |

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group

might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Depending on the number of domains and mailboxes on the server running the POP3 service, this procedure might require considerable CPU resources.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage E-mail Domains

# Manage Mailboxes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage mailboxes**

- Create a mailbox
- Delete a mailbox
- List mailboxes
- Lock a mailbox
- Unlock a mailbox
- Change a password
- View mailbox statistics
- Create a quota file
- Migrate an encrypted password file user account to an Active Directory user account

# Create a mailbox

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a mailbox**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, click the *computer_name* node.

   **Where?**

   - POP3 Service/*computer_name*

3. In the details pane, right-click the domain in which you want to create a new mailbox, point to **New**, and then click **Mailbox**.
4. Check the **Create associated user for this mailbox** box if you are using local Windows accounts authentication or Active Directory integrated authentication, unless a user account already exists with the same name as the mailbox you want to create. If the check box is already selected, clear it only if an account already exists with the same name as the mailbox you want to create.
5. In **Mailbox Name**, type a name for the mailbox.
6. In **Password**, type the password for the mailbox, and then retype it in **Confirm Password**.

**Important**

- If you are using Active Directory integrated authentication and have multiple POP3 e-mail domains, when creating mailboxes, you should consider whether the mailbox you are creating has the same name as an existing mailbox in another POP3 e-mail domain.

  Each mailbox is associated with an Active Directory user account, which has both a user logon name and pre-Windows 2000 user logon name. The user logon name is the name of the Active Directory user account and the POP3 mailbox. Usually, the pre-Windows 2000 user logon name is the same as the user logon name. If you try to create a mailbox and user account when an existing user account has the same pre-Windows 2000 user logon name, a naming conflict will occur.

  Active Directory does not support multiple accounts with the same pre-Windows 2000 user logon name. If there is a naming conflict, the mailbox name and e-mail address will not be affected, but the account's pre-Windows 2000 logon name will be modified to prevent any naming conflict with the existing account. If you are using Secure Password Authentication and a naming conflict has occurred, the pre-Windows 2000 logon name must be used for e-mail client authentication.

  To determine the modified logon name, go to the **Active Directory Users and Computers** snap-in located in the **Administrative Tools** menu. Go to the **Users** folder, right-click the user account and then click **Properties**. Click the **Account** tab and the modified account name will be displayed in **User logon name (pre-Windows 2000)**. You must note the pre-Windows 2000 logon name and provide it to the user. For more information about configuring e-mail

clients to use Secure Password Authentication, see Related Topics.

- Depending on your authentication method, certain characters are not allowed in mailbox names. Periods are allowed, but they must not be either leading or trailing. The following example illustrates the correct use of a period:

  john.someone@example.com

  The following table lists the characters that are prohibited for each type of authentication.

| Authentication method | Prohibited characters |
|---|---|
| Active Directory integrated authentication | @ ( ) / \ [ ] : ; , " < > * = \| ? + |
| Local Windows accounts authentication | @ ( ) / \ [ ] : ; , " < > * = \| ? + |
| Encrypted password file authentication | @ ( ) / \ [ ] : ; , " < > * = \| ? |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.

- The **Create associated user for this mailbox** option is only available if the server is using Active Directory integrated authentication or local Windows accounts authentication.

- The maximum length for a mailbox name is 20 characters for local Windows accounts authentication and 64 characters for encrypted password file authentication or Active Directory integrated authentication. The minimum length is 1 character.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.

   - If you are using Active Directory integrated authentication or local Windows accounts authentication, type:

     **winpop add***username@domain_name* [**/createuser:***new_user's_password*]

   - If you are using encrypted password file authentication, type:

     **winpop add***username@domain_name***/createuser:***new_user's_password*

| Value | Description |
|---|---|
| **winpop add** | Creates a mailbox in the domain. |
| *username@domain_name* | Specifies the user name for the corresponding mailbox and domain name in which to create the mailbox. |

| | |
|---|---|
| **/createuser** | Creates a corresponding user in the Active Directory or local computer.<br><br>If you are using Active Directory integrated or local Windows accounts authentication, you must use this command when creating a mailbox, unless a user account exists already with the same as the mailbox you want to create.<br><br>If you are using encrypted password file authentication, **/createuser** is still required even though a user account is not created. If you have configured your server with this option as a default, this parameter will have no effect. |
| *new_user's_password* | Specifies the password that will be used to access the mailbox. |

**Important**

- If you are using Active Directory integrated authentication and have multiple POP3 e-mail domains, when creating mailboxes, you should consider whether the mailbox you are creating has the same name as an existing mailbox in another POP3 e-mail domain.

  Each mailbox is associated with an Active Directory user account, which has both a user logon name and pre-Windows 2000 user logon name. The user logon name is the name of the Active Directory user account and the POP3 mailbox. Usually, the pre-Windows 2000 user logon name is the same as the user logon name. If you try to create a mailbox and user account when an existing user account has the same pre-Windows 2000 user logon name, a naming conflict will occur.

  Active Directory does not support multiple accounts with the same pre-Windows 2000 user logon name. If there is a naming conflict, the mailbox name and e-mail address will not be affected, but the account's pre-Windows 2000 logon name will be modified to prevent any naming conflict with the existing account. If you are using Secure Password Authentication and a naming conflict has occurred, the pre-Windows 2000 logon name must be used for e-mail client authentication.

  To determine the modified logon name, go to the **Active Directory Users and Computers** snap-in located in the **Administrative Tools** menu. Go to the **Users** folder, right-click the user account and then click **Properties**. Click the **Account** tab and the modified account name will be displayed in **User logon name (pre-Windows 2000)**. You must note the pre-Windows 2000 logon name and provide it to the user. For more information about configuring e-mail clients to use Secure Password Authentication, see Related Topics.

- Depending on your authentication method, certain characters are not allowed in mailbox names. Periods are allowed, but they cannot be either leading or trailing. The following example illustrates the correct use of a period:

  john.someone@example.com

  The following table lists the characters that are prohibited for each type of authentication.

| **Authentication method** | **Prohibited characters** |
|---|---|
| Active Directory integrated authentication | @ ( ) / \ [ ] : ; , " < > * = \| ? + |
| Local Windows accounts authentication | @ ( ) / \ [ ] : ; , " < > * = \| ? + |
| Encrypted password file authentication | @ ( ) / \ [ ] : ; , " < > * = \| ? |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- The maximum length for a mailbox name is 20 characters for local Windows accounts authentication and 64 characters for encrypted password file authentication or Active Directory integrated authentication. The minimum length is 1 character.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage Mailboxes
Configure Outlook Express for Secure Password Authentication

# Delete a mailbox

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To delete a mailbox

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open POP3 service.
2. In the console tree, click the domain containing the mailbox you want to delete.

   **Where?**

   - POP3 Service/*computer_name*/*domain_name*

3. In the details pane, right-click the mailbox you want to delete, and then click **Delete**.
4. If the mailbox has an associated Active Directory or local Windows accounts user, a dialog box appears. To delete the user account from the Active Directory or local computer Security Accounts Manager (SAM), select **Also delete the user account associated with this mailbox**.

### Caution

- Deleting a mailbox deletes the mail store directory for the mailbox and all e-mail that is stored in that directory.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **winpop delete** *username@domain_name* **/deleteuser**

| Value | Description |
|---|---|

| | |
|---|---|
| **winpop delete** | Deletes a mailbox from a domain.<br><br>If there is a corresponding user in the Active Directory or local computer Security Accounts Manager (SAM), a prompt appears, confirming the deletion of the user entry in addition to the mailbox. |
| *username@domain_name* | Specifies the mailbox name at the corresponding domain, for example, someone@example.com. |
| **/deleteuser** | Deletes the corresponding user from the Active Directory or SAM. |

**Caution**

- Deleting a mailbox deletes the mail store directory for the mailbox and all e-mail that is stored in that directory.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage Mailboxes

# List mailboxes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To list mailboxes**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, click the appropriate domain. The list of mailboxes appears in the details pane.

   **Where?**

   - POP3 Service/*computer_name*/*domain_name*

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **winpop list***domain_name*

| Value | Description |
|---|---|
| **winpop list** | Lists all mailboxes in the domain. |
| *domain_name* | Specifies the name of the domain for which you want a mailbox list, such as: example.com. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage Mailboxes

# Lock a mailbox

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To lock a mailbox**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. To display the list of mailboxes, in the console tree, click the appropriate domain.

   **Where?**

   - POP3 Service/*computer_name*/*domain_name*

3. In the details pane, right-click the mailbox you want to lock, and then click **Lock**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- When a mailbox is locked, it continues to receive incoming e-mail, which is delivered to the mail store. The user is unable to connect to the server and retrieve e-mail, however.
- While a mailbox is locked, the only restriction is that the user cannot connect to the server. An administrator can still perform all administrative tasks, such as deleting a mailbox or changing the password for a mailbox.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **winpop lock***username@domain_name*

| Value | Description |
|---|---|
| **winpop lock** | Locks the specified mailbox account to prevent the user from retrieving e-mail. |

| | |
|---|---|
| *username@domain_name* | The mailbox account that you want to lock; for example, someone@example.com |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- When a mailbox is locked, it continues to receive incoming e-mail, which is delivered to the mail store. The user is unable to connect to the server and retrieve e-mail, however.

- While a mailbox is locked, the only restriction is that the user cannot retrieve e-mail from the mail server. An administrator can still perform all administrative tasks, such as deleting a mailbox or changing the password for a mailbox.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage Mailboxes

# Unlock a mailbox

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To unlock a mailbox**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, click the appropriate domain to display the list of mailboxes.

   **Where?**

   - POP3 Service/*computer_name*/*domain_name*

3. In the details pane, right-click the mailbox you want to unlock, and then click **Unlock**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- If an individual mailbox is locked because it is part of a locked domain, you cannot unlock it.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **winpop unlock***username@domain_name*

| Value | Description |
|---|---|
| **winpop unlock** | Unlocks the specified mailbox account if it has been previously locked. |
| *username@domain_name* | Specifies the mailbox account that you want to unlock. |

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group

might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- If an individual mailbox is locked because it is part of a locked domain, you cannot unlock it.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.
- To view the complete syntax for this command, at a command prompt, type:

  **winpop help unlock**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage Mailboxes

# Change a password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a password**

1. Open Command Prompt.

2. Type:

   **winpop changepwd***username@domain_namenew_password*

| Value | Description |
|---|---|
| **winpop changepwd** | Changes the password of a mailbox account. |
| *username@domain_name* | Specifies the mailbox account for the password you want to change. |
| *new_password* | Specifies the new password for the mailbox account you want to change |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Depending on your authentication method there are other ways you can change a user's password.

  If you are using Active Directory integrated authentication you can use the Active Directory Users and Computers snap-in to change a user's password.

  If you are using local Windows accounts authentication you can use the Computer Management snap-in to change a user's password.

  For more information, see Related Topics.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage Mailboxes
Strong passwords

Reset the password for a local user account
Reset a user password

# View mailbox statistics

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view mailbox statistics**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open POP3 service.
2. In the console tree, click the appropriate domain.

   **Where?**

   - POP3 Service/*computer_name*/*domain_name*

   Mailbox statistics appear in the details pane.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open the POP3 service snap-in, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **POP3 Service**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **winpop stat***domain_name*

| Value | Description |
|---|---|
| **winpop stat** | Displays statistics for the domain, including the number of messages in account, the lock/unlock status of the mailbox, size of account, and whether the user is currently connected. |
| *domain_name* | Specifies the domain that you want to view statistics for, for example, *example.com*. |

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Manage Mailboxes

# Create a quota file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a quota file**

1. Open Command Prompt.

2. Type:

   **winpop createquotafile***username@domain* [**/user:***username*]

| Value | Description |
|---|---|
| **winpop createquotafile** | Specifies to create a quota file. |
| *username@domain* | Specifies the user for which a quota file should be created. |
| **/user:** *username* | Specifies to use an existing user account's quota file as a reference to create the quota file. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.

- If you are using Active Directory integrated authentication or local Windows accounts authentication, a quota file is created by default when a mailbox is created.

- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring disk quotas for the POP3 service
Winpop

# Migrate an encrypted password file user account to an Active Directory user account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To migrate an encrypted password file user account to an Active Directory user account**

1. Open Command Prompt.
2. Type:

   **winpop migrateToAD***UserName@DomainName*

| Value | Description |
|---|---|
| **winpop migrateToAD** | Migrates an encrypted password file user account to an Active Directory user account. The encrypted password file user account's password becomes the Active Directory user account's password. |
| *UserName@DomainName* | Specifies the existing encrypted password file user account. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to perform this procedure.
- The user's mailbox is not migrated. For the new Active Directory user account to be able to send and receive e-mail, you must change the authentication method to Active Directory integrated authentication and create a new mailbox for the Active Directory user account.
- To view the complete syntax for this command, at a command prompt, type:

  **winpop help**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Encrypted password file authentication
Active Directory integrated authentication

# Configure the Simple Mail Transfer Protocol Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure the Simple Mail Transfer Protocol service**

- Start or stop the Simple Mail Transfer Protocol service
- Configure the Simple Mail Transfer Protocol service for e-mail relaying
- Configure e-mail clients for e-mail relaying

# Start or stop the Simple Mail Transfer Protocol service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start or stop the Simple Mail Transfer Protocol service**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Services.
2. In the details pane, right-click **Simple Mail Transfer Protocol (SMTP)**, and then click the appropriate service state.

### Note

- To open Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Services**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **net** {**start**|**stop**} **smtpsvc**

| Value | Description |
|---|---|
| **net** {**start**\|**stop**} | Starts or stops a service. |
| **smtpsvc** | Specifies the Simple Mail Transfer Protocol (SMTP) service. |

### Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command Prompt**.
- To view the complete syntax for this command, at a command prompt, type:

  **net** {**start**|**stop**} **/?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Simple Mail Transfer Protocol (SMTP)
Configure the Simple Mail Transfer Protocol Service

# Configure the Simple Mail Transfer Protocol service for e-mail relaying

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the Simple Mail Transfer Protocol service for e-mail relaying**

1. To open Internet Information Services (IIS), click **Start**, click **Run**, type **inetmgr**, and then click **OK**. If the Internet Information Services (IIS) snap-in does not appear, verify that you have installed IIS.

2. In the console tree, right-click **Default SMTP Virtual Server**, and then click **Properties**.

   **Where?**

   - Internet Information Services/*computer_name*/Default SMTP Virtual Server

3. Click **Access**, click **Authentication**, select **Integrated Windows Authentication**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open IIS Manager, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Internet Information Services (IIS) Manager**.

- This procedure is necessary only if you are using Active Directory integrated authentication or local Windows accounts authentication, and you want to enable e-mail relaying. If you are using encrypted password file authentication, see the topic on Simple Mail Transfer Protocol (SMTP) in Related Topics for information about e-mail relaying.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure the Simple Mail Transfer Protocol Service
Simple Mail Transfer Protocol (SMTP)

# Configure e-mail clients for e-mail relaying

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure e-mail clients for e-mail relaying**

1. Click **Start**, point to **All Programs**, and then click **Outlook Express**.
2. Click **Tools**, and then click **Accounts**.
3. Click the **Mail** tab, click your POP3 service mail server, and then click **Properties**.
4. In **Outgoing Mail Server**, click **My server requires authentication**, click **Settings**, and then select **Log on using**.
5. In **Account name**, type the name of your POP3 mailbox without the domain. For example, if your mailbox is someone@example.com, you would type:

   **someone**

6. In **Password**, type the password for your POP3 mailbox, and then click **OK**.

## Notes

- For this procedure, you must use Active Directory integrated or local Windows accounts authentication. If you are using encrypted password file authentication and want to configure e-mail relaying, see the topic on Simple Mail Transfer Protocol (SMTP) in Related Topics.
- This procedure uses Outlook Express as an example. If you are not using Outlook Express, check the documentation for your e-mail client for information about configuring e-mail clients to authenticate to the Simple Mail Transfer Protocol (SMTP) service.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Simple Mail Transfer Protocol (SMTP)
Configure the Simple Mail Transfer Protocol Service

# Configure E-mail Clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure e-mail clients**

- Configure Outlook Express for e-mail services
- Configure Outlook Express for Secure Password Authentication

# Configure Outlook Express for e-mail services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Outlook Express for e-mail services**

1. Click **Start**, point to **All Programs**, and then click **Outlook Express**.
2. On the **Tools** menu, click **Accounts**.
3. Click **Add**, and then click **Mail**.
4. Follow the instructions in the **Internet Connection Wizard** until you reach the **Internet Mail Logon** page.
5. In **Account name**, type your POP3 service user name, followed by the domain, for example, someone@example.com.
6. In **Password**, type the password associated with your POP3 service e-mail account.
7. If your mail server is configured to require SPA, select **Log on using Secure Password Authentication (SPA)**, and then complete the wizard.

### Important

- You must perform this procedure on e-mail clients that will use the POP3 service.

### Notes

- If you are using an e-mail client other than Outlook Express, see your product documentation for information about how to configure your e-mail client to use the user name and domain name.
- Be sure to update your e-mail client and client operating system with the latest service pack available.
- For additional information about configuring Outlook Express for Secure Password Authentication (SPA), see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Configure E-mail Clients
Configure Outlook Express for Secure Password Authentication

# Configure Outlook Express for Secure Password Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Outlook Express for Secure Password Authentication**

1. Click **Start**, point to **All Programs**, and then click **Outlook Express**.
2. On the **Tools** menu, click **Accounts**.
3. In **Internet Accounts**, click the **Mail** tab, click the name of your POP3 e-mail account, and then click **Properties**.
4. Click the **Servers** tab, and then click **Log on using Secure Password Authentication**.
5. In **Account name**, type your POP3 service user name not including the domain. For example, if your e-mail address is someone@example.com, you would type:

   **someone**

   and then click **OK**. If a naming conflict occurred when the mailbox was created, the user name is the pre-Windows 2000 logon name. For more information about the pre-Windows 2000 logon name, see Notes.

## Notes

- Secure Password Authentication (SPA) only supports Active Directory integrated authentication and local Windows accounts authentication.
- Be sure to update your e-mail client software and client operating system with the latest service pack available.
- If you are using an e-mail client other than Outlook Express, check your product documentation for information about how to configure your e-mail client to use Secure Password Authentication (SPA) (also known as NTLM Authentication).
- If a dialog box prompts you for your credentials after you configure Outlook Express for SPA, enter your credentials, as described in the following table:

| Value | Description |
|---|---|
| **User Name** | Your user name, not including the domain. For example, if your e-mail address is someone@example.com, you would type:<br>**someone**<br>If a naming conflict occurred when the mailbox was created, the user name is the pre-Windows 2000 logon name. |
| **Password** | Your password. |
| **Domain** | For Active Directory integrated authentication, this is the network basic input/output system (NetBIOS) name of the domain. For local Windows accounts authentication, this is the name of the local computer. |

- If you are using Active Directory integrated authentication and Secure Password Authentication and a naming conflict occurred when the mailbox was created, the pre-Windows 2000 logon name must be used for e-mail client authentication. To determine the modified logon name, go to the **Active Directory Users and Computers** snap-in in **Administrative Tools**. Go to the **Users** folder, right-click the user account, and then click **Properties**. Click the **Account** tab and the modified account name will appear in **User logon name (pre-Windows 2000)**. You must note the pre-Windows 2000 logon name and provide it to the user.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
POP3 service administration tools
Configure E-mail Clients

# E-mail Services Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

The following topics provide background information about the POP3 service:

- POP3 service overview
- Understanding the POP3 Service
- Using the POP3 Service

# POP3 service overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**POP3 service overview**

The POP3 service is an e-mail service that retrieves e-mail. Administrators can use the POP3 service to store and manage e-mail accounts on the mail server.

When the POP3 service is installed on the mail server, users can connect to the mail server and retrieve e-mail to their local computer using an e-mail client that supports the POP3 protocol (such as Microsoft Outlook). The POP3 service is used with the Simple Mail Transfer Protocol (SMTP) service, which sends outgoing e-mail.

**E-mail transfer and retrieval overview**




The diagram illustrates how e-mail is transferred from sender to recipient, and how it is then retrieved to the recipient's local computer.

The sender's client computer connects to the Internet through an Internet Service Provider (ISP). Using an e-mail client, the sender sends an e-mail. The e-mail is picked up according to the SMTP protocol and is transferred to the sender's ISP, which then routes the e-mail out onto the Internet.

The e-mail moves through the Internet and is relayed through a varying number of intermediate servers on its way to the recipient. When the e-mail reaches the recipient's ISP, it is dropped into the recipient's mailbox.

When the recipient's computer connects to the recipient's ISP, the e-mail is then transferred from the ISP to the recipient's e-mail client on the local computer, according to the POP3 protocol. The POP3 service is the mechanism that allows users to retrieve their e-mail from the mail server.

# Understanding the POP3 Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding the POP3 service**
The POP3 service is an e-mail service that retrieves e-mail. Administrators can use the POP3 service to store and manage e-mail accounts on the mail server.

When the POP3 service is installed on the mail server, users can connect to the mail server and retrieve and download e-mail to their local computer. They do so by using an e-mail client that supports the POP3 protocol (such as Microsoft Outlook). The POP3 service is used with the Simple Mail Transfer Protocol (SMTP) service, which sends outgoing e-mail.

- E-mail concepts
- POP3 protocol
- Simple Mail Transfer Protocol (SMTP)
- POP3 Service Authentication Methods
- Mail store

# E-mail concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**E-mail concepts**

**Components of the POP3 e-mail system**

The three components of the POP3 e-mail system are the POP3 e-mail client, the Simple Mail Transfer Protocol (SMTP) service, and the POP3 service. These are each described in the following table.

| Component | Description |
|---|---|
| POP3 e-mail client | The POP3 e-mail client software that is used to read, compose, and manage e-mail.<br><br>The POP3 e-mail client retrieves e-mail from the mail server and transfers it to the user's local computer so it can then be managed by the user. An example of an e-mail client that supports the POP3 protocol is Microsoft Outlook Express. |
| SMTP service | The e-mail transfer system that routes e-mail from the sender to the recipient using the SMTP protocol.<br><br>The POP3 service uses the SMTP service as the e-mail transfer system. E-mail is composed by the user on a POP3 e-mail client. Then, after the user connects to the mail server using an Internet or network connection, the SMTP service picks up and transfers the e-mail across the Internet to the recipient's mail server. |
| POP3 service | The e-mail retrieval system that downloads e-mail from the mail server to the user's local computer using the POP3 protocol.<br><br>The POP3 protocol controls the connection between a user's POP3 e-mail client and the server where e-mail is stored. |

Administrators manage the POP3 service on three organizational levels: mail servers, e-mail domains, and mailboxes. These are described in the following table.

| Category | Description |
|---|---|
| Mail server | The computer where the POP3 service is installed. Users connect to it to retrieve their e-mail. |
| E-mail domains | The e-mail domain must be a registered domain name and it must match the Mail eXchanger (MX) record created by your Internet Service Provider (ISP). |
| Mailboxes | Each mailbox corresponds to a user who is a member of the e-mail domain, such as someone@example.com. The user's mailbox corresponds to a directory in the mail store in which the user's e-mail is stored until it is retrieved. |

**A POP3 e-mail system**




In the diagram, e-mail is sent to someone@example.com, as illustrated by step 1. The e-mail is picked up by the SMTP service and is sent out on the Internet (step 2).

The e-mail domain, example.com, is then resolved to a mail server on the Internet, mailserver1.example.com (step 3). Mailserver1.example.com is a mail server running the POP3 service that receives incoming e-mail for the e-mail domain, example.com.

In step 4, the e-mail for someone@example.com is received by mailserver1.example.com and is moved into the mail store directory (step 5), where e-mail intended for someone@example.com is stored.

In step 6, the user, someone, checks for e-mail by connecting to the mail server running the POP3 service. The POP3 protocol transmits the user and password authentication credentials for the user, someone. The POP3 service verifies these credentials, and then either accepts or denies the connection.

If the connection is accepted, all of the e-mail for the user, someone, that is being stored in the mail store is downloaded from the mail server to someone's local computer (step 7). The mail is usually then deleted from the mail store.

**Network security**

The POP3 and SMTP protocols are not encrypted. If someone gains access to the network on which the server running the POP3 service is running, they could potentially read users' e-mail.

To increase network security, implement Internet Protocol security (IPSec).

IPSec is a framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks through the use of cryptographic security services. For more information about IPSec, see Internet Protocol Security (IPSec).

# POP3 protocol

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**POP3 protocol**

The Post Office Protocol 3 (POP3) is a standard protocol for retrieving e-mail. The POP3 protocol controls the connection between a POP3 e-mail client and a server where e-mail is stored. The POP3 service uses the POP3 protocol for retrieving e-mail from a mail server to a POP3 e-mail client.

The POP3 protocol has three process states for handling the connection between the mail server and the POP3 e-mail client: the authentication state, the transaction state, and the update state.

During the authentication state, the POP3 e-mail client that is connected to the server must be authenticated before users can retrieve their e-mail. If the user name and password that are provided by the e-mail client match those on the server, the user is authenticated and proceeds to the transaction state. If not, the user receives an error message and is not allowed to connect to retrieve e-mail.

To prevent any damage to the mail store after the client has been authenticated, the POP3 service locks the user's mailbox. New e-mail that is delivered to the mailbox after the user has been authenticated (and the mailbox has been locked) is not available to download until after the connection has been terminated. Also, only one client can connect to a mailbox at a time; additional connection requests to the mailbox are rejected.

During the transaction state, the client sends POP3 commands and the server receives and responds to them according to the POP3 protocol. Any client requests the server receives that do not comply with POP3 protocol are ignored, and an error message is sent back.

The update state closes the connection between the client and server. It is the last command the client transmits.

After the connection is closed, the mail store is updated to reflect the changes made while the user was connected to the mail server. For example, after the user successfully retrieves e-mail, the retrieved e-mail is marked for deletion and then deleted from the mail store, unless the user's e-mail client is configured to do otherwise.

# Simple Mail Transfer Protocol (SMTP)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Simple Mail Transfer Protocol (SMTP)**
Simple Mail Transfer Protocol (SMTP) controls the way e-mail is transported and delivered across the Internet to the destination server. SMTP receives and sends e-mail between servers. The SMTP service is installed by default with the POP3 service to provide complete e-mail services.

The SMTP service is automatically installed on the computer where the POP3 service is installed to allow users to send outgoing e-mail. When you create a domain using the POP3 service, the domain is also added to the SMTP service to allow mailboxes in that domain to send outgoing e-mail. The SMTP service on the mail server receives incoming mail and transfers the e-mail to the mail store.

**E-mail relay**
E-mail relay occurs when users who are not members of the e-mail domain use a mail server with SMTP to send e-mail. SMTP mail servers that are not configured to prevent open relay are often abused by those who want to send large volumes of unsolicited commercial e-mail. The Microsoft SMTP service is configured by default to prevent e-mail relaying.

If you want to enable e-mail relay, however, you have the following two options, depending on which authentication method you are using:

- If you are using Active Directory integrated authentication or local Windows accounts authentication, you can configure the mail server to require authentication before accepting outgoing e-mail.

  For more information on configuring the mail server, see Configure the Simple Mail Transfer Protocol service for e-mail relaying. If you configure the SMTP service to authenticate users trying to send outgoing mail, you must also configure the users' e-mail clients. For more information about configuring e-mail clients to allow relaying, see Configure e-mail clients for e-mail relaying.

- If you are using encrypted password file and want to configure e-mail relaying, you must configure the mail server to allow relay based on either Internet Protocol (IP) address or e-mail domain name.

**Note**

- When an e-mail is undeliverable, the Simple Mail Transfer Protocol (SMTP) service returns it to the sender with a non-delivery report (NDR). If the NDR cannot be delivered to the sender, a copy of the message is put in the Badmail directory. To prevent the possibility of the operating system running out of disk space (if the Badmail directory becomes too large), you must move the SMTP Badmail directory to a volume other than the one on which the operating system is installed.

# POP3 Service Authentication Methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**POP3 service authentication methods**

The POP3 service provides three different authentication methods for verifying users connecting to the mail server.

You must choose an authentication method before creating any e-mail domains on the mail server. You can change the authentication method only if there are no existing e-mail domains on the mail server.

The following table describes when to use each authentication method:

| Authentication method | When to use this authentication method |
|---|---|
| Local Windows accounts authentication | • Your mail server is not a member of an Active Directory domain;<br>-AND-<br>• You want to store user accounts on the server on which the POP3 service is installed. |
| Active Directory integrated authentication | • Your mail server is a domain controller or a member of an Active Directory domain. |
| Encrypted password file authentication | • Your mail server is not using Active Directory;<br>-OR-<br>• You do not want to have user accounts for the POP3 service on the local computer. |

For more information about the POP3 service authentication methods, see the following topics:

- Local Windows accounts authentication
- Active Directory integrated authentication
- Encrypted password file authentication

**Notes**

- If the computer on which the POP3 service is running is either a member of an Active Directory domain or a domain controller, the default authentication method is Active Directory integrated authentication. Otherwise, the default is local Windows accounts authentication.
- If the computer where the POP3 service is running is a domain controller, the available authentication methods are Active Directory integrated authentication and encrypted password file authentication. If the computer where the POP3 service is running is a member server in an Active Directory domain, all three authentication methods are available. Otherwise, the available authentication methods are local Windows accounts authentication and encrypted password file authentication.

# Local Windows accounts authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local Windows accounts authentication**
If you are not using Active Directory, but you want to create user accounts on the local computer where the POP3 service is installed, you can use local Windows accounts authentication to link the POP3 service into the local system user accounts.

Use local Windows accounts authentication if:

- Your mail server is not a member of an Active Directory domain.

  -and-

- You want to have user accounts on the same local computer as the POP3 service.

Local Windows accounts authentication integrates the POP3 service into the local computer's Security Accounts Manager (SAM). With Security Accounts Manager, users who have user accounts on the local computer can use the same user name and password to be authenticated by both the POP3 service and the local computer.

You can use local Windows accounts authentication to support multiple domains on the server, but user names across domains must be unique. For example, you cannot have a user named someone@example.com and a user named someone@northwindtraders.com.

**POP3 Users Security Accounts Manager group**
If a mailbox is created with a corresponding user account, the user account is added to the "POP3 Users" local group. Members of the "POP3 Users" group cannot log on locally to the server even though they have a user account on the server.

The restriction on local logon is enforced through the computer's local security policy. It increases server security by restricting local logon rights only to authorized users. Being unable to log on locally to the server does not affect the users' ability to use the POP3 service.

For more information on local groups and modifying local group membership, see Local Users and Groups.

For more information on local security policy, see Local Security Policy.

**E-mail client authentication**
Local Windows accounts authentication supports both plaintext and Secure Password Authentication (SPA) e-mail client authentication.

Because plaintext transmits the user's credentials in an unsecured, unencrypted format, the use of plaintext is not recommended. SPA, however, does require that e-mail clients transmit both the user name and password using secure authentication; it is therefore recommended over plaintext authentication. For more information, see Configure the mail server to require Secure Password Authentication.

### Note

- To install Active Directory on the computer where the POP3 service is installed, see the section on upgrading from local Windows accounts to Active Directory integrated authentication in Active Directory integrated authentication. Installing Active Directory on a member server without following the recommended procedure might cause the POP3 service to work incorrectly.

# Active Directory integrated authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory integrated authentication**

You can use Active Directory integrated authentication to incorporate the POP3 service into your existing Active Directory domain. If mailboxes are created that correspond to existing Active Directory user accounts, users can use their existing Active Directory domain user name and password to send and receive e-mail. For more information about Active Directory, see Active Directory.

Use Active Directory integrated authentication if:

- The server on which you want to install POP3 service is a member of an Active Directory domain or is an Active Directory domain controller.

You can use Active Directory integrated authentication to support multiple POP3 e-mail domains, and you can use the same user name across different POP3 e-mail domains. For example, you can have a user named someone@example.com and a user named someone@northwindtraders.com.

If you use Active Directory integrated authentication and have multiple POP3 e-mail domains, when you create a mailbox, be sure to consider whether the new mailbox has the same name as an existing mailbox in another POP3 e-mail domain. Each mailbox is associated with an Active Directory user account that has both a user logon name and pre-Windows 2000 user logon name. The user logon name is the name of the Active Directory user account and the POP3 mailbox.

Usually, the pre-Windows 2000 user logon name is the same as the user logon name. If, however, you try to create a mailbox and user account when there is an existing user account with the same pre-Windows 2000 user logon name, a naming conflict will occur. Active Directory does not support multiple accounts with the same pre-Windows 2000 user logon name.

If a naming conflict occurs, the mailbox name and e-mail address are not affected, but the account's pre-Windows 2000 logon name is modified to prevent any naming conflict with the existing account.

If you use Secure Password Authentication and a naming conflict has occurred, the pre-Windows 2000 logon name must be used for e-mail client authentication. To determine the modified logon name, go to the **Active Directory Users and Computers** snap-in located in the **Administrative Tools** menu. Go to the **Users** folder, right-click the user account, and then click **Properties**. Click the **Account** tab. The modified account name will appear in **User logon name (pre-Windows 2000)**.

You must note the pre-Windows 2000 logon name and provide it to the user. For more information about configuring e-mail clients to use Secure Password Authentication, see Configure Outlook Express for Secure Password Authentication.

**E-mail client authentication**

Active Directory integrated authentication supports both plaintext and Secure Password Authentication (SPA) e-mail client authentication. Because plaintext transmits the user's credentials in an unsecured, unencrypted format, however, the use of plaintext authentication is not recommended. SPA does require e-mail clients to transmit both the user name and password using secure authentication; it is therefore recommended over plaintext authentication. For more information, see Configure the mail server to require Secure Password Authentication.

**Upgrading from local Windows accounts to Active Directory integrated authentication**

If you are upgrading your mail server to a domain controller, and you were using local Windows accounts authentication, you must complete the steps described in the following table before you can configure the POP3 service for Active Directory integrated authentication.

| Step | Reference |
|---|---|
| Delete all existing e-mail domains in the POP3 service. | For more information on deleting e-mail domains, see Delete a domain. |
| Change the authentication method to Active Directory integrated. | For more information on changing the authentication method, see Set the authentication method. |

| | |
|---|---|
| Recreate the domains and associated mailboxes. | • For more information on creating e-mail domains, see Create a domain<br>• For more information on creating mailboxes, see Create a mailbox. |

**Notes**

- Installing Active Directory on a member server without following the recommended procedure can prevent the POP3 service from working correctly.
- If you are using Active Directory integrated authentication, you must log on to the Active Directory domain, not the local computer, to administer the POP3 service.

# Encrypted password file authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encrypted password file authentication**
Encrypted password file authentication is ideal for large scale deployments where there is not an existing Active Directory deployment, and there will potentially be too many accounts to manage easily from a single local computer.

Use encrypted password file authentication if:

- You are not using Active Directory.

  -Or-

- You do not want to create users on the local computer.

With encrypted password file authentication, you can use an identical user name across different domains, but you cannot assign the same user name to multiple mailboxes within a domain. For example, you cannot have two mailboxes named someone@example.com, but you can have someone@example.com and someone@northwindtraders.com.

Encrypted Password file authentication creates an encrypted file with the user's password, which is then stored in the user's mailbox directory on the server. During authentication, the password the user supplies is encrypted and compared to the encrypted file stored on the server. If the encrypted password matches the encrypted password stored on the server, the user is authenticated.

**Migrating encrypted password file user accounts to Active Directory user accounts**
You can migrate user accounts created under encrypted password file authentication to Active Directory user accounts using the winpop.exe command-line tool. For more information, see Migrate an encrypted password file user account to an Active Directory user account.

**E-mail client authentication**
Encrypted Password file authentication supports both plaintext and Authenticated POP (APOP) e-mail client authentication. Plaintext transmits the user's credentials in an unsecured, unencrypted format. Therefore, the use of plaintext authentication is not recommended.

Because APOP e-mail client authentication requires e-mail clients to transmit both the user name and password using secure authentication, it is recommended over plaintext authentication. For more information, see Configure the mail server to require Secure Password Authentication.

# Mail store

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Mail store**
The mail store is a directory or Universal Naming Convention (UNC) path where the POP3 service stores all e-mail until users retrieve it to their client computer.

The basic structure of the mail store, or mail root, is a directory on the local hard disk where all e-mail is stored.

When a domain is created, the POP3 service creates a corresponding directory in the directory that has been designated for the mail store. For each user with a mailbox in that domain, POP3 creates a directory in the domain directory. E-mail that a user receives is stored as an individual file within the user's directory until the user retrieves it using a POP3 e-mail client.

The following is an example of the path to an e-mail message in the mail store:

C:\inetpub\mailroot\mailbox\example.com\P3_someone.mbx\P347865.eml

where mailroot corresponds to the mail store directory, example.com to the domain directory, P3_somone.mbx to the directory for a mailbox named someone and P347865.eml to a single saved e-mail message.

The directory and file permissions for each directory in the mail store are identical. When you configure the mail store, the permissions are set so that only local or domain administrators and the local network service, which the POP3 service is configured to run under, are assigned permissions to the directories. No other user is assigned read/write permissions.

The mail store's functionality depends on having adequate hard disk space available. To ensure the mail store's functionality, you should develop a disk-space requirement estimate based on the number of users on the server, the volume of e-mail that they will receive, and the average size of the e-mail they will receive.

In addition, you can protect the server from situations where the mail store's disk usage might increase unexpectedly by implementing disk quotas. Disk quotas monitor and control disk space that is used on NTFS file system volumes. For more information, see Configuring disk quotas for the POP3 service.

**Notes**

- Because the mail store can potentially use large amounts of disk space, you should either set a disk quota limit on the volume of the mail store (to control its disk space usage) or set it to use a volume other than the one where the operating system is installed. This will prevent the possibility of the operating system running out of disk space if the mail store becomes too large. For more information, see Set the mail store. For more information on disk quotas, see Configuring disk quotas for the POP3 service.

- The mail store must be configured to use either a directory on the local hard disk or a UNC path; other storage options, such as mapped drives, are not supported.

- You cannot set the mail store to the root directory of the hard disk, for example C:\, or to a directory in which files are currently in use.

- If you restore the mail store from a backup or move it to a new location, you must reset the permissions on the mail store directory using the command-line procedure described in Set the mail store.

- If you transfer the mail store to a new directory, you must move the mail store directory to ensure the directory retains the correct ownership; copying the mail store will not work.

- Physical access to a server is a high security risk. To maintain a more secure environment, restrict physical access to the server where the mail store resides.

# Using the POP3 Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using POP3 service**

This section provides information on how to use the POP3 service:

- Contacting your Internet Service Provider
- POP3 service administration tools
- POP3 service commands overview
- Configuring disk quotas for the POP3 service
- Remote administration of POP3 service

# Contacting your Internet Service Provider

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Contacting your Internet Service Provider**

Before installing and configuring the POP3 service, you must contact your Internet Service Provider (ISP) and complete the steps described in the following table.

| Step | Description | Additional reference |
|---|---|---|
| Register an e-mail domain name. | The e-mail domain name must be a registered domain name, and it must match the MX record created by your ISP. If you do not have have an e-mail domain name, contact your ISP for assistance registering a domain name. | |
| Have your ISP create a Mail eXchanger (MX) record for your e-mail domain that matches the name of your mail server running the POP3 service. | An MX record provides e-mail routing to a mail exchanger host for e-mail sent to the domain name.<br><br>For example, in the domain example.com, the mail server might be mailserver1.example.com.<br><br>The ISP would create an MX record for example.com that points to mailserver1.example.com.<br><br>An e-mail that is sent to someone@example.com is routed to mailserver1.example.com. This is the server where the POP3 service is running and to which users connect to retrieve e-mail. | To change the name of your server to a title that reflects its role (for example, mailserver1), see Rename a computer. |
| Obtain information to configure your mail server for static addressing. | Your ISP must provide a separate static IP address for each mail server on which you intend to install the POP3 service.<br><br>To obtain a static Internet Protocol (IP) address, you must also obtain the subnet mask, default gateway, and preferred Domain Name System (DNS) server from your ISP. Each mail server must then be configured to use an individual static IP address. | For more information on how to configure your mail server with a static IP address, see Configure TCP/IP for static addressing. |

# POP3 service administration tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**POP3 service administration tools**
The POP3 service provides three methods for local administration and two methods for remote administration.

**Local administration**
You can administer the POP3 service on the local computer from the Microsoft Management Console (MMC) snap-in, the command line, or the Web Interface for Remote Administration. For more information, see Using Web Interface for Remote Administration.

**Remote administration**
You can administer the POP3 service on a remote computer from the Microsoft Management Console (MMC) snap-in or the Web Interface for Remote Administration. You can also set up Terminal Services on the remote computer, which would then make the same local computer administration options available. For more information, see Remote administration of POP3 service.

**Configuring and managing the POP3 service**
The following tables describe the administrative options you can use to configure and manage the POP3 service.

- General service settings
- Domain administration options
- Mailbox administration options
- Service state settings

**General service settings**
For information on how to configure these settings, see Set Up the POP3 Service and Change Service Settings.

| Service parameter | Description |
|---|---|
| Port | Specifies the port on which the POP3 service monitors for POP3 e-mail clients and responds to connection requests. Do not change the port from its default of 110 unless you have a specific need to do so. You must also inform all users if you are using a nonstandard port because most POP3 e-mail clients configure port 110 as the default for a POP3 connection. |
| Event logging level | Specifies how much information is stored by the POP3 service in Event Viewer. Logging information is useful for auditing and troubleshooting. The logging levels are:<br>• None - No logging.<br>• Low - Only critical events are logged.<br>• Medium - Critical and warning events are logged.<br>• High - Critical, warning, and informational events are logged. |

| | |
|---|---|
| Authentication | Specifies the authentication method used to verify users who connect to the mail server. The POP3 service supports three authentication methods:<br>• Local Windows accounts authentication<br>• Active Directory integrated authentication<br>• Encrypted password file authentication<br><br>For more information, see POP3 Service Authentication Methods. |
| Mailroot | Specifies the directory on the local hard disk where e-mail is stored. For more information, see Mail store. |
| Sockets | Specifies the number of sockets.<br><br>A socket receives the network request from the e-mail client and then handles the transaction. Each time a POP3 e-mail client connects to the mail server, a socket is used. The POP3 service has four configurable settings that control the sockets: maxsockets, minsockets, backlogsockets, and thresholdsockets.<br><br>For example, on a system where maxsockets has been set to 20, up to 20 simultaneous connections can be made. If minsockets has been set to 10, when the POP3 service starts, 10 sockets are created immediately. The minsocket value determines how many connection requests the service can respond to on startup, and, based upon the number of sockets to create on startup, how much memory is initially allocated to the service.<br><br>New sockets are created based upon the value of thresholdsockets. If thresholdsockets is 5, when the number of available sockets on the server is less than 5, 5 new sockets are created. This increases the total sockets on the system to 15. When the 5 new sockets are filled, 5 additional sockets are created.<br><br>This process continues until maxsockets is reached. If maxsockets is 20, when 20 sockets are created and filled, the value of backlogsockets determines how many subsequent connections are accepted, but not responded to. If backlogsockets is 5, the server will accept a total of 25 connections: the 20 connections that fill the available sockets and 5 additional connections that are queued until a socket is available. Connection attempts subsequent to maxsockets + backlogsockets are denied. |
| Threads | Specifies the number of threads.<br><br>A thread is a sequence of program code that is scheduled to run on the processor. A process can have several threads running simultaneously. The operating system coordinates this process by assigning each thread a scheduling priority. With the POP3 service, you can change the number of threads created per processor by configuring the thread count setting. The number of threads per processor multiplied by the number of processors equals the number of simultaneous requests the POP3 service can process.<br><br>The number of threads per processor can affect performance when traffic is high. Adding more threads can increase throughput, but adding too many threads will decrease server performance. The default settings are appropriate for most typical environments. For high-load environments, such as dedicated mail servers, the administrator should experiment with the number of threads for best results. |
| Quotas | Specifies disk quotas that are implemented.<br><br>You can implement disk quotas to control the amount of hard disk space allocated for each user's mailbox. Disk quotas are important because they ensure that the server is protected from situations where the mail store might suddenly and rapidly increase in size. For more information, see Configuring disk quotas for the POP3 service. |

**Domain administration options**

For information on how to set these options to manage a domain, see Manage E-mail Domains.

| **Domain administration** | **Description** |
|---|---|

| | |
|---|---|
| Lock/Unlock | Locks or unlocks a domain. Unlocking the domain allows all users in the domain to retrieve their e-mail; locking the domain prevents them from doing so.<br><br>When a domain is locked, incoming e-mail for the domain is still received and delivered to the appropriate user directory in the mail store, and outgoing mail is still sent. All users attempting to download mail from the server are denied a connection, however.<br><br>While a domain is locked, you can still perform administrative tasks such as deleting a domain. If a user is connected to the mail server when the domain is locked, the mailbox remains unlocked until the user disconnects, after which the mailbox is locked. |
| Stat | Displays mailbox statistics, including number of messages in a mailbox, lock/unlock status, size of directory, and whether a user is currently connected |
| List | Lists all mailboxes in the domain. |
| New domain | Creates a new domain and adds it to the existing POP3 directory structure. |
| Delete a domain | Deletes an existing domain. If you delete an existing domain, you also delete all mailboxes and stored e-mail for that domain from the mail store. |

**Mailbox administration options**

For information on how to set these options to manage mailboxes, see Manage Mailboxes.

| **Mailbox administration** | **Description** |
|---|---|
| New mailbox | Creates a new mailbox in an existing domain. To receive e-mail, every user in the domain must have a unique mailbox. Both Active Directory integrated authentication and local Windows accounts authentication support creating an associated user for each mailbox to link POP3 authentication with the existing domain or local computer. |
| Delete a mailbox | Deletes a mailbox and removes the mailbox from the domain. After a mailbox has been deleted, all e-mail sent to that mailbox is returned to the sender as nondeliverable receipt. When you delete a mailbox, you can specify that the corresponding user account also be deleted. |
| Lock or unlock a mailbox | Locks or unlocks a mailbox.<br><br>When the mailbox is unlocked, the user can connect to the server and retrieve e-mail. When it is locked, the mailbox continues to receive incoming e-mail, which is delivered to the mail store, but the user is not allowed to connect to the server and retrieve e-mail. The administrator can still perform administrative tasks, however, such as deleting a mailbox or changing a mailbox's password.<br><br>If a user is connected to the mail server when the mailbox is locked, the mailbox remains unlocked until the user disconnects, after which the mailbox is locked. |
| Change a mailbox password | Resets the mailbox password. If a user loses or forgets his or her password, you must reset the mailbox password to something new; it is impossible to reveal an existing password. |

**Service state settings**

The POP3 service is a Windows service with full support for Start, Stop, Pause, Resume, and Restart commands. For more information on how to change the POP3 service state, see Change the POP3 service state.

| **Service state** | **Description** |
|---|---|
| Start | Starts the POP3 service from a stopped or paused state and allows POP3 clients to connect to the server and retrieve e-mail. |
| Stop | Stops the POP3 service. All current connections are terminated. |
| Pause | Pauses the POP3 service so that no new connections are accepted. Existing connections are not affected. |

| | |
|---|---|
| Resume | Starts the POP3 service from a paused state to allow POP3 clients to connect to the server and retrieve e-mail. |
| Restart | Restarts the POP3 service. All current connections are terminated. |

# POP3 service commands overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**POP3 service commands overview**

You can use the POP3 service command application, Winpop.exe, to manage the POP service and modify parameter settings. With Winpop.exe, you can use all the POP3 service parameters available in the POP3 service Microsoft Management Console (MMC) snap-in, as well as additional administrative settings that are only configurable using Winpop.exe.

For more information, see:

- Winpop

# Configuring disk quotas for the POP3 service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring disk quotas for the POP3 service**

**Important**

- You can configure disk quotas only on NTFS file system partitions.

You can use disk quotas to control and limit the amount of disk space individual mailboxes on the mail server can use. This ensures that individual mailboxes, and the mail store in general, do not use excessive or unanticipated amounts of disk space and adversely affect the performance of the server where the POP3 service is running.

For example, if the mail server suddenly receives a large volume of unsolicited e-mail, the mail store expands rapidly and might use all of the available disk space on the hard disk. If you implement quotas, the mail store only expands to the quota limit that you specified. As a result, no more mail is accepted by the server, and the rest of the server still functions normally.

If you are using Active Directory integrated authentication or local Windows accounts authentication, the e-mail delivered to a POP3 service mailbox will have file ownership assigned to the mailbox user by default. A quota file is created in the mailbox directory that contains the security identifier (SID) of the user account associated with the mailbox. File ownership is then assigned to the user account that corresponds to the SID contained in the quota file. The SID is also used by the NTFS file system disk quota system to enforce the quota limits specified on the user account matching the SID. All e-mail transferred to the mailbox's mail store directory is marked with the SID contained in the quota file; this marks the e-mail so it can be monitored by the quota system.

For more information, see Disk Quotas.

**createquotafile command**

If you are using encrypted password file authentication, there is no valid user account for the mailbox that the quota system can use. You can use the **createquotafile /user** command, however, to manually associate a given mailbox with a valid user account that is configured to have a disk quota. This association is for disk quota purposes only and is separate from mailbox authentication. If you are using Active Directory integrated authentication or local Windows accounts authentication, a quota file is created by default when you create a mailbox.

For more information about the **createquotafile** command, see Create a quota file.

**Configuring domain disk quotas**

Although quotas are designed to be implemented on a per-mailbox basis, you can create a domain-wide disk quota.

To create a disk quota for a domain, you must create a new mailbox and user account with an associated quota. The quota file for the new mailbox functions as a template that you can then copy into all of the other mailbox directories in the domain to create a domain-wide quota policy.

The following table describes the steps for creating a domain disk quota.

| Step | Reference |
|---|---|
| Create a new mailbox and user account. When you create the new mailbox, you must also create an associated user account for the mailbox. This account will be used to create a domain-wide quota. Other accounts in the domain can be configured to reference this account and its associated quota. All accounts associated with this account will have their disk usage aggregated under a single quota limit, which will create a domain-wide quota limit. | Create a mailbox |
| Enable disk quotas for the partition on which the mail store is configured. | Enable disk quotas |
| Create a quota to be used as the domain quota and assign it to the domain quota account. | Add new quota entries |

| | |
|---|---|
| Do one of the following:<br>• Copy the quota file from the mail store directory of the domain account to the corresponding mail store directory of all the mailboxes in the domain.<br>• Or, using the **winpop createquotafile** command and the **/user:** switch, associate one or more accounts in the domain with the domain account and its quota. For more information on the **createquotafile** command, see Create a quota file. | N/A |

**Notes**

- When a mailbox quota is exceeded, the user is not notified. E-mail intended for the user is not accepted, and a Non-Delivery Report (NDR), a notice that the e-mail was not delivered to the recipient, is returned to the sender.
- Be sure that users configure their e-mail client to delete from the server any e-mail that has been successfully retrieved. If users leave successfully retrieved e-mail on the server, they can quickly exceed their quota. Users are likely to be unaware of the disk usage and impact of old e-mail stored on the server.
- You cannot set quota limits on the accounts of Administrators or members of the Administrators group.

# Remote administration of POP3 service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote administration**
There are three approaches you can use to administer the POP3 service remotely: remote Web administration, Terminal Services, and Microsoft Management Console.

- **Remote Web administration**

  You can use the Web to connect remotely to a server running the POP3 service. The POP3 service is integrated into the Web Interface for Remote Administration. You can use a Web browser to exert full administrative control of a computer over the Internet. For more information, see Using Web Interface for Remote Administration.

  **Note**

  - This content is not available in this preliminary release.
  - This content is not available in this preliminary release.

- **Remote administration using Terminal Services**

  You can use a Terminal Services client to connect remotely to a server that is running the POP3 service. With Terminal Services, you have full administrative control over the computer and the POP3 service. For more information, see Remote Administration using Terminal Services.

- **Remote administration using Microsoft Management Console (MMC)**

  Using the POP3 service MMC snap-in, you can connect to another computer that is running the POP3 service. For security reasons, to manage the POP3 service on a remote computer, you must be a member of the Administrators group on the remote computer, or you must have been delegated the appropriate authority. For more information about connecting to a remote server, see Connect to a remote server running the POP3 service.

**Note**

- To perform this procedure on a remote computer, you must be a member of the Administrators group on the remote computer, or you must have been delegated the appropriate authority.

# Troubleshooting E-mail services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Users are unable to connect to the POP3 server to retrieve or send e-mail
- I am using local Windows accounts authentication. The POP3 service will not start after I install Active Directory on the member server.
- Users are unable to access their mailboxes after the mail store directory is restored from a backup or moved to a new location, or event ID 1022 appears in the event log.
- I am using local Windows accounts authentication and a remote mail store. A security identifier (SID) appears instead of an account name as the owner of the files in the remote mail store.
- Deleting a user account from the Active Directory Users and Computers snap-in does not delete the user's corresponding POP3 mailbox.
- When using Secure Password Authentication and Active Directory integrated authentication, users are unable to connect.
- I am unable to perform administrative tasks when using Active Directory integrated authentication, or the following error message appears:"The specified domain either does not exist or could not be contacted."

**Users are unable to connect to the POP3 server to retrieve or send e-mail**
**Cause:** The user's e-mail client has not been configured correctly, or the user has incorrect logon credentials.

**Solution:** Review the procedures to configure e-mail clients.

- If you are not using Secure Password Authentication, see Configure Outlook Express for e-mail services.
- If you are using Secure Password Authentication, see Configure Outlook Express for Secure Password Authentication.

**I am using local Windows accounts authentication. The POP3 service will not start after I install Active Directory on the member server.**
**Cause:** When Active Directory is installed on a computer, the local Security Accounts Manager (SAM) is disabled. Local Windows accounts authentication depends upon SAM, however. Without SAM, the POP3 service will not start if it is configured to use local Windows accounts authentication.

**Solution:** After you install Active Directory on the server, configure the POP3 service to use Active Directory integrated authentication as described in the following table:

Step

Reference

Delete all existing e-mail domains in the POP3 service.

For more information on deleting e-mail domains, see Delete a domain.

Change the authentication method to Active Directory integrated.

For more information on changing the authentication method, see Set the authentication method.

Recreate the domains and associated mailboxes.

- For more information on creating e-mail domains, see Create a domain.

- For more information on creating mailboxes, see Create a mailbox.

**See also:** Active Directory integrated authentication.

**Users are unable to access their mailboxes after the mail store directory is restored from a backup or moved to a new location, or event ID 1022 appears in the event log.**
**Cause:** The discretionary access control list (DACL) that is set on the mail store directory and all subdirectories has been lost.

**Solution:** Reset the DACL on the mail store directory. The DACL will then be propagated to all subdirectories.

1. Open Command Prompt.

2. Type:

   **winpop set mailroot***path_to_mail_store*

   **Note**

   - The **set mailroot** command might take a while to complete if the mail store is large.

**See also:** Mail store.

**I am using local Windows accounts authentication and a remote mail store. A security identifier (SID) appears instead of an account name as the owner of the files in the remote mail store.**
**Cause:** The SID displayed corresponds to a user account on the computer on which the POP3 service is running. Because that account does not exist on the remote computer, there is no user account to be referenced, and the SID appears instead.

**Solution:** This is a known issue.

**Deleting a user account from the Active Directory Users and Computers snap-in does not delete the user's corresponding POP3 mailbox.**
**Cause:** The Active Directory Users and Computers snap-in is designed to function this way.

**Solution:** To delete both the user account and the POP3 mailbox, use the POP3 service snap-in. For more information, see Delete a mailbox.

Or, if you delete the user account from the Active Directory Users and Computers snap-in, you can delete the POP3 mailbox manually.

To delete a mailbox manually, go to the directory that is configured to be the mail store. Delete the directory that corresponds to the POP3 mailbox.

For more information on how to determine the directory that corresponds to the POP3 mailbox, see Mail store.

**When using Secure Password Authentication and Active Directory integrated authentication, users are unable to connect.**
**Cause:** The user's logon credentials are incorrect.

**Solution:** Verify that the user is using the correct logon name and password. If a naming conflict occurred when the user's mailbox was created, the user must use the pre-Windows 2000 logon name as his or her user name when connecting to the server running the POP3 service. For more information about naming conflicts, see Active Directory integrated authentication.

- To verify the user's logon name, go to the **Active Directory Users and Computers** snap-in located in the **Administrative Tools** menu. Go to the **Users** folder, right-click the user account and then click **Properties**. Click the **Account** tab. The user name that must be used will appear in **User logon name (pre-Windows 2000)**.

- If the user's password is incorrect, and the user does not remember the correct password, reset the password and provide the user with a new password. For more information, see Change a password.

**I am unable to perform administrative tasks when using Active Directory integrated authentication, or the following error message appears:"The specified domain either does not exist or could not be contacted."**
**Cause:** You are not logged on to the Active Directory domain. If you are using Active Directory integrated authentication, you

must log on to the Active Directory domain, not the local computer, to administer the POP3 service.

**Solution:** Log on to the Active Directory domain using your administrator credentials.

# Background Intelligent Transfer Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Background Intelligent Transfer Service**
Background Intelligent Transfer Service (BITS) helps you transfer large amounts of data without degrading network performance. It does this by transferring data in small chunks, utilizing unused bandwidth as it becomes available, and reassembling the data at the destination. BITS is supported on Microsoft® Windows Server 2003 family operating systems and on Microsoft® Windows 2000 Servers.

BITS downloads and uploads both file names, as Universal Resource Locators (URLs), and data. Any transfers that are interrupted by network disconnections or by computer shutdowns and restarts resume when network connectivity is re-established.

BITS Server Extensions is part of Internet Information Services (IIS) and has the following components:

- BITS management console snap-in, for configuring and managing BITS by using Microsoft Management Console (MMC).
- BITS server extension ISAPI, which is used for BITS data transfer.

If you want client computers to be able to use BITS to upload files to a server, you must configure an IIS server with BITS Server Extensions installed and activated.

For more information, see the following topics and Web sites:

- To install BITS Server Extensions, see Install BITS Server Extensions.
- To configure BITS Server Extensions, see Configure BITS Server Extensions.
- To review BITS-specific security information, see Background Intelligent Transfer Service Concepts.
- To review the BITS PSDK, see Background Intelligent Transfer Service.

# Background Intelligent Transfer Service Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This topic describes how to use BITS notifications, and also provides security-related information for BITS.

**Using BITS notifications**

BITS Server Extensions provides a way to notify Web applications of the arrival of data from a client computer to the server, and for that server to generate a reply to the client computer. You can configure this mechanism for each virtual directory on the server by using controls in the **Allow notifications** section of the **BITS Server Extensions** tab.

When **Allow notifications** is checked, two controls become available for configuration: **Notification type** and **Notification URL**.

**Notification type** requires that you select either **Send the file name** or **Send the data**.

- **Send the file name** passes only the name of the file and tells the server the location the location of that uploaded file, as well as where to record a reply. Your application is responsible for opening the request (uploaded) data file and for creating any response data.

- **Send the data** passes all of the data and replies directly between the server and application and specifies the locations for the data and replies.

When BITS Server Extensions receives all of the uploaded data from the application on the client computer, it generates a notification reply in the **Notification URL** that you specify. For **Send the file name**, BITS writes the response to a URL. For **Send the data**, BITS writes the data and response to a temporary file for the client computer to download.

BITS uses IIS security mechanisms. BITS notifications are generated solely in the security context specified for the client system. This context is set in **Authentication Methods**, which is accessed from the **Directory Security** tab of the **Properties** page for the virtual directory.

**Securing virtual directories**

IIS allows a user to set Run script and Execute permissions for a virtual directory separately from Write permissions. Customarily, administrators disallow writing to the virtual directory to prevent clients from uploading executable scripts or programs to a directory and then executing them.

However, when a virtual directory has been enabled for BITS uploads, clients can write to this directory by means of BITS uploads, even though write permissions on the virtual directory have been disallowed. To protect the server, BITS Server Extensions disables all Run script and Execute permissions on a virtual directory that has been upload-enabled for BITS. To keep the virtual directory secure, BITS Server Extensions does not respond to any client requests for uploading data until permissions are disabled. If these permissions are enabled at any time, BITS Server Extensions denies all upload requests until the permissions are disabled again. Whenever BITS Server Extensions denies a request, it writes an IIS log entry.

For more information about securing virtual directories, see the online Help.

**Using BITS with SSL**

BITS Server Extensions is installed with a default configuration that does not secure BITS traffic during data replication. If you want to use BITS to transfer sensitive data, you should do so over a secure connection, such as Secure Sockets Layer (SSL).

For general information about SSL, see the online Help.

# Install BITS Server Extensions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install BITS Server Extensions**

1. Click **Start** and then click **Control Panel**.

2. Select **Add or Remove programs** and then click **Add or Remove Windows Components**.

3. From the **Windows Components Wizard**, select **Application Server**, and then click **Details**.

4. In the **Application Server** dialog box, check **Internet Information Services (IIS)**, and then click **Details**.

5. In the **Internet Information Services (IIS)** dialog box, check **Background Intelligent Transfer Service (BITS) Server Extensions**.

   By default, both the BITS server extension ISAPI and the BITS management console snap-in are installed. To deselect either one, click **Details** in the Internet Information Services (IIS) dialog box to open the Background Intelligent Transfer Service (BITS) server extension dialog box.

6. Close the **IIS** and **Application Server** dialog boxes to return to the **Windows Component Wizard**. Then click **Next** to begin the installation of BITS Server Extensions as part of Application Server. The wizard completes the installation.

## Notes

- To install BITS Server Extensions, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate credentials. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **runas** command to perform this procedure. For Windows 2000 Server family products, see the online Help documentation or go to Windows 2000 product Documentation and select Help for the appropriate product.

- You might need your Windows Server 2003 family CD available when you install BITS Server Extensions, in case the installation process requires additional files.

- You can install BITS on Windows 2000 Server by installing the Platform SDK Redistributables for Windows 2000. Windows 2000 Server uses IIS 5. For help using IIS 5, see Windows 2000 Server Documentation for Internet Information Services.

- To install BITS on a client computer, open the SUPPORT\TOOLS directory on the Windows XP Professional CD, disk image, or network directory, and double-click **SUPTOOLS.EXE**. The BITS tool file is Bitsadmin.exe.

- Internet Information Services is installed on the server you are using when the console opens. If you want to connect to another computer, add that computer in the **Connect to Computer** dialog box by right-clicking **Internet Information Services** in the console tree and selecting **Connect**.

# Configure BITS Server Extensions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure BITS Server Extensions**

1. Click **Start**, click **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
2. Click to expand the nodes under the server where BITS Server Extensions is installed, and then click to expand **Web Sites**.
3. Right-click the virtual directory for which you want to configure BITS Server Extensions, and then select **Properties**.
4. In the **Properties** dialog box, select the **BITS Server Extensions** tab.

The following sections explain the settings on the **BITS Server Extensions** tab.

- Allow clients to transfer data to this virtual directory
- Custom settings
- Enable server farm support
- Allow notifications
- Schedule Cleanup
- Run Cleanup Now

**Allow clients to transfer data to this virtual directory**
Enables uploads from BITS clients into the virtual directory. If this check box is not checked, BITS clients attempting an upload are denied access.

### Use default settings

Precludes configuring BITS Server Extensions for the virtual directory. The remaining settings cannot be configured unless you select **Customize settings**.

### Note

- Configuration settings can be inherited. If the directory is a subdirectory of another virtual directory and the parent virtual directory defines a custom configuration, the configuration of the parent virtual directory is used instead of the default settings.

### Customize settings

Allows you to configure this property page according to the settings specified in **Custom settings**.

**Custom settings**
Permits selection of maximum file size and allows you to specify a time limit after which incomplete jobs can be deleted.

### Maximum file size

Defines the maximum size of a file that can be uploaded in a single BITS job. If this setting is not enabled, clients can upload files of any size.

### Delete incomplete jobs after

If enabled, permits incomplete upload files to be deleted by the cleanup process if no progress has been made within the

specified time period. Deletions are performed only when the cleanup process is scheduled to run, or when you manually run a cleanup. See **Schedule Cleanup** for information about scheduling the cleanup process. See **Run Cleanup Now** for information about manually running a cleanup.

#### Enable server farm support

If enabled, instructs the client to try to reconnect to the server of choice instead of the original upload URL used by the client. Server farms can be configured two ways: by using a shared volume for the upload virtual directory, or by using server local storage for the uploaded files. If you want to use local storage, this option enables the client computer to reconnect to the same server computer in the server farm. Uploading to server farms and using local storage are not recommended without enabling and configuring **Enable server farm support**.

### Reconnect to IP address

Specifies the public IP address to use for reconnecting to this server. A server usually has two IP addresses. One points outward and is reachable by client systems. The other belongs to the inner (protected) network. You should specify the outward-pointing IP address here.

### Notes

- You can specify a DNS name here only if each computer in your Web farm has an independent DNS name that can be resolved by client computers.
- **Enable server farm support** must be checked in order to configure this setting.

### Use original IP address after

Specifies the amount of time during which a client computer must reconnect to the server specified in **Reconnect to IP address**. If the client computer does not connect to the server during this time, the client computer must attempt to connect to the original DNS name specified in the upload URL.

#### Allow notifications

Enables notifications for this virtual directory. Both Notification type and Notification URL settings are required if this setting is enabled. For additional information about the notification settings on the **BITS Server Extensions** tab on the **Properties** page for the virtual directory, see **Using BITS notifications** in Related Topics.

### Notification type

You can send either the file name or the data. If **Send the file name** is selected, the server passes the path for the completed upload file to the URL specified by the Notification URL setting. The server returns a file to the client computer as a response. If **Send the data** is selected, the server passes the entire upload file to the URL specified by the Notification URL as a POST. The response is sent to the client computer.

### Notification URL

Specifies the URL to pass the completed upload file to via a POST request if the Notification type is either Send the file contents or Send the file name. The URL can be absolute or self-relative. If the URL is self-relative, the original request URL is used as the base.

#### Schedule Cleanup

Configures the frequency with which BITS Server Extensions scans the virtual directory for incomplete upload files. If an incomplete file is found, it is checked against the time-out value set in **Delete incomplete jobs after**. If the time-out period has elapsed since progress was made on this file, the upload is cancelled and any partial file contents are deleted. This setting is not affected by IIS property inheritance and applies only to the virtual directory.

#### Run Cleanup Now

Immediately starts a cleanup by scanning the virtual directory for incomplete upload files. If an incomplete file is found, it is checked against the time-out value set in **Delete incomplete jobs after**. If the time-out period has elapsed since progress was made on this file, the upload is cancelled and any partial file contents are deleted. To remove all incomplete files from this virtual directory, set the time-out value to 0 before running the cleanup.

### Notes

To install BITS Server Extensions, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate credentials. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using the **runas** command to perform this

procedure. For Windows 2000 Server family products, see the online Help documentation or go to Windows 2000 Product Documentation and select Help for the appropriate product.

- After you configure settings on the **BITS Server Extensions** tab and click **Apply** or **OK**, a security message appears, indicating that BITS will shut down all Run script and Execute permissions for this virtual directory. When you click **OK**, the **Run Cleanup Now** button becomes enabled. For more information about this security message, see **Setting permissions on virtual directories** in Related Topics.

- BITS Server Extensions must be in the Allowed state in order to function. To determine the state, expand **Web Service Extension** in the IIS console tree. Locate the extension state in the right panel. Click **Allowed** if the extension is in the Prohibited state.

- For information about virtual directories, see Help for Internet Information Services.

**See Also**
**Other Resources**
Using BITS notifications request/response headers
Setting permissions on virtual directories

# Disks and Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disks and data**
This section covers:

- Backing Up and Recovering Data
- Managing Data Storage
- Managing Files and Folders
- Managing Disks and Volumes

# Backing Up and Recovering Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backing up and recovering data**
This section covers:

- Backup
- Shadow Copies for Shared Folders
- Accessing Previous Versions of Shared Files and Folders
- Disaster Recovery

# Backup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backing up and restoring data**

Backup is a utility that you can use to back up and restore user data.

- Before backing up data, see Checklists for Backing Up and Restoring Data.
- For tips on backing up and restoring data, see Best practices for Backup.
- For help with specific tasks, see Backup How To ....
- For general background information, see Backup Concepts.
- For problem-solving information, see Backup Troubleshooting.
- For more information about the recovery options that are available, see Repair overview.

# Checklists for Backing Up and Restoring Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**


- Checklist: Backing up data
- Checklist: Restoring data

# Checklist: Backing up data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Backing up data**

| Step | Reference |
|---|---|
| Review key concepts and prepare to back up data. | Backup Concepts |
| Verify that you are logged on as a member of the Backup Operators group or Administrators group. | Identify members of a local group |
| Verify that you can connect to all shared folders on other computers that need to be backed up. Note: you cannot back up another computer's System State or another domain controller. | |
| If you are using an external storage device, verify that it is compatible with products in the Windows Server 2003 family. Confirm that it is connected directly to the computer that is performing the backup and that the computer is turned on. | Check the hardware compatibility information in Support resources. |
| If you are backing up to a tape, insert the required tape into the tape drive. If you are backing up to a disk drive, verify that there is enough available space. | |
| If you are backing up an Encrypted File System (EFS), first back up the designated recovery agent's EFS private key (the PKCS #12 file) to ensure the successful recovery of encrypted data in case of a disaster such as a system failure. | Export a certificate with the private key |
| To back up files manually, use the Backup or Restore Wizard or click the **Backup** tab in the Backup UI to select files to backup. | Back up files to a file or a tape |
| To back up files automatically, use the schedule feature in backup. | Schedule a backup |

# Checklist: Restoring data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Restoring data**

| Step | Reference |
|---|---|
| Verify that you are logged on as a member of the Backup Operators group or Administrators group. | Identify members of a local group |
| Verify that you can connect to all shared folders on other computers that need to be restored. | |
| Insert the required tape(s) into the tape drive or library. | |
| If you are restoring the System State to a running domain controller, start the server in Directory Services Restore Mode. | Follow the steps in Start the computer in Safe Mode, and then click **Directory Services Restore Mode**. |
| To restore files manually, click the **Restore and Manage Media** tab to select files to restore or use the Backup or Restore Wizard. | Restore files from a file or a tape |
| If you are restoring an Encrypted File System (EFS) on a system where the private key for the encrypted data is somehow inaccessible (for example, on a computer that is not part of a network) or is corrupted or lost, import the designated recovery agent's EFS private key (the PKCS #12 file). | Import a certificate |

# Best practices for Backup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for Backup**

- **Develop backup and restore strategies and test them.**

  With a good plan, you can quickly recover your data if it is lost.

- **Train appropriate personnel.**

  In minimum-security and medium-security networks, assign backup rights to one user and restore rights to a different user. Train personnel with restore rights to perform all of the restore tasks if the administrator is unavailable.

  In a high-security network, make sure that only administrators can restore files.

- **Back up all data on the system and boot volumes and the** System State**.**

  Back up the data on all volumes and the System State data at the same time. This precaution prepares you for the unlikely event of a disk failure.

- **Create an Automated System Recovery backup set.**

  Always create an Automated System Recovery (ASR) backup set when the operating system changes, for example, whenever you install new hardware and drivers or apply a service pack. With an ASR backup set, you can more easily recover from a system failure. Also, backup all data volumes at the same time; ASR protects only the system, so data volumes must be backed up separately.

- **Create a backup log.**

  Always choose to create a backup log for each backup and print them for reference. Keep a book of logs to make it easier to locate specific files. The backup log is helpful when restoring data; you can print it or read it from any text editor. Also, if the tape containing the backup set catalog is corrupted, the printed log can help you locate a file.

- **Retain copies.**

  Keep at least three copies of the media. Keep at least one copy off-site in a properly-controlled environment.

- **Perform trial restorations.**

  Perform a trial restoration periodically to verify that your files were properly backed up. A trial restoration can uncover hardware problems that do not show up when you verify software.

- **Secure devices and media.**

  Secure both the storage device and the backup media. It is possible for someone to access the data from a stolen medium by restoring the data to another server for which they are an administrator.

- **Back up your server cluster effectively.**

  To be able to effectively restore your server cluster in the event of application data or quorum loss, or individual node or complete cluster failure, follow these steps when preparing backups:

  1. Perform an Automated System Recovery (ASR) backup on each node in the cluster.

2. Back up the cluster disks from each node.

3. Back up each individual application (for example, MS Exchange or MS SQL Server) running on the nodes.

For more information, see Backing up and restoring server clusters. For more information on backing up applications in a cluster, see the documentation for that application.

# Backup How To ...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Back Up Data
- Restore Data
- Set Options
- Maintain Media
- Perform System Repair and Recovery

# Back Up Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Back up data**

- Back up files to a file or a tape
- Back up System State data
- Schedule a backup

# Back up files to a file or a tape

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up files to a file or a tape**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use this wizard or go to the next step to work in **Advanced Mode**.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Backup** tab, and then on the **Job** menu, click **New**.

4. Select the files and folders you want to back up by clicking the box to the left of a file or folder.

5. In **Backup destination**, do one of the following:

   - Choose **File** if you want to back up files and folders to a file. This is selected by default.
   - Choose a tape device if you want to back up files and folders to a tape.

6. In **Backup media or file name**, do one of the following:

   - If you are backing up files and folders to a file, type a path and file name for the backup (.bkf) file, or click the **Browse** button to find a file.
   - If you are backing up files and folders to a tape, choose the tape you want to use.

7. Select any backup options you want, such as the backup type and the log file type, by clicking the **Tools** menu, and then clicking **Options**. When you have finished selecting backup options, click **OK**.

8. Click **Start Backup**, and then make any changes to the **Backup Job Information** dialog box.

9. If you want to set advanced backup options such as data verification or hardware compression, click **Advanced**. When you have finished setting advanced backup options, click **OK**.

10. Click **Start Backup** to start the backup operation.

**Caution**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

**Notes**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- If the Backup or Restore Wizard does not start by default, you can still use it to back up files by clicking **Wizard Mode** on the **Welcome** tab.

- To back up and restore Microsoft SQL Server database files, it is recommended that you use the built-in backup and restore utilities that SQL has. For more information, see the Microsoft SQL Server documentation.

- Some tape drives may not support hardware compression.

- System State data contains most elements of a system's configuration, but it might not include all of the information that you require to recover your system from a failure. Therefore, it is recommended that you back up all boot and system volumes, including the System State, when you back up your system.

- You can only back up the System State data on a local computer. You cannot back up the System State data on a remote computer.

- Backup does not support the use of backup media such as CD-RW (compact disc rewritable), CD-R (compact disc recordable), and DVD-R (digital video disc recordable). To save backups to this type of media, back up to a file and copy that file to the CD. You can use Backup to restore from a CD.

- Backup files usually have the extension .bkf. You can use any extension you like, but it is highly recommended that you use .bkf, which has file associations that will ensure that the backup file is recognizable.

- Backup operators and administrators can back up and restore encrypted files and folders without decrypting the files or folders.

- If you have Windows Media Services running on your computer, and you want to back up the files associated with these services, see "Running Backup with Windows Media Services" in the Windows Media Services online documentation. You must follow the procedures outlined in the Windows Media Services online documentation before you can back up or restore files associated with Windows Media Services.

**Using a command line**

1. Open Command Prompt.

2. To backup to a file or tape, type:

   **ntbackup backup***@bks file name***/J** {"*backup job name*"} **/F** {"*file name*"} **/T** {"*tape name*"} **/P** {"*pool name*"} **/G** {"*guid name*"} **/N** {"*media name*"} **/A**

| Value | Description |
|---|---|
| *@bks file name* | Specifies the name of the backup selection file (.bks file) to be used for this backup operation. The @ character must precede the name of the backup selection file. A backup selection file contains information on the files and folders you have selected for backup. You have to create the file using the graphical user interface (GUI) version of Backup. Alternatively, you could supply the path to the drive or file to backup, for example, D:\. |
| **/J** {"*backup job name*"} | Specifies the job name to be used in the backup report. The job name usually describes the files and folders you are backing up in the current backup job. |

| | |
|---|---|
| **/F** {"*file name*"} | Specifies the logical disk path and file name of the backup file. You must not use the following switches with this switch: **/P /G /T**. |
| **/T** {"*tape name*"} | Overwrites or appends to this tape. |
| **/P** {"*pool name*"} | Specifies the media pool from which you want to use media. This is usually a subpool of the Backup media pool, such as 4mm DDS. If you select this value, you must not use the following switches: **/A /G /F /T** |
| **/G** {"*guid name*"} | Overwrites or appends to this tape. Do not use this switch in conjunction with **/P**. |
| **/N** {"*media name*"} | Specifies the new tape name. You must not use **/A** with this switch. |
| **/A** | Performs an append operation. Either **/G** or **/T** must be used in conjunction with this switch. Do not use this switch in conjunction with **/P**. |

For information on additional file and tape backup options, see the Ntbackup.

**Examples:**
To backup drive D: to file C:\backup.bkf, with job name "Backup Job 1" type the following at the command line:

**ntbackup.exe backup D:\ /j "Backup Job 1" /f "C:\Backup.bkf"**

For a tape backup, to create a backup job named "Backup Job 1" that backs up all the files and folders listed in the backup selection file c:\backup.bks to a tape drive, type:

**ntbackup backup @C:\backup.bks /n "Media created 11/5/2001 at 1:25 PM" /d "Set created 11/5/2001 at 1:25 PM" /j "Backup Job 1" /p "4mm DDS"**

**Important**

- If you don't specify options at the command prompt, they will default to those specified in the Backup application.
- Before using the command line to perform backup operations, it is recommended that you use the user interface to create the initial backup selection file, settings and a scheduled task. You can then use that scheduled task to view the command syntax. Copy the information from the scheduled task, make any changes that are necessary to create a script or command line syntax, then delete the scheduled task if it is not needed otherwise.
- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.
- To back up and restore Microsoft SQL Server database files, it is recommended that you use the built-in backup and restore utilities that SQL has. For more information, see the Microsoft SQL Server documentation.
- Some tape drives may not support hardware compression.

**Notes**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at a command prompt, type:

  **ntbackup /?**

- If you do not specify the other Backup options, **ntbackup** will use the default values of the Backup program for the backup type, verification setting, logging level, hardware compression, and any other settings.

- System State data contains most elements of a system's configuration, but it might not include all of the information that you require to recover your system from a failure. Therefore, it is recommended that you back up all boot and system volumes, including the System State, when you back up your system.

- You can only back up the System State data on a local computer. You cannot back up the System State data on a remote computer.

- Backup files usually have the extension .bkf. You can use any extension that you want, but it is highly recommended that you use .bkf, which has file associations that will ensure that the backup file is recognizable.

- Backup operators and administrators can back up and restore encrypted files and folders without decrypting the files or folders.

- If you have Windows Media Services running on your computer, and you want to back up the files associated with these services, see "Running Backup with Windows Media Services" in the Windows Media Services online documentation. You must follow the procedures outlined in the Windows Media Services online documentation before you can back up or restore files associated with Windows Media Services.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Network Monitor
Backup
Checklists for Backing Up and Restoring Data
Permissions and user rights required to back up and restore
Using batch files to back up data
Identify members of a local group

# Back up System State data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up System State data**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use this wizard or go to the next step to work in **Advanced Mode**.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Backup** tab, then click the box next to **System State** and any other items you would like to backup. It is recommended that you backup all boot and system volumes along with the System State.

## Notes

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.
- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
- If the Backup or Restore Wizard does not start by default, you can still use it to back up the System State data by clicking **Wizard Mode** on the **Welcome** tab.
- System State data contains most elements of a system's configuration, but it might not include all of the information that you require to recover your system from a failure. Therefore, it is recommended that you back up all boot and system volumes, including the System State, when you back up your system.
- You can only back up the System State data on a local computer. You cannot back up the System State data on a remote computer.

**Using a command line**

1. Open Command Prompt.

2. To backup the System State data, type:

   **ntbackup backup systemstate**

| Value | Description |
|---|---|
| **systemstate** | Specifies that you want to back up the System State data. When you select this option, the backup type will be forced to copy. |

For information on additional backup options, see the Ntbackup.

For example, to create a backup job named "Backup Job 1" that backs up the System State data to the file C:\backup.bkf, type:

**ntbackup backup systemstate /J "Backup Job 1" /F "C:\backup.bkf"**

All other options will default to those specified in the Backup program.

### Notes

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- System State data contains most elements of a system's configuration, but it might not include all of the information that you require to recover your system from a failure. Therefore, it is recommended that you back up all boot and system volumes, including the System State, when you back up your system.

- To view the complete syntax for this command, at a command prompt, type:

  **ntbackup /?**

- If you do not specify the other Backup options, **ntbackup** will use the default values of the Backup program for backup type, verification setting, logging level, hardware compression, and any other settings.

- You can only back up the System State data on a local computer. You cannot back up the System State data on a remote computer.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System State data
Using batch files to back up data
Identify members of a local group

# Schedule a backup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To schedule a backup**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use this wizard or go to the next step to work in **Advanced Mode**.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Backup** tab, and then on the **Job** menu, click **New**.

4. Select the files and folders you want to back up by clicking the box to the left of a file or folder.

5. Select **File** or a tape device in **Backup destination**.

6. In **Backup media or file name**, type a path and file name for the backup file, or select a tape.

7. Select any backup options you want, such as the backup type and the log file type, by clicking the **Tools** menu, and then clicking **Options**. When you have finished selecting backup options, click **OK**.

8. On the **Job** menu select **Save Selections** to save your selections as a backup job file (.bks).

9. Click **Start Backup** and make any changes you want to the **Backup Job Information** dialog box.

10. If you want to set advanced backup options such as data verification or hardware compression, click **Advanced**. When you have finished selecting advanced backup options, click **OK**.

11. Click **Schedule** in the **Backup Job Information** dialog box.

12. In the **Set Account Information** dialog box, enter the user name and password that you want the scheduled backup to run under.

13. In the **Scheduled Job Options** dialog box, in **Job name**, type a name for the scheduled backup job, and then click **Properties** on the **Schedule data** tab to set the date, time, and frequency parameters for the scheduled backup. When you have finished, click **OK**, and then click **OK** again.

**Caution**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

**Notes**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- If you are scheduling a tape backup, you may have to use Removable Storage to make sure that your tape is available in the Backup media pool. For more information on Removable Storage, see Related Topics.

- You must have the Task Scheduler service running before you can schedule a backup. To do this, open the command prompt window and type **net start schedule**. You can also use Services in the Computer Management administrative tool to start, stop, and view the status of services.

- System State data contains most elements of a system's configuration, but it might not include all of the information that you require to recover your system from a failure. Therefore, it is recommended that you back up all boot and system volumes, including the System State, when you back up your system.

- If you want to back up your System State data, you must select **System State** in **Click to select the check box for any drive, folder, or file that you want to back up**.

- You can only back up the System State data on a local computer. You cannot back up the System State data on a remote computer.

- You can change the settings of a scheduled backup job after you schedule it by clicking the **Schedule Jobs** tab and then clicking the backup icon that is displayed on the calendar.

- In the **Scheduled Job Options** dialog box, you can delete a scheduled backup from the Task Scheduler by clicking **Delete**.

- Backup operators and administrators can back up and restore encrypted files and folders without decrypting the files or folders.

- If you have Windows Media Services running on your computer, and you want to back up the files associated with these services, see "Running Backup with Windows Media Services" in the Windows Media Services online documentation. You must follow the procedures outlined in the Windows Media Services online documentation before you can back up or restore files associated with Windows Media Services.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Permissions and user rights required to back up and restore
Identify members of a local group
Schedule Tasks

# Restore Data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Restore data**

- Restore files from a file or a tape
- Restore System State data

# Restore files from a file or a tape

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore files from a file or a tape**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Restore and Manage Media** tab, select the files and folders you want to restore by clicking the box to the left of a file or folder.

4. In **Restore files to**, do one of the following:

   - Click **Original location** if you want the backed up files and folders to be restored to the folder or folders they were in when they were backed up. Skip to step 6.

   - Click **Alternate location** if you want the backed up files and folders to be restored to a folder that you designate. This option will preserve the folder structure of the backed up data; all folders and subfolders will appear in the alternate folder you designate.

   - Click **Single folder** if you want the backed up files and folders to be restored to a folder that you designate. This option will not preserve the folder structure of the backed up data; the files will appear only in the folder that you designate.

5. If you selected **Alternate location** or **Single folder**, type a path for the folder under **Alternate location**, or click the **Browse** button to find the folder.

6. On the **Tools** menu, click **Options**, click the **Restore** tab, and then do one of the following:

   - Click **Do not replace the file on my computer** if you do not want the restore operation to copy over files that are already on your hard disk.

   - Click **Replace the file on disk only if the file on disk is older** if you want the restore operation to replace older files on your disk with newer files from your backup.

   - Click **Always replace the file on my computer** if you want the restore operation to replace files on your disk regardless of whether the backup files are newer or older.

   - Click **OK** to accept the restore options you have set.

7. Click **Start Restore**.

8. If you want to change any of the advanced restore options, such as restoring security settings and junction point data, click **Advanced**. When you are done setting advanced restore options click **OK**.

9. Click **OK** to start the restore operation.

**Caution**

- Choosing **Always replace the file on my computer** might cause you to lose data if a file you are working on is contained in the backup.

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- You can also use the Restore Wizard to restore files by clicking **Restore Wizard** on the **Tools** menu.

- To back up and restore Microsoft SQL Server database files, it is recommended that you use the built-in backup and restore utilities in SQL. For more information, see the Microsoft SQL Server documentation.

- To restore the System State data on a domain controller, you must first start your computer in Directory Services Restore Mode. This will allow you to restore the SYSVOL directory and the Active Directory directory service database.

- You can only restore the System State data on a local computer. You cannot restore the System State data on a remote computer.

- Backup operators and administrators can back up and restore encrypted files and folders without decrypting the files or folders.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set advanced restore options
Checklist: Restoring data
Update a catalog
Identify members of a local group

# Restore System State data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore System State data**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Restore and Manage Media** tab, click the box next to **System State**. This will restore the System State data along with any other data you have selected for the current restore operation.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- You can also use the Restore Wizard to restore the System State data by clicking **Restore Wizard** on the **Tools** menu.

- If you are restoring the System State data to a domain controller, you must choose whether you want to perform a primary restore, an authoritative restore or a nonauthoritative restore. The default method of restoring the System State data to a domain controller is nonauthoritative (normal). In this mode, any component of the System State that is replicated to with another domain controller, such as the Active Directory directory service or the File Replication service (including the SYSVOL directory), will be brought up to date by replication after you restore the data. For example, if the last backup was performed a week ago, and the System State is restored using the default restore method (nonauthoritative), any changes made subsequent to the backup operation will be replicated from the domain controllers.

  In some cases, you may not want to replicate the changes that have been made subsequent to the last backup operation. In other words, there may be instances where you want all replicas to have the same state as the backed up data. To achieve this state, you must perform an authoritative restore.

  For example, you have to perform an authoritative restore if you inadvertently delete users, groups, or organizational units from the Active Directory directory service, and you want to restore the system so that the deleted objects are recovered and replicated. To do this, you need to run the Ntdsutil utility after you have restored the data but before you restart the domain controller. This utility lets you mark objects as authoritative, which will ensure that any replicated or distributed data that you have restored is properly replicated or distributed throughout your organization. The Ntdsutil command line utility can be run from the command prompt. Help for the Ntdsutil utility can also be found at the command prompt by typing **ntdsutil /?**. For more information , see Related Topics.

- To restore the System State data on a domain controller, you must first start your computer in a special startup option called Directory Services Restore Mode. This will allow you to restore the SYSVOL directory and Active Directory directory service database. To access Directory Services Restore Mode, press F8 during startup and select it from the list of startup options. For more information on startup options, see Related Topics.

- You can only restore the System State data on a local computer. You cannot restore the System State data on a remote computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
System State data
Back up System State data
Startup options
Update a catalog
Permissions and user rights required to back up and restore
Ntdsutil

# Set Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set options**

- Set backup options
- Set the backup type
- Exclude file types from a backup
- Set advanced backup options
- Set restore options
- Set advanced restore options

# Set backup options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set backup options**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Tools** menu, then click **Options**.

4. On the **General** tab, set the options you want.

**Notes**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

| Item | Description |
|---|---|
| Compute selection information before backup and restore operations | Estimates the number of files and bytes that will be backed up or restored during the current backup or restore operation. This information is calculated and displayed before the backup or restore begins. |
| Use the catalogs on the media to speed up building restore catalogs on disk | Specifies that you want to use the on-media catalog to build the on-disk catalog for restore selections. This is the fastest way to build an on-disk catalog. However, if you want to restore data from several tapes, and the tape with the on-media catalog is missing or you want to restore data from media that is damaged, do not select this option. Backup will then scan your entire backup set (or as much of it as you have) and build an on-disk catalog. This could take several hours if your backup set is very large. |
| Verify data after the backup completes | Checks for disk errors that may have occurred while files were being copied to tape. This might substantially increase the time it takes to perform a backup. It is recommended that you only verify backups of data files; system backups are difficult to verify due to the large number of changes that happen to system files on a continual basis. Some data files that were in use during the backup might also cause verification errors, but you can generally ignore these errors. If there are a large number of verification errors, there may be a problem with the media or the file you are using to back up data. If this occurs, use different media or designate another file and run the backup operation again. |
| Back up the contents of mounted drives | Backs up the data that is on a mounted drive. If you select this option, and you back up a mounted drive, the data that is on the mounted drive will be backed up. If you do not select this option, and you back up a mounted drive, only the path information for the mounted drive will be backed up. |

| | |
|---|---|
| Show alert message when I start Backup and Removable Storage is not running | Displays a dialog box when you start Backup and Removable Storage is not running. Backup will then start Removable Storage automatically. If you primarily back up data to a file, and you save the file to a floppy disk, a hard disk, or any type of removable disk, do not check this box. If you primarily back up data to a tape, or other media that is managed by Removable Storage, check this box. |
| Show alert message when I start Backup and there is recognizable media available | Displays a dialog box when you start Backup and there is new media available. If you primarily back up data to a file, and you save the file to a floppy disk, a hard disk, or any type of removable disk, do not check this box. If you primarily back up data to a tape or other media that is managed by Removable Storage, check this box. For more information on media pools and Removable Storage, see Related Topics. |
| Show alert message when new media is inserted | Displays a dialog box when new media is detected by Removable Storage. If you primarily back up data to a file and you save the file to a floppy disk, a hard disk, or any type of removable disk, do not check this box. If you primarily back up data to a tape, or other media that is managed by Removable Storage, check this box. |
| Always allow use of recognizable media without prompting | Automatically moves new media that is detected by Removable Storage to the Backup media pool. If you usually back up data to a file, and you save the file to a floppy disk, a hard disk, or any type of removable disk, do not check this box. If you want all new media to be available to the Backup program only, and no one is using Removable Storage to manage new media, check this box. |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set advanced backup options
Using batch files to back up data

# Set the backup type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To set the backup type

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Tools** menu, then click **Options**.

4. On the **Backup Type** tab, in **Default Backup Type**, select the type of backup you want to perform.

### Notes

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- There are five backup types: copy, daily, differential, incremental, and normal. For more information about these backup types, see Related Topics.

#### Using a command line

1. Open Command Prompt.

2. To set the backup type, type:

   **ntbackup backup /M** {*backuptype*}

Backup type must be one of the following: normal, copy, differential, incremental, or daily. For information on additional backup options, see the Ntbackup.

For example, to create a normal backup job named "Backup Job 1" that backs up the System State data to the file C:\backup.bkf, type:

**ntbackup backup systemstate /M normal /J "Backup Job 1" /F "C:\backup.bkf"**

### Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to

prevent losing data. Some file systems might not support all of the features of other file systems.

- To view the complete syntax for this command, at a command prompt, type:

  **ntbackup /?**

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set advanced backup options
Using Network Monitor
Backup
Checklists for Backing Up and Restoring Data
Types of backup
Using batch files to back up data

# Exclude file types from a backup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To exclude file types from a backup**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Tools** menu, then click **Options**.

4. On the **Exclude Files** tab, do one of the following:

   - Click **Add new** under the **Files excluded for all users** list if you want to exclude files that are owned by all users.

   - Click **Add new** under the **Files excluded for user** list if you want to exclude only files that you own.

5. In the **Add Excluded Files** dialog box, do one of the following:

   - If you want to exclude a registered file type, click the file type in **Registered file type**.

   - If you want to exclude a custom file type, enter a period and then the one, two, or three letter file extension in **Custom file mask**.

6. Type a path in **Applies to path** if you want to restrict the excluded file type to a specific folder or hard disk drive, and then click **OK**.

## Notes

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- If you restrict excluded files to a specific path (folder), the files will be restricted from all subfolders of that path unless you clear the **Applies to all subfolders** check box.

- You can change how a file is excluded from a backup set by clicking the file in **Files excluded for all users** or **Files excluded for user**, and then clicking **Edit**. The **Add Excluded Files** dialog box appears.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using Network Monitor

Backup
Checklists for Backing Up and Restoring Data
Set backup options
Set advanced backup options

# Set advanced backup options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set advanced backup options**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Backup** tab and select the files and folders you want to back up.

4. Click **Start Backup**.

5. On the **Backup Job Information** dialog box, click **Advanced**.

6. Set the advanced backup options you want, and then click **OK**. See the Notes section for a description of each option.

## Notes

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

| Item | Description |
|---|---|
| Back up data that is in Remote Storage | Backs up data that has been designated for Remote Storage. If you choose this option, Remote Storage reparse points (placeholder files) are backed up. If you do not select this option, Remote Storage reparse points are not backed up. You can restore Remote Storage data only to an NTFS volume. |
| Verify data after backup | Checks for disk errors that may have occurred while files were being copied to tape. This might substantially increase the time it takes to perform a backup. |
| If possible, compress the backup data to save space | Compresses the data you are backing up so that you can save more data on a tape. If this option is disabled, you do not have a tape drive on your computer or your tape drive cannot handle compressed data. |
| Automatically back up System protected Files with the System State | Backs up all of the system files that are in your systemroot directory in addition to the boot files that are included with the System State data. This option will substantially increase the size of your backup job, but select it if you have recently updated drivers or changed other operating system files. For example, if you choose this option on Windows XP Professional, it will add over 200 megabytes to the size of your backup job. |
| Backup Type | Determines how your data is backed up. Backup types include copy, daily, differential, incremental, and normal. |

#### Using a command line

1. Open Command prompt.

2. To set the advanced backup options, type:

   **ntbackup backup/V:{yes|no} /HC:{on|off}**

| Value | Description |
|---|---|
| **/V: {yes\|no}** | Specifies whether or not the backup data is verified after the backup is complete. |
| **/HC: {on\|off}** | Specifies whether or not hardware compression is used (if it is available on the backup media). Usually only tape drives have the ability to compress data. |

For other backup options, see information about the Ntbackup command.

For example, to create a backup job named "Backup Job 1" that backs up data on drive D:\ to the file C:\backup, type:

**ntbackup backup D:\ /J "Backup Job 1" /F "C:\backup.bkf"**

All other options will default to those specified in the Backup program.

### Notes

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **ntbackup /?**

- If you do not specify the other Backup options, **ntbackup** will use the default values of the Backup program for the backup type, verification setting, logging level, hardware compression, and any other settings.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set backup options
Exclude file types from a backup
Using batch files to back up data

# Set restore options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set restore options**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Tools** menu, click **Options**.

4. On the **Restore** tab, do one of the following:

   - Click **Do not replace the file on my computer** if you do not want the restore operation to copy over files that are already on your disk.

   - Click **Replace the file on disk only if the file on disk is older** if you want the restore operation to replace older files on your disk with newer files from your backup.

   - Click **Always replace the file on my computer** if you want the restore operation to replace files on your disk regardless whether the backup files are newer or older.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set advanced restore options
Set the backup type
Checklist: Restoring data

# Set advanced restore options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set advanced restore options**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. Click the **Restore and Manage Media** menu and select the files to restore.

4. Click the **Start Restore** button.

5. On the **Confirm Restore** dialog box, click **Advanced**.

6. Set the advanced restore options you want, and then click **OK**. See the Notes section for a description of each option.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

| Item | Description |
|---|---|
| **Restore security** | Restores security settings for each file and folder. Security settings include permissions, audit entries, and ownership. This option is available only if you have backed up data from an NTFS volume used in the Windows Server 2003 family of operating systems, and you are restoring it to an NTFS volume used in Windows Server 2003 family operating systems. |
| **Restore junction points, and restore file and folder data under junction points to the original location** | Restores junction points on your hard disk as well as the data that the junction points point to. If you do not select this check box, the junction points will be restored as common directories and the data your junction points point to will not be accessible. Also, if you are restoring a mounted drive and you want to restore the data that is on the mounted drive, you must select this check box. If you do not select this check box, you will only restore the folder containing the mounted drive. |
| **When restoring replicated data sets, mark the restored data as the primary data for all replicas** | Performs a primary restore. Ensures that restored File Replication service (FRS) data is replicated to your other servers. Select this option only when restoring the *first* replica set to the network. Do not use this option if one or more replica sets have already been restored. |

| | |
|---|---|
| **Restore the Cluster Registry to the quorum disk and all other nodes** | Ensures that the cluster quorum database is restored and replicated on all nodes in a server cluster. If you select this option, Backup will stop the Cluster service on all other nodes of the server cluster after the node that was restored is rebooted. The entire server cluster will therefore be down during an authoritative restore of the data on the cluster quorum disk. For more information, see Related Topics. |
| **Preserve existing volume mount points** | Prevents the restore operation from writing over any volume mount points that you have created on the partition or volume you are restoring data to. This option is primarily applicable when you are restoring data to an entire drive or partition. For example, if you are restoring data to a replacement drive and you have partitioned, formatted the drive, and restored volume mount points, select this option so your volume mount points are not restored. If you are restoring data to a partition or drive that you have just reformatted and you want to restore the old volume mount points, do not select this option. |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore files from a file or a tape
Set restore options
Checklist: Restoring data
Restore the cluster database on a local node
Authoritative, primary, and normal restores

# Maintain Media

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Maintain media**
Backup provides a way to perform some simple management of backup files and tapes. For advanced management of backup media, see Removable Storage.

- Format a tape
- Retension a tape
- Mark a tape as free
- Delete a catalog
- Update a catalog

# Format a tape

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To format a tape**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. On the **Restore and Manage Media** tab, right-click the tape you want to format in **Backup Identification Label**.

4. Click **Format**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- Not all drives support using Backup to format a tape.

- If you do not have a tape drive, this feature is not available.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Mark a tape as free
Removable Storage
Maintain Media

# Retension a tape

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To retension a tape**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. On the **Restore and Manage Media** tab, click the tape you want to retension.

4. On the **Tools** menu, click **Media Tools**, and then click **Retension**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- You can also access the **Retension** command by right-clicking a tape on the **Restore and Manage Media** tab.

- Not all drives require or support the retension option.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Removable Storage
Maintain Media

# Mark a tape as free

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To mark a tape as free**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. On the **Restore and Manage Media** tab, click the tape you want to erase.

4. On the **Tools** menu, click **Media Tools**, and then click **Mark as free** .

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- Do not use this option if you need to erase tapes for legal reasons. **Mark as free** will not destroy the data on the tape. You must use another method to erase the tape.

- You can only mark a tape as free if it is in the tape drive or in a loader slot of a tape changer.

- Marking a tape as free moves the tape from the backup media pool to the free media pool.

- You can also use Removable Storage to manage tapes.

- You can also access the **Mark as free** command by right-clicking a tape on the **Restore and Manage Media** tab.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Format a tape
Removable Storage
Maintain Media

# Delete a catalog

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a catalog**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. On the **Restore and Manage Media** tab, right-click the backup set that contains the catalog you want to delete.

4. Click **Delete catalog** if the backup set has been saved to a file, or click **Delete from system** if the backup set has been saved to a tape.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- Use this procedure only if you have a damaged or lost tape or file and you want to delete it from the on-disk catalog.

- If you click **Delete catalog**, this procedure will delete the on-disk catalog for the selected backup set. It will not delete the backup file.

- If you click **Delete catalog from system**, this procedure will delete the on-disk catalog for the selected backup set, and it will delete the tape from the Removable Storage database if the tape is offline. It will not delete the on-media catalog that is on the tape.

- You can restore a catalog for a file or a tape by clicking the **Restore and Manage Media** tab, right-clicking **File** or the tape you want to catalog, and then clicking **Catalog file** or **Catalog**. For more information on creating a catalog, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Removable Storage
Format a tape

# Update a catalog

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To update a catalog**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. On the **Restore and Manage Media** tab, right-click the media containing the backup set you want to catalog.

4. Click **Catalog**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- This procedure creates an on-disk catalog. It does not create an on-media catalog.

- If you are creating a catalog of a backup set that is saved on a tape, do one of the following before you perform this procedure:

  - If you have all of the tapes in the backup set and the tapes are not damaged or corrupted, click the **Tools** menu, click **Options**, click the **General** tab, and then make sure the **Use the catalogs on the media to speed up building restore catalogs on disk** check box is checked.

  - If you are missing a tape in the backup set or a tape is damaged or corrupted, click the **Tools** menu, click **Options**, click the **General** tab, and then make sure the **Use the catalogs on the media to speed up building restore catalogs on disk** check box is cleared. This will ensure that the catalog is complete and accurate; however, it might take a long time to create the catalog.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Removable Storage
Maintain Media

# Perform System Repair and Recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Perform system repair and recovery**

- Create an Automated System Recovery set using Backup
- Recover from a system failure using Automated System Recovery
- Install the Recovery Console as a startup option
- Delete the Recovery Console
- Start the computer and use the Recovery Console

# Create an Automated System Recovery set using Backup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an Automated System Recovery set using Backup**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use the Backup or Restore Wizard to create an Automated System Recovery (ASR) set by answering **All information on this computer** in the **What do you want to backup?** section. Otherwise, you can go to the next step to create an ASR set in **Advanced Mode**.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. On the **Tools** menu, click **ASR Wizard**.

4. Follow the instructions that appear on your screen.

### Important

- You will need a blank 1.44 MB floppy disk to save your system settings and media to contain the backup files. If your computer does not have a floppy drive, see Backup Troubleshooting.

### Notes

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- This procedure will only back up those system files that are necessary for starting your system. You will need to backup your data separately.

- After creating the ASR set, label this floppy disk and backup media carefully and keep them together. To use the backup media, you must have the floppy disk that you created with that set of media. You cannot use a floppy disk created at a different time or with a different set of media. You must also have your Setup CD available at the time you perform Automated System Recovery.

- Keep the ASR set in a secure location. The ASR set contains information about your systems configuration that could be used to do harm to your system.

- If you are backing up a server cluster, run the Automated System Recovery Preparation Wizard on all nodes of the cluster and make sure that the Cluster service is running when you start each Automated System Recovery backup. Make sure that one of the nodes on which you run the Automated System Recovery Preparation Wizard is listed as the owner of the quorum resource while the wizard is running.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Recover from a system failure using Automated System Recovery
Backing up and restoring server clusters

# Recover from a system failure using Automated System Recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To recover from a system failure using Automated System Recovery**

1. Make sure you have the following available before you begin the recovery procedure:

   - Your previously created Automated System Recovery (ASR) floppy disk.
   - Your previously created backup media.
   - The original operating system installation CD.
   - If you have a mass storage controller and you are aware that the manufacturer has supplied a separate driver file for it (different from driver files available on the Setup CD), obtain the file (on a floppy disk) before you begin this procedure.

2. Insert the original operating system installation CD into your CD drive.
3. Restart your computer. If you are prompted to press a key to start the computer from CD, press the appropriate key.
4. If you have a separate driver file as described in step 1, use the driver as part of Setup by pressing F6 when prompted.
5. Press F2 when prompted at the beginning of the text-only mode section of Setup. You will be prompted to insert the ASR floppy disk you have previously created.
6. Follow the directions on the screen.
7. If you have a separate driver file as described in step 1, press F6 (a second time) when prompted after the system reboots.
8. Follow the directions on the screen.

## Notes

- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.
- Automated System Recovery will not restore your data files. See Related Topics for more information on backing up and restoring your data files.
- If you are restoring a server cluster in which all nodes failed and the quorum disk cannot be restored from backup, use Automated System Recovery on each node in the original cluster to restore the disk signatures and the partition layout of the cluster disks (quorum and nonquorum). For more information about backing up and restoring server clusters, see the Related Topics link at the bottom of this page.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Create an Automated System Recovery set using Backup

Back Up and Restore Server Clusters

Backup overview

# Install the Recovery Console as a startup option

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install the Recovery Console as a startup option**

1. With Windows running, insert the Setup CD into your CD-ROM drive.
2. Click **Start** and select **Run**.
3. Type the following where *D*: is the CD-ROM drive letter:

   On 32-bit versions of Windows Server 2003: `D:\i386\winnt32.exe /cmdcons`

   - or -

   On x64-based versions of Windows Server 2003: `D:\amd64\winnt32.exe /cmdcons`
4. Follow the instructions on the screen.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- You can not install the Recovery Console on Itanium-based computers.
- To run the Recovery Console, you must restart your computer and select the Recovery Console option from the list of available operating systems.
- To see the commands available on the Recovery Console, type **help** at the console prompt.
- If your computer will not start or you are using an Itanium-based computer, you can run the Recovery Console from the Setup CD. See Start the computer and use the Recovery Console for information on running the Recovery Console when your computer will not start.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovery Console commands
Recovery Console overview
Delete the Recovery Console
Start the computer and use the Recovery Console
Working with MMC console files

# Delete the Recovery Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete the Recovery Console**

1. Open My Computer.
2. Double-click the hard drive on which you installed the Recovery Console.
3. On the **Tools** menu, click **Folder Options**.
4. Click the **View** tab.
5. Click **Show hidden files and folders**, clear the **Hide protected operating system files** check box, and then click **OK**. If you receive a warning about displaying protected system files, click **Yes** to continue.
6. At the root directory, delete the \Cmdcons folder.
7. At the root directory, delete the file Cmldr.
8. Click **Start**, then **Control Panel**, then click **System**.
9. On the **Advanced** tab under **Startup and Recovery**, select **Settings**.
10. On the **Startup and Recovery** dialog box, under **System startup**, select **Edit** to edit the startup options file manually. Boot.ini will display in Notepad.
11. Remove the entry for the Recovery Console. It will look similar to this:

    C:\cmdcons\bootsect.dat="Microsoft Windows Recovery Console" /cmdcons

12. Save the file and close it.

### Caution

- Modifying the Boot.ini file incorrectly may prevent your computer from restarting. Be sure to delete only the entry for the Recovery Console.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- You can only install the Recovery Console feature on x86-based computers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovery Console overview
Repair overview

# Start the computer and use the Recovery Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the computer and use the Recovery Console**

1. Insert the Setup CD and restart the computer from the CD drive.
2. When the text-based part of Setup begins, follow the prompts; choose the repair or recover option by pressing **R**.
3. If you have a dual-boot or multiple-boot system, choose the installation that you need to access from the Recovery Console.
4. When prompted, type the password for the local administrator account.
5. At the system prompt, type the relevant Recovery Console commands; type **help** for a list of commands or **help***commandname* for help about a specific command.
6. To exit the Recovery Console and restart the computer, type **exit**.

**Important**

- Because the Recovery Console is quite powerful, it is recommended for use only by advanced users or administrators.
- In most cases, you will need the password for the built-in administrator account to use the Recovery Console.
- You can set the **Recovery console: Allow automatic administrative logon** setting in the Security Options section of Local Security Settings. Changing this setting will allow you to start the Recovery Console without having to provide the administrator password. For more information, see "Recovery console: Allow automatic administrative logon" and "Edit local security settings" in Related Topics at the bottom of this page.

  Be careful when enabling this option. This will allow anyone with physical access to a computer to boot into the Recovery Console and access the Recovery Console commands.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovery Console commands
Recovery Console overview
Repair overview
Edit local security settings
Recovery console: Allow automatic administrative logon

# Backup Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about the Backup utility.

- Backup overview
- Understanding Backup
- Administering Backup

# Backup overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backup overview**

The Backup utility helps you protect data from accidental loss if your system's hardware or storage media fails. For example, you can use Backup to create a duplicate copy of the data on your hard disk and then archive the data on another storage device. The backup storage medium can be a logical drive such as your hard disk, a separate storage device such as a removable disk, or an entire library of disks or tapes organized and controlled by a robotic changer. If the original data on your hard disk is accidentally erased or overwritten, or becomes inaccessible because of a hard disk malfunction, you can easily restore the data from the archived copy.

Backup creates a volume shadow copy of your data to create an accurate copy of the contents of your hard drive, including any open files or files that are being used by the system. Users can continue to access the system while the Backup utility is running without risking loss of data.

For more information on volume shadow copy, see Volume shadow copy overview.

Using Backup, you can:

- Archive selected files and folders on your hard disk.
- Restore the archived files and folders to your hard disk or any other disk you can access.
- Use Automated System Recovery to save and restore all the system files and configuration settings that you need to recover from a complete system failure.
- Make a copy of any Remote Storage data and any data stored in mounted drives.
- Make a copy of your computer's System State.
- Create a log of what files were backed up and when the backup was performed.
- Make a copy of your computer's system partition, boot partition, and the files needed to start up your system in case of computer or network failure.
- Schedule regular backups to keep your archived data up-to-date.
- Backup also performs simple media management functions, for example, formatting. More advanced management tasks such as mounting and dismounting a tape or disk are done by a service called Removable Storage. For more information on Removable Storage, see Removable Storage overview.

**Caution**

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.
- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.

# Understanding Backup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Backup**
This section covers:

- Types of backup
- Volume shadow copy overview
- Permissions and user rights required to back up and restore
- System State data
- Authoritative, primary, and normal restores
- Repair overview
- Recovery Console overview
- Automated System Recovery (ASR) overview

# Types of backup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Types of backup**
The Backup utility supports five methods of backing up data on your computer or network.

**Copy backup**
A copy backup copies all the files you select, but does not mark each file as having been backed up (in other words, the archive attribute is not cleared). Copying is useful if you want to back up files between normal and incremental backups because copying does not affect these other backup operations.

**Daily backup**
A daily backup copies all the files that you select that have been modified on the day the daily backup is performed. The backed-up files are not marked as having been backed up (in other words, the archive attribute is not cleared).

**Differential backup**
A differential backup copies files that have been created or changed since the last normal or incremental backup. It does not mark files as having been backed up (in other words, the archive attribute is not cleared). If you are performing a combination of normal and differential backups, restoring files and folders requires that you have the last normal as well as the last differential backup.

**Incremental backup**
An incremental backup backs up only those files that have been created or changed since the last normal or incremental backup. It marks files as having been backed up (in other words, the archive attribute is cleared). If you use a combination of normal and incremental backups, you will need to have the last normal backup set as well as all incremental backup sets to restore your data.

**Normal backup**
A normal backup copies all the files you select and marks each file as having been backed up (in other words, the archive attribute is cleared). With normal backups, you only need the most recent copy of the backup file or tape to restore all of the files. You usually perform a normal backup the first time you create a backup set.

Backing up your data using a combination of normal backups and incremental backups requires the least amount of storage space and is the quickest backup method. However, recovering files can be time-consuming and difficult because the backup set might be stored on several disks or tapes.

Backing up your data using a combination of normal backups and differential backups is more time-consuming, especially if your data changes frequently, but it is easier to restore the data because the backup set is usually stored on only a few disks or tapes.

For more information, see Set the backup type.

# Volume shadow copy overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Volume shadow copy overview**

With Backup, you can create shadow copy backups of volumes and exact copies of files, including all open files. For example, databases that are held open exclusively and files that are open due to operator or system activity are backed up during a volume shadow copy backup. In this way, files that have changed during the backup process are copied correctly.

Shadow copy backups ensure that:

- Applications can continue to write data to the volume during a backup.
- Files that are open are no longer omitted during a backup.
- Backups can be performed at any time, without locking out users.

Some applications manage storage consistency differently while files are open, which can affect the consistency of the files in the backup. For critical applications, consult the application documentation or your provider for information about the recommended backup method. When in doubt, shutdown the application before performing a backup.

It is highly recommended that you perform backups when server usage is low. For example, perform backups late at night.

For more information, see Set advanced backup options.

# Permissions and user rights required to back up and restore

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Permissions and user rights required to back up and restore**
You must have certain permissions and user rights to back up files and folders. If you are an administrator or a backup operator in a local group, you can back up any file and folder on the local computer to which the local group applies. Likewise, if you are an administrator or backup operator on a domain controller you can back up any file and folder locally on any computer in the domain or any computer on a domain with which you have a two-way trust relationship. However, if you are not an administrator or a backup operator, and you want to back up files, then you must be the owner of the files and folders that you want to back up, or you must have one or more of the following permissions for the files and folders you want to back up: Read, Read and execute, Modify, or Full Control.

You must also be certain that there are no disk-quota restrictions that may restrict your access to a hard disk. These restrictions make it impossible for you to back up data. You can check whether you have any disk-quota restrictions by right-clicking the disk you want to save data to, clicking **Properties**, and then clicking the **Quota** tab.

You can also restrict access to a backup file by selecting **Allow only the owner and the Administrator access to the backup data** in the **Backup Job Information** dialog box. If you select this option, only an administrator or the person who created the backup file will be able to restore the files and folders.

**Notes**

- You can only use Backup locally; you cannot backup a remote computer.
- You can only back up and restore System State data on a local computer. You cannot back up and restore System State data on a remote computer even if you are an administrator on the remote computer.

For more information, see Identify members of a local group and Best practices for permissions and user rights.

# System State data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System State data**
With Backup, you can back up and restore the following system components to back up the System State:

| Component | When this component is included in System State? |
|---|---|
| Registry | Always |
| COM+ Class Registration database | Always |
| Boot files, including the system files | Always |
| Certificate Services database | If it is a Certificate Services server |
| Active Directory directory service | If it is a domain |
| SYSVOL directory | Only if it is a domain controller |
| Cluster service information | If it is within a cluster |
| IIS Metadirectory | If it is installed |
| System files that are under Windows File Protection | Always |

Backup refers to these system components as the System State data. The exact system components that make up your computer's System State data depend on the computer's operating system and configuration.

**Windows XP Professional**
The System State data includes only the registry, COM+ Class Registration database, files under Windows File Protection, and boot files.

**Windows Server 2003 family operating systems**
The System State data includes the registry, COM+ Class Registration database, files under Windows File Protection, and system boot files. Depending on the configuration of the server, other data may be included in the System State data. For example, if the server is a certificate server, the System State will also contain the Certificate Services database. If the server is a domain controller, Active Directory and the SYSVOL directory are also contained in the System State data.

**Backing up System State data**
When you choose to back up or restore the System State data, all of the System State data that is relevant to your computer is backed up or restored; you cannot choose to back up or restore individual components of the System State data. This is due to dependencies among the System State components. However, you can restore the System State data to an alternate location. If you do this, only the registry files, SYSVOL directory files, Cluster database information files, and system boot files are restored to the alternate location. The Active Directory directory services database, Certificate Services database, and COM+ Class Registration database are not restored if you designate an alternate location when you restore the System State data.

**Restoring Active Directory**
If you have more than one domain controller in your organization and the Active Directory directory service is replicated to any of these other servers, you may have to authoritatively restore any Active Directory data that you want to restore. To do this, run the Ntdsutil utility after you have restored the System State data, but before you restart the server on the network. Run this utility only while the server is in Directory Services Restore Mode. With the Ntdsutil utility, you can mark Active Directory objects for authoritative restore. This utility ensures that any replicated or distributed data that you restore is properly replicated or distributed throughout your organization.

For example, if you inadvertently delete or modify objects stored in the Active Directory directory service, and those objects are replicated or distributed to other servers, you will need to authoritatively restore those objects so they are replicated or distributed to the other servers. If you do not authoritatively restore the objects, they will never get replicated or distributed to your other servers because they will appear to be older than the objects currently on your other servers. Using the Ntdsutil utility to mark objects for authoritative restore ensures that the data you want to restore gets replicated or distributed throughout your organization. On the other hand, if your system disk has failed or the Active Directory database is corrupted,

then you can simply restore the data nonauthoritatively without using the Ntdsutil utility.

The Ntdsutil command line utility can be run from the command prompt. Help for the Ntdsutil utility can also be found at the command prompt by typing **ntdsutil /?**.

**Notes**

- System State data contains most elements of a system's configuration, but it might not include all of the information that you require to recover your system from a failure. Therefore, it is recommended that you back up all boot and system volumes, including the System State, when you back up your system.

- Although you cannot change which components of the System State are backed up, you can back up all of the protected system files with the System State data by setting advanced backup options.

- To restore the System State data on a domain controller, you must first start your computer in Directory Services Restore Mode. This will allow you to restore the SYSVOL directory and the Active Directory.

- You can only back up and restore the System State data on a local computer. You cannot back up and restore the System State data on a remote computer.

For more information, see Back up System State data.

# Authoritative, primary, and normal restores

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authoritative, primary, and normal restores**
In Backup, distributed services data that is part of the System State data, such as the Active Directory directory service data, can be restored using one of three restore methods:

- Primary restore
- Normal (nonauthoritative) restore
- Authoritative restore

To understand how each restore method works, it is important to understand how the Backup utility backs up data for distributed services. When you back up the System State data on a domain controller, you are backing up all Active Directory data that exists on that server (along with other system components such as the SYSVOL directory and the registry). To restore these distributed services to that server, you must restore the System State data. However, the number and configuration of domain controllers in your system will dictate the type of restore method you choose. For example, if you need to roll back replicated Active Directory changes, but have more than one domain controller in your organization, you will need to perform an authoritative restore to ensure that your restored data gets replicated to all of your servers. However, if you need to restore Active Directory data on a standalone domain controller or on the first of several domain controllers, you will need to perform a primary restore. If you need to restore Active Directory data on just one domain controller in a system where Active Directory data is replicated across several domain controllers, you can use a normal restore if your restored data does not have to be replicated to all your servers.

**Primary restore**
Use this type of restore when the server you are trying to restore is the only running server of a replicated data set (for example, the SYSVOL and FRS are replicated data sets). Select primary restore only when restoring the **first** replica set to the network. Do not use primary if one or more replica sets have already been restored. Typically, perform a primary restore only when all the domain controllers in the domain are lost, and you are trying to rebuild the domain from backup.

| Distributed Data | Reason for using Primary Restore of System State Data |
|---|---|
| Active Directory | Restoring a standalone domain controller.Restoring the first of several domain controllers. |
| SYSVOL | Restoring a standalone domain controller.Restoring the first of several domain controllers. |
| Replica Sets | Restoring the first replica set. |

To perform a primary restore, select **When restoring replicated data sets, mark the restored data as the primary data for all replicas** in the advanced options. For more information, see Set advanced restore options.

**Normal restore**
During a normal restore operation, Backup operates in nonauthoritative restore mode. That is, any data that you restore, including Active Directory objects, will have their original update sequence number. The Active Directory replication system uses this number to detect and propagate Active Directory changes among the servers in your organization. Because of this, any data that is restored nonauthoritatively will appear to the Active Directory replication system as though it is old, which means the data will never get replicated to your other servers. Instead, if newer data is available from your other servers, the Active Directory replication system will use this to update the restored data. To replicate the restored data to the other servers, you must use an authoritative restore.

| Distributed Data | Reason for using Normal Restore of System State Data |
|---|---|
| Active Directory | Restoring a single domain controller in a replicated environment. |
| SYSVOL | Restoring a single domain controller in a replicated environment. |
| Replica Sets | Restoring all but the first replica sets (that is, sets 2 through *n*, for *n* replica sets). |

**Authoritative restore**

To authoritatively restore Active Directory data, you need to run the Ntdsutil utility after you have restored the System State data but before you restart the server. The Ntdsutil utility lets you mark Active Directory objects for authoritative restore. When an object is marked for authoritative restore its update sequence number is changed so that it is higher than any other update sequence number in the Active Directory replication system. This will ensure that any replicated or distributed data that you restore is properly replicated or distributed throughout your organization.

For example, if you inadvertently delete or modify objects stored in the Active Directory directory service, and those objects are replicated or distributed to other servers, you will need to authoritatively restore those objects so they are replicated or distributed to the other servers. If you do not authoritatively restore the objects, they will never get replicated or distributed to your other servers because they will appear to be older than the objects currently on your other servers. Using the Ntdsutil utility to mark objects for authoritative restore ensures that the data you want to restore gets replicated or distributed throughout your organization. On the other hand, if your system disk has failed or the Active Directory database is corrupted, then you can simply restore the data nonauthoritatively without using the Ntdsutil utility.

You can run the Ntdsutil command-line utility from the command prompt. For more information about using ntsustil to perform an authoritative restore, see Ntdsutil. Help for the Ntdsutil utility is also available through the command prompt by typing **ntdsutil /?**.

| Distributed Data | Reason for using Authoritative Restore of System State Data |
|---|---|
| Active Directory | Rolling back or undoing changes. |
| SYSVOL | Resetting data. |
| Replica Sets | Rolling back or undoing changes. |

**Caution**

- When you restore the System State data, and you do not designate an alternate location for the data, Backup will erase the System State data that is currently on your computer and replace it with the System State data you are restoring.

**Notes**

- To restore the System State data on a domain controller, you must first start your computer in Directory Services Restore Mode. This will allow you to restore the SYSVOL directory and the Active Directory. For more information on starting your computer in Directory Services Restore Mode, see Startup options.

- You can only restore the System State data on a local computer. You cannot restore the System State data on a remote computer.

For more information, see Set advanced restore options.

# Repair overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Repair overview**
The Microsoft® Windows Server 2003 family includes several features to ensure that your computer and the applications and devices installed on it work correctly. These features help you solve the problems that can result from adding, deleting, or replacing files that your operating system, applications, and devices require to function.

**Important**

- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.

The recovery feature or features that you use depend on the type of problem or failure you encounter. Use the following list to determine which features to use to address a specific problem.

- Your personal data files are missing or have been corrupted, or you want to revert to a previous version of a file.
- After updating a device driver, you can log on, but the system is unstable.
- After installing a new device, the system is unstable.
- After installing an application, the system is unstable or an application is not working properly.
- Operating system does not start (the logon screen does not appear).

**Your personal data files are missing or have been corrupted, or you want to revert to a previous version of a file.**
If you have saved a copy of the data file in another location, you can copy it from that location back to your hard disk. You can make a back-up copy at any time by copying the file to another location on your computer, to external storage such as a removable disk, or to a different computer.

**Feature: Backup Utility**

**When to use it:** When you need to restore a version of a data file that you saved using Backup.

**What it does:** Restores files from backups.

**For more information**, see Restore files from a file or a tape.

**After updating a device driver, you can log on, but the system is unstable.**

**Feature: Device Driver Roll Back**

**When to use it:** When the only change you want to undo is the updating of a device driver (other than a printer driver). You must be logged on as an administrator.

**What it does:** Reinstalls the driver you were using previously and restores any driver settings that were changed when you added the new driver. Affects no other files or settings. Printer drivers cannot be restored with Device Driver Roll Back.

**For more information**, see Roll back to the previous version of a driver.

**After installing a new device, the system is unstable.**

**Feature: Disable the Device**

**When to use it:** When you suspect that one or more specific hardware devices are causing the problem. You must be logged on as an administrator.

**What it does:** Disables the hardware device and its drivers.

**For more information**, see Disable a device.

**After installing an application, the system is unstable or an application is not working properly.**
Under some circumstances, the computer will prompt you to repair or reinstall an application from a network location or from

the original set-up media (for example, the program CD). If this happens, follow the directions on the screen to repair your program. If you are not prompted with this information, you can use one of the following features to repair your program:

**Feature: Add or Remove Programs**

**When to use it:** When you suspect that one or more specific programs are causing the problem.

**What it does:** Removes the programs you specify.

**For more information**, see Windows interface administrative tool reference A-Z: Add or Remove Programs.

**Operating system does not start (the logon screen does not appear).**

**Feature: Last Known Good Configuration startup option**

**When to use it:** When you suspect that a change you made to your computer before restarting might be causing the failure.

**What it does:** Restores the registry settings and drivers that were in effect the last time the computer started successfully.

**For more information**, see Start the computer using the last known good configuration.

**Feature: Recovery Console**

**When to use it:** If using the Last Known Good Configuration startup option is unsuccessful and you cannot start the computer in Safe Mode. This method is recommended only if you are an advanced user who can use basic commands to identify and locate problem drivers and files. To use the Recovery Console, restart the computer with the installation CD for the operating system in the CD drive. When prompted during text-mode setup, press R to start the Recovery Console.

**What it does:** From the Recovery Console, you can access the drives on your computer. You can then make any of the following changes so that you can start your computer:

- Enable or disable device drivers or services.
- Copy files from the installation CD for the operating system, or copy files from other removable media. For example, you can copy an essential file that had been deleted.
- Create a new boot sector and new master boot record (MBR). You might need to do this if there are problems starting from the existing boot sector.
-

**For more information**, see Recovery Console overview.

**Feature: Automated System Recovery**

**When to use it:** After all other attempts at recovery fail, or if you have replaced a damaged system hard drive. You must have previously used Backup to create an Automated System Recovery set.

**What it does:** Automated System Recovery restores all disk signatures, volumes, and partitions on the disks required to start the computer. Automated System Recovery then installs a simplified installation of Windows and automatically starts a restoration using the backup created by the Automated System Recovery Preparation Wizard.

**For more information**, see Automated System Recovery (ASR) overview.

**Feature: Windows installation CD**

**When to use it:** When you cannot start the computer in Safe Mode, you cannot use either the Last Known Good Configuration startup option or the Recovery Console successfully, and you have no backup to use with Automated System Recovery.

**What it does:** Reinstalls the operating system. You will then need to reinstall your applications and restore your data files from backups.

# Recovery Console overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovery Console overview**
If Safe Mode and other startup options do not work, consider using the Recovery Console. This method is recommended only if you are an advanced user who can use basic commands to identify and locate problem drivers and files. In addition, you will need the password for the built-in administrator account to use the Recovery Console.

Using the Recovery Console, you can enable and disable services, format drives, read and write data on a local drive (including drives formatted to use NTFS), and perform many other administrative tasks. The Recovery Console is particularly useful if you need to repair your system by copying a file from a floppy disk or CD-ROM to your hard drive, or if you need to reconfigure a service that is preventing your computer from starting properly.

**Important**

- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.

There are two ways to start the Recovery Console:

- If you are unable to start your computer, you can run the Recovery Console from your Setup CD.
- As an alternative, you can install the Recovery Console on your computer to make it available in case you are unable to restart Windows. You can then select the Recovery Console option from the list of available operating systems on when you start your computer. You cannot install the Recovery Console on an Itanium-based computer.

After you start the Recovery Console, you will have to choose which installation you want to log on to (if you have a dual-boot or multiple-boot system), and you will have to log on with the password for the built in administrator account for that installation.

The console provides commands you can use to do simple operations such as changing to a different directory or viewing a directory, as well as more powerful operations such as fixing the boot sector. You can access Help for the commands in the Recovery Console by typing **help** at the Recovery Console command prompt.

For more information about securing your computer, see Best practices for security.

**Note**

- Installation of the Recovery Console provides an alternate method for logging on to the computer. This could potentially provide a method for an intruder to gain access to the computer by circumventing many of the built-in safeguards. Consider the potential security effects before installing the Recovery Console.

# Automated System Recovery (ASR) overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automated System Recovery (ASR) overview**
With Automated System Recovery (ASR), you can create ASR sets on a regular basis as part of an overall plan for system recovery in case of system failure. Use ASR as a last resort in system recovery, only after you have exhausted other options such as the startup options Safe Mode and Last Known Good Configuration. For more information about these and other recovery options, see Startup options.

ASR is a recovery option that has two parts: ASR backup and ASR restore. You can access the backup portion through the Automated System Recovery Preparation Wizard located in Backup. The Automated System Recovery Preparation Wizard backs up the System State data, system services, and all disks associated with the operating system components. It also creates a floppy disk, which contains information about the backup, the disk configurations (including basic and dynamic volumes), and how to accomplish a restore.

You can access the restore part of ASR by pressing F2 when prompted in the text mode portion of setup. ASR reads the disk configurations from the floppy disk and restores all of the disk signatures, volumes and partitions on the disks required to start your computer (at a minimum). (It will attempt to restore all of the disk configurations, but under some circumstances, it may not be able to). ASR then installs a simple installation of Windows and automatically starts to restore from backup using the backup ASR set created by the Automated System Recovery Preparation Wizard.

**Notes**

- ASR does not include data files. Back up data files separately on a regular basis and restore them after the system is working.
- Using **All information on this computer** in the simple wizard also creates an ASR floppy disk and an ASR set.
- ASR supports FAT16 volumes up to 2.1 GB only. ASR does not support 4 GB FAT16 partitions that use a cluster size of 64K. If your system contains 4-GB FAT16 partitions, convert them from FAT16 to NTFS before using ASR.

# Administering Backup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Backup**
This section covers:

- Managing Backup from the command line
- Using batch files to back up data

# Managing Backup from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Backup from the command line**

In addition to using the Backup Wizard, you can use the following command line utility to manage Backup:

- Ntbackup

For information about other command-line utilities, see Command-line reference A-Z.

# Using batch files to back up data

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using batch files to back up data**

You can perform backup operations from batch files using the **ntbackup** command followed by various command line parameters. However, there are two important limitations to using batch files to back up your data:

- Using the Ntbackup command, you can back up entire folders only. You cannot designate individual files for backup. However, you can designate a backup selection file (.bks file) from the command line, which contains a list of files you want to back up. You must use the GUI version of the Backup utility to create backup selection files.

- The **ntbackup** command does not support the use of wildcard characters. For example, typing ***.txt** will not back up files with a .txt extension.

For more information about using Backup, see "Backup and Restore" at the Microsoft Windows Resource Kits Web site.

# Backup Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section provides problem-solving guidelines for the Backup utility.

What problem are you having?

- Backup does not back up to a CD device.
- I cannot locate one of my old backup logs.
- I want to do an Automated System Recovery (ASR) backup, but my computer does not have a floppy drive.
- I want to do an Automated System Recovery (ASR) restore, but I do not have the original floppy disk created during the ASR backup.

**Backup does not back up to a CD device.**
**Cause:** Backup does not support the use of backup media such as CD-RW (compact disc rewritable), CD-R (compact disc recordable), and DVD-R (digital video disc recordable).

**Solution:** Back up to an alternate media type, such as a tape or hard disk. If you plan to use CD-RW, CD-R, or DVD-R disks as backup media, you must backup to a file and copy that file to the CD.

**Note**

- You can use Backup to restore from a CD-R, CD-RW, or DVD-R disk.

**I cannot locate one of my old backup logs.**
**Cause:** Backup saves only the last 10 backup reports (log files).

**Solution:** Copy the backup log files that you want to archive to another location. By default, Backup saves the backup log files in the Documents and Settings*user name*\Local Settings\Application Data\Microsoft\Windows NT\NTbackup\Data folder, where *user name* is the name of the user who performed the backup.

**I want to do an Automated System Recovery (ASR) backup, but my computer does not have a floppy drive.**
**Cause:** When you restore your system files in an ASR restore, Backup requires that you insert a floppy disk containing the ASR backup files.

**Solution:** Perform an ASR backup on the computer without the floppy drive. Copy the asr.sif and asrpnp.sif files located in the %systemroot%\repair directory to another computer with a floppy drive, then copy those files onto a floppy disk.

**Important**

- Before running the ASR restore, you must attach a floppy drive to your computer.

**I want to do an Automated System Recovery (ASR) restore, but I do not have the original floppy disk created during the ASR backup.**
**Cause:** When you restore your system files in an ASR restore, Backup requires that you insert a floppy disk containing the ASR backup files asr.sif and asrpnp.sif.

**Solution:** When you perform an ASR backup, the files asr.sif and asrpnp.sif are copied to a floppy disk and also to the ASR backup set (on tape or other backup media). If you lose the floppy disk, first use the Backup utility to locate and catalog the corresponding ASR backup set, then restore the asr.sif and asrpnp.sif files (located in the %systemroot%\repair directory on the backup media) to a floppy disk. You can use this newly created ASR floppy disk to perform the ASR restore.

**See also:** Update a catalog; Restore files from a file or a tape

# Shadow Copies for Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shadow Copies of Shared Folders**
Shadow Copies of Shared Folders provides point-in-time copies of files that are located on shared resources, such as a file server. With Shadow Copies of Shared Folders, you can view shared files and folders as they existed at points of time in the past.

You can use shadow copies when you want to recover files that were accidentally deleted or overwritten, as well as to compare different versions of files.

- Before deploying Shadow Copies of Shared Folders, see Checklist: Deploying Shadow Copies of Shared Folders.
- For tips on using Shadow Copies of Shared Folders, see Best practices for Shadow Copies of Shared Folders.
- For help with specific tasks, see Shadow Copies for Shared Folders How To....
- For general background information, see Shadow Copies for Shared Folders Concepts.
- For problem-solving information, see Troubleshooting Shadow Copies for Shared Folders.

# Checklist: Deploying Shadow Copies of Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Deploying Shadow Copies of Shared Folders**

| Step | Reference |
|---|---|
| Review key concepts. | Shadow Copies for Shared Folders overview |
| Create shared resources. | Share a folder or drive |
| Decide where to store the shadow copies. | Determining storage options for shadow copies |
| Determine the best schedule for taking shadow copies. | Defining a schedule for shadow copies |
| Enable Shadow Copies of Shared Folders on a volume. | Enable Shadow Copies for Shared Folders |
| Deploy the client software. | Deploying the client software for shadow copies |

# Best practices for Shadow Copies of Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Select a separate volume on another disk as the storage area for shadow copies.**

- Select a storage area on a volume that is not being shadow copied. Using a separate volume on another disk provides two advantages. First, it eliminates the possibility that high I/O load will cause shadow copies to be deleted. Second, this configuration provides better performance. This is the recommended configuration for heavily used file servers.

  For more information, see Change the settings for Shadow Copies of Shared Folders.

**Consider how your clients will be using a shared resource before you enable Shadow Copies of Shared Folders and set scheduling options.**

- Adjust the shadow copy schedule to fit the work patterns of your clients.

  For more information about scheduling shadow copies, see Defining a schedule for shadow copies.

**Do not enable shadow copies on volumes that use mount points.**

- The mounted drive will not be included when shadow copies are taken. Enable shadow copies only on volumes without mount points or when you do not want the shared resources on the mounted volume to be shadow copied.

  For more information, see Using NTFS mounted drives.

**Do not enable shadow copies on dual-booting computers.**

- If you have enabled dual-booting into previous versions of Windows (such as Windows NT 4.0), the shadow copies which persist while restarting the older version might be corrupted and unusable when the computer is started again with Windows Server 2003 .

**Use Backup to perform regular backups of your file server.**

- Shadow Copies of Shared Folders is not a replacement for performing regular backups. Use a backup utility in coordination with Shadow Copies of Shared Folders as your best preparation for restoring.

  For more information, see Backup.

**Do not schedule copies to occur more often than once per hour.**

- The default schedule is 7:00 A.M. and 12:00 noon, Monday through Friday. If you decide that you need copies to be taken more often, verify that you have allotted enough storage space and that you do not take copies so often that server performance degrades. There is also an upper limit of 64 copies per volume that can be stored before the oldest copy is deleted. If shadow copies are taken too often, this limit might be reached very quickly, and older copies could be lost at a rapid rate.

  For more information, see Defining a schedule for shadow copies.

**Before deleting a volume that is being shadow copied, delete the scheduled task for creating shadow copies.**

- If the volume is deleted without deleting the shadow copy task, the scheduled task will fail and an Event ID: 7001 error will be written to the event log. Delete the task before deleting the volume to avoid filling the event log with these errors.

For more information, see Change the schedule for creating shadow copies.

**Use an allocation unit size of 16 kilobytes (KB) or larger when formatting a source volume on which Shadow Copies of Shared Folders will be enabled.**

- If you plan to defragment the source volume on which Shadow Copies of Shared Folders is enabled, it is recommended that you set the cluster allocation unit size to be 16 KB or larger when you initially format the source volume. If you do not, the number of changes caused by defragmentation can cause previous versions of files to be deleted.

- If you require NTFS file compression on the source volume, you cannot use an allocation unit size larger than 4 KB. In this case, when you defragment a volume that is very fragmented, you may lose older shadow copies faster than expected.

  For more information about formatting a volume, see Format a basic volume.

# Shadow Copies for Shared Folders How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Enable Shadow Copies for Shared Folders
- Change the settings for Shadow Copies of Shared Folders
- Change the schedule for creating shadow copies
- Delete a shadow copy
- Revert an entire volume shadow copy to a previous version
- Disable Shadow Copies for Shared Folders

# Enable Shadow Copies for Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Shadow Copies of Shared Folders**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**, select **All Tasks**, and click **Configure Shadow Copies**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click the volume where you want to enable Shadow Copies of Shared Folders, and then click **Enable**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- For information about enabling shadow copies in a server cluster, see Enable Shadow Copies for Shared Folders in a cluster. Do not use Cluster Administrator or cluster.exe to make changes to the settings or schedule for shadow copies of shared folders, use the shadow copies interface.
- You can also access Shadow Copies of Shared Folders by right-clicking a fixed NTFS volume in **My Computer** or in the Disk Management snap-in details pane (in **Computer Management**). Click **Properties**, and then click the **Shadow Copies** tab.
- You can use the Disk Management snap-in to manage Shadow Copies of Shared Folders on remote servers as well. In the **Computer Management** console tree, right-click **Disk Management**, point to **All Tasks**, and then click **Configure Shadow Copies**.
- By enabling Shadow Copies of Shared Folders, you also create a default schedule and settings for subsequent shadow copies. To make changes to these defaults, click **Settings**.
- You can only enable Shadow Copies of Shared Folders on a per-volume basis; that is, you cannot select specific shared folders and files on a volume to be copied or not copied.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Change the settings for Shadow Copies of Shared Folders

# Change the settings for Shadow Copies of Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the settings for Shadow Copies of Shared Folders**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**, select **All Tasks**, and click **Configure Shadow Copies**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click the volume where you want to make changes, and then click **Settings**.
4. In the **Settings** dialog box, change the settings as appropriate.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- There is a limit of 64 shadow copies per volume that can be stored. When this limit is reached, the oldest shadow copy will be deleted and cannot be retrieved.
- The following table lists the settings that you can configure.

| Setting | Description |
|---|---|
| **Storage area** | Specifies the volume where the shadow copies are stored. The default setting is the volume that contains the source files. It is recommended that you use a separate volume on another disk to provide better performance for heavily-used file servers.<br><br>You can only change the storage volume when there are no shadow copies present. If you need to change the storage volume for a volume that has already been enabled, you must first delete all of the shadow copies on that volume. |
| **Maximum size** | Specifies the maximum amount of space that will be used to store shadow copies on a particular volume. The default size is 10% of the size of the volume that contains the source files that are being shadow copied. If the shadow copies are stored on a separate volume from the source files, change this default to reflect the amount of the storage volume that you want to dedicate to shadow copies.<br><br>The storage limit must be at least 300 megabytes (MB); with this limit, only a single shadow copy can be stored. If you set a restrictive storage limit, verify that the number of shadow copies that you have scheduled will fit within those restraints. If shadow copies are deleted prematurely due to storage limits, you may be defeating the purpose of enabling Shadow Copies of Shared Folders. Click **Details** to compare the amount of space used on your storage volume with the maximum storage limit. |

| | |
|---|---|
| **Schedule** | Specifies settings for creating a schedule to take shadow copies on a regular basis. Schedule shadow copies at a time of day that will work best for your users. The default schedule is 7:00 A.M. and 12:00 noon, Monday through Friday. For more information on creating a schedule for Shadow Copies of Shared Folders, see "Defining a schedule for shadow copies" in Related Topics. |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Enable Shadow Copies for Shared Folders
Disable Shadow Copies for Shared Folders

# Change the schedule for creating shadow copies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the schedule for creating shadow copies**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**, select **All Tasks**, and then click **Configure Shadow Copies**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. Click the volume where you want to make changes to the schedule, and then click **Settings**.
4. In the **Settings** dialog box, click **Schedule**. Change the settings in the task scheduler as appropriate.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- The default schedule is 7:00 A.M. and 12:00 noon, Monday through Friday.
- It is recommended that you do not schedule copies to be taken more frequently than once an hour.
- There is a limit of 64 shadow copies per volume that can be stored. When this limit is reached, the oldest shadow copy will be deleted and cannot be retrieved.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Enable Shadow Copies for Shared Folders
Disable Shadow Copies for Shared Folders
Task Scheduler Concepts
Change the settings for Shadow Copies of Shared Folders
Best practices for Shadow Copies of Shared Folders

# Delete a shadow copy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a shadow copy**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**, select **All Tasks**, and then click **Configure Shadow Copies**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. In the **Shadow Copies** dialog box, click the shadow copy that you want to delete and click **Delete Now**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Shadow copies are not included in backups; once they are deleted, they cannot be restored.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Enable Shadow Copies for Shared Folders
Disable Shadow Copies for Shared Folders

# Revert an entire volume shadow copy to a previous version

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To revert an entire volume shadow copy to a previous version**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**, select **All Tasks**, and then click **Configure Shadow Copies**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

3. In the **Shadow Copies** dialog box, click the shadow copy that you want to revert to and click **Revert**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Shadow copies are not included in backups. Once they are deleted, they cannot be restored.
- You can use the **Revert** button to revert an entire volume back to an earlier point in time. Reverting the entire volume will undo any changes that were made to files and folders on the volume since the selected snapshot was taken. If you revert a volume, you cannot undo the action later. You should not use **Revert** on volumes that contain data stores like Active Directory or Exchange.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shadow Copies for Shared Folders How To...
Shadow Copies for Shared Folders overview

# Disable Shadow Copies for Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable Shadow Copies of Shared Folders**

1. Open Computer Management.
2. In the console tree, right-click **Shared Folders**, select **All Tasks**, and then click **Configure Shadow Copies**.

   **Where?**

   - Computer Management\System Tools\Shared Folders

3. Click the volume where you want to disable Shadow Copies of Shared Folders, and then click **Disable**.

### Caution

- Disabling Shadow Copies of Shared Folders will delete all existing shadow copies on the volume, as well as the scheduled task, if one exists.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Change the settings for Shadow Copies of Shared Folders
Deploying Shadow Copies for Shared Folders

# Shadow Copies for Shared Folders Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about creating shadow copies of shared folders:

- Shadow Copies for Shared Folders overview
- Deploying Shadow Copies for Shared Folders

# Shadow Copies for Shared Folders overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shadow Copies of Shared Folders overview**

Shadow Copies of Shared Folders provides point-in-time copies of files that are located on shared resources, such as a file server. With Shadow Copies of Shared Folders, you can view shared files and folders as they existed at points of time in the past. Accessing previous versions of your files, or shadow copies, is useful because you can:

- **Recover files that were accidentally deleted**. If you accidentally delete a file, you can open a previous version and copy it to a safe location.
- **Recover from accidentally overwriting a file**. If you accidentally overwrite a file, you can recover a previous version of the file.
- **Compare versions of file while working**. You can use previous versions when you want to check what has changed between two versions of a file.

You can access the server portion of Shadow Copies of Shared Folders through the **Shadow Copies** tab of the **Local Disk Properties** dialog box. For more information about how to use Shadow Copies of Shared Folders, see Shadow Copies for Shared Folders How To....

The client software for Shadow Copies of Shared Folders is located on the server and must be deployed to client computers. You can access the client view of shadow copies through the **Previous Versions** tab of the **Properties** dialog box of the shared file or folder. For more information about how to deploy Shadow Copies of Shared Folders, see Deploying Shadow Copies for Shared Folders.

**Notes**

- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.
- When you restore a file, the file permissions will not be changed. Permissions will remain the same as they were before the restore. When you recover a file that was accidently deleted, the file permissions will be set to the default permissions for the directory.
- For information on enabling and using shadow copies of shared folders in a server cluster, see Using Shadow Copies of Shared Folders in a server cluster.
- Creating shadow copies is not a replacement for creating regular backups.
- When storage area limits are reached, the oldest shadow copy will be deleted and cannot be retrieved.
- There is a limit of 64 shadow copies per volume that can be stored. When this limit is reached, the oldest shadow copy will be deleted and cannot be retrieved.
- Shadow copies are read-only. You cannot edit the contents of a shadow copy.
- You can only enable Shadow Copies of Shared Folders on a per-volume basis; that is, you cannot select specific shared folders and files on a volume to be copied or not copied.

# Deploying Shadow Copies for Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying Shadow Copies of Shared Folders**

This section provides general background information about Shadow Copies of Shared Folders:

- Determining storage options for shadow copies
- Defining a schedule for shadow copies
- Deploying the client software for shadow copies

# Determining storage options for shadow copies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Determining storage options for shadow copies**

Before deploying Shadow Copies of Shared Folders, it is recommended that you create a plan that specifies where the shadow copies will be stored and what the storage limits will be.

The minimum amount of storage space that you can specify is 300 megabytes (MB). The default storage size is 10% of the source volume (the volume being copied). If the shadow copies are stored on a separate volume, change the default to reflect the space available on the storage volume instead of the source volume. Remember that when the storage limit is reached, older versions of the shadow copies will be deleted and cannot be restored.

When determining the amount of space to allocate for storing shadow copies, you must consider both the number and size of files that are being copied, as well as the frequency of changes between copies. For example, 100 files that only change monthly require less storage space than 10 files that change daily.

Remember that to change the storage volume, you must delete the shadow copies. If you delete shadow copies, you lose the existing file change history that is kept on the original storage volume. To avoid this problem, verify that the storage volume that you initially select is large enough to handle your growing business needs.

**Caution**

- There is a limit of 64 shadow copies that can be stored on a volume. When this limit is reached, the oldest shadow copy is deleted and cannot be retrieved.

  For more information, see Change the settings for Shadow Copies of Shared Folders.

- If you are using a basic disk as a storage area for shadow copies and you intend to convert the disk into a dynamic disk, it is important to take the following precaution to avoid data loss. If the disk is a non-boot volume and is a different volume from where the original files reside, you must first dismount and take offline the volume containing the original files before you convert the disk containing shadow copies to a dynamic disk. You must bring the volume containing the original files back online within 20 minutes, otherwise, you will lose the data stored in the existing shadow copies. If the shadow copies are located on a boot volume, you can convert the disk to dynamic without losing shadow copies.

  You can use the **mountvol** command with the **/p** option to dismount the volume and take it offline. You can mount the volume and bring it online using the **mountvol** command or the Disk Management snap-in.

**Note**

- For information about enabling shadow copies in a server cluster, see Enable Shadow Copies for Shared Folders in a cluster. Do not use Cluster Administrator or cluster.exe to make changes to the settings or schedule for shadow copies of shared folders; use the shadow copies interface instead.

# Defining a schedule for shadow copies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Defining a schedule for shadow copies**

When you enable Shadow Copies of Shared Folders, a default schedule is also created. This schedule might work for your users, but it is important to consider the particular work habits of your users before defaulting to this schedule. You might want to consider the following questions:

- Do most of your users work from 9 A.M. to 5 P.M., Monday through Friday, in a single time zone?
- Are users spread out across time zones? Does this mean that you need more copies throughout the day to cover all of the work hours?
- How much work time can your users afford to lose? Will the default schedule of copies twice a day (7:00 A.M. and 12:00 noon, Monday through Friday) provide adequate coverage?
- How often can copies be taken without you needing to add additional storage?
- How many days of history can be retained before reaching the limit of 64 shadow copies?

You might want to create an initial schedule and deploy to a small group to test whether your scheduled shadow copies stay within the storage limits you set. Also, consider polling your users to find out their work habits and when they think that a shadow copy would be beneficial. This kind of information can help you determine the best schedule for your specific users.

**Note**

- It is recommended that you do not schedule copies to be taken more than once per hour.

For more information, see Change the schedule for creating shadow copies.

# Deploying the client software for shadow copies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying the client software for shadow copies**

The client software for Shadow Copies of Shared Folders is available from several locations:

- The Windows XP client is installed with Windows°XP Service Pack°2 (SP2).
- Windows 2000 (Service Pack 3 and later), Windows 98 Second Edition, and Windows XP clients are available at the Microsoft Web site.
- The client is also available on the server in the \\%systemroot%\system32\clients\twclient directory.

You can distribute the client software in a variety of ways. Consider the various options before deployment. There are several tools included in Windows Server 2003 operating systems, such as Group Policy, that can make deploying and maintaining the client software easier.

- We recommend that you place the client software on a shared resource and send an e-mail to users about what the client does and how to install it.

| **Example e-mail:** |
| --- |
| There is a new feature called Previous Versions that is available for employees who use shared folders for corporate documents. If you work with files that are located in shared folders on your network, you may be able to access previous versions of your files. Accessing previous versions of your files is useful because you can:<br>• **Recover files that were accidentally deleted.** If you accidentally delete a file, you can open a previous version and copy it to a safe location.<br>• **Recover from accidentally overwriting a file.** If you accidentally overwrite a file, you can recover a previous version of the file.<br>• **Compare versions of file while working.** You can use Previous Versions when you want to check what has changed between two versions of a file.<br>This feature has been enabled on the following file server: *server**sharedresource*. Copies are scheduled to be taken at 7:00 A.M. and 12:00 noon, Monday through Friday. Remember that these are copies of the files that were taken at these times. If you have worked on the file, you will lose all the work you did after the copy was taken. Saving your work frequently is still the best way to ensure your work is not lost.<br>To install the software for accessing previous versions on your desktop, go to: *server**sharedresource* and double-click twclient.msi. |

- Use Group Policy to assign the client to users within your network.

  For more information about deploying client software using Group Policy, see "Group Policy Software Deployment Background" at the Microsoft Windows Resource Kits Web site.

- Use Systems Management Server (SMS) or another software management application to deploy the client software. For more information about using Systems Management Server to deploy the client .msi package, see Deploying Windows Installer Setup Packages with SMS 2.0 at the Microsoft Web site.

**Note**

- The client software refers to shadow copies as "previous versions"; this was done to make the concept clearer for end

users.

For more information, see Accessing Previous Versions of Shared Files and Folders.

# Troubleshooting Shadow Copies for Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section provides problem-solving guidelines for deploying Shadow Copies of Shared Folders.

What problem are you having?

- The Properties dialog box does not contain a Previous Versions tab (for example, on the shared resource or local volume). I cannot the list of previous versions that have been created.
- The Event Log is filled with Event ID: 7001.
- Shadow copies are being deleted, but I still have plenty of storage space available.

**The Properties dialog box does not contain a Previous Versions tab (for example, on the shared resource or local volume). I cannot the list of previous versions that have been created.**
**Cause:** You cannot view previous versions of files on a local shared resource. You must have a working network connection and be viewing the volume as a shared resource.

**Solution:** Verify that your network connection is available. Then, create a connection to the shared resource on your local computer as if it were in another location on the network. For example: To view previous versions of files in the shared folder \Documents on the local computer \\Server1, you would need to create a network connection to \\Server1\Documents and access the previous versions as if they were on a remote computer. Right-click this network connection and select **Properties** to see the **Previous Versions** tab.

**The Event Log is filled with Event ID: 7001.**
**Cause:** A volume that was scheduled for shadow copies has been removed or is unavailable. The task scheduler is attempting the task and every time the task fails, it writes an error to the event log.

**Solution:** Delete the scheduled task using the Task Scheduler. For more information, see Remove a scheduled task.

If you are removing a volume that is being shadow copied, delete the scheduled task before deleting the volume so the event log is not filled with errors.

**Shadow copies are being deleted, but I still have plenty of storage space available.**
**Cause:** There is an upper limit of 64 shadow copies per volume that can be stored.

**Solution:** Adjust the schedule for taking shadow copies to a rate which retains a long enough history of the files without reaching the limit of 64 shadow copies.

**See also:** Defining a schedule for shadow copies; Determining storage options for shadow copies

# Accessing Previous Versions of Shared Files and Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing previous versions of shared files and folders**

With Previous Versions, you can access previous versions of files and folders that are shared on your network.

- For general background information, see Accessing previous versions overview.
- For help with specific tasks, see Previous Versions How To....
- For problem-solving information, see Troubleshooting Previous Versions.

# Accessing previous versions overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Accessing previous versions overview**

If you work with files that are located in shared folders on your network, you may be able to access previous versions of your files. Accessing previous versions of your files is useful because you can:

- **Recover files that were accidentally deleted**. If you accidentally delete a file, you can open a previous version and copy it to a safe location.
- **Recover from accidentally overwriting a file**. If you accidentally overwrite a file, you can recover a previous version of the file.
- **Compare versions of file while working**. You can use previous versions when you want to check what has changed between two versions of a file.

To access previous versions of files, your server administrator must enable Shadow Copies of Shared Folders on the file server. Shadow copies are copies of files that are located on the server and appear as previous versions. If you would like to have Shadow Copies of Shared Folders enabled, contact your administrator.

**Notes**

- For you to access previous versions of files, the shared files must be located on a computer running an operating system in the Microsoft® Windows Server 2003 family that has Shadow copies of Shared Folders enabled.
- Previous versions are read-only. You cannot make changes to a previous version of the file as it exists on the server.
- Previous versions are periodically deleted and can disappear at any time. If you think you might want a previous version of a file, it is recommended that you copy it to a safe place.
- Please contact your server administrator if you want to know when previous versions are created or deleted.

For more information, see Previous Versions How To....

# Previous Versions How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- View a previous version of a file
- Copy a previous version of a file to a new location
- Restore a previous version of a file
- Recover a file that was accidentally deleted

# View a previous version of a file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view a previous version of a file**

1. Locate the file that you want to view a previous version of (on the network), right-click the file, and click **Properties**. The **Properties** dialog box will appear.

2. On the **Previous Versions** tab, click the version you want to view and click **View**.

## Notes

- If the **Previous Versions** tab does not appear in the **Properties** dialog box, Shadow Copies of Shared Folders might not be enabled on the file server. Shadow copies are copies of files that are located on the server and appear as previous versions. Please contact your administrator for more information.

- If there are no previous versions listed, the file has not changed since the oldest copy was taken.

- Previous versions are read-only. You cannot make changes to the previous version of the file as it exists on the server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Accessing previous versions overview
Troubleshooting Previous Versions

# Copy a previous version of a file to a new location

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy a previous version of a file to a new location**

1. Locate the file that you want to copy (on the network), right-click the file, and click **Properties**. The **Properties** dialog box will appear.

2. On the **Previous Versions** tab, click the version of the file that you want to copy and click **Copy**.

3. In the **Copy Items** dialog box, click a new location to place the file and click **Copy**.

## Notes

- If the **Previous Versions** tab does not appear in the **Properties** dialog box, Shadow Copies of Shared Folders might not be enabled on the file server. Shadow copies are copies of files that are located on the server and appear as previous versions. Please contact your administrator for more information.

- If there are no previous versions listed, the file has not changed since the oldest copy was taken.

- When you use **Copy**, the file permissions will be set to the default permissions for the directory where the copy of the file is placed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Accessing previous versions overview
Troubleshooting Previous Versions

# Restore a previous version of a file

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore a previous version of a file**

1. Locate the file that you want to restore (on the network), right-click the file, and click **Properties**. The **Properties** dialog box will appear.

2. On the **Previous Versions** tab, click the version of the file that you want to restore and click **Restore**. A warning message about restoring a previous version will appear. Click **Yes** to complete the procedure.

## Caution

- Restoring a previous version will delete the current version. If you choose to restore a previous version of a folder, the folder will be restored to its state at the date and time of the version you selected. You will lose any changes that you have made to files in the folder since that time. If you do not want to delete the current version of a file or folder, use **Copy** to copy the previous version to a different location.

## Notes

- If the **Previous Versions** tab does not appear in the **Properties** dialog box, Shadow Copies of Shared Folders might not be enabled on the shared resource. Shadow copies are copies of files that are located on the server and appear as previous versions. Please contact your administrator for more information.

- If there are no previous versions listed, the file has not changed since the oldest copy was taken.

- When you restore a previous version of a folder, files in the current folder that were not contained in the previous version of the folder will not be erased.

- When you restore a file, the file permissions will not be changed; permissions will remain the same as they were before the restore. When you copy a previous version of a file, the permissions will be set to the default permissions for the directory where the copy of the file is placed.

- Restoring a large directory puts a heavy load on the file server and can result in previous versions being deleted. Try to restore individual files instead of entire directories whenever possible.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Accessing previous versions overview
Troubleshooting Previous Versions

# Recover a file that was accidentally deleted

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To recover a file that was accidentally deleted**

1. Locate the folder where the deleted file was stored (on the network), right-click the folder, and click **Properties**. The **Properties** dialog box will appear.

2. On the **Previous Versions** tab, double-click the most recent version of the folder that contains the file that you want to recover. A list of files that are stored in that previous version will appear.

3. Right-click the file that was accidentally deleted and click **Copy**.

4. Right-click the location where you want to place the file and click **Paste**.

## Notes

- If the **Previous Versions** tab does not appear in the **Properties** dialog box, Shadow Copies of Shared Folders might not be enabled on the shared resource. Shadow copies are copies of files that are located on the file server and appear as previous versions. Please contact your administrator for more information.

- When you copy a file, file permissions will be set to the default permissions for the directory where the copy of the file is placed.

- When you recover a file that was accidentally deleted, you are copying a previous version. You will lose any changes made to the file after the last previous version was created.

- To recover a deleted folder, go to a previous version of the parent folder to view and select the folder to copy.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Accessing previous versions overview
Troubleshooting Previous Versions

# Troubleshooting Previous Versions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
This section provides problem-solving guidelines for accessing previous versions of files and folders.

What problem are you having?

- I do not see a Previous Versions tab in the Properties dialog box.
- There are no files listed in the Previous Versions tab of the Properties dialog box.
- The previous version of the file I want to view is not listed anymore in the Previous Versions tab of the Properties dialog box.
- I am having problems viewing a previous version of a file.
- I cannot locate previous versions of files that are stored in my coworker's shared folder.

**I do not see a Previous Versions tab in the Properties dialog box.**
**Cause:** Shadow Copies of Shared Folders is not enabled on the file server.

**Solution:** Contact your administrator about enabling Shadow Copies of Shared Folders on the file server.

**There are no files listed in the Previous Versions tab of the Properties dialog box.**
**Cause:** The file has not changed.

**Solution:** The file must be changed to view a previous version.

**The previous version of the file I want to view is not listed anymore in the Previous Versions tab of the Properties dialog box.**
**Cause:** Previous versions are periodically deleted due to space limitations on the server.

**Solution:** Contact your administrator about providing space for more previous versions on the file server.

**I am having problems viewing a previous version of a file.**
**Cause:** Previous versions are read-only, and some applications do not work well with read-only files.

**Solution:** Use **Copy** instead to copy the file to another location and open it from there to view.

**I cannot locate previous versions of files that are stored in my coworker's shared folder.**
**Cause:** The shared folder must be located on a computer running an operating system in the Windows Server 2003 family that has shadow copies enabled.

**Solution:** Move the files you share with your coworker to a shared folder on a computer running an operating system in the Windows Server 2003 family that has shadow copies enabled.

# Disaster Recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disaster recovery**
Disaster recovery includes information about how to prepare for system problems, how to use options that can help start a system that otherwise would not start, and how to use the repair and recovery options.

- To safeguard your system, see Checklist: Safeguarding your system.
- To recover your system, see Checklist: Recovering a system that will not start.
- For tips about using disaster recovery, see Best practices for disaster recovery.
- For help with specific tasks, see Disaster Recovery How To....
- For general background information, see Disaster Recovery Concepts.
- For problem-solving instructions, see Troubleshooting Disaster Recovery.

# Checklist: Safeguarding your system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Safeguarding your system**

You can take a number of steps as precautions against the possibility of disk failure or other serious system problems.

| Step | Reference |
|---|---|
| Follow the best practices for security and keep your computer in a secure location. | Best practices for security |
| Specify what you want to do if the system stops unexpectedly. For example, you can specify that you want your computer to restart automatically. | Specify what happens when the system stops unexpectedly |
| Perform regular backups (including backups of your boot and system partitions and your System State). Keep your backup media in a safe place where it is not accessible to a malicious user. | Back up files to a file or a tape |
| Use the Automated System Recovery Preparation Wizard to create a backup of your system settings and your local system partition. Keep your Automated System Recovery (ASR) set in a safe place where it is not accessible to a malicious user. | Create an Automated System Recovery set using Backup |
| Keep the installation CD in a safe place where you can easily find it but it is not accessible to a malicious user. Then, if needed, you can start the computer from the CD and use the Recovery Console or Automated System Recovery. | Start the computer and use the Recovery Console<br><br>Recover from a system failure using Automated System Recovery |
| Install the Recovery Console as a startup option. | Install the Recovery Console as a startup option |

# Checklist: Recovering a system that will not start

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Recovering a system that will not start**

| Step | Reference |
|---|---|
| Try Safe Mode and related startup options, which you can use to start the system with only the minimal necessary services. | Startup options |
| If other startup options do not help, consider using the Recovery Console. This option is recommended only for advanced users or administrators. | Recovery Console overview<br><br>Start the computer and use the Recovery Console |
| If startup options and the Recovery Console do not work in your situation and if you made appropriate advance preparations, try using your original installation disk and an Automated System Recovery (ASR) set to restore your system. | Recover from a system failure using Automated System Recovery |

# Best practices for disaster recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Use Emergency Management Services if applicable.**

  To ensure access to the server regardless of the condition of the network drivers or operating system files, consider setting up Emergency Management Services, a new feature in the Windows Server 2003 family. With Emergency Management Services, you can remotely manage a server in emergency situations that would typically require a local keyboard, mouse, and monitor, such as when the network is unavailable or the server is not functioning properly. Emergency Management Services has specific hardware requirements, and is available only for products in the Windows Server 2003 family.

  For more information, see Emergency Management Services.

- **Create a plan for performing regular backup operations.**

  Review and incorporate the best practices for Backup into a plan for backing up all of your files on a regular basis.

  For more information, see Best practices for Backup.

- **Keep the installation CD where you can easily find it.**

  If needed, you can start the computer from the installation CD and use the Recovery Console or Automated System Recovery.

  For more information, see Start the computer and use the Recovery Console or Recover from a system failure using Automated System Recovery.

- **Install the Recovery Console as a startup option.**

  You can install the Recovery Console on your computer to make it available in case you are unable to restart Windows. You can then select the Recovery Console option from the list of available operating systems on startup. You cannot install the Recovery Console on an Itanium-based computer.

  For more information, see Install the Recovery Console as a startup option.

- **Specify startup and recovery options.**

  Specify what you want the operating system to do if the computer stops unexpectedly. For example, you can specify that you want your computer to restart automatically and you can control logging options.

  For information about these options, see Specify what happens when the system stops unexpectedly.

# Disaster Recovery How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Recover a System That Will Not Start
- Create an Automated System Recovery set using Backup
- Install the Recovery Console as a startup option
- Delete the Recovery Console

# Recover a System That Will Not Start

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recover a system that will not start**

- Start the computer in Safe Mode
- Start the computer using the last known good configuration
- Start the computer at a command prompt
- Start the computer and use the Recovery Console
- Recover from a system failure using Automated System Recovery

# Start the computer in Safe Mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the computer in Safe Mode**

1. Click **Start**, and then click **Shut Down**.

2. Select **Restart**, and then click **OK**.

3. Press F8 at the indicated time:

   - For an x86-based computer: When a screen of text appears and then disappears (the screen of text may include a memory test, lines about the BIOS, and other lines), press F8. There may be a prompt that tells you when to press F8.

   - For an Itanium architecture-based computer: After making your selection from the boot menu, press F8. There may be a prompt that tells you when to press F8.

   For information about configuring system settings that affect whether a prompt for F8 appears and how long it appears, see "To specify the default operating system for startup" in Related Topics at the bottom of this page.

4. Use the arrow keys to highlight the appropriate Safe Mode option, and then press ENTER.

   NUM LOCK must be off before the arrow keys on the numeric keypad will function.

5. Use the arrow keys to highlight an operating system, and then press ENTER.

## Notes

- In Safe Mode, Windows only uses basic files and drivers (that is, mouse, monitor, keyboard, mass storage, base video, default system services, and no network connections). You can choose the Safe Mode with Networking option, which loads all of the above files and drivers plus the essential services and drivers to start networking, or the Safe Mode with Command Prompt option, which is exactly the same as Safe Mode except that a command prompt is started instead of the user interface. You can also choose Last Known Good Configuration startup option, which starts your computer using registry information that was saved at the last shutdown.

- Safe Mode helps you diagnose problems. If a symptom does not reappear when you start in Safe Mode, you can eliminate the default settings and minimum device drivers as possible causes. If a newly added device or a changed driver is causing problems, you can use Safe Mode to remove the device or reverse the change.

- There are circumstances where Safe Mode will not be able to help you, such as when Windows system files required to start the system are corrupted or damaged. In this case, the Recovery Console may help you.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Startup options
Start the computer at a command prompt

Start the computer using the last known good configuration
Specify the default operating system for startup
Recovery Console overview
Repair overview

# Start the computer using the last known good configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the computer using the last known good configuration**

1. Click **Start**, and then click **Shut Down**.

2. Select **Restart**, and then click **OK**.

3. Press F8 at the indicated time:

   - For an x86-based computer: When a screen of text appears and then disappears (the screen of text may include a memory test, lines about the BIOS, and other lines), press F8. There may be a prompt that tells you when to press F8.

   - For an Itanium architecture-based computer: After making your selection from the boot menu, press F8. There may be a prompt that tells you when to press F8.

   For information about configuring system settings that affect whether a prompt for F8 appears and how long it appears, see "To specify the default operating system for startup" in Related Topics at the bottom of this page.

4. Use the arrow keys to select **Last Known Good Configuration**, and then press ENTER.

   NUM LOCK must be off before the arrow keys on the numeric keypad will function.

5. Use the arrow keys to highlight an operating system, and then press ENTER.

## Notes

- Choosing the Last Known Good Configuration startup option provides a way to recover from problems such as a newly added driver that may be incorrect for your hardware. It does not solve problems caused by corrupted or missing drivers or files.

- When you choose Last Known Good Configuration, only the information in registry key HKLM\System\CurrentControlSet is restored. Any changes you have made in other registry keys remain.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the computer in Safe Mode
Start the computer at a command prompt
Specify the default operating system for startup
Startup options
Repair overview

# Start the computer at a command prompt

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the computer at a command prompt**

1. Click **Start**, and then click **Shut Down.**

2. Select **Restart**, and then click **OK**.

3. Press F8 at the indicated time:

   - For an x86-based computer: When a screen of text appears and then disappears (the screen of text may include a memory test, lines about the BIOS, and other lines), press F8. There may be a prompt that tells you when to press F8.

   - For an Itanium architecture-based computer: After making your selection from the boot menu, press F8. There may be a prompt that tells you when to press F8.

   For information about configuring system settings that affect whether a prompt for F8 appears and how long it appears, see "To specify the default operating system for startup" in Related Topics at the bottom of this page.

4. Use the arrow keys to select **Safe Mode with Command Prompt**, and then press ENTER.

   NUM LOCK must be off before the arrow keys on the numeric keypad will function.

5. Use the arrow keys to highlight an operating system, and then press ENTER.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the computer in Safe Mode
Startup options
Specify the default operating system for startup
Start the computer using the last known good configuration
Repair overview

# Start the computer and use the Recovery Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the computer and use the Recovery Console**

1. Insert the Setup CD and restart the computer from the CD drive.
2. When the text-based part of Setup begins, follow the prompts; choose the repair or recover option by pressing **R**.
3. If you have a dual-boot or multiple-boot system, choose the installation that you need to access from the Recovery Console.
4. When prompted, type the password for the local administrator account.
5. At the system prompt, type the relevant Recovery Console commands; type **help** for a list of commands or **help***commandname* for help about a specific command.
6. To exit the Recovery Console and restart the computer, type **exit**.

**Important**

- Because the Recovery Console is quite powerful, it is recommended for use only by advanced users or administrators.
- In most cases, you will need the password for the built-in administrator account to use the Recovery Console.
- You can set the **Recovery console: Allow automatic administrative logon** setting in the Security Options section of Local Security Settings. Changing this setting will allow you to start the Recovery Console without having to provide the administrator password. For more information, see "Recovery console: Allow automatic administrative logon" and "Edit local security settings" in Related Topics at the bottom of this page.

  Be careful when enabling this option. This will allow anyone with physical access to a computer to boot into the Recovery Console and access the Recovery Console commands.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovery Console commands
Recovery Console overview
Repair overview
Edit local security settings
Recovery console: Allow automatic administrative logon

# Recover from a system failure using Automated System Recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To recover from a system failure using Automated System Recovery**

1. Make sure you have the following available before you begin the recovery procedure:

   - Your previously created Automated System Recovery (ASR) floppy disk.
   - Your previously created backup media.
   - The original operating system installation CD.
   - If you have a mass storage controller and you are aware that the manufacturer has supplied a separate driver file for it (different from driver files available on the Setup CD), obtain the file (on a floppy disk) before you begin this procedure.

2. Insert the original operating system installation CD into your CD drive.
3. Restart your computer. If you are prompted to press a key to start the computer from CD, press the appropriate key.
4. If you have a separate driver file as described in step 1, use the driver as part of Setup by pressing F6 when prompted.
5. Press F2 when prompted at the beginning of the text-only mode section of Setup. You will be prompted to insert the ASR floppy disk you have previously created.
6. Follow the directions on the screen.
7. If you have a separate driver file as described in step 1, press F6 (a second time) when prompted after the system reboots.
8. Follow the directions on the screen.

## Notes

- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.
- Automated System Recovery will not restore your data files. See Related Topics for more information on backing up and restoring your data files.
- If you are restoring a server cluster in which all nodes failed and the quorum disk cannot be restored from backup, use Automated System Recovery on each node in the original cluster to restore the disk signatures and the partition layout of the cluster disks (quorum and nonquorum). For more information about backing up and restoring server clusters, see the Related Topics link at the bottom of this page.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Create an Automated System Recovery set using Backup

Back Up and Restore Server Clusters

Backup overview

# Create an Automated System Recovery set using Backup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create an Automated System Recovery set using Backup**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use the Backup or Restore Wizard to create an Automated System Recovery (ASR) set by answering **All information on this computer** in the **What do you want to backup?** section. Otherwise, you can go to the next step to create an ASR set in **Advanced Mode**.

2. Click the Advanced Mode link on the Backup or Restore Wizard.

3. On the **Tools** menu, click **ASR Wizard**.

4. Follow the instructions that appear on your screen.

### Important

- You will need a blank 1.44 MB floppy disk to save your system settings and media to contain the backup files. If your computer does not have a floppy drive, see Backup Troubleshooting.

### Notes

- You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- This procedure will only back up those system files that are necessary for starting your system. You will need to backup your data separately.

- After creating the ASR set, label this floppy disk and backup media carefully and keep them together. To use the backup media, you must have the floppy disk that you created with that set of media. You cannot use a floppy disk created at a different time or with a different set of media. You must also have your Setup CD available at the time you perform Automated System Recovery.

- Keep the ASR set in a secure location. The ASR set contains information about your systems configuration that could be used to do harm to your system.

- If you are backing up a server cluster, run the Automated System Recovery Preparation Wizard on all nodes of the cluster and make sure that the Cluster service is running when you start each Automated System Recovery backup. Make sure that one of the nodes on which you run the Automated System Recovery Preparation Wizard is listed as the owner of the quorum resource while the wizard is running.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Recover from a system failure using Automated System Recovery
Backing up and restoring server clusters

# Install the Recovery Console as a startup option

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install the Recovery Console as a startup option**

1. With Windows running, insert the Setup CD into your CD-ROM drive.
2. Click **Start** and select **Run**.
3. Type the following where *D*: is the CD-ROM drive letter:

   On 32-bit versions of Windows Server 2003: `D:\i386\winnt32.exe /cmdcons`

   - or -

   On x64-based versions of Windows Server 2003: `D:\amd64\winnt32.exe /cmdcons`
4. Follow the instructions on the screen.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- You can not install the Recovery Console on Itanium-based computers.
- To run the Recovery Console, you must restart your computer and select the Recovery Console option from the list of available operating systems.
- To see the commands available on the Recovery Console, type **help** at the console prompt.
- If your computer will not start or you are using an Itanium-based computer, you can run the Recovery Console from the Setup CD. See Start the computer and use the Recovery Console for information on running the Recovery Console when your computer will not start.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovery Console commands
Recovery Console overview
Delete the Recovery Console
Start the computer and use the Recovery Console
Working with MMC console files

# Delete the Recovery Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete the Recovery Console**

1. Open My Computer.
2. Double-click the hard drive on which you installed the Recovery Console.
3. On the **Tools** menu, click **Folder Options**.
4. Click the **View** tab.
5. Click **Show hidden files and folders**, clear the **Hide protected operating system files** check box, and then click **OK**. If you receive a warning about displaying protected system files, click **Yes** to continue.
6. At the root directory, delete the \Cmdcons folder.
7. At the root directory, delete the file Cmldr.
8. Click **Start**, then **Control Panel**, then click **System**.
9. On the **Advanced** tab under **Startup and Recovery**, select **Settings**.
10. On the **Startup and Recovery** dialog box, under **System startup**, select **Edit** to edit the startup options file manually. Boot.ini will display in Notepad.
11. Remove the entry for the Recovery Console. It will look similar to this:

    C:\cmdcons\bootsect.dat="Microsoft Windows Recovery Console" /cmdcons

12. Save the file and close it.

### Caution

- Modifying the Boot.ini file incorrectly may prevent your computer from restarting. Be sure to delete only the entry for the Recovery Console.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- You can only install the Recovery Console feature on x86-based computers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Recovery Console overview
Repair overview

# Disaster Recovery Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

- Repair overview
- Understanding Disaster Recovery
- Recovery Console commands

# Repair overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Repair overview**
The Microsoft® Windows Server 2003 family includes several features to ensure that your computer and the applications and devices installed on it work correctly. These features help you solve the problems that can result from adding, deleting, or replacing files that your operating system, applications, and devices require to function.

**Important**

- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.

The recovery feature or features that you use depend on the type of problem or failure you encounter. Use the following list to determine which features to use to address a specific problem.

- Your personal data files are missing or have been corrupted, or you want to revert to a previous version of a file.
- After updating a device driver, you can log on, but the system is unstable.
- After installing a new device, the system is unstable.
- After installing an application, the system is unstable or an application is not working properly.
- Operating system does not start (the logon screen does not appear).

**Your personal data files are missing or have been corrupted, or you want to revert to a previous version of a file.**
If you have saved a copy of the data file in another location, you can copy it from that location back to your hard disk. You can make a back-up copy at any time by copying the file to another location on your computer, to external storage such as a removable disk, or to a different computer.

**Feature: Backup Utility**

**When to use it:** When you need to restore a version of a data file that you saved using Backup.

**What it does:** Restores files from backups.

**For more information**, see Restore files from a file or a tape.

**After updating a device driver, you can log on, but the system is unstable.**

**Feature: Device Driver Roll Back**

**When to use it:** When the only change you want to undo is the updating of a device driver (other than a printer driver). You must be logged on as an administrator.

**What it does:** Reinstalls the driver you were using previously and restores any driver settings that were changed when you added the new driver. Affects no other files or settings. Printer drivers cannot be restored with Device Driver Roll Back.

**For more information**, see Roll back to the previous version of a driver.

**After installing a new device, the system is unstable.**

**Feature: Disable the Device**

**When to use it:** When you suspect that one or more specific hardware devices are causing the problem. You must be logged on as an administrator.

**What it does:** Disables the hardware device and its drivers.

**For more information**, see Disable a device.

**After installing an application, the system is unstable or an application is not working properly.**
Under some circumstances, the computer will prompt you to repair or reinstall an application from a network location or from

the original set-up media (for example, the program CD). If this happens, follow the directions on the screen to repair your program. If you are not prompted with this information, you can use one of the following features to repair your program:

**Feature: Add or Remove Programs**

**When to use it:** When you suspect that one or more specific programs are causing the problem.

**What it does:** Removes the programs you specify.

**For more information**, see Windows interface administrative tool reference A-Z: Add or Remove Programs.

**Operating system does not start (the logon screen does not appear).**

**Feature: Last Known Good Configuration startup option**

**When to use it:** When you suspect that a change you made to your computer before restarting might be causing the failure.

**What it does:** Restores the registry settings and drivers that were in effect the last time the computer started successfully.

**For more information**, see Start the computer using the last known good configuration.

**Feature: Recovery Console**

**When to use it:** If using the Last Known Good Configuration startup option is unsuccessful and you cannot start the computer in Safe Mode. This method is recommended only if you are an advanced user who can use basic commands to identify and locate problem drivers and files. To use the Recovery Console, restart the computer with the installation CD for the operating system in the CD drive. When prompted during text-mode setup, press R to start the Recovery Console.

**What it does:** From the Recovery Console, you can access the drives on your computer. You can then make any of the following changes so that you can start your computer:

- Enable or disable device drivers or services.
- Copy files from the installation CD for the operating system, or copy files from other removable media. For example, you can copy an essential file that had been deleted.
- Create a new boot sector and new master boot record (MBR). You might need to do this if there are problems starting from the existing boot sector.
-

**For more information**, see Recovery Console overview.

**Feature: Automated System Recovery**

**When to use it:** After all other attempts at recovery fail, or if you have replaced a damaged system hard drive. You must have previously used Backup to create an Automated System Recovery set.

**What it does:** Automated System Recovery restores all disk signatures, volumes, and partitions on the disks required to start the computer. Automated System Recovery then installs a simplified installation of Windows and automatically starts a restoration using the backup created by the Automated System Recovery Preparation Wizard.

**For more information**, see Automated System Recovery (ASR) overview.

**Feature: Windows installation CD**

**When to use it:** When you cannot start the computer in Safe Mode, you cannot use either the Last Known Good Configuration startup option or the Recovery Console successfully, and you have no backup to use with Automated System Recovery.

**What it does:** Reinstalls the operating system. You will then need to reinstall your applications and restore your data files from backups.

# Understanding Disaster Recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding disaster recovery**
This section covers:

- Startup options
- Recovery Console overview
- Automated System Recovery (ASR) overview

# Startup options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Startup options**
If your computer will not start, you may be able to start it in Safe Mode. In Safe Mode, Windows uses default settings (VGA monitor, Microsoft mouse driver, no network connections, and the minimum device drivers required to start Windows).

For example, if your computer will not start after you install new software, you may be able to start it with minimal services in Safe Mode and then change your computer settings or remove the newly installed software that is causing the problem. You can reinstall the service pack or the entire operating system, if necessary.

If a symptom does not reappear when you start in Safe Mode, you can eliminate the default settings and minimum device drivers as possible causes.

**Caution**

- The startup options can potentially be used to perform malicious actions on your computer. To protect yourself from this kind of abuse, make sure your computer is in a secure location at all times.

The startup options are:

- Safe Mode
- Safe Mode with Networking
- Safe Mode with Command Prompt
- Enable Boot Logging
- Enable VGA Mode
- Last Known Good Configuration
- Directory Service Restore Mode
- Debugging Mode
- Disable Automatic Restart On System Failure

**Safe Mode**
Starts using only basic files and drivers (mouse, except serial mice; monitor; keyboard; mass storage; base video; default system services; and no network connections). If your computer does not start successfully using Safe Mode, you may need to use the Recovery Console feature to repair your system.

**Safe Mode with Networking**
Starts using only basic files and drivers, plus network connections.

**Safe Mode with Command Prompt**
Starts using only basic files and drivers. After logging on, the command prompt is displayed instead of the Windows desktop, **Start** menu, and Taskbar.

**Enable Boot Logging**
Starts while logging all the drivers and services that were loaded (or not loaded) by the system to a file. This file is called ntbtlog.txt and it is located in the %systemroot% directory. Safe Mode, Safe Mode with Networking, and Safe Mode with Command Prompt add to the boot log a list of all the drivers and services that are loaded. The boot log is useful in determining the exact cause of system startup problems.

**Enable VGA Mode**
Starts the computer using currently installed video driver at the lowest resolution. The basic video driver is always used when you start in Safe Mode (either Safe Mode, Safe Mode with Networking, or Safe Mode with Command Prompt).

**Last Known Good Configuration**
Starts using the registry information and drivers that Windows saved at the last logon. Any driver settings or other systems setting changes made since the last successful logon will be lost. Use only in cases of incorrect configuration.

**Directory Service Restore Mode**
This option is for the server operating systems and is only used in restoring the SYSVOL directory and the Active Directory directory service on a domain controller.

**Debugging Mode**
Starts while sending debugging information through a serial cable to another computer. The serial cable must be connected to the COM1 port with the baud rate set at 115,200.

If you are using, or have used, Remote Installation Services to install Windows on your computer, you may see additional options related to restoring or recovering your system using Remote Installation Services.

**Disable Automatic Restart on System Failure**
Specifies that the computer will not automatically restart whenever the system unexpectedly stops. Use this option to keep your system from restarting while you troubleshoot system failures. Starting with this option is the same as opening the **System** control panel, selecting the **Advanced** tab, clicking **Settings** under **Startup and Recovery Settings** and on the **Startup and Recovery** dialog box, clearing the **Automatically restart** check box under **System Failure**.

### Important

- Even if you disable the local administrator account, it will still be available when you use Safe Mode.
- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.

# Recovery Console overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovery Console overview**
If Safe Mode and other startup options do not work, consider using the Recovery Console. This method is recommended only if you are an advanced user who can use basic commands to identify and locate problem drivers and files. In addition, you will need the password for the built-in administrator account to use the Recovery Console.

Using the Recovery Console, you can enable and disable services, format drives, read and write data on a local drive (including drives formatted to use NTFS), and perform many other administrative tasks. The Recovery Console is particularly useful if you need to repair your system by copying a file from a floppy disk or CD-ROM to your hard drive, or if you need to reconfigure a service that is preventing your computer from starting properly.

**Important**

- Physical access to a server is a high security risk. To maintain a more secure environment, you must restrict physical access to all servers and network hardware.

There are two ways to start the Recovery Console:

- If you are unable to start your computer, you can run the Recovery Console from your Setup CD.
- As an alternative, you can install the Recovery Console on your computer to make it available in case you are unable to restart Windows. You can then select the Recovery Console option from the list of available operating systems on when you start your computer. You cannot install the Recovery Console on an Itanium-based computer.

After you start the Recovery Console, you will have to choose which installation you want to log on to (if you have a dual-boot or multiple-boot system), and you will have to log on with the password for the built in administrator account for that installation.

The console provides commands you can use to do simple operations such as changing to a different directory or viewing a directory, as well as more powerful operations such as fixing the boot sector. You can access Help for the commands in the Recovery Console by typing **help** at the Recovery Console command prompt.

For more information about securing your computer, see Best practices for security.

**Note**

- Installation of the Recovery Console provides an alternate method for logging on to the computer. This could potentially provide a method for an intruder to gain access to the computer by circumventing many of the built-in safeguards. Consider the potential security effects before installing the Recovery Console.

# Automated System Recovery (ASR) overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Automated System Recovery (ASR) overview**
With Automated System Recovery (ASR), you can create ASR sets on a regular basis as part of an overall plan for system recovery in case of system failure. Use ASR as a last resort in system recovery, only after you have exhausted other options such as the startup options Safe Mode and Last Known Good Configuration. For more information about these and other recovery options, see Startup options.

ASR is a recovery option that has two parts: ASR backup and ASR restore. You can access the backup portion through the Automated System Recovery Preparation Wizard located in Backup. The Automated System Recovery Preparation Wizard backs up the System State data, system services, and all disks associated with the operating system components. It also creates a floppy disk, which contains information about the backup, the disk configurations (including basic and dynamic volumes), and how to accomplish a restore.

You can access the restore part of ASR by pressing F2 when prompted in the text mode portion of setup. ASR reads the disk configurations from the floppy disk and restores all of the disk signatures, volumes and partitions on the disks required to start your computer (at a minimum). (It will attempt to restore all of the disk configurations, but under some circumstances, it may not be able to). ASR then installs a simple installation of Windows and automatically starts to restore from backup using the backup ASR set created by the Automated System Recovery Preparation Wizard.

**Notes**

- ASR does not include data files. Back up data files separately on a regular basis and restore them after the system is working.
- Using **All information on this computer** in the simple wizard also creates an ASR floppy disk and an ASR set.
- ASR supports FAT16 volumes up to 2.1 GB only. ASR does not support 4 GB FAT16 partitions that use a cluster size of 64K. If your system contains 4-GB FAT16 partitions, convert them from FAT16 to NTFS before using ASR.

# Recovery Console commands

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Recovery Console commands**
The following commands can be used with the Recovery Console. For more information about each command, type the command name followed by **/?** in the Recovery Console.

For more information about the Recovery Console, see "Tools for Troubleshooting" at the Microsoft Windows Resource Kits Web site.

- **Attrib**

  Changes the attributes of a file or directory.

- **Batch**

  Executes the commands specified in the text file.

- **Bootcfg**

  Boot file (boot.ini) configuration and recovery.

- **ChDir (Cd)**

  Displays the name of the current directory or changes the current directory.

- **Chkdsk**

  Checks a disk and displays a status report.

- **Cls**

  Clears the screen.

- **Copy**

  Copies a single file to another location.

- **Delete (Del)**

  Deletes one or more files.

- **Dir**

  Displays a list of files and subdirectories in a directory.

- **Disable**

  Disables a system service or a device driver.

- **Diskpart**

  Manages partitions on your hard drives.

- **Enable**

Starts or enables a system service or a device driver.

- **Exit**

  Exits the Recovery Console and restarts your computer.

- **Expand**

  Extracts a file from a compressed file.

- **Fixboot**

  Writes a new partition boot sector onto the system partition.

- **Fixmbr**

  Repairs the master boot record of the partition boot sector.

- **Format**

  Formats a disk.

- **Help**

  Displays a list of the commands you can use in the Recovery Console.

- **Listsvc**

  Lists the services and drivers available on the computer.

- **Logon**

  Logs on to an Windows operating system installation.

- **Map**

  Displays the drive letter mappings.

- **Mkdir (Md)**

  Creates a directory.

- **More (Type)**

  Displays a text file.

- **Rename (Ren)**

  Renames a single file.

- **Rmdir (Rd)**

  Deletes a directory.

- **Set**

Displays and sets environment variables.

- **Systemroot**

  Sets the current directory to the systemroot directory of the system you are currently logged on to.

# Troubleshooting Disaster Recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

**Important**

- For more information about resolving computer issues, see Repair overview.

What problem are you having?

- My computer stops and displays a long message beginning with the word STOP.
- My computer supports starting directly from the CD-ROM drive. However, when I try to start from the setup compact disc, nothing happens.

**My computer stops and displays a long message beginning with the word STOP.**
**Cause:** Windows has encountered an error from which it cannot recover, and to prevent file damage, has stopped operating.

**Solution:** For general instructions on troubleshooting STOP messages, see Troubleshooting Stop messages: general strategies. For specific instructions on troubleshooting STOP messages, see Troubleshooting specific Stop messages.

**My computer supports starting directly from the CD-ROM drive. However, when I try to start from the setup compact disc, nothing happens.**
**Cause:** To start your computer directly from the setup compact disc, your computer's BIOS or EFI must support the El Torito Bootable CD-ROM (no emulation mode) format. Your BIOS or EFI may not support that format.

**Solution:** Check with your computer manufacturer if you are unsure whether your computer supports this feature. The BIOS or EFI specifies the boot order. Make sure that the CD-ROM is listed as the first boot device.

# Managing Data Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing data storage**
This section covers:

- Removable Storage
- Remote Storage

# Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Removable Storage**
Removable Storage makes it easy for you to track your removable storage media (tapes and optical disks) and to manage the libraries, such as changers and jukeboxes, that contain them.

- Before you begin using Removable Storage, see the Checklist: Configure Removable Storage.
- For tips about using Removable Storage, see Best practices for Removable Storage.
- For step-by-step help with specific tasks, see Removable Storage How To....
- For background information about Removable Storage and how it works, see Removable Storage Concepts.
- For problem-solving guidelines, see Troubleshooting Removable Storage.

# Checklist: Configure Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Configure Removable Storage**

| Step | Reference |
| --- | --- |
| Verify that you have a supported library. | See the hardware compatibility information in the Windows Catalog. You can access the Windows Catalog from Support resources. |
| Install and test the library. | Read the library's operation guide and follow the manufacturer's recommendations. Pay particular attention to slot numbering and supported media types, and note whether the drives require periodic cleaning. |
| Create new media pools to support applications, as needed. | Create a new media pool. |
| Move sufficient media to a free media pool. | Move media to another media pool. |
| Set user or group permissions for Removable Storage as needed. | Change user permissions for Removable Storage. |

# Best practices for Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Use Run as to run Removable Storage as an Administrator**.

  To maintain system security, log in as a member of the Users group and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking Run as, and choosing an account in the local Administrators group.

- **Keep tapes and disks secure**.

  A malicious user could use Removable Storage tapes and disks to gain access to sensitive data. To maintain security, keep tapes and disks in a secure location when not in use, and secure your libraries by assigning them sufficient permissions to prevent unauthorized users from tampering with or ejecting the tapes and disks.

- **Verify your library hardware is supported first before using Removable Storage**.

  Before adding a new library to a Removable Storage system, you should review the hardware compatibility information in the Windows Catalog. You can access the Windows Catalog from Support resources.

- **Use the Cleaner Management Wizard to clean tape drives**.

  If your library's drives support cleaning, inject cleaning tapes using the Cleaner Management Wizard.

- **Insert media into an automated library before use**.

  Before you create a new media pool, insert media into each automated library you want Removable Storage to control. If a library is empty, Removable Storage cannot detect the media type.

- **Keep your media count under 1,000 per system installation**.

  For performance reasons, do not manage more than 1,000 tapes or disks in a single Removable Storage system.

- **Turn on operator requests only when necessary**.

  Turning on operator requests enables programs on a remote computer to redirect and display operator requests on your computer. Turn off operator requests when it is not needed.

  You can turn on or turn off operator requests by selecting or clearing the **Display operator request and progress dialogs** and **Use notification area icon for pending operator requests** check boxes as described in Change how operator requests are displayed.

- **Give users the appropriate level of access**.

  The following table shows the media pool permissions applicable to specific types of users depending on their location. The Control permission on the Root pool enables users to create media pools. In addition, to create sub-pools, users must have the Control permission on the parent pool. Only those users who must have the ability of creating media pools should have Control permissions added to the Root pool.

| Location | User Type | Use | Control | Modify | View Permissions | Modify Permissions |
|---|---|---|---|---|---|---|
| Computer | User | X | | | | |

| | | | | | | |
|---|---|---|---|---|---|---|
| | Power User | X | X | | | |
| | Backup Operator | X | X | | | |
| | System | X | X | X | X | X |
| Media Library | User | X | | | | |
| | Power User | X | X | | | |
| | Backup Operator | X | X | | | |
| | System | X | X | X | X | X |
| Root Pool | User | | | | | |
| | Power User | | X | | | |
| | Backup Operator | | X | | | |
| | System | | X | | | |
| System Pool | User | X | | | | |
| | Power User | X | X | | | |
| | Backup Operator | X | X | | | |
| | System | X | X | X | X | X |

# Removable Storage How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- View or set Removable Storage properties
- Change the way Removable Storage displays
- Manage Libraries
- Manage Media Pools
- Manage Tapes and Disks
- Manage Operator Requests and Queued Work
- Manage Security for Removable Storage

# View or set Removable Storage properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view or set Removable Storage properties**

1. Open Removable Storage.
2. In the console tree, right-click **Removable Storage**, and then click **Properties**.
3. In the **Properties** dialog box, review the settings and make any changes.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- To view or set properties for specific media pools, library drives, media, or other items, click the applicable folder in the console tree, right-click the specific item in the details pane, and then click **Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing libraries
Managing media pools
Managing tapes and disks
Managing operator requests and queued work
Managing security for Removable Storage
Working with MMC console files

# Change the way Removable Storage displays

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the way Removable Storage displays**

1. Open Removable Storage.
2. Verify that **Removable Storage** is selected in the console tree.
3. On the **View** menu, click **Full** (for the full view) or **Simple** (for a simplified view).

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- The Full view is the default view for the Windows Server 2003 family.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Removable Storage How To...

# Manage Libraries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage libraries**

- Enable or disable a library
- Enable or disable a drive
- Change the library media type
- Inventory a library
- Set the default inventory method
- Open a library door
- Set the library door and port time-out
- Clean a stand-alone drive
- Clean an automated library

# Enable or disable a library

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable a library**

1. Open Removable Storage.
2. In the console tree, double-click **Libraries**.
3. Right-click the library you want to enable or disable, and then click **Properties**.

   The library **Properties** dialog box appears.
4. On the **General** tab, verify that the **Enable library** check box is selected.

   Or, to disable the library, clear the **Enable library** check box.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- Disabling a library makes it unavailable for use.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable a drive
Enable or disable a tape or disk
Working with MMC console files

# Enable or disable a drive

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable a drive**

1. Open Removable Storage.
2. In the console tree, click **Drives**.

   **Where?**

   - Removable Storage *(computer name)*/Libraries/*applicable library*/Drives

3. In the details pane, right-click the applicable drive, and then click **Properties**.
4. On the **General** tab, verify that the **Enable drive** check box is selected.

   Or, to disable the drive, clear the **Enable drive** check box.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- Disabling a drive makes it unavailable for use.
- This procedure applies only to automated libraries. Since stand-alone drive libraries consist of a single drive, you enable or disable the entire library for such cases. For information on how to enable or disable a library, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable a library
Enable or disable a tape or disk
Working with MMC console files

# Change the library media type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the library media type**

1. Open Removable Storage.
2. In the console tree, double-click **Libraries**.
3. Right-click the library you want to change, and then click **Properties**.
4. In the *library***Properties** dialog box, click the **Media** tab, and then click **Change**.
5. In the **Change Media Types** dialog box, do the following:
   - To add a new media type, in **Available types**, click the entry you want to add, and then click **Add**.
   - To remove an existing media type, in **Selected types**, click the entry you want to remove, and then click **Remove**.
   - Repeat these steps as necessary.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing tapes and disks
Working with MMC console files

# Inventory a library

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To inventory a library**

1. Open Removable Storage.
2. In the console tree, double-click **Libraries**.
3. Right-click the library you want to inventory, and then click **Inventory**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- This procedure applies only to automated libraries.
- The type of inventory performed will be the default for that library.
- An inventory is automatically performed each time a library door has been accessed.
- The inventory is performed according to the default inventory method: either Fast, Full, or None. For information on how to change the default inventory method, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the default inventory method
Working with MMC console files

# Set the default inventory method

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the default inventory method**

1. Open Removable Storage.
2. In the console tree, double-click **Libraries**.
3. Right-click the library you want to set an inventory method for, and then click **Properties**.
4. On the **General** tab, in **Inventory Method**, click the method you want to use: **None**, **Fast**, or **Full**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- A fast inventory only checks for storage slots that have changed status between being occupied and unoccupied. A full inventory with bar-coded media will read bar-code information. Otherwise, the tape or disk is mounted and the on-media identifier is read.
- The selected inventory method is automatically performed each time a library door has been accessed.
- A full inventory is always performed when a mount operation fails, regardless of inventory method selected. To disable this feature, clear the **Perform full inventory on mount failure** check box.
- This procedure applies only to automated libraries.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the default inventory method
Inventory a library
Working with MMC console files

# Open a library door

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To open a library door**

1. Open Removable Storage.
2. In the console tree, double-click **Libraries**.
3. Right-click the applicable library, and then click **Door Access**.
4. When prompted, open the door.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- You can use the library door to insert or remove media. If the library has an inject/eject port, you can use that to prevent a re-inventory of the media. Make sure you use the correct slots in the library.
- If you exchange media in a library that are not bar-coded, you must perform a full inventory afterward. The media exchange is not detected during a fast inventory and a media ID mismatch will occur the next time the media in the applicable slot is mounted.
- This procedure applies only to automated libraries.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the library door and port time-out
Working with MMC console files

# Set the library door and port time-out

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the library door and port time-out**

1. Open Removable Storage.
2. In the console tree, double-click **Libraries**.
3. Right-click the applicable library, and then click **Properties**.
4. In the **Properties** dialog box, click the **Components** tab.
5. Under **Doors**, in **Time-out**, enter a new value (if necessary) for the door time-out interval (in minutes).
6. Under **Ports**, in **Time-out**, enter a new value (if necessary) for the inject/eject port time-out interval (in minutes).

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- This procedure applies only to automated libraries.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Open a library door
Working with MMC console files

# Clean a stand-alone drive

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To clean a stand-alone drive**

1. Open Removable Storage.
2. In the console tree, click **Drives**.

   **Where?**

   - Removable Storage/Libraries/*applicable library*/Drives

3. In the details pane, right-click the drive, and then click **Mark as Clean**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- To clean a stand-alone drive, you manually insert a cleaning cartridge into the stand-alone drive. This procedure is performed only to notify Removable Storage that a stand-alone drive has been manually cleaned so that a record can be kept in the Removable Storage database.
- This procedure applies only to cleaning the drive in a stand-alone drive library. For information on how to clean drives in an automated library, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Clean an automated library
Working with MMC console files

# Clean an automated library

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To clean an automated library**

1. Open Removable Storage.
2. In the console tree, right-click the library you want to clean.

   **Where?**

   - Removable Storage *(computer name)*/Libraries/*applicable library*

3. Click **Cleaner Management**.
4. Follow the instructions in the Cleaner Management wizard.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- To complete this procedure, you must insert a cleaning cartridge into the library when prompted. Do not insert a cleaning cartridge using the same method used for inserting regular tapes (such as through the insert/eject port). Follow the instructions in the Cleaner Management wizard.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Clean a stand-alone drive
Working with MMC console files

# Manage Media Pools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage media pools**

- Create a new media pool
- Delete an application media pool
- Automatically draw media from a free media pool
- Automatically return media to a free media pool
- Set the allocation limit for a media pool

# Create a new media pool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new media pool**

1. Open Removable Storage.

2. If Removable Storage is in Full view, in the console tree, right-click **Media Pools**, and then click **Create Media Pool**. Or, to create a new media pool within another media pool, right-click the applicable media pool and then click **Create Media Pool**.

   If Removable Storage is in Simple view, on the **View** menu, click **Full**, right-click **Media Pools**, and then click **Create Media Pool**. Or, to create a new media pool within another media pool, on the **View** menu, click **Full**, right-click the applicable media pool and then click **Create Media Pool**.

3. On the **General** tab, in **Name**, type a name for the new media pool; then, in **Description**, type a relevant description.

4. Under **Media Information**, click **Contains media of type**, and select the appropriate media type from the list.

5. Under **Allocation/Deallocation policy**, do one or more of the following:

   - To automatically draw unused media from a free media pool when needed, select the **Draw media from Free media pool** check box.

   - To automatically return media to a free media pool when no longer needed, select the **Return media to Free media pool** check box.

   - To set an allocation limit for media in this media pool, select the **Limit reallocations** check box, and then change the default value as necessary.

6. On the **Security** tab, if the user(s) is not listed in the **Group or user names** box, click **Add**, and then complete the **Select Users, Computers, or Groups** dialog box.

7. In the **Permissions for***user name* box, select the **Allow** check box for the Use, Control, and Modify permissions.

**Notes**

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.

- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.

- You can create media pools within application media pools only. You cannot create new media pools within free, unrecognized, or import media pools.

- To create a media pool or a sub-pool beneath the parent media pool, a user must have Control permissions on the Root pool. The Control permission enables users to create media pools.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete an application media pool
Change the way Removable Storage displays
Working with MMC console files

# Delete an application media pool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete an application media pool**

1. Open Removable Storage.
2. If Removable Storage is in Full view, in the console tree, double-click **Media Pools**.

   If Removable Storage is in Simple view, on the **View** menu, click **Full**, and then double-click **Media Pools**.
3. Right-click the application media pool you want to delete, and then click **Delete**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- You cannot delete free, import, or unrecognized media pools.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a new media pool
Change the way Removable Storage displays
Working with MMC console files

# Automatically draw media from a free media pool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To automatically draw media from a free media pool**

1. Open Removable Storage.
2. If Removable Storage is in Full view, in the console tree, double-click **Media Pools**.

   If Removable Storage is in Simple view, on the **View** menu, click **Full**, and then double-click **Media Pools**.
3. Right-click the applicable media pool, and then click **Properties**.
4. On the **General** tab, select the **Draw media from Free media pool** check box.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- If you do not select the **Draw media from Free media pool** check box, you must manually move tapes or disks from a free media pool whenever this media pool needs additional media. For information on how to manually move media to a free media pool, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Move media to another media pool
Automatically return media to a free media pool
Create a new media pool
Change the way Removable Storage displays
Working with MMC console files

# Automatically return media to a free media pool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To automatically return media to a free media pool**

1. Open Removable Storage.
2. If Removable Storage is in Full view, in the console tree, double-click **Media Pools**.
3. Right-click the applicable media pool, and then click **Properties**.

   If Removable Storage is in Simple view, on the **View** menu, click **Full**, and then double-click **Media Pools**.
4. On the **General** tab, select the **Return media to Free media pool** check box.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- If you do not select the **Return media to Free media pool** check box, you must manually move tapes or disks in this application media pool to a free media pool when no longer needed. After you do this, another application can use those tapes or disks. For information on how to manually move media to a free media pool, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Move media to another media pool
Automatically draw media from a free media pool
Create a new media pool
Change the way Removable Storage displays
Working with MMC console files

# Set the allocation limit for a media pool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the allocation limit for a media pool**

1. Open Removable Storage.
2. If Removable Storage is in Full view, in the console tree, double-click **Media Pools**.

   If Removable Storage is in Simple view, on the **View** menu, click **Full**, and then double-click **Media Pools**.
3. Right-click the applicable media pool, and then click **Properties**.
4. Under **Media information**, click **Contains media of type**.
5. Under **Allocation/Deallocation policy**, select the **Limit reallocations** check box, and then change the default value as necessary.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a new media pool
Change the way Removable Storage displays
Working with MMC console files

# Manage Tapes and Disks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage tapes and disks**

- Move media to another media pool
- Enable or disable a tape or disk
- Insert a tape or disk into an automated library
- Eject a tape or disk from a stand-alone drive
- Eject a tape or disk from an automated library
- Mount a tape or disk
- Dismount a tape or disk from an automated library

# Move media to another media pool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move media to another media pool**

1. Open Removable Storage.
2. If Removable Storage is in Simple view, on the **View** menu, click **Full**.
3. In the console tree, click **Contents**.

   **Where?**

   - Removable Storage *(computer name)*/Libraries/*applicable library*/Contents

4. In the details pane, drag the applicable tape or disk from the details pane over to the applicable media pool in the console tree.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- When moving media to a free media pool, make sure the data on the tape or disk you are moving is no longer needed. Data will be automatically deleted when the tape or disk is moved to a free media pool. To manually delete data before moving media to a free media pool, right-click the tape or disk and then click **Free**.
- You cannot move read-only media, such as compact discs, to a free media pool.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Automatically return media to a free media pool
Automatically draw media from a free media pool
Change the way Removable Storage displays
Working with MMC console files

# Enable or disable a tape or disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable a tape or disk**

1. Open Removable Storage.
2. In the console tree, click **Contents**.

   **Where?**

   - Removable Storage *(computer name)*/Libraries/*applicable library*/Contents

3. In the details pane, right-click the tape or disk, and then click **Properties**.
4. On the **Media** tab, verify that the **Enable media** check box is selected.

   Or, to disable the tape or disk, clear the **Enable media** check box.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- Disabling the tape or disk makes it unavailable for use.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable or disable a drive
Enable or disable a library
Working with MMC console files

# Insert a tape or disk into an automated library

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To insert a tape or disk into an automated library**

1. Open Removable Storage.
2. In the console tree, double-click **Libraries**.
3. Right-click the applicable library, and then click **Inject**.
4. Follow the instructions in the Media Inject wizard.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- Use this procedure only with an automated library. For a stand-alone drive library, you must manually insert the tape or disk.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Eject a tape or disk from an automated library
Working with MMC console files

# Eject a tape or disk from a stand-alone drive

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To eject a tape or disk from a stand-alone drive**

1. Open Removable Storage.
2. In the console tree, double-click **Libraries**.
3. Right-click the applicable library, and then click **Eject**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- Use this procedure only with a stand-alone drive library. For information on how to eject a tape or disk from an automated library, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Eject a tape or disk from an automated library
Working with MMC console files

# Eject a tape or disk from an automated library

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To eject a tape or disk from an automated library**

1. Open Removable Storage.
2. In the console tree, click **Contents**.

   **Where?**

   - Removable Storage *(computer name)*/Libraries/*applicable library*/Contents

3. In the details pane, right-click the applicable tape or disk, and then click **Eject**.
4. Follow the instructions in the Media Eject wizard.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Eject a tape or disk from a stand-alone drive
Insert a tape or disk into an automated library
Working with MMC console files

# Mount a tape or disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To mount a tape or disk**

1. Open Removable Storage.
2. In the console tree, click **Contents**.

   **Where?**

   - Removable Storage *(computer name)*/Libraries/*applicable library*/Contents

3. In the details pane, right-click the applicable tape or disk, and then click **Mount**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dismount a tape or disk from an automated library
Working with MMC console files

# Dismount a tape or disk from an automated library

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To dismount a tape or disk from an automated library**

1. Open Removable Storage.
2. In the console tree, click **Drives**.

   **Where?**

   - Removable Storage *(computer name)*/Libraries/*applicable library*/Drives

3. In the details pane, right-click the drive from which you want to dismount the tape or disk, and then click **Dismount**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Mount a tape or disk
Working with MMC console files

# Manage Operator Requests and Queued Work

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage operator requests and queued work**

- Respond to an operator request
- Change how operator requests are displayed
- Delete operator requests
- Cancel pending operations in the work queue
- Change the mount order in the work queue

# Respond to an operator request

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To respond to an operator request**

1. Open Removable Storage.
2. On the **View** menu, click **Full**, and in the console tree, click **Operator Requests**.
3. In the details pane, right-click the applicable request, and then do one of the following:
   - To complete the request, click **Complete**, and then perform the action when prompted.
   - To cancel the request, click **Refuse**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- Removable Storage stores operator requests for approximately 72 hours after they have been completed or cancelled.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete operator requests
Change how operator requests are displayed
Working with MMC console files

# Change how operator requests are displayed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change how operator requests are displayed**

1. Open Removable Storage.
2. In the console tree, right-click **Removable Storage** and then click **Properties**.
3. On the **General** tab, under **Options**, do the following:
   - To display operator requests in a pop-up window, select the **Display operator request dialogs** check box; otherwise, clear the check box (the default).
   - To display an icon on the Windows taskbar when there are operator requests waiting to be serviced, select the **Use notification area icon for pending operator requests** check box; otherwise, clear the check box (the default).

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Respond to an operator request
Delete operator requests
Working with MMC console files

# Delete operator requests

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete operator requests**

1. Open Removable Storage.
2. On the **View** menu, click **Full**, and in the console tree, right-click **Operator Requests**, and then click **Properties**.
3. To have all requests automatically deleted, in the **Operator Requests Properties** dialog box, verify that the **Automatically delete completed requests** check box is selected.
4. To automatically delete failed requests also, verify that **Delete failed requests** is selected.

   Or, to save failed requests, click **Keep failed requests**.
5. In **Delete after**, change the time interval as necessary for automatically deleting requests.
6. To restore default settings, click **Default**.
7. To delete all requests now, click **Delete All Now**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- Restoring the default settings will automatically delete all requests, including failed requests, every 72 hours.
- To delete an individual operator request, right-click the applicable request in the details pane, and then click **Delete**.
- This procedure also applies for deleting items from the work queue. Right-click **Work Queue** in the console tree instead.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Respond to an operator request
Change how operator requests are displayed
Working with MMC console files

# Cancel pending operations in the work queue

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To cancel pending operations in the work queue**

1. Open Removable Storage.
2. On the **View** menu, click **Full**, and in the console tree, click **Work Queue**.
3. In the details pane, right-click the applicable item in the work queue, and then click **Cancel Request**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- You can also delete an item from the work queue by right-clicking the applicable item, and then clicking **Delete**.
- You cannot delete work items that are in progress, except for door opening requests.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the mount order in the work queue
Working with MMC console files

# Change the mount order in the work queue

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the mount order in the work queue**

1. Open Removable Storage.
2. On the **View** menu, click **Full**, and in the console tree, click **Work Queue**.
3. In the details pane, right-click the applicable mount operation, and then click **Re-order Mounts**.
4. In the **Change Mount Order** dialog box, do one of the following:
   - To make the selected mount operation the next to be completed, click **Move to the front of the queue**.
   - To make the selected mount operation the last to be completed, click **Move to the end of the queue**.
   - To move the selected mount operation to a specific position in the work queue, click **Make it number**, and then click the up or down arrow to move the item ahead or behind in the queue.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- You can re-order only queued mount operations. You cannot re-order mount operations that are completed or in progress.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Cancel pending operations in the work queue
Working with MMC console files

# Manage Security for Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage security for Removable Storage**

- Change user permissions for Removable Storage
- Add users for access to Removable Storage

# Change user permissions for Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change user permissions for Removable Storage**

1. Open Removable Storage.
2. In the console tree, right-click the specific item you want to change user permissions for, and then click **Properties**.
3. In the **Properties** dialog box, click the **Security** tab.
4. Click the name of the user or group whose permissions you want to change, and then do one of the following:
   - To change specific permissions, select or clear the **Allow** and **Deny** check boxes for each of the three levels of permissions: **Use**, **Modify**, and **Control**, as applicable.
   - To deny all permissions, click **Remove**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- You can apply user permissions to libraries, to media pools and their media, or to Removable Storage as a whole.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add users for access to Removable Storage
Working with MMC console files

# Add users for access to Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add users for access to Removable Storage**

1. Open Removable Storage.
2. In the console tree, right-click the specific item you want to add user access to, and then click **Properties**.
3. On the **Security** tab, click **Add**.
4. In the **Select Users, Computers, or Groups** dialog box, in **Enter the object name to select**, type the name of the user or group, and click **Check Names** to verify the name against the user database. You can add users from a different domain by clicking **Locations** and choosing a domain from the list.
5. Click the name of the user or group that you just added, and then do one of the following:
   - To change specific permissions, select or clear the **Allow** and **Deny** check boxes for each of the three levels of permissions: **Use**, **Modify**, and **Control**, as applicable.
   - To deny all permissions, click **Remove**.

## Notes

- To complete this procedure, you need the appropriate administrative credentials for the affected Removable Storage item. To maintain system security, rather than logging in as an administrator, log in as a user without administrative credentials, and run Removable Storage as an administrator by right-clicking the Computer Management icon, clicking **Run As**, and choosing an account in the local Administrators group.
- To open Removable Storage, click **Start** and then click **Control Panel**. Double-click **Administrative Tools**, and then double-click **Computer Management**. In the console tree, double-click **Storage** and then click **Removable Storage**.
- You can add user access to libraries, to media pools and their media, or to Removable Storage as a whole.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change user permissions for Removable Storage
Working with MMC console files

# Removable Storage Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

These topics provide background information about Removable Storage.

- Removable Storage overview
- Understanding Removable Storage
- Using Removable Storage
- Removable Storage Resources

# Removable Storage overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Removable Storage overview**

Removable Storage makes it easy for you to track your removable storage media (tapes and optical disks) and to manage the hardware libraries, such as changers and jukeboxes, that contain them. Removable Storage labels, catalogs, and tracks media; controls library drives, slots, and doors; and provides drive-cleaning operations.

Removable Storage works in concert with your data-management programs, such as Backup. You use data-management programs to manage the actual data stored on the media. Removable Storage makes it possible for multiple programs to share the same storage media resources, which can reduce your costs.

Removable Storage organizes all the media in your libraries into different media pools. Removable Storage also moves media between media pools in order to provide the amount of data storage your applications require.

Removable Storage does not provide volume management, such as for media siding or striping. Also, Removable Storage does not provide file management, such as for data backup or disk-extender operations. These services are done by data-management applications such as Backup or Remote Storage. Remote Storage is not available on computers running Windows XP Professional, Windows Server 2003, Web Edition; or Windows Server 2003, Standard Edition.

You must run all your data-management programs on the same computer that connects to your library. Removable Storage does not support multiple data-management programs running on different computers that are connected to the same library.

# Understanding Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Removable Storage**
This section covers:

- Removable Storage components
- Libraries
- Media pools
- Media identification
- Media states

# Removable Storage components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Removable Storage components**
Removable Storage consists of an administrative interface implemented as a Microsoft Management Console (MMC) snap-in, a Windows service with an API, and a database. The Removable Storage service provides media services to data-management programs. The following figure illustrates the relationship between Removable Storage, data management programs, libraries, and Windows.

You can use the Removable Storage snap-in to perform the following tasks:

- Create media pools and set media pool properties.
- Insert and eject media in an automated library.
- Mount and dismount media.
- View the state of media and libraries.
- Clean tape drives.
- Perform library inventories.
- Set security permissions for users.

**Note**

- You do not use the Removable Storage API or database to perform administrative tasks. For more information on the Removable Storage API and database, including how to program to them, see "Managing Data with Removable and Remote Storage" at the Microsoft Windows Resource Kits Web site.

# Libraries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Libraries**

In its simplest form, a library consists of data storage media and the device used to read from and write to the media. The group of libraries and associated media that are managed by a Removable Storage installation is called a *Removable Storage system*. There are two major types of libraries:

- **Automated libraries** are automated units that hold multiple tapes or disks, and some have multiple drives. These libraries are sometimes called *changers* or *jukeboxes*, and commonly use robotic subsystems to move media stored in the library's storage slots. Automated libraries can also consist of other hardware components that are managed by Removable Storage, such as doors, inject/eject ports, cleaner cartridges, and bar-code readers.

- **Stand-alone drive libraries**, or *stand-alone drives*, are single-drive, nonautomated units, such as tape or CD-ROM drives, that hold a single tape or disk. With these, you manually insert a tape or disk into the unit.

Not only can Removable Storage manage multiple libraries, but it can also track offline media that are not currently contained in a library. These media might be anywhere: on a shelf or in a drawer. Every tape or disk that is offline belongs to a library.

# Media pools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Media pools**

A media pool is a collection of tapes or disks to which the same management properties apply. All media in a Removable Storage system belong to a media pool, and each media pool holds only one type of media. Data management programs use media pools to gain access to specific tapes or disks within a library.

Using media pools, you can define properties that apply to a set of media. This is useful because Removable Storage allows multiple programs to share the same media within a single library. A library can include media from different media pools, each with different properties. A single media pool can span multiple libraries. You can also create hierarchies of media pools, or media pools that contain other media pools. For example, you can create a media pool for each specific media type required by a program, and then create another media pool that contains this collection of media pools. Media pools can contain either media or other media pools, but not both.

A Removable Storage system provides two classes of media pools: system and application.

System media pools include free media pools, unrecognized media pools, and import media pools. Removable Storage creates one free, one unrecognized, and one import media pool for each media type in your Removable Storage system. The system media pools are used to hold media that are not currently being used by an application.

Application media pools are created by data management programs such as Backup and Remote Storage, the latter for computers running Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; the 64-bit versions of Windows Server 2003, Enterprise Edition; or the 64-bit versions of Windows Server 2003, Datacenter Edition.

The different media pools are defined as follows:

- **Unrecognized media pools** contain blank (new) media and media that Removable Storage does not recognize. You should immediately move a new tape or disk from an unrecognized media pool to a free media pool so that the tape or disk can be used by applications, or remove it from the library.

  Unrecognized media are automatically deleted from the Removable Storage database when they are ejected from a library.

- **Import media pools** contain media that Removable Storage recognizes in the database but that have not been used before in a particular Removable Storage system. For example, media in an import media pool could be media from one office location that are introduced into a Removable Storage system at another office location.

  You can move media from import media pools to free media pools or application media pools so applications can use them.

- **Free media pools** contain media that are not currently in use by applications and do not contain useful data. Media in free media pools are available for use by applications.

  Application media pools can be configured to automatically draw media from free media pools when there are not sufficient media available in a particular application media pool. If this configuration is not implemented, you must manually move media from a free media pool when needed.

- **Application media pools**, created by data management applications (and by you), determine what media can be accessed by which applications.

  Media in an application media pool are controlled by that application or by an administrator. An application can use more than one media pool, and more than one application can share a single media pool. For example, Backup might use one media pool for full-backup and another media pool for incremental backup, each containing a different media type.

  There can be any number of application media pools in a Removable Storage system. Media that are currently reserved for use by an application, called *allocated media*, cannot be moved between media pools. Allocation controls how the media are used by applications.

The following illustration shows how media can move from one media pool to another.



# Media identification

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Media identification**
Removable Storage uses two identification methods to track media and perform inventories: on-media identifiers and bar codes.

**On-media identifiers**
On-media identifiers are electronically recorded on media the first time a tape or disk is inserted into a library. The next time the medium is inserted into the library, Removable Storage uses the on-media identifier to recognize and track the tape or disk.

An on-media identifier has two parts: a label type and a label ID. The label type identifies the specific format used to record information on the tape or disk, such as Microsoft Tape Format (MTF). The label ID is a unique identifier for the tape or disk in the Removable Storage system. If Removable Storage does not recognize the label type, it puts the medium in an unrecognized media pool. If Removable Storage recognizes the label type but does not recognize the label ID, it puts the media in an import media pool. If Removable Storage recognizes the label type and the label ID, it updates the Removable Storage database to reflect that the medium is online and available for use.

**Note**

- Removable Storage does not record on-media identifiers for read-only or write-once optical media, such as CD-ROM, DVD-ROM, CD-R, DVD-R, or WORM disks. Instead, Removable Storage uses the volume and serial number information that is already associated with the disk.

**Bar codes**
If your library supports bar codes, Removable Storage can use them to identify media. Media with bar codes also have on-media identifiers, and Removable Storage can use either one. Using bar codes to track media is generally much faster because each tape or disk does not have to be mounted to read the on-media identifier.

# Media states

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Media states**

Each tape or disk in a Removable Storage system has a state associated with it that determines what operations can be performed with the media. When media are inserted into a library, Removable Storage uses the on-media identifier or bar code to determine its state.

There are two sets of media states that govern media movement and usage by Removable Storage: physical states and side states.

**Physical states**

The physical states identify operations performed that involve media location or movement. There are five possible physical states:

| Physical state | Description |
|---|---|
| Idle | Tape or disk is currently in a storage slot in an automated library or is shelved offline. |
| In-use | Tape or disk is in the process of being moved. |
| Loaded | Tape or disk is mounted in a drive and data is available to be read or written. |
| Mounted | Tape or disk is mounted in a drive but data is not yet available to be read or written. |
| Unloaded | Tape or disk has been dismounted and is ready to be removed from a drive. |

**Side states**

A tape or disk side is where actual information is stored. Side states identify operations performed that involve media usage rather than location or movement. There are nine possible side states:

| Side state | Description |
|---|---|
| Allocated | Side is reserved for use by an application. The side is not available to any other application. |
| Available | Side is currently available for use by an application. |
| Completed | Side is in use but cannot be used for write operations. The side is full. |
| Decommissioned | Side is no longer available for use, as it has reached its allocation maximum. |
| Imported | Side's label type is recognized but its label ID is not recognized by Removable Storage. |
| Incompatible | The media type of the side is not compatible with the library. The tape or disk should be removed from the library. |
| Reserved | Side is unavailable for allocation except by a single application. This applies only to two-sided media where one side has already been allocated. |
| Unprepared | Side has been placed in a free media pool but does not yet have a free media label. This is a temporary state. |
| Unrecognized | Side's label type and label ID are not recognized by Removable Storage. |

There is a relationship between side states and media pools: Import media pools can only hold media sides in the imported state. Unrecognized media pools can only hold media sides that are in the unrecognized state. Free media pools can only hold media sides that are in the available (and transitional unprepared) state.

# Using Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Removable Storage**
This section covers:

- Managing libraries
- Managing media pools
- Managing tapes and disks
- Managing operator requests and queued work
- Managing security for Removable Storage
- Managing Removable Storage from the command line

# Managing libraries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing libraries**

Removable Storage attempts to automatically configure all libraries when you first start Removable Storage and when you add or remove a library from a computer that is running Removable Storage. Removable Storage can always automatically configure stand-alone drives.

You can use the automatic configuration feature of Removable Storage for automated libraries only if you connect all drives on the same SCSI bus as the associated media changer, or if all drives have unique serial numbers. Also, you must use automated libraries that support drive-element address reporting. For information on whether your library supports this feature, consult the library manufacturer.

**Note**

- Because not all library configurations support the automatic configuration feature, Removable Storage provides a method for manually configuring libraries. However, do not use this method unless it is absolutely necessary because manually configured libraries do not recognize system configuration changes. They also involve registry changes. For information on how to configure your libraries manually, see "Managing Data with Removable and Remote Storage" at the Microsoft Windows Resource Kits Web site.

Removable Storage provides comprehensive control and management of libraries. You can enable or disable entire libraries or only specific drives within a library. You can change the media type for a library. You can also open the door on an automated library and set various time-out values.

For information on how to perform a specific task, click:

- Enable or disable a library
- Enable or disable a drive
- Change the library media type
- Open a library door
- Set the library door and port time-out

**Inventorying libraries**

You can create an inventory of all media in a library. Removable Storage supports two inventory methods:

- **Fast inventory.** Removable Storage checks the storage slots and reads the on-media identifier on media in slots that were previously empty.
- **Full inventory.** Removable Storage mounts each tape or disk in your library and reads the on-media identifier.

For information on how to perform an inventory task, click:

- Inventory a library
- Set the default inventory method

**Cleaning drives**

Removable Storage can manage the cleaning of an entire automated library or the drive for a stand-alone drive. Each tape-based automated library can contain one cleaning cartridge. You can use Removable Storage to clean a drive, assign or release a cleaning cartridge slot, and insert or eject a cleaning cartridge. Removable Storage maintains a usage count for each cleaning

cartridge. When a cleaning cartridge reaches its maximum usage count, Removable Storage generates an operator request. If you eject a cleaning cartridge before it has reached its maximum usage count, Removable Storage deletes the usage count information.

For information on how to perform a specific task, click:

- Clean a stand-alone drive
- Clean an automated library

# Managing media pools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing media pools**
After configuring your libraries, you can configure your media pools. Before creating a new media pool, you should answer the following questions:

- What media types are available in your Removable Storage system?
- Should media be automatically drawn from a free media pool when there are not sufficient media in an application media pool? (It is recommended that you implement this configuration.)
- Should media be automatically returned to a free media pool when they are no longer needed by an application?

For information on how to perform a specific task, click:

- Create a new media pool
- Delete an application media pool
- Automatically draw media from a free media pool
- Automatically return media to a free media pool
- Set the allocation limit for a media pool

**Note**

- You cannot delete system media pools.

# Managing tapes and disks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing tapes and disks**

After you create and configure your media pools, you can move media to a specific media pool. It is important that you have a sufficient amount of media in a free media pool so that applications can use media in this pool when needed. For information on how to move media to a free media pool, see Move media to another media pool.

Next, you can use a data management application such as Backup to allocate media to the application media pool created for (or by) that application. Allocating media to an application media pool reserves all media in that media pool for exclusive use by that application. See the documentation for your particular data management application for information on how to reserve media for that application.

Removable Storage provides comprehensive control of tapes or disks in your library. You can enable or disable specific media in your library, insert and eject media, and mount and dismount media.

A tape or disk can be inserted in or ejected from an automated library using either a library door or an inject/eject port. A library door allows unrestricted access to the media inside. You can add or remove media directly from a storage slot using a door access. An inject/eject port allows controlled access to the media inside. You insert a tape or disk in the port and then the library uses the drive bay to move the media to a storage slot.

A tape or disk can be mounted in or dismounted from a particular drive in a library. You can use the same media many times before dismounting them. You can also mount and dismount the same media many times before deallocating them.

For information on how to perform a specific task, click:

- Enable or disable a tape or disk
- Insert a tape or disk into an automated library
- Eject a tape or disk from a stand-alone drive
- Eject a tape or disk from an automated library
- Mount a tape or disk
- Dismount a tape or disk from an automated library

**Notes**

- Removable Storage requires compact discs to be formatted with the CDFS file system, and rewritable optical media (such as MO disks) to be formatted with either the FAT, FAT32, or NTFS file system.
- Windows does not directly support the use of media that can be written to only once, such as WORM, CD-R, and DVD-R disks. If you plan to use these, you must acquire other software that supports these technologies.

# Managing operator requests and queued work

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing operator requests and queued work**

The work queue lists all library requests, or work items, whether initiated by an application or by Removable Storage. A request to mount a tape in a library, for example, results in a mount work item listed in the work queue.

The following table describes the states that a work item can have:

| Work state | Description |
|---|---|
| Completed | The work item has been completed successfully. |
| Failed | Removable Storage was unable to complete the work item. |
| In Process | Removable Storage is actively operating on the work item. |
| Queued | The work item is ready to be carried out, but Removable Storage has not yet begun to operate on it. |
| Waiting | One or more resources needed to satisfy the request are currently in use. |
| Cancelled | A user has cancelled the work item. |

An operator request is a message that asks you to perform a specific task. Operator requests can be issued by Removable Storage or by a program that is aware of Removable Storage, such as Backup. Removable Storage can generate an operator request when any of the following occur:

- An application initiates a mount request for a tape or disk that is offline.
- There are no available media online. That is, a program requires media, and none exist in the appropriate application media pool or in a free media pool.
- A library has failed and requires servicing.
- A drive needs cleaning and no usable cleaner cartridges are available in the library.

You can respond to operator requests by either completing or refusing the operator request. When you refuse an operator request, Removable Storage notifies the application that generated the request. By default, Removable Storage stores an operator request in the work queue for 72 hours after you respond to the request. For information on how to respond to an operator request, see Respond to an operator request.

The following table describes the states that an operator request can have:

| Operator request state | Description |
|---|---|
| Completed | You have indicated, or Removable Storage has detected, that the operator request was completed. |
| Refused | You have indicated that the operator request will not be performed. |
| Submitted | Removable Storage or an application has submitted an operator request for action. |

You can change how operator requests are displayed on your computer, delete operator requests, cancel any pending operation listed in the work queue, and change the order of mount requests in the work queue. For information on how to perform a specific task, click:

- Change how operator requests are displayed
- Delete operator requests

- Cancel pending operations in the work queue
- Change the mount order in the work queue

# Managing security for Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing security for Removable Storage**
Removable Storage provides user-level security for media pools and libraries. Media pool security controls user access to media. Library security controls user access to library drives, slots, and inject/eject ports.

User permissions determine which users or groups of users can access media pools and libraries. Before assigning permissions for a particular media pool or library, it is a good idea to use Active Directory Users and Computers to create a user group containing the accounts of users who need access to a media pool or library. It is also a good idea to review the security tips listed in the Removable Storage Best practices for Removable Storage.

For information on how to perform a specific task, click:

- Change user permissions for Removable Storage
- Add users for access to Removable Storage

# Managing Removable Storage from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Removable Storage from the command line**

You can use the command-line utility for Removable Storage to manage media resources and to write batch scripts for applications that do not currently support the Removable Storage API (are not aware of Removable Storage). You can run the command-line utility from a command prompt by typing **rsm**.

For help using the Removable Storage command line utility, see Rsm.

For information about other command-line utilities, see Command-line reference A-Z.

# Removable Storage Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

The *Microsoft* Windows *Resource Kit* contains advanced information on how to configure, administer, optimize, and troubleshoot Removable Storage. For more information, see "Managing Data with Removable and Remote Storage" at the Microsoft Windows Resource Kits Web site.

# Troubleshooting Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Windows cannot detect the library after you install it.
- Removable Storage has stopped the automatic configuration process for an automated library.
- Removable Storage was unable to automatically configure an automated library.
- The library is configured correctly but is malfunctioning.
- When configuring a media pool, Removable Storage cannot detect the media type.
- You cannot set or change security for Removable Storage.
- I cannot access the Computer Management extension snap-ins on a remote computer.

**Windows cannot detect the library after you install it.**
**Cause:** The library may not be supported by Windows.

**Solution:** Verify that the library is supported by Windows. For a comprehensive list of the libraries supported by Windows operating systems, see the hardware compatibility information in the Windows Catalog. You can access the Windows Catalog from Support resources. Make sure you have configured the library according to any guidelines listed for the applicable library.

**Cause:** The library may not be properly connected to your computer.

**Solution:** Verify that the library is properly connected to your computer. If an automated library is connected to a SCSI host adapter (by means of a SCSI cable), make sure there are no SCSI ID conflicts with other devices in your computer, such as your hard drive or CD-ROM drive. In addition, verify that all cables are installed and terminated properly in accordance with the guidelines for SCSI operation, especially that the maximum length for such cables is not exceeded.

**Removable Storage has stopped the automatic configuration process for an automated library.**
**Cause:** Occasionally, Windows will configure a drive inside an automated library before the library itself has been configured. In this situation, Removable Storage will complete the configuration process once Windows is done configuring the library.

**Solution:** Verify that Windows has recognized the library and associated drives and has configured the device drivers correctly.

**Removable Storage was unable to automatically configure an automated library.**
**Cause:** The library may not support drive-element address reporting. Consequently, Removable Storage cannot map each drive in the library to a particular drive bay.

**Solution:** Use only libraries that support drive-element address reporting, and connect all library drives on the same SCSI bus as the associated drive bay. For information on whether your library supports drive-element address reporting, consult the library manufacturer.

**Solution:** If Removable Storage still cannot automatically configure your library correctly, you will need to manually configure it. However, do not use this method unless it is absolutely necessary because manually configured libraries do not recognize system configuration changes and also involve registry changes. For information on how to configure your libraries manually, see "Managing Data with Removable and Remote Storage" at the Microsoft Windows Resource Kits Web site.

**The library is configured correctly but is malfunctioning.**
**Cause:** Many such problems can be caused by hardware device errors.

**Solution:** See the Windows system event log for details. You can use Event Viewer to view the error messages stored in the event log. Look for specific Removable Storage error messages as well as for device, drive, and controller error messages to

pinpoint the source of the problem. Contact the library manufacturer once you have ruled out Removable Storage as the source of the problem.

**When configuring a media pool, Removable Storage cannot detect the media type.**
**Cause:** The library may not contain any tapes or disks.

**Solution:** Before you create a new media pool, insert media into each automated library you want Removable Storage to control.

**You cannot set or change security for Removable Storage.**
**Cause:** Removable Storage requires that you have administrative credentials for these operations.

**Solution:** Verify that you are logged on with administrative credentials.

**I cannot access the Computer Management extension snap-ins on a remote computer.**
**Cause:** The Remote Registry service is not running on the remote computer.

**Solution:** Ensure that the Remote Registry service is started on the remote computer.

**See also:** Start, stop, pause, resume, or restart a service. You need appropriate credentials on the remote computer to start the service.

### Note

- You might also receive this error message if the remote computer is running Windows 95. Computer Management does not support remote access to computers that are running Windows 95.

# Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Storage**
Remote Storage allows you to extend disk space on your server using a tape or magneto-optical disk library. Remote Storage manages the disk volumes on your server and automatically copies infrequently-used files to a tape or disk library. When the amount of free disk space falls below the defined level, Remote Storage then frees up the space on your disks.

Remote Storage is not included on Windows Server 2003, Standard Edition. If you are using Windows 2000 Server with Remote Storage, you cannot upgrade to Windows Server 2003, Standard Edition. You can upgrade to Windows Server 2003, Enterprise Edition, because Remote Storage is included in this product.

- Before you begin using Remote Storage, see Checklist: Setting up Remote Storage.
- For tips about using Removable Storage, see Best practices for Remote Storage.
- For help with specific tasks, see Remote Storage How To....
- For general background information see Remote Storage Concepts.
- For problem-solving instructions, see Troubleshooting Remote Storage.

# Checklist: Setting up Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Setting up Remote Storage**

Remote Storage is not installed by default when you run Windows Server 2003, Enterprise Edition Setup, Windows Server 2003, Datacenter Edition Setup, or Setup for the 64-bit versions of Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition. You can select Remote Storage for installation during Setup, or you can manually install Remote Storage through **Add or Remove Programs** in **Control Panel**.

| Step | Reference |
|---|---|
| Log on to the local computer with administrative credentials. | Default local groups. |
| Verify that you have a supported tape or disk library and that Removable Storage can detect this library. Remote Storage can use any magneto-optical disk or 4mm, 8mm, or DLT tape library. | For more information about compatible hardware, see Support resources. |
| Using Removable Storage, verify that a sufficient number of tapes or magneto-optical disks for Remote Storage to use are assigned to either the Free or Remote Storage media pool. | View or set Removable Storage properties. |
| Verify that the volumes that will be managed have been formatted with the NTFS file system for Windows 2000 or the Windows Server 2003 family. | Convert command reference. |
| If you intend to use data compression, compress the volume data beforehand. | See the product documentation for your applicable software. |
| If you intend to use content indexing, index the volume data beforehand. | See the product documentation for your applicable software. |
| Install Remote Storage and start the Remote Storage Setup wizard. | Set up Remote Storage. |
| To maintain security, log out, log back in using an account in a different security group, and run Remote Storage as an administrator. To do this, right-click **Remote Storage**, click Run as, and choose an account in the Administrators group. | Access Control. |
| Select volumes for management. | Follow the Remote Storage Setup Wizard. |
| Set the free-space level for volumes that you choose to manage. | Follow the Remote Storage Setup Wizard. |
| Set the criteria for copying files to remote storage. | Follow the Remote Storage Setup Wizard. |
| Select the type of tape or disk to use for remote storage. | Follow the Remote Storage Setup Wizard. |
| Set the file-copy schedule. | Follow the Remote Storage Setup Wizard. |
| Set the number of media copies that you want to make. | Set the number of media copies. |

# Best practices for Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

The following list represents some best practices for Remote Storage. You should implement this list whenever possible.

- **Backup files on managed volumes as well as the Remote Storage database files.** Remote Storage relies on backups to protect applicable files cached on your local volumes, in addition to the Remote Storage database and other program data files in the System32\RemoteStorage folder. You should keep track of when data backups occur and which tapes are used. For information on Backup, see Backup overview.

  For more information on protecting data managed by Remote Storage, see Protecting and recovering data from loss.

- **Backup files as an administrator when backing up Remote Storage database files.** Remote Storage requires that you run system state backups as a member of the Administrators group in order to backup the Remote Storage database files. If you run a system state backup as a member of the Backup Operators group, it will not include the Remote Storage database files.

- **Make copies of your remote storage data.** You should make multiple copies of the tapes or disks that Remote Storage uses and regularly update them. In this way, if you later encounter a problem with your original tape or optical disk (the *master media*), you can recreate it using a copy you have made. Doing this ensures no loss of data.

- **Keep full media copies off site**. If possible, keep full media copies off site or rotate media copies between your business sites for safekeeping. (Non-full media copies are required on site for synchronization.)

- **Do not use the Remote Storage command-line utility before you configure Remote Storage.** You must configure Remote Storage using the Remote Storage Setup wizard before you use the **rss.exe** command-line utility, or you will not be able to use the utility to successfully manage Remote Storage.

- **Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage.** In addition, do not add a volume that contains directories that are part of an FRS replica set to Remote Storage. Otherwise, you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

- **Validate your managed volumes regularly.** You should schedule validation of your managed volumes on a regular basis. Validation ensures that all files on managed volumes point to the correct data in remote storage.

- **Discontinue managing all volumes before uninstalling Remote Storage**. Before you uninstall Remote Storage, you should discontinue managing all volumes by selecting **Recall copied data from storage** in the Remove Volume Management wizard. Otherwise, you may lose your file data.

- **Think before you delete a file from a managed volume**. If you delete a file from a managed volume, the associated data can no longer be recalled from remote storage. Be careful not to delete such files if you might want to recall the data at a later time.

- **Do not use Remote Storage to manage excessively full volumes.** Remote Storage requires free space on each volume to store the metadata for each file it manages. If there is not enough space on the volume, Remote Storage will behave unpredictably or fail to work. There is no way to predict how much space you'll need for the metadata, so use your best judgment and clear space on excessively full volumes before attempting to manage them.

- **Do not format a managed volume**. When a managed volume is formatted, Remote Storage cannot locate the volume anymore, and does not know if the volume is temporarily unavailable or was permanently removed (which is the case when a volume is being formatted).

- **Load balance your scheduled tasks**. When managing several large-capacity volumes, you should schedule tasks to run on different days and at different times to spread out the workload on your server. For example, all tasks for volume A could be scheduled to run on Tuesdays and Thursdays at 2:00 am; all tasks for volume B could be scheduled to run on Mondays, Wednesdays, and Fridays at 2:00 am.

- **Avoid scheduling tasks at system restart**. To prevent excess workload on your server, you should not schedule any tasks to start at system restart.

- **Do not change the drive letter assigned to a managed volume.** If you must change a volume's drive letter, stop managing it with Remote Storage first.

- **Log in as a domain user when managing volumes**. Remote Storage only notifies you of file recalls when you are logged in as a domain user. If you log onto your local computer only, you will not receive these recall notifications.

- **Avoid deleting the USN journal, if it has been created.** If someone has created an update sequence number (USN) journal on the computer using the **fsutil** utility, deleting it will degrade the performance of Remote Storage.

# Remote Storage How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Set up Remote Storage
- View or set properties for Remote Storage
- Monitor tasks for Remote Storage
- View events for Remote Storage
- Manage Local Disk Volumes
- Manage Remote Storage Media

# Set up Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up Remote Storage**

1. If Remote Storage is not yet installed on your computer, do the following:

    - Open Add or Remove Programs in Control Panel.
    - Click **Add or Remove Windows Components**.
    - Select the **Remote Storage** check box, and then click **Next**.
    - When asked if you want to restart your computer now, click **Yes**.

2. Open Remote Storage.

3. Follow the instructions that appear in the **Remote Storage Setup wizard**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Add or Remove programs, click **Start**, click **Control Panel**, and then double-click **Add or Remove programs**.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**. Remote Storage will not be available until after you have installed it.
- The first time you open Remote Storage after it is installed, the Remote Storage Setup wizard guides you through the steps required to fully configure it for operation.
- Think carefully about the type of media you want to use for remote storage because this cannot be changed later.
- You should not install Remote Storage on a Windows server cluster. Remote Storage will not failover to the other computer node.
- Certain Windows components require configuration before they can be used. If you installed one or more of these components but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components Wizard, click **Components**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage Local Disk Volumes
Working with MMC console files

# View or set properties for Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view or set properties for Remote Storage**

1. Open Remote Storage.
2. In the console tree, click the applicable item. The objects are displayed in the details pane.
3. Right-click the object that you want to change, and then click **Properties**.
4. Review settings and make changes, as needed.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage Local Disk Volumes
Manage Remote Storage Media
Working with MMC console files

# Monitor tasks for Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To monitor tasks for Remote Storage**

1. Open Task Scheduler.
2. Click the applicable task. The status of the task is displayed in the **Status** column.

## Notes

- To perform this procedure, you must be a member of the Administrators, Backup Operators, or Server Operators group on the local computer. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Task Scheduler, click **Start**, click **Control Panel**, and then double-click **Scheduled Tasks**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Remote Storage How To...
Discontinue managing a volume
Copy files to remote storage
Create free space immediately
Change the file-copy schedule
Windows interface administrative tool reference A-Z: Task Scheduler
Task Scheduler
Working with MMC console files

# View events for Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view events for Remote Storage**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then double-click **Event Viewer**.
3. Click a log, and in the details pane, double-click an entry.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- For more information about Event Viewer, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Event Viewer overview
Event Viewer
Working with MMC console files

# Manage Local Disk Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage local disk volumes**


- Add a volume for management
- Discontinue managing a volume
- Set the desired free space for a volume
- Specify basic file-selection criteria
- Add a file rule
- Change priority of file rules
- Set the runaway recall limit
- Set the maximum number of drives to use simultaneously
- Change the file-copy schedule
- Set advanced schedule options
- Copy files to remote storage
- Validate files
- Create free space immediately

# Add a volume for management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a volume for management**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**.
3. Right-click **Managed Volumes**, point to **New**, and then click **Managed Volume(s)**.
4. Follow the instructions that appear in the Add Volume Management wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- You can use the Add Volume Management wizard any time after you run the Remote Storage Setup wizard.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Discontinue managing a volume
Working with MMC console files

# Discontinue managing a volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To discontinue managing a volume**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Managed Volumes**.
3. In the details pane, right-click the applicable volume, and then click **Remove**.
4. Follow the instructions that appear in the Remove Volume Management wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- If you have adequate disk space to recall all data from remote storage, but the **recall copied data from storage** option is not available, you may need to validate the files on the volume first. Validating the volume updates the statistics used by Remote Storage to determine if enough disk space is available to recall data from remote storage. For information on how to validate the files on a managed volume, see Related Topics.
- For information on how to monitor the status of this task, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Validate files
Add a volume for management
Monitor tasks for Remote Storage
Working with MMC console files
Managing local volumes

# Set the desired free space for a volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the desired free space for a volume**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Managed Volumes**.
3. In the details pane, right-click the applicable volume, and then click **Settings**.
4. On the **Settings** tab, in **Desired free space**, click the UP or DOWN ARROW to increase or decrease the value.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- Using this procedure, you can define the free space for the volume. When the available space on the volume falls below this setting, Remote Storage frees up space on the volume by removing cached data from files that have been copied to remote storage.
- You use this procedure to configure Remote Storage to automatically free up space on your volume when the actual free space falls below the **Desired Free Space** setting. For information on how to override this setting and immediately free up space on your volume, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create free space immediately
Working with MMC console files

# Specify basic file-selection criteria

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify basic file-selection criteria**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Managed Volumes**.
3. In the details pane, right-click the applicable volume, and then click **Settings**.
4. On the **Settings** tab, under **File criteria**, in **Larger than**, click the UP or DOWN ARROW to set the file size value.
5. Under **File criteria**, in **Not accessed in**, click the UP or DOWN ARROW to change the time value.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- If insufficient free space is created on your managed volume because not enough files are being copied to remote storage, you may need to decrease these values.
- For information on how to specify more complex criteria for copying files to remote storage based on file rules, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a file rule
Change priority of file rules
Working with MMC console files

# Add a file rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a file rule**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Managed Volumes**.
3. In the details pane, right-click the applicable volume, and then click **Include/Exclude Rules**.
4. On the **Include/Exclude Rules** tab, click **Add**.
5. In **Path**, type the folder path to the file (do not include the drive letter).
6. In **File type**, type the wildcard character and file extension (for example: *.jpg).
7. To exclude matching files when the rule is applied, click **Exclude matching files**.

   Or, to include matching files when the rule is applied, click **Include matching files**.
8. To apply the rule to any matching files stored in subordinate folders, select the **Apply rule to subfolders** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- For information on how to copy files to remote storage based also on file size and last access date, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change priority of file rules
Specify basic file-selection criteria
Working with MMC console files

# Change priority of file rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change priority of file rules**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Managed Volumes**.
3. In the details pane, right-click the applicable volume, and then click **Include/Exclude Rules**.
4. On the **Include/Exclude Rules** tab, select the applicable rule.
5. Click the UP ARROW to give the rule a higher priority.

   Or click the DOWN ARROW to give the rule a lower priority.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- The position of a file rule on the list determines which rule is applied first. Remote Storage applies the first rule on the list that matches a file in a given folder.
- There are a number of predefined rules provided when Remote Storage is first set up. You cannot move some of these predefined rules up or down the list. You can however, move all rules that you define up or down the list.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a file rule
Working with MMC console files

# Set the runaway recall limit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the runaway recall limit**

1. Open Remote Storage.
2. In the console tree, right-click **Remote Storage**, and then click **Properties**.
3. On the **Limits** tab, change the value in **Maximum number of successive recalls** by clicking the UP or DOWN ARROWS.
4. To exclude all users with administrative credentials from this limit, select the **Limit does not apply to members of the Administrators group** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- The runaway recall limit prohibits a user, or an application, from recalling files from remote storage more than the specified number of times in succession with less than 10 seconds between each recall. Virus checking, data backup, indexing, and search applications typically read all files on your computer and can cause excessive file recalls to occur otherwise.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify basic file-selection criteria
Set the desired free space for a volume
Working with MMC console files

# Set the maximum number of drives to use simultaneously

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the maximum number of drives to use simultaneously**

1. Open Remote Storage.
2. In the console tree, right-click **Remote Storage**, and then click **Properties**.
3. On the **Limits** tab, change the value in **Maximum number of removable media drives to use** by clicking the UP or DOWN ARROWS.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- This limit sets the maximum number of tape or magneto-optical drives in the selected device that Remote Storage can read from or write to simultaneously.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the runaway recall limit
Working with MMC console files

# Change the file-copy schedule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the file-copy schedule**

1. Open Remote Storage.
2. In the console tree, right-click **Remote Storage**, click **Properties**, and then click the **Schedule** tab. This displays the current schedule.
3. Click **Change Schedule**. This displays the **File Copy Schedule** dialog box.
4. In **Schedule Task**, click the arrow, and then select an interval.
5. In **Start time**, set the time by clicking the UP or DOWN ARROW.
6. Under **Schedule Task Daily**, in **Every**, change the value by clicking the UP or DOWN ARROW, if necessary.
7. To set multiple schedules, select the **Show multiple schedules** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- You use this procedure to automatically copy files to remote storage based on a schedule. For information on how to immediately copy files from specific volumes to remote storage, see Related Topics.
- For information on how to monitor the status of this task, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set advanced schedule options
Copy files to remote storage
Monitor tasks for Remote Storage
Working with MMC console files

# Set advanced schedule options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set advanced schedule options**

1. Open Remote Storage.
2. In the console tree, right-click **Remote Storage**, and then click **Properties**.
3. On the **Schedule** tab, click **Change Schedule**.
4. In the **Remote Storage File Copy Schedule** dialog box, click **Advanced**.
5. In the **Advanced Schedule Options** dialog box, in **Start Date**, select a calendar date.
6. To set an end date, select the **End Date** check box, click the DOWN ARROW, and select a calendar date.
7. To repeat tasks, select the **Repeat task** check box and then do the following, as needed:
   - To change the time interval, in **Every**, click the UP or DOWN ARROW to set the value, and then click the DOWN ARROW and select the units.
   - To change the time, in **Until**, click **Time**, and click the UP or DOWN ARROW.
   - In **Until**, click **Duration**, and then change the settings in **hour(s)** and **minute(s)** by clicking the UP and DOWN ARROWS.
8. To stop a task if it is still running at the time previously specified, select the **If the task is still running, stop it at this time** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- For information on how to monitor the status of this task, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the file-copy schedule
Monitor tasks for Remote Storage
Working with MMC console files

# Copy files to remote storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To copy files to remote storage**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Managed Volumes**.
3. In the details pane, right-click the applicable volume, click **All Tasks**, and then click **Copy Files to Remote Storage**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- You use this procedure to immediately copy files from specific volumes to remote storage. For information on how to automatically copy files to remote storage based on a schedule, see Related Topics.
- For information on how to monitor the status of this task, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the file-copy schedule
Monitor tasks for Remote Storage
Working with MMC console files

# Validate files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To validate files**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Managed Volumes**.
3. In the details pane, right-click the applicable volume, click **All Tasks**, and then click **Validate Files**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- Validation checks to make sure that files on managed volumes still point to valid data in remote storage and updates the volume statistics.
- For information on how to monitor the status of this task, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor tasks for Remote Storage
Working with MMC console files

# Create free space immediately

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create free space immediately**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Managed Volumes**.
3. In the details pane, right-click the applicable volume, click **All Tasks**, and then click **Create Free Space**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- This procedure frees up space on your volume by removing cached data from files that have been copied to remote storage. If all cached data has already been removed, no additional free space can be created. In addition, recall activity will increase as data is no longer cached locally.
- You use this procedure to immediately free up space on your volume. For information on how to configure Remote Storage to automatically free up space on your volume when the available space falls below a specified amount, see Related Topics.
- For information on how to monitor the status of this task, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the desired free space for a volume
Monitor tasks for Remote Storage
Working with MMC console files

# Manage Remote Storage Media

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage remote storage media**

- Set the number of media copies
- Synchronize media copies
- Delete media copies
- Recreate the media master

# Set the number of media copies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the number of media copies**

1. Open Remote Storage.
2. In the console tree, right-click **Remote Storage**, and then click **Properties**.
3. On the **Media Copies** tab, change the value in **Number of media copy sets** by clicking the UP or DOWN ARROW.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- Copies of remote storage media can be created only if two or more drives are available.
- You can make up to three copies of remote storage media. You should synchronize media copies with the original media (called the media master) on a regular basis to ensure no loss of data. For information on how to synchronize media copies, see Related Topics.
- Reducing the number of media copies does not cause any media to be deallocated, it only limits the copy sets that Remote Storage will display and access. To deallocate the media, select all the media and then delete the media copies. For information on how to delete media copies, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Synchronize media copies
Delete media copies
Working with MMC console files

# Synchronize media copies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To synchronize media copies**

1. Open Remote Storage.
2. In the console tree, right-click **Media** and then click **Synchronize Media Copies**.
3. Follow the instructions that appear in the Media Copy Synchronization wizard.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- Media copies can be created only if two or more drives are available.
- For information on how to monitor the status of this task, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitor tasks for Remote Storage
Set the number of media copies
Delete media copies
Working with MMC console files

# Delete media copies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete media copies**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Media**.
3. In the details pane, right-click one or more of the media shown, and then click **Media Copies**.
4. On the **Media Copies** tab, do one or more of the following:
   - To delete the first copy of the media master, click **Delete Copy 1**.
   - To delete the second copy of the media master, click **Delete Copy 2**.
   - To delete the third copy of the media master, click **Delete Copy 3**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- Deleting a media copy removes it from the Remote Storage database, but it does not erase the data from the tape or disk on which it is stored. This data can still be read by a malicious user who has physical access to the tape or disk. To ensure the security of the data, keep tapes or disks in a secure location.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set the number of media copies
Synchronize media copies
Working with MMC console files

# Recreate the media master

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To recreate the media master**

1. Open Remote Storage.
2. In the console tree, double-click **Remote Storage**, and then click **Media**.
3. In the details pane, right-click one of the media shown, and then click **Properties**.
4. Click the **Recovery** tab, and then click **Re-create Master**.
5. Follow the instructions that appear on-screen.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Remote Storage, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Remote Storage**.
- You should recreate the media master only if you experience errors when you recall files from remote storage.
- Make sure the media copy you selected to use to recreate the media master is the most current version.
- Using Removable Storage (not to be confused with Remote Storage), you should verify that there are tapes or disks in a free media pool of equal or greater capacity to the media copy you are using to recreate the media master. For information on how to view Removable Storage properties, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
View or set Removable Storage properties
Working with MMC console files

# Remote Storage Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides background information about Remote Storage.

- Remote Storage overview
- Understanding Remote Storage
- Using Remote Storage
- Remote Storage Resources

# Remote Storage overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Storage overview**
Remote Storage makes it easy for you to extend disk space on your server computer without adding more hard disks. Remote Storage automatically copies eligible files on your local volumes to a library of magnetic tapes or magneto-optical disks. Remote Storage then monitors the amount of space available on your local volumes.

File data is cached locally so that it can be accessed quickly as needed. When the amount of available space on a managed volume dips below the level you need, Remote Storage automatically removes the content from the cached files, providing the disk space you need. When data is removed from a file, the disk space used by the file is reduced to zero. Data from cached files is not removed until more disk space is needed. When you need to open a file whose data has been removed, the data is automatically recalled from remote storage.

Since removable tapes or disks in a library are less expensive per megabyte (MB) than hard disks, this can be an economical way to provide both maximum data storage and optimal local disk performance.

**Levels of data storage**
Remote Storage data storage is hierarchical, with two defined levels. The upper level, called local storage, includes the NTFS disk volumes of the computer running Remote Storage. The lower level, called remote storage, is on the automated media library or stand-alone tape or disk drive that is connected to the server computer.

Remote Storage supports all SCSI class 4mm, 8mm, DLT, and magneto-optical devices that are supported by Removable Storage. Using Remote Storage with Exabyte 8200 tape libraries is not recommended. Remote Storage does not support QIC tape libraries or rewritable compact disc and DVD formats.

**Retrieval of stored files**
When you need to access a file on a volume managed by Remote Storage, you simply open the file as usual. If the data for the file is no longer cached on your local volume, Remote Storage recalls the data from a tape library. Because this can take more time than usual, Remote Storage removes the data only from those files on your local volumes that you are least likely to need, based on criteria that you set.

**Coordination with other tools**
Remote Storage uses Removable Storage to access the applicable tapes contained in libraries. Remote Storage also works with Backup for data recovery, and with Task Scheduler to schedule file copy operations. For information on Removable Storage, see Removable Storage. For information on Backup, see Backup overview. For information on Task Scheduler, see Task Scheduler.

Remote Storage also provides certain data-recovery features, including the ability to generate multiple copies of data in remote storage.

### Note

- Using Remote Storage to copy files to remote storage is not the same as backing up your files. You should follow a regular schedule of data backups, including backing up the contents of the local volumes that Remote Storage manages. You should also backup the Remote Storage database and other program files located in the System32\RemoteStorage folder.

### Caution

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. In addition, do not add a volume that contains directories that are part of an FRS replica set to Remote Storage. Otherwise, you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

# Understanding Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Remote Storage**
This section covers:

- Basic operation of Remote Storage
- How Remote Storage affects the operating system
- Managed volumes
- Remote Storage and Removable Storage

# Basic operation of Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Basic operation of Remote Storage**

The following steps describe how to use Remote Storage to manage space on your local disk volumes.

Before Remote Storage can manage your volumes, you must verify that a sufficient number of tapes or disks have been moved to a free media pool using Removable Storage. For more information on managing tapes or disks using Removable Storage, see Managing tapes and disks.

- You determine which local disk volumes you want Remote Storage to manage. You then define the amount of free space you want maintained on the volumes, the file-selection criteria for copying files to remote storage, and the type of media to use for remote storage.

- Remote Storage then copies files to remote storage that meet the criteria, leaving the original data stored, or cached, on the local volumes.

- Remote Storage compares the amount of available space on a managed volume with the amount you defined above. If the actual free disk space is less than the amount that you defined, Remote Storage removes the data on your local volumes, reducing the disk space used by the files to zero. This step continues until the free space setting that you defined is met or when all data has been removed from cached files on your managed volumes.

- When you want to access a file, you simply open the file on your local volume in the usual way. Remote Storage automatically recalls the data from remote storage and copies the data back to your local volume. The type of library and media used determines how much time delay you experience.

Recalling a file does not erase it from the tape or disk on which it is stored. This data can still be read by a malicious user who has physical access to the tape or disk. To ensure the security of the data, keep tapes or disks in a secure location.

# How Remote Storage affects the operating system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Remote Storage affects the operating system**

Remote Storage runs as a service and uses Removable Storage to access the applicable tapes or disks in a library that are used for remote storage. You administer Remote Storage from Microsoft Management Console (MMC), using the Remote Storage snap-in to perform all tasks, or from the command line using the Rss command. The following figure illustrates the relationship between Remote Storage and your server.




Remote Storage uses Windows security to grant or deny access for storage management. Only user accounts that have administrative credentials can administer Remote Storage; however, users with appropriate permissions can open the files on volumes that Remote Storage manages. To maintain system security, log in as a member of a different security group and run Remote Storage as an administrator. To do this, right-click **Remote Storage**, click Run as, and choose an account in the Administrators group.

Remote Storage adds a **Remote Storage** tab to the volume **Properties** dialog box of Windows Explorer and to the Disk Management snap-in on the local computer. This tab shows local and remote storage statistics and settings. While you can remotely manage volumes from other computers, the **Remote Storage** tab only appears on the computer that the Remote Storage services are running on.

Programs that use file size, such as Disk Quota, are not affected by Remote Storage operations. When Remote Storage frees up disk space by removing data from cached files that have been copied, the disk space used by file becomes zero. However, the *logical* size and the date and time (created, last modified, last accessed) attributes of the file remain unchanged.

When using Windows File Manager (Winfile.exe), files that Remote Storage manages are not displayed. This affects client computers running versions of Windows prior to Windows 2000.

Only one Remote Storage installation can exist on a server.

"You cannot remotely administer Remote Storage from a computer running Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, to a computer running Windows 2000."

**Important**

- Do not remotely administer Remote Storage on a computer running a different-language version of the operating system.

# Managed volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managed volumes**

The local disk volumes on your server computer that are under Remote Storage control are called managed volumes. For example, you can configure drive C as a managed volume. If you have supplementary hard disks on other drives, such as D, E, and so forth, you can also configure these as additional managed volumes.

Disk volumes managed by Remote Storage must be nonremovable and must be formatted with NTFS.

Remote Storage automatically copies data from managed volumes to remote storage according to a schedule and to rules that you specify. You configure the amount of free disk space you want to maintain on the managed volumes and the criteria used to copy specific files to remote storage.

The following figure illustrates how Remote Storage copies files from local storage to remote storage.




**Step 1** - Files copied to remote storage and cached locally




**Step 2** - Data removed from cached local files when actual free space is less than the desired free space setting

# Remote Storage and Removable Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote Storage and Removable Storage**

Remote Storage uses Removable Storage to access the applicable media in a library that are used for remote storage.

Remote Storage can support a multiple-drive automated library, but only if all the drives within it are identical. Within such an automated library, Remote Storage can access two or more drives at the same time. If the automated library contains two or more drives, Remote Storage can simultaneously copy data to, and recall data from, two or more different tapes or disks.

All remote storage tapes or disks used by Remote Storage exist in a single application media pool that is automatically created during Remote Storage Setup. You use Removable Storage to verify that sufficient media have been moved to a free media pool, so that Remote Storage can use tapes or disks from that pool if needed. You cannot move tapes or disks from the Remote Storage application media pool to another application media pool.

Remote Storage can support only a single tape or disk type for use as remote storage. You specify which type is supported when you run the Remote Storage Setup wizard. You cannot change this type later.

Recalling data from remote storage may take more time than usual. The amount of this time delay depends on several factors:

- Whether or not a tape or disk is already mounted in a drive.
- The speed of a library in mounting a tape or disk.
- Availability of the tapes or disks and drives in a library.
- Whether or not other file recalls are pending.

The following figure illustrates how Remote Storage recalls data from remote storage to local storage. Note that when data is recalled, it is copied back to the original volume and the file again takes up disk space.




Remote Storage does not support the use of QIC tapes as remote storage media.

Media used by Remote Storage contain only data that has been copied from managed volumes. However, the library can contain additional media used by other applications, such as Backup.

Remote Storage does not support recalling data for a volume that has been managed at different times by two different Remote Storage installations.

# Using Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Remote Storage**
This section covers:

- Setting up Remote Storage
- Managing local volumes
- Managing remote storage media
- Protecting and recovering data from loss
- Managing Remote Storage from the command line

# Setting up Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting up Remote Storage**

Remote Storage is not installed by default when you install the operating system. You can set Remote Storage for installation when you set up your operating system, or afterwards using Add or Remove Programs in Control Panel. For information on how to set up Remote Storage after you have installed the operating system, see Set up Remote Storage.

You should not install Remote Storage on a Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition computer in a server cluster configuration. Remote Storage will not fail over to the other computer node.

Before you can use Remote Storage, you must satisfy three requirements:

- You must be logged on with administrative credentials.
- You must verify that a sufficient number of tapes or disks are in a free media pool. (Use Removable Storage to perform this task.)
- You must verify that volumes managed by Remote Storage are formatted with the version of NTFS used in Windows 2000 and the Windows Server 2003 family. For more information, see Convert.

If you intend to use data compression on managed volumes, compress the volumes beforehand. Also, if you intend to use content indexing, run your indexing software beforehand. See the product documentation for your particular software

You must open Remote Storage using an account in the Administrators group. To maintain system security, log in using an account in a different security group and run Remote Storage as an administrator. To do this, right-click **Remote Storage**, click Run as, and choose an account in the Administrators group.

Remote Storage is administered in Microsoft Management Console (MMC) using the Remote Storage snap-in. The first time you open the Remote Storage snap-in, the Remote Storage Setup wizard guides you through the steps required to fully configure Remote Storage for operation. After you have configured Remote Storage with the wizard, you can use either the Remote Storage snap-in or the Remote Storage command-line tool, Rss, to administer Remote Storage.

# Managing local volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing local volumes**

Remote Storage manages physical disk volumes on the local computer by copying selected files on those volumes to a tape or disk library according to the schedule and rules that you specify. The original files are left cached on your local volumes. You can use Remote Storage to automatically copy files, following a schedule that you specify, or you can immediately copy files to tape or disk.

Volumes managed by Remote Storage must be nonremovable (fixed) media. You initially select which volumes Remote Storage manages when you complete the Remote Storage Setup wizard. For information on how to add an additional volume for Remote Storage management, see Add a volume for management.

**Setting the desired free space**

You can specify the amount of free space that you want to maintain on each local volume. This setting establishes a threshold value.

Remote Storage continuously monitors the actual amount of free space on your local volumes. If the amount of actual free space falls below the Desired Free Space setting, Remote Storage removes the cached data from files on your local volumes that have been copied to tape or disk. The disk space used by files on your local volumes becomes zero.

You set the free space initially when you complete the Remote Storage Setup wizard. This value is used for each volume that Remote Storage manages. For information on how to change the Desired Free Space setting, or to set a different value for each volume, see Set the desired free space for a volume.

You can also free up disk space manually. Doing this immediately removes the cached data from files on your local volumes. For information on how to immediately free up disk space, see Create free space immediately.

**Specifying file-selection criteria and rules**

You can specify the criteria governing which files you want Remote Storage to copy to remote storage. All files that meet the criteria are eligible to be copied to tape or disk. You specify criteria based on minimum file size and time elapsed since the files were last accessed.

You initially specify file criteria when you complete the Remote Storage Setup wizard. These criteria are used for each volume that Remote Storage manages. For information on how to change the criteria, or to specify separate criteria for each volume, see Specify basic file-selection criteria.

You can also specify file inclusion and exclusion rules that affect file selection for copying to remote storage. These rules can include or exclude matching files based on:

- A specified folder, including all subfolders within it.
- A specified folder, excluding subfolders.
- File name extension.
- File name, including wildcard characters.

Remote Storage provides a predefined list of file inclusion and exclusion rules that you can use. You cannot edit or change the order of some of these predefined rules. You can, however, move and edit all rules that you create.

Remote Storage uses the first rule in the rule list that matches a file in a given folder. Changing the order of a rule in the list changes the priority of that rule.

For information on how to perform specific tasks, click:

- Add a file rule.
- Change priority of file rules.

### Note

- By default, all hidden, system, extended attribute, encrypted, and sparse files are excluded from the file rule list. These files cannot be copied to remote storage.

**Changing the file-copy schedule**

Remote Storage provides a default schedule used for file copying when you complete the Remote Storage Setup wizard. For information about how to change the file-copy schedule and about advanced schedule options, see Change the file-copy schedule.

You can also copy files to remote storage at any time, regardless of the schedule. For information on how to immediately copy eligible files, see Copy files to remote storage.

### Note

- Scheduled recalling of data in remote storage is not supported.

**Setting the maximum number of drives to access simultaneously**

You can specify how many drives in your multiple-drive device Remote Storage can access at the same time for file copy and file recall operations. The number you should specify depends on your individual needs. Increasing the limit can improve performance when you are managing multiple hard disks, and when several users at once try to access files that have been copied to different remote media. On the other hand, if you have another application (like Backup) that accesses removable storage, you may want to set the limit to less than the total number of drives in your device, so that Remote Storage cannot prevent the other application from running by using all of the available drives concurrently.

Setting the limit to a number greater than the total number of drives in your device has the same effect as setting it to the total number of drives.

**Setting the runaway recall limit**

Remote Storage provides a default value to the maximum number (called the runaway recall limit) of successive file-recalls that a user, or an application, can make on any managed volume in the same session. Once a file recall is completed, if the user or application requests another recall within 10 seconds, the count is increased. File-recalls copy the data back to local volumes, thereby decreasing the available free disk space on your computer. Virus checking, data backup, indexing, and search applications typically read all files on your computer and can also cause a large number of file-recalls. For information on how to change the runaway recall limit, see Set the runaway recall limit.

**Validating files**

Validation ensures that all the files on your managed volumes point to valid and correct data in remote storage. It also recalculates volume statistics. Validation is automatically performed two hours after a backup program is used to restore a remote storage file. Validation also detects files that have been moved from one local volume to another or that have been modified.

Validation should be performed on a regular basis to validate all files on a given volume and to correct any inconsistencies. Validation should also be performed after restoring files on a local volume or after disk errors have occurred on a volume. For information on how to validate the files on a managed volume, see Validate files.

The validation operation is integrated with and can be monitored using Scheduled Tasks. For more information on how to monitor Remote Storage tasks, see Monitor tasks for Remote Storage.

**Discontinuing volume management**

You can select one of the two following ways in which Remote Storage discontinues managing a volume:

- Remote Storage discontinues managing the local volume but leaves data in remote storage. Data that is already in remote storage can be recalled as before. No automatic freeing of disk space takes place.

- Remote Storage recalls data from remote storage and then discontinues managing the local volume. You should only select this option if your local volume has adequate disk space for the recalled data.

  If you have adequate disk space to recall data from remote storage, but the option of recalling all data is disabled, click the **Refresh** button in order to update volume statistics. After you have updated the volume statistics, the option should be enabled.

For information on how to discontinue managing a volume using the Remove Volume Management wizard, see Discontinue managing a volume.

These operations are integrated with and can be monitored using Scheduled Tasks. For more information on how to monitor Remote Storage tasks, see Monitor tasks for Remote Storage.

**Important**

- Before you uninstall Remote Storage, you should discontinue managing all volumes by clicking **Recall copied files from remote storage** in the Remove Volume Management wizard. Otherwise, you may lose your file data.

**Note**

- When a managed volume is formatted, the volume will still be displayed in the Remote Storage snap-in, but the only operation you can perform on it is to discontinue managing it. This occurs because Remote Storage cannot locate the volume anymore, and does not know if the volume is temporarily unavailable or was permanently removed (which is the case when a volume is formatted).

**Modifying files on managed volumes**

There are a number of considerations that must be taken into account before you modify or move a file on a managed volume. This section addresses those issues.

If a file on a managed volume is deleted, lost, or destroyed, you must retrieve a copy of the file from a backup tape or disk to gain access to the file. If you delete a file from a managed volume, the associated data can no longer be recalled from remote storage. Be careful not to delete such files if you might want to recall the data at a later time. You can protect such files using Backup. For information on Backup, see Backup overview.

Copying or moving a file between managed volumes causes the associated data in remote storage to be recalled. You must use a backup program to back up and then restore a file before you move the file between managed volumes to avoid causing a recall. All applicable volumes must be formatted using the version of NTFS used in Windows 2000 and the Windows Server 2003 family. The local volume on which the moved file now resides must be managed by Remote Storage for disk space to be automatically freed, and you should also validate the volume afterward.

**Note**

- Renaming a file on a managed volume does not recall the associated data from remote storage.

# Managing remote storage media

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing remote storage media**

Remote Storage supports all SCSI class optical disk or 4mm, 8mm, and DLT tape libraries. Using Remote Storage with Exabyte 8200 tape libraries is not recommended. Remote Storage does not support QIC tape libraries.

All tapes or magneto-optical disks used by Remote Storage exist in a single application media pool. This pool is automatically created at the completion of the Remote Storage Setup wizard. You must select the appropriate media type to be used for remote storage in the Remote Storage Setup wizard. The media type you specify cannot be changed later.

Verify that sufficient tapes or disks have been moved to a free media pool so that Remote Storage can use the tapes if needed. You cannot move media from the Remote Storage media pool to another application media pool.

Offline media are tapes or disks that are not currently inside a library and are therefore inaccessible by Remote Storage. Since these tapes or disks remain in the same media pool, Removable Storage is responsible for tracking their location. For information on Removable Storage, see Removable Storage.

If a file that has been copied to remote storage is currently on offline media, you will receive a message from Removable Storage to manually insert the applicable tape or disk. If you do not do so within a couple of minutes or so, the recall will fail.

### Note

- Remote Storage is optimized for copying large files. If a managed volume has many small files that are eligible for copying to remote storage, the data transfer rate may not match the rating given for the tape library.

**Creating media copies**

Remote Storage can back up the data on tapes or disks that it uses. The original tape or disk collection used to hold all files copied from local storage is called the media master set. Remote Storage can automatically create copies from the media master. The collection of copies for a media master is called a media copy set. You can specify up to three media copy sets. For information on how to set the number of media copy sets, see Set the number of media copies.

### Notes

- Files that are copied to remote storage are not protected by mirroring or striping. The only way to obtain redundancy for data in remote storage is to create copies of your media masters.

- Media copies can be created only if two drives are available for use in a tape library.

Remote Storage will first use available tapes or disks in the Remote Storage media pool, then use tapes or disks from a free media pool, for media copies. The free media pool is a collection of unused media that can be used when there is not enough storage space available in the Remote Storage media pool.

**Synchronizing media copies**

You should update, or *synchronize*, media copies on a regular basis. This ensures that data on a media copy matches that on the media master. Synchronization updates media copies for all the media masters; however, you can only synchronize one media copy set at a time. Files cannot be managed or recalled while a media copy set is being synchronized, even if additional drives are available. For information on how to synchronize media copies, see Synchronize media copies.

Media copies can be synchronized later as needed if they are removed from a library.

**Recreating the media master**

When data on a media master becomes corrupted or is lost, Remote Storage can recreate the media master using data from a media copy. For information on how to recreate the media master, see Recreate the media master.

# Protecting and recovering data from loss

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Protecting and recovering data from loss**
Keep the following points in mind if you are restoring files on a managed volume using Backup or other data backup application:

- You cannot restore Remote Storage files to a different computer.
- You cannot restore Remote Storage files to a volume that is not formatted with the version of NTFS used by Windows 2000 and the Windows Server 2003 family.
- You should not restore Remote Storage files to other NTFS volumes on the same computer unless you are sure that those volumes are also managed by Remote Storage.
- Remote Storage uses the same data format as Backup. In an emergency you can use Backup to restore files. For more information, see Backup.

For advanced information on how to protect and recover data that is managed or needed by Remote Storage, see "Managing Data with Removable and Remote Storage" at the Microsoft Windows Resource Kits Web site.

# Managing Remote Storage from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Remote Storage from the command line**

You can use the command-line utility for Remote Storage to copy and recall files, to configure and manage Remote Storage, and to write batch scripts for applications that will allow them to access Remote Storage directly. You can run the command-line utility from a command prompt by typing **Rss**.

For help using the Remote Storage command line utility, see Rss.

For information about other command-line utilities, see Command-line reference A-Z.

# Remote Storage Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

**Windows Server 2003 Family Resource Kit**
The *Microsoft* Windows Server 2003 *Family Resource Kit* contains advanced information on how to configure, administer, optimize, and troubleshoot Remote Storage. For more information, see "Managing Data with Removable and Remote Storage" at the Microsoft Windows Resource Kits Web site.

# Troubleshooting Remote Storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
Keep in mind that Remote Storage is used to move data between your local computer and tape libraries, and depends on Removable Storage to access specific tapes within these libraries. Consequently, Removable Storage must be properly configured beforehand and must be running for Remote Storage to function correctly.

What problem are you having?

- You cannot install or run Remote Storage.
- You cannot find the Remote Storage snap-in.
- Remote Storage could not find any valid media type to use.
- Remote Storage cannot manage a particular volume or find a volume that you want to manage.
- You have set a file-copying schedule but no files have been copied to remote storage.
- You can no longer recall files from remote storage.
- File recalls are slow.
- There are excessive file recalls occurring.

**You cannot install or run Remote Storage.**
**Cause:** Remote Storage requires that you have administrative credentials.

**Solution:** Verify that you are running Remote Storage with administrative credentials. To maintain system security, log in using an account in a different security group and run Remote Storage as an administrator. To do this, right-click **Remote Storage**, click Run as, and choose an account in the Administrators group.

**You cannot find the Remote Storage snap-in.**
**Cause:** Remote Storage is not installed by default when the operating system is installed.

**Solution:** Install Remote Storage on your computer using **Add or Remove Programs**.

**Remote Storage could not find any valid media type to use.**
**Cause:** Removable Storage has not detected the media type yet.

**Solution:** Wait until the Removable Storage snap-in displays the available media that can be used.

**Cause:** You do not have a supported tape or disk library.

**Solution:** Verify that you have a supported tape or disk library. Remote Storage supports all SCSI class magneto-optical disk or 4mm, 8mm, and DLT tape libraries. Remote Storage does not support QIC tape libraries. For tape and disk libraries supported by Removable Storage, see the information about compatible hardware at the Support resources page.

**Remote Storage cannot manage a particular volume or find a volume that you want to manage.**
**Cause:** Remote Storage can only manage nonremovable (fixed), local NTFS volumes; other file systems or volumes will not be displayed in the Remote Storage Setup wizard.

**Solution:** Format all volumes that will be managed with the version of the NTFS file system used in Windows 2000 or the Windows Server 2003 family. For help converting volumes to NTFS, see Convert.

**You have set a file-copying schedule but no files have been copied to remote storage.**
**Cause:** There may be no files that meet the specified criteria.

**Solution:** Change the file rules and other file-copy criteria as necessary.

**You can no longer recall files from remote storage.**
**Cause:** Remote Storage uses the runaway recall limit to prohibit a user, or an application, from recalling files from remote storage more than a specified number of times in succession. Virus checking, data backup, indexing, and search applications typically read all files on your computer, and can cause the runaway recall limit to be quickly reached.

**Solution:** Increase this setting as necessary or close the application causing the excessive recalls.

**File recalls are slow.**
**Cause:** Remote Storage may take five minutes or longer to recall a file from a tape-based library.

**Solution:** Set the free space value lower so data in cached files are not removed as often. In this way, data will not have to be recalled from remote storage as often.

**There are excessive file recalls occurring.**
**Cause:** Virus checking, data backup, indexing, and search applications typically read all files on your computer.

In particular, if you use Find (on the Edit menu) in Microsoft Office 95 or later to search for content on a managed volume, all files in remote storage are automatically recalled.

**Solution:** Take into account use of applications that cause excessive files recalls on your managed volumes when planning your activities.

**Cause:** If Microsoft Word (.doc), Microsoft PowerPoint (.ppt), or HTML files have been copied to remote storage, all files in remote storage are recalled if you install Microsoft Office 97 on the computer running Remote Storage. This occurs at the end of the Office 97 installation.

**Solution:** To stop the recalls, you must use Windows Task Manager (press CTRL-ALT-DEL) to end the Findfast.exe task.

# Managing Files and Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing files and folders**
This section covers:

- Shared Folders
- Distributed File System (DFS)
- Encrypting File System
- File compression overview
- Using Indexing Service
- Offline Files
- Synchronization Manager

# Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shared Folders**
You can use the Shared Folders snap-in to manage shared resources over a network. With Shared Folders you can control user access permissions, session activity, and shared resource properties.

- For tips about using Shared Folders, see Best practices for Shared Folders.
- For help with specific tasks, see Shared Folders How To....
- For general background information, see Shared Folders Concepts.
- For problem-solving instructions, see Troubleshooting Shared Folders.

# Best practices for Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Assign permissions to groups, not user accounts.**

- Assigning permissions to groups simplifies management of shared resources, because you can then add users to or remove them from the groups without having to reassign permissions. To deny all access to a shared resource, deny the Full Control permission.

**Assign the most restrictive permissions that still allow users to perform required tasks.**

- For example, if users need only to read information in a folder, and they will never delete, create, or change files, assign the Read permission.

**If users log on locally to access shared resources, such as on a terminal server, set permissions by using NTFS file system permissions or access control.**

- Share permissions apply only to users who access shared resources over the network; they do not apply to users who log on locally. For this situation, use NTFS and access control. For more information, see Set, view, change, or remove permissions on files and folders.

**Organize resources so that objects with the same security requirements are located in the same folder.**

- For example, if users require the Read permission for several application folders, store the application folders in the same parent folder. Then, share the parent folder, rather than sharing each individual application folder. Note that if you need to change the location of an application, you may need to reinstall it.

**When you share applications, organize all shared applications in one folder.**

- Organizing all applications in one shared folder simplifies administration, because there is only one location for installing and upgrading software.

**To prevent problems with accessing network resources, do not deny permissions to the Everyone group.**

- The Everyone group includes anyone who has access to network resources, including the Guest account, with the exception of the Anonymous Logon group. For more information, see Default security settings for groups and Differences in default security settings.

**Avoid explicitly denying permissions to a shared resource.**

- It is usually necessary to explicitly deny permissions only when you want to override specific permissions that are already assigned.

**Limit membership in, and assign the Full Control permission to, the Administrators group.**

- This enables administrators to manage application software and to control user rights.

**In most cases, do not change the default permission (Read) for the Everyone group.**

- The Everyone group includes anyone who has access to network resources, including the Guest account. In most cases, do not change this default unless you want users to be able to make changes to the files and objects in the shared

resource. For more information about share permissions, see Share permissions.

**Grant access to users by using domain user accounts.**

- On computers running Windows XP Professional that are connected to a domain, grant access to shared resources through domain user accounts, rather than through local user accounts. This centralizes the administration of share permissions.

**Use centralized data folders.**

- With centralized data folders, you can manage resources and back up data easily.

**Use intuitive, short labels for shared resources.**

- This ensures that the shared resources can be easily recognized and accessed by users and all client operating systems.

**Use a firewall.**

- A firewall protects shared resources from access through the Internet. In Windows XP and in the Windows Server 2003 family, you can take advantage of new firewall capabilities. For more information, see Internet Connection Firewall. Instead of Internet Connection Firewall, computers running Windows XP with Service Pack 2 (SP2) and computers running Windows Server 2003 with Service Pack 1 (SP1), use Windows Firewall. For more information, see Help: Windows Firewall.

# Shared Folders How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Share a folder or drive
- Set permissions on a shared resource
- View information about shared resources
- Manage properties of a shared resource
- Disconnect a user or users
- Close an open file or resource
- Stop sharing a resource
- Limit the number of users of a shared resource
- Configure offline settings for a shared resource

# Share a folder or drive

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To share a folder or drive**

- Using Shared Folders
- Using Windows Explorer
- Using a command line

**Using Shared Folders**

1. Open Computer Management.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. On the **Action** menu, click **New Share**.
4. Follow the steps in the Share a Folder Wizard, and then click **Finish**.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using Windows Explorer**

1. Open Windows Explorer.
2. Right-click the folder or drive to be shared, and then click **Sharing and Security**.
3. Click **Share this folder**.
4. Set any other options that you want, and then click **OK**.

**Note**

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **net share***sharename=drive:path*

| Value | Description |
| --- | --- |
| **net share** | Creates, deletes, or displays shared resources. |
| *sharename=drive:path* | The network name of the shared resource and the absolute path of its location. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at the command prompt, type:

  **net help share**

**Important**

- When you create a shared resource, determine how to secure that resource. There are several methods that you can use to secure shared resources. For more information, see Related Topics.

**Notes**

- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- You can use Shared Folders to manage shared resources on both local and remote computers. For information about how to connect to another computer, see Related Topics. With Windows Explorer and the command line, you can manage shared resources on your local computer only.
- You can hide the shared resource from users by typing **$** as the last character of the shared resource name (the **$** then becomes part of the resource name). Users can map a drive to this shared resource, but they cannot see the shared resource when they browse to it in Windows Explorer, or in My Computer on the remote computer, or when they use the **net view** command on the remote computer. For more information about this command, see Related Topics.
- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Share permissions
Publish a shared folder
Securing shared resources
Microsoft Management Console
Connect to another computer
**Other Resources**
Net share
Net view

# Set permissions on a shared resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions on a shared resource**

- Using Shared Folders
- Using Windows Explorer

**Using Shared Folders**

1. To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. In the details pane, right-click the shared resource that you want to set permissions for, and then click **Properties**.
4. On the **Share Permissions** tab, make any of the following changes, and then click **OK**:

   - To assign permissions to a user or group for a shared resource, click **Add**. In the **Select Users, Computers, or Groups** dialog box, look for or type the user or group name, and then click **OK**.
   - To revoke access to the shared resource, click **Remove**.
   - To set individual permissions for the user or group, in the **Permissions for***group or user* box, select the **Allow** or **Deny** check boxes.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using Windows Explorer**

1. Open Windows Explorer.
2. Right-click the shared folder or drive that you want to set permissions for, and then click **Sharing and Security**.
3. On the **Sharing** tab, click **Permissions**, make any of the following changes, and then click **OK**:

   - To assign permissions to a user or group for a shared resource, click **Add**. In the **Select Users, Computers, or Groups** dialog box, look for or type the user or group name, and then click **OK**.
   - To revoke access to a shared resource, click **Remove**.
   - To set individual permissions for the user or group, in the **Permissions for** *group or user* box, select the **Allow** or **Deny** check boxes.

### Note

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

### Important

- Share permissions apply only to users who gain access to the resource over the network. They do not apply to users who log on locally, such as on a terminal server. In these cases, use access control on the NTFS file system to set permissions. For more information, see Related Topics.

### Notes

- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- You can use Shared Folders to manage shared resources on both local and remote computers. For information about how to connect to another computer, see Related Topics. With Windows Explorer and the command line, you can manage shared resources on your local computer only.

- When permissions have been assigned both to the shared resource and at the file system level, the more restrictive permission always applies.

- It is usually easier to assign permissions to groups and then add users to groups, rather than assigning identical permissions to individual users.

- If you change permissions on special shared resources, such as ADMIN$, the default settings may be restored when the Server service is stopped and restarted or when the computer is restarted. Note that this does not apply to user-created shared resources whose share name ends in $. For more information about special shared resources, see Related Topics.

- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "How to configure file sharing in Windows XP," in the Microsoft Knowledge Base.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Special shared resources
Set, view, change, or remove permissions on files and folders
Microsoft Management Console
Connect to another computer
**Other Resources**
Net session
Send a console message to all connected computers
Send a console message to a computer that you are managing

# View information about shared resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view information about shared resources**

- Using Shared Folders
- Using a command line

**Using Shared Folders**

1. Open Computer Management.
2. In the console tree, click one of the following folders:
   - To view information about the shared resources on the selected computer, click **Shares**.
   - To view information about open sessions, click **Sessions**.
   - To view information about open files, click **Open Files**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares, Sessions, Open Files

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using a command line**

1. Open Command Prompt.
2. Type one of the following:

   **net share**

   **net session**

   **net file**

| Value | Description |
|---|---|
| **net share** | Displays information about all shared resources on the local computer. |
| **net session** | Displays information about all open sessions between the local computer and other computers on the network. |
| **net file** | Displays information about all open files on shared resources. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for these commands, at the command prompt, type one of the following:

  **net help share**

  **net help session**

  **net help file**

## Notes

- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- When you administer another computer remotely, your connection appears as an open named pipe. It cannot be closed.

- Only files that are open on the selected computer are visible.

- For information about how to connect to another computer, see Related Topics.

- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Microsoft Management Console
Connect to another computer
**Other Resources**
Net share
Net session
Net file
Send a console message to all connected computers
Send a console message to a computer that you are managing

# Manage properties of a shared resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage properties of a shared resource**

- Using Shared Folders
- Using Windows Explorer

**Using Shared Folders**

1. Open Computer Management.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management\System Tools\Shared Folders\Shares

3. In the details pane, right-click the shared resource that you want to manage, and then click **Properties**.
4. Make the changes that you want, and then click **OK**.

### Note

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using Windows Explorer**

1. Open Windows Explorer.
2. Right-click the shared folder or drive that you want to manage, and then click **Sharing and Security**.
3. Make the changes that you want, and then click **OK**.

### Notes

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- You can use Shared Folders to manage shared resources on both local and remote computers. For information about how to connect to another computer, see Related Topics. With Windows Explorer and the command line, you can manage shared resources on your local computer only.
- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Shared Folders How To...
Microsoft Management Console
Connect to another computer

# Disconnect a user or users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disconnect a user or users**

- Using Shared Folders
- Using a command line

**Using Shared Folders**

1. Open Computer Management.
2. In the console tree, click **Sessions**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Sessions

3. Do one of the following:

   - To disconnect all users, on the **Action** menu, click **Disconnect All Sessions**.
   - To disconnect one user, in the details pane, right-click the user name, and then click **Close Session**.
   - To disconnect multiple users, press the CTRL key while clicking the user names, right-click any user name, and then click **Close Session**.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- When you administer another computer remotely, your connection appears as an open named pipe. It cannot be closed.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **net session***computername***/delete**

| Value | Description |
|---|---|
| **net session** | Lists or disconnects inbound sessions on the local computer. |
| *computername* | Identifies the computer for which you want to list or disconnect sessions. |
| **/delete** | Ends sessions that are originated by *computername*, and closes all open files that are opened from *computername*. If *computername* is omitted, all inbound sessions and associated files are closed. |

### Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at the command prompt, type:

  **net help session**

### Caution

- If you disconnect users who are accessing shared resources, they might lose data. You may want to warn connected users before you disconnect them.

### Notes

- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- For information about how to connect to another computer, see Related Topics.

- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Microsoft Management Console
Connect to another computer
**Other Resources**
Net session
Send a console message to all connected computers
Send a console message to a computer that you are managing

# Close an open file or resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To close an open file or resource**

- Using Shared Folders
- Using a command line

**Using Shared Folders**

1. Open Computer Management.
2. In the console tree, click **Open Files**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Open Files

3. Do one of the following:

   - To disconnect all open resources, on the **Action** menu, click **Disconnect All Open Files**.
   - To disconnect one open file or resource, in the details pane, right-click the file name, and then click **Close Open File**.
   - To disconnect multiple open files or resources, press the CTRL key while clicking the file names, right-click any selected file name, and then click **Close Open File**.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using a command line**

- Open Command Prompt.
- Type:

  **net file***ID***/close**

| Value | Description |
|---|---|
| **net file** | Displays the names of all open shared files on a server and the number of file locks, if any, on each file. |
| *ID* | The dynamically assigned ID of the shared resource. |
| **/close** | Closes the shared resource and removes file locks. Type this command from the server where the file is shared. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can use the **openfiles** command to determine the file ID of the shared resource. For more information, see Related Topics.

- To view the complete syntax for this command, at the command prompt, type:

  **net help file**

### Caution

- If you close open files or resources, users might lose data. You may want to warn connected users before you close an open file or resource. For more information, see Related Topics.

### Notes

- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

- You can use Shared Folders to manage shared resources on both local and remote computers. For information about how to connect to another computer, see Related Topics. With Windows Explorer and the command line, you can manage shared resources on your local computer only.

- When you administer another computer remotely, your connection appears as an open named pipe. It cannot be closed.

- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Microsoft Management Console
Openfiles
Connect to another computer
**Other Resources**
Net file
Send a console message to all connected computers
Send a console message to a computer that you are managing
Send a console message to all connected computers

# Stop sharing a resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop sharing a resource**

- Using Shared Folders
- Using Windows Explorer
- Using a command line

**Using Shared Folders**

1. Open Computer Management.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. In the details pane, right-click the folder that you want to stop sharing, and then click **Stop Sharing**.

   To stop sharing multiple instances of a folder or drive, press the CTRL key while you click the instances, right-click any instance of the folder or drive, and then click **Stop Sharing**.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using Windows Explorer**

1. Open Windows Explorer.
2. Right-click the shared folder or drive that you want to stop sharing, and then click **Sharing and Security**.
3. On the **Sharing** tab, do one of the following, and then click **OK**:

   - To stop sharing all instances of the shared resource, click **Do not share this folder**.
   - To stop sharing only one instance of a folder or drive that is shared multiple times, in **Share name**, click the folder or drive that you want to remove, and then click **Remove Share**.

**Note**

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **net share***sharename***/delete**

| Value | Description |
|---|---|
| **net share** | Creates, deletes, or displays shared resources. |
| *sharename* | The network name of the shared resource. |
| **/delete** | Stops sharing the resource. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at the command prompt, type:

  **net help share**

**Caution**

- If you stop sharing open resources, users might lose data. You may want to warn connected users before you close an open file or resource. For more information, see Related Topics.
- In most cases, it is recommended that you not stop sharing the special shared resources that are created by the system, such as ADMIN$. These special shared resources are automatically recreated with their original settings when the computer is restarted.
- The special shared resources NETLOGON and SYSVOL are required on domain controllers. Removing them causes a loss of functionality on all clients that the domain controller serves.
- The IPC$ special shared resource is required for the Server service to operate, and it cannot be deleted.

**Notes**

- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- You can use Shared Folders to manage shared resources on both local and remote computers. For information about how to connect to another computer, see Related Topics. With Windows Explorer and the command line, you can manage shared resources on your local computer only.
- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview

Microsoft Management Console
Connect to another computer
**Other Resources**
Net share
Send a console message to all connected computers
Send a console message to a computer that you are managing

# Limit the number of users of a shared resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To limit the number of users of a shared resource**

- Using Shared Folders
- Using Windows Explorer
- Using a command line

**Using Shared Folders**

1. Open Computer Management.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. In the details pane, right-click the shared folder or drive that you want to limit the number of users for, and then click **Properties**.
4. Under **User limit**, click **Maximum allowed**, and then click **OK**, or click **Allow this number of users**.

   If you click **Allow this number of users**, type or click a maximum number of users, and then click **OK**.

**Note**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using Windows Explorer**

1. Open Windows Explorer.
2. Right-click the shared folder or drive that you want to limit the number of users for, and then click **Sharing and Security**.
3. Under **User limit**, click **Maximum allowed**, and then click **OK**, or click **Allow this number of users**.

   If you click **Allow this number of users**, type or click a maximum number of users, and then click **OK**.

**Note**

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **net share***sharename***/users:***number*

| Value | Description |
|---|---|
| **net share** | Creates, deletes, or displays shared resources. |
| *sharename* | The network name of the shared resource and its absolute path. |
| **/users:** *number* | Sets the maximum number of users who can simultaneously access the shared resource. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at the command prompt, type:

  **net help share**

**Notes**

- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- You can use Shared Folders to manage shared resources on both local and remote computers. For information about how to connect to another computer, see Related Topics. With Windows Explorer and the command line, you can manage shared resources on your local computer only.
- The maximum number of users allowed is limited by software licensing limits. Software licensing limits preempt any limits that can be set in Shared Folders.
- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Microsoft Management Console
Connect to another computer
**Other Resources**
Net share
Send a console message to all connected computers
Send a console message to a computer that you are managing

# Configure offline settings for a shared resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure offline settings for a shared resource**

- Using Shared Folders
- Using Windows Explorer
- Using a command line

**Using Shared Folders**

1. Open Computer Management.
2. In the console tree, click **Shares**.

   **Where?**

   - Computer Management/System Tools/Shared Folders/Shares

3. In the details pane, right-click the shared resource that you want to configure offline settings for, and then click **Properties**.
4. On the **General** tab, click **Offline Settings**.
5. In the **Offline Settings** dialog box, click the option that you want, and then click **OK**.

### Note

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using Windows Explorer**

1. Open Windows Explorer.
2. Right-click the shared folder or drive that you want to configure offline access for, and then click **Sharing and Security**.
3. On the **Sharing** tab, click **Offline Settings**.
4. In the **Offline Settings** dialog box, click the option that you want, and then click **OK**.

### Note

- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.

**Using a command line**

1. Open Command Prompt.
2. Type:

**net share***sharename***/cache:manual**

or

**net share***sharename***/cache:documents**

or

**net share***sharename***/cache:programs**

or

**net share***sharename***/cache:none**

| Value | Description |
|---|---|
| **net share** | Manages shared resources. |
| *sharename* | The network name of the shared resource in *drive:path* format. To display information about a shared folder, type **net share** with only the share name. |
| **/cache:manual** | Allows users to manually select which documents and programs are stored offline. |
| **/cache:documents** | Automatically stores documents offline. |
| **/cache:programs** | Automatically stores documents and programs offline. |
| **/cache:none** | Prevents users from storing any files offline. |

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view the complete syntax for this command, at the command prompt, type:

  **net help share**

**Important**

- When you create a new shared resource, offline access is allowed by default, which means that secure shared resources can be stored offline on potentially insecure computers. For best security, do not allow users to store files offline. If you do, be sure to set permissions appropriately by using either share permissions or access control on the NTFS file system. For more information, see "Set permissions on a shared resource" and "Set, view, change, or remove permissions on files and folders" in Related Topics.

**Notes**

- You must be logged on as a member of the Administrators group, Server Operators group, or Power Users group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.
- For more information about offline files, see Related Topics.
- For information about how to connect to another computer, see Related Topics.
- File sharing options and offline access may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft

Knowledge Base.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Offline settings for shared resources
Set permissions on a shared resource
Set, view, change, or remove permissions on files and folders
Microsoft Management Console
Connect to another computer
Offline Files
**Other Resources**
Net share
Send a console message to all connected computers
Send a console message to a computer that you are managing

# Shared Folders Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Shared Folders:

- Shared Folders overview
- Understanding Shared Folders
- Administering Shared Folders

# Shared Folders overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Shared Folders overview**
You can use Shared Folders to view a summary of connections and resource use for local and remote computers. Shared Folders replaces resource-related components in Control Panel in Microsoft Windows NT® 4.0 Server.

With Shared Folders, you can:

- Create, view, and set permissions for shared resources.
- View a list of all users who are connected over a network to the computer and disconnect one or all of them.
- View a list of files that are opened by remote users and close one or all of the open files.

You must be logged on as a member of the Administrators group, Server Operators group, or the Power Users group to use Shared Folders. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure. File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

For information about publishing a shared folder in Active Directory, see Publish a shared folder.

The subfolders in Shared Folders contain information, arranged in columns, about all the shared resources, sessions, and open files on the computer. If Services for Macintosh is installed, information about resources that are managed by that service may appear as well. The column headings in these folders are defined as follows:

**Shares**
The following table shows the information that the Shares folder contains about the shared resources that are available on the computer.

| Column | Description |
|---|---|
| **Share Name** | Lists the shared resources that are available on the computer. In some cases, a connection to a printer is monitored here as a connection to a named pipe. A shared resource can be a shared directory, a named pipe, a shared printer, or a resource of an unrecognized type. |
| **Folder Path** | Displays the path to the shared resource. |
| **Type** | Displays the type of network connection: Windows, NetWare, or Macintosh. |
| **# Client Connections** | Displays the number of users who are connected to the shared resource. |
| **Description** | Describes the shared resource. |

**Sessions**
The following table shows the information that the Sessions folder contains about all the network users who are connected to the computer.

| Column | Description |
|---|---|
| **User** | Lists the network users who are connected to the computer. |
| **Computer** | Displays the computer name of the connected user's computer. |
| **Type** | Displays the type of network connection: Windows, NetWare, or Macintosh. |
| **# Open Files** | Displays the number of resources that have been opened on this computer by this user. |

| | |
|---|---|
| **Connected Time** | Displays the hours and minutes that have elapsed since this session was established. |
| **Idle Time** | Displays the hours and minutes that have elapsed since this user last initiated an action. |
| **Guest** | Specifies whether this user is connected to this computer as a guest (displayed as Yes or No). |

**Open Files**

The following table shows the information that the Open Files subfolder contains about all the open files on the computer.

| Column | Description |
|---|---|
| **Open File** | Lists the names of open files. An open file could be a file, a named pipe, a print job in a print spooler, or a resource of an unrecognized type. In some cases, a print job is shown here as an open named pipe. |
| **Accessed By** | The name of the user who has opened the file or accessed the resource. |
| **Type** | The type of network connection: Windows, NetWare, or Macintosh. |
| **# Locks** | Displays the number of application-initiated file locks on the resource. |
| **Open Mode** | Displays the permission that was granted when the resource was opened. |

# Understanding Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Shared Folders**
This section covers:

- Share permissions
- Offline settings for shared resources
- Special shared resources

# Share permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Share permissions**
A shared resource provides access to applications, data, or a user's personal data. You can assign or deny permissions for each shared resource.

You can control access to shared resources with a variety of methods. You can use share permissions, which are simple to apply and manage. Or, you can use access control on the NTFS file system, which provides more detailed control of the shared resource and its contents. You can also use a combination of these methods. If you use a combination of these methods, the more restrictive permission always applies. For example, if the share permission is set to Everyone = Read (which is the default), and the NTFS permission allows users to make changes to a shared file, the share permission applies, and the user is not allowed to change the file.

It is not always necessary to explicitly deny a permission to a shared resource. Denying permission is usually necessary only when you want to override specific permissions that are already assigned.

For information about how to set NTFS file permissions, see Set, view, change, or remove permissions on files and folders. For information about best practices for working with Shared Folders, including assigning permissions, see Best practices for Shared Folders.

**Important**

- In the Windows Server 2003 family, when you create a new shared resource, the Everyone group is automatically assigned the Read permission, which is the most restrictive.

Share permissions:

- Apply only to users who gain access to the resource over the network. They do not apply to users who log on locally, such as on a terminal server. In these cases, use access control on NTFS to set permissions.
- Apply to all files and folders in the shared resource. If you want to provide a more detailed level of security to the subfolders or objects in a shared folder, use access control on NTFS.
- Are the only way to secure network resources on FAT and FAT32 volumes, because NTFS permissions are not available on FAT or FAT32 volumes.
- Specify the maximum number of users who are allowed to access the shared resource over the network. This is in addition to the security provided by NTFS.

You can assign the following types of access permissions to shared folders or drives:

**Read**
Read is the default permission that is assigned to the Everyone group. Read allows:

- Viewing file names and subfolder names
- Viewing data in files
- Running program files

**Change**
Change is not a default permission for any group. The Change permission allows all Read permissions, plus:

- Adding files and subfolders

- Changing data in files
- Deleting subfolders and files

**Full Control**

Full Control is the default permission that is assigned to the Administrators group on the local computer. Full Control allows all Read and Change permissions, plus:

- Changing permissions (NTFS files and folders only)

# Offline settings for shared resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Offline settings for shared resources**

You can use Offline Files to work with shared resources even when you are not connected to the network. When you create a shared resource, you can specify if and how files in the shared resource are cached (stored) locally on client computers when other users access the files.

To make shared network resources available offline, Offline Files stores a version of the shared resources in a reserved portion of disk space on the client computer, which is called a file system cache. The client computer can access this cache whether or not the client computer is connected to the network.

You can choose one of the following caching options for each shared resource:

**Only the files and programs that users specify will be available offline.**
This option gives users control over which files are available offline.

**All files and programs that users open from the share will be automatically available offline.**
This option allows all the files that users open from the shared folder to be automatically available offline.

If you select the **Optimized for performance** check box, all programs will be automatically cached so that they can be run locally. This option is especially useful for file servers that host applications, because it reduces network traffic and improves server scalability.

**Files or programs from the share will not be available offline.**
This option prevents users from storing files offline.

**Important**

- When you create a new shared resource, offline access is allowed by default, which means that secure shared resources can be stored offline on potentially insecure computers. For best security, do not allow users to store files offline. If you do, be sure to set permissions appropriately using either share permissions or access control on the NTFS file system. For more information, see Set permissions on a shared resource and Set, view, change, or remove permissions on files and folders.

Offline access may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

For more information about Offline Files, see Offline Files overview. For more information about how to configure offline access for shared resources, see Configure offline settings for a shared resource.

# Special shared resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Special shared resources**

Depending on the configuration of your computer, some or all of the following special shared resources are created for administrative and system use. These shared resources are not visible from My Computer, but you can view them by using Shared Folders. In most cases, it is recommended that you not delete or modify special shared resources.

You can hide other shared resources from users by typing **$** as the last character of the shared resource name (the **$** then becomes part of the resource name). These shared folders are hidden from Windows Explorer like special shared resources, but otherwise they are not special. For more information about how to share a resource, see Share a folder or drive.

If you change permissions on special shared resources, such as ADMIN$, the default settings may be restored when the Server service is stopped and restarted or when the computer is restarted. Note that this does not apply to user-created shared resources whose share name ends in $.

You may see some or all of the following administrative shared resources in the Shares folder in Shared Folders. Other applications may create additional special shared resources. To view special shared resources in the Shares folder, see View information about shared resources.

***drive letter*$**
A shared resource that enables administrators to connect to the root directory of a drive.

**ADMIN$**
A resource that is used during remote administration of a computer. The path of this resource is always the path to the system root (the directory in which the operating system is installed, for example, C:\Windows).

**IPC$**
A resource that shares the named pipes that are essential for communication between programs. You use IPC$ during remote administration of a computer and when you view a computer's shared resources. You cannot delete this resource.

**NETLOGON**
A required resource that is used on domain controllers. Removing this shared resource causes a loss of functionality on all client computers that the domain controller serves.

**SYSVOL**
A required resource that is used on domain controllers. Removing this shared resource causes a loss of functionality on all client computers that the domain controller serves.

**PRINT$**
A resource that is used during remote administration of printers.

**FAX$**
A shared folder on a server that is used by fax clients in the process of sending a fax. The shared folder is used to temporarily cache files and to access cover pages that are stored on the server.

# Administering Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Shared Folders**
This section covers:

- Securing shared resources
- Managing shared folders from the command line

# Securing shared resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing shared resources**

You can control access to shared resources with a variety of methods. You can use share permissions, which are simple to apply and manage. You can use access control on the NTFS file system, which provides more detailed control of a shared resource and its contents. You can also use a combination of these methods. If you use a combination of these methods, the more restrictive permission always applies. For example, if the share permission is set to Everyone = Read (which is the default), and the NTFS permission allows users to make changes to a shared file, the share permission applies, and the user is not allowed to change the file.

For more information about how to set NTFS file permissions, see Set, view, change, or remove permissions on files and folders. For more information about how to set share permissions, see Set permissions on a shared resource. For more information about best practices for working with Shared Folders, see Best practices for Shared Folders.

Fixed disks on your computer, such as drive C or drive D, are automatically shared using the syntax *drive letter$*, such as C$ or D$. These drives do not appear with the hand icon that indicates sharing in My Computer or Windows Explorer, and they are also hidden when users connect to your computer remotely.

If your computer is not protected by a firewall, and someone knows the user name and password of any member of the Administrators group, Backup Operators group, or Server Operators group, that person has the same access to your computer as an administrator.

**Important**

- To keep your drives secure, use a strong password for all accounts. For more information, see Strong passwords.
- For best security, you can also rename the Administrator account. For more information about how to do this, see Accounts: Rename administrator account.
- If you change permissions on special shared resources, such as ADMIN$, the default settings may be restored when the Server service is stopped and restarted or when the computer is restarted. Note that this does not apply to user-created shared resources whose share name ends in $. For more information about special shared resources, see Special shared resources.

# Managing shared folders from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing shared folders from the command line**

In addition to Shared Folders and Windows Explorer, you can use the command-line utilities in the following table to manage shared resources.

| Name | Use to |
|---|---|
| net file | View and control resources that are shared on the network. The Server service must be running for you to use this command. |
| net config | View the maximum number of users who can access a shared resource and the maximum open files per session. |
| Net use | Connect a computer to or disconnect a computer from a shared resource. |
| net session | Manage server connections. |
| net share | Create, delete, manage, and display shared resources. |
| net view | Display information about the domains, computers, or resources that are shared by the specified computer, including the offline client caching settings. |
| Net help | View help for network commands. |

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

# Troubleshooting Shared Folders

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- I cannot access shared resources.
- I cannot view shared resource information in Shared Folders.

**I cannot access shared resources.**
**Cause:** The network connection was lost.

**Solution:** Verify network settings.

**See also:** Troubleshooting network and dial-up connections

**Cause:** Share permissions or file system permissions prevent access.

**Solution:** Verify that the share permissions and file system permissions on the computer to which you are connecting allow access to network resources and that you have access to these resources. If you can access Shared Folders on the remote computer, check the Sessions folder to determine the user rights with which you are logged on to the remote computer. You may be connected as Guest.

You can also verify that the Full Control permission is not denied to the Everyone group. If the Full Control permission is denied to the Everyone group, no one is able to access the shared resource, even if you grant Full Control to the Administrators group.

**See also:** View information about shared resources; Set, view, change, or remove permissions on files and folders

**I cannot view shared resource information in Shared Folders.**
**Cause:** You are not logged on to the specified computer as a member of the Administrators group, Server Operators group, or Power Users group.

**Solution:** Log on to the specified computer with the appropriate user rights.

**Note**

- File sharing options may be limited if simple file sharing is enabled. For more information about simple file sharing, see article Q304040, "Description of File Sharing and Permissions in Windows XP," in the Microsoft Knowledge Base.

# Distributed File System (DFS)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Distributed File System**

This topic describes Distributed File System (DFS) and how to administer it.

- Before creating a distributed file system, see Checklist: Creating a distributed file system.
- For recommendations on implementing DFS, see Best practices for Distributed File System (DFS).
- For help with specific tasks, see Distributed File System (DFS) How To....
- For general background information, see Distributed File System (DFS) Concepts.

# Checklist: Creating a distributed file system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a distributed file system**

| Step | Reference |
|---|---|
| Review distributed file system concepts | Distributed File System (DFS) Concepts |
| Create a DFS root for a distributed file system | Create a DFS root |
| (Optional) If using a domain DFS root, add a root target | Add a DFS root target |
| Add a DFS link to a distributed file system | Add a DFS link |
| Add targets to a DFS link | Add a target |
| (Optional) Configure replication for targets | Configure DFS Replication |
| Manage security on shared folders used by DFS | Securing shared resources |

# Best practices for Distributed File System (DFS)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Use DFS with the NTFS file system only.**

  Do not install DFS on a FAT file system. DFS on the Windows Server 2003 family can only be used with the NTFS file system. For information on converting a volume to NTFS, see Convert.

- **Set the** replication policy**.**

  To maintain a balanced load on your server, you should consider the schedule of synchronization. This requires that you understand the following:

  - The topology between the participating servers.
  - How network performance may be impacted, based on available bandwidth.
  - The quantity of replication traffic likely to occur.
  - Replication schedules for the distributed file system, as well as the replication schedules for any other distributed file systems in the same site.

    Automatic file replication through the File Replication service (FRS) is only available with domain DFS.

- **Do not create FRS replica sets on a volume that is managed by** Remote Storage**.**

  Doing so might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

- **Perform regular administration as follows:**

  After you have created a DFS console for the DFS roots that you administer, save the console file for future use.

  Periodically, perform a status check on common targets to ensure that the targets are still accessible.

- **Ensure that file system permissions are set correctly.**

  In order for a user to access files in a DFS target folder, the user must be allowed access to the target and to all parent folders. In addition, the user must be allowed access to the server hosting the DFS root, and to the link and the link's parent folders. For more information on setting permissions, see Access control overview.

- **Monitor anti-virus or defragmentation software carefully.**

  Anti-virus and defragmentation software can interfere with successful file replication. If you are running anti-virus or defragmentation software on any of the computers that host your domain DFS root or targets, disable it or monitor it carefully until initial file replication is complete.

- **Consider the needs of your users.**

Include commonly-used network shares in your DFS namespace.

For links that point to targets with content that frequently changes,use a shorter cache time-out value to ensure that clients maintain an updated list of targets.

# Distributed File System (DFS) How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Create a DFS root
- Add a DFS root target
- Show a DFS root
- Add a DFS link
- Filter the DFS link display
- Enable or disable DFS referral
- Add a target
- Check status of a DFS root, DFS link, or target
- Hide a DFS root
- Delete a DFS root, root target, link, or target
- Configure DFS Replication

# Create a DFS root

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a DFS root**

1. Open Distributed File System.
2. On the **Action** menu, click **New Root**.
3. Follow the instructions that appear in the New Root wizard.

## Notes

- In order to make changes to an existing DFS namespace, you must be a member of the Administrators group on the server hosting the namespace. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- If the folder that you specify to correspond to the root does not exist, or has not been shared, you will be prompted to create a new shared folder. You can then type the name of a new folder, or type the name of an existing folder that has not yet been shared. DFS will create the folder (if necessary), then share the folder.
- You can host multiple DFS roots on a single computer when using Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition.
- On server clusters, the name you give a non-clustered DFS root on a node's local storage must be different from the name you give a clustered DFS root on the cluster storage.
- To remove a root, select the root, open the **Action** menu, and click **Delete Root**.
- Any member server or domain controller can host a DFS root. You must have be a member of the Administrators group on a server to create a DFS root on that server. In addition, you need domain administrator credentials to create a domain root.
- If your DFS namespace includes roots and root targets which exist on both computers running one of the Windows Server 2003 family of products and computers running Windows 2000, you must administer those roots from a computer running a product in the Windows Server 2003 family, or from a computer running the Windows Server 2003 Administration Tools Pack. You can not administer the DFS namespace from a computer running Windows 2000.
- Only NTFS volumes can host DFS roots. You cannot host a DFS root on a FAT or FAT32 partition.
- When working with shared folders, use NTFS and file sharing permissions to secure the folder and any subfolders so that only authorized users can access them.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Creating a DFS root
Delete a DFS root, root target, link, or target
Add a DFS root target
Securing shared resources
Set permissions on a shared resource
Set, view, change, or remove permissions on files and folders

# Add a DFS root target

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a DFS root target**

1. Open Distributed File System.
2. Select the domain root for which you want to add a root target.
3. On the **Action** menu, click **New Root Target**.
4. Follow the instructions that appear in the New Root wizard.

## Notes

- In order to make changes to an existing DFS namespace, you must be a member of the Administrators group on the server hosting the namespace. As a security best practice, consider using Run as to perform this procedure. You must be a member of the Administrators group on a server to create a root target on that server. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- To remove a root target, select the root target, open the **Action** menu, and click **Remove Target**.
- You can only add root targets to domain DFS roots.
- If the folder that you specify to correspond to the root target does not exist or has not been shared, you will be prompted to create a new shared file. You can then type the name of a new folder, or type the name of an existing folder that has not yet been shared. DFS will create the folder (if necessary), then share the folder. Be aware that although the name of the folder itself does not have to be the same as the root, the folder will automatically be shared with the same name as the root.
- If you add a shared folder on a Windows 2000 computer as a root target, it may not display in the list of root targets immediately.
- If your DFS namespace includes roots and root targets which exist on both computers running one of the Windows Server 2003 family of products and computers running Windows 2000, you must administer those roots from a computer running a product in the Windows Server 2003 family, or from a computer running the Windows Server 2003 Administration Tools Pack. You can not administer the DFS namespace from a computer running Windows 2000.
- When working with shared folders, use NTFS and file sharing permissions to secure the folder and any subfolders so that only authorized users can access them.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a DFS root
Delete a DFS root, root target, link, or target

Securing shared resources
Set permissions on a shared resource
Set, view, change, or remove permissions on files and folders

# Show a DFS root

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To show a DFS root**

1. Open Distributed File System.
2. Click **Distributed File System**.
3. On the **Action** menu, click **Show Root**.
4. In the **Root or host server** field, type in the UNC name for an existing DFS root, or expand the list of trusting domains and click the DFS root, then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- To specify the DFS root, you can type either *HostDomain**DFSRootName* or *HostServer**DFSRootName*, where *HostDomain* is the name of the domain on which the root resides (for domain DFS), *HostServer* is the name of the server on which the root resides (for stand-alone DFS), and *DFSRootName* is the name of the root.
- You can also show a root by typing the name of the host server alone. If the host server has a single root, DFS will show the root. If the host server has multiple roots, DFS will display a list of all the roots on the server and allow you to pick the root or roots to display. You can select more than one root by holding down the SHIFT or CTRL keys as you click.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Add a DFS link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a DFS link**

1. Open Distributed File System.
2. In the console tree, click the DFS root.
3. On the **Action** menu, click **New Link**.
4. Type a name for the new link.
5. Type a path for the new link or click **Browse** to select from the list of available shared folders.
6. Type a comment, if you would like to use it to further identify or describe the link.
7. Type the duration for which a reference to this link will be cached on a DFS client, then click **OK**.

## Notes

- In order to make changes to an existing DFS namespace, you must be a member of the Administrators group on the server hosting the namespace. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- The folder that you specify as the target must be an existing shared folder.
- To delete a link, select the link, open the **Action** menu, and click **Delete Link**.
- If you later delete a target from this link using **dfscmd /remove** on the command line, you must specify the same name you used to create the link. For example, if you create the link using its DNS name, you must specify the DNS name when using **dfscmd**, rather than the NetBIOS name. Consider using a consistent naming standard when adding links, so it will be easier to remember later.
- When working with shared folders, use NTFS and file sharing permissions to secure the folder and any subfolders so that only authorized users can access them.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a target
Delete a DFS root, root target, link, or target
Adding a DFS link
Securing shared resources
Set permissions on a shared resource
Set, view, change, or remove permissions on files and folders

# Filter the DFS link display

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To filter the DFS link display**

1. Open Distributed File System.
2. In the console tree, click the DFS root.
3. On the **Action** menu, click **Filter Links**.
4. In the **Type the maximum number of links** box, type the maximum number of links to display in the console tree.
5. Select **Filter links by name**, and then do one of the following:
   - To display all links with names starting with a certain text string, select **Starts with** in the list box, and then type the text string on which you would like to filter.
   - To display all links with names containing a certain text string, select **Contains** in the list box, and then type the text string on which you want to filter.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- To display all DFS links, type a number greater than or equal to the total number of existing links in the **Type the maximum number of links** box, and select **Do not filter links by name**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a DFS link
Add a DFS link
Delete a DFS root, root target, link, or target

# Enable or disable DFS referral

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable DFS referral**

1. Open Distributed File System.
2. In the console tree, click the DFS link that points to the target for which you would like to change the referral status.
3. Click the desired target. On the **Action** menu, select **Enable or Disable Referral**.

## Notes

- In order to make changes to an existing DFS namespace, you must be a member of the Administrators group on the server hosting the namespace. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- When you disable DFS referral for a target, DFS will not direct users to the target folder until DFS referral is enabled.
- If DFS referral is enabled for a target, selecting **Enable or Disable Referral** will disable it. If DFS referral is disabled, the action will enable it.
- The **DFS Referral** column and **Status** column in the details pane provide information about the referral status of each target. The **DFS Referral** column indicates if DFS referral is enabled and the **Status** column indicates if a particular target is available. If a target can not accessed, the **Status** column will read **Offline**. To see the current status for a target, click it and select **Check Status** from the **Action** menu.
- When working with shared folders, use NTFS and file sharing permissions to secure the folder and any subfolders so that only authorized users can access them.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Check status of a DFS root, DFS link, or target
Checking target status
Securing shared resources
Set permissions on a shared resource
Set, view, change, or remove permissions on files and folders

# Add a target

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a target**

1. Open Distributed File System.
2. In the console tree, click the DFS link to which you want to assign a target.
3. On the **Action** menu, click **New Target**.
4. Under **Path to target (shared folder)**, enter the name of a shared folder, or click **Browse** to select from the list of available shared folders.
5. If you are adding the target to a domain root, click **Add this target to the replication set** to enable automatic file replication for the target.
6. If you are not enabling automatic file replication, you must manually copy files to the new target from other targets.

## Notes

- In order to make changes to an existing DFS namespace, you must be a member of the Administrators group on the server hosting the namespace. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- The folder that you specify as the target must be an existing shared folder.
- To remove a target, select the target, open the **Action** menu, and click **Remove Target**.
- This procedure for adding a target should not be used for adding the first target. The first target is added when you create the DFS link.
- If a link points to another domain DFS, it cannot point to any other targets. To add a target in this case, you must add it to the targeted DFS and not to the link itself.
- Automatic file replication through the File Replication service is only available with domain DFS.
- When working with shared folders, use NTFS and file sharing permissions to secure the folder and any subfolders so that only authorized users can access them.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a DFS root, root target, link, or target
Adding a target
Securing shared resources

Set permissions on a shared resource
Set, view, change, or remove permissions on files and folders

# Check status of a DFS root, DFS link, or target

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To check status of a DFS root, DFS link, or target**

1. Open Distributed File System.
2. In the DFS administration tool, click the DFS root, DFS link, or target whose status you want to check.
3. On the **Action** menu, click **Check Status**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Adding a target
Checking target status

# Hide a DFS root

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To hide a DFS root

1. Open Distributed File System.
2. In the console tree, click the DFS root.
3. On the **Action** menu, click **Hide Root**.
4. A dialog box will appear, asking if you want to hide the root. Click **Yes**.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- The root is removed from the DFS administration tool. The distributed file system itself is not affected. User access to the root is not affected by this procedure.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Delete a DFS root, root target, link, or target

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a DFS root, root target, link, or target**

1. Open Distributed File System.
2. Click the DFS root, root target, link, or target that you want to delete. You might have to navigate through the console tree to reveal the object you would like to delete.
3. On the **Action** menu, click **Delete Root**, **Delete Link**, or **Remove Target**, depending on the nature of the object you would like to delete.
4. A dialog box will appear asking you for confirmation. Click **Yes**.

### Important

- Deleting a root is not a recoverable operation. Deleting a root removes the root from the administration tool and removes the entire DFS structure associated with the root, but does not delete any data. If it is a stand-alone root, the procedure removes the DFS configuration data from the host computer. If it is a domain root, the procedure removes the configuration data from each host computer and from Active Directory.
- Deleting a root, root target, link, or target does not delete any associated shared folders from the file system, and it does not change any permissions associated with files or folders. If you no longer want users to have access to data in these shared folders, use NTFS or file sharing permissions to restrict access.

### Notes

- In order to make changes to an existing DFS namespace, you must be a member of the Administrators group on the server hosting the namespace. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- If you remove the only target of a link, the link will be deleted.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a DFS root
Add a DFS root target
Add a DFS link
Add a DFS link
Add a target
Securing shared resources
Set permissions on a shared resource
Set, view, change, or remove permissions on files and folders

# Configure DFS Replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure replication**

- View replication properties
- Initiate file replication
- Show or hide replication information for a list of targets
- Enable or disable replication for a target
- Modify replication topology
- Set the replication schedule
- Exclude files or subfolders from replication
- Remove files or subfolders from the exclusion list

# View replication properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view replication properties**

1. Open Distributed File System.
2. Click the domain DFS root or link for which replication has been configured.
3. On the **Action** menu, click **Properties**.
4. Click the **Replication** tab.

## Caution

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. Otherwise you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- This procedure is applicable only after you have initiated replication.
- Automatic file replication through the File Replication service is only available with domain DFS.
- From here, you can view and change the replication topology, view and change the replication schedule, and exclude files and subfolders from replication.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a DFS root
Modify replication topology
Set the replication schedule
Exclude files or subfolders from replication

# Initiate file replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To initiate file replication**

1. Open Distributed File System.
2. Click the domain root or link for which you would like to initiate replication.
3. On the **Action** menu, select **Configure Replication**.
4. Follow the instructions that appear in the Configure Replication wizard.

## Caution

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. Otherwise you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

- Do not enable replication on Distributed File System (DFS) roots.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- Automatic file replication through the File Replication service is only available with domain DFS.
- A domain DFS root or link must have at least two targets before you can begin replication.
- To stop replication on a link, click the link, click the **Action** menu, and then click **Stop Replication**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a DFS root
View replication properties

# Show or hide replication information for a list of targets

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To show or hide replication information for a list of targets**

1. Open Distributed File System.
2. Click the DFS link for which you want to display or hide replication information.
3. On the **Action** menu, click **Show Replication Information** (if the "File Replication" column is currently hidden and you want to display it) or **Hide Replication Information** (if the "File Replication" column is currently visible and you want to hide it).

### Caution

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. Otherwise you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- Automatic file replication through the File Replication service is only available with domain DFS.
- The "File Replication" column shows whether each individual target participates in replication.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a DFS root
Initiate file replication
View replication properties
Modify replication topology

# Enable or disable replication for a target

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable replication for a target**

1. Open Distributed File System.
2. Click the DFS link for which you want to manage replication.
3. If the "File Replication" column is not visible in the list of targets, display it by clicking the **Action** menu and clicking **Show Replication Information.**
4. Click the target for which you would like to change the replication status.
5. On the **Action** menu, click **Enable Replication** (if replication is currently disabled and you want to enable it) or **Disable Replication** (if replication is currently enabled and you want to disable it).

   If you are enabling replication for the target, you will be prompted to designate a staging folder.

## Caution

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. Otherwise you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- Automatic file replication through the File Replication service is only available with domain DFS.
- You cannot disable replication on a target that is serving as the hub in a hub-and-spoke replication system without first designating a different target as a hub.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a DFS root
Initiate file replication
View replication properties

Modify replication topology

# Modify replication topology

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify replication topology**

1. Open Distributed File System.
2. Click the root or link for which you want to manage replication.
3. On the **Action** menu, click **Properties**.
4. On the **Replication** tab, click the **Customize** button.
5. In the **Topology** drop-down list, click the type of replication topology you want to use. A list of server connections will appear under **Connections**.
6. Under **Connections**, select or clear the check boxes to enable or disable replication connections between servers.
7. If you choose a custom topology, you must add and delete connections manually. To add a connection, click the **New** button, or, to delete an existing connection, click **Delete**.
8. To change the synchronization priority for initial synchronization or restores, click the **Priorities** button.

## Caution

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. Otherwise you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

## Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- Automatic file replication through the File Replication service is only available with domain DFS.
- The **Replication** tab is only available for domain DFS roots or links for which you have already configured replication.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Show a DFS root
Set the replication schedule

# Set the replication schedule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the replication schedule**

1. Open Distributed File System.
2. Click the domain DFS root or link for which you want to manage replication. On the **Action** menu, click **Properties**.
3. On the **Replication** tab, click **Schedule**. The replication schedule appears, with blue representing the times when replication is available, and white representing the times when replication is not available.
4. On the replication schedule, click to select an hour or drag to select a range of hours.
5. Enable or disable replication for the selected time period or periods as follows:
   - To enable replication, click **Replication Available**.
   - To disable replication, click **Replication Not Available**
6. Repeat the previous two steps as necessary, then click **OK**.

### Caution

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. Otherwise you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- Automatic file replication through the File Replication service is only available with domain DFS.
- Clicking on the name of a day of the week will select the entire day. Clicking directly underneath an hour label will select that hour every day. Clicking **All** will select the entire schedule.
- This method allows you to set the schedule for all existing connections. To set the schedule for individual connections, on the **Replication** tab, click **Customize**, select the connection or connections for which you would like to change the schedule, and click **Schedule**.
- Modifying the schedule for all existing connections will override any changes you have made to the schedule for individual connections.

- To override the replication schedule for the initial synchronization of inbound connections and during file restoration operations, on the **Replication** tab, click **Customize**, click **Priorities** and then check the **Override Schedule** check box for the appropriate inbound connection.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a DFS root
View replication properties
Setting the initial synchronization priorities for inbound connections

# Exclude files or subfolders from replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To exclude files or subfolders from replication**

1. Open Distributed File System.
2. Click the domain DFS root or link for which you want to manage replication. On the **Action** menu, click **Properties**.
3. On the **Replication** tab, next to **File filter** (for files) or **Subfolder filter** (for subfolders), click **Edit**.
4. Type the name of a file or subfolder you want to exclude from replication.
5. Click **Add** to add the file or folder name to the exclusion list, and then click **OK**.

### Caution

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. Otherwise you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- Automatic file replication through the File Replication service is only available with domain DFS.
- When typing the file or folder name to exclude from replication, you can use the asterisk (*) or the question mark (?) as wildcard characters. For example, **.tmp* matches all files with the .tmp extension.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a DFS root
Remove files or subfolders from the exclusion list

# Remove files or subfolders from the exclusion list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove files or subfolders from the exclusion list**

1. Open Distributed File System.
2. Click the domain DFS root or link for which you want to manage replication. On the **Action** menu, click **Properties**.
3. On the **Replication** tab, next to **File filter** or **Subfolder filter**, click **Edit**.
4. Click the file or folder name or specification you would like to remove from the list.
5. Click **Remove** to remove the item from the exclusion list, and then click **OK**.

### Caution

- Do not create File Replication service (FRS) replica sets on a volume that is managed by Remote Storage. Otherwise you might severely impact system performance and possibly cause data loss within your media library.

  FRS might need to periodically read every file in the replica set to send the file contents to another computer. This causes FRS to recall all files that Remote Storage has sent to secondary storage, which might take a long time (hours or days). If you use tape for your secondary storage, remember FRS recalls files in directory order rather than media order, so the excessive number of tape seeks performed by FRS will likely ruin the tapes and cause data loss.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open DFS, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Distributed File System**.
- Automatic file replication through the File Replication service is only available with domain DFS.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Show a DFS root
Exclude files or subfolders from replication

# Distributed File System (DFS) Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Distributed File System overview
- Understanding Distributed File System
- Using Distributed File System

# Distributed File System overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Distributed File System overview**
With Distributed File System (DFS), system administrators can make it easy for users to access and manage files that are physically distributed across a network. With DFS, you can make files distributed across multiple servers appear to users as if they reside in one place on the network. Users no longer need to know and specify the actual physical location of files in order to access them.

For example, if you have marketing material scattered across multiple servers in a domain, you can use DFS to make it appear as though all of the material resides on a single server. This eliminates the need for users to go to multiple locations on the network to find the information they need.

**Reasons for using DFS**
You should consider implementing DFS if:

- You expect to add file servers or modify file locations.
- Users who access targets are distributed across a site or sites.
- Most users require access to multiple targets.
- Server load balancing could be improved by redistributing targets.
- Users require uninterrupted access to targets.
- Your organization has Web sites for either internal or external use.

**DFS types**
You can implement a distributed file system in either of two ways, either as a stand-alone root distributed file system, or as a domain distributed file system.

**Accessing DFS targets from other computers**
In addition to the server-based DFS component of the Windows Server 2003 family, there is a client-based DFS component. The DFS client caches a referral to a DFS root or a DFS link for a specific length of time, defined by the administrator.

The DFS client component runs on a number of different Windows platforms. For information on DFS client versions and their associated platforms, see Platform compatibility.

For more information, see Understanding Distributed File System and Using Distributed File System.

# Understanding Distributed File System

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Distributed File System**
This section covers:

- Distributed File System features
- Distributed File System mapping
- Distributed File System and security
- Platform compatibility

# Distributed File System features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Distributed File System features**
The Distributed File System (DFS) provides several important features, described in the following sections.

**Easy access to files**
A distributed file system makes it easier for users to access files. Users need only go to one location on the network to access files, even though the files may be physically spread across multiple servers.

Also, when you change the physical location of a target, user access to the folder is not impacted. They still access the folder in the same way as before, because the location of the file looks the same.

Users no longer need multiple drive mappings to access their files.

Finally, scheduled file server maintenance, software upgrades and other tasks that normally require taking the server offline can be accomplished without disrupting user access. This is particularly useful for Web servers. By selecting the root for the Web site as a DFS root, you can move resources within the distributed file system without breaking any HTML links.

**Availability**
Domain DFS ensures that users retain access to their files, in two ways:

First, the Windows Server 2003 operating system automatically publishes the DFS mapping to Active Directory. This ensures that the DFS namespace is always visible to users on all servers in the domain.

Second, as an administrator, you can replicate your DFS roots and your targets. Replication means that you can duplicate DFS roots and targets on multiple servers in the domain. In this way, users can still access their files, even if one of the physical servers on which those files reside becomes unavailable.

**Server load balancing**
A DFS root can support multiple targets that are physically distributed across a network. This is useful, for example, if you have a file that you know will be accessed heavily by your users. Rather than all users physically accessing this file on a single server, and thus taxing the server, DFS ensures that user access to the file is distributed to multiple servers. To users, however, the file resides in one location on the network.

**File and folder security**
Because the shared resources DFS manages use standard NTFS and file sharing permissions, you can use preexisting security groups and user accounts to ensure that only authorized users have access to sensitive data.

# Distributed File System mapping

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Distributed File System mapping**

A Distributed File System (DFS) mapping consists of a DFS root, one or more DFS links, and referrals to one or more targets.

The domain server on which a DFS root resides is known as a host server. You can replicate a DFS root by creating root targets on other servers in the domain. This provides file availability when the host server becomes unavailable.

To users, a DFS mapping provides unified and transparent access to the network resources they need. To system administrators, a DFS mapping is a single DNS namespace: With domain DFS, the DNS names for the DFS root targets resolve to the host servers for the DFS root.

Because the host server for a domain distributed file system is a member server within a domain, the DFS mapping is automatically published to Active Directory by default, thus providing synchronization of DFS topologies across host servers. This, in turn, provides fault tolerance for the DFS root and supports optional replication of targets.

You can expand a DFS mapping by adding a DFS link to the DFS root. The only constraint on the number of hierarchical levels in a DFS mapping is imposed by the Windows Server 2003 family limit of 260 characters for any file path. A new DFS link can refer to a target with or without subfolders, or to an entire Windows Server 2003 family volume. If you have adequate permissions, you can also access any local subfolders that exist in or are added to a target.

# Distributed File System and security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Distributed File System and security**

Aside from granting the necessary permissions, Distributed File System (DFS) service does not implement any additional security measures, beyond what the Windows Server 2003 family system provides. The permissions assigned to a DFS root or DFS link determine who can add a new DFS link.

Permissions to a target are not related to the DFS topology. For example, suppose there is a DFS link named MarketingDocs and you have appropriate permissions to access a particular target to which MarketingDocs points. In this case, you can then access all other targets in the set of targets, regardless of whether or not you have permissions to access those other targets. However, having permission to access these targets does determine whether or not you have access to any of the information within the targets. Such access is determined by standard Windows Server 2003 family security controls.

In summary, security is enforced by the underlying file system when a user tries to access a target and its contents. Thus, a FAT volume provides share-level security on files, while an NTFS volume provides full Windows Server 2003 family security. To maintain security, you should use NTFS and file sharing permissions to secure any shared folders used by DFS so that only authorized users can access them.

For more information on setting permissions, see Securing shared resources. For security recommendations for DFS, see Best practices for Distributed File System (DFS).

# Platform compatibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Platform compatibility**

The Windows Server 2003 family supports targets on the following platforms.

| Platform | Host DFS Client | Host DFS Root |
|---|---|---|
| MS-DOS, Windows 3.x, Windows for Workgroups, and NetWare servers | No | No |
| Windows 95 | Yes, download client | No |
| Windows 98 | Yes, stand-alone client included; download client for domain | No |
| Windows NT 4.0 and Service Pack 3 | Yes, stand-alone client included | Yes, stand-alone root DFS only |
| Windows 2000 and the Windows Server 2003 family | Yes, client included | Yes, stand-alone and domain DFS |

**Note**

- To support synchronization of targets, the referenced resource for a target must be located on a Windows Server 2003 family NTFS partition.

# Using Distributed File System

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Distributed File System**

This section covers:

- Creating a DFS root
- Adding a DFS root target
- Adding a DFS link
- Adding a target
- Using File Replication Service
- Checking target status
- Managing Distributed File System from the command line

# Creating a DFS root

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating a DFS root**
You can create a DFS root on an NTFS partition residing on a computer in the Windows Server 2003 family. When creating a DFS root, you have the option of establishing either a stand-alone DFS root or a domain DFS root.

A stand-alone DFS root:

- Does not use Active Directory.
- Cannot have more than one root-level target.
- Does not support automatic file replication using the File Replication service.
- Supports fault tolerance through Server Clusters.

A domain DFS root:

- Must be hosted on a domain member server.
- Has its DFS namespace automatically published to Active Directory.
- Can have more than one root-level target.
- Supports automatic file replication through FRS.
- Supports fault tolerance through FRS.

Using the DFS administration tool, you specify a target to assign as the DFS root. In addition to accessing this target, users can access any subfolders of this target.

You can host multiple DFS roots on a single computer when using Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.

On server clusters, the name you give a non-clustered DFS root on a node's local storage must be different from the name you give a clustered DFS root on the cluster storage.

Large domain-based DFS namespaces might cause significantly increased network traffic due to the size of the DFS Active Directory object. As a result, it is recommended that you use fewer than 5,000 DFS links for a domain root. The largest recommended namespace size for a stand-alone root is 50,000 links on servers running Windows Server 2003.

If your DFS namespace includes roots and root targets which exist on both computers running one of the Windows Server 2003 family of products and computers running Windows 2000, you must administer those roots from a computer running a product in the Windows Server 2003 family, or from a computer running the Windows Server 2003 Administration Tools Pack. You can not administer the DFS namespace from a computer running Windows 2000.

For more information, see Create a DFS root.

# Adding a DFS root target

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding a DFS root target**

Having more than one root target for a domain root helps ensure that if the host server becomes unavailable for any reason, the distributed file system associated with that DFS root is still available to domain users.

To add a root target, click the root and select **New Root Target** from the **Action** menu. This invokes the New Root Wizard.

**Note**

- You can also delete a root target.

# Adding a DFS link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding a DFS link**

A Distributed File System (DFS) link is an element in the DFS namespace that lies below the root and maps to one or more targets. Instead of seeing the target's actual name and physical location, users accessing a DFS namespace see the link name listed as a folder under the root. Because link names are not constrained by the target's name or location, you can create link names that are meaningful to the user. For example, you could create a link named "Marketing" that points to one or more targets that contain marketing information regardless of the actual name and location of those targets.

Large domain-based DFS namespaces might cause significantly increased network traffic due to the size of the DFS Active Directory object. As a result, it is recommended that you use fewer than 5,000 DFS links for a domain root. The largest recommended namespace size for a stand-alone root is 50,000 links on servers running Windows Server 2003.

For more information, see Add a DFS link.

# Adding a target

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding a target**

For each DFS link, you can create a set of targets to which the DFS link points. Within a set of targets, you add the first folder to the set when you create the DFS link. Subsequent targets are added using the DFS administration tool's **New Target** dialog box.

Within a set of targets, you can choose which folders will participate in replication. If you add targets for domain DFS roots, you also need to do the following:

- Set the replication policy among the targets.
- Check the replication status.

For information on setting replication policy, see Using File Replication Service.

For information on checking replication status, see Checking target status.

# Using File Replication Service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using the File Replication service**
This section covers:

- Configuring replication topology
- Configuring the replication schedule
- Excluding files and folders from replication
- Setting the initial synchronization priorities for inbound connections

# Configuring replication topology

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring replication topology**

The word "topology" typically refers to the relationship between components in a network. The replication topology describes the logical connections that FRS uses to replicate files between computers.

For each DFS root or DFS link, you can specify one of four basic topology types:

- In a ring topology, files replicate from one computer to another in a circular configuration, with each computer connected to the computers on either side of it in the ring.
- In a hub and spoke topology, one computer is designated as the hub and is connected to each of the other computers, called spokes. Files replicate from the hub computer to the spoke computers and vice versa, but files do not replicate directly between two spoke computers. You must choose which computer will serve as the hub.
- In a full mesh topology, every computer is connected to every other computer. A file created on one computer replicates directly to all the others.
- In a custom topology, you create the connections between computers yourself.




**Ring Topology**




**Hub and Spoke Topology**




**Full Mesh Topology**

When you select a topology, you can further customize it by selectively enabling or disabling the connections between computers. You can completely disable the relationship between two computers to prevent files from replicating between them at all, or you can make it one-way by disabling the connection from the first computer to the second while leaving the opposite connection available.

Many factors can affect the type of topology that is appropriate for your network, including bandwidth, security, geographic, and organizational considerations.

### Note

- Automatic file replication is available only for domain roots.

# Configuring the replication schedule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring the replication schedule**

When you first configure a /DFS root or a DFS link for replication, the File Replication service (FRS) schedules replication to take place around the clock. Whenever you copy a file to a target that participates in replication, FRS replicates the file to the other targets using the connections specified in the replication topology.

You might not want replication to take place automatically at certain times. For example, if certain servers are taken off the network every Friday evening from 9 to 11 for maintenance and restarting, you might want to halt replication to and from those servers at 9 on Friday and resume it two hours later. When a connection returns to replication, FRS immediately reconciles any differences between the two shares.

FRS lets you schedule the times when replication will or will not be available. You can modify the schedule for individual connections or for all connections at once. The **Schedule** dialog box displays the schedule for every hour in the week and lets you select blocks of time during which replication will or will not take place. Clicking the name of a weekday will select the entire day. Clicking a box at the top of the schedule to the right of the word **All** will select the same hour on every day throughout the week. Clicking **All** will select the entire schedule.

You can access the schedule in two ways. Clicking the **Schedule** button on the Replication tab of a root's or link's **Properties** sheet lets you view and change the schedule for all of the connections belonging to that root or link. If you select one or more connections in the **Customize Topology** dialog box and click the **Schedule** button, you can change the schedule for the selected connections only.

### Notes

- Automatic file replication through the File Replication service is only available with domain DFS.
- Changing the schedule for all connections will overwrite any changes you have made to the schedule for individual connections.

# Excluding files and folders from replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Excluding files and folders from replication**
Sometimes you may not want to replicate every file or subfolder contained within a DFS root or DFS link. For example, if a shared folder used as a target contains several temporary files created by applications for caching or record-keeping purposes, you might not want to replicate those files to all of the targets that participate in replication.

For each root and link, you can use the File Replication service to specify files and subfolders to exclude from replication. The exclusion rules apply to every target or root target designated as belonging to the selected root or link, and also apply to subfolders beneath the shared folders themselves.

You can use * and ? as wildcard characters when entering file or folder names. For example, *.tmp matches all files with the .tmp extension.

For specific instructions on excluding files and folders from replication, see:

- Exclude files or subfolders from replication
- Remove files or subfolders from the exclusion list

**Note**

- Automatic file replication is only available for domain roots.

# Setting the initial synchronization priorities for inbound connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting the initial synchronization priority for inbound connections**

It is often advantageous to control the sequence of synchronization during replication. Using Distributed File System (DFS) snap-in and the File Replication Service (FRS), you can configure synchronization priorities. You can determine which inbound connection will go first in the initial synchronization that occurs when you add a new member to a replica set or when you restore one of your replica set members. For example, you can adjust synchronization priorities to ensure that computers across a high-bandwidth network connection synchronize first, followed by synchronization with computers that are communicating through a lower-bandwidth connection. By adjusting the synchronization priorities, you can often reduce the time needed to perform the synchronization process.

Using FRS, you can choose from three synchronization priorities:

- High: All connections marked High must successfully synchronize before a synchronization is attempted on Medium-priority connections.
- Medium: At least one Medium priority connection must successfully complete an initial synchronization before any synchronization is attempted on Low-priority connections. FRS attempts to synchronize all connections in the Medium priority level, but only one must be successful before FRS attempts to synchronize Low priority connections.
- Low: FRS attempts to synchronize with each computer on a Low priority connection one time, but any failure will not delay other synchronization attempts.

Low is the default synchronization priority for all connections. Typically, you should configure high-bandwidth or local connections to have a higher synchronization priority than low-bandwidth connections. Therefore, depending on the scenario, a priority of either Medium or High is appropriate for your local connections. However, be aware that if you have a High priority connection that fails, no other synchronization will take place.

The following two scenarios describe how you might use synchronization priorities:




In the first example, you are using a hub and spoke topology. Computer X at a branch site fails and needs to be restored. Computer X has inbound connections from Computers A and B, which are located at a hub site. Computer X also has an inbound connection from Computer Y, which is located at the same branch site as Computer X. You should set the synchronization priority from Computer Y to Computer X to High. Set synchronization priorities from Computers A and B at the hub site to Computer X at the branch site to Low.

The reason for these settings is that under normal conditions, you would not want to do a full synchronization over the slow link between Computer X and Computer A or Computer B. A synchronization from the hub site would take too long and consume all of the bandwidth of the slow link. Instead, the recovery of Computer X must occur first from the local Computer Y. Computer X can then try to synchronize from either Computer A or Computer B, but if communication problems prevent the synchronization, Computer X can still come online with the same content as Computer Y.




For the second example, you are using a custom topology that consists of a two-level hub and spoke architecture. Computer A is at the hub site. The second level consists of Computers J, K, L, and M at the regional offices. At the bottom level are several departmental computers in branch offices.

Computer A at the hub site has high bandwidth bi-directional connections to Computers J, K, L and M. At the regional offices, Computers J, K, L and M are connected in a medium bandwidth bi-directional ring. The branch office computers are each connected to two of the regional computers (J, K, L and M) by a low-bandwidth connection.

Computers J, K, L and M are all expected to be up to date. Therefore, for Computer A, you should set the inbound connections from Computers J, K, L and M to Medium priority. You should also set the priority for the ring of connections between the regional office computers to Medium. The priority from the branch office computers to Computers J, K, L and M should be set to Low.

Consider Computer K. Its inbound connections from A, J and L should all be set to Medium priority because a synchronization with any one of them would be sufficient for Computer K to come online. If the hub site (Computer A) is unavailable, Computers J, K, L and M can still complete a recovery without first synchronizing with Computer A. The justification for the Low priority from the branch office computers is the same as described in the first example.

**Overriding the schedule for initial synchronization**

Typically, synchronization will only take place at times defined in the synchronization schedule that you configure. However, you can override the schedule for initial synchronization or the initial creation of replica set member. By selecting the Override Schedule check box in the Priorities dialog box, you can cause the initial replication to occur immediately.

For more information on setting the priority, see Modify replication topology. For more information on setting the replication schedule, see Set the replication schedule.

# Checking target status

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checking target status**

You can check whether the targets that you have assigned are being referenced properly by DFS. The status check indicates one of the following conditions:

Checking status on a DFS root or DFS link will reveal one of the following icons:

- A blue check mark on the root or link indicates that the root or link can be reached, and all of its targets can be reached.
- A yellow exclamation point on the root or link indicates that the root or link can be reached, but not all of its targets can be reached, either because DFS referral is disabled on that target, or some other problem is preventing access to the target.
- A red cross on the root or link indicates that the root or link cannot be reached.

Checking status on a target will result in the **Status** column in the details pane indicating **Online** or **Offline**. If the status of a target is **Offline**, that indicates that the target cannot be reached.

# Managing Distributed File System from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing DFS from the command line**

You can use the **dfscmd** command-line utility to manage and configure DFS and DFS roots, links, and targets, and to write batch scripts so that other applications can interact with DFS. You can run the command-line utility from a command prompt by typing **dfscmd**.

For help using the DFS command line utility, see Dfscmd.

For information about other command-line utilities, see Command-line reference A-Z.

# Synchronization Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Synchronization Manager**

Using Synchronization Manager, you can control when your offline files are synchronized with files on the network.

- For general background information, see Synchronization Manager overview.
- For help with synchronizing offline items, see Synchronizing offline items.
- For help with scheduling synchronization, see Scheduling offline items to synchronize.
- For help with synchronizing items when users log on or log off, see Synchronizing offline items when you log on to or log off from a computer.
- For help with synchronizing items while your computer is idle, see Synchronizing offline items while your computer is idle.

# Synchronization Manager overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Synchronization Manager overview**
Using Synchronization Manager, you can control when your offline files are synchronized with files on the network. This ensures that you have the latest information from your network or from the Internet when you need it, while minimizing the disruption to work on your system.

You can use Synchronization Manager to automatically synchronize the information that is available to you offline in the following ways:

- Every time you log on or log off your computer
- At specific intervals while your computer is idle
- At scheduled times

You can apply these options in various combinations to offline files for each shared resource.

Synchronization Manager compares items on the network to those that you opened or updated while working offline, and then makes the most current version available to both the computer and the network. Items that you can synchronize include individual files, folders and their contents, and offline Web pages. Synchronization Manager provides a single location where you can go to synchronize shared files that are available offline, regardless of the type of file or program that was used to make the file available offline.

For more information about using offline items, see Offline Files and Synchronization Manager.

**Note**

- You must share files before you can use them with Offline Files. Once shared, the files appear in the list of items that are available for synchronization.

# Synchronizing offline items

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To synchronize offline items**

1. Open Synchronization Manager.
2. Under **Select the check box for any items you want to synchronize**, select the check box for each offline item that you want to synchronize, such as a folder on a mapped network drive or an Internet Explorer offline Web page.
3. Click **Synchronize**.

## Notes

- To open Synchronization Manager from the command line, click **Start**, click **Run**, type **mobsync**, and then click **OK**.
- You must share files before you can use them with Offline Files. Once shared, the files appear in the list of items that are available for synchronization. For more information about working with Offline Files, see Related Topics.
- To synchronize individual files, folders, or Web pages in Microsoft Windows Explorer, My Computer, or Microsoft Internet Explorer, click the item that you want to synchronize, and then, on the **Tools** menu, click **Synchronize**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Offline Files
Synchronization Manager overview
Synchronizing offline items when you log on to or log off from a computer
Synchronizing offline items while your computer is idle
Scheduling offline items to synchronize

# Scheduling offline items to synchronize

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To schedule offline items to synchronize**

1. Open Synchronization Manager.
2. Click **Setup**, and then click the **Scheduled** tab.
3. Click **Add** to start the Scheduled Synchronization Wizard, and then follow the steps in the wizard.

## Notes

- To open Synchronization Manager from the command line, click **Start**, click **Run**, type **mobsync**, and then click **OK**.
- To synchronize individual files, folders, or Web pages in Windows Explorer, My Computer, or Internet Explorer, click the item that you want to synchronize, and then, on the **Tools** menu, click **Synchronize**.
- To change or remove an existing synchronization schedule, on the **Scheduled** tab, click the existing schedule, and then click either **Edit** or **Remove**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Synchronization Manager overview
Synchronizing offline items when you log on to or log off from a computer
Synchronizing offline items while your computer is idle
Synchronizing offline items

# Synchronizing offline items when you log on to or log off from a computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To synchronize offline items when you log on to or log off from a computer**

1. Open Synchronization Manager.
2. Click **Setup**, and then click the **Logon/Logoff** tab.
3. In **When I am using this network connection**, click the network connection that you want to use.
4. In **Synchronize the following checked items**, select the check box for the offline items that you want to synchronize, such as a folder on a mapped network drive or an Internet Explorer offline Web page.
5. Under **Automatically synchronize the selected items**, do any of the following:
   - To synchronize items when you log on to your computer, select the **When I log on to my computer** check box.
   - To synchronize items when you log off of your computer, select the **When I log off my computer** check box.
   - To request permission before synchronizing your offline items, select the **Ask me before synchronizing the items** check box.

## Notes

- To open Synchronization Manager from the command line, click **Start**, click **Run**, type **mobsync**, and then click **OK**.
- To synchronize individual files, folders, or Web pages in Windows Explorer, My Computer, or Internet Explorer, select the item that you want to synchronize, and then, on the **Tools** menu, click **Synchronize**.
- By choosing when offline items are synchronized, you can better manage the work on your computer and on the network. For example, when connected to the network by modem, synchronization time is minimized by clearing all automatic synchronization options and manually synchronizing individual files.
- When using offline files, it is recommended that you always synchronize when you log on to your computer. This ensures that changes that were made on your computer are synchronized with changes that were made on the network while you were disconnected.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Synchronization Manager overview
Synchronizing offline items while your computer is idle
Scheduling offline items to synchronize
Synchronizing offline items

# Synchronizing offline items while your computer is idle

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To synchronize offline items while your computer is idle**

1. Open Synchronization Manager.
2. Click **Setup**, and then click the **On Idle** tab.
3. In **When I am using this network connection**, click the network connection that you want to use.
4. In **Synchronize the following checked items**, select the check box for the offline items that you want to synchronize, such as a folder on a mapped network drive or an Internet Explorer offline Web page.
5. Select the **Synchronize the selected items while my computer is idle** check box.

## Notes

- To open Synchronization Manager from the command line, click **Start**, click **Run**, type **mobsync**, and then click **OK**.
- To synchronize individual files, folders, or Web pages in Windows Explorer, My Computer, or Internet Explorer, click the item that you want to synchronize, and then, on the **Tools** menu, click **Synchronize**.
- You can create a different synchronization setup for each of your network connections.
- To specify whether to synchronize offline items while your computer is running on batteries, open Synchronization Manager, click **Setup**, and then click the **On Idle** tab. In **When I am using this network connection**, click the network connection that you want to use, click **Advanced**, and then select the items that you want to synchronize and for which you want to prevent synchronization when the computer is on battery power.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Synchronizing offline items when you log on to or log off from a computer
Scheduling offline items to synchronize
Synchronizing offline items

# Managing Disks and Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing disks and volumes**

- Disk Management
- Disk Defragmenter
- Disk Quotas

# Disk Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk Management**
You use Disk Management in Windows XP Professional and Windows Server 2003 operating systems to perform disk-related tasks, such as creating and formatting partitions and volumes, and assigning drive letters. On computers running Windows Server 2003 operating systems, you can also use Disk Management to perform advanced tasks, such as creating and repairing fault tolerant volumes. In addition, you can use the **DiskPart** command, along with other command-line utilities, to perform Disk Management tasks. For information about Disk Management-related command-line utilities, see Managing disks and volumes from the command line.

- Checklist: Adding a new disk.
- To find features that have been moved from where they were located in previous versions of Windows, see New ways to do familiar tasks for Disk Management.
- For tips about using Disk Management, see Best practices for Disk Management.
- For help with specific tasks, see Disk Managment How To....
- For general background information, see Disk Management Concepts.
- For problem-solving instructions, see Troubleshooting Disk Management.

# Checklist: Adding a new disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Adding a new disk**

| Step | Reference |
|---|---|
| Review key concepts. | Disk Management overview; Best practices for Disk Management |
| Confirm that you are a member of the Backup Operators group or the Administrators group.<br><br>If you have access to the administrator account and password, it is a security best practice to log on as a user that is not a member of the Administrators group and run Disk Management by right-clicking the Computer Management snap-in icon, and then clicking Run as. | Default local groups; Why you should not run your computer as an administrator; Run a program with administrative credentials |
| Add the new disk. | Adding a disk |
| Create partitions or volumes. | Partition styles;Create a partition or logical drive |
| Format volumes on the disk. | Choosing a file system: NTFS, FAT, or FAT32;Format a basic volume;Format a dynamic volume |

# New ways to do familiar tasks for Disk Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for administering disks and volumes. The user interface for performing these tasks is different in this version of Windows from the way it was in Windows NT version 4.0.

### Note

- There are no significant user interface differences between Windows 2000, Windows XP, and Windows Server 2003 operating systems for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows use |
|---|---|---|
| Format a volume | **Disk Administrator** in Administrative tools. | Computer Management (Local).<br><br>For more information, see Format a dynamic volume. |
| Add a new disk | **Disk Administrator** in Administrative tools. | Computer Management (Local).<br><br>For more information, see Adding a disk |
| Create a volume set | **Disk Administrator** in Administrative tools. | Computer Management (Local). For more information, see Create a spanned volume. |
| Create a stripe set | **Disk Administrator** in Administrative tools. | Computer Management (Local).<br><br>For more information, see Create a striped volume. |
| Create a mirror set | **Disk Administrator** in Administrative tools. | Computer Management (Local).<br><br>For more information, see Create a mirrored volume. |
| Create a stripe set with parity | **Disk Administrator** in Administrative tools. | Computer Management (Local).<br><br>For more information, see Create a RAID-5 volume. |

# Best practices for Disk Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

### Best practices

#### Use the appropriate credentials

- When using Disk Management, confirm that you are a member of the Backup Operators group or the Administrators group.

  If you have access to the administrator account and password, it is a security best practice to log on as a user that is not a member of the Administrators group and run Disk Management by right-clicking the Computer Management snap-in icon, and then clicking Run as.

  For more information about why you should not run your computer as an administrator, see Why you should not run your computer as an administrator.

#### Back up data

- Because deleting or creating partitions or volumes destroys any existing data, be sure to back up the disk contents beforehand. As with any major change to disk contents, you should back up the entire contents of the hard disk before working with partitions or volumes, even if you do not plan to make any changes to one or more of your partitions or volumes.

  For more information on how to back up and restore data, see Backup.

#### Format volumes using the NTFS file system

- Many features in Windows 2000 and Windows Server 2003 operating systems, such as file and folder permissions, encryption, large volume support, and sparse file management, require the NTFS file system format.

  For more information about file system compatibility and limitations, see Choosing a file system: NTFS, FAT, or FAT32.

#### Use dynamic disks

- Several Disk Management tasks, including creating fault-tolerant disks, can be performed only with dynamic disks. You can use dynamic disks to create and delete simple, spanned, striped, mirrored, and RAID-5 volumes. Dynamic disks cannot be directly accessed by computers running MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT, or Windows XP Home Edition. As a result, you cannot start these operating systems on dynamic disks. Shared folders on dynamic disks are available across a network to computers running all of these operating systems, however.

  For more information about dynamic disks, see Dynamic disks and volumes.

#### Use disks and disk controllers that are the same model and size, and from the same manufacturer

- When creating or repairing mirrored or RAID-5 volumes, it is best to use disks that are the same model and size and from the same manufacturer. This ensures that the disk geometries are the same and simplifies the processes of creating a new mirrored or RAID-5 volume or replacing a failed disk. It is also recommended that you have spare disks and disk controllers available so that if a disk or disk controller fails, you can quickly replace the faulty disk or disk controller with one of the same type.

#### Use Event Viewer to check the system log when mirrored or RAID-5 volumes fail

- When a mirrored or RAID-5 volume fails, it is sometimes due to a faulty disk controller. Changing one of the disks that make up a mirrored or RAID-5 volume will not fix the volume if the problem is being caused by a faulty disk controller. Always use Event Viewer to check the system log to determine the true nature of the disk-related problem before changing a disk or disk controller.

  For more information on how to use Event Viewer to check the system log, see Event Viewer.

# Disk Managment How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Manage Disks
- Manage Basic Volumes
- Manage Dynamic Volumes

# Manage Disks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage disks**


- Initialize new disks
- Move disks to another computer
- View disk properties
- Update disk information
- Refresh drive letter, file system, and volume information
- Reactivate a missing or offline dynamic disk
- Change a basic disk into a dynamic disk
- Change a dynamic disk back to a basic disk
- Change a master boot record disk into a GUID partition table disk
- Change a GUID partition table disk into a master boot record disk
- Manage disks on a remote computer
- Remove a missing dynamic disk

# Initialize new disks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To initialize new disks**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the disk you want to initialize, and then click **Initialize Disk**.
4. In the **Initialize Disk** dialog box, select the disk(s) to initialize.

   On x86-based computers running Windows Server 2003 with Service Pack 1 (SP1), Itanium-based computers, or x64-based computers, you can select whether to use the master boot record (MBR) or GUID partition table (GPT) partition style.

   The disk is initialized as a basic disk.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- New disks appear as **Not Initialized**. Before you can use a disk, you must first initialize it. If you start Disk Management after adding a disk, the Initialize Disk Wizard appears so you can initialize the disk.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Update disk information
Disk status descriptions
Working with MMC console files

# Move disks to another computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move disks to another computer**

1. Print this procedure or write down the steps before attempting to move disks from one computer to another.

2. Use Disk Management to make sure the status of the volumes on the disks is **Healthy**. If the status is not **Healthy**, you should repair the volumes before you move the disks. For more information about repairing disks, see Related Topics.

   **How?**

   a. Open Computer Management (Local).

   b. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.

      Volume status is displayed in the **Volume List** and **Graphical View**:

      In the **Volume List** view, volume status is displayed in the **Status** column.

      In the **Graphical View**, volume status is displayed under the volume size and file system information.

3. Uninstall the disks you want to move using Device Manager.

   **How?**

   a. Open Computer Management (Local).

   b. In the console tree, click **Computer Management**, click **System Tools**, and then click **Device Manager**.

   c. In the device list, double-click **Disk drives**.

   d. Right-click the disks you want to uninstall, and then click **Uninstall**.

   e. In the **Confirm Device Removal** dialog box, click **OK**.

4. For basic disks, skip to step 5. For dynamic disks, use Disk Management to remove the disks you want to move.

   **How?**

   a. Open Computer Management (Local).

   b. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.

   c. Right-click the disks that you want to move, and then click **Remove Disk**.

   d. If the disks are external, unplug them from the computer. If the disks are internal, turn off the computer and then physically remove the disks.

5. If the disks are external, plug them into the computer. If the disks are internal, make sure the computer is turned off and then physically install the disks in that computer.

6. Start the computer that contains the disks you moved.

7. Follow the instructions in the **Found New Hardware** dialog box.

   If the **Found New Hardware** dialog box does not appear, click **Start**, point to **Settings**, click **Control Panel**, and then click **Add Hardware** to start the Add Hardware Wizard.

8. Use Disk Management to detect the new disks.

   **How?**

   a. Open Computer Management (Local).

   b. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.

   c. Click **Action**, and then click **Rescan Disks**.

   d. Right-click any disk marked **Foreign**, click **Import Foreign Disks**, and then follow the instructions on your screen.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- When moved to another computer, basic volumes receive the next available drive letter on that computer. Dynamic volumes retain the drive letter they had on the previous computer. If a dynamic volume did not have a drive letter on the previous computer, it does not receive a drive letter when moved to another computer. If the drive letter is already used on the computer where they are moved, the volume receives the next available drive letter. If an administrator has used the **mountvol /n** or the **diskpart automount** commands to prevent new volumes from being added to the system, volumes moved from another computer are prevented from being mounted and from receiving a drive letter. To use the volume, you must manually mount the volume and assign it a drive letter using Disk Management or the **DiskPart** and **mountvol** commands.

- If you are moving spanned, striped, mirrored, or RAID-5 volumes, it is highly recommended that you move all disks containing the volume together. Otherwise, the volumes on the disks cannot be brought online and will not be accessible except to delete them.

- Do not move multidisk volumes from computers running Windows NT 4.0 to computers running Windows XP Professional or Windows Server 2003 operating systems. These operating systems do not support mirror sets, stripe sets, stripe sets with parity, or volume sets. Instead, you must back up and then delete these volumes before you move the disks, and then create new dynamic volumes and restore the data to the new volumes.

- Every time you remove or import disks to a computer, you must click **Action**, click **Rescan Disks**, and then verify that the disk information is correct.

- You can move multiple disks from different computers to a computer by installing the disks, opening Disk Management, right-clicking any of the new disks, and then clicking **Import Foreign Disks**. When importing multiple disks from different computers, always import all of the disks from one computer at a time. For example, if you want to move disks from two computers, import disks from the first computer and then import disks from the second computer.

- Disk Management describes the condition of the volumes on the disks before they are imported. Review this information carefully. If there are any problems, this will tell you what will happen to each volume on these disks once the disks have been imported. For more information, see Related Topics.

- You can only move GUID partition table (GPT) disks that are used for data storage among x86-based computers running Windows Server 2003 with Service Pack 1 (SP1), Itanium-based computers, or x64-based computers. If you move a GPT disk containing the Windows operating system from an Itanium-based computer to an x86-based or x64-based computer, you can access the data, but you cannot boot from that operating system.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Volume status descriptions
Disk status descriptions
Troubleshooting Disk Management
Working with MMC console files

# View disk properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view disk properties**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. In the graphical view or disk list, right-click a disk and then click **Properties**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- The **Properties** dialog box lists information about the physical disk and the volumes it contains. If you want to view the properties of one of the volumes listed, click the volume and then click **Properties**.
- You can view the disk's partitioning style on the **Volumes** tab in the **Properties** dialog box. Disks with the MBR partitioning style are displayed as Master Boot Record (MBR) disks, and disks with the GPT partitioning style are displayed as GUID Partition Table (GPT) disks.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

4. At the **DISKPART** prompt, type:

   **select disk** *n*

   Select the disk *n* that you want to display additional details about.

5. At the **DISKPART** prompt, type:

   **detail disk**

   **DiskPart** displays information about the disk, such as the disk's type, volumes on the disk, drive letter, file system, size, and status.

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** *n* | Selects the specified disk *n*, and gives it focus. |
| **detail disk** | Displays the properties of the selected disk and the volumes on that disk. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- For more information about **DiskPart**, see related topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
DiskPart
Manage Disks

# Update disk information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To update disk information**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Click **Action**, and then click **Rescan Disks**.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Use Device Manager to scan your computer's hardware for changes. Then, use the **Rescan Disks** command if your disk configuration has changed but the changes do not appear in Disk Management. For more information about Device Manager, see Related Topics.
- When Disk Management rescans disks, it scans all attached disks for disk configuration changes. It also updates information on removable media, CD-ROM drives, basic volumes, file systems, and drive letters.
- Rescanning disks can take several minutes, depending on the number of hardware devices installed.
- If you add a new disk to your computer and Disk Management does not detect it after you click **Rescan Disks**, you might need to restart your computer. New disks appear as **Not Initialized**. Before you can use a disk, you must first initialize it. For instructions describing how to initialize a disk, see Related Topics. If you start Disk Management after adding a disk, the Initialize Disk Wizard appears so you can initialize the disk.
- If you create partitions using another program or utility and Disk Management does not detect the changes, you must close and then restart Disk Management.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **rescan**

| Value | Description |
|---|---|
| **rescan** | Locates new disks that may have been added to the computer. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Use Device Manager to scan your computer's hardware for changes. Then, use the **DiskPart rescan** command if your disk configuration has changed but the changes do not appear in DiskPart. For more information about Device Manager, see Related Topics.

- When DiskPart rescans disks, it scans all attached disks for disk configuration changes. It also updates information on removable media, CD-ROM drives, basic volumes, file systems, and drive letters.

- Rescanning disks can take several minutes, depending on the number of hardware devices installed.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Initialize new disks
Refresh drive letter, file system, and volume information
Change a basic disk into a dynamic disk
Partition styles
Working with MMC console files
DiskPart
Reinstall a Plug and Play device
Device Manager overview

# Refresh drive letter, file system, and volume information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To refresh drive letter, file system, and volume information**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Click **Action**, and then click **Refresh**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- When you click the **Refresh** command, Disk Management refreshes drive letter, file system, volume, and removable media information, and also checks to see if unreadable volumes are now readable. However, it does not scan hardware. To update hardware information, click **Rescan Disks** instead.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Update disk information
Working with MMC console files

# Reactivate a missing or offline dynamic disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To reactivate a missing or offline dynamic disk

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the disk marked **Missing** or **Offline**, and then click **Reactivate Disk**.

   The disk should be marked **Online** after the disk is reactivated.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- A dynamic disk may become **Missing** when it is corrupted, powered down, or disconnected.
- A dynamic disk may become **Offline** if it is corrupted or intermittently unavailable. A dynamic disk may also become **Offline** if you attempt to import a foreign (dynamic) disk and the import fails. An error icon appears on the **Offline** disk. Only dynamic disks display the **Missing** or **Offline** status.
- Only dynamic disks can be reactivated.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the missing or offline disk that you want to bring online. Missing disks are numbered M0, M1, M2, etc. and offline disks are numbered 0, 1, 2, etc.

4. At the **DISKPART** prompt, type:

   **select disk***n*

5. At the **DISKPART** prompt, type:

   **online**

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** *n* | Sets the focus to the specified disk, where *n* is the disk number. |
| **online** | Brings an offline disk or volume with focus online. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- A dynamic disk may become **Missing** when it is corrupted, powered down, or disconnected.

- A dynamic disk may become **Offline** if it is corrupted or intermittently unavailable. A dynamic disk may also become **Offline** if you attempt to import a foreign (dynamic) disk and the import fails. Only dynamic disks display the **Offline** status.

- Only dynamic disks can be reactivated.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Partition styles
Disk status descriptions
Working with MMC console files
DiskPart

# Change a basic disk into a dynamic disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a basic disk into a dynamic disk**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the basic disk you want to convert, click **Convert to Dynamic Disk**, and then follow the instructions on your screen.

   If you do not see this menu item, you might be right-clicking a volume instead of a disk, the disk might have been previously converted to a dynamic disk, or the computer is a portable computer. (Dynamic disks are not supported on portable computers, removable disks, detachable disks that use Universal Serial Bus (USB) or IEEE 1394 (also called FireWire) interfaces, or on disks connected to shared SCSI buses.) In addition, you cannot convert cluster disks connected to shared SCSI or Fibre Channel buses to dynamic. The Cluster service supports basic disks only.)

### Caution

- If you are using a basic disk as a storage area for shadow copies and you intend to convert the disk into a dynamic disk, it is important to take the following precaution to avoid data loss. If the disk is a non-boot volume and is a different volume from where the original files reside, you must first dismount and take offline the volume containing the original files before you convert the disk containing shadow copies to a dynamic disk. You must bring the volume containing the original files back online within 20 minutes, otherwise, you will lose the data stored in the existing shadow copies. If the shadow copies are located on a boot volume, you can convert the disk to dynamic without losing shadow copies.

  You can use the **mountvol** command with the **/p** option to dismount the volume and take it offline. You can mount the volume and bring it online using the **mountvol** command or the Disk Management snap-in.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- For additional information about converting basic disks to dynamic disks, see Related Topics.
- After you convert a basic disk to a dynamic disk, you cannot change the dynamic volumes back to partitions. Instead, you must delete all dynamic volumes on the disk and then use the **Convert To Basic Disk** command. If you want to keep your data, you must first back it up or move it to another volume. For more information about the **Convert To Basic Disk** command, see Related Topics.

- Before you convert disks, close any programs that are running on those disks.

- Once converted, a dynamic disk will not contain basic volumes (primary partitions or logical drives), nor can it be accessed by MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT, or Windows XP Home Edition operating systems. Dynamic disks can only be accessed with Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems.

- When you convert a basic disk to a dynamic disk, any existing partitions or logical drives on the basic disk become simple volumes on the dynamic disk.

- Do not convert disks to dynamic that contain multiple installations of Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems. It is likely that you will no longer be able to start the second installation.

- For the conversion to succeed, any master boot record (MBR) disks to be converted must contain at least 1 MB of space for the dynamic disk database. Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems, automatically reserves this space when creating partitions or volumes on a disk, but disks with partitions or volumes created by other operating systems may not have this space available. (This space may exist even if it is not visible in Disk Management.)

**Using a command line**

1. Open Command prompt.

2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the disk you want to convert to dynamic.

4. At the **DISKPART** prompt, type:

   **select disk***n*

5. At the **DISKPART** prompt, type:

   **convert dynamic**

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **convert dynamic** | Converts a basic disk into a dynamic disk. Any existing partitions on the disk become simple volumes. |

**Caution**

- If you are using a basic disk as a storage area for shadow copies and you intend to convert the disk into a dynamic disk, it is important to take the following precaution to avoid data loss. If the disk is a non-boot volume and is a different volume

from where the original files reside, you must first dismount and take offline the volume containing the original files before you convert the disk containing shadow copies to a dynamic disk. You must bring the volume containing the original files back online within 20 minutes, otherwise, you will lose the data stored in the existing shadow copies. If the shadow copies are located on a boot volume, you can convert the disk to dynamic without losing shadow copies.

You can use the **mountvol** command with the **/p** option to dismount the volume and take it offline. You can mount the volume and bring it online using the **mountvol** command or the Disk Management snap-in.

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- For additional information about converting basic disks to dynamic disks, see Related Topics.

- After you convert a basic disk to a dynamic disk, you cannot change the dynamic volumes back to partitions. Instead, you must delete all dynamic volumes on the disk and then use the **DiskPart convert basic** command. If you want to keep your data, you must first back it up or move it to another volume. For more information about **DiskPart**, see Related Topics.

- Before you convert disks, close any programs that are running on those disks.

- Once converted, a dynamic disk will not contain partitions or logical drives, nor can it be accessed by MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT, or Windows XP Home Edition operating systems. Dynamic disks can be accessed only with Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems.

- When you convert a basic disk to a dynamic disk, any existing partitions or logical drives on the basic disk become simple volumes on the dynamic disk.

- Do not convert disks to dynamic that contain multiple installations of Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems. It is likely that you will no longer be able to start the second installation.

- For the conversion to succeed, any master boot record (MBR) disks to be converted must contain at least 1 MB of space for the dynamic disk database. Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems, automatically reserves this space when creating partitions or volumes on a disk, but disks with partitions or volumes created by other operating systems may not have this space available. (This space may exist even if it is not visible in Disk Management.)

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Converting a basic disk to dynamic
Change a dynamic disk back to a basic disk
Support for removable media
Partition styles

Working with MMC console files
Dynamic disks and volumes
DiskPart

# Change a dynamic disk back to a basic disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a dynamic disk back to a basic disk**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Back up all volumes on the disk you want to convert from dynamic to basic.

   For more information about backing up volumes, see Related Topics.

2. Open Computer Management (Local).

3. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.

4. Right-click each volume on the dynamic disk you want to convert to a basic disk, and then click **Delete Volume** for each volume on the disk.

5. When all volumes on the disk have been deleted, right-click the disk, and then click **Convert To Basic Disk**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- The disk must not have any volumes nor contain any data before you can change it back to a basic disk. If you want to keep your data, back it up or move it to another volume before you convert the disk to a basic disk.

- Once you change a dynamic disk back to a basic disk, you can create only partitions and logical drives on that disk.

**Using a command line**

1. Back up all volumes on the disk you want to convert from dynamic to basic.

   For more information about backing up volumes, see Related Topics.

2. Open Command Prompt.

3. Type:

   **diskpart**

4. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the dynamic disk you want to convert to a basic disk.

5. At the **DISKPART** prompt, type:

   **select disk***n*

6. At the **DISKPART** prompt, type:

   **detail disk***n*

7. For each volume on the disk, at the **DISKPART** prompt, type:

   **select volume=***n*

   and then type:

   **delete volume**

8. At the **DISKPART** prompt, type:

   **select disk***n*

   Specify the disk number of the disk that you want to convert to a basic disk.

9. At the **DISKPART** prompt, type:

   **convert basic**

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **detail disk** | Displays the properties of the selected disk and the volumes on that disk. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **delete volume** | Deletes the selected volume. You cannot delete the system volume, boot volume, or any volume that contains the active paging file or crash dump (memory dump). |
| **convert basic** | Converts an empty dynamic disk into a basic disk. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- The disk must not have any volumes nor contain any data before you can change it back to a basic disk. If you want to keep your data, back it up or move it to another volume before you convert the disk to a basic disk.
- Once you change a dynamic disk back to a basic disk, you can create only partitions and logical drives on that disk.
- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Partition styles
Working with MMC console files
DiskPart
Backup overview
Back up files to a file or a tape

# Change a master boot record disk into a GUID partition table disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a master boot record disk into a GUID partition table disk**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Back up or move the data on the basic master boot record (MBR) disk you want to convert into a GUID partition table (GPT) disk. If the disk does not contain any partitions or volumes, skip to step 2.

   For more information about backing up volumes, see Related Topics.

2. Open Computer Management (Local).

3. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.

4. If the disk does not contain any partitions or volumes, skip to step 5. Otherwise, right-click any volumes on the disk and then click **Delete Partition** or **Delete Volume**.

5. Right-click the MBR disk that you want to change into a GPT disk, and then click **Convert to GPT Disk**.

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- You can change a disk from MBR to GPT as long as the disk does not contain any partitions or volumes.

- You cannot use the GPT partitioning style on removable media, or on cluster disks that are connected to shared SCSI or Fibre Channel buses used by Cluster service.

- Itanium-based computers, x86-based computers running Windows Server 2003 with Service Pack 1 (SP1), and x64-based computers support both MBR and GPT disks.

**Using a command line**

1. Back up or move all data on the basic master boot record (MBR) disk you want to convert into a GUID partition table (GPT) disk. If the disk does not contain any partitions or volumes, skip to step 2.

   For more information about backing up volumes, see Related Topics.

2. Open Command Prompt.

3. Type:

   **diskpart**

   If the disk does not contain any partitions or volumes, skip to step 7.

4. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the volume that you want to delete.

5. At the **DISKPART** prompt, type:

   **select volume** *n*

6. At the **DISKPART** prompt, type:

   **delete volume**

7. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the disk that you want to convert to a GUID partition table (GPT) disk.

8. At the **DISKPART** prompt, type:

   **select disk** *n*

9. At the **DISKPART** prompt, type:

   **convert gpt**

**Important**

- You can only convert basic MBR disks without any partitions or volumes into a GPT disk. Back up your data and then delete all partitions or volumes before converting the disk.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **delete volume** | Deletes the selected volume. You cannot delete the system volume, boot volume, or any volume that contains the active paging file or crash dump (memory dump). |
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |

| | |
|---|---|
| **select disk** *n* | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **convert gpt** | Converts an empty basic disk with the master boot record (MBR) partition style into a basic disk with the GUID partition table (GPT) partition style. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can change a disk from MBR to GPT as long as the disk does not contain any partitions or volumes.

- You cannot use the GPT partitioning style on removable media, or on cluster disks that are connected to shared SCSI or Fibre Channel buses used by Cluster service.

- Itanium-based computers, x86-based computers running Windows Server 2003 with Service Pack 1 (SP1), and x64-based computers support both MBR and GPT disks.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Partition styles
Working with MMC console files
Change a GUID partition table disk into a master boot record disk
DiskPart
Backup overview
Back up files to a file or a tape

# Change a GUID partition table disk into a master boot record disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change a GUID partition table disk into a master boot record disk**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Back up or move all volumes on the basic GUID partition table (GPT) disk you want to convert into a master boot record (MBR) disk. If the disk does not contain any partitions or volumes, skip to step 2.

   For more information about backing up volumes, see Related Topics.

2. Open Computer Management (Local).

3. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.

4. If the disk does not contain any partitions or volumes, skip to step 5. Otherwise, right-click the any volumes on the disk and then click **Delete Volume**.

5. Right-click the GPT disk that you want to change into an MBR disk, and then click **Convert to MBR disk**.

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- You can change a disk from GPT to MBR as long as the disk is empty.

- Itanium-based computers, x86-based computers running Windows Server 2003 with Service Pack 1 (SP1), and x64-based computers support both MBR and GPT disks.

**Using a command line**

1. Back up or move all volumes on the basic GUID partition table (GPT) disk you want to convert into a master boot record (MBR) disk. If the disk does not contain any partitions or volumes, skip to step 2.

   For more information about backing up volumes, see Related Topics.

2. Open Command Prompt.

3. Type:

**diskpart**

If the disk does not contain any partitions or volumes, skip to step 7.

4. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the volume that you want to delete.

5. At the **DISKPART** prompt, type:

   **select volume** *n*

6. At the **DISKPART** prompt, type:

   **delete volume**

7. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the disk that you want to convert to an MBR disk.

8. At the **DISKPART** prompt, type:

   **select disk** *n*

9. At the **DISKPART** prompt, type:

   **convert mbr**

**Important**

- The disk must be empty to convert it to an MBR disk. Back up your data and then delete all partitions or volumes before converting the disk.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **delete volume** | Deletes the selected volume. You cannot delete the system volume, boot volume, or any volume that contains the active paging file or crash dump (memory dump). |
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** | Selects the specified disk, where *n* is the disk number, and gives it focus. |

| | |
|---|---|
| **convert mbr** | Converts an empty basic disk with the GUID Partition Table (GPT) partition style to a basic disk with the master boot record (MBR) partition style. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can change a disk from GPT to MBR as long as the disk is empty.

- Itanium-based computers, x86-based computers running Windows Server 2003 with Service Pack 1 (SP 1), and x64-based computers support both MBR and GPT disks.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Partition styles
Working with MMC console files
Change a master boot record disk into a GUID partition table disk
DiskPart
Backup overview
Back up files to a file or a tape

# Manage disks on a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage disks on a remote computer**

1. Open Computer Management (Local).
2. In the console tree, right-click **Computer Management (Local)**, and then click **Connect to another computer**.
3. In **Another computer**, type the name of the computer that you want to manage remotely.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

  The user of the client computer must also be a member of the Backup Operators group or Administrators group on the remote computer.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- If Windows Firewall (included in Windows XP Service Pack 2 (SP2) and in Windows Server 2003 Service Pack 1 (SP1)) is enabled on either the local or remote computer, you must add the following settings to the **Exceptions** tab in Windows Firewall on the appropriate computer to perform this procedure:

| | **Windows XP SP2** | **Windows Server 2003 SP1** |
|---|---|---|
| Client computer (local) | Dmremote.exe<br>TCP 135 | Dmremote.exe<br>Vdsldr.exe<br>TCP 135 |
| Server computer (remote) | Dmadmin.exe<br>TCP 135 | Dmadmin.exe<br>Vds.exe<br>TCP 135 |

- Mirrored volumes, RAID-5 volumes, and GUID partition table (GPT) disks are not available on all versions of Windows. The types of volumes and disks you can create depend on the version of Windows running on the remote computer, not the local computer. For example, you can use a local computer running Windows XP Professional to create mirrored and RAID-5 volumes on a remote computer that is running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.

- You can use a computer running Windows XP Professional to manage disks on a remote computer running Windows 2000 or Windows Server 2003 operating systems and vice versa.

- You can use an x86-based or x64-based computer to manage an Itanium-based computer and vice versa.

- If you use a computer running Windows XP Professional or Windows Server 2003 of operating systems to manage a

remote computer running Windows 2000, the property pages for IEEE 1394 and universal serial bus (USB) detachable disks do not show all of the information.

- If you use a computer running Windows 2000 to manage a remote x86-based computer running Windows Server 2003 with Service Pack 1 (SP 1), an x64-based computer, or an Itanium-based computer, you cannot manage GUID partition table (GPT) disks or dynamic disks in the remote computer if one of the disks in the disk group is a GPT disk.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a striped volume
Create a spanned volume
Create a simple volume
Create a RAID-5 volume
Create a mirrored volume
Create and test a mirrored system or boot volume
Create a partition or logical drive

# Remove a missing dynamic disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To remove a missing dynamic disk

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click each volume on the dynamic disk marked **Missing**, and then click **Delete Volume** to remove all volumes on the **Missing** disk.
4. Right-click the dynamic disk marked **Missing**, and then click **Remove Disk**.

### Important

- Do not delete the volumes from a **Missing** dynamic disk unless you are certain that the disk will not be powered up, reconnected, repaired or otherwise returned to the computer. Otherwise, you will not be able to access those volumes when the disk is returned to the computer.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- A dynamic disk may become **Missing** when it is corrupted, powered down, or disconnected.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the missing disk that you want to remove from the disk list. Missing disks are

numbered M0, M1, M2, etc.

4. At the **DISKPART** prompt, type:

   **select disk***n*

5. At the **DISKPART** prompt, type:

   **delete disk**

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** *n* | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **delete disk** | Deletes a missing dynamic disk from the disk list. |

**Important**

- Do not delete the volumes from a **Missing** dynamic disk unless you are certain that the disk will not be powered up, reconnected, repaired or otherwise returned to the computer. Otherwise, you will not be able to access those volumes when the disk is returned to the computer.

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- A dynamic disk may become **Missing** when it is corrupted, powered down, or disconnected.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Move disks to another computer
Disk status descriptions
Troubleshooting Disk Management
Working with MMC console files

# Manage Basic Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage basic volumes**

- View the properties of a volume
- Assign, change, or remove a drive letter
- Create a mounted drive
- Mark a partition as active (32-bit only)
- Format a basic volume
- Create a partition or logical drive
- Delete a partition or logical drive
- Extend a basic volume

# View the properties of a volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the properties of a volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the partition, logical drive, or other volume, and then click **Properties**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can only view the properties for local volumes.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   **DiskPart** displays a list of volumes in the computer and information about them, such as their drive letter, file system, size, and status.

4. At the **DISKPART** prompt, type:

   **select volume***n*

   Select the volume *n* that you want to display additional details about.

5. At the **DISKPART** prompt, type:

**detail volume**

**DiskPart** displays information about the disk on which the volume is located, such as the disk's size, status, free space, and whether the disk is dynamic or whether it uses the GUID Partition Table (GPT) partitioning scheme.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** *n* | Selects the specified volume *n*, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. You can view the numbers of all volumes on the computer by using the **list volume** command. |
| **detail volume** | Displays the disks on which the current volume resides. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- For more information about **DiskPart**, see related topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Assign, change, or remove a drive letter
Create a mounted drive
DiskPart

# Assign, change, or remove a drive letter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign, change, or remove a drive letter**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click a partition, logical drive, or volume, and then click **Change Drive Letter and Paths**.
4. Do one of the following:
    - To assign a drive letter, click **Add**, click the drive letter you want to use, and then click **OK**.
    - To modify a drive letter, click it, click **Change**, click the drive letter you want to use, and then click **OK**.
    - To remove a drive letter, click it, and then click **Remove**.

### Important

- Be careful when making drive-letter assignments because many MS-DOS and Windows programs make references to a specific drive letter. For example, the path environment variable shows specific drive letters in conjunction with program names.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- A computer can use up to 26 drive letters. Drive letters A and B are typically reserved for a floppy disk drives. If the computer does not have a floppy disk drive, you can assign drive letters A and B to removable drives, hard disk drives, or mapped network drives. Hard disk drives are typically assigned drive letters C through Z, while mapped network drives are assigned drive letters in reverse order (Z through C).
- You cannot change the drive letter of the system volume or boot volume.
- An error message may appear when you attempt to assign a letter to a volume, CD-ROM drive, or other removable media device, possibly because it is in use by a program in the system. If this happens, close the program accessing the volume or drive, and then click the **Change Drive Letter and Paths** command again.

- You can also mount a local drive at an empty folder on an NTFS volume using a drive path instead of a drive letter. For more information, see Related Topics.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the simple volume whose drive letter you want to assign, change, or remove.

4. At the **DISKPART** prompt, type:

   **select volume** *n*

   Select the volume, where *n* is the volume's number, whose drive letter you want to assign, change, or remove.

5. At the **DISKPART** prompt, type one of the following:

   - **assign letter=***L*

     Where *L* is the drive letter you want to assign or change.

   - **remove letter=***L*

     Where *L* is the drive letter you want to remove.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **assign letter=** *L* | Assigns a drive letter, *L*, to the volume with focus. If no drive letter is specified, then the next available drive letter is assigned. If the drive letter is already in use, an error is generated. |
| **remove letter=** *L* | Removes the drive letter, *L*, from the volume with focus. If no drive letter or mount point is specified, then **DiskPart** removes the first drive letter or mount point it encounters.<br><br>The **remove** command can be used to change the drive letter associated with a removable drive. You cannot remove the drive letters on system, boot, or paging volumes. In addition, you cannot remove the drive letter for an OEM partition, any GPT partition with an unrecognized GUID, or any of the special, non-data, GPT partitions such as the EFI system partition. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of

the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- A computer can use up to 26 drive letters. Drive letters A and B are typically reserved for a floppy disk drives. If the computer does not have a floppy disk drive, you can assign drive letters A and B to removable drives, hard disk drives, or mapped network drives. Hard disk drives are typically assigned drive letters C through Z, while mapped network drives are assigned drive letters in reverse order (Z through C).

- You cannot change the drive letter of the system volume or boot volume.

- An error message may appear when you attempt to assign a letter to a volume, CD-ROM drive, or other removable media device, possibly because it is in use by a program in the system. If this happens, close the program accessing the volume or drive, and then click the **Change Drive Letter and Paths** command again.

- You can also mount a local drive at an empty folder on an NTFS volume using a drive path instead of a drive letter. For more information, see Related Topics.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a mounted drive
Using NTFS mounted drives
Working with MMC console files
DiskPart

# Create a mounted drive

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a mounted drive**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the partition or volume you want to mount, and then click **Change Drive Letter and Paths**.
4. Do one of the following:
    - To mount a volume, click **Add**. Click **Mount in the following empty NTFS folder**, type the path to an empty folder on an NTFS volume, or click **Browse** to locate it.
    - To unmount a volume, click it and then click **Remove**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- When you mount a local drive at an empty folder on an NTFS volume, Disk Management assigns a drive path to the drive rather than a drive letter. Drive paths are available only on empty folders on basic or dynamic NTFS volumes.
- If you are administering a local or remote computer, you can browse NTFS folders on that computer.
- When using server clusters, make sure that you create unique mounted drives so that they do not conflict with existing local drives on any node in the cluster.
- When using server clusters, do not create mounted drives between disks on the cluster storage device (cluster disks) and local disks.
- When using server clusters, do not create mounted drives from the cluster disk that contains the quorum resource (the quorum disk). You can, however, create a mounted drive from the quorum disk to a clustered disk.
- When using server clusters, mounted drives from one cluster disk to another must be in the same cluster resource group, and must be dependent on the root disk.

For more information on server clusters, see Server clusters overview.

- When creating mounted drives, use Event Viewer to check the system log for any Cluster service errors or warnings indicating mount point failures. These errors would be listed as ClusSvc in the Source column and Physical Disk Resource in the Category column.

  For more information on how to use Event Viewer to check the system log, see Event Viewer.

  For more information on how to resolve errors with mounted drives on server clusters, see Cluster disk and drive connection problems.

**Using a command line**

1. Open Command Prompt.

2. From the NTFS drive\folder that you want to mount elsewhere, type:

   **diskpart**.

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the simple volume that you want to mount elsewhere.

4. At the **DISKPART** prompt, type:

   **select volume** *n*

   Select the simple volume, *n*, that you want to mount elsewhere.

5. At the **DISKPART** prompt, type:

   **assign** [**mount=***path*]

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **assign** | Assigns a drive letter or mount point to the volume with focus. If no drive letter or mount point is specified, then the next available drive letter is assigned. If the drive letter or mount point is already in use, an error is generated.<br><br>Using the **assign** command, you can change the drive letter associated with a removable drive.<br><br>You cannot assign drive letters to system volumes, boot volumes, or volumes that contain the paging file. In addition, you cannot assign a drive letter to an Original Equipment Manufacturer (OEM) partition, EFI system partition, or any GPT partition other than a basic data partition. |
| **mount=** *path* | Specifies an empty, existing NTFS folder where the mounted drive will reside. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of

the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Drive paths are available only on empty folders on basic or dynamic NTFS volumes.
- To modify a drive path, remove it and then create a new drive path using the new location. You cannot modify the drive path directly.
- When using server clusters, make sure that you create unique mounted drives so that they do not conflict with existing local drives on any node in the cluster.
- When using server clusters, do not create mounted drives between disks on the cluster storage device (cluster disks) and local disks.
- When using server clusters, do not create mounted drives from the cluster disk that contains the quorum resource (the quorum disk). You can, however, create a mounted drive from the quorum disk to a clustered disk.
- When using server clusters, mounted drives from one cluster disk to another must be in the same cluster resource group, and must be dependent on the root disk.

  For more information on server clusters, see Server clusters overview.
- When creating mounted drives, use Event Viewer to check the system log for any Cluster service errors or warnings indicating mount point failures. These errors would be listed as ClusSvc in the Source column and Physical Disk Resource in the Category column.

  For more information on how to use Event Viewer to check the system log, see Event Viewer.

  For more information on how to resolve errors with mounted drives on server clusters, see Cluster disk and drive connection problems.
- For more information about **DiskPart**, see Related Topics.
- You can also create a mounted drive using the **mountvol** command. For more information about **mountvol**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign, change, or remove a drive letter
Using NTFS mounted drives
Working with MMC console files
DiskPart
Mountvol
Using the Event Viewer snap-in
Event Viewer overview
View an event log

# Mark a partition as active (32-bit only)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To mark a partition as active (32-bit only)**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the primary partition you want to mark as active, and then click **Mark Partition as Active**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Changing or deleting an active partition might cause your computer not to start.
- You cannot mark a logical drive as active.
- You can only mark a partition as active on a master boot record (MBR) disk.
- A computer can only have one active partition per disk.
- In the list and graphical views, the system, active, and boot partitions are marked as **System**, **Active**, or **Boot**, respectively.
- The names commonly used for partitions containing the startup and operating system files are system and boot partitions, respectively.
- The system partition must be a primary partition that has been marked as active for startup purposes and must be located on the disk that the computer accesses when starting up the system. There can be only one active system partition on a disk at a time. You can have multiple basic disks and each can have one active partition. However the computer will only start from one specific disk. If you want to use another operating system, you must first mark its system partition as active before restarting the computer.
- The system partition can never be part of a striped volume, spanned volume, or RAID-5 volume.
- You cannot mark an existing dynamic volume as active. However, you can convert a basic disk containing the active partition to a dynamic disk. Once the disk is converted, the partition becomes a simple volume that is active. If the active

partition is not the current system or boot partition it becomes a simple volume and loses its entry in the partition table, so it can no longer be active.

#### Using a command line

1. Open Command Prompt.

2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list partition**

   Make note of the number of the partition that you want to mark as active.

4. At the **DISKPART** prompt, type:

   **select partition***n*

   Select the partition, *n*, you want to mark as active.

5. At the **DISKPART** prompt, type:

   **active**

### Important

- **DiskPart** verifies only that the partition is capable of containing an operating system's startup files. **DiskPart** does not check the contents of the partition. If you mistakenly mark a partition as "active" and it does not contain the operating system's startup files, your computer might not start.

| Value | Description |
|---|---|
| **list partition** | Displays the partitions listed in the partition table of the current disk. |
| **select partition** | Selects the specified partition and gives it focus. If no partition is specified, the **select** command lists the current partition with focus. |
| **active** | On basic disks, marks the partition with focus as active. This informs the basic input/output system (BIOS) or Extensible Firmware Interface (EFI) that the partition or volume is a valid system partition or system volume. Only partitions can be marked as active. |

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Changing or deleting an active partition might cause your computer not to start.

- You cannot mark a logical drive as active.

- You can only mark a partition as active on a master boot record (MBR) disk.
- A computer can only have one active partition per disk.
- The names commonly used for partitions containing the startup and operating system files are system and boot partitions, respectively.
- The system partition must be a primary partition that has been marked as active for startup purposes and must be located on the disk that the computer accesses when starting up the system. There can be only one active system partition on a disk at a time. You can have multiple basic disks and each can have one active partition. However the computer will only start from one specific disk. If you want to use another operating system, you must first mark its system partition as active before restarting the computer.
- The system partition can never be part of a striped volume, spanned volume, or RAID-5 volume.
- You cannot mark an existing dynamic volume as active. However, you can convert a basic disk containing the active partition to a dynamic disk. Once the disk is converted, the partition becomes a simple volume that is active. If the active partition is not the current system or boot partition it becomes a simple volume and loses its entry in the partition table, so it can no longer be active.
- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change a basic disk into a dynamic disk
Partition styles
Working with MMC console files
DiskPart

# Format a basic volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To format a basic volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the partition, logical drive, or basic volume you want to format (or reformat), and then click **Format**.
4. Select the options you want, and then click **OK**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You cannot format system, boot, OEM, or unknown partitions.
- Before you enable shadow copies on a volume, we recommend that you select a cluster size of 16 KB or larger. If you do not, the number of changes caused by defragmenting the volume might cause previous versions of files to be deleted.
- File compression is supported only on NTFS volumes with cluster sizes 4 KB and smaller.
- If you select the **Perform a quick format** check box, the files are removed from the disk but the disk is not scanned for bad sectors. Use this option only if this disk has been previously formatted and you are sure the disk is not damaged.

**Using a command line**

1. Open Command prompt.
2. Type:

   **format***volume*[**/fs:***filesystem*] [**/v:***label*] [**/q**] [**/a:***unitsize*] [**/c**]

| Value | Description |
|---|---|

| | |
|---|---|
| *Volume* | Specifies the mount point, volume name, or drive letter (followed by a colon), of the drive you want to format. If you do not specify any of the following command-line options, **format** uses the volume type to determine the default format for the disk. |
| **/fs:** *filesystem* | Specifies the file system to use FAT, FAT32, or NTFS. Floppy disks can use only the FAT file system. |
| **/v:** *label* | Specifies the volume label. If you omit the **/v** command-line option or use it without specifying a volume label, **format** prompts you for the volume label after the formatting is completed. Use the syntax **/v:** to prevent the prompt for a volume label. If you format more than one disk by using one **format** command, all of the disks will be given the same volume label. |
| **/q** | Performs a quick format. Deletes the file table and the root directory of a previously formatted volume but does not perform a sector by sector scan for bad areas. You should use the **/q** command-line option to format only previously formatted volumes that you know are in good condition. |
| **/a:** *unitsize* | Specifies the cluster size, also known as allocation unit size, to use on FAT, FAT32, or NTFS volumes. Use one of the following values for *unitsize*. If *unitsize* is not specified, it will be chosen based on volume size.<br><br>**512** - Creates 512 bytes per cluster.<br><br>**1024** - Creates 1024 bytes per cluster.<br><br>**2048** - Creates 2048 bytes per cluster.<br><br>**4096** - Creates 4096 bytes per cluster.<br><br>**8192** - Creates 8192 bytes per cluster.<br><br>**16K** - Creates 16 kilobytes per cluster.<br><br>**32K** - Creates 32 kilobytes per cluster.<br><br>**64K** - Creates 64 kilobytes per cluster. |
| **/c** | NTFS only. Files created on the new volume will be compressed by default. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You cannot format system, boot, OEM, or unknown partitions.

- Compression is supported only on NTFS volumes.

- Quick format removes files from the disk but does not scan the disk for bad sectors. Use this option only if this disk has been previously formatted and you are sure the disk is not damaged.

- Before you enable shadow copies on a volume, we recommend that you select a cluster size of 16 KB or larger. If you do not, the number of changes caused by defragmenting the volume might cause previous versions of files to be deleted.

- File compression is supported only on NTFS volumes with cluster sizes 4 KB and smaller.

- For Windows NT 4.0, Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems, the cluster size of FAT16 volumes from 2 GB through 4 GB is 64 KB, which can create compatibility issues with some applications. For example, setup programs do not compute free space properly on a volume with 64 KB clusters and

cannot run because of a perceived lack of free space. For this reason, either FAT32 or NTFS must be used on volumes larger than 2 GB. **Format** displays a warning and asks for a confirmation before formatting a volume that has 64 KB clusters using FAT16.

- For more information about **format**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Mark a partition as active (32-bit only)
Assign, change, or remove a drive letter
Create a mounted drive
Basic disks and volumes
Working with MMC console files
Format

# Create a partition or logical drive

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a partition or logical drive**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click an unallocated region of a basic disk, and then click **New Partition**, or right-click free space in an extended partition, and then click **New Logical Drive**.
4. In the New Partition wizard, click **Next**, click **Primary partition**, **Extended partition**, or **Logical drive**, and then follow the instructions on your screen.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can create primary partitions, extended partitions, and logical drives only on basic disks. You should create basic volumes instead of dynamic volumes if this computer also runs MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, or Windows XP Home Edition.
- On a master boot record (MBR) disk, you can create up to four primary partitions, or three primary partitions, one extended partition, and unlimited logical drives.
- On a GUID partition table (GPT) disk, you can create up to 128 primary partitions.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

Make note of the disk number of the disk on which you want to create a primary or extended partition.

4. At the **DISKPART** prompt, type:

   **select disk***n*

   Select the disk *n* where you want to create the primary or extended partition.

5. At the **DISKPART** prompt, type one of the following:

   **create partition primary** [**size=***n*] [**offset=***n*] [**ID=***byte* | *GUID*] [**noerr**]

   or

   **create partition extended** [**size=***n*] [**offset=***n*] [**noerr**]

   or

   **create partition logical** [**size=***n*] [**offset=***n*] [**noerr**]

6. At the **DISKPART** prompt, type:

   **assign letter**=*D*

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **create partition primary** | Creates a primary partition on the current basic disk. After you create the partition, the focus automatically shifts to the new partition. The partition does not receive a drive letter; you must use the **assign** command to assign a drive letter to the partition. |
| **create partition extended** | Creates an extended partition on the current drive. After the partition has been created, the focus automatically shifts to the new partition. Only one extended partition can be created per disk. This command fails if you attempt to create an extended partition within another extended partition. You must create an extended partition before you can create logical drives. |
| **create partition logical** | Creates a logical drive in the extended partition. After the partition has been created, the focus automatically shifts to the new logical drive. |
| **size=** *n* | The size of the partition in megabytes (MB). If no size is given, the partition continues until there is no more unallocated space in the current region. The size is cylinder snapped; the size is rounded to the closest cylinder boundary. For example, if you specify a size of 500 MB, the partition would be rounded up to 504 MB. |
| **offset=** *n* | The byte offset at which to create the partition. If no offset is given, the partition will start at the beginning of the first free space on the disk. For master boot record (MBR) disks, the offset is cylinder snapped; the offset is rounded to the closest cylinder boundary. For example, if you specify an offset that is 27 MB and the cylinder size is 8 MB, the offset is rounded to the 24 MB boundary. |

| | |
|---|---|
| **ID=** *byte* \| *GUID* | Intended for Original Equipment Manufacturer (OEM) use only.<br><br>**Caution**<br><br>• Creating partitions with this parameter might cause your computer to crash or be unable to start up. Unless you are an OEM or an IT professional experienced with GPT disks, do not create partitions on GPT disks using the **ID=***byte* \| *GUID* parameter. Instead, always use the **create partition efi** command to create EFI System partitions, the **create partition msr** command to create Microsoft Reserved partitions, and the **create partition primary** command (without the **ID=***byte* \| *GUID* parameter) to create primary partitions on GPT disks.<br><br>For MBR disks, you can specify a partition type byte for the partition. If no partition type byte is specified on an MBR disk, the **create partition primary** command creates a partition of type 0x6. Any partition type byte can be specified with the **ID=***byte* \| *GUID* parameter. **DiskPart** does not check the partition type byte for validity, nor does it perform any other checking of the ID parameter.<br><br>For GPT disks you can specify a partition type GUID for the partition you want to create:<br><br>• EFI System partition: c12a7328-f81f-11d2-ba4b-00a0c93ec93b<br><br>• Microsoft reserved partition: e3c9e316-0b5c-4db8-817d-f92df00215ae<br><br>• Basic data partition: ebd0a0a2-b9e5-4433-87c0-68b6b72699c7<br><br>• LDM Metadata partition on a dynamic disk: 5808c8aa-7e8f-42e0-85d2-e1e90434cfb3<br><br>• LDM Data partition on a dynamic disk: af9b60a0-1431-4f62-bc68-3311714a69ad<br><br>If no partition type GUID is specified, the **create partition primary** command creates a basic data partition. Any partition type can be specified with the **ID=***byte* \| *GUID* parameter. **DiskPart** does not check the partition GUID for validity, nor does it perform any other checking of the ID parameter. |
| **noerr** | For scripting only. When an error is encountered, specifies that **DiskPart** continues to process commands as if the error did not occur. Without the **noerr** parameter, an error causes **DiskPart** to exit with an error code. |
| **assign letter**=*D* | Assigns a drive letter, *D*, to the volume with focus. If no drive letter or mount point is specified, the next available drive letter is assigned. If the drive letter or mount point is already in use, an error is generated. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can create primary partitions, extended partitions, and logical drives only on basic disks. You should create basic volumes instead of dynamic volumes if this computer also runs MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, or Windows XP Home Edition.

- On a master boot record (MBR) disk, you can create up to four primary partitions, or three primary partitions, one extended partition, and unlimited logical drives.

- On a GUID partition table (GPT) disk, you can create up to 128 primary partitions, but no extended partitions or logical drives.

- You cannot create partitions on removable media using **DiskPart**.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change a basic disk into a dynamic disk
Create a mounted drive
Format a basic volume
Assign, change, or remove a drive letter
Delete a partition or logical drive
Create a partition or logical drive
Mark a partition as active (32-bit only)
Partition styles
Working with MMC console files
DiskPart

# Delete a partition or logical drive

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a partition or logical drive**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the partition, logical drive, or basic volume you want to delete, and then click **Delete Partition**.

   If the menu item reads **Delete Volume** instead of **Delete Partition**, the volume is a dynamic volume. For instructions on how to delete a dynamic volume, see Related Topics.

### Caution

- When you delete a partition, all data on that deleted partition or logical drive is lost.
- You cannot recover deleted partitions or logical drives.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You cannot delete the system volume, boot volume, or any volume that contains the active paging file or crash dump (memory dump), or OEM partition.
- You cannot delete an extended partition unless the extended partition is empty.

  All the logical drives in an extended partition be deleted before you can delete the extended partition.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the disk from which you want to delete the partition.

4. At the **DISKPART** prompt, type:

   **select disk***n*

   Select the disk *n* from which you want to delete the partition.

5. At the **DISKPART** prompt, type:

   **list partition**

   Make note of the number of the partition that you want to delete.

6. At the **DISKPART** prompt, type:

   **select partition***n*

   Select the partition *n* that you want to delete.

7. At the **DISKPART** prompt, type:

   **delete partition**

**Caution**

- Deleting a partition on a dynamic disk can delete all dynamic volumes on the disk, thus leaving the disk in a corrupt state. To delete a dynamic volume, use the **delete volume** command instead. For instructions on how to delete a dynamic volume, see Related Topics.

- When you delete a partition, all data on that deleted partition or logical drive is lost.

- You cannot recover deleted partitions or logical drives.

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **list partition** | Displays the partitions listed in the partition table of the current disk. |
| **select partition** *n* | Selects the specified partition and gives it focus. If no partition is specified, the **select** command lists the current partition with focus. |
| **delete partition** | On a basic disk, deletes the partition with focus. You cannot delete the system partition, boot partition, or any partition that contains the active paging file or crash dump (memory dump). |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You cannot delete the system volume, boot volume, or any volume that contains the active paging file or crash dump (memory dump), or OEM partition.

- You cannot delete an extended partition unless the extended partition is empty.

  All the logical drives in an extended partition be deleted before you can delete the extended partition.

- You cannot delete partitions on removable media using **DiskPart**.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Basic disks and volumes
Working with MMC console files
DiskPart
Delete a dynamic volume

# Extend a basic volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To extend a basic volume**

1. Open Command Prompt.
2. Type:

   **diskpart**
3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the basic volume you want to extend.
4. At the **DISKPART** prompt, type:

   **select volume***n*

   Selects the basic volume, *n*, you want to extend into contiguous, empty space on the same disk.
5. At the **DISKPART** prompt, type:

   **extend** [**size=***n*]

   Extends the selected volume by size=*n* megabytes (MB).

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **extend** | Extends the volume with focus into next contiguous unallocated space. For basic volumes, the unallocated space must be on the same disk as, and must follow (be of higher sector offset than) the partition with focus. A dynamic simple or spanned volume can be extended to any empty space on any dynamic disk. Using this command, you can extend an existing volume into newly created space.<br><br>If the partition was previously formatted with the NTFS file system, the file system is automatically extended to occupy the larger partition. No data loss occurs. If the partition was previously formatted with any file system format other than NTFS, the command fails with no change to the partition.<br><br>You cannot extend the current system or boot partitions. |
| **size=** *n* | The amount of space, in megabytes (MB), to add to the current partition. If you do not specify a size, the disk is extended to take up all of the next contiguous unallocated space. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the

computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To extend a basic volume, it must either not be formatted with a file system or it must be formatted with the NTFS file system.
- You can only extend a basic volume onto the same disk.
- You can only extend a basic volume if it is followed by contiguous unallocated space.
- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Basic disks and volumes
Partition styles
Manage Basic Volumes
Working with MMC console files
DiskPart

# Manage Dynamic Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage dynamic volumes**

- View the properties of a volume
- Assign, change, or remove a drive letter
- Create a mounted drive
- Format a dynamic volume
- Delete a dynamic volume
- Reactivate a dynamic volume
- Manage Simple Volumes
- Manage Spanned Volumes
- Create a striped volume
- Manage Mirrored Volumes
- Manage RAID-5 Volumes

# View the properties of a volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the properties of a volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the partition, logical drive, or other volume, and then click **Properties**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can only view the properties for local volumes.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   **DiskPart** displays a list of volumes in the computer and information about them, such as their drive letter, file system, size, and status.

4. At the **DISKPART** prompt, type:

   **select volume***n*

   Select the volume *n* that you want to display additional details about.

5. At the **DISKPART** prompt, type:

**detail volume**

**DiskPart** displays information about the disk on which the volume is located, such as the disk's size, status, free space, and whether the disk is dynamic or whether it uses the GUID Partition Table (GPT) partitioning scheme.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** *n* | Selects the specified volume *n*, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. You can view the numbers of all volumes on the computer by using the **list volume** command. |
| **detail volume** | Displays the disks on which the current volume resides. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- For more information about **DiskPart**, see related topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Assign, change, or remove a drive letter
Create a mounted drive
DiskPart

# Assign, change, or remove a drive letter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign, change, or remove a drive letter**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click a partition, logical drive, or volume, and then click **Change Drive Letter and Paths**.
4. Do one of the following:
    - To assign a drive letter, click **Add**, click the drive letter you want to use, and then click **OK**.
    - To modify a drive letter, click it, click **Change**, click the drive letter you want to use, and then click **OK**.
    - To remove a drive letter, click it, and then click **Remove**.

### Important

- Be careful when making drive-letter assignments because many MS-DOS and Windows programs make references to a specific drive letter. For example, the path environment variable shows specific drive letters in conjunction with program names.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- A computer can use up to 26 drive letters. Drive letters A and B are typically reserved for a floppy disk drives. If the computer does not have a floppy disk drive, you can assign drive letters A and B to removable drives, hard disk drives, or mapped network drives. Hard disk drives are typically assigned drive letters C through Z, while mapped network drives are assigned drive letters in reverse order (Z through C).
- You cannot change the drive letter of the system volume or boot volume.
- An error message may appear when you attempt to assign a letter to a volume, CD-ROM drive, or other removable media device, possibly because it is in use by a program in the system. If this happens, close the program accessing the volume or drive, and then click the **Change Drive Letter and Paths** command again.

- You can also mount a local drive at an empty folder on an NTFS volume using a drive path instead of a drive letter. For more information, see Related Topics.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the simple volume whose drive letter you want to assign, change, or remove.

4. At the **DISKPART** prompt, type:

   **select volume** *n*

   Select the volume, where *n* is the volume's number, whose drive letter you want to assign, change, or remove.

5. At the **DISKPART** prompt, type one of the following:

   - **assign letter=***L*

     Where *L* is the drive letter you want to assign or change.

   - **remove letter=***L*

     Where *L* is the drive letter you want to remove.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **assign letter=** *L* | Assigns a drive letter, *L*, to the volume with focus. If no drive letter is specified, then the next available drive letter is assigned. If the drive letter is already in use, an error is generated. |
| **remove letter=** *L* | Removes the drive letter, *L*, from the volume with focus. If no drive letter or mount point is specified, then **DiskPart** removes the first drive letter or mount point it encounters.<br><br>The **remove** command can be used to change the drive letter associated with a removable drive. You cannot remove the drive letters on system, boot, or paging volumes. In addition, you cannot remove the drive letter for an OEM partition, any GPT partition with an unrecognized GUID, or any of the special, non-data, GPT partitions such as the EFI system partition. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of

the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- A computer can use up to 26 drive letters. Drive letters A and B are typically reserved for a floppy disk drives. If the computer does not have a floppy disk drive, you can assign drive letters A and B to removable drives, hard disk drives, or mapped network drives. Hard disk drives are typically assigned drive letters C through Z, while mapped network drives are assigned drive letters in reverse order (Z through C).

- You cannot change the drive letter of the system volume or boot volume.

- An error message may appear when you attempt to assign a letter to a volume, CD-ROM drive, or other removable media device, possibly because it is in use by a program in the system. If this happens, close the program accessing the volume or drive, and then click the **Change Drive Letter and Paths** command again.

- You can also mount a local drive at an empty folder on an NTFS volume using a drive path instead of a drive letter. For more information, see Related Topics.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a mounted drive
Using NTFS mounted drives
Working with MMC console files
DiskPart

# Create a mounted drive

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a mounted drive**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the partition or volume you want to mount, and then click **Change Drive Letter and Paths**.
4. Do one of the following:
    - To mount a volume, click **Add**. Click **Mount in the following empty NTFS folder**, type the path to an empty folder on an NTFS volume, or click **Browse** to locate it.
    - To unmount a volume, click it and then click **Remove**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- When you mount a local drive at an empty folder on an NTFS volume, Disk Management assigns a drive path to the drive rather than a drive letter. Drive paths are available only on empty folders on basic or dynamic NTFS volumes.
- If you are administering a local or remote computer, you can browse NTFS folders on that computer.
- When using server clusters, make sure that you create unique mounted drives so that they do not conflict with existing local drives on any node in the cluster.
- When using server clusters, do not create mounted drives between disks on the cluster storage device (cluster disks) and local disks.
- When using server clusters, do not create mounted drives from the cluster disk that contains the quorum resource (the quorum disk). You can, however, create a mounted drive from the quorum disk to a clustered disk.
- When using server clusters, mounted drives from one cluster disk to another must be in the same cluster resource group, and must be dependent on the root disk.

For more information on server clusters, see Server clusters overview.

- When creating mounted drives, use Event Viewer to check the system log for any Cluster service errors or warnings indicating mount point failures. These errors would be listed as ClusSvc in the Source column and Physical Disk Resource in the Category column.

  For more information on how to use Event Viewer to check the system log, see Event Viewer.

  For more information on how to resolve errors with mounted drives on server clusters, see Cluster disk and drive connection problems.

**Using a command line**

1. Open Command Prompt.

2. From the NTFS drive\folder that you want to mount elsewhere, type:

   **diskpart**.

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the simple volume that you want to mount elsewhere.

4. At the **DISKPART** prompt, type:

   **select volume** *n*

   Select the simple volume, *n*, that you want to mount elsewhere.

5. At the **DISKPART** prompt, type:

   **assign** [**mount=***path*]

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **assign** | Assigns a drive letter or mount point to the volume with focus. If no drive letter or mount point is specified, then the next available drive letter is assigned. If the drive letter or mount point is already in use, an error is generated.<br><br>Using the **assign** command, you can change the drive letter associated with a removable drive.<br><br>You cannot assign drive letters to system volumes, boot volumes, or volumes that contain the paging file. In addition, you cannot assign a drive letter to an Original Equipment Manufacturer (OEM) partition, EFI system partition, or any GPT partition other than a basic data partition. |
| **mount=** *path* | Specifies an empty, existing NTFS folder where the mounted drive will reside. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of

the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- Drive paths are available only on empty folders on basic or dynamic NTFS volumes.
- To modify a drive path, remove it and then create a new drive path using the new location. You cannot modify the drive path directly.
- When using server clusters, make sure that you create unique mounted drives so that they do not conflict with existing local drives on any node in the cluster.
- When using server clusters, do not create mounted drives between disks on the cluster storage device (cluster disks) and local disks.
- When using server clusters, do not create mounted drives from the cluster disk that contains the quorum resource (the quorum disk). You can, however, create a mounted drive from the quorum disk to a clustered disk.
- When using server clusters, mounted drives from one cluster disk to another must be in the same cluster resource group, and must be dependent on the root disk.

  For more information on server clusters, see Server clusters overview.
- When creating mounted drives, use Event Viewer to check the system log for any Cluster service errors or warnings indicating mount point failures. These errors would be listed as ClusSvc in the Source column and Physical Disk Resource in the Category column.

  For more information on how to use Event Viewer to check the system log, see Event Viewer.

  For more information on how to resolve errors with mounted drives on server clusters, see Cluster disk and drive connection problems.
- For more information about **DiskPart**, see Related Topics.
- You can also create a mounted drive using the **mountvol** command. For more information about **mountvol**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign, change, or remove a drive letter
Using NTFS mounted drives
Working with MMC console files
DiskPart
Mountvol
Using the Event Viewer snap-in
Event Viewer overview
View an event log

# Format a dynamic volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To format a dynamic volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the volume you want to format (or reformat), and then click **Format**.
4. Select the options you want, and then click **OK**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You cannot format the system or boot volume.
- Before you enable shadow copies on a volume, we recommend that you select a cluster size of 16 KB or larger. If you do not, the number of changes caused by defragmenting the volume might cause previous versions of files to be deleted.
- File compression is supported only on NTFS volumes with cluster sizes 4 KB and smaller.
- If you select the **Perform a quick format** check box, the files are removed from the disk but the disk is not scanned for bad sectors. Use this option only if this disk has been previously formatted and you are sure the disk is not damaged.
- Only the NTFS file system is available for dynamic volumes in Disk Management. However, you can format dynamic volumes with either the FAT or FAT32 file systems using the **format** command.

**Using a command line**

1. Open Command prompt.
2. Type:

   **format***volume*[**/fs:***filesystem*] [**/v:***label*] [**/q**] [**/a:***unitsize*] [**/c**]

| Value | Description |
| --- | --- |
| *Volume* | Specifies the mount point, volume name, or drive letter of the drive you want to format. If you do not specify any of the following command-line options, **format** uses the volume type to determine the default format for the disk. |
| **/fs:** *filesystem* | Specifies the file system to use FAT, FAT32, or NTFS. Floppy disks can use only the FAT file system. |
| **/v:** *label* | Specifies the volume label. If you omit the **/v** command-line option or use it without specifying a volume label, **format** prompts you for the volume label after the formatting is completed. Use the syntax **/v:** to prevent the prompt for a volume label. If you format more than one disk by using one **format** command, all of the disks will be given the same volume label. |
| **/q** | Performs a quick format. Deletes the file table and the root directory of a previously formatted volume but does not perform a sector by sector scan for bad areas. You should use the **/q** command-line option to format only previously formatted volumes that you know are in good condition. |
| **/a:** *unitsize* | Specifies the cluster size, also known as allocation unit size, to use on FAT, FAT32, or NTFS volumes. Use one of the following values for *unitsize*. If *unitsize* is not specified, it will be chosen based on volume size.<br><br>**512** - Creates 512 bytes per cluster.<br><br>**1024** - Creates 1024 bytes per cluster.<br><br>**2048** - Creates 2048 bytes per cluster.<br><br>**4096** - Creates 4096 bytes per cluster.<br><br>**8192** - Creates 8192 bytes per cluster.<br><br>**16K** - Creates 16 kilobytes per cluster.<br><br>**32K** - Creates 32 kilobytes per cluster.<br><br>**64K** - Creates 64 kilobytes per cluster. |
| **/c** | NTFS only. Files created on the new volume will be compressed by default. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You cannot format the system or boot volume.

- Quick format removes files from the disk but does not scan the disk for bad sectors. Use this option only if this disk has been previously formatted and you are sure the disk is not damaged.

- Before you enable shadow copies on a volume, we recommend that you select a cluster size of 16 KB or larger. If you do not, the number of changes caused by defragmenting the volume might cause previous versions of files to be deleted.

- File compression is supported only on NTFS volumes with cluster sizes 4 KB and smaller.

- For Windows NT 4.0, Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems, the cluster size of FAT16 volumes from 2 GB through 4 GB is 64 KB, which can create compatibility issues with some applications. For example, setup programs do not compute free space properly on a volume with 64 KB clusters and cannot run because of a perceived lack of free space. For this reason, either FAT32 or NTFS must be used on volumes larger than 2 GB. **Format** displays a warning and asks for a confirmation before formatting a volume that has 64 KB clusters using FAT16.

- For more information about **format**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign, change, or remove a drive letter
Create a mounted drive
Dynamic disks and volumes
Working with MMC console files
Format
DiskPart

# Delete a dynamic volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To delete a dynamic volume

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the dynamic volume you want to delete, and then click **Delete Volume.**

   If the menu item reads **Delete Partition** instead of **Delete Volume**, the volume is a partition or logical drive. For instructions on how to delete a partition or logical drive, see Related Topics.

### Caution

- Use caution when deleting a volume. Deleting a simple volume, spanned volume, striped volume, mirrored volume, or RAID-5 volume deletes all the data on the volume as well the volume itself.
- You cannot recover deleted volumes.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You cannot delete the system volume, boot volume, or any volume that contains the active paging file or crash dump (memory dump).
- Dynamic simple, spanned, striped, mirrored, and RAID-5 volumes are not available on Windows XP Home Edition. Dynamic mirrored and RAID-5 volumes are not available on Windows XP Professional or Windows XP 64-bit Edition (Itanium).

#### Using a command line

1. Open Command prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the dynamic volume you want to delete.

4. At the **DISKPART** prompt, type:

   **select volume***n*

   Select the dynamic volume, where *n* is the volume's number, that you want to delete.

5. At the **DISKPART** prompt, type:

   **delete volume**

   Deletes the selected dynamic volume.

**Caution**

- Use caution when deleting a volume. Deleting a simple volume, spanned volume, striped volume, mirrored volume, or RAID-5 volume deletes all the data on the volume as well the volume itself.

- You cannot recover deleted volumes.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **delete volume** | Deletes the selected volume. You cannot delete the system volume, boot volume, or any volume that contains the active paging file or crash dump (memory dump). |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You cannot delete the system volume, boot volume, or any volume that contains the active paging file or crash dump (memory dump).

- Dynamic simple, spanned, striped, mirrored, and RAID-5 volumes are not available on Windows XP Home Edition. Dynamic mirrored and RAID-5 volumes are not available on Windows XP Professional or Windows XP 64-bit Edition (Itanium).

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dynamic disks and volumes
Working with MMC console files
DiskPart
Delete a partition or logical drive

# Reactivate a dynamic volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To reactivate a dynamic volume

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. If the disk that contains the volume does not display the **Online** status, right-click the disk and then click **Reactivate Disk**.
4. Right-click the dynamic volume with **Failed** status, and then click **Reactivate Volume**.

   The volume should return to **Healthy** status.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- A dynamic disk may become **Missing** when it is corrupted, powered down, or disconnected.
- A dynamic disk may become **Offline** if it is corrupted or intermittently unavailable. A dynamic disk may also become **Offline** if you attempt to import a foreign (dynamic) disk and the import fails. Only dynamic disks display the **Offline** status.
- Only dynamic disks can be reactivated.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

Make note of the disk number of the missing or offline dynamic disk that you want to bring online. Missing disks are numbered M0, M1, M2, etc. and offline disks are numbered 0, 1, 2, etc.

4. At the **DISKPART** prompt, type:

   **select disk***n*

5. At the **DISKPART** prompt, type:

   **online**

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** *n* | Sets the focus to the specified disk, where *n* is the disk number. |
| **online** | Brings an offline disk or volume with focus online. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- If the disk does not return to **Online** status, select the volume and then try the **online** command again.

- A dynamic disk may become **Missing** when it is corrupted, powered down, or disconnected.

- A dynamic disk may become **Offline** if it is corrupted or intermittently unavailable. A dynamic disk may also become **Offline** if you attempt to import a foreign (dynamic) disk and the import fails. Only dynamic disks display the **Offline** status.

- Only dynamic disks can be reactivated.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using spanned volumes
Volume status descriptions
Troubleshooting Disk Management
Working with MMC console files
DiskPart

# Manage Simple Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage simple volumes**

- Create a simple volume
- Extend a simple or spanned volume

# Create a simple volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a simple volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the unallocated space on the dynamic disk on which you want to create the simple volume, and then click **New Volume**.
4. In the New Volume Wizard, click **Next**, click **Simple**, and then follow the instructions on your screen.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can create simple volumes on dynamic disks only.
- You can extend a volume only if it does not have a file system or if it is formatted using the NTFS file system. You cannot extend volumes formatted using FAT or FAT32.
- You cannot extend a system volume, boot volume, striped volume, mirrored volume, or RAID-5 volume.
- You can extend simple or extended volumes that are not system or boot volumes, as long as there is available disk space. This includes volumes created as dynamic volumes or created as basic volumes and then converted to dynamic on Windows XP Professional or on Windows Server 2003 operating systems.
- If you upgraded from Windows 2000 to Windows XP Professional or Windows Server 2003 operating systems, you cannot extend a simple or spanned volume that was originally created as a basic volume and converted to a dynamic volume on Windows 2000.
- You can extend a simple volume onto additional dynamic disks to create a spanned volume. Spanned volumes cannot be mirrored.
- If you extend a spanned volume, you cannot delete any portion of it without deleting the entire spanned volume.
- Simple volumes are not fault tolerant.

- Simple volumes can only be accessed by Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems. You should create basic volumes on basic disks if the computer also runs MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, or Windows XP Home Edition because these operating systems cannot access dynamic volumes.

**Using a command line**

1. Open Command prompt.

2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the disk where you want to create a simple volume.

4. At the **DISKPART** prompt, type:

   **create volume simple** [**size=***n*] [**disk=***n*]

5. At the **DISKPART** prompt, type:

   **assign letter**=*D*

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **create volume simple** | Creates a simple volume. After you create the volume, the focus automatically shifts to the new volume. |
| **size=** *n* | The size of the volume in megabytes (MB). If no size is given, the new volume takes up the remaining free space on the disk. |
| **disk=** *n* | The dynamic disk on which to create the volume. If no disk is given, the current disk is used. |
| **assign letter**=*D* | Assigns a drive letter, *D*, to the volume with focus. If no drive letter or mount point is specified, the next available drive letter is assigned. If the drive letter or mount point is already in use, an error is generated. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can create simple volumes on dynamic disks only.

- You can extend a volume only if it does not have a file system or if it is formatted using the NTFS file system. You cannot

extend volumes formatted using FAT or FAT32.

- You cannot extend a system volume, boot volume, striped volume, mirrored volume, or RAID-5 volume.
- You can extend simple or extended volumes that are not system or boot volumes, as long as there is available disk space. This includes volumes created as dynamic volumes or created as basic volumes and then converted to dynamic on Windows XP Professional or on Windows Server 2003 operating systems.
- If you upgraded from Windows 2000 to Windows XP Professional or Windows Server 2003 operating systems, you cannot extend a simple or spanned volume that was originally created as a basic volume and converted to a dynamic volume on Windows 2000.
- You can extend a simple volume onto additional dynamic disks to create a spanned volume. Spanned volumes cannot be mirrored.
- If you extend a spanned volume, you cannot delete any portion of it without deleting the entire spanned volume.
- Simple volumes are not fault tolerant.
- Simple volumes can only be accessed by Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems. You should create basic volumes on basic disks if the computer also runs MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, or Windows XP Home Edition because these operating systems cannot access dynamic volumes.
- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using simple volumes
Extend a simple or spanned volume
Add a mirror to an existing simple volume
Create a mounted drive
Format a dynamic volume
Assign, change, or remove a drive letter
Delete a dynamic volume
Working with MMC console files
DiskPart

# Extend a simple or spanned volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To extend a simple or spanned volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the simple or spanned volume you want to extend, click **Extend Volume**, and then follow the instructions on your screen.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can extend a volume only if it does not have a file system or if it is formatted using the NTFS file system. You cannot extend volumes formatted using FAT or FAT32.
- You cannot extend a system volume, boot volume, striped volume, mirrored volume, or RAID-5 volume.
- You can extend simple or extended volumes that are not system or boot volumes, as long as there is available disk space. This includes volumes created as dynamic volumes or created as basic volumes and then converted to dynamic on Windows XP Professional or Windows Server 2003 operating systems.
- If you upgraded from Windows 2000 to Windows XP Professional or Windows Server 2003 operating systems, you cannot extend a simple or spanned volume that was originally created as a basic volume and converted to a dynamic volume on Windows 2000.
- You can extend a simple volume onto additional dynamic disks to create a spanned volume. Spanned volumes cannot be mirrored.
- If you extend a spanned volume, you cannot delete any portion of it without deleting the entire spanned volume.

**Using a command line**

1. Open Command prompt.
2. Type:

**diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the simple volume you want to extend onto another disk.

4. At the **DISKPART** prompt, type:

   **select volume***n*

   Selects the simple volume, *n*, you want to extend onto another disk.

5. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the number of the disk that you want to extend the simple volume onto.

6. At the **DISKPART** prompt, type:

   **extend** [**size=***n*] [**disk=***n*]

   Extends the selected volume onto disk *n*, and make the extension size=*n* megabytes (MB) in size.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **extend** | Extends the volume with focus into next contiguous unallocated space. For basic volumes, the unallocated space must be on the same disk as, and must follow (be of higher sector offset than) the partition with focus. A dynamic simple or spanned volume can be extended to any empty space on any dynamic disk. Using this command, you can extend an existing volume into newly created space.<br><br>If the partition was previously formatted with the NTFS file system, the file system is automatically extended to occupy the larger partition. No data loss occurs. If the partition was previously formatted with any file system format other than NTFS, the command fails with no change to the partition.<br><br>You cannot extend the current system or boot partitions. |
| **disk=** *n* | The dynamic disk on which to extend the volume. An amount of space equal to **size=***n* is allocated on the disk. If no disk is specified, the volume is extended on the current disk. |
| **size=** *n* | The amount of space, in megabytes (MB), to add to the current partition. If you do not specify a size, the disk is extended to take up all of the next contiguous unallocated space. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of

the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- You can extend a volume only if it does not have a file system or if it is formatted using the NTFS file system. You cannot extend volumes formatted using FAT or FAT32.
- You cannot extend a system volume, boot volume, striped volume, mirrored volume, or RAID-5 volume.
- You can extend simple or extended volumes that are not system or boot volumes, as long as there is available disk space. This includes volumes created as dynamic volumes or created as basic volumes and then converted to dynamic on Windows XP Professional or Windows Server 2003 operating systems.
- If you upgraded from Windows 2000 to Windows XP Professional or Windows Server 2003 operating systems, you cannot extend a simple or spanned volume that was originally created as a basic volume and converted to a dynamic volume on Windows 2000.
- You can extend a simple volume onto additional dynamic disks to create a spanned volume. Spanned volumes cannot be mirrored.
- If you extend a spanned volume, you cannot delete any portion of it without deleting the entire spanned volume.
- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a mounted drive
Format a dynamic volume
Assign, change, or remove a drive letter
Delete a dynamic volume
Using simple volumes
Using spanned volumes
Working with MMC console files
Dynamic disks and volumes
DiskPart

# Manage Spanned Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage spanned volumes**

- Create a spanned volume
- Extend a simple or spanned volume

# Create a spanned volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To create a spanned volume

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the unallocated space on one of the dynamic disks where you want to create the spanned volume, and then click **New Volume**.
4. In the New Volume Wizard, click **Next**, click **Spanned**, and then follow the instructions on your screen.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can create spanned volumes on dynamic disks only.
- You need at least two dynamic disks to create a spanned volume.
- You can extend a spanned volume onto a maximum of 32 dynamic disks.
- Spanned volumes cannot be mirrored.
- Spanned volumes are not fault tolerant.

#### Using a command line

1. Open Command prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

Make note of the disk number of the disk where you want to create a simple volume.

4. At the **DISKPART** prompt, type:

   **create volume simple** [**size=***n*] [**disk=***n*]

5. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the simple volume you want to extend onto another disk.

6. At the **DISKPART** prompt, type:

   **select volume***n*

   Selects the simple volume, *n*, you want to extend onto another disk.

7. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the number of the disk that you want to extend the simple volume onto.

8. At the **DISKPART** prompt, type:

   **extend** [**size=***n*] [**disk=***n*]

   Extends the selected volume onto disk *n*, and make the extension size=*n* megabytes (MB) in size.

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **create volume simple** | Creates a simple volume. After you create the volume, the focus automatically shifts to the new volume. |
| **size=** *n* | The size of the volume in megabytes (MB). If no size is given, the new volume takes up the remaining free space on the disk. |
| **disk=** *n* | The dynamic disk on which to create the volume. If no disk is given, the current disk is used. |
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |

| | |
|---|---|
| **extend** | Extends the volume with focus into next contiguous unallocated space. For basic volumes, the unallocated space must be on the same disk as, and must follow (be of higher sector offset than) the partition with focus. A dynamic simple or spanned volume can be extended to any empty space on any dynamic disk. Using this command, you can extend an existing volume into newly created space.<br><br>If the partition was previously formatted with the NTFS file system, the file system is automatically extended to occupy the larger partition. No data loss occurs. If the partition was previously formatted with any file system format other than NTFS, the command fails with no change to the partition.<br><br>You cannot extend the current system or boot partitions. |
| **disk=** *n* | The dynamic disk on which to extend the volume. An amount of space equal to **size=***n* is allocated on the disk. If no disk is specified, the volume is extended on the current disk. |
| **size=** *n* | The amount of space, in megabytes (MB), to add to the current partition. If you do not specify a size, the disk is extended to take up all of the next contiguous unallocated space. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can create simple volumes on dynamic disks only.

- You can extend a volume only if it does not have a file system or if it is formatted using the NTFS file system. You cannot extend volumes formatted using FAT or FAT32.

- You cannot extend a system volume, boot volume, striped volume, mirrored volume, or RAID-5 volume.

- You can extend simple or extended volumes that are not system or boot volumes, as long as there is available disk space. This includes volumes created as dynamic volumes or created as basic volumes and then converted to dynamic on Windows XP Professional or Windows Server 2003 operating systems.

- If you upgraded from Windows 2000 to Windows XP Professional or Windows Server 2003 operating systems, you cannot extend a simple or spanned volume that was originally created as a basic volume and converted to a dynamic volume on Windows 2000.

- You can extend a simple volume onto additional dynamic disks to create a spanned volume. Spanned volumes cannot be mirrored.

- If you extend a spanned volume, you cannot delete any portion of it without deleting the entire spanned volume.

- Simple volumes are not fault tolerant.

- Simple volumes can only be accessed by Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems. You should create basic volumes on basic disks if the computer also runs MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, or Windows XP Home Edition because these operating systems cannot access dynamic volumes.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Create a mounted drive
Format a dynamic volume
Assign, change, or remove a drive letter
Delete a dynamic volume
Extend a simple or spanned volume
Using spanned volumes
Working with MMC console files
Dynamic disks and volumes
DiskPart

# Extend a simple or spanned volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To extend a simple or spanned volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the simple or spanned volume you want to extend, click **Extend Volume**, and then follow the instructions on your screen.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can extend a volume only if it does not have a file system or if it is formatted using the NTFS file system. You cannot extend volumes formatted using FAT or FAT32.
- You cannot extend a system volume, boot volume, striped volume, mirrored volume, or RAID-5 volume.
- You can extend simple or extended volumes that are not system or boot volumes, as long as there is available disk space. This includes volumes created as dynamic volumes or created as basic volumes and then converted to dynamic on Windows XP Professional or Windows Server 2003 operating systems.
- If you upgraded from Windows 2000 to Windows XP Professional or Windows Server 2003 operating systems, you cannot extend a simple or spanned volume that was originally created as a basic volume and converted to a dynamic volume on Windows 2000.
- You can extend a simple volume onto additional dynamic disks to create a spanned volume. Spanned volumes cannot be mirrored.
- If you extend a spanned volume, you cannot delete any portion of it without deleting the entire spanned volume.

**Using a command line**

1. Open Command prompt.
2. Type:

**diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the simple volume you want to extend onto another disk.

4. At the **DISKPART** prompt, type:

   **select volume***n*

   Selects the simple volume, *n*, you want to extend onto another disk.

5. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the number of the disk that you want to extend the simple volume onto.

6. At the **DISKPART** prompt, type:

   **extend** [**size=***n*] [**disk=***n*]

   Extends the selected volume onto disk *n*, and make the extension size=*n* megabytes (MB) in size.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **extend** | Extends the volume with focus into next contiguous unallocated space. For basic volumes, the unallocated space must be on the same disk as, and must follow (be of higher sector offset than) the partition with focus. A dynamic simple or spanned volume can be extended to any empty space on any dynamic disk. Using this command, you can extend an existing volume into newly created space.<br><br>If the partition was previously formatted with the NTFS file system, the file system is automatically extended to occupy the larger partition. No data loss occurs. If the partition was previously formatted with any file system format other than NTFS, the command fails with no change to the partition.<br><br>You cannot extend the current system or boot partitions. |
| **disk=** *n* | The dynamic disk on which to extend the volume. An amount of space equal to **size=***n* is allocated on the disk. If no disk is specified, the volume is extended on the current disk. |
| **size=** *n* | The amount of space, in megabytes (MB), to add to the current partition. If you do not specify a size, the disk is extended to take up all of the next contiguous unallocated space. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of

the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- You can extend a volume only if it does not have a file system or if it is formatted using the NTFS file system. You cannot extend volumes formatted using FAT or FAT32.
- You cannot extend a system volume, boot volume, striped volume, mirrored volume, or RAID-5 volume.
- You can extend simple or extended volumes that are not system or boot volumes, as long as there is available disk space. This includes volumes created as dynamic volumes or created as basic volumes and then converted to dynamic on Windows XP Professional or Windows Server 2003 operating systems.
- If you upgraded from Windows 2000 to Windows XP Professional or Windows Server 2003 operating systems, you cannot extend a simple or spanned volume that was originally created as a basic volume and converted to a dynamic volume on Windows 2000.
- You can extend a simple volume onto additional dynamic disks to create a spanned volume. Spanned volumes cannot be mirrored.
- If you extend a spanned volume, you cannot delete any portion of it without deleting the entire spanned volume.
- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a mounted drive
Format a dynamic volume
Assign, change, or remove a drive letter
Delete a dynamic volume
Using simple volumes
Using spanned volumes
Working with MMC console files
Dynamic disks and volumes
DiskPart

# Create a striped volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a striped volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click unallocated space on one of the dynamic disks on which you want to create the striped volume, and then click **New Volume**.
4. In the New Volume Wizard, click **Next**, click **Striped**, and then follow the instructions on your screen.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You need at least two dynamic disks to create a striped volume. You can create a striped volume onto a maximum of 32 disks.
- Striped volumes are not fault tolerant and cannot be extended or mirrored.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the disk on which you want to create a striped volume.

4. At the **DISKPART** prompt, type:

**create volume stripe** [**size=***n*] [**disk=***n*[,[*n*,...]]

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **create volume stripe** | Creates a striped volume on the specified disks. After you create the volume, the focus automatically shifts to the new volume. |
| **size=** *n* | The amount of disk space, in megabytes (MB), that the volume will occupy on each disk. If no size is given, the new volume takes up the remaining free space on the smallest disk and an equal amount of space on each subsequent disk. |
| **disk=** *n* | The dynamic disks on which to create the volume. An amount of space equal to **size=***n* is allocated on each disk. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You need at least two dynamic disks to create a striped volume. You can create a striped volume onto a maximum of 32 disks.

- Striped volumes are not fault tolerant and cannot be extended or mirrored.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a mounted drive
Format a dynamic volume
Assign, change, or remove a drive letter
Delete a dynamic volume
Using striped volumes
Working with MMC console files
DiskPart

# Manage Mirrored Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage mirrored volumes**

- Create a mirrored volume
- Create and test a mirrored system or boot volume
- Add a mirror to an existing simple volume
- Break a mirrored volume into two volumes
- Remove a mirror from a mirrored volume
- Reconnect the disk and repair the mirrored volume
- Reactivate a mirrored volume
- Replace a failed mirror with a new mirror on another disk

# Create a mirrored volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a mirrored volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the unallocated space on one of the dynamic disks on which you want to create the mirrored volume, and then click **New Volume**.
4. In the New Volume Wizard, click **Next**, click **Mirrored**, and then follow the instructions on your screen.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can create mirrored volumes only on computers running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.
- You need two dynamic disks to create a mirrored volume.
- You can mirror an existing simple volume.
- Mirrored volumes are fault tolerant and use RAID-1, which provides redundancy by creating two identical copies of a volume.
- Mirrored volumes cannot be extended.
- Both copies (mirrors) of the mirrored volume share the same drive letter.
- You can create mirrored, striped, spanned and RAID-5 volumes using a combination of MBR and GPT disks, but the MBR cylinder-alignment restriction may cause some difficulties in creating mirrored volumes. Always mirror the MBR disk to the GPT disk to avoid cylinder-alignment difficulties.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk numbers for the dynamic disks that will make up the mirrored volume.

4. At the **DISKPART** prompt, type:

   **select disk** *n*

5. At the **DISKPART** prompt, type:

   **create volume simple** [**size=***n*] [**disk=***n*]

6. At the **DISKPART** prompt, type:

   **add disk** *n*

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **create volume simple** | Creates a simple volume. After you create the volume, the focus automatically shifts to the new volume. |
| **size=** *n* | The size of the volume in megabytes (MB). If no size is given, the new volume takes up the remaining free space on the disk. |
| **disk=** *n* | The dynamic disk on which to create the volume. If no disk is given, the current disk is used. |
| **add disk=** *n* | Mirrors the simple volume with focus to the specified disk, where *n* is the disk number. The specified disk must have unallocated space at least as large as the size of the simple volume you want to mirror. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can create mirrored volumes only on computers running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.

- You need at least two dynamic disks to create a mirrored volume.

- Mirrored volumes are fault tolerant and use RAID-1, which provides redundancy by creating two identical copies of a volume.

- Mirrored volumes cannot be extended.

- Both copies (mirrors) of the mirrored volume share the same drive letter.

- You can create mirrored, striped, spanned and RAID-5 volumes using a combination of MBR and GPT disks, but the MBR cylinder alignment-restriction may cause some difficulties in creating mirrored volumes. Always mirror the MBR disk to the GPT disk to avoid cylinder-alignment difficulties.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using mirrored volumes
Create a mounted drive
Format a dynamic volume
Assign, change, or remove a drive letter
Delete a dynamic volume
Working with MMC console files
DiskPart

# Create and test a mirrored system or boot volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create and test a mirrored system or boot volume (32-bit only)**
To ensure that your x86-based computer can load Windows if one of the disks or controllers fails, you can mirror the system volume. This makes an exact copy of the volume that contains the hardware-specific files needed to load Windows. You can also mirror the boot volume. This makes an exact copy of the volume that contains the Windows operating system.

- To create a mirrored system or boot volume
- To test a mirrored system or boot volume

**To create a mirrored system or boot volume**

1. Print this procedure or write down the steps before attempting to create a mirrored system or boot volume.
2. Read and understand all of the information in the **Important** and **Notes** sections.
3. Open Computer Management (Local).
4. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
5. If the system or boot volume already resides on a dynamic disk, skip to step 6. Otherwise, you must convert the disk to a dynamic disk before you can mirror it.

   **How?**

   a. Right-click the disk that contains the system or boot volume and click **Convert to Dynamic Disk**.

      **Caution**

      Do not convert basic disks to dynamic disks if they contain multiple operating systems (for example, the disk is set up to dual boot with another operating system). After the disk is converted to dynamic you can start the operating system that you used to convert the disk, but you will not be able to start the other operating system(s) on the disk.

   b. In the **Convert to Dynamic Disk** dialog box, select the basic disk you want to convert to dynamic and then click **OK**.

   c. In the **Disks to Convert** dialog box, click **Convert**.

   d. In the **Disk Management** dialog, click **Yes**.

   e. In the **Convert Disk to Dynamic** dialog box, click **Yes**.

   f. In the **Confirm** dialog box, click **OK** to restart your computer.

   g. After your computer restarts, log on. On the **System Settings Change** dialog box, click **Yes** to restart your computer again.

   h. After your computer restarts, log on.

   i. Open Computer Management (Local).

   j. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.

Continue to step 6.

6. Right-click the system or boot volume and click **Add Mirror**.

7. In the **Add Mirror** dialog box, select an unallocated dynamic disk with enough free space to mirror the system or boot volume and click **Add Mirror**.

   Disk Management resynchronizes both disks. When resynchronization is complete, the system or boot volume is a fault-tolerant mirrored volume.

**Important**

- Because a problem might go undetected until you try to restart the computer from the second disk, you must create a Windows startup floppy disk that you can use to start Windows on the remaining mirror. In addition, you should recreate your startup floppy disk after you add or remove volumes on the disk that contains the boot volume, after you convert the disk to dynamic, or after you install a Windows service pack.

- Consider creating an Automated System Recovery (ASR) set that you can use to restore the computer in case both halves of the mirrored system or boot volume become damaged. For more information about how to create an Automated System Recovery set, see Related Topics.

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- For additional information about converting basic disks to dynamic disks, see Related Topics.

- After you convert a basic disk to a dynamic disk, you cannot change the dynamic volumes back to partitions. Instead, you must delete all dynamic volumes on the disk and then use the **Convert To Basic Disk** command. If you want to keep your data, you must first back it up or move it to another volume. For more information about the **Convert To Basic Disk** command, see Related Topics.

- Before you convert disks, close any programs that are running on those disks.

- Once converted, a dynamic disk will not contain basic volumes (primary partitions or logical drives), nor can it be accessed by MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT, or Windows XP Home Edition operating systems. You can access dynamic disks only with computers running Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems.

- When you convert a basic disk to a dynamic disk, any existing partitions or logical drives on the basic disk become simple volumes on the dynamic disk.

- Do not convert basic disks to dynamic disks if they contain multiple operating systems (for example, the disk is set up to dual boot with another operating system). After the disk is converted to dynamic you will be able to start the operating system that you used to convert the disk, but you will not be able to start the other operating system(s) on the disk.

- For the conversion to succeed, any master boot record (MBR) disks to be converted must contain at least 1 MB of free space at the end of the disk for the dynamic disk database. Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems automatically reserve this space when creating partitions or volumes on a disk, but disks with partitions or volumes created by other operating systems might not have this space available. (This space might

exist even if it is not visible in Disk Management.)

- You can create mirrored volumes only on computers running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.
- You need two dynamic disks to create a mirrored volume.
- Mirrored volumes are fault tolerant and use RAID-1, which provides redundancy by creating two identical copies of a volume.
- Mirrored volumes cannot be extended.
- Both copies (mirrors) of the mirrored volume share the same drive letter and mount points.
- You can mirror the system and boot volumes to a different disk on the same controller or to a different disk on a different controller. Using one controller per disk is referred to as duplexing.
- Before you mirror the system or boot volume in an x86-based computer, you must note the following guidelines:
    - Using Advanced Technology Attachment (ATA) disks for a mirrored system volume is supported, but the recovery procedure is more complicated when the master disk on the primary integrated device electronics (IDE) channel fails. In this case, you must move the disk with the remaining mirror to the primary IDE channel and set its jumper to master position.
    - Do not mirror the system volume by using an ATA disk with a SCSI disk because startup problems can occur if one of the disks fails.
    - If you use duplexed SCSI controllers, use identical controllers from the same manufacturer.
    - To ensure that the computer can start from the remaining mirror, you must test the mirrored system volume before a failure.
- For more information about how to mirror the system or boot volume, see "Disk Management" at the Microsoft Resource Kits Web site.

**To test a mirrored system or boot volume**

1. Print this procedure or write down the steps before testing a mirrored system or boot volume.
2. Read and understand all of the information in the **Notes** section.
3. Shut down the computer, and then disconnect or turn off one of the disks to simulate the failure of a disk.
4. Restart the computer by using the remaining mirror.
5. After you verify that Windows starts correctly, shut down the computer and reconnect the disk.
6. Restart the computer. When the boot menu appears, select the mirror on the disk that remained connected.
7. Open Computer Management (Local).
8. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
9. Right-click the disk that has any of the volumes marked **Failed Redundancy**, and then click **Reactivate Disk**.
10. After the mirrors are resynchronized, shut down the computer and repeat steps 1 through 9 using the other disk.

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- In Steps 4 and 6, the Windows can start from the remaining mirror if all of the following conditions are met:

    - In SCSI systems, the remaining disk must be the first disk specified by the SCSI controller. If other disks are present and the remaining mirror is on the same SCSI controller as the failed disk, the BIOS for that controller must support specifying the boot device from any disk on the controller. If other disks are present and the remaining mirror is on a different controller, the original controller must support turning off the BIOS, and the mirror's controller must support specifying the boot device from any disk on the controller.

    - In SCSI systems, both SCSI controllers must have translation options set identically, either both enabled or both disabled.

    - In IDE systems, the disk that contains the remaining mirror must be the master disk on the primary IDE controller. You can move the remaining disk to this position if the master disk on the primary controller fails. Set the jumper to the master position, if necessary.

    - The disk remaining after a failure must contain a valid system volume that is marked active and must contain a valid Boot.ini file along with Ntldr and Ntdetect.com.

- When you mirror the boot volume, operating systems in the Windows Server 2003 family usually automatically update Boot.ini. If Windows cannot update Boot.ini, as when the BIOS for the second SCSI controller is disabled, you see a message instructing you to update Boot.ini manually. Therefore, to ensure that the computer can start from the remaining mirror, you need to verify that the information in Boot.ini is correct before a failure occurs, especially if you manually make changes to Boot.ini.

  For example, if you use the Signature syntax to specify the disk that contains the boot volume and that disk fails, the computer does not start because the system is looking for the disk with this particular signature. For this reason, you need to use Multi syntax to specify the disk that contains the boot volume. For more information about Boot.ini, Signature syntax, and Multi syntax, see "Troubleshooting Startup" at the Microsoft Resource Kits Web site.

- When you break a mirrored system or boot volume using the Disk Management snap-in, you must select the disk in the mirrored volume that was used to start the computer. Disk Management retains the volume on the selected disk, which keeps its drive letter and becomes a simple volume. The other disk in the mirrored system or boot volume receives the next available drive letter and becomes a new simple volume.

  When breaking a mirrored system or boot volume using the DiskPart command, however, you must specify the disk in the mirrored system or boot volume that was not used to start the computer. The disk used to start the computer keeps the mirrored volume's drive letter and becomes a simple volume, while the specified disk receives the next available drive letter and becomes a simple volume.

- For more information about how to test a mirrored the system or boot volume, see "Disk Management" at the Microsoft Resource Kits Web site.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Converting a basic disk to dynamic
Change a dynamic disk back to a basic disk
Support for removable media
Partition styles
Dynamic disks and volumes
Using mirrored volumes
Create a mounted drive
Format a dynamic volume
Assign, change, or remove a drive letter
Delete a dynamic volume
Working with MMC console files
DiskPart
Using Network Monitor
Backup
Checklists for Backing Up and Restoring Data
Automated System Recovery (ASR) overview
Create an Automated System Recovery set using Backup

# Create or break a mirrored boot volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create or break a mirrored boot volume (Itanium-based computers only)**
To ensure that your Itanium-based computer can load Windows if one of the disks or controllers fails, you can add a mirror to the boot volume, which makes an exact copy of the volume that contains the Windows operating system.

- To create a mirrored boot volume
- To break a mirrored boot volume

**To create a mirrored boot volume**

1. Print this procedure or write down the steps before attempting to create a mirrored system or boot volume.
2. Read and understand all of the information in the **Important** and **Notes** sections. In this procedure, disk 0 is the boot disk and disk 1 is the disk that you will mirror to the boot disk.
3. Before performing this procedure, make sure that the boot entries in NVRAM are in the following order:
   - EFI
   - OEM (optional)
   - MSR
   - Unknown OEM (platform flag is not set)
4. Open **Command Prompt** (click **Start**, click **Run**, and then type **cmd**).
5. Type:

   **diskpart**
6. At the DISKPART prompt, type:

   **list disk**

   Make note of the disk numbers for the dynamic disks that will make up the mirrored volume.
7. At the DISKPART prompt, type:

   **select disk 1**

   Disk 1 is an empty disk that you will mirror to the boot disk.
8. At the DISKPART prompt, type:

   **clean**
9. At the DISKPART prompt, type:

**convert gpt**

10. Delete the Microsoft Reserved (MSR) partition. At the DISKPART prompt, type:

    **select partition 1**

11. At the DISKPART prompt, type:

    **delete partition override**

12. Determine the size of the Extensible Firmware Interface (EFI) partition and the MSR partition on the boot disk. At the DISKPART prompt, type:

    **select disk 0**

13. List the partitions on disk 0. At the DISKPART prompt, type:

    **list partition**

14. Switch to disk 1. At the DISKPART prompt, type:

    **select disk 1**

15. Create an EFI partition of the same size as that on disk 0. At the DISKPART prompt, type:

    **create partition efi size=**

16. Create an MSR partition of the same size as that on disk 0. At the DISKPART prompt, type:

    **create partition msr size=**

17. Mount the EFI partition on disk 0. At the DISKPART prompt, type:

    **select disk 0**

18. At the DISKPART prompt, type:

    **select partition 1**

19. At the DISKPART prompt, type:

    **assign letter=m**

20. Mount the EFI partition on disk 1. At the DISKPART prompt, type:

    **select disk 1**

21. At the DISKPART prompt, type:

    **select partition 1**

22. At the DISKPART prompt, type:

    **assign letter=n**

23. At the DISKPART prompt, type:

**exit**

24. Format the EFI on disk 1. At the command prompt, type:

    **format n:/fs:FAT /q /y**

25. Copy all contents from the EFI on disk 0 to the EFI on disk 1. Use the **/h** option to copy the hidden files. At the command prompt, type:

    **xcopym:*.* n:\ /s /h**

26. Convert disk 0 and disk 1 to dynamic disks. At the command prompt, type:

    **diskpart**

27. At the DISKPART prompt, type:

    **select disk=1**

28. At the DISKPART prompt, type:

    **convert dynamic**

29. At the DISKPART prompt, type:

    **select disk=0**

30. At the DISKPART prompt, type:

    **convert dynamic**

    DiskPart displays the following message:

    **You must reboot your computer to complete this operation.**

31. At the DISKPART prompt, type:

    **exit**

32. Restart the computer. Wait for Windows to load and to prompt you to restart the computer a second time to complete the conversion of the boot disk to dynamic.

33. Add a mirror to the boot volume. At the command prompt, type:

    **diskpart**

34. At the DISKPART prompt, type:

    **select volume=c**

    Drive letter C is the boot volume, which is mirrored to disk 1 in the next step.

35. At the DISKPART prompt, type:

    **add disk=1**

Wait for the volume to synchronize.

36. At the DISKPART prompt, type:

    **exit**

37. Verify that there are 2 boot entries. At the command prompt, type:

    **bootcfg**

    The first entry is to boot to the first disk with the original EFI on disk 0. The second entry is to boot to the second disk with the original EFI on disk 0.

38. Clone the EFI partition. At the command prompt, type:

    **bootcfg /list**

    Write down the first EFI partition's globally unique identifier (GUID). You will need this in step 40.

39. Clone the EFI partition. At the command prompt, type:

    **bootcfg /list 1**

    Write down the second EFI partition's GUID. You will need this in step 40.

40. At the command prompt, type:

    **bootcfg /clone /sg** *GUID 1***/tg** *GUID 2***/d+ clone_entry**

    Use double quotes if the description (**/d** argument) has spaces.

41. Verify that the **bootcfg /clone** command created two entries in NVRAM. At the command prompt, type:

    **bootcfg**

    There should be 2 new entries, 4 total: two entries that use the EFI on disk 0, and two entries that use the EFI on disk 1.

**Important**

- Consider creating an Automated System Recovery (ASR) set that you can use to restore the computer in case both halves of the mirrored system or boot volume become damaged. For more information about how to create an Automated System Recovery set, see Related Topics.

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- For additional information about converting basic disks to dynamic disks, see Related Topics.

- After you convert a basic disk to a dynamic disk, you cannot change the dynamic volumes back to partitions. Instead, you must delete all dynamic volumes on the disk and then use the **Convert To Basic Disk** command. If you want to keep

your data, you must first back it up or move it to another volume. For more information about the **Convert To Basic Disk** command, see Related Topics.

- Before you convert disks, close any programs that are running on those disks.

- Once converted, a dynamic disk will not contain basic volumes (primary partitions or logical drives), nor can it be accessed by MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT, or Windows XP Home Edition operating systems. You can access dynamic disks only with computers running Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems.

- When you convert a basic disk to a dynamic disk, any existing partitions or logical drives on the basic disk become simple volumes on the dynamic disk.

- Do not convert basic disks to dynamic disks if they contain multiple operating systems (for example, if the disk is set up to dual boot with another operating system). After the disk is converted to dynamic, you will be able to start the operating system that you used to convert the disk, but you will not be able to start the other operating system(s) on the disk.

- For the conversion to succeed, any master boot record (MBR) disks to be converted must contain at least 1 MB of free space at the end of the disk for the dynamic disk database. Windows 2000, Windows XP Professional, and Windows Server 2003 operating systems automatically reserve this space when creating partitions or volumes on a disk, but disks with partitions or volumes created by other operating systems might not have this space available. (This space might exist even if it is not visible in Disk Management.)

- If there are any known original equipment manufacturer (OEM) partitions on the disk, they remain as basic partitions when you convert the disk to dynamic.

- You can create mirrored volumes only on computers running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.

- You need two dynamic disks to create a mirrored volume.

- Mirrored volumes are fault tolerant and use RAID-1, which provides redundancy by creating two identical copies of a volume.

- Mirrored volumes cannot be extended.

- Both copies (mirrors) of the mirrored volume share the same drive letter and mount points.

- You can mirror the system and boot volumes to a different disk on the same controller or to a different disk on a different controller. Using one controller per disk is referred to as duplexing.

- Before you mirror the system or boot volume in an x86-based computer, you must note the following guidelines:

    - Using Advanced Technology Attachment (ATA) disks for a mirrored system volume is supported, but the recovery procedure is more complicated when the master disk on the primary integrated device electronics (IDE) channel fails. In this case, you must move the disk with the remaining mirror to the primary IDE channel and set its jumper to master position.

    - Do not mirror the system volume by using an ATA disk with a SCSI disk because startup problems can occur if one of the disks fails.

    - If you use duplexed SCSI controllers, use identical controllers from the same manufacturer.

    - To ensure that the computer can start from the remaining mirror, you must test the mirrored system volume before a failure.

**To break a mirrored boot volume**

When you no longer require a mirrored boot volume or when a disk failure occurs, you can use this procedure to break the mirrored boot volume.

1. Remove the disk that is no longer on the system. At the command prompt, type:

   **diskpart**

2. At the DISKPART prompt, type:

   **list disk**

   Note the disk number of the disk that is missing. You will need this disk number in step 4.

3. At the DISKPART prompt, type:

   **select volume=c**

4. At the DISKPART prompt, type:

   **break disk=***disk number* **nokeep**

5. At the DISKPART prompt, type:

   **list volume**

6. Verify that disk 0 is removed from the mirror and the existing disk is now a simple volume.

7. At the DISKPART prompt, type:

   **exit**

8. Clean up the boot configuration entries from the deleted disk. At the command prompt, type:

   **bootcfg**

   Note the entries in the boot list that are no longer valid (for example, ID 1) These are the entries that boot from the original EFI partition, not the new EFI partition.

**See Also**

**Concepts**

Converting a basic disk to dynamic
Partition styles
Dynamic disks and volumes
Using mirrored volumes
Format a dynamic volume
Assign, change, or remove a drive letter
Delete a dynamic volume
DiskPart
Checklists for Backing Up and Restoring Data
Automated System Recovery (ASR) overview
Create an Automated System Recovery set using Backup

# Add a mirror to an existing simple volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a mirror to an existing simple volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the simple volume you want to mirror, click **Add Mirror**, and then follow the instructions on your screen.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can mirror volumes only on computers running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.
- Any existing simple volume can be mirrored onto another dynamic disk, as long as there is sufficient unallocated space on that disk. If you do not have a dynamic disk with enough unallocated space, the **Add Mirror** menu item is unavailable. (To verify you have enough space, right-click the disk, click **Properties**, and then check the size in **Unallocated Space**. This size may be slightly smaller than shown in the graphical and list views.)
- You need at least two dynamic disks to create a mirrored volume.
- Mirrored volumes are fault tolerant and use RAID-1, which provides redundancy by creating two identical copies of a volume.
- Mirrored volumes cannot be extended.
- Both copies (mirrors) of the mirrored volume share the same drive letter.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the simple volume that you want to mirror onto another disk.

4. At the **DISKPART** prompt, type:

   **select volume***n*

   Select the simple volume, *n*, you want to mirror onto another disk.

5. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the disk on which you want to mirror the volume.

6. At the **DISKPART** prompt, type:

   **add disk***n*

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **add disk** *n* | Mirrors the simple volume with focus to the specified disk. The specified disk must have unallocated space at least as large as the size of the simple volume you want to mirror. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can mirror volumes only on computers running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.

- You need at least two dynamic disks to create a mirrored volume.

- Mirrored volumes are fault tolerant and use RAID-1, which provides redundancy by creating two identical copies of a volume.

- Mirrored volumes cannot be extended.

- Both copies (mirrors) of the mirrored volume share the same drive letter.
- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using mirrored volumes
Break a mirrored volume into two volumes
Working with MMC console files
DiskPart

# Break a mirrored volume into two volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To break a mirrored volume into two volumes**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click one of the volume copies in the mirrored volume, and then click **Break Mirrored Volume**.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- When you break a mirrored volume, the two volume copies that made up the mirrored volume become two independent simple volumes. These volumes are no longer fault tolerant. One of the volumes retains the drive letter or mount point, while the other volume is assigned the next available drive letter.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the number of the mirrored volume.

4. At the **DISKPART** prompt, type:

   **select volume** *n*

   Select the mirrored volume you want to break into two separate volumes.

5. At the **DISKPART** prompt, type:

**detail volume**

Make note of the disk number of the mirrored volume you want to break into two separate volumes.

6. At the **DISKPART** prompt, type:

   **break disk** *n*

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **detail volume** | Displays the disks on which the current volume resides. |
| **break disk** *n* | Applies to dynamic disks only. Breaks a mirrored volume into two simple volumes. These volumes are no longer fault tolerant. The disk specified by *n* receives the focus so you can assign it a drive letter. If the specified disk is the current system or boot disk, the command fails. The other disk retains the drive letter and any mount points that belonged to the mirrored volume. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Remove a mirror from a mirrored volume
Working with MMC console files
DiskPart

# Remove a mirror from a mirrored volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To remove a mirror from a mirrored volume

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the mirror you want to remove, click **Remove Mirror**, and then follow the instructions on your screen.

### Important

- Once you remove a mirror from a mirrored volume, the removed mirror becomes unallocated space and the remaining mirror becomes a simple volume that is no longer fault tolerant. All of the data on the removed mirror is deleted.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can only regenerate mirrored volumes with **Failed Redundancy** status. If the mirrored volume has **Failed** status, try to return it to **Failed Redundancy** status by using the **Reactivate Disk** command on the remaining disks.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the volume number of the mirrored volume.

4. At the **DISKPART** prompt, type:

**select volume** *n*

5. At the **DISKPART** prompt, type:

   **detail volume**

   Make note of the disk number that specifies the half of the mirrored volume you want to remove and convert to free space.

6. At the **DISKPART** prompt, type:

   **break disk** *n* [**nokeep**]

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** *n* | Selects the specified volume and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. You can view the numbers of all volumes on the computer by using the **list volume** command. |
| **detail volume** | Displays the disks on which the current volume resides. |
| **break disk** *n* | Applies to dynamic disks only. Breaks a mirrored volume into two simple volumes. These volumes are no longer fault tolerant. The disk specified by *n* receives the focus so you can assign it a drive letter. If the specified disk is the current system or boot disk, the command fails. The other disk retains the drive letter and any mount points that belonged to the mirrored volume.<br><br>Using the **nokeep** parameter, the specified disk, *n*, is deleted and converted to free space, while the other half of the mirror is retained as a simple volume. Neither volume receives focus. |
| **Nokeep** | Specifies that only one of the mirrored volumes retained; the other simple volume is deleted and converted to free space. Neither volume receives the focus. |

**Important**

- Once you remove a mirror from a mirrored volume, the removed mirror becomes unallocated space and the remaining mirror becomes a simple volume that is no longer fault tolerant. All of the data on the removed mirror is deleted.

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can only regenerate mirrored volumes with **Failed Redundancy** status. If the mirrored volume has **Failed** status, try to return it to **Failed Redundancy** status by using the **DiskPart online** command on the remaining disks.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Break a mirrored volume into two volumes
Working with MMC console files
DiskPart

# Reconnect the disk and repair the mirrored volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reconnect the disk and repair the mirrored volume**

A disk containing part of a mirrored volume has been disconnected. Try to reconnect and reactivate the disk.

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Make sure that the physical disk is turned on, plugged in, and attached to the computer. If necessary, turn on or reattach the physical disk.
2. Open Computer Management (Local).
3. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
4. Right-click the Missing or Offline disk, and then click **Reactivate Disk**.

   If the disk does not return to **Online** status or if the volume does not return to **Healthy** status, see Replace a failed mirror with a new mirror on another disk.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using a command line**

1. Make sure that the physical disk is turned on, plugged in, and attached to the computer. If necessary, turn on or reattach the physical disk.
2. Open Command Prompt.
3. Type:

   **diskpart**

4. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number for the disk with **Offline** status.

5. At the **DISKPART** prompt, type:

**select disk***n*

Select the disk with **Offline** status.

6. At the **DISKPART** prompt, type:

   **online**

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **online** | Brings an offline disk or volume with focus online. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disk status descriptions
Volume status descriptions
Troubleshooting Disk Management
Working with MMC console files
DiskPart

# Reactivate a mirrored volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reactivate a mirrored volume**

If the I/O errors are transient, you can try to reactivate the disk.

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the partially failed disk, and then click **Reactivate Disk**.

   The mirrored volume's status should change to **Regenerating**, then **Healthy**.

   If the disk containing part of the mirrored volume does not return to **Healthy** status, see Replace a failed mirror with a new mirror on another disk.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number for the disk with **Offline** status.

4. At the **DISKPART** prompt, type:

   **select disk***n*

   Select the disk with **Offline** status.

5. At the **DISKPART** prompt, type:

   **online**

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **online** | Brings an offline disk or volume with focus online. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disk status descriptions
Volume status descriptions
Troubleshooting Disk Management
Working with MMC console files
DiskPart

# Replace a failed mirror with a new mirror on another disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To replace the failed mirror with a new mirror on another disk**
If the disk containing part of the mirrored volume is missing and the volume does not return to **Healthy** status, you should replace the failed mirror with a new mirror on another disk.

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the mirror on the **Missing** or **Offline** disk, and then click **Remove Mirror**. Follow the instructions on your screen.
4. Right-click the volume to be re-mirrored, and then click **Add Mirror**. Follow the instructions on your screen.

### Important

- Before installing a new disk, break the mirror and then examine the system log to see if the disk or the disk controller is failing. If the failing mirror is on a controller that is failing, installing a new disk on the failing controller will not solve the problem. For instructions describing how to open Event Viewer to examine the system log, see Related Topics.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To replace a mirror in the mirrored volume, you must have a dynamic disk with unallocated space that is at least as large as the region to repair. If you do not have a dynamic disk with enough unallocated space, the **Add Mirror** command is unavailable. (To verify you have enough space, right-click the disk, click **Properties**, and then check the size in **Unallocated space**. This size may be slightly smaller than shown in the graphical and list views.)
- If you are replacing one of the disks in a mirrored system volume, be sure to use an identical disk to the one you are replacing. Otherwise, startup problems might occur if the disks use different geometries or if the system volumes are at different offsets on the disks. Always test your configuration by attempting to start the operating system from each volume to ensure that the operating system can actually start if one of the disks fails. For more information about mirroring the system volume, see "Disk Management" at the Microsoft Windows Resource Kits Web site.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the volume number of the mirrored volume.

4. At the **DISKPART** prompt, type:

   **select volume** *n*

5. At the **DISKPART** prompt, type:

   **detail volume**

   Make note of the disk number that specifies the half of the mirrored volume you want to remove and convert to free space.

6. At the **DISKPART** prompt, type:

   **break disk** *n* **[nokeep]**

7. At the **DISKPART** prompt, type:

   **select volume** *n*

   Select the volume that contains the volume you want to mirror.

8. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the dynamic disk on which you want to mirror the volume.

9. At the **DISKPART** prompt, type:

   **add disk** *n*

**Important**

- Before installing a new disk, break the mirror and then examine the system log to see if the disk or the disk controller is failing. If the failing mirror is on a controller that is failing, installing a new disk on the failing controller will not solve the problem. For instructions describing how to open Event Viewer to examine the system log, see Related Topics.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **select volume** *n* | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |

| | |
|---|---|
| **detail volume** | Displays the disks on which the current volume resides. |
| **break disk** *n* | Applies to dynamic disks only. Breaks a mirrored volume into two simple volumes. These volumes are no longer fault tolerant. The disk specified by *n* receives the focus so you can assign it a drive letter. If the specified disk is the current system or boot disk, the command fails. The other disk retains the drive letter and any mount points that belonged to the mirrored volume.<br><br>Using the **nokeep** parameter, the specified disk, *n*, is deleted and converted to free space, while the other half of the mirror is retained as a simple volume. Neither volume receives focus. |
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **add disk** *n* | Mirrors the simple volume with focus to the specified disk. The specified disk must have unallocated space at least as large as the size of the simple volume you want to mirror. |
| **Nokeep** | Specifies that only one of the mirrored volumes retained; the other simple volume is deleted and converted to free space. Neither volume receives the focus. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To replace a mirror in the mirrored volume, you must have a dynamic disk with unallocated space that is at least as large as the region to repair.

- If you are replacing one of the disks in a mirrored system volume, be sure to use an identical disk to the one you are replacing. Otherwise, startup problems might occur if the disks use different geometries or if the system volumes are at different offsets on the disks. Always test your configuration by attempting to start the operating system from each volume to ensure that the operating system can actually start if one of the disks fails. For more information about mirroring the system volume, see "Disk Management" at the Microsoft Windows Resource Kits Web site.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using the Event Viewer snap-in
Event Viewer overview
View an event log
Using mirrored volumes
Disk status descriptions
Volume status descriptions
Troubleshooting Disk Management
Working with MMC console files
DiskPart

# Manage RAID-5 Volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage RAID-5 volumes**

- Create a RAID-5 volume
- Reconnect the disk and repair the RAID-5 volume
- Replace a disk region in the RAID-5 volume
- Reactivate a RAID-5 disk

# Create a RAID-5 volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a RAID-5 volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the unallocated space on one of the dynamic disks where you want to create the RAID-5 volume, and then click **New Volume**.
4. In the New Volume Wizard, click **Next**, click **RAID-5**, and then follow the instructions on your screen.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- You can create RAID-5 volumes only on computers running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.
- You need at least three (but no more than 32) dynamic disks to create a RAID-5 volume.
- RAID-5 volumes provide fault tolerance at a cost of only one additional disk for the volume. For example, if you use three 10-GB disks to create a RAID-5 volume, the volume will have a 20-GB capacity. The remaining 10-GB is used for parity.
- RAID-5 volumes cannot be extended or mirrored.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

Make note of the disk number of the disks where you want to create a RAID-5 volume.

4. At the **DISKPART** prompt, type:

   **create volume raid** [size=*n*] disk=*n*,*n*,*n*

   Creates a RAID-5 volume of *n* MB on the selected disks.

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **create volume raid** | Creates a RAID-5 volume using three or more specified dynamic disks. After you create the volume, the focus automatically shifts to the new volume. |
| **size=** *n* | The amount of disk space, in megabytes (MB), that the volume will occupy on each disk. If no size is given, the largest possible RAID-5 volume will be created. The disk with the smallest available contiguous free space determines the size for the RAID-5 volume and the same amount of space is allocated from each disk. The actual amount of usable disk space in the RAID-5 volume is less than the combined amount of disk space because some of the disk space is required for parity. |
| **disk=** *n* | The dynamic disks on which to create the volume. An amount of space equal to **size=***n* is allocated on each disk. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can create RAID-5 volumes only on computers running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.

- You need at least three (but no more than 32) dynamic disks to create a RAID-5 volume.

- RAID-5 volumes provide fault tolerance at a cost of only one additional disk for the volume. For example, if you use three 10-GB disks to create a RAID-5 volume, the volume will have a 20-GB capacity. The remaining 10-GB is used for parity.

- RAID-5 volumes cannot be extended or mirrored.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using RAID-5 volumes
Create a mounted drive

Format a dynamic volume
Assign, change, or remove a drive letter
Delete a dynamic volume
Working with MMC console files
DiskPart

# Reconnect the disk and repair the RAID-5 volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reconnect the disk and repair the RAID-5 volume**

A disk containing part of a RAID-5 volume has been disconnected. Try to reconnect and reactivate the disk.

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Make sure that the physical disk is turned on, plugged in, and attached to the computer. If necessary, turn on or reattach the physical disk.
2. Open Computer Management (Local).
3. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
4. Right-click the Missing or Offline disk, and then click **Reactivate Disk**.

   If the disk containing part of the RAID-5 volume cannot be reactivated and the volume does not return to the **Healthy** status, you should replace the failed disk region in the RAID-5 volume, see Replace a disk region in the RAID-5 volume

### Important

- Examine the system log to see if the disk or the disk controller is failing. If the failing disk is on a controller that is failing, installing a new disk on the failing controller will not solve the problem. For instructions describing how to open Event Viewer to examine the system log, see Related Topics.

### Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using a command line**

1. Make sure that the physical disk is turned on, plugged in, and attached to the computer. If necessary, turn on or reattach the physical disk.
2. Open Command Prompt.
3. Type:

   **diskpart**

4. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number for the disk with **Offline** status.

5. At the **DISKPART** prompt, type:

   **select disk***n*

   Select the disk with **Offline** status.

6. At the **DISKPART** prompt, type:

   **online**

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** | Selects the specified disk, where *n* is the disk number, and gives it focus. |
| **online** | Brings an offline disk or volume with focus online. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using the Event Viewer snap-in
Event Viewer overview
View an event log
Disk status descriptions
Volume status descriptions
Troubleshooting Disk Management
Working with MMC console files
DiskPart

# Replace a disk region in the RAID-5 volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To replace a disk region in the RAID-5 volume**
If the disk containing part of the RAID-5 volume cannot be reactivated and the volume does not return to the **Healthy** status, you should replace the failed disk region in the RAID-5 volume.

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the portion of the RAID-5 volume on the failed disk, click **Repair Volume**, and then follow the instructions on your screen.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To replace a disk region in the RAID-5 volume, you must have a dynamic disk with unallocated space that is at least as large as the region to repair. If you do not have a dynamic disk with enough unallocated space, the **Repair Volume** command is unavailable. (To verify that you have enough space, right-click the disk, click **Properties**, and then check the size in **Unallocated Space**. This size may be slightly smaller than shown in the graphical and list views.)
- When a member of a RAID-5 volume fails in a severe manner (such as a loss of power or a complete hard disk failure), computers running Windows Server 2003 operating systems can regenerate the data from the remaining members of the RAID-5 volume.
- If the RAID-5 failure is due to a power or cabling failure on a single device, you can regenerate the data within the failed member of the RAID-5 volume once the hardware state is restored.
- The RAID-5 volume will not display **Healthy** status in Disk Management until regeneration is complete.
- You can only regenerate RAID-5 volumes with **Failed Redundancy** status. If the RAID-5 volume has **Failed** status, try to return it to **Failed Redundancy** status by using the **Reactivate Disk** command on the remaining disks.

**Using a command line**

1. Make sure that the physical disk is turned on, plugged in, and attached to the computer. If necessary, turn on or reattach the physical disk.

2. Open Command Prompt.

3. Type:

   **diskpart**

4. At the **DISKPART** prompt, type:

   **list volume**

   Make note of the volume number of the RAID-5 volume with **Failed Rd** status.

5. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the number of the dynamic disk you want to use to repair the RAID-5 volume.

6. At the **DISKPART** prompt, type:

   **select volume** *n*

7. At the **DISKPART** prompt, type:

   **repair disk=***N*

   Using the **list volume** command, you can check the status of the RAID-5 volume while it is being regenerated. During regeneration, the volume status is **Rebuild**, and after regeneration, the volume status is **Healthy**.

| Value | Description |
|---|---|
| **list volume** | Displays a list of basic and dynamic volumes on all disks. |
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select volume** | Selects the specified volume, where *n* is the volume number, and gives it focus. If no volume is specified, the **select** command lists the current volume with focus. You can specify the volume by number, drive letter, or mount point path. On a basic disk, selecting a volume also gives the corresponding partition focus. |
| **repair disk=** *N* | Repairs the RAID-5 volume with focus by replacing the failed RAID-5 member with the specified dynamic disk, *N*. The specified dynamic disk must have free space greater than or equal to the total size of the failed RAID-5 member. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To replace a disk region in the RAID-5 volume, you must have a dynamic disk with unallocated space that is at least as large as the region to repair. If you do not have a dynamic disk with enough unallocated space, the **Repair Volume** command is unavailable. (To verify that you have enough space, right-click the disk, click **Properties**, and then check the size in **Unallocated Space**. This size may be slightly smaller than shown in the graphical and list views.)

- When a member of a RAID-5 volume fails in a severe manner (such as a loss of power or a complete hard disk failure), computers running Windows Server 2003 operating systems can regenerate the data from the remaining members of the RAID-5 volume.

- If the RAID-5 failure is due to a power or cabling failure on a single device, you can regenerate the data within the failed member of the RAID-5 volume once the hardware state is restored.

- The RAID-5 volume will not display **Healthy** status in Disk Management until regeneration is complete.

- You can only regenerate RAID-5 volumes with **Failed Redundancy** status. If the RAID-5 volume has **Failed** status, try to return it to **Failed Redundancy** status by using the **Reactivate Disk** command on the remaining disks.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Reconnect the disk and repair the RAID-5 volume
Reactivate a RAID-5 disk
Manage RAID-5 Volumes

# Reactivate a RAID-5 disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To reactivate a RAID-5 disk**

If the I/O errors are transient, you can try to reactivate the disk.

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Computer Management (Local).
2. In the console tree, click **Computer Management (Local)**, click **Storage**, and then click **Disk Management**.
3. Right-click the partially failed disk, and then click **Reactivate Disk**.

   The RAID-5 volume's status should change to **Regenerating**, then **Healthy**.

   If the disk containing part of the RAID-5 volume does not return to **Healthy** status, see Replace a disk region in the RAID-5 volume.

## Notes

- To perform this procedure on a local computer, you must be a member of the Backup Operators group or Administrators group on the local computer, or you must have been delegated the appropriate authority. To perform this procedure remotely, you must be a member of the Backup Operators group or Administrators group on the remote computer. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **diskpart**

3. At the **DISKPART** prompt, type:

   **list disk**

   Make note of the disk number of the missing or offline disk that you want to bring online. Missing disks are numbered M0, M1, M2, etc. and offline disks are numbered 0, 1, 2, etc.

4. At the **DISKPART** prompt, type:

   **select disk***n*

5. At the **DISKPART** prompt, type:

   **online**

| Value | Description |
|---|---|
| **list disk** | Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table (GPT) partition style. The disk marked with an asterisk (*) has focus. |
| **select disk** *n* | Sets the focus to the specified disk, where *n* is the disk number. |
| **online** | Brings an offline disk or volume with focus online. |

**Notes**

- To perform this procedure on a local computer, you must be a member of the Backup Operators group, Administrators group, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- A dynamic disk may become **Missing** when it is corrupted, powered down, or disconnected.

- A dynamic disk may become **Offline** if it is corrupted or intermittently unavailable. A dynamic disk may also become **Offline** if you attempt to import a foreign (dynamic) disk and the import fails. Only dynamic disks display the **Offline** status.

- Only dynamic disks can be reactivated.

- For more information about **DiskPart**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using the Event Viewer snap-in
Event Viewer overview
View an event log
Disk status descriptions
Volume status descriptions
Troubleshooting Disk Management
Working with MMC console files
DiskPart

# Disk Management Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about Disk Management.

- Disk Management overview
- Understanding Disk Management
- Using Disk Management
- Disk Management Resources

# Disk Management overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 8 Beta

**Disk Management overview**

Disk Management is a system utility for managing hard disks and the volumes or partitions that they contain. With Disk Management, you can initialize disks, create volumes, format volumes with the FAT, FAT32, or NTFS file systems, and create fault-tolerant disk systems. Disk Management enables you to perform most disk-related tasks without restarting the system or interrupting users; most configuration changes take effect immediately.

Disk Management, which replaces the Disk Administrator utility used in Microsoft® Windows® NT 4.0, offers many features, including:

- **Simplified tasks and intuitive user interface.** Disk Management is easy to use. Menus that are accessible from the right mouse button display the tasks you can perform on the selected object, and wizards guide you through creating partitions or volumes and initializing or converting disks. For more information, see The Disk Management window.

- **Basic and dynamic disk storage.** Basic disks contain basic volumes, such as primary partitions and logical drives in extended partitions. Use basic disks on portable computers or when you plan to install multiple operating systems in different partitions on the same disk. For more information, see Basic disks and volumes.

  Dynamic disks contain dynamic volumes that offer more features than basic disks, such as the ability to create fault-tolerant volumes on the Microsoft® Windows® 2000 Server family or Microsoft® Windows Server 2003 family of operating systems. You can extend dynamic volumes (except system or boot volumes), mirror dynamic volumes, and add new dynamic disks without restarting the computer. For more information, see Dynamic disks and volumes.

- **Local and remote disk management.** By using Disk Management, you can manage any remote computer running Windows 2000, Microsoft® Windows® XP Professional, or Windows Server 2003 operating systems on which you are a member of the Backup Operators group or Administrators group. For more information about managing disks and volumes remotely, see Remote disk management.

- **Mounted drives.** You can use Disk Management to connect, or mount, a local drive at any empty folder on a local NTFS-formatted volume. Mounted drives make data more accessible and give you the flexibility to manage data storage based on your work environment and system usage. Mounted drives are not subject to the 26-drive limit imposed by drive letters, so you can use mounted drives to access more than 26 drives on your computer. For more information, see Using NTFS mounted drives.

- **Support for MBR and GPT disks.** Disk Management offers support for master boot record (MBR) disks in x86-based computers running Windows Server 2003 with Service Pack 1 (SP 1), x64-based computers, and Itanium-based computers. You can create mirrored, striped, spanned and RAID-5 volumes using a combination of MBR and GPT disks, but the MBR cylinder-alignment restriction may cause some difficulties in creating mirrored volumes. Always mirror the MBR disk to the GPT disk to avoid cylinder-alignment difficulties. For more information, see Partition styles.

- **Support for storage area networks (SANs).** For better interoperability in storage area network (SAN) environments, on Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, volumes on new basic disks are not automatically mounted and assigned drive letters by default when added to the system. You must manually mount basic volumes and assign the drive letters using Disk Management or the **DiskPart** and **mountvol** commands. Dynamic disks, removable media devices (such as Zip or Jaz drives), and optical discs (such as CD-ROM or DVD-RAM) are always automatically mounted by the system when they are detected.

- **Manage disks at the command line.** Use the command-line tool **DiskPart** to perform disk-related tasks at the command line as an alternative to using Disk Management. With **DiskPart**, you can create scripts to automate tasks, such as creating volumes or converting disks to dynamic. For more information, see Managing disks and volumes from the command line.

# Understanding Disk Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Disk Management**
This section covers:

- The Disk Management window
- Partition styles
- Basic disks and volumes
- Dynamic disks and volumes
- Remote disk management

# The Disk Management window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The Disk Management window**

The Disk Management window displays disks and volumes in a list view and in a graphical view. The following figure depicts the disk and volume status columns in Disk Management's graphical view:




In the **Status** column of the list view, you can view the status of a disk or volume. The status also appears in the graphical view of each disk or volume. If a disk or volume status indicates an error condition, it is important to repair the disk or volume as quickly as possible to avoid data loss. If the status is not **Healthy** for volumes or **Online** for disks, see Troubleshooting Disk Management.

You can customize the list or graphical views by changing the type of information that is displayed in the top and bottom panes and by selecting the colors and patterns used to display volumes and disk regions. To change the way objects are displayed in the Disk Management window, click **View** and then select the options you want.

For more information about disk and volume statuses, see:

- Disk status descriptions
- Volume status descriptions

# Disk status descriptions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk status descriptions**
One of the following disk status descriptions will always appear in the graphical view of the disk and in the **Status** column of the disk in list view.

Audio CD

Foreign

Initializing

Missing

No Media

Not Initialized

Online

Online (Errors)

Offline

Unreadable

**Audio CD**
The **Audio CD** status occurs when you have an audio CD in your CD-ROM or DVD-ROM drive.

**Foreign**
The **Foreign** status occurs when you move a dynamic disk to the local computer from another computer running Windows 2000, Windows XP Professional, Windows XP 64-bit Edition (Itanium), or the Windows Server 2003 family of operating systems. **Foreign** status can also occur on computers running Windows XP Home Edition that are configured to dual-boot with another operating system that uses dynamic disks (such as Windows 2000 Professional or Windows XP Professional). Dynamic disks are not supported on Windows XP Home Edition or on portable computers. A warning icon appears on disks that display the **Foreign** status.

To access data on the disk, you must add the disk to your computer's system configuration. To add a disk to your computer's system configuration, import the foreign disk (right-click the disk and then click **Import Foreign Disks**). Any existing volumes on the foreign disk become visible and accessible when you import the disk. For instructions describing how to move and import disks, see Move disks to another computer.

However, you cannot access data on the disk if you are running Windows XP Home Edition. To use the disk on Windows XP Home Edition, you must convert it to a basic disk, which destroys all data on the disk.

**Caution**

- Do not convert a dynamic disk to a basic disk unless you are certain that you no longer need the data on that disk. Converting a dynamic disk to a basic disk destroys all data on the disk.

To convert the disk to an empty basic disk, see Change a dynamic disk back to a basic disk.

In some cases, a disk that was previously connected to the system can display the **Foreign** status. Configuration data for dynamic disks is stored on all dynamic disks, so the information about which disks are owned by the system is lost when all dynamic disks fail.

For instructions describing how to fix disks with **Foreign** status, see Troubleshooting Disk Management.

**Initializing**
The **Initializing** status is a temporary status that occurs when you convert a basic disk into a dynamic disk. When initialization is complete, the disk's status changes to **Online**.

**Missing**
The **Missing** status occurs when a dynamic disk is corrupted, turned off, or disconnected. Instead of appearing in the status column, the **Missing** status is displayed as the disk name. After you reconnect or turn on the missing disk, open Disk Management, right-click the missing disk, and then click **Reactivate Disk**. If there is more than one missing disk in the disk group, Disk Management will attempt to reactivate all of them.

For instructions describing how to fix disks with **Missing** status, see Troubleshooting Disk Management.

**No Media**
The **No media** status occurs when the CD-ROM, DVD-ROM, or removable drive is empty. The status changes to **Online** or **Audio CD** when you insert the appropriate media into the drive. Only CD-ROM, DVD-ROM, or removable disk types display the **No Media** status.

**Not Initialized**
The **Not Initialized** status occurs when a disk does not contain a valid signature. After you install a new disk, Windows XP Professional or the Windows Server 2003 family of operating systems must write a master boot record (MBR) or GUID partition table (GPT) before you can create partitions on the disk. When you first start Disk Management after installing a new disk, a wizard appears that provides a list of the new disks detected. If you cancel the wizard before the disk signature is written, the disk status remains **Not Initialized** until you right-click the disk and then click **Initialize Disk**. The disk status briefly changes to **Initializing** and then **Online** status.

For instructions describing how to fix disks with **Not Initialized** status, see Troubleshooting Disk Management.

**Online**
The **Online** status occurs when a basic or dynamic disk is accessible and has no known problems. This is the normal disk status. No user action is required.

**Online (Errors)**
The **Online (Errors)** status occurs when I/O errors are detected on a region of a dynamic disk. A warning icon appears on the dynamic disk with errors.

If the I/O errors were temporary, (e.g., due to a loose cable that is no longer loose) the disk returns to the **Online** status when you reactivate it.

For instructions describing how to fix disks with **Online (Errors)** status, see Troubleshooting Disk Management.

**Offline**
The **Offline** status occurs when a dynamic disk is not accessible. The dynamic disk may be corrupted or intermittently unavailable. An error icon appears on the offline dynamic disk.

If the disk status is **Offline** and the disk's name changes to **Missing**, the disk was recently available on the system but can no longer be located or identified. The missing disk may be corrupted, powered down, or disconnected.

**To bring a disk that is Offline and Missing back online**
Repair any disk, controller, or cable problems and make sure that the physical disk is turned on, plugged in, and attached to the computer. In Disk Management, right-click the disk and then click **Reactivate Disk** to bring the disk back online.

If the disk status remains **Offline** and the disk name remains **Missing**, and you determine that the disk has a problem that cannot be repaired, you can remove the disk from the system (using the **Remove Disk** command). However, before you can remove the disk, you must delete all volumes (or mirrors) on the disk. You can save any mirrored volumes on the disk by removing the mirror instead of the entire volume. Deleting a volume destroys the data in the volume, so you should remove a disk only if you are absolutely certain that the disk is permanently damaged and unusable.

**To bring a disk that is Offline and is still named Disk # (not Missing) back online**
In Disk Management, right-click the disk and then click **Reactivate Disk** to bring the disk back online. If the disk status remains **Offline**, check the cables and disk controller, and make sure that the physical disk is healthy. Correct any problems and try to reactivate the disk again. If the disk reactivation succeeds, any volumes on the disk should automatically return to the **Healthy** status.

For instructions describing how to fix disks with **Offline** status, see Troubleshooting Disk Management.

**Unreadable**
The **Unreadable** status occurs when either a basic or dynamic disk is not accessible. The disk may have experienced hardware failure, corruption, or I/O errors. The disk's copy of the system's disk configuration database may be corrupted. An error icon appears on disks that display the **Unreadable** status.

Disks may display the **Unreadable** status while they are spinning up or when Disk Management is rescanning all of the disks

on the system. In some cases, an unreadable disk has failed and is not recoverable. For dynamic disks, the **Unreadable** status usually results from corruption or I/O errors on part of the disk, rather than failure of the entire disk. You can rescan the disks (click **Action** and then click **Rescan Disks**) or restart the computer to see if the disk status changes.

For instructions describing how to fix disks with **Unreadable** status, see Troubleshooting Disk Management.

# Volume status descriptions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Volume status descriptions**

In Disk Management, one of the following volume status descriptions will always appear in the graphical view of the volume and in the **Status** column of the volume in list view.

- Failed
- Failed Redundancy
- Formatting
- Healthy
- Regenerating
- Resynching
- Unknown

**Failed**

The **Failed** status occurs when a basic or dynamic volume cannot be started automatically, or the disk is damaged. This status also occurs after importing a volume with **Data Incomplete** status. An error icon appears on the failed volume. Unless the disk or file system can be repaired, the **Failed** status indicates data loss.

If the volume is a basic volume, make sure that the underlying physical disk is turned on, plugged in, and attached to the computer. No other user action is possible for basic volumes.

If the volume is a dynamic volume, make sure the underlying dynamic disks are online. If the disks are not online, you should reactivate them. If this succeeds, the volume automatically restarts and returns to the **Healthy** status. For instructions describing how to bring a disk online, see Reactivate a missing or offline dynamic disk.

If the underlying dynamic disk returns to the **Online** status, but the dynamic volume does not return to the **Healthy** status, you can reactivate the volume manually. For instructions describing how to reactivate a volume, see Reactivate a dynamic volume.

If the volume is a mirrored or RAID-5 volume with old data, bringing the underlying disk online automatically reactivates the volume. If the disks that contain current data are disconnected, you should bring those disks online first (to allow the data to become synchronized). Otherwise, reactivate the mirrored or RAID-5 volume manually, and then run Chkdsk.exe.

For instructions describing how to reactivate a volume, see Reactivate a dynamic volume. To run Chkdsk.exe, click **Start**, click **Run**, type **chkdsk**, and then click **OK**.

For instructions describing how to fix volumes with **Failed** status, see Troubleshooting Disk Management.

**Failed Redundancy**

The **Failed Redundancy** status occurs when the data on a mirrored or RAID-5 volume is no longer fault tolerant because one of the underlying disks is not online. A warning icon appears on the volume with failed redundancy.

Volumes with the **Failed Redundancy** status often display an associated substatus (in parentheses). A volume may only display one substatus at a time. The substatus is displayed according to the order in which they appear in the table below. For example, if you have only one volume that serves as the boot volume, system volume, active volume, page file, and crash dump, its status is displayed as **Failed Redundancy (System)**. However, if an error occurs on a mirrored or RAID-5 volume, the **(At Risk)** substatus takes precedence.

| Failed Redundancy (substatuses) | Description |
|---|---|
| System | Indicates that the volume is the system volume. |
| Boot | Indicates that the volume is the boot volume. |
| Page File | Indicates that the volume holds the paging file. For more information about the paging file, see System Properties overview. |
| Crash Dump | Indicates that the volume holds the crash dump, also known as a memory dump. The memory dump is used to record the contents of the system memory when Windows XP Professional or Windows Server 2003 operating systems stop unexpectedly. For more information about the different types of memory dumps available, see Specify what happens when the system stops unexpectedly. |
| At Risk | Indicates that the data on a mirrored or RAID-5 volume is no longer fault tolerant because one of the disks has failed and I/O errors have been detected on the remaining dynamic disks. If an I/O error is detected on any part of a disk, a warning icon appears on all dynamic volumes on the disk. Dynamic volumes that are not fault tolerant display the **Healthy (At Risk)** status.<br><br>When the volume status is **Failed Redundancy (At Risk)**, the underlying disk's status is usually **Online (Errors)**. You should reactivate the disk to return it to **Online** status, and then the volume status should change to **Failed Redundancy**. |

You can continue to access the volume using the remaining online disks, but if another disk that contains the volume fails, you will lose the volume and its data. To avoid such loss, you should attempt to repair the volume as soon as possible:

- If an underlying disk is not online, reactivate the disk (right-click the disk and then click **Reactivate Disk**). If disk reactivation succeeds, the volume automatically repairs itself and returns to the **Healthy** status. A mirrored volume repairs itself by resynchronizing the data in its mirrors. A RAID-5 volume repairs itself by regenerating its parity and data.

- If the disk returns to the **Online** status, but the volume does not return to the **Healthy** status, you can reactivate the volume manually (right-click the volume and then click **Reactivate Volume**).

- If the disk does not return to the **Online** status and the volume does not return to the **Healthy** status, there may be something wrong with the disk. You should replace the failed mirror or RAID-5 disk region. To replace the failed mirror in a mirrored volume, right-click the failed mirror and then click **Remove Mirror**, and then right-click the other volume and click **Add Mirror** to create a new mirror on another disk. To replace the failed disk region in a RAID-5 volume, right-click the RAID-5 volume and then click **Repair Volume**.

For more information about how to fix mirrored or RAID-5 volumes with **Failed Redundancy** status, see Troubleshooting Disk Management.

**Formatting**

The **Formatting** status is a temporary status that occurs only while a volume is being formatted with a file system. As the volume is being formatted, the amount of the volume formatted is shown as a percentage. When formatting is complete, the volume status changes to **Healthy**.

**Healthy**

The **Healthy** status is the normal volume status on both basic and dynamic volumes when the volume is accessible and has no known problems. No user action is required.

Volumes with the **Healthy** status often display an associated substatus (in parentheses). A volume can display only one substatus at a time. The substatus is displayed according to the order in which they appear in the table below. For example, if you have only one volume which serves as the boot volume, system volume, active volume, page file, and crash dump, its status is displayed as **Healthy (System)**. However, if an error occurs on a dynamic volume, the **(At Risk)** substatus takes precedence.

| Healthy (substatuses) | Description |
|---|---|
| System | Indicates that the volume is the system volume. |

| | |
|---|---|
| Boot | Indicates that the volume is the boot volume. |
| Page File | Indicates that the volume holds the paging file. For more information about the paging file, see System Properties overview. |
| Active | Indicates that the volume is an active volume on a basic disk. For more information about making a volume the active volume, see Mark a partition as active (32-bit only). |
| Crash Dump | Indicates that the volume holds the crash dump, also known as a memory dump. The memory dump is used to record the contents of the system memory when Windows XP Professional or Windows Server 2003 operating systems stop unexpectedly. For more information about the different types of memory dumps available, see Specify what happens when the system stops unexpectedly. |
| Hibernation Partition | Indicates that the volume is an original equipment manufacturer (OEM) hibernation partition, which is used on some laptops to save the current state of the system during hibernation. |
| GPT protective Partition | Indicates that the volume is a GUID partition table (GPT) disk. The GPT protective partition contains the protective MBR to prevent disk utilities that are not compatible with GPT disks from accidentally destroying GPT partitions. |
| EFI System Partition | Indicates that the volume is the Extensible Firmware Interface (EFI) system partition on a GPT disk. |
| EISA Configuration | Indicates that the volume is an original equipment manufacturer (OEM) partition. |
| Unknown Partition | Indicates that the partition is not recognized.<br><br>Partitions on master boot record (MBR) or GUID partition table (GPT) disks with **Healthy (Unknown Partition)** status might be an unrecognized original equipment manufacturer (OEM) partition, or non-Windows operating system partition. You cannot format, assign drive letters or mount points to, or access data on partitions with **Healthy (Unknown Partition)** status. You can, however, delete these partitions using Disk Management or the DiskPart command. For instructions describing how to delete partitions, see Delete a partition or logical drive. |
| At Risk | Indicates that the dynamic volume is currently accessible, but I/O errors have been detected on the underlying dynamic disk. If an I/O error is detected on any part of a dynamic disk, all volumes on the disk display the **Healthy (At Risk)** status and a warning icon appears on the volume.<br><br>When the volume status is **Healthy (At Risk)**, an underlying disk's status is usually **Online (Errors)**. You should reactivate the underlying disk to return it to **Online** status, and then the volume should return to the **Healthy** status. If the **Healthy (At Risk)** status persists, the disk might be failing. Back up the data and replace the disk as soon as possible.<br><br>For instructions describing how to fix volumes with **Healthy (At Risk)** status, see Troubleshooting Disk Management. |

**Regenerating**

The **Regenerating** status occurs when a missing disk in a RAID-5 volume is reactivated, when an offline disk in a RAID-5 volume is reactivated, when a failed RAID-5 volume is reactivated, when disks in a RAID-5 volume are imported, and when data and parity are being regenerated for a RAID-5 volume. No user action is required. When regeneration is complete, the volume's status returns to **Healthy**. You can access a RAID-5 volume while data and parity regeneration is in progress.

**Resynching**

The **Resynching** status occurs when creating a mirror or restarting a computer with a mirrored volume, when an offline disk in a mirrored volume is reactivated, when disks in a mirrored volume are imported, and when a mirrored volume is being resynchronized so that both mirrors contain identical data. No user action is required. When resynchronization is complete, the mirrored volume's status returns to **Healthy**. Resynchronization might take some time, depending on the size of the mirrored volume. Although you can access a mirrored volume while resynchronization is in progress, you should avoid making configuration changes (such as breaking a mirror) during resynchronization.

Volumes with the **Resynching** status often display an associated substatus (in parentheses). The table below lists the possible substatuses displayed by volumes with **Resynching** status:

| Resynching (substatuses) | Description |
| --- | --- |
| System | Indicates that the volume is the system volume. |
| Boot | Indicates that the volume is the boot volume. |
| Page File | Indicates that the volume holds the paging file. For more information about the paging file, see System Properties overview. |
| Crash Dump | Indicates that the volume holds the crash dump, also known as a memory dump. The memory dump is used to record the contents of the system memory when Windows XP Professional or Windows Server 2003 operating systems stop unexpectedly. For more information about the different types of memory dumps available, see Specify what happens when the system stops unexpectedly. |

**Unknown**

The **Unknown** status occurs when the boot sector for the volume is corrupted (possibly due to a virus) and you can no longer access data on the volume.

For instructions describing how to fix volumes with **Unknown** status, see Troubleshooting Disk Management.

When importing disks, all volumes on those disks display the **OK** condition in the **Foreign Disk Volumes** dialog box unless there are problems with the volumes. When importing mirrored or RAID-5 volumes, these problems might be **Data Incomplete**, **Data Not Redundant**, or **Stale Data**.

**Data Incomplete**

The **Data Incomplete** status is displayed in the **Foreign Disk Volumes** dialog box, and occurs when data spans multiple disks, but not all of the disks were moved. Data on this volume will be destroyed unless you move the remaining disks that contain this volume, and then import all of the disks together. You cannot import the missing disks at a later time to restore the data.

For instructions describing how to fix volumes with **Data Incomplete** status, see Troubleshooting Disk Management.

**Data Not Redundant**

The **Data Not Redundant** status is displayed in the **Foreign Disk Volumes** dialog box when you import all but one of the disks in a mirrored or RAID-5 volume. In Disk Management, the imported half of the mirrored volume receives **Failed Redundancy** status, while the disk that holds the half of the mirror that was not imported receives **Missing** status. RAID-5 volumes receive **Failed Redundancy** status.

To prevent the **Data Not Redundant** status from occurring, connect all disks that belong to the mirrored or RAID-5 volume to the computer at the same time, and then import all of the disks together. For mirrored volumes, you can import the **Missing** disk at a later time to restore redundancy.

For instructions describing how to fix volumes with **Data Not Redundant** status, see Troubleshooting Disk Management.

**Stale Data**

The **Stale Data** status is displayed in the **Foreign Disk Volumes** dialog box, and occurs when a mirrored or RAID-5 volume has stale mirror information, stale parity information, or I/O errors.

For instructions describing how to fix mirrored or RAID-5 volumes with **Stale Data** status, see Troubleshooting Disk Management.

# Partition styles

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Partition styles**

Partition style refers to the method that Windows XP Professional and Windows Server 2003 operating systems use to organize partitions on the disk. All x86-based and x64-based computers use the partition style known as master boot record (MBR). MBR contains a partition table that describes where the partitions are located on the disk. Because MBR is the only partition style available on x86-based computers without Windows Server 2003 with Service Pack 1 (SP1), you do not need to choose this style. It is used automatically.

Itanium-based computers, x86-based computers running Windows Server 2003 with Service Pack 1 (SP 1), and x64-based computers can use either the MBR partition style or the GUID partition table (GPT) partition style. There are some differences between MBR and GPT partition styles, but most disk-related tasks remain the same. Basic disks and dynamic disks work the same way as in Windows 2000, and these storage types are available on disks that use either partition style. For more information about basic disks and volumes, see Basic disks and volumes. For more information about dynamic disks and volumes, see Dynamic disks and volumes. For more information about GPT disks, see GUID partition table.

In order to start, x86-based and x64-based computers must have an MBR disk with an active partition. You can also install GPT disks on x86-based computers with Windows Server 2003 with Service Pack 1 (SP1), and on x64-based computers, but you cannot start an operating system from them. Itanium-based computers require a GPT disk that contains an Extensible Firmware Interface (EFI) system partition and the files necessary to start the computer. You can also install MBR disks on Itanium-based systems, but you cannot start an operating system from them. For more information about EFI and EFI System partitions, see Extensible Firmware Interface.

In order to more easily differentiate between disks that use the MBR and GPT partition styles, Disk Management labels disks that use the master boot record partition style as MBR disks, while disks that use the GUID partition table partition style are labeled GPT disks.

The following table depicts storage types and partition styles in Windows XP Professional and Windows Server 2003 operating systems:

| Operating system | Basic volumes | Dynamic simple, spanned, and striped volumes | Dynamic mirrored and RAID-5 volumes | MBR boot volume | MBR data volume | GPT boot volume | GPT data volume |
|---|---|---|---|---|---|---|---|
| Windows XP Home Edition | X | | | X | X | | |
| Windows XP Professional | X | X | | X | X | | |
| Windows Server 2003, Standard Edition | X | X | X | X | X | | |
| Windows Server 2003, Enterprise Edition | X | X | X | X | X | | |
| Windows Server 2003, Datacenter Edition | X | X | X | X | X | | |
| Windows Server 2003, Standard Edition with Service Pack 1 (SP1) | X | X | X | X | X | | X |
| Windows Server 2003, Enterprise Edition with Service Pack 1 (SP1) | X | X | X | X | X | | X |
| Windows Server 2003, Datacenter Edition with Service Pack 1 (SP1) | X | X | X | X | X | | X |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| Windows XP 64-bit Edition (Itanium) | X | X | | | X | X | X |
| Windows Server 2003, Enterprise Edition for Itanium-based Systems | X | X | X | | X | X | X |
| Windows Server 2003, Datacenter Edition for Itanium-based Systems | X | X | X | | X | X | X |
| Windows XP Professional x64 Edition | X | X | | X | X | | X |
| Windows Server 2003, Standard x64 Edition | X | X | X | X | X | | X |
| Windows Server 2003, Enterprise x64 Edition | X | X | X | X | X | | X |
| Windows Server 2003, Datacenter x64 Edition | X | X | X | X | X | | X |

Disk Management requires NTFS on all dynamic disks and GUID partition table (GPT) disks. You can, however, use the **format** command to format these disks using FAT or FAT32. For more information about the **format** command, see Format.

For more information about file system compatibility and limitations, see Choosing a file system: NTFS, FAT, or FAT32.

# Basic disks and volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Basic disks and volumes**

A basic disk is a physical disk that contains primary partitions, extended partitions, or logical drives. Partitions and logical drives on basic disks are known as basic volumes. You can only create basic volumes on basic disks.

The number of partitions you can create on a basic disk depends on the disk's partition style:

- On master boot record (MBR) disks, you can create up to four primary partitions per disk, or you can create up to three primary partitions and one extended partition. Within the extended partition, you can create unlimited logical drives.

- On GUID partition table (GPT) disks, you can create up to 128 primary partitions. Because GPT disks do not limit you to four partitions, you do not need to create extended partitions or logical drives.

You can add more space to existing primary partitions and logical drives by extending them into adjacent, contiguous unallocated space on the same disk. To extend a basic volume, it must be formatted with the NTFS file system. You can extend a logical drive within contiguous free space in the extended partition that contains it. If you extend a logical drive beyond the free space available in the extended partition, the extended partition grows to contain the logical drive as long as the extended partition is followed by contiguous unallocated space. For instructions describing how to extend a basic volume, see Extend a basic volume.

Always use basic volumes, instead of dynamic volumes, on computers running MS-DOS, Windows 95, Windows 98, Windows Millinnium Edition, Windows NT 4.0, or Windows XP Home Edition that are configured to dual-boot with Windows XP Professional or Windows Server 2003 operating systems. These operating systems cannot access the data stored on dynamic volumes.

Windows XP Professional and the Windows Server 2003 operating systems do not support multidisk basic volumes created using Windows NT 4.0 or earlier, such as volume sets, mirror sets, stripe sets, or stripe sets with parity. For information about using Windows NT 4.0 multidisk basic volumes, see Using Windows NT 4.0 multidisk storage.

For more information about dynamic volumes, see Dynamic disks and volumes.

# Dynamic disks and volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dynamic disks and volumes**

Dynamic disks provide features that basic disks do not, such as the ability to create volumes that span multiple disks (spanned and striped volumes), and the ability to create fault tolerant volumes (mirrored and RAID-5 volumes). All volumes on dynamic disks are known as dynamic volumes.

There are five types of dynamic volumes: simple, spanned, striped, mirrored, and RAID-5. Mirrored and RAID-5 volumes are fault tolerant and are available only on computers running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems. You can, however, use a computer running Windows XP Professional to remotely create mirrored and RAID-5 volumes on these operating systems.

Regardless of whether the dynamic disk uses the master boot record (MBR) or GUID partition table (GPT) partition style, you can create up to 2,000 dynamic volumes, although the recommended number of dynamic volumes is 32 or less.

For information about how to manage dynamic volumes, see Manage Dynamic Volumes.

**Considerations when using dynamic disks and dynamic volumes**

When using dynamic volumes, the following considerations apply:

- **Installing Windows Server 2003 operating systems**. You can perform a fresh installation of Windows Server 2003 operating systems on a dynamic volume only if that volume was converted from a basic boot volume or basic system volume. If the dynamic volume was created from unallocated space on a dynamic disk, you cannot install Windows Server 2003 operating systems on that volume. This setup limitation occurs because Setup for Windows Server 2003 recognizes only dynamic volumes that have an entry in the partition table. You can, however, extend the volume (if it is a simple or spanned volume).

  Do not convert basic disks to dynamic disks if they contain multiple installations of Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems. After the conversion, it is unlikely that you will be able to start the computer using that operating system.

  For information about basic volumes, see Basic disks and volumes.

- **Portable computers and removable media**. Dynamic disks are not supported on portable computers, removable disks, detachable disks that use Universal Serial Bus (USB) or IEEE 1394 (also called FireWire) interfaces, or on disks connected to shared SCSI buses. If you are using a portable computer and right-click a disk in the graphical or list view in Disk Management, you will not see the option to convert the disk to dynamic.

- **Boot and system partitions**. You can convert a basic disk containing the system or boot partitions to a dynamic disk. After the disk is converted, these partitions become simple system or boot volumes (after restarting the computer). You cannot mark an existing dynamic volume as active. You can convert a basic disk containing the boot partition (which contains the operating system) to a dynamic disk. After the disk is converted, the boot partition becomes a simple boot volume (after restarting the computer).

- **Mirroring the boot and system volumes**. If you convert the disk containing the boot and system partitions to a dynamic disk, you can mirror the boot and system volumes onto another dynamic disk. Then, if the disk containing the boot and system volumes fails, you can start the computer from the disk containing the mirrors of these volumes. For more information, see Create and test a mirrored system or boot volume.

- **Converting dynamic disks to basic disks**. After you convert a basic disk into a dynamic disk, you cannot change the dynamic volumes back to partitions. Instead, you must move or back up your data, delete all dynamic volumes on the disk and then convert the disk. For more information, see Change a dynamic disk back to a basic disk.

- **Shadow copies storage area**. If you are using a basic disk as a storage area for shadow copies and you intend to convert the disk into a dynamic disk, it is important to take the following precaution to avoid data loss. If the disk is a non-boot volume and is a different volume from where the original files reside, you must first dismount and take offline

the volume containing the original files before you convert the disk containing shadow copies to a dynamic disk. You must bring the volume containing the original files back online within 20 minutes, otherwise, you will lose the data stored in the existing shadow copies. If the shadow copies are located on a boot volume, you can convert the disk to dynamic without losing shadow copies.

You can use the **mountvol** command with the **/p** option to dismount the volume and take it offline. You can mount the volume and bring it online using the **mountvol** command or the Disk Management snap-in.

# Remote disk management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Remote disk management**

Using Disk Management, you can remotely manage disks on another computer in your network. To manage disks that are connected to remote computers, you must be a member of the Backup Operators group or Administrators group on both the local and remote computers. In addition, the user account and both computers must be members of the same domain or within trusted domains.

Mirrored volumes, RAID-5 volumes, and GUID partition table (GPT) disks are not available on all versions of Windows XP or Windows Server 2003 operating systems. The types of volumes and disks you can create depend on the operating system running on the remote computer, not the local computer. For example, you can use a local computer running Windows XP Professional to create mirrored and RAID-5 volumes on a remote computer that is running Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or Windows Server 2003 operating systems.

For information on administering the disk system on another computer, see Manage disks on a remote computer.

# Using Disk Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Disk Management**
This section covers:

- Adding a disk
- Using NTFS mounted drives
- Support for removable media
- Using Basic Disks
- Using Dynamic Disks
- Managing disks and volumes from the command line

# Adding a disk

Updated: December 5, 2012

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding a disk**

**Note**

This topic does not apply to Windows Server® 2008, Windows Server® 2008 R2, or Windows® 7. For information on adding disks when using these operating systems, see the following:

- Overview of Disk Management
- Manage Disks
- Adding Hard Disk Drives

When you attach a new disk to your computer, you must initialize the disk before you can create volumes or partitions. When you first start Disk Management after installing a new disk, a wizard appears that provides a list of the new disks detected by the operating system. When you complete the wizard, the operating system initializes the disk(s) by writing a disk signature, the end of sector marker (also called signature word), and a master boot record (MBR) or GUID partition table (GPT) on the disk. If you cancel the wizard before the disk signature is written, the disk status remains **Not Initialized**.

You can use either basic storage or dynamic storage. Use basic storage, which uses basic volumes, if you want to create partitions and logical drives on the disk or if you want to use the disk with other operating systems. Later, you can convert basic storage to dynamic storage, which uses dynamic volumes. Dynamic volumes cannot be accessed by computers running MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, or Windows XP Home Edition that are configured to dual-boot with Windows XP Professional or Windows Server 2003 operating systems.

There are five types of dynamic volumes: simple, spanned, striped, mirrored, and RAID-5. You can create simple, spanned, or striped volumes on computers running Windows 2000 Professional, Windows XP Professional, or Windows XP 64-bit Edition (Itanium). You can create all five types of dynamic volumes on computers running Windows 2000 Server or Windows Server 2003 operating systems. You can, however, use a computer running Windows XP Professional to remotely create mirrored and RAID-5 volumes on these operating systems.

For information about which partition style or volume type you can use on computers running Windows XP or Windows Server 2003 operating systems, see Partition styles.

For information on how to initialize a new disk, see Initialize new disks.

To install a disk that is already initialized, see Update disk information.

To change the storage type for which a disk is initialized, see Change a basic disk into a dynamic disk or Change a dynamic disk back to a basic disk.

# Using NTFS mounted drives

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using NTFS mounted drives**

NTFS mounted drives are a convenient way to add additional volumes to a computer when you no longer have any drive letters remaining. In addition, you can add more space to a volume without having to recreate the volume on another, larger disk by mounting other disks as folders on the volume. Growing the folder structure with mounted drives gives the volume a scalable, larger capacity. You can use Disk Management to mount a local volume at any empty folder on a local NTFS volume. You can format a mounted drive with any file system supported by Windows XP Professional or Windows Server 2003 operating systems.

When you mount a local drive at an empty folder on an NTFS volume, Disk Management assigns a drive path to the drive rather than a drive letter. Mounted drives are not subject to the 26-drive limit imposed by drive letters, so you can use mounted drives to access more than 26 drives on your computer. Windows ensures that drive paths retain their association to the drive, so you can add or rearrange storage devices without the drive path failing.

For example, if you have a hard disk drive with the drive letter D and an NTFS-formatted volume with the drive letter C, you can mount the hard disk drive at an empty folder with the following path: C:\datadisk. You can then access the hard disk drive directly through the path C:\datadisk. Also, you can choose to remove the drive letter D from the hard disk drive and continue to access it through the mounted drive path.

Mounted drives make data more accessible and give you the flexibility to manage data storage based on your work environment and system usage. For example, you can:

- Make the C:\Users folder a mounted drive with NTFS disk quotas in order to track or constrain disk usage without doing the same on drive C.
- Make the C:\Temp folder a mounted drive to provide additional disk space for temporary files.
- Move the My Documents folder to another, larger drive when space is low on drive C, and mount it as C:\My Documents.

When using NTFS mounted drives with server clusters:

- Make sure that you create unique mounted drives so that they do not conflict with existing local drives on any node in the cluster.
- Do not create mounted drives between disks on the cluster storage device (cluster disks) and local disks.
- Do not create mounted drives from the cluster disk that contains the quorum resource (the quorum disk). You can, however, create a mounted drive from the quorum disk to a clustered disk.
- Mounted drives from one cluster disk to another must be in the same cluster resource group, and must be dependent on the root disk.

  For more information on server clusters, see Server clusters overview.
- Use Event Viewer to check the system log for any Cluster service errors or warnings indicating mount point failures. These errors would be listed as ClusSvc in the Source column and Physical Disk Resource in the Category column.

  For more information on how to use Event Viewer to check the system log, see Event Viewer.

  For more information on how to resolve errors with mounted drives on server clusters, see Cluster disk and drive connection problems.

# Support for removable media

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Support for removable media**

In the **Graphical View** in Disk Management, removable media devices are labeled **Removable**.

Most removable media devices, such as Zip and Jaz drives and optical discs (for example, DVD-RAM), already have a primary partition equal to the total size of the media. If the removable media has a primary partition, you cannot mark the partition as active or delete it. You cannot create extended partitions, logical drives, or dynamic volumes on removable media devices.

In addition, you cannot use the GUID partition table (GPT) partitioning style on removable media.

**Important**

- For better interoperability in storage area network (SAN) environments, on Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, volumes on new basic disks are not automatically mounted and assigned drive letters by default when added to the system. You must manually mount basic volumes and assign the drive letters using Disk Management or the **DiskPart** and **mountvol** commands. Dynamic disks, removable media devices (such as Zip or Jaz drives), and optical discs (such as CD-ROM or DVD-RAM) are always automatically mounted by the system when they are detected.

For information about the other types of storage you can use with Disk Management, see Using Disk Management.

# Using Basic Disks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using basic disks**

This section covers:

- Converting a basic disk to dynamic
- Using Windows NT 4.0 multidisk storage

# Converting a basic disk to dynamic

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Converting a basic disk to dynamic**
Setup and Disk Management ensure that disks initialized by Windows Server 2003 operating systems can be converted to dynamic disks. Dynamic disks provide features that basic disks do not, such as the ability to create multidisk volumes, such as striped, mirrored, and RAID-5 volumes.

Before you convert disks, close any programs that are running on those disks. If you convert a boot disk, or if a volume or partition is in use on the disk you attempt to convert, you must restart the computer for the conversion to succeed. The conversion may fail if you change the disk layout of a disk to be converted or if the disk has I/O errors during the conversion. After you convert a basic disk into a dynamic disk, any existing partitions on the basic disk become (dynamic) simple volumes. You cannot change the dynamic volumes back to partitions.

**Caution**

- If you are using a basic disk as a storage area for shadow copies and you intend to convert the disk into a dynamic disk, it is important to take the following precaution to avoid data loss. If the disk is a non-boot volume and is a different volume from where the original files reside, you must first dismount and take offline the volume containing the original files before you convert the disk containing shadow copies to a dynamic disk. You must bring the volume containing the original files back online within 20 minutes. Otherwise, you will lose the data stored in the existing shadow copies. If the shadow copies are located on a boot volume, you can convert the disk to dynamic without losing shadow copies.

  You can use the **mountvol** command with the **/p** option to dismount the volume and take it offline. You can mount the volume and bring it online using the **mountvol** command or the Disk Management snap-in.

Dynamic disks cannot be directly accessed by MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT, or Windows XP Home Edition, so you cannot start these operating systems on dynamic disks. However, shared folders on dynamic disks are available across a network to computers running all of these operating systems.

Access to dynamic disks is further restricted by the partition style used on the dynamic disk:

- Only x86-based computers running Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems, and Itanium-based computers running the 64-bit version of Windows Server 2003, Enterprise Edition, or the 64-bit version of Windows Server 2003, Datacenter Edition, can access dynamic MBR disks.

- Only Itanium-based computers, x86-based computers running Windows Server 2003 with Service Pack 1 (SP1), and x64-based computers can access dynamic GUID partition table (GPT) disks.

For instructions on how to convert a basic disk into a dynamic disk, see Change a basic disk into a dynamic disk.

For more information about dynamic disks and volumes, see Dynamic disks and volumes.

# Using Windows NT 4.0 multidisk storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Windows NT 4.0 multidisk storage**

Windows Server 2003 operating systems do not support mirror sets, stripe sets, stripe sets with parity, or volume sets created in Windows NT 4.0. These multidisk volumes must be backed up and deleted, or otherwise converted to dynamic disks prior to installing Windows Server 2003 operating systems. If you did not back up the volumes before you upgraded to Windows Server 2003, you can access these volumes using the FTOnline utility. You can install the FTOnline utility using the Support Tools setup program located in the \Support\Tools folder on the Windows XP Professional and on the Windows Server 2003 operating system CDs.

In Windows Server 2003 operating systems, you can use dynamic disks to create the equivalent types of multidisk storage found on basic disks in Windows NT 4.0.

| Windows NT 4.0 Storage Type | Windows Server 2003 family Storage Type |
|---|---|
| mirror set | mirrored volume |
| stripe set | striped volume |
| stripe set with parity | RAID-5 volume |
| volume set | spanned volume |

Mirrored volumes and RAID-5 volumes are not available on Windows XP Home Edition or Windows XP Professional and cannot be extended.

For more information about dynamic disks, see Using Dynamic Disks.

# Using Dynamic Disks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using dynamic disks**


- Using simple volumes
- Using spanned volumes
- Using striped volumes
- Using mirrored volumes
- Using RAID-5 volumes

# Using simple volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using simple volumes**

A simple volume is a portion of a physical disk that functions as though it were a physically separate unit. Simple volumes are the dynamic storage equivalent of primary partitions in Windows NT 4.0 or earlier. When you have only one dynamic disk, they are the only kind of volume that you can create.

You can create simple volumes only on dynamic disks. Simple volumes (and the data they contain) cannot be accessed by, or created on, computers running MS-DOS, Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0, or Windows XP Home Edition that are configured to dual-boot with Windows XP Professional or Windows Server 2003 operating systems. If you want computers running these operating systems to be able to access the data, store the data on basic volumes.

You can increase the size of an existing simple volume by extending the volume onto unallocated space on the same disk or a different disk. To extend a simple volume, the volume must be unformatted or formatted with the version of NTFS used in Windows 2000 or Windows Server 2003 operating systems. With Windows Server 2003, simple volumes can be extended unless they are system partitions, boot partitions, or simple volumes that were formerly partitions on basic disks that were converted to dynamic using Windows 2000. By extending the simple volume on the same disk, the volume remains a simple volume and you can still mirror it.

You can also extend a simple volume to regions on other dynamic disks on the same computer. When you extend a simple volume to one or more other disks, it becomes a spanned volume. After a spanned volume is extended, no portion of it can be deleted without deleting the entire spanned volume. Spanned volumes cannot be mirrored.

For procedures on creating and extending simple volumes, see Manage Simple Volumes.

**Assigning drive letters**

You can create more than 26 volumes with Windows, but you cannot assign more than 26 drive letters for accessing these volumes. Drive letters A and B are typically reserved for floppy disk drives. If the computer does not have a floppy disk drive, you can assign drive letters A and B to removable drives, hard disk drives, or mapped network drives. Hard disk drives are typically assigned drive letters C through Z, while mapped network drives are assigned drive letters in reverse order (Z through C).

You should be careful when making drive-letter assignments because many programs for MS-DOS and Windows make references to a specific drive letter. The path environment variable shows specific drive letters with program names.

Volumes created after the 26th drive letter has been used must be accessed using volume mount points, as described in Using NTFS mounted drives.

**Formatting and labeling volumes**

Before you can store files and directories on the volumes that you have created, you must first format each volume for use with the file system that you want to work with. You can also assign descriptive volume labels at this time.

If a volume uses NTFS formatting, you can enable data compression. Before you can format a volume, you must assign it either a drive letter or a mount point, if it does not already have one.

To format a volume, see Format a dynamic volume.

**Deleting volumes**

Before deleting volumes, you need to make sure the information on them has been backed up onto another storage medium and verified, or that the data is no longer needed. Windows XP and Windows Server 2003 operating systems do not allow you to delete the system volume or boot volume.

# Using spanned volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using spanned volumes**

Spanned volumes combine areas of unallocated space from multiple disks into one logical volume, allowing you to more efficiently use all of the space and all drive letters on a multiple-disk system.

When you need to create a volume but do not have enough unallocated space for the volume on a single disk, you might be able to create a volume of sufficient size by combining sections of unallocated space from multiple disks into one spanned volume. The areas of unallocated space used to create spanned volumes can be different sizes. Spanned volumes are organized so that the space allocated to the volume on one disk gets filled up and then, starting at the next disk, the space allocated to the volume on that disk gets filled up.

Spanned volumes allow you to get more data on a disk without using mount points. By combining the space used by multiple disks into one spanned volume, you can free drive letters for other uses and enable the creation of a large volume for file system use.

Increasing the capacity of an existing volume is called *extending*. Existing spanned volumes formatted with the NTFS file system can be extended by the amount of unallocated space on all disks. However, after a spanned volume is extended, no portion of it can be deleted without deleting the entire spanned volume. Disk Management formats the new area without affecting any existing files on the original spanned volume. You cannot extend spanned volumes formatted with the FAT file system.

Before making any changes to spanned volumes, you should first back up all the information on the volume.

For procedures on working with spanned volumes, see Manage Spanned Volumes.

**Caution**

- Spanned volumes cannot be mirrored and do not offer fault tolerance. If one of the disks containing a spanned volume fails, the entire volume fails and all the data on it is lost.

You can create spanned volumes on dynamic disks only. You need at least two dynamic disks to create a spanned volume. MS-DOS, Windows 95, Windows 98, Windows NT 4.0, Windows XP Home Edition, and other operating systems lacking dynamic storage capability cannot recognize any spanned volumes created by Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems. Therefore, if you create a spanned volume on a dual-boot computer, the disks that make up that volume become unusable by those operating systems.

# Using striped volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using striped volumes**

Striped volumes are created by combining areas of free space on two or more disks into one logical volume. Striped volumes use RAID-0, which stripes data across multiple disks. Striped volumes cannot be extended or mirrored, and do not offer fault tolerance. If one of the disks containing a striped volume fails, the entire volume fails. When creating striped volumes, it is best to use disks that are the same size, model, and manufacturer.

With a striped volume, data is divided into blocks and spread in a fixed order among all the disks in the array, similar to spanned volumes. Striping writes files across all disks so that data is added to all disks at the same rate.

Despite their lack of fault tolerance, striped volumes offer the best performance of all the Windows disk management strategies and provide increased I/O performance by distributing I/O requests across disks. For example, striped volumes offer improved performance when:

- Reading from or writing to large databases.
- Collecting data from external sources at very high transfer rates.
- Loading program images, dynamic-link libraries (DLLs), or run-time libraries.

MS-DOS, Windows 95, Windows 98, Windows Millinnium Edition, Windows NT 4.0, Windows XP Home Edition, and other operating systems lacking dynamic storage capability cannot recognize any striped volumes created by Windows 2000, Windows XP Professional, or Windows Server 2003 operating systems. Therefore, if you create a striped volume on a dual-boot computer, that volume becomes unusable by those operating systems.

For instructions on creating a striped volume, see Create a striped volume.

# Using mirrored volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using mirrored volumes**

A mirrored volume is a fault-tolerant volume that provides data redundancy by using two copies, or mirrors, of the volume to duplicate the data stored on the volume. All data written to the mirrored volume is written to both mirrors, which are located on separate physical disks.

If one of the physical disks fails, the data on the failed disk becomes unavailable, but the system continues to operate using the unaffected disk. When one of the mirrors in a mirrored volume fails, you must break the mirrored volume to expose the remaining mirror as a separate volume with its own drive letter. You can then create a new mirrored volume with unused free space of equal or greater size on another disk. When creating mirrored volumes, it is best to use disks that are the same size and model and from the same manufacturer. You cannot extend a mirrored volume to increase the size of the volume later.

Because dual-write operations can degrade system performance, many mirrored volume configurations use duplexing, where each disk in the mirrored volume resides on its own disk controller. A duplexed mirrored volume has the best data reliability because the entire input/output (I/O) subsystem is duplicated. This means that if one disk controller fails, the other controller (and thus the disk on that controller) continues to operate normally. If you do not use two controllers, a failed controller makes both mirrors in a mirrored volume inaccessible until the controller is replaced.

When you mirror the system or boot volumes, you can make the configuration more fault tolerant by using a separate disk controller for each disk in the mirrored volume. This enables your computer to survive hard-disk or disk-controller failures. When creating mirrored volumes, it is best to use disks that are the same size and model, and from the same manufacturer. If you are using duplexing, it is recommended that you use identical disks and controllers, especially if you plan to mirror the system or boot volumes. For more information, see Create and test a mirrored system or boot volume. On Itanium-based computers, you cannot mirror the Extensible Firmware Interface (EFI) system partition on GUID partition table (GPT) disks. Instead, you must use the **bootcfg** command to clone the EFI system partition.

**Important**

- When mirroring the system volume, always test to make sure you can start the operating system from each mirror if one of the disks fails. To help prevent startup problems, always use identical disks and controllers.

You can create mirrored volumes using a combination of master boot record (MBR) and GUID partition table (GPT) disks, but the MBR cylinder-alignment restriction may cause some difficulties in creating mirrored volumes. Always mirror the MBR disk to the GPT disk to avoid cylinder-alignment difficulties.

For step-by-step procedures on how to create, break, or repair a mirrored volume, or to add or remove a mirror from an existing mirrored volume, see Manage Mirrored Volumes.

For information about which operating systems support mirrored volumes, see Partition styles.

For more information about mirrored volumes, see "Disk Management" at the Microsoft Windows Resource Kits Web site.

# Using RAID-5 volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using RAID-5 volumes**

A RAID-5 volume is a fault-tolerant volume with data and parity striped intermittently across three or more physical disks. If a portion of a physical disk fails, you can recreate the data that was on the failed portion from the remaining data and parity. RAID-5 volumes are a good solution for data redundancy in a computer environment where most activity consists of reading data.

You can create a RAID-5 volume using hardware- or software-based solutions. With hardware-based RAID, an intelligent disk controller handles the creation and regeneration of redundant information on the disks that make up the RAID-5 volume. Windows Server 2003 operating systems provide software-based RAID, where the creation and regeneration of redundant information on the disks in the RAID-5 volume is handled by Disk Management. In either case, data is stored across all members in the disk array.

In general, hardware-based RAID offers performance advantages over software-based RAID because hardware-based RAID incurs no overhead on the system processor. For example, you can improve data throughput significantly by implementing RAID-5 through hardware that does not use system software resources. This is accomplished by using more disks at a given capacity than would typically be available in a conventional storage solution. Read and write performance and total storage size can be further improved by using multiple disk controllers.

RAID-5 volumes have better read performance than mirrored volumes. When a member is missing, however, such as when a disk has failed, the read performance is degraded by the need to recover the data with the parity information. Nevertheless, this strategy is recommended over mirrored volumes for programs that require redundancy and are primarily read-oriented. Write performance is reduced by the parity calculation. A write operation also requires three times more memory than a read operation during normal operation. Moreover, when a volume fails, reading requires at least three times more memory than before the failure. Both conditions are caused by the parity calculation.

RAID-5 volumes include one parity block per stripe. Therefore, you must use at least three disks, rather than two, to allow for the parity information. Parity stripes are distributed across all the volumes to balance the input/output (I/O) load. When regenerating a RAID-5 volume, the parity information with the data on the healthy disks is used to recreate the data on the failed disk.

For step-by-step information about how to create and repair RAID-5 volumes, see Manage RAID-5 Volumes.

For information about which operating systems support RAID-5 volumes, see Partition styles.

For more information about mirrored volumes, see "Disk Management" at the Microsoft Windows Resource Kits Web site.

# Managing disks and volumes from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing disks and volumes from the command line**
In addition to using the Disk Management snap-in, you can use command-line utilities to manage disks and volumes.

For example, you can use:

- chkdsk to check disks for errors and repair any errors found.
- Convert to change FAT or FAT32 volumes into NTFS volumes.
- DiskPart to extend basic or dynamic volumes, add or break mirrors, assign or remove a disk's drive letter, create or delete partitions and volumes, convert basic disks to dynamic disks, or convert master boot record (MBR) disks to GUID partition table (GPT) disks, import disks and bring offline disks and volumes online.
- Format to format a volume or mounted drive with a file system.
- Fsutil to perform many NTFS file system related tasks, such as managing disk quotas, dismounting a volume, or querying volume information. Because **fsutil** is quite powerful, it should be used only by advanced users who have a thorough knowledge of Windows XP Professional or Windows Server 2003 operating systems.
- Mountvol to mount a volume at an NTFS folder or unmount the volume from the NTFS folder.

For information about other command-line utilities, see Command-line reference A-Z.

# Disk Management Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

- You can learn more about Disk Management in "Disk Management" at the Microsoft Resource Kits Web site. The Microsoft Windows Resource Kits provide information technology professionals with the technical information and tools they need to successfully deploy, manage, and support Windows operating systems.

# Troubleshooting Disk Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 8 Beta

**Troubleshooting**
What problem are you having?

- Partitions on basic disks added to the system do not appear in the Disk Management volume list view.
- A basic disk's status is Not Initialized.
- A basic or dynamic disk's status is Unreadable.
- A dynamic disk's status is Foreign.
- A dynamic disk's status is Online (Errors).
- A dynamic disk's status is Offline or Missing.
- A basic or dynamic volume's status is Failed.
- A basic or dynamic volume's status is Unknown.
- A dynamic volume's status is Data Incomplete.
- Cannot manage striped, mirrored, or RAID-5 volumes using Disk Management or DiskPart.
- A dynamic volume's status is Healthy (At Risk).
- A mirrored volume's status is Data Not Redundant.
- A mirrored volume's status is Failed.
- A mirrored volume's status is Failed Redundancy.
- A mirrored volume's status is Stale Data.
- A RAID-5 volume's status is Data Not Redundant.
- A RAID-5 volume's status is Failed.
- A RAID-5 volume's status is Failed Redundancy.
- A RAID-5 volume's status is Stale Data.

**Partitions on basic disks added to the system do not appear in the Disk Management volume list view.**
**Cause:** For better interoperability in storage area network (SAN) environments, on Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, volumes on basic disks added to the system are not automatically mounted and assigned drive letters by default.

Dynamic disks appear as **Foreign** when they are added to the system. To use the volumes, you must import the **Foreign** disks and then bring the volumes **Online**.

Removable media devices (such as Zip or Jaz drives), and optical discs (such as CD-ROM or DVD-RAM) are always

automatically mounted by the system.

**Solution:** Manually mount the basic volumes by assigning drive letters, or by creating mount points using Disk Management or the **DiskPart** or **mountvol** commands.

For instructions describing how to assign a drive letter, see Assign, change, or remove a drive letter.

For more information about the **DiskPart** or **mountvol** commands, see DiskPart or Mountvol.

**A basic disk's status is Not Initialized.**

**Cause:** The disk does not contain a valid signature. After you install a new disk, the operating system must write a disk signature, the end of sector marker (also called signature word), and a master boot record or GUID partition table before you can create partitions on the disk. When you first start Disk Management after installing a new disk, a wizard appears that provides a list of the new disks detected by the operating system. If you cancel the wizard before the disk signature is written, the disk status remains **Not Initialized**.

**Solution:** Initialize the disk. The disk status briefly changes to **Initializing** and then **Online** status. For instructions describing how to initialize a disk, see Initialize new disks.

For more information about disk status descriptions, see Disk status descriptions.

**A basic or dynamic disk's status is Unreadable.**

**Cause:** The basic or dynamic disk is not accessible and might have experienced hardware failure, corruption, or I/O errors. The disk's copy of the system's disk configuration database might be corrupted. An error icon appears on disks that display the **Unreadable** status.

Disks might also display the **Unreadable** status while they are spinning up or when Disk Management is rescanning all of the disks on the system. In some cases, an unreadable disk has failed and is not recoverable. For dynamic disks, the **Unreadable** status usually results from corruption or I/O errors on part of the disk, rather than failure of the entire disk.

**Solution:** Rescan the disks or restart the computer to see if the disk status changes. For instructions describing how to rescan the disks, see Update disk information.

For more information about disk status descriptions, see Disk status descriptions.

**A dynamic disk's status is Foreign.**

**Cause:** The **Foreign** status occurs when you move a dynamic disk to the local computer from another computer running Windows 2000, Windows XP Professional, Windows XP 64-bit Edition (Itanium), or the Windows Server 2003 family of operating systems. A warning icon appears on disks that display the **Foreign** status.

In some cases, a disk that was previously connected to the system can display the **Foreign** status. Configuration data for dynamic disks is stored on all dynamic disks, so the information about which disks are owned by the system is lost when all dynamic disks fail.

**Solution:** Add the disk to your computer's system configuration so that you can access data on the disk. To add a disk to your computer's system configuration, import the foreign disk (right-click the disk and then click **Import Foreign Disks**). Any existing volumes on the foreign disk become visible and accessible when you import the disk.

For more information about disk status descriptions, see Disk status descriptions.

**A dynamic disk's status is Online (Errors).**

**Cause:** The dynamic disk has I/O errors on a region of the disk. A warning icon appears on the dynamic disk with errors.

**Solution:** If the I/O errors are temporary, reactivate the disk to return it to **Online** status. For instructions describing how to bring the disk back online, see Reactivate a missing or offline dynamic disk.

For more information about disk status descriptions, see Disk status descriptions.

**A dynamic disk's status is Offline or Missing.**

**Cause:** An **Offline** dynamic disk might be corrupted or intermittently unavailable. An error icon appears on the offline dynamic disk.

If the disk status is **Offline** and the disk's name changes to **Missing**, the disk was recently available on the system but can no longer be located or identified. The missing disk may be corrupted, powered down, or disconnected.

**Solution:**

**To bring a disk that is Offline and Missing back online**

Repair any disk, controller, or cable problems and make sure that the physical disk is turned on, plugged in, and attached to the computer. Next, use the **Reactivate Disk** command to bring the disk back online. For instructions describing how to reactivate missing or offline disks, see Reactivate a missing or offline dynamic disk.

If the disk status remains **Offline** and the disk name remains **Missing**, and you determine that the disk has a problem that cannot be repaired, you can remove the disk from the system (right-click the disk and then click **Remove Disk**). However, before you can remove the disk, you must delete all volumes (or mirrors) on the disk. You can save any mirrored volumes on the disk by removing the mirror instead of the entire volume. Deleting a volume destroys the data in the volume, so you should remove a disk only if you are absolutely certain that the disk is permanently damaged and unusable.

**To bring a disk that is Offline and is still named Disk # (not Missing) back online, try one or more of the following:**

- In Disk Management, right-click the disk and then click **Reactivate Disk** to bring the disk back online. If the disk status remains **Offline**, check the cables and disk controller, and make sure that the physical disk is healthy. Correct any problems and try to reactivate the disk again. If the disk reactivation succeeds, any volumes on the disk should automatically return to the **Healthy** status.

- In Event Viewer, check the event logs for any disk related errors such as “no good config copies”. If the event logs contain the "no good config copies" error, contact Microsoft Product Support Services.

- Try moving the disk to another computer. If you can get the disk to go **Online** on another computer, the problem is most likely due to the configuration of the computer where the disk does not go **Online**.

- Try moving the disk to another computer that has dynamic disks. Import the disk on that computer and then move it back to the computer where it would not go **Online**.

For more information about disk status descriptions, see Disk status descriptions.

**A basic or dynamic volume's status is Failed.**
**Cause:** The basic or dynamic volume cannot be started automatically, the disk is damaged, or the file system is corrupt. Unless the disk or file system can be repaired, the **Failed** status indicates data loss.

**Solution:** If the volume is a basic volume with **Failed** status, make sure that the underlying physical disk is turned on, plugged in, and attached to the computer. No other user action is possible for basic volumes.

If the volume is a dynamic volume with **Failed** status, make sure the underlying disks are online. If not, return the disks to the **Online** status. If this succeeds, the volume automatically restarts and returns to the **Healthy** status. For instructions describing how to bring a disk online, see Reactivate a missing or offline dynamic disk.

If the dynamic disk returns to the **Online** status, but the dynamic volume does not return to the **Healthy** status, you can reactivate the volume manually. For instructions describing how to reactivate a volume, see Reactivate a dynamic volume.

If the dynamic volume is a mirrored or RAID-5 volume with old data, bringing the underlying disk online will not automatically restart the volume. If the disks that contain current data are disconnected, you should bring those disks online first (to allow the data to become synchronized). Otherwise, restart the mirrored or RAID-5 volume manually, and then run the Error-checking tool or Chkdsk.exe.

For instructions describing how to reactivate a volume, see Reactivate a dynamic volume. For instructions describing how to run the Error-checking tool, see Detecting and repairing disk errors. To run Chkdsk.exe, click **Start**, click **Run**, type **chkdsk**, and then click **OK**.

For more information about volume status descriptions, see Volume status descriptions.

**A basic or dynamic volume's status is Unknown.**
**Cause:** The **Unknown** status occurs when the boot sector for the volume is corrupted (possibly due to a virus) and you can no longer access data on the volume. The **Unknown** status also occurs when you install a new disk but do not successfully complete the wizard to create a disk signature.

**Solution:** Initialize the disk. For instructions describing how to initialize a disk, see Initialize new disks.

For more information about volume status descriptions, see Volume status descriptions.

**A dynamic volume's status is Data Incomplete.**
**Cause:** You moved some but not all of the disks in a multidisk volume. Data on this volume will be destroyed unless you

move and import the remaining disks that contain this volume.

**Solution:** Move all of the disks that comprise the multidisk volume to the computer and then import the disks. For instructions describing how to move and import disks, see Move disks to another computer.

If you no longer require the multidisk volume, you can import the disk and create new volumes on it. Right-click the volume with **Failed** or **Failed Redundancy** status and then click **Delete Volume**. Right-click the disk and then click **New Volume**.

For more information about volume status descriptions, see Volume status descriptions.

**A dynamic volume's status is Healthy (At Risk).**
**Cause:** Indicates that the dynamic volume is currently accessible, but I/O errors have been detected on the underlying dynamic disk. If an I/O error is detected on any part of a dynamic disk, all volumes on the disk display the **Healthy (At Risk)** status and a warning icon appears on the volume.

When the volume status is **Healthy (At Risk)**, an underlying disk's status is usually **Online (Errors)**.

**Solution:** Return the underlying disk to the **Online** status. Once the disk is returned to **Online** status, the volume should return to the **Healthy** status. If the **Healthy (At Risk)** status persists, the disk might be failing. Back up the data and replace the disk as soon as possible. For instructions describing how to bring the disk back online, see Reactivate a missing or offline dynamic disk.

For more information about volume status descriptions, see Volume status descriptions.

**Cannot manage striped, mirrored, or RAID-5 volumes using Disk Management or DiskPart.**
**Cause:** Some non-Microsoft disk management products replace Microsoft Logical Disk Manager (LDM) for advanced disk management, which can disable the LDM.

**Solution:** If you are using non-Microsoft disk management software that has disabled LDM, you must contact the vendor on the non-Microsoft disk management software for support or assistance in troubleshooting problems with the disk configuration.

**A mirrored volume's status is Data Not Redundant.**
**Cause:** One half, but not both halves, of a mirrored volume was imported. In Disk Management, the imported half of the mirrored volume receives **Failed Redundancy** status, while the disk that holds the half of the mirror that was not imported receives **Missing** status. You can import the **Missing** disk at a later time to restore redundancy.

**Solution:** Move all disks that span this volume to the new computer at the same time, and then import all of the disks together. For instructions describing how to move and import disks, see Move disks to another computer.

If you no longer require a mirrored volume, you can break the mirror and retain the imported half of the mirror as a simple volume. Right-click the mirrored volume with **Failed Redundancy** status and then click **Remove Mirror**. In the **Remove Mirror** dialog box, click the disk with **Missing** status and then click **Remove Mirror**. The imported half of the mirror with **Failed Redundancy** status becomes a simple volume with **Healthy** status. For more information about how to break a mirrored volume, see Remove a mirror from a mirrored volume.

For more information about volume status descriptions, see Volume status descriptions.

For more information about disk status descriptions, see Disk status descriptions.

**A mirrored volume's status is Failed.**
**Cause:** The mirrored volume cannot be started automatically, or the disk is damaged. Unless the disk can be repaired, the **Failed** status indicates data loss.

**Solution:** Make sure that the underlying physical disks are turned on, plugged in, and attached to the computer. Make sure that the underlying physical disks are online. If not, return the disks to the **Online** status. If this succeeds, the volume automatically restarts and changes to **Failed Redundancy** status. For instructions describing how to bring a disk online, see Reactivate a missing or offline dynamic disk.

If the dynamic disk returns to the **Online** status, but the dynamic volume does not return to the **Healthy** status, you can reactivate the volume manually. For instructions describing how to reactivate a volume, see Reactivate a dynamic volume.

If the dynamic volume is a mirrored volume with old data, bringing the underlying physical disk online will not automatically restart the volume. If the disks that contain current data are disconnected, you should bring those disks online first (to allow the data to become synchronized). Otherwise, restart the mirrored volume manually, and then run the Error-checking tool or **Chkdsk.exe**.

For instructions describing how to reactivate a volume, see Reactivate a dynamic volume. For instructions describing how to run the Error-checking tool, see Detecting and repairing disk errors. To run Chkdsk.exe, click **Start**, click **Run**, type **chkdsk**, and then click **OK**.

For more information about volume status descriptions, see Volume status descriptions.

**A mirrored volume's status is Failed Redundancy.**
**Cause:** One or both of the members of the mirrored volume has failed, and the volume is no longer fault tolerant. To avoid data loss, you should attempt to repair the volume as soon as possible.

**Solution:** If the dynamic disk's status is **Offline** or **Missing** (it cannot be accessed at all), an icon **(X)** appears in the graphical view of the Missing or Offline disk. For instructions describing how to repair the mirrored volume, see Reconnect the disk and repair the mirrored volume.

If the disk's status is **Online (Errors)**, an icon **(!)** appears in the graphical view of the disk. For instructions describing how to reactivate the disk, see Reactivate a mirrored volume.

For more information about volume status descriptions, see Volume status descriptions.

For more information about disk status descriptions, see Disk status descriptions.

**A mirrored volume's status is Stale Data.**
**Cause:** If you attempt to import disks that contain a mirrored volume, but one of the volumes on the disks contains stale mirror information, the status **Stale Data** appears in the **Import Foreign Disks** dialog box. This problem can occur when the volumes on the disks you moved had a status other than **Healthy** before you moved them.

**Solution:** Move the disks back to the computer that they came from. In Disk Management, click **Action**, and then click **Rescan Disks**; or restart the computer to make sure all of the disks are installed correctly.

If the volume showed the **Failed Redundancy** status before the disks were moved, the volume will still show the **Failed Redundancy** status when the disks are moved back. Make sure that the status of the disk with stale mirror information is not **Offline**. If the status is **Offline**, check for any hardware problems, reconnect the disk if necessary, and then bring the disk back online. For instructions describing how to bring a mirrored volume back online, see Reconnect the disk and repair the mirrored volume.

If the volume still shows the **Failed Redundancy** status, reactivate the volume. Disk Management will attempt to bring the volume back online. You will see a warning message indicating that Disk Management might not be able to bring back all of the data on the volume. For instructions describing how to reactivate a volume, see Reactivate a dynamic volume.

If Disk Management successfully brings the volumes back online, the status of the volume becomes **Healthy** and any mirrored volumes on the disk will be resynchronized. Now you can move all of the disks that contain this volume to another computer, attach them, and import all of the disks at the same time.

For more information about volume status descriptions, see Volume status descriptions.

For more information about disk status descriptions, see Disk status descriptions.

**A RAID-5 volume's status is Data Not Redundant.**
**Cause:** You moved some but not all of the disks in a RAID-5 volume. If you choose to import all but one of the disks, the RAID-5 volume receives the **Failed Redundancy** status. If more than one disk is missing from the RAID-5 volume, the volume receives the **Failed** status.

**Solution:** Move all disks that comprise this volume to the new computer at the same time, and then import all of the disks together. For instructions describing how to move and import disks, see Move disks to another computer.

If you imported all but one of the disks in the RAID-5 volume, the volume receives **Failed Redundancy** status. Right-click the volume and then click **Repair Volume**. Select another suitable dynamic disk to repair the RAID-5 volume. After the RAID-5 volume returns to **Healthy** status, you can right-click the missing disk and then click **Remove Disk**.

If you no longer require the RAID-5 volume, you can import the disks and create new volumes on them. The RAID-5 volume receives the **Failed Redundancy** or **Failed** status. Right-click the volume and then click **Delete Volume**. Right-click each disk and then click **New Volume**.

For more information about volume status descriptions, see Volume status descriptions.

**A RAID-5 volume's status is Failed.**
**Cause:** The RAID-5 volume cannot be started automatically, or the disk is damaged. Unless the disk can be repaired, the

**Failed** status indicates data loss.

**Solution:** Make sure that the underlying physical disk is turned on, plugged in, and attached to the computer. Make sure that the underlying physical disks are online. If not, return the disks to the **Online** status. If this succeeds, the volume automatically restarts and returns to the **Healthy** status. For instructions describing how to bring a disk online, see Reactivate a missing or offline dynamic disk.

If the RAID-5 volume returns to the **Online** status, but the volume does not return to the **Healthy** status, you can reactivate the volume manually. For instructions describing how to reactivate a volume, see Reactivate a dynamic volume.

If the dynamic volume is a RAID-5 volume with old data, bringing the underlying disk online will not automatically restart the volume. If the disks that contain current data are disconnected, you should bring those disks online first (to allow the data to become synchronized). Otherwise, restart the RAID-5 volume manually, and then run the Error-checking tool or **Chkdsk.exe**.

For instructions describing how to reactivate a volume, see Reactivate a dynamic volume. For instructions describing how to run the Error-checking tool, see Detecting and repairing disk errors. To run Chkdsk.exe, click **Start**, click **Run**, type **chkdsk**, and then click **OK**.

For more information about volume status descriptions, see Volume status descriptions.

**A RAID-5 volume's status is Failed Redundancy.**
**Cause:** One of the members of the RAID-5 volume has failed and the volume is no longer fault tolerant. To avoid data loss, you should attempt to repair the volume as soon as possible.

**Solution:** If the dynamic disk's status is **Offline** or **Missing** (it can not be accessed at all), an icon **(X)** appears in the graphical view of the Missing or Offline disk. For instructions describing how to repair a RAID-5 volume, see Reconnect the disk and repair the RAID-5 volume.

If the disk's status is **Online (Errors)**, an icon **(!)** appears in the graphical view of the partially failed disk. For instructions describing how to reactivate the disk, see Reactivate a RAID-5 disk.

For more information about volume status descriptions, see Volume status descriptions.

For more information about disk status descriptions, see Disk status descriptions.

**A RAID-5 volume's status is Stale Data.**
**Cause:** If you attempt to import disks that contain a RAID-5 volume, but one of the volumes on the disks contains stale parity information, the status **Stale Data** appears in the **Import Foreign Disks** dialog box. This problem can occur when the volumes on the disks you moved had a status other than **Healthy** before you moved them.

**Solution:** Move the disks back to the computer that they came from. In Disk Management, click **Action**, and then click **Rescan Disks**; or restart the computer to make sure all of the disks are installed correctly.

If the volume showed the **Failed Redundancy** status before the disks were moved, the volume will still show the **Failed Redundancy** status when the disks are moved back. Make sure that the status of the disk that contained the stale parity information is not **Offline**. If the status is **Offline**, check for any hardware problems, reconnect the disk if necessary, and then bring the disk back online. For instructions describing how to bring a RAID-5 volume back online, see Reactivate a RAID-5 disk.

If the volume still shows the **Failed Redundancy** status, reactivate the volume. Disk Management will attempt to bring the volume back online. You will see a warning message indicating that Disk Management might not be able to bring back all of the data on the volume. For instructions describing how to reactivate a volume, see Reactivate a dynamic volume.

If Disk Management successfully brings the volumes back online, the status of the volume becomes **Healthy** and any RAID-5 volumes will regenerate parity. Now you can move all of the disks that contain this volume to another computer. However, if you move all but one of the disks, then the volume will show the **Data Not Redundant** status in the **Import Foreign Disk** dialog box. You should attach the missing disks and import all of the disks at the same time.

For more information about volume status descriptions, see Volume status descriptions.

For more information about disk status descriptions, see Disk status descriptions.

# Disk Defragmenter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk Defragmenter**

Disk Defragmenter is a system utility for analyzing local volumes and locating and consolidating fragmented files and folders. You can also defragment disks from a command line using the Defrag command.

- Before defragmenting files or folders, see Checklist: Defragmenting volumes.
- For tips about using Disk Defragmenter, see Best practices for Disk Defragmenter.
- For help with specific tasks, see Disk Defragmenter How To....
- For general background information, see Disk Defragmenter Concepts.
- For problem-solving instructions, see Troubleshooting Disk Defragmenter.

# Checklist: Defragmenting volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Defragmenting volumes**

| Step | Reference |
|---|---|
| Review the overview. | Disk Defragmenter overview |
| Review best practices. | Best practices for Disk Defragmenter |
| Confirm that you are a member of the Administrators group.<br><br>If you have access to the administrator account and password, it is a security best practice to log on as a user that is not a member of the Administrators group and run Disk Defragmenter by right-clicking the Disk Defragmenter icon, and then clicking Run as. | Default local groups; Why you should not run your computer as an administrator; Run a program with administrative credentials |
| Analyze the volume. | Analyze a volume |
| Defragment the volume. | Defragment a volume |
| Review the report. | Analysis and defragmentation reports |

# Best practices for Disk Defragmenter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Analyze before defragmenting a volume**

- Analyze volumes before defragmenting them. After analyzing a volume, a dialog box tells you the percentage of fragmented files and folders on the volume and recommends whether to defragment the volume. Analyze volumes regularly and defragment them only when Disk Defragmenter recommends it. A good guideline is to analyze volumes at least once a week. If you seldom need to defragment volumes, analyze volumes monthly instead of weekly.

**Analyze after large numbers of files are added**

- Volumes might become excessively fragmented when users add a large number of files or folders, so be sure to analyze volumes after this happens. Generally, volumes on busy file servers should be defragmented more often than those on single-user workstations.

**Ensure that your volume has at least 15% free space**

- A volume must have at least 15% free space for Disk Defragmenter to completely and adequately defragment it. Disk Defragmenter uses this space as a sorting area for file fragments. If a volume has less than 15% free space, Disk Defragmenter will only partially defragment it. To increase the free space on a volume, delete unneeded files or move them to another disk.

**Defragment during low-usage periods**

- Defragment file server volumes during low-volume usage periods to minimize the effect that the defragmentation process has on file server performance. The time that Disk Defragmenter takes to defragment a volume depends on several factors, including the size of the volume, the number of files on the volume, the number of fragmented files, and available system resources.

**Defragment after installing software or installing Windows**

- Defragment volumes after installing software or after performing an upgrade or clean install of Windows. Volumes often become fragmented after installing software, so running Disk Defragmenter helps to ensure the best file system performance.

# Disk Defragmenter How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Analyze a volume
- Defragment a volume

# Analyze a volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To analyze a volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Disk Defragmenter.
2. Click the volume that you want to check for fragmented files and folders, and then click **Analyze**.

   When Disk Defragmenter finishes analyzing the volume, it displays its analysis of the volume in **Estimated disk usage before defragmentation**. It also displays a dialog box that tells you whether you need to defragment the volume.
3. Click **View Report** to view the analysis report, which displays more detailed volume information and a list of the most fragmented files.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Disk Defragmenter, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Disk Defragmenter**.
- Running the **defrag** command and Disk Defragmenter are mutually exclusive. If you are using Disk Defragmenter to defragment a volume and you run the **defrag** command at a command prompt, the **defrag** command fails. Conversely, if you run the **defrag** command, and open Disk Defragmenter, the defragmentation options in Disk Defragmenter are unavailable.

**Using a command line**

1. Open Command Prompt.
2. Type **defrag***volume* [**/a**] [**/v**]

| Value | Description |
|---|---|
| *volume* | The drive letter or a mount point of the volume to be analyzed. |
| **/a** | Analyzes the volume, displays a summary of the analysis report, and indicates whether you should defragment the volume. |
| **/v** | Displays the complete analysis report. |

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group

might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, type **defrag /?**

- While **defrag** is analyzing and defragmenting a volume, it displays a blinking cursor. When **defrag** is finished analyzing and defragmenting the volume, it displays the analysis report, the defragmentation report, or both reports, and then exits to the command prompt.

- By default, **defrag** displays a summary of both the analysis and defragmentation reports if you do not specify the **/a** or **/v** parameters.

- You can send the reports to a text file by typing >*filename.txt*, where *filename.txt* is a file name you specify. For example:

  **defrag** *volume* **/v** >*filename.txt*

- Running the **defrag** command and Disk Defragmenter are mutually exclusive. If you are using Disk Defragmenter to defragment a volume and you run the **defrag** command at a command prompt, the **defrag** command fails. Conversely, if you run the **defrag** command, and open Disk Defragmenter, the defragmentation options in Disk Defragmenter are unavailable.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Why volumes become fragmented
Defragment a volume
Analysis and defragmentation reports
Defrag

# Defragment a volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To defragment a volume**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Disk Defragmenter.
2. Click the volume that you want to defragment, and then click **Defragment**.

   After the defragmentation is complete, Disk Defragmenter displays the results in **Estimated disk usage after defragmentation**.
3. Click **View Report** to view the Defragmentation Report, which displays detailed information about the volume you defragmented.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Disk Defragmenter, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Disk Defragmenter**.
- You should analyze volumes before defragmenting them. This tells you whether you need to take the time to defragment them.
- A volume must have at least 15% free space for Disk Defragmenter to completely and adequately defragment it. Disk Defragmenter uses this space as a sorting area for file fragments. If a volume has less than 15% free space, Disk Defragmenter will only partially defragment it. To increase the free space on a volume, delete unneeded files or move them to another disk.
- You cannot defragment volumes that the file system has marked as dirty, which indicates possible corruption. You must run **chkdsk** on a dirty volume before you can defragment it. You can determine if a volume is dirty by using the **fsutil dirty query** command. For more information about **chkdsk** and **fsutil dirty**, see Related Topics.
- The time that Disk Defragmenter takes to defragment a volume depends on several factors, including the size of the volume, the number and size of files on the volume, the percentage of fragmentation in the volume, and available system resources.
- You can defragment only local file system volumes, and you can run only one instance of Disk Defragmenter at a time.
- To stop or temporarily interrupt defragmenting a volume, click **Stop** or **Pause**, respectively.
- If you start Disk Defragmenter while performing a backup on the same volume, Disk Defragmenter stops.
- Running the **defrag** command and Disk Defragmenter are mutually exclusive. If you are using Disk Defragmenter to

defragment a volume and you run the **defrag** command at a command prompt, the **defrag** command fails. Conversely, if you run the **defrag** command, and open Disk Defragmenter, the defragmentation options in Disk Defragmenter are unavailable.

#### Using a command line

1. Open Command Prompt.

2. Type **defrag***volume* [**/v**]

| Value | Description |
|---|---|
| *volume* | The drive letter or a mount point of the volume to be defragmented. |
| **/v** | Displays the complete analysis and defragmentation reports. |

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, type **defrag /?**

- A volume must have at least 15% free space for **defrag** to completely and adequately defragment it. **Defrag** uses this space as a sorting area for file fragments. If a volume has less than 15% free space, **defrag** will only partially defragment it. To increase the free space on a volume, delete unneeded files or move them to another disk.

- You cannot defragment volumes that the file system has marked as dirty, which indicates possible corruption. You must run **chkdsk** on a dirty volume before you can defragment it. You can determine if a volume is dirty by using the **fsutil dirty query** command. For more information about **chkdsk** and **fsutil dirty**, see Related Topics.

- While **defrag** is analyzing and defragmenting a volume, it displays a blinking cursor. When **defrag** is finished analyzing and defragmenting the volume, it displays the analysis report, the defragmentation report, or both reports, and then exits to the command prompt.

- By default, **defrag** displays a summary of both the analysis and defragmentation reports if you do not specify the **/a** or **/v** parameters.

- You can send the reports to a text file by typing >*filename.txt*, where *filename.txt* is a file name you specify. For example:

  **defrag***volume***/v** >*filename.txt*

- To stop the defragmentation process, at the command prompt, press CTRL+C.

- If you run **defrag** while performing a backup on the same volume, **defrag** stops.

- Running the **defrag** command and Disk Defragmenter are mutually exclusive. If you are using Disk Defragmenter to defragment a volume and you run the **defrag** command at a command prompt, the **defrag** command fails. Conversely, if you run the **defrag** command, and open Disk Defragmenter, the defragmentation options in Disk Defragmenter are unavailable.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Working with MMC console files
Analyze a volume
Analysis and defragmentation reports
Defrag
chkdsk
Fsutil
Fsutil: dirty

# Disk Defragmenter Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Disk Defragmenter overview
- Understanding Disk Defragmenter
- Resources for Disk Defragmenter

# Disk Defragmenter overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk Defragmenter overview**

Disk Defragmenter and the Defrag command analyze local volumes and consolidate fragmented files and folders so that each occupies a single, contiguous space on the volume. As a result, your system can access files and folders and save new ones more efficiently. By consolidating your files and folders, Disk Defragmenter also consolidates a volume's free space, making it less likely that new files will be fragmented. The process of consolidating fragmented files and folders is called defragmentation.

The amount of time that defragmentation takes depends on several factors, including the size of the volume, the number and size of files on the volume, the amount of fragmentation, and the available local system resources. You can find the fragmented files and folders before defragmenting them by analyzing the volume first. You can then see how fragmented the volume is and decide whether you would benefit from defragmenting the volume. For step-by-step instructions describing how to analyze or defragment a volume, see Analyze a volume and Defragment a volume.

Disk Defragmenter can defragment volumes that are formatted with the file allocation table (FAT) file system, the FAT32 file system, and the NTFS file system.

# Understanding Disk Defragmenter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Disk Defragmenter**

This section covers:

- Why volumes become fragmented
- Why files are not moved to some regions of NTFS-formatted volumes
- Disk Defragmenter window
- Analysis and defragmentation reports

# Why volumes become fragmented

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Why volumes become fragmented**

Volumes become fragmented as users create and delete files and folders, install new software, or download files from the Internet. Computers typically save files in the first contiguous free space that is large enough for the file. If a large enough free space is not available, the computer saves as much of the file as possible in the largest available space and then saves the remaining data in the next available free space, and so on.

After a large portion of a volume has been used for file and folder storage, most of the new files are saved in pieces across the volume. When you delete files, the empty spaces left behind fill in randomly as you store new ones.

The more fragmented the volume is, the slower the computer's file input/output performance will be.

# Why files are not moved to some regions of NTFS-formatted volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Why files are not moved to some regions of NTFS-formatted volumes**

On volumes formatted with the NTFS file system, Windows reserves a portion of the free space called the master file table (MFT) Zone, as a space where the MFT can grow. Windows stores all the information necessary to retrieve files from the volume in the MFT.

Disk Defragmenter does not move files to the MFT Zone because it is reserved by Windows. However, Windows allows other files to be saved in the space reserved for the MFT Zone when no other disk space is available.

# Disk Defragmenter window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk Defragmenter window**

The Disk Defragmenter window consists of two main areas. The upper portion lists the volumes on the local computer. The lower portion displays a graphical representation of the amount of fragmentation on the volume in colored horizontal displays called **Estimated disk usage before defragmentation** and **Estimated disk usage after defragmentation**.

Each bar under **Estimated disk usage before defragmentation** and **Estimated disk usage after defragmentation** represents a grouping of dozens or even hundreds of clusters, depending on the size of the volume and the cluster. The displays are not large enough to depict each disk cluster in a volume as a vertical colored bar.

Disk Defragmenter does not attempt to consolidate all of the free space on a volume. Although free space fragmented into hundreds of pieces adversely affects performance, free space split into a few pieces does not. Having all of the free space consolidated in a single location provides very little performance benefit. Because fragmented files, contiguous files, files that cannot be relocated from their current positions on the disk (unmovable files), and free space can all reside in the same group of clusters, the color of each vertical bar is determined by the following rules:

| Color | Description |
|---|---|
| Red | Most of the clusters are part of a fragmented file. |
| Blue | Most of the clusters are contiguous files with clusters in the group that contain only free space and contiguous clusters. |
| Green | Most of the clusters are part of a file that cannot be moved from its current location.<br><br>On volumes formatted using the file allocation table (FAT), FAT32, or NTFS file system, the green areas typically represent paging files, which are also called swap files. On volumes formatted using the NTFS file system, however, they may also represent space used by the NTFS Change journal and the NTFS log file. |
| White | Most of the clusters are free space and contiguous clusters. |

By comparing **Estimated disk usage before defragmentation** to **Estimated disk usage after defragmentation**, you can view the improvement in your volume after defragmenting it. The display shows less detail than the reports, so if you compare the results, you might notice discrepancies. For precise, numerical figures, see Analysis and defragmentation reports.

For step-by-step instructions describing how to analyze or defragment a volume, see Analyze a volume and Defragment a volume.

# Analysis and defragmentation reports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Analysis and defragmentation reports**

You can view the Analysis and Defragmentation Reports after analyzing or defragmenting a volume. The Analysis and Defragmentation Reports provide detailed information about the volume that was scanned for fragmented files, including the volume size, the amount of free space available, the number of fragmented files and folders, and the average number of fragments per file.

For instructions on how to analyze a volume and view the report, see Analyze a volume. For instructions on how to defragment a volume and view the report, see Defragment a volume.

**Analysis Report**

You can view or print the Analysis Report only after you analyze a volume. The **Analysis Report** dialog box displays detailed information about the volume that was scanned for fragmented files and folders, including the path and name of the most fragmented files on the volume, and the number of fragments those files are in, the volume size, the amount of free space available, and the average number of fragments per file. If you use these files frequently, the impact on your system performance might be greater than indicated in **Average fragments per file**.

The average number of fragments per file is a good indicator of how fragmented the files on the volume are. The following table describes the possible averages and what they indicate:

| Average fragments per file | Description |
|---|---|
| 1.00 | Most or all files are contiguous. |
| 1.10 | About ten percent of the files are in fragments of two or more pieces. |
| 1.20 | About twenty percent of the files are in fragments of two or more pieces. |
| 1.30 | About thirty percent of the files are in fragments of two or more pieces. |
| 2.00 | Most or all of the files are in fragments of two or more pieces. |

**Defragmentation Report**

You can view or print the Defragmentation Report only after you defragment a volume. The **Defragmentation Report** dialog box displays detailed information about the volume that was defragmented. The second half of the Defragmentation Report lists any files that were not defragmented.

If the volume is formatted with the NTFS file system and a file is open and writable, the defragmentation report might show small files (for example, 1 KB in size) with a large number of fragments. This is because when a writable file is open, NTFS attempts to allocate additional space to help prevent fragmentation as the file grows. Because Disk Defragmenter does not allow the extra space to be moved and consolidated while the file is open, the extra space is shown as additional fragments in the report.

# Resources for Disk Defragmenter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

- Disk Defragmenter was developed by Executive Software International in cooperation with Microsoft. For more information about Executive Software International, see the Executive Software International Web site.

- You can learn more about Disk Defragmenter in "Troubleshooting Disks and File Systems" at the Microsoft Windows Resource Kits Web site. The Microsoft Windows Resource Kits provide information technology professionals with the technical information and tools they need to successfully deploy, manage, and support Windows operating systems.

  Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Troubleshooting Disk Defragmenter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- Computer is running slowly or opening programs on the hard disk has become slow.
- The display and report do not agree.
- There appear to be unmovable files on volumes other than the system volumes and boot volumes.
- The defragmentation report on an NTFS volume shows small files with a large number of fragments.

**Computer is running slowly or opening programs on the hard disk has become slow.**
**Cause:** The volumes on your hard disk may have become excessively fragmented, possibly because you have recently created a large number of files or folders or have installed new software.

**Solution:** Analyze and then defragment the volumes on your hard disk if necessary. For more information, see Best practices for Disk Defragmenter.

**The display and report do not agree.**
**Cause:** The **Estimated disk usage before defragmentation** and **Estimated disk usage after defragmentation** displays provide a general view of a volume's fragmentation. The displays show less detail than the Analysis and Defragmentation reports, so if you compare the results you might notice discrepancies. For precise, numerical figures, always use the reports.

**Solution:** Use the **Estimated disk usage before defragmentation** and **Estimated disk usage after defragmentation** displays only for a general idea of the volume's fragmentation. Use the Analysis and Defragmentation reports for precise, numerical figures. For more information, see Analysis and defragmentation reports.

**There appear to be unmovable files on volumes other than the system volumes and boot volumes.**
**Cause:** Paging files appear as unmovable files in Disk Defragmenter. On NTFS volumes the NTFS change journal also appears as an unmovable file.

**Solution:** This is intentional because paging files cannot be moved and therefore cannot be defragmented. For more information, see Why files are not moved to some regions of NTFS-formatted volumes.

**The defragmentation report on an NTFS volume shows small files with a large number of fragments.**
**Cause:** When a file is open for write access, the NTFS file system attempts to allocate additional space to help prevent fragmentation as the file grows. When you defragment a volume, Disk Defragmenter does not allow the extra space to be moved and consolidated while the file is open. The extra space is shown as additional fragments in the report.

**Solution:** Close the file to release the additional fragments and reduce the size of the file. The **Size on disk** value (right-click the file and then click **Properties**) reports the size that you would obtain if you closed the file.

# Disk Quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk quotas**
You use disk quotas to monitor and limit disk space use on NTFS volumes.

- Before setting up disk quotas, see Checklist: Setting up disk quotas.
- For a list of best practices for disk quotas, see Best practices for Disk Quotas.
- For help with specific tasks, see Disk Quotas How To....
- For general background information, see Disk Quotas Concepts.
- For problem-solving instructions, see Troubleshooting Disk Quotas.

# Checklist: Setting up disk quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Setting up disk quotas**

| Step | Reference |
| --- | --- |
| Review key concepts. | Disk quotas overview |
| Confirm that you are a member of the Administrators group. | Default local groups |
| Enable disk quotas. | Enable disk quotas |
| Set disk quotas for an NTFS volume. | Assign default quota values |

# Best practices for Disk Quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

**Apply appropriate disk quota limits**

- Apply disk quota limits according to realistic disk space requirements for your users. Start by classifying users by the amount of disk space that you anticipate they will require (for example, users who work with scanned photographs or artwork may require a large amount of disk space). Next, structure your volumes by those classes, and then use disk quotas to limit the amount of disk space allowed for the users on each volume.

  Ensure that you have enough disk space to accommodate both existing users and future additions. Each file stored on the volume can use up to 64 kilobytes (KB) of NTFS metadata that is not applied to a user's quota limit. To avoid running out of disk space, ensure that there is sufficient disk space to accommodate this metadata.

**Set default limits**

- Set moderately restrictive default limits for all user accounts, and then modify the limits to allow more disk space to users who work with large files. It may be better to increase disk quota limits for a few user accounts than to force some users to work with a quota limit that does not meet their needs. User Policies occupy at least 2 megabytes (MB) of disk space, so you should set the default limits for all user accounts above 2 MB.

**Allow for growth**

- Set quota limits at a level that will allow for growth in disk space use. When quota increases are warranted, increase them by a percentage of the current quota setting.

**Delete users' quota entries**

- When a user no longer needs the use of a particular volume (for example, when he or she leaves the company), delete the user's quota entry for that volume and then delete, move, or take ownership of any files that the user owns. This helps you eliminate the use of volume space by unneeded files.

**Limit installation to administrators**

- If you enable disk quota settings on the same volume where the Windows operating system and other programs are installed, it is important that the user accounts that are used to install programs and Windows components have a sufficiently high disk quota. You can ensure this in one of two ways:

  - Install all programs and Windows components using the local Administrator account. This account has no quota limit.

  - Configure the disk quota to unlimited (do not limit disk usage) or to a sufficiently large amount of disk space for the user accounts that you use to install programs and Windows components.

# Disk Quotas How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Manage Disk Quotas
- Manage Disk Quota Entries

# Manage Disk Quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage disk quotas**

- View disk quota settings
- Enable disk quotas
- Disable disk quotas
- Deny disk space to users who exceed their limit
- Assign default quota values
- Manage disk quotas on a remote computer

# View disk quota settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view disk quota settings**

1. Open My Computer.
2. Right-click the disk volume for which you want to view disk quota settings, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable disk quotas
Disable disk quotas

# Enable disk quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable disk quotas**

1. Open My Computer.
2. Right-click the disk volume for which you want to enable disk quotas, and then click **Properties**.
3. In the **Properties** dialog box, click the Quota tab.
4. On the **Quota** tab, click the **Enable quota management** check box.
5. Select one or more of the following options, and then click **OK**:

**Deny disk space to users exceeding quota limit**

Users who exceed their quota limit receive an "insufficient disk space" error from Windows and cannot write additional data to the volume without first deleting or moving some existing files from it.

Individual programs determine their own error handling for this condition. To the program, it appears that the volume is full. If you clear this check box, users can exceed their quota limit. Enabling quotas and not limiting disk space use are useful when you do not want to deny users access to a volume but want to track disk space use on a per-user basis. You can also specify whether or not to log an event when users exceed either their quota warning level or their quota limit.

**Limit disk space to**

Enter the amount of disk space that new users of the volume are allowed to use and the amount of disk space that has to be used before an event is written to the system log. Administrators can view these events in Event Viewer. You can use decimal values (for example, 20.5) For the disk space and warning levels, select the appropriate units from the drop-down list (for example, KB, MB, GB). For more information about Event Viewer, see Related Topics.

**Log event when a user exceeds their quota limit**

If quotas are enabled, an event is written to the system log on the local computer whenever users exceed their quota limit. Administrators can view these events in Event Viewer, filtering for disk event types.

By default, quota events are written hourly to the system log on the local computer. You can change the interval at which quota events are written to the system log on the local computer using the **fsutil behavior** command. For more information about Event Viewer or the **fsutil behavior** command, see Related Topics.

**Log event when a user exceeds their warning level**

If quotas are enabled, an event is written to the system log on the local computer whenever users exceed their quota warning level. Administrators can view these events in Event Viewer, filtering for disk event types.

By default, quota events are written hourly to the system log on the local computer. You can change the interval at which quota events are written to the system log on the local computer using the **fsutil behavior** command. For more information about Event Viewer or the **fsutil behavior** command, see Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.

- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.

- Event Viewer builds a historical, chronological record of which users exceeded their quota warning level and quota limits, and when they exceeded them. However, it does not provide information about which users are currently over their quota warning level. To find that information, click the **Quota Entries** button.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable disk quotas
View disk quota settings
Assign default quota values
Event Viewer overview
Fsutil: behavior

# Disable disk quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable disk quotas**

1. Open My Computer.
2. Right-click the disk volume for which you want to disable disk quotas, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, clear the **Enable quota management** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable disk quotas
Disable disk quotas
View disk quota settings

# Deny disk space to users who exceed their limit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To deny disk space to users who exceed their limit**

1. Open My Computer.
2. Right-click the volume for which you want to deny disk space use, and then click **Properties**.
3. In the **Properties** dialog box, click the Quota tab.
4. On the **Quota** tab, select the **Enable quota management** check box, select the **Deny disk space to users exceeding quota limit** check box, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- When the disk quota system is configured to deny disk space to users of the volume who exceed their assigned quota limit, users cannot write additional data to the volume without first deleting or moving some existing files from the volume. The response of the system to a user exceeding the quota limit depends on each individual program. To the program, it will appear that the volume is full.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign default quota values
Modify user disk space limits and warning levels

# Assign default quota values

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To assign default quota values**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open My Computer.
2. Right-click the volume for which you want to assign default quota values, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, select the **Enable quota management** check box, and then select the **Limit disk space to** option.
5. Type numeric values for the disk space limit and warning levels, select the appropriate units from the drop-down lists, and then click **OK**.

   You can use decimal values (for example, 20.5 MB).

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- When you enable disk quotas on a volume, any users with write access to the volume who have not exceeded their quota limit can store data on the volume. The first time a user writes data to a quota-enabled volume, default values for disk space limit and warning level are automatically assigned by the quota system.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **fsutil quota modify** [*VolumePathname*] [*threshold*] [*limit*] [*username*]

| Value | Description |
|---|---|
| *VolumePathname* | The drive letter (followed by a colon). |

| | |
|---|---|
| *Threshold* | The limit, in bytes, at which warnings are issued. |
| *limit* | The maximum allowed disk usage, in bytes. |
| *username* | The domain and user name of the user whose disk usage you want to limit. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- In order to assign default quota values, the volume must be formatted with the NTFS file system.

- For more information about **fsutil**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Deny disk space to users who exceed their limit
Quotas and users
Quotas and volumes
Fsutil: quota

# Manage disk quotas on a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage disk quotas on a remote computer**

1. Right-click **My Computer**, and then click **Map Network Drive**.
2. In the **Map Network Drive** dialog box, in the **Folder** box, type the path to the volume on the remote computer for which you want to manage disk quotas, and then click **Finish**.
3. After you connect to the remote computer's volume, open **My Computer**.
4. Right-click the remote volume for which you want to manage disk quotas, and then click **Properties**.
5. In the **Properties** dialog box, click the **Quota** tab.
6. On the **Quota** tab, click **Quota Entries**.

   The **Quota Entries** window appears. Each row in the list contains information for a user of the volume.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume's **Properties** dialog box.
- To manage disk quotas on a remote computer volume, you must create a share from the remote volume's root directory and you must be a member of the Administrators group on the remote computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disk quotas overview
Local and remote implementation of disk quotas
Share a folder or drive

# Manage Disk Quota Entries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage disk quota entries**

- View disk quota information for volume users
- Add new quota entries
- Delete quota entries
- Modify user disk space limits and warning levels
- Find a quota entry in the Quota Entries window
- Sort quota entries
- Change column order in the Quota Entries window
- Create quota reports
- Import quota settings from another volume
- Export quota settings to another volume

# View disk quota information for volume users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view disk quota information for volume users**

Using the Windows interface

Using a command line

**Using the Windows interface**

1. Open My Computer.
2. Right-click the volume for which you want to view quota information, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, click **Quota Entries**.

   The Quota Entries window appears. Each row contains information for a volume user.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- When the Quota Entries window is first opened, volume quota information must be obtained from the network domain controller. This information is filled in dynamically as it is obtained, and then it is saved so that it is immediately available the next time the Quota Entries window is opened.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **fsutil quota query** [*volumepathname*]

| Value | Description |
|---|---|
| *Volumepathname* | The drive letter for the volume, for example, C: |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- In order to view quota information, the volume must be formatted with the NTFS file system.
- For more information about **fsutil**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sort quota entries
Modify user disk space limits and warning levels
Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries
Fsutil: quota

# Add new quota entries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add new quota entries**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open My Computer.
2. Right-click the volume for which you want to add new disk quota entries, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, click **Quota Entries**.
5. In the Quota Entries window, on the **Quota** menu, click **New Quota Entry**.
6. In the **Select Users** dialog box, in the **Enter the object names to select** text box, type the domain or workgroup name, followed by a back slash and the user name of the user for which you want to impose quotas. Click **OK**.
7. In the **Add New Quota Entry** dialog box, specify one of the following options, and then click **OK**:

   **Do not limit disk usage.** Tracks disk space usage without limiting disk space.

   **Limit disk space to.** Activates fields for limiting disk space and setting warning levels. Type a numeric value in the text field, and then select a disk space limit unit from the drop-down list. You can use decimal values (for example, 20.5 MB).

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- You use the **New Quota Entry** feature to set up quota limits and warning levels before the user actually writes data to the volume. This is useful when you do not want to use the default disk space limit and warning level values for a particular user. Typically, the default disk space limit and warning level values established by the volume administrator are sufficient for new users of the volume.

**Using a command line**

1. Open Command Prompt.
2. Type:

**fsutil quota modify** [*VolumePathname*] [*threshold*] [*limit*] [*domain\username*]

| Value | Description |
|---|---|
| *VolumePathname* | The drive letter (followed by a colon). |
| *threshold* | The limit, in bytes, at which warnings are issued. |
| *limit* | The maximum allowed disk usage, in bytes. |
| *domain\username* | The domain and user name of the user whose disk usage you want to limit. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- In order to add new quota entries, the volume must be formatted with the NTFS file system.

- For more information about **fsutil**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sort quota entries
Modify user disk space limits and warning levels
Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries
Fsutil: quota

# Delete quota entries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete quota entries**

1. Open My Computer.
2. Right-click the volume for which you want to delete quota entries, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, click **Quota Entries**.
5. In the Quota Entries window, click the entries for the users you want to delete, and then on the **Quota** menu, click **Delete Quota Entry**.
6. If the **Disk Quota** dialog box appears, click **Yes**, click the files or folders that you want to take action on, click one of the following buttons, and then click **Close**.

| Button | Action |
|---|---|
| **Delete** | Deletes selected files from the volume. You cannot delete folders. |
| **Take ownership** | Gives you ownership of the selected files or folders on the volume. |
| **Move** | Moves selected files to a specified volume. You cannot move folders. |

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- You can filter the files and folders list to show only files or only folders using the **Show folders only** and **Show files only** check boxes. When working with large numbers of files and folders, it is best to first delete, take ownership, or move files, and then to take ownership of folders.
- User quota entries cannot be deleted if a user still owns files or folders on the volume. Before the record can be deleted, all files owned by that user on the volume must either be deleted or moved to another volume, or you must take ownership of the files and folders.
- You are prompted to delete or move files, or take ownership of files and folders, only if the user still has files or folders on the volume at the time that you want to delete the quota entry.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Sort quota entries
Modify user disk space limits and warning levels
Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries

# Modify user disk space limits and warning levels

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify user disk space limits and warning levels**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open My Computer.
2. Right-click the volume for which you want to modify quota values, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, click **Quota Entries**.
5. In the Quota Entries window, click the entries for the users whose options you want to modify, and on the **Quota** menu, click **Properties**.
6. In the **Quota Settings** dialog box, modify any of the following options you want to change, and then click **OK**:

   **Do not limit disk usage.** Click to track disk space usage without limiting disk space.

   **Limit disk space to.** Activates fields for limiting disk space and setting warning levels. Type a numeric value in the text field, and select a disk space limit unit from the drop-down list. You can use decimal values (for example, 20.5 MB).

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **fsutil quota modify** [*VolumePathname*] [*threshold*] [*limit*] [*domain\username*]

| Value | Description |
|---|---|
| *VolumePathname* | The drive letter (followed by a colon). |

| | |
|---|---|
| *threshold* | The limit, in bytes, at which warnings are issued. |
| *limit* | The maximum allowed disk usage, in bytes. |
| *domain\username* | The domain and user name of the user whose disk usage you want to limit. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- In order to set user disk space limits and warning levels, the volume must be formatted with the NTFS file system.

- For more information about **fsutil**, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sort quota entries
Modify user disk space limits and warning levels
Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries
Fsutil: quota

# Find a quota entry in the Quota Entries window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find a quota entry in the Quota Entries window**

1. Open My Computer.
2. Right-click the volume for which you want to find quota entries, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, click **Quota Entries**.
5. In the Quota Entries window, on the **Edit** menu, click **Find**.
6. In the **Find Quota Entry** dialog box, enter the user name that you want to find as either *domain_name\user_name* or *user_name@domain_name*, and then click **OK**.

   Found entries appear highlighted in the Quota Entries window.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- You can find more than one volume user by repeatedly using the **Find Quota Entry** dialog box. Each found entry remains highlighted in the Quota Entries window.
- After you have used the **Find Quota Entry** dialog box to find one or more user quota entries, you can modify the settings for all of the found entries by clicking **Properties** on the **Quota** menu.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sort quota entries
Modify user disk space limits and warning levels
Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries

# Sort quota entries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To sort quota entries

1. Open My Computer.
2. Right-click the volume for which you want to sort quota entries, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, click **Quota Entries**.
5. In the Quota Entries window, on the **View** menu, click **Arrange Items**, and then click one of the following options:

   **By Folder**
   Sorts according to folder.

   **By User Name**
   Sorts according to the full name.

   **By Logon Name**
   Sorts according to logon name.

   **By Status**
   Sorts according to quota status.

   **By Amount Used**
   Sorts according to the amount of space the users have used on the volume.

   **By Quota Limit**
   Sorts according to disk space limits.

   **By Warning Level**
   Sorts according to warning levels.

   **By Percent Used**
   Sorts according to the percentage of disk space limit used.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- You can click a column's title in the **Quota Entries** window to sort based on the entries in that column. To reverse the order, click the column title again.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sort quota entries
Modify user disk space limits and warning levels
Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries

# Change column order in the Quota Entries window

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change column order in the Quota Entries window**

1. Open My Computer.
2. Right-click the volume for which you want to view quota entries, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, click **Quota Entries**.
5. In the Quota Entries window, click the column header you want to move, and then drag it horizontally to a new position in the list.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- When the Quota Entries window is closed, the ordering and size of the columns are saved and automatically restored the next time the window is opened.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sort quota entries
Modify user disk space limits and warning levels
Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries

# Create quota reports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create quota reports**

1. Open My Computer.
2. Right-click the volume for which you want to create a report, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, click **Quota Entries**.
5. While the **Quota Entries** window is open, open a document in the program that you will use to create the report; for example, a Microsoft Excel spreadsheet.
6. In the **Quota Entries** window, click the entries for the users you want to include in your quota report, and drag the rows to the program you are using to create the report; for example, you can drag the objects to a Microsoft Excel spreadsheet or a Microsoft Word document. The transferred data retains the same column and entry order as in the **Quota Entries** window.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- The supported Clipboard formats for reports are: Rich Text Format, Comma Separated Value, CF_UNICODETEXT, and CF_TEXT.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sort quota entries
Modify user disk space limits and warning levels
Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries

# Import quota settings from another volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import quota settings from another volume**

1. Open My Computer.
2. Right-click the volume to which you want to import quota settings, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, click **Quota Entries**.
5. In the Quota Entries window, on the **Quota** menu, click **Import**.
6. In the **Import Quota Settings** dialog box, select the name of the file that contains the quota settings you want to import, and then click **Open**.

   When you are importing quota records, a dialog box appears if imported settings will overwrite existing settings for a volume user. Specify whether you want to overwrite the existing user settings.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- You can open two instances of the Quota Entries window: one for the volume from which you are importing the quota records, and one for the volume into which you want to import quota records. Then you can drag quota records from the source volume's Quota Entries window to the destination volume's Quota Entries window.
- For entries that already exist on the destination volume, the settings transferred are the user's disk space limit and quota warning level values. If the entries do not exist, new entries are created.
- Disk quota settings are saved in a file format that is exclusive to the Disk Quotas utility and that is used only for importing and exporting quota settings from one volume to another. The **Import Quota Settings** dialog box does not show the file name extension for the settings file, even if the folder that contains it shows file name extensions for all files.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sort quota entries
Modify user disk space limits and warning levels

Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries

# Export quota settings to another volume

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export quota settings to another volume**

1. Open My Computer.
2. Right-click the volume from which you want to export quota settings, and then click **Properties**.
3. In the **Properties** dialog box, click the **Quota** tab.
4. On the **Quota** tab, make sure the **Enable quota management** check box is selected, and then click **Quota Entries**.
5. Click the user quota records you want to export.
6. On the **Quota** menu, click **Export**.
7. In the **Export Quota Settings** dialog box, specify a destination folder, type the file name for the saved settings, and then click **Save**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open My Computer, click **Start**, and then click **My Computer**.
- If the volume is not formatted with the NTFS file system, or if you are not a member of the Administrators group, the **Quota** tab is not displayed in the volume`s **properties** dialog box.
- Alternatively, you can open two instances of the **Quota Entries** window: one for the source disk volume from which you want to export quota records, and one for the destination volume into which you are exporting the quota records. Then you can drag quota records from the source volume's **Quota Entries** window to the destination volume's **Quota Entries** window.
- For entries that already exist on the destination volume, the settings transferred are the user's disk space limit and quota warning level values. If the entries do not exist, new entries are created.
- Disk quota settings are saved in a file format that is exclusive to the Disk Quotas utility and that is used only for importing and exporting quota settings from one volume to another. The **Export Quota Settings** dialog box does not reveal the file name extension for the settings file, even if the folder that contains it shows file name extensions for all files.
- You can drag a selected quota setting to a folder where an export file is automatically created.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Sort quota entries
Modify user disk space limits and warning levels
Import quota settings from another volume
Export quota settings to another volume
Delete quota entries
Create quota reports
Change column order in the Quota Entries window
Add new quota entries

# Disk Quotas Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about disk quotas:

- Disk quotas overview
- Understanding Disk Quotas
- Using Disk Quotas
- Resources for Disk Quotas

# Disk quotas overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk quotas overview**

Disk quotas track and control disk space usage for NTFS volumes. Administrators can configure Windows to:

- Prevent further disk space use and log an event when a user exceeds a specified disk space limit, that is, the amount of disk space a user is allowed to use.
- Log an event when a user exceeds a specified disk space warning level, that is, the point at which a user is nearing his or her quota limit.

When you enable disk quotas, you can set two values: the disk quota limit and the disk quota warning level. For example, you can set a user's disk quota limit to 500 megabytes (MB), and the disk quota warning level to 450 MB. In this case, the user can store no more than 500 MB of files on the volume. If the user stores more than 450 MB of files on the volume, you can configure the disk quota system to log a system event. You must be a member of the Administrators group to administer quotas on a volume. For instructions on setting disk quota values, see Assign default quota values.

You can specify that users can exceed their quota limit. Enabling quotas and not limiting disk space use are useful when you do not want to deny users access to a volume but want to track disk space use on a per-user basis. You can also specify whether or not to log an event when users exceed either their quota limit or their quota warning level.

When you enable disk quotas for a volume, volume usage is automatically tracked for all users from that point on.

You can enable quotas on local volumes, network volumes, and removable drives as long as they are formatted with the NTFS file system. In addition, network volumes must be shared from the volume's root directory and removable drives must be shared. Volumes formatted with the version of NTFS used in Windows NT 4.0 are upgraded automatically by Windows Setup.

You cannot use file compression to prevent users from exceeding their quota limits because compressed files are tracked based on their uncompressed size. For example, if you have a 50 MB file that is 40 MB after it is compressed, Windows counts the file's original 50 MB size toward the quota limit.

# Understanding Disk Quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding disk quotas**
This section covers:

- Quotas and users
- Quotas and volumes
- Quota entry updates
- Quotas and converted NTFS volumes
- Local and remote implementation of disk quotas

# Quotas and users

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quotas and users**

Because disk quotas monitor volume use by individual user, each user's utilization of disk space does not affect the disk quotas for other users of the same volume. For example, if Volume F has a quota limit of 500 megabytes (MB) and a user saves 500 megabytes (MB) worth of files to Volume F, that user cannot write additional data to the volume without first deleting or moving some existing files from it. However, each of the other users can continue to save up to 500 MB worth of files on that volume as long as there is sufficient free space.

Disk quotas are based on file ownership and are independent of the folder location of the user's files within the volume. For example, if users move their files from one folder to another on the same volume, their volume space usage does not change. However, if users copy their files to a different folder on the same volume, their volume space usage doubles. Or, if another user creates a 200 kilobyte (KB) file and you take ownership of that file, that user's disk use decreases by 200 KB and your disk use increases by 200 KB.

For instructions on enabling disk quotas, see Enable disk quotas.

# Quotas and volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quotas and volumes**

Disk quotas apply only to volumes and are independent of folder structures of volumes and their layout on physical disks.

If a volume has multiple folders, the quotas assigned to that volume apply equally to all of the folders. For example, if \\Production\QA and \\Production\Public are shared folders on Volume F, files stored by users in these folders cannot use more disk space than the quota limit set for Volume F.

If a single physical disk contains multiple volumes, and you apply quotas to each of the volumes, each volume quota applies only to its specific volume. For example, if you share two different volumes as Volume F and Volume G, quotas are tracked independently for the two volumes, even if they exist on the same physical disk.

If a volume spans multiple physical disks, the same quota for that volume applies to the entire spanned volume. For example, if Volume F has a quota limit of 500 megabytes (MB), then users cannot save more than 500 MB to Volume F, regardless of whether Volume F resides on one physical disk or whether it spans three.

For instructions on enabling disk quotas, see Enable disk quotas.

# Quota entry updates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quota entry updates**

Within the NTFS file system, volume usage information is stored by user security ID (SID), not user account name. When the Quota Entries window is opened for the first time, disk quotas must obtain user account names from either the network domain controller or Local Users and Groups, match them to the SIDs for current volume users, and populate the entries in the **Name** column with user names. These names appear in this field as they are obtained from the domain controller or Local Users and Groups. This process begins immediately the first time you view quota entries.

After the names have been obtained, they are saved to cache files in the DiskQuota folder in the user's Application Data directory (\Application Data\Microsoft\Windows NT\DiskQuota) so that they are immediately available the next time the Quota Entries window is opened. However, because these cache files might go several days without being updated by Windows, the Quota Entries window might not reflect changes you have made to the domain's user accounts list since viewing quota entries.

To obtain current user names, press F5. Windows then refreshes the user names for each user from the network domain controller or Local Users and Groups. Depending on the number of volume users and the current network speed, this process could take a few minutes. During this process, if you want to see a particular user's name more quickly, you can select that entry to move it to the top of the list.

For instructions, see View disk quota information for volume users.

# Quotas and converted NTFS volumes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quotas and converted NTFS volumes**

Disk quotas work with all NTFS volumes on computers running the Windows 2000, the Windows XP, or the Windows Server 2003 family of operating systems. However, files on volumes that have been converted from FAT or FAT32 to NTFS automatically belong to the administrator, and so quota on these files are charged against the administrator account. This is seldom a problem because administrators have unlimited volume use. This issue only applies to files that existed before conversion to NTFS; files created after conversion belong to the appropriate user.

Because disk quotas are based on file ownership, any changes to a volume that affect ownership status of the files on it, including file system conversion, might affect disk quotas for the volume. Therefore, you should understand the changes to ownership that the conversion might cause before converting existing volumes from one file system to another. For more information about file ownership issues, see Deleting disk quota entries.

For more information about file system compatibility and limitations, see Choosing a file system: NTFS, FAT, or FAT32.

# Local and remote implementation of disk quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local and remote implementation of disk quotas**

You can enable disk quotas on NTFS volumes residing on both local computers and remote computers. You can use quotas to limit the amount of volume space available to different users who log on to the local computer and limit volume usage by remote users.

You can use quotas to make sure that:

- Multiple users who log on to the same computer do not interfere with the other users' ability to work.
- Disk space on public servers is not monopolized by one or more users.
- Users do not use excessive disk space on a shared folder on your personal computer.

To enable quotas on remote computer volumes, those volumes must be shared from the volume's root directory and you must be a member of the Administrators group on the remote computer. In addition, those volumes must be formatted NTFS and exist on a computer running Windows 2000, Windows XP, or the Windows Server 2003 family of operating systems. For instructions describing how to enable quotas on remote computers, see Manage disk quotas on a remote computer.

When you implement disk quotas on a system volume, take into account the disk space used by the Windows files. Depending on the free space available on the volume, you might have to set either a high quota limit or no limit at all for the user who installed Windows. If Windows was installed by an administrator, you do not need to do this because administrators and members of the Administrators group have an unlimited quota limit.

For instructions describing how to enable quotas, see Enable disk quotas.

# Using Disk Quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using disk quotas**

This section covers:

- Enabling disk quotas
- Deleting disk quota entries
- Tracking disk space use
- Managing disk quotas from the command line

# Enabling disk quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling disk quotas**

If you are a member of the Administrators group, you can enable quotas on NTFS volumes. On volumes that already contain files, Windows calculates the disk space used by all users who have copied, saved, or taken ownership of files on the volume up to that point. The quota limit and warning level are then applied to all current users based on those calculations, and to users who begin using the volume from that point on. You can then set different quotas, or disable quotas, for individual or multiple users. You can also set quotas for specific users who have not yet copied, saved, or taken ownership of files on the volume.

For example, if \\Production\Public is shared from the root directory of an NTFS volume, you might want to set a quota limit of 50 megabytes (MB) for all users of \\Production\Public, while making sure two users who work with larger files on the server have a 100 MB limit. If both of these users already have files stored on \\Production\Public, you can select both users and set their quota limit to 100 MB. However, if one or both users do not have files stored on the server when you enable quotas, you need to select the users in the Quota Entries window and then set their quota limit to a value higher than the default for new users.

For instructions about enabling disk quotas, see Enable disk quotas. For instructions describing how to enable quotas on remote computers, see Manage disk quotas on a remote computer.

# Deleting disk quota entries

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deleting disk quota entries**

When a user no longer needs to use a volume on which quotas are enabled, you can eliminate that user's warning level and quota limit by deleting his or her entry from the Quota Entries window. However, you can do so only after all of the volume files owned by that user have either been deleted or removed from the volume, or after another user has taken ownership of the files.

Because disk quotas track files by owner, locating these files and then moving, deleting, or taking ownership of them is made simple with a dialog box that appears when you delete a quota entry. For instructions describing how to delete quota entries, see Delete quota entries.

For instructions describing how to take ownership of a file, see Take ownership of a file or folder.

# Tracking disk space use

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tracking disk space use**

You can specify a disk quota limit without denying users further disk space use. This is useful when you do not want to deny users access to a volume, but want to track disk space use on a per-user basis. You can also specify whether to log an event when users exceed either their quota warning level or their quota limit. Tracking disk space use is helpful if you want to:

- Keep historical records of disk space use.
- Show additional disk space needs.
- Show distribution of disk space on a per-user basis.
- Notify users when they exceed their allotted disk space.

For instructions on documenting disk space use, see Create quota reports.

# Managing disk quotas from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing disk quotas from the command line**
In addition to using **My Computer** to manage disk quotas (right-click a volume, click **Properties**, and then click the **Quota** tab), you can use the **fsutil** command-line utility to manage disks and volumes.

**Important**

- **Fsutil** is an advanced tool that is only intended to be used by advanced users who have a thorough understanding of Windows.

You can use:

- Fsutil: quota to create and modify disk quotas for a specific user, query disk quotas for a specific volume, search the system log for any quota violations, or to perform these tasks using scripts.

- Fsutil: behavior to change how often quota events are written to the system log.

For information about other command-line utilities, see Command-line reference A-Z.

# Resources for Disk Quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

- You can learn more about Disk Quotas in "File Systems" at the Microsoft Windows Resource Kits Web site. The Microsoft Windows Resource Kits provide information technology professionals with the technical information and tools they need to successfully deploy, manage, and support Windows operating systems.

# Troubleshooting Disk Quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The Quota tab does not appear on the Properties dialog box.
- Cannot delete a quota entry.
- A user gets an "insufficient disk space" message when trying to add files to a volume.
- Users exceeded their quota limits when Windows NT 4.0 was running.
- Quota entries appear to be outdated.
- Cannot enable quotas for a volume on a remote computer.

**The Quota tab does not appear on the Properties dialog box.**
**Cause:** Disk quotas can only be set by members of the Administrators group on volumes formatted with the NTFS file system. Therefore, the **Quota** tab does not appear if you are not a member of the Administrators group, if the volume is not shared from the volume's root directory, or if the volume is not formatted with NTFS.

**Solution:** If you want to track or limit disk space usage, ensure that you are a member of the Administrators group, that the volume is shared from the volume's root directory, and that the volume is formatted with NTFS.

**See also** Quotas and converted NTFS volumes.

**Cannot delete a quota entry.**
**Cause:** You cannot delete a quota entry for a user account until all files that the user owns have been removed from the volume or another user has taken ownership of the files.

**Solution:** Move, delete, or take ownership of all files that the user owns on that volume.

**See also** Delete quota entries.

**A user gets an "insufficient disk space" message when trying to add files to a volume.**
**Cause:** The user has exceeded the quota limit.

**Solution:** Increase the user's quota limit, move or delete files from the volume, or clear the **Deny disk space to users exceeding quota limit** check box in Disk Quotas.

**See also** Enable disk quotas.

**Users exceeded their quota limits when Windows NT 4.0 was running.**
**Cause:** Disk quotas work under the Windows 2000, the Windows XP, and the Windows server 2003 families of operating systems only. For example, if you set disk quota limits on a computer that is configured to run both Windows NT 4.0 and Windows Server 2003, Standard Edition, users can exceed their quota limit while Windows NT 4.0 is running. If this happens, the next time that Windows Server 2003, Standard Edition, is started, users who have exceeded their quota limit while Windows NT 4.0 was running cannot add files to the volume to which the quota is applied.

**Solution:** To enable users to add more files to the volume while Windows Server 2003, Standard Edition, is running, see the previous issue. To prevent users from exceeding their quota limit on a volume under Windows NT 4.0, change permissions on the volume so that users do not have Write permission.

**See also** File and folder permissions.

**Quota entries appear to be outdated.**
**Cause:** The **Enable quota management** check box on the **Quota** tab for the volume is not selected. Windows retains any pre-existing quota information but cannot update it unless quotas are enabled.

**Solution:** Enable disk quotas. For instructions describing how to enable disk quotas, see Enable disk quotas.

**Cannot enable quotas for a volume on a remote computer.**
**Solution:** Establish a connection to the remote volume and then enable quotas. For instructions on how to establish a connection to a remote volume and manage quotas on that volume, see Manage disk quotas on a remote computer. For instructions on how to enable disk quotas, see Enable disk quotas.

# Hardware

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware**
This section covers:

- Hardware overview
- Driver signing for Windows
- Hardware compatibility
- Designed for Microsoft Windows logo
- Device Manager

In addition, for information about managing printers and faxes, see Printers and Faxes.

# Hardware overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware overview**

Hardware includes any device that is connected to your computer and is controlled by your computer's microprocessor. This includes equipment that was connected to your computer when it was manufactured, as well as peripheral equipment that you added later. The following are examples of devices:

- Universal Serial bus (USB)
- IEEE 1394
- Modems
- Disk drives
- CD-ROM drives
- Digital video disc (DVD) drives
- Printers
- Network adapters
- Keyboards
- Video adapter cards
- Monitors
- Game controllers

Devices (Plug and Play and non-Plug and Play) can be connected to your computer in several ways. Some devices, such as network adapters and sound cards, are connected to expansion slots inside your computer. Other devices, such as printers and scanners, are connected to ports on the outside of your computer. Some devices, known as PC Cards, connect only to PC Card slots on a portable computer.

For a device to work properly with Windows, software known as a device driver must be installed on the computer. Each device is supported by one or more device drivers, which are typically supplied by the device manufacturer. However, some device drivers are included with Windows. If the device is Plug and Play, Windows can automatically detect it and install the appropriate device drivers.

If the device is not automatically installed by Windows, the Found New Hardware Wizard will appear and ask you to insert any media (such as compact discs or floppy disks) that were provided with the device. Non-Plug and Play devices are installed using the Add Hardware Wizard in Control Panel. If you want to manually install device drivers, you must use Device Manager. Before manually installing device drivers, you should consult the device documentation provided by the manufacturer.

**Important**

- You must be logged on as an administrator or as a member of the Administrators group in order to install or configure a device if:
- The device driver does not have the Designed for Windows Logo or a digital signature.

- Further action is required to install the device, requiring Windows to display a user interface.
- The device driver is not already on your computer.
- You need to configure a device using Device Manager.
- Your computer is connected to a network; network policy settings may also prevent you from installing hardware.

**Note**

- For the latest updates, including security fixes, service packs, new Help files, and device drivers, see Windows Update Web site.

**See Also**
**Concepts**
Hardware profiles overview
Using Internet Explorer
**Other Resources**
Hardware types overview
Installing devices overview

# Driver signing for Windows

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Driver signing for Windows**

Device drivers and operating system files included with Windows have a Microsoft digital signature. The digital signature indicates that a particular driver or file has met a certain level of testing and that it was not altered or overwritten by another program's installation process. Device drivers for hardware products that have the **Designed for Microsoft Windows XP** or **Designed for Microsoft Windows Server 2003** logos have a digital signature from Microsoft which indicates that the product was tested for compatibility with Windows and has not been altered since testing. For more information about the **Designed for Microsoft Windows** logo, see **Related Topics**.

Depending on how your administrator configured your computer, Windows either ignores device drivers that are not digitally signed, displays a warning when it detects device drivers that are not digitally signed (the default behavior), or prevents you from installing device drivers that do not have digital signatures.

Windows includes the following features to ensure that your device drivers and system files remain in their original, digitally-signed state:

- Windows File Protection
- System File Checker
- File Signature Verification

**See Also**

**Concepts**

System Properties Overview - Control Panel
System File Checker (sfc)
Designed for Microsoft Windows logo
Designed for Microsoft Windows logo

**Other Resources**

Set file signature verification options

# Hardware compatibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware compatibility**

One of the most important steps to take before running Setup on a server is to confirm that your hardware is compatible with products in the Windows Server 2003 family. You can do this by running a preinstallation compatibility check from the Setup CD or by checking the hardware compatibility information available from Microsoft. Also, as part of confirming hardware compatibility, check to see that you have obtained updated hardware device drivers and an updated system BIOS (or for an Itanium architecture-based computer, the correct firmware).

Regardless of whether you run a preinstallation compatibility check, Setup checks hardware and software compatibility at the beginning of an upgrade or new installation and displays a report if there are incompatibilities.

**Running a preinstallation compatibility check from the Setup CD**

You can run a hardware and software compatibility check from the Setup CD. The compatibility check does not require you to actually begin an upgrade or an installation. To run the check, insert the Setup CD in the CD-ROM drive and, when a display appears, follow the prompts for checking system compatibility. You will be offered the option to download the latest Setup files (through Dynamic Update) when you run the check. If you have Internet connectivity, it is recommended that you allow the download.

For more information about downloading the latest Setup files, see Using Dynamic Update for updated drivers and other setup files.

Another way to run the compatibility check is to insert the Setup CD in the CD-ROM drive, open a command prompt, and type

***d*:\i386\winnt32 /checkupgradeonly**

where *d* represents the CD-ROM drive.

**Checking hardware and software compatibility information available from Microsoft**

Another way to confirm that your hardware and software are designed for products in the Windows Server 2003 family is to click the appropriate links for compatibility information in Support resources.

Note that one type of hardware no longer supported is the microchannel bus.

**Checking drivers and system BIOS or firmware**

Check that you have obtained updated drivers for your hardware devices and that you have the latest system BIOS (for an x86-based or x64-based computer) or the correct firmware (for an Itanium architecture-based computer). The device manufacturers can help you obtain these items. For information about the Advanced Configuration and Power Interface (ACPI) standard for the BIOS, see Understanding the ACPI BIOS for an x86-based computer.

Finally, if you have devices that do not use Plug and Play, or you are aware that your Plug and Play devices are not implemented exactly to the standards, consider taking a device inventory of the hardware devices in your computer. For more information, see Taking an inventory of devices that do not use Plug and Play.

# Designed for Microsoft Windows logo

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Designed for Microsoft Windows logo**
Hardware products displaying the **Designed for Microsoft Windows XP** or **Designed for Microsoft Windows Server 2003** logos have been tested for compatibility with Microsoft Windows operating systems through use of Microsoft-provided testing procedures. For the best performance, Microsoft recommends using hardware products that display the **Designed for Microsoft Windows XP** or **Designed for Microsoft Windows Server 2003** logos on the external packaging and on the device itself.

Device drivers for hardware products with the **Designed for Microsoft Windows XP** or **Designed for Microsoft Windows Server 2003** logos have a digital signature from Microsoft, indicating that the product was tested for compatibility with Windows and has not been altered since testing. For information on accessing the latest updates, including security fixes, service packs, new Help files, and device drivers, see Support resources. You may also want to contact the hardware vendor directly and ask for a device driver that has met the **Designed for Microsoft Windows XP** or **Designed for Microsoft Windows Server 2003** logo requirements. For more information, visit the Microsoft Windows Logo Program on the Microsoft Web site.

**Caution**

- Microsoft strongly recommends that you use only device drivers with the **Designed for Microsoft Windows XP** or **Designed for Microsoft Windows Server 2003** logos. Installing device drivers that have not been digitally-signed by Microsoft might disable the system, allow viruses onto your computer, or otherwise impair the correct operation of your computer either immediately or in the future.

**See Also**
**Concepts**
Driver signing for Windows
System Properties Overview - Control Panel
**Other Resources**
Set file signature verification options

# Device Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Device Manager**
You can use Device Manager to update the drivers (or software) for hardware devices, modify hardware settings, and troubleshoot problems.

- For help with specific tasks, see Device Manager How To....
- For general background information, see Device Manager Concepts.

**Caution**

- Incorrectly altering your hardware configuration can damage your system. Be sure to read the appropriate instructions before making changes to your hardware configuration.

# Printers and Faxes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printers and faxes**
This section covers:

- Using Fax
- Printing

# Using Fax

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Fax**
Fax provides a complete fax solution for your organization. As a fax administrator you can share fax resources between multiple users, and manage network fax resources from a central location. Using Fax Service Manager you can configure fax devices, and share fax printers to enable network users to send and receive faxes through remote fax printer connections. You can set up routing policies for incoming faxes, and specify outgoing fax rules to route outbound faxes to specific device groups based on dialing destinations. You can set up archiving to provide access to faxes that have been previously sent or received, and configure activity logging to track the use of fax resources.

Fax users can send, receive, and manage faxes using a local fax device attached to their computer, or through a remote fax printer connection. Fax Console provides a central point for sending and receiving faxes, and for monitoring and managing incoming, outgoing, and archived faxes. Fax is not installed by default during Windows Setup.

For more information about Fax Service Manager, see Fax Service Manager Help. For more information about Fax Console, see Fax Console Help.

**To install Fax**

1. Open the Windows Components Wizard.

2. In the Windows Components Wizard, select the **Fax Services** check box, and then follow the instructions on your screen.

## Notes

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.

- To open Fax Console after you install Fax, click **Start**, point to **All programs**, point to **Accessories**, point to **Communications**, point to **Fax**, and then click **Fax Console**.

- Fax sharing is not supported on Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Printing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing**
The ability to share printers is a major advantage of networking with the Microsoft® Windows Server 2003 family. By putting one high-end printer on a network, you can spread its cost and usage among all users.

- Before setting up a printer, see Checklist: Setting up a printer.
- If you are familiar with Microsoft® Windows NT® version 4.0, see New ways to do familiar Printing tasks.
- For tips on printing, see Printing Best practices.
- For procedures on specific tasks, see Printing How To...
- For general background information, see Printing Concepts
- For problem-solving instructions, see Printing Troubleshooting

# Checklist: Setting up a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Setting up a printer**

| | Step | Reference |
|---|---|---|
| | If you are new to printing administration, review key concepts. | Printing Concepts |
| | If you are familiar with printing administration, read about new features. | New printing features |
| | Plan your printing environment. | Planning for Printing |
| | Verify that your printer is in the Windows Catalog. | Support resources |
| | Physically install the printer, and verify that it is ready to print. | Printer's vendor documentation |
| | Log on as a member of the Administrators group. | |
| | If you set up a printer that is connected directly to a computer (typically through a parallel port) and the printer connects through a USB or IEEE 1394 port, the printer will be detected and configured automatically. | Add a printer attached to your computer |
| | If you set up a printer that is attached directly to your network (by way of a network adapter), first install the communication protocol and network services it uses. | Add a printer attached directly to the network |
| | Share a printer with other users and advertise it (if not done while adding the printer). | Share a Printer with Users Across the Network |
| | Assign printer permissions. | Set or remove permissions for a printer |
| | Set the printer's installable options. | Set installable options for a printer |
| | Set the default printing preferences. | Set printing defaults |
| | Install additional printer drivers for other computers and operating systems to allow automatic driver distribution to clients other than Windows XP or Windows 2000. | Add printer drivers for other versions of Windows |
| | Set up network clients to use the printer. | Connect clients to a printer |
| | For Windows NT 4.0, Windows 98, and Windows 95 clients, install the appropriate printer driver on the print server. | Connect Windows NT 4.0, Windows 95, and Windows 98 clients |
| | For MS-DOS, Windows 3.*x*, and OS/2 clients, install a printer driver on the client. | Connect MS-DOS, Windows 3.x, and OS/2 clients |
| | For other clients, install the network services that the clients use. | Connect clients to a printer |
| | Set advanced printer tasks. | Use Advanced Options |
| | Enable printer location tracking. | Enable printer location tracking |
| | Set Group Policy for printers. | Set Group Policy for printers |

# New ways to do familiar Printing tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**
This table lists common tasks for printing. The user interface for performing these tasks is different in the Windows Server 2003 family than it was in Windows NT 4.0.

### Note

- There are no significant user interface differences between Windows 2000, Windows XP, and the Windows Server 2003 family for this component.

| If you want to | In Windows NT 4.0 use | In this version of Windows |
|---|---|---|
| Add a printer that uses the TCP/IP protocol | Add Printer Wizard, **Add Port**, and click **LPR port**. | Printers and Faxes, Add Printer Wizard.<br><br>For more information, see Add a printer attached directly to the network. |
| Set printer availability, priority, and spool options | Printer's property page, **Schedule** tab | Printers and Faxes, Printer's property page, **Advanced** tab.<br><br>For more information, see Set printer properties. |
| Set a printer separator page or print processor, or add a new driver | Printer's property page, **General** tab | Printers and Faxes, Printer's property page, **Advanced** tab.<br><br>For more information, see Set printer properties. |
| Start and stop the Print Spooler service | **Services** in Control Panel | Computer Management.<br><br>For more information, see Start and stop the print spooler. |

# Printing Best practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- **Planning your printing environment.**

  Consider the physical location of your users and their printing needs to determine the number and types of printers and print servers you need, and where to physically locate them.

- **Using the search for printers utility.**

  The best and fastest way to find and connect to a printer is by using the search for printers utility. From the **Start** menu click **Search**, then click **For Printers**, and then type the criteria for the printers you are searching for. For more information see, Connect to a printer published in Active Directory.

- **Installing additional printer drivers.**

  If you share printers with clients other than Windows XP or Windows 2000 clients, install the required additional printer drivers for these clients.

- **Using the same guidelines that apply to any shared resource.**

  Create a local *printer_name* Users group with Print permission, and then put global groups in the local group.

- **Removing Print permission from the default group Everyone.**

  Instead, assign the Print permission to the built-in group Users. This limits printer use to those users in the domain for which you have created accounts.

- **Distributing the administrative load.**

  If security is not an issue, assign the *printer_name* Users group the Manage Documents or Manage Printers print permission.

- **Maintaining physical security.**

  Secure the printer in a locked room if it is used for printing confidential information. When printing confidential documents, you may want to set the paper source to manual so the document will not print until the user is physically at the printer. Let only members of the Administrators group manage the printer.

- **Maintaining data security.**

  To help mitigate client-to-server eavesdropping, consider deploying IPSec. For more information, see Internet Protocol Security (IPSec). To help mitigate server-to-print device document interception, deploy an IPv4 or IPv6-compatible printer. You can use a printer that provides password authentication (that is, the printer does not print until the user enters a password that was sent with the print job). You can also print to a local printer.

- **Choosing a physical location for printing pools.**

  For printing pools, place the printers physically close to each other so users do not have to check separate locations for their printed documents.

- **Managing printer traffic.**

Create multiple logical printers with different schedules to reduce printer traffic during peak hours. Have users send large documents, such as accounting reports, to a logical printer that is available only at night so that those documents wait until off-peak hours to be printed.

- **Maintaining records.**

  Document information about printers and the users who have the ability to administer them.

- **Auditing printers.**

  Use the audit feature to keep track of changes made by administrators who manage shared printers.

- **Using network adapters.**

  It is better to use printers that connect directly to the network with a network adapter than to use a parallel port (LPT) printer that connects directly to print servers. A printer attached with a parallel port prints slower, and the parallel port requires considerably more CPU time.

- **Using the printer troubleshooter.**

  To solve problems quickly, use the printing troubleshooter. Also, you can use the networking (TCP/IP) troubleshooter for networking problems.

- **Creating a custom folder for printers.**

  To ease printer administration, create a custom folder for the printers and for the folders of any print server you manage regularly. You can also create a shortcut for that folder on the Desktop.

# Printing How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Set Up a Printer
- Use Advanced Options
- Manage Printers
- Manage and Set Print Servers
- Connect clients to a printer
- Manage Documents

# Set Up a Printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set up a printer**

- Add a printer attached to your computer
- Add a printer attached directly to the network
- Set or remove permissions for a printer
- Specify your default printer
- Print a test page
- Install new or updated printer drivers
- Add printer drivers for other versions of Windows
- Set printing defaults
- Set personal printing preferences
- Choose a separator page
- Set printer properties
- Share a Printer with Users Across the Network
- Set Printing Preferences
- Configure Printer Options

# Add a printer attached to your computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a printer attached to your computer**

1. Connect the printer to the appropriate port on your computer according to the printer manufacturer's documentation, and verify that it is ready to print.

2. Although Windows automatically detects and installs most printers, you might need to provide additional information to complete the installation. Choose from the following, depending on the type of printer you have.

   Install a parallel port (LPT) attached printer

   Install a USB or IEEE 1394 printer

   Install an infrared printer

**Install a parallel port (LPT) attached printer**
If your printer attaches to your computer with a parallel port (LPT port), connect the printer to your computer, and then open Printers and Faxes. Double-click **Add Printer** to start the Add Printer Wizard, and then click **Next**. Click **Local printer attached to this computer**, make sure the **Automatically detect and install my Plug and Play printer** check box is selected, and then click **Next**. Depending on the printer you are installing, a Found New Hardware message or the Found New Hardware Wizard appears to notify you that the printer has been detected and that installation has begun. Follow the instructions on the screen to complete the printer installation. (You do not need to continue with the following instructions.)

Alternatively, after connecting your printer, you can start or restart your computer to allow your operating system to automatically detect and start the Found New Hardware Wizard.

The printer icon will be added to your Printers and Faxes folder.

**Install a USB or IEEE 1394 printer**
If you install a universal serial bus (USB) or IEEE 1394 printer, your Windows Server 2003 family operating system detects it and automatically starts the Found New Hardware Wizard. You do not need to shut down or restart your computer; just follow the instructions on the screen to finish the setup.

The printer icon will be added to your Printers and Faxes folder.

**Install an infrared printer**
If you install an infrared printer, it is automatically detected and installed. You do not need to shut down or restart your computer. First, make sure that the infrared-enabled computer and the printer are turned on. Then, position the computer and printer within approximately 1 meter of each other to establish a wireless infrared connection. After several seconds, the computer recognizes the printer, and the appropriate drivers are installed on the computer. If the drivers are not installed, you will be prompted to insert a manufacturer's disk containing the drivers. (You do not need to continue with the rest of the steps in this topic.)

The printer icon will be added to your Printers and Faxes folder.

If your computer is not infrared-capable, you can attach an infrared transceiver to your serial (COM) port. For more information, see Related Topics.

1. If you could not install your printer using Plug and Play, or if the printer is attached to your computer with a serial (COM) port, then open Printers and Faxes.

2. Double-click **Add Printer** to start the Add Printer Wizard, and then click **Next**.

3. Clear the **Automatically detect and install my Plug-n-Play printer** check box to avoid having to wait for the completion of another printer search.

4. Click **Local printer attached to this computer**, and then click **Next**.

5. Follow the instructions on the screen to finish setting up the printer by selecting a printer port, selecting the manufacturer and model of your printer, and typing a name for your printer.

   In Windows Server 2003 family products, the Add Printer Wizard shares the printer and publishes it in Active Directory by default, unless you select **Do not share this printer** in the wizard's **Printer Sharing** screen. The Add Printer Wizard does not share the printer automatically on a computer running Windows XP Professional. You need to select **Share as** to share and publish the printer.

### Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

- Infrared functionality is not present on the following operating systems: Windows XP 64-bit Edition (Itanium); 64-bit version of Windows Server 2003, Datacenter Edition; and 64-bit version of Windows Server 2003, Enterprise Edition.

- If you add and set up a Plug-and-Play printer (for example, USB, IEEE 1394, LPT, or Infrared), you do not need to have administrative credentials. However, to add and set up a non-Plug-and-Play printer connected directly to your computer, you must be logged on as an administrator or a member of the Administrators group. If your computer is connected to a network, network policy settings may also prevent you from completing this procedure.

- If you intend to share the printer with client computers running operating systems other than Windows XP or Windows 2000, you need to install the appropriate printer drivers for these clients on the print server. When clients running Windows NT 4.0, Windows 95, and Windows 98 connect to the printer, the server automatically downloads the correct driver to the client.

- When you are adding a new printer that is connected to a computer and the Add Printer Wizard prompts you to select the printer port, you normally select from the **Existing** list one of the parallel (LPT) ports. For some plotters you might need to select one of the serial (COM) ports.

- The following Group Policy settings can change the default behavior of the Windows Server 2003 family Server Add Printer Wizard:

  **Allow printers to be published** is enabled by default; you can disable it to prevent printers from being published.

  **Automatically publish new printers in the Active Directory** is enabled by default; you can disable it to prevent the Add Printer Wizard from automatically publishing printers when adding a new printer.

  **Browse the network to find printers** is enabled by default; you can disable it to prevent the Add Printer Wizard from browsing the network for shared printers.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Share your printer
Publish a printer in Active Directory
Add a printer attached directly to the network
Add printer drivers for other versions of Windows
Establishing a naming convention for printers
Set Group Policy for printers
Add a new infrared device to a serial port

# Add a printer attached directly to the network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a printer attached directly to the network**

1. Open Printers and Faxes.
2. Double-click **Add Printer** to start the Add Printer Wizard, and then click **Next**.
3. Click **Local printer attached to this computer**, clear the **Automatically detect and install my Plug and Play printer** check box, and then click **Next**.
4. Follow the instructions on the screen to finish setting up the printer by selecting a printer port, selecting the manufacturer and model of your printer, and typing a name for your printer.
5. When the Add Printer Wizard prompts you to select the printer port, click **Create a new port**.
6. From the list, click the appropriate port type and follow the instructions. (By default, only **Local Port** and **Standard TCP/IP Port** appear in the list.)

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- To add a new port type, see Related Topics.
- To add and set up a printer connected directly to the network, you must be logged on as a member of either the Power Users group or the Administrators group.
- The Add Printer Wizard shares the printer and publishes it in Active Directory by default, unless you select **Do not share this printer** in the **Printer Sharing** screen. In Windows 2000 Professional, the Add Printer Wizard does not share the printer automatically; you must select **Share as** to share and publish the printer.
- You must clear the **Automatically detect and install my Plug and Play printer** check box because the printer is attached directly to the network, not to a Windows 2000 computer.
- If you intend to share the printer with clients other than those running Windows XP or Windows 2000, you need to install the appropriate printer drivers for these clients on the print server. When clients running Windows NT 4.0, Windows 95, and Windows 98 connect to the printer, the server automatically downloads the correct driver to the client.
- The following Group Policy settings can change the default behavior of the Add Printer Wizard:

  **Allow printers to be published** is enabled by default; you can disable it to prevent printers from being published.

  **Automatically publish new printers in the Active Directory** is enabled by default; you can disable it to prevent the Add Printer Wizard from automatically publishing printers when adding a new printer.

  **Browse the network to find printers** is enabled by default; you can disable it to prevent the Add Printer Wizard from browsing the network for shared printers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set Ports and Protocols
Share your printer
Publish a printer in Active Directory
Add printer drivers for other versions of Windows
Establishing a naming convention for printers
Set Group Policy for printers

# Set or remove permissions for a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set or remove permissions for a printer**

1. Open Printers and Faxes.
2. Right-click the printer for which you want to set permissions, click **Properties**, and then click the **Security** tab.
3. Do one of the following:
   - To change or remove permissions from an existing user or group, click the name of the user or group.
   - To set up permissions for a new user or group, click **Add**. In **Select Users, Computers, or Groups**, type the name of the user or group you want to set permissions for, and then click **OK** to close the dialog box.
4. In **Permissions**, click **Allow** or **Deny** for each permission you want to allow or deny, if necessary. Or, to remove the user or group from the permissions list, click **Remove**.

## Notes

- To change device settings, you must have the Manage Printers permission. For information about printing security permissions, see Related Topics.
- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- To view or change the underlying permissions that make up Print, Manage Printers, and Manage Documents, click the **Advanced** button.
- A printer must be shared in order for the permission settings to affect the users and groups listed.
- You can also view the permissions assigned to you by clicking the group you belong to on the **Security** tab. For information on finding out what group you belong to, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Share your printer
Identify members of a local group
Assigning printer permissions

# Specify your default printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify your default printer**

The default printer is used when you click **Print** on the **File** menu of many Windows-based programs.

1. Open Printers and Faxes.

2. Right-click the printer you want to use as the default printer, and then click **Set as Default Printer**.

   A check mark appears next to the printer icon in **Printers and Faxes**.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

- For more information on administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Connect to a printer on a network
Share your printer
Add a printer attached to your computer
Prnmngr.vbs

# Print a test page

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print a test page**

To print a test page, you must have the Print permission, which is one of several printer permissions. You can also print a test page when you first install a printer. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the printer you want to test, and then click **Properties**.
3. On the **General** tab, click **Print Test Page**.
4. Click **OK** if the test page printed correctly. If the test page did not print correctly, click **Troubleshoot** to open the Print Troubleshooter for help resolving the printing problem.

**Notes**

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open printer properties from the **Printers and Faxes** folder by clicking a printer, and then clicking **Set printer properties** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- Another way to open the Printing Troubleshooter is to click **Troubleshoot printing** under **See Also** on the left side of Printers and Faxes. This option is available only if folders are set to look like a Web page, and if a printer is not currently selected. For more information, click **Related Topics**.
- For more information on administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Set or remove permissions for a printer
Choosing a printer name
Assigning printer permissions
Prnqctl.vbs

# Install new or updated printer drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install new or updated printer drivers**

1. Open Printers and Faxes.
2. Right-click the printer for which you want to change drivers, and then click **Properties**.
3. On the **Advanced** tab, click **New Driver** to install a new or updated version of the printer driver using the Add Printer Driver Wizard.
4. Click **Next**, and then do one of the following:
   - Click the appropriate printer manufacturer and printer model if the new or updated driver is on the list.
   - Click **Have Disk** if the printer driver is not included in the list or if you have received a new or updated driver on CD or diskette from the printer manufacturer. Type the path where the driver is located, and then click **OK**.
5. Click **Next**, and then follow the instructions on the screen to finish installing the printer driver.

## Notes

- To install new or updated printer drivers, you must be logged on to your computer as a member of the Administrators group. When using Windows 2000 Professional, you may be able to install new or updated printer drivers when logged on as a member of the Power Users group, depending on the components required by the printer driver.
- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- If the printer driver you want to use already exists on the print server, you can install it by selecting it in the **Driver** list.
- For information about updating printer drivers and other system files using Windows Update, see Related Topics.
- For more information about administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assigning printer permissions
Add printer drivers for other versions of Windows
Printer drivers overview
Prndrvr.vbs

# Add printer drivers for other versions of Windows

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add printer drivers for other versions of Windows**
The following procedure describes how to install additional drivers on the print server, as a distribution point for clients running older versions of Windows.

If you share a printer with users running other versions of Windows (Windows 95, Windows 98, or Windows NT 4.0), you can install additional printer drivers on your computer so those users can connect to your printer without being prompted to install the drivers missing from their systems. The drivers are located on the Support CD. Printer drivers for Windows NT 3.1 and Windows NT 3.5 are not included.

1. Open Printers and Faxes.
2. Right-click the printer for which you want to install the additional drivers, and then click **Properties**.
3. On the **Sharing** tab, click **Additional Drivers**.
4. In the **Additional Drivers** dialog box, select the check boxes for the additional environments and operating systems you need, and then click **OK**.

### Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open printer properties in Printers and Faxes by clicking a printer, and then clicking **Set printer properties** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- A client running Windows NT 4.0, Windows 2000, or Windows XP Home Edition checks the printer driver each time it reconnects to the network. A Windows NT 3.*x* client checks each time the spooler service on the client is started. If the driver is not current, a copy of the new driver is downloaded automatically. The printer driver for a Windows 95-based or Windows 98-based client is not automatically kept current. If you update the driver on the print server, you must manually install the driver on Windows 95-based and Windows 98-based clients.
- For more information on administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Share your printer
Install new or updated printer drivers
Prndrvr.vbs

# Set printing defaults

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set printing defaults**

You can change a printer's default printing preferences only if you have the Manage Printers permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the printer for which you want to set the default printing preferences, and then click **Properties**.
3. On the **Advanced** tab, click **Printing Defaults**.
4. Click the **Layout** and **Paper/Quality** tabs and the **Advanced** button to see the available choices, and specify the new default settings.

**Notes**

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- The preferences you set become the default settings for any user who connects to the printer. These preferences include the page orientation, the tray from which paper is fed to the printer, and the number of copies printed. Individual users can customize these settings by right-clicking the printer in Printers and Faxes, and then clicking **Printing Preferences**.
- You can also open printer properties in Printers and Faxes by clicking a printer, and then clicking **Set printer properties** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Set printer properties
Set personal printing preferences
Assigning printer permissions

# Set personal printing preferences

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set personal printing preferences**

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **Printing Preferences**.
3. Click the **Layout** and **Paper/Quality** tabs and the **Advanced** button to see the available choices, and make changes as needed.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- When you set personal printing preferences using Printers and Faxes, the default settings are changed for all documents you print to that printer. These preferences include the page orientation, the tray from which paper is sent, and the number of copies printed. Making these changes requires that you have the Print permission, which is one of several printer permissions. For more information, click **Related Topics**.
- You can change printing preferences for individual documents from within a software program. To do so, click **Page Setup** (or **Print Setup**) on the program's **File** menu, and then make selections. To set the printing default for all users connected to the printer, see Set printing defaults.
- You can also open printing preferences in Printers and Faxes by clicking a printer, and then clicking **Select printing preferences** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set printing defaults
Specify paper size
Set the number of copies to print
Set printing resolution
To print TrueType fonts as graphics
Print on one or both sides of the paper
Set page orientation
Assigning printer permissions

# Choose a separator page

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To choose a separator page**

To set up separator pages, you must have the Manage Printers permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **properties**.
3. On the **Advanced** tab, click **Separator Page**.
4. Click **Browse** and find your systemroot\System32 folder.
5. Specify one of the separator page files that Windows provides in the System32 folder:
   - Pcl.sep, which prints a page after switching the printer to PCL printing.
   - Sysprint.sep, which prints a page after switching the printer to PostScript printing.
   - Pscript.sep, which does not print a page after switching the printer to Postscript printing.

### Note

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Using separator pages

Assigning printer permissions

# Set printer properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set printer properties**

Printer properties cover many categories of options, including those related to printer drivers, printer ports, printer sharing, setting printer permissions for other users, deciding when the printer is available on a network, printing a test page, and so on.

Changing the printer properties affects only the selected printer, and can only be made if you have the Manage Printers permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the printer for which you want to change settings, and then click **Properties**.
3. Click each tab to see the available choices, and change options as needed.

### Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open printer properties in Printers and Faxes by clicking a printer, and then clicking **Set printer properties** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- For more information on administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Set printing defaults
To set print server properties
Share your printer
To change the parallel port time-out settings
Install new or updated printer drivers
Set or remove permissions for a printer
Assigning printer permissions
Prncnfg.vbs

# Share a Printer with Users Across the Network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Share a printer with users across the network**

- Share your printer
- Stop sharing your printer
- Publish a printer in Active Directory

# Share your printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To share your printer**

1. Open Printers and Faxes.
2. Right-click the printer you want to share, and then click **Sharing**.
3. On the **Sharing** tab, click **Share this printer**, and then type a name for the shared printer.

   If you share the printer with users on different hardware or different operating systems, click **Additional Drivers**. Click the environment and operating system for the other computers, and then click **OK** to install the additional drivers.

   If you are logged on to a Windows 2000 or Windows Server 2003 family domain, you can make the printer available to other users on the domain by clicking **List in the Directory** to publish the printer in the Directory.
4. Click **OK**, or if you have installed additional drivers, click **Close**.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.
- Printers are not shared by default when you install them on computers running Windows XP Professional, but you can choose to share any printer you install on your computer. (On any Windows Server 2003 family product, the printer is shared by default when you add the printer.)
- When you publish a printer in Active Directory, other users logged onto the domain will be able to search for the printer based on its location and features such as how many pages it prints per minute and whether color printing is supported.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set or remove permissions for a printer
Choosing a printer name

# Stop sharing your printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop sharing your printer**

1. Open Printers and Faxes.
2. Right-click the printer you want to stop sharing, and then click **Sharing**.
3. On the **Sharing** tab, click **Do not share this printer**.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Share your printer
Connect to a printer on a network

# Publish a printer in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To publish a printer in Active Directory**

1. Open Printers and Faxes.
2. Right-click the printer you want to publish, and then click **Sharing**.
3. On the **Sharing** tab, click **Share this printer**, and then type a name for the shared printer.
4. Select the **List in the Directory** check box to publish the printer in Active Directory.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- Only a shared printer can be published. If you stop sharing a printer, the printer is no longer published.
- If you add a printer using the Add Printer Wizard and click **Share this printer**, the printer will be published by default if the Group Policy settings **Automatically publish new printers in Active Directory** and **Allow printers to be published** are enabled, as they are by default.
- You must open Active Directory Users and Groups to publish printers that are managed through print servers that are running Windows NT 4.0 or an earlier operating system.
- Printers you add using the Add Printer Wizard are shared by default if you add them from a computer running a product in the Windows Server 2003 family. However, printers are not shared by default if you add them using the Add Printer Wizard from a computer running Windows XP.
- You must have the Manage Printers permission for the printer you want to share or to publish.
- You can create printer names that contain spaces, special characters, and more than eight characters. However, some clients do not recognize or correctly handle such names. The entire qualified name of a shared printer (for example \\PRINTER2\PSCRIPT) must contain 32 or fewer characters.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manually publish a printer in Active Directory
Set Group Policy for printers

# Set Printing Preferences

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set printing preferences**

- Set page orientation
- Specify paper size
- Set printing resolution
- Print on one or both sides of the paper
- Print multiple pages per sheet
- Set the number of copies to print

# Set page orientation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set page orientation**

To change printing preferences, you must have the Print permission, which is one of several printer permissions. For more information, click **Related Topics**.

Any changes you make in **Printing Preferences** will alter the default document settings for the printer and will affect all documents you send to the printer. To change print settings for individual documents from within a program, click **Page Setup** (or **Print Setup**) on the program`s **File** menu, and then make selections.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **Printing Preferences**.
3. On the **Layout** tab, under **Orientation**, click one of the following options:
   - **Portrait** to print vertically on the page.
   - **Landscape** to print horizontally on the page.
   - **Rotated Landscape** to rotate the print 90 degrees counterclockwise on the page.

   A representation of your selection will appear in the display area of the dialog box.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open printing preferences in Printers and Faxes by clicking a printer, and then clicking **Select printing preferences** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- The **Orientation** options appear in the dialog box only if your printer supports this feature. The specific options available depend on the printer you are using.
- To set the printing default for all users connected to the printer, see Set printing defaults.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set printing defaults
Assigning printer permissions

# Specify paper size

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify paper size**

To change printing preferences, you must have the Print permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **Printing Preferences**.
3. Click **Advanced**, and under **Paper/Output** click **Paper Size**.
4. Click the paper size (form) you want.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open printing preferences in Printers and Faxes by clicking a printer, and then clicking **Select printing preferences** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- Any changes you make in **Printing Preferences** will alter the default document settings for the printer and will affect all documents you send to the printer. To change print settings for individual documents from within a program, click **Page Setup** (or **Print Setup**) on the program`s **File** menu, and then make selections.
- To set the printing default for all users connected to the printer, see Set printing defaults.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Match a form to a printer tray
Set printing defaults
Assigning printer permissions

# Set printing resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set printing resolution**

To create a document, your printer applies tiny dots of ink to the paper. The more dots per inch (dpi) your printer sends, the crisper and better the results, or the *higher its resolution*. For most printers, selecting a higher printer resolution increases the time it takes for the file to print.

Printer resolution settings are not available if your printer does not support the ability to change the resolution. Also, you must have the Print permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **Printing Preferences**.
3. Click **Advanced**, and under **Graphic**, click **Print Quality**.
4. Click the resolution option you want.

**Notes**

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- In Windows, printer resolution is specified using two numbers. The first represents horizontal resolution and the second represents vertical resolution. For example, 240 X 144 specifies a resolution of 240 dots per inch across (horizontally) and 144 dots per inch down (vertically). Changing the resolution does not change the quality of the printed text unless you are using a LaserJet printer and printing TrueType fonts as graphics.
- Any changes you make in **Printing Preferences** will alter the default document settings for the printer and will affect all documents you send to the printer. To change print settings for individual documents from within a program, click **Page Setup** (or **Print Setup**) on the program`s **File** menu, and then make selections.
- To set the printing default for all users connected to the printer, see Set printing defaults.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
To print TrueType fonts as graphics
Set printing defaults
Assigning printer permissions

# Print on one or both sides of the paper

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print on one or both sides of the paper**

To change printing preferences, you must have the Print permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **Printing Preferences**.
3. On the **Layout** tab, under **Print on Both Sides**, click one of the following options:
   - **None** to print on only one side of the paper.
   - **Flip on Long Edge** to print on both sides of the paper with pages that are bound on the left edge.
   - **Flip on Short Edge** to print on both sides of the paper with pages that are bound on the top edge.

   A representation of your selection will appear in the display area of the dialog box.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- The options available on the **Layout** tab depend on the printer you are using. For example, the **Print on Both Sides** options appear only if your printer supports this feature.
- You can also open printing preferences in Printers and Faxes by clicking a printer, and then clicking **Select printing preferences** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- For information about setting the printing default for all users connected to the printer, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set printing defaults
Assigning printer permissions

# Print multiple pages per sheet

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print multiple pages per sheet**

To print multiple pages on one sheet of paper, your printer must support this feature, and you must have the Print permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **Printing Preferences**.
3. On the **Layout** tab, click the number of pages you want to print on a sheet in **Pages Per Sheet**.

   A representation of your selection will appear in the display area of the dialog box.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- Any changes you make in **Printing Preferences** will alter the default document settings for the printer and will affect all documents you send to the printer. To change print settings for individual documents from within a program, click **Page Setup** (or **Print Setup**) on the program`s **File** menu, and then make selections.
- For information about setting the printing default for all users connected to the printer, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Set printing defaults

Assigning printer permissions

# Set the number of copies to print

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the number of copies to print**
To change printing preferences, you must have the Print permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **Printing Preferences**.
3. Click **Advanced**.
4. Under **Paper/Output**, specify the **Copy Count**.

   If your printer supports collation, the pages will be collated by default when you print multiple copies.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open printing preferences in Printers and Faxes by clicking a printer, and then clicking **Select printing preferences** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- Any changes you make in **Printing Preferences** will alter the default document settings for the printer and will affect all documents you send to the printer. To change print settings for individual documents from within a program, click **Page Setup** (or **Print Setup**) on the program`s **File** menu, and then make selections.
- For information about setting the printing default for all users connected to the printer, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set printing defaults
Assigning printer permissions

# Configure Printer Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure printer options**

- Set installable options for a printer
- Set printer memory
- Match a form to a printer tray
- Ensure that complex pages print properly
- Install or remove external PCL soft fonts
- Make cartridge fonts available

# Set installable options for a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set installable options for a printer**

1. Open Printers and Faxes.
2. Right-click the printer you want to set options for, and then click **Properties**.
3. On the **Device Settings** tab, under **Installable Options**, click any option to set it.

## Notes

- To change device settings, you must have the Manage Printers permission. For information about printing security permissions, see Related Topics.
- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- Only installable options supported by the specific printer will be shown in the **Installable Options** list.
- Administrators need to set the installable options that actually exist for the specific printer. Users will only be able to see and choose from the installable options the administrator sets as available.

  Examples of installable options include: printing on both sides of a page (duplex unit), envelope feeder, multiple bin mailbox trays, stapler, internal disk, and others.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set printer memory
Assigning printer permissions

# Set printer memory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set printer memory**

To change the amount of printer memory in Windows, your printer must support the ability to change the amount of printer memory. Otherwise, the options described here will not appear in Windows. Also, you must have the Manage Printers permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. If you do not know already, find out how much memory is installed in your printer. Many printers - especially laser printers - have a self-test feature that will print a document containing the current memory configuration. For more information on how to determine your printer's memory amount, see your printer's operating manual.

2. Open Printers and Faxes.

3. Right-click the icon for the printer you are using, and then click **properties**.

4. On the **Device Settings** tab, under **Installable Options**, click **Printer Memory**, and then click the number that matches the amount of memory installed in your printer or plotter.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

- Your printer has memory chips for storing information that is sent to it. When you install a printer, Windows reads the amount of printer memory so the printer and Windows work properly together. When you install additional memory into your printer, you need to change the printer memory setting in Windows so that the two match. If Windows is configured to require more or less memory than actually resides in your printer, Windows might display a message.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Using plotters
Assigning printer permissions

# Match a form to a printer tray

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To match a form to a printer tray**

Many printers have upper and lower trays that contain different sizes of paper. Matching a form to a source ensures that documents always print on the correct paper size when the form is selected from within a program. To make these changes, you must have the Manage Printers permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **properties**.
3. On the **Device Settings** tab, under **Form To Tray Assignment**, click a printer tray.
4. Click one of the available paper sizes for the selected tray.
5. Repeat steps 3 and 4 for each paper source.

### Note

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Specify paper size

Assigning printer permissions

# Ensure that complex pages print properly

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To ensure that complex pages print properly**

To change device settings, you must have the Manage Printers permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **properties**.
3. On the **Device Settings** tab, click **Page Protection**, and then click **On**.

   When page protection is active, the printer creates the entire printed page in memory before starting to print. To use page protection, your printer must have at least 1 megabyte (MB) of optional memory available. If your printer does not support page protection, this option is not available.

### Note

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Match a form to a printer tray
Set printer memory
Assigning printer permissions

# Install or remove external PCL soft fonts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install or remove external PCL soft fonts**

This procedure enables you to install PCL format fonts on Hewlett-Packard LaserJet, DeskJet, or compatible printers. External fonts can be PCL soft fonts or cartridge fonts. Refer to your font documentation to check the format of your fonts.

You must also have the Manage Printers permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the icon for the printer you are using, and then click **properties**.
3. On the **Device Settings** tab, click **External Fonts**, and then click **Properties**.
4. Do one of the following:
   - To install external PCL soft fonts, in **New Soft Font Directory**, type the path to the folder or drive where the fonts are located, and then click **Open**. Click the fonts you want to install in **New Soft Fonts**, and then click **Add** after each selection.
   - To remove external PCL soft fonts, in **Installed Soft Fonts**, click the fonts you want to remove, and then click **Delete**.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- If you need to add or remove TrueType fonts, use the Fonts folder in Control Panel.
- You can also open printer properties in Printers and Faxes by clicking a printer, and then clicking **Set printer properties** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Printer fonts overview
Make cartridge fonts available
Assigning printer permissions

# Make cartridge fonts available

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make cartridge fonts available**

1. Make sure that the cartridges are properly inserted into the cartridge slots on your printer.
2. Open Printers and Faxes.
3. Right-click the icon for the printer you are using, and then click **properties**.
4. On the **Device Settings** tab, under **Installed Font Cartridges**, click the slot number in which the cartridge is installed.

   The **Installed Font Cartridges** options appear only if your printer supports this feature. To change this option, you must have the Manage Printers permission, which is one of several printer permissions.

### Note

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Printer fonts overview
Assigning printer permissions

# Use Advanced Options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use advanced options**

- Schedule alternate printing times
- Set different print priority to different groups
- Enable Postscript and PCL on the same printer
- Keep a copy of everything you print
- Set up a printer to print to multiple devices
- Add a color profile to a printer
- Enable printer location tracking
- Start and stop the print spooler
- Run printing commands from the command prompt

# Schedule alternate printing times

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To schedule alternate printing times**

1. Open Printers and Faxes.
2. Right-click the printer you want to set, and then click **Properties**.
3. Click the **Advanced** tab, and then click **Available from**.
4. To set the time period that the printer will be available, do one of the following:
   - Click the up or down arrows.
   - Or, type a start and end time, such as **6:00 PM To 6:00 AM**.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- By default, printers are set to be always available.
- For more information on using different printing times for a printer, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Scheduling alternate printing times
Set different print priority to different groups
Enable Postscript and PCL on the same printer

# Set different print priority to different groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set different print priority to different groups**

1. Open Printers and Faxes.
2. Right-click the printer you want to set, click **Properties**, and then click the **Advanced** tab.
3. In **Priority**, click the up or down arrows, and then click **OK**.

   Or, type a priority level, where 1 is the lowest level and 99 is the highest, and then click **OK**.
4. Click **Add Printer** to add a second logical printer for the same physical printer. For instructions, see Related Topics.
5. Click the **Advanced** tab.
6. In **Priority**, set a priority higher than that of the first logical printer.
7. Instruct the regular group of users to use the first logical printer name and the group with higher priority to use the second logical printer name. Set the appropriate permissions for the different groups.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You gain no benefits just from setting a priority to a printer. You need to set at least two different logical printers for the same physical printer to take advantage of this option.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a printer attached to your computer
Add a printer attached directly to the network

# Enable Postscript and PCL on the same printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Postscript and PCL on the same printer**

1. Open Printers and Faxes.

2. Double-click **Add Printer**, and then click **Next**. For instructions on how to add a printer, see Related Topics.

3. When the Add Printer Wizard prompts you to select the printer manufacturer and model, from the **Manufacturers** list, click the correct name; from the **printers** list, click the appropriate model, with either Postscript (PS) or PCL.

   For example, for HP LaserJet 8100 Series, click **HP LaserJet 8100 Series PS** to set the printer with a Postscript driver.

   Follow the instructions on the screen to finish setting up the printer with the Add Printer Wizard.

4. Double-click **Add Printer** again to add a second logical printer for the same physical printer. For instructions on how to add a printer, see Related Topics.

5. When the Add Printer Wizard prompts you to select the printer manufacturer and model, select the same manufacturer name again. For model, from the **Printers** list, click the model for the other printer's language.

   For the same HP LaserJet 8100 Series as in the previous example, click **HP LaserJet 8100 Series PCL 5e** to set this logical printer with a PCL driver.

6. Follow the instructions on the screen to finish setting up the printer.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

- Users with Postscript documents will send their documents to be printed to the Postscript logical printer (the first printer in the previous example); users with Printer Control Language (PCL) documents will use the PCL logical printer (the second printer in the previous example).

- Use the same procedure to allow printing to any printer that supports multiple languages.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a printer attached to your computer
Add a printer attached directly to the network

# Keep a copy of everything you print

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To keep a copy of everything you print**

1. Open Printers and Faxes.
2. Right-click the printer you want to set, click **Properties**, and then click the **Advanced** tab.
3. Select the **Keep printed documents** check box.
4. If you have compatibility problems with your printing, clear the **Enable advanced printing features** check box.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- If you keep printed documents in the spool, carefully monitor the spool folder and the available disk space.
- To purge all printed documents from the spool, clear the **Keep printed documents** check box (default).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Set up a printer to print to multiple devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up a printer to print to multiple devices**

1. Open Printers and Faxes.
2. Right-click the printer you are using, and then click **Properties**.
3. On the **Ports** tab, select the **Enable printer pooling** check box.
4. Click each port where the printers you want to pool are connected.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- With printer pooling, the printers must be the same type of printer using the same printer driver.
- The previous procedure assumes the printers you want to pool are already in the Printers and Faxes folder.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Add a color profile to a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a color profile to a printer**

1. Open Printers and Faxes.
2. Right-click the printer that you want to associate with a color profile, click **Properties**, and then click the **Color Management** tab.
3. Click **Add** to open the **Add Profile Association** dialog box.
4. Locate the new color profile you want to associate with the printer.
5. Click the new profile, and then click **Add**.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- Color profiles communicate the color characteristics of a device to the color management system. Associating the correct color profile with all of your publishing tools helps to ensure consistent color application throughout the publishing process.
- The Color Management tab is only visible for color printers that support color management.
- For more information about a profile in the **Add Profile Association** dialog box, right-click the profile, and then click **Properties**.
- Color profiles are usually installed in the systemroot\System32\Spool\Drivers\Color folder.
- Windows provides a color profile named sRGB Color Space Profile.icm, which you can use with any device that supports Image Color Management 2.0.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the rendering intent for viewing or printing an image or page

# Enable printer location tracking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable printer location tracking**

1. Open Active Directory Sites and Services.
2. Create the sites you want, if you haven't done so already. To create a site, see Related Topics.
3. Associate a subnet for this newly created site, if you haven't done so already. To associate a subnet for the site, see Related Topics.
4. Right-click each subnet you created, and then click **Properties**.
5. Click the **Location** tab, enter the location of the subnet in the following format based on your location naming convention, and then click **OK**:

   *location/location*

6. Using Group Policy, double-click the **Pre-populate printer search location** policy setting, and then click to enable it.

   The following example demonstrates how to set the policy to a group (organizational unit in Active Directory terms) of users that use that printer. However, you can set the policy to a different object using the appropriate Active Directory utility.

   - Start Group Policy from Active Directory Users and Computers. To apply Group Policy to a domain or organizational unit, see Related Topics.
   - Right-click the organizational unit you want to apply the policy to, and then click **Properties**.
   - Click the **Group Policy** tab, and then click **New** to create a new object. Type the name of the new object, and then click **OK**.
   - Select the newly created object and then click **Edit**.
   - Expand **Computer Configuration**, then **Administrative Templates**, and then **Printers**.
   - Double-click the policy setting **Pre-populate printer search location**, and then click to enable it. When you search the directory service for printers, it automatically fills in your location field using the value specified in the **Location** tab of the Subnet properties page. If this policy is disabled (default), the location will be empty.

7. Open Printers and Faxes.
8. Right-click the printer, and then click **Properties**.
9. Type the printer location, or click **Browse** to find it. You probably want to be more precise in describing the printer location than the subnet location. For example, for the subnet location US/NYC, you might enter **US/NYC/Floor42/Room4207**.
10. Repeat the last two steps for each printer in Active Directory.

**Notes**

- To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.
- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- To enable printer location tracking, you must have the following:
    - A working directory service with more than one site or more than one subnet.
    - A network IP addressing scheme that corresponds roughly to the physical layout of your enterprise.
    - Client computers that can query the directory service.
    - A subnet object for each site. You can create them using Active Directory Sites and Services.
    - Your own naming convention for printer location.
    - Administrative credentials.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a site
Associate a subnet with a site
Establishing a naming convention for printer locations
Manage Organizational Units
Manage Sites

# Start and stop the print spooler

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start and stop the print spooler**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. In the details pane, double-click **Print Spooler**.
4. Do one of the following depending on the task you want to do:

   Click **Stop** to stop the service.

   Click **Start** to start the service.

## Notes

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- When you stop the Print Spooler service, it might stop other dependent services.
- The print spooler runs under the Local Computer system account by default.
- You can also start and stop the Print Spooler service through a command-prompt command. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Run printing commands from the command prompt
Enable or disable logging of printing events

# Run printing commands from the command prompt

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run printing commands from the command prompt**

1. Open Command Prompt.
2. At the command prompt, you can run the following printing-related commands:

   - To print a text file or display the contents of a print queue.
   - To display or control print jobs and printer queues.
   - To connect to, or disconnect from, a shared printer.
   - To start the print spooler service.
   - To stop the print spooler service.

**To print a text file or display the contents of a print queue.**
For command syntax, see Print.

Example:

To print the text file Autoexec.bat to a printer on parallel port LPT2:

**print /d:LPT2: C:\autoexec.bat**

**To display or control print jobs and printer queues.**
For command syntax, see Net print.

Examples:

To view the print queue of MyPrinter1 on print server W2KPRN1:

**net print \\W2KPRN1\MyPrinter1**

To delete print job number 2 on MyPrinter1 on print server W2KPRN1:

**net print \\W2KPRN1 2/delete**

**To connect to, or disconnect from, a shared printer.**
For command syntax, see Net use.

Examples:

To redirect print output for the LPT1 port to MyPrinter1 on print server W2KPRN1:

**net use LPT1: \\W2KPRN1\MyPrinter1**

To display information about the LPT1 port:

**net use LPT1:**

**To start the print spooler service.**
For command syntax, see Net start.

Example:

To start the Print Spooler service:

**net start** *spooler*

**To stop the print spooler service.**
For command syntax, see Net stop.

Example:

To stop the Print Spooler service:

**net stop** *spooler*

Note that other dependent services might need to be stopped as well. The system prompts you to confirm, while listing the other service names.

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- Command-prompt commands can be useful if you need to connect users to shared printers through a login script or any other automated batch or a script procedure.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Change the rendering intent for viewing or printing an image or page

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the rendering intent for viewing or printing an image or page**

1. In the graphics program in which you are creating your image or page, on the **File** menu, click **Color Management**.
2. Select the **Enable Color Management** check box, and then click **Basic color management**.
3. In **Rendering Intent**, click the rendering intent you want to use.

**Note**

- Your graphics program must support Image Color Management 2.0 in order to implement color management options.

# Manage Printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage printers**

- Remotely manage printers
- Manage printers from a browser
- Manually publish a printer in Active Directory
- Set Group Policy for printers
- Block the installation of kernel-mode printer drivers
- Take ownership of a printer
- Track printer usage
- Monitor print queue performance
- Search for Printers

# Remotely manage printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remotely manage printers**

1. Double-click **My Network Places**, and then locate the print server for the printers you want to manage.

2. Double-click the print server, double-click the **Printer folder** icon on that server, and then click a printer.

   **Important**

   - To facilitate stronger network security, remote printer management is not available by default. To enable remote printer management, in Group Policy, you must enable the Allow Print Spooler to Accept Client Connections policy. For more information, see Related Topics.

3. Change the print server, printer, or printing preference settings as required.

## Notes

- You can change printer settings only if you have Manage Printers permission.

- If you are having trouble finding your printer and you are logged on to a Windows Server 2003 family domain, you can find your printer by clicking **Start**, pointing to **Search**, and then clicking **Printers**.

- You can also drag the Printers and Faxes folder of any print server (or specific printers from any Printers and Faxes folder) to the Printers and Faxes folder on your computer. This is an easy way to manage any printer without having to search for it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set printer properties
Set printing defaults
To set print server properties
Set or remove permissions for a printer
Set Group Policy for printers

# Manage printers from a browser

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage printers from a browser**

1. From Internet Explorer, or any other browser, type the following URL:

   **http://***PrintServerName*/**printers/**

   Or, type a specific printer URL:

   **http://***PrintServerName/PrinterName/*

2. In **All Printers on***PrintServerName*, click the printer you want to manage.

   **Important**

   - To facilitate stronger network security, remote printer management is not available by default. To enable remote printer management, in Group Policy, you must enable the Allow Print Spooler to Accept Client Connections policy. For more information, see Related Topics.

3. In *PrinterName* **on** *PrintServerName*, you can click any function on the left pane to stop, resume, or cancel a specific document or all documents. You can also click a specific document in the queue to see its properties.

**Notes**

- An administrator can disable Internet printing with the Group Policy setting **Disable Web-based Printing**.

- For Internet printing you must have Internet Information Services (IIS) installed on a computer running a Windows Server 2003 family operating system (this is installed by default) or Peer Web Services (PWS) on a computer running Windows XP Professional.

- You can also access the Web-based view from the **Printers** folder if **Enabled Web view content in folder** is selected (the folder has the Web-based view in the left pane). If you click any printer in the right pane, you can select **Get More Info** in the left pane to view the printer's content in Web-based view.

  In addition, you can use the left pane to jump to the Microsoft support Web site, and, if available, the printer's vendor support Web site.

- You can manage any printer connected to a computer running a Windows Server 2003 family operating system from any other computer running a Windows Server 2003 family operating system, when Web-based printing is enabled. Web-based printing is enabled by default.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set or remove permissions for a printer
Choosing a printer name
Set Group Policy for printers

# Manually publish a printer in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manually publish a printer in Active Directory**

1. Open **Active Directory Users and Computers**.
2. In **Name**, right-click the container object folder in which you want to publish the printer.
3. Point to **New**, and then click **Printer**.
4. Type the UNC name of the printer that you want to publish in the directory.

## Notes

- To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
- To perform this procedure, you must be a member of the Print Operators group, Domain Admins group, or the Enterprise Admins group in Active Directory, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- Group Policy settings for printers determine whether printers published with a print server running Windows NT 4.0 or an earlier operating system should be pruned if the print server becomes unavailable.
- You can also use the **Pubprn.vbs** script, provided in the System32 folder, to publish printers with print servers running Windows NT 4.0 or an older operating system.
- The printer you want to publish must be shared.
- You must have directory service administration rights to publish printers.
- You must have the Manage Printers permission for the printer you want to share or to publish.
- You can create printer names that contain spaces, special characters, and more than eight characters. However, some clients do not recognize or correctly handle such names. The entire qualified name of a shared printer (for example, \\PRINTER2\PSCRIPT) must contain 32 or fewer characters.

**See Also**
**Concepts**
Group Policy and printing
Publish a printer in Active Directory

# Set Group Policy for printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set Group Policy for printers**

1. Start Group Policy according to the object you want to set printer policy to. For more information on how to start Group Policy, see Related Topics.

2. After selecting the properties page of the object you want to set printer policy to, select the Group Policy node.

   - If you want to set policies that apply only to computers, expand the **Computer Configuration** node, and then expand **Administrative Templates**.

   - If you want to set policies that apply only to users, expand the **User Configuration** node, expand **Administrative Templates**, and then expand **Control Panel**.

3. Double-click **Printers** to open a listing of policies.

4. Double-click the printer policy you want to set.

5. On the **Policy** tab, enable or disable the policy by selecting or clearing the appropriate radio button. With some policies, you might need to enter additional information.

### Note

- If you do not want to change the current state of the policy setting, leave it as it is (not configured) to save processing time.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Group Policy overview
Group Policy and printing

# Block the installation of kernel-mode printer drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To block the installation of kernel-mode printer drivers**

1. Open Group Policy.
2. Under **Local Computer Policy**, double-click **Computer Configuration**.
3. Double-click **Administrative Templates**, and then double-click **Printers**.
4. Right-click **Disallow installation of printers using kernel mode drivers**, and then click **properties**.
5. On the **Setting** tab, click **Enabled**, and then click **OK**.

### Important

- This policy does not affect existing kernel-mode printer drivers. Kernel-mode printer drivers that were installed before you set this policy will not be disabled or removed. You can upgrade existing non-kernel-mode printer drivers whether you set this policy or not.

### Notes

- To open Group Policy so that you can edit the local Group Policy object, click **Start**, click **Run**, type **gpedit.msc**, and then press ENTER.
- In Windows Server 2003, Standard Edition and Windows Server 2003, Enterprise Edition, installation of kernel-mode drivers is blocked by default. In Windows XP Professional, installation of kernel-mode drivers is not blocked by default.
- You must click **Enabled** to block the installation of kernel-mode printer drivers. Clicking **Not Configured** or **Disabled** allows kernel-mode printer drivers to be installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Printer drivers overview
Install new or updated printer drivers

# Take ownership of a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To take ownership of a printer**

1. Open Printers and Faxes.
2. Right-click the printer icon, and then click **Properties**.
3. Click the **Security** tab, and then click **Advanced**.
4. Click the **Owner** tab, click your user account under **Change owner to**, and then click **OK** twice.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- If you are a member of the Administrators group and you want this group to take ownership of the printer, click the Administrators group.
- By default, the user who installed the printer owns it.
- A user or a member of a group who has Manage Printers permission for the printer can take ownership.
- By default, members of the Administrators, Print Operators, Server Operators, and Power Users groups have Manage Printers permission, which allows them to take ownership of a printer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Track printer usage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To track printer usage**

1. Open Printers and Faxes.
2. Right-click the printer you want to audit, and then click **Properties**.
3. Click the **Security** tab, click the **Advanced** button, and then click the **Auditing** tab.
4. Click **Add**, and then click the user or group whose printer access you want to audit.
5. In the **Access** column, click successful and failed printing events you want to audit, and then click **OK** twice.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- **Successful** means you want to audit all successful attempts to perform this action.
- **Failed** means you want to audit all failed attempts to perform this action.
- You can view the audit information in the security log file.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Monitor print queue performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To monitor print queue performance**

1. Open Performance.
2. Click the plus sign (**+**) on the Performance toolbar.
3. From the **Select counters from computer** list box, click the print server computer.
4. From the **Performance object** list box, click **Print Queue**.
5. Click the counters you want to monitor.

## Notes

- To open Performance, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Performance**.
- To view a detailed explanation of a counter, select it, and then click **Explain**.
- You can also add System Monitor to any new Microsoft Management Console (MMC).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Search for Printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Search for printers**


- Search for a printer
- Filter your search results
- Use advanced options to search for a printer
- Add or remove columns in the search results list
- Create a custom printer search

# Create a custom printer search

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a custom printer search**

1. Search for printers in Active Directory. For more information, see Related Topics.

2. After the search results appear, you can refine the search as necessary by using any of the search options. In particular, the **Advanced** tab can be a powerful way to build custom printer queries.

   To change the appearance of search results to the way you want them to appear to users, click **View**.

3. Click **File**, and then click **Save Search**. Select where you want the search to be saved, type a name for the search, and then click **Save**.

## Notes

- As administrator, you can save search results into users profiles, which allow you to distribute the search results file to users' Printers and Faxes folders, desktops, or any other location on users' computers.

- You and your users can further narrow search results by using the **Filter** option on the **View** menu, which is available only when you click **Details**. However, filtering criteria are not saved by **Save Search**.

- Users can use the saved search file by double-clicking the file name. The saved search runs in the background, and only the result of the saved search appears.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set or remove permissions for a printer
Choosing a printer name

# Connect clients to a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connect clients to a printer**


- Connect to a printer published in Active Directory
- Connect to a printer on a network
- Connect to a printer with a browser
- Connect Windows NT 4.0, Windows 95, and Windows 98 clients
- Connect MS-DOS, Windows 3.x, and OS/2 clients
- Connect Macintosh clients to a printer
- Print from an LPR client to a printer
- Connect to a printer on a NetWare server
- Connect to a printer on a Macintosh server

# Connect to a printer published in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a printer published in Active Directory**

1. Find the printer you want to connect to. For instructions on how to search for a printer, see Related Topics.

2. In the details pane, right-click the printer you want to use, and then click **Connect**, or double-click the printer name from the list.

## Notes

- You can search for printers in Active Directory only if you are logged on to a Windows Server 2003 family or Windows 2000 domain, and your client is running Windows XP Professional, Windows 2000, Windows 95, or Windows 98.

- Windows 95 and Windows 98 clients must run the directory service to query Active Directory.

- Windows NT 4.0 clients can connect to a shared printer on a print server by browsing My Network Places, or by using the Add Printer Wizard.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect Windows NT 4.0, Windows 95, and Windows 98 clients

# Connect to a printer on a network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a printer on a network**

1. Open Printers and Faxes.
2. Double-click **Add Printer** to start the Add Printer Wizard, and then click **Next**.
3. Click **A network printer, or a printer attached to another computer**, and then click **Next**.
4. Connect to the desired printer using one of the following three methods.

   **Searching for it in the Active Directory**

   If you are not logged on to a domain running Active Directory, this method will not be available.

   a. Click **Find a printer in the Directory**, and then click **Next**.

      If needed, change the default printer location listed in the **Location** field by clicking the Location **Browse** button, and then choosing the appropriate location.

   b. Click **Find Now**.

   c. In the list that appears, click the printer you want to connect to, and then click **OK**.

   **Typing the printer name or browsing for it**

   a. Click **Connect to this printer**, or click **Next** to browse for a printer.

   b. Do one of the following:

      • Type the printer name using the following format, and then click **Next**:

      *****printserver_name******share_name*

      • Browse for it on the network by clicking **Next**, locating the printer in **Shared printers**, and then clicking **Next**.

   **Connecting to an Internet or intranet printer**

   Connecting to a printer using its URL allows you to connect to printers across the Internet, provided you have permission to use that printer. If you cannot connect to your printer using the general URL format above, please see your printer's documentation or contact your network administrator.

   a. Click **Connect to a printer on the Internet or on a home or office network**.

   b. Type the URL to the printer using the following format, and then click **Next**:

      http://*printserver_name*/Printers/*share_name*/.printer

5. Follow the instructions on the screen to finish connecting to the printer.

The icon for the printer appears in your Printers and Faxes folder.

### Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also connect to a printer by dragging the printer from the Printers and Faxes folder on the print server and dropping it into your Printers and Faxes folder, or by right-clicking the icon and then clicking **Connect**.
- After you have connected to a shared printer on the network, you can use it as if it were attached to your computer.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a printer on a NetWare server

# Connect to a printer with a browser

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a printer with a browser**

1. Open Internet Explorer, or open Printers and Faxes.
2. In the Address bar, type one of the following:
   - If you do not know the printer name, type the following, using the format: **http://***PrintServerName***/printers/**

     For example, type **http://PrintSrv1/printers** to receive a page listing all the printers, and then click the printer you want.
   - If you know the printer name, type its URL using the format: **http://***PrintServerName/PrinterName/*

     For example, type **http://PrintSrv1/ColorPrinter5/** to go directly to that printer's page.
3. When viewing the printer's page, click **Connect** to connect to that printer.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- A server running a Windows Server 2003 family operating system automatically copies the appropriate printer driver to the client computer.
- The administrator can disable Internet printing using the Group Policy setting **Disable Web-based Printing** for printers.
- For Internet printing, you must have Internet Information Services (IIS) installed on the server (or Peer Web Services (PWS) on a computer running Windows XP Professional .
- You must use Internet Explorer version 4.0 or higher to connect to a printer.
- After you connect to a shared printer from a Web browser, you can use it as if it were attached to your computer. You can also add it to your browser Favorites.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a printer on a network

# Connect Windows NT 4.0, Windows 95, and Windows 98 clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect Windows NT 4.0, Windows 95, and Windows 98 clients**

1. Open Printers and Faxes.
2. Double-click **Add Printer**, and then click **Next**.
3. Click **A network printer, or a printer attached to another computer**, and then click **Next**.
4. Click the printer in **Connect to this printer**.

   Or, in **Name**, type the name of a printer on your network using the following format:

   *server_name**share_name*

5. Follow the instructions on the screen to finish connecting to the printer.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- After you complete these procedures, the printer driver downloads automatically, and the icon for that printer appears in your Printers and Faxes folder.
- To connect a printer to a Windows NT 4.0, Windows 95, or Windows 98 client, you must install the print drivers for these clients on the print server. During connection, the print driver downloads to the client without the need for the original compact disc or discs.
- To install additional printer driver for previous Windows NT clients and for Windows 95 and Windows 98 clients, see Related Topics.
- If directory services is not installed, double-click **My Network Places** to locate the print server and the printer you want to connect to. Then, right-click the printer and select **Install**.
- If directory services is installed (directory service clients are available for Windows 95 and Windows 98 only), you can use Active Directory to find the printer you want to connect to. Right-click the printer's name, and then click **Connect**.
- For Windows NT 4.0 clients, the print driver is automatically downloaded if it does not exist. If it already exists, it is updated on the print server. For Windows 95 and Windows 98 clients, the print driver is downloaded automatically only the first time (if it does not exist). When an updated driver is installed on the print server, the printer driver must be updated manually from the print server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a printer on a network
Add printer drivers for other versions of Windows

Manage printers from a browser
Connect to a printer with a browser

# Connect MS-DOS, Windows 3.x, and OS/2 clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect MS-DOS, Windows 3.x, and OS/2 clients**

1. Install the appropriate printer driver on the client computer.

2. Redirect an output port (such as LPT1) of the client computer to the logical printer (in the form of \\PrintServerName\sharename) on the Windows Server 2003 family print server, using the command specific to its operating system.

   Typically, the command for clients that add the **Clients for Microsoft Networks** option is:

   **net use lptx***server_name******share_name*

## Notes

- Other Microsoft clients include Windows 3.x, Windows for Workgroups, MS-DOS, and OS/2 with the LAN Manager 2.x client installed.

- If the client is not connected to the network, install the network client on the client computer first, and then verify active connection to the network.

- Fonts and forms available on a Windows Server 2003 family operating system might not be available on client computers running other operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Connect Macintosh clients to a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect Macintosh clients to a printer**

1. Verify that the print server is running the AppleTalk protocol, and Print Server for Macintosh.
2. Set up a user account for Macintosh print jobs. For more information, see Related Topics.
3. From a Macintosh client computer, connect to the printer using Macintosh Chooser.

## Notes

- Macintosh clients can connect to any printer on the Windows Server 2003 family network that runs the AppleTalk protocol.
- Install the printer driver for Macintosh on each client that you want to connect to the printer.
- Macintosh clients can search for printers on print server computers running a Windows Server 2003 family operating system by browsing for all printer addresses using Macintosh Chooser. Macintosh users do not have access to Active Directory.
- If your Macintosh uses a protocol other than AppleTalk, such as TCP/IP, verify that both the clients and the print server run that protocol.
- Because AppleTalk does not support client user names or passwords, Macintosh clients cannot identify themselves on a Windows Server 2003 family network, and the print server cannot impose user-level security on them. If a Macintosh client is physically able to send a job to the printer, it has implicit permission. However, an administrator can set user-level permission for all Macintosh print clients as a group.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a printer on a Macintosh server
Set up a user account for Macintosh print jobs
Connecting clients to a printer

# Print from an LPR client to a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print from an LPR client to a printer**

1. Install the appropriate printer driver on the client.

2. Install Print Services for Unix on the Windows Server 2003 family print server, and then start the Line Printer Daemon (LPD) service (if not started automatically) from Services.

3. Use the Line Printer Remote (LPR) utility on your client to connect to the logical printer on the print server. Refer to your client operating system documentation for the correct command. Typically the LPR command is in the following format:

   **lpr -S***server_name***-P***queue_name file_name*

## Notes

- Both client and server must run TCP/IP, and the client must have an RFC-compliant version of LPR.

- Some hosts do not support LPR. Refer to your operating system documentation to make sure your host does support LPR.

- If documents sent by LPR services on a UNIX client do not print correctly on the print server, try reconfiguring the LPR software to use a lowercase control command.

- The Lpdsvc service is independent of the Lprmon, which allows a computer running a Windows Server 2003 family operating system (and all clients who can access this computer) to print to a printer connected to a UNIX system.

- If you send a document from a Windows XP client to another LPD-enabled target (for example, a UNIX host), the Windows XP client must run Print Services for Unix.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Connect to a printer on a NetWare server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a printer on a NetWare server**

1. Open Printers and Faxes.
2. Double-click **Add Printer**, and then click **Next**.
3. Click **A network printer, or a printer attached to another computer**, and then click **Next**.
4. Click **Connect to this printer**.
5. In **Name**, type the name of a printer in the following format:

   *****servername******sharename*

   Or, to find the NetWare printer in **Shared printers**, click **Next**.
6. Follow the remaining instructions in the Add Printer Wizard to finish connecting to the NetWare printer.

   The icon for that printer appears in your Printers and Faxes folder.

### Important

- The IPX/SPX protocol is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family. Due to its dependency on the IPX/SPX protocol, Client Service for NetWare is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family.

### Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- To connect to a printer on a NetWare server, you must have Client Service for NetWare and NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) installed. For information on how to do this, see Related Topics.
- If a printer driver is not available locally for a NetWare printer, you are prompted to install one.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to a printer attached to a NetWare server
Add a port

# Connect to a printer on a Macintosh server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a printer on a Macintosh server**

1. Open Printers and Faxes.

2. Double-click **Add Printer** to start the Add Printer Wizard, and then click **Next**.

3. Click **Local printer attached to this computer**, and then click **Next**.

4. Follow the instructions on the screen to finish setting up the printer by selecting a printer port, selecting the manufacturer and model of your printer, and typing a name for your printer.

   When the wizard prompts you to select the printer port, click **Create a new port**. From the drop-down list, click **AppleTalk Printing Device**, and then click **Next**.

   When you're prompted for available AppleTalk printing devices, select a zone and a printer, and then click **OK**.

   When you're asked whether you want to capture the printer, you normally click **Yes**; see the notes following this procedure.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

- Capturing the printer is the recommended procedure, because a captured printer accepts print jobs only from print server computers running a Windows Server 2003 family operating system. This provides the administrator complete control over the printer. However, if the printer is also being used by a Macintosh server, you cannot capture the printer.

- Because AppleTalk does not support client user names or passwords, Macintosh clients cannot identify themselves on a Windows Server 2003 family network, and the print server cannot impose user-level security on them. If a Macintosh client is physically able to send a job to the printer, it has implicit permission. However, an administrator can set user-level permission for all Macintosh print clients as a group.

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect Macintosh clients to a printer
Connecting clients to a printer

# Manage and Set Print Servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage and set print servers**

- Add printer drivers for other clients
- Update printer drivers
- Remove printer drivers
- List all files on a printer driver
- Change the spool folder location
- Enable or disable logging of printing events
- Enable or disable notifying users when printing completes
- Create a custom printer form
- Set Ports and Protocols

# Add printer drivers for other clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add printer drivers for other clients**

1. Open Printers and Faxes.
2. On the **File** menu, click **Server Properties**, and then click the **Drivers** tab.
3. Click **Add**, click **Next**, and then follow the instructions on the Add Printer Driver Wizard.
4. After selecting the printer's manufacturer and model, you are prompted to select the environment and operating systems for which you want to install drivers.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can add printer drivers for other platforms (such as Alpha-based or x86-based), or for other operating systems, using the **Sharing** tab on each printer's properties page. If you need to add additional drivers for multiple printers, it might be easier to use the **Drivers** tab on the properties page of the print server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Update printer drivers
Remove printer drivers
List all files on a printer driver

# Update printer drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To update printer drivers**

1. Open Printers and Faxes.
2. On the **File** menu, click **Server Properties**, and then click the **Drivers** tab.
3. In the **Installed printer drivers** list, click the driver you want to update, and then click **Replace**.
4. Click **Yes**. The driver files are copied for all the installed environments and versions. The system might prompt you for the path to the drivers.

### Note

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add printer drivers for other clients
Remove printer drivers
List all files on a printer driver

# Remove printer drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove printer drivers**

1. Open Printers and Faxes.
2. On the **File** menu, click **Server Properties**, and then click the **Drivers** tab.
3. In the **Installed printer drivers** list, click the driver you want to remove, click **Remove**, and then click **Yes** to confirm.

### Note

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Update printer drivers
Add printer drivers for other clients
List all files on a printer driver

# List all files on a printer driver

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To list all files on a printer driver**

1. Open Printers and Faxes.
2. On the **File** menu, click **Server Properties**, and then click the **Drivers** tab.
3. In the **Installed printer drivers** list, click the driver you want, and then click **Properties**.

   The names and descriptions of all the files that are part of the specific driver are listed.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can view details on any of the files by clicking the file and then clicking **Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Update printer drivers
Remove printer drivers
Add printer drivers for other clients

# Change the spool folder location

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the spool folder location**

1. Open Printers and Faxes.
2. On the **File** menu, click **Server Properties**, and then click the **Advanced** tab.
3. In the **Spool folder** window, enter the path and the name of the new default spool folder for this print server, and then click **Apply** or **OK**.
4. Stop and restart the spooler service, or reboot the server.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- If the server functions as the print server for a large number of printers or for print jobs that are typically very long, you should consider moving the spool folder to another location.
- The spool folder might require a large amount of disk space to spool all the jobs.
- The spool folder should not be located at the root directory (generally C:\).
- If the new folder does not exist, this procedure creates it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Setting the location of the spool folder
Enable or disable logging of printing events

# Enable or disable logging of printing events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To enable or disable logging of printing events

1. Open Printers and Faxes.
2. On the **File** menu, click **Server Properties**, and then click the **Advanced** tab.
3. To enable or disable writing spooler error events to the system log file, select or clear the **Log spooler error events** check box.
4. To enable or disable writing spooler warning events to the system log file, select or clear the **Log spooler warning events** check box.
5. To enable or disable writing spooler information events to the system log file, select or clear the **Log spooler information events** check box.
6. Stop and restart the spooler service, or reboot the server.

### Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- By default, all printing events are written to the system log file. If you want to reduce the amount of data that is logged, disable the logging of information events. If you want to further reduce the amount, disable the warning events. Disabling the logging of spooler error events is not recommended.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Enable or disable notifying users when printing completes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable notifying users when printing completes**

1. Open Printers and Faxes.
2. Click **File**, click **Server Properties**, and then click the **Advanced** tab.
3. To enable or disable notifying users when their print jobs finish printing, select or clear the **Notify when remote documents are printed** check box.
4. Stop and restart the spooler service, or reboot the server.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- When the Group Policy setting **Notify Computer, not User** is enabled, the computer that originated the print job is notified, even if the user who submitted the print job has logged off and another user has logged on.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Create a custom printer form

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a custom printer form**

1. Open Printers and Faxes.
2. On the **File** menu, click **Server Properties**.
3. On the **Forms** tab, select the **Create a new form** check box.
4. In **Form name**, type a name for the new form.
5. Under **Form description (measurements)**, type the dimensions you want for the new form.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- Printer forms define a paper size and type used by a print device. Forms are provided for all standard paper sizes. You can define a custom paper form for nonstandard paper sizes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Set Ports and Protocols

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Set ports and protocols**

- Add a port
- Delete a port
- Add a local port
- Add a standard TCP/IP port
- Add an LPR port
- Add an AppleTalk printing device port

# Add a port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a port**

1. Open Printers and Faxes.
2. On the **File** menu, click **Server Properties**.
3. Click the **Ports** tab, and then click **Add Port**.
4. In the **Available port types** list, click a port name to select it. By default, only **Local Port** and **Standard TCP/IP Port** appear in the **Available port types** list.

   If you want to add a Line Printer Remote (LPR) port, you need to install the optional networking component, **Print Services for Unix**. Once you do so, **LPR Port** will display as one of the available port types. For more information, click **Related Topics**.

   To add other port types, click **New Port Type**, type the path to the compact disc containing the necessary files, and click **OK**.
5. After clicking the port type in the **Printer Ports** dialog box, click **New Port**.
6. Type a name that describes this port, and then click **OK**.
7. On the **Printer Ports** dialog box, click **Close**.
8. If necessary, configure the port by clicking **Configure Port**. Different types of ports have different characteristics that you can configure. After you have finished configuring the port, click **OK**.
9. Click **Close** to close the **Print Server Properties** dialog box.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- For a list of ports you can add, and information about adding them, click **Related Topics**.
- If you add a port using the Add Printer Wizard, the port you add will appear in the list.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a port
Add a printer attached directly to the network
Add a standard TCP/IP port
Add a local port
Add an LPR port
Add an AppleTalk printing device port

## Add optional networking components

# Delete a port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a port**

For this procedure, you must have the Manage Printers permission, which is one of several printer permissions.

1. Open Printers and Faxes.
2. On the **File** menu, click **Server Properties**, and then click the **Ports** tab.
3. In the list of ports, click the port you want to delete, and then click **Delete Port**.
4. Click **Yes** to confirm.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- If you are unable to delete the port and you know that you have printer permissions to do so, it may be that one or more of your printers is set up to print to that port. Reassign the printer or printers to another port, and then try to delete the port again. Deleting a port removes it from the available port list. Any printer that communicates through that port is disconnected.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Add a port
Add a local port
Add an LPR port
Add a standard TCP/IP port

# Add a local port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a local port**

1. Open Printers and Faxes.
2. Double-click **Add Printer** to open the Add Printer Wizard, and then click **Next**.
3. Click **Local printer attached to this computer**, clear the **Automatically detect and install my Plug and Play printer** check box, and then click **Next**.
4. Click **Create a new port**, click **Local Port**, and then click **Next**.
5. Type the port name, and then click **OK**.
6. Follow the instructions on the screen to finish adding a printer with the Add Printer Wizard.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also add ports using the **Print Server Properties** dialog box.
- You can configure the time-out for parallel (LPT) ports by selecting any LPT port and then clicking **Configure Port**.
- In a network environment, it is better to use a printer that is connected directly to the network through a network adapter than through a parallel port.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a port
Add a printer attached directly to the network
Add a standard TCP/IP port
Add a port
To change the parallel port time-out settings
Add an LPR port
Add an AppleTalk printing device port

# Add a standard TCP/IP port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a standard TCP/IP port**

1. Open Printers and Faxes.
2. Double-click **Add Printer** to open the Add Printer Wizard, and then click **Next**.
3. Click **Local printer attached to this computer**, clear the **Automatically detect and install my Plug and Play printer** check box, and then click **Next**.
4. Click **Create a new port**, and then click **Standard TCP/IP Port**.
5. Click **Next** to run the Add Standard TCP/IP Printer Port Wizard.
6. Follow the instructions on the screen to finish installing the TCP/IP printer.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- Most network printers support the TCP/IP protocol. The standard TCP/IP port simplifies connecting to remote printers using the TCP/IP protocol. A computer that acts as a print server must run the TCP/IP protocol in order to print.
- To connect to a print server running NetWare, you need to install Client Services for NetWare.
- You can also add ports using the **Print Server Properties** dialog box. For more information, click **Related Topics**.
- For print servers that need to communicate with host computers (such as UNIX or VAX computers), it is best to install an LPR port. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add a port
Add an LPR port
Print from an LPR client to a printer

# Add an LPR port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an LPR port**

1. Open Printers and Faxes.
2. Double-click **Add Printer** to open the Add Printer Wizard, and then click **Next**.
3. Click **Local printer attached to this computer**, clear the **Automatically detect and install my Plug and Play printer** check box, and then click **Next**.
4. Click **Create a new port**, and then click **LPR Port**.

   If **LPR Port** is not available, click **Cancel** to stop the wizard. To add the LPR port, you need to install the optional networking component, **Print Services for Unix**. For more information, click **Related Topics**.
5. Click **Next**, and then provide the following information:

   - In **Name or address of server providing LPD**, type the Domain Name System (DNS) name or Internet Protocol (IP) address of the host for the printer you are adding. The host may be the direct-connect TCP/IP printing device or the UNIX computer to which the printing device is connected. The DNS name can be the name specified for the host in the Hosts file.
   - In **Name of printer or print queue on that server**, type the name of the printer as it is identified by the host, which is either the direct-connect printer itself or the UNIX computer.
6. Follow the instructions on the screen to finish installing the TCP/IP printer.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- The Line Printer Remote (LPR) port is best suited to servers that need to communicate with host computers (such as UNIX or VAX computers) by way of RFC 1179.
- A network-connected printer must have a card that supports LPD for LPR printing to work properly.
- You can also add ports using the **Print Server Properties** dialog box. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a port
Add a printer attached directly to the network
Add a standard TCP/IP port
Add a local port
Add a port

Add an AppleTalk printing device port
Add optional networking components

# Add an AppleTalk printing device port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add an AppleTalk printing device port**

1. Open Printers and Faxes.
2. Double-click **Add Printer**, and then click **Next**.
3. Click **Local printer attached to this computer**, clear the **Automatically detect and install my Plug and Play printer** check box, and then click **Next**.
4. Click **Create a new port**, click **AppleTalk Printing Devices**, and then click **Next**.
5. In the **AppleTalk Printing Devices** list, click the AppleTalk printer you want to add, and then click **OK**.
6. Click **Yes** to capture this AppleTalk printer.

   **Important**

   - If you capture a printer, some users may be unable to print to it.

7. Follow the instructions on the screen to finish installing the AppleTalk printer.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also add ports on the **Ports** tab in the **Print Server Properties** dialog box. To open the **Print Server Properties** dialog box, click **Start**, point to **Settings**, and then click **Printers and Faxes**. On the **File** menu, click **Server Properties**.
- **AppleTalk Printing Devices** is available only when AppleTalk protocol is installed. If **AppleTalk Printing Devices** is not available, click **Cancel** to close the Add Printer Wizard and follow the procedures to add optional networking components (see Related Topics).
- Your system must use the AppleTalk protocol if you want to add an AppleTalk printing device. You might need other Macintosh services if you want to coexist with an AppleTalk network system. For more information, see Related Topics.
- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Delete a port
Add a printer attached directly to the network
Add a standard TCP/IP port
Add a local port
Add an LPR port

Add a port
Add optional networking components
Manage AppleTalk Printing

# Manage Documents

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage documents**

- Print a document
- View documents waiting to print
- Change the printing priority of documents waiting to print
- Cancel printing a document
- Pause or resume printing a document
- Restart printing a document
- Transfer documents to another printer
- Pause or resume printing all documents
- Cancel printing all documents

# Print a document

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To print a document**

1. Open the document you want to print.
2. On the **File** menu of the program you are using, click **Print**.

## Notes

- In the **Print** dialog box, you can see the number of documents waiting to print on any printer installed on your computer by right-clicking the printer in **Select Printer** and then clicking **Open**.
- If you are logged on to a domain running Active Directory, you can click **Find Printer** to search for a printer on the network with the capabilities you want, such as high-speed or color printing. For more information on searching for a printer, see Related Topics.
- To store a document as a file rather than sending it to the printer, click the **Print to file** check box in the **Print** dialog box.
- For easy access to your printer, you can create a shortcut to it on the desktop. You can double-click the shortcut to open the print queue and view the documents waiting to print. For more information on how to create a shortcut, see Related Topics.
- You can print a document without opening it by dragging its icon to a printer in the Printers and Faxes folder or to a shortcut on your desktop.
- While a document is printing, a printer icon appears next to the clock in the notification area on the taskbar. When this icon disappears, it means that your document has finished printing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Print queue overview
View documents waiting to print
Specify your default printer

# View documents waiting to print

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view documents waiting to print**

1. Open Printers and Faxes.
2. Double-click the printer for which you want to view the documents waiting to print.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- The print queue displays information about a document such as print status, owner, and number of pages to be printed. From the print queue, you can cancel or pause printing for any document you send to the printer.
- You can also open the print queue by double-clicking the printer icon in the notification area.
- Another way to open the print queue is to click the printer in Printers and Faxes, and then click **See what's printing** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Cancel printing a document
Change the printing priority of documents waiting to print
Pause or resume printing a document
Restart printing a document
Print queue overview

# Change the printing priority of documents waiting to print

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the printing priority of documents waiting to print**

You can change the printing priority of documents waiting to print only if you have the Manage Documents permission, which is one of several printer permissions. Once a document has started printing, any printing priority changes you make will not affect that document. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Double-click the printer you are using to open the print queue.
3. Right-click the document you want to move in the print order, and then click **Properties**.
4. On the **General** tab, drag the **Priority** slider to raise or lower the document priority. The higher the number, the higher the priority.

   When you click **OK**, the documents will print in the order you want even though your changes may not be reflected in the list of documents.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open the print queue in Printers and Faxes by clicking a printer, and then clicking **See what's printing** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Cancel printing a document
Pause or resume printing a document
Restart printing a document
Print queue overview
Assigning printer permissions

# Cancel printing a document

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To cancel printing a document**

By default, all users can pause, resume, restart, and cancel printing of their own documents. However, to manage documents printed by other users, you must have the Manage Documents permission, which is one of several printer permissions.

1. Open Printers and Faxes.
2. Double-click the printer you are using to open the print queue.
3. Right-click the document you want to stop printing, and then click **Cancel**.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open the print queue in Printers and Faxes by clicking a printer, and then clicking **See what's printing** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- You can cancel the printing of more than one document by holding down the CTRL key and then clicking each document you want to cancel. Once all the documents are selected, right click one of them and click **Cancel**.
- You can also open the print queue by right-clicking the printer icon in the notification area.
- For more information on administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Pause or resume printing a document
Restart printing a document
Change the printing priority of documents waiting to print
Cancel printing all documents
Prnjobs.vbs

# Pause or resume printing a document

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To pause or resume printing a document**

By default, all users can pause, resume, restart, and cancel printing of their own documents. However, to manage documents printed by other users, you must have the Manage Documents permission, which is one of several printer permissions.

1. Open Printers and Faxes.
2. Double-click the printer you are using to open the print queue.
3. Right-click the document for which you want to pause or resume printing.
4. Do one of the following:
   - To pause printing, click **Pause**.

     The document will not print until you resume printing.
   - To resume printing, click **Resume**.

     The document will begin printing. However, if higher priority documents are waiting to print, they will print first.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open the print queue in Printers and Faxes by clicking a printer, and then clicking **See what's printing** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- In general, once a document has started printing, it will finish printing even if you pause it.
- For more information on administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Cancel printing a document
Restart printing a document
View documents waiting to print
Prnjobs.vbs

# Restart printing a document

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restart printing a document**

When you restart printing a document, it starts printing again from the beginning. If higher priority documents are in the print queue, they will print first.

By default, all users can pause, resume, restart, and cancel printing of their own documents. However, to manage documents printed by other users, you must have the Manage Documents permission, which is one of several printer permissions.

1. Open Printers and Faxes.
2. Double-click the printer you are using to open the print queue.
3. Right-click the document for which you want to restart printing, and then click **Restart**.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open the print queue in Printers and Faxes by clicking a printer, and then clicking **See what's printing** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Pause or resume printing a document
Cancel printing a document
Change the printing priority of documents waiting to print

# Transfer documents to another printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To transfer documents to another printer**

Use this procedure to send your documents to an alternate printer when your printer has failed and your documents are stuck in the print queue. To redirect documents to another printer, you must have the Manage Printers permission for both printers, which is one of several printer permissions. For more information, click **Related Topics**.

Since the documents you send to the printer are converted into that printer's language, you must redirect the documents to the same type of printer, such as a laser, ink-jet, or dot-matrix.

1. Open Printers and Faxes.
2. Double-click the printer that holds the document(s) you want to redirect.
3. On the **Printer** menu, click **Properties**.
4. Click the **Ports** tab, and then do one of the following:

   **To send documents to another printer on the same print server**

   - Click the port to which the other printer is assigned and then click **OK**.

   **To send documents to a printer on a different print server**

   The printer must already exist and be shared on the print server.

   - Click **Add Port**, click **Local Port**, and then click **New Port**. Type the name of the other print server and the share name of the shared printer in the following format:

     *****other_print_server******share_name*

### Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open printer properties in Printers and Faxes by clicking a printer, and then clicking **Set printer properties** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- The document currently printing cannot be transferred to another printer.
- If the printer holding the print jobs is offline or has a paper jam, the document showing the error cannot be moved to a new port.
- If a print job shows an error, all subsequent print jobs will be held until the document is deleted from the print queue.
- For more information on administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Print a document
Assigning printer permissions
Prnport.vbs

# Pause or resume printing all documents

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To pause or resume printing all documents**

To pause, or resume printing on the printer, you must have the Manage Printers permission, which is one of several printer permissions.

1. Open Printers and Faxes.

2. Right-click the printer for which you want to pause or resume printing, and then click **Pause Printing**.

   When you pause printing, the document currently printing will finish, and then no documents will print until you resume printing.

   A check mark appears next to **Pause Printing** when the printer is paused. To clear the check mark and resume printing, click **Pause Printing** again.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.

- You can also open the print queue in Printers and Faxes by clicking a printer, and then clicking **See what's printing** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.

- For more information on administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

View documents waiting to print
Change the printing priority of documents waiting to print
Assigning printer permissions
Prnjobs.vbs

# Cancel printing all documents

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To cancel printing all documents**

To cancel, pause, or resume printing on the printer, you must have the Manage Printers permission, which is one of several printer permissions. For more information, click **Related Topics**.

1. Open Printers and Faxes.
2. Right-click the printer for which you want to cancel printing, and then click **Cancel All Documents**.
3. When prompted to cancel all documents, click **Yes**.

   In some cases, the document currently printing will finish, and then all other documents are cancelled.

## Notes

- To open Printers and Faxes, click **Start**, and then click **Printers and Faxes**.
- You can also open the print queue in Printers and Faxes by clicking a printer, and then clicking **See what's printing** under **Tasks** on the left side of the window. This option is available only if folders are set to look like a Web page and a printer is selected. For more information, click **Related Topics**.
- For more information on administering printers from the command line, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

View documents waiting to print
Pause or resume printing all documents
Assigning printer permissions
Prnjobs.vbs

# Printing Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about planning, implementing, and operating printers and print servers.

- Printing Overview
- Understanding Printing
- Using Printing
- Printing Resources

# Printing Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing overview**

Using products in the Microsoft® Windows Server 2003 family, you can share printing resources across your entire network. Clients on a variety of computers and operating systems can send print jobs to printers attached locally to a print server computer running a Windows Server 2003 family operating system, across the Internet, or to printers connected to the network using internal or external network adapters, or another server.

Products in the Windows Server 2003 family support several advanced printing features. For example, you can administer a print server computer running a Windows Server 2003 family operating system that exists anywhere on your network. Another advanced feature is that you do not have to install a printer driver on a Microsoft® Windows® XP client computer to enable it to use a printer. The driver is downloaded automatically when the client connects to a print server computer running a Windows Server 2003 family operating system.

Although you can, of course, print locally from a server running a product in the Windows Server 2003 family or a computer running Windows XP to a locally attached printer, this Help is primarily concerned with printing across the network with a computer running a Windows Server 2003 family product functioning as a print server.

You should make sure to distinguish between a printer (the device that does the actual printing) and a logical printer (its software interface on the print server). When you initiate a print job, it is spooled on the logical printer, which is also called a printer, before it is sent to the printer itself.

See New printing features for a list of Windows Server 2003 family printing enhancements.

# New printing features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New features**

### Printing setup

For the administrator, products in the Windows Server 2003 family provide improved setup tools for common network configurations. For example, the new standard port (described later in this topic), simplifies installation of most TCP/IP printers on your network.

### Installing printer drivers on all nodes in a cluster

Previously, you had to manually install printer drivers on every node of a cluster. Using products in the Windows Server 2003 family, this is no longer necessary. You can now install printer drivers on one node and the drivers are automatically propagated to the other nodes. The procedure for adding a printer driver to a cluster is the same as the procedure for adding printer drivers to any other computer. For more information on adding printer drivers, see Install new or updated printer drivers.

### Printing administration

The Windows Server 2003 family operating systems provide improved remote administration by adding remote port administration. You can now perform a full remote administration and configuration (including remote port administration) of printers from any computer running a Windows Server 2003 family operating system.

### Printers folder

The **Printers** folder has been renamed to **Printers and Faxes** for all Windows Server 2003 family and Windows XP products. All fax-related tasks are available through this folder.

### Internet printing

The Windows Server 2003 family printing architecture is now seamlessly integrated with the Internet. For the end user, the Windows Server 2003 family offers printing across the Internet.

- You can print from Windows XP clients to print server computers running a Windows Server 2003 family operating system by using a Uniform Resource Locator (URL).
- You can use a browser to manage printers. You can pause, resume, or delete a print job, and view the printer and print job's status.
- You can connect to printers on your network using Web point-and-print for single-click installation of a shared printer. You can click the shared printer to install the required printer driver on your local computer. To facilitate stronger network security, Web point-and-print is not enabled on the print server by default because it is possible to include malicious code in a printer driver. To enable Web point-and-print functionality, in Group Policy, you must enable the Point-and-Print Restrictions policy. For more information, see Set Group Policy for printers

### Directory services

By default, a Windows Server 2003 family server makes all shared printers in the domain available as objects in Active Directory. When you list, or *publish*, shared printers in the directory, you can use search tools to quickly locate the most convenient printing resources. For example, you can search by feature (such as color printing) or by location (such as a specific floor in a building).

### Standard port monitor

The new standard port monitor connects a print server computer running a Windows Server 2003 family operating system to network-interface printers that use the TCP/IP protocol. It replaces LPRMON for TCP/IP printers connected directly to the network or through a network adapter. The new standard port 50 is faster than LPRMON for many print jobs. LPRMON is still required for a printer connected to a UNIX or VAX host.

### Print queue monitoring

You can monitor the performance of a local or a remote printer using the System Monitor's new Print Queue object. Counters can be set up for a variety of performance criteria such as bytes printed per second, job errors, and total pages printed.

### User settings

The ability to change personal document default settings has now been extended to Windows XP Professional clients and products in the Windows Server 2003 family. Previously, only Windows 95 and Windows 98 users could modify their document settings after they were initially set by an administrator.

### Printing from an application

When printing from an application, the standard print dialog box has been enhanced and improved. You can now search for printers in Active Directory, or even add printers.

### Kernel-mode printer driver blocking

Using a Windows Server 2003 family product, administrators can block the installation of kernel-mode printer drivers by using Group Policy. For more information on how to block kernel-mode printer driver installation, see Block the installation of kernel-mode printer drivers.

### Automate printing administration using scripts

Windows Management Instrumentation (WMI) is the management API created by Microsoft to allow all system components to be monitored and controlled, either locally or remotely. The WMI Print provider allows you to manage print servers, print devices, and other printing-related objects from the command line. With WMI Print provider, you can use Visual Basic (VB) scripts to perform administrative printer functions. You can create your own scripts or modify one of the VB scripts that are provided with a Windows Server 2003 family product. Either way, you can manage local and remote printers using these scripts.

# Understanding Printing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding printing**
As a system administrator responsible for a Windows Server 2003 family printing environment, you need to understand the printing process and the options available for print servers and clients to communicate within computers running Windows, with computers running other operating systems, and across the Internet.

This section covers:

- The printing process
- Internet printing
- Printing and print servers
- Printer drivers overview
- Printer fonts overview
- Print processors and data types
- Choosing and configuring a port
- Installing printers with Plug and Play
- Find printers overview
- Printer properties overview
- Assigning printer permissions
- Enabling printer location tracking
- Group Policy and printing

# The printing process

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The printing process**

The following is an overview of the operations performed on a document submitted to a printer from a Windows XP client. The printer is attached to a computer running a Windows Server 2003 family operating system. (Some processes are slightly different for non-Windows brand operating system print clients.)




1. A user on a Windows XP client computer chooses to print a document.

2. If the document is submitted from a Windows application, the application calls the graphics device interface (GDI), which calls the printer driver associated with the target printer. Using the document information from the application, the GDI and the driver exchange data to render the print job in the printer language of the printer, and then pass it to the client-side print spooler. If the client is using an operating system other than Windows, or a non-Windows application, another component replaces the GDI to perform a similar task.

3. The client computer delivers the print job to the print server. For Windows XP, Windows 2000, or Windows NT 4.0 clients, the client-side spooler makes a remote procedure call (RPC) to the server side of the spooler, which uses the print router to poll the remote print provider on the client side. The remote print provider initiates another RPC to the server spooler, which receives the print job over the network.

4. On the print server, print jobs from Windows XP, Windows 2000, or Windows NT 4.0 clients are of the enhanced metafiles (EMF) data type. Many other applications use the RAW (ready to print) data type.

5. The router on the server passes the print job to the local print provider on the server (component of the spooler), which spools the print job (writes it to the disk).

6. The local print provider polls the print processor. The print processor recognizes the job's data type and receives the print job. The print processor then converts the print job according to its data type.

7. If the target printer is defined on the client computer, the print server service decides whether the server's spooler should alter the print job or assign a different data type. The print job then passes to the local print provider, which writes it to the disk.

8. Control of the print job is passed to the separator page processor, which adds a separator page, if specified, to the front of the job.

9. The job is despooled to the print monitors. For bidirectional printers, a language monitor handles the two-way communication between the sender and the printer, and then it passes the print job to a port monitor. If the printer is not bidirectional, the print job goes directly to the port monitor, which sends it to the target printer (or, to another network print server).

10. The printer receives the print job, converts each page into bitmap format, and prints it.

# Internet printing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet printing**
Internet printing provides a way to submit print jobs to a printer across the Internet. For example, you can send your company's new catalog directly to your publisher's printer, providing you have your publisher's printer name and the proper permission. You can also use this technology with your local area network (LAN) or intranet.




- To install a printer from the Internet you use the printer's URL as the name of the printer. Administrators can also use the URL format within an intranet.
- For a computer running a Windows Server 2003 family operating system and Windows XP to process print jobs that contain URLs, it must be running Microsoft Internet Information Services (IIS).
- Internet printing uses Internet Printing Protocol (IPP) as its low-level protocol, which is encapsulated within HTTP, using it as a carrier. When accessing a printer through a browser, the system first attempts to connect using RPC (on Intranets and LANs), which is fast and efficient.
- Print server security is provided by IIS. To support all browsers and all Internet clients, the administrator must choose basic authentication. Alternatively, the administrator can use Microsoft challenge/response or Kerberos authentication, both of which are supported by Microsoft Internet Explorer.
- You can manage printers from any browser, but you must use Internet Explorer version 4.0 or higher to be able to connect to a printer using a browser.

# Printing and print servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing and print servers**

When you add a printer that is attached directly to the network through a network adapter, you can implement printing in two ways: adding a printer directly to each user's computer without the use of a print server computer, or adding the printer once to a print server computer and then connecting each user to the printer through the print server computer.

### Printing without using a print server

For the first method, consider the following scenario: A small workgroup network has only a few computers and a printer that is connected directly to the network. Each user on the network adds the printer to his or her Printers and Faxes folder without sharing the printer, and sets his or her own driver setting.

The disadvantage of this configuration is the users do not know the true state of the printer. Each computer has its own print queue displaying only those print jobs sent from that computer. You cannot determine where your print job is in relation to all the print jobs from other computers. Another drawback is that error messages--such as paper jams or empty paper trays--appear only on the print queue for the current print job. Finally, all the processes on a document submitted for printing are done on that one computer.

### Printing with a print server

The second way to set up printing is to have one computer running a Windows Server 2003 family operating system function as a print server. That computer adds the printer and shares it with the other users. A computer running Windows XP Professional can also function as a print server but with some limitations; it cannot support Macintosh or NetWare services, and it is limited to only 10 connections within the same local area network (LAN).

Printing with a print server provides the following advantages:

- The print server manages the printer driver settings.
- A single print queue appears on every computer connected to the printer, letting each user see where his or her print job is in relation to others waiting to print.
- Because error messages appear on all computers, everyone knows the true state of the printer.
- Some processing is passed from the client computer to the print server.
- You can have a single log for administrators wanting to audit the printer events.

The only disadvantage of using a print server is that it requires a computer to function as the print server. However, it does not need to be a dedicated computer; typically, print servers are implemented on servers that also perform other duties.

# Printer drivers overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printer drivers overview**
A printer driver is software used by computer programs to communicate with printers and plotters. Printer drivers translate the information you send from the computer into commands that the printer understands. Usually, printer drivers are not compatible across platforms, so various drivers must be installed on the print server to support different hardware and operating systems. For example, if you are running Windows XP and you share a printer with users running Windows 3.1, you might need to install multiple printer drivers.

In general, printer drivers are made up of three types of files:

- Configuration or printer interface file
- Data file
- Printer graphics driver file

**Configuration or printer interface file**
Displays the **Properties** and **Preferences** dialog boxes when you configure a printer. This file has a .dll extension.

**Data file**
Provides information about the capabilities of a specific printer, including its resolution capability, whether it can print on both sides of the page, and what size paper it can accept. This file could have a .dll, .pcd, .gpd, or .ppd extension.

**Printer graphics driver file**
Translates device driver interface (DDI) commands into commands that a printer can understand. Each driver translates a different printer language. For example, the file Pscript.dll translates the PostScript printer language. This file has a .dll extension.

These files, which are usually accompanied by a Help file, work together to make printing possible. For example, when you install a new printer, the configuration file reads the data file and displays the available printer options. When you print, the graphics driver file queries the configuration file about your selections so that it can create the proper printer commands.

For information about installing printer drivers, see Install new or updated printer drivers.

# Printer fonts overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printer fonts overview**

Most programs that support printing allow you to choose from a number of different fonts. Printers offering the widest assortment of fonts include laser, ink-jet, and dot-matrix printers. They support a fixed set of resident fonts, which you can expand by installing font cartridges or loading fonts from software.

- Internal fonts (resident fonts) are used most often by laser, dot-matrix, and ink-jet printers. Pre-loaded into the printer's memory (ROM), internal fonts are always available for printing.
- Cartridge fonts are stored in a cartridge or a card that plugs into the printer.
- Downloadable fonts are sets of characters sent from the computer to a printer's memory when needed for printing. Also called soft fonts, downloadable fonts are used most often with laser printers and other page printers, although some dot-matrix printers also use them. Clients that use downloadable fonts and that print to Windows print servers should install the fonts locally to increase the printing speed.

For each document you print, Windows may need to send the required screen and downloadable fonts to the printer. To increase printing speed, use fonts that do not need to be downloaded (such as internal fonts or cartridge fonts). Some printers have the ability to keep a list of downloadable fonts. If your printer supports this feature, enable it to increase printing speed. Not all printers can use all three types of fonts. Pen plotters, for example, cannot normally use downloadable fonts. For more information on the types of fonts you can use, see your printer's operating manual.

For information about installing fonts, see Install or remove external PCL soft fonts.

# Print processors and data types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Print processors and data types**

The print processor tells the spooler to alter a job according to the document data type. It works in conjunction with the printer driver to send the spooled print jobs from the hard drive to the printer.

Software vendors occasionally develop their own print processors to support custom data types. Normally, the print processor does not require any settings or intervention from administrators.

**Data types**

The Windows printing process normally supports five data types. The two most commonly used data types, enhanced metafile (EMF) and ready to print (RAW), affect performance in different ways on both the client computer and the print server computer.

**EMF**

EMF, or enhanced metafile, is the default datatype with most Windows-based programs. With EMF, the printed document is altered into a metafile format that is more portable than RAW files and usually can be printed on any printer. EMF files tend to be smaller than RAW files that contain the same print job. Regarding performance, only the first portion of a print job is altered, or rendered on the client computer, but most of the impact is on the print server computer, which also helps the application on the client computer to return control to the user faster.

**RAW**

RAW is the default data type for clients other than Windows-based programs. The RAW data type tells the spooler not to alter the print job at all prior to printing. With this data type, the entire process of preparing the print job is done on the client computer.

The following table shows the five different data types supported by the default Windows print processor.

| Data type | Directions to spooler | Use |
|---|---|---|
| | | |
| RAW | Print the document with no changes. | This is the data type for all clients not based on Windows. |
| RAW [FF appended] | Append a form-feed character (0x0C), but make no other changes. (A PCL printer omits the document's last page if there is no trailing form-feed.) | Required for some applications. It can be set as the default in the **Print Processor** dialog box. |
| RAW [FF auto] | Check for a trailing form-feed and add one if it is not already there, but make no other changes. | Required for some applications. It can be set as the default in the **Print Processor** dialog box. |
| EMF 1.00x | Treat the document as an enhanced metafile (EMF) rather than the RAW data that the printer driver puts out. | EMF documents are created by products in the Windows Server 2003 family. |
| TEXT | Treat the entire job as ANSI text and add print specifications using the print device's factory defaults. | This is useful when the print job is simple text and the target print device cannot interpret simple text. |

# Choosing and configuring a port

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing and configuring a port**

The term print monitor is used to describe two types of print monitors: language monitor and port monitor. (See the The printing process to see how the print monitor fits into the printing process.) A language monitor is needed if the printer supports bidirectional printing, a two-way communication between the printer and the spooler running on the print server. The spooler then can access configuration and status information from the printer. The more important monitor is the port monitor, which controls the I/O port to the printer. A printer port is an interface through which a printer can communicate with the computer. You can add printer ports using the Add Printer Wizard while adding a printer, from the **Ports** tab in the printer's property page, or through the print server property page. See Set Ports and Protocols for the complete procedures on ports.

The process of selecting a port and the subsequent options that are displayed depends on how your printer is connected to the server or network and what software (including protocols) you have installed.

If the printer is physically attached to the print server, select the appropriate local port. LPT1 through LPT3 represent parallel ports; COM1 through COM4 represent serial ports. When a client prints to a printer port denoted as FILE, the client is prompted for the file name, and the output file is stored on the client computer. If you decide to add a new local port, you can enter one of the following:

- A file name, such as **C:***dir**filename*. All jobs sent to this port are written to the named file, and each new job overwrites the last one.
- The share name of a printer, such as *****server**printer* (URLs are not accepted). Jobs sent to this port are transferred over the network to the named share by the network redirector.
- NUL. This specifies the null port, which you can use to test whether network clients can send jobs. Jobs sent to NUL are deleted without wasting paper or delaying real print jobs.
- IR. Use this port to connect to infrared-enabled printers meeting Infrared Data Association (IrDA) specifications. If your hardware does not support IR, it will not be listed on the **Ports** tab.

Certain ports are not listed in the **Ports** tab unless a printer that requires one of them is installed. Universal Serial Bus (USB) and IEEE 1394 printers are Plug and Play compatible; when you plug a printer into the correct physical port (and USB or IEEE 1394) enabled in BIOS, the correct port monitor is installed automatically. Your operating system detects the device, displays its settings on the screen, and prompts you to confirm. You might be required to insert a CD-ROM containing driver files, as described in Add a local port.

If the printer is attached directly to the network, add a port when prompted by the Add Printer Wizard for a port selection. The following table shows the port options available and explains under what conditions the port is available. You must have administrative credentials to add a port.

| Additional port | Enables clients to print to | Available |
|---|---|---|
| Local port | Printers connected to a parallel port, serial port, specific file name, universal naming convention (UNC) name, or the NUL port. | By default; see Add a local port. |
| Standard TCP/IP port | TCP/IP printers connected directly to the network. | By default; see Add a standard TCP/IP port. |
| AppleTalk printing devices | AppleTalk printers. | When AppleTalk protocol is installed; see Add an AppleTalk printing device port. |

| | | |
|---|---|---|
| LPR port | TCP/IP printers connected to a UNIX (or VAX) server. | When Print Services for Unix is installed; see Add an LPR port. |
| Port for NetWare | NetWare printing resources. | When NWLink protocol and Client Services for NetWare are installed. |

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

The port you need might not appear in the wizard if its prerequisite network protocol or services are not installed. To add a network protocol or service, see Add a network component.

**Print Services for Unix**

With UNIX, a program on one computer can use a Line Printer Remote (LPR) service to send a document to a print spooler service on another computer. The receiving service is usually called a Line Printer Daemon (LPD).

Print Services for Unix provides both LPR and LPD services, which act independently on a print server computer running a Windows Server 2003 family operating system. LPD is run on the print server computer to allow the server to accept jobs from LPR clients such as UNIX systems. LPR is an option for computers running Windows XP Professional to submit jobs to UNIX servers running LPD.

- LPDSVC on the print server receives documents from native LPR utilities running on client UNIX computers.

- LPRMON on the print server sends print jobs to native LPD processes on UNIX computers to which printers are attached.

By default, the LPD service is set to start manually. To have it start automatically, change the startup options for the TCP/IP print server service in Services. You need a Windows Server 2003 family operating system compact disc to install Print Services for Unix on a computer running a Windows Server 2003 family operating system. You need the Windows XP Professional compact disk to install the services on a computer running Windows XP Professional.

# Installing printers with Plug and Play

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing printers with Plug and Play**
Plug and Play provides automatic detection and quick installation of many hardware devices, including printers. Nearly every printer on the market today is a Plug and Play printer. However, the way printers are detected and the level of user intervention required for the installation process differs, and depends on the way the printer is connected to the computer.

**Universal serial bus (USB) and IEEE 1394 printers**
Printers using the latest connector technology, such as printers with a universal serial bus (USB) port or IEEE 1394 port, are detected immediately. When you insert the jack into the port, your operating system detects the printer and starts the installation process without any user intervention. Note that a USB or 1394 port is installed automatically upon detection of such device. If no such device is connected to your computer, these ports are not listed in the existing port list.

**Parallel (LPT port) printers**
Printers that connect with a parallel port (LPT port) cannot be detected upon the physical connection. After you physically connect the parallel cable and verify that the printer is turned on and ready to print, you need to add the printer with the Add Printer Wizard, making sure that the **Automatically detect and install my Plug and Play printer** check box is selected. This starts the Plug and Play detection system. Or, you can restart your system to cause the startup process to detect the printer and start the New Found Hardware Wizard.

### Note

- If you add a parallel printer with the Add Printer Wizard while the **Automatically detect and install my Plug and Play printer** check box is cleared, the next time you restart your computer (even after several days or weeks working with that printer) Windows detects a new hardware device and starts the New Found Hardware Wizard to add that printer.

**Infrared-enabled (IR port) printers**
Infrared-enabled printers are detected in the same manner as printers with USB or IEEE 1394: automatically upon physical connection. You should place the printer approximately one yard (one meter) away from the infrared transceiver for it to be detected.

Note that the IR port, like a USB port, is not present unless Windows detects and installs it. However, if your system is not infrared-enabled and you add an infrared transceiver, you need to add the IR port manually.

**Printers that are not Plug and Play**
The following printers (as well as other devices that connect the same way) are not Plug and Play and are not detected or installed automatically. These printers must be installed using the Add Printer Wizard:

- Printers connected through a serial (COM) port.
- Printers that are connected directly to the network with a network adapter.

For more information on using the Add Printer Wizard, see Add a printer attached to your computer.

# Printer properties overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printer properties overview**

You use a printer's properties pages, or tabs, to set and customize a printer's behavior, provided you have the appropriate printer permission. The printer's properties tabs include a standard set of tabs that always appear for all printers, and optional tabs that might appear based on the printer's driver. For example, color printers have an additional tab called **Color Management**.

The first table summarizes the standard set of the printer's properties tabs, which always appear. The second table shows some examples of optional printer's properties tabs that might appear.

| Tab | Use to |
|---|---|
| | |
| General | Change printer's name<br>Set printer's location and comment<br>Print a test page<br>Set personal printing preferences |
| Sharing | Share a printer, or stop sharing<br>Publish a printer in Active Directory<br>Install additional drivers |
| Ports | Add, delete, and configure a port<br>Enable printer pooling |
| Advanced | Schedule printer availability<br>Assign printer priority<br>Change printer's driver<br>Set spooling options<br>Set printer's default preferences<br>Set a separator page |
| Security | Assign permissions for a printer<br>Audit printer use<br>Take ownership of a printer |
| Device Settings | Set printer memory<br>Assign a print form to a tray<br>Choose font types<br>Set what optional features are installed |

Here are some examples of optional properties tabs that might appear depending on the type of printer (as indicated by its driver), the situations in which they appear, and what you use them to do. Other optional tabs not listed in the table might also appear.

| Tab | Appears with | Use to |
| --- | --- | --- |
| | | |
| Color Management | Color printers | Set color profile |
| About | Some drivers | Printer manufacturer information (when manufacturer assists with driver) |
| Utilities | Some drivers | Printer manufacturer utilities |
| User Information | Fax drivers | Information to be used on the fax cover page |
| Printer Commands | Generic printers (text only) | Set control codes for a generic printer |
| Font Selection | Generic printers (text only) | Set control codes for fonts |

# Assigning printer permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Assigning printer permissions**
When a printer is installed on a network, default printer permissions are assigned that allow all users to print, and allow select groups to manage the printer, the documents sent to it, or both. Because the printer is available to all users on the network, you might want to limit access for some users by assigning specific printer permissions. For example, you could give all nonadministrative users in a department the Print permission and give all managers the Print and Manage Documents permissions. In this way, all users and managers can print documents, but managers can also change the print status of any document sent to the printer.

Windowsprovides three levels of printing security permissions: Print, Manage Printers, and Manage Documents. When multiple permissions are assigned to a group of users, the least restrictive permissions apply. However, when Deny is applied, it takes precedence over any permission. The following is a brief explanation of the types of tasks a user can perform at each permission level.

**Print**
The user can connect to a printer and send documents to the printer. By default, the Print permission is assigned to all members of the Everyone group.

**Manage Printers**
The user can perform the tasks associated with the Print permission and has complete administrative control of the printer. The user can pause and restart the printer, change spooler settings, share a printer, adjust printer permissions, and change printer properties. By default, the Manage Printers permission is assigned to members of the Administrators and Power Users groups.

By default, members of the Administrators and Power Users groups have full access, which means that the users are assigned the Print, Manage Documents, and Manage Printers permissions.

**Manage Documents**
The user can pause, resume, restart, cancel, and rearrange the order of documents submitted by all other users. The user cannot, however, send documents to the printer or control the status of the printer. By default, the Manage Documents permission is assigned to members of the Creator Owner group.

When a user is assigned the Manage Documents permission, the user cannot access existing documents currently waiting to print. The permission will only apply to documents sent to the printer after the permission is assigned to the user.

**Deny**
All of the preceding permissions are denied for the printer. When access is denied, the user cannot use or manage the printer or adjust any of the permissions.

**Printing permissions assigned to groups**
Windowsassigns printer permissions to six groups of users. These groups include Administrators, Creator Owner, Everyone, Power Users, Print Operators, and Server Operators. By default, each group is assigned a combination of the Print, Manage Documents, and Manage Printers permissions as shown in the following table.

| Group | Print | Manage Documents | Manage Printers |
|---|---|---|---|
| Administrators | X | X | X |
| Creator Owner | | X | |
| Everyone | X | | |
| Power Users | X | X | X |
| Print Operators | X | X | X |
| Server Operators | X | X | X |

The Print Operators and Server Operators groups are located only on domain controllers.

| |
|---|
| Members of this group can manage, create, share, and delete printers and print queues. Members of this group can load and unload device drivers on the server. Users who can load and unload device drivers also have the ability to load malicious code on the server. As a security best practice, only add trusted users to this group. |

Each permission consists of a group of special rights that allow the user to perform specific tasks. The following table summarizes the level of access associated with each of the printing security permissions.

| **Tasks permitted** | **Print** | **Manage Documents (applies to documents only)** | **Manage Printers** |
|---|---|---|---|
| Print | X | | X |
| Manage Printers | | | X |
| Manage Documents | | X | |
| Read Permissions | X | X | X |
| Change Permissions | | X | X |
| Take Ownership | | X | X |

For more information, see Set Group Policy for printers.

# Enabling printer location tracking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling printer location tracking**

In network environments, most users do not need to know the physical location of files, databases, and Web sites. However, they do need to know the physical location of hardware devices such as printers, scanners, copiers, and facsimile (fax) machines. The printer location tracking feature helps users find hardware devices by storing their locations in Active Directory. When this feature is enabled, the location string is displayed automatically for each printer.

Your network environment must have the following to be able to use this feature:

- A working directory service with more than one site or more than one subnet.
- A network IP addressing scheme that corresponds roughly to the physical layout of your enterprise.
- Client computers that can query the directory service.
- A subnet object for each site. You can create them using Active Directory Sites and Services.
- Administrative credentials.

If these requirements are met, you need to do the following to enable this feature:

1. **Establish a naming convention for your organization.**

   For more information, see Establishing a naming convention for printer locations.

2. **Create a subnet object in the directory service.**

   For location tracking to work, you must have more than one subnet or site. You also need a subnet object for each subnet that a client computer might use. The naming of subnets is defined by the directory service. For more information, see Create a subnet.

   Note that to set and manage the directory service site and subnet you need to use Active Directory Sites and Services.

3. **Set the location attribute of the sites and subnets.**

   After you have created the site and subnet objects, you need to set the location string for each subnet. Use the **Location** tab of the subnet property page to set the location string. Note that the directory service site name does not have to be the same as the location name.

4. **Enable location tracking.**

   You need to use Group Policy to enable location tracking. For information on how to do this, see Enable printer location tracking.

5. **Set the location of each printer.**

   Add the location for each printer in each printer's property page. You can either type the location attribute, or you can click to browse for it.

   Consider adding more information about the location of the printer, such as its room number. For information on how to do this, see Enable printer location tracking.

6. **Test your settings.**

Always test your settings by trying them on several printers, and on different print servers from different locations.

**How an end user uses location tracking**

Location tracking is used whenever a user queries the directory service. To start the query, click **Start**, click **Search**, and then click **Find Printers**. You can also open the query screen from the Add Printer Wizard while connecting to a printer connection, by clicking **Find a printer in the Directory**. If location tracking is enabled, the system first determines where the client computer is physically located in the organization. While this is in progress, the location field of the query form displays **Checking**. After the location has been determined, it is entered into **Location**. If the location cannot be determined, **Location** is left blank.

When the user clicks **Find Now**, Active Directory lists all printers matching the user's query that are located in the user location. Users can change the location field by clicking **Browse** to search for printers in other locations. For example, suppose a company is located in a building with several floors, where each floor is configured as a subnet. If a user located in the first floor fails to locate a color printer in **Floor 1**, then the user can change the location to **Organization 1/Floor 2** or even to **Organization 1** to increase the scope of the search.

Note that the location field is not automatically available for users running Windows 95, Windows 98, or Windows NT 4.0 without a directory service client. However, they can still query the directory service for printers in a specific location, if they know the exact location attribute and type it in the location field.

# Group Policy and printing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group Policy and printing**

Administrators use Group Policy to define the computer environment that is provided to a group of users and computers, including the applications that are available to them, the applications that appear on their desktop and on the **Start** menu options, who can modify their desktop and who cannot, and so on. For more information about Group Policy see Group Policy overview.

Administrators can also use Group Policy to change the default behavior of the printing environment and to provide computers and users a standard set of preferences. For example, an administrator can restrict some groups or users from adding or deleting printers, or prevent everyone from using Internet printing or publishing their printers in Active Directory.

The following list shows examples of setting policies on printers.

- To disable Web-based printing for every client on the print server:

  Open Group Policy. Expand the following: Computer Configuration, Administrative Templates, Printers. Select **Disable Web-based Printing** to enable the policy.

- To enable Web-based printing for every client on the print server, except Client1:

  Select **Disable Web-based Printing** to disable the policy for everyone on the print server. Then, add the policy as a snap-in for Client1 as the Group Policy object. Here, clear **Disable Web-based Printing** to enable the policy.

Other usages of Group Policy for printing are as follows:

- Administrators can change the default behavior of the Add Printer Wizard. For more information, see Add a printer attached to your computer, or Add a printer attached directly to the network.

- Administrators can change the way printers are published in Active Directory. For more information, see Publish a printer in Active Directory, and Manually publish a printer in Active Directory

- Administrators can disable Internet printing. For more information, see Manage printers from a browser, and Connect to a printer with a browser

Several other Group Policy settings change the default behavior of printing. For information on these settings, click **Explain** on the setting's dialog box. For more information about Group Policy and printing, see Set Group Policy for printers

# Using Printing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using printing**

As with any other operating system, building an effective printing environment with the Windows Server 2003 family requires careful planning. Many variables can affect the performance of your printing resources and the satisfaction of users. The wide variety of printing equipment and strategies for deploying it add further considerations for decision-making.

This section describes a number of different printing scenarios and provides information on fine-tuning your printing performance.

- Planning for Printing
- Administering Printing

# Planning for Printing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for printing**

As administrator of a computer running a member of the Windows Server 2003 family, you first need to examine your printing workload and estimate the capacity needed to satisfy demand under varying conditions. You also need to establish a naming convention that will ease installation, usage, and support of your printing environment.

Next, you need to determine how many print servers your network requires, and how many printers you need to allocate to each server. Finally, you must decide which printers to buy, which computers to use as print servers, where to locate printers, and how to manage printer traffic.

This section covers:

- Choosing a printer name
- Establishing a naming convention for printers
- Establishing a naming convention for printer locations
- Sizing and choosing printers
- Sizing and choosing print servers

# Choosing a printer name

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing a printer name**

Windows XP and Windows Server 2003 family products support the use of long printer names. This allows you to create printer names that contain spaces and special characters. However, if you share a printer over a network, some clients will not recognize or correctly handle the long names, and users may experience problems printing. Also, some programs cannot print to printers with names longer than 32 characters.

For shared printers, the entire qualified name (including the server name, for example \\PRINTER2\PSCRIPT) must be 32 or fewer characters.

- If you share a printer with a variety of clients on a network, use 32 or fewer characters for printer names, and do not include spaces or special characters in these names.

- If you share a printer with MS-DOS computers, do not use more than eight characters for the printer's share name. You can lengthen the name by adding a period followed by no more than three characters, but you cannot use spaces in the name.

Several Windows version 3.*x* programs cannot print to a printer if its name contains more than the specified number of characters and will generate an Access Violation or other message if printing is attempted. Other programs may not print to any printer, even printers with short names, if the default printer's name is too long. To resolve these problems, rename the printers used by these programs with shorter names and make one of the renamed printers the default printer.

# Establishing a naming convention for printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing a naming convention for printers**

Setting a standard for printer names is a well-accepted practice, and many organizations establish it early in their printing planning. However, in the Windows Server 2003 family environment, the printer location tracking feature shifts the importance of naming standards from printer names to location names. When printer location tracking is enabled, setting a solid location name standard is imperative to its successful implementation. The feature also allows for much more liberty in naming printers. For more information about setting printer location names, see Establishing a naming convention for printer locations.

However, printer location tracking is not suited for every organization, and it is disabled by default. If you do not enable printer location tracking, then you should carefully plan your organization's naming standard for printers. You need to create a simple, easy-to-implement naming standard that fits your own printer deployment plan.

For example, if you have two identical printers in the mail room, your naming convention needs to provide a way to identify each of the printers, perhaps by adding a number at the end of each printer name. You also want a naming standard that provides users with information about the printer. Therefore, to continue with the previous example, you might name the printers Printer-01, and Printer-02. The naming standard might also tell users what kind of printed output to expect. You could call a color laser printer ColorLaser-01, and a color inkjet printer ColorInk-01.

If you do not enable printer location tracking, consider adding the printer location as part of your printer naming standard. However, remember that printers are likely to move, which will require you to change their names.

For other considerations and issues with setting printer names, see Choosing a printer name.

# Establishing a naming convention for printer locations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Establishing a naming convention for printer locations**

To use printer location tracking, described in Enable printer location tracking, you need to set a naming convention for printers' locations using the following rules:

- Location names are in the form *name***/***name***/***name***/***name***/**... (The slash (**/**) must be the dividing character.)
- A name can consist of any characters except for the slash (**/**).
- The number of levels to a name is limited to 256.
- The maximum length of *name* is 32 characters.
- The maximum length of an entire location name is 260 characters.

Because location names are used by end users, they should be simple and easy to recognize. Avoid using special names known only to facilities management. For better readability, avoid using special characters in a name, and keep names to a maximum of 32 characters to make sure that the whole name string is visible in the user interface.

**Example:**

Consider an organization with multiple locations in North America and Europe. The naming scheme for the organization might be as follows:

| **Top level** | **Level** | **Level** | **Level** |
|---|---|---|---|
| | | | |
| NorthAmerica | | | |
| | NewYork | | |
| | | HQ | |
| | | | Floor1 |
| | | | Floor2 |
| | | Buffalo | |
| | | | Design |
| | | | Plant |
| | | | Admin |
| | Vancouver | | |
| | | Building1 | |
| | | Building2 | |
| | Dallas | | |
| Europe | | | |
| | Paris | | |

| | | | |
|---|---|---|---|
| | | Design | |
| | London | | |
| | | ... | |
| | Madrid | | |

Note that the tree varies in depth depending on the complexity of the organization and the amount of detail available in the IP network. The naming convention for this example includes more levels than for a smaller organization located in a single city, or a single building. The full name of the design building in Buffalo, New York is NorthAmerica/NewYork/Buffalo/Design; for the design building in Paris, the full name is Europe/Paris/Design, and the smaller office in Dallas has the name as NorthAmerica/Dallas.

**Setting standards when printer location tracking is disabled**

As mentioned previously, you can add information to **Location** in the **General** tab of the printer property page even if printer location tracking is disabled. However, this has the potential to make printers difficult for users to locate. When users search for printers on the tenth floor, they need to know exactly what to type in the **Find Printers** property page when they click **Start**, then **Search**, and then **Find Printers**. When printer location tracking is enabled, the **Location** field in the **Find printer** property page is filled in automatically.

# Sizing and choosing printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sizing and choosing printers**

When you plan the printing strategy for your organization, you need to estimate the number and type of printers you need both now and in the future. Consider the following:

- Determine how printing resources will be divided and allocated. High-volume printers generally have more features but affect more users if they break down.
- Consider the printer features you need, such as color, duplex printing, envelope feeders, multi-bin mailboxes, internal disks, and staples. Determine who needs these features and where these users are physically located.
- Although cost is a factor, laser printers might be a good choice for both black-and-white and color printing. However, many less expensive laser printers do not support larger page size, which you might require.
- Generally, you will experience fewer maintenance problems if you match printing volume with a printer's duty cycle (the number of pages the printer was designed to print per month).
- Consider what type of graphics your organization needs. Windows, TrueType, and other technologies make it possible to print complex and sophisticated graphics and fonts on most printers.
- Consider printing speed requirements. Generally, printers that attach directly to the network with network adapters offer faster throughput than printers attached with parallel buses. However, print throughput rates also depend on network traffic, the type of network adapter, and the protocol used, not just the type of the printer.
- To confirm that your printer is designed for this operating system, see the compatibility information in Support resources.

**Deciding where to locate printers**

You want to place printers close to the people who will be using them. However, you also need to determine printers' locations in relation to the location of the print server and users' computers within your network. Another goal should be to minimize the impact of printing on your network environment.

Check your network infrastructure and try to prevent print jobs from hopping through multiple internetwork devices. If you have a group of users with high-volume printing needs, you can isolate them with their own printer located on their segment of the network, minimizing the effect on other users.

# Sizing and choosing print servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sizing and choosing print servers**

Consider the following when sizing and choosing your print servers:

- A print server can be any computer running a product in the Windows Server 2003 family or Windows XP Professional.
- A print server running Windows XP Professional is limited to 10 concurrent connections from other computers and does not support printing from Macintosh and NetWare clients.
- The minimum requirements for a Windows Server 2003 family server for file and print services will be sufficient for a print server that manages just a few printers with low printing throughput. Managing a large number of printers or many large documents requires more memory, disk space, and potentially a more powerful computer.
- In situations where many large documents are submitted for printing at the same time, the print server must have enough disk space to spool all of the documents. If you need to keep a copy of everything you print, you must provide additional disk space.
- If you use the Windows Server 2003 family server for both file and print sharing, file operations have a higher priority, so printing will not slow access to files. That is especially true if the printers are connected directly to the server; parallel and serial ports are always the primary bottleneck.
- Printers that attach directly to the print server, typically through the parallel port, require more CPU time. It is always preferable to use printers that are attached directly to the network with a network adapter.
- To maximize print server throughput, and to manage many printers, or printing of many large documents, you should dedicate a Windows Server 2003 family server only for printing.
- To increase the throughput of a print server, you can change the location of the print spool folder. For best results, move the spool folder to a dedicated disk drive that does not share any files, including page files with the operating system.

# Administering Printing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering printing**
You can administer printers and print servers locally or over the network. Administrative tasks include managing documents in a print queue, managing individual printers on print servers, and managing the print servers themselves.

This section covers:

- Managing documents
- Managing printers
- Managing printer drivers
- Setting document printing defaults
- Connecting clients to a printer
- Setting device-specific properties
- Using separator pages
- Using plotters
- Printing extended or international characters
- Managing printing from the command line
- Managing Printing Traffic

# Managing documents

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing documents**

When managing a document, you can:

- Pause or resume printing.
- Restart the document from the beginning.
- Delete a document.
- View and change various document settings, such as the document priority and the person notified when the document is done. You can also view--but not change--the form type, paper source, page orientation, and number of copies. These options are available only to Windows XP and Windows 2000 clients using a printer on the network.

For information on managing the documents in a print queue, see Manage Documents.

# Managing printers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing printers**
When you manage printers, your actions affect the entire printer, not individual documents. Typical tasks include pausing, resuming, and purging the queue for a printer.

**Redirecting documents**
If a printer becomes faulty, you might need to send documents in its queue to another printer. This prevents users from having to resubmit print jobs that are already in the queue. You can redirect documents to a printer on the same print server or on a different print server. However, both printers must use the same printer driver.

To send documents to a different printer, see Transfer documents to another printer.

**Administering printers on the network**
Any network user can check the status of a printer from any computer across the network. However, only users who have Manage Printers or Manage Documents permission for the printer can manage documents other than their own. If you do not have the correct permission, certain options are unavailable. Also, when you attempt to view certain properties on computers running Windows NT 4.0 or earlier, an error message appears.

You can also manage printers from a Web browser, if Internet printing is enabled.

For more information about accessing a printer on the network, see Remotely manage printers. For more information about managing a printer from a Web browser, see Manage printers from a browser.

**Printers and Active Directory**
The easiest way to manage, locate, and connect to printers is through Active Directory. Users can also search for specific printers--for example, all color printers in building 10--and even save search results for future use. Administrators can create custom searches for different groups of users, and then save the result files on users' desktops, or at any other location. For more information on how to do custom searches for printers, see Create a custom printer search.

By default, all printers that are added to a computer running a Windows Server 2003 family operating system and shared are also published in Active Directory. Administrators can also manually publish printers, as described in Manually publish a printer in Active Directory.

**Auditing a printer**
To increase the level of security of some specific printers, or as a tool for troubleshooting, you can audit a printer's activity. For more information about auditing a printer, see Track printer usage.

**Taking ownership of a printer**
To change print administrators (for example, if the current print administrator left your organization) you would need to take ownership of the printer in order to manage it. For more information about taking ownership of a printer, see Take ownership of a printer.

**Setting Group Policy for printers**
If you want to define a standard printing environment for your organization, you can use a Group Policy setting for groups of users. Each new user that joins the organization has the same credentials as other users in that group without any intervention from the administrator. For more information on Group Policy settings, see Set Group Policy for printers.

# Managing printer drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing printer drivers**

If your network contains a mixture of Alpha-based and x86-based computers, you can install printer drivers for each one on each print server. This ensures that documents originating from Windows XP, Windows 2000, Windows NT 4.0, Windows 95, or Windows 98 clients running on any of the hardware platforms can use all printers.

For example, suppose your network contains the following: Windows 95 clients, x86-based clients running Windows XP or Windows 2000 and Windows NT 4.0, and Alpha-based clients running Windows NT Workstation 4.0 and 3.51. If you create a shared printer on an x86-based computer, you should install these four printer drivers in addition to the x86-based Windows Server 2003 family printer driver that is installed by default for the printer you have selected:

- Windows 95.
- Windows NT 4.0 x86-based.
- Windows NT 4.0 Alpha-based.
- Windows NT 3.5 or 3.51 Alpha-based.

When installed on a computer which is acting as a print server, the Windows Server 2003 family operating system determines whether incoming print requests are Alpha-based or x86-based and automatically sends the appropriate driver to the client.

# Setting document printing defaults

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting document printing defaults**
Device-specific properties such as resolution and memory specify a printer's hardware characteristics. Document printing defaults describe how that hardware can be used to perform a printing task. Unless you override them, a program uses the default document properties that are set for a printer.

Typical document default properties include:

- Page orientation
- One-sided or two-sided printing
- Paper source
- Type of form (paper size)

You can set other preferences as well. To view the document printing preferences, see Set Printing Preferences.

**Setting default printing preferences**
As administrator, you can set the initial defaults for a printer. When a user connects to the printer, the user's defaults are set to these global settings. For more information on setting the printing default preferences, see Set printing defaults.

**Setting personal printing preferences**
Individual users can change default document properties and set their own personal printing preferences. These preferences include the same options as the default printing preferences except that they apply only to that user instead of to the printer. As with Windows NT 4.0, a user can also change the printing defaults from the print dialog box of the application. In such a case, the new defaults apply when a program prints. Note that changes made by an administrator after a user has established a connection to a printer have no effect on that user. For more information on setting personal printing preferences, see Set personal printing preferences.

# Connecting clients to a printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connecting clients to a printer**
After you have added a printer to a computer running a Windows Server 2003 family operating system, you can set up clients to access it. The tasks that you perform to make sure users can print depend on the operating system used by the client computers on your network.

Windows XP**, Windows 2000, Windows NT, Windows 95, and Windows 98 clients**

For Windows XP, Windows 2000, Windows NT 4.0, Windows 95, and Windows 98 clients, you do not need to do anything further. The users connect to the shared printer by clicking **Connection to this printer** in the Add Printer Wizard. Or, if Active Directory is enabled, they can just search for the printer they want to add, right-click it, and then click **Connect**. In both cases, the printer driver is automatically copied to the client computer.

Windows 2000 and Windows NT 4.0 clients check the printer driver and printer configuration each time they connect. Windows NT 3.*x* clients check each time the spooler service on the client is started. If the driver is not current, a copy of the new driver is downloaded automatically. The printer driver for a client running Windows 95 or Windows 98 is not automatically kept current. If you update the driver on the print server, you must manually install the driver on clients running Windows 95 or Windows 98. However, administrators should update the drivers on the print server, making it easier for these clients to update their drivers manually from the print server instead of requiring floppy disks or CD-ROMs. For more information, see Connect to a printer on a network, or Connect Windows NT 4.0, Windows 95, and Windows 98 clients.

**Windows 3.*x* , Windows for Workgroups, and MS-DOS clients**

MS-DOS and all the different versions of previous Windows operating systems support printing from MS-DOS-based applications and Windows-based applications, by using 16-bit printer drivers that must be installed on each computer. Typically, they send jobs to the print server by using one of the following network client redirectors:

- LAN Manager
- Microsoft Network Client 3.0 for Windows
- Windows for Workgroups built-in redirector

The network redirector can send jobs using NWLink or TCP/IP. The print server receives the job through the server service, which typically does not alter the print job. For more information, see Connect MS-DOS, Windows 3.x, and OS/2 clients.

**NetWare Clients**

NetWare clients can install network client software on each client computer to provide direct access to NetWare servers, and to Windows Server 2003 family servers and printers. For example, on Windows 95 and Windows 98, clients that are connected only to the NetWare server need to add the Client for Microsoft Networks to add connectivity to computers running Windows Server 2003 operating systems.

**Note**

- The IPX/SPX protocol is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

**Macintosh Clients**

Macintosh clients send print jobs over AppleTalk to a print server computer running a Windows Server 2003 family operating system. Print Server for Macintosh must be installed on the print server.

**Note**

- Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

### UNIX Clients

UNIX clients can connect to a print server computer running a Windows Server 2003 family operating system, by using the Line Printer Remote (LPR) monitor. This might be problematic because not all UNIX systems support LPR specifications, and different versions of UNIX use different LPR command syntax. For more information, see Print from an LPR client to a printer.

# Setting device-specific properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting device-specific properties**
Device-specific properties describe the physical configuration of a printer, such as how much memory it has, which paper trays are loaded, and so on. When you add a new printer, make sure the device-specific properties match the settings of the printer. Although default settings work for many printing needs, some special options, such as those available with PostScript printer drivers, require specific settings.

**Setting installable options**
As an administrator, you have the ability to specify which options are actually installed on the printer. Many printers today support options such as duplex units (to allow printing on both sides of a paper), internal hard disks, envelope feeders, staplers, multi-bin mailbox trays, and others.

To set the options available to users, see Set installable options for a printer.

**Setting printer memory**
Because page printers must store an entire page in memory, they require relatively large amounts of memory. If you are using a page printer, such as a laser printer, make sure that the amount of memory available in the device matches the value shown in the **Device Settings** tab. If the printer has substantially more or less memory than what is shown in the **Device Settings** tab, print throughput can suffer. For example, your operating system might try to download more fonts to the printer than it can reasonably handle. (Running a printer self-test usually tells you how much RAM the printer contains.)

To change the amount of memory specified for a printer, see Set printer memory.

**Assigning forms to paper trays**
If a printer has multiple trays that hold different types of forms, you can assign a form to a specific paper tray. Users can then select the form from within the program that they are using. When they print a document, the print job is routed to the correct paper tray.

Windows-based programs can use different forms within a document. For example, a user might use Envelope for the first page, Letterhead for the second page, and Letter for the third and following pages.

The default form setting for a paper tray is Letter. To set forms to trays, see Match a form to a printer tray.

**Choosing font types**
Products in the Windows Server 2003 family support three types of fonts:

- Device fonts actually reside in the hardware of your printer. They can be built into the printer itself or can be provided by a font cartridge or font card.

- Screen fonts are fonts (including OpenType, TrueType, raster, and vector fonts) that can be translated for output to the printer. OpenType extends the TrueType font standard but is compatible with previous versions of Windows and any application that supports TrueType fonts.

- Downloadable or soft fonts are sets of characters stored on disk and sent to a printer's memory when needed for printing a document. Clients that use downloadable fonts and that print to computers running Windows Server 2003 family operating systems should install the fonts locally.

Required screen and downloadable fonts are sent to the printer for each document. To improve printing times, use device fonts, which are already present at the printer.

Not all printers can use all three types of printer fonts. Pen plotters, for example, usually cannot use downloadable or print raster screen fonts.

# Using separator pages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using separator pages**

A separator or banner page typically states who sent the document to the printer and gives the date and time of printing. You can use one of the standard separator pages or create a custom page. Windows provides three separator page files, which are located in the systemroot\System32 folder.

The following table shows the names of the separator page files, the type of printer with which it is compatible, and the purpose of each page. You can edit or customize any of these separator pages.

| File name | Printer language | Purpose |
|---|---|---|
| Pcl.sep | Printer Control Language | Switches the printer to PCL printing and prints a separator page before each document. Pcl.sep is compatible with PCL printers, but may not work unless the printer also supports PJL. |
| Pscript.sep | PostScript | Switches the printer to PostScript printing, but does not print a separator page before each document. Pscript.sep is compatible with PostScript printers that also support PJL. |
| Sysprint.sep | PostScript | Switches the printer to PostScript printing and prints a separator page before each document. Sysprint.sep is compatible with PostScript printers. |

For more information, see Choose a separator page.

**Note**

- The separator page files provided by Windows may not work with some printers.

# Using plotters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using plotters**

The following are some considerations when printing to a plotter:

- In this version of Windows, 7-bit serial connections to plotters are not supported, so you must use 8-bit connections.
- If you notice that some lines are not printing on a pen plotter, try setting the plotter's control panel so that installed pens are set at the minimum width supported by the hardware. Be aware, however, that some plotters will not print thin lines unless the thinnest pen is selected.
- Pen plotters can take longer to clear their buffers than other plotters, so you may want to increase the port time-out setting to avoid receiving error messages.
- Windows plotter printer driver for most HP DesignJet series plotters are designed for use with margins set to Normal, paper size set to Software, and autorotation set to Off. When troubleshooting problems, reset the plotter to the factory defaults and then change only the necessary settings.
- All plotters supported by this version of Windows use the HP-GL/2 language. Some plotters supported by Windows can be set to receive a choice of plotter languages; however, to work with Windows, HP-GL/2 must be selected. (This version of Windows does not support the HPGL language.)
- When printing to plotter models that offer a choice of roll or sheet media, the form to tray assignment must be set to the actual paper size available to the plotter. Since selecting a media size or type not featured on the plotter is not supported, doing so may result in misaligned output.

# Printing extended or international characters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Printing extended or international characters**

In addition to the 128 standard ASCII characters you can type using your keyboard, you can use extended or international characters by using the Character Map program.

When Windows prints a document using printer fonts, each character your program sends to the printer is translated to the appropriate resident font. If your printer supports that character, the character prints. Otherwise, another character prints instead. Check your printer manual and experiment with your printer to determine which extended characters your printer supports. This limitation applies only to the printer's built-in fonts. TrueType and OpenType fonts provided in Windows will print all extended characters, unless the font has been mapped to a resident font on the printer. In that case, the TrueType or OpenType font will print only the characters supported by the printer.

For information about changing fonts or using the Character Map program, see Printer fonts overview.

# Managing printing from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing printing from the command line**

You can manage your printers at the command line by using:

- Print to print a text file or to display the contents of a print queue.
- Net print to display or control print jobs and printer queues.
- Net start spooler to start the Spooler service.
- Lpr to print a file to a computer running an LPD server.
- Lpq to obtain status of a print queue on a computer running the LPD server.

You can obtain additional functionality through scripts by using:

- Prncnfg.vbs to get and set printer configurations, or to rename a printer.
- Prndrvr.vbs to add, delete, and list printer drivers.
- Prnjobs.vbs to pause, resume, cancel, and list print jobs.
- Prnqctl.vbs to print a test page, pause or resume a printer, and clear a printer queue.
- Prnmngr.vbs to add, delete, and list printer connections. Prnmngr.vbs can also be used for obtaining and setting the default printer.
- Prnport.vbs to add, delete, and list standard TCP/IP ports. Prnport.vbs can also be used to obtain and set the port configuration.

For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools. For more information about printing, see Printing.

# Managing Printing Traffic

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing printing traffic**

With a computer running a member of the Windows Server 2003 family, you can use several methods to improve the efficiency of your printing operations and manage your printing traffic:

- Scheduling alternate printing times
- Assigning different priority levels
- Setting print spooler options
- Creating a printing pool
- Setting the location of the spool folder
- Monitoring print queue performance

# Scheduling alternate printing times

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Scheduling alternate printing times**

One way to maximize the use of printers is to schedule alternate printing times for long documents or certain types of documents. For example, if printer traffic is heavy during the day, you can postpone printing of long documents by routing them to a printer that prints only during off-hours. The print spooler continues to accept documents, but it does not send them to the destination printer until the designated start time.

Instead of dedicating an actual print device for only off-hour printing, which is not an efficient use of resources, you can set different logical printers for the same print device and configure each with different available times. One printer might be available from 6:00 PM to 6:00 AM, while the other is available 24 hours a day. You can then tell users to send long documents to the printer available only during off-hours, and all other documents to the printer available all the time. To set up printer availability times, see Schedule alternate printing times.

You can also differentiate the two printers with different priority levels to print certain types of documents before any other types, regardless of time of day. To set priority levels, see Set different print priority to different groups.

# Assigning different priority levels

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Assigning different priority levels**

You can expedite documents that need to be printed immediately. Documents sent by users with high priority levels can bypass a queue of lower priority documents waiting to be printed. If two logical printers are associated with the same printer, the Windows Server 2003 family operating system routes documents with the highest priority level to the printer first.

To take advantage of this print priority system, create multiple logical printers for the same printer. Assign each a different priority level, and then create a group of users that corresponds to each printer. For example, users in Group1 might have access rights to a priority 1 printer, users in Group2 might have access rights to a printer with priority 2, and so on.

For a complete procedure on how to set different priority levels using multiple logical printers, see Set different print priority to different groups.

# Setting print spooler options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting print spooler options**

Although you can control whether to use spooling or not, you'll typically want to spool documents you send to printing. Spooling documents expedites the printing process, and allows the application that sent the document to release control back to the user much faster. All the spooling options are set through the **Advanced** tab of the printer's property page. The following table shows other options you can set to affect how the spooler operates.

| Option | Description |
|---|---|
| | |
| Spool print documents so program finishes printing faster | Either select this option or the **Print directly to the printer** option. If you choose this option, the documents spool. This option has two related options that you must choose between (see the next two options). |
| Start printing after last page is spooled | The printer does not print a document until it is completely spooled. This is useful for documents that are assigned a low priority. Documents that are assigned a higher priority start printing immediately. |
| Start printing immediately | The printer starts printing a document before it is completely spooled, which means it is printed sooner, and the application you are printing from releases control back to you faster. |
| Print directly to the printer | The document does not spool, which decreases printing time. Select this option only for a nonshared printer. This might be useful for other programs that use their own spooling process. |
| Hold mismatched documents | Documents that do not match the configuration of the printer are not printed. This prevents errors resulting from documents that use paper sizes different from letter size. |
| Print spooled documents first | A spooled document is printed before a partially spooled document. |
| Keep documents after they are printed | Documents remain in the print spooler after they are printed, and can be quickly resubmitted for printing. You must carefully watch for disk space if you choose this option. |
| Enable advanced printing features | When enabled, the document is rendered using metafile datatype (EMF) and advanced features such as Page Order, Booklet Printing, and Pages Per Sheet are available. If compatibility problems occur, try to disable this option. |

For more information, see Set printer properties.

If you choose to set different spooling options on the same printer for different groups of users, you need to set multiple logical printers for the same physical printer.

# Creating a printing pool

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating a printing pool**

You can create a printing pool to automatically distribute print jobs to the next available printer. A printing pool is one logical printer connected to multiple printers through multiple ports of the print server. The printer that is idle receives the next document sent to the logical printer. This is useful in a network with a high volume of printing because it decreases the time users wait for their documents. A printing pool also simplifies administration because multiple printers can be managed from the same logical printer on a server.

With a printing pool created, the user prints a document without having to find out which printer is available. The logical printer checks for an available port and sends documents to ports in the order that they are added. Adding the port connected to the fastest printer first ensures that documents are sent to the printer that can print the fastest before they are routed to slower printers in the printing pool.

Before setting a printing pool, consider the following:

- All printers in a pool must use the same driver.
- Because users will not know which printer in the pool prints a given document, make sure all printers in the pool are located in the same place.

For procedures on how to set a printing pool, see Set up a printer to print to multiple devices.

# Setting the location of the spool folder

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting the location of the spool folder**

By default the spool folder is located at systemroot**\System32\Spool\Printers**. However, this drive also holds the Windows system files. Because these files are frequently accessed by the operating system, performance of both Windows and the printing functions might be slowed.

If your print server serves only one or two printers with low traffic volumes, the default location of the spool folder is sufficient. However, for high-volume printing requirements, to support large numbers of printers, or to support large print jobs, you should relocate the spool folder. For best results, move the spool folder to a drive that has its own controller, which reduces printing's impact on the rest of the operating system.

For information on relocating the print server spool folder, see Change the spool folder location.

For more information on changing the spool folder location for a single printer, see the Windows Server 2003 Resource Kit.

# Monitoring print queue performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring print queue performance**

You can monitor the performance and operation of printers across the network using System Monitor. Counters can be set up for a variety of performance criteria such as Bytes Printed/sec, Job Errors, and Total Pages Printed.

When you select **Print Queue** in the **Performance object** text box, you are provided with a list of all available performance counters for that object. The one you will probably use most is the **Jobs** counter, which is selected by default.

Other counters help you determine the current load on the print server or provide you with statistics and accumulative printing performance since the last time the system was started.

In System Monitor, click the plus sign (**+**) to add counters, and then click **Explain** to see an explanation of the counter.

For information on how to use System Monitor to monitor print queues, see Monitor print queue performance.

# Printing Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

- Windows Catalog at Support resources
- Microsoft TechNet at Microsoft TechNet
- Microsoft TechNet compact disks
- Microsoft Product Support Services at the Microsoft Web site
- Request for Comments Web site

**Note**

- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

# Printing Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

Products in the Windows Server 2003 family have a modular printing architecture with each module responsible for different types of tasks. This modularity provides a great deal of flexibility, allowing you to support different clients on different operating systems and submit print jobs from different environments and applications. However, that flexibility and variety of options can make troubleshooting more challenging.

This section contains information on solving printing problems.

- General troubleshooting information
- Troubleshooting common printing problems

# General troubleshooting information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**General troubleshooting information**

Successful troubleshooting depends on your ability to identify the source of the problem. Printing problems occur because of trouble with:

- The printing device.
- The connection between the printer and the network.
- Other printing components in the Windows Server 2003 family products.
- The network, protocols, and other communication components.

When solving printing problems, a good first step is to identify the process creating the problem. You can often determine this by testing different scenarios. For example, if a network client cannot print to a certain printer and other users can, then you know that the problem is not with the physical printer and its connection to the network.

**Troubleshooting printing problems: quick list**

A quick way to solve most problems associated with printing is to follow these four steps:

- Verify the physical printer is operational. If other users can print, it is probably not a problem with the printer, or with the print server.
- Verify that the printer on the print server is using the correct printer driver. If print clients are using other operating systems, make sure you installed all necessary drivers for the other platforms.
- Verify the print server is operational, make sure there is enough disk space for spooling, and make sure the print spool service is running.
- If print clients are using operating systems other than Windows XP , verify the client computer has the correct printer driver.

**Troubleshooting physical problems**

- Check the physical printer and make sure it is in the ready state (ready to print). With most printers, you can print a test page from the printer console itself to verify correct operation.
- Check the cable connecting the printer and make sure it is connected properly. If the printer is connected directly to the network with a network adapter, check the light on the card, which indicates network connectivity.

**Troubleshooting communication problems**

- Make sure you can communicate with the printer over the network. As an example, if a printer communicates over the TCP/IP protocol, you can use the **ping** command to verify connectivity.
- Verify communication to the print server from other computers.
- Check that any service required for the printer and the client submitting the print job is working properly. For example, for a Macintosh client, verify that the AppleTalk protocol is working.
- Make sure the print spool service is running on the print server.

**Troubleshooting Windows Server 2003 family printing**

The modularity of the printing architecture makes diagnosing problems fairly easy. Windows Server 2003 family printing across a network consists of seven processes that always occur. By testing each process, you can generally identify the problem.

1. The administrator adds a printer and shares it.

    - Check the property of the logical printer; make sure the driver is correct for that printer.
    - Using the Add Printer Wizard, you can add another logical printer for the same physical printer to quickly identify whether the problem is with the logical printer.
    - If you cannot browse printer connections, or cannot find the printer port, you might have a network problem.
    - Verify you are logged on as administrator, or as a member of the Administrator group.

2. A network client connects to that share.

    - Check the property of the logical printer on the client computer; make sure the driver is correct for that printer.
    - Using the Add Printer Wizard, you can add another logical printer for the same physical printer to quickly identify whether the problem is with the logical printer.
    - Check the users permissions to print to that printer, and check which Group Policy settings for printers are enabled.
    - If you cannot browse printer connections, you might have a network problem.

3. The client application creates a print job.

    - Check whether the document the client is trying to print consists of text only or includes graphics. Check the printer driver and the fonts settings.
    - Check whether there is a problem with the separator page selection.
    - Try to reproduce the same print job from another client; if it prints properly, the problem is most likely not caused by this process.

4. The client system sends the print job to the printer share on the print server.

    - Check the network transport; for example, TCP/IP, NWLink, or AppleTalk.
    - Check other network components that are required to print; for example Print Server for Macintosh or Print Services for Unix.

    **Note**

Print Server for Macintosh is not available on the 64-bit versions of the Windows operating systems.

- Make sure the network is working properly for the client. Check the network adapter and intermediate network devices.
- Make sure the print spool folder on the print server has not run out of disk space.
- Try to send the print job to the spooler in another way, or from another application.

1. The print server receives, spools, and--if necessary--alters the print job.

    - Check disk space for spooling.
    - Make sure the problem does not involve an incorrect data type, such as RAW or PScript1.
    - Try to change an EMF print job to RAW to see whether it solves the problem.

2. The print server sends the print jobs to the printer.

    - Check network transport (such as TCP/IP or NWLink) and other network devices.
    - Check the physical printer connection to the network.
    - Make sure the printer is ready to print.
    - Try to print from other network clients.

3. The printer interprets the job and produces hardcopy output.

    - Make sure the printer is ready to print.
    - Try to print a test page from the printer console.

# Troubleshooting common printing problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting common printing problems**
What problem are you having?

- A printer connected to a computer does not print.
- A printer connected to the network does not print.
- Printing from an Outlook client with multiple languages is slow.
- Users on Windows 95 or Windows 98 cannot connect to a printer.
- Users get an Access Denied message when trying to configure a printer from within an application.
- The document does not print completely, or comes out garbled.
- Hard disk problems occur, and the document does not reach the print server.
- Documents on the print server will not print and cannot be deleted.
- An application on a 16-bit version of Windows (such as Windows for Workgroups) gives an out-of-memory error on startup.
- A document sent to the printer from an MS-DOS–based program on a Windows XP, Windows 2000, or Windows NT 4.0 client does not print.
- A driver is not listed.
- A new print driver downloaded from the Internet needs to be installed.
- Microsoft does not have the printer driver for the printer.
- With printer location tracking enabled, some printers cannot be located.
- The company has changed its organization, and as a result the naming scheme has to change.
- My organization has a small network without any subnets.

**A printer connected to a computer does not print.**
**Cause:** A problem might exist with the physical printer, the print driver, a print server, or the application you are trying to print from.

**Solution:** Try the following steps to identify and solve the problem.

- Verify the physical printer is in ready state, and that the correct default printer is set.
- Try to print a test page.
- Try printing from the command line (on non-PostScript printers only). Type **Dir > LPT1** at the command prompt.

- Try printing from Notepad. This verifies that the printer driver is correct, and confirms that the problem lies with the application. If you cannot print from Notepad, the problem lies with the printer driver.

- Check the available disk space on the system drive. If there is not enough room to spool the job, larger jobs might fail where smaller jobs do not.

**A printer connected to the network does not print.**
**Cause:** A problem might exist with the physical printer, the logical printer setting on the client computer, the application you are trying to print from, network protocols, or hardware.

**Solution:** Try the following to isolate and solve the problem:

- Verify basic network connectivity. Check user rights, protocols, share names, and so on to determine whether you can see the server. Try to copy a file to the server; if you cannot access the server, you might not be able to access the printer.

- On the **Advanced** tab, make sure the printer availability time is correct.

- Create a local printer and redirect the port to the network server. Using the Add Printer Wizard, click Local, and then type the server and printer name for the printer (*****servername******printername*). This determines whether you can copy files from the print server.

- Check the available disk space on the print server. If there is not enough room to spool the job, larger jobs might fail where smaller jobs do not.

**Printing from an Outlook client with multiple languages is slow.**
**Cause:** Multiple languages are installed on the client computer, but not on the print server computer.

**Solution:** If clients use additional languages on their computers, you need to add these languages to print server.

Users on Windows 95 or Windows 98 cannot connect to a printer.

**Cause:** Print drivers for Windows 95 and Windows 98 are not the same as for the Windows Server 2003 family, and are not installed on the print server.

**Solution:** Make sure the user has Print permission, and make sure the proper printer driver is installed on the print server for that operating system.

**Users on Windows 95 or Windows 98 cannot connect to a printer.**
**Cause:** Print drivers for Windows 95 and Windows 98 are not the same as for the Windows Server 2003 family, and are not installed on the print server.

**Solution:** Make sure the user has Print permission, and make sure the proper printer driver is installed on the print server for that operating system.

**Users get an Access Denied message when trying to configure a printer from within an application.**
**Cause:** You do not have the appropriate permission to change printer configuration.

**Solution:** You need Manage Printer permission to change the printer's setup.

**The document does not print completely, or comes out garbled.**
**Cause:** The printer's driver is either corrupted or incorrect.

**Solution:** Verify or reinstall the correct printer driver on the client computer.

**Hard disk problems occur, and the document does not reach the print server.**
**Cause:** The hard disk might not have enough space for spooling the document.

**Solution:** Make sure the hard disk has enough disk space, or relocate the spool folder to another volume.

**Documents on the print server will not print and cannot be deleted.**
**Cause:** The print spooler might be stalled.

**Solution:** On the print server, try to stop and restart the Print Spooler service.

**An application on a 16-bit version of Windows (such as Windows for Workgroups) gives an out-of-memory error on startup.**

**Cause:** A default printer is not selected.

**Solution:** Add a printer and set it as the default printer.

**A document sent to the printer from an MS-DOS–based program on a Windows XP, Windows 2000, or Windows NT 4.0 client does not print.**
**Cause:** Some MS-DOS programs will not print until you close the program.

**Solution:** Close the application you tried to print from. In addition, make sure the correct printer driver is installed on the client computer.

**A driver is not listed.**
**Cause:** The printer was not on the market when the Windows Server 2003 family was released, or a newer version of the print driver has only recently become available.

**Solution:** Windows printer drivers are developed through cooperation between Microsoft and the independent hardware vendor (IHV) that manufactures the print device. A product in the Windows Server 2003 family supports all print devices listed in the Windows Catalog; see Support resources.

For new and updated print drivers, see Microsoft printing support at the Microsoft Web site. You should also check Windows Update.

**A new print driver downloaded from the Internet needs to be installed.**
**Cause:** The print driver is not listed in the Add Printer Wizard's list of printers, but instead was downloaded from the Internet.

**Solution:**

1. Download the printer driver to an empty folder.

2. Expand the files by typing the name of the executable followed by the **-d** switch. For example, if the name of the file is Nt4prn.exe, type the following at a command prompt:

   **nt4prn -d**

3. Read the Readme.txt and the License.txt files.

4. Click **Start**, point to **Settings**, and then click **Printers and Faxes**.

5. Double-click **Add Printer** to start the Add Printer Wizard, click **Next**, click **Local printer attached to this computer**, and then click **Next**.

6. Click the port where your printer is attached, and then click **Next**.

7. Click **Have Disk**, enter the path from Step 1, and then click **OK**.

8. Click the appropriate printer, click **Next**, and then continue with the Add Printer Wizard.

**Microsoft does not have the printer driver for the printer.**
**Cause:** The printer might be too new, and Microsoft does not have the driver ready on its Web site.

**Solution:** Find a printer driver emulation that can provide the correct output. Check with the printer manufacturer to see if they have a Windows Server 2003 family printer driver.

**With printer location tracking enabled, some printers cannot be located.**
**Cause:** After printer location tracking is enabled, the default behavior is that the user will find only those printers whose location attribute matches the naming convention.

**Solution:** You need to set the location string again for this printer. For information on how to do this, see Enable printer location tracking.

**The company has changed its organization, and as a result the naming scheme has to change.**
**Solution:** First, use Active Directory Sites and Services to update the sites and subnets. Next, you need to update all the printer location fields for the affected printers in the organization. You can create an ADSI script to expedite the updates.

**My organization has a small network without any subnets.**

**Solution:** Active Directory cannot locate a node in the tree of your directory without having a subnet associated with it. If your organization is small enough, then this is not a problem, and users can search for any printer in the organization. But for larger organizations, or ones with many printers, you can still allow users to browse the location hierarchy as an aid to locate a printer. To do this, you need to create artificial subnets in Active Directory to correspond to the hierarchy you want to create.

# Availability and Scalability

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Availability and scalability**
This section covers:

- Windows Clustering
- Enterprise Memory Architecture

**Information about other technologies that enhance availability**
For information about other operating system technologies that enhance availability for servers running Microsoft® Windows® Server 2003, see the following:

- For information about using an Uninterruptible Power Supply (UPS) that connects to your computer using a serial cable, see Using an Uninterruptible Power Supply.
- For information about dynamic volumes, including RAID-5 volumes which provide disk fault tolerance, see Disk Management.
- For information about monitoring tools, see Monitoring and Status Tools.
- For information about backup tools, see Backup.
- For information about troubleshooting tools and strategies, see Troubleshooting Strategies.
- For information about preparing for recovery of a system, see Disaster Recovery.
- For detailed information about setting system recovery options, see Specify what happens when the system stops unexpectedly.

**Information about practices that enhance availability**
Maintaining highly available systems requires not only good choices in technology, but careful attention to design and operational practices. For more information about practices that enhance availability, see the following:

- By following best practices for the technologies you use, you can prevent problems before they start, increasing system availability. For information about best practices for specific technologies, see the list in Best practices: pointers to recommendations and tips.
- By following best practices for security, you protect against attacks that could decrease system availability. For more information, see Best practices for security.
- By understanding the deployment and management practices described in the Windows Resource Kits, you can design, test, and operate systems in ways that maximize availability. For more information, see the following:
    - For information about browsing or obtaining the Windows Resource Kits, see Using the Windows Deployment and Resource Kits.
    - For information about deploying Windows, see "Planning for Deployment" at the Microsoft Windows Resource Kits Web site.
    - For information about planning for high availability, see "Planning Server Deployments" at the Microsoft Windows Resource Kits Web site.

- For detailed information about disks and volumes, server clusters, and network load balancing clusters, see the "Server Management Guide" at the Microsoft Windows Resource Kits Web site.

# Windows Clustering

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Clustering**

Windows Clustering provides three different, but complementary, clustering technologies. The clustering technologies, which ship in a number of different products, can be used separately or combined to provide scalable and highly-available services.

| **Clustering Technology** | **Network Load Balancing (NLB) clusters** | **Component Load Balancing (CLB) clusters** | **Server clusters** |
|---|---|---|---|
| Available in ... | Microsoft® Windows Server™ 2003, Web Edition; Microsoft® Windows Server® 2003, Standard Edition; Microsoft® Windows Server™ 2003, Enterprise Edition; and Microsoft® Windows Server™ 2003, Datacenter Edition | Microsoft Application Center 2000 | Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition |
| Maximum number of nodes | 32 | 12 | 8 |
| Application | Load balancing Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) traffic | Single point of management and configuration for Web farms | Failover and failback of applications |
| Specialized hardware required? | No<br>**Note**<br>• If you use teaming network adapters, you must select network adapters listed in the Windows Catalog. For more information, see Network Load Balancing system requirements and the compatibility information in Support resources. | No | Yes<br>To confirm that your server cluster hardware is designed for Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, see the compatibility information in Support resources. |
| Typical deployments | Web servers, Microsoft Internet Security and Acceleration (ISA) server, virtual private networks, Windows Media™ servers, Mobile Information servers, Terminal Services | Web farms | MS SQL Server, MS Exchange Server, file and print servers, Message Queuing |
| Stateful or stateless? | Stateless | Stateless | Stateful |

**Important**

- Microsoft will not support the configuration of server clusters and Network Load Balancing clusters on the same server. For more information about how the Windows Clustering technologies can be combined in a multitiered approach to provide highly available services, see "Planning for High Availability" in the *Microsoft Windows Server 2003 Deployment Kit* at the Microsoft Windows Resource Kits Web site. In addition, see "Server Clusters and Network Load Balancing" in the *Microsoft Windows Server 2003 Resource Kit* at the Microsoft Windows Resource Kits Web site.

1. **Network Load Balancing clusters.** Network Load Balancing clusters provide scalability and high availability for TCP- and UDP-based services and applications by combining up to 32 servers running Windows Server 2003, Web Edition; Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition, into a single cluster. By using Network Load Balancing to build a group of cloned, or identical, clustered computers, you can enhance the availability of these servers: Web and File Transfer Protocol (FTP) servers, ISA servers (for proxy servers and firewall services), virtual private network (VPN) servers, Windows Media servers, Terminal Services over your corporate LAN.

   You can install Network Load Balancing clusters through Network Connections or by using the Network Load Balancing Manager. For more information about Network Load Balancing clusters, see Network Load Balancing Overview.

2. **Component Load Balancing clusters.** Component Load Balancing clusters provide high scalability and availability by enabling COM+ applications (for example, a shopping cart application on an e-commerce Web site) to be distributed across multiple servers. For more information, see the documentation for Microsoft Application Center 2000 in Microsoft TechNet at the Microsoft Web site.

   **Important**

   - Component Load Balancing clusters is a feature of Microsoft Application Center 2000. It is not a feature of Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition.

3. **server clusters.** server clusters provide high availability for applications through the failover of resources. server clusters focus on preserving client access to applications and system services, such as Microsoft Exchange for messaging, Microsoft SQL Server for database applications, and file and print services.

   Server clusters can combine up to eight nodes. In addition, a cluster cannot be made up of nodes running both Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. In server clusters with more than two nodes, all nodes must run Windows Server 2003, Datacenter Edition, or Windows Server 2003, Enterprise Edition, but not both.

   By default, all clustering and administration software files are automatically installed on your computer when you install any operating system in the Windows Server 2003 family. For more information about server clusters, see Understanding Server Clusters.

For more information about Network Load Balancing clusters and server clusters, see the following topics:

- For conceptual and procedural information on Network Load Balancing clusters, see Network Load Balancing Clusters.
- For conceptual and procedural information on server clusters, see Server Clusters.
- For guidelines on securing a Network Load Balancing cluster, see Network Load Balancing Best practices.
- For guidelines on securing a server cluster, see Best practices for securing server clusters.

# Network Load Balancing Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Load Balancing clusters**
This section of Windows Clustering covers Network Load Balancing (NLB), one of two clustering technologies offered as part of Windows. For a general overview of this subject, see Windows Clustering. You can use Network Load Balancing to enhance the scalability and availability of server applications. Applications that can take advantage of Network Load Balancing include Web services like HTTP and File Transfer Protocol (FTP) (using Internet Information Services (IIS)), firewall and proxy (using ISA 2000), virtual private networking, Windows Media Services, Mobile Information Server, and Terminal Services. Network Load Balancing can also help you scale your server's performance to keep up with the increasing demands of your Internet-based clients.

A new Network Load Balancing feature is Network Load Balancing Manager. Network Load Balancing Manager allows you to create, configure, and manage all hosts of a Network Load Balancing cluster from a single computer. For more information, see Network Load Balancing key features.

- Before you enable and configure Network Load Balancing, it is vital that you review the checklist, including the sections on cluster parameters, host parameters and port rules. It is highly recommended that you read through these three topics before configuring Network Load Balancing.

  Checklist: Enabling and configuring Network Load Balancing

- To find features that have been changed in the Windows Server 2003 family of products, see New ways to do familiar Network Load Balancing tasks.

- For tips about using Network Load Balancing, see Network Load Balancing Best practices.

- For help with specific tasks, see Network Load Balancing How To....

- For general background information, see Network Load Balancing Concepts.

- For troubleshooting information, see Network Load Balancing Troubleshooting.

# Checklist: Enabling and configuring Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Enabling and configuring Network Load Balancing**

| Step | Reference |
|---|---|
| Read "Best practices." | Network Load Balancing Best practices |
| Read "Network Load Balancing system requirements." | Network Load Balancing system requirements |
| Read "Planning your Network Load Balancing cluster." | Planning Your Network Load Balancing Cluster |
| (Optional) Install a second network adapter. | Single network adapter limitations |
| Have ready the cluster's full Internet name. | See Full internet name in Network Load Balancing parameters |
| Have ready the cluster's primary IP address. | See IP address on the Cluster Parameters tab in Network Load Balancing parameters;<br>Single network adapter;<br>Multiple network adapters |
| Have ready the subnet mask for the cluster. | See Subnet mark in Network Load Balancing parameters |
| Have ready the current IP address for each host in the cluster. | See IP address on the Host Parameters tab in Network Load Balancing parameters;<br>Single network adapter;<br>Multiple network adapters |
| Have ready the subnet mask for each host in the cluster. | See Subnet mark in Network Load Balancing parameters |
| Read "Cluster parameters." | See Cluster parameters in Network Load Balancing parameters |
| Carefully consider how to configure the following:<br>• **multicast support**<br>• **remote password**<br>• **confirm password**<br>• **remote control** | See Cluster Operation Mode, Remote password, Confirm password, and Remote control in Network Load Balancing parameters. |
| Read "Host parameters." | See Host parameters in Network Load Balancing parameters |

| | |
|---|---|
| Carefully consider how to configure the following:<br>• **priority (ID)**<br>• **initial state** | See Priority ID and Initial host state in Network Load Balancing parameters |
| Read "Port rules." | See Port rules in Network Load Balancing parameters |
| Carefully consider how to configure the following:<br>• **port range**<br>• **protocols**<br>• **Filtering mode**<br>• **Affinity**<br>• **load weight**<br>• **handling priority** | See Port range, Protocols, Filtering mode, Affinity, Load weight, and Handling priority in Network Load Balancing parameters |
| Enable and configure Network Load Balancing. | Enable Network Load Balancing |

# New ways to do familiar Network Load Balancing tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar Network Load Balancing tasks**

The following table lists common tasks for configuring and managing Network Load Balancing. The user interface for performing these tasks is different in the Windows Server 2003 family of products from the way it was in Windows 2000. Although the Windows 2000 procedures are still valid, in most cases the new procedures are the recommended methods.

| **If you want to** | **In Windows 2000 use** | **In the Windows Server 2003 family of products use** |
|---|---|---|
| Enable Network Load Balancing | Network and Dial-up Connections on each host. | Network Load Balancing Manager. For more information, see Enable Network Load Balancing. |
| Configure cluster parameters | Network and Dial-up Connections on each host. | Network Load Balancing Manager. For more information, see Configure cluster parameters. |
| Add a host | Network and Dial-up Connections on each host. | Network Load Balancing Manager. For more information, see Add a host. |
| Configure host parameters | Network and Dial-up Connections on each host. | Network Load Balancing Manager. For more information, see Configure host parameters. |
| Create, edit, or remove a port rule | Network and Dial-up Connections on each host. | Network Load Balancing Manager. For more information, see Create a new port rule, Edit an existing port rule, and Remove a port rule. |
| Perform basic Network Load Balancing commands: Stop, Start, Suspend, Resume, Drain, Drainstop | **Wlbs.exe**. | either Network Load Balancing Manager, **nlb.exe**, or **wlbs.exe**. For more information, see Perform Basic Network Load Balancing Operations and Nlb. |
| Disable Network Load Balancing | Network and Dial-up Connections on each host. | Network Load Balancing Manager. For more information, see Remove a host from a Network Load Balancing cluster and Remove all hosts from a Network Load Balancing cluster. |

# Network Load Balancing Best practices

Updated: December 21, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Note**

- The general security and configuration guidance in this topic also applies to Network Load Balancing (NLB) clusters in Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012. However, for detailed operational guidance about NLB in Windows Server 2008 or Windows Server 2008 R2, see Network Load Balancing.
- The NLB functionality in Windows Server 2012 is generally the same as in Windows Server 2008 R2. However, some task details are changed in Windows Server 2012. For information on new ways to do tasks, see Common Management Tasks and Navigation.

**Best practices**

**Properly secure the Network Load Balancing hosts and the load balanced applications.**

- Network Load Balancing does not provide additional security for the load balanced hosts and cannot be used as a firewall. It is therefore important to properly secure the load balanced applications and hosts. Security procedures can typically be found in the documentation for each particular application. For example, if you are using Network Load Balancing to load balance a cluster of Internet Information Services (IIS) servers, you should follow the procedures and guidelines for securing IIS. To view the IIS 6.0 product documentation, install IIS and then see Microsoft Internet Information Services, or install IIS and then open the IIS User Interface (the IIS snap-in) and click **Help**.
- The Network Load Balancing subnet must be physically protected from intrusion by unauthorized computers and devices in order to avoid interference from unauthorized heartbeat packets.
- If you use the optional host list with Network Load Balancing Manager, ensure that only users in the local Administrators group have access to the host list file.

  For other general information on best practices for securing servers, see Best practices for security and the Microsoft Windows Resource Kits Web site. For tips about installing IIS, see Installing IIS.

**Use two or more network adapters in each cluster host whenever possible. Two network adapters, however, are not required.**

- If the cluster is operating in unicast mode (the default), Network Load Balancing cannot distinguish between single adapters on each host. Therefore, any communication among cluster hosts is not possible unless each cluster host has at least two network adapters.
- You can configure Network Load Balancing on more than one network adapter. However, if you use a second network adapter to address this best practice, make sure that you install Network Load Balancing on only one adapter (called the cluster adapter).

  For more information, see Multiple network adapters.

**Use only the TCP/IP network protocol on the cluster adapter.**

- Do not add any other protocols (for example, IPX) to this adapter.

**Use Network Load Balancing Manager.**

- You can configure many Network Load Balancing options through either Network Load Balancing Manager or the **Network Load Balancing Properties** dialog box accessed through **Network Connections**. However, Network Load

Balancing Manager is the preferred method. Using both Network Load Balancing Manager and **Network Connections** together to change Network Load Balancing properties can lead to unpredictable results.

**Do not enable Network Load Balancing remote control.**

- The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. It is highly recommended that you do not enable remote control and instead use Network Load Balancing Manager or other remote management tools such as Windows Management Instrumentation (WMI).




If you choose to enable remote control, it is vital that you restrict access by specifying a strong remote control password. It is also imperative that you use a firewall to protect the Network Load Balancing UDP control ports (the ports that receive remote control commands) in order to shield them from outside intrusion. By default, these are ports 1717 and 2504 at the cluster IP address. Use remote control only from a secure, trusted computer within your firewall.

For more information on the remote control parameter, see Remote control in Network Load Balancing parameters. For more information about strong passwords, see Strong passwords.

**Enable Network Load Balancing Manager logging.**

- You can configure Network Load Balancing manager to log each Network Load Balancing Manager event. This log can be very useful in troubleshooting problems or errors when using Network Load Balancing Manager. Enable Network Load Balancing Manager logging by clicking **Log Settings** in the Network Load Balancing Manager **Options** menu. Check the **Enable logging** box and specify a name and location for the log file.

  The Network Load Balancing Manager log file contains potentially sensitive information about the Network Load Balancing cluster and hosts, so it must be properly secured. By default, the log file inherits the security settings of the directory in which it is created, so you might have to change the explicit permissions on the file to restrict read and write access to those individuals who don't need full control of the file. Be aware that the individual using Network Load Balancing Manager does require full control of the log file. For more information, see Access Control How To....

**Verify that the following is true for cluster parameters, port rules and host parameters:**

- Cluster parameters and port rules for each unique virtual IP address are identical across all hosts.

  Each unique virtual IP address must be configured with the same port rules across all hosts that service that Virtual IP address. However, if you have multiple virtual IP addresses configured on a host, each of those virtual IP addresses can have a different set of port rules.

- Port rules are set for all ports used by the load-balanced application. For example, FTP uses port 20, port 21, and ports 1024-65535).

  Always click **Add** after setting a port rule. Otherwise, the port rule will not appear in the list of rules, and the rule will not take effect.

- The dedicated IP address is unique and the cluster IP address is added to each cluster host.

- Affinity is set to **Single** or **Class C** when you are using **UDP** or **Both** for your protocol setting.

For more information, see Cluster parameters, Host parameters, and Port rules in Network Load Balancing parameters.

**Verify that any given load-balanced application is started on all cluster hosts on which the application is installed.**

- Network Load Balancing does not start or stop applications.

**Verify that the following is true for the dedicated IP address and the cluster IP address:**

- Except in the case of a virtual private network (VPN),both the dedicated IP address and the cluster IP address must be entered during setup in the **Network Load Balancing Properties** dialog box and also in the **Internet Protocol (TCP/IP) Properties** dialog box.Make sure that the addresses are the same in both places. However, if you are configuring a VPN load balancing cluster, you should not configure the dedicated IP address. On a VPN, only the cluster IP address should be present on each of the cluster hosts because clients running Windows 95, Windows 98, or Windows NT 4.0 may be unable to connect to the cluster if the dedicated IP address is configured on the Network Load Balancing cluster hosts.

  If you omit this step, the cluster will converge and appear to be working properly, but the cluster host will not accept and handle cluster traffic.

- The dedicated IP address is always listed first (before the cluster IP address) in the **Internet Protocol (TCP/IP) Properties** dialog box. This will ensure that responses to connections originating from a host will return to the same host. For more information, see Set up TCP/IP for Network Load Balancing.

- Both the dedicated IP address and the cluster IP address must be static IP addresses. They cannot be DHCP addresses.

**Ensure that all hosts in a cluster belong to the same subnet and that the cluster's clients are able to access this subnet.**

- No cluster interconnect is used by Network Load Balancing other than the subnet in which the cluster is located. You should therefore not connect two network adapters in an effort to create a system area network (SAN) for which there is no need.

**Perform moves of a cluster host according to the following guidelines:**

- If you move a cluster host from one cluster to another on the same subnet by changing the cluster IP address, first remove Network Load Balancing as described in Remove a host from a Network Load Balancing cluster, and then reenable Network Load Balancing after changing the IP address.

  This will prevent you from experiencing an IP address conflict.

**Verify that all cluster hosts are operating in either unicast or multicast mode, one or the other, but not both.**

**Always begin Network Load Balancing command-line commands with "nlb.exe".**
For more information, see Nlb.

**Do not enable Network Load Balancing on a computer that is part of a server cluster.**

- Network Load Balancing can interfere with server cluster's use of network adapters and Microsoft does not support this configuration. Instead, use separate Network Load Balancing and server clusters.

For more information on server clusters, see Windows Clustering.

**Avoid uninstalling Network Load Balancing.**

- There is typically no need to uninstall this feature. Network Load Balancing is an integral part of the products in the Windows Server 2003 family and does not need to be installed or uninstalled separately.

# Network Load Balancing How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Create and Manage Network Load Balancing Clusters
- Configure Network Load Balancing Parameters
- Perform Basic Network Load Balancing Operations
- Perform Advanced Network Load Balancing Procedures

# Create and Manage Network Load Balancing Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and Manage Network Load Balancing clusters**

- Enable Network Load Balancing
- Set up TCP/IP for Network Load Balancing
- Add a host
- Connect to existing clusters from Network Load Balancing Manager using a single host
- Connect to existing clusters from Network Load Balancing Manager using a host list
- Enable Network Load Balancing Manager logging
- Remove a host from a Network Load Balancing cluster
- Remove all hosts from a Network Load Balancing cluster
- Stop managing a cluster from Network Load Balancing Manager

# Enable Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Network Load Balancing**

1. Open Network Load Balancing Manager.

2. Right-click **Network Load Balancing Clusters**, and then click **New Cluster**.

3. Enter the cluster's IP address and other cluster information and click **Next**.

4. If necessary, add additional virtual IP addresses that will be used by this cluster by clicking **Add**, entering the virtual IP address and Subnet mask information, then clicking **OK**. Repeat this process for each additional virtual IP address, then click **Next**.

5. If necessary, add appropriate port rules and then click **Next**.

6. Type the name of a host that will be a member of your cluster and click **Connect**. After you click **Connect** the network adapters that are available on the host that you typed will be listed at the bottom of the dialog box. Click the network adapter that you want to use for Network Load Balancing and then click **Next**. The IP address configured on this network adapter will be the dedicated IP address for this host.

   **Important**

   - To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

7. Configure the remaining host parameters as appropriate, then click **Finish**.

8. Add additional hosts as appropriate.

**Important**

- To ensure that Network Load Balancing Manager is displaying the most recent host information, right-click the cluster and click **Refresh**. This step is necessary because the host properties that Network Load Balancing Manager displays are a copy of the host properties that were configured the last time Network Load Balancing Manager connected to that host. When you click **Refresh**, Network Load Balancing Manager reconnects to the cluster and displays updated information.

**Notes**

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

- If Network Load Balancing Manager is not used to enable Network Load Balancing, you must manually add the cluster's IP address to the TCP/IP settings on each host. For more information, see "Set up TCP/IP for Network Load Balancing" in Related Topics below.

- To install and configure Network Load Balancing, you must use an account that is in the Administrators group on each host. As you install and configure each host, if you are not using an account in the Administrators group, you will be prompted to provide the logon credentials for such an account. You can set up an account that Network Load Balancing Manager will use by default by clicking **Credentials** in the Network Load Balancing Manager **Options** menu. It is recommended that this account is an account not used for any other purpose. Be sure to keep the password from expiring on the account (follow your organization's policies for password renewal).

- When you add additional hosts, they will automatically inherit the cluster port rules from the initial host.

- If Network Load Balancing clients are accessing a cluster through a router, be sure that the router meets the following requirements:

    - Accepts a dynamic Address Resolution Protocol (ARP) reply that maps the unicast IP address to its unicast media access control (MAC) address

    - In multicast mode, accepts an ARP reply that has a MAC address in the payload of the ARP structure. If your router does not meet these requirements, you can also create a static ARP entry in the router. Certain routers require a static ARP entry because they do not support the resolution of unicast IP addresses to multicast MAC addresses.

- The parameters set in this procedure are recorded in the registry on each host.

- Network Load Balancing also supports unattended setup. For information about unattended setup, see "Planning for unattended setup" in Related Topics.

- When installing Network Load Balancing, you might expect to find a setup application (such as Setup.exe). There is no such application. The Network Load Balancing driver is installed as a standard Windows networking driver.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Checklist: Enabling and configuring Network Load Balancing
Virtual clusters
Create a new port rule
Configure cluster parameters
Configure host parameters
Configure Port Rules
Set up TCP/IP for Network Load Balancing
Network Load Balancing parameters
Planning for unattended Setup

# Set up TCP/IP for Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set up TCP/IP for Network Load Balancing**

1. Open Network Connections.

   **Important**

   - These steps should be used to configure TCP/IP only if you choose not to use Network Load Balancing Manager to create your cluster. When configuring Network Load Balancing, you can use either Network Load Balancing Manager, or you can access the Network Load Balancing Properties dialog box through Network Connections as described in this procedure. Using Network Load Balancing Manager is the preferred method. For more information on using Network Load Balancing Manager to create a cluster, see "Enable Network Load Balancing" in Related Topics.

2. Right-click **Local Area Connection**, and then click **Properties**.

3. In the **Local Area Connection Properties** dialog box, click **Internet Protocol (TCP/IP)**, and then click **Properties**.

4. Click **Use the following IP address**, and then enter the host's dedicated IP address in the space for **IP address**.

5. Enter subnet mask and default gateway information for your TCP/IP network.

6. Save your settings by clicking **OK**.

7. On the host you just added, return to Network Connections.

8. Right-click the network adapter that is using Network Load Balancing and then click **Properties**.

9. In the **Local Area Connection Properties** dialog box, click **Internet Protocol (TCP/IP)**, and then click **Properties**.

10. Click **Advanced**, and then, under **IP addresses**, click **Add**.

11. Enter the cluster IP address in the space for the IP address, followed by the appropriate subnet mask, and then click **Add**.

    This IP address corresponds to the cluster's primary IP address that you entered in the **Network Load Balancing Properties** dialog box under **Cluster parameters** when first enabling Network Load Balancing.

12. If you need to configure additional default gateways or virtual IP addresses for your cluster, you can do so here, in the **Advanced TCP/IP Settings** dialog box.

**Important**

- It is essential that you enter the dedicated IP address first, in the **Internet Protocol (TCP/IP) Properties** dialog box (not in the **Advanced TCP/IP Settings** dialog box), so that all outbound connections made from this host (for example, Telnet or FTP) are initiated with this address.

- TCP/IP is the only network protocol that should be present on the cluster adapter. You must not add any other protocols (for example, IPX) to this adapter.

- Be aware that if you are using Network Load Balancing to load balance Point-to-Point Tunneling Protocol (PPTP) servers,

you must use two network adapters on each cluster host (although Network Load Balancing does not need to be installed on both adapters). When load balancing PPTP, the hosts can not have a dedicated IP addresses configured.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.
- You should repeat this procedure for each network adapter on which you want to install Network Load Balancing.
- Both the dedicated IP address and the cluster's primary IP address must be static IP addresses. They cannot be DHCP (Dynamic Host Configuration Protocol) addresses.
- Regardless of the mode in which Network Load Balancing is operating (unicast or multicast), the steps described previously are the same.
- Although unicast mode is preferable in many situations, the dedicated IP address, will not be addressable by the other cluster hosts unless multicast support has been enabled in the **Network Load Balancing Properties** dialog box under **Cluster parameters**.
- If only one **Local Area Connection** icon is displayed in **Network Connections**, it means that you have only one network adapter installed on your computer. Additional network adapters are represented by additional **Local Area Connection** icons.
- If you use two network adapters you should supply an IP address for the default gateway for only one of the adapters (the one with a route to the largest network, for example to the internet).

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Load Balancing parameters
Single network adapter
How Network Load Balancing works
Enable Network Load Balancing

# Add a host

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a host**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.

   **Important**

   - To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

3. Right-click the cluster where you want to add the host and choose **Add Host To Cluster**.
4. Enter the host's name and click **Connect**.
5. After you click **Connect** the network adapters that are available on the host that you typed will be listed at the bottom of the dialog box. Click the network adapter that you want to use for Network Load Balancing and then click **Next**. The IP address configured on this network adapter will be the dedicated IP address for this host.
6. Configure the remaining host parameters as appropriate, then click **Finish**.

**Important**

- To ensure that Network Load Balancing Manager is displaying the most recent host information, right-click the cluster and click **Refresh**. This step is necessary because the host properties that Network Load Balancing Manager displays are a copy of the host properties that were configured the last time Network Load Balancing Manager connected to that host. When you click **Refresh**, Network Load Balancing Manager reconnects to the cluster and displays updated information.

  Similarly, if a host is added to a cluster through the host's local **Network Connections** rather than through Network Load Balancing Manager, the host will not be visible by Network Load Balancing Manager until you **Refresh** the connection as described.

**Notes**

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- When you add additional hosts after the first host is added, the subsequent hosts will automatically inherit the cluster port rules from the initial host.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Configure host parameters
Configure cluster parameters
Configure Port Rules
Network Load Balancing parameters
Nlbmgr

# Connect to existing clusters from Network Load Balancing Manager using a single host

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to existing clusters from Network Load Balancing Manager using a single host**

1. Open Network Load Balancing Manager.
2. Right-click **Network Load Balancing Clusters** and click **Connect to Existing**.
3. Enter the name of one of the cluster hosts and then click **Connect**.
4. After you click **Connect** the Network Load Balancing clusters that exist on that host will be listed at the bottom of the dialog box. Click the cluster name that you want to manage from Network Load Balancing Manager and then click **Finish**.

   **Important**

   - To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

## Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as. If you attempt to connect to an existing Network Load Balancing cluster, either using a single host or using a host list, and are not a member of the Administrators group on any of the cluster's hosts, those hosts will not be displayed in Network Load Balancing Manager.

  If the Network Load Balancing Manager log displays "Access denied, error connecting to "*host*"" you must reconnect to the cluster using a logon name that is a member of the Administrators group on all hosts. You can do this by logging on with a different logon name, using Run as, or providing an alternate logon name for Network Load Balancing Manager to use by clicking **Credentials** in the Network Load Balancing Manager **Options** menu.

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- This procedure allows you to connect to a cluster by entering the name of a single host. After you select a particular cluster, that cluster and all hosts in the cluster will be displayed in Network Load Balancing Manager.
- When using this procedure, certain network problems can prevent Network Load Balancing Manager from connecting to one or more of the other hosts in the cluster. Similarly, if you have misconfigured the cluster by using duplicate priorities on multiple hosts, Network Load Balancing Manager will not be able to connect to those misconfigured hosts. In these cases, some hosts will not appear in Network Load Balancing Manager. If all of the hosts do not appear in Network Load Balancing Manager, try using a host list instead. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Connect to existing clusters from Network Load Balancing Manager using a host list
Stop managing a cluster from Network Load Balancing Manager
Remove a host from a Network Load Balancing cluster
Remove all hosts from a Network Load Balancing cluster
Uninstall Network Load Balancing
Nlbmgr

# Connect to existing clusters from Network Load Balancing Manager using a host list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to existing clusters from Network Load Balancing Manager using a host list**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Network Load Balancing Manager.
2. In the **File** menu, click **Load Host List**.
3. Navigate to and select your host list text file and click **Open**.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- For security reasons, ensure that only users in the local Administrators group have access to the host list file.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.

  If you attempt to connect to an existing Network Load Balancing cluster, either using a single host or using a host list, and are not a member of the Administrators group on any of the cluster's hosts, those hosts will not be displayed in Network Load Balancing Manager.

  If the Network Load Balancing Manager log displays "Access denied, error connecting to "*host*"" you must reconnect to the cluster using a logon name that is a member of the Administrators group on all hosts. You can do this by logging on with a different logon name, using Run as, or providing an alternate logon name for Network Load Balancing Manager to use by clicking **Credentials** in the Network Load Balancing Manager **Options** menu.

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- The host list text file is a text file that you create using Notepad or any other similar application. The file should list the name of each host. Each host name should be separated by a line break. You can include comments in the host list by preceding the comment with a ";".
- This procedure allows you to connect to multiple hosts and all Network Load Balancing clusters on those hosts. The clusters and hosts will then be displayed in Network Load Balancing Manager. This is particularly useful when a host's Network Load Balancing network adapter is having connectivity problems, but another network adapter is available on that host. You can also use this procedure to detect and diagnose problems such as cluster hosts converging separately.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **nlbmgr /hostlist** *host-list*

| Value | Description |
|---|---|
| host-list | The fully qualified path to a text file that you create using Notepad or any other similar application. The file should list the name of each host. Each host name should be separated by a line break. |

**Important**

- For security reasons, ensure that only users in the local Administrators group have access to the host list file.

**Notes**

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

  If you attempt to connect to an existing Network Load Balancing cluster, either using a single host or using a host list, and are not a member of the Administrators group on any of the cluster's hosts, those hosts will not be displayed in Network Load Balancing Manager.

  If the Network Load Balancing Manager log displays "Access denied, error connecting to "*host*"" you must reconnect to the cluster using a logon name that is a member of the Administrators group on all hosts. You can do this by logging on with a different logon name, using Run as, or providing an alternate logon name for Network Load Balancing Manager to use by clicking **Credentials** in the Network Load Balancing Manager **Options** menu.

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- To view the complete syntax for this command, at a command prompt, type:

  **nlbmgr /help**

- Ensure that only users in the local Administrators group have access to the host list file.

- You can include comments in the host list by preceding the comment with a ";".

- This procedure allows you to connect to multiple hosts and all Network Load Balancing clusters on those hosts. The clusters and hosts will then be displayed in Network Load Balancing Manager. This is particularly useful when a host's Network Load Balancing network adapter is having connectivity problems, but another network adapter is available on

that host. You can also use this procedure to detect and diagnose problems such as cluster hosts converging separately.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a single host
Stop managing a cluster from Network Load Balancing Manager
Remove a host from a Network Load Balancing cluster
Remove all hosts from a Network Load Balancing cluster
Uninstall Network Load Balancing
Nlbmgr

# Enable Network Load Balancing Manager logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Network Load Balancing Manager logging**

1. Open Network Load Balancing Manager.
2. Click **Log Settings** in the **Options** menu.
3. Check **Enable logging** and specify a name and location for the log file.

## Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- Network Load Balancing Manager logging is different from the logging of events generated by the Network Load Balancing service. Network Load Balancing Manager logging records only events related to the use of Network Load Balancing Manager. To view events related to the Network Load Balancing service, you must use Event Viewer. For more information, see Related Topics below.
- You should always manage a Network Load Balancing cluster from a trusted and secure computer, particularly when logging is enabled. The Network Load Balancing Manager log file contains potentially sensitive information about the Network Load Balancing cluster and hosts, so it must be properly secured. By default, the log file inherits the security settings of the directory in which it is created, so you might have to change the explicit permissions on the file to restrict read and write access to those individuals who do not need full control of the file. Be aware that the individual using Network Load Balancing Manager does require full control of the log file. For more information, see Related Topics below.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Network Load Balancing Manager
Event logging
Access Control How To...
Nlbmgr

# Remove a host from a Network Load Balancing cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a host from a Network Load Balancing cluster**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click the host on which you want to disable Network Load Balancing and choose **Delete Host**.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

  If you can not use Network Load Balancing Manager, you can disable Network Load Balancing through **Network Connections Properties** on each host.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**. Right-click the **Local Area Connection** on which Network Load Balancing is installed, and then click **Properties**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Nlbmgr
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Remove all hosts from a Network Load Balancing cluster

# Remove all hosts from a Network Load Balancing cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove all hosts from a Network Load Balancing cluster**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click the cluster on which you want to disable Network Load Balancing and choose **Delete Cluster**.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

  If you can not use Network Load Balancing Manager, you can disable Network Load Balancing through **Network Connections Properties** on each host.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**. Right-click the **Local Area Connection** on which Network Load Balancing is installed, and then click **Properties**.
- After performing this procedure, all existing connections to the hosts will be lost and the cluster will no longer exist.
- When Network Load Balancing is enabled for use with a network adapter in multicast mode, a manually configured media access control (MAC) address, which may have been optionally assigned to the adapter, will be removed after Network Load Balancing is later disabled. When the adapter is subsequently reinitialized (typically when the computer is restarted), the network adapter will revert to its factory-assigned MAC address. To avoid this problem, be sure to reassign the adapter's manually configured MAC address after disabling Network Load Balancing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable Network Load Balancing
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Remove all hosts from a Network Load Balancing cluster
Nlbmgr

# Stop managing a cluster from Network Load Balancing Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop managing a cluster from Network Load Balancing Manager**

1. Open Network Load Balancing Manager.
2. Right-click the cluster that you want to remove from Network Load Balancing Manager and choose **Remove from view**.

## Notes

- Performing this task does not require you to have administrative credentials. Therefore, as a security best practice, consider performing this task as a user without administrative credentials.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To ensure that Network Load Balancing Manager is displaying the most recent host information, right-click the cluster and click **Refresh**. This step is necessary because the host properties that Network Load Balancing Manager displays are a copy of the host properties that were configured the last time Network Load Balancing Manager connected to that host. When you click **Refresh**, Network Load Balancing Manager reconnects to the cluster and displays updated information.
- This procedure only removes visibility to the cluster from Network Load Balancing Manager. The cluster still exists and the Network Load Balancing service on each of the cluster's hosts is unaffected.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Remove a host from a Network Load Balancing cluster
Remove all hosts from a Network Load Balancing cluster
Uninstall Network Load Balancing
Nlbmgr

# Configure Network Load Balancing Parameters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure Network Load Balancing parameters**

- Configure cluster parameters
- Enable multicast support
- Enable Internet Group Management Protocol (IGMP) support
- Configure host parameters
- Configure Port Rules

# Configure cluster parameters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure cluster parameters**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click the cluster and choose **Cluster Properties**.
4. Click the **Cluster Parameters** tab.
5. On the **Cluster Parameters** tab, specify values for **IP address**, **Subnet mask**, **Full Internet name**, **Cluster Operation Mode**, **Remote control**, **Remote password** and **Confirm password** using information from the enabling and configuring checklist.

**Important**

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To ensure that Network Load Balancing Manager is displaying the most recent host information, right-click the cluster and click **Refresh**. This step is necessary because the host properties that Network Load Balancing Manager displays are a copy of the host properties that were configured the last time Network Load Balancing Manager connected to that host. When you click **Refresh**, Network Load Balancing Manager reconnects to the cluster and displays updated information.
- The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. It is highly recommended that you do not enable remote control and instead use Network Load Balancing Manager or other remote management tools such as Windows Management Instrumentation (WMI).




If you choose to enable remote control, it is vital that you restrict access by specifying a strong remote control password. It is also imperative that you use a firewall to protect the Network Load Balancing UDP control ports (the ports receiving remote control commands) in order to shield them from outside intrusion. By default, these are ports 1717 and 2504 at the cluster IP address. Use remote control only from a secure, trusted computer within your firewall. For more information on the remote control parameter, see Remote control in Network Load Balancing parameters. For more information about strong passwords, see Strong passwords.

**Notes**

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

- You only need to perform this procedure on a single host. All hosts will automatically inherit the cluster configuration from the initial host.

- The parameters set in the properties dialog box are recorded in the registry on each host.

- Changes to Network Load Balancing parameters are applied when you click **OK** in the properties dialog box. Clicking **OK** stops Network Load Balancing (if it is running), reloads the parameters, and then restarts cluster operations.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Configure host parameters
Configure Port Rules
Network Load Balancing parameters
Checklist: Enabling and configuring Network Load Balancing

# Enable multicast support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable multicast support**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click the cluster, and then click **Cluster Properties**.
4. On the **Cluster Parameters** tab, in **Cluster operation mode**, select **Multicast**. If appropriate, you can also enable Internet Group Management Protocol (IGMP) support by clicking the **IGMP multicast** check box.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- You can only enable IGMP support when your cluster is configured for multicast mode.
- If you enable IGMP support, the allowed multicast IP address is restricted to the standard class D range; that is, 224.0.0.0 to 239.255.255.255.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Network Load Balancing parameters
Multiple network adapters
Single network adapter
Nlbmgr

# Enable Internet Group Management Protocol (IGMP) support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Internet Group Management Protocol (IGMP) support**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click the cluster, and then click **Cluster Properties**.
4. On the **Cluster Parameters** tab, click the **IGMP multicast** check box to enable IGMP support on all hosts in the cluster.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- You can only enable IGMP support when your cluster is configured for multicast mode.
- If you enable IGMP support, the allowed multicast IP address is restricted to the standard class D range; that is, 224.0.0.0 to 239.255.255.255.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Multiple network adapters
Single network adapter
Enable multicast support
Nlbmgr

# Configure host parameters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure host parameters**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click the host that you want to configure, and then click **Host Properties**.
4. Click the **Host Parameters** tab.
5. Specify values for **Priority (unique host identifier)**, **IP address**, and **Subnet mask**, and select an initial host state.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To ensure that Network Load Balancing Manager is displaying the most recent host information, right-click the cluster and click **Refresh**. This step is necessary because the host properties that Network Load Balancing Manager displays are a copy of the host properties that were configured the last time Network Load Balancing Manager connected to that host. When you click **Refresh**, Network Load Balancing Manager reconnects to the cluster and displays updated information.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- The parameters set in the **Network Load Balancing Properties** dialog box are recorded in the registry on each host.
- Changes to Network Load Balancing parameters are applied when you click **OK** in the **Network Load Balancing Properties** dialog box. Clicking **OK** stops Network Load Balancing (if it is running), reloads the parameters, and then restarts cluster operations.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list

Connect to existing clusters from Network Load Balancing Manager using a single host
Configure cluster parameters
Configure Port Rules
Network Load Balancing parameters
Checklist: Enabling and configuring Network Load Balancing

# Configure Port Rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configure port rules**

Before you proceed, it is vital that you review the port rules section of Checklist: Enabling and configuring Network Load Balancing. It is highly recommended that you read through this section before configuring Network Load Balancing.

- Create a new port rule
- Edit an existing port rule
- Remove a port rule

# Create a new port rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new port rule**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click the cluster, and then click **Cluster Properties**.
4. Click the **Port Rules** tab.
5. Click **Add**.
6. Using information from the enabling and configuring checklist, specify values for the following:
   - **Cluster IP address**, which is the virtual IP address that you want this rule to apply to. Enter a specific virtual IP address to create a virtual cluster, or check **All** to apply the rule to all virtual IP addresses.
   - **Port range**
   - **Protocols**
   - **Filtering mode** and as appropriate, values for **Affinity**, and **Load weight**, and then click **OK**.

**Important**

- The number and type of rules must be exactly the same for each host in the cluster. This is best achieved by using Network Load Balancing Manager to configure port rules (see second note below). If you are using Network Load Balancing Manager, when you add additional hosts, they will automatically inherit the cluster port rules from the initial host.

  If a host attempts to join the cluster with a different number of rules, or with different rules from the other hosts, it is not accepted as part of the cluster. The rest of the cluster then continues to handle the traffic as before. At the same time, a message is entered into the Windows event log. When this happens, consult the event log to determine which host is in conflict with the other cluster hosts, resolve the conflict, and then restart Network Load Balancing on that host.

- In order to allow Network Load Balancing to properly handle IP fragments, you should avoid using **None** when selecting UDP or Both for your protocol setting.

- You cannot add more than 32 port rules to a Network Load Balancing cluster.

**Notes**

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

- If all of your hosts are running a product in the Windows Server 2003 family, you can specify port rules to apply to only specific IP addresses. This is useful when creating virtual clusters. If you are not running a Windows Server 2003 family product, you must apply the port rules to all IP addresses. For more information about virtual clusters, see "Understanding Virtual Clusters."

- The list of all currently installed port rules is sorted by port range.

- The parameters set in the properties dialog box are recorded in the registry on each host.

- Changes to Network Load Balancing parameters are applied when you click **OK** in the properties dialog box. Clicking **OK** stops Network Load Balancing (if it is running), reloads the parameters, and then restarts cluster operations.

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

- To ensure that Network Load Balancing Manager is displaying the most recent host information, right-click the cluster and click **Refresh**. This step is necessary because the host properties that Network Load Balancing Manager displays are a copy of the host properties that were configured the last time Network Load Balancing Manager connected to that host. When you click **Refresh**, Network Load Balancing Manager reconnects to the cluster and displays updated information.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Configure cluster parameters
Configure host parameters
Network Load Balancing parameters
Edit an existing port rule
Remove a port rule
Checklist: Enabling and configuring Network Load Balancing
Understanding virtual clusters
Virtual clusters
Nlbmgr

# Edit an existing port rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To edit an existing port rule**

1. Open Network Load Balancing Manager.
2. Right-click the cluster and then click **Cluster Properties**.
3. Click the **Port Rules** tab.
4. In the Defined port rules list, click a rule then click **Edit**.
5. Modify the cluster IP address that you want this rule to apply to, the port range, protocols and filtering mode parameters as required, and then click **OK**.

   For more information, see See Also.

### Important

- The number and type of rules must be exactly the same for each host in the cluster. This is best achieved by using Network Load Balancing Manager to configure port rules (see second note below). If you are using Network Load Balancing Manager, when you add additional hosts, they will automatically inherit the cluster port rules from the initial host.

  If a host attempts to join the cluster with a different number of rules, or with different rules from the other hosts, it is not accepted as part of the cluster. The rest of the cluster then continues to handle the traffic as before. At the same time, a message is entered into the Windows event log. When this happens, consult the event log to determine which host is in conflict with the other cluster hosts, resolve the conflict, and then restart Network Load Balancing on that host.

- You cannot add more than 32 port rules to a Network Load Balancing cluster.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To ensure that Network Load Balancing Manager is displaying the most recent host information, right-click the cluster and click **Refresh**. This step is necessary because the host properties that Network Load Balancing Manager displays are a copy of the host properties that were configured the last time Network Load Balancing Manager connected to that host. When you click **Refresh**, Network Load Balancing Manager reconnects to the cluster and displays updated information.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

- If all of your hosts are running a product in the Windows Server 2003 family, you can specify port rules to apply to only specific IP addresses. This is useful when creating virtual clusters. If you are not running a Windows Server 2003 family product, you must apply the port rules to all IP addresses. For more information about virtual clusters, see "Understanding Virtual Clusters."

- The list of all currently installed port rules is sorted by port range.

- The parameters set in the properties dialog box are recorded in the registry on each host.

- Changes to Network Load Balancing parameters are applied when you click **OK** in the properties dialog box. Clicking **OK** stops Network Load Balancing (if it is running), reloads the parameters, and then restarts cluster operations.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure cluster parameters
Configure host parameters
Checklist: Enabling and configuring Network Load Balancing
Network Load Balancing parameters
Create a new port rule
Remove a port rule
Understanding virtual clusters
Nlbmgr

# Remove a port rule

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove a port rule**

1. Open Network Load Balancing Manager.
2. Right-click the cluster, and then click **Cluster Properties**.
3. Click the **Port Rules** tab.
4. Click the rule you want to remove, and then click **Remove**.

**Important**

- The number and type of rules must be exactly the same for each host in the cluster. This is best achieved by using Network Load Balancing Manager to configure port rules (see second note below). If you are using Network Load Balancing Manager, when you add additional hosts, they will automatically inherit the cluster port rules from the initial host.

  If a host attempts to join the cluster with a different number of rules, or with different rules from the other hosts, it is not accepted as part of the cluster. The rest of the cluster then continues to handle the traffic as before. At the same time, a message is entered into the Windows event log. When this happens, consult the event log to determine which host is in conflict with the other cluster hosts, resolve the conflict, and then restart Network Load Balancing on that host.

- You cannot add more than 32 port rules to a Network Load Balancing cluster.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To ensure that Network Load Balancing Manager is displaying the most recent host information, right-click the cluster and click **Refresh**. This step is necessary because the host properties that Network Load Balancing Manager displays are a copy of the host properties that were configured the last time Network Load Balancing Manager connected to that host. When you click **Refresh**, Network Load Balancing Manager reconnects to the cluster and displays updated information.

**Notes**

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- If all of your hosts are running a product in the Windows Server 2003 family, you can specify port rules to apply to only specific IP addresses. This is useful when creating virtual clusters. If you are not running a Windows Server 2003 family product, you must apply the port rules to all IP addresses. For more information about virtual clusters, see "Understanding Virtual Clusters."

- The list of all currently installed port rules is sorted by port range.
- The parameters set in the properties dialog box are recorded in the registry on each host.
- Changes to Network Load Balancing parameters are applied when you click **OK** in the properties dialog box. Clicking **OK** stops Network Load Balancing (if it is running), reloads the parameters, and then restarts cluster operations.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure cluster parameters
Configure host parameters
Checklist: Enabling and configuring Network Load Balancing
Create a new port rule
Edit an existing port rule
Understanding virtual clusters
Nlbmgr

# Perform Basic Network Load Balancing Operations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Perform basic Network Load Balancing operations**

- Immediately stop handling all Network Load Balancing cluster traffic
- Restart handling Network Load Balancing cluster traffic
- Suspend Network Load Balancing
- Resume Network Load Balancing
- Disable and block all Network Load Balancing traffic handling for specific ports
- Enable Network Load Balancing traffic handling for specific ports
- Disable new traffic handling for specific ports
- Disable new traffic handling for specific ports and stop Network Load Balancing

# Immediately stop handling all Network Load Balancing cluster traffic

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To immediately stop handling all Network Load Balancing cluster traffic**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. If you want to stop cluster operations on all cluster hosts, right-click the cluster, and then point to **Control Hosts**.

   Or, if you want to stop the cluster service on only a particular host, right-click the host, and then point to **Control Host**.
4. Click **Stop**

Cluster operations are immediately stopped and all existing connections are immediately closed.

**Important**

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

**Notes**

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- When you use the **stop** command, either through Network Load Balancing Manager or the command prompt, any client connections already in progress are interrupted. To avoid interrupting active connections, consider using the **drainstop** cluster-control command, which allows the host to continue servicing active connections, but disables all new traffic to that host. For more information, see "Disable new traffic handling and stop Network Load Balancing" in Related Topics.
- For information on how to restart cluster operations, see "To start Network Load Balancing" in Related Topics.
- The **stop** command differs from the **suspend** command in that **suspend** stops Network Load Balancing on the host and also suspends all Network Load Balancing cluster-control commands on the host, except for the **resume** and **query** commands. **Stop** only stops Network Load Balancing on the host, but does not affect the other Network Load Balancing cluster-control commands. For more information on the **suspend** command, see "To suspend Network Load Balancing".

**Using a command line**

1. Open Command Prompt.

2. Type:

   **nlb.exe stop**

   Cluster operations are immediately stopped on the local cluster host and all existing connections are immediately closed.

   Or, if you want to stop cluster operations on all cluster hosts, type:

   **nlb.exe stop***cluster IP address* | *cluster name*

### Notes

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- When you use the **stop** command, either through Network Load Balancing Manager or the command prompt, any client connections already in progress are interrupted. To avoid interrupting active connections, consider using the **drainstop** cluster-control command, which allows the host to continue servicing active connections, but disables all new traffic to that host. For more information, see "Disable new traffic handling and stop Network Load Balancing" in Related Topics.

- For information on how to restart cluster operations, see "To start Network Load Balancing" in Related Topics.

- The **stop** command differs from the **suspend** command in that **suspend** stops Network Load Balancing on the host and also suspends all Network Load Balancing cluster-control commands on the host, except for the **resume** and **query** commands. **Stop** only stops Network Load Balancing on the host, but does not affect the other Network Load Balancing cluster-control commands. For more information on the **suspend** command, see "To suspend Network Load Balancing" in Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Suspend Network Load Balancing
Nlb
Disable new traffic handling for specific ports and stop Network Load Balancing
Restart handling Network Load Balancing cluster traffic
Nlbmgr

# Restart handling Network Load Balancing cluster traffic

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restart handling Network Load Balancing cluster traffic**

Using the Windows interface

Using a command line

**Using the Windows interface**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. If you want to start cluster operations on all cluster hosts, right-click the cluster, and then point to **Control Hosts**.

   Or, if you want to start the cluster service on only a particular host, right-click the host, and then point to **Control Host**.
4. Click **Start**

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- This command can be used after a previous **stop** or **suspend** command. It restarts cluster operations. It does not reenable the use of cluster-control commands, including remote control commands, which have been disabled by a previous **suspend** command. For information on how to reenable cluster-control commands, see "To resume Network Load Balancing" in Related Topics.
- You should not use the **start** command to form a new cluster. For more information, see "Enable Network Load Balancing".

**Using a command line**

1. Open Command Prompt.
2. Type:

   **nlb.exe start**

   Or, if you want to start cluster operations on all cluster hosts, type:

   **nlb.exe start***cluster IP address* | *cluster name*

## Notes

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- This command can be used after a previous **stop** or **suspend** command. It restarts cluster operations. It does not reenable the use of cluster-control commands, including remote control commands, which have been disabled by a previous **suspend** command. For information on how to reenable cluster-control commands, see "To resume Network Load Balancing" in Related Topics.

- You should not use the **start** command to form a new cluster. For more information, see "Enable Network Load Balancing".

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Nlb
Disable new traffic handling for specific ports and stop Network Load Balancing
Immediately stop handling all Network Load Balancing cluster traffic
Enable Network Load Balancing traffic handling for specific ports
Disable and block all Network Load Balancing traffic handling for specific ports
Enable Network Load Balancing
Resume Network Load Balancing
Nlbmgr

# Suspend Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To suspend Network Load Balancing**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. If you want to suspend cluster operations on all cluster hosts, right-click the cluster and then point to **Control Hosts**.

   Or, if you want to suspend the cluster service on only a particular host, right-click the host, and then point to **Control Host**.
4. Click **Suspend**

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- The **suspend** command suspends all cluster operations until the **resume** and **start** commands are issued.

  The purpose of the **suspend** command is to override any remote control commands that might be issued. All subsequent cluster-control commands except **resume** and **query** are ignored.
- This command can be performed on the whole cluster or remotely to a single host.
- The **stop** command differs from the **suspend** command in that **suspend** stops Network Load Balancing on the host and also suspends all Network Load Balancing cluster-control commands on the host, except for the **resume** and **query** commands. **Stop** only stops Network Load Balancing on the host, but does not affect the other Network Load Balancing cluster-control commands. For more information on the **stop** command, see "Immediately stop handling all Network Load Balancing cluster traffic" in Related Topics.

**Using a command line**

1. Open Command Prompt.

2. Type:

   **nlb.exe suspend**

   Or, if you want to suspend cluster operations on all cluster hosts, type:

   **nlb.exe suspend***cluster IP address* | *cluster name*

### Notes

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- The **suspend** command suspends all cluster operations until the **resume** and **start** commands are issued.

  The purpose of the **suspend** command is to override any remote control commands that might be issued. All subsequent cluster-control commands except **resume** and **query** are ignored.

- This command can be performed on the whole cluster or remotely to a single host.

- The **stop** command differs from the **suspend** command in that **suspend** stops Network Load Balancing on the host and also suspends all Network Load Balancing cluster-control commands on the host, except for the **resume** and **query** commands. **Stop** only stops Network Load Balancing on the host, but does not affect the other Network Load Balancing cluster-control commands. For more information on the **stop** command, see "Immediately stop handling all Network Load Balancing cluster traffic" in Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Nlb
Resume Network Load Balancing
Restart handling Network Load Balancing cluster traffic
Immediately stop handling all Network Load Balancing cluster traffic
Enable Network Load Balancing traffic handling for specific ports
Disable and block all Network Load Balancing traffic handling for specific ports
Nlbmgr

# Resume Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To resume Network Load Balancing

- Using the Windows interface
- Using a command line

#### Using the Windows interface

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click a host or a cluster, and then point to **Control Host** or **Control Hosts**.
4. Click **Resume**

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- Performing this command at the cluster level resumes Network Load Balancing on all hosts within the cluster. Performing this command at the host level resumes Network Load Balancing on the specific host only.
- This command can be used after a previous **suspend** command. It reenables the use of cluster-control commands, including remote control commands. It does not restart cluster operations. For information on how to restart cluster operations, see "Restart handling Network Load Balancing cluster traffic" in Related Topics.
- This command can be performed on the whole cluster or remotely to a single host.

#### Using a command line

1. Open Command Prompt.
2. Type:

   **nlb.exe resume**

   Or, if you want to reenable cluster-control commands on all cluster hosts, type:

**nlb.exe resume***cluster IP address* | *cluster name*

## Notes

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- This command can be used after a previous **suspend** command. It reenables the use of cluster-control commands, including remote-control commands. It does not restart cluster operations. For information on how to restart cluster operations, see "Restart handling Network Load Balancing cluster traffic" in Related Topics.

- This command can be performed on the whole cluster or remotely to a single host.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Nlb
Suspend Network Load Balancing
Restart handling Network Load Balancing cluster traffic
Nlbmgr

# Disable and block all Network Load Balancing traffic handling for specific ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable and block all Network Load Balancing traffic handling for specific ports**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click a host or a cluster, and then click **Control Ports**.
4. Select the port rule that contains the ports that you want to disable and then click **Disable**.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- This command disables and immediately blocks all traffic handling for the rule whose port range contains the specified port. All ports specified by the port rule are affected.
- All active connections on the specified hosts are blocked. To maintain active connections, use the drain function instead.
- This has no effect if the specified hosts have not started cluster operations.
- For information on how to enable specific ports, see "Enable Network Load Balancing traffic handling for specific ports" in Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

**nlb.exe disable***vip*[:*port*|:all] | all[:*port*|:all] *cluster*[:*host*]| all [*local*|*global*]

For more information about the parameters that you can use with this command, see "Nlb" in Related Topics.

## Notes

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- This command disables and immediately blocks all traffic handling for the rule whose port range contains the specified port. All ports specified by the port rule are affected.

- All active connections on the specified hosts are blocked. To maintain active connections, use the drain function instead.

- This has no effect if the specified hosts have not started cluster operations.

- For information on how to enable specific ports, see "Enable Network Load Balancing traffic handling for specific ports" in Related Topics.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Nlb
Disable new traffic handling for specific ports and stop Network Load Balancing
Immediately stop handling all Network Load Balancing cluster traffic
Enable Network Load Balancing traffic handling for specific ports
Disable and block all Network Load Balancing traffic handling for specific ports
Nlbmgr

# Enable Network Load Balancing traffic handling for specific ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Network Load Balancing traffic handling for specific ports**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click a host or a cluster, and then click **Control Ports**.
4. Select the port rule that contains the ports that you want to enable and then click **Enable**.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- This command enables traffic handling for the rule whose port range contains the specified port. All ports specified by the port rule are affected.
- This has no effect if the specified hosts have not started cluster operations.
- For information on how to disable specific ports, see "Disable and block all Network Load Balancing traffic handling for specific ports" in Related Topics.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **nlb.exe enable***vip*[:*port*|:all] | all[:*port*|:all] *cluster*[:*host*]| all [*local*|*global*]

For more information about the parameters that you can use with this command, see "Nlb" in Related Topics.

### Notes

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- This command enables traffic handling for the rule whose port range contains the specified port. All ports specified by the port rule are affected.

- This has no effect if the specified hosts have not started cluster operations.

- For information on how to disable specific ports, see "Disable and block all Network Load Balancing traffic handling for specific ports" in Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Nlb
Disable new traffic handling for specific ports and stop Network Load Balancing
Restart handling Network Load Balancing cluster traffic
Nlbmgr

# Disable new traffic handling for specific ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable new traffic handling for specific ports**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click a host or a cluster, and then click **Control Ports**.
4. Select the port rule that contains the ports that you want to drain and then click **Drain**.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open Command Prompt. From the command window, type **NLBMgr.exe**, and then press **ENTER**.
- Draining disables new traffic handling for the rule whose port range contains the specified port. New connections to the specified hosts are not allowed, but all active connections are maintained. To disable active connections, use the disable command instead.
- Performing this command at the cluster level drains the specified ports on all hosts within the cluster. Performing this command at the host level drains the ports on the specific host only.
- This command has no effect if the specified hosts have not started cluster operations.

**Using a command line**

1. Open Command Prompt.
2. Type:

   **nlb.exe drain***vip*[:*port*|:all] | all[:*port*|:all] *cluster*[:*host*]| all [*local*|*global*]

   For more information about the parameters that you can use with this command, see "Nlb" in Related Topics.

**Notes**

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- This command disables new traffic handling for the rule whose port range contains the specified port. All ports specified by the port rule are affected.

  If **all** is specified for the port, this command is applied to the ports covered by all port rules.

- Draining can be performed on the entire cluster or remotely to a single host.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Nlb
Disable new traffic handling for specific ports and stop Network Load Balancing
Restart handling Network Load Balancing cluster traffic
Nlbmgr

# Disable new traffic handling for specific ports and stop Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable new traffic handling for specific ports and stop Network Load Balancing**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster. For more information, see "Manage existing clusters from Network Load Balancing Manager" in Related Topics.
3. Right-click a host or a cluster, and then point to **Control Host** or **Control Hosts**.
4. Click **Drainstop**

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open Command Prompt. From the command window, type **nlbmgr.exe**, and then press **ENTER**.
- Performing this command at the cluster level performs the command on all hosts within the cluster. Performing this command at the host level performs the command on the specific host only.
- This command disables all new traffic handling on the specified hosts. While draining, hosts continue to service opened connections and stop their cluster operations when there are no more active connections. Draining mode can be terminated by explicitly stopping cluster mode with the stop command or by restarting new traffic handling with the start command
- While draining, hosts remain in the cluster and stop their cluster operations when there are no more active connections. To stop draining mode, explicitly stop cluster mode with the **stop** command, or restart new traffic handling with the **start** command.

**Using a command line**

1. Open Command Prompt.

2. Type **nlb.exe drainstop**.

   For more information about the parameters that you can use with this command, see "Nlb" in Related Topics.

## Notes

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To run nlb.exe from a remote computer, you must enable remote control on the host that is running Network Load Balancing. The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. Only use remote control on a secure computer within your firewall. Because of the many security risks remote control presents, it is recommended that you do not enable the remote control feature. Instead, use other remote management tools such as Network Load Balancing Manager or Windows Management Instrumentation (WMI).

- This command disables all new traffic handling on the specified hosts. While draining, hosts continue to service opened connections and stop their cluster operations when there are no more active connections. Draining mode can be terminated by explicitly stopping cluster mode with the stop command or by restarting new traffic handling with the start command

- This command can be performed on the whole cluster or remotely to a single host.

- While draining, hosts remain in the cluster and stop their cluster operations when there are no more active connections. To stop draining mode, explicitly stop cluster mode with the **stop** command, or restart new traffic handling with the **start** command.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Nlb
Immediately stop handling all Network Load Balancing cluster traffic
Enable Network Load Balancing traffic handling for specific ports
Disable and block all Network Load Balancing traffic handling for specific ports
Restart handling Network Load Balancing cluster traffic
Nlbmgr

# Perform Advanced Network Load Balancing Procedures

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Perform advanced Network Load Balancing procedures**

- Uninstall Network Load Balancing
- Install Network Load Balancing if it was previously uninstalled
- Adjust convergence parameters
- Limit switch port flooding

# Uninstall Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To uninstall Network Load Balancing**

1. Log on as an administrator on the local computer.

   **Important**

   - If possible, you should avoid uninstalling Network Load Balancing. There is typically no need to uninstall this feature, and doing so might adversely affect the operation of other system components. Removing Network Load Balancing from the network adapter, as described in "Remove a host from a Network Load Balancing cluster" in Related Topics below is all that is needed to ensure that Network Load Balancing is not being used by the host.
   - If you do not follow these steps in the exact order shown, your network will experience IP address conflicts.

2. Open Network Connections.
3. Right-click the **Local Area Connection** on which Network Load Balancing is installed, and then click **Properties**.
4. Click **Internet Protocol (TCP/IP)**, and then click **Properties**.
5. In the **Internet Protocol (TCP/IP) Properties** dialog box, remove the cluster's primary IP address. If the cluster's primary IP address is in the **Advanced TCP/IP Settings** dialog box, click **Advanced**.
6. In the **Advanced TCP/IP Settings** dialog box, remove the cluster's primary IP address and any other cluster IP addresses that you might have entered here, and then click **OK**.
7. In the **Internet Protocol (TCP/IP) Properties** dialog box, click **OK**.
8. In the **Local Area Connection Properties** dialog box, click **OK**.
9. Right-click the **Local Area Connection** on which Network Load Balancing is installed, and then click **Properties**.
10. Click **Network Load Balancing**, and then click **Uninstall**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Remove a host from a Network Load Balancing cluster

Remove all hosts from a Network Load Balancing cluster
Enable Network Load Balancing

# Install Network Load Balancing if it was previously uninstalled

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Network Load Balancing if it was previously uninstalled**

1. Open Network Connections.

   **Important**

   - To avoid networking problems, follow these steps in the exact order shown.

2. Right-click the **Local Area Connection** on which Network Load Balancing is to be installed, and then click **Properties**.

3. Click **Install**.

4. In the **Select Network Component Type** dialog box, click **Service**, and then click **Add**.

5. In the **Select Network Service** dialog box, click **Network Load Balancing**, and then click **OK**.

   **Important**

   - Do not select the check box for Network Load Balancing at this point. When you configure the cluster using Network Load Balancing Manager (in the next step), Network Load Balancing Manager will automatically enable this service.

6. Use Network Load Balancing Manager to configure the cluster. For more information, see "Enable Network Load Balancing" in Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Network Connections, click **Start**, click **Control Panel**, and then double-click **Network Connections**.

- If Network Load Balancing clients are accessing a cluster through a router, be sure that the router meets the following requirements:

  - Accepts a dynamic Address Resolution Protocol (ARP) reply mapping the unicast IP address to its unicast media access control (MAC) address

  - In multicast mode, accepts an ARP reply that has a MAC address in the payload of the ARP structure. If your router does not meet these requirements, you can also create a static ARP entry in the router. Cisco routers require a static ARP entry because they do not support the resolution of unicast IP addresses to multicast MAC addresses.

- If only one **Local Area Connection** icon is displayed in **Network Connections**, you have only one network adapter installed on your computer. Additional network adapters are represented by additional **Local Area Connection** icons.

- When installing Network Load Balancing, you might expect to find a setup application (such as Setup.exe). There is no

such application. The Network Load Balancing driver is installed as a standard Windows networking driver.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure cluster parameters
Configure host parameters
Configure Port Rules
Set up TCP/IP for Network Load Balancing
Checklist: Enabling and configuring Network Load Balancing
Network Load Balancing parameters
Enable Network Load Balancing

# Adjust convergence parameters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To adjust convergence parameters**

1. Click **Start**, click **Run**, and then, in **Open**, type **regedit**.
2. Click the registry key:

   HKEY_LOCAL_MACHINE\System\CurrentControlSet

   \Services\WLBS\Parameters\Interface*Network adapter ID*
3. Edit the **AliveMsgPeriod** or **AliveMsgTolerance** parameters as necessary.
4. Restart the computer.

### Caution

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- The purpose of this procedure is to allow you to control both the message exchange period and the number of missed messages required to initiate convergence. For more information, see Related Topics.
- The **AliveMsgPeriod** value holds the message exchange period in milliseconds, and the **AliveMsgTolerance** value specifies how many exchanged messages from a host can be missed before the cluster initiates convergence.

  These numbers should reflect your installation's requirements. A longer message exchange period reduces the heartbeat traffic needed to maintain fault tolerance, but it increases the time that it takes for Network Load Balancing to stop sending network messages to an offline host. Likewise, increasing the number of message exchanges prior to convergence reduces the number of unnecessary convergence initiations due to network congestion, but it too increases the time for an offline host to stop receiving network traffic.

  Using the default values, 5 seconds are needed to discover a missing host, and another 5 seconds are needed for the cluster to redistribute the load. A total of 10 seconds to stop sending network traffic to an offline host should be acceptable for most TCP/IP applications. This configuration incurs very low networking overhead.
- WLBS stands for Windows NT Load Balancing Service, the former name of Network Load Balancing in Windows NT 4.0. For reasons of backward compatibility, WLBS continues to be used in certain instances.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

How Network Load Balancing works

Nlbmgr

# Limit switch port flooding

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To limit switch port flooding**

1. If multicast support has been enabled during setup, open Network Load Balancing Manager.
2. If Network Load Balancing Manager does not already list the cluster, connect to the cluster.
3. Right-click the cluster, and then click **Cluster Properties**.
4. On the **Cluster Parameters** tab, click the **IGMP multicast** check box to enable Internet Group Management Protocol (IGMP) support on all hosts in the cluster.

### Important

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

### Notes

- When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring, or you must have been delegated the appropriate authority. As a security best practice, consider using Run as to perform this procedure. If you are configuring a cluster or host by running Network Load Balancing Manager from a computer that is not part of the cluster, you do not need to be a member of the Administrators group on that computer. For more information, see Default local groups, Default groups, and Using Run as.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.
- If the cluster hosts are directly connected to a hub to receive client requests and IGMP support is not enabled, incoming client traffic is automatically sent to all switch ports, and may produce switch flooding.
- You can only enable IGMP support when your cluster is configured for multicast mode. To enable multicast mode, see "Enable multicast support" in Related Topics below.
- Windows NT Load Balancing Service (WLBS) is the former name of Network Load Balancing in Windows NT 4.0. For backward compatibility reasons, WLBS continues to be used in certain instances.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Connect to existing clusters from Network Load Balancing Manager using a host list
Connect to existing clusters from Network Load Balancing Manager using a single host
Remove all hosts from a Network Load Balancing cluster
Enable multicast support

Nlb
Nlbmgr

# Network Load Balancing Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Network Load Balancing Overview
- Understanding Network Load Balancing
- Planning Your Network Load Balancing Cluster
- Using Network Load Balancing
- Network Load Balancing Resources

# Network Load Balancing Overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Load Balancing overview**
This section covers:

- Introduction to Network Load Balancing
- Network Load Balancing key features
- Network Load Balancing system requirements

# Introduction to Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introduction to Network Load Balancing**

The Network Load Balancing (NLB) service enhances the availability and scalability of Internet server applications such as those used on Web, FTP, firewall, proxy, VPN, and other mission-critical servers. A single computer running Windows can provide a limited level of server reliability and scalable performance. However, by combining the resources of two or more computers running one of the products in the Windows Server 2003 family into a single cluster, Network Load Balancing can deliver the reliability and performance that Web servers and other mission-critical servers need. The following diagram depicts two connected Network Load Balancing clusters. The first cluster consists of two hosts and the second cluster consists of four hosts:




Each host runs separate copies of the desired server applications, such as that for a Web, FTP, and Telnet server. Network Load Balancing distributes incoming client requests across the hosts in the cluster. The load weight to be handled by each host can be configured as necessary. You can also add hosts dynamically to the cluster to handle increased load. In addition, Network Load Balancing can direct all traffic to a designated single host, called the default host.

Network Load Balancing allows all of the computers in the cluster to be addressed by the same set of cluster IP addresses (but also maintains their existing unique, dedicated IP addresses).

For load-balanced applications, when a host fails or goes offline, the load is automatically redistributed among the computers still operating. Applications with a single server have their traffic redirected to a specific host. When a computer fails or goes offline unexpectedly, active connections to the failed or offline server are lost. However, if you bring a host down intentionally, you can use the drainstop command to service all active connections prior to bringing the computer offline. In either case, when ready, the offline computer can transparently rejoin the cluster and regain its share of the workload.

**Note** If you plan to use Network Load Balancing in a 64-bit environment, you must use the 64-bit Network Load Balancing version. If you do not, the cluster will fail to form.

**Overview of Network Load Balancing configuration**

Network Load Balancing runs as a Windows networking driver. Its operations are transparent to the TCP/IP networking stack. The following diagram shows the relationship between Network Load Balancing and other software components in a typical configuration of a Network Load Balancing host:




**Database access from load-balanced server applications**

Some server applications access a database that is updated by client requests. When these applications are load balanced in

the cluster, these updates need to be properly synchronized. Each host can use local, independent copies of databases that are merged offline as necessary. Alternatively, the clustered hosts can share access to a separate, networked database server. A combination of these approaches can also be used. For example, static Web pages can be replicated among all clustered servers to ensure fast access and complete fault tolerance. However, database requests would be forwarded to a common database server that handles updates for multiple Web servers.

Some mission-critical applications might require the use of highly available database engines to ensure complete fault tolerance for the service. It is recommended that you deploy cluster-aware database software to deliver highly available and scalable database access within an overall clustering scheme. One such example of this is Microsoft SQL Server, which can be deployed with the Cluster service in a server cluster. The Cluster service ensures that if one node fails, a remaining node assumes the responsibilities of the failed computer, thus providing almost continuous service to Microsoft SQL Server clients. It is able to do this because the computers in the server cluster make use of a cluster storage device. For more information on the Cluster service and how it works with Network Load Balancing, see Updated technical information.

**Notes**

- It is important to distinguish between the two cluster solutions under discussion. The first, Network Load Balancing, is intended primarily to load balance incoming TCP/IP traffic. The computers participating in this solution form one type of cluster. The second, the Cluster service, is intended primarily to provide failover service from one computer to another. The computers participating in this solution form a different type of cluster. Moreover, the Network Load Balancing cluster would most commonly be running Web server applications. In contrast, the Cluster service would most commonly be running database applications (when used in conjunction with Network Load Balancing). Network Load Balancing and the Cluster service can not both be active on the same computer, but by joining the two cluster solutions together to function in a complementary fashion, the user creates an overall clustering scheme, as shown in the following diagram:




- For more information on how Network Load Balancing achieves fault tolerance and scalability, see How Network Load Balancing works.

# Network Load Balancing key features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Load Balancing key features**

**Scalability**

- Load balances requests for individual TCP/IP services across the cluster.
- Supports up to 32 computers in a single cluster.
- Load balances multiple server requests, from either the same client, or from several clients, across multiple hosts in the cluster.
- Fully pipe-lined implementation ensures high performance and low overhead.

  **Note**

  - During packet reception, the Network Load Balancing fully pipe-lined implementation overlaps the delivery of incoming packets to TCP/IP and the reception of other packets by the NDIS driver. This increases the overall processing speed and reduces latency because TCP/IP can process a packet while the NDIS driver receives a subsequent packet. It also reduces the overhead required for TCP/IP and the NDIS driver to coordinate their actions, and in many cases it eliminates an extra memory copy of packet data. During packet sending, Network Load Balancing also enhances throughput and reduces latency and overhead by increasing the number of packets that TCP/IP can send with one NDIS call. To achieve these performance enhancements, Network Load Balancing allocates and manages a pool of packet buffers and descriptors, that it uses to overlap the actions of TCP/IP and the NDIS driver.

**High Availability**

- Automatically detects and recovers from a failed or offline computer.
- Automatically balances the network load when hosts are added or removed.
- Recovers and redistributes the workload within 10 seconds.

**Manageability**

- You can manage and configure multiple Network Load Balancing clusters and the cluster hosts from a single computer using Network Load Balancing Manager.
- You can specify the load balancing behavior for a single IP port or group of ports using port management rules.
- If you use the same set of load balanced servers for multiple applications or Web sites, by using virtual clusters you can define different port rules for each Web site based on the destination virtual IP address.
- Using optional single-host rules, you can direct all client requests to a single host, in effect using Network Load Balancing to route client requests to a particular host running specific applications.
- You can block undesired network access to certain IP ports.
- You can enable Internet Group Management Protocol (IGMP) support on the cluster hosts to control switch flooding when operating in multicast mode.




- Network Load Balancing logs all actions and cluster changes in the Windows event log.

- Using shell commands or scripts, you can remotely start, stop, and control Network Load Balancing actions from any networked computer that is running Windows.

**Ease of use**

- Network Load Balancing is installed as a standard Windows networking driver component.

- Network Load Balancing requires no hardware changes to enable and run.

- Network Load Balancing Manager allows you to create new Network Load Balancing clusters and configure and manage clusters and all of the cluster's hosts from a single remote or local computer. Ideally you should use a second network adapter when managing the cluster from a local computer.

- Network Load Balancing lets clients access the cluster with a single logical Internet name and virtual IP address (also know as the cluster IP address) while retaining individual names for each computer. Network Load Balancing allows multiple virtual IP addresses for multihomed servers (although in the case of virtual clusters, the servers do not need to be multihomed in order to have multiple virtual IP addresses).

- Network Load Balancing can be bound to multiple network adapters allowing you to configure multiple independent clusters on each host. Support for multiple network adapters is different from virtual clusters in that virtual clusters allow you to configure multiple clusters on a single network adapter.




- You do not have to modify server applications to run in a Network Load Balancing cluster.

- If a cluster host fails and then is subsequently brought back online, Network Load Balancing can be configured to automatically add that host to the cluster. The added host will then be able to start handling new server requests from clients.

- You can take computers offline for preventive maintenance without disturbing cluster operations on the other hosts.

# Network Load Balancing system requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Load Balancing system requirements**

Network Load Balancing is designed to work as a standard networking device driver in the Microsoft® Windows Server 2003 family of products. Because Network Load Balancing provides clustering support for TCP/IP-based server applications, TCP/IP must be installed in order to take advantage of Network Load Balancing functionality. The current version of Network Load Balancing is designed for use with Ethernet network adapters. It is not compatible with asynchronous transfer mode (ATM), ATM local area network (LAN) emulation, or token ring. It has been tested on 10 megabits per second (Mbps), 100 Mbps, and gigabit Ethernet networks with a wide variety of network adapters, including teaming network adapters listed as compatible with the Windows Server 2003 family of products. For information regarding compatible hardware, click the appropriate link in Support resources.

On x86-based computers, Network Load Balancing uses between 750 kilobytes (KB) and 27 MB of RAM per network adapter during operation, using the default parameters and depending on the network load. The parameters can be modified to allow up to 84 MB memory to be used. Typical memory usage ranges between 750 KB and 2 MB.

On x64-based computers, Network Load Balancing uses between 825 KB and 32.3 MB of RAM per network adapter during operation, using the default parameters and depending on the network load. The parameters can be modified to allow up to 102 MB memory to be used. Typical memory usage for x64-based computers ranges between 825 KB and 2.5 MB.

For optimum cluster performance, you should install a second network adapter on each Network Load Balancing host. In this configuration, the first network adapter handles the network traffic addressed to the server as part of cluster, while the second network adapter is used for intra-host communication.

It is important to understand that it is possible to use Network Load Balancing with only a single network adapter. For more details, see Single network adapter limitations.

**Note**

- Network Load Balancing can operate in two modes: unicast and multicast. Unicast support is enabled by default. To enable multicast support, see Enable multicast support.

**Using a router**

If Network Load Balancing clients are accessing a cluster through a router when the cluster has been configured to operate in multicast mode, be sure that the router meets the following requirements:

- Accepts an Address Resolution Protocol (ARP) reply that has one media access control (MAC) address in the payload of the ARP structure, but appears to arrive from a station with another MAC address, as determined by the Ethernet header

- Accepts an ARP reply for a unicast IP address with a multicast MAC address in the payload of its ARP structure

These conditions allow the router to map the cluster IP addresses to the corresponding MAC address. If your router does not meet these requirements, you can also create a static ARP entry in the router. Cisco routers require a static ARP entry because they do not support the resolution of unicast IP addresses to multicast MAC addresses.

# Understanding Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Network Load Balancing**


- How Network Load Balancing works
- Understanding virtual clusters

# How Network Load Balancing works

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How Network Load Balancing works**

Network Load Balancing provides high availability and scalability of servers using a cluster of two or more host computers working together. Internet clients access the cluster using a either a IP address or a set of addresses. The clients are unable to distinguish the cluster from a single server. Server applications do not identify that they are running in a cluster. However, a Network Load Balancing cluster differs significantly from a single host running a single server application because it can provide uninterrupted service even if a cluster host fails. The cluster can also respond more quickly to client requests than a single host.

Network Load Balancing delivers high availability by redirecting incoming network traffic to working cluster hosts if a host fails or is offline. Existing connections to an offline host are lost, but the Internet services remain available. In most cases (for example, with Web servers), client software automatically retries the failed connections, and the clients experience a delay of only a few seconds in receiving a response.

Network Load Balancing delivers scaled performance by distributing the incoming network traffic among one or more virtual IP addresses (the cluster IP addresses) assigned to the Network Load Balancing cluster. The hosts in the cluster then concurrently respond to different client requests, even multiple requests from the same client. For example, a Web browser might obtain each of the multiple images in a single Web page from different hosts within a Network Load Balancing cluster. This speeds up processing and shortens the response time to clients.

Network Load Balancing enables all cluster hosts on a single subnet to concurrently detect incoming network traffic for the cluster IP addresses. On each cluster host, the Network Load Balancing driver acts as a filter between the cluster adapter driver and the TCP/IP stack in order to distribute the traffic across the hosts.

Network Load Balancing employs a fully distributed algorithm to statistically map incoming clients to the cluster hosts based on their IP address and port. No communication between the hosts is necessary for this process to occur. When inspecting an arriving packet, all hosts simultaneously perform this mapping to quickly determine which host should handle the packet. The mapping remains invariant unless the number of cluster hosts changes. The Network Load Balancing filtering algorithm is much more efficient in its packet handling than centralized load-balancing applications, which must modify and retransmit packets.

**Distribution of cluster traffic**

Network Load Balancing controls the distribution of TCP and UDP traffic from the Internet clients to selected hosts within a cluster as follows: After Network Load Balancing has been configured, incoming client requests to the cluster IP addresses are received by all hosts within the cluster. Network Load Balancing filters incoming datagrams to specified TCP and UDP ports before these datagrams reach the TCP/IP protocol software. Network Load Balancing manages the TCP and UDP protocols within TCP/IP, controlling their actions on a per-port basis.

In multicast mode, Network Load Balancing can limit switch flooding by providing Internet Group Management Protocol (IGMP) support. Network Load Balancing does not control any incoming IP traffic other than TCP and UDP traffic for specified ports and IGMP traffic in multicast mode. It does not filter other IP protocols (for example, ICMP or ARP), except as described below. Be aware that you should expect to see duplicate responses from certain point-to-point TCP/IP applications (such as ping) when the cluster IP address is used. If required, these applications can use the dedicated IP address for each host to avoid this behavior.

**Convergence**

To coordinate their actions, Network Load Balancing hosts periodically exchange heartbeats within the cluster (for more information, see Internet Group Management Protocol (IGMP)). IP multicasting allows the hosts to monitor the status of the cluster. When the state of the cluster changes (such as when hosts fail, leave, or join the cluster), Network Load Balancing invokes a process known as convergence, in which the hosts exchange a limited number of messages to determine a new, consistent state of the cluster and to designate the host with the highest host priority as the new default host. When all cluster hosts have reached consensus on the correct new state of the cluster, they record the completion of convergence in the Windows event log. This process typically takes less than 10 seconds to complete.

During convergence, the remaining hosts continue to handle incoming network traffic. Client requests to working hosts are unaffected. At the completion of convergence, the traffic destined for a failed host is redistributed to the remaining hosts. Load-balanced traffic is repartitioned among the remaining hosts to achieve the best possible new load balance for specific TCP or UDP ports.

If a host is added to the cluster, convergence allows this host to receive its share of the load-balanced traffic. Expansion of the cluster does not affect ongoing cluster operations and is achieved transparently to both Internet clients and to server applications. However, it might affect client sessions that span multiple TCP connections when client affinity is selected, because clients might be remapped to different cluster hosts between connections. For more information on affinity, see Network Load Balancing and stateful connections.

Network Load Balancing assumes that a host is functioning properly within the cluster as long as it exchanges heartbeats with other cluster hosts. If other hosts do not receive a response from any member for several heartbeat exchanges, they initiate convergence to redistribute the load that would have been handled by the failed host.

You can control both the message exchange period and the number of missed messages required to initiate convergence. The default values are respectively set to 1,000 milliseconds (1 second) and 5 missed message exchange periods. Because these parameters are not usually modified, they are not configurable through the **Network Load Balancing Properties** dialog box. They can be adjusted manually in the registry as necessary. The procedures for this are described in Adjust convergence parameters.

# Understanding virtual clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding virtual clusters**

It is becoming increasingly common to host multiple applications or Web sites on a single Network Load Balancing (NLB) cluster. When doing so, you may require independent load balancing policies, defined through port rules, applied to each of these applications or sites.

In the Windows Server 2003 family of products you can configure multiple Network Load Balancing clusters on the same network adapter and then apply specific port rules to each of those IP addresses. These are referred to as "virtual clusters."

You can also use virtual clusters to block network traffic to a specific host for a specific application, without affecting traffic for other applications on that host. You can also use virtual clusters to limit each application, Web site, or virtual IP address to a specific subset of computers within your primary cluster. See the diagram below for an example of this use of virtual clusters.




This diagram depicts a four host Network Load Balancing cluster. Through the use of virtual clusters and IP address-specific load weight, Network Load Balancing directs network traffic as follows:

- Users accessing Web Site A (IP address *nnn.nnn.nnn*.1) are directed to any of the four hosts
- Users accessing Web site B (IP address *nnn.nnn.nnn*.2), which is a virtual cluster, are directed to hosts 1 and 2
- Users accessing Web site C (IP address *nnn.nnn.nnn*.3), which is a virtual cluster, are directed to hosts 3 and 4

For more information on how this is configured, see Virtual clusters and Create a new port rule.

# Planning Your Network Load Balancing Cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning your Network Load Balancing cluster**

This section covers the issues you must consider as you plan how to set up Network Load Balancing in your organization.

- Identifying Your Risks
- Choosing a Network Load Balancing Model
- Capacity Planning
- Optimizing Network Load Balancing Clusters

# Identifying Your Risks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Identifying your risks**

When you identify network risks, you identify the possible failures that can interrupt access to network resources. A single point of failure is any component in your environment that would block data or applications if it failed. Single points of failure can be hardware, software, or external dependencies, such as power supplied by a utility company and dedicated wide area network (WAN) lines.

In general, you provide maximum reliability when you:

- Minimize the number of single points of failure in your environment.
- Provide mechanisms that maintain service when a failure occurs.

In the case of Network Load Balancing, you also provide maximum reliability when you:

- Load balance only applications appropriate to Network Load Balancing.
- Make sure that the application servers are properly configured for the applications they are running.

This section covers the following areas:

- Addressing your risks
- Determining which applications to use with Network Load Balancing

For more information about monitoring recommendations for high-availability systems, see the "System Performance and Troubleshooting Guide" at the Microsoft Windows Resource Kits Web site. For more information about maximizing the availability of a Network Load Balancing cluster, see the "Server Management Guide" at the Microsoft Windows Resource Kits Web site.

# Addressing your risks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Addressing your risks**

A principal goal of Network Load Balancing is to provide increased reliability. A cluster of two or more computers ensures that if one computer fails, another computer is available to continue processing client requests.

However, Network Load Balancing is not designed to protect all aspects of your workflow in all circumstances. For example, Network Load Balancing is not an alternative to backing up data. Network Load Balancing protects only the availability of data, not the data itself. Also, it does not guard against a power outage that would disable the entire cluster.

Windows Server 2003, Standard Edition ; Windows Server 2003, Enterprise Edition ; and Windows Server 2003, Datacenter Edition , have built-in features that protect certain computer and network processes during failure. These features include disk mirroring (RAID Level 1) and disk striping [AMC25A] with parity (RAID Level 5). When planning your Network Load Balancing environment, look for places where these features can help you in ways that Network Load Balancing cannot.

The following table lists common points of failure in a Windows Server 2003, Standard Edition ; Windows Server 2003, Enterprise Edition ; or Windows Server 2003, Datacenter Edition , environment, and describes whether the point of failure can be protected, either by Network Load Balancing or by other means.

These features include RAID-1, disk mirroring, RAID-5, and disk striping with parity.

| **Failure point** | **Network Load Balancing solution** | **Other solutions** |
|---|---|---|
| Network hub | n/a | Redundant networks |
| Power outage | n/a | • Uninterruptable power supply (UPS)<br>• Generator |
| server connection | Failover | n/a |
| Disk | Failover | Hardware or software RAID, to ensure against the loss of data on a specific computer and to provide uninterrupted service |
| Other server hardware, such as CPU or memory | Failover | Spare components, such as motherboards and SCSI controllers (any spare components should exactly match the original components, including network and SCSI components) |
| server software, such as the operating system or specific applications | n/a | server clusters (only computers running Windows Server 2003, Enterprise Edition , or Windows Server 2003, Datacenter Edition ) |
| IP-level failures, such as IIS service interruptions | Failover | n/a |
| WAN links, such as routers and dedicated lines | n/a | Redundant links, to provide secondary access to remote connections |
| Dial-up connection | n/a | Multiple modems |
| Application or Service failure | n/a | Microsoft Application Center, available as a separate product, contains tools to monitor specific applications and services. |

# Determining which applications to use with Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Determining which applications to use with Network Load Balancing**

Many applications work with Network Load Balancing. This section offers guidelines for determining which applications might be suitable.

In general, Network Load Balancing can load balance any application or service that uses TCP/IP as its network protocol and is associated with a specific Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port. For more information, see Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

Some examples are:

| Protocol | Examples |
|---|---|
| HTTP over TCP/IP | Web servers, such as Microsoft Internet Information Services: Port 80. |
| HTTPS over TCP/IP | HTTP over SSL (Secure Sockets Layer), for encrypting Web traffic: Port 443. |
| FTP over TCP/IP | FTP (File Transfer Protocol): Port 20, port 21, and ports 1024-65535. |
| TFTP over TCP/IP | TFTP servers (used by such applications as BOOTP): port 69. |
| SMTP over TCP/IP | SMTP (Simple Mail Transfer Protocol), used by such applications as Microsoft Exchange: Port 25. |
| RDP over TCP/IP | RDP (Remote Datagram Protocol), used by Terminal Services: Port 3389. |

In addition, Network Load Balancing is in common use with:

- Web servers using round-robin DNS.
- Virtual private network (VPN) servers.
- Streaming media servers.
- Print Services.
- Microsoft Internet Security and Acceleration Server (ISA).

# Choosing a Network Load Balancing Model

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing a Network Load Balancing model**
Network Load Balancing can be configured using one of four different models. This section describes the models, and sets forth the advantages and disadvantages of each.

- Single network adapter in unicast mode
- Multiple network adapters in unicast mode
- Single network adapter in multicast mode
- Multiple network adapters in multicast mode

# Single network adapter in unicast mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Single network adapter in unicast mode**
This model is suitable for a cluster in which ordinary network communication among cluster hosts is not required and in which there is limited dedicated traffic from outside the cluster subnet to specific cluster hosts.

**Advantages**

- Only one network adapter is required. It is not necessary to install a second adapter.
- This is the most straightforward configuration because unicast mode is the default.
- This model works with all routers.

**Disadvantages**

- Ordinary network communication among cluster hosts is not possible.
- Network Load Balancing itself does not affect network performance and a second network adapter is not a requirement. However, under certain conditions, a second adapter can improve overall network performance. For more information, see Multiple network adapters.

Regardless of the number of network adapters you use and how the adapters are configured, the cluster heartbeat traffic always originates from the network load balancing network adapter.

**Important**

- Network Load Balancing does not support a mixed unicast/multicast environment within a single cluster. Within each cluster, all network adapters in that cluster must be either multicast or unicast; otherwise, the cluster will not function properly.

  There is no restriction on the number of network adapters that can be bound to network load balancing on each computer: Different hosts can have a different number of adapters, but you can never have more than one adapter on a computer be part of the same cluster.

  Different adapters on each computer can be part of different clusters, and they can all reside on the same subnet. You can also host multiple clusters on the same network adapter. For more information, see Understanding virtual clusters.

- Before you install Network Load Balancing, it is vital that you review Checklist: Enabling and configuring Network Load Balancing, including the sections on cluster parameters, host parameters, and port rules. It is highly recommended that you read through these sections before configuring Network Load Balancing.

For more information, see also Single network adapter and Single network adapter limitations.

# Multiple network adapters in unicast mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multiple network adapters in unicast mode**
This model is suitable for a cluster in which ordinary network communication among cluster hosts is necessary or desirable. It is also appropriate when you want to separate the traffic used to manage the cluster from the traffic occurring between the cluster and client computers.

**Advantages**

- Ordinary network communication among cluster hosts is permitted.
- This model works with all routers.

**Disadvantages**

- This model requires a second network adapter.

Regardless of the number of network adapters you use and how the adapters are configured, the cluster heartbeat traffic always originates from the network load balancing network adapter.

**Important**

- Network Load Balancing does not support a mixed unicast/multicast environment within a single cluster. Within each cluster, all network adapters in that cluster must be either multicast or unicast; otherwise, the cluster will not function properly.

  There is no restriction on the number of network adapters that can be bound to network load balancing on each computer: Different hosts can have a different number of adapters, but you can never have more than one adapter on a computer be part of the same cluster.

  Different adapters on each computer can be part of different clusters, and they can all reside on the same subnet. You can also host multiple clusters on the same network adapter. For more information, see Understanding virtual clusters.

- Before you install Network Load Balancing, it is vital that you review Checklist: Enabling and configuring Network Load Balancing, including the sections on cluster parameters, host parameters, and port rules. It is highly recommended that you read through these sections before configuring Network Load Balancing.

For more information, see Multiple network adapters.

# Single network adapter in multicast mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Single network adapter in multicast mode**
This model is suitable for a cluster in which ordinary network communication among cluster hosts is necessary or desirable but in which there is limited dedicated traffic from outside the cluster subnet to specific cluster hosts.

**Advantages**

- Only one network adapter is required. It is not necessary to install a second adapter.
- This model permits ordinary network communication among cluster hosts.
- You can enable Internet Group Management Protocol (IGMP) support on the cluster hosts to control switch flooding when operating in multicast mode. For more information, see Limit switch port flooding.

**Disadvantages**

- Network Load Balancing itself does not affect network performance and a second network adapter is not a requirement. However, under certain conditions, a second adapter can improve overall network performance.
- Some routers might not support the use of a multicast media access control (MAC)address. This only affects the Network Load Balancing/MAC address (not all MAC addresses) and only when dynamic ARP replies are sent by the cluster to the router, not all MAC addresses.

In this model, the computer can also handle traffic from inside the subnet if the IP datagram does not carry the same MAC address as on the cluster adapter. For more information, see Address Resolution Protocol (ARP).

Regardless of the number of network adapters you use and how the adapters are configured, the cluster heartbeat traffic always originates from the network load balancing network adapter.

**Important**

- Network Load Balancing does not support a mixed unicast/multicast environment within a single cluster. Within each cluster, all network adapters in that cluster must be either multicast or unicast; otherwise, the cluster will not function properly.

  There is no restriction on the number of network adapters that can be bound to network load balancing on each computer: Different hosts can have a different number of adapters, but you can never have more than one adapter on a computer be part of the same cluster.

  Different adapters on each computer can be part of different clusters, and they can all reside on the same subnet. You can also host multiple clusters on the same network adapter. For more information, see Understanding virtual clusters.

- Before you install Network Load Balancing, it is vital that you review Checklist: Enabling and configuring Network Load Balancing, including the sections on cluster parameters, host parameters, and port rules. It is highly recommended that you read through these sections before configuring Network Load Balancing.

For more information, see Single network adapter, Single network adapter limitations, and Using a router in Network Load Balancing system requirements.

# Multiple network adapters in multicast mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multiple network adapters in multicast mode**
This model is suitable for a cluster in which ordinary network communication among cluster hosts is necessary and in which there is heavy dedicated traffic from outside the cluster subnet to specific cluster hosts.

**Advantages**

- Because there are at least two network adapters, overall network performance is typically enhanced.
- Ordinary network communication among cluster hosts is permitted.

**Disadvantages**

- This model requires a second network adapter.
- Some routers might not support the use of a multicast media access control (MAC)address. This only affects the Network Load Balancing/MAC address (not all MAC addresses) and only when dynamic ARP replies are sent by the cluster to the router. not all MAC addresses.

For more information, see Multiple network adapters and Using a router in Network Load Balancing system requirements.

**Important**

- Network Load Balancing does not support a mixed unicast/multicast environment within a single cluster. Within each cluster, all network adapters in that cluster must be either multicast or unicast; otherwise, the cluster will not function properly.

  There is no restriction on the number of network adapters that can be bound to network load balancing on each computer: Different hosts can have a different number of adapters, but you can never have more than one adapter on a computer be part of the same cluster.

  Different adapters on each computer can be part of different clusters, and they can all reside on the same subnet. You can also host multiple clusters on the same network adapter. For more information, see Understanding virtual clusters.

- Before you install Network Load Balancing, it is vital that you review Checklist: Enabling and configuring Network Load Balancing, including the sections on cluster parameters, host parameters, and port rules. It is highly recommended that you read through these sections before configuring Network Load Balancing.

# Capacity Planning

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capacity planning**
This section helps you determine the number of servers you need and how they should be configured. Your planning should be based on the Network Load Balancing model you choose to implement. For more information, see Choosing a Network Load Balancing Model.

- Sizing Network Load Balancing clusters
- Determining server capacity requirements

# Sizing Network Load Balancing clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sizing Network Load Balancing clusters**

Cluster size, defined as the number of cluster hosts participating in the cluster (up to 32), is based on the number of computers required to meet the anticipated client load for a given application.

If, for example, you determine that you need six computers running Microsoft Internet Information Services to meet the anticipated client demand for Web services, then Network Load Balancing will run on all six computers, and your cluster will consist of six cluster hosts.

As a general rule, add servers until the servers can comfortably meet their demands without saturating the network. The exact number depends on the nature of the application. When the cluster subnet approaches saturation, add clusters using round-robin DNS. The maximum cluster size is determined by network capacity. The minimum cluster size is determined by server capacity. For example, if two extremely powerful servers can take the place of six less powerful servers, your entire cluster will consist of only two cluster hosts. In addition, you should always make sure that there is enough extra server capacity so that if one server fails, the server next in line will not fail, which would create a cascading series of server failures.

### Notes

- If you deploy two or more clusters, consider putting them on individual switches so that incoming cluster traffic is handled separately. For more information on using Network Load Balancing hosts on switches, see article Q238219, "How Network Load Balancing Hosts Converge When Connected to a Layer 2 Switch," and article Q193602, "Configuration Options for WLBS Hosts Connected to a Layer 2 Switch," in the Microsoft Knowledge Base.

- Before load balancing an application in a Network Load Balancing cluster, review the application license, or check with the application vendor. Each application vendor sets its own licensing policies for applications running on clusters.

# Determining server capacity requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Determining server capacity requirements**

After you determine the size of your cluster, you are ready to determine how individual cluster hosts should be configured. In general, this is based on the applications being load balanced and the anticipated client demand on those applications.

Some server applications (such as file and print servers) are extremely disk intensive and require very large disk capacities and fast I/O. Some applications (such as Microsoft Exchange) are CPU intensive and require very fast processers to work at maximum efficiency. Other applications (such as Microsoft SQL Server) require large amounts of memory. You should consult the documentation for each application that you will be running to determine how the servers in your cluster should be configured.

You should also bear in mind that while it might be possible to substitute two or three very powerful servers for a larger number of less powerful computers, deploying the larger number might actually be a better approach. With more servers, the client load can be more widely distributed, so that if one or more servers fail, the impact on clients is reduced.

# Optimizing Network Load Balancing Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Optimizing Network Load Balancing clusters**

Because Windows uses an adaptable architecture, it is largely self-tuning. It is able to allocate resources as needed to meet changing usage requirements. However, you can make certain hardware and configuration choices to improve cluster performance. These choices are explained in the following sections:

- Focusing on the right hardware resources
- Optimizing the paging file size and location
- Optimizing network performance
- Using additional TCP/IP services

# Focusing on the right hardware resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Focusing on the right hardware resources**

The goal in tuning Network Load Balancing and the applications it load balances is to determine which hardware resource will experience the greatest demand, and then to adjust the configuration to relieve that demand and maximize total throughput.

For example, if the primary role of the cluster is to provide high availability of file and print services, high disk use will be incurred due to the large number of files being opened and closed. File and print services also cause a heavy load on network adapters because of the large amount of data that is being transferred. It is important to make sure that your network adapter and cluster subnet can handle the load. In this scenario, RAM typically does not carry a heavy load (although memory usage can be heavy if a large amount of RAM is allocated to a file system cache). Processor use is also typically low in this environment. In such cases, memory and processor use usually do not need the optimizing that other components need.

In contrast, a server-application environment (such as e-commerce sites running COM+ applications) is much more processor and RAM intensive than a typical file or print server environment because much more processing is taking place at the server. In these cases, it is best to use high-end multiprocessor computers. The disk and network loads tend to be used less because the amount of data being sent is smaller. Network Load Balancing uses few system resources either for host-to-host communication or for the operation of the cluster itself.

# Optimizing the paging file size and location

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Optimizing the paging file size and location**

The location and size of the paging file can greatly affect performance. Putting the paging file on a separate, fast, low-use drive can boost performance. The size setting of the paging file is also critical to performance. A paging file that must constantly expand and shrink requires additional processor and I/O overhead. Usually, you should set the size of your paging file at two or two-and-a-half times the amount of installed physical memory if applications such as Microsoft SQL Server or Microsoft Exchange Server are present. If you set the same value for the minimum and maximum size of your paging file, the amount of disk defragmentation is reduced.

# Optimizing network performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Optimizing network performance**

If the cluster hosts are directly connected to a hub in order to receive client requests, incoming client traffic is automatically sent to all ports.

To limit switch port flooding (for example, when computers outside the cluster also share the hub), see Limit switch port flooding. The steps outlined make all incoming client traffic flow through one switch port and thereby conserve switch bandwidth. In this configuration, you are advised to consider adding a second, dedicated network adapter for each host connected to an individual switch port. The use of two network adapters per host on the cluster subnet helps to direct network traffic through the cluster hosts. Incoming client traffic flows through the hub for simultaneous delivery to all hosts, while outgoing traffic flows directly to the switch ports. Be sure to enter a gateway IP address in the **Internet Protocol (TCP/IP) Properties** dialog box in **Network Connections** for the dedicated network adapter, to cause outgoing network traffic to flow through that network adapter. At the same time, be sure to remove the default gateway for the cluster adapter, if one has been configured. It is essential that the default gateway be configured for either the cluster adapter or the dedicated adapter but not both.

# Using additional TCP/IP services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using additional TCP/IP services**

WINS, DHCP, and DNS Services use additional server resources, which must be considered before implementing any of these services on a server also operating as a cluster host. It is highly recommended that you do not run WINS, DHCP, or DNS Services on cluster hosts. In particular, WINS can require substantial overhead during peak logon times in large networks. In such cases, make sure that your server has the necessary capacity to perform adequately.

# Using Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Network Load Balancing**
This section covers:

- Network Load Balancing Manager
- Single network adapter
- Single network adapter limitations
- Multiple network adapters
- Network Load Balancing parameters
- Virtual clusters
- Upgrading Network Load Balancing
- Adjusting multicast parameters
- Handling LAN failures
- Error detection and handling
- Event logging
- Network Load Balancing and stateful connections
- Managing Network Load Balancing from the command line

# Network Load Balancing Manager

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Load Balancing Manager**
Network Load Balancing (NLB) Manager allows you to create and manage Network Load Balancing clusters from a single computer. By centralizing Network Load Balancing administration tasks, Network Load Balancing Manager helps eliminate many common configuration errors.

Using Network Load Balancing Manager, you can configure and manage your clusters and all cluster hosts from a single computer, and you can also replicate the cluster configuration to other hosts.

You can perform the following procedures with Network Load Balancing Manager:

- Create new clusters
- Add hosts to an existing cluster and, by doing so, automatically enable Network Load Balancing on that host and and automatically configure cluster parameters on that host
- Remove hosts from a cluster and, by doing so, automatically disable Network Load Balancing on that host
- Configure all host and cluster properties
- Replicate the cluster configuration to other hosts
- Diagnose problems with existing Network Load Balancing clusters.

As part of the Windows Server 2003 Administration Tools Pack, Network Load Balancing Manager can also be installed and used on Windows XP Professional. You can use Network Load Balancing Manager on Windows XP Professional only to manage Network Load Balancing clusters on remote computers running the Windows Server 2003 family of products. You can not install the Network Load Balancing service itself on Windows XP Professional.

Network Load Balancing Manager uses Windows Management Instrumentation (WMI) which in turn has a dependency on remote procedure call (RPC) and Distributed Component Object Model (DCOM). To understand how to enable DCOM across a firewall, refer to Using Distributed COM with Firewalls .

When using Network Load Balancing Manager, you must be a member of the Administrators group on the host that you are configuring. However, you do not have to me a member of the Administrators group on the computer from which you are running Network Load Balancing Manager.

For more information on Network Load Balancing, see Network Load Balancing Clusters, and Network Load Balancing Overview.

**Important**

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

**Network Load Balancing Manager compared to other deployment and management tools**
Network Load Balancing Manager is similar to other application management tools, such as Microsoft Application Center 2000, but Network Load Balancing Manager does not completely replace these other tools. Although you can use either Network Load Balancing Manager or a similar tool to manage Network Load Balancing clusters, you can not use Network Load Balancing Manager to perform additional tasks that are often available through these other tools.

Network Load Balancing Manager is used to centrally create and manage the configuration of Network Load Balancing clusters only. It is not used to administer and deploy applications. In contrast, other management tools are designed to create, deploy, and manage large scale Web applications. These tools typically include advanced monitoring, logging, and scripting capabilities which are not included with Network Load Balancing Manager.

For example:

- Network Load Balancing Manager provides limited monitoring of the status of the Network Load Balancing hosts, but does not provide any alerts or notifications. Other management tools often provide real-time monitoring, notification and alert capabilities allowing you to view the health, performance and availability of clusters and other applications.

- Scripting is not supported in Network Load Balancing Manager but is often used by other management tools for monitoring and automating various administrative tasks.

- Network Load Balancing Manager can log general Network Load Balancing cluster information to a text-based log file, while other tools often can provide detailed logging on a variety of applications.

Network Load Balancing Manager is a tool targeted at very specific tasks related to Network Load Balancing clusters and you should not attempt to use Network Load Balancing Manager for other administrative purposes.

# Single network adapter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Single network adapter**

Network Load Balancing can operate in two modes: unicast and multicast.

In unicast mode, the cluster's media access control (MAC) address is assigned to the computer's network adapter, and the network adapter's built-in MAC address is not used.

In multicast mode, the cluster's MAC address is assigned to the computer's network adapter, but the network adapter's built-in address is retained so that both addresses are used, the first for client-to-cluster traffic and the second for network traffic specific to the computer.

The following describes the principal characteristics of each mode when a single network adapter is used.

**Single network adapter in unicast mode**

- The adapter's own, unique MAC address is disabled.

  The cluster MAC address (which Network Load Balancing automatically generates) replaces this address.

- The adapter becomes, in effect, the cluster adapter.

- Both the dedicated IP address and the cluster IP address resolve to the cluster MAC address.

- Because all cluster hosts share the same MAC address, and because the original MAC address is not used, ordinary network communication between this host and the other cluster hosts is not possible. However, the computer can still handle traffic originating from outside the subnet in which the cluster is located, and from inside the subnet if the IP datagram does not carry the same MAC address as on the cluster adapter. For more information, see Address Resolution Protocol (ARP).

**Single network adapter in multicast mode**

- Network Load Balancing automatically generates a cluster MAC address for the network adapter.

- The adapter's own, unique MAC address is retained.

- The cluster IP address resolves to the cluster MAC address.

- The dedicated IP address resolves to the original MAC address.

**Important**

- Network Load Balancing does not support a mixed unicast/multicast environment within a single cluster. Within each cluster, all network adapters in that cluster must be either multicast or unicast; otherwise, the cluster will not function properly.

  However, there is no restriction on the number of network adapters: Different hosts can have a different number of adapters.

- To start Network Load Balancing Manager, open a Command prompt. From the command window, type **NLBMgr**, and then press ENTER.

Unicast support is enabled by default. To enable multicast support, see Enable multicast support.

If multicast support is enabled, a single network adapter can handle both client-to-cluster traffic and dedicated traffic without difficulty. However, the use of a second network adapter might improve performance by handling dedicated traffic concurrently with the original network adapter.

Multicast support, with or without multiple network adapters, has several advantages, such as eliminating the inherent limitations of a single network adapter, as described in Single network adapter limitations. Network Load Balancing handles the resolution of the cluster's primary IP address to its associated multicast cluster MAC address within the ARP on your router. In rare cases, the use of a multicast MAC address might not be supported by the ARP implementation on a given router. If this problem arises, the cluster will not be accessible from outside the local subnet. To correct this problem, you must create a static ARP entry within the router. Refer to the documentation for your router to determine how to create a static ARP entry. For more information, see Network Load Balancing system requirements.

If you disable multicast support (causing the host to revert to unicast mode), Network Load Balancing automatically instructs the driver belonging to the cluster adapter to override the adapter's unique, built-in network address and to change its MAC address to the cluster's MAC address. This is the address used on all cluster hosts. You do not need to manually configure the network adapter to recognize this address. (Note that some network adapters do not support changing their MAC addresses. If you experience this problem, you must install a network adapter that does.)

If you are using multiple network adapters, see Multiple network adapters.

# Single network adapter limitations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Single network adapter limitations**

The limitations described in this topic apply only when multicast support is not enabled.

If you do not enable multicast support, you should consider installing a dedicated network adapter to achieve optimal performance and the full range of networking functionality. Although a successful cluster can be implemented with only one network adapter and without activating multicast support, this approach has two limitations:

- Ordinary network communication among cluster hosts is not possible.
- Network traffic intended for any individual computer within the cluster generates additional networking overhead for all computers in the cluster.

Further, if you are working from a computer that has a single network adapter that is bound to Network Load Balancing in unicast mode, you can not use Network Load Balancing Manager on this computer to configure and manage other hosts.

These limitations might not be important for many installations. If the cluster hosts are accessed only from outside the cluster (because host-to-host communication cannot take place) and noncluster communications are infrequent (for example, short nightly updates of the Web server content to individual computers in the cluster), the use of a single network adapter should be sufficient. Also, it is always possible to test and demonstrate Network Load Balancing using a single network adapter prior to installing additional network adapters for operational purposes.

# Multiple network adapters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multiple network adapters**
Network Load balancing can use multiple network adapters in two different and distinct ways. The first use is to create multiple Network Load Balancing clusters; achieved by binding Network Load Balancing to each of the multiple adapters.

The second use for multiple network adapters is to use one or more adapters for handling cluster traffic, but also configure at least one adapter to be used for server management. The management adapter does not handle cluster traffic. Further information on this this second use of multiple network adapters is described below.

**Multiple network adapters in unicast mode**
In unicast mode, the cluster's media access control (MAC) address is assigned to the network adapter that is functioning as the cluster adapter, and the cluster adapter's built-in MAC address is not used. Because of this, host to host communication (other than Network Load Balancing heartbeat-related traffic) is not possible on the cluster adapter.

The following characteristics apply:

- Adapters used as cluster adapters
    - The cluster IP address resolves to the adapter's cluster MAC address (which Network Load Balancing automatically generates).
    - If the cluster adapter's dedicated IP address is used, this IP address also resolves to the cluster MAC address.
    - The cluster adapter handles client-to-cluster traffic using the virtual IP address. However host-to-host traffic within the same cluster is not possible on this adapter because the MAC address for each host in a particular cluster is the same. For more information, see Address Resolution Protocol (ARP).
- Adapter used as the management adapter
    - Network Load Balancing in not bound to this adapter.
    - This adapter's IP address resolves to the adapter's built-in MAC address.
    - This adapter handles network traffic specific to the computer, including traffic from both inside and outside the subnet in which the computer is located.

**Multiple network adapters in multicast mode**
In multicast mode, the cluster's MAC address is assigned to network adapters that are functioning as cluster adapters, but the cluster adapters' built-in address is retained so that both addresses are used, the first for client-to-cluster traffic and the second for network traffic specific to the computer.

This following characteristics apply:

- Adapters used as cluster adapters
    - The cluster IP address resolves to the adapter's cluster MAC address (which Network Load Balancing automatically generates).
    - If the cluster adapter's dedicated IP address is used, this IP address resolves to the adapter's original, built-in MAC address.
    - In multicast mode, the cluster adapters can handle both client-to-cluster traffic and traffic specific to the computer, including all traffic from both inside and outside the subnet in which the computer is located.

- Adapter used as the management adapter
    - Network Load Balancing in not bound to this adapter.
    - This adapter's IP address resolves to the adapter's built-in MAC address.
    - This adapter can be used to handle the network traffic specific to the computer, including traffic from both inside and outside the subnet in which the computer is located.

**Important**

- Network Load Balancing does not support a mixed unicast/multicast environment within a single cluster. Within each cluster, all network adapters in that cluster must be either multicast or unicast; otherwise, the cluster will not function properly.

  However, there is no restriction on the number of network adapters: Different hosts can have a different number of adapters.

Network Load Balancing handles the resolution of the cluster's IP address to its associated multicast cluster media access control address within the Address Resolution Protocol (ARP) on your router. In rare cases, the use of a multicast MAC address might not be supported by the ARP implementation on a given router. If this problem arises, the cluster will not be accessible from outside the local subnet. To correct this problem, you must create a static ARP entry within the router. Refer to the documentation for your router to determine how to create a static ARP entry. For more information, see Using a router in Network Load Balancing system requirements.

If you disable multicast support (causing the host to revert to unicast mode), Network Load Balancing automatically instructs the driver belonging to the cluster adapter to override the adapter's unique, built-in network address and to change its media access control address to the cluster's media access control address. This is the address used on all cluster hosts. You do not need to manually configure the network adapter to recognize this address. (Note that some network adapters do not support changing their media access control addresses. If you experience this problem, you must install a network adapter that does.)

**Notes**

- Multicast support is not enabled by default. To enable it, see Enable multicast support. However, even if you do not enable multicast support, you are advised to consider using at least two network adapters (with one network adapter dedicated to handling client-to-cluster traffic) in order to achieve optimum performance and the full range of networking functionality. For more information, see Optimizing network performance.
- If you are using single network adapter, see Single network adapter and Single network adapter limitations.

# Network Load Balancing parameters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Load Balancing parameters**

This topic describes the parameters that are specific to each host in a Network Load Balancing cluster and the parameters for a Network Load Balancing cluster as a whole. It also describes the parameters that control how the cluster functions. These are called port rules. Values entered for the entire cluster must be the same for all host computers in the cluster. You enter parameters in the **Network Load Balancing Properties** dialog box and they are recorded in the registry.

The parameters are contained in the following tabs on the **Network Load Balancing Properties** dialog box:

| Tab | These parameters |
|---|---|
| Cluster parameters | Apply to the entire cluster. |
| Host parameters | Apply to a specific host. |
| Port rules | Control how the cluster functions. |

**Cluster parameters**

| Parameter | Description |
|---|---|
| **IP address** on the **Cluster Parameters** tab | This parameter specifies the cluster's primary IP address in standard Internet dotted notation (for example, w.x.y.z). The address is a virtual IP address and must be set identically for all hosts in the cluster. This IP address is used to address the cluster as a whole, and it should be the IP address that maps to the full Internet name that you specify for the cluster. This should be the primary IP address and subnet mask for the cluster. If you want to add multiple IP addresses to the cluster, you enter the additional IP addresses in the TCP/IP properties dialog box or in the Cluster IP Addresses dialog box in Network Load Balancing Manager.<br><br>**Important**<br><br>• Typically, both the dedicated IP address and the cluster IP address, entered during setup in the **Network Load Balancing Properties** dialog box, must also be entered in the **Internet Protocol (TCP/IP) Properties** dialog box. Make sure that the addresses are the same in both places. However, if you are configuring a virtual private network (VPN) load balancing cluster, you should not configure the dedicated IP address. On a VPN, only the cluster IP address should be present on each of the cluster hosts. |
| **Subnet mask** on the **Cluster Parameters** tab | This parameter denotes the subnet mask for the IP address specified. The mask is entered in standard Internet dotted notation (for example, 255.255.255.0). |
| **Full Internet name** on the **Cluster Parameters** tab | This parameter specifies a full Internet name for the Network Load Balancing cluster (for example, cluster.microsoft.com). This name is used for the cluster as a whole and should be the same for all hosts in the cluster. If you alias several names for the cluster, the primary (main) name should be entered here. In any case, this name should be resolvable to the cluster's primary IP address through your DNS server or Hosts file. |

| | |
|---|---|
| **Network address** on the **Cluster Parameters** tab | This parameter specifies the network address (media access control [MAC] address) for the network adapter to be used for handling client-to-cluster traffic.<br><br>If multicast support is disabled (causing the host to revert to unicast mode), Network Load Balancing automatically instructs the driver belonging to the cluster adapter to override the adapter's unique, built-in network address and to change its MAC address to the cluster's MAC address. This is the address used on all cluster hosts. You do not need to manually configure the network adapter to recognize this address.<br><br>**Note**<br><br>• Some network adapters might not allow the built-in network address to be modified. If you experience this problem, you must obtain and install a different network adapter that supports this functionality.<br><br>If you have other Network Load Balancing clusters on one local subnet, each cluster needs to use a different network address. When you select a different primary IP address for each cluster, Network Load Balancing automatically ensures that the clusters use unique network addresses. |
| **Cluster Operation Mode** on the **Cluster Parameters** tab | These parameters specify whether or not a multicast MAC address should be used for cluster operations. If multicast is enabled, Network Load Balancing converts the cluster MAC address belonging to the cluster adapter into a multicast address. It also ensures that the cluster's primary IP address resolves to this multicast address as part of the ARP protocol. At the same time, the adapter can now use its original, built-in MAC address that, in unicast mode, was disabled.<br><br>**Note**<br><br>• If Network Load Balancing clients are accessing a cluster through a router when the cluster has been configured to operate in multicast mode, be sure that the router meets the following requirements:<br><br>  • Accepts an ARP reply that has one MAC address in the payload of the ARP structure but appears to arrive from a station with another MAC address, as identified by the Ethernet header.<br><br>  • In multicast mode, accepts an ARP reply that has a multicast MAC address in the payload of the ARP structure.<br><br>  This allows the router to map the cluster's primary IP address and other multihomed addresses to the corresponding MAC address. If your router does not meet these requirements, you can also create a static ARP entry in the router. Cisco routers require a static ARP entry because they do not support the resolution of unicast IP addresses to multicast MAC addresses.<br><br>In multicast mode, the IGMP multicast check box enables Internet Group Management Protocol (IGMP) support for limiting switch flooding by limiting traffic to "Network Load Balancing ports" only. That is, enabling IGMP support ensures that traffic intended for a Network Load Balancing cluster passes through only those ports serving the cluster hosts and not all switch ports.<br><br>**Notes**<br><br>• You must first enable multicast support before enabling IGMP support.<br><br>• You must enable IGMP support on the network adapter also.<br><br>• IGMP support is not enabled by default. To enable it, see Enable Internet Group Management Protocol (IGMP) support.<br><br>• To avoid switch flooding in multicast mode, see Limit switch port flooding.<br><br>If you do not select multicast support (causing the host to revert to unicast mode), Network Load Balancing automatically instructs the driver belonging to the cluster adapter to override the adapter's unique, built-in network address and to change its MAC address to the cluster's MAC address. This is the address used on all cluster hosts. You do not need to manually configure the network adapter to recognize this address. (Note that some network adapters do not support changing their MAC addresses. If you experience this problem, you must |

| | |
|---|---|
| | install a network adapter that does.)<br><br>**Important**<br><br>• Network Load Balancing does not support a mixed unicast/multicast environment within a single cluster. Within each cluster, all network adapters in that cluster must be either multicast or unicast; otherwise, the cluster will not function properly.<br><br>However, there is no restriction on the number of network adapters: Different hosts can have a different number of adapters.<br><br>**Note**<br><br>• Multicast support is not enabled by default. To enable it, see Enable multicast support. However, if you do not enable multicast support, you are advised to consider using at least two network adapters (with one network adapter dedicated to handling client-to-cluster traffic) in order to achieve optimum performance and the full range of networking functionality. For more information, see Multiple network adapters and Optimizing network performance. |
| **Remote Control** on the **Cluster Parameters** tab | This parameter specifies whether remote control operations are enabled. When enabled, remote, networked computers can control cluster operations by using the NLB.exe cluster-control application. Remote control is disabled by default.<br>Once remote control has been enabled, remote access can be restricted by specifying a remote control password.<br><br>Network Load Balancing remote control commands will not work correctly if they are sent from a computer that has Internet Protocol security (IPSec) configured such that the remote control traffic is encrypted by IPSec. For more information, see Internet Protocol Security (IPSec).<br><br>**Caution**<br><br>• The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. It is highly recommended that you do not enable remote control and instead use Network Load Balancing Manager or other remote management tools such as Windows Management Instrumentation (WMI).<br><br>  <br><br>If you choose to enable remote control, it is vital that you restrict access by specifying a strong remote control password. It is also imperative that you use a firewall to protect the Network Load Balancing UDP control ports (the ports that receive remote control commands) in order to shield them from outside intrusion. By default, these are ports 1717 and 2504 at the cluster IP address. Use remote control only from a secure, trusted computer within your firewall. For more information about strong passwords, see Strong passwords. |

| | |
|---|---|
| **Remote password** on the **Cluster Parameters** tab | This parameter specifies a password used for restricting access to the cluster from remote, networked computers that use the NLB.exe cluster-control application. The password consists of a string of alphanumeric characters. The password must be entered as the value of this parameter and also a second time for confirmation in the **Confirm password** parameter.<br><br>After a password has been accepted, subsequent NLB.exe remote control operations must use the **/PASSW** command-line parameter and submit the password. Clearing both fields disables use of the remote control password.<br><br>If you choose to enable remote control it is vital that you restrict access by specifying a strong remote control password. It is also imperative that you use a firewall to protect the Network Load Balancing UDP control ports (the ports receiving remote control commands) in order to shield them from outside intrusion. By default, these are ports 1717 and 2504 at the cluster IP address. For more information about strong passwords, see Strong passwords.<br><br>The remote control password is not used to restrict control operations from a cluster host. |
| **Confirm password** on the **Cluster Parameters** tab | This parameter specifies the password entered in the **Remote password** field. It is used to confirm proper entry of this password. Clearing both fields disables use of a remote control password. |

**Host parameters**

| Parameter | Description |
|---|---|
| **Interface** on the **Host Parameters** tab | This parameter only appears when using Network Load Balancing Manager and is configured when you add the host to the cluster. The parameter specifies the host's network adapter that will use network load balancing in context of the current cluster. |
| **Priority (Unique host ID)** on the **Host Parameters** tab | This parameter specifies a unique ID for each host.<br><br>The host with the lowest numerical priority among the current members of the cluster handles all of the cluster's network traffic that is not covered by a port rule. You can override these priorities or provide load balancing for specific ranges of ports by specifying rules in the **Port rules** tab of the **Network Load Balancing Properties** dialog box.<br><br>If a new host joins the cluster and its priority conflicts with another host in the cluster, the host is not accepted as part of the cluster. The rest of the cluster will continue to handle the traffic as before. A message describing the problem is written to the Windows event log. For more information on error logging and cluster operations, see How Network Load Balancing works, Event logging, and Error detection and handling. |

| | |
|---|---|
| **IP address** on the **Host Parameters** tab | This parameter specifies this host's unique IP address used for network traffic not associated with the cluster (for example, Telnet access to a specific host within the cluster). It should be entered in standard Internet dotted notation (for example, w.x.y.z). This IP address is used to individually address each host in the cluster and hence should be unique for each host. The dedicated IP address should always be entered first in TCP/IP properties.<br><br>Network Load Balancing references the dedicated IP address only when a single network adapter is used to handle both client-to-cluster traffic and other network traffic that must go specifically to the dedicated IP address. Network Load Balancing ensures that all traffic to the dedicated IP address is unaffected by the Network Load Balancing current configuration, including:<br><br>• When this host is running as part of the cluster<br><br>• When Network Load Balancing is disabled due to parameter errors in the registry<br><br>**Important**<br><br>• Typically, both the dedicated IP address and the cluster IP address, entered during setup in the **Network Load Balancing Properties** dialog box, must also be entered in the **Internet Protocol (TCP/IP) Properties** dialog box. Make sure that the addresses are the same in both places. However, if you are configuring a virtual private network (VPN) load balancing cluster, you should not configure the dedicated IP address. On a VPN, only the cluster IP address should be present on each of the cluster hosts.<br><br>• The dedicated IP address must be a static IP address. It cannot be a DHCP address. |
| **Subnet mask** on the **Host Parameters** tab | This parameter denotes the subnet mask for the IP address specified. The mask is entered in standard Internet dotted notation (for example, 255.255.255.0). |
| **Initial host state** on the **Host Parameters** tab | This parameter specifies whether Network Load Balancing will start and whether the host will immediately join the cluster when the operating system is started. For example, you might want to start other services manually and in a specific order before starting Network Load Balancing. Hosts can be commanded to join and leave the cluster dynamically by using the **start** and **stop** commands in Network Load Balancing command-line control. If the **Retain suspended state after computer starts** check box is selected, when the host is shut down while in a suspended state, the host will remain suspended when Windows is started.<br><br>For more information on command-line control, see Managing Network Load Balancing from the command line. |

**Port rules**

To maximize control of various types of TCP/IP traffic, you can set up port rules to control how each port's cluster network traffic is handled. The method by which a port's network traffic is handled is called its filtering mode. There are three possible filtering modes: **Multiple hosts**, **Single host**, and **Disabled**.

You can also specify that a filtering mode apply to a numerical range of ports. You do this by defining a port rule with a set of configuration parameters that define the filtering mode. Each rule consists of the following configuration parameters:

- The virtual IP address that the rule should be applied.
- The TCP or UDP port range for which this rule should be applied.
- The protocols for which this rule should apply, including TCP, UDP, or both.
- The filtering mode that specifies how the cluster handles traffic described by the port range and protocols.

In addition, you can select one of three options for client affinity: None, Single, and Class C. Single and Class C are used to ensure that all network traffic from a particular client be directed to the same cluster host. In order to allow Network Load Balancing to properly handle IP fragments, you should avoid using **None** when selecting **UDP** or **Both** for your protocol setting.

**Important**

- The number and type of rules must be exactly the same for each host in the cluster.
- You cannot add more than 32 port rules to a Network Load Balancing cluster.
- If a host attempts to join the cluster with a different number of rules from the other hosts, it is not accepted as part of the cluster and the rest of the cluster continues to handle the traffic as before. At the same time, a message is entered into the Windows event log. When this happens, consult the event log to determine which host has a conflicting number of rules, resolve the conflict, and restart Network Load Balancing on this host.
- The rules entered on each host in the cluster must have matching cluster IP addresses, port ranges, protocol types, and filtering modes.

  If Network Load Balancing detects an inconsistent rule among the hosts in the cluster, it records a message in the Windows event log. When this happens, consult the event log to determine the host in question and which rule is responsible for the problem, fix it, and restart Network Load Balancing on this host.

  For more information on error logging and cluster operations, see Error detection and handling and Event logging.
- When using Network Load Balancing to load balance VPN traffic such as PPTP/GRE and IPSEC/L2TP, you must configure the port rules that govern the ports handling the VPN traffic (TCP port 1723 for PPTP and UDP port 500 for IPSEC) to use either Single or Class C affinity.

**Note**

- By default, all cluster network traffic not governed by port rules is handled by the host with the highest host priority among the current members of the cluster. This single host handles all of the cluster network traffic, with another host taking over the traffic in the event that the highest priority host fails or goes offline. This default behavior ensures that Network Load Balancing does not affect cluster network traffic for ports that you do not specifically manage with the Network Load Balancing load-balancing mechanisms. It also provides high availability in the handling of your cluster network traffic.

The **Port Rules** tab shows only the summary of existing port rules. If you click **Add** or **Edit**, the following parameters will be available:

| Parameter | Description |
|---|---|
| **Cluster IP address** on the **Add/Edit Port Rules** dialog box | This parameter specifies the cluster IP address that the port rule should cover. If this parameter is left blank and **All** is selected instead, the port rule is a global port rule and will cover all cluster IP addresses associated with that particular Network Load Balancing cluster.<br>**Note**<br>• If a cluster IP address is specified, the port rule overrides any conflicting global port rule for that particular cluster IP address. |
| **Port range** on the **Add/Edit Port Rules** dialog box | This parameter specifies the TCP/UDP port range that a port rule should cover. Port numbers in a range of 0 to 65,535 are currently supported. The default port range is 0 to 65,535.<br>**Note**<br>• Rules for a single port are encoded as a range having the same starting and ending port numbers. |

| | |
|---|---|
| **Protocols** on the **Add/Edit Port Rules** dialog box | This parameter lets you choose the specific TCP/IP protocol that a port rule should cover: TCP, UDP, or both. Only the network traffic for the specified protocol is affected by the rule. Traffic not affected by the port rule is handled by the default host. |
| **Filtering mode** on the **Add/Edit Port Rules** dialog box | • **Multiple hosts.** This parameter specifies that multiple hosts in the cluster handle network traffic for the associated port rule. This filtering mode provides scaled performance in addition to fault tolerance by distributing the network load among multiple hosts. You can specify that the load be equally distributed among the hosts or that each host handle a specified load weight.<br><br>• **Single host.** This parameter specifies that network traffic for the associated port rule be handled by a single host in the cluster according to the specified handling priority. This filtering mode provides port specific fault tolerance for the handling of network traffic.<br><br>• **Disable this port range.** This parameter specifies that all network traffic for the associated port rule be blocked. In this case, the Network Load Balancing driver filters all corresponding network packets or datagrams. This filtering mode lets you block network traffic addressed to a specific range of ports. |

| | |
|---|---|
| **Affinity** on the **Add/Edit Port Rules** dialog box | This parameter is applicable only for Multiple host filtering mode.<br><br>• The **None** option specifies that multiple connections from the same client IP address can be handled by different cluster hosts (no client affinity). In order to allow Network Load Balancing to properly handle IP fragments, you should avoid using **None** when selecting **UDP** or **Both** for your protocol setting.<br><br>• The **Single** option specifies that Network Load Balancing should direct multiple requests from the same client IP address to the same cluster host. This is the default setting for affinity.<br><br>You can optionally modify Network Load Balancing client affinity to direct all client requests from a TCP/IP Class C address range (instead of a single IP address) to a single cluster host by enabling the **Class C** option instead of the **Single** option. This feature ensures that clients that use multiple proxy servers to access the cluster can have their TCP connections directed to the same cluster host. The use of multiple proxy servers at the client's site causes requests from a single client to appear to originate from different computers. Assuming that all of the client's proxy servers are located within the same Class C address range, enabling the **Class C** option ensures that client sessions are properly handled. If you do not need this capability, use the **Single** option instead to maximize scaled performance when using client affinity.<br><br>• **Class C** affinity specifies that Network Load Balancing direct multiple requests from the same TCP/IP Class C address range to the same cluster host.<br><br>Enabling **Class C** affinity instead of **Single** affinity ensures that clients that use multiple proxy servers to access the cluster have their TCP connections directed to the same cluster host. The use of multiple proxy servers at the client's site causes requests from a single client to appear to originate from different computers. Assuming that all of the client's proxy servers are located within the same Class C address range, **Class C** affinity ensures that client sessions are properly handled. If you do not need this capability, use **Single** affinity instead to maximize scaled performance.<br><br>Enabling either **Single** or **Class C** affinity ensures that only one cluster host handles all connections that are part of the same client session. This is important if the server application running on the cluster host maintains session state (such as "server cookies") between connections.<br><br>At the same time, it is important to realize that this does not preserve session state with back-end databases in which many different transactions are occurring involving many different computers. Once the connection is ended, session state also ends.<br><br>Disabling affinity allows for improved load balancing because it allows multiple connections from the same client to be handled concurrently by different cluster hosts. To maximize scaled performance, disable client affinity (using the **None** option) when it is not needed. However, in order to allow Network Load Balancing to properly handle IP fragments, you should avoid using **None** when selecting **UDP** or **Both** for your protocol setting.<br><br>**Important**<br><br>• When using Network Load Balancing to load balance VPN traffic such as PPTP/GRE and IPSEC/L2TP, you must configure the port rules that govern the ports handling the VPN traffic (TCP port 1723 for PPTP and UDP port 500 for IPSEC) to use either Single or Class C affinity. |

| | |
|---|---|
| **Load Weight** on the **Add/Edit Port Rules** dialog box | This parameter is applicable only for Multiple host filtering mode. You can configure this parameter only when you open the port rules dialog box through **Host Properties**. This parameter is not configurable when you open the port rules dialog box through **Cluster Properties**.<br><br>When using Multiple hosts filtering mode, this parameter specifies the relative amount of load-balanced network traffic that this host should handle for the associated port rule. Allowed values range from 0 (zero) to 100. To prevent a host from handling any network traffic, set the load weight to 0 (zero).<br><br>The actual fraction of traffic handled by each host is computed as the local load weight divided by the sum of all load weights across the cluster.<br><br>You can specify different load weights for each host in the cluster by using the **Load weight** parameter. You can specify that all hosts distribute the network load equally by using the Equal load distribution parameter instead of the **Load weight** parameter. |
| **Handling priority** on the **Add/Edit Port Rules** dialog box | This parameter is applicable only for Single host filtering mode. You can configure this parameter only when you open the port rules dialog box through **Host Properties**. This parameter is not available when you open the port rules dialog box through **Cluster Properties**<br><br>When Single host filtering mode is being used, this parameter specifies the local host's priority for handling the networking traffic for the associated port rule. The host with the highest handling priority (lowest numerical value) for this rule among the current members of the cluster will handle all of the traffic for this rule. The allowed values range from 1, the highest priority, to the maximum number of hosts allowed (32). This value must be unique for all hosts in the cluster.<br><br>Although this parameter is displayed in the **Defined port rules** list, you configure this parameter on the **Host Parameters** tab. |

**Notes**

- The parameters set in the **Network Load Balancing Properties** dialog box are recorded in the registry.
- Changes to Network Load Balancing parameters are applied when you click **OK** in the **Network Load Balancing Properties** dialog box. Clicking **OK** stops Network Load Balancing (if it is running), reloads the parameters, and then restarts cluster operations.
- WLBS stands for Windows NT Load Balancing Service, the former name of Network Load Balancing in Windows NT 4.0. For reasons of backward compatibility, WLBS continues to be used in certain instances.

# Virtual clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual clusters**

You create a virtual cluster by configuring multiple virtual IP addresses (where each IP address typically corresponds to different Web sites or applications hosted on the cluster) on a single network adapter and then configuring different port rules for each virtual IP address. In this way, on each host you are able to have IP address-specific:

- Port Range
- Protocols
- Affinity
- Load weight
- Filtering mode

By using a load weight of zero for specific virtual IP addresses, you can also define a virtual cluster that limits applications, Web sites, or virtual IP addresses to a specific subset of hosts within your Network Load Balancing (NLB) cluster. See the diagram below for an example of this use of virtual clusters.




| IP Address | Load Weights | | | |
|---|---|---|---|---|
| xxx.xxx.xxx.1 (Web site A) | 25 | 25 | 25 | 25 |
| xxx.xxx.xxx.2 (Web site B) | 50 | 50 | 0 | 0 |
| xxx.xxx.xxx.3 (Web site C) | 0 | 0 | 50 | 50 |

This diagram depicts a four host Network Load Balancing cluster. Through the use of virtual clusters defined by IP address-specific load weights, Network Load Balancing directs network traffic as follows:

- Users accessing Web Site A (IP address *nnn.nnn.nnn*.1) are directed to any of the four hosts
- Users accessing Web site B (IP address *nnn.nnn.nnn*.2), which is a virtual cluster, are directed to hosts 1 and 2
- Users accessing Web site C (IP address *nnn.nnn.nnn*.3), which is a virtual cluster, are directed to hosts 3 and 4

This is accomplished by setting the load weights as follows:

- Load weight = 0 on hosts 3 and 4 for IP address xxx.xxx.xxx.2
- Load weight = 0 on hosts 1 and 2 for IP address xxx.xxx.xxx.3.

Be aware that to use virtual clusters, all hosts in the cluster must be running one of the Windows Server 2003 family of products.

For more information on creating port rules, see Create a new port rule.

# Upgrading Network Load Balancing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrading Network Load Balancing**
Individual computers participating in a Network Load Balancing cluster can be dynamically upgraded to a newer version of Network Load Balancing without interrupting service on the cluster as a whole.

To upgrade a cluster host, simply upgrade the computer to the current version of Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. All settings are retained. It is therefore not necessary to reconfigure Network Load Balancing parameters after you upgrade. In addition, the host automatically rejoins the cluster when the upgrade is complete.

Network Load Balancing clusters can function successfully in a mixed environment of hosts running either Network Load Balancing for the Windows server 2003 family of products or Windows 2000, Windows NT Load Balancing Service (WLBS) for Windows NT 4.0, or Convoy Cluster Software for Windows NT 4.0. After upgrading to Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition, you should run your cluster in a mixed environment for a period of time to ensure your cluster continues to operate as expected on your upgraded hosts. However, you will not be able to use the new Network Load Balancing features, such as virtual clusters, until all of your hosts are upgraded.

**Note**

- Windows NT Load Balancing Service (WLBS) is the former name of Network Load Balancing in Windows NT 4.0. Convoy Cluster Software is the former name of WLBS.

# Adjusting multicast parameters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adjusting multicast parameters**

Network Load Balancing uses one of two methods to enable the cluster adapters on all cluster hosts to simultaneously receive network traffic for the cluster's IP addresses:

- When multicast support is enabled in the **Network Load Balancing Properties** dialog box, Network Load Balancing adds a multicast media access control (MAC) address to the cluster adapters on all cluster hosts. At the same time, the cluster adapters retain their original MAC addresses.

- When multicast support is disabled in the **Network Load Balancing Properties** dialog box (causing the host to revert to unicast mode), Network Load Balancing instructs the driver belonging to the cluster adapter to change the adapter's MAC address to the cluster MAC address. This is the same address used on all cluster hosts. (Note that some network adapters do not support changing their MAC addresses.) When this occurs, the cluster adapters are no longer addressable by their original MAC addresses.

**Advantages of using multicast support**

Multicast support has several advantages, such as removing the inherent limitations of a single network adapter, as explained in Single network adapter limitations. Network Load Balancing handles the resolution of the cluster's IP address to its associated multicast cluster MAC address used in ARP on your router. Note that in rare cases the use of a multicast MAC address might not be supported by the ARP implementation on a given router. If this problem arises, the cluster will not be accessible from outside the local subnet. To correct this problem, you must create a static ARP entry within the router. To determine how to create a static ARP entry, see the documentation for your router. For more information, see Using a router in Network Load Balancing system requirements.

**Avoiding switch port flooding**

Some network adapter drivers experience increased CPU load percentages (by 5 percent or more depending on network traffic) when handling multicast MAC addresses instead of unicast addresses. If your cluster hosts have unusually high CPU load percentages when using Network Load Balancing multicast support, disable this feature to see if it is the source of the problem. If switch flooding occurs, that is, your non-cluster hosts receive traffic destined only for the cluster hosts, then you can enable Internet Group Management Protocol (IGMP) support. For more information on avoiding switch port flooding, see Limit switch port flooding. For more information on enabling IGMP support, see Enable Internet Group Management Protocol (IGMP) support.

### Notes

- Multicast support is not enabled by default. To enable it, see Enable multicast support.

- When Network Load Balancing is enabled for use with a network adapter in multicast mode, a manually configured MAC address that may have been optionally assigned to the adapter will be removed after Network Load Balancing is later disabled. When the adapter is subsequently reinitialized (typically when the computer is restarted), the network adapter will revert to its factory-assigned MAC address. To avoid this problem, be sure to reassign the adapter's manually configured MAC address after disabling Network Load Balancing. For more information, see Remove a host from a Network Load Balancing cluster.

# Handling LAN failures

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Handling LAN failures**

If a Network Load Balancing cluster is split into multiple subnets that cannot communicate with each other, each part will reconverge as an individual cluster ready to handle all of the traffic for the cluster IP address. (This problem is distinct from a networking failure in which a single host goes offline. In that case, the host just leaves the cluster.)

If multiple subnets with clusters remain active (continue to handle incoming requests), a conflict arises when the networking failure is corrected and the cluster rejoins. In this case, some outstanding connections might have to be closed in order to establish a consistent state for the rejoined cluster. This situation occurs infrequently. It is far more probable that only one subnet will have access to the Internet, in which case no conflict arises because the other subnets do not provide service during the network outage.

# Error detection and handling

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Error detection and handling**

Network Load Balancing error-handling mechanisms are designed to minimize the possibility of disrupting the cluster's service to client requests, while allowing the cluster's parameters and member hosts to be dynamically changed. For example, you can add hosts to the cluster, remove them for maintenance, add port rules, and modify rule Network Load Balancing parameters, all without interrupting service.

The Network Load Balancing driver checks its configuration parameters entered in the **Network Load Balancing Properties** dialog box for consistency and validity. Network Load Balancing will not start cluster operations if a problem, such as an invalid IP address, is detected. All detected errors are reported in the Windows event log. For more information, see Event logging.

When a local host joins the cluster, it exchanges its primary parameter values, such as its host priority and port rules, with the other hosts in the cluster. If an inconsistency is detected among the hosts, Network Load Balancing recovers from the error according to the type of inconsistency detected, as follows:

| Inconsistency | Network Load Balancing recovery method |
|---|---|
| Duplicate host priorities | The host joining the cluster handles no traffic until convergence is complete. The cluster does not converge to a consistent state until duplicate hosts are no longer detected. |
| Different number of port rules | The host joining the cluster handles no traffic until convergence is complete. The cluster does not converge to a consistent state until all hosts have the same number of rules. |
| Inconsistent port rule | The host joining the cluster handles no traffic until convergence is complete. The cluster does not converge to a consistent state until the rule is consistent across all hosts in the cluster. |

If inconsistencies are discovered when a host attempts to join the cluster, the host should be taken out of the cluster using the **nlb.exe stop** command. Locate and correct the problem using the **Network Load Balancing Properties** dialog box on each host prior to returning the host to the cluster using the **nlb.exe start** command.

You can add port rules or update parameters by taking each host out of the cluster in turn, updating its parameters, and then returning it to the cluster. During this process, other cluster hosts detects inconsistencies in the rules and handles them as described previously. This permits the cluster to maintain service for client requests while the changes are made. For example, if a rule is added, it does not take effect until all hosts have been updated and have rejoined the cluster. If a rule is modified, the rule is not used until it is made consistent on all hosts.

**Note**

- WLBS stands for Windows NT Load Balancing Service, the former name of Network Load Balancing in Windows NT 4.0. For reasons of backward compatibility, WLBS continues to be used in certain instances.

# Event logging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Event logging**

Network Load Balancing uses Windows event logging to record information about important events. You can view the event log using Event Viewer.

Network Load Balancing classifies logged events into the following categories:

| Event category | Description |
|---|---|
| Information events | Records normal state transitions, such as driver startup, the completion of convergence, and the completion of cluster-control actions. These events are recorded for informational purposes only. |
| Warning events | Records abnormal occurrences that do not impede cluster operations. These events can indicate a configuration problem, such as an invalid dedicated IP address or other parameter. They can also indicate a temporary condition, such as limited driver resources. Cluster operations are not disrupted when warning events occur, and are not considered errors. |
| Error events | Records major problems or errors that might require immediate attention, such as significant cluster configuration parameter problems, detected inconsistencies, or unexpected execution errors. Error events usually disrupt cluster operations and might require individual hosts to be stopped and reconfigured. Note that errors can occur under normal operations when you update the cluster by adding or modifying port rules. For more information, see Error detection and handling. |

Event log records provide crucial information when you are analyzing and solving problems. Use the **nlb.exe display** command to display the most recent event log records produced by Network Load Balancing. This information should be copied into the text of any e-mail message sent to Microsoft requesting technical support. For more information, see Nlb.

**Network Load Balancing Manager logging**

Network Load Balancing Manager does not use Windows event logging and therefore you can not use Event Viewer to review messages from Network Load Balancing Manager. You instead must explicitly enable Network Load Balancing Manager logging and define a specific log file. Enable Network Load Balancing Manager logging by clicking **Log Settings** in the Network Load Balancing Manager **Options** menu. Check the **Enable logging** box and specify a name and location for the log file.

You should always manage a Network Load Balancing cluster from a trusted and secure computer, particularly when logging is enabled. The Network Load Balancing Manager log file contains potentially sensitive information about the Network Load Balancing cluster and hosts, so it must be properly secured. By default, the log file inherits the security settings of the directory in which it is created, so you might have to change the explicit permissions on the file to restrict read and write access to those individuals who don't need full control of the file. Be aware that the individual using Network Load Balancing Manager does require full control of the log file. For more information, see Access Control How To....

### Notes

- WLBS stands for Windows NT Load Balancing Service, the former name of Network Load Balancing in Windows NT 4.0. For reasons of backward compatibility, WLBS continues to be used in certain instances.

- For information on how to interpret messages, refer to the Microsoft Windows Resource Kits Web site.

# Network Load Balancing and stateful connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Load Balancing and stateful connections**

Application servers maintain two kinds of stateful connections:

- Interclient state: A state whose updates must be synchronized with transactions performed for other clients, such as merchandise inventory at an e-commerce site.
- Intraclient state: A state that must be maintained for a given client throughout a session (that can span multiple connections), such as a shopping cart process at an e-commerce site.

Network Load Balancing should not be used to scale applications that directly update interclient state, such as Microsoft SQL Server, because these applications generally were not designed to permit multiple instances to simultaneously access a shared database and synchronize updates. Instead, Network Load Balancing should be used to scale stateless front-end services, such as Microsoft Internet Information Services, that might access a shared back-end database server.

However, Network Load Balancing can be used to scale applications that manage intraclient state within a session that spans multiple connections. When client affinity is enabled, Network Load Balancing directs all TCP connections to the same cluster host. This allows session state to be maintained in host memory. Client/server applications that embed state within cookies or push it to a back-end database do not need client affinity to be maintained.

# Managing Network Load Balancing from the command line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing Network Load Balancing from the command line**

After you have installed and configured Network Load Balancing, you can control its operations and modify parameter settings using the Network Load Balancing control application, Wlbs.exe, that is installed in the *systemroot*\System32 folder. To simplify and centralize system administration, you can run Nlb.exe either on the cluster hosts or on any remote computer that can access the cluster over a local or wide area network. Certain actions, such as modifying parameters, can be performed only on the cluster hosts. Because it is a shell-based application, this application can be incorporated into administrative scripts.

The command line for Nlb.exe has the following syntax:

**Nlb** *command* [*cluster*[**:***host*] [*remote options*] ]

For more information on the Network Load Balancing control application, see Nlb. For information about using Network Load Balancing Manager from the command line, see Nlbmgr. For information about other command-line utilities, see Command-line reference A-Z. For more information about manageability, see Management Strategies and Tools.

**Controlling remote access**

You can help protect the cluster from unauthorized access by remote computers by using a strong remote control password. The password is assigned in the **Network Load Balancing Properties** dialog box for each cluster host to be protected. When password protection is enabled, remote invocations of Nlb.exe must include the **/PASSW** option followed by the password. If no password value is entered on the command line, the shell displays a prompt for a password.

**Caution**

- The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. It is highly recommended that you do not enable remote control and instead use Network Load Balancing Manager or other remote management tools such as Windows Management Instrumentation (WMI).




If you choose to enable remote control, it is vital that you restrict access by specifying a strong remote control password. It is also imperative that you use a firewall to protect the Network Load Balancing UDP control ports (the ports that receive remote control commands) in order to shield them from outside intrusion. By default, these are ports 1717 and 2504 at the cluster IP address. Use remote control only from a secure, trusted computer within your firewall. For more information on the remote control parameter, see Remote control in Network Load Balancing parameters. For more information about strong passwords, see Strong passwords.

**Note**

- WLBS stands for Windows NT Load Balancing Service, the former name of Network Load Balancing in Windows NT 4.0. For reasons of backward compatibility, WLBS continues to be used in certain instances.

# Network Load Balancing Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

The following links point to information that might help you better understand Network Load Balancing within the overall context of Windows Clustering.

- For information on the solutions that make up Windows Clustering, refer to Windows Clustering.
- For information on Server clusters, refer to Server clusters overview and the Microsoft Windows Resource Kits Web site.
- For information on server clusters and how they work with Network Load Balancing clusters, refer to Updated technical information.
- For a detailed description of how a business has used Network Load Balancing clusters and server clusters to improve performance and availability of their Web site and other services, refer to the scenario at the Microsoft Web site.

Maximizing the uptime of server systems requires careful attention to planning, design, and operations. For more information about monitoring recommendations for high-availability systems, see the "System Performance and Troubleshooting Guide" at the Microsoft Windows Resource Kits Web site. For more information about maximizing the availability of a Network Load Balancing cluster, see the "Server Management Guide" at the Microsoft Windows Resource Kits Web site and "High availability with clustering and NLB" at the Microsoft Windows Resource Kits Web site.

Microsoft maintains up-to-date product documentation at the Microsoft Web site . This site has information that is more current than the Help documentation on your product CD.

Microsoft Product Support Services provides a searchable database of technical support information for Microsoft products. You can search here for general information about Network Load Balancing or for information related to other Microsoft products. You can get to the Microsoft Product Support Services Web site by clicking the appropriate link in Support resources.

For information about unattended setup, refer to Planning for unattended Setup.

# Network Load Balancing Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2012

**Troubleshooting**
What problem are you having?

- After installing Network Load Balancing and restarting a cluster host, a message appears: "The system has detected an IP address conflict with another system on the network..."
- There is no response when using ping to access the cluster's IP address from an outside network.
- There is no response when using ping to access a host's dedicated IP addresses from another cluster host.
- When attempting to use Network Load Balancing Manager to connect to a host in your cluster, you receive the error "Host unreachable".
- When using Telnet or attempting to browse a computer outside the cluster from a cluster host, there is no response.
- When invoking the Network Load Balancing remote control commands from a computer outside the cluster, there is no response from one or more cluster hosts.
- When using the dedicated IP address of a host to specify it as a target for a remote control command, there is no reply. However, specifying the host by its priority (ID) works.
- Connectivity to the cluster is denied to some, but not all users.
- You cannot view or change the Network Load Balancing properties using net config and Windows Management Instrumentation (WMI).
- An unusual number of TCP connections to the cluster's IP address are being reset by either the server or the client.
- Virtual Private Network (VPN) calls fail when you make a change that causes convergence (such as adding a host, removing a host, or draining a host).
- After the cluster hosts start, they begin converging but never complete convergence.
- The cluster moves in and out of a converged state.
- After the cluster hosts start, Network Load Balancing reports that convergence has finished, but more than one host is a default host.
- Network Load Balancing is not load balancing applications, and the default host handles all network traffic.
- Traffic alternates unexpectedly between the cluster hosts breaking TCP connections.
- Network traffic does not appear to load balance evenly among the cluster hosts.
- When you are using Network Load Balancing with Microsoft Internet Security and Acceleration (ISA) Server, one cluster host might log blocked packets that are directed to the dedicated Internet Protocol (IP) address of another host.
- You are unable to create a Network Load Balancing cluster in a 64-bit environment.

**After installing Network Load Balancing and restarting a cluster host, a message appears: "The system has detected an IP address conflict with another system on the network..."**

**Possible Cause**

The same IP address exists somewhere else on the network.

**Solution**

Choose a new IP address, or remove the duplicate address.

**Possible Cause**

You have configured different cluster operation modes (Unicast and Multicast) on the hosts causing two different MAC addresses to map to the same IP address.

**Solution**

Ensure that all hosts are configured with the same cluster operation mode. For more information see Enable Network Load Balancing.

**Possible Cause**

You configured the cluster's IP address before Network Load Balancing was bound to the network adapter.

**Solution**

Remove the cluster's IP address from TCP/IP properties, enable Network Load Balancing on the proper adapter as described in Enable Network Load Balancing and then configure the cluster's IP address.

**Possible Cause**

You added the cluster's IP address to a network adapter that has not been enabled for Network Load Balancing.

**Solution**

Remove the cluster's IP address from incorrect adapter's TCP/IP properties, enable Network Load Balancing on the proper adapter as described in Enable Network Load Balancing and then configure the cluster's IP address.

**Possible Cause**

You added the cluster's IP address in the TCP/IP dialog box, but you did not add the IP address correctly using the Network Load Balancing dialog box.

**Solution**

If you do not use Network Load Balancing Manager to configure your cluster, you must manually configure TCP/IP with the cluster's IP address. For more information, see Set up TCP/IP for Network Load Balancing.

**There is no response when using ping to access the cluster's IP address from an outside network.**

Verify that you can use **ping** to access the dedicated IP addresses for the cluster hosts from a computer outside the router. If this test fails and you are using multiple network adapters, the problem is unrelated to Network Load Balancing. If you are using a single network adapter for both the dedicated and cluster IP addresses, consider the following causes:

**Possible Cause**

You did not add the cluster's IP address correctly in the TCP/IP properties.

**Solution**

If you do not use Network Load Balancing Manager to configure your cluster, you must manually configure TCP/IP with the cluster's IP address. For more information, see Set up TCP/IP for Network Load Balancing

**Possible Cause**

If you are using multicast support, you might find that your router has difficulty resolving the primary IP address into a multicast media access control (MAC) address using the ARP protocol.

**Solution**

To check this, verify that you can use **ping** to access the cluster from a client on the cluster's subnet and to access the cluster

hosts' dedicated IP addresses from a computer outside the router. If these tests work properly, the router is probably at fault. You should be able to add a static ARP entry to the router to circumvent the problem. You can also turn off Network Load Balancing multicast support and use a unicast network address without a hub.

**Possible Cause**

When using Network Load Balancing in either multicast or unicast mode, routers need to be able to accept proxy ARP responses (IP-to-network address mappings that are received with a different network source address in the Ethernet frame).

**Solution**

Make sure that your router has proxy ARP support turned on. You can also set a static ARP entry to keep proxy ARP support disabled in the router.

**Possible Cause**

Internet control message protocol (ICMP) to the cluster is blocked by a router or firewall.

**Solution**

Allow ICMP traffic through the firewall or router. Be aware that this may expose your system to additional security risk.

**There is no response when using ping to access a host's dedicated IP addresses from another cluster host.**
**Possible Cause**

You did not add the host's dedicated IP addresses correctly in TCP/IP properties.

**Solution**

You must manually configure TCP/IP with the host's dedicated IP address. For more information, see Set up TCP/IP for Network Load Balancing

**Possible Cause**

When using Network Load Balancing in either multicast or unicast mode, routers need to be able to accept proxy ARP responses (IP-to-network address mappings that are received with a different network source address in the Ethernet frame).

**Solution**

Make sure that your router has proxy ARP support turned on. You can also set a static ARP entry to keep proxy ARP support disabled in the router.

**Possible Cause**

Internet control message protocol (ICMP) to the cluster is blocked by a router or firewall.

**Solution**

Allow ICMP traffic through the firewall or router. Be aware that this may expose your system to additional security risk.

**When attempting to use Network Load Balancing Manager to connect to a host in your cluster, you receive the error "Host unreachable".**
**Cause:** Internet control message protocol (ICMP) to the host is either blocked by a router or firewall, or disabled on the host's network adapter.

**Solution:** Enable ICMP on the host's network adapter or allow ICMP traffic through the firewall or router. Be aware that this may expose your system to additional security risk. You can also use Network Load Balancing Manager's **/noping** option. For more information, see Nlbmgr.

**When using Telnet or attempting to browse a computer outside the cluster from a cluster host, there is no response.**
**Cause:** Verify that you can use **ping** to access the computer outside the cluster. If this test is successful you might not have listed the host's dedicated IP address first in the TCP/IP properties.

**Solution:** When you must manually configure TCP/IP with the host's dedicated IP addresses and the cluster's IP address, the dedicated IP addresses must be listed first. For more information, see Set up TCP/IP for Network Load Balancing.

If ping fails to access the computer outside of the cluster refer to the problems labeled "There is no response when using **ping** to access the cluster's IP address from an outside network" and "There is no response when using **ping** to access a host's dedicated IP addresses from another cluster host" described earlier in this troubleshooting information.

**When invoking the Network Load Balancing remote control commands from a computer outside the cluster, there is no response from one or more cluster hosts.**

**Possible Cause**

Remote control commands are not being sent to the cluster's IP address. Commands must be sent to the cluster's primary IP address that was assigned in the **Network Load Balancing Properties** dialog box.

**Solution**

Be sure that you send remote commands to the correct IP address.

**Possible Cause**

The remote control traffic is being encrypted by Internet Protocol security (IPSec). Network Load Balancing remote control commands will not work correctly if they are sent from a computer that has IPSec configured such that the remote control traffic is encrypted by IPSec.

**Solution**

Disable IPSec.

For more information, see Internet Protocol Security (IPSec).

**Possible Cause**

Remote control is not enabled.

**Solution**

Enable remote control in the **Network Load Balancing Properties** dialog box.

For more information, see Configure cluster parameters.

**Caution**

- The Network Load Balancing remote control option presents many security risks, including the possibility of data tampering, denial of service and information disclosure. It is highly recommended that you do not enable remote control and instead use Network Load Balancing Manager or other remote management tools such as Windows Management Instrumentation (WMI).




If you choose to enable remote control, it is vital that you restrict access by specifying a strong remote control password. It is also imperative that you use a firewall to protect the Network Load Balancing UDP control ports (the ports receiving remote control commands) in order to shield them from outside intrusion. By default, these are ports 1717 and 2504 at the cluster's IP address. Use remote control only from a secure, trusted computer within your firewall. For more information on the remote control parameter, see Remote control in Network Load Balancing parameters. For more information about strong passwords, see Strong passwords.

**Possible Cause**

Network Load Balancing UDP control ports are protected incorrectly by a firewall. By default, remote control commands are sent to UDP ports 1717 and 2504 at the cluster IP address.

**Solution**

Be sure that these ports have not been blocked incorrectly by a router or firewall. You can also change the port number by modifying the corresponding Network Load Balancing parameter.

**Possible Cause**

You used an incorrect password when attempting to use remote control.

**Solution**

Use the remote control password that you configured when you enabled the Network Load Balancing cluster.

**When using the dedicated IP address of a host to specify it as a target for a remote control command, there is no reply. However, specifying the host by its priority (ID) works.**
**Cause:** None of the hosts have a dedicated IP addresses.

**Solution:** Assign a dedicated IP address to each host. For more information, see Configure host parameters.

**Connectivity to the cluster is denied to some, but not all users.**
**Possible Cause**

An application being load balanced in not responding.

**Solution**

This is an application specific problem not related to Network Load Balancing. Refer to your application's documentation for correction. You may need to stop and restart the application.

**Possible Cause**

If your cluster is configured for unicast mode, a switch might have learned the Network Load Balancing network adapter's MAC address.

**Solution**

Clear the switch's port to MAC address mapping.

**Possible Cause**

The cluster's IP address was not added to TCP/IP on one or more of the hosts.

**Solution**

If you do not use Network Load Balancing Manager to configure your cluster, you must manually configure TCP/IP with the cluster's IP address. For more information, see Set up TCP/IP for Network Load Balancing.

**Possible Cause**

A host is leaving the cluster because of a drainstop or stop command, but convergence did not complete correctly.

**Solution**

Wait for convergence to complete. If convergence does not complete, see the problem titled "After the cluster hosts start, they begin converging but never complete convergence" later in this troubleshooting guide.

**You cannot view or change the Network Load Balancing properties using net config and Windows Management Instrumentation (WMI).**
**Cause:** In order to view or change **Network Load Balancing properties**, you must be a member of the Administrators group.

**Solution:** Log on as a user who is in the local Administrators group of the computer that is running Network Load Balancing.

**An unusual number of TCP connections to the cluster's IP address are being reset by either the server or the client.**
**Possible Cause**

HTTP keep-alives are enabled on the Network Load Balancing hosts and keep-alive enabled clients are connecting to the cluster.

**Solution**

Disable HTTP keep-alives. For more information on HTTP keep-alives and Internet Information Services (IIS), refer to the IIS documentation set. To view the IIS documentation set from your desktop, install IIS then click the Start button, click Run, and type the following command in the Open text box: **%windir%\help\iisrv.chm**.

**Possible Cause**

Low system resources on the server are causing TCP itself to reject the connections.

**Solution**

Free up system resources by for example, adding additional system memory or closing unnecessary applications.

**Possible Cause**

The cluster has diverged into two separately converged clusters, causing more than one node to claim ownership of every connection.

**Solution**

Remove the two clusters then recreate a single cluster.

**Virtual Private Network (VPN) calls fail when you make a change that causes convergence (such as adding a host, removing a host, or draining a host).**
**Cause:** When using Network Load Balancing to load balance VPN traffic such as PPTP/GRE and IPSEC/L2TP, you must configure the port rules that govern the ports handling the VPN traffic (TCP port 1723 for PPTP and UDP port 500 for IPSEC) to use either Single or Class C affinity.

**Solution:** Configure the port rules governing ports 500 and 1723 to use either Single or Class C affinity. For more information see Network Load Balancing parameters.

**After the cluster hosts start, they begin converging but never complete convergence.**
**Possible Cause**

Either a different number of port rules or incompatible port rules on different cluster hosts were entered. This will inhibit convergence.

**Solution**

Open the **Network Load Balancing Properties** dialog box on each cluster host and verify that all hosts have identical port rules.

**Possible Cause**

You have a bad network adapter or cable.

**Solution**

Use the ping command to test connectivity. Ping the host's fully qualified domain name. You can also learn more about the problem by pinging your domain controller by IP address, and by pinging other network servers by name and IP address.

**Possible Cause**

You have mismatched duplex settings on a switch or hub.

**Solution**

Confirm that duplex settings in each of your switches and hubs is configured appropriately.

**Possible Cause**

The dedicated IP address that you used for one of the hosts already exists somewhere else on the network.

**Solution**

Choose a new IP address, or remove the duplicate address.

**Possible Cause**

You defined one or more port rules that are applicable to only specific IP addresses, but your cluster contains hosts running Windows 2000.

**Solution**

If you use IP address specific port rules, your cluster must be running a product in the Windows Server 2003 family on all hosts.

**Possible Cause**

You have configured different cluster operation modes (Unicast and Multicast) on the hosts.

**Solution**

Ensure that all hosts are configured with the same cluster operation mode. For more information see Enable Network Load Balancing.

**The cluster moves in and out of a converged state.**
**Cause:** Heartbeats are being missed due to intermittent network connectivity caused by a bad network adapter or cable, or other network problems

**Solution:** Use the ping command to test connectivity. Ping the host's fully qualified domain name. You can also learn more about the problem by pinging your domain controller by IP address, and by pinging other network servers by name and IP address.

**After the cluster hosts start, Network Load Balancing reports that convergence has finished, but more than one host is a default host.**
**Possible Cause**

The cluster hosts have become members of different subnets, so that all hosts are not accessible on the same network.

**Solution**

Be sure that all cluster hosts can communicate with each other.

**Possible Cause**

A layer three switch is being used.

**Solution**

Put a layer two switch between the hosts and the layer three switch.

**Possible Cause**

A break in a redundant switch caused the cluster to separate into two separate clusters, creating two separate default hosts.

**Solution**

Remove the two clusters then recreate a single cluster.

**Possible Cause**

Your switch is configured to reject broadcast packets.

**Solution**

Configure your switch to accept broadcast packets (be aware that this might introduce certain security risks), or configure your Network Load Balancing cluster to use multicast mode.

**Possible Cause**

One host in unable to send or receive heartbeats.

**Solution**

Use the ping command to test connectivity to each of the hosts. Ping the hosts' fully qualified domain name.

**Possible Cause**

A host is plugged into the wrong port on the switch.

**Solution**

Use the correct port on the switch.

**Network Load Balancing is not load balancing applications, and the default host handles all network traffic.**
**Possible Cause**

A port rule is missing. By default, Network Load Balancing directs all incoming network traffic not governed by port rules to the default host. This ensures that the application you do not want load balanced behaves properly.

**Solution**

To load balance an application across the cluster, create a port rule on every cluster host for the TCP/IP port(s) serviced by the application.

**Possible Cause**

You added a second host to a single host cluster, but the second host is not configured correctly. The cluster never converges and the original host continues to handle all of the traffic.

**Solution**

Carefully review and if necessary correct each of the settings on the second host, for example cluster IP address, dedicated IP address and port rules.

**Possible Cause**

If your cluster is configured for unicast mode, a switch might have learned the Network Load Balancing network adapter's MAC address.

**Solution**

Clear the switch's port to MAC address mapping.

**Possible Cause**

A proxy server is sending all connections using a single IP address to your cluster in single affinity mode.

**Solution**

Configure your proxy server to use multiple IP addresses.

**Traffic alternates unexpectedly between the cluster hosts breaking TCP connections.**
**Cause:** Unicast network addresses are causing problems with the switching hub. If you are using a switching hub to interconnect the cluster hosts, you must use Network Load Balancing multicast support; otherwise, the switch is likely to behave erratically when the same unicast network is used on multiple switch ports.

**Solution:** Check that you have selected multicast support in the **Network Load Balancing Properties** dialog box. If you do not want to use multicast support, you can interconnect the cluster hosts with a hub or coaxial cable instead of a switch.

**Network traffic does not appear to load balance evenly among the cluster hosts.**
**Cause:** The network traffic is coming from a limited number of IP addresses, possibly due to the use of a proxy server.

**Solution:** Configure your proxy server to use multiple IP addresses.

**Notes**

- These topics describe several common problems that you might encounter when installing and initially using Network Load Balancing. Each topic describes the likely reasons for each problem and one or more suggested remedies.

- These topics assume that your system and applications meet the minimum requirements for Network Load Balancing. For more information, see Checklist: Enabling and configuring Network Load Balancingand Determining which applications to use with Network Load Balancing.

- Many of the problems described here can be avoided, diagnosed and corrected by using Network Load Balancing Manager. For more information, see Network Load Balancing Manager.

- For additional troubleshooting support, go the Microsoft Product Support Services Web site by clicking on the appropriate link at Support resources and performing a search on "Network Load Balancing", "Nlb", or "Windows Server 2003."

- You should test your network and all network adapters for proper operation before installing Network Load Balancing. Be sure to follow all installation steps and check that the cluster parameters and port rules are identically set for all

cluster hosts. If a problem occurs, always check the Windows event log for a message from the Network Load Balancing driver. For more information, see Cluster parameters, Host parameters, and Port rules in Network Load Balancing parameters.

**When you are using Network Load Balancing with Microsoft Internet Security and Acceleration (ISA) Server, one cluster host might log blocked packets that are directed to the dedicated Internet Protocol (IP) address of another host.**
**Cause:** One of the cluster hosts is configured with a host priority identifier equal to 1.

**Solution:** Do not configure any cluster host with a host priority identifier of 1. Instead, use numbers that are greater than 1. For more information, see Configure host parameters.

**You are unable to create a Network Load Balancing cluster in a 64-bit environment.**
**Cause:** You might not be running the appropriate Network Load Balancing version for the environment you are using. Network Load Balancing cannot form a cluster when the 32-bit version of Network Load Balancing is used on a 64-bit computer. This issue might have gone undetected because 32-bit Network Load Balancing components (nlb.exe, wlbs.exe, and nlbmgr.exe) appear to run correctly in the 64-bit environment.

**Solution:** If you plan to use a 64-bit environment, you must use the 64-bit Network Load Balancing version.

# Server Clusters

Updated: April 10, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server clusters**
A server cluster is a group of independent computer systems, known as nodes, running Microsoft® Windows Server™ 2003, Enterprise Edition or Microsoft® Windows Server™ 2003, Datacenter Edition, and working together as a single system to ensure that critical applications and resources remain available to clients. The nodes in a cluster remain in constant communication through the exchange of periodic messages, called heartbeats. If one of the nodes becomes unavailable as a result of failure or maintenance, another node immediately begins providing service (a process known as failover).

Server clusters can combine up to eight nodes. In addition, a cluster cannot be made up of nodes running both Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition since the different operating systems may be running incompatible versions of the Cluster service. In server clusters with more than two nodes, all nodes must run Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, but not both. However, a server cluster can be operated with some nodes running the Microsoft® Windows® 2000 operating system and others running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.

Server clusters can be set up as one of three different cluster model configurations:

- **Single node server clusters** can be configured with, or without, external cluster storage devices. For single node clusters without an external cluster storage device, the local disk is configured as the cluster storage device.
- **Single quorum device server clusters** have two or more nodes and are configured so that every node is attached to one or more cluster storage devices. The cluster configuration data is stored on a single cluster storage device.
- **Majority node set server clusters** have two or more nodes but the nodes may or may not be attached to one or more cluster storage devices. The cluster configuration data is stored on multiple disks across the cluster and the Cluster service makes sure that this data is kept consistent across the different disks.

It is recommended that you understand the advantages and limitations of the different cluster models before you configure your server cluster. For example, a majority node set cluster can tolerate fewer simultaneous node failures than an equivalent single quorum device cluster.

For more information about the three cluster models, see Choosing a Cluster Model.

For cluster storage, you can use parallel SCSI, Fibre Channel, Serial Attach SCSI (SAS), or iSCSI. For details about which types of storage can be used with a specific Windows Server 2003 operating system, see Server clusters overview.

Server clusters enable users and administrators to access and manage the nodes as a single system rather than as separate computers.

- Before creating a server cluster, see Checklists: Creating Server Clusters.
- Before installing resources for a server cluster, see Checklists: Installing Server Cluster Resources.
- To find features that have been changed in Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, see New ways to do familiar Server Cluster tasks.
- For guidelines on securing server clusters, see Best practices for securing server clusters.
- For tips about using server clusters, see Best practices for configuring and operating server clusters.
- For help with specific tasks, see Server Cluster How To....
- For general background information, see Server Cluster Concepts.
- For problem-solving instructions, see Server Cluster Troubleshooting.

**Note**

- This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

# Checklists: Creating Server Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists: Creating server clusters**


- Checklist: Planning and creating a server cluster
- Checklist: Creating a new group
- Checklist: Creating a new resource
- Checklist: Creating a server cluster virtual server
- Checklist: Creating a server cluster file server
- Checklist: Creating a server cluster print server
- Checklist: Creating a clustered IIS Web or FTP service
- Checklist: Creating a server cluster application server
- Checklist: Validating your clustering system

# Checklist: Planning and creating a server cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Planning and creating a server cluster**

**Caution**

- If you are using a shared storage device, when you power on and start the operating system, it is of vital importance that only one node has access to the cluster disk. Otherwise the cluster disks can become corrupted. To prevent the corruption of the cluster disks, shut down all but one cluster node, or use other techniques (for example, Logical Unit Number or LUN masking, selective presentation, or zoning) to protect the cluster disks, before creating the cluster. Once the Cluster service is running properly on one node, the other nodes can be added and configured simultaneously.

Follow these steps to plan for and create a server cluster. You can use the tutorial for Cluster Installation at the Microsoft Web site to help you create a server cluster. If you are upgrading your server cluster, review the cluster upgrade documentation.

**Planning issues to consider before you create your server cluster**

| Step | Reference |
|---|---|
| Review basic cluster concepts. You will need to understand nodes, failover and failback, groups, and resources before planning your server cluster. | Server Cluster Concepts<br><br>Server clusters overview |
| Review security best practices for server clusters. | Best practices for securing server clusters |
| Review limitations of server clusters. | Limitations of server clusters |
| Choose your cluster model. | Choosing a Cluster Model |
| Choose applications that can run on your server cluster. | Choosing applications to run on a server cluster |
| Organize your cluster applications and resources into groups based on resource dependencies, failover policies, and administrative convenience. | Planning your groups |
| After you choose a cluster model and determine how to group your applications and resources, you are ready to determine the hardware capacity required for each server in the cluster. | Determining server-capacity requirements |
| Plan your network hardware. | Planning your networking hardware<br><br>Best practices for configuring and operating server clusters |
| Plan your system disks.<br><br>It is recommended that the system disks on all your cluster nodes use the same drive letter. | |

| | |
|---|---|
| Plan your cluster disks.<br><br>To guard against the failure of your cluster disks, it is highly recommended that you implement a hardware RAID solution.<br><br>If your cluster disk hardware includes a SCSI bus, it is important to note that the SCSI host adapters must each have a unique SCSI ID (usually the numbers 6 and 7). In addition, the SCSI bus must be properly terminated.<br><br>By default, the New Server Cluster Wizard will automatically select the smallest disk that is larger than 500 MB as the quorum resource. If no disk larger than or equal to 500 MB is available, the New Server Cluster Wizard will select the largest disk that is at least 50 MB as the quorum resource but will generate a warning. For optimal performance of the NTFS file system, make sure that the disk containing the quorum resource is at least 500 megabytes (MB) in size. It is strongly recommended that you install the quorum resource on a separate disk, by itself, so that other applications cannot affect it. Also, ensure that the quorum disk has sufficient free space so that the Cluster service can start and continue to write to the quorum log. | Hardware redundant array of independent disks (RAID)<br><br>The documentation for your storage hardware. |

**Actions to take when creating a server cluster**

| Step | Reference |
|---|---|
| Choose your cluster hardware.<br><br>Microsoft supports only complete cluster systems chosen from the Windows Catalog; it is recommended that you do not mix and match hardware components when building your cluster. Check to ensure that your hardware, including your cluster disks, is compatible with Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. | Support resources<br><br>Best practices for configuring and operating server clusters |
| Verify the availability of domain services and add your servers to the domain. | Choosing a domain model |
| Install and configure your cluster disks.<br><br>Turn on your cluster nodes, but do not allow the operating system to start on any of the nodes. Connect your shared storage buses, completing the initial installation steps in the BIOS or EFI firmware configuration screens for your particular storage device by following the directions in the hardware documentation.<br><br>Allow the operating system to start on one node and configure the cluster disks as basic disks. Format the disks using the NTFS file system. Assign drive letters to or create mounted drives for the cluster disks on the shared bus. If you assign drive letters, the disks must have the same drive letter on all nodes in the cluster.<br><br>**Caution**<br><br>• Cluster disks on a shared bus must be partitioned as Master Boot Record (MBR) and not as GUID Partition Table (GPT) disks. | The documentation for your storage hardware.<br><br>General storage configuration tips for server clusters<br><br>Connect the shared storage buses and assign drive letters |
| Configure network hardware. | Configuring cluster network hardware |
| Configure IP address settings for your network. | Configuring cluster network hardware<br><br>Private network addressing options |

| | |
|---|---|
| Gather the information required for configuring the first node in your cluster.<br><br>You will need the following:<br><br>1. A static IP address for the cluster.<br>2. Computer accounts for the cluster nodes.<br>3. A user account for the Cluster service.<br>4. A name for the cluster, that is, the name administrators will use for connections to the cluster. (The actual applications running on the cluster will typically have different network names.)<br><br>**Caution**<br><br>• Computer accounts for the nodes must be in the same domain. It is strongly recommended that the Cluster service account be in the same domain as well.<br>• It is strongly recommended that you do not use the same account for the Cluster service and applications in the cluster (for example, Microsoft SQL Server). If you do use the same account, you may not be able to later change the Cluster service password without disrupting your cluster applications.<br>• The cluster name must be different from the domain name, from all computer names or virtual server names in the domain, and from other cluster names in the domain. | Domain accounts and the Cluster service<br><br>The Cluster service and TCP/IP |
| Create a cluster on the first node.<br><br>Use Cluster Administrator to create a cluster on the first node.<br><br>**Note**<br><br>• You cannot use the Configure Your Server or Manage Your Server tools to create a cluster or administer cluster nodes.<br><br>You can use the tutorial for Cluster Installation to complete the New Server Cluster Wizard. | Create a cluster<br><br>Microsoft Web site |
| Verify that you can see the cluster disks.<br><br>Use Disk Management to verify that you can see the cluster disks.<br><br>Use Cluster Administrator to verify that the disks are all online. | Disk Management<br><br>Bring a group online |
| Allow the remaining nodes to boot to the operating system then add them to the cluster.<br><br>Use Cluster Administrator to add one or more nodes to the cluster.<br><br>You can use the tutorial for Cluster Installation to complete the New Server Cluster Wizard. | Add additional nodes to the cluster<br><br>Microsoft Web site |
| Validate the cluster disks.<br><br>Use Disk Management to verify that all nodes see the same number of cluster disks.<br><br>Use Cluster Administrator to failover the cluster disks to each node in the cluster and verify that they come back online. | Disk Management<br><br>Test whether group resources can fail over |

| | |
|---|---|
| Configure the cluster network role and type.<br><br>By default, the New Server Cluster Wizard will set up all cluster networks as mixed (public-and-private) networks. Optionally, you can change this so that the networks are set up as private networks (that is, for node-to-node communication only) or public networks (that is, for client access only). | Change how the cluster uses a network |
| Configure the cluster network properties.<br><br>Place the private network at the top of the Network Priority list for internal node-to-node communication in the cluster. | Change network priority for communication between nodes |
| Create your resource groups and resources. | Checklist: Creating a new group<br><br>Checklist: Creating a new resource |
| Validate your clustering system. | Checklist: Validating your clustering system |

# Checklist: Creating a new group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a new group**

| Step | Reference |
|---|---|
| Review resource groups. | Server Cluster groups |
| Determine the name you will assign to the group. This name must be different from any other group name in the cluster. | Not applicable |
| (Optional) Determine the text description for the group. | Not applicable |
| (Optional) Determine the preferred owner for the group. | Not applicable |
| Use the New Group Wizard to create the group. | Create a new group in a cluster |
| Determine the group's failover/failback policy. | Failover and failback<br><br>Determining failover and move policies for groups<br><br>Set group failover policy<br><br>Set group failback policy |

# Checklist: Creating a new resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a new resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Determine the name of your resource. This name must be different from any other resource name in the cluster. | Not applicable |
| Determine the type of your resource. | Resource types |
| Determine the group to which your resource will belong. | Server Cluster groups |
| Determine if the resource will run in a separate resource monitor. | Resource Monitors |
| List the possible owners of the resource. | Failover and failback |
| Determine the resource dependencies. | Resource dependencies |
| Determine the restart policy for the resource. | Specify the restart policy for a resource<br><br>Failover and failback |
| List properties specific to your resource. | Checklists: Installing Server Cluster Resources |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Creating a server cluster virtual server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a server cluster virtual server**

| Step | Reference |
| --- | --- |
| Review the concepts behind virtual servers. | Virtual servers |
| Use the Cluster Application Configuration Wizard to create your virtual server. | Start the Cluster Application Wizard |
| Configure resource dependencies in your virtual server. | Setting resource properties |

# Checklist: Creating a server cluster file server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a server cluster file server**

| Step | Reference |
|---|---|
| Review server cluster components. | Server Cluster Components |
| Review cluster resources. | Server Cluster Resources |
| Determine what type of file share is appropriate for your situation.<br><br>**Note**<br><br>• If you plan to maintain home folders for a large number of users on a server cluster, you do not have to create a file share resource for every home folder (which may slow down your server cluster's performance). Instead you can create a single file share for the root home folder and use the Cluster service to share the subfolders to the individual users. | File Share resource type<br><br>Using a server cluster with large numbers of file shares |
| Plan your resource groups. | Planning your groups |
| Create a new server cluster. | Checklist: Planning and creating a server cluster |
| Create a virtual server for your file server and a file share resource. | Checklist: Creating a new group<br><br>Checklist: Creating a server cluster virtual server<br><br>Checklist: Creating a new resource<br><br>Share subfolders on a File Share resource<br><br>Create a cluster-managed encrypted file share |

# Checklist: Creating a server cluster print server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a server cluster print server**

| Step | Reference |
|---|---|
| Review server cluster components. | Server Cluster Components |
| Review cluster resources. | Server Cluster Resources |
| Review documentation for the Print Spooler resource type. | Print Spooler resource type |
| Plan your resource groups. | Planning your groups |
| If a server cluster does not already exist, create a new server cluster. | Checklist: Planning and creating a server cluster |
| Create a group, a virtual server for your print server, and a printer spooler resource. | Checklist: Creating a new group<br>Checklist: Creating a server cluster virtual server<br>Checklist: Creating a new resource |

# Checklist: Creating a clustered IIS Web or FTP service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a clustered IIS Web or FTP service**

**Important**

- It is recommended that you review the information on Network Load Balancing clusters before using a server cluster to create a high-availability Web server. Depending on your needs, using a Network Load Balancing cluster to load balance the incoming client HTTP traffic might be a better solution. For more information on Network Load Balancing clusters, see Network Load Balancing Clusters.

  For the most current information on server clusters and how they work with Network Load Balancing, see Updated technical information.

| Step | Reference |
|---|---|
| Review server cluster components. | Server Cluster Components |
| Review cluster resources. | Server Cluster Resources |
| Review Internet Information Services (IIS) concepts. | IIS documentation |
| Install the Microsoft Distributed Transaction Coordinator (DTC) resource. | Distributed Transaction Coordinator<br><br>Checklist: Installing a Distributed Transaction Coordinator resource |
| Install Internet Information Services (IIS). | IIS documentation |
| Decide which internet service you will provide: FTP or Web.<br><br>If you are creating an FTP server, install the FTP service and default FTP site.<br><br>**Note**<br><br>• A Web site is installed by default when you first install IIS. | IIS documentation |
| Plan your resource groups. | Planning your groups |
| Create a new cluster. | Create a cluster |
| Add nodes to the cluster. | Add additional nodes to the cluster |
| Create a group and a virtual server for your clustered IIS service. | Checklist: Creating a new group<br><br>Checklist: Creating a server cluster virtual server |

| | |
|---|---|
| Create a Generic Script Resource for the IIS service in the virtual server resource group you created above.<br><br>Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition include 2 example Visual Basic scripts that you can use for this purpose, **clusweb.vbs** and **clusftp.vbs**. By default, these are located in the %systemroot%\system32\inetsrv directory.<br><br>**Important**<br><br>• It is recommended that the Generic Script resource configured for IIS be dependent on both a Network Name resource (so that remote clients can access it) and a Physical Disk (or other storage class device) resource. The Network Name will, in turn, be dependent on the IP Address (that will correspond to the virtual root).<br><br>• If you use a Generic Script resource to host multiple Web sites, it is recommended that you run the same version of IIS on all nodes in the cluster.<br><br>**Note**<br><br>• You can run the clusweb.vbs and clusftp.vbs scripts remotely from a non-Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition computer by pointing to the %systemroot%\system32\inetsrv Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition directory where they are located. | Checklist: Installing a Generic Script resource<br><br>Setting resource properties Modify resource dependencies |
| Configure IIS bindings.<br><br>To configure an FTP site to handle requests sent to a virtual server, ensure that the FTP site is configured to listen to the IP address of the virtual server. You can do this by assigning an IP address to an FTP site.<br><br>To configure a Web site to handle requests sent to a virtual server, ensure that the Web site is configured to listen to the IP address and/or name of the virtual server. The name of the virtual server can be used as the Host Header value for the Web site. You can do this by assigning an IP address and/or Host Header values to a Web site. | IIS documentation |
| Synchronize IIS on all nodes by using the **IIScnfg.vbs /copy** command to copy the IIS configuration to other nodes. | IIS documentation |

# Checklist: Creating a server cluster application server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Creating a server cluster application server**

| Step | Reference |
| --- | --- |
| Review server cluster components. | Server Cluster Components |
| Review cluster resources. | Server Cluster Resources |
| Review the requirements for your application. | Your application documentation<br><br>Checklist: Installing a Generic Application resource |
| Plan your resource groups. | Planning your groups |
| Create a group, a virtual server for your application, and a generic application resource. | Checklist: Creating a new group<br><br>Checklist: Creating a server cluster virtual server<br><br>Checklist: Creating a new resource |

# Checklist: Validating your clustering system

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Validating your clustering system**

| Step | Reference |
|---|---|
| One by one, take each resource group offline then bring it back online.<br><br>**Note**<br><br>• The default Cluster Group (containing the quorum resource) will not go offline. | Take a group offline<br><br>Bring a group online |
| Test that resources can fail over by moving resource groups from one node to another. | Test whether group resources can fail over |
| Test the failover policies for the resource groups. | Test group failover policy |
| Test that the cluster can recover from node failures. | Test node failure |
| Test that the cluster can recover from private network failures. | Test private network failures |
| Review your system logs for any server cluster errors recorded by the Cluster service.<br><br>In Event Viewer, server cluster errors logged by the Cluster service show up under the Source column as "ClusSvc". | Event Viewer |

# Checklists: Installing Server Cluster Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists: Installing server cluster resources**

- Checklist: Installing a DHCP Service resource
- Checklist: Installing a Distributed File System (DFS) root resource
- Checklist: Installing a Distributed Transaction Coordinator resource
- Checklist: Installing a File Share resource
- Checklist: Installing a Generic Application resource
- Checklist: Installing a Generic Script resource
- Checklist: Installing a Generic Service resource
- Checklist: Installing an IP Address resource
- Checklist: Installing a Local Quorum resource
- Checklist: Installing a Majority Node Set resource
- Checklist: Installing a Message Queuing resource
- Checklist: Installing a Network Name resource
- Checklist: Installing a Physical Disk resource
- Checklist: Installing a Print Spooler resource
- Checklist: Installing a Volume Shadow Copy Service Task resource
- Checklist: Installing a WINS Service resource

# Checklist: Installing a DHCP Service resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a DHCP Service resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Review the concepts behind Dynamic Host Configuration Protocol (DHCP). | DHCP |
| Plan common resource settings. | Checklist: Creating a new resource |
| Configure location of DHCP database and backup of database files. | DHCP Service and WINS Service resource types |
| If you already use DHCP, and you want to retain your existing DHCP database files, move the files to the new DHCP database before bringing the new DHCP server online. | Backing up the DHCP database |
| Make certain that you install the DHCP Service on all nodes. | Install a DHCP server |
| Use the New Resource Wizard to create the resource.<br>**Important**<br>• When you create a DHCP Service resource on a node, the Cluster service sets the startup type for the DHCP service on that node to **Manual**. The Cluster service also stops the DHCP service on all the other nodes, but, on those nodes not hosting the DHCP Service resource, the Cluster service leaves the DHCP service startup type set to **Automatic**. This DHCP service startup type will be set to Manual on any of the other nodes only when the DHCP Service resource fails over. In addition, if you subsequently evict a node from the cluster, the DHCP service will not be restarted on that evicted node. | Create a new resource<br>Configure how a service is started |

# Checklist: Installing a Distributed File System (DFS) root resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Distributed File System (DFS) root resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Plan common resource settings. | Checklist: Creating a new resource |
| Review Distributed File System (DFS) information. | Distributed File System (DFS) |
| Using the DFS Microsoft Management Console (MMC) snap-in, make sure that a DFS root does not already exist. | Show a DFS root |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Installing a Distributed Transaction Coordinator resource

Updated: January 20, 2011

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Distributed Transaction Coordinator resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Plan common resource settings. | Checklist: Creating a new resource |
| Create or select the group you want to install the Distributed Transaction Coordinator (DTC) resource into.<br><br>By default, the DTC resource will install into the first group it locates that has an IP, Name, and Physical Disk resource, unless you identify a specific group. | Checklist: Creating a new group<br><br>Distributed Transaction Coordinator |
| Use the New Resource Wizard to create the resource.<br><br>Make the Distributed Transaction Coordinator resource dependent on the Physical Disk resource. | Create a new resource |

| **Important** |
|---|
| In Windows Server 2003 Service Pack 1 (SP1) and later service packs, DTC provides the administrator with additional settings that control the level of authentication required between computers involved in transactions. However, there are limitations in the way you can use these settings in a clustered environment. You cannot use the **Mutual Authentication Required** transaction mode with computers that are in a clustered environment, or any computers that are negotiating transactions with such computers. In that context, you can use the **Incoming Caller Authentication Required** transaction mode instead. For more information, see Distributed Transaction Coordinator (http://go.microsoft.com/fwlink/?LinkId=209176) and New functionality in the Distributed Transaction Coordinator service in Windows Server 2003 Service Pack1 and in Windows XP Service Pack 2 (http://go.microsoft.com/fwlink/?LinkId=209177). |

# Checklist: Installing a File Share resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a File Share resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies.<br><br>**Important**<br><br>• It is recommended that the File Share resource be dependent on a Physical Disk and a Network Name resource. | Server Cluster groups |
| Review the behavior of the file share resource type. | File Share resource type |
| Plan common resource settings. | Checklist: Creating a new resource |
| Determine whether you will use the advanced properties of the File Share resource type. | File Share resource type |
| Use the New Resource Wizard to create the resource. | Create a new resource |
| Configure file share permissions using Cluster Administrator. | Using Cluster Administrator<br><br>Start Cluster Administrator in Microsoft Management Console |

# Checklist: Installing a Generic Application resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Generic Application resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Plan common resource settings. | Checklist: Creating a new resource |
| Install your application on all nodes.<br>**Note**<br>• Before attempting to cluster your application, contact the application vendor for specific installation instructions or recommendations on configuring it optimally on your server cluster. | Generic Application resource type |
| Note the command line for your application. | The documentation for your application. |
| Determine the directory in which your application needs to run. | The documentation for your application. |
| Determine whether your application can display windows on the node on which it runs. | The documentation for your application. |

| | |
|---|---|
| Determine which application registry keys need to be replicated.<br><br>**Information on registry key replication**<br><br>Server clusters provides the ability for Generic Service and Generic Application resource types to specify registry keys below HKEY_LOCAL_MACHINE that need be replicated to all nodes in the cluster.<br><br>Some services or applications might update registry information while they are running by use of Win32 registry functions. The Cluster service provides registry checkpointing for noncluster-aware applications that may use these functions. Registry checkpointing for these resources occurs under the following conditions:<br><br>• When you specify a new registry key for the resource, the specified key is checkpointed.<br>• When the resource goes online, the registry keys are updated with the previously checkpointed information.<br>• When the resource is brought offline, all the checkpoints associated with this resource are saved.<br>• When the resource is online and changes are made to the registry key which is registered with the cluster server for replication, the Cluster service makes sure the updates are written to a checkpoint maintained on the quorum device.<br><br>If you manually update these registry keys while the application or service is offline, the changes may not be replicated or may be lost. To prevent this from happening, make any manual changes while the service or application resource is online.<br><br>Cluster aware applications need to make use of cluster registry functions as provided in the cluster API set when working with registry keys under HKEY_LOCAL_MACHINE\Cluster. Refer to the Windows Clustering documentation in the Platform Software Development Kit (SDK) for more information on this topic. | The documentation for your application. |
| Determine whether a dependent network name needs to be substituted for the computer name.<br><br>Applications can make application program interface (API) calls to determine the name of the computer on which they are running. Making the generic application dependent on a network name resource causes the **Use Network Name** check box in **Cluster Administrator** to be selected. In such a case, the network name is returned instead of the computer name when the application makes the API call. This enables applications that are not cluster aware to continue to appear to be running on the same computer even after they move to a different cluster node. | The documentation for your application.<br><br>For more information on API calls for server clusters, see the Microsoft Platform Software Development Kit (SDK). |
| Determine whether any application-specific registry keys need to be replicated on all cluster nodes.<br><br>As described previously, this pertains to applications that are not cluster-aware. An application can store data in the registry. If the application is moved but the registry data is not moved with it, the application's state is lost. If the cluster administrator specifies the root of the registry key, it enables the Cluster service to "watch" that key. The Cluster service captures the data in the key and any changes that occur. When the generic application resource is moved, the Cluster service moves the registry data to the new resource host and then brings the generic application resource online. | The documentation for your application. |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Installing a Generic Script resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Generic Script resource**

| Step | Reference |
| --- | --- |
| Review server cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Plan common resource settings. | Checklist: Creating a new resource |
| Note the name of the generic script along with the path to the script file. | For more information on developing generic scripts, see the Microsoft Platform Software Development Kit (SDK). |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Installing a Generic Service resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Generic Service resource**

| Step | Reference |
|---|---|
| Review server cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Plan common resource settings. | Checklist: Creating a new resource |
| Note the exact service name as it appears in the registry under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services**. | The documentation for your service. |
| Note the startup parameters of your service. | The documentation for your service. |
| Ensure that the service is installed on all nodes. | Not applicable |
| Determine which service registry keys need to be replicated on all nodes.<br><br>**Information on registry key replication**<br><br>Server clusters provides the ability for Generic Service and Generic Application resource types to specify registry keys below HKEY_LOCAL_MACHINE that need to be replicated to all nodes in the cluster.<br><br>Some services or applications may update registry information while they are running by use of Win32 registry functions. The Cluster service provides registry checkpointing for noncluster-aware applications that may use these functions. Registry checkpointing for these resources occurs under the following conditions:<br><br>• When you specify a new registry key for the resource, the specified key is checkpointed.<br>• When the resource goes online, the registry keys are updated with the previously checkpointed information.<br>• When the resource is brought offline, all the checkpoints associated with this resource are saved.<br>• When the resource is online, and changes are made to the registry key that is registered with the cluster server for replication, the Cluster service ensures that the updates are written to a checkpoint maintained on the quorum device.<br><br>If you manually update these registry keys while the application or service is offline, the changes may not be replicated or may be lost. To prevent this from happening, make any manual changes while the service or application resource is online.<br><br>Cluster-aware applications need to make use of cluster registry functions as provided in the cluster application program interface (API) set when working with Cluster registry keys under HKEY_LOCAL_MACHINE. Refer to the Windows Clustering documentation in the Platform Software Development Kit (SDK) for more information on this topic. | The documentation for your service.<br><br>Checklist: Installing a Generic Application resource |

| | |
|---|---|
| Determine whether a dependent network name needs to be substituted for the computer name.<br><br>Applications can make application program interface (API) calls to determine the name of the computer on which they are running. Making the generic application dependent on a network name resource causes the **Use Network Name** check box in **Cluster Administrator** to be selected. In such a case, the network name is returned instead of the computer name when the application makes the API call. Thus, applications that are not cluster aware continue to appear to be running on the same computer even after they move to a different cluster node. | The documentation for your application. |
| Determine whether any application-specific registry keys need to be replicated on all cluster nodes.<br><br>This pertains to applications that are not cluster aware. An application may store data in the registry. If the application is moved, but the registry data is not moved with it, the application's state is lost. If the cluster administrator specifies the root of the registry key, it enables the Cluster service to "watch" that key. The Cluster service captures the data in the key and any changes that occur. When the generic application resource is moved, the Cluster service moves the registry data to the new resource host and then brings the generic application resource online. | The documentation for your application. |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Installing an IP Address resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing an IP Address resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Review TCP/IP concepts. | TCP/IP Concepts |
| Plan common resource settings. | Checklist: Creating a new resource |
| Obtain an Internet Protocol (IP) address and subnet mask from your network administrator.<br><br>This address must be a static IP address not already in use on the network. Make sure that the subnet mask is the same subnet mask as the one for the associated cluster network. | IP Addressing and Routing |
| Determine whether or not NetBIOS needs to be enabled for this IP address.<br><br>If NetBIOS for the IP address is enabled, it allows a dependent network name to be published using NetBIOS. If it is not enabled, the dependent network name will only be published using DNS. | Disable WINS/NetBT name resolution<br><br>DNS |
| Determine the network on which to create this resource.<br><br>**Important**<br><br>• It is recommended that you create IP Address resources on a network accessible by clients, that is, a public or mixed network. | Server cluster networks |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Installing a Local Quorum resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Local Quorum resource**

| Step | Reference |
|---|---|
| Review server cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Review the behavior of the Local Quorum resource type.<br><br>**Note**<br><br>• The Local Quorum resource can be used as the quorum resource only for a single-node cluster. | Local Quorum resource type |
| Plan common resource settings. | Checklist: Creating a new resource |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Installing a Majority Node Set resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Majority Node Set resource**

| Step | Reference |
|---|---|
| Review server cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Review the behavior of the Majority Node Set resource type and the cluster configurations that use this resource type.<br><br>**Note**<br><br>• Only one Majority Node Set resource can exist in a server cluster. | Majority Node Set resource type<br><br>Model 3: Majority node set server cluster configuration |
| Plan common resource settings. | Checklist: Creating a new resource |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Installing a Message Queuing resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Message Queuing resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Review Message Queuing concepts. | Microsoft Message Queuing |
| Plan common resource settings. | Checklist: Creating a new resource |
| Make sure that Microsoft Distributed Transaction Coordinator (DTC) is already installed. | Distributed Transaction Coordinator<br><br>Checklist: Installing a Distributed Transaction Coordinator resource |
| Create a Physical Disk or other storage class device resource. | Checklist: Installing a Physical Disk resource |
| Create a Network Name resource. | Checklist: Installing a Network Name resource |
| Enable Kerberos authentication for that Network Name resource. | Enable Kerberos authentication for virtual servers |

| | |
|---|---|
| Use the New Resource Wizard to create the Message Queuing resource.<br><br>Make the Message Queuing resource dependent on the Network Name and Physical Disk resources created earlier.<br><br>**Caution**<br><br>• You can have multiple Message Queuing resources in a server cluster. However, you can only have one Message Queuing resource per resource group.<br><br>• If you do have multiple Message Queuing resources in a server cluster, message activity may not function as expected unless you increase the system view space memory pool size of each node by 4 MB for each instance of the Message Queuing resource. For more information, see the troubleshooting topic "The Message Queuing resource fails to handle message activity correctly which may result in resource failures" in Group and resource failure problems.<br><br>**Notes**<br><br>• The Message Queuing resource does not have to be located in the same resource group as the group containing the DTC resource.<br><br>• The Message Queuing resource can come online even if the DTC resource is not online provided that there are no transactions to recover. If the DTC resource is offline, transactions are disabled; once the DTC resource comes back online, transactions are enabled. | Create a new resource |

# Checklist: Installing a Network Name resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Network Name resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Review networking concepts. | Network Services |
| Plan common resource settings. | Checklist: Creating a new resource |
| Choose an Internet Protocol (IP) address resource to be associated with the network name.<br><br>If an IP address resource does not already exist, create one. If necessary, obtain the IP address from your network administrator. This must be a static IP address.<br><br>When the Network Name resource is brought online, the IP address resource with which it is associated is included when the name is registered with the various name resolution services. | Create a new resource<br><br>IP Addressing and Routing<br><br>Disable WINS/NetBT name resolution<br><br>DNS |
| Choose the name you want to publish on the network, following any naming conventions that might be in effect. | Not applicable |
| Use the New Resource Wizard to create the network name resource.<br><br>Make the network name resource dependent on the IP address resource.<br><br>**Note**<br><br>• Optionally, you can enable Kerberos authentication and specify that DNS registration must succeed for the network name resource. | Create a new resource<br><br>Modify resource dependencies<br><br>Virtual servers |
| Optionally, use the **cluster** command with the **/registerdns** option to register the resource with a DNS server in a way that does not interrupt cluster availability. | Cluster |

# Checklist: Installing a Physical Disk resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Physical Disk resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Review disk storage concepts. | Disk Management |
| Plan common resource settings. | Checklist: Creating a new resource |
| Make sure that the disk has a signature. | Disk Management |
| Make sure that the disk is configured as basic and not dynamic. | Disk Management |
| Make sure that all partitions on the disk are formatted using the NTFS file system. | Disk Management |
| Make sure that all partitions have a drive letter assigned or is set up as a mounted drive.<br>**Important**<br>• If you use drive letters, assign drive labels that match the drive letters and make sure that you assign the same drive letter to each node in the cluster. | Disk Management<br>Install local storage buses and devices |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Installing a Print Spooler resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Print Spooler resource**

| Step | Reference |
| --- | --- |
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups<br><br>Setting resource properties |
| Review network printing concepts. | Printing |
| Review print spooler resource type information. | Print Spooler resource type |
| Make sure that your printer can be used with TCP/IP. | Create a cluster-managed printer |
| Plan common resource settings. | Checklist: Creating a new resource |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# Checklist: Installing a Volume Shadow Copy Service Task resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a Volume Shadow Copy Service Task resource**

| Step | Reference |
|---|---|
| Review the general concepts for cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Review the behavior of the Volume Shadow Copy Service Task resource type. | Volume Shadow Copy Service resource type |
| Plan common resource settings. | Checklist: Creating a new resource |
| Create or select the resource group into which you want to install the Volume Shadow Copy Service Task resource. | Checklist: Creating a new group |
| Use the New Resource Wizard to create the resource.<br><br>If necessary, make the Volume Shadow Copy Service Task resource dependent on other resources in the resource group.<br><br>**Caution**<br><br>• If the name you provide for your Volume Shadow Copy Service Task resource is the same as an existing scheduled task, that scheduled task will be deleted when the Volume Shadow Copy Service Task resource is brought online.<br><br>**Important**<br><br>• If you want to create a Volume Shadow Copy Service Task resource to manage shadow copies on a cluster node, use this procedure instead: Enable Shadow Copies for Shared Folders in a cluster. Using Cluster Administrator or cluster.exe to completely configure and manage shadow copies in a cluster is not a supported scenario.<br><br>• The scheduled task managed by the Volume Shadow Copy Service Task resource runs under the Local System account with a default working directory of *systemroot*\system32. | Create a new resource<br><br>Setting resource properties |

# Checklist: Installing a WINS Service resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installing a WINS Service resource**

| Step | Reference |
|---|---|
| Review the concepts behind cluster resources. | Server Cluster Resources |
| Review resource groups and resource dependencies. | Server Cluster groups |
| Review Windows Internet Naming Resource (WINS) concepts. | WINS |
| Plan common resource settings. | Checklist: Creating a new resource |
| Configure location of WINS database and backup of database files. | DHCP Service and WINS Service resource types |
| If you already use WINS, and you want to retain your existing WINS database files, move the files to the new WINS database before bringing the new WINS server online. | The WINS database<br><br>Back up the WINS database |
| Make certain that you install the WINS Service resource on all nodes. | Not applicable |
| Use the New Resource Wizard to create the resource. | Create a new resource |

# New ways to do familiar Server Cluster tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

The following table lists common tasks for configuring server clusters in Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. The user interface for performing these tasks is different from the way it was in Windows 2000.

| If you want to | In Windows 2000 use | In Windows Server 2003 family use |
|---|---|---|
| Create a new cluster. | **Add or Remove Programs** to install the Cluster service then use Cluster Administrator to add the first node to the cluster. | Cluster Administrator and select the **Create new cluster** option. For more information, see Checklist: Planning and creating a server cluster and Create a cluster. |
| Add nodes to a cluster. | **Add or Remove Programs** to install the Cluster service on each server to be added to the cluster then use Cluster Administrator locally to add each node, one by one, to the cluster. | Cluster Administrator and select the **Add nodes to cluster** option to add one or more nodes all at once. You can also add nodes locally or remotely. For more information, see Add additional nodes to the cluster. |
| Restore the cluster quorum for all nodes in the cluster. | Cluster Administrator to stop the Cluster service on all nodes. Then use Backup to restore the System State data, and finally the Clusrest.exe tool to copy the restored data to the quorum disk. | Backup to restore the System State data and select the advanced restore **When restoring the Cluster Quorum, mark the restored data as Authoritative** option. For more information, see Restore the cluster database on a local node. |
| Uninstall the Cluster service. | **Add or Remove Programs** to uninstall the Cluster service on each node. | Cluster Administrator to evict the last node in the cluster. |
| Create a Web or FTP server. | Cluster Administrator to create an IIS instance resource. Then configure the IIS bindings and use the **IISSYNC** tool to replicate the IIS configuration to all nodes. | Cluster Administrator to create a Generic Script resource for running the IIS service. Then configure the IIS bindings and use the IIS Backup/Restore feature to replicate the IIS configuration to all nodes. For more information, see Checklist: Creating a clustered IIS Web or FTP service. |
| Change network roles, names, and priorities. | The New Server Cluster Wizard and change these network properties during setup. | Cluster Administrator after setup to change the properties of the cluster networks. For more information, see Change how the cluster uses a network, Change the name or description of a network, and Change network priority for communication between nodes. |
| Delete a resource. | Cluster Administrator to first take the resource offline before deleting it. | Cluster Administrator to delete it. For more information, see Delete a resource. |
| Change the password for the Cluster service account. | Cluster Administrator to first stop the Cluster service on all nodes. Then use Active Directory Users and Computers to change the password for the account before restarting the Cluster service and bringing the nodes back online. | Cluster.exe commands to dynamically change the password on the domain controller and the cluster nodes without bringing down the entire cluster. For more information, see Change the Cluster service account password. |

# Server Cluster Best Practices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices**

- Best practices for configuring and operating server clusters
- Best practices for securing server clusters

# Best practices for configuring and operating server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for configuring and operating server clusters**
The following guidelines will help you effectively use a server cluster:

**Secure your server cluster.**
To prevent your server cluster from being adversely affected by denial of service attacks, data tampering, and other malicious attacks, it is highly recommended that you plan for and implement the security measures detailed in Best practices for securing server clusters.

**Check that your server cluster hardware is listed in the Windows Catalog.**
For Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, Microsoft supports only complete server cluster systems chosen from the Windows Catalog. To see if your system or hardware components, including your cluster disks, are compatible, see Support resources. For a geographically dispersed cluster, both the hardware and software configuration must be certified and listed in the Windows Catalog.

The network interface controllers (NICs) used in certified cluster configurations must be selected from the Windows Catalog.

It is recommended that your cluster configuration consist of identical storage hardware on all cluster nodes to simplify configuration and eliminate potential compatibility problems.

**Partition and format disks before adding the first node to your cluster.**
Partition and format all disks on the cluster storage device before adding the first node to your cluster. You must format the disk that will be the quorum resource. All partitions on the cluster storage device must be formatted with NTFS (they can be either compressed or uncompressed), and all partitions on one disk are managed as one resource and move as a unit between nodes.

**Important**

- Cluster disks on the cluster storage device must be partitioned as master boot record (MBR) and not as GUID partition table (GPT) disks.

**Correctly set up your server cluster's networks.**
Follow the guidelines below to reduce network problems in your server cluster:

- Use identical network adapters in all cluster nodes, that is, make sure each adapter is the same make, model, and firmware version.

- Use at least two interconnects. Although a server cluster can function with only one interconnect, at least two interconnects are necessary to eliminate a single point of failure and are required for the verification of original equipment manufacturer (OEM) clusters.

- Reserve one network exclusively for internal node-to-node communication (the private network). Do not use teaming network adapters on the private networks.

- Set the order of the network adapter binding as follows:

  1. External public network

  2. Internal private network (Heartbeat)

  3. [Remote Access Connections]

  For more information, see Modify the protocol bindings order.

- Manually set the speed and duplex mode for multiple speed adapters to the same values and settings. If the adapters are connected to a switch, ensure that the port settings of the switch match those of the adapters. For more information, see Change network adapter settings.

- Use static IP addresses for each network adapter on each node.

- For private networks, define the TCP/IP properties for static IP addresses following the guidelines at Private network addressing options. That is, specify a class A, B, or C private address.

- Do not configure a default gateway or DNS or WINS server on the private network adapters. Also, do not configure private network adapters to use name resolution servers on the public network; otherwise, a name resolution server on the public network might map a name to an IP address on the private network. If a client then received that IP address from the name resolution server, it may fail to reach the address because no route from the client to the private network address exists.

- Configure WINS and/or DNS servers on the public network adapters. If Network Name resources are used on the public networks, set up the DNS servers to support dynamic updates; otherwise, the Network Name resources may not fail over correctly. For more information, see Configure TCP/IP settings.

- Configure a default gateway on the public network adapters. If there are multiple public networks in the cluster, configure a default gateway on only one of these. For more information, see Configure TCP/IP settings.

- Clearly identify each network by changing the default name. For example, you could change the name of the private network connection from the default *Local Area Connection* to *Private Cluster Network*.

- Change the role of the private network from the default setting of **All communications (mixed network)** to **Internal cluster communications only (private network)** and verify that each public network is set to **All communications (mixed network)**. For more information, see Change how the cluster uses a network.

- Place the private network at the top of the Network Priority list for internal node-to-node communication in the cluster. For more information, see Change network priority for communication between nodes.

**Do not install applications into the default Cluster Group.**

**Do not delete or rename the default Cluster Group or remove any resources from that resource group.**
The default Cluster Group contains the settings for the cluster and some typical resources that provide generic information and failover policies. This group is essential for connectivity to the cluster. It is therefore very important to keep application resources out of the default Cluster Group and so prevent clients from connecting to the Cluster Group's IP address and network name resources. If a resource for an application is added to this group and the resource fails, it may cause the cluster group to fail also, therefore reducing the overall availability of the entire cluster. It is highly recommended that you create separate resource groups for application resources.

For more information, see Planning your groups and Checklist: Planning and creating a server cluster.

**Back up your server cluster.**
To be able to effectively restore your server cluster in the event of application data or quorum loss, or individual node or complete cluster failure, follow these steps when preparing backups:

1. Perform an Automated System Recovery (ASR) backup on each node in the cluster.

2. Back up the cluster disks from each node.

3. Back up each individual application (for example, Microsoft Exchange Server or Microsoft SQL Server) running on the nodes.

### Note

- By default, Backup Operators do not have the user rights necessary to create an Automated System Recovery (ASR)

backup on a cluster node. However, Backup Operators can perform this procedure if that group is added to the security descriptor for the Cluster service. You can do that using Cluster Administrator or cluster.exe. For more information, see Give a user permissions to administer a cluster and Cluster.

For more information, see Backing up and restoring server clusters. For more information on backing up applications in a cluster, see the documentation for that application.

**Maintain a backup of the RAID controller.**
In a single quorum device server cluster, the RAID controller is a single point of failure. Always maintain a backup of the RAID controller configuration in case the RAID controller is replaced.

**Do not use APM/ACPI Power saving features.**
APM/ACPI Power saving features must not be enabled on server cluster members. A cluster member that turns off disk drives or enters "system standby" or "hibernate" mode can initiate a failure in the cluster. If multiple cluster nodes have power saving enabled, this can result in the entire cluster becoming unavailable.

Cluster members must use any power scheme that sets the **Turn off hard disks** option to **Never**, for example, the **Always On** power scheme. For more information on choosing a power scheme (located under **Power Options** in Control Panel), see Choose a power scheme.

For cluster nodes without Terminal Services installed, see Configure the Always On power scheme without Terminal Services installed.

For cluster nodes with Terminal Services installed, see Configure the Always On power scheme with Terminal Services installed.

**Note**

- Installing Terminal Services on a system reduces the power management options available to the user. The **System standby** and **System hibernates** options are not available.

**Give the Cluster service account full rights to administer computer objects if Kerberos authentication is enabled for virtual servers.**
If you enable Kerberos authentication for a virtual server's Network Name resource, the Cluster service account does not need full access rights to the computer object associated with that Network Name resource. The Cluster service can use the default access rights given to members of the authenticated users group, but certain operations (for example, renaming the computer object) will be restricted. It is recommended that you work with your domain administrator to set up appropriate administration rights and permissions for the Cluster service account.

For more information, see information about Kerberos authentication in Virtual servers.

**Do not install scripts used by Generic Script resources on cluster disks.**
It is recommended that you install script files used by Generic Script resources on local disks, not on cluster disks. Incorrectly written script files can cause the cluster to stop responding. Installing the script files on a local disk makes it easier to recover from this scenario. For guidelines on writing scripts for the Generic Script resource, see the *Microsoft Platform Software Development Kit* (SDK). For information on troubleshooting Generic Script resource issues, see article Q811685, "A Server Cluster with a Generic Script Resource Stops Responding" in the Microsoft Knowledge Base.

# Best practices for securing server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Best practices for securing server clusters**
Use the following guidelines to help keep your server cluster from being adversely affected by denial of service attacks, data tampering, and other malicious attacks:

**Apply the latest software updates.**
Apply the latest software updates, including security updates and service packs. Windows Server 2003 Service Pack 1 (SP1) contains a large number of improvements to security. We highly recommend that the version of Windows Server 2003 that you run is no earlier than Windows Server 2003 with SP1.

**Obtain the latest information about security.**
Review the Microsoft Web site from time to time for information about security. For example, review one or more of the following:

- Information about Server Security on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkID=59992).
- Knowledge Base article 891597, "How to apply more restrictive security settings on a Windows Server 2003-based cluster server" on the Microsoft Support Web site (http://go.microsoft.com/fwlink/?LinkID=59993).

**Restrict physical access to the server cluster and related infrastructure components to trusted personnel.**
Only allow authorized personnel to have physical access to the server cluster nodes and related infrastructure components (for example, the network and storage components). For more information about standard best practices for securing servers, see Best practices for security.

**Secure your cluster networks.**
**For public or mixed networks**: Ensure that all servers (for example, infrastructure servers such as DNS or WINS servers) that are attached to the cluster subnet (or any other subnet) are secure and trusted. Also, use a firewall to protect your cluster from unauthorized client access.

**For private networks**: Physically disjoint the private network(s) from other networks. Specifically, do not use a router, switch, or bridge to join a private cluster network to any other network. Do not include other network infrastructure or application servers on the private network subnet.

If heartbeat messages that are carried by private or mixed networks are disrupted (for example, by unauthorized users flooding the cluster networks), the cluster's reaction to these false failures may affect not only groups with IP address resources, but also the health of the cluster nodes. Repeated failover of critical cluster resources (for example, resources associated with Microsoft Exchange Server) may cause clients to lose access to the cluster.

**If you use Windows Firewall, use the Security Configuration Wizard to configure it.**
The Security Configuration Wizard simplifies the process of configuring Windows Firewall for cluster nodes. For more information, see Using Windows Firewall with a server cluster.

**Administer cluster nodes remotely only from trusted, secure computers.**
If the remote computers are compromised, untrustworthy or malicious code can be started unknowingly by the cluster administration tools and adversely affect the cluster itself.

**Make the Cluster service account a member of the local Administrators group on all nodes, but do not make it a member of the domain administrator group.**
By giving the minimal possible administration rights and permissions to the Cluster service account, you avoid potential security issues if that account is compromised. For more information about the administration rights and permissions required by the Cluster service account, see Change the account under which the Cluster service runs.

**Take steps to protect the Cluster service account, and do not use the Cluster service account to administer the cluster.**
The Cluster service account has elevated administration rights and permissions. Limit the use of the account, restrict the number of people who know the account name and password, and from time to time, change the password.

For more information about using another user account to administer a cluster, see Give a user permissions to administer a cluster.

For detailed information about the rights and permissions required for the Cluster service account, see Change the account under which the Cluster service runs.

For information about changing the password of the Cluster service account, which you can do without any downtime, see Change the Cluster service account password.

**Use different accounts for the Cluster service and applications in the cluster.**
This will limit exposure to the application and not the cluster account if an application account is compromised. In addition, if you use cluster.exe to change the password for the Cluster service account and one or more applications also use that account, your cluster applications may not function correctly. For more information, see Change the Cluster service account password.

**Use different Cluster service accounts for multiple clusters.**
This will limit exposure to one cluster only if a Cluster service account is compromised.

**Remove the Cluster service account after evicting all the nodes in your cluster.**
The cluster administration tools will not automatically delete the Cluster service account when all nodes have been evicted from the cluster. The Cluster service account has a high level of administration rights and permissions, and that can present potential security issues if it is compromised; it is highly recommended that you remove this account from the local Administrators group if you no longer have use for it. However, because the administrative rights and permissions for the Cluster service account are granted locally on each cluster node and not domain wide, the impact of a compromised account is limited to just the cluster nodes. For more information about deleting user accounts, see Delete a local user account.

**For clusters in an Active Directory domain, enable Kerberos authentication for Network Name resources.**
Kerberos authentication is much more secure than the alternative, NTLM authentication. Note that when you enable Kerberos authentication, you must add certain rights and permissions to the account that the Cluster service creates for the Network Name resource, and possibly to the Cluster service account itself. For more information, see Knowledge Base article 307532, "How to troubleshoot the Cluster service account when it modifies computer objects," on the Microsoft Support Web site (http://go.microsoft.com/fwlink/?LinkId=59994).

**Audit security-related events in the cluster.**
To ensure that only trusted personnel have user rights and permissions to administer the cluster and to track additions of unauthorized user accounts, audit changes to the local Administrators group. Note that, by default, when you create a server cluster or add nodes, the Cluster service account is added to the local Administrators group on each node. For more information about auditing security events, see Auditing Security Events.

**Limit and audit access to shared data (for example, files and folders on cluster disks).**
If you want to audit access to shared data, enable auditing on all cluster nodes. For more information, see Securing shared data in a cluster.

**Limit client access to cluster resources.**
Use Windows Server 2003 family security features to control client access to cluster resources as described in Limiting client access to cluster resources. Note that when you create user and group accounts through which you control access to cluster resources, you must use domain-level accounts, not local accounts, so that appropriate access is available regardless of which node currently owns a clustered resource.

**Limit access to the quorum disk and ensure that the quorum disk always has sufficient free space.**
To help reduce the risk of unauthorized reads and writes to the quorum disk, it is highly recommended that you give access to the quorum disk only to the Cluster service account and members of the local Administrators group. For the Cluster service to start and continue to write to the quorum log, the quorum disk must have sufficient free space. It is recommended that the quorum disk be at least 500 megabytes (MB) in size. For more information, see Checklist: Planning and creating a server cluster and Disk resource security.

**Ensure that all applications running in the cluster are from a trusted source and that access to the cluster disks is restricted only to applications that are managed as a cluster resource.**
For more information, see Disk resource security.

**Secure script files called by Generic Script resources.**
Use NTFS file-level security for execute permissions on script files and permissions for APIs called in those scripts. For more information, see Limiting client access to cluster resources.

**Ensure that applications called by Generic Application resources are from a trusted source and that files, registry checkpoints, and other resources needed for those applications are in a secure location.**
Applications launched as a Generic Application resource will run under the context of the Cluster service account, with its elevated administration rights and permissions. Therefore, ensure that such applications are from a trusted source. In addition, we recommend that in Cluster Administrator, when configuring the parameters for a Generic Application resource, you do not

select **Allow application to interact with desktop** unless it is necessary. For more information, see Limiting client access to cluster resources.

**Secure the cluster configuration log files created when remotely administering a cluster.**
After using the New Server Cluster Wizard and the Add Nodes Wizard from a remote computer, the cluster configuration log file generated by those wizards is saved to the remote computer (at %systemroot%\system32\LogFiles\Cluster\ClCfgSrv.log). This log file contains important information about the cluster. Restrict access to that file to cluster administrators and the Cluster service account. For more information, see Set, view, change, or remove permissions on files and folders.

| **Note** |
| --- |
| If you do not have administrative rights and permissions on the node on which you are using the New Server Cluster Wizard or Add Nodes Wizard, the log file will be written to the local %Temp% directory. |

**Do not copy files to the local cluster directories in a majority node set cluster.**
When using the majority node set model, each node maintains a copy of the quorum database in the cluster directories located at %systemroot%\Cluster\MNS.%ResourceGUID%$\%ResourceGUID%$\MSCS. If you put files in the directories below the \Cluster directory, and then delete the Majority Node Set resource, those files will be deleted by the Cluster service.

**Do not change the default security settings on the HKEY_LOCAL_MACHINE system registry subtree.**
By default, only members of the Administrators group in the Builtin folder and the local system account have Full Control of the HKEY_LOCAL_MACHINE system registry subtree. If this registry subtree is compromised, some resources (for example, the Generic Script resource) may fail to start. For more information about securing the system registry, see Maintain Registry Security.

For more information about security in a server cluster, see Managing Security in a Cluster.

# Server Cluster How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**


- Create a Cluster and Manage Nodes
- Create and Administer File and Print Server Clusters
- Manage Accounts and Domains in a Cluster
- Manage the Cluster Configuration and Diagnostic Logs
- Perform Advanced Administrative Tasks
- Manage Server Cluster Groups
- Manage Resources
- Use Cluster Administrator
- Manage Networks
- Manage Cluster Hardware
- Back Up and Restore Server Clusters
- Test Cluster Failures and Failover Policies

# Create a Cluster and Manage Nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create a cluster and manage nodes**

- Create a cluster
- Add additional nodes to the cluster
- Evict a node from the cluster
- Repair a node without evicting it from the cluster
- Pause a node
- Resume a node
- Stop the Cluster service
- Start the Cluster service
- Set the Cluster service to only start manually

# Create a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a cluster**

1. Open Cluster Administrator.
2. In the **Open Connection to Cluster** dialog box that appears, in **Action**, select **Create new cluster**, then click **OK**.
3. The New Server Cluster Wizard appears. Click **Next** to continue.

   **Important**

   - During the cluster creation process (using the **Quorum** button on the **Proposed Cluster Configuration** page) you will be able to select a quorum resource type (that is, a Local Quorum resource, Physical Disk or other storage class device resource, or Majority Node Set resource). For information on how these quorum resource types relate to the different cluster models, see "Choosing a cluster model" in Related Topics.

4. Upon completion of the New Server Cluster Wizard, click **Finish**.

**Caution**

- If you are using a shared storage device, when you power on and start the operating system, it is of vital importance that only one node has access to the cluster disk. Otherwise the cluster disks can become corrupted. To prevent the corruption of the cluster disks, shut down all but one cluster node, or use other techniques (for example, Logical Unit Number or LUN masking, selective presentation, or zoning) to protect the cluster disks, before creating the cluster. Once the Cluster service is running properly on one node, the other nodes can be added and configured simultaneously.
- When creating a server cluster, all nodes in the cluster must have the same system default locale. That is, the language and country or region selected during the installation of the operating system on each node must be the same. For more information on changing the system default locale, see "To change number, currency, time, and date settings" in Related Topics.

**Important**

- When you create a cluster, Physical Disk resources will be automatically created for cluster disks that use drive letters. However, if your cluster disks use mount points, Physical Disk resources will not be automatically created for those disks; you must manually create disk resources for cluster disks that use mount points. For more information on how to do this, see "Checklist: Installing a Physical Disk resource" in Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- You can also open the New Server Cluster Wizard from the **File** menu, by clicking **New**, then clicking **Cluster**.
- You can safely cancel and restart the installation process at any time.

- A majority node set server cluster is best used in controlled, targeted scenarios, as part of a cluster solution offered by your Original Equipment Manufacturer (OEM), Independent Software Vendor (ISV), or Independent Hardware Vendor (IHV). The single quorum device server cluster, the standard cluster model available in Windows NT and Windows 2000, will still cover the vast majority of your cluster deployment needs.

- Of the default resource types, only the Physical Disk, Local Quorum, or Majority Node Set resource can be a quorum resource. However, third party vendors can supply other storage class resource types that are quorum capable.

  If you want to change the disk used for the quorum resource after creating the cluster, see "Use a different disk for the quorum resource" in Related Topics.

- To create and configure a cluster after an unattended setup, run a script to invoke the **cluster /create:** command and supply all the necessary configuration information on the command line. For information about unattended setup, see "Planning for unattended setup" in Related Topics. For more information on the cluster command line syntax, see "Cluster" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Quorum resource
Checklist: Installing a Local Quorum resource
Choosing a Cluster Model
Planning for unattended Setup
Cluster
Checklist: Installing a Physical Disk resource
Use a different disk for the quorum resource

# Add additional nodes to the cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add additional nodes to the cluster**

1. Open Cluster Administrator.
2. In the **Open Connection to Cluster** dialog box that appears, in **Action**, select **Add nodes to cluster**. Then in **Cluster or server name**, either type the name of an existing cluster, select a name from the drop-down list box, or click **Browse** to search for an available cluster.
3. Click **OK** to continue.
4. The Add Nodes Wizard appears. Click **Next** to continue.
5. Upon completion of the Add Nodes Wizard, click **Finish**.

### Caution

- If you are using a shared storage device, when you power on and start the operating system, it is of vital importance that only one node has access to the cluster disk. Otherwise the cluster disks can become corrupted. To prevent the corruption of the cluster disks, shut down all but one cluster node, or use other techniques (for example, Logical Unit Number or LUN masking, selective presentation, or zoning) to protect the cluster disks, before creating the cluster. Once the Cluster service is running properly on one node, the other nodes can be added and configured simultaneously.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- You can also add a node to an existing cluster from the **File** menu, by first clicking **Open Connection** to connect to an existing cluster. Then, you can open the Add Nodes Wizard from the **File** menu, by clicking **New**, then **Node**.
- You can safely cancel and restart the installation process.
- If the computer you are logged onto is already a node on an existing cluster, Cluster Administrator will connect to that cluster by default.
- After you have added new nodes to a server cluster, use Cluster Administrator to fail over the cluster disks and verify that they come back online. For information on how to do this, see "To test whether group resources can fail over" in Related Topics.
- To add nodes to an existing cluster after an unattended setup, run a script to invoke the **cluster /add:** command and supply all the necessary configuration information on the command line. For information about unattended setup, see "Planning for unattended setup" in Related Topics. For more information on the cluster command line syntax, see "Cluster" in Related Topics.
- Allow the process of adding or evicting nodes to complete before changing the quorum in any way. For example, if you are adding a node, wait until that process is complete before changing from a single quorum device cluster to a majority node set cluster.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Test whether group resources can fail over
Planning for unattended Setup
Cluster

# Evict a node from the cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To evict a node from the cluster**

1. Close all applications on the node that is to be removed.
2. Open Cluster Administrator.
3. Review the failover policies of the resource groups on the node to be evicted.

   For information on how to do this, see "Determining failover and move policies for groups" in Related Topics.
4. If necessary, move all groups on the node to be evicted to another node or nodes in the cluster.

   For information on how to do this, see "To move a group to another node" in Related Topics.
5. Select the node to be evicted.
6. Stop the Cluster service if the node being evicted is not the last node in the cluster. That is, on the **File** menu, click **Stop Cluster Service**.

   For information on how to do this, see "To stop the Cluster service" in Related Topics.

   If the node being evicted is the last node in the cluster, proceed to the next step.
7. On the **File** menu in Cluster Administrator, click **Evict Node**.

**Caution**

- Prior to evicting a node, check with application vendors to determine whether your applications require unclustering.
- Application-specific data may be lost if you do not close your applications before you perform this procedure.

**Important**

- Evicting the last node from the server cluster will not automatically delete the Cluster service account. It is highly recommended that you remove this account from the local Administrators group if you no longer have use for it. However, because the administration rights and permissions for the Cluster service account are granted locally on each cluster node and not domain wide, the impact of a compromised account is limited to just the cluster nodes. For more information, see Best practices for securing server clusters.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- In some cases you may be unable to stop the Cluster service on a node or evict the last node from the cluster, for example, if you are administering a mixed-version server cluster from a Windows 2000 node and try to evict a Windows

Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition node. In these cases, you can manually remove the configuration of the Cluster service on that node by typing **cluster node***node name***/force**[**cleanup**] at the command line. For more information on the **cluster node** command-line options, see Related Topics.

- When you evict a node from the cluster, that node is configured as a stand-alone node (that is, not a cluster node). The clustering software remains installed on the node. You can re-add the evicted node to this cluster or create a new cluster on that node at any time.

- Typically, you remove a node from a cluster when testing, redeploying, or replacing a node.

- Allow the process of adding or evicting nodes to complete before changing the quorum in any way. For example, if you are evicting a node, wait until that process is complete before changing from a single quorum device cluster to a majority node set cluster.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Repair a node without evicting it from the cluster
Evict a node from the cluster
Specify preferred owners of a group
Specify which nodes can own a resource
Stop the Cluster service
Determining failover and move policies for groups
Move a group to another node
Cluster node

# Repair a node without evicting it from the cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To repair a node without evicting it from the cluster**

1. Open Cluster Administrator.
2. In the console tree, click the node requiring repair.
3. On the **File** menu, click **Stop Cluster Service**.
4. If the node must be rebooted after being repaired, set the Cluster service to start manually.

   For more information, see "To set the Cluster service to only start manually" in Related Topics.
5. Perform repair operations on the node.
6. On the **File** menu, click **Start Cluster Service** to restart the Cluster service on that node.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Evict a node from the cluster
Set the Cluster service to only start manually

# Pause a node

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To pause a node**

1. Open Cluster Administrator.
2. In the console tree, click the node.
3. On the **File** menu, click **Pause Node**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- When you pause a node, existing groups and resources stay online, but additional groups and resources cannot be brought online on the node. Pausing a node is used for maintenance and testing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Repair a node without evicting it from the cluster
Resume a node

# Resume a node

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To resume a node**

1. Open Cluster Administrator.
2. In the console tree, click the node.
3. On the **File** menu, click **Resume Node**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Pause a node
Repair a node without evicting it from the cluster

# Stop the Cluster service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To stop the Cluster service**

1. Open Cluster Administrator.
2. In the console tree, click the node.
3. On the **File** menu, click **Stop Cluster Service**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. In addition, your account must be added to the security descriptor for the Cluster service. To do this, you can use Cluster Administrator or cluster.exe. For more information, see "To give a user permissions to administer a cluster" and "Cluster" in Related Topics. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- When you stop the Cluster service on a node, you prevent clients from accessing cluster resources through that node. In this event, all groups move to another node (if the failover policies allow it).

  This set of circumstances may also affect the availability of the cluster. For example, if the node to which the groups failed over fails, and your cluster consists of only two nodes, there will be no other node available to which the groups can then fail over.
- This procedure, stopping the Cluster service, can also be performed in Computer Management, which may be necessary if Cluster Administrator fails. You can also start the Cluster service in Computer Management.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Evict a node from the cluster
Start the Cluster service
Cluster
Give a user permissions to administer a cluster

# Start the Cluster service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the Cluster service**

1. Open Cluster Administrator.
2. In the console tree, click the node.
3. On the **File** menu, click **Start Cluster Service**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. In addition, your account must be added to the security descriptor for the Cluster service. To do this, you can use Cluster Administrator or cluster.exe. For more information, see "To give a user permissions to administer a cluster" and "Cluster" in Related Topics. As a security best practice, consider using Run as to perform this procedure. For more information, see Default local groups, Default groups, and Using Run as.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- You can also start the Cluster service from Computer Management. For more information, see "To set the Cluster service to only start manually" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Stop the Cluster service
Set the Cluster service to only start manually
Cluster
Give a user permissions to administer a cluster

# Set the Cluster service to only start manually

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the Cluster service to only start manually**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. In the details pane, click **Cluster Service**.
4. On the **Action** menu, click **Properties**.
5. In **Startup type**, select **Manual**, and then click **OK**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- To start the Cluster service even if there are problems with the quorum resource, you can use **fixquorum** as a start parameter. That is, in **Start parameters**, type **/fixquorum**, and then click **Start**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start the Cluster service
Start, stop, pause, resume, or restart a service

# Create and Administer File and Print Server Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Create and administer file and print server clusters**

- Create a cluster-managed printer
- Create a cluster-managed file share
- Create a cluster-managed encrypted file share
- Share subfolders on a File Share resource
- Set client-side caching for a File Share resource
- Enable Shadow Copies for Shared Folders in a cluster

# Create a cluster-managed printer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a cluster-managed printer**

1. Open Cluster Administrator.

2. Create a virtual server, that is, a resource group containing a Network Name resource (where the Network Name resource is the name of a virtual server), an IP address resource, and a Physical Disk resource, if these do not already exist.

   For more information on creating a new resource group and resources, see "To create a new group" and "To create a new resource" in Related Topics.

   For more information, see "Virtual servers" in Related Topics.

3. Create a Print Spooler resource in that virtual server and make it dependent on the Network Name and Physical Disk resources.

   **Important**

   - These two resources are required dependencies for the Print Spooler resource.

   For more information, see "Setting resource dependencies" and "To modify resource dependencies" in Related Topics.

4. Bring the virtual server containing the Print Spooler resource online.

   For more information, see "To bring a group online" in Related Topics.

5. Install printer drivers for your printer on the virtual server.

   **To install printer drivers on the virtual server:**

   a. Click **Start** and then click **Run**.

   b. In the **Run** dialog box, type *MyVirtualServer* where *MyVirtualServer* is the name of the virtual server that the print spooler resource depends on.

      The virtual server opens.

   c. Open **Printers and Faxes**, right-click and select **Server Properties**, then click the **Drivers** tab.

   d. Install the printer drivers.

   **Important**

   The port monitor must be installed on all nodes in the cluster before installing printer drivers on the virtual server. For the Windows Server 2003 family, the port monitor for the Standard TCP/IP Port is installed by default with the operating system. For two other types of ports, you must first install the appropriate port monitor on each node before installing the printer drivers, printer port, and printer queue on the virtual server.

   a. LPR Ports: The LPR port monitor files are shipped on the Windows Server 2003 family media. For more information on installing the LPR port monitor and adding it to the list of **Available port types** for a node, see "To

add a port" in Related Topics.

b. Custom IHV Ports: Printer drivers from Independent Hardware Vendors (IHVs) may include a custom IHV port monitor that must first be installed on each node. To install the IHV port monitor on each node, install the driver package on each virtual server in the cluster. Contact your IHV to verify that the third party port monitor is compatible with the Cluster service before installing on the cluster nodes.

For more information, see "To add printer drivers for other versions of Windows" in Related Topics.

6. Install a printer port for your printer on the virtual server.

   **To install a printer port on the virtual server:**

   a. Click **Start** and then click **Run**.

   b. In the **Run** dialog box, type *MyVirtualServer* where *MyVirtualServer* is the name of the virtual server that the print spooler resource depends on.

      The virtual server opens.

   c. Open Printers and Faxes, right-click and select **Server Properties**, then click the **Ports** tab.

   d. Add a port.

   For more information, see "To add a port" in Related Topics.

7. Add a print queue to the virtual server.

   **To install a print queue on the virtual server:**

   a. Open the remote printers folder of the virtual server.

   b. Use the **Add Printer** Wizard to add a new print queue.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

- You can use the **Add Printer** Wizard to combine steps 5 through 7 above into a single step by adding the printer drivers, printer port, and print queue successively. However, before using the **Add Printer** Wizard in this way, make sure that the appropriate port monitor is already installed on all nodes in the cluster.

- Not all printer ports may be used in a server cluster. LPR and Standard TCP/IP Port, both provided with the Windows Server 2003 family, may be used in a server cluster. Check your printer documentation to verify that your printer can use one of these two ports.

- To share out the printer, you must share out the print queue from the virtual server itself and not from the individual nodes.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
File and Printer Sharing for Microsoft Networks
Create a new group in a cluster
Create a new resource
Virtual servers
Setting resource properties
Modify resource dependencies
Bring a group online
Add printer drivers for other versions of Windows
Add a port

# Create a cluster-managed file share

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a cluster-managed file share**

1. Using Windows Explorer, create a folder on the cluster disk.

   **Caution**

   - Do not create a share for this folder.

2. Verify that the Cluster service account has Full Control rights to this folder for the NTFS file system.

   For more information on setting folder permissions, see "To set, view, change, or remove file and folder permissions" in Related Topics.

3. Open Cluster Administrator.

4. Create a Physical Disk resource for the drive on which the folder is located, a Network Name resource, and an IP Address resource, if these do not already exist, and place them in the same resource group.

5. Create a File Share resource in the same resource group as the resources above.

   By default, Cluster Administrator will create a **Normal** File Share resource.

6. Optionally, click **Advanced** to change the File Share resource type from **Normal** to **DFS root** or **Share subdirectories**.

   For more information on these three File Share resource types, see "File Share resource type" in Related Topics.

   **Important**

   - If the file share is configured as a DFS root, it must be dependent on a Network Name resource.

7. Although it is not required, it is recommended that you set up the File Share resource so that it has a dependency on the Physical Disk resource as well as the Network Name resource.

   For more information on creating a new resource and setting up a dependency, see "To create a new resource" in Related Topics.

8. In the details pane, click the File Share resource that you created, then on the **File** menu, click **Properties**.

9. On the **Parameters** tab in **Path**, type the path to the folder whose subfolders you want to share.

10. To configure file share permissions, click **Permissions**.

    For more information about setting permissions, including NTFS permissions, see "To set permissions on a shared resource" in Related Topics.

11. On the **File** menu, click **Bring online**.

**Caution**

- Always use the Cluster Administration tools and not Windows Explorer to manage these file shares. Changes made using Windows Explorer will be lost when these file shares fail over to other nodes in the cluster.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set client-side caching for a File Share resource
File Share resource type
Create a new resource
Set, view, change, or remove permissions on files and folders
Set permissions on a shared resource

# Create a cluster-managed encrypted file share

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a cluster-managed encrypted file share**

1. Open Cluster Administrator.
2. Create a cluster-managed file share.

   For more information, see "To create a cluster-managed file share" in Related Topics.
3. In the console tree, double-click **Groups**, and then click the group that contains the File Share resource that you created in step 2.
4. In the details pane, click the Network Name resource in that you will use to connect to this file share.
5. On the **File** menu, click **Take Offline**.
6. On the **File** menu, click **Properties**.
7. On the **Parameters** tab, select **Enable Kerberos Authentication**, and then click **OK**.
8. In the details pane, click the Network Name resource for the file share.
9. On the **File** menu, click **Bring Online**.
10. On the domain controller, configure all accounts used by the cluster nodes to be trusted for delegation. You must also configure the computer account used by the Network Name resource on which you enabled Kerberos to be trusted for delegation.

    For more information, see "To allow a user to be trusted for delegation" in Related Topics.

    **Important**

    - The group containing the file share resource is a virtual server because it contains a Network Name and an IP Address resource. All accounts used to connect to this virtual server must also be trusted for delegation.
11. On the domain controller, enable roaming profiles for all users or service accounts that store data in the folder managed by the File Share resource.

    For more information, see "To create a roaming user profile" in Related Topics.

    **Important**

    - Enabling roaming profiles ensures that the encryption keys for a folder are still valid even when the File Share resource managing that folder fails over.
12. Enable encryption on the folder managed by the File Share resource.

    For more information, see "To encrypt a file or folder" in Related Topics.

**Note**

- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Create a cluster-managed file share
Virtual servers
Allow a user to be trusted for delegation
Create a roaming user profile
Encrypting File System overview
Encrypt a file or folder

# Share subfolders on a File Share resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To share subfolders on a File Share resource**

1. Open Cluster Administrator.

2. Create a File Share resource.

   For more information on creating a new resource, see Related Topics.

3. In the console tree, double-click **Groups**, and then click the group that contains the File Share resource that you created in step 2.

4. In the details pane, click the File Share resource that you created. On the **File** menu, click **Properties**.

5. On the **Parameters** tab in **Path**, type the path to the folder whose subfolders you want to share.

6. To configure file share permissions, click **Permissions**.

   For more information about setting permissions, including NTFS permissions, see "To set permissions on a shared resource" in Related Topics.

7. Click **Advanced**, and then click **Share subdirectories** to enable the combo file share option that allows File Share resources to share subfolders. If you want to prevent users from seeing the subdirectory shares when they browse the network, click **Hide subdirectories shares**, and then click **OK**.

   For more information about file and printer sharing, see "File and Printer Sharing for Microsoft networks" in Related Topics.

8. On the **File** menu, click **Bring online**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

- The combo file share option permits a File Share resource to share subfolders. Subfolders of a combo file share are shared using the folder name. You cannot specify a name for a subfolder share. The only way to change a subfolder file share name is to change the subfolder name (for example, using Windows Explorer). Changing the name of a shared subfolder breaks the file share temporarily. The subfolder file share will be unavailable for the time specified by the **File Share resource Is Alive** poll interval. The subfolder will be shared again the next time the **File Share resource Is Alive** poll is called.

- If you select the **Hide subdirectories shares** option you do not have to type $ as the last character when specifying the file share name. The File Share resource will automatically hide the subdirectories when you select the **Hide subdirectories shares** option.

- For more information, see article Q256926, "Implementing Home Folders on a Server" in the Microsoft Knowledge Base.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
File Share resource type
Using a server cluster with large numbers of file shares
Create a new resource
Set client-side caching for a File Share resource
Set permissions on a shared resource
File and Printer Sharing for Microsoft Networks

# Set client-side caching for a File Share resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set client-side caching for a File Share resource**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the appropriate File Share resource.
4. On the **File** menu, click **Properties**.
5. Click the **Parameters** tab, then click the **Caching** button.
6. Select the **Allow caching of files in this shared folder** check box to enable client-side caching for this File Share resource.
7. In **Setting**, select the appropriate client-side caching setting as shown in the table below.

| Caching setting in Cluster Administrator | Command line option in cluster.exe (cluster resource) | Description |
|---|---|---|
| **Allow caching of files in this shared folder** is not enabled | **CSCCache** is set to 48 | Client-side caching is disabled for all files on the file share. |
| **Automatic caching of programs and documents** | **CSCCache** is set to 32 | Client-side caching is disabled for those files that are only opened, but not modified. |
| **Automatic caching of documents** | **CSCCache** is set to 16 | Changes in all files opened on the client will be reconciled between the client and the server cluster. |
| **Manual caching of documents** | **CSCCache** is set to 0 | Changes in only those files specifically identified on the file share will be reconciled between the client and the server cluster. |

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- Client-side caching allows network users to access files on clustered file shares, even when the client computer is disconnected from the network, by caching these offline files on the client computer. You can also set the client-side caching for a File Share resource at the command line by using the corresponding **cluster resource** command option as shown in the table above. For more information, see "Cluster resource" in Related Topics.
- For more information on setting up your computer to use offline files, see "Sharing a drive or folder overview" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Cluster resource
Server Cluster Resources
Resource types
Server Cluster groups
Nodes
Sharing a drive or folder

# Enable Shadow Copies for Shared Folders in a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Shadow Copies of Shared Folders in a cluster**

1. On the cluster node where you want to enable Shadow Copies of Shared Folders, create a cluster-managed file share.

2. Set up the File Share resource so that it has a dependency on the Physical Disk resource that manages the shadow copies and designated storage volumes, as well as the Network Name resource for the file share.

   For more information, see "To create a cluster-managed file share" in Related Topics below.

   **Important**

   - If the shadow copy volume and the designated storage volume are on separate disks, the Physical Disk resources that manage those volumes must both be in the same resource group.

3. Open Computer Management.

4. In the console tree, right-click **Shared Folders**, select **All Tasks**, and click **Configure Shadow Copies**.

   **Where?**

   - Computer Management/System Tools/Shared Folders

5. Click the volume where you want to enable Shadow Copies of Shared Folders, and then click **Enable**.

   **Important**

   - The Computer Management tool automatically generates a scheduled task to create the shadow copies on the storage volume and a Volume Shadow Copy Service Task resource to manage that scheduled task. It also creates a dependency on the Physical Disk resource that manages the storage volume containing the shadow copies. For more information on the default settings for the Volume Shadow Copy Service Task resource created when you enable Shadow Copies of Shared Folders, see "Using Shadow Copies of Shared Folders in a server cluster" in Related Topics below.

**Important**

- If the shadow copy volume and the designated storage volume are on separate disks, enabling Shadow Copies of Shared Folders will automatically make the Physical Disk resource that manages the shadow copy volume dependent on the Physical Disk resource that manages the storage volume.

- When you enable shadow copies, the shadow copy volumes (and all other resources directly or indirectly dependent on the disk) go offline for a brief period while the resource dependencies are created. The shadow copy volumes are not accessible to applications and users while they are offline.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- You cannot enable Shadow Copies of Shared Folders for the quorum resource, although you can enable Shadow Copies of Shared Folders for a File Share resource.

- You can also access Shadow Copies of Shared Folders by right-clicking a fixed NTFS volume in **My Computer** or in the Disk Management snap-in details pane (in **Computer Management**). Click **Properties**, and then click the **Shadow Copies** tab.

- You can use the Disk Management snap-in to manage Shadow Copies of Shared Folders on remote servers as well. In the **Computer Management** console tree, right-click **Disk Management**, point to **All Tasks**, and then click **Configure Shadow Copies**.

- By enabling Shadow Copies of Shared Folders, you also create a default schedule and settings for subsequent shadow copies. To make changes to these defaults, click **Settings**.

- You can only enable Shadow Copies of Shared Folders on a per-volume basis; that is, you cannot select specific shared folders and files on a volume to be copied or not copied.

- The scheduled task that is automatically created when you enable shadow copies on a cluster node will only appear in the Scheduled Tasks folder when the associated Volume Shadow Copy Service Task resource is online. If you want to change the schedule, use Computer Management; do not use the Scheduled Tasks folder.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Shared Folders overview
Change the settings for Shadow Copies of Shared Folders
Shadow Copies for Shared Folders
Volume Shadow Copy Service resource type
Create a new resource
Using Shadow Copies of Shared Folders in a server cluster
Create a cluster-managed file share
Resource dependencies

# Manage Accounts and Domains in a Cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage accounts and domains in a cluster**

- Change the account under which the Cluster service runs
- Change the Cluster service account password
- Give a user permissions to administer a cluster
- Move a server cluster to another domain

# Change the account under which the Cluster service runs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the account under which the Cluster service runs**

1. Stop the Cluster service on all nodes:

   - Make sure the account has membership in the local Administrators group on all nodes.
   - Open Local Security Policy and Grant the following rights to the account, or to the local Administrators group, on all nodes:

   **Where?** Security Settings/Local Policies/User Rights Assignments

   - **Act as part of the operating system**
   - **Back up files and directories**
   - **Restore files and directories**
   - **Adjust memory quotas for a process**
   - **Log on as a service**
   - **Increase scheduling priority**

   By default, the Cluster service account inherits the following user rights as a result of being a member of the local Administrators group:

   - **Manage auditing and security log**
   - **Debug programs**
   - **Impersonate a client after authentication**

   If your organization has removed these user rights from the default set of privileges assigned to the local Administrators group, you need to specifically assign these user rights to the Cluster service account.

   - Open Computer Management.
   - In Computer Management, double-click **Services and Applications**, and then click **Services**.
   - In the details pane, click **Cluster Service**.
   - On the **Action** menu, click **Stop**.

2. Repeat step 1 on all other nodes.

3. In the details pane of one node, double-click **Cluster service**.

4. On the **Log On** tab, type the account name in **This account**, type the password in **Password**, and then confirm the password and click **OK**.

5. On the **Action** menu, click **Start**.

6. Repeat steps 2, 3, 4, and 5 on all other nodes.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- To open Local Security Policy, click **Start**, point to **Settings**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Local Security Policy**. Select **Show Advanced User Rights** to view all the available security settings.

- The Cluster service on all nodes must be stopped and restarted during this procedure. The Cluster service must use the same account and password at all times on all nodes within the cluster.

- You must use Active Directory Users and Computers to view account properties if the account is in a Windows 2000 or Windows Server 2003 family domain. You must use User Manager for the Domain to view account properties if the account is in a Windows NT 4.0 domain. User Manager for the Domain can be installed using the client-based Network Administration tools by running Setup.exe from the Windows NT Server 4.0 media in the Clients\Srvtools\Winnt\ directory or by running Usrmgr.exe from the Clients\Srvtools\Winnt\I386 or Clients\Srvtools\Winnt\Alpha directory.

- For more information on granting account rights, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Manage Users, Groups, and Computers
Start, stop, pause, resume, or restart a service
Domain accounts and the Cluster service
Updated technical information

# Change the Cluster service account password

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the Cluster service account password**

1. Log on to a computer in the domain. The account you use to log in must belong to the local administrative group on all the cluster nodes.

   **Important**

   - To be able to change the Cluster service account password, you must have appropriate security administrative rights and permissions to be able to traverse directories in the cluster domain even if you do not have user rights on the traversed directory (Bypass traverse checking). By default, this user right is enabled for all users, but some administrators may disable it for everyone.

     For more information, see "Bypass traverse checking" in Related Topics.

   - For multiple clusters that use the same Cluster service account, it is highly recommended that you change the password for all the clusters at the same time.

   - You can only use this procedure to change the Cluster service account if applications (for example, Microsoft SQL Server 2000) do not use the same account as the Cluster service account. If you use this procedure on a server cluster where the Cluster service and one or more applications use the same account, your cluster applications may not function correctly. To resolve this problem, change the application account password (see your application's documentation for details).

2. Open Command Prompt

3. Type the following command, using the table below as a guide:

   **cluster/cluster:***cluster_name1*[,*cluster_name2*, ...] **/changepass**[**word**][**:***new_password*[,*old_password*]] [**/skipdc**] [**/force**] [**/***options*]

   For more information on the other optional password change cluster.exe commands, see "Cluster" in Related Topics.

| Argument | Description |
|---|---|
| **/cluster:***cluster_name1*[,*cluster_name2*, ...] | Identifies the cluster(s) for the account password change. If multiple clusters are specified in one operation, they must all use the same Cluster service account. If some nodes are unavailable, the account password is not changed on any of the nodes nor on the domain controller. |
| **/changepass**[**word**] [**:***new_password*[,*old_password*]] | Changes the Cluster service account password on the domain controller and all the cluster nodes from *old_password* to *new_password*. If not supplied as part of the command, you will be prompted to provide the *new_password* and/or the *old_password*. You can specify a blank password by supplying two double quotes ("""). |
| **/skipdc** | Changes the Cluster service account password only on the cluster nodes. You can use this command if you previously issued a **cluster /changepass** command for multiple clusters and the Cluster service account password was not updated on all the clusters, for whatever reason. If you use this argument, you do not have to supply the *old_password* parameter. |

| | |
|---|---|
| **/force** | Forces the execution of the password change command on the available nodes in the cluster even if some of the cluster nodes are not available.<br><br>**Important**<br><br>• Any node that is down when the password is changed will be unable to later join the cluster until the Cluster service account password for that node is updated manually using Computer Management. |

For the first example, to change the password for clusters CLUS1 and CLUS2 from B2PartTwo! to WordTestPassPartOne! on the domain controller and on all cluster nodes, type:

**cluster /cluster:CLUS1,CLUS2 /changepass:WordTestPassPartOne!, B2PartTwo!**

For the second example, if 2 of the nodes in cluster CLUS1 were down, you could force this password change on the domain controller and on all available cluster nodes in CLUS1 by typing:

**cluster /cluster:CLUS1 /changepass:WordTestPassPartOne!, B2PartTwo! /force**

## Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To view the complete syntax for this command, at a command prompt, type:

  **cluster /changepass /?**

- If you do not use the **/force** command in a cluster where some of the nodes are unavailable, the Cluster service password for any available cluster nodes or the domain controller will not be updated.

- You cannot use this method to change the Cluster service account password if some nodes in the cluster are running Windows NT or Windows 2000, while others are running Windows Server 2003 family operating systems.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Cluster
Reset a user password
Domain accounts and the Cluster service
Bypass traverse checking

# Give a user permissions to administer a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To give a user permissions to administer a cluster**

1. Open Cluster Administrator.
2. Select the cluster.
3. On the **File** menu, click **Properties**.
4. On the **Security** tab, give users permissions to administer the cluster.

### Important

- It is recommended that you allow only trusted personnel to administer a cluster.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- To administer a cluster using Cluster Administrator, you must use an account that is a member of the local Administrators group on each node. The account should be a domain account and can be a member of other groups, such as global groups.
- You can use an account that is a domain user account or a member of a global group to remotely or locally administer a cluster if you use cluster.exe or cluster scripting tools. For more information, see Cluster and the Microsoft Platform Software Development Kit (SDK).
- The local Administrator account and local system account always have access to the cluster.
- When you perform this procedure, the server on which you run Cluster Administrator must run one of the following versions of the operating system and service pack:
    - Windows Server 2003
    - Windows Server 2003 with Service Pack 1
    - A product in the Windows 2000 Server family with Service Pack 2 or later

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Domain accounts and the Cluster service

Change the account under which the Cluster service runs

# Move a server cluster to another domain

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move a server cluster to another domain**

1. Create a user account for the Cluster service in the new domain.

   For more information, see "To create a new user account" in Related Topics.

2. Open Local Security Policy.

   **Where?**

   - Security Settings/Local Policies/User Rights Assignments

3. Grant that account the following rights:

   - **Act as part of the operating system**
   - **Back up files and directories**
   - **Restore files and directories**
   - **Adjust memory quotas for a process**
   - **Log on as a service**
   - **Increase scheduling priority**

   By default, the Cluster service account inherits the following user rights as a result of being a member of the local Administrators group:

   - **Manage auditing and security log**
   - **Debug programs**
   - **Impersonate a client after authentication**

   If your organization has removed these user rights from the default set of privileges assigned to the local Administrators group, you need to specifically assign these user rights to the Cluster service account.

   **Important**

   - The new Cluster service account must have administrative rights and permissions on all nodes in the cluster.

4. Stop the Cluster service on all nodes.

   For more information, see "To stop the Cluster service" in Related Topics.

5. Power down all nodes except one.

6. Move the node to another domain.

   For more information, see "To join a domain" in Related Topics.

7. Set the Cluster service to start manually.

   For more information, see "To set the Cluster service to only start manually" in Related Topics.

8. Restart the node.

9. On the node, change the service account used by the Cluster service to log on to the domain using the user account you just created.

   **How?**

   a. Open Computer Management.

   b. In the console tree, click **Services**.

      **Where?** Computer Management/Services and Applications/Services

   c. In the details pane, click **Cluster Service**.

   d. On the **Action** menu, click **Properties**.

   e. On the **Log On** tab, select **This account**, and then supply the new account and password.

10. Start the Cluster service.

    For more information, see "To start the Cluster service" in Related Topics.

11. Verify that cluster node is operating correctly.

    **How?**

    a. Bring all resources online.

       For more information, see "To bring a resource online" in Related Topics.

    b. Ping the IP address of the cluster, the individual nodes, and any resource containing an IP address (for example, a virtual server) from all clients. That is, on each client, type the following at the command prompt:

       **ping** xxx.xxx.xxx.xxx

       where xxx.xxx.xxx.xxx is the IP address.

    **Important**

    - At this point, you can still cancel the move by moving this node back into the old domain and starting those nodes that have not yet been moved.

12. Power up another node and repeat steps 5-10 to migrate the other nodes, one by one, to the new domain.

### Caution

- Use this procedure only if your cluster is extremely simple; otherwise recreate your existing cluster in the new domain.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Quorum resource
Checklist: Installing a Local Quorum resource
Create a new user account
Stop the Cluster service
Join a domain
Set the Cluster service to only start manually
Start the Cluster service
Bring a resource online

# Manage the Cluster Configuration and Diagnostic Logs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage the cluster configuration and diagnostic logs**

- View the cluster configuration log
- View the cluster diagnostic log
- Change the size of the cluster diagnostic log

# View the cluster configuration log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the cluster configuration log**

1. Open Command Prompt
2. Locate the ClCfgSrv.log file.

   By default, the log file for the New Server Cluster Wizard and the Add Nodes Wizard can be found at %systemroot%\system32\LogFiles\Cluster\ClCfgSrv.log.
3. Use any text editor to view the log file.

### Important

- After using the New Server Cluster Wizard and the Add Nodes Wizard from a remote computer, the cluster configuration log file generated by those wizards is saved to the remote computer. This log file may contain company-sensitive data such as a complete description of the cluster configuration, including IP addresses. This information could be used by malicious users to attack the cluster; it is recommended that you restrict access to that file to cluster administrators and the Cluster service account. For more information, see Best practices for securing server clusters.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- If you do not have administrative rights on the node on which you are using the New Server Cluster Wizard or Add Nodes Wizard, the log file will be written to the local %Temp% directory.
- This log file contains all of the errors, warnings, and informational messages that have been generated by the New Server Cluster Wizard or the Add Nodes Wizard when creating a new cluster or when the cluster membership changes (for example, if a new node is added to the cluster).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Additional New Server Cluster Wizard information
Checklist: Planning and creating a server cluster

# View the cluster diagnostic log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the cluster diagnostic log**

1. Open Command Prompt
2. Locate the cluster.log file.

   By default, the log file for the Cluster service can be found at %systemroot%\Cluster\cluster.log.
3. Use any text editor to view the log file.

   For more information about interpreting the entries in the cluster log, see "Interpreting the Cluster Log" at the Microsoft Windows Resource Kits Web site.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the size of the cluster diagnostic log
Managing the Cluster Service from the Command Line

# Change the size of the cluster diagnostic log

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the size of the cluster diagnostic log**

1. Open Control Panel.
2. Double-click **System**.
3. On the **Advanced** tab, click **Environment Variables**.
4. Under **System variables**, click **New**.
5. In the **New System Variable** dialog box, in **Variable name**, type **ClusterlogSize**, and in **Variable value**, type the size in megabytes that you want the cluster log to be.
6. Click **OK** until you exit the **System Properties** page.
7. Stop and restart the Cluster service. For more information on starting and stopping the Cluster service, see Related Topics.

## Notes

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- The default diagnostic cluster log size is 8 megabytes (MB), but can be changed in the manner described previously.
- When the cluster log file is full, the Cluster service overwrites the first half of the log with the second half. Therefore, if the cluster log size is *n*MB, then only (*n*/2)MB of that log is guaranteed to be valid.
- You must set ClusterlogSize as a System variable, not a user variable.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Stop the Cluster service
Start the Cluster service

# Perform Advanced Administrative Tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Perform advanced administrative tasks**

- Use a different disk for the quorum resource
- Recover from a corrupted quorum log or quorum disk
- Migrate a multinode cluster to a single-node configuration
- Force quorum in a majority node set server cluster
- Enable Kerberos authentication for virtual servers

# Use a different disk for the quorum resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To use a different disk for the quorum resource**

1. Choose the disk that you will use for the quorum resource. If the disk has two or more NTFS partitions, ensure that all partitions on the disk are assigned drive letters.
2. Open Cluster Administrator.
3. If one does not already exist, create a Physical Disk or other storage-class resource for the new disk.

   For more information, see "Checklist: Installing a Physical Disk resource" in Related Topics.
4. In the console tree, click the Resources folder.
5. In the details pane, click the resource that you will use for the quorum resource.
6. On the **File** menu, click **Take Offline.**
7. In the console tree, click the cluster name.
8. On the **File** menu, click **Properties**.
9. On the **Quorum** tab, click **Quorum resource**, and then select a new disk or storage-class resource you want to use as the quorum resource for the cluster.
10. In **Partition**, if the disk has more than one partition, click the partition where you want the cluster specific data kept.
11. In **Root path**, type the path to the folder on the partition; for example:

    \MSCS\
12. In the console tree, click the Resources folder.
13. In the details pane, click the new quorum resource.
14. On the **File** menu, click **Bring Online**.

**Important**

- If Cluster Administrator does not open, it may be because the Cluster service is not started. To start the Cluster service, see "To set the Cluster service to only start manually" in Related Topics.
- If the quorum resource becomes corrupted, you cannot move it until it is repaired. For more information, see "Node-to-node connectivity problems" in Related Topics.
- If the Cluster service is stopped on any other node, or the node is shut down, that node will not be able to form a cluster. Only nodes that are online when this change is made can form the cluster. However, nodes that are offline can still join the cluster and will be able to form a cluster after they have joined the cluster at least once.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

- Of the default resource types, only the Physical Disk, Local Quorum, or Majority Node Set resource can be a quorum resource. However, third party vendors can supply other storage class resource types that are quorum capable.

- When you are making changes to the quorum, do not add or evict nodes at the same time.

- If you designate a resource as the quorum resource without first taking that resource offline, restart a node afterwards. This will ensure that normal checkpointing is initiated for the quorum resource. Checkpointing is the saving of an extra copy of cluster configuration data and is normally done at four-hour intervals.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign, change, or remove a drive letter
Quorum resource
Set the Cluster service to only start manually
Checklist: Installing a Physical Disk resource
Node-to-node connectivity problems

# Recover from a corrupted quorum log or quorum disk

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To recover from a corrupted quorum log or quorum disk**

1. If the Cluster service is running, open Computer Management.
2. In the console tree, double-click **Services and Applications**, and then click **Services**.
3. In the details pane, click **Cluster Service**.
4. On the **Action** menu, click **Stop**.
5. Repeat steps 1, 2, 3, and 4 for all nodes.
6. If you have a backup of the quorum log, restore the log by following the instructions in "Backing up and restoring server clusters" in Related Topics.
7. If you do not have a backup, select any given node. Make sure that **Cluster Service** is highlighted in the details pane, and then on the **Action** menu, click **Properties**.

   Under **Service status**, in **Start parameters**, specify **/fixquorum**, and then click **Start**.
8. Switch from the problematic quorum disk to another quorum resource.

   For more information, see "To use a different disk for the quorum resource" in Related Topics.
9. In Cluster Administrator, bring the new quorum resource disk online.

   For information about how to do this, see "To bring a resource online" in Related Topics.
10. Click **Start**, click **Run**, and type a command with the following syntax:

    **cluster** [*ClusterName*] **res** *QuorumDiskResourceName* **/maint:on**
11. Run Chkdsk, using the switches **/f** and **/r**, on the quorum resource disk to determine whether the disk is corrupted.

    For more information on running Chkdsk, see "Chkdsk" in Related Topics.
12. Click **Start**, click **Run**, and type a command with the following syntax:

    **cluster** [*ClusterName*] **res** *QuorumDiskResourceName* **/maint:off**
13. If no corruption is detected by Chkdsk, it is likely that the log was corrupted. Proceed to step 15.

    If corruption is detected on the disk, check the system log in Event Viewer for possible hardware errors.

    Resolve any hardware errors before continuing.
14. Stop the Cluster service after Chkdsk is complete, following the instructions in steps 1-4.
15. Make sure that **Cluster Service** is highlighted in the details pane. On the **Action** menu, click **Properties**.

    Under **Service status**, in **Start parameters**, specify **/resetquorumlog**, and then click **Start**.

This restores the quorum log from the node's local database.

**Important**

- The Cluster service must be started by clicking **Start** on the service control panel. You cannot click **OK** or **Apply** to commit these changes as this does not preserve the **/resetquorumlog** parameter.

16. Restart the Cluster service on all other nodes.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.

- The quorum disk must be formatted with the NTFS file system.

- If none of the nodes are running, or one node fails while you are changing the quorum resource, only the running nodes are able to form the cluster, and the offline node is only able to join the cluster. After the offline node has joined the cluster, all nodes are again able to form or join the cluster. This design prevents the offline node from forming the cluster using the old quorum resource.

- Optionally, after step 13 above, you can use the ClusterRecovery tool, available in the *Microsoft Windows Server 2003 Resource Kit*, to restore the registry checkpoint files.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Start, stop, pause, resume, or restart a service
Quorum resource
Backing up and restoring server clusters
Use a different disk for the quorum resource
Run a disk maintenance tool such as Chkdsk on a Physical Disk resource
Bring a resource online
chkdsk

# Migrate a multinode cluster to a single-node configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To migrate a multinode cluster to a single-node configuration**

1. If your multinode cluster has been upgraded from Windows 2000 or Windows NT 4.0, open Command Prompt, and type the following to create the Local Quorum resource type:

   **cluster restype "local quorum" /create /dll:clusres.dll /type:"local quorum"**

   If your multinode cluster is not an upgraded cluster, skip this step.

2. Open Cluster Administrator.

3. Create a Local Quorum resource on the node you want to configure as the single-node cluster.

   For more information on creating a new resource, see "To create a new resource" in Related Topics.

   For more information on the Local Quorum resource type, see "Local Quorum resource type" in Related Topics.

4. Remove all nodes from the cluster except the node you want to configure as the single-node cluster.

   For more information on removing a node, see "To evict a node from the cluster" in Related Topics.

5. Verify that the single node now owns all resource groups by clicking on **Groups** in the console tree and reviewing the **Owner** column in the details pane.

6. Change the quorum resource to use the Local Quorum resource.

   In the console tree, click the cluster name. On the **File** menu, click **Properties**. On the **Quorum** tab, click **Quorum resource**, and then select **Local Quorum**.

7. Select the appropriate partition.

8. Select the appropriate directory (root path). By default, the cluster database for a single-node cluster is stored at %systemroot%\Cluster\MSCS.

**Important**

- If the quorum resource becomes corrupted, you cannot move it until it is repaired. For more information, see "Node-to-node connectivity problems" in Related Topics.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

- Of the default resource types, only the Physical Disk, Local Quorum, or Majority Node Set resource can be a quorum resource. However, third party vendors can supply other storage class resource types that are quorum capable.

- Cluster disks on a shared bus must be partitioned as master boot record (MBR) and not as GUID partition table (GPT) disks.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Quorum resource
Checklist: Installing a Local Quorum resource
Create a new resource
Local Quorum resource type
Evict a node from the cluster
Node-to-node connectivity problems

# Force quorum in a majority node set server cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To force quorum in a majority node set server cluster

1. Open Cluster Administrator.

2. Stop the Cluster service on all remaining nodes.

   For more information on stopping the Cluster service, see "To stop the Cluster service" in Related Topics.

3. Open Registry Editor.

4. Create a String value called ForceQuorum under the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ClusSvc\Parameters subkey

5. Set the value of ForceQuorum to a comma-separated list of the names of the nodes that are to be part of the quorum set.

   For example, if the cluster that has lost quorum contains nodes "Node5", "Node6" and "Node7", then set up the ForceQuorum registry key as "Node5,Node6,Node7".

   **Caution**

   - If you do not list all the nodes stopped in step 2 in the ForceQuorum key, then you must either power off the nodes not listed or physically disconnect them from the network before performing the next step.

6. Start the Cluster service on the remaining nodes.

   For more information on starting the Cluster service, see "To start the Cluster service" in Related Topics.

### Important

- When you force quorum in a majority node set cluster, an error message from the Service Control Manager may appear, indicating that the Cluster service failed to start. You can safely ignore this message at this time; the Cluster service will continue startup after the end of the startup timeout period for Service Control Manager.

- Do not change your cluster configuration while the cluster is operating in this interim state. For example, do not add or remove nodes, or install or modify applications running in the cluster. Changing the cluster configuration in this interim operating state may result in application data inconsistencies.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your

account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Quorum resource
Checklist: Installing a Local Quorum resource
Stop the Cluster service
Start the Cluster service
Forcequorum command

# Enable Kerberos authentication for virtual servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Kerberos authentication for virtual servers**

1. Open Cluster Administrator.
2. In the console tree, click the Groups folder.
3. Select the appropriate virtual server.
4. In the details pane, click the Network Name resource for that virtual server.
5. On the **File** menu, click **Take Offline**.
6. On the **File** menu, click **Properties**.
7. On the **Parameters** tab, select **Enable Kerberos Authentication**, and then click **OK**.
8. In the details pane, click the Network Name resource for that virtual server.
9. On the **File** menu, click **Bring Online**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- Do not rename a Network Name resource and change the Kerberos authentication setting at the same time. Instead, allow time for changes to replicate after performing each action. Otherwise, clients might not be able to connect to the virtual server that uses the Network Name resource.
- If you enable Kerberos authentication for a virtual server's Network Name resource, make sure that the Cluster service account has appropriate rights to the computer object associated with that Network Name resource. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Virtual servers
Cluster resource

# Manage Server Cluster Groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage groups**

- Create a new group in a cluster
- Delete a group in a cluster
- Bring a group online
- Take a group offline
- Change the properties of a group
- Set group failover policy
- Set group failback policy
- Specify preferred owners of a group
- Move a group to another node
- Rename a group

# Create a new group in a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new group**

1. Open Cluster Administrator.
2. On the **File** menu, point to **New**, and then click **Group**.
3. In the New Group Wizard, in **Name**, type a name for the new group.
4. In **Description**, type any comments you want, and then click **Next**.
5. Under **Available nodes**, click the nodes you want to be the preferred owners for the group, and then click **Add**.

   You can also leave **Preferred owners** empty. This means that it does not matter to which node the group fails back after failover takes place. For more information, see "Failover and failback" in Related Topics.
6. To change the priority of an owner, select the owner, click **Move Up** or **Move Down**, and then click **Finish**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The name that you give the group is for administrative purposes only. It is not the same as a network name, which is the name that clients can use to connect to the group. However, make sure this name is different from any other group name in the cluster.
- Cluster Administrator lists the group description and the group name in the details pane.
- All resources within the new group fail over together. To have the group fail back to a certain node, specify that node as the preferred owner and enable failback. You can balance groups among all nodes to maximize the performance of the cluster. However, you can choose not to have a preferred owner if the location of the group does not greatly affect performance.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster groups
Failover and failback

# Delete a group in a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a group**

1. Open Cluster Administrator.
2. In the console tree, click the Groups folder.
3. In the details pane, click the group you want to remove.
4. Delete or move all resources in the group.

   For more information on deleting or moving resources, see Related Topics.
5. On the **File** menu, click **Delete**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster groups
Delete a resource
Move a resource to another group

# Bring a group online

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To bring a group online**

1. Open Cluster Administrator.
2. In the console tree, double-click the Groups folder.
3. In the details pane, click the group you want.
4. On the **File** menu, click **Bring Online**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- Resources in a group come online in the order of their dependencies.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Take a group offline
Server Cluster groups
Resource dependencies

# Take a group offline

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To take a group offline**

1. Open Cluster Administrator.
2. In the console tree, double-click the Groups folder.
3. In the details pane, click the group you want.
4. On the **File** menu, click **Take Offline**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- Resources in a group go offline in the order of their dependencies.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Bring a group online
Server Cluster groups
Resource dependencies

# Change the properties of a group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the properties of a group**

1. Open Cluster Administrator.
2. In the console tree, click the Groups folder.
3. In the details pane, click the group you want.
4. On the **File** menu, click **Properties**.
5. Click each tab and set the various properties for the group.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The settings on the tabs give general information about the group and define how it performs during failover and failback.
- Every new cluster includes at least one default group: Cluster Group. This group contains the settings for the default cluster and some typical resources that provide generic information and failover policies. Do not delete or rename the Cluster Group or move the resources in it.

  It is important to keep the Cluster Group separate. If another resource is added to this group and the resource fails, it may cause the Cluster Group to fail also, making it impossible to access the cluster using the cluster name.

  If the quorum resource is located in the Cluster Group and that resource group is moved, state change operations for other resource groups will be blocked until the Cluster Group move is completed.

  The availability of the quorum resource takes priority over failover policies set for other resources in the Cluster Group.
- You can use the default groups as templates for creating your own groups. For example, you might add an IP Address resource, Network Name resource, and a Generic Script IIS FTP service resource to Disk Group 1, and then rename the group FTP Group.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Failover and failback
Setting group properties

# Set group failover policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set group failover policy**

1. Open Cluster Administrator.
2. In the console tree, click the Groups folder.
3. In the details pane, click the appropriate group.
4. On the **File** menu, click **Properties**.
5. On the **Failover** tab, type values for **Threshold** and **Period**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The failover policy for a group is the maximum number of times (**Threshold**) that the group is allowed to fail over in the specified number of hours (**Period**) before it is left in a failed state. If a group fails over more often than this, the Cluster service leaves it offline. For example, if a group failover threshold is set to 5 and its failover period to 3, the Cluster service will fail over the group at most five times within a three-hour period. The sixth time a resource in the group fails, the Cluster service fails all other resources in the group and leaves the entire group offline instead of failing over the group.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Failover and failback
Server Cluster groups

# Set group failback policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set group failback policy**

1. Open Cluster Administrator.
2. In the console tree, click the Groups folder.
3. In the details pane, click the appropriate group.
4. On the **File** menu, click **Properties**.
5. On the **Failback** tab, click **Prevent failback** or **Allow failback**.

   If you click **Allow failback**, then either click **Immediately**, or click **Failback between** and set the time interval.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- To set the time interval for **Failback between**, enter numbers between 0 and 23 for the beginning and end of the interval. If the first number is greater than the second, the interval ends on the following day. The numbers correspond to the local time of the cluster group, as read on a 24-hour clock.

  It is important to set the failback time because you may not want failback to occur during hours of peak usage.
- If no preferred owners are specified for the group, then failback does not occur.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Specify preferred owners of a group
Server Cluster groups
Failover and failback

# Specify preferred owners of a group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify preferred owners of a group**

1. Open Cluster Administrator.
2. In the console tree, click the Groups folder.
3. In the details pane, click the appropriate group.
4. On the **File** menu, click **Properties**.
5. On the **General** tab, next to **Preferred owners**, click **Modify**.
6. In the **Modify Preferred Owners** dialog box, enter any changes you want to make:
   - To add one or more preferred owners, under **Available nodes**, not **Preferred owners**, click the nodes you want to add, and then click the right arrow.
   - To remove a preferred owner, under **Preferred owners**, click the nodes you want to remove, and then click the left arrow.
   - To change the priority of a preferred owner, click the node, and then click the up or down arrow.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- If more than one node is listed under **Preferred owners**, failover priority is determined by the order of the list. It also controls failback behavior. If failback is enabled, the group will fail back to the highest-priority node that is available. For more information, see "Determining failover policies for groups" in Related Topics.
- If no preferred owners are specified, failback does not occur.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Setting group properties
Server Cluster groups
Determining failover and move policies for groups

# Move a group to another node

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move a group to another node**

1. Open Cluster Administrator.
2. In the console tree, double-click the Groups folder.
3. In the console tree, click the group you want to move.
4. On the **File** menu, click **Move Group**.
5. In clusters with more than two nodes, select **Best Possible** or select the node to which you want to move the group.

   For more information on the logic used for the **Best Possible** option, see "Determining failover and move policies for groups" in Related Topics.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- After the transfer, the new node owns all resources in the group. Verify that the **Owner** column in the details pane reflects this change.
- A list of the nodes in the cluster appears only in clusters containing three or more nodes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Move a resource to another group
Nodes
Server Cluster groups
Determining failover and move policies for groups

# Rename a group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To rename a group

1. Open Cluster Administrator.
2. In the console tree, expand the Groups folder.
3. In the details pane, click the group you want to rename.
4. On the **File** menu, click **Rename**.
5. In the details pane, type a new name for the group.

   This name is for administrative purposes only.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The group name must be different from any other group name in the cluster.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Manage Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage resources**

- Create a new resource
- Delete a resource
- Bring a resource online
- Take a resource offline
- Run a disk maintenance tool such as Chkdsk on a Physical Disk resource
- View or set resource properties
- View or set resource type properties
- Specify the restart policy for a resource
- Modify resource dependencies
- Specify which nodes can own a resource
- Move a resource to another group
- Rename a resource
- Initiate a resource failure

# Create a new resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create a new resource**

1. Open Cluster Administrator.
2. In the console tree, double-click the Groups folder.
3. In the details pane, click the group to which you want the resource to belong.
4. On the **File** menu, point to **New**, and then click **Resource**.
5. In the New Resource Wizard, type the appropriate information in **Name** and **Description**, click the appropriate information in **Resource type** and **Group**, and then click **Next**.
6. Add or remove possible owners of the resource, and then click **Next**.
7. To add dependencies, under **Available resources**, click a resource, and then click **Add**.

   Or, to remove dependencies, under **Resource dependencies**, click a resource, and then click **Remove**.
8. Repeat step 7 for any other resource dependencies, and then click **Next**.
9. Set resource properties in the *resource***Parameters** dialog box, where *resource* is the name of the resource type.

   Different resource types contain different configuration information in their respective dialog boxes.

   For more information on setting resource properties, see Related Topics.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- Before adding a resource to your cluster, you must verify that the following are true:
  - The type of resource is either one of the basic types provided with the Windows Server 2003 family or a custom resource type provided by the resource vendor.
  - A group already exists within the cluster to which your resource will belong.
  - All dependent resources have been created.
- A resource that fails frequently can affect its associated Resource Monitor and other resources interacting with that Resource Manager. For this reason, a separate Resource Monitor is recommended for any resource that has failed repeatedly in the past.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster Resources
Resource types
Standard Resource Types
View or set resource properties
Setting resource properties
Resource dependencies
Resource Monitors

# Delete a resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete a resource**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the resource you want to remove.
4. On the **File** menu, click **Delete**.

   **Caution**

   - When you delete a resource, Cluster Administrator also deletes all the resources that have a dependency on the deleted resource.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The Cluster service requires that a resource be offline before deleting it. If the resource is online, Cluster Administrator will automatically take the resource offline before deleting it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster Resources
Resource types
Take a resource offline

# Bring a resource online

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To bring a resource online**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the resource you want.
4. On the **File** menu, click **Bring Online**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- Bringing a resource online causes all resources on which the resource depends to be brought online.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster Resources
Resource types
Take a resource offline

# Take a resource offline

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To take a resource offline**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the resource you want.
4. On the **File** menu, click **Take Offline**.

   **Caution**

   - Taking a resource offline causes all resources that depend on that resource to be taken offline.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster Resources
Resource types
Bring a resource online

# Run a disk maintenance tool such as Chkdsk on a Physical Disk resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To run a disk maintenance tool such as Chkdsk on a Physical Disk resource**

1. Determine whether an application is dependent on the Physical Disk resource and whether that application will fail (or will not respond smoothly) if it cannot access or write to the disk. If so, pause or stop the application, or notify users that the application will become unavailable.

2. Open a command prompt.

3. Turn on maintenance mode by typing a command with the following syntax:

   **cluster** [*ClusterName*] **res** *DiskResourceName* **/maint:on**

4. Confirm that maintenance mode is on by typing a command with the following syntax:

   **cluster** [*ClusterName*] **res** *DiskResourceName* **/maint**

5. Run your disk maintenance tool, for example, Chkdsk.

6. When the maintenance tool has finished running, type a command with the following syntax:

   **cluster** [*ClusterName*] **res** *DiskResourceName* **/maint:off**

7. Confirm that maintenance mode is off by typing a command with the following syntax:

   **cluster** [*ClusterName*] **res** *DiskResourceName* **/maint**

8. Check the state of any application that is dependent on the Physical Disk resource. If necessary, restart the applications and inform users that the applications are available.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- You can use maintenance mode only on a Physical Disk resource, and only when that resource is online.

- Maintenance mode is designed to allow maintenance tools to run on a Physical Disk resource without being interrupted by failover. While maintenance mode is on, the resource appears to be online, but your maintenance tool might make the disk unavailable to applications (for example, when you run Chkdsk with the **/f** option to fix disk errors, the disk becomes unavailable). While maintenance mode is on, failures from Looks Alive and Is Alive polling are ignored and therefore the running of the maintenance tool on the resource does not trigger a failover.

- Maintenance mode will remain on until you turn it off or until one of the following occurs:

  - The resource goes offline or fails.

  - The node on which the resource is running restarts or loses communication with other nodes (which causes failover of all resources on that node).

- As soon as any maintenance tools (for example, Chkdsk) have finished running, turn the maintenance mode off.
- An event is logged when maintenance mode is turned on or off.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**Formatting legend**

| Format | Meaning |
| --- | --- |
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Physical Disk resource type
Cluster resource
chkdsk

# View or set resource properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view or set resource properties**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the appropriate resource.
4. On the **File** menu, click **Properties**.
5. Click each tab and set the various properties for the resource.

   For more information, see "Setting resource properties" in Related Topics.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster Resources
Resource types
Server Cluster groups
Nodes
Setting resource properties

# View or set resource type properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view or set resource type properties**

1. Open Cluster Administrator.
2. In the console tree, double-click the Cluster Configuration folder.
3. In the console tree, click the Resource Types folder.
4. In the details pane, click the appropriate resource type.
5. On the **File** menu, click **Properties**, and view or set the various properties for the resource.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Resource types
Setting resource type properties

# Specify the restart policy for a resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify the restart policy for a resource**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the resource you want.
4. On the **File** menu, click **Properties**.
5. On the **Advanced** tab, make the changes you want.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The settings on the **Advanced** tab specify the restart policy for a resource.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Resource failure
Failover and failback

# Modify resource dependencies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To modify resource dependencies**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the resource you want.
4. On the **File** menu, click **Properties**.
5. On the **Dependencies** tab, click the resource whose dependencies you want to set, click **Modify**, and then make the changes you want.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Planning your groups
Failover and failback
Resource dependencies
Setting resource properties

# Specify which nodes can own a resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify which nodes can own a resource**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the appropriate resource.
4. On the **File** menu, click **Properties**.
5. On the **General** tab, next to **Possible owners**, click **Modify**.
6. Specify which nodes you want to be possible owners.
   - To add a possible owner, under **Available nodes**, click a node, and then click the right arrow.
   - To remove a possible owner, under **Possible owners**, click a node, and then click the left arrow.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- A possible owner is a node that is capable of running a resource.
- If you want to make a group capable of failing over, it must have more than one possible owner.
- If a resource has no possible owners, it will not come online. This can be useful if you want to keep a resource offline but you do not want to delete it and lose its configuration settings.
- If the resource DLL for a particular resource type is not available on some nodes, these nodes are automatically removed from the possible owners list for all the resources of this type.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster Resources
Failover and failback
Nodes

# Move a resource to another group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To move a resource to another group**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the resource you want to move.

   **Caution**

   - When you move a resource to another group, all resources that depend on it or on which it depends are also moved.

4. On the **File** menu, point to **Change Group**, and then click the name of the group to which you want to transfer ownership. Alternatively, you can select the resource you want to move, then drag and drop it to another group.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster Resources
Resource types
Server Cluster groups

# Rename a resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a resource**

1. Open Cluster Administrator.
2. In the console tree, click the Resources folder.
3. In the details pane, click the resource you want to rename.
4. On the **File** menu, click **Rename**.
5. Type a new name for the resource.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- Changing the name of a resource does not affect the failover policy for that resource or for the other resources in the group to which it belongs. Also, changing the resource name does not affect the dependencies that other resources might have on the resource.
- The resource name is only used for administrative purposes.
- The resource name must be unique within the same cluster; other clusters may use the same resource name.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server Cluster Resources
Resource types
Failover and failback
Virtual servers
Take a resource offline

# Initiate a resource failure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To initiate a resource failure**

1. Open Cluster Administrator.
2. In the console tree, click the Groups folder, and then click the group that contains the resource to fail.
3. In the details pane, click the resource you want.
4. On the **File** menu, click **Initiate Failure**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- When a resource fails, it automatically restarts if a restart policy is defined for it.
- If a resource fails repeatedly, the entire group to which the resource belongs can fail over, depending on the failover policy defined for the group.
- This command is useful for testing the restart and failover policies of the resource and group.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Resource failure
Setting group properties
Specify the restart policy for a resource
Set group failover policy

# Use Cluster Administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Use Cluster Administrator**

- Connect to a cluster with Cluster Administrator
- Change the maximum quorum-log size
- Rename a cluster
- Start the Cluster Application Wizard
- Start Cluster Administrator in Microsoft Management Console

# Connect to a cluster with Cluster Administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect to a cluster with Cluster Administrator**

1. Open Cluster Administrator.
2. If you are already connected to an existing cluster, on the **File** menu, click **Open Connection**.

   If you are not connected to a cluster, the **Open Connection to Cluster** dialog box will appear.
3. In **Action**, select **Open connection to cluster**. Then in **Cluster or server name**, type the network name or IP address of an existing cluster, or select a name from the drop-down list box, or click **Browse** to search for an available cluster.

### Caution

- If you use Cluster Administrator or cluster.exe to manage a server cluster, the locale settings for the user account must match the system default locale on the computer used to manage the cluster. For more information on changing the system default locale, see "To change number, currency, time, and date settings" in Related Topics.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- When you start Cluster Administrator, connections left open when you closed the previous session are restored. If no connections were left open, Cluster Administrator prompts you to type a cluster name, a node, or an IP address before it displays any information. If you do not want to open old connections, you can start Cluster Administrator from the command line with the **cluadmin /noreconnect** command. For more information on administering a cluster from the command line, see "Managing a server cluster from the command line" in Related Topics.
- You can simultaneously connect to multiple clusters, but be aware that multiple connections can affect the performance of Cluster Administrator.

  To prevent this problem, you can start a new instance of Cluster Administrator for each cluster you want to administer.
- Use the network name or IP address of either the cluster or one of its nodes to open a connection to a cluster. If Cluster Administrator cannot find the cluster by its name, type the name of an active node in the cluster. You can also connect to the cluster by opening Cluster Administrator on a cluster node and entering a PERIOD (.) in **Cluster or server name**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing a Server Cluster from the Command Line

# Change the maximum quorum-log size

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the maximum quorum-log size**

1. Open Cluster Administrator.
2. In the console tree, click the cluster.
3. On the **File** menu, click **Properties**.
4. Click the **Quorum** tab.
5. In **Reset quorum log at**, type the maximum size (in increments of 64 kilobytes) that the quorum log file can be before the cluster resets it.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The cluster records all changes to the cluster database in the quorum log file. When the quorum log attains the specified size, the cluster saves the database and resets the log file.
- By default, the maximum size of the quorum log is 4096 kilobytes. You may want to increase the size of the quorum log if the cluster checkpoint file stored on the quorum disk is larger than 4096 kilobytes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Use a different disk for the quorum resource
Quorum resource

# Rename a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To rename a cluster**

1. Open Cluster Administrator.
2. In the console tree, click the cluster you want to rename.
3. On the **File** menu, click **Rename**.
4. Type the new name for the cluster next to its icon.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The cluster name cannot be the same as the computer name of a cluster node, and it cannot conflict with any other server, cluster, or domain name in use on your network.
- The new name does not take effect until the Cluster Name resource is taken offline and brought back online. The Cluster Name resource is found in the Cluster Group resource group in a newly created cluster.
- The cluster name can also be changed by displaying the properties of the Cluster Name resource, selecting the **Parameters** tab, then clicking on the **Rename** button.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Bring a resource online
Take a resource offline

# Start the Cluster Application Wizard

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start the Cluster Application Wizard**

1. Open Cluster Administrator.
2. On the **File** menu, click **Configure Application**.
3. Follow the instructions in the wizard and provide all the information required to run the application on a server cluster, for example, the virtual server to host the application, the resource to manage your application, resource dependencies, failover and failback policies, and so on as shown in the table below.

| Information required | What it is used for |
|---|---|
| Internet protocol (IP) address | To create a new virtual server for your application, you need an IP address. You do not need an IP address if your application runs in an existing virtual server. |
| Virtual server resource group | Your clients will access your application through a virtual server: a cluster resource group with an IP address and a network name. Using the wizard, you can create a new resource group for a virtual server, use an existing resource group, or use an existing virtual server. |
| Resource type for your application | You can create a resource to manage your application using the wizard. You need to know the resource type for the resource. |
| Application resource name | If you use the wizard to create a resource for your application, you must name the resource. |
| Application resource dependencies | If the application requires other resources to run, you can create resource dependencies using the wizard. |

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Standard Resource Types
Virtual servers
Resource dependencies

# Start Cluster Administrator in Microsoft Management Console

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start Cluster Administrator in Microsoft Management Console**

1. Open Computer Management.
2. In the console tree, double-click to expand **Services and Applications**.
3. Right-click **Cluster**, and then click **Manage**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- Cluster Administrator is not a Microsoft Management Console (MMC) snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing Cluster Administrator on a remote computer
Cluster Hardware Configuration

# Manage Networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage networks**

- Change the name or description of a network
- Enable a network for cluster use
- Disable a network for cluster use
- Change how the cluster uses a network
- Change the description of a network interface
- Change network priority for communication between nodes
- Troubleshoot client-to-virtual server connectivity

# Change the name or description of a network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the name or description of a network**

1. Open Cluster Administrator.
2. In the console tree, double-click to expand **Cluster Configuration**, and then click **Networks**.
3. In the details pane, click the network.
4. On the **File** menu, click **Properties**.
5. To change the name of the network, in **Name**, type a new name.
6. To describe the network, in **Description**, type a new description.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The description provides information about the network. It appears in the details pane of Cluster Administrator when you click the Networks folder. The network description is optional.
- When the name of a network is changed in the cluster, the name of the associated network connection is changed in the Network Connections folder on every node in the cluster.

  Also, if the name of a network connection in the Network Connections folder on any node is changed, the associated network in the cluster will also be changed, along with the network connection on all other nodes.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Server cluster networks
Network Configuration in a Server Cluster
Setting network properties

# Enable a network for cluster use

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable a network for cluster use**

1. Open Cluster Administrator.
2. In the console tree, double-click to expand **Cluster Configuration**, and then click **Networks**.
3. In the details pane, click the network you want to enable.
4. On the **File** menu, click **Properties**.
5. Select the **Enable this network for cluster use** check box.
6. Specify how you want the network to be used by the cluster:
    - To use the network for communications with clients and between nodes, click **All communications (mixed network)**.
    - To use the network only for communications between nodes, click **Internal cluster communications only (private network)**.
    - To use the network only for communications with clients, click **Client access only (public network)**.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- If multiple network interfaces exist on the same subnet, only the first one can be used. Others will be ignored by the Cluster service.
- With Windows Server 2003 with Service Pack 1, if you use Windows Firewall, use the Security Configuration Wizard to configure it. For more information, see "Using Windows Firewall with a Server Cluster" in Related Links.
- If you are protecting your network with a firewall other than Windows Firewall, take note of the following list of ports needed by server clusters:
    - 135 (RPC endpoint mapper/DCOM).
    - 135 (RPC endpoint mapper over UDP).
    - 3343 (used by the Cluster Network Driver).
    - 445 (SMB).
    - 139 (NetBIOS session service).

- For cluster nodes running many services, ports 5000-5099 (or more) might need to be open for remote RPC connectivity to Cluster Administrator. When these ports are closed, event log error 1721 might occur when you connect to a cluster through Cluster Administrator (unless you just type a period in Cluster Administrator).

  This problem can occur because the Cluster service requires at least 100 ports for communication through remote procedure calls (RPC). The number of ports available to the Cluster service can become too small when other services are using some of the necessary ports. Such services include Windows DNS service, Windows Internet Name Service (WINS), the Microsoft SQL Server service, and others.

- If the nodes are separated by a firewall, ports 8011-8031 must be open for internode RPC connectivity. Otherwise, errors in the cluster log will indicate that a "Sponsor" is not available. These errors occur because there are not enough ports available for RPC communication between a node that is trying to join the cluster and a node that can "sponsor" that node.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Disable a network for cluster use
Change how the cluster uses a network
Server cluster networks
Network Configuration in a Server Cluster
Setting network properties
Using Windows Firewall with a server cluster

# Disable a network for cluster use

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To disable a network for cluster use**

1. Open Cluster Administrator.
2. In the console tree, double-click to expand **Cluster Configuration**, and then click **Networks**.
3. In the details pane, click the network you want to disable.
4. On the **File** menu, click **Properties**.
5. Clear the **Enable this network for cluster use** check box.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- After you disable a network for cluster use, the cluster cannot use the network for any communications between nodes or with clients.
- Before you disable a network for cluster use, verify that IP address resources are not configured to use that network. If IP address resources are configured to use the network, you will be unable to clear the **Enable this network for cluster use** check box.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable a network for cluster use
Change how the cluster uses a network
Network Configuration in a Server Cluster
Server cluster networks
Setting network properties

# Change how the cluster uses a network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change how the cluster uses a network**

1. Open Cluster Administrator.
2. In the console tree, double-click to expand **Cluster Configuration**, and then click **Networks**.
3. In the details pane, click the appropriate network.
4. On the **File** menu, click **Properties**.
5. Under **Enable this network for cluster use**, specify how you want the network to be used by the cluster:
   - To use the network for communication with clients and between nodes, click **All communications (mixed network)**.
   - To use the network only for communications between nodes, click **Internal cluster communications only (private network)**.
   - To use the network only for communications with clients, click **Client access only (public network)**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- It is strongly recommended that at least two networks be configured for communication among nodes (either private or mixed). For more information, see "Best practices for configuring and operating server clusters" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enable a network for cluster use
Disable a network for cluster use
Network Configuration in a Server Cluster
Server cluster networks
Setting network properties
Best practices for configuring and operating server clusters

# Change the description of a network interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the description of a network interface**

1. Open Cluster Administrator.
2. In the console tree, double-click to expand **Cluster Configuration**, and then click **Network Interfaces**.
3. In the details pane, click the interface for the network on a node.
4. On the **File** menu, click **Properties**.
5. In **Description**, type information about the network interface on the node.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- The description can provide information about the network interface on a specific node. Your information appears in the details pane when you click the Network Interfaces folder. The description is optional. However, it is strongly recommended that the network interface description reflects its role.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change the name or description of a network
Change how the cluster uses a network
Network Configuration in a Server Cluster
Server cluster networks
Setting network interface properties

# Change network priority for communication between nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change network priority for communication between nodes**

1. Open Cluster Administrator.
2. In the console tree, click the cluster.
3. On the **File** menu, click **Properties**.
4. Click the **Network Priority** tab.
5. Under **Networks used for internal cluster communications**, click a network.
6. Click **Move Up** to increase the network priority.

   - or -

   Click **Move Down** to lower the network priority.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- If multiple networks are configured as private or mixed, you can specify which one to use for internode communication. It is recommended that private networks have higher priority than mixed networks. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change how the cluster uses a network
Setting network properties
Network Configuration in a Server Cluster
Server cluster networks
Setting network properties

# Troubleshoot client-to-virtual server connectivity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To troubleshoot client-to-virtual server connectivity**

Follow the steps below to troubleshoot client-virtual server connectivity issues.

1. Make sure the IP address and dependent Network Name resources are online.

2. Use **ipconfig /all** to examine the TCP/IP configuration for all adapters. Make sure the IP address of the resource is registered to the correct adapter and that each adapter has a primary address. If the resource's IP address is not registered but the resource is online, take the resource offline and then bring it back online. If the address is still not registered with the correct adapter, consult the diagnostic cluster log.

3. If NetBIOS has been enabled for the IP address resource, verify that the network name is properly associated with its dependent IP address resource.

   Use **nbtstat** to display the node's NetBIOS name registrations. For example, type the following at the Windows Server 2003 family command prompt:

   **nbtstat -n**

   The output from this command displays the IP addresses that are registered with NetBIOS. Cluster IP address resources are listed under the heading, **\Device\NetBt_IF***x*, where *x* is a number. Under **NetBIOS Local Name Table**, you will find the NetBIOS names that are associated with the IP addresses. For each dependent network name, there are two entries: *Name***<20>** and *Name***<00>**, where *Name* is the virtual network name.

   The following is an example of this output:

   **\Device\NetBt_If2:**

   **Node ipAddress: [172.31.248.167] Scope Id: []**

   **NetBIOS Local Name Table**

   | Name | Type | Status |
   |---|---|---|
   | CLWCLUS | <20> Unique | Registered |
   | CLWCLUS | <00> Unique | Registered |

   If the NetBIOS names are unregistered, or if there are registration conflicts, use **nbtstat -rr** to release the old NetBIOS names and renew their registrations. For example, type the following at the Windows Server 2003 family command prompt:

   **nbtstat -rr**

4. Verify that DNS is available and functioning correctly.

   Use the **netdiag** command line tool (available in the *Microsoft Windows Server 2003 Resource Kit*) to evaluate the registration of DNS. For example, type the following at the Windows Server 2003 family command prompt:

   **netdiag /test:DNS /d:***xxx.xxx.xxx.xxx***/debug**

   where *xxx.xxx.xxx.xxx* is the IP address of the virtual server.

   If the output suggests that connectivity to the DNS server is suspect or DNS registration for the virtual name may have

failed, check to see that the network connections on the virtual server are configured with the correct IP addresses of the DNS servers to be used for name resolution.

5. Try to ping the group's IP address from all clients and cluster nodes.

   For example, on each client, and all cluster nodes, type the following at the Windows Server 2003 family command prompt:

   **ping** *xxx.xxx.xxx.xxx*

   where *xxx.xxx.xxx.xxx* is the IP address of the virtual server.

   This tests for physical connectivity between the clients and cluster node, and between the cluster nodes.

6. Try to ping the virtual server network name from each client and cluster node.

   For example, at the Windows Server 2003 family command prompt on each client and on all cluster nodes, type the following:

   **ping** *Netname*

   where *Netname* is the network name for the virtual server.

   This test uses the appropriate name-resolution mechanism (DNS or WINS). If you can ping the IP address but ping returns **unknown host** when you ping the network name, then the client is unable to resolve the network name to an IP address. Check connectivity to the name server or servers by pinging their IP addresses (**ipconfig all** will list the addresses for each network adapter). If DNS is in use on the network but the name returned is not a fully qualified domain name (FQDN), (that is, just the NetBIOS name is returned), then connectivity to the DNS server is suspect or DNS registration for the virtual name may have failed.

7. If the virtual server has a File Share resource, try to connect to the resource using **net use**.

   For example, at the Windows Server 2003 family command prompt, type the following:

   **net use** * *Virtualserver**Sharename*

   where *Virtualserver* is the network name for the virtual server and *Sharename* is the share name for the File Share resource.

   If you are unable to connect to the resource, you may not have the appropriate permissions set for the resource. For more information, see "Set permissions on a shared resource" in Related Topics.

8. If clients are unable reconnect to the virtual server after a failover or move, verify that the router/switch between the client and the cluster is able to correctly handle the ARP cache flush protocol. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Ping
Net use
Nbtstat
Ipconfig
Address Resolution Protocol (ARP)
Set permissions on a shared resource

**Other Resources**

Net services overview

# Manage Cluster Hardware

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Manage cluster hardware**

- Install local storage buses and devices
- Connect the shared storage buses and assign drive letters
- Add drives on the shared storage bus
- Add a network adapter
- Install an additional shared storage bus
- Replace the redundant array of independent disks (RAID) controller
- Configure the Always On power scheme without Terminal Services installed
- Configure the Always On power scheme with Terminal Services installed

# Install local storage buses and devices

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install local storage buses and devices**

1. Move all groups off the node on which you want to add the hardware.

   For more information, see "To move a group to another node" in Related Topics.

2. Shut down Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition on the node to which you want to add the hardware, and then turn off the computer.

3. Install and configure the card and devices for the new local storage bus, and then install the appropriate drivers.

4. Turn on the computer and restart Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition on the node to which you added the hardware.

5. If you installed a new drive, partition and format it, assigning new drive letters or mounted drives as appropriate.

6. If you want to use Encrypted File System (EFS) on this local device, enable encryption on the device (using either the drive letter or mount point). Then, create roaming profiles for all users or service accounts that store data on that device.

   For more information, see "To encrypt a file or folder" and "To create a roaming user profile" in Related Topics.

7. In Cluster Administrator, verify that the node has rejoined the cluster.

8. Move the groups referenced in step 1 to the node with the new hardware, as appropriate for your cluster configuration.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign, change, or remove a drive letter
Using NTFS mounted drives
Move a group to another node
Encrypting File System overview
Encrypt a file or folder
Create a roaming user profile

# Connect the shared storage buses and assign drive letters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To connect the shared storage buses and assign drive letters**

1. Install Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition on all nodes.

   **Caution**

   - All nodes in a cluster must be running the same version of the operating system, that is, either Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. You cannot run both Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition on nodes in the same cluster.

2. Prepare the controllers for the shared storage bus, following the instructions in your storage bus owner's manual.

3. Install the controllers for the shared storage bus, but do not connect them to the shared storage bus.

4. Connect the shared storage devices to the shared buses, and then connect the shared storage buses to all nodes.

5. Start the operating system on the first node on which you intend to create the cluster.

6. Format the disks on the cluster storage device as NTFS.

   For more information, see "To format a basic volume" in Related Topics.

   **Caution**

   - Each disk on the cluster storage device must be a basic disk, not a dynamic disk.

7. Partition the disks on the cluster storage device.

   For more information, see "To create a partition or logical drive" in Related Topics.

8. Do one of the following:

   - Assign drive letters to the disks on the cluster storage device.

   - Assign mount points to the disks on the cluster storage device.

     For more information, see "To assign, change, or remove a drive letter" and "To create a mounted drive" in Related Topics.

   **Caution**

   - Make sure that you create unique mounted drives so that they do not conflict with existing local drives on any node in the cluster.

   - Do not create mounted drives between disks on the cluster storage device (cluster disks) and local disks.

   - Do not create a mounted drive from a clustered disk to the cluster disk that contains the quorum resource (the

quorum disk). You can, however, create a mounted drive from the quorum disk to a clustered disk.

- Mounted drives from one cluster disk to another must be in the same cluster resource group, and must be dependent on the root disk.

9. Label all partitions to match their drive letters or mount points assigned in the step above.

10. Use Cluster Administrator to add the first node to the cluster.

    For more information, see "To add the first node to the cluster" in Related Topics.

    **Caution**

    - Completely create the cluster on one node before you join additional nodes to the cluster.

    **Important**

    - When you create a cluster, Physical Disk resources will be automatically created for cluster disks that use drive letters. However, if your cluster disks use mount points, Physical Disk resources will not be automatically created for those disks; you must manually create disk resources for cluster disks that use mount points. For more information on how to do this, see "Checklist: Installing a Physical Disk resource" in Related Topics.

11. Join one or more additional nodes to the cluster.

    For more information, see "To add additional nodes to the cluster" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add additional nodes to the cluster
Physical Disk resource type
Assign, change, or remove a drive letter
Format a basic volume
Create a partition or logical drive
Create a mounted drive
Create a cluster
Add additional nodes to the cluster
Checklist: Installing a Physical Disk resource

# Add drives on the shared storage bus

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add drives on the shared storage bus**

1. Turn off all nodes.

   **Important**

   - Unless your shared drive supports hot swapping, you must shut down and turn off all cluster nodes to add drives. If your shared drive supports hot swapping, turn off all but one node.

2. Add or replace drives on the shared storage bus, following the manufacturer's recommendations on configuring a new drive.

3. Turn on the first node, if it is not already on.

4. On the first node, use Computer Management to:

   - Write a signature to the new disk.
   - Make sure the disk is a basic disk and not a dynamic disk.
   - Format the drive as NTFS.
   - Partition the drive.

     For more information, see "To create a partition or logical drive" in Related Topics.

   - Assign drive letters or mount points to the partitions on the disk that do not conflict with any other local or clustered drive letters or mount points already in use.

     For more information, see "To assign, change, or remove a drive letter" and "To create a mounted drive" in Related Topics.

   - Label each partition to match its corresponding drive letter or mount point.

5. Create a Physical Disk or other storage-class type resource for the new disk.

   For information on how to do this, see "Checklist: Installing a Physical Disk resource", and "To create a new resource" in Related Topics.

6. Bring the new resource online.

   For information on how to do this, see "To bring a resource online" in Related Topics.

7. Restart all remaining nodes.

8. Make sure the new drive is functioning correctly by placing the Physical Disk resource in a separate group and successively failing that group over to the other nodes in the cluster.

   For information on how to do this, see "To test whether group resources can fail over" in Related Topics.

**Notes**

- To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
- On a 2-node SCSI bus, the SCSI host adapters must each have a unique SCSI ID (usually the numbers 6 and 7). In addition, the SCSI bus must be properly terminated.
- On a Fibre Channel bus, devices can be added at any time.
- Cluster disks on a shared bus must be partitioned as Master Boot Record (MBR) and not as GUID partition table (GPT) disks.
- For a quorum resource, you cannot enable Shadow Copies of Shared Folders.
- For mounted drives, do not create mount points between clustered and local disks and do not create mount points to the Quorum disk. If you assign mount points from one shared disk to another, the disks must be in the same resource group and the mounted drive must be dependent on the root disk.
- You can enable Encrypted File System (EFS) on this cluster device and make encrypted directories accessible to clients. For more information, see "To create a cluster-managed encrypted file share" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign, change, or remove a drive letter
Create a partition or logical drive
Format a basic volume
Physical Disk resource type
Node-to-node connectivity problems
Checklist: Installing a Physical Disk resource
Create a new resource
Bring a resource online
Test whether group resources can fail over
Create a mounted drive
Create a cluster-managed encrypted file share

# Add a network adapter

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add a network adapter**

1. Shut down the operating system on the node to which you want to add the hardware, and then turn off the computer.
2. Install the adapter.
3. Turn on the computer and restart the operating system on the node to which you added the hardware.
4. Load the driver for the network adapter.

### Note

- Confirm that your hardware is designed for Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition by reviewing the compatibility information in Support resources.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Install an additional shared storage bus

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install an additional shared storage bus**

1. Shut down the operating system on all nodes and turn off all computers.
2. Prepare one controller for the shared storage bus, referring to your storage bus owner's manual for instructions.

   Install the controllers for the shared storage bus, but do not connect them to the shared storage bus.

   For SCSI controllers, make sure that the SCSI controllers for the shared SCSI bus are using different SCSI IDs. SCSI controller IDs are preset to 7. Reset one SCSI controller ID to 6.
3. Connect the shared storage devices to the shared bus, and then connect the shared storage bus to all nodes.
4. Start the operating system on one node. Turn on a second node, and press F8 when the operating system Loader screen appears (not allowing the operating system to start).
5. On the first node, use Disk Management to assign drive letters to the disks available on the new shared storage bus.

   **Where?**

   - Computer Management/Storage/Disk Management
6. Make sure the Cluster service is running on this node.
7. Create new resource groups if needed and install a Physical Disk or other storage-class type resource for each new disk.

   For more information see, "Checklist: Creating a new group", "To create a new group", "Checklist: Creating a new resource", and "To create a new resource" in Related Topics.
8. Shut down the first node, but do not turn off the computer.
9. On a second node, select the operating system and press ENTER.
10. Restart the first node.

**Caution**

- Do not run the operating system on all nodes at the same time until you have completed this procedure.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assign, change, or remove a drive letter
Checklist: Creating a new group
Create a new group in a cluster

Checklist: Creating a new resource
Create a new resource

# Replace the redundant array of independent disks (RAID) controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To replace the redundant array of independent disks (RAID) controller**

- Replacing the RAID controller and updating controller configuration is hardware dependent. Refer to your hardware vendor for information on replacing your RAID controller and the effect that removal of the controller from the shared bus will have on termination.

**Note**

- Replacing the RAID controller may affect the quorum disk. It is therefore advisable to back up the individual nodes and the cluster disks before replacing the RAID controller. If the quorum disk is affected, it can then be restored from the backup.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Hardware redundant array of independent disks (RAID)
Node-to-node connectivity problems

# Configure the Always On power scheme without Terminal Services installed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the Always On power scheme without Terminal Services installed**

1. Open Control Panel.
2. Double-click **Power Options**.
3. In the **Power Options Properties** dialog box, under **Power schemes**, verify that **Always On** is selected.
4. Under **Settings for Always On power scheme**, verify that **Turn off hard disks** is set to **Never**.
5. Verify that **System standby** is set to **Never**.
6. Verify that **System hibernates** is set to **Never**.

### Important

- APM/ACPI Power saving features must not be enabled on server cluster members. A cluster member that turns off disk drives or enters "system standby" or "hibernate" mode can initiate a failure in the cluster. If multiple cluster nodes have power saving enabled, this can result in the entire cluster becoming unavailable.

### Note

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure the Always On power scheme with Terminal Services installed
**Other Resources**
Power Options overview

# Configure the Always On power scheme with Terminal Services installed

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the Always On power scheme with Terminal Services installed**

1. Open Control Panel.
2. Double-click **Power Options**.
3. In the **Power Options Properties** dialog box, under **Power schemes**, verify that **Always On** is selected.
4. Under **Settings for Always On power scheme**, verify that **Turn off hard disks** is set to **Never**.

### Important

- APM/ACPI Power saving features must not be enabled on server cluster members. A cluster member that turns off disk drives can initiate a failure in the cluster. If multiple cluster nodes have power saving enabled, this can result in the entire cluster becoming unavailable.

### Notes

- To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.
- Installing Terminal Services on a system reduces the power management options available to the user. The **System standby** and **System hibernates** options are not available.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configure the Always On power scheme without Terminal Services installed
**Other Resources**
Power Options overview

# Back Up and Restore Server Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Back up and restore server clusters**

- Back up cluster disk signatures and partition layouts
- Back up the cluster quorum
- Back up data on cluster nodes
- Restore the cluster database on a local node
- Restore the contents of a cluster quorum disk for all nodes in a cluster
- Restore a damaged cluster node using Automated System Recovery

# Back up cluster disk signatures and partition layouts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up cluster disk signatures and partition layouts**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click **Advanced Mode** on the Backup or Restore Wizard.

3. On the **Tools** menu, click **ASR Wizard**.

4. Follow the instructions that appear on your screen.

### Important

- You will need a blank 1.44 MB floppy disk to save your system settings and media to contain the backup files. If your computer does not have a floppy drive, see "Troubleshooting" in Related Topics.

### Notes

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- By default, Backup Operators do not have the user rights necessary to create an Automated System Recovery (ASR) backup on a cluster node. However, Backup Operators can perform this procedure if that group is added to the security descriptor for the Cluster service. You can do that using Cluster Administrator or cluster.exe. For more information, see Give a user permissions to administer a cluster and Cluster.

- Automated System Restore not only backs up cluster disk signatures and partition layouts, it also backs up the system state and other data associated with the operating system.

- You later need to restore the signature of the quorum disk, for example, if you experience a complete system failure, and the signature of the quorum disk has changed since you last backed up, or if you experience Event ID 1034 errors when replacing a disk.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Backing up and restoring server clusters
Backup Troubleshooting

# Back up the cluster quorum

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up the cluster quorum**

1. Open Cluster Administrator.
2. If you are not already connected to a cluster, connect to one.

   For more information, see "To connect to a cluster with Cluster Administrator" in Related Topics.
3. In the console tree, click the cluster node.
4. On the **File** menu, click **Properties** and verify that the Cluster service is running on the node on which you are performing the backup.
5. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.
6. Click **Advanced Mode** on the Backup or Restore Wizard.
7. Click the **Backup** tab, then in **Click to select the check box for any drive, folder, or file that you want to back up** click the box next to **System State**.

   This will back up the contents of the cluster quorum along with other system state data.

   For more information, see "To back up System State data" in Related Topics.

### Important

- It is highly recommended that you use the Backup utility as described in this procedure to perform regular backups of the cluster quorum. The cluster quorum contains information that is indispensible to the operation of the server cluster, for example, the current cluster configuration, application registry checkpoints, and cluster recovery log.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.
- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
- During backup, the following error may be displayed: "Completed with Skipped Files." This error may be ignored. The Mscs directory from the Cluster quorum drive is still successfully backed up.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Backing up and restoring server clusters

Connect to a cluster with Cluster Administrator

Back up System State data

# Back up data on cluster nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To back up data on cluster nodes**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Follow the wizard instructions, selecting **All information on this computer** on the **What to Back Up** page.

   This will back up all data on the node and all the cluster disks currently owned by that node. Cluster disks not owned by that node will not be backed up.

### Important

- You will need a blank 1.44 MB floppy disk to save your system settings and media to contain the backup files. If your computer does not have a floppy drive, see "Troubleshooting the Backup Utility" in Related Topics.

- If a cluster disk owned by the node being backed up fails over to another node during the backup process, the backup set will not contain a full backup of that disk.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Backing up and restoring server clusters
Backup Troubleshooting

# Restore the cluster database on a local node

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore the cluster database on a local node**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click **Advanced Mode** on the Backup or Restore Wizard.

3. Click the **Restore and Manage Media** tab, and then in **Click to select the check box for any drive, folder, or file that you want to restore**, click the box next to **System State**.

4. Click **Start Restore**.

   This will stop the Cluster service on the local node and restore the cluster configuration (the cluster registry file ClusDB located in the *systemroot*\cluster subdirectory) for that node.

   For more information on backup and restore procedures, see "Backing up and restoring data" in Related Topics.

5. Reboot the node when prompted.

   Rebooting the operating system replaces the cluster configuration file currently in use with the restored copy.

6. Verify that the Cluster service restarted on the local node.

   **Where?**

   - Computer Management/Services and Applications/Services

**Important**

- If the signature of the quorum disk has changed since you last backed up, instead of restoring the System State only, use Automated System Recovery in the Backup or Restore Wizard to restore the quorum disk signature along with the System State.

  For information on how to use Automated System Recovery, see "To recover from a system failure using Automated System Recovery" in Related Topics.

- If you do not have a backup set available, you can manually restore the cluster database on the local node by copying the *systemroot*\cluster\CLUSDB file from another node. If the copy operation fails, follow the troubleshooting steps in General administrative problems.

**Notes**

- You must be an administrator or a backup operator to restore System State data. For more information on permissions or user rights, see Related Topics.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- You can only restore the cluster quorum disk on a local node. You cannot restore the cluster quorum disk on a remote computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Backing up and restoring server clusters
System State data
Set advanced restore options
Using Network Monitor
Backup
Checklists for Backing Up and Restoring Data
Recover from a system failure using Automated System Recovery

# Restore the contents of a cluster quorum disk for all nodes in a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore the contents of a cluster quorum disk for all nodes in a cluster**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled.

2. Click **Advanced Mode** on the Backup or Restore Wizard.

3. Click the **Restore and Manage Media** tab, and then in **Click to select the check box for any drive, folder, or file that you want to restore**, click the box next to **System State**.

4. Click **Start Restore**.

5. On the **Confirm Restore** dialog box, click **Advanced**.

6. Select **Restore the Cluster Registry to the quorum disk and all other nodes**, then click **OK**.

7. This action will stop the Cluster service on the local node and restore the cluster configuration (the cluster registry file ClusDB located in the *systemroot*\cluster subdirectory) for that node.

   For more information on backup and restore procedures, see "Backing up and restoring data" in Related Topics.

8. Reboot the node when prompted.

   This will restart the Cluster service on the local node and stop the Cluster service on all other nodes in the cluster. Backup will then copy the restored quorum data from the restarted node to the cluster quorum disk and all other nodes in the cluster.

9. Start the Cluster service on the other nodes in the cluster.

   For more information on starting the Cluster service, see "Start the Cluster service" in Related Topics.

10. Verify that the Cluster service restarted on all nodes.

    **Where?**

    - Computer Management/Services and Applications/Services

**Important**

- Backup will stop the Cluster service on all other nodes of the server cluster after the node that was restored reboots. The entire server cluster will therefore be down while the restored quorum data is copied to the quorum disk and the other nodes in the cluster.

**Notes**

- You must be an administrator or a backup operator to restore System State data. For more information on permissions or user rights, see Related Topics.

- To start Backup, click **Start**, point to **All programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Backing up and restoring server clusters
System State data
Set advanced restore options
Using Network Monitor
Backup
Checklists for Backing Up and Restoring Data
Start the Cluster service

# Restore a damaged cluster node using Automated System Recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To restore a damaged cluster node using Automated System Recovery**

1. Make sure you have the following on hand before you begin the recovery procedure:

   - Your previously created Automated System Recovery (ASR) floppy disk and ASR backup media. The ASR disk and media contain the backed up cluster disk signatures, partition layouts, and cluster quorum (as part of the system state).
   - Your previously created backup media containing the data on the cluster nodes.
   - The original operating system installation CD.
   - If you have a mass storage controller and you are aware that the manufacturer has supplied a separate driver file for it (different from driver files available on the Setup CD), obtain the file (on a floppy disk) before you begin this procedure.

2. Insert the original operating system installation CD into the CD drive of the damaged cluster node.

   **Important**

   - If you are restoring a disk signature to a damaged cluster disk, power down all other cluster nodes except the one on which you are performing the ASR restore. This cluster node must have exclusive rights to the damaged cluster disk.

3. Restart your computer. If you are prompted to press a key to start the computer from CD, press the appropriate key.
4. If you have a separate driver file as described in step 1, to use the driver as part of Setup, press F6 when prompted.
5. Press **F2** when prompted during the text-only mode section of Setup. You will be prompted to insert the ASR floppy disk you have previously created.
6. Follow the directions on the screen.
7. If you have a separate driver file as described in step 1, press F6 (a second time) when prompted after the system reboots.
8. Follow the directions on the screen.
9. Once you have completed all the restore steps, the restored node can rejoin the cluster.

**Notes**

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- This procedure is useful if your cluster node is damaged because of corrupt or missing system files or Cluster registry files, hard disk failure, or changes to disk signatures due to a disk replacement (Event ID 1034 errors).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Restore a damaged cluster node using Automated System Recovery
Back up cluster disk signatures and partition layouts
Using Network Monitor
Backup
Checklists for Backing Up and Restoring Data

# Test Cluster Failures and Failover Policies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Test cluster failures and failover policies**


- Test group failover policy
- Test whether group resources can fail over
- Test private network failures
- Test node failure

# Test group failover policy

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test group failover policy**

1. Open Cluster Administrator.
2. In the console tree, double-click the Groups folder.
3. In the details pane, click the group.
4. On the **File** menu, click **Properties**.
5. On the **Failover** tab, set the **Threshold** to 0, and then click **OK**.
6. In the Cluster Administrator console tree, click the group.
7. In the details pane, click a resource.
8. On the **File** menu, click **Properties**.
9. On the **Advanced** tab, select the **Affect the group** check box, if it is not already selected.
10. On the Advanced tab, set the **Threshold** to 0, and then click **OK**.
11. On the **File** menu, click **Initiate Failure**.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Set group failover policy
Failover and failback
Server Cluster groups
View or set resource properties

# Test whether group resources can fail over

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test whether group resources can fail over**

1. Open Cluster Administrator.
2. In the console tree, double-click the Groups folder.
3. In the console tree, click a group.
4. On the **File** menu, click **Move Group**.

   On a multinode cluster server, when using **Move Group**, select the node to move the group to.

   Make sure the **Owner** column in the details pane reflects a change of owner for all of the group's dependencies.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Resource dependencies
Server Cluster groups

# Test private network failures

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test private network failures**

1. Open Cluster Administrator.

2. Verify that there are at least two networks set up for internal cluster communication and determine which network has the highest priority for internal cluster communications.

   For more information, see "To change network priority for communication between nodes" in Related Topics. When you follow that procedure, Cluster Administrator will only list those networks that support internal cluster communication.

3. If there are two or more internal networks, continue to step 4, otherwise stop here and do not continue.

4. On one node, unplug the cable for the highest priority network identified above.

5. In the console tree, click the cluster node and verify that the state of the node is still **Up** (that is, actively participating in all cluster operations).

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

- To open Cluster Administrator, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Cluster Administrator**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Change network priority for communication between nodes

# Test node failure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test node failure**

- Test node failure using one of the following four methods:

  - Stop the Cluster service on the node.

    This prevents clients from accessing cluster resources through that node. In this event, all resources owned by this node fail over to other nodes in the server cluster. For more information on stopping the Cluster service, see Related Topics.

  - Turn off the power on the node.

    This tests the ability of the server cluster to fail over all the resources that were owned by the powered down node.

  - On one node, unplug the network cables for all networks enabled for intracluster communications.

    This test disables all intracluster network interfaces for the node and prevents the node from exchanging heartbeats with the Cluster Network Driver on the other active nodes. In this event, the Cluster service then initiates failover of the resources owned by the node to the other nodes in the server cluster.

### Note

- Disconnecting the SCSI or fiber cable between the host bus adapter and the cluster storage device is not a supported node failure test.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Stop the Cluster service
Nodes

# Server Cluster Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section provides general background information about server clusters:

- Server clusters overview
- Understanding Server Clusters
- Planning Your Server Cluster
- Using Server Clusters
- Administering Server Clusters
- Resources for Server Cluster information

# Server clusters overview

Updated: April 10, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server clusters overview**

A server cluster is a group of independent computer systems, known as nodes, working together as a single system to ensure that critical applications and resources remain available to clients. These nodes must be running Microsoft® Windows Server™ 2003, Enterprise Edition or Microsoft® Windows Server™ 2003, Datacenter Edition. Clustering allows users and administrators to access and manage the nodes as a single system rather than as separate computers. For more information about nodes, see Nodes.

A server cluster can consist of up to eight nodes and may be configured in one of three ways: as a single node server cluster, as a single quorum device server cluster, or as a majority node set server cluster. For more information about these three server cluster models, see Choosing a Cluster Model.

Every node may be attached to one or more cluster storage devices. For most versions of Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, the choices for cluster storage include iSCSI, Serial Attached SCSI, parallel SCSI, and Fibre Channel. The following table provides details about the storage you can use with each version of the operating system, along with the maximum number of nodes you can have with each storage type:

| **Description of Operating System** | **Storage** | **Maximum Nodes** |
|---|---|---|
| Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition<br><br>Versions:<br>• Windows Server 2003<br>• Windows Server 2003 with Service Pack 1 (SP1)<br>• Windows Server 2003 R2<br><br>Platforms:<br>• x86 and x64 (not Itanium) | Parallel SCSI | 2 |
| Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition<br><br>Versions:<br>• Windows Server 2003<br>• Windows Server 2003 with SP1<br>• Windows Server 2003 R2<br><br>Platforms:<br>• x86, x64, and Itanium | Fibre Channel | 8 |

| | | |
|---|---|---|
| Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition<br><br>Versions:<br>• Windows Server 2003 with SP1<br>• Windows Server 2003 R2<br><br>Platforms:<br>• x86, x64, and Itanium | iSCSI<br><br>or<br><br>Serial Attached SCSI (SAS) | 8 |

The following figure shows a four-node cluster that provides high availability for four different types of applications or services:




A server cluster runs several pieces of software that fall into two categories: the software that makes the cluster run (clustering software) and the software that you use to administer the cluster (administrative software). By default, all clustering and administration software files are automatically installed on your computer when you install any operating system in the Microsoft® Windows Server 2003 family of products.

**Important**

- Only computers running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition can be cluster nodes.

**Clustering software**

The clustering software enables the nodes of a cluster to exchange specific messages that trigger the transfer of resource operations at the appropriate times. There are two main pieces of clustering software: the Resource Monitor and the Cluster service. The Resource Monitor facilitates communication between the Cluster service and application resources. The Cluster service runs on each node in the cluster and controls cluster activity, communication between cluster nodes, and failure operations. When a node or application in the cluster fails, the Cluster service responds by restarting the failed application or dispersing the work from the failed system to the remaining nodes in the cluster. For more information about this process, see Failover and failback.

**Administrative software**

Administrators use cluster management applications to configure, control, and monitor clusters. The Windows Server 2003 family provides Cluster Administrator for this purpose. Any computer running Microsoft Windows NT® version 4.0 Service Pack 3 or later, regardless of whether it is a cluster node, can install Cluster Administrator. By default, a copy of Cluster Administrator is automatically installed on your computer when you install Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition. For information about remote administration, see Installing Cluster Administrator on a remote computer.

You can also create, configure, and administer clusters using the **cluster** command. For more information, see Administering

Server Clusters. You can use or create custom administration tools developed using the Cluster automation interfaces. For more information about the Cluster automation interfaces, see the *Microsoft Platform Software Development Kit* (SDK), available separately.

Administrators organize cluster resources into functional units, called groups, and assign these groups to individual nodes. If a node fails, the Cluster service transfers the groups that were being hosted by the node to other nodes in the cluster. This transfer process is called failover. The reverse process, failback, occurs when the failed node becomes active again, and the groups that were failed over to other nodes are transferred back to the original node. For more information about resources, see Server Cluster Resources. For more information about groups, see Server Cluster groups.

# Understanding Server Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding server clusters**

Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition use clustering to provide high-availability application support. A server cluster is a group of independent computer systems, called nodes, working together as a single system. Server clusters allow users and administrators to access and manage the nodes as a single system rather than as separate computers.

The following sections describe how Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition use server clusters to provide high availability, easy manageability, and enhanced scalability:

- Cluster application types
- Server Cluster Components
- Cluster Objects
- Server Cluster Networking

# Cluster application types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster application types**

Applications that run in a server cluster fall into one of four categories:

- **Cluster-unaware applications**

  These types of applications do not interact with the server cluster at all but can still fail over. Failure detection is limited. The Cluster service protects these applications mainly against hardware failures.

- **Cluster-aware applications**

  These types of applications are characterized by superior failure detection. The Cluster service can protect these applications not only against hardware but also against software failures.

- **Cluster management applications**

  These types of applications, which include Cluster Administrator and Cluster.exe, allow administrators to manage and configure clusters. For more information, see Server Cluster Components.

- **Custom resource types**

  Resource types provide customized cluster management and instrumentation for applications, services, and devices. For more information, see Resource types.

For more information on the concepts in this topic, see:

- Cluster Administrator

- Administering Server Clusters

# Server Cluster Components

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server cluster components**

Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition use the server cluster components to create server clusters that provide high availability, easy manageability, and enhanced scalability. These components work together to manage the cluster objects.




The diagram shows how the cluster components relate to three types of applications and to each other within a single cluster node.

For definitions of these application types, see Cluster application types.

For more information on cluster objects, see Cluster Objects.

For more information on the individual server cluster components, see:

- Cluster service
- Resource Monitors
- Resource dynamic-link libraries (DLLs) in clusters running Windows Server 2003
- Cluster Administrator
- Cluster Automation Server
- Cluster database
- Network and disk drivers

# Cluster service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster service**
The Cluster service manages all cluster-specific activity. The Cluster service is installed at the same time as the operating system. One instance of the Cluster service runs on every node in a cluster. The Cluster service handles the following operations:

- Manages cluster objects and configuration
- Manages local restart policy
- Coordinates with other instances of the Cluster service in the cluster
- Handles event notification
- Facilitates communication among other software components
- Performs failover operations

For more information on cluster objects, see Cluster Objects.

# Resource Monitors

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource Monitors**

A Resource Monitor is a cluster component that acts as an intermediary between the Cluster service and a resource dynamic-link library (DLL). When the Cluster service makes a request of a resource, the Resource Monitor transfers that request to the appropriate resource DLL. When a resource DLL must report status or notify the Cluster service of an event, the Resource Monitor ensures that the information is delivered successfully. Resource Monitors are passive communication layers; they do not initiate operations.

The Resource Monitor runs in a process separate from the Cluster service for the following reasons:

- To protect the Cluster service from resource failures.
- To allow the Resource Monitor to take action if the Cluster service fails. The Resource Monitor detects Cluster service failure and responds by taking all of the resources and groups on the affected node offline.

Every cluster node runs one or more Resource Monitors. By default, the Cluster service starts only one Resource Monitor to interact with all of the resources hosted by the node. However, this default can be overridden by an administrator using either Cluster Administrator or another management application. Assign potentially problematic resource DLLs to their own Resource Monitor to prevent faulty behavior from affecting other resources and other Resource Monitors.

- For more information on the Cluster service, see Cluster service.
- For more information on Cluster Administrator, see Cluster Administrator.
- For more information on the resource DLLs, see Resource dynamic-link libraries (DLLs) in clusters running Windows Server 2003.

# Resource dynamic-link libraries (DLLs) in clusters running Windows Server 2003

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource dynamic-link libraries (DLLs)**

This topic provides a brief description of resource DLLs in server clusters running Windows Server 2003.

Resource DLLs contain instructions that manage cluster resources of one or more types. The most important role of the resource DLL is to detect application failures.

When the Cluster service issues a request to a resource, that request is passed to the Resource Monitor. The Resource Monitor, in turn, calls an entry point function in the resource DLL responsible for the resource.

Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition define several resource types and provide resource DLLs to manage them. However, a server cluster is not limited to the resource types provided with the operating system. Developers can create their own resource DLLs to replace or augment the resource DLLs provided with the operating system. Custom resource DLLs have detailed knowledge of the resources they manage and are used by resource monitors to determine the resource state (dead/alive).

# Cluster Administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster Administrator**
Administrators use cluster management applications to configure, control, and monitor clusters. Cluster Administrator is one such application and exists for this purpose. On computers that are not cluster nodes, for example Windows XP Professional, you install Cluster Administrator using Windows Server 2003 Administration Tools. On computers running the Windows Server 2003, Standard Edition family of products, Cluster Administrator is installed, along with the Cluster service, when you install the operating system.

**Important**

- Only computers running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition can be cluster nodes.

For more information, see Installing Cluster Administrator on a remote computer.

You use Cluster Administrator to create the cluster, join nodes to the cluster, manage cluster objects, establish groups, initiate failover, handle maintenance, and monitor cluster activity through a convenient graphical interface. Developers providing their own resource types can write an extension that enables Cluster Administrator to manage their custom types. For more information on developing extensions, see the Microsoft Platform Software Development Kit (SDK).

For more information on using Cluster Administrator to manage your server cluster, see Using Cluster Administrator.

# Cluster Automation Server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster Automation Server**

The Cluster Automation Server exposes a set of Component Object Model (COM) objects to any scripting language that supports automation, such as Visual Basic or the Windows Script Host (WSH). By enabling object-oriented design and the use of high-level languages, the Cluster Automation Server simplifies the process of creating a cluster-management application.

For information on how to write scripts using the Cluster Automation Server, see the *Microsoft Platform Software Development Kit (SDK)*. For more information on WSH, see Windows Script Host. For more information on cluster-management applications, see Cluster application types.

# Cluster database

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster database**
The cluster registry is stored in the cluster database on each node. By default, the cluster database for a single quorum device server cluster is located at %systemroot%\Cluster\CLUSDB; for a majority node set server cluster, the cluster database is located at %systemroot%\Cluster\MNS.%ResourceGUID%$\%ResourceGUID%$\MSCS. A copy of the cluster registry is also stored in the CHKxxx.TMP file on the quorum disk. The quorum log maintains consistency between the two versions of the cluster registry. The cluster database resides in the Windows Server 2003 family registry on each cluster node. It contains information about all physical and logical elements in a cluster, including cluster objects, their properties, and configuration data.

**Important**

- Do not use the **net use** command to connect to the share containing the cluster database files (%systemroot%\Cluster\MNS.%ResourceGUID%$\%ResourceGUID%$\MSCS) in a majority node set server cluster. Otherwise, the Majority Node Set resource (the quorum resource) may be unable to connect to the share and may fail to come online.

Each node's Cluster service maintains a consistent, updated image of the cluster database through the following mechanisms:

- **Global updates**. The Cluster service replicates changes to the cluster database uniformly across all nodes.

  Changes to the cluster database are replicated across all nodes automatically every 4 hours and also when any of the following events occur:

  1. A resource is taken online.
  2. A resource is taken offline.
  3. The Resource Monitor detects a change to one of the registry keys.

- **Periodic checkpointing**. The Cluster service on the node that owns the quorum resource regularly saves a copy of the cluster database to ensure consistency.

- **Quorum resource**. The Cluster service ensures that the quorum resource's recovery log contains the most recent cluster database information. Any changes between checkpoints are logged to the quorum disk if the Cluster service fails to replicate it among all nodes. The main purpose of the quorum resource is to maintain cluster consistency.

For more information on the concepts in this topic, see:

- Nodes
- Cluster Objects
- Cluster service
- Quorum resource

# Network and disk drivers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network and disk drivers**

**Network drivers**
Each node in the cluster runs an instance of the Cluster Network Driver, Clusnet.sys. The Cluster Network Driver is responsible for:

- Providing highly available intracluster communication by monitoring and leveraging all network paths between nodes.
- Routing messages for intracluster traffic.
- Detecting communication failure. Each node's Cluster Network Driver exchanges periodic messages, called heartbeats, with the Cluster Network Driver on the other active nodes. If a node fails to respond to a heartbeat message, the Cluster Network Driver on the node that detects the failure notifies the Cluster service that the corresponding network interface has failed. If all other intracluster network interfaces for that node fail, the Cluster service will then initiate failover of the node. Note that failure of a node to respond to a heartbeat message can be caused by node failure, network interface failure, or network failure.

The Cluster Network Driver on each node uses the Internet Assigned Numbers Authority (IANA)-assigned port of 3343. For information on removing malicious processes bound to the cluster network ports, see information about the Cluster service failing to start and returning ERROR_SHARING_VIOLATION in General administrative problems.

**Disk drivers**
Each node in the cluster runs an instance of the Cluster Disk Driver, Clusdisk.sys. The Cluster Disk Driver is responsible for:

- Maintaining exclusive ownership of cluster disks. Only the node that owns the physical disk resource can access the disk. All other nodes cannot access this disk.
- Placing reservations on disks so that they cannot be accessed by another node. If a node owns a particular disk and that disk is online, then reservations placed by Clusdisk will prevent other nodes from accessing that disk. For disks that are offline, Clusdisk will prevent disk access by blocking I/O on the local system.

# Cluster Objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster objects**
Cluster objects are the physical and logical units managed by the Cluster service. For more information on Cluster service, see Cluster service. Each object is associated with:

- One or more properties, or attributes, that define the object and its behavior within the cluster.
- A set of cluster-control codes used to manipulate the object's properties.
- A set of object-management functions used to manage the object through the Cluster service.

This section covers:

- Server cluster networks
- Server cluster network interfaces
- Nodes
- Resource types
- Server Cluster groups
- Server Cluster Resources

# Server cluster networks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server cluster networks**

A network (sometimes called an interconnect) performs one of the following roles in a cluster:

- A *private network* carries internal cluster communication. The Cluster service authenticates all internal communication, but administrators who are particularly concerned about security can restrict internal communication to physically secure networks.
- A *public network* provides client systems with access to cluster application services. IP Address resources are created on networks that provide clients with access to cluster services.
- A *mixed* (public-and-private) *network* carries internal cluster communication and connects client systems to cluster application services.
- A network that is not enabled for use by the cluster (that is, neither public nor private) carries traffic unrelated to cluster operation.

Make certain that the private network is at the top of the Network Priority list for internal node-to-node communication in the cluster. For more information, see Change network priority for communication between nodes.

**Tolerating network failure**

The Cluster service uses all available private and mixed networks for internal communication. Configure multiple networks as private or mixed to protect the cluster from a single network failure.

If there is only one such network available and it fails, the cluster nodes stop communicating with each other. When two nodes are unable to communicate, they are said to be partitioned. After two nodes become partitioned, the Cluster service automatically shuts down on one node to guarantee the consistency of application data and the cluster configuration. This can lead to the unavailability of all cluster resources.

For example, if each node has only one network adapter, and the network cable on one of the nodes fails, each node (because it is unable to communicate with the other) attempts to take control of the quorum resource. There is no guarantee that the node with a functioning network connection will gain control of the quorum resource. If the node with the failed network cable gains control, the entire cluster is unavailable to network clients.

However, if each node has at least two networks and both networks are configured for internal communication (for example both are mixed networks), the Cluster service can tolerate network failures. In this scenario, the Cluster service can detect a public network adapter failure and fail over all resources that depend on that adapter (through its IP address) to a node where this network is available. This is accomplished because the private network is still functioning properly. If, on the other hand, an adapter on the private network fails, the Cluster service can use the public network for internal communication. This is accomplished because the public network is mixed, allowing both internal and client traffic.

Note that there must be a unique subnet for each network in the Cluster.

**Node-to-node communication**

The Cluster service does not use public-only networks for internal communication, even if a public network is the only available network. For example, suppose a cluster has Network A configured as private and Network B configured as public, and Network A fails. The Cluster service does not use Network B because it is public only; thus, the nodes stop communicating and the cluster breaks apart.

For more information on the concepts in this topic, see:

- Nodes
- Quorum resource
- Network state management and failure detection

- Network communication between nodes
- Manage Networks
- Server cluster network interfaces

# Server cluster network interfaces

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network interfaces**

A network interface is a card or other network adapter that connects a computer to a network. The Cluster service keeps track of all network interfaces in a server cluster. This allows you to view the state of all cluster network interfaces from a cluster management application, such as Cluster Administrator. For more information, see Network state management and failure detection.

Windows Server 2003 family operating systems automatically detect the addition and removal of network interfaces. For more information, see Plug and Play network support.

# Nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Nodes**

A server cluster node is a system that has a working installation of Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition and the Cluster service. By definition, a node is always a member of a server cluster. For more information on the Cluster service, see Cluster service.

Nodes have the following characteristics:

- Every node is attached to one or more cluster storage devices. Each cluster storage device attaches one or more disks. The disks store all of the cluster's configuration and resource data. Each disk can be owned by only one node at any point in time, but ownership can be transferred between nodes. The result is that each node has access to all cluster configuration data.
- Every node communicates with the other nodes in the cluster through one or more physically independent networks, known as interconnects. Network adapters, also called network interfaces, attach nodes to networks.
- Every node in the cluster can detect another system joining or leaving the cluster.
- Every node in the cluster can detect the resources that are running locally and the resources that are running on the other cluster nodes.
- All nodes in the cluster are grouped under a common name, the cluster name, which is used when accessing and managing the cluster.

The following states describe a node's participation in server cluster operation:

| State | Meaning |
|---|---|
| Down | The node is not actively participating in cluster operations. |
| Joining | The node is in the process of becoming an active participant in the cluster operations. |
| Paused | The node is actively participating in cluster operations but cannot take ownership of resource groups and cannot bring resources online. |
| Up | The node is actively participating in all cluster operations. |
| Unknown | The state cannot be determined. |

An active node can host cluster groups. For more information about groups, see Server Cluster groups.

When running Cluster Administrator for the first time on a server node, the administrator can choose whether that node forms its own cluster or joins an existing cluster. When a node is powered up, the Cluster service already installed on that node searches for other active nodes on networks enabled for internal cluster communications.

**Forming a new cluster**

If a cluster does not already exist, the node attempts to form its own cluster by gaining control of the quorum resource. If the node gains control of the quorum resource, the node forms its own cluster and uses the recovery logs in the quorum resource to update its cluster database. The Cluster service maintains a consistent, updated copy of the cluster database on all active nodes

For more information about the quorum resource, see Quorum resource. For more information on the cluster database, see Cluster database.

**Joining an existing cluster**

A node can join an existing cluster if it can communicate with another cluster node. If a cluster exists and the joining node finds

an active node, it attempts to join that node's cluster. If it succeeds, the Cluster service then validates the node's name and verifies version compatibility. If the validation process succeeds, the node joins the cluster. The node updates its copy of the cluster database from that on the other active nodes.

# Resource types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource types**

Cluster resources are categorized by type. The Windows Server 2003 family defines several types of resources and provides resource dynamic-link libraries (DLLs) to manage the types. Developers are encouraged to create new types to fit the specific needs of their cluster-aware applications and services. By creating a new resource type, developers can have more control over how the Cluster service manages their application or service.

For more information on the resource types provided with the operating system, see Standard Resource Types.

For more information on the concepts in this topic, see:

- Server Cluster Resources
- Resource dynamic-link libraries (DLLs) in clusters running Windows Server 2003

# Server Cluster groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Groups**

A group is a collection of cluster resources with the following characteristics:

- Groups define the units of failover. That is, when one resource in a group fails and it is necessary to move the resource to an alternate node, all of the resources in the group are moved to the alternate node. For more information, see Failover and failback.

- A group is always owned by one node at any point in time. Likewise, a resource is always owned by a single group. These relationships ensure that all of a group's members reside on the same node.

For more information on resources or nodes, see Server Cluster Resources or Nodes.

Groups enable resources to be combined into larger logical units. Typically, a group is made up of related or dependent resources, such as applications and their associated peripherals and data. Indeed, groups define dependency boundaries for resources: resources cannot depend on resources in other groups. Because most application-level resources have some dependency on storage, most groups are configured around storage-class resources. A simple rule of thumb is this: If a resource uses shared storage resources, the resource belongs in the same group as that storage resource.

However, groups can also be established with resources that are unrelated and nondependent to balance the load or for administrative convenience. For more information on dependent resources, see Resource dependencies. For more information on setting resource dependencies, see Modify resource dependencies.

Every group maintains a prioritized list of the nodes that can act as its host. The preferred nodes list is generated from two sources:

- The Cluster service produces a list of preferred nodes for a group from the list of possible owners that is maintained by the group's resources. For more information on Cluster service, see Cluster service.

- Administrators can add nodes to the list through a cluster-management application, such as Cluster Administrator. For more information on using Cluster Administrator to specify the preferred owners of a group, see Specify preferred owners of a group.

To fully exploit the processing power of a cluster, establish at least as many groups as there are nodes in the cluster and evenly distribute these groups across the cluster.

The following topics discuss important concepts related to groups:

- Virtual servers

- Failover and failback

# Virtual servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Virtual servers**

To access a network application or resource in a nonclustered environment, network clients must connect to a physical server (that is, a specific computer on the network identified by a unique network name and Internet protocol (IP) address). If that server fails, access to the application or resource is impossible.

Through server clusters, Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition enable the creation of virtual servers. Unlike a physical server, a virtual server is not associated with a specific computer and can be failed over like a group. If the node hosting the virtual server fails, clients can still access its resources using the same server name.

A virtual server is a group that contains:

- A Network Name resource.
- An IP Address resource.
- All other resources, including applications, to be accessed by the clients of the virtual server.
- Other virtual servers.

For more information on groups, see Server Cluster groups.

A virtual server acts like a physical server in the following ways:

- Allows access to network resources.
- Is published to network clients under a unique server name.
- Is associated with a network name and an IP address.

**Kerberos authentication for virtual servers**

Server clusters can maintain an Active Directory computer object for virtual servers in a cluster. This allows clients accessing these virtual servers to use the security features provided by Kerberos as well as NTLM. Applications that can use Active Directory, such as Message Queuing on a virtual server, can now publish information to these computer objects.

Limitations include:

- You cannot apply Group Policy to a virtual server or the applications running on the virtual server. That is, when a virtual computer object is part of an Organizational Unit that has an associated Group Policy object, that Group Policy is not applied to the virtual server object.
- The Cluster service account is limited to creating a maximum of 10 virtual server computer objects in Active Directory by default. For information on overriding this default, see **Granting additional permissions to the Cluster service account** below.
- You cannot rename a Network Name resource and change the Kerberos authentication setting at the same time. Instead, you must allow time for changes to replicate after performing each action.

By default, Kerberos authentication for virtual servers is not enabled. You can use Cluster Administrator or cluster.exe commands to enable this feature and to configure the Kerberos properties for the network name resource of the virtual server. For more information on configuring these Kerberos properties, see Enable Kerberos authentication for virtual servers and see the listing for RequireKerberos in Cluster resource.

**The Cluster service account and Active Directory access rights**

By default, the Cluster service account is a member of the Authenticated Users group, and as such, will have the user right**Add workstations to domain**. For this user right, the default quota limitation on the number of computer accounts that can be created is 10. The value of the quota can be changed by the domain administrator.

**Note**

- This quota applies to the Cluster service account and not to the individual server cluster. If the same Cluster service account is used for multiple servers clusters, then the quota applies to the total number of computer accounts in all the clusters.

The default permissions granted to the Cluster service account when the computer object is created as a result of using this **Add workstations to domain** user right do not allow the Cluster service account to rename or disable the virtual computer object. These operations are attempted when the name of the Network Name resource is changed or if Kerberos authentication is in the process of being disabled.

**Granting additional permissions to the Cluster service account**

There are a number of techniques that can be used to overcome the limitation on the number of virtual server computer objects that the Cluster service account can create. All of these techniques include granting the Cluster service account one or more permissions on objects in Active Directory. At a minimum, the Cluster service account needs to be able to:

- Reset the password on the virtual computer object.
- Write the DnsHostName attribute on the object.
- Write the ServicePrincipalName attribute on the object.
- Create Message Queuing Configuration objects if the network name resource is the dependent resource of an Message Queuing resource.

The ability to write all properties is needed if it is necessary to allow the virtual computer object to be renamed or disabled. The following table identifies the permissions needed by the Cluster service account for a virtual computer object.

| Permission | Description | Reference |
|---|---|---|
| **Add workstations to domain** | This is given by default to the Cluster service account (and all domain accounts). With this access right, the Cluster service is allowed to create computer objects in Active Directory. | Add workstations to domain |
| **Write all Properties** | Allows changes to the properties of the computer objects. For example, if you need to rename the computer object, you must give the Cluster service this access right. | |
| **Create Computer Objects** | Overrides the default limit of 10 virtual server computer objects in Active Directory. | Add workstations to domain |

You cannot rename the computer object for a virtual server using Active Directory tools. Instead, you must use Cluster Administrator or cluster.exe to rename the Network Name for the virtual server. The Cluster service will then automatically change the name of the computer object.

For more information on the concepts in this topic, see:

- Server Cluster Resources
- Nodes
- Active Directory Overview

# Failover and failback

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Failover and failback**

**Failover**
If an individual application in a server cluster fails (but the node does not), the Cluster service typically tries to restart the application on the same node. If that fails, it moves the application's resources and restarts them on another node of the server cluster. This process is called *failover*. The Cluster Administrator can use a graphical console to set various recovery policies, such as dependencies between applications, whether or not to restart an application on the same server, and whether or not to automatically rebalance, or fail back, workloads when a failed server comes back online.

The Cluster service attempts to fail over a group when:

- The node currently hosting the group becomes inactive for any reason.
- One of the resources within the group fails, and it is configured to affect the group. For more information on configuring the failure of a resource to affect the group, see View or set resource properties.
- You force failover. For more information on forcing a resource to fail over, see Initiate a resource failure.

A failover attempt consists of the following steps:

1. The Cluster service takes all the resources in the group offline in an order determined by the group's dependency hierarchy: dependent resources first, followed by the resources on which they depend. For example, if an application depends on a Physical Disk resource, the Cluster service takes the application offline first, allowing the application to write changes to the disk before the disk is taken offline.

   The Cluster service takes a resource offline by invoking, through a Resource Monitor, the resource DLL that manages the resource. If the resource does not shut down within a specified time limit, the Cluster service forcibly terminates the resource. For more information on setting the time-out value for a resource, see Specify the restart policy for a resource.

2. When all of the resources are offline, the Cluster service attempts to transfer the group to the node that is listed next on the group's list of preferred host nodes. For more information on how the Cluster service determines which node to failover or failback to, see Determining failover and move policies for groups.

3. If the Cluster service successfully moves the group to another node, it tries to bring all of the group's resources online. Failover is complete when all of the group's resources are online on the new node.

The Cluster service continues attempting to fail over a group until it succeeds or until a defined number of attempts have been made within a given time span. A group's failover policy specifies the maximum number of failover attempts that can occur in an interval of time. If the Cluster service exceeds this limit, it concludes that the group cannot be brought online anywhere in the cluster and stops trying to fail over the group. For information on how to set group failover policy, see Set group failover policy.

**Ways to control failover policy**

- Define how the Cluster service detects and responds to the failure of individual resources in the group. For more information, see Resource failure.
- Control the order in which the Cluster service takes resources offline by establishing dependency relationships between resources. For more information, see Resource dependencies.
- Specify time-out, failover threshold, and failover period for resources. The time-out controls how long the Cluster service waits for the resource to shut down. The failover threshold and period control how many times the Cluster service attempts to fail over a resource in a particular period of time. For more information on how to set these values, see

Specify the restart policy for a resource.

- Specify a possible owner list for resources. The possible owner list for a resource controls which cluster nodes are allowed to host the resource. For more information on how to set the possible owner list for a resource, see Specify which nodes can own a resource.

For more information on the Cluster service, see Cluster service.

**Failback**
When a node becomes inactive for any reason, the Cluster service fails over any groups hosted by the node. When the node becomes active again, the Cluster service can fail back the groups originally hosted by the node.

The Cluster service fails back a group using the same procedures it performs during failover. That is, the Cluster service takes all of the resources in the group offline, moves the group, and then brings all of the resources in the group online.

You can set failback to occur during a specific time period. It is important to set the failback time because you may not want failback to occur during hours of peak usage.

For more information on how to configure a group's failback policy, see Set group failback policy.

# Server Cluster Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server cluster resources**
A cluster resource is any physical or logical component that has the following characteristics:

- Can be brought online and taken offline.
- Can be managed in a server cluster.
- Can be hosted (owned) by only one node at a time.

To manage resources, the Cluster service communicates to a resource dynamic-link library (DLL) through a Resource Monitor. When the Cluster service makes a request of a resource, the Resource Monitor calls the appropriate entry-point function in the resource DLL to check and control the resource's state. For more information on Resource Monitor or resource DLL, see Resource Monitors or Resource dynamic-link libraries (DLLs) in clusters running Windows Server 2003.

**Dependent resources**
A dependent resource is one that requires another resource to operate. For example, a network name must be associated with an IP address. Because of this requirement, a Network Name resource is dependent on an IP Address resource. Dependent resources are taken offline before the resources upon which they depend are taken offline, likewise, they are brought online after the resources on which they depend are brought online. A resource can specify one or more resources on which it is dependent. Resource dependencies also determine bindings. For example, clients will be bound to the particular IP address that a Network Name resource is dependent on.

When creating resource dependencies, take into account the following: some dependencies are strictly required, while others are not required but make sense to do so. For example, a file share that is not a DFS root has no required dependencies, but if the Disk resource that holds the file share fails, the file share will be inaccessible to users. Therefore, it is logical to make the file share dependent on the Disk resource.

A resource can also specify a list of nodes on which it is able to run. Possible nodes and dependencies are important considerations when administrators organize resources into groups. For more information on preferred nodes and groups, see Server Cluster groups. For more information on nodes, see Nodes.

The following topics describe important concepts related to resources:

- Quorum resource
- Resource failure
- Resource dependencies
- Cluster resource visibility

# Quorum resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Quorum resource**
In every cluster, a single resource is designated as the quorum resource. This resource maintains the configuration data necessary for recovery of the cluster. This data, in the form of recovery logs, contains details of all of the changes that have been applied to the cluster database. This provides node-independent storage for cluster configuration and state data.

**Quorum resource and the cluster database**
The cluster database is an integral part of the formation of a server cluster. When a node joins or forms a cluster, the Cluster service must update the node's private copy of the cluster database. When a node joins an existing cluster, the Cluster service can retrieve the data from the other active nodes. However, when a node forms a cluster, no other node is available. The Cluster service uses the quorum resource's recovery logs to update the node's cluster database. To ensure cluster unity, Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition use the quorum resource to ensure that only one set of active, communicating nodes is allowed to operate as a cluster. A node can form a cluster only if it can gain control of the quorum resource. A node can join a cluster or remain in an existing cluster only if it can communicate with the node that controls the quorum resource.

For more information on the cluster database, see Cluster database.

For more information on cluster resources, see Server Cluster Resources.

**Scenarios**
To illustrate the importance of the quorum resource, consider these situations in a simple two-node cluster:

- Node 1 fails. Node 2 continues operating, writing changes to the cluster database. But before node 1 can be restarted, node 2 fails. When node 1 becomes active, it must update its private copy of the cluster database with the changes made by node 2. The Cluster service uses the quorum resource's recovery logs to perform the update.

- The networks providing communication between nodes 1 and 2 fail. Each node assumes the other node has failed and attempts to continue operating as the cluster. If both could succeed, the result would be two separate clusters using the same cluster name and competing for the same resources. Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition use quorum resource ownership to prevent this problem and maintain cluster unity. In this scenario, the node that gains control of the quorum resource is allowed to form a cluster, and the other fails over its resources and becomes inactive.

The quorum resource can be any resource with the following attributes:

- The resource must enable a single node to gain physical control of it and defend its control.

- The resource must provide physical storage that can be accessed by any node in the cluster.

- The resource must use the NTFS file system.

**Types of quorum resources**
There are two resources that come with the operating system that can act as a quorum resource for multinode clusters. They are:

- The Physical Disk resource

- The Majority Node Set resource

For single-node clusters, the Local Quorum resource acts as a quorum resource. However, other developers can create their own quorum resource types for any resources that meet the arbitration and storage requirements. For more information, see the Microsoft Software Development Kit (SDK).

# Resource failure

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource failure**

The Cluster service considers a resource to be failed when the resource is not operational and cannot be restarted on the current host node. The Cluster service detects resource failure as follows:

- At periodic intervals, the Cluster service checks to see whether the resource appears operational. It does this by periodically invoking the Resource Monitor. The Resource Monitor, in turn, relies on the resource dynamic-link library (DLL) for each resource to implement a procedure to detect whether the resource is functioning properly. The resource DLL communicates the results back through the Resource Monitor to the Cluster service. You can specify how frequently the Cluster service checks for failed resources by setting the Looks Alive and Is Alive poll intervals. The Cluster service requests a more thorough check of the resource's state at each Is Alive interval than it does at each Looks Alive interval; therefore, the Is Alive poll interval is typically longer than the Looks Alive poll interval. For more information on how to set poll intervals, see Specify the restart policy for a resource.

  For more information on Resource Monitors, see Resource Monitors.

- If the resource DLL reports that the resource is not operational, the Cluster service attempts to restart the resource. You can specify the number of times the Cluster service can attempt to restart a resource in a given time interval. If the Cluster service exceeds the maximum number of restart attempts within the specified time period, and the resource is still not operational, the Cluster service considers the resource to be failed. For more information, see Set group failover policy.

You can configure whether a failed resource causes the group that contains the resource to be failed over to another node. For more information, see the section about setting advanced resource properties in Setting resource properties.

For more information on cluster resources, see Server Cluster Resources.

If the failed resource is configured to cause the group that contains the resource to fail over to another node, Cluster service will attempt a failover. If the number of failover attempts exceeds the group's threshold and the resource is still in a failed state, Cluster service will attempt a restart of the resource. The restart attempt will be made after a period of time specified by the resource's Retry Period On Failure property, a property common to all resources. It first tries to restart and then to fail over a resource.

Even though the unit value of this property is in milliseconds, choose a value that is in the order of minutes. Also, choose a value that is greater or equal to the value of the resource's restart period property and enforce this rule.

# Resource dependencies

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resource dependencies**
A dependent resource requires, or depends on, another resource to operate. For example, if a Generic Application resource requires access to clustered physical storage, it would depend on a Physical Disk resource.

The following terms describe resources in a dependency relationship:

- A dependent resource depends on other resources (the dependencies).
- A dependency is a resource on which another resource depends.
- A dependency tree is a series of dependency relationships such that resource A depends on resource B, resource B depends on resource C, and so on.

Resources in a dependency tree obey the following rules:

- A dependent resource and all of its dependencies must be in the same group.
- The Cluster service takes a dependent resource offline before any of its dependencies are taken offline, and brings a dependent resource online after all of its dependencies are online, as determined by the dependency hierarchy.
- Resource dependencies determine bindings. For example, clients will be bound to the particular IP address that a Network Name resource is dependent on.

Some of the resource types provided with the operating system have required dependencies. For more information, see Setting resource properties.

For more information on the concepts in this topic, see:

- Server Cluster Resources
- Resource types

# Cluster resource visibility

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster resource visibility**

When users browse the network, they can see more resources associated with a network name than you might have intended.

This happens because the browsing the network is a two-stage process. First, you locate the computer associated with a network name. Second, you ask that computer to list all of the resources that it owns. If that computer is a cluster node that hosts multiple virtual servers, you can get unintended results. For example, suppose the node called ClusterNode hosts two virtual servers, PrintServer and FileServer. In the PrintServer resource group, you have a network printer named ClusterPrinter, and in the FileServer resource group, you have a file share named ClusterFileShare. When users browse the network, all three network names (ClusterNode, PrintServer, and FileServer) are associated with the same cluster node, and that cluster node owns two resources (ClusterPrinter and ClusterFileShare). Therefore, when users browse the network, they see ClusterPrinter and ClusterFileShare associated with all three network names.

This allows users to circumvent your intended usage of network names. In the previous example, you place PrintServer and ClusterPrinter in the same resource group because you want both the network name and the printer resource to fail over as a unit to other cluster nodes; users who connect to ClusterPrinter through the network name PrintServer can maintain the connection in the event of a failover. However, users in the previous example can connect to ClusterPrinter through the network names ClusterNode and FileServer as well. If users do this, they will lose their connection to ClusterPrinter in the event of a failover because the resource ClusterPrinter will no longer belong to the computer named by ClusterNode or FileServer.

You can minimize the potential for problems by:

- Using cluster and node names that are easily identifiable by administrators, but not users.
- Defining virtual server names that are easily identifiable by users.
- Associating the virtual server names with the resources hosted by the virtual server (for example, ClusterPrintServer is a good name for a virtual server that hosts a printer named ClusterPrinter).

For more information on virtual servers, see Virtual servers.

# Server Cluster Networking

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server cluster networking**
This section covers:

- The Cluster service and TCP/IP
- The Cluster service and name resolution
- Network communication between nodes

# The Cluster service and TCP/IP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The Cluster service and TCP/IP**
The Cluster service requires the TCP/IP suite for its internal and client communication.

TCP/IP is an industry-standard suite of protocols, providing communication in a heterogeneous environment. Although TCP/IP is best known for its access to the Internet and Internet resources, TCP/IP is also well suited as an enterprise networking protocol because it provides interoperability between different types of computer systems. For more information on TCP/IP, see TCP/IP Core Networking Service.

**The Cluster service and DHCP**
The Windows Server 2003 family supports the use of DHCP addresses as public or private node addresses (you configure public and private node addresses through the Network Connections folder, not through cluster management software). The Cluster service uses the Plug and Play network support to handle the events that occur when a DHCP-allocated Internet protocol (IP) address changes. For more information about public and private networks, see Server cluster networks.

However, it is highly recommended that you use static IP addresses for each network adapter on each node. Using DHCP leases for node addresses has some disadvantages:

- First, when the address lease expires for a network, the node cannot communicate over that network until it obtains a new address. This makes the node's availability dependent on the availability of the DHCP server; the node cannot communicate over the network until it obtains its address from the DHCP server.

- Second, if an address lease expires on a network configured for public communication (communication between cluster nodes and clients), it might trigger failover for resource groups that used an IP Address resource on that network.

**Caution**

- The Cluster service does not support the use of IP addresses assigned from a DHCP server for the cluster administration address (which is an IP address resource associated with the cluster name) or any other IP address resources. If possible, it is recommended that you use static IP addresses for all public and private node addresses.

# The Cluster service and name resolution

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The Cluster service and name resolution**

Clients use name resolution services to find network resources. Clients use these same services to find resources in a server cluster. The Cluster service also uses these services when a node joins one or more other nodes to form a cluster.

In both cases, name resolution is the process of translating computer names to Internet protocol (IP) addresses. You can use the Windows Internet Name Service (WINS), the Domain Name System (DNS), or both. You can also use IP broadcast name resolution. However, because IP broadcast name resolution increases network traffic, and is ineffective in routed networks, the following sections on the Cluster service and name resolution assume you are using WINS or DNS.

The Cluster service registers a Cluster resource network name and its associated Host (A) and PTR records with a DNS server if the network adapter associated with the resource has a DNS server address specified and the DNS zone accepts dynamic updates. If the zone is also integrated with Active Directory, the zone can be configured to accept only secure updates. In a secure Active Directory-integrated zone, a Cluster resource network name has the Cluster service account associated with the Host (A) and PTR records. This registration can be deleted only by the Cluster service account or the DNS administrator. If a cluster is decommissioned but might have its name reused with a different service account, the Cluster name records must be deleted from the DNS server. If not, new records will not be registered with the DNS server and the resource will fail to go online. The Event Viewer will contain an entry indicating the resource failure along with the text "DNS operation refused" in the log entry.

For more information on resource records, see Resource records reference.

For more information on dynamic updates, see Understanding zones and zone transfer.

For more information on zones and security, see the section about secure dynamic update in Dynamic update.

The Cluster service does not require any special configuration of either WINS or DNS. If you use WINS, all cluster network resource names that have associated IP addresses are registered with WINS.

Do not use static WINS mappings for cluster resources. If you do, the clustering software cannot update the WINS entries in the event of a failover.

For more information on static versus dynamic WINS entries, see Using static mappings.

# Network communication between nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network communication between nodes**
Communication between the nodes of a cluster enables the Cluster service to detect node failures and status changes, and to manage the cluster as a single entity. You can set up multiple networks for this node-to-node communication so as to eliminate a single point of failure.

**Heartbeats**
The Cluster service, running on each node of the cluster, keeps track of the current state of the nodes within a cluster and determines when a group and its resources fail over to an alternate node. This communication takes the form of messages that are sent regularly between each node's Cluster service. These messages are called heartbeats.

In addition to heartbeat traffic, some network traffic occurs whenever changes are made to the cluster. For example:

1. During startup, node-to-node communication occurs for the purpose of synchronizing the cluster databases that reside on each node. The cluster database stores configuration information about the cluster resources.

2. Node-to-node communication to verify node failures during cluster configuration changes.

As the number of nodes in a server cluster increases, the node-to-node communication rises significantly. In Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, for server clusters with 3 or more nodes, multiple unicast messages for two classes of intracluster traffic are replaced by single multicast messages. This reduces the intracluster traffic, resulting in lower network bandwidth consumption and improved node performance. The configuration of the server cluster must meet the following conditions for heartbeat multicasting:

- The number of nodes that are members of the server cluster (rather than the number of nodes that are currently up and actively participating in the cluster) must be 3 or greater.

- All the nodes in the server cluster must be running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.

**Important**

- Both conditions above refer to nodes configured in the cluster membership rather than nodes that are currently up and actively participating in the cluster.

- 2-node server clusters use unicast, not multicast, messaging for all intracluster traffic.

- If you operate a server cluster of 3 or more nodes as a mixed version cluster (that is, Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition is installed on some nodes, and Windows 2000 on others), then the cluster as a whole will send unicast, not multicast, intracluster messages.

The two kinds of intracluster traffic affected by heartbeat multicasting are:

- Heartbeat messages sent between nodes.

- Node-to-node communication to verify node failures during cluster configuration changes.

# Planning Your Server Cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning your server cluster**
This section covers the options to consider as you plan how to set up server clusters to work most efficiently in your organization:

- Choosing a Cluster Model
- Cluster deployment and operation options
- Storage configuration options
- Planning your networking hardware
- Addressing risks using server clusters
- Limitations of server clusters
- Choosing applications to run on a server cluster
- Determining failover and move policies for groups
- Choosing a domain model
- Capacity Planning for Server Clusters
- Planning for Fault-Tolerant Disks
- Planning for unattended Setup
- Optimizing Your Clusters

# Choosing a Cluster Model

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing a cluster model**

Clusters can be categorized in three configuration models in order of increasing complexity. Each cluster model uses a different quorum resource type. This section describes each model and gives examples of the types of applications that are suited to each. The three models are:

| Cluster Model | Typical Deployments | Where is the Cluster Configuration Data Maintained? |
|---|---|---|
| Single node cluster | Development and testing platforms. | The quorum resource maintains the definitive copy of the cluster configuration data either on a cluster storage device or locally on the single node. Uses a Local Quorum resource type. |
| Single quorum device cluster | Highly available applications on a cluster storage device. | The quorum resource maintains the definitive copy of the cluster configuration data on the single cluster storage device connected to all nodes. Uses a Physical Disk resource type (or other storage class resource type). |
| Majority node set cluster | Geographically dispersed server clusters. | Each node maintains its own copy of the cluster configuration data. The quorum resource ensures that the cluster configuration data is kept consistent across the nodes. Uses a Majority Node Set resource type.<br>**Important**<br>• A majority node set cluster is best used in controlled, targeted scenarios, as part of a cluster solution offered by your Original Equipment Manufacturer (OEM), Independent Software Vendor (ISV), or Independent Hardware Vendor (IHV). The single quorum device server cluster, the standard cluster model available in Windows NT and Windows 2000, will still cover the vast majority of your cluster deployment needs. |

- Model 1: Single node server cluster configuration
- Model 2: Single quorum device server cluster configuration
- Model 3: Majority node set server cluster configuration

For information on cluster deployment options, see Cluster deployment and operation options.

# Model 1: Single node server cluster configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Model 1: Single node server cluster configuration**
Model 1 shows how you can use the virtual server concept with applications on a single node server cluster.

This cluster model does not make use of failover. It is merely a way to organize resources on a server for administrative convenience and for the convenience of your clients. The main advantage of this model is that both administrators and clients can readily see descriptively named virtual servers on the network rather than having to navigate a list of actual servers to find the file shares they need. You can also use this model to provide a platform for developing and testing cluster-aware applications.

Other advantages of this model include the following:

- The Cluster service automatically restarts the various application and dependent resources after a computer has been restored following a resource failure. This is useful for applications that benefit from an automatic restart function but do not have their own mechanisms for accomplishing it.

- You can cluster the single node with other nodes at a future time, and the resource groups will already be in place. After you configure failover policies for the groups, the virtual servers are ready to operate.

The following figure represents an example of a single-node cluster that does not make use of failover.




For example, you can use this model to locate all the file and print resources in your organization on a single computer, establishing separate groups for each department. When clients from one department need to connect to the appropriate file or print share, they can find the share as easily as they would find an actual computer.

**Note**

- Some applications, such as Microsoft SQL Server versions 6.5 and 7.0, cannot be installed on a single-node cluster.

# Model 2: Single quorum device server cluster configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Model 2: Single quorum device server cluster configuration**
Model 2 shows how you can configure multiple nodes in a single quorum device server cluster.

In this cluster model the quorum resource maintains the cluster configuration data on a single cluster storage device. This is the traditional cluster model that was available in previous versions of Windows.

The following figure represents an example of a multinode server cluster that uses a quorum device configuration.




**Note**

- An inaccurate perception of this model is that the cluster storage device can be a single point of failure. This is not true, especially for clusters that use Storage Area Networks (SANs) where there are often multiple, redundant paths from the cluster nodes to the storage device. It is recommended that you implement a hardware RAID solution for your cluster storage devices to guard against disk failure.

# Model 3: Majority node set server cluster configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Model 3: Majority node set server cluster configuration**

Model 3 shows how you can configure multiple nodes in a majority node set server cluster.

In this cluster model, each node maintains its own copy of the cluster configuration data. The Majority Node Set resource, acting as the Quorum Resource, ensures that the cluster configuration data is kept consistent across the different nodes. The quorum traffic travels unencrypted over Server Message Block (SMB) file shares. The main advantage of this model is that cluster nodes do not have to be physically connected to one or more cluster disks.

**Important**

- Use a majority node set cluster only in targeted scenarios. Do not configure your cluster as a majority node set cluster unless it is part of a cluster solution offered by your Original Equipment Manufacturer (OEM), Independent Software Vendor (ISV), or Independent Hardware Vendor (IHV). The single quorum device server cluster, the standard cluster model available in Windows NT and Windows 2000, will still cover the vast majority of your cluster deployment needs.

For detailed instructions on creating a majority node set cluster, see Create a cluster.

The following figure represents an example of a majority node set server cluster.




For example, in an 8-node cluster, you could use this model to physically locate 5 nodes of the server cluster in one data center (for example, in Seattle) and the other 3 nodes in another data center some distance away (for example, in Los Angeles) while still presenting a single point of access to clients.

Other advantages of this model include the following:

- If an application inadvertently corrupts one of the local quorum disks or if you need to take a quorum disk offline for maintenance, the cluster still remains available.

- You can create clusters without cluster disks. This configuration is useful if you need to host applications that can failover, but where there is some other, application-specific way, to replicate or mirror data between nodes.

**Number of nodes in a majority node set cluster**

The main disadvantage of this model is that, for *n*-node clusters, if more than half the nodes fail at any one time, then the cluster itself fails. The cluster is then said to have lost quorum. This is in contrast to the *n*-node single quorum device server cluster model which will not fail even if only a single node remains in the cluster. The table below shows the maximum number of node failures allowed in a majority node set cluster before the cluster fails.

| Number of nodes, *n*, in the cluster | Maximum number of node failures before the cluster loses quorum |
|---|---|
| 1 | 0 |
| 2 | 0 |
| 3 | 1 |
| 4 | 1 |

| | |
|---|---|
| 5 | 2 |
| 6 | 2 |
| 7 | 3 |
| 8 | 3 |

**Note**

- If the server cluster loses quorum, you can restore it by forcing the quorum. For more information, see Force quorum in a majority node set server cluster.

# Cluster deployment and operation options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster deployment and operation options**

For the single quorum device and majority node set server cluster models, there are many ways of deploying and operating your server cluster to best match the needs of your organization. These modes of operation (specified by defining the failover policies of the server cluster) are outlined in the table below.

For more information on these different modes of operation and example scenarios, see the Windows Deployment Guide.

For more information on cluster models, see Choosing a Cluster Model.

| Operation Mode | Description | Implementation Guide |
|---|---|---|
| N-node Failover Pairs | In this mode of operation, each application is set to fail over only between 2 specified nodes. | Limit the possible owner list for each resource group to a given pair of nodes.<br><br>For more information, see Specify which nodes can own a resource. |
| Hot-Standby Server<br><br>/N+I | Hot-standby server operation mode reduces the overhead of failover pairs by consolidating the "spare" node for each pair into a single node, providing a hot-standby server that is capable of running the applications from each node pair in the event of a failure. This mode of operation is also referred to as *active/passive*.<br><br>For larger clusters, N+I mode provides an extension of the hot-standby server mode where N cluster nodes host applications and I cluster nodes are spare nodes. | Use a combination of the preferred owners list and the possible owners list. Set the preferred node to the node that the application will run on by default and set the possible owners for a given resource group to the preferred node and the spare node.<br><br>For more information, see Specify preferred owners of a group and Specify which nodes can own a resource.<br><br>For N+I configurations, see article Q296799, "How to Configure Windows Clustering Groups for Hot Spare Support" in the Microsoft Knowledge Base. |
| Failover Ring | In this mode of operation, each node in the cluster runs an application instance. In the event of a failure, the application on the failed node is moved to the next node in sequence. | Define the order of failover for a given resource group using the preferred owner list. Choose a node order, with the preferred node list for each group starting at a different node.<br><br>For more information, see Specify preferred owners of a group. |
| Random | For large clusters running multiple applications, the best policy in some cases is to allow the server cluster to choose the fail over node at random. | Do not define a preferred owners list for the resource groups. Each resource group that has an empty preferred owners list will be failed over to a random node in the cluster in the event that the node currently hosting that group fails. |

# Storage configuration options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Storage configuration options**
You can configure your cluster storage devices in two ways when you use the cluster management tools (Cluster Administrator or cluster.exe) to create a new cluster or add nodes by selecting either the Typical (full) Configuration or Advanced (minimum) Configuration options.

**Typical Configuration**
The Typical Configuration option is appropriate for most installations and results in a completely configured cluster. When you choose this configuration option, the cluster configuration wizards (Add Nodes Wizard and New Server Cluster Wizard) automatically:

- Configure all disks on the shared bus as cluster disks to be managed by the cluster.
- Create cluster disk resources for all identified cluster disks.

### Important

- The storage validation heuristics used in the Typical Configuration option assumes that all nodes see the same target IDs and Logical Unit Numbers (LUNs). For some complicated storage devices, using the Typical Configuration option could result in an error during set up. For example, when you create the first node of the cluster, the Typical Configuration option could fail if some disks, which are not on the same storage bus as the system disk, cannot be remounted. The Typical Configuration option could also fail when nodes are added to the cluster if the target IDs and LUNs do not match on all the cluster nodes. For more information about target IDs and LUNs, see the documentation for your storage devices.

**Advanced Configuration**
If you do not want the cluster configuration wizards to automatically set up your cluster storage, select the Advanced Configuration option. This configuration option is appropriate for clusters that have complex storage configurations or for configurations in which you do not want the cluster to manage all of the storage.

**Creating a new cluster**
When you create a cluster and select this option, the New Server Cluster Wizard automatically:

- Configures only one disk to be managed by the cluster.
- Provides you with the option to create either a Local Quorum resource, a Physical Disk (or other storage-class type) resource, or a Majority Node Set resource as the quorum resource, and then create the quorum resource.

  ### Important

  - If you do not select a quorum resource type, the local quorum resource will be selected by default; however, you will be able to later switch to another quorum resource type. A local quorum cluster can only support a single node; additional nodes cannot be joined to a local quorum cluster. For more information on the three types of quorum resources, see Choosing a Cluster Model.

For information on using cluster.exe to create a cluster using the Advanced Configuration option, see Cluster.

Once the initial cluster is set up, you can configure additional cluster disks using Cluster Administrator or cluster.exe.

**Adding nodes to a cluster**
When you add nodes to a cluster and select this option, the Add Nodes Wizard does not configure the cluster storage devices. Each device on the shared bus remains as a storage device that is local to the node. You can later use Cluster Administrator or cluster.exe to configure these local storage devices as cluster storage devices.

For more information on configuring cluster storage, see Checklist: Installing a Physical Disk resource.

**Recommendations**

The following table summarizes these configuration options:

| **Storage scenario** | **Configuration option** | **Reason** |
|---|---|---|
| SCSI or Fiber channel | Typical Configuration | Simple configuration where the storage devices have the identical identity as seen from the Device Manager MMC snap-in on each node. |
| Fiber channel switched fabric with multiple switches (or other complex storage configurations) | Advanced Configuration (Add Node Wizard) | One storage unit on the fabric may have a different identity (target ID, LUN) on each node in the cluster. This configuration is illustrated in the diagram below.   |
| Private storage | Advanced Configuration (New Server Cluster Wizard) | One or more of the storage devices (on a different storage bus than the system disk) will not be managed by the cluster storage. |

It is important that you correctly configure the storage topology (for example, SCSI, Fibre Channel, Storage Area Networks) and the storage interconnects (for example, multiple paths) used in your server cluster. Before deploying your server cluster, contact your hardware vendors to ensure that your particular cluster storage configuration is supported at the hardware level. For descriptions of supported cluster storage topologies, best practices for deploying and managing cluster storage, and a list of cluster storage-related Knowledge Base articles, see the cluster storage information at the Microsoft Web site.

# Planning your networking hardware

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning your networking hardware**

Implementing the correct networking infrastructure is integral to trouble-free operation of your server cluster. This section offers the following guidelines for planning your networking hardware:

**Requirements**

- Your network hardware must be selected from the Windows Catalog. For information about hardware compatibility and the Windows Catalog, see Support resources.
- The nodes of a cluster must be connected by two or more local area networks (LANs); at least two networks are required to prevent a single point of failure. A server cluster whose nodes are connected by only one network is not a supported configuration. The adapters, cables, hubs, and switches for each network must fail independently. This usually implies that the components of any two networks must be physically independent.
- At least two networks must be configured to handle **All communications (mixed network)** or **Internal cluster communications only (private network)**. For more information, see Change how the cluster uses a network.
- Each network connecting the nodes of a cluster must be configured as a single IP subnet. The IP subnet numbers for the networks connecting the cluster nodes must be different. If a cluster node has multiple adapters attached to the same cluster network, only one adapter will be used by the Cluster service; the others will be ignored. Multiple adapters connected on one node to the same network will not provide fault tolerance or load balancing and is not recommended.
- Teaming network adapters on all cluster networks concurrently is not supported. At least one of the cluster networks that is enabled for internal communication between cluster nodes (typically a private network) must not be teamed. The use of teaming network adapters on other cluster networks is acceptable; however, if communication problems occur on a teamed network, Microsoft Product Support Services will require that teaming be disabled. If this action resolves the problem or issue, then you must seek further assistance from the manufacturer of the teaming solution.
- Use a virtual LAN (VLAN) to connect the nodes of a cluster as long as all of the nodes view it as the same IP subnet. The one-way communication latency between any pair of cluster nodes on the VLAN must be less than 0.5 second. The VLAN must fail independently of all other LANs or VLANs used to connect the nodes of a cluster. Again, this usually implies that the VLAN must be constructed of components that are physically independent from other networks used in the cluster. Both the hardware and software configuration of a VLAN must be certified and listed on the Windows Catalog.

**Recommendations**

- Make sure that the network adapters used in all cluster nodes are identical, that is, each is the same make, model, and firmware version.
- Reserve one network exclusively for internal communications between nodes; that network must not be teamed.

# Addressing risks using server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Addressing risks using server clusters**

When you configure a cluster, identify the possible failures that can interrupt access to resources. A single point of failure is any component in your environment that would block data or applications if it failed. Single points of failure can be hardware, software, or external dependencies, such as power supplied by a utility company and dedicated wide area network (WAN) lines.

In general, you provide maximum reliability when you:

- Minimize the number of single points of failure in your environment.
- Provide mechanisms that maintain service when a failure occurs.

With Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, you can use server clusters and new administrative procedures to provide increased reliability. However, server clusters are not designed to protect all components of your workflow in all circumstances. For example, clusters are not an alternative to backing up data; clusters protect only availability of data, not the data itself.

The following table lists common points of failure and describes whether the point of failure can be protected, either by server clusters or by other means.

| Failure point | server cluster solution | Other solutions |
|---|---|---|
| Network hub | Redundant networks. For more information, see Configuring cluster network hardware. | -- |
| Utility company power | -- | Uninterruptible power supply (UPS).<br><br>Put each cluster node on a separate circuit. |
| server connection | Failover | -- |
| Disk | -- | Hardware RAID, to protect the data on the disk. For more information, see Hardware redundant array of independent disks (RAID). |
| Other server hardware, such as CPU or memory. | Failover | -- |
| server software, such as the operating system or specific applications. | Failover | -- |
| Wide area network (WAN) links, such as routers and dedicated lines. | -- | Redundant links over the WAN, to provide secondary access to remote connections. |
| Dial-up connection | -- | Multiple modems. |
| Client computer within your organization | -- | Configuring multiple clients for the same level of access. If one client fails, you still have access through other clients. |
| Authentication of the cluster service account | -- | Redundant networks.<br><br>Configure individual nodes as domain controllers. |

To further increase the availability of network resources and prevent the loss of data:

- Consider having replacement disks and controllers available at your site. Always make sure that any spare parts you keep on hand exactly match the original parts, including network and SCSI components. The cost of two spare SCSI controllers can be a small fraction of the cost of having hundreds of clients unable to use data.

- Consider providing UPS protection for individual computers and the network itself, including hubs, bridges, and routers. Computers running Windows Server 2003, Standard Edition support UPS. Many UPS solutions provide power for 5 to 20 minutes, which is long enough for the operating system to do an orderly shutdown when power fails.

# Limitations of server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Limitations of server clusters**

This topic covers important limitations of server clusters.

- **Disk configuration**

  You cannot configure cluster storage devices as dynamic disks or spanned volumes (volume sets) if they will be used as cluster resources.

  Disks to be used for cluster storage must be configured as basic disks and must be formatted using the NTFS file system. Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition do not support the use of dynamic disks or the features of dynamic disks, such as spanned volumes (volume sets), for cluster storage. You cannot change an existing cluster disk, such as the quorum disk, to a dynamic disk. You can, however, use the DiskPart.exe utility to extend the volume of a basic cluster disk. for more information, see DiskPart.

  **Note**

  - Various third party products allow for dynamic disks in a cluster. For more information, search the Microsoft Knowledge Base using the keywords *server cluster dynamic disks*.

  You cannot use Remote Storage on the cluster storage device. You cannot enable Shadow Copies of Shared Folders for the quorum resource, although you can enable Shadow Copies of Shared Folders for a File Share resource.

  You can enable write caching on disks on any of the products in the Windows Server 2003 family. It is highly recommended that you do not enable this feature on cluster disks; otherwise, cached data will be lost during failover.

  Disable write back cache (write through only) on an internal redundant array of independent disks (RAID) controller because the data in the cache will be lost during failover. An example of an internal controller is a Peripheral Component Interconnect (PCI) card inside the node. Enable write cache on an external RAID controller. An external RAID controller is usually inside the disk cabinet and the data in the cache will fail over.

  You cannot use software RAID. You can use hardware RAID to protect the data on your cluster disk.

  You must partition cluster disks on the cluster storage device as master boot record (MBR) and not as GUID partition table (GPT) disks.

- **Network configuration**

  The Cluster service functions only with IP-based protocols.

  The private and public networks must be on the same subnet.

# Choosing applications to run on a server cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing applications to run on a server cluster**
Many, but not all, applications can be adapted to run on a server cluster. Of those that can, not all need to be set up as cluster resources. This section offers guidelines for making these decisions.

**Important**

- All applications running in a server cluster must be from a trusted source and all files, registry checkpoints, and other resources needed for those applications must be in a secure location. For more information, see Best practices for securing server clusters.

Three criteria determine whether an application can adapt to server clustering failover mechanisms:

- The application must use an IP-based protocol.

  Client/server applications must use an IP-based protocol (TCP, UDP, DCOM, Named Pipes, or RPC over TCP/IP) for their network communications to run on a server cluster. Any application that uses only NetBEUI or IPX protocols cannot take advantage of cluster failover.

- The application must be able to specify where the application data is stored.

  Any application you run on a server cluster must be able to store its data in a configurable location, that is, on the disks attached to shared buses. Some applications that cannot store their data in a configurable location can still be configured to fail over. However, in such cases access to the application data is lost at failover because the data is available only on the disk of the failed node.

- Client applications that connect to the server application must retry and recover from temporary network failures.

  During failover, client applications will experience a temporary loss of network connectivity. If the client application is configured to recover from temporary network connection problems, it will be able to continue operating after a server failover.

Applications that can be failed over can be further divided into two groups: those that support the Cluster API and those that do not.

Applications that support the Cluster API (for example, Microsoft SQL Server 2000) are defined as *cluster aware*. These applications can register with the Cluster service to receive status and notification information, and they can use the Cluster API to administer clusters.

**Important**

- Only Microsoft SQL Server 2000 Failover Clusters and subsequent versions are supported in a server cluster running Windows Server 2003, Enterprise Editionor Windows Server 2003, Datacenter Edition.

Applications that do not support the Cluster API are defined as *cluster unaware*. If cluster-unaware applications meet the TCP/IP and remote storage criteria, they can still be used in a cluster and often can be configured to fail over.

In either case, applications that keep significant state information in memory are not the best applications for clustering because information that is not stored on disk is lost at failover.

Finally, it is important to note that in order for an application to be certified by Microsoft as cluster compliant, the application must meet certain requirements. Those requirements are as follows:

**Summary of Cluster service requirements**
**Note**

- Applications that do not meet these requirements are eligible for the Windows Server 2003, Standard Editioncertification only.

**Rationale**

A server cluster is a group of independent servers managed as a single system for higher availability. Cluster service is a set of system services in Windows Server 2003, Enterprise Editionand Windows Server 2003, Datacenter Editionthat enable you to form server clusters by connecting multiple servers together, making them appear to network clients as a single, highly available system.

Cluster service can automatically detect the failure of an application or server, and restart the application, either on the same server if it is still alive, or on another surviving server.

These requirements help ensure that your application will run properly with Cluster service enabled, so that:

- Your server application can fail over to other servers.
- The client side of your application properly handles failure of the server application.

**Customer benefits**

Customers who run your application in a clustered environment can achieve higher availability because your application can continue to provide service during both planned downtime (such as hardware and software upgrades) and unplanned outages (such as hardware or software failure).

When one of the systems, or nodes, in the cluster fails or becomes unavailable, Cluster service transfers its workload to another system in the cluster. Users experience only a momentary pause in service. You can also configure Cluster service to provide failback so that when the failed server comes back online, the workload is rebalanced across the server cluster.

**Requirements**

- Applications must be able to be installed on up to eight nodes for certification on Windows Server 2003, Enterprise Editionor Windows Server 2003, Datacenter Edition.
- Applications must support failover to all cluster members.
- Clients must survive failure of the server application without failing or affecting the stability of the system.

**How to comply with Cluster service requirements**

- Applications must be able to be installed on up to eight nodes for certification on Windows Server 2003, Enterprise Editionor Windows Server 2003, Datacenter Edition.

  **Note**

  - Make certain your application setup does not make any assumptions about the number of nodes in the cluster. Verify that it enumerates all nodes in the cluster and allows installing your application on any node, even if the disks where your application data is stored are not physically located on that node.

- Applications must support failover to all cluster nodes.

  When a node in the cluster fails, Cluster service will move the resource group from that node to a new node. A resource group is a collection of resources that provide services to clients and can depend on each other.

  Resources that represent the primary functionality of your application must be able to start up (come online) on any other node in the cluster. After a failover is complete, very that clients are able to access all data exposed by primary functions.

  **Note**

  - Cluster service operates under a "shared nothing" architecture in which each server owns its own disk resources.

In the event of a server failure, ownership of the clustered disk is transferred from one server to another. For applications to properly support failover, the application's data must be stored on the clustered disk.

- Clients must survive failure of the server application without failing or affecting the stability of the system.

  Any client that ships with your server application must gracefully handle both cluster node failures and application failures. Cluster and application failures may cause clients to temporarily lose their connection to the server application (see the next item). Your client must survive both the failure of the server application and node failure as follows:

  - When a connection to the server application is lost, your client application must not stop responding or compromise the stability of the client operating system.

  - After the failover is complete and the application is restarted on a cluster node, your client must reconnect to the cluster in one of the following ways:

Reestablish the lost connection without user intervention and with no loss of data

--or--

Offer the user a chance to reconnect and retry the operation that failed; for example, prompt the user to refresh the data in the client

- If the server application is not able to restart, the client must inform the user that the connection could not be reestablished.

Connections to the server application can be lost for any of the following reasons:

- The application fails and is then restarted on the same node.

- The application fails and is restarted on a new node.

- The node fails, and all resources fail over to a new node.

- The administrator moves the resource group containing the application to a new node.

- The administrator shuts down the server application.

- All nodes in the cluster fail.

- The client's network connection to the cluster is interrupted, even though the cluster and the server application are still running.

These failures might be exposed to the client application as application time-outs, invalid handles, network failures, and connection time-outs.

**Development guidelines**
The guidelines in this section are not requirements that will be tested individually for certification. However, following these guidelines will help you meet the requirements described previously:

- Use TCP/IP protocol

  Services that communicate with clients (as well as their clients) must use TCP/IP to be able to take advantage of IP address failover provided by Cluster service. servers that do not communicate with clients do not need to use TCP/IP.

- Make sure that applications use a virtual server name and IP address to connect to the node hosting the server application.

  Clients communicating with the cluster resources must use the virtual server IP address or virtual server network name

to support failover.

If the server application publishes a network name or IP address to clients, it must publish an IP virtual server IP address or network name. A server application that depends on a computer name or an IP address is supposed to use a network name or an IP address of a virtual server that is used by clients to access this application. Make sure that your server application does not fail to restart on another node because the computer name on this node is different.

The following code example illustrates how to set the server application environment as part of your resource DLL online routine.

```
//
// Create the new environment with the simulated net name when the
// services queries GetComputerName.
//
if ( ! ClusWorkerCheckTerminate( pWorker ) )
{
nStatus = ResUtilSetResourceServiceEnvironment(
YOUR_SERVICE_NAME,
pResourceEntry->hResource,
g_pfnLogEvent,
pResourceEntry->hResourceHandle
);
if ( nStatus != ERROR_SUCCESS )
{
break;
} // if: error setting the environment for the service
}
```

**About IP address failover**

Client applications use a virtual server IP address to access services running on a server Cluster. A virtual server is a cluster resource group containing an IP address and a network name. A virtual server can be brought online on any node in the cluster. However, it appears to clients accessing it as the same physical computer.

The IP address of the virtual server has to be configured as a cluster resource in the same resource group where the server application was created. In case of a node failure, all resource groups running on this node are moved to another node in the cluster. The IP address of the virtual server is now available on another node and all connections with the clients can be reestablished.

- Upon failure, clients must preserve user data.

  The client application must be able to reconnect and resume an operation in the event of cluster node failure. It must either offer the user a chance to retry the connection, or it must retry the connection automatically until it succeeds or can determine that the server application could not be brought online.

  In case of a node failure, all resource groups running on the failed node are moved to another node in the cluster. Cluster service requires some time to bring all resources online and restart services on the other node. The time needed to fail over a server application depends on many factors. The most significant is the time required to restart the application.

- The location of application data must be configurable.

  Cluster service can fail over only disks managed by the cluster that are on the storage bus shared among all nodes in the cluster. Make sure that your application setup allows for selecting the drive and installing application data on any drive. Cluster-aware setup should allow installing the application data only on a shared drive managed by the cluster.

- Checkpoint either automatically or manually the state information required for a clean restart.

  If a server application maintains any state information required for a clean restart, make sure that it checkpoints this state information frequently to a cluster disk managed by a cluster. It is supposed to use this data to recover quickly after a failure.

- Upon failure, an application can be restarted and, if applicable, recover to the last checkpoint.

  The server application must recover from a node failure. Make sure that a sudden node failure (for example, a power blackout) will not leave your application in a state in which it cannot restart.

  After a node failure, Cluster service moves the server application running on the node to another node along with any other resources it may depend on. The server application must restart, recover, and resume operation in the time you specified in your product literature.

- At least one instance of the application can run as a cluster resource.

  Cluster service manages applications as cluster resources. A cluster resource is a physical or logical entity that can be owned by a node, brought online and taken offline, moved between nodes, and managed as a server cluster object. A resource can be owned only by a single node at any point in time. A resource is associated with, and managed by, a resource type.

  If the resource supports it, Cluster service can manage multiple instances of the same resource, but it is acceptable to support only one instance.

  To take advantage of clustering, your application has to be configured as a cluster resource. Make sure you can create at least one instance of your application. Your application must function properly as a cluster resource. Cluster service must be used to start your application (bring it online) and stop it (take it offline).

- Can be configured at least as a generic service or application.

  The monitoring and failover capabilities of Cluster service can be extended to support any application. Cluster service uses resource DLLs to extend its failover support to other resource types.

  Applications that do not offer an application-specific resource DLL can still take advantage of clustering by using a generic application or generic service resource type. These resource types offer failover protection against most failures, notably node failure. However, they cannot detect your application failures. If your application stops responding, Cluster service will not be able to detect this failure and either restart or fail over your application.

  It is acceptable to use a generic application or a generic service type to manage your application as a cluster resource.

**How to pretest applications for Cluster service requirements**

**How to pretest whether your application is cluster ready.**
If your application's setup is cluster aware, use setup to configure all nodes.

If your application's setup is not cluster aware, install your server application on at least two nodes in the cluster. Use the Cluster Administrator console to create a virtual server and configure your server application as a generic service or application. Use the Cluster Administrator console to move your resource to either node in the cluster. If your application is cluster-ready, it is supposed to come online on any node in the cluster. Clients should be able to access the service provided by your application no matter which node hosts it.

For certification on Windows Server 2003, Enterprise Editionor Windows Server 2003, Datacenter Edition, repeat this procedure for three-node, four-node, five-node, six-node, seven-node, and eight-node configurations.

**How to pretest whether your application supports failover**

1. After you have installed the application on all nodes in the cluster, run functionality tests to verify that the application is fully functional and stable.

2. Cause the node running your application to fail so that failover of the application is triggered.

   Following are suggestions for triggering failure:

   - **Hardware failure.** Simulate by doing a hard reset.
   - **Operating system failure.** Simulate by emitting a CTRL+C command followed by a ".reboot" command from a remote kernel debugger.
   - **Application failure.** Simulate using the End Process feature in Task Manager or Process Viewer (Pview.exe in the Platform SDK).

     Note that normal shutdown of the computer is not a valid test for failover because the application will have the opportunity to shut down gracefully.

3. Verify that the application restarts on a new node in the cluster.
4. Run functionality tests to verify that all functionality is again available on the new node. The application must have access to all data to which it previously had access.
5. For testing on Windows Server 2003, Enterprise Editionor Windows Server 2003, Datacenter Edition, repeat steps 2 through 4 to verify that the application subsequently fails over to each of the remaining nodes.

**How to pretest whether clients you provide survive failure and subsequent restart of the server application**
Cause the server application to fail using each of the following scenarios:

- Shut down the server application using the normal shutdown sequence and leave all nodes in the cluster running.
- Terminate the application process (do not use the normal shutdown sequence), but leave the node running.
- Stop the node.

  Do not use normal shutdown. The node and application must not have time to exit gracefully. Following are suggestions for various failure modes:

  - **Hardware failure.** Simulate by doing a hard reset.
  - **Operating system failure.** Simulate by emitting a CTRL+C command followed by a ".reboot" command from a remote kernel debugger.
  - **Application failure.** Simulate using the End Process feature in Task Manager or Process Viewer (Pview.exe in the Platform SDK).

For each case:

- Verify that the client does not stop responding or lose stability when the server application fails.
- After the server application is restarted, either on the same node or a different node, verify that the client either:

  - Reestablishes the connection with no user intervention and no loss of user data.

    --or--

  - Prompts the user to retry the connection and the application then establishes the connection.

**Note**

- If you need to manually configure the client to access the server application, configure it to use the virtual server, not the node name.

**How to pretest whether clients you provide survive failure without subsequent restart of the server application**

- Cause the application to fail in a way that does not allow it to restart on the cluster. To do this, you can take the resource offline, or stop all nodes.
- Verify that the clients do not stop responding or lose stability.
- Verify that the clients notify the user in a reasonable time that the connection to the server application was lost.
- Verify that the client can be closed without failing or affecting the stability of the client's workstation, and that the user can preserve data, if appropriate.

# Determining failover and move policies for groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Determining failover and move policies for groups**
For every resource group, the Cluster service maintains a prioritized list of the nodes that are supposed to act as its host. To fully exploit the processing power of a cluster, it is recommended that you understand how the Cluster service uses this prioritized list to move, failover, and failback resource groups. This will allow you to better balance groups among all nodes and maximize the performance of your server cluster.

**Failover policies**
You can assign the failover policies for each group of resources in your cluster. These policies determine exactly how a group behaves when failover occurs. You can choose which policies are most appropriate for each resource group you set up. For more information on setting failover policies, see Setting group properties.

Failover policies for groups include three settings:

Failover timing

Preferred nodes list

Failback timing

**Failover timing**
You can set a group for immediate failover when a resource fails, or you can instruct the Cluster service to try to restart the failed resource a number of times before failover occurs. If it is possible that the resource failure can be overcome by restarting all resources within the group, then set the Cluster service to restart the group.

Verify that the failover threshold does not exceed the number of nodes in the cluster. For example, if the cluster contains three nodes, set the failover threshold to three.

**Preferred nodes list**
The Cluster service maintains an ordered, internal list of available nodes for resource group moves, failovers, and failbacks. Optionally, you can set up a prioritized list of nodes that are preferred owners for a group so that the Cluster service can choose the best available node on that list rather than choose a node at random. This list, known as the preferred nodes list, is useful if one or more of the nodes is better equipped to host the group.

### Important

- When setting up a preferred nodes list for a resource group, it is highly recommended that you define a **complete** preferred nodes list, that is, one listing all nodes in your server cluster, in order of priority.

**Failback timing**
You can set a group to fail back to its preferred node as soon as the Cluster service detects that the failed node has been restored, or you can instruct the Cluster service to wait until a specified hour of the day, such as after peak business hours.

### Important

- Failback only occurs when you have defined a preferred nodes list for a resource group and failback is allowed for that resource group.

  If you specify that a group failback to a preferred node and then restart the node to test the failback policy you set, the resource group will not failback. A resource group will not failback when a node is restarted after a planned shutdown and restart. To test the failback policy, you must press the reset button on the node. For more information on testing failback policies, see Test node failure.

**Moving groups**
You can use Cluster Administrator or cluster.exe to manually move a resource group from one node to another. If you specify the destination node for the move, the Cluster service will move the group to the node you specify. However, if you do not specify the destination node, but instead select the **Best Possible** option, the Cluster service will select a destination node. The

logic that the Cluster service uses to select the destination node will be different for each of these scenarios:

- You **have** specified a list of nodes as preferred owners for that resource group.
- You **have not** specified a list of nodes as preferred owners for that resource group.

For more information on moving resource groups, see Move a group to another node.

The table below summarizes the move, failover, and failback logic that the Cluster service uses depending on whether the user defines a preferred nodes list.

| Scenario | Move group | Resource or node failures | Failback enabled |
|---|---|---|---|
| Preferred Nodes List Defined | Fails over to the **first** node on the list that is available. | Fails over to the **next** node on the list. | Fails back to the **original** node. |
| Preferred Nodes List Not Defined | Fails over to a randomly selected node. | Fails over to a randomly selected node. | No failback. |

The section below provides details on the logic behind manual moves, failovers, and failbacks.

**Moves, failovers, and failbacks**

**Preferred nodes list not defined**
If you have not set up a preferred owners list for a group, the Cluster service creates an internal list of nodes ordered by their node IDs. In general, this is the same as the order in which the nodes were added to the server cluster. It may, however, be ordered differently if nodes have been evicted from the server cluster. When a resource group or node fails and you have not defined a preferred list of owners for that group, the Cluster service randomly selects a node from the internal list of nodes and attempts to transfer the group to that node.

If you manually move a resource group off a node without specifying a destination node, the Cluster service transfers the group to a randomly selected node.

**Preferred nodes list defined**
If you define a complete preferred nodes list for a group (that is, one listing all the nodes in the cluster), then the Cluster service uses this defined list as its internal list. However, if you define a partial preferred nodes list for a group, then the Cluster service uses this defined list as its internal list and appends any other installed nodes not on the preferred list, ordered by their node IDs. For example, if you created a 5-node cluster (installing the nodes in the order *Node1*, *Node2*, *Node3*, *Node4*, and *Node5*) and defined *Node3*, *Node4*, and *Node5* as preferred owners for the resource group, PRINTGR1, the Cluster service would maintain this ordered list for PRINTGR1: *Node3*, *Node4*, *Node5*, *Node1*, *Node2*. How the Cluster service uses this list depends on whether the resource group move is due to a resource/node failure or a manual move group request.

### Important

- When setting up a preferred nodes list for a resource group, it is recommended that you define a **complete** preferred nodes list, that is, one listing all nodes in your server cluster, in order of priority.

**Preferred lists and resource or node failures**
For resource group or node failures, the group fails over to the node **next** to the current owner on the preferred nodes list. In the example above, if the resource group PRINTGR1 on *Node3* fails, then the Cluster service would fail that group over to the next node on the list, *Node4*. If you allow failback for that group, then when *Node3* comes up again, the Cluster service will fail back PRINTGR1 to that node.

**Preferred lists and manual moves**
If you manually move a resource group without specifying a destination node, the Cluster service fails over that group to the **first** node listed on the preferred nodes list that is available.

# Choosing a domain model

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing a domain model**

Nodes in a server cluster can be either member servers or domain controllers. However, in either case, all nodes must belong to the same domain.

If you configure your cluster nodes as domain controllers, you must first make sure you have the hardware to support them. For more information, see Capacity Planning for Server Clusters.

If you configure all of the cluster nodes as member servers, then the availability of the cluster will depend on the availability of the domain controller. The Cluster service requires that the nodes be able to contact the domain controller to function correctly, so the domain controller must be highly available. To avoid a single point of failure, the domain must have at least two domain controllers. For more information on increasing availability, see Addressing risks using server clusters.

At least two domain controllers must be configured to be global catalog servers.

**DNS considerations**

If DNS is used to resolve names in the domain, then you must deploy at least two Domain Name System (DNS) servers. If your cluster nodes are the only domain controllers in their domain, they must be configured as Domain Name System (DNS) servers. Make sure that the DNS servers support dynamic updates. For more information on DNS, see DNS.

Each domain controller and cluster node must be configured with a primary and at least one secondary DNS server. If the domain controllers are also DNS Servers, then make sure each points to itself for primary DNS resolution and to the other DNS servers for secondary resolution.

**Performance considerations**

Take into account the additional overhead that is incurred by the domain controller services. In large networks running on Windows Server 2003 family operating systems, substantial resources can be required by domain controllers for performing directory replication and server authentication for clients. For this reason, many applications, such as Microsoft SQL Server and Message Queuing, recommend that, for best performance, you not install the application on domain controllers. However, if you have a very small network in which account information rarely changes and in which users do not log on and off frequently, you can use domain controllers as cluster nodes.

# Capacity Planning for Server Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Capacity planning**
After you assess your clustering needs, you are ready to determine how many servers you need and with what specifications, such as memory and hard disk storage.

This section covers:

- Planning your groups
- Determining server-capacity requirements

# Planning your groups

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning your groups**
You can take six steps to organize your applications and other resources into groups:

**List all your server-based applications**
Most groups contain one or more applications. Make a list of all applications in your environment, regardless of whether or not you plan to use them with the Cluster service. Your total capacity needs are determined by the sum of the total number of groups (virtual servers) that will run in your environment and the total amount of other software you will run independently of groups.

**Sort the list of applications**
Determine which of your applications can use failover. For more information, see Choosing applications to run on a server cluster.

Also list applications that will reside on cluster nodes but that will not use the failover feature because it is inconvenient, unnecessary, or impossible to configure the applications for failover. Although you do not set failover policies for these applications or arrange them in groups, they still use a portion of the server capacity.

Before clustering an application, review the application license, or check with the application vendor. Each application vendor sets its own licensing policies for applications running on clusters.

**List all other resources**
Determine which hardware, connections, and operating-system software a server cluster can protect in your network environment. For more information, see Addressing risks using server clusters.

For example, the Cluster service can fail over print spoolers, to protect client access to printing services. Another example is a file-server resource, which you can set to fail over, maintaining client access to files. In both cases, capacity is affected, such as the random access memory (RAM) required to service the clients.

**List all dependencies for each resource**
When you create this list, include all resources that support the core resources. For example, if a Web-server application fails over, the Web addresses and cluster storage devices containing the files for that application must also fail over if the Web server is to function. All these resources must be in the same group. This ensures that the Cluster service keeps interdependent resources together at all times.

Two guidelines for quantifying resources are:

- A resource and its dependencies must be together in a single group.
- A resource cannot span groups.

For example, if several applications depend on a particular resource, you must include all of those applications with that resource in a single group.

**Make preliminary grouping decisions**
Another factor in the way you assign groups is administrative convenience. For example, you might put several applications into one group because viewing those particular applications as a single entity makes it easier to administer the network.

A common use of this technique is to combine file-sharing resources and print-spooling resources in a single group. All dependencies for those applications must also be in the group. You can give this group a unique name for the part of your organization it serves, such as "Accounting File&Print." Whenever you need to intervene with the file- and print-sharing activities for that department, you would look for this group in Cluster Administrator.

Another common practice is to put applications that depend on a particular resource into a single group. For example, suppose a Web-server application provides access to Web pages, and that those Web pages provide result sets that clients access by querying a Microsoft SQL Server 2000 database application, through the use of Hypertext Markup Language (HTML) forms. By putting the Web server and the Microsoft SQL Server 2000 database in the same group, the data for both core applications can reside on a specific disk volume. Because both applications exist within the same group, you can also create an Internet

protocol (IP) address and network name specifically for this resource group.

**Make final grouping assignments**
After you list the resources that you want to group together, assign a different name to each group, and create a dependency tree. A dependency tree is useful for visualizing the dependency relationships between resources.

To create a dependency tree, first write down all the resources in a particular group. Then draw arrows from each resource to each resource on which the resource directly depends.

A direct dependency between resource A and resource B means that there are no intermediary resources between the two resources. An indirect dependency occurs when a transitive relationship exists between resources. If resource A depends on resource B and resource B depends on resource C, there is an indirect dependency between resource A and resource C. However, resource A is not directly dependent on resource C.

The figure shows the resources in a final grouping assignment in a dependency tree.




In the diagram, the File Share resource depends on the Network Name resource, which, in turn, depends on the IP Address resource. However, the File Share resource does not directly depend on the IP Address resource.

For information on using Cluster Administrator to set resource dependencies, see Setting resource properties.

# Determining server-capacity requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Determining server-capacity requirements**

After you choose a cluster model, determine how to group your resources, and determine the failover policies required by each resource, you are ready to determine the hardware capacity required for each server in the cluster. This topic explains the criteria for choosing computers for use as cluster nodes.

- **Cluster storage requirements**

  Each node in a cluster must have enough persistent storage capacity to store permanent copies of all applications and other resources required to run all groups. Calculate this for each node as if all of these resources in the cluster were running on that node, even if some or all of those groups run on other nodes most of the time. Plan these disk space allowances so that either node can efficiently run all resources during failover.

- **CPU requirements**

  Failover can strain the CPU processing capacity of a node when it takes control of the resources from a failed node. Without proper planning, the CPU of a surviving node can be pushed beyond its practical capacity during failover, slowing response time for users. Plan your CPU capacity on each node so that it can accommodate new resources without unreasonably affecting responsiveness.

- **Random access memory (RAM) requirements**

  When planning your capacity, make sure that each node in your cluster has enough RAM to run all applications that might run on either node. Also, make sure your paging files are set appropriately for each node's RAM.

  For more information about paging files, see Optimizing the paging file size and location and Change the size of the virtual memory paging file.

**Note**

- You can use the Windows DNA Performance Kit to test and tune the performance of your applications in your server cluster. For more information about the Windows DNA Performance Kit or about how to download the kit, see the Windows DNA Performance Kit link in Web Resources at the Microsoft Windows Resource Kits Web site.

# Planning for Fault-Tolerant Disks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for fault-tolerant disks**
Many groups include disk resources for disks on shared buses. In some cases, these are simple physical disks, but in other cases they are complex disk subsystems containing multiple disks. Resources depend on the physical disk resource that represents the storage device. An unrecoverable failure of a disk resource results in certain failure of all resources that depend on that disk.

For these reasons, you might decide to use special methods to protect your disks and disk subsystems from failure. One common solution is the use of a hardware-based redundant array of independent disks (RAID). RAID support ensures the high availability of data contained on disk sets in your clusters. Some of these hardware-based solutions are considered fault tolerant, which means that data will not be lost if a member of the disk set fails.

**Important**

- You cannot use software fault-tolerant disk sets for cluster storage.

This section covers:

- Hardware redundant array of independent disks (RAID)
- Error recovery

# Hardware redundant array of independent disks (RAID)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Hardware redundant array of independent disks (RAID)**
The Windows Catalog contains many different hardware RAID configurations for clusters. Because many hardware RAID solutions provide power, bus, and cable redundancy within a single cabinet, and track the state of each component in the hardware RAID firmware, they provide data availability with multiple redundancy, protecting against multiple points of failure. Hardware RAID solutions can also use an on-board processor and cache.

Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition can use these disks as standard disk resources.

For more information on the different hardware RAID configurations, see the Microsoft Windows Catalog at Support resources.

**Important**

- You cannot use software fault-tolerant disk sets for cluster storage.

# Error recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Error recovery**

With transaction logging and recovery, the NTFS file system ensures that the volume structure will not be corrupted, so all files remain accessible after a system failure. The NTFS file system also uses the recovery technique called cluster remapping. When the operating system returns a bad-sector error to the NTFS file system, NTFS dynamically replaces the disk cluster containing the bad sector and allocates a new disk cluster for the data. If the error occurs during a read, NTFS returns a read error to the calling program, and the data is lost (unless protected by RAID fault tolerance). When the error occurs during a write, NTFS writes the data to the new disk cluster, and no data is lost. NTFS puts the address of the disk cluster containing the bad sector in its Bad Sector file so that the bad sector is not reused.

Even with transaction logging and recovery and disk cluster remapping, user data can be lost due to hardware failure if you do not use a fault-tolerant-disk solution. For more information, see Planning for Fault-Tolerant Disks.

# Planning for unattended Setup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for unattended Setup**

This topic provides general information about unattended Setup. For detailed instructions on running unattended Setup (also called automated installation), see "Automating and Customizing Installations" at the Microsoft Windows Server System Web site. Also see the Windows XP Professional Resource Kit at the Microsoft Windows XP Resource Kits Web site. For information about the Windows Deployment and Resource Kits, see Using the Windows Deployment and Resource Kits.

To simplify the process of setting up a product in the Windows Server 2003 family on multiple computers, you can run Setup unattended. To do this, you create and use an *answer file*, a customized script that answers the Setup questions automatically. Then, you run Winnt32.exe or Winnt.exe with the appropriate options for unattended Setup. Choose the command according to the operating system that is running when you start unattended Setup:

- To start unattended Setup on a computer running MS-DOS or Windows 3.*x*, use Winnt.exe (with the appropriate options).
- To start unattended Setup on a computer running Windows 95, Windows 98, Windows Millennium Edition, Windows NT, Windows 2000, Windows XP, or a product in the Windows Server 2003 family, use Winnt32.exe (with the appropriate options). With Windows NT 4.0, before starting unattended Setup, apply Service Pack 5 or later.

To learn about unattended Setup, see the following sources:

- **For detailed instructions on running unattended Setup (also called automated installation):** See the Windows Deployment and Resource Kits listed in the first paragraph of this topic.
- **To view the command options available for Winnt.exe:** On a computer running Windows 3.*x* or MS-DOS, insert the Setup CD for a product in the Windows Server 2003 family in the CD-ROM drive and open the command prompt. Then, change to the CD-ROM drive, change to the I386 directory, and type:

  **winnt /?**
- **To use an x86-based or x64-based computer to view the command options available for Winnt32.exe:** On a computer running Windows 95, Windows 98, Windows Millennium Edition, Windows NT, Windows 2000, Windows XP, or a product in the Windows Server 2003 family, insert the Setup CD for Windows Server 2003, Standard Edition, in the CD-ROM drive, and open the command prompt (click **Start**, click **Run**, and then type **cmd**). Then, change to the CD-ROM drive, change to the I386 directory, and type:

  **winnt32 /?**
- **To use an Itanium-based computer to view the command options available for Winnt32.exe:** On an Itanium-based computer running Windows XP 64-Bit Edition; the Itanium-based version of Windows Server 2003, Enterprise Edition; or the Itanium-based version of Windows Server 2003, Datacenter Edition, insert the Setup CD for the product in the CD-ROM drive and open the command prompt (click **Start**, click **Run**, and then type **cmd**). Then, change to the CD-ROM drive, change to the IA64 directory, and type:

  **winnt32 /?**

For information about using Winnt.exe or Winnt32.exe for interactive setup (not unattended Setup), see How to install and upgrade the operating system.

# Optimizing Your Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Optimizing your clusters**

Because Cluster service uses an adaptable architecture, it is able to allocate resources as needed to meet changing usage requirements. However, there are some hardware and configuration choices you can make to improve cluster performance. These choices are explained in the following sections:

- Focusing on the right hardware resources for server clusters
- Optimizing the paging file size and location for server clusters
- Using additional TCP/IP services for server clusters

# Focusing on the right hardware resources for server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Focusing on the right hardware resources**

The goal in tuning the Cluster service and hosted applications is to determine which hardware resource will experience the greatest demand, and then to adjust the configuration to relieve that demand and maximize total throughput. Structure a system so that its resources are used efficiently.

For example, if the primary role of the cluster is to provide high availability of file and print services, high disk I/O usage will be incurred due to the large number of files being opened and closed. File and print services also cause a heavy load on network adapters because of the large amount of data that is being transferred. It is important to make sure that your network adapter can handle the load. In this scenario, memory typically does not carry a heavy load (although memory usage can be heavy if a large amount of system memory is allocated to file-system cache). Processor utilization is also typically low in this environment. In such cases, memory and processor utilization usually do not need the optimizing that other components need.

In contrast, a server-application environment (such as one running Microsoft SQL Server, Microsoft Exchange Server, Microsoft Systems Management Server, and Microsoft Systems Network Architecture (SNA) Server) is much more processor- and memory-bound than a typical file and print server environment because much more processing is taking place at the server. In these cases, it is best to use high-end multiprocessor computers. The disk and network loads tend to be less used because the amount of data being sent over the wire and to the disk is smaller. The Cluster service itself uses few of the system resources either for intracluster communication or for the operation of the cluster itself.

# Optimizing the paging file size and location for server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Optimizing the paging file size and location**

The location and size of the paging files can greatly affect performance. Paging files cannot be located on the shared bus, but if multiple unshared drives are available on the individual servers, placing the paging file on a separate, fast, low-use drive can boost performance. The size setting of the paging file is also critical to performance. A paging file that must constantly expand and shrink requires additional processor and I/O overhead.

# Using additional TCP/IP services for server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using additional TCP/IP services**

WINS, DHCP, and DNS Services use additional system resources, which you must consider before implementing any of these services on a cluster. In particular, WINS can require substantial overhead during peak logging on times in large networks. In such cases, make sure that your system has the necessary capacity to perform adequately.

**Important**

- DNS is not clustered, but the node acts as the server of the service. If you run DNS on a node of a cluster, make sure to implement backup servers to provide redundancy, just as you would normally do on your network.

# Using Server Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using server clusters**

This section provides information on how to use a server cluster to provide high availability application support.

- Server cluster rules
- Applying administrative procedures to server clusters
- Cluster automation interface
- General storage configuration tips for server clusters
- Backing up and restoring server clusters
- Uninstalling a server cluster
- Cluster Hardware Configuration
- Standard Resource Types
- Managing Security in a Cluster
- Using Cluster Administrator
- Using a server cluster with large numbers of file shares
- Using Shadow Copies of Shared Folders in a server cluster
- Setting Properties
- Network Management

# Server cluster rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Server cluster rules**

Following these rules will help you solve the most common problems associated with clusters.

- Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition must be installed on each node.
- All nodes of the cluster must belong to the same Windows Server 2003 family operating systems domain; see Choosing a domain model.
- When you first create the cluster or add nodes to the cluster, you must be logged on to the domain that has administrative user rights and permissions on all nodes.
- Each cluster node can belong to only one cluster.
- The shared drive must be formatted using the NTFS file system.
- Cluster disks on a shared bus must be partitioned as master boot record (MBR) and not as GUID partition table (GPT) disks.
- Do not change the computer name of a cluster node after configuring the first node of your cluster.
- Do not repartition disks on the shared bus without first deleting the corresponding disk resources.
- If you start the node using an operating system other than Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, do not run any diagnostic tools that make low-level writes to a physical disk.
- Do not reassign the drive letters of the system disks on the nodes.
- Verify that either the WINS or DNS service is accessible from all nodes.
- Verify that TCP/IP is installed on all nodes and clients.
- Verify that all cluster nodes can detect one another over the network. Use the TCP/IP **ping** utility or the Windows Server 2003 family operating system **net view** command.
- Check all event logs on all affected computers for error messages.

For information about how to obtain product support, see Technical support options.

# Applying administrative procedures to server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Applying administrative procedures to server clusters**
Typically, an administrator follows a set of server operating procedures when performing administrative tasks on a server, such as restarting a server, backing up applications and data, and installing new hardware or software. For a server cluster, reexamine these operating procedures; some will have to be modified in a server cluster environment.

**Reexamining procedures that take a server cluster offline**
Because you install Cluster Administrator on more than one computer, the tasks of administering servers and performing common server operating procedures are less disruptive to users. There are some administrative tasks that make a server unavailable to users. A server cluster, however, has another node that can continue to make the server resources available to users while you are performing administrative tasks. Therefore, you might not have to wait for nonpeak hours, when:

- A server must be restarted.
- New hardware must be installed.
- New software must be installed.

**Reexamining backups on a server cluster**
Make sure that your backup procedures account for clustering issues. In a server cluster, the following procedure is recommended:

1. Perform an Automated System Recovery (ASR) backup on each node in the cluster.

   **Note**

   - By default, Backup Operators do not have the user rights necessary to create an Automated System Recovery (ASR) backup on a cluster node. However, Backup Operators can perform this procedure if that group is added to the security descriptor for the Cluster service. You can do that using Cluster Administrator or cluster.exe. For more information, see Give a user permissions to administer a cluster and Cluster.

2. Back up the cluster disks from each node.
3. Back up each individual application (for example, Microsoft Exchange or Microsoft SQL Server) running on the nodes.

For more information, see Backing up and restoring server clusters.

# Cluster automation interface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster automation interface**

The cluster automation interface is a Component Object Model (COM) interface to the Cluster API. You can use the Cluster Automation Interface with the Windows Script Host (WSH) to automate many administrative tasks.

For more information on the Cluster API, see the Microsoft Platform Software Development Kit (SDK).

# General storage configuration tips for server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**General storage configuration tips for server clusters**

Before adding or removing hardware from a server cluster, make sure the drive letter assigned to the system partition (typically C or D) is fixed. The drive-letter assignment must remain constant because references to this drive letter are recorded in the system registry, and your server cluster is affected by these references.

It is also necessary for drive letters on the cluster storage device to remain constant and for the drive-letter assignments to be the same on each cluster node. In addition, drive letters for the cluster storage device must not conflict with any local drive letters.

If you assign drive letters to the shared cluster storage devices, the Windows Server 2003 family operating system limits the total number of such devices to 23. Mounted drives are not subject to this limit, so you can use mounted drives to access more than 23 shared cluster storage devices in your server cluster.

For more information on using mounted drives, see Using NTFS mounted drives.

For information about how to obtain product support, see Technical support options.

# Backing up and restoring server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Backing up and restoring server clusters**
Performing regular backups of your server cluster is imperative for high availability. This topic explains how you can use the Backup or Recovery Wizard to back up cluster nodes, describes ten cluster failure scenarios, and offers data restore solutions for each scenario using the Backup or Restore Wizard and recovery utilities from the *Microsoft Windows Server 2003 Resource Kit*.

For more information on backup and restore procedures, see Backing up and restoring data.

**Backing up cluster data**
In a server cluster, there are four groups of data critical to the proper operation of the cluster; the disk signatures and partitions of the cluster disks, the cluster quorum data, the data on the cluster disks, and the data on the individual cluster nodes.

- Cluster disk signatures and partitions
- Cluster quorum data
- Data on the cluster disks
- Data on the individual cluster node

**Cluster disk signatures and partitions**
Before you begin to back up any data on the server cluster nodes, make sure you backup the cluster disk signatures and partitions using Automated System Recovery in the Backup Wizard. This step is necessary if you later need to restore the signature of the quorum disk, for example, if you experience a complete system failure, and the signature of the quorum disk has changed since you last backed up.

### Note

- By default, Backup Operators do not have the user rights necessary to create an Automated System Recovery (ASR) backup on a cluster node. However, Backup Operators can perform this procedure if that group is added to the security descriptor for the Cluster service. You can do that using Cluster Administrator or cluster.exe. For more information, see Give a user permissions to administer a cluster and Cluster.

For information, see Back up cluster disk signatures and partition layouts.

**Cluster quorum data**
When you back up data on a server cluster node, make sure you also back up the cluster quorum. The cluster quorum is important because it contains the current cluster configuration, application registry checkpoints, and the cluster recovery log.

You can use the Backup Wizard to back up the cluster quorum data if you perform a System State backup from any node provided the Cluster service is running on that node.

For information, see Back up the cluster quorum.

**Data on the cluster disks**
To back up all cluster disks owned by a node, perform a full backup from that node.

You can also back up this data through a network connection to a hidden administrative file share. For example, you might use the New Resource Wizard to create FBackup$, GBackup$, and HBackup$ file shares for the root of drives F, G, and H, respectively. These shares would not appear in the browse list and could be configured to allow access only to members of the Backup Operators group.

For information on backing up data on the cluster disks, see Back up data on cluster nodes.

### Important

- If a cluster disk owned by the node being backed up fails over to another node during the backup process, the backup set will not contain a full backup of that disk.

**Important**

- You can only backup a cluster disk on a local node. You cannot backup a cluster disk on a remote computer.

**Data on the individual cluster nodes**

After you back up the cluster quorum disk on one node, it is not necessary to back up the quorum on the remaining cluster nodes. However, you may want to back up the clustering software, cluster administrative software, system state, and application data on the remaining nodes.

**Important**

- If you back up the system state for a node, you will also automatically back up the quorum data as long as the Cluster service is running on that node.

For information on backing up data on individual cluster nodes, see Back up data on cluster nodes.

**Cluster failure and restore scenarios**

This section describes ten failure scenarios that will require restoring your cluster. The type of failure you experience determines the steps you must follow.

- Scenario 1—Cluster Disk Data Loss
- Scenario 2—Cluster Quorum Corruption
- Scenario 3—Cluster Quorum Loses Checkpoints
- Scenario 4—Cluster Disk Corruption or Failure
- Scenario 5—Cluster Quorum Disk Failure
- Scenario 6—Single Cluster Node Corruption or Failure
- Scenario 7—Cluster Quorum Rollback
- Scenario 8—Complete Cluster Failure
- Scenario 9—Majority Node Set Cluster Failure
- Scenario 10—Application Data Loss in a Server Cluster

**Scenario 1—Cluster Disk Data Loss**

If you have lost files and folders on one of your cluster disks, but not on the disk containing the cluster quorum, you can use the Backup or Restore Wizard to restore that data.

**Important**

- You must restore the cluster disk data from the node that owns the cluster disk.

For information, see Restore files from a file or a tape.

**Scenario 2—Cluster Quorum Corruption**

**Symptom:** The cluster nodes can boot up, but the Cluster service fails to start because the quorum resource cannot come online.

If this problem results from corrupted files on the quorum disk, try starting the Cluster service by opening a command prompt

and typing **net start clussvc /resetquorumlog**. This creates a new quorum log file, using information stored in the cluster database on the local node. For additional information about recovering from cluster quorum corruption, see Recover from a corrupted quorum log or quorum disk. If the cluster quorum disk needs to be replaced, see Scenario 5, below. For a majority node set cluster, see Scenario 9, below.

**Scenario 3—Cluster Quorum Loses Application Checkpoints**
**Symptom:** Some resources fail to come online and the application checkpoints are out of date.

If you have recovered from quorum corruption by creating a new quorum log as described in Scenario 2 above, you may need to restore the matching checkpoints before the quorum resource can come back online.

**Using the Microsoft Windows Server 2003 Resource Kit tools**

- Use the ClusterRecovery utility. For more information, see the *Microsoft Windows Server 2003 Resource Kit* and the Help for the ClusterRecovery utility.

**Using Windows Server 2003 family tools**

- Locate backup sets of the cluster quorum disk and system disk. The cluster quorum disk backup will contain other system state data. These backups need to be a matched set, that is make sure they were taken at the same time.
- Restore data on the system disk. For information, see Restore files from a file or a tape.
- Restore the cluster quorum. For information, see Restore the contents of a cluster quorum disk for all nodes in a cluster.

**Scenario 4—Cluster Disk Corruption or Failure**
**Symptom:** A cluster disk cannot come online. Resources that depend on that cluster disk will not be able to come online.

First, see if you can run a diagnostic utility from the disk manufacturer to determine the condition of the disk. If the cluster disk is corrupted or the disk hardware fails, you can restore the disk more quickly by using utilities in the *Microsoft Windows Server 2003 Resource Kit*. If you do not have access to these tools, you can still restore your cluster disk using the Backup and Recovery utilities included with Windows Server 2003 family operating systems.

**Using the Microsoft Windows Server 2003 Resource Kit tools**

- Use the ClusterRecovery utility. For more information, see the *Microsoft Windows Server 2003 Resource Kit* and the Help for the ClusterRecovery utility.
- Use NTBackup along with the Confdisk utility (from the *Microsoft Windows Server 2003 Resource Kit*) to restore the data on the cluster disk. For more information, see the *Microsoft Windows Server 2003 Resource Kit*.

**Using Windows Server 2003 tools**

- If necessary, replace the cluster disk. For information, see Install local storage buses and devices.
- Stop the Cluster service on all nodes of the cluster.
- Locate the data backup set for the node that owns the cluster disk. Also, locate the Automated System Recovery backup set for that node, if it is available. Perform an Automated System Recovery restore on a node. Use ASR as a last resort in system recovery, only after you have exhausted other options. For more information, see Restore a damaged cluster node using Automated System Recovery.
- After the restored node comes back online, restart the Cluster service on the remaining nodes.

**Scenario 5—Cluster Quorum Disk Failure**
**Symptom:** The cluster nodes can boot up, but the Cluster service fails to start because the quorum resource cannot come online. Entries in the Event Log indicate hardware failures.

First, try starting the Cluster service by opening a command prompt and typing **net start clussvc /fixquorum**. This starts the Cluster service with all resources offline, including the quorum resource. Then you can try switching to a new quorum

resource, with or without using the Clusterrecovery utility in the Windows Server 2003 Resource Kit. For more information, see Fixquorum command.

If the cluster quorum disk (the disk containing the quorum resource) fails, you can replace it more quickly by using utilities in the *Microsoft Windows Server 2003 Resource Kit*. If you do not have access to these tools, you can still replace your cluster quorum disk using the backup and restore utilities shipped with Windows Server 2003 family operating systems.

**Using the Microsoft Windows Server 2003 Resource Kit tools**

- Use NTBackup along with the Confdisk utility (from the *Microsoft Windows Server 2003 Resource Kit)* to restore the data on the cluster disk.

- Use the ClusterRecovery utility. For more information, see the *Microsoft Windows Server 2003 Resource Kit* and the Help for the ClusterRecovery utility.

**Using Windows Server 2003 family tools**

- If necessary, replace the cluster quorum disk. For information, see Install local storage buses and devices.

- Locate the data backup set for the cluster quorum. Also, locate the Automated System Recovery backup set for the node you used to create the cluster quorum backup set, if it is available. Perform an Automated System Recovery restore on any node in the cluster. Use ASRas a last resort in system recovery, only after you have exhausted other options. For more information, see Restore a damaged cluster node using Automated System Recovery.

- Restore data on the node. For information, see Restore files from a file or a tape.

**Scenario 6—Single Cluster Node Corruption or Failure**

**Symptom:** The node cannot join the cluster.

If the Event Log indicates that the cluster database on the local node is corrupted, you can perform a System State restore on that node to replace the local cluster database. For information, see Restore the cluster database on a local node. Alternatively, you can copy the latest checkpoint file (CHKxxx.TMP) from the quorum disk to the %systemroot%\Cluster\ directory, rename it as file CLUSDB, and restart the Cluster service on that node.

If a single node fails in the cluster due to system disk or other hardware failure, follow these steps to rebuild the node and rejoin the cluster:

- After verifying that all cluster resource groups have been successfully moved to other nodes, repair or replace the failed hardware. For information, see To Move a group to another node and To Manage Cluster Hardware.

- Perform an Automated System Recovery restore on the failed node to rebuild the node. For information, see To Restore a damaged cluster node using Automated System Recovery.

- If you have other files or application data for that node backed up on on tape or other backup medium, you can restore that now. For information, see To Restore files from a file or a tape and Scenario 8 below.

- For each cluster group and resource, verify that the newly recovered node appears as a possible owner in Cluster Administrator, then move a resource group to the newly recovered node and verify that the move is successful. For information, see To test Test whether group resources can fail over.

  **Note**

  - If you do not have an Automated System Recovery backup of the node, you can evict that node and add a new node to the cluster. For more information, see To Evict a node from the clusterand To Add additional nodes to the cluster.

**Scenario 7—Cluster Quorum Rollback**

If recent changes to your cluster have resulted in the cluster not functioning as expected, you can use the Backup or Restore Wizard to roll back your cluster to a previous configuration. For example, if a number of resources have mistakenly been

deleted from the cluster configuration, you can roll it back, using a backup that contains those resources.

For information, see Restore the contents of a cluster quorum disk for all nodes in a cluster.

**Scenario 8—Complete Cluster Failure**

**Symptom:** None of the nodes can boot up.

If all nodes fail in a cluster and the quorum disk cannot be repaired, follow these steps:

- Use Automated System Recovery on one node in the original cluster, choosing a node that was backed up recently and that was active in the cluster at the time it was backed up. This restores the disk signatures, the partition layout of the cluster disks (quorum and nonquorum), and the cluster configuration data. Do not start other nodes until the first node is restored. For more information, see To Restore a damaged cluster node using Automated System Recovery.

- Restore other nodes. For more information, see Restore a damaged cluster node using Automated System Recovery.

- Restore your applications and application data from backup data sets.

### Important

- If you do not have an Automated System Recovery backup of each node, you cannot restore the cluster. Instead, you must recreate your cluster from scratch. For more information, see Checklist: Planning and creating a server cluster.

**Scenario 9—Majority Node Set Cluster Failure**

The methods for restoring a majority node set cluster are the same as for restoring other clusters. However, in a majority node set cluster, if some of the nodes fail, and the cluster loses quorum, you can force the remaining nodes to form a quorum and restart the cluster. For more information, see To Force quorum in a majority node set server cluster.

### Note

- On a majority node set cluster, the cluster database is not stored on a cluster disk central to all nodes, but is instead stored locally on each node at %systemroot%\Cluster\MNS.%ResourceGUID%$\%ResourceGUID%$\MSCS\.

**Scenario 10—Application Data Loss in a Server Cluster**

When restoring application data in a server cluster, follow the instructions provided in the documentation that shipped with your application.

### Important

- If you are backing up Microsoft Exchange Server, a newer version of NTBackup.exe may be available from the Exchange Server section of the Microsoft Web site. Otherwise, you can use the version of NTBackup.exe that is included with Windows Server 2003 family operating systems.

# Uninstalling a server cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Uninstalling a server cluster**

You can uninstall or remove a node of a server cluster at any time. Before uninstalling a node, take all groups offline or move them to one or more other nodes. Also, close Cluster Administrator.

For instructions on how to remove a node from a server cluster, see Evict a node from the cluster.

# Cluster Hardware Configuration

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster hardware configuration**
This section covers:

- Fibre Channel hardware configuration
- Network Configuration in a Server Cluster

# Fibre Channel hardware configuration

Updated: April 10, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Fibre Channel hardware configuration**

When using Fibre Channel for a cluster, note the following details:

- The Fibre Channel systems are listed in a separate section of the Windows Catalog at Support resources. To list all compatible Fibre Channel devices, select **All** in **Device Type**, then in **Model** search for **Fibre Channel**.
- Your Fibre Channel disks are managed just like any other disk connected to a storage bus.
- The Fibre Channel controller ID is preconfigured with a unique ID. This simplifies your hardware setup because you do not need to manually configure each cluster node with a separate ID.

The figure shows a Fibre Channel cluster with a fiber switch or hub. Fiber optic cable runs between a node's Fibre Channel controller and the fiber switch or hub. The device bus between the switch or hub is fiber.




For more information about the types of storage you can use in a server cluster running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, see Server clusters overview.

# Network Configuration in a Server Cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network configuration in a server cluster**
This section covers:

- Configuring cluster network hardware
- Private network addressing options

# Configuring cluster network hardware

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring cluster network hardware**

**Planning your cluster network hardware and network infrastructure**
Before you configure your cluster network hardware, make sure that your networking hardware and domain infrastructure meets all the requirements for a server cluster. For more information, see Planning your networking hardware and Choosing a domain model.

**Network Interface Controller (NIC) configuration**
Follow the guidelines below to correctly configure your cluster networking hardware:

**Requirements**

- Your network adapters must use the same configuration settings.
- Each network is on a single subnet.
- At least one of the cluster networks enabled for internal communication (for example, either a private or mixed network) must be configured using static addresses.

**Recommendations**

- If an adapter is set to use Auto Select, it may drop packets while automatically negotiating the network settings. Manually set the speed and duplex mode for multiple speed adapters to the same values and settings (set these properties on the **Advanced** tab of the network adapter's Property Page). If the adapters are connected to a switch, ensure that the port settings of the switch match those of the adapters. For more information, see Change network adapter settings.
- Use static IP addresses for all adapters. Choose private IP addresses following the guidelines at Private network addressing options (that is, specify a class A, B, or C private address for each). For public networks, also use static IP addresses. It is recommended that you do not use DHCP leases for public network addresses because cluster operations might be disrupted. For more information, see The Cluster service and TCP/IP.

  The use of Automatic Private IP Addressing (APIPA) to configure a cluster network is not supported. APIPA is not designed for use in computers that are attached to multiple networks. For more information, see TCP/IP configuration methods.
- Do not configure a default gateway or DNS or WINS server on the private network adapters.
- Configure WINs and/or DNS servers and a default gateway on the public network adapters. If Network Name resources are used on the public networks, set up the DNS server to support dynamic updates; otherwise the Network Name resources may not fail over correctly. For more information, see Configure TCP/IP settings.
- Set the order of the network adapter binding as follows:

  1. External public network
  2. Internal private network (Heartbeat)
  3. [Remote Access Connections]

  For more information, see Modify the protocol bindings order.

- Clearly identify each network by changing the default name. For example, you could change the name of the private network connection from the default *Local Area Connection* to *Private Cluster Network*.

  For more information on networks and network roles, see Server cluster networks.

**Configuring logically multihomed network adapters**

A logically multihomed network adapter is one that has two IP addresses assigned to it. These adapters can be used only for node-to-node cluster communication if their primary addresses are on the same IP subnet. If the primary addresses on all nodes are not on the same IP subnet, reorder the IP addresses assigned to the adapter by using the Network Connections folder.

# Private network addressing options

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Private network addressing options**

If an interconnect connects only the cluster nodes and does not support any other network clients, you can assign it a private Internet protocol (IP) network address instead of using one of your enterprise's official IP network addresses.

By agreement with the Internet Assigned Numbers Authority (IANA), several IP networks are always left available for private use within an enterprise. These reserved numbers are:

- 10.0.0.0 through 10.255.255.255 (Class A)
- 172.16.0.0 through 172.31.255.255 (Class B)
- 192.168.0.0 through 192.168.255.255 (Class C)

You can use any of these networks or one of their subnets to configure a private interconnect for a cluster. For example, address 10.0.0.1 can be assigned to the first node with a subnet mask of 255.0.0.0. Address 10.0.0.2 can be assigned to a second node, and so on. Do not specify any default gateway, DNS or WINS servers for this network.

Ask your network administrator which of the private networks or subnets you can use within your enterprise before configuring your cluster.

Never route these private network addresses. For more information on private IP network addresses, see Resources for Server Cluster information.

# Standard Resource Types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Standard resource types**
Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition come with numerous resource types that you can use to manage cluster resources.

Internet Information Services (IIS) can be installed on a server cluster. For more information on how IIS runs on a server cluster, view the IIS online documentation.

If you have IIS installed, you can view the online documentation by typing **http://localhost/iisHelp/** in your browser address bar and pressing ENTER.

You can install Message Queuing on a server cluster. For more information on the Message Queuing and Message Queuing Triggers resource types and how they run on a server cluster, see Microsoft Message Queuing.

You can install Distributed Transaction Coordinator (DTC) on a server cluster. For more information on the Distributed Transaction Coordinator resource type and how it runs on a server cluster, see Distributed Transaction Coordinator.

**Note**

- IIS Server Instance, Simple Mail Transfer Protocol (SMTP) Server Instance, and Network News Transfer Protocol (NNTP) Server Instance resource types, previously available on Windows 2000 cluster nodes, are not supported on Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition cluster nodes.

This section also covers the following resource types and topics:

- Physical Disk resource type
- DHCP Service and WINS Service resource types
- Print Spooler resource type
- File Share resource type
- Internet Protocol Address resource type
- Local Quorum resource type
- Majority Node Set resource type
- Network Name resource type
- Generic Application resource type
- Generic Script resource type
- Generic Service resource type
- Volume Shadow Copy Service resource type

# Physical Disk resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Physical Disk resource type**

You use the Physical Disk resource type to manage disks that are on a cluster storage device.

You can assign drive letters to the disks that are on a cluster storage device or create mounted drives. If you assign drive letters, the disks must have the same drive letters on all cluster nodes.

To prevent data corruption, it is important that only one node at a time use a cluster disk. Normally, the Cluster service prevents more than one node from using a disk at the same time. However, there are two circumstances in which the ordinary cluster safeguards do not work: when you install a new disk, and when you remove a disk from control by the cluster.

For information on replacing disks in a server cluster, see Updated technical information and article Q305793, "How to Replace a Disk with Windows 2000 or Windows Server 2003 family Clusters" in the Microsoft Knowledge Base.

# DHCP Service and WINS Service resource types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**DHCP Service and WINS Service resource types**

Dynamic Host Configuration Protocol (DHCP) and Windows Internet Name Service (WINS) provide network services to clients. DHCP leases Internet protocol (IP) addresses to clients. This lets you centrally manage IP addresses and related configuration details. For more information on DHCP, see DHCP. WINS provides a dynamic database service that can register and resolve NetBIOS names to IP addresses used on your network. For more information on WINS, see WINS.

Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition provide resource types that allow you to manage DHCP and WINS services in a cluster. Each of these resources stores a database on disk. In order for these resources to fail over, the database must be kept on a disk that is on cluster storage.

You must configure a location for the DHCP and WINS database files and for the backup of these files.

The path to the DHCP or WINS database files and backup of these files must:

- End with a backslash (\).
- Use a drive letter of a disk resource on which the DHCP or WINS resource depends.
- Cannot begin with *systemroot*, although you can use other environment variables.

The path for the DHCP or WINS database will be of the form f:\dhcp\ or f:\wins\.

The path to the backup of the DHCP database files will be of the form f:\dhcp\backup\.

For WINS, you provide a root path, and the backup of the WINS database files is created in a subdirectory of this path. The path will be of the form f:\wins\.

# Print Spooler resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Print Spooler resource type**

You use the Print Spooler resource to cluster your print services. You can only cluster a printer that is connected to your servers by way of a network. If the printer is connected by way of a local port, such as LPT1, there is no way to fail over control of that printer in the case of failure of the server to which it is connected.

You cluster print services using the Print Spooler resource. Multiple Print Spooler resources can exist in a cluster, but there can be no more than one Print Spooler resource per group. Clients access clustered printers in exactly the same way they access nonclustered printers. Clients can use either the network name or Internet protocol (IP) address as the server name.

When a group containing a Print Spooler resource fails over to another node, the document that is currently being spooled to the printer is restarted from one of the other nodes after failover. When you move a Print Spooler resource or take it offline, the Cluster service waits until all jobs that are spooling to the printers are finished (or until the configured wait time has elapsed). Documents that are spooling from an application to a Print Spooler resource are discarded and must be respooled (or reprinted) to the Print Spooler resource if the group containing the Print Spooler resource fails over before the application has finished spooling.

In addition to creating a Print Spooler resource, you must properly configure your printer ports and printer drivers to function on a cluster. To properly set up a printer on a cluster, you must install printer drivers and a printer port on the same virtual server as the Print Spooler resource. The Cluster service will replicate the printer drivers from the virtual server to all cluster nodes the first time the Print Spooler resource fails over to another node.

For more information on the entire printing process in Windows Server 2003 family operating systems, see Understanding Printing. For detailed instructions on how to configure a printer to operate in a cluster, see Create a cluster-managed printer.

# File Share resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File Share resource type**

If you want to use a server cluster to provide a high availability file server you need to use the File Share resource type. You can use this resource type in three major ways:

- Normal file share

  You can use the File Share resource type to provide a basic file share. Using this method, a single file folder is published to the network under a single name and functions much as an Server Message Block (SMB) share does.

- Share or hide subdirectories

  You can also use the File Share resource to publish several network names: one each for a file folder and all of its immediate subfolders. This is an efficient way to create large numbers of related file shares on a single file server. An example application for this is creating a file share for each user who has files on the server. For more information on why you would want to share subdirectories, see Using a server cluster with large numbers of file shares. You can also prevent users from seeing the subdirectory shares when they browse the network. For more information on sharing or hiding subdirectories, see Share subfolders on a File Share resource.

  **Important**

  - When you share subdirectories, do not use the same name for subfolders under different folders (for example, \folder1\name1 and \folder2\name1). The first shared subdirectory to come online will remain online, but the second shared subdirectory will not be initialized and will not come online.

- DFS root

  You can use the File Share resource type to create a resource that manages a stand-alone DFS root. Fault-tolerant DFS roots cannot be managed by this resource. The DFS root File Share resource has required dependencies on a network name and an Internet Protocol (IP) address. The network name can be either the cluster name or any other network name for a virtual server.

  A cluster-managed DFS root is very different from an Active Directory (or domain-based) DFS root. If the data that needs to be replicated is not changed very often, using a domain-based DFS root with replicas can be a better choice than a cluster-managed DFS root for providing high availability. If the data set changes frequently, replication is usually undesirable, and a cluster-managed DFS root is recommended.

  In Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, server clusters support multiple standalone DFS roots. These multiple DFS roots can exist in multiple resource groups and each group can be hosted on a different node.

  **Important**

  - Clusters running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition on some nodes and Windows 2000 on others will only support one DFS root per cluster.

  - Do not make DFS configuration changes (for example, by adding new links) while operating a mixed-version cluster.

  If you create multiple DFS roots in a single cluster, be aware of the following restrictions:

  - The name you specify for one DFS root cannot overlap with the names for other DFS roots in the same cluster. For example, if you named one DFS file share c:\dfsroots\root1, you could not share out other roots using names

derived from the first DFS root, for example, c:\dfsroots, c:\dfsroots\root1, or c:\dfsroots\root1\root2.

- The DFS roots must be hosted on NTFS partitions.
- The name you give a non-clustered DFS root on a node's local storage must be different from the name you give a clustered DFS root on the cluster storage.

For more information on DFS in a server cluster, see article Q301588, "Using DFS on a Server Cluster to Maintain a Single Namespace" in the Microsoft Knowledge Base.

For more information on DFS, see Distributed File System (DFS).

When you create a new File Share resource, permissions for the Everyone group for that file share are automatically set by default to Read only. You can use Cluster Administrator or Cluster.exe to change the default permissions. For more information about using Cluster Administrator for this purpose, see View or set resource properties. For more information about using Cluster.exe for this purpose, see information about private property names for a File Share resource in Cluster resource.

You can also cache files from a cluster file share onto client computers. Clients can then have access to these files even when the client computer is disconnected from the network. For more information, see Set client-side caching for a File Share resource.

For more information on creating a highly available file share, see Create a cluster-managed file share.

You can also encrypt the data on your clustered file shares using Encrypting File System (EFS). For more information, see Create a cluster-managed encrypted file share.

# Internet Protocol Address resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet Protocol Address resource type**

You use the Internet Protocol (IP) Address resource type to manage IP network addresses as cluster resources. When an IP Address resource is included in a group with a Network Name resource, the group can be accessed by network clients as a virtual server.

For more information on virtual servers, see Virtual servers.

# Local Quorum resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Local Quorum resource type**

You use the Local Quorum resource type to manage the system disk on the local node as the quorum resource. The Local Quorum resource cannot fail over to another node. If you create a new cluster on a node that is not connected to a cluster disk , the New Server Cluster Wizard will, by default, create a single-node cluster using the Local Quorum resource .

One scenario where a single-node cluster can be useful is in the development and testing of cluster-aware applications or custom resource types. For more information on how to create a cluster-aware application or custom resource type, see the *Microsoft Platform Software Development Kit* (SDK).

You can migrate an existing multinode cluster to a single-node cluster by creating a Local Quorum resource on one node, evicting all other nodes, and using the Local Quorum resource as the quorum resource. This can be useful if you need to troubleshoot cluster problems, for example on the cluster disk.

# Majority Node Set resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Majority Node Set resource type**
You use the Majority Node Set resource type to manage cluster configuration data that may or may not be on a cluster storage device. For example, the Majority Node Set resource type can manage cluster configuration data that is actually stored on multiple nodes in a cluster at the same physical location or in a geographically dispersed cluster. The Majority Node Set resource ensures that the cluster configuration data is kept consistent across the different nodes. In contrast, in a single quorum device server cluster, the quorum resource maintains the definitive copy of the cluster configuration data on the cluster storage device.

**Important**

- Only one Majority Node Set resource can exist in a server cluster.
- Do not use the **net use** command to connect to the share containing the cluster database files (%systemroot%\Cluster\MNS.%ResourceGUID%$\%ResourceGUID%$\MSCS) in a majority node set server cluster. Otherwise, the Majority Node Set resource (the quorum resource) may be unable to connect to the share and may fail to come online.

For detailed instructions on creating a majority node set cluster, see Create a cluster.

For information on differences between a single quorum device server cluster and a majority node set server cluster, see Choosing a Cluster Model.

# Network Name resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network Name resource type**

You use the Network Name resource type to provide an alternate computer name for an entity that exists on a network. When included in a group with an Internet Protocol (IP) Address resource, a Network Name resource provides an identity to the group, allowing the group to be accessed by network clients as a virtual server.

For more information on virtual servers, see Virtual servers.

# Generic Application resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Generic Application resource type**
You use the Generic Application resource type to manage cluster-unaware applications.

**When to use the Generic Application resource type**
Not all applications will function effectively in a cluster. The minimum requirements for an cluster-unaware application are:

- The application must be able to store its data in a configurable location.
- The application must use TCP/IP to connect with clients.
- Clients must reconnect in the event of intermittent network failure.

**Installing a generic application**
If your cluster-unaware application is to work during failover, you must make sure the application can run from all cluster nodes. There are two strategies for accomplishing this:

1. Install the application onto a cluster disk. Because all nodes use the same drive letters for cluster disks, this solution guarantees that the application can be started from all nodes using the same command line.
2. Install separate copies of the application on all cluster nodes. Make sure that the application is installed in a folder on each node that has the same drive letter and path as all other nodes. Although this solution uses more disk space, it is the only way you can install an application that will allow you to perform a rolling upgrade of that application.

# Generic Script resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Generic Script resource type**

You use the Generic Script resource type to manage Windows Server 2003 family operating system scripts as cluster resources. Similar to the Generic Application resource type, the Generic Script resource type provides only the most basic functionality. If the script is running, it is presumed to be online. For more information on creating a generic script using Windows Script Host, see the *Microsoft Platform Software Development Kit* (SDK). To provide greater functionality, you can define a custom resource type. For more information on how to create a custom resource type, see the *Microsoft Platform Software Development Kit* (SDK).

# Generic Service resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Generic Service resource type**

You use the Generic Service resource type to manage Windows Server 2003 family operating system services as cluster resources. Similar to the Generic Application resource type, the Generic Service resource type provides only the most basic functionality. For example, the failure of a Generic Service resource is determined by a query of the Service Controller. If the service is running, it is presumed to be online. To provide greater functionality, you can define a custom resource type. For more information on how to create a custom resource type, see the *Microsoft Platform Software Development Kit* (SDK).

# Volume Shadow Copy Service resource type

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Volume Shadow Copy Service Task resource type**

With the Volume Shadow Copy Service Task resource type, you can create jobs in the Scheduled Task folder that must be run on the node that is currently hosting a particular resource group. In this way, you can define a scheduled task that can failover from one cluster node to another. However, in the Microsoft® Windows Server 2003 family of products, the Volume Shadow Copy Service Task resource type has limited capabilities for scheduling tasks and serves primarily to support Shadow Copies of Shared Folders in a server cluster. If you need to extend the capabilities of this resource type, consider using the Generic Script resource type. For more information on the Generic Script resource, see Checklist: Installing a Generic Script resource.

For more information on shadow copies, see Shadow Copies for Shared Folders.

For more information on using the Volume Shadow Copy Service Task resource to manage shadow copies, see Using Shadow Copies of Shared Folders in a server cluster.

# Managing Security in a Cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing security in a cluster**
The Cluster service relies on Windows Server 2003 family security at all levels for securing communication between nodes, and controlling authentication and access, including client access to the cluster and cluster resources. You can use Cluster Administrator to specify which users or groups can:

- Administer the cluster.
- Access files or folders on File Share resources.

For more information on setting server cluster administration permissions, see Give a user permissions to administer a cluster.

For information on how to change the account or the password for the Cluster service, see Change the account under which the Cluster service runs and Change the Cluster service account password.

This section covers common security tasks:

- Limiting client access to cluster resources
- Securing shared data in a cluster
- Domain accounts and the Cluster service
- Disk resource security
- Using Windows Firewall with a server cluster

# Limiting client access to cluster resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Limiting client access to cluster resources**

The following table lists the Windows Server 2003 family security features you can use to control client access to cluster resources.

**Important**

- In general, grant permissions to the individual applications and services as you would normally. However, always use Cluster Administrator to set File Share resource security policies. File Share security policies configured using Windows Explorer are lost when the File Share resource fails over.

| Resource type | Windows Server 2003 family security |
|---|---|
| Dynamic Host Configuration Protocol (DHCP) Service | Windows Server 2003 family DHCP security. |
| File Share | NTFS or share-level security configured with a cluster management application. NTFS file permissions strongly preferred. Only NTFS file permissions for subfolder shares. For information about administering DFS shares, see the DFS documentation. |
| Generic Application | Windows Server 2003 family network authentication and NTFS security applied to the shared file system.<br><br>Note that when using Cluster Administrator to configure the parameters for a Generic Application resource, we recommend that you avoid selecting **Allow application to interact with desktop** unless it is necessary. |
| Generic Service | Determined by Windows Server 2003 family service configuration. |
| Generic Script | NTFS file-level security for execute permissions on files and permissions for APIs called in the script. |
| IP Address | N/A |
| Distributed Transaction Coordinator | Controlled by Microsoft Distributed Transaction Coordinator (MS DTC). |
| Message Queue Service | Controlled by Message Queue Service. |
| Network Name | N/A |
| Physical Disk/Local Quorum | NTFS file-level security. |
| Print Spooler | Windows Server 2003 family network and print-level permissions. |
| WINS Service | Windows Server 2003 family WINS security. |

# Securing shared data in a cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Securing shared data in a cluster**

**Securing data on cluster disks**
You can control access to folders and files of File Share resources on a user or group basis. You must set access to File Share resources through Cluster Administrator. The account used to start the Cluster service must have appropriate access to the directories with File Share resources. At minimum, this account must have NTFS Read permissions.

By default, access to cluster file shares is disabled to anonymous users. To allow anonymous access to specific file shares you can either enable Kerberos authentication on the Network Name resource associated with the file share or change the local security policy setting. For more information on configuring these Kerberos properties, see Enable Kerberos authentication for virtual servers.

**Important**

- If you change file share permissions using Windows Explorer or My Computer instead of using **Permissions** on the **Parameters** tab in Cluster Administrator, the permissions are lost when the resource is taken offline.

When setting access permission on File Share resources, specify permissions based on, at minimum, a local Administrator group account; do not specify permissions based on local user or local group accounts. On member servers, local users and local groups have security context only on the local computer. The security context of these accounts and groups is meaningless when failed from one node to another. For this reason, you cannot give local users or local groups permissions to administer the cluster. The single exception to this rule is the local Administrators group. This is not a problem on domain controllers because the local accounts and groups have security context on all member servers in the domain.

**Important**

- The Cluster service must always have full access to the folder on the quorum disk containing the quorum log (the folder is named MSCS by default). (The quorum log file is used to write all cluster state and configuration changes that cannot be committed to the other nodes.)

**Encrypting shared data**
You can further secure the contents of the data on your cluster disks by using Encrypting File System (EFS). EFS prevents unauthorized users that may gain physical access to your cluster disks from reading the data on those disks. For more information, see Create a cluster-managed encrypted file share.

**Auditing access to shared data**
Use standard security features to audit access to files and folders that reside on drives on the cluster storage. Use either My Computer or Windows Explorer to set file or directory auditing. You cannot do this using Cluster Administrator. For more information on auditing, see Auditing Security Events.

Security events are written to the security log on the node that owns the File Share resource.

Audit settings are specific to a node. Therefore, if you enable auditing on one node but not on another, you only see security events when the node with auditing enabled owns the File Share resource. If you want to audit access to shared data, it is a best practice to enable auditing on all cluster nodes.

**Taking ownership of files or folders**
Use either My Computer or Windows Explorer to set file or directory ownership. You cannot do this using Cluster Administrator.

# Domain accounts and the Cluster service

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Domain accounts and the Cluster service**

You must log on using an account that is a member of an Administrators group to create the first node in a cluster or add additional nodes. All nodes must have joined a domain (instead of a workgroup) and be members of the same domain. When you create a cluster or join nodes to a cluster, you specify the domain user account under which the Cluster service runs. It is best if the Cluster service user account is a new account. The account must have local administrative rights and permissions on the cluster nodes. In addition, make sure that applications (for example, Microsoft SQL Server 2000) do not use the same account as the Cluster service account. Be sure to keep the password from expiring on the account (follow your organization's policies for password renewal).

When you install the operating system on a member server, Setup adds the Cluster service account as a member of the local Administrators group. When you install the operating system on a domain controller, Setup gives cluster-administrative permissions to the Domain Administrators group.

When you use Cluster Administrator to administer a cluster from a remote location, such as from another server or a computer running Windows XP Professional, or to administer a cluster locally from a node, you must use an account that is a member of the local Administrators group on all nodes. For more information, see Give a user permissions to administer a cluster.

**Note**

- You can use an account that is a domain user account or a member of a global group to remotely or locally administer a cluster if you use cluster.exe or cluster scripting tools. For more information, see Cluster and the Microsoft Platform Software Development Kit (SDK).

If you have administrative permissions on all nodes of the cluster, you can fully administer the cluster. However, if you have only specific permissions to administer the cluster, you cannot change the cluster description, change cluster security, or register extensions to Cluster Administrator.

For information on changing the account used to run the Cluster service, see Change the account under which the Cluster service runs.

For information on changing the password for the Cluster service account, see Change the Cluster service account password.

# Disk resource security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Disk resource security**

There are two disk subsystems in a cluster environment: The first is the shared bus or buses, including the cluster disks, and the second is theWindows Server 2003 family operating system files. Keep these two systems separate. Do not put Windows Server 2003 family system files on the cluster disks. In addition, you can locate Windows Server 2003 family system files on either file allocation table (FAT) or NTFS drives, but you must locate the cluster disks on the shared bus only on NTFS-formatted drives.

**Securing access to the cluster disks**

Restrict access to the cluster disks to applications that are managed as a cluster resource. If there are other applications on a cluster node that are not managed by the cluster, install those applications either on the system disk or on a local disk. To help reduce the risk of unauthorized reads and writes to the quorum disk, it is recommended that you give access to the quorum disk only to the Cluster service account and members of the local Administrators group.

**Securing disk resources on cluster disks**

All standard Windows Server 2003 family file security options apply to the disk resources on the cluster disk. You can use Windows Explorer to set permissions on disk resources.

# Using Windows Firewall with a server cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Windows Firewall with a Server Cluster**
For a server cluster running Windows Server 2003 Service Pack 1, if you enable Windows Firewall, you must carry out the following steps on each cluster node:

- Learn about how to install and run the Security Configuration Wizard, an optional component that can be installed through Control Panel, by opening a command prompt and typing **hh scwhelp.chm**.

- Install and run the Security Configuration Wizard. Select the appropriate server roles, client features, and administration options for that server. Select only what is necessary, but be sure to include the following:

  - Among server roles, select **Cluster server** and any server roles that correspond to the clustered resources for that node. For example, for a node with a File Share resource, select **Cluster server** and **File server**.

  - Among client features, select those necessary for the node to function. Examples of client features are **Domain member**, **Microsoft networking client**, **DHCP client**, **DNS client**, and **WINS client**.

  - Among administration and other options, select **Join a cluster** and any other appropriate options.

  The Security Configuration Wizard helps ensure that ports needed by the Cluster service are open and the Cluster service and other necessary services can operate normally.

- In the Services snap-in, make sure that the Recovery action for the Cluster service is set to **Restart the Service** (the default setting). For more information, see Set up recovery actions to take place when a service fails.

  Note that this step will set the recovery action but will not immediately restart the Cluster service. If necessary, after taking the step, type **net start clussvc** to restart the service.

**See Also**
**Concepts**
Help: Windows Firewall

# Using Cluster Administrator

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Cluster Administrator**
Cluster Administrator shows you information about the groups and resources on all of your clusters and specific information about the clusters themselves. A copy of Cluster Administrator is automatically installed on your computer when you install Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition.

**Important**

- Cluster nodes must be running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. You cannot create a cluster or add nodes to an existing cluster on computers running Windows Server 2003, Standard Edition.

For remote administration, you can install copies of Cluster Administrator on other computers on your network running Windows server 2003 family operating systems. Cluster Administrator is available through Windows server 2003 Administration Tools, included on the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition compact disc sets. The remote and local copies of Cluster Administrator will be identical. It is also possible to administer a Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition server cluster remotely from a computer running Windows NT 4.0 Service Pack 3 or later using the Windows NT server, Enterprise Edition 4.0, Cluster Administrator tool or from a computer running Windows 2000.

This section describes the uses of Cluster Administrator, including the ways in which it changes procedures you might have been familiar with in a noncluster environment.

- Identifying failovers
- Installing Cluster Administrator on a remote computer

# Identifying failovers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Identifying failovers**

Failover can occur in one of two situations: An entire node fails, or a resource on a node fails (such as a network interface). Node failure is easily identifiable in Cluster Administrator. When you are connected to a cluster, Cluster Administrator displays a list of all nodes and their current status. Resource failure that does not bring down a node, but does result in group failover, can be harder to detect. Use Cluster Administrator or another cluster management application to periodically check the owner of all cluster groups to see if a group has failed over.

It is important to use Cluster Administrator to routinely monitor the status of all clusters and check for failover activities that diminish either the performance or the availability of resources. Failover is supposed to be easily tolerated. Your work flow will not suffer if all applications are working, failover policies have been correctly set, and you have provided sufficient capacity for all situations. If any of these three criteria are not met, you must take action to meet them. For information on how to plan your server cluster, failover policies, and capacity requirements, see Planning Your Server Cluster.

Serious failover situations can affect performance and availability. If a node fails and its companion node is unable to serve clients efficiently, the situation must be resolved immediately. If groups persistently fail over and fail back without surviving on either node for very long, availability to clients is completely lost.

**Responding to a failover**

When investigating possible problems with failovers, answer the following questions:

- How serious is the problem?

  For example, other resources could be suffering due to the increased load placed on the alternate node.

- What caused the failover?

  Many types of failure can cause failovers. Usually, these occur at the operating system and hardware levels. For help with determining the exact type of a failure, see Server Cluster Troubleshooting.

  To review what happened prior to and during a group failure, run Event Viewer and look at the application log, system log, and the security log (if applicable). This helps you determine what types of errors occurred.

- Are any changes needed to prevent this kind of event in the future?

  If your performance has seriously suffered as a result of a failover, review the capacity planning issues presented in Capacity Planning for Server Clusters.

  If you determine that the nodes did not efficiently transfer control of various groups and resources during failover, review the failover policies you have set. For example, the Cluster service could be attempting to bring a File Share resource online before the corresponding Physical Disk resource was brought online. This type of error is easily fixed by configuring dependencies.

The Cluster service records error messages about resource failure in the system log. The system log is found in Event Viewer. The **Event Properties** dialog box contains error message information. Information about a resource failure might appear as follows:

```
Date: 9/8/2001

Time: 3:22:49 PM

Type: Error

Source: ClusSvc

Category: Failover Mgr

Event ID: 1069
```

```
User: N/A

Computer: DTCENTER3

Description: "Cluster resource 'Disk V:' in Resource Group 'SQL' failed.

For more information, see Help and Support Center at http://go.microsoft.com/fwlink/events.asp."
```

Other applications can keep track of such errors and notify appropriate parties as necessary.

For more information about diagnosing and resolving specific kinds of failures and errors, see Server Cluster Troubleshooting. For more information on resource dependencies, see Server Cluster groups.

For information on how to solve failover problems, see Server Cluster Troubleshooting.

# Installing Cluster Administrator on a remote computer

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Cluster Administrator on a remote computer**
You can use Cluster Administrator to remotely administer a server cluster. Cluster Administrator runs on any computer that is running Windows NT 4.0 Service Pack 3 or later, Windows 2000, or a Windows Server 2003 family operating system. You can use Cluster Administrator on any computer running Windows 2000 or Windows Server 2003 family operating system to remotely create, join nodes, or manage a server cluster. On computers running Windows NT 4.0 Service Pack 3 or later, you can use Cluster Administrator just to remotely manage a server cluster.

**Important**

- Computers used to remotely administer a server cluster must be secure and restricted to trusted personnel. For more information, see Best practices for securing server clusters.
- When you remotely manage a cluster using Cluster Administrator or cluster.exe, you must verify that NetBIOS over TCP/IP (NetBT) is enabled on the client. For more information on enabling NetBIOS over TCP/IP, see Configure TCP/IP settings.

You use Windows Server 2003 Administration Tools to install Cluster Administrator on computers that are running Windows Server 2003 family operating systems that are not cluster nodes. On Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition, Cluster Administrator is installed when you install the operating system. For more information on Windows Server 2003 Administration Tools, see Windows Server 2003 Administration Tools Pack Overview.

You cannot use Windows Server 2003 Administration Tools to install Cluster Administrator on a computer running Windows NT 4.0. However, if you have the version of Cluster Administrator that was provided with Windows NT Server, Enterprise Edition 4.0, you can use that version from any Windows NT 4.0, Service Pack 3 or later computer to administer a cluster running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. Because Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition provide resource types that were not supported in Windows NT 4.0, you might see error messages if you administer a cluster running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition from a computer running Windows NT 4.0.

If you are adding or removing users from the list of those who can administer a cluster, or otherwise changing the cluster security descriptor, the server from which you run Cluster Administrator must run one of the following versions of the operating system and service pack:

- Windows Server 2003
- Windows Server 2003 with Service Pack 1
- A product in the Windows 2000 Server family with Service Pack 2 or later

**Caution**

- When remotely managing a server cluster from another computer, the system default locale on that remote computer can be different from the system default locale on all the nodes in the cluster only if the names of the cluster and all the cluster nodes use English language characters. For more information on changing the system default locale, see Change number, currency, time, and date settings.

# Using a server cluster with large numbers of file shares

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using a server cluster with large numbers of file shares**

Every cluster resource you create has a performance impact; the Cluster service must periodically poll all resources to determine whether they have failed. Because of this, you may notice that your cluster's performance decreases as you increase the number of resources that it manages.

If you want to use a cluster to create a high availability file server that will have hundreds of file shares, consider managing those file shares using a single File Share resource. You can do this if the folders that you want to share on the network are all subfolders of a common folder. For example, if you use your file server to support hundreds of users, and each user has a private folder that is a subfolder of the Users folder, you can manage individual file shares for each user using a single File Share resource. You also have the option to hide the file shares that you create for each user so that they will not be visible when browsing. For information on sharing the subfolders of a folder and hiding the files shares you create, see Share subfolders on a File Share resource.

The maximum number of resources supported by the default Resource Monitor is 1674 (the default Resource Monitor creates a maximum of 62 threads, each thread monitoring up to 27 resources). Additional Resource Monitors can each monitor one resource.

For more information on how the Cluster service checks for failed resources, see Resource failure.

# Using Shadow Copies of Shared Folders in a server cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Shadow Copies of Shared Folders in a server cluster**
Shadow Copies of Shared Folders provides point-in-time copies of files that are located on shared resources, such as a file server. You can use Shadow Copies of Shared Folders as part of your overall disaster recovery strategy to quickly recover files that were accidentally deleted or overwritten, as well as to compare different versions of files. However, there are several important considerations to keep in mind if you want to use Shadow Copies of Shared Folders on a cluster disk:

- Cluster-managed shadow copies can only be created in a single quorum device cluster on a disk with a Physical Disk resource. In a single node cluster or majority node set cluster without a shared cluster disk, shadow copies can only be created and managed locally.
- You cannot enable Shadow Copies of Shared Folders for the quorum resource, although you can enable Shadow Copies of Shared Folders for a File Share resource.
- The recurring scheduled task that generates volume shadow copies must run on the same node that currently owns the storage volume.
- The cluster resource that manages the scheduled task must be able to fail over with the Physical Disk resource that manages the storage volume.

The Volume Shadow Copy Service Task resource type can be used to manage shadow copies in a server cluster. The Volume Shadow Copy Service Task resource used in this way has a required dependency on the Physical Disk resource.

For more information on shadow copies, see Shadow Copies for Shared Folders.

For more information on the Volume Shadow Copy Service Task resource type, see Volume Shadow Copy Service resource type.

**Best practices for using shadow copies in a cluster**

- To improve server performance, use different disks for the shadow copy volume and the storage volume.

  **Note**

  - You cannot create storage volume associations between distinct volumes on a disk managed by a single Physical Disk resource.

- If you have multiple shadow copy volumes, use a single, separate disk for the storage volume.
- Do not enable shadow copies in a deployed cluster. When you enable shadow copies, the shadow copy volumes (as well as all resources that directly or indirectly depend on the disk) go offline for a brief period while the resource dependencies are created. The shadow copy volumes are not accessible to applications and users while they are offline.
- Enable shadow copies after you have created your cluster. If you enable shadow copies before clustering your servers, you must disable, then reenable, shadow copies.

  **Note**

  - If you cluster a disk containing a previously-created shadow copy, the shadow copy may stop functioning once that disk's Physical Disk resource fails over.

**Using Cluster Administrator or cluster.exe to manage shadow copies**
**Important**

- Using Cluster Administrator or cluster.exe to completely configure and manage shadow copies in a cluster is not a supported scenario. The only supported scenario for configuring shadow copies in a cluster is described in Enable Shadow Copies for Shared Folders in a cluster.

For more information on using Cluster Administrator to create a Volume Shadow Copy Service Task resource, see Checklist: Installing a Volume Shadow Copy Service Task resource.

For more information on using cluster.exe to create a Volume Shadow Copy Service Task resource, see Cluster resource.

**Enabling shadow copies on a cluster node**

When you enable shadow copies on a cluster node (for example, by using the Configure Shadow Copy user interface through the Computer Management MMC ), a Volume Shadow Copy Service Task resource and a scheduled task for creating that shadow copy are automatically generated and configured. You do not need to explicity create the Volume Shadow Copy Service Task resource using Cluster Administrator or cluster.exe. In addition, the Configure Shadow Copy user interface automatically sets up the required resource dependencies. The default properties of the scheduled task and Volume Shadow Copy Service Task resource are described in the table below. For more information, see Enable Shadow Copies for Shared Folders in a cluster.

| **Scheduled task property** | **Volume Shadow Copy Service Task resource Cluster Administrator property (cluster.exe property)** | **Default setting** |
|---|---|---|
| **Name of task** | **Name of resource (TaskName)** | ShadowCopyVolume{*VolumeGUID*}<br><br>**Note**<br><br>• The name of the scheduled task is the same as the name of the Volume Shadow Copy Service Task resource associated with that task. |
| **Run** | **Command to run/Command parameters (ApplicationName/ApplicationParams)** | %systemroot%\system32\vssadmin.exe Create Shadow /AutoRetry=5 /For=\\?\Volume{VolumeGUID}\<br><br>For more information, see Vssadmin. |
| **Creator** | n/a | Cluster service |
| **Start in** | **Start in** | %systemroot%\system32\ |
| **Run as** | n/a | Local System |
| **Schedule** | **Schedule (TriggerArray)** | The default settings used by Shadow Copies of Shared Folders.<br><br>For more information, see Defining a schedule for shadow copies. |

# Setting Properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting properties**

After setting up your cluster, you might need to change the configuration of some resources and groups from time to time. To make changes to a selected group or resource, on the **File** menu, click **Properties**.

For detailed procedures on setting properties for groups, resources, networks, and network interfaces, see Server Cluster How To....

This section covers:

- Setting group properties
- Setting resource properties
- Setting network properties
- Setting network interface properties
- Setting node properties
- Setting resource type properties

# Setting group properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting group properties**

Using the group **Properties** dialog box, you can:

- View the state of the group (whether it is online or offline) and view which node currently owns the group.
- Change the group name, description, and preferred owner.
- Set the failover policy (using the **Failover** tab).
- Set the failback policy (using the **Failback** tab).

Do not update the properties of a cluster object on multiple nodes simultaneously. For more information on multiple and simultaneous property updates, see the Microsoft Platform Software Development Kit (SDK).

**Setting failover policy**

Use the **Failover** tab to set failover policy for the group. You can set the following:

- The failover threshold
- The failover period
- Whether or not a failure of one resource affects other resources in the group.

  By default, if one or more resources fail in a group, the Cluster service fails all other resources in the group.

The failover threshold is the number of times the group can fail over within the number of hours specified by the failover period. For example, if a group failover threshold is set to 5 and its failover period to 3, the Cluster service will fail over the group at most five times within a three-hour period. The sixth time a resource in the group fails, the Cluster service fails all other resources in the group and leaves the entire group offline instead of failing over the group. If the **Effect the Group** failover option is disabled, and a resource in the group fails for a sixth time, the Cluster service leaves the resource in the offline state instead of failing all other resources in the group. For example, if the **Effect the Group** option was disabled and a Network Name resource failed six times in three hours, on the sixth failure the Network Name would be left offline. An IP Address resource, however, would be left online.

**Setting failback policy**

By default, groups are set not to fail back. Unless you manually configure your group to fail back after failing over, it continues to run on the alternate node after the failed node comes back online.

When you configure a group to automatically fail back to the preferred node, you specify whether you want the group to fail back as soon as the preferred node is available or to fail back only during specific hours that you define. This option is useful if you want the failback to occur after peak business hours or if you want to make sure the preferred node is able to support the group when it does come back online.

The group must be configured to have a preferred owner to fail back. You can specify a preferred owner on the **General** tab of the group **Properties** dialog box.

For instructions on how to open the **Properties** dialog box, see Change the properties of a group.

# Setting resource properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting resource properties**

Using the resource **Properties** dialog box, you can:

- View or change the resource name.
- View or change the resource description and possible owners.
- Specify whether the resource runs in a separate memory space.
- View the resource type, group ownership, resource state, and which node currently owns the resource.
- View preexisting dependencies and modify resource dependencies.
- Specify whether to restart a resource and, if so, what settings to use to restart.
- Specify how often you want the Cluster service to check whether the resource is in the online state (by setting the **Looks Alive** and **Is Alive** polling intervals).
- Specify the amount of time that a resource in a pending state (Online Pending or Offline Pending) has to either resolve its status or report that it is making progress to the Resource Monitor before the Cluster service puts the resource in Failed status.
- Set specific resource parameters.

The **General**, **Dependencies**, and **Advanced** tabs are the same for every resource. Some resource types support additional tabs.

Do not update the properties of a cluster object on multiple nodes simultaneously. For more information on multiple and simultaneous property updates, see the Microsoft Platform Software Development Kit (SDK).

**Setting resource dependencies**

The Cluster service uses the dependencies list when bringing resources online and offline. For example, if a group is brought online in which a physical disk and a file share are located, it is futile to try to bring the file share online before bringing the physical disk (on which the file share is located) online. It is important to remember this when you configure resources because the groups function properly only if these dependencies are correctly configured.

The following table shows resources and their dependencies. The resources in the right column must be configured before you create the resource.

| Resource | Dependencies (required or recommended) |
|---|---|
| DHCP Service | Physical Disk or other storage class device (on which the files are located)<br><br>Internet Protocol (IP) Address (for client access to the DHCP server)<br><br>Network Name (for client access to the DHCP server) |
| Distributed Transaction Coordinator | Physical Disk or other storage class device |

| | |
|---|---|
| File Share | A file share that is configured as a DFS root has the following required dependency:<br>• Network Name<br><br>A file share that is not DFS has no required dependencies. It is recommended that file share resources have the following dependencies:<br>• Physical Disk<br>• Network Name |
| Generic Application | Network Name |
| Generic Script | None |
| Generic Service | Network Name |
| IP Address | None |
| Local Quorum | None |
| Majority Node Set | None |
| Message Queuing | Physical Disk or other storage class device<br><br>Network Name (so that remote clients can access it) |
| Network Name | IP Address (that you want to correspond to the name) |
| Physical Disk | None |
| Print Spooler | Physical Disk or other storage class device (on which the spool is located)<br><br>Network Name (so that remote clients can access it) |
| Volume Shadow Copy Service Task | A Volume Shadow Copy Service Task resource has no required dependencies. However, Volume Shadow Copy Service Task resources that manage scheduled tasks for shadow copies have a required dependency on the Physical Disk resources for the volumes on which the volume shadow copies are located. |
| WINS Service | Physical Disk or other storage class device (on which the files are located)<br><br>Network Name (for client access to the WINS server)<br><br>**Note**<br>• The Network Name will, in turn, be dependent on the IP Address which is necessary for client access to the WINS server. |

The New Resource Wizard enforces the use of resource dependencies.

**Setting advanced resource properties**

Use the **Advanced** tab of the resource **Properties** dialog box to:

- Specify whether the Cluster service attempts to restart a resource or allows the resource to fail.

- If you want the Cluster service to restart the resource, you must specify whether to affect the group. If you want the group to be failed over when this resource fails, select the **Affect the group** check box, and then enter the values you want in **Threshold** and **Period**. These values determine how many times the Cluster service attempts to restart the resource before failover and the amount of time in which the threshold number of restart attempts takes place. If the **Affect the group** check box is not selected, a failure of this resource never causes the group to fail over.

- Specify how often the Cluster service makes a cursory check (Looks Alive) or in-depth check (Is Alive) to determine if a resource is in the online state.

- Click **Use resource type value** to use the default number for the resource type.

- Specify the amount of time that a resource in a pending state (Online Pending or Offline Pending) has to either resolve its status or report that it is making progress to the Resource Monitor before the Cluster service puts the resource in Failed status.

**Setting resource parameters**

Most resources have a resource-specific **Parameters** tab in their **Properties** dialog box. The following table lists each resource and its configurable parameters.

| Resource | Configurable parameters |
|---|---|
| DHCP Service | You can enter the path to the DHCP database files and to the backup of the DHCP database files.<br><br>You can enter the path to the audit log file stored by the DHCP service. |
| Distributed Transaction Coordinator | None. |
| File Share | You can use the **Parameters** tab for the **File Share Properties** dialog box to set permissions on the file share and limit the number of simultaneous users.<br><br>**Important**<br><br>• When you create a new File Share resource, permissions for the Everyone group for that file share are automatically set by default to Read only. You can use Cluster Administrator or cluster.exe to change the default permissions. For more information, see View or set resource properties.<br><br>You can also use the **Parameters** tab to set the name of the file share (because clients will detect it in their browse or explore lists), set the file share comment, and set the path to the shared files. |
| Generic Application | You can use the **Parameters** tab in the **Generic Application Properties** dialog box to specify the application command line, current directory, whether it uses the Network Name for the computer name, and whether or not it can interact with the desktop.<br><br>**Important**<br><br>• It is recommended that you carefully consider the ways in which you interact with the server before you enable the Generic Application to interact with the desktop. This option is not valid in a number of scenarios, for example, administering a cluster either in a terminal server session or remotely using Emergency Management Services and an out-of-band connection. |
| Generic Script | You can use the **Parameters** tab for the **Generic Script Properties** dialog box to specify the path to the Visual Basic (VB) script file. |
| Generic Service | You can use the **Parameters** tab for the **Generic Service Properties** dialog box to specify the service name, startup parameters, and whether it uses the Network Name for the computer name. |

| | |
|---|---|
| IP Address | You can configure the IP address and subnet mask. You must also specify which cluster network to use and enable NetBIOS for the IP address. |
| Local Quorum | None. |
| Majority Node Set | None. |
| Message Queuing | None. |
| Network Name | You can configure the computer name parameter, enable Kerberos authentication, and indicate whether DNS registration failures are tolerated for the Network Name resource. |
| Physical Disk | You can specify only which drive to use for the Physical Disk resource. You cannot change this after the resource is created. |
| Print Spooler | You can configure the path to the Print Spooler folder and the job completion time-out for the Print Spooler resource. |
| Volume Shadow Copy Service Task | You can use the **Parameters** tab for the **Volume Shadow Copy Service Task Properties** dialog box to specify the command to run for the task (a combination of the **Command to run** and **Command parameters** fields), the working directory, and the schedule for running that command.<br>**Important**<br>• These parameters are automatically set when you enable shadow copies on cluster nodes. For more information, see Enable Shadow Copies for Shared Folders in a cluster. |
| WINS Service | You can enter the path to the WINS database files and to the backup of the WINS database files. |

# Setting network properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting network properties**

After setting up your cluster, you might need to change the configuration of the networks that support client-to-cluster communication and node-to-node communication. For detailed procedures on setting network properties, see Manage Networks.

Using the network **Properties** dialog box, you can:

- View or change the network name.
- View or change the network description.
- Enable or disable the network for cluster use. If you enable the network for cluster use, you can specify whether to use the network for only client-to-cluster communication, only node-to-node communication, or both.
- View the network's state.
- View the network ID.
- View the subnet mask.

Do not update the properties of a cluster object on multiple nodes simultaneously. For more information on multiple and simultaneous property updates, see the Microsoft Platform Software Development Kit (SDK).

**Setting network priorities**

Cluster nodes can communicate over multiple networks. You can change the network's priority to specify the order in which the nodes will attempt to communicate over the networks. For detailed instructions on changing the network priority, see Change network priority for communication between nodes.

# Setting network interface properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting network interface properties**

After setting up your cluster, you might need to change the configuration of the network interfaces. For detailed procedures on setting network interface properties, see Manage Networks.

Using the network interface **Properties** dialog box, you can:

- View the node and network object name.
- View or change the network interface description.
- View the network interface's state.
- View the name of the network adapter for the network interface.
- View the network interface's Internet Protocol (IP) address.

Do not update the properties of a cluster object on multiple nodes simultaneously. For more information on multiple and simultaneous property updates, see the Microsoft Platform Software Development Kit (SDK).

# Setting node properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting node properties**

Using the node **Properties** dialog box, you can:

- View the node object name.
- View or change the node description.
- View the node's state. A node's state can be:
  - **Down.** The node is not actively participating in cluster operations.

    **Joining.** The node is in the process of becoming an active participant in the cluster operations.

    **Paused.** The node is actively participating in cluster operations but cannot take ownership of resource groups and cannot bring resources online.

    **Unknown.** The state could not be determined.

    **Up.** The node is actively participating in all cluster operations.
- View the version and build number of the Cluster service.

Do not update the properties of a cluster object on multiple nodes simultaneously. For more information on multiple and simultaneous property updates, see the Microsoft Platform Software Development Kit (SDK).

# Setting resource type properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting resource type properties**

Using the resource type **Properties** dialog box, you can:

- View or change the resource type display name. The display name is used to manage the resource type in the cluster administration tools.
- View or change the description of the resource type.
- View the nodes on which the resource type can run.
- View or change the default poll intervals (in milliseconds) for resources of this type. The poll intervals are Looks Alive and Is Alive.

  The Looks Alive poll interval specifies how often the Cluster service checks the state of the resource to determine whether or not the resource appears to be active. You can use the default value for the resource type or specify a number in milliseconds.

  The Is Alive poll interval specifies how often the Cluster service checks whether the resource is active. You can use the default value for the resource type or specify a number in milliseconds.
- View the type name of the resource type. The type name is used internally by APIs and resource dynamic-link libraries (DLLs).
- View the DLL file that implements the behavior of this resource type.
- View whether resources of this type can be used as a quorum resource. The only quorum-capable resource types currently provided by Microsoft are the Physical Disk, Majority Node Set, and Local Quorum resource types. For more information on quorum resources, see Quorum resource. For more information on the Physical Disk, Majority Node Set, and Local Quorum resource types, see Physical Disk resource type, Majority Node Set resource type, and Local Quorum resource type.

Do not update the properties of a cluster object on multiple nodes simultaneously. For more information on multiple and simultaneous property updates, see the Microsoft Platform Software Development Kit (SDK).

# Network Management

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network management**

This topic describes features of server clusters that affect network management:

- Plug and Play network support
- Name coordination between Cluster service and Network and Dial-up Connections
- Network Failure Detection and Recovery

# Plug and Play network support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Plug and Play network support**

Using the Plug and Play technology built into Windows Server 2003 family operating systems, the Cluster service detects the addition and removal of TCP/IP networking components. For example, the Cluster service can detect the following events:

- Addition or removal of network adapters
- Binding or unbinding the TCP/IP protocol stack to a network adapter
- Addition or release of a DHCP-assigned or static Internet protocol (IP) address

When a new network adapter is added to the computer and bound to TCP/IP using a subnet not already in use by the Cluster service, the Cluster service adds a network interface object to the cluster. If the network interface is connected to a new network, the Cluster service also adds a new network object to the cluster. You do not need to restart the Cluster service for these changes to take place. By default, the newly created interface and network objects are used for mixed communication, that is for both client and internal cluster communication.

Similarly, if a TCP/IP stack is removed, the Cluster service removes the corresponding network interface object from the cluster. If that network interface is the last interface connected to a network, the Cluster service removes the network object from the cluster.

If one network adapter is removed from a node, all IP Address resources dependent on that adapter will attempt to failover to another node. If all network adapters in the cluster are removed (that is, reconfigured or connected to a different physical network), all dependent IP Address resources go offline and are "orphaned," that is, their associated network is empty. If another network adapter is then added and assigned to the same subnet as the orphaned IP Address resource, the resource is reassociated with the "new" network; it is new in the sense that it has a different GUID and is a new instance of a network that has the same properties as the "old" network.

For more information on configuring network objects, see Setting network properties. For more information on configuring network interface objects, see Setting network interface properties.

# Name coordination between Cluster service and Network and Dial-up Connections

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Name coordination between Cluster service and Network Connections**
In the Windows Server 2003 family, each physical network adapter in a computer is represented by a named icon in the Network Connections folder. The Cluster service coordinates the names of its network interface objects with the names of the icons in the Network Connections folder.

**Automatic renaming of networks**
Windows Server 2003 family operating systems automatically ensure that both a server cluster and the Network Connections folder use the same name to reference any individual network interface. For example, if you use any cluster administration tool to rename a cluster network object, the Windows Server 2003 family operating system automatically renames the corresponding icon in the Network Connections folder. Similarly, if you rename a network connection icon in the Network Connections folder, the Windows Server 2003 family operating system renames the matching cluster network object, if one exists.

The Windows Server 2003 family operating system ensures that cluster nodes also use the same name for the network object and corresponding icons in the Network Connections folder. Therefore, if you rename a cluster network object on one node of a cluster, the name of the network object on the other nodes and the icon name in the Network Connections folders on all nodes also change. The operating system does not notify you when it renames a related cluster network object or icon in a Network Connections folder.

The following table summarizes the actions the Windows Server 2003 family operating system takes in response to any renaming you perform.

| Action | Result |
|---|---|
| Set up a new cluster. | The Windows Server 2003 family operating system uses the names of the icons in the Network Connections folder as the default names for cluster network objects. If you choose to rename any of the cluster network objects, the operating system renames the corresponding icon in the Network Connections folder. |
| Add a new node to an existing cluster. | The Windows Server 2003 family operating system renames icons in the Network Connections folder to match the names of the cluster network objects. |
| Rename a network object using a cluster administration tool. | The corresponding icon in the Network Connections folder is renamed on all cluster nodes. |
| Rename an icon in the Network Connections folder. | The Windows Server 2003 family operating system renames the corresponding cluster network object. It renames the corresponding icon in the Network Connections folder on the other node. |

**Network name conflicts**
When you change the name of a network object or an icon in the Network Connections folder on one node of a cluster, you introduce the possibility of a naming conflict. For example, assume you have a cluster consisting of Node 1 and Node 2, and you change the name of a cluster network object on Node 1 from *Public Network* to *Corporate Network*. If an icon in the Network Connections folder on either Node 1 or Node 2 is already named Corporate Network, then the renaming attempt creates a name conflict.

Windows Server 2003 family operating systems prevent name conflicts by appending a number to any existing icon or network object names that would cause a name conflict with whatever new name you enter. In the previous example, the operating system renames the existing **Corporate Network** icon to **Corporate Network (2)**. This feature ensures that whenever you attempt to rename a network object or an icon in the Network Connections folder, your renaming attempt succeeds without generating any name conflicts.

# Network Failure Detection and Recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network failure detection and recovery**
The Cluster service is able to detect and isolate network failures. The Cluster service uses its improved network failure detection to improve failure recovery actions.

This section describes the Cluster service network state management and failure recovery process:

- Network state management and failure detection
- Network-failure recovery

# Network state management and failure detection

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network state management and failure detection**

Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition use sophisticated algorithms to monitor the health of network interfaces in a server cluster. The following tables summarize the states that a network interface and a network can be in.

**Network interface states**

| State | Description |
|---|---|
| Failed | The node that owns the network interface is active; however, it cannot communicate through its network interface. The Cluster service has determined that the error is isolated to the network interface.<br><br>Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition use several methods to determine whether failures are isolated to network interfaces. First, the Cluster service uses a heartbeat to help detect network problems. Second, the operating system can receive hardware failure notifications from a network adapter that supports the Network Driver Interface Specification (NDIS). Using network adapters that support newer NDIS features, for example, media sense, improves cluster network state management. Finally, the Cluster service attempts to ping external hosts to estimate the scope of a network interface failure. This can cause a temporary increase in network traffic while the operating system investigates the extent of a network failure.<br><br>**Note**<br><br>• When a failed network adapter is replaced, the Cluster service will determine that the new card is the interface for an existing network if the card is on the same subnet. |
| Unavailable | The network is disabled for cluster use, or the node associated with this network interface is down. |
| Unreachable | The node cannot communicate through this network interface. The reason for the failure is unknown. |
| Up | The network interface is active and can communicate with all other interfaces on the network (except those currently in Failed or Unavailable states). Up is the normal operational state. |
| Unknown | The state cannot be determined. |

**Network states**

| State | Description |
|---|---|
| Down | The network has failed. None of the active nodes in the cluster can communicate with one another using this network. |
| Unavailable | The network is disabled for cluster use, or all of the cluster nodes attached to this network are inactive. |
| Partitioned | The network has partially failed. Some active cluster nodes cannot communicate with one another over this network. |
| Up | The network is functioning normally. Up is the normal operational state. |
| Unknown | The state cannot be determined. |

# Network-failure recovery

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Network-failure recovery**

When a network interface enters the Failed state, the Cluster service triggers the failure of all IP Address resources that use that network interface. When an IP Address resource fails, the Cluster service treats it like the failure of any other resource: All resources that depend on the IP Address resource fail, and the Cluster service uses the failover policy that you have set to determine whether or not to fail over the resource group.

The Cluster service does not trigger the failure of IP Address resources when a network interface enters the Unreachable state. If the network interface is in the Unreachable state, the Cluster service has not isolated the network problem enough to implement an appropriate recovery policy. When a network interface enters the Unreachable state, investigate the state of the network to determine whether or not to cause a manual failover of resource groups.

# Administering Server Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering server clusters**

You can use cluster commands to administer a server cluster and the Cluster service from the command prompt.

- Managing a Server Cluster from the Command Line
- Managing the Cluster Service from the Command Line

# Managing a Server Cluster from the Command Line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing a server cluster from the command line**
You can use cluster commands to create, configure, and administer server clusters from the command prompt. You can also call the program Cluster.exe from command scripts to automate many cluster administration tasks. Cluster.exe is provided on all computers running Windows Server 2003 family operating systems.

You can use Cluster.exe to create, configure, and administer clusters from a node of the cluster, from nodes of other clusters, or from other computers running Windows NT 4.0 with Service Pack 3 or later.

**Caution**

- When using cluster.exe to manage a server cluster, the locale settings for the user account must match the system default locale on the computer used to manage the cluster. If you remotely manage a server cluster from another computer, the system default locale on that remote computer can be different from the system default locale on all the nodes in the cluster only if the names of the cluster and all the cluster nodes use English language characters. For more information on changing the system default locale, see Change number, currency, time, and date settings.

**Important**

- When you remotely manage a cluster using Cluster Administrator or cluster.exe, you must verify that NetBIOS over TCP/IP (NetBT) is enabled on the client. For more information on enabling NetBIOS over TCP/IP, see Configure TCP/IP settings.

You can also start the Cluster Administration tool at a command prompt using the Cluadmin.exe command.

For information about other command-line utilities, see Command-line reference. For more information about manageability, see Management Strategies and Tools.

The following rules apply to cluster.exe and its subcommands when entering commands from the command prompt:

- You must use quotation marks around all names that contain spaces and special characters. Examples of special characters are: the minus sign (-), the slash mark (/), quotation marks ("), the colon (:), the comma (,), and the equal sign (=). For other examples, see the cluster commands in Related Topics.

- When setting properties to True or False, specify 1 (one) for True and 0 (zero) for False.

- You specify the local cluster when you enter a PERIOD (.) in the command for cluster name.

- You can specify multiple options for one command line. Cluster.exe processes options from left to right. If an option fails, the command will stop executing at the failed option.

- Use double quotation marks (") in a string that has two consecutive double quotation marks (""). The following example shows double quotation marks in a string:

  **cluster resource AResource /priv aProp="a """**

This section describes cluadmin.exe as well as the cluster.exe commands, command abbreviations, and property lists:

- Cluadmin

- Cluster

- Cluster node

- Cluster group
- Cluster network
- Cluster netinterface
- Cluster resource
- Cluster resourcetype
- Cluster command abbreviations
- Property lists

# Cluadmin

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluadmin**
The cluadmin command is used to connect to a server cluster. Used without parameters, **cluadmin** starts Cluster Administrator, the tool used to configure and manage server clusters.

**Syntax**
**cluadmin** [**/norecon**[**nect**]] [*IP Address*] [**.**] [{*ClusterName* | *NodeName*}]

**Parameters**
**/norecon[nect]**
Starts Cluster Administrator without restoring all previous cluster connections left open when you closed the previous session.

***IP Address***
Connects to either a node in the cluster or the cluster using the IP address.

**[ .]**
Connects to the local node.

**{*ClusterName* | *NodeName*}**
Connects to the cluster using a NetBIOS name. You can use either one of the node names or the cluster name.

**/?**
Displays help at the command prompt.

**Remarks**

- Using the **/norecon**[**nect**] option

  If you use this parameter, you will be prompted to enter the name of the cluster that you want to connect to.

- Using the [**.**] option

  If the cluster name or cluster Internet Protocol (IP) address does not come online, you can use this parameter to make a connection to the cluster.

- Using the {*ClusterName* | *NodeName*} option

  To use this parameter, the name resolution service must be functioning correctly.

**Examples**
To connect to cluster node NODE1, type:

**cluadmin NODE1**

To start Cluster Administrator without opening old cluster connections, type:

**cluadmin /norecon**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |

| | |
|---|---|
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Command-line reference A-Z

Command shell overview

**Other Resources**

Cluster commands overview

# Cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster**
The cluster command is used to create a new cluster or administer an existing cluster.

**Syntax**
**cluster** [[**/cluster:**]*ClusterName*] **/***Option*

**Parameters**
***ClusterName* /create /ipaddr[ess]:*Static_IP_Address*[,*SubnetMask*,*NetworkConnectionName*] /pass[word]:*Password* /user:{*Domain**UserName* | *UserName@Domain*} [/node:*NodeName*] [/verb[ose]] [/unattend[ed]] [/min[imum]]**
Creates a cluster. Descriptions of the optional parameters are provided in the table below.

***ClusterName* /add[nodes][:*NodeName1*[,*NodeName2*, ...]] [/pass[word]:*Password*] [/verb[ose]] [/unattend[ed]] [/min[imum]]**
Adds one or more nodes to an existing cluster.

**/ren[ame]:*ClusterName***
Renames a cluster.

**/cluster: *ClusterName1*[,*ClusterName2*, ...] /changepass[word][:*NewPassword*[,*OldPassword*]] [/skipdc] [/force] [/test] [/quiet] [/verb[ose]] [/unattend[ed]] [/help]**
For one or more clusters, changes the Cluster service account password on the domain controller and all the cluster nodes from *OldPassword* to *NewPassword*. If not supplied as part of the command, you will be prompted to provide the *NewPassword* or the *OldPassword*. You can specify a blank password by supplying two double quotes ("""). If multiple clusters are specified in one operation, they must all use the same Cluster service account. Descriptions of the optional parameters are provided in the table below.

**/ver[sion]**
Displays the Cluster service version number.

**/listnetpri[ority]**
Lists the priority of all internal (private and mixed) networks.

**/setnetpri[ority]:*NetworkName1*[,*NetworkName2*, ...]**
Sets the priority order for all internal networks. If you use this option, you must supply all the internal networks in the cluster.

**/quorum[resource][:*ResourceName*] [/path:*Drive*:*Path*\] [/maxlogsize:*MaximumSize_in_Kilobytes*]**
Changes the name or location of the quorum resource or the size of the quorum log. If you change the location of the quorum resource, do not omit the drive letter, the colon, or any backslashes. For example, for a quorum resource on drive f: in a folder called Quorumfolder, you must specify the location with **/path:f:\Quorumfolder** not **/path:f:Quorumfolder**.

If no options are supplied, this command will display the name of the quorum resource, the path, and maximum log size.

**/reg[admin]ext:*AdminExtension_DLL* [,*AdminExtension_DLL* ...]**
Registers a Cluster Administrator extension DLL with the cluster.

**/unreg[admin]ext:*AdminExtension_DLL* [,*AdminExtension_DLL* ...]**
Unregisters a Cluster Administrator extension DLL from the cluster.

**/list:[*DomainName*]**
Displays a list of clusters in the computer's domain or a specified domain. Do not use the cluster name with this command-line option.

**/prop[erties] [*PropertyList*]**
Views or sets cluster common properties.

**/priv[properties] [*PropertyList*]**
Views or sets cluster private properties.

**/registerdns**
Registers existing Network Name resources with a DNS server in a way that does not interrupt cluster availability.

**/wiz[ard]**
Starts the Cluster Configuration Wizard. This command must be used with either the **/create** or **/add**[**nodes**] commands.

**Important**

- You cannot supply the subnet mask or network connection name when you use the **/wiz**[**ard**] command to create a cluster. The program will use the information from the network connection with the same subnet mask as the cluster to supply the subnet mask, and network connection name.

**/setfail[urereactions][:*NodeName*[,*NodeName* ...]]**
Resets the restart actions settings that the Services Microsoft Management Console snap-in uses for the Cluster service to the original values used at setup.

**/?**
Displays help at the command prompt.

**Remarks**

- Using the **cluster** command

  The cluster name is optional except when using the **/changepass**[**word**] command. If you do not specify a cluster, Cluster.exe attempts to connect to the cluster running on the node that is running Cluster.exe.

  If the name of your cluster is also a cluster command or its abbreviation, such as "cluster" or "c", use **/cluster:** to explicitly specify the cluster name. For a list of all the cluster commands, see Related Topics.

- Using command-line options with **cluster /create**

  Unless you are invoking the Cluster Configuration Wizard (by using the **/wiz**[**ard**] option), you must supply a name and static IP address for the cluster and the user name for the Cluster service account. If you do not use the **/unattend**[**ed**] option, you must also supply the password for the user account used to run the Cluster service on the first node. If possible, the program will use the information from the network connection with the same subnet mask as the node to supply the subnet mask, and network connection name. The cluster will be created in the same domain to which the first server belongs. See the table below for details on the other optional parameters.

- Using the **cluster** command for unattended setup.

  You can create and configure a cluster during an unattended setup by including the cluster.exe command in the [GuiRunOnce] section of the Unattend.txt file. The cluster.exe command must include all the command line options required for creating and configuring a cluster. For more information on unattended installs, see the Microsoft Windows Corporate Deployment Tools User's Guide (deploy.chm) in the support folder on your installation media.

- Using the **/prop**[**erties**] [*PropertyList*] command-line option

  For more information on property lists, see Related Topics.

- Using the **/priv**[**properties**] [*PropertyList*] command-line option

  For more information on property lists, see Related Topics.

- Using the **/setfail**[**urereactions**][**:***NodeName*[**,***NodeName* ...]] command-line option

  This option resets the Cluster service restart settings for one or more nodes to the following values:

| **Service Restart Option** | **Setting** |
|---|---|
| First failure | Restart the service |
| Second failure | Restart the service |

| | |
|---|---|
| Subsequent failure | Restart the service |
| Reset fail count after *N* days | 0 days |
| Restart service after *N* minutes | 1 minute |

- Command options for creating a cluster

| Option | Description | Default |
|---|---|---|
| [**/node:***NodeName*] | The name of the server to be configured as the first cluster node. | The command will use the name of the current computer if you do not supply a node name. |
| [**/min**[**imum**]] | This option creates a cluster using the Advanced (minimum) Configuration option. Use this option when creating clusters with complex storage configurations. For more information, see Storage configuration options. | n/a |
| [**/verb**[**ose**]] | This option writes all event messages to the cluster log. Use this option for troubleshooting cluster configuration issues. | n/a |
| **/unattend**[**ed**] | If you use this option, you will not be prompted for passwords not supplied at the command prompt. | n/a |

- Command options for changing the Cluster service account password

| Option | Description | Default |
|---|---|---|
| [**/skipdc**] | This option will change the Cluster service account password only on the cluster nodes. It will not change the password on the domain controller. Use this option to complete a previous password update that was only partially completed. For example, use this option if you previously issued a **cluster /changepass** command for multiple clusters and the Cluster service account password was not updated on all the clusters, for whatever reason.<br><br>**Note** If you use this option, you do not have to supply the *OldPassword*. | n/a/ |
| [**/force**] | This option will force execution of the password change command on the available nodes in a cluster even if some of the cluster nodes are not available (that is, some nodes are in the Down, Joining, or Unknown states). | n/a |
| [**/test**] | This option will open a connection to the specified cluster(s), verify that an appropriate domain controller is available, check the version of the operating system installed on each node, and verify that all the cluster nodes are online and use the same Cluster service account. | n/a |
| [**/quiet**] | This option will result in no output, unless an error occurs. | n/a |
| [**/verb**[**ose**]] | This option outputs detailed information about the cluster nodes (for example, node names and Cluster service account names) to the screen. Use this option for troubleshooting password change issues. | n/a |
| **/unattend**[**ed**] | If you use this option, you will not be prompted for passwords not supplied at the command prompt. | n/a |
| [**/help**] | This option displays help for the **cluster /changepass** command only. | n/a |

- Cluster common property names

The following table describes the common **cluster** property names, their uses, and valid settings:

| Common property names | Use |
|---|---|

| | |
|---|---|
| **AdminExtensions** | Describes one or more class identifiers (CLSIDs) of Cluster Administrator extension for the cluster object. |
| **DefaultNetworkRole** | Specifies the default role to assign to networks. |
| **Description** | Specifies administrative description of the cluster. |
| **Security** | Describes the Windows Server 2003 family-compatible security descriptor. |
| **Security Descriptor** | Describes Windows NT-SP4 and the Windows Server 2003 family-compatible security descriptor. |
| **Groups\AdminExtensions** | Describes one or more class identifiers (CLSIDs) of Cluster Administrator extensions for the groups. |
| **NetworkInterfaces\AdminExtensions** | Describes one or more class identifiers (CLSIDs) of Cluster Administrator extensions for the netinterfaces. |
| **Networks\AdminExtensions** | Describes one or more class identifiers (CLSIDs) of Cluster Administrator extensions for the networks. |
| **Nodes\AdminExtensions** | Describes one or more class identifiers (CLSIDs) of Cluster Administrator extensions for the nodes. |
| **Resources\AdminExtensions** | Describes one or more class identifiers (CLSIDs) of Cluster Administrator extensions for the resources. |
| **ResourceTypes\AdminExtensions** | Describes one or more class identifiers (CLSIDs) of Cluster Administrator extensions for the resource types. |
| **EnableEventLogReplication** | Specifies if Event Log Replication is enabled or disabled cluster wide. (True (1) means enabled; False (0) means disabled.) |
| **QuorumArbitrationTimeMin**<br>**QuorumArbitrationTimeMax** | Determines the minimum and maximum time, in seconds, a node is allowed to spend arbitrating for the quorum resource before giving up. The range of values allowed is 0 to 3600 seconds for **QuorumArbitrationTimeMin** and 1 to 3600 seconds for **QuorumArbitrationTimeMax**.<br><br>For a single quorum device server cluster, the default values are 7 seconds for **QuorumArbitrationTimeMin** and 60 seconds for **QuorumArbitrationTimeMax**.<br><br>For a majority node set server cluster, the default values are 7 seconds for **QuorumArbitrationTimeMin** and 180 seconds for **QuorumArbitrationTimeMax**. |
| **DisableGroupPreferredOwnerRandomization** | If you do not specify a preferred owner list for resource groups, the Cluster service randomizes the list it uses internally to decide which node to fail over the groups to, as described in Determining failover and move policies for groups. That is, by default, this property is set to FALSE (0). If you want to disable this randomization and instead follow the internal ordered list of nodes for failover, set the **DisableGroupPreferredOwnerRandomization** property to be TRUE (or 1). |
| **EnableEventDeltaGeneration** | When a set of events in the Event Log is replicated across cluster nodes, this property generates a delta event in the Event Log to indicate the time difference between nodes. True (1), the default, means enabled; False (0) means disabled. |
| **ClusSvcHeartbeatTimeout** | Used by the Cluster Network driver, ClusNet.sys, to monitor the health of the Cluster service. |

| | |
|---|---|
| **HangRecoveryAction** | Describes the recovery action taken if ClusNet.sys detects a failure in the Cluster service. |
| **EnableResourceDllDeadlockDetection**<br><br>**ResourceDllDeadlockTimeout**<br><br>**ResourceDllDeadlockThreshold**<br><br>**ResourceDllDeadlockPeriod** | Used by the Cluster service to monitor if the entry points to the Resource DLLs are deadlocked (a condition that occurs when two processes are each waiting for the other to complete before proceeding). |

- Using cluster private property names

  To view the private properties for a cluster, use the **/priv**[**properties**] parameter without any options. For example, to see the private properties for the cluster OpsClust, type the following at a command prompt:

  **cluster opsclust /priv**

  The following table describes a **cluster** private property for cryptographic checkpoints. This private property is useful in two scenarios:

  - If there is a third-party resource or application that uses a third-party cryptographic provider that is not supplied by Microsoft.

    For information on troubleshooting this issue, see "A third-party resource fails to come online in a mixed-version cluster or while upgrading a cluster" in Group and resource failure problems.

  - If you have security concerns about the cryptographic checkpoint data that is written to the quorum when you import resource data to a cluster node before bringing the resource online.

  You can use this private property to change the encryption levels for any of the cryptographic providers (supplied by third-party developers or by Microsoft) that are used by the Cluster service.

| **Private property names** | **Use** |
|---|---|

| | |
|---|---|
| **"*CSP*"** | Sets the encryption levels (key lengths) for a cryptographic key that is used to export (encrypt) and import (decrypt) resource data (cluster and cluster application cryptographic checkpoints). The imported and exported resource data is saved to the quorum. The cryptographic key is generated by a cryptographic provider that uses the RC2 block encryption method.<br><br>Usage: **"*CSP*"**=*key_length,effective_key_length*:**MULTISTR**<br><br>"*CSP*" is the name of the cryptographic provider. For example, one of the standard cryptographic providers supplied with the Windows Server 2003 family is "Microsoft Enhanced Cryptographic Provider v1.0." For more information, see CLUSCTL_RESOURCE_ADD_CRYPTO_CHECKPOINT in the *Microsoft Platform Software Development Kit* (SDK).<br><br>*key_length* is the RC2 key length, in bits.<br><br>*effective_key_length* is the RC2 effective key length, in bits<br><br>For example, to change the RC2 key length for a third-party cryptographic provider named "Database CSP v2.0" in cluster opsclust to 128 bits, type:<br><br>**cluster opsclust /priv "Database CSP v2.0"=128,40:MULTISTR**<br><br>**Notes**<br><br>For this property to take effect, the resource DLL must add the cryptographic checkpoint for the property. This is because the cryptographic key is exported and saved to the quorum when the checkpoint is added for the very first time; prior to the resource coming online the key is imported from the file saved in the quorum. Review the documentation for your cryptographic provider to obtain the correct procedure for adding the cryptographic checkpoint.<br><br>Review the documentation for your cryptographic provider to obtain valid values for the following RC2 encryption algorithm parameters: *key_length* and *effective_key_length*. |

**Examples**
To rename a cluster called OpsClust to OpsClust1, type:

**cluster opsclust /rename:opsclust1**

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**See Also**
**Concepts**
Command-line reference A-Z
Command shell overview
Property lists
**Other Resources**
Cluster commands overview

# Cluster node

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster node**

The cluster node command is used to administer a node in a server cluster or to view its status. Used without parameters, **cluster node** defaults to the **/status** command-line option and displays the status for all nodes is displayed.

**Syntax**

**cluster**[[**/cluster:**]*ClusterName*] **node** *NodeName* **/***Option*

**Parameters**

**/stat[us]**
Displays the cluster node status (for example, Up, Down, Paused).

**/force[cleanup] [/wait[:*Timeout_in_Seconds*]]**
Manually restores the configuration of the Cluster service on the specified node to its original state.

**/start [/wait[:*Timeout_in_Seconds*]]**
Starts the Cluster service on the specified node.

**/stop [/wait[:*Timeout_in_Seconds*]]**
Stops the Cluster service on the specified node.

**/pause**
Pauses a node.

**/resume**
Resumes a paused node.

**/evict [/wait[:*Timeout_in_Seconds*]]**
Evicts a node from a cluster.

**/listint[erfaces]**
Lists the node's network interfaces.

**/prop[erties] [*PropertyList*]**
Views or sets node common properties.

**/priv[properties] [*PropertyList*]**
Views or sets node private properties.

**/?**
Displays help at the command prompt.

**Remarks**

- Using the **cluster node** command

  The cluster name is optional. If the name of your cluster is also a cluster command or its abbreviation, such as "node", use **/cluster:** to explicitly specify the cluster name. For a list of all the cluster commands, see Related Topics.

  With **/force**[**cleanup**], **/start**, **/stop**, and **/evict,** the command-line option **/wait**[:*Timeout_in_Seconds*] specifies how long Cluster.exe waits before canceling the command if it does not successfully complete. If a time-out is not specified, Cluster.exe waits indefinitely or until the node state changes.

  The node name (node_name) is optional for all commands except for **/listint**[**erfaces**] or when you use the **/prop**[**erties**] or **/priv**[**properties**] commands to set common or private properties.

- Using the **/stat**[**us**] command-line option

  The node name is optional. The status of all nodes is displayed if you do not provide a node name.

- Using the **/evict** command-line option

  If both the node and the cluster are up, this command evicts the node and removes the cluster configuration from the node. If the cluster is up, but the node is down, this command only evicts the node. The Cluster service will then remove the cluster configuration for the node when it detects that the node no longer belongs to the cluster.

- Using the **/force**[**cleanup**], **/start**, **/stop**, **/pause**, **/resume**, or **/evict** command-line options

  If the node name is not specified, the command will take effect on the local node, that is, the node where the command is run.

- Using the **/prop**[**erties**] [*PropertyList*] command-line option

  The node name is optional. If you do not specify a node name, properties for all nodes are displayed. For more information on property lists, see Related Topics.

- Using the **/priv**[**properties**] [*PropertyList*] command-line option

  The node name is optional. If you do not provide a node name, private properties for all nodes are displayed . For more information on property lists, see Related Topics.

- Using the **/stat**[**us**], **/prop**[**erties**], or **/priv**[**properties**] command-line options

  If the node name is not specified, the command will display results for all nodes in the cluster.

- Using cluster node common property names

  The following table describes the common **cluster node** property names, their uses, and valid settings.

| Common property names | Use |
|---|---|
| **Description** | Describes a node. |
| **EventEnableLogReplication** | Specifies if Event Log Replication is enabled or disabled on a node. (True (1) means enabled; False (0) means disabled.) If the cluster Event Log Replication is disabled, Event Log Replication on a node is disabled, regardless of property setting. |
| **NodeName** | Displays the name of the node. |
| **NodeHighestVersion**<br>**NodeLowestVersion** | Displays the highest and lowest versions of the Cluster service with which a node is compatible. |
| **MajorVersion**<br>**MinorVersion** | Displays the version of the operating system installed on the node (for example, **MajorVersion**=5 and **MinorVersion**=00 indicates a version number of 5.00). |
| **BuildNumber** | Displays the build number of the operating system installed on the node. |
| **CSDVersion** | Displays the version of the latest service pack installed on the node. |

- Using cluster node private property name

  By default, there are no private properties for **cluster node**. Software vendors can add private properties to extend cluster functionality.

**Examples**

To display the status of a node named OpsNode1, type:

**cluster node opsnode1 /status**

**Formatting legend**

| Format | Meaning |
| --- | --- |
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Command-line reference A-Z

Command shell overview

Property lists

**Other Resources**

Cluster commands overview

# Cluster group

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster group**

The cluster group command is used to create a new cluster group or administer an existing group. Used without parameters, **cluster group** defaults to the **/status** command-line option and displays the status for a group.

**Syntax**

**cluster** [[**/cluster:**]*ClusterName*] **group** *GroupName* **/***Option*

**Parameters**

**/node: *NodeName***
Specifies the node that owns the group.

**/stat[us]**
Displays the status of a group (online, offline, or Partially Online).

**/create**
Creates a new group.

**/delete**
Deletes a group.

**/ren[ame]:*NewGroupName***
Renames a group.

**/move[to][:*NodeName*] [/wait[:*Timeout_in_Seconds*]]**
Moves a group to another node.

**/on[line][:*NodeName*] [/wait[:*Timeout_in_Seconds*]]**
Brings a group online.

**/off[line][:*NodeName*] [/wait[:*Timeout_in_Seconds*]]**
Takes a group offline.

**/prop[erties] [*PropertyList*]**
Views or sets the group common properties.

**/priv[properties] [*PropertyList*]**
Views or sets the group private properties.

**/listowners**
Displays a list of preferred owners.

**/setowners: *NodeName*[,*NodeName* ...]**
Specifies a preferred owner.

**/?**
Displays help at the command prompt.

**Remarks**

- Using the **cluster group** command

  The cluster name is optional. If the name of your cluster is also a cluster command or its abbreviation, such as "cluster" or "group," use **/cluster:** to explicitly specify the cluster name. For a list of all the cluster commands, see Related Topics.

  With **/move**[**to**], **/online**, and **/offline**, the command-line option **/wait**[**:***Timeout_in_Seconds*] specifies how long Cluster.exe waits before canceling the command if it does not successfully complete. If you do not use the **/wait** option with the **/move**[**to**] command, Cluster.exe waits 20 seconds. If you do not use the **/wait** option with the **/offline** or **/online** commands, Cluster.exe waits 10 seconds. For the **/move**[**to**], **/online**, or **/offline** commands, if you use **/wait** without specifying a time-out period, Cluster.exe waits indefinitely or until the group state changes.

If a node name is not specified, the Cluster service will automatically select a node following the move and failover logic described in Determining failover and move policies for groups.

- Using the **/stat**[**us**] command-line option

  The group name is optional. If you do not provide a group name, status for all groups is displayed.

- Using the **/prop**[**erties**] [*PropertyList*] command-line option

  The group name is optional. Properties for all groups are displayed if group name is not provided. For more information on property lists, see Related Topics.

- Using the **/priv**[**properties**] [*PropertyList*] command-line option

  The group name is optional. Private properties for all groups are displayed if you do not provide a group name. For more information on property lists, see Related Topics.

- Using cluster group common property names

  The following table describes the common **cluster group** property names, their uses, and valid settings:

| **Common property names** | **Use** |
|---|---|
| **Name** | Name of the cluster group object. |
| **Description** | Describes a group. |
| **PersistentState** | Describes the last known persistent state of a group (1, (true) means online; 0 (false) means offline). |
| **FailoverThreshold** | Specifies the number of times the Cluster service attempts to failover a group before it concludes that the group cannot be brought online anywhere in the cluster. |
| **FailoverPeriod** | Specifies the interval (in hours) over which the Cluster service attempts to fail over a group. |
| **AutoFailbackType** | Set to **ClusterGroupPreventFailback (0)** to prevent failback. Set to **ClusterGroupAllowFailback (1)** to allow failback. |
| **FailbackWindowStart** | Specifies the start time (on a 24-clock) for failback of a group to its preferred node. These values must be between 0 (midnight) and 23 (11:00 p.m.) in local time. For immediate failback, set to -1. |
| **FailbackWindowEnd** | Specifies the end time (on a 24-hour clock) for failback of a group to its preferred node. These values must be between 0 (midnight) and 23 (11:00 p.m.) in local time. For immediate failback, set to -1. |
| **AntiAffinityClassNames** | Specifies a name for a class of groups that should not be hosted on the same cluster node. This property (a string of alphanumeric characters) can be used to configure N+I server clusters. For more information on N+I server clusters, see Cluster deployment and operation options. For more information on using this property, see article Q296799, "How to Configure Windows Clustering Groups for Hot Spare Support," in the Microsoft Knowledge Base. |
| **LoadBalState** | This property is not used at this time by the Cluster service, but is reserved for future use. |

- Using Cluster group private property names

  By default, there are no private properties for **cluster group**. Software vendors can add private properties to extend cluster functionality.

**Examples**

To move a group called Disk Group 1 on the OpsClus1 cluster from OpsNode1 to OpsNode2, type:

**cluster opsclust group "disk group 1" /moveto:opsnode2**

To specify that the Cluster service not place groups "disk group 1" and "disk group 2" together on the same node, supply the same string "SEP1" for groups "disk group 1" and "disk group 2". Similarly, to specify that the Cluster service not place groups "disk group 1" and "file share 1" together on the same node, supply the same string "SEP2" for groups "disk group 1" and "file share 1". That is, type:

**cluster opsclust group "disk group 1" /prop AntiAffinityClassNames="SEP1","SEP2"**

**cluster opsclust group "disk group 2" /prop AntiAffinityClassNames="SEP1"**

**cluster opsclust group "file share 1" /prop AntiAffinityClassNames="SEP2"**

**Formatting legend**

| **Format** | **Meaning** |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Command-line reference A-Z

Command shell overview

Property lists

**Other Resources**

Cluster commands overview

# Cluster network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster network**
The cluster network command is used to view or change the properties of all the networks in the cluster. Used without parameters, **cluster network** defaults to the **/status** command-line option and displays the status of all cluster networks.

**Syntax**
**cluster** [[**/cluster:**]*ClusterName*] **net**[**work**] *NetworkName* **/***Option*

**Parameters**
**/stat[us]**
Displays the status for cluster network.

**/prop[erties] [*PropertyList*]**
Views or sets the network common properties.

**/priv[properties] [*PropertyList*]**
Views or sets the network private properties.

**/ren[ame]:*NewName***
Renames the network.

**/listint[erfaces]**
Lists network interfaces associated with the network.

**/?**
Displays help at the command prompt.

**Remarks**

- Using the **cluster net**[**work**] command

  The cluster name is optional. If the name of your cluster is also a cluster command or its abbreviation, such as "cluster" or "network," use **/cluster:** to explicitly specify the cluster name. For a list of all the cluster commands, see Related Topics.

- Using the **/stat**[**us**] command-line option

  The network name is optional. If you do not provide a network name, the status of all cluster networks is displayed.

- Using the **/prop**[**erties**] [*PropertyList*] command-line option

  The network name is optional. If you do not provide a network name the properties for all networks are displayed. For more information on property lists, see Related Topics.

- Using the **/priv**[**properties**] [*PropertyList*] command-line option

  The network name is optional. Private properties for all networks are displayed if you do not provide a network name. For more information on property lists, see Related Topics.

- Using cluster network common property names

  The following table describes the common **cluster network** property names, their uses, and valid settings.

| Property name | Use |
|---|---|

| | |
|---|---|
| **Name** | The name of the network. The cluster uses the same name as the Network Connections folder. For more information, see Name coordination between Cluster service and Network and Dial-up Connections. |
| **Address** | Provides the address for the entire network or subnet. |
| **AddressMask** | The subnet mask for the network. |
| **Description** | A description of the network. |
| **Role** | The communication role of the network. For more information on network roles, see Server cluster networks. |

- For more information about the private properties for **cluster network**, see Q307962, "Multicast Support Enabled for the Cluster Heartbeat," in the Microsoft Knowledge Base.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Command-line reference A-Z
Command shell overview
Property lists
**Other Resources**
Cluster commands overview

# Cluster netinterface

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster netinterface**
The cluster netinterface command is used to view or change the properties of all the network interface devices in the cluster. Used without parameters, **cluster netinterface** defaults to the **/status** command-line option and displays the status for cluster network interface.

**Syntax**
**cluster** [[**/cluster:**]*ClusterName*] **netint**[**erface**] **/node:***NodeName* **/net**[**work**]**:***NetworkName* **/***Option*

**Parameters**

**/stat[us]**
Displays the status for cluster network interface.

**/prop[erties] [*PropertyList*]**
Views or sets the network interface common properties.

**/priv[properties] [*PropertyList*]**
Views or sets the network interface private properties.

**/?**
Displays help at the command prompt.

**Remarks**

- Using the **cluster netint**[**erface**] command

  The cluster name is optional. If the name of your cluster is also a cluster command or its abbreviation, such as "cluster" or "c," use **/cluster:** to explicitly specify the cluster name. For a list of all the cluster commands, see Related Topics. The commands **/node:***node name***/network:***network name* are both required. You cannot supply one without supplying the other.

- Using the **/stat**[**us**] command-line option

  The network name and node name are optional. If you do not provide the network name and node name, the status of all cluster network interfaces is displayed.

- Using the **/prop**[**erties**] [*PropertyList*] command-line option

  The network name and node name are optional. Properties for all network interfaces are displayed if you do not provide a network name and node name. For more information on property lists, see Related Topics.

- Using the **/priv**[**properties**] [*PropertyList*] command-line option

  The network name and node name are optional. If you do not provide a network name and node name, the private properties for all network interfaces are displayed. For more information on property lists, see Related Topics.

- Using cluster netinterface common properties

  The following table describes the common **cluster netinterface** property names, their uses, and valid settings:

| Name | Use |
|---|---|
| **Adapter** | Uniquely identifies the adapter in the cluster. |
| **AdapterID** | A Globally Unique Identifier (GUID) for the adapter in the cluster. |

| | |
|---|---|
| **Address** | The primary network address used by the network interface. |
| **Description** | Optional description of the network interface. |
| **Name** | Name of the cluster network interface object. |
| **Network** | The name of the network to which the network interface is connected. |
| **Node** | The name of the node on which the network interface is installed. |

- Using cluster netinterface private properties

  By default, there are no private properties for **cluster netinterface**. Software vendors can add private properties to extend cluster functionality.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**
**Concepts**
Command-line reference A-Z
Command shell overview
**Other Resources**
Cluster commands overview

# Cluster resource

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster resource**

The cluster resource command is used to create a new cluster resource or administer an existing resource. Used without parameters, **cluster resource** defaults to the **/status** command-line option and displays the status of a cluster resource.

**Syntax**

**cluster** [[**/cluster:**]*ClusterName*] **res**[**ource**] *ResourceName* **/***Option*

**Parameters**

**/stat[us]**
Displays the status of a cluster resource (online, offline, or failed).

**/create /group: *GroupName* /type: *ResourceType* [/separate]**
Create a new resource in a specified group.

**/delete**
Deletes a resource.

**/ren[ame]:*NewResourceName***
Renames a resource.

**/addowner: *NodeName***
Adds a node name to the list of possible owners.

**/removeowner: *NodeName***
Removes a node name from the list of possible owners.

**/listowners**
Displays a list of possible owners.

**/move[to]:*Group***
Moves the resource to a different group.

**/prop[erties] [*PropertyList*]**
Views or sets the resource common properties.

**/priv[properties] [*PropertyList*]**
Views or sets the resource private properties.

**/fail**
Initiates resource failure.

**/on[line] [/wait[:*Timeout_in_Seconds*]]**
Brings the resource online.

**/off[line] [/wait[:*Timeout_in_Seconds*]]**
Takes the resource offline.

**/maint[enancemode][:{on|off}]**
Views or controls maintenance mode (used when running a disk maintenance tool).

**/listdep[endencies]**
List the dependencies for a resource.

**/adddep[endency]:*Resource***
Adds a dependency for a resource.

**/removedep[endency]:*Resource***
Removes a dependency for a resource.

**/addcheck[points]:*Key*[*Subkey* ...][,*Key*[*Subkey* ...] ...]**
Adds a registry checkpoint for a resource.

**/removecheck[points]:*Key*[*Subkey* ...][,*Key*[*Subkey* ...] ...]**
Removes a registry checkpoint for a resource.

**/check[points]**
Lists registry checkpoints for a resource.

**/addcryptocheck[points]:*Type**Name**Key* ... [,*Type**Name**Key* ...]**
Adds a cryptographic key checkpoint for a resource.

**/removecryptocheck[points]:*Type**Name**Key* ... [,*Type**Name**Key* ...]**
Removes a cryptographic key checkpoint for a resource.

**/cryptocheck[points]**
Lists the cryptographic key checkpoints for a resource.

**/?**
Displays help at the command prompt.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

**Remarks**

- Using the **cluster res**[**ource**] command

  The cluster name is optional. If the name of your cluster is also a cluster command or its abbreviation, such as "cluster" or "resource," use **/cluster:** to explicitly specify the cluster name. For a list of all the cluster commands, see Related Topics.

  With **/on**[**line**] and **/off**[**line**] the command-line option **/wait**[**:***Timeout_in_Seconds*] specifies how long Cluster.exe waits before canceling the command if it does not successfully complete. If you do not use the **/wait** option, Cluster.exe waits 10 seconds. If no time-out is specified, Cluster.exe waits indefinitely or until the resource state changes.

- Using the **/stat**[**us**] command-line option

  The resource name is optional. If you do not provide a resource name, the status of all cluster resources is displayed.

- Using the **/create** command-line option

  When you use cluster.exe to create a new File Share resource, permissions for the Everyone group for that file share are automatically set by default to Full Control. In contrast, if you use Cluster Administrator to create a new File Share resource, permissions for the Everyone group are set by default to Read-only. You can change the default permissions by specifying the Security private property as described below.

- Using the **/separate** command-line option

  Use **/separate** to specify that the resource should run in a separate Resource Monitor.

- Using the **/prop**[**erties**] [*PropertyList*] command-line option

  The resource name is optional. If you do not provide a resource name, properties for all resources are displayed. For more information on property lists, see Related Topics.

- Using the **/priv**[**properties**] [*PropertyList*] command-line option

  Resource name is optional. If you do not provide a resource name, private properties for all resources are displayed. For more information on property lists, see Related Topics.

- Using the **/maint**[**enancemode**][**:**{**on**|**off**}] command-line option

Maintenance mode can be run only on a disk resource and only when the resource is online. Maintenance mode allows a disk maintenance tool, such as Chkdsk, to finish running without being interrupted by failover. Maintenance mode will remain on until you turn it off, until the resource goes offline or fails, or until the node on which the resource is running restarts or loses communication with other nodes (which causes failover of all resources on that node). As soon as the maintenance tool has finished running, turn maintenance mode off. Type **/maint** to check whether maintenance mode is on or off.

- Using the **/check**[**points**] command-line option

  The resource name is optional. Checkpoints for all cluster resources are displayed if you do not provide a resource name.

- Using the **/addcryptocheck**[**points**] or **/removecryptocheck**[**points**] command-line options

  The argument for a cryptographic checkpoint is the path to a cryptographic key container (that is, *Type*****Name******Key*). *Type* is a number, *Name* is the name of the cryptographic provider, and *Key* is the key container name (usually created by an application).

- Using the **/cryptocheck**[**points**] command-line option

  The resource name is optional. If you do not provide a resource name, checkpoints for all cluster resources are displayed. The argument for a cryptographic checkpoint is the path to a cryptographic key container (that is, *Type*****Name******Key*). *Type* is a number, *Name* is the name of the cryptographic provider, and *Key* is the key container name (usually created by an application).

- Using cluster resource common property names

  The following table describes the common **cluster resource** property names, their uses, and valid settings.

| Common property names | Use |
|---|---|
| **Name** | Name of the cluster resource object. |
| **Description** | Describes a resource. |
| **Type** | Specifies the resource type. |
| **DebugPrefix** | Specifies the appropriate debugger for the resource. For more information, see the *Microsoft Platform Software Development Kit*. |
| **SeparateMonitor** | Indicates whether or not a resource shares a Resource Monitor. Valid values are True (1) and False (0). |
| **PersistentState** | Describes the last known persistent state of a resource. In order to save the resource's **PersistentState** property, the resource must be offline. |
| **LooksAlivePollInterval** | Recommended interval in milliseconds that the Cluster service polls a resource to determine if it appears operational. If a value is not specified, a default value is taken from the **LooksAlivePollInterval** property for the resource type. |
| **IsAlivePollInterval** | Interval (in milliseconds) that the Cluster service polls a resource to determine if it is operational. If a value is not specified, a default value is taken from the **IsAlivePollInterval** property for the specific resource type. **IsAlivePollInterval** cannot be zero. |

| | |
|---|---|
| **RestartAction** | Describes the action to perform if the resource fails. The choices are:<br><br>**ClusterResourceDontRestart (0)**: Do not restart following a failure.<br><br>**ClusterResourceRestartNoNotify (1)**: If the resource exceeds its restart threshold within its restart period, the Cluster service does not attempt to failover the group to another node.<br><br>**ClusterResourceRestartNotify (2)**: If the resource exceeds its restart threshold within its restart period, the Cluster service attempts to fail over the group to another node.<br><br>If you do not enter a value, the default allows the resource to affect the group, which can cause the group to fail over to another system. |
| **RestartThreshold** | Specifies how many times the Cluster service attempts to restart the resource in the **RestartPeriod** before failing over the group. |
| **RestartPeriod** | Specifies the amount of time allowed for the restart attempts to reach the **RestartThreshold** before the Cluster service fails over the group. |
| **PendingTimeout** | The amount of time that a pending online or pending offline resource has to either resolve its status or report that it is making progress to the Resource Monitor before the Cluster service puts the resource in Failed status. The default value is 3 minutes. |
| **RetryPeriodonFailure** | Specifies the interval of time (in milliseconds) that a resource should remain in a failed state before the Cluster service attempts to restart it. |
| **LoadBalStartupInterval**<br><br>**LoadBalSampleInterval**<br><br>**LoadBalMinProcessorUnits**<br><br>**LoadBalMinMemoryUnits** | These properties are not used at this time by the Cluster service, but are reserved for future use. |

- Using cluster resource private property names

  To view the private properties for a resource, use the **/priv**[**properties**] parameter without any options. For example, to see the disk signature and other private properties for the Disk L: resource, type:

  **cluster resource "disk l:" /priv**

  The following tables describe some important private **cluster res**[**ource**] property names for File Share, Physical Disk, Network Name, Local Quorum, and Volume Shadow Copy Service Task resources, their uses, and valid settings.

**File Share resource**

| Private property names | Use |
|---|---|

| | |
|---|---|
| **CSCCache** | Controls client side caching for file shares on a server cluster. Client side caching allows network users to access files on file shares, even when the client computer is disconnected from the network, because the files have been cached on the client computer. The choices are:<br><br>**CSCCache** is set to 0. Changes in only those files specifically identified on the file share will be reconciled between the client and the server cluster.<br><br>**CSCCache** is set to 16. Changes in all files opened on the client will be reconciled between the client and the server cluster.<br><br>**CSCCache** is set to 32. Client side caching is disabled for those files that are only opened, but not modified.<br><br>**CSCCache** is set to 48. Client side caching is disabled for all files on this file share.<br><br>For example, to automatically reconcile changes in all files on the shared file resource FLSH1 opened on the client, type:<br><br>**cluster resource "FLSH1" /priv CSCCache=16** |
| **Security** | Sets file share-level permissions for file shares on a server cluster.<br><br>Usage: **Security=***Security group or user,access type,access mask***:security**<br><br>Access types are grant, deny, set, or revoke. An access mask is required for the types grant, deny, or set. Access masks are R for read, C for change, and F for full.<br><br>For example, to set read-only access to File Share resource FSHR1 for the Everyone group, type:<br><br>**cluster resource "FSHR1" /priv Security=Everyone,set,R:security** |

**Physical Disk resource**

| Private property names | Use |
|---|---|
| **Signature**<br><br>**DiskInfo**<br><br>**MountVolumeInfo** | Displays the disk signature and layout details of the disk resource. |
| **SkipChkdsk**<br><br>**ConditionalMount** | Controls if and how the chkdsk program runs on a corrupt shared cluster disk or inaccessible quorum disk. The choices are:<br><br>**SkipChkdsk** is set to 1. The Cluster service does not run Chkdsk against the drive and mounts the disk for immediate use. This setting overrides any setting for **ConditionalMount**.<br><br>**SkipChkdsk** set to 0, and **ConditionalMount** is set to 1. The Cluster service fails the disk resource and leaves it offline.<br><br>**SkipChkdsk** set to 0, and **ConditionalMount** is set to 1. The Cluster service runs the command **chkdsk /f** against the disk and then mounts it.<br><br>For example, to bypass **chkdsk** for a corrupt disk resource Disk Q: and mount it directly, type:<br><br>**cluster resource "Disk Q:" /priv SkipChkdsk=1** |

**Network Name resource**

| Private property names | Use |
|---|---|

| | |
|---|---|
| **RequireDNS** | Controls whether or not the network name must be registered with a DNS server so that the Network Name resource can be brought online. The choices are:<br><br>**RequireDNS** is set to 1. The DNS Host (A) record must be registered in DNS in order for the Network Name resource to be brought online. If the DNS server accepts dynamic updates but the DNS Host (A) record could not be updated, the Network Name resource will fail. If the DNS server does not accept dynamic updates, or if there are no DNS servers associated with the network, the Network Name resource will still come online.<br><br>**RequireDNS** is set to 0. Disables this option. The network name does not have to be registered in DNS in order for the Network Name resource to be brought online. |
| **RequireKerberos** | Controls whether or not the Network Name resource can be brought online without Kerberos authentication. The choices are:<br><br>**RequireKerberos** is set to 1. The next time this resource is brought online, clients using this resource will be authenticated using Kerberos. That is, a computer object for the Network Name resource will be created or updated before that resource comes online.<br><br>**RequireKerberos** is set to 0. Disables this option, that is, no computer object is created for the Network Name resource. This is the default. The Cluster service will try to detect if a computer object already exists for the Network Name resource; if a computer object is present, the resource will fail to come online. If this property were previously set to 1, setting it back to zero will disable the computer object created for this Network Name resource, if any.<br><br>Do not rename a Network Name resource and change the Kerberos authentication setting at the same time. Instead, allow time for changes to replicate after performing each action. |
| **CreatingDC** | Displays the fully qualified DNS name of the domain server on which the computer object for the Network Name resource was created. |
| **StatusDNS** | Displays the error status when the Network Name resource is registered with the DNS servers. If the value is zero, the action was successful; otherwise the value returned is the error code encountered. To view the text for the error code, *ErrorCode*, from the command prompt, type, **net helpmsg** *ErrorCode*. |
| **StatusKerberos** | Displays the error status when the computer object associated with the Network Name resource is created or updated. If the value is zero, the action was successful; otherwise the value returned is the error code encountered. |
| **StatusNetBIOS** | Displays the error status when the Network Name resource is registered with NetBIOS. If the value is zero, the action was successful; otherwise the value returned is the error code encountered. |

**Local Quorum resource**

| Private property names | Use |
|---|---|
| **Path** | Displays the default location for the quorum files if the Local Quorum resource is chosen as the quorum resource. You can override the default location of quorum files by explicitly specifying another local file path or a Server Message Block (SMB) path denoting a file share.<br><br>**Important**<br><br>If the path is not local, quorum logging is turned off. This can result in an out-of-date cluster database. |

| | |
|---|---|
| **Debug** | Specifies whether or not a server cluster that uses a Local Quorum resource as the quorum resource can have more than one node. The choices are:<br><br>Debug is set to 0. This is the default and specifies that the server cluster using a Local Quorum resource as the quorum resource can have only one node.<br><br>Debug is set to 1. This specifies that the server cluster using a Local Quorum resource as the quorum resource can have more than one node (that is, it can be a multi-node server cluster).<br><br>**Important**<br><br>If you have a multi-node server cluster that uses a Local Quorum resource as the quorum resource, you cannot set the Debug property to 0. |

**Volume Shadow Copy Service Task resource**

| Private property names | Use |
|---|---|
| **TaskName** | Specifies the name of the task created in the Scheduled Task folder.<br><br>**Caution**<br><br>If the name you provide for TaskName is the same as an existing scheduled task, that scheduled task will be deleted when the Volume Shadow Copy Service Task resource is brought online. |
| **ApplicationName** | Specifies the path to the executable program to be run for the scheduled task. |
| **ApplicationParams** | Specifies the parameters for the executable program to be run for the scheduled task. This property is used in combination with the **ApplicationName** property to create the full command line that will be run. For example, if the **ApplicationName** property is C:\WINDOWS\SYSTEM32\defrag.exe and the **ApplicationParams** property is -a C:, then the command run when the scheduled task is executed is: C:\WINDOWS\SYSTEM32\defrag.exe -a C:. |
| **TriggerArray** | Displays the schedule information for the scheduled task. This property cannot be specified using cluster.exe, but can be specified using Cluster Administrator or the Cluster APIs. For more information, see Checklist: Installing a Volume Shadow Copy Service Task resource. |

**Examples**

To create a Print Spooler resource in the Ops Print Spoolers group, type:

**cluster resource /create /group:"Ops Print Spoolers" /Type:"print spooler"**

To require that the Network Name resource CLUS N1 be authenticated by Kerberos before being brought online, type:

**cluster resource "CLUS N1" /priv RequireKerberos=1**

To give user accounts USR1 and USR2 in domain DOMAIN1, read-only access to File Share FSHR1, type:

**cluster resource "FSHR1" /priv Security=DOMAIN1\USR1,set,R,DOMAIN1\USR2,set,R:security**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| `Courier font` | Code or program output |

**See Also**

**Concepts**

Command-line reference A-Z

Command shell overview

Property lists

**Other Resources**

Cluster commands overview

# Cluster resourcetype

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster resourcetype**

The cluster resourcetype command is used to create a new cluster resource type or administer an existing resource type. Used without parameters, **cluster resourcetype** defaults to the **/list** command-line option and displays available installed cluster resource types.

**Syntax**

**cluster** [[**/cluster:**]*ClusterName*] **resourcetype** *ResourceType_DisplayName* **/***Option*

**Parameters**

**/list**
Lists the available installed cluster resource types.

**/listowners**
Displays a list of preferred owners.

**/create /Dll[Name]:*DLL_Name* [/type:*TypeName*] [/IsAlive:*Interval_in_Milliseconds*] [/LooksAlive:*Interval_in_Milliseconds*]**
Creates a resource type.

**/delete [/type]**
Deletes a resource type.

**/properties [*PropertyList*]**
Views or sets the resourcetype common properties.

**/privproperties [*PropertyList*]**
Views or sets the resourcetype private properties.

**/?**
Displays help at the command prompt.

**Remarks**

- Using the **cluster resourcetype** command

  The cluster name is optional. If the name of your cluster is also a cluster command or its abbreviation, such as "cluster" or "c," use **/cluster:** to explicitly specify the cluster name. For a list of all the cluster commands, see Related Topics.

- Using the **/list** command-line option

  Resource type display name is optional. If you do not provide a resource type display name, all installed resource types are listed.

- Using the **/delete** [**/type**] command-line option

  If the resource DLL cannot be accessed by the Cluster service, specify the resource type name with the /type command-line option instead of specifying the resource type display name.

- Using the **/prop**[**erties**] [*PropertyList*] command-line option

  Resource type display name is optional. Properties for all resourcetypes are displayed if you do not provide a resource type display name. For more information on property lists, see Related Topics.

- Using the **/priv**[**properties**] [*PropertyList*] command-line option

  The resource type display name is optional. If you do not provide a resource type display name, the private properties for

all resourcetypes are displayed. For more information on property lists, see Related Topics.

- Using cluster resourcetype common property names

  The following table describes the common **cluster resourcetype** property names, their uses, and valid settings.

| Common Property Names | Use |
|---|---|
| **Name** | Changes a resource type display name. |
| **Description** | Changes the text that describes a resource type. |
| **DllName** | Specifies the name of the dynamic-link library (DLL) for a specific resource type. |
| **DebugPrefix** | Specifies the appropriate debugger for the resource type. For more information, see the *Microsoft Platform Software Development Kit*. |
| **DebugControlFunctions** | Specifies if debugger breakpoint is set at the resource control function. (True (1) it is set; False (0) it is not set.) |
| **AdminExtensions** | Describes one or more class identifiers (CLSID) for Cluster Administrator extensions. |
| **LooksAlivePollInterval** | Specifies the interval (in milliseconds) that the Cluster service polls resources of a particular resource type to determine if the resources appear operational. |
| **IsAlivePollInterval** | Specifies the interval (in milliseconds) that the Cluster service polls resources of a particular resource type to determine if the resources are operational. |

- Using cluster resourcetype private property name

  By default, there are no private properties for **cluster resourcetype**. Software vendors can add private properties to extend cluster functionality.

**Examples**

To display the properties of the Generic Application resource type on the OpsClust1 cluster, type:

**cluster opsclust1 resourcetype "generic application" /properties**

**Formatting legend**

| Format | Meaning |
|---|---|
| Italic | Information that the user must supply |
| Bold | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**See Also**

**Concepts**

Command-line reference A-Z

Command shell overview

Property lists

**Other Resources**

Cluster commands overview

# Cluster command abbreviations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster command abbreviations**

The following table lists keyword abbreviations supported by Cluster.exe.

| Keyword | Supported abbreviations |
|---|---|
| **addcheckpoints** | **addchk**, **addcheck**, **addchkpoints** |
| **addcryptocheckpoints** | **addcryptochk**, **addcryptocheck**, **addcryptochkpoints** |
| **adddependency** | **adddep** |
| **changepassword** | **changepass** |
| **checkpoints** | **chk**, **check**, **chkpoints** |
| **cluster** | **c** |
| **cryptocheckpoints** | **cryptochk**, **cryptocheck**, **cryptochkpoints** |
| **delete** | **del** |
| **dllname** | **dll** |
| **expandstring** | **expandstr** |
| **forcecleanup** | **force** |
| **help** | **?** |
| **listdependencies** | **listdep** |
| **listinterfaces** | **listint**, **listinterface** |
| **moveto** | **move** |
| **multistring** | **multistr** |
| **network** | **net** |
| **netinterface** | **netint** |
| **offline** | **off** |
| **online** | **on** |
| **privproperties** | **priv**, **privprop**, **privprops** |
| **properties** | **prop**, **props** |
| **quorumresource** | **quorum** |
| **removecheckpoints** | **removechk**, **removecheck**, **removechkpoints** |
| **removecryptocheckpoints** | **removecryptochk**, **removecryptocheck**, **removecryptochkpoints** |
| **removedependency** | **removedep** |
| **removeowner** | **remowner** |

| | |
|---|---|
| **rename** | **ren** |
| **resource** | **res** |
| **resourcegroup** | **resgroup**, **group** |
| **resourcetype** | **restype**, **type** |
| **status** | **stat**, **state** |
| **string** | **str** |
| **version** | **ver** |

**See Also**

**Concepts**

Command-line reference A-Z

Command shell overview

**Other Resources**

Cluster commands overview

# Property lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Property lists**

You specify property lists on the Cluster.exe command line in the following ways:

- *Name* **=** *Value*[**,***Value* ...][**:***Format*][*Name*=*Value*[**,***Value* ...][**:***Format*] ...]

  For example:

  **cluster opsclust node mynode /priv prop1=string1,string2:multistring prop2=0xa,0xb:binary**

- **:** *PropertyName1*[**,PropertyName2,** ...]**/usedefault**

  For example:

  **cluster opsclust node mynode /priv:prop1,prop2 /usedefault**

To revert to the default value of a common property or to delete a private property, use the **/usedefault** command-line option.

The following table describes the available format options.

| Format | Description | Example |
|---|---|---|
| **Binary** | Separated list of values from 0 to 255 | myBinary = 12,0xab,255:binary |
| **DWord** | Double word | myDWord=2346:dword |
| **String** | Double quotation marks are required around strings with spaces | myString="string value":string |
| **ExpandString** | Expanded by Cluster.exe | myExpandString=^%windir^%:expandstr |
| **MultiString** | List of strings separated by commas | myMultiString="string value 1", "string value 2":multistr |
| **Security** | Set access permissions. Access types are grant, deny, set, or revoke. An access mask is required for the types grant, deny, or set. Access masks are R for read, C for change, and F for full. | mySecurity= devgroup \ user1, grant, RC, devgroup \ user2, revoke, everyone, set, F:security |
| **ULarge** | Unsigned 64-bit value | myULarge = 12345678901234567890:ularge |

**See Also**

**Concepts**

Command-line reference A-Z

Command shell overview

**Other Resources**

Cluster commands overview

# Managing the Cluster Service from the Command Line

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing the Cluster service from the command line**
You can use cluster service commands to troubleshoot or debug the Cluster service from the command prompt. The Cluster service command line can only be run locally from the cluster directory:

%systemroot%\cluster

Do not use the **clussvc** command under normal conditions, but only as a temporary diagnostic tool if the Cluster service fails to start (because of a system-related error, for example).

**Important**

- If you use the clussvc command to start the Cluster service in a server cluster with 3 or more nodes, the Cluster service will be unable to authenticate any intracluster messages sent in multicast mode. The intracluster traffic in the server cluster will then revert to unicast mode. However, this will not affect the ability of any node to participate in the server cluster. For information, see Network communication between nodes.

- Do not change the Cluster service account password while running the Cluster service from the command line. For information on changing the Cluster service account password, see Change the Cluster service account password

For information about other command-line utilities, see Command-line reference. For more information about manageability, see Management Strategies and Tools.

The following rules apply when entering commands from the command prompt:

- When entering commands from the command prompt, you can specify multiple options for one command line. Clussvc.exe processes options from left to right. If an option fails, the command stops running at the failed option.

- To use the Cluster service commands, you must be logged on with administrative rights.

- To stop the Cluster service, at the command prompt, press CTRL+C.

The command line for starting the Cluster service has the following syntax:

**clussvc /debug** [*command*]

The Cluster service runs in the command prompt window and displays all output recorded in the cluster log in the window.

| Cluster service command | Action |
|---|---|
| **/debugresmon** | Enables the debugging of the resource dynamic-link libraries (DLLs) that are loaded by the resource monitor process. |
| **/fixquorum** | Allows the Cluster service to start up, despite problems with the quorum device. |
| **/resetquorumlog** | If the quorum log file is not found or is corrupted, creates a new quorum log file based on information in the local node's cluster database file. If the quorum log file is found and is not corrupted, this command has no effect. |
| **/norepevtlogging** | Allows no replication of event log entries. |

| | |
|---|---|
| **/forcequorum** | Restores quorum for a majority node set server cluster that has lost quorum.<br><br>**Note**<br><br>• It is recommended that you use the method described in Force quorum in a majority node set server cluster instead of the clussvc /debug /forcequorum command option described here. |

For more information on commands for the Cluster service, see:

- Debugresmon command
- Fixquorum command
- Resetquorumlog command
- Norepevtlogging command
- Forcequorum command

# Debugresmon command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Debugresmon command**

The **debugresmon** command helps you to debug the resource monitor process and, therefore, the resource dynamic-link libraries (DLLs) that are loaded by the resource monitor. You can use any standard Windows-based debugger, such as Microsoft Developer Studio.

The syntax is:

**clussvc /debug /debugresmon**

**Note**

- The **debugresmon** command can be used only when starting the cluster service in the command prompt window. The **debugresmon** command cannot be used as a start parameter when starting the cluster service in the Computer Management console.

Just before the resource monitor process is started up, the cluster service issues the message:

"Waiting for debugger to connect to the resmon process *x*"

where *x* is the process identification number (PID) of the resource monitor process.

After the resource monitor process has started and you have attached a debugger to the resource monitor process, the Cluster service continues with its initialization.

For more information on the Resource Monitor, see Resource Monitors.

For more information on resource DLLs, see Resource dynamic-link libraries (DLLs) in clusters running Windows Server 2003.

For more information on debugging resource DLLs, see the Microsoft Software Development Kit (SDK).

# Fixquorum command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Fixquorum command**

The **fixquorum** command allows you to start the cluster service even if there are problems with the quorum resource. Only one node at a time can be started with this command. You cannot join any other nodes to the node started using this command.

The syntax is:

**clussvc /debug /fixquorum**

Preferably, you can use the **fixquorum** command as a start parameter when starting the cluster service in the Computer Management console. For more information, see Set the Cluster service to only start manually.

When the **fixquorum** command is issued on a particular node, the Cluster service starts, but all the resources, including the quorum resource, remain offline. This allows you to then manually bring the quorum resource online and more easily diagnose quorum device failures. You can use this command, for example, if you need to switch from a problematic quorum disk to another quorum resource. For more information, see Use a different disk for the quorum resource.

# Resetquorumlog command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resetquorumlog command**

The **resetquorumlog** command creates a new quorum log file, if missing or corrupted, using information stored in the local node's cluster database and creates a new registry checkpoint file. The new quorum file is created using information in the cluster database located in %systemroot%\cluster\CLUSDB. If the quorum log file is not missing or corrupt, this command has no effect.

The syntax is:

**clussvc /debug /resetquorumlog**

The **resetquorumlog** command can also be set as a start parameter when starting the cluster service in the Computer Management console.

**Caution**

- The information retrieved from the cluster database might be outdated. Therefore, use this command only if there are no backups of the quorum log file available. You must use the Backup or Restore Wizard even if you do not have a backup of the quorum log. For more information on restoring server cluster data using the Backup or Restore Wizard, see Backing up and restoring server clusters.

  For more information on recovering from a corrupted quorum log, see Recover from a corrupted quorum log or quorum disk.

**Note**

- The registry checkpoint file created by the **resetquorumlog** command has the format CHKxxx.TMP.

# Norepevtlogging command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Norepevtlogging command**

The **norepevtlogging** command prevents replication of those events recorded in the event log. This command is useful in reducing the amount of information displayed in the command window by filtering out events already recorded in the event log.

The syntax is:

**clussvc /debug /norepevtlogging**

The **norepevtlogging** command can also be set as a start parameter when starting the cluster service in the Computer Management console.

For example, to start the cluster service and log those events not recorded in the event log to a local file, Debugnorep.log:

```
clussvc /debug /norepevtlogging > c:\DEBUGNOREP.log\
```

For more information on event logs, see Event Viewer.

# Forcequorum command

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Forcequorum command**

The **forcequorum** command restores quorum for a majority node set server cluster that has lost quorum. This command is useful if a partition of a geographically dispersed server cluster (for example all the nodes in one location) has lost communication with the nodes at another location.

The syntax is:

**clussvc /debug /forcequorum** *node list*

For example, to start the cluster service on nodes Node1, Node2, and Node3, type:

```
clussvc /debug /forcequorum Node1,Node2,Node3
```

**Important**

- When you force quorum in a majority node set cluster, an error message from the Service Control Manager may appear, indicating that the Cluster service failed to start. You can safely ignore this message at this time; the Cluster service will continue startup after the end of the startup timeout period for Service Control Manager.

**Note**

- It is recommended that you use the method described in Force quorum in a majority node set server cluster instead of the clussvc /debug /forcequorum command option described here.

# Resources for Server Cluster information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

- For a good technical reference on clustering technologies, see Pfister, Gregory F. *In Search of Clusters*. Upper Saddle River, NJ: Prentice Hall, 1998.

- For information on Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition, see the Microsoft Web site.

- Microsoft maintains up-to-date Windows Server 2003 family product documentation at the Microsoft Web site. This site has information that is more current than the Help documentation on your product CD.

- For technical support information on server clusters, see Updated technical information.

- The Microsoft Windows Resource Kits contain additional technical, planning, and deployment information on server clusters. For more information about server cluster deployments, see "Planning for High Availability" at the Microsoft Windows Resource Kits Web site.

- For a detailed description of how a business has used Network Load Balancing clusters and server clusters to improve performance and availability of their Web site and other services, refer to the clustering scenario at the Microsoft Web site.

- Requests for Comments (RFCs) are an evolving series of reports, proposals for protocols, and protocol standards used by the Internet community. Standards are defined in RFCs published by the Internet Engineering Task Force (IETF) and other working groups. For more information on private IP network addresses, see RFC 1918: Address Allocation for Private Internets.

  You can obtain RFCs from the RFC Editor Web site.

  Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.

- Cluster Administrator supports command-line parameters. If you start Cluster Administrator from the command line, you can connect to one or more clusters by specifying the cluster name, the name of one of the cluster nodes, or the name of a cluster IP address.

  When you specify the clusters to connect to using the Cluster Administrator command line, Cluster Administrator will not try to restore your most recent cluster connections. You can also prevent Cluster Administrator from restoring your most recent cluster connections by specifying the **cluadmin /noreconnect** command at the command line.

  For more information on administering a cluster from the command line, see Managing a Server Cluster from the Command Line.

- The following links point to information that may help you better understand server clusters within the overall context of Windows Clustering.

  For information on the solutions that make up Windows Clustering, see Windows Clustering.

  For information on Network Load Balancing clusters, see Network Load Balancing Clusters.

  For information on server clusters and how they work with Network Load Balancing clusters, see Updated technical information.

- Server clusters support unattended setup, a feature of Windows Server 2003 family operating systems.

For information about unattended setup, see Planning for unattended Setup.

# Server Cluster Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**

This section provides symptom-based diagnostic procedures for troubleshooting clustering problems. You can also use the Cluster Diagnostics and Verification Tool (ClusDiag) to run diagnostic tests. These tests will help you verify the functionality of a cluster and troubleshoot cluster failures using the generated log files. ClusDiag is available in the *Microsoft Windows Server 2003 Resource Kit* and at the Microsoft Download Center at the Microsoft Web site . For information about how to obtain product support, see Technical support options.

This section covers:

- Additional New Server Cluster Wizard information
- Cluster installation problems
- Cluster disk and drive connection problems
- Node-to-node connectivity problems
- Client-to-cluster connectivity problems
- Group and resource failure problems
- File and print server cluster problems
- General administrative problems

# Additional New Server Cluster Wizard information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Additional New Server Cluster Wizard Information**
Problems that occur when you use the New Server Cluster Wizard to create a new server cluster are usually the result of misconfigured cluster networks and hardware or permissions issues with the Cluster service account. You can use the information in this topic to help diagnose and troubleshoot these problems.

This topic contains overviews of clustering technology, additional information on error and event log messages, step-by-step implementation guides, scenarios, top troubleshooting issues, and relevant Knowledge Base (KB) articles from the Microsoft Product Support Services Web site.

**Top Troubleshooting Issues**

- Cluster installation problems
- Cluster disk and drive connection problems
- Node-to-node connectivity problems
- Client-to-cluster connectivity problems
- Group and resource failure problems
- General administrative problems

**Error and Event Log Messages**
To search for additional information on error and event log messages, click **Error and Event Log Messages** in Help and Support Center.

**Top Knowledge Base (KB) articles**
For information about server clusters, see the following key articles in the Microsoft Knowledge Base.

| Problem Area | Key KB Article # | Title |
|---|---|---|
| Understanding virtual servers | Q257932 | Using Microsoft Cluster Server to Create a Virtual Server |
| Event 1034 occurs when replacing a shared SCSI disk | Q280425 | Recovering from an Event ID 1034 on a Server Cluster |
| Network configuration errors | Q259267 | Microsoft Cluster Service Installation Resources |
| Creating clustered home folders | Q256926 | Implementing Home Folders on a Server Cluster |
| Setting up clustered printers | Q197046 | How to Set up a Print Spooler on Microsoft Cluster Server |
| Starting the Cluster service | Q262962 | Cluster Service Does Not Start with Event ID 7000: Not Responding to Control Function |
| Accessing cluster disks | Q275704 | Cluster Disks May Be Incorrectly Accessible from All Nodes |
| Setting up "heartbeat" configurations for private networks | Q258750 | Recommended Private "Heartbeat" Configuration on Cluster Server |
| Setting up Cluster service account permissions | Q269229 | How to Manually recreate the Cluster Service Account |

- For a list of the top server cluster support issues, see the Microsoft Web site
- To find other relevant articles, search the Microsoft Knowledge Base using the keywords *server cluster*.

**Help Topics**

- For a checklist of the steps to follow when creating a new server cluster or adding a node to an existing cluster, see Checklist: Planning and creating a server cluster.
- For an overview of server clusters, see Server clusters overview.
- For an overview of the three complementary Microsoft Windows Clustering technologies (server clusters, Network Load Balancing, and Component Load Balancing), see Windows Clustering.

**Step-by-Step Implementation Guides and Scenarios**

- For a tutorial on Cluster Installation, see the Microsoft Web site.
- For example deployments of Windows Clustering, see "Planning for High Availability" at the Microsoft Windows Resource Kits Web site.

# Cluster installation problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster installation problems**

For help in diagnosing and solving problems with creating a cluster or adding one or more nodes to an existing cluster, see Troubleshooting cluster node installations.

For more information about troubleshooting cluster installation issues, see Knowledge Base article 266274, "How to Troubleshoot Cluster Service Startup Issues" in the Microsoft Knowledge Base.

For information about how to obtain product support, see Technical support options.

# Cluster disk and drive connection problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cluster disk and drive connection problems**
What problem are you having?

- When the physical disks are not powering up or spinning, Cluster service cannot initialize any quorum resources.
- The Cluster service fails to start and generates an Event ID 1034 in the Event log after you replace a failed hard disk, or change drives for the quorum resource.
- Drive on the shared storage bus is not recognized.
- Configuration cannot be accessed through Disk Management.
- SCSI or fibre channel storage devices do not respond.
- Disk groups do not move or stay online pending after move.
- Disks do not come online or Cluster service does not start when a node is turned off.
- Drives do not fail over or come online.
- Mounted drives disappear, do not fail over, or do not come online.
- The cluster quorum disk (containing the quorum resource) becomes disconnected from all nodes in a cluster and you are later unable to add the nodes back to the cluster.

**When the physical disks are not powering up or spinning, Cluster service cannot initialize any quorum resources.**
**Cause:** Cables are not correctly connected, or the physical disks are not configured to spin when they receive power.

**Solution:** After checking that the cables are correctly connected, check that the physical disks are configured to spin when they receive power.

**The Cluster service fails to start and generates an Event ID 1034 in the Event log after you replace a failed hard disk, or change drives for the quorum resource.**
**Cause:** If a hard disk is replaced, or the bus is reenumerated, the Cluster service may not find the expected disk signatures, and consequently may fail to mount the disk.

**Solution:** Write down the expected signature from the Description section of the Event ID 1034 error message. Then follow these steps:

1. Backup the server cluster.
2. Set the Cluster service to start manually on all nodes, and then turn off all but one node.
3. If necessary, partition the new disk and assign a drive letter.
4. Use the confdisk.exe tool (available in the *Microsoft Windows Server 2003 Resource Kit*) to write that signature to the disk.
5. Start the Cluster service and bring the disk online
6. If necessary, restore the cluster configuration information.

7. Turn on each node, one at a time.

For information on replacing disks in a server cluster, see Knowledge Base article Q305793, "How to Replace a Disk with Windows 2000 or Windows Server 2003 family Clusters" in the Microsoft Knowledge Base.

**Drive on the shared storage bus is not recognized.**
**Cause:** Scanning for storage devices is not disabled on each controller on the shared storage bus.

**Solution:** Verify that scanning for storage devices is disabled on each controller on the shared storage bus.

Many times, the second computer you turn on does not recognize the shared storage bus during the BIOS scan if the first computer is running. This situation can manifest itself in a "Device not ready" error being generated by the controller, or in substantial delays during startup.

To correct this, disable the option to scan for devices on the shared controller.

### Note

- This symptom can manifest itself as one of several errors, depending on the attached controller. It is normally accompanied with a one- to two-minute start delay and an error indicating the failure of some device.

**Configuration cannot be accessed through Disk Management.**
Under normal cluster operations, the node that owns a quorum resource locks the drive storing the quorum resource, preventing the other nodes from using the device. If you find that the cluster node that owns a quorum resource cannot access configuration information through Disk Management, the source of the problem and the solution might be one of the following:

**Cause:** A device does not have physical connectivity and power.

**Solution:** Reseat controller cards, reseat cables, and make sure the drive spins up when you start.

**Cause:** You attached the cluster storage device to all nodes and started all the nodes before installing the Cluster service on any node.

**Solution:** After you attach all servers to the cluster drives, you must install the Cluster service on one node before starting all the nodes. Attaching the drive to all the nodes before you have the cluster installed can corrupt the file system on the disk resources on the shared storage bus.

**SCSI or fibre channel storage devices do not respond.**
**Cause:** The SCSI bus is not properly terminated.

**Solution:** Make sure that the SCSI bus is not terminated early and that the SCSI bus is terminated at both ends.

**Cause:** The SCSI or fibre channel cable is longer than the specification allows.

**Solution:** Make sure that the SCSI or fibre channel cable is not longer than the cable specification allows.

**Cause:** The SCSI or fibre channel cable is damaged.

**Solution:** Make sure that the SCSI or fibre channel cable is not damaged. (For example, check for bent pins and loose connectors on the cable and replace it if necessary.)

**Disk groups do not move or stay online pending after move.**
**Cause:** Cables are damaged or not properly installed.

**Solution:** Check for bent pins on cables and make sure that all cables are firmly anchored to the chassis of the server and drive cabinet.

**Disks do not come online or Cluster service does not start when a node is turned off.**
**Cause:** If the quorum log is corrupted, the Cluster service cannot start.

**Solution:** If you suspect the quorum resource is corrupted, see the information on the problem "Quorum log becomes corrupted" in Node-to-node connectivity problems.

**Drives do not fail over or come online.**
**Cause:** The drive is not on a shared storage bus.

**Solution:** If drives on the shared storage bus do not fail over or come online, make sure the disk is on a shared storage bus, not on a nonsystem bus.

**Cause:** If you have more than one local storage bus, some drives in **Shared cluster disks** will not be on a shared storage bus.

**Solution:** If you do not remove these drives from **Shared cluster disks**, the drives do not fail over, even though you can configure them as resources.

**Shared cluster disks** is in the Cluster Application Wizard.

**Mounted drives disappear, do not fail over, or do not come online.**
**Cause:** The clustered mounted drive was not configured correctly.

**Solution:** Look at the Cluster service errors in the Event Log (ClusSvc under the Source column). You need to recreate or reconfigure the clustered mounted drive if the description of any Cluster service error is similar to the following:

Cluster disk resource "*disk resource*": Mount point "*mount drive*" for target volume "*target volume*" is not acceptable for a clustered disk because *reason*. This mount point will not be maintained by the disk resource.

When recreating or reconfiguring the mounted drive(s), follow these guidelines:

- Make sure that you create unique mounted drives so that they do not conflict with existing local drives on any node in the cluster.
- Do not create mounted drives between disks on the cluster storage device (cluster disks) and local disks.
- Do not create a mounted drive from a clustered disk to the cluster disk that contains the quorum resource (the quorum disk). You can, however, create a mounted drive from the quorum disk to a clustered disk.
- Mounted drives from one cluster disk to another must be in the same cluster resource group, and must be dependent on the root disk.

For more information on viewing the Event Log, see View Event Logs.

For more information on creating mounted drives in a server cluster, see Add drives on the shared storage bus.

**The cluster quorum disk (containing the quorum resource) becomes disconnected from all nodes in a cluster and you are later unable to add the nodes back to the cluster.**
See the solution for the identically titled problem in Node-to-node connectivity problems.

For information about how to obtain product support, see Technical support options.

It is important that you correctly configure the storage topology (for example, SCSI, Fibre Channel, Storage Area Networks) and the storage interconnects (for example, multiple paths) used in your server cluster. Before deploying your server cluster, contact your hardware vendors to ensure that your particular cluster storage configuration is supported at the hardware level. For descriptions of supported cluster storage topologies, best practices for deploying and managing cluster storage, and a list of cluster storage-related Knowledge Base articles, see the cluster storage information at the Microsoft Web site.

# Node-to-node connectivity problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Node-to-node connectivity problems**
What problem are you having?

- I cannot complete a cluster on the first node.
- When the resources fail over and the nodes do not detect each other, there is no connectivity between the nodes or with the cluster storage device.
- Quorum resource does not start.
- Quorum resource fails.
- Quorum log becomes corrupted.
- Additional node cannot join the cluster.
- Nodes cannot connect to the cluster drives.
- The cluster quorum disk (containing the quorum resource) becomes disconnected from all nodes in a cluster and you are later unable to add the nodes back to the cluster.

**I cannot complete a cluster on the first node.**
**Cause:** Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition is incorrectly installed.

**Solution:** Make sure that Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition is correctly installed.

**Cause:** Your hardware is not listed in the Cluster category on the Microsoft Windows Catalog.

**Solution:** Make sure that your hardware is listed on the Microsoft Windows Catalog. See the Windows Catalog at Support resources. Search for "Cluster".

If any of the hardware you are using for your cluster is not on this list, consider replacing those components with hardware that is listed.

**Cause:** You have chosen individual components from the Windows Catalog instead of systems.

**Solution:** Do not choose individual *components* of your cluster from the Microsoft Windows Catalog. Microsoft supports only *systems* chosen from this list.

**Cause:** Your primary Internet protocol (IP) address is invalid.

**Solution:** If the node uses DHCP to obtain noncluster IP addresses, use Ipconfig.exe to verify that you have a valid primary IP address for all network adapters. If the second IP address listed (the subnet mask) is 0.0.0.0, your primary address is invalid.

**When the resources fail over and the nodes do not detect each other, there is no connectivity between the nodes or with the cluster storage device.**
**Cause:** The Remote Procedure call (RPC) service is not running.

**Solution:** On each node, use **Services** in Control Panel to confirm that the RPC service is running.

**Cause:** The nodes do not have RPC connectivity.

**Solution:** Verify that the nodes have RPC connectivity.

You can determine this by using a network analyzer (such as Network Monitor), or you can use RPCPing (available on the Microsoft Exchange Server CD)

**Quorum resource does not start.**
**Cause:** The resource is not physically connected to the server.

**Solution:** Make sure that the resource is physically connected to the server.

**Quorum resource fails.**
**Cause:** The disk on the shared bus holding the quorum resource has failed.

**Solution:** If the disk on the shared bus holding the quorum resource fails and cannot be brought online, the Cluster service cannot start. To correct this situation, you must use the **fixquorum** option to start the Cluster service on a single node, and then use Cluster Administrator to configure the Cluster service to use a different disk on the shared bus for the quorum resource.

When **fixquorum** is specified, the Cluster service starts without a quorum resource, and does not bring the quorum disk online. A node cannot join a cluster when the Cluster service is running with the **fixquorum** option.

For instructions on how to change the disk that the Cluster service uses for the quorum resource, see Use a different disk for the quorum resource.

**Quorum log becomes corrupted.**
**Cause:** This may occur for a variety of reasons.

**Solution:** If the quorum log is corrupted, the Cluster service attempts to correct the problem by resetting the log file. In this case, the Cluster service writes the following message in the system log:

The log file [*name*] was found to be corrupt. An attempt will be made to reset it.

If the quorum log cannot be reset, the Cluster service cannot start.

If the Cluster service fails to detect that the quorum log is corrupted, the Cluster service may fail to start. In this case, there may be an "ERROR_CLUSTERLOG_CORRUPT" message in the system log.

To correct this, you must use the **noquorumlogging** option when starting the Cluster service to temporarily run the Cluster service without quorum logging. You can then correct the disk corruption and delete the quorum log, as necessary. When **noquorumlogging** is specified, the Cluster service brings the quorum disk online, but disables quorum logging. You can then run Chkdsk on your quorum disk to detect and correct disk corruption.

For instructions on how to recover from a corrupted quorum log or quorum disk, see Recover from a corrupted quorum log or quorum disk.

**Additional node cannot join the cluster.**
**Cause:** The cluster configuration on the node may not have been completely removed if the node was previously evicted.

**Solution:** At a command prompt, type **cluster** [*cluster name*] **node***node name***/forcecleanup**.

When an additional node fails to join a cluster, improper name resolution is often the cause. The problem may exist because of invalid data in the WINS cache. You may also have the wrong binding on the WINS or DNS Services for the additional node.

If WINS or DNS is functioning correctly on all nodes:

**Cause:** You may not be using the proper cluster name, node name, or IP address.

**Solution:** Confirm that you are using the proper cluster name, node name, or IP address.

When joining a cluster, you can specify the cluster name, the computer name of the first node, or the IP address of either the cluster or the first node.

**Cause:** The Cluster Name resource may not have started.

**Solution:** Confirm that the Cluster Name resource started.

Use Cluster Administrator on the first node to ensure that the Cluster Name resource is running.

**Cause:** The Cluster service may not be running on the first node.

**Solution:** Confirm that the Cluster service is running on the first node and that all resources within the Cluster Group are online before installing a second node.

The Cluster service may not have yet started when you attempted to join the cluster.

**Cause:** Network connectivity may not exist between the nodes.

**Solution:** Confirm that network connectivity exists between the nodes.

Make sure TCP/IP is properly configured on all nodes.

**Cause:** You may not have IP connectivity to the cluster address.

**Solution:** Confirm that you have IP connectivity to the cluster address and that the IP address is assigned to the correct network.

If you cannot ping the IP address of the cluster, run Cluster Administrator on the first node and ensure the cluster IP Address resource is running. Also, use Cluster Administrator to ensure that the cluster has a valid IP address and subnet mask (click **Cluster Group**, right-click **Cluster IP Address**, and click **Properties**), and that the IP address does not conflict with an IP address that is already in use on the network. If the address is not valid, change it, take the Cluster IP Address resource offline, and then bring it online again. If the IP address is not assigned to the correct network, use Cluster Administrator to correct the problem.

If your cluster nodes use DHCP to obtain noncluster IP addresses, use Ipconfig.exe to verify that you have a valid primary IP address for the networks in question. If the second IP address listed (the subnet mask) is 0.0.0.0, your primary address is invalid.

**Cause:** The network role may have changed from **Use for all communication**.

**Solution:**Confirm that the network role has not been changed from **Use for all communication**.

A network role is initialized to **Use for all communications** but can be changed by the administrator. After verifying that you have IP connectivity to the cluster address and that the IP address is assigned to the correct network, use Cluster Administrator to confirm that at least one of the networks connected between the nodes is enabled to **Use for all communications**. To verify the role in Cluster Administrator, open the Networks folder, right-click **Network**, and then click **Properties**. **Use for all communications** enables the network to have both node-to-node communication and client-to-cluster communication.

For more information on networks, see Server cluster networks. For more information on network roles, see Configuring cluster network hardware. To use Cluster Administrator to reset the network's role to **Use for all communications**, see Enable a network for cluster use.

**Nodes cannot connect to the cluster drives.**

**Cause:** The same drive letters may not have been assigned to the cluster drives on all nodes.

**Solution:** Confirm that the cluster drives are assigned the same drive letters on all nodes.

To do so, run Disk Management on each node and make sure that identical drive letters are assigned to all cluster drives.

**Where?**

- Computer Management/Storage/Disk Management

**Cause:** The node may not be physically connected to the cluster drive.

**Solution:** Confirm that the node is physically connected to the cluster drive.

If it is not, shut down all nodes and the cluster drive. Connect the nodes to the shared storage bus. Then, start the cluster drive and start the first node. After the Cluster service starts on the first node, start the additional nodes, and attempt to connect to the cluster drive.

**Cause:** If you are using a shared SCSI bus, the SCSI devices may not have unique IDs.

**Solution:** Verify that each SCSI device has a unique ID. SCSI controller IDs are preset to seven. Reset one SCSI controller ID to six.

**Cause:** The controllers on the shared storage bus may not be correctly configured.

**Solution:** Confirm that the controllers on the shared storage bus are correctly configured (with all cards configured to transfer data at the same rate).

**Cause:** The devices and controllers may not match.

**Solution:** Confirm that your devices and controllers match.

For example, do not use a wide connection controller on one node and a narrow connection controller on another node. It is also recommended that all fibre channel controllers be homogeneous, so do not use different brands of controllers together.

**The cluster quorum disk (containing the quorum resource) becomes disconnected from all nodes in a cluster and you are later unable to add the nodes back to the cluster.**
**Cause:** The cluster configuration on the nodes may not have been completely removed.

**Solution #1:** Use the cluster.exe node /force[cleanup] command to evict the nodes from the cluster. For more information, see Evict a node from the cluster.

**Solution #2:** Use the Cluster service fixquorum start up parameter to start the Cluster service. Only one node at a time can be started with this command. You cannot join any other nodes to the node started using this command. For more information, see Recover from a corrupted quorum log or quorum disk.

For information about how to obtain product support, see Technical support options.

# Client-to-cluster connectivity problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Client-to-cluster connectivity problems**
What problem are you having?

- Clients cannot connect to virtual servers.
- Clients cannot access resources.
- A node cannot communicate on a network.
- Clients cannot access a group that has failed over.
- Clients cannot attach to a cluster file share resource.
- Clients cannot access a cluster resource.
- Client can detect all nodes but cannot detect a virtual server.
- Clients cannot connect to a virtual server and the system event logs on the client computers contain Kerberos authentication errors.

**Clients cannot connect to virtual servers.**
**Cause:** The client might not be using the correct network name or IP address to access the cluster.

**Solution:** Make sure that the client is accessing the cluster using the correct network name or IP address.

**Cause:** The client might not have the TCP/IP protocol correctly installed and configured.

**Solution:** Make sure that the client has the TCP/IP protocol correctly installed and configured.

Depending on the application being accessed, the client can address the cluster by specifying either the resource network name or the IP address. In the case of the network name, you can verify proper name resolution by checking the NetBT cache (using the Nbtstat.exe utility) to determine whether the name had been previously resolved. Also, confirm proper WINS configuration, both at the client and through the WINS Administrator.

If the client is accessing the resource through a specific IP address, ping the IP address of the cluster resource and cluster nodes from a command prompt.

For more information on Nbtstat.exe, see Nbtstat.

For instructions on how to troubleshoot connectivity problems, see Troubleshoot client-to-virtual server connectivity.

**Cause:** The client may be attempting to connect to an incorrect node after a failover.

**Solution:** Make sure that the routers are configured correctly. For more information, see article Q244331, "MAC Address Changes for Virtual Server During a Failover with Clustering" in the Microsoft Knowledge Base.

**Clients cannot access resources.**
**Cause:** Both your private network and the primary network are down.

**Solution:** If your server cluster nodes are multihomed and use a private network to communicate, the primary network can be down and the Cluster service can still function normally.

If the private network is lost between nodes, resources that are configured to do so will fail over to the node that has ownership of the quorum resource. Each of the other nodes will stop its Cluster service.

**Cause:** A node might have lost connectivity with the network.

**Solution:** If you suspect a node has lost connectivity with the network:

- Confirm that the connection configuration settings are correct for each connection. To do this, open **Control Panel**, then double-click **Network Connections**. Right-click a connection and click **Properties**.
- Check WINS or DNS to make sure that they are properly configured.
- Confirm that the static IP addresses being used by the Cluster service are still in place and are not being used by other resources on the network.

**A node cannot communicate on a network.**
**Cause:** If a node cannot communicate on a network, it might lack a physical connection to the other nodes, or the other nodes might not be connected to one another.

If the cabling has failed, the Cluster service might not have received the heartbeat of the node and so failed the resources owned by that node over to other nodes in the cluster. The node can no longer communicate to clients on a network.

**Solution:** Check the hub and local cabling, both to the network and to the cluster disks.

**Clients cannot access a group that has failed over.**
**Cause:** There might not be a physical connection between the cluster nodes and the group that failed over has also failed to come online.

**Solution -- step 1:** Confirm that you have a physical connection between the cluster nodes.

- Make sure that the network cabling has not failed or been damaged.

  If the cabling has failed, the Cluster service might have failed over the resources to another node.
- Ensure that the storage device cabling has not failed.
- If the nodes are separated by one or more hubs, check the connectivity through all hubs.

**Solution -- step 2:** Bring the group back online.

**Clients cannot attach to a cluster file share resource.**
**Cause:** WINS or DNS might not be correctly configured.

**Solution:** Make sure that WINS or DNS are correctly configured.

If you are using WINS, run WINS Manager on the WINS server; then, on the **Mappings** menu, click **Show Database**. Make sure that each node and the cluster are registered in the WINS database and that the registrations are active.

For more information on WINS and WINS Manager, see the *Networking Guide* in the *Microsoft Windows Server 2003 Resource Kit*.

**Cause:** Your security policies might not allow the client to access the file share.

**Solution:** Make sure that your security policies allow the client to access the file share.

The client must have the right to log on to the file share, or the Guest account must be enabled for the client to have access.

**Clients cannot access a cluster resource.**
**Cause:** Either the IP Address resource or Network Name resource for the group in which the resource is contained is not online.

**Solution:** Check the group dependencies; the resource is supposed to be dependent on either the IP Address or the Network Name resource. Ensure that the IP Address and Network Name resources are online in the resource group. From the client computer, try to ping the IP addresses of the virtual server and individual nodes.

**Cause:** Either the client or the cluster computer is not configured for either WINS or DNS.

**Solution:** Make sure that the cluster nodes are configured properly for name resolution using either WINS or DNS, and make sure that clients are configured to use the same form of name resolution.

**Cause:** The client is attempting to access the cluster from a different subnet, and DNS is not correctly configured.

**Solution:** Configure the cluster nodes and client computer to use WINS or DNS. If you use DNS, add a DNS address record for the cluster in the DNS database.

**Client can detect all nodes but cannot detect a virtual server.**
**Cause:** The virtual server might not have its own IP Address and Network Name resources.

**Solution:** Make sure that the virtual server has its own IP Address and Network Name resources, and that both resources are online.

Your servers must have an address and a name on the network for any other server or client to properly recognize that the servers are on the network.

**Cause:** One or more nodes are not correctly configured to use WINS or DNS.

**Solution:** Make sure that all nodes are correctly configured to use WINS or DNS.

**Clients cannot connect to a virtual server and the system event logs on the client computers contain Kerberos authentication errors.**
**Cause:** The cluster might have used a previously-created computer object for the virtual server. When the virtual server's Network Name resource is brought online for the first time after Kerberos support has been enabled, the cluster changes the object's password. Active Directory then replicates this change to all domain controllers in the cluster. However, because of replication latencies, clients connecting to this virtual server might receive a Kerberos ticket encrypted with the old password. This can result in Kerberos authentication errors.

**Solution:** Wait until the new password has been replicated to all domain controllers, or force replication using the **repadmin** and **replmon** command-line tools, available in the \Support\Tools folder on your operating system installation CD.

For information on monitoring and troubleshooting replication issues, see Troubleshooting replication.

For information about how to obtain product support, see Technical support options.

# Group and resource failure problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Group and resource failure problems**
What problem are you having?

- A resource fails, but is not brought back online.
- You cannot bring a resource online.
- You cannot bring the default physical disk resource online in Cluster Administrator.
- In Disk Management, you do not see the disk for the group that is online on that node.
- You are unable to manually move a group, or it does not fail over to another node when it is supposed to.
- A group failed over but did not fail back.
- The entire group failed and has not restarted.
- All nodes are functioning, but resources fail back repeatedly.
- The Cluster service does not successfully fail over resources.
- You fail over a resource group from one node to another, but it automatically fails back.
- The Network Name resource fails when you change to a system locale that is different than the input language used by the Network Name resource.
- The Message Queuing resource fails to handle message activity correctly which may result in resource failures.
- A third-party resource fails to come online in a mixed-version cluster or while upgrading a cluster.

**A resource fails, but is not brought back online.**
**Cause:** A resource may depend on another resource that has failed.

**Solution:** In the resource **Properties** dialog box, make sure that the **Do not restart** check box is clear. If the resource needs another resource to function, and if the second resource fails, confirm that the dependencies are correctly configured.

**You cannot bring a resource online.**
**Cause:** The resource is not properly installed.

**Solution:** Make sure the application or service associated with the resource is properly installed.

**Cause:** The resource is not properly configured.

**Solution:** Make sure the properties are set correctly for the resource.

**Cause:** The resource is not compatible with server clusters.

**Solution:** Not all applications can be configured to fail over in a cluster. For more information, see Choosing applications to run on a server cluster.

**Cause:** The resource is generating a specific error.

**Solution:** Review the system Event Log (look for ClusSvc entries under the Source column) to see if that resource is generating a specific error message.

**You cannot bring the default physical disk resource online in Cluster Administrator.**
Most cluster configuration problems result from improper configuration of the shared storage bus or the restart of servers.

**Cause:** You may not have restarted the servers after installing the Cluster service.

**Solution:** Make sure that you restarted all servers after installing the Cluster service.

When the servers are restarted, the signature of each disk in the cluster storage is read, and the registries are updated with the signature information.

**Cause:** There may be hardware errors or transport problems.

**Solution:** Make sure that there are no hardware errors or transport problems.

Using Event Viewer (on the **Start** menu, under **Programs** and **Administrative Tools (Common)**), look in the event log for disk I/O error messages or indications of problems with the communications transport.

**Cause:** You may not have waited long enough for the registries to be updated.

**Solution:** Make sure that you waited long enough for the registries to be updated.

Cluster Administrator takes a backup of the registry when it starts up. However, it can take up to a minute after the second server restarts for the disk signatures to be written to the registries. Wait a minute, and then click **Refresh**.

**Cause:** One or more adapters on the shared storage bus are configured incorrectly.

**Solution:** Make sure that the adapters are configured correctly.

**Cause:** The shared storage bus exceeds the maximum cable length.

**Solution:** Make sure that the shared storage bus does not exceed the maximum cable length.

**Cause:** The disk is not supported.

**Solution:** Make sure that the disk hardware or firmware revision level is not outdated.

**Cause:** The bus adapter is not supported, or the adapter hardware or firmware revision level is outdated.

**Solution:** Make sure that the bus adapter is supported, and that the adapter hardware or firmware revision level is current.

**Cause:** If you move your storage bus adapter to another I/O slot, add or remove bus adapters, or install a new version of the bus adapter driver, the cluster software may not be able to access disks on your shared storage bus

**Solution:** To accommodate these changes, make sure that your shared storage bus adapter has been properly reconfigured.

**Cause:** The operating system is incorrectly configured to access the shared storage bus.

**Solution:** Verify that the operating system can detect the shared storage bus adapter.

**In Disk Management, you do not see the disk for the group that is online on that node.**
**Where?**

- Computer Management/Storage/Disk Management

**Cause:** You may not be looking at the right disks.

**Solution:** Make sure that you are looking at the right disks.

If you have not labeled your disks or assigned fixed drive letters to them, you may not recognize which disks are part of the cluster and which ones are not. Label your disks in a meaningful manner and assign fixed drive letters to all partitions.

**Cause:** There may have been hardware problems.

**Solution:** Make sure that there have not been any hardware problems.

Run Event Viewer and check for disk I/O error messages or indications of hardware problems.

**You are unable to manually move a group, or it does not fail over to another node when it is supposed to.**
**Cause:** The fail over node may not be designated as a possible owner for all resources in the group that you want to fail over.

**Solution:** Make sure that the fail over node is designated as a possible owner for all resources in the group you want to fail over.

Check the ownership configuration in the group resource **Properties** dialog box. If the node is not set as a possible owner for all resources in the group, the node cannot own the group, so failover will not occur. To fix this, make the node a possible owner for all resources in the group.

**Cause:** A resource in the group may be continually failing.

**Solution:** Determine if a resource in the group is continually failing.

If the node can, it will bring the resource back up without failing over the group. If the resource continually fails but does not fail over, make sure that the resource property **Restart and affect the group** is selected. Also, check the **Restart Threshold** and **Restart Period** settings, which are also in the resource **Properties** dialog box.

**A group failed over but did not fail back.**
**Cause:** The group will only fail back if the node the group was running on itself failed and then rejoined the cluster. If the group, but not the node, failed, then the group will fail over to another node, but will not fail back to the original node.

**Cause:** The failback policies of both the group and the resources may not be properly configured.

**Solution:** Make sure that the **Prevent failback** check box is clear in the group **Properties** dialog box. If the **Allow failback** check box is selected, be sure to wait long enough for the group to fail back. Check these settings for all affected resources within a group. Because groups fail over as a whole, one resource that is prevented from failing back affects the entire group.

**Cause:** The node to which you want the group to fail back is not configured as the preferred owner of the group.

**Solution:** Make sure that the node to which you want the group to fail back is configured as the preferred owner of the group. If not, the Cluster service leaves the group on the node to which they failed over.

**The entire group failed and has not restarted.**
**Cause:** A node is offline.

**Solution:** Make sure that the node is not offline.

If the node on which the group had been running is offline, check that another node is a possible owner of the group and of every resource in the group.

**Cause:** The group has failed repeatedly.

**Solution:** The group may have exceeded its failover threshold or its failover period. Try to bring the resources online individually (following the correct sequence of dependencies) to determine which resource is causing the problem. Or, create a temporary resource group (for testing purposes) and move the resources to it, one at a time.

**All nodes are functioning, but resources fail back repeatedly.**
**Cause:** Power may be intermittent or failing.

**Solution:** Ensure that your power is not intermittent or failing. You can correct this by using an uninterruptable power supply (UPS) or, if possible, by changing power companies.

**The Cluster service does not successfully fail over resources.**
**Cause:** Cluster storage device is not properly configured.

**Solution:** Verify that the cluster storage device is properly configured and that all cables are properly connected.

**You fail over a resource group from one node to another, but it automatically fails back.**
**Cause:** One or more resources fail to come online on the new node.

**Solution:** Use a process of elimination to determine which resource is failing to come online. For more information, see article Q303431, "Explanation of Why Server Clusters Do Not Verify that Resources will Work Properly on All Nodes" in the Microsoft Knowledge Base.

**The Network Name resource fails when you change to a system locale that is different than the input language used by the Network Name resource.**
**Cause:** The system locale must be the same on all nodes of a cluster and on the computer used to connect to the cluster.

**Solution:** Change the system locale. For more information, see Connect to a cluster with Cluster Administrator.

**The Message Queuing resource fails to handle message activity correctly which may result in resource failures.**

**Cause:** Each instance of Message Queuing on a server maps 4 MB of the system view space when handling message activity. This results in a default limit of three active, working instances of Message Queuing on a cluster node. In a server cluster with three Message Queuing resources, a node could have four concurrent Message Queuing services running (the service running on the local node plus the three services associated with the Message Queuing resources.) In this scenario, message activity could be limited, resulting in resource failures.

**Solution:** Increase the system view space memory pool on each node of a server cluster with three or more Message Queuing resources. (We also recommend that you increase the system view space memory pool even for nodes running fewer than three Message Queuing resources.)

- Open Registry Editor.
- Open the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management**.
- Create a new DWORD value called **SystemViewSize**.
- Calculate and enter the data for this DWORD value using the following formula: (16 + (the number of Message Queuing resources x 4)).

  For example, the calculation result for a cluster with three Message Queuing resources is 28.
- Reboot each node.

**A third-party resource fails to come online in a mixed-version cluster or while upgrading a cluster.**
**Cause:** If a resource uses a cryptographic provider not supplied by Microsoft to export (encrypt) and import (decrypt) resource data (cluster and cluster application cryptographic checkpoints), the default encryption key lengths may be different in the Windows 2000 and the Windows Server 2003 family operating systems. The result is that the resource might fail to come online and the cluster and event logs might contain cryptographic checkpoint synchronization errors for that resource.

**Solution:** Use the cluster.exe "CSP" private property to set the key length and effective key length for the third-party cryptographic provider that encrypts and decrypts data for the failing resource type.

- Open Command Prompt.
- Type **cluster***ClusterName***"***CSP***"**=*key_length,effective_key_length***:MULTISTR**

  *ClusterName* is the name of the cluster, *CSP* is the name of the cryptographic provider, and *key_length* and *effective_key_length* are the key and effective key lengths for the RC2 encryption algorithm, in bits. For more information on using cluster.exe, see Cluster.
- Depending on the resource, either bring the resource online or recreate the resource to add the new cryptographic checkpoint.

**Note**

- Review the documentation for your cryptographic provider to obtain valid values for the following RC2 encryption algorithm parameters: *key_length* and *effective_key_length*. Also review the cryptographic provider documentation for the correct procedure for adding the cryptographic checkpoint.

For information about how to obtain product support, see Technical support options.

# File and print server cluster problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**File and print server cluster problems**
What problem are you having?

- Clients cannot attach to a cluster file share resource.
- Clients cannot attach to a cluster printer share resource.
- After a resource failure, the clustered print server does not come back online.

**Clients cannot attach to a cluster file share resource.**
**Cause:** WINS or DNS might not be correctly configured.

**Solution:** Make sure that WINS or DNS are correctly configured.

If you are using WINS, run WINS Manager on the WINS server; then, on the **Mappings** menu, click **Show Database**. Make sure that each node and the cluster are registered in the WINS database and that the registrations are active.

For more information on WINS and WINS Manager, see the *Networking Guide* in the *Microsoft Windows Server 2003 Resource Kit*.

**Cause:** Your security policies might not allow the client to access the file share.

**Solution:** Make sure that your security policies allow the client to access the file share.

The client must have the right to log on to the file share, or the Guest account must be enabled for the client to have access.

**Clients cannot attach to a cluster printer share resource.**
**Cause:** Not all printer ports may be used in a server cluster.

**Solution:** Use cluster-aware ports, for example LPR or Standard TCP/IP ports.

The LPR and Standard TCP/IP Port, both provided with Windows Server 2003 family operating systems, can be used in a server cluster. Check your printer documentation to verify that your printer can use one of these two ports.

For more information, see Create a cluster-managed printer.

**After a resource failure, the clustered print server does not come back online.**
**Cause:** The share name for the printer may also used by another printer in the cluster.

**Solution:** Choose a unique share name for the printer.

We recommend that you do not use the same share name for different printers, even if they are in a different resource group and associated with a different Print Spooler resource. In the event of a failure, the same node in the cluster may own both resource groups. If this occurs, printers that share a common name will not be available.

For information about how to obtain product support, see Technical support options.

# General administrative problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**General administrative problems**
**What problem are you having?**


The Cluster service fails and the node cannot detect the network.

An IP address added to a group in the cluster fails.

An IP address resource is unresponsive when taken offline, for example you are unable to query its properties.

You receive the error: "RPC server is unavailable."

Cluster Administrator cannot open a connection to a node.

An application starts but cannot be closed.

A resource group has failed over but will not fail back.

All nodes appear to be functioning correctly, but you cannot access all of the drives from one node.

Cluster Administrator update delays.

Cluster Administrator stops responding when a node fails.

Cannot connect to cluster from recent file list.

Node performance is sluggish and node fails.

The cluster log contains numerous resource informational messages (for example, Entered LooksAlive, Entered Open, Entered Offline).

The Cluster service fails to start and returns an error code of ERROR_SHARING_VIOLATION (32) with event ID 1144 (NM_EVENT_REGISTER_NETWORK_FAILED).

You cannot manually restore the cluster database on a local node by copying the systemroot\cluster\CLUSDB file from another node.


**The Cluster service fails and the node cannot detect the network.**

In this case, you probably have a configuration problem. Check the following:

- **Cause:** Have you made any configuration changes recently?

  **Solution:** If the node was recently configured, or if you have installed some resource that required you to restart the computer, make sure that the node is still properly configured for the network.

- **Cause:** Is the node properly configured?

  **Solution:** Check that the server is properly configured for TCP/IP. Also check that the appropriate services are running. If the node recently failed, there is an instance of failover; but, if the other nodes are misconfigured as well, the failover will be inadequate and client access will fail.

**An IP address added to a group in the cluster fails.**

- **Cause:** The Internet protocol (IP) address is not unique.

  **Solution:** The IP address must be different from every other group IP address and every other IP address on the network.

- **Cause:** The IP address is not a static IP address.

**Solution:** The IP addresses must be statically assigned outside of a DHCP scope, or they must be reserved by the network administrator.

**An IP address resource is unresponsive when taken offline, for example you are unable to query its properties.**

- **Cause:** You may not have waited long enough for the resource to go offline.

  **Solution:** If an IP Address resource is unresponsive when taken offline, make sure that you wait long enough for the resource to go offline.

  Certain resources take time to go offline. For example, it can take up to three minutes for the IP Address resource to go fully offline.

**You receive the error: "RPC server is unavailable."**

- **Cause:** The server may not be operational, or the Cluster service and the RPC services may not be running.

  **Solution:** If you receive the error "RPC Server is unavailable," make sure the server is operational and that both the Cluster service and the RPC services are running. Also, check the name resolution of the cluster; it is possible that you are using the wrong name or that the name is not being properly resolved by WINS or DNS.

**Cluster Administrator cannot open a connection to a node.**

- **Cause:** The node may not be running.

  **Solution:** If Cluster Administrator cannot open a connection to a node, make sure that the node is running. If it is, confirm that both the Cluster service and the RPC services are running.

**An application starts but cannot be closed.**

- **Cause:** You may not have taken a resource offline using Cluster Administrator.

  **Solution:** When you bring resources online using Cluster Administrator, you must also take those resources offline using Cluster Administrator; do not attempt to close or exit the application from the application interface.

**A resource group has failed over but will not fail back.**

- **Cause:** The hardware and network configurations may not be valid.

  **Solution:** Make sure that the hardware and network configurations are valid.

  If any interconnect fails, failover can occur because the Cluster service does not detect a heartbeat, or it may not even register that the node was ever online. In this case, the Cluster service fails over the resources to the other nodes in the server cluster, but it cannot fail back because that node is still down.

- **Cause:** The resource group may not have been configured to fail back immediately, or you are not troubleshooting the problem within the allowable failback hours for the resource.

  **Solution:** Make sure that the resource group is configured to fail back immediately, or that you are troubleshooting the problem within the allowable failback hours for the resource group.

  A group can be configured to fail back only during specified hours. Often, administrators prevent failback during peak business hours. To check this, use Cluster Administrator to view the resource failback policy.

- **Cause:** You restarted the node to test the failover policy for the group instead of pressing the reset button.

**Solution:** Make sure that you press the reset button on the node. The resource group will not failback to the preferred node if you shutdown, then restart the node. For more information on testing failback policies, see Test node failure.

**All nodes appear to be functioning correctly, but you cannot access all of the drives from one node.**

- **Cause:** The shared drive may not be functioning.

  **Solution:** Confirm that the shared drive is still functioning.

  Try to access the drive from another node. If you can do that, check the cable from the device to the node that you cannot perform the access. If the cable is not the problem, restart the computer and then try again to access the device. If you cannot access the drive, check your configuration.

- **Cause:** The drive has completely failed.

  **Solution:** Determine (from another node) whether the drive is functioning at all. You may have to restart the drive (by restarting the computer) or replace the drive.

  The hard disk with the resource or a dependency for the resource may have failed. You may have to replace a hard disk. You may also have to reinstall the cluster.

**Cluster Administrator update delays.**

- **Cause:** If you run Cluster Administrator from a remote computer, it may not display the correct (updated) cluster state when the cluster network name fails over from one node to another node. This can result in Cluster Administrator displaying a node as being online, when it is actually offline.

  **Solution:** To work around this problem, restart Cluster Administrator.

  You can avoid this problem by connecting to clusters through node names. However, if the node you are connected to fails, Cluster Administrator stops responding until the RPC connection times out.

**Cluster Administrator stops responding when a node fails.**

- **Cause:** The Cluster Administrator may be slow in doing dynamic updates.

  **Solution:** If Cluster Administrator stops responding when a node fails, make sure that Cluster Administrator is not just slow in doing dynamic updates. If the Cluster service is running on a remaining node, Cluster Administrator is either not responding or is updating very slowly. There are two ways to see if the Cluster service is running on a remaining node:

- Use the TCP/IP Ping utility to ping the cluster name on a remaining node.

- In **Control Panel**, double-click **Services**, and check whether the Cluster service is running.

**Cannot connect to cluster from recent file list.**

- **Cause:** Files listed in the Cluster Administrator recent file list (both on the **File** menu and in the **Open Connection to Cluster** dialog box) have the cluster name appended to the path. For example, instead of Webclust1, the recent file list may list C:\Windows\Cluster\Webclust1. This problem occurs when Microsoft Visual C++ version 5.0 is installed.

  **Solution:** To work around this problem, manually type the cluster name when you open the connection.

**Node performance is sluggish and node fails.**

- **Cause:** CPU may be overloaded.

  **Solution:** Check that your node is not processor-bound. That is, that the CPU is not running at 100-percent utilization. If

you try to run too many resources for the node capacity, you can overload the CPU.

Also, review the size of your paging file. If the paging file is too small, the Cluster service can detect this as a node failure and fail over the groups.

**The cluster log contains numerous resource informational messages (for example, Entered LooksAlive, Entered Open, Entered Offline).**

- **Cause:** One or more of your Generic Script resources fills the cluster log with multiple copies of Entered LooksAlive, Entered Open, Entered Offline messages.

  **Solution:** When creating a script for a Generic Script resource, do not use the LogInformation method when calling the LooksAlive function. For more information, see the Microsoft Platform Software Development Kit (SDK).

**The Cluster service fails to start and returns an error code of ERROR_SHARING_VIOLATION (32) with event ID 1144 (NM_EVENT_REGISTER_NETWORK_FAILED).**

- **Cause:** The Internet Assigned Numbers Authority (IANA)-assigned port (3343) used by the cluster network driver (ClusNet) is bound to another process, preventing the Cluster service from starting.

  **Solution:** Use port scanning and process termination utilities to identify and end the process that is bound to port 3343.

To do this:

1. Open Command Prompt.

2. Navigate to the %systemroot%\system32 directory.

3. Type **netstat -a -o**.

   This will display all listening and connected ports and the process ID of each process bound to that port. Port 3343 will appear for each cluster network on the node.

   **Notes**

   - The **-a** option indicates that all connections and listening ports are supposed to be displayed. Server clusters uses UDP so the ports are normally in listening mode rather than in connections.

   - The **-o** option indicates that the owning process ID is supposed to be displayed.

4. Type **tasklist**.

   This will display the IDs for all the processes running on the node, including the process ID that matches the Cluster service (ClusSvc.exe).

5. Type **taskkill /pid** ID to terminate the process(es) bound to port 3343 that do not match the ID for the Cluster service.

**You cannot manually restore the cluster database on a local node by copying the systemroot\cluster\CLUSDB file from another node.**

- **Cause:** If the cluster registry hive is already locked and loaded by the Cluster service, the operating system will prevent you from copying a local CLUSDB file or overwriting an existing CLUSDB file on another node.

  **Solution:** Stop the Cluster service. Then unload the **HKEY_LOCAL_MACHINE\Cluster** hive before restoring the cluster database file.

To do this:

1. Open **Command Prompt**.

2. Type **net stop clussvc** to stop the Cluster service.

3. Use the Registry Editor to unload the hive under **HKEY_LOCAL_MACHINE\Cluster**. For more information, see Unload a hive from the registry.

   The operating system will now allow you to copy the CLUSDB file from a node and manually restore it to another node.

   For information about how to obtain product support, see Technical support options.

# Enterprise Memory Architecture

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enterprise Memory Architecture**
Enterprise Memory Architecture (EMA) increases the amount of physical memory available to applications. EMA supports two types of memory enhancement: application memory tuning--also known as 4-gigabyte tuning (4GT)--and Physical Address Extension (PAE) X86. Neither is enabled by default. Both application memory tuning and PAE must be manually enabled by editing the Boot.ini file.

- Before enabling application memory tuning, see Checklist: Enabling application memory tuning.
- Before enabling PAE, see Checklist: Enabling Physical Address Extension (PAE).
- For help with specific tasks, see Enterprise Memory Architecture How To....
- For general background information, see Enterprise Memory Architecture Concepts.
- For problem-solving instructions, see Enterprise Memory Architecture Troubleshooting.

### Notes

- This feature is not required on Microsoft® Windows® XP 64-bit Edition (Itanium) or the 64-bit versions of the Microsoft® Windows Server 2003 family.
- This topic does not apply to Windows Server 2003, Web Edition.

# Checklist: Enabling application memory tuning

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Enabling application memory tuning**

| Step | Reference |
|---|---|
| Review the concepts behind application memory tuning. | Application memory tuning |
| Read Enterprise Memory Architecture hardware requirements. | Enterprise Memory Architecture hardware and software requirements |
| Determine whether your computer supports application memory tuning by consulting the hardware compatibility information for products in the Windows Server 2003 family, which you can find by clicking the appropriate link in Support resources. Also, contact your vendor for additional compatibility information. | See the hardware compatibility information in Support resources.<br><br>Every computer listed has been tested with the maximum amount of memory the computer supports. To ascertain this amount, contact the appropriate vendor. If the amount is 2 gigabytes (GB) or more, and if the computer meets all other requirements, you can use application memory tuning with the computer. |
| If your hardware is unsupported, do one of the following:<br>• Update your hardware.<br>If you update your hardware, first back up your system.<br>• Or select different hardware on which to enable application memory tuning. | If necessary, contact the hardware vendor about updating your hardware.<br><br>Consult your user documentation for information on backing up your system. |
| Enable application memory tuning. | Enable application memory tuning |

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Checklist: Enabling Physical Address Extension (PAE)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Enabling Physical Address Extension (PAE) X86**

| Step | Reference |
|---|---|
| Review the concepts behind Physical Address Extension. | Physical Address Extension (PAE) overview;<br><br>Physical Address Extension (PAE) and application memory tuning |
| Read Enterprise Memory Architecture hardware requirements. | Enterprise Memory Architecture hardware and software requirements |
| Determine whether your hardware supports PAE by consulting the hardware compatibility information for products in the Windows Server 2003 family, which you can find by clicking the appropriate link in Support resources. Also, contact your vendor for additional compatibility information. | See the hardware compatibility information in Support resources.<br><br>Every computer listed has been tested with the maximum amount of memory the computer supports. To ascertain this amount, contact the appropriate vendor. If the amount is more than 4 gigabytes (GB), and if the computer meets all other requirements, PAE can be used with the computer. |
| Determine whether your network and storage adapters support PAE by consulting the hardware compatibility information, which you can find by clicking the appropriate link in Support resources. | See the hardware compatibility information in Support resources. |
| If your hardware is unsupported, do one of the following:<br><br>• Update your hardware.<br><br>If you update your hardware, first back up your system.<br><br>• Or select different hardware on which to enable application memory tuning. | If necessary, contact the hardware vendor about updating your hardware.<br><br>Consult your user documentation for information on backing up your system. |
| Enable Physical Address Extension (PAE) X86. | Enable Physical Address Extension (PAE) |

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Enterprise Memory Architecture How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

- Enable application memory tuning
- Enable Physical Address Extension (PAE)
- Configure the memory dump settings for PAE

**Notes**

- This feature is not needed on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.
- This topic does not apply to Windows Server 2003, Web Edition.

# Enable application memory tuning

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To enable application memory tuning

1. Open Windows Explorer.
2. On the **Tools** menu, click **Folder Options**.
3. On the **View** tab, click **Show hidden files and folders**, clear the **Hide protected operating system files** check box, and then click **OK**. If you are presented with a warning dialog box, click **Yes** to continue.
4. In the root folder (for example, C:), locate the Boot.ini file and remove its read-only attribute.
5. Open the Boot.ini file, and then add the **3GB** parameter to the ARC path, as shown in the following example for Windows Server 2003, Datacenter Edition:

   ```
   multi(0)disk(0)rdisk(0)partition(2)\%systemroot%="Windows Server 2003, Datacenter Edition" /3GB
   ```

6. On the **File** menu, click **Save**.
7. Restore the read-only attribute to the Boot.ini file.
8. For the change to take effect, restart the computer.

### Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- This topic does not apply to Windows Server 2003, Web Edition.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Enable Physical Address Extension (PAE)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Physical Address Extension (PAE) X86**

1. Open Windows Explorer.
2. On the **Tools** menu, click **Folder Options**.
3. On the **View** tab, click **Show hidden files and folders**, clear the **Hide protected operating system files** check box, and then click **OK**. If you are presented with a warning dialog box, click **Yes** to continue.
4. In the root folder (for example, C:), locate the Boot.ini file and remove its read-only attribute.
5. Open the Boot.ini file, and then add the **/PAE** parameter to the ARC path, as shown in the following example for Windows Server 2003, Datacenter Edition:

   ```
   multi(0)disk(0)rdisk(0)partition(2)\%systemroot%="Windows Server 2003, Datacenter Edition" /PAE
   ```

6. On the **File** menu, click **Save**.
7. Restore the read-only attribute to the Boot.ini file.
8. For the change to take effect, restart the computer.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open Windows Explorer, click **Start**, point to **All programs**, point to **Accessories**, and then click **Windows Explorer**.
- This topic does not apply to Windows Server 2003, Web Edition.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Configure the memory dump settings for PAE

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure the memory dump settings for PAE**

1. Open System Properties.
2. On the **Advanced** tab, click **Settings** in the **Startup and Recovery** box.
3. Under **Write debugging information**, select **Kernel memory dump**.
4. Under **Dump File**, you can accept the default, **%systemroot%\MEMORY.DMP**, or type the path and name of a memory dump file of your own choosing.

## Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.
- To open System properties, click **Start**, click **Control Panel**, and then double-click **System**.
- This topic does not apply to Windows Server 2003, Web Edition.
- The drive specified in step 4 needs to have at least 2,060 megabytes (MB) of free disk space if your system contains 4 gigabytes (GB) or more of physical memory and you specify **Kernel memory dump** under **Write debugging information**.
- Make sure the paging file size is at least 2,060 MB if your system contains 4 GB or more of physical memory and you specify **Kernel memory dump** under **Write debugging information**. For instructions on changing your paging file size, see Related Topics.
- For more information about memory dump file creation, see "Resources" in Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enterprise Memory Architecture Resources
Change the size of the virtual memory paging file

# Enterprise Memory Architecture Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section covers:

- Enterprise Memory Architecture overview
- Understanding Enterprise Memory Architecture
- Using Enterprise Memory Architecture
- Enterprise Memory Architecture Resources

**Notes**

- This feature is not required on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.
- This topic does not apply to Windows Server 2003, Web Edition.

# Enterprise Memory Architecture overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enterprise Memory Architecture overview**
Enterprise Memory Architecture (EMA) supports two types of memory enhancement: application memory tuning, also known as 4-gigabyte tuning (4GT), and Physical Address Extension (PAE) X86.

If you are new to Enterprise Memory Architecture, see:

- Application memory tuning
- Changing the default paging file size on large-memory computers

For additional information, see:

- Enable application memory tuning
- Enabling application memory tuning support in your applications
- Enable Physical Address Extension (PAE)
- Enterprise Memory Architecture hardware and software requirements

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Understanding Enterprise Memory Architecture

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Enterprise Memory Architecture**
This section covers:

- Physical Address Extension (PAE) and application memory tuning
- Application memory tuning
- Physical Address Extension (PAE) overview

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Physical Address Extension (PAE) and application memory tuning

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Physical Address Extension (PAE) X86 and application memory tuning**
Physical Address Extension (PAE) X86 and application memory tuning, also known as 4-gigabyte tuning (4GT), are complementary technologies for computers that use Intel processors.

Application memory tuning provides more of the computer's virtual memory to applications by providing less virtual memory to the operating system. PAE X86 allows software using the Address Windowing Extensions (AWE) API set running on a computer with an Intel Pentium Pro processor or later and either Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, to map more physical memory into the application's virtual address space.

For example, if you have a computer with a Pentium Pro processor or later and 16 gigabytes (GB) of physical memory (the maximum amount supported by application memory tuning), you can use application memory tuning to allocate 3 GB of virtual memory to applications (leaving 1 GB for the operating system). A given application can then lock the physical memory. Using the AWE API set and the PAE X86 kernel, the remaining 12 GB of physical memory is available to the application for caching purposes. Consequently, the application can access up to 15 GB of physical memory.

For more information, see:

- Application memory tuning
- Physical Address Extension (PAE) overview

**Important**

- On computers with more than 16 GB of physical memory using application memory tuning, physical memory above the 16 GB address is ignored by the operating system. For more information, see Enable application memory tuning.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Application memory tuning

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application memory tuning**
With application memory tuning, also known as 4-gigabyte tuning (4GT), memory-intensive applications running on Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, can use up to 50 percent more virtual memory. This capability is intended for applications that run on powerful computers with 2 to 4 gigabytes (GB) of physical RAM and that can take advantage of a larger virtual address space. The application must also be configured to use application memory tuning.

Application memory tuning makes more of the computer's virtual memory available to applications by making less virtual memory available to the operating system. Applications can use this feature to address 3 GB of virtual memory instead of the 2 GB normally allocated to each application process.

For more information, see Enabling application memory tuning support in your applications.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Physical Address Extension (PAE) overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Physical Address Extension (PAE) X86 overview**

Physical Address Extension (PAE) X86 allows software using the Address Windowing Extensions (AWE) API set and running on a computer with an Intel Pentium Pro processor or later, and more than 4 gigabytes (GB) of physical memory to map more physical memory into the application's virtual address space.

Applications not using the AWE API set can also benefit from PAE X86 because the operating system uses the larger physical memory to reduce paging and thus increase performance. This can also benefit consolidation servers hosting multiple applications.

Applications that manipulate large amounts of data achieve better performance by keeping data in memory instead of on disk. For example, PAE X86 can significantly improve the performance of the following types of applications:

- Databases, such as Microsoft SQL/E 7.0 or later.
- Scientific and engineering applications, such as computational fluid dynamics.
- Statistical analysis applications that do extensive data mining.

**Updates to PAE X86**

The following updates have been made to support the addition of Data Execution Prevention (DEP), which is also known as no-execute page protection:

- PAE is automatically enabled on computers running Windows Server 2003 with Service Pack 1 (SP1) and Windows XP with Service Pack 2 (SP2) when DEP is enabled on a computer with a processor that supports the no-execute page protection feature.
- When PAE mode is enabled on Windows Server 2003, Standard Edition with SP1 and Windows XP with SP2, physical address space is limited to 4 GB. Limiting physical address space to 4 GB helps prevent driver compatibility issues with PAE mode.

Only specific hardware will support PAE X86, therefore, this feature is not enabled upon initial installation of the operating system. For more information on what hardware supports PAE X86, refer to the hardware compatibility information for products in the Windows Server 2003 family, which you can find by clicking the appropriate link in Support resources. For more information about how to enable PAE X86, see Enable Physical Address Extension (PAE).

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Using Enterprise Memory Architecture

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Enterprise Memory Architecture**
This section covers:

- Enterprise Memory Architecture hardware and software requirements
- Enabling application memory tuning support in your applications
- Changing the default paging file size on large-memory computers

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Enterprise Memory Architecture hardware and software requirements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enterprise Memory Architecture hardware and software requirements**

The table below shows the hardware and software requirements for the two types of memory enhancement that EMA supports: application memory tuning and Physical Address Extension (PAE) X86.

| Technology | Hardware and software requirements |
|---|---|
| Application memory tuning, also known as 4-gigabyte tuning (4GT) | • X86-based processor<br>• 2 gigabytes (GB) or more of random access memory (RAM) |
| Physical Address Extension (PAE) X86 | • Pentium Pro processor or later<br>• More than 4 GB of RAM<br>• 450 NX or compatible chip set and support, or later<br>• The Hardware Abstraction Layer (HAL) must support PAE mode and the proper access to RAM.<br>• Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition |

Contact your hardware vendor for additional information.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Enabling application memory tuning support in your applications

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling application memory tuning support in your applications**
The changes to support application memory tuning, also known as 4-gigabyte tuning (4GT), are done at both the system and application levels. For background information, see Application memory tuning.

**System changes**
After you have installed Windows Server 2003, Datacenter Edition, you must modify the Boot.ini file to enable application memory tuning. To do this, add the **/3GB** parameter to the ARC path, as shown in the following example for Windows Server 2003, Datacenter Edition:

```
multi(0)disk(0)rdisk(0)partition(2)\%systemroot%="Windows Server 2003, Datacenter Edition" /3GB
```

For more information, see Enable application memory tuning.

**Application changes**
No new application programming interfaces (APIs) are required for application memory tuning support. However, it would be ineffective to automatically provide every application with a 3-gigabyte (GB) address space, therefore this feature is not enabled upon initial installation of the operating system.

Executable files that can use the 3-GB address space are required to have the bit IMAGE_FILE_LARGE_ADDRESS_AWARE set in their image header. Some manufacturers preconfigure their applications to use application memory tuning, making it unnecessary for you to make any changes. For more information, see your application documentation and contact your application vendor to determine if they support Large Address Awareness (which allows them to address 3 GB of virtual memory instead of the 2 GB normally allocated to each application process), or whether you can enable Large Address Awareness in their application.

### Notes

- When modifying executable files to make use of application memory tuning, be aware that it is not necessary to modify the associated DLLs.
- This topic does not apply to Windows Server 2003, Web Edition.

# Changing the default paging file size on large-memory computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Changing the default paging file size on large-memory computers**

If you have 4 gigabytes (GB) or more of memory on your server, you may want to change the default size of your paging file. On x86-based computers with 4 GB or more of physical memory there should be a paging file sufficiently large to capture a kernel memory dump. This is true even if there is little use of the paging file due to the majority of the memory being used by an Address Windowing Extensions (AWE) application.

When setting up the paging file on a large memory system, make the minimum and maximum page file sizes the same (for example, set **Initial size (MB)** to 2,050 MB and set **Maximum size (MB)** to 2,050 MB). This prevents the page file from getting fragmented on the hard drive, which might cause problems when writing a kernel memory dump to disk. A paging file size of at least 2,050 MB is the minimum size necessary for computers with 4 GB or more of physical memory to capture a complete kernel memory dump. The paging file can be much larger (up to three times the physical RAM is typical) if the system is being used for application consolidation. For instructions on changing the paging file size, see Change the size of the virtual memory paging file.

In computers with eight or more processors that have the maximum amount of RAM installed, you might be able to gain some incremental performance improvement by splitting the paging file into multiple paging files. Each paging file must be on a separate physical disk and for reliability, each disk should be part of a hardware RAID-5 volume.

For more information about dividing the paging file into multiple paging files, see "Troubleshooting Performance" at the Microsoft Windows Resource Kits Web site.

**Note**

- This topic does not apply to Windows Server 2003, Web Edition.

# Enterprise Memory Architecture Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resources**

Microsoft Product Support Services provides a searchable database of technical support information for Microsoft products. You can search here for general information about Enterprise Memory Architecture or for information related to other Microsoft products. You can get to the Microsoft Product Support Services Web site by clicking the appropriate link in Support resources.

### Note

- This topic does not apply to Windows Server 2003, Web Edition.

# Enterprise Memory Architecture Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting**
What problem are you having?

- The computer will not start after PAE is enabled.
- After PAE is enabled, the computer runs for a time and then displays a Stop error.
- System Properties, Microsoft System Diagnostics (WinMSD) or other system utilities report 256MB less than the total physical memory installed to the system.
- The writing of memory dumps stops or stalls when the Physical Addressing Extension (PAE) option is enabled, or a written memory dump is corrupt or incomplete.
- You cannot enable application memory tuning (4-gigabyte tuning) and Physical Address Extension (PAE) X86 on 64-bit versions of the Windows Server 2003 operating systems running on Itanium architecture-based or x64-based computers.

**The computer will not start after PAE is enabled.**
**Cause:** Your hardware may not support PAE.

**Solution:** When starting, run in Safe Mode, which disables PAE, and then remove the **/PAE** parameter from the Boot.ini file.

Also, determine whether your hardware supports PAE by consulting the hardware compatibility information for products in the Windows Server 2003 family, which you can find by clicking the appropriate link in Support resources. Contact your vendor for additional compatibility information.

You are also advised to contact your hardware vendor.

For more information, see Checklist: Enabling Physical Address Extension (PAE), and Enterprise Memory Architecture hardware and software requirements.

**After PAE is enabled, the computer runs for a time and then displays a Stop error.**
**Cause:** Your hardware may not support PAE.

**Solution:** Determine whether your hardware supports PAE by consulting the hardware compatibility information for products in the Windows Server 2003 family, which you can find by clicking the appropriate link in Support resources. Also contact your vendor for additional compatibility information.

If your hardware does not support PAE, or if the problem persists and you continue to believe that PAE might be the cause, remove the **/PAE** parameter from the Boot.ini file. To remove the **/PAE** parameter, restart your computer in Safe Mode, which disables PAE, and then remove the parameter from the Boot.ini file. You start your computer in Safe Mode by pressing F8, when indicated, during startup.

It is also recommended that you contact your hardware vendor.

For more information, see Checklist: Enabling Physical Address Extension (PAE), and Enterprise Memory Architecture hardware and software requirements.

**System Properties, Microsoft System Diagnostics (WinMSD) or other system utilities report 256MB less than the total physical memory installed to the system.**
**Cause:** 256MB of physical memory will appear to be missing on servers using certain hardware. The system ROM may reserve the upper 256MB memory region for memory mapped I/O devices. The amount of reserved physical memory may increase to 512MB depending on the number of I/O devices installed to the system. For example, a system with 4GB of actual physical memory may appear to have 3.84 GB total physical memory.

**Solution:** No correction is necessary. However, to see all memory, add a **/PAE** switch to the system's **boot.ini** file. For more information, see Enable Physical Address Extension (PAE).

**The writing of memory dumps stops or stalls when the Physical Addressing Extension (PAE) option is enabled, or a written memory dump is corrupt or incomplete.**
**Cause:** You might have misconfigured your recovery settings, your paging file might not be large enough for the memory dump, or you might lack available disk space.

**Solution:** Ensure that you are writing kernel-only debug information. Also, make sure the paging file size is at least 2,060 MB if your system contains 4 GB or more of physical memory and you have clicked **Kernel memory dump** under **Write debugging information**. For more information, see Configure the memory dump settings for PAE.

**See Also:** Article Q299515, "Crashdump Does Not Work on a Computer with More Than 4 Gigabytes of Memory," in the Microsoft Knowledge Base.

**You cannot enable application memory tuning (4-gigabyte tuning) and Physical Address Extension (PAE) X86 on 64-bit versions of the Windows Server 2003 operating systems running on Itanium architecture-based or x64-based computers.**
**Cause:** Not applicable - these features are not available on 64-bit versions of the Windows Server 2003 operating systems running on Itanium architecture-based or x64-based computers.

**Solution:** Not applicable - these features are not available on 64-bit versions of the Windows Server 2003 operating systems running on Itanium architecture-based or x64-based computers.

For more information about how to obtain product support, see Support resources.

### Note

- This topic does not apply to Windows Server 2003, Web Edition.

# Technologies for Application and Development Support

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application and development support**
This section covers:

- The .NET Framework
- Message Queuing
- UDDI Services overview
- Using Component Services
- Using Data Sources (ODBC)
- Windows Management Instrumentation Control
- Using MSDN

# The .NET Framework

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**The .NET Framework**
The .NET Framework is a new computing platform designed to simplify application development in the highly distributed environment of the Internet. The .NET Framework has several main components: the common language runtime, the .NET Framework class library, and the runtime hosts. The common language runtime is the foundation of the .NET Framework. It manages code execution and provides core services such as memory management, thread management, and remote execution. The .NET Framework class library is a comprehensive, object-oriented collection of reusable classes that you can use to develop a wide range of applications, including ASP.NET applications and XML Web services. The .NET Framework also provides several runtime hosts, including Windows Forms and ASP.NET, which work directly with the runtime to provide a scalable, server-side environment for managed code. The .NET Framework is installed by default with Windows Server 2003 operating systems.

For more information about the .NET Framework, see .NET Framework at the Microsoft Web site.

### Notes

- This feature is not available on the Itanium-based versions of the Windows operating systems.

- For more information about using the .NET Framework, open a command prompt, and then type:

  **start %systemroot%\Microsoft.NET\Framework\<*LCID*>\Admin.chm**

  where <*LCID*> is the locale identifier. For example, if you are using an English language installation of the .NET Framework, type:

  **start %systemroot%\Microsoft.NET\Framework\1033\Admin.chm**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

# Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing**
Message Queuing (also known as MSMQ) is a messaging infrastructure and a development tool for creating distributed messaging applications for Microsoft® Windows® operating systems.

Message Queuing is of interest to both system administrators (for installing and managing the infrastructure) and software developers (for creating messaging applications). This documentation is written for system administrators and discusses management of the infrastructure provided by Message Queuing. This documentation does not discuss application development.

For the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed on it. The term *Message Queuing client* can refer to either an *independent client* or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed. Such computers are referred to as having independent client functionality or dependent client functionality. The following links provide detailed information:

- Before installing or upgrading Message Queuing, see Checklists.
- To find features that have been moved in Windows Server 2003 family operating systems, see New ways to do familiar Message Queuing tasks.
- For instructions on specific administration tasks, see Message Queuing How To....
- To learn more about Message Queuing features, planning, installing, and upgrading Message Queuing, and background information to help you understand Message Queuing, see Message Queuing Concepts.
- For troubleshooting help and links to other Message Queuing resources, see Troubleshooting and Additional Resources.

# Using Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Message Queuing**

Message Queuing (also known as MSMQ) is a messaging infrastructure and environment for building and running distributed, loosely coupled, messaging applications. As part of the application server role, Message Queuing provides features for developing and managing Web-based applications. Message Queuing applications running on source computers send messages to queues, and applications running on target computers retrieve messages from queues. Such applications communicate across the Internet and other heterogeneous networks, and with computers that might be temporarily offline.

Message Queuing can be installed in various configurations. Computers configured as Message Queuing independent clients run the Message Queuing service, host queues, send and receive messages, and do not require synchronous access to a Message Queuing server. Message Queuing servers provide additional functionality, including:

- Directory access for client computers running earlier versions of Message Queuing on Windows NT 4.0, Windows 98, or Windows 95 computers.
- Support for dependent clients, which are a limited Message Queuing client configuration that requires synchronous access to a server for all messaging functions.
- Routing capabilities that enable you to route messages within a site or between different sites when there is no direct connectivity between client computers.

When you install Message Queuing you choose the subcomponents for your required configuration. The Message Queuing component is not installed by default during Windows Setup. For more information about Message Queuing, see Message Queuing Help.

**To install Message Queuing with independent client functionality**

1. Open the Windows Components Wizard.
2. In the Windows Components Wizard, select the **Application Server** check box, and then click **Details**.
3. Select the **Message Queuing** check box.
4. If you want to change the default set of optional subcomponents that will be installed, click **Details**, and then:
    - If the computer belongs to a domain and you want it to operate in workgroup mode, clear the **Active Directory Integration** check box.

**Notes**

- To open the Windows Components Wizard, click **Start**, click **Control Panel**, double-click **Add or Remove programs**, and then click **Add/Remove Windows Components**.
- Use the following administrative tools to manage Message Queuing:
    - Computer Management, which provides access to local queue and message information. To open Computer Management, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Computer Management**.
    - Active Directory Users and Computers, and Active Directory Sites and Services, which provide access to information in Active Directory for routing configurations, computers, public queues, and other Message Queuing objects. To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click

**Administrative Tools**, and then double-click **Active Directory Users and Computers**. To open Active Directory Sites and Services, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

# Checklists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklists**

This section provides checklists for installing Message Queuing client and server software on new computers. For details, see the following links:

- Checklist: Installation for Message Queuing clients
- Checklist: Installation for Message Queuing servers

# Checklist: Installation for Message Queuing clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installation for Message Queuing clients**
Use this checklist when installing Message Queuing clients on Windows Server 2003 family computers that do not have previous versions of Message Queuing installed. Message Queuing is not installed by default during Windows Server 2003 family operating system setup.

| Step | Reference |
|---|---|
| Review key concepts and determine the client functionality needed. | Message Queuing features |
| Review the installation requirements. | Installation overview |
| Review the problems that may arise during installation and their solutions. | Troubleshooting installation problems |
| Log on with local administrative permissions (as a member of the local Administrators group). For best security practice, log off as administrator after installation is complete. | Installation permissions |
| For a dependent client, verify that a Message Queuing computer with at least independent client functionality is available, preferably in the same site in which you are installing. Any such computer operating in domain mode can serve as a supporting server. Following installation, Message Queuing computers are not configured to support dependent clients, and this feature must be enabled. Note that a computer in workgroup mode cannot act as a supporting server for a dependent client. | Enable a supporting server; Installing dependent clients |
| Run the Windows Components Wizard and select the subcomponents that provide the applicable client functionality. | Install independent clients; Install dependent clients |
| Verify that the Message Queuing service has started. | Start or stop services for Message Queuing |
| Send a test message. Note that this feature must be enabled following Message Queuing installation. | Enable route tracking and test messages; Send test messages |

**Notes**

- New Message Queuing 3.0 features, such as HTTP transport, Message Queuing Triggers, and sending messages to multiple destinations, are not available to dependent clients. Thus, distribution lists, multiple-element format names, and multicast messaging are not supported for dependent clients.

- Dependent clients will no longer be supported in future versions of Message Queuing.

# Checklist: Installation for Message Queuing servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Checklist: Installation for Message Queuing servers**

Use this checklist when installing Message Queuing servers on computers running on Windows Server 2003 family computers that do not have previous versions of Message Queuing installed. Message Queuing is not installed by default during Windows Server 2003 family operating system setup.

| Step | Reference |
|---|---|
| Review key concepts and determine the server functionality needed. | Message Queuing features |
| Review installation requirements. | Installation overview |
| Review the problems that may arise during installation and their solutions. | Troubleshooting installation problems |
| Log on with the appropriate administrative permissions. | Installation permissions |
| Before installing a Message Queuing server on a computer in a particular site, verify that there is at least one available domain controller in that site or in a nearby site. | Installation overview |
| For each Windows Server 2003 or Windows 2000 domain controller, verify that the computer is trusted for delegation. | Make a domain controller trusted for delegation |
| To reduce replication traffic, consider configuring at least one computer from the Windows Server 2003 or Windows 2000 domain controller in each site as a global catalog server. | Configure a domain controller as a global catalog server |
| If this server will be installed in domain mode and act as a supporting server for dependent clients, enable this feature. | Enable a supporting server |
| If this server will support Windows 2000 Message Queuing clients in the same site, verify that there is a domain controller running the Message Queuing directory service (Downlevel Client Support) or Message Queuing 2.0 in the site, or promote this computer to a domain controller and select the **Downlevel Client Support** subcomponent during installation. | Installing Message Queuing servers |
| If this server will support MSMQ 1.0 clients in the same site, verify that there is a domain controller running the Message Queuing directory service (Downlevel Client Support) or Message Queuing 2.0 and configured as a global catalog server from the domain of each such client in the site. Or, if applicable, promote this computer to a domain controller, configure it as a global catalog server, and select the **Downlevel Client Support** subcomponent during installation. | Installing Message Queuing servers |
| Run the Windows Components Wizard and select the subcomponents that provide the services needed. | Install Message Queuing servers |

| | |
|---|---|
| When prompted during the first or a subsequent installation of a Message Queuing server with **Downlevel Client Support** on a domain controller in your organization, weaken security if you want to support MSMQ 1.0 clients or clients logged on with local user accounts. For best security practice, only weaken security if necessary. | Enabling weakened security |
| If this server will be deployed for Internet messaging, configure HTTPS authentication. | HTTPS authentication |
| Verify that the Message Queuing service has started. | Start or stop services for Message Queuing |
| If this server has been enabled for routing, create a routing link for this computer to serve as a site gate to route messages to another site. | Create routing links |
| Send a test message. Note that this feature must be enabled following Message Queuing installation. | Enable route tracking and test messages; Send test messages |

# New ways to do familiar Message Queuing tasks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**New ways to do familiar tasks**

These tables compare different user interfaces used for common tasks in Message Queuing in:

- Windows Server 2003 family operating systems and Windows NT 4.0.
- Windows Server 2003 family operating systems and Windows 2000.

For a table that lists terminology and architectural differences between previous versions and the current version of Message Queuing, see Upgrade overview.

**Windows Server 2003 family operating systems and Windows NT 4.0**

| If you want to | In Windows NT 4.0 use | In Windows Server 2003 family |
|---|---|---|
| Set access control, set access auditing, or change ownership for Message Queuing | MSMQ Explorer | Open Active Directory Users and Computers. For more information, see Message Queuing security overview. |
| Change the Message Queuing service account | **Services** in Control Panel | Open Computer Management. For more information, see Change the Message Queuing service account. |
| Configure routing links | MSMQ Explorer | Open Active Directory Sites and Services. For more information, see Routing links. |
| For cross-platform messaging, create foreign sites or add foreign computers | MSMQ Explorer | Open Active Directory Sites and Services. For more information, see Foreign computers and sites. |
| Configure Message Queuing computers and properties, or set computer quotas or properties | MSMQ Explorer | Open Active Directory Users and Computers. For more information, see Administration overview. |
| Set DNS properties for Message Queuing computers | **Network** in Control Panel | Open Network Connections in Control Panel. For more information, see Set DNS properties for Message Queuing. |
| Create, delete, or configure queues, or set queue quotas or properties | MSMQ Explorer | Open Computer Management or open Active Directory Users and Computers. For more information, see Administering Queues. |
| View or delete messages, or view message properties | MSMQ Explorer | Open Computer Management or open Active Directory Users and Computers. For more information, see Managing messages. |

**Windows Server 2003 family operating systems and Windows 2000**

| If you want to | In Windows 2000 use | In Windows Server 2003 family |
|---|---|---|
| Specify a new site for mobile computers. | **Message Queuing** in Control Panel | Open Computer Management. For more information, see Specify a new site for mobile computers. |

| | | |
|---|---|---|
| Change the supporting server for a dependent client. Note that Message Queuing computers are not enabled by default to support dependent clients. | **Message Queuing** in Control Panel | Open Computer Management. For more information, see Enable a supporting server and Change the supporting server for a dependent client. |
| Specify the drives and folders for storing message files, message logs, and transaction logs. | **Message Queuing** in Control Panel | Open Computer Management. For more information, see Store messages on multiple disk drives. |
| Change, register, view, and remove certificates. | **Message Queuing** in Control Panel | Open Computer Management. For more information, see Authentication for Message Queuing. |
| To renew cryptographic keys for Message Queuing. | **Message Queuing** in Control Panel | Open Computer Management. For more information, see Renew cryptographic keys for Message Queuing. |

**Notes**

- Message Queuing 3.0 does not support a mixed-mode domain environment. That is, an environment where MSMQ 1.0 controller servers running on a Windows NT 4.0 computer exist in a Windows Server 2003 domain, together with Windows Server 2003 domain controllers. After you upgrade Windows NT 4.0 controller servers to the Windows Server 2003 family operating system, and follow the Message Queuing upgrade procedures, the result is a Windows Server 2003 domain environment, with no MQIS servers remaining, and no replication service.

- Until this process is completed, and as long as the domain environment is in mixed-mode, Message Queuing directory service functionality (Downlevel Client Support component) cannot be added to Windows Server 2003 domain controllers. In Message Queuing 3.0, directory service functionality is only required to provide directory services for Message Queuing objects to Message Queuing client computers running MSMQ 1.0 and 2.0. Message Queuing 3.0 servers and independent clients running on Windows Server 2003 family computers can access Active Directory directly.

- The Message Queuing directory service is not installed by default. To install it, you must select the **Downlevel Client Support** check box during Message Queuing 3.0 installation. You cannot add this functionality in a mixed-mode environment.

- A mixed-mode Windows 2000 domain environment comprised of MSMQ 1.0 controller servers and Windows 2000 domain controllers can still exist. In such an environment, Windows Server 2003 domain controllers can exist, but the Downlevel Client Support component cannot be installed on them.

- MSMQ Explorer does not run on Windows Server 2003 family or Windows 2000 computers.

- There are no tools that allow you to manage an MSMQ 1.0 enterprise from a Windows Server 2003 family or Windows 2000 computer.

# Message Queuing How To...

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**How to...**

This section provides instructions on completing Message Queuing administration tasks:

- For information on how to install and uninstall Message Queuing, configure independent and dependent clients, and configure Message Queuing in clusters, see Install Message Queuing.
- For instructions on managing routing links, and site gates, see Administer a Message Queuing Network.
- For information on Message Queuing in Active Directory, see Administer Message Queuing in Active Directory.
- For information on administering tasks on Message Queuing computers, see Administer Message Queuing Computers.
- To learn about security tasks, including managing Message Queuing object security, authentication, and encryption, see Administer Security for Message Queuing.
- For instructions on managing and configuring queues and queue aliases, see Administer Queues for Message Queuing.
- For information on distribution list tasks, see Administer Distribution Lists for Message Queuing.
- To learn about message properties, including setting default message lifetime, and purging messages, see Administer Messages for Message Queuing.
- For information on creating and managing triggers and rules, modifying default trigger values, and configuring Message Queuing trigger resources for clusters, see Administer Triggers for Message Queuing.
- For instructions on viewing Message Queuing events, external transactions, and performance counters, and to learn to enable message route tracking and test messages, see Monitor and Test Message Queuing.

# Install Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Install Message Queuing**

This section shows you how to:

- Configure a domain controller as a global catalog server
- Make a domain controller trusted for delegation
- Set installation permissions to enable upgrade from MSMQ 1.0
- Install Message Queuing servers
- Install independent clients
- Enable network DTC access
- Enable a supporting server
- Install dependent clients
- Configure Message Queuing resources for server clusters
- Configure Message Queuing with HTTP for server clusters
- Uninstall Message Queuing

# Configure a domain controller as a global catalog server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure a domain controller as a global catalog server**

1. Open Active Directory Sites and Services.
2. In the console tree, right-click **NTDS Settings**.

   **Where?**

   - Active Directory Sites and Services/Sites/*YourApplicableSite*/Servers/*YourApplicableServer*/NTDS Settings

3. Click **Properties**.
4. On the **General** page, select the **global catalog** check box.

**Notes**

- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
- Each Windows Server 2003 family site or Windows 2000 site containing MSMQ 1.0 clients running on Windows NT 4.0, Windows 98, or Windows 95 computers must contain at least one Windows Server 2003 domain controller that is running the Message Queuing directory service (Downlevel Client Support) and is configured as a global catalog server or at least one Windows 2000 domain controller that is running a Message Queuing server and is configured as a global catalog server from the domain of each such client. This is required to enable such clients to send queries programmatically to Message Queuing servers.
- When a Windows NT 4.0 primary enterprise controller (PEC) is migrated, it is automatically promoted to a domain controller configured as a global catalog server.
- The first domain controller in a Windows Server 2003 family forest is automatically configured as a global catalog server.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing on a domain controller
Working with MMC console files

# Make a domain controller trusted for delegation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To make a domain controller trusted for delegation**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click the applicable domain controller.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/Domain Controllers/*YourDomainController*

4. Click **Properties**.
5. On the **General** page, select the **Trust computer for delegation** check box (if needed).

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- If this property is disabled for a domain controller, the Message Queuing service cannot run.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing on a domain controller
Working with MMC console files

# Set installation permissions to enable upgrade from MSMQ 1.0

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set installation permissions to enable upgrade from MSMQ 1.0**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click **Computers** (or applicable organizational unit).

   **Where?**

   - Active Directory Users and Computers/*YourDomain/Computers* (or applicable organizational unit)

3. Click **Delegate control** to start the Delegation of Control Wizard, and then click **Next**.
4. On the **Users or Groups** page, click **Add**.
5. In the **Select Users, Computers, or Groups** dialog box, ensure that **Users, Groups, or Build-in Security Principals** appears in **Select this object type**. If it does not, click **Object Types**, and select these objects, clearing any other check boxes that are selected. Then click **OK**.
6. In **Enter the object name to select**, do the following:

   - To allow authenticated users permission to install Message Queuing, under **Name**, type **Authenticated Users**, and then click **OK**.
   - To allow everyone permission to install Message Queuing, under **Name**, type **Everyone**, and then click **OK**.

7. On the **Users or Groups** page, click **Next**.
8. Select **Create a custom task to delegate**, and then click **Next**.
9. On the **Active Directory Object Type** page, click **Only the following objects in the folder**, select the **Computer objects** check box, and then click **Next**.
10. On the **Permissions** page, under **Show these permissions**, clear the **General** check box, and then select the **Creation/deletion of specific child objects** check box.
11. Under **Permissions**, ensure that the **Create MSMQ Configuration objects** check box and the **Delete MSMQ Configuration objects** check box are selected, and that the other check boxes are cleared. Click **Next**, and then click **Finish**.

**Notes**

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Use the Delegation of Control Wizard to grant the proper permissions to users that want to install MSMQ 1.0 on a computer that is part of a domain that runs Windows Server 2003 family or Windows 2000.
- On the **Security** page of the **Properties** dialog box, the check boxes for permissions granted using the Delegation of

Control Wizard or using inherited permissions are dimmed.

- The Delete MSMQ Configuration Objects permission is required to uninstall Message Queuing.

- When Message Queuing is installed on computers running Windows 95, Windows 98, and some versions of Windows NT 4.0 (without a hotfix) or Windows 2000 (prior to Service Pack 4), or when these computers join a Windows Server 2003 domain, they require a given GUID when creating MSMQ configuration objects. By default a Windows Server 2003 forest does not support the addition of objects with a given GUID. To overcome this problem, it is recommended that you upgrade the client computer to an operating system that does not require creating an object with a given GUID. These include Windows NT 4.0 (with a hotfix), and Windows 2000 (with Service Pack 4).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installation permissions
Working with MMC console files

# Install Message Queuing servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install Message Queuing servers**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, select the **Application Server** check box, click **Details**, select the **Message Queuing** check box, and then click **Details**.
4. On the **Message Queuing** page, you can select the following subcomponents for your Message Queuing server:
   - To provide message routing services, select the **Routing Support** check box.
   - To install the Message Queuing directory service, which provides access to Message Queuing objects in Active Directory for MSMQ 1.0 clients running on Windows 98, Windows 95, Windows Millennium Edition and Windows NT 4.0 computers, and for MSMQ 2.0 clients running on Windows 2000 computers, select the **Downlevel Client Support** check box.
   - If the computer belongs to a domain and you want it to operate in workgroup mode, clear the **Active Directory Integration** check box. In this case, you cannot install Routing Support.
   - If you want to install Message Queuing Triggers, select the **Triggers** check box.
   - If you intend to send messages by HTTP transport, select the **MSMQ HTTP Support** check box.
5. Click **OK**, and then click **Next**.
6. If prompted, type the DNS or NetBIOS name of an MSMQ 1.0 PEC or PSC (in a Windows NT 4.0 enterprise) in the local site, and then click **Next**.
7. Follow the remaining instructions in the wizard.

**Notes**

- To open Add or Remove Programs, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Add or Remove Programs**.
- Message Queuing 3.0 servers can be installed only on computers running on platforms from the Windows Server 2003 family.
- Only Message Queuing 3.0 clients, not servers, can be installed in a Windows NT 4.0 enterprise.
- During setup, an attempt is made to automatically detect the nearest Windows Server 2003 or Windows 2000 domain controller in the local site or in a nearby site. If a suitable domain controller cannot be found, Setup assumes that you are installing a Message Queuing client in a Windows NT 4.0 environment and prompts you for the name of a controller server (PEC or PSC) before Setup can continue.
- The following subcomponents are selected by default for 32-bit servers: Active Directory Integration and Common (consisting of Core functionality and Local Storage).

- In a Windows Server 2003 site, to provide access to Message Queuing objects in Active Directory for Message Queuing 1.0 clients on Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0 computers, or Message Queuing 2.0 clients running on Windows 2000 computers, you must install a Message Queuing server with the **Downlevel Client Support** check box selected on a Windows Server 2003 domain controller in that site, or you must install a Message Queuing server on a Windows 2000 domain controller in that site. For MSMQ 1.0 clients, you must also configure the applicable domain controller as a global catalog server, and each MSMQ 1.0 client must have such a domain controller from its own domain in the site.

- Although the **Downlevel Client Support** component can be installed on all Message Queuing servers, it is relevant only for domain controllers. If it is installed on a computer that is not a domain controller, it only becomes active if the server is converted to a domain controller.

- You can install a Message Queuing server without Routing Support in workgroup mode if the computer is part of a Windows Server 2003 family workgroup. In this case, only private queues can be accessed and direct connectivity is required to send messages. For more information, see Related Topics.

- If you are installing a Message Queuing server with Downlevel Client Support on a domain controller and security has not been weakened in your organization to support MSMQ 1.0 clients or clients logged on with local user accounts, you will be prompted during installation about whether or not to weaken security. Weaken security only if you want to support such clients.

- Do not attempt to use Message Queuing 3.0 files to create a Message Queuing server on a computer that does not belong to the Windows Server 2003 family.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing Message Queuing servers
Choosing installation subcomponents

# Install independent clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install independent clients**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, select the **Application Server** check box, click **Details**, select the **Message Queuing** check box, and then click **Details**.
4. On the **Message Queuing** page, you can select the following subcomponents for the independent client computer:
   - If the computer belongs to a domain and you want it to operate in workgroup mode, clear the **Active Directory Integration** check box. In this case, you cannot install Routing Support.
   - If you want to install Message Queuing Triggers, select the **Triggers** check box.
   - If you intend to send messages by HTTP transport, select the **MSMQ HTTP Support** check box.
5. Click **Next**.
6. Follow the remaining instructions in the wizard.
7. If prompted to restart your computer, click **Yes** to allow Setup to complete.

**Notes**

- To open Add or Remove Programs, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Add or Remove Programs**.
- The following subcomponents are selected by default for 32-bit servers: Active Directory Integration and Common (consisting of Core functionality and Local Storage).
- If a domain controller cannot be found automatically in the local site or in a nearby site, the installation will continue, and Message Queuing will be installed in workgroup mode. In the context of Message Queuing, any mode of operation that does not allow access to a directory service is called workgroup mode, even if the applicable computer belongs to a domain.
- You can also install an independent client in workgroup mode if the computer is part of a Windows Server 2003 family workgroup. In this case, direct connectivity is required to send messages. For more information, see Related Topics.
- Do not attempt to use Message Queuing 3.0 files to create an independent client on a computer that does not belong to the Windows Server 2003 family.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Installing independent clients

Choosing installation subcomponents

# Enable network DTC access

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable network DTC access**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, select the **Application Server** check box, and then click **Details**.
4. Select **Enable network DTC access**, and then click **OK**.

### Note

- To open Add or Remove Programs, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Add or Remove Programs**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing dependent clients

# Enable a supporting server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable a supporting server**

1. Open Computer Management.
2. In the console tree, click **Services and Applications**.

   **Where?**

   - Computer Management/Services and Applications

3. In the details pane, right-click **Message Queuing**, and then click **Properties**.
4. On the **Server Security** tab, clear the selection **Disable dependent client support to prevent this computer acting as a supporting server for dependent clients**, and then click **OK**.

**Important**

- Before enabling a supporting server, you must enable network access for Microsoft Distributed Transaction Coordinator (DTC). For instructions, see Enable network DTC access

**Notes**

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- For best security practice, only turn on this feature if it is required.
- Following installation, dependent client support is disabled by default for Message Queuing computers.
- To reset server security settings on this tab to the default settings following installation, click **Restore Defaults**.
- Dependent clients cannot be installed on computers in workgroup mode or on 64-bit computers.
- New Windows Server 2003 family Message Queuing features, such as HTTP transport, Message Queuing Triggers, and sending messages to multiple destinations, are not available to dependent clients. Thus, distribution lists, multiple-element format names, and multicast messaging are not supported for dependent clients.
- Dependent clients will no longer be supported in future versions of Message Queuing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing independent clients
Installing dependent clients

# Install dependent clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To install dependent clients**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, select the **Application Server** check box, click **Details**, select the **Message Queuing** check box, and then click **Details**.
4. On the **Message Queuing** page, clear the **Active Directory Integration**, **Triggers**, and **MSMQ HTTP Support** check boxes, select the **Common** check box, and then click **Details**.
5. Clear the **Local Storage** check box, click **OK**, click **OK**, and then click **Next**.
6. If prompted, type the DNS or NetBIOS name of a supporting server in the local site. If the specified supporting server cannot be found:
   - Click **Yes** to complete setup. Note that the dependent client cannot send or receive messages until a connection with a Message Queuing server is established.
   - Click **No** to specify a different Message Queuing server.

**Notes**

- To open Add or Remove Programs, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Add or Remove Programs**.
- The following subcomponents are selected by default for 32-bit servers: Active Directory Integration and Common (consisting of Core functionality and Local Storage).
- By default, following installation, Message Queuing servers are not enabled to support dependent clients. For instructions on enabling a Message Queuing computer to act as a supporting server, see Enable a supporting server.
- Any Message Queuing server operating in domain mode with at least independent client functionality can serve as the supporting server for a dependent client. Setup automatically attempts to locate such a Message Queuing server in the same site or in a nearby site. If one cannot be located, Setup prompts you for the name of such a server. If the server specified is not available, Setup displays a message stating that the server is not reachable. You can choose to complete setup without a specified server, allowing you to install a dependent client before you install its supporting Message Queuing server. For best security practice, do not configure servers as supporting servers unless this functionality is required.
- Dependent clients cannot be installed on computers in workgroup mode or on 64-bit computers.
- New Windows Server 2003 family Message Queuing features, such as HTTP transport, Message Queuing Triggers, and sending messages to multiple destinations, are not available to dependent clients. Thus, distribution lists, multiple-element format names, and multicast messaging are not supported for dependent clients.
- Dependent clients will no longer be supported in future versions of Message Queuing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Installing dependent clients

# Configure Message Queuing resources for server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Message Queuing resources for server clusters**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.

2. To create a Message Queuing resource:

   - When the **Open Connection to Cluster** dialog box appears, in the **Cluster or server name** list box, select the applicable cluster or server. If you select a server, the name of the cluster to which it belongs will appear in the console tree (console tree location is: ClusterName/Groups/YourGroup).

   - In the console tree, right-click the cluster group to which you want the Message Queuing resource to belong.

   - On the **File** menu, point to **New**, and then click **Resource**.

   - In the New Resource Wizard, type the applicable information in **Name** and **Description**, select **Message Queuing** in the **Resource type** list box, and then click **Next**.

   - On the **Possible Owners** page, if you want to add nodes to the automatically chosen list under **Possible owners**, select them from the list under **Available nodes**, and then click **Add**. If you want to remove nodes from the automatically chosen list under **Possible owners**, select them in this list, click **Remove**, and then click **Next**.

   - On the **Dependencies** page, under **Available resources**, select the name of the Network Name resource in the group and click **Add**, and then select the name of the Physical Disk resource in the group and click **Add**. Then click **Finish**.

3. If Message Queuing has the Active Directory Integration component installed, enable Kerberos authentication for the Network Name resource on which the Message Queuing resources depends. This is not enabled by default. To enable, do the following:

   - To bring the Network Name resource offline, in the console tree for the cluster group, right-click the Network Name resource selected as a dependency in the details pane, and then click **Take Offline**.

   - Right-click the Network Name resource again, and then click **Properties**.

   - On the **Parameters** tab, select the **Enable Kerberos Authentication** check box, to enable authentication.

   - To bring the Network Name resource online, in the console tree for the cluster group, right-click the Network Name resources in the details pane, and then click **Bring Online**.

   **Caution**

   - Other resources in the group may also depend on the Network Name resource. These will be taken offline automatically when the Network Name resource goes offline. When you bring the Network Name resource online, bring these other resources online too.

4. To bring the Message Queuing resource online, in the console tree, click the **Resources** folder. In the details pane, click the Message Queuing resource. On the **File** menu, click **Bring Online**.

5. To manage the Message Queuing resource from Computer Management, create a Generic Application resource in the same cluster group and configure it to launch the Computer Management MMC snap-in:

    - In the Cluster Administrator console tree, click the cluster group to which the Message Queuing resource belongs.
    - On the **File** menu, point to **New**, and then click **Resource**. The New Resource Wizard launches.
    - In the New Resource Wizard, type **Computer Management** in the **Name** text box and select **Generic Application** in the **Resource type** list box. Then click **Next**.
    - On the **Possible Owners** page, select possible owners to match those selected for the Message Queuing resource. Then click **Next**.
    - On the **Dependencies** page, under **Available resources**, select the name of the Network Name resource in the group and click **Add**, and then select the Message Queuing resource previously created and click **Add**. Then click **Next**.
    - On the **Generic Application Parameters** page, in **Command line**, type **mmc compmgmt.msc**.
    - In **Current directory**, type **%windir%/system32**.
    - Select the **Allow application to interact with desktop** check box, select the **Use Network Name for computer name** check box, click **Next**, and then click **Finish**.
    - Bring the resource online using the procedure above. The Computer Management snap-in will appear on the desktop of the computer node that owns the resource.

**Notes**

- This task is performed on a computer that is running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition as part of a server cluster, and has the Message Queuing service already installed.
- If you want to use triggers, HTTP messaging, or external transactions in your cluster, before creating a Message Queuing resource in any cluster group hosted on a node of a server cluster, you must configure a Microsoft Distributed Transaction Coordinator (DTC) resource in the cluster. For instructions on creating a cluster, see Create a cluster, in Server Clusters Help.
- Errors occurring during the performance of this task are recorded in a Windows event log. For example, a "device not ready" error indicates that Message Queuing was not installed on the node.
- All Message Queuing applications that run in the context of a cluster group (whether cluster-aware or not) must depend on the Message Queuing resource in that group, and use the Network Name resource in that group. Repeat the steps stated above for the Computer Management MMC snap-in to properly configure Generic Application resources for "cluster-unaware" applications.
- Once created, do not rename resources for Message Queuing.
- If the virtual Message Queuing service managed by the Network Name resource will access Active Directory, you must enable Kerberos authentication on the Network Name resource on which the Message Queuing resource depends. Note that this feature is not enabled by default.
- You can also enable Kerberos authentication from the command line by typing the following command *cluster res "Network_Name_resource_name" /priv RequireKerberos=1:DWORD*
- When you select the Physical Disk resource in the group, Message Queuing allocates its storage in the *MSMQ\STORAGE*

folder on the shared disk. After storage has been allocated, you cannot modify the folder location.

- To uninstall Message Queuing on a computer that is part of a server cluster, the Message Queuing virtual server must be deleted from the node before uninstalling Message Queuing on the physical node. For example, if Message Queuing is installed on NodeA of a server cluster, you need to identify all Message Queuing resources on this cluster, ensure that NodeA does not own any cluster groups containing Message Queuing resources, and remove NodeA from the list of possible owners of each Message Queuing resource. Then you uninstall Message Queuing from NodeA.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Message Queuing servers
Administering server cluster nodes
Installing Message Queuing in a server cluster

# Configure Message Queuing with HTTP for server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure Message Queuing with HTTP for server clusters**

1. Open Computer Management.

2. In the console tree, click to expand **Services and Applications**.

   **Where?**

   - Computer Management/Services and Applications

3. Click **Internet Information Services (IIS) Manager**, click to expand **Web Sites**, right-click **Default Web Site**, and then click **Properties**.

4. To assign to the site all IP addresses that are not assigned to other sites, on the **Web Site** tab, in **Web site identification**, select **Default Web Site**, and in IP Address select **All Unassigned**.

5. To assign the explicit IP address of the virtual server as the default Web site, on the **Web Site** tab, in **Web Site Identification**, click **Advanced**, and then in **Advanced Multiple Web Site Configuration**, click **Add** to add the IP address of the virtual server.

### Note

- This task is performed on a computer that is running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition as part of a server cluster, and has the Message Queuing service already installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
HTTP/HTTPS messaging in server clusters
Message Queuing in server clusters
Administering server cluster nodes

# Uninstall Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To uninstall Message Queuing**

1. Open Add or Remove Programs in Control Panel.
2. Click **Add/Remove Windows Components**.
3. In the Windows Components Wizard, select the **Application Server** check box, click **Details**, clear the **Message Queuing** check box, and then click **Next**.
4. When asked if you want to permanently delete the queues and messages created by Message Queuing, click **Yes**.

## Notes

- To open Add or Remove Programs, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Add or Remove Programs**.
- When uninstalling Message Queuing in a multi-domain environment, wait for Active Directory replication to occur and check that the msmq object for the applicable computer was removed before attempting to reinstall Message Queuing on the same computer.
- Uninstalling Message Queuing does not remove Message Queuing files from the \Windows\System32 and \Windows\System32\Drivers folders.
- To uninstall Message Queuing on a computer that is part of a server cluster, the Message Queuing virtual server must be deleted from the node before uninstalling Message Queuing on the physical node.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Uninstalling Message Queuing

# Administer a Message Queuing Network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer a Message Queuing Network**
This section shows you how to:

- Create routing links
- Change the cost for a routing link
- Change the site gates for a routing link
- Administer Cross-Platform Messaging

# Create routing links

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create routing links**

1. Open Active Directory Sites and Services.
2. On the **View** menu, click **Show Services Node**.
3. In the console tree, right-click **MsmqServices**.

   **Where?**

   - Active Directory Sites and Services/Services/MsmqServices

4. Point to **New**, and then click **MSMQ Routing Link**.
5. On the **New Object - Routing Link** dialog box, in **Site 1** and in **Site 2**, click the arrows that open the lists and select the two sites between which you want to route messages.
6. In **Routing Link Cost**, enter a cost.

## Notes

- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
- A routing link defines the route for connectivity between two sites. For more information on routing links, how they are used to route messages, and routing-link costs, see Related Topics.
- Costs can range from 1 to 999999.
- By default, users in the Enterprise Admins group and Domain Admins group can create or modify routing links. For information on how to set permissions for routing links, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring routing links
Routing with Message Queuing servers
Message Queuing routing overview
Set permissions for routing link objects
Working with MMC console files

# Change the cost for a routing link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the cost for a routing link**

1. Open Active Directory Sites and Services.
2. On the **View** menu, click **Show Services Node**.
3. In the console tree, click **MsmqServices**.

   **Where?**

   - Active Directory Sites and Services/Services/MsmqServices

4. In the details pane, right-click the applicable routing link, and then click **Properties**.
5. On the **General** page, in **Routing Link Cost**, specify a new cost.

## Notes

- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
- A routing link defines the route for connectivity between two sites. For more information on routing links, how they are used to route messages, and routing-link costs, see Related Topics.
- All routing links are assigned a default cost of 1. Change routing-link costs only in the case where alternative routes are possible and you want to enforce one route over another route using link costs. In that case, you can assign costs that reflect the differences in the communications costs.
- Costs can range from 1 to 999999.
- By default, only users in the Enterprise Admins group can create or modify routing links if there are multiple domains in your organization. For users in the Domain Admins group to be able to create or modify routing links, they will need to be granted the Create All Child Objects permission for the MsmqServices object. For information on how to set permissions for routing links, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring routing links
Routing with Message Queuing servers
Message Queuing routing overview
Set permissions for routing link objects
Working with MMC console files

# Change the site gates for a routing link

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the site gates for a routing link**

1. Open Active Directory Sites and Services.
2. On the **View** menu, click **Show Services Node**.
3. In the console tree, click **MsmqServices**.

   **Where?**

   - Active Directory Sites and Services/Services/MsmqServices

4. In the details pane, right-click the applicable routing link, and then click **Properties**.
5. On the **Site Gates** page, do one of the following:

   - To add a Message Queuing routing server as a site gate to the routing link, in **Available servers**, select a server from the list, and then click **Add**.
   - To remove a Message Queuing routing server from the routing link, in **Site gates**, select the applicable server from the list, and then click **Remove**.

6. Repeat step 5 as needed.

## Notes

- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
- A routing link defines the route for connectivity between two sites, and a site gate is a Message Queuing routing server or one of several such servers through which the message traffic between two sites is configured to pass. For more information on routing links, site gates, and how they are used to route messages, see Related Topics.
- A site gate can belong to either or both of the sites connected by a routing link. A routing link can have any number of site gates or no site gate.
- By default, only users in the Enterprise Admins group can create or modify routing links if there are multiple domains in your organization. For users in the Domain Admins group to be able to create or modify routing links, they will need to be granted the Create All Child Objects permission for the MsmqServices object. For information on how to set permissions for routing links, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Configuring routing links
Routing with Message Queuing servers
Routing links
Intersite routing
Set permissions for routing link objects
Working with MMC console files

# Administer Cross-Platform Messaging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Cross-Platform Messaging**
This section shows you how to:

- Create foreign sites
- Add foreign computers to a site

# Create foreign sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create foreign sites**

1. Open Active Directory Sites and Services.
2. On the **View** menu, click **Show Services Node**.
3. In the console tree, right-click **MsmqServices**.

   **Where?**

   - Active Directory Sites and Services/Services/MsmqServices.

4. Click **New Foreign Site**.
5. In the **New Object – Foreign Site** dialog box, type a name for the new site under **Foreign site name**.

**Notes**

- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
- A routing link defines the route for connectivity between any two sites, including a foreign site and a non-foreign site. A site gate is a Message Queuing routing server or one of several such servers through which the message traffic between two sites is configured to pass. For more information, see Related Topics. To provide connectivity to a foreign site, you must perform the following task:
- Create a routing link between the foreign site and a Windows Server 2003 family or Windows 2000 Message Queuing (non-foreign) site that contains a Message Queuing server running a connector application. For more information, see Related Topics.

  1. Define this Message Queuing server as a site gate for the routing link. For more information, see Related Topics.
  2. Open Active Directory Users and Computers.
  3. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
  4. In the console tree, right-click the **msmq** object of the computer running the connector application.
  5. Click **Properties**.
  6. On the **Sites** page, select the foreign site under **All sites**, and then click **Add** to add this site to the list of sites where this server resides and to add it to the foreign site.

- To delete a foreign site, in the Active Directory Sites and Services console tree, right-click the name of the foreign site under **Sites**, and then click **Delete**. Before deleting a foreign site, you may need to refresh the list of sites. To do this, right-click **Sites**, and then click **Refresh**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring cross-platform messaging
Configuring routing links
Routing with Message Queuing servers
Foreign computers and sites
Intersite routing
Set permissions for routing link objects
Working with MMC console files

# Add foreign computers to a site

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add foreign computers to a site**

1. Open Active Directory Sites and Services.
2. On the **View** menu, click **Show Services Node**.
3. In the console tree, right-click **MsmqServices**.
4. Click **New Foreign Computer** to open the **New Object – Foreign Computer** dialog box.
5. In the **Computer name** text box, type the name of the foreign computer.
6. In the **Foreign site** list box, select the name of the foreign site to which the computer will be added from the list of available foreign sites.

**Notes**

- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
- This procedure creates two objects in Active Directory for the foreign computer. One is a computer object in the Computers container. This object is similar to an ordinary computer object. The other is an msmq (MSMQ Configuration) object, which is created under the computer object and contains a pointer to the foreign site in the **MSMQ-Sites** attribute.
- You must create a foreign site before you can add a foreign computer to the site. For information on how to create a foreign site, see Related Topics.
- Typing the name of a foreign computer in the **New Foreign Computer** dialog box does not guarantee that messages can be sent to it.
- A routing link defines the route for connectivity between two sites. A site gate is a Message Queuing routing server or one of several such servers through which the message traffic between two sites is configured to pass. For more information on routing links, site gates, and how they are used to route messages, see Related Topics.
- A Message Queuing routing server that is running a connector application and has been configured as a site gate for a routing link between its site and the foreign site must also be added to the foreign site. To do this, perform the following task:

    1. Open Active Directory Users and Computers.
    2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
    3. In the console tree, right-click **msmq** (console tree location is: Active Directory Users and Computers/*YourDomain*/Domain Controllers/*ApplicableServer*/msmq).
    4. Click **Properties**.
    5. On the **Sites** page, select the foreign site under **All sites**, and then click **Add** to add this site to the list of sites where this server resides and to add it to the foreign site.

- For Windows 2000 computers running a connector application, after adding the foreign site to the list of sites, the foreign site will not appear in the list until you restart the Message Queuing service on the domain controller of the domain to which the Message Queuing server running the connector application belongs.

- Messages cannot be sent over HTTP transport to foreign computers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring cross-platform messaging
Configuring routing links
Routing with Message Queuing servers
Foreign computers and sites
Intersite routing
Set permissions for routing link objects
Working with MMC console files

# Administer Message Queuing in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Message Queuing in Active Directory**

This section shows you how to:

- Remove Message Queuing objects from Active Directory

# Remove Message Queuing objects from Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove Message Queuing objects from Active Directory**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer*/**msmq**

4. Click **Delete**. When prompted, click **Yes**.
5. Open Active Directory Sites and Services.
6. In the details pane, right-click **Msmq Settings**.

   **Where?**

   - Active Directory Sites and Services/**Sites**/*YourSite*/**Servers**/ *YourServer*/**MSMQ Settings**

7. Click **Delete**. When prompted, click **Yes**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
- It is best to remove Message Queuing objects from Active Directory by uninstalling Message Queuing software from the actual computer. It is recommended that you perform the procedure above only if you cannot uninstall Message Queuing software from a computer, such as during a hardware malfunction, if you no longer have physical access to the computer, or if Message Queuing Setup has failed. For information on how to uninstall Message Queuing software from a computer, see Related Topics.
- When you remove Message Queuing objects from Active Directory, all queues and messages are also removed. However, actual Message Queuing software, including registry values and storage files, remains on the computer.
- If you remove objects for an in-routing or out-routing server, the server is also removed from the routing lists of all clients using that server.
- If you remove objects for a site gate, the server is also removed from the associated routing link.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Uninstalling Message Queuing
Message Queuing and Active Directory
Working with MMC console files

# Administer Message Queuing Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Message Queuing Computers**
This section shows you how to:

- Join a domain from a workgroup
- Add sites to a Message Queuing server's list of sites
- Remove sites from a Message Queuing server's list of sites
- Specify a new site for mobile computers
- Change in-routing or out-routing servers
- Change the supporting server for a dependent client
- View the dependent client list for a server
- Set the message storage size for computers
- Set the journal storage size for computers
- Store messages on multiple disk drives
- Set DNS properties for Message Queuing
- Start or stop services for Message Queuing

# Join a domain from a workgroup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To join a domain from a workgroup**

1. Open System in Control Panel.
2. On the **Computer Name** page, click **Change**.
3. Under **Member of**, click **Domain**, type the domain name, and then click **OK**.
4. When prompted, provide a valid user name and password.
5. When prompted, restart your computer to apply the changes.

### Caution

- When you join a Message Queuing computer in workgroup mode to a domain, or move a Message Queuing computer to a different forest, all messages in outgoing transactional queues on the computer will be discarded, or placed in the dead-letter queue if negative source journaling is enabled.

### Notes

- To open System, click **Start**, point to **Settings**, and then click **Control Panel**.
- Use this procedure if you are running Message Queuing in a workgroup environment and want to join (or rejoin) a domain. When you join a domain, the computer can access public queues listed in Active Directory.
- For more information on this procedure, see Join a domain in the System online Help file.
- To perform this procedure, the Active Directory Integration subcomponent must be selected during setup. If the Active Directory Integration subcomponent is not installed, Message Queuing can run only in workgroup mode.
- You can use this procedure to join a workgroup from a domain. In that case, click **Workgroup** under **Member of** and type the name of the workgroup. You will be disjoined from the domain.
- You can use this procedure to move a computer from one domain to another domain within your forest. In this case, if you leave the msmq (MSMQ Configuration) object in the old domain, the computer will continue to use the old object. Otherwise, it will create a new msmq object in the new domain.
- After joining a domain from a workgroup or another domain, some time (as much as 20-30 minutes) is required for replication of the changes in Active Directory.
- When a computer running Message Queuing in workgroup mode joins a Windows NT 4.0 domain, the installation subcomponent for Active Directory Integration is removed. After the computer has been joined to the domain and rebooted, open **Add or Remove Programs** in Control Panel, and add the Active Directory Integration subcomponent.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Deploying in workgroup mode

Working with MMC console files

# Add sites to a Message Queuing server's list of sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add sites to a Message Queuing server's list of sites**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer*/**msmq**

4. Click **Properties**.
5. Click the **Sites** tab, and in **All Sites**, select a site to add to the list of sites where the server resides, and then click **Add**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- In most cases, Message Queuing automatically determines the sites in which servers reside. This procedure is needed to add sites manually to the list of sites where a server resides. The applicable servers are Message Queuing servers with routing enabled, which are also called routing servers. Any routing server that resides in a specific site can be configured as an in-routing or out-routing server for computers that reside in that site. Such routing servers can be configured as site gates in routing links to that site. A routing server can reside in multiple sites, and it can be configured as an in-routing server, an out-routing server, or a site gate, independently in each site in which it resides.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Adding or removing sites from a server list
Routing with Message Queuing servers
Working with MMC console files

# Remove sites from a Message Queuing server's list of sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To remove sites from a Message Queuing server's list of sites**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer/*msmq

4. Click **Properties**.
5. Click the **Sites** tab, and in ***ComputerName* Sites**, select a site to remove from the list of sites where the server resides, and then click **Remove**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- In most cases, Message Queuing automatically determines the sites in which servers reside. This procedure is needed to remove sites manually from the list of sites where a server resides. The applicable servers are Message Queuing servers with routing enabled, which are also called routing servers. Any routing server that resides in a specific site can be configured as an in-routing or out-routing server for computers that reside in that site. Such routing servers can be also configured as site gates in routing links to that site. A routing server can reside in multiple sites, and it can be configured as an in-routing server, an out-routing server, or a site gate, independently in each site in which it resides.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Add sites to a Message Queuing server's list of sites
Adding or removing sites from a server list
Working with MMC console files

# Specify a new site for mobile computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify a new site for mobile computers**

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Click **Properties**.
4. In the **Message Queuing Properties** dialog box, click the **Mobile** tab, and in **New site**, click the arrow that opens the list and select the site to which the mobile computer will move.
5. When prompted to restart Windows, click **Restart Now**.

**Note**

- To open Computer Management, right-click **My Computer**, and then click **Manage**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Moving mobile computers between sites
Working with MMC console files

# Change in-routing or out-routing servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change in-routing or out-routing servers**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer/*msmq

4. Click **Properties**.
5. Click the **Routing** tab.
6. To change an in-routing server used, under **In-routing servers**, click the arrow that opens the list. Select a new Message Queuing server or select **none** to remove a particular Message Queuing server used as an in-routing server.
7. To change an out-routing server used, under **Out-routing servers**, click the arrow that opens the list. Select a new Message Queuing server or select **none** to remove the Message Queuing server used as an out-routing server.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- This task is applicable only to independent clients and Message Queuing servers without routing enabled that are installed on a nondomain controller, but can be performed from any computer with suitable permissions.
- You can install a Message Queuing server with routing enabled on a domain controller, but it may impair performance.
- When a client is moved to another site, these routing settings are disabled until the client returns to its original site.
- Message Queuing servers used as in-routing or out-routing servers must be in the same site as the client. The same computer can be used as both the in-routing server and the out-routing server for a client.
- If used, it is recommended that you assign more than one in-routing or out-routing server per client. This provides load balancing and failure recovery.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Assigning in-routing and out-routing servers

# Change the supporting server for a dependent client

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the supporting server for a dependent client**

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Click **Properties**.
4. On the **Message Queuing Client** page in the **Message Queuing Properties** dialog box, type the name of the supporting Message Queuing server.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- This task is performed from the dependent client.
- Message Queuing 3.0 dependent clients cannot be installed with a Windows NT Server 4.0 computer as a supporting server.
- Any Message Queuing server operating in domain mode with at least independent client functionality, preferably in the local site, but also in a nearby site, can serve as the supporting server for a dependent client.
- By default, following installation, Message Queuing servers are not enabled to support dependent clients. For instructions on enabling a Message Queuing server to act as a supporting server, see Enable a supporting server.
- To enable a Message Queuing server to act as a supporting server for a dependent client, you must change the security configuration of the Microsoft Distributed Transaction Coordinator (DTC) to support network transactions. For details, see Enable network DTC access.
- If you have the early version of a dependent client running on Windows NT 4.0 and you want to change its supporting server to a Message Queuing server running on Windows 2000, the dependent client must first have Windows NT 4.0 Service Pack 4 or higher installed or transactional messaging will not function.
- New Message Queuing 3.0 features, such as HTTP transport, Message Queuing Triggers, and sending messages to multiple destinations, are not available to dependent clients. Thus, distribution lists, multiple-element format names, and multicast messaging are not supported for dependent clients.
- Dependent clients will no longer be supported in future versions of Message Queuing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Dependent clients

Working with MMC console files

# View the dependent client list for a server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view the dependent client list for a server**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer*/msmq

4. Click **Properties**.
5. Click the **Dependent Clients** tab, and then click **Refresh** to ensure that the list is current.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- You can only view those dependent clients that are currently active (such as currently sending messages), which may not be the total number of dependent clients in your organization.
- By default, following installation, Message Queuing servers are not enabled to support dependent clients. For instructions on enabling a Message Queuing server to act as a supporting server, see Enable a supporting server.
- For information on how to change the supporting Message Queuing server for a dependent client, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Dependent clients
Working with MMC console files

# Set the message storage size for computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the message storage size for computers**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer*/msmq

4. Click **Properties**.
5. On the **General** tab, in **Storage limits**, select the **Limit message storage to (KB)** check box, and then type a maximum total size (in kilobytes) for messages stored in all queues on the selected computer.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- This procedure is not applicable to computers operating in workgroup mode or virtual servers in clusters. You can set the computer quota of a workgroup computer using the **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache\MachineQuota** registry key. To change the quota for a virtual server, you need to change the analogous registry key for clustered servers. After setting the registry, restart the Message Queuing service for the changes to take effect.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

- Following a new installation of Message Queuing, or an upgrade from Windows 2000, the default limit for these storage sizes is 8 gigabytes (GB).
- When the computer quota is reached, no messages can be delivered to any queue or journal on the computer. No negative acknowledgment message is generated, and an error is returned for messages sent locally, while a remote computer tries to resend until the cumulative size of all messages in the queues drops below the specified limit.
- For information on how to set the journal storage size limit (or quota) for a computer, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Setting message quotas
Administering Messages
Set the journal storage size for queues
Working with MMC console files

# Set the journal storage size for computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the journal storage size for computers**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer*/msmq

4. Click **Properties**.
5. On the **General** page, in **Storage Limits**, select the **Limit journal storage to (KB)** check box, and then type the total size (in kilobytes) of all journal messages that can be stored in the computer journal on the selected computer.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- This procedure is not applicable to computers operating in workgroup mode or virtual servers in clusters. You can set the computer quota of a workgroup computer using the **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache\MachineJournalQuota** registry key. To change the quota for a virtual server, you need to change the analogous registry key for clustered servers. After setting the registry, restart the Message Queuing service for the changes to take effect.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

- Following a new installation of Message Queuing, or an upgrade from Windows 2000, the default limit for these storage sizes is 8 gigabytes (GB).
- Automatically created folders used to store copies of messages are called *journals*. Messages stored in journals are called *journal messages*. When the journal storage limit (or quota) is reached, additional messages cannot be placed in the computer journal on the computer until the cumulative size of the journal messages in all journals on that computer drops below the specified limit.
- When you set the journal storage size limit for a computer, you automatically set equal quotas for the computer journal (the Journal messages queue), the dead-letter queue, and the transactional dead-letter queue.
- The journal messages stored in source or computer journals can include any outgoing nontransactional message. The journal messages stored in target or queue journals can include any incoming message, including transactional messages.
- For information on how to set the message storage limit (or quota) for a computer, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Setting message quotas
Working with MMC console files

# Store messages on multiple disk drives

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To store messages on multiple disk drives**

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Click **Properties**.
4. On the **Storage** page, type a different drive letter and path as applicable in **Message files folder**, **Message logger folder**, or **Transaction logger folder**.
5. If prompted to stop the Message Queuing service, click **Yes**.
6. If prompted to restart Windows, click **Restart Now**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- By default, Message Queuing stores all message files, log files, and transaction log files in the \Windows\System32\Msmq\Storage folder on the drive on which Message Queuing is installed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Storing message files on multiple disks
Working with MMC console files

# Set DNS properties for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set DNS properties for Message Queuing**

1. Open Network Connections in Control Panel.
2. Right-click **Local Area Connection**, and then click **Properties**.
3. In the **Local Area Connection Properties** dialog box, in the components list, click **Internet Protocol (TCP/IP)**.
4. Click **Properties**.
5. On the **General** page, do any of the following:

   - To specify that the DNS server will be located automatically, on the **General** page, select **Obtain DNS server address automatically**.
   - To specify one or two DNS servers, on the **General** page, select **Use the following DNS server addresses**. Then, in **Preferred DNS server**, type the IP address of the preferred DNS server, and in **Alternate DNS server**, type the IP address of an optional alternate DNS server.
   - To specify more than two DNS servers, on the **General** page, click **Advanced**. On the **Advanced TCP/IP Settings** page, click the **DNS** tab. Under **DNS server addresses, in order of use**, click **Add**; in **DNS server**, type the address of a DNS server and click **Add**. Repeat this step for each DNS server that you want to specify. You can use the **Up** and **Down** buttons to change the order of the DNS servers in the list. You can also click **Edit** to change the address of the selected DNS server in the list, or you can click **Remove** to delete it from the list.

## Notes

- To open Network Connections, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Network Connections**.
- This task is performed only if you are operating in a DNS-only environment (an environment that has no WINS server). Message Queuing requires DNS to be configured to correctly resolve computer names.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Configuring DNS settings for Message Queuing
Working with MMC console files

# Start or stop services for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To start or stop services for Message Queuing**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. In the details pane, right-click the applicable service, and then do one of the following, if applicable (options that are not applicable are dimmed):

   - To start a service, click **Start**.
   - To stop a service, click **Stop**.
   - To pause a service, click **Pause**.
   - To resume a service that has been paused, click **Resume**.
   - To restart a service, click **Restart**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- You can view the status of a service by looking in the **Status** column in the details pane.
- This procedure is applicable to the Message Queuing service, and the Message Queuing Triggers service.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing the Message Queuing service account
Working with MMC console files

# Administer Security for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Security for Message Queuing**
This section shows you how to:

- Set permissions for computer and queue objects
- Set permissions for routing link objects
- Set permissions for MSMQ Settings objects
- Set permissions for Message Queuing files
- Set permissions for distribution lists
- Change ownership of computer and queue objects
- Change the Message Queuing service account
- Audit access to computer and queue objects
- Enable Message Queuing to access SACL information
- Audit access to routing link objects
- Enable hardened MSMQ mode
- Enable secured remote read
- Weaken security using ADSI Edit
- Change the security settings for Internet Information Services
- Administer Authentication and Encryption for Message Queuing

# Set permissions for computer and queue objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions for computer and queue objects**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. Do one of the following:
   - To grant Message Queuing-specific permissions for a computer (an msmq object), in the console tree, right-click **msmq** (console tree location is: Active Directory Users and Computers/*YourDomain/YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer/*msmq).
   - To grant Message Queuing-specific permissions for a queue, right-click the applicable queue. (console tree location is: Active Directory Users and Computers\ *YourDomain*\ *YourOrganizationalUnit* (such as Computers or Domain Controllers)\ *YourComputer*\ msmq\ *YourQueueFolder* (Private Queues for a private queue)\ *YourQueue*).
4. Click **Properties**.
5. On the **Security** page, set permissions for the object specified in Step 3, as needed:
   - To grant permissions for this object to a group or user appearing under **Group or user names**, select the applicable group or user, and then in **Permissions for *GroupOrUser***, select the check boxes in the **Allow** column following the names of the applicable permissions.
   - To deny a group or user permissions for this object, select the applicable group or user in **Group or user names**, and then in **Permissions for *GroupOrUser***, select the check boxes in the **Deny** column following the names of the applicable permissions.
   - To add a new group or user for access, click **Add**. In the **Select Users, Computers, or Groups** dialog box, click **Object Types**, select the **Group** and/or **Users** check box as appropriate, clear the remaining check boxes, and click **OK**. In **Enter the object name to select**, type the name of a group or user or the names of several groups or users separated by semicolons, and click **OK**. Or, click **Advanced** to search for groups or users, enter the applicable parameters, click **Find Now**, select the group or user, click **OK**, and then click **OK** again. Then, select the group or user you just added and select the applicable check boxes as needed.

**Notes**

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- The queue objects for the queues residing on a particular computer are child objects of the msmq object of the applicable computer. For example, to create a queue, a user must have the Create All Child Objects permission for the msmq object under which the queue will be created.
- You can grant or deny permissions for an object even to the Administrators group.
- This procedure cannot be used to set permissions for a private queue on a remote computer.

**Important**

- The default permission is that everyone can send messages to a queue. For tighter security, you can change the default security permissions for the queue. You can also specify properties for greater security when you create a queue, for example, to accept authenticated messages only. For instructions, see Allow only authenticated messages on queues.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Message Queuing and Active Directory
Access control for Message Queuing
Working with MMC console files

# Set permissions for routing link objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions for routing link objects**

1. Open Active Directory Sites and Services MsmqServices node.

   **Where?**

   - Click **Start**, click **Programs**, click **Administrative Tools**, and click **Active Directory Sites and Services** to launch the **Active Directory Sites and Services** dialog box.
   - Click **Active Directory Sites and Services** in the console tree.
   - Click the **View** menu of the **Active Directory Sites and Services** dialog box and select the option to **Show Services Node** if this option is not already enabled.
   - Click to expand the **Services** node in the console tree.
   - Right-click **MsmqServices**, and then click **Properties**.

2. On the **Security** page of the MsmqServices node, set permissions as needed:

   - To grant permissions for this object to a group or user appearing under **Group or user names**, select the applicable group or user, and then in **Permissions for *GroupOrUser***, select the check boxes in the **Allow** column.
   - To deny a group or user permissions for this object to a group or user appearing under **Group or user names**, select the applicable group or user appearing under **Group or user names**, select the applicable group or user, and then in **Permissions for *GroupOrUser***, select the check boxes in the **Deny** column.
   - To add a new group or user for access, click **Add** and in the **Select Users, Computers, or Groups** dialog box, click **Object Types**, select the appropriate Object type, clear the remaining check boxes, and click **OK**. In **Enter the object name to select**, type the name of a group or user or the names of one or more groups or users separated by semicolons, and then click **OK**. Or, click **Advanced** to search for groups or users, enter the applicable parameters, click **Find Now**, select the group or user, click **OK**, and then click **OK** again. Then select the group or user you just added and select the applicable check boxes in **Permissions for *GroupOrUser*** as needed.

| **Note** |
| --- |
| Each routing link object is a child object of the MsmqServices object. Thus, the **Create all child objects** and **Delete all child objects** permissions for the MsmqServices object are needed to create, modify, and delete routing links. By default, only the Enterprise Admins group is granted both of these permissions. |

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Message Queuing and Active Directory

Access control for Message Queuing
Configuring routing links
Working with MMC console files

# Set permissions for MSMQ Settings objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions for MSMQ Settings objects**

1. Open Active Directory Sites and Services.

2. In the console tree, select the applicable server object.

   **Where?**

   - Active Directory Sites and Services/Sites/*YourSiteObject*/Servers/*YourServerObject*

3. Right-click the **MSMQ Settings** object in the details pane, and then click **Properties**.

4. On the **Security** page, set permissions as needed:

   - To grant permissions for this object to a group or user appearing under **Group or user names**, select the applicable group or user, and then in **Permissions for *GroupOrUser***, select the check boxes in the **Allow** column following **Create All Child Objects** and **Delete All Child Objects**.

   - To deny a group or user permissions for this object, select the applicable group or user in **Group or user names**, and then select the check boxes in the **Deny** column following **Create All Child Objects** and **Delete All Child Objects**.

   - To add a new group or user for access, click **Add**. In the **Select Users, Computers, or Groups** dialog box, click **Object Types**, select the **Group** and/or **Users** check box as appropriate, clear the remaining check boxes, and click **OK**. In **Enter the object names to select**, type the name of a group or user or the names of several groups or users separated by semicolons, and then click **OK**. Or, click **Advanced** to search for groups or users, enter the applicable parameters, click **Find Now**, select the group or user, click **OK**, and then click **OK** again. Then select the group or user that you just added and select the applicable check boxes as needed.

**Note**

- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Message Queuing and Active Directory
Access control for Message Queuing
Working with MMC console files

# Set permissions for Message Queuing files

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions for Message Queuing files**

1. On the desktop, right-click **My Computer**, and then click **Explore**.
2. Locate and right-click the folder used to store the message files, log files, and transaction log files (usually **Storage** under \Windows\System32\Msmq), and then click **Properties**.
3. On the **Security** page, in **Group or user names**, click **Administrators**.
4. In **Permissions for Administrators**, under **Allow**, select the **Full Control** check box.
5. In **Group or user names**, select any group or user, including the Administrators group, that you do not want to access the Message Queuing storage files directly, and click **Remove**. Repeat this step as needed.

### Note

- You can set both file-level and folder-level permissions only on NTFS-formatted drives. You cannot set folder-level permissions on FAT-formatted drives. (The **Security** tab is not available for folders on FAT-formatted drives.)

### Important

- By default, only members of the Administrators group of the local computer have permissions for the folder used to store message files, log files, and transaction log files. For best security practice, it is not recommended that you widen these default permissions, or you may open Message Queuing for access by unauthorized persons.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Access control for Message Queuing
Working with MMC console files

# Set permissions for distribution lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set permissions for distribution lists**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. To grant permissions for a distribution list, in the console tree, right-click the applicable distribution list (group).

   **Where?**

   - Active Directory Users and Computers/*YourDomain/*Users/*YourDistributionList(Group)*

4. On the **Security** page, set permissions for the object specified in Step 3, as needed:

   - To grant permissions for this object to a group or user appearing under **Group or user names**, select the applicable group or user, and then in **Permissions for *GroupOrUser***, select the check boxes in the **Allow** column following the names of the applicable permissions.
   - To deny a group or user permissions for this object, select the applicable group or user in **Group or user names**, and then select the check boxes in the **Deny** column following the names of the applicable permissions.
   - To add a new group or user for access, click **Add**. In the **Select Users, Computers, or Groups** dialog box, click **Object Types**, select the **Group** and/or **Users** check box as appropriate, clear the remaining check boxes, and click **OK**. In **Enter the object name to select**, type the name of a group or user or the names of several groups or users separated by semicolons, and then click **OK**. Or, click **Advanced** to search for groups or users, enter the applicable parameters, click **Find Now**, select the group or user, click **OK**, and then click **OK** again. Then, select the group or user you just added and select the applicable check boxes as needed.

### Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Message Queuing distribution lists are distribution group objects in Active Directory.

### Important

- Distribution lists are not created or managed with Message Queuing tools or APIs. For best security practice, consider carefully the permissions granted to distribution list objects.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

View or change properties of distribution lists
Access control for Message Queuing
Working with MMC console files

# Change ownership of computer and queue objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change ownership of computer and queue objects**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. To change the ownership of a computer (an msmq object), in the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer/*msmq

   Or, to change the ownership of a queue, right-click the applicable queue.

   **Where?**

   - Active Directory Users and Computers\ *YourDomain*\ *YourOrganizationalUnit* (such as Computers or Domain Controllers)\ *YourComputer*\ msmq\ *YourQueueFolder* (Private Queues for a private queue)\ *YourQueue*

4. Click **Properties**.
5. On the **Security** page, click **Advanced**.
6. On the **Owner** page, in **Change Owner to**, click a new user or group, and then click **OK**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- By default, the computer account is the owner of the msmq object, unless routing or the Downlevel Client Support component is enabled on the server. In that case, the creator is the owner. For the queue object, the creator is the owner by default.
- When a local user (a user not logged on to a Windows Server 2003 family domain) creates an object, the Everyone group is granted **Full Control** permissions for the object.
- There can be only one owner per Message Queuing computer or queue.
- This procedure cannot be used to change the ownership of a private queue on a remote computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Message Queuing and Active Directory

Access control for Message Queuing

Working with MMC console files

# Change the Message Queuing service account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the Message Queuing service account**

1. Open Computer Management.
2. In the console tree, click **Services**.

   **Where?**

   - Computer Management/Services and Applications/Services

3. In the details pane, double-click **Message Queuing**.
4. On the **Log On** page, click **This account**, and then click **Browse**. In **Name**, type the name of the user account, and then click **OK**.
5. In **Password**, type the account password.
6. In **Confirm Password**, type the account password again.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- This procedure is also applicable to the Message Queuing Triggers service. In this case, in Step 3, double-click **Message Queuing Triggers** in the details pane.
- Running the Message Queuing service under user accounts is not supported on domain controllers. Thus, the Message Queuing Downlevel Client Support service, which enables pre-Windows Server 2003 family clients to access objects in Active Directory, must run under the Local System account on domain controllers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Message Queuing security overview
Managing the Message Queuing service account
Working with MMC console files

# Audit access to computer and queue objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To audit access to computer and queue objects**

1. Open Active Directory Users and Computers.

2. Highlight **Active Directory Users and Computers**, and on the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.

3. To audit access to a computer (an msmq object), in the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer/*msmq

   Or, to audit access to a queue, right-click the applicable queue.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer/*msmq/*YourQueueFolder* (Private Queues for a private queue)/*YourQueue*

4. Click **Properties**.

5. On the **Security** page, click **Advanced**.

6. On the **Auditing** page of the **Advanced Security Settings** dialog box, click **Add**.

7. In the **Select Users, Computers, or Groups** dialog box, click **Object Types**, select the **Group** and/or **Users** check box as appropriate, clear the remaining check boxes, and click **OK**. In **Enter the object name to select**, type the name of a group or user whose access you want to audit or the names of several such groups or users separated by semicolons and click **OK**. Or, click **Advanced** to search for groups or users, enter the applicable parameters, click **Find Now**, select the group or user, click **OK**, and then click **OK** again.

8. In the **Auditing Entry** dialog box, in **Apply onto**, select the applicable objects for which user access will be audited. Then, under **Access**, select the appropriate check boxes next to the applicable entries and click **OK**. Or, to clear all audit entries, click **Clear All**, and then click **OK**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.

- This procedure cannot be used to audit access to a private queue on a remote computer.

- In addition to auditing access to computer and queue objects, you must also establish an audit policy for the local computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Auditing Message Queuing objects
Message Queuing and Active Directory
Working with MMC console files

# Enable Message Queuing to access SACL information

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable Message Queuing to access SACL information**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click **Domain Controllers**.

   **Where?**

   - Active Directory Users and Computers/Domain Controllers

3. Click **Properties**.
4. Click the **Group Policy** tab.
5. Click the default Group Policy, and then click **Edit**.
6. In the Group Policy object snap-in, open **Computer Configuration**, open **Windows Settings**, open **Security Settings**, open **Local Policies**, right-click **User Rights Assignment**, and then click **Open**.
7. Double-click **Manage Auditing and Security Log**, and then click **Add User or Group**.
8. Add the computer account name of the computer running the Message Queuing service, and then click **OK**.
9. Wait for about 5 minutes to allow changes to take effect in Active Directory, and then restart the computer.

### Note

- After the Message Queuing service is enabled to access SACL information, you need to establish an auditing policy. For information, see Define or modify auditing policy settings for an event category.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Auditing Message Queuing objects
Message Queuing and Active Directory
Working with MMC console files

# Audit access to routing link objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To audit access to routing link objects**

1. Open Active Directory Sites and Services.
2. Highlight **Active Directory Sites and Services**, and on the **View** menu, click **Show Services Node**.
3. In the console tree, click **MsmqServices**.

   **Where?**

   - Active Directory Sites and Services/Services/MsmqServices

4. In the details pane, right-click the applicable routing link, and then click **Properties**.
5. On the **Security** page, click **Advanced**.
6. On the **Auditing** page of the **Advanced Security Settings** dialog box, click **Add**.
7. In the **Select Users, Computers, or Groups** dialog box, click **Object Types**, select the **Group** and/or **Users** check box as appropriate, clear the remaining check boxes, and click **OK**. In **Enter the object name to select**, type the name of a group or user whose access you want to audit or the names of several such groups or users separated by semicolons and click **OK**. Or, click **Advanced** to search for groups or users, enter the applicable parameters, click **Find Now**, select the group or user, click **OK**, and click **OK** again.
8. In the **Auditing Entry** dialog box, in **Apply onto**, select the applicable objects for which user access will be audited. Then, under **Access**, select the appropriate check boxes next to the applicable entries and click **OK**. Or, to clear all audit entries, click **Clear All**, and then click **OK**.

## Notes

- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
- The information generated by an audit is stored in the security log. For information on how to use Event Viewer to view information in the security log, see Related Topics.
- This procedure can be used to audit access to any Message Queuing object under, and including, the MsmqServices object.
- In addition to auditing access to computer and queue objects, you must also establish an audit policy for the local computer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Auditing Message Queuing objects
Message Queuing and Active Directory
Working with MMC console files

# Enable hardened MSMQ mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable hardened MSMQ mode**

1. Open Computer Management.
2. In the console tree, click **Services and Applications**.

   **Where?**

   - Computer Management/Services and Applications

3. In the details pane, right-click **Message Queuing**, and then click **Properties**.
4. On the **Server Security** tab, select **Enable Hardened MSMQ mode to secure this computer on the Internet**, and then click **OK**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- Following installation, hardened MSMQ mode is not enabled by default.
- To reset server security settings on this tab to the default settings following installation, click **Restore Defaults**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Message Queuing in hardened MSMQ mode
Working with MMC console files

# Enable secured remote read

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable secured remote read**

1. Open Computer Management.
2. In the console tree, click **Services and Applications**.

   **Where?**

   - Computer Management/Services and Applications

3. In the details pane, right-click **Message Queuing**, and then click **Properties**.
4. On the **Server Security** tab, select **Enable secured remote receiving. In this mode, only servers running Windows Server 2003 can receive messages remotely from queues on this computer**, and then click **OK**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- To reset server security settings on this tab to the default settings following installation, click **Restore Defaults**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Message Queuing security overview
Working with MMC console files

# Weaken security using ADSI Edit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To weaken security using ADSI Edit**

1. Click **Start**, click **Run**, type **adsiedit.msc**, and then press **Enter**.
2. In the console tree, right-click **CN=MsmqServices**.

   **Where?**

   - Configuration Container/CN=Configuration,.../CN=Services/CN=MsmqServices

3. Click **Properties**.
4. On the **Attribute Editor** page, in **Attributes**, select *mSMQNameStyle*, and then click **Edit**.
5. In **Boolean Attribute Editor**, select **True** to weaken security, or **False** to tighten security.
6. Click **OK**.

## Notes

- Windows Server 2003 family Support Tools, which is provided on the Windows Server 2003 family compact disc, must be installed before you can use ADSI Edit. To install these tools, perform the following steps:

  1. Insert the Windows Server 2003 family CD in your CD-ROM drive.
  2. Open the Support folder, and then open the Tools folder.
  3. Double-click **Suptools** to start the install program.
  4. Follow the instructions in the Windows Support Tools Setup Wizard, and select **Complete** on the **Select An Installation Type** page.

- This procedure is used to weaken Active Directory security to support MSMQ 1.0 users, Message Queuing 2.0 clients logged on with local user accounts, and certain operating configurations.
- Only users with the Write permission for the MsmqServices object can perform this procedure. By default, only members of the Domain Admins or Enterprise Admins groups of the root domain and members of the SYSTEM group (local system services running on a domain controller) have the Write permissions for this object.
- Dependent clients cannot run under a local user account. Also, any computer that sends queries about Message Queuing objects to Active Directory on a domain controller directly, rather than through the Message Queuing directory service, will not be able to access Active Directory when it logs on using a local user account even if weakened security for Active Directory is enabled.
- After performing this procedure, for the change to take effect, you must restart all instances of the Message Queuing Downlevel Client Support service running on Windows Server 2003 domain controllers and Message Queuing servers running on Windows 2000 domain controllers.

- For best security practice, it is recommended that Active Directory security not be weakened unless necessary.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Enabling weakened security
Working with MMC console files

# Change the security settings for Internet Information Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the security settings for Internet Information Services**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.

2. On the **File** menu, click **Add/Remove Snap-in**.

3. In the **Add/Remove Snap-in** dialog box, click **Add**, select **Internet Information Services**, click **Add**, click **Close**, and then click **OK**.

4. In the console tree, right-click **MSMQ** virtual directory.

   **Where?**

   - Console Root/Internet Information Services/*YourInternet Information ServicesComputer/*Web Sites/Default Web Site/MSMQ

5. Click **Properties**, and then click the **Directory Security** tab.

6. Under **Authentication and access control**, click **Edit**.

7. To use a domain user account instead of the Internet Information Services local user for Anonymous Access, do the following:

   - In the **User Name** edit box, type the name of a domain user with the required permissions for the Message Queuing operation.

   - Type the user password, and then click **OK**.

   - Retype the password in the **Confirm Password** dialog, and then click **OK**.

8. To disable anonymous access and trust the Internet Information Services computer for delegation, do the following:

   - In the **Enable Anonymous Access** group, clear the **Enable anonymous access** check box, and click **OK**.

   - Open Active Directory Users and Computers.

   - On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.

   - In the console tree, right-click the name of your Internet Information Services computer.

     **Where?**

     Active Directory Users and Computers\ *YourDomain*\ Computers\ *YourInternet Information ServicesComputer*

   - Click **Properties**.

   - On the **General** page, select the **Trust computer for delegation** check box, and then click **OK**

9. In the **Internet Information Services** snap-in, right-click *YourInternet Information ServicesComputer*.

   **Where?**

   Console Root\ Internet Information Services\ *YourInternet Information ServicesComputer*

10. Restart Internet Information Services by selecting **All Tasks \ Restart Internet Information Services...**, and then click **OK**.

## Notes

- By default, Internet Information Services (IIS) impersonates a special Internet Information Services *local* user account. This account is local to the IIS computer (unless it is a domain controller). It cannot be authenticated by other computers in the network, and is treated as an anonymous user. Active Server Page (ASP) applications and scripts run under IIS and, by default, impersonate the IIS local user for any Message Queuing operations, including queries to Active Directory. Since anonymous users no longer belong to the Everyone group in Windows Server 2003 family operating systems, these queries fail. ASP applications and scripts cannot locate, create, and delete queues by default. In addition anonymous users cannot open a queue for remote read, and you cannot read messages from queues not belonging to the IIS computer using ASP by default.

- For the changes to take effect, you must stop and start the World Wide Web Publishing service.

- Trusting the computer for delegation allows multiple hops to communicate with a domain controller, assuming that the user account is trusted for delegation.

### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authentication for Message Queuing
Message Queuing security overview
Working with MMC console files

# Administer Authentication and Encryption for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Authentication and Encryption for Message Queuing**
This section shows you how to:

- Obtain client certificates
- Enable client certificates
- Register certificates for Message Queuing
- Remove certificates for Message Queuing
- View certificates for Message Queuing
- Renew certificates for Message Queuing
- Renew cryptographic keys for Message Queuing
- Obtain server certificates for HTTPS messaging
- Administer Server Authentication

# Obtain client certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To obtain client certificates**

1. In your Web browser, open the form at http://*servername*/certsrv for requesting a certificate from your CA, where *servername* is the name of the Web server where the CA that you want to access is located.

2. Click **Request a certificate**, click **Advanced certificate requests**, and then click **Create and submit a request to this CA**.

3. Type the information requested and any other options required, including:

   - Select **Mark keys as exportable**.

   - Select **Store certificate in the local computer certificate store**.

4. Click **Submit**, and then do one of the following:

   - If you see the **Certificate Pending** Web page, check on the pending certificate request.

   - If you see the **Certificate Issued** Web page, click **Install this certificate**.

5. If you are finished using the Certificate Services Web pages, close Internet Explorer.

Once you have obtained the client certificate, verify that it is in your Personal store and is valid. To do this:

1. Open the Certificates MMC snap-in.

2. Locate the installed client certificate in the **Personal** store under **Certificates (Local Computer)**, and then double-click the certificate to verify it.

Once you have ensured that the client certificate is valid, copy it to the MSMQ certificates store, as follows:

1. Open the Certificates MMC snap-in.

2. Locate the installed client certificate in the **Personal** store under **Certificates (Local Computer)**.

   **Where?**

   - Console Root/Certificates (Local Computer)/Personal/Certificates

3. Drag the certificate with your mouse to move it to the **MSMQ\Personal** store under **Certificates - Service (Message Queuing) on Local Computer** in the Certificates snap-in.

   **Where?**

   - Console Root/Certificates - Service (Message Queuing) on Local Computer/MSMQ\Personal/Certificates

**Notes**

- To open the Certificates snap-in, click **Start**, click **Run**, type **mmc**, and then click **OK**. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**. Select **Certificates**, and then click **Add**. Select **Service account**, click **Next**, select **Local Computer (the computer this console is running on)**, and then click **Next**. In **Service Account**, select **Message Queuing**, and then click **Finish**. When you finish working in the MMC console, you can save the settings so that the custom MMC is available next time you open the console.

- To check on a pending certificate request, in the CA Web page, click **View the status of a pending certificate request**. If you see pending certificates, select the certificate request you want to check, and then click **Next**. If the status is **Still pending**, you must wait for the administrator of the CA to issue the certificate. If the status is **Issued**, to install the certificate, click **Install this certificate**. If the status is **Denied**, contact the CA administrator for more information.

- Note that before requesting a client certificate, you might need to make the CA Web page a trusted site for Internet Explorer. To do this, on the **Tools** menu in Internet Explorer, click **Internet Options**. On the **Security** tab, click **Trusted Sites**, and then click **Sites**. Uncheck **Require server verification (https:) for all sites in the zone**. Type **http://*servername***, and then click **Add**.

- It is not necessary to have the client certificate locally on the receiving computer, but the receiving computer must be able to access the client certificate.

- This task can be performed only from a Message Queuing server installed on a Windows Server 2003 family domain controller.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authentication for Message Queuing
Working with MMC console files

# Enable client certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To enable client certificates

1. On the receiving computer, open the Internet Information Services snap-in as follows: click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Internet Information Services (IIS)**.

2. In the console tree, right-click **MSMQ**.

   **Where?**

   - Console Root/Internet Information Services/*YourComputer*/Web Sites/Default Web Site/MSMQ

3. Click **Properties**.

4. Select the **Directory Security** tab, and then under **Secure Communications**, click **Edit**.

5. To require that a secure link be used for a site resource, click **Require secure channel (SSL)**.

6. Under **Client certificates**, select one of the following to enable client certificate authentication:

   - To enable clients to access a resource with a client certificate, click **Accept client certificates**. Note that a client certificate is not required.

   - To ensure that users without a valid client certificate will be denied access to the site resource, click **Require client certificates**.

   - To grant access to users with or without a client certificate, click **Ignore client certificates**.

7. To enable client certificate mapping, in the **Secure Communications dialog box,** select **Enable client certificate mapping** if not already selected, and then click **Edit**. For instructions on creating a mapping, see the IIS online Help file.

### Notes

- By default, anonymous access is enabled for the **MSMQ** virtual directory using the account IUSR_computername. To prevent both authenticated and unauthenticated HTTP messages from being sent, disable anonymous access. For instructions, see Change the security settings for Internet Information Services.

- By default, IUSR_computername user has write permission for the physical MSMQ directory (default is %SystemDrive%\Inetpub\wwwroot\msmq), and **Ignore client certificates** is selected, thus allowing everyone to send HTTPS messages. You can control access to **MSMQ** virtual directory by configuring security settings for the physical Message Queuing directory.

- Requiring a secure channel means that a user cannot connect to a site resource without using a secure link (that is, the link's URL must begin with https://).

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Authentication for Message Queuing

Working with MMC console files

# Register certificates for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To register certificates for Message Queuing**

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Click **Properties**.
4. In the **Message Queuing Properties** dialog box, click the **User Certificate** tab, and then under **User certificates**, click **Register**.
5. In the **Personal Certificates** dialog box, click a certificate, and then click **Register**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- Your computer must have access to Active Directory to perform this task.
- During setup, an internal user certificate is automatically created for you on the local computer and is registered in Active Directory the first time you log on to the local computer in a domain provided your computer can communicate with a domain controller when you log on.
- If no internal certificate exists, such as when the internal certificate is removed, this procedure creates a new internal certificate, which you can register. You can also use this procedure to register an external certificate after you import it.
- If a certificate registered on the local computer is removed from a remote computer, the record of the certificate is removed from Active Directory, but the certificate will still exist on the local computer.
- By default, users have permission to register certificates for Message Queuing. However, if default user permissions are changed, this might affect your ability to register certificates. For registering certificates, the user object requires the Write Personal Information permission in the Active Directory.
- Active Directory sets a multi-valued attribute limit of approximately 800 user certificates for a specific user account. This limit is usually exceeded when obsolete user certificates have not been deleted from Active Directory. If multiple certificates exist for a user account, only the latest is used, and obsolete certificates can be deleted. For instructions, see Remove certificates for Message Queuing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Authentication for Message Queuing
Remove certificates for Message Queuing
Renew certificates for Message Queuing
Working with MMC console files

# Remove certificates for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

#### To remove certificates for Message Queuing

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Click **Properties**.
4. In the **Message Queuing Properties** dialog box, click the **User Certificate** tab, and then under **User certificates**, click **Remove**.
5. In the **Personal Certificates** dialog box, click the applicable user certificate, and then click **Remove**.

### Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- A list of all certificates registered for the user in Active Directory is displayed in the **Personal Certificates** dialog box. There might be a certificate for the same user on more than one computer, including the computer from which you are currently running the MMC snap-in. The list will take the form of **domain\user, computer name**.
- You can remove any user certificate that is listed. However, if a registered certificate for a computer is removed using another computer, the certificate is removed from Active Directory, but will still exist on the local computer.
- Active Directory sets a multi-valued attribute limit of approximately 800 user certificates for a specific user account. This limit is usually exceeded when obsolete user certificates have not been deleted from Active Directory. If multiple certificates exist for a user account (*user account, computer_name*), indicating obsolete entries, then only the latest certificate is used, and the others can be deleted. For example:
  - DOMAINA\user1, computer3
  - DOMAINA\user1, computer3
  - DOMAINA\user1, computer3
- You can check which is the latest entry for **DOMAINA\user1, computer3**, and delete the other **computer3** entries. To check the latest entry, click the required certificate, and then click **View Certificate**. On the **Details** tab, look at the **Valid from** field. Note that only the certificate with the most recent **Valid from** date is in use. The others are obsolete.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Authentication for Message Queuing
Register certificates for Message Queuing
Renew certificates for Message Queuing
Working with MMC console files

# View certificates for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view certificates for Message Queuing**

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Click **Properties**.
4. In the **Message Queuing Properties** dialog box, click the **User Certificate** tab, and then under **User certificates**, click **View**.
5. In the **Personal Certificates** dialog box, click a certificate, and then click **View Certificate**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- A list of all certificates registered for the user in Active Directory is displayed in the **Personal Certificates** dialog box. There might be a certificate for the same user on more than one computer (including the computer from which you are currently running the MMC snap-in. The list will take the form of **domain\user, computer name**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authentication for Message Queuing
Remove certificates for Message Queuing
Register certificates for Message Queuing

# Renew certificates for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To renew certificates for Message Queuing**

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Click **Properties**.
4. In the **Message Queuing Properties** dialog box, click the **User Certificate** tab, and then under **Internal certificate**, click **Renew**.

   A warning appears stating that authenticated messages that you have already sent using the current internal certificate may be rejected as not authenticated.

5. If rejecting those messages is acceptable, click **Yes**. Or, if it is not acceptable, click **No** and renew your internal certificate some other time.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- When the certificate is successfully renewed, a message appears stating that the internal certificate was successfully renewed and registered in Active Directory.
- This task is not relevant for Message Queuing computers installed in workgroup mode.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authentication for Message Queuing
Remove certificates for Message Queuing
Register certificates for Message Queuing

# Renew cryptographic keys for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To renew cryptographic keys for Message Queuing**

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Click **Properties**.
4. In the **Message Queuing Properties** dialog box, click the **Service Security** tab, and under **Cryptographic keys**, click **Renew**.

   A warning appears stating that received messages may be encrypted with a cryptographic key different from the one that your computer currently has. You will not be able to read those messages. You will be asked whether you want to continue.

5. If not being able to read those messages is acceptable, click **Yes**. Or, if it is not acceptable, click **No** and renew your cryptographic key at some other time.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- Only the user who installed Message Queuing has the permissions needed to renew the cryptographic keys. If Message Queuing is installed in workgroup mode, and then the computer joins a domain, only domain administrators are allowed to renew the cryptographic keys.
- This task does not apply to dependent clients and Message Queuing computers installed in workgroup mode.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Renew certificates for Message Queuing
Authentication for Message Queuing

# Obtain server certificates for HTTPS messaging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To obtain server certificates for HTTPS messaging**

1. On the server requiring a certificate, to request a certificate, open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the server where the certification authority (CA) you want to access is located.
3. Click **Request a certificate**, and then click **Advanced certificate request**.
4. To submit a request using a form, click **Create and submit a certificate request to this CA**, and then click **Next**.
5. Fill in any identifying information requested and any other options you require, including:
   - In **Name**, enter the name of the computer that requires the certificate.
   - In **Type of Certificate Needed**, choose **Server Authentication Certificate**.
   - In **Key Options**, select **Use the local machine store**.
   - In **Key Options**, select **Marks keys as exportable** if you want to export the private key at a later time.
   - Leave other options at default values.
6. Click **Submit**.
7. When you see the Certificate Service Web Page, click **Install this certificate** to add the certificate to your certificate store.
8. If you are finished using the Certificate Services Web pages, close Internet Explorer.
9. After you have the certificate, open Computer Management.
10. In the console tree, click **Default Web Site**.

    **Where?**

    - Computer Management/Services and Applications/Internet Information Services (IIS)/Web Sites
11. Right-click **Default Web Site**, and then click **Properties**.
12. On the **Directory Security** tab, click **Server Certificate**.
13. In the Web Server Certificate Wizard, do the following:
    - On the **Server Certificate** page, select **Assign an existing certificate**.
    - In **Available Certificates**, select the required certificate, and then click **Next**.

- In the **SSL Port** page, click **Next**.
- In **Certificate Summary**, click **Next**, and then click **Finish**.

### Note

- After this procedure is complete, you can view the certificate. To do this, right-click **Default Web Site**, and then click **Properties**. On the **Directory Security** tab, click **View Certificate**.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
HTTPS authentication
Working with MMC console files

# Administer Server Authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Server Authentication**
This section shows you how to:

- Export key certificates from Windows NT 4.0
- Create server certificates for Message Queuing
- Import server certificates for Message Queuing
- Change server certificates for Message Queuing
- Import a CA certificate from Microsoft Certificate Services

# Export key certificates from Windows NT 4.0

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To export key certificates from Windows NT 4.0**

1. On your MSMQ 1.0 controller server with Windows NT 4.0 Option Pack installed, click **Start**, point to **Programs**, point to **Windows NT 4.0 Option Pack**, point to **Microsoft Information Server**, and then click **Internet Services Manager**.

2. In the console tree, right-click **Default Web Site**.

   **Where?**

   - Console Root/Internet Information Services/*YourComputer*/Default Web Site

3. Click **Properties**.

4. Click the **Directory Security** tab, and then under **Secure Communications**, click **Key Manager**.

5. In the Key Manager console, click **Certificate**.

   **Where?**

   - Local Computer/MSMQ/Certificate

6. On the **Key** menu, point to **Export Key** and click **Backup File**.

7. Follow the instructions in the wizard and type a name for the backup file.

## Notes

- For more information about using Internet Services Manager, click the **Action** menu after the console appears, and then click **Help**.

- This task applies only if you are upgrading an MSMQ 1.0 controller server to a Message Queuing server.

- You must export the key certificate before you upgrade the MSMQ 1.0 controller server. Then you can import the certificate to the upgraded server. You need to import a server (key) certificate only for each Message Queuing server that will reply to Active Directory queries from MSMQ clients running on Windows NT 4.0, Windows 98, or Windows 95.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authentication for Message Queuing
Working with MMC console files

# Create server certificates for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create server certificates for Message Queuing**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. In the **Add/Remove Snap-in** dialog box, click **Add**, select **Certificates**, and then click **Add**.
4. In the **Certificates snap-in** dialog box, click **Computer account**, click **Next**, click **Finish**, click **Close**, and then click **OK**.
5. In the console tree, double-click **Certificates**, right-click **Personal**, point to **All Tasks**, and then click **Request New Certificate**.
6. Follow the instructions that appear in the Certificate Request Wizard. When prompted to type a friendly name for the certificate, type the name of the local Message Queuing server.
7. After you have obtained the certificate, open Computer Management.
8. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

9. Click **Properties**.
10. In the **Message Queuing Properties** dialog box, click the **Service Security** tab, and then in the **Server certificates** group box, click **Select a Certificate**.
11. Under **Issued to**, select the certificate with the name of the applicable Message Queuing server, and then click **OK**.
12. When prompted to restart Windows, click **Restart Now**.

**Notes**

- You can open the Certificates snap-in after you have added it to a console and saved the console settings in an .msc file as follows: click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click the name of the saved .msc file.
- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- This task can be performed only from a Message Queuing server installed on a Windows Server 2003 family domain controller.
- You need to create a server certificate only for each Message Queuing server that replies to Active Directory queries from MSMQ clients running on Windows NT 4.0, Windows 95, or Windows 98.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authentication for Message Queuing
Import server certificates for Message Queuing
Working with MMC console files

# Import server certificates for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import server certificates for Message Queuing**

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. In the **Add/Remove Snap-in** dialog box, click **Add**, select **Internet Information Services**, click **Add**, click **Close**, and then click **OK**.
4. In the console tree, right-click **MSMQ**.

   **Where?**

   - Console Root/Internet Information Services/*YourComputer*/Web Sites/Default Web Site/MSMQ

5. Click **Properties**.
6. Click the **Directory Security** tab, and then click **Server Certificate**. In the Web Server Certificate Wizard, click **Next**.
7. Click **Import a certificate from a Key Manager backup file**, and then click **Next**.
8. Enter the path and file name for the backup file you previously created when you exported the certificate from Windows NT 4.0, and then click **Next**.
9. Type the proper password, click **Next**, and then click **Finish**.
10. After you have obtained the certificate, open Computer Management.
11. In the console tree, right-click **Message Queuing**.

    **Where?**

    - Computer Management/Services and Applications/Message Queuing

12. Click **Properties**.
13. In the **Message Queuing Properties** dialog box, click the **Service Security** tab, and then in the **Server certificates** group box, click **Select a Certificate**.
14. Under **Issued to**, select the certificate with the name of the Message Queuing server to which you imported the certificate, and then click **OK**.

**Notes**

- You can also open the Internet Information Services snap-in as follows: click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Internet Services Manager**. Or, you can open this snap-in after you have added it to a console and saved the console settings in an .msc file as follows: click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click the name of the saved .msc file.

- To open Computer Management, right-click **My Computer**, and then click **Manage**.

- This task applies only if you are upgrading an MSMQ 1.0 controller server to a Message Queuing server installed on a Windows Server 2003 family domain controller.

- You must first export the key certificate before you upgrade the computer to Windows Server 2003 family. Then you can import the server certificate. You need to import a server (key) certificate only for each Message Queuing server that will reply to Active Directory queries from MSMQ clients running on Windows NT 4.0, Windows 98, or Windows 95.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authentication for Message Queuing
Working with MMC console files

# Change server certificates for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change server certificates for Message Queuing**

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Click **Properties**.
4. In the **Message Queuing Properties** dialog box, click the **Service Security** tab, and then in the **Server certificate** group box, click **Select a Certificate**.
5. In the **Choose MSMQ Server Certificate** dialog box, select the applicable certificate, and then click **OK**.
6. When prompted to restart Windows, click **Restart Now**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- This task can be performed only from a Message Queuing server installed on a Windows Server 2003 family domain controller.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authentication for Message Queuing
Create server certificates for Message Queuing
Working with MMC console files

# Import a CA certificate from Microsoft Certificate Services

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To import a CA certificate from Microsoft Certificate Services**

1. On the sender computer requiring a CA certificate, open Internet Explorer.
2. In Internet Explorer, connect to http://*servername*/certsrv, where *servername* is the name of the server where the certification authority (CA) you want to access is located.
3. Click **Download a CA certificate, certificate chain, or CRL**.
4. Do one of the following:
   - If you want to trust all the certificates issued by this certification authority (CA), click **Install this CA certificate chain**.
   - If the certification authority has been renewed, you will have the choice of which version of the CA certificate you want to download:
   - Select the encoding method you want for the CRL: **DER** or **Base 64**.
   - Under **CA Certificate**, click the CA certificate you want to download, and then click **Download CA certificate** or click **Download CA certificate chain**.
   - In **File Download**, click **Open this file from its current location**, and then click **OK**.
   - When the Certificate dialog box appears, click **Install this certificate**.
   - In the Certificate Import wizard, click **Automatically select the certificate store based on the type of certificate**. This places the CA certificate in your Trusted Root store.
5. If you are finished using the Certificate Services Web pages, close Internet Explorer.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
HTTPS authentication
Working with MMC console files

# Administer Queues for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Queues for Message Queuing**

This section shows you how to:

- Create queues
- Allow only authenticated messages on queues
- Change the privacy level for queues
- Find queues
- Delete queues
- Purge all messages in a queue
- Take outgoing queues offline
- Pause or resume an outgoing queue
- Change queue labels
- Change queue type IDs
- Change the base priority of queues
- Specify or view queue multicast addresses and ports
- Set the message storage size for queues
- Manage queues on other computers
- Manage queues using Active Directory Users and Computers
- Administer Journals for Message Queuing
- Administer Queue Aliases for Message Queuing

# Create queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create queues**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue folder.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (such as Public Queues or Private Queues)

3. Point to **New**, and then click **Public Queue** (or **Private Queue**) to open the **New Object – Public Queue** dialog box (or the **New Private Queue** dialog box).
4. Under **Queue name**, type a name for the new queue. Or, if you want to create a queue to accept transactional messages, type a name and select the **Transactional** check box.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.
- You cannot create queues in the Outgoing Queues or System Queues folder. These queues are created automatically.
- The queue objects for the queues residing on a particular computer are child objects of the msmq object of the applicable computer. Thus, to create a queue, a user must have the Create All Child Objects permission for the msmq object under which the queue will be created.
- You can create public queues on remote computers using the Active Directory Users and Computers snap-in provided you have domain administrative permissions and are logged on using a domain user account.
- When you create a queue, you are prompted to specify the minimal information needed. Each queue has many more properties (attributes) that can be set. After a queue is created, the values of many additional attributes can be viewed and modified. For information on how to do this, see Related Topics.
- When you create a queue, the default permission is that everyone can send to this queue. For tighter security, you can change the default security permissions for the queue. You can also specify properties for the queue when you create it. For example, you can specify to accept authenticated messages only. For instructions, see Allow only authenticated messages on queues.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**

Administering Queues
Manage queues using Active Directory Users and Computers
Working with MMC console files

# Allow only authenticated messages on queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To allow only authenticated messages on queues**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (such as Public Queues or Private Queues)/*YourQueue*

3. Click **Properties**.
4. On the **General** page, select the **Authenticated** check box. Or, to accept any messages on the queue, clear the **Authenticated** check box.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- If you select the **Authenticated** check box, the queue accepts only authenticated messages.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Authentication for Message Queuing
Manage queues using Active Directory Users and Computers
Change the privacy level for queues
Working with MMC console files

# Change the privacy level for queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the privacy level for queues**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (such as Public Queues or Private Queues)/*YourQueue*

3. Click **Properties**.
4. On the **General** page, in **Privacy Level**, do one of the following:

   - To have the queue accept both encrypted and unencrypted messages, click **Optional** (the default).
   - To have the queue accept only encrypted messages, click **Body**.
   - To have the queue accept only unencrypted messages, click **None**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.
- Privacy levels are used to specify whether a queue accepts encrypted messages, unencrypted messages, or both by selecting the privacy level for the queue. If the queue privacy level is set to None, the queue accepts only unencrypted messages. If the queue privacy level is set to Body, the queue accepts only encrypted messages. If the queue privacy level is set to Optional, the queue accepts both unencrypted and encrypted messages. You can check whether a message sent to the queue was encrypted or unencrypted by using Active Directory Users and Computers (or programmatically by checking the PROPID_M_PRIV_LEVEL message property).

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Encryption for Message Queuing
Manage queues using Active Directory Users and Computers
Allow only authenticated messages on queues
Working with MMC console files

# Find queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To find queues**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click any folder or object shown.
4. Click **Find**.
5. In **Find**, click the arrow that opens the list, and then click **Message Queuing Queue**.
6. On the **Find Message Queuing Queue** page, in **Label** or **Type ID**, type the appropriate queue information, and then click **Find Now**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- If you leave **Label** and **Type ID** empty, a list of all queues will be displayed. If you enter information in both **Label** and **Type ID**, only those queues that satisfy both criteria will be displayed.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Queues
Manage queues using Active Directory Users and Computers
Working with MMC console files

# Delete queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete queues**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (such as Public Queues or Private Queues)/*YourQueue*

3. Click **Delete**.
4. When prompted, click **Yes** to delete the queue.

### Important

- When you delete a queue, all messages within the queue are also deleted.

### Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- You cannot delete queues in the Outgoing Queues or System Queues folder.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.
- When you delete a queue, the triggers associated with that queue are not deleted. If you later re-create a queue of the same name, any existing triggers will be used by this queue. However, this re-created queue will not be monitored by the Message Queuing Triggers service until the computer is restarted, or the associated triggers disabled, and then enabled again. For instructions on setting the enabled or disabled property for a trigger, see View or change properties of triggers and rules for triggers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Queues
Queues
Manage queues using Active Directory Users and Computers
Working with MMC console files

# Purge all messages in a queue

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To purge all messages in a queue**

1. Open Computer Management.
2. In the console tree, right-click **Queue messages** (or **Journal messages**).

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (such as Public Queues or Private Queues)/*YourQueue/*Queue messages/Journal messages

3. Point to **All Tasks**, and then click **Purge**.
4. When prompted, click **Yes**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- You cannot purge individual messages from a queue.
- To increase performance, messages in journals and dead-letter queues must be purged on a regular basis.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing messages
Administering Messages
Working with MMC console files

# Take outgoing queues offline

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To take outgoing queues offline**

1. Open Computer Management.
2. In the console tree, right-click **Message Queuing**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing

3. Point to **All Tasks**, and do one of the following:

   - To take Message Queuing offline from the network, click **Take Message Queuing Offline**.
   - To bring Message Queuing back online, click **Bring Message Queuing Online**.

4. When prompted, click **Yes**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- When you take Message Queuing offline from the network, you pause the operation of all outgoing queues on the local computer. This is useful when you communicate over a slow, expensive connection and want to defer sending messages in all outgoing queues on the computer to a later time.
- You can also pause the operation of a single outgoing queue. For more information, see Related Topics.
- This task can be performed only by using the Computer Management snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Queues
Start or stop services for Message Queuing
Working with MMC console files

# Pause or resume an outgoing queue

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To pause or resume an outgoing queue**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/Outgoing Queues/*YourOutgoingQueue*

3. Point to **All Tasks**.
4. Click **Pause**. Or, to resume the operation of the queue, click **Resume**.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- When you pause the operation of an outgoing queue, you postpone the sending of messages from that queue to the corresponding destination queue. This way, you can send all pending messages to a destination queue at a later time.
- You can also pause the operation of all outgoing queues by taking Message Queuing offline from the network. For more information, see Related Topics.
- The icons of outgoing queues that are offline are marked by an X on a circular red background in the Computer Management console tree.
- This task can be performed only by using the Computer Management snap-in.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Queues
Start or stop services for Message Queuing
Working with MMC console files

# Change queue labels

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change queue labels**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (Public Queues or Private Queues)/*YourQueue*

3. Click **Properties**.
4. On the **General** page, in **Label**, type the new label for the queue.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- A queue label is a string with a maximum length of 124 Unicode characters that can be used to identify the queue. For example, queues that accept messages sent by a specific application may all have the label "Messages from my application."
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Queues
Manage queues using Active Directory Users and Computers
Working with MMC console files

# Change queue type IDs

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change queue type IDs**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (Public Queues or Private Queues)/*YourQueue*

3. Click **Properties**.
4. On the **General** page, in **Type ID**, type the new queue type ID.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- The queue type ID specifies the function or type of service provided by the queue. The queue type ID must be a valid GUID in the form {*dddddddd-dddd-dddd-dddd-dddddddddddd*}, where *d* is a hexadecimal text digit. The use of the queue type ID is optional and depends on the requirements of the Message Queuing application.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Queues
Manage queues using Active Directory Users and Computers
Working with MMC console files

# Change the base priority of queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change the base priority of queues**

1. Open Computer Management.
2. In the console tree, right-click the applicable public queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/Public Queues/ *YourQueue*

3. Click **Properties**.
4. On the **General** page, in **Base Priority**, type a new value or use the UP ARROW or DOWN ARROW keys to enter a new value.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- The base priority can be set from -32768 to 32767. The default priority setting is 0.
- The base priority (also known as queue priority) can be set only for public queues. Private queues do not support base priority.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Queues
Manage queues using Active Directory Users and Computers
Working with MMC console files

# Specify or view queue multicast addresses and ports

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To specify or view queue multicast addresses and ports**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (such as Public Queues or Private Queues)/*YourQueue*

3. Click **Properties**, and then click **Multicast**.
4. Under **Multicast address**, type the multicast address and port number in the *IP_Address*:*PortNumber* syntax.

**Notes**

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- A multicast address and a port number can be associated with any public or private queue. However, in workgroup mode, multicast addresses can be associated only with private queues, and in domain mode, a Windows Server 2003 family domain controller must be available to associate multicast addresses with public queues.
- No spaces or leading zeros are allowed in the multicast address that you type.
- IP multicast addresses must consist of four bytes and must lie in the class D range from 224.0.0.0 to 239.255.255.255. However, only certain ranges of addresses in this range are available for sending multicast messages. For more information, see Related Topics.
- You can also specify or view the multicast addresses and port numbers of queues using the Active Directory Users and Computers snap-in. For more information, see Related Topics.
- After setting the multicast properties of a remote queue, it may take up to 15 minutes for the change to take effect.
- On a computer with multiple NICs, a problem will arise when the computer tries to send multicast messages by means of these multiple NICs. In effect, Message Queuing will randomly choose a NIC for sending multicast messages. If the private NIC is chosen, multicast recipients will not receive the message. To work around this issue, you can specify the source IP address to be used for sending multicast messages. To do this, create the string value **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\MulticastBindIP**, in the usual IP address format. After setting the registry, restart the Message Queuing service for the changes to take effect. For more information, see IP multicasting.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

- Be aware of the security implications of multicasting. Messages cannot be encrypted, and messages could be read or used by unauthorized persons. Ensure that the address specified is used only for Message Queuing multicasting.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Multiple Destination Messaging
Working with MMC console files

# Set the message storage size for queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the message storage size for queues**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (such as Public Queues or Private Queues)/*YourQueue*

3. Click **Properties**.
4. Select the **Limit message storage to (KB)** check box, and then type the total size, in kilobytes, of all queue messages that can be stored in the queue.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- Any changes made will not take effect until the Message Queuing service is stopped and restarted.
- The message storage size limit, or quota, for a queue specifies the maximum cumulative size of all messages contained in the applicable queue. This quota can be any integer between 0 and 4,294,967,294.
- When the message storage size limit for a destination queue is reached, new incoming messages are discarded, and a negative acknowledgment message indicating that the destination queue is full is generated. Messages can be placed in the queue only after one or more messages are removed from the queue.
- If you clear the **Limit message storage to (KB)** check box, the message size will be limited only by the available disk space on the local computer.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Setting message quotas
Manage queues using Active Directory Users and Computers
Set the journal storage size for queues
Working with MMC console files

# Manage queues on other computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage queues on other computers**

1. Open Computer Management.
2. Right-click **Computer Management**, and then click **Connect to another computer**.
3. In the **Select Computer** dialog box, select the computer on which you want to manage queues, and then click **OK**.
4. Perform the desired task as you would for the local computer.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- To manage public queues, you must have access to Active Directory.
- You cannot manage private queues located in a remote computer.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Queues
Working with MMC console files

# Manage queues using Active Directory Users and Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage queues using Active Directory Users and Computers**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click the applicable queue.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer/*msmq/*YourQueueFolder* (such as Private Queues for a private queue)/*YourQueue/YourMessage*

4. Complete the desired task as you would using Computer Management.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- This method can be used only if the local computer belongs to a Windows Server 2003 family or Windows 2000 domain and can contact a domain controller.
- You can create, delete, and change the properties of public queues on remote computers using the Active Directory Users and Computers snap-in only if you have domain administrative permissions and are logged on using a domain user account.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Queues
Working with MMC console files

# Administer Journals for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Journals for Message Queuing**

This section shows you how to:

- Enable or disable target journaling
- Set the journal storage size for queues

# Enable or disable target journaling

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable or disable target journaling**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (Public Queues or Private Queues)/*YourQueue*

3. Click **Properties**.
4. On the **General** page, in **Journal**, select the **Enabled** check box. Or, to disable target journaling for the queue, clear the **Enabled** check box.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- Target journaling is enabled or disabled on a queue basis.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Target journaling
Working with MMC console files

# Set the journal storage size for queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the journal storage size for queues**

1. Open Computer Management.
2. In the console tree, right-click the applicable queue.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (Public Queues or Private Queues)/*YourQueue*

3. Click **Properties**.
4. On the **General** page, in **Journal**, select the **Limit journal storage to (KB)** check box, and then type the total size (in kilobytes) of all journal messages that can be stored in the queue journal.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- The journal storage size limit, or quota, for a queue specifies the maximum cumulative size of all journal messages contained in the applicable queue journal. This quota can be any integer between 0 and 4,294,967,294.
- The journal storage size limit is applicable only when target journaling is enabled.
- When the journal storage size limit is reached, additional messages can no longer be placed in the queue journal until the cumulative size of the journal messages in the queue journal drops below the specified limit.
- If you clear the **Limit journal storage to (KB)** check box, the journal size will be limited by the computer quota of the local computer.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Source journaling
Working with MMC console files

# Administer Queue Aliases for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Queue Aliases for Message Queuing**

This section shows you how to:

- Create queue aliases
- Add descriptions to queue aliases
- View or change properties of queue aliases
- Delete queue aliases

# Create queue aliases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create queue aliases**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then on the **View** menu, click **Advanced Features**.
3. In the console tree, right-click the applicable container.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourContainer* (such as Computers or Domain Controllers)
4. Point to **New**, and then click **Message Queuing Queue Alias** to open the **New Object - Message Queuing Queue Alias** dialog box.
5. Under **Queue alias name**, type the alias.
6. Under **Format name**, type the applicable format name.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- The format name typed can be any public, private, or direct single-element format name. The validity of the format name typed and the existence of the queues to which the queue alias points are not tested during this procedure. For examples of valid format names, see **Queue Names** in Related Topics.
- Multiple aliases can be created for the same physical queue.
- When you create a queue alias in Active Directory, you are prompted to specify the minimal information needed to establish the object in Active Directory. Each queue alias has many more properties (attributes) that can be set. After a queue alias is created, the values of many additional attributes can be viewed and modified. For information on how to do this, see Related Topics.
- A queue alias can be used to point to a private queue and publish a private queue. For best security practice, create a queue alias with the narrowest security permissions possible.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Queue aliases
Working with MMC console files

# Add descriptions to queue aliases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To add descriptions to queue aliases**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then on the **View** menu, click **Advanced Features**.
3. In the console tree, click the applicable container.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourContainer* (such as Computers or Domain Controllers)
4. In the details pane, double-click the applicable queue alias or right-click the queue alias, and then click **Properties**.
5. Under **Description**, type a description for the alias queue.

## Notes

- To see the new queue alias description, press F5 to refresh the pane.
- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Queue aliases
Working with MMC console files

# View or change properties of queue aliases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view or change properties of queue aliases**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then on the **View** menu, click **Advanced Features**.
3. In the console tree, click the applicable container.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourContainer* (such as Computers or Domain Controllers)
4. In the details pane, double-click the applicable queue alias or right-click the queue alias, and then click **Properties**.
5. View or change the queue alias properties as needed.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- The format name typed can be any public, private, or direct single-element format name. The validity of the format name typed and the existence of the queues to which the queue alias points are not tested during this procedure.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Queue aliases
Working with MMC console files

# Delete queue aliases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete queue aliases**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then on the **View** menu, click **Advanced Features**.
3. In the console tree, click the applicable container.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourContainer* (such as Computers or Domain Controllers)
4. In the details pane, click the applicable queue alias, and then press the DELETE key.

**Note**

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Queue aliases
Working with MMC console files

# Administer Distribution Lists for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Distribution Lists for Message Queuing**

This section shows you how to:

- Create distribution lists
- View, add, or delete elements in a distribution list
- View or change properties of distribution lists
- Get a list of distribution lists containing a given element
- Delete distribution lists

# Create distribution lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create distribution lists**

1. Open Active Directory Users and Computers.
2. In the console tree, double-click the domain node.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*

3. Right-click the folder in which you want to add the group, point to **New**, and then click **Group**.
4. In the **New Object – Group** dialog box, type the name of the distribution list under **Group Name**, select the group scope that you want to use, and then select the group type **Distribution**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- You can choose to give your distribution list universal, global, or domain local scope, as follows:
  - **Universal**. Members of universal groups can include other groups and accounts from any domain in the domain tree or forest and can be granted permissions in any domain in the domain tree or forest.
  - **Global**. Members of global groups can include other groups and accounts only from the domain in which the group is defined and can be granted permissions in any domain in the forest.
  - **Domain local**. Members of domain local groups can include other groups and accounts from Windows Server 2003, Windows 2000, or Windows NT domains and can be granted permissions only within a domain.
- Message Queuing distribution lists are distribution group objects in Active Directory.
- When you create a distribution list (group) in Active Directory, you are prompted to specify the minimal information needed to establish the object in Active Directory. Each distribution list (group) has many more properties (attributes) that can be set. After a distribution list (group) is created, the values of many additional attributes can be viewed and modified. For information on how to do this, see Related Topics.
- For information on how to add and remove members from distribution lists (groups), see Related Topics.
- With distribution lists, a single operation sends the same message to multiple destinations, and does so by duplicating packets from senders to receivers. For best security practice, consider security permissions for the distribution list object in Active Directory — who is authorized to manage it and use it.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Distribution lists

Working with MMC console files

# View, add, or delete elements in a distribution list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view, add, or delete elements in a distribution list**

1. Open Active Directory Users and Computers.

2. In the console tree, double-click the domain node.

3. Click the folder that contains the distribution group.

4. In the console tree, right-click the applicable distribution list (group).

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/Folder/*YourDistributionList(Group)*

5. Click **Properties**.

6. On the **Members** page, you can view the list of members in the distribution list (group) selected. The explicit members of a distribution list can be public queues, queue aliases, and other distribution lists. You can also delete elements (members) from the distribution list (group) as follows: click the applicable element (member), click **Remove**, and then click **OK**.

7. To add a public queue to a distribution list (group), do the following:

   In the console tree, right-click the applicable queue (console tree location is: Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer*/msmq/*YourPublicQueue*).

   - Click **Properties**.

   - On the **Member of** page, click **Add**.

   - On the **Select Groups** dialog box, in **Name**, type the name of a distribution list (group) or the names of several distribution lists (groups) separated by semicolons, click **OK**, and then click **OK** again. Or click **Advanced** to search for distribution lists (groups), enter the applicable parameters, click **Find Now**, select the applicable distribution lists (groups), click **OK**, click **OK** again, and then click **OK** a third time.

8. To add private queues or URL-named queues to a distribution list (group), do the following:

   - Create an alias for each private queue or URL-named queue that you want to add to the distribution list (group). For information on how to create an alias for a private queue or a URL-named queue, see Related Topics.

   - In the console tree, right-click the distribution list (group) to which you want to add a queue alias or queue aliases (console tree location is: Active Directory Users and Computers/*YourDomain*/Users/*YourDistributionList(Group)*).

   - Click **Properties**.

   - On the **Members** page, click **Add**. In the **Select Groups** dialog box, click **Look For**, select the **Other objects** check box, clear the remaining check boxes, and click **OK**. In **Name**, type the name of a queue alias or the names of

several queue aliases separated by semicolons, click **OK**, and then click **OK** again. Or click **Advanced** to search for queue aliases, enter the applicable parameters, click **Find Now**, select the queue alias, click **OK**, click **OK** again, and then click **OK** a third time.

9. To add a queue alias to one or several distribution lists (groups), do the following:

   - In the console tree, right-click the applicable queue alias (console tree location is: Active Directory Users and Computers/*YourDomain*/*YourContainer* (such as Computers or Domain Controllers)/*YourAlias*).
   - Click **Properties**.
   - On the **Member of** page, click **Add**.
   - On the **Select Groups** dialog box, in **Name**, type the name of a distribution list (group) or the names of several distribution lists (groups) separated by semicolons, click **OK**, and then click **OK** again. Or click **Advanced** to search for distribution lists (groups), enter the applicable parameters, click **Find Now**, select the applicable distribution lists (groups), click **OK**, click **OK** again, and then click **OK** a third time.

10. To add a distribution list (group) as a member of another distribution list (group), do the following:

    - In the console tree, right-click the distribution list (group) that you want to be a new member of another distribution list (group) (console tree location is: Active Directory Users and Computers/*YourDomain*/Users/*YourDistributionList(Group)*).
    - Click **Properties**.
    - On the **Member of** page, click **Add**.
    - On the **Select Groups** dialog box, in **Name**, type the name of a distribution list (group) or the names of several distribution lists (groups) separated by semicolons, click **OK**, and then click **OK** again. Or click **Advanced** to search for distribution lists (groups), enter the applicable parameters, click **Find Now**, select the applicable distribution lists (groups), click **OK**, click **OK** again, and then click **OK** a third time.

### Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Message Queuing distribution lists are distribution group objects in Active Directory.
- The explicit members of distribution lists can be public queues, queue aliases, and other distribution lists, but not private queues or URL-named queues. Private queues and URL-named queues can be included in a distribution list by using queue aliases. For information on how to create an alias for a private queue or a URL-named queue, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Distribution lists
Working with MMC console files

# View or change properties of distribution lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view or change properties of distribution lists**

1. Open Active Directory Users and Computers.
2. In the console tree, double-click the domain node.
3. Click the folder that contains the distribution group.
4. In the console tree, right-click the applicable distribution list (group).

   **Where?**

   - Active Directory Users and Computers/*YourDomain/*Folder/*YourDistributionList(Group)*

5. Click **Properties**.
6. View or change the distribution list (group) properties as needed.

| **Property** | **Tab** |
|---|---|
| Group name | **General** |
| Description | **General** |
| Group scope | **General** |
| Group type | **General** |
| Distribution list (group) elements (members) | **Members** |
| Folders containing the members in Active Directory | **Members** |
| Distribution lists (groups) of which the applicable distribution list (group) is a member | **Member of** |
| Administrative permissions | **Security** |

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Message Queuing distribution lists are distribution group objects in Active Directory.
- The explicit members of distribution lists can be public queues, queue aliases, and other distribution lists, but not private queues or URL-named queues. Private queues and URL-named queues can be included in a distribution list by using queue aliases.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Distribution lists

Working with MMC console files

# Get a list of distribution lists containing a given element

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To get a list of distribution lists containing a given element**

1. Open Active Directory Users and Computers.
2. In the console tree, double-click the domain node.
3. Click the folder that contains the distribution group.
4. In the console tree, do one of the following:
   - Right-click the applicable distribution list (group)

     Console tree location is: Active Directory Users and Computers/*YourDomain*/Folder/*YourDistributionList(Group)*)
   - Right-click the applicable public queue

     Console tree location is: Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer*/msmq/*YourPublicQueue*)
   - Right-click the applicable queue alias

     Console tree location is: Active Directory Users and Computers/*YourDomain*/*ContainerForAliases*/*YourAlias*)
5. Click **Properties**.
6. On the **Member of** page, you can view the list of distribution lists (groups) of which the applicable object is a member. For information on how to add the object selected to another distribution list (group), see Related Topics. You can remove the applicable object from a distribution list (group) of which it is a member as follows: click the name of the host distribution list (group), and then click **Remove**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Message Queuing distribution lists are distribution group objects in Active Directory.
- The explicit members of distribution lists can be public queues, queue aliases, and other distribution lists, but not private queues or URL-named queues. Private queues and URL-named queues can be included in a distribution list by using queue aliases.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Distribution lists

# Delete distribution lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete distribution lists**

1. Open Active Directory Users and Computers.
2. In the console tree, double-click the domain node.
3. Click the folder that contains the distribution group.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/*Folder

4. In the details pane, right-click the applicable distribution list (group), and then click **Delete** on the menu. Alternatively, you can click the applicable distribution list (group), and then press **Delete**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- Message Queuing distribution lists are distribution group objects in Active Directory.
- When a distribution list is deleted, it is automatically removed from all the distribution lists of which it is a member.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Distribution lists
Working with MMC console files

# Administer Messages for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Messages for Message Queuing**

This section shows you how to:

- View message properties
- Set the default lifetime for messages
- Manage messages on other computers
- Manage messages using Active Directory Users and Computers

# View message properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view message properties**

1. Open Computer Management.
2. In the console tree, right-click **Queue messages** (or **Journal messages**).

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/*YourQueueFolder* (such as Public Queues or Private Queues)/*YourQueue/*Queue messages/Journal messages

3. In the details pane, right-click a message, and then click **Properties**.
4. Click each tab to view the various message properties and other relevant information.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Messages
Working with MMC console files

# Set the default lifetime for messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To set the default lifetime for messages**

1. Open Active Directory Sites and Services.
2. On the **View** menu, click **Show Services Node**.
3. In the console tree, right-click **MsmqServices**.

   **Where?**

   - Active Directory Sites and Services/Services/MsmqServices

4. Click **Properties**.
5. On the **General** page, type a new value and select new units as needed.

### Notes

- To open Active Directory Sites and Services, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Sites and Services**.
- Note that the default lifetime for transactional messages is 4 days.
- The default lifetime of a message is also known as the time-to-reach-queue (TTRQ).

### Important

- For best security practice, it is recommended that messages not be set with an infinite lifetime to ensure tighter control over messages that are traveling through the network. Also, MSMQ can request that a negative acknowledgment message is returned if a timer expires, and Message Queuing sends the appropriate negative acknowledgment message to the administration queue specified by the message.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Messages
Working with MMC console files

# Manage messages on other computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage messages on other computers**

1. Open Computer Management.
2. Right-click **Computer Management**, and then click **Connect to another computer**.
3. In the **Select Computer** dialog box, select the computer on which you want to manage messages, and then click **OK**.
4. Perform the desired task as you would for the local computer.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- You can also perform this task using the Active Directory Users and Computers snap-in. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Managing messages
Working with MMC console files

# Manage messages using Active Directory Users and Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To manage messages using Active Directory Users and Computers**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click the applicable message.

   **Where?**

   - Active Directory Users and Computers/*YourDomain/YourDomain/YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer/*msmq/*YourQueueFolder* (such as Private Queues for a private queue)/*YourQueue/YourMessage*

4. Complete the desired task as you would using Computer Management.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- This method can be used only if the local computer belongs a Windows Server 2003 family or Windows 2000 domain and can contact a domain controller.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Messages
Working with MMC console files

# Administer Triggers for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administer Triggers for Message Queuing**

This section provides instructions for creating and managing triggers:

- Create triggers
- Create and attach rules to triggers
- View or change properties of triggers and rules for triggers
- Delete triggers and rules for triggers
- Change default configuration values for triggers
- Configure trigger resources for server clusters

# Create triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create triggers**

1. Open Computer Management.
2. In the console tree, right-click **Triggers**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/Triggers/Triggers

3. Point to **New**, and then click **Trigger**.
4. In **Trigger name** on the first page of the **New Trigger** dialog box, type a name for the new trigger.
5. If you want to specify a user-defined queue as the *monitored queue*, in **Queue path name**, type the path name of the queue with which the trigger will be associated in the *ComputerName**QueueName* syntax, or the .*QueueName* syntax for the local computer. (Do not use format names.) Otherwise, click the name of an automatically created queue on your computer.
6. Click **Peeking**, **Retrieval**, or **Transactional retrieval** to specify the type of message processing.
7. Click the applicable check box to configure the trigger as enabled or disabled.
8. Click the applicable check box to configure the trigger as serialized or not serialized, and then click **Next**.
9. On the second page of the **New Trigger** dialog box, to attach an existing rule to the trigger, select a rule in the **Rules** list, and then click **Attach**
10. To remove a rule from a trigger, select a rule in the **Attached Rules** list, and then click **Detach**.
11. To change the order in which rules are tested for the triggers, click the **Up** and **Down** buttons.
12. When you have the list of attached rules that you want, click **Finish**.

**Notes**

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- Triggers are not installed by default during Message Queuing installation. For information on installing triggers, see Installing Message Queuing triggers.
- This procedure can also be performed by opening the **New Trigger** dialog box from the Triggers folder under any monitored queue in the console tree as follows: right-click **Triggers**, point to **New**, and then click **Trigger**. In this case, the queue automatically becomes the monitored queue.
- Regardless of where the creation of a trigger is initiated, all the triggers created on the local computer, which is called the *trigger set*, are listed in the details pane of the Computer Management snap-in when **Triggers** is selected in the console tree under the Triggers folder. By default, for each trigger, the following is displayed in the details pane:

    - Trigger name
    - Name of the monitored queue
    - Number of rules attached to the trigger
    - Whether the trigger is enabled
    - Whether the trigger is serialized
    - Message processing type

- Before creating a trigger, verify that the account under which the Message Queuing Triggers service is running has been granted the Peek Message permission and the Receive Message permission for all existing queues that will be associated with triggers having the retrieval or transactional-retrieval processing type. For more information, see Related Topics.

- After you type the name of a user queue to which the trigger will be associated in **Queue path name**, the existence of the queue is verified. If a queue with the name typed cannot be found, an error message appears.

- Triggers can be serialized or nonserialized. A serialized trigger tests the conditions of all the rules associated with it on each message and performs all the applicable actions before moving on to the next message in the order in which the messages arrive. However, the throughput is lower than for nonserialized triggers because the messages are processed one at a time. Nonserialized triggers permit higher throughput but may fail to perform a given operation on the messages in the order in which they arrive. This occurs, for example, when multiple messages are processed simultaneously. If you create multiple triggers for a particular queue and at least one of them is serialized, all the triggers will behave as serialized triggers, even if they are not defined as serialized. Triggers with the transactional-retrieval processing type can only be serialized. If message processing order is significant, it is recommended that you use only a single trigger per queue and enable the serialized option, because the order of trigger firing is random with respect to their GUIDs.

- By default, the Message Queuing Triggers service runs under the Network Service account. Note that when the Triggers service is running under the Network Service account, it does not interact with the desktop. If this functionality is required, you can configure the Triggers service to run under the Local System account and enable interaction with the desktop in Computer Management. For best security practice, do not change the default account to a less secure account unless absolutely necessary. The Triggers service must run under an account with administrator permissions.

- Executable files that are invoked by triggers are invoked with the same security privileges as those assigned to the Triggers service.

- After a trigger has been created, rules can be attached to it. For more information, see Related Topics.

- When you delete a queue, the triggers associated with that queue are not deleted. If you later re-create a queue of the same name, any existing triggers will be used by this queue. However, this re-created queue will not be monitored by the Message Queuing Triggers service until the computer is restarted, or the associated triggers disabled, and then enabled again. For instructions on setting the enabled or disabled property for a trigger, see View or change properties of triggers and rules for triggers.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Triggers
Triggers overview

# Create and attach rules to triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To create and attach rules to triggers**

1. Open Computer Management.
2. In the console tree, right-click **Rules**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/Triggers/Rules

3. Point to **New**, and then click **Rule**.
4. On the first page of the **New Rule** dialog box, in **Rule name**, type a name for the new rule, and in **Rule description**, type an optional description.
5. Click **Next**.
6. In **New condition**, enter a condition by selecting the initial string in the list box, typing the applicable final string in the **Filter value/string** text box below, and clicking **Add**. The complete condition then appears in **Condition list**. You can create additional conditions by repeating this operation.
7. When you finish creating the conditions that you want to create, click **Next**.
8. To specify whether the trigger will invoke a stand-alone executable or a COM component (the default), click the applicable button.
9. To specify that the trigger invoke a COM component, in **COM component invocation**, select **Invoke COM component**. In **Component ProgID**, type the ID of the COM component, and in **Method name**, type the name of the method.
10. To specify that the trigger invoke a stand-alone executable, in **Standalone executable invocation**, select **Invoke standalone executable(EXE)**. In **Executable path**, type the full path of the executable file.
11. To add parameters for either a COM component or a stand-alone executable, click **Parameters**. Select a parameter in the **Parameter** list box. If the parameter is a numeric literal or a string literal, type its value in **Literal value**. Click **Add** to add the parameter to the **Invocation parameters** list. Repeat this operation for each parameter that you want to add, and use the **Up** and **Down** buttons to change the order in which the parameters are transferred to the program. Click **OK**, and then click **Finish**.
12. To attach existing rules to a trigger, do the following:

    - In the console tree, click **Triggers**.

      Console tree location is: Computer Management/Services and Applications/Message Queuing/Triggers

      Or, click **Triggers** under the applicable monitored queue.

      Console tree location is: Computer Management/Services and Applications/Message Queuing/ *YourQueueFolder* (Public Queues, Private Queues, or System Queues)/ *YourMonitoredQueue*/Triggers

- In the details pane, right-click the applicable trigger, click **Properties**, and then click the **Attached Rules** tab. Select the rule you want to attach to the trigger from the **Rules** list, and then click **Attach** to move it to the **Attached Rules** list. Then click **OK**.

### Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.

- All the rules created on the local computer, which is called the *rule set*, are listed in the details pane of the Computer Management snap-in when **Rules** is selected in the console tree.

- A rule is composed of one or more conditions and an action that is executed when the conditions are true. Each condition is applied to a message when it arrives at the monitored queue, and the result is either true or false. After being created, the conditions applied to a particular rule are listed in **Condition list** on the second page of the **New Rule** dialog box as well as on the **Conditions** page of the **Properties** dialog box for that rule. A condition can be eliminated from this list by selecting it and clicking **Remove**. Multiple conditions must all be true for an action to be executed. If no condition is specified, a true value will always be returned. This is useful when you want a trigger to fire unconditionally every time a message arrives at the monitored queue.

- The actions performed when a condition is true are expressed either in the form of an executable file and parameters (arguments) or in the form of a COM component, a method, and parameters. You can also invoke other types of files by making them parameters of an .exe file. For example, you can invoke an .msc file by specifying mmc.exe as the executable file and including the full path of the .msc file as an invocation parameter of the string literal type. You can specify one or more parameters to be passed per action. The supported parameters are listed in the list box on the **Invocation Parameters** dialog box.

- If your trigger invokes an executable file with command-line arguments, these arguments can be included as parameters of the string literal type. However, you must add a separate parameter for each command-line argument.

- You can attach multiple rules to a trigger. In this case, the rules are executed in the order in which they are displayed on the **Attached Rules** page of the **Properties** dialog box for that rule. You can change the priority (order of execution) of the rules in the list. To do so, use the **Up** and **Down** buttons on this page to change the order of priority of each rule as needed to give the rule a higher or lower priority, respectively. You can also detach rules on this page by selecting each rule that you want to detach and clicking **Detach**. For more information, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Triggers
Trigger rules
Working with MMC console files

# View or change properties of triggers and rules for triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view or change properties of triggers and rules for triggers**

1. Open Computer Management.
2. In the console tree, click **Triggers** or **Rules**, whichever is applicable.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/Triggers/Triggers (or Rules)

3. In the details pane, right-click the applicable trigger or rule, and then click **Properties**.
4. View or change trigger properties as needed. Note that trigger names cannot be changed.

| Trigger Property | Tab |
|---|---|
| Trigger name (view only) | **General** |
| Monitored queue (path name for a user queue) | **General** |
| Message processing type | **General** |
| Enabled or disabled | **General** |
| Serialized or nonserialized | **General** |
| Rules attached to the trigger | **Attached Rules** |
| Rules available for the trigger | **Attached Rules** |

Or, view or change rule properties as needed. Note that rule names cannot be changed.

| Rule property | Tab |
|---|---|
| Rule name (view only) | **General** |
| Rule description | **General** |
| Conditions applied to the rule | **Rule Condition** |
| Whether the rule launches a stand-alone executable or a COM component | **Rule Action** |
| Path of a stand-alone executable | **Rule Action** |
| COM component ProgID | **Rule Action** |
| Method chosen for a COM component | **Rule Action** |
| Parameters used and their order. For types of parameters, see Trigger rules. | **Invocation Parameters** |

**Notes**

- To open Computer Management, right-click **My Computer**, and then click **Manage**.

- Triggers are not installed by default during Message Queuing installation. For information on installing triggers, see Installing Message Queuing triggers.

- The **Properties** dialog box for a specific trigger has **General** and **Attached Rules** tabs. The **Properties** dialog box for a specific rule has **General**, **Rule Condition**, and **Rule Action** tabs, and there is an **Invocation Parameters** dialog box, which can be accessed from the **Rule Action** page.

- The monitored queue of an existing trigger cannot be changed.

- When you delete a queue, the triggers associated with that queue are not deleted. If you later re-create a queue of the same name, any existing triggers will be used by this queue. However, this re-created queue will not be monitored by the Message Queuing Triggers service until the computer is restarted, or the associated triggers disabled, and then enabled again.

- A rule consists of one or more conditions and an action that is executed when the conditions are true. Each condition is applied to a message when it arrives at the monitored queue, and the result is either true or false. The conditions applied to a particular rule are listed in **Condition list** on the **Rule Condition** page of the **Properties** dialog box for that rule. A condition can be eliminated from this list by selecting it and clicking **Remove**. New conditions are composed by selecting a supported initial string in the **New condition** list box and typing a concluding string in the text box below it. Multiple conditions must all be true for an action to be executed. If no condition is specified, a true value will always be returned. This is useful when you want a trigger to fire unconditionally every time a message arrives at the monitored queue.

- The actions performed when the conditions are true are expressed either in the form of an executable file and parameters (arguments) or in the form of a COM component, method, and parameters. You can specify one or more arguments to be passed per action. The supported parameters are listed in the list box on the **Invocation Parameters** dialog box, which can be accessed by clicking **Parameters** on the **Rule Action** page of the **Properties** dialog box for the particular rule.

- The rules displayed on the **Attached Rules** tab are disabled by default, and you must attach them to the trigger to enable them. You can attach multiple rules to a trigger. In this case, the rules are executed in the order in which they are displayed in the **Attached Rules** list. You can change the priority (order of execution) of the rules in the list. To do so, use the **Up** and **Down** buttons on this page to give the rule a higher or lower priority, respectively. You can also detach rules on this page by selecting each rule that you want to detach and clicking **Detach**. For more information, see See Also.

- All triggers and rules for triggers created on the local computer are displayed in the details pane of the Computer Management snap-in when the applicable folder is selected in the console tree.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Triggers
Triggers overview
Trigger rules

# Delete triggers and rules for triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To delete triggers and rules for triggers**

1. Open Computer Management.
2. In the console tree, click **Triggers** or **Rules**, whichever is applicable.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/Triggers/Triggers (or Rules)

3. In the details pane, right-click the applicable trigger or rule, and then click **Delete**. Alternatively, you can click the applicable trigger or rule, and then press the DELETE key.

## Notes

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- All triggers and rules for triggers created on the local computer are displayed in the Computer Management snap-in.
- To detach rules from a trigger without deleting them, right-click the applicable trigger in the details pane, and then click **Properties**. Click the **Attached Rules** tab, and then select each rule that you want to detach and click **Detach**. For more information, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Triggers
Triggers overview
Trigger rules
Working with MMC console files

# Change default configuration values for triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To change default configuration values for triggers**

1. Open Computer Management.
2. In the console tree, click to expand **Message Queuing**, and then right-click **Triggers**.

   **Where?**

   - Computer Management/Services and Applications/Message Queuing/Triggers

3. Click **Properties**, and then change the values in **Initial number of threads**, **Max number of threads**, and **Default message body size** as needed.

**Notes**

- To open Computer Management, right-click **My Computer**, and then click **Manage**.
- All configuration information as well as trigger and rule definitions are stored in the local registry.
- The **Default message body** value is the size of the buffer that the Message Queuing Triggers service will allocate for an incoming message. It is recommended that this value be set equal to a size that will accommodate 90% of all messages that are expected to arrive in monitored queues. If this value is set too small, it will force the service to reallocate these buffers for every message. If it is set too large, an excessive amount of memory will be allocated.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Administering Triggers
Triggers overview
Modifying default trigger values
Working with MMC console files

# Configure trigger resources for server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To configure trigger resources for server clusters**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.

2. To create a Message Queuing Triggers resource:

   - When the **Open Connection to Cluster** dialog box appears, in the **Cluster or server name** list box, select the applicable cluster or server. (If you select a server, the name of the cluster to which it belongs will appear in the console tree.)

   - In the console tree, right-click the cluster group to which you want the Message Queuing Triggers resource to belong (console tree location is: *ClusterName/*Groups/*YourGroup*).

   - On the **File** menu, point to **New**, and then click **Resource**.

   - In the New Resource Wizard, type the applicable information in **Name** and **Description**, select **Message Queuing Triggers** in the **Resource type** list box, and then click **Next**.

   - On the **Possible Owners** page, select the set of possible owners to match that of the Message Queuing resource in the same group, and then click **Next**.

   - On the **Dependencies** page, in **Available resources**, select the name of the Message Queuing resource in the group, and then click **Add**.

   - Select the name of the Network Name resource in the group, click **Add**, and then click **Finish**.

3. If Message Queuing has the Active Directory Integration component installed, enable Kerberos authentication for the Network Name resource on which the Message Queuing resources depend:

   - Bring the Network Name resource offline as follows: In the console tree for the cluster group, right-click the Network Name resource in the details pane, and then click **Take Offline**.

   - Right-click the Network Name resource again, and then click **Properties**.

   - On the **Parameters** tab, select the **Enable Kerberos Authentication** check box to enable authentication. The default setting is that Kerberos authentication is not enabled.

   - Bring the Network Name resource online as follows: In the console tree for the cluster group, right-click the Network Name resources in the details pane, and then click **Bring Online**.

   **Caution**

   - Other resources in the group may also depend on the Network Name resource. These will be taken offline automatically when the Network Name resource goes offline. When you bring the Network Name resource online, bring these other resources online too.

4. To bring the Message Queuing Triggers resource online, in the console tree, click the **Resources** folder. In the details

pane, click the Message Queuing Triggers resource. On the **File** menu, click **Bring Online**.

5. To manage the Message Queuing Triggers resource from Computer Management, create a Generic Application resource in the same cluster group, and configure it to launch the Computer Management MMC snap-in:

   - In the Cluster Administrator console tree, click the cluster group to which the Message Queuing Triggers resource belongs.
   - On the **File** menu, point to **New**, and then click **Resource**. The New Resource Wizard launches.
   - In the New Resource Wizard, type **Computer Management** in the **Name** text box and select **Generic Application** in the **Resource type** list box. Then, click **Next**.
   - On the **Possible Owners** page, select possible owners to match those selected for the Message Queuing Triggers resource. Then, click **Next**.
   - On the **Dependencies** page, under **Available resources**, select the name of the Network Name resource in the group and click **Add**, and then select the Message Queuing Triggers resource previously created and click **Add**. Then, click **Next**.
   - On the **Generic Application Parameters** page, in **Command line**, type **mmc compmgmt.msc**
   - In **Current directory**, type **%windir%/system32**.
   - Select the **Allow application to interact with desktop** check box, select the **Use Network Name for computer name** check box, click **Next**, and then click **Finish**.
   - Bring the resource online using the procedure above. The Computer Management snap-in will appear on the desktop of the computer node that owns the resource.

### Notes

- This task is performed on a computer that is running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition as part of a server cluster, has Message Queuing with Triggers already installed, and has a Message Queuing resource already configured.
- Errors occurring during the performance of this task are recorded in a Windows event log. For example, a "device not ready" error indicates that Message Queuing was not installed on the node.
- Once created, do not rename resources for Message Queuing Triggers.
- The trigger and rule definitions for the cluster resource created are stored in subkeys in the local Windows registry under a special key. The Message Queuing Triggers resource is checkpointed with this registry key and the subkeys under it so that when the resource is moved to a new node (during failover, for example), the required keys are propagated to the local registry on the new node. After failover, the Triggers service can thus continue to process the incoming messages in each monitored queue and invoke the applicable stand-alone executable or COM component according to the rules defined.
- You can use the Generic Application resource created with a dependency on the Message Queuing resource to manage Triggers in the applicable resource group from a local Computer Management snap-in. For more information on these resources, see Related Topics.

#### Information about functional differences

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**

**Concepts**

Configure Message Queuing resources for server clusters
Installing Message Queuing in a server cluster
Message Queuing servers
Installing Message Queuing triggers
Working with MMC console files

# Monitor and Test Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitor and Test Message Queuing**
This section shows you how to:

- View Message Queuing events
- View external transactions
- Track Message Queuing service counters
- Track Message Queuing queue counters
- Enable route tracking and test messages
- Track message routing
- Send test messages
- Test connectivity using MQPing

# View Message Queuing events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view Message Queuing events**

1. Open Event Viewer.
2. In the console tree, click **Application**.
3. In the details pane, click the **Source** column to gather the **MSMQ** and **MSMQTriggers** events together.
4. Scroll until you see **MSMQ** and **MSMQTriggers** events listed in the column.

## Notes

- To open Event Viewer, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Event Viewer**.
- When viewing events in Event Viewer, you can sort, filter, and search for events, based on event properties.
- Message Queuing service events can be located with the same procedure, by clicking **System** in the console tree, instead of **Application**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitoring events
Working with MMC console files

# View external transactions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To view external transactions**

1. Open Component Services.
2. In the console tree, double-click **Component Services**, double-click **Computers**, double-click **My Computer**, and then double-click **Distributed Transaction Coordinator**.
3. To view the status of external transactions, double-click **Transaction List**.
4. To view the statistics of external transactions, double-click **Transaction Statistics**.

## Notes

- To open the Component Services snap-in, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Component Services**.
- The Microsoft Distributed Transaction Coordinator (DTC) must be running to implement external transactions.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitoring transactional messages
Working with MMC console files

# Track Message Queuing service counters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To track Message Queuing service counters**

1. Open System Monitor.
2. In the details pane, click the **Add** (+) button on the toolbar.
3. In the **Add Counters** dialog box, click **Use local computer counters** to view counters on the computer that you are using. Or, to view counters for another computer, click **Select counters from computer**, and then click the applicable computer.
4. In **Performance object**, click the arrow that opens the list, and then select **MSMQ Service**.
5. Click **All counters**. Or, click **Select counters from list**, and then select the counters that you want to monitor.
6. Click **Add**.
7. To view the definition of a counter, select a counter name, and then click **Explain**.
8. Click **Close**.

   The monitor graphs appear in the details pane.

## Notes

- To open System Monitor, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Performance**.
- To delete a counter from those tracked, in System Monitor, choose the required counter in the details pane, and then click the **Delete** (X) button on the toolbar.
- You can also use this procedure to track other categories of Message Queuing performance objects, including MSMQ Session, MSMQ Incoming HTTP Traffic, MSMQ Outgoing HTTP Session, MSMQ Outgoing Multicast Session, and MSMQ Incoming Multicast Session objects.
- You can view performance data from System Monitor in Chart, Report, or Histogram views, and you can log the performance data to a file. For more information on System Monitor, see System Monitor overview.
- For more information on monitoring Message Queuing, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitoring Message Queuing performance
Working with MMC console files

# Track Message Queuing queue counters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To track Message Queuing queue counters**

1. Open System Monitor.
2. In the details pane, click the **Add** (+) button on the toolbar.
3. In the **Add Counters** dialog box, click **Use local computer counters** to view counters on the computer that you are using. Or, to view counters for another computer, click **Select counters from computer**, and then click the applicable computer.
4. In **Performance object**, click the arrow that opens the list, and then select **MSMQ Queue**.
5. Click **All counters**. Or, click **Select counters from list**, and then select the counters that you want to monitor.
6. Click **All instances**. Or, click **Select instances from list**, and then select the instances you want to monitor.
7. Click **Add**.
8. To view the definition of a counter, select a counter name, and then click **Explain**.
9. Click **Close**.

   The monitor graph appears in the details pane.

## Notes

- To open System Monitor, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Performance**.
- To delete a counter from those tracked, in System Monitor, choose the required counter in the details pane, and then click the **Delete** (X) button on the toolbar.
- You can view performance data from System Monitor in Chart, Report, or Histogram views, and you can log the performance data to a file. For more information on System Monitor, see System Monitor overview.
- For more information on monitoring Message Queuing, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitoring Message Queuing performance
Working with MMC console files

# Enable route tracking and test messages

Updated: September 8, 2006

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To enable route tracking and test messages**

1. Click **Start**, click **Run**, type **regedit**, and then click **OK**.
2. Click to expand **My Computer**, click to expand **HKEY_LOCAL_MACHINE**, click to expand **SOFTWARE**, and then click to expand **Microsoft**.
3. Click to expand **MSMQ**, click to expand **Parameters**, and then click **Security**.
4. Right-click the details pane, point to **New**, and then click **DWORD Value**.
5. Type **EnableReportMessages** as the name of the new value, and then press ENTER.
6. Double-click the new **EnableReportMessages** value to change the value data. To enable the sending of test messages, enter a Decimal value of 1 in **Value data**. To disable the sending of test messages, enter a Decimal value of 0 in **Value data**. This is the default value.
7. Click **OK** to save the value. After setting the registry, restart the Message Queuing service for the changes to take effect.

## Caution

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

## Notes

- For instructions on sending test messages after the feature is enabled, see Send test messages.
- By enabling this feature, you also enable the message route tracking feature. For more details, see Track message routing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sending test messages

# Track message routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To track message routing**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourComputer/*msmq

4. Click **Properties**.
5. Click the **Diagnostics** tab, and then click **Define Tracking**. The **Message Route Tracking** dialog box appears.
6. In **Path name of report queue**, click the applicable report queue. Or, to create a new queue for report messages, click **New** and type the queue path name (such as *MyComputer**MyQueue* or .*MyQueue*).
7. Click **Track all messages** or click **Track only test message**.

**Notes**

- The tracking message route feature is disabled by default after installing Message Queuing. For instructions on turning on this feature, see Enable route tracking and test messages.
- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- After sending test messages is enabled, you can also send a test message from the Computer Management snap-in. To do this, open **Computer Management**, click **Services and Applications**, right-click **Message Queuing**, and then click **Properties**. On the **Diagnostics** tab, click **Define Tracking**.
- Any public queue can serve as a report queue. For information on how to create a new queue for report messages before performing this procedure, see Related Topics.
- If you click **Track all messages** and if the sending application has set the message route tracking flag and the destination queue is located on another computer, all messages sent by such an application from the source computer have their routes tracked. All report messages sent along the route of each message are then sent to the report queue on the source computer. Tracking occurs both for messages sent by the source computer and for messages routed through a Message Queuing server with routing enabled.
- If you click **Track only test message**, report messages are sent only for test messages. For information on how to send a test message, see Related Topics.
- Messages sent locally to a queue on the same computer are not tracked.
- Messages sent over HTTP/HTTPS transport and messages sent to multiple destinations using IP multicast, distribution

lists, or multiple-element format names cannot be tracked in the current version of Message Queuing.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Tracking messages and routes
Send test messages
Working with MMC console files

# Send test messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To send test messages**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourSourceComputer*/msmq

4. Click **Properties**.
5. On the **Diagnostics** page, click **Send Test**.
6. In **Destination queue**, click the applicable test queue, and then click **Send**. Or, to create a new queue to send test messages to, click **New Queue**, type the full queue path (such as *DestinationComputer**DestinationQueue*), click **OK**, and then click **Send**.
7. Click **Close** to close the dialog box.

**Notes**

- The ability to send a test message is disabled by default after installing Message Queuing. For instructions on turning on this feature, see Enable route tracking and test messages.
- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- After sending test messages is enabled, you can also send a test message from the Computer Management snap-in. To do this, open **Computer Management**, click **Services and Applications**, right-click **Message Queuing**, and then click **Properties**. On the **Diagnostics** tab, click **Send Test**.
- The computer whose msmq object you select is the source computer used to send a test message.
- You can send test messages only to test queues. All test queues are public queues and have a queue type ID of {55EE8F33-CCE9-11CF-B108-0020AFD61CE9}.
- You cannot use a transactional queue as a test queue. Only transactional messages can be sent to transactional queues, and test messages are not transactional. Similarly, you cannot change the type ID of a queue requiring authentication to that of a test queue and send test messages to the queue, because the messages will not be delivered (unless the queue is located on the source computer).
- Because all test messages are sent with the message route tracking flag set, if the source computer has message route tracking enabled and if the test queue is located on another computer, the route of such test messages is tracked and report messages are sent to the report queue on the source computer. For information on how message route tracking is enabled, see Related Topics.

- If you create a queue and immediately send a test message to it, the receiving computer may not receive the test message because of replication delays. In this case, wait several minutes and send the test message again.

- You can verify that the sent test messages arrived at the destination queue by viewing the properties of the messages in the destination queue or by viewing the contents of the report queue on the source computer if message route tracking is enabled. For information on how to view the properties of messages in a particular queue, see Related Topics.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Sending test messages
Working with MMC console files

# Test connectivity using MQPing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**To test connectivity using MQPing**

1. Open Active Directory Users and Computers.
2. On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**.
3. In the console tree, right-click **msmq**.

   **Where?**

   - Active Directory Users and Computers/*YourDomain*/*YourOrganizationalUnit* (such as Computers or Domain Controllers)/*YourCompute/*msmq

4. Click **Send MQPing message**.

## Notes

- To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.
- You can also send an MQPing message as follows: right-click **msmq**, click **Properties**, click the **Diagnostics** tab, and then click **MQPing**.
- When you send an MQPing message to a computer, Message Queuing waits 45 seconds to receive a response from the computer being tested. A message appears, notifying you whether MQPing succeeded or failed. You may need to send several MQPing messages to determine whether or not a computer is connected.
- You can also send an MQPing message from the Computer Management snap-in. To do this, open **Computer Management**, click **Services and Applications**, right-click **Message Queuing**, and then click **Properties**. On the **Diagnostics** tab, click **MQPing**.

**Information about functional differences**

- Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings. For more information, see Viewing Help on the Web.

**See Also**
**Concepts**
Monitoring connectivity
Working with MMC console files

# Message Queuing Concepts

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Concepts**

This section is designed to help you understand Message Queuing concepts, and provide background information about installing and administering a Message Queuing network:

- For an overview of Message Queuing features, Message Queuing servers and clients, and Message Queuing with Active Directory, see Introducing Message Queuing.
- For background information about Message Queuing features, see Understanding Message Queuing.
- For information on Message Queuing administration, see Administering Message Queuing.

# Introducing Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Introducing Message Queuing**

This section provides you with an overview of Message Queuing and a summary of Message Queuing features. Information about Message Queuing computers (servers and clients) and Message Queuing objects in Active Directory is also provided. For details, see the following links:

- For a general introduction to Message Queuing, see Message Queuing overview.
- For information about new and legacy features for Message Queuing 3.0, see Message Queuing features.
- For information on Message Queuing and Active Directory, see Message Queuing and Active Directory.
- For an overview of Message Queuing in clusters, see Distributing Message Queuing in a network.
- To learn about the different ways that Message Queuing can be deployed on computers, including information about Message Queuing server, and independent and dependent clients, see Message Queuing Computers.

# Message Queuing overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing overview**

Message Queuing (also known as MSMQ) is a messaging infrastructure and a development tool for creating distributed messaging applications for Microsoft Windows operating systems. Applications developed for Message Queuing send messages to queues, which are temporary storage locations, from which messages can proceed to their final destination as conditions permit. Such applications can communicate across heterogeneous networks and can send messages between computers that may be temporarily unable to connect to one another. Message Queuing provides guaranteed message delivery, efficient routing, security, support for sending messages within transactions, and priority-based messaging. Software products with these features are often referred to in the industry as message-queuing software, store and forward software, or message-oriented middleware.

Message Queuing is not a database and does not provide database functionality, but can be used together with a database, for example for message storage.

With Message Queuing, end users can communicate across networks and computers that may be offline, regardless of the current state of the network and computers. System administrators can use Message Queuing to efficiently manage large, complex networks of computers and message queues. Through Message Queuing, MIS decision makers get more reliable communication and more efficient use of network resources. Developers can focus on business logic instead of networking issues because Message Queuing effectively provides guaranteed network communication.

The current version of Message Queuing introduces a number of new features. For more details, see Message Queuing features. There are also a number of terminology and administrative changes from previous versions. If you are upgrading from a previous version, see the Upgrade overview for a list of changes.

**Message Queuing naming**

In its earliest versions (for Windows NT 4.0, Windows 95, and Windows 98), Message Queuing was called Microsoft Message Queuing Server, or MSMQ 1.0. Similarly, the version for Windows 2000 is officially called Message Queuing 2.0 and is often referred to as MSMQ 2.0 in external literature and other external sources.

The current version of Message Queuing, which is based on the original MSMQ technology and conceptions and is supplemented with numerous new features, is officially called Message Queuing 3.0. This documentation uses the current official name *Message Queuing* as a generic term for all the versions. This is not intended to discourage the use of the already familiar term *MSMQ*. The term MSMQ is used widely by veteran users and in external sources, and it is recommended that all users be familiar with both terms.

# Message Queuing features

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing features**

Following is a summary of the major features of Message Queuing 3.0:

- **Delivering messages over HTTP transport**. New in Message Queuing 3.0, SRMP (SOAP Reliable Messaging Protocol), an XML-based messaging protocol, has been developed for delivering high Quality of Service (QoS) messages. The direct, public, and private format names of administration and response queues can be included in messages sent over HTTP transport. Similarly, the names of administration and response queues in HTTP format can be included in messages sent over ordinary (non-HTTP) transport. For more information about these features, see HTTP/HTTPS messages. For more information on format names, see Queue names.

- **Triggers**. Now a component in Message Queuing 3.0, rather than an add-on module. Triggers provides a mechanism that associates the arrival of each incoming message at a queue with a response that depends on the contents of the message and may invoke either a COM component or a stand-alone executable program. Business rules can be defined and invoked in response to such messages without any additional programming. For more information about this feature, see Triggers.

- **Sending messages to multiple destinations**. New in Message Queuing 3.0, clients are able to send the same message to multiple recipient queues. Lists of destination queues can be specified explicitly by means of distribution group objects (distribution lists) in Active Directory as well as in the form of multiple-element format names. Support for ensuring that messages sent to distribution lists and multiple-element format names will reach queues on downlevel computers is provided. In addition, message delivery to IP multicast destinations using the PGM protocol is supported. For more information about these features, see Multiple Destination Messaging.

- **Message lookup**. Message Queuing 3.0 now provides a way to peek at or retrieve a specific message without using cursors to navigate through the queue until the message sought is located. This functionality is based on a 64-bit lookup identifier that is assigned to each message when it is placed in a destination queue. For more information, see Managing messages.

- **Active Directory Integration**. Message Queuing 3.0 improves integration with, and extension of, the Active Directory service to store all configuration, security, and status information. Message Queuing clients now use the Lightweight Directory Access Protocol (LDAP) to access domain controllers and global catalog servers for Message Queuing-specific information in Active Directory directly without assistance from a Message Queuing server on a domain controller.

- **Microsoft Management Console (MMC) support**. From Message Queuing 3.0, Message Queuing is now administered more completely using snap-ins hosted in an MMC console. Administrative tasks formerly performed in Control Panel are now performed in snap-ins. For more information, see New ways to do familiar Message Queuing tasks.

- **Workgroup support**. Message Queuing can be installed in workgroup mode on computers belonging to a Windows Server 2003 workgroup, rather than to a domain. In addition, a computer on which Message Queuing is installed in a workgroup can later join a domain, and then separate from the domain. For more information about belonging to a workgroup or a domain, see Deploying in a domain environment and Deploying in workgroup mode.

- **Active/active cluster support**. Message Queuing now fully supports the active/active paradigm in a server cluster, which means that Message Queuing can run on all nodes in a server cluster simultaneously. Message Queuing triggers are also integrated with active/active cluster support. For more information on the active/active model in server clusters, see Distributing Message Queuing in a network.

- **Windows CE support**. A special version of a Message Queuing client is preinstalled on handheld and palm-sized computers running the Windows CE 3.0 (or later) operating system. Windows CE also now supports the Message Queuing SRMP protocol, for HTTP messaging.

- **Message backup and restore**. Message storage files, log files, transaction log files, and registry settings can be backed up and restored in case of computer failure. For more information, see Backing up and restoring messages.

- **Message prioritization**. Message prioritization allows urgent or important messages to be sent before less important messages, so you can guarantee adequate response time for critical applications at the expense of less important applications. For more information on sending messages with different priorities, see Message priority.

- **Guaranteed message delivery**. Messages can be stored on a disk-based queue, and then later forwarded to provide guaranteed delivery.

- **Sending messages within transactions**. Using transactional capabilities, you can couple several related actions in a single transaction, ensure messages are delivered in order, ensure messages are delivered only once, and confirm that messages were successfully retrieved from the destination queue. For more information on sending messages within transactions, see Transactional messaging.

- **Dynamic queue creation**. You can create queues or change queue properties on the fly without affecting messaging applications.

- **Message routing**. Message Queuing provides message routing based on the physical topology of the network, session concentration needs, and transport connectivity. Session concentration facilitates the efficient usage of slow communication links. For more information on this subject, see Routing with Message Queuing servers.

- **Cross-platform integration**. Message Queuing can be used across a wide variety of hardware platforms using connectivity products provided by other vendors. For more information on messaging with different platforms, see Cross-Platform Messaging.

These features make Message Queuing ideal for the implementation of semi-independent client/server systems (such as order entry, accounting, and inventory applications), batch processing, queue-based client/server systems (first-come, first-served resource access), and for migration from legacy systems.

# Message Queuing and Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing and Active Directory**

In a domain environment, one or more Message Queuing objects are created in Active Directory during setup for all the types of Message Queuing computers except dependent clients. These objects store various information associated with Message Queuing. The access of users to Message Queuing objects is controlled by assigning security descriptors to the objects. A security descriptor lists the users and groups that are granted or denied access to an object and the specific permissions assigned to those users and groups. For more information about the control of access to Message Queuing objects, see Access control for Message Queuing.

Note that if you install Message Queuing on a computer that is in a domain environment, and later re-image the computer, the new installation of Message Queuing on the computer will create a new Queue Manager globally unique identifier (GUID). This has the following effects:

1. The public queue data from the old installation of Message Queuing will no longer be usable.

2. The new installation of Message Queuing does not recognize the Message Queuing objects in Active Directory that belong to the old installation. This prevents Message Queuing from working in Domain mode, which means that features that require Domain mode, such as access to public queues, authentication, encryption, and routing, will not be available. For more information, see Deploying in a domain environment.

To resolve the latter issue, the domain administrator must delete the Message Queuing objects in Active Directory for the old installation of Message Queuing. Note that this will also delete all public queues for the computer that was re-imaged. For more information about deleting Message Queuing objects, see Remove Message Queuing objects from Active Directory.

Under certain conditions, the name of an object can be changed by the directory service. This can occur if multiple objects with the same name are created simultaneously on different domain controllers. For information about how Active Directory handles object naming and conflicts, see Object names, in the Active Directory online Help file.

The following table lists those objects that are created by Message Queuing by default for each type of Message Queuing computer and that are displayed in the MMC snap-ins. For general information about MMC and snap-ins, see Microsoft Management Console.

| Message Queuing computer | Object name | Object class |
|---|---|---|
| Message Queuing server without directory service functionality or routing services on a domain controller | msmq | MSMQ-Configuration |
| Message Queuing server with directory service functionality on a domain controller | msmq; MSMQ-Settings | MSMQ-Configuration; MSMQ-Settings |
| Message Queuing server with routing services | msmq; MSMQ-Settings | MSMQ-Configuration; MSMQ-Settings |
| Message Queuing server on a nondomain controller | msmq | MSMQ-Configuration |
| Independent client | msmq | MSMQ-Configuration |
| Dependent client | none | not applicable |

When you manually create an object in Active Directory, you are prompted to specify the minimal information needed to establish the object in Active Directory. Each object has many more properties (attributes) that can be set. After an object is created, the values of many of these additional attributes can be viewed and modified in the **Properties** dialog box.

In addition, some information for Message Queuing is also stored in Active Directory in non-Message Queuing objects. Specifically, the User object contains about user certificates, the applicable computer object (above the msmq object) contains certificate information, information related to foreign sites (for cross-platform messaging) is contained in the Sites object, and distribution lists are defined as distribution group objects.

For advanced information about object attributes, see the Active Directory topic in the Microsoft Platform Software Development Kit (SDK). To access the Message Queuing SDK on the Web, refer to the MSDN Library Web site(http://msdn.microsoft.com/library/default.asp).

# Distributing Message Queuing in a network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Distributing Message Queuing in a network**
There are several techniques that you can employ to make a Message Queuing solution highly available to users in a network, and more scaleable.

**Implementing Message Queuing in clusters**
One solution might be to implement the Message Queuing service in clusters. Message Queuing triggers can also be implemented in a cluster. A cluster is a set of computers, known as nodes or hosts, that work together under a common identity to provide a service. All clusters provide load balancing, increased availability, and scalability, but it is important to distinguish the two main cluster solutions. The first, Network Load Balancing, is intended primarily to load balance incoming TCP/IP traffic. Network Load Balancing clusters distribute client connections over multiple servers, providing scalability and high availability for client requests for services and applications, but they do not have failover capabilities. For more information, see Message Queuing in Network Load Balancing clusters.

The second cluster type, a server cluster, is intended primarily to provide failover service from one computer to another. The computers participating in this solution form a different type of cluster. The Network Load Balancing cluster would most commonly be running Web server applications. In contrast, a server cluster would most commonly be running database applications (when used in conjunction with Network Load Balancing). A server cluster appears to network clients as a single, highly available system. For more information, see Message Queuing in server clusters.

Network Load Balancing and the Cluster service can not both be active on the same computer, but by joining the two cluster solutions together to function in a complementary fashion, an overall clustering scheme can be created.

**Web farms**
To increase the availability and scalability of Web server resources, a Web farm is often deployed as a collection of independent servers, to which all HTTP incoming requests are directed. Network Load Balancing clusters, network address translation (NAT), and DNS round robin are some of the methods typically used to deploy Web farm capability.

**DNS round-robin distribution**
Microsoft DNS server supports a process referred to as round-robin distribution, a common session-based load balancing technique that balances the number of clients connecting to an Internet site supported by multiple servers. Round-robin distribution is specified by RFC 1794. With DNS round-robin, a single logical DNS name (such as www.microsoft.com) maps to several IP addresses. When a request is made to resolve the DNS name, the DNS server sends back one of the addresses from the list. Each time a request is made for this DNS name, the DNS server rotates the addresses in round-robin fashion to distribute the set of users across a set of servers. Thus, different users requesting the same DNS name are given different IP addresses to balance the load.

This technique distributes incoming requests evenly among the available IP addresses but does not fully balance the load because it is not interactive. This is because the DNS server does not monitor the loading of each IP address nor does it check whether a particular server is currently running. Round-robin distribution can be a useful starting point for small groups of servers. If you are using round-robin distribution, closely monitor your servers so that you can quickly remove any failed machines from the distribution list.

**Network address translation**
Network address translation (NAT) is a protocol that allows a network with private addresses to access information on the Internet through an IP translation process. NAT can be used to configure a small network to share a single connection to the Internet.

# Message Queuing Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing Computers**
This section provides information about different ways in which Message Queuing can be configured on a computer. For details, see the following links:

- For information on the functionality provided by configuring a computer as a Message Queuing server, see Message Queuing servers.
- For information on the functionality and advantages of deploying independent clients, see Independent clients.
- For information on dependent client requirements and considerations, see Dependent clients.

For the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed on it. The term *Message Queuing client* can refer to either an *independent client*, or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed. Such computers are referred to as having independent client functionality or dependent client functionality.

# Message Queuing servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing servers**
For the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed on it. The term *Message Queuing client* can refer to either an *independent client*, or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed, and such computers are referred to as having independent client functionality, or dependent client functionality.

Message Queuing servers also provide the following features:

- Directory access for all client computers in the Windows Server 2003 family domain or Windows 2000 domain that are running Message Queuing on Windows NT 4.0, Windows 98, or Windows 95 computers.
- Support for dependent clients.
- Routing capabilities.

**Directory access**
Message Queuing servers enable client computers running MSMQ 1.0 clients running on Windows NT 4.0, Windows 98, or Windows 95 computers, and Message Queuing 2.0 clients running on Windows 2000 computers, to view and change information listed in the Active Directory directory service. To do this, the Message Queuing directory service (Downlevel Client Support) must be installed on a Windows Server 2003 family domain controller. For more information, see Installation requirements and restrictions.

Message Queuing independent clients running on Windows Server 2003 family computers can access Active Directory directly without the aid of a Message Queuing server.

**Routing capability**
Message Queuing servers installed with routing services enabled provide message routing and intermediate store-and-forward message queuing. Such Message Queuing servers can route messages within a site or between different sites when there is no direct connectivity between client computers. Message Queuing servers that provide message routing functions within the same site are called *in-routing servers* (InRS) or *out-routing servers* (OutRS). Intersite connectivity is defined by a *routing link*. A server through which the message traffic from one site to another can be configured to pass in a routing link is called a *site gate*. A site gate can belong to either or both of the sites connected by a routing link. A routing link can be configured with any number of site gates or no site gate. For more information, see Routing with Message Queuing servers.

For information on sites, see Deploying in a domain environment.

**Dependent client support**
Every dependent client in your network needs a supporting server, which can be any Message Queuing server with or without routing enabled. By default, following installation, Message Queuing servers are not enabled to support dependent clients. For instructions on enabling a Message Queuing server to act as a supporting server, see Enable a supporting server.

# Independent clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Independent clients**

For the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed. The term *Message Queuing client* can refer to either an *independent client*, or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed, and such computers are referred to as having independent client functionality, or dependent client functionality.

Computers defined as running with independent client functionality do not have the following features:

- They do not have the Routing component enabled, and thus do not have the message routing capabilities of a Message Queuing server with this component.
- They do not have the Downlevel Client Support component enabled, and thus cannot provide Active Directory access for MSMQ 1.0 clients running on Windows NT 4.0, Windows 98, or Windows 95 computers, or MSMQ 2.0 clients running on Windows 2000.

Message Queuing computers configured with independent client functionality run the Message Queuing service, host queues, send and receive messages, and can operate while disconnected from the network. They do not require synchronous access to a Message Queuing server to send or receive messages, or to query Active Directory. However, if you choose to install an independent client in a workgroup environment or any other environment that does not provide a directory service, such as Active Directory, direct connectivity is required to send messages. Any Windows Server 2003 family computer running Message Queuing and operating in domain mode, including independent clients, can act as a supporting server for dependent clients.

If an independent client is disconnected from the network, it can continue to generate messages destined for other computers. The independent client stores these messages locally and automatically sends them as soon as the network connection is reestablished.

**Independent client advantages**

The advantages of deploying independent clients in your organization are as follows:

- They are not dependent on a Message Queuing server to send or receive messages, and synchronous communications with a Message Queuing server are not required.
- The messaging load can be distributed across several Message Queuing servers by configuring independent clients to use specific Message Queuing servers (in-routing servers and out-routing servers).
- They support mobility; independent clients can move from site to site.

**Mobile usage**

Mobile independent clients, such as those installed on laptop computers, can send and receive messages when they move around. When an independent client connects to a new Windows Server 2003 family site, it automatically detects the domain controllers in its new domain and queries the directory service to determine the new site in which it resides. After the new site is established, all messages that were sent to the client while it was disconnected from the network are automatically rerouted from the previous site to the new site. Note that mobile users that run Message Queuing applications while disconnected from the network can place an additional load on the network. This happens when they attempt to send or receive a large number of messages when they reconnect to the network. This temporary load can be transferred to additional Message Queuing servers to prevent reduced performance.

A Message Queuing 3.0 computer that belongs to a Windows NT 4.0 enterprise and is capable of connecting to MSMQ 1.0 controller servers is not capable of automatically revealing its new site. In this case, you can provide the name of the new site either before disconnecting the computer from the network or after reconnecting to the network. Selecting a new site enables your computer to connect to the closest controller server and thus optimizes access to the MQIS database. If you specify the new site before disconnecting from the network, messages sent to the computer while it is disconnected are routed directly to the new site while the computer is en route.

As stated previously, clients can be configured to use Message Queuing servers (in-routing and out-routing servers) for intrasite routing. When such a client moves to a different site, those in-routing or out-routing servers are not used for routing messages in the new site. When the independent client returns to its original site, those same Message Queuing servers are again used for routing messages.

**Windows CE usage**

A special version of the independent client is also available on handheld and palm-sized computers that are running Windows CE 3.0 or later. Because of capacity requirements and lack of RPC support, the following limitations apply to the Windows CE version of the independent client:

- No Active Directory support. Windows CE independent clients cannot query Message Queuing servers running on Windows Server 2003 family domain controllers.
- Limited message routing support.
- Limited public queue support.
- No message authentication or encryption. All messages sent to or received by other Message Queuing computers travel as plaintext.
- No remote retrieval of messages. Messages stored on another Message Queuing computer cannot be retrieved from the Windows CE independent client.
- No multi-message transactions.
- No cross-platform support. Messages cannot be sent to computers running other messaging systems.

# Dependent clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dependent clients**

For the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed. The term *Message Queuing client* can refer to either an *independent client*, or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed, and such computers are referred to as having independent client functionality, or dependent client functionality.

Message Queuing servers with dependent client functionality require synchronous access to a Message Queuing server, called a *supporting server*, for all messaging functions. The supporting server can be any Message Queuing computer operating in domain mode with at least independent client functionality, preferably in the local site, but also in a nearby site. Note that computers configured as independent clients cannot provide Active Directory access for MSMQ 1.0 clients. These clients require Message Queuing servers running on a domain controller.

A dependent client relies on its supporting server to perform all messaging functions on its behalf, such as hosting queues, storing messages, sending messages, and receiving messages. If the supporting server goes offline, the dependent client will be unable to send and receive messages. By default, following installation, Message Queuing computers are not enabled to support dependent clients. For instructions on enabling a Message Queuing server to act as a supporting server, see Enable a supporting server. If you want to specify a Message Queuing server to act as a supporting server for a dependent client, you must also change the security configuration of the Microsoft Distributed Transaction Coordinator (DTC) on the Windows Server 2003 family computer to enable network transactions. The DTC is automatically installed on each computer during Windows Server 2003 family operating system installation for local transactions, and you need to enable network DTC. For instructions, see Enable network DTC access.

When you install a dependent client, you designate its supporting server. You can also later change the supporting server. Note that if you have the early version of a dependent client running on Windows NT 4.0 and you want to change its supporting server to a Message Queuing server running on Windows 2000 or Windows Server 2003 family, the dependent client must have Windows NT 4.0 Service Pack 4 or higher installed or else transactional messaging will not function. For instructions on changing the supporting server for a dependent client, see Change the supporting server for a dependent client.

For information on how to view the list of dependent clients currently supported by a Message Queuing server, see View the dependent client list for a server.

**Dependent client considerations**

Decide whether or not you want to deploy dependent clients in your organization. The advantages of using dependent clients are as follows:

- Server disk space is used to store messages. When necessary, you upgrade only the server resources (such as memory and hard disk space) to improve Message Queuing performance, instead of upgrading all the client computers.
- With fewer points of administration, you have fewer computers to back up and fewer journals and dead-letter queues to monitor.
- Fewer resources are used if the dependent client is frequently offline and the server is not.

Dependent clients also have the following disadvantages:

- Because the Message Queuing service runs on the supporting server, encrypted messages sent to or received by dependent clients travel between the dependent client and supporting server as plaintext.
- Queue access time is slower for a dependent client than it is for an independent client. This is because queue access is through an RPC-based connection.
- There may be performance issues if multiple dependent clients connect to their supporting server at the same time to send or receive messages.
- Dependent client queues are located on their supporting server, and it is not obvious which queues belong to which

dependent clients.

- When dependent client software is uninstalled, the associated queues are not automatically deleted from the supporting server.

- Because communication from a dependent client goes through the supporting server, if the communication to the supporting server is slow or expensive (such as between two sites), all dependent client communications are also slow or expensive.

- Dependent clients cannot run under a local user account.

- New Windows Server 2003 family Message Queuing features, such as HTTP transport, Message Queuing Triggers, and sending messages to multiple destinations, are not available to dependent clients. Thus, distribution lists, multiple-element format names, and multicast messaging are not supported for dependent clients.

- Dependent clients cannot be installed on computers in workgroup mode or on 64-bit computers.

- Transactional messaging is not supported for dependent clients running on Windows NT 4.0, Windows Milleninum Edition (ME), and Windows 98, Windows 95 computers that use a supporting Message Queuing server running on Windows XP Professional or Windows Server 2003 family operating systems.

# Understanding Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Understanding Message Queuing**
This section provides detailed information that you will need to understand before administering a Message Queuing network. For details, see the following links:

- To learn about message routing within and between sites, and deploying routing on Message Queuing servers, see Message Queuing routing.
- For information on security concepts that you must understand when administering a Message Queuing network, see Message Queuing security.
- To learn about different queue types, queue format names, and queue aliases, see Queues.
- For information about message types and properties, see Messages.
- For information about Message Queuing triggers, and trigger rules and conditions, see Triggers.
- For information about one-to-many messaging, see Multiple Destination Messaging.
- For an overview of using Message Queuing over the Internet and through firewalls, see Internet Messaging.
- For information on implementing Message Queuing in clusters, see Message Queuing in Clusters.
- For details of messaging between systems running Message Queuing and systems running different messaging applications, see Cross-Platform Messaging.

# Message Queuing routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing**

This section provides information about routing, and how you can deploy routing-enabled Message Queuing servers in your messaging network. For details, see the following links:

- For an overview of message routing, see Message Queuing routing overview.
- For information about routing messages within a site, see Intrasite routing.
- For information about routing messages between sites, see Intersite routing.
- To learn about routing links and routing-link costs, see Routing links.
- For information on deploying routing on Message Queuing servers, see Routing with Message Queuing servers.

# Message Queuing routing overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing overview**

A Message Queuing client always tries to establish a direct connection, or *session*, with the destination computer using the underlying IP network protocol. If a direct connection is not possible or not allowed, Message Queuing servers with routing enabled (routing servers) can temporarily store messages and subsequently forward them to the destination computer or to another routing server. In this way, they establish the route taken by messages as they proceed from the source computer to the destination computer. Specifically, message routing can occur when one or both of the following conditions exist:

- A session cannot be established between the source computer and the destination computer (for example, if the destination computer is offline).
- Specific Message Queuing routing servers are designated to store and forward messages on behalf of the source or destination computer.

**Message routing between networks**

Message Queuing can be used to send messages to and receive messages from external networks, such as other organizations or other Windows Server 2003 family forests. Message Queuing computers can use various service ports to communicate with external networks through properly configured firewalls that allow communication over those ports. For more information, see Internet messaging overview.

Although clients within a site that use the same network protocol are assumed to have direct connectivity, using in-routing and out-routing servers to route messages within a site may reduce the bandwidth requirements in your network. Clients can be configured to use in-routing and out-routing servers to provide message routing on their behalf. In that case, all messages sent or received by a client are routed first through such a Message Queuing server. The term client here refers to any Message Queuing computer that cannot route messages on its own, such as independent clients and Message Queuing servers without routing enabled. This feature is not applicable to dependent clients. Dependent clients require a supporting server to send and receive messages on their behalf. A supporting server without routing enabled, rather than its dependent clients, can be configured to use in-routing and out-routing servers.

If site gates are used to route messages between sites, only those assigned Message Queuing servers can route messages between the sites.

**Note**

- Message routing is not applicable in a workgroup environment. Thus, in-routing servers, out-routing servers, and site gates cannot be used, and direct connectivity between the source and destination computers is required.

# Intrasite routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Intrasite routing**

Message routing within a site is called *intrasite routing*. In the simplest case, the source and destination computers are not configured to use in-routing or out-routing servers. In this case, as a first choice, Message Queuing tries to deliver a message to the destination computer directly (one hop), and as a second choice, it delivers the message to one randomly chosen Message Queuing routing server that resides in the site (two hops).




This diagram shows one possibility for routing a message from one computer to another computer within a site. Although the number of hops can vary, here, a message sent from one computer in the site to another computer in the same site makes two hops: one hop to get to the Message Queuing routing server and one hop to get to the destination computer. Because each hop is within a site, the performance overhead of each hop is minimal.

**In-routing and out-routing servers**

In the context of intrasite routing, independent clients and Message Queuing servers without routing enabled are functionally equivalent. These computers can be configured to use in-routing servers (InRS) or out-routing servers (OutRS) to store and forward messages on their behalf. Dependent clients cannot be configured to use in-routing or out-routing servers because all messaging functions are performed by a supporting server.

If an independent client or a Message Queuing server without routing enabled is configured to use an out-routing server, every outgoing message sent by that computer is routed through that server. Likewise, if such a computer is configured to use an in-routing server, every message sent to it is routed through that server. Such a computer can use the same Message Queuing routing server as both its in-routing and out-routing server.

All in-routing and out-routing servers must be in the original site of the computer that they serve. When a mobile client moves to another site, these servers are not used for message routing until the client returns to its original site.

**Network bandwidth**

Intrasite routing can reduce the usage of network bandwidth by providing session concentration. For example, if you use a star topology within a site and each independent client usually communicates directly with all other independent clients, you can reduce your bandwidth usage through session concentration by configuring each client to use in-routing and out-routing servers.




In this figure, the same Message Queuing routing server is configured as both an in-routing server and an out-routing server for each client.

# Intersite routing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Intersite routing**

Message routing between different sites, called *intersite routing,* requires connectivity between those sites. Connectivity between different sites can be defined by logical communications links, called *routing links*. If the message traffic between two specific sites can be configured to pass through a specific Message Queuing routing server or one of several such servers, that routing server is called a *site gate*. Routing links can use site gates to route messages between two sites.

If your organization has more than one Windows Server 2003 family site, you must create routing links to guarantee that messages can be routed between different Windows Server 2003 family sites. The route taken by messages between sites can be determined by the routing-link costs. All routing links are assigned a default cost of 1.

Link costs can be used to select the best intersite message route and may be based on the speed or the monetary cost of the underlying physical communication link. After you determine the number of Windows Server 2003 family sites in your forest, decide if you need to specify different costs for routing links. For information on how to change routing-link costs, see Change the cost for a routing link.

For information on sites, see Deploying in a domain environment.

In a domain environment with more than one Windows Server 2003 family site, you must create routing links to guarantee that messages can be routed between different sites. In the absence of a routing link, Message Queuing will usually attempt to establish a direct connection with the destination computer in an unconnected site.

Each routing link is assigned a relative cost, which may reflect the speed or the monetary cost of the underlying physical communication link. The default value is 1, and costs can range from 1 to 999999. Change routing-link costs only in cases where alternative intersite routes are possible over the existing routing links and you want to enforce one route over another route using the costs of the individual links.

**Notes**

- A site gate can belong to either or both of the sites connected by a routing link. A routing link can be configured with any number of site gates or no site gate.
- Do not confuse routing links with site links. Routing links are used by Message Queuing servers to route messages between sites. Site links are used by domain controllers to replicate Active Directory information between sites.

When a message is sent from an independent client in one site to a computer in another site, the source computer first tries to connect directly to the destination computer, and only as a second choice, delivers the message to a routing server in its site, which delivers the message to a site gate. This process continues until the message reaches the destination computer.

**Session reduction**

Intersite routing typically reduces the number of sessions created by clients. In the following figure, the routing beyond the site is transparent to the computers within the site, with the exception of the site gate. This simplifies the message route and reduces site-to-site sessions. Only the designated Message Queuing server can establish a session between different sites.



# Routing links

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing links**

If your organization has more than one Windows Server 2003 family site, you must create routing links to guarantee that messages can be routed between different Windows Server 2003 family sites. The route taken by messages between sites can be determined by the routing-link costs. All routing links are assigned a default cost of 1. Change routing-link costs only in the case where alternative routes are possible over the existing routing links and you want to enforce one route over another route using the costs of the individual links. For more information on configuring routing links, see Configuring routing links.

Link costs can be used to select the best intersite message route and may be based on the speed or the monetary cost of the underlying physical communication link. After you determine the number of Windows Server 2003 family sites in your forest, decide if you need to specify different costs for routing links. For information on how to change routing-link costs, see Change the cost for a routing link.

For information on sites, see Deploying in a domain environment.

**Routing link example**

A routing link that connects sites A and B has four site gates (SG1-SG4). If SG1 and SG2 belong to site A, SG3 belongs to both sites A and B, and SG4 belongs to site B, messages can be routed through one of the site gates of site A (SG1, SG2 or SG3) and then through one of the site gates of site B (SG3 or SG4). However, if the messages are routed through SG3, which belongs to both sites, they will reach the destination computer in one less hop. Thus, messages are routed preferentially through SG3. In the general case, if a routing link is associated with only a pair of site gates, one in each site, only those servers can route messages between the designated sites in that routing link.

Multiple routing links can be used to route messages between sites in your network for load balancing or redundancy. In this case, if you want to enforce one route over another route using link costs (for example, to specify a default link), you can change the routing-link costs to determine the route that messages take.

For example, if the cost associated with the A-B routing link is set to 3 and the cost associated with the other three routing links (B-C, B-D, and C-D) is 1, messages routed from site A to C always travel from site A to site B and then to site C. However, if the cost associated with the A-B and B-C routing links is 3 and the cost associated to the C-D and B-D routing links is 1, messages routed from site A to C always travel from site A to site B, to site D, and then to site C.




The route that a message takes to get from site A to site C depends on the costs of the routing links.

| Routing link | Cost |
|---|---|
| A-B | 3 |
| B-C | 1 |
| B-D | 1 |
| C-D | 1 |

In this situation, messages routed from site A to site C always travel from site A to site B and then to site C.

| Routing link | Cost |
| --- | --- |
| A-B | 3 |
| B-C | 3 |
| B-D | 1 |
| C-D | 1 |

In this situation, messages routed from site A to site C always travel from site A to site B, then to site D, and finally to site C.

# Routing with Message Queuing servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Routing with Message Queuing servers**

The number of Message Queuing servers with routing enabled (routing servers) that you deploy in your organization can depend on the following factors:

- The number of sites among which you are routing messages.
- How frequently client computers are offline. Message Queuing servers with routing services are needed for message store-and-forward capability when the source and destination computers may not be online at the same time.

It is recommended that at least one routing server be installed in each Windows Server 2003 site to ensure that messages are forwarded successfully. For performance and scalability reasons, it is recommended that multiple Message Queuing routing servers be installed in each site.

A source computer will successfully forward a message to a destination computer when a direct connection is established and when the destination site has no site gate. If the destination site has a site gate, the messages will not be forwarded and will accumulate at the source.

It is recommended that in-routing and out-routing servers be available for mobile clients that are mostly disconnected from the network or connect to the network through a remote access server.

There is no technical restriction or limitation on installing Message Queuing routing servers on domain controllers. However, because the main roles of domain controllers include authenticating users, handling queries to Active Directory, and maintaining domain security, it is recommended that you carefully consider how installing Message Queuing routing servers on domain controllers may influence their performance and security.

By providing session concentration, the use of Message Queuing routing servers may reduce operating costs. Such servers can reduce bandwidth needs within a site and reduce the number of sessions between different sites. Intrasite routing is considered relatively fast and inexpensive, while intersite routing is considered relatively slow and expensive.

In addition, when Message Queuing routing servers are available, messages do not accumulate in independent clients, which generally have a small storage capacity, but can be delivered to an out-routing server, which stores them until the destination computer becomes online.

Message Queuing routing servers that provide intrasite message routing are called *in-routing servers* (InRS) or *out-routing servers* (OutRS). Message Queuing servers providing message routing on a routing link, which defines connectivity between sites, are called *site gates*. For more information on intrasite routing, see Intrasite routing. For more information on intersite routing, see Intersite routing.

The relative cost of using a given route to reach a destination computer is commonly measured in *hops*, where each hop represents the delivery of a message from one computer to another. If there are multiple routes with the same destination, the route with the lowest *hop count* is the best.

# Message Queuing security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security**

Message Queuing uses security features provided by Windows Server 2003 family operating systems, including access control and auditing for Message Queuing objects, message authentication, and encryption. For details, see the following links:

- For a general overview of security features for Message Queuing, see Message Queuing security overview.
- For information on implementing authentication to check who sent messages, and message integrity, see Authentication for Message Queuing.
- For information on encrypting private messages, see Encryption for Message Queuing.
- For information on restricting operations of Message Queuing objects to specific groups or users, see Access control for Message Queuing.
- For information on auditing access operations of Message Queuing objects, see Auditing Message Queuing objects.
- Message Queuing can be deployed in hardened MSMQ mode to enhance the security of Message Queuing computers using HTTP messaging over the Internet. For more information, see Message Queuing in hardened MSMQ mode.

# Message Queuing security overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Security overview**

Message Queuing takes advantage of the various built-in security features of the Windows Server 2003 family operating systems. Specifically, Message Queuing uses access control, message authentication, encryption, and auditing for security. In Message Queuing 3.0, 128-bit encryption is supported in addition to 40-bit encryption. By managing security properties for objects, you can set permissions, assign ownership, and monitor user access.

Access control is used to restrict user access to Message Queuing objects in Active Directory and is implemented by assigning security descriptors to objects. Message Queuing objects include the MsmqServices, msmq, Queue, routing link, and MSMQ Settings objects. The msmq object for the local computer is also known as the MSMQ configuration object, for example, in names of specific permissions for this object and in the Delegation of Control Wizard. A security descriptor lists the users and groups that are granted or denied access to an object as well as the specific permissions assigned to those users and groups. For more information, see Access control for Message Queuing.

Security queries to Active Directory and other directory services, in which the client and server are online and talk to one another, are implemented using the Kerberos V5 security protocol (between Windows Server 2003 computers) or server authentication (for compatibility with MSMQ 1.0 running on Windows NT 4.0).

Message authentication, which verifies the identity of the sender of a message to the receiver, is implemented using certificates. Messages are authenticated asynchronously without the sender and the receiver communicating with one another. For more information, see Authentication for Message Queuing.

Encryption is implemented using both public/private key (asymmetric) and secret key (symmetric) algorithms. Encryption is used by Message Queuing applications to encrypt messages sent between Message Queuing computers. For more information, see Encryption for Message Queuing.

Security auditing is used to record which users attempt to access Message Queuing objects in Active Directory. The security descriptor for an object specifies the various security events to be audited for the object. For more information, see Auditing Message Queuing objects.

**Message security**

Message Queuing ensures the security of messages sent from a source computer to a destination computer. Applications can perform a *local read* or a *remote read* when accessing queues. A local read is performed when the receiving application accesses a destination queue that resides on the same computer. In this case, message security can be guaranteed.

A remote read is performed when the contents of a queue are accessed from a computer other than the one on which the queue is located. For remote reading, Message Queuing 3.0 uses encrypted RPC by default. This feature is available when a Windows Server 2003 family client computer does a remote read against a Message Queuing computer running on Windows Server 2003 family computers, Windows 2000, or Windows NT 4.0. Note that in situations where encrypted RPC cannot be used, (for example, where a workgroup computer is part of the remote read process) the message will be passed to the remote computer as clear text and message security is not guaranteed. A clear text message that has reached its destination queue can be read only by users that have the necessary access rights to read messages from the queue.

**Secured remote read**

Message Queuing 3.0 on Windows Server 2003 family computers provides a new secure remote read interface that enhances the security of the old remote read interface used by Message Queuing running on earlier versions of Windows operating systems, and allows Message Queuing servers to expose remote-read functionality in a more secure way.

Following Message Queuing installation, Message Queuing 3.0 on Windows Server 2003 family computers listens for both the old remote read interface and the new secure remote read interface, and the following default settings are applied:

- Message Queuing 3.0 clients on Windows Server 2003 family computers in the same forest as the Message Queuing 3.0 server will use the secure remote read interface with an encrypted channel.

- Message Queuing 3.0 clients across forests on Windows Server 2003 family computers in non-trusted domains will use the secure remote read interface. By default, the Message Queuing 3.0 server requires domain clients to establish an encrypted channel, and such a channel cannot be established between non-trusted domains. Thus remote read requests from such clients will be rejected. To modify this default behavior and allow the Message Queuing server to accept

domain clients that do not establish an encrypted channel, create a DWORD value **Security\NewRemoteReadServerAllowNoneSecurityClient** in the registry and set it to 1. Creation of this registry key causes clients from non-trusted domains to be validated using Anonymous logon credentials. After creating this registry key, the Anonymous logon account must be granted Peek or Receive permissions in order to accommodate remote read requests for clients from non-trusted domains. After implementing this registry key and granting permissions to the Anonymous logon account, Message Queuing server accepts Peek or Receive requests from everyone without authentication checks. Therefore, these changes should only be implemented when absolutely necessary.

- Message Queuing 3.0 workgroup clients on Windows Server 2003 family computers will use the secure remote read interface. Workgroup clients cannot establish an encrypted channel for remote reading, and by default, the Message Queuing server accepts workgroup clients on a non-encrypted channel using Anonymous logon credentials. In this scenario, the Anonymous logon account must be granted Peek or Receive permissions in order to accommodate remote read requests from workgroup clients. To modify this default behavior so that the Message Queuing 3.0 server rejects workgroup clients, create a DWORD value **Security\NewRemoteReadServerDenyWorkgroupClient** in the registry and set it to 1.

- MSMQ 1.0 and MSMQ 2.0 clients, and Message Queuing 3.0 clients on Windows XP computers, will use the old remote read interface. If you enable your Message Queuing 3.0 server to use only the secure remote read interface, the computer does not listen on the old remote read interface, and remote reads from these clients are not supported.

You can enable your Message Queuing 3.0 server to use only secured remote reading mode in the Server Security tab of Message Queuing properties in Computer Management. In secured remote reading mode, your computer will only listen on the new secure remote read interface, and not on the old remote read interface. The effect of this is that only Message Queuing 3.0 servers on Windows Server 2003 family computers can remotely receive messages from queues on your computer, and remote reads from MSMQ 1.0 clients, MSMQ 2.0 clients, and Message Queuing 3.0 clients running on Windows XP computers are not supported. For instructions on enabling your server to use only the new secured mode, see Enable secured remote read.

# Authentication for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authentication for Message Queuing**
The following sections describe the various types of authentication available for Message Queuing.

**Server authentication**
*Server authentication* is used by clients to authenticate queries sent to domain controllers and Message Queuing servers. A client can verify that the query results have not been tampered with and that the results were returned by the correct domain controller or Message Queuing server. Security queries to Active Directory and other directory services, in which the client and server are online and talk to one another, are implemented using the Kerberos V5 security protocol (between Windows Server 2003 family and/or Windows 2000 computers), or with server certificates (for compatibility with MSMQ 1.0 running on Windows NT 4.0). The method used depends on the version of Message Queuing (and hence the operating system) running on the server and the client. For more information on server authentication, see Managing server authentication.

**Message authentication**
*Message authentication*, which confirms that a message has not been tampered with and can be used to verify the identity of the sender of a message to the receiver, is implemented using certificates. Messages are authenticated asynchronously without the sender and the receiver communicating with one another. The use of message authentication is optional. Message Queuing applications are not required to use authentication when sending or receiving messages. Note that the use of message authentication slows the transfer of messages. When used, however, Message Queuing applications that send or receive authenticated messages must use the cryptographic service provider installed by default with Windows Server 2003 family, which is needed for both message authentication and encryption. For more information on message encryption, see Encryption for Message Queuing. For more information on message authentication, see Managing message authentication.

For more information on authentication for Windows Server 2003 family, see Authentication, in the Windows Help file.

For more information on certificates for Windows Server 2003 family, see Certificates overview in the Windows Help file.

**HTTPS authentication**
Using HTTPS to send messages ensures authentication of users requesting access to Message Queuing resources by means of a Web server, by establishing a Secure Sockets Layer (SSL) connection for secure communication between a sender and recipient. Authentication for an SSL connection is done using server and client certificates. For more information, see HTTPS authentication.

# Encryption for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Encryption for Message Queuing**

In the context of Message Queuing, encryption is the process for encoding and decoding messages, ensuring they cannot be read or used by unauthorized persons. Message Queuing implements message encryption through the use of symmetric and asymmetric (public/private) key encryption algorithms. The encryption implementation used is based on the CryptoAPI (Microsoft Cryptographic Application Program Interface), which in turn uses a cryptographic service provider (CSP) installed on the computer.

The use of message encryption is optional. Message Queuing applications are not required to use encryption when sending or receiving messages. Note that the use of message encryption slows the transfer of messages. When used, however, Message Queuing applications that send or receive encrypted messages must use the CSP installed by default with Windows Server 2003 family, which is needed for both message authentication and encryption. For more information on message authentication, see Authentication for Message Queuing.

Although other CSPs can be used for the encryption of data that does not involve Message Queuing, the CSP provided with Windows Server 2003 family must be used for any Message Queuing application that sends or receives encrypted messages. Non-Message Queuing applications can continue to use other CSPs for encryption.

**How message encryption works**

Messages are encrypted by the Message Queuing service on the source computer and decrypted by the Message Queuing service on the destination computer. Decrypted messages are stored in the destination queue as plaintext. Message encryption is designed to prevent unauthorized viewing of messages in transit.

When message encryption is used, performance is reduced, although typically not as much as when using message authentication. Performance losses due to encryption are most significant when you send encrypted messages to many different computers. However, when you send multiple messages to the same destination computer, only the first message sent takes significantly longer than usual to send.

To see why message throughput is reduced, it is important to understand the encryption mechanism that Message Queuing uses. On the source computer, Message Queuing does the following:

1. Creates a secret key.
2. Encrypts the message body by using the secret key.
3. Encrypts the secret key by using the public key of the destination computer.
4. Attaches the encrypted secret key to the encrypted message.

On the destination computer, Message Queuing does the following:

1. Decrypts the secret key by using its private key (of the key pair).
2. Decrypts the message body using the secret key.

Secret keys can be used for both encryption and decryption, hence they are called symmetric. Public keys can be used for encryption only, while private keys can be used for decryption only.

Asymmetric key encryption has more overhead than symmetric key encryption. Sending the first message incurs the most overhead because the secret key is encrypted with the public key of the destination computer. By default, Message Queuing stores the encrypted secret key on the source computer for 10 minutes. If you send additional messages to the same computer within 10 minutes after sending the first message, Message Queuing reuses the same secret key. Consequently, the overhead for each subsequent message will be just the overhead needed to encrypt the message body using the secret key. In this case, Message Queuing does not need to encrypt the secret key with the public key of the destination computer again. The same performance results are realized on the destination computer. Upon receiving the first message, Message Queuing decrypts

the secret key by using its own private key. Message Queuing then stores the decrypted secret key and can reuse it later for subsequent messages without any additional overhead.

You can specify whether a queue accepts encrypted messages, unencrypted messages, or both by selecting the *privacy level* for the queue. If the queue privacy level is set to **None**, the queue accepts only unencrypted messages. If the queue privacy level is set to **Body**, the queue accepts only encrypted messages. If the queue privacy level is set to **Optional**, the queue accepts both unencrypted and encrypted messages. You can check whether a message sent to the queue was encrypted or unencrypted using Active Directory Users and Computers (or programmatically by checking the PROPID_M_PRIV_LEVEL message property).

For information on how to set the privacy level for a queue, see Change the privacy level for queues.

**Notes**

- When messaging across different messaging platforms, encrypted messages sent to foreign computers stay encrypted on the intermediary Message Queuing server running the connector application. The connector application can forward the encrypted messages to the foreign computer or can decrypt the messages first before forwarding.

- Because encryption requires queries to Active Directory, it cannot be used in workgroup or offline mode.

**Cryptographic keys**

Public/private cryptographic key pairs for a computer are created for you during Message Queuing Setup and can be renewed using the **Message Queuing Properties** dialog box, which can be opened from Computer Management. All Message Queuing computers except dependent clients use cryptographic keys.

For more information on encryption keys for Windows Server 2003 family, see Generating encryption keys and certificate requests.

Cryptographic key renewal is a security feature used by Message Queuing. It is a good idea to renew your cryptographic keys periodically, such as once a year. It is also recommended that you renew your cryptographic keys any time you suspect they have been compromised. In particular, this will reduce the chance that your private key has been compromised.

When you renew the cryptographic keys for a Message Queuing computer, any messages that were encrypted using the previous public keys cannot be decrypted. This occurs only until the new public keys are replicated in Active Directory to all Message Queuing servers and the source computers can start using the new public keys.

For information on how to renew your cryptographic keys, see Renew cryptographic keys for Message Queuing.

# Access control for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Access control for Message Queuing**

Access control is used to restrict user access to Message Queuing objects in Active Directory and can be implemented by assigning security descriptors to objects. A security descriptor lists the users and groups that are granted or denied access to an object and the specific permissions assigned to those users and groups. This part of the security descriptor is known as a discretionary access control list (DACL). By setting the permissions on an object, the *owner* of the object controls which access is allowed. The Message Queuing objects include computer (msmq), queue, routing link, and MSMQ Settings objects. For more information on Message Queuing objects, where they are created, and where they are located in Active Directory, see Message Queuing and Active Directory.

Permissions can also be used to restrict users from sending messages to, or retrieving messages from, a particular queue on a computer. Although messages are not objects in Active Directory, they are protected through the security descriptor of the queue object.

Because Message Queuing provides asynchronous messaging, the source and destination computers do not need to be online at the same time. In this case, Message Queuing can implement access control for offline users through the use of sender security IDs (SIDs). For example, because a queue can restrict access to itself, the sending application must attach the sender's SID to any message directed to that queue. The Message Queuing service on the destination computer then checks the SID to verify that the sender has the proper permissions to access the queue.

**Note**

- You must also be granted certain permissions to be able to install and uninstall Message Queuing. For more information, see Installation permissions.

For more information on access control for the Windows Server 2003 family, see Access Control, in the Windows Help file.

# Auditing Message Queuing objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Auditing Message Queuing objects**

You can use auditing to record which users attempt to access Message Queuing objects, the type of operation attempted, and whether that access succeeded or failed. The security descriptor for an object specifies the various access events to be audited for the object. This part of the security descriptor is known as a system access control list (SACL). More specifically, a SACL specifies the following:

- The users or groups to audit when accessing the object.
- The access events to be audited for each group or user.
- A Success or Failure attribute for each access event, based on the permissions granted to each group or user in the object's discretionary access control list (DACL).

For more information on auditing for Windows Server 2003 family, see Auditing overview, in the Windows Help file.

The security log records audited operations, called *security events*, which you can view using Event Viewer. For more information on Event Viewer, see Event Viewer, in the Windows Help file.

**Configuring auditing**

To implement auditing, you must establish an audit policy. An audit policy specifies categories of security-related events that you want to audit. For instructions, see Define or modify auditing policy settings for an event category, in the Windows Help file. To audit "object open" security events such as *Peek Message*, *Receive Message*, and *Send Message*, you grant the Message Queuing service of a specific computer security permissions to access SACL information for its public queue. For more information, see Enable Message Queuing to access SACL information.

Next, you select the access entries to be audited for success or failure for a user or group on a specific object or resource. For information on how to set auditing options for a Message Queuing computer or queue object, see Audit access to computer and queue objects. For information on how to set auditing options for a routing link object, see Audit access to routing link objects.

If auditing options are specified for a Message Queuing object but an audit policy has not been established on the Message Queuing server that carries out the operation, a security event is not created.

**Security events**

Security events are created in the security log on the computer that performs the operation, which is not necessarily the computer that is the *owner* of the object. Because of this, audit messages for a single Message Queuing object can be scattered across a network. For example, in this figure, if the public test queue located on client 1 in the San Francisco site is deleted by client 2 in the Seattle site, the audit message is created on server 2 of the Seattle site. The following diagram illustrates this:



# Message Queuing in hardened MSMQ mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing in hardened MSMQ mode**
The purpose of Message Queuing in hardened MSMQ mode (a new feature in Message Queuing 3.0) is to enhance the security of Message Queuing computers running on the Internet. Hardened MSMQ mode is intended to support scenarios that employ only HTTP (SRMP) messages.

Message Queuing in hardened MSMQ mode imposes the following restrictions:

- The Message Queuing service does not listen directly to any ports. Thus, only messages arriving from remote computers through IIS are accepted and placed in their destination queues. Remote RPC calls are ignored (remote reading is blocked), and dependent clients are not supported.

- Attempts to open remote queues with non-HTTP format names and send messages to them succeed. However, all the outgoing queues created on the local computer that correspond to remote destination queues specified by non-HTTP format names are in the locked state. Messages, including acknowledgment and response messages, that reside in locked outgoing queues are not transmitted to their destinations.

- Messages intended for remote destination queues with non-HTTP format names, including remote queues designated by non-HTTP elements of a multiple-element format name, remain in the locked outgoing queues until hardened MSMQ mode is canceled and the Message Queuing service is restarted.

In hardened MSMQ mode, messages can still be placed in local queues with non-HTTP format names and received from them.

**Security considerations in hardened MSMQ mode**
When deploying a domain controller with the Downlevel Client Support component in a perimeter network, be aware that this component listens to the RPC interface, and that it is not governed by hardened MSMQ mode. For reducing the attack surface, you may want to remove the Message Queuing Downlevel Client Support component from the domain controller.

For instructions on enabling or canceling hardened MSMQ mode on the local computer, see Enable hardened MSMQ mode.

# Queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Queues**

In the context of Message Queuing, queues are temporary storage locations for different types of messages. The following sections provide information about different types of queues that can be deployed in a messaging infrastructure. Queues can be created by applications and by Message Queuing. Queues created by applications or by administrators are referred to as application queues. Queues created by Message Queuing are known as system queues. For details, see the following links:

- For information on the different types of application-generated queues used by Message Queuing, see Application Queues.
- For information on journal queues, outgoing queues, dead-letter queues, and other types of system queues created and used by Message Queuing, see System Queues.
- For information on the names that can be used to designate Message Queuing queues, see Queue names.
- For information on referencing queues by queue alias names, see Queue aliases.
- For information on target journaling, see Target journaling.
- For information on source journaling, see Source journaling.

# Application Queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Application Queues**
Queues that are created by an application or by the Message Queuing administrator are referred to as application queues. For information on different types of application queues, see the following links:

- Destination queues are any application queue that sending applications send messages to, or that receiving applications receive messages from. For more information, see:
    - Public and private queues
    - Transactional and nontransactional queues
- Administration queues are application-generated queues used to store system-generated negative and positive acknowledgment messages that are created by Message Queuing or connector applications. Any available nontransactional queue can be used as an administration queue. For more information, see Administration queues.
- Response queues are application-generated queues used to store application-generated response messages that are typically returned by an application that is reading messages from a queue. Any available queue can be specified as a response queue. For more information, see Response queues.
- Report queues are application-generated queues used to store report messages generated by Message Queuing or connector applications when tracing is requested by the sending application. When a message is sent with tracing enabled, a report message is generated each time the message leaves or arrives at a Message Queuing computer. For more information, see Report queues.

# Public and private queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Public and private queues**

**Private queues**
*Private queues* are queues that are not published in Active Directory and are displayed only on the local computer that contains them. Private queues have the following features:

- Message Queuing registers private queues locally by storing a description of the queue in the LQS (local queue storage) directory on the local computer. In MSMQ 2.0, and Message Queuing 3.0, the default location is %windir%\system32\msmq\storage\lqs. Note that a description of each public queue created on the local computer is also stored locally in a separate file in the LQS folder.

- Private queues are registered on the local computer, not in the directory service, and typically cannot be located by other Message Queuing applications.

- Private queues are accessible only by Message Queuing applications that know the full path name, the direct format name, or the private format name of the queue, as follows:

  - **Path name:***ComputerName*\private$*QueueName*.

  - **Path name on local computer:** \private$*QueueName*.

  - **Direct format name:**: DIRECT=*ComputerAddress*\PRIVATE$*PrivateQueueName*.

  - **Private format name:** PRIVATE=*ComputerGUID**QueueNumber*.

  For more information on path names and format names, see Queue names.

- Private queues have the advantage of no directory service overhead, making them quicker to create, no latency in accessing them, and no replication overhead.

- Private queues are not dependent on the directory service, and thus can be created and deleted when the directory service is not working. This is useful for offline operation.

- One way that private queues can be exposed to other applications is by setting a message property. To distribute the location of a private queue, an application can send a format name of the private queue as the response queue property of a message.

- In a workgroup environment, only private queues are available.

- For a Message Queuing application on the local computer to access a private queue, the application only needs to know the path name of the queue. For a Message Queuing application on a remote computer, the application requires a direct or private format name to access such a queue. To read messages from the queue, the remote Message Queuing application requires the Receive Message permission for that queue. By default, everyone has permissions to send messages to public and private queues.

**Public queues**
In a domain environment, *public queues* are queues that are published in Active Directory and hence are replicated throughout your Windows Server 2003 family forest. Note that only the properties for these queues are replicated, not the actual queues themselves or their contents. Any computer within your forest can potentially access information regarding public queues provided the user has sufficient permissions to access the applicable queue objects. Generally speaking, any user in the forest with access to Active Directory and the Send To permission for a given public queue can send messages to it.

In an Active Directory environment, defining public queues ensures that queues are registered in the directory service, and that their registration is backed up. They are persistent and available to other applications.

**Creating local public queues**

By default, in Active Directory, the Create Child Object permission that is required to create local public queues is granted to:

- The Domain Administrator account
- The domain user account of the local administrator that installed Message Queuing
- Local System and Network Service accounts

To permit all other users to create local public queues by default, Message Queuing uses Message Queuing service elevation of privilege. The effect of this is that if the user account does not have Active Directory permissions for this action, the user account privilege is elevated and the local public queue is created using Message Queuing service's Local System account. This is the default Message Queuing setting. If you want to disable this behavior, and restrict the creation of local public queues, set the following key:
**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\EnableCreateQueueThroughService**. To disable the key, set the DWORD value to 0. After setting the registry, restart the Message Queuing service for the changes to take effect.

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

Default security access for public queues gives everyone permission to send messages to a public queue. Specific permissions must be granted for read access.

The advantage of public destination queues is that they are registered in the directory service, where they can be located by other Message Queuing applications. Public queues are persistent, and their registration information can be backed up on the enterprise, making them good for long-term use.

# Transactional and nontransactional queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Transactional and nontransactional queues**
A transactional queue is one that only contains transactional messages, which are messages sent within a transaction. The requirements for sending and retrieving messages are based on the type of queue used (transactional or nontransactional) and the queue's location (local or remote).

**Sending messages**
When using transactions to send messages, the messages can be sent to any transactional queue regardless of where the queue is located, but only transactional queues can accept messages sent in the context of a transaction. In a similar way, nontransactional queues can only accept messages sent outside the context of a transaction. Note that only transactional messages are delivered with Exactly-Once In-Order Delivery (EOD) guarantees. For more information on transactional messaging, see Transactional messaging.

**Receiving messages**
When using a transaction to retrieve messages, the queue must be a local transactional queue. You cannot retrieve a message from a remote transactional queue or from a local nontransactional queue within a transaction. Message Queuing does not provide support for transactional remote-read operations (accessing the contents of a queue from a computer other than the one on which the queue is located). However, this same functionality can be achieved by using a transactional read-response application.

Note that nontransactional receive operations are allowed from any queue (transactional or nontransactional, local or remote). In particular, a nontransactional receive operation from a transactional queue is essentially a transaction consisting of a single receive operation.

**Transactional read-response applications**
A transactional read-response application is used when you need to retrieve transactional messages from a remote transaction queue. It uses transactions to redirect the messages from the remote queues on the sending computer to local queues on the receiving computer.

A read-response application resides on a server that contains two transactional queues:

- An input queue for collecting messages from sending applications.
- A read-request queue for collecting read-requests from the receiving application.

The read-response application does the following in a single transaction:

- Retrieves a request made by the receiving application from the read-request queue.
- Retrieves a message sent by the sending application from the input queue.
- Forwards the message from the input queue to the receiving application that made the request.

In this way, the receiving application gets its messages within a transaction, without needing to perform a transactional remote-read operation.

The receiving application sends a request message to the read-request queue on the server, specifying the response queue to which the input message will be sent. This request can be made within the context of its own transaction. Note that it is independent of the transaction used to send the message back to the receiving application.

# Administration queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administration queues**

*Administration queues* are application-generated queues used to store system-generated negative and positive acknowledgment messages that are created by Message Queuing or connector applications. They are specified programmatically in the original messages by the sending application. Any available nontransactional queue can be specified as an administration queue. Administration queues do not contain administrative messages, which are stored in internal private queues. For more information on acknowledgment messages, see Message types.

The system-generated acknowledgment messages returned to these queues can indicate whether a message arrived at its destination queue, whether it was retrieved from the destination queue, or both. Each acknowledgment message contains information that describes what triggered the acknowledgment and to which message the acknowledgment refers. When requesting acknowledgment messages, the sending application must specify the queue to be used as the administration queue and the acknowledgment level of the message. Typically, an administration queue is a local queue so that the sending application can read the acknowledgment messages locally.

# Response queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Response queues**

*Response queues* are application-generated queues used to store application-generated response messages that are typically returned by an application that is reading messages from a queue. Response queues are specified programmatically by the sending application when a message is sent. Any available queue can be specified as a response queue. Response messages are not the same as administrative response messages, which are stored in an internal private queue.

The response messages returned to these queues must be understood by the application returning the response messages and the application reading the response messages. Message Queuing has no control over what information is sent in response messages. When requesting response messages, the sending application must specify the queues that it wants to use as response queues. Typically, response queues are local queues so that the sending application can read the response messages locally.

When a response queue and an administration queue are both needed, their functionality can be combined into a single queue. However, because all administration queues must be nontransactional, this queue will accept only nontransactional messages. For more information on administration queues, see Administration queues.

# Report queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Report queues**

*Report queues* are application-generated queues used to store report messages generated by Message Queuing or connector applications when tracing is requested by the sending application. They indicate the route that a message took to its destination on a remote computer or the route that a test message took to a test queue on a remote computer. There can be only one report queue per computer.

All report queues must have the following label and queue type identifier:

- Label: MQReport Queue
- Queue type identifier (GUID):

```
{55EE8F32-CCE9-11CF-B108-0020AFD61CE9}
```

Report queues are typically created by the Message Queuing administrator using Active Directory Users and Computers, but they can also be generated by an application as long as the correct label and queue type identifier is specified when the queue is created. Message Queuing uses the label and queue type identifier to locate and open the appropriate queue when sending report messages.

Report queues are not limited to report messages. Application-generated messages can be sent to report queues as well. When a message is sent with tracing enabled, a report message is generated each time the message leaves or arrives at a Message Queuing computer. Report messages are sent to a report queue only if message route tracking is enabled on the host computer and the message route tracking flag is set. Typically a single report queue is specified by the source computer where the message is sent. However, each Message Queuing computer can specify its own report queue.

# System Queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**System Queues**

Queues that are created by Message Queuing are referred to as system queues. Applications can read the messages in system queues, but cannot address messages to them. For information on different types of system queues, see the following links:

- Queue journals contain copies of all messages removed from a queue. Computer journals contain copies of all messages sent from a computer. For more information, see Journals.
- Message Queuing uses internal private queues as interim queues for storing and forwarding messages in transit to a destination queue. These queues are not published in Active Directory. Outgoing queues are local internal queues, generated automatically to store messages sent to remote queues. For more information, see Internal private queues and outgoing queues.
- Dead-letter queues contain messages that cannot be delivered. Message Queuing provides a transactional dead-letter queue and a nontransactional dead-letter queue. For more information, see Dead-letter queues.
- Connector queues are used for cross-platform messaging. For more information, see Connector queues.

# Journals

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Journals**

Queue and computer journals are system queues that are automatically created by Message Queuing. When a sending application sends a message, a copy of the message is optionally stored in the computer journal of each computer it passes through. When a receiving application removes the message from the destination queue, a copy of the message is optionally sent to the queue journal. Applications can only read the messages in queue and computer journals. They cannot send messages to queue or computer journals. In contrast, Message Queuing never removes messages from queue or computer journals. It is up to the application reading the messages to clear the messages from the queues. If an application uses journaling and no provisions are made to clean up the journals used, Message Queuing disk storage will quickly be consumed.

**Computer journals**

A computer journal is created whenever a Message Queuing computer is added to a Windows Server 2003 enterprise. It is used to store copies of all messages successfully sent from the computer. Message Queuing stores copies of the messages sent from the computer whenever positive source journaling is specified by the sending application. For more information, see Source journaling.

Computer journals can be referenced by a machine format name or a direct format name. The following lines show the syntax of these format names:

MACHINE=ComputerGUID;JOURNAL

DIRECT=AddressSpecification\SYSTEM$;JOURNAL (Introduced in MSMQ 2.0)

For instructions on setting the computer journal size, see Set the journal storage size for computers.

A computer journal is not created for computers operating in workgroup mode or virtual servers in clusters. You can set the computer quota of a workgroup computer using the **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache\MachineJournalQuota** registry key. To change the quota for a virtual server, you need to change the analogous registry key for clustered servers. After setting the registry, restart the Message Queuing service for the changes to take effect.

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

**Queue journals**

Queue journals are created whenever an application queue is created. Message queuing creates the queue journal on the same computer where the messages for the application-generated queue are stored.

Message Queuing stores copies of the messages removed from the application queue whenever target journaling is enabled on the queue. This does not include messages removed from the queue when their time-to-be-received elapses or messages purged from the queue in the Message Queuing MMC snap-in. For more information, see Target journaling.

For instructions on setting the queue journal size, see Set the journal storage size for queues.

# Internal private queues and outgoing queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internal private queues and outgoing queues**

The following sections describe internal private queues and outgoing queues.

**Internal private queues**

*Private queues* are queues that are not published in Active Directory. Private queues are accessible only by Message Queuing applications that know the full path name or format name of the queue. Internal private queues are not listed in the Private Queues folder, and any information on these queues and their contents can be obtained only programmatically. You cannot delete internal private queues.

**Outgoing queues**

*Outgoing queues* are local internal queues that are used to store messages sent to remote queues. Messages can be stored in outgoing queues under offline conditions and then sent to target queues on remote computers when connectivity is re-established. Outgoing queues are listed in the Outgoing Queues folder. These queues are generated automatically and cannot be created or deleted manually. Outgoing queues can be in any of the following states:

- **Connected**. The outgoing queue is connected to the queue manager on the computer. This state indicates the local queue manager is sending messages and that the messages will be going out immediately.
- **Disconnected**. The outgoing queue is not connected nor is it in the process of being connected. This state indicates that the queue manager on the computer is not sending messages to the destination queue nor in the process of connecting the applicable outgoing queue.
- **Disconnecting**. The queue manager is in the process of disconnecting the applicable outgoing queue, or the computer is in the process of disconnecting from the applicable session. The queue manager cannot send messages to the destination queue in this state.
- **Inactive**. The applicable outgoing queue is in a state in which the queue manager does not attempt to connect it to any session.
- **Local connection**. The queue is a local queue on the computer. This state indicates that sending messages to the queue does not require a network connection.
- **Locked**. The outgoing queue is in the state defined by hardened MSMQ mode for outgoing queues that correspond to destination queues specified by non-HTTP format names. Messages in locked outgoing queues are not transmitted to their destinations.
- **Need Validation**. The queue was opened when the queue manager was offline.
- **On hold**. The outgoing queue is paused.
- **Waiting**. The queue manager is in the process of connecting the outgoing queue.

# Dead-letter queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Dead-letter queues**

Undeliverable and expired nontransactional messages are stored in nontransactional dead-letter queues. Undeliverable and expired transactional messages are stored in transactional dead-letter queues. Such messages are stored in the applicable queues on the computer on which the message expired (or failed to be delivered or forwarded). This can be the source computer, the target computer, or an intermediate Message Queuing server.

There is one nontransactional dead-letter queue and one transactional dead-letter queue on each Message Queuing server and independent client. These queues are created automatically and are listed in the System Queues folder of the applicable computer in MMC snap-ins.

Messages are considered undeliverable when the following conditions apply:

- The destination queue is unknown by the sender.
- The message exceeds its maximum number of allowed hops.
- The queue storage limit, or quota, is reached.
- A nontransactional message is sent to a transactional queue or vice versa.
- For transactional messages, the sending or receiving transaction failed. However, when the receiving transaction fails, these messages are not placed in the transactional dead-letter queue.

Messages are considered expired when:

- The lifetime of messages on the network, called the time-to-reach-queue (TTRQ), elapses. This can be set programmatically by a Message Queuing application.
- Their time-to-be-received (TTBR) elapses. This is always set programmatically by a Message Queuing application.

When a sending application is not allowed to send messages to a particular queue because of a lack of the necessary access rights, those messages are not stored in the nontransactional dead-letter or transactional dead-letter queues. However, negative acknowledgment messages may be generated if the original messages are so configured.

Messages in nontransactional dead-letter and transactional dead-letter queues are always written to disk and are therefore recoverable.

For more information on transactional messaging, see Transactional messaging.

# Connector queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Connector queues**

*Connector queues* receive messages sent to destinations on foreign computers. Such queues reside on Message Queuing servers running a connector application, which enables the foreign computers to communicate with other Message Queuing computers in your network. Such servers typically have two connector queues: one used for transactional messages and one used for nontransactional messages. When a Message Queuing client sends a message to a foreign computer, the message is first routed to a connector queue. The connector application then retrieves, converts, and forwards the message to the foreign computer.

Connector queues are created implicitly. They are not published in Active Directory and are transparent to the Message Queuing applications sending messages to a foreign queue. From the standpoint of applications, sending a message to a foreign queue is identical to sending a message to a Message Queuing queue. A *foreign queue* is a part of the foreign messaging system but is exposed to the Message Queuing environment.

For more information on communicating with foreign computers, see Foreign computers and sites.

# Queue names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Queue names**

The following sections describe the various ways to designate Message Queuing queues. Queues are specified using *path names*, *format names*, and *queue aliases*. For more information about the types of queues mentioned here, see Queues.

**Path names**

The path name of a queue consists of the name of the computer hosting the queue and the name of the queue separated by a slash in the form *ComputerName**QueueName*, except for a public queue located on the local computer where the path name is in the form .*QueueName*. The names of private queues are prefixed by the string private$ separated by a slash. Thus, the path name of a private queue has the form *ComputerName*\private$*QueueName* or .\private$*QueueName* for a private queue located on the local computer.

**Format names**

Format names are used mostly by Message Queuing applications to perform tasks programmatically. However, they are also used in some administrative tasks (for example, in creating queue aliases and multiple-element format names), and administrators must have an understanding of format names to make correct decisions. The following sections describe most of the types of format names used in Message Queuing.

**Direct format names**

Direct format names are used to reference public or private queues without accessing the directory service. Message Queuing can use the information provided by direct format names to send messages directly to their destination. Thus, direct format names can be used to send and receive messages in workgroup mode, to send messages to computers on the Internet, and to send messages directly to a computer.

**Note** Direct format names that specify the HTTP or HTTPS protocol cannot be used to peek at or receive messages, only to send them.The following is the general format of direct format names. Note that public, private, and (for MSMQ 2.0 or later clients and servers) computer system queues can be accessed directly.

Direct format names contain the string `DIRECT=` followed by the location of the queue and its name. The following is the general format of direct format names for public, private, system, and URL-named queues (note that the format name string is a null-terminated Unicode string):

```
DIRECT=ComputerAddress\PublicQueueName

DIRECT=ComputerAddress\PRIVATE$\PrivateQueueName

DIRECT=ComputerAddress\SYSTEM$;SystemQueueName

DIRECT=HTTP://URL_Address/msmq/PublicQueueName

DIRECT=HTTPS://URL_Address/msmq/PublicQueueName

DIRECT=HTTPS://URL_Address/msmq/private$\PrivateQueueName
```

The computer address can be specified in two forms: as the network address of the target computer, which includes the network protocol TCP, or as any name for the target computer that is supported natively by the underlying operating system, where OS is used to indicate that the computer's native protocol will be used. The following examples illustrate these two forms:

```
DIRECT=TCP:157.18.3.1\MyPublicQueue

DIRECT=OS:pc001.microsoft.com\MyPublicQueue
```

In the case of a URL-named queue, the queue name is preceded by the string `msmq` separated by a slash. For example, the following format name can be used to reference "MyPublicQueue"

```
DIRECT=HTTP://URL_Address/msmq/MyPublicQueue
```

**Public format names**

Public format names and direct format names are used to reference public queues. When a public format name is used, Message Queuing uses its internal routing algorithm to define the route to the destination queue. When a direct format name is used, the message is sent directly to the destination queue.

When Message Queuing detects a public format name, it queries the directory service to determine what computer is currently hosting the queue, what protocol the host computer uses, and any other information it needs to rout a message to the queue.

Public format names contain the string `PUBLIC=` followed by the identifier assigned to the queue when it was created. This identifier is the GUID listed for the queue object in Active Directory. The following is the general format used to reference a public queue and its associated queue journal:

```
PUBLIC=QueueGUID
```

```
PUBLIC=QueueGUID;JOURNAL
```

A public format name is strictly equivalent to the identifier of the queue except that a public format name is formatted as a string rather than as binary data. The following is an example of a public format name string (note that the format name string is a null-terminated Unicode string).

```
PUBLIC=308FB580-1EB2-11CA-923B-08002B1075A7
```

**Private format names**

Private format names and direct format names are used to reference private queues. When a private format name is used, Message Queuing uses its internal routing algorithm to define the route to the destination queue. When a direct format name is used, the message is sent directly to the destination queue.

When Message Queuing detects a private format name, it does not refer to the directory service for information about the queue. However, it does use the directory service to look up information about the computer for routing purposes.

Private format names contain the string `PRIVATE=` followed by the identifier of the computer where the queue is registered and a hexadecimal number that identifies the queue. The following is the general format used to reference a private queue and its associated journal:

```
PRIVATE=ComputerGUID\QueueNumber
```

```
PRIVATE=ComputerGUID\QueueNumber;JOURNAL
```

The following is an example of a private format name string (note that the format name string is a null-terminated Unicode string).

```
PRIVATE=ae0c5671-f190-12ce-ae10-00dd0114290f\0000000d
```

**Distribution list format names**

Distribution list format names are used to reference distribution lists (group objects) stored in Active Directory. The distribution list identifier is the GUID listed for the distribution list (group) object in Active Directory.

Distribution list format names contain the string `DL=` followed by the distribution list identifier. This identifier is the GUID listed for the distribution list (group) object in Active Directory. The following is the general format used to reference a distribution list with optional inclusion of the domain name:

```
DL=DistributionListGUID[@DomainName]
```

The following is an example of a distribution list format name string (note that the format name string is a null-terminated Unicode string).

```
DL=ae0c5671-f190-12ce-ae10-00dd0114290f
```

To retrieve information about a distribution list, Message Queuing tries to resolve the DL format name to the corresponding Active Directory object in three steps:

1. If the domain name is included, it tries to bind to the Active Directory group object using this domain name.

2. If this attempt fails or the domain name is not specified, it allows Active Directory to use serverless binding to find the best domain controller in the default domain.

3. If this also fails, it tries to query the global catalog.

A distribution list format name can be an element of a multiple-element format name. In this respect, the analogy to a queue format name is preserved.

For more information about distribution lists, see Distribution lists.

**Multiple-element format names**

A multiple-element format name is formed as a concatenation of one or more public, private, direct, or distribution list format names separated by commas. Thus, different kinds of format names used in Message Queuing can be used together as elements of a multiple-element format name. The following example shows a multiple-element format name that contains a direct format name, a public format name, and a distribution list format name.

```
DIRECT=ComputerAddress\QueueName,PUBLIC=QueueGUID,DL=DL_GUID
```

Note that a multiple-element format name containing an element that points to a public queue or a distribution list cannot be used when there is no access to Active Directory.

Multiple-element format names cannot contain the format names of read-only queues, such as queue journals, computer journals, or dead-letter queues. An error is returned if the format name of a read-only queue is included in the multiple-element format name.

For more information about multiple-element format names, see Multiple-element format names.

For more information about these and other types of format names, such as machine and connector format names, refer to the Message Queuing Software Development Kit (SDK). For information about accessing the Message Queuing SDK, see Additional Message Queuing resources.

# Queue aliases

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Queue aliases**

Message Queuing 3.0 provides support for referencing queues with an Active Directory object called a queue alias, which can be used to reference queues that might not be listed in Active Directory.

Queue aliases are published in Active Directory, and in this respect, they resemble public queues.

Using a queue alias provides the following benefits:

- A queue alias can point to a public, private, or direct single-element format name used to reference a queue. Messages are sent to the queue or queues named in the corresponding format name. The validity of the format name and the existence of the queue to which a queue alias points are not tested when the alias is created. A queue alias thus points to a single queue, but it does not function as a queue because messages are not sent to it, but only to the queue or queues named in the corresponding format name. For more information about format names, see Queue names.

- A queue alias can serve as a tool to reference a private queue and URL-named queues included in distribution lists. While queue aliases can point to any queue, including private and URL-named queues, distribution lists can only include explicit references to destinations published in Active Directory; that is, to public queues, queue aliases, and other distribution lists, but not to private and URL-named queues. Thus, a queue alias can serve as a tool for including private and URL-named queues in distribution lists.

- When a queue alias object is deleted, the alias is automatically removed from all distribution lists that reference the alias. A queue referenced by a queue alias can be changed without having to change a distribution list that references it.

- A queue referenced by the queue alias can be changed without changing the alias reference. The alias can be used as a constant reference, as in a distribution list, even though the queue it represents may change.

- Queue aliases can be used to reference a queue that is not listed in the directory service. This includes private queues or queues from another organization.

- Queue aliases can be used to send HTTP messages by referencing the destination queue using a direct format name.

- Multiple aliases can be created for a single physical queue.

User interfaces are provided in Active Directory Users and Computers for creating and deleting queue aliases as well as for viewing and changing their properties. For more information about performing these tasks, see Create queue aliases, View or change properties of queue aliases, and Delete queue aliases.

The queue alias is a new object in Windows Server 2003 family. Thus, it can be used only in Windows Server 2003 family domains and cannot be accessed by computers running earlier versions of Windows.

# Target journaling

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Target journaling**

*Target journaling* is the process of storing copies of incoming messages on the computer hosting the applicable destination queue. Target journaling is configured on a per-queue basis. If you enable target journaling, Message Queuing sends a copy of any message you retrieve from a queue to the corresponding queue journal. Target journaling lets you keep track of messages even after they have been retrieved. Note that this does not include messages purged from the queue. For more information, see Purge all messages in a queue. Peeking at messages does not remove them from the queue. For more details, see Managing messages.

When enabling target journaling, consider the following:

- Target journaling is enabled by setting the journaling level of the queue. This can be set both when the queue is created or later.
- You must periodically clear messages from the queue journal to remove messages that are no longer needed. Message stored in the queue journal count against the quota of the journal and the quota of the computer where the queue journal resides (the computer quota is set by the administrator).
- If the journal or computer quota is exceeded, Message Queuing will not send copies of the messages to the journal when the messages are received, and those messages will be lost.
- A copy of each message retrieved from the applicable queue is transferred to the corresponding queue journal. This does not include messages removed from the queue when their time-to-reach-queue lifetime expires. For more information, see Message timers.

A queue journal is automatically created for each queue. It is listed as *Journal messages* under the applicable queue in Computer Management and Active Directory Users and Computers. Target journaling is disabled by default. For information on how to enable target journaling, see Enable or disable target journaling.

# Source journaling

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Source journaling**

*Source journaling* is the process of storing copies of outgoing messages on the source computer. It is configured on a per-message basis and is a property that is set programmatically by a Message Queuing application. Source journaling allows you to track the messages sent by a computer, and can be used to track messages that were sent successfully, messages that could not be delivered, or both. Source journaling is disabled by default.

There are two types of sources journaling, *positive source journaling* and *negative source journaling*.

**Positive source journaling**

You can track messages that were successfully sent by the local computer, so that when you send the message to remote destination queues, a copy of the message is placed in the source computer's journal if the message is sent successfully. This is known as *positive source journaling*. Note that this does not mean the message has successfully reached its destination queues, only that the local computer was able to send the message.

When messages are sent to local destination queues, the message can be placed in one of two journal queues:

- If the local destination queue has not enabled target journaling, a copy of the message is saved to the computer's journal queue only after the message has been removed from the local destination queue. This is the default setting.

- If the local destination queue has enabled target journaling, a copy of the message is saved to the queue's journal after the message is removed from the queue. The message is not, in this case, stored in the computer's journal queue. Message Queuing never stores more than one copy of a message on a single computer at one time.

**Negative source journaling**

You can track messages that were not delivered. This type of journaling is known as *negative source journaling*. When you specify this type of source journaling, Message Queuing reacts differently depending on the type of message you send:

- For nontransactional messages, a copy of the message is placed in the dead-letter queue of any computer that cannot deliver the message. This can be the source computer, or any routing server used to pass the message on to its destination. Failure indicates that the source queue manager on the computer cannot pass the message to the next hop.

- For transactional messages, a copy of the message is placed in the transactional dead-letter queue of the source computer only if Message Queuing does not confirm that the message was retrieved from its destination queue.

**Source journaling for transactional messages**

For transactional messages sent to transactional queues, journaling can be applied by setting the journal property of a transactional message. The following property values can be set programmatically:

- The default value. Source journaling is disabled. Message Queuing does not store copies of the message in the computer journal on success nor in the applicable dead-letter queue on failure.

- *Negative source journaling* is requested. The message is stored in the applicable dead-letter queue on failure, but no copy is stored in the source computer journal on success.

- *Positive source journaling* is requested. A copy of the message is stored in the source computer journal if the message was successfully delivered to the target computer, but no copy is stored in the applicable dead-letter queue on failure.

- Both *negative source journaling* and *positive source journaling* are requested. Copies of the message are stored in the source computer journal on success and in the applicable dead-letter queue on failure. This is the recommended setting.

**Differences in source journaling for transactional messages between MSMQ 1.0 and later versions**

MSMQ 1.0 treated source journaling for transactional messaging differently than MSMQ 2.0 and Message Queuing 3.0.

MSMQ 1.0 only allows you to specify *positive source journaling* to choose if you want to save copies of messages sent successfully. MSMQ 1.0 always uses *negative source journaling*, and places a copy of a non-confirmed transactional message in the transactional dead-letter queue.

MSMQ 2.0 and later versions allow complete control over source journaling. (If you are running MSMQ 1.0 applications on an MSMQ 2.0 computer, you can force the computer to use the MSMQ 1.0 semantics by setting the **XactDeadLetterAlways** registry entry to 1. After setting the registry, restart the Message Queuing service for the changes to take effect.

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

**Managing source journals**

When you store messages in a journal or dead-letter queue, clear these queues periodically to remove messages that are no longer needed. Messages stored in journal and dead-letter queues count against the quota for the computer where the queues reside (the computer quota is set by the administrator).

When source journaling is enabled, a copy of the outgoing message is placed in the single computer journal of the source computer, which is listed as the *Journal messages* queue in the System Queues folder in Computer Management and Active Directory Users and Computers. This occurs either when the message arrives at the computer hosting the destination queue or when the message reaches a Message Queuing routing server that is one hop closer to the target computer. A Message Queuing server that routes messages for clients does not store copies of those messages in its computer journal.

Outgoing messages sent by dependent clients are placed in the computer journal on the supporting Message Queuing server.

# Messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Messages**

In the context of Message Queuing, messages consist of a body and a set of properties. Message bodies can contain text or any form of binary information, and can be encrypted. Most message properties are specified programmatically by Message Queuing applications. Some properties, such as the sender's security ID (SID) and the message ID, are specified by Message Queuing. Message properties cannot be encrypted. A message, including the body and all specified properties, can be no larger than 4 megabytes (MB) in size. For details, see the following links:

- For information about different types of message classes, see Message types.
- For information about message properties and the default lifetime of a message, see Message properties.
- For information about setting default lifetimes for messages, see Message timers.
- For information about how message priority is determined, see Message priority.
- For information about express and recoverable message delivery, see Message delivery methods.
- To learn about the characteristics of transactional messages, see Transactional messaging.
- To learn about message storage and message storage quotas, see Message storage.

# Message types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message types**

In the context of Message Queuing, messages can be of different types, and consist of a body and a set of properties. Message types, also called message classes, indicate the message function or disposition. The following message types are supported:

- **Normal messages**. All messages created by Message Queuing applications are normal messages. Normal messages can be sent to public, private, and dead-letter queues of both the nontransactional and transactional types, and copies of normal messages, which are themselves normal messages, can be sent to journals when source or target journaling is enabled. Normal messages do not contain status information; such information is contained in acknowledgment, report, and response messages.

- **Acknowledgment messages**. Acknowledgment messages report the status of normal messages. These messages are nontransactional messages that are sent to the administration queues specified in the original messages by the sending application. Acknowledgment messages can be either positive or negative. A positive acknowledgment message indicates that the original message arrived at the target queue and that it was received (removed by the receiving application). A negative acknowledgment message indicates that the original message did not reach the target queue before its time-to-reach-queue (TTRQ) elapsed, did not reach the target queue for a specified reason, that it was not received before its time-to-be-received (TTBR) elapsed, or was not received for a specified reason.

- **Report messages**. Report messages indicate the route followed by a message as it proceeds from a source computer with tracking enabled to its destination on a remote computer. Report messages are sent to the report queue on a source computer for all messages sent from that computer with the message route tracking flag set or for only test messages sent from it.

- **Response messages**. Response messages are messages sent by a receiving application to response queues specified by a sending application.

For more information on the various types of queues that these messages can be sent to, see Queues.

For more information on the various reasons why messages do not reach the target queue or are not received that can be specified in the message class, refer to the Message Queuing Software Development Kit (SDK).

# Message properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message properties**
In the context of Message Queuing, messages consist of a body and a set of properties, and can be of various types. Message bodies can contain text or any form of binary information, and can be encrypted.

Every message has properties. Some message properties are read-only, and the majority can only be modified programmatically. For more information, refer to the Message Queuing Software Development Kit (SDK).

**Default message lifetime**
Through the Computer Management snap-in, you can set the default lifetime of messages in your network, also known as the time-to-reach-queue (TTRQ) value. Messages are considered expired when the TTRQ value is reached. If you do not specify a TTRQ value, the following default setting apply:

- 4 days for messages sent from the following types of computers:
    - Computers running Message Queuing 3.0 in workgroup mode.
    - Clients running Message Queuing 3.0 that belong to a completely new Windows Server 2003 domain.
    - Clients running Message Queuing 3.0 that join an existing domain provided weakened security was not enabled for that forest.
- 90 days for messages sent from the following types of computers:
    - Any computer in a migrated enterprise and computers running MSMQ 1.0 or MSMQ 2.0 in Windows 2000 and Windows NT 4.0.

For information on how to set the lifetime of messages on your network, see Set the default lifetime for messages.

You can view the properties of messages, and you can choose which message property columns appear by default in the details pane (usually the right pane) of the snap-in and in what order. For instructions on how to view message properties, see View message properties. You can choose the columns to be displayed by selecting a message in the Computer Management snap-in, and then clicking the **View** menu, and selecting **Add/Remove Columns**. The following message properties are available for selection.

| Message property | Explanation |
|---|---|
| Admin Queue | Specifies the name of the administration queue used for system-generated acknowledgment messages. |
| Admin Queue Name Length | Indicates the length (in Unicode characters) of the format name buffer of the administration queue. |
| Application-specific | Specifies application-generated information such as single integer values or application-defined message classes. |
| Authenticated | Indicates whether the message was authenticated at the request of the sending application. |
| Correlation | Specifies an application-generated correlation identifier that the receiving application can use to sort messages. |
| Delivery Mode | Specifies how Message Queuing delivers the message to the queue, in either recoverable, or express mode. For more information, see Message delivery methods. |
| Destination Queue Format Name | Indicates the name of the destination queues to which the message was sent. |

| | |
|---|---|
| Destination Queue Format Name Length | Indicates the length of the format name buffer (in Unicode characters) allocated by the receiving application. |
| Encrypted | Specifies whether the message body of a private message was encrypted. |
| Encryption Algorithm | Specifies the encryption algorithm used to encrypt the message body of a private message. |
| Hash Algorithm | Identifies the hashing algorithm that Message Queuing uses when authenticating the message. |
| Label | Specifies the label of the message. |
| Label Length | Identifies the length (in Unicode characters) of the message label buffer. |
| Message ID | Indicates a unique identifier for each message in your enterprise. This includes application-generated and system-generated messages sent by Message Queuing. |
| Message Lookup ID | Indicates the lookup identifier of the message. The receiving application can use this identifier to peek at or retrieve the message from the queue without navigating the queue. |
| Priority | Indicates the message priority. A higher number indicates a higher priority message. The default value is 3. |
| Recipient Queue | Specifies the queue to which application-generated response messages are returned. |
| Recipient Queue Name Length | Indicates the length (in Unicode characters) of the response queue format name buffer. This is used by the receiving application to allocate a buffer for the response queue specified by the sending application. |
| Response Queue | Specifies the response queues used for returning application-generated response messages. |
| Response Queue Format Name | Specifies the response queues used for returning application-generated response messages. |
| Response Queue Format Name Length | Indicates the length (in Unicode characters) of the format name buffer allocated by the receiving application. |
| Response Queue Name Length | Indicates the length (in Unicode characters) of the format name buffer of the response queue format name buffer. |
| Sender ID | Identifies who sent the message, and used primarily by the receiving queue manager when authenticating a message. |
| Sender ID exists | Identifies who sent the message, and used primarily by the receiving queue manager when authenticating a message. |
| Sender ID Length | Indicates the length of the Sender ID. |
| Size | Indicates the size of the message body. |
| Source Computer | Provides the globally unique identifier (GUID) of the computer that sent the message. |
| Time Arrived | Indicates the date and time that the message was received. |
| Time Sent | Indicates the date and time that the message was sent by the source. |
| Tracked | Specifies whether Message Queuing traces the route of the message. |

# Message timers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message timers**

Message Queuing provides two timers to help you maintain better control of your messages: a *time-to-reach-queue* and a *time-to-be-received* timer.

The time-to-reach-queue (TTRQ) timer determines how long a message has to reach the queue manager of the target queue. Typically, this timer is set to a value less than the time-to-be-received setting.

The time-to-be-received timer determines how long a message remains in the system, starting from the time the message is sent to the time it is removed from the target queue.

Through the Computer Management snap-in, you can set the TTRQ value. If you do not specify a TTRQ value, the following default settings apply:

- 4 days for messages sent from the following types of computers:
    - Computers running Message Queuing 3.0 in workgroup mode.
    - Clients running Message Queuing 3.0 that belong to a completely new Windows Server 2003 domain.
    - Clients running Message Queuing 3.0 that join an existing domain provided weakened security was not enabled for that forest.
- 90 days for messages sent from the following types of computers:
    - Any computer in a migrated enterprise and computers running MSMQ 1.0 or MSMQ 2.0 in Windows 2000 and Windows NT 4.0.

For information on how to set the lifetime of messages on your network, see Set the default lifetime for messages.

The default setting for the time-to-be-received timer is that no limit is set.

When both timers are used and the time-to-be-received timer is set to a smaller value than the TTRQ timer, the time-to-be-received timer takes precedence. For example, if the time-to-be-received timer is set to 3 minutes and the TTRQ timer is set to 10 minutes, Message Queuing will discard the message after 3 minutes. A message is not allowed to remain in the system longer than the time specified by the time-to-be-received timer.

When a timer expires, Message Queuing discards the message. However, Message Queuing can also send a copy of the discarded message to the applicable dead-letter queue or it can send a negative acknowledgment message to an administration queue.

Discarded messages can be sent to dead-letter queues if the journal property of the message specifies a dead-letter queue. Copies of nontransactional messages are sent to the dead-letter queue on the computer where the timer expired. Copies of transactional messages are sent to the transactional dead-letter queue on the source machine.

Acknowledgment messages are returned if a timer expires and the acknowledgment property of the message specifies full or negative acknowledgments. Message Queuing sends the appropriate negative acknowledgment message to the administration queue specified by the message.

# Message priority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message priority**
The order in which messages are sent to a queue on the destination computer, their *priority*, is determined first by the queue priority (also known as the base priority), and then by the message priority. The queue priority for public queues listed in Active Directory can be set to any number from -32768 to 32767. The default queue priority is 0.

Message priority affects how Message Queuing handles the message while it is in route, as well as where the message is placed in the queue. Higher priority messages are given preference during routing, and inserted toward the front of the queue. Message Queuing applications can set the message priority for nontransactional messages in the range from 0 to 7, where 7 represents the highest priority. The default value is 3. Messages with the same priority are placed in the queue according to their arrival time. Message Queuing automatically sets the priority level of transactional messages to 0.

Messages are sent to public queues based on the sum of the destination queue's priority and each message's priority. Messages with the highest resultant priority value are sent first.

Private queues do not support queue priority. Messages are sent to private queues based only on message priority.

# Message delivery methods

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message delivery methods**
Two basic message delivery methods are supported: express and recoverable. The choice between the two is essentially a choice between better performance with minimal resource use (express delivery), and reliability and recovery after a failure (recoverable delivery). The message delivery method is specified programmatically by the sending application. Recoverable messaging can also be further differentiated into transactional and nontransactional messaging.

**Express messaging**
Using express messaging means that messages are stored in RAM memory during routing and delivery, providing extremely fast performance but no recoverability if any computer that the message passes through fails. Notably, express messages are lost whenever the Message Queuing service is stopped.

Too many express messages accumulating in RAM may cause paging out to disk, leading to performance degradation. Thus, it is not recommended to store more express messages than physical memory can hold.

Express messages can, however, survive a network failure. For example, if a Message Queuing application sends express messages through a Message Queuing server and the network link between the Message Queuing server and the destination computer fails, the Message Queuing server continues to store the messages in memory. However, if the Message Queuing server fails before the network link is restored, the express messages are lost.

**Caution**

- If you are running Message Queuing in a server cluster, express messages will be lost during a failover from one node to another.

**Recoverable messaging**
Using recoverable messaging means that messages are written to disk during routing and delivery, making delivery somewhat slower than the use of express messaging, but ideal when failures cannot be tolerated or when computers shut down while messages remain in queues (such as for mobile clients on laptop computers). If a computer fails or is shut down while sending a message, the message is stored on disk. When the computer is restarted and the Message Queuing service restarts, the sending process is automatically resumed.

All recoverable messages are stored on disks in memory-mapped files. Performance gains can be realized by employing several physical disks to store recoverable messages. For more information, see "Disk considerations" in Planning overview.

# Transactional messaging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Transactional messaging**

Sending and receiving applications can indicate that they want to send or retrieve messages within the context of a transaction, known as transactional messaging. Sending or retrieving messages outside the context of a transaction is referred to as nontransactional messaging. Transactional message is only used when you want to perform several tasks in such as way that all tasks, including non-Message Queuing operations, will all succeed or all fail.

Transactions are characterized by their ACID (atomicity, consistency, isolation, durability) properties:

- **Atomicity.** Atomicity refers to the all-or-nothing behavior of a transaction. When a transaction consists of a series of actions, all the actions are treated as a single operation that is either completed successfully or not executed at all.
- **Consistency.** Transactions transform a system from one valid state to another.
- **Isolation.** Multiple transactions that are executed concurrently have no effect on one another during their execution. It appears to each transaction that the other transactions are executed either before or after it.
- **Durability.** After transactions are executed, their effects are permanent despite system or network failures.

When transactional messaging is used, the sending or receiving application has the opportunity to commit to the transaction (all the operations have succeeded), or to abort the transaction (one of the operations failed) so all changes are rolled back. When a transaction is aborted all operations are rolled back to the state when the transaction was invoked.

When sending messages within a transaction, all the messages are sent together and in the order they were sent, or they are not sent at all. Message Queuing guarantees that all messages sent within the transaction will arrive exactly once if they can be delivered, and in the order that they were sent. When retrieving messages within a transaction, all the messages are removed from the queue or they are placed back in the queue in their original position in the queue. Note that transactional messages cannot be sent to a multicast address, and that peeking at messages make no sense in the context of a transaction.

In Message Queuing, a transaction can encompass multiple sending and receiving operations. However, the same message cannot be sent and received within a single transaction. In a case where a transaction consists of sending multiple messages, either all the messages are transmitted in the order in which they were sent (if the transaction is committed), or none of the messages are transmitted (if the transaction is aborted). If multiple messages are sent to the same queue, Message Queuing guarantees that they arrive in the destination queue exactly once and in the order in which they were sent within the transaction. For more information, see Sending transactional messages to multiple destinations.

**Message order**

Messages that are to be delivered in the order in which they were sent belong to the same *message stream*. Each message in a stream carries a *sequence number*, identifying the message in the stream, and a previous number, which equals the sequence number of a previous message in the stream. The message whose sequence number equals the previous number of a message must be present in the destination queue before the message can be delivered. The first message in a stream always carries a previous number equal to zero.

Likewise, a transaction can consist of removing multiple messages from a queue. The messages are removed from the queue only if and when the transaction is committed; otherwise, they are returned to the queue and can be subsequently read during another transaction.

In addition, Message Queuing guarantees that messages sent in consecutive transactions initiated from the same computer to the same destination queue will arrive in the order in which the transactions were committed. Moreover, as in the other cases of sending transactional messages, each message, if it arrives at the destination queue, will arrive exactly once.

**Transaction performance**

Transactional messaging is slower than nontransactional messaging because messages are written to disk more than once and greater computer resources are required. Transactional messages can be sent only to transactional queues on the local computer or a remote computer (and nontransactional messages cannot be sent to transactional queues). When messages are received within a transaction, they can be retrieved only from a local transactional queue. However, messages can be retrieved

in nontransactional operations from local and remote transactional queues.

In the general case, there are two ways to implement transactional messaging, as follows:

- As external transactions
- As internal transactions

**External transactions**
External transactions are used if resources other than those that Message Queuing can provide need to be part of the transaction. They are coordinated by the Microsoft Distributed Transaction Coordinator (DTC), which is now part of Component Services (formerly known as Transaction Server). Message Queuing then serves as a *resource manager* under the control of the DTC. Using the DTC as the *transaction manager* for Message Queuing, transaction-based applications can be implemented, ensuring that all actions either all succeed or all fail. For example, a Message Queuing application can send a message and update a database in the same transaction. The DTC ensures that either both actions succeed (both actions are committed) or neither action is executed (both actions are aborted). The DTC is automatically installed on each computer during Windows Server 2003 family operating system installation for local transactions. To enable network DTC required for Message Queuing transactions across the network, see Enable network DTC access.

For more information on the DTC, see Using Component Services.

You can use external transactions to send or receive a message that is part of any COM-compliant resource operation and to send or receive multiple messages that are part of a single transaction.

**COM+ Transactions**
When your application is running in the COM+ (Component Services) environment, Message Queuing can send and receive messages within a COM+ transaction by using the current COM+ transaction object if it is available. Message Queuing, COM+ (Component Services), and DTC work together for this type of transaction.

**XA-Compliant Transactions**
Message Queuing can work in transaction systems that have an XA-compliant transaction manager. Because of this, DTC is used as a transaction process monitor and Message Queuing is a resource manager.

**Internal transactions**
Because they have less overhead, internal transactions execute much faster than external transactions coordinated by the DTC. Use internal transactions if there is no need to coordinate Message Queuing operations with operations involving external resources that require the use of the DTC. Message Queuing provides a transaction coordinator that supports only itself as a resource manager.

You can use internal transactions to do the following:

- Ensure that a message is delivered only once (called *exactly-once delivery*).
- Ensure that all messages sent from one computer to another are delivered in order (called *in-order delivery*).
- Using acknowledgments, confirm that the transaction was successful (messages were received). You can also use acknowledgments to confirm that messages reached (or were removed from) the destination queue.

There is a special type of internal transaction that Message Queuing provides when sending a single message:

**Single-message transactions**
When sending a single message, Message Queuing provides a single-message transaction that implicitly uses the Message Queuing internal transaction, and consists of a single operation performed on a single message. A single-message transaction can be used to request guaranteed exactly-once delivery of a single message sent to a transactional queue. The single-message transaction provides the best performance of all transaction types for sending a message. In addition to exactly-once delivery, a single-message transaction is faster than an external transaction.

For more information on transaction types, see "Microsoft Platform Software Development Kit (SDK)" in Additional Message Queuing resources.

**Order acknowledgment messages for transactional HTTP messages**
Transactional messages are to be delivered exactly once and in the order in which they are sent. When transactional messages

are sent over HTTP from the internal network of an enterprise to the Internet, the internal name of the sender's order queue must be translated into a logical address using a mapping to direct the order acknowledgment messages from the receiving computer outside the sender's enterprise, back to the sender's order queue. For more information, see Store and forward servers.

# Message storage

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message storage**
Message storage size limits, or quotas, specify the maximum cumulative storage size for all messages contained in a particular queue or on a computer. There are additional storage size limits, or quotas, that specify the maximum cumulative storage size for all journal messages contained in a particular queue journal or for messages contained in the computer journal and the dead-letter queues.

By understanding how messages are stored locally, you can determine the disk space being used to store messages, and you can inspect the properties of local queues.

The following sections discuss these concepts.

**Quotas for messages**
Quotas for messages can be divided into computer quotas and queue quotas. These quotas specify the cumulative size for all messages (including journal messages) contained in a particular queue or on a computer.

**Computer quotas**
Computer quotas specify the cumulative limit for all messages stored in destination queues, outgoing queues, and queue journals on a computer. Computer quotas are not applied towards internal private queues. Unlike queue quotas, when a computer quota is reached, no messages can be delivered to any queue or journal on the computer. No negative acknowledgment message is generated, and an error is returned for messages sent locally, while a remote computer tries to resend until one or more messages are removed (read) from a queue on the computer. Computer quotas are set for Message Queuing servers and independent clients. Computer quotas set on Message Queuing servers that support dependent clients apply both to the server and to all its dependent clients. Computer quotas for message storage and journal storage are set separately. For instructions on setting the message storage size for a computer, see Set the message storage size for computers. For instructions on setting the journal storage size for computers, see Set the journal storage size for computers. Following a new installation of Message Queuing, or an upgrade from Windows 2000, the default limit for these storage sizes is 8 gigabytes (GB).

**Queue quotas**
A queue quota specifies the cumulative storage size (number of bytes) available for all messages stored in a particular queue, whereas a computer quota limits the total storage available for all messages in destination queues, outgoing queues, and queue journals on the computer. When a queue quota is reached, messages can no longer be sent to the queue until one or more messages are retrieved (removed from the queue by a Message Queuing application). Queue quotas can be set for all queues. Queue quotas for message queues and journal queues are set separately. For instructions on setting the quotas for regular messages, see Set the message storage size for queues. For instructions on setting the storage size for the computer journal and dead-letter queues, see Set the journal storage size for queues.

Queue quotas and computer quotas can be set independently, and only on Message Queuing servers and independent clients. Message Queuing servers and independent clients can store messages only up to a certain limited maximum number on each computer, which depends on several factors. Quotas set on Message Queuing servers that support dependent clients apply both to the server and to all its dependent clients, since dependent clients do not store messages locally. Thus, a Message Queuing server supporting several dependent clients can hold no more than its own maximum number of messages divided by the number of dependent clients. For example, if the computer quota for the server is 50 megabytes (MB) and one dependent client creates 50 MB of messages, the server and the other dependent clients supported by the server cannot send messages until the total message storage size drops below that quota.

The computer quota is enforced first, regardless of the cumulative queue quotas. For example, if you specify a 10 MB quota for each of six queues on a computer and a 50 MB quota for the computer, Message Queuing enforces the 50 MB computer quota first before each queue quota is reached. However, each queue quota still prevents any one queue from storing more than 10 MB of messages.

**Local message storage**
The files used to store messages are located in the Msmq\Storage and Msmq\Storage\Lqs folders (under \Windows\System32 by default).

Messages are stored in memory-mapped files (.mq files) on the hard disk of the source computer, destination computer, or intermediate Message Queuing server with routing services enabled. Disk space is allocated for both express and recoverable

messages in 4 MB increments as needed. For each recoverable message, the memory-mapped message file is written to disk, ensuring that the message is recoverable. For express messages, memory-mapped files are written to disk only when the computer does not have sufficient resources to hold the message in memory.

The computer sending the message creates the memory-mapped files. Message Queuing servers used to route messages also store messages in memory-mapped files. If source journaling is used, the memory-mapped files are located on the source computer. If target journaling is used, they are located on the destination computers. Message Queuing servers store all message files for the dependent clients they support.

To help safeguard your system against attacks by intruders, deleting, editing, or adding any file under the Msmq\Storage folder generates an audit entry in the security log.

The memory-mapped files can be stored using any file system supported by Windows Server 2003 family. However, because the NTFS file system for Windows Server 2003 family provides more reliable file recovery after a power failure or system crash, it is highly recommended for Message Queuing servers. In addition, it is recommended that you back up the folders where the message storage files are kept for fault tolerance. If the hard disk that contains the message files fails and if the data cannot be recovered using a backup program or the message backup/restore utility provided with Message Queuing, the messages will be lost. For more information on the message backup and restore utility, see Backing up and restoring messages.

**Files in the Msmq\Storage folder**
Files with an .mq extension represent memory-mapped files in which messages are stored. The .mq file name extension is prefixed with a letter representing the type of message it is currently storing, as follows:

- **r** — Express messages (Express messages are normally stored in RAM, not on disk. This file is used when there is no more available RAM.)
- **p** — Recoverable messages.
- **j** — Journal messages.
- **l** — Message logger file for interpreting recoverable and journal message files.

An .mq file does not represent one particular queue. Messages from multiple queues can be stored in one .mq file, and messages from a single queue can be stored in multiple .mq files. However, a single message cannot span multiple files. This is the reason for the 4 MB limit in message size.

The QMLog file is a history of all transactional processing. When the log is full, the oldest data are overwritten as new entries are made. To allow such continuous cyclic writing to the QMLog file, the message sequence and the state of all active transactions are periodically saved to the following files, respectively:

- MQInSeqs.lg1
- MQInSeqs.lg2
- MQTrans.lg1
- MQTrans.lg2

The MQInSeqs.* files are referred to as the message logger files, and the MQTrans.* files are referred to as the transaction logger files.

The *.lg1 and *.lg2 files are written to in an alternating pattern to ensure that one file in each pair is recoverable if the other is corrupted during a write operation.

**Files in the Msmq\Storage\Lqs folder**
Files in the Lqs (local queue storage) folder are text files that describe properties of the local queues.

File names with the *nnnnnnnn.hhhhhhhh* format represent private queues. The *nnnnnnnn* portion is the queue number (which is included in the direct format name of the queue). The *hhhhhhhh* extension is a hash of the queue path and is used for quick lookup.

File names with the *aaaaaaaabbbbccccddddeeeeeeeeeeee.hhhhhhhh* format represent local public queues. The left-hand 32-character portion of the file name is the globally unique identifier (GUID) of the queue. The *hhhhhhhh* extension is a hash of the queue path.

Warning

- If you delete the files in the Lqs folder, you will destroy the internal private queues, in which messages needed for internal Message Queuing operations are stored and without which the Message Queuing service cannot start running. Check for Message Queuing events in Event Viewer. The most common symptoms of the loss of these queues are events with IDs 2020 and 2052.

- The message store location must only be modified using the **Storage** tab, in **Computer Management**, **Services and Applications**, Message Queuing properties.

# Triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Triggers**

Message Queuing Triggers is a service that you can use to associate the arrival of each incoming message at a queue with a response. The response depends on the contents of the message and may invoke either a COM component or a stand-alone executable program. For details, see the following links:

- For an overview of what triggers are and how to define them, see Triggers overview.
- For more detailed information on conditions and actions you can specify when creating rules for triggers, see Trigger rules.

# Triggers overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Triggers overview**

Message Queuing Triggers is a service that allows you to associate the arrival of each incoming message at a queue with a response that depends on the contents of the message and may invoke either a COM component or a stand-alone executable program. You can use Message Queuing Triggers to define business rules that can be invoked in response to such messages without any additional programming. Ordinarily, to provide such message-handling functionality, application developers would have to create a receiving application that invokes specific behavior on a per-message basis. By using Message Queuing Triggers, specific behavior can be invoked at the queue level. Application developers no longer need to write any infrastructure code to provide such message-handling functionality. Triggers can be associated with queues located on the local computer and a remote computer.

Triggers are not installed by default during Message Queuing installation. For information on installing triggers, see Installing Message Queuing triggers.

**Triggers security**

By default, the Message Queuing Triggers service runs under the Network Service account. Services that are running under this account interact with the network using the credentials of the computer account, and this account has the same level of access to resources and objects as members of the Users group accounts. You can specify a less secure account for the Triggers service; for example, the Local System account, allowing access to the entire domain. Note that the Triggers service running under the Network Service account does not interact with the desktop. If this functionality is required, you can configure the Triggers service to run under the Local System account and enable interaction with the desktop in Computer Management.

The account used must be able to access the queues with which triggers are associated. Specifically, the account must be granted the **Log On As A Service** permission and the **Peek Message** permission for all applicable queues. If the message processing type of a trigger is "retrieval" or "transactional retrieval," the user must be granted the **Receive Message** permission for that queue. If you change the security settings for a queue to enable a trigger to monitor it, you must stop and start Triggers to update the service's configuration before the trigger will start to work. Note that executable files invoked by triggers are invoked with the same security privileges as those assigned to the Triggers service.

The Triggers service must run under an account with administrator permissions.

**Serialized and nonserialized triggers**

Message Queuing triggers can be serialized or nonserialized. The type of triggers you define for a specific queue will affect how the Message Queuing Triggers service processes incoming messages. If you create multiple triggers for a particular queue and at least one of them is serialized, all the triggers will behave as serialized triggers, even if they are not defined as serialized. Triggers with the transactional-retrieval processing type can only be serialized.

A serialized trigger processes messages in the order in which they arrive, testing each message for the conditions of all the rules associated with it and performing all the applicable actions before moving on to the next message. When serialized triggers are defined, the throughput of the Message Queuing Triggers service is lower than for nonserialized triggers because the messages are processed one at a time and the service must wait until all the operations performed on each message are completed.

Nonserialized triggers process messages faster, but the messages are not necessarily processed in the order in which they arrived. Nonserialized triggers provide a faster throughput because the Message Queuing Triggers service does not wait for all the operations performed on one message to be completed before starting to process the next message. It is recommended that you use only one trigger with the serialized option because the order of trigger firing is random with respect to their GUIDs.

**Trigger elements**

Each trigger is associated with a specific *monitored queue* on the local computer and is invoked by the Message Queuing Triggers service for every Message Queuing message arriving at that queue. A trigger is configured to obey one or more *rules*. A rule is defined by one or more *conditions* and an *action* in the form of an executable file or COM component that is executed when all the conditions associated with the rule are evaluated as true.

A *trigger* is an automatic mechanism for implementing a specific action in response to a message event. In the context of Message Queuing Triggers, an event is defined as a set of one or more conditions that may be satisfied by a message arriving in a queue. The collection of all triggers on a particular computer is called the *trigger set*.

A *condition* is a criterion that is applied to a message when it arrives at a monitored queue. A condition can test only for a property of the message. Every condition or set of multiple conditions is associated with an action.

An *action* results when a condition or a set of multiple conditions is true. An action can be expressed as an executable and a parameter list, or as a COM component, a method, and a parameter list. You can also invoke other types of files by making them parameters of an .exe file. For example, you can invoke an .msc file by specifying mmc.exe as the executable file and the full path of the .msc file as a parameter.

A *rule* is a condition-action pair. A rule specifies the action that will be taken when a specific condition or set of conditions is true. The collection of all rules (for all triggers) on a particular computer is called the *rule set*.

# Trigger rules

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Trigger rules**

All rules created on the local computer are displayed in the Computer Management snap-in. You can attach multiple rules to a trigger. If you do so, the rules are executed in the order in which they are displayed in the **Attached rules** list on the second page of the **New Trigger** dialog box or on the **Attached Rules** tab of the ***YourTrigger*** **Properties** dialog box. You can change the priority (order of execution) of rules in this list by selecting a rule and clicking the applicable button. The rules displayed in the **Rules** list are disabled by default. To enable a rule for use with a trigger, you must attach the rule to the applicable trigger. For more information, see View or change properties of triggers and rules for triggers.

A rule is composed of one or more conditions and an action that is executed when the conditions are true. Each condition is applied to every message arriving at the monitored queue, and the result is either true or false. Multiple conditions must all be true for an action to be executed. If no condition is specified, a true value will always be returned. This is useful when you want a trigger to fire unconditionally every time a message arrives at the monitored queue. The conditions applied to a particular rule are listed in the **Condition list** on the **Conditions** page of the **Properties** dialog box for that rule. New rules are composed by selecting a supported initial string in the **New condition** list box and typing a concluding string in the text box below it. These conditions are summarized in the following table.

| Condition | Description |
|---|---|
| Message label contains/does not contain | Tests whether the label of the newly arrived message contains (or does not contain) the string you specified. This string is case-sensitive. |
| Message body contains/does not contain | This applies only to messages with a string body type (VT_BSTR, VT_LPWSTR). Otherwise, the condition will always be false, and no action is invoked. |
| Message priority equals/does not equal/is greater than/is less than | Tests whether the priority of the newly arrived message satisfies the mathematical expression you specified. |
| Application-specific message number (AppSpecific) equals/does not equal/is greater than/is less than | Tests whether an application-specific message number satisfies the mathematical expression you specified. |
| Message source computer ID (SourceMachineGuid) equals/does not equal | Tests whether the source computer GUID equals (or does not equal) the value you specified. |

The actions performed when a condition or set of conditions is true are expressed either in the form of a stand-alone executable file and parameters (arguments), or in the form of a COM component, a method, and parameters. You can specify one or more parameters to be passed per action.

The supported parameters are:

- Message ID (as variant)
- Message label (as string)
- Message body (as variant)
- Message body (as string)
- Message priority (as number)
- Message sending time (as date variant)
- Message arrival time (as date variant)

- Message correlation ID (as variant)
- Message lookup ID (as variant)
- Queue path name (as string)
- Queue format name (as string)
- Response queue format name (as string)
- Administration queue format name (as string)
- Application-specific message number (AppSpecific as long)
- Trigger name (as string)
- Trigger ID (as string)
- Message source computer ID (SourceMachineGuid as string)
- String literal
- Numeric literal

These parameters are presented in the list box under **Parameter** on the **Invocation Parameters** dialog box, which can be accessed by clicking **Parameters** on the third page of the **New Rule** dialog box, as well as on the **Action** page of the **Properties** dialog box for a particular rule. Note that the Message Queuing Triggers service runs under the Network Service account, and that executable files invoked by triggers are invoked with the same security privileges as those assigned to the Triggers service. If your rule invokes a COM component, make sure that it can be successfully invoked from the security context used by the Triggers service, is properly registered, and has suitable dependencies.

By default, the Message Queuing Triggers service runs under the Network Service account. Services that are running under this account interact with the network using the credentials of the computer account, and this account has the same level of access to resources and objects as members of the Users group accounts. For heightened security, you can change the default Triggers service account to a more secure setting; for example, the Local Service account, which limits access to the local computer. You can also specify a less secure account for the Triggers service; for example, the Local System account, allowing access to the entire domain. Note that the Triggers service running under the Network Service account does not interact with the desktop. If this functionality is required, you can configure the Triggers service to run under the Local System account, and enable interaction with the desktop in Computer Management.

If your trigger invokes an executable file with command-line arguments, these arguments can be included as parameters of the string literal type. However, you must add a separate parameter for each command-line argument.

You can also use parameters to invoke other types of files that can be included as a command-line argument of an .exe file. For example, you can invoke an .msc file by specifying mmc.exe as the executable file and the full path of the .msc file as an invocation parameter of the string literal type.

Triggers and rules are created and modified by mqtrig.dll, but the service is implemented by mqtgsvc.exe. Therefore, information about the state of any new or modified trigger or rule must be passed to the service. This is done by sending a notification message to a specific internal private queue. The service reads each message sent to this queue, updates its internal data accordingly, and then deletes the message.

For more information on how to create rules, see Create and attach rules to triggers.

# Multiple Destination Messaging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multiple Destination Messaging**

Message Queuing 3.0 includes a new model for one-to-many messaging. In previous versions of MSMQ, a single message could be sent to a single queue. Now, one-to-many messaging enables the same message to be sent to multiple recipient queues. This can be implemented in a number of ways. For details, see the following links:

- To learn about implementing one-to-many messaging using an Active Directory distribution group object to encapsulate a collection of destination queues that can serve as a distribution list in Message Queuing, see Distribution lists.
- To learn how messages can be sent programmatically to multiple destinations specified in the form of a multiple-element format name, see Multiple-element format names.
- For information on how to send messages to a single destination IP address, to be received and processed by multiple IP hosts, see IP multicasting.

# Distribution lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Distribution lists**

Message Queuing 3.0 allows sending applications to send a single message to a list of destination queues. That list can be constructed dynamically in the form of a multiple-element format name or can be published and kept in Active Directory in the form of a distribution list. For more information on multiple-element format names, see Multiple-element format names.

A distribution list is a set of target queues and/or distribution lists that can be used to do the following:

- Send transactional and nontransactional messages to multiple destinations.
- Specify one or more administration queues for sending application-generated acknowledgment messages, when acknowledgment messages are requested.
- Specify one or more response queues for sending application-generated response messages, when response messages are requested.

**Distribution lists and Active Directory**

Distribution lists are public lists that are published as distribution group objects in Active Directory. For more information, see Message Queuing objects. The explicit members of a distribution list must be destinations published in Active Directory; that is, public queues, queue aliases, and other distribution lists. Private queues and URL-named queues cannot be elements of a distribution list.

Private queues and URL-named queues, however, can be referenced indirectly in a distribution list through queue aliases. Queue aliases provide additional benefits when they are used in distribution lists. When a queue alias object is deleted, the alias is automatically removed from all distribution lists that reference the alias. A queue referenced by a queue alias can be modified without having to change a distribution list that references it. For more information on queue aliases, see Queue aliases.

Since distribution lists are objects in Active Directory, they cannot be used in workgroup or offline mode, in which there is no access to Active Directory. In these modes, messages can be sent to multiple destinations by using multiple-element format names, provided they do not contain an element that points to a public queue or a distribution list.

**Distribution list security**

Because distribution lists are objects in Active Directory, security attributes can be attached to control access to them, and their properties can be queried in Active Directory. In addition, they support encryption similar to queues. Each copy of a message can be encrypted using the public key of the destination computer for that copy.

A security descriptor is assigned to each queue and distribution list. This security descriptor lists the users and groups that are granted or denied access to the queue or distribution list, and the specific permissions assigned to those users and groups. The following specific permissions are associated with distribution lists:

- **Send To permission**. Allows a group or user to send messages to the applicable distribution list. By default, this permission is granted to everyone.
- **Add/Remove Self as Member permission for a public queue or queue alias**. Allows a group or user to add or remove the applicable public queue or queue alias as a member in a distribution list. All users are granted this permission by default.
- **Add/Remove Self as Member permission for a distribution list**. Allows a group or user to add or remove the applicable distribution list as a member in another distribution list. Only the owner of a distribution list has the Add/Remove Self as Member permission for it.

When a queue or distribution list member is added to a distribution list, the set of users with the Send To permission for the distribution list containing the new member is not altered. Thus, in particular, when a distribution list with more restricted access is added to a distribution list with less restricted access, it becomes possible for clients of the expanded distribution list

to access the new distribution list with more restricted access. This does not create a security risk because each distribution list can still control the set of users allowed to add members to it.

When a message is sent to a distribution list, Message Queuing verifies that the sender has the Send To permission to each of the top-level members (queues or distribution lists) in the distribution list. This is a compromise for achieving network traffic optimization. When a message is received, Message Queuing verifies that the sender has the Send To permission to the top-level members and to the destination queue itself.

To achieve security in sending operations, member containment in a distribution list is defined to be bidirectional: a member references its container and the container references its members. To respect the security restrictions on members, a member cannot be added to a member of a distribution list without permission from that member and its prospective container. In other words, both the container and the member must agree. For example, if a user wants to insert distribution list DL2 into DL1, the user must have the Add/Remove-Element permission for DL1 and the Add/Remove-Container permission for DL2. However, for the user to remove DL2 from DL1, only one of these permissions is required.

You can use Active Directory Users and Computers to create and delete distribution lists as distribution group objects, to add and delete members, and to view their properties. For more information on performing these tasks, see Create distribution lists, Delete distribution lists, and View, add, or delete elements in a distribution list.

Distribution lists can be referenced programmatically through distribution list format names. For more information on using distribution list format names and their syntax, refer to the Message Queuing Software Development Kit (SDK).

**Distribution list scope**

Distribution lists can be domain-scoped, global, or universal:

- **Universal**. Members of universal groups can include other groups and accounts from any domain in the domain tree or forest and can be granted permissions in any domain in the domain tree or forest.

- **Global**. Members of global groups can include other groups and accounts only from the domain in which the group is defined and can be granted permissions in any domain in the forest. Thus, distribution lists with global scope can include queues and distribution lists from the local domain, and queues and global distribution lists from other domains in the forest. They can be moved to another domain.

- **Domain-scoped**. Members of domain local groups can include other groups and accounts from Windows Server 2003, Windows 2000, or Windows NT domains and can be granted permissions only within a domain. Thus domain-scoped distribution lists can only contain queues and distribution lists from the local domain as members, and cannot be moved to another domain.

**Sending messages to distribution lists**

When a message is sent to a distribution list, Active Directory is queried to ascertain its members and to translate any queue aliases into queues. Thus, the distribution list is expanded to a list of target queues. Distribution list members can be added or removed dynamically during the expansion process. Then a copy of the message with the same message ID is sent to each of those queues. If a distribution list contains multiple references to the same destination, duplicate references are ignored when the destinations are opened for sending messages. For example, if the members of distribution list D1 include queue Q1 and another distribution list D2, it may turn out that D2 also contains Q1 or D1. If this is the case, the second reference to Q1 or D1 is ignored. Thus, recursive cycles specified by distribution lists containing other distribution lists are broken.

Distribution lists may reference a mixture of transactional and nontransactional target queues. When a nontransactional message is sent to such a mixture of targets, the message is allowed to be sent to all the transactional and nontransactional target queues on the local or a remote computer. However, when the nontransactional message reaches a transactional queue, it is rejected and transferred to the nontransactional dead-letter queue on the computer hosting the transactional queue for storage, and a negative acknowledgement message indicating that a nontransactional message was sent to a transactional queue is sent to the administration queue specified in the original message.

**Sending transactional messages to distribution lists**

Message Queuing guarantees that all messages sent within a transaction will arrive exactly once and in the order in which they were sent. Message Queuing also guarantees that all messages sent in consecutive transactions from a given computer to the same destination queue will also arrive exactly once and in the order in which they were sent. This means that the messages sent within the first transaction will arrive before those sent within a subsequent transaction.

These transactional semantics are also supported for distribution lists. Messages sent to the same distribution list within a transaction will be delivered to the distribution list exactly once and in the order in which they were sent. In addition, messages sent within multiple transactions from the same source computer to the same distribution list will be delivered to the

distribution list exactly once and in the order in which they were sent. Thus, the delivery of transactional messages to distribution lists and to queues is identical.

This equivalence, however, does not extend to the queues referenced by distribution lists. There is no guarantee of in-order delivery to a queue for transactional messages sent to different distribution lists, even if they contain the same elements (queues). Nevertheless, the messages will be delivered exactly once. In the case of delivery failures, Message Queuing places the original message with an indication of the first failed delivery in the source computer's transactional dead-letter queue.

Consider the following examples:

- Example: In-order delivery to the same distribution list:
    - Distribution list DL1 = Queue Q1
    - Send message M1 to DL1
    - Send message M2 to DL1
- In this case, Message Queuing guarantees both in-order and exactly-once delivery.
- Example: Out of order delivery to different distribution lists:
    - Distribution list DL1 = Queue Q1
    - Distribution list DL2 = Queue Q1
    - Send message M1 to DL1
    - Send message M2 to DL2
- In this case, even though both distribution lists contain the same elements, messages to Q1 will not necessarily arrive in-order to Q1. However, exactly-once delivery is guaranteed, so that M1 and M2 will arrive once and only once.

In contrast, using multiple-element format names can provide this in-order guarantee for messages sent to different distribution lists, since each message sent to a multiple-element format name containing a list of target queues. For more information, see Multiple-element format names.

**Dynamic distribution lists and transactions**
Because distribution lists are dynamic, recipients can be added while transactional messages are in transit. Such new recipients will start to receive transactional messages on the next appropriate transaction boundary. Only messages sent after a new queue joins a distribution list will reach the new member.

**Distribution list acknowledgements**
When a message is sent to a distribution list, potentially a set of different acknowledgements can be generated, since some target queues might receive the message successfully, while others may not. A sender should expect at least *n* acknowledgment messages for a message sent to a distribution list with *n* ultimate recipients when acknowledgment messages are requested. More than *n* acknowledgment messages should be expected when a single message generates a positive acknowledgment message upon arrival at its destination and a positive or negative read acknowledgment message. Note that all distribution lists and embedded distribution lists are expanded and duplicate recipients are eliminated before any messages are sent. However, to optimize message traffic, acknowledgment messages of the same class and generated for the same original message are combined into a single acknowledgment message containing a list of the queue GUIDs that acknowledged the message.

# Multiple-element format names

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Multiple-element format names**

Message Queuing 3.0 allows a sending application to send a single message to a list of destination queues. That list can be constructed dynamically in the form of a multiple-element format name or can be published and kept in Active Directory in the form of a distribution list. For more information on distribution lists, see Distribution lists.

A multiple-element format name can be constructed to do the following:

- Send transactional and nontransactional messages to multiple destinations.
- Specify multiple administration queues for sending acknowledgment messages.
- Specify multiple response queues for sending response messages.

A multiple-element format name is formed as a concatenation of one or more public, private, direct, or distribution list format names separated by commas (for more information on format names, see Format names in Queue names). Thus, different kinds of format names used in Message Queuing can be used together as elements of a multiple-element format name. The following example shows a multiple-element format name that contains a direct format name, a public format name, and a distribution list format name.

```
DIRECT=ComputerAddress\QueueName,PUBLIC=QueueGUID,DL=DL_GUID
```

Multiple-element format names can be used to send messages to multiple destinations even in workgroup or offline mode provided they do not contain an element that points to a public queue or a distribution list. For example, if a multiple-element format name references a private queue and a distribution list, no message will be sent to either target when there is no access to Active Directory.

Multiple-element format names can be used to reference multiple administration and response queues. However, they cannot contain format names of read-only queues, such as queue journals, computer journals, or dead-letter queues. An error is returned if a format name of a read-only queue is included in a multiple-element format name.

To enable downlevel compatibility, a list of explicit destinations is derived from every distribution list or multiple-element format name used to send a message in the source computer, and this information is included along with the distribution list or multiple-element format name in each message packet that is sent. For example, when a message is sent to distribution list D1 that references three queues (Q1, Q2, and Q3), the message packet sent to the three computers contains the explicit destinations Q1, Q2, and Q3 as well as the original distribution list D1. If any of the destination computers is running Message Queuing 2.0 or earlier, it will use the explicit destinations Q1, Q2, or Q3. Any computer that is running Message Queuing 3.0 will use D1 and disregard the repeated destinations.

**Sending transactional messages to multiple-element format names**

When a transactional message is sent to a multiple-element format name, appropriate guaranteed positive or negative acknowledgment messages can be returned to the sending application. In addition, Message Queuing guarantees that messages sent in consecutive transactions initiated from the same computer to the same destination queue will arrive exactly once and in the order in which the transactions were committed regardless of the way in which the target queue was addressed: either as an element of a multiple-element format name or as an individual queue.

When a transactional message is sent to a multiple-element format name, Message Queuing logically constructs a single transaction that consists of a series of discrete operations for sending the message to each of the target queues contained in the multiple-element format name and then either commits or aborts the entire transaction. This means that sending a message to a multiple-element format name within a transaction is semantically equivalent to sending a series of messages to individual queues within a single transaction.

Consider, for example, a multiple-element format name N1, which contains two queues Q1 and Q2, and a second multiple-element format name N2, which contains Q2 and Q3. Now consider a transaction consisting of the following three actions:

1. Sending message M1 to N1.

2. Sending message M2 to N2.

3. Sending message M3 to Q2.

Messages M1, M2, and M3 will arrive at queue Q2 exactly once and in the order in which they were sent.

However, the guarantee of exactly-once and in-order delivery applies only to identical top-level elements of multiple-element format names and does not extend to queues that are members of different distribution lists referenced by multiple-element format names. Consider the case of two distribution lists D1 and D2, each of which contains queue Q1 as a member, a multiple-element format name N1, which contains D1 and another queue Q2 as elements, and another multiple-element format name N2, which contains D2 and Q2 as elements. If message M1 is sent to N1 and message M2 is sent to N2 within the same transaction, the messages will arrive at Q2 in the order in which they were sent, but they will not necessarily arrive at Q1 in that order. Nevertheless, the messages will be delivered to both queues exactly once.

# IP multicasting

Updated: March 19, 2010

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**IP multicasting**

Message Queuing provides a mechanism to send and receive IP multicast traffic using the Pragmatic General Multicast (PGM) driver, which is installed and configured as part of the Core subcomponent during every installation of Message Queuing. Messages can be sent to a single destination IP address and then received and processed by multiple IP hosts. A host listens for a specific IP multicast address and receives all packets sent to it. The set of hosts listening on a specific IP multicast address is called the host group.

Using this method, a single copy of the message is sent on the network, and this message can then be received and processed by numerous destination applications. This method is especially useful if the same message has to be delivered to large numbers of receiving applications.

The sender specifies a target IP multicast address. An IP multicast address property is attached to a queue. Subsequently, every message that is sent to that multicast address is inserted into the queue by Message Queuing. Applications continue to read messages from the queue using standard Message Queuing features. The Message Queuing service on a specific host listens to all multicast addresses to which its local queues are bound, routes incoming messages to the appropriate queues, duplicates messages that are sent to multiple queues, and dynamically starts and stops listening to a multicast address as local queues become bound and unbound.

Be aware of the security implications of multicasting. Messages cannot be encrypted, and messages could be read or used by unauthorized persons. Ensure that the specified address is used only for Message Queuing multicasting.

The following are some of the specific features of IP multicast messaging:

- Only one copy of the message is sent, and it is received and processed only by computers, or hosts, that are listening for it.
- Hosts can join and leave the host group dynamically at any time.
- There are no limitations on the size of a host group.
- A host can send messages to an IP multicast address without belonging to the corresponding host group.
- Transactional sending is not supported.

A multicast address and a port number in the *IP_Address*: *PortNumber* syntax can be associated with any public or private queue. For public queues, this information is stored in Active Directory. For a private queue, this information is stored on the computer where messages sent to the queue are stored. Therefore, in workgroup mode, multicast addresses can be associated only with private queues, and in domain mode, a Windows Server 2003 family domain controller must be available to associate multicast addresses with public queues.

IP multicast addresses must be in the class D range from 224.0.0.0 to 239.255.255.255, which corresponds to setting the first four high-order bits equal to 1110. However, only certain ranges of addresses in this range are available for sending multicast messages. For the latest list of reserved multicast addresses, see the link to Internet Assigned Number Authority (IANA) Internet Multicast Addresses at Windows 2000 Resource Kits - Web Resources (http://go.microsoft.com/fwlink/?LinkID=291).

For information about how to specify or view the multicast address and port number for a specific queue, see Specify or view queue multicast addresses and ports.

No additional security mechanisms are provided for multicast messaging. Security is achieved through authorization and authentication based on the mechanisms for one-to-one messaging. Each target queue is protected by authentication based on sender and access based on sender and receiver. Encryption and decryption are not supported.

**Multicast messaging with multiple network adapters**

On a computer with multiple network adapters, a problem will occur when the computer tries to send multicast messages by means of these multiple network adapters. In effect, Message Queuing randomly chooses a network adapter for sending multicast messages. If the private network adapter is chosen, multicast recipients will not receive the message. To work around

this issue, you can specify the source IP address to be used for sending multicast messages. To do this, create the **REG_SZ** value **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\MulticastBindIP** in the usual IP address format. After setting the registry, restart the Message Queuing service for the changes to take effect.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry. |

**Multicast messaging on a cluster node**

For multicast messaging on a cluster node, if no IP address is specified in the **MulticastBindIP** key, the IP address is derived in the following way:

- On a virtual server, the virtual server's IP address that is defined for the IP address resource is used.
- If the cluster node has only one IP address, that address is used.
- If an IP address is specified in the registry, that address is used. If the IP address that is specified cannot be found on the node, an event is issued and the registry entry is ignored.
- If a node has multiple IP addresses, the cluster application programming interface (API) is used to make two lists—a first list of all private (not static) IP addresses, and a second list of all the node's IP addresses. These two lists are then compared to select an address that is in the second list but not in the first list. An informational event is issued with the IP address that is chosen. If the only IP addresses available are the addresses in the first list, one of those addresses is used and an event is issued.

To specify an IP address, create a **REG_SZ** value **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\ClusterBindIP** in the usual IP address format. After editing the registry, restart the Message Queuing service for the changes to take effect.

| ⚠**Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry. |

# Internet Messaging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012

**Internet Messaging**

Message Queuing 3.0 provides Internet-messaging support. This section provides information about the basic concepts and features introduced in Message Queuing 3.0 to provide this support. Message Queuing sends messages over HTTP transport to URL-named queues. Message Queuing 3.0 provides Internet-messaging support with the following features:

- For an overview of Internet messaging, see Internet messaging overview.
- Message Queuing provides messaging over HTTP as an optional transport protocol. For more information, see HTTP/HTTPS messages.
- HTTPS provides access control and authentication when messaging with HTTP. For more information, see HTTPS authentication.
- When messages are sent over HTTP transport, external logical names or aliases can be mapped to internal physical queue locations, protecting internal resources and providing a solution for messaging through firewalls. For more information, see Store and forward servers.
- HTTP messages are formatted using the SOAP Reliable Messaging Protocol (SRMP) extensions of the SOAP protocol. To read more about the Message Queuing implementation of the SOAP protocol, see SRMP messages.
- For information on deploying HTTP messaging in Network Load Balancing clusters, see HTTP/HTTPS messaging in Network Load Balancing clusters.
- For information on HTTP messages in server clusters, see HTTP/HTTPS messaging in server clusters.

# Internet messaging overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet messaging overview**

Message Queuing 3.0 provides support for Internet and intranet messaging, with the following features:

- Reference queues using URL addresses.
- Send and receive HTTP messages to destination queues on an intranet or over the Internet in XML-based, SOAP Reliable Messaging Protocol (SRMP) format.
- Authenticate messages sent over HTTPS.
- Support for messaging through firewalls and perimeter networks.
- Web farm support, including a facility for mapping between external logical queue names and internal queue names, so that internal names do not need to be exposed to message senders.
- Support for deployment of Message Queuing in hardened MSMQ mode, to enhance the security of Message Queuing computers on the Internet.

**HTTP messages**

Message Queuing 3.0 includes an optional HTTP Support component, to send and receive messages in an intranet or over the Internet with the HTTP/HTTPS protocol, in addition to the Message Queuing protocol. Messages can be sent over HTTP/HTTPS to destination queues on an intranet within an enterprise or over the Internet, and the sender and recipient can be in different forests or organizations. HTTP/HTTPS is automatically chosen as the transport for any Message Queuing message sent using an HTTP/HTTPS format name. For more information, see HTTP/HTTPS messages.

**SRMP messages**

SRMP is an extension of the SOAP Message Exchange Protocol (SMXP) that defines a reliable end-to-end messaging protocol for sending HTTP messages. Message properties are transmitted in XML, rather than in binary form. Message Queuing transforms message properties to specific SRMP and SMXP elements, and then on the receiver side, maps these elements to the respective Message Queuing properties before storing the message. For more information, see SRMP messages.

**HTTPS authentication**

Using HTTPS to send messages ensures authentication of users requesting access to Message Queuing resources by means of a Web server, and establishes a Secure Sockets Layer (SSL) connection for secure communication. A direct format name containing HTTPS as the protocol name invokes a secure HTTP transport through an SSL channel. SSL uses certificates to allow clients and Web servers to authenticate each other, and encryption keys to establish a secure communication connection. For more information, see HTTPS authentication.

**Message Queuing in hardened MSMQ mode**

To enhance the security of Message Queuing computers in an Internet environment and when sending and receiving HTTP messages, Message Queuing can be deployed in hardened MSMQ mode. Hardened MSMQ mode imposes restrictions on non-HTTP messaging. For more information, see Message Queuing in hardened MSMQ mode.

**Messaging with perimeter networks**

Most organizations connect to the Internet using a proxy server and a perimeter network protected by firewalls. Incoming messages pass through the firewall, and outgoing messages are sent through a proxy server. The proxy server is a mediator between the organization's intranet and the Internet world. It provides caching and decoupling. Message Queuing uses an HTTP proxy to connect to Web addresses for delivering HTTP messages. Message Queuing thus needs to know the name of the proxy server and the applicable port to which it must connect. At the same time, there may be a list of target names for which the proxy connection is bypassed; that is, Web addresses that do not need to be accessed through the proxy server. An attempt is made to establish a direct connection with such targets. All of the information regarding proxy servers is stored in registry keys. For more information, see Configuring firewalls for Message Queuing and Configuring Message Queuing with a proxy server.

**Message Queuing with ISA Server**
To use Message Queuing with ISA Server, install ISA Server Service Pack 1 or later. For more information, see ISA Server Service Pack 1.(http://www.microsoft.com/isaserver/downloads/sp1.asp)

**Web farm support**
For increased scalability and availability, Message Queuing with HTTP Support installed integrates with common network load-balancing mechanisms. Message Queuing applications can be deployed on the nodes of a Network Load Balancing cluster to receive and process messages sent to the cluster. When messages are sent over HTTP transport to a queue at a virtual URL address, a TCP session is established with one of the nodes in the cluster. For more information, see Message Queuing in Network Load Balancing clusters.

Similar Web farms can also be implemented using other load-balancing mechanisms, such as DNS round-robin distribution, and network address translation (NAT). For information on load-balancing Web farms, see Distributing Message Queuing in a network.

When messages are sent over HTTP transport, the network names of the computers within an organization and the names of the queues hosted on them do not need to be disclosed to message senders. Message Queuing provides a store and forward feature to map external and internal names, protecting internal resources. Mapping is enabled on a store-and-forward Message Queuing server, positioned in the front-end of the enterprise, where all incoming HTTP requests are directed. If messages meet mapping criteria, they are forwarded further into the enterprise. For more information, see Store and forward servers.

# HTTP/HTTPS messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**HTTP/HTTPS messages**

Message Queuing includes support for delivering messages over HTTP, in addition to the Message Queuing protocol. Messages can be sent over HTTP transport to destination queues on an intranet within an enterprise or over the Internet, and the sender and recipient can be in different forests or organizations. For Message Queuing to send messages over HTTP, you must install the Message Queuing HTTP Support subcomponent. This subcomponent installs Internet Information Services (IIS) and creates an IIS extension for Message Queuing. For more information, see Choosing installation subcomponents. This subcomponent is not installed during a default Message Queuing installation.

Message Queuing formats messages using the XML-based SOAP Reliable Messaging Protocol (SRMP), an open general-purpose extension of the SOAP Message Exchange Protocol (SMXP). SRMP defines an open implementation-independent protocol for the reliable one-way delivery of SOAP 1.1 messages. For more information, see SRMP messages. Using SRMP, applications can also be written to create messages that are indistinguishable from Message Queuing-generated HTTP messages on computers that do not have HTTP Support installed, such as computers running MSMQ 1.0 or MSMQ 2.0. These messages can be sent as HTTP packets and accepted by queues hosted on computers with HTTP Support installed. They can then be received as HTTP messages by applications running on those computers.

To send messages to a destination queue over HTTP, the target computer must have Message Queuing's HTTP Support subcomponent and IIS installed. Access to the queue can be restricted to specific users by changing the permissions of the virtual directory in the IIS snap-in.

For secure HTTP messaging, HTTPS provides authentication and a Secure Sockets Layer (SSL) connection between a sender and recipient. For more information, see HTTPS authentication.

**HTTP message formats**

Destination queues for HTTP messages are opened using direct format names that include the URL address of the target computer, the virtual directory name, and the queue name separated by slashes. The default virtual directory name is msmq, but Message Queuing can be configured by IIS to use a different virtual directory. The format names used to send messages over the Internet must include the IP address or complete DNS name of the target computer, while the computer name is sufficient for sending messages within the enterprise. Examples of these direct format names include:

- DIRECT=HTTP://157.18.3.1/msmq/MyQueue (Single computer or cluster)
- DIRECT=HTTP://Mike01/msmq/MyQueue (Within the enterprise)
- DIRECT=HTTP://mail.www.northwindtraders.com/msmq/MyQueue (Over the Internet)
- DIRECT=HTTPS://mail.www.northwindtraders.com/msmq/MyQueue (Over the Internet)
- DIRECT=HTTP://193.2.34.5/msmq/MyQueue (Over the Internet)

The public, private, and direct format names of administration or response queues can be included in messages sent over HTTP. Conversely, the names of administration or response queues in HTTP format can be included in messages sent over an ordinary (non-HTTP) transport. However, the use of distribution lists and multiple-element format names for these purposes is not supported.

For information on direct format names, see "Direct format names" in Queue names.

**Pipelining**

Sending a new request before finishing to read all the responses from previous requests is called pipelining. Because proxies do not always support pipelining and because IIS HTTPS currently does not support it, the pipeline modes of HTTP and HTTPS delivery are controlled by two subkeys, namely HttpPipeLine and HttpsPipeline (of type DWORD), under the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters registry entry. When they are set equal to 1, pipelining is enabled. By default, HTTP delivery works in pipeline mode, while HTTPS delivery does not. After setting the registry, restart the Message Queuing service for the changes to take effect.

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

# HTTPS authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**HTTPS authentication**

Using HTTPS to send messages ensures authentication of users requesting access to Message Queuing resources by means of a Web server, by establishing a Secure Sockets Layer (SSL) connection for secure communication between a sender and recipient. The sender is always regarded as the SSL client, and the recipient is regarded as the SSL server, regardless of whether the computers are running Message Queuing server or client software. Note that authentication for the purpose of establishing an SSL session is not the same as message authentication, which confirms that a message has not been tampered with and can be used to verify the identity of the message sender. For information about message authentication, see Managing message authentication.

Two kinds of certificates are used in HTTPS authentication:

- **Server certificates**. This certificate contains information about the server that allows a client to identify the server before sharing sensitive information.
- **Client certificates**. This certificate contains personal information about the user and identifies the SSL client (the sender) to the server.

**Server certificates**

Before an SSL connection can be established for sending messages, the recipient computer requires a server certificate that resides in the **Personal** store under **Certificates (Local Computer)** in the Certificates snap-in on the recipient computer.

For instructions on obtaining a server certificate from a CA and installing it, see Obtain server certificates for HTTPS messaging.

A server certificate obtained from a CA can be issued to the NetBIOS name or full DNS name of a computer. However, when HTTPS messages are sent, the destination specified in messages must be identical to the computer name in the recipient's server certificate. Thus, when only the NetBIOS name of the destination (for example, with a direct format name of the form DIRECT=HTTPS://msmqserv/msmq/private$/q1) is used to send messages within the same forest over an intranet to a computer whose server certificate contains the full DNS name, authentication fails. Consequently, when Message Queuing tries to create an SSL session, it receives a warning that the name in the certificate does not match the name of the destination, and the messages sent in this way accumulate in the outgoing queue.

**Verifying server certificates on the sender computer**

IIS on the recipient side must send the recipient's server certificate to the sender for authentication. This server certificate, which contains the signature of the certification authority (CA), the recipient's public key, additional information about the recipient, and an expiration date, must be from a CA that the sender trusts. To authenticate the recipient computer, the sender verifies that it trusts the CA and validates the signature in the recipient's server certificate.

When a sender computer trusts a CA, such as VeriSign or Microsoft Certificate Services, it has a certificate from that CA containing the CA's signature and public key in the **Trusted Root Certification Authorities** store under **Certificates (Local Computer)** in the Certificates snap-in. For instructions on importing a CA certificate on the sender computer from Microsoft Certificate Services, see Import a CA certificate from Microsoft Certificate Services.

**Client certificates**

An additional optional security component can be required for SSL sessions if IIS on the recipient side also requests the sender's client certificate for authentication. Client certificates are also obtained from a trusted CA and are stored in the MSMQ\Personal store under Certificates - Service (Message Queuing) on Local Computer in the Certificates snap-in on the sending computer. For information about obtaining and installing a client certificate, see Obtain client certificates.

**Mapping client certificates**

After a Web resource is configured to require an SSL channel, it still needs to be protected from unauthorized access. Since the anonymous user account has default write permissions for the physical MSMQ directory in IIS, and default IIS behavior is to ignore client certificate authentication, any user can send HTTPS messages. After client certificates are enabled, you can further protect content by creating mappings that relate the information contained in a client certificate to a Windows user account. For instructions, see Enable client certificates. Mapping is flexible, with one-to-one mapping to map individual client certificates to accounts, or many-to-one mapping to accept a wide range of certificates that meet specific criteria. For more information about mapping client certificates, see IIS online Help. After this is done, you can restrict or allow access to the MSMQ virtual

directory by configuring security settings for the corresponding physical directory.

When all the conditions imposed by IIS are satisfied, the sender (SSL client) and recipient (SSL server) create and exchange SSL session keys, which are used to encrypt all packets sent over the SSL connection. For more information about certificates and SSL connections, see IIS online Help.

**Access control with HTTPS**

For messages sent with the https transport (DIRECT=HTTPS) there are two unrelated levels of authentication used.

### Hop to Hop Authentication

Hop to hop authentication is performed on the wire between hops through the use of an ISAPI filter (mqis.dll) installed on the destination IIS Server as follows:

- The Message Queuing service on hop N-1 authenticates the identity of IIS on hop N via the server certificate installed on the IIS Server.

- The IIS Server on hop N authenticates the identity of hop N-1 using client certificate mapping. In the absence of client certificate mapping, the IIS Server on hop N sets the identity of the Message Queuing service on hop N-1 to the guest account configured for the IIS Server. The IIS Server on hop N then performs an access check of the specified account against the MSMQ web directory.

Hop to hop authentication is performed before the message ever reaches the MSMQ service on hop N.

| **Note** |
| --- |
| Hop to hop authentication is performed per hop, if store and forward director (SFD) redirection is used then each hop may authenticate differently. |

### End to End Authentication

Once hop-to-hop authentication has been successfully negotiated, the Message Queuing service on the destination computer authenticates the user that sent the message as follows:

- The destination Message Queuing service checks the message for a signature.

- If the message is signed then the destination Message Queuing service authenticates the user as the account specified in the message signature.

  | **Note** |
  | --- |
  | Specify whether a message should be authenticated and what digital signature is used by setting the appropriate value for the PROPID_M_AUTH_LEVEL property. |

- If the message is not signed then the destination computer authenticates the message using anonymous access credentials.

For more information about how users sending authenticated messages over HTTP/HTTPS transport can be validated, see Authenticating HTTP/HTTPS messages.

# Store and forward servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Store and forward servers**
When messages are sent over HTTP transport, the network names of the computers within an organization and the names of the queues hosted on them do not need to be disclosed to message senders. Message Queuing 3.0 can map external and internal names, protecting internal resources, and providing a solution for messaging through firewalls. Mapping is enabled on a store and forward server, positioned in the front-end of the enterprise, where all incoming HTTP requests are directed. If messages meet mapping criteria, they are forwarded further into the enterprise. For instance, only messages arriving by means of an HTTPS session, or messages whose sender can be authenticated, or messages selected by another security criterion will be forwarded. This ensures that only authorized messages are ever presented to corporate servers.

**Logical and physical addresses**
Using this mapping facility, HTTP messages can be redirected from a logical address, accessible from the Internet, to a corresponding internal physical address. These are defined as follows:

- **Logical address**. Contained in URL format in the **<to>** subelement of the **<path>** element of the SRMP Message.
- **Physical address can be any of the following**:
  - A queue on the target computer specified in the logical address, to which the message must be written by the Message Queuing service.
  - A queue on a computer in the internal network. To deliver the message, an HTTP/HTTPS request call is made by the Message Queuing service handling the message.

Mappings are created manually by including **<Redirections>** entries in XML *mapping files*. Each **<Redirections>** element contains 0 or more **<redirection>** elements, each of which contains exactly one **<from>** subelement and exactly one **<to>** subelement. Each **<redirection>** element describes a map from a logical address (**<from>**), to a physical address (**<to>**). These XML mapping files must be stored in the \Windows\System32\Msmq\Mapping folder. A sample mapping template is provided in this folder.

The following example illustrates an XML mapping file:

```
<Redirections>
    <redirection>
      <from>http://www.fabrikam.com/accounts</from>
      <to>http://Accounts01/msmq/private$/DestQ</to>
    </redirection>
  </Redirections>
```

When an SRMP message reaches the computer specified as the logical address, the Message Queuing service uses a *local resolution algorithm* to map the logical address to the physical address.

**Local resolution algorithm**
A *local resolution algorithm* defines how the local Message Queuing service on the target computer specified in the logical address resolves the logical address into its physical address.

**Transparent mapping**
By default, if a message arriving from a remote computer address cannot be matched with a physical address in the mapping file, the message is discarded. However, if transparent mapping is enabled, a physical address for the message is set to the message's logical address. To enable transparent mapping, create the **MQS_TransparentSFD** registry value as indicated in the **Outbound message redirection** section of this topic.

The local resolution algorithm works as follows:

- If the logical address specifies a public or private queue on the local computer, and the queue exists, the physical address is set to the logical address of the queue.

- Otherwise, the Message Queuing service searches through all the **<redirection>** elements in its local mapping files to try to match the logical address with a **<from>** subelement in one of the local mapping files, but the corresponding **<to>** subelement does not indicate a local public or private queue, it sets the physical address to match this **<to>** subelement, and forwards the message to the remote computer specified in this **<to>** subelement. This process is called *redirection*.

- If transparent mapping is enabled, and the logical address of a message arriving from a remote computer cannot be resolved by any of the above methods, the physical address of the message is set to its logical address and the message is stored on the local computer.

- If transparent mapping is not enabled, and the logical address of a message arriving from a remote computer cannot be resolved by any of the above methods, the message is discarded.

- If the logical address of a message arriving from a local application cannot be resolved by any of the above methods, the physical address will be set to the logical address.

Hence if in the example above http://Accounts01/msmq/private$/DestQ does not point to a local private queue, then any SRMP message that is sent to the logical address specified will be forwarded to the computer Accounts01, where it will be written into the local private queue DestQ if such a local queue exists (subject to the normal authorization process) or it will be discarded silently if such a queue does not exist. If this **<redirection>** element is present in a mapping file on some other computer, the physical address of the message will be set to http://Accounts01/msmq/private$/DestQ and the message will be forwarded to this URL.

This provides a means to put HTTP messages sent to a logical address into an interim queue on the store and forward server in the receiving side of a perimeter network. From here they can be directed to the corresponding queue on a computer in the internal network.

**Using wildcards**

Limited use of regular expressions in **<from>** elements are supported. Asterisk-terminated URLs can define a redirection from multiple logical addresses to a single physical address. Only this type regular expression can be used, and only in the **<from>** element:

```
<Redirections>
  <redirection>
    <from>https://www.fabrikam.*</from>
    <to>http://Accounts01/msmq/private$/DestQ</to>
  </redirection>
</Redirections>
```

In this example, any message with a logical address that starts with the string https://www.fabrikam. will be redirected to the physical address in the <to> element. Thus, for example, messages with logical addresses https://www.fabrikam.com/accounts and https://www.fabrikam.net/data will be both redirected to the physical address http://Accounts01/msmq/private$/DestQ.

This functionality is based on software provided by J. Kanza. For copyright information for this software, see Message Queuing: External software copyright and permission notice.

**Order acknowledgment messages for transactional HTTP messages**

Transactional messages are to be delivered exactly once and in the order in which they are sent. When transactional messages (streamed HTTP messages) are sent over HTTP from the internal network of an enterprise to the Internet, the physical address of the sender's order queue must be translated into a logical address using a mapping, to direct the order acknowledgment messages (stream receipts) from the receiving computer (outside the sender's enterprise), back to the sender's order queue.

Mappings between logical addresses for sending streamed HTTP messages from local applications and logical addresses for returning order acknowledgment messages (stream receipts) from the recipient computers are created manually by including **<StreamReceiptSetup>** entries in an XML mapping file that is stored on the local computer. Each **<StreamReceiptSetup>** entry contains 0 or more **<setup>** elements, each of which contains exactly one **<LogicalAddress>** subelement and exactly one **<StreamReceiptURL>** subelement. In each **<setup>** entry, the **<StreamReceiptURL>** element specifies the URL that will be included in streamed (transactional) HTTP messages sent from a local application to a logical address matching the URL specified in the **<LogicalAddress>** element.

When a streamed HTTP message is sent from a local application, the local Message Queuing service searches the **<setup>** entries in the mapping file. If a URL matching the logical address of the message is found in the **<LogicalAddress>** element of one of the **<setup>** entries, the URL specified in the accompanying **<StreamReceiptURL>** element is inserted into the

**<sendReceiptsTo>** element of the **<Stream>** entry in the SOAP envelope of the message.

Each **<setup>** element also defines an implicit mapping that redirects every HTTP message arriving at the local computer with the destination specified in **<StreamReceiptURL>** to the physical address of the local order queue. Such messages are normally stream receipts (order acknowledgment messages).

A **<StreamReceiptSetup>** entry may contain one **<default>** element, which defines the default stream receipt URL that the local Message Queuing service inserts into every outgoing streamed (transactional) HTTP message, unless the logical address of the message matches the **<LogicalAddress>** element in one of the **<setup>** elements.

The following is an example of an XML mapping file that specifies the correspondence between a logical address for sending streamed HTTP messages from local applications and the logical address for returning stream receipts from the recipient computer along with the default URL for returning stream receipts:

```
<StreamReceiptSetup xmlns="msmq-streamreceipt-mapping.xml">
  <setup>
    <LogicalAddress>
      https://http://www.fabrikam.com/accounts
    </LogicalAddress>
    <StreamReceiptURL>
      https://www.fabrikam.com/StreamReceipts
    </StreamReceiptURL>
  </setup>
  <default>
    http://Accounts01/msmq/PRIVATE$/order_queue$
  </default>
</StreamReceiptSetup>
```

The limited use of regular expressions in the context of **<LogicalAddress>** elements is supported. Asterisk-terminated URLs can define the correspondence of multiple logical addresses to a single stream receipt URL. This is the only type of regular expression supported in **<StreamReceiptSetup>** entries, and it is allowed only in the **<LogicalAddress>** element.

In following example, any message with a logical address that starts with the string http://www.fabrikam.com/* will use the stream receipt URL specified in the **<StreamReceiptURL>** element.

```
<StreamReceiptSetup xmlns="msmq-streamreceipt-mapping.xml">
  <setup>
    <LogicalAddress>
      http://www.fabrikam.com/*
    </LogicalAddress>
    <StreamReceiptURL>
      http://www.fabrikam.com/StreamReceipts
    </StreamReceiptURL>
  </setup>
</StreamReceiptSetup>
```

**Outbound message redirection**

The number of HTTP/HTTPS sessions established and the number of client certificates needed can be minimized by routing all outbound messages (messages sent out of the enterprise) through a dedicated intermediate host. This dedicated intermediate host is called a transparent store and forward server (TSFD).

A computer with HTTP Support installed is configured as an intermediate host TSFD by creating the following REG_DWORD value in the registry and setting it to a value of 1:

```
HKLM\SOFTWARE\Microsoft\MSMQ\Parameters\MachineCache\MQS_TransparentSFD
```

After setting the registry, restart the Message Queuing service for the changes to take effect. After setting the registry, restart the Message Queuing service for the changes to take effect.

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

Any other value or a lack of this registry key indicates that the computer will not function as a TSFD and that all messages sent to remote queues will be discarded if the computer does not have a regular XML mapping file for them.

For a source computer to be able to redirect its messages through an intermediate host, a special XML configuration file needs

to be added. This configuration file defines which messages will be redirected through an intermediate host and which will not. This XML configuration file must be stored in the \Windows\System32\Msmq\Mapping folder and is loaded when the Message Queuing service starts.

The following is an example of an XML configuration file that specifies the correspondence between the address of a destination queue and an intermediate host:

```
<outbound_redirections xmlns="msmq_outbound_mapping.xml">
  <redirection>
    <destination>http://target_host/msmq/private$/DestQ</destination>
    <through>http://intermediate_host/msmq</through>
  </redirection>
</outbound_redirections>
```

The **<outbound_redirections>** element contains 0 or more **<redirection>** subelements, each of which in turn contains exactly one **<destination>** subelement and exactly one **<through>** subelement. Each **<redirection>** element describes a redirection (or mapping) of the physical address of the destination queue given in the **<destination>** subelement to the virtual directory on the intermediate host given in the **<through>** subelement.

MSMQ 3.0 supports a limited use of regular expressions in a **<destination>** element. Users can use an asterisk-terminated URL to define a redirection of multiple physical addresses of destination queues to a single address of a virtual directory on an intermediate host.

In the following example, all messages sent using the HTTP protocol will be redirected through the intermediate host:

```
<redirection>
  <destination>http://*</destination>
  <through>http://intermediate_host/msmq</through>
</redirection>
```

Each **<outbound_redirections>** element may contain 0 or more **<exception>** subelements. You can use **<exception>** elements to exclude messages sent to specific queues from being redirected by outbound **<redirection>** rules.

MSMQ 3.0 supports a limited use of regular expressions in an **<exception>** element with the same restrictions as for the previously described **<destination>** element.

In the following example, all messages sent to private queues on special_host will not be redirected through the intermediate host:

```
<exception>http://special_host/msmq/private$/*</exception>
```

You can use the keyword **local_names** in an **<exception>** element to exclude messages sent to any internal (intranet) host from being redirected by outbound **<redirection>** rules.

In the following example, no messages sent to internal hosts will be redirected through the intermediate host:

```
<exception>local_names</exception>
</outbound_redirections>
```

# SRMP messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**SRMP messages**

Internet messaging is increasingly using the medium of XML, particularly the XML messaging framework of SOAP 1.1. The Web Services Routing Protocol (WS-Routing) supports both one-way and request/response message exchange paradigms. However, it is explicitly not a goal of WS-Routing to provide reliability. The SOAP Reliable Messaging Protocol (SRMP) is an open general-purpose extension of WS-Routing that adds reliability to WS-Routing and SOAP.

Message Queuing uses both WS-Routing and SRMP elements as well as elements defined in the MSMQ namespace to specify its message properties in HTTP messages. Some Message Queuing message properties have equivalents defined by WS-Routing, and others have equivalents defined by SRMP.

When sending an HTTP message, Message Queuing automatically transforms the individual binary message properties to specific WS-Routing elements, SRMP elements, and elements of the MSMQ namespace, and includes them in the header portion of the *SOAP envelope*. On the receiver side, Message Queuing maps these XML elements to the respective binary message properties before storing the message. The end-to-end semantics for transmitting message properties over a non-HTTP transport are preserved for messages sent over an HTTP transport.

The header portion of the SOAP envelope in an HTTP message contains two required entries and up to seven optional entries in the following order:

- **<path>**

  (Required.) Contains WS-Routing subelements.

- **<properties>**

  (Required.) Contains SRMP subelements.

- **<services>**

  (Optional.) Contains SRMP subelements.

- **<Stream>**

  (Optional.) Contains SRMP subelements. This entry is generated only in transactional messages.

- **<streamReceipt>**

  (Optional.) Contains SRMP subelements. A **<streamReceipt>** entry is generated only in order acknowledgment messages (stream receipts) sent for transactional messages (streamed HTTP messages).

- **<deliveryReceipt>**

  (Optional.) Contains SRMP subelements. A **<deliveryReceipt>** entry is generated only in acknowledgment messages sent in response to a request to indicate that the original message reached the destination queue.

- **<commitmentReceipt>**

  (Optional.) Contains SRMP subelements. A **<commitmentReceipt>** entry is generated only in acknowledgment messages sent in response to a request to indicate whether or not the original message is received from the destination queue before its time-to-be-received timer expires.

- **<Msmq>**

(Optional.) Contains subelements defined in the MSMQ namespace.

- **<Signature>**

  (Optional.) Contains subelements that specify an XML digital signature.

For more information about SRMP elements and the equivalent Message Queuing properties, see the Message Queuing Software Development Kit (SDK). For details of how to access the SDK, see Additional Message Queuing resources.

Applications can be written to create messages that are indistinguishable from Message Queuing-generated HTTP messages on computers that do not have HTTP Support installed, such as computers running MSMQ 1.0 or MSMQ 2.0. These messages can be sent as HTTP packets and accepted by queues hosted on computers with HTTP Support installed. They can then be received as HTTP messages by applications running on those computers.

# HTTP/HTTPS messaging in Network Load Balancing clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**HTTP/HTTPS messaging in Network Load Balancing clusters**

Message Queuing can be implemented in Network Load Balancing clusters to receive and process messages sent to the cluster. To implement Message Queuing in a Network Load Balancing cluster, Message Queuing is installed with the same set of subcomponents on all nodes in a cluster. So that if you install Message Queuing with HTTP support on one node, install Message Queuing with the same functionality on all other nodes.

All computers in the cluster are addressed by the same virtual cluster IP address or set of virtual cluster IP addresses, and when messages are sent over HTTP transport to a queue at a cluster IP address, a TCP session is established with one of the hosts in the cluster. Messages can be sent only to destinations specified by direct format names that include a URL or IP address of the cluster.

In a Web farm implementation, such as a Network Load Balancing cluster, successive HTTP connections from the same source might be directed to two different servers. This means that end-to-end reliable messaging is difficult, and Message Queuing does not guarantee exactly-once delivery through a Web farm, although a best-effort attempt is made by means of a lightweight nonpersistent duplicate removal mechanism. This issue can be addressed by the use of Message Queuing's store and forward feature. Store and forward mapping uses XML files located on Message Queuing computers, with Message Queuing usually installed on each host in the Network Load Balancing cluster. These XML files maintain a mapping between the logical external URL or IP address to which the message is sent, and the physical internal address of a queue. For every incoming message arriving at a Network Load Balancing host running Message Queuing with store and forward functionality, the Message Queuing service will read the logical address to which it was sent. If a message meets mapping criteria, it is forwarded further into the enterprise. For more information, see Store and forward servers.

Similar Web farms and store and forward functionality can also be implemented using DNS round-robin distribution, and network address translation (NAT). For an overview of these, see Distributing Message Queuing in a network.

# HTTP/HTTPS messaging in server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**HTTP/HTTPS messaging in server clusters**

This section applies only to users running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition in a server cluster configuration. To implement Message Queuing in a server cluster, Message Queuing is installed with the same set of subcomponents on all nodes in a server cluster. This means that if you install Message Queuing with HTTP Support on one node, install Message Queuing with the same functionality on all other nodes.

When you install Message Queuing with the HTTP component on a cluster node, IIS will also be installed on that node, if it is not already installed. You then define cluster groups for Message Queuing, and create a virtual server - a cluster group that will contain the Message Queuing resource, plus other resources upon which the Message Queuing resource depends. For further information, see Message Queuing in server clusters. Each virtual server is associated with a network name and an IP address. To support these virtual servers, IIS must assign a Web site for all virtual server IP addresses. To do this, configure the default Web site setting to *All Unassigned*, which assigns to the default Web site all IP addresses that are not assigned to a specific site. Alternatively, you can specify the IP address of a particular virtual server as the default Web site. For instructions on setting these options, see Configure Message Queuing with HTTP for server clusters.

**Failover for the Message Queuing resource with HTTP**

If the IIS service on the cluster node goes down, this does not cause the Message Queuing cluster group to fail over. This means that in such a case the Message Queuing resource will continue to work on the node, but HTTP messaging cannot be used. For more information on forcing failover of cluster groups when IIS goes down, see the Windows Server 2003 family Help file, or the IIS Help file.

# Message Queuing in Clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing in Clusters**
This section provides information about messaging over clusters, and how to administer message queuing in clusters. For details, see the following links:

- For details on deploying Message Queuing in a server cluster, see Message Queuing in server clusters.
- For information on implementing Message Queuing on a Network Load Balancing cluster, see Message Queuing in Network Load Balancing clusters.

# Message Queuing in server clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing in server clusters**
The Message Queuing service can be implemented in a server cluster, which appears to network clients as a single, highly available system. Such clusters are used to increase reliability and fault tolerance through failover as well as to achieve load balancing. Message Queuing server clusters can be created only on computers running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. For these clusters, the Message Queuing service is installed on every node in the cluster. These instances of the Message Queuing service are needed to support Message Queuing resources in the virtual servers. In addition, although they are unaware of clustering, they are needed, even in a clustered system, to serve the many applications and system components that run outside the context of any virtual server and call Message Queuing APIs.

**Cluster groups and virtual servers**
A cluster group is a collection of cluster resources with the following characteristics:

- Groups define the units of failover. That is, when one resource in a group fails and it is necessary to move the resource to an alternate node, all of the resources in the group are moved to the alternate node.

- A group is always owned by one node at any point in time. Likewise, a resource is always owned by a single group. These relationships ensure that all of a group's members reside on the same node.

To access a network application or resource in a nonclustered environment, network clients must connect to a physical server (that is, a specific computer on the network identified by a unique network name and Internet Protocol (IP) address). If that server fails, access to the application or resource is impossible. Through server clusters, Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition enable the creation of virtual servers. Unlike a physical server, a virtual server is not associated with a specific computer and can be failed over like a group. If the node hosting the virtual server fails, clients can still access its resources using the same server name. A virtual server is a group that contains a network name resource, IP address resources, other virtual servers, and other resources - including applications like Message Queuing, to be accessed by the clients of the virtual server.

When you form cluster groups for Message Queuing, you must create a virtual server. The Message Queuing resource is dependent on a Physical Disk resource, to store message and queue data, and a Network Name resource, so that remote clients can access it. This virtual server is created on each node of the cluster using standard cluster tools (Cluster Administrator is the standard GUI tool) or cluster-creating APIs. If you are using triggers, HTTP messaging, or external transactions on your cluster, requiring a Distributed Transaction Coordinator (DTC) resource, a single DTC resource is created only in one of the virtual servers. Then you create a Message Queuing resource using Cluster Administrator in some or all of the existing cluster groups. The first time that a Message Queuing resource on any virtual server is brought online, the cluster Network Name resource creates a pseudo-computer object, and the Message Queuing resource creates the Message Queuing objects under it in Active Directory. Each virtual server functions similarly to a physical computer, so a Message Queuing server running in the context of a virtual server provides services similar to a Message Queuing server running on a physical computer. In particular, queues can be created on a virtual server, and messages can be sent to them. Such queues are addressed using the *VirtualServerName\QueueName* syntax.

For more information on Cluster Administrator, see Using Cluster Administrator, in the Cluster Administrator Help file.

For more information on creating virtual servers for Message Queuing using Cluster Administrator and on creating a DTC resource in a server cluster, see Installing Message Queuing in a server cluster.

Note that when you select the Physical Disk resource for the Message Queuing cluster group, Message Queuing allocates its storage in the *MSMQ\STORAGE* folder on the shared disk. After storage has been allocated, you cannot modify the folder location.

Message Queuing clients can communicate with a standard Message Queuing server running on a node of a server cluster, or if Message Queuing applications are cluster-aware, they can run on a Message Queuing server running in the context of a virtual server.

In the active/active model, which is supported by all versions of Message Queuing following MSMQ 1.0, there can be multiple Message Queuing virtual servers running (active) on a single node. Thus, in a simplified example, if there are two nodes in a server cluster and Message Queuing is installed on both of them and if these two nodes host a total of three virtual servers,

any of these virtual servers can fail over to the other node.

In the event of a failover, the resources of a group on one node, initially the preferred node, are taken offline and then brought back online on another node. The data stored on the physical disk remains unchanged because only the ownership of the Physical Disk resource is transferred to the new node. Upon failback, the resources are moved back logically to the preferred node.

Because the Message Queuing Triggers service is cluster-aware and supports the active/active paradigm, when the Message Queuing service fails over to another node, the Triggers service will fail over along with it. For this to occur, you must create a Message Queuing Triggers resource with resource dependencies on the Message Queuing resource and the Network Name resource on the applicable node using Cluster Administrator. When failover occurs, the trigger and rule definitions, which are stored in the Windows registry, can then propagate between cluster nodes along with the other Message Queuing keys. After failover, the Triggers service can thus continue to process the incoming messages in each monitored queue and invoke the applicable stand-alone executable or COM component according to the rules defined.

Virtual servers are managed using standard snap-ins. The Users and Computers snap-in can be used to manage the Message Queuing resources within a virtual cluster just as it is used to manage Message Queuing computers in a network. You can also open the Computer Management snap-in from within a virtual server for local management of the Message Queuing resource.

In a server cluster, the Message Queuing servers installed on the nodes must all provide the same set of services to enable failover. The types of services provided by Message Queuing running in the context of a virtual server depend on the configuration of the Message Queuing servers installed on the host nodes. For example, if Message Queuing servers with routing services enabled are installed on the nodes of a server cluster, these services will also be available in the context of the virtual servers hosted on those nodes.

Note that the **Downlevel Client Support** component cannot run on a virtual server. However, computers requiring this component can query a remote Message Queuing server with this component, running on a domain controller.

**IP addressing with multiple network interface cards on a cluster node**
On the physical node of a cluster, there are usually two network interface cards with different IP addresses. One of these network interface cards is used only for internal cluster communication. To ensure that Message Queuing does not use the private IP address of an internal cluster Network Interface Card, which would result in messaging failure, Message Queuing maintains a list of all private IP addresses on a cluster node. When Message Queuing starts, it checks this list to make sure that the IP address of an internal network interface card is not selected for messaging. Alternatively, you can specify an IP address on the cluster node that must be used for messaging, instead of allowing Message Queuing to randomly select an IP address from those available.

The IP address used for messaging is derived by applying the following conditions:

- On a virtual server, the virtual server's IP address defined for the IP address resource is used.
- Else if the cluster node has only one IP address, that is used.
- Else if an IP address is specified in the registry, that is used. If the IP address specified cannot be found on the node, an event will be issued and the registry entry ignored.
- Else if a node has multiple IP addresses, the cluster API is used to make two lists - a first list of all private (not static) IP addresses, and a second list of all the node's IP addresses. These two lists are then compared, to select an address that is in the second list but not the first. An informational event is issued with the IP address chosen. If the only IP addresses available are those in the first list, one of those is used, and an event issued.

To specify an IP address, create a string value **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\ClusterBindIP**, in the usual IP address format. After editing the registry, restart the Message Queuing service for the changes to take effect.

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

Note that after restarting a node computer, the Message Queuing service running on the node does not restart automatically, and you need to restart it manually. However, Message Queuing resources that were online do come back online

automatically.

**IP addressing for multicast messaging on a cluster node**
On a cluster node with multiple network interface cards, a problem will arise when the computer tries to send multicast messages by means of these multiple network interface cards. In effect, Message Queuing will randomly choose a network interface card for sending multicast messages. If the private network interface card is chosen, multicast recipients will not receive the message. To work around this issue, you can specify the source IP address to be used for sending multicast messages. To do this, create the string value **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\MulticastBindIP**, in the usual IP address format. After setting the registry, restart the Message Queuing service for the changes to take effect.

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

Note that for multicast messaging on a cluster node, if no IP address is specified in the **MulticastBindIP**, then the IP address to be used will be derived from the procedure detailed previously in the section entitled **IP addressing with multiple network interface cards on a cluster node**.

**Memory management for Message Queuing resources in clusters**
If you have multiple Message Queuing resources in a server cluster, message activity may not function as expected unless you increase the system view space memory pool size of each node by 4 MB for each instance of the Message Queuing resource. This is recommended for cluster nodes with one or two Message Queuing resources and necessary for cluster nodes with three or more Message Queuing resources: To do this:

- Open Registry Editor using **regedit**.
- Open the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management**.
- Create a new string value called **SystemViewSize**.
- Calculate and specify this value as follows:
    - Use this formula: (16 + (the number of Message Queuing resources x 4)).
    - For example, the value for a cluster with three Message Queuing resources equals 28.
- Reboot each node.

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

For more information on server clusters for Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition, see Server Clusters, in the Windows Help file.

# Message Queuing in Network Load Balancing clusters

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing in Network Load Balancing clusters**

Message Queuing services can be implemented in an Network Load Balancing cluster. Network Load Balancing clusters distribute client connections over multiple servers, providing scalability and high availability for client requests for services and applications, but they do not have failover capabilities. In its most common implementation, Network Load Balancing distributes incoming TCP/IP traffic between multiple Internet Information Services (IIS) servers, which comprise a Web farm. Network Load Balancing can combine up to 32 servers running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition into a single load-balancing cluster. Servers can be added without shutting down the cluster.

The heart of Network Load Balancing is the WLBS driver, which is loaded into each member server, or host, in the cluster. This driver includes the statistical mapping algorithm that the cluster hosts collectively use to determine which host handles each incoming request.

Under this type of clustering, if a host fails or goes offline, Network Load Balancing automatically reconfigures the cluster to redirect subsequent client requests to hosts that are still viable members of the cluster. Any client connections that were open when the host failed or went offline are ended. On retry, the client connection is routed to a viable host. The downtime for client connections is less than 10 seconds.

Network Load Balancing allows all computers in the cluster to be addressed by the same set of cluster Internet Protocol (IP) addresses, while also maintaining their existing, dedicated IPs. This feature transparently distributes client requests among the hosts in the cluster, using virtual Internet Protocol (IP) addresses. You must run IIS 5.0 or another TCP/IP service on each host, and the hosts must serve the same content so that any of them can handle any request. You can copy updated pages to local disks on the hosts, or you can use commercial file-replication software to perform updates. Network Load Balancing allows you the option of specifying that all connections from the same client IP address be handled by a particular server (unless, of course, that server fails). It also permits you to allocate all requests from a Class C address range to a single server.

When you set up a Web farm using Network Load Balancing, each server is in constant communication with all the others. They exchange performance statistics and divide the responsibilities of handling incoming requests. Every incoming request is seen by every server in the farm, and Network Load Balancing has its own algorithm to determine which server will handle each request.

When messages are sent over HTTP transport to a queue at a virtual URL address, a TCP session is established with one of the nodes in the cluster. Messages can be sent only to destinations specified by direct format names. In the case of express and recoverable messages, the receiver is not required to maintain any internal state, and Message Queuing does not guarantee exactly-once delivery through a Web farm, although a best-effort attempt is made by means of a lightweight nonpersistent duplicate removal mechanism. When transactional messages are sent over HTTP transport, in-order delivery is guaranteed by including an SRMP stream header entry that includes information for message sequencing. For more information, see Store and forward servers.

When a session between a client and a node is disconnected and the client is routed to a different node for the next session, the messages whose delivery was not acknowledged are resent by the client and accepted by the new node with the possible creation of duplicates.

Message Queuing applications can be deployed on the nodes of an Network Load Balancing cluster to receive and process messages sent to the cluster.

Similar Web farms can also be implemented using DNS round-robin distribution, and network address translation (NAT).

# Cross-Platform Messaging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Cross-Platform Messaging**

This section provides information about messaging between systems running Message Queuing and those running a different messaging software. These disparate systems communicate by means of a connector application. For details, see the following links:

- Foreign computers and sites
- MSMQ-MQSeries Bridge

# Foreign computers and sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Foreign computers and sites**

Message Queuing clients can communicate with foreign computers using a connector application. Foreign computers run operating systems that do not support Message Queuing or use other messaging systems. Connector applications run on Message Queuing servers and enable foreign computers to communicate with Message Queuing clients. A Message Queuing server running a connector application typically has two additional queues, called connector queues: one used for transactional messages and one used for nontransactional messages. For a definition of transactional messages, see Transactional messaging. When a Message Queuing client sends a message to a foreign computer, the message is first routed to a connector queue on the Message Queuing server and is then forwarded to the foreign computer.

Examples of connector applications are Microsoft MSMQ-MQSeries Bridge, which allows connectivity with computer platforms running IBM MQSeries.

Microsoft MSMQ-MQSeries Bridge extends the functionality of Message Queuing to provide seamless, transactional connectivity with IBM MQSeries Version 2.x and Version 5.x platforms. Microsoft MSMQ-MQSeries Bridge currently ships as part of Microsoft Host Integration Server 2000. For more information on Microsoft Host Integration Server 2000, refer to the Microsoft Host Integration Server Web site.(http://www.microsoft.com/hiserver/)

**Foreign sites**

To communicate with foreign computers, you must first create a foreign site, add the foreign computers to the foreign site, and create queues on the foreign computers. All of these are Message Queuing objects in Active Directory whose existence in the physical world is not essential. Then, you must create a routing link between the foreign site and a Windows Server 2003 family or Windows 2000 Message Queuing (non-foreign) site that contains a Message Queuing server running a connector application, define this Message Queuing server as a site gate for the routing link, and add the foreign site to the list of sites where this computer resides.

Foreign computers can be assigned to the same foreign site if the computers are all serviced by the same connector application. Connector applications run on Message Queuing servers and enable foreign computers to communicate with other Message Queuing computers in your network.. When a Message Queuing computer sends a message to a foreign computer, the message is routed to the Message Queuing server running the connector application in the foreign site because this server is also a site gate in the routing link connecting to the foreign site. When the server identifies the message destination as a queue on a foreign computer, it places it in a connector queue. The message is subsequently processed by the connector application, which communicates with the foreign computer.

You must create a routing link between the foreign site and each Windows Server 2003 family or Windows 2000 site that will communicate with the foreign site. The Message Queuing routing server that runs the connector application in the foreign site must also function as a site gate in each routing link connecting to the foreign site. Site gates are always associated with a particular routing link. Note that if the routing server running the connector application, and functioning as a site gate, is unavailable, site gates in other sites cannot route messages to the applicable connector queue for transmission to computers in the foreign site.

For information on how to create a foreign site, see Create foreign sites.

For information on how to add foreign computers to a foreign site and how to add a Message Queuing server running a connector application to a foreign site, see Add foreign computers to a site.

For information on how to create and configure a routing link, see Routing links.

# MSMQ-MQSeries Bridge

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**MSMQ-MQSeries Bridge**

Microsoft MSMQ-MQSeries Bridge is a connector application used to interface computers running Message Queuing applications and computers running IBM MQSeries applications. To send a message from a Message Queuing computer to a MQSeries computer, a destination queue on the foreign (MQSeries) computer must already exist. For more information on configuring Microsoft MSMQ-MQSeries Bridge and foreign computers running IBM MQSeries, see the documentation provided with Microsoft MSMQ-MQSeries Bridge.

Microsoft MSMQ-MQSeries Bridge currently ships as part of the Microsoft Host Integration Server 2000. For more information, refer to the Microsoft Host Integration Server Web site (http://www.microsoft.com/hiserver/).

# Administering Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Message Queuing**

This section provides conceptual information about the administrative tasks for Message Queuing:

- For an overview of the tools provided to administer your Message Queuing network, see Administration overview.
- For information about planning and deployment, installing Message Queuing, and upgrading computers running MSMQ 1.0 and MSMQ 2.0, see Planning, Installing, and Upgrading Message Queuing.
- For information on administering sites, configuring routing links, messaging through firewalls, and configuring DNS settings, see Administering a Message Queuing Network.
- For information on Message Queuing objects in Active Directory, see Administering Message Queuing in Active Directory.
- For information on assigning Message Queuing servers for routing, administering server cluster nodes, and joining a workgroup computer to a domain, see Administering Message Queuing Computers.
- To learn more about security tasks, see Administering Message Queuing Security.
- To learn about administering queues, see Administering Queues.
- For more on administering and managing distribution lists, see Administering Distribution Lists.
- For background information about message tasks, see Administering Messages.
- To learn more about defining and processing triggers, see Administering Triggers.
- For information about monitoring Message Queuing and performance tuning, see Monitoring and Testing Message Queuing.

# Administration overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administration overview**

You administer your Message Queuing network and computers primarily using snap-ins displayed in the Microsoft Management Console (MMC). For a few administrative tasks, you use the **Message Queuing Properties** dialog box, which can be opened from Computer Management. The following table shows the administrative tools used to manage Message Queuing.

| Administrative tool name | Type of tool | Type of tasks used for |
|---|---|---|
| Computer Management | MMC snap-in | Most workgroup-based tasks; managing queues and messages; some local tasks; managing user certificates and cryptographic keys |
| Active Directory Users and Computers | MMC snap-in | Most domain-based tasks; managing computers, queues, and messages |
| Active Directory Sites and Services | MMC snap-in | Some domain-based tasks; managing routing links, and foreign sites for cross-platform messaging. |

For general information about using MMC and snap-ins, see Microsoft Management Console.

MSMQ Explorer was the administrative tool used to administer MSMQ objects on MSMQ 1.0 computers, but it will not run on Windows 2000, Windows XP, or Windows Server 2003 family computers. For Windows XP and Windows Server 2003 family computers, Message Queuing is administered using the Computer Management, Active Directory Users and Computers, and Active Directory Sites and Services snap-ins in MMC consoles. This section does not discuss how to administer security for Message Queuing. For more information, see Message Queuing security overview.

**Using snap-ins to administer Message Queuing**

Three snap-ins are used to manage Message Queuing. The Computer Management snap-in gives you access to local queue and message information. The Active Directory Users and Computers snap-in and the Active Directory Sites and Services snap-in give you access to information in Active Directory for routing configurations, computers, queues, and other Message Queuing objects throughout your Windows Server 2003 family forest. For more information on the objects associated with Message Queuing, see Message Queuing and Active Directory.

Most Message Queuing objects appear in the console tree (usually the left pane) of the MMC console displaying the snap-ins. When you want to view the properties for a particular Message Queuing object, you can usually right-click the object in the console tree, and then click **Properties**. The details pane (usually the right pane) of the MMC console displaying the snap-ins typically shows the name, type, and properties (using columns) of Message Queuing objects.

Active Directory replication can delay the display of information updates, so changes you make using a snap-in do not appear until after replication takes place. For example, when you add a queue to a computer in site C from a computer in site B, you do not see the change in the applicable snap-in until the change has been replicated from site C to site B. You can press F5 after making a change to refresh the information displayed in the snap-in.

**Computer Management snap-in**

The Computer Management snap-in can be used to perform many of the daily administrative tasks involving queues and messages on computers. In fact, it is the only snap-in available to manage Message Queuing in a workgroup environment (not part of a Windows Server 2003 family or Windows 2000 domain).

You can use the Computer Management snap-in to perform the following tasks:

- Create and view private queues and set queue properties.
- View cached public queues (if the local computer has at one time belonged to a domain and had access to Active Directory).
- Pause and resume the operation of an outgoing queue.

- Pause and resume the operation of all outgoing queues for a computer.
- View message properties and purge outgoing messages.
- Manage queues and messages on other computers.

In addition, you can use the **Message Queuing Properties** dialog box, which can be opened from the Computer Management snap-in, to perform the following local administrative tasks for Message Queuing:

- Specify which local disk drives and folders on which to store message and log files.
- Register, remove, renew, and view user certificates.
- Renew cryptographic keys.
- Specify a new site to which messages will be forwarded for a mobile computer. This applies only to Message Queuing 3.0 computers that belong to a Windows NT 4.0 enterprise and are capable of connecting to MSMQ 1.0 controller servers to access the MQIS database.

Outgoing queues are visible only when using the Computer Management snap-in. They are not displayed in any other snap-in.

The Computer Management snap-in allows you to manage queues and messages on other computers (provided you have administrative permissions for the other computers). For more information on how to manage queues and messages on other computers, see Manage queues on other computers.

You can also manage the queues and messages for all computers in your workgroup (or domain) by creating an MMC console (.msc) file that displays all the applicable computers. For information about creating console files, see Creating consoles.

**Note**

- You manage Message Queuing in the context of a virtual server in a cluster using the Computer Management snap-in only. In this case, a cluster group object having the name of the Network Name resource of the virtual server is displayed. Message Queuing objects are displayed under this object.

**Active Directory Users and Computers snap-in**

You can use the Active Directory Users and Computers snap-in to perform the following domain-based administrative tasks for Message Queuing involving Active Directory:

- View all Message Queuing computers and set computer-specific properties.
- Create, view, delete, and find all public queues and set queue properties.
- Create, view, and delete local private queues and set queue properties.
- View message properties and purge messages.
- Send MQPing messages, test messages, and configure message route tracking.
- Create and delete queue aliases as well as view and change their properties.

The shortcut link used in the procedure topics (step 1) to open the Active Directory Users and Computers snap-in opens MMC in author mode. This snap-in must be initially opened in author mode so that you can select the proper view options for displaying all Message Queuing objects in Active Directory. You will note the subsequent step in the procedures:

"On the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**."

You can save these changes so that these views are always enabled by default by creating your own .msc file for this snap-in. In such a case, the first time you open the .msc file that you previously created, on the **File** menu, click **Options**, and then select

one of the three user mode entries in **Console mode**. Now the console will subsequently open in the user mode selected with the correct view enabled. Opening a snap-in console in one of the user modes (rather than in author mode) is safer from an administrative standpoint because no one else can potentially change the behavior or the look of the console.

**Active Directory Sites and Services snap-in**

You can use the Active Directory Sites and Services snap-in to perform the following domain-based administrative tasks for Message Queuing involving Active Directory:

- Create and delete routing links as well as view and set their properties (including the site gates).
- Create foreign sites and foreign computers as well as view and set their properties.
- Set the default lifetime of messages on your network (known as the time-to-reach-queue setting).

The shortcut link used in the applicable procedure topics (step 1) to open the Active Directory Sites and Services snap-in opens MMC in author mode. This snap-in must be initially opened in author mode so that you can select the proper view option for displaying Message Queuing objects in Active Directory. You will note the subsequent step in the procedures:

"On the **View** menu, click **Show Services Node**."

Again, you can save these changes so that this view is always enabled by default by creating an .msc file for this snap-in in the manner described previously.

To use Message Queuing to send and receive messages, relevant Message Queuing applications must first be designed, developed, and deployed.

# Planning, Installing, and Upgrading Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning, Installing, and Upgrading Message Queuing**

This section discusses planning and deployment issues, installation options, and upgrade strategies. For more information, see:

- Planning for Message Queuing
- Installing Message Queuing
- Upgrading to Message Queuing

# Planning for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning for Message Queuing**

This section suggests ways to plan for and deploy Message Queuing in your organization. Planning and deployment options are included. The section on business scenarios provides simple overviews for deploying Message Queuing in different organizations. For details, see the following links:

- For an overview of planning Message Queuing deployment, see Planning overview.
- For information on licensing considerations for a Message Queuing network, see Licensing considerations.
- Message Queuing can be deployed in either a domain or a workgroup environment. In the context of Message Queuing, a domain environment includes domain controllers that provide a directory service such as Active Directory. For more information, see Deploying in a domain environment.
- A workgroup does not provide or have access to such a directory service, and can be run in an environment that is not part of a domain. For more information about workgroup deployment, see Deploying in workgroup mode.
- Message Queuing 3.0 supports the sending of messages over HTTP. For more information on deploying in an Internet environment, see Internet deployment considerations.
- For an overview of how different business situations can use Message Queuing, see Business scenarios.
- For information on how Message Queuing and Windows Firewall work together to secure your server from unsolicited incoming traffic, see Windows Firewall and Message Queuing

# Planning overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Planning overview**

When planning a Message Queuing deployment, it is important to consider performance requirements. After you have assessed these requirements, you can make informed decisions about your hardware requirements and upgrading strategies. Generally, performance of Message Queuing applications, like other applications, can be improved with the following strategies:

- Increasing the amount of system memory for each server.
- Upgrading the processor or adding additional processors.
- Increasing the number, speed, or size of the disks.
- Replacing disk drives and volume sets with striped drive arrays.
- Limiting the number of other services running on the computer by moving nonessential software to other computers. A server that runs additional services needs greater hardware resources than it would otherwise.

**Disk considerations**

Consider the following:

- **Number of disks**. Since Windows Server 2003 family operating systems and Message Queuing are able to perform many disk operations in parallel, using separate physical disks (as opposed to single disks with multiple partitions) will result in performance improvements for most Message Queuing applications. Applications that use recoverable or transactional messages will see the greatest improvement because these messages are constantly being written to disk and retrieved from disk.
- **Disk type**. Not all hard disks deliver equivalent performance. In particular, there are high-performance disks that use hardware striping and battery-protected write-through disk controllers that delay write operations until they can be performed most efficiently. These disks will significantly improve messaging performance, especially when sending or receiving recoverable or transactional messages. Note also that for the best messaging performance on Message Queuing servers, put the messaging (*.mq) files, message log files, and transaction log files on separate physical disks.
- **Disk size**. When deciding disk size for your computers, consider the space requirements for Message Queuing files, the type of messaging (recoverable, transactional or express), the total number of messages that are likely to accumulate on the server at any time, and the total size of those messages. It can be difficult to anticipate message storage requirements. However, if you use Message Queuing servers for session concentration within a site (in-routing and out-routing servers) or in routing links (site gates), you must allow for additional message storage space on those servers. Also, servers that support dependent clients are likely to need additional disk space.
- **Number of disk drives**. Adding disk drives ensures that disk access has a minimal impact on messaging performance. However, if your CPU (processor) usage is at or near 100%, adding disk drives does not improve messaging performance. To determine whether your computers are limited by disk access or by processing power, use System Monitor to track the **% Processor Time** counter (located in the Processor object) and the **Avg. Disk Queue Length** counter (located in the DiskLength object). If the sustained **% Processor Time** value is above 75% during the time messages are being sent, adding processing capacity may improve messaging performance. If the **Avg. Disk Queue Length** value for any drive is greater than 0.6 when messages are being sent, additional disks may improve messaging performance. For more information on monitoring Message Queuing performance, see Monitoring Message Queuing performance.

**Memory considerations**

Consider the following:

- **Storing messages**. To maximize messaging performance, there must be enough memory on a given computer to hold all of the messages that are expected to accumulate in its queues under normal operation. Messages may accumulate on the source computers, if destination computers are unreachable. On destination computers, messages will accumulate if receiving applications are not running, or are unable to keep up with the arrival rate of messages.

- **Calculating RAM requirements**. To calculate the amount of RAM required to hold all messages, you must consider message sizes. For example, when sending 20,000 messages that are approximately 1 kilobytes (KB) each with typical header sizes, allow at least 23 megabytes (MB) (20,000 × 1 KB + 20,000 × 150) of available RAM beyond normal system requirements. Note that this applies only to cases where all messages actually accumulate on a computer. If messages are normally retrieved (removed from queues) as quickly as they are delivered, significantly less RAM will be required.

**Rolling out Message Queuing**

A rollout process can consist of the following phases:

- **Setting administrative policy**. When setting out such a policy, consider how access control will be established in your organization, how administrative duties will be delegated, how work will be documented, and how backups and emergencies will be handled.

- **Conducting lab tests**. Lab testing helps you to determine the load that Message Queuing places on your servers, and tests your in-house Message Queuing applications. Before conducting such tests, ensure that your computers meet the minimum hardware requirements for running Message Queuing, and that your Message Queuing applications are functioning properly.

- **Conducting one or more pilot rollouts**. After Message Queuing has been tested in a lab, you can deploy in a limited environment. It is recommended that the pilot rollout be introduced on a limited scale, such as 15 to 50 users. You can then make any needed adjustments to your administrative policy and to your Message Queuing applications for the final rollout. If you make extensive changes to the final rollout plan, consider doing a second pilot rollout.

- **Conducting the final rollout**. After analyzing the results of the pilot rollouts, you are ready to finalize the plans for the full-scale rollout of Message Queuing in your organization. The final rollout plan is an extension of the pilot planning process, with the added steps of documenting, budgeting for, and carrying out the final logistics.

# Licensing considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Licensing considerations**

Message Queuing counts Windows Server 2003 family client access licenses (CALs). If no CALs are available, Message Queuing applications can fail to send messages and to open all queues. Message Queuing enforces the following rules for counting CALs:

- An independent client running on Windows XP Professional can have no more than 10 concurrent Message Queuing sessions and remote-read sessions with other independent clients running on Windows XP Professional.
- A dependent client does not count CALs, but its supporting server counts any sessions or queues opened on behalf of the dependent client. All Message Queuing sessions, remote-read sessions, and queues opened by dependent clients are counted. The number of dependent clients that can be supported is limited by the number of server CALs on the server.
- All sessions used by a Message Queuing computer are counted only once.

For more information on licensing for Windows Server 2003 family operating systems, see Licensing, in the Windows Help file.

# Deploying in a domain environment

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying in a domain environment**

During installation, if you choose to install Message Queuing with Active Directory Integration and the installation computer is part of a domain, Message Queuing objects are created in Active Directory. For more information about installation options, see Installation overview. For more information about Message Queuing objects in Active Directory, see Message Queuing and Active Directory.

An Active Directory domain includes one or more domain controllers that store directory information for the domain. The collection of all domains in a Windows Server 2003 network is called a forest, and one or more domains in the forest can be designated as a global catalog, with directory information for other domains in the forest, in addition to its own. Active Directory employs a multi-master model, so that any Windows Server 2003 family domain controller can read from or write to objects stored in Active Directory. For more information, see Introduction to Active Directory, in Active Directory Help.

**Accessing Active Directory**

In a domain environment, computers running Windows XP, or a Windows Server 2003 operating system, and Message Queuing can directly access Message Queuing objects in Active Directory. On the other hand, Windows 2000 computers running Message Queuing require either a Windows Server 2003 domain controller running the Downlevel Client Support Component to provide the Message Queuing Directory Service, or a Windows 2000 domain controller hosting a Message Queuing server in its site and its domain to access such objects in Active Directory. In addition to this requirement, a Windows NT 4.0, Windows 98, or Windows 95 computer running MSMQ 1.0 also requires that this domain controller be configured as a global catalog server.

**Note**

- When receiving a user's request to Active Directory, Active Directory Service Interfaces (ADSI) resolves the query by attempting to locate a domain controller in the user's domain. This can potentially be a problem for users who log on to remote computers in a different domain from their own. If the remote computer's domain is not connected to the user's domain, the Active Directory query will fail.

**Sites**

In a domain environment, a Message Queuing network can be divided into different Windows Server 2003 family *sites*, which are interconnected using *routing links*. Sites map the physical structure of your network, whereas domains generally map the logical structure of your organization, independently of each other. There may be multiple domains in a single site as well as multiple sites in a single domain. For more information, see Message Queuing and Active Directory. For more general information on sites in Windows Server 2003 family, see Sites overview, in the Active Directory Help file.

In the context of Message Queuing, a site can consist of the following:

- Windows Server 2003 domain controllers, which hold configuration and status information in the Active Directory and replicate such information between sites using site links. Active Directory employs what is called a multi-master model, which means that any Windows Server 2003 domain controller can read from or write to objects stored in Active Directory.

- Windows Server 2003 family computers, all of which use the same network protocol (typically IP). Any two computers in a site that use the same network protocol have direct connectivity with each other. Such connectivity also implies fast and cheap communication.

- Collections of associated subnets, each with a single IP subnet address.

When a Message Queuing client communicates with a Message Queuing server, it directs the request to a predetermined server in the same site, if one is available.

**Site considerations**

All sites are determined during the planning and installation of Windows Server 2003 family computers in your organization. There is not a separate site topology as there was in previous versions of Message Queuing. This simplifies your site planning.

However, it is recommended that you consider the factors listed in the following table.

| Site consideration | Requirement or suggestion |
|---|---|
| Communication link | The communication link between sites must be permanent. |
| Available network bandwidth | The network bandwidth must be able to support the volume of messages transmitted within the site. |
| Performance | The number of domain controllers in each site affects performance. Each additional domain controller in a site generates more network traffic for Active Directory replication throughout that site. You need to weigh the benefits of redundancy against the performance overhead of increased replication. |
| Performance and connectivity | Configuring a domain controller as a global catalog server in each site will enhance performance and reduce connectivity issues. |
| Organizational | Group together users who work in the same site to improve overall performance of the system, reduce network traffic, and reduce resource use. |

**Note**

- There is no longer a special definition of a site in the context of Message Queuing as there was for MSMQ 1.0. When upgrading from MSMQ 1.0, each MSMQ 1.0 site is mapped to a Windows Server 2003 family site, and for cross-platform messaging, each foreign connected network is mapped to a foreign site. If you are upgrading from MSMQ 1.0, see Upgrade overview for a list of changes.

**Offline mode**

A Message Queuing computer can also belong to a domain but be temporarily unable to communicate with a domain controller. This is known as *offline mode* and occurs when:

- The computer itself is offline.
- All the domain controllers in its site are offline.
- An attempt is made to access a remote computer and the remote computer is temporarily unable to query a domain controller for authentication.

In all these cases, the computer's configuration as a member of a domain is maintained in the registry, and normal domain operation resumes as soon as a domain controller becomes available. While in offline mode, messages can be sent directly by using *direct format names*. When other format names are used, the messages are stored on the local computer and are sent as soon as a domain controller becomes available. For more information, see "Format names" in Queue names.

**Managing computers in a domain under the Local User account**

Computers that are part of a domain can be managed under the Local User account. To do this, the Message Queuing Directory Service must be running on the domain controller, and you need to add the **EnableLocalUser** key (DWORD=1), to the registry of the local computer, at location
**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Parameters\EnableLocalUser**

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

# Deploying in workgroup mode

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Deploying in workgroup mode**
For Message Queuing, any mode of operation that does not allow access to Active Directory is known as *workgroup mode*. Note that this is regardless of whether a computer belongs to a domain, and that computers that belong to a domain can operate in workgroup mode. Computers in workgroup mode require direct connectivity with a destination computer for messaging.

**Computers that belong to a domain**
If you are installing Message Queuing on a computer that is a member of a Windows Server 2003 or Windows 2000 domain, you can specify whether the computer will operate in domain mode, by installing the Active Directory Integration subcomponent, or in workgroup mode, by not installing the subcomponent. In the context of Message Queuing, workgroup mode can be defined as any mode of operation that does not allow access to a directory service, even if the applicable computer belongs to a domain.

The default installation setting is that the Active Directory Integration component is installed, and when your computer is part of a domain, Setup will attempt to contact a Windows Server 2003 or Windows 2000 Server domain controller in the local site or a nearby site. If successful, the computer will operate in domain mode, meaning that access to Active Directory will be enabled. If one of these domain controllers cannot be located, you will be prompted for the name of an MSMQ 1.0 PEC or PSC (in a Windows NT 4.0 enterprise) in the local site. If you do not supply such a name, and select **Next** in the setup process, Message Queuing will be installed in workgroup mode. In such a situation, where the Active Directory Integration subcomponent is installed but the computer is unable to access a domain controller, the next time you restart such a computer, or stop and restart the Message Queuing service, Message Queuing will automatically attempt to contact a domain controller in the local domain. If a domain controller is available and contact is successful, the computer will then operate in domain mode.

**Computers that do not belong to a domain**
For computers that are not part of a domain and have no access to Active Directory, during default setup they are installed as independent clients in workgroup mode. The Active Directory Integration subcomponent is still installed, and if the workgroup computer later joins a domain, this subcomponent is activated. Conversely, a Message Queuing computer can be part of a domain and then join a workgroup. Such a computer can later rejoin the same domain.

The following restrictions exist when using Message Queuing in workgroup mode:

- Computers in workgroup mode require direct connectivity with a destination computer, and only support direct message transmission. Messages sent by such computers cannot be routed.
- There is no access to Active Directory. As such, you can create and manage private queues only on a local computer. You can view the list of private queues on another computer, and the messages in them, using the Computer Management snap-in. For information on managing queues on remote computers, see Manage queues on other computers. You cannot view or manage public queues or any other information in Active Directory. You can, however, send messages to or retrieve messages from private queues if there is direct connectivity.
- Internal certificates cannot be used for sending authenticated messages; an external certificate must be used. For more information on user certificates, see Managing user certificates.
- Messages cannot be encrypted. For more information on encryption, see Encryption for Message Queuing.
- There is no support for dependent clients. For more information on restrictions on the deployment of dependent clients, see Dependent clients.
- Cross-platform messaging is not supported. For more information, see Cross-Platform Messaging.

**Joining a Windows NT 4.0 domain**
When a computer running Message Queuing in workgroup mode joins a Windows NT 4.0 domain, the installation subcomponent for Active Directory Integration is removed. After the computer has been joined to the domain and rebooted,

open the Control Panel, double-click **Add or Remove Programs**, click **Add/Remove Windows Components**, and in the Windows Components Wizard, select **Application Server**, click **Details**, select **Message Queuing**, click **Details**, and select **Active Directory integration**.

# Internet deployment considerations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Internet deployment considerations**

Message Queuing 3.0 has increased support for Internet messaging and includes the following features:

- HTTP as an optional transport protocol
- An XML-based SOAP extension that defines a reliable end-to-end messaging protocol
- Support for load-balancing, Web farms, and firewalls

If your network will be sending or receiving messages through the Internet, consider the following:

- **Network load balancing**. For increased scalability and availability, Message Queuing integrates with common network load-balancing mechanisms, namely, Network Load Balancing (NLB), NAT, and DNS round-robin. For more information, see Message Queuing in Network Load Balancing clusters.
- **Messaging through perimeter networks**. Most organizations connect to the Internet using a proxy server and a perimeter network protected by firewalls. Incoming messages pass through the firewall, and outgoing messages are sent through a proxy server. The proxy server is a mediator between the organization's intranet and the Internet world. It provides caching and decoupling. Message Queuing uses an HTTP proxy to connect to Web addresses for delivering HTTP messages. Message Queuing thus needs to know the name of the proxy server and the applicable port to which it will connect. At the same time, there may be a list of target names for which the proxy connection is bypassed; that is, Web addresses that do not need to be accessed through the proxy server. An attempt is made to establish a direct connection with such targets. All of the information regarding proxy servers is stored in registry keys. For more information, see Configuring firewalls for Message Queuing and Configuring Message Queuing with a proxy server.
- **Deploying Message Queuing in hardened MSMQ mode**. This new feature enhances the security of Message Queuing computers running on the Internet. Hardened MSMQ mode is intended to support scenarios that employ only HTTP (SRMP) messages. For more information, see Message Queuing in hardened MSMQ mode.

For information on Internet messaging, see Internet Messaging.

# Business scenarios

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Business scenarios**

The following scenarios show how different business situations could benefit from using Message Queuing. They also serve as representative examples of the different options available when you use Message Queuing.

Because the differences between using Message Queuing servers and clients affect implementation details, not actual business scenarios, the scenarios have been kept simple. It is recommended that application developers see the Message Queuing SDK for representative samples of Message Queuing applications.

**Floor automation scenario**




You can use Message Queuing to control the automated manufacturing of any product that is produced by a numerically controlled industrial machine that is, in turn, controlled through a computer. In the example in the accompanying figure, Message Queuing is used to collect work orders and manage the information flow for production.

The following Message Queuing features are critical to the success of floor automation:

- **Guaranteed message delivery**. Ensures that each message is delivered and helps ensure a reliable messaging system.
- **Message prioritization**. Efficiently routes messages. For example, although quality assurance (QA) messages need to be delivered eventually, they should not obstruct the efficient delivery of time-critical messages.
- **Connectionless messaging**. Enables all computers and machines to interact as a single system.
- **Multiple-platform support**. Integrates personal computers and legacy systems. For example, an accounting package can reside on an IBM mainframe MVS system and can communicate with Message Queuing applications residing on personal computers using Microsoft MSMQ-MQSeries Bridge. For more information, see Foreign computers and sites.
- **Performance**. Because messages can be delivered quickly, they can control real-time production processes, such as cooking temperature.
- **Asynchronous delivery**. Makes it possible for faster devices to continue processing without having to wait for slower devices to respond.

**Ticket sales scenario**



The accompanying figure shows how multiple clients can access a ticket sales Web server using the Internet. Web clients view HTML pages that show ticket availability, pricing, and seating diagrams.

A business selling tickets to events can use the following Message Queuing features to improve its system:

- **Internet connection**. Accepts orders over the Internet by using a Message Queuing application and COM components in conjunction with Active Server Pages (ASP) and Internet Information Services (IIS). This is done through Web forms, which provide security and multiple-platform support.

- **Connectionless messaging**. Account updating and ticket shipping, performed asynchronously.

- **Multiple-platform support**. Integrates personal computers and legacy systems. For example, the ticket sales server might be running on a UNIX computer. For more information, see Foreign computers and sites.

**Publish-and-subscribe scenario**

Message Queuing can be used to simplify the development of typical publish-and-subscribe programs, such as news wire services and financial data subscription services. Although a server program must continuously send updated data of a specified type to the client, the server does not have to resend the data if the client is offline or unavailable.




The accompanying figure shows a publish-and-subscribe scenario that keeps a mobile sales staff current on the status of available goods. A Message Queuing application can also be used to take orders, and the two systems can be integrated. However, to simplify the details of the example, this publish-and-subscribe scenario covers only the dissemination of information.

In this figure, mobile clients retrieve data in the field. Regional servers retrieve data from the main server only for the types of information to which the clients subscribe. Each client then receives only information it subscribes to receive. This reduces both intersite communication and mobile communication costs.

In this scenario, the regional servers and the data center are located in different sites. Messages are routed between the data center and the sites along routing links in which the main server in the data center and the regional servers are configured to be site gates.

The following Message Queuing features are critical to the success of traditional publish-and-subscribe programs:

- **Connectionless messaging and message prioritization**. Places the burden of message delivery on Message Queuing. You do not have to deal with information that expires (which, therefore, is no longer useful) or with information that cannot be delivered in a timely fashion.

- **Security**. Ensures that the information is delivered to the correct client, and ensures that services are not used without payment.

**Banking scenario**




Financial transactions require both security and transaction-based programming. Automated teller machine (ATM) transactions also require connectionless messaging. In the accompanying figure, each ATM has its own local database. Even if the local, regional, or headquarters database is temporarily offline, the ATM can still process transactions.

In this scenario, the regional data centers, each of which has direct connections to multiple local data centers that communicate directly with several ATMs, are located in different sites. Messages are routed between the headquarters data center and the sites along routing links, in which the main server in the headquarters data center and the servers in the regional data centers are configured to be site gates.

This banking scenario uses these Message Queuing features:

- **Connectionless messaging**. For communication between ATMs and regional centers.
- **Transactional messaging**. For sending messages that occur synchronously with the updating of the financial databases. Messages are delivered only once.
- **Security**. For ensuring the integrity of the system. Would-be thieves cannot impersonate an ATM, and transactions cannot be monitored by persons attempting to steal account numbers and personal identification numbers.

**Stock brokerage scenario**




In the example shown here, customers communicate by telephone with traders who make trade requests on their behalf. The trade server performs the following tasks in the following order:

1. Queries the account server to verify that the transaction can be completed.
2. Verifies the transaction with the risk-management server.
3. Places the trade request with the stock exchange server.

The stock exchange server makes the stock transaction and then asynchronously updates the customer account and risk-management server.

This stock transaction scenario uses these Message Queuing features:

- **Security**. For message authentication, encryption, and access auditing to ensure a secure environment.
- **Multiple-platform support**. Integrates personal computers and legacy systems. For more information, see Foreign computers and sites.

- **Transactions (exactly once delivery)**. For guaranteed message delivery, ensuring accurate accounting.
- **Connectionless messaging**. For updating accounts and making risk assessments asynchronously.

# Windows Firewall and Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Firewall and Message Queuing**

Windows Firewall monitors and restricts information that travels between your computer and a network or the Internet. Windows Firewall drops all incoming traffic that has not been sent in response to a request of the computer (solicited traffic) or specified as traffic from a program or port that has been added to the exceptions list. This behavior of Windows Firewall provides a level of protection from malicious users and programs that use unsolicited incoming traffic to attack computers. By default, Windows Firewall is disabled for Windows Server 2003 Service Pack 1 (SP1).

When Windows Firewall is enabled, there is more protection for Windows computers, however, it can impair some types of communication and affect applications such as Message Queuing. To help overcome this, Windows Server 2003 SP1 allows administrators to define an exceptions list of applications by specifying the path to the file name of each application.

**Configuring Windows Firewall for Message Queuing**

In Windows Server 2003 Service Pack 1 (SP1), Windows Firewall is disabled by default for all connections including LAN (wired and wireless), dial-up, and virtual private network (VPN) connections. New connections also have Windows Firewall disabled by default. Windows Firewall is enabled by default only when it is explicitly enabled by using Network Setup Wizard or Internet Connection Wizard or if Windows Firewall/Internet Connection Service is enabled when you apply Windows Server 2003 Service Pack 1 (SP1).

By default, when you install Message Queuing or upgrade Message Queuing, Setup automatically adds Message Queuing to the exceptions list in Windows Firewall. You can remove Message Queuing from the exceptions list by running the following command:

**netsh firewall delete allowedprogram program=** *%windir%\system32\mqsvc.exe* **profile** *=all*

By default, when installing or upgrading to Windows Server 2003 SP1, the Downlevel Client Support component (mqdssvc.exe) is added to the exceptions list. This service is separate from the Message Queuing service (mqsvc.exe) and is not installed by default. To remove mqdssvc.exe from the exceptions list, run the following command:

**netsh firewall delete allowedprogram program=** *%windir%\system32\mqdssvc.exe* **profile** *=all*

To add Message Queuing to the exceptions list, run the following command:

**netsh firewall add allowedprogram program=** *%windir%\system32\mqsvc.exe* **name=** *msmq* **mode=** *enable* **scope=** *all* **profile** *=all*

To add Message Queuing Downlevel Client Support component (mqdssvc.exe) to the exceptions list, run the following command:

**netsh firewall add allowedprogram program=** *%windir%\system32\mqdssvc.exe* **name=** *mqds* **mode=** *enable* **scope=** *all* **profile** *=all*

To view the exceptions list, run the following command:

**netsh firewall show allowedprogram ENABLE**

When running Message Queuing with Windows Firewall enabled, in order to view public queues on remote computers by using Active Directory Users and Computers, it might be necessary to add Remote Administration to the exceptions list.

**netsh firewall set service type=** *remoteadmin* **mode=** *enable* **scope=** *all*

When installing MSMQ HTTP Support, setup automatically installs Internet Information Services (IIS), which is a World Wide Web Publishing service, and creates an IIS extension for Message Queuing (called MSMQ). Refer to the IIS documentation for information about configuring IIS for Windows Firewall.

You can also use the Security Configuration Wizard included with Windows Server 2003 SP1 to add and remove Message Queuing to and from the exceptions list. For more information, seeSecurity Configuration Wizard Overview.

For computers running Windows Server 2003 SP1 that are joined to a domain, you can also configure Windows Firewall settings by using Active Directory, and the Windows Firewall settings in Computer Configuration Group Policy of Windows 2000 or Windows Server 2003. The domain policy will override the local Windows Firewall settings.

# Installing Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Message Queuing**

This section explains the requirements for installing Message Queuing software on computers running Windows Server 2003 family operating systems. It also explains how to install new Message Queuing software remotely, using a script. For details, see the following links:

- For general installation information, see Installation overview.
- For information on installation requirements and restrictions, see Installation requirements and restrictions.
- For information on permissions required to install Message Queuing components, see Installation permissions.
- To help you select which Message Queuing subcomponents to install, see Choosing installation subcomponents.
- For information on integrating Message Queuing with Active Directory, see Active Directory Integration subcomponent.
- For information on installing Message Queuing servers, independent clients, and dependent clients, see Installing Message Queuing Computers.
- For a guide to installing Message Queuing triggers, see Installing Message Queuing triggers.
- For advice on installing Message Queuing in a server cluster configuration on Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition computers, see Installing Message Queuing in a server cluster.
- For information about using Windows unattended setup feature to install Message Queuing software, see Unattended setup.
- For help with uninstalling Message Queuing software components, see Uninstalling Message Queuing.

# Installation overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installation overview**

Message Queuing is not installed by default during installation of Windows Server 2003 family operating systems, but the files needed to install it and its subcomponents are copied to the computer during this process. When you install Message Queuing, an additional 6.7 megabytes (MB) of space is required for Message Queuing folders. You install Message Queuing on a computer by adding it as a Windows component after installation of the operating system is complete. The Message Queuing component can be found in **Add or Remove Programs**, **Add/Remove Windows Components**, **Application Server**, **Message Queuing**.

To configure Message Queuing, you choose the Message Queuing subcomponents you require by selecting **Message Queuing**, clicking **Details**, and then selecting or clearing the applicable check boxes.

The Message Queuing subcomponent tree has the following form:

Message Queuing

├Common

│ ├Core

│ └Local Storage

├Active Directory Integration

├Triggers

├MSMQ HTTP Support

├Routing Support

└Downlevel Client Support

For more detailed information on these components, see Choosing installation subcomponents.

Message Queuing can be installed in two modes:

- **Active Directory mode**. If you install Message Queuing with Active Directory Integration and the installation computer is part of a domain, Message Queuing objects are created in Active Directory. This mode supports direct message transmission, and messaging using path names. For more information on these types of messaging, see Queue names. For more information on Active Directory domain mode, see Deploying in a domain environment.

- **Workgroup mode**. Any mode of operation that does not allow access to Active Directory is known as workgroup mode, regardless of whether a computer belongs to a domain. Computers in this mode require direct connectivity with a destination computer, and only support direct message transmission. Messages sent by such computers cannot be routed. For more information, see Deploying in workgroup mode.

Note that during a default installation, the Active Directory Integration subcomponent is installed. If the computer is in the workgroup mode, the presence of this subcomponent allows for the possibility that the computer might later join a domain. In that case, the computer will try to integrate with Active Directory and attempt to create a Message Queuing object in Active Directory, each time the Message Queuing service is started. For more information, see Active Directory Integration subcomponent.

By choosing the subcomponents you require, you can decide on Message Queuing functionality for your Windows Server 2003 family computer, as follows:

- **Dependent client**. Dependent clients require the Local Storage subcomponent of the Common component to be selected. Dependent clients rely on a supporting server for all their messaging functions and cannot use new Message Queuing 3.0 features. It is recommended that you install independent clients, rather than dependent clients, where possible.

- **Independent client**. These clients can operate in workgroup mode or Active Directory domain mode. Computers configured as independent clients run the Message Queuing service, host queues, send and receive messages, operate while disconnected from the network, access Active Directory directly, and can act as a supporting server for dependent clients. However, for an independent client in workgroup mode, direct connectivity is required to send messages, and the computer cannot act as a supporting server for dependent clients. To differentiate between independent clients and Message Queuing servers:

    - Independent clients do not provide the message routing capabilities of a Message Queuing server with the Routing component enabled.

    - Independent clients do not have the Downlevel Client Support component enabled, and cannot provide Active Directory access for MSMQ 1.0 clients running on Windows NT 4.0, Windows 98, or Windows 95 computers, or Message Queuing 2.0 clients running on Windows 2000 computers.

**Default installation settings**

A default installation installs the Common (Core and Local Storage) and Active Directory Integration components. Installing the Active Directory Integration component means that if your computer is part of a domain, and has access to Active Directory, your computer will operate in domain mode following installation. If your computer is in workgroup mode, or in a domain but without Active Directory access, the Active Directory Integration component attempts to access a domain controller, and work in domain mode each time the Message Queuing service is restarted. This allows you to join your workgroup computer to a domain and work in domain mode at any time following installation.

If you are installing Message Queuing on a computer that is a member of a Windows Server 2003 family domain or Windows 2000 domain, you can still specify whether the computer will operate in domain mode (by ensuring that Active Directory Integration subcomponent is selected) or in workgroup mode (by canceling the selection of the Active Directory Integration subcomponent).

During any installation or upgrade, there are certain requirements and restrictions that apply and must be met before you can install Message Queuing. For more information, see Installation requirements and restrictions. For more information on installation permissions, see Installation permissions.

**Unattended setup**

Windows unattended setup feature can be used to install Message Queuing software remotely. For more information, see Unattended setup.

# Installation requirements and restrictions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installation requirements and restrictions**

**Installation requirements**
Regardless of whether you are performing a new installation of Message Queuing or upgrading from previous versions, the following requirements must be met before you can successfully install Message Queuing on a computer running a Windows Server 2003 family operating system:

- You must be logged on with the proper administrative permissions for the computer on which you want to install or upgrade Message Queuing. For more information on the permissions required for installation, see Installation permissions.

- It is recommended that you have network connectivity if you want to install or upgrade Message Queuing on a computer that is part of a Windows Server 2003 family or Windows 2000 domain. However, Message Queuing can be installed offline.

- To support MSMQ 1.0 clients on Windows NT 4.0, Windows 98, or Windows 95 computers, and MSMQ 2.0 clients on Windows 2000 computers in a Windows Server 2003 domain, you must install the Message Queuing Downlevel Client Support component on a Windows Server 2003 domain controller or install Message Queuing 2.0 on a Windows 2000 domain controller in the local domain.

- If a domain controller in the local site or in a nearby site cannot be accessed or does not exist, the installation process can continue. In this case, if the Active Directory Integration subcomponent is not selected, Message Queuing will be installed in workgroup mode, even if the computer on which Message Queuing is being installed belongs to a domain. If the Active Directory component is selected, the installation process continues, and Message Queuing will attempt to create the Active Directory object every time the Message Queuing service is restarted.

**Installation restrictions**
To successfully install or upgrade Message Queuing, the following restrictions must be observed:

- Do not install Message Queuing simultaneously on multiple Windows Server 2003 family domain controllers. In Active Directory, creating an object with the same name at the same time on multiple domain controllers may result in objects' names being changed. This can happen because it takes a certain amount of time for changes to be replicated among domain controllers.

- Do not manually create Message Queuing objects (such as the msmq object, for example) in Active Directory before you install or upgrade Message Queuing.

In addition, the following restrictions apply to specific subcomponent configurations:

- A dependent client cannot be installed in workgroup mode. Thus, the Core subcomponent cannot be installed without the Local Storage subcomponent in workgroup mode.

- After Message Queuing is installed, the Local Storage subcomponent cannot be added or removed.

- The domain controller that an independent client will attempt to contact during installation cannot be defined in advance in attended setup. If you want to specify a pre-Windows Server 2003 family domain controller, you must run Setup in unattended mode and include a **ControllerServer** entry in the [msmq] section.

- No additional subcomponents can be installed on a dependent client. If you want to add a subcomponent, you must first uninstall Message Queuing and install it again as an independent client.

- The Downlevel Client Support component cannot be installed if the computer does not belong to a domain.
- The Downlevel Client Support component cannot be installed if no domain controller is available, as in *offline mode*.
- The Routing Support component cannot be installed if the computer does not belong to a domain.
- The Routing Support component cannot be installed if no domain controller is available, as in *offline mode*.

Finally, there may be certain restrictions on creating queues on computers that are specified during Message Queuing Setup. For more information, see Access control for Message Queuing.

**Installing Message Queuing using a Windows NT 4.0 user account**
When you install Message Queuing on a computer running Windows Server 2003 family and you are logged on using a Windows NT 4.0 domain user account, it is recommended that you have at least one Message Queuing server running on a Windows Server 2003 family domain controller that is in the same domain in which you are installing. Otherwise, you will be prompted to enter the name of the primary site controller (PSC) in a Windows NT 4.0 domain during setup.

**Installing MSMQ 1.0 in a Windows Server 2003 family domain**
The following applies if you are installing MSMQ 1.0 on a computer running Windows NT 4.0, Windows 98, or Windows 95 in a Windows Server 2003 family domain:

When prompted for the name of a PSC during setup, you must enter the name of a domain controller instead. This Windows Server 2003 family domain controller must be running the Message Queuing directory service (Downlevel Client Support), must be configured as a global catalog server, and must be in the same site and the same domain in which you are installing MSMQ 1.0.

If you plan to install MSMQ 1.0 on computers running Windows NT 4.0 in a new Windows Server 2003 family domain, you select the Downlevel Client Support subcomponent during installation of the first domain controller in the new domain. If you are installing MSMQ 1.0 on computers running Windows NT 4.0 in an existing Windows Server 2003 family domain, you verify that the Everyone group has been added to the "Pre-Windows Server 2003 family Compatible Access" group for the domain. To do this, open Active Directory Users and Computers. In the console tree, double-click **Builtin**, right-click **Pre-Windows Server 2003 family Compatible Access**, click **Properties**, and on the **Members** page, check to see if the Everyone group is listed. If not, click **Add**, and then select **Everyone**.

Also, dependent clients running Windows NT 4.0 must have Windows NT 4.0 Service Pack 4 (SP4) or higher installed. SP4 or higher is required for transactional messaging.

In addition, for computers running Windows 98 or Windows 95, you must be granted the Create Computer Objects permission for the Computers container object in Active Directory, as the computer object for such computers does not exist in Active Directory.

To allow users to install MSMQ 1.0 in a Windows Server 2003 family domain, you must use the Delegation of Control Wizard to grant specific permissions to all users in the domain. For more information, see Set installation permissions to enable upgrade from MSMQ 1.0.

# Installation permissions

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installation permissions**

The following table shows the administrative permissions required to install different Message Queuing software components and to create various Message Queuing computers. By default, logging on using an account with these permissions allows the applicable Message Queuing objects to be created in Active Directory during Message Queuing installation.

| Message Queuing computer | Permission level required |
|---|---|
| Message Queuing server on a domain controller | domain administrative permissions (or member of the Domain Admins group) |
| Message Queuing server on a nondomain controller with Routing Support | enterprise administrative permissions (or member of the Enterprise Admins group). |
| Message Queuing server on a nondomain controller without Routing Support | local administrative permissions (or member of the local Administrators group) |
| Independent client | local administrative permissions (or member of the local Administrators group) |
| Dependent client | local administrative permissions (or member of the local Administrators group) |

For more information on administrative groups for Windows Server 2003 family, see Default groups, in the Active Directory Help file.

**Notes**

- When installing a Message Queuing server on a nondomain controller with routing enabled, domain administrative permissions are sufficient, provided Setup can contact a Windows Server 2003 or Windows 2000 domain controller and provided you are granted the specific permissions for objects listed in the table that follows.

- When installing Message Queuing in a server cluster, the user account under which the cluster service is running must be granted permissions to create computer objects and Message Queuing objects. This can be done only on a Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. See the following table for specific permissions needed to create specific objects.

If you do not want to grant users such general or wide-ranging permissions to install Message Queuing, you can grant permissions to create those objects specifically required to create different Message Queuing computers. The following table lists the specific child objects in Active Directory that you must be granted permission to create for Message Queuing installation to complete.

| Message Queuing computer | Specific permissions required | For which object |
|---|---|---|
| Message Queuing server (on a domain controller or on a nondomain controller with routing enabled) | Create MSMQ Configuration Objects | applicable domain controller object located in Active Directory Users and Computers |
| Message Queuing server (on a domain controller or on a nondomain controller with routing enabled) | Create All Child Objects | Servers container object located in Active Directory Sites and Services |
| Message Queuing server (on a nondomain controller with no routing) | Create MSMQ Configuration Objects | applicable computer object located in Active Directory Users and Computers |

| | | |
|---|---|---|
| Independent client | Create MSMQ Configuration Objects | applicable computer object located in Active Directory Users and Computers |
| Dependent client | none | not applicable |

You can grant these specific permissions to specific users or to all users in a domain using the Delegation of Control Wizard.

For more information on Message Queuing objects, where they are created, and where they are located in Active Directory, see Message Queuing and Active Directory.

For more information on permissions for objects, see Access control for Message Queuing.

For an example of how to grant permissions using the Delegation of Control Wizard, see Set installation permissions to enable upgrade from MSMQ 1.0.

**Notes**

- To install Message Queuing, you may also need the Delete MSMQ Configuration Objects permission, because if Setup finds an msmq (MSMQ Configuration) object in Active Directory, it must delete it before creating a new one.

- To uninstall Message Queuing, you must be granted the Delete MSMQ Configuration Objects permission.

# Choosing installation subcomponents

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Choosing installation subcomponents**

The following table lists the subcomponents of Message Queuing, the 32-bit Windows Server 2003 family operating systems on which they can be installed, their dependencies on other subcomponents, and their names in the .inf file.

| Component or subcomponent name | Description and setup steps | Suitable platforms | Depends on | Name in the .inf file |
|---|---|---|---|---|
| Message Queuing | Parent for all subcomponents | All Windows Server 2003 family operating systems except Windows XP Home Edition | DTC | msmq |
| Common | Performs common setup tasks, parent for the Core and Local Storage subcomponents | All Windows Server 2003 family operating systems. | none | msmq_Common |
| Core | Performs the most fundamental setup operations, creates and sets registry keys, and provides dependent client functionality | All Windows Server 2003 family operating systems. | none | msmq_Core |
| Local Storage | Creates and installs the Message Queuing service and provides independent client functionality | All Windows Server 2003 family operating systems. | Core | msmq_LocalStorage |
| Active Directory Integration | Creates an msmq object and (if needed) an MSMQ Setting object in Active Directory and thereby provides integration with Active Directory whenever the computer belongs to a domain | All Windows Server 2003 family operating systems. | Local storage | msmq_ADIntegrated |
| Triggers | Installs the Message Queuing Triggers service | All Windows Server 2003 family operating systems. | Local storage | msmq_TriggersService |
| MSMQ HTTP Support | Creates an IIS extension for Message Queuing (called MSMQ) | All Windows Server 2003 family operating systems. | Local storage, IIS (the World Wide Web Publishing service) | msmq_HTTPSupport |
| Routing Support | Creates an MSMQ Setting object in Active Directory and provides Message Queuing routing server functionality | Windows Server 2003 family operating systems. | Active Directory Integration | msmq_RoutingSupport |
| Downlevel Client Support | Installs the Message Queuing directory service on a domain controller | All Windows Server 2003 family operating systems. | Active Directory Integration | msmq_MQDSService |

The Message Queuing subcomponent tree has the following form:

Message Queuing

├ Common

│ ├ Core

│ └Local Storage

├Active Directory Integration

├Triggers

├MSMQ HTTP Support

├Routing Support

└Downlevel Client Support

**Note**

- On 64-bit platforms, the Core and Local Storage subcomponents are combined in a single subcomponent, called Common, which in this case, provides independent client functionality.

**Default options**

When you select Message Queuing in the Windows Components Wizard on computers running on platforms from the Windows Server 2003 family, the following subcomponents are selected by default:

- Common (consisting of Local Storage and Core)
- Local Storage
- Active Directory Integration

This means that, by default, an independent client capable of integrating with Active Directory (whenever the computer belongs to a domain) will be installed on computers running Windows Server 2003 family operating systems.

If a computer belongs to a workgroup, the Active Directory Integration subcomponent is selected by default. In workgroup mode, the presence of this subcomponent means that the computer will allow for the possibility that it joined a domain and will try to integrate with Active Directory each time the Message Queuing service is started (that is, it will attempt to create an msmq object in Active Directory, and so on). As long as the computer remains in workgroup mode, no msmq object is created in Active Directory, only direct messages can be sent, and only private queues can be created or deleted.

**Computers running on 64-bit platforms**

A dependent client cannot be installed on a computer running on a 64-bit Windows Server 2003 family platform. Therefore, we can simplify the list of Message Queuing subcomponents for 64-bit computers by combining the Core and Local Storage subcomponents into a single Common subcomponent, and we do not need two levels of Message Queuing subcomponents. The subcomponent tree for 64-bit computers has the following form:

Message Queuing

├Common

├Active Directory Integration

├Triggers

├MSMQ HTTP Support

├Routing Support

└Downlevel Client Support

The Routing Support and Downlevel Client Support subcomponents are available only on platforms from the Windows Server 2003 family. Furthermore, Downlevel Client Support can be installed only on a domain controller. If only the Core subcomponent is selected, an independent client will be installed.

The following table shows the subcomponent dependencies for 64-bit computers.

| Subcomponent name | Depends on | Comment |
|---|---|---|

| | | |
|---|---|---|
| Message Queuing | DTC | parent |
| Core | none | none |
| Active Directory Integration | Core | none |
| Triggers | Core | none |
| MSMQ HTTP Support | Core, IIS (the World Wide Web Publishing service) | none |
| Routing Support | Active Directory Integration | only on servers |
| Downlevel Client Support | Active Directory Integration | only on servers |

**Registry keys**

A registry key is created in **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSMQ\Setup** for each subcomponent installed, and its value is set equal to 0x1. It indicates that the applicable subcomponent was installed successfully on the computer. The absence of a key means that the subcomponent was not installed. The following list gives the names of these registry keys:

- msmq_Core
- msmq_LocalStorage
- msmq_ADIntegrated
- msmq_TriggersService
- msmq_HTTPSupport
- msmq_RoutingSupport
- msmq_MQDSService

# Active Directory Integration subcomponent

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Active Directory Integration subcomponent**
During a default installation, the Active Directory Integration subcomponent is installed, regardless of whether your computer is a member of a domain.

If you are installing Message Queuing on a computer that is a member of a Windows Server 2003 or Windows 2000 domain, you can specify whether the computer will operate in domain mode, by installing the Active Directory Integration subcomponent, or in workgroup mode, by not installing the subcomponent. In the context of Message Queuing, workgroup mode can be defined as any mode of operation that does not allow access to a directory service, even if the applicable computer belongs to a domain. The workgroup mode of operation requires direct connectivity to send messages.

If you are installing Message Queuing on a computer that belongs to a domain and you select the Active Directory Integration subcomponent (as in the default configuration), Setup will attempt to contact a Windows Server 2003 or Windows 2000 Server domain controller in the local site or a nearby site. If successful, the computer will operate in domain mode, meaning that access to Active Directory will be enabled and an msmq object will be created in Active Directory. Note that a Windows Server 2003 or Windows 2000 Serverdomain controller can only be found automatically. If one of these domain controllers cannot be located, you will be prompted for the name of an MSMQ 1.0 PEC or PSC (in a Windows NT 4.0 enterprise) in the local site. If you do not supply such a name, and select **Next**, Message Queuing will be installed in workgroup mode. In such a situation, where your computer belongs to a domain and has the Active Directory Integration subcomponent installed, but is temporarily unable to access a domain controller, the computer will work in workgroup mode. However, the next time you restart such a computer, or stop and restart the Message Queuing service, Message Queuing will automatically attempt to contact a domain controller in the local domain. If a domain controller is available and contact is successful, the computer will then operate in domain mode.

If your computer does not belong to a domain, default installation of the Active Directory Integration subcomponent allows for the possibility that the computer might later join a domain.

For more information on the characteristics of workgroup mode, domain mode, and offline mode, see Deploying in a domain environment, and Deploying in workgroup mode.

For information on security permissions required by Message Queuing with the Active Directory Integration component on a computer that belongs to a domain, see Installation permissions.

.

# Installing Message Queuing Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Message Queuing Computers**

This section provides information about installing different types of Message Queuing computers, and the platforms that support them. For details, see the following links:

- For a summary of what types of Message Queuing computers can be installed on different Windows Server 2003 family computers, see Installation platforms for Message Queuing computers.
- For information on installing Message Queuing servers, and the necessary requirements, see Installing Message Queuing servers.
- For information on installing Message Queuing on a domain controller, see Installing on a domain controller.
- For details of requirements for installing independent clients, see Installing independent clients.
- To learn about installing dependent clients with a supporting server, see Installing dependent clients.

Note that for the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed on it. The term *Message Queuing client* can refer to either an *independent client* or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed. Such computers are referred to as having independent client functionality or dependent client functionality.

# Installation platforms for Message Queuing computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installation platforms for Message Queuing computers**

The following table shows which types of Message Queuing computers can be created on which Windows Server 2003 family operating systems.

| Type of Message Queuing computer | Windows Server 2003 family operating systems |
|---|---|
| Message Queuing server | All Windows Server 2003 family operating systems |
| Independent client | All Windows Server 2003 family operating systems |
| Dependent client | All 32-bit Windows Server 2003 family operating systems |

**Notes**

- For the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed. The term *Message Queuing client* can refer to either an *independent client* or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed, and such computers are referred to as having independent client functionality or dependent client functionality.

- A dependent client cannot be installed on a computer operating in workgroup mode or on a 64-bit computer.

# Installing Message Queuing servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Message Queuing servers**

This topic covers the installation of Message Queuing server software on computers running on platforms from the Windows Server 2003 family that do not have previous versions of Message Queuing installed.

For the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed. The term *Message Queuing client* can refer to either an *independent client* or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed, and such computers are referred to as having independent client functionality or dependent client functionality.

Before installing a Message Queuing server on a Windows Server 2003 family computer in a particular site, verify that there is at least one Windows Server 2003 or Windows 2000 domain controller trusted for delegation in the local site or in a nearby site. To reduce replication traffic, configure a Windows Server 2003 or Windows 2000 domain controller in each site as a global catalog server. During setup, an attempt is made to automatically detect the nearest Windows Server 2003 or Windows 2000 domain controller in the local site or in a nearby site. If a suitable domain controller cannot be found, Setup assumes that you are installing a Message Queuing client in a Windows NT 4.0 environment and prompts you for the name of a controller server (PEC or PSC) before Setup can continue.

If a Windows Server 2003 or Windows 2000 domain controller is not available, you can also start the installation again and install Message Queuing without the Active Directory Integration subcomponent. If this option is chosen, only private queues can be created and managed, direct connectivity is needed to send messages, and the computer will never try to access Active Directory. In the context of Message Queuing, any mode of operation that does not allow access to a directory service is called workgroup mode, even if the applicable computer belongs to a domain.

Alternatively, Message Queuing can be installed on a computer that belongs to a workgroup and has not joined a domain. In this case, Setup will install Message Queuing in workgroup mode. In this mode of operation, only private queues can be created and managed, and direct connectivity is needed to send messages. However, in this case, the Active Directory Integration subcomponent will be installed, but the computer will not be able to access Active Directory and will continue to operate in workgroup mode until it joins a domain. After the computer joins a domain, the next time you restart it, or stop and restart the Message Queuing service, Message Queuing will automatically attempt to access Active Directory through a domain controller in the local domain and create an msmq object.

Similarly, if a Message Queuing server is taken offline or if all the domain controllers from its domain in the local and nearby sites are offline, the computer will not have access to Active Directory and will operate in offline mode. As in workgroup mode, only private queues can be created and managed and direct connectivity is needed to send messages. Each time the computer is restarted, Message Queuing will attempt to locate its domain controller. If successful, Message Queuing will return to the domain mode of operation, which enables access to Active Directory.

Before you can install a Message Queuing server, the following requirements must first be met:

- You must install a Windows Server 2003 family platform from the server family.
- To install a Message Queuing server on a Windows Server 2003 family domain controller, you must be logged on with domain administrative permissions or belong to the Domain Admins group.
- To install a Message Queuing server with routing services on a nondomain controller, you must be logged on with enterprise administrative permissions or belong to the Enterprise Admins group.
- To install Message Queuing, you may also need the Delete MSMQ Configuration Objects permission, because if Setup finds an msmq (MSMQ Configuration) object in Active Directory, it must delete the object before creating a new one.
- To uninstall Message Queuing, you must be granted the Delete MSMQ Configuration Object permission.

For more information on the administrative permissions required for installation, see Installation permissions.

**Installing additional Message Queuing servers**

In a domain environment, when Message Queuing servers are installed on additional Windows Server 2003 family computers in the same site, Setup automatically attempts to locate a Windows Server 2003 domain controller, or Windows 2000 domain controller, in the local or a nearby site. If one cannot be located, Setup assumes that you are installing a Message Queuing client in a Windows NT 4.0 environment and prompts you for the name of a controller server (PEC or PSC) before Setup can continue.

For information on installing Message Queuing servers on additional Windows 2000 computers in the same site, see the version of this topic in the online Help file for Message Queuing in a Windows 2000 computer.

If you are installing a Message Queuing server with routing services in a mixed-mode domain environment, you must specify the name of a Windows Server 2003 domain controller running the Message Queuing directory service (Downlevel Client Support) or a Windows 2000 domain controller hosting a Message Queuing server in the same site (not the name of an MSMQ 1.0 PSC) if prompted during setup.

You can also choose to install a Message Queuing server without routing on a nondomain controller in a workgroup environment (or any other environment that does not provide a directory service such as Active Directory). In this case, direct connectivity is required to send messages.

# Installing on a domain controller

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing on a domain controller**

If Message Queuing is being installed only on Windows Server 2003 family computers in a particular site, there is no need to install Message Queuing on a domain controller in that site because Message Queuing clients on Windows Server 2003 family computers can access Active Directory directly. However, if you intend to install Message Queuing on a Windows 2000 computer in a site, you must install Message Queuing with Downlevel Client Support on a Windows Server 2003 domain controller in the site, or install a Message Queuing server on a Windows 2000 domain controller in the site, or promote the Windows 2000 computer on which you want to install a Message Queuing server to a domain controller. In addition, even in a pure Windows Server 2003 family enterprise, enabling routing services for Message Queuing servers running on domain controllers may degrade performance. It is therefore preferable to install such Message Queuing servers on nondomain controllers. If you need to install Message Queuing on a domain controller (because you want to run an application over Message Queuing), it is recommended that you install a Message Queuing server without routing enabled (a server with independent client functionality).

To create a server that will provide access to Active Directory for Message Queuing 1.0 clients on Windows 95, Windows 98, Windows Millennium Edition, Windows NT 4.0 computers, or Message Queuing 2.0 clients running on Windows 2000 computers, promote the computer to a Windows Server 2003 domain controller. For information on how to promote a computer to a Windows Server 2003 family domain controller, see the Active Directory Help topic on installing on a domain controller.

Furthermore, before you can install MSMQ 1.0 on clients running on Windows NT 4.0, Windows 98, or Windows 95 computers, in a Windows Server 2003 domain, there must be at least one domain controller running the Message Queuing directory service (Downlevel Client Support component) or Message Queuing 2.0 that is configured as a global catalog server in the local site from the domain of each computer on which you want to install MSMQ 1.0. For information on how to configure a domain controller as a global catalog server, see Configure a domain controller as a global catalog server. For more information on the global catalog in Active Directory, see The role of the global catalog. Note that the Message Queuing Downlevel Client Support service must run under the default LocalSystem account.

For each applicable domain controller, also verify that it is trusted for delegation before installing Message Queuing. This property is enabled by default. If this property is disabled, the Message Queuing service cannot run. For information on how to verify that this property is enabled, see Make a domain controller trusted for delegation.

If you install a Message Queuing server with Downlevel Client Support on a domain controller and security has not been weakened in your forest to support MSMQ 1.0 clients or Message Queuing 2.0 clients logged on with local user accounts, Setup prompts you about whether or not to weaken security. This mode of operation is optional. However, to effectively support computers running MSMQ 1.0 and computers running Message Queuing 2.0 that are logged on using a local user account, security for Active Directory needs to be weakened. Nevertheless, dependent clients cannot run under a local user account, and any computer that sends queries about Message Queuing objects to Active Directory on a domain controller directly, rather than through the Message Queuing directory service, will not be able to access Active Directory when it logs on using a local user account even if the security for Active Directory is weakened. For more information, see Enabling weakened security.

For information on how to install a Message Queuing server after all requirements are met, see Install Message Queuing servers.

**Important**

- When you make a change from a computer running a Windows Server 2003 operating system against directory service objects belonging to computers running Windows 95, Windows 98, Windows NT 4.0, or Windows 2000, these earlier computers are unable to receive notification about the change from the Windows Server 2003 family computer. To resolve this, Active Directory detects by default that the object belongs to an earlier computer and attempts to make the change without accessing Active Directory directly. To override this default behavior:
    - For a single computer, set the DisableDownlevelNotifications key in the registry.
    - For an enterprise-wide system, set the PROPID_E_CSP_NAME attribute in the msmq enterprise object.

- Each Windows Server 2003 family site containing MSMQ 1.0 clients running on Windows NT 4.0, Windows 98, or Windows 95 computers must contain at least one Windows Server 2003 domain controller that is running the Message Queuing directory service (Downlevel Client Support) and is configured as a global catalog server or at least one Windows 2000 domain controller that is running Message Queuing and is configured as a global catalog server from the domain of each such client. This is required to enable such clients to communicate with Message Queuing servers. For information on how to configure a computer as a global catalog server, see Configure a domain controller as a global catalog server.

# Installing independent clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing independent clients**

This topic covers the installation of Message Queuing servers with independent client functionality on computers. Computers configured as independent clients can operate in workgroup mode or Active Directory domain mode. These computers run the Message Queuing service, host queues, send and receive messages, operate while disconnected from the network, access Active Directory directly, and can act as a supporting server for dependent clients. However, for an independent client in workgroup mode, direct connectivity is required to send messages, and the computer cannot act as a supporting server for dependent clients. For more information, see Independent clients.

For the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed. The term *Message Queuing client* can refer to either an *independent client* or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed, and such computers are referred to as having independent client functionality or dependent client functionality.

Before installing a Message Queuing server with independent client functionality, note the following:

- To install Message Queuing in domain mode, you must have network connectivity. During a default installation, the Active Directory Integration subcomponent is installed, and Setup will automatically attempt to contact a Windows Server 2003 or Windows 2000 domain controller. If this attempt is successful, Message Queuing is installed in domain mode. If such a domain controller cannot be found, you will be prompted to enter the name of an MSMQ 1.0 PEC or PSC in the local site (for a Windows NT 4.0 enterprise). If you do not enter a name, Message Queuing will automatically be installed in workgroup mode. In the context of Message Queuing, any mode of operation that does not allow access to a directory service is called workgroup mode, even if the applicable computer belongs to a domain. The workgroup mode of operation requires direct connectivity to send messages.

- During a default installation, the Active Directory Integration component is installed. If your computer is in a domain, but without Active Directory access, this component will attempt to access a domain controller and work in domain mode each time the computer, or the Message Queuing service, is restarted. If your computer is not part of a domain and is in workgroup mode, installing this component provides you with the option of joining the computer to a domain at a later time.

- You can also install your Message Queuing server with independent client functionality in workgroup mode by canceling the selection of the Active Directory Integration component that is selected as part of the default installation settings.

- You must be logged on with local administrative permissions (as a member of the local Administrators group). For more information on the permissions required for installation, see Installation permissions.

- If you are installing a Message Queuing server with independent client functionality on a Windows Server 2003 family computer in a Windows NT 4.0 enterprise, and you want to specify an MSMQ 1.0 PEC during installation, you can only do this from unattended setup. For more information, see Unattended setup.

- To install Message Queuing, you may also need the Delete MSMQ Configuration Objects permission, because if Setup finds an msmq (MSMQ Configuration) object in Active Directory, it must delete the object before creating a new one.

- To uninstall Message Queuing, you must be granted the Delete MSMQ Configuration Object permission.

- A Message Queuing server configured with independent client functionality does not have the Routing subcomponent enabled. In addition, the Downlevel Client Support component is not enabled, and independent clients cannot provide Active Directory access for MSMQ 1.0 clients running on Windows NT 4.0, Windows 98, or Windows 95 computers, or Message Queuing 2.0 clients running on Windows 2000 computers.

For instructions on how to install an independent client after all requirements are met, see Install independent clients.

# Installing dependent clients

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing dependent clients**

This topic describes the installation of Message Queuing servers with dependent client functionality. It is recommended that you install independent clients wherever possible.

For the purposes of this documentation, the term *Message Queuing server* refers to a Windows Server 2003 family computer that has Message Queuing installed. The term *Message Queuing client* can refer to either an *independent client*, or a *dependent client*. Both of these terms are used to describe Message Queuing servers that have particular Message Queuing components installed, and such computers are referred to as having independent client functionality, or dependent client functionality.

Before you can install Message Queuing with dependent client functionality, the following requirements must first be met:

- You must have network connectivity and be able to contact a supporting server.
- Dependent clients cannot be installed on 64-bit computers, or on Windows Server 2003 family computers operating in workgroup mode.
- You must be logged on with local administrative permissions as a member of the local Administrators group. For more information on the permissions required for installation, see Installation permissions.
- You must have a Windows Server 2003 family computer installed with at least independent client functionality, to act as a supporting server. This computer must be configured to support dependent clients.
- Note that computers configured as independent clients cannot provide Active Directory access for MSMQ 1.0 clients. These clients require Message Queuing servers running on a domain controller to act as a supporting server.

**Supporting servers for dependent clients.**

A dependent client requires server support for all its messaging functionality. The Message Queuing server that supports a dependent client (by storing queues and messages) is referred to as the supporting server. Note that following installation, Message Queuing servers are not configured to support dependent clients. In addition, to enable a Windows Server 2003 family computer to act as the supporting server for dependent clients, you must change the security configuration of the Microsoft Distributed Transaction Coordinator (DTC) to enable network DTC. For instructions on enabling dependent client support and changing the DTC configuration, see Enable a supporting server.

During setup of a dependent client, you must specify the name of such a Message Queuing server in the same site or in a nearby site. If no name is specified, setup cannot continue. Dependent clients and their supporting servers do not have to be in the same site. However, because all communication with the dependent client goes through the supporting Message Queuing server, if the communication link to the supporting server is slow or expensive, all dependent client communication will also be slow or expensive.

For information on how to install a dependent client after all installation requirements have been met, see Install dependent clients.

# Installing Message Queuing triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Message Queuing triggers**

Message Queuing Triggers is an integral part of Message Queuing and can be installed during Message Queuing Setup by selecting the Triggers subcomponent. Triggers and rules are defined and managed in the Computer Management snap-in. Diagnostic data generated by the Message Queuing Triggers service are recorded in the Application Log and can be viewed using Event Viewer.

The support for Message Queuing Triggers can be installed by Message Queuing Setup during the installation of a Message Queuing server or independent client on computers running platforms from the Windows Server 2003 family. Message Queuing Triggers cannot be installed on a dependent client.

To install Message Queuing Triggers, select the Triggers subcomponent of Message Queuing during the installation of a Message Queuing server or an independent client. If the check box for this subcomponent is selected, the Triggers service is installed automatically and started. If this check box is not selected, the Triggers service can be installed later by running **Add or Remove Programs** and selecting **Triggers** in the Windows Components Wizard.

An answer file for unattended setup can contain an **msmq_TriggersService = ON** entry, which corresponds to selecting the **Triggers** check box on the **Message Queuing** page of the Windows Components Wizard.

Setup requires the presence of four files for the Triggers service in the \Windows\System32 folder: mqtrig.dll, mqtgsvc.exe, mqgentr.dll, and mqtgclus.dll.

By default, the Message Queuing Triggers service runs under the Network Service account. Services that are running under this account interact with the network using the credentials of the computer account, and this account has the same level of access to resources and objects as members of the Users group accounts. For heightened security, you can change the default Triggers service account to a more secure setting; for example, the Local Service account, which limits access to the local computer. You can also specify a less secure account for the Triggers service; for example, the Local System account, allowing access to the entire domain. Note that the Triggers service running under the Network Service account does not interact with the desktop. If this functionality is required, you can configure the Triggers service to run under the Local System account, and enable interaction with the desktop in Computer Management.

Message Queuing Triggers installs with a set of default configuration values that are stored in the registry. These values can be changed later using the Computer Management snap-in. For more information on how to change the default configuration, see Change default configuration values for triggers.

Message Queuing Triggers is cluster-aware and supports the active/active model. When the Message Queuing service fails over to another node, the Triggers service can fail over along with it. For this to occur, you must create a Message Queuing Triggers resource with resource dependencies on the Message Queuing resource and the Cluster Name resource on the applicable node using Cluster Administrator, and you must bring it online. The trigger and rule definitions for the cluster resource are stored in the local Windows registry. The Message Queuing Triggers resource is checkpointed with this registry key and the subkeys under it so that when the resource is moved to a new node (during failover, for example), the required keys propagate to the local registry on the new node. After failover, the Triggers service can thus continue to process the incoming messages in each monitored queue and invoke the applicable stand-alone executable or COM component according to the rules defined. For more information on server clusters and the active/active model, see Distributing Message Queuing in a network. For more information on creating and configuring Message Queuing Triggers resources in a virtual server, see Configure trigger resources for server clusters.

# Installing Message Queuing in a server cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Installing Message Queuing in a server cluster**

This section applies only to users running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition in a server cluster configuration. Installing Message Queuing in a server cluster will provide active/active queue failover in case of computer failure or malfunction.

Installing the required instances of the Message Queuing service on each node of a server cluster before creating and configuring a Message Queuing resource is similar to installing Message Queuing on any other server. However, there are several additional requirements and considerations to keep in mind when installing Message Queuing in a server cluster, including the following:

- Install Message Queuing with the same set of subcomponents on all nodes in a server cluster. For example, if you install a Message Queuing server with Routing Support on one node, install Message Queuing servers with Routing Support on the other nodes as well.
- Create one Message Queuing resource for each cluster group, using Cluster Administrator. To open Cluster Administrator, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.
- If you intend to use triggers, HTTP/HTTPS messaging, or external transactions on your cluster, before creating a Message Queuing resource in any cluster group hosted on a computer in a server cluster, do the following:
    - Cluster Microsoft Distributed Transaction Coordinator (DTC). This must be done using Cluster Administrator, and not by running comclust.exe. DTC is installed by default on Windows Server 2003 family operating systems.
    - Configure the DTC, create a DTC resource, and bring it online. Note that only one DTC resource is created for an entire server cluster. For more information on creating a DTC resource, see Checklist: Installing a Distributed Transaction Coordinator resource, in the Server Clusters Help file. For instructions on bringing a DTC resource online, see Bring a resource online, in the Server Clusters Help file.
    - After the DTC resource is online, install Message Queuing on each cluster node.
- After the preceding tasks are completed, you can create Message Queuing resources on any virtual server in the cluster. When you create a Message Queuing resource in a cluster group, configure it with a Network Name resource and a Disk resource using Cluster Administrator. For instructions on configuring Message Queuing and Message Queuing Triggers resources for server clusters, see Configure Message Queuing resources for server clusters, and Configure trigger resources for server clusters.

**Notes**

- It is recommended that Message Queuing be installed on the physical computer before creating a Message Queuing resource.
- You must not create more than one Message Queuing resource in a single cluster group, because the behavior of such multiple resources is undefined.
- To uninstall Message Queuing on a computer that is part of a server cluster, the Message Queuing virtual server must be deleted from the node before uninstalling Message Queuing on the physical node computer.
- On a cluster, the Message Queuing service start configuration is set to manual. Thus, after restarting a node computer, the Message Queuing service running on the node does not restart automatically. However, Message Queuing resources that were online do come back online automatically.
- For more information on configuring resources using Cluster Administrator, see Server Cluster Resources, in the Server

Clusters Help file.

- For information on upgrading Message Queuing in a server cluster, see Upgrading Message Queuing in a server cluster.
- For instructions on how to start the Message Queuing service, see Start or stop services for Message Queuing.

# Unattended setup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Unattended setup**

You can use Windows unattended setup to install Message Queuing software by remotely using an answer file. It is recommended that this file be located either in the Message Queuing installation folder on the source computer from which you are running unattended setup or in the Windows system folder (typically C:\Windows) on the target computer on which you are installing Message Queuing.

For more information on unattended setup for Windows Server 2003 family, see Planning for unattended Setup in the Windows Help file.

During Message Queuing unattended setup, no text or dialog boxes are displayed. Errors are logged to the Msmqinst.log file located in the Windows system folder (typically C:\Windows) of the target computer.

Because the name of an online Message Queuing server is required for some installation types, it is easiest to create and run separate unattended setup scripts for computers in each Windows Server 2003 family or Windows 2000 site. You can use a single script file for multiple sites, but if you do so, you must create separate file entries for each computer that is not in the specified (default) site.

To run Message Queuing unattended setup, type the following command at the command prompt in the Windows system folder:

**sysocmgr.exe /i:sysoc.inf /u:** *path and file name of answer file*

**Answer file entries**

The answer file specifies various installation settings. The applicable settings for Message Queuing are described below.

Setup looks in the [Components] section to determine what to do for each of the Message Queuing subcomponents. In this section **ON** means to install the subcomponent, and **OFF** means to uninstall it, while the absence of an entry means that the state of the subcomponent will not be changed. The keys appearing in entries in this section are as follows:

**msmq_Core**

An **msmq_Core = ON** entry calls for installation of the Core subcomponent, which performs the most fundamental setup operations: creates and sets registry keys, and provides dependent client functionality. It is mandatory if you want to install Message Queuing. To install a dependent client, you must specify a supporting server. The supporting server can be any Message Queuing server operating in domain mode with at least independent client functionality, preferably in the local site, but also in a nearby site. Specify its name using the **SupportingServer** key in the [msmq] section. No additional subcomponents can be installed on a dependent client. To perform an unattended uninstallation of Message Queuing software, include the entry **msmq_Core = OFF**. All message and log files are removed when Message Queuing is uninstalled. However, no Message Queuing files are removed from the \Windows\System32 and \Windows\System32\Drivers folders.

**msmq_LocalStorage**

An **msmq_LocalStorage = ON** entry calls for installation of the Local Storage subcomponent, which installs the Message Queuing service and provides independent client functionality. This entry is mandatory if you are installing any additional subcomponent. Local Storage cannot be installed without the Core subcomponent. Thus, if you are installing Local Storage, there is no need to include an **msmq_Core = ON** entry. This subcomponent will be installed automatically. After the Local Storage subcomponent is installed, it cannot be removed without completely uninstalling Message Queuing.

**msmq_ADIntegrated**

The entry **msmq_ADIntegrated = ON** calls for installation of the Active Directory Integration subcomponent, which creates an msmq object and (if needed) an MSMQ Setting object in Active Directory and thereby provides integration with Active Directory whenever the computer belongs to a domain. If you want to use a specific pre-Windows 2000 domain controller, you must specify it using the **ControllerServer** key in the [msmq] section. Active Directory Integration cannot be installed without the Core and Local Storage subcomponents. Thus, if you are installing Active Directory Integration, there is no need to include **msmq_Core = ON** and **msmq_LocalStorage = ON** entries. The applicable subcomponents will be installed automatically.

You can include an **msmq_ADIntegrated = ON** entry if you are installing Message Queuing in a workgroup environment or any environment that does not have access to Active Directory. In this case, the **ControllerServer** entry is ignored and the

computer will never attempt to connect to a domain controller to access Active Directory. Such a computer requires direct connectivity to other computers to send and receive messages. This entry is applicable only if you are installing a configuration with independent client functionality on a nondomain controller.

**msmq_TriggersService**

This key specifies whether the Message Queuing Triggers service will be installed. If you include an **msmq_TriggersService = ON** entry, it will be installed. Conversely, if you include an **msmq_TriggersService = OFF** entry, it will be removed. Setting this key equal to **ON** is equivalent to selecting the Triggers subcomponent during the installation of Message Queuing with independent client functionality. Thus, if you are installing Triggers, there is no need to include **msmq_Core = ON** and **msmq_LocalStorage = ON** entries. These subcomponents will be installed automatically.

**msmq_HTTPSupport**

This key specifies whether MSMQ HTTP Support, which installs Internet Information Services (IIS) and creates an IIS extension for Message Queuing (MSMQ), will be installed. If the answer file includes an **msmq_HTTPSupport = ON** entry, an IIS extension will be created. Conversely, if you include an **msmq_HTTPSupport = OFF** entry, the extension will be removed. MSMQ HTTP Support cannot be installed without the Core and Local Storage subcomponents. Thus, if you are installing MSMQ HTTP Support, there is no need to include **msmq_Core = ON** and **msmq_LocalStorage = ON** entries. The applicable subcomponents will be installed automatically.

**msmq_RoutingSupport**

The entry **msmq_RoutingSupport = ON** calls for installation of the Routing Support subcomponent, which creates an MSMQ Setting object in Active Directory and provides routing services. This key can be used only when you are installing Message Queuing on computers running on platforms from the Windows Server 2003 family. The Routing Support subcomponent can only be installed on computers that have access to Active Directory. Routing Support cannot be installed without the Core, Local Storage, and Active Directory Integration subcomponents. Thus, if you are installing Routing Support, there is no need to include **msmq_Core = ON**, **msmq_LocalStorage = ON**, and **msmq_ADIntegrated = ON** entries. The applicable subcomponents will be installed automatically. Routing Support can be uninstalled by including an **msmq_RoutingSupport = OFF** entry.

**msmq_MQDSService**

Include a **msmq_MQDSService = ON** entry if you need to install the Message Queuing directory service on a domain controller to provide access to Message Queuing objects in Active Directory for downlevel clients. Setting this key equal to **ON** is equivalent to selecting the Downlevel Client Support subcomponent during the installation of Message Queuing on a domain controller. This key is relevant only to platforms from the Windows Server 2003 family. To install Downlevel Client Support, Message Queuing must be integrated with Active Directory.

The keys appearing in entries in the [msmq] section are as follows:

**ControllerServer**

This key specifies the NetBIOS or DNS name of the MSMQ 1.0 site controller hosting a Message Queuing server that the computer will try to contact during Setup. Specifying a Windows 2000 domain controller or a Windows Server 2003 domain controller will generate an error. This key is applicable only to an independent client or a routing server being installed on a domain computer and only when **msmq_ADIntegrated** is set equal to ON. It is optional, but if it is not included in the answer file, Setup will try to find a suitable domain controller automatically through AD.lib. If an entry is specified, unattended setup will use the entry and will not attempt to find the computer automatically. If it fails, however, unattended setup will not continue.

Do not include an entry with this key if you are installing Message Queuing on a domain controller.

**SupportingServer**

This key is set equal to the NetBIOS or DNS name of a Message Queuing server that will be used to support the dependent client being installed. This entry is required when you are creating a dependent client.

**SupportLocalAccountsOrNT4**

A **SupportLocalAccountsOrNT4 = TRUE** entry provides support for computers running MSMQ 1.0 and computers running Message Queuing 2.0 that are logged on using a local user account. This entry enables weakened security for Active Directory on a domain controller, which is then replicated to all other domain controllers in every domain in your forest. Thus, it must be included only for the first instance of Message Queuing with Downlevel Client Support on a domain controller in your forest, and Downlevel Client Support must be installed on all domain controllers. Then, computers requiring this support will be able

to query Active Directory and view the properties of Message Queuing objects. Otherwise, such computers will not be able to query Active Directory about Message Queuing objects. This key is optional, and if it is set to **FALSE**, weakened security is removed.

For the security implications of weakening Active Directory security, see Enabling weakened security.

### Notes

- Dependent clients cannot run under a local user account.

- Any computer that sends queries about Message Queuing objects to Active Directory on a domain controller directly, rather than through the Message Queuing directory service, will not be able to access Active Directory when it logs on using a local user account even if weakened security for Active Directory is enabled.

### Site

Include a line with this key if you are installing a Message Queuing routing server. Setup may not be able to automatically resolve the proper site if subnets in your network are not configured correctly. In such a case, this entry becomes mandatory. Set the key equal to the name of the default Active Directory site containing the domain controller with which this server will communicate.

#### Sample answer file

You can use the following sample answer file as a template. Use a text editor to appropriately edit the entries.

### Note

- Unattended setup ignores any file entries that begin with a semicolon (**;**). You can use these to place comments in your answer file.

```
[Version]
Signature = "$Windows NT$"

[Global]
FreshMode = Custom
MaintenanceMode = RemoveAll
UpgradeMode = UpgradeOnly

[Components]
msmq_Core = ON | OFF
msmq_LocalStorage = ON | OFF
msmq_ADIntegrated = ON | OFF
msmq_TriggersService = ON | OFF
msmq_HTTPSupport = ON | OFF
msmq_RoutingSupport = ON | OFF
msmq_MQDSService = ON | OFF

[Msmq]
ControllerServer = MyDomainController
SupportingServer = MyDomainControllerOrRoutingServer
SupportLocalAccountsOrNT4 = FALSE | TRUE
Site = DefaultFirstSiteName
```

Be sure to test your unattended setup script before deployment.

# Uninstalling Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Uninstalling Message Queuing**
To uninstall Message Queuing software from a computer, it is recommended that you first be granted the Delete MSMQ Configuration Object permission. If you do not have this permission, you can still uninstall Message Queuing, but this object will not be deleted from Active Directory. For more information on installation permissions, see Installation permissions.

**Running Message Queuing Setup**
You uninstall Message Queuing using **Add or Remove Programs** in Control Panel. When you uninstall Message Queuing from a computer, Message Queuing Setup removes all of the queues (including private queues) and messages for that computer from Active Directory and from the local computer. However, no Message Queuing files are removed from the \Windows\System32 and \Windows\System32\Drivers folders.

When you uninstall a Message Queuing server, all the applicable computer objects (the MSMQ Configuration and MSMQ Settings objects) are also deleted, and the name of the server is removed from the message routing lists of clients using that server.

For information on how to uninstall Message Queuing software, see Uninstall Message Queuing.

You can also uninstall certain individual Message Queuing subcomponents by running **Add or Remove Programs** and clearing the check boxes for the applicable subcomponents in the Windows Components Wizard. For more information on the Message Queuing subcomponents, see Choosing installation subcomponents.

**Removing Message Queuing objects from Active Directory**
If you cannot uninstall Message Queuing software from a computer, such as during a hardware malfunction or if you no longer have physical access to a computer, you can remove Message Queuing objects for the computer from Active Directory. When you remove the Message Queuing objects from Active Directory, all queues and messages on the local computer are also removed.

For more information on Message Queuing objects, where they are created, and where they are located in Active Directory, see Message Queuing and Active Directory.

For information on how to remove objects for Message Queuing in Active Directory, see Remove Message Queuing objects from Active Directory.

### Note

- You must not delete the MsmqServices object unless you are removing Message Queuing from every computer in your organization. This object is the "enterprise" object for Message Queuing.

# Upgrading to Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrading to Message Queuing**

This section explains the requirements for upgrading earlier versions of Message Queuing software (MSMQ 1.0 and MSMQ 2.0) on computers that are being upgraded to Windows Server 2003 family operating systems from Windows NT Server 4.0 or Windows 2000 Server. For details, see the following links:

- For an overview of upgrading, see Upgrade overview.
- For information on upgrading on clusters, see Upgrading Message Queuing in a server cluster.

# Upgrade overview

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrade overview**
This topic provides an overview of upgrading from MSMQ 1.0 on Windows NT Server 4.0 computers, and from Message Queuing 2.0, which runs on Windows 2000, to Message Queuing 3.0.

**Upgrading from MSMQ 2.0 on Windows 2000 computers**
MSMQ 2.0, running on a Windows 2000 computer will automatically be updated to Message Queuing 3.0 during the upgrade to a Windows Server 2003 family operating system.

**Upgrading from MSMQ 1.0 on Windows NT 4.0 computers**
For computers running Windows NT Server 4.0, the Message Queuing software is also automatically upgraded to Message Queuing 3.0, with the exception of MSMQ 1.0 controller servers. Message Queuing is upgraded for controller servers using the Message Queuing migration tool. The upgrade entails upgrading MSMQ 1.0 controller servers, and migrating all the information stored in the MQIS database to Active Directory. After the upgrade, PECs will be upgraded to a Windows Server 2003 family operating system, promoted to domain controllers, and running Message Queuing 3.0. Other MSMQ 1.0 controller servers will be upgraded to Windows Server 2003 family computers, running as Message Queuing 3.0 routing servers.

Note that *mixed-mode operation* is not permitted in Message Queuing 3.0. In other words, a configuration of MSMQ 1.0 controller servers running on Windows NT 4.0 computers, together with Message Queuing domain controllers in a Windows Server 2003 domain environment, is not supported.

For more information on this upgrade scenario, see the white paper at Microsoft Message Queuing(http://www.microsoft.com/).

# Upgrading Message Queuing in a server cluster

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Upgrading Message Queuing in a server cluster**
This topic applies only to Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition users who need to upgrade an existing server cluster that is running an earlier version of Message Queuing software. For more information on how Message Queuing functions in a server cluster, see Distributing Message Queuing in a network.

**Upgrading from Windows 2000 on a cluster containing Message Queuing resources**
When upgrading Message Queuing resources from a Windows 2000 cluster, do the following:

- Confirm that your computer hardware configuration is compatible with Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.
- Take all Message Queuing resources in the cluster offline using Cluster Administrator.
- Upgrade all of the cluster nodes to Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. Before upgrading, use Cluster Administrator to remove resources that are not supported by these operating systems.
- Bring the Network Name resources back online, ensure that the Distributed Transaction Coordinator (DTC) is online, and then bring the Message Queuing resources online. Note that DTC is required if you want to use external transactions, triggers, or HTTP messaging with Message Queuing applications on your cluster.

For instructions on managing cluster resources, see Manage Resources, in the Server Clusters Help file.

Note that during a rolling upgrade of a cluster, Windows Server 2003 family computers do not support failover of the Message Queuing service from Windows Server 2003, Standard Edition nodes to Windows 2000 nodes.

For more information on creating and configuring Message Queuing Triggers resources in a virtual server, see Configure trigger resources for server clusters.

**Upgrading MSMQ 1.0 in a server cluster**
You can also upgrade from Windows NT Server 4.0 to Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition on a cluster that contains DTC and Message Queuing resources. There is a Knowledge Base (KB) article available with detailed instructions for this upgrade scenario. To search for this KB article, see Searching the Knowledge Base (http://www.microsoft.com/)

# Administering a Message Queuing Network

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering a Message Queuing Network**
This section provides conceptual information about common tasks for administering a Message Queuing network. Instructions for completing individual tasks can be found in the How To... section in this online Help file. For details about administering Message Queuing networks, see the following links:

- For information on moving Windows Server 2003 family computers between sites, see Moving mobile computers between sites.
- To learn about messaging between different kinds of messaging systems, see Configuring cross-platform messaging.
- For information on adding and removing sites from a server list, see Adding or removing sites from a server list.
- For information on routing links, see Configuring routing links.
- To learn about configuring your company's firewall to allow Message Queuing messages through, see Configuring firewalls for Message Queuing.
- To learn about configuring DNS settings so that Message Queuing can resolve computer names correctly, see Configuring Message Queuing with a proxy server.

# Moving mobile computers between sites

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Moving mobile computers between sites**

Windows Server 2003 family computers can be moved to other Windows Server 2003 family sites. When you move a computer to a new site, Message Queuing automatically detects the domain controllers in its new site and uses them for information services. It then updates the directory service with the new site in which the computer resides. After the new site is established, all messages that were sent to the computer while it was disconnected are automatically rerouted from the previous site to the new site. Whenever the computer connects at the new site, its messages are routed directly to it at that site. If the computer is a Message Queuing routing server, it automatically recovers its full functionality in the new site.

A Windows Server 2003 family computer that uses an MSMQ 1.0 site controller as an information server is not capable of automatically revealing its new site. In this case, you can provide the name of the new site either before disconnecting the computer from the network or after reconnecting to the network. Selecting a new site enables your computer to connect to the closest controller server and thus optimizes access to the MQIS database. If you specify the new site before disconnecting from the network, messages sent to the computer while it is disconnected are routed directly to the new site while the computer is en route. If the computer's new site is not specified, Message Queuing will attempt to use a site controller in its former site. Any message sent to that computer will be routed to its previous site. This applies only to Message Queuing independent clients and Message Queuing servers without routing enabled. Message Queuing servers that have routing enabled (provide routing services to clients) are not supported in an environment of MSMQ 1.0 site controllers.

For information on how to specify a new site for mobile computers, see Specify a new site for mobile computers.

# Configuring cross-platform messaging

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring cross-platform messaging**

For cross-platform messaging to communicate with foreign computers, you must first create a foreign site, add the foreign computers to the foreign site, and create queues on the foreign computers. All of these entities are Message Queuing objects in Active Directory. Then you must create a routing link between the foreign site and a Windows Server 2003 site that contains a Message Queuing server running a connector application, define this Message Queuing server as a site gate for the routing link, and add the foreign site to the list of sites where this computer resides.

Note that for Windows Server 2003 family computers running a connector application, after adding the foreign site to the list of sites, the foreign site will not appear in the list until you restart the Message Queuing service on the domain controller of the domain to which the Message Queuing server running the connector application belongs.

Foreign computers can be assigned to the same foreign site if the computers are all serviced by the same connector application. Connector applications run on Message Queuing servers and enable foreign computers to communicate with other Message Queuing computers in your network. A common Message Queuing connector application is Microsoft MSMQ-MQSeries Bridge. When a Message Queuing computer sends a message to a foreign computer, the message is routed to the Message Queuing server running the connector application in the foreign site because this server is also a site gate in the routing link connecting to the foreign site. When the server identifies the message destination as a queue on a foreign computer, it places it in a connector queue. The message is subsequently processed by the connector application, which communicates with the foreign computer.

You must create a routing link between the foreign site and each Windows Server 2003 family or Windows 2000 site that will communicate with the foreign site. The Message Queuing routing server that runs the connector application in the foreign site must also function as a site gate in each routing link connecting to the foreign site. Site gates are always associated with a particular routing link. Note that if the routing server running the connector application, and functioning as a site gate, is unavailable, site gates in other sites cannot route messages to the applicable connector queue for transmission to computers in the foreign site.

For information on how to create a foreign site, see Create foreign sites.

For information on how to add foreign computers to a foreign site and how to add a Message Queuing server running a connector application to a foreign site, see Add foreign computers to a site.

For information on how to create and configure a routing link, see Routing links.

# Adding or removing sites from a server list

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Adding and removing sites from a server list**

In most cases, Message Queuing determines the sites in which servers reside automatically. Message Queuing also offers procedures for manually adding and removing sites from the list of sites where a Message Queuing server with routing enabled, which is also known as a routing server, resides. Any routing server that resides in a specific site can be configured as an in-routing or out-routing server for computers that reside in that site. Such routing servers can be also configured as site gates in routing links to that site. A routing server can reside in multiple sites, and it can be configured as an in-routing server, an out-routing server, or a site gate independently in each site in which it resides.

For information on how to add sites to the list of sites where a Message Queuing server with routing enabled resides, see Add sites to a Message Queuing server's list of sites.

For information on how to remove sites from the list of sites where a Message Queuing server with routing enabled resides, see Remove sites from a Message Queuing server's list of sites.

To learn more about sites and subnets, see Sites overview, in the Active Directory Help file.

# Configuring routing links

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring routing links**

Message Queuing uses *routing links* to route messages between sites. When a new Windows Server 2003 site is created, you must also create a corresponding routing link to guarantee that messages can always be sent to other sites. However, you will not be prompted to create a routing link when a new site is created. Create multiple routing links to provide load balancing and redundancy. For more information, see Routing links.

For general information on routing, see Message Queuing routing overview. For instructions on creating a routing link, see Create routing links.

By default, users in the Enterprise Admins group and Domain Admins group can create or modify routing links. For information on how to set permissions for routing links, see Set permissions for routing link objects.

You can change routing-link costs where alternative routes are possible and you want to enforce one route over another route using link costs. You can assign costs that reflect the differences in bandwidth or communications costs between the different links. Costs can range from 1 to 999999.

For example, two sites in one city, called A and B, are connected by a high-speed link, and an overseas site, called C, is connected to site B by a low-speed link. In this case, you can define the cost between A and B as 1, and the cost between B and C as 2. For information on how to change routing-link costs, see Change the cost for a routing link.

A routing link may use Message Queuing servers with routing enabled, called *site gates*, to route messages between two sites. A routing link can have any number of site gates or no site gate, and each site gate can be in either or both of the sites connected by it. If a routing link is associated with a pair of site gates, one in each site, only those servers are used to transfer messages on the routing link on behalf of clients. It is recommended that all sites in your organization be interconnected using routing links.

After a routing link has been created, you can double-click on the link and change the link's description.

For information on how to change the Message Queuing servers used for a routing link, see Change the site gates for a routing link.

# Configuring firewalls for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring firewalls for Message Queuing**
The following sections describe configurations for messaging through firewalls.

**HTTP messaging through firewalls**
The use of HTTP as a transport enables MSMQ to leverage existing firewall support, without the need for Message Queuing-specific firewall configurations, and it is recommended that HTTP be used for messaging through firewalls. For more information on HTTP messaging, see HTTP/HTTPS messages.

**Messaging through firewalls using the Message Queuing protocol**
To send messages across firewalls using the native Message Queuing protocol, specific configuration settings are required to allow Message Queuing computers located on either side of the firewall to communicate with each other. Note that for best security practice, it is recommended that HTTP messaging be used as a solution for messaging through firewalls. These specific configuration settings include allowing Message Queuing computers located on an external network to be able to connect to Message Queuing computers on your internal network over the Internet. This is achieved by opening the following service ports on your firewall.

| | |
|---|---|
| TCP | 1801 |
| RPC | 135, 2101, 2103, 2105 |
| UDP | 3527, 1801 |

**Important**

- If Message Queuing computers located on an external network require access to Active Directory, it is strongly recommended that you set up a virtual private network (VPN) connection with the Point-to-Point Tunneling Protocol (PPTP). Otherwise, your network security could be compromised. For more information on VPN connections and PPTP, see the Point-to-Point Tunneling Protocol (PPTP) in the Windows Help file.

**Notes**

- With direct format names, only the TCP port needs to be opened. This provides for two-way sending, but only local reading. If you need to read from remote queues, or if you need to query Active Directory for information regarding public queues or other Message Queuing objects, you must also open the RPC ports.

- RPC ports 2103 and 2105 can be incremented by 11 if the initial choice of an RPC port is in use when the Message Queuing service starts. Message Queuing queries RPC port 135 to discover the 2*xxx*-series ports.

- Before establishing a connection with another Message Queuing computer, Message Queuing sends a ping message on UDP port 3527 to verify that the other computer responds, and only then does it attempt to establish a connection. If you do not want to open this port on the firewall, you can disable this feature by setting registry key value **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\msmq\parameters\UsePing** to zero. This might incur some performance cost as ping verification is disabled. After setting the registry, restart the Message Queuing service for the changes to take effect.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

The following firewall configurations allow you three different operating modes for remote Message Queuing computers located on external networks.

**Sending messages only**
For remote clients to be able to send messages, you must allow these clients access to TCP port 1801 on your firewall.

In this configuration, remote clients cannot access Active Directory nor will messages be routed by Message Queuing servers through your internal network. This means that remote clients must be able to directly connect to the destination computer on your internal network over this port. Remote clients can then send messages using a direct format name for the destination queue.

This does not apply to dependent clients.

**Sending messages with Active Directory access**
In addition to opening TCP port 1801, allowing access to RPC ports 135 and 2101 permits remote clients access to Active Directory. RPC port 135 is used for handshaking between a remote client and a Message Queuing server. Message Queuing servers also use RPC port 2101 for communicating with each other.

Assuming that multicast network packets are allowed to pass through your firewall, allowing access to UDP port 1801 also permits remote clients to send a broadcast message to automatically determine their site.

This does not apply to dependent clients.

**Sending and receiving messages with Active Directory access**
In addition to opening TCP port 1801 and RPC ports 135 and 2101, allowing access to RPC ports 2103 and 2105 also permits remote clients to access queues and retrieve messages on Message Queuing computers located on your internal network.

Access to RPC ports 2103 and 2105 also allows remote dependent clients to send and receive messages using their supporting server on your network.

# Configuring Message Queuing with a proxy server

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring Message Queuing with a proxy server**
Most organizations connect to the Internet using a proxy server for dispatching outgoing messages and a perimeter network protected by firewalls for routing incoming messages. A proxy server acts as a liaison between the intranet and the Internet. It creates a protected channel for passing information between the private network and the Internet and provides caching. Message Queuing uses an HTTP proxy to connect to URL addresses for delivering HTTP messages. Message Queuing thus needs to know the name of the proxy server and the applicable port to which it must connect. At the same time, there may be a list of target names for which the proxy connection is bypassed; that is, URL addresses that do not need to be accessed through the proxy server. An attempt is made to establish a direct connection with such targets. A Message Queuing application can only send messages over HTTP through a proxy server. No other type of Message Queuing messages can be sent. All of the information regarding proxy servers is stored in registry keys, at **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\Internet Settings**. After setting the registry, restart the Message Queuing service for the changes to take effect.

**Caution**

- Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.

You can use the Proxy Configuration utility to configure a proxy server. To configure a proxy server, type the following command at the command prompt:

```
Proxycfg.exe -p ProxyServerName:PortNumber BypassedTargets
```

Note that wildcards can be used to define a list of bypassed targets.

The proxy server configuration is read whenever an outgoing queue is created to send messages to a new URL, so you do not need to stop and restart the Message Queuing service for new proxy settings to take affect. However, if you change the proxy server configuration using the Proxy Configuration utility and you want an active outgoing queue to use the new settings, you must pause and resume that outgoing queue.

# Administering Message Queuing in Active Directory

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Message Queuing in Active Directory**

This section provides information about Message Queuing objects found in Active Directory, and their location. For details, see the following links:

- For information about each Message Queuing object in Active Directory, see Message Queuing objects.
- For details of where each Message Queuing object is located in the Active Directory snap-ins, see Message Queuing object locations.

# Message Queuing objects

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing objects**
This topic provides an overview of all Message Queuing objects in Active Directory. For more detailed information on Message Queuing objects, see the Active Directory topic in the Microsoft Platform Software Development Kit (SDK). To open the Active Directory Start Page in the Platform SDK section of the MSDN Library, in **Contents**, click **Active Directory**.

**MsmqServices object**
An MsmqServices object is created when the first domain controller of your forest is configured. There can be only one MsmqServices object in a forest. The MsmqServices object is of the MSMQ-Enterprise class type. The MsmqServices object contains organization-wide configuration information for Message Queuing.

**Routing link object**
A routing link object with a user-defined name is created for each routing link that you create between two Message Queuing sites and contains connectivity information. Routing link objects are of the MSMQ-Routing-Link class type.

**MSMQ Settings object**
An MSMQ Settings object is created for each Message Queuing server that is installed on a Windows Server 2003 family domain controller and for each Message Queuing server with routing services. The MSMQ Settings object contains configuration information, such as what Message Queuing-specific services the server provides. The MSMQ Settings object is of the MSMQ-Settings class type.

**Message Queuing object**
An msmq object is created for each independent client and for each Message Queuing server, and it contains configuration information for the local computer. The msmq object is of the MSMQ-Configuration class type. The msmq object is the parent object of the queue objects under it. The msmq object is also known as the MSMQ Configuration object, for example, in names of specific permissions for this object and in the Delegation of Control Wizard.

**Queue object**
A queue object with a user-defined name is created for each queue and contains information associated with the particular queue. Queue objects are of the MSMQ-Queue class type. Each queue object contains a Queue messages folder and a Journal messages folder.

**Distribution list object**
Distribution lists are essentially distribution group objects in Active Directory. They can contain explicit references only to destinations published in Active Directory; that is, to public queues, queue aliases, and other distribution lists, but not to private and URL-named queues. Queue aliases are published in Active Directory and can point to any queue, including private and URL-named queues. Thus, queue aliases pointing to private and URL-named queues can be included in a distribution list.

**Queue alias object**
A queue alias object with a user-defined name is created for each queue alias and contains information associated with the particular queue alias, such as the string comprising the alias and the format name to which the alias points. Queue alias objects are of the MSMQ-Custom-Recipient class type. For more information on queue aliases, see Queue aliases.

**Distribution group object**
Distribution lists are essentially distribution group objects in Active Directory. They can contain explicit references only to destinations published in Active Directory; that is, to public queues, queue aliases, and other distribution lists, but not to private and URL-named queues. Queue aliases are published in Active Directory and can point to any queue, including private and URL-named queues. Thus, queue aliases pointing to private and URL-named queues can be included in a distribution list.

A distribution group object with a user-defined name is created for each distribution list published in Active Directory and contains information associated with the particular distribution list, such as the elements comprising the distribution list, which can be public queues, queue aliases, and other distribution lists. Distribution group objects are of the Group class type. For more information on distribution lists, see Distribution lists.

# Message Queuing object locations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Message Queuing object locations**

This table describes where Message Queuing objects are located in Active Directory.

| Object name | Location |
|---|---|
| MsmqServices object | The MsmqServices object is displayed in the console tree (usually the left pane) of the Active Directory Sites and Services snap-in in the following location:<br>• Active Directory Sites and Services<br>• Services<br>• MsmqServices |
| Routing link object | Routing link objects are located under the MsmqServices object in the details pane of the Active Directory Sites and Services snap-in:<br>• Active Directory Sites and Services<br>• Services<br>• MsmqServices<br>• Routing link object |
| MSMQ Settings object | The MSMQ Settings object is displayed in the details pane (usually the right pane) of the Active Directory Sites and Services snap-in when the applicable server object is selected. Server objects are found in the following location:<br>• Active Directory Sites and Services<br>• Sites<br>• *YourSiteObject*<br>• Servers<br>• *YourServerObject* |

| | |
|---|---|
| msmq object | The msmq object and child container objects are displayed in the console tree of the Active Directory Users and Computers snap-in in the following location:<br>• Active Directory Users and Computers<br>• *YourDomainObject*<br>• *YourOrganizationalUnitObject* (such as Computers or Domain Controllers)<br>• *YourComputerObject*<br>• **msmq**<br>• *YourQueueFolder* (such as Private Queues and System Queues)<br>• *YourQueueObject*<br>• Queue messages Journal messages |
| Queue object | Queue objects are located under the msmq object of the computer on which they reside (except for dependent clients, whose Queue objects are located under the msmq object of the supporting server). |
| Queue alias object | Queue alias objects are displayed in the console tree of the Active Directory Users and Computers snap-in in the following location:<br>• Active Directory Users and Computers<br>• *YourDomainObject*<br>• *YourContainer* (such as Computers or Domain Controllers)<br>• *YourQueueAliasObject* |
| Distribution group object | Distribution group objects are displayed in the console tree of the Active Directory Users and Computers snap-in in the following location:<br>• Active Directory Users and Computers<br>• *YourDomainObject*<br>• Users<br>• *YourDistributionGroupObject* |

# Administering Message Queuing Computers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Message Queuing Computers**
This section provides conceptual information about common administration tasks for Message Queuing computers. Instructions for completing individual tasks can be found in the How To... section in this online Help file. For details about administering Message Queueing computers, see the following links:

- To learn about administering server cluster nodes, see Administering server cluster nodes.
- For information on assigning Message Queuing servers as in-routing and out-routing servers, see Assigning in-routing and out-routing servers.
- For information on joining workgroup computers to a domain, see Joining a domain from a workgroup.
- For information on configuring DNS settings so that Message Queuing can resolve computer names, see Configuring DNS settings for Message Queuing.

# Administering server cluster nodes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering server cluster nodes**

You administer Message Queuing running in a server cluster node the same way as you would a standard Message Queuing server, using the same snap-ins.

In addition, you can also locally administer Message Queuing running in the context of a virtual server in a cluster by using the Computer Management snap-in. A virtual server is a cluster group that contains a Disk resource and a Network Name resource. A virtual server functions similarly to a server running on a physical computer, so a Message Queuing server running in the context of a virtual server provides services similar to the services provided by a Message Queuing server running on a physical computer.

To administer Message Queuing running in the context of a virtual server, the associated node on which the virtual server is running must be online and available. For information on how to create a Generic Application resource for managing a Message Queuing cluster resource from a local Computer Management snap-in, see Configure Message Queuing resources for server clusters.

# Assigning in-routing and out-routing servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Assigning in-routing and out-routing servers**

Message Queuing servers, called in-routing servers or out-routing servers, can be used to route messages within a site. Each independent client (or Message Queuing server without routing enabled) can be configured to use up to three in-routing servers and out-routing servers to route messages on its behalf.

It is recommended that in-routing and out-routing servers be used for mobile clients that are mostly disconnected from the network or connect to the network through a remote access server.

The same Message Queuing server can serve as both an in-routing server and out-routing server. However, all such servers must be located in the original site of the client (the site the client was installed in). When a mobile client moves to another site, these in-routing and out-routing servers are not used for message routing until the client returns to its original site.

For redundancy, it is recommended that you configure clients to use more than one in-routing or out-routing server.

For information on how to specify or change the in-routing or out-routing servers that a client can use, see Change in-routing or out-routing servers.

# Joining a domain from a workgroup

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Joining a domain from a workgroup**

Message Queuing can be used in an environment that is not part of a Windows Server 2003 domain, such as a computer that is part of a Windows Server 2003 family operating system workgroup, or a Windows 2000 workgroup. If Message Queuing is run on a computer that is part of a workgroup, that computer can later join a domain. Conversely, a Message Queuing computer can be part of a domain, and then join a workgroup. Such a computer can then rejoin the same or a new domain at a later time. A Message Queuing computer running on Windows Server 2003 family operating systems can join any new Windows Server 2003 family or Windows 2000 domain. For a Message Queuing computer running on a Windows Server 2003 family computer to be able to join a new domain, a domain controller must be available in its site. It is also advisable that the domain controller be configured as a global catalog server. In addition, Message Queuing computers can join only Windows Server 2003 and Windows 2000 domains. For information on how to join a domain from a workgroup, see Join a domain from a workgroup.

In Windows Server 2003, directory service functionality is configured by the user. When a Message Queuing computer has been configured with directory service functionality, this functionality will be provided whenever that computer is promoted to a domain controller.

**Caution**

- When you join a Message Queuing computer in workgroup mode to a domain, or move a Message Queuing computer to a different forest, all messages in outgoing transactional queues on the computer will be discarded, or placed in the dead-letter queue if negative source journaling is enabled.

**Joining a Windows NT 4.0 domain**

When a computer running Message Queuing in workgroup mode joins a Windows NT 4.0 domain, the installation subcomponent for Active Directory Integration is removed. After the computer has been joined to the domain and rebooted, open the Control Panel, double-click **Add or Remove Programs**, click **Add/Remove Windows Components**, and in the Windows Components Wizard, select **Application Server**, click **Details**, select **Message Queuing**, click **Details**, and select **Active Directory integration**.

For information on the operating restrictions of running Message Queuing in a nondomain environment, see Deploying in workgroup mode.

# Configuring DNS settings for Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Configuring DNS settings for Message Queuing**

If you are operating in an environment that requires the resolution of DNS domain names, such as an environment in which queries are sent to Active Directory, the DNS settings (for example, the DNS server addresses) must be configured properly on each computer for Message Queuing to resolve computer names correctly.

It is recommended that you verify that each computer is set with a fully qualified DNS name. Otherwise, some Message Queuing operations may fail. This property is typically set automatically by the Net Logon service on each computer in a Windows Server 2003 domain. To check this property, open Active Directory Users and Computers. In the console tree, right-click the applicable computer object, click **Properties**, and on the **General** page, inspect **DNS Name**.

For information on how to configure DNS properties for Message Queuing, see Set DNS properties for Message Queuing.

# Administering Message Queuing Security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Message Queuing Security**

This section provides background information about security tasks for Message Queuing. Instructions for completing individual tasks can be found in the How To... section in this online Help file. For details about administering security, see the following links:

- For background information about the Message Queuing service and the Local System account, see Managing the Message Queuing service account.
- To learn about weakened security for Active Directory, and when to enable it, see Enabling weakened security.
- For information about message authentication, see Managing message authentication.
- To learn about authentication for Message Queuing using user certificates, see Managing user certificates.
- For information about server authentication, see Managing server authentication.
- For details of how HTTP/HTTPS messages are authenticated, see Authenticating HTTP/HTTPS messages.
- For information about IIS security, see Resolving IIS Anonymous Access settings.

# Managing the Message Queuing service account

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing the Message Queuing service account**
Services run in their own security context. To establish a security context, a service needs to log on to the system. The Message Queuing service runs by default under the LocalSystem account, and you can change this default account if necessary.

The Message Queuing Downlevel Client Support service is installed on Message Queuing servers that provide Active Directory Integration to Windows 95, Windows 98, and Windows NT 4.0 computers running MSMQ 1.0, and Windows 2000 computers running MSMQ 2.0. This service must run under the default LocalSystem account.

The following are features of the LocalSystem account:

- The LocalSystem account is a member of the local Administrators group.
- The LocalSystem account already has the required rights to run as a service.
- The LocalSystem account already has the Generate Security Audits privilege.

The default LocalSystem account is a member of the Authenticated Users group, and is granted the Read permission for all Message Queuing objects in Active Directory. This permission allows these users to read the attributes of Message Queuing objects in Active Directory, but does not control the rights of users to receive (and remove) messages from queues. Change the account that the Message Queuing service runs under only if you do not want to grant the Read permission to all members of the Authenticated Users group. In this case, to run the Message Queuing service under a specific user account, you need to do the following:

- Grant the Read permission for all Message Queuing objects in Active Directory to the user account.
- Add the user account to the local Administrators group.

If you must provide greater functionality you may consider granting the user account the following privileges:

- Log on as a service
- Act as part of the operating system
- Replace a process level token
- Increase quotas
- Generate security audits

In addition, if you use a local user account to run Message Queuing applications on Windows 2000 independent clients, such clients cannot authenticate Message Queuing servers that run on domain controllers. Such domain controllers provide access to Active Directory. If authenticating those Message Queuing servers is a requirement, you must use a domain user account instead.

For information on how to change the account the Message Queuing service runs under, see Change the Message Queuing service account.

**Triggers service security**
By default, the Message Queuing Triggers service runs under the Network Service account. Services that are running under this account interact with the network using the credentials of the computer account, and this account has the same level of access to resources and objects as members of the Users group accounts. For heightened security, you can change the default Triggers service account to a more secure setting; for example, the Local Service account, which limits access to the local computer. You can also specify a less secure account for the Triggers service; for example, the LocalSystem account, allowing

access to the entire domain. The account used must be able to access the queues with which triggers are associated. Specifically, the account must be granted the **Log On As A Service** permission as well as the **Peek Message** permission for all applicable queues. If the message processing type of a trigger is "retrieval" or "transactional retrieval," the user must be granted the **Receive Message** permission for that queue. If you change the security settings for a queue to enable a trigger to monitor it, you must stop and start Triggers to update the service's configuration before the trigger will start to work. Note that executable files invoked by triggers are invoked with the same security privileges as those assigned to the Triggers service.

# Enabling weakened security

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling weakened security**

Message Queuing servers running on Windows Server 2003 domain controllers, or Windows 2000 domain controllers can operate using weakened security for Active Directory. If you install Message Queuing with the Message Queuing Downlevel Client Support component on a domain controller, a dialog will appear, asking you if you want to enable or disable weakened security. Enable weakened security if any of the following operating configurations apply to your organization:

- A domain environment where users running MSMQ 1.0 (on Windows NT 4.0, Windows 98, or Windows 95) access the Message Queuing Downlevel Client Support service on Windows Server 2003 domain controllers, and Windows 2000 domain controllers for directory services. This applies even if such users are logged on with Windows Server 2003 or Windows 2000 domain accounts.

- A Windows Server 2003 family domain, or Windows 2000 domain, where users running Message Queuing (on Windows Server 2003 family operating systems, Windows XP, or Windows 2000) are logged on with Windows NT 4.0 domain accounts.

- An environment where Message Queuing 2.0 users are logged on with local user accounts. Note that:

    - Dependent clients cannot run under a local user account.

    - Any computer that sends queries about Message Queuing objects to Active Directory on a domain controller directly, rather than through a Message Queuing server, will not be able to access Message Queuing objects in Active Directory when it logs on with a local user account, even if the security for Active Directory is weakened.

If weakened security is enabled, a Message Queuing server running on a Windows Server 2003 domain controller or a Windows 2000 domain controller will access Active Directory within its own security context. This access will be under the LocalSystem account on the applicable domain controller. You can view the properties of Active Directory objects, including msmq (MSMQ configuration) objects and queue objects. You can also check security settings for operating systems requiring the Downlevel Client Support service or Windows XP clients. You cannot bypass any security restrictions on creating objects or setting object security.

For users running MSMQ 1.0, when the MSMQ service starts an RPC call to a Message Queuing server running on a Windows Server 2003 domain controller, the call is impersonated as an anonymous logon in the default security context. To allow such an anonymous user access to Active Directory, domain security is weakened by not impersonating this call. Consequently, all queries for Message Queuing objects in Active Directory are accepted by Windows Server 2003 Message Queuing servers. This means only that the properties of Message Queuing objects can be viewed. It does not mean that messages can be retrieved (removed) from public queues.

You can also check how Message Queuing is configured and modify this setting after installation using ADSI Edit. For information on how to do this, see Weaken security using ADSI Edit.

It is also possible to support MSMQ 1.0 users (and other configurations discussed) without weakening security. In this case, you must grant the Everyone group the List Content permission on all computer objects in each domain. This, however, greatly compromises the domain security.

# Managing message authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing message authentication**

*Message authentication*, which confirms that a message has not been tampered with and can be used to verify the identity of the sender of a message to the receiver, is implemented using certificates. Messages are authenticated asynchronously without the sender and the receiver communicating with one another.

The use of message authentication is optional. Message Queuing applications are not required to use authentication when sending or receiving messages. It should be noted that the use of message authentication slows the transfer of messages. When used, Message Queuing applications that send or receive authenticated messages must use the cryptographic service provider installed by default with Windows Server 2003 family operating systems, and which is needed for both message authentication and encryption. For more information on message encryption, see Encryption for Message Queuing.

For more information on authentication for Windows Server 2003 family, see Authentication, in Windows help.

For more information on certificates for Windows Server 2003 family, see Certificates overview in Windows help.

In addition to confirming that no tampering has been done to a message, message authentication can be used to verify the identity of the user who sent the message. Authentication is done at the request of the sending application, which indicates that a message is to be authenticated by setting its authentication level.

An authenticated message is sent with a *digital signature* and a *certificate*. Digital signatures are created and validated using public-key (asymmetric) encryption. In Message Queuing, the sender uses the private key to create a digital signature, and the receiver uses the public key in the certificate attached to the message to validate the digital signature.

Because Message Queuing is an asynchronous messaging system, the source computer and the destination computer are not assumed to be online and connected at the same time.

The sender of a message cannot always be authenticated. For example, before you can authenticate the sender of a message sent between computers in the same forest with access to Active Directory, a user certificate must be registered in Active Directory. Registering a certificate in Active Directory associates the certificate with the owner's security ID (SID). You can also use the internal user certificate created on the local computer during setup, which is automatically registered the first time you log on to the computer.

Then, when an authenticated message is sent, a Message Queuing application on the source computer creates a digital signature for the message, attaches the digital signature, the SID, and the applicable certificate to the message, and sends the message. The Message Queuing service on the destination computer validates the digital signature and then queries the applicable directory service to see whether the SID/certificate pair is registered. If the SID/certificate pair is registered, the message is considered authenticated and is placed in the destination queue.

When the sending application indicates that a message is to be authenticated by setting its authentication level, the Message Queuing runtime code does the following:

1. Retrieves the applicable certificate. By default the internal certificate is retrieved. However an external certificate is used if one is provided on the sender certificate property. For more information, see Managing user certificates.

2. Obtains the private signing key for the certificate used. For Windows Server 2003 family computers, certificates are stored in certificate stores. Certificates can also be found in a public certificate store, but Message Queuing will not accept these. In the case of external certificates, Message Queuing searches for the private key in the personal certificate store (the external certificate must be registered there). For internal certificates, Message Queuing locates the private key internally.

3. Computes a hash value from a fixed limited portion of the message using a hash algorithm. The default hash algorithm is the Secure Hash Algorithm (SHA).

4. Encrypts the hash value using private key obtained in step 2 to create a digital signature.

5. Attaches the digital signature, the security ID (SID), the name of the hash algorithm, and a digital certificate containing the

public cryptographic key for the source computer to the message.

On the destination computer, the Message Queuing service does the following:

1. Computes a hash value from a fixed limited portion of the message using the same hash algorithm as the source computer.

2. Decrypts the digital signature attached to the message using the public cryptographic key from the digital certificate to obtain a second hash value.

3. Compares the two hash values to confirm that the message was not tampered with during delivery.

4. If the values match, locates the user certificate in Active Directory and verifies that the user who registered the certificate matches the sender (the SID) listed on the message. If the hash values do not match, the message is discarded, and a negative acknowledgment is returned to the sending application, if such an acknowledgment was requested.

In this case, the digital signature uniquely identifies the user who sent the message as well as the message itself. Because no user can digitally sign a message with another user's identity, no user can deny having sent a message if it contains a validated digital signature that can be identified with a specific sender.

You can also send authenticated messages to connector applications. The intermediary Message Queuing server (running the connector application) can validate authenticated messages.

You can specify whether or not a queue accepts only authenticated messages. For more information, see Allow only authenticated messages on queues.

**Authentication in workgroup mode**
When authenticating messages when the receiving computer is in workgroup mode, you cannot work with user certificates registered in Active Directory to map the certificate to the respective user security ID (SID). Instead, do the following:

On the sender computer:

- The user gets an external certificate from a CA that is trusted by both the sender and the receiver, particularly the receiver.

- Set Message Queuing to use this certificate when sending the message (use **MQGetSecurityContextEx**).

- Message Queuing finds the private key and attaches the signature, but does not add the SID (use **MQMSG_SENDERID_TYPE_NONE**).

On the receiver computer, the procedure is the same as that described above for an Active Directory computer, except that:

- The Message Queuing service does not attempt to locate the certificate in Active Directory.

- The user must verify the certificate by verifying the certificate chain (use **CertGetCertificateChain** and **CertVerifyCertificateChainPolicy**).

For more information on internal certificates, see Managing user certificates.

For more information on cryptographic keys used by Message Queuing, see "Cryptographic keys" in Encryption for Message Queuing.

# Managing user certificates

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing user certificates**

There are two types of user certificates, or public-key certificates, that Message Queuing can use: *internal certificates* and *external certificates*. An internal certificate is automatically created during Message Queuing Setup for the user installing Message Queuing. External user certificates contain information that is supplied by a certification authority (CA). Unless stated otherwise, the information in this section applies to both types of user certificates.

You manage user certificates for a Message Queuing computer using the **Message Queuing Properties** dialog box in Computer Management. For information on how to view all user certificates that are registered in Active Directory, see View certificates for Message Queuing. All certificates have a validity period. After the end of the validity period, the certificate is no longer considered an acceptable or usable credential. Also, if Message Queuing is installed on a computer operating in a workgroup, and then the computer joins a domain, you must manually renew the internal certificate. For information on how to renew internal user certificates, see Renew certificates for Message Queuing.

User certificates can be stored in Active Directory. If you reinstall Windows Server 2003 family on a computer (rather than upgrade an existing installation), you must delete the user's previous (now obsolete) certificate. For information on how to delete a user certificate, see Remove certificates for Message Queuing. Next, you must register the new certificate manually. For information on how to register a user certificate, see Register certificates for Message Queuing.

**Notes**

- If permissions are used to control access to queues, user certificates must be registered in Active Directory so that Message Queuing can map the certificate to the respective SID of the user. Although it is possible to send unauthenticated messages to queues that restrict access, doing so could compromise security as the content of such messages could be tampered with or read while en route.

- In the case of a cluster, a user profile, which is needed to store a user certificate in the registry, will not be created for a user account that is used only for a cluster service until it logs onto a node of the cluster at least once. After a user profile is created, the user must manually register a user certificate for Message Queuing.

- Active Directory sets a multi-valued attribute limit of approximately 800 user certificates for a specific user account. This limit is usually exceeded when obsolete user certificates have not been deleted from Active Directory. If multiple certificates exist for a user account, only the latest is used, and obsolete certificates can be deleted. For instructions, see Remove certificates for Message Queuing.

**External certificates**

External user certificates contain information that is supplied by a certification authority (CA), rather than the SID, for verifying a sender's identity. The information in an external certificate is guaranteed by the CA that created the certificate. The Message Queuing service on the destination computer still validates the digital signature when a message is received but does not validate the external certificate itself. It is the responsibility of the receiving application to validate an external certificate before acting on the message. For more information on certification authorities, see Certification Authorities.

An external certificate can be registered in Active Directory and then used to authenticate the sender's SID when authenticated messages are sent to other computers in the same forest, precisely as in the case of an internal certificate, but with a possibly higher level of security.

If you register an external certificate in Active Directory and thereby associate an SID with the certificate, you can send messages to queues that restrict access. This is because the Message Queuing service verifies the sender's SID. Registering an external certificate thus provides an additional, although optional, measure of authentication, which can be used by the receiving application.

External certificates are required for sending authenticated messages to environments other than Windows Server 2003 family or Windows 2000. Other reasons you may want to use an external certificate include the following:

- You want to control who can send authenticated messages by controlling the use of certificates. Internal certificates are created automatically during the installation of Message Queuing and can be used by everyone.

- You want to send messages to someone who does not have access to Active Directory.
- You belong to a workgroup, as opposed to a domain.

In cases where SIDs cannot be used to authenticate the sender, as in the case of a message sent between computers operating in workgroup mode, an external certificate can be used for message authentication. In such cases, when an authenticated message is sent, the Message Queuing service on the source computer creates a digital signature for the message, attaches the digital signature and the applicable external certificate to the message, and sends the message. The Message Queuing service on the destination computer validates the digital signature.

More specifically, on the source computer, the Message Queuing service does the following:

1. Computes a hash value from a fixed limited portion of the message using a hash algorithm.
2. Encrypts the hash value using the sender's private cryptographic key to create a digital signature.
3. Attaches the digital signature, the name of the hash algorithm, and a digital certificate containing the sender's public cryptographic key to the message.

On the destination computer, the Message Queuing service does the following:

1. Computes a hash value from a fixed limited portion of the message using the same hash algorithm as the source computer.
2. Decrypts the digital signature attached to the message using the public cryptographic key from the digital certificate to obtain a second hash value.
3. Compares the two hash values to confirm that the message was not tampered with during delivery.
4. If the values match, validates the external certificate (this step is optional).

One way to obtain an external certificate is to use Internet Explorer 5.0 or later to request a certificate from a certification authority. For example, you can contact one of the following representative CA vendors:

- AT&T
- InternetMCI Mail
- Keywitness Canada
- VeriSign Commercial Software Publishers

In general, you can obtain different levels of external certificates. Some external certificates do not prove your identity. They prove only that you are the same person each time you send a message. This is sufficient for subscription-based applications. Other levels of external certificates prove that you are who you claim to be.

# Managing server authentication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing server authentication**

Server authentication is used by clients to authenticate queries sent to domain controllers and Message Queuing servers. A client can verify that the query results have not been tampered with and that the results were returned by the correct domain controller or Message Queuing server. The method used for server authentication depends on the version of Message Queuing (and hence the operating system) running on the server and the client. The following table shows the method used for each situation.

| Client | Server | Server authentication method |
|---|---|---|
| Message Queuing client running on Windows Server 2003 family | Windows Server 2003 or Windows 2000 domain controller | Kerberos V5 |
| Message Queuing client running on Windows 2000 | Message Queuing server running on Windows Server 2003 family or Windows 2000 | Kerberos V5 |
| Message Queuing client on Windows 2000 | MSMQ 1.0 controller server running on Windows NT 4.0 | Server certificate |
| MSMQ 1.0 client running on Windows NT 4.0, Windows 98, or Windows 95 | Message Queuing server running on Windows Server 2003 family or Windows 2000 | Server certificate |

**Using server certificates**

Server certificates are used to authenticate some Message Queuing servers. You need to create and register a server certificate for the following:

- Message Queuing servers that reply to Active Directory queries from MSMQ 1.0 clients running on Windows NT 4.0, Windows 98, or Windows 95.
- MSMQ 1.0 controller servers running on Windows NT 4.0 that reply to queries from Message Queuing clients on Windows 2000 computers.

Note that Message Queuing clients running on Windows Server 2003 family or Windows 2000 use Kerberos V5 to authenticate Windows Server 2003 family or Windows 2000 Servers, and do not require server certificates.

For information on how to create a server certificate, see Create server certificates for Message Queuing. For information on how to change an existing server certificate, see Change server certificates for Message Queuing.

Server certificates used to authenticate Message Queuing servers must be installed separately from any user (client) certificates that may be installed for message authentication. However, if properly configured, the same certificate can be used by Message Queuing to authenticate servers and by Internet Information Services (IIS) for secure Web (HTTPS) communication.

# Authenticating HTTP/HTTPS messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Authenticating HTTP/HTTPS messages**

When messages sent over HTTP/HTTPS transport are authenticated, the sender attaches an XML digital signature and a user certificate, which contains the user's name, the user's public key, and the signature of the issuer. A user can obtain an external user certificate from a certification authority (CA), such as VeriSign or Microsoft Certificate Services, or use the internal user certificate that is created and registered on the local computer during setup and then first time that a user logs on to the computer. If an external user certificate is used, the user must register that certificate in Active Directory in the local forest using the Message Queuing UI.

When messages are sent over HTTPS transport, additional security requirements must be satisfied before an SSL session is established between the sender and the recipient. For more information, see HTTPS authentication.

When a message requiring authentication is sent within the same forest, the recipient decrypts the digital signature, and tries to locate the user certificate in Active Directory according to its GUID. If the user certificate is not found in Active Directory, Message Queuing tests the certificate to ascertain whether it is self-signed (an internal user certificate). If it is found in Active Directory, the message is inserted into the queue and marked as authenticated. However, if the user certificate is not found in Active Directory and it is not self-signed, the sender is regarded as an anonymous user and is allowed to access the destination queue accordingly.

When messages are sent over HTTP/HTTPS transport between different forests, the recipient cannot search for the certificate in Active Directory in the sender's forest. To overcome this problem and use the same user certificate to authenticate messages sent between different forests, perform the following steps:

1. Export the sender's user certificate to a file.
2. Send the file to the recipient's forest.
3. In the recipient's forest, create a pseudo-user to represent the sender or the sender's group.
4. Import the certificate to the pseudo-user's personal certificate store.
5. Log on as the pseudo-user.
6. Use the Message Queuing UI to register the certificate in Active Directory to the pseudo-user.
7. Grant the pseudo-user permissions to write to Message Queuing queues.

The recipient can now authenticate messages sent between forests by querying Active Directory in the local forest as in the case of messages sent over HTTP/HTTPS transport within a single forest.

For information about how to register a user certificate, see Register certificates for Message Queuing.

# Resolving IIS Anonymous Access settings

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Resolving IIS Anonymous Access settings**

By default, Internet Information Services (IIS) impersonates a special IIS local user account. Since this account is local to the IIS computer (unless it is a domain controller), it cannot be authenticated by other computers in the network and is treated as an anonymous user.

Active Server Page (ASP) applications and scripts run under IIS, and by default impersonate the IIS local user for any Message Queuing operations, including queries to Active Directory. Since anonymous users no longer belong to the Everyone group in Windows Server 2003 family, these queries fail, and ASP applications and scripts cannot locate, create, and delete queues. In addition, an anonymous user cannot open a queue for remote read, and you cannot read messages from queues that do not belong to the IIS computer using ASP.

You can resolve this issue in either of the following ways:

- **Using a domain user account instead of the IIS local user account for Anonymous Access**. The account must have the required permissions for the Message Queuing operation, for example access to Active Directory.

- **Disable Anonymous Access**. Trust the IIS computer for delegation to allow multiple hops in communicating with a domain controller. Note that this option will not resolve remote read limitations.

For instruction, see Change the security settings for Internet Information Services.

# Administering Queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Queues**

This section provides information about managing and configuring queues for Message Queuing. You can use the Computer Management snap-in or the Active Directory Users and Computers snap-in to administer queues. For more information, see Administration overview. Instructions for completing individual tasks can be found in the How To... section in this online Help file. For details about administering queues, see:

- Managing queues
- Managing queue properties
- Setting queue quotas
- Enabling target journaling

# Managing queues

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing queues**
The Computer Management snap-in or the Active Directory Users and Computers snap-in is used to manage and administer queues. For more information, see Administration overview.

**Creating queues**
There may be certain restrictions on creating queues on computers that are determined during Message Queuing Setup. For more information, see Access control for Message Queuing.

For information on how to create, delete, or find a queue, see Create queues, Delete queues, and Find queues.

**Viewing queues**
The operation of viewing the contents of a queue from the same computer on which the queue is located is called *local* reading. The operation of accessing the contents of a queue from a computer other than the one on which the queue is located is called *remote* reading.

You can view the contents of only those queues that you have sufficient permissions to open. Therefore, there can be public queues whose contents are not visible in the snap-ins. For more information, see Access control for Message Queuing.

By default, any user with the *Get Properties* permission for the msmq object can view the queue status of local outgoing queues. To limit this default, you can do either of the following:

- Change users' *Get Properties* permissions for the msmq object.
- In the registry, add the dword **RestrictAdminAPI** with a registry value of 1, under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Msmq\Parameters. This restricts the ability to view the queue status of active outgoing queues to local administrators, regardless of security permissions for the msmq object.

  **Caution**

  - Incorrectly editing the registry may severely damage your system. It is recommended that you back up any valuable data on the computer before making changes to the registry.
- After setting the registry, restart the Message Queuing service for the changes to take effect.

**Viewing internal private queues and outgoing queues**
There are certain internal private queues that are used exclusively for internal operation, and the Active Directory Users and Computers and Computer Management snap-ins do not display any internal private queues. For more information, see Internal private queues and outgoing queues.

Outgoing queues on a computer are used to store messages temporarily before the messages are sent to their destinations. Outgoing queues are visible only when using the Computer Management snap-in. However, transactional messages in outgoing queues cannot be viewed from Computer Management.

**Pausing and resuming queues**
You can pause and then later resume the operation of a local outgoing queue, thereby deferring the sending of messages from that queue to the corresponding destination queue. This way, you can send all pending messages to a destination queue at a later time. For information on how to pause or resume an outgoing queue, see Pause or resume an outgoing queue.

You can also take Message Queuing offline from the network. When you take Message Queuing offline from the network, you pause the operation of all outgoing queues on the local computer. This is useful when you communicate over a slow, expensive connection and want to defer sending messages in all outgoing queues on the computer to a later time. For information on how to take Message Queuing offline from the network for a computer, see Take outgoing queues offline.

**Managing queues with triggers**
When you delete a queue, the triggers associated with that queue are not deleted. If you later re-create a queue of the same name, any existing triggers will be used by this queue. However, this re-created queue will not be monitored by the Message

Queuing Triggers service until the computer is restarted, or the associated triggers disabled and then enabled again. For instructions on setting the enabled or disabled property for a trigger, see View or change properties of triggers and rules for triggers.

# Managing queue properties

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing queue properties**

The Computer Management snap-in or the Active Directory Users and Computers snap-in is used to administer queues. For more information, see Administration overview.

Messages are sent to public queues based on the sum of each message's priority value and the destination queue's priority, or base priority, value. Messages with the highest numerical total are sent first. Messages are sent first by base priority, and then by message priority. Private queues do not support base priority, so messages are sent only according to message priority.

For information on how to change the base priority for a queue, see Change the base priority of queues.

The queue type ID specifies the type of service provided by the queue. The queue type ID must be a valid GUID in the form {*dddddddd-dddd-dddd-dddd-dddddddddddd*}, where *d* is a hexadecimal text digit. It can be used to search for queues of a specific type. After you have located queues by type ID, you can use the name retrieved to open or delete the queue. The use of the queue type ID is optional.

For information on how to change the queue type ID, see Change queue type IDs.

Each queue can have a label, which is a string with a maximum length of 124 Unicode characters that can be used to identify the queue. For example, queues that accept messages sent by a specific application may all have the label "Messages from my application." For information on how to change the queue label, see Change queue labels.

You can specify whether a queue accepts encrypted messages, unencrypted messages, or both by selecting the privacy level for the queue. If the queue privacy level is set to **None**, the queue accepts only unencrypted messages. If the queue privacy level is set to **Body**, the queue accepts only encrypted messages. If the queue privacy level is set to **Optional**, the queue accepts both unencrypted and encrypted messages. You can check whether a message sent to the queue was encrypted or unencrypted by using Active Directory Users and Computers (or programmatically by checking the PROPID_M_PRIV_LEVEL message property).

For information on how to set the privacy level for a queue, see Change the privacy level for queues.

You can also specify whether a queue accepts only authenticated messages. For more information on how to set this option, see Allow only authenticated messages on queues. For general information on message encryption, see Encryption for Message Queuing.

# Setting queue quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting queue quotas**

The Computer Management snap-in or the Active Directory Users and Computers snap-in is used to administer queues. For more information, see Administration overview.

Queue quotas specify the cumulative limit for messages stored in a particular queue. When the message quota for a destination queue is reached, new incoming messages are discarded, and a negative acknowledgment message indicating that the destination queue is full is generated. Messages can be placed in the queue only after one or more messages are removed (read) from the queue.

Queue quotas can be set only for Message Queuing servers and independent clients. Queue quotas set on Message Queuing servers that support dependent clients apply both to the server and to all its dependent clients.

For background information on quotas, see "Quotas for messages" in Message storage.

For information on how to set a queue quota for normal messages, see Set the message storage size for queues.

For information on how to set a queue quota for journal messages, see Set the journal storage size for queues.

# Enabling target journaling

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Enabling target journaling**

Target journaling is configured on a queue basis. If it is enabled, a copy of each incoming message is placed in the applicable queue journal when a message is removed (read) from the destination queue.

For information on how to enable target journaling for a queue, see Enable or disable target journaling.

# Administering Distribution Lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Distribution Lists**

A distribution list is a set of one or more distribution lists, or target queues, which can be used to do the following:

- Send transactional and nontransactional messages to multiple destinations.
- Specify one or more administration queues for sending application-generated acknowledgment messages, when acknowledgment messages are requested.
- Specify one or more response queues for sending application-generated response messages, when response messages are requested.

For details about administering distribution lists, see the following links:

- To learn about creating distribution lists, and setting properties and security permissions, see Creating and managing distribution lists.
- To learn about managing queues, queue aliases, and other distribution list members of a distribution list, see, Managing distribution list elements.

# Creating and managing distribution lists

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating and managing distribution lists**

Distribution lists are public lists that are published as distribution group objects in Active Directory. You can use Active Directory Users and Computers to create and delete distribution lists as distribution group objects, and to view their properties. For more information on performing these tasks, see Create distribution lists, Delete distribution lists, and View, add, or delete elements in a distribution list. For more conceptual information on distribution lists, see Distribution lists.

When you create a distribution list (group) in Active Directory, you are prompted to specify the minimal information needed to establish the object in Active Directory. Each distribution list (group) has many more properties (attributes) that can be viewed and set, as follows.

| Property | Details |
|---|---|
| Group name | Name of distribution list object |
| Description | Description of distribution list object |
| Group scope | Distribution lists can be domain-scoped, global, or universal |
| Group type | The type of group for this object is *distribution* |
| Distribution list (group) elements (members) | Elements that make up the distribution group |
| Folders containing the members in Active Directory | Location of group elements |
| Distribution lists (groups) of which the applicable distribution list (group) is a member | If the distribution list is in turn the member of another distribution list |
| Administrative permissions | Security for the distribution list |

**Setting distribution list scope**

Distribution lists can be specified as "Global", "Domain-scope" or "Universal":

- Domain-scope distribution lists can only contain queues and distribution lists from the local domain as members, and cannot be moved to another domain.
- Distribution lists with global scope can include queues and distribution lists from the local domain, and queues and global distribution lists from other domains in the forest. They can be moved to another domain.
- Distribution lists with universal scope can include other groups and accounts from any domain in the domain tree or forest and can be granted permissions in any domain in the domain tree or forest.

**Setting distribution list security**

Because distribution lists are objects in Active Directory, security attributes can be attached to control access to them. A security descriptor is assigned to each queue and distribution list. This security descriptor lists the users and groups that are granted or denied access to the queue or distribution list, and the following specific permissions are associated with distribution lists:

- **Send To** permission allows a group or user to send messages to the applicable distribution list. By default, this permission is granted to everyone.
- **Add/Remove Self as Member** permission for a public queue or queue alias allows a group or user to add or remove the applicable public queue or queue alias as a member in a distribution list. All users are granted this permission by default.

- **Add/Remove Self as Member** permission for a distribution list allows a group or user to add or remove the applicable distribution list as a member in another distribution list. Only the owner of a distribution list has the Add/Remove Self as Member permission for it.

For more information on distribution list security, see Distribution lists.

Distribution lists can be referenced programmatically through distribution list format names. For more information on using distribution list format names and their syntax, refer to the Message Queuing Software Development Kit (SDK).

# Managing distribution list elements

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing distribution list elements**

Distribution lists are public lists that are published as distribution group objects in Active Directory. You can use Active Directory Users and Computers to add and delete members of a distribution list. For more information, see View, add, or delete elements in a distribution list. For more conceptual information on distribution lists, see Distribution lists.

The explicit members of a distribution list must be destinations published in Active Directory, including:

- Public queues. For more information on public queues, see Public and private queues.
- Queue aliases. For more information on queue aliases, see Queue aliases.
- Other distribution lists.

You cannot explicitly use private queues or URL-named queues as members of your distribution lists, but they can be included in a distribution list, and referenced indirectly by using queue aliases. For instructions, see Create queue aliases.

You can also identify all distribution lists for a given element, and remove a distribution list, queue, or queue alias from one or more distribution lists if required. For instructions, see Get a list of distribution lists containing a given element.

# Administering Messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Messages**

This section provides information about configuring and managing messages. Instructions for completing individual tasks can be found in the How To... section in this online Help file. For details about administering messages, see the following links:

- For information on message lookup and purging messages, see Managing messages.
- For information on sending test messages and tracking message routes, see Tracking messages and routes.
- To learn more about backing up messages, see Backing up and restoring messages.
- To learn about setting queue quotas and computer quotas for messages, see Setting message quotas.
- For information about storing messages on multiple disks, see Storing message files on multiple disks.
- To learn about sending transactional messages to multiple destinations, see Sending transactional messages to multiple destinations.

# Managing messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Managing messages**

A message body can contain text or binary information. If a message body contains text information, this information is displayed. If a message body contains binary information, hexadecimal data is displayed along with its text translation. Message bodies can be viewed on the **Body** dialog box for the message.

To retrieve messages from a queue, you must first be granted the Receive Messages permission for the queue. It is not sufficient to have the Read permission granted for the queue. For information on how to set permissions for queues, see Access control for Message Queuing.

**Message lookup**

Applications can be written to peek at or retrieve a specific message without navigating through the queue to find the message. This functionality is based on a 64-bit lookup identifier that is assigned to each message when it is placed in a destination queue. The lookup identifier of a message is unique within the queue and is always greater than the lookup identifier of the preceding message and less than the lookup identifier of the next message. Note that this functionality is designed to help you look at specific messages you have received. It is not a random search mechanism.

Both messages sent by applications and messages sent by Message Queuing have a lookup identifier. The local queue manager generates the lookup identifier and sets the lookup identifier properties of the message (**PROPID_M_LOOKUPID** and **MSMQMessage.LookupId**) whenever it places a message in a queue. Message Queuing generates a lookup identifier for messages placed in user- and application-generated queues, such as destination queues, administration queues, response queues, and report queues, as well as for copies of messages placed in system-generated journals, dead-letter queues, and connector queues. If a message is sent to several destination queues, or if a copy of a message is stored in a computer journal or queue journal, each copy of the message will be assigned its own lookup identifier when it is placed in its respective queue.

Message Queuing has an API function and methods of the **MSMQQueue** object that enable applications to perform the following tasks programmatically without using cursors:

- Peeking at the first message in the queue.
- Peeking at the message preceding the message specified by the message lookup identifier supplied.
- Peeking at the message specified by the message lookup identifier supplied.
- Peeking at the message immediately following the message specified by the message lookup identifier supplied.
- Peeking at the last message in the queue.
- Retrieving the first message in the queue.
- Retrieving the message preceding the message specified by the message lookup identifier supplied.
- Retrieving the message specified by the message lookup identifier supplied.
- Retrieving the message immediately following the message specified by the message lookup identifier supplied.
- Retrieving the last message in the queue.

When API function calls are made to peek at or retrieve messages, messages can be processed asynchronously, synchronously, or within a transaction. However, the method calls process the messages synchronously.

A receiving application can use the lookup identifier of a message to go directly to that message and then continue navigating the queue from that point toward the front or the end of the queue. Lookup identifiers differ from cursors in this respect. When cursors are used, the application must start at the front of the queue and navigate towards the end of the queue.

**Purging messages**

You can purge all messages in any queue, including outgoing queues. The only exception to this is system queues. Note that you cannot purge individual messages in a queue. For instructions on purging all messages in a queue, see Purge all messages in a queue.

# Tracking messages and routes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Tracking messages and routes**
You can track the passage of messages from a source computer to a target queue by sending test messages. You can also track the route taken in your network by test messages or all messages sent from the source computer to test queues on remote computers by enabling the message route tracking feature. Neither of these features is enabled by default following Message Queuing installation. For instructions on enabling these features, see Enable route tracking and test messages.

**Sending a test message**
You can send a test message from an independent client or a Message Queuing server to a test queue on any other independent client or Message Queuing server within the same Windows forest. To do so, after you have enabled test messages, you must first create a test queue, or change the type ID of an existing queue, so that you can send test messages to the test queue on the destination computer. For more information, see Sending test messages.

Report messages, which track the routes of messages through your network, are always sent to the report queue specified by the source computer for test messages if message tracking is enabled on the source computer.

For information on how to send a test message, see Send test messages.

**Message route tracking**
You can use the message route tracking feature to track the route that messages take through your network. When all the conditions for route tracking are satisfied, a report message is sent to the report queue for the source computer when the message leaves its source computer, when it arrives at its destination computer, and whenever the message arrives at or leaves a Message Queuing server with routing enabled.

After message route tracking is enabled, there are two basic requirements for its use. First, you must enable message route tracking and create or specify a report queue for the source computer. Second, the route tracking flag (property) in the message must be set programmatically by the sending application. In addition, only messages sent to a remote computer are tracked, and messages sent over HTTP/HTTPS transport as well as messages sent to multiple destinations using IP multicast, distribution lists, or multiple-element format names cannot be tracked in the current version of Message Queuing.

For information on message route tracking, see Track message routing. For information on setting properties programmatically, refer to the Message Queuing Software Development Kit (SDK).

Only one queue can be specified in the header of a message as the report queue for that message. If a route-tracked message with a report queue specified in its header subsequently passes through a Message Queuing server with routing enabled that also has message route tracking enabled, Message Queuing sends a conflict report message to the report queue for that server when the message leaves it.

For example, consider three computers, called A, B, and C, where message route tracking is enabled for computer A. In this example, the Message Queuing application on computer A sends a message to computer C through out-routing server B. If the report queue for computer A is called QA, a report message is sent to QA when the following events occur:

- The message leaves computer A.
- The message arrives at out-routing server B.
- The message leaves out-routing server B.
- The message arrives at computer C.

If out-routing server B is also configured with message route tracking enabled, with report messages being sent to QB, this queue receives a conflict report message when the message leaves out-routing server B. The original source queue QA, however, continues to receive the regular tracking messages.

In a second example, computer A does not have message route tracking enabled, but route tracking is enabled on computers B and C. In this case, the message starts reporting its movements when it leaves out-routing server B. A total of two report messages are sent to QB (when the message leaves computer B and when the message arrives at computer C), but no conflict

report messages are sent.

In a third example, neither computer A nor computer B has message route tracking enabled, but route tracking is enabled on computer C. In this case, no report messages are sent before the message reaches the destination computer, but one report message is sent to QC when the message arrives at its destination. No conflict report messages are sent.

Report messages consist of a label and a body. The label has the following format:

```
gggg:dddd:hh sent from <computer> to <address> at <time>
```

```
gggg:dddd:hh received by <computer> to <address> at <time>
```

where *gggg* is the first four hexadecimal digits of the globally unique identifier (GUID) of the source queue, *dddd* is the internal message identifier, and *hh* is the hop count.

The internal message identifier is the message sequence number, which is part of the entire message identifier. The message identifier is not used in its entirety because it is too big to be easily read. It is a 16-byte source GUID and a 4-byte sequence number.

When conflict report messages are sent, the labels contain the string "Report Message Conflict."

The body of a regular report message contains the ID of the message, the target queue, the destination of the next hop (if any), and the hop count in an XML format. For example, the body of the report message sent to QA when the message leaves the out-routing server in the first example will have the following form:

```
<MESSAGE ID>message_ID</MESSAGE ID>
```

```
<TARGET QUEUE>format_name_of_target_queue_on_computer_C</TARGET QUEUE>
```

```
<NEXT HOP>IP_address_of_computer_C</NEXT HOP>
```

```
<HOP COUNT>1</HOP COUNT>
```

The body of a conflict report message contains has an additional line containing the original source queue. For example, the body of the "Report Message Conflict" message sent to QB when the report message just described is sent to QA in the first example will have the following form:

```
<ORIGINAL QUEUE>format_name_of_QA</ORIGINAL QUEUE>
```

```
<MESSAGE ID>message_ID</MESSAGE ID>
```

```
<TARGET QUEUE>format_name_of_target_queue_on_computer_C</TARGET QUEUE>
```

```
<NEXT HOP>IP_address_of_computer_C</NEXT HOP>
```

# Backing up and restoring messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

You can use the message backup and restore command-line utility, Mqbkup, provided with Message Queuing to protect message storage files, log files, transaction log files, and registry settings against computer failure.

**Backing up and restoring messages**

You must first create a folder to use for backup before using the utility. You run the utility from a command prompt by typing **mqbkup**.

The syntax of the utility is:

**mqbkup {-b** *folderpath* **| -r** *folderpath* **| -y | -c** *resource name* **| -?}**

where:

**-b** *folderpath* backs up the files to the specified folder,

**-r** *folderpath* restores the files from the specified folder,

**-y** don't prompt (Yes to all questions).

**-c** cluster resource; please specify the MSMQ resource name.

**-?** displays help for the utility.

Before restoring messages, you must close Control Panel and all Message Queuing applications.

Certificate and cryptographic key information is not backed up using this utility.

| **Warning** |
| --- |
| When you run mqbkup, run wbadmin to backup any public queue information. |

After a restore operation, you must register a new user certificate. For information on how to register a certificate, see Register certificates for Message Queuing.

This utility stops and restarts the Message Queuing service.

# Setting message quotas

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Setting message quotas**
Message storage limits can be set by configuring the following:

**Set computer quota**
This quota specifies the cumulative limit for all messages stored in destination queues, outgoing queues, and queue journals on a computer. After a computer quota is reached, no messages can be delivered to any queue or journal on the computer. The computer quotas can be set for Message Queuing servers and independent clients. Following a new installation of Message Queuing, or an upgrade from Windows 2000, the default limit for these storage sizes is 8 gigabytes (GB). For more information, see Message storage. Computer quotas are set separately to limit message storage size and journal size. When you set the journal storage size limit for a computer, you automatically set equal quotas for the computer journal (the Journal messages queue), the dead-letter queue, and the transactional dead-letter queue. For instructions on setting the message storage limit for a computer, see Set the message storage size for computers. For instructions on setting the journal storage size for computers, see Set the journal storage size for computers.

**Set queue quota**
This quota specifies the cumulative storage size (number of bytes) available for all messages stored in a particular queue, whereas a computer quota limits the total storage available for all messages in destination queues, outgoing queues, and queue journals on the computer. When a queue quota is reached, messages can no longer be sent to the queue until one or more messages are retrieved (removed from the queue by a Message Queuing application). Queue quotas can be set for all queues. Queue quotas for journal queues and other queues are set separately. For instructions on setting the quotas for regular messages, see Set the message storage size for queues. For instructions on setting the journal storage size for a queue, see Set the journal storage size for queues.

# Storing message files on multiple disks

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Storing message files on multiple disks**

Independent clients and Message Queuing servers can improve their messaging performance by storing the various files used to implement messaging across several physical disks. Message Queuing can use three different disks to implement message functions. For the best messaging performance on Message Queuing servers, put the messaging (*.mq) files, message log files, and transaction log files on separate physical disks.

For information on how to store messaging files on multiple disks, see Store messages on multiple disk drives.

# Sending transactional messages to multiple destinations

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sending transactional messages to multiple destinations**
Both distribution lists and multiple-element format names have support for transactional messaging. When a message is sent to a distribution list or a multiple-element format name within a transaction, a copy of the message is sent to each applicable outgoing queue. Failure in sending the message to any of these outgoing queues will abort the entire transaction. A distribution list or a multiple-element format name may reference a mixture of transactional and nontransactional target queues. When a transactional message is sent to such a mixture of targets, the message is allowed to be sent to all the transactional and nontransactional target queues on the local or a remote computer. However, when the message reaches a nontransactional queue, it is rejected and transferred to the transactional dead-letter queue on the computer hosting the nontransactional queue for storage, and a negative acknowledgement message indicating that a transactional message was sent to a nontransactional queue is sent to the administration queue specified in the original message.

**Exactly-once and in-order delivery**
Message Queuing guarantees that transactional messages will be delivery in-order and only once. The Exactly-Once In-Order Delivery (EOD) protocol is an end-to-end protocol, guaranteeing that messages are not duplicated or delivered out-of-order. To achieve this, the sender attaches information to each message that allows the receiver to eliminate duplicate and out-of-order messages. All messages that have order relation between them are said to be in the same *message stream*. Each message in a stream carries a sequence number, identifying the message in the stream, and a previous number identifying the preceding message in the same message stream. The first message in a stream always carries a previous number equal to zero.

Each receiver filters the newly arriving messages and accepts only those messages that arrive in-order and are not duplicates. To implement this, the receiver uses the sequence numbers and previous numbers attached to the messages. The receiver needs to keep track of the sequence number of the last message accepted in each message stream, there is no need to keep message history or a larger database. The receiver acknowledges by sending the sequence number of the last message accepted back to the sender. This lets the sender know that the receiver accepted all messages up to the acknowledged message.

**Sending messages**
The sender tags each message with a monotonically increasing sequence number and links it to the previous message using the previous number. When there is no previous message in the message stream the sender places a zero in the previous number field. This ensures that the receiver will accept the message. A message that is discarded at the sender side (for example, when the time to receive the message before time-out has elapsed) is removed from the message stream, and the next message is linked to the message preceding the removed message.

The sender needs to keep resending the same messages until an "order-acknowledgment" arrives from the receiver. The acknowledgment carries a sequence number indicating that all messages up to that number have safely arrived in-order and exactly-once at the target. The sender can thus safely discard all messages up-to and including the indicated sequence number.

To overcome crash scenarios where the receiver cannot recover its database (last accepted message sequence number for a given message stream), the sender must assign the first message in the queue a previous number of zero. On the other hand, after a sequence number has been attached to a message, it must never be changed.

**Receiving Messages**
The receiver keeps track of the last accepted message sequence number for each message stream. It uses this number to filter out arriving messages and accept only those that arrive in-order and are not duplicates. The receiver initiates this last number with a zero for each new message stream. The simplest algorithm the receiver must implement looks like the following:

```
if(message.sequence <= last_accepted_sequence)
{
    Reject Message
}

if(message.previous > last_accepted_sequence)
{
    Reject Message
}
```

```
Accept Message

last_accepted_sequence = message.sequence;
```

When a new message stream has been created, the receiver occasionally sends an "order-acknowledgment" containing the last accepted sequence number. This acknowledgment signals the sender that the target has accepted all messages in the message stream up to the indicated sequence number. The receiver keeps sending these acknowledgments for active message streams; that is, message streams for which messages continue to arrive even if these messages are not accepted.

The receiver sends these "order-acknowledgments" to the acknowledgment return address indicated by the incoming message. All messages within the same message stream must have the same acknowledgment return address.

It is recommended that the quality of service for acknowledgment messages be as lightweight as possible because the receiver maintains the state to regenerate the acknowledgment persistently (last sequence number).

For more information on sending messages within transactions, see Transactional messaging.

For information on sending transactional messages to distribution lists, see Distribution lists.

For information on sending transactional messages to multiple-element format names, see Multiple-element format names.

# Administering Triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Administering Triggers**

This section provides conceptual information about administering Message Queuing triggers. Instructions for completing individual tasks can be found in the How To... section in this online Help file. For details about Message Queuing triggers, see the following links:

- To learn about defining message processing types for triggers, see Defining message processing types.
- For information on how trigger messages are processed, see Processing trigger messages.
- To learn about creating triggers, see Creating triggers.
- To learn about default trigger values, see Modifying default trigger values.
- For an example, see Implementing triggers.

# Defining message processing types

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Defining message processing types**

A message processing type is defined for each trigger. The available message processing types are peeking, retrieval, and transactional retrieval.

When peeking is selected as the message processing type for a given trigger and the Message Queuing Triggers service is running, it peeks at each message arriving at the monitored queue, tests it for the conditions of all the rules associated with the queue, and performs all the applicable actions, but does not remove it from the queue.

When retrieval is selected as the message processing type, the Message Queuing Triggers service peeks at each message arriving at the monitored queue, tests it for the conditions of all the rules associated with the queue, and invokes all the applicable actions. If at least one condition is satisfied, and no invoked action fails, the message is removed from the queue.

The transactional retrieval type of message processing is similar to the retrieval type, but in this case the message is returned to the queue if any of the operations involved in retrieving the message cannot be completed and the retrieval transaction is aborted.

When messages arrive at a monitored queue, Message Queuing does not limit access to these messages to the Message Queuing Triggers service. Thus, it is possible for another Message Queuing application to open the queue and remove messages before the Message Queuing Triggers service manages to process those messages.

# Processing trigger messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Processing trigger messages**

The Message Queuing Triggers service starts to process the messages in the queue as soon as the service starts running. The service starts at the beginning of the queue and processes all the messages in the queue. After processing all the existing messages, the service waits at the end of the queue for a new message. If the service is stopped and restarted, it will process all messages remaining in the queue again.

**Processing messages having the same priority**

When all the messages sent to a monitored queue have the same priority, new messages are always placed at the end of the queue. After processing all the existing messages in the queue, the Message Queuing Triggers service will activate the trigger each time a new message arrives at the end of the queue. The illustration below shows a queue with four messages and portrays a trigger with the peeking type of message processing. The Message Queuing Triggers service already peeked at message 1, message 2, message 3, and message 4 and will continue to peek at messages as they arrive at the end of the queue.




**Processing messages having different priorities**

When messages with different priorities arrive in a queue, a new message can be placed anywhere in the queue. For example, when a high-priority message is placed somewhere in the queue, the Message Queuing Triggers service processes the new message next regardless of where it is placed in the queue and then continues to process all the low-priority messages that are in the queue.

In the next illustration, which also portrays a trigger with the peeking type of message processing, the Message Queuing Triggers service makes one pass through the queue peeking at message 1, message 2, message 3, and message 4. Then, when message 5, a new message of priority 1, arrives at the queue, it is placed between message 2 and message 3. The service goes back and peeks at message 5 and then continues to peek at message 3 and message 4 again.




Note that in this case, the trigger is activated twice for message 3 and twice for message 4. You can avoid messages being read more than once by using a transactional queue (where all messages have a priority of 0). For more information on transactional queues, see Transactional messaging.

**Processing messages within transactions**

When a message is retrieved from within a transaction, Message Queuing does not roll back the cursor if the transaction is aborted. For example, given a queue with two messages, message 1 and message 2, if you remove message 1 from within a transaction, Message Queuing will move the cursor to message 2. However, if the transaction is aborted for any reason, message 1 is inserted back into the queue, but the cursor remains at message 2.

For the peeking type of message processing in a queue containing one message when a COM component invoked by a trigger rule receives the message invoking it and the same COM component is configured to be part of an external transaction, if any operation in the transaction cannot be completed, the transaction is aborted and the message is returned to the queue. Then, if the trigger is serialized, the Triggers service waits until the process performed by the COM component ends (aborts) before starting to process a message again. In this case, the queue cursor moves on to wait for the next message, and the trigger does not invoke the COM component again for the same message. However, if the trigger is not serialized, as soon as the Triggers service finishes peeking at the message, it sends a request to peek at the next message and waits for a message to appear in

the queue. Therefore, when the transactional COM component is invoked, there is already a pending peeking operation in the system. When the message is returned to the empty queue in the aborted transaction, it arrives at the queue in the position to which the queue cursor is pointing. In this case, the returned message is processed again, the COM component is invoked for the same message, and a perpetual loop exists until a new message arrives or all the operations in the transaction can be completed.

**Sample implementation of Message Queuing Triggers**

Suppose you create a new trigger named PrivateTrigger, which is associated with the private queue MyComputer\MyQueue and has two rules attached to it: MyRule1 and MyRule2.

Each of these rules defines a condition (such as "Message label contains marketing") and an action to be taken when the condition is true (such as "open MSPaint.exe" with the parameter Message label).

When a message arrives at MyComputer\MyQueue, the Message Queuing Triggers service is notified. Next, for the trigger, PrivateTrigger, defined for that queue, the rules MyRule1 and MyRule2, which are attached to that trigger, are applied. Last, Microsoft Paint is opened if all the conditions for those rules are satisfied (are true).

For more information about implementing the Message Queuing Triggers service, see Installing Message Queuing triggers.

# Creating triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Creating triggers**

Triggers are created on the local computer in Triggers under the Message Queuing Triggers folder in the console tree of the Computer Management snap-in or in Triggers under any queue that is already monitoring a trigger by right-clicking **Triggers**, pointing to **New**, and then clicking **Trigger**. In the latter case, the queue selected automatically becomes the *monitored queue*. On the first page of the **New Trigger** dialog box, you must type a name for the new trigger in **Trigger name**. The *monitored queue* can be a user-defined queue registered in Active Directory, whose path name must be typed in **Queue path name**, or one of the automatically created queues on the local computer, which is selected by clicking the applicable button. Regardless of where the creation of a trigger is initiated, every trigger created on the local computer is displayed in the Computer Management snap-in under **Triggers** and under the applicable monitored queue. You then select the *message processing type*, and finally you indicate whether the trigger will be *enabled* and whether it will be *serialized*. For more information on the message processing type, enabled triggers, and serialized triggers, see Triggers overview.

By default, the Message Queuing Triggers service runs under the Network Service account. For heightened security, you can change the default Triggers service account to a more secure setting. You can also specify a less secure account for the Triggers service. The account used must be able to access the queues with which triggers are associated. Note that executable files invoked by triggers are invoked with the same security privileges as those assigned to the Triggers service. If your rule invokes a COM component, make sure that it can be successfully invoked from the security context used by the Triggers service, is properly registered, and has suitable dependencies. For more information, see Triggers overview.

You can attach an existing rule to the new trigger by selecting a rule and then associating it with the trigger. The order in which the rules are tested (their priorities) can be specified. For more information, see Create triggers.

Before creating a trigger, verify that the account under which the Message Queuing Triggers service is running has been granted the **Peek Message** permission for all existing queues that will be associated with triggers. For more information, see Set permissions for computer and queue objects.

**Managing triggers and queues**

When you delete a queue, the triggers associated with that queue are not deleted. If you later re-create a queue of the same name, any existing triggers will be used by this queue. However, this re-created queue will not be monitored by the Message Queuing Triggers service until the computer is restarted, or the associated triggers disabled and then enabled again. For instructions on setting the enabled or disabled property for a trigger, see View or change properties of triggers and rules for triggers.

# Modifying default trigger values

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Modifying default trigger values**

The Message Queuing Triggers service has a set of default configuration values. You can modify these values after the service is installed by using the Computer Management snap-in.

Configuration information as well as the trigger and rule definitions are stored in the local registry.

You can modify any of the following configuration settings:

- **Initial number of threads**. The default value is 5.
- **Maximum number of threads**. The default value is 20.
- **Default message body size**. The default value is 2048. This parameter is the size of the buffer that Message Queuing Triggers allocates for an incoming message. It is recommended that the value of this setting be sufficient for a buffer that can accommodate 90% of all messages that are expected to arrive at the monitored queue. An excessively low value will force the service to reallocate this buffer for every message that arrives with a larger message body. An excessively high value will waste resources.

When you change a configuration setting, you will need to stop and restart the service before the new setting will take effect. For more information on how to change configuration settings, see Change default configuration values for triggers.

# Implementing triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Implementing triggers**

Suppose you create a new trigger named PrivateTrigger, which is associated with the private queue MyComputer\MyQueue and has two rules attached to it: MyRule1 and MyRule2.

Each of these rules defines a condition (such as "Message label contains marketing") and an action to be taken when the condition is true (such as "open MSPaint.exe" with the parameter Message label).

When a message arrives at MyComputer\MyQueue, the Message Queuing Triggers service is notified. Next, for the trigger, PrivateTrigger, defined for that queue, the rules MyRule1 and MyRule2, which are attached to that trigger, are applied. Last, Microsoft Paint is opened if all the conditions for those rules are satisfied (are true).

For more information about implementing the Message Queuing Triggers service, see Installing Message Queuing triggers.

# Monitoring and Testing Message Queuing

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring and Testing Message Queuing**

This section provides information on monitoring Message Queuing, so that you can tune performance and identify issues. For details, see the following links:

- Monitoring events
- Monitoring Message Queuing performance
- Monitoring transactional messages
- Monitoring connectivity
- Sending test messages

# Monitoring events

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring events**

You can monitor Message Queuing events using Event Viewer.

Message Queuing logs information about its operation to the Windows Server 2003 family event logs. Message Queuing events are located in the Source column under "MSMQ." Most events are written to the Application Log. Security-related events are written to the security log, and service-related events are written to the system log.

For information on how to view Message Queuing events, see View Message Queuing events.

For more information on Event Viewer, see Event Viewer.

There are three types of application events for Message Queuing in the Windows Server 2003 family Application Log, as follows:

- **Informational events**. Infrequent, significant events that describe successful operations.
- **Warnings**. Not necessarily significant events, but indicative of potential problems.
- **Errors**. Events generated for significant problems.

Some application events are not logged at each occurrence. For example, if a Message Queuing server cannot establish a session with the destination computer, an event is logged for the first failure only and another event is logged after the session is established.

There are two types of security-related events for Message Queuing in the security log, as follows:

- **Success Audit**. Generated when selected objects are successfully accessed.
- **Failure Audit**. Generated when access to a selected object is denied.

For more information on access auditing, see Auditing Message Queuing objects.

# Monitoring Message Queuing performance

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring performance**
You can monitor the performance of Message Queuing by using the performance counters provided. You can view these performance counters using System Monitor. The default performance counters are grouped under the following objects:

- MSMQ Service
- MSMQ Queue
- MSMQ Session
- MSMQ Incoming HTTP Traffic
- MSMQ Outgoing HTTP Session
- MSMQ Incoming Multicast Session
- MSMQ Outgoing Multicast Session

For instructions on tracking counters, see Track Message Queuing service counters, and Track Message Queuing queue counters.

For more information about System Monitor, see System Monitor overview.

**Notes**

- Performance counters are not available when multiple instances of Message Queuing are running in a server cluster.
- Only the MSMQ Service and MSMQ Queue objects are always visible.

**MSMQ Service performance object**
The MSMQ Service performance object is provided on Message Queuing servers and independent clients. The MSMQ service object contains performance counters that you can use to monitor Message Queuing statistics for a selected computer.

The value you see for memory usage includes the memory space occupied by all messages in local queues on the computer. Because recoverable messages are stored on disk, this memory space is backed up by memory-mapped files on disk, not from the Windows Server 2003 family paged memory pool. For more information about how messages are stored, see "Local message storage" in Message storage.

**MSMQ Queue performance object**
The MSMQ Queue performance object contains performance counters that you can use to monitor message statistics for a selected queue. There can be an instance of the object for each open queue and for each queue containing messages on the computer where the queue resides up to the maximum number of instances allowed on that computer. In addition, there can be an instance for each outgoing queue opened for every destination queue on the computer where the outgoing queue resides up to the maximum number of instances allowed on that computer.

The name given a queue counter is the format name of the queue, unless the queue is an outgoing queue, in which case the name is the format name without the "PRIVATE=", "PUBLIC=", or "DIRECT=" prefix. For example, the following outgoing queues have the following queue counter names.

| Outgoing queue | Queue counter name |
|---|---|
| DIRECT=HTTP://fake/msmq/q1 | http:\\fake\msmq\q1 |

| | |
|---|---|
| DIRECT=OS:fake\q1 | os:fake\q1 |
| DIRECT=TCP:10.0.0.0\q1 | tcp:10.0.0.0\q1 |

Note that there can potentially be multiple outgoing queues (and counter instances) for the same queue, opened with different format names. If a connector application is running on the computer, a counter is available for the connector queue also; the queue name is displayed as **foreign=***GUID*.

Under **Instance**, only the first 97 queues created on the computer, including private queues, are listed. The queue counters for a specific computer appear for each instance. The **Computer Queues** instance represents the source or computer journal and the dead-letter queue for the computer. The **Bytes in Queue** and **Messages in Queue** counters monitor that computer's dead-letter queue. The **Bytes in Journal** and **Messages in Journal** counters then monitor the computer journal (Journal messages in the System Queues folder).

An outgoing queue is identifiable because the computer name listed in the counter instance matches the name of the computer on which the corresponding destination queue resides.

**MSMQ Session performance object**

The MSMQ Session performance object is visible only when the computer has established an active session with another Message Queuing computer. This object contains counters that can be used to monitor statistics about active sessions between computers running Message Queuing. Multiple instances of this object may be listed, one for each session. There can be an instance of this object on both computers involved in each session up to the maximum number of instances allowed on each computer. **Instance** lists the IP address of the computer with which there is an active session.

**Performance objects for HTTP messages**

The MSMQ Incoming HTTP Traffic and MSMQ Outgoing HTTP Session performance objects include counters that monitor messages sent over HTTP transport. These objects contain counters that monitor statistics about incoming HTTP messages, the rates at which the messages and message bytes arrive at the computer in an incoming HTTP connection, the total number of HTTP messages that arrived in the connection, and the total number of message bytes that arrived in the connection as well as counters that monitor the statistics about active outgoing HTTP sessions between the computer and other computers, the rates at which message bytes and messages are sent from the computer in a specific outgoing session, the total number of HTTP messages sent in the session, and the total number of message bytes sent in the session.

**Performance objects for multicast messaging**

The MSMQ Incoming Multicast Session and MSMQ Outgoing Multicast Session performance objects contain counters that monitor messages sent in multicast sessions. These objects contain counters that monitor statistics about active incoming multicast sessions, the rates at which the messages and message bytes arrive at the computer in a specific incoming multicast session, the total number of multicast messages that arrived in the session, and the total number of message bytes that arrived in the session as well as counters that monitor the statistics about active outgoing multicast sessions between the computer and other computers, the rates at which message bytes and messages are sent from the computer in a specific outgoing session, the total number of multicast messages sent in the session, and the total number of message bytes sent in the session.

**Improving messaging performance**

If available processing resources (CPU time) are not limiting messaging performance, it is possible that independent clients and Message Queuing servers can improve messaging performance by storing the various message and log files used to implement messaging across multiple physical disks.

To determine whether your computers are limited by disk access or by processing power, track the **% Processor Time** counter (located in the Processor object) and the **Avg. Disk Queue Length** counter (located in the DiskLength object). If the sustained **% Processor Time** value is above 75 percent during the time messages are being sent, adding processing capacity may improve messaging performance. If the **Avg. Disk Queue Length** value for any drive is greater than 0.6 when messages are being sent, additional disks may improve messaging performance.

For information about how to store the message and log files on separate physical disks, see Store messages on multiple disk drives.

# Monitoring transactional messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring transactional messages**

You can monitor Message Queuing transactions using the Component Services snap-in when the Microsoft Distributed Transaction Coordinator (DTC) is running.

Message Queuing uses the DTC, installed by default with Windows Server 2003 family, for implementing external transactions. You can use the Component Services snap-in to monitor statistics for transactional messages that are sent or received within external transactions and to monitor statistics for all transactions managed by the DTC. For more information, see View external transactions.

When you monitor external transactions, the details pane opened by clicking **Transaction Statistics** in the console tree of the Component Services snap-in provides useful information on active, committed, and aborted transactions. The DTC must be running to perform external transactional messaging for Message Queuing.

For more information on transactional messaging, see Transactional messaging.

For more information on the Microsoft DTC, see Using Component Services.

# Monitoring connectivity

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Monitoring connectivity**

Message Queuing offers an MQPing test, which has some similarity to the **ping** command supported by the TCP/IP protocol. You can use MQPing to determine whether the Message Queuing service is running on the computer being tested and whether that computer is currently connected to the network. When a computer running Message Queuing receives an MQPing request, it returns a response message to a response queue on the computer that sent the MQPing request. Thus, unlike the **ping** command, MQPing does not test for direct IP connectivity between the two computers, but verifies that messages can be sent between them with allowance for the possibility of routing.

For information on how to use MQPing, see Test connectivity using MQPing.

# Sending test messages

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Sending test messages**

The Computer Management snap-in or the Active Directory Users and Computers snap-in is used to administer queues. For more information, see Administration overview.

A test message is a message that is generated by the Message Queuing service on a specified source computer and sent to a target queue on a specified destination computer. The source computer can be any independent client or Message Queuing server, and the specified destination computer hosting the target queue can be any other independent client or Message Queuing server within the same forest. The target queue must be a test queue; that is, a nontransactional public queue with a queue type ID set equal to {55EE8F33-CCE9-11CF-B108-0020AFD61CE9}.

Test messages are sent to verify that messages can be sent from a given source computer and delivered to public queues on a given destination computer and, if message route tracking is enable, to track the route taken by messages sent from the source computer to public queues on the destination computer. For more information on tracking the routes taken by messages, see Tracking messages and routes.

You can either create a new test queue or transform an existing public queue into a test queue. Most existing public queues can be made to function as a test queue by changing the type ID to that of a test queue. However, it is useless to change the type ID of a transactional queue or a queue requiring message authentication to convert it into a test queue, because test messages, which are nontransactional and are not authenticated, will not be delivered to a transactional queue or to a queue requiring message authentication (unless the queue is located on the source computer).

Sending test messages is not enabled by default following Message Queuing installation, and must be enabled before you can send a test message. For instructions, see Enable route tracking and test messages. For information on how to send a test message, see Send test messages.

# Troubleshooting and Additional Resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting and Additional Resources**

This section provides help on troubleshooting Message Queuing and links to additional Message Queuing resources.

For help with installation, deployment, and administration issues, see:

- Troubleshooting installation problems
- Troubleshooting administration problems
- Troubleshooting triggers

For information and links to other Message Queuing resources, see:

- Additional Message Queuing resources

# Troubleshooting installation problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting installation problems**
This topic provides guidelines for solving problems associated with installing new Message Queuing software or running Message Queuing Setup.

When encountering installation problems, always inspect the setup log files first. This is especially important when you perform an unattended setup, because you will not know if something went wrong. The Message Queuing setup log file, Msmqinst.log, is located in the Windows folder. If this file does not identify the problem, look at the Windows Server 2003 family setup log file, Setuperr.log, which is located in the Windows folder (typically C:\Windows).

What problem are you having?

- You have completed Windows Server 2003 family Setup, but cannot find Message Queuing.
- You cannot install Message Queuing, or you receive an "Access is Denied" error.
- You have been prompted for the name of a server during Message Queuing Setup in a Windows Server 2003 family enterprise.
- You receive a "Message Queuing application is running" error while uninstalling Message Queuing.
- You are installing a Message Queuing computer as a dependent client and cannot configure a supporting server for it.

**You have completed Windows Server 2003 family Setup, but cannot find Message Queuing.**
**Cause:** Message Queuing is not installed by default during Windows Server 2003 family Setup.

**Solution:** Using **Add or Remove Programs** in Control Panel, install Message Queuing after Windows Server 2003 family is installed.

**See also:** Installation overview

**You cannot install Message Queuing, or you receive an "Access is Denied" error.**
**Cause:** To install Message Queuing clients, you must be logged on with local administrative permissions for the computer on which you want to install software. To install Message Queuing servers or Message Queuing servers with Downlevel Client Support on Windows Server 2003 family domain controllers, you need to have domain administrative permissions. To install Message Queuing servers with Routing Support on nondomain controllers, you need to have enterprise administrative permissions.

**Solution:** Verify that you are logged on with the proper permissions for the computer you want to install Message Queuing on.

**Solution:** Allow Active Directory replication to complete, and then try again.

**See also:** Installation permissions

**You have been prompted for the name of a server during Message Queuing Setup in a Windows Server 2003 family enterprise.**
**Cause:** For a Windows Server 2003 family computer, Setup was unable to automatically locate a Windows Server 2003 family domain controller in the local site or a nearby site.

**Solution:** Verify that both the computer on which you are running Message Queuing Setup and the domain controller have network connectivity.

**See also:** Installation overview

**You receive a "Message Queuing application is running" error while uninstalling Message Queuing.**
**Cause:** A Message Queuing application is currently running.

**Solution:** Close all Message Queuing applications.

**See also:** Start or stop services for Message Queuing

**You are installing a Message Queuing computer as a dependent client and cannot configure a supporting server for it.**
**Cause:** Following installation, the default setting is that Message Queuing server computers do not support dependent clients.

**Solution:** Enable the required Message Queuing server to act as a supporting server for dependent clients. For instructions, see Enable a supporting server.

# Troubleshooting administration problems

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting administration problems**
This topic provides guidelines for solving problems associated with viewing or setting properties, accessing objects in Active Directory, or using the snap-ins for Message Queuing.

What problem are you having?

- You cannot find Message Queuing under Services and Applications in Computer Management.
- You cannot find any objects for Message Queuing in Active Directory.
- You cannot access any objects for Message Queuing in Active Directory.
- You are unable to establish a connection with Active Directory.
- Changes you made do not show up in Active Directory.
- Windows NT 4.0 clients and local users cannot access Active Directory.
- You cannot access the contents of a particular queue.
- You cannot create queues.
- You cannot send messages to another site.
- You cannot send or receive authenticated messages.
- Messages sent to a queue are rejected.
- You cannot retrieve messages from a queue.
- Messages can no longer be delivered to a particular computer or queue.
- Express messages are missing or lost.
- Transactional messaging does not function for dependent clients.
- A test message cannot be sent to a particular queue.
- You cannot access messaging files stored locally on disk.
- You cannot locate the audit messages for a particular operation.
- There are problems communicating with IPX computers.
- Dependent clients cannot run Message Queuing applications.
- Resources for Message Queuing have timed out in a server cluster.
- A foreign site to which a Windows Server 2003 Message Queuing server with a connector application belongs does not appear in the list of sites for that server.

- You cannot register a user certificate for Message Queuing.

**You cannot find Message Queuing under Services and Applications in Computer Management.**
**Cause:** The Message Queuing service is not running.

**Solution:** In the Computer Management console tree, under **Services and Applications**, click **Services**. In the details pane, under **Name**, right-click **Message Queuing**, and then click **Start**. Restart Computer Management.

**You cannot find any objects for Message Queuing in Active Directory.**
**Cause:** You have not selected the necessary view options for these snap-ins.

**Solution:** In Active Directory Users and Computers, on the **View** menu, click **Users, Groups, and Computers as containers**, and then click **Advanced Features**. Message Queuing objects are typically located under the Computers container or the Domain Controllers container.

When Message Queuing is running in offline mode, you can see public queues, but not messages, queue properties, private queues, system queues, and triggers.

**Solution:** In Active Directory Sites and Services, on the **View** menu, click **Show Services Node**. Message Queuing objects are under the Servers object and under the Services object.

So that these views are always enabled by default, you can save these changes by creating your own .msc file for these snap-ins. In such a case, the first time you open a .msc file that you previously created, on the **File** menu, click **Options**, and then select one of the three user mode entries in **Console mode**. Now the console will subsequently open in a user mode with the correct views enabled. Opening a snap-in console in one of the user modes (rather than in author mode) is safer from an administrative standpoint, because no one else can potentially change the behavior or the look of the console.

**See also:** Message Queuing and Active Directory

**You cannot access any objects for Message Queuing in Active Directory.**
**Cause:** The Message Queuing service may have been changed to run under a specific user account. To provide access to objects in Active Directory for earlier Windows MSMQ clients, this service must run under the Local System account on Windows Server 2003 or Windows 2000 domain controllers.

**Solution:** Verify that the Message Queuing service is running under the Local System account on the applicable Message Queuing server.

**You are unable to establish a connection with Active Directory.**
**Cause:** The Message Queuing server with which you are communicating may be offline, or the Message Queuing service may be stopped.

**Solution:** Verify that the Message Queuing server with which you are communicating is online and that the Message Queuing service is running.

**Cause:** Your computer may be offline, your network adapter may be malfunctioning, or your computer may not have an IP address assigned (or may have an IP address of 0.0.0.0.).

**Solution:** Verify that your computer is online, your network adapter is functioning properly, and TCP/IP is properly configured.

**Cause:** The DNS server is offline, is not configured properly, or the Message Queuing server's IP address has recently changed and has not been replicated yet to the DNS server.

**Solution:** Verify that the DNS server is online and properly configured.

**Cause:** The site that you are in does not contain a Message Queuing server running on a Windows Server 2003 or Windows 2000 domain controller that is configured as a global catalog server. This affects MSMQ 1.0 users running Windows NT 4.0 and all users running under a local user account.

**Solution:** Verify the Message Queuing server you communicate with is properly configured.

**Cause:** You may have insufficient permissions to access Active Directory.

**Solution:** Grant yourself the proper permissions to perform this operation.

**Changes you made do not show up in Active Directory.**

**Cause:** Replication can delay information updates displayed in the snap-ins, so that changes you make using the snap-ins do not appear until after replication takes place.

**Solution:** Press F5 after making the change, or wait at least 15 minutes.

**Windows NT 4.0 clients and local users cannot access Active Directory.**
**Cause:** Weakened security was not selected during an installation of Message Queuing with Downlevel Client Support on a Windows Server 2003 domain controller or an installation of a Message Queuing server on a Windows 2000 domain controller in your forest. To effectively support MSMQ 1.0 clients running Windows NT 4.0 and users running under a local user account, security for Active Directory must be weakened. (Dependent clients cannot run under a local user account. Any computer that sends queries about Message Queuing objects to Active Directory on a domain controller directly, rather than through the Message Queuing directory service (Downlevel Client Support), will not be able to access Active Directory when it logs on with a local user account, even if the security for Active Directory is weakened.)

**Solution:** Use the ADSI Editor provided with the Windows Server 2003 family Resource Kit to weaken security.

**See also:** Message Queuing security overview

**You cannot access the contents of a particular queue.**
**Cause:** You may not have the proper permissions for viewing the queue. Message Queuing displays all queues that are viewable by default. However, you can view the contents of only those queues that you have proper permissions to open. Therefore, there may be public queues whose contents are not viewable in the snap-ins.

**Solution:** Ensure that the Peek Message or Receive Message permissions for the queue are granted to you. To view the contents of a queue journal, you must have the Peek Journal permission granted for the queue.

**See also:** Managing queues

**You cannot create queues.**
**Cause:** You may have insufficient permissions to perform this operation.

**Solution:** Grant yourself the proper permissions to perform this operation.

**See also:** Access control for Message Queuing

**You cannot send messages to another site.**
**Cause:** A routing link may not be defined between the two sites, or a designated site gate may be unavailable. As a result, messages can be sent by an application, but they cannot leave the outgoing queue.

**Solution:** Create a routing link between the two sites and define site gates for the routing link. It is recommended that you create multiple routing links between all sites for redundancy.

**See also:** Routing links

**You cannot send or receive authenticated messages.**
**Cause:** The user certificate for the computer may have been removed, expired, or not registered in Active Directory. The user certificate is required for sending authenticated messages. The user certificate is installed for you during Message Queuing Setup and is registered in Active Directory the first time you log on.

**Solution:** Use the **Message Queuing Properties** dialog box, which can be opened from Computer Management, to renew the certificate and then register it.

**Cause:** The Microsoft Base cryptographic service provider (CSP) provided with Windows Server 2003 family is not being used by Message Queuing applications. The Microsoft Base CSP is required for all Message Queuing applications that send or receive authenticated messages.

**Solution:** Verify that Message Queuing applications use the Microsoft Base (or Enhanced) CSP when sending or receiving messages.

**See also:** Authentication for Message Queuing

**Messages sent to a queue are rejected.**
**Cause:** The queue may restrict access.

**Solution:** Make sure that the queue is located in the same domain as the computer that sent the message or is in a trusted domain.

**Cause:** The default permission is that everyone can send messages to a queue. However if this default setting is changed, the

Message Queuing service on the receiving computer performs an access check to check the sender's SID (security ID) context against the SD (security descriptor) of the destination queue. If the access check is carried out and the sender has permissions to send to the queue, the message is accepted. If the access check is carried out and the sender does not have permissions, it is rejected. If the access check itself cannot be carried out and fails, the message is also rejected.

**Solution:** Check that the permissions settings for the queue. If the default setting to allow everyone to send messages to a queue has been modified, check that the message sender has permissions to send to the queue. If the access check itself cannot be completed (this might be indicated in Event Viewer), ensure that Message Queuing has access to the TokenGroupsGlobalAndUniversal attribute of the sender's user object. To provide this access, add the Message Queuing service account to the Windows Authorization Access Group, which has access to this attribute. Note that only users with domain administration permissions can add members to this group. You can add the Message Queuing service account to the group in one of two ways. You can manually add the relevant accounts to the Windows Authorization Access Group, repeating the operation for each Message Queuing computer requiring this permission. Alternatively, as a less secure practice, you can add the Authenticated Users group to the Windows Authorization Access Group. This grants every authenticated user, including the Message Queuing service on any computer, access to the TokenGroupsGlobalAndUniversal attribute for all users, and requires no further manual administration.

**See also:** Access control for Message Queuing

**You cannot retrieve messages from a queue.**
**Cause:** To retrieve messages from a queue, you must be granted the Receive Messages permission for the queue. It is not sufficient to have the Read permission granted for the queue.

**Solution:** Grant yourself the Receive Messages permission for the queue.

**See also:** Managing messages

**Messages can no longer be delivered to a particular computer or queue.**
**Cause:** The computer quota or the queue quota for the particular computer may have been exceeded.

You can verify this in the following two ways:

- By requesting negative acknowledgment messages and checking their class. (This technique is applicable to the quotas of destination queues, not to computer quotas.)

- By using performance counters to compare the queue or computer quota to the actual volume of messages.

**Solution:** Process and remove messages from the applicable queues or increase the appropriate quota value.

**See also:** Setting message quotas

**Express messages are missing or lost.**
**Cause:** Undelivered express messages are lost whenever the Message Queuing service is stopped on the computer where they reside in volatile memory, or the computer shuts down or fails.

**Solution:** Use recoverable or transactional messages. (Note that only transactional messages provide a complete guarantee against data loss.)

**See also:** Backing up and restoring messages

**Transactional messaging does not function for dependent clients.**
**Cause:** MSMQ 1.0 dependent clients running Windows NT 4.0 whose supporting server is running Windows Server 2003 family or Windows 2000 may not have Windows NT 4.0 Service Pack 4 (SP4) installed. SP4 or later is required for transactional messaging.

**Solution:** Install Windows NT 4.0 SP4 or later on the applicable dependent clients.

**Cause:** The Microsoft Distributed Transaction Coordinator (DTC) may not be running. The DTC is required for external transactions.

**Solution:** Start the DTC.

**A test message cannot be sent to a particular queue.**
**Cause:** Test messages can be sent from the MMC snap-in only to test queues, which have a queue type ID of {55EE8F33-CCE9-11CF-B108-0020AFD61CE9}.

**Solution:** Verify that the queue to which you want to send the test message has the correct queue type ID. Do not change a transactional queue or a queue requiring authentication to a test queue, because any test messages sent will not be delivered to such a queue. Note that you can send test messages programmatically to queues with any queue type ID.

**See also:** Sending test messages

**You cannot access messaging files stored locally on disk.**
**Cause:** By default, only members of the Administrators group for the local computer have full access to all locally stored message files, log files, and local queue storage (LQS) files.

**Solution:** Grant yourself administrative permissions for the computer, or add yourself to the local Administrators group.

**See also:** Access control for Message Queuing

**You cannot locate the audit messages for a particular operation.**
**Cause:** Audit messages are stored in the security log on the Message Queuing server that performs the operation, which may not be the server that owns the object. In addition, audit messages are created in the security log only when a queue is opened, not each time a message is received or sent.

**Solution:** Use Event Viewer to view information in the security log of the appropriate computer.

**See also:** Auditing Message Queuing objects

**There are problems communicating with IPX computers.**
**Cause:** The Message Queuing service no longer supports the IPX protocol.

**Solution:** Add the IP protocol to any computers that use only the IPX protocol.

**Dependent clients cannot run Message Queuing applications.**
**Cause:** You may be logged on using a local user account or other unauthenticated account. Message Queuing applications cannot run under unauthenticated accounts.

**Solution:** Log on using a domain user account, which is authenticated.

**Resources for Message Queuing have timed out in a server cluster.**
**Cause:** Cluster resources have a default time-out setting of three minutes for coming online. Message Queuing may take longer than three minutes, depending on recovery scenarios and the number of messages stored.

**Solution:** Increase the default time-out setting for Message Queuing resources.

**A foreign site to which a Windows Server 2003 Message Queuing server with a connector application belongs does not appear in the list of sites for that server.**
**Cause:** Adding a foreign site to the list of sites to which the Windows Server 2003 Message Queuing server with a connector application belongs does not ensure that the foreign site appears in the server's site list.

**Solution:** For a Windows Server 2003 Message Queuing server running connector applications, after adding the foreign site to the list of sites for the server, restart the Message Queuing service on the domain controller of the domain to which the Message Queuing server running the connector application belongs.

**You cannot register a user certificate for Message Queuing.**
**Cause:** By default, users have permission to register certificates for Message Queuing. However, if default user permissions are changed, this might affect your ability to register certificates.

**Solution:** Check that the user object is granted Write Personal Information permissions in the Active Directory schema.

**Cause:** Active Directory sets a multi-valued attribute limit of approximately 800 user certificates for a specific user account. This limit is usually exceeded when obsolete user certificates have not been deleted from Active Directory. If multiple certificates exist for a user account, only the latest is used.

**Solution:** Manually delete obsolete user certificates. For instructions, see Remove certificates for Message Queuing.

# Troubleshooting triggers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Troubleshooting triggers**
This topic discusses the most common problems encountered when using the Message Queuing Triggers service.

What problem are you having?

- The Message Queuing Triggers service fails to start.
- A trigger is defined, but it does not appear to be firing.
- The processing of messages by the Message Queuing Triggers service slows disproportionately as more messages arrive at a monitored queue.
- I cannot see my triggers application on the desktop, although I see in Task Manager that the application is running.

**The Message Queuing Triggers service fails to start.**
**Cause:** The Message Queuing Triggers service has been configured to run under a domain user account without proper administrative permissions.

**Solution:** Configure the Message Queuing Triggers service to run under the default Network Service account, which has sufficient permissions to run the Message Queuing Triggers service, or grant yourself the necessary permissions. The Triggers service must run under an account with administrator permissions.

**Cause:** The Message Queuing Triggers service has failed. In-process Component Object Model (COM) components that are invoked can sometimes fail and thereby cause the Message Queuing Triggers service to fail.

**Solution:** Rectify the condition that causes the COM component to fail, and then restart the Message Queuing Triggers service.

**A trigger is defined, but it does not appear to be firing.**
**Cause:** The trigger is currently disabled.

**Solution:** Verify that the trigger has been enabled.

**See also:** View or change properties of triggers and rules for triggers

**Cause:** The account under which the Message Queuing Triggers service is running has not been granted the Peek Message permission for the monitored queue. Or, if the message processing type of the trigger is "retrieval" or "transactional retrieval," the account has not been granted the Receive Message permission for that queue.

**Solution:** Change the permissions for the queues, and then restart the Message Queuing Triggers service.

**See also:** Set permissions for computer and queue objects

**Cause:** The Message Queuing Triggers service is not running.

**Solution:** Verify that this service has been started.

**Cause:** The Message Queuing Triggers service has been paused.

**Solution:** Verify that this service is not in a paused state.

**Cause:** The trigger conditions are not all satisfied. If any trigger condition evaluates as false, the trigger will not fire.

**Solution:** Verify that all conditions of the rule are true.

**Cause:** The Message Queuing Triggers service cannot be initialized when running under a non-administrator account.

**Solution:** By default, the Message Queuing Triggers service runs under the Network Service account. If you have changed this default, ensure that the account you have specified has administrator permissions.

**The processing of messages by the Message Queuing Triggers service slows disproportionately as more messages arrive at a monitored queue.**

**Cause:** This is normal when the Message Queuing Triggers service runs under heavy loads. It occurs because the maximum number of threads is reached.

**Solution:** You can accelerate the processing of messages by increasing the maximum number of allowable threads. You can increase the number of working threads in the Computer Management snap-in. To do this, open Computer Management, click **Services and Applications**, click **Message Queuing**, right-click **Triggers**, and then click **Properties**. Increase the maximum number of threads in **Max number of threads**.

**I cannot see my triggers application on the desktop, although I see in Task Manager that the application is running.**

**Cause:** The Triggers service runs by default under the Network Service account, and interaction with the desktop functionality is only available for the Local System account.

**Solution:** Configure the Triggers service to run under the Local System account and enable it to interact with the desktop. To do this, in Computer Management, click **Services**, and in the details pane, double-click **Message Queuing Triggers**. Click the **Log On** tab, click **Local System** and select the **Allow service to interact with desktop** check box.

# Additional Message Queuing resources

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Additional Message Queuing resources**
This section provides sources of additional information for Message Queuing and other software products that work with Message Queuing.

**Microsoft Platform Software Development Kit**
The Message Queuing Software Development Kit (SDK) is available as part of the Microsoft Platform SDK in the Microsoft Developer Network (MSDN) Library. The Message Queuing SDK includes background information on developing Message Queuing applications and a reference section that thoroughly describes the functions, properties, structures, and COM components supplied and used by Message Queuing. The Message Queuing SDK also contains sample code and program files for representative types of applications.

The Microsoft Platform SDK is released on a regular basis in the MSDN Library, which is available both on CDs and on the Web. To access the Message Queuing SDK on the Web, refer to the MSDN Library Web site. (http://msdn.microsoft.com/library/default.asp)

**Web Sites**

- For the latest information on Message Queuing, see the Message Queuing Web page.(http://www.microsoft.com/msmq)
- For the latest Frequently Asked Questions (FAQs) on Message Queuing, see the Message Queuing FAQ Web page. (http://support.microsoft.com/msmq/faq.htm)
- For the information on Message Queuing available from Microsoft Product Support Services, see the Microsoft Product Support Services home page.(http://support.microsoft.com/directory/default.asp)

**Internet newsgroups**
You can subscribe to the following newsgroups by following the instructions provided with your Web browser or newsreader software for connecting to an Internet news server:

- microsoft.public.msmq.deployment
- microsoft.public.msmq.interop
- microsoft.public.msmq.networking
- microsoft.public.msmq.performance
- microsoft.public.msmq.programming
- microsoft.public.msmq.security
- microsoft.public.msmq.setup

**Books**
If you are interested in reading books that pertain to application development for Message Queuing, see the following list:

- Dickman, Alan. *Designing Applications With MSMQ: Message Queuing for Developers*. Addison-Wesley Publishing Co., 1998.
- Homer, Alex, and David Sussman. *Professional MTS and MSMQ Programming with VB and ASP*. Wrox Press Inc., 1998.
- Sessions, Roger. Chapter 11 in *COM and DCOM: Microsoft's Vision for Distributed Objects*. John Wiley & Sons, 1998.

# Using Data Sources (ODBC)

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using Data Sources (ODBC)**
You can use Data Sources Open Database Connectivity (ODBC) to access data from a variety of database management systems. For example, if you have a program that accesses data in a SQL database, Data Sources (ODBC) will let you use the same program to access data in a Visual FoxPro database. To do this, you must add software components called drivers to your system. Data Sources (ODBC) helps you add and configure these drivers.

### Notes

- To open Data Sources (ODBC), click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Data Sources (ODBC)**.

- For information about using Data Sources (ODBC), in the **ODBC Data Source Administrator** dialog box, click **Help**.

# Windows Management Instrumentation Control

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Windows Management Instrumentation Control**

- For help on using WMI Control, see Windows Management Instrumentation Control How To....
- To understand more about WMI Control, see Windows Management Instrumentation Control Concepts.

# Using MSDN

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Using MSDN**

The Microsoft Developer Network (MSDN) is a complete resource for developers. You can find downloads, guidelines for hardware and driver compatibility testing, and how-to guides to help you build device drivers for Windows operating systems. For more information, see MSDN at the Microsoft Web site.

# Windows Deployment and Resource Kits

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Windows Deployment and Resource Kits**
The Microsoft® Windows® Deployment and Resource Kits provide IT professionals with the technical information and tools they need to successfully deploy and operate Windows.

See the topics in this Help and Support site for the content in the Microsoft® Windows® Server 2003 Deployment Kit, the Microsoft® Windows® Server 2003 Resource Kit, and Microsoft® Windows® XP Professional Resource Kit Documentation. For more information about purchasing the kits, or to browse the documentation on the Web, go to the Windows Deployment and Resource Kits Web site. This Web site also contains other important documentation developed by the Windows Resource Kit teams, including the Windows Server 2003 System Administration Scripting Guide and Windows 2000 documentation.

**Important**

- Some of the links on this page are to English-only content on the Web, regardless of the language of your operating system. In addition, if you install a Windows Deployment or Resource Kit on a non-English operating system, you will see English content mixed with non-English content in the Help and Support Center table of contents.

**Using Windows Server 2003 Help and Support Center from Windows XP Professional**
Windows Server 2003 Help and Support Center contains technical content for the IT professional that you won't find in Windows XP Professional. If you administer servers running Windows Server 2003 from a computer running Windows XP Professional, you should install Windows Server 2003 Help and Support Center. When you do this, your installed Windows Deployment and Resource Kits can be viewed and searched from within the Windows Server 2003 Help and Support Center. This way you have easy access to a more complete documentation set for server deployment and administration.

To install Windows Server 2003 Help and Support Center on a computer running Windows XP Professional, follow the steps in the Help topic Install Help content from a Windows Server 2003 server operating system onto a computer running Windows XP Professional. Windows Server 2003 family Help content is available on the Windows Server 2003 Deployment Kit companion CD, the Windows Server 2003 Resource Kit companion CD, and the Windows Server 2003 operating system CD.

# Microsoft Windows Server 2003 Deployment Kit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Windows Server 2003 Deployment Kit**
The Microsoft® Windows Server™ 2003 Deployment Kit provides guidelines and recommended processes for designing and deploying Windows Server 2003 family technologies to meet your business needs and IT goals.

**Important**

- Some of the links on this page are to English-only content on the Web, regardless of the language of your operating system. In addition, if you install a Windows® Deployment or Resource Kit on a non-English operating system or on an operating system with a Multilingual User Interface Pack (MUI), you will see some English content mixed with non-English content in Help and Support Center.

**Windows Server 2003 Deployment Kit**
For more information about the Windows Server 2003 Deployment Kit, see the Windows Server 2003 Deployment Kit Web site.

Planning for Deployment

Provides guidelines for creating a deployment project plan, setting up a deployment test lab, creating software and hardware inventories, choosing a licensing method, and conducting a deployment pilot project.

Deploying Network Services

Provides guidelines to help you build a TCP/IP network, deploy DHCP and DNS, and deploy WINS. This book also includes information about wireless networking, network security, and remote access connectivity.

Designing and Deploying Directory and Security Services

Provides guidelines to help you design and deploy directory services, including upgrading and restructuring domains. This book also includes planning and deployment guidelines for distributed security services, including authentication, access control, public key infrastructures, and smart cards.

Planning Server Deployments

Provides guidelines to help you design and deploy file servers, print servers, server clusters, Network Load Balancing clusters, and terminal services. This book also provides information on planning for storage, high availability, and remote server management.

Automating and Customizing Installations

Provides deployment planning information for automated installations, multilingual deployments, and compatibility testing, and includes design information for unattended installations, image-based installations, and Remote Installation Services (RIS) installations.

Designing a Managed Environment

Provides guidelines to help you design and deploy Group Policy infrastructure, security policy, Group Policy-based software deployment, update management, and a user data and settings management strategy.

Deploying Internet Information Services (IIS) 6.0

Provides prescriptive guidance on designing and deploying IIS 6.0 solutions for medium to large application server environments.

Tools and Technical References

Provides the complete set of Resource Kit tools and tools help, plus a technical reference on the Windows registry.

For more information about purchasing the Windows Deployment and Resource Kits, or to browse additional Windows

Deployment and Resource Kit documentation on the Web, go to the Windows Deployment and Resource Kits Web site.

# Microsoft Windows Server 2003 Resource Kit

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Windows Server 2003 Resource Kit**
Looking for The *Microsoft® Windows Server™ 2003 Resource Kit*? See the Windows Server 2003 Technical Reference documentation, which provides comprehensive information about the technologies included in the Microsoft® Windows Server™ 2003 operating system. This set of Technical Reference documentation is designed to assist IT planners and administrators by providing the foundational information about the technology components of the operating system

**Important**

- Some of the links on this page are to English-only content on the Web, regardless of the language of your operating system. In addition, if you install a Windows Deployment or Resource Kit on a non-English operating system or on an operating system with a Multilingual User Interface Pack (MUI), you will see some English content mixed with non-English content in Help and Support Center.

**Windows Server 2003 Technical Reference**
Technical Reference documentation for the Windows Server 2003 family of operating systems includes the following technology collections:

Active Directory Collection

This collection provides in-depth technical reference information about the Windows Server 2003 Active Directory® directory service and its components. These components include logical structure, physical structure, data storage, replication, and service publication. In addition, this collection includes technical reference information about the tools that are used to install, upgrade, and migrate an Active Directory environment.

Core Operating System Collection

This collection describes the core Windows Server 2003 operating system technologies. These technologies provide basic functionality such as printing services, data recovery, and troubleshooting.

Group Policy Collection

This collection provides in-depth technical information about Windows Server 2003 Group Policy and its supporting technologies, including core Group Policy, Group Policy components, and Group Policy administrative tools. These technologies enable you to deliver, apply, and centrally manage one or more configurations, including Administrative Templates (registry-based policy), folder redirection, scripts, security, and software installation settings, to a set of targeted users and computers within an Active Directory environment.

High Availability and Scalability Collection

This collection describes the Windows Server 2003 high availability technologies that are designed to reduce unscheduled outages and mitigate the impact of scheduled downtime for servers. This collection also describes the Windows Server 2003 scalability technologies used to increase or decrease computing capacity in order to address the capacity needs of an organization.

Networking Collection

This collection provides in-depth technical reference information about Windows Server 2003 technologies that make up a networking infrastructure based on the TCP/IP suite of protocols, services, and other industry-standard networking protocols and technologies.

Storage Technologies Collection

This collection describes the Windows Server 2003 storage technologies that are designed to help you store, access, and manage data on servers.

Windows Security Collection

This collection describes Windows Server 2003 security technologies, including authentication, authorization, and public key infrastructure (PKI), and provides detailed information about security policy settings.

### Internet Information Services 6.0 Technical Reference

Provides information about running IIS as a Web application platform, ensuring that IIS stays secure, monitoring and tuning IIS, administering the server programmatically, and capitalizing on built-in IIS scalability features to manage large-scale deployments.

### Tools and Technical References

Provides the complete set of Resource Kit tools and tools help, plus a technical reference on the Windows registry.

For more information about purchasing the Windows Deployment and Resource Kits, or to browse additional Windows Deployment and Resource Kit documentation on the Web, go to the Windows Deployment and Resource Kits Web site.

# Microsoft Windows XP Professional Resource Kit Documentation

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

**Microsoft Windows XP Professional Resource Kit Documentation**
Microsoft® Windows® XP Professional Resource Kit Documentation is a comprehensive technical resource for installing, configuring, and supporting Microsoft® Windows® XP Professional and Windows XP 64-Bit Edition in networks that use Windows 2000 Server, Windows NT® Server 4.0, and other server systems.

### Important

- All of the links on this page are to English-only content on the Web, regardless of the language of your operating system.

  The links on this page are to Windows XP Professional Resource Kit Documentation on the Web. If you have installed Windows XP Professional Resource Kit Documentation, you can access this information locally on your computer from the **Start** menu.

**Windows XP Professional Resource Kit Documentation**
To see a browsable version of Windows XP Professional Resource Kit Documentation on the Web, go to the Windows XP Professional Resource Kit Web site, or use the links below. To view specific chapters online, click a chapter link.


**Part I Deployment and Installation**
Ch. 1 Planning Deployments

Ch. 2 Automating and Customizing Installations

Ch. 3 Multilingual Solutions for Global Business

Ch. 4 Supporting Installations

**Part II Desktop Management**
Ch. 5 Managing Desktops

Ch. 6 Managing Files and Folders

Ch. 7 Supporting Mobile Users

Ch. 8 Configuring Remote Desktop

Ch. 9 Managing Devices

Ch. 10 Managing Digital Media

Ch. 11 Enabling Printing and Faxing

Ch. 12 Disk Management

Ch. 13 File Systems

Ch. 14 Backup and Restore

**Part III Security**
Ch. 15 Logon and Authentication

Ch. 16 Authorization and Access Control

Ch. 17 Encrypting File System

**Part IV Networking**
Ch. 18 Connecting Clients to Windows Networks

Ch. 19 Configuring TCP/IP

Ch. 20 Configuring IP Addressing and Name Resolution

Ch. 21 Connecting Remote Offices

Ch. 22 Configuring Telephony and Conferencing

**Part V Interoperability**
Ch. 23 Interoperating with UNIX

Ch. 24 Interoperating with NetWare

Ch. 25 Interoperating with IBM Host Systems

**Part VI System Troubleshooting**
Ch. 26 Troubleshooting Concepts and Strategies

Ch. 27 Troubleshooting Disks and File Systems

Ch. 28 Troubleshooting Startup

**Part VII Appendices**
Appendix A System Files Reference

Appendix B User Rights

Appendix C Common Stop Messages for Troubleshooting

Appendix D Tools for Troubleshooting

Appendix E Security Event Messages

Appendix F Device Manager Error Codes

Appendix G Differences with Windows XP Home Edition

Appendix H Differences with Windows XP 64-Bit Edition

Appendix I Accessibility for People with Disabilities

Appendix J Well-Known Security Identifiers

For more information about purchasing the Windows Deployment and Resource Kits, or to browse additional Windows Deployment and Resource Kit documentation on the Web, go to the Windows Deployment and Resource Kits Web site.

# R2: Product Help (R2 only)

Applies To: Windows Server 2003 R2

The content in this section applies to features released in Windows Server 2003 R2.

**In this section**

Microsoft Services for Network File System

DFS Management

Installing Windows Server 2003 R2

Active Directory Federation Services (ADFS)

Active Directory Application Mode (ADAM)

Storage Manager for SANs

File Server Resource Manager

Identity Management for UNIX: Welcome

Welcome to Subsystem for UNIX-based Applications

Print Management

# Microsoft Services for Network File System

Applies To: Windows Server 2003 R2

Microsoft® Services for Network File System (NFS) provides a file sharing solution for enterprises that have a mixed Windows and UNIX environment. With Microsoft Services for NFS, you can transfer files between computers running Microsoft Windows Server™ 2003 R2 and UNIX operating systems using the NFS protocol.

Microsoft Services for NFS is an update to the NFS components that were previously available in Windows® Services for UNIX 3.5. Microsoft Services for NFS includes the following new features:

- Active Directory Lookup. The Identity Management for UNIX Active Directory schema extension, available in Windows Server 2003 R2, includes UNIX user identifier (UID) and group identifier (GID) fields, which enables NFS clients to look up UNIX identity information directly from Active Directory.

- 64-bit support. Microsoft Services for NFS components can be installed on all editions of Windows Server 2003 R2, including 64-bit editions.

- Enhanced server performance.

- Support for UNIX special devices (mknod).

Microsoft Services for NFS enables you to support a mixed environment of Windows-based and UNIX-based operating systems. It also allows you to update your company's computers while supporting older technology during the transition phase. The following scenarios are examples of how enterprises can benefit from deploying Microsoft Services for NFS.

- Allow UNIX clients to access resources on computers running Windows Server 2003 R2. Your company may have UNIX clients accessing resources, such as files, on UNIX file servers. To take advantage of new Windows Server 2003 features, such as Shadow Copies for Shared Folders, you can move resources from your UNIX servers to computers running Windows Server 2003 R2 and then configure Microsoft Services for NFS to enable access by UNIX clients that are running NFS software. All of your UNIX clients will be able to access the resources using the NFS protocol with no changes required.

- Allow computers running Windows Server 2003 R2 to access resources on UNIX file servers. Your company may have a mixed Windows and UNIX environment with resources, such as files, stored on UNIX file servers. You can use Microsoft Services for NFS to enable computers running Windows Server 2003 R2 to access these resources when the file servers are running NFS software.

- Take advantage of new 64-bit hardware. You can run Microsoft Services for NFS components on 64-bit editions of Windows Server 2003 R2.

# Help for Microsoft Services for Network File System

Applies To: Windows Server 2003 R2


What's new in Microsoft Services for NFS

Best practices for configuring and using Microsoft Services for NFS

Installing and removing Microsoft Services for NFS components

Administering Microsoft Services for NFS

Microsoft Services for NFS troubleshooting

Microsoft Services for NFS command-line tools

# What's new in Microsoft Services for NFS

Applies To: Windows Server 2003 R2

Microsoft Services for Network File System (NFS) provides many improvements that make it easier than ever to integrate Windows and UNIX networks effectively.

**Active Directory Lookup**
The Identity Management for UNIX Active Directory schema extension, available in Microsoft Windows Server 2003 R2, includes UNIX user identifier (UID) and group identifier (GID) fields, which enables Server for NFS and Client for NFS to look up Windows-to-UNIX user account mappings directly from Active Directory. Identity Management for UNIX simplifies Windows-to-UNIX user account mapping management in Active Directory.

| **Note** |
| --- |
| After this release, the User Name Mapping component will no longer ship with Microsoft Services for NFS. |

**64-bit support**
You can install Microsoft Services for NFS components on all editions of Windows Server 2003 R2, including 64-bit editions.

**Enhanced server performance**
Microsoft Services for NFS includes a file filter driver, which significantly reduces common server file access latencies.

**UNIX special device support**
Microsoft Services for NFS supports UNIX special devices (mknod).

**Enhanced UNIX support**
Microsoft Services for NFS now supports the following versions of UNIX:

- Sun Microsystems Solaris version 9
- Red Hat Linux version 9
- IBM AIX version 5L 5.2
- Hewlett Packard HP-UX version 11i

**Deprecated features**
To streamline and simplify Microsoft Services for NFS, the following features were removed from this release:

- Gateway for NFS
- Server for PCNFS
- All PCNFS components of Client for NFS

# Best practices for configuring and using Microsoft Services for NFS

Applies To: Windows Server 2003 R2

**Use the correct utilities to administer Microsoft Services for NFS.**
You cannot use Microsoft Services for Network File System (NFS) or command-line utilities to administer earlier versions of Microsoft Services for NFS. Also, you cannot use earlier versions of Microsoft Services for NFS or command-line utilities to administer newer versions of Microsoft Services for NFS.

You can use Microsoft Services for NFS to administer NFS components on a remote computer if both computers are running a Windows Server 2003 R2 operating system.

**Open firewall ports.**
Microsoft Services for NFS requires several ports, which must be open on Windows Firewall and other firewalls. When prompted, make sure you allow Microsoft Services for NFS components.

For information about changing port settings, see Microsoft Services for NFS troubleshooting.

**See Also**
**Concepts**
Server for NFS best practices
User Name Mapping best practices
Client for NFS best practices

# Server for NFS best practices

Applies To: Windows Server 2003 R2

**General best practices**

**Provide user-level security.**

With Server for NFS, you can control access by users and groups to Network Files System (NFS) resources. You must configure and populate Active Directory with user identifier (UID) and group identifier (GID) information or install User Name Mapping on one computer in your network to associate Windows user accounts with UNIX user accounts. In some cases you may also need to install Server for NFS Authentication, see Server for NFS Authentication for more details.

**Secure files.**

Server for NFS only supports storage volumes that are formatted for NTFS. NTFS volumes enable you to provide file-level security by allowing and denying file access to specific users and groups. When you want anonymous users to access files, make sure directory and file permissions provide appropriate access for anonymous users. In addition, when sharing the directory, use host-level NFS access control when creating the NFS share to provide an additional level of security to protect the files in the directory.

**Secure new drives.**

When you add a new drive to a computer running Server for NFS, be sure to modify the permissions protecting the root directory of the drive to ensure untrusted users, including Everyone users, cannot write to the directory. This will prevent untrusted users from being able to compromise Server for NFS security by protecting shared directories on the drive.

**Allow users to disconnect before stopping the Server for NFS service.**

Before you stop Server for NFS or uninstall it, notify users who are connected to NFS shares that you are about the stop the service. You can then stop the service after the users have had the opportunity to close open files and disconnect from shared directories.

**Use naming conventions to identify shares with EUC encoding.**

If a directory is shared with one type of Extended UNIX Code (EUC) encoding (such as EUC-JP), and a client that is configured to use a different EUC encoding (such as EUC-TW) attempts to connect to the shared directory, unexpected results can occur. To prevent this, establish a naming convention to use when sharing directories with EUC encoding so users of client computers can know how shared directories are encoded.

**Protect configuration files.**

If you create configuration files (such as a character translation file or an audit log file), be sure to protect them with a discretionary access control list (DACL) that grants Full Control to the built-in System account and to the Administrators group. The DACL should contain no other entries.

**Best practices for running Server for NFS in a server cluster**

**Stop Server for NFS before stopping the server cluster.**

To ensure proper operation of Server for NFS in a server cluster, when stopping the server cluster, first stop Server for NFS, then stop the server cluster.

**Ensure share availability when a node fails.**

To ensure that a shared directory cluster resource will be available after the node containing the resource fails, make the cluster resource dependent on the appropriate physical disk resource.

**Use the appropriate tool to manage Network File System (NFS) share cluster resources.**

Although you can use Windows Explorer to view properties of an NFS share cluster resource, you should not use it to change those properties. You should only use Cluster Administrator or the **cluster** command to create and administer shared NFS directories on a server cluster.

**Avoid conflicting share names.**

Ensure that the share name of every shared directory on the server cluster is unique. Otherwise, if one node in the cluster fails over to another node, and if shared directories on the two nodes have the same name, then only one of the shared directories will be available.

**Ensure the availability of audit logs.**

Do not designate a shared disk resource as the location of the Server for NFS audit log. Only one node in the cluster can own the log file. This means that if the ownership of a group is moved to another node, the audit events for the remaining shares on the original node cannot be recorded in the file. To ensure the availability of Server for NFS audit logs, you should log events to the Event Viewer event log.

**Move file shares or take them offline before stopping Server for NFS.**

When Server for NFS is stopped on a cluster node hosting active NFS share resources, the Cluster service responds as though this were a failure of one of its managed resources. Consequently, the Cluster service will try to restart the service in an attempt to keep the resource available. Before you attempt to stop Server for NFS on a server cluster node, either move all groups containing NFS file shares to another node on the cluster, or take all NFS file shares on the node offline.

**Take resources offline before modifying.**
Be sure to take an NFS share directory resource offline before modifying its properties. Failing to do so might produce unexpected results.

**Administer Server for NFS only from computers in a trusted domain.**
To ensure that Server for NFS configuration changes are properly replicated, always use a computer that belongs to a trusted domain when you administer Server for NFS running on a cluster. This is necessary because Server for NFS configuration changes are not properly replicated among nodes in a cluster if you run Microsoft Services for NFS Administration or **nfsadmin** on a computer that belongs to a domain that is not trusted by the domain of the cluster.

**Restart the Server for NFS service after the Cluster service restarts.**
If the Cluster service must be restarted on a node in the cluster, stop and then restart the Server for NFS service on that node. This will ensure that Server for NFS configuration changes will be replicated properly among the nodes of the cluster.

**Choose the appropriate sharing mode.**
When you create an NFS share directory resource, you have the option of sharing the root of the specified directory or all the subdirectories in the directory. Before choosing this option, first decide whether you will need to control access to the subdirectories or change their share names. If you need to do this, you must share the subdirectories separately because when you create an NFS share directory resource to share the subdirectories, you cannot set access permissions, allow (or deny) anonymous access, or change the share name of the subdirectories separately.

If you choose to share all the subdirectories in the directory, make sure that the permissions protecting the directory do not allow untrusted users to create subdirectories. Otherwise, a malicious user can overwhelm the Cluster service by creating a very large number of subdirectories that would automatically be shared as a cluster resource.

**Use the command line properly when creating or modifying NFS share cluster resources.**
If you use the **cluster** command to create and modify NFS share cluster resources, the following considerations apply:

- When you create the cluster resource using the cluster **command**, you can set some nonprivate properties, but not all.
- When using the **cluster** command to set properties, do not rely on default values because they might be invalid for an NFS share cluster resource. Always set property values explicitly.
- To set or view permissions on an NFS share cluster resource, use Cluster Administrator. You can also use the **nfsshare** command to view permissions, but not to set them.

**Use hard mounts.**
Users of client computers should be instructed to use hard mounts when mounting NFS directories shared on the server cluster. This will ensure that, if the node sharing the directory fails, the client computer will not time out before the failover to another node is complete.

**Use the correct virtual server name.**
Users of client computers should be instructed to mount NFS directories shared on the server cluster using the virtual server name from the same group as the shared directory. Using a virtual server name from another group, or the node name, allows the client computer to mount the directory, but the mount can be lost in case of a failover.

# User Name Mapping best practices

Applies To: Windows Server 2003 R2

**Install User Name Mapping on a domain controller.**
If you have a large number of user maps, installing User Name Mapping on a domain controller improves performance by reducing network traffic when maps are first created or refreshed.

**Create a User Name Mapping server pool.**
You can use DNS round robin to create a pool of computers running User Name Mapping. This will provide improved performance on wide area networks as well as provide failover capability when one of the servers is no longer available. For more information, see Creating a User Name Mapping pool.

**Configure User Name Mapping on a server cluster.**
You can use a User Name Mapping server cluster to provide high availability of User Name Mapping. To ensure proper operation of User Name Mapping on a server cluster, when stopping the server cluster, first stop User Name Mapping, then stop the server cluster. For more information, see Configuring User Name Mapping on a server cluster.

**Make sure User Name Mapping can download users from all domains.**
In order to download users from a Windows domain, make sure that the User Name Mapping server belongs to a domain that is trusted by the Windows domain which it is trying to access.

If Active Directory was installed with permissions compatible with the Windows Server 2000 Server option, do one of the following:

- If the computer running User Name Mapping is a member of the domain, add the computer to the **Pre-Windows 2000 Compatible Access** security group.

- If the computer running User Name Mapping does not belong to the domain (that is, it belongs to a trusted domain), add the special group **Everyone** to the **Pre-Windows 2000 Compatible Access** security group of the Active Directory domain.

**Refresh data whenever a user is added or changed.**
To ensure that the user will have immediate access to Network File System (NFS) resources, refresh the User Name Mapping database immediately after you add a user or otherwise make changes to the user's Windows or UNIX accounts that would affect the user's mapping. For information about refreshing the database, see Refreshing maps.

**Place password and group files on the User Name Mapping server.**
If User Name Mapping is configured to use password and group files, these files must be located on a hard drive on the server to ensure that User Name Mapping can access the files whenever it refreshes the mapping database.

**Use appropriate permissions to protect password and group files.**
Protect password and group files with permissions that allow access only by appropriate users. We recommend that the permissions list contain only entries that grant Full Access to SYSTEM and the Administrators group. Also, do not change the permissions that Microsoft Services for NFS applies to other User Name Mapping configuration files.

**Back up complex maps whenever you make changes.**
To avoid loss of complex advanced maps in case of system difficulty, or to aid in transferring maps to another server, be sure to back up your User Name Mapping data whenever you change advanced maps.

**Ensure consistency of group mapping.**
To ensure proper file access, Windows and UNIX groups that are mapped to each other should contain the same users, and the members of the Windows and UNIX groups should be properly mapped to each other.

**Specify the computers that can access User Name Mapping.**
User Name Mapping requires you to identify the computers that can access User Name Mapping in the .maphosts file. (If the list in this file is empty, only the computer that is running User Name Mapping can access the service.) To maintain a high level of security, you should explicitly specify the computers that can access User Name Mapping, rather than using the plus sign (+) by itself to grant access to all computers. For more information, see Securing access to the User Name Mapping server.

# Client for NFS best practices

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**Use the correct command prompt.**
If you are running Client for NFS and using a command prompt, you should use the 32-bit version (Cmd.exe), which is available through the **Start** menu. Do not use the 16-bit version (Command.com). You must run Nfsadmin.exe from an elevated privilege command prompt to change Client for NFS configuration settings.

**Avoid using hard mounts.**
Regardless of the mount type, Client for NFS will not indefinitely try to mount an NFS shared resource. Once the shared resource has been successfully mounted, if the mount type is a hard mount and the Network File System (NFS) server becomes unavailable, Client for NFS will retry accessing the shared resource until the NFS server is once again available. As a result, Windows applications that are trying to access the shared resource will appear to stop responding.

| **Note** |
| --- |
| The **umount** command will disconnect a hard mount — even if the server is unavailable. |

# Installing and removing Microsoft Services for NFS components

Applies To: Windows Server 2003 R2


Installing Microsoft Services for NFS components

Removing Microsoft Services for NFS

# Installing Microsoft Services for NFS components

Applies To: Windows Server 2003 R2

**To install Microsoft Services for Network File System (NFS) components**

1. Click **Start**, point to **Settings**, and click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**. The Windows Component Wizard appears.
4. From the list of components, double-click **Other Network File and Print Services**.
5. Select the **Microsoft Services for NFS** box and click **OK**.
6. To select individual components, click **Details**. Then, select the check boxes of the components to install and click **OK**.
7. Click **Next** and follow the on-screen prompts to complete the installation.

# Removing Microsoft Services for NFS

Applies To: Windows Server 2003 R2

**To remove Microsoft Services for Network File System (NFS) components**

1. Click **Start**, point to **Settings**, and click **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**. The Windows Component Wizard appears.
4. From the list of components, double-click **Other Network File and Print Services**.
5. Select the **Microsoft Services for NFS** box and click **OK**.
6. To remove individual components, click **Details**. Then, select the check boxes of the components to remove and click **OK**.
7. Click **Next** and follow the on-screen prompts to remove the component.

# Administering Microsoft Services for NFS

Applies To: Windows Server 2003 R2


Using the Microsoft Services for NFS snap-in

Server for NFS administration

User Name Mapping administration

Client for NFS administration

# Using the Microsoft Services for NFS snap-in

Applies To: Windows Server 2003 R2

**To access Microsoft Services for Network File System**

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Microsoft Services for Network File System**.
2. Use this documentation to help you to administer Microsoft Services for Network File System (NFS).

# Server for NFS administration

Applies To: Windows Server 2003 R2


Understanding the Server for NFS component

Starting and stopping Server for NFS

Configuring Server for NFS

Securing Server for NFS

Optimizing Server for NFS performance

Using file systems with NFS

Managing NFS shares

Managing NFS client groups

Using Microsoft Services for NFS with server clusters

Server for NFS Authentication

# Understanding the Server for NFS component

Applies To: Windows Server 2003 R2

With Server for NFS, a computer running the Microsoft Windows Server 2003 R2 operating system can act as a Network File System (NFS) server. Users can then share files in a mixed environment of computers, operating systems, and networks.

Users on computers running NFS client software can gain access to directories (called shares) on the NFS server by connecting (mounting) those directories to their computers. From the viewpoint of the user on a client computer, the mounted files are indistinguishable from local files.

UNIX computers follow advisory locking for all lock requests. This means that the operating system does not enforce lock semantics on a file, and applications that check for the existence of locks can use these locks effectively. However, Server for NFS implements mandatory locks even for those locking requests that are received through NFS. This ensures that locks acquired through NFS are visible through the server message block (SMB) protocol and to applications accessing the files locally. Mandatory locks are enforced by the operating system.

# Starting and stopping Server for NFS

Applies To: Windows Server 2003 R2


Starting Server for NFS

Stopping Server for NFS

# Starting Server for NFS

Applies To: Windows Server 2003 R2

**To start Server for NFS from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS**, and then click **Start**.

**To start Server for NFS from the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server** [*ComputerName*] **start**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer on which you want to start Server for NFS. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsadmin server /?** |

# Stopping Server for NFS

Applies To: Windows Server 2003 R2

**To stop Server for NFS from the Windows interface**

1. Open Microsoft Services for Network File System: To open Microsoft Services for NFS, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS**, and then click **Stop**.

| **Note** |
| --- |
| Before attempting to stop Server for NFS running on a node in a server cluster, either move all groups containing Network File System (NFS) file shares to another node or take all NFS file shares offline. |

**To stop Server for NFS from the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server** [*ComputerName*] **stop**

| Argument | Description |
| --- | --- |
| *ComputerName* | The name of the computer on which you want to stop Server for NFS. |

| **Note** |
| --- |
| Before attempting to stop Server for NFS running on a node in a server cluster, either move all groups containing Network File System (NFS) file shares to another node or take all NFS file shares offline.<br><br>To view the complete syntax for this command, at a command prompt, type: **nfsadmin server /?** |

# Configuring Server for NFS

Applies To: Windows Server 2003 R2


Specifying how Server for NFS obtains Windows user and group information

Setting the lock grace period

Releasing locks

Setting up audit logging

Specifying Server for NFS TCP and UDP settings

Enabling NFS version 3

Configuring file-name character translation

Configuring the authentication renewal period

Configuring case sensitivity for file and folder names

Running Server for NFS on a server cluster

# Specifying how Server for NFS obtains Windows user and group information

Applies To: Windows Server 2003 R2

In a UNIX environment, authentication is the process of providing a user identifier (UID) and group identifier (GID) to a user who presents a valid user name and password. UNIX hosts provide these identifiers when a user logs on to the host. When the user logs on, the user's name and password are compared to those in a password file. If they correspond with a user name and password in the file, the server returns a corresponding UID and GID. The UID and GID identify the user for Network File System (NFS) browsing and mounting operations.

Server for NFS makes it possible for clients to access NFS resources without separately logging on to Server for NFS. Instead, the first time you try to access an NFS resource, Server for NFS looks up your UNIX UID and GID information in either Active Directory or User Name Mapping. If your UNIX UID and GID information is mapped to a Windows user and group accounts, the Windows names are returned to Server for NFS, which then uses the Windows user and group names to grant file access.

If your UNIX UID and GID information is not mapped, then Server for NFS will deny file access.

**Specifying how Server for NFS obtains Windows user and group information**
You can perform this procedure using the following methods:

- Using the Windows interface
- Using a command line

### To specify how Server for NFS obtains UID and GID information using the Windows interface

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Microsoft Services for NFS**, and then click **Properties**.
4. Specify how Server for NFS obtains user identifier (UID) and group identifier (GID) information for users:
   - To use Active Directory Lookup, select the **Active Directory Lookup** check box, type the domain name, and then click **Apply**.
   - To use User Name Mapping, type the name of the mapping server you want to use for authentication, and then click **Apply**.

| **Note** |
| --- |
| Before using Active Directory Lookup, you must install and populate the Identity Management for UNIX Active Directory schema extension, included in Windows Server 2003 R2, or have an equivalent schema which includes UNIX UID and GID fields. Identity Management for UNIX greatly simplifies this aspect of Windows-to-UNIX user account management in Active Directory.<br><br>You can specify the Internet Protocol (IP) address of the User Name Mapping server instead of the name of the server.<br><br>Before using User Name Mapping, the computer running Server for NFS must be listed in the .maphosts file on the computer running User Name Mapping. For more information, see Securing access to the User Name Mapping server |

### To specify how Server for NFS obtains UID and GID information using the command line

1. Open the command prompt.
2. Specify how Server for NFS obtains user identifier (UID) and group identifier (GID) information for users:

- To use User Name Mapping, at a command prompt, type:

  **nfsadmin mapping [***ComputerName***] config adlookup=yes addomain=***DomainName***

- To use User Name Mapping, at a command prompt, type:

  **nfsadmin mapping [***ComputerName***] config mapsvr=***ServerName***

| Argument | Description |
|---|---|
| *DomainName* | The name of the domain controller you want to use. |
| *ComputerName* | The name of the computer you want to manage. |
| *ServerName* | The name of the mapping server you want to use. |

**Note**

To view the complete syntax for this command, at a command prompt, type: **nfsadmin mapping /?**

Before using Active Directory Lookup, you must install and populate the Identity Management for UNIX Active Directory schema extension, included in Windows Server 2003 R2, or have an equivalent schema which includes UNIX UID and GID fields. Identity Management for UNIX simplifies this aspect of Windows-to-UNIX user account management in Active Directory.

Before using User Name Mapping, the computer running Server for NFS must be listed in the .maphosts file on the computer running User Name Mapping. For more information, see Securing access to the User Name Mapping server

# Setting the lock grace period

Applies To: Windows Server 2003 R2

**To set the lock grace period using the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS** and then click **Properties**.
4. Click the **Locking** tab.
5. In **Waiting Period**, type the number of seconds within which Network File System (NFS) clients can reestablish file locks after a network or server service interruption..
6. To save this setting, click **Apply**.

**To set the lock grace period using the command line**

1. Open the command prompt.
2. At a command prompt, type:

   **nfsadmin server [Computer** *Name* **] config lockperiod=** *LockingGracePeriod*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *LockingGracePeriod* | The time period (in seconds) within which users can reestablish locks after the server restarts. |

| Note |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsadmin server /?** |

# Releasing locks

Applies To: Windows Server 2003 R2

**To release locks using the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS**, and then click **Properties**.
4. Click the **Locking** tab.
5. In the list of clients, click the client whose locks you want to release, click **Release Locks**, and then click **Apply**.

**To release locks using the command line**

1. Open the command prompt.
2. To list all locks, type:

   **nfsadmin server [Computer** *Name* **] –l**

3. To release locks, type:

   **nfsadmin server [Computer** *Name* **] –r [C** *lient***| all]**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *Client* | The client whose locks you want to release. |
| **All** | Releases all locks. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**nfsadmin server /?** |

# Setting up audit logging

Applies To: Windows Server 2003 R2

**To set up logging using the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS**, and then click **Properties**.
4. Click the **Audit Logging** tab.
5. To record audited events in the Event Viewer application log, select **Log events to event log**.
6. To record audited events in a file, select **Log events to the following text file**. In the text box, type the name of the log file, or click **Browse** to find the file, and then in **Maximum file size**, increase the maximum size for the log file, if desired (default: 64MB).
7. In **Server for NFS Events**, select the events to audit from the following options:
   - **Mount and unmount shares.**
   - **Lock and unlock files**.
   - **Read** files.
   - **Write** files.
   - **Create** files.
   - **Delete** files.
   - **All events listed above**.
8. To save the settings, click **Apply**.

| **Note** |
| --- |
| By default, no events are audited. If recording audited events to a text file, the log file must be on your local computer. To stop logging an event, clear its check box. If you choose to record Server for NFS audited events in the event log, the string corresponding to the file name could be truncated due to a system limitation.<br><br>One solution to this problem is to record audited events to a text file. On a server cluster, it is recommended that you record audited events to the Event Viewer because its data replicates to all the nodes in the cluster. In this scenario, recording audited events to a text file in addition to the event log can help associate complete file names with a particular audit log entry. |

# Specifying Server for NFS TCP and UDP settings

Applies To: Windows Server 2003 R2

**To specify Server for NFS TCP and UDP settings using the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System.**
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS**, and then click **Properties**.
4. Click the **Server Settings** tab.
5. Select the desired transport protocol setting from the **Transport protocol** drop-down list, and then click **Apply**.

| **Important** |
|---|
| These changes will not take effect until Server for NFS is restarted. For information about how to stop and start Server for NFS, see Starting and stopping Server for NFS. |

**To specify Server for NFS TCP and UDP settings using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server [** *ComputerName* **] config [** *Protocol* **]= [** *Transport* **]**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *Protocol* | The name of the Network File System (NFS) protocol you want to configure. See the list of "Protocol types" below. |
| *Transport* | The transport or combination of transports to enable for the given NFS protocol. See the list of "Transport types" below. |

**Protocol types:**

| Protocol | Description |
|---|---|
| PortmapProtocol | Specifies the Portmap protocol. |
| MountProtocol | Specifies the Mount protocol. |
| NLMProtocol | Specifies the Network Lock Manager protocol. |
| NSMProtocol | Specifies the Network Status Manager protocol. |

**Transport types:**

| Transport | Description |
|---|---|
| TCP | Enables only the TCP transport. |
| UDP | Enables only the UDP transport. |

| | |
|---|---|
| TCP+UDP | Enables both the TCP and UDP transports. |

**Important**

These changes will not take effect until Server for NFS is restarted. For information about how to stop and start Server for NFS, see Starting and stopping Server for NFS.

**Note**

To view the complete syntax for this command, at a command prompt, type: **nfsadmin server /?**

# Enabling NFS version 3

Applies To: Windows Server 2003 R2

**To enable NFS Version 3 using the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **AdministrativeTools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS**, and then click **Properties**.
4. Click the **Server Settings** tab.
5. Select the **Enable NFS v3 support** check box, and then click **Apply**.

| **Important** |
|---|
| These changes will not take effect until Server for NFS is restarted. For information about how to stop and start Server for NFS, see Starting and stopping Server for NFS. |

**To enable NFS Version 3 using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server [** *remoteMachineName* **] config enablev3={yes | no}**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to configure. |

| **Important** |
|---|
| These changes will not take effect until Server for NFS is restarted. For information about how to stop and start Server for NFS, see Starting and stopping Server for NFS. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**nfsadmin server /?**<br>NFS v3 support is enabled by default. |

**See Also**
**Concepts**
Starting and stopping Server for NFS

# Configuring file-name character translation

Applies To: Windows Server 2003 R2

**To configure file-name character translation using the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS**, and then click **Properties**.
4. Click the **Filename Handling** tab.
5. Select the **Translate file names**check box.
6. In the **Character translation file** box, type the fully qualified path of the file containing the character-translation map, or click **Browse** to find the file, and then click **Apply**.

| **Important** |
|---|
| These changes will not take effect until Server for NFS is restarted. For information about how to stop and start Server for NFS, see Starting and stopping Server for NFS. |

| **Important** |
|---|
| To prevent unauthorized users from tampering with the file-name character translation file and causing Server for NFS to fail, ensure that the file is protected with a discretionary access control list (DACL) that grants Full Control to the built-in System account and to the Administrators group. The DACL should contain no other entries. |

**To configure file-name character translation using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server [** *ComputerMachineName* **] config translationfile=** *file*

| Argument | Description |
|---|---|
| *remoteMachineName* | The name of the computer you want to configure. |
| *file* | The fully qualified path and name of the file containing the character-translation map list. If you do not specify a file, character translation will be disabled. |

| **Important** |
|---|
| These changes will not take effect until Server for NFS is restarted. For information about how to stop and start Server for NFS. |

| **Important** |
|---|
| To prevent unauthorized users from tampering with the file-name character translation file and causing Server for NFS to fail, ensure that the file is protected with a discretionary access control list (DACL) that grants Full Control to the built-in System account and to the Administrators group. The DACL should contain no other entries. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**nfsadmin server /?** |

**See Also**

**Concepts**

Starting and stopping Server for NFS

# Configuring the authentication renewal period

Applies To: Windows Server 2003 R2

**To configure how often authentication is renewed using the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS**, and then click **Properties**.
4. Click the **Server Settings** tab.
5. Do one of the following:
   - To prevent Server for NFS from periodically renewing authentication, click **Do not renew authentication**.
   - To configure Server for NFS to renew authentication periodically, click **Renew authentication every:** and type how many seconds will elapse before authentication will be renewed. The default is 600 seconds.
6. Click **Apply**.

**To configure how often authentication is renewed using the command line**

1. Open the command prompt.
2. To configure Server for NFS to renew authentication periodically, type the following at the command prompt:

   **nfsadmin server [** *ComputerName* **] config renewauth=yes**

   **nfsadmin server [** *ComputerName* **] config renewauthinterval=** *seconds*
3. To configure Server for NFS to not renew authentication periodically, type the following at the command prompt:

   **nfsadmin server [** *ComputerName* **] config renewauth=no**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to configure. |
| *seconds* | The number of seconds that will elapse before authentication is renewed. The default is 600. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsadmin server /?** |

# Configuring case sensitivity for file and folder names

Applies To: Windows Server 2003 R2

**To configure case sensitivity for file and folder names using the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System.**
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NFS**, and then click **Properties**.
4. Click the **Filename Handling** tab.
5. Do one of the following:
   - To enable case-sensitive file and directory name lookups, select the **Enable case sensitive lookups** check box.
   - To disable case-sensitive file and directory name lookups, clear the **Enable case sensitive lookups** check box.
6. Click **Apply**.

**Important**

These changes will not take effect until Server for NFS is restarted. For information about how to stop and start Server for NFS, see Starting and stopping Server for NFS.

You also need to disable Windows kernel case-insensitivity in order for Server for NFS to support case-sensitive file names. You can disable Windows kernel case-insensitivity by clearing the following registry key to 0:

HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\kernel

DWORD "obcaseinsensitive"

**To configure case sensitivity for file and folder names using the command line**

1. Open the command prompt.
2. To enable case sensitivity, type the following:

   **nfsadmin server [** *ComputerName* **] config casesensitivelookups=yes**
3. To disable case sensitivity and optionally specify the case of file names returned by Server for NFS, at a command prompt, type the following:

   **nfsadmin server [** *ComputerName* **] config casesensitivelookups=no [ntfscase={upper | lower | preserve}]**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to configure. |

**Important**

These changes will not take effect until Server for NFS is restarted. For information about how to stop and start Server for NFS, see Starting and stopping Server for NFS.

**Note**

The ntfscase option sets the case sensitivity for the NTFS file system. The default case sensitivity is preserve (preserve case).

To view the complete syntax for this command, at a command prompt, type:

**nfsadmin server /?**

**See Also**
**Concepts**
Starting and stopping Server for NFS

# Running Server for NFS on a server cluster

Applies To: Windows Server 2003 R2


Sharing directories as cluster resources

Modifying shared directory cluster resources

Deleting shared directory cluster resources

# Sharing directories as cluster resources

Applies To: Windows Server 2003 R2

**To share a directory as a cluster resource using the Windows interface**

1. Open Cluster Administrator: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.
2. In the console tree (usually the left pane), double-click the **Groups** folder.
3. In the details pane (usually the right pane), click the group to which you want the resource to belong.
4. On the **File** menu, point to **New**, and then click **Resource**.
5. In the New Resource Wizard, type the appropriate information in **Name** and **Description**. In **Resource type**, click **NFS Sharing**. Click **Next**.
6. Add or remove possible owners of the resource, and then click **Next**.
7. To add or remove dependencies, do one of the following:
   - To add dependencies, under **Available resources**, click a resource, and then click **Add**.
   - To remove dependencies, under **Resource dependencies**, click a resource, and then click **Remove**.
8. In **Share Name**, type the name of the shared folder.
9. In **Share Path**, type the fully qualified path of the folder to be shared.
10. Do one of the following:
    - To share the folder and its contents as a single shared folder, click **Share Root Only**.
    - To share all folders (except hidden folders) in the folder specified in the previous step as separate shared folders, click **Share Subdirectories Only**.
11. To allow access by users who have not been authenticated, click **Allow anonymous access**.
12. To change the user identifier (UID) and group identifier (GID) values assigned to anonymous users, type the new values in the **Anonymous UID** and **Anonymous GID** boxes.
13. In the **Encoding** list, select the type of encoding to be used for file and directory names in the shared directory. In addition to the default American National Standards Institute (ANSI) encoding, the following encoding is available:
    - BIG5 (Chinese)
    - EUC-JP (Japanese)
    - EUC-KR (Korean)
    - EUC-TW (Chinese)
    - GB2312-80 (Simplified Chinese)
    - KSC5601 (Korean)
    - SHIFT-JIS (Japanese)

    If this option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:

- Japanese: SHIFT-JIS
- Korean: KS_C_5601-1987
- Simplified Chinese: GB2312-80
- Traditional Chinese: BIG5

14. To specify permissions for the client computers that can access the shared folder, click **Permissions**.
15. Click **Finish**.

| **Important** |
| --- |
| Before adding a resource to your cluster, you must verify the following: |

- The NFS Share resource type is available. This resource type was created when you installed Server for NFS on the cluster node.
- A group already exists within the cluster to which your resource will belong.
- All dependent resources have been created.

| **Note** |
| --- |
| In Windows Server 2003 operating systems, the Everyone group does not include anonymous users by default. For information about providing anonymous access when Server for NFS is running on these operating systems, see Allowing anonymous access to resources by NFS clients. |

**To share a directory as a cluster resource using the command line**

1. Open the command prompt.
2. Type the following:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /create /restype:"NFS Share" /group:"** *GroupName* **" /privproperties ShareName="** *ShareName* **":STRING Path="** *SharePath* **":STRING ShareSubDirs=** *ShareType* **:DWORD GlobalPerm=** *PermType* **:DWORD Encoding=** *EncodingType* **:DWORD AnonymousAccessAllowed=** *AccessType* **:DWORD**

3. Then, type:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /online /wait[:** *Timeout* **]**

| **Argument** | **Description** |
| --- | --- |
| *ClusterName* | The name of the cluster you want to manage. |
| *ResourceName* | The name of the cluster resource. |
| *GroupName* | The name of the group to which you want to add the resource. |
| *ShareName* | The name of the shared directory you want to create. This value is ignored if ShareType is **0x1** |
| *SharePath* | The fully qualified path of the directory to be shared. |

| | |
|---|---|
| *ShareType* | **0x0** if only the directory specified by SharePath is to be shared. **0x1** if the subdirectories (except hidden subdirectories) in the directory specified by SharePath are to be shared individually. |
| *PermType* | Specifies root-access and read-write permission combinations. PermType must be one of the following values:<br>**0x1**—No access<br>**0x9**—No Access + Root<br>**0x2**—Read only<br>**0xa**—Read only + Root<br>**0x4**—Read/write<br>**0xc**—Read/write + Root |
| *EncodingType* | Specifies the default encoding used for file and directory names and, if used, must be set to one of the following:<br>• **0x7**—ANSI<br>• **0x5**—BIG5 (Chinese)<br>• **0x0**—EUC-JP (Japanese)<br>• **0x1**—EUC-KR (Korean)<br>• **0x2**—EUC-TW (Chinese)<br>• **0x6**—GB2312-80 (Simplified Chinese)<br>• **0x4**—KS_C_5601-1987 (Korean)<br>• **0x3**—SHIFT-JS (Japanese)<br><br>If this is option is set to **0x7**, the encoding scheme is ANSI or, on systems configured for non-English locales, the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:<br>• Japanese: SHIFT-JIS<br>• Korean: KS_C_5601-1987<br>• Simplified Chinese: GB2312-80<br>• Traditional Chinese: BIG5 |
| *AccessType* | **0x0**if anonymous access to the shared directory is not allowed. **0x1** if anonymous access is allowed. |
| *Timeout* | The number of seconds **cluster** waits before canceling the command if it does not successfully complete. If **/wait** is specified without this value, **cluster** waits indefinitely or until the resource state changes. |

| **Important** |
|---|
| Before adding a resource to your cluster, you must verify the following: |

- The NFS Share resource type is available. This resource type was created when you installed Server for NFS on the cluster node.
- A group already exists within the cluster to which your resource will belong.
- All dependent resources have been created.

| **Note** |
| --- |
| All arguments listed above are required. Using default values might produce unexpected results.<br><br>You must use the Windows interface to set permissions on shared directories.<br><br>If the attempt to share the directory fails, **cluster** displays an error message, and detailed information about the failure is recorded in the cluster log.<br><br>You cannot set the user identifier (UID) and group identifier (GID) values assigned to anonymous users using the command line.<br><br>In Windows Server 2003 operating systems, the Everyone group does not include anonymous users by default. For information about providing anonymous access when Server for NFS is running on these operating systems, see Allowing anonymous access to resources by NFS clients.<br><br>To view the complete syntax for this command, at a command prompt, type:<br><br>**cluster /?** |

# Modifying shared directory cluster resources

Applies To: Windows Server 2003 R2

**To modify a shared directory cluster resource using the Windows interface**

1. Open Cluster Administrator: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.
2. In the console tree (usually the left pane), double-click the **Groups** folder.
3. In the details pane (usually the right pane), double-click the group that contains the resource you want to modify, and then click the resource.
4. On the **File** menu, click **Take Offline**.
5. When the resource is offline, on the **File** menu, click **Properties**.
6. Change the properties you want to modify, and then click **OK**.
7. On the **File** menu, click **Bring Online**.

**To modify a shared directory cluster resource using the command line**

1. Open the command prompt.
2. Type the following:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /offline /wait[:** *Timeout* **]**
3. Then, type:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /** *Option*
4. Then, type:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /online /wait[:** *Timeout* **]**

| Argument | Description |
|---|---|
| *ClusterName* | The name of the cluster you want to manage. |
| *ResourceName* | The name of the cluster resource. |
| *Timeout* | The number of seconds **cluster** waits before canceling the command if it does not successfully complete. If **/wait** is specified without this value, **cluster** waits five minutes or until the resource state changes. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **cluster /?** |

**Cluster Resource Options**

The following table describes the available cluster resource options:

| Option | Use to... |
|---|---|
| **/status** | Display the status of a cluster resource (Online, Offline, or Failed). The resource name is optional. If you do not provide a resource name, the status of all cluster resources is displayed. |
| **/create /group:***group name***/type:***resource type***[/separate]** | Create a new resource in a specified group. Use **/separate** to specify that the resource should run in a separate Resource Monitor. |
| **/delete** | Delete a resource. |

| | |
|---|---|
| **/rename:***new resource name* | Rename a resource. |
| **/addowner:** *node name* | Add a node name to the list of possible owners. |
| **/removeowner:node name** | Remove a node name from the list of possible owners. |
| **/listowners** | Display a list of possible owners. |
| **/moveto:** *group* | Move the resource to a different group. |
| **/properties [** *property list* **]** | View or set the resource common properties. The resource name is optional. If you do not provide a resource name, properties for all resources are displayed. |
| **/privproperties [** *property list* **]** | View or set the resource private properties. Resource name is optional. If you do not provide a resource name, private properties for all resources are displayed. |
| **/fail** | Initiate resource failure. |
| **/online [/wait[:** *timeout in seconds* **]]** | Bring the resource online. |
| **/offline [/wait[:** *timeout in seconds* **]]** | Take the resource offline. |
| **/listdependencies** | List the dependencies for a resource. |
| **/adddependency:** *resource* | Add a dependency for a resource. |
| **/removedependency:** *resource* | Remove a dependency for a resource. |
| **/addcheckpoints:** *key* **[** *subkey* **...][,** *key* **[** *subkey* **...]...]** | Add a registry checkpoint for a resource. |
| **/removecheckpoints:** *key* **[** *subkey* **...][,** *key* **[** *subkey* **...]...]** | Remove a registry checkpoint for a resource. |
| **/checkpoints** | List registry checkpoints for a resource. The resource name is optional. Checkpoints for all cluster resources are displayed if you do not provide a resource name. |
| **/addcryptocheckpoints:** *key* **[** *subkey* **...][,** *key* **[** *subkey* **...]...]** | Add a cryptographic key checkpoint for a resource. |
| **/removecryptocheckpoints:** *key* **[** *subkey* **...][,** *key* **[** *subkey* **...]...]** | Remove a cryptographic key checkpoint for a resource. |
| **/cryptocheckpoints** | List cryptographic key checkpoints for a resource. The resource name is optional. If you do not provide a resource name, checkpoints for all cluster resources are displayed. |
| **/?** or **/help** | Display **cluster resource** syntax. |

**Cluster resource common property names**

The following table describes the common cluster resource property names, their uses, and valid settings:

| **Common property names** | **Use** |
|---|---|
| **Description** | Describes a resource. |
| **Type** | Specifies the resource type. |
| **DebugPrefix** | Specifies the appropriate debugger for the resource. For more information, see the Microsoft Platform Software Development Kit. |
| **SeparateMonitor** | Indicates whether a resource shares a Resource Monitor. Valid values are True (1) and False (0). |

| | |
|---|---|
| **PersistentState** | Describes the last known persistent state of a resource. In order to save the resource's PersistentState property, the resource must be offline. |
| **LooksAlivePollInterval** | Recommended interval (in milliseconds) that the Cluster service polls a resource to determine whether it appears to be operational. If a value is not specified, a default value is taken from the LooksAlivePollInterval property for the resource type. |
| **IsAlivePollInterval** | Interval (in milliseconds) that the Cluster service polls a resource to determine whether it is operational. If a value is not specified, a default value is taken from the IsAlivePollInterval property for the specific resource type. IsAlivePollInterval cannot be zero. |
| **RestartAction** | Describes the action to perform if the resource fails. The choices are as follows:<br><br>ClusterResourceDontRestart (0): Do not restart following a failure.<br><br>ClusterResourceRestartNoNotify (1): If the resource exceeds its restart threshold within its restart period, the Cluster service does not attempt to failover the group to another node.<br><br>ClusterResourceRestartNotify (2): If the resource exceeds its restart threshold within its restart period, the Cluster service attempts to failover the group to another node.<br><br>If you do not enter a value, the default allows the resource to affect the group, which can cause the group to failover to another system. |
| **RestartThreshold** | Specifies how many times the Cluster service attempts to restart the resource in the RestartPeriod before failing over the group. |
| **RestartPeriod** | Specifies the amount of time allowed for the restart attempts to reach the RestartThreshold before the Cluster service fails over the group. |
| **PendingTimeout** | Specifies the amount of time that a Pending Online or Pending Offline resource has to resolve its status before the Cluster service puts the resource in Offline or Failed status.<br><br>The default value is three minutes. |

# Deleting shared directory cluster resources

Applies To: Windows Server 2003 R2

**To delete a shared directory cluster resource using the Windows interface**

1. Open Cluster Administrator: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**
2. In the console tree (usually the left pane), double-click the **Groups** folder.
3. In the details pane (usually the right pane), double-click the group that contains the resource you want to delete, and then click the resource.
4. On the **File** menu, click **Take Offline**.
5. When the resource is offline, on the **File** menu, click **Delete**.

**To delete a shared directory cluster resource using the command line**

1. Open the command prompt.
2. Type the following:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /offline /wait[:** *Timeout* **]**

3. Then, type:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /delete**

| Argument | Description |
|---|---|
| *ClusterName* | The name of the cluster you want to manage. |
| *ResourceName* | The name of the cluster resource. |
| *Timeout* | The number of seconds **cluster** waits before canceling the command if it does not successfully complete. If **/wait** is specified without this value, **cluster** waits indefinitely or until the resource state changes. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**cluster /?** |

# Changing the directory cache memory setting

Applies To: Windows Server 2003 R2

Server for NFS maintains a cache of directory information that it can search instead of directories on the physical disk. By default, this cache is configured to use 128 kilobytes (KB) of memory per directory that Server for NFS caches. However, you can reduce the amount of memory that the directory cache uses. You can do this by configuring the folder cache size using the command line.

**To configure the folder cache size using the command line**

1. Open the command prompt.
2. Type the following:

   **nfsadmin server [** *remoteMachineName* **] config dircache=** *size*

| Value | Description |
|---|---|
| *remoteMachineName* | The name of the computer you want to configure. |
| *size* | The size in kilobytes (KB) to set the directory cache to. The default size is 128 KB. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**nfsadmin server /?** |

# Securing Server for NFS

Applies To: Windows Server 2003 R2


Allowing root access to resources by NFS clients

Allowing anonymous access to resources by NFS clients

Using Windows security descriptors with Server for NFS

Using DACLs with Server for NFS

Using setuid, setgid, and sticky bits with Server for NFS

# Allowing root access to resources by NFS clients

Applies To: Windows Server 2003 R2

You can allow UNIX users to have root access to resources. Some of these methods include:

- Assigning the client computer root access when the share is exported. See Securing access to the User Name Mapping server.

- Mapping (through either Active Directory Lookup or User Name Mapping) the UNIX user root (UID 0) to the Windows user Administrator. Also, mapping the group to which root belongs to the Windows group Administrators.

For example, consider a situation in which the directory Testdir is owned by Administrators. This directory was created with an account that is a member of the Windows group Administrators. The permissions listed were set after the directory was created. At the UNIX client, the permissions look like this:

```
ls ld /testdir
drwxrwxr-x 2 root root 64 Nov 3 12:11 /testdir/
```

When the directory is owned by Administrators, Server for NFS reports a user identifier (UID) and group identifier (GID) pair of 0,0.

# Allowing anonymous access to resources by NFS clients

Applies To: Windows Server 2003 R2

By default, Server for NFS does not allow anonymous users to access a shared directory. When you share a directory, you can allow anonymous access to the directory and you can change the default anonymous UID and GID values to the UID and GID of any valid UNIX user and group accounts. If you change the anonymous UID and anonymous GID for a share, those values will be used when reporting the owner of a file owned by a Windows user which is not mapped to UNIX user, even if anonymous access is not allowed. For more information about configuring anonymous access for a shared directory, see Allowing anonymous access to an NFS share.

# Using Windows security descriptors with Server for NFS

Applies To: Windows Server 2003 R2

Server for NFS uses the Windows security descriptor when implementing Network File System (NFS) access permissions. The security descriptor is the structure that governs security assignments for Windows. The security descriptor contains the following components:

- File owner
- File group
- Discretionary access control list (DACL)
- System access control list (SACL)

| **Note** |
| --- |
| The SACL is used for auditing and does not affect file permissions. |

Within the security descriptor, the file owner and group refer to security identifiers (SIDs). A SID can be thought of as the internal representation for an individual user or group. The primary reason for using SIDs is to differentiate among accounts across different domains that share the same account name. Even though the names are the same, they represent different accounts and can be given different permissions to the same file.

For more information, see Using DACLs with Server for NFS.

# Using DACLs with Server for NFS

Applies To: Windows Server 2003 R2

The discretionary access control list (DACL) within the security descriptor provides the core of Windows security. The DACL is a list of entries that grant or deny certain rights to specific users or groups. A list entry is called an access control entry (ACE). Each ACE consists of the following:

- A security identifier (SID) to identify a particular user or group.
- An access list that specifies the permissions allowed or denied for the user or group.

The following is an example of a DACL:

- DACL: Mrjones Full Control (All)
- ToolGroup:Read(RX)
- Everyone:Read (RX)

In this DACL, Mrjones has read, write, and execute access to the file. Members of the group ToolGroup have read-and-execute access. Members of the group Everyone (all users) have read-and-execute access.

The following rules govern access to a file:

- If no DACL is present, everyone is granted full access.
- If a DACL is present, but contains no entries, everyone is denied access.
- The file owner always has the ability to change the DACL.

In turn, these rules apply to the DACL:

- DACL entries are searched sequentially.
- All permissions are implicitly denied.
- Once a permission has been denied, it cannot be granted.
- Once a permission has been granted, it cannot be denied.

For more information, see Using Windows security descriptors with Server for NFS.

# Using setuid, setgid, and sticky bits with Server for NFS

Applies To: Windows Server 2003 R2

Users of client computers can use the chmod utility to set the setuid (set-user-identifier-on-execution), setgid (set-group-identifier-on-execution), and sticky file mode bits on files or directories that are stored on an NTFS file system partition and shared through Server for NFS. When the file or directory is subsequently accessed by a UNIX-based client, the standard semantics for these bits will apply. For example, an executable file that has the setuid bit set will execute under the user ID of the file's owner, not the user ID of the user who is executing the file.

Typically, when the setuid or setgid bit is set on a file, the owner or group of the file is changed to the owner or group under whose ID the file will be executed. Unless a user has the right to restore files or directories, Windows security allows a user to take ownership of a file (if the file's permissions allow it), but not to transfer ownership to a second user. Consequently, to use chown or chgrp to change the owner or group of a file to another user or group, you must have the privilege to restore files and directories. By default, this privilege is assigned to members of the Administrators and Backup Operators groups, although it can be assigned to other groups or to individual users. In addition, the account of the user running chown or chgrp and the user or group to whom ownership is being transferred must be properly mapped through Active Directory Lookup or User Name Mapping.

Some UNIX-based Network File System (NFS) servers apply special interpretations or restrictions for the setuid, setgid, and sticky bits. Some versions of UNIX, for example, enforce mandatory locking on a directory with the setgid bit set but no execute permissions. Server for NFS does not implement special interpretations or restrictions on using these bits.

**Security issue**
Typical UNIX setuid and setgid behavior present a known security issue in UNIX. An intruder can exploit setuid and setgid behavior by overwriting an executable file with a Trojan horse, and then running the executable file. The executable file runs with the rights of the owner, instead of running as the intruder.

To prevent this known security issue, by default, setuid and setgid are disabled and will not affect the UID or GID settings if the file or directory has both of the following characteristics:

- One or both of the following bits is set: setgid or setuid.
- The file or directory is group writable, group executable, world writable, or world executable.

Some customers may find this security measure problematic because it is different from the typical behavior of UNIX, although the typical behavior of UNIX setuid and setgid is not specified as part of the Network File System (NFS) protocol.

Procedures

**To change setuid and setgid behavior**

1. Open the Registry Editor.
2. Set the following registry key:

   ```
   HKEY_Local_Machine\System\CurrentControlSet\Services\NfsSvr\Parameters\SafeSetUidGidBits = (DWORD)
   ```

   - A value of 1 causes the bits the safer setuid and setgid behavior.
   - A value of 0 causes the standard UNIX behavior.

# Optimizing Server for NFS performance

Applies To: Windows Server 2003 R2


Adding performance counters

Monitoring and tuning performance

Changing the directory cache memory setting

# Adding performance counters

Applies To: Windows Server 2003 R2

**To add performance counters**

1. Open **Performance**: click **Start**, point to **Programs** or **All Programs**, point to **Control Panel**, point to **Administrative Tools**, and then click **Performance**.
2. Click the plus-sign (+) button.
3. In the **Performance** object list, click **Server for NFS**, and then do one of the following:
   - Click **All counters**, and then click **Add**.
   - Or, click **Select counters from list**, click the counters you want to add, and then click **Add**.

| **Note** |
| --- |
| There is no command-line method for this procedure. |

# Monitoring and tuning performance

Applies To: Windows Server 2003 R2

Server for NFS includes several tools for monitoring Network File System (NFS) use and for probing the network for other NFS servers and resources. These tools include Windows System Monitor, rpcinfo, showmount, and nfsstat.

You can use Windows System Monitor to display NFS statistics, including the number and type of requests received and the number of bytes read and written.

You can use the **rpcinfo** utility to:

- List all programs registered with the port mapper on a local or remote host (**rpcinfo -p**).
- Find all network nodes that have a specified program and version registered with the port mapper (**rpcinfo -b**). You must specify a program name and number and version.
- Use the Transmission Control Protocol (TCP) to call the specified program (**rpcinfo -t**). You must specify a node name and program name. If a version is not specified, the program calls all versions.
- Use the User Datagram Protocol (UDP) to call the specified program (**rpcinfo -u**). You must specify a node name and program name. If a version is not specified, the program calls all versions.

You can use the **showmount** utility to:

- List all file systems exported on the specified NFS server (**showmount -e**).
- List all remotely mounted clients and the directories each has mounted on the specified NFS server (**showmount -a**).
- List all directories remotely mounted on the specified NFS server (**showmount -d**).

You can use the **nfsstat** utility to display the number of NFS v2, NFS v3, and Mount v3 calls made to the server since the counters were set to 0, either when the service started or when the counters were reset using **nfsstat –z**.

# Using file systems with NFS

Applies To: Windows Server 2003 R2


Understanding NFS and file names

Using CDFS with NFS

Using NTFS file systems with NFS

# Understanding NFS and file names

Applies To: Windows Server 2003 R2

Because Network File System (NFS) has its origins on UNIX-based computers, the conventions used for naming files and directories on NFS file systems differ in several respects from those used by Windows. This topic covers three of these differences and describes how you can configure Server for NFS to deal appropriately with them. These differences are:

- Special characters
- Case sensitivity
- Hidden files

In addition, this topic also discusses directory- and file-name encoding for client computers localized for East Asian languages.

**Special characters**
Every file system has characters that cannot be used in directory and file names, usually because those characters are reserved for use in path and command syntax. For example, a file stored on the NTFS file system cannot contain a colon (:) character because that character is used to separate the drive letter from the rest of the path in a fully qualified path name. UNIX file systems do not use letters to identify drives, so this syntax is not used, and consequently, colon characters can appear in UNIX file names.

Server for NFS stores files on Windows-based file systems on behalf of NFS clients, which assume UNIX file system conventions. Because of this, you can use Server for NFS to specify how characters in directory and file names will be converted (mapped) when files are moved between Server for NFS and client computers. To accomplish this, you must create a file that defines how these characters are to be mapped, and then specify where Server for NFS can find the file containing these character maps. For information about creating the character map file, see Overview of file-name character translation. To learn how to configure Server for NFS to use this file, see Configuring file-name character translation.

**Case sensitivity**
Windows and UNIX computers typically differ in the way they use the case of letters in directory and file names.

UNIX computers are usually case sensitive when handling directory and file names. That is, a UNIX computer will distinguish between a file named MYFILE and a file named MyFile, allowing both to coexist in the same directory and accessing them only when the case used in the file name matches the stored file name exactly.

Windows-based computers, on the other hand, are case preserving, but not case sensitive. That is, they will allow a file to be named MyFile or MYFILE and will preserve the case difference, but will not allow two files with the same name, differing only in letter case, to be stored in the same directory. Nor do they use case as a criterion when locating files. In other words, either MyFile or MYFILE can exist in the same directory, but not both, and either file can be accessed using the name MYFILE or MyFile.

By default, Server for NFS is case sensitive when matching directory-access and file-access requests to directories and files in its shared directories, and it preserves the case of the names of directories and files when creating files on behalf of NFS client computers. You can instruct Server for NFS not to be case sensitive when returning directory and file names. For more information, see Configuring case sensitivity for file and folder names.

**Hidden files**
Files can be hidden (that is, not displayed when the directory they are store in is listed) on both Windows and UNIX computers. In the case of Windows-based computers, files are hidden by setting a special attribute for the file. UNIX file systems simply do not list files whose names begin with a period (.) character. By default, Server for NFS does not create files with names beginning with a period as hidden files in the Windows file system.

**Directory- and file-name encoding**
Client computers that are localized for an East Asian language require special support by NFS servers to be able to use characters in that language when creating directory and file names. This support, called encoding, determines how two-byte character codes exchanged between the server and a client computer will be interpreted. You can configure an NFS share to provide a default language-specific encoding scheme for all clients that access the share; you can also specify an encoding scheme for each client in the permissions list of the share. For more information, see Setting default encoding for an NFS share and Creating an NFS share.

# Using CDFS with NFS

Applies To: Windows Server 2003 R2

Unlike the NTFS file system, the CD-ROM file system (CDFS) does not provide any type of access control. The mount permissions granted by the Network File System (NFS) server administrator govern all access control. The CDFS file system always shares data as a read-only resource, regardless of the type of share permission you give it.

The characteristics of the CDFS file system include the following:

- Because this file system was designed around single-user computers, there is no such thing as a file owner.
- Just as no file owner exists, neither does a file group.
- There is no support for file links. This file system uses a single-name/single-file format with no support for multiple directory entries referring to the same file.

Because of these limitations, the following conventions are a part of the Server for NFS implementation:

- The file owner and group are always returned as the owner and group given in the UNIX verification of the request. This means that every individual user is reported as the owner of all files. Having different users of the same file both reported as owners can seem confusing at first, but it is in keeping with the concept that anyone with access to the file system owns it. Any attempt to modify the owner or group of a file fails. If no UNIX verification is passed, the anonymous user identifier (UID) and the anonymous group identifier (GID) are reported. By default, the anonymous UID and anonymous GID are both –2 (65534 when used as an unsigned 2-byte integer), although these values can be changed for a share.
- Permissions are reported as read-and-execute for owner, group, and other (in UNIX notation:

  ```
  -r-xr-xr-x
  ```

Any attempt to modify data or attributes on the shared resource fails.

# Using NTFS file systems with NFS

Applies To: Windows Server 2003 R2

Under UNIX, a file has a single owner and a single group, with certain permissions (read, write, and/or execute) granted to each. Separate permissions are also granted for others, meaning any other users who are not members of the group that owns the file.

The NTFS file system, however, allows permissions to be assigned to many users or groups independently. For example, an NTFS file can specify that User A is to have read-and-write access, while User B is to have only read access. The same file can also have certain permissions for one or more groups. The objective of Server for NFS is to provide a consistent, intuitive mapping of file permissions that UNIX clients can work with, and that users can understand.

# Managing NFS shares

Applies To: Windows Server 2003 R2


Accessing a Windows NFS share from a UNIX client

Overview of file-name character translation

Specifying the NFS protocol to use

Creating an NFS share

Setting default encoding for an NFS share

Adding clients and groups to the NFS share permissions list

Configuring access to an NFS share for a client computer or group

Allowing anonymous access to an NFS share

Configuring DFS shares for NFS

Configuring shared folders on a server cluster

Resolving file permission conflicts

Enabling advisory file locking

# Accessing a Windows NFS share from a UNIX client

Applies To: Windows Server 2003 R2

**To access a Windows NFS share from a UNIX client**

1. Open the UNIX command prompt.
2. Create a directory that will act as the mount point for the Windows Network File System (NFS) share.
3. Type the following command:

   **mount [-f nfs] [-o** *options* **]** *remote_host_name* **:/** *directory mount_point*

| Argument | Description |
|---|---|
| *options* | Options include **hard** for hard mounts, **soft** for soft mounts, and more. For more information, refer to the UNIX man pages. |
| *remote_host_name* **:/** *directory* | The name of the computer and directory to mount. |
| *mount_point* | The local directory that you created which will act as a mount point. |

| **Note** |
|---|
| To view the complete syntax for this command, at the UNIX command prompt, type: **man mount** |

# Overview of file-name character translation

Applies To: Windows Server 2003 R2

Although Windows and UNIX file systems do not allow certain characters in file names, the characters that are prohibited by each operating system are not the same. For example, a valid Windows file name cannot contain a colon (:), but a valid UNIX file name can. If a UNIX user attempts to create a file in a Network File System (NFS) share on Server for NFS, and that file contains an illegal character in its name, the attempt will fail.

To prevent this problem, use file-name character mapping to replace characters that are not allowed in a file system. To enable file-name character translation, create a text file that maps Windows to UNIX characters, and then configure Server for NFS to specify the path and name of the translation file.

A file-name character translation file is simply a list of mapped characters in the format:

**0x** *nn* **0x** *nn* **: 0x** *nn* **0x** *nn* **[ ;** *comment* **]**

where *nn* is the hexadecimal value of one byte of a two-byte Unicode character, and *comment* is an optional comment, which is preceded by a semicolon (;). Everything from the semicolon to the end of the line is ignored. The first character in the entry is the character on the client (UNIX) computer, while the second is the character used on the Windows-based Server for NFS computer.

For example, the following maps the UNIX colon (:) to a Windows dash (-):

**0x00 0x3a : 0x00 0x2d ; replace client : with - on server**

If this entry appears in the file-name character translation file on a Server for NFS computer, when an NFS client creates a file named re:salesquotas on a Server for NFS share, Server for NFS names the file re-salesquotas. This is the name that appears in Windows Explorer when local or remote Windows users display the contents of the shared folder. NFS clients that list the contents of the shared directory, however, see the file name as re:salesquotas. Character translation takes place for all files shared by the server, regardless of whether they were created by an NFS client.

Do not map a given character to more than one character. Mapping a character to two or more different characters, as shown in the following example, will produce unexpected results:

**0x00 0x11 : 0x00 0x22**

**0x00 0x11 : 0x00 0x33**

Do not map a period (.) because it is used in the file-name syntax of both Windows and UNIX file systems.

After you have created the file-name character translation file, you must specify its location.

| **Important** |
| --- |
| To prevent unauthorized users from tampering with the file-name character translation file and causing Server for NFS to fail, ensure that the file is protected with a discretionary access control list (DACL) that grants Full Control to the built-in System account and to the Administrators group. The DACL should contain no other entries. |

# Specifying the NFS protocol to use

Applies To: Windows Server 2003 R2

Network File System (NFS) is a network protocol that, in its most widely adopted forms, is defined in RFC 1094 for NFS version 2 and in RFC 1813 for NFS version 3. By default, Server for NFS supports clients using either version. You can override this, however, and configure Server for NFS to allow access only by version 2 client computers. See Enabling NFS version 3.

NFS relies on two transport protocols: Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). Server for NFS supports both protocols. You can specify the transport protocol supported by Server for NFS. For more information about this,

see Specifying Server for NFS TCP and UDP settings.

# Creating an NFS share

Applies To: Windows Server 2003 R2

**To create a share using the Windows interface**

1. Open Windows Explorer: click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
2. In the details pane, right-click the shared directory you want to manage.
3. Click **NFS Sharing**.
4. Select **Share this Folder**.
5. Type a name for the share in the **Share Name** text box.
6. Click **Apply**.

**To create a share using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsshare ShareName=** *drive:path*

| Argument | Description |
|---|---|
| *sharename* | The name of the directory to share. |

| Note |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>*nfsshare /?*<br>This command only works with Windows Server 2003 operating systems. |

# Setting default encoding for an NFS share

Applies To: Windows Server 2003 R2

**To set default encoding for a shared directory using the Windows interface**

1. Open Windows Explorer: click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
2. In the details pane, right-click the shared directory you want to manage.
3. Click **Sharing**.
4. Click **NFS Sharing**.
5. In the **Encoding** list, select the type of encoding to be used for file and directory names in the shared directory. In addition to the default American National Standards Institute (ANSI) encoding, the following encoding is available:
   - BIG5 (Chinese)
   - EUC-JP (Japanese)
   - EUC-KR (Korean)
   - EUC-TW (Chinese)
   - GB2312-80 (Simplified Chinese)
   - KSC5601 (Korean)
   - SHIFT-JIS (Japanese)
6. If this option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:
   - Japanese: SHIFT-JIS
   - Korean: KS_C_5601-1987
   - Simplified Chinese: GB
   - Traditional Chinese: BIG5

**To set default encoding for a shared directory using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsshare -o encoding={big5|euc-jp|euc-kr|euc-tw|gb2312-80|ksc5601|shift-jis}***sharename*

| Argument | Description |
|---|---|
| *sharename* | The name of the directory to share. |

| | |
|---|---|
| **-o encoding=** *type* | Specifies the default encoding. You must specify one of the following types:<br>• **big5** (Chinese)<br>• **euc-jp** (Japanese)<br>• **euc-kr** (Korean)<br>• **euc-tw** (Chinese)<br>• **gb2312-80** (Simplified Chinese)<br>• **ksc5601** (Korean)<br>• **shift-jis** (Japanese)<br><br>If you do not specify an encoding type, the default encoding scheme is ANSI or, on systems configured for non-English locales, the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:<br>• Japanese: SHIFT-JIS<br>• Korean: KS_C_5601-1987<br>• Simplified Chinese: GB2312-80<br>• Traditional Chinese: BIG5 |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsshare /?** |

# Adding clients and groups to the NFS share permissions list

Applies To: Windows Server 2003 R2

**To add clients and groups to the NFS share permissions list using the Windows interface**

1. Open Windows Explorer: click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
2. In the details pane, right-click the shared directory you want to manage.
3. Click **Sharing**.
4. Click **NFS Sharing**.
5. If the directory is not already shared, click **Share this folder**, and then, in the **Share name** box, type a name for the share.
6. Click **Permissions**, click **Add**, and then do one of the following:
7. Do either of the following:
   - In the **Names** list, click the clients and groups you want to add, and then click Add.
   - Or, in the **Add Names** box, type the names of clients or groups you want to add, separating names in the list with a semicolon (;).
8. To see a list of the members of a group, in the **Names** list, click a group, click **Members**.
9. In the **Type of Access** list, click the type of access you want to allow the selected clients and groups.
10. Select **Allow Root Access** if you want a user identified as root to have access other than as an anonymous user. By default, the user identifier (UID) root user is coerced to the anonymous UID.
11. In the **Encoding** list, click the type of directory- and file-name encoding to be provided to the selected clients and groups. In addition to the default American National Standards Institute (ANSI) encoding, the following encoding is available:
    - BIG5 (Chinese)
    - EUC-JP (Japanese)
    - EUC-KR (Korean)
    - EUC-TW (Chinese)
    - GB2312-80 (Simplified Chinese)
    - KSC5601 (Korean)
    - SHIFT-JIS (Japanese)

| **Note** |
| --- |
| If this option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales: |

- Japanese: SHIFT-JIS
- Korean: KS_C_5601-1987

- Simplified Chinese: GB
- Traditional Chinese: BIG5

12. To close the dialog box, click **OK**.

**To add clients and groups to the NFS share permissions list using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsshare -o***option***[=***host***[:***host***]...]***sharename*

| Argument | Description |
|---|---|
| *sharename* | The name of the shared directory. |
| **-o***option* | Specifies one or more options to set and must be set to at least one of the following:<br>• **big5** (Chinese)<br>• **euc-jp** (Japanese)<br>• **euc-kr** (Korean)<br>• **euc-tw** (Chinese)<br>• **gb2312-80** (Simplified Chinese)<br>• **ksc5601** (Korean)<br>• **shift-jis** (Japanese)<br><br>If this option is not set, the default encoding scheme is used.<br>**noroot[=** *host* **[:** *host* **]...]**<br>Denies root access to the shared directory by the clients or client groups specified by host. Separate client and group names with a colon (:). If you do not specify a host, root access is denied to all clients and groups for which root access is not explicitly granted.<br>**ro[=** *host* **[:** *host* **]...]**<br>Provides read-only access to the shared directory by the clients or client groups specified by host. Separate client and group names with a colon (:). If you do not specify a host, read-only access is granted to all clients and groups for which access is not explicitly granted or denied.<br>**root[=** *host* **[:** *host* **]...]**<br>Provides root access to the shared directory by the clients or client groups specified by host. Separate client and group names with a colon (:). If you do not specify a host, root access is granted to all clients and groups for which root access is not explicitly denied.<br>**rw[=** *host* **[:** *host* **]...]**<br>Provides read/write access to the shared directory by the clients or client groups specified by host. Separate client and group names with a colon (:). If you do not specify a host, read/write access is granted to all clients and groups for which access is not explicitly granted or denied. |

**Note**

To view the complete syntax for this command, at a command prompt, type:

**nfsshare /?**

By default, when a directory is shared, all client computers are granted read-only access, and anonymous access is not allowed.

# Configuring access to an NFS share for a client computer or group

Applies To: Windows Server 2003 R2

**To configure access to an NFS share for a client or group using the Windows interface**

1. Open Windows Explorer: click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
2. In the details pane, right-click the shared directory you want to manage.
3. Click **Sharing**.
4. Click **NFS Sharing**, and then, if the folder has not already been shared, click **Share this folder**.
5. Click **Permissions**.
6. Click the client or group for which you want to set permissions.
7. In the **Type of Access** list, click the type of access you want to allow the client or group.
8. Select **Allow Root Access** if you want a user identified as root to have access other than as an anonymous user. By default, the user identifier (UID) root user is coerced to the anonymous UID.
9. In the **Encoding** list, click the type of directory- and file-name encoding to be provided to the client or group. In addition to the default American National Standards Institute (ANSI) encoding, the following encoding is available:

   - BIG5 (Chinese)
   - EUC-JP (Japanese)
   - EUC-KR (Korean)
   - EUC-TW (Chinese)
   - GB2312-80 (Simplified Chinese)
   - KSC5601 (Korean)
   - SHIFT-JIS (Japanese)

   | **Note** |
   | --- |
   | If this option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales: |

   - Japanese: SHIFT-JIS
   - Korean: KS_C_5601-1987
   - Simplified Chinese: GB
   - Traditional Chinese: BIG5

10. Click **OK** twice.

**To configure access to an NFS share for a client or group using the command line**

1. Open the command prompt.

2. At the command prompt, type:

   **nfsshare -o***option***[=***host***[:***host***]...]***sharename*

| Argument | Description |
|---|---|
| *sharename* | The name of the shared directory. |
| **-o** | Specifies one or more options to set and must be set to at least one of the following:<br>**{big5\|euc-jp\|euc-kr\|euc-tw\|gb2312-80\|ksc5601\|shift-jis}=** *host* **[:** *host* **]...**<br>Specifies the file- and directory-name encoding for individual client computers or client groups. The following encoding schemes are available:<br>• **big5** (Chinese)<br>• **euc-jp** (Japanese)<br>• **euc-kr** (Korean)<br>• **euc-tw** (Chinese)<br>• **gb2312-80** (Simplified Chinese)<br>• **ksc5601** (Korean)<br>• **shift-jis** (Japanese)<br>If this option is not set, the default encoding scheme is used.<br>**na**<br>Specifies that no client or group is granted access to the NFS share unless you explicitly grant it access.<br>**noroot[=** *host* **[:** *host* **]...]**<br>Denies root access to the shared directory by the clients or client groups specified by host. Separate client and group names with a colon (:). If host is not specified, root access is denied to all clients and groups for which root access is not explicitly granted.<br>**removeclient[=** *host* **[:** *host* **]...]**<br>Removes the specified clients from the list of permissions of the shared directory. Separate client and group names with a colon (:).<br>**ro[=** *host* **[:** *host* **]...]**<br>Provides read-only access to the shared directory by the clients or client groups specified by host. Separate client and group names with a colon (:). If host is not specified, read-only access is granted to all clients and groups for which access is not explicitly granted or denied.<br>**root[=** *host* **[:** *host* **]...]**<br>Provides root access to the shared directory by the clients or client groups specified by host. Separate client and group names with a colon (:). If host is not specified, root access is granted to all clients and groups for which root access is not explicitly denied.<br>**rw[=** *host* **[:** *host* **]...]**<br>Provides read/write access to the shared directory by the clients or client groups specified by host. Separate client and group names with a colon (:). If host is not specified, read/write access is granted to all clients and groups for which access is not explicitly granted or denied. |

**Note**

To view the complete syntax for this command, at a command prompt, type:

**nfsshare /?**

By default, when a directory is shared, all client computers are granted read-only access, and anonymous access is not allowed.

# Allowing anonymous access to an NFS share

Applies To: Windows Server 2003 R2

**To allow anonymous access to an NFS share using the Windows interface**

1. Open Windows Explorer: click **Start**, point to **Programs** or **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
2. In the details pane, right-click the shared directory you want to manage.
3. Click **Sharing**.
4. Click **NFS Sharing**.
5. Select **Allow anonymous access**.
6. To specify a nondefault value for the anonymous user identifier (UID) or anonymous group identifier (GID), type the value in the **Anonymous UID** or **Anonymous GID** box.
7. Click **Apply**.

| **Note** |
|---|
| The anonymous UID and anonymous GID values for a share will be used when reporting the owner of a file owned by an unmapped user, even if Allow Anonymous Access is disabled.<br><br>In Windows Server 2003 operating systems, the Everyone group does not include anonymous users by default. For information about providing anonymous access, see Allowing anonymous access to resources by NFS clients. |

**To allow anonymous access to an NFS share using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsshare -o***option*=*value*...*sharename*

| Argument | Description |
|---|---|
| *sharename* | The name of the shared directory. |
| **-o** | Specifies one or more options to be set and must be set to at least one of the following:<br><br>**anon={yes \| no}**<br><br>Specifies whether anonymous (unmapped) users can access the shared directory. The default is no.<br><br>**anonuid=** *uid*<br><br>Specifies that anonymous (unmapped) users will access the share directory using uid as their user identifier (UID). The default is -2. The anonymous UID will be used when reporting the owner of a file owned by an unmapped user, even if anonymous access is disabled.<br><br>**anongid=** *gid*<br><br>Specifies that anonymous (unmapped) users will access the share directory using gid as their group identifier (GID). The default is -2. The anonymous GID will be used when reporting the owner of a file owned by an unmapped user, even if anonymous access is disabled. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsshare /?**<br><br>In Windows Server 2003 operating systems, the Everyone group does not include anonymous users by default. For information about providing anonymous access, see Allowing anonymous access to resources by NFS clients. |

# Configuring DFS shares for NFS

Applies To: Windows Server 2003 R2

You can use Distributed File System (DFS) to connect to a Network File System (NFS) share.

**To configure DFS shares for NFS**

1. Set up an NFS server and create an NFS share.
2. Install Client for NFS on the computer that is running DFS.
3. On the DFS server, add a link in DFS that points to the share on the NFS server.
4. To access the NFS share from any computer, connect to the DFS server as you would normally connect to any Windows file server.
5. Browse to the link you created.

# Configuring shared folders on a server cluster

Applies To: Windows Server 2003 R2


Configuring an NFS share as a cluster resource

Modifying an NFS shared cluster resource

Deleting an NFS shared cluster resource

# Configuring an NFS share as a cluster resource

Applies To: Windows Server 2003 R2

**To share a directory as a cluster resource using the Windows interface**

1. Open Cluster Administrator: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.
2. In the console tree (usually the left pane), double-click the **Groups** folder.
3. In the details pane (usually the right pane), click the group to which you want the resource to belong.
4. On the **File** menu, point to **New**, and then click **Resource**.
5. In the New Resource wizard, type the appropriate information in **Name** and **Description**, and in **Resource type**, click **NFS Sharing**. Click **Next**.
6. Add or remove possible owners of the resource, and then click **Next**.
7. To add or remove dependencies, do one of the following:
   - To add dependencies, under **Available resources**, click a resource, and then click **Add**.
   - To remove dependencies, under **Resource dependencies**, click a resource, and then click **Remove**.
8. In **Share Name**, type the name of the shared folder.
9. In **Share Path**, type the fully qualified path of the folder to be shared.
10. Do one of the following:
    - To share the folder and its contents as a single shared folder, click **Share Root Only**.
    - To share all folders (except hidden folders) in the folder specified in the previous step as separate shared folders, click **Share Subdirectories Only**.
11. To allow access by users who have not been authenticated, click **Allow anonymous access**.
12. To change the user identifier (UID) and group identifier (GID) values assigned to anonymous users, type the new values in the **Anonymous UID** and **Anonymous GID** boxes.
13. In the **Encoding** list, select the type of encoding to be used for file and directory names in the shared directory. In addition to the default American National Standards Institute (ANSI) encoding, the following encoding is available:
    - BIG5 (Chinese)
    - EUC-JP (Japanese)
    - EUC-KR (Korean)
    - EUC-TW (Chinese)
    - GB2312-80 (Simplified Chinese)
    - KSC5601 (Korean)
    - SHIFT-JIS (Japanese)

    **Note**

| |
|---|
| If this option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales: |

- Japanese: SHIFT-JIS
- Korean: KS_C_5601-1987
- Simplified Chinese: GB2312-80
- Traditional Chinese: BIG5

14. To specify permissions for the client computers that can access the shared folder, click **Permissions**.
15. Click **Finish**.

| **Important** |
|---|
| Before adding a resource to your cluster, you must verify the following: |

- The NFS Share resource type is available. This resource type was created when you installed Server for NFS on the cluster node.
- A group already exists within the cluster to which your resource will belong.
- All dependent resources have been created.

| **Note** |
|---|
| In Windows Server 2003 operating systems, the Everyone group does not include anonymous users by default. For information about providing anonymous access when Server for NFS is running on these operating systems, see Allowing anonymous access to resources by NFS clients. |

**To share a directory as a cluster resource using the command line**

1. Open the command prompt.
2. Type the following:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /create /restype:"NFS Share" /group:"** *GroupName* **" /privproperties ShareName="** *ShareName* **":STRING Path="** *SharePath* **":STRING ShareSubDirs=** *ShareType* **:DWORD GlobalPerm=** *PermType* **:DWORD Encoding=** *EncodingType* **:DWORD AnonymousAccessAllowed=** *AccessType* **:DWORD**

3. Then, type:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /online /wait[:** *Timeout* **]**

| Argument | Description |
|---|---|
| *ClusterName* | The name of the cluster you want to manage. |
| *ResourceName* | The name of the cluster resource. |
| *GroupName* | The name of the group to which you want to add the resource. |
| *ShareName* | The name of the shared directory you want to create. This value is ignored if ShareType is **0x1** |

| | |
|---|---|
| *SharePath* | The fully qualified path of the directory to be shared. |
| *ShareType* | **0x0** if only the directory specified by SharePath is to be shared. **0x1** if the subdirectories (except hidden subdirectories) in the directory specified by SharePath are to be shared individually. |
| *PermType* | Specifies root-access and read-write permission combinations. PermType must be one of the following values:<br>**0x1**—No access<br>**0x9**—No Access + Root<br>**0x2**—Read only<br>**0xa**—Read only + Root<br>**0x4**—Read/write<br>**0xc**—Read/write + Root |
| *EncodingType* | Specifies the default encoding used for file and directory names and, if used, must be set to one of the following:<br>• **0x7**—ANSI<br>• **0x5**—BIG5 (Chinese)<br>• **0x0**—EUC-JP (Japanese)<br>• **0x1**—EUC-KR (Korean)<br>• **0x2**—EUC-TW (Chinese)<br>• **0x6**—GB2312-80 (Simplified Chinese)<br>• **0x4**—KS_C_5601-1987 (Korean)<br>• **0x3**—SHIFT-JS (Japanese)<br>If this is option is set to 0x7, the encoding scheme is ANSI or, on systems configured for non-English locales, the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:<br>• Japanese: SHIFT-JIS<br>• Korean: KS_C_5601-1987<br>• Simplified Chinese: GB2312-80<br>• Traditional Chinese: BIG5 |
| *AccessType* | **0x0** if anonymous access to the shared directory is not allowed. **0x1** if anonymous access is allowed. |
| *Timeout* | The number of seconds **cluster** waits before canceling the command if it does not successfully complete. If **/wait** is specified without this value, **cluster** waits indefinitely or until the resource state changes. |

**Important**

Before adding a resource to your cluster, you must verify the following:

- The **NFS Share** resource type is available. This resource type was created when you installed Server for NFS on the cluster node.
- A group already exists within the cluster to which your resource will belong.
- All dependent resources have been created.

**Note**

In Windows Server 2003 operating systems, the Everyone group does not include anonymous users by default. For information about providing anonymous access when Server for NFS is running on these operating systems, see Allowing anonymous access to resources by NFS clients.

To view the complete syntax for this command, at a command prompt, type:

**cluster /?**

All arguments listed above are required. Using default values might produce unexpected results.

You must use the Windows interface to set permissions on shared directories.

If the attempt to share the directory fails, **cluster** displays an error message, and detailed information about the failure is recorded in the cluster log.

You cannot set the user identifier (UID) and group identifier (GID) values assigned to anonymous users using the command line.

# Modifying an NFS shared cluster resource

Applies To: Windows Server 2003 R2

**To modify a shared directory cluster resource using the Windows interface**

1. Open Cluster Administrator: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.
2. In the console tree (usually the left pane), double-click the **Groups** folder.
3. In the details pane (usually the right pane), double-click the group that contains the resource you want to modify, and then click the resource.
4. On the **File** menu, click **Take Offline**.
5. When the resource is offline, on the **File** menu, click **Properties**.
6. Change the properties you want to modify, and then click **OK**.
7. On the **File** menu, click **Bring Online**.

**To modify a shared directory cluster resource using the command line**

1. Open the command prompt.
2. Type the following:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /offline /wait[:** *Timeout* **]**
3. Then, type:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /** *Option*
4. Then, type:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /online /wait[:** *Timeout* **]**

| Argument | Description |
|---|---|
| *ClusterName* | The name of the cluster you want to manage. |
| *ResourceName* | The name of the cluster resource. |
| *Timeout* | The number of seconds **cluster** waits before canceling the command if it does not successfully complete. If **/wait** is specified without this value, **cluster** waits indefinitely or until the resource state changes. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**cluster /?** |

# Deleting an NFS shared cluster resource

Applies To: Windows Server 2003 R2

**To delete a shared directory cluster resource using the Windows interface**

1. Open Cluster Administrator: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.
2. In the console tree, double-click the **Groups** folder.
3. In the details pane, double-click the group that contains the resource you want to delete, and then click the resource.
4. On the **File** menu, click **Take Offline**.
5. When the resource is offline, on the **File** menu, click **Delete**.

**To delete a shared directory cluster resource using the command line**

1. Open the command prompt.
2. Type the following:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /offline /wait[:** *Timeout* **]**

3. Then, type:

   **cluster [** *ClusterName* **] resource "** *ResourceName* **" /delete**

| Argument | Description |
|---|---|
| *ClusterName* | The name of the cluster you want to manage. |
| *ResourceName* | The name of the cluster resource. |
| *Timeout* | The number of seconds **cluster** waits before canceling the command if it does not successfully complete. If **/wait** is specified without this value, **cluster** waits indefinitely or until the resource state changes. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**cluster /?** |

# Resolving file permission conflicts

Applies To: Windows Server 2003 R2

UNIX uses three levels of file permission settings: user permissions, group permissions, and permissions for everyone else. Additionally, root has full permissions to all files, regardless of the file settings.

In Windows, the administrator can access all files, but cannot modify the files without first modifying the permissions. Consequentially, automated activities such as virus checking can change the UNIX permissions and potentially cause problems for UNIX clients.

To resolve this problem, you must modify the KeepInheritance registry value, add the Windows group to the share, and configure Access Control Entries (ACEs) to apply to newly created directories.

**To modify the KeepInheritance registry value**

1. Open the Registry Editor: click **Start**, and then click **Run**. Then, enter **regedit** and click **OK**.
2. Set the following registry key to **1**:

   `HKLM\Software\Microsoft\Server for NFS\CurrentVersion\Mapping\KeepInheritance`

   | **Note** |
   | --- |
   | Setting this registry key can affect how permissions appear in UNIX. |

**To configure ACEs to apply to newly created directories**

1. Open Windows Explorer: click **Start**, point to **Programs or All Programs**, point to **Accessories**, and then click **Windows Explorer**.
2. Right-click a folder and select **Properties**.
3. Click the **Security** tab.
4. Click **Advanced**.
5. Select the **Allow inheritable permissions** check box.
6. Click **OK**.

# Enabling advisory file locking

Applies To: Windows Server 2003 R2

In Windows, file locking is mandatory. In UNIX, file locking is advisory.

By default, Server for NFS file shares use mandatory file locking. This is generally not a problem, but can affect some UNIX applications.

To configure Windows to use advisory file locking, you must clear the EnableSMBLocking registry key.

**To modify the Advisory Locking registry value**

1. Open the Registry Editor.
2. Set the following registry key to **0**:

   `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NfsSvr\NlmNsm\EnableSMBLocking`

| **Caution** |
| --- |
| Setting this registry value turns off mandatory locks only for Network File System (NFS) lock requests. Windows applications continue to use mandatory locks. As a result, if one NFS client obtains a read/write lock on a file, then other NFS clients that are also using NLM locks before accessing a file for read/write will honor that lock. However, any Windows application can still modify the file, either through a server message block (SMB) share, or when running on the computer where the file is located, because advisory locks do not actually lock the file on the server. Additionally, although an NFS client may have obtained a lock on a file, a Windows application (either running locally or accessing the file through a CIFS share) can get a mandatory lock on the file, thereby preventing the NFS client from accessing the file, even though the NFS client has a read/write lock on the file. |

# Managing NFS client groups

Applies To: Windows Server 2003 R2


Creating an NFS client group

Renaming an NFS client group

Deleting an NFS client group

Adding an NFS client computer to an NFS client group

Deleting an NFS client computer from an NFS client group

# Creating an NFS client group

Applies To: Windows Server 2003 R2

**To create an NFS client group using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server [** *ComputerName* **] creategroup** *GroupName*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *GroupName* | The name of the group to create. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**nfsadmin server /?**<br>You cannot give a group the same name as a client in the group. |

# Renaming an NFS client group

Applies To: Windows Server 2003 R2

**To rename an NFS client group using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server [** *ComputerName* **] renamegroup** *OldName NewName*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *OldName* | The current name of the group you want to change. |
| *NewName* | The new name for the group. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**nfsadmin server /?**<br>You cannot give a group the same name as a client in the group. |

# Deleting an NFS client group

Applies To: Windows Server 2003 R2

**To delete an NFS client group using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server [** *ComputerName* **] deletegroup** *GroupName*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *GroupName* | The name of the group to delete. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**nfsadmin server /?** |

# Adding an NFS client computer to an NFS client group

Applies To: Windows Server 2003 R2

**To add an NFS client computer to an NFS client group using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server [** *ComputerName* **] addmembers***GroupName Member***[,***Member***...]**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *GroupName* | The name of the group to which you want to add members. |
| *Member* | The host name or IP address of the member you want to add. |

**Note**

To view the complete syntax for this command, at a command prompt, type: **nfsadmin server /?**

You cannot give a group the same name as a client in the group.

You can add a range of IP addresses as a group member, for example, 127.0.0.*. However, you cannot remove a single computer in the range specified in this manner. You must remove the entire range.

# Deleting an NFS client computer from an NFS client group

Applies To: Windows Server 2003 R2

**To delete an NFS client computer from an NFS client group using the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **nfsadmin server [** *ComputerName* **] deletemembers***GroupNameMember***[,***Member***...]**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *GroupName* | The name of the group from which you want to delete members. |
| *Member* | The host name or IP address of the member you want to delete. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsadmin server /?** |

# Using Microsoft Services for NFS with server clusters

Applies To: Windows Server 2003 R2


Understanding how Server for NFS works with server clusters

Using Server for NFS on a server cluster

# Understanding how Server for NFS works with server clusters

Applies To: Windows Server 2003 R2

A server cluster is a group of independent computer systems, known as nodes, that work together as a single system to ensure that mission-critical applications and resources remain available to clients. Every node is attached to one or more cluster storage devices. Clustering allows users and administrators to access and manage the nodes as a single system rather than as separate computers.

Server for NFS is specifically designed to make use of the server cluster. Because of this, when Server for NFS is installed on a server cluster, you can share a directory as a cluster resource. This ensures that if the node sharing the directory fails, another node in the cluster will resume sharing the directory until the failed node is restored. Running Server for NFS on a server cluster also allows you to provide static load balancing by distributing file shares among two or more nodes in the cluster. For more information about Windows clustering, see "Server Clusters" in Windows Help.

# Using Server for NFS on a server cluster

Applies To: Windows Server 2003 R2

Because Server for NFS is specifically designed to make use of server clusters, you can use Cluster Administrator to configure Network File System (NFS) shared directories as cluster resources. This provides for high availability of those shared directories. This ensures that if the cluster node sharing the directory should be shut down or fail, another node in the cluster will assume sharing the NFS shared directory until the first node restarts. In addition to enhancing availability, distributing NFS shares among nodes in the cluster can provide for static load balancing among several servers. For more information about sharing NFS directories on a server cluster, see Configuring shared folders on a server cluster.

Because the NFS shared directory is the cluster resource (instead of the Server for NFS service itself), ownership of individual NFS shares can move easily from one node to the other, as circumstances require. However, configuration changes to Server for NFS (such as the name of the User Name Mapping server) on one node are automatically replicated to instances of Server for NFS running on other cluster nodes. This simplifies administration of multiple servers because you only have to configure Server for NFS on one computer.

| **Note** |
| --- |
| To ensure proper operation of Server for NFS in a server cluster, when stopping the server cluster, first stop Server for NFS, then stop the server cluster. |

# Server for NFS Authentication

Applies To: Windows Server 2003 R2

When a UNIX user attempts to access a file shared by Server for NFS, Server for NFS uses either Active Directory Lookup or User Name Mapping to obtain the corresponding Windows user name of that UNIX user. When the Windows user name is obtained, Server for NFS then passes this information to either a domain controller or the security authority of the local server, depending on the type of account (domain or local):

- If the Windows user name is a domain account, then the domain controller authenticates the user with Kerberos extensions called Services-For-User (S4U).
- If the Windows user name is a local account, then the local security authority needs the assistance of Server for NFS Authentication. Without Server for NFS Authentication, the local security authority cannot authenticate the user and access will be denied.

| **Note** |
| --- |
| Active Directory Lookup always associates UNIX users to Windows domain user accounts. If Server for NFS is configured to only use Active Directory Lookup, then you will not need to install Server for NFS Authentication on the local computer. If you want to associate UNIX users with Windows user accounts, you must use User Name Mapping. |

| **Important** |
| --- |
| Domain controllers running Windows 2000 and earlier operating systems do not support S4U authentication. UNIX users associated with Windows domain accounts on Windows 2000 and earlier domain controllers will be denied access unless you download and install a prior version of Server for NFS Authentication. For more information, see the Microsoft Web site. |

**To install Server for NFS Authentication**

1. In Control Panel, click **Add or Remove Programs**. Then click **Add/Remove Windows Components**.
2. From the list of components, in the **Windows Components Wizard** dialog box, select **Other Network File and Print Services**, and click **Details**.
3. From the list of subcomponents, select **Microsoft Services for NFS**, and click **Details** again.
4. Select **Server For NFS Authentication**, and click **OK**.
5. Click **OK** again and then **Next**, and after the configuration of the new component has completed, click **Finish**.

# User Name Mapping administration

Applies To: Windows Server 2003 R2

User Name Mapping associates Windows and UNIX user names for Client for NFS and Server for NFS. This allows users to connect to Network File System (NFS) resources without having to log on to UNIX and Windows systems separately.

| **Note** |
| --- |
| Most of the functionality of User Name Mapping has been replaced by Active Directory Lookup. Active Directory Lookup enables Client for NFS and Server for NFS to obtain user identifier (UID) and group identifier (GID) information directly from Active Directory. For information about storing UNIX user data in Active Directory, refer to the accompanying Identity Management for UNIX documentation. For information about enabling Active Directory Lookup, see Specifying how Server for NFS obtains Windows user and group information or Specifying how Client for NFS retrieves UNIX-style identity data. |

This section includes the following topics:


Understanding the User Name Mapping component

Starting and Stopping User Name Mapping

Configuring User Name Mapping

Securing access to the User Name Mapping server

Managing maps

Mapping groups

# Understanding the User Name Mapping component

Applies To: Windows Server 2003 R2

User Name Mapping acts as a single clearinghouse that provides centralized user mapping services for Client for NFS and Server for NFS.

User Name Mapping lets you create maps between Windows and UNIX user and group accounts even though the user and group names in both environments may not be identical. Perhaps most importantly, User Name Mapping lets you maintain a single mapping database for the entire enterprise. This makes it easy to configure account mapping for multiple computers running Microsoft Services for Network File System (NFS).

In addition to one-to-one mapping between Windows and UNIX user and group accounts, User Name Mapping permits one-to-many mapping. This lets you associate multiple Windows accounts with a single UNIX account. This can be useful, for example, when you do not need to maintain separate UNIX accounts for individuals and would rather use a few accounts to provide different classes of access permission.

You can use simple maps, which map Windows and UNIX accounts with identical names. You can also create advanced maps to associate Windows and UNIX accounts with different names, which you can use in conjunction with simple maps. For information about simple and advanced maps, see Understanding user maps.

User Name Mapping can obtain UNIX user, password, and group information from one or more Network Information Service (NIS) servers or from password and group files located on a local hard drive. The password and group files can be copied from a UNIX host or from a NIS server. See Configuring User Name Mapping for more information about specifying the source for UNIX user information.

User Name Mapping periodically refreshes its mapping database from the source databases, ensuring that it is always kept up-to-date as changes occur in the Windows and UNIX name spaces. You can also refresh the database anytime you know the source databases have changed. For information about how to refresh the database automatically and manually, see Configuring the refresh interval for maps and Refreshing maps.

You can back up and restore User Name Mapping data at any time. Because the database is backed up to a file, you can use that file to copy the mapping database to another server. This provides redundancy for the sake of fault tolerance. To learn how to back up and restore the mapping database, see Backing up and restoring maps.

| **Note** |
| --- |
| If you obtain information from multiple NIS domains, it is assumed that each domain has unique users and user identifiers (UIDs). User Name Mapping does not perform any checks. |

# Starting and Stopping User Name Mapping

Applies To: Windows Server 2003 R2


Starting User Name Mapping

Stopping User Name Mapping

# Starting User Name Mapping

Applies To: Windows Server 2003 R2

**To start User Name Mapping from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **User Name Mapping**, and then click **Start Service**.

**To start User Name Mapping from the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] start**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer on which you want to start User Name Mapping. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**mapadmin /?** |

# Stopping User Name Mapping

Applies To: Windows Server 2003 R2

**To stop User Name Mapping from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **User Name Mapping**, and then click **Stop Service**.

**To stop User Name Mapping from the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] stop**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer on which you want to stop User Name Mapping. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br>**mapadmin /?** |

# Configuring User Name Mapping

Applies To: Windows Server 2003 R2


Selecting the UNIX user name source

Configuring the refresh interval for maps

Creating a User Name Mapping pool

Configuring User Name Mapping on a server cluster

# Selecting the UNIX user name source

Applies To: Windows Server 2003 R2

**To specify the UNIX user name source using the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **User Name Mapping**, and then click **Properties**.
4. Do one of the following:
   - Select **Use Network Information Service (NIS)**, or
   - Select **Use Password and Group files**, and in the **Password file path and name** and the **Group file path and name** boxes type or browse to the password and group files.
5. You can use the default refresh interval provided or change it by typing a new day, hour and minute interval.
6. To save the settings, click **Apply**.

| **Note** |
| --- |
| If you use password and group files, the files must be stored on the local computer that is running User Name Mapping. Typically, these files are the /etc/passwd and /etc/group files that are copied from a UNIX server and then periodically updated by the administrator of the computer that is running User Name Mapping.<br><br>If you are using password and group files and need to map users from multiple UNIX servers, you must combine the multiple password files so that each user is assigned a unique user identifier (UID) and combine the multiple group files so that each group is assigned a unique group identifier (GID).<br><br>You should not set up User Name Mapping to use password and group files which are using Microsoft Services for Network File System (NFS) to set up maps that use password and group files that are stored on a remote computer.<br><br>The refresh interval determines how often User Name Mapping refreshes password and group information from the NIS server or the password and group files. |

**To specify the UNIX user name source using the command line**

1. Open the command prompt.
2. To select the UNIX user name source, do one of the following:
   - To use Network Information Service (NIS), at a command prompt, type:

   **mapadmin [** *ComputerName* **] [–u** *User* **] [–p** *Password* **] config –d** *ntDomain* **–y** *nisDomain* **[–n** *nisServer* **]**

   - To use password and group files, at a command prompt, type:

   **mapadmin [** *ComputerName* **] [–u** *User* **] [–p** *Password* **] config –d** *ntDomain* **–f** *Password/GroupDirectory*

3. To select simple or advanced maps, at a command prompt, type:

   **mapadmin [** *ComputerName* **] [–u** *User* **] [–p** *Password* **] config –i {yes|no}**

4. To specify the refresh interval for password and group information, at a command prompt, type:

   **mapadmin [** *ComputerName* **] [–u** *User* **] [–p** *Password* **] config –r** *dddd:hh:mm*

| **Argument** | **Description** |
| --- | --- |

| | |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |
| *ntDomain* | The name of the Windows domain. |
| *Password/GroupDirectory* | The fully qualified path of the directory where the password and group files are located on the computer that is running User Name Mapping. |
| *nisDomain* | The name of the NIS domain. |
| *nisServer* | The name of the NIS server (optional). |
| *dddd:hh:mm* | How often, in days, hours, and minutes, User Name Mapping will refresh password and group information from the NIS server or the password and group files. |

**Note**

If you use password and group files, the files must be stored on the local computer that is running User Name Mapping. Typically, these files are the /etc/passwd and /etc/group files that are copied from a UNIX server and then periodically updated by the administrator of the computer that is running User Name Mapping.

If you are using password and group files and need to map users from multiple UNIX servers, you must combine the multiple password files so that each user is assigned a unique user identifier (UID) and combine the multiple group files so that each group is assigned a unique group identifier (GID).

You cannot use the **mapadmin** command to set up mappings that use password and group files on a remote computer.

The refresh interval determines how often User Name Mapping refreshes password and group information from the NIS server or from the password and group files.

# Configuring the refresh interval for maps

Applies To: Windows Server 2003 R2

**To configure the refresh interval from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **User Name Mapping** and click **Properties**.
4. Click the **Configuration** tab.
5. In **Days, Hours**, and **Minutes**, type the refresh interval. The minimum interval is five minutes.
6. To save the settings, click **Apply**.

**To configure the refresh interval from the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] config –r [[***days***:]***hh***:]***mm*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |
| *days:hh:mm* | The refresh interval for updating from the Windows and Network Information Service (NIS) databases in days, hours, and minutes. The minimum interval is five minutes. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **mapadmin /?** |

# Creating a User Name Mapping pool

Applies To: Windows Server 2003 R2

In larger networks, especially networks that extend across a wide area network (WAN), it is often desirable to run User Name Mapping on multiple computers. In the case of a WAN, running User Name Mapping on computers on the local area networks (LANs) as other computers running Microsoft Services for Network File System (NFS) can help ensure a high level of performance. Even if your network is not geographically dispersed, you might want to run User Name Mapping on more than one computer to improve performance and availability.

When you run User Name Mapping on multiple computers, you can create a pool of servers using Domain Name System (DNS) round robin. Doing this provides a number of benefits:

- You can enhance performance by locating computers running User Name Mapping near other computers running Microsoft Services for NFS.
- You can achieve load balancing by alternating which servers respond to mapping requests.
- High availability can be ensured by providing redundant servers so that if one fails, other servers in the pool can continue to respond to mapping requests.

The first step in creating a User Name Mapping server pool is to install User Name Mapping on two or more servers. These servers must be configured and maintained identically to ensure that they provide consistent responses to mapping requests. The easiest way to do this is to use batch (.cmd) files that repeat the same **mapadmin** command for each server in the pool.

For example, you could install User Name Mapping on three servers named UNM1, UNM2, and UNM3. You would then create a batch file that would apply the same **mapadmin** command and arguments to all three servers:

```
REM POOLADMIN.CMD

REM Used to administer three User Name Mapping servers: UNM1, UNM2, and UNM3

REM Syntax: POOLADMIN mapadmin_argument[...]

mapadmin UNM1 %0 %1 %2 %3 %4 %5 %6 %7 %8 %9 -u UNM1\administrator -p admUNM1pwd

mapadmin UNM2 %0 %1 %2 %3 %4 %5 %6 %7 %8 %9 -u UNM2\administrator -p admUNM2pwd

mapadmin UNM3 %0 %1 %2 %3 %4 %5 %6 %7 %8 %9 -u UNM3\administrator -p admUNM3pwd
```

Use a file like this, or a similar technique, to configure and administer all the User Name Mapping servers in the pool, both before and after the pool has been created.

Once you have set up and configured the servers that will belong to the pool, the next step is to create the pool itself. You do this by configuring a DNS server to map a single name to the Internet Protocol (IP) addresses of all three User Name Mapping servers. For example, you could map the name UNMpool to the IP addresses for UNM1, UNM2, and UNM3. Whenever a client computer asks the DNS server to resolve the name UNMpool, the DNS server will return all three IP addresses, the first of which is the one most clients will use. Because the DNS server rotates (in a round-robin fashion) the order in which it returns the three IP addresses, this automatically provides for load balancing.

The final step in setting up a User Name Mapping pool is to configure computers on the network running Client for NFS or Server for NFS to use the virtual server name of the pool (UNMpool in the example). Because the DNS server rotates the order of the IP addresses it returns when it resolves this name, each of these User Name Mapping clients will receive a different primary IP address than the one given to the previous client, thereby automatically providing load balancing among the three User Name Mapping servers. In addition, these Microsoft Services for NFS components are designed to send mapping requests to each IP address it receives if requests to the previous IP address times out. This provides for automatic failover if one of the servers in the pool should fail.

As noted earlier, it is essential that all servers that belong to a pool are configured and maintained identically. If password and group files are used to provide UNIX account information for mapping users, for example, identical password and group files must be copied to and maintained on all servers in the pool. Similarly, if you use .maphosts files to control access to servers in the pool, you must maintain identical .maphosts files on all servers in the pool.

# Configuring User Name Mapping on a server cluster

Applies To: Windows Server 2003 R2

A server cluster is a group of independent computer systems, known as nodes, working together as a single system to ensure that mission-critical applications and resources remain available to clients. Every node is attached to one or more cluster storage devices. Clustering allows users and administrators to access and manage the nodes as a single system rather than as separate computers.

User Name Mapping is specifically designed to make use of the server cluster and the server cluster's ability to replicate configuration settings among its nodes. This means that when User Name Mapping is installed on the nodes of a cluster, all changes to the configuration of User Name Mapping on one node are automatically replicated to User Name Mapping on all other nodes of the cluster. Consequently, you can manage User Name Mapping on all the nodes in the cluster simply by administering User Name Mapping on any one node. This ensures that user and group mappings and other settings will be identical on all servers in the cluster, thereby allowing you to easily achieve static load balancing by distributing User Name Mapping clients among the nodes of the cluster.

**To install User Name Mapping on a cluster**

1. Create one cluster group for each node with its own Internet Protocol (IP) address and network name resource, and assign the preferred owner for each group so that it is primarily owned by a different node. This allows you to achieve static load balancing.
2. Install User Name Mapping on the nodes of the cluster one at a time. Before installing User Name Mapping on a node, move all groups away from the node.
3. When User Name Mapping has been installed on all nodes, return all groups to their primary owners.
4. Configure User Name Mapping on any node. Settings changed on one node are automatically replicated to all other nodes in the cluster.
5. Configure each client computer to use the virtual server name from one of the cluster groups as its User Name Mapping server. This will ensure that client computers will continue to be served by User Name Mapping in case of node failure. To provide static load balancing, evenly distribute clients among the virtual servers of the cluster.

**Note**

To ensure proper operation of User Name Mapping on a server cluster, when stopping the server cluster, first stop User Name Mapping, then stop the server cluster.

To run User Name Mapping on a server cluster, you must install User Name Mapping on the nodes of an existing cluster. Adding a stand-alone computer that is already running User Name Mapping as a new node in an existing cluster is not supported.

When running User Name Mapping on a server cluster, do not create maps using local Windows user and group accounts. Because the security identifiers (SIDs) of local accounts are recognized only by the server where those accounts were created, other nodes in the cluster will not be able to resolve those accounts when the server fails. Instead, you should create maps using domain accounts only.

If the cluster service is restarted on a node, you must also restart the User Name Mapping service. Until you do, configuration changes will not be replicated to and from User Name Mapping on the node.

User Name Mapping configuration changes are not properly replicated among nodes in a cluster if you run Microsoft Services for Network File System (NFS) or **mapadmin** on a computer that belongs to a domain that is not trusted by the domain of the cluster. To ensure that User Name Mapping configuration changes are properly replicated, always administer User Name Mapping running on a cluster using a computer that belongs to a trusted domain.

You must maintain identical .maphosts files on all servers in the pool. For more information about the .maphosts file, see Securing access to the User Name Mapping server.

If you use password and group files to provide user and group information for User Name Mapping running on a cluster, you must store the password and group files on local disk drives on each node, and the system administrator must update these files as required. In addition, if you restore mapping settings from a backup file, you must also copy the restored password and group files to the appropriate location on all of the nodes in the server cluster and ensure that all nodes in the cluster are configured to use password and group files located in the same local path. For more information about backing up and restoring User Name Mapping data, see Backing up and restoring maps.

If you are running User Name Mapping as well as Server for NFS on a server cluster, configure each node of the cluster to use **localhost** as its User Name Mapping server.

# Securing access to the User Name Mapping server

Applies To: Windows Server 2003 R2

The authentication mechanism used by the Network File System (NFS) protocol is based on trust relationships between host computers. That is, an NFS server accepts the user identifiers (UIDs) and group identifiers (GIDs) from a given client computer because the server is configured to recognize the client computer as trustworthy.

User Name Mapping acts as an intermediary between NFS servers and NFS clients on a network containing UNIX hosts and Windows-based computers. To maintain the implicit trust relationship between NFS client and host computers, it is essential that you control which computers can access User Name Mapping and hence act as a trusted host on the NFS network.

You control which computers can access User Name Mapping by editing a file named .maphosts in the %windir%\msnfs directory of the computer User Name Mapping. By adding entries in this file, you can easily permit or deny access by any other computer to User Name Mapping. If the .maphosts file is present but not empty, then only those computers allowed access by entries in the file can access User Name mapping. If the .maphosts file is present but empty (the default), no computers except the computer running User Name Mapping itself can access User Name Mapping. If the .maphosts file is not present, no computers (including the computer running User Name Mapping) can access User Name Mapping.

The .maphosts file consists of a list of one or more of the elements described in the following table, each on a separate line. When matching a computer making a request against the elements in the list, User Name Mapping searches from the top down until it finds a match. Consequently, order is essential.

| Element | Description |
|---|---|
| *host* | Specifies one or more computers that can access User Name Mapping. The host can be specified as an Internet Protocol (IP) address (IP version 4), or as a host name that resolves to one or more IP addresses. |
| *host –* | Specifies one or more computers that are denied access to User Name Mapping. The host can be specified as an IP address (IP version 4), or as a host name that resolves to one or more IP addresses. Note that there must be at least one blank space between host and the dash (–). |
| **+** | Specifies that all computers can access User Name Mapping unless disallowed by an earlier entry in the list. All entries in the list following this element are ignored. |
| **-** | Specifies that all computers are disallowed access to User Name Mapping unless allowed by an earlier entry in the list. All entries in the list following this element are ignored. |

The # character can be used to add comments to the file; everything on a line following the # character is ignored.

Access by the computer running User Name Mapping (the local host) is always allowed unless specifically denied in the .maphosts file, or the .maphosts file is not present. In case of a multihomed User Name Mapping server, access to the service running on other nodes from any of the nodes is also allowed, unless explicitly denied.

**Examples**

The following example file allows access only by computers named mktg_nfs and sales_nfs:

```
mktg_nfs

sales_nfs

-
```

The following example file allows access by all computers except computers named rogue_host and malicious_host:

```
rogue_host -

malicious_host -

+
```

In the next example, rogue_host is denied access even though it appears in the list without a minus sign (–) because it follows a minus sign on its own line:

```
mktg_nfs
```

sales_nfs

-

rogue_host

# Managing maps

Applies To: Windows Server 2003 R2


Refreshing maps

Backing up and restoring maps

Mapping users

# Refreshing maps

Applies To: Windows Server 2003 R2

**To refresh user data from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **User Name Mapping** and select **Properties**.
4. In **Days**, **Hours**, and **Minutes**, type the refresh interval. The minimum interval is five minutes.
5. To refresh the user data now, click **Synchronize Now**.
6. To save the settings, click **Apply**.

**To refresh user data from the command line**

1. Open the command prompt.
2. At the command prompt, type:

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] config –r [[***days***:]***hh***:]***mm*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |
| *days:hh:mm* | The refresh interval for updating from the Windows and Network Information Service (NIS) databases in days, hours, and minutes. The minimum interval is five minutes. |

| **Note** |
|---|
| There is no command-line method for immediately refreshing maps.<br><br>To view the complete syntax for this command, at a command prompt, type: **mapadmin /?** |

# Backing up and restoring maps

Applies To: Windows Server 2003 R2

**To backup and restore data from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. To back up maps, right-click **User Name Mapping** and select **Backup Maps**.

   Type the fully qualified path and name of the file to which you want to back up the maps, or click **Browse** to find the file. When you have identified or found the file, click **Backup**.
4. To restore maps, right-click **User Name Mapping** and select **Restore Maps**.

   Type the fully qualified path and name of the file from which you want to restore maps, or click **Browse** to find the file. When you have identified or found the file, click **Restore**.

   The following settings and maps are backed up and restored:

   - Configuration type (either Network Information Service [NIS], or password and group files)
   - Domain maps
   - Advanced user and group maps
   - Password and group files, if applicable

   | **Note** |
   | --- |
   | If you back up settings and maps on a remote computer that is configured to use password and group files, the password and group maps are not backed up. |

   If you are using password and group files as the source for UNIX user and group names, and if you are running User Name Mapping in a server cluster or a server pool, you must perform the following tasks after restoring maps:

   - You must copy the restored password and group files to the appropriate directories on all nodes in the server cluster or pool. On a server cluster, the files must be copied to directories with identical path names on each node in the cluster.
   - After copying the files on a server pool, you must configure each node in the pool to use the files in the directories where you copied the files.
   - After copying the files on a server cluster, you must configure User Name Mapping on one node to use the files in the directory where you copied the files, and then click Apply. This will cause the updated settings to be replicated to other nodes in the cluster.

**To backup and restore data from the command line**

1. Open the command prompt.
2. At the command prompt, type one of the following:

   - To create a backup file and perform the backup, type:

     **mapadmin [***ComputerName***] [–u***User***] [–p***Password***] backup***file***

   - To restore, type:

**mapadmin** [*ComputerName*] [**–u***User*] [**–p***Password*] **restore** *file*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |
| *File* | The fully qualified path and name of the file to be used for backup. |

The following settings and maps are backed up and restored:

- Configuration type (either Network Information Service (NIS), or password and group files)
- Domain maps
- Advanced user and group maps
- Password and group files, if applicable

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **mapadmin /?**<br>If you back up settings and maps on a computer configured to use password and group files, the password and group maps are not backed up. |

If you are using password and group files as the source for UNIX user and group names, and if you are running User Name Mapping in a server cluster or a server pool, you must perform the following tasks after restoring maps:

- You must copy the restored password and group files to the appropriate directories on all nodes in the server cluster or pool. On a server cluster, the files must be copied to directories with identical path names on each node in the cluster.
- After copying the files on a server pool, you must configure each node in the pool to use the files in the directories where you copied the files.
- After copying the files on a server cluster, you must configure User Name Mapping on one node to use the files in the directory where you copied the files, and then click **Apply**. This will cause the updated settings to be replicated to other nodes in the cluster.

# Mapping users

Applies To: Windows Server 2003 R2


Understanding user maps

Viewing user maps

Creating a simple user map

Creating an advanced user map

Determining the file GID

Determining the file UID

Removing a user map

# Understanding user maps

Applies To: Windows Server 2003 R2

User Name Mapping acts as a single clearinghouse that provides centralized user mapping services for Client for NFS and Server for NFS.

User Name Mapping enables you create maps between Windows and UNIX user and group accounts even though the user and group names in both environments may not be identical. Perhaps most importantly, User Name Mapping lets you maintain a single mapping database for the entire enterprise. This makes it easy to configure authentication for multiple computers running Microsoft Services for NFS.

In addition to one-to-one mapping between Windows and UNIX user and group accounts, User Name Mapping permits one-to-many mapping. This lets you associate multiple Windows accounts with a single UNIX account. This can be useful, for example, when you do not need to maintain separate UNIX accounts for individuals and would rather use a few accounts to provide different classes of access permission.

You can use simple maps, which map Windows and UNIX accounts with identical names. You can also create advanced maps to associate Windows and UNIX accounts with different names, which you can use in conjunction with simple maps. For information about simple and advanced maps, see Understanding user maps.

User Name Mapping can obtain UNIX user, password, and group information from one or more Network Information Service (NIS) servers or from password and group files located on a local hard drive. The password and group files can be copied from a UNIX host or from a Windows-based computer that is running Server for NIS. See Configuring User Name Mapping for more information about specifying the source for UNIX user information.

User Name Mapping periodically refreshes its mapping database from the source databases, ensuring that it is always kept up-to-date as changes occur in the Windows and UNIX name spaces. You can also refresh the database anytime you know the source databases have changed. For information about how to refresh the database automatically and manually, see Configuring the refresh interval for maps and Refreshing maps.

You can back up and restore User Name Mapping data at any time. Because the database is backed up to a file, you can use that file to copy the mapping database to another server. This provides redundancy for the sake of fault tolerance. To learn how to back up and restore the mapping database, see Backing up and restoring maps.

If you obtain information from multiple NIS domains, it is assumed that each domain have unique users and user identifiers (UIDs). User Name Mapping does not perform any checks.

# Viewing user maps

Applies To: Windows Server 2003 R2

**To view user maps from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Expand **User Name Mapping** and select **User Maps**.

| **Note** |
|---|
| Maps are displayed in the **Mapped users** and **Mapped groups** lists in the order in which the maps are evaluated. Primary maps are always listed and evaluated first, followed by secondary maps. To change the order in which maps are evaluated within the list of secondary maps, right-click **User Name Mapping** and click **Properties**. Then, click the **Simple Maps** tab and use the **Move up** or **Move down** buttons to reposition the map within the list. |

**To view user maps from the command line**

1. Open the command prompt.
2. At the command prompt, type one of the following:

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] list {–all|–advanced|–simple}**

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] list –wu***Windows_User_Name***

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] list –uu***unix_User_Name***

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] list –wg***WindowsGroupName***

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] list –ug***unixGroupName***

| **Argument** | **Description** |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative credentials on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative credentials on the computer, if different from the current user. |
| **–all** | View details of both simple and advanced maps. |
| **–advanced** | View all advanced maps for users and groups. |
| **–simple** | View the simple-mapped domains. |
| *Windows_User_Name* | The Windows user for which you want to see details. |
| *Unix_User_Name* | The UNIX user for which you want to see details. |
| *WindowsGroupName* | The Windows group for which you want to see details. |
| *UnixGroupName* | The UNIX group for which you want to see details. |

| **Note** |
|---|

Maps are displayed in the order in which the maps are evaluated. Primary maps are always listed and evaluated first, followed by secondary maps. You can use the Windows interface to change the order in which secondary maps are evaluated. For information about how to do this, see "Using the Windows interface" in this topic.

To view the complete syntax for this command, at a command prompt, type: **mapadmin /?**

# Creating a simple user map

Applies To: Windows Server 2003 R2

**To create a simple user map from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **User Name Mapping** and select **Properties**.
4. Click the **Simple Maps** tab.
5. Select **Use simple maps**.
6. Click **Add**.
7. If you are creating a simple map using Network Information Service (NIS) data, in the **NIS domain name** box, type the name of the NIS domain. If the master server of the NIS domain is located on a different subnet than the User Name Mapping server, then in the **NIS Server name** box, type the name of the master server.
8. Select the Windows domain for which you are creating a simple map.
9. Click **OK**.
10. To add another simple map between a Windows and an NIS domain, repeat the previous step.
11. To remove a simple map between a Windows and an NIS domain, select the map, and then click **Remove**.
12. To reposition a simple map between a Windows and an NIS domain in the list, select the map, and then click **Move up** or **Move down**.
13. To save the settings, click **Apply**.

| **Note** |
| --- |
| You can create multiple simple maps between Windows and NIS domains. Maps are evaluated from the top of the list of maps to the bottom. |

**To create a simple user map from the command line**

1. Open the command prompt.
2. At the command prompt, do one of the following:
   - To supply required information for authentication using password and group files, type:

     **mapadmin [***ComputerName***] [–u***User***] [–p***Password***] adddomainmap –d***ntDomain***–f***Password***/***GroupDirectory***

   - To supply required information for Network Information Service (NIS) authentication, type:

     **mapadmin [***ComputerName***] [–u***User***] [–p***Password***] adddomainmap –d***ntDomain***–y***nisDomain***

| **Argument** | **Description** |
| --- | --- |
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |

| | |
|---|---|
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |
| *ntDomain* | The name of the Windows domain. |
| *Password/GroupDirectory* | The fully qualified path of the directory on the User Name Mapping computer where the password and group files are located. |
| *nisDomain* | The name of the NIS domain. |

**Note**

Password and group files must be stored on the computer that is running User Name Mapping. Typically, these files are the /etc/passwd and /etc/group files that are copied from a UNIX host and then periodically updated by the administrator of the computer that is running User Name Mapping.

To view the complete syntax for this command, at a command prompt, type:

**mapadmin /?**

# Creating an advanced user map

Applies To: Windows Server 2003 R2

**To create an advanced map from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Expand **User Name Mapping**.
4. Right-click **User Maps** and select **Create Map**.
5. In the **Windows domain name** list, click the domain for which you want to list users, and then click **List Windows Users**.
6. If you are obtaining UNIX user information from a Network Information Service (NIS) domain, in the **NIS Domain name** box, type the name of the NIS domain from which you are getting account information. If the master server of the NIS domain is located on a different subnet than the User Name Mapping server, then in the **NIS Server name box**, type the name of the master server.
7. Click **List UNIX Users**.
8. In the **Windows Account** and **UNIX Account** lists, click the two accounts you want to map to each other, and then click **Add**.
9. To save the settings, click **Apply**.

| **Note** |
| --- |
| As you add maps, new maps are listed below existing maps. This does not reflect the order in which the maps are evaluated, however. To view the actual order, click **Apply**, press **F5** to refresh the User Name Mapping page, and then reopen the list of advanced maps. Primary maps are always listed and evaluated first, followed by secondary maps. To change the order in which maps are evaluated within the list of secondary maps, select a map and then click **Move up** or **Move down** to reposition the map within the list. |

**To create an advanced map from the command line**

1. Open the command prompt.
2. At the command prompt, to supply required information for authentication using password and group files, type:

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] add –wu***WindowsDomain******UserName***–uu***UserName***[–setprimary]**

| Argument | Description |
| --- | --- |
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |
| *WindowsDomain* **** *UserName* | The name of the Windows user. |
| *UserName* | The name of the UNIX user, specified as domain\user_name or user_name@domain, where domain is the name of the Network Information Service (NIS) domain or **PCNFS**, if you are using password and group files. |

| | |
|---|---|
| **–setprimary** | Make this the primary mapping. |

**Note**

The command-line method maps one pair of users or groups at a time. To map more users or groups, repeat the command.

When you use **mapadmin** to create multiple maps for a UNIX user, use the **–setprimary** option to specify the primary map. If necessary, you can use the **mapadmin setprimary** command later to change the primary map.

Maps cannot be evaluated in the order in which they are created. Primary maps are always evaluated before secondary maps.
To view the complete syntax for this command, at a command prompt, type: **mapadmin /?**

# Determining the file GID

Applies To: Windows Server 2003 R2

**To determine the GID of a file**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Double-click **My Network Places**, and then double-click **Entire Network**.
3. Double-click **NFS Network**.
4. Double-click the local area network (LAN) whose servers you want to view.
5. Double-click the server whose shared directories you want to view.
6. Locate and right-click the file or folder and click **Properties**. The file information and group identifier (GID) are displayed.

| **Note** |
| --- |
| When you double-click **NFS Network**, it might take several seconds for the local area network (LAN) to appear, depending on the broadcast network configuration and the number of computers in the broadcast. |

# Determining the file UID

Applies To: Windows Server 2003 R2

**To determine the UID of a file**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Double-click **My Network Places**, and then double-click **Entire Network**.
3. Double-click **NFS Network**.
4. Double-click the local area network (LAN) whose servers you want to view.
5. Double-click the server whose shared directories you want to view.
6. Locate and right-click the file or folder and click **Properties**. The file information and user identifier (UID) are displayed.

| **Note** |
| --- |
| When you double-click **NFS Network**, it might take several seconds for the local area network (LAN) to appear, depending on the broadcast network configuration and the number of computers in the broadcast. |

# Removing a user map

Applies To: Windows Server 2003 R2

**To remove a user map from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **User Name Mapping** and click **Properties**.
4. Click the **Simple Maps** tab.
5. To disable simple maps, clear the **Use simple maps** box.
6. To remove a simple map between a Windows and an NIS domain, select the map, and then click **Remove**.
7. To save the settings, click **Apply**.

**To remove a user map from the command line**

1. Open the command prompt.
2. At the command prompt, do one of the following:
   - To remove a specific map between a Windows domain and an NIS domain, type:

     **mapadmin [***ComputerName***] [–u***User***] [–p***Password***] removedomainmap –d***WindowsDomain***–y***nisDomain*

   - To remove all simple maps, type:

     **mapadmin [***ComputerName***] [–u***User***] [–p***Password***] removedomainmap -all**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |
| *WindowsDomain* | The Windows domain of the map to be deleted. |
| *nisDomain* | The NIS domain of the map to be deleted. |
| **-all** | Deletes all simple maps. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **mapadmin /?** |

# Mapping groups

Applies To: Windows Server 2003 R2


Understanding group maps

Viewing group maps

Creating a simple group map

Creating an advanced group map

Removing a group map

# Understanding group maps

Applies To: Windows Server 2003 R2

User Name Mapping acts as a single clearinghouse that provides centralized user mapping services for Client for NFS and Server for NFS.

User Name Mapping enables you to create maps between Windows and UNIX user and group accounts even though the user and group names in both environments may not be identical. Perhaps most important, User Name Mapping lets you maintain a single mapping database for the entire enterprise. This makes it easy to configure authentication for multiple computers running Microsoft Services for Network File System (NFS).

In addition to one-to-one mapping between Windows and UNIX user and group accounts, User Name Mapping permits one-to-many mapping. This lets you associate multiple Windows accounts with a single UNIX account. This can be useful, for example, when you do not need to maintain separate UNIX accounts for individuals and would rather use a few accounts to provide different classes of access permission.

You can use simple maps, which map Windows and UNIX accounts with identical names. You can also create advanced maps to associate Windows and UNIX accounts with different names, which you can use in conjunction with simple maps. For information about simple and advanced maps, see Understanding user maps.

User Name Mapping can obtain UNIX user, password, and group information from one or more Network Information Service (NIS) servers or from password and group files located on a local hard drive. The password and group files can be copied from a UNIX host or from a Windows-based computer that is running Server for NIS. See Configuring User Name Mapping for more information about specifying the source for UNIX user information.

User Name Mapping periodically refreshes its mapping database from the source databases, ensuring that it is always kept up-to-date as changes occur in the Windows and UNIX name spaces. You can also refresh the database anytime you know the source databases have changed. For information about how to refresh the database automatically and manually, see Configuring the refresh interval for maps and Refreshing maps.

You can back up and restore User Name Mapping data at any time. Because the database is backed up to a file, you can use that file to copy the mapping database to another server. This provides redundancy for the sake of fault tolerance. To learn how to back up and restore the mapping database, see Backing up and restoring maps.

If you obtain information from multiple NIS domains, it is assumed that each domain have unique users and user identifiers (UIDs). User Name Mapping does not perform any checks.

# Viewing group maps

Applies To: Windows Server 2003 R2

**To view group maps from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Expand **Username Mapping Server** and select **Group Maps**.

| **Note** |
| --- |
| Maps are displayed in the **Mapped users** and **Mapped groups** lists in the order in which the maps are evaluated. Primary maps are always listed and evaluated first, followed by secondary maps. To change the order in which maps are evaluated within the list of secondary maps, right-click **Username Mapping Server** and click **Properties**. Then, click the **Simple Maps** tab and use the **Move up or Move down** buttons to reposition the map within the list. |

**To view group maps from the command line**

1. Open the command prompt.
2. At the command prompt, type one of the following:

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] list {–all|–advanced|–simple}**

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] list –wg***WindowsGroupName***

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] list –ug***UnixGroupName***

| **Argument** | **Description** |
| --- | --- |
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative credentials on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative credentials on the computer, if different from the current user. |
| **–all** | View details of both simple and advanced maps. |
| **–advanced** | View all advanced maps for users and groups. |
| **–simple** | View the simple-mapped domains. |
| *WindowsGroupName* | The Windows group for which you want to see details. |
| *UnixGroupName* | The UNIX group for which you want to see details. |

| **Note** |
| --- |
| Maps are displayed in the order in which the maps are evaluated. Primary maps are always listed and evaluated first, followed by secondary maps. You can use the Windows interface to change the order in which secondary maps are evaluated. For information about how to do this, see "Using the Windows interface" in this topic.<br><br>To view the complete syntax for this command, at a command prompt, type: **mapadmin /?** |

# Creating a simple group map

Applies To: Windows Server 2003 R2

**To create a simple group map from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **User Name Mapping** and select **Properties**.
4. Click the **Simple Maps** tab.
5. Select **Use simple maps**.
6. Click **Add**.
7. If you are creating a simple map using Network Information Service (NIS) data, in the **NIS domain name** box, type the name of the NIS domain. If the master server of the NIS domain is located on a different subnet than the User Name Mapping server, then in the **NIS Server name** box, type the name of the master server.
8. Select the Windows domain for which you are creating a simple map.
9. Click **OK**.
10. To add another simple map between a Windows and an NIS domain, repeat the previous step.
11. To remove a simple map between a Windows and an NIS domain, select the map, and then click **Remove**.
12. To reposition a simple map between a Windows and an NIS domain in the list, select the map, and then click **Move up** or **Move down**.
13. To save the settings, click **Apply**.

| **Note** |
| --- |
| You can create multiple simple maps between Windows and NIS domains. Maps are evaluated from the top of the list of maps to the bottom. |

**To create a simple group map from the command line**

1. Open the command prompt.
2. At the command prompt, do one of the following:
   - To supply required information for authentication using password and group files, type:

     **mapadmin [***ComputerName***] [–u***User***] [–p***Password***] adddomainmap –d***ntDomain***–f***Passwd***/***GroupDirectory*

   - To supply required information for Network Information Service (NIS) authentication, type:

     **mapadmin [***ComputerName***] [–u***User***] [–p***Password***] adddomainmap –d***ntDomain***–y***nisDomain*

| Argument | Description |
| --- | --- |
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative credentials on the computer, if different from the current user. |

| | |
|---|---|
| *Password* | The password for the person who has administrative credentials on the computer, if different from the current user. |
| *ntDomain* | The name of the Windows domain. |
| *Password* **/** *GroupDirectory* | The fully qualified path of the directory on the User Name Mapping computer where the password and group files are located. |
| *nisDomain* | The name of the NIS domain. |

**Note**

Password and group files must be stored on the computer that is running User Name Mapping. Typically, these files are the /etc/passwd and /etc/group files that are copied from a UNIX host and then periodically updated by the administrator of the computer that is running User Name Mapping.

To view the complete syntax for this command, at a command prompt, type: **mapadmin /?**

# Creating an advanced group map

Applies To: Windows Server 2003 R2

**To create an advanced group map from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Expand **User Name Mapping**.
4. Right-click **Group Maps** and select **Create Map**.
5. In the **Windows domain** list, select the domain for which you want to list users, and then click **List Windows Groups**.
6. If you are obtaining UNIX user information from a Network Information Service (NIS) domain, in the **NIS Domain name** box, type the name of the NIS domain from which you are getting account information. If the master server of the NIS domain is located on a different subnet than the User Name Mapping server, then in the **NIS Server name box,** type the name of the master server.
7. Click **List UNIX Groups**.
8. In the **Windows Group** and **UNIX Group** lists, click the two accounts you want to map to each other and click **Add**.
9. To save the settings, click **Apply**.

| **Note** |
|---|
| As you add maps, new maps are listed below existing maps. This does not reflect the order in which the maps are evaluated, however. To view the actual order, click **Apply**, press **F5** to refresh the User Name Mapping page, and then reopen the list of advanced maps. Primary maps are always listed and evaluated first, followed by secondary maps. To change the order in which maps are evaluated within the list of secondary maps, select a map and then click **Move up** or **Move down** to reposition the map within the list. |

**To create an advanced group map from the command line**

1. Open the command prompt.
2. At the command prompt, to supply required information for authentication using password and group files, type:

   **mapadmin [** *ComputerName* **] [–u***User***] [–p***Password***] add –wg***WindowsGroup***–ug***UnixGroup***[–setprimary]**

| **Argument** | **Description** |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative credentials on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative credentials on the computer, if different from the current user. |
| *WindowsGroup* | The name of the Windows group. |
| *UnixGroup* | The name of the UNIX group. |
| **– setprimary** | Make this the primary mapping. |

| **Note** |
|---|

The command-line method maps one pair of groups at a time. To map more groups, repeat the command.

When you use **mapadmin** to create multiple maps for a UNIX user, use the **–setprimary** option to specify the primary map. If necessary, you can use the **mapadmin setprimary** command later to change the primary map.

Maps cannot be evaluated in the order in which they are created. Primary maps are always evaluated before secondary maps. To view the complete syntax for this command, at a command prompt, type: **mapadmin /?**

# Removing a group map

Applies To: Windows Server 2003 R2

**To remove a group map from the Windows interface**

1. Open Microsoft Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **User Name Mapping** and click **Properties**.
4. Click the **Simple Maps** tab.
5. To disable simple maps, clear the **Use simple maps** box.
6. To remove a map between a Windows and an NIS domain, select the map, and then click **Remove**.
7. To save the settings, click **Apply**.

**To remove a group map from the command line**

1. Open the command prompt.
2. At the command prompt, type one of the following:
   - To remove a specific map between a Windows domain and an NIS domain, at a command prompt, type:

     **mapadmin [***ComputerName***] [–u***User***] [–p***Password***] removedomainmap–d***WindowsDomain***–y***nisDomain***

   - To remove all simple maps, at a command prompt, type:

     **mapadmin [***ComputerName***] [–u***User***] [–p***Password***] removedomainmap -all**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| *User* | The name of the user who has administrative privileges on the computer, if different from the current user. |
| *Password* | The password for the person who has administrative privileges on the computer, if different from the current user. |
| *WindowsDomain* | The Windows domain of the map to be deleted. |
| *nisDomain* | The NIS domain of the map to be deleted. |
| **-all** | Deletes all simple maps. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **mapadmin /?** |

# Client for NFS administration

Applies To: Windows Server 2003 R2

Client for NFS enables Windows-based computers to access files on UNIX Network File System (NFS) servers. Users can connect to and disconnect from an NFS share as though it were any other shared drive on their network.


Understanding the Client for NFS component

Using Client for NFS

Configuring Client for NFS

Mounting an NFS shared resource to a drive letter

Viewing NFS mount options and properties

Viewing NFS server properties

# Understanding the Client for NFS component

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

Client for NFS gives a Windows-based computer the ability to act as a client to a Network File System (NFS) server. With Client for NFS, you can access files in a mixed environment of computers, operating systems, and networks.

Accessing files on NFS servers requires each Windows user to also have a UNIX-style identity. To enable users to access NFS shared resources using Client for NFS, you can include UNIX-style identity attributes in Active Directory.

Windows Server 2003 R2 domain controllers include the appropriate UNIX-style identity attributes from Internet Engineering Task Force (IETF) Request for Comments (RFC) 2307 (http://go.microsoft.com/fwlink/?LinkId=70027) in the Active Directory schema. They also include Identity Management for UNIX, an optional component, which facilitates managing UNIX-style identities in Active Directory.

Besides retrieving UNIX-style identity data from Active Directory, Client for NFS can also retrieve UNIX-style identity data from a User Name Mapping server. User Name Mapping is an optional component in Windows Server 2003 R2, Server Appliance Kit, and Services for UNIX, and it may already be available on your network. We do not recommend using User Name Mapping for new Client for NFS deployments.

Once Client for NFS is installed and configured, you can access files on the NFS server by mounting those files to your computer using either the UNIX **server:/export** format or the Windows **\\server\share** format. You can also browse shared folders by using the **\\server\share** format in Windows Explorer.

# Using Client for NFS

Updated: July 6, 2012

Applies To: Windows Server 2003 R2


- Starting Client for NFS
- Stopping Client for NFS
- Mounting an NFS shared resource to a drive letter
- Unmounting an NFS shared resource from a drive letter
- Configuring NFS permissions on a file or folder
- Viewing NFS mount options and properties
- Viewing NFS server properties

# Starting Client for NFS

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**To start Client for NFS from the Windows interface**

1. Open Services for Network File System: click **Start**, point to **Programs or All Programs**, point to **Administrative Tools**, and then click **Services for Network File System (NFS)**.
2. If necessary, connect to the computer that you want to manage.
3. Right-click **Client for NFS**, and then click **Start Service**.

**To start Client for NFS from the command line**

1. Open an elevated privilege command prompt.
2. At a command prompt, type:

   **nfsadmin client** [*ComputerName*] **start**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer on which you you want to start Client for NFS. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsadmin client /?** |

# Stopping Client for NFS

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**To stop Client for NFS from the Windows interface**

1. Open Services for Network File System: click **Start**, point to **Programs or All Programs**, point to **Administrative Tools**, and then click **Services for Network File System (NFS)**.
2. If necessary, connect to the computer that you want to manage.
3. Right-click **Client for NFS**, and then click **Stop**.

**To stop Client for NFS from the command line**

1. Open an elevated privilege command prompt.
2. At a command prompt, type:

   **nfsadmin client** [*ComputerName*] **stop**

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer on which you you want to stop Client for NFS. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsadmin client /?** |

# Configuring Client for NFS

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

- Specifying how Client for NFS retrieves UNIX-style identity data
- Selecting the network protocol
- Setting the default mount type
- Specifying default UNIX permissions for new files

# Specifying how Client for NFS retrieves UNIX-style identity data

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

Access to Network File System (NFS) file servers requires UNIX-style user and group identities, which are not the same as Windows user and group identities. To enable users to access NFS shared resources, Client for NFS can retrieve UNIX-style identity data from Active Directory (if the schema includes the appropriate attributes), or from a User Name Mapping server. If Active Directory does not include UNIX-style identity attributes and a User Name Mapping server is not available on your network, then Client for NFS will attempt to access NFS resources anonymously.

NFS servers often limit anonymous access to read-only operations, or simply do not allow anonymous access.

**To specify how Client for NFS retrieves UNIX-style identity data using the Windows interface**

1. Open Services for Network File System: click **Start**, point to **Programs or All Programs**, point to **Administrative Tools**, and then click **Services for Network File System (NFS)**.
2. If necessary, connect to the computer that you want to manage.
3. Right-click **Services for NFS**, and then click **Properties**.
4. Specify how Client for NFS retrieves UNIX-style identity data:
   - To use Active Directory, select the **Active Directory domain name** check box, type the domain name, and then click **Apply**.
   - To use User Name Mapping, select the User Name Mapping check box, type the name of the mapping server, and then click **Apply**.

| **Note** |
| --- |
| Before using User Name Mapping, the computer running Client for NFS must be listed in the .maphosts file on the computer running User Name Mapping. |

**To specify how Client for NFS retrieves UNIX-style identity data using the command line**

1. Open an elevated privilege command prompt.
2. Specify how Client for NFS retrieves UNIX-style identity data:
   - To use Active Directory, at a command prompt, type:

     **nfsadmin mapping** [*ComputerName*] **config adlookup=yes addomain=***DomainName*

   - To use User Name Mapping, at a command prompt, type:

     **nfsadmin mapping** [*ComputerName*] **config maplookup=yes mapsvr=***ServerName*

| Argument | Description |
| --- | --- |
| *DomainName* | The name of the domain controller that you want to use. |
| *ComputerName* | The name of the computer that you want to manage. |
| *ServerName* | The name of the mapping server that you want to use. |

| **Note** |
| --- |
| Before using User Name Mapping, the computer running Client for NFS must be listed in the .maphosts file on the computer running User Name Mapping. |

# Selecting the network protocol

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**To set the network protocol using the Windows interface**

1. Open Services for Network File System: click **Start**, point to **Programs or All Programs**, point to **Administrative Tools**, and then click **Services for Network File System (NFS)**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Client for NFS**, and then click **Properties**.
4. Click the **Client Settings** tab.
5. From the **Transport Protocol** list box, select the type of protocol that Client for NFS uses.
6. To save this setting, click **Apply**.

**To set the network protocol using the command line**

1. Open an elevated privilege command prompt.
2. At a command prompt, type:

   **nfsadmin client** [*ComputerName*] **config protocol**={**TCP** | **UDP** | **TCP+UDP**}

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer that you want to manage. |
| **protocol**={**TCP** \| **UDP** \| **TCP+UDP**} | Set to TCP for Transmission Control Protocol (TCP), UDP for User Datagram Protocol (UDP), or TCP+UDP for both. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsadmin client /?** |

# Setting the default mount type

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**To set the type of mount using the Windows interface**

1. Open Services for Network File System: click **Start**, point to **Programs or All Programs**, point to **Administrative Tools**, and then click **Services for Network File System (NFS)**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Client for NFS**, and then click **Properties**.
4. Click the **Client Settings** tab.
5. To use soft mounts, select **Use soft mounts** and specify the number of times that Client for NFS will attempt to make a connection in **Number of times to retry**. To use hard mounts, select **Use hard mounts**.
6. Specify the amount of time between connection attempts in **Retry interval**.
7. To save these settings, click **Apply**.

**To set the type of mount using the command line**

1. Open an elevated privilege command prompt.
2. At a command prompt, specify the mount type by typing:

   **nfsadmin client** [*ComputerName*] **config mtype=**{**hard** | **soft**}

3. Specify the soft mount retries by typing:

   **nfsadmin client** [*ComputerName*] **config retry=***n*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer that you want to manage. |
| **mtype=**{**hard** \| **soft**} | The type of mount. |
| **retry=** *n* | The number of times Client for NFS will attempt to make a connection for a soft mount. This value must be in the range 1–10. The default is 1. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsadmin client /?** |

# Specifying read and write buffers

Applies To: Windows Server 2003 R2

**To set the buffer sizes using the Windows interface**

1. Open Services for Network File System: click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Services for Network File System**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Client for NFS** and then click **Properties**.
4. Click the **Client Settings** tab.
5. From the **Read buffer size** list box, select the size of the read buffer.
6. From the **Write buffer size** list box, select the size of the write buffer.
7. To save these settings, click **Apply**.

**To set the buffer sizes using the command line**

1. Open the command prompt.
2. At a command prompt, specify the size of the read buffer by typing:

   **nfsadmin client [** *ComputerName* **] config rsize=** *n*
3. Specify the size of the write buffer by typing:

   **nfsadmin client [** *ComputerName* **] config wsize=** *n*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer you want to manage. |
| **config rsize=** *n* | The size of the read buffer, in kilobytes. This value must be 1, 2, 4, 8, 16, or 32; the default is 32 |
| **config wsize=** *n* | The size of the write buffer, in kilobytes. This value must be 1, 2, 4, 8, 16, or 32; the default is 32. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **nfsadmin client /?** |

# Specifying default UNIX permissions for new files

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

When a file is created on a Network File System (NFS) server, a default set of permissions is applied to the file to control access to the file. Client for NFS uses the following default permissions for newly created files on NFS shared resources:

- User:rwx (read, write, and execute)
- Group:r–x (read and execute)
- Others:r–x (read and execute)

**To specify default file permissions using the Windows interface**

1. Open Services for Network File System: click **Start**, point to **Programs or All Programs**, point to **Administrative Tools**, and then click **Services for Network File System (NFS)**.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Client for NFS**, and then click **Properties**.
4. Click the **File Permissions** tab, and select the permissions to use as defaults.
5. To save these settings, click **Apply**.

**To specify default file permissions using the command line**

1. Open an elevated privilege command prompt.
2. At a command prompt, type the following:

   **nfsadmin client** [*ComputerName*] **fileaccess=***ogw*

| Argument | Description |
|---|---|
| *ComputerName* | The name of the computer that you want to manage. |
| *ogw* | o = Code for owner permissions<br>g = Code for group permissions<br>w = Code for permissions for others |

Codes for permissions are shown in the following table.

| Code | Description |
|---|---|
| **0** | None |
| **1** | Execute |
| **2** | Write |
| **3** | Execute and write |
| **4** | Read |
| **5** | Read and execute |

| | |
|---|---|
| **6** | Read and write |
| **7** | Read, write, and execute |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type **nfsadmin client /?** |

# Adding a LAN to the NFS network

Applies To: Windows Server 2003 R2

**To add a LAN to the NFS network**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Double-click **My Network Places**, and then double-click **Entire Network**.
3. Right-click **NFS Network**, and then click **Add/Remove NFS LANs**.
4. Click **Add LAN**.
5. Type the required information, and then click **OK**.

| **Note** |
| --- |
| You must be a member of the Administrators group to add a local area network (LAN) to the Network File System (NFS) network.<br><br>There is no command-line method for this procedure.<br><br>If necessary, contact your UNIX network administrator for the information needed to complete this procedure.<br><br>When you open NFS Network, it might take several seconds for the new LAN to appear, depending on the broadcast network configuration and the number of computers in the broadcast.<br><br>Do not attempt to create a LAN with the name Default LAN or Favorite LAN. |

# Deleting a LAN from the NFS network

Applies To: Windows Server 2003 R2

**To delete a LAN from the NFS network**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Double-click **My Network Places**, and then double-click **Entire Network**.
3. Right-click **NFS Network**, and then click **Add/Remove NFS LANs**.
4. In the list of local area networks (LANs), click the one you want to delete, and then click **Delete**.

| **Note** |
| --- |
| You must be a member of the Administrators group to delete a LAN from the Network File System (NFS) network.<br><br>There is no command-line method for this procedure. |

# Adding a server to the Favorite LAN

Applies To: Windows Server 2003 R2

**To add a server to the favorite LAN**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Double-click **My Network Places**, and then double-click **Entire Network**.
3. Double-click **NFS Network**.
4. Right-click **Favorite LAN**, and then click **Add/Remove Hosts**.
5. Click **Add Host**.
6. Type the name of the host or the Internet Protocol (IP) address of the server to access.

| **Note** |
| --- |
| You must be a member of the Administrators group to delete a local area network (LAN) from the Network File System (NFS) network.<br><br>There is no command-line method for this procedure.<br><br>When you double-click **NFS Network**, it might take several seconds for the LAN to appear, depending on the broadcast network configuration and the number of computers in the broadcast. |

# Deleting a server from the Favorite LAN

Applies To: Windows Server 2003 R2

**To delete a server from the Favorite LAN**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Double-click **My Network Places**, and then double-click **Entire Network**.
3. Double-click **NFS Network**.
4. Right-click **Favorite LAN**, and then click **Add/Remove Hosts**.
5. Click the name of the host you want to delete, and then click **Delete**.

| **Note** |
| --- |
| You must be a member of the Administrators group to delete a local area network (LAN) from the Network File System (NFS) network.<br><br>There is no command-line method for this procedure.<br><br>When you double-click **NFS Network**, it might take several seconds for the LAN to appear, depending on the broadcast network configuration and the number of computers in the broadcast. |

# Opening an NFS file or folder

Applies To: Windows Server 2003 R2

**To browse the NFS network**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Double-click **NFS Network**.
3. Double-click the local area network (LAN) whose servers you want to view.
4. Double-click the server whose shared directories you want to view.
5. Double-click the shared directory containing the directory or file you want to open.
6. Double-click the directory or file you want to open.

| **Note** |
| --- |
| When you double-click **NFS Network**, it might take several seconds for the LAN to appear, depending on the broadcast network configuration and the number of computers in the broadcast. |

# Browsing the NFS network

Applies To: Windows Server 2003 R2

**To browse the NFS network**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Double-click **My Network Places**, and then double-click **Entire Network**.
3. Double-click **NFS Network**.
4. Double-click the local area network (LAN) whose servers you want to view.
5. Double-click the server whose shared directories you want to view.
6. Double-click the shared directory containing the directory or file you want to open.
7. Double-click the directory or file you want to open.

| **Note** |
| --- |
| When you double-click **NFS Network**, it might take several seconds for the LAN to appear, depending on the broadcast network configuration and the number of computers in the broadcast. |

# Mounting an NFS shared resource to a drive letter

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**To mount an NFS shared resource to a drive letter using the command line**

1. Open a command prompt (this does not need to be an elevated command prompt).
2. At a command prompt, type the following:

   **mount** [**–o** *Options*] *ComputerName:/ShareName* {*DeviceName* | *}

   or

   **mount** [**–o** *Options*] *ComputerName**ShareName* {*DeviceName* | *}

| Argument | Description |
|---|---|
| **–o rsize=** *n* | The read buffer size, in kilobytes. Acceptable values are 1, 2, 4, 8, 16, and 32. The default is 32 kilobytes (KB). |
| **–o wsize=** *n* | The write buffer size, in kilobytes. Acceptable values are 1, 2, 4, 8, 16, and 32. The default is 32 KB. |
| **–o timeout=** *n* | The time-out for a remote procedure call (RPC), in seconds. Acceptable values are 0.8, 0.9, and any integer in the range 1-60. The default is 0.8. |
| **–o retry=** *n* | The number of retries for a soft mount. Acceptable values are integers in the range 1-10. The default is 1. |
| **–o mtype=**{**soft** \| **hard**} | Soft or hard mount. Regardless of the mount type, the mount will return if it cannot immediately mount the shared resource. However, once the shared resource has been successfully mounted, if the mount type is a hard mount and the Network File System (NFS) server becomes unavailable, Client for NFS will retry accessing the shared resource until the NFS server is once again available. As a result, Windows applications that are trying to access the shared resource will appear to stop responding. |
| **–o anon** | Mounts as anonymous user. |
| **–o nolock** | Disables locking. This option might improve performance if you only need to read files. |

| | |
|---|---|
| **-o lang=**{**euc-jp** \| **euc-tw** \| **euc-kr** \| **shift-jis** \| **big5** \| **ksc5601** \| **gb2312-80** \| **ansi**} | Specifies the default encoding used for file and directory names and, if used, must be set to one of the following:<br>• ansi<br>• big5 (Chinese)<br>• euc-jp (Japanese)<br>• euc-kr (Korean)<br>• euc-tw (Chinese)<br>• gb2312-80 (Simplified Chinese)<br>• ksc5601 (Korean)<br>• shift-jis (Japanese)<br><br>If this option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:<br>Japanese: SHIFT-JIS<br>Korean: KS_C_5601-1987<br>Simplified Chinese: GB2312-80<br>Traditional Chinese: BIG5 |
| **-o fileaccess=** *mode* | Specifies the default permission mode of new files that are created on the NFS shared resource. Specify the mode as a three-digit number in the form ogw, where o, g, and w are each a digit that represents the access granted to the file owner, group, and the world. The digits must be in the range 0–7 with the following meaning:<br>• 0: No access<br>• 1: x (execute access)<br>• 2: w (write access)<br>• 3: wx (write and execute access)<br>• 4: r (read access)<br>• 5: rx (read and execute access)<br>• 6: rw (read and write access)<br>• 7: rwx (read, write, and execute access) |
| **-o casesensitive** | Forces case-sensitive file lookup (similar to that of UNIX-based NFS clients). |
| *ComputerName* | The name of the server with the shared resource to mount. |

| | |
|---|---|
| *ShareName* | The name of the shared directory to mount. |
| *DeviceName* \| ***** | The drive letter to assign to the mounted shared resource. The asterisk (*) means to use the next available letter. |

**Note**

To view the complete syntax for this command, at a command prompt, type: **mount /?**

When Mounting an NFS shared resource to a drive letter, we recommend the first format shown (*ComputerName*:/*ShareName*).

You can also use the **net use** command to mount an NFS shared resource to a drive letter. For more information about using the **net use** command, see Windows Help or, at a command prompt, type: **net use /?**

# Unmounting an NFS shared resource from a drive letter

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**To unmount an NFS shared resource from a drive letter using the command line**

1. Open a command prompt (this does not need to be an elevated command prompt).
2. At a command prompt, type the following:

   **umount** [**–f**] {**–a** | *Drive*}

| Argument | Description |
|---|---|
| **–f** | Forces the unmount of a Network File System (NFS) network drive. |
| **–a** | Unmounts all NFS network drives. |
| *Drive* | The drive letter that is mapped to an NFS shared resource. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type: **umount /?** |

# Configuring NFS permissions on a file or folder

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**To configure NFS permissions using Windows Explorer**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Right-click the name of the file or directory.
3. Click **Properties**, and then click the **NFS Attributes** tab.
4. In **File attributes (mode)**, select or clear the check boxes corresponding to the mode bits that you want to set or clear, click **Apply**, and then click **OK**. To close the dialog box without making any changes, click **Cancel**.

| **Note** |
| --- |
| Only the owner of a file or a user with appropriate permissions can change the mode of a file.<br><br>There is no command-line method for this procedure. |

# Viewing NFS mount options and properties

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**To view NFS mount options using Windows Explorer**

1. Open Windows Explorer: click **Start**, point to **Programs or AllPrograms**, point to **Accessories**, and then click **Windows Explorer**.
2. Right-click the name of the file or directory, right-click the drive that is assigned to a Network File System (NFS) mount, and then click **Properties**.
3. To view the NFS mount options for the drive, click **NFS Mount Options**.
4. To view the NFS properties for the mounted directory, click **NFS Attributes**.

| **Note** |
| --- |
| If you are logged on with an account that maps to a UNIX account on the NFS server with appropriate administrative credentials, you can change the UNIX file-mode permissions of the mounted directory when viewing the NFS properties. |

**To view NFS mount options using the command line**

1. Open a command prompt (this does not need to be an elevated command prompt).
2. At a command prompt, type the following:

   **mount**

| **Note** |
| --- |
| The mount command without arguments displays all mounted volumes and their options.<br><br>To view the complete syntax for this command, at a command prompt, type: **mount /?** |

# Viewing NFS server properties

Updated: July 6, 2012

Applies To: Windows Server 2003 R2

**To view NFS server properties using the command line**

1. Open a command prompt (this does not need to be an elevated command prompt).
2. At a command prompt, type the following:

   **showmount -e***Server*

| Argument | Description |
|---|---|
| *Server* | The name of the computer for which you want to view properties. |

| Note |
|---|
| To view the complete syntax for this command, at a command prompt, type: **showmount /?** |

# Microsoft Services for NFS troubleshooting

Applies To: Windows Server 2003 R2

**General Issues**
**I need to modify ports used by Microsoft Services for NFS.**
**Cause**
Port conflicts or security requirements necessitate port changes.

**Solution**
Log on to servers running Microsoft Services for Network File System (NFS) and change the following ports from the Registry Editor:

**NLM and NSM ports**
To modify the NLM and NSM ports, change the following registry settings:

HKLM\SYSTEM\CurrentControlSet\Services\NfsSvr\NlmNsm

DWORD "NlmPort"

DWORD "NsmPort"

**Mount and NFS ports**
To modify the mount protocol port, change the following registry setting:

HKLM\SYSTEM\CurrentControlSet\Services\NfsSvr\Parameters

DWORD "MountPort"

DWORD "NfsPort

**Base port**
The base port specifies the first of four sequential ports used by the mapping server. To modify the base port and the subsequent mapping server ports, change the following registry setting:

HKLM\SYSTEM\CurrentControlSet\Services\Mapsvc\CurrentVersion

DWORD "BasePort"

**Microsoft Services for NFS component troubleshooting**
Troubleshooting Server for NFS

Troubleshooting Client for NFS

Troubleshooting User Name Mapping

# Troubleshooting Server for NFS

Applies To: Windows Server 2003 R2

What problem are you having?

**Issues**
**Despite appearing to have appropriate permissions, the user cannot access folder or file.**
**Cause**
The user belongs to one or more groups that are not mapped consistently, Active Directory is not accessible, or User Name Mapping is not running.

**Solution**
To ensure proper file access, ensure that Windows and UNIX groups that are mapped to each other, in either Active Directory or User Name Mapping, contain the same users, and that the members of the Windows and UNIX groups are properly mapped to each other. Also, ensure that Active Directory is accessible or the User Name Mapping service is running on the designated server.

**Authenticated users cannot access Network File System (NFS) resources.**
**Cause**
Active Directory Lookup or User Name Mapping is not properly configured to work with this computer.

**Solution**
If you are using Active Directory Lookup, ensure that Microsoft Services for Network File System (NFS) is pointing to the proper Active Directory domain.

If you are using User Name Mapping, ensure that the .maphosts file on the computer running User Name Mapping specifies the names or Internet Protocol (IP) addresses of computers that can map user accounts by using User Name Mapping. For more information, see Securing access to the User Name Mapping server. If users cannot access NFS resources intermittently and configuring the .maphosts file does not solve the problem, it might be that too many client computers are trying to access User Name Mapping simultaneously. See "Error message: Unable to perform the requested operation as the mapping service cannot be contacted" in Troubleshooting User Name Mapping.

**Users, including properly mapped users, cannot change the current directory to a shared directory or create files in the directory, even though anonymous access to the directory is allowed.**
**Cause**
Either the NFS client does not support NFS version 3, or Server for NFS is not configured to support NFS 3. Also, the discretionary access control list (DACL) protecting the shared directory does not have an entry for Everyone, and so the access mode bit for Other is reported as 0. Because NFS 2 clients rely on directory mode settings rather than performing a separate access check for the shared directory, the client incorrectly fails the attempted access.

**Solution**
Do one of the following:

- To the DACL protecting the directory being shared, add an entry that grants Everyone read or read/write access, as appropriate.
- Ensure that the NFS client supports NFS 3, and enable NFS 3 support by Server for NFS.
- The shared directory is not accessible even though it is listed as available.

**Cause**
The directory was moved after it was shared.

**Solution**
Return the directory to its original location, or stop sharing the directory and then reshare it.

**All files created on Server for NFS are owned by Anonymous.**
**Cause**
Authentication is not configured properly.

**Solution**
Ensure that mappings are set up correctly in Active Directory or in User Name Mapping and that Server for NFS is correctly configured to use either Active Directory Lookup or User Name Mapping. Also, be sure all Windows Server 2000 domain

controllers are properly configured. For more information, see Server for NFS Authentication

**Files created by a new user are owned by Anonymous.**
**Cause**
If you are using User Name Mapping, then User Name Mapping and Server for NFS have not yet refreshed data from the Network Information Service (NIS) server. Typically, User Name Mapping refreshes data from NIS once an hour, and Server for NFS refreshes data from User Name Mapping once an hour.

**Solution**
The new user should wait at least two hours before attempting to access or create files on Server for NFS, or the administrator of the computer running User Name Mapping can refresh the mapping database.

**A user cannot write to a file.**
**Cause**
File permissions or attributes do not allow write access to the file or its directory.

**Solution**
If the directory is owned by the Administrators group, make an individual user account the owner of the directory. Ensure that the user's UNIX account is mapped to a valid Windows account and that the NTFS file system permissions of the directory and file allow write access to the Windows user account. Ensure that the read-only attribute is not set on the file or the directory.

**Users of Japanese UNIX systems cannot view file names in Japanese.**
**Cause**
Extended UNIX character (EUC) set is not enabled.

**Solution**
Configure the share to use the appropriate character encoding.

**Server for NFS configuration settings are not replicated across nodes in a server cluster.**
**Cause**
The Cluster service is not running, was not running when Server for NFS started, or failed after Server for NFS started.

**Solution**
Take all NFS shares owned by the node offline, or move the cluster groups containing NFS shares to another node. Stop Server for NFS, start the cluster service if needed, and then restart Server for NFS. Return the NFS shares online or move the cluster groups back to the node.

**Server for NFS fails to stop on a server cluster.**
**Cause**
This is by design. When an NFS share is online on a cluster node, the cluster service automatically restarts Server for NFS to keep the shares online.

**Solution**
Before stopping Server for NFS on a server cluster node, take all NFS shares owned by the node offline, or move the cluster groups containing NFS shares to another node.

**An NFS network share on a server cluster node fails to come online.**
**Cause**
An NFS share on that node with the same alias or path already exists on the node.

**Solution**
Make sure that the shared path and alias are unique across the cluster. Also, avoid having nonclustered NFS shares on a server cluster node.

**Cause**
The user who installed the cluster service does not have read permission on the directory that is shared, so the path cannot be validated.

**Solution**
Grant read access to the directory for the user who installed the cluster service.

**Cause**
The disk resource containing the directory being shared is offline, so the cluster service cannot verify the path of the share.

**Solution**
Bring the disk resource online, and then bring the NFS share online. We recommend that the NFS share resource be made dependent on the disk resource that contains the shared folder.

**Cause**

The disk is inaccessible due to hardware error.

**Solution**
Take the NFS shares on the disk offline. Ensure that the disk is accessible from all cluster nodes, and then bring the NFS shares online.

**Cause**
There are a large number of subdirectories under a subdirectories-only share, and the resource times out before all the shares are created when the resource is coming online.

**Solution**
Increase the time-out interval for the resource.

**Creating or modifying an NFS share fails with the error: The share path specified does not exist, or you are trying to modify the properties of an online share.**
**Cause**
The specified directory does not exist.

**Solution**
Make sure that the directory exists and that the path is correct.

**Cause**
The share is online.

**Solution**
Take the share offline, make the required modifications, and then bring the share online.

**Cause**
The disk resource containing the shared directory is offline, so the cluster service cannot verify the path of the share.

**Solution**
Bring the disk online, make the necessary modifications to the NFS share resource, and then bring the NFS share online.

**User Name Mapping is properly configured, but users are not being mapped correctly.**
**Cause**
Server for NFS is not set to use the correct User Name Mapping server.

**Solution**
Ensure that the specified User Name Mapping server is valid. If the User Name Mapping server is on a server cluster, ensure that the following conditions exist:

- User Name Mapping is installed on all cluster nodes
- User Name Mapping data is being replicated to all nodes in the cluster
- Server for NFS is using the name of a Network Name cluster resource as the User Name Mapping server, not localhost or the name of any of the cluster nodes

**Cause**
Server for NFS has not received updated maps from the mapping server. If Server for NFS and User Name Mapping are on different computers, this occurs once every 30 minutes.

**Solution**
Force Server for NFS to refresh the maps by doing one of the following:

- In Microsoft Services for Network File System, click **Server for NFS**, and then click **Apply**.
- Use the **nfsadmin server** command to execute an operation, such as by setting a value to its current value.
- Restart Server for NFS.

**Cause**
Account changes on the Windows domain controller or the Network Information Service (NIS) server have not been received by User Name Mapping.

**Solution**
Force Server for NFS to refresh the maps by doing one of the following:

- In Microsoft Services for Network File System, click **Server for NFS**, and then click **Apply**.
- Use the **nfsadmin server** command to execute an operation, such as by setting a value to its current value.
- Restart Server for NFS.

**Cause**
Local accounts on a cluster node were mapped to UNIX user accounts. Local accounts are not valid on all nodes in a cluster.

**Solution**
Ensure that all Windows accounts mapped to UNIX accounts on User Name Mapping running on a cluster are Windows domain accounts.

**Cause**
The passwd and group files are not at the same location on all nodes of the cluster, or on a network drive.

**Solution**
Ensure that the passwd and group files are identical and at identical locations on local disks of all nodes.

**Cause**
The Server for NFS server is not on the list of permitted User Name Mapping server clients.

**Solution**
Ensure that the .maphosts files on all User Name Mapping server cluster nodes are identical to permit the node running Server for NFS to obtain maps from the User Name Mapping server.

**Cause**
The server running User Name Mapping has failed.

**Solution**
Correct the cause of failure and then restart User Name Mapping on the server.

**Cause**
A Windows account mapped to a UNIX account is disabled or no longer exists.

**Solution**
If the Windows account exists, but is disabled, enable it. If the account does not exist, create a new account and, if necessary, recreate the corresponding advanced map.

**Cause**
A Windows user account has not been granted the credentials to log on to the network.

**Solution**
Grant the required credentials to the Windows user account, and then force Server for NFS to refresh the maps by doing one of the following:

- In Microsoft Services for Network File System, click **Server for NFS**, and then click **Apply**.
- Use the **nfsadmin server** command to execute an operation, such as by setting a value to its current value.
- Restart Server for NFS.

**Cause**
Windows and UNIX groups mapped to each other do not contain the same members.

**Solution**
Ensure that all Windows users in a group are mapped to UNIX users in the corresponding UNIX group, and that all UNIX users in a group are mapped to users in the corresponding Windows group.

**Cause**
User Name Mapping settings are not properly replicated on all nodes in a server cluster.

**Solution**
Ensure that User Name Mapping on a server cluster is configured properly to allow replication on all nodes.

**A group containing an NFS share fails to come online on a particular cluster server node.**
**Cause**
Server for NFS is not installed on the node.

**Solution**
Install Server for NFS on the node.

**Cause**
The node is not configured as the preferred owner of the group.

**Solution**
Configure the group's properties to make the node the preferred owner of the group.

**Cause**
One of the resources in the group does not list the node as a possible owner even though the group specifies the node as a preferred owner.

**Solution**
Configure the resource's properties to specify the node as a possible owner.

**The showmount –e command for a virtual server lists all shares on the node rather than those in the same group as the virtual server.**
**Cause**
This is by design. There is only one instance of Server for NFS running on a node that will enumerate all shares on that node. It does not distinguish between shares in different cluster groups.

**Solution**
Maintain different groups on different nodes.

**The root user is not granted proper permissions.**
**Cause**
The share does not have root access enabled.

**Solution**
Right-click the shared directory, click **Properties**, click **NFS Sharing**, click **Permissions**, and then click **Allow root access**.

**Cause**
The computer from which the root user is accessing the share is not permitted root access.

**Solution**
Right-click the shared directory, click **Properties**, click **Permissions**, and then do one of the following:

- Grant root access to ALL MACHINES.
- Grant root access to a client group containing the computer.
- Grant root access to the computer itself.

**Cause**
The root user does not have read/write permission.

**Solution**
Grant the appropriate permissions to the Windows user mapped to the root user. Right-click the shared directory, click **Properties**, click **Permissions**, and then click **Root access allowed**.

**Cause**
The root user account is not properly mapped to a Windows user account.

**Solution**
Map the root user to a Windows account in the Administrators or Domain Admins group, and map the root user's group to the same Windows group.

**Cause**
Server for NFS has not received updated maps from User Name Mapping.

**Solution**
Force Server for NFS to refresh the maps by doing one of the following:

- In Microsoft Services for Network File System, click **Server for NFS**, and then click **Apply**.
- Use the **nfsadmin server** command to execute an operation, such as by setting a value to its current value.
- Restart Server for NFS.

**Cause**
The root user's user identifier (UID) is not 0. Server for NFS grants root access only to a UNIX user with a UID of 0.

**Solution**
Change the root user's UID to 0.

**Anonymous access fails on Windows Server 2003 operating systems.**
**Cause**
Local security policy is not set to enable Everyone permissions to apply to anonymous users (the default).

**Solution**
Use the Local Security Policy manager to enable **Network Access: Let Everyone permissions to apply to anonymous users** in Security Options in Local Policies.

# Troubleshooting Client for NFS

Applies To: Windows Server 2003 R2

What problem are you having?

**Issues**
**I need Client for NFS to only use "privileged" ports.**
**Cause**
Some Network File System (NFS) servers have a concept of "privileged" ports (meaning ports numbered lower than 1024), and will not accept client connections from ports above that range.

**Solution**
Log onto computers running Client for NFS and set the following registry key to 1 and restart Client for NFS:

HKLM\Software\Microsoft\Client for NFS\CurrentVersion\default

DWORD "UseReservedPorts"

**An error occurred while trying to communicate with the Client for NFS service.**
**Cause**
The Client for NFS service is not started, cannot contact Active Directory or the User Name Mapping server, or is running with credentials that are mapped incorrectly.

**Solution**
In Microsoft Services for Network File System, you can ensure that the service is started by right-clicking **Client for NFS**, and then clicking **Start**. In Microsoft Services for Network File System, ensure that the name of the correct Active Directory domain or User Name Mapping server is specified. Check with the administrator of the Active Directory or the User Name Mapping server to ensure that your Windows user account is properly mapped to your UNIX user account.

**System error 1722 has occurred. The remote procedure call (RPC) server is unavailable.**
**Cause**
The Client for NFS service is not started.

**Solution**
In Microsoft Services for Network File System, right-click **Client for NFS**, and then click **Start**.

**An attempt to mount a shared directory produces a network error 67, the network name was not found.**
**Cause**
The directory is not shared, or the server is unable to resolve the client name.

**Solution**
Use **showmount –e***server* to display the names of shared directories on the server. If this fails with a remote procedure call (RPC) time-out, contact the administrator of the Network File System (NFS) server to ensure that the NFS and portmap services are started on the server. Ensure that the Internet Protocol (IP) address of the client is specified in either the server /etc/hosts file, or in the DNS database, as appropriate.

**I received the message RPC: Timed out or RPC: Port mapper failure – RPC: Timed out.**
**Cause**
Either the port mapper or Network File System (NFS) server daemons are not running on the server, or Client for NFS performance settings need to be adjusted.

**Solution**
Ensure that the services are started. If the port mapper and NFS server daemons are running on the server, in Microsoft Services for Network File System, click **Client for NFS**, and then increase the maximum number of retries and the interval between retries.

**A file I created on a mounted Network File System (NFS) drive is shown as having been created by nobody:nogroup or by –2:–1, or the user identifier (UID) is a very large number (2**32-2).**
**Cause**
Your Windows user name is not properly mapped in Active Directory or User Name Mapping.

**Solution**
Use the server running User Name Mapping for your Windows user account that is not properly mapped. In Microsoft Services for NFS, check to ensure that the name of the correct Active Directory domain or User Name Mapping server is specified. Check with the administrator of the Active Directory or User Name Mapping server to ensure that your Windows user account

is properly mapped to your UNIX user account.

**When I attempt to map a network drive, I am asked for a user name and password, and then the attempt to map the drive fails even though I supply valid credentials.**
**Cause**
Either an Active Directory domain or a User Name Mapping server has not been designated for use by Microsoft Services for NFS.

**Solution**
Specify the correct Active Directory domain or User Name Mapping server. For more information, see Specifying how Client for NFS retrieves UNIX-style identity data.

**I cannot access files on a mapped network drive.**
**Cause**
Neither an Active Directory domain nor a User Name Mapping server has been designated for use by Microsoft Services for NFS.

**Solution**
Specify the correct Active Directory domain or User Name Mapping server. For more information, see Specifying how Client for NFS retrieves UNIX-style identity data.

**I cannot mount a shared root directory (/) on a UNIX computer.**
**Cause**
When a root directory (/) on a UNIX computer is shared, Client for NFS cannot mount it by using the standard path syntax (net use * \\server\ or net use server:/).

**Solution**
To mount the root directory, enter this command:

**net use * ** *server* **\!**

The share name must be an exclamation mark (!).

**MS-DOS-based or earlier Windows-based applications cannot access files on NFS, or a message appears indicating that a 16-bit application cannot be run.**
**Cause**
By default, Client for NFS does not convert UNIX long file names to the MS-DOS–style 8.3 format.

**Solution**
Use Registry Editor (Regedit.exe) to ensure that the following entry is set to a value of 1 (the default):

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Client for NFS\CurrentVersion\Users\Default\Defaults\8.3Format
```

This causes Client for NFS to convert all UNIX file names to the 8.3 format required by MS-DOS-based applications and by applications designed to run with Windows 3.x and earlier. Setting this entry to a value of 0 improves Client for NFS performance somewhat because Client for NFS does not convert file names.

| ⚠Caution |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

**I received a delayed write error message.**
**Cause**
Either the NFS server is not responding or the Client for NFS performance options are not properly adjusted.

**Solution**
Use the **showmount –e***server* command to determine whether the NFS server is running properly.

**Attempting to mount with the net use command produces the error message: 1788: The trust relationship between the primary domain and the trusted domain failed.**
**Cause**
You are attempting to access a computer in an untrusted domain, and either Client for NFS has not started on your computer or the share is not available from the NFS server.

**Solution**
In Microsoft Services for Network File System, right-click **Client for NFS**, and then click **Start**. Use **showmount –e***server* to display the names of shared directories on the server. If this fails with an RPC time-out, contact the administrator of the NFS

server to ensure that the NFS and portmap services are started on the server and that the shared directory is available.

**Even after changing the default permissions for new files, new files on a mount are still created with the previous permissions.**
**Cause**
This is by design. New default permissions are only applied to new mounts.

**Solution**
Unmount and then remount the shared directory.

**After restarting the computer, the mount parameters of a persistent mount are changed.**
**Cause**
You used the **net use** command to mount the directory.

**Solution**
Use the **mount** utility to mount the directory.

**Authenticated users cannot access NFS resources, or all NFS shares are accessed as Anonymous.**
**Cause**
User Name Mapping is not properly configured to work with this computer.

**Solution**
Ensure that the .maphosts file on the computer running User Name Mapping specifies the names or Internet Protocol (IP) addresses of computers that can map user accounts by using User Name Mapping. For more information, see Securing access to the User Name Mapping server. If users cannot access NFS resources intermittently and configuring the .maphosts file does not solve the problem, it could be that too many client computers are trying to access User Name Mapping simultaneously. See "Error message: Unable to perform the requested operation as the mapping service cannot be contacted" in Troubleshooting User Name Mapping.

# Troubleshooting User Name Mapping

Applies To: Windows Server 2003 R2

What problem are you having?

**Issues**
**I cannot download Windows users from a Windows domain.**
**Cause**
You do not have adequate credentials, or the computer running User Name Mapping does not belong to a domain trusted by the Windows user domain.

**Solution**
Log on as a user in a domain trusted by the domain to which the computer running User Name Mapping belongs and that belongs to the Administrators group on the computer running User Name Mapping.

**I cannot download Windows users from an Active Directory domain.**
**Cause**
When installing the Active Directory domain controller, you chose the **Permissions compatible with Windows 2000 Server** option.

**Solution**
Do one of the following:

- If the computer running User Name Mapping is a member of the domain, add the computer to the security group called Pre-Windows 2000 Compatible Access.

- If the computer running User Name Mapping does not belong to the domain (that is, it belongs to a trusted domain), add the special group Everyone to the Pre-Windows  000 Compatible Access security group of the Active Directory domain.

You must restart the domain controller after adding the computer or group to the security group.

If, for some reason, you cannot add Everyone or the computer that is running User Name Mapping to the Pre-Windows 2000 Compatible Access security group, you can enable User Name Mapping to enumerate users in the domain by running the User Name Mapping service with an account that belongs to the Domain Admins group. This method is less secure than the other two solutions, however, and should be avoided if possible.

**I cannot download UNIX users from a Network Information Service (NIS) server.**
**Cause**
You have specified only the domain name, and either the NIS server is on a different subnet, or the NIS server is down.

**Solution**
Ensure that the NIS server is running. Ensure that the computer running User Name Mapping and the NIS server are on the same subnet, or specify the Internet Protocol (IP) address of the NIS server.

**Users of a particular Windows-based computer cannot access Network File System (NFS) shares.**
**Cause**
The Windows-based computer is not granted access in the .maphosts file to User Name Mapping.

**Solution**
Add the Windows-based computer to the .maphosts file. Make sure it does not appear after any – entry.

**Mapped users cannot access Network File System (NFS) resources.**
**Cause**
The .maphosts file has not been modified to specify trusted computers.

**Solution**
Ensure that the .maphosts file specifies the names or Internet Protocol (IP) addresses of computers that can map user accounts by using User Name Mapping. For more information, see Securing access to the User Name Mapping server.

If users still cannot access NFS resources, see "Error message: Unable to perform the requested operation as the mapping service cannot be contacted" later in this topic.

**Error message: Unable to perform the requested operation as the mapping service cannot be contacted.**
**Cause**

The maximum number of concurrent client connections has been exceeded.

**Solution**
If this message appears frequently when you are attempting to administer User Name Mapping, or if properly mapped users can only access NFS resources intermittently and other troubleshooting steps have failed, you can increase the maximum number of concurrent client connections that are allowed by creating a DWORD registry value named RpcMaxConcurrentConnectionsPerIp in the following location:

**HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft \MSNFS**

The value must be between 16 (the default) and 64.

| **Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

If the problem persists, or if NFS access performance noticeably degrades, consider establishing a User Name Mapping server pool. For more information, see Creating a User Name Mapping pool.

**If you click Reload immediately after creating and saving a map, the new map is not shown.**
**Cause**
Maps are saved to the registry, which notifies User Name Mapping of the new map. Because this process can take a minute or more, if you click **Reload** during that time, you will not see current data.

**Solution**
Wait at least a minute, and then click **Reload** again.

**A user is unable to access NFS resources after being added or changed.**
**Cause**
Maps have not been refreshed with new or changed data in the Windows or UNIX user databases.

**Solution**
After adding or changing a user in the Windows domain, the Network Information Service (NIS) domain, or the password and group files, update User Name Mapping. On the **Configuration** tab, click **Synchronize Now**, and then click **Apply**.

**User Name Mapping configuration settings are not replicated across nodes in a server cluster.**
**Cause**
The Cluster service is not running, was not running when User Name Mapping started, or failed after User Name Mapping started.

**Solution**
Stop User Name Mapping, start the Cluster service if needed, and then restart User Name Mapping.

If replication consistently fails after the computer restarts, it might be because User Name Mapping is starting before the Cluster service completes its startup procedures. In this case, use the Services snap-in to make User Name Mapping dependent on the Cluster service.

**Restoring password and group files on a server cluster node does not recreate the files on other nodes.**
**Cause**
User Name Mapping restores password and group files on one node only.

**Solution**
Copy the password and group files to all other cluster server nodes.

**Maps are not updated to reflect changes in Windows user and group accounts or in NIS or password and group files.**
**Cause**
The User Name Mapping server has not received the changes. By default, User Name Mapping refreshes its maps once every 24 hours.

**Solution**
In Microsoft Services for Network File System, click **User Name Mapping**, click **Map Maintenance**, and then click **Synchronize Now**.

# Microsoft Services for NFS command-line tools

Applies To: Windows Server 2003 R2

The following command-line tools are available for Microsoft Services for Network File System (NFS):

- mapadmin
- mount
- nfsadmin
- nfsshare
- nfsstat
- rpcinfo
- showmount
- umount

For more information about a specific command, type the command name followed by **/?** at the command prompt.

# mapadmin

Applies To: Windows Server 2003 R2

You can use **Mapadmin** to manage User Name Mapping for Microsoft Services for Network File System.

**Syntax**

**mapadmin** [*computer*] [-u *user* [-p *password*]]

**mapadmin** [*computer*] [-u *user* [-p *password*]] {start | stop}

**mapadmin** [*computer*] [-u *user* [-p *password*]] **config** *option*[...]

**mapadmin** [*computer*] [-u *user* [-p *password*]] **add -wu** *WindowsUser* **-uu** *UNIXUser* [**-setprimary**]

**mapadmin** [*computer*] [-u *user* [-p *password*]] **add -wg** *WindowsGroup* **-ug** *UNIXGroup* [**-setprimary**]

**mapadmin** [*computer*] [-u *user* [-p *password*]] **setprimary -wu** *WindowsUser* [-uu *UNIXUser*]

**mapadmin** [*computer*] [-u *user* [-p *password*]] **setprimary -wg** *WindowsGroup* [-ug *UNIXGroup*]

**mapadmin** [*computer*] [-u *user* [-p *password*]] **delete** *option*[...]

**mapadmin** [*computer*] [-u *user* [-p *password*]] **list** *option*[...]

**mapadmin** [*computer*] [-u *user* [-p *password*]] **backup** *filename*

**mapadmin** [*computer*] [-u *user* [-p *password*]] **restore** *filename*

**mapadmin** [*computer*] [-u *user* [-p *password*]] **adddomainmap -d** *WindowsDomain* {-y *NISDomain* | -f *path*}

**mapadmin** [*computer*] [-u *user* [-p *password*]] **removedomainmap -d** *WindowsDomain* **-y** *NISDomain*

**mapadmin** [*computer*] [-u *user* [-p *password*]] **removedomainmap -all**

**mapadmin** [*computer*] [-u *user* [-p *password*]] **listdomainmaps**

**Description**

The **mapadmin** command-line utility administers User Name Mapping on the local or remote computer running Microsoft Services for Network File System. If you are logged on with an account that does not have administrative credentials, you can specify a user name and password of an account that does.

In addition to specific command arguments, **mapadmin** accepts the following arguments and options:

***computer***
Specifies the remote computer running the User Name Mapping service that you want to administer. You can specify the computer using a Windows Internet Name Service (WINS) name or a Domain Name System (DNS) name, or by Internet Protocol (IP) address.

**-u *user***
Specifies the user name of the user whose credentials are to be used. It might be necessary to add the domain name to the user name in the form *domain******username*.

**-p *password***
Specifies the password of the user. If you specify the **-u** option but omit the **-p** option, you are prompted for the user's password.

The specific action that **mapadmin** performs depends on the command argument you specify:

**start**
Starts the User Name Mapping service.

**stop**
Stops the User Name Mapping service.

**config**
Specifies general settings for User Name Mapping. The following options are available with this command argument:

**-r** *dddd:hh:mm*

Specifies the refresh interval for updating from the Windows and NIS databases in days, hours, and minutes. The minimum interval is 5 minutes.

**-i** *yes* | *no*

Turns simple mapping on (**yes**) or off (**no**). By default, simple mapping is on.

**add**

Creates a new mapping for a user or group. The following options are available with this command argument:

| Term | Definition |
|---|---|
| **-wu** *name* | Specifies the name of the Windows user for which a new mapping is being created. |
| **-uu** *name* | Specifies the name of the UNIX user for which a new mapping is being created. |
| **-wg** *group* | Specifies the name of the Windows group for which a new mapping is being created. |
| **-ug** *group* | Specifies the name of the UNIX group for which a new mapping is being created. |
| **-setprimary** | Specifies that the new mapping is the primary mapping. |

**setprimary**

Specifies which mapping is the primary mapping for a UNIX user or group with multiple mappings. The following options are available with this command argument:

| Term | Definition |
|---|---|
| **-wu** *name* | Specifies the Windows user of the primary mapping. If more than one mapping for the user exists, use the **-uu** option to specify the primary mapping. |
| **-uu** *name* | Specifies the UNIX user of the primary mapping. |
| **-wg** *group* | Specifies the Windows group of the primary mapping. If more than one mapping for the group exists, use the **-ug** option to specify the primary mapping. |
| **-ug** *group* | Specifies the UNIX group of the primary mapping. |

**delete**

Removes the mapping for a user or group. The following options are available for this command argument:

| Term | Definition |
|---|---|
| **-wu** *user* | The Windows user for which the mapping will be deleted, specified as *WindowsDomain\UserName*. You must specify either the **-wu** or the **-uu** option, or both. If you specify both options, the particular mapping identified by the two options will be deleted. If you specify only the **-wu** option, all mappings for the specified user will be deleted. |
| **-wg** *group* | The Windows group for which the mapping will be deleted, specified as *WindowsDomain\GroupName*. You must specify either the **-wg** or the **-ug** option, or both. If you specify both options, the particular mapping identified by the two options will be deleted. If you specify only the **-wg** option, all mappings for the specified group will be deleted. |
| **-uu** *user* | The UNIX user for whom the mapping will be deleted, specified as *UserName*. You must specify either the **-wu** or the **-uu** option, or both. If you specify both options, the particular mapping identified by the two options will be deleted. If you specify only the **-uu** option, all mappings for the specified user will be deleted. |

| | |
|---|---|
| **-ug** *group* | The UNIX group for which the mapping will be deleted, specified as *GroupName*. You must specify either the **-wg** or the **-ug** option, or both. If you specify both options, the particular mapping identified by the two options will be deleted. If you specify only the **-ug** option, all mappings for the specified group will be deleted. |

**list**

Displays information about user and group mappings. The following options are available with this command argument:

| Term | Definition |
|---|---|
| **-all** | Lists both simple and advanced mappings for users and groups. |
| **-simple** | Lists all simple mapped users and groups. |
| **-advanced** | Lists all advanced mapped users and groups. Maps are listed in the order in which they are evaluated. Primary maps, marked with an asterisk (*), are listed first, followed by secondary maps, which are marked with a carat (^). |
| **-wu** *name* | Lists the mapping for a specified Windows user. |
| **-wg** *group* | Lists the mapping for a Windows group. |
| **-uu** *name* | Lists the mapping for a UNIX user. |
| **-ug** *group* | Lists the mapping for a UNIX group. |

**backup**

Saves User Name Mapping configuration and mapping data to the file specified by *filename*.

**restore**

Replaces configuration and mapping data with data from the file (specified by *filename*) that was created using the **backup** command argument.

**adddomainmap**

Adds a simple map between a Windows domain and an NIS domain or password and group files. The following options are available for this command argument:

| Term | Definition |
|---|---|
| **-d** *WindowsDomain* | Specifies the Windows domain to be mapped. |
| **-y** *NISDomain* | Specifies the NIS domain to be mapped.<br><br>**-n *nisServer***<br>Specifies the NIS server for the NIS domain specified with the **-y** option. |
| **-f** *path* | Specifies the fully qualified path of directory containing the password and group files to be mapped. The files must be located on the computer being managed, and you cannot use **mapadmin** to manage a remote computer to set up maps based on password and group files. |

**removedomainmap**

Removes a simple map between a Windows domain and an NIS domain. The following options and argument are available for this command argument:

| Term | Definition |
|---|---|
| **-d** *WindowsDomain* | Specifies the Windows domain of the map to be removed. |

| | |
|---|---|
| **-y** *NISDomain* | Specifies the NIS domain of the map to be removed. |
| **-all** | Specifies that all simple maps between Windows and NIS domains are to be removed. This will also remove any simple map between a Windows domain and password and group files. |

**listdomainmaps**

Lists the Windows domains that are mapped to NIS domains or password and group files.

If you do not specify a command argument, **mapadmin** displays the current settings for User Name Mapping.

For all options that specify a user or group name, the following formats can be used:

For Windows users, use the form *domain**username*, *computer**username*, *computer**username*, or *computer**username*

For Windows groups, use the form *domain**groupname*, *computer**groupname*, *computer**groupname*, or *computer**groupname*

For UNIX users, use the form *NISdomain**username*, *username***@***NISdomain*, username@PCNFS, or PCNFS*username*

For UNIX groups, use the form *NISdomain**groupname*, *groupname***@***NISdomain*, groupname@PCNFS, or PCNFS*groupname*

**See Also**
**Concepts**
User Name Mapping administration

# mount

Applies To: Windows Server 2003 R2

You can use **mount** to mount Network File System (NFS) network shares.

**Syntax**
mount [-o <Option>[...]] [-u:<UserName>] [-p:{<Password> | *}] {\\<ComputerName>\<ShareName> | <ComputerName>:/<ShareName>} {<DeviceName> | *}

**Description**
The **mount** command-line utility mounts the file system identified by *ShareName* exported by the NFS server identified by *ComputerName* and associates it with the drive letter specified by *DeviceName* or, if an asterisk (*****) is used, by the first available driver letter. Users can then access the exported file system as though it were a drive on the local computer. When used without options or arguments, **mount** displays information about all mounted NFS file systems.

The **mount** utility is available only if Client for NFS is installed.

The following options and arguments can be used with the **mount** utility.

| Term | Definition |
|---|---|
| -o rsize=<buffersize> | Sets the size in kilobytes of the read buffer. Acceptable values are 1, 2, 4, 8, 16, and 32; the default is 32 KB. |
| -o wsize=<buffersize> | Sets the size in kilobytes of the write buffer. Acceptable values are 1, 2, 4, 8, 16, and 32; the default is 32 KB. |
| -o timeout=<seconds> | Sets the time-out value in seconds for a remote procedure call (RPC). Acceptable values are 0.8, 0.9, and any integer in the range 1-60; the default is 0.8. |
| -o retry=<number> | Sets the number of retries for a soft mount. Acceptable values are integers in the range 1-10; the default is 1. |
| -o mtype={soft \| hard} | Sets the mount type (default is **soft**). Regardless of the mount type, **mount** will return if it cannot immediately mount the share. Once the share has been successfully mounted, however, if the mount type is **hard**, Client for NFS will continue to try to access the share until it is successful. As a result, if the NFS server is unavailable, any Windows program trying to access the share will appear to stop responding, or "hang," if the mount type is **hard**. |
| -o anon | Mounts as an anonymous user. |
| -o nolock | Disables locking (default is **enabled**). |
| -o casesensitive | Forces file lookups on the server to be case sensitive. |

| | |
|---|---|
| -o fileaccess=<mode> | Specifies the default permission mode of new files created on the NFS share. Specify *mode* as a three-digit number in the form *ogw*, where *o*, *g*, and *w* are each a digit representing the access granted the file's owner, group, and the world, respectively. The digits must be in the range 0-7 with the following meaning:<br>• 0: No access<br>• 1: x (execute access)<br>• 2: w (write access)<br>• 3: wx<br>• 4: r (read access)<br>• 5: rx<br>• 6: rw<br>• 7: rwx |
| -o lang={euc-jp\|euc-tw\|euc-kr\|shift-jis\|big5\|ksc5601\|gb2312-80\|ansi} | Specifies the default encoding used for file and directory names and, if used, must be set to one of the following:<br>• **ansi**<br>• **big5** (Chinese)<br>• **euc-jp** (Japanese)<br>• **euc-kr** (Korean)<br>• **euc-tw** (Chinese)<br>• **gb2312-80** (Simplified Chinese)<br>• **ksc5601** (Korean)<br>• **shift-jis** (Japanese)<br><br>If this option is set to **ansi** on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:<br>• Japanese: SHIFT-JIS<br>• Korean: KS_C_5601-1987<br>• Simplified Chinese: GB2312-80<br>• Traditional Chinese: BIG5 |

| | |
|---|---|
| -o sec={sys\|krb5\|krb5i} | Specifies the security flavor. This can be set to one of the following values:<br>• **sys** (Use RPC's AUTH_SYS authentication protocol)<br>• **krb5** (Use Kerberos v5 authentication)<br>• **krb5i** (Use Kerberos v5 authentication with integrity checking) |
| -u:<UserName> | Specifies the user name to use for mounting the share. If *username* is not preceded by a backslash (****), it is treated as a UNIX user name. |
| -p:<Password> | The password to use for mounting the share. If you use an asterisk (*****), you will be prompted for the password. |

**Note**

If you make a persistent connection with **mount**, you must use umount to delete the connection. Neither the **net use** command nor Microsoft Windows Explorer will delete these connections.

**See Also**

**Concepts**

umount

# nfsadmin

Applies To: Windows Server 2003 R2

You can use **nfsadmin** to manage Server for NFS and Client for NFS.

**Syntax**

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] -l

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] -r {*client* | all}

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] {start | stop}

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] config *Option* [...]

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] creategroup *Name*

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] listgroups

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] deletegroup *Name*

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] renamegroup *OldName NewName*

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] addmembers *Name Host* [...]

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] listmembers

**nfsadmin server** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] deletemembers *Group Host* [...]

**nfsadmin client** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] {start | stop}

**nfsadmin client** [ *ComputerName* ] [ -u *UserName* [-p *Password*]] config *Option* [...]

**Description**

The **nfsadmin** command-line utility administers Server for NFS or Client for NFS on the local or remote computer running Microsoft Services for Network File System (NFS). If you are logged on with an account that does not have the required privileges, you can specify a user name and password of an account that does. The action performed by **nfsadmin** depends on the command arguments you supply.

In addition to service-specific command arguments and options, **nfsadmin** accepts the following:

***ComputerName***
Specifies the remote computer you want to administer. You can specify the computer using a Windows Internet Name Service (WINS) name or a Domain Name System (DNS) name, or by Internet Protocol (IP) address.

**-u *UserName***
Specifies the user name of the user whose credentials are to be used. It might be necessary to add the domain name to the user name in the form *domain**UserName*

**-p *Password***
Specifies the password of the user specified using the **-u** option. If you specify the **-u** option but omit the **-p** option, you are prompted for the user's password.

**Administering Server for NFS**

Use the **nfsadmin server** command to administer Server for NFS. The specific action that **nfsadmin server** takes depends on the command option or argument you specify:

**-l**
Lists all locks held by clients.

**-r {*client* | all}**
Releases the locks held by a *client* or, if **all** is specified, by all clients.

**start**
Starts the Server for NFS service.

**stop**

Stops the Server for NFS service.

**config**

Specifies general settings for Server for NFS. You must supply at least one of the following options with the **config** command argument:

**mapsvr=*server***

Sets *server* as the User Name Mapping server for Server for NFS. Although this option continues to be supported for compatibility with previous versions, you should use the **sfuadmin** utility instead.

**auditlocation={eventlog | file | both | none}**

Specifies whether events will be audited and where the events will be recorded. One of the following arguments is required.

**eventlog**

Specifies that audited events will be recorded only in the Event Viewer application log.

**file**

Specifies that audited events will be recorded only in the file specified by **config fname**.

**both**

Specifies that audited events will be recorded in the Event Viewer application log as well as the file specified by **config fname**.

**none**

Specifies that events will not be audited.

**fname=*file***

Sets the file specified by *file* as the audit file. The default is %sfudir%\log\nfssvr.log

**fsize=*size***

Sets *size* as the maximum size in megabytes of the audit file. The default maximum size is 7 MB.

**audit=[+|-]mount [+|-]read [+|-]write [+|-]create [+|-]delete [+|-]locking [+|-]all**

Specifies the events to be logged. To start logging an event, type a plus sign (**+**) before the event name; to stop logging an event, type a minus sign (**-**) before the event name. If the sign is omitted, the plus sign is assumed. Do not use **all** with any other event name.

**lockperiod=*seconds***

Specifies the number of seconds that Server for NFS will wait to reclaim locks after a connection to Server for NFS has been lost and then reestablished or after the Server for NFS service has been restarted.

**Portmapprotocol={TCP | UDP | TCP+UDP}**

Specifies which transport protocols Portmap supports. The default setting is **TCP+UDP**.

**mountprotocol={TCP | UDP | TCP+UDP}**

Specifies which transport protocols mount supports. The default setting is **TCP+UDP**.

**nfsprotocol={TCP | UDP | TCP+UDP}**

Specifies which transport protocols Network File System (NFS) supports. The default setting is **TCP+UDP**

**nlmprotocol={TCP | UDP | TCP+UDP}**

Specifies which transport protocols Network Lock Manager (NLM) supports. The default setting is **TCP+UDP**.

**nsmprotocol={TCP | UDP | TCP+UDP}**

Specifies which transport protocols Network Status Manager (NSM) supports. The default setting is **TCP+UDP**.

**enableV3={yes | no}**

Specifies whether NFS version 3 protocols will be supported. The default setting is **yes**.

**renewauth={yes | no}**

Specifies whether client connections will be required to be reauthenticated after the period specified by **config renewauthinterval**. The default setting is **no**.

**renewauthinterval=*seconds***

Specifies the number of seconds that elapse before a client is forced to be reauthenticated if **config renewauth** is set to **yes**. The default value is 600 seconds.

**dircache=*size***
Specifies the size in kilobytes of the directory cache. The number specified as *size* must be a multiple of 4 between 4 and 128. The default directory-cache size is 128 KB.

**translationfile=[*file*]**
Specifies a file containing mapping information for replacing characters in the names of files when moving them from Windows-based to UNIX-based file systems. If *file* is not specified, then file name character translation is disabled. If the value of **translationfile** is changed, you must restart the server for the change to take effect.

**dotfileshidden={yes | no}**
Specifies whether files that are created with names beginning with a period (.) will be marked as hidden in the Windows file system and consequently hidden from NFS clients. The default setting is **no**.

**casesensitivelookups={yes | no}**
Specifies whether directory lookups will be case sensitive (requiring exact matching of character case).

You also need to disable Windows kernel case-insensitivity in order for Server for NFS to support case-sensitive file names. You can disable Windows kernel case-insensitivity by clearing the following registry key to 0:

HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\kernel

DWORD "obcaseinsensitive"

**ntfscase={lower | upper | preserve}**
Specifies whether the case of characters in the names of files in the NTFS file system will be returned in lowercase, uppercase, or in the form stored in the directory. The default setting is **preserve**. This setting cannot be changed if **casesensitivelookups** is set to **yes**.

**creategroup *name***
Creates a new client group, giving it the specified *name*.

**listgroups**
Displays the names of all client groups.

**deletegroup *name***
Removes the client group specified by *name*.

**renamegroup *OldName NewName***
Changes the name of the client group specified by *OldName* to *NewName*

**addmembers *Name Host*[...]**
Adds *Host* to the client group specified by *Name*.

**listmembers *Name***
Lists the host computers in the client group specified by *Name*.

**deletemembers *Group Host*[...]**
Removes the client specified by *Host* from the client group specified by *Group*.

If you do not specify a command option or argument, **nfsadmin server** displays the current Server for NFS configuration settings.

#### Administering Client for NFS

Use the **nfsadmin client** command to administer Client for NFS. The specific action that **nfsadmin client** takes depends on the command argument you specify:

**start**
Starts the Client for NFS service.

**stop**
Stops the Client for NFS service.

**config**
Specifies general settings for Client for NFS. You must supply at least one of the following options with the **config** command argument:

**fileaccess=*mode***
Specifies the default permission mode for files created on Network File System (NFS) servers. The *mode* argument consists

of a three digits from 0 to 7 (inclusive) representing the default permissions granted the user, group, and others (respectively). The digits translate to UNIX-style permissions as follows: 0=none, 1=x, 2=w, 3=wx, 4=r, 5=rx, 6=rw, and 7=rwx. For example, **fileaccess=750** gives rwx permission to the owner, rx permission to the group, and no access permission to others.

**mapsvr=*server***
Sets *server* as the User Name Mapping server for Client for NFS. Although this option continues to be supported for compatibility with previous versions, you should use the **sfuadmin** utility instead.

**mtype={hard | soft}**
Specifies the default mount type. For a hard mount, Client for NFS continues to retry a failed RPC until it succeeds. For a soft mount, Client for NFS returns failure to the calling application after retrying the call the number of times specified by the **retry** option.

**retry=*number***
Specifies the number of times to try to make a connection for a soft mount. This value must be from 1 to 10, inclusive. The default is 1.

**timeout=*seconds***
Specifies the number of seconds to wait for a connection (remote procedure call). This value must be 0.8, 0.9, or an integer from 1 to 60, inclusive. The default is 0.8.

**Protocol={TCP | UDP | TCP+UDP}**
Specifies which transport protocols the client supports. The default setting is **TCP+UDP**

**rsize=*size***
Specifies the size, in kilobytes, of the read buffer. This value can be 0.5, 1, 2, 4, 8, 16, or 32. The default is 32.

**wsize=*size***
Specifies the size, in kilobytes, of the write buffer. This value can be 0.5, 1, 2, 4, 8, 16, or 32. The default is 32.

**perf=default**
Restores the following performance settings to default values:

- **mtype**
- **retry**
- **timeout**
- **rsize**
- **wsize**

**fileaccess=*mode***
Specifies the default permission mode for files created on Network File System (NFS) servers. The *mode* argument consists of a three digits from 0 to 7 (inclusive) representing the default permissions granted the user, group, and others (respectively). The digits translate to UNIX-style permissions as follows: 0=none, 1=x, 2=w, 3=wx, 4=r, 5=rx, 6=rw, and 7=rwx. For example, **fileaccess=750** gives rwx permission to the owner, rx permission to the group, and no access permission to others.

**SecFlavors=[+|-]sys [+|-]krb5 [+|-]krb5i**
Specifies (**+**) or removes (**-**) permitted security flavors. This can be set to any combination of the following:

- **sys** (Use RPC's AUTH_SYS authentication protocol)
- **krb5** (Use Kerberos v5 authentication)
- **krb5i** (Use Kerberos v5 authentication with integrity checking)

If you do not specify a command option or argument, **nfsadmin client** displays the current Client for NFS configuration settings.

# nfsshare

Applies To: Windows Server 2003 R2

You can use **nfsshare** to control Network File System (NFS) shares.

**Syntax**

nfsshare <ShareName>=<Drive:Path> [-o <Option=value>...]

nfsshare {<ShareName> | <Drive>:<Path> | * } /delete

**Description**

Without arguments, the **nfsshare** command-line utility lists all Network File System (NFS) shares exported by Server for NFS. With *ShareName* as the only argument, **nfsshare** lists the properties of the NFS share identified by *ShareName*. When *ShareName* and *Drive***:***Path* are provided, **nfsshare** exports the folder identified by *Drive***:***Path* as *ShareName*. When the **/delete** option is used, the specified folder is no longer made available to NFS clients.

**Options**

The **nfsshare** command accepts the following options and arguments:

| Term | Definition |
|---|---|
| -o anon={yes \| no} | Specifies whether anonymous (unmapped) users can access the shared directory. The default is **no**. |
| -o rw[=<Host>[:<Host>]...] | Provides read-write access to the shared directory by the hosts or client groups specified by *Host*. Separate host and group names with a colon (**:**). If *Host* is not specified, all hosts and client groups (except those specified with the **ro** option) have read-write access. If neither the **ro** nor the **rw** option is set, all clients have read-write access to the shared directory. |
| -o ro[=<Host>[:<Host>]...] | Provides read-only access to the shared directory by the hosts or client groups specified by *Host*. Separate host and group names with a colon (**:**). If *Host* is not specified, all clients (except those specified with the **rw** option) have read-only access. If the **ro** option is set for one or more clients, but the **rw** option is not set, only the clients specified with the **ro** option have access to the shared directory. |

| | |
|---|---|
| -o encoding={big5\|euc-jp\|euc-kr\|euc-tw\|gb2312-80\|ksc5601\|shift-jis} | Specifies the default encoding used for file and directory names and, if used, must be set to one of the following:<br>• **big5** (Chinese)<br>• **euc-jp** (Japanese)<br>• **euc-kr** (Korean)<br>• **euc-tw** (Chinese)<br>• **gb2312-80** (Simplified Chinese)<br>• **ksc5601** (Korean)<br>• **shift-jis** (Japanese)<br><br>If this is option is not set, the default encoding scheme is ANSI or, on systems configured for non-English locales, the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:<br>• Japanese: SHIFT-JIS<br>• Korean: KS_C_5601-1987<br>• Simplified Chinese: GB2312-80<br><br>Traditional Chinese: BIG5 |
| -o anongid=<gid> | Specifies that anonymous (unmapped) users will access the share directory using *gid* as their group identifier (GID). The default is -2. The anonymous GID will be used when reporting the owner of a file owned by an unmapped user, even if anonymous access is disabled. |
| -o anonuid=<uid> | Specifies that anonymous (unmapped) users will access the share directory using *uid* as their user identifier (UID). The default is -2. The anonymous UID will be used when reporting the owner of a file owned by an unmapped user, even if anonymous access is disabled. |
| -o root[=<Host>[:<Host>]...] | Provides root access to the shared directory by the hosts or client groups specified by *Host*. Separate host and group names with a colon (**:**). If *Host* is not specified, all clients have root access. If the **root** option is not set, no clients have root access to the shared directory. |
| -o sec=[<SecurityFlavor>[:<SecurityFlavor>]...] | Specifies permitted security flavors. This can be set to any combination of the following:<br>• **sys** (Use RPC's AUTH_SYS authentication protocol)<br>• **krb5** (Use Kerberos v5 authentication)<br>• **krb5i** (Use Kerberos v5 authentication with integrity checking) |
| /delete | If *ShareName* or *Drive***:***Path* is specified, deletes the specified share. If * is specified, deletes all NFS shares. |

| **Note** |
|---|
| To view the complete syntax for this command, at a command prompt, type:<br><br>**nfsshare /?** |

**See Also**

**Concepts**

Managing NFS shares

# nfsstat

Applies To: Windows Server 2003 R2

You can use **nfsstat** to display or reset counts of calls made to Server for NFS.

**Syntax**
**nfsstat** [-z]

**Description**
When used without the **-z** option, the **nfsstat** command-line utility displays the number of NFS V2, NFS V3, and Mount V3 calls made to the server since the counters were set to 0, either when the service started or when the counters were reset using **nfsstat -z**.

# rpcinfo

Applies To: Windows Server 2003 R2

You can use **rpcinfo** to list programs on remote computers.

**Syntax**

List all programs registered with the port mapper:

**rpcinfo** -p [*Node*]

Request a response from network nodes that have a specified program:

**rpcinfo** -b *Program Version*

Use Transmission Control Protocol (TCP) to call a program:

**rpcinfo** -t *Node Program* [*Version*]

Use User Datagram Protocol (UDP) to call a program:

**rpcinfo** -u *Node Program* [*Version*]

**Description**

The **rpcinfo** command-line utility makes a remote procedure call (RPC) to an RPC server and reports what it finds.

The **rpcinfo** utility accepts the following options:

**-p [*Node*]**
Lists all programs registered with the port mapper on the specified host. If you do not specify a node (computer) name, the program queries the port mapper on the local host.

**-b *Program Version***
Requests a response from all network nodes that have the specified program and version registered with the port mapper. You must specify both a program name or number and a version number.

**-t *Node Program* [*Version*]**
Uses the TCP transport protocol to call the specified program. You must specify both a node (computer) name and a program name. If you do not specify a version, the program calls all versions.

**-u *Node Program* [*Version*]**
Uses the UDP transport protocol to call the specified program. You must specify both a node (computer) name and a program name. If you do not specify a version, the program calls all versions.

# showmount

Applies To: Windows Server 2003 R2

You can use **showmount** to display mounted directories.

**Syntax**

**showmount** { -e | -a | -d} [*Server*]

**Description**

The **showmount** command-line utility displays information about mounted file systems exported by Server for NFS on the computer specified by *Server*. If *Server* is not provided, **showmount** displays information about the computer on which the **showmount** command is run.

You must provide one of the following options:

**-e**

Displays all file systems exported on the server.

**-a**

Displays all Network File System (NFS) clients and the directories on the server each has mounted.

**-d**

Displays all directories on the server that are currently mounted by NFS clients.

**See Also**

**Concepts**

Server for NFS administration

# umount

Applies To: Windows Server 2003 R2

You can use **umount** to remove Network File System (NFS)–mounted drives.

**Syntax**
**umount** [-f] [{-a | *DriveLetter*:[...] | *NetworkMount*[...]}]

**Description**
The **umount** command-line utility disconnects the specified NFS-mounted drive. You must supply at least one of the following options or arguments.

**-f**
Forces deletion of Network File System (NFS) network drives.

**-a**
Deletes all NFS network drives. If there are active connections, **umount** prompts you for confirmation unless you also use the **-f** option.

***DriveLetter***
The letter of the logical drive to be disconnected.

***NetworkMount***
The network mount point to be disconnected. This mount must have been created using the **net use** Windows command-line utility without specifying a drive letter.

**See Also**
**Concepts**
mount

# DFS Management

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use the Distributed File System (DFS) Management snap-in to manage DFS Namespaces and DFS Replication. When used together, these two technologies offer simplified, fault-tolerant access to files, load sharing, and WAN-friendly replication.

- To install the DFS Management snap-in, see Installing the DFS Management Snap-in and DFS Replication Service.
- For help with specific tasks, see Creating and Managing Namespaces and Configuring and Managing DFS Replication.
- For information about property pages in the DFS Management snap-in, see Namespace Properties and DFS Replication Properties.
- For more information about the DFS solution in the Microsoft® Windows Server™ 2003 R2 operating systems, see the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=39468).

# Installing the DFS Management Snap-in and DFS Replication Service

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

If you plan to use DFS Replication, you must first update the Active Directory® directory service schema to install the Active Directory classes and attributes used by DFS Replication. To upgrade the schema, on the schema operations master, run **adprep.exe /forestprep**. The adprep.exe command-line tool is available in the cmpnents\r2\adprep folder on the second Windows Server 2003 R2 operating system installation CD.

After the Windows Server 2003 R2 operating system is installed, and the schema is updated as necessary, you can install the Distributed File System (DFS) components by using the following procedures.

| **Note** |
| --- |
| Two methods for installing Distributed File System are provided. The method using the **Manage Your Server** file server tool allows you to install Distributed File System as a part of the file server role. This method also installs other file server tools, such as the **File Server Resource Manager** and **File Server Management**. The method using **Add or Remove Programs** allows you to install either the DFS Management snap-in, the DFS Replication service, or both. |

Use the following procedure if the file server role has not been added.

**To install Distributed File System using Manage Your Server**

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Manage Your Server**.
2. Click **Add or remove a role**.
3. In the Configure Your Server Wizard, click **Next**. On the **Server Role** page, click **File Server**, and then click **Next.** Click **Next** again.
4. In the Add File Server Role Wizard, click **Next**, and then click **Replicate data to and from this server**. Follow the steps in the wizard to complete the installation. You will be asked to restart the server when the installation is complete.

Use the following procedure to install the DFS Management snap-in, the DFS Replication service, or both, without installing the File Server role.

**To install Distributed File System using Add or Remove Programs**

1. In Control Panel, double-click **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**, highlight **Distributed File System**, and then click **Details**.
3. Do one of the following:
   - To install the Distributed File System management tools without installing the DFS Replication service, click **DFS Management**.
   - To install the DFS Replication service without installing the management tools, click **DFS Replication Service**.

# Namespaces

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- New DFS Namespaces Terminology and Features in Windows Server 2003 R2
- Namespaces Overview
- Getting Started With Namespaces
- Creating and Managing Namespaces
- Namespace Properties

# New DFS Namespaces Terminology and Features in Windows Server 2003 R2

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- Updated namespace terminology
- New namespace features

# Updated namespace terminology

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

To make namespace concepts easier to understand, namespace terminology has been simplified. The following table describes the previous and updated terms.

| **Previous Term** | **Updated Term** | **Definition** |
|---|---|---|
| Link | Folder | Any folder that appears in the namespace after *ServerOrDomainName**RootName*. A folder can have optional folder targets. |
| Link target | Folder target | A Universal Naming Convention (UNC) path of a shared folder or another namespace that is associated with a folder in a namespace. |
| Not defined | Namespace root | The top-level folder in the namespace. For example, *ServerOrDomainName**RootName* is the namespace root. |
| Root | Namespace | A virtual tree of folders that begins with *ServerOrDomainName**RootName*. |
| Root server | Namespace server | A server that hosts a namespace. |

# New namespace features

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The Distributed File System (DFS) technology in the Microsoft Windows Server 2003 operating systems has been renamed to DFS Namespaces. Although the underlying service and basic functionality are unchanged, there are a number of new namespace settings exposed in the Windows Server 2003 R2 operating systems. These settings are as follows:

**Target priority**

When a client accesses a namespace, the client receives a referral that contains a list of targets associated with the namespace root or folder. These targets are listed according to the current ordering method for the namespace or folder. To fine-tune how particular targets are ordered, you can specify whether a server appears first or last in a referral. Assigning target priority is useful in many scenarios, such as “hot-standby” scenarios where one server is considered the server of last resort. In this scenario, you specify that the standby server always appears last in referrals, and clients will fail over to this server only if all the other servers fail or become unavailable due to network outages. For more information about target priority, see Target priority.

**Client failback**

Client failover in DFS Namespaces is the process in which clients attempt to access another server in a referral after one of the servers fails or is removed from the namespace. Unless client failback is configured, clients will continue using the server they failed over to unless the client is restarted or the client's referral cache is cleared. When client failback is configured, and clients have the appropriate client failback hotfix installed, clients will fail back to a preferred, local server when it is restored. For more information about client failback, see Referral properties.

**Better delegation**

You can easily delegate the ability to create domain-based namespaces and manage individual stand-alone and domain-based namespaces. The DFS Management snap-in sets the appropriate permissions on either the DFS Namespace configuration objects in Active Directory or in the namespace server’s registry (depending on the namespace type). For more information about delegation, see Security requirements for creating and managing namespaces.

**Ability to restructure the namespace**

Renaming or moving folders in the namespace is easy when you use the DFS Management snap-in. You can restructure the namespace to correct mistakes or to adjust the hierarchy as business needs change or as new folders are added to the namespace. You can also move namespace folders by using the updated version of the command-line tool, Dfscmd.exe.

# Namespaces Overview

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- Introduction to namespaces
- Types of namespaces
- Namespace folders and folder targets
- Namespace servers
- Namespace limits
- Namespace clients
- Security requirements for creating and managing namespaces
- Considerations for hosting namespaces on servers running a mix of Windows Server 2003 and Windows 2000 Server

# Introduction to namespaces

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

With DFS Namespaces technology, you can group shared folders located on different servers by transparently connecting them to one or more namespaces. A *namespace* is a virtual view of shared folders in an organization. When a user views the namespace, the folders appear to reside on a single hard disk. Users can navigate the namespace without needing to know the server names or shared folders hosting the data.

The path to a namespace is similar to a Universal Naming Convention (UNC) path to a shared folder, such as \\Server1\Public\Software\Tools. If you are familiar with UNC paths, you know that in this example, the shared folder Public and its subfolders Software and Tools are all hosted on Server1. Now, assume that you want to give users a single place to locate data, but you want to host data on different servers for availability and performance purposes. To do this, you can deploy a namespace similar to the one shown in the following figure. The elements of this namespace are described after the figure.




- **Namespace server**. A namespace server hosts a namespace. The namespace server can be a member server or domain controller.

- **Namespace root**. The namespace root is the starting point of the namespace. In the previous figure, the name of the root is Public, and the namespace path is \\Contoso\Public. This type of namespace is known as a domain-based namespace because it begins with a domain name (for example, Contoso) and its metadata is stored in the Active Directory® directory service. Although a single namespace server is shown in the previous figure, a domain-based namespace can be hosted on multiple namespace servers to increase the availability of the namespace.

- **Folder**. Folders help build the namespace hierarchy. Folders can optionally have folder targets. When users browse a folder with targets in the namespace, the client computer receives a referral that transparently redirects the client computer to one of the folder targets.

- **Folder targets**. A folder target is a UNC path of a shared folder or another namespace that is associated with a folder in a namespace. In the previous figure, the folder named Tools has two folder targets, one in London and one in New York, and the folder named Training Guides has a single folder target in New York. A user who browses to \\Contoso\Public\Software\Tools is transparently redirected to the shared folder \\LDN-SVR-01\Tools or \\NYC-SVR-01\Tools, depending on which site the user is in.

**See Also**
**Other Resources**
Microsoft Web site

# Types of namespaces

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When creating a namespace, you create either a stand-alone namespace or a domain-based namespace. The characteristics of each namespace type are described in the following table.

| Characteristic | Domain-based Namespace | Stand-Alone Namespace |
|---|---|---|
| Path to namespace | *NetBIOSDomainName**RootName*<br>*DNSDomainName**RootName* | *ServerName**RootName* |
| Where namespace information is stored | In Active Directory and in a memory cache on each namespace server | In the registry and in a memory cache on the namespace server |
| Namespace size recommendations | We recommend that you keep the size of the namespace object in Active Directory to less than 5 megabytes (MB) by using fewer than 5,000 folders with targets in domain-based namespaces. (The namespace object is where namespace metadata is stored in Active Directory.) | The largest recommended namespace size for a stand-alone namespace is 50,000 folders with targets. |
| Supported methods to ensure namespace availability | Use multiple namespace servers to host the namespace. (The namespace servers must be in the same domain.) | Create a stand-alone namespace on a server cluster using the Cluster Administrator snap-in. |
| Support for using DFS Replication to replicate folder targets | Supported | Supported |

Use the following guidelines to choose a namespace type.

Choose stand-alone namespaces if:

- Your organization does not use Active Directory.
- You need to create a single namespace with more than 5,000 folders with targets. (If you can divide your folder targets among two or more namespaces, a domain-based namespace is an option.)
- You want to ensure the availability of the namespace by using a server cluster.

Choose domain-based namespaces if:

- You want to ensure the availability of the namespace by using multiple namespace servers.
- You want to hide the name of the namespace server from users because you plan to replace the namespace server or migrate the namespace to another server in the future.

**See Also**
**Other Resources**
Microsoft Web site

# Namespace folders and folder targets

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Folders are the primary elements of a namespace. They appear after the namespace root (*ServerOrDomainName**RootName*) and help build the hierarchy of the namespace. When you create a folder using the DFS Management snap-in, you type a name for the folder and specify whether to add any folder targets. A folder target is a Universal Naming Convention (UNC) path to one of the following locations:

- A shared folder
- A folder within a shared folder
- A path to another namespace

You can specify multiple folder targets to increase the redundancy of the folder. If one of the folder targets is not available, the client will attempt to access one of the other available targets. (The order in which clients attempt to access targets is defined by the ordering method described in Referral properties.)

When adding folder targets, it is important to follow these guidelines:

- If you specify a domain-based namespace as a folder target (either the namespace root or a folder within that namespace), you cannot specify alternate folder targets. (Windows Server 2003 R2 operating systems enforce this restriction.)
- If you plan to specify a stand-alone namespace as a folder target (either the namespace root or a folder within that namespace), you can specify alternate folder targets that are either stand-alone namespace roots or folders within the stand-alone namespace. Do not specify domain-based namespace roots or shared folders as alternate targets.

**See Also**
**Other Resources**
Microsoft Web site

# Namespace servers

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A namespace server is a domain controller or member server that hosts a namespace. The number of namespaces you can host on a server is determined by the operating system running on the namespace server.

Servers running the following operating systems can host a single namespace:

- Windows Server 2003 R2, Standard Edition
- Windows Server 2003, Web Edition
- Windows Server 2003, Standard Edition
- Any version of Windows® 2000 Server

| **Important** |
| --- |
| To create multiple domain-based namespaces on a server running Windows Server 2003, Standard Edition, install the hotfix described in article 903651 in the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink?linkid=53476). |

The following servers can host multiple namespaces:

- Windows Server 2003 R2, Enterprise Edition
- Windows Server 2003 R2, Datacenter Edition
- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition

The following table describes additional factors to consider when choosing servers to host a namespace.

**Guidelines for Servers That Host Namespaces**

| Server Hosting Stand-Alone Namespaces | Server Hosting Domain-based Namespaces |
| --- | --- |
| • Must contain an NTFS volume to host the namespace.<br>• Can be a member server or domain controller.<br>• Can be a clustered file server. | • Must contain an NTFS volume to host the namespace.<br>• Must be a member server or domain controller in the domain in which the namespace is configured. (This requirement applies to every namespace server that hosts a given domain-based namespace.)<br>• The namespace cannot be a clustered resource in a server cluster. However, you can locate the namespace on a server that also functions as a node in a server cluster if you configure the namespace to use only local resources on that server. |

**See Also**
**Other Resources**
Microsoft Web site

# Namespace limits

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When adding folders to namespaces, keep the following recommendations in mind:

- For stand-alone namespaces, we recommend a maximum of 50,000 folders with targets.
- For domain-based namespaces, we recommend keeping the namespace object in Active Directory to 5 megabytes (MB) or less (equivalent to about 5,000 folders with targets).

| **Note** |
| --- |
| Folders without folder targets do not count against these limits. |

To make it easy for you to check the size of your namespace, the namespace size is shown in the **Namespace Properties** dialog box. To view this, right-click the namespace in the console tree, and then click **Properties**.

**See Also**
**Other Resources**
Microsoft Web site

# Namespace clients

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Clients running versions of the following operating systems can access namespaces:

- Windows Server 2003 R2
- Windows Server 2003
- Windows Storage Server 2003
- Windows XP
- Windows Preinstallation Environment (Windows PE) (Can access stand-alone namespaces, but cannot access domain-based namespaces.)
- Windows 2000 Server
- Windows 2000 Professional
- Windows NT® Server 4.0 with Service Pack 6a
- Windows NT Workstation 4.0 with Service Pack 6a

**See Also**
**Other Resources**
Microsoft Web site

# Security requirements for creating and managing namespaces

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The following table describes the groups that can perform basic namespace tasks by default, and the method for delegating the ability to perform these tasks.

| Task | Groups That Can Perform This Task By Default | Delegation Method |
|---|---|---|
| Create a domain-based namespace | Domain Admins group in the domain where the namespace is configured | Right-click the **Namespace** node in the console tree, and then click **Delegate Management Permissions**. You must also add the user to the Local Administrators group on the namespace server. |
| Add a namespace server to a domain-based namespace | Domain Admins group in the domain where the namespace is configured | Right-click the domain-based namespace in the console tree, and then click **Delegate Management Permissions**. You must also add the user to the Local Administrators group on the namespace server to be added. |
| Manage a domain-based namespace | Local Administrators group on each namespace server | Right-click the domain-based namespace in the console tree, and then click **Delegate Management Permissions**. |
| Create a stand-alone namespace | Local Administrators group on the namespace server | Add the user to the local Administrators group on the namespace server. |
| Manage a stand-alone namespace* | Local Administrators group on the namespace server | Right-click the stand-alone namespace in the console tree, and then click **Delegate Management Permissions**. |
| Create a replication group or enable DFS Replication on a folder | Domain Admins group in the domain where the namespace is configured | Right-click the **Replication** node in the console tree, and then click **Delegate Management Permissions**. |

* Delegating management permissions to manage a stand-alone namespace does not grant the user the ability to view and manage security using the **Delegation** tab. This issue occurs because the DFS Management snap-in cannot retrieve the discretionary access control lists (DACLs) for the stand-alone namespace from the registry. To enable the snap-in to display delegation information, you must follow the steps in article 314837 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink?linkid=46803).

**See Also**
**Other Resources**
Microsoft Web site

# Considerations for hosting namespaces on servers running a mix of Windows Server 2003 and Windows 2000 Server

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

It is possible to host namespaces on domain controllers and member servers running a mix of Windows 2000 Server, Windows Server 2003 R2, and Windows Server 2003 without Service Pack 1 (SP1), but some functionality is disabled or available inconsistently, depending on the operating systems on the servers. Some examples of mixed-mode behavior are as follows:

- If domain controllers or namespace servers are running Windows Server 2003 without SP1 or without Windows Server 2003 R2, they cannot provide referrals that support target priority or client failback.

- If domain controllers or namespace servers are running Windows 2000 Server, they cannot provide referrals that support target server priority or client failback; nor can they order targets by lowest cost in referrals. Additional configuration is required to enable these namespace servers and domain controllers to detect the site of each target server in the namespace. For details, see Update the static site table on namespace servers running Windows 2000 Server.

- If the DFS Management snap-in connects to a namespace server that is not running Windows Server 2003 with SP1 or R2 versions of Windows Server 2003, none of the new configuration settings (such as client failback or target priority) can be enabled. Renaming or moving namespace folders will not work, and delegation will not be effective.

**See Also**
**Other Resources**
Microsoft Web site

# Getting Started With Namespaces

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**How to...**

- Deploy a namespace for publishing content
- Increase the availability of a namespace

# Deploy a namespace for publishing content

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

DFS Namespaces and DFS Replication can be used to publish documents, software, and line-of-business data to users throughout an organization. Although DFS Replication alone is sufficient to distribute data, you can use DFS Namespaces to configure the namespace so that a folder in the namespace is hosted by multiple servers, each of which holds a replica of the published content in that folder. This increases data availability and distributes the client load across servers.

When browsing a folder in the namespace, users are not aware that the folder is hosted by multiple servers. When a user opens the folder, the client computer is automatically referred to a server in its own site; if no same-site servers are available, you can configure the namespace to refer the client to a server that has the lowest connection cost as defined in Active Directory. In addition, using a DFS Namespaces enhancement known as target priority, you can specify the priority of servers so that a specific server is always placed first or last in the list of servers (known as a referral) the client receives when accessing a folder with targets in the namespace.

To configure a namespace for publication, perform the following procedures as specified:

1. Create a namespace
2. Create a folder in a namespace
3. Add folder targets
4. Set the ordering method for targets in referrals
5. Set target priority to override referral ordering (Optional)
6. Enable client failback (Optional)
7. Replicate folder targets using DFS Replication (Optional)

# Increase the availability of a namespace

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

To ensure that the namespace and its folders are still available if a server fails, you can add multiple namespace servers (for domain-based namespaces) and multiple folder targets for each folder. You can optionally use DFS Replication to keep the contents of the folder targets synchronized.

The procedures for increasing the availability of a namespace are as follows:

1. Add namespace servers to a domain-based namespace
2. Add folder targets
3. Replicate folder targets using DFS Replication (Optional)

| **Note** |
| --- |
| To increase the availability of a stand-alone namespace, you must use the Cluster Administrator snap-in to create a stand-alone namespace on a server cluster. |

# Creating and Managing Namespaces

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- Deploying a Namespace
- Increasing the Availability of a Namespace
- Optimizing a Namespace
- Managing Namespaces Hosted on Multiple Server Operating Systems

# Deploying a Namespace

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**How to...**

- Create a namespace
- Create a folder in a namespace
- Delegate management permissions for an existing namespace
- Delegate the ability to create new domain-based namespaces

# Create a namespace

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To create a namespace**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, right-click the **Namespaces** node, and then click **New Namespace**.
3. Follow the instructions in the New Namespace Wizard.

**Note**

- Security requirements for performing this task are described in Security requirements for creating and managing namespaces.

**See Also**
**Concepts**
Delegate the ability to create new domain-based namespaces
Namespace servers
Types of namespaces
Namespace folders and folder targets
**Other Resources**
Microsoft Web site

# Create a folder in a namespace

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To create a folder in a namespace**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a namespace or a folder within a namespace, and then click **New Folder**.
3. In the **Name** box, enter the name of the new folder.
4. To add one or more folder targets to the folder, click **Add** and specify the Universal Naming Convention (UNC) path of the folder target, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Local Administrators group on each server that hosts the namespace, or you must have been delegated the ability to manage an existing namespace.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
Namespace folders and folder targets
**Other Resources**
Microsoft Web site

# Delegate management permissions for an existing namespace

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To delegate management permissions for an existing namespace**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click the namespace for which you want to delegate management permissions, and then click **Delegate Management Permissions**.
3. Type user names or group names separated by commas, click **Check Names** to verify the objects, and then click **OK**.

**Note**

- Security requirements for performing this task are described in Security requirements for creating and managing namespaces.

**See Also**
**Other Resources**
Microsoft Web site

# Delegate the ability to create new domain-based namespaces

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To delegate the ability to create new domain-based namespaces**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. Right-click the **Namespaces** node, and then click **Delegate Management Permissions**.
3. Click **Add** to add users or groups to the list of delegated users and groups.
4. Type user names or group names separated by commas, click **Check Names** to verify the objects, and then click **OK**.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
**Other Resources**
Microsoft Web site

# Increasing the Availability of a Namespace

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**How to...**

- Add namespace servers to a domain-based namespace
- Add folder targets
- Replicate folder targets using DFS Replication

# Add namespace servers to a domain-based namespace

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can increase the availability of a domain-based namespace by specifying additional namespace servers to host the namespace.

**To add a namespace server to a domain-based namespace**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a domain-based namespace, and then click **Add Namespace Server**.
3. Enter the path to another server or click **Browse** to locate a server.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority. You must also be a member of the local Administrators group on the servers to be added to the namespace.
- This procedure is not applicable for stand-alone namespaces because stand-alone namespaces can only support a single namespace server. To increase the availability of a stand-alone namespace, you must use the Cluster Administrator snap-in to create the namespace on a server cluster.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
Namespace servers
Considerations for hosting namespaces on servers running a mix of Windows Server 2003 and Windows 2000 Server
**Other Resources**
Microsoft Web site

# Add folder targets

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A folder target is a Universal Naming Convention (UNC) path of a shared folder or another namespace that is associated with a folder in a namespace. Adding multiple folder targets increases the availability of the folder in the namespace.

**To add a folder target**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a folder, and then click **Add Folder Target**.
3. Type the path to the folder target or click **Browse** to locate the folder target.
4. If the folder is replicated using DFS Replication, you can specify whether to add the new folder target to the replication group.

**Note**

- To perform this procedure, you must be a member of the Local Administrators group on each server that hosts the namespace, or you must have been delegated the ability to manage an existing namespace.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
Replicate folder targets using DFS Replication
**Other Resources**
Microsoft Web site

# Replicate folder targets using DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use DFS Replication to keep the contents of folder targets in sync so that users see the same files regardless of which folder target the client computer is referred to.

**To replicate folder targets using DFS Replication**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a folder that has two or more folder targets, and then click **Replicate Folder**.
3. Follow the instructions in the Replicate Folder Wizard.

**Important**

- Replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate management permissions for an existing namespace
Introduction to DFS Replication
DFS Replication requirements
Replication groups and replicated folders
Replication topologies
Replication schedules and bandwidth
What to expect during initial replication
**Other Resources**
Microsoft Web site

# Optimizing a Namespace

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**How to...**

- Rename or move a folder
- Enable or disable referrals to a folder target
- Change the amount of time that clients cache referrals
- Set the ordering method for targets in referrals
- Set target priority to override referral ordering
- Enable client failback
- Optimize namespace polling

# Rename or move a folder

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

### To rename a folder

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a folder, and then click **Rename Folder**.
3. Type the new name in the **New Name** field.

### To move a folder

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a folder, and then click **Move Folder**.
3. Enter the folder path in the **New location** field or click **Browse** to locate a folder.

### Notes

- To perform this procedure, you must be a member of the Local Administrators group on each server that hosts the namespace, or you must have been delegated the ability to manage an existing namespace.
- You can also drag and drop folders in the console tree.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
**Other Resources**
Microsoft Web site

# Enable or disable referrals to a folder target

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A referral is an ordered list of targets that a client computer receives from a domain controller or namespace server when the user accesses a namespace root or folder with targets in the namespace. After the client receives the referral, the client attempts to access the first target in the list; if the target is not available, the client attempts to access the next target.

By disabling a folder target's referral, you can prevent the client from ever being directed to that folder target when accessing the folder in the namespace.

**To enable or disable referrals to a folder target**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, click a folder with targets, and then click the **Folder Targets** tab in the details pane.
3. Right-click the folder target, and then click one of the following options:
   - To enable the referral, select **Enable Folder Target**.
   - To disable the referral, select **Disable Folder Target**.

**Note**

- To perform this procedure, you must be a member of the Local Administrators group on each server that hosts the namespace, or you must have been delegated the ability to manage an existing namespace.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
Referral properties
**Other Resources**
Microsoft Web site

# Change the amount of time that clients cache referrals

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A referral is an ordered list of targets that a client computer receives from a domain controller or namespace server when the user accesses a namespace root or folder with targets in the namespace. You can adjust how long clients cache a referral before requesting a new one.

**To change the amount of time that clients cache namespace root referrals**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a namespace, and then click **Properties**.
3. On the **Referrals** tab, in the **Cache duration (in seconds)** box, type the amount of time (in seconds) that clients cache namespace root referrals. The default setting is 300 seconds or five minutes.

**To change the amount of time that clients cache folder referrals**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a folder that has targets, and then click **Properties**.
3. On the **Referrals** tab, in the **Cache duration (in seconds)** box, type the amount of time (in seconds) that clients cache folder referrals in the **Cache duration (in seconds)** box. The default setting is 1800 seconds or 30 minutes.

**Note**

- To perform this procedure, you must be a member of the Local Administrators group on each server that hosts the namespace, or you must have been delegated the ability to manage an existing namespace.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
Referral properties
**Other Resources**
Microsoft Web site

# Set the ordering method for targets in referrals

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A referral is an ordered list of targets that a client computer receives from a domain controller or namespace server when the user accesses a namespace root or folder with targets in the namespace. After the client receives the referral, the client attempts to access the first target in the list; if the target is not available, the client attempts to access the next target.

Targets in the client's site are always listed first in a referral. Targets outside of the client's site are listed according to the ordering method. The ordering methods are described in Referral properties.

Use the following procedure to set the ordering method on the namespace root.

**To set the ordering method for targets in namespace root referrals**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a namespace, and then click **Properties**.
3. On the **Referrals** tab, select an ordering method.

Folders with targets inherit the ordering method from the namespace root. You can override the ordering method by using the following procedure.

**To set the ordering method for targets in folder referrals**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a folder with targets, and then click **Properties**.
3. On the **Referrals** tab, select the **Exclude targets outside of the client's site** check box.

**Note**

- To perform this procedure, you must be a member of the Local Administrators group on each server that hosts the namespace, or you must have been delegated the ability to manage an existing namespace.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
Referral properties
**Other Resources**
Microsoft Web site

# Set target priority to override referral ordering

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A referral is an ordered list of targets that a client computer receives from a domain controller or namespace server when the user accesses a namespace root or folder with targets in the namespace. Each target in a referral is ordered according to the ordering method for the namespace root or folder. To refine how targets are ordered, you can set priority on individual targets. For example, you can specify that the target is first among all targets, last among all targets, or first (or last) among all targets of equal cost.

To set target priority on a root target for a domain-based namespace, use the following procedure.

**To set target priority on a root target for a domain-based namespace**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, click the namespace whose root targets you want to set priority on.
3. In the details pane, on the **Namespace Servers** tab, right-click the root target whose priority you want to change, and then click **Properties**.
4. On the **Advanced** tab, click **Override referral ordering**, and then click the priority you want.

To set target priority on a folder target, use the following procedure.

**To set target priority on a folder target**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, click the folder whose targets you want to set priority on.
3. In the details pane, on the **Folder Targets** tab, right-click the folder target whose priority you want to change, and then click **Properties**.
4. On the **Advanced** tab, click **Override referral ordering**, and then click the priority you want.

**Note**

- To perform this procedure, you must be a member of the Local Administrators group on each server that hosts the namespace, or you must have been delegated the ability to manage an existing namespace.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
Target priority
**Other Resources**
Microsoft Web site

# Enable client failback

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A referral is an ordered list of targets that a client computer receives from a domain controller or namespace server when the user accesses a namespace root or folder with targets in the namespace. If a target becomes unavailable, you can configure clients to fail back to the target after it is restored.

**To enable client failback for a namespace root**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a namespace, and then click **Properties**.
3. On the **Referrals** tab, click the **Clients fail back to preferred targets** check box.

Folders with targets inherit client failback settings from the namespace root. If client failback is disabled on the namespace root, you can use the following procedure to enable client failback on a folder with targets.

**To enable client failback for a folder with targets**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a folder with targets, and then click **Properties**.
3. On the **Referrals** tab, click the **Clients fail back to preferred targets** check box.

**Notes**

- To perform this procedure, you must be a member of the Local Administrators group on each server that hosts the namespace, or you must have been delegated the ability to manage an existing namespace.
- For failback to work, client computers must run the operating systems and hotfixes described in Referral properties.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
Referral properties
**Other Resources**
Microsoft Web site

# Optimize namespace polling

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

To maintain a consistent domain-based namespace across namespace servers, it is necessary for namespace servers to periodically poll Active Directory to obtain the most current namespace data. You can use this procedure to optimize how this polling occurs.

**To optimize namespace polling**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Namespaces** node, right-click a namespace, and then click **Properties**.
3. On he **Advanced** tab, select whether the namespace will be optimized for consistency or scalability.

**Note**

- To perform this procedure, you must be a member of the Local Administrators group on each server that hosts the namespace, or you must have been delegated the ability to manage an existing namespace.

**See Also**
**Concepts**
Security requirements for creating and managing namespaces
Delegate management permissions for an existing namespace
Polling properties
**Other Resources**
Microsoft Web site

# Managing Namespaces Hosted on Multiple Server Operating Systems

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**How to...**

- Update the static site table on namespace servers running Windows 2000 Server

# Update the static site table on namespace servers running Windows 2000 Server

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

If you want referrals from namespace servers and domain controllers running Microsoft Windows 2000 Server operating systems to be ordered according to site information, you can use the /UpdateWin2kStaticSiteTable parameter in Dfsutil.exe to update the static site information for all root targets and folder targets in the namespace object in Active Directory.

**To update the static site table**

1. Install Dfsutil.exe, which is part of the Windows Support Tools. You can install the Windows Support Tools by running Suptools.msi in the \Support\Tools folder on your Windows Server 2003 R2 operating system CD.
2. At the command prompt, type the following command:

   **Dfsutil /root:<Namespacepath> /UpdateWin2kStaticSiteTable**

   For example, to update the namespace \\Contoso\Public, type:

   **Dfsutil /root:\\Contoso\Public /UpdateWin2kStaticSiteTable**

**Notes**

- This procedure applies only to domain-based namespaces.
- Using this parameter increases the size of the namespace object in Active Directory, possibly making it exceed the recommended 5-MB size limit.
- This command must be run each time a target server's site information is changed, such as after the target is moved to another site or after the target's site name changes.

**See Also**
**Concepts**
Considerations for hosting namespaces on servers running a mix of Windows Server 2003 and Windows 2000 Server
**Other Resources**
Microsoft Web site

# Namespace Properties

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- Referral properties
- Polling properties
- Target priority

# Referral properties

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The following options are available when you view the properties of a namespace root or folder with targets.

**Amount of time clients cache referrals**

**When this option is available:** Available when the DFS Management snap-in connects to a namespace server that is running a version of one of the following operating systems: Windows Server 2003 R2, Windows Server 2003 with Service Pack 1 (SP1), Windows Server 2003, and Windows 2000 Server.

**Supported namespace types:** Domain-based; stand-alone.

**Supported namespace server operating systems:** Windows Server 2003 and Windows 2000 Server.

**Supported clients:** All supported namespace clients.

**Where configured:** Namespace properties, Referrals tab; Folder properties, Referrals tab.

For clients that are not running Windows XP with Service Pack 2 (SP2) or Windows Server 2003 with SP1, the cache duration determines the earliest time that a client computer will request a new referral if the existing cached referral expires before it is accessed again. Clients that use a cached referral will renew the cache duration value of the referral each time a file or folder is accessed using the referral and, thus, use the referral indefinitely until the client's referral cache is cleared or the client is restarted.

This behavior is different for clients running Windows XP with SP2 or Windows Server 2003 with SP1. Specifically, the cache duration is not renewed each time a client accesses a target using a cached referral. Instead, the referral expires after the cache duration value lapses.

One benefit of this change is that namespace clients running Windows XP with SP2 or Windows Server 2003 with SP1 will more quickly discover changes to namespaces and namespace folders. For example, if the folder targets of a folder named Current are changed daily, namespace clients without Windows XP with SP2 or Windows Server 2003 with SP1 would refresh the duration value each time they access the Current folder, causing them to continue to reference stale folder targets well beyond the duration value associated with the initial referral request.

To change the amount of time that clients cache referrals, see Change the amount of time that clients cache referrals.

**Ordering methods for targets in referrals**

**When this option is available:** Availablewhen the DFS Management snap-in connects to a namespace server that is running a version of one of the following operating systems: Windows Server 2003 R2 or Windows Server 2003 with SP1.

**Supported namespace types:** Domain-based; stand-alone.

**Supported namespace server operating systems:** Windows Server 2003; namespace servers running Windows 2000 Server cannot provide referrals based on lowest cost.

**Supported clients:** All supported namespace clients.

**Where configured:** Namespace properties, Referrals tab; Folder properties, Referrals tab.

When a client computer attempts to access a namespace, a domain controller or namespace server provides a referral to the client. The referral contains a list of target servers that are sorted according to the currently configured ordering method. When a client accesses the namespace root or a folder with targets in the namespace, the client attempts to access the first target at the top of the referral; the client moves to the next target if the prior target was not available.

The three ordering methods are:

**Random order**

In this method, targets are ordered as follows:

1. Targets in the same Active Directory site as the client are listed in random order at the top of the referral.

2. Next, targets outside of the client's site are listed in random order.

If no same-site target servers are available, the client computer is referred to a random target server no matter how expensive the connection is or how distant the target is.

**Lowest cost**

In this method, targets are ordered as follows:

1. Targets in the same site as the client are listed in random order at the top of the referral.

2. Next, targets outside of the client's site are listed in order of lowest cost to highest cost. Referrals with the same cost are grouped together and within each group the targets are listed in random order.

| **Note** |
| --- |
| Site link costs are not shown in the DFS Management snap-in. To view site link costs, use the Active Directory Sites and Services snap-in. |

**Exclude targets outside of the client's site**

In this method, the referral contains only the targets that are in the same site as the client. These same-site targets are listed in random order. If no same-site targets exist, the client does not receive a referral and cannot access that portion of the namespace.

| **Note** |
| --- |
| Targets that have target priority set to "First among all targets" or "Last among all targets" are still listed in the referral, even if the ordering method is set to **Exclude targets outside of the client's site**. |

To change the ordering method, see Set the ordering method for targets in referrals.

**Clients failback to preferred targets**

**When this option is available:** Availablewhen the DFS Management snap-in connects to a namespace server that is running a version of one of the following operating systems: Windows Server 2003 R2 or Windows Server 2003 with SP1.

**Supported namespace types:** Domain-based; stand-alone.

**Supported namespace server operating systems:** Windows Server 2003 with SP1.

**Supported clients:** Windows XP with SP2 and the Windows XP client failback hotfix; Windows Server 2003 with SP1 and the Windows Server 2003 client failback hotfix. These hotfixes are described in article 898900 in the Microsoft Knowledge Base on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=53202).

**Where configured:** Namespace properties, Referrals tab; Folder properties, Referrals tab.

Client failover in DFS Namespaces is the process in which a client attempts to access another target in a referral after one of the targets becomes unavailable due to a server or network failure, or due to the target being removed from the namespace. After failover, the client might access a non-optimal target, such as a target with a high site link cost with respect to the client, and the client will continue to access the remote server even after the local server is restored. If you want clients to fail back to a preferred local server when it is restored, select the **Clients fail back to preferred targets**check box.

| **Note** |
| --- |
| When a preferred local server is available, all new file accesses will be to the local server. Files that are currently open will be accessed from the non-preferred server until that file is closed. |

To enable client failback, see Enable client failback.

**See Also**
**Other Resources**
Microsoft Web site

# Polling properties

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**When this option is available:** Available when the DFS Management snap-in connects to a namespace server that is running a version of Windows Server 2003 R2 or Windows Server 2003 with Service Pack 1 (SP1).

**Supported namespace types:** Domain-based only.

**Where configured:** Namespace properties, Advanced tab.

To maintain a consistent domain-based namespace across namespace servers, it is necessary for namespace servers to periodically poll Active Directory to obtain the most current namespace data.

You can optimize how this polling occurs by choosing one of the following options:

**Optimize for consistency**
**Supported namespace server operating systems:** Windows Server 2003 and Windows 2000 Server.

In this mode, which is enabled by default, the Distributed File System service attempts to keep the domain-based namespace highly consistent on all namespace servers. This consistency is achieved as follows:

- Namespace servers running Windows Server 2003 and Windows 2000 Server poll the primary domain controller (PDC) emulator every hour by default.

- Namespace servers running Windows Server 2003 poll the PDC emulator after the namespace changes to obtain the updated namespace data without needing to wait for Active Directory replication to replicate the updated namespace object to all domain controllers. However, this process can cause an increased load on the PDC emulator in namespaces that have more than 16 namespace servers and in namespaces that change frequently.

This mode is recommended only for domain-based namespaces that are hosted by fewer than 16 namespace servers. For namespaces hosted by more than 16 namespace servers, we recommend choosing the **Optimize for scalability** option.

**Optimize for scalability**
**Supported namespace server operating systems:** Windows Server 2003.

This mode, also known as root scalability mode, allows organizations to use more than the recommended 16 namespace servers for hosting a domain-based namespace in consistency mode. When root scalability mode is enabled, namespace servers do not send change notification messages to other namespaces servers when the namespace changes, nor do they poll the PDC emulator every hour. Instead, they poll their closest domain controller every hour to discover updates to the namespace.

Root scalability mode reduces network traffic to the PDC emulator at the expense of faster updates to all namespaces servers. Updates are still made to the namespace object in Active Directory on the PDC emulator, but namespace servers do not discover those changes until the updated namespace object replicates (using Active Directory replication) to the closest domain controller for each namespace server. After root scalability mode is enabled, clients might have an inconsistent view of the namespace until the most up-to-date version of the namespace object is replicated to all domain controllers and the namespaces servers poll a domain controller to discover the changes.

| **Note** |
| --- |
| Namespace servers running Windows 2000 Server continue to poll the PDC emulator every hour by default, even if this option is selected. |

To change the polling method, see Optimize namespace polling.

**See Also**
**Other Resources**
Microsoft Web site

# Target priority

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**When this option is available:** Availablewhen the DFS Management snap-in connects to a namespace server that is running a version of Windows Server 2003 with Service Pack 1 (SP1) or Windows Server 2003 R2.

**Supported namespace types:** Domain-based (folder targets and root targets); stand-alone (folder targets only).

**Supported namespace server operating systems:** Windows Server 2003 with SP1 and Windows Server 2003 R2.

**Supported clients:** All supported namespace clients.

**Where configured:** Namespace server properties, Advanced tab; Folder Target properties, Advanced tab.

By assigning target priority for individual root targets or folder targets, you can specify whether a target server appears first or last in a referral. Assigning target priority is useful in many scenarios, such as "hot standby" scenarios where one server is considered the server of last resort. In this scenario, you can specify that the standby server always appears last in referrals, and clients will fail over to this server only if all the other servers fail or become unavailable due to network outages. If multiple servers are configured for always first (or last), then those servers will be sorted based on the overall prioritization (random or lowest cost).

To change target priority, see Set target priority to override referral ordering.

# Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- DFS Replication Overview
- Getting Started With DFS Replication
- Configuring and Managing DFS Replication
- DFS Replication Properties

# DFS Replication Overview

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- Introduction to DFS Replication
- Replication groups and replicated folders
- Security requirements for setting up and managing DFS Replication
- DFS Replication requirements
- DFS Replication scalability guidelines
- What to expect during initial replication

# Introduction to DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

DFS Replication, the successor to the File Replication service (FRS) introduced in Windows 2000 Server operating systems, is a new, state-based, multimaster replication engine that supports replication scheduling and bandwidth throttling. DFS Replication uses a new compression algorithm known as remote differential compression (RDC). RDC is a "diff-over-the wire" client-server protocol that can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, and re-arrangements of data in files, enabling DFS Replication to replicate only the changed file blocks when files are updated.

| **Note** |
| --- |
| For more current information about DFS Replication, see DFS Namespaces and DFS Replication Overview. |

DFS Replication uses many sophisticated processes to keep data synchronized on multiple servers. Before you begin using DFS Replication, it is helpful to understand the following concepts.

- DFS Replication is a multimaster replication engine. Any change that occurs on one member is replicated to all other members of the replication group.
- DFS Replication detects changes on the volume by monitoring the update sequence number (USN) journal, and DFS Replication replicates changes only after the file is closed.
- DFS Replication uses a staging folder to stage a file before sending or receiving it. For more information about staging folders, see Staging folders and Conflict and Deleted folders.
- DFS Replication uses a version vector exchange protocol to determine which files need to be synchronized. The protocol sends less than 1 kilobyte (KB) per file across the network to synchronize the metadata associated with changed files on the sending and receiving members.
- When a file is changed, only the changed blocks are replicated, not the entire file. The RDC protocol determines the changed file blocks. Using default settings, RDC works for any type of file larger than 64 KB, transferring only a fraction of the file over the network.
- DFS Replication uses a conflict resolution heuristic of last writer wins for files that are in conflict (that is, a file that is updated at multiple servers simultaneously) and earliest creator wins for name conflicts. Files and folders that lose the conflict resolution are moved to a folder known as the Conflict and Deleted folder. You can also configure the service to move deleted files to the Conflict and Deleted folder for retrieval should the file or folder be deleted. For more information, see Staging folders and Conflict and Deleted folders.
- DFS Replication is self-healing and can automatically recover from USN journal wraps, USN journal loss, or loss of the DFS Replication database.
- DFS Replication uses a Windows Management Instrumentation (WMI) provider that provides interfaces to obtain configuration and monitoring information from the DFS Replication service.

**See Also**
**Other Resources**
Microsoft Web site

# Replication groups and replicated folders

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Replication groups and replicated folders are two important components of DFS Replication. They are illustrated in the following figure.




As this figure shows, a *replication group* is a set of servers, known as *members*, that participates in the replication of one or more replicated folders. A *replicated folder* is a folder that is kept synchronized on each member. In the figure, there are two replicated folders, Projects and Proposals. As data changes in each replicated folder, the changes are replicated across connections between the members of the replication group. The connections between all members form the replication topology.

Creating multiple replicated folders in a single replication group simplifies the process of deploying replicated folders because the topology, schedule, and bandwidth throttling for the replication group are applied to each replicated folder. To deploy additional replicated folders, you can use Dfsradmin.exe or a brief wizard to define the local path and permissions for the new replicated folder. Each replicated folder also has its own settings, such as file and subfolder filters, so that you can filter out different files and subfolders for each replicated folder.

The replicated folders stored on each member can be located on different volumes in the member, and the replicated folders do not need to be shared folders or part of a namespace, though the DFS Management snap-in makes it easy to share replicated folders and optionally publish them in an existing namespace.

**See Also**
**Other Resources**
Microsoft Web site

# Security requirements for setting up and managing DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The following table describes the groups that can perform basic DFS Replication tasks by default and the method for delegating the ability to perform these tasks.

| Task | Users or Groups That Can Perform This Task By Default | Delegation Method |
|---|---|---|
| Create a replication group or enable DFS Replication on a folder that has folder targets | Domain Admins group in the domain where the replication group will be created. | Right-click the **Replication** node in the console tree, and then click **Delegate Management Permissions**. |
| Administer a replication group | Domain Admins group in the domain where the replication group is configured, or the creator of the replication group. | Right-click the replication group in the console tree, and then click **Delegate Management Permissions**. |
| Add a server to a replication group[1, 2] | If the server is a member server, the user must be a member of the local Administrators group of the server to add.<br>If the server is a domain controller, the user must be a member of the Domain Admins group in the domain where the server is located. | Add the user to local Administrators group of the member server to add, or add the user to the Domain Admins group of the domain controller to add. |

[1]Assumes that the user has been delegated the ability to administer the replication group.

[2]The server to be added must be online.

If you plan to delegate the ability to create and administer replication groups, note the following two important considerations:

- If you delegate to a user or group the ability to create replication groups, and you later remove the user or group from the delegation list, there is no change to the security settings on existing replication groups.

- If you delegate to a user or group the ability to administer a specific replication group, and you later remove the user or group from the delegation list, there is no change to the security settings on any existing configuration data. For example, if the user who is being removed had created a connection in the replication group, then the user would still have permissions to edit that connection because the user is the owner of the Active Directory object that contains the configuration information for the connection.

**See Also**
**Other Resources**
Microsoft Web site

# DFS Replication requirements

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Before you can deploy DFS Replication, you must configure servers as follows:

- The Active Directory schema must be updated to include the new DFS Replication objects.

  To upgrade the schema, on the schema operations master, run **adprep.exe /forestprep**. The Adprep.exe command-line tool is available in the Cmpnents\R2\Adprep folder on the second Windows Server 2003 R2 installation CD.

- The servers that will participate in DFS Replication must run a Windows Server 2003 R2 operating system. After you install Windows Server 2003 R2, you must install the DFS Replication service on each server that will take part in replication, and you must install the DFS Management snap-in on one server to manage replication. For installation instructions, see Installing the DFS Management Snap-in and DFS Replication Service.

- Antivirus software must be compatible with DFS Replication; contact your antivirus software vendor to check for compatibility.

- Servers in a replication group must be in the same forest. You cannot enable replication across servers in different forests.

- Replicated folders must be stored on NTFS volumes.

- On server clusters, replicated folders must be located in the local storage of a node because the DFS Replication service is not designed to work in a coordinated way with cluster components, and the service will not fail over to another node.

**See Also**
**Other Resources**
Microsoft Web site

# DFS Replication scalability guidelines

Updated: May 5, 2010

Applies To: Windows Server 2003 R2

DFS Replication is a highly scalable replication engine that was built to efficiently replicate large sets of content over wide area networks. The following list provides a set of scalability guidelines that have been tested by Microsoft for Windows Server 2003 R2.

| **Note** |
| --- |
| For Windows Server 2008 R2 and Windows Server 2008, there is no longer a limit to the number of replication groups, replicated folders, connections, or replication group members. For more information, see the What are the supported limits of DFS Replication? topic in the DFS Replication: Frequently Asked Questions (FAQ). |

- Each server can be a member of up to 256 replication groups.
- Each replication group can contain up to 256 replicated folders.
- Each server can have up to 256 connections (for example, 128 incoming connections and 128 outgoing connections).
- On each server, the number of replication groups multiplied by the number of replicated folders multiplied by the number of simultaneously active connections must be kept to 1024 or fewer.
- A replication group can contain up to 256 members.
- A volume can contain up to 8 million replicated files, and a server can contain up to 1 terabyte of replicated files.
- The maximum tested file size is 64 gigabytes.

| **Note** |
| --- |
| You can successfully deploy configurations that exceed any or all of these tested limits. However, it is important to test large configurations and verify that there is adequate space in the staging folders before using them in production environments. In addition, you might experience increased latency during replication. See the following Web pages for additional information: Understanding DFS Replication "Limits" (http://go.microsoft.com/fwlink/?LinkId=70576); More on DFS Replication Limits (http://go.microsoft.com/fwlink/?LinkId=70575). Also see Staging folders and Conflict and Deleted folders. |

**See Also**
**Other Resources**
Microsoft Web site

# What to expect during initial replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When you first set up replication, you must choose a primary member. Choose the member that has the most up-to-date files that you want replicated to all other members of the replication group, because the primary member's content is considered "authoritative." This means that during initial replication, the primary member's files will always win the conflict resolution that occurs when the receiving members have files that are older or newer compared to the same files on the primary member.

The following concepts will help you better understand the initial replication process:

- Initial replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

- Initial replication always occurs between the primary member and the receiving replication partners of the primary member. After a member has received all files from the primary member, then that member will replicate files to its receiving partners as well. In this way, replication for a new replicated folder starts from the primary member and then progresses out to the other members of the replication group.

- When receiving files from the primary member during initial replication, the receiving members that contain files that are not present on the primary member move those files to their respective DfsrPrivate\PreExisting folder. If a file is identical to a file on the primary member, the file is not replicated. If the version of a file on the receiving member is different from the primary member’s version, the receiving member's version is moved to the Conflict and Deleted folder and remote differential compression (RDC) can be used to download only the changed blocks.

- To determine whether files are identical on the primary member and receiving member, DFS Replication compares the files using a hash algorithm. If the files are identical, only minimal metadata is transferred.

- After the initialization of the replicated folder, the "primary member" designation is removed. (Initialization takes place after all files that exist before DFS Replication picks up the configuration are added to the DFS Replication database.) That member is then treated like any other member and its files are no longer considered authoritative over other members that have completed initial replication. Any member that has completed initial replication is considered authoritative over members that have not completed initial replication.

**See Also**
**Other Resources**
Microsoft Web site

# Getting Started With DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- Publish data using DFS Replication
- Collect data for backup purposes using DFS Replication

# Publish data using DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

To publish data using DFS Replication, create a replication group and choose a hub-and-spoke topology with either one or two hub servers (for redundancy). In this topology, you can add data to the hub server, and the data replicates out to the spoke servers. Use the following procedure to create a replication group:

- Create a replication group

To customize the replication group, use the following procedures as appropriate:

- Edit the replication group schedule
- Create or edit a custom connection schedule
- Enable or disable replicating a replicated folder on a specific member
- Enable or disable RDC for a specific connection
- Edit the replication filters for a replicated folder
- Edit the quota size of the staging folder and Conflict and Deleted folder

After you configure the replication group, you can create a namespace to enable clients to fail over to the hub server if the branch server fails or otherwise becomes unavailable. You can also configure the namespace so that clients fail back to the branch server after it is restored.

To configure a namespace for publication, use the following procedures in the order specified:

1. Publish a replicated folder in an existing namespace
2. Set the ordering method for targets in referrals
3. Enable client failback

# Collect data for backup purposes using DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

To collect data for backup or other purposes, use the New Replication Group Wizard to set up a replication group for data collection. This option sets up a two-way replication between two servers, such as a branch server and a hub server. Administrators at the hub office can use backup software to back up the branch server's data from a hub server, eliminating the often error-prone process of having end users performing the backups at branch offices that are not staffed by trained information technology (IT) personnel.

| **Important** |
| --- |
| The data collection topology allows data to replicate from the hub server to the branch server. If you do not want changes made at the hub server to replicate to the branch server, set up shared folder permissions to prevent data from changing on the hub server, or disable the hub-to-branch connection after you set up the replication group. |

| **Important** |
| --- |
| Using DFS Replication for data collection is not a replacement for performing regular backups. |

Use the following procedure to create a replication group for data collection purposes:

- Create a replication group

To customize the replication group, use the following procedures as appropriate:

- Edit the replication group schedule
- Enable or disable replication for a specific connection
- Edit the replication filters for a replicated folder
- Edit the quota size of the staging folder and Conflict and Deleted folder

After you configure the replication group, you can create a namespace to simplify the way users access the published data and increase the availability of the data.

To configure a namespace for data collection, perform the following procedures in the order specified:

1. Publish a replicated folder in an existing namespace
2. Set the ordering method for targets in referrals
3. Set target priority to override referral ordering (Optional)
4. Enable client failback (Optional)

# Configuring and Managing DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- Deploying DFS Replication
- Managing DFS Replication

# Deploying DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**How to...**

- Create a replication group
- Add a replicated folder to a replication group
- Add a member to a replication group
- Create a connection
- Create a topology
- Delegate the ability to create replication groups
- Delegate the ability to manage an existing replication group

# Create a replication group

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

### To create a replication group

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, right-click the **Replication** node, and then click **New Replication Group**.
3. Follow the instructions in the New Replication Group Wizard.

### Important

- When you create a new replication group, replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and each member's long polling interval (60 minutes).

### Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority to create replication groups and be a member of the local Administrators group on each member.

- To check the status of replication, or to troubleshoot replication, create a diagnostic report. For more information, see Create a diagnostic report for DFS Replication.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Introduction to DFS Replication
DFS Replication requirements
Replication groups and replicated folders
Replication topologies
Replication schedules and bandwidth
What to expect during initial replication
**Other Resources**
Microsoft Web site

# Add a replicated folder to a replication group

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

### To create a replicated folder

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, right-click a replication group, and then click **New Replicated Folders**.
3. Follow the instructions in the New Replicated Folders Wizard.

### Important

- Replication of the new replicated folder does not begin immediately. The new DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

### Notes

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority for the replication group and be a member of the local Administrators group on each member. If you are the creator of the replication group, you can also perform this procedure if you are a member of the local Administrators group on each member.

- To check the status of replication, or to troubleshoot replication, create a diagnostic report. For more information, see Create a diagnostic report for DFS Replication.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
Replication groups and replicated folders
**Other Resources**
Microsoft Web site

# Add a member to a replication group

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add a member to a replication group**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, right-click a replication group, and then click **New Member**.
3. Follow the instructions in the New Member Wizard.

**Important**

- Replication does not begin immediately on the new member. The DFS Replication settings for the new member must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency, the short polling interval (5 minutes) on the new member, and the long polling interval (60 minutes) on existing members.

**Notes**

- Security requirements for performing this task are described in Security requirements for setting up and managing DFS Replication.
- To check the status of replication, or to troubleshoot replication, create a diagnostic report. For more information, see Create a diagnostic report for DFS Replication.

**See Also**
**Other Resources**
Microsoft Web site

# Create a connection

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To create a connection**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, right-click the replication group that you want to create a new connection in, and then click **New Connection**.
3. Select the sending and receiving members, and select the schedule to use for the connection.
4. To create a second connection for two-way replication between the sending and receiving members, select **Create a second connection in the opposite direction**.

**Important**

- Replication does not begin immediately on the new connection. The new connection settings must be replicated to all domain controllers, and the sending and receiving members must poll their closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
**Other Resources**
Microsoft Web site

# Create a topology

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

### To create a topology

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, right-click the replication group to create a new topology for, and then click **New Topology**.
3. Follow the instructions in the New Topology Wizard.

### Important

- Replication in the new topology does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

### Note

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
Replication topologies
**Other Resources**
Microsoft Web site

# Delegate the ability to create replication groups

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To delegate the ability to create replication groups**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, right-click the **Replication** node, and then click **Delegate Management Permissions**.
3. To add a user or group to the delegation list, click **Add**.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the appropriate authority.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
**Other Resources**
Microsoft Web site

# Delegate the ability to manage an existing replication group

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To delegate the ability to manage an existing replication group**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, right-click the replication group, and then click **Delegate Management Permissions**.
3. Type the name of the user or group that you want to delegate management permissions to, and then click **OK**.

**Notes**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

- To view the delegation list for a replication group, select the replication group in the console tree, and then click the **Delegation** tab in the details pane.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
**Other Resources**
Microsoft Web site

# Managing DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**How to...**

- Edit the replication group schedule
- Create or edit a custom connection schedule
- Enable or disable replication for a specific connection
- Enable or disable replicating a replicated folder on a specific member
- Enable or disable RDC for a specific connection
- Enable or disable topology verification
- Repair a disconnected topology
- Edit the replication filters for a replicated folder
- Edit the quota size of the staging folder and Conflict and Deleted folder
- Share a replicated folder
- Publish a replicated folder in an existing namespace
- Create a diagnostic report for DFS Replication

# Edit the replication group schedule

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To edit the replication group schedule**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, right-click the replication group whose schedule you want to edit, and then click **Edit Replication Group Schedule**.

**Important**

- Schedule changes are not applied immediately. The new schedule must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain the changes. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
Replication schedules and bandwidth
**Other Resources**
Microsoft Web site

# Create or edit a custom connection schedule

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To create or edit a custom connection schedule**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, click the replication group that contains the connection you want to edit.
3. In the details pane, click the **Connections** tab.
4. Right-click the connection, click **Properties**, and then click the **Schedule** tab.
5. To specify a custom schedule, select **Custom connection schedule**, and then click **Edit Schedule** to edit the schedule that applies only to that connection.

**Important**

- Schedule changes are not applied immediately. The new schedule must be replicated to all domain controllers, and the sending and receiving members for the connection must poll their closest domain controller to obtain the changes. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
Replication schedules and bandwidth
**Other Resources**
Microsoft Web site

# Enable or disable replication for a specific connection

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

### To enable or disable replication for a specific connection

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, click the replication group that contains the connection you want to edit.
3. In the details pane, click the **Connections** tab.
4. Do one of the following:
   - To disable a connection, right-click the connection, and then click **Disable**.
   - To enable a connection, right-click the connection, and then click **Enable**.

### Important

- Connection changes are not applied immediately. The new connection settings must be replicated to all domain controllers, and the sending and receiving members must poll their closest domain controller to obtain the changes. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

### Note

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
**Other Resources**
Microsoft Web site

# Enable or disable replicating a replicated folder on a specific member

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To enable or disable replicating a replicated folder to a specific member**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, click the replication group that contains the membership you want to enable or disable.
3. In the details pane, click the **Memberships** tab.
4. Select the replicated folder on the specific member, and then do one of the following:
   - To disable the replicated folder on the member, right-click the replicated folder, and then click **Disable**.
   - To enable the replicated folder on the member, right-click the replicated folder, and then click **Enable**.

**Important**

- Before you use this procedure, it is important to understand the ramifications of enabling and disabling memberships. After a disabled membership is enabled, the member must complete an initial replication of the replicated folder. Initial replication will cause about 1 KB of data to be transferred for each file or folder in the replicated folder, and any updated or new files present on the member will be moved to the DfsrPrivate\PreExisting folder on the member and be replaced with authoritative files from another member.
- If all memberships are disabled for any replicated folder, then the first member that is enabled will automatically be made the primary member. If the discretionary access control list (DACL) on the local path for this member contains any inherited access-control entries (ACE), these ACEs will be converted to explicit ACEs with the same permissions so that all members will end up with the same DACL on the replicated folder root.
- Membership changes are not applied immediately. The membership changes must be replicated to all domain controllers, and the member must poll its closest domain controller to obtain the changes. The amount of time this takes depends on Active Directory replication latency and the short polling interval (5 minutes) on the member.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
What to expect during initial replication
**Other Resources**
Microsoft Web site

# Enable or disable RDC for a specific connection

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Remote differential compression (RDC) is a "diff-over-the-wire" protocol that can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, re-arrangements of data in files, enabling DFS Replication to replicate only the changes when files are updated.

Because disabling RDC can help conserve disk input/output (I/O) and CPU resources, you might want to disable RDC on a connection if the sending and receiving members are in a local area network (LAN), and bandwidth use is not a concern. However, in a LAN environment where bandwidth is contended, RDC can be beneficial when transferring large files.

**To enable or disable RDC for a specific connection**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, click the replication group that contains the connection you want to edit.
3. In the details pane, click the **Connections** tab.
4. Right-click the connection, click **Properties**, and then on the **General** tab, select or clear the **Use remote differential compression (RDC)** check box.

**Important**

- Changes to RDC settings are not applied immediately. The RDC changes must be replicated to all domain controllers, and the sending and receiving members must poll their closest domain controller to obtain the changes. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
**Other Resources**
Microsoft Web site

# Enable or disable topology verification

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Topology verification checks whether the topology is fully connected. You can view the status of the topology by selecting a replication group, clicking the **Connections** tab in the details pane, and viewing the status that appears in the details pane.

**To enable or disable topology validation**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, right-click the **Replication** node, and then do one of the following:
    - To disable topology verification, click **Disable Topology Verification**.
    - To enable topology verification, click **Enable Topology Verification**.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
Repair a disconnected topology
**Other Resources**
Microsoft Web site

# Repair a disconnected topology

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To repair a disconnected topology**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, click the replication group that contains the disconnected topology.
3. In the details pane, on the **Connections** tab, click the error message at the top of the tab to open the Disconnected Topology dialog box.
4. To view disconnected members, select a replicated folder from the **Replicated folder** list, and then expand the **Disconnected Members**.
5. Click **Close**.
6. Repair the topology by performing any of the following procedures, as appropriate:
   - Create a connection
   - Edit the replication group schedule
   - Create or edit a custom connection schedule
   - Enable or disable replication for a specific connection
   - Enable or disable replicating a replicated folder on a specific member

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
**Other Resources**
Microsoft Web site

# Edit the replication filters for a replicated folder

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To edit the replication filters for a replicated folder**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, click the replication group that contains the replicated folder whose filters you want to edit.
3. In the details pane, on the **Replicated Folders** tab, right-click a replicated folder, and then click **Properties**.
4. On the **General** tab, edit the existing filters or add new filters.

**Important**

- Filter changes are not applied immediately. The updated filters must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain the changes. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
File and subfolder filters
**Other Resources**
Microsoft Web site

# Edit the quota size of the staging folder and Conflict and Deleted folder

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can edit the quota of the staging folder and Conflict and Deleted folder on a per-replicated folder, per-member basis.

**To edit the quota size of the staging folder and Conflict and Deleted folder**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, click the replication group that contains the replicated folder whose quotas you want to edit.
3. In the details pane, on the **Memberships** tab, right-click the replicated folder on the member whose quota you want to edit, and then click **Properties**.
4. On the **Advanced** tab, adjust the quotas as necessary.

**Important**

- Changes to staging folder quotas and Conflict and Deleted folder quotas are not applied immediately. The updated quotas must be replicated to all domain controllers, and the member must poll its closest domain controller to obtain the changes. The amount of time this takes depends on Active Directory replication latency and the short polling interval (5 minutes) on the member.

**Note**

- To perform this procedure, you must be a member of the Domain Admins group in Active Directory, or you must have been delegated the ability to manage an existing replication group. If you are the creator of the replication group, you can also perform this procedure.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Delegate the ability to manage an existing replication group
Staging folders and Conflict and Deleted folders
**Other Resources**
Microsoft Web site

# Share a replicated folder

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To share a replicated folder**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, click the replication group that contains the replicated folder you want to share.
3. In the details pane, on the **Replicated Folders** tab, right-click the replicated folder you want to share, and then click **Share and Publish in Namespace**.
4. In the Share and Publish Replicated Folder Wizard, select the **Share the replicated folder** option, and then follow the steps in the wizard.

**Note**

- To perform this procedure, you must be a member of the local Administrators group on each server where the replicated folder is shared.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
**Other Resources**
Microsoft Web site

# Publish a replicated folder in an existing namespace

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To publish a replicated folder**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, click the replication group that contains the replicated folder you want to share.
3. In the details pane, on the **Replicated Folders** tab, right-click the replicated folder you want to share, and then click **Share and Publish in Namespace**.
4. In the Share and Publish Replicated Folder Wizard, select the **Share and publish the replicated folder in a namespace** option, and then follow the steps in the wizard.

**Notes**

- To perform this procedure, you must meet the security requirements for managing DFS Replication and namespaces. If folders will be shared as part of the Share and Publish Replicated Folder Wizard, you must also be a member of the local Administrators group on the servers where each folder is shared.

- If you do not have an existing namespace, you can create one in the **Namespace Path** page in the Share and Publish Replicated Folder Wizard. To create the namespace, click **Browse** in the **Namespace Path** page, and then click **New Namespace**.

**See Also**
**Concepts**
Security requirements for setting up and managing DFS Replication
Security requirements for creating and managing namespaces
Introduction to namespaces
**Other Resources**
Microsoft Web site

# Create a diagnostic report for DFS Replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To create a diagnostic report for DFS Replication**

1. Click **Start**, point to **Administrative Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, right-click the replication group that you want to create a diagnostic report for, and then click **Create Diagnostic Report**.
3. Follow the instructions in the Diagnostic Report Wizard.

**Notes**

- To create a diagnostic report, you must be a member of the local Administrators group on each server that you prepare a report for.

- The amount of time necessary to generate a diagnostic report will vary based on a number of factors, including: DFS Replication health; the number of replicated folders; available server resources (for example, CPU, memory, and so forth); WAN availability (connectivity, bandwidth, and latency); and the chosen reporting options (backlog and file system enumeration). Because of the potential delay in creating diagnostic reports, we recommend creating diagnostic reports for no more than 50 servers at a time.

**See Also**
**Other Resources**
Microsoft Web site

# DFS Replication Properties

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section covers:

- Replication schedules and bandwidth
- Replication topologies
- Staging folders and Conflict and Deleted folders
- File and subfolder filters

# Replication schedules and bandwidth

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The following sections describe options for configuring replication schedules and bandwidth usage in DFS Replication.

**Replication schedules**
The replication schedule determines the days and times at which replication occurs. You can configure the schedule at 15-minute intervals 24 hours a day, seven days a week.

There are two types of schedules:

- **Replication group schedule**. This schedule applies to all connections in the replication group except connections that have a custom connection schedule. You can edit the replication group schedule by viewing the properties of the replication group.

- **Custom connection schedule**. This is a unique schedule that is applied to an individual connection. You can edit the custom schedule by viewing the properties of a connection.

Because members can exist in different time zones, it is important to understand how the schedule is affected by time zones, and whether daylight savings time is in effect. Replication is always initiated by the receiving member; therefore, the schedule reflects the time at which a receiving member initiates replication with a sending member. The receiving members can interpret the schedule in one of the following ways:

- **Universal Coordinated Time (UTC)**. This option causes the receiving member to treat the schedule as an absolute clock. For example, a schedule that begins at 0800 UTC is the same for any location, regardless of time zone or whether daylight savings time is in effect for a receiving member. For example, assume that you set replication to begin at 0800 UTC. A receiving member in Eastern Standard Time would begin replicating at 3:00 A.M. local time (UTC - 5), and a receiving member in Rome would begin replicating at 9:00 A.M. local time (UTC + 1). Note that the UTC offset shifts when daylight savings time is in effect for a particular location.

- **Local time of receiving member**. This option causes the receiving member to use its local time to start and stop replication. Local time is determined by the time zone and daylight savings time status of the receiving member. For example, a schedule that begins at 8:00 A.M. will cause every receiving member to begin replicating when the local time is 8:00 A.M. Note that daylight savings time does not cause the schedule to shift. If replication starts at 9 A.M. before daylight savings time, replication will still start at 9 A.M. when daylight savings time is in effect.

When the schedule is open, replication occurs as files are changed, created, or deleted. When the schedule closes, replication stops, regardless of whether all changed or new files have replicated.

| **Important** |
| --- |
| Schedule changes are not applied immediately. The new schedule must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain the changes. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member. |

**Bandwidth usage**
You can configure a different bandwidth usage for each interval of the schedule. The bandwidth setting is an absolute value, not a percentage of available bandwidth. The bandwidth used is not kept at the specified bandwidth setting for every one-second interval. Instead, the bandwidth used can occur in bursts at any given second, but over a period of time, the average bandwidth used is equal to the bandwidth setting.

**See Also**
**Other Resources**
Microsoft Web site

# Replication topologies

Updated: September 2, 2009

Applies To: Windows Server 2003 R2

The replication topology consists of the logical connections that DFS Replication uses to replicate files among servers. When you set up a replication group, you can choose from three topologies:

- **Hub and spoke**. This topology requires three or more members; otherwise this option is unavailable. For each spoke member, you can choose a required hub member and an optional second hub member for redundancy. This optional hub ensures that a spoke member can still replicate if one of the hub members is unavailable. If you specify two hub members, the hub members will have a full-mesh topology between them.

- **Full mesh**. In this topology, every member replicates with all other members of the replication group. This topology works well when there are ten or fewer members in the replication group. We recommend against using a full mesh topology if you have more than ten members in the replication group.

- **No topology**. Choose this option if you want to create connections yourself after you finish the New Replication Group Wizard or the Replicate Folder Wizard. No replication will take place until you create the connections.

When choosing a topology, keep in mind that two one-way connections are created between the members you choose. These two connections allow data to flow in both directions. For example, in a hub and spoke topology, data will flow from the hub members to the spoke members and from the spoke members to the hub members.

**See Also**
**Other Resources**
Microsoft Web site

# Staging folders and Conflict and Deleted folders

Updated: February 25, 2010

Applies To: Windows Server 2003 R2

**Staging folders**

DFS Replication uses staging folders to act as caches for new and changed files to be replicated from sending members to receiving members. The sending member begins staging a file when it receives a request from the receiving member. The process involves reading the file from the replicated folder and building a compressed representation of the file in the staging folder. This is the *staged file*. After being constructed, the staged file is sent to the receiving member; if remote differential compression [RDC] is used, only a fraction of the staging file might be replicated. The receiving member downloads the data and builds the file in its staging folder. After the file has completed downloading on the receiving member, DFS Replication decompresses the file and installs it into the replicated folder.

Each replicated folder has its own staging folder, which by default is located under the local path of the replicated folder in the DfsrPrivate\Staging folder. The default size of each staging folder is 4,096 MB. This is not a hard limit, however. It is only a quota that is used to govern cleanup and excessive usage based on high and low watermarks (90 percent and 60 percent of staging folder size, respectively). For example, when the staging folder reaches 90 percent of the configured quota the oldest staged files are purged until the staging folder reaches 60 percent of the configured quota.

It is important to note that the staging folder quota does not determine the largest file that can be replicated. In other words, it is possible to replicate a file that is larger than the configured quota of a staging folder. The large file is placed in the staging folder, and the staging folder cleanup process is triggered when the file is finished staging and space usage is at or above the high watermark. If the cleanup fails because the large file is still in the process of being replicated to receiving members, the cleanup process will be retried later and eventually the large file will be purged from the staging folder.

Although you can adjust the size of each staging folder, you must take the following factors into account while doing so:

- For good operational performance, increasing the quota size of a staging folder is recommended when you have multiple large files that change frequently. We also recommend that you increase the staging folder quota on hub members that have many replication partners.

- The size of each staging folder on a member is cumulative per volume. For example, if you have three replicated folders on a member on the same volume, it is possible for DFS Replication to use 12 GB or more for staging purposes. However, staging space is not preallocated; disk space is only used when staged files are present.

  | **Note** |
  | --- |
  | In this example, staging cleanup might not be triggered when the staging files occupy 10.8 GB (0.9*4 GB*3), because staging space is distributed across the three replicated folders. If free disk space is a concern, you might need to configure the staging quota to be lower than the default quota when several replicated folders share staging space on the same volume. This ensures that staging cleanup is triggered. |

- If the size of a staging folder is below 90 percent of configured capacity (the high watermark), then staged files are kept in the folder and can be used in case new members are added.

- Staged files are purged (based on a "least recently used" algorithm) when a staging folder reaches 90 percent of the configured quota (the high watermark). Files are purged until the staging size falls below the low watermark (default is 60 percent) of the configured quota.

- If a staging folder quota is configured to be too small, DFS Replication might consume additional CPU and disk resources to regenerate the staged files. Replication might also slow down because the lack of staging space can effectively limit the number of concurrent transfers with partners.

- For the initial replication of existing data on the primary member, it is important that you size the staging folder quota large enough so that if multiple large files are blocked in staging due to partners not being able to download the files, the remaining files can continue replicating. To properly size the staging folder for initial replication, you must take into account the size of the files to be replicated. At a minimum, the staging folder quota for servers running Windows Server 2003 R2 should be at least the size of the nine largest files in the replicated folder. To improve performance, set

the size of the staging folder quota as close as possible to the size of the replicated folder.

**Tip**

To determine the size of the largest files in a replicated folder, you can either use Windows Explorer (sort by size and add the nine largest file sizes to get the minimum staging folder size), or, from a Windows PowerShell™ command prompt, use this Windows PowerShell command where <*replicatedfolderpath*> is the path to the replicated folder:

```
Get-ChildItem <replicatedfolderpath> -recurse | Sort-Object length -descending  | select-object -first 9
| ft name,length -wrap -auto
```

To install Windows PowerShell on a server running Windows Server 2003, see Windows PowerShell 1.0 English-Language Installation Package for Windows Server 2003 (KB926139) in the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=184089).

- During normal operation, if the event that indicates the staging quota is over its configured size (event ID 4208 in the DFS Replication event log) is logged multiple times in a day, increase the staging quota by 20 percent.

**Conflict and Deleted folders**

DFS Replication uses a "last-writer wins" method for determining which version of a file to keep when a file is modified on two or more members. The losing file is stored in the Conflict and Deleted folder on the member that resolves the conflict. This member might not be the member where the changes originated.

Each replicated folder has its own Conflict and Deleted folder, which is located under the local path of the replicated folder in the DfsrPrivate\ConflictandDeleted folder. The quota size of the Conflict and Deleted folder is 660 MB by default. Like the staging folder, the Conflict and Deleted folder has high and low watermarks (90 percent and 60 percent of Conflict and Deleted folder quota, respectively) that govern cleanup and excessive usage of the folder. And like the staging folder, the size of each Conflict and Deleted folder on a member is cumulative per volume.

The files in the Conflict and Deleted folder are renamed and are accessible only by members of the local Administrators group. The access control lists (ACLs) on the conflict files are preserved to help ensure security. You can view a log of conflict files, as well as their original file names, by viewing the ConflictandDeletedManifest.xml file in the DfsrPrivate folder.

The Conflict and Deleted folder can also be used to store files that are deleted from replicated folders. You can configure this option by viewing the properties of a membership. Deleted files are treated like conflict files in that they are purged when the Conflict and Deleted folder reaches 90 percent of the configured quota. They are also renamed and listed in the ConflictandDeletedManifest.xml file.

**See Also**

**Other Resources**

Microsoft Web site

# File and subfolder filters

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can configure file and subfolder filters to prevent files and subfolders from replicating. Both types of filters are set on a per-replicated folder basis. You exclude subfolders by specifying their name or by using the asterisk (*) wildcard. You exclude files by specifying their name or by using the asterisk (*) wildcard to specify file names and extensions.

By default, no subfolders are excluded. The default file filters exclude the following files from replication:

- File names starting with a tilde (~) character
- Files with .bak or .tmp extensions

Regardless of which filters are set, the following types files are always excluded from replication:

- NTFS mounted drives.
- Files encrypted by using encrypting file system (EFS).
- Any reparse points except those associated with DFS Namespaces. If a file has a reparse point used for Hierarchical Storage Management (HSM) or Single Instance Store (SIS), DFS Replication replicates the underlying file but not the reparse point.
- Files on which the temporary attribute has been set.

When creating filters, note the following restrictions:

- You cannot create file or subfolder filters by specifying a full path, such as C:\Replicatedfolder\Temp or C:\Replicatedfolder\file.dat.
- You cannot specify a comma in a filter because commas are used as delimiters.

| **Important** |
| --- |
| Filter changes are not applied immediately. The updated filters must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain the changes. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member. |

It is also important to understand what happens when filters are created and removed.

- **When a filter is created**. After a member detects a new filter, the member scans its database and removes any file records of files that match the filter. Because the files are no longer listed in the database, any future changes to the files are ignored.
- **When a filter is removed**. After a member detects that a filter has been removed, the member scans the file system, adds records for all files that match the removed filter, and replicates the files.

**See Also**
**Other Resources**
Microsoft Web site

# Active Directory Federation Services (ADFS)

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) is based on the emerging, industry-supported Web Services Architecture, which is defined in WS-* specifications. ADFS helps you use single sign-on (SS0) to authenticate users to multiple, related Web applications over the life of a single online session. ADFS accomplishes this by securely sharing digital identity and entitlement rights across security and enterprise boundaries.

- For help with specific tasks, see ADFS How To.
- For general background information, see ADFS Concepts.
- For problem-solving instructions, see ADFS troubleshooting.

# ADFS How To

Updated: August 22, 2005

Applies To: Windows Server 2003 R2


- Deploy ADFS
- Complete the Add Resource Partner Wizard
- Complete the Add Account Partner Wizard
- Complete the Add Account Store Wizard
- Complete the Add Applications Wizard
- Configure the Account Store Priority
- Manage ADFS

# Deploy ADFS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

For information about how to deploy Active Directory Federation Services (ADFS), look for ADFS planning and deployment content on the Windows Server 2003 R2 Roadmap page (http://go.microsoft.com/fwlink/?LinkId=45560) of the Microsoft Windows Server 2003 TechCenter Web site.

# Complete the Add Resource Partner Wizard

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use the Add Resource Partner Wizard to add a new resource partner manually or by importing a policy file. Use this wizard to add a resource partner that will provide a Web application to users who have accounts in your account store (Active Directory or Active Directory Application Mode (ADAM)).

**Add a resource partner manually**
You can use the following procedure to add a resource partner manually.

**To add a resource partner manually**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. In the console tree, right-click **Resource Partners**, point to **New**, and then click **Resource Partner**.
4. On the **Welcome to the Add Resource Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, ensure that **No** is selected, and then click **Next**.
6. On the **Resource Partner Details** page, do the following, and then click **Next**:
   - In **Display name**, type the name of the resource partner.
   - In **Federation Service URI**, type the Uniform Resource Identifier (URI) for the resource partner.
   - In **Federation Service endpoint URL**, type the Uniform Resource Locator (URL) of the Federation Service.
7. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
8. On the **Resource Partner Identity Claims** page, select one or more identity claims to share with the resource partner, and then click **Next**:
   - If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

| **Note** |
| --- |
| If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable. This is because UPN claims are required for this scenario. |

   - If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.
   - If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.
9. If you selected **UPN Claim** as an identity claim, on the **Select UPN Suffix** page, do one of the following, and then click

**Next**:

- To pass all UPN suffixes through without replacing them, click **Pass all UPN suffixes through unchanged**.
- To replace all UPN suffixes with a different suffix, click **Replace all UPN suffixes with the following**, and then type the suffix that you want to use to replace all UPN suffixes.

10. If you selected **E-mail Claim** as an identity claim, on the **Select E-mail Suffix** page, do one of the following, and then click **Next**:
    - To pass all e-mail suffixes without replacing them, click **Pass all E-mail suffixes through unchanged**.
    - To replace all e-mail suffixes with a different suffix, click **Replace all E-mail suffixes with**, and then type the suffix that you want to use to replace all e-mail suffixes.

| **Note** |
| --- |
| Common name claims require no additional information. |

11. On the **Enable this Resource Partner** page, if you do not want to enable the resource partner now, clear the **Enable this resource partner** check box, and then click **Next**.
12. To add the new resource partner and close the wizard, click **Finish**.

**Add a resource partner by importing a policy file**

You can use the following procedure to add a resource partner by importing a policy file.

**To add a resource partner by importing a policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. Right-click **Resource Partners**, point to **New**, and then click **Resource Partner**.
4. On the **Welcome to the Add Resource Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, do the following, and then click **Next**:
    - Click **Yes**.
    - In **Partner interoperability policy file**, browse to or type the location of the resource partner policy file.
6. On the **Federation Scenario** page, do one of the following, and then click **Next**:
    - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**.
    - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
7. On the **Resource Partner Identity Claims** page, select one or more identity claims that the account partner will provide to the resource partner, and then click **Next**:
    - If the resource partner requires UPN claims to make authorization decisions, select the **UPN Claim** check box.

| **Note** |
| --- |

| |
|---|
| If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable. This is because UPN claims are required for this scenario. |

- If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.
- If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.

8. If you selected **UPN Claim** as an identity claim, on the **Select UPN Suffix** page, select one of the following, and then click **Next**.
   - To pass all UPN suffixes through without replacing them, click **Pass all UPN suffixes through unchanged**.
   - To replace all UPN suffixes with a different suffix, click **Replace all UPN domain suffixes with the following**, type the suffix that you want to use to replace all UPN suffixes, and then click **Add**.
9. If you selected **E-mail Claim** as an identity claim, on the **Select E-mail Suffix** page, do one of the following, and then click **Next**:
   - To pass all e-mail suffixes without replacing them, click **Pass all e-mail suffixes through unchanged**.
   - To replace all UPN suffixes with a different suffix, click **Replace all E-mail suffixes with**, and then type the suffix that you want to use to replace all e-mail suffixes.

| **Note** |
|---|
| Common name claims require no additional information. |

10. On the **Enable this Resource Partner** page, if you do not want to enable the resource partner now, clear the **Enable this resource partner** check box, and then click **Next**.
11. To add the new resource partner and close the wizard, click **Finish**.

**Related Information**

- Partner organizations
- Federation scenarios
- Federation trusts
- Federation Service
- Claims

# Complete the Add Account Partner Wizard

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use the Add Account Partner Wizard to add a new account partner manually or by importing a policy file. This will allow user accounts in the account partner to access Web applications that are protected by this Federation Service.

**Add an account partner manually**

You can use the following procedure to add an account partner manually.

**To add an account partner manually**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. Right-click **Account Partners**, point to **New**, and then click **Account Partner**.
4. On the **Welcome to the Add Account Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, click **No**, and then click **Next**.
6. On the **Account Partner Details** page, do the following, and then click **Next**
   - In **Display name**, type the display name of the account partner.
   - In **Federation Service URI**, type the Uniform Resource Identifier (URI) of the Federation Service.
   - In **Federation Service endpoint URL**, type the Uniform Resource Locator (URL) of the Federation Service.
7. On the **Account Partner Verification Certificate** page, type the path to the verification certificate, or browse to it, and then click **Next**.
8. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**, and then go to step 10.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust, click **Federated Web SSO with Forest Trust**.
9. On the **Federated Web SSO with Forest Trust** page, do one of the following, and then click **Next**:
   - To accept users in all domains that are trusted by the account partner, click **All Active Directory domains and forests**. Any user that can authenticate to the account partner will be accepted.
   - To accept user accounts that are located in some of the domains that are trusted by the account partner, click **The following Active Directory domains and forests**. Then, in **New, trusted Active Directory domain or forest**, type the name of a domain or forest, and then click **Add**. Only users from the specified domains will be accepted.
10. On the **Account Partner Identity Claims** page, select one or more identity claims to share with the resource partner, and then click **Next**:
    - If the resource partner requires user principal name (UPN) claims to make authorization decisions, select the **UPN Claim** check box.

**Important**

| |
|---|
| When UPN claims or e-mail claims are used to make authorization decisions, it is essential that each account partner use a unique UPN suffix or e-mail suffix. If two account partners use the same UPN suffix or e-mail suffix, it may not be possible to uniquely identify users. This condition might allow a user from one account partner to receive the permissions that are intended for a user in another account partner. This condition might also introduce a significant security weakness because an administrator could intentionally create user accounts to impersonate users from one of your other account partners. |

| **Note** |
|---|
| If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable. This is because UPN claims are required for this scenario. |

- If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.
- If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.

11. If you selected **UPN Claim** as an identity claim, on the **Accepted UPN Suffixes** page, click **All UPN suffixes** (this option is available only if you selected the **Federation Web SSO with Forest Trust** option), or type the accepted suffix, click **Add**, and then click **Next**.
12. If you selected **E-mail Claim** as an identity claim, on the **Accepted E-mail Suffixes** page, click **All E-mail suffixes** (this option is available only if you selected the **Federation Web SSO with Forest Trust** option), or type the accepted suffix, click **Add**, and then click **Next**.

| **Note** |
|---|
| Common name claims require no additional information. |

13. On the **Enable this Account Partner** page, if you do not want to enable the account partner now, clear the **Enable this account partner** check box, and then click **Next**.
14. To add the new account partner and close the wizard, click **Finish**.

**Add an account partner by importing a policy file**

You can use the following procedure to add an account partner by importing a policy file.

**To add an account partner by importing a policy file**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **Partner Organizations**.
3. Right-click **Account Partners**, point to **New**, and then click **Account Partner**.
4. On the **Welcome to the Add Account Partner Wizard** page, click **Next**.
5. On the **Import Policy File** page, do the following, and then click **Next**:
   - Click **Yes**.
   - In **Partner interoperability policy file**, browse to or type the location of the account partner policy file.
6. On the **Federation Scenario** page, do one of the following, and then click **Next**:
   - If you are establishing a federated trust with another organization or you do not want to use an existing forest trust, click **Federated Web SSO**, and then go to step 8.
   - If you are establishing a federated trust within the same organization when both sides already share a forest trust,

click **Federated Web SSO with Forest Trust**.

7. On the **Federation Web SSO with Forest Trust** page, do one of the following, and then click **Next**:

   - To accept users in all domains that are trusted by the account partner, click **All Active Directory domains and forests**. Any user that can authenticate to the account partner will be accepted.

   - To accept user accounts that are located in some of the domains that are trusted by the account partner, click **The following Active Directory domains and forests**. Then, in **New, trusted Active Directory domain or forest**, type the name of the domain or forest, and then click **Add**. Only users from the specified domains will be accepted.

8. On the **Account Partner Identity Claims** page, select one or more identity claims that this partner will provide, and then click **Next**:

   - If the resource partner requires UPN claims to make authorization decisions, select the **UPN Claim** check box.

| **Important** |
| --- |
| When UPN claims or e-mail claims are used to make authorization decisions, it is essential that each account partner use a unique UPN suffix or e-mail suffix. If two account partners use the same UPN suffix or e-mail suffix, it may not be possible to uniquely identify users. This condition might allow a user from one account partner to receive the permissions that are intended for a user in another account partner. This condition might also introduce a significant security weakness because an administrator could intentionally create user accounts to impersonate users from one of your other account partners. |

| **Note** |
| --- |
| If you selected the **Federated Web SSO with Forest Trust** scenario, the **UPN Claim** option is selected and not configurable. This is because UPN claims are required for this scenario. |

   - If the resource partner requires e-mail claims to make authorization decisions, select the **E-mail Claim** check box.

   - If the resource partner requires common name claims to make authorization decisions, select the **Common Name Claim** check box.

9. If you selected **UPN Claim** as an identity claim, on the **Accepted UPN Suffixes** page, click **All UPN suffixes** (this option is available only if you selected the **Federation Web SSO with Forest Trust** option), or type the accepted suffixes, click **Add**, and then click **Next**.

10. If you selected **E-mail Claim** as an identity claim, on the **Accepted E-mail Suffixes** page, click **All E-mail suffixes** (this option is available only if you selected the **Federation Web SSO with Forest Trust** option), or type the accepted suffix, click **Add**, and then click **Next**.

| **Note** |
| --- |
| Common name claims require no additional information. |

11. On the **Enable this Account Partner** page, if you do not want to enable the account partner now, clear the **Enable this account partner** check box, and then click **Next**.

12. To add the new account partner and close the wizard, click **Finish**.

**Rename an imported account partner**

You can use the following procedure to rename an imported account partner.

**To rename an imported account partner**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.

2. In the console tree, double-click **Federation Service**, **Trust Policy**, **Partner Organizations**, and **Account Partners**.
3. Right-click the account partner, and then click **Rename**.
4. Type a new name for account partner.

**Related Information**

- Partner organizations
- Federation scenarios
- Federation trusts
- Federation Service
- Claims

# Complete the Add Account Store Wizard

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use the Add Account Store Wizard to add new Active Directory account stores and Active Directory Application Mode (ADAM) account stores to a Federation Service trust policy. New account stores are added to the end of the account store lists by default. Account stores are used to authenticate users and provide user identity information to Active Directory Federation Services.

**Create an Active Directory account store**

You can use the following procedure to create an Active Directory account store.

**To create an Active Directory account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **My Organization**.
3. Right-click **Account Stores**, point to **New**, and then click **Account Store**.
4. On the **Welcome to the Add Account Store Wizard** page, click **Next**.
5. On the **Account Store Type** page, click **Active Directory**, and then click **Next**.

   | **Note** |
   | --- |
   | You can have only one Active Directory store that is associated with a Federation Service. If the **Active Directory** option is not available, it is because an Active Directory store has already been created for this Federation Service. |

6. If you do not want to enable this account store now, on the **Enable this Account Store** page, clear the **Enable this account store** check box, and then click **Next**.
7. To add the new account store and close the wizard, click **Finish**.

**Create an ADAM account store**

You can use the following procedure to create an ADAM account store.

**To create an ADAM account store**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **My Organization**.
3. Right-click **Account Stores**, point to **New**, and then click **Account Store**.
4. On the **Welcome to the Add Account Store Wizard** page, click **Next**.
5. On the **Account Store Type** page, click **Active Directory Application Mode (ADAM)**, and then click **Next**.
6. On the **ADAM Store Details** page, do the following, and then click **Next**:
   - In **Account store display name**, type the friendly name of the account store.
   - In **Account store URI**, type the Uniform Resource Identifier (URI) for the ADAM account store.
7. On the **ADAM Server Settings** page, do the following, and then click **Next**:
   - In **ADAM server name or IP address**, type the name or Internet Protocol (IP) address of the ADAM server.
   - In **Port number**, type the TCP/IP port number for the account service.
   - In **LDAP search base distinguished name**, type the distinguished name, for example, DC=adatum,DC=com.

- In **User name LDAP attribute**, type the name of the user name attribute, for example, **userPrincipalName**.

8. On the **Identity Claims** page , select one or more identity claims that will be provided by the account store, and then click **Next**:
   - If the account store provides UPN identity claims, select the **User Principal Name (UPN)** check box, and then type the Lightweight Directory Access Protocol (LDAP) attribute name.
   - If the account store provides e-mail identity claims, select the **E-mail** check box, and then type the LDAP attribute name.
   - If the account store provides common name identity claims, select the **Common name** check box, and then type the LDAP attribute name.
9. If you do not want to enable this account store now, on the **Enable this Account Store** page, clear the **Enable this account store** check box, and then click **Next**.
10. To add the new account store and close the wizard, click **Finish**.

**Related Information**

- Account stores
- Federation Service
- Claims

# Complete the Add Applications Wizard

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use the Add Application Wizard to add either a claims-aware application or a Windows NT token–based application to the Federation Service trust policy.

**Add a claims-aware application**

You can use the following procedure to add a claims-aware application to the Federation Service trust policy.

**To add a claims-aware application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **My Organization**.
3. Right-click **Applications**, point to **New**, and then click **Application**.
4. On the **Welcome to the Add Application Wizard** page, click **Next**.
5. On the **Application Type** page, click **Claims-aware application**, and then click **Next**.
6. On the **Application Details** page, do the following, and then click **Next**:
   - In **Application display name**, type the name of the application.
   - In **Application URL**, type the Uniform Resource Locator (URL) of the application.

| **Important** |
| --- |
| This URL must match the return URL that is configured on the ADFS Web Agent for this application. |

7. On the **Accepted Identity Claims** page, select each identity claim type that the application will use to make authorization decisions, and then click **Next**:
   - If the application requires user principal name (UPN) identity claims to make authorization decisions, select the **User principal name (UPN)** check box.
   - If the application requires e-mail identity claims to make authorization decisions, select the **E-mail** check box.
   - If the application requires common name identity claims to make authorization decisions, select the **Common name** check box.
8. If you do not want to enable the claims-aware application now, on the **Enable this Application** page, clear the **Enable this application** check box, and then click **Next**.
9. To add the new claims-aware application and close the wizard, click **Finish**.

**Add a Windows NT token–based application**

You can use the following procedure to add a Windows NT token–based application to the Federation Service trust policy.

**To add a Windows NT token–based application**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **My Organization**.
3. Right-click **Applications**, point to **New**, and then click **Application**.
4. On the **Welcome to the Add Application Wizard** page, click **Next**.

5. On the **Application Type** page, click **Windows NT token–based application**, and then click **Next**.
6. On the **Application Details** page, do the following, and then click **Next**:
   - In **Application display name**, type the name of the application.
   - In **Application URL**, type the Uniform Resource Locator (URL) of the application.

| **Important** |
| --- |
| This URL must match the return URL that is configured on the ADFS Web Agent for this application. |

7. On the **Accepted Identity Claim** page, select each identity claim type that the application will use to make authorization decisions, and then click **Next**:
   - If the application requires user principal name (UPN) identity claims to make authorization decisions, click **User principal name (UPN)**.
   - If the application requires e-mail identity claims to make authorization decisions, click **E-mail**.
8. If you do not want to enable the Windows NT token–based application now, on the **Enable this Application** page, clear the **Enable this application** check box, and then click **Next**.
9. To add the new Windows NT token–based application and close the wizard, click **Finish**.

**Related Information**

- Claims-aware applications
- Windows NT token-based applications
- Federation Service
- Claims

# Configure the Account Store Priority

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When you add more than one account store, account store priority determines the order in which account stores will be used by Active Directory Federation Services (ADFS). ADFS will attempt to authenticate users beginning with the first account store. Only if authentication fails for a user will ADFS attempt to authenticate the user with the next account store in the list. ADFS will attempt to authenticate a user until the user is successfully authenticated or until all account stores have been attempted. ADFS will stop attempting additional account stores as soon as one account store successfully authenticates a user.

User identity attributes that are provided by account stores are not cumulative. ADFS is aware only of user identity attributes that are provided by the account store that is used to authenticate the user.

**Configure the account store priority**
You can use the following procedure to configure the account store priority.

**To configure the account store priority**

1. Click **Start**, point to **Administrative Tools**, and then click **Active Directory Federation Services**.
2. In the console tree, double-click **Federation Service**, **Trust Policy**, and **My Organization**.
3. Right-click **Account Stores**, and then click **Store Priority**.
4. In the **Account Store Prioritization** dialog box, select an account store, use the **Up** and **Down** buttons to move the account store in the priority list, and then click **OK**.

**See Also**
**Concepts**
Account stores
Federation Service

# Manage ADFS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage ADFS**

For information about how to manage Active Directory Federation Services (ADFS) after you deploy it, look for ADFS operations content on the Windows Server 2003 R2 page (http://go.microsoft.com/fwlink/?LinkId=45560) of the Microsoft Windows Server 2003 TechCenter Web site.

# ADFS Concepts

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ADFS Concepts**


- ADFS Overview
- Understanding ADFS Components
- Understanding ADFS Authentication and Authorization
- Understanding Certificates Used by ADFS
- ADFS resources

# ADFS Overview

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ADFS Overview**


- Introduction to ADFS
- ADFS server roles
- Federation scenarios
- Partner organizations
- Federation trusts
- ADFS requirements
- Terminology used in ADFS

# Introduction to ADFS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

ADFS is a component in Microsoft® Windows Server™ 2003 R2 that provides Web single-sign-on (SSO) technologies to authenticate a user to multiple Web applications over the life of a single online session. ADFS accomplishes this by securely sharing digital identity and entitlement rights, or "Claims," across security and enterprise boundaries.

ADFS is not:

- A database or repository for employee or customer identity data.
- An extension of the Active Directory™ directory service schema.
- A type of Windows domain or forest trust.

ADFS in Windows Server 2003 R2 supports the WS-Federation Passive Requestor Profile (WS-F PRP).

**Key features of ADFS**

The following are some of the key features of ADFS in Windows Server 2003 R2:

- Federation and Web SSO

  When an organization uses the Active Directory™ directory service, it currently experiences the benefit of SSO functionality through Windows-integrated authentication within the organization's security or enterprise boundaries. ADFS extends this functionality to Internet-facing applications, which enables customers, partners, and suppliers to have a similar, streamlined, Web SSO user experience when they access the organization's Web-based applications. Furthermore, federation servers can be deployed in multiple organizations to facilitate business-to-business (B2B) federated transactions between partner organizations. For more information about ADFS federation, see Federation scenarios.

- Web Services (WS)-* interoperability

  ADFS provides a federated identity management solution that interoperates with other security products that support the WS-* Web Services Architecture. ADFS does this by employing the federation specification of WS-*, called WS-Federation. The WS-Federation specification makes it possible for environments that do not use the Windows identity model to federate with Windows environments. For more information about WS-* specifications, see ADFS resources.

- Extensible architecture

  ADFS provides an extensible architecture that supports the Security Assertion Markup Language (SAML) token type and Kerberos authentication (in the Federated Web SSO with Forest Trust scenario). ADFS can also perform claim mapping, for example, modifying claims using custom business logic as a variable in an access request. Organizations can use this extensibility to modify ADFS to coexist with their current security infrastructure and business policies. For more information about modifying claims, see Claim mapping.

**Extending Active Directory to the Internet**

Active Directory serves as a primary identity and authentication service in many organizations. With Windows Server 2003 Active Directory, forest trusts can be created between two or more Windows Server 2003 forests to provide access to resources that are located in different business units or organizations. For more information about forest trusts, see How Domain and Forest Trusts Work on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=35356).

However, there are scenarios in which forest trusts are not a viable option. For example, access across organizations may need to be limited to only a small subset of individuals, not every member of a forest.

By employing ADFS, organizations can extend their existing Active Directory infrastructures to provide access to resources that are offered by trusted partners across the Internet. These trusted partners can include external third parties or other

departments or subsidiaries in the same organization.

ADFS is tightly integrated with Active Directory. ADFS retrieves user attributes from Active Directory, and it authenticates users against Active Directory. ADFS also uses Windows Integrated Authentication.

ADFS works with both Active Directory and Active Directory Application Mode (ADAM). Specifically, ADFS works with both enterprise-wide deployments of Active Directory or instances of ADAM. When it works with Active Directory, ADFS can take advantage of the strong authentication technologies in Active Directory, including Kerberos, X.509 digital certificates, and smart cards. When it works with ADAM, ADFS uses Lightweight Directory Access Protocol (LDAP) Bind as a means to authenticate users. For more information about how ADFS works with Active Directory and ADAM, see Account stores.

ADFS supports distributed authentication and authorization over the Internet. ADFS can be integrated into an organization's or department’s existing access management solution to translate the terms that are used in the organization into claims that are agreed on as part of a federation. ADFS can create, secure, and verify the claims that move between organizations. It can also audit and monitor the activity between organizations and departments to help ensure secure transactions.

# ADFS server roles

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) can operate only when the servers running Windows Server 2003 R2 are configured appropriately. Depending on the environment in your organization, specific ADFS server roles must be deployed. The following sections describe the server roles that can be used to provide an ADFS federated identity management solution.

**Federation server**
Federation servers host the Federation Service component of ADFS. They are used to route authentication requests that are made from user accounts in other organizations (in Federated Web Single-Sign-On (SSO) scenarios) or from clients that can be located anywhere on the Internet (in the Web SSO scenario). For more information about the different ADFS scenarios, see Federation scenarios.

Federation servers also host a security token service that issues tokens that are based on the credentials (for example, user name and password) that are presented to it. After the credentials are verified (by the user logging on), Claims for the user are collected through examination of the attributes for the user that are stored in Active Directory or Active Directory Application Mode (ADAM).

In Federated Web SSO scenarios, claims can then be modified by claim mappings for a specific resource partner. The claims are built into a token that is sent to a federation server in the resource partner. After a federation server in the resource partner receives the claims as incoming claims, it maps them into their organization claims. The organization claims are then built into a new token that is sent to the ADFS Web Agent.

The role that a federation server plays in either of the Federated Web SSO scenarios (Federated Web SSO or Federated Web SSO with Forest Trust) can depend on whether your organization is designated as the account partner or the resource partner:

- Federation servers in the account partner are used to log on local user accounts in either an Active Directory store or an Active Directory Application Mode (ADAM) store. Federation servers also issue initial security tokens that the local user accounts can use to access Web-based applications that are hosted in the resource partner. In addition, federation servers in the account partner issue cookies to users to maintain login status. These cookies include claims for those users. These cookies enable SSO capabilities so that users do not have to enter credentials each time that they visit different Web-based applications in the resource partner.

- Federation servers at the resource partner validate the security tokens that are issued by the federation servers at the account partner. Federation servers at the resource partner also issue security tokens that are meant for the Web-based applications in the resource partner. In addition, federation servers in the resource partner issue cookies to the user accounts, which come from the account partner. These cookies enable SSO capabilities so that users do not have to log in again at their federation servers in the account partner when users attempt to access different Web-based applications at the resource partner.

For more information about the account and resource partners, see Partner organizations.

**Federation server proxy**
Federation server proxies host the Federation Service Proxy component of ADFS. Federation server proxies can be deployed in an organization's perimeter network (also known as a demilitarized zone, extranet, or screened subnet) to forward requests to federation servers that are not accessible from the Internet.

| **Note** |
| --- |
| Although you can deploy separate servers to host the Federation Service Proxy component, it is not necessary to deploy a separate server to act as a federation server proxy in the intranet forest of either the account partner or the resource partner. A federation server performs this role automatically. |

The role that a federation server proxy plays in your organization can depend on whether your organization is the account partner or the resource partner:

- Federation server proxies at the account partner act as proxies for user logons to federation servers that are located in the intranet. Federation server proxies also act as proxies for security tokens that are issued by the account partner

federation server for both its own tokens and those tokens that are destined for resource partners.

- Federation service proxies in the resource partner act as proxies for user security tokens, which are issued from federation servers in both the account partner and the resource partner, to Web-based applications in the resource partner.

**Web server**

In ADFS, Web servers in the resource forest host the ADFS Web Agent component to provide secure access to the Web applications that are hosted on those Web servers. The ADFS Web Agent manages security tokens and authentication cookies that are sent to a Web server. The Web server requires a relationship with a Federation Service so that all authentication tokens come from that Federation Service.

The ADFS Web Agent supports two types of applications: claims-aware applications and Windows NT token–based applications. For information about these types of applications, see Controlling Access to Web-based Applications.

# Federation scenarios

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) supports federated identity scenarios that use Web Services Federation (WS-Federation), WS-Federation Passive Requestor Profile (WS-F PRP), and WS-Federation Passive Requestor Interoperability Profile specifications. The ADFS solution helps administrators deal with federated identity management challenges by making it possible for organizations to securely share a user's identity information over Federation trusts. The following three descriptions of deployment scenarios illustrate how you can use ADFS server roles to federate identities, depending on the needs of your organization.

**Federated Web SSO**

The ADFS Federated Web Single-Sign-On (SSO) scenario involves secure communication that often spans multiple firewalls, perimeter networks, and name resolution servers, in addition to the entire Internet routing infrastructure. Communication over a federated Web SSO environment can help foster more efficient and secure online transactions between organizations that are joined by federation trust relationships.

As shown in the following illustration, a federation trust relationship can be established between two businesses. In this scenario, federation servers route authentication requests from user accounts in Tailspin Toys to Web-based applications that are located in the network of Online Retailer.




Federation servers authenticate requests from trusted partners based on the credentials of the partners. Representations of the credentials are exchanged in the form of security tokens.

For enhanced security, federation server proxies can be used to relay requests to federation servers that are not directly accessible from the Internet.

**Federated Web SSO with Forest Trust**

The ADFS Federated Web SSO with Forest Trust scenario involves two Active Directory forests in a single organization, as shown in the following illustration. One of the forests is located in the organization's perimeter network (also known as a demilitarized zone, extranet, or screened subnet). The other forest is located in the internal network. A one-way, forest trust is established so that the forest in the perimeter network trusts the forest in the internal network. Federation servers are deployed in both networks. A federation trust is established so that accounts in the internal forest can be used to access a Web-based application in the perimeter network, whether the accounts access the site from the intranet forest or from the Internet.




**A. Datum Corp.**

In this scenario, external users, such as customers, can access the Web application by authenticating to the external account federation server, which is located in the perimeter network. External users have user accounts in the perimeter-network Active Directory forest. Internal users, such as employees, can also access the Web application by authenticating to the internal account federation server, which is located in the internal network. Internal users have accounts in the internal Active Directory forest.

If the Web-based application is a Windows NT token-based application, the ADFS Web Agent that is running on the Web application server intercepts requests and creates Windows NT security tokens, which are required by the Web application to make authorization decisions. For external users, this is possible because the Web server that hosts the Windows NT token-based application is joined to the domain in the external forest. For internal users, this is enabled through the forest trust relationship that exists between the perimeter forest and the internal forest.

If the Web-based application is a claims-aware application, the ADFS Web Agent that is running on the Web application server does not have to create Windows NT security tokens for the user. The ADFS Web agent can expose the claims that come across, which allows the application to make authorization decisions based on the contents of the security token that is provided by the account federation server. As a result when it deploys claims-aware applications, the Web server does not have to be joined to the domain, and the external-forest-to-internal-forest trust is not required.

**Web SSO**

In the ADFS Web SSO scenario, users must authenticate only once to access multiple Web-based applications. In this scenario all users are external, and no federation trust exists. Because the Web servers must be Internet accessible and also joined to the Active Directory domain, they are connected to two networks; that is, they are multihomed. The first network is Internet facing (the perimeter network) to provide the needed connectivity. The second network contains the Active Directory forest (the protected network), which is not directly Internet accessible. The federation server proxy is also multihomed to provide the necessary connectivity to the federation server and the Internet. In this scenario, placing the federation server on a network that is not directly accessible from the Internet greatly reduces the risk to the federation server.


Firewall
Protected Network
AD DS
Federation server
Federation server proxy
Perimeter Network
App1
App2
App3
AD FS-enabled Web servers
Internet
Client
Tailspin Toys

# Partner organizations

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When you plan for cross-organizational (federation-based) collaboration using Active Directory Federation Services (ADFS), you first determine whether your organization will host a Web resource to be accessed by other organizations across the Internet — or the reverse. This determination affects how you deploy ADFS, and it is fundamental in the planning of your ADFS infrastructure.

For Federation scenarios such as Federated Web Single Sign-On (SSO) and Federated Web SSO with Forest Trust (but not the Web SSO scenario), ADFS uses terms such as "account partner" and "resource partner" to help differentiate the organization that hosts the accounts (the account partner) from the organization that hosts the Web-based resources (the resource partner). The term "Federation trusts" is used in ADFS to characterize the one-way, nontransitive relationship that is established between the account partner and the resource partner.

The following sections explain some of the concepts that are related to account partners and resource partners.

**Account partner**

An account partner represents the organization in the federation trust relationship that physically stores user accounts in either an Active Directory store or an Active Directory Application Mode (ADAM) store. The account partner is responsible for collecting and authenticating a user's credentials, building up claims for that user, and packaging the claims into security tokens. These tokens can then be presented across a federation trust to access Web-based resources that are located in the resource partner organization.

In other words, an account partner represents the organization for whose users the account-side Federation Service issues security tokens. The Federation Service in the account partner organization authenticates local users and creates security tokens that are used by the resource partner in making authorization decisions.

In relation to Active Directory, the account partner in ADFS is conceptually equivalent to a single Active Directory forest whose accounts need access to resources that are physically located in another forest. Accounts in this example forest can access resources in the resource forest only when an external trust or forest trust relationship exists between the two forests and the resources to which the users are trying to gain access have been set with the proper authorization permissions.

| **Note** |
| --- |
| This analogy is meant strictly to emphasize how the relationship between account and partner organizations in ADFS is similar, in concept, to the relationship between an account forest and a resource forest in Active Directory. External trusts and forest trusts are not required for ADFS to function. |

**Producing claims that go to the resource partner**

A claim is a statement that a server makes (for example, name, identity, key, group, privilege, or capability) about a client. An account partner produces claims that the resource partner Federation Service consumes. The following list describes the different types of claims that can be configured in the account partner:

- UPN claim

  When you configure the account partner, you can specify a list of user principal name (UPN) domains and suffixes that may be accepted from the account partner. If a UPN identity is received whose domain part is not in the list, the request is rejected.

- E-mail claim

  When you configure the account partner, you can specify a list of e-mail domains and suffixes that may be accepted from the account partner. As with the UPN claim, if an e-mail identity is received whose domain part is not in the list, the request is rejected.

- Common name claim

  When you configure the account partner, you can specify whether common name claims can be received from the account partner. This type of claim may not be mapped; it is simply passed through if it is enabled.

- Group claims

  When you configure the account partner, you can specify a set of incoming group claims that may be accepted from the partner. You can then associate each possible incoming group with an organization group claim. Note that this creates a group mapping. If an incoming group is encountered that has no mapping, it is discarded.

- Custom claims

  When you configure the account partner, you can specify a set of incoming names of custom claims that are accepted from the partner. You can then map each possible incoming name to an organization custom claim. Note that this creates a name mapping. If an incoming custom claim is encountered that has no mapping, it is discarded.

**Resource partner**

A resource partner is the second organizational partner in the federation trust relationship. A resource partner is where the Web servers that host one or more Web-based applications (the resources) reside. The resource partner trusts the account partner to authenticate users. Therefore, to make authorization decisions, the resource partner consumes the claims that are packaged in security tokens coming from users in the account partner.

In other words, a resource partner represents the organization whose Web servers are protected by the resource-side Federation Service. The Federation Service at the resource partner uses the security tokens that are produced by the account partner to make authorization decisions for Web servers that are located in the resource partner.

To function as an ADFS resource, Web servers in the resource partner organization must have the ADFS Web Agent component of ADFS installed. Web servers that function as an ADFS resource can host either claims-aware applications or Windows NT token-based applications. For more information about the two types of applications, see Controlling Access to Web-based Applications.

| **Note** |
| --- |
| If the application that is hosted on the Web server is a Windows NT token–based application, a resource account may be required for the Active Directory forest in the resource partner organization. |

In relation to Active Directory, the resource partner is conceptually equivalent to a single forest whose resources are made available over an external trust or forest trust relationship to accounts that are physically stored in another forest.

| **Note** |
| --- |
| This analogy is meant strictly to emphasize how the relationship between account and partner organizations in ADFS is similar, in concept, to the relationship between an account forest and a resource forest in Active Directory. External trusts and forest trusts are not required for ADFS to function. |

**Consuming claims that come from the account partner**

A resource partner consumes claims that the account partner Federation Service produces and packages in security tokens. The following list describes how claims can be sent to the resource partner:

- UPN claim

  When you configure the resource partner, you can specify whether a UPN claim is to be sent to the resource partner. You can also specify a suffix mapping so that any suffix is mapped into a specified outgoing suffix. For example, julianp@sales.tailspintoys.com can be mapped to julianp@tailspintoys.com. Note that only one outgoing suffix may be specified.

- E-mail claim

  When you configure the resource partner, you can specify whether an e-mail claim is to be sent to the resource partner. You can also specify a suffix mapping so that any suffix is mapped into a specified suffix. For example, vernettep@sales.tailspintoys.com can be mapped to vernettep@tailspintoys.com. Note that only one outgoing suffix may be specified.

- Common name claim

  When you configure the resource partner, you can specify whether common name claims can be sent to the resource partner. This type of claim may not be mapped; it is simply passed through to the resource partner if it is enabled.

- Group claims

  When you configure the resource partner, you can specify a set of outgoing group claims that will be accepted by the resource partner. You can then associate each possible outgoing group claim to organization group claims. Note that this creates a set of group mappings. Organization group claims that do not match an outgoing group claim are not created.

- Custom claims

  When you configure the resource partner, you can specify a set of outgoing custom claims that are accepted by the resource partner. You can map each possible outgoing custom claim to a custom claim. Note that this creates a set of name mappings. Organization custom claims that do not match an outgoing custom claim are not created.

#### Enhanced identity privacy

**Enhanced identity privacy** is an optional setting that can be configured on a resource partner in the trust policy. If the **Enhanced identity privacy** option is enabled, this setting hashes the user-name portion of outgoing UPN claims and e-mail claims. It substitutes the common name with a random value.

The purpose of this feature is to prevent:

- The resource partner from correlating identity claims to personally identifiable user information.

- Collusion between partners in correlating identity claims to personally identifiable user information. This setting creates a unique hash per partner so that identity claim values are different across different trusting realm partners but consistent across sessions for a single partner.

- Simple dictionary attacks against the hash by salting the user value with data that is in the trust policy — data that is not known by the resource partners.

#### How enhanced identity privacy works

If the **Enable enhanced identity privacy** setting is enabled, the user's identity is transformed by the account Federation Service before it is sent to the resource realm. The account Federation Service does this by hashing a combination of the privacy key (for salting), the resource partner Uniform Resource Identifier (URI), and the identity claim itself. Therefore, the resulting transformed identity will be the same when the following values are the same:

- Privacy key

- Resource partner URI

- Identity claim (UPN or e-mail)

The hashing algorithm is case sensitive. Therefore, the hashed results are different if the case of the URI or the user-name portion of the UPN claim or e-mail claim is different.

# Federation trusts

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use Active Directory Federation Services (ADFS) to enable efficient and secure online transactions between Partner organizations that are joined by federation trust relationships. In other words, a federation trust is the embodiment of a business-level agreement or partnership between two organizations.

As shown in the following illustration, you can establish federation trust relationships between two partner organizations when both of the organizations deploy at least one ADFS federation server and they configure their Federation Service settings appropriately. The one-way arrow signifies the direction of the trust, which — like the direction of Windows trusts — always points to the account side of the forest. This means that authentication flows from the account partner organization to the resource partner organization.




| Note |
| --- |
| No communication occurs over the network between the account Federation Service and the resource Federation Service. |

After you create the federation trust, users who are located in the account partner organization can send authentication requests successfully through the federation trust to the Web server in the resource partner organization. A federation trust is created when the account partner organization and the resource partner organization both install the Federation Service component of ADFS and they both use the Active Directory Federation Services snap-in to configure the account partner and resource partner appropriately.

If one side of a federation trust (either the account partner or the resource partner) is not configured or if it is configured incorrectly by the administrator for either organization, the federation trust will not be created successfully. For detailed guidance about how to create federation trusts, look for ADFS step-by-step or deployment content on the Windows Server 2003 R2 page (http://go.microsoft.com/fwlink/?LinkId=45560) of the Microsoft Windows Server 2003 TechCenter Web site.

| Note |
| --- |
| Federation trusts are not used in the Web Single-Sign-On (SSO) scenario. For more information about the Web SSO scenario, see Federation scenarios. |

# ADFS requirements

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) has the following hardware and software requirements.

**Hardware requirements**

- Processor speed: 133 MHz for x86-based computers
- Recommended minimum RAM: 256 MB
- Free disk space for setup: 10 MB

**Software requirements**

ADFS relies on server functionality that is built into the Windows Server 2003 R2 operating system. The Federation Service, Federation Service Proxy, and ADFS Web Agent components cannot run on earlier operating systems. This section describes the software requirements for each ADFS component. It also describes the overall software configurations that are necessary for ADFS in your network environment.

| **Note** |
| --- |
| The Federation Service and Federation Service Proxy components cannot coexist on the same computer, and they can be installed only on computers running Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Datacenter Edition. |

**Federation Service**

Computers running the Federation Service must have the following software installed:

- Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Datacenter Edition
- Internet Information Services (IIS)
- Microsoft ASP.NET 2.0
- Microsoft .NET Framework 2.0
- A default Web site that is configured with Transport Layer Security and Secure Sockets Layer (TLS/SSL)
- A certificate for the Federation Service

| **Note** |
| --- |
| This certificate is used for signing tokens. It must be a digital signing X.509 certificate. |

**Active Directory and ADAM account store requirements**

ADFS requires the presence of user accounts in Active Directory or Active Directory Application Mode (ADAM) for the account Federation Service. Active Directory domain controllers or computers hosting the account stores must have the following software installed:

- Windows Server 2003

Or

- Windows Server 2003 R2

Or

- Windows 2000 with Service Pack 4 (SP4) with critical updates

ADFS does not require schema changes or functional-level modifications to Active Directory. To ensure that ADAM works with ADFS, install the version of ADAM that comes with Windows Server 2003 R2.

**Federation Service Proxy**
Computers running the Federation Service Proxy must have the following software installed:

- Windows Server 2003 R2, Enterprise Edition, or Windows Server 2003 R2, Datacenter Edition
- IIS
- ASP.NET 2.0
- Microsoft .NET Framework 2.0
- A default Web site configured with TLS/SSL

**ADFS Web Agent**
Computers running the ADFS Web Agent must have the following software installed:

- Windows Server 2003 R2, Standard Edition; Windows Server 2003 R2, Enterprise Edition; or Windows Server 2003 R2, Datacenter Edition
- IIS
- ASP.NET 2.0
- Microsoft .NET Framework 2.0
- After the ADFS Web Agent installation is complete, at least one Web site in IIS must be configured with TLS/SSL so that federated users can access Web-based applications that are hosted on the Web server.

**Trusted certification authorities**
Because both TLS/SSL and token signing rely on digital certificates, certification authorities (CAs) are an important part of ADFS. Public CAs, such as VeriSign, Inc., represent a mutually trusted third party that allows the identity of the bearer of a certificate to be identified. You can use enterprise CAs, such as Microsoft Certificate Services, for providing token signing and other internal certificate services.

If a client is presented with a server's authentication certificate, the client computer verifies that the CA that issued the certificate is in the client's list of trusted CAs and that the CA has not revoked that certificate. This verification ensures that the client has reached the intended server. When a certificate is used for verifying signed tokens, the client uses the certificate to verify that the token was issued by the correct federation server and that the token has not been tampered with.

**TCP/IP network connectivity**
For ADFS to function, TCP/IP network connectivity must exist between the client; a domain controller; and the computers that host the Federation Service, the Federation Service Proxy (when it is used), and the ADFS Web Agent.

**DNS**
For the purpose of authenticating users in the intranet, internal Domain Name System (DNS) servers in the intranet forest must be configured to return the canonical name (CNAME) of the internal server that is running the Federation Service. For best results, do not use Hosts files with DNS.

**Web browser**
Although any current Web browser with JScript enabled should work as an ADFS client, only Internet Explorer 6, Internet Explorer 5 or 5.5, Mozilla Firefox, and Safari on Apple Macintosh have been tested by Microsoft. For performance reasons, it is highly recommended that JScript be enabled. Cookies must be enabled, or at least trusted, for the federation servers and Web applications that are being accessed.

# Terminology used in ADFS

Updated: March 27, 2008

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) uses terminology from several different technologies, including certificate services, Internet Information Services (IIS), Active Directory, Active Directory Application Mode (ADAM), and Web Services (WS-*). The following table describes these terms.

| Term | Description |
|---|---|
| account partner | A federation partner that is trusted by the Federation Service to provide security tokens. The account partner issues these tokens to its users (that is, users in the account partner realm) so that they can access Web-based applications in the resource partner. |
| Active Directory Federation Services (ADFS) | A Windows Server 2003 R2 component that provides Web single-sign-on (SSO) technologies to authenticate a user to multiple Web applications over the life of a single online session. ADFS accomplishes this by securely sharing digital identity and entitlement rights across security and enterprise boundaries. ADFS in Windows Server 2003 R2 supports the WS-Federation Passive Requestor Profile (WS-F PRP). |
| claim | A statement that an issuer makes (for example, name, identity, key, group, privilege, or capability) about a client. |
| claim mapping | The act of mapping, removing or filtering, or passing claims between various claim sets. |
| claims-aware application | A Microsoft ASP.NET application that performs authorization based on the claims that are present in an ADFS security token. |
| client account partner discovery Web page | The Web page that is used to interact with the user to determine which account partner the user belongs to when ADFS cannot automatically determine which of the account partners should authenticate the user. |
| client logoff Web page | When ADFS performs a logoff operation, a Web page that is executed to provide visual feedback to the user that the logoff has occurred. |
| client logon Web page | When ADFS collects client credentials, a Web page that is executed to perform the user interaction. The client logon Web page may use any necessary business logic to determine the type of credentials to collect. |
| federation | A pair of realms or domains that have established a federation trust. |
| Federation Service | A security token service that is built into Windows Server 2003 R2. The Federation Service provides tokens in response to requests for security tokens. |
| Federation Service Proxy | A proxy to the Federation Service in the perimeter network (also known as a demilitarized zone, extranet, or screened subnet). The Federation Service Proxy uses WS-F PRP protocols to collect user credential information from browser clients and Web applications and send the information to the Federation Service on their behalf. |
| organization claims | Claims in intermediate or normalized form within an organization's namespace. |
| passive client | A Hypertext Transfer Protocol (HTTP) browser, capable of broadly supported HTTP, that can make use of cookies. ADFS in Windows Server 2003 R2 supports only passive clients, and it adheres to the WS-F PRP specification. |

| | |
|---|---|
| resource partner | A federation partner that trusts the Federation Service to issue claims-based security tokens. The resource partner contains published Web-based applications that users in the account partner can access. |
| security token | A cryptographically signed data unit that expresses one or more claims. |
| security token service (STS) | A Web service that issues security tokens. An STS makes assertions based on evidence that it trusts, to whoever trusts it (or to specific recipients). To communicate trust, a service requires proof, such as a signature to prove knowledge of a security token or set of security tokens. A service itself can generate tokens or it can rely on a separate STS to issue a security token with its own trust statement. This forms the basis of trust brokering. In ADFS, the Federation Service is an STS. |
| server farm | In ADFS, a collection of load-balanced federation servers, federation server proxies, or Web servers hosting the ADFS Web Agent. |
| single sign-on (SSO) | An optimization of the authentication sequence to remove the burden of repeated logon actions by an end user. |
| token-signing certificate | An X509 certificate whose associated public/private key pair is used to provide integrity for security tokens. |
| Uniform Resource Identifier (URI) | A compact string of characters that identifies an abstract resource or physical resource. URIs are explained in RFC 2396 (http://go.microsoft.com/fwlink/?LinkId=48289). In ADFS, URIs are used to uniquely identify partners and account stores. |
| Web Services (WS-*) | The specifications for a Web Services Architecture that is based on industry standards such as Simple Object Access Protocol (SOAP); XML; Web Service Description Language (WSDL); and Universal Description, Discovery, and Integration (UDDI). WS-* provides a foundation for delivering complete, interoperable business solutions for the extended enterprise, including the ability to manage federated identity and security.<br><br>The Web services model is based on the idea that enterprise systems are written in different languages, with different programming models, which run on and are accessed from many different types of devices. Web services are a means of building distributed systems that can connect and interact with one another easily and efficiently across the Internet, regardless of what language they are written in or what platform they run on. |
| Web Services Security (WS-Security) | A series of specifications that describes how to attach signature and encryption headers to SOAP messages. In addition, WS-Security describes how to attach security tokens, including binary security tokens such as X.509 certificates and Kerberos tickets, to messages. In ADFS, WS-Security is used when Kerberos signs security tokens. |
| Windows NT token–based application | A Windows application that relies on a Windows NT token to perform authorization of users. |
| WS-Federation | A specification that defines a model and set of messages for brokering trust and the federation of identity and authentication information across different trust realms.<br><br>The WS-Federation specification identifies two sources of identity and authentication requests across trust realms: active requestors, such as SOAP-enabled applications, and passive requestors, which are defined as HTTP browsers capable of supporting broadly supported HTTP, for example, HTTP 1.1. |

| | |
|---|---|
| WS-Federation Passive Requestor Profile (WS-F PRP) | An implementation of the WS-Federation specification that proposes a standard protocol for how passive clients (such as Web browsers) apply the federation framework. Within this protocol, Web service requestors are expected to understand the new security mechanisms and be capable of interacting with Web service providers. |

# Understanding ADFS Components

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Understanding ADFS Components**


- Federation Service
- Federation Service Proxy
- ADFS Web Agent
- Active Directory Federation Services snap-in
- Active Directory Federation Services Proxy snap-in

# Federation Service

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The Federation Service is a component of Active Directory Federation Services (ADFS) that can be installed independently from other ADFS components. The Federation Service functions as a security token service. The act of installing the Federation Service component on a computer makes that computer a federation server. It also makes the Active Directory Federation Services snap-in available in the **Administrative Tools** menu on that computer.

The Federation Service is designed to use Active Directory to provide tokens in response to requests for security tokens. This enables Active Directory domains and forests to function as:

- Identity providers that can federate with compliant account partners and resource partners. As an identity provider, the Federation Service can project Active Directory identities across the Internet to interact with applications at compliant service providers.

- Service providers that can federate with compliant account partners and resource partners. As a service provider, the Federation Service can allow identities from other organizations to access a partner's Windows-based and ASP.NET-based applications.

- Security token providers for applications that are compliant with the WS-Federation Passive Requestor Profile (WS-F PRP) specification.

As an account partner, the Federation Service allows Active Directory users to access resources at partner organizations. In response to a request from a resource partner, the Federation Service collects and verifies user credentials against Active Directory or Active Directory Application Mode (ADAM). The Federation Service can then populate a set of organization claims based on the Lightweight Directory Access Protocol (LDAP) attributes of the user account. The organization claims are then mapped to appropriate claims for the resource partner and packaged into a security token that is signed by the Federation Service's token-signing certificate. The resultant security token is posted as the response to the resource partner's original request. The resource partner then uses the token to allow access for the user.

As a resource partner, the Federation Service plays the opposite role. When a user attempts to access an ADFS-protected application, the Federation Service determines which account partner should authenticate the user. It then sends an authentication request to that partner. When the user returns with a security token, the Federation Service verifies that the token has been correctly signed by the partner. It then extracts the claims from the token. The claims are mapped to organization claims, and the filtering policy for the specific application is applied. The filtered organization claims are packaged into a security token that is either signed by the Federation Service's token-signing certificate or protected by a Kerberos session key for the Web application. The resultant security token is posted back to the original application Uniform Resource Locator (URL). The application then uses the token to allow access for the user.

ADFS uses the WS-F PRP protocol to carry Claims in security tokens that are issued by the Federation Service to the Web application. For more information about the WS-F PRP specification, see ADFS resources.

These claims are populated initially from Account stores, either Active Directory or ADAM account stores. The Federation Service issues tokens based on the credentials that are presented. After the credentials are verified using the account store, the claims for the user are generated according to the rules of the trust policy. The obtained inbound claims are mapped into outbound claims that are appropriate for a resource partner. The resulting claim mappings are added to a security token that is issued to the resource partner. For more information, see Claim mapping.

After the token is verified, an authentication cookie is issued and written to the client browser. Each time that the client needs to be authenticated, this cookie is used by the Federation Service so that the client does not have to enter credentials again, thereby enabling single sign-on (SSO). For more information, see Cookies used by ADFS.

**Federation Service Web pages**

The Federation Service provides a Web page that prompts the user to select an appropriate account partner to which the user can authenticate. The Federation Service also provides a Web page that prompts for the user's credentials, such as a user name and password for forms-based authentication. A Web page is also provided that supports Windows Integrated Authentication.

Behind the Web pages, the Federation Service provides a Microsoft ASP.NET Web service that processes requests from the client or the federation server proxy. The federation server proxy is located in the perimeter network. It acts as an intermediary

between an Internet client and a Federation Service in the intranet. For more information about the role of the federation server proxy, see ADFS server roles.

There are two basic types of requests to which the Federation Service responds:

- Requests to issue security tokens
- Requests to retrieve trust policy data

**Account partner discovery**

Account partner discovery is the process by which a client can identify what account partner it prefers for authentication in the event that more than a single account partner is configured. The Federation Server presents this choice to the client browser as a drop box containing the account partner names as they are configured in the trust policy.

One mechanism that you can use to avoid account partner discovery is to include the **`whr`** parameter in the query string for the resource being accessed, for example,

`https://webserver/testapp/testpage.aspx?whr=urn:federation:<`*accountpartner*`>`, where <*accountpartner*> indicates the account partner realm of the client.

When the **`whr`** parameter is used, the Web server removes the parameter and writes a cookie to the client browser to remember this setting for future requests. Then, the request proceeds in the same way as if it had not been provided.

# Federation Service Proxy

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The Federation Service Proxy is a component of Active Directory Federation Services (ADFS) in Windows Server 2003 R2 that can be installed independently from other ADFS components. The Federation Service Proxy functions as a proxy in a perimeter network (also known as a demilitarized zone, extranet, or screened subnet) for the Federation Service. The act of installing the Federation Service Proxy component on a computer makes that computer a federation server proxy. It also makes the Active Directory Federation Services Proxy snap-in available on that computer in the **Administrative Tools** menu.

A federation server proxy participates in the WS-Federation Passive Requestor Profile (WS-F PRP) protocol by communicating with a protected Federation Service on the client's behalf. When the federation server proxy is protecting an account partner, it collects user credential information from browser clients. When the federation server proxy is protecting a resource partner, it relays requests by and for Web applications to the Federation Service.

The federation server proxy also stores Hypertext Transfer Protocol (HTTP) cookies on clients when necessary to facilitate single sign-on (SSO). For more information, see Cookies used by ADFS. The federation server proxy writes all three types of cookies: authentication cookies, account partner cookies, and sign-out cookies. For more information about the federation server proxy role, see ADFS server roles.

# ADFS Web Agent

Updated: August 17, 2006

Applies To: Windows Server 2003 R2

The Active Directory Federation Services (ADFS) Web Agent is a component of ADFS. It is used to consume security tokens and either allow or deny a user access to a Web application. To accomplish this, the Web server requires a relationship with a resource Federation Service so that it can direct the user to the Federation Service as needed.

The ADFS Web Agent can be used for two different types of applications:

- Claims-aware applications: A Microsoft ASP.NET application that is written to published ADFS objects that allow the querying of ADFS security token Claims. The application makes authorization decisions based on these claims. Security token failures for this type of application result in the client seeing an "Access Denied" message and events in the Federation Service event log. For more information, see Claims-aware applications.

- Windows NT token–based applications: An application that uses Windows-based authorization mechanisms. The ADFS Web Agent supports conversion from an ADFS security token to an impersonation-level Windows NT access token. For more information, see Windows NT token-based applications.

The Web server also stores Hypertext Transfer Protocol (HTTP) cookies on clients where necessary to facilitate single sign-on (SSO). The ADFS Web Agent comprises two separate components:

- ADFS Web Agent ISAPI Extension

- ADFS Web Agent Authentication Service

**ADFS Web Agent ISAPI Extension**

The ADFS Web Agent ISAPI Extension is an Internet Server Application Programming Interface (ISAPI) extension that you can use to configure information in the Internet Information Services (IIS) metabase. In IIS Manager you can use the **ADFS Web Agent** tabs on the **Web Sites** and **Default Web Sites** property pages to administer policy and certificates that verify the ADFS security token and cookies.

The ADFS Web Agent properties in the following table are inheritable. These properties are required on an IIS resource if the ISAPI extension is going to support the WS-Federation Passive Requestor Profile (WS-F PRP) protocol.

| Properties | Description |
|---|---|
| Federation Service URL | The Uniform Resource Locator (URL) of the Federation Service. This URL is required so that it may be queried for trust information. |
| Cookie path | The path that is specified when the authentication cookie is written. |
| Cookie domain | The domain for which the cookie is valid. |
| Return URL | The URL that the token from the Federation Service comes back to after authentication at the Federation Service. This URL should match the Audience element of the token. The checking against the Audience element is performed by the Windows service. |

**Optional debug trace logging configuration**

Optional debug trace logging is configured for the ISAPI extension through the registry **DWORD** value in the following path: **HKLM\Software\Microsoft\ADFS\WebServerAgent\DebugPrintLevel**. If the value is present, a log in the directory that is specified at **HKLM\Software\Microsoft\ADFS\LogFilesPath** with the name ifsext_<*ApplicationPoolName*>.log is produced. A notification on the registry key **HKLM\Software\Microsoft\ADFS\WebServerAgent** is registered by the extension so that a change in the debug level may be made without restarting the IIS worker process. The following flag values are used to control the verbosity of the log file:

- Error – 0x1
- Warning – 0x2
- Trace - 0x4

**ADFS Web Agent Authentication Service**

The ADFS Web Agent Authentication Service validates incoming tokens and cookies. It runs as Local System to generate a token by using either Service-for-User (S4U) or the ADFS authentication package. However, the IIS application pool is not required to run as Local System.

The ADFS Web Agent Authentication Service has interfaces that may be called only with local remote procedure call (LRPC), not RPC. This service returns an impersonation Windows NT access token if it is given an ADFS security token or an ADFS cookie.

# Active Directory Federation Services snap-in

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The Active Directory Federation Services (ADFS) Microsoft Management Console (MMC) snap-in is installed when you install the Federation Service component in **Add or Remove Programs** in Windows Server 2003 R2, Enterprise Edition. You can use the ADFS snap-in to:

- Configure the Federation Service or federation server farm.
- Manage the trust policy that is associated with your Federation Service:
  - Administer Active Directory or Active Directory Application Mode (ADAM) Account stores.
  - Manage account partners and resource partners that will trust your Partner organizations.
  - Manage the applications that are going to be protected by ADFS.
  - Manage Claims, Certificates used by federation servers, and ADFS-protected Web applications.

Settings that you configure in the ADFS snap-in are stored partly in the Web.config file, which is located in the Federation Service virtual directory, and partly in the trust policy file. You can edit the Web.config file directly and push it out to different servers, or you can use the ADFS snap-in to modify the settings.

The trust policy file should not be edited manually. Instead, edit the trust policy file by using the ADFS snap-in, or edit it programmatically by using the ADFS object model.

| **Note** |
| --- |
| Scripting support is provided in the ADFS object model. |

When you open the ADFS snap-in, the snap-in reads the Web.config file from the Federation Service virtual directory and notes the location of the trust policy file. The snap-in then presents a console tree hierarchy representing the Federation Service and all aspects of the trust policy, including organization claims, partners, account stores, and applications. Each item in this console tree hierarchy has context-menu options that you can use to view, modify, add, and delete trust policy entities.

**Federation Service node**

The **Federation Service** node in the console tree of the ADFS snap-in represents the local Federation Service that is assigned to the federation server on which you are viewing the snap-in. You control the local federation server configuration through this node in the ADFS snap-in. The local federation server configuration is different from the trust policy configuration in that the trust policy configuration is shared among all the federation servers in the federation server farm. The local configuration is stored in the Web.config file, and it includes the following items:

- The friendly name for the Federation Service
- The path to the trust policy file
- The local certificate to be used for signing tokens
- The Microsoft ASP.NET Web pages
- The debug logging level

# Active Directory Federation Services Proxy snap-in

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The Active Directory Federation Services (ADFS) Proxy Microsoft Management Console (MMC) snap-in is installed when you install the Federation Service Proxy component in **Add or Remove Programs** in Windows Server 2003 R2, Enterprise Edition. You can use the ADFS Proxy snap-in to:

- Configure the Federation Service that this Federation Service Proxy is acting as a proxy for.
- Determine how to collect user credential information from browser clients and Web applications.

The settings in the ADFS Proxy snap-in are stored in the Web.config file in the Federation Service Proxy virtual directory.

**Federation Service Proxy node**

The **Federation Service Proxy** node in the snap-in's console tree hierarchy represents the current Federation Service Proxy settings. You control the local federation server proxy configuration through this node in the ADFS Proxy snap-in. The actual client authentication certificate with the private key is present in the local computer personal certificate store.

ADFS distinguishes between the local configuration for a federation server proxy and the trust policy configuration that is shared among all servers in the federation server farm. The local proxy configuration is stored in the Web.config file, and it includes the following items:

- The Federation Service Uniform Resource Locator (URL)
- The client authentication certificate to be used by the federation server proxy for Transport Layer Security and Secure Sockets Layer (TLS/SSL) communication with the Federation Service
- Microsoft ASP.NET Web pages

# Understanding ADFS Authentication and Authorization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Understanding ADFS Authentication and Authorization**


- Account stores
- Cookies used by ADFS
- Controlling Access to Web-based Applications
- Understanding Claims

# Account stores

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) uses account stores to log on users and extract security claims for those users. You can configure multiple account stores for a single Federation Service and define their priority. The Federation Service uses Lightweight Directory Access Protocol (LDAP) to communicate with account stores.

ADFS supports two types of account stores: Active Directory account stores and Active Directory Application Mode (ADAM) account stores.

**Active Directory account stores**
ADFS is tightly integrated with Active Directory. ADFS retrieves user attributes and authenticates users against Active Directory. ADFS also uses Windows Integrated Authentication and security tokens that Active Directory creates.

For a user to log on to Active Directory, a user name must be in the user principal name (UPN) format (user@adatum.com) or in the Security Accounts Manager (SAM) account name format (adatum\user). After the user is logged on and impersonated, user security IDs (SIDs) are enumerated from the access token. The SIDs are then mapped to organization group Claims.

E-mail claims, common name claims, and custom claims can be extracted from user object attributes that are defined in Active Directory when the Federation Service account is used to perform an LDAP search of an object.

The Federation Service account must have access to the user object. If the user object resides in a domain different from the domain where the Federation Service account resides, the former domain must have an Active Directory domain trust in place to the latter domain.

There is no direct way of determining whether any given user name exists in Active Directory and in all directories that are trusted (either directly or transitively) by Active Directory. Active Directory returns an authoritative failure only if the logon attempt fails as a result of policy restrictions. Examples of policy restriction failures include the following:

- The account is disabled.
- The account password has expired.
- The account is not allowed to log on to this computer.
- The account has logon time restrictions and is not allowed to log on at this time.

Otherwise, Active Directory account store logon failures are always nonauthoritative, and the next-priority account store is tried. For more information about account store logon failures, see ADFS troubleshooting.

The following is the list of configuration items for the Active Directory account store:

- A setting indicating whether this entry is enabled
- The display name for the trusted Active Directory account store
- Optional: A list of SIDs that map to particular organization group claims
- Optional: Configuration defining how to extract organization claims from user object attributes in Active Directory:
    - E-mail identity claim or common name identity claim
    - Custom organization claim

The name of a user object attribute defines the claim name when a custom claim is built.

**ADAM account stores**

ADAM provides data storage and retrieval for directory-enabled applications, without the dependencies that are required for Active Directory. ADAM provides much of the same functionality as Active Directory, but it does not require the deployment of domains or domain controllers. Similar to the way in which ADFS uses Active Directory account store information, ADFS also retrieves user attributes and authenticates users with ADAM.

| **Note** |
| --- |
| ADFS cannot authenticate ADAM accounts that use parentheses as part of the account name. Accounts that have an open parenthesis in the user name will cause an LDAP search failure as a result of the user name forming an invalid LDAP filter. |

Security claims are obtained when the Federation Service account is used to perform an LDAP search of the object. This is a two-step process:

- First, the user object is found by means of a search for the object whose configured attribute is equal to the supplied user name. The Federation Service account uses Kerberos or NTLM encryption to protect this communication.

  | **Note** |
  | --- |
  | This process requires the ADAM server to be joined to a domain that trusts the domain that the Federation Service is a member of. |

- Next, the user credentials are validated by binding to the found user object with the supplied password using LDAP. If Transport Layer Security and Secure Sockets Layer (TLS/SSL) are configured for the ADAM account store properties in the trust policy, the user credentials are protected.

  | **Important** |
  | --- |
  | It is strongly recommended that the traffic between the ADAM server and the federation server be protected by TLS/SSL or other means, such as Internet Protocol security (IPsec). |

If more than one object is returned from the LDAP query with the supplied user name, this is considered an authentication failure. For more information about authentication failures, see ADFS troubleshooting.

Because the presence of the user object is always determined, if the user account exists in the ADAM store, the ADAM store logon is authoritative. Therefore, no additional account stores are attempted. The following is the list of configuration items for the ADAM store:

- A setting indicating whether this entry is enabled
- The display name for the trusted ADAM account store
- The ADAM server name
- An LDAP user object attribute containing the user name
- Optional: Port number. If the port number is not specified, it defaults to 389.
- Optional: Search base. If a subtree is specified, searches are performed on the specified subtree. Otherwise, the entire directory tree is searched.
- Optional: Client timeout: the maximum time that Federation Service waits for a response from the ADAM server before timing out the connection. The default is five seconds.
- Optional: Configuration defining how to extract organization claims:
  - UPN/e-mail/common name identity claim: The name of a defined user object attribute whose value is extracted as the UPN/e-mail/common name claim value

- Group organization claim: The name of a user object attribute and its value that, if present, results in the particular group claim

- Custom organization claim: The name of a defined user object attribute whose value is extracted as the particular custom claim value

If this configuration is absent, the directory is used only to log on the user — without extracting claims.

**Determining the priority order of user logon requests**
When a logon request is made to either Active Directory or ADAM through an ADFS client, it is passed immediately to the specified account store. However, if the account store Uniform Resource Identifier (URI) is not specified, each of the stores is tried in priority order to log on the user. The authentication result is returned if:

- There is only one store configured, and credential verification information is returned.

- The store URI was specified in the logon request, and credential verification information is returned.

- The result of authentication by one of the stores was authoritative.

- Authentication by one of the stores was successful.

**Disabling account stores**
Each account store can be marked as enabled or disabled. If an account store is disabled, it does not participate in any account-store-related operations. Cookies with claims that originate from an account store that is currently disabled are discarded or deleted, and the client is directed to the logon page.

# Cookies used by ADFS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) uses the following three types of cookies:

- Authentication cookie
- Account partner cookie
- Sign-out cookie

**Authentication cookie**

Authentication cookies can be issued by both the Federation Service and the ADFS Web Agent. The ADFS Web Agent takes the ADFS security token that it receives and uses that token as the cookie value. The benefit at the Web server is that it does not need keying material. The Federation Service publishes all the information that is necessary to validate its tokens.

At the Federation Service, the security token in a cookie holds the organization Claims for the client. The organization claims may be mapped to outgoing claims for a particular resource. The ADFS Web Agent can also authenticate and use cookies that are issued by the Federation Service. The Web server receives a cookie when the client comes to the Web server. Then, the ADFS Web Agent can authenticate this cookie and use the claims that it contains. For more information about how the Federation Service uses tokens, claims, and authentication cookies, see Federation Service.

The authentication cookie facilitates single sign-on (SSO). After the Federation Service validates the client once, the authentication cookie is written to the client. The Federation Service produces and consumes the contents of the authentication cookie, and they are opaque to federation server proxies. Further authentication takes place through use of the cookie rather than through repeated collection of the client credentials. For more information about the federation server proxies, see ADFS server roles.

The following illustration shows the contents of an authentication cookie and the ADFS components that use the authentication cookie. The ADFS Web Agent comprises both the ADFS Web Agent Authentication Service and the ADFS Web Agent ISAPI Extension.




The authentication cookie is always a session cookie. The authentication cookie is signed but not encrypted, which is one reason why use of Transport Layer Security and Secure Sockets Layer (TLS/SSL) in ADFS is mandatory.

**Account partner cookie**

The account partner cookie facilitates SSO. After interactive account partner membership discovery occurs, if the account partner cookie has a valid token, the cookie is written to the client. Further interactions use the information in this cookie rather

than prompting the client for account partner membership information again. The account partner cookie is set as a result of the account partner discovery process. For more information about account partner discovery, see Federation Service.

The account partner cookie is a long-lived, persistent cookie. It is neither signed nor encrypted.

**Sign-out cookie**

The sign-out cookie facilitates sign-off. Whenever the Federation Service issues a token, the token's resource partner or target server is added to the sign-out cookie. When it receives a sign-off request, the Federation Service or Federation Service Proxy sends requests to each of the token target servers asking them to clean up any authentication artifacts, such as cached cookies, that the resource partner or Web server may have written to the client. In the case of a resource partner, it sends a cleanup request to any application Web servers that the client has used.

The sign-out cookie is always a session cookie. It is neither signed nor encrypted.

# Controlling Access to Web-based Applications

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Controlling Access to Web-based Applications**


- Claims-aware applications
- Windows NT token-based applications

# Claims-aware applications

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Claims are statements (for example, name, identity, key, group, privilege, or capability) made about users — and understood by both partners in an Active Directory Federation Services (ADFS) federation — that are used for authorization purposes in an application. A claims-aware application is a Microsoft ASP.NET application that has been written using the ADFS library. This type of application is fully capable of using ADFS Claims to make authorization decisions directly. A claims-aware application accepts claims that the Federation Service sends in ADFS security tokens. For more information about how the Federation Service uses security tokens and claims, see Federation Service.

Claim mapping is the act of mapping, removing or filtering, or passing inbound claims into outbound claims. Claim mapping does not occur when claims are sent to an application. Instead, only the organization claims that are specified by the Federation Service administrator are sent to the application. (Organization claims are claims in intermediate or normalized form within an organization's namespace.)

By default, no organization claims that are marked as auditable are sent. All nonauditable claims are sent. Therefore, the Federation Service administrator will opt in for auditable claims to be sent to the application and will opt out for nonauditable claims.

The following list describes the organization claims that can be used by claims-aware applications:

- Identity claims (UPN/e-mail/common name)

  When you configure the application, you specify which of these identity claims will be sent to the application. No mapping or filtering is performed.

- Group claims

  When you configure the application, you specify the organization group claims that will be sent to the application. Organization group claims that are not designated to be sent to the application will be discarded.

- Custom claims

  When you configure the application, you specify the organization custom claims that will be sent to the application. Organization custom claims that are not designated to be sent to the application will be discarded.

**Claims-aware authorization**

Claims-aware authorization consists of a Hypertext Transfer Protocol (HTTP) module and objects for querying the claims that are carried in the ADFS security token. Claims-aware authorization is supported only for Microsoft ASP.NET applications.

The HTTP module processes ADFS protocol messages based on configuration settings in the Web application's Web.config file. The Web pages perform authentication and authorization tasks. The HTTP module also authenticates cookies and obtains claims from the cookies.

# Windows NT token-based applications

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A Windows NT token–based application is an Internet Information Services (IIS) application that has been written to use traditional Windows native authorization mechanisms. This type of application is not prepared to consume Active Directory Federation Services (ADFS) Claims.

Windows NT token–based applications may be used only by Windows users from the local realm or any realm that is trusted by the local realm — that is, only by users who can log on to the computer with Windows NT token–based authentication techniques.

| **Note** |
| --- |
| In federation scenarios, this means that resource accounts may be required. |

The ADFS security token that is sent to Windows NT token–based applications contains an e-mail claim for the user. It may also contain custom claims with the security IDs (SIDs) of the user account. The Web server generates a Windows impersonation-level access token. An impersonation-level access token captures the security information for a client process, thereby allowing a service to "impersonate" the client process in security operations.

For Windows NT token–based Web applications, the following process order determines how a Windows NT token is created:

1. If the Security Assertion Markup Language (SAML) token contains SIDs in the SAML advice element, the SIDs are used to generate the Windows NT token.

2. If the SAML token does not contain SIDs and instead contains a user principal name (UPN) identity claim, the UPN claim is used to generate the Windows NT token.

3. If the SAML token does not contain SIDs and the UPN identity claim in the e-mail identity claim is present, it is interpreted as a UPN and it is used to generate the Windows NT token.

This behavior is irrespective of whether the UPN or e-mail identity claim is specified as the identity claim that is used to generate the Windows NT token when the trust policy entry for the Web application is created in the Federation Service.

**Traditional Windows-based authorization**

Support for converting an ADFS security token into an impersonation-level Windows NT access token requires a number of components:

- Internet Server Application Programming Interface (ISAPI) extension: This component checks for ADFS cookies, checks for ADFS security tokens from the Federation Service, performs the appropriate protocol redirects, and writes the necessary Cookies used by ADFS to make ADFS work.

- ADFS authentication package: The ADFS authentication package generates an impersonation-level access token, given a UPN for a domain account. The package requires that the caller have the Trusted Computing Base (TCB) privilege.

- **ADFS Web Agent** tabs in IIS Manager: You can use the ADFS Web Agent tabs to administer policy and certificates for verifying the ADFS security token and cookies.

- ADFS Web Agent Authentication Service: The ADFS Web Agent Authentication Service runs as Local System to generate a token by either using Service-for-User (S4U) or using the ADFS authentication package. However, the IIS application pool is not required to run as Local System. The ADFS Web Agent Authentication Service has interfaces that may be called only with local remote procedure call (LRPC), not remote procedure call (RPC). The service returns an impersonation-level Windows NT access token if it is given an ADFS security token or ADFS cookie.

- ADFS Web Agent ISAPI Filter: Certain traditional IIS Web applications use an ISAPI filter that may modify incoming data, such as Uniform Resource Locators (URLs). If this is the case, the ADFS Web Agent ISAPI Filter must be enabled and configured as the highest priority filter. This filter is not enabled by default.

# Understanding Claims

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Understanding Claims**


- Claims
- Claim mapping

# Claims

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Claims are statements (for example, name, identity, key, group, privilege, or capability) made about users — and understood by both partners in an Active Directory Federation Service (ADFS) federation — that are used for authorization purposes in an application.

The Federation Service brokers trust between many disparate entities. It is designed to allow the trusted exchange of arbitrary claims containing arbitrary values. These claims are then used by the receiving party to make authorization decisions.

There are three ways that claims flow through the Federation Service:

- From the account store to the account Federation Service to the resource partner
- From the account partner to the resource Federation Service to the application resource
- From the account store to a Federation Service to the application resource

The Federation Service can be configured to act in all three of these roles. Therefore, one single Federation Service may facilitate all three communication flows.

There are three types of claims that are supported by the Federation Service: identity claims, group claims, and custom claims. The following table describes each of these claim types in more detail.

| **Claim type** | **Description** |
|---|---|
| Identity | UPN, e-mail, and common name in ADFS are referred to as identity claim types.<br>• UPN — Indicates a Kerberos-style user principal name (UPN), for example: user@realm. Only one claim may be the UPN type. Even if multiple UPN values must be communicated, only one may be of the UPN type. Additional UPNs may be configured as custom claim types.<br>• E-mail — Indicates RFC 2822–style e-mail names of the form user@domain. Only one claim may be the e-mail type. Even if multiple e-mail values must be communicated, only one may be of e-mail type. Additional e-mails may be configured as custom claim types.<br>• Common name — Indicates an arbitrary string that is used for personalization. Examples include John Smith or Tailspin Toys Employee. Only one claim may have the common name type. It is important to note that there is no mechanism for guaranteeing the uniqueness of the common name claim. Therefore, use caution when using this claim type for authorization decisions. |
| Group | Indicates membership in a group or role. Administrators define individual claims that have the group type “Group claims.” For example, you might define the following set of group claims: [Developer, Tester, Program Manager]. Each group claim is a separate unit of administration for claim population and mapping. It is useful to think of the value of a group claim as a Boolean value indicating membership. |
| Custom | Indicates a claim that contains custom information about a user, for example, an employee ID number. |

If more than one of the three identity claim types is present in a token, the identity claims are prioritized in the following order:

1. UPN
2. E-mail

3. Common name

At least one of these identity claim types must be present for a token to be issued.

**Auditing claims**
Some group claims and custom claims may be designated as auditable. When auditing is enabled, the audit allows the name of the claim to be exposed in the security event log, but the value of the claim is omitted. An example of an auditable claim is Social Security Number. The claim name Social Security Number is exposed, but the actual number value that is stored in that claim is not exposed. The claim value is not audited when the claim is produced or mapped.

| **Note** |
| --- |
| Identity claim types are always auditable. |

**Claim producers and consumers**
The way claims are used depends on the claim producer or consumer. Claims are either inbound or outbound. ADFS supports the following claim producers and consumers:

- Active Directory account stores
- ADAM account stores
- Account partners
- Resource partners
- Claims-aware applications
- Windows NT token–based applications

**Active Directory account store**
The Active Directory account store is a claim producer that represents authentication for the Federation Service. Specifically, the Federation Service may log on users from its domain, from domains that are directly trusted by its domain, from domains in the same forest as its domain, and from domains in forests that have forest trusts with the domain's forest.

The Active Directory account store is available only if the Federation Service is joined to a domain.

- UPN claim: When you configure the Active Directory account store, the UPN claim is enabled automatically.
- E-mail claim: When you configure the Active Directory account store, you can specify what Lightweight Directory Access Protocol (LDAP) user attributes, if any, contain the user's e-mail address.
- Common name claim: When you configure the Active Directory account store, you can specify what LDAP user attributes, if any, contain the user's common name.
- Group claims: you can assign Windows users and groups directly to the organization group claims using the object picker.
- Custom claims: When you configure the Active Directory account store , you can specify what LDAP user attributes contain claim values and then assign each attribute name to an organization custom claim.

**ADAM account store**
The ADAM account store is a claim producer that represents authentication for the Federation Service.

- UPN claim: When you configure the ADAM account store, you can specify the LDAP user attribute, if any, containing the user's UPN.

- E-mail claim: When you configure the ADAM account store, you can specify the LDAP user attribute, if any, containing the user's e-mail address.

- Common name claim: When you configure the ADAM account store, you can specify the LDAP user attribute, if any, containing the user's common name.

You must assign at least one identity claim type to the ADAM account store for the Federation Service to allow that store to be enabled.

- Group claims: When you configure the ADAM account store, you can specify the LDAP user attribute containing the user's LDAP groups or any other attribute that could function as a group, such as Title (if groups are based on job role), and then assign each possible LDAP group to an organization group.

- Custom claims: When you configure the ADAM account store, you can specify the LDAP user attributes containing claim values. You then assign each attribute name to an organization custom claim.

# Claim mapping

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Active Directory Federation Services (ADFS) uses the WS-Federation Passive Requestor Protocol (WS-F PRP) that carries Claims in security tokens that are issued by the Federation Service. The claims are populated initially from Account stores, either Active Directory account stores or Active Directory Application Mode (ADAM) account stores.

In addition, the Federation Service can map those claims when they go out to a federation partner or when they come in from a federation partner. Claim mapping is the act of mapping, removing or filtering, or passing inbound claims into outbound claims. Claim mapping may be different for each federation partner. Defining the population and mapping of these claims is important for the configuration of the federation. Claim mappings use string comparisons that are case sensitive. The following illustration shows the claim mapping process.




**Organization claim sets**
All incoming claims are mapped into organization claims. Organization claims are claims in intermediate or normalized form within an organization's namespace. All internal Federation Service actions are performed on the organization claim set. Organizational claims are consumed by resource applications.

With organization claims, mappings do not have to be administered individually between any two organizations that need to communicate. Each organization defines a single mapping either to or from its organization claims. This reduces the administrative complexity of ADFS. For example, if the federation has:

- **x** account partners
- **y** resource applications

The federation has **x** + **y** claim mappings. This is reduced from a potential **x** × **y** claim mappings. As a concrete example, if a Federation Service has:

- 3 account partners
- 7 resource applications

The federation needs only 10 claim mappings, as compared to 21 claim mappings, when mapping occurs directly from the incoming claims to the outgoing claims.

**E-mail**
E-mail claim types always map to e-mail claim types. As part of this mapping, on the account Federation Service, the domain suffix may be mapped to a constant value. On the resource Federation Service, the domain suffix may be filtered against a list of constant values.

The following example describes an ADFS federation between two organizations, Tailspin Toys and Adventure Works. In this example, Tailspin Toys is the account partner, and Adventure Works is the resource partner.

- Tailspin Toys, acting as an account Federation Service, maps the e-mail organization claim to the outgoing e-mail claim for Adventure Works. As part of that mapping, it maps all e-mail suffixes to tailspintoys.com. Given the organization e-mail claim (e-mail=jsmith@sales.tailspintoys.com), the outgoing e-mail claim is (e-mail=jsmith@tailspintoys.com).

- Adventure Works, acting as a resource Federation Service, maps the incoming Tailspin Toys e-mail claim to the e-mail organization claim, and, as part of that mapping, it filters the suffix list against tailspintoys.com. Therefore, an incoming Tailspin Toys e-mail claim (email=jsmith@tailspintoys.com) is accepted, but an incoming Tailspin Toys e-mail claim (e-mail=jsmith@adventure-works.com) is rejected.

**UPN**
User principal name (UPN) claim types always map to UPN claim types. They are subject to suffix mappings and filtering in the same way that e-mail claims are. However, because Active Directory allows UPNs without the @ symbol, the account Federation Service attaches the @ symbol, followed by the suffix if there is a UPN suffix mapping defined. Otherwise, if any suffix is passed through, the Federation Service passes the UPN as is, without the @ symbol. On the resource side, if any UPN suffix is allowed, the UPN without the @ symbol is accepted. Otherwise, if a specific UPN suffix is allowed, the UPN without the @ symbol is rejected.

**Common name**
Common name claim types always map to common name claim types. They are subject to no additional rules.

**Custom**
Custom claim types always map to other custom claim types. For example, given an incoming claim set of (UPN, Custom=[EmployeeNumber, TaxPayerID]) and an organization claim set of (UPN, Custom=[Employee, SSN]), you can create mappings from EmployeeNumber to Employee and from TaxPayerID to SSN.

**Group**
Group claim types always map to other group claim types. For example, given an incoming claim set of (UPN, Group=[One, Two, Three]) and an organization claim set of (UPN, Group=[X,Y,Z]), you can create mappings from One to Y, from Two to X, and from Three to Z.

**Group-to-UPN mapping**
In addition to the standard mappings described in the previous sections, you may also use a special group-to-UPN claim mapping. The group-to-UPN claim mapping is supported only on the resource Federation Service when claims are incoming from an account partner. In this case, UPN claim types do not map to UPN claim types. Instead, you provide an ordered list of group-to-UPN claim mappings.

For example, the group-to-UPN list might be:

1. Dev to developers@internal.tailspintoys.com

2. Test to testers@internal.tailspintoys.com

3. PM to progmgrs@internal.tailspintoys.com

Given an incoming claim set of (Common name=John Smith, Group=[Dev]) the organization claim set contains (Common name=John Smith, UPN=developers@internal.tailspintoys.com). Remember that the list is ordered. Therefore, a claim set of (Common name=John Smith, Group=[Dev,PM]) results in (Common name=John Smith, UPN=developers@internal.tailspintoys.com). Also, if the incoming claim has a UPN, the UPN is overwritten. This special mapping rule specifically supports group-based resource accounts that access legacy resources. The order of the group-to-UPN mappings is specified in the trust policy for the Federation Service.

# Understanding Certificates Used by ADFS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Understanding Certificates Used by ADFS**


- Certificates used by federation servers
- Certificates used by federation server proxies
- Certificates used by Web servers

# Certificates used by federation servers

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Token-signing certificates**

Each federation server uses a token-signing certificate to digitally sign all security tokens that it produces. Because each security token is digitally signed by the account partner, the resource partner can verify that the security token was in fact issued by the account partner and that it was not modified. This helps prevent attackers from forging or modifying security tokens to gain unauthorized access to resources.

Digital signatures on security tokens are also used within the account partner when more than one federation server is used. In this situation the digital signatures verify the origin and integrity of security tokens that are issued by other federation servers within the account partner. The digital signatures are verified with verification certificates.

| **Note** |
| --- |
| Each token-signing certificate contains a private key that is associated with the certificate. |

**Verification certificates**

Verification certificates are used to verify that a security token was issued by a valid federation server and that it was not modified. Verification certificates are actually the token-signing certificates of other federation servers.

To verify that a security token was issued by a given federation server and not modified, the federation server must have a verification certificate for the federation server that issued the security token. For example, if federation server A issues a security token and sends the security token to federation server B, federation server B must have a verification certificate (federation server A's token-signing certificate) for federation server A.

| **Note** |
| --- |
| Unlike a token-signing certificate, a verification certificate does not contain the private key that is associated with the certificate. |

**SSL server authentication certificates**

The federation server uses Secure Sockets Layer (SSL) server authentication certificates to secure Web services traffic for communication with Web clients or the federation server proxy. These certificates are requested and installed through the Internet Information Services (IIS) snap-in.

For more information about certificates, see Public Key Infrastructure for Windows Server 2003 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=19936).

# Certificates used by federation server proxies

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**SSL client authentication certificates**
Each federation server proxy uses a Secure Sockets Layer (SSL) client authentication certificate to authenticate to the Federation Service. Any certificate with client authentication extended key usage (EKU) can be used as a federation server proxy client authentication certificate. A copy of the federation server proxy client authentication certificate is stored on both the federation server proxy and in the trust policy of the federation server. However, only the federation server proxy stores the private key that is associated with the federation server proxy client authentication certificate.

**SSL server authentication certificates**
The federation server proxy uses SSL server authentication certificates to secure Web services traffic for communication with Web clients. These certificates are requested and installed through the Internet Information Services (IIS) snap-in.

For more information about certificates, see Public Key Infrastructure for Windows Server 2003 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=19936).

# Certificates used by Web servers

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Each Web server that hosts the ADFS Web Agent uses Secure Sockets Layer (SSL) server authentication certificates to securely communicate with Web clients. These certificates are requested and installed through the Internet Information Services (IIS) snap-in.

For more information about certificates, see Public Key Infrastructure for Windows Server 2003 on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=19936).

# ADFS resources

Updated: September 13, 2007

Applies To: Windows Server 2003 R2

You can find detailed documentation about Active Directory Federation Services (ADFS) on the ADFS home page (http://go.microsoft.com/fwlink/?LinkId=79542).

For more information about ADFS or other ADFS-related information, see the following Web resources:

- Web Services and Other Distributed Technologies Developer Center on the MSDN Web site (http://go.microsoft.com/fwlink/?LinkId=44189)
- Web Services Protocol Workshops on the MSDN Web site (http://go.microsoft.com/fwlink/?LinkId=44190)
- Web Services Specifications on the MSDN Web site (http://go.microsoft.com/fwlink/?LinkId=44191)
- Web Services Interoperability Organization (WS-I) (http://go.microsoft.com/fwlink/?LinkId=34328)

# ADFS troubleshooting

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Setup issues**
**You receive an Internet Explorer error page with the message "This page cannot be displayed" "Cannot find server or DNS Error."**
There are a few things that can cause this problem:

- Verify that Windows Firewall is turned off.
- Verify that all the Federation Services and application Web servers have a server authentication certificate issued to the default Web site.
- If there is an external account partner Federation Service Proxy involved, verify that the correct Federation Service host name was used during installation.
- If you are using a Windows NT token-based application, verify that the Federation Service Uniform Resource Locator (URL) on the **ADFS Web Agent** tab of **Web Sites Properties** is configured correctly.

**When you try to connect to the application, you get an Internet Explorer error page with the message "This page cannot be found" "HTTP Error 404 – File or directory not found."**
This issue might be caused by the following configuration problems:

- Verify that the Web application is properly configured in Internet Information Services (IIS).
- Verify that the Web application URL is properly named in the Active Directory Federation Services Snap-in.
- Verify that Microsoft ASP.NET is installed on the application Web server and in the Federation Service.
- If you are connecting to a Windows NT token–based application that uses ASP and you receive the 404 error after supplying your credentials, in the IIS Manager console tree, click the local computer, click **Web Service Extensions**, and then verify that **Active Server Pages** is set to **Allowed**.

**After setting up a Windows NT token–based application, you attempt to connect to it but you are not prompted to choose a host realm and login credentials.**
Verify that the virtual directory of the Windows NT token–based application is set up to use the lfsext.dll Internet Server Application Programming Interface (ISAPI) extension.

**Where are the ADFSSetup.log, ADFSOCM.log, and ADFSMSI.log files located?**
They are located in %SYSTEMROOT%\.

**Enable logging**
**Registry settings to enable logging on the account federation server**
The account federation server uses an authentication package for mapping client certificates. To enable logging for the account federation server authentication package, perform the following tasks in order:

1. If it is not already installed, install the Federation Service component of Active Directory Federation Services (ADFS).
2. Set the following registry key:
   **[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\WebSso\Parameters]**

   **"DebugLevel"=dword:ffffffff**

**Registry settings to enable logging on the Web server for the:**
**ADFS Web Agent Authentication Package**
The ADFS Web Agent authentication package is used by Windows NT token-based applications for generating tokens when Service-for-User (S4U) is not available.

**[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\WebSso\Parameters]**

**"DebugLevel"=dword:ffffffff**

**ADFS Web Agent ISAPI Extension**
The ADFS Web Agent ISAPI Extension handles the protocols that are used by ADFS to authenticate requests.

**[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\ADFS\WebServerAgent]**

**"DebugPrintLevel"=dword:ffffffff**

**ADFS Web Agent Authentication Service**
The ADFS Web Agent Authentication Service validates incoming tokens and cookies.

**[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IFSSVC\Parameters]**

**"DebugPrintLevel"=dword:ffffffff**

**Where are these logs located?**
They are located in %SYSTEMDRIVE%\ADFS\logs.

**Configuration issues**
The following section covers some of the known issues with ADFS configuration.

**Condition: server error**
**Error:** The token request for application with URL https://... cannot be fulfilled because the URL does not identify any known trusting application

**Solution:** This error is returned by the resource Federation Service when the application URL does not identify any known application. Make sure that the application has been added to the trust policy for the Federation Service. For more information about how to do this, see Complete the Add Applications Wizard.

For a claims-aware application, verify that the return URL is typed correctly in the application's web.config file and that it matches the application URL that is specified in the trust policy of the Federation Service.

For a Windows NT token-based application, verify that the return URL is typed correctly on the **ADFS Web Agent** tab of IIS and that it matches the application URL in the trust policy of the Federation Service.

**Condition: validation error**
**Error:** Validation of viewstate MAC failed. If this application is hosted by a Web farm or cluster, ensure that **<**machineKey**>** configuration specifies the same validationKey and validation algorithm.

**AutoGenerate** cannot be used in a cluster. An unhandled exception occurred during the execution of the current Web request. Review the stack trace for more information about the error and where it originated in the code.

Or

**Error:** An unhandled exception was generated during the execution of the current web request. Information regarding the origin and location of the exception can be identified using the exception stack trace below.

**Solution:** Using a text editor, add the following setting to the Web.config file on the computer hosting either the Federation Service, Federation Service Proxy, or ADFS Web Agent that will be farmed:

```
<system.web>

<machineKey>

<machineKey validationKey="specify key for the appropriate algorithm"

decryptionKey="specify key"

validation="SHA1|MD5|3DES"/>
```

Or

**Solution:** Add the following element in the <system.web> section of the Web.config file on the computers hosting the Federation Service, Federation Service Proxy, or ADFS Web Agent that are set up in the farm:

```
<pages enableViewStateMac="false"/>
```

**Condition: The account Federation Service is configured with the wrong cookie path.**

Behavior: The Federation Service tries to authenticate the user each time, and no sign-on occurs. The event logs on the federation server show that the user has been authenticated successfully. However, the user never sees the login succeed.

**Solution:** To resolve this issue, correct the cookie path that is specified in the trust policy.

The cookie path must match a prefix of the request Uniform Resource Identifier (URI). It is important to note that some browsers treat "/path" or "/path1/samp" as a prefix match of "/path1/sample" while others do not allow matches that consume only parts of the individual words. These strict implementations accept only a subset of those matches that are allowed by the first implementation, for example, "/path1" or "/path1/sample".

**Condition: A Windows NT token–based application cannot create a Windows access token for the user using the security token from the resource Federation Service as a result of one of the following conditions:**

- The only identity claim is the common name claim.
- Services-for-User (S4U) fails because there is no resource account corresponding to the user principal name (UPN).

**Behavior:** The Web application returns an anonymous user token.

**Condition: The UPN name for forms authentication with the user@domain NetBIOS name can be different from the UPN for integrated authentication and Transport Layer Security / Secure Sockets Layer (TLS/SSL) client authentication.**
**Behavior:** In the case of integrated and TLS/SSL client authentication, the UPN is always the explicit UPN that is configured in Active Directory for the user object. However, for forms authentication, the UPN can be whatever the user enters in the user name field. Therefore, the UPN can be in the form:

username@netbiosname of domain or username@domaindnsname

**Condition: The Web server application should not be configured with any URLs that are not Secure Hypertext Transfer Protocol (HTTPS).**
**Behavior:** If the application contains links that are not HTTPS, the browser will not send the browser authentication cookies for those URLs. Single sign-on (SSO) is not possible because the ADFS cookies are secure cookies. In cases where the Web server is hosting only Windows NT token–based applications, the Web server will redirect to the HTTPS URL. Therefore, this behavior may not be noticed in most cases.

**Condition: If the Federation Service or Web server is running on a domain controller, APPIDENTITY should be provided with the Write permission to the Temporary ASP.NET files directory.**
(%windir%\Microsoft.NET\Framework\v1.1.4322\Temporary ASP.NET files)

APPIDENTITY refers to the identity of the application pool with which the Federation Service or Federation Service Proxy is configured to run.

**Condition: The Federation Service log file gets filled, and the Federation Service starts to fail all requests.**
**Solution:** Increase the log size of the Federation Service log to something large, such as 16 megabytes (MB).

**Condition: A forest trust is created between an extranet forest and an intranet forest, and a user has successfully authenticated as an intranet domain user to the extranet forest. However, the user cannot authenticate as a user from a domain that is a child domain of the intranet forest.**
**Cause:** The Federation Service maintains a cache of forest trust information that is refreshed every eight hours.

**Solution:** Recycle the Federation Service application pool. This ensures that no authentication requests are dropped.

**Duplicate group or custom claim names**
As a best practice, do not create duplicate group or custom claims using case-insensitive names. For example, do not create an organization group claim using the names “Smith” and “smith”. If similar claim names are typed but they have case-sensitive differences, ADFS will treat those claims differently.

**Active Directory node in the ADFS snap-in:**
**Extracting multiple Active Directory users to a single organizational group claim displays ‘*’ under the account store data field in the details pane.**
**Solution:** This is by design. To verify the list of added Active Directory users, open the properties dialog box for the organizational group claim.

**Account Partners node in the ADFS snap-in:**
**Object picker under the Groups tab in the UPN claim mapping dialog box returns a selected user’s UPN as NULL in the UPN edit field.**
**Solution:** This is expected behavior. If the selected user is not assigned with a UPN explicitly, the object picker returns NULL, and the UPN field shows a blank value. Type the UPN manually in the field.

**Applications node in the ADFS snap-in:**
**A common name claim for a Windows NT token-based application cannot be enabled in the Add Applications Wizard.**
**Solution:** The user cannot enable a common name identity claim for a Windows NT token–based application in the Add Applications Wizard. This can be fixed in the properties dialog box for the common name identity claim in the Added Applications details pane.

**ADAM issues**
**After user accounts are created in Active Directory Application Mode (ADAM) and the trust policy is configured with information about the ADAM store, the Federation Service is not able to validate users in the ADAM store.**
**Solution**: Use caution while creating user accounts using the ADAM ADSI Edit snap-in. Always create a user account with a password. If you create a user account without a password, use ADSI Edit to reset the password for the user account. Most importantly, check the value of the **msDS-UserAccountDisabled** property of the user account. This property should not have the value **True**. The value should be either **False** or **Not set**. If the value of **msDS-UserAccountDisabled** property is **True**, it means that the user account is disabled and the Federation Service cannot validate credentials for this ADAM user account.

**You have enabled an ADAM account store, but the Federation Service is not able to retrieve any claims.**
If the Federation Service is running as Local System, you must add the machine account of the computer hosting the Federation Service to the Readers group in the ADAM store.

If the Federation Service is running as Network Service, you must add the domain account to the Readers group in the ADAM store.

# Active Directory Application Mode (ADAM)

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Active Directory®Application Mode**

- ADAM How To...
- ADAM Concepts
- Administering ADAM
- ADAM troubleshooting and frequently asked questions (FAQs)
- ADAM Administration Tools
- Glossary

# ADAM How To...

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**How to ...**

- Install or Remove ADAM
- Create or Remove an ADAM Instance
- Manage an ADAM Instance
- Manage Replication, Sites, and Configuration Sets
- Manage Schemas and Directory Partitions
- Manage Users and Groups
- Manage Access Control
- Monitor and Maintain ADAM
- Synchronize Data from Active Directory to an ADAM Instance

# Install or Remove ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Install or remove ADAM**

- Install ADAM
- Remove ADAM

# Install ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To install ADAM**

1. Open **Add or Remove Programs** from **Control Panel**.
2. Click **Add/Remove Windows Components**.
3. Click **Active Directory Services** under **Components**, and then click **Details**.
4. Select the **Active Directory Application Mode (ADAM**) check box under **Subcomponents of Active Directory Services**, and then click **OK**.
5. Click **Next** on the **Windows Components** page.
6. Review the message that appears. Based on the contents of message, do one of the following:
   - Click **Finish** if the message, "You have successfully completed the ADAM Setup Wizard" appears.
   - If an error message appears, make a note of the error, click **Finish**, and then review the Active Directory Application Mode (ADAM) event messages in **Event Viewer**.

## Notes

- To open **Add or Remove Programs**, click **Start**, click **Control Panel**, and then click **Add or Remove Programs**.
- If an error occurs in the Active Directory Application Mode Setup Wizard before the **Completing the Active Directory Application Mode Setup Wizard** page, first review the error message that describes the cause of the problem. Then, click **Start**, click **Run**, and type the following:

  **%windir%\Debug\adamsetup.log**

**See Also**
**Concepts**
Adaminstall
Install an ADAM replica from media

# Remove ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To remove ADAM**

1. Open **Add or Remove Programs** from **Control Panel**.
2. Click **Add/Remove Windows Components**.
3. Click **Active Directory Services** under **Components**, and then click **Details**.
4. Clear the **Active Directory Application Mode (ADAM)** check box under **Sub Components of Active Directory Services**, and then click **Next**.

**Notes**

- To open **Add or Remove Programs**, click **Start**, click **Control Panel**, and then click **Add or Remove Programs**.
- You can sort programs in **Add or Remove Programs** by selecting different options in **Sort by**.
- You might be prompted to restart the computer after removing **Active Directory Application Mode (ADAM)**. If you choose not to restart the computer, future installations of ADAM might fail.
- It is not possible to remove ADAM if there is an ADAM instance running on the computer.

# Create or Remove an ADAM Instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Create or Remove an ADAM Instance**

- Use the ADAM Setup Wizard
- Perform an unattended install of an ADAM instance
- Install an ADAM replica from media
- Remove an ADAM instance
- Perform an unattended removal of an ADAM instance

# Use the ADAM Setup Wizard

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Use the ADAM Setup Wizard**


- Starting the ADAM Setup Wizard
- Selecting a setup option
- Specifying an instance name
- Specifying communication ports
- Selecting an application directory partition option
- Specifying file locations
- Selecting a service account
- Selecting an administrator account
- Importing LDIF files during setup
- Reviewing ADAM Setup Wizard selections
- Finishing the installation

# Starting the ADAM Setup Wizard

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To start the ADAM Setup Wizard**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **Create an ADAM instance**.
2. On the **Welcome to the Active Directory Application Mode Setup Wizard** page, click **Next**.

**Note**

- You can use the ADAM Setup Wizard to install a new, unique ADAM instance or to install a replica ADAM instance.

**See Also**
**Concepts**
Adaminstall
Install an ADAM replica from media

# Selecting a setup option

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the **Setup Options** page, you must choose whether to install a unique Active Directory Application Mode (ADAM) instance or join an existing configuration set.

**To select a setup option**

- On the **Setup Options** page, do one of the following, and then click **Next**:
  - Click **A unique instance**.
  - Click **A replica of an existing instance**.

### Notes

- The following table summarizes the two choices presented on this page.

| Setup option | Result |
|---|---|
| A unique instance | If you choose to install a unique ADAM instance, the ADAM instance that you install will not be a part of any existing configuration set, and it will not have any replication partners. The ADAM Setup Wizard will create new, default copies of the configuration and schema directory partitions for the new ADAM instance that you are creating. Any schema extensions that you have added to any other ADAM instance will not be available in this new ADAM instance. |
| A replica of an existing instance | If you choose to install a replica of an existing ADAM instance, the ADAM instance that you are installing will contain replicated copies of the configuration and schema directory partitions (including any schema extensions) from the instance from which you choose to replicate. In addition, you have an opportunity in the ADAM Setup Wizard to selectively replicate any application directory partitions that exist in the source configuration set. You must specify a running ADAM instance as the source from which to replicate. To speed up the replication process for the new instance, you can use a restored copy of a previously backed up ADAM instance as the data source. |

- An ADAM instance consists of a single copy of the ADAM directory service, along with its associated directory store, assigned Lightweight Directory Access Protocol (LDAP) and Secure Sockets Layer (SSL) ports, and application event log. You can run multiple ADAM instances simultaneously on a single computer. For more ADAM terminology, see Related Topics.
- You can only join an ADAM instance to a configuration set during ADAM installation.

**See Also**
**Concepts**
Glossary
Understanding ADAM replication and configuration sets
Install an ADAM replica from media

# Specifying an instance name

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the **Instance Name** page, you must provide an easily recognizable name for the Active Directory Application Mode (ADAM) instance that you are installing. This name is used on the local computer to uniquely identify this ADAM instance.

The name that you specify appears in the following places:

- Event Viewer
- Services snap-in
- Add or Remove Programs

**To specify an instance name**

- On the **Instance Name** page, type a name for the ADAM instance that you are installing, and then click **Next**.

### Notes

- The Active Directory Application Mode Setup Wizard presents a default instance name of instance*X*, where *X* equals 1 for the first ADAM instance that is installed on the computer, 2 for the second ADAM instance that is installed on the computer, and so on.
- The name that you specify must meet the following requirements:
  - It must be unique with respect to other ADAM instances running on the same computer.
  - It must be no longer than 44 characters.
  - It must use characters only from the ranges of a through z, A through Z, or 0 through 9.
  - The name "ntds" cannot be used.
- The instance name that you specify is used to construct the service name, the service display name, and the service description, as follows:

| Name supplied during setup | Service name | Service display name | Service description |
|---|---|---|---|
| instance1 | ADAM_instance1 | instance1 | Blank (No default description is provided.) |

- You can modify the service display name, as well as the service description, after installation. For more information, see Related Topics.

**See Also**
**Concepts**
View or modify the service display name of an ADAM instance
View or modify the service description of an ADAM instance
Start, stop, pause, resume, or restart an ADAM instance

# Specifying communication ports

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the **Ports** page, you can specify the communications ports that the Active Directory Application Mode (ADAM) instance will use to communicate with directory-enabled applications. Directory-enabled applications locate an ADAM instance by using a combination of the name of the computer on which the ADAM instance is running and a communication port number. An ADAM instance can communicate using Lightweight Directory Access Protocol (LDAP) or Secure Sockets Layer (SSL). You must provide a port number for each.

If you are not sure what communication port numbers to use, accept the default values.

**To specify communication ports**

1. On the **Ports** page, type a port number in **LDAP port number**.
2. Type a port number in **SSL port number**, and then click **Next**.

### Notes

- By default, ADAM uses the default communication port numbers of 389 and 636 for LDAP and SSL, respectively. If either of these ports is already in use on the destination computer, ADAM specifies the first available port, starting at 50000.
- If you plan to run ADAM and Active Directory concurrently on the same computer, you must assign ports other than 389 and 636 to ADAM. Active Directory requires the use of ports 389 and 636 (as well as ports 3268 and 3269, if the domain controller running Active Directory is configured as a global catalog).
- For more information about how directory-enabled applications can locate a given ADAM instance, see Related Topics.
- For information about changing the communication ports of an ADAM instance, see Related Topics.

**See Also**
**Concepts**
Modify the communication ports used by an ADAM instance
Running ADAM on Windows XP Professional
Administering computers running ADAM
Administering an ADAM instance

# Selecting an application directory partition option

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the **Application Directory Partition** page, you can create a new application directory partition during setup for Active Directory Application Mode (ADAM). An application directory partition holds data that is specific to one or more directory-enabled applications. The ADAM Setup Wizard does not require that you create an application directory partition. You can also create application directory partitions at any time after ADAM installation.

**To select an application directory partition option**

- On the **Application Directory Partition** page, do one of the following, and then click **Next**:
  - Click **No, do not create an application directory partition**.
  - Click **Yes, create an application directory partition**, and then specify a distinguished name for the new partition. ADAM supports both X.500-style and Domain Name System (DNS)-style distinguished names. You can use any of the distinguished name attribute types in the following table when specifying an application partition in ADAM. You must include at least one attribute in the distinguished name that you specify.

| Distinguished name attribute | Type |
|---|---|
| C | Country/region |
| CN | Container |
| DC | Domain component |
| L | Locality |
| O | Organization |
| OU | Organizational unit |

For example:

CN=Partition1,DC=Woodgrove,DC=COM

**Important**

- If you want to add organizational unit (OU) containers in your application directory partition, limit the distinguished name that you specify to the following object classes: OU, C, O, and DC. By default, OUs can only be added under these object classes. For example, you can add an OU to o=Microsoft,c=US. You cannot add an OU to cn=test,o=Microsoft,c=US.

**Note**

- You can also create application directory partitions after ADAM installation. For more information, see Related Topics.

**See Also**
**Concepts**

Delete an application directory partition
Add an application directory partition
Understanding ADAM data and data stores
Administering directory partitions

# Specifying file locations

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the **File Locations** page, you can view and change the default installation directories for Active Directory Application Mode (ADAM) data (.dit) files and recovery (.log) files.

By default, ADAM data files and recovery files are installed on the destination computer in the following directory:

\%ProgramFiles%\Microsoft ADAM*instancename*\data

where *instancename* equals the ADAM instance name that you specified on the **Instance Name** page.

By default, ADAM program files and administration tools are installed in the following directory:

\%windir%\ADAM

**To specify install file locations for ADAM data or recovery files**

- On the **File Locations** page, do one of the following:
  - Click **Next** to accept the default file locations.
  - Type new file locations in **Data files** or **Data recovery files**, and then click **Next**.

### Notes

- If the directory paths that you specify do not already exist, the ADAM Setup Wizard creates them for you.
- The ADAM Setup Wizard does not accept environmental variables in file location paths.

# Selecting a service account

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the **Service Account Selection** page in the ADAM Setup Wizard, you must select a service account for use by the Active Directory Application Mode (ADAM) instance. The account that you select determines the security context in which the ADAM instance runs. Changing the service account after installation may require some additional configuration.

In most cases, the ADAM Setup Wizard defaults to the Network Service account as the service account. The Network Service account is a special, built-in account, with authority similar to that of an authenticated user account. The name of the account is NT AUTHORITY\NetworkService. The Network Service account has limited access to the local computer and authenticated access (as the computer account) to network resources. This limited access helps safeguard your system if individual services or processes are compromised. Services that run as the Network Service account access network resources using the credentials of the computer account.

For more information about selecting a service account and ADAM service account requirements, see Selecting an ADAM service account.

**To select a service account**

- On the **Service Account Selection** page, do one of the following, and then click **Next**:
  - Click **Network service account**.
  - Click **This account**, click **Browse**, and then select a local or domain account.

### Notes

- On Windows XP Professional, workstation and domain user accounts that are used as the ADAM service account must be members of the local Administrators group.
- When you install ADAM on Windows XP Professional computers that are joined to a workgroup, you must edit the **forceguest** registry key to allow users to bind to ADAM successfully. For more information, see Enable binding to ADAM instances running on Windows XP Professional computers joined to a workgroup.
- For additional considerations when running ADAM on Windows XP Professional, see Running ADAM on Windows XP Professional.
- To enable auditing for an ADAM instance running under a service account other than the Network Service account, you must grant the **Generate security audits** right to the account that is used as the ADAM service account.
- To enable a workstation or domain user account as a service account, you must grant the **Log on as a service** right to the account that is used as the ADAM service account. For more information, see Add the Log on as a service right to an account.

**See Also**
**Concepts**
Running ADAM on Windows XP Professional
Administering computers running ADAM
Selecting an ADAM service account
Administering an ADAM instance

# Selecting an administrator account

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the **ADAM Administrators** page, you must select a user or group to become the default administrator for this Active Directory Application Mode (ADAM) instance. This user or group that you select will have full administrative control of the ADAM instance.

By default, the ADAM Setup Wizard specifies the currently logged on user. You can change this selection to any local or domain account or group on your network.

**To select an administrator account**

- On the **ADAM Administrators** page, do one of the following, and then click **Next**:
  - Click **Currently logged on user**.
  - Click **This account**, click **Browse**, and then select a local or domain account or group.

### Notes

- If you specify a group to become the ADAM administrator, every member of that group will have administrative control of the ADAM instance.
- The user or group that you select automatically becomes a member of the Administrators group in the configuration directory partition of the ADAM instance. For more information about users and groups in ADAM, see Related Topics.
- Do not specify built-in groups or built-in accounts, such as DOMAIN\Administrators. Instead, if you want to specify a group, specify a domain group, such as *domainname*\Domain Admins, where *domainname* represents the name of your domain.

**See Also**
**Concepts**
Understanding ADAM users and groups
Administering users and groups

# Importing LDIF files during setup

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the **Import LDIF Files** page, you can specify that the optional ms-InetOrgPerson.ldf, ms-User.ldf, ms-UserProxy.ldf, and MS-AZMan.ldf LDAP Data Interchange Format (LDIF) files be imported. Together, these files contain several user class schema definitions, along with objects for use with Windows Authorization Manager, that can be imported into the schema of the new ADAM instance that is currently being installed.

**To import LDIF files during setup**

- On the **Import LDIF Files** page, do one of the following, and then click **Next**:
  - If you do not want to import or use the optional user class schema definitions, click **Do not import LDIF files for this instance of ADAM**.
  - If you want to use the optional user class schema definitions or Windows Authorization Manager with your ADAM instance, click **Import the selected LDIF files for this instance of ADAM**; in **Available files**, click one or more of the LDAP Data Interchange Format (LDIF) files to import; and then click **Add**.

### Notes

- The following table describes each of the optional .ldf files.

| .ldf file | User classes | Import this file if ... |
|---|---|---|
| MS-User.LDF | • **Person**<br>• **Organizational-Person**<br>• **User** | You want to create user objects in the ADAM directory, but you do not want to create users of the **InetOrgPerson** class (as defined in RFC 2798). |
| MS-InetOrgPerson.LDF | • **Person**<br>• **Organizational-Person**<br>• **User**<br>• **inetOrgPerson** | You want to create user objects in the ADAM directory, and you want to create users of the **InetOrgPerson** class (as defined in RFC 2798). |
| MS-UserProxy.LDF | • **User-Proxy** | You want to create proxy objects in ADAM for use in bind redirection. For more information about proxy objects and bind redirection, see Understanding ADAM bind redirection. |
| MS-AZMan.LDF | • Not applicable | You want to use Windows Authorization Manager with ADAM. For information about role-based authorization in Windows Server 2003, see "Dynamic Groups in Windows Server 2003 Authorization Manager" at the Microsoft MSDN Web site. |

- If you choose not to import these LDIF files during setup, you can import them at a later time, using the **ldifde** command-line tool. For more information, see Related Topics.
- This wizard page appears only during the installation of a unique ADAM instance and not during the installation of a

replica ADAM instance.

- You must be logged on as an account that is specified as the ADAM administrator to import these schema class definitions into the ADAM schema.

**See Also**
**Concepts**
Import the user classes supplied with ADAM
Administering users and groups
Understanding ADAM users and groups
ADAM Ldifde

# Reviewing ADAM Setup Wizard selections

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the **Ready to Install** page, you can review the ADAM setup options that you have selected.

**To review or change your selections in the ADAM Setup Wizard**

1. In the **Selections** box, use the scroll bar to review your selections.
2. To change any of your selections:
   a. Click **Back** until you reach the wizard page that contains the setting that you want to change.
   b. After making the change, click **Next** until you reach the **Ready to Install** page again.
3. To make any additional changes, repeat step 2.
4. Click **Next**.

**Note**

- After you click **Next**, the ADAM Setup Wizard automatically begins copying files and configuring the new ADAM instance.

# Finishing the installation

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

If the Active Directory Application Mode Setup Wizard completes successfully, the **Completing the Active Directory Application Mode Setup Wizard** page displays the message "You have successfully completed the ADAM Setup Wizard." If the Active Directory Application Mode Setup Wizard does not complete successfully, an error message appears.

Errors can occur either during the setup process or during directory service initialization.

**To finish the installation**

- Review the message that appears. Based on the contents of message, do one of the following:
  - If "You have successfully completed the ADAM Setup Wizard" appears, click **Finish**.
  - If an error message appears, make a note of the error, click **Finish**, and then review the Active Directory Application Mode (ADAM) event messages in **Event Viewer**.

### Notes

- If an error occurs in the Active Directory Application Mode Setup Wizard before the **Completing the Active Directory Application Mode Setup Wizard** page, first review the error message that describes the cause of the problem. Then, click **Start**, click **Run**, and type either of the following:

  **%windir%\Debug\adamsetup.log**

  **%windir%\Debug\adamuninstall.log**

  These files contain information that can help you troubleshoot the cause of the ADAM setup failure.

- For both installs and uninstalls, ADAM uses a registry key to return error codes and messages back to the caller (for example, an .msi package). ADAM Setup creates this registry key and associated values only if errors or warnings occur. This registry key is **HKLM\Software\Microsoft\Windows\CurrentVersion\ADAM_Installer_Results**, and it includes the possible values in the following table.

| Key | Contents |
|---|---|
| **ADAMInstallErrorCode** | The numeric error code that caused the installer to fail |
| **ADAMInstallErrorMessage** | A message that is associated with the install error code |
| **ADAMInstallWarnings** | Messages that are associated with the install warnings |
| **ADAMUninstallErrorCode** | The numeric error code that caused the uninstall to fail |
| **ADAMUninstallErrorMessage** | A message that is associated with the uninstall error code |
| **ADAMUninstallWarnings** | Messages that are associated with the uninstall warnings |

# Perform an unattended install of an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

In addition to using the Active Directory Application Mode Setup Wizard, you can install Active Directory Application Mode (ADAM) without user intervention. An unattended ADAM installation requires an answer file (Answer.txt) that contains a set of preconfigured installation options.

**To perform an unattended install of an ADAM instance**

1. Using any text editor, create a new text file. (As an option, you can copy and paste the sample answer file below into your answer file.)
2. Specify the installation parameters, as described in the table below.
3. Save your answer file.
4. At a command prompt (or in a batch or script file), change to the drive and directory that contains the ADAM setup files, and then type the following command:

   **Adamsetup /c:"adaminstall.exe /answer:** *drive:\pathname\filename.txt* **"**

   where *drive:\pathname\filename.txt* represents the drive, path, and file name of your answer file. (The double quotation marks in the command are required.)

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- The following table describes the parameters that can be used in an ADAM answer file. These parameters are not case sensitive. In other words, you can specify either **InstallType** or **installtype** in your answer file. However, case is preserved on the values that you specify for the **instancename** and **servicepassword** parameters.

| Parameter | Possible values | Default behavior if the parameter is not present in the answer file |
|---|---|---|
| **InstallType**<br>Valid for all installations.<br>Optional. | **Unique** — Creates a unique instance of ADAM.<br>**Replica** — Creates an instance of ADAM by replicating all or part of an existing ADAM instance, either over the network or from restored backup media.<br>When values are also specified in the answer file for the **ReplicationDataSourcePath** and **ReplicationLogSourcePath** parameters, and when **InstallType** is set to **Replica**, ADAM setup installs an ADAM replica instance from restored backup media. If values for those parameters are not present, ADAM setup installs an ADAM replica instance over the network. For more information about installation from media, see Install an ADAM replica from media.<br>**Tools** — Installs ADAM tools only.<br>Any other value — Returns "Invalid installation type specified in InstallType" error message. | Same behavior as **Unique** |

| | | |
|---|---|---|
| **ShowOrHideProgressGUI**<br>Valid for all installations.<br>Optional. | **Show** — ADAM setup displays progress information during installation.<br>**Hide** — ADAM setup does not display progress information during installation. | Same behavior as **Hide** |
| **InstanceName**<br>Valid for all installations.<br>Optional. | An ADAM instance name must meet the following requirements:<br>• It must be unique with respect to other ADAM instances running on the same computer.<br>• It must be no longer than 44 characters.<br>• It must use characters only from the ranges of a through z, A through Z, or 0 through 9. | The instance will be named Instance*n*, where *n* is the lowest number greater than 0 and Instance*n* is unique on the local computer. |
| **ApplicationPartitionsToReplicate**<br>Valid only for replica installations.<br>Optional. | Specify the distinguished names of the application partitions to replicate from the source ADAM instance.<br>The following example specifies three application partitions to replicate:<br>ApplicationPartitionsToReplicate = "CN=my,O=partition" "DC=partition2" "CN=embed qu\"ote in DN"<br>To replicate all application partitions from the source ADAM instance, specify "*" as the value. Any value that is specified for **ApplicationPartitionsToReplicate** is ignored if **InstallType** is not set to **Replica**. | No application partitions are replicated. |
| **ReplicationDataSourcePath**<br>Valid only for replica installations.<br>When a value for this parameter is present, ADAM setup attempts an installation from media. If the value for this parameter is not valid, ADAM setup writes an error to the setup log. | Specify the directory path to a restored instance of ADAM data. Any value that is specified for **ReplicationDataSourcePath** is ignored if **InstallType** is not set to **Replica** or if a value is not also specified for **ReplicationLogSourcePath**. | Application data is replicated over the network, rather than from a restored backup of an ADAM instance. If a value is specified for this parameter, but not for **ReplicationLogSourcePath**, an error occurs. |
| **ReplicationLogSourcePath**<br>Valid only for replica installations.<br>When a value for this parameter is present, ADAM setup attempts an installation from media. If the value for this parameter is not valid, ADAM setup writes an error to the setup log. | Specify the directory path to the log file for a restored instance of ADAM. Any value that is specified for **ReplicationLogSourcePath** is ignored if **InstallType** is not set to **Replica** or if a value is not also specified for **ReplicationDataSourcePath**. | Application data is replicated over the network, rather than from a restored backup of an ADAM instance. If a value is specified for this parameter, but not for **ReplicationDataSourcePath**, an error occurs. |
| **LocalLDAPPortToListenOn**<br>Required for all installations. | 389 or any unused port number between 1025 and 65535, inclusive.<br>Any other value — Returns "Invalid local LDAP port specified." | Defaults to 389. Or, if 389 is not available, defaults to the first available port equal to or greater than 50000. |

| | | |
|---|---|---|
| **LocalSSLPortToListenOn**<br>Required for all installations. | 636 or any unused port number between 1025 and 65535, inclusive.<br>Any other value — Returns "Invalid local SSL port specified." | Defaults to 636. Or, if 636 is not available, defaults to the first available port equal to or greater than 50000. |
| **SourceServer**<br>Required for replica installations. | A valid DNS name or network basic input/output system (NetBIOS) name.<br>Any other value — If the installation type is replica, returns "Invalid syntax for replication source server," else do nothing. | If the installation type is replica, returns "Replication source server not specified," else do nothing. |
| **SourceLDAPPort**<br>Required for replica installations. | Either 389 or a number between 1025 and 65535.<br>Any other value — If the installation type is replica, returns "Invalid replication source port specified," else do nothing. | If the installation type is replica, returns "Replication source port not specified," else do nothing. |
| **NewApplicationPartitionToCreate**<br>Valid for installations of new, unique ADAM instances.<br>Optional. | A valid distinguished name — Creates an application partition with the specified name.<br>"" (empty string) — No new application partition will be created.<br>Any other value — If the installation type is unique, returns "Invalid application partition syntax in NewApplicationPartitionToCreate," else do nothing. | Same behavior as "" |
| **DataFilesPath**<br>Valid for all installations.<br>Optional. | A syntactically correct path name, which may include unresolved environment variables that do not contain existing ADAM files.<br>Any other value — Returns "Invalid path in DataFilesPath." | Store data files in Program Files\Microsoft ADAM*instancename*\data. |
| **LogFilesPath**<br>Valid for all installations.<br>Optional. | A syntactically correct path name, which may include unresolved environment variables that do not contain existing ADAM files.<br>Any other value — Returns "Invalid path in LogFilesPath." | Stores log files in Program Files\Microsoft ADAM*instancename*\data |
| **ServiceAccount**<br>Valid for all installations.<br>Optional. | A valid DNS domain name, followed by a backslash and the account or group name.<br>A valid NetBIOS domain name, followed by a backslash and the account name.<br>A valid user principal name (UPN).<br>A valid account name only. (This is generally not recommended, because resolving an account name that is not accompanied by a domain name requires additional processing.)<br>Any other value — Returns "Invalid user specified in ServiceAccount." | This instance of ADAM will run under the Network Service account. |

| | | |
|---|---|---|
| **AddPermissionsToServiceAccount**<br>Valid for all installations.<br>Optional. | **Yes** — ADAM setup attempts to add the **Log on as a service** right to the account that is specified as the service account.<br>Any other value — ADAM setup does not attempt to add the **Log on as a service** right to the account that is specified as the service account. | ADAM setup does not attempt to add the **Log on as a service** right to the account that is specified as the service account. |
| **ServicePassword**<br>Valid for all installations.<br>Required, unless **ServiceAccount** is the Network Service account. | Any string of characters, including the empty string. | If **ServiceAccount** is the Network Service account, do nothing; otherwise, returns "No password specified in ServicePassword." |
| **Administrator**<br>Valid for all installations.<br>Optional. | A valid DNS domain name, followed by a backslash and the account name.<br>Do not specify built-in groups or built-in accounts, such as DOMAIN\Administrators. Instead, if you want to specify a group, specify a domain group, such as *domainname*\Domain Admins, where *domainname* represents the name of your domain.<br>A valid NetBIOS domain name, followed by a backslash and the account name.<br>A valid UPN.<br>A valid account name only. (This is generally not recommended, because resolving an account name that is not accompanied by a domain name requires additional processing.)<br>Any other value — Returns "Invalid user specified in Administrator." | The currently logged on user will have administrator permissions on this instance of ADAM. |
| **ShowInAddRemovePrograms**<br>Valid for all installations.<br>Optional. | **Show** — Lists the ADAM instance in **Add or Remove Programs**.<br>**Hide** — Does not list the ADAM instance in **Add or Remove Programs** | The installed ADAM instance is shown in **Add or Remove Programs**. |
| **ImportLDIFFiles** | The optional .ldf files (ms-User.ldf, ms-InetOrgPerson.ldf, ms-UserProxy.ldf, and ms-azman.ldf) that you want to import into the ADAM schema. The file names must be enclosed in double quotation marks, delimited with a space (" "). | None of the optional .ldf files are imported. |
| **SourceUserName**<br>**SourcePassword**<br>Valid for replica installations.<br>Optional. | The user name and password of an account that has administrative rights for an existing configuration set. Use these parameters when installing a replica that is being joined to the configuration set. | The replica is joined to the configuration set using the credentials of the logged on user. |

- The following represents a sample answer file for use with **adaminstall**.

```
[ADAMInstall]

; The following line specifies to install a unique ADAM instance.
```

```
InstallType=Unique

; The following line specifies the name to be assigned to the new instance.
InstanceName=MyFirstInstance

; The following line specifies the communications port to use for LDAP.
LocalLDAPPortToListenOn=389

; The following line specifies an application partition to create
NewApplicationPartitionToCreate="o=microsoft,c=us"

; The following line specifies the directory to use for ADAM data files.
DataFilesPath=C:\Program Files\Microsoft ADAM\instance1\data

; The following line specifies the directory to use for ADAM log files.
LogFilesPath=C:\Program Files\Microsoft ADAM\instance1\data

; The following line specifies the .ldf files to import into the ADAM schema.
ImportLDIFFiles="ms-inetorgperson.ldf" "ms-user.ldf"
```

- For both installs and uninstalls, ADAM uses a registry key to return error codes and messages back to the caller, for example, an .msi package. ADAM Setup creates this registry key and associated values only if errors or warnings occur. This registry key is **HKLM\Software\Microsoft\Windows\CurrentVersion\ADAM_Installer_Results**, and it includes the possible values in the following table.

| Key | Contents |
|---|---|
| **ADAMInstallErrorCode** | The numeric error code that caused the installer to fail |
| **ADAMInstallErrorMessage** | A message that is associated with the install error code |
| **ADAMInstallWarnings** | Messages that are associated with the install warnings |
| **ADAMUninstallErrorCode** | The numeric error code that caused the uninstall to fail |
| **ADAMUninstallErrorMessage** | A message that is associated with the uninstall error code |
| **ADAMUninstallWarnings** | Messages that are associated with the uninstall warnings |

**See Also**
**Concepts**
Starting the ADAM Setup Wizard
Install ADAM
Manually register ADAM service principal names in Active Directory
Adaminstall

# Install an ADAM replica from media

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Install an ADAM replica from media**
When you install an Active Directory Application Mode (ADAM) replica from media, you use a restored backup of an ADAM instance as the data source, rather than another ADAM instance.

- Using the Active Directory Application Mode Setup Wizard
- Using ADAM unattended install

**Using the Active Directory Application Mode Setup Wizard**
**To install an ADAM replica from media**

1. Restore a backup copy of the ADAM instance from which you want to install to an alternate location. (Do not restore the backup to the original location of the ADAM instance.)
2. At a command prompt, type the following:

   *d:\path* **adaminstall /adv**

   where *d:\path* represents the location of Adaminstall.exe.
3. Follow the steps in the Active Directory Application Mode Setup Wizard.

### Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- To view Help for any given wizard page, click **Help** in the Active Directory Application Mode Setup Wizard.
- You can install an ADAM replica only from media created by restoring a backup copy of an ADAM instance. For information about restoring an ADAM instance, see Related Topics.
- When you restore an ADAM instance for use in a replica installation from media, you must restore the files to an alternate location, rather than to the original location from which they were backed up. If the replica that you are creating uses a different service account than the service account that the original instance uses, you must assign permission to read all the directories and files in the alternate location to the service account of the replica instance before you run Adaminstall.exe.
- After you restore ADAM files from a backup to an alternate location, the Adamntds.dit file and Edb*.log files will be nested in the specified alternate location. For example, if you specify c:\restore_dir as the restore location for the ADAM files, Adamntds.dit and the Edb*.log files will be located at c:\restore_dir\Program Files\Microsoft ADAM*instancename*\data, where *instancename* represents the ADAM instance that was restored.
- For information about how to back up an ADAM instance, see Related Topics.
- For more information about ADAM unattended installs, see Related Topics.

**Using ADAM unattended install**
**To install an ADAM replica from media**

1. Restore a backup copy of the ADAM replica from which you want to install to an alternate location. (Do not restore the backup to the original location of the ADAM instance.)

2. Create an ADAM unattended install answer file that contains the answer file attributes and values in the following table, in addition to any other answer file parameters that you want to set, as described in Perform an unattended install of an ADAM instance.

| Attribute | Value |
|---|---|
| **InstallType** | **Replica** |
| **ReplicationDataSourcePath** | The directory location to which you restored the ADAM replica data files |
| **ReplicationLogSourcePath** | The directory location to which you restored the ADAM replica log files |
| **ApplicationPartitionsToReplicate** | The application directory partitions (from the restored files) that you want to install. Setting **ApplicationPartitionsToReplicate** equal to ***** installs all application directory partitions that exist in the restored files. Setting a specific value for this parameter installs the specified application directory partitions, along with any application directory partitions that are not part of the restored files but that do exist on other ADAM instances in the configuration set. |

3. At a command prompt (or in a batch or script file), change to the drive and directory that contains the ADAM setup files, and then type the following command:

   **adamsetup /c:"adaminstall.exe /answer:** *drive:\pathname\filename.txt*"

   where *drive:\pathname\filename.txt* represents the drive, path, and file name of your answer file.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- You can install an ADAM replica only from media created by restoring a backup copy of an ADAM instance. For information about restoring an ADAM instance, see Related Topics.
- When you restore an ADAM instance for use in a replica install from media, you must restore the files to an alternate location, rather than to the original location from which they were backed up.
- After you restore ADAM files from a backup into an alternate location, the Adamntds.dit file and Edb*.log files will be nested in the alternate location that you specified. For example, if you specified **c:\restore_dir** as the restore location for the ADAM files, Adamntds.dit and the Edb*.log files will be located at c:\restore_dir\Program Files\Microsoft ADAM*instancename*\data, where *instancename* represents the ADAM instance that was restored.
- For information about how to back up an ADAM instance, see Related Topics.
- For more information about ADAM unattended installs, see Related Topics.
- For more information about unattended install parameters, see Related Topics.

**See Also**
**Concepts**
Perform an unattended install of an ADAM instance
Adaminstall
Back up an ADAM instance to a file or a tape
Restore an ADAM instance

# Remove an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To remove an ADAM instance**

1. Open Add or Remove Programs.
2. Click **Change or Remove Programs**, and then click the ADAM instance that you want to remove.
3. Click **Remove**.

**Notes**

- To open Add or Remove Programs, click **Start**, click **Control Panel**, and click **Add or Remove Programs**.
- You can sort programs in **Add or Remove Programs** by selecting different options in **Sort by**.
- It is not necessary to restart the computer after removing an ADAM instance.

# Perform an unattended removal of an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can remove an ADAM instance without using Add or Remove Programs, by performing an unattended removal of Active Directory Application Mode (ADAM). An unattended removal of ADAM can only be performed under the security context of the ADAM administrator. If you are not currently logged on as the ADAM administrator, the ADAM administrator credentials can be specified through a text answer file.

**To perform an unattended removal of an ADAM instance**

- At a command prompt (or in a batch or script file), type the following command:

  **adamuninstall /i:** *instancename* **/answer:***d:\dirpath\filename*

  where *instancename* represents the name of the ADAM instance to be removed, and *d:\dirpath\filename* represents the drive, path and file name of an optional answer file containing the credentials of the ADAM administrator for the specified instance.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.

- To pass the credentials of the ADAM administrator to **adamuninstall** from the command line, create a text answer file containing the following syntax:

  **Administrator=***domain\account*

  **AdministratorPassword=***password*

  where *domain\account* represents the domain and account name of the ADAM administrator, and *password* represents the corresponding password.

- The following represents a sample answer file for use with **adamuninstall**:

```
Administrator=cohovineyard\administrator
AdministratorPassword=adminpassword
```

- For both installs and uninstalls, ADAM uses a registry key to return error codes and messages back to the caller, for example, an .msi package. ADAM Setup creates this registry key and associated values only if errors or warnings occur. This registry key is **HKLM\Software\Microsoft\Windows\CurrentVersion\ADAM_Installer_Results**, and it includes the possible values in the following table.

| Key | Contents |
|---|---|
| **ADAMInstallErrorCode** | The numeric error code that caused the installer to fail |
| **ADAMInstallErrorMessage** | A message that is associated with the install error code |
| **ADAMInstallWarnings** | Messages that are associated with the install warnings |
| **ADAMUninstallErrorCode** | The numeric error code that caused the uninstall to fail |

| | |
|---|---|
| **ADAMUninstallErrorMessage** | A message that is associated with the uninstall error code |
| **ADAMUninstallWarnings** | Messages that are associated with the uninstall warnings |

**See Also**
**Concepts**
Remove an ADAM instance
Manually remove ADAM service principal names from Active Directory
Adamuninstall

# Manage an ADAM Instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage an ADAM instance**

- Connect and bind to an ADAM instance using ADAM ADSI Edit
- Connect and bind to an ADAM instance using Ldp.exe
- Browse the directory tree using ADAM ADSI Edit
- Browse the directory tree using Ldp.exe
- List the directory partitions, or naming contexts, of an ADAM instance or configuration set
- List the ADAM instances installed on a computer
- Start, stop, pause, resume, or restart an ADAM instance
- Add the Log on as a service right to an account
- Modify the communication ports used by an ADAM instance
- Modify the service account used by an ADAM instance
- View or modify the service display name of an ADAM instance
- View or modify the service description of an ADAM instance
- Enable binding to ADAM instances running on Windows XP Professional computers joined to a workgroup
- Allow anonymous LDAP binding to an ADAM instance
- Manually register ADAM service principal names in Active Directory
- Manually remove ADAM service principal names from Active Directory

# Connect and bind to an ADAM instance using ADAM ADSI Edit

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To connect and bind to an ADAM instance using ADAM ADSI Edit**

1. Click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
2. In the console tree, click **ADAM ADSI Edit**.
3. On the **Action** menu, click **Connect to**.
4. In **Connection name**, type a name for this connection.
5. In **Server name**, type the DNS name, network basic input/output system (NetBIOS) name, or Internet Protocol (IP) address of the computer on which the ADAM instance is running.
6. In **Port**, type the Lightweight Directory Access Protocol (LDAP) communication port in use by the ADAM instance to which you want to connect.
7. Under **Connect to the following node**, do one of the following:
   - Click **Distinguished name (DN) or naming context**, and then specify the distinguished name to which you want to connect.
   - Click **Well-known naming context**, and then click **Configuration**, **RootDSE**, or **Schema**.
8. Under **Connect using these credentials**, do one of the following:
   - Click **The account of the currently logged on user**.
   - Click **This account**, and then specify the user name and password of the account with which you want to connect.

**Notes**

- To create additional connections to ADAM instances, on the **Action** menu, click **Connect to** for each new connection.
- When you run ADAM on computers running Windows XP Professional, you cannot set or modify the password for any principal in ADAM using ADAM ADSI Edit.
- The default communication port for LDAP is 389.
- ADAM ADSI Edit does not support Secure Sockets Layer (SSL) connections to ADAM.
- To connect to an ADAM instance running on the local computer, type **localhost** as the server name.

# Connect and bind to an ADAM instance using Ldp.exe

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To connect and bind to an ADAM instance using Ldp.exe**

1. Open **Ldp**.
2. On the **Connection** menu, click **Connect**.
3. In **Server**, type the DNS name, network basic input/output system (NetBIOS) name, or Internet Protocol (IP) address of the computer on which the ADAM instance is running.
4. In **Port**, type the Lightweight Directory Access Protocol (LDAP) or Secure Sockets Layer (SSL) communication port number in use by the ADAM instance to which you want to connect, and then click **OK**.
5. On the **Connection** menu, click **Bind**.
6. Do one of the following:
   - To bind using the credentials you logged on with, click **Bind as currently logged on user**.
   - To bind using a domain user account, click **Bind using credentials**,type the user name, password, and domain name (or the computer name, if you are using a local workstation account) of the account that you are using, and then click **OK.**
   - To bind using just a user name and password, click **Simple bind**, type the user name and password of the account that you are using, and then click **OK**.
   - To bind using an advanced method (NTLM, DPA, negotiate, or digest), click **Advanced (***method***)**, click **Advanced**, in **Method**, select the desired method, set other options as needed, and then click **OK**.
7. When you are finished specifying bind options, click **OK**.

### Notes

- To open **Ldp**, click **Start**, point to **All Programs**, point to **ADAM**, click **ADAMTools Command Prompt**, and then type **ldp** at the command prompt.
- The default communication port for LDAP is 389. The default communication port for SSL is 636.
- To connect to an ADAM instance running on the local computer, type **localhost** as the server name.

**See Also**
**Concepts**
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using Ldp.exe
Ldp

# Browse the directory tree using ADAM ADSI Edit

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To browse the directory tree using ADAM ADSI Edit**

1. Connect and bind to a directory partition using ADAM ADSI Edit.
2. In the console tree, double-click *connectionname*[*hostname*], where *connectionname* is the name of the connection and *hostname* is the name of the computer on which the directory partition is stored.
3. In the console tree, double-click a directory partition object to view its top-level containers.
4. In the console tree, double-click a top-level container to view the next level of objects in that container.
5. Do one of the following:
   - Continue double-clicking the next lowest container level to continue down a directory tree branch.
   - To view a directory object at the current directory level, in the details pane, double-click the directory object.

## Notes

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- To connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.
- ADAM ADSI Edit does not support Secure Sockets Layer (SSL) connections to ADAM.
- For more information about ADAM ADSI Edit, see Related Topics.

**See Also**
**Concepts**
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using Ldp.exe
ADAM ADSI Edit

# Browse the directory tree using Ldp.exe

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To browse the directory tree using Ldp.exe**

1. Connect and bind to an ADAM instance using Ldp.exe.
2. On the **View** menu, click **Tree**.
3. In the **Tree View** dialog box, do one of the following:
   - To view all directory partitions on the ADAM instance, click **OK**.
   - To view only a specific directory partition on the ADAM instance, type the distinguished name of the directory partition in **BaseDN**, and then click **OK**.
4. In the console tree, double-click a directory partition object to view its top-level containers.
5. Double-click a top-level container to view the next level of objects in that container.
6. Do one of the following:
   - Continue double-clicking the next lowest container level to continue down a directory tree branch.
   - To view a directory object at the current directory level, double-click the object in the console tree, and then, in the details pane, view object details.

**Notes**

- To open Ldp.exe, click **Start**, point to **All Programs**, point to **ADAM**, click **ADAM Tools Command Prompt**, and then type **ldp** at the command prompt.
- To connect and bind to an ADAM instance using Ldp.exe, see Related Topics.
- For more information about Ldp, see Related Topics.

**See Also**
**Concepts**
Connect and bind to an ADAM instance using Ldp.exe
Browse the directory tree using Ldp.exe
Browse the directory tree using ADAM ADSI Edit
Ldp

# List the directory partitions, or naming contexts, of an ADAM instance or configuration set

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To list the directory partitions of an ADAM instance or configuration set**

1. Open an ADAM tools command prompt.
2. At the command prompt, type:

   **dsmgmt**
3. At the **dsmgmt:** prompt, type:

   **partition management**
4. At the **partition management:** prompt, type:

   **connections**
5. At the **server connections:** prompt, type:

   **connect to server** *computername:portnumber*

   where *computername* is the computer on which the ADAM instance is running and *portnumber* is the Lightweight Directory Access Protocol (LDAP) or Secure Sockets Layer (SSL) communications port used by the ADAM instance.
6. At the **server connections:** prompt, type:

   **quit**
7. At the **partition management:** prompt, type:

   **list**
8. After viewing the list of directory partitions, at the **partition management:** prompt, type:

   **quit**

   and then at the **dsmgmt:** prompt, type:

   **quit**

**Note**

- To open an ADAM Tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- For help with **dsmgmt**, type **?** at a **dsmgmt** prompt.

**See Also**
**Concepts**
Dsmgmt
Administering directory partitions

# List the ADAM instances installed on a computer

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To list the ADAM instances installed on a computer**

1. Open an ADAM tools command prompt.
2. At the command prompt, type:

   **dsdbutil**

3. At the **dsdbutil:** prompt, type:

   **list instances**

4. To exit the **dsdbutil** administration tool, at the **dsdbutil:** prompt, type:

   **quit**

## Notes

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.

- The information about the ADAM instances that are listed by this procedure includes the following: ADAM instance name, Lightweight Directory Access Protocol (LDAP) port, Secure Sockets Layer (SSL) port, installation directory, database file, log file directory, and service state.

**See Also**
**Concepts**
Dsdbutil
Administering an ADAM instance
List the directory partitions, or naming contexts, of an ADAM instance or configuration set

# Start, stop, pause, resume, or restart an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Start, stop, pause, resume, or restart an ADAM instance**

- Using the graphical user interface
- Using the command prompt

**Using the graphical user interface**
**To start, stop, pause, resume, or restart an ADAM instance**

1. Open **Services**.
2. In the details pane, do one of the following:
   - Click the ADAM instance that you want to manage. On the **Action** menu, click **Start**, **Stop**, **Pause**, **Resume**, or **Restart**.
   - Right-click the service, and then click **Start**, **Stop**, **Pause**, **Resume**, or **Restart**.

### Notes

- To open **Services**, click **Start**, click **Run**, type **services.msc**, and then click **OK**.
- By default, an ADAM instance is configured to start automatically.

**Using the command prompt**
**To start, stop, pause, resume, or restart an ADAM instance**

1. Open a Command Prompt.
2. Do one of the following:
   - To start an ADAM instance, type the following at the command prompt:

     **net start** *instance_name*

     where *instance_name* represents the ADAM instance that you want to start.
   - To stop an ADAM instance, type the following at the command prompt:

     **net stop** *instance_name*

     where *instance_name* represents the ADAM instance that you want to stop.

### Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- You cannot pause or resume an ADAM instance from a command prompt.
- By default, an ADAM instance is configured to start automatically.

- For more information about the **net** command, type **net /?** at a command prompt, or see Help and Support Center in Windows Server 2003.

# Add the Log on as a service right to an account

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Add the Log on as a service right to an account**
**To add the "Log on as a service" right to an account on your local computer**

1. Open Local Security Policy.
2. In the console tree, double-click **Local Policies**, and then click **User Rights Assignments**.
3. In the details pane, double-click **Log on as a service**.
4. Click **Add User or Group**, and then add the appropriate account to the list of accounts that possess the **Log on as a service** right.

**To add the "Log on as a service" right to an account for a Group Policy object, when you are on a workstation or server that is joined to a domain**

1. Click **Start**, point to **Run**, type **mmc**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. In **Add/Remove Snap-in**, click **Add**, and then, in **Add Standalone Snap-in**, double-click **Group Policy Object Editor**.
4. In **Select Group Policy Object**, click **Browse**, browse to the Group Policy object (GPO) that you want to modify, click **OK**, and then click **Finish**.
5. Click **Close**, and then click **OK**.
6. In the console tree, click **User Rights Assignment**.

   Where?

   - *GroupPolicyObject* [*ComputerName*] Policy
   - Computer Configuration
   - Windows Settings
   - Security Settings
   - Local Policies
   - User Rights Assignment

7. In the details pane, double-click **Log on as a service**.
8. If the security setting has not yet been defined, select the **Define these policy settings** check box.
9. Click **Add User or Group**, and then add the appropriate account to the list of accounts that possess the **Log on as a service** right.

**To add the "Log on as a service" right to an account for a Group Policy object, when you are on a domain controller or on a computer that has the Windows Server 2003 Administration Tools Pack installed**

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the domain or organizational unit (OU) for which you want to edit security settings.
3. Click **Properties**, and then click the **Group Policy** tab.

4. In **Group Policy Object Links**, click the Group Policy object for the domain or OU for which you want to edit security settings, and then click **Edit**.
5. In the console tree, click **User Rights Assignment**.

   Where?

   - *GroupPolicyObject* [*ComputerName*] Policy
   - Computer Configuration
   - Windows Settings
   - Security Settings
   - Local Policies
   - User Rights Assignment

6. Double-click **Log on as a service** in the details pane.
7. If this security setting has not yet been defined, select the **Define these policy settings** check box.
8. Click **Add User or Group**, and then add the appropriate account to the list of accounts that possess the **Log on as a service** right.

**To add the "Log on as a service" right to an account for only domain controllers, when you are on a domain controller**

1. Open **Domain Controller Security Policy**.
2. In the console tree, click **User Rights Assignment**.

   Where?

   - Security Settings
   - Local Policies
   - User Rights Assignment

3. In the details pane, double-click **Log on as a service**.
4. If this security setting has not yet been defined, select the **Define these policy settings** check box.
5. Click **Add User or Group**, and then add the appropriate account to the list of accounts that possess the **Log on as a service** right.

### Notes

- To open **Local Security Policy**, click **Start**, point to **Control Panel**, point to **Administrative Tools**, and then double-click **Local Security Policy**.
- To open **Domain Controller Security Policy**, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Domain Controller Security Policy**.
- When you change a security setting, that setting will take effect in the next refresh of settings.
- The security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes.

# Modify the communication ports used by an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To modify the communications ports used by an ADAM instance**

1. Open an ADAM tools command prompt.
2. At the command prompt, type:

   **net stop** *instancename*

   where *instancename* represents the service name of the ADAM instance on which you want to change the communication port.
3. At the command prompt, type:

   **dsdbutil**
4. At the **dsdbutil:** prompt, type:

   **activate instance** *instancename*

   where *instancename* represents the service name of the ADAM instance on which you want to modify the communication port.
5. At the **dsdbutil:** prompt, do one of the following:

   - To change the Lightweight Directory Access Protocol (LDAP) communication port, type:

     **LDAP port** *portnumber*

     where *portnumber* represents a communication port (either 389 or in the range of 1025 through 65535) that you want to assign to LDAP.

   - To change the Secure Sockets Layer (SSL) communication port, type:

     **SSL port** *portnumber*

     where *portnumber* represents the communication port (either 636 or in the range of 1025 through 65535) that you want to assign to SSL.

6. To exit **dsdbutil**, at the **dsdbutil:** prompt, type:

   **quit**
7. At the command prompt, type:

   **net start** *instancename*

   to restart the ADAM instance.

**Notes**

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- Allowable communication port numbers in ADAM include 389, 636, and any port number in the range of 1025 through 65535.

**See Also**
**Concepts**
Using ADAM administration tools

Running ADAM on Windows XP Professional
Administering computers running ADAM

# Modify the service account used by an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To modify the service account used by an ADAM instance**

1. Open an ADAM tools command prompt.
2. At the command prompt, type:

   **net stop** *instancename*

   where *instancename* represents the service name of the ADAM instance on which you want to change the service account.
3. At the command prompt, type:

   **dsdbutil**
4. At the **dsdbutil:** prompt, type:

   **activate instance** *instancename*

   where *instancename* represents the service name of the ADAM instance on which you want to change the service account.
5. At the **dsdbutil:** prompt, type:

   **change service account** *accountname password*

   where *accountname password* represents the account name and password of the account to be used as the ADAM service account.
6. To exit **dsdbutil**, at the **dsdbutil:** prompt, type:

   **quit**
7. At the command prompt, type:

   **net start** *instancename*

   to restart the ADAM instance.

**Notes**

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- When you specify a workstation or domain user account as the service account, the specified account must possess the **Log on as a service** right. To assign the **Log on as a service** right to an account, you must edit the security setting on the appropriate Group Policy object (GPO). For more information, see Related Topics.
- The service account that you select must have permission to read and write the ADAM data and log files located in %ProgramFiles%\Microsoft ADAM*instancename*.
- The service account that you select may depend on whether the ADAM instance participates in a configuration set and also on the replication security level. For more information, see Related Topics.

**See Also**
**Concepts**
Add the Log on as a service right to an account
Administering replication and configuration sets
Understanding ADAM replication and configuration sets
Selecting an ADAM service account

# View or modify the service display name of an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To view or modify the service display name of an ADAM instance**

1. Open Command Prompt.
2. Do one of the following:
   - To view the service display name, type:

     **sc** *servername* **getdisplayname** *servicename*

     where *servername* represents the name of the computer on which the ADAM instance is running, and *servicename* represents the service name of the ADAM instance.

   - To modify the service display name, type:

     **sc** *servername* **config** *servicename* **displayname= "***display_name***"**

     where *servername* represents the name of the computer on which the ADAM instance is running, *servicename* represents the service name of the ADAM instance that you want to modify, and *display_name* represents the new service display name that you want to assign. (A space character (" ") is required after the equal sign (=) and before *display_name*).

### Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- If the ADAM instance is running on the local computer, you can omit the *servername* parameter.
- By default, the ADAM instance name is *instancename*, where *instancename* is the name that you provided during ADAM installation.
- By default, the ADAM service name is ADAM_*instancename*, where *instancename* is the name that you provided during ADAM installation.

**See Also**
**Concepts**
Administering an ADAM instance
View or modify the service description of an ADAM instance

# View or modify the service description of an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To view or modify the service description of an ADAM instance**

1. Open Command Prompt.
2. Do one of the following:
   - To view the service description, type:

     **sc** *servername* **qdescription** *serviceename*

     where *servername* represents the name of the computer on which the ADAM instance is running, and *servicename* represents the service name of the ADAM instance.

   - To modify the service description, type:

     **sc** *servername* **description** *servicename* "*service_description*"

     where *servername* represents the name of the computer on which the ADAM service is running, *servicename* represents the service name of the ADAM instance, and *service_description* represents the new description.

### Notes

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- By default, the ADAM service description is empty.
- By default, the ADAM service name is ADAM_*instancename*, where *instancename* is the name that you provided during ADAM setup.

**See Also**
**Concepts**
Administering an ADAM instance
View or modify the service display name of an ADAM instance

# Enable binding to ADAM instances running on Windows XP Professional computers joined to a workgroup

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To enable binding to ADAM instances running on Windows XP Professional computers joined to a workgroup**

1. Open Registry Editor.
2. In Registry Editor, navigate to the registry subkey HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa.

   Where?

   - HKEY_LOCAL_MACHINE
   - System
   - CurrentControlSet
   - Control
   - LSA

3. In the details pane, right-click **forceguest**, and then click **Modify**.
4. In **Value data**, type **0**, and then click **OK**.

**Caution**

- Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You can also use the Last Known Good Configuration startup option if you encounter problems after manual changes have been applied.

**Note**

- To open Registry Editor, click **Start**, click **Run**, type **regedit**, and then click **OK**.

**See Also**
**Concepts**
View or modify the service display name of an ADAM instance
View or modify the service description of an ADAM instance
Start, stop, pause, resume, or restart an ADAM instance
Administering an ADAM instance

# Allow anonymous LDAP binding to an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To allow anonymous LDAP binding to an ADAM instance**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of the ADAM instance on which you want to allow anonymous Lightweight Directory Access Protocol (LDAP) binding.
3. In the console tree, double-click the configuration directory partition (CN=Configuration,CN={*GUID*}), double-click the services container (CN=Services), double-click the Windows NT container (CN=Windows NT), right-click the directory service container (CN=Directory Service), and then click **Properties**.
4. In **Attributes**, click **dsHeuristics**, and then click **Edit**.
5. In **Value**, modify the value of the seventh character in the attribute (counting from the left) to 2, as follows:

   0000002001001
6. Click **OK** twice.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- To connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.
- ADAM does not allow anonymous LDAP binding by default.

**See Also**
**Concepts**
Connect and bind to an ADAM instance using ADAM ADSI Edit
Administering an ADAM instance
ADAM ADSI Edit
ADAM troubleshooting and frequently asked questions (FAQs)

# Manually register ADAM service principal names in Active Directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To manually register ADAM service principal names in Active Directory**

1. Do one of the following:
   - Log on as a member of the Domain Admins group, and then open a command prompt.
   - Using **run as**, open a command prompt as a member of the Domain Admins group.
2. Change the current directory to:

   Program Files\Microsoft ADAM*instancename*\data

   where *instancename* represents the name of the ADAM instance.
3. At the command prompt, type:

   *dnsdomainname*

   where *dnsdomainname* represents the name of the *dnsdomainname*.bat script file that ADAM Setup created for manually registering service principal names (SPNs) in Active Directory. ADAM uses the DNS domain name of the computer on which ADAM is running to name this file.

**Notes**

- To open a command prompt, click **Start**, point to **All programs**, point to **Accessories**, and then click **Command prompt**.
- SPN registration does not apply to ADAM instances running on computers that are joined to a workgroup, rather than to a domain.
- If you do not register the SPNs for an ADAM instance running on a computer that is joined to a domain, you can still run the ADAM instance, but you will not be able to run the ADAM instance as part of a configuration set for which the replication security level has been set to 2. For more information about replication security levels, see Related Topics.
- The *dnsdomainname*.bat file appears in the directory approximately one minute after ADAM installation is complete.
- If the *dnsdomainname*.bat file does not appear after approximately one minute, the attempt of ADAM setup to register SPNs for the ADAM instance completed successfully.
- The *dnsdomainname*.bat file contains **repadmin /writespn** commands, similar to the commands in the following table.

```
repadmin.exe /writespn hostname.microsoft.com
ADD "CN=adam,CN=Users,DC=microsoft,DC=com"
E3514235-4B06-11D1-AB04-00C04FC2DCD2-ADAM/netbioshostname:389

repadmin.exe /writespn hostname.microsoft.com
ADD "CN=adam,CN=Users,DC=microsoft,DC=com"
E3514235-4B06-11D1-AB04-00C04FC2DCD2-ADAM/dnshostname:389

repadmin.exe /writespn hostname.microsoft.com
ADD "CN=adam,CN=Users,DC=microsoft,DC=com" ldap/netbioshostname:389
```

```
repadmin.exe /writespn hostname.microsoft.com
ADD "CN=adam,CN=Users,DC=microsoft,DC=com" ldap/dnshostname:389
```

**See Also**

**Concepts**

Administering ADAM service principal names

Manually remove ADAM service principal names from Active Directory

Repadmin

Understanding ADAM replication and configuration sets

# Manually remove ADAM service principal names from Active Directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To manually remove ADAM service principal names from Active Directory**

1. Do one of the following:
   - Log on as a member of the Domain Admins group, and then open ADAM ADSI Edit.
   - Using **run as**, open ADAM ADSI Edit as a member of the Domain Admins group.
2. Connect and bind to the domain directory partition of the Active Directory domain in which the ADAM instance resides.
3. Do one of the following:
   - If the Network Service account is specified as the ADAM service account, navigate to the computer object on the computer on which ADAM is installed.
   - If a domain user account is specified as the ADAM service account, navigate to the domain user object of the domain user account.
4. Right-click the computer object or the domain user object, and then click **Properties**.
5. In **Attributes**, click **ServicePrincipalName**, and then click **Edit**.
6. For each ADAM value that is listed in **Values**, click the value, and then click **Remove**.

**Note**

- If service principal names (SPN) in Active Directory exist for the ADAM instance being uninstalled, and **adamuninstall** cannot remove the SPNs automatically, **adamuninstall** creates a .bat file that can be used to remove the SPNs manually. This .bat file is located at %windir%\Debug*domain-servicename*.bat, where *domain* represents the Active Directory domain in which the ADAM instance resides, and *servicename* represents the service name of the ADAM instance being removed. (By default, the service name of an ADAM instance is ADAM_*instancename*, where *instancename* is the instance name that you provide during ADAM setup.)

**See Also**
**Concepts**
Administering ADAM service principal names
Manually register ADAM service principal names in Active Directory
Repadmin
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit

# Manage Replication, Sites, and Configuration Sets

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage replication, sites, and configuration sets**

- Configure replication frequency within a site
- Configure replication availability on a site link
- Add a site object and servers container
- Move an ADAM instance to a site object
- Add a site link object
- Add a site object to a site link object
- Set the cost on a site link object
- Delete a site object
- Identify or transfer the naming master role
- Modify the replication security level of a configuration set

# Configure replication frequency within a site

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To configure replication frequency within a site**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of the ADAM instance for which you want to configure the replication schedule.
3. In the console tree, double-click the connection, double-click the configuration partition, double-click the sites container (CN=Sites), and then click the site container in which the instance resides.
4. In the details pane, right-click **CN=NTDS Site Settings**, and then click **Schedule**.
5. In the **Schedule** dialog box, select the block of time that you want to schedule, and then click **None**, **Once per Hour**, **Twice per Hour**, or **Four Times per Hour** as the replication frequency.
6. When you are finished updating the schedule, click **OK**.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to a directory partition on an ADAM instance using ADAM ADSI Edit, see Related Topics.
- The replication frequency schedule ensures that ADAM instances within a site periodically check with one another to determine if any changes have occurred. Regardless of the replication frequency that is specified, ADAM instances continue to replicate changes through update notifications. For more information, see Related Topics.

**See Also**
**Concepts**
Understanding ADAM sites and replication
Administering replication and configuration sets
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit

# Configure replication availability on a site link

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To configure replication availability on a site link**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of an ADAM instance belonging to the site link for which you want to configure replication availability.
3. In the console tree, click the configuration directory partition, double-click the sites container (CN=Sites), double-click the Inter-site Transports container (CN=Inter-site Transports), and then click the IP container (CN=IP).
4. In the details pane, right-click the site link object (by default, CN=DEFAULTIPSITELINK) for which you want to schedule availability, and then click **Schedule**.
5. In the **Schedule** dialog box, select the block of time that you want to schedule, and then click **Replication Not Available** or **Replication Available**.
6. When you are finished updating the schedule, click **OK**.

## Notes

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to the configuration partition of an ADAM instance using ADAM ADSI Edit, see Related Topics.
- By default, replication is always available on a given site link.
- By default, all sites belong to the default site link object (CN=DEFAULTIPSITELINK).

**See Also**
**Concepts**
Understanding ADAM sites and replication
Administering replication and configuration sets
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit

# Add a site object and servers container

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add a site object and servers container**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of an ADAM instance in the configuration set for which you want to create a site.
3. In the console tree, double-click the configuration directory partition (CN=Configuration,CN={*GUID*}), right-click the sites container (CN=Sites), point to **New**, and then click **Object**.
4. In **Select a class**, click **site**, and then click **Next**.
5. In **Value**, type a common name (CN) for the new site, and then click **Next**.
6. To set values for additional attributes, click **More attributes**, set any additional attribute values, and then click **Finish**.
7. In the console tree, double-click the sites container (CN=Sites).
8. In the console tree, right-click the new site, point to **New**, and then click **Object**.
9. In **Select a class**, click **serversContainer**, and then click **Next**.
10. In **Value**, type a common name (CN) for the server container, and then click **Next**.
11. To set values for additional attributes, click **More attributes**, and then set any additional attribute values.
12. After setting all desired attributes for the servers container, click **Finish**.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.
- To add ADAM instances to a site, the site must have a servers container into which the ADAM instances can be added.

**See Also**
**Concepts**
Understanding ADAM sites
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
Administering sites
Move an ADAM instance to a site object
Add a site object to a site link object

# Move an ADAM instance to a site object

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To move an ADAM instance to a site object**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of an ADAM instance in the configuration set that contains the site object.
3. In the console tree, double-click the connection, double-click the configuration directory partition, and then double-click the sites container (CN=Sites).
4. In the sites container, double-click the site containing the ADAM instance that you want to move, and then double-click the servers container of the site.
5. In the servers container, right-click the ADAM instance that you want to move, and then click **Move**.
6. In the **Move** dialog box, double-click **Sites**, double-click the site to which you want to move the ADAM instance, click the servers container of the site, and then click **OK**.

## Notes

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.
- Before you can move an ADAM instance to a site, the site must have a servers container to which you can move the ADAM instance.
- The ADAM instance does not appear in its new location until you right-click the destination servers container and then click **Refresh**.

**See Also**
**Concepts**
Understanding ADAM sites
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
Add a site object and servers container

# Add a site link object

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add a site link object**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of an ADAM instance in the configuration set for which you want to create a site link.
3. In the console tree, double-click the configuration directory partition, double-click the sites container (CN=Sites), and then double-click the Inter-Site Transports container (CN=Inter-Site Transports).
4. In the Inter-Site Transports container, right-click the IP container (CN=IP), point to **New**, and then click **Object**.
5. In **Select a class**, click **siteLink**, and then click **Next**.
6. In **Value**, type a name for the common name (CN) of the new site link object, and then click **Next**.
7. In **Value**, type a distinguished name of a site to be placed in the site link object, and then click **Next**.
8. To set more attributes on the site link object, click **More attributes**.
9. To create the new site link object, click **Finish**.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.

**See Also**
**Concepts**
Understanding ADAM sites
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
Administering sites

# Add a site object to a site link object

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add a site object to a site link object**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of an ADAM instance in the configuration set in which the site and the site link reside.
3. In the console tree, double-click the configuration directory partition, double-click the sites container (CN=Sites), and then double-click the Inter-Site Transports container (CN=Inter-Site Transports).
4. In the console tree, click the IP container (CN=IP).
5. In the details pane, right-click the site link to which you want to add a site, and then click **Properties**.
6. In **Attributes**, click **SiteList**, and then click **Edit**.
7. In **Value to add**, type the distinguished name of the site that you want to add to the site link, click **Add**, and then click **OK**.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.

**See Also**
**Concepts**
Understanding ADAM sites
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
Administering sites

# Set the cost on a site link object

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To set the cost on a site link object**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of an ADAM instance in the configuration set in which the site link object that you want to modify resides.
3. In the console tree, double-click the configuration directory partition, double-click the sites container (CN=Sites), and then double-click the Inter-Site Transports container (CN=Inter-Site Transports).
4. In the console tree, click the IP container (CN=IP).
5. In the details pane, right-click the site link for which you want to set the cost, and then click **Properties**.
6. In the **Properties** dialog box, in **Attributes**, click **cost**, and then click **Edit**.
7. In **Value**, type a cost for the site link.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.
- The default cost for a site link is 100.
- The cost of a site link represents the relative cost of using the site link for replication. The higher the cost of a site link, the lower the preference in ADAM for using the site link for replication.

**See Also**
**Concepts**
Understanding ADAM sites
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
Administering sites
Add a site link object

# Delete a site object

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To delete a site object**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of an ADAM instance in the configuration set in which the site object resides.
3. In the console tree, double-click the configuration directory partition, double-click the sites container (CN=Sites), right-click the site that you want to delete, and then click **Delete**.

### Notes

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.

**See Also**
**Concepts**
Understanding ADAM sites
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
Understanding ADAM sites and replication
Administering sites

# Identify or transfer the naming master role

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To identify or transfer the naming master role**

1. Open an ADAM tools command prompt.
2. At the command prompt, type:

   **dsmgmt**
3. At the **dsmgmt:** command prompt, type:

   **roles**
4. At the **fsmo maintenance:** command prompt, type:

   **connections**
5. At the **server connections:** command prompt, type:

   **connect to server** *servername:portnumber*

   where *servername:portnumber* is the computer name and communications port number of the ADAM instance that you want to use as the new naming master.
6. At the **server connections:** command prompt, type:

   **quit**
7. At the **fsmo maintenance:** command prompt, type:

   **transfer naming master**

**Note**

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- For help with **dsmgmt**, type **?** at a **dsmgmt** prompt.

**See Also**
**Concepts**
Administering an ADAM instance
Dsmgmt

# Modify the replication security level of a configuration set

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To modify the replication security level of a configuration set**

1. Open ADAM ADSI Edit.
2. Connect and bind to the configuration directory partition of an ADAM instance in the configuration set that you want to modify.
3. In the console tree, double-click the connection.
4. In the console tree, right-click the configuration directory partition, and then click **Properties**.
5. In **Attributes**, click **msDS-ReplAuthenticationMode**, click **Edit**, and then, in the **Value** box, set the appropriate value (2, 1, or 0).

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to the configuration directory partition of an ADAM instance using ADAM ADSI Edit, see Related Topics.
- The default replication security level for a new ADAM instance is 1, unless a local workstation user account is specified as the ADAM service account, in which case the replication security level defaults to 0.
- If the replication security level is set to 2, all replicating ADAM instances must register service principal names (SPNs) in Active Directory. For more information, see Related Topics.
- The values for **msDS-ReplAuthenticationMode** and their corresponding replication security levels are described in the following table.

| Replication security level | Value | Description | Default environment |
|---|---|---|---|
| Mutual authentication required | 2 | Kerberos authentication (using SPNs) is required. If Kerberos authentication fails, the ADAM instances will not replicate. | The configuration set is fully contained within an Active Directory domain, forest, or forest trust. |
| Negotiated | 1 | Kerberos authentication (using SPNs) is attempted first. If Kerberos fails, NTLM authentication is attempted. If NTLM fails, the ADAM instances will not replicate. | The configuration set contains Windows NT 4.0 member servers. |
| Negotiated pass-through | 0 | All ADAM instances in the configuration set must use identical service account names and passwords. | The configuration set includes computers that are joined to one or more workgroups or to multiple domains or forests without trust relationships. |

**See Also**
**Concepts**
Understanding ADAM replication and configuration sets
Browse the directory tree using ADAM ADSI Edit
Connect and bind to an ADAM instance using ADAM ADSI Edit
Administering ADAM service principal names

# Manage Schemas and Directory Partitions

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage schemas and directory partitions**

- Install the ADAM Schema snap-in
- Import or export directory objects using ldifde
- Import or export directory objects using csvde
- Manually configure the replication authentication method
- Deactivate a schema object class or attribute
- Identify or transfer the schema master role
- Seize the ADAM schema master role
- Add a new schema class or attribute definition in ADAM
- Add an application directory partition
- Delete an application directory partition

# Install the ADAM Schema snap-in

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To install the ADAM Schema snap-in**

1. Click **Start**, click **Run**, type **mmc /a**, and then click **OK**.
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. In **Available Standalone Snap-ins**, click **ADAM Schema**, click **Add**, click **Close**, and then click **OK**.
4. To save this console, on the **File** menu, click **Save**.
5. In **Save in**, point to the %systemroot%**\system32** directory.
6. In **File name**, type **adamschmmgmt.msc**, and then click **Save**.
7. To create a shortcut on your **Start** menu:
   a. Right-click **Start**, click **Open All Users**, double-click the **Programs** folder, and then double-click the **ADAM** folder.
   b. On the **File** menu, point to **New**, and then click **Shortcut**.
   c. In **Type the location of the item**, type **adamschmmgmt.msc**, and then click **Next**.
   d. In **Type a name for this shortcut**, type **ADAM Schema**, and then click **Finish**.

**Caution**

- Modifying the schema is an advanced operation that is best performed by experienced programmers and system administrators. For detailed information about modifying the schema, see the "Active Directory Programmer's Guide" at the Microsoft MSDN Web site.

# Import or export directory objects using ldifde

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To import or export directory objects using ldifde**

1. Open an ADAM tools command prompt.
2. Do one of the following:
   - To import directory objects, at the command prompt, type:

     **ldifde -i -f** *filename* **-s** *servername:port* **-m -a** *username domain password*

   - To export directory objects, at the command prompt, type:

     **ldifde -e -f** *filename* **-s** *servername:port* **-m -a** *username domain password*

**Notes**

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- The following table contains the parameters in step 2 and other commonly used **ldifde** parameters. For more information about **ldifde** parameters, type **ldifde /?** at a command prompt, or see Related Topics.

| Parameter | Description |
|---|---|
| **-i** | Perform an import. |
| **-e** | Perform an export. |
| **-f** | File to import or export. |
| **-s** | Host name and port of the ADAM instance. |
| **-m** | Ignore (that is, do not import or export) attributes that are only used by Active Directory.<br>You can use this parameter when exporting directory objects from an existing Active Directory forest and importing them into ADAM. |
| **-a** | Account credentials. If they are not provided, **ldifde** uses the credentials of the currently logged on user. |
| **-h** | Allows importing of passwords using simple authentication and security layer (SASL) encryption. |
| **-c** *String1 String2* | Replaces all occurrences of *String1* with *String2*. With ADAM, you can use the constants **#schemaNamingContext** and **#configurationNamingContext** in place of the distinguished names of the schema directory partition and configuration directory partitions when replacing strings in .ldf files. |

- For specific information about importing the user class definitions that are supplied with ADAM, see Related Topics.
- For complete syntax information for the **ldifde** command, see Related Topics.

**See Also**
**Concepts**
Import the user classes supplied with ADAM

ADAM Ldifde
ADAM Csvde

# Import or export directory objects using csvde

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To import or export directory objects using csvde**

1. Open an ADAM tools command prompt.
2. Do one of the following:
   - To import directory objects, at the command prompt, type:

     **csvde -i -f** *filename* **-s** *servername:port* **-m -a** *username domain password*

   - To export directory objects, at the command prompt, type:

     **csvde -e -f** *filename* **-s** *servername:port* **-m -a** *username domain password*

## Notes

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- The following table contains the parameters in step 2 and other commonly used **csvde** parameters. For more information about **csvde** parameters, type **csvde /?** at a command prompt, or see Related Topics.

| Parameter | Description |
|---|---|
| **-i** | Perform an import. |
| **-e** | Perform an export. |
| **-f** | File to import. |
| **-s** | Host name and port of the ADAM instance. |
| **-m** | Ignore (that is, do not import or export) attributes that are only used by Active Directory.<br>You can use this parameter when exporting directory objects from an existing Active Directory forest and importing them into ADAM. |
| **-a** | Account credentials. If they are not provided, **ldifde** uses the credentials of the currently logged on user. |
| **-h** | Allows importing of passwords using simple authentication security layer (SASL) encryption. |
| **-c** *String1 String2* | Replaces all occurrences of *String1* with *String2*. With ADAM, you can use the constants **#schemaNamingContext** and **#configurationNamingContext** in place of the distinguished names of the schema directory partition and configuration directory partitions when replacing strings within .ldf files. |

- For specific information about importing the user class definitions that are supplied with ADAM, see Related Topics.
- For complete syntax information for the **csvde** command, see Related Topics.

**See Also**
**Concepts**
Import the user classes supplied with ADAM

ADAM Ldifde
ADAM Csvde

# Manually configure the replication authentication method

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To manually configure the replication authentication method**

1. Open ADAM ADSI Edit.
2. Connect and bind to the ADAM instance in the configuration set that you want to configure.
3. In the console tree, double-click the configuration directory partition.
4. In the console tree, right-click the configuration directory partition, and then click **Properties**.
5. In **Attributes**, click **msDS-ReplAuthenticationMode**, click **Edit**, and then do one of the following:
   - To force the use of service accounts for replication authentication between configuration set members, type **1** in the **Value** box, and then click **OK**.
   - To allow ADAM to automatically determine how replication authentication occurs between configuration set members, type **0** in the **Value** box, and then click **OK**.

### Notes

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.
- For information about how ADAM configures replication authentication between configuration set members, see Related Topics.

**See Also**
**Concepts**
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
Administering replication and configuration sets

# Deactivate a schema object class or attribute

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To deactivate a schema object class or attribute**

1. Open ADAM Schema.
2. In the console tree, click **ADAM Schema**.
3. On the **Action** menu, click **Change ADAM Server**.
4. In **ADAM server**, specify the DNS name, network basic input/output system (NetBIOS) name, or Internet Protocol (IP) address of the computer on which the ADAM instance that you want to configure is running.
5. In **Port**, type the Lightweight Directory Access Protocol (LDAP) or Secure Sockets Layer (SSL) port number of the ADAM instance to which you want to connect.
6. Under **Connect using these credentials**, do one of the following, and then click **OK**:
   - Click **The account of the currently logged on user**.
   - Click **This account**, and then select or type the user name and password of the account with which you want to connect.
7. Do one of the following, depending on whether you want to deactivate an object class or an attribute:
   - To deactivate an object class, in the console tree, double-click **ADAM Schema**, and then click **Classes**. In the details pane, right-click the class that you want to deactivate, and then click **Properties**.
   - To deactivate an attribute, in the console tree, double-click **ADAM Schema**, and then click **Attributes**. In the details pane, right-click the attribute that you want to deactivate, and then click **Properties**.
8. On the **General** tab, clear the **Class is active** check box or the **Attribute is active** check box, click **Yes** to confirm that you want to deactivate the class or attribute, and then click **OK**.

### Notes

- To open the ADAM Schema snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Schema**.
- If the ADAM Schema snap-in is not installed, see Related Topics.

**See Also**
**Concepts**
Install the ADAM Schema snap-in
Understanding ADAM schema
Administering the ADAM schema

# Identify or transfer the schema master role

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To identify or transfer the schema master role**

1. Open ADAM Schema.
2. On the **Action** menu, click **Change ADAM Server**.
3. In **ADAM server**, specify the DNS name, network basic input/output system (NetBIOS) name, or Internet Protocol (IP) address of an ADAM instance in the configuration set for which you want to identify or transfer the schema master role.
4. In **Port**, type the Lightweight Directory Access Protocol (LDAP) port number of the ADAM instance to which you want to connect.
5. Under **Connect using these credentials**, do one of the following, and then click **OK**:
   - Click **The account of the currently logged on user**.
   - Click **This account**, and then select or type the user name and password of the account with which you want to connect.
6. In the console tree, right-click **ADAM Schema**, and then click **Operations Master**.
7. Do one of the following:
   - To identify the current schema master, note the ADAM instance that is listed in **Current Schema Master**.
   - To transfer the schema master role to the currently targeted ADAM instance, click **Change**.

## Notes

- To open the ADAM Schema snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Schema**. If the ADAM Schema snap-in is not installed, see Related Topics.
- Each configuration set has only one schema master. To identify the schema master in a different configuration set, target the appropriate configuration set before clicking **Operations Master**.
- For help with **dsmgmt**, type **?** at a **dsmgmt** prompt.

**See Also**
**Concepts**
Install the ADAM Schema snap-in
Understanding ADAM schema
Dsmgmt
Administering an ADAM instance
Administering the ADAM schema

# Seize the ADAM schema master role

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To seize the schema master role**

1. Open an ADAM tools command prompt.
2. At the command prompt, type:

   **dsmgmt**
3. At the **dsmgmt:** command prompt, type:

   **roles**
4. At the **fsmo maintenance:** command prompt, type:

   **connections**
5. At the **server connections:** command prompt, type:

   **connect to server** *computername:portnumber*

   where *computername:portnumber* is the computer name and communications port number of the ADAM instance that you want to use as the new schema master.
6. At the **server connections:** command prompt, type:

   **quit**
7. At the **fsmo maintenance:** command prompt, type:

   **seize schema master**

**Caution**

- Do not seize the schema master role if you can transfer it instead. Seizing the schema master role is a drastic step that should be considered only if the current operations master will never be available again. For more information about transferring operations master roles, see Related Topics.

**Note**

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- For help with **dsmgmt**, type **?** at a **dsmgmt** prompt.

**See Also**
**Concepts**
Install the ADAM Schema snap-in
Understanding ADAM schema
Dsmgmt
Administering an ADAM instance
Identify or transfer the schema master role
Administering the ADAM schema

# Add a new schema class or attribute definition in ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add a new schema class or attribute definition**

1. Open ADAM Schema.
2. In the console tree, double-click **ADAM Schema**.
3. Do one of the following:
   - To add a class definition, in the console tree, right-click **Classes**, click **Create Class**, and then provide the requested information.
   - To add an attribute definition, in the console tree, right-click **Attributes**, click **Create Attribute**, and then provide the requested information.

## Notes

- To open the ADAM Schema snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Schema**.
- If the ADAM Schema snap-in is not installed, see Related Topics.

**See Also**
**Concepts**
Install the ADAM Schema snap-in
Understanding ADAM schema

# Add an application directory partition

Updated: March 1, 2013

Applies To: Windows Server 2003 R2

**To add an application directory partition**

1. Connect and bind to an ADAM instance with ADAM Ldp.exe.
2. On the **Browse** menu, click **Add child**.
3. In **Dn**, type a distinguished name for the application partition.
4. Under **Edit entry**, type **ObjectClass** in the **Attribute** box and **container** in the **Values** box, and then click **Enter**.
5. Under **Edit entry**, type **instanceType** in the **Attribute** box and **5** in the **Values** box, and then click **Enter**.
6. Click **Run**.

   After the new application directory partition is added, the following appears in the results pane:

   Added {*distinguished name*}

   where *distinguished name* is the distinguished name that you typed in step 3.
7. Click **Close**.

**Notes**

- To connect and bind to an ADAM instance with Ldp.exe, see Related Topics.
- Application directory partitions can be of any object class, but they must have a value of **5** for the **instanceType** attribute.

**See Also**
**Concepts**
Connect and bind to an ADAM instance using Ldp.exe
Understanding ADAM data and data stores
Ldp

# Delete an application directory partition

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To delete an application directory partition**

1. Open ADAM ADSI Edit.
2. Connect and bind to the ADAM instance.
3. Double-click the configuration directory partition (CN=Configuration,CN={*GUID*}).
4. In the console tree, click the partitions container (CN=Partitions).
5. In the details pane, right-click the application directory partition that you want to delete, and then click **Delete**.

### Caution

- The only way to retrieve a deleted application directory partition is to restore the entire ADAM instance from a backup copy that was made before the deletion of the partition.

### Notes

- To start the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.

**See Also**
**Concepts**
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit

# Manage Users and Groups

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage users and groups**

- Import the user classes supplied with ADAM
- Add an ADAM user to the directory
- Add an ADAM group to the directory
- Add or remove members to or from an ADAM group
- Disable or enable an ADAM user
- Set or modify the password of an ADAM user
- Add an organizational unit to the directory
- Allow ADAM users to be created in the configuration partition

# Import the user classes supplied with ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To import the user classes supplied with ADAM**

1. Open an ADAM tools command prompt.
2. At the command prompt, type:

   **ldifde -i -f** *filename* **-s** *computername:port* [**-b** *username domain password*] **-k -j . -c "CN=Schema,CN=Configuration,DC=X" #schemaNamingContext**

   where *filename* represents the name of one of the .ldf files that is supplied with ADAM (ms-User.ldf, ms-InetOrgPerson.ldf, or ms-UserProxy.ldf), *computername:port* represents the computer name and port number of an ADAM instance, and *username domain password* represents the account with which to run the command.

   You should see results similar to the following:

```
Connecting to "localhost"
Logging in as current user using SSPI
Importing directory from file "ms-user.ldf"
Loading entries.............................
65 entries modified successfully.

The command has completed successfully
```

## Notes

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- The following table contains the parameters in step 2 and other commonly used **ldifde** parameters. For more information about **ldifde** parameters, type **ldifde /?** at a command prompt, or see Related Topics.

| Parameter | Description |
|---|---|
| **-i** | Perform an import. |
| **-f** | File to import or export. |
| **-s** | Host name and port of the ADAM instance. |
| **-b** | Security credentials to use during the operation. |
| **-k** | Continue the operation in the event of errors. |
| **-j** | Create a log file in the specified directory. In this case, the current ("**.**") directory. |
| **-c "CN=Schema,CN=Configuration,DC=X" #schemaNamingContext** | Do not modify this string.<br>The **-c** parameter replaces a specified string in the .ldf file with a different string during import. The distinguished name that is specified in the .ldf file (for example, **CN=Schema,CN=Configuration,DC=X**) is replaced with the distinguished name of the schema directory partition for your particular ADAM instance, as passed by the **#schemaNamingContext** constant. |

| Replace ... | With ... |
| --- | --- |
| *computername:port* | The host name and port number of your ADAM instance. You can use "localhost" for *computername* if you are administering the local computer. |
| *username domain password* | The account, domain, and password of an account being used as an ADAM administrator account. |

- ADAM includes three .ldf files containing user classes that you can import. These files are ms-User.ldf, ms-InetOrgPerson.ldf, and ms-UserProxy.ldf, and they are located in the %windir%\adam directory.
- As an alternative to using **ldifde**, you can import the optional ADAM user classes during ADAM setup. For more information, see Related Topics.
- If you do not specify user credentials using the **-b** parameter, **ldifde** uses the credentials of the currently logged on user.
- For complete syntax information for the **ldifde** command, see Related Topics.

**See Also**
**Concepts**
Understanding ADAM users and groups
ADAM Ldifde
Importing LDIF files during setup

# Add an ADAM user to the directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add an ADAM user to the directory**

1. Open ADAM ADSI Edit.
2. Connect and bind to the ADAM instance and directory partition to which you want to add the user.
3. In the console tree, double-click the directory partition to which you want to add the user.
4. In the console tree, right-click the container to which you want to add a user, point to **New**, and then click **Object**.
5. In **Select a class**, click the class that you want to use (**user**, **inetOrgPerson**, **person**, or **OrganizationalPerson**), and then click **Next**.
6. In **Value**, type a value for the common name (CN) attribute of the new user, and then click **Next**.
7. If you want to set values for additional attributes, click **More attributes**.
8. After setting any additional attributes for the new user, click **Finish**.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.
- By default, an ADAM user account is enabled when the user account is created. However, no initial password is set on an ADAM user account that is created using ADAM ADSI Edit. On ADAM instances running on Windows Server 2003, where local or domain password policy restrictions are in effect, the ADAM user account is disabled by default. Before you can enable the user account, you must set a password for it that meets the password policy restrictions that are in effect. This does not apply to ADAM instances running on Windows XP Professional. For more information, see Related Topics.
- Any object class can be used as a security principal in ADAM, if the object class definition contains the **SecurityPrincipal** static auxiliary class and the **unicodePwd** attribute. For more information, see Related Topics.
- The **user**, **inetOrgPerson**, and **OrganizationalPerson** object classes are not available until you import the ADAM user class definitions into the schema. For more information, see Related Topics.

**See Also**
**Concepts**
Import the user classes supplied with ADAM
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
Administering users and groups
Set or modify the password of an ADAM user

# Add an ADAM group to the directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add an ADAM group to the directory**

1. Open ADAM ADSI Edit.
2. Connect and bind to the ADAM instance to which you want to add a group.
3. In the console tree, double-click the directory partition to which you want to add the group.
4. In the console tree, right-click the container to which you want to add the group, point to **New**, and then click **Object**.
5. In **Select a class**, click **Group**, and then click **Next**.
6. In **Value**, type a common name (CN) for the new group, and then click **Next**.
7. In **groupType**, type **2147483650** (representing an account group) for the new group, and then click **Next**.
8. If you want to set values for additional attributes, click **More attributes**.
9. After setting all desired attributes for the new group, click **Finish**.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- When you type a value for the **groupType** attribute, 2147483650 (equivalent to 0x80000002 hexidecimal) represents a group type of "account."
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.

**See Also**
**Concepts**
Understanding ADAM users and groups
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit

# Add or remove members to or from an ADAM group

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add or remove members to or from an ADAM group**

1. Open ADAM ADSI Edit.
2. Connect and bind to the ADAM instance containing the group that you want to modify.
3. In the console tree, double-click the directory partition containing the group that you want to modify.
4. Right-click the group that you want to modify, and then click **Properties**.
5. In **Attributes**, click **Member**, and then click **Edit**.
6. For each ADAM security principal that you want to add to the group, click **Add ADAM account**, type the distinguished name of the new member, and then click **OK**.
7. For each Windows security principal you want to add to the group, click **Add Windows account**, type the account name of the new member, and then click **OK**.
8. For each group member that you want to remove from the group, click the member that you want to remove, and then click **Remove**.
9. After making the changes that you want to the group, click **OK** twice.

### Notes

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.

**See Also**
**Concepts**
Understanding ADAM users and groups
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
ADAM ADSI Edit

# Disable or enable an ADAM user

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To disable or enable an ADAM user**

1. Open ADAM ADSI Edit.
2. Connect and bind to an ADAM instance.
3. Browse to the ADAM user that you want to disable or enable, right-click that user, and then click **Properties**.
4. In **Attributes**, click **msDS-UserAccountDisabled**, and then click **Edit**.
5. Do one of the following, and then click **OK**:
   - To disable the ADAM user, click **True**.
   - To enable the ADAM user, click either **False** or **Not set**.

## Notes

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- To connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.
- By default, an ADAM user is enabled when the user is created. However, if you assign a new ADAM user a password that does not meet the password policy restrictions in effect on the local server or domain, that ADAM user will be disabled by default. This note does not apply to ADAM instances running on Windows XP Professional. For more information, see Related Topics.
- If the ADAM user that you want to disable or enable is currently logged on to the ADAM instance, that user must log off for the new setting to take effect.

**See Also**
**Concepts**
Understanding ADAM users and groups
Import the user classes supplied with ADAM
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit
Administering users and groups

# Set or modify the password of an ADAM user

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Set or modify the password of an ADAM user**

- Using ADAM ADSI Edit
- Using Ldp over an encrypted, non-SSL connection
- Using Ldp over an SSL connection

**Using ADAM ADSI Edit**
**To set or modify the password of an ADAM user**

1. Open ADAM ADSI Edit.
2. Connect and bind to the directory partition containing the ADAM user for whom you want to set or modify the password.

   To complete this task, you can bind to the directory partition as an ADAM user or as a Windows Security principal.

   To bind to the directory partition as an ADAM user, complete the following:

   - Click **Start**, point to **Administrative Tools**, and then click **ADSI Edit**.
   - On the **Action** menu, click **Connect to**. The **Connection Settings** dialog box appears.
   - In **Select or type a domain or server: (Server | Domain[:port])**, type the Domain Name System (DNS) name, NetBIOS name, or IP address of the computer on which the ADAM instance is running, followed by a colon (:) and the LDAP communication port that the ADAM instance to which you want to connect is using.For example, if ADAM is running on the local computer, you can type **localhost:389**.
   - Under **Connection point**, click **Select or type a Distinguished Name or Naming Context**, and then type the name of your directory partition.
   - Click **Advanced**, and then click **Specify Credentials**.
   - Under **Connect using these credentials**, type the DN (distinguished name) and the password of your ADAM user, make sure **Simple bind authentication** checkbox is checked, and then click OK.

     | **Note** |
     | --- |
     | Make sure that your ADAM user is enabled, by setting its MSds-UserAccountDisabled attribute to FALSE. |

   To bind to the directory partition as a Windows security principal, complete the following:

   - Click **Start**, point to **Administrative Tools**, and then click **ADSI Edit**.
   - On the **Action** menu, click **Connect to**. The **Connection Settings** dialog box appears.
   - In **Select or type a domain or server: (Server | Domain[:port])**, type the Domain Name System (DNS) name, NetBIOS name, or IP address of the computer on which the ADAM instance is running, followed by a colon (:) and the LDAP communication port that the ADAM instance to which you want to connect is using.For example, if ADAM is running on the local computer, you can type **localhost:389**.
   - Under **Connection point**, click **Select or type a Distinguished Name or Naming Context**, and then type the

name of your directory partition.

- Click **Advanced**, and then click **Specify Credentials**.
- Under **Connect using these credentials**, type the domain, user name, and password of your Windows principal, and then click OK.

3. Browse to the directory object representing the ADAM user, and then right-click the directory object.
4. Click **Reset password**, and then type a password for the user in **New password** and in **Confirm password**.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- For information about how to connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.

**Using Ldp over an encrypted, non-SSL connection**
**To set or modify the password of an ADAM user**

1. Open Ldp.
2. On the **Options** menu, click **Connection Options**.
3. In **Option Name**, click **LDAP_OPT_ENCRYPT**.
4. In **Value**, type **1**, click **Set**, and then click **Close**.
5. Connect and bind to the ADAM instance as a Windows security principal, and then view the directory partition containing the ADAM user for whom you want to set a password.

   To connect and bind to the ADAM instance as a Windows security principal, complete the following:

   - On the **Connection** menu, click **Connect**. In **Server**, type the DNS name, NetBIOS name, or IP address of the computer on which the ADAM instance is running, and in **Port**, type the LDAP communication port that the ADAM instance to which you want to connect is using.
   - On the **Connection** menu, click **Bind**.
   - Under **Bind type**, select **Bind with credentials**. Type the name of your Windows security principal in the **User** field, along with the password of this account in the **Password** field, and the domain name that this account belongs to in the **Domain** field, and then click OK.

6. Right-click the ADAM user, and then click **Modify**.
7. In **Attribute**, type **userpassword**, and then in **Value**, type a password for the account.
8. Click **Enter**, and then click **Run**. The details pane displays a message similar to the following:

```
***Call Modify...
ldap_modify_s(ld, 'CN=Mary Baker,O=Microsoft,C=US',[1] attrs);
Modified "CN=Mary Baker,O=Microsoft,C=US".
```

**Notes**

- To open Ldp, click **Start**, point to **All Programs**, point to **ADAM**, click **ADAM Tools Command Prompt**, and then type **ldp** at the command prompt.
- You can also use ADAM ADSI Edit to set or modify passwords: right-click the directory object representing the ADAM security principal in ADAM ADSI Edit, and then click **Reset Password**.

- For information about how to connect and bind to an ADAM instance using Ldp, see Related Topics.

- By default, an ADAM instance running on Windows Server 2003 automatically enforces any local or domain password policies. If you set a password for an ADAM user that does not meet the requirements of the password policy in effect, the user account will be disabled.

- The ADAM user for whom you set or modify the password must use the new password the next time that the user logs on.

- This procedure applies to any object class that is used as a security principal in ADAM. Any object class in ADAM can be used as a security principal, as long as the object class definition contains the **msDS-bindableobject** auxiliary class and the **unicodePwd** attribute. For more information, see Related Topics.

- The **user**, **person**, **inetOrgPerson**, and **OrganizationalPerson** object classes are not available by default in the ADAM schema. You must first import them. For more information, see Related Topics.

**Using Ldp over an SSL connection**
**To set or modify the password of an ADAM user**

1. Install a server certificate on the computer where the ADAM instance is running, and install a matching client certificate on the computer from which you administer the ADAM instance.

2. Open Ldp.

3. Connect and bind to the ADAM instance (selecting **SSL** in the **Connect** dialog box) containing the user for whom you want to set or modify a password. For more information, see Appendix A: Configuring LDAP over SSL Requirements for AD LDS (http://go.microsoft.com/fwlink/?LinkId=160191).

   To complete this procedure, connect and bind to the ADAM instance as an ADAM user by completing the following:

   - On the **Connection** menu, click **Connect**. In **Server**, type the DNS name, NetBIOS name, or IP address of the computer on which the ADAM instance is running, and in **Port**, type the SSL communication port that the ADAM instance to which you want to connect is using.

   - On the **Connection** menu, click **Bind**.

   - Under **Bind type**, select **Simple Bind**, type the DN of your ADAM user in the **User** field, along with the password that you just assigned to this account in the **Password** field, and then click OK.

4. Right-click the ADAM user, and then click **Modify**.

5. In **Attribute**, type **userpassword**, and then in **Value**, type a password for the account.

6. Click **Enter**, and then click **Run**. The details pane displays a message similar to the following:

```
***Call Modify...
ldap_modify_s(ld, 'CN=Mary Baker,OU=Beta
users,O=Microsoft,C=US',[1] attrs);
Modified "CN=Mary Baker,OU=Beta users,O=Microsoft,C=US".
```

### Notes

- To open Ldp, click **Start**, point to **All Programs**, point to **ADAM**, click **ADAM Tools Command Prompt**, and then type **ldp** at the command prompt.

- For information about how to connect and bind to an ADAM instance using Ldp, see Related Topics.

- Establishing Secure Sockets Layer (SSL) connections requires the presence of certificates on the server and clients. For information about installing certificates, see Related Topics.

- By default, an ADAM instance running on Windows Server 2003 automatically enforces any local or domain password policies. If you set a password for an ADAM user that does not meet the requirements of the password policy in effect, the user account will be disabled.

- If the ADAM user is currently logged on, that user must log off for the new password to take effect.

- This procedure applies to any object class that is used as a security principal in ADAM. Any object class in ADAM can be used as a security principal, as long as the object class definition contains the **SecurityPrincipal** static auxiliary class and the **unicodePwd** attribute. For more information, see Related Topics.

- The **user**, **person**, **inetOrgPerson**, and **OrganizationalPerson** object types are not available by default in the ADAM schema. You must first import them. For more information, see Related Topics.

**See Also**
**Concepts**
Understanding ADAM users and groups
Import the user classes supplied with ADAM
Administering an ADAM instance
Connect and bind to an ADAM instance using Ldp.exe
Connect and bind to an ADAM instance using ADAM ADSI Edit

# Add an organizational unit to the directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add an organizational unit to the directory**

1. Open ADAM ADSI Edit.
2. Connect and bind to the directory partition of the ADAM instance to which you want to add an organizational unit (OU).
3. In the console tree, double-click the directory partition, right-click the container to which you want to add the OU, point to **New**, and then click **Object**.
4. In **Select a class**, click **organizationalUnit**, and then click **Next**.
5. In **Value**, type a name for the new OU, and then click **Next**.
6. If you want to set values for additional attributes, click **More attributes**.
7. After you set all desired attributes for the new OU, click **Finish**.

**Notes**

- To open the ADAM ADSI Edit snap-in, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
- To connect and bind to an ADAM instance using ADAM ADSI Edit, see Related Topics.
- By default, OUs can only be added under OU (OU=), country/region (C=), organization (O=) or domain-DNS (DC=) object classes. For example, you can add an OU to o=Microsoft,c=US. You cannot add an OU to cn=test,o=Microsoft,c=US. You can, however, update the schema definition of the OU object class to allow other superiors.

**See Also**
**Concepts**
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit

# Allow ADAM users to be created in the configuration partition

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use this procedure to allow users to be created in the configuration partition so that Active Directory Application Mode (ADAM) users can be ADAM administrators.

**To allow ADAM users to be created in the configuration partition**

1. Click **Start**, point to **Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.
2. Double-click the CN=Configuration,CN=Guid container.
3. Double-click the CN=Services object, and then double-click the CN=**Windows NT** object.
4. Right-click CN=**Directory Service**, and then click **Properties**.
5. In the **Attributes** list, click **msDS-Other-Settings**, and then click **Edit**.
6. In the **Values** list, click each value that contains the text **ADAMAllowADAMSecurityPrincipalsInConfigPartition=0**, and then click **Remove**.
7. In the **Value to add box**, type the following:

   **ADAMAllowADAMSecurityPrincipalsInConfigPartition=1**, and then click **Add**.
8. Click **OK** two times, and then close ADAM ADSI Edit.

# Manage Access Control

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage access control**

- View or set permissions on a directory object

# View or set permissions on a directory object

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To view or set permissions on a directory object**

1. Open an ADAM tools command prompt.
2. At the command prompt, do one of the following:
   - To list the effective permissions on a directory object, type:

     **dsacls ***hostname:portnumber\object_dn*

     where *hostname* represents the name of the computer on which the ADAM instance that holds the directory object is running, *portnumber* represents the communications port number on which the ADAM instance communicates, and *object_dn* represents the distinguished name of the directory object.

     Example:

     ```
     dsacls \\localhost:389\O=Microsoft,C=US
     ```

   - To grant permissions on a directory object, type:

     **dsacls ***hostname:portnumber******object_dn* **/G** *user_or_group***:***Permissions*

     where *hostname* represents the name of the computer on which the ADAM instance that holds the directory object is running, *portnumber* represents the communications port number on which the ADAM instance communicates, *object_dn* represents the distinguished name of the directory object, *user_or_group* represents the user or group for whom the permissions apply, and *Permissions* represents the permissions to grant.

     Example:

     ```
     dsacls \\localhost:389\cn=Object1, cn=container1,O=Microsoft,C=US /G
     CN=inetuser1,O=Microsoft,C=US:SD
     ```

   - To deny permissions on a directory object, type:

     **dsacls ***hostname:portnumber******object_dn* **/D** *user_or_group***:***PermissionStatement*

     where *hostname* represents the name of the computer on which the ADAM instance that holds the directory object is running, *portnumber* represents the communications port number on which the ADAM instance communicates, *object_dn* represents the distinguished name of the directory object, *user_or_group* represents the user or group for whom the permissions apply, and *PermissionStatement* represents the permissions to deny.

     Example:

     ```
     dsacls \\localhost:389\cn=Object1, cn=container1,O=Microsoft,C=US /D
     CN=inetuser1,O=Microsoft,C=US:SD
     ```

**Notes**

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- For a complete description of all parameters that apply to **dsacls**, including the setting of inheritance, type **dsacls /?** at the ADAM tools command prompt, or see Related Topics.

- A directory object residing on multiple replicas of a given directory partition possesses the same permissions on all of the replica partitions.

**See Also**
**Concepts**
Understanding ADAM access control
Understanding ADAM users and groups
Dsacls

# Monitor and Maintain ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Monitor and maintain ADAM**

- Back up an ADAM instance to a file or a tape
- Restore an ADAM instance
- View an ADAM event log

# Back up an ADAM instance to a file or a tape

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Back up an ADAM instance to a file or a tape**

- Using the graphical user interface
- Using the command prompt

**Using the graphical user interface**
**To back up an ADAM instance to a file or a tape**

1. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use this wizard, or you can go to the next step to work in **Advanced Mode**.

2. In the Backup or Restore Wizard, click the link for **Advanced Mode**.
3. Click the **Backup** tab, and then, on the **Job** menu, click **New**.
4. To select the ADAM instance folders that you want to back up, select the check box to the left of the folders. By default, ADAM database and log files are located in %ProgramFiles%\Microsoft ADAM*instancename* where *instancename* equals the ADAM instance name.

   | **Important** |
   | --- |
   | If you cannot select adamntds.dit file for backup, open the registry editor and then delete "ADAM (*instanceName*)Write" key value from the following registry key on the ADAM server:HKEY_LOCAL_MACHINE\CurrentControlSet\Control\BackupRestore\FilesNotToBackup |

5. In **Backup destination**, do one of the following:

   - If you want to back up files and folders to a file, click **File**.
   - If you want to back up files and folders to a tape, select a tape device.

   If no tape drive exists, the **Backup destination** option is unavailable, and it is automatically set to **File**.

6. In **Backup media or file name**, do one of the following:

   - If you are backing up files and folders to a file, type a path and file name for the backup (.bkf) file, or click the **Browse** button to find a file.
   - If you are backing up files and folders to a tape, select the tape that you want to use.

7. To select any other backup options that you want, such as the backup type and the backup log type, on the **Tools** menu, click **Options**.
8. Click **Start Backup**, and then make any changes in the **Backup Job Information** dialog box.
9. If you want to set advanced backup options, such as data verification or hardware compression, click **Advanced**.
10. Click **Start Backup** to start the backup operation.

| **Caution** |
| --- |

| |
|---|
| If you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. |

| **Important** |
|---|
| Before backing up the ADAM folder of an instance, use Services to confirm that the instance is started. To start Services, click **Start**, point to **Administrative Tools**, and then click **Services**. |

**Notes**

- To start Backup, click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
- If Backup does not start in wizard mode by default, you can start the wizard by clicking **Wizard Mode** on the **Welcome** tab.
- Some tape drives may not support hardware compression.
- Backup does not support the use of backup media such as CD-RW (compact disc rewritable), CD-R (compact disc recordable), and DVD-R (digital video disc recordable). To save backups to this type of media, back up to a file, and then copy that file to the CD. You can use Backup to restore from a CD.
- Backup files usually have the file name extension .bkf. You can use any extension you like, but it is highly recommended that you use .bkf, which has file associations that ensure that the backup file is recognizable.
- Backup operators and administrators can back up and restore encrypted files and folders without decrypting the files or folders.

**Using the command prompt**
**To back up an ADAM instance to a file or a tape**

1. Open Command Prompt.
2. To back up to a file or tape, type:

   **ntbackup backup** *@bks file name* **/J** {"*backup job name*"} **/F** {"*file name*"} **/T** {"*tape name*"} **/P** {"*pool name*"} **/G** {"*guid name*"} **/N** {"*media name*"} **/A**

| Value | Description |
|---|---|
| *@bks file name* | The name of the backup selection file (.bks file) to be used for this backup operation. The @ character must precede the name of the backup selection file. A backup selection file contains information about the files and folders that you have selected for backup. You have to create the file using the graphical user interface (GUI) version of Backup. As an alternative, you can supply the path to the drive or file to back up, for example, %ProgramFiles%\Microsoft ADAM\instance1. |
| **/J** {"*backup job name*"} | The job name to be used in the backup report. The job name usually describes the files and folders that you are backing up in the current backup job. |
| **/F** {"*file name*"} | The logical disk path and file name of the backup file. Do not use in conjunction with **/P /G /T**. |
| **/T** {"*tape name*"} | Overwrites or appends to this tape. |
| **/P** {"*pool name*"} | The media pool from which you want to use media. This is usually a subpool of the Backup media pool, such as **4mm DDS**. Do not use in conjunction with **/A /G /F /T** |
| **/G** {"*guid name*"} | Overwrites or appends to this tape. Do not use in conjunction with **/P**. |
| **/N** {"*media name*"} | The new tape name. Do not use in conjunction with **/A**. |

| | |
|---|---|
| **/A** | Performs an append operation. Either **/G** or **/T** must be used in conjunction with this switch. Do not use in conjunction with **/P**. |

For information about additional file and tape backup options, see the Help topic on the **ntbackup** command-line utility in Help and Support Center in Windows Server 2003.

**Examples:**

To back up the files and folders that are listed in the backup selection file C:\adaminstance.bks to file C:\backup.bkf, with the job name "Backup Job 1," type the following at the command line:

**ntbackup backup @C:\adaminstance.bks /j "Backup Job 1" /f "C:\Backup.bkf"**

For a tape backup, to create a backup job named "Backup Job 1" that backs up all the files and folders that are listed in the backup selection file C:\backup.bks to a tape drive, type:

**ntbackup backup @C:\backup.bks /n "Media created 11/5/2003 at 1:25 PM" /d "Set created 11/5/2003 at 1:25 PM" /j "Backup Job 1" /p "4mm DDS"**

**Important**

- Before backing up the ADAM folder of an instance, use Services to confirm that the instance is started. To start Services, click **Start**, point to **Administrative Tools**, and then click **Services**.
- If you do not specify options at the command prompt, the options default to the values that are specified in the Backup application.
- Before you use the command line to perform backup operations, it is recommended that you use the Backup GUI to create the initial backup selection file and settings and a scheduled task. You can then use that scheduled task to view the command syntax. Copy the information from the scheduled task, make any changes that are necessary to create a script or command-line syntax, and then delete the scheduled task if it is not needed.
- Some tape drives may not support hardware compression.

**Notes**

- To view the complete syntax for this command, at a command prompt, type:

  **ntbackup /?**

- If you do not specify the other Backup options, **ntbackup** uses the default values of the Backup program for the backup type, verification setting, logging level, hardware compression, and any other settings.
- Backup files usually have the file name extension .bkf. You can use any extension that you want, but it is highly recommended that you use .bkf, which has file associations that ensure that the backup file is recognizable.
- Backup operators and administrators can back up and restore encrypted files and folders without decrypting the files or folders.

# Restore an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Restore an ADAM instance**

- Restore ADAM data on an ADAM instance that belongs to a configuration set
- Restore a corrupted database that is preventing the ADAM instance from starting
- Restore an ADAM instance that does not belong to a configuration set
- Restore an ADAM instance that belongs to a configuration set

**Restore ADAM data on an ADAM instance that belongs to a configuration set**
**To authoritatively restore ADAM data on an ADAM instance that belongs to a configuration set**

1. If it is running, stop the ADAM instance for which data will be restored.
2. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use this wizard, or click the **Advanced Mode** link to work in advanced mode.
3. Click the **Restore and Manage Media** tab.
4. In the details pane, select the ADAM instance that you want to restore by selecting the check box to the left of the folder that represents the ADAM instance that you want to restore.
5. In **Restore files to**, click **Original location**.
6. On the **Tools** menu, click **Options**, click the **Restore** tab, click **Always replace the file on my computer**, and then click **OK**.
7. Click **Start Restore**.
8. After the restore is complete, close Backup.
9. Open an ADAM tools command prompt.
10. At the command prompt, type:

    **dsdbutil**
11. At the **dsdbutil:** prompt, type:

    **activate instance** *instancename*

    where *instancename* represents the service name of the ADAM instance on which you want to restore data.
12. At the **dsdbutil:** prompt, type:

    **authoritative restore**
13. At the **authoritative restore:** prompt, type one of the commands in the following table.

| Command | Meaning |
|---|---|
| **restore database** | Performs authoritative restore of the entire directory database |

| | |
|---|---|
| **restore object** *dn* | Performs authoritative restore of the directory object whose distinguished name is represented by *dn* |
| **restore subtree** *dn* | Performs authoritative restore of the directory subtree whose distinguished name is represented by *dn* |

**Notes**

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- To start Backup, click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
- If Backup does not start in wizard mode by default, you can start the wizard by clicking **Wizard Mode** on the **Welcome** tab.
- For information about stopping an ADAM instance, see Related Topics.
- This procedure can be used if you want the data being restored to authoritatively override existing data in the configuration set.
- For more information about the **authoritative restore** command in **dsdbutil**, at the **authoritative restore:** prompt, type **?**.
- When you restore ADAM instances on Windows XP Professional, you can restore only one ADAM instance at a time.
- If you accidentally start a restore of an ADAM instance over a currently running ADAM instance, it is recommended that you immediately restart the computer, stop the ADAM instance, and then perform the restoration again.

**Restore a corrupted database that is preventing the ADAM instance from starting**
**To restore a corrupted database that is preventing the ADAM instance from starting**

1. If it is running, stop the ADAM instance for which data will be restored.
2. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use this wizard, or click the **Advanced Mode** link to work in advanced mode.
3. Click the **Restore and Manage Media** tab.
4. In the details pane, select the ADAM instance that you want to restore by selecting the check box to the left of the folder that represents the ADAM instance that you want to restore.
5. In **Restore files to**, click **Original location**.
6. On the **Tools** menu, click **Options**, click the **Restore** tab, click **Always replace the file on my computer**, and then click **OK**.
7. Click **Start Restore**.
8. After the restore is complete, close Backup.
9. Restart the ADAM instance.

**Notes**

- To start Backup, click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
- If Backup does not start in wizard mode by default, you can start the wizard by clicking **Wizard Mode** on the **Welcome** tab.

- If the ADAM instance fails to start after you perform this procedure, perform the appropriate procedure for completely restoring an ADAM instance, depending on whether your ADAM instance belongs to a configuration set or is unique.

- When you restore ADAM instances to Windows XP Professional, you can restore only one ADAM instance at a time.

- If you accidentally start a restore of an ADAM instance over a currently running ADAM instance, it is recommended that you immediately restart the computer, stop the ADAM instance, and then perform the restoration again.

**Restore an ADAM instance that does not belong to a configuration set**
**To completely restore an ADAM instance that does not belong to a configuration set**

1. Using the Active Directory Application Mode Setup Wizard or an ADAM unattended install, create an ADAM instance, specifying the same settings that you used during your original ADAM installation. However, do not create an application directory partition during setup.
2. Stop the ADAM instance you just created, as follows:
   a. Click **Start**, point to **Administrative Tools**, and then click **Services**.
   b. In Services, right-click the ADAM instance, and then click **Stop**.
3. Open Backup.

   The Backup or Restore Wizard starts by default, unless it is disabled. You can use this wizard, or click the **Advanced Mode** link to work in advanced mode.
4. Click the **Restore and Manage Media** tab.
5. In the details pane, select the ADAM instance that you want to restore by selecting the check box to the left of the folder that represents the ADAM instance that you want to restore.
6. In **Restore files to**, click **Original location**.
7. On the **Tools** menu, click **Options**, click the **Restore** tab, click **Always replace the file on my computer**, and then click **OK**.
8. Click **Start Restore**.
9. After the restore is complete, close Backup.
10. Stop the ADAM instance you just created, as follows:
    a. Click **Start**, point to **Administrative Tools**, and then click **Services**.
    b. In Services, right-click the ADAM instance, and then click **Stop**.

**Notes**

- To start Backup, click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.

- If Backup does not start in wizard mode by default, you can start the wizard by clicking **Wizard Mode** on the **Welcome** tab.

- For information about backing up an ADAM instance, see Related Topics.

- When you restore ADAM instances to Windows XP Professional, you can restore only one ADAM instance at a time.

- If you accidentally start a restore of an ADAM instance over a currently running ADAM instance, it is recommended that you immediately restart the computer, stop the ADAM instance, and then perform the restoration again.

**Restore an ADAM instance that belongs to a configuration set**
**To completely restore an ADAM instance that belongs to a configuration set**

1. On one of the remaining ADAM instances in the configuration set, delete the server object that represents the ADAM instance that you want to restore, as follows:

   a. Open an ADAM tools command prompt.

   b. At the command prompt, type **dsmgmt**.

   c. At the **dsmgmt:** prompt, type **metadata cleanup**.

   d. At the **metadata cleanup:** prompt, type **select operation target**.

   e. At the **select operation target:** prompt, type **connections**.

   f. At the **server connections:** prompt, type **connect to server***computername:portnumber*, where *computername:portnumber* represents the ADAM instance to which you are connecting.

   g. At the **server connections:** prompt, type **q**.

   h. At the **select operation target:** prompt, type **list sites**, and identify the number corresponding to the site in which the server object that you want to delete resides.

   i. At the **select operation target:** prompt, type **select site** *n*, where *n* represents the number that is identified in the previous step.

   j. Type **list servers in site**, and identify the full distinguished name (also known as DN) of the server whose server object you want to delete.

   k. At the **select operation target:** prompt, type **q**.

   l. At the **metadata cleanup:** prompt, type **q**.

   m. At the **dsmgmt:** prompt, type **q**.

   n. At the command prompt, type **dsmgmt**.

   o. At the **dsmgmt:** prompt, type **metadata cleanup**.

   p. At the **metadata cleanup:** prompt, type **remove selected server***servername***on***computername:portnumber*

      where *servername* represents the distinguished name of the server object that you want to remove and *computername:portnumber* represents the Domain Name System (DNS) name of the ADAM instance to which you are connecting.

   q. Click **Yes** to confirm the deletion of the server object.

   r. In addition, you must delete the server object from the sites container:

      CN=Servers, CN=*sitename*, CN=Sites, CN=Configuration, CN={*GUID*}

      where *sitename* represents the name of the site in which the server object exists, and *GUID* represents the globally unique identifier (GUID) of the ADAM instance.

2. When you complete the previous steps, install an ADAM replica from media.

**Notes**

- To start Backup, click **Start**, point to **All Programs**, point to **Accessories**, point to **System Tools**, and then click **Backup**.
- If Backup does not start in wizard mode by default, you can start the wizard by clicking **Wizard Mode** on the **Welcome** tab.
- For help with **dsmgmt**, type **?** at a **dsmgmt** prompt.
- If you accidentally start a restore of an ADAM instance over a currently running ADAM instance, it is recommended that you immediately restart the computer, stop the ADAM instance, and then perform the restoration again.

**See Also**
**Concepts**
Back up an ADAM instance to a file or a tape
Install an ADAM replica from media
Connect and bind to an ADAM instance using ADAM ADSI Edit
Browse the directory tree using ADAM ADSI Edit

# View an ADAM event log

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To view an ADAM event log**

1. Open Event Viewer.
2. In the console tree, click the ADAM log that you want to view.
3. In the details pane, view the list of individual events.
4. If you want to see more details about a specific event, in the details pane, double-click the event.

## Notes

- To open Event Viewer, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Event Viewer**.
- To refresh the view in Event Viewer, on the **Action** menu, click **Refresh**. When you open a log, Event Viewer displays the current information for the log. While you view the log, the display is not updated unless you refresh the Event Viewer window. If you switch to another log and then return to the first log, the display for the first log is automatically refreshed. When you refresh the view of an event log, only the view is updated; no changes are made to the log. The **Refresh** command is not available for archived logs, because those files can no longer be updated.
- To enable auditing for an ADAM instance running under a service account other than the Network Service account, you must grant the **Generate security audits** right to the ADAM service account.

**See Also**
**Concepts**
Administering an ADAM instance
Running ADAM on Windows XP Professional
Administering computers running ADAM

# Synchronize Data from Active Directory to an ADAM Instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Active Directory to ADAM Synchronizer is a command-line tool that synchronizes data from an Active Directory forest to a configuration set of an ADAM instance.

**To synchronize data from Active Directory to an ADAM instance**

1. Open an ADAM tools command prompt.
2. At the command prompt, type:

   **adamsync /install** *server name*:*port .xml file name*

### Notes

- To open an ADAM tools command prompt, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.
- The following table contains the parameters in step 2 and other commonly used **adamsync** parameters. For more information about **adamsync** parameters, type **adamsync /?** at a command prompt, or see Related Topics.

| Parameter | Description |
|---|---|
| **/?** | Displays command-line options. |
| **/l** or **/list** | Displays all available configurations of Active Directory to ADAM Synchronizer. |
| **/d** or **/delete** *configuration_dn* | Deletes the specified configuration. |
| **/i** or **/install** *input_file* | Installs the configuration that is contained in the specified input file. |
| **/download** *configuration_dn output_file* | Creates an output file in XML containing the specified configuration. |
| **/export** *configuration_dn output_file* | Saves the current configuration to the specified output file. |
| **/sync** *configuration_dn* | Synchronizes the specified configuration |
| **/reset** *configuration_dn* | Resets the replication cookie of the specified configuration. |
| **/mai** *configuration_dn* | Configures the specified configuration as the authoritative instance. |
| **/fs** *configuration_dn* | Performs a full replication synchronization for the specified configuration. |
| **/ageall** *configuration_dn* | Performs an aging run for the specified configuration. |
| **/so** *configuration_dn object DN* | Performs a replication synchronization for the specified object in the specified configuration. Use the distinguished name of the object. |
| **/passPrompt** | Prompts for user credentials. |

- For specific information about importing the user class definitions that are supplied with ADAM, see Related Topics.

- For complete syntax information for the **adamsync** command, see Related Topics.
- You must have Read or Dirsync access to the objects or partitions in the Active Directory forest that you want to synchronize.
- You must have full control of an application directory partition on an ADAM instance to run this command.

**See Also**
**Concepts**
Glossary
Adamsync

# ADAM Concepts

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ADAM Concepts**


- ADAM Overview
- Understanding ADAM architecture
- Understanding ADAM data and data stores
- Understanding ADAM schema
- Understanding ADAM replication and configuration sets
- Understanding ADAM sites
- Understanding ADAM sites and replication
- Understanding ADAM service publication
- Understanding ADAM LDAP referrals
- Understanding ADAM bind redirection
- Understanding ADAM users and groups
- Understanding ADAM access control
- ADAM resources

# ADAM Overview

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Getting started with ADAM**

- Introducing ADAM
- ADAM requirements
- Comparing ADAM to Active Directory

# Introducing ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Introducing ADAM**

For organizations that require flexible support for directory-enabled applications, Microsoft has developed Active Directory Application Mode (ADAM). ADAM is a Lightweight Directory Access Protocol (LDAP) directory service. You can run ADAM on servers running Microsoft® Windows Server™ 2003 and also on clients running Microsoft Windows® XP Professional.

| **Note** |
| --- |
| To run ADAM on clients running Windows XP Professional, you must install the latest service packs and hot fixes. |

ADAM provides data storage and retrieval for directory-enabled applications, without the dependencies that are required for the Active Directory® directory service. ADAM provides much of the same functionality as Active Directory, but it does not require the deployment of domains or domain controllers. You can run multiple instances of ADAM concurrently on a single computer, with an independently managed schema for each ADAM instance.

**What's new in ADAM**

The following features are new to ADAM in Windows Server 2003 R2:

- Users can be created in the configuration partition so that ADAM users can be ADAM administrators. For more information, see Allow ADAM users to be created in the configuration partition.
- Active Directory to ADAM Synchronizer tool. This tool synchronizes objects from Active Directory to an ADAM instance. For more information, see Synchronize Data from Active Directory to an ADAM Instance and Adaminstall.
- ADAM users can bind to an ADAM instance by using digest authentication. This authentication method uses the default credentials of the server applications and eliminates the need to keep a plaintext version of the application's password in memory. Digest bind is supported in Ldp.
- Active Directory Schema Analyzer tool. This tool helps to migrate the Active Directory schema to ADAM. For more information, see ADSchemaAnalyzer.
- Newer version of LDP tool with access control list (ACL) editor. For more information about Ldp, see Ldp.
- ADAM can now chain user password requests in ADAM to the user object in Active Directory so that a password is changed in both directory services. When a user in ADAM who is also a user in Active Directory attempts to change the user password in ADAM, that change is treated the same as a user password change in Active Directory. Both the old and new password must be provided (except for Active Directory administrators, who only need to supply the new password), and the new password must meet any password policies that are set in Active Directory. Active Directory performs all policy checking.

**Microsoft directory technologies**

With the introduction of Active Directory Application Mode, Microsoft provides a choice of directory services. Both ADAM and Active Directory build on the same core Microsoft directory service technologies, but they address different needs within an organization.

**Active Directory.** Active Directory provides directory services for both the Windows network operating system (NOS) and for directory-enabled applications. For the NOS, Active Directory stores critical information about the network infrastructure, users and groups, network services, and so on. In this role, Active Directory must adhere to a single schema throughout an entire forest.

**ADAM.** ADAM provides directory services specifically for directory-enabled applications. ADAM does not require or rely on Active Directory domains or forests. However, in environments where Active Directory exists, ADAM can use Active Directory for the authentication of Windows security principals.

**Notes**

- Directory services (such as ADAM and Active Directory) and relational databases both provide data storage and retrieval, but they differ in their optimization. Directory services are optimized for read processing, while relational databases are optimized for transaction processing. In general, consider implementing a directory service if your application reads data more frequently than it writes data. Consider implementing a relational database if your application writes or modifies data more frequently than it reads data, such as with transactional applications.

- ADAM and Active Directory can run concurrently in the same network. In addition, ADAM can support both domain and workgroup users simultaneously, as shown in the figure below. For a comparison between ADAM and Active Directory, see Comparing ADAM to Active Directory.




**Directory-enabled applications**

A directory-enabled application uses a directory, rather than a database, flat file, or other data storage structure, to hold its data. Many off-the-shelf applications, as well as many custom applications, use a directory-enabled design. Examples of types of applications that often use a directory-enabled design include Customer Relationship Management (CRM) applications, Human Resource (HR) applications, and global address book applications.

For information about developing directory-enabled applications for ADAM, see the ADAM Software Developer's Kit (SDK) and ADAM resources.

# ADAM requirements

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Requirements**
The following requirements apply to Active Directory Application Mode (ADAM).

**Platform requirements**
ADAM requires one the following platforms:

- Microsoft® Windows Server™ 2003, Standard Edition
- Microsoft Windows Server™ 2003, Enterprise Edition
- Microsoft Windows Server™ 2003, Datacenter Edition
- The 64-bit version of Windows Server 2003, Enterprise Edition
- The 64-bit version of Windows Server 2003, Datacenter Edition
- Microsoft Windows® XP Professional Service Pack 1 (SP1)
- Microsoft Windows® XP 64-Bit Edition Service Pack 1 (SP1)

| **Note** |
| --- |
| There are some functional considerations for running ADAM on Windows XP Professional. For more information, see Running ADAM on Windows XP Professional. In addition, you must also install the latest service packs and hot fixes to run ADAM on Windows XP Professional. For more information, see Windows Update at the Microsoft Web site. |

**Workgroup and domain requirements**
ADAM does not require a forest, domain, or domain controller. You can install ADAM on computers that are configured as any of the following:

- Domain controllers
- Domain members
- Stand-alone computers

**Support for LDAP clients**
ADAM supports any Lightweight Directory Access Protocol (LDAP) client that is written to the LDAP v3 technical specification.

**Support for ADSI clients**
ADAM supports Active Directory Service Interfaces (ADSI) for clients running the following:

- Windows XP Professional
- Microsoft Windows® XP Home Edition
- Windows Server 2003 family
- Microsoft Windows® 2000

| Note |
| --- |
| For complete ADSI functionality with Windows XP Professional, you must install hotfix 817583. For more information about this hotfix, see article 817583 in the Microsoft Knowledge Base. |

# Comparing ADAM to Active Directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Comparing ADAM to Active Directory**

The following table illustrates the functional differences and similarities between Active Directory Application Mode (ADAM) and Active Directory.

| Feature | ADAM | Active Directory |
|---|---|---|
| Supports multiple schemas per server | Yes | No |
| Supports multiple directory instances per server | Yes | No |
| Runs on Windows XP Professional | Yes | No |
| Runs on member servers | Yes | No |
| Supports X.500 naming for top-level directory partitions | Yes | No |
| Supports installing, starting, and stopping without a reboot | Yes | No |
| Group Policy | No | Yes |
| Global catalog | No | Yes |
| IntelliMirror desktop management | No | Yes |
| Automated software distribution | No | Yes |
| Domain trusts and forest trusts | No | Yes |
| Public key infrastructure (PKI)/X.509 | No | Yes |
| Supports DNS service (SRV) resource records | No | Yes |
| Supports Lightweight Directory Access Protocol (LDAP) application programming interface (API) | Yes | Yes |
| Supports Active Directory Service Interfaces (ADSI) API | Yes | Yes |
| Supports Messaging API (MAPI) | No | Yes |
| Delegated administration | Yes | Yes |
| Multimaster replication | Yes | Yes |
| InetOrgPerson | Yes | Yes |
| LDAP over Secure Sockets Layer (SSL) | Yes | Yes |
| Attribute-level security | Yes | Yes |
| LDAP access control list (ACL) support | Yes | Yes |
| Microsoft Identity Integration Server 2003 compatibility | Yes | Yes |
| Extensible schema | Yes | Yes |
| Supports application directory partitions | Yes | Yes |
| Supports installation of a replica from media | Yes | Yes |

| | | |
|---|---|---|
| Supports 64-bit servers | Yes | Yes |
| Supports concurrent LDAP binding | Yes | Yes |

# Understanding ADAM architecture

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Architecture**

Active Directory Application Mode (ADAM) uses the same architecture—and even the same code base—as Active Directory. ADAM provides a hierarchical data store, a directory service component, and interfaces that clients can use to communicate with the directory service. ADAM can run on any computer running Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; Windows Server 2003, Datacenter Edition; or Windows XP Professional. ADAM does not require a domain controller or a DNS server. For information about considerations when running ADAM on Windows XP Professional, see Running ADAM on Windows XP Professional.

The following table lists the basic ADAM components, along with their Active Directory counterparts.

| Component | ADAM | Active Directory | Purpose |
|---|---|---|---|
| Directory data store | Adamntds.dit | Ntds.dit | Provides a hierarchical data store for directory data. |
| Directory service | Dsamain.exe, Adamdsa.dll | Ntdsa.dll (Runs within Lsass.exe) | Processes requests from directory clients and other directory services. |
| Directory service interfaces | Lightweight Directory Access Protocol (LDAP), replication | LDAP, replication, Messaging API (MAPI) | Provide interfaces to the directory service for making directory-related requests. |
| Directory server | Windows Server 2003, Windows XP Professional | Microsoft® Windows® 2000 Server or Windows Server 2003 domain controller | Provides the platform on which the directory service runs. |

The following illustration shows the relationship between the components of ADAM. For information about the ADAM directory data store, see Understanding ADAM data and data stores.




**Directory service**

The ADAM directory service manages the directory data store. It responds to directory requests from directory clients, as well as from other directory services. The ADAM directory service runs in the security context of the account that is specified as the ADAM service account. The ADAM directory service provides all of the following functions:

- Authentication of directory users
- Fulfillment of data requests
- Data synchronization between directory servers (through multimaster replication)
- Data management

**Notes**

- ADAM manages replication by grouping ADAM instances into configuration sets. For more information about ADAM replication, see Understanding ADAM replication and configuration sets.

- The ADAM directory service allows or denies access to clients based on credentials that are provided by the clients. ADAM supports all the same authentication, or binding methods, as Active Directory. For more information about authentication and data security, see Understanding ADAM access control.

- ADAM service account requirements vary, depending on the computer on which ADAM is installed and also on the replication scenario. For more information, see Selecting an ADAM service account.

**ADAM instances**
In ADAM, a service instance (or, simply, *instance*) refers to a single running copy of the ADAM directory service. Unlike Active Directory, multiple copies of the ADAM directory service can run simultaneously on the same computer. Each instance of the ADAM directory service has a separate directory data store, a unique service name, and a unique service description that is assigned during installation. For more information about ADAM instances, see Administering an ADAM instance.

**Directory service interfaces**
ADAM provides directory service interfaces for communication with directory clients and other ADAM instances. ADAM provides a Lightweight Directory Access Protocol (LDAP) interface for clients for authentication and for making directory service requests. ADAM provides a remote procedure call (RPC) interface for replication with other ADAM servers and for administration.

**Service dependencies**
ADAM does not require the DNS service or the File Replication service (FRS). Clients address ADAM directly, using the name or Internet Protocol (IP) address of the computer on which ADAM is running, along with the communication port number used by the ADAM instance. Because ADAM does not require or rely on domains or domain controllers, ADAM replication does not require FRS.

# Understanding ADAM data and data stores

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Data and data stores**

Active Directory Application Mode (ADAM) stores directory data in a hierarchical, file-based directory store. By default, the directory store for a given ADAM instance is located in the following:

%ProgramFiles%\Microsoft ADAM*instancename*\data\adamntds.dit

Each ADAM instance has a separate directory store. Any data operation that you perform in ADAM, including searches, reads, imports, exports, indexing, replication, backups, and restores, uses the ADAM directory store in some way. The directory store is organized into logical directory partitions, also called *naming contexts*.

| **Note** |
| --- |
| Like Active Directory, Active Directory Application Mode uses the Extensible Storage Engine (ESE) to manage its directory store. |

**Directory partitions**

The ADAM directory store is organized into logical directory partitions, or naming contexts. Three different types of directory partitions exist: configuration, schema, and application. Each ADAM directory store must contain a single configuration directory partition and a single schema directory partition, and it can contain zero or more application directory partitions. The following table summarizes these partition types, their contents, and ADAM requirements.

| Partition type | Distinguished name | Required/optional | Contents |
| --- | --- | --- | --- |
| Configuration | CN=Configuration,CN={*GUID*} | Required | ADAM configuration information |
| Schema | CN=Schema,CN=Configuration,CN={*GUID*} | Required | Class and attribute definitions |
| Application | Assigned during partition creation | Optional | Application data |

**Configuration directory partition**

The configuration directory partition holds information about ADAM replication scheduling and replica sets, information that defines the other partitions in the replication set, information about the users and groups in the replica set, and other information.

**Schema directory partition**

The schema directory partition holds the definitions for the type of data that can be held by the directory store. The ADAM directory service relies on the definitions in the schema partition for maintaining data consistency. In addition, applications can refer to the schema partition to determine the type of data that the ADAM instance allows. You can extend the schema to allow ADAM to hold data that is specific to a particular application. For more information about ADAM schemas, see Understanding ADAM schema.

| **Note** |
| --- |
| A common configuration partition and a common schema partition are replicated to ADAM service instances that belong to the same configuration set. For more information about replication, see Understanding ADAM replication and configuration sets. |

**Application directory partitions**

Application directory partitions hold the data that is used by your applications. You can create an application partition during ADAM setup or at any time after installation. Depending on your application, you might extend the schema manually or your application might automatically extend the schema for you. Typically, you manage data in a given application directory partition through your application. After the application directory partition is created, ADAM holds the application partition reference objects in CN=Partitions,CN=Configuration.

# Understanding ADAM schema

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Schema**

An Active Directory Application Mode (ADAM) schema defines, using object classes and attributes, the kinds of objects and data that can be created and stored in an ADAM directory. Each ADAM configuration set has its own independently manageable schema, which is stored in the schema directory partition. In keeping with the ADAM design concepts of simplicity and flexibility, the base (or default) ADAM schema contains only the classes and attributes that are needed to start an ADAM instance. The schema can be extended with new classes and attributes, either by administrators or by the applications themselves. In addition, unneeded schema classes and attributes can be deactivated. As with all objects in the directory, access control lists (ACLs) protect schema objects, so that only authorized users can alter the schema. Every object in an ADAM directory is an instance of an object class that is defined in a schema, as shown in the following illustration.




**Object classes**

An object class represents a category of objects, such as users, printers, or application programs, that share a set of common characteristics. The definition for each object class contains a list of the attributes that can be used to describe instances of the class. For example, the User class has attributes such as **givenName**, **surname**, and **streetAddress**. The list of attributes for a class is divided into those attributes that an object of that class must contain and additional attributes that an object may contain. The definition of each class also lists the classes whose objects can be parents of objects of a given class.

For example, the following table shows the definition for the ADAM user class.

| Attribute | Value |
|---|---|
| **Dn** | CN=User,CN=Schema,CN=Configuration |
| **objectClass** | top; classSchema; |
| **cn** | User; |
| **distinguishedName** | CN=User,CN=Schema,CN=Configuration; |
| **subClassOf** | organizationalPerson; |
| **mayContain** | audio; carLicense; departmentNumber; displayName; employeeNumber; employeeType; givenName; homePostalAddress; jpegPhoto; labeledURI; photo; preferredLanguage; roomNumber; secretary; uid; userPKCS12; userSMIMECertificate; x500uniqueIdentifier; |
| **rDNAttID** | cn; |
| **adminDisplayName** | User; |
| **adminDescription** | User; |
| **objectClassCategory** | 1; |
| **lDAPDisplayName** | user; |

| | |
|---|---|
| **name** | User; |
| **objectGUID** | dac9093a-d2aa-408a-81bb-0fe8179165da; |
| **schemaIDGUID** | bf967aba-0de6-11d0-a285-00aa003049e2; |
| **objectCategory** | CN=Class-Schema,CN=Schema,CN=Configuration; |
| **defaultObjectCategory** | CN=Person,CN=Schema,CN=Configuration; |

**Attributes**

The schema also defines each attribute. The definition for each attribute includes unique identifiers for the attribute, the syntax for the attribute, optional range limits for the attribute values, whether the attribute can have only one value or multiple values, and whether the attribute is indexed. The directory schema defines each attribute exactly once. Each attribute can then be referenced by multiple object classes. For example, the **description** attribute is defined once and then referenced by many object classes.

The following table shows the attribute definition for **telephone number**, a representative schema attribute.

| Attribute | Value |
|---|---|
| **objectClass** | top; attributeSchema; |
| **cn** | Telephone-Number; |
| **distinguishedName** | CN=Telephone-Number,CN=Schema,CN=Configuration; |
| **instanceType** | 0x4 = ( IT_WRITE ); |
| **whenCreated** | 11/12/2002 13 |
| **22** | 14 Pacific Standard Time Pacific Daylight Time; |
| **whenChanged** | 11/12/2002 13 |
| **22** | 14 Pacific Standard Time Pacific Daylight Time; |
| **uSNCreated** | 217; |
| **attributeID** | 2.5.4.20; |
| **attributeSyntax** | 2.5.5.12; |
| **isSingleValued** | TRUE; |
| **rangeLower** | 1; |
| **rangeUpper** | 64; |
| **uSNChanged** | 217; |
| **showInAdvancedViewOnly** | TRUE; |
| **adminDisplayName** | Telephone-Number; |
| **adminDescription** | Telephone-Number; |
| **oMSyntax** | 64; |
| **searchFlags** | 0; |
| **lDAPDisplayName** | telephoneNumber; |
| **name** | Telephone-Number; |

| | |
|---|---|
| **objectGUID** | bf19d7eb-ea0f-4f2d-af67-f439a037d8a4; |
| **schemaIDGUID** | bf967a49-0de6-11d0-a285-00aa003049e2; |
| **attributeSecurityGUID** | 77b5b886-944a-11d1-aebd-0000f80367c1; |
| **systemOnly** | FALSE; |
| **systemFlags** | 0x10 = ( FLAG_SCHEMA_BASE_OBJECT ); |
| **isMemberOfPartialAttributeSet** | TRUE; |
| **objectCategory** | CN=Attribute-Schema,CN=Schema,CN=Configuration; |

**Single-valued and multivalued attributes**

Attributes can be single-valued or multivalued. An instance of a single-valued attribute can contain only a single value. An instance of a multivalued attribute can contain multiple values, which must all use the same syntax. Each value of a multivalued attribute must be unique.

| **Note** |
|---|
| A multivalued attribute stores its values in random order. Therefore, you should not make any directory decisions—programmatic or manual—based on the order of the values in a multivalued attribute. |

**Indexed attributes**

Indexing attributes helps to improve the performance of queries that are based on the indexed attribute. Both single-valued and multivalued attributes can be indexed; classes cannot be indexed. Attributes can be marked for indexing through their schema definition. Indexing an attribute also makes it possible for users to use wildcards (*) as prefixes and suffixes when they specify a search string. When you mark an attribute as indexed, all instances of the attribute are added to the index, not just the instances that are members of a particular class. Indexing attributes, particularly multivalued attributes, can negatively affect replication and object creation time, as well as directory database size. Therefore, you should only index commonly used attributes.

For general information about schemas, see "Active Directory Schema" at the Microsoft MSDN Web site.

# Understanding ADAM replication and configuration sets

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Replication and configuration sets**
Active Directory Application Mode (ADAM) uses replication to provide fault tolerance and load balancing for directory services. ADAM uses a type of replication called multimaster replication. Through replication, ADAM copies directory data updates that are made to a directory partition on one ADAM instance to other ADAM instances that hold copies of the same directory partition. ADAM instances that hold copies of the same directory partition or partitions form a logical grouping called a configuration set.

**Multimaster replication**
Multimaster replication simply means that you can make changes to directory data on any ADAM instance. ADAM replicates these changes to other members of the configuration set automatically. Multimaster replication is characterized by loose data consistency with convergence. When you make changes to data on a given directory partition at one ADAM instance, replicas of that directory partition that are stored on other ADAM instances become inconsistent with the most up-to-date replica of the directory partition (the partition where the changes were made). However, as changes get replicated through the configuration set, all partition replicas once again become identical; that is, they converge to the most recent data.

**Configuration sets**
ADAM instances replicate data based on participation in a configuration set. All ADAM instances that are joined to the same a configuration set must replicate a common configuration directory partition and a common schema directory partition. ADAM instances in a configuration set can also replicate any number of application directory partitions. ADAM instances in a configuration set are not required to replicate all application directory partitions in the configuration set. A single ADAM instance can replicate all—or any subset of—the application directory partitions in its configuration set. An ADAM instance cannot, however, replicate an application directory partition from a different configuration set. For more information about the types of directory partitions in ADAM, see Understanding ADAM data and data stores.

The following illustration shows an example of two ADAM configuration sets, each with two ADAM instances. As shown in the illustration, a single computer can run multiple ADAM instances, each in a different configuration set.




**Note**

An ADAM instance can be joined to a configuration set only during installation of the instance.

**Preventing replication conflicts**
What if two different users make changes to the same data on replicas of the same directory partition on two different ADAM instances? In such a case, each ADAM instance attempts to replicate the changes, creating a conflict. To resolve this conflict, replication partners that receive these conflicting changes examine the attribute data that is contained in the changes, each of which holds a version and a time stamp. ADAM instances accept the change with the higher version and discard the other change. If the versions are identical, ADAM instances accept the change with the more recent time stamp.

If two or more values in a multivalued attribute on an object are updated simultaneously on two different ADAM instances,

only one of the updated values will be replicated. In other words, simultaneous updates to a multivalued attribute that occur on two different ADAM instances are considered in conflict, even if the updates apply to different values within the multivalued attribute. The only exception to this rule is for linked value attributes (such as group memberships), which do allow for simultaneous updates to different values within the linked value attribute.

**Replication topology**

Knowledge Consistency Checker (KCC), a process that runs as part of each ADAM instance, automatically constructs the most efficient topology for replication traffic to follow, based on the network. The KCC regularly recalculates the replication topology to adjust for any network changes that occur in the environment.

| **Note** |
| --- |
| An ADAM configuration set maintains its own replication topology, separate from any Active Directory replication topology that might also exist. Directory partitions cannot be replicated between ADAM instances and Active Directory domain controllers. |

**Ensuring replication security**

To ensure replication security, ADAM authenticates replication partners before replication, and replication authentication always occurs over a secure channel. ADAM uses Security Support Provider Interface (SSPI) to establish the appropriate authentication security level between replication partners. The method used for replication authentication within a configuration set depends on the value of the **msDS-ReplAuthenticationMode** attribute on the configuration directory partition. After replication partners have successfully authenticated, all replication traffic between the two partners is encrypted.

The following table describes the security levels for replication authentication and the corresponding **msDS-ReplAuthenticationMode** attribute value for each security level. The default replication security level for a new, unique ADAM instance is 1, unless a local workstation user account is specified as the ADAM service account. If a local workstation account is specified as the ADAM service account, the replication security level is set to 0.

| **Replication security level** | **Authentication mode** | **Description** | **Supported environments** |
| --- | --- | --- | --- |
| Mutual authentication with Kerberos | 2 | Kerberos authentication, using service principal names (SPNs), required. If Kerberos authentication fails, the ADAM instances will not replicate. | Configuration set must be fully contained in an Active Directory domain, forest, or forest trust. |
| Negotiated | 1 | Kerberos authentication (using SPNs) is attempted first. If Kerberos fails, NTLM authentication is attempted. If NTLM fails, the ADAM instances will not replicate. | Configuration set can contain Microsoft® Windows NT® 4.0 member servers. |
| Negotiated pass-through | 0 | All ADAM instances in the configuration set use an identical account name and password as the ADAM service account. | Configuration set can include computers that are joined to one or more workgroups or that are joined to multiple domains or forests without trust relationships. |

For information about modifying the replication security level of a configuration set, see Modify the replication security level of a configuration set.

To help maintain ADAM replication security, the following best practices are recommended:

- Use the highest level of replication security that your environment can support.
- In Active Directory environments, run ADAM on member servers, rather than on domain controllers, whenever possible.
- If you run ADAM on a domain controller in an Active Directory environment, do not use the Network Service account as the ADAM service account. Instead, use a domain user account that does not have administrative privileges.
- In workgroup and Windows NT 4.0 environments, do not use an account with administrative privileges as an ADAM service account.
- Use separate configuration sets for applications with strict isolation requirements.

For more information about ADAM replication requirements for service accounts, see Selecting an ADAM service account. For more information about service principal names (SPNs) and replication security, see Administering ADAM service principal names.

# Understanding ADAM sites

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Sites**
Active Directory Application Mode (ADAM) uses topology information, which is stored as site objects and site link objects in the configuration directory partition, to build the most efficient replication topology for a configuration set. You can use sites in ADAM to represent the physical structure, or *topology*, of your network. You can use ADAM ADSI Edit to define site objects and site link objects.

**Site objects**
You can create a site object in ADAM to represent each area of your network (typically, a building or a group of buildings) in which all computers are connected by high-speed bandwidth. You can then move the directory objects of ADAM instances that are located in each area of the network into the corresponding site object. By default, the directory objects of each ADAM instance that you create belong to a single default site (CN=Default-First-Site-Name,CN=Sites,CN=Configuration,CN={*GUID*}).

For information about establishing single or multiple sites, see Administering sites.

**Site link objects**
You can use site link objects in the directory to represent low-speed bandwidth connections between the sites on your network. Site link objects help you tune replication on your network; you can use them to determine the relative cost of replication across each site link. By default, the default site (CN=Default-First-Site-Name) belongs to the default site link object (CN=DEFAULTIPSITELINK,CN=IP,CN=Inter-Site Transports,CN=Sites,CN=Configuration,CN={GUID} ). You can add new sites that you create to the default site link object. You can also create additional site link objects to represent each of the actual site links on your network.

**Site link costs**
Each site link object in the directory has an associated cost. You can edit the cost attribute on a site link object to specify the relative cost of replication across that site link. ADAM takes into consideration site link costs when it builds the replication topology. ADAM considers a site link with a higher cost to be less desirable than a site link with a lower cost. When multiple paths for replication are available, ADAM prefers the site link with the lower cost.

For more information about how sites affect replication, see Understanding ADAM sites and replication.

# Understanding ADAM sites and replication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Sites and replication**
In Active Directory Application Mode (ADAM), the replication topology is the set of physical connections that is used to replicate directory updates among ADAM instances in a configuration set. ADAM builds the most efficient replication topology for a configuration set based on the site or sites in which members of the configuration set reside. ADAM treats replication within a site differently than replication between sites.

**Building the intersite replication topology**
One ADAM server in each site, called the intersite topology generator, is assigned to build the topology. Together, these intersite topology generators build a least-cost spanning tree algorithm that eliminates redundant replication paths between sites. The intersite replication topology is updated regularly to respond to any changes that occur in the network. You can control intersite replication through the information that you provide when you create your site links.

**Determining when intersite replication occurs**
Active Directory preserves bandwidth between sites by minimizing the frequency of replication and by allowing you to schedule the availability of site links for replication. By default, intersite replication across each site link occurs every 180 minutes (3 hours). You can adjust this frequency to match your specific needs. Be aware that increasing this frequency increases the amount of bandwidth that is used by replication. In addition, you can schedule the availability of site links for use by replication. By default, a site link is available to carry replication traffic 24 hours a day, 7 days a week. You can limit this schedule to specific days of the week and times of day. You can, for example, schedule intersite replication so that it only occurs after normal business hours. For more information, see Configure replication availability on a site link.

**Building the intrasite replication topology**
The Knowledge Consistency Checker (KCC) on each ADAM instance automatically builds the most efficient replication topology for intrasite replication, using a bidirectional ring design. This bidirectional ring topology attempts to create at least two connections to each ADAM instance (for fault tolerance) and no more than three hops between any two ADAM instances (to reduce replication latency). To help prevent connections of more than three hops, the topology can include shortcut connections across the ring. The KCC updates the replication topology regularly.

**Determining when intrasite replication occurs**
Intrasite replication is optimized for speed, rather than bandwidth, because bandwidth within a site is assumed to be high-speed. Replication within a site occurs automatically on the basis of change notification. Intrasite replication begins when a directory update occurs. By default, the source ADAM instance waits 15 seconds and then sends an update notification to its closest replication partner. If the source ADAM instance has more than one replication partner, subsequent notifications go out by default at 3-second intervals to each partner. After receiving notification of a change, a partner ADAM instance sends a directory update request to the source ADAM instance. The source ADAM instance responds to the request with a replication operation. The 3-second notification interval prevents the source ADAM instance from being overwhelmed with simultaneous update requests from its replication partners.

If no directory updates have occurred within a given time period, intrasite replication still occurs, based on a scheduled interval. By default, this scheduled interval is once per hour. For information about modifying this time period, see Configure replication frequency within a site.

# Understanding ADAM service publication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ADAM service publication**
In Active Directory environments, *service publication* refers to the ability of a service to publish information about itself in the directory and to the ability of clients to discover that information and locate the service. When a computer on which Active Directory Application Mode (ADAM) is running is joined to an Active Directory domain, ADAM attempts to create service connection point (SCP) objects in Active Directory.

**Active Directory and service connection points**
In Active Directory network environments, services can publish information about their existence using **serviceConnectionPoint** objects in the directory. These objects contain binding information that client applications use to find and connect to instances of the service. To access a service, a client does not have to know about specific computers; the objects in Active Directory include this information.

When an ADAM instance is running in an Active Directory environment, it makes a best-effort attempt to publish updated information about itself in Active Directory. This attempt is a "best-effort" attempt because the attempt to create SCPs does not always succeed, and failure to create SCPs does not prevent the ADAM instance from running or accepting client connections.

The attempt to create SCPs succeeds if:

- The computer on which the ADAM service is running is joined to a domain, and
- The ADAM service account possesses Create Child rights on the computer object where the **serviceConnectionPoint** object is to be created.

By default, ADAM runs as NetworkService, and the **serviceConnectionPoint** object is created under the computer object that represents the computer on which ADAM is running. If the SCP object already exists, ADAM updates the object with any new information about the ADAM instance.

A **serviceConnectionPoint** object contains the information in the following table.

| serviceConnectionPoint attribute | Attribute contents |
|---|---|
| **Keywords** | • Windows 2000 Lightweight Directory Access Protocol (LDAP) object identifier (also known as OID)<br>• ADAM LDAP object identifier<br>• Globally unique identifier (GUID) of the **NTDS Settings** object under the computer object of the ADAM instance<br>• Site of the ADAM instance<br>• Instance name of the ADAM instance<br>• Operation master role (schema or naming), if held by the ADAM instance<br>• Distinguished name and GUID of the configuration directory partition of the ADAM instance<br>• Distinguished name and GUIDs of all application directory partitions on the ADAM instance |

| | |
|---|---|
| **serviceBindingInformation** | • LDAP connection point (in the form of ldap:*computername:ldapport*)<br>• Secure Sockets Layer (SSL) connection point (in the form of ldaps:*computername:sslport*) |

| **Note** |
|---|
| The **Keywords** attribute is a multivalued attribute. |

By default, the global catalog in Active Directory contains the contents of the **Keywords** attribute of an SCP object. As a result, a client can locate an object on which an SCP is created, even if the object is located in a different domain than the client.

**Client searches**

Client applications can find ADAM instances by searching the SCP attributes that are contained in the global catalog. For information about performing this search programmatically, see "How Clients Find and Use a Service Connection Point" at the Microsoft MSDN Web site. (Unlike the example shown at the Microsoft MSDN Web site, client applications need to perform only one search to retrieve the information necessary for ADAM binding.) Client applications can search for the ADAM object identifier, configuration partition GUID, ADAM instance GUID, ADAM instance name, or any directory partition. Client applications may also perform load balancing by selecting an ADAM instance at random when a search returns more than one applicable ADAM instance.

**Service publication and discovery without SCPs**

ADAM does not require that SCPs be published to run successfully. ADAM operates successfully with or without SCPs. Examples of environments without SCPs include workgroup environments and ADAM instances running under service accounts that do not possess sufficient privileges to create SCPs. In these cases, client applications can rely on Domain Name System (DNS) to resolve the host name of a computer on which ADAM is running. Note, however, that ADAM instances, unlike domain controllers in Active Directory, do not create service (SRV) records in DNS.

# Understanding ADAM LDAP referrals

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ADAM LDAP referrals**

In accordance with RFC 2251, Active Directory Application Mode (ADAM) supports the use of referrals. With referrals you can use ADAM to direct LDAP clients to a more appropriate location in response to a query from the client. Also in accordance with RFC 2251, ADAM can return referrals in response to any LDAP operation request, except "unbind" and "abandon" requests, which have no responses. ADAM maintains referral data in **crossRef** objects that are stored in the partitions container of the configuration directory partition.

ADAM creates and maintains **crossRef** objects for all directory partitions in the configuration set, including the configuration and schema partitions. These **crossRef** objects are used to generate referrals in response to queries that request data about objects that exist in the configuration set but that are not contained in the ADAM instance handling the request.

At a minimum, at least one Uniform Resource Locator (URL) must be present in the referral. When possible, ADAM returns at least three referrals, to provide referral redundancy to clients in distributed environments. A referral is not returned for a **singleLevel** or **wholeSubtree** search in which the search scope spans multiple naming contexts and several different servers must be contacted to complete the operation. Instead, continuation references are returned.

If a **crossRef** object with a matching directory partition exists, ADAM uses it for referrals. However, if a **crossRef** object does not exist, but the queried namespace uses DNS-style (DC=) naming, ADAM constructs a DNS name to be used as a referral. If the queried namespace does not use DNS-style naming, a superior reference is used, if one exists. A superior reference can be configured by setting the **superiorDnsRoot** attribute on the configuration partition **crossRef** object. If a superior reference is not configured (it is not configured by default), an operational error is returned.

**Generating referrals from a crossRef object**

If the **dNSRoot** attribute has been set on the **crossRef** object (**dNSRoot** is not set by default), ADAM returns the DNS name as the referral for the queried namespace. (By default, the **dNSRoot** attribute is set to NULL.) Then, the client can chase the referral through registered DNS records for each host that contains the queried namespace.

If the **dNSRoot** attribute is NULL, the referral is generated based on other information about the **crossRef**, as follows:

- ADAM refers to the **msDS-NC-Replica-Locations** attribute to find the **NTDS Settings** objects representing the instances hosting this partition (for example: CN=NTDS Settings,CN=Machine1$instance1,CN=Servers,CN=Default-First-Site-Name,CN=Sites,CN=Configuration,CN={<GUID>}). If more than three ADAM instances host the directory partition, a randomly selected set of three instances is chosen.

- From the **NTDS Settings** object (the authoritative source within the directory for communications port information), ADAM retrieves the **msDS-PortLDAP** attribute for the communications port.

- From the parent (**Server Object**), ADAM retrieves the **dNSHostName** attribute (instance1.machine1.example.com). **dNSHostName** is the authoritative source for the server host name, and it can either be a DNS host name or a host Internet Protocol (IP) address. If the server name changes, this object can be relied on to have the updated information.

- ADAM then constructs the referral based on the server and communications port information. ADAM returns multiple referrals for each request to ensure that the client can pursue the request, even if one or more hosts represented in the referrals cannot be reached. ADAM selects a random number of locations (three by default) and continues to use those defaults for five minutes, after which another three hosts are selected. Both the number of hosts and number of minutes are configurable. If the number of hosts specified is greater than the actual number of instances holding the directory partition, all of the hosts are returned as referrals, and they will continue to be used indefinitely. This information is kept in the crossref cache, and it is refreshed anytime the crossref cache is rebuilt. The crossref cache is automatically rebuilt when **msds-nc-replica-locations** changes, that is, when an ADAM instance is added or removed.

| **Note** |
| --- |
| You can modify the values in the following table on the **msDS-Other-Settings** attribute of the directory service object, or on a **crossRef** object, to configure the referral behavior of an ADAM instance. |

| Value name | Default | Meaning |
| --- | --- | --- |
| **MaxReferrals** | 3 | Maximum number of referrals returned by ADAM |
| **ReferralRefreshInterval** | 5 minutes | Time interval at which referral information is updated |
| **SelfReferralOnly** | 0 | Setting the value to 1 prevents ADAM from providing referral information |

Values that you set on a **crossRef** object take precedence over any values that you set on the directory service object.

**Generating referrals without crossRef objects**

If a **crossRef** does not exist for the search base, ADAM determines the DNS name if the distinguished name that is presented in the request uses DNS-style (DC=) syntax. Otherwise, ADAM uses the **supref** attribute of the configuration directory partition to generate a referral to the **Superior Referral**, if one is configured. This attribute is NULL by default, and it must be configured explicitly. Otherwise, no referrals are generated.

# Understanding ADAM bind redirection

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ADAM bind redirection**

Active Directory Application Mode (ADAM) bind redirection is designed for use primarily with legacy applications that are unable to authenticate directly against Active Directory but still need to use ADAM as the application data store. Through bind redirection, ADAM can accept a bind request from an application and redirect this bind request to Active Directory, based on the contents of a proxy object. A proxy object in ADAM represents an Active Directory security principal, and it can be augmented to store additional data related to that security principal that is specific to the application. Through bind redirection, applications can take advantage of the identity store of Active Directory without any additional overhead. For example, if an administrator disables an account in Active Directory, that account can no longer be used to bind to ADAM. In addition, applications retain the flexibility of using ADAM as an application data store.

**Proxy objects for bind redirection**

Bind redirection relies on the existence of stub, or *proxy*, objects in ADAM, each of which represents an Active Directory security principal. To implement bind redirection, you first create a proxy object schema definition to store data that is specific to your application. In this definition, you include the **msDS-bindProxy** auxiliary class. The **msds-BindProxy** class possesses a single "must contain" attribute, **ObjectSid**, which holds the security ID (SID) of the associated Active Directory security principal. You can initially populate ADAM with proxy objects representing Active Directory users through a synchronization mechanism, such as Microsoft Identity Integration Server 2003, or through the construction and importation of .ldf files into ADAM.

**Processing bind requests**

When a user sends a bind request to ADAM, ADAM does the following based on the type of request:

For simple bind requests:

- If the object class to which the user requests a bind possesses **msDS-bindableObject** as a static auxiliary class, ADAM processes the bind request directly, generating a security context that is based on the ADAM user's SID.

- If, instead, the object class to which the user requests a bind possesses the **msDS-bindProxy** as an auxiliary class, ADAM redirects the bind request to Active Directory. ADAM then generates a security context that includes the token that is returned by Active Directory in response to the redirection, along with the ADAM groups in which the principal is a member.

For simple authentication and security layer (SASL) bind requests:

- The bind request goes directly to the **AcceptSecurityContext** application programming interface (API) call. ADAM generates a security context that is based on the token that is returned from **AcceptSecurityContext**.

**Designing applications for bind redirection**

Although bind redirection is intended for preexisting, legacy applications, it can also be used with new applications. However, for new and modifiable applications, it is recommended that you write the application to use Windows security principals, rather than proxy objects, for authentication. With this type of design, you can maintain a single sign on (SSO) solution and prevent the duplication of accounts.

**New and modifiable applications**

For applications that are currently in development, and for applications that can be updated, it is recommended that you write your applications to authenticate directly against Active Directory, using Windows security principals. This type of design provides the highest degree of security, because passwords flow directly from the client to Active Directory, rather than through ADAM. ADAM can then be used simply as an application-specific data store. Proxy objects in ADAM can hold application data that is specific to each Windows security principal, and an attribute on the ADAM proxy object can be used to uniquely link each proxy object to a particular Windows security principal. A unique identifier on a Windows security principal, such as a security ID (SID) or a globally unique identifier (GUID), can be used for this linking. If your application is unable to authenticate against Active Directory, and you need to synchronize directory contents between ADAM and Active Directory, you can use a synchronization service, such as Microsoft Identity Integration Server (MIIS). However, this option introduces a delay between the originating update and the synchronization of the update.

**Legacy applications**
In some cases, you may be running already deployed, directory-enabled applications that cannot be modified. In these cases, your application may send a flat string to the directory, expecting the return of a distinguished name string with which to bind. For this scenario, the ADAM proxy object can support a unique distinguished name, and it can also contain a flag indicating that the object represents a security principal in Active Directory. Based on the presence of the flag, when a user attempts to bind to the proxy object, ADAM redirects the bind request to Active Directory, rather than processing the request locally.

**Password security**
Because bind redirection can only be used in conjunction with simple bind requests, any password that is sent to ADAM as part of a bind redirection is sent in a plaintext, unencrypted format. Therefore, it is strongly recommended that you implement bind redirection only when using an encrypted connection, for example, Secure Sockets Layer (SSL), which is the default setting in ADAM.

## Notes

- ADAM should be used to store data that is specific to a particular application, while Active Directory should be used to hold more widely used data.
- Only simple bind requests can be redirected.
- You cannot modify the password of an Active Directory security principal through a bind redirection.
- Bind direction only supports security principals from the forest to which the computer is joined or to a trusted domain or forest.
- If both **msDS-bindProxy** and **msDS-bindableObject** are specified as auxiliary classes on a single object, **msDS-bindableObject** takes precedence.

# Understanding ADAM users and groups

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Users and groups**

Active Directory Application Mode (ADAM) relies on users and groups to provide and control access to directory data. ADAM supports the simultaneous use of both Windows users and ADAM users. ADAM provides four default, role-based groups. You can create additional ADAM groups as necessary. Both Windows users and ADAM users can be members of ADAM groups. To create ADAM users in ADAM, you must first import the user object class definitions that are provided with ADAM, or you can supply your own user object definitions.

**Default groups**

ADAM provides four default, role-based groups: Administrators, Instances, Readers, and Users. These groups reside in the configuration partition and in each application partition, but not in the schema partition. Within a configuration set, ADAM replicates these groups, along with all other directory data.

The following three groups reside in the CN=Roles container of each directory partition:

- Administrators (CN=Administrators, CN=Roles)
- Readers (CN=Readers, CN=Roles)
- Users (CN=Users, CN=Roles)

The following group resides only in the CN=Roles container of the configuration directory partition:

- Instances (CN=Instances, CN=Roles)

The following table lists the default ADAM groups, along with the default members and default access assigned to each group.

| Group | Default members | Default access |
|---|---|---|
| **Administrators**<br><br>(CN=Administrators, CN=Roles) | **Configuration partition:**<br><br>ADAM administrators assigned during ADAM setup<br><br>**Application partitions:**<br><br>Administrators group from the configuration partition | Full access to all partitions |
| **Instances**<br><br>(CN=Instances, CN=Roles) | All instances | |
| **Readers**<br><br>(CN=Readers, CN=Roles) | None | Read access to the partition |
| **Users**<br><br>(CN=Users, CN=Roles) | **Configuration partition:**<br><br>Transitively, all ADAM users<br><br>**Application partitions:**<br><br>Transitively, all ADAM users created in the partition | None |

Groups from the configuration partition can be assigned permissions in any partition of an ADAM instance or configuration set. Groups from an application partition can be assigned permissions only in that application partition. Windows security principals can be assigned permissions in any application partition.

**Security principals**

The term *security principal* refers to any object that has a security identifier (SID) and that can be assigned permissions to directory objects. In ADAM, security principals can reside in ADAM, on a local computer, or in an Active Directory domain.

| **Note** |
| --- |
| You can retrieve an active user's individual and group SIDs by explicitly querying the **tokenGroups** attribute on the **rootDSE**. |

**ADAM security principals**

ADAM does not include any default security principals. However, ADAM does provide importable schema extensions that you can use to create users in ADAM. Users created from these user classes can be used as security principals. In addition, you can make any object class in the ADAM schema a security principal by adding the **msDS-bindableobject** auxiliary class and the **unicodePwd** attribute to the schema definition of an object class. Each ADAM security principal must be assigned an account and password, which ADAM uses for authentication.

For more information about authentication in ADAM, see Understanding ADAM access control. For information about the optional user classes provided with ADAM, see Administering users and groups.

**Windows security principals**

ADAM allows the use of Windows security principals for authentication and access control. Local Windows users and groups, as well as domain users and groups, can be used with ADAM. In addition, you can add Windows security principals membership to ADAM groups as members. By default, the security principal that you specify as the ADAM administrator during ADAM setup becomes a member of the Administrators group in the configuration partition. For Windows security principals, ADAM relies on the Local Security Authority (LSA) on the local computer (for local accounts) or the LSA on a domain controller (for domain accounts) for authentication.

For more information about authentication in ADAM, see Understanding ADAM access control.

# Understanding ADAM access control

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Access control**

The term *access control* describes how a directory service provides and secures access to directory data. Access control in Active Directory Application Mode (ADAM) consists of two parts. First, ADAM authenticates the identity of users requesting access to the directory, allowing only successfully authenticated users into the directory. Second, ADAM uses security descriptors, called access control lists (ACLs), on directory objects to determine to which objects an authenticated user has access.

**Authentication**

Users, or *security principals*, request directory data from ADAM through directory-enabled applications, which in turn make requests to ADAM using Lightweight Directory Access Protocol (LDAP). Before making a request for data, the directory-enabled application must present the user's credentials to ADAM for authentication, or *binding*. This request includes a user name, password, and—depending on the type of bind—a domain name or computer name.

ADAM can accept authentication, or *bind*, requests from both ADAM security principals and Windows (local and domain) security principals. ADAM security principals are authenticated directly by ADAM. Local Windows security principals are authenticated by the local computer. Domain security principals must authenticated by an Active Directory domain controller. The following table summarizes supported authentication methods in ADAM, based on the type of security principal.

| Security principal | Accepted credential formats | Authenticated by |
|---|---|---|
| ADAM | • Distinguished name<br>• Canonical name<br>• **displayName** attribute value<br>• User principal name (UPN) | ADAM instance |
| Workstation account | • Machine name\account | Local Security Authority (LSA) on the local computer |
| Domain account | • Domain name\account<br>• UPN | LSA on a domain controller from the user's domain |

| **Note** |
|---|
| A UPN that is used in ADAM does not require the use of the "@" symbol. |

The possible authentication paths in ADAM are shown in the figure below.




For more information about binding to LDAP directories, see Establishing an LDAP Session (http://go.microsoft.com/fwlink/?LinkId=99557).

**ADAM security principals**

ADAM security principals are users or other objects that are created and reside in the ADAM directory. ADAM security principals can be created based on the importable user classes (including **inetOrgPerson**) that are provided with ADAM. Or, ADAM security principals can be created from any object class in the schema that contains the **SecurityPrincipal** static auxclass and the **unicodePwd** attribute. ADAM assigns each ADAM security principal a unique security ID (SID), which is guaranteed to be unique across all ADAM naming contexts. For more information about ADAM security principals, see Understanding ADAM users and groups.

**ACLs**

Permissions on ADAM directory objects are assigned to users and groups using access control lists (ACLs). An ACL contains a list of users and groups that have been assigned permissions on a directory object, and it also contains the types of permissions (read, write, and so on) that they have been assigned. (Each individual entry in an ACL is called an access control entry (ACE).) If an ACL on a directory object does not explicitly assign permissions to a user, or to any groups that a user is a member of, the user will be denied access to that object.

For the most part, ACLs in ADAM function the same as in Active Directory. ACEs differ between ADAM and Active Directory in that ADAM validates SIDs as they are added to the ACE, while Active Directory allows all user SIDs. ADAM user SIDs are allowed only if one of the following conditions is true:

- The ADAM user is in the same naming context as the object to which the ACE is added.
- The object is in the schema naming context.
- The object is in the configuration naming context.

**Security context**

After a user has been successfully authenticated, ADAM can construct the security context for the authenticated user. This security context specifies the SID of the user, along with a list of all of the groups to which the user belongs.

The security context for an authenticated Windows security principal includes:

- The user's SID
- The user's Windows group memberships
- The user's ADAM group memberships (including those groups in which the user is a direct member and those groups in which the user is a member through recursive group expansion)

The security context for an authenticated ADAM security principal includes:

- The user's SID
- SIDs of each of the following ADAM groups:
    - Everyone
    - AuthUsers
    - Users group from the partition in which the ADAM user resides
    - Users group from the configuration partition
    - Groups in which the user is a member through recursive group expansion

By comparing an authenticated user's security context with the ACL on the directory object to which the user has requested access, ADAM can determine whether to grant the user access to the object.

**Retrieving user SIDs**
You can retrieve an active user's SIDs by explicitly querying the **tokenGroups** attribute on the **rootDSE**. The **tokenGroups** attribute is a multivalued attribute. The first value represents the user's individual SID, with the values following representing the user's ADAM and Windows group SIDs.

For more information about how access control is determined in Active Directory and ADAM, see Controlling Access to Objects in Active Directory Domain Services (http://go.microsoft.com/fwlink/?LinkId=99558).

# ADAM resources

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Resources**

Along with the Active Directory Application Mode Administrator's Guide, these additional resources can help you get the most out of your Active Directory Application Mode (ADAM) deployments.

**Microsoft Web site**

- "Understanding LDAP" (http://go.microsoft.com/fwlink/?LinkId=99559).
- "Active Directory Domain Services" at the Microsoft MSDN Web site.
- "Lightweight Directory Access Protocol" (http://go.microsoft.com/fwlink/?LinkId=99560).
- For information about installing the latest service packs or hot fixes, see Windows Update at the Microsoft Web site. Also, see Using Windows Update.
- For information about binding to Lightweight Directory Access Protocol (LDAP) directories, see "Establishing an LDAP Session" (http://go.microsoft.com/fwlink/?LinkID=99557).
- For information about access in Active Directory—and in ADAM, see "Controlling Access to Active Directory Objects in Active Directory Domain Services" (http://go.microsoft.com/fwlink/?LinkId=99558).
- For information about the Active Directory schema, see "Active Directory Schema" at the Microsoft MSDN Web site.
- For information about developing directory-enabled applications, see "Design Guide for Directory-Enabled Applications" (http://go.microsoft.com/fwlink/?LinkId=99561).
- For information about role-based authorization in Windows Server 2003, see "Dynamic Groups in Windows Server 2003 Authorization Manager" (http://go.microsoft.com/fwlink/?LinkId=99562).
- For information about setting up and using certificate services, see "Certificate Services How To" at the Microsoft TechNet Web site.

**Other Resources**

- RFC 2849: The LDAP Data Interchange Format (LDIF)—Technical Specification at the Internet Engineering Task Force (IETF) FTP site (http://go.microsoft.com/fwlink/?LinkId=63904)
- RFC 3377: Lightweight Directory Access Protocol (v3): Technical Specification at the Internet Engineering Task Force (IETF) FTP site (http://go.microsoft.com/fwlink/?LinkId=63908)

| **Note** |
| --- |
| Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here. |

# Administering ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering ADAM**


- Using ADAM administration tools
- Administering computers running ADAM
- Running ADAM on Windows XP Professional
- Administering an ADAM instance
- Administering directory partitions
- Administering sites
- Administering replication and configuration sets
- Administering ADAM service publication
- Administering ADAM service principal names
- Administering the ADAM schema
- Administering users and groups
- Administering access control
- Backing up and restoring ADAM
- Selecting an ADAM service account

# Using ADAM administration tools

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Using ADAM administration tools**

Active Directory Application Mode (ADAM) comes with several administration tools. The ADAM Setup Wizard installs ADAM administration tools when you install a new ADAM instance. In addition, you can use the ADAM Setup Wizard to install the ADAM administration tools only, without installing a new ADAM instance. For more information, see Selecting components to install.

To open a command prompt from which to run ADAM administration tools, click **Start**, click **All Programs**, point to **ADAM**, and then click **ADAM Tools Command Prompt**.

The following table describes ADAM administration tools and their purpose.

| Tool | Purpose |
|---|---|
| ADAM ADSI Edit | Microsoft Management Console (MMC) snap-in for general ADAM administration, including replication scheduling. |
| ADAM Schema | MMC snap-in for administration of the ADAM schema. |
| Adamsetup | Command that launches the ADAM Setup Wizard and, through command-line parameters, that can be used to perform unattended ADAM installs. |
| Csvde | Command-line tool for importing and exporting directory contents, such as schema definitions and directory objects, that are stored as comma-separated value (.csv) files. |
| Dsmgmt | Command-line tool for general ADAM administration. |
| Dsacls | Command-line tool for displaying and changing permissions (access control entries (ACEs)) in the access control list (ACL) of objects in ADAM. |
| Dsdbutil | Command-line tool for managing ADAM database (.dit) files. |
| Dsdiag | Command-line tool for general ADAM monitoring. |
| Ldifde | Command-line tool for importing and exporting directory contents, such as schema definitions and directory objects, that are stored as LDAP Data Interchange Format (LDIF) files. |
| Ldp | Graphical user interface (GUI) tool for general administration of a Lightweight Directory Access Protocol (LDAP) directory service. |
| Repadmin | Command-line tool for monitoring replication. |

For information about how to use the ADAM ADSI Edit snap-in, see Connect and bind to an ADAM instance using ADAM ADSI Edit and Browse the directory tree using ADAM ADSI Edit.

For information about how to use Ldp, see Connect and bind to an ADAM instance using Ldp.exe and Browse the directory tree using Ldp.exe.

For more information about ADAM administration tools, see ADAM Administration Tools.

# Administering computers running ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering computers running ADAM**

Directory clients locate an Active Directory Application Mode (ADAM) instance using the name or Internet Protocol (IP) address of the computer running ADAM and the communications port in use by the ADAM instance. In addition, in Active Directory environments, clients can locate ADAM instances through service connection points (SCPs) that are published in Active Directory. Modifying the computer name, Domain Name System (DNS) name, network basic input/output system (NetBIOS) name, or IP address of a computer on which an ADAM instance is running may affect the directory clients' ability to locate that ADAM instance.

| **Note** |
| --- |
| ADAM can be installed on computers running Windows Server 2003 or Windows XP Professional. For considerations when running ADAM on Windows XP Professional, see Running ADAM on Windows XP Professional. |

**Locating ADAM instances**

For a directory-enabled application to communicate with an ADAM instance, the application must specify the NetBIOS name, DNS name, or IP address of the computer on which ADAM is running. In addition, the application must specify the Lightweight Directory Access Protocol (LDAP) or Secure Sockets Layer (SSL) communication port being used by the ADAM instance. For example, a directory-enabled application referencing an ADAM instance running on a computer with a DNS name of ADAMSrv1.microsoft.com—with a NetBIOS name of ADAMSrv and an IP address of 192.168.0.1 and using communication port 389 for LDAP connections—can address the ADAM instance using one of the following:

- ADAMSrv1.microsoft.com:389
- ADAMSrv:389
- 192.168.0.1:389

**Locating ADAM instances in Active Directory environments**

When you run ADAM in an Active Directory environment, each ADAM instance creates and maintains SCP objects about itself in Active Directory. For more information, see Administering ADAM service publication.

**Server object names**

By default, the relative distinguished name of a server object in ADAM is *servername*$*instancename*, where *servername* represents the computer on which the ADAM instance is running and *instancename* represents the service name of the ADAM instance.

**Modifying settings on a computer running ADAM**

You initially specify the computer name, DNS name, NetBIOS name, and IP address when you set up the computer. You initially specify the communication port numbers used by a given ADAM instance when you install the ADAM instance.

You can change these settings on a computer running ADAM at any time. However, when you update any of these settings, you may also need to reconfigure directory-enabled applications to properly address the ADAM instance. For example, if the host name of the computer on which ADAM is running changes, your directory-enabled application may need to be updated accordingly.

# Running ADAM on Windows XP Professional

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Running ADAM on Windows XP Professional**
In addition to running Active Directory Application Mode (ADAM) on domain controllers, member servers, or workgroup servers running Windows Server 2003, you can also run ADAM on computers running Windows XP Professional. Running ADAM on Windows XP Professional introduces some functional considerations, as described in this topic.

**Using ADAM on Windows XP Professional as a personal data store**
For use on computers running Windows XP Professional, ADAM has been designed primarily as a personal data store. It is recommended that you run ADAM on Windows Server 2003 if you need scalability to support a growing number of simultaneously connected clients or if you intend to use ADAM on a centralized directory server.

**Running ADAM on Windows XP Professional in a workgroup**
When you install ADAM on a computer running Windows XP Professional that is joined to a workgroup, you must set the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\forceguest** to 0. (The default is 1.) Otherwise, users connecting to ADAM over the network are forced to a security context of Guest, and binds to ADAM fail. For more information, see Enable binding to ADAM instances running on Windows XP Professional computers joined to a workgroup.

**Summary of functional considerations when running ADAM on Windows XP Professional**
The following table summarizes all other functional considerations related to running ADAM on Windows XP Professional.

| ADAM feature | Windows Server 2003 | Windows XP Professional |
|---|---|---|
| Support for multiple processors | Supported. | Not supported. When multiple-processor computers are present, ADAM uses only one process at a time to prevent resource monopolization. |
| Maximum cache value | No limit. | Limited to one half of physical RAM. |
| Data and log file locations | Data and log files can reside on different drives. | Data and log files must reside on the same drive. |
| Password policies | Supported. | Not supported. |
| Auditing | Supported. For more information, see Administering an ADAM instance. | Not supported. |
| ADAM service account requirements | See Selecting an ADAM service account. | See Selecting an ADAM service account. |
| Stopping and restarting an ADAM instance | An ADAM instance can be stopped and then restarted immediately. | If a Lightweight Directory Access Protocol (LDAP) client is connected when you shut down the ADAM instance, you may need to wait as long as two minutes before restarting the ADAM instance. This time is required to allow any existing connections to clear. |
| Backing up ADAM | Multiple ADAM instances running on the same computer can be backed up simultaneously. | For computers running multiple ADAM instances, only one ADAM instance can be backed up at a time. |

| | | |
|---|---|---|
| Full ADSI functionality | No hotfixes are required. | For complete Active Directory Service Interfaces (ADSI) functionality with Windows XP Professional, you must install hotfix 817583. For more information about this hotfix, see article 817583 in the Microsoft Knowledge Base. |

For information about general ADAM requirements, see ADAM requirements.

# Administering an ADAM instance

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering an ADAM instance**

Each ADAM instance runs as an independent, and separately administered, service on a computer. You can configure the account under which an ADAM instance runs, stop and restart an ADAM instance, and change the ADAM instance service display name and service description. In addition, you can enable Secure Sockets Layer (SSL) connections in ADAM by installing certificates. In Active Directory environments, each ADAM instance attempts to register a service principal name (SPN) in the directory on the account object that corresponds to the service account on which the ADAM instance is running (or, in the case of the ADAM instance running as the Network Service account, on the computer object corresponding to the computer on which the ADAM instance is running). This SPN is used for replication authentication. Depending on the network environment into which you install ADAM, you may have to create SPNs manually.

For information about tasks related to ADAM instances, see Manage an ADAM Instance. For information about SPNs, see Administering ADAM service principal names.

**ADAM service account**

The service account that is used by an ADAM instance determines the access that the ADAM instance has on the local computer and on other computers in the network. ADAM instances also use the service account to authenticate other ADAM instances in their configuration set, to ensure replication security. You determine the ADAM service account during ADAM installation. For information about ADAM service account requirements, see Selecting an ADAM service account.

To modify the ADAM service account for an ADAM instance, you can use the **change service account** command within the **Dsdbutil** command line tool. For more information, see Dsdbutil.

In addition, you may need to assign the **Log on as a service** right to the account that you specify as the ADAM service account. For more information, see Add the Log on as a service right to an account.

To enable auditing for an ADAM instance for which the service account is something other than the Network Service account, you must assign the **Generate security audits** right to the ADAM service account.

**ADAM instance name**

During setup, you assign a name to the ADAM instance, which is used in the creation of the file directory structure and registry keys for ADAM. In addition, the name that you assign is used to create the service name, service display name, and service description, as shown in the example in the following table.

| Name supplied during setup | Service name | Service display name | Service description |
|---|---|---|---|
| instance1 | ADAM_instance1 | instance1 | Blank (No default description provided.) |

The name that you specify for an ADAM instance during setup must meet the following requirements:

- It must be unique with respect to other ADAM instances running on the same computer.
- It must be no longer than 44 characters.
- It must use characters only from the ranges of a through z, A through Z, or 0 through 9.
- The name "ntds" cannot be used.

The service display name appears in Add or Remove Programs and in the Services snap-in. The service description appears in the Services snap-in. You can modify the service display name and the service description at any time after installation.

For information about modifying the service display name, see View or modify the service display name of an ADAM instance. For information about modifying the service description, see View or modify the service description of an ADAM instance.

**Starting and stopping an ADAM instance**

The service for an ADAM instance can be stopped and restarted independently, without restarting the computer on which ADAM is running. In addition, the service for an ADAM instance can be configured to start either automatically or manually, and it can also be disabled.

For information about stopping and starting ADAM, see Start, stop, pause, resume, or restart an ADAM instance. For information about assigning the **Log on as a service** right, see Add the Log on as a service right to an account.

| **Note** |
| --- |
| When you run ADAM on a computer running Windows XP Professional, you may need to wait up to two minutes after stopping an ADAM instance before restarting the instance. ADAM requires this time to close any existing connections. |

#### Event logs and auditing

Each ADAM instance running on a computer writes events to a separate event log. The ADAM event log, which has a display name that matches the name of the ADAM instance, can be viewed in Event Viewer.

To enable auditing for an ADAM instance running under a service account other than the Network Service account, you must grant the **Generate security audits** right to the ADAM service account.

To enable replication auditing for an ADAM instance, you must modify the registry key:

**HKLM\System\CurrentControlSet\Services\instance_name\Parameters**

where *instance_name* represents the name of the ADAM instance on which you want to audit replication. The following table describes the values in the registry key that control replication auditing. To enable replication auditing, set one or both of the values to 1.

| Registry key value | Data type | Meaning |
| --- | --- | --- |
| **Audit Access in Replication** | DWORD | Provides a summary of the replication operations that are occurring. |
| **Audit Objects in Replication** | DWORD | Audits the changes to individual objects and attributes. |

| **Note** |
| --- |
| ADAM also creates setup and uninstall logs in the *%systemroot%*\debug directory. These log files are called Adamsetup.log and Adamuninstall.log. |

#### Using certificates with ADAM

To enable Secure Sockets Layer (SSL) connections to ADAM, you must have certificates installed on the computer running ADAM and also on the clients connecting to ADAM. You can obtain certificates from a trusted certification authority (CA). Or, for internal use, you can issue certificates from your own internal CA, by installing and using certificate services on Windows 2000 Server or Windows Server 2003.

For information about setting up and using certificate services, see "Certificate Services" on the Microsoft TechNet Web site.

In addition, on computers running Windows XP Professional that need to establish SSL connections to an ADAM instance, you must install the hotfix described in article 817583 in the Microsoft Knowledge Base Web site.

| **Note** |
| --- |
| When using certificates with ADAM, you must assign the ADAM service account read access to the certificates that are installed on the computer where the ADAM instance is running. The certificate store is located in the following directory: Documents and Settings\All Users\Application Data\Microsoft\Crypto\RSA\MachineKeys. |

# Administering directory partitions

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering directory partitions**

Three types of directory partition exist in Active Directory Application Mode (ADAM): configuration, schema, and application. Configuration and schema directory partitions are created automatically during installation. Application directory partitions can be created either during or after installation. Every ADAM directory partition has a unique distinguished name. Distinguished names for directory partitions in ADAM support both Domain Naming System (DNS)-style and X.500-style naming. An ADAM instance does not specify a naming context by default, but it can be configured to provide a default naming context to clients. ADAM configuration and schema directory partitions can be viewed and managed using ADAM administration tools, which are installed during ADAM setup. Application directory partitions are typically managed through directory-enabled applications.

For information about tasks related to directory partitions, see Manage Schemas and Directory Partitions. For more information about types of directory partitions, see Understanding ADAM data and data stores.

**Creating directory partitions**

The schema and configuration directory partitions are created automatically during installation of an ADAM instance. (Configuration and schema partitions can only be created during ADAM installation.) If you choose to join the ADAM instance being installed to an existing configuration set, the schema and configuration directory partitions from an ADAM instance in the configuration set are replicated to the new instance. If you install a new ADAM instance, the installation process creates default versions of the schema and configuration directory partitions for you. If you install an ADAM instance from media, the schema and configuration directory partitions for the new ADAM instance are replicas of the directory partitions from the media.

When you install a new ADAM instance, you can also choose to replicate one or more application directory partitions from an existing ADAM instance. Or, you can specify the creation of a new application partition. When you create a new application directory partition during installation, you must specify a unique distinguished name for the partition. The new application directory partition consists only of a partition container object with the distinguished name that you specify.

**Naming directory partitions**

Each ADAM directory partition has its own, unique distinguished name. ADAM supports both DNS-style and X.500-style names for top-level directory partitions, including the distinguished name components that are listed in the following table.

| Distinguished name attribute | Meaning |
|---|---|
| C= | Country/region |
| CN= | Common name |
| DC= | Domain component |
| L= | Location |
| O= | Organization |
| OU= | Organizational unit |

**Specifying a default naming context**

By default, an ADAM instance does not provide a default naming context. You can, however, configure ADAM to provide a default naming context, by specifying a value for the **msDS-defaultNamingContext** attribute on the **NTDSA** object, at:

This **msDS-defaultNamingContext** attribute contains no value by default. If a value is set in the attribute, that value is returned to a client through the **defaultNamingContext** attribute of the **rootDse** object when the **rootDse** object is read.

When a bind is requested on the **rootDSE** object, ADAM returns the value of **defaultNamingContext**, if the **msDS-defaultNamingContext** attribute on the **NTDSA** object is set to the distinguished name of a top-level container of an instantiated naming context. Otherwise, no value is returned.

**Importing and exporting directory objects**

You can use a number of methods to import and export directory objects, both to and from the schema directory partition, as well as to and from application directory partitions. You can use the **ldifde** command-line tool to import and export .ldf files.

You can use the **csvde** command-line tool to import and export .csv (comma-delimited) files. For more information, see Import or export directory objects using ldifde and Import or export directory objects using csvde.

**Administering application directory partitions**

You can create an application directory partition during setup. You can also create application directory partitions at any time after installation, either through ADAM administration tools or through your directory-enabled application. In production environments, you typically manage your application directory partitions and the data that they contain through your directory-enabled applications.

For information about creating application directory partitions, see Add an application directory partition.

# Administering sites

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering sites**

You can optimize the replication efficiency and reduce the administrative overhead of your network by establishing sites appropriately. The most effective number of sites depends on the physical design of your network. When you first create a new configuration set for Active Directory Application Mode (ADAM), a single, default site (called Default-First-Site-Name) is created that represents your entire network. A single site can be very efficient for a single location network that is completely connected by high-speed bandwidth. If your network contains multiple geographic locations that communicate over low-speed, wide area network (WAN) connections, establishing multiple sites gives you more detailed control of replication behavior, and it reduces network traffic on the WAN that connects the sites. You can administer site objects and site link objects through ADAM ADSI Edit.

For information about sites, see Understanding ADAM sites. For information about tasks related to sites, see Manage Replication, Sites, and Configuration Sets.

**Bandwidth as an important consideration**

Within a site, bandwidth affects how efficiently replication can work. The frequency with which intrasite replication occurs requires high-speed bandwidth to function most effectively. Before you create a new site, make sure that high-speed bandwidth connects all computers in the site candidate. Any area where domain controllers are connected by 10 megabits per second (Mbps) or more of bandwidth is a good site candidate.

**When to establish a single site**

If you have a single local area network (LAN) consisting of a single location, or if your network consists of a series of closely located locations (as in a campus), in which all locations are connected by a high-speed backbone, using the default site as the single site in replication topology can simplify replication management. In a single-site topology, all replication on your network occurs as intrasite replication, which requires no manual replication configuration. In a single-site design, all ADAM instances also remain very current with respect to directory changes, because directory updates are replicated almost immediately.

**When to establish multiple sites**

When your network consists of multiple geographic locations connected by a WAN, establishing separate sites for each location provides the following benefits:

- Efficient use of WAN bandwidth for replication
- Detailed control of replication behavior

Physically separate network locations typically communicate over WAN connections, which are most often characterized by low-speed bandwidth. By creating a separate site for each physical location on your network, you ensure that ADAM instances communicating over WAN connections use intersite replication, which is specifically designed for efficiency over low-bandwidth connections. When you create a site object, you also need to create a servers container within the site object. To move an ADAM instance to a new site object, you move the ADAM instance into the servers container of the site. For more information, see Add a site object and servers container and Move an ADAM instance to a site object.

Site link objects represent the network links between the sites on your network. When you create a new site object in ADAM, add the site object to a site link object. Until you add the new site to a site link object, any ADAM instances in the site will not be able to replicate to ADAM instances outside the site. You can add a new site object to the default site link object. Or, you can create additional site link objects, giving you more detailed control over the relative cost of replication over different WAN connections on your network.

With multiple sites, you have more detailed control of replication behavior through configurable intersite replication settings. These settings include the relative cost of different replication paths and the availability of connections for use by replication. For more information, see Configure replication availability on a site link and Set the cost on a site link object.

# Administering replication and configuration sets

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering replication and configuration sets**
In Active Directory Application Mode (ADAM), replication within a site occurs automatically, and it does not require any configuration beyond the construction of configuration sets. For replication between sites, you can specify a schedule for replication. For example, you may want replication to occur only during the night, when your network is not busy. If the ADAM instance replication partners do not all belong to Active Directory domains, you must use identical local service account names and passwords to enable replication.

### Notes

- You can only join an ADAM instance to a configuration set using the ADAM Setup Wizard or during an unattended ADAM installation. After an ADAM instance is installed, it cannot be joined to a configuration set.
- ADAM does not support replication using Simple Mail Transfer Protocol (SMTP).

**Scheduling replication**
For replication between ADAM instances within a site, you can schedule the amount of time that elapses after the last change notification before a default replication event occurs. For more information, see Configure replication frequency within a site.

For replication between ADAM instances in different sites, you can schedule the availability of the site link that exists between the two sites. For more information, see Configure replication availability on a site link.

For more information about replication within and between sites, see Understanding ADAM sites and replication.

**Authenticating ADAM replication partners**
Considerations for replication authentication in ADAM include the authentication level of the configuration set, which is determined by the domain, forest, or workgroup environment, and the service accounts that are in use by the ADAM instances in the configuration set. For information about authentication levels, see Understanding ADAM replication and configuration sets. For information about service account requirements, see Selecting an ADAM service account.

**Operations master roles**
Each ADAM instance belonging to a given configuration set operates with equal authority. However, certain activities that are performed with a given configuration set should only be performed by a single ADAM instance. These activities include schema management, to prevent conflicting schema changes, and instance naming management, to prevent conflicting directory partition names. The ADAM instances that are responsible for these activities are called, respectively, the schema master and the naming master. ADAM assigns these roles automatically. You can also reassign these roles, as may be required, for example, if an ADAM instance that holds one of the operations master roles stops responding or goes offline. For information about reassigning operations master roles, see Identify or transfer the schema master role and Identify or transfer the naming master role.

**Schema master**
The schema master instance controls all updates and modifications to the schema within a given configuration set. To update the schema, you must be a member of the ADAM Administrators group on the configuration set. You must also be able to connect successfully to the schema master. Each configuration set can contain only one schema master.

**Naming master**
The ADAM instance that holds the naming master role controls the addition and removal of directory partitions, or *naming contexts*, in a given configuration set. Each configuration set can contain only one naming master.

# Administering ADAM service publication

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering ADAM service publication**
In Active Directory environments, Active Directory Application Mode (ADAM) uses service connection point (SCP) objects to publish ADAM service information in the Active Directory directory store. Considerations for administering ADAM SCP objects include the following:

- Manually creating SCPs when necessary
- Removing SCP objects when necessary
- Modifying permissions on SCP objects when necessary
- Checking the ADAM event log for SCP errors

**Creating SCPs**
By default in Active Directory environments, the ADAM service creates SCPs in Active Directory when the ADAM service starts. To successfully create SCPs, the ADAM service account must have sufficient rights in Active Directory. If SCP creation fails, ADAM writes an event to the ADAM event log regarding the failure.

### Notes

- The default location for the ADAM SCP object is under the computer object that represents the computer on which ADAM is running. This default location can be altered by specifying a different location on the **SCPPublishingService** object.
- When it is used as the ADAM service account on a computer that is joined to a domain, the Network Service account usually has sufficient rights to create SCPs in Active Directory. However, when it is used on an ADAM instance running on a domain controller, the Network Service account does not have sufficient rights to create SCPs in Active Directory.

**Removing SCPs**
An ADAM instance deletes its SCP from Active Directory when you remove the ADAM instance from the computer. Removing the SCP requires sufficient administrative privileges. If SCP removal fails, client applications may be directed to a nonexistent ADAM instance.

**Modifying SCPs**
An ADAM instance checks and updates, if necessary, its SCP each time the ADAM instance starts. At startup, the ADAM instance searches the global catalog for its own globally unique identifier (GUID) and retrieves the distinguished name of the SCP object. The ADAM instance then binds to that distinguished name and updates the SCP object as necessary. In addition, the ADAM instance reviews the SCP object on an hourly basis by default to confirm its validity, particularly regarding any directory partitions that have been added to or removed from the ADAM instance since the SCP object was last updated.

| **Note** |
| --- |
| You can modify the default time interval at which ADAM reviews the SCP object, by adding a value named **Server information update interval (mins)** to the registry key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services*instancename*\Parameters, and setting this value to the time interval (in minutes) that you want to use. |

For more information about publishing and using SCPs, see "About Service Publication" on the MSDN Web site at http://go.microsoft.com/fwlink/?LinkId=48288.

**Managing ADAM SCPs**
One efficient way to manage SCPs for ADAM is to create an SCP container for your ADAM configuration set in Active Directory. In this container, place the computer objects of the computers on which ADAM instances are running. In addition, create a group called, for example, "ADAM instances," and assign permissions on the SCP container to the group. Also, delegate control

of the group to the assigned ADAM administrator. Then, each time the ADAM administrator installs a new ADAM instance under a new service account, the ADAM administrator can simply add the new service account to the "ADAM instances" group, and the ADAM instance creates and maintains its SCPs transparently.

**Checking the ADAM event log**

SCP-related errors do not prevent an ADAM instance from functioning properly. ADAM does, however, report SCP-related errors in the ADAM event log so that errors can be resolved. ADAM reports SCP-related errors as follows:

- When an SCP creation or update fails, the ADAM instance reports the error and points to a .ldf file that can be used to manually resolve the problem. Using this .ldf file to create or update the SCPs requires administrator privileges in Active Directory.

- When ADAM fails to remove an SCP object, ADAM reports this failure both in the audit log of the appropriate domain controller and in the ADAM event log.

# Administering ADAM service principal names

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering ADAM service principal names**
When they run in an Active Directory domain or forest, Active Directory Application Mode (ADAM) instances attempt to register service principal names (SPNs) in Active Directory for use in Kerberos authentication during replication. Each ADAM instance—the first time that the instance starts—attempts to register values to the **Service-Principal-Name** attribute on the account object that represents the service account in use by the ADAM instance. Or, if the ADAM instance uses the Network Service account as the ADAM service account, the ADAM instance attempts to register values to the **Service-Principal-Name** attribute on the computer object that represents the computer on which the ADAM instance is running. The values that the ADAM instance attempts to register include the following:

- E3514235-4B06-11D1-AB04-00C04FC2DCD2-ADAM/netbiosname:port
- E3514235-4B06-11D1-AB04-00C04FC2DCD2-ADAM/dnshostname:port
- Ldap/netbiosname:port
- Ldap/dnshostname:port

In addition, if the ADAM instance uses a reserved communications port (LDAP_PORT, LDAP_GC_PORT, LDAP_SSL_PORT, or LDAP_SSL_GC_PORT), the instance attempts the following SPN registrations:

- Ldap/netbiosname
- Ldap/dnshostname

Wldap32.dll does not use a port-formatted SPN to establish a connection to an ADAM instance using a reserved port. For example, if ADAM instance1 on machine1 uses port 3268 for LDAP, the SPN that Wldap32.dll creates for the connection is ldap\machine1. However, if ADAM instance2 on machine2 uses port 1026 for LDAP, the SPN that Wldap32.dll creates for the connection is ldap\machine2:1026.

| **Note** |
| --- |
| When multiple ADAM instances are running on a single computer, and two or more of the instances each use a reserved port, those ADAM instances must run under the same ADAM service account for the SPNs to be registered correctly. |

SPN registration attempts occur in the security context of the ADAM service account. If the ADAM service account that you specify is not the Network Service account, or if it is a domain user account that does not belong to the Domain Administrators group, this SPN registration fails. ADAM reports this failure as event ID 2516 in the event log for the ADAM instance. If this failure occurs, ADAM setup also reports event ID 2535, which describes how to register the SPNs for ADAM manually.

To register SPNs for ADAM manually in Active Directory, use the *dnsdomainname*.bat script file that ADAM setup creates in the data directory of the ADAM instance (Program Files\Microsoft ADAM*instancename*\data), where *dnsdomainname* represents the name of the DNS domain in which the ADAM instance resides.

**Notes**

- The *.bat file shows up in the directory about one minute after ADAM installation completes.
- The *.bat file contains **repadmin /writespn** commands, similar to the following:

```
repadmin.exe /writespn hostname.microsoft.com
ADD "CN=adam,CN=Users,DC=microsoft,DC=com"
```

```
E3514235-4B06-11D1-AB04-00C04FC2DCD2-ADAM/netbioshostname:389

repadmin.exe /writespn hostname.microsoft.com
ADD "CN=adam,CN=Users,DC=microsoft,DC=com"
E3514235-4B06-11D1-AB04-00C04FC2DCD2-ADAM/dnshostname:389

repadmin.exe /writespn hostname.microsoft.com
ADD "CN=adam,CN=Users,DC=microsoft,DC=com" ldap/netbioshostname:389

repadmin.exe /writespn hostname.microsoft.com
ADD "CN=adam,CN=Users,DC=microsoft,DC=com" ldap/dnshostname:389
```

This script registers the appropriate SPNs in Active Directory, and it must be run by a member of the Domain Administrators group in the domain.

| **Note** |
| --- |
| When SPNs for ADAM instances in a configuration are registered in Active Directory, ADAM replication authentication uses Kerberos. Otherwise, ADAM uses negotiated replication authentication. For more information about replication authentication, see Understanding ADAM replication and configuration sets. |

**Removing ADAM SPNs from Active Directory**

When you remove an ADAM instance, make sure that its associated SPNs are also removed from Active Directory. These SPNs reside on the computer object of the computer where the ADAM instance is installed (if the Network Service account is specified for the ADAM service account) or on the domain user object (if a domain user account is specified for the ADAM service account).

**Forcing the use of SPNs in an Active Directory environment**

To force the use of SPNs and Kerberos for replication authentication in a a configuration set, you can modify the replication security level of the configuration set to equal 2. For more information, see Modify the replication security level of a configuration set.

| **Caution** |
| --- |
| If you set replication security level equal to 2, and SPNs are not registered or properly configured, replication will fail. |

# Administering the ADAM schema

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering the ADAM schema**
You can administer the Active Directory Application Mode (ADAM) schema in a number of ways, including extending the schema with new classes and attributes, deactivating classes and attributes, and reactivating classes and attributes. You can use the ADAM Schema snap-in to administer the ADAM schema.

For information about tasks related to administering the schema, see Manage Schemas and Directory Partitions.

**Extending the schema**
Before you can use ADAM to support a particular directory-enabled application, you may need to modify the ADAM schema with classes and attributes that are required by that application. You can easily modify the ADAM schema by using a number of methods, including:

- Automatically, through a directory-enabled application.
- Semiautomatically, using Ldifde.exe and the importation of .ldf files.
- Programmatically, through Active Directory Service Interfaces (ADSI).
- Manually, through Ldp and the ADAM Schema snap-in.

| **Note** |
| --- |
| Before you can use the ADAM Schema snap-in, you must first install it. For more information, see Install the ADAM Schema snap-in. |

For environments that require the ability to create users in the ADAM directory, ADAM includes an importable .ldf file that contains four user object classes. For information about importing the ADAM user .ldf file, see Import the user classes supplied with ADAM.

**Deactivating classes and attributes**
Before deactivating a class, consider the following:

- You can deactivate a class only if that class is not specified as a **subClassOf**, **auxiliaryClass**, **systemAuxiliaryClass**, **possSuperiors**, or **systemPossSuperiors** of any existing active class.
- You cannot use a defunct class in definitions of new classes, and you cannot add a defunct class to existing class definitions.
- You cannot create objects that are instances of a defunct class or modify existing instances of a defunct class. However, the instances of the defunct class become available again for modification when the defunct class is reactivated.

Before deactivating an attribute, consider the following:

- You can deactivate an attribute only if the attribute is not specified as a **systemMustContain**, **mustContain**, **systemMayContain**, **mayContain**, or **rdnAttId** of any existing active class.
- You cannot use a defunct attribute in definitions of new classes, and you cannot add a defunct attribute to existing class definitions.
- You cannot read, modify, or delete instances of a defunct attribute that are present in existing objects. However, the instances of the defunct attribute become available when the defunct attribute is reactivated.

- To purge the directory of instances of an attribute, you must delete the instances before deactivating the attribute.

# Administering users and groups

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering users and groups**

You can administer users and groups in Active Directory Application Mode (ADAM) through the ADAM ADSI Edit snap-in or through your directory-enabled applications. For information about how to perform tasks related to users and groups in ADAM, see Manage Users and Groups.

**ADAM users**

To create users in ADAM, you must first import the optional user classes that are provided with ADAM into the ADAM schema. These user classes are provided in importable .ldf files (ms-User.ldf, ms-InetOrgPerson.ldf, ms-UserProxy.ldf), which you can find in the directory %windir%\adam on the computer where ADAM is installed. The following table describes the contents of each of the optional .ldf files.

| .ldf file | User classes | Import this file if ... |
|---|---|---|
| ms-user.ldf | • **Person**<br>• **Organizational-Person**<br>• **User** | You want to create user objects in the ADAM directory, but you do not want to create users of the **inetOrgPerson** class (as defined in RFC 2798). |
| ms-inetorgperson.ldf | • **Person**<br>• **Organizational-Person**<br>• **User**<br>• **inetOrgPerson** | You want to create user objects in the ADAM directory, and you want to create users of the **inetOrgPerson** class (as defined in RFC 2798). |
| ms-userproxy.ldf | • **User-Proxy** | You want to create proxy objects in ADAM for use in bind redirection. For more information about proxy objects and bind redirection, see Understanding ADAM bind redirection. |

For information about the schema definitions for each of the user object classes, see "Active Directory Schema" on the Microsoft MSDN Web site.

For more information about importing user classes, see Import the user classes supplied with ADAM. For information about creating ADAM users, see Add an ADAM user to the directory.

You can also disable ADAM users. For more information about disabling ADAM users, see Disable or enable an ADAM user.

**Setting and modifying passwords**

You can set and modify passwords for ADAM security principals over Secure Sockets Layer (SSL) connections (using Ldp.exe) or over encrypted, non-SSL connections (using ADAM ADSI Edit or Ldp.exe). To establish an SSL connection to ADAM, you must install certificates on the computer running ADAM and on all the clients. For more information about certificates, see Administering an ADAM instance. To make SSL connections to an ADAM instance, you must use Ldp.exe; ADAM ADSI Edit does not support SSL connections.

For information about setting and modifying passwords, see Set or modify the password of an ADAM user.

| **Note** |
|---|
| On computers running Windows XP Professional that need to establish SSL connections to an ADAM instance, you must install the hotfix described in article 817583 in the Microsoft Knowledge Base. |

By default, an ADAM instance running on Windows Server 2003 automatically enforces any local or domain password policies that exist. If you create a new ADAM user, and if you assign a password to that user that does not meet the requirements of the password policy that is in effect, the user will be disabled.

**Password policy settings and account lockouts**
When ADAM runs on a computer running Windows Server 2003, it supports and enforces the password policy settings and account lockout settings that are provided by Windows Server 2003, including the following:

- Minimum age
- Maximum age
- Complexity
- History
- Too many failed logon attempts
- Disabling and enabling of accounts

If the server on which ADAM is running belongs to a workgroup, the server's local password policy settings and account lockout settings are implemented. If the server on which ADAM is running belongs to a domain, the password policy settings and account lockout settings from Active Directory are implemented. Password policy settings are not enforced on ADAM instances that run on a computer running Windows XP Professional.

You can disable the enforcement of password policy settings in ADAM by setting **ADAMDisablePasswordPolicies**, a value in the attribute **msDS-Other-Settings** on CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,CN=GUID, to 1.

For more information about password policy settings and account lockout settings in Windows Server 2003, see Help and Support Center in Windows Server 2003.

**ADAM groups**
In addition to the four role-based groups that are provided by default in each directory partition in ADAM, you can add your own groups. You can add both ADAM users and Windows users to the ADAM groups that you create.

For more information, see Add an ADAM group to the directory and Add or remove members to or from an ADAM group.

| **Note** |
| --- |
| In ADAM, it is possible for an ADAM administrator, or for a user with sufficient access to the Administrators group, to remove member accounts from the ADAM Administrators group, possibly leaving ADAM without any valid administrators. To recover from this scenario, the assigned ADAM administrator, as the owner of the Administrators group, can repopulate the ADAM Administrators group with the appropriate accounts. |

**Organizing ADAM users and groups with organizational units**
To keep your ADAM users and groups organized, you may want to place users and groups in organizational units (OUs). In Active Directory and in ADAM, as well as in other Lightweight Directory Access Protocol (LDAP)-based directories, OUs are the most commonly used method for keeping users and groups organized. For information about creating OUs in ADAM, see Add an organizational unit to the directory.

| **Note** |
| --- |
| By default, OUs can only be added under other OU (OU=), country/region (C=), organization (O=), or domain-DNS (DC=) object classes. For example, you can add an OU to O=Microsoft,C=US. You cannot add an OU to CN=test,O=Microsoft,C=US. You can, however, update the schema definition of the OU object class to allow other superiors. |

# Administering access control

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering access control**

You can govern access control in Active Directory Application Mode (ADAM) at the directory partition level by assigning users memberships to the role-based groups that are located on each partition. You can customize access control in ADAM on an object-by-object basis using the **dsacls** command-line tool. Other considerations for administering access control in ADAM include the assignment of permissions to Windows security principals and the best practice of avoiding the assignment of permissions to built-in security principals.

For more information about the **dsacls** command-line tool, see Dsacls.

**Default ACLs**

By default, ADAM provides access control list (ACL) settings on each directory partition. Also by default, all objects on a given partition inherit these default ACLs. To customize access control, you can:

- Modify the default ACLs on the partition objects.
- Set new ACLs on individual objects.
- Change inheritance settings on individual objects.

For a table describing the default access that is provided to each of the default ADAM groups, see Understanding ADAM users and groups.

For information about viewing and setting object permissions, see View or set permissions on a directory object.

| **Note** |
| --- |
| In ADAM, it is possible for an ADAM administrator, or for a user with sufficient access to the Administrators group, to remove member accounts from the ADAM Administrators group, possibly leaving ADAM without any valid administrators. To recover from this scenario, the assigned ADAM administrator, as the owner of the Administrators group, can repopulate the ADAM Administrators group with the appropriate accounts. |

**Custom authorization mechanisms**

Depending on your applications, you may have applications that use a custom authorization scheme. With a custom authorization scheme, ADAM still performs the task of authenticating users, but the directory-enabled application manages its own set of authorization information. Even with a custom authorization scheme in place, ADAM continues to enforce the ACLs on ADAM objects. Therefore, if your application implements its own authorization scheme, you may want to assign more relaxed permissions on all ADAM objects, so that the ADAM permissions do not interfere with your application's authorization scheme. If you do relax permissions in ADAM in this way, be sure to prevent clients from gaining access to ADAM through Lightweight Directory Access Protocol (LDAP), thus bypassing your custom authorization mechanisms.

**Role-based authorization**

With Windows Server 2003, one way that a directory-enabled application can control access to directory objects is through role-based authorization, which is an alternative to ACL-based access control. In this scenario, a directory-enabled application does not expose LDAP but instead controls all access to the directory data. ADAM is used simply as a data store. For more information about role-based authorization in Windows Server 2003, see the Microsoft MSDN Web site.

| **Note** |
| --- |
| Using ADAM as an authorization store requires that you import ms-azman.ldf into the ADAM schema. You can import ms-azman.ldf through the ADAM Setup Wizard during installation or by using the **ldifde** command-line tool. For more information, see Importing LDIF files during setup. |

**Windows security principals**

By default, authenticated Windows security principals in ADAM can only read objects in the schema directory partition. To enable authenticated Windows security principals to read any other objects, you can assign permissions on objects to the well-

known security ID (SID) **authenticated user**. You can assign Read permissions for an entire directory partition by making **authenticated user** a member of the Readers group on that directory partition. Or, you can assign Read permissions on an object-by-object basis, using **dsacls**.

**Best practices for assigning permissions**

Although it is possible to assign permissions on objects to built-in security principals, it is recommended as a best practice that you avoid doing so. Built-in security principals represent user groups that are defined by the security system. These built-in groups take the form BUILTIN*groupname*. For example, the Builtin Administrators group (SID S-1-5-32-544) exists on every computer. If permission assignments for a particular object are made to the Builtin Administrators group, the intent of this assignment cannot always be determined. In other words, ADAM cannot always tell whether a security principal requesting access to an object belongs to the Builtin Administrators group for which a permission assignment is intended.

It is also possible to assign permissions on objects to well-known security principals. Well-known security principals represent special identities that are defined within the security system, including Everyone, Local System, Principal Self, Authenticated User, Creator Owner, and so on. It is generally considered acceptable to assign permissions on objects to any of the following well-known security principals:

- Anonymous Logon (after considering the possible security effects)
- Authenticated Users (S-1-5-11)
- Network (S-1-5-2)
- Interactive (S-1-5-4)
- Creator Owner (S-1-3-0)
- Self (S-1-5-10)

As a best practice, avoid assigning permissions on objects to the following well-known security principals:

- Network Service (S-1-5-20)
- Local Service (S-1-5-19)
- Local System (S-1-5-18)

The assignment of permissions on objects to any other well-known security principal should be made on a case-by-case basis, after considering the needs of the application.

# Backing up and restoring ADAM

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Backing up and restoring ADAM**
Active Directory Application Mode (ADAM) data and log files should be backed up regularly to ensure the continued availability of data to applications and users in the event of a system failure.

**Backing up ADAM**
By default, each instance of ADAM running on an ADAM server stores its database file, Adamntds.dit, and the associated log files in %program files%\Microsoft adam*instancename*\data. These files should be included as part of the regular backup plan of your organization. You can back up the directory stores using Backup or any Windows Logo Program, third-party backup utility. For information about backing up an ADAM instance, see Back up an ADAM instance to a file or a tape.

| **Note** |
| --- |
| You can use a restored backup of an ADAM instance to install an ADAM replica from media. For more information, see Install an ADAM replica from media. |

**Restoring ADAM**
When you restore a database to an existing ADAM instance, you must stop the ADAM instance before you run the restore operation. In addition, it is recommended that you move (or delete) the existing database and log files from the ADAM instance before the restore operation.

**Authoritative restore**
If objects in the directory are inadvertently deleted or modified, and if those objects are replicated in a configuration set, you must authoritatively restore those objects so that the correct version of the objects are replicated. To authoritatively restore directory data, run the **dsdbutil** utility after you have restored the data but before you restart the ADAM instance. With **dsdbutil**, you can mark directory objects for authoritative restore. When an object is marked for authoritative restore, its update sequence number is changed so that the number is higher than any other update sequence number in the configuration set. This ensures that any data you restore is properly replicated throughout the configuration set.

For more information about restoring an ADAM instance, including authoritative restores, see Restore an ADAM instance.

## Notes

- If you restore an ADAM backup over a running ADAM instance, Backup leaves the restored files in a pending state, and it does not write the files to disk until the computer is rebooted. In this situation, any directory changes that are made to the running ADAM instance after Backup is run are lost.

- If you accidentally start a restore of an ADAM instance over a currently running ADAM instance, it is recommended that you immediately restart the computer, stop the ADAM instance, and then perform the restoration again.

# Selecting an ADAM service account

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Selecting an ADAM service account**

On the **Service Account Selection** page of the ADAM Setup Wizard, you must select a service account for use by the Active Directory Application Mode (ADAM) instance. The account that you select determines the security context in which the ADAM instance runs. Changing the service account after installation may require some additional configuration.

| **Note** |
| --- |
| The first ADAM instance in a configuration set determines the default replication authentication method. |

**Service account requirements**

ADAM runs as a service, and it requires a service account. ADAM service account requirements depend on the Windows workgroup or domain environment into which you install ADAM, as well as the computer on which ADAM is running.

For ADAM instances that are joined to a configuration set, the service account is also used to authenticate against other ADAM instances in the configuration set for replication. The type of authentication that is used between replication partners is determined by the environment in which ADAM is running and by the service accounts in use. For more information, see Understanding ADAM replication and configuration sets.

The following table outlines ADAM service account requirements.

| **Security context** | **Service account for first ADAM instance** | **Service account for replica ADAM instances** | **Default replication authentication method**** |
| --- | --- | --- | --- |
| Workgroup | Network Service | Replica ADAM instances not allowed | Not applicable |
| Workstation user | Workstation user | Negotiated pass-through* | Not applicable |
| Windows 2000 domain or forest<br><br>-or-<br><br>Windows Server 2003 domain or forest | Network Service<br><br>-or-<br><br>Domain user | Network Service<br><br>-or-<br><br>Domain user | Negotiated |
| Workstation user | Workstation user | Negotiated pass-through* | Not applicable |
| Domain user | Domain user<br><br>-or-<br><br>Network Service | Negotiated | Not applicable |
| Windows NT 4.0 domain | Workstation user<br><br>Domain user | Workstation user<br><br>Domain user | Negotiated pass-through*<br><br>Negotiated |

*When a workstation user account is used on the first ADAM instance in a configuration set, all subsequent ADAM instances in the same configuration set must use an identical local workstation account name and password as the ADAM service account.

**When the Network Service account is used as the ADAM service account, the replication authentication mode is set to Negotiated by default.

**Notes**

- The Network Service account is a special, built-in account, with authority similar to that of an authenticated user account. The name of the account is NT AUTHORITY\NetworkService. The Network Service account has limited access to the local computer and authenticated access (as the computer account) to network resources. This limited access helps safeguard

your system if individual services or processes are compromised. Services that run as the Network Service account access network resources using the credentials of the computer account.

- To enable auditing for an ADAM instance running under a service account other than the Network Service account, you must grant the **Generate security audits** right to the account that is used as the ADAM service account.

- To enable a workstation or domain user account as a service account, you must grant the **Log on as a service** right to the account that is used as the ADAM service account. For more information, see Add the Log on as a service right to an account.

- The account that is used as the ADAM service account must be able to create, read, and modify files in the directory %ProgramFiles%\Microsoft ADAM*instancename*\data.

For more information, see Understanding ADAM architecture.

# ADAM troubleshooting and frequently asked questions (FAQs)

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Troubleshooting and frequently asked questions (FAQs)**

| |
|---|
| **The installation or removal of Active Directory Application Mode (ADAM) fails to complete successfully.**<br><br>If no screen message appears and ADAM setup fails to complete successfully, you can view the setup log at:<br><br>• %windir%\Debug\adamsetup.log<br><br>If no screen message appears and ADAM removal fails to complete successfully, you can view the uninstall log at:<br><br>• %windir%\Debug\adamuninstall.log |
| **A directory-enabled application cannot find an ADAM instance.**<br><br>Be sure to reference the correct communications port number (that is, 389 or 636) when specifying an ADAM instance. |
| **The ADAM service will not start.**<br><br>If the service account that is specified for ADAM is a workstation or a domain user account, make sure that the account possesses the **Run as a service** right. |
| **In an attempt to connect to an ADAM instance using the ADAM Schema snap-in, the following error message appears: "The schema FSMO holder could not be found. Schema modifications can only be made on the schema FSMO holder."**<br><br>You can only modify the schema on the ADAM instance that holds the schema master role. To determine the holder of the schema master role for a configuration set, see Identify or transfer the schema master role. |
| **How do I set or modify the password of an ADAM user?**<br><br>For information about modifying passwords, see Set or modify the password of an ADAM user. |
| **I can't bind to an ADAM instance that runs on a computer running Windows XP Professional.**<br><br>When the computer that runs ADAM is joined to a workgroup, you must set the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\forceguest** to 0. (The default is 1.) Otherwise, users connecting to ADAM over the network are forced to a security context of Guest, and binds to ADAM fail. For more information, see Enable binding to ADAM instances running on Windows XP Professional computers joined to a workgroup.<br><br>For information about other functional considerations when running ADAM on Windows XP Professional, see Running ADAM on Windows XP Professional. |
| **I can't connect to an ADAM instance over an SSL connection.**<br><br>To enable SSL connections, you must first install certificates on the computer running ADAM and on all client computers. For information about setting up and using certificate services, see "Certificate Services" on the Microsoft TechNet Web site.<br><br>On computers running Windows XP Professional that need to establish SSL connections to an ADAM instance, you must install the hotfix that is described in article 817583 in the Microsoft Knowledge Base Web site.<br><br>For information about other functional considerations when running ADAM on Windows XP Professional, see Running ADAM on Windows XP Professional. |

**I can't create an organizational unit (OU) in my application directory partition.**

OUs can only be created under the following objects:

- OU (ou=)
- country/region (c=)
- organization (o=)
- domain (dc=)

For example, you can create an OU in an application partition named O=Microsoft,C=US, but you cannot create an OU in an application partition named l=Microsoft,C=US.

**I can't locate the security IDs (SIDs) for an ADAM or Windows security principal.**

You can retrieve an active user's individual and group SIDs by explicitly querying the **tokenGroups** attribute on the **rootDSE**.

**Do multivalued attributes in ADAM have a limit to the number of values that they can hold?**

ADAM (as with Active Directory running at Windows Server 2003 forest functional level) places no limits on the number of linked attribute values that a multivalued attribute can hold. For nonlinked values, it is recommended that a multivalued attribute contain no more than 1,500 values.

For information about linked attributes, see "Linked Attributes" on the Microsoft MSDN Web site. For information about **member** as an example of a multivalued attribute in Active Directory, see "Enumerating Groups That Contain Many Members" on the Microsoft MSDN Web site.

**How can ADAM be configured to support anonymous LDAP binds?**

ADAM does not accept anonymous bind requests by default. To enable anonymous LDAP operations in ADAM, you must set the seventh character of the **dsHeuristics** value to 2.

You can find the **dsHeuristics** attribute on the **Directory Service** object in the configuration directory partition (CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,CN={*GUID*}).

Modify the attribute to the following value:

0000002001000

In addition, assign permissions so that anonymous users have access to the appropriate objects in the directory. To grant the Read permission on all objects in a given directory partition to anonymous users, you can simply add the built-in security principal Anonymous (from the local computer) to the Readers group on that directory partition.

# ADAM Administration Tools

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ADAM administration tools**

- ADAM ADSI Edit
- ADAM Csvde
- ADAM Ldifde
- Adaminstall
- Adamsync
- Adamuninstall
- ADSchemaAnalyzer
- Dsacls
- Dsdbutil
- Dsdiag
- Dsmgmt
- Ldp
- Repadmin

# ADAM ADSI Edit

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ADAM ADSI Edit**
ADAM ADSI Edit runs as a Microsoft Management Console (MMC) snap-in. To start ADAM ADSI Edit, click **Start**, point to **All Programs**, point to **ADAM**, and then click **ADAM ADSI Edit**.

The following sections describe user interface (UI) options in the ADAM ADSI Edit snap-in, depending on the node that you click in the console tree or the command that you click in the **Action** menu.

- ADAM ADSI Edit Node
- Connection Node
- Object Node

**ADAM ADSI Edit Node**
To view the following UI options, in the console tree, click **ADAM ADSI Edit**, and then, on the **Action** menu, click:

**Connect to**
The following dialog box appears:

### Connection Settings

You can use the following options in the **Connection Settings** dialog box to create a connection to a directory partition on an ADAM instance:

**Connection name.** Type a display name for the connection. The connection will be listed in the console tree using the **connection name** that you provide.

**Server name.** The computer name, or DNS host name, of the ADAM instance to which you want to connect.

**Port.** The Lightweight Directory Access Protocol (LDAP) communications port number in use by the ADAM instance to which you want to connect. The default LDAP port is 389.

| **Note** |
| --- |
| ADAM ADSI Edit does not support Secure Sockets Layer (SSL) connections. |

**Distinguished name (DN) or naming context.** Click this option if you want to specify the distinguished name of the directory partition to which you want to connect.

**Well-known naming context.** Click this option if you want to connect to one of the well-known naming contexts, which include **Configuration**, **RootDSE**, and **Schema**.

**The account of the currently logged on user.** Click this option if you want to connect to the ADAM instance using the security context of the account that you are currently using.

**This account.** Click this option if you want to specify a different account to use when connecting to the ADAM instance.

**Refresh**
Updates the information about connections and objects that is displayed in ADAM ADSI Edit.

**Connection node**
To view the following commands, in the console tree, click a connection node for an already established connection, and then, on the **Action** menu, click:

**Remove**
Removes the currently selected connection node from ADAM ADSI Edit. This command only affects the contents of the ADAM ADSI Edit console tree; the command does not delete any objects from the ADAM instance.

**Update Schema Now**
Reloads the schema information from ADAM into the cache of the local computer.

**New**
Creates a new query when you point to **New** and then click **Query**.

**Refresh**
Updates the information about connections and objects that is displayed in ADAM ADSI Edit.

**Object node**
To view the following commands, in the console tree, click a directory object in a directory partition in an ADAM instance, and then, on the **Action** menu, click:

**Move**
Moves the object to another container in ADAM. Opens a dialog box that you can use to select the destination container.

**New Connection from Here**
Creates a new connection and adds the new connection to the console tree.

**New**
Creates a new child object in the selected container when you point to **New** and then click **Object**. Opens a set of chained dialog boxes that begins with the class of the object. If you do not have permission to create an object in the selected container, no classes are listed. After a class is selected, a dialog box opens for each required attribute. In the final dialog box, click **More** to view and edit any optional attributes.

**Delete**
Deletes the selected object from ADAM. A dialog box appears asking you to confirm the deletion. This command does not appear in the menu if you do not have permissions to delete the object from ADAM.

**Rename**
Changes the name of the object in ADAM.

**Refresh**
Updates the information that is displayed in ADAM ADSI Edit about connections and objects.

**Properties**
Opens a dialog box in which you can edit the attribute values of an object.

**Reset Password**
This option is available only for objects that contain a **unicodePwd** attribute. This option provides a dialog box in which the value of the **unicodePwd** attribute can be reset.

# ADAM Csvde

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Csvde**

Imports and exports data from Active Directory using files that store data in the comma-separated value (CSV) format. You can also support batch operations based on the CSV file format standard.

**Syntax**

**csvde** [**-i**] [**-f** *FileName*] [**-s** *ServerName*] [**-c** *String1 String2*] [**-v**] [**-j** *Path*] [**-t** *PortNumber*] [**-d** *BaseDN*] [**-r** *LDAPFilter*] [**-p** *Scope*] [**-l** *LDAPAttributeList*] [**-o** *LDAPAttributeList*] [**-g**] [**-m**] [**-n**] [**-k**] [**-a** *UserDistinguishedName Password*] [**-b** *UserName Domain Password*]

**Parameters**

**-i**
Specifies the import mode. If the import mode is not specified, the default mode is export.

**-f *FileName***
Identifies the import or export file name.

**-s *ServerName***
Specifies the domain controller to perform the import or export operation.

**-c *String1 String2***
Replaces all occurrences of *String1* with *String2*. This is generally used when you are importing data from one domain to another and the distinguished name of the export domain (*String1*) needs to be replaced with the distinguished name of the import domain (*String2*).

**-v**
Sets verbose mode.

**-j *Path***
Sets the log file location. The default is the current path.

**-t *PortNumber***
Specifies a Lightweight Directory Access Protocol (LDAP) port number. The default LDAP port is 389. The global catalog port is 3268.

**-d *BaseDN***
Sets the distinguished name of the search base for data export.

**-r *LDAPFilter***
Creates an LDAP search filter for data export.

**-p *Scope***
Sets the search scope. Search scope options are **Base**, **OneLevel**, or **SubTree**.

**-l *LDAPAttributeList***
Sets the list of attributes to appear in the results of an export query. If this parameter is omitted, all attributes are returned.

**-o *LDAPAttributeList***
Sets the list of attributes to be omitted from the results of an export query. This is typically used when exporting objects from Active Directory and then importing them into another LDAP-compliant directory. If the attributes are not supported by another directory, you can omit the attributes from the result set using this option.

**-g**
Omits paged searches.

**-m**
Omits attributes that cannot be written, such as the **ObjectGUID** and **objectSID** attributes.

**-n**
Omits export of binary values.

**-k**
Ignores errors during the import operation, and continues processing. The following is a complete list of ignored errors: "Object already exists," "Constraint violation," and "Attribute or value already exists."

**-a *UserDistinguishedName Password***
Sets the command to run using the supplied *UserDistinguishedName* and *Password*. By default, the command runs using the credentials of the user who is currently logged on to the network.

**-b *UserName Domain Password***
Sets the command to run as *UserName Domain Password*. By default, the command will run using the credentials of the user who is currently logged on to the network.

**-?**
Displays the command menu.

**Remarks**

- Applications such as Microsoft Excel are capable of reading and saving data in the CSV format. In addition, Microsoft Exchange Server administration tools are also capable of importing and exporting data using the CSV format, as are many other non-Microsoft tools.

  The CSV format consists of one or more lines of data, with each value separated by a comma. The first line (sometimes referred to as the header) of the CSV file must contain the names of each attribute in the same order as the data in any line following the first line.

  Example:

  *CN*,*FirstName*,*SurName*,*Description*

  *FirstUserLogonName*,*1stUserFirstName*,*1stUserSurName*,*Manager*

  *SecondUserLogonName*,*2ndUserFirstName*,*2ndUserSurName*,*President*

- You can use **csvde -r** to create an LDAP search filter for data export. For example, the following filter exports all users with a particular surname: **csvde -r (and(objectClass=***User***)(sn=***SurName***))**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

# ADAM Ldifde

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Ldifde**

Creates, modifies, and deletes directory objects. You can also use **ldifde** to extend the schema, export user and group information to other applications or services, and populate Active Directory Application Mode (ADAM) with data from other directory services.

**Syntax**

**ldifde** [**-i**] [**-f** *FileName*] [**-s** *ServerName*] [**-c** *String1 String2*] [**-v**] [**-j** *Path*] [**-t** *PortNumber*] [**-d** *BaseDN*] [**-r** *LDAPFilter*] [**-p** *Scope*] [**-l** *LDAPAttributeList*] [**-o** *LDAPAttributeList*] [**-g**] [**-m**] [**-n**] [**-k**] [**-a** *UserDistinguishedName Password*] [**-b** *UserName Domain Password*] [**-?**]

**Parameters**

**-i**
Specifies import mode. If not specified, the default mode is export.

**-f *FileName***
Identifies the import or export file name.

**-s *ServerName***
Specifies the computer on which to perform the import or export operation. By default, **ldifde** will run on the computer on which **ldifde** is installed.

**-c *String1 String2***
Replaces all occurrences of *String1* with *String2*. This is generally used when importing data from one domain to another and the distinguished name of the export domain (*String1*) needs to be replaced with that of the import domain (*String2*).

**-v**
Sets verbose mode.

**-j *Path***
Sets the log file location. The default is the current path.

**-t *PortNumber***
Specifies a Lightweight Directory Access Protocol (LDAP) port number. The default LDAP port is 389. The global catalog port is 3268.

**-d *BaseDN***
Sets the distinguished name of the search base for data export.

**-r *LDAPFilter***
Creates an LDAP search filter for data export. For example, to export all users with a particular surname, you can use the following filter: **-r (and(objectClass=***User***)(sn=***SurName***))**

**-p *Scope***
Sets the search scope. Search scope options are **Base**, **OneLevel**, or **SubTree**.

**-l *LDAPAttributeList***
Sets the list of attributes to return in the results of an export query. If this parameter is omitted, all attributes are returned.

**-o *LDAPAttributeList***
Sets the list of attributes to omit from the results of an export query. This is typically used when exporting objects from Active Directory and then importing them into another LDAP-compliant directory. If attributes are not supported by another directory, you can omit the attributes from the result set using this option.

**-g**
Omits paged searches.

**-m**
Omits attributes that cannot be written, such as the **ObjectGUID** and **objectSID** attributes.

**-n**

Omits export of binary values.

**-k**
Ignores errors during the import operation, and continues processing. The following is a complete list of ignored errors:

- Object is already a member of the group.
- Object class violation (meaning that the specified object class does not exist), if the object being imported has no other attributes.
- Object already exists.
- Constraint violation.
- Attribute or value already exists.
- No such object.

**-a *UserDistinguishedName Password***
Sets the command to run using the supplied *UserDistinguishedName* and *Password*. By default, the command will run using the credentials of the user who is currently logged on to the network. Cannot be used with the **-b** option.

**-b *UserName Domain Password***
Sets the command to run using the supplied *UserName Domain Password*. By default, the command will run using the credentials of the user who is currently logged on to the network. Cannot be used with the **-a** option.

**-?**
Displays the command menu.

**Remarks**

- With the **-c** parameter, you can use the constants **#schemaNamingContext** and **#configurationNamingContext** in place of the distinguished names of the schema directory partition and configuration directory partition.
- When you create the import file to use with the **ldifde** command, use a **changeType** value to define the type of changes the import file will contain. The **changeType** values in the following table are available.

| Value | Description |
|---|---|
| **add** | Specifies that new content is contained in the import file. |
| **modify** | Specifies that existing content has been modified in the import file. |
| **delete** | Specifies that content has been deleted in the import file. |

The following is an example of an LDAP Data Interchange Format (LDIF) import file format using the **add** value:

**DN: CN=***SampleUser***,DC=***DomainName*

**changetype: add**

**CN:** *SampleUser*

**description:** *DescriptionOfFile*

**objectClass:** *User*

**userPrincipalName:** *SampleUser*

**Examples**
To retrieve only the distinguished name, common name, first name, surname, and telephone number of the returned objects, type:

**-l** *DistinguishedName*, **CN,** *GivenName*, **SN,** *Telephone*

To omit the globally unique identifier (GUID) for an object, type:

**-o** *whenCreated*, *whenChanged*, *objectGUID*

**Formatting legend**

| **Format** | **Meaning** |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

# Adaminstall

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Adaminstall**
Performs an unattended installation of Active Directory Application Mode (ADAM) using a preconfigured setup script file.

**Syntax**
**adaminstall** [**/q**] [**/answer:***filename*]

**Parameters**
**/q**
Runs **adamsetup** without displaying progress information on the screen.

**/answer: *filename***
Specifies an unattended ADAM installation, using the answer file that is indicated by *filename*.

**/?**
Displays help at the command prompt.

**Examples**
The following example shows an unattended installation of ADAM using a setup script file called answer.txt:

**adaminstall /answer:** *drive:\pathname***answer.txt**

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**See Also**
**Concepts**
Perform an unattended install of an ADAM instance
Adamuninstall

# Adamsync

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Adamsync**
Synchronizes objects from Active Directory to an Active Directory Application Mode (ADAM) instance.

**Syntax**
**adamsync** [**/?**] [**/l**] [**/d** *configuration_dn*] [**/i** *input_file*] [**/download** *configuration_dn output_file*] [**/export** *configuration_dn output_file*] [**/sync** *configuration_dn*] [**/reset** *configuration_dn*] [**/mai** *configuration_dn*] [**/fs** *configuration_dn*] [**/ageall** *configuration_dn*] [**/so** *configuration_dn object_dn*] [**/passPrompt**]

**Parameters**

**/?**
Displays command-line options.

**/l or /list**
Displays all available configurations of Active Directory to ADAM Synchronizer.

**/d or /delete *configuration_dn***
Deletes the specified configuration.

**/i or /install *input_file***
Installs the configuration contained in the specified input file.

**/download *configuration_dn output_file***
Creates an output file in XML containing the specified configuration.

**/export *configuration_dn output_file***
Saves the current configuration to the specified output file.

**/sync *configuration_dn***
Synchronizes the specified configuration.

**/reset *configuration_dn***
Resets the replication cookie of the specified configuration.

**/mai *configuration_dn***
Configures the specified configuration as the authoritative instance.

**/fs *configuration_dn***
Performs a full replication synchronization for the specified configuration.

**/ageall *configuration_dn***
Performs an aging run for the specified configuration.

**/so *configuration_dn object_dn***
Performs a replication synchronization for the specified object in the specified configuration. Use the distinguished name of the object.

**/passPrompt**
Prompts for user credentials.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |

| | |
|---|---|
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**See Also**
**Concepts**
Perform an unattended install of an ADAM instance
Adaminstall
Remove an ADAM instance

# Adamsync Configuration File XML Reference

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The following topics provide reference information for XML elements that are used in the configuration file for Active Directory to ADAM Synchronizer.


- account-domain
- aging
- attributes
- base-dn
- configuration
- description
- exclude
- frequency
- include
- num-objects
- object-filter
- query
- schedule
- schtasks-cmd
- security-mode
- source-ad-account
- source-ad-name
- source-ad-partition
- source-object-class
- synchronizer-state
- target-dn
- target-object-class
- user-proxy

# account-domain

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the domain name to use for authentication on the source forest.

**Usage**

```
<account-domain/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<configuration>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

**Remarks**
The <source-ad-account> element and the <account-domain> element combine to create the account domain and name to use for alternate credentials. For example, if the account to use is "fabrikam\jsmith", <account-domain> would contain "fabrikam" and <source-ad-account> would contain "jsmith".

# aging

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains elements that describe how often an aging search is performed and how many objects to search for in the aging run.

**Usage**

```
<aging/>
```

**Attributes**
None.

**Child Elements**
<frequency>

<num-objects>

**Parent Elements**
<schedule>

**Child Element Sequence**
The tree shows the ordering and number of child elements. Child elements in a Sequence must appear in the order given and child elements in a Choice are mutually exclusive. Sequences and Choices can be nested. Element names are followed by the minimum and maximum occurrences.

```
frequency (Minimum = 0; Maximum = 1)
num-objects
```

**Element Information**

| | |
|---|---|
| Can be empty | No |

**Remarks**
An aging search is performed to determine if the objects in ADAM have been deleted from Active Directory. This is accomplished by searching for the ADAM objects in Active Directory because this is the only consistent way to determine if an object has been deleted in Active Directory.

# attributes

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains elements to specify the attributes to include or exclude in a search.

**Usage**

```
<attributes/>
```

**Attributes**
None.

**Child Elements**
<include>

<exclude>

**Parent Elements**
<query>

**Child Element Sequence**
The tree shows the ordering and number of child elements. Child elements in a Sequence must appear in the order given and child elements in a Choice are mutually exclusive. Sequences and Choices can be nested. Element names are followed by the minimum and maximum occurrences.

```
Sequence
    include (Minimum = 0; Maximum = many)
    exclude (Minimum = 0; Maximum = many)
Sequence
```

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

**Remarks**
If this element does not contain any child elements, then all attributes are included in the search.

Attributes can either be included or excluded, but not both. If one or more attributes are included, only those object attributes are included in the search. If one or more attributes are excluded, all attributes are included in the search except those that are specifically excluded.

# base-dn

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the distinguished name of the base object where the search will begin. The search is performed only on this object and objects that exist below it in the directory tree.

**Usage**

```
<base-dn/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<query>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# configuration

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The <configuration> element is the root element that contains the configuration data that is required by Active Directory to ADAM Synchronizer.

**Usage**

```
<configuration/>
```

**Attributes**
None.

**Child Elements**
<description>

<security-mode>

<source-ad-name>

<source-ad-partition>

<source-ad-account>

<account-domain>

<target-dn>

<user-proxy>

<query>

<schedule>

<synchronizer-state>

**Parent Elements**
None.

**Child Element Sequence**
The tree shows the ordering and number of child elements. Child elements in a Sequence must appear in the order given and child elements in a Choice are mutually exclusive. Sequences and Choices can be nested. Element names are followed by the minimum and maximum occurrences.

```
description (Minimum = 0; Maximum = 1)
security-mode (Exactly 1)
source-ad-name (Exactly 1)
source-ad-partition (Exactly 1)
source-ad-account (Exactly 1)
account-domain (Minimum = 0; Maximum = 1)
target-dn (Minimum = 0; Maximum = 1)
user-proxy (Minimum = 0; Maximum = 1)
query (Exactly 1)
schedule (Minimum = 0; Maximum = 1)
synchronizer-state
```

**Element Information**

| | |
|---|---|
| Can be empty | No |

# description

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains a string that describes the configuration file. This string can be any length.

**Usage**

```
<description>
```

**Attributes**

None.

**Child Elements**

None.

**Parent Elements**

<configuration>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# exclude

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the Lightweight Directory Access Protocol (LDAP) display name of an attribute to exclude in the query during synchronization.

**Usage**

```
<exclude/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<attributes>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# frequency

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the number of synchronization runs that will occur before an aging search will be performed. This value must be a positive integer.

**Usage**

```
<frequency/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<aging>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

**Remarks**
Setting this value to "1" will verify the objects on every synchronization run. Setting this value to "2" will verify the objects on every other synchronization run.

# include

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the LDAP display name of an attribute to include in the query during synchronization.

**Usage**

```
<include/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<attributes>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# num-objects

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the number of objects to search for in each aging search.

**Usage**

```
<num-objects/>
```

**Attributes**

None.

**Child Elements**

None.

**Parent Elements**

<aging>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# object-filter

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the Lightweight Directory Access Protocol (LDAP) query string.

**Usage**

```
<object-filter/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<query>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

**Remarks**
Characters that are reserved for use by XML, such as "<", ">" and "&" must be properly formatted to prevent the XML from being invalidated. For example, the "&" character is replaced by the "&" string. This is known is a character reference.

# query

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains child elements that describe the query that is used to search Active Directory for objects to synchronize.

**Usage**

```
<query/>
```

**Attributes**
None.

**Child Elements**
<base-dn>

<object-filter>

<attributes>

**Parent Elements**
None.

**Child Element Sequence**
The tree shows the ordering and number of child elements. Child elements in a Sequence must appear in the order given and child elements in a Choice are mutually exclusive. Sequences and Choices can be nested. Element names are followed by the minimum and maximum occurrences.

```
base-dn (Minimum = 1; Maximum = many)
object-filter
attributes (Minimum = 0; Maximum = 1)
```

**Element Information**

| | |
|---|---|
| Can be empty | No |

**Remarks**
All filtering is performed on the client side of the operation, not the server side.

# schedule

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains elements that describe how often to verify deleted objects and changed attribute values in Active Directory.

**Usage**

```
<schedule/>
```

**Attributes**
None.

**Child Elements**
<aging>

<schtasks-cmd>

**Parent Elements**
None.

**Child Element Sequence**
The tree shows the ordering and number of child elements. Child elements in a Sequence must appear in the order given and child elements in a Choice are mutually exclusive. Sequences and Choices can be nested. Element names are followed by the minimum and maximum occurrences.

```
aging (Minimum = 0; Maximum = 1)
schtasks-cmd (Minimum = 0; Maximum = 1)
```

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# schtasks-cmd

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Reserved for future use.

**Usage**

```
<schtasks-cmd/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
None.

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# security-mode

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Defines the security mode of the synchronization.

**Usage**

```
<security-mode/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<configuration>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

**Remarks**
This element can contain either "partition" or "object". Any other value is invalid.

If this element contains "partition", the user account must have Replicate Directory Changes permissions on the source partition to synchronize objects.

If this element contains "object", the user account requires only readaccess to the objects to be synchronized.

# source-ad-account

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the account name to use for authentication on the source forest. If this element is empty, the synchronizer will use the credentials of the account that is currently logged on.

**Usage**

```
<source-ad-account/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<configuration>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

**Remarks**
The <source-ad-account> element and the <acocunt-domain> element combine to create the account domain and name to use for alternate credentials. For example, if the account to use is "fabrikam\jsmith", <account-domain> would contain "fabrikam" and <source-ad-account> would contain "jsmith".

Because passwords should never be stored as plain text in an XML file, the ADAM synchronizer will request the password from the user when the configuration is installed. This password will be encrypted and stored by the synchronizer for use when the synchronization is performed.

# source-ad-name

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the name of the Active Directory forest that is the source of the synchronization.

**Usage**

```
<source-ad-name/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<configuration>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

**Remarks**
If this element contains a fully-qualified distinguished name, it will be resolved using the domain controller locator. If an LDAP bindable name is specified, such as a DNS name, this object will be bound to.

# source-ad-partition

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the distinguished name of the Active Directory partition that is the source of the synchronization.

**Usage**

```
<source-ad-partition>
```

**Attributes**

None.

**Child Elements**

None.

**Parent Elements**

<configuration>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# source-object-class

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the Lightweight Directory Access Protocol (LDAP) display name of the schema class on the source that should be converted when objects of this type are copied to Active Directory Application Mode (ADAM).

**Usage**

```
<source-object-class/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<user-proxy>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

**Remarks**
These elements should be user principals only.

Object classes that are derived from this class are not converted. You must specify a separate <source-object-class> element for each object class to be converted.

# synchronizer-state

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains information that is reserved for use by the synchronizer. You should not modify any information in this element.

**Usage**

```
<synchronizer-state/>
```

**Attributes**

None.

**Child Elements**

None.

**Parent Elements**

<configuration>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# target-dn

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the distinguished name of the Active Directory Application Mode (ADAM) partition that is the destination of the synchronization.

**Usage**

```
<target-dn/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<configuration>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

# target-object-class

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the Lightweight Directory Access Protocol (LDAP) display name of the schema class to which the objects will be converted.

**Usage**

```
<target-object-class/>
```

**Attributes**
None.

**Child Elements**
None.

**Parent Elements**
<user-proxy>

**Element Information**

| | |
|---|---|
| Can be empty | Yes |

**Remarks**
All objects of the classes that are specified by the preceding <source-object-class> elements will be converted to this type during the conversion.

This class should have all the necessary attributes for each <source-object-class> in its schema definition.

# user-proxy

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Contains the list of source object classes that should be converted during the synchronization.

**Usage**

```
<user-proxy/>
```

**Attributes**
None.

**Child Elements**
<source-object-class>

<target-object-class>

**Parent Elements**
<configuration.

**Child Element Sequence**
The tree shows the ordering and number of child elements. Child elements in a Sequence must appear in the order given and child elements in a Choice are mutually exclusive. Sequences and Choices can be nested. Element names are followed by the minimum and maximum occurrences.

```
Sequence
    source-object-class (Minimum = 1; Maximum = many)
    target-object-class (Exactly 1)
Sequence
```

**Element Information**

| | |
|---|---|
| Can be empty | No |

**Remarks**
This element contains one or more <source-object-class> elements, followed by a single <target-object-class> element. All objects of the class specified by the <source-object-class> elements will be converted into the type specified by the <target-object-class> element.

Object classes that are derived from the classes specified by the <source-object-class> elements are not converted. You must specify a separate <source-object-class> element for each object class to be converted.

The <source-object-class> elements should contain only user principles.

If no conversion is needed, it is recommended that you include a single, empty <source-object-class> element in this element.

# Adamuninstall

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Adamuninstall**
Uninstalls an Active Directory Application Mode (ADAM) instance.

**Syntax**
**adamuninstall** [**/q**] [**/answer:***filename*] [**/force**] {**/i:***instancename* | **/tools**}

**Parameters**

**/q**
Runs **adamuninstall** without displaying progress information on the screen.

**/answer: *filename***
Specifies the name of the answer file, represented by *filename*, that contains preconfigured values for ADAM uninstall.

This answer file can be used to specify the account and password of the ADAM administrator, using the following syntax:

**Administrator=***domain\account*

**AdministratorPassword=***password*

where *domain\account* represents the account of the ADAM administrator and *password* represents the corresponding password.

**/force**
Causes **adamuninstall** to ignore any warnings and continue the removal of a given ADAM instance. **Adamuninstall** automatically skips any operations that generate warnings.

**/i: *instancename***
Uninstalls an ADAM instance named *instancename*, and also uninstalls the ADAM administration tools.

**/tools**
Uninstalls the ADAM administration tools only.

**Remarks**

- If service principal names (SPN) in Active Directory exist for the ADAM instance that is being uninstalled, and **adamuninstall** cannot remove the SPNs automatically, **adamuninstall** creates a .bat file that can be used to remove the SPNs manually. This .bat file is located at %windir%\Debug*domain-servicename*.bat, where *domain* represents the Active Directory domain in which the ADAM instance resides, and *servicename* represents the service name of the ADAM instance. (By default, the service name of an ADAM instance is ADAM_*instancename*, where *instancename* is the ADAM instance name that you provide during ADAM setup.)

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

**See Also**

**Concepts**

Perform an unattended install of an ADAM instance

Adaminstall

Remove an ADAM instance

# ADSchemaAnalyzer

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ADSchemaAnalyzer**

You can use ADSchemaAnalyzer to help migrate the Active Directory schema to ADAM. To start ADSchemaAnalyzer, click **Start**, point to **All Programs**, point to **ADAM**, click **ADAM Tools Command Prompt**, and then at the command prompt type: **adschemaanalyzer**.

The following sections describe commands on the ADSchemaAnalyzer menus:

- File
- Schema
- Tools

**File**

The following table describes commands on the **File** menu.

| Command | Description |
|---|---|
| **Load target schema** | Opens a dialog box that you can use to authenticate to a specified Lightweight Directory Access Protocol (LDAP) server. Type a user name and password of an account that has permissions to the LDAP server. If you do a simple bind with an empty password, you are connected with anonymous credentials. Select the server type to load the target schema from. To select an LDAP Data Interchange Format (LDIF) file to load, click **Load LDIF**. |
| **Load base schema** | Opens a dialog box that you can use to authenticate to a specified LDAP server. Type a user name and password of an account that has permissions to the LDAP server. If you do a simple bind with an empty password, you are connected with anonymous credentials. Select the server type to load the base schema from. To select an LDIF file to load, click **Load LDIF**. |
| **Create LDIF file** | Produces an LDIF file that contains all selected schema elements. |

**Schema**

The following table describes commands on the **Schema** menu.

| Command | Description |
|---|---|
| **Hide present elements** | Hides all elements that are present on the target server. |
| **Show present elements** | Shows all elements that are present on the target server. |
| **Mark all non-present elements as included** | Marks all elements that are not present on the target server as included for export. |

**Tools**

The following table describes commands on the **Tools** menu.

**Options**

Opens a dialog box that you can use to set options for ADSchemaAnalyzer. The following table describes options on the tabs in the **Options** dialog box.

| Tab | Description |
|---|---|

| | |
|---|---|
| **Dependencies graph** | The following are the options on the **Dependencies graph** tab.<br>• **AttributeSecurityGUID (attribute to propset reference)** - Select this option to display attributes that belong to a propset.<br>• **AppliesTo (propset to class reference)** - Select this option to display classes that reference a given attribute or class.<br>• **Show Dependents node** - Select this option to show the Dependents node in the tree for each element. |
| **LDIF generation** | The following are the options on the **LDIF generation** tab.<br>• **Don't update present elements** - Use this option if you do not want to update present elements in the LDIF file.<br>• **Update with references to new elements only** - Use this option to update the LDIF file with references to new elements.<br>• **Update with references to new and present elements** - Use this option to update the LDIF file with references to new and present elements.<br>• **Allowed to write systemPossSureriors and systemAuxClasses as possSuperiors and auxClasses, if schema circular dependencies exist** - Use this option if there are schema circular dependencies and you want to import the schema by writing **systemPossSuperiors** as **possSuperiors** and **systemAuxClasses** as **auxClasses**. |
| **Miscellaneous** | The following are the options on the **Miscellaneous** tab.<br>• **Marking an element as present marks all its subelements as present** - Select this option to mark all subelements as present automatically when you mark an element as present in the tree.<br>• **Debug output** - Select this option to debug the output of the LDIF file. |

Right-click schema elements in the tree to mark them as **Auto**, **Included**, **Excluded**, or **Present**:

- **Auto** automatically marks an element as included or excluded in the export. If an element is marked as **Auto (included)**, you can right-click that element and then click **Why auto included?** to see the reverse dependency tree for the element.

- **Included** marks the element so that it is included in the export. ADSchemaAnalyzer marks all related elements, such as **superclasses**, aux classes, must/may contains, **defaultObjectCategory**, and **possSuperiors**. ADSchemaAnalyzer includes propsets for included attributes and back-links for links.

- **Excluded** marks the element so that it is not included in the export. You can block certain paths in the dependency graph. For example, you might want to import **domainDns**, but not **samAccountDomain** (which is an **auxClass** of **domainDns**). You can exclude a complete element, such as the **samAccountDomain** class, or you can exclude a relationship; for example, you can remove the **auxClass** reference from the **domainDns** class. If you exclude a relationship, any other classes that reference that element continue to include it.

- **Present** means that the element is present on the target server. By default, the top class is marked as present.

# Dsacls

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Dsacls**
Displays and changes permissions (access control entries (ACEs)) in the access control list (ACL) of objects in Active Directory Application Mode (ADAM).

The ACEs that you add by using **dsacls** must be object-specific permissions that override the default partition permissions that are defined in the ADAM schema. Do not add ACEs unless you are well informed about security for ADAM objects.

If you specify an object without additional parameters, **dsacls** displays the ACEs in the ACL.

**Syntax**
**dsacls "**[*****Computer*****]*ObjectDN***"** [**/A**] [**/D** *PermissionStatement* [*PermissionStatement*]...] [**/G** *PermissionStatement* [*PermissionStatement*]...] [**/I:**{**T** | **S** | **P**}] [**/N**] [**/P:**{**Y** | **N**}] [**/R** {*User* | *Group*} [{*User* | *Group*}]...] [**/S** [**/T**]] [**/resetDefaultACL**] [**/resetDefaultSACL**] [**/takeOwnership**] [**/simple**] [**/domain:** *domain*] [**/user:** *username*] [**/passwd:** {*password* | ***** }] [**/?**]

**Parameters**

**"[*Computer*\]*ObjectDN*"**
Identifies the Active Directory object to investigate. Type the distinguished name of the object. To specify an object on a remote computer, type the computer name followed by the distinguished name. This parameter must be enclosed in quotation marks.

Example:

**"CN=Kim Akers,CN=Users,DC=domain,DC=test,DC=microsoft,DC=com"**

or

**"\\Server01\CN=Kim Akers,CN=Users,DC=domain,DC=test,DC=microsoft,DC=com"**

**/A**
Adds ownership and auditing information to the display.

**/D PermissionStatement [*PermissionStatement*]...**
Denies the specified permissions to the user or group.

You can deny permissions to multiple users in each **/D** command.

Example:

**/D Domain1\User1:CCDC Domain1\User2:DC;computer**

**/G PermissionStatement [*PermissionStatement*]...**
Grants specified permissions to the user or group.

You can grant permissions to multiple users in each **/G** command

Example:

**/G Domain1\User1:CCDC Domain1\User2:DC;computer**

**/I:{T | S | P}**
Specifies the objects to which the permissions are applied. This parameter determines whether the permissions are inheritable. **T** is the default.

| Value | Description |
|---|---|
| T | This object and subobjects |
| S | Subobjects only |
| P | Propagate inheritable permissions one level only |

**/N**
Provides that the specified ACE replaces the ACEs in the ACL. By default, the ACE is added to the ACL.

**/P:{Y | N}**
Determines whether the object is protected so that it cannot inherit permissions from its parent objects. If you omit this parameter, the inheritance properties of the object are not changed.

| Value | Description |
|---|---|
| **Y** | The object is protected, and it cannot inherit permissions. |
| **N** | The object is not protected, and it can inherit permissions. |

**Note**

- This parameter changes a property of the object, not of an ACE. To determine whether an ACE is inheritable, use the **/I** parameter.

**/R {*User* | *Group*} [{*User* | *Group*}]...**
Deletes all ACEs for the specified users or groups. *User* can be specified as *User@Domain* or as *Domain\User*. *Group* can be specified as *Group@Domain* or as *Domain\Group*.

You can delete ACEs for multiple users and groups in a single **/R** parameter

Example:

**/R Domain1\User1 Domain1\User2**

**/S**
Restores the security on the object to the default for that object class, as defined in the Active Directory schema.

**/T**
Restores the security on the tree of objects to the default for each object class. This parameter is valid only with the **/S** parameter.

**/resetDefaultDACL**
Restores the DACL on the object to the default setting for the object, based on the schema definition for the object class.

**/resetDefaultSACL**
Restores the SACL on the object to the default setting, based on the schema definition for the object class.

**/takeOwnership**
Assigns object owernship to the account under which Dsacls is running.

**/simple**
Specifies to bind to the directory server using an LDAP simple bind.

**/domain: *domain***
Specifies to bind to the directory server using an account from the domain represented by *domain*.

**/user: *username***
Specifies to bind to the directory server using the account represented by *username*.

**/passwd: *password***
Specifies to bind to the directory server using the password represented by *password*. If no password is specified, the user is prompted for a password.

**/?**
Displays help for **dsacls**.

**Syntax for PermissionStatement**
*PermissionStatement* values use the following format:

{*User* | *Group*}**:***Permissions*[**;**{*ObjectType* | *Property*}][**;***InheritedObjectType*]

**Parameters**
**{*User* | *Group*}**

Specifies the user or group to whom the rights apply. *User* can be specified as a distinguished name, *User@Domain*, or *Domain**User*. *Group* can be specified as a distinguished name, *Group@Domain*, or *Domain**Group*.

***Permissions***

Type one or more of the values in the following tables (without spaces).

**Generic Permissions**

| Value | Description |
|---|---|
| **GR** | Generic Read |
| **GE** | Generic Execute |
| **GW** | Generic Write |
| **GA** | Generic All |

**Specific Permissions**

| Value | Description |
|---|---|
| **SD** | Delete. |
| **DT** | Delete an object and all of its children. |
| **RC** | Read security information. |
| **WD** | Change security information. |
| **WO** | Change owner information. |
| **LC** | List the children of an object. |
| **CC** | Create a child object.<br><br>If {*ObjectType* \| *Property*} is not specified to define a specific child-object type, this applies to all types of child objects; otherwise, it applies to the specified child-object type. |
| **DC** | Delete a child object.<br><br>If {*ObjectType* \| *Property*} is not specified to define a specific child-object type, this applies to all types of child objects; otherwise, it applies to the specified child-object type. |
| **WS** | Write to self object.<br><br>This is meaningful only on Group objects and when {*ObjectType* \| *Property*} is a "member." |
| **RP** | Read property.<br><br>If {*ObjectType* \| *Property*} is not specified to define a specific property, this applies to all properties of the object; otherwise, it applies to the specified property of the object. |
| **WP** | Write property.<br><br>If {*ObjectType* \| *Property*} is not specified to define a specific property, this applies to all properties of the object; otherwise, it applies to the specified property of the object. |
| **CA** | Control access right.<br><br>If {*ObjectType* \| *Property*} is not specified to define the specific extended right for control access, this applies to all meaningful control accesses on the object; otherwise, it applies to the specific extended right for that object. |

| | |
|---|---|
| **LO** | List the object access.<br><br>Can be used to grant list access to a specific object if List Children (**LC**) is not granted to the parent as well. Can also be denied on specific objects to hide those objects if the user/group has **LC** permission on the parent.<br><br>**Note**<br><br>• Active Directory does not enforce this permission by default. Active Directory must be configured to check for this permission. |

**{*ObjectType* | *Property*}**

Limits the permission to the specified object type or property. Type the display name of the object type or the property. If an object type or property is not specified, the permission applies to all object types and properties.

For example, the following command permits the user to create all types of child objects:

**/G Domain\User:CC**

However, the following command permits the user to create only child computer objects:

**/G Domain\User:CC;computer**

***InheritedObjectType***

Limits inheritance of the permission to the specified type of object. Type the display name of the object type. If an object type is not specified, the permission can be inherited by all object types. This parameter is used only when permissions are inheritable.

For example, the following command permits all types of objects to inherit the permission:

**/G Domain\User:CC**

However, the following command permits only user objects to inherit the permission:

**/G Domain\User:CC;;user**

**Examples**

**SDRCWDWO;;user**

Delete, read security information. Change security information and change ownership permissions on objects of the type "user."

**CCDC;group;**

Create child and delete child permissions to create or delete objects of the type "group."

**RPWP;telephonenumber;**

Read property and write property permissions on the telephone number property.

**Formatting legend**

| Format | Meaning |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

# Dsdbutil

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Dsdbutil**
Provides management facilities for Active Directory Application Mode (ADAM). This tool is intended for use by experienced administrators.

| **Note** |
| --- |
| Some **dsdbutil** commands can be used only against ADAM instances that are not currently running. |

The syntax for each command is listed below.

- activate instance
- authoritative restore
- change service account
- files
- LDAP port
- list instances
- popups
- quit
- semantic database analysis
- set DSRM Password
- SSL port

**activate instance**
Specifies the ADAM instance to be managed. At the **dsdbutil:** prompt, type the command using the syntax below. You cannot activate an ADAM instance if that ADAM instance is currently running.

**Syntax**
**activate instance** *instancename*

***Instancename***
Specifies the name of an ADAM instance to be managed.

**authoritative restore**
Restores an ADAM instance to a specific point in time by marking objects in ADAM as being authoritative with respect to their replication partners. At the **authoritative restore:** prompt, type any of the parameters listed under "Syntax."

**Syntax**
{**restore database**|**restore database verinc** *%d*|**restore subtree** *%s*|**restore subtree** *%s* **verinc** *%d*|**restore object** *%s*|**restore object** *%s* **verinc** *%d*}

**Parameters**
**restore database**
Marks all of Adamntds.dit (all directory partitions that are held by the ADAM instance) as authoritative. The schema cannot be authoritatively restored.

**restore database verinc *%d***
Marks all of Adamntds.dit (all directory partitions that are held by the ADAM instance) as authoritative, and increments the version number as follows. If **verinc** is not provided, the version number is increased by the age (in days) of the object being restored, multiplied by 100,000. If **verinc** is provided, the version number is increased by the age (in days) of the object being restored, multiplied by *%d*. Use this option only to authoritatively restore over a previous, incorrect, authoritative restore, such as an authoritative restore completed using a backup that contains incorrect data.

***%d***
A numeric variable, such as the replication delay time period.

**restore subtree *%s***
Marks the subtree and all children of the subtree as being authoritative. The subtree is defined by using the full distinguished name of the object.

**restore subtree *%s* verinc *%d***
Marks the subtree and all children of the subtree as being authoritative, and increments the version number as follows. If **verinc** is not provided, the version number is increased by the age (in days) of the object being restored, multiplied by 100,000. If **verinc** is provided, the version number is increased by the age (in days) of the object being restored, multiplied by *%d*. The subtree is defined by using the full distinguished name of the object. Use this option only to authoritatively restore over a previous, incorrect, authoritative restore, such as an authoritative restore completed using a backup that contains incorrect data.

**restore object *%s***
Marks the object as being authoritative. The object is specified by using the full distinguished name of the object.

**restore object *%s* verinc *%d***
Marks the object as being authoritative, and increments the version number as follows. If **verinc** is not provided, the version number is increased by the age (in days) of the object being restored, multiplied by 100,000. If **verinc** is provided, the version number is increased by the age (in days) of the object being restored, multiplied by *%d*. The object is specified by using the full distinguished name of the object. Use this option only to authoritatively restore over a previous, incorrect, authoritative restore, such as an authoritative restore completed using a backup that contains incorrect data.

***%s***
An alphanumeric variable, such as the name of an ADAM instance.

**Quit**
Takes you back to the previous menu, or exits the utility.

**? or help**
Displays help at the command prompt.

#### change service account
Sets the service account for the ADAM instance. At the **dsdbutil:** prompt, type the command using the syntax listed below.

##### Syntax
**change service account** *account password*

***account***
The account to be used as the new ADAM service account.

***password***
The password for the account to be used as the new ADAM service account.

#### files
Provides commands for managing directory service data and log files. The data file is called Adamntds.dit. At the **files:** prompt, type any of the parameters listed under "Syntax."

##### Syntax
{**compact to** *%s*|**header**|
**info**|**integrity**|**move DB to** *%s*|**move logs to** *%s*|**recover**|**set path backup** *%s*|**set path db** *%s*|**set path logs** *%s*|**set path working dir** *%s*

##### Parameters
**compact to *%s* (where *%s* identifies an *empty target directory*)**
Compacts the existing data file, and writes the compacted file to the specified directory. The directory can be remote; that is, it can be mapped by means of the **net use** command or by similar means. After compaction is complete, archive the old data file, and move the newly compacted file back to the original location of the data file.

**header**
Writes the header of the Adamntds.dit data file to the screen. This command can help you analyze database problems.

**info**
Analyzes and reports the free space for the disks that are installed in the system, reads the registry, and then reports the sizes of the data and log files. The directory service maintains the registry, which identifies the location of the data files, log files, and directory service working directory.

**integrity**
Performs an integrity check on the data file, which can detect any kind of low-level database corruption. It reads every byte of the data file; therefore, it can take a long time to process large databases. Note that you should always run **recover** before performing an integrity check.

**move DB to *%s* (where *%s* identifies a target directory)**
Moves the Adamntds.dit data file to the new directory that is specified by *%s*, and updates the registry so that, upon service restart, the directory service uses the new location.

**move logs to *%s* (where *%s* identifies a target directory)**
Moves the directory service log files to the new directory that is specified by *%s*, and updates the registry so that, upon service restart, the directory service uses the new location.

**recover**
Performs a soft recovery of the database. Soft recovery scans the log files and ensures that all committed transactions in the log files are also reflected in the data file. The **ntbackup** program truncates the log files appropriately. Logs are used to ensure that committed transactions are not lost if the system fails or if there is an unexpected power loss. Essentially, transaction data is written first to a log file and then to the data file. When you restart after failure, you can rerun the log to reproduce the transactions that are committed but that have not been written to the data file.

**set path backup *%s* (where *%s* identifies a target directory)**
Sets the disk-to-disk backup target to the directory that is specified by *%s*. The directory service can be configured to perform an online, disk-to-disk

backup at scheduled intervals.

**set path db *%s* (where *%s* identifies a target directory)**
Updates the part of the registry that identifies the location and file name of the data file. Use this command only to rebuild an ADAM instance that has lost its data file and that is not being restored by means of normal restoration procedures.

**set path logs *%s* (where *%s* identifies a target directory)**
Updates the part of the registry that identifies the location of the log files. Use this command only to rebuild an ADAM instance that has lost its log files and that is not being restored by means of normal restoration procedures.

**set path working dir *%s* (where *%s* identifies a target directory)**
Sets the part of the registry that identifies the directory service's working directory to the directory that is specified by *%s*.

***%s***
An alphanumeric variable, such as a domain or ADAM instance name.

**quit**
Takes you back to the previous menu, or exits the utility.

**? or help**
Displays help at the command prompt.

| **Caution** |
| --- |
| Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. |

| **Note** |
| --- |
| Active Directory Application Mode is implemented on top of an indexed sequential access method (ISAM) table manager. This is the same table manager that is used by Microsoft Exchange Server, the file replication service, the security configuration editor, the certificate server, Windows Internet Name Service (WINS), and other Windows components. The version of the database that Windows Server 2003 uses is called extensible storage engine (ESENT). ESENT is a transacted database system that uses log files to support rollback semantics to ensure that transactions are committed to the database. Ideally, data and log files should be located on separate drives to improve performance and support recovery of the data if a disk fails. |

**LDAP port**
Sets the Lightweight Directory Access Protocol (LDAP) communications port for the ADAM instance. At the **dsdbutil:** prompt, type the command using the syntax listed below.

**Syntax**
**LDAP port** *portnumber*

***portnumber***
Specifies the new communications port number to be used for LDAP on the ADAM instance.

**list instances**
Lists information about the ADAM instances installed on the current computer. To run this command, type **list instances** at the **dsdbutil:** prompt. This command does not require or accept any parameters.

**popups**
Enables or disables interactive pop-up messages. To enable interactive pop-up messages, at the **dsdbutil:** prompt, type **popups on**. To disable pop-up messages, at the **dsdbutil:** prompt, type **popups off**.

**quit**
Closes the **dsdbutil** command-line tool.

**semantic database analysis**
Analyzes data with respect to ADAM semantics. At the **semantic database analysis:** prompt, type any of the parameters listed under "Syntax."

**Syntax**
{**get** *%d*|**go**|**verbose** *%s*}

**Parameters**
**get *%d***
Retrieves the record number *%d* from Ntds.dit.

**go**
Starts the semantic analysis of Ntds.dit. A report is generated and written to a file named Dsdit.dmp.*n* in the current directory, where *n* is an integer that is incremented each time that you carry out the command.

**verbose *%s***
Toggles verbose mode on or off.

***%d***
A numeric variable, such as a replication delay time period.

***%s***

An alphanumeric variable, such as a domain or ADAM instance name.

**quit**
Takes you back to the previous menu, or exits the utility.

**? or help**
Displays help at the command prompt.

**Remarks**

- Unlike the file management commands described earlier, which test the integrity of the database with respect to ESENT database semantics, the semantic analysis analyzes data with respect to ADAM semantics. It generates reports on the number of records present, including deleted and phantom records.

| **Note** |
| --- |
| End users should not use this command except, when Microsoft requests them to use it as an aid in fault diagnosis. |

**set DSRM Password**
This command does not apply to ADAM.

**SSL port**
Sets the Secure Sockets Layer (SSL) communications port for the ADAM instance. At the **dsdbutil:** prompt, type the command using the syntax listed below.

**Syntax**
**SSL port** *portnumber*

***portnumber***
Specifies the new communications port number to be used for SSL on the ADAM instance.

**Remarks**

- If the variable has spaces in it, enclose it in parentheses, instead of in quotation marks, as follows:

  **connect to server (***xxx yyy***)**

**Formatting legend**

| Format | Meaning |
| --- | --- |
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

# Dsdiag

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Dsdiag**
Analyzes the state of the Active Directory directory service, or the Active Directory Application Mode (ADAM) directory service, and reports any problems, to assist in troubleshooting. **Dsdiag** provides detailed information about how to identify abnormal behavior in the system.

**Dsdiag** consists of a framework for running tests and a series of tests to verify different functional areas of the system. This framework selects which directory services are tested, according to scope directives from the user.

| **Note** |
| --- |
| Some **Dsdiag** tests apply only to Active Directory. |


- Syntax
- Examples
- Formatting legend


**Syntax**
**dsdiag /s:***adamserver* [**/n:***NamingContext*] [**/u:***Domain\UserName* **/p:{*** | *Password* | **""}**] [{**/a** | **/e**}] [{**/q** | **/v**}] [**/i**] [**/f:***LogFile*] [**/ferr:***ErrLog*] [**/c** [**/skip:***Test*]] [**/test:***Test*][**/d**][**/fix**] [{**/h** | **/?**}]

**Parameters**

**/s: *adamserver***
Uses *adamserver* as the home server. This parameter is required.

**/n: *Partition***
Uses *Partition* as the application directory partition to test. Domains may be specified in network basic input/output system (NetBIOS), Domain Name System (DNS), or distinguished name formats.

**/u: *Domain\UserName* /p:{* | *Password* | ""}**
Uses *Domain\UserName* credentials for binding, with *Password* as the password. Use **""** for an empty or null password, or use ***** to prompt for the password.

**/a**
Tests all the servers on this site.

**/e**
Tests all the servers in the entire enterprise. Overrides **/a**.

**/q**
(Quiet) Prints only error messages.

**/v**
(Verbose) Prints extended information.

**/d**
(Debug) Prints configuration information for the entire enterprise, and displays verbose output information. This parameter can be useful for discovering detailed information about a directory server.

**/i**
Ignores superfluous error messages.

**/fix**
Only affects the *MachineAccount* test. This parameter causes the test to repair the service principal names (SPNs) on the machine account object of the directory server.

**/f: *LogFile***
Redirects all output to *LogFile*. The **/f** parameter operates independently of **/ferr**.

**/ferr: *ErrLog***
Redirects fatal error output to a separate file called *ErrLog*. The **/ferr** parameter operates independently of **/f**.

**/c**
(Comprehensive) Runs all tests except **DcPromo** and **RegisterInDNS**, including nondefault tests. Optionally, this parameter can be used with **/skip** to skip specified tests. The following tests are not run by default:
**TopologyCutoffServersOutboundSecureChannels**

**/skip: *Test***
Skips the specified test. This parameter must be used with **/c**. This parameter should not be run in the same command with **/test**. The only test that cannot be skipped is **Connectivity**.

**/test: *Test***
Runs only this test. The nonskippable test **Connectivity** is also run. This parameter should not be run in the same command with **/skip**.

| **Note** |
| --- |
| All tests except **DcPromo** and **RegisterInDNS** must be run on computers that have been promoted to directory server. |

Valid tests are as follows:

**Connectivity**
Tests whether directory servers are DNS registered, can be pinged, and have Lightweight Directory Access Protocol/remote procedure call (LDAP/RPC) connectivity.

**Replications**
Checks for timely replication and any replication errors between directory servers.

**Topology**
Checks that Knowledge Consistency Checker (KCC) has generated a fully connected topology for all directory servers.

**CutoffServers**
Checks for any servers that are not receiving replications because their partners are not available.

**NCSecDesc**
Checks that the security descriptors on the application directory partition heads have appropriate permissions for replication.

**NetLogons**
Checks that the appropriate logon privileges exist to allow replication to proceed.

**Advertising**
Checks whether each directory server is advertising itself in the roles it should be capable of fulfilling. This test fails if the Net Logon service has stopped or failed to start.

**KnowsOfRoleHolders**
Checks whether the directory server can contact the servers that hold the five operations master (also known as flexible single master operations, or FSMO) roles.

**Intersite**
Checks for failures that would prevent or temporarily hold up intersite replication, and tries to predict how long it will take before the KCC is able to recover.

**Caution**

- Results of this test are often not valid.

**FSMOCheck**
Checks that the directory server can contact a Key Distribution Center (KDC), Time Server, Preferred Time Server, primary directory server (primary domain controller (PDC)), and global catalog server. This test does not test any of the servers for operations master roles.

**RidManager**
Checks whether the RID master is accessible and if it contains the proper information.

**MachineAccount**
Checks whether the machine account is properly registered and the services are advertised.

**Services**
Checks whether the appropriate directory server services are running.

**OutboundSecureChannels**
Checks that secure channels exist from all the directory servers in the domain to the domains that are specified by **/testdomain**. The **/nositerestriction** parameter prevents the test from being limited to the directory servers in the site.

**ObjectsReplicated**
Checks that **Machine Account** and **DSA** objects have replicated. Use **/objectdn:***dn* with **/n:***nc* to specify an additional object to check.

**frssysvol**
Checks that File Replication service (FRS) SYSVOL is ready.

**kccevent**
Checks that KCC is completing without errors.

**systemlog**
Checks that the system is running without errors.

**DcPromo**
This command does not apply to ADAM.

**RegisterInDNS**
This commmand does not apply to ADAM.

**DeadCRTest**
Looks for cross-references that appear to be left over from a failed creation of an application directory partition.

**CheckSDRefDom**
Checks that all application directory partitions have appropriate security descriptor reference domains.

**VerifyReplicas**
Verifies that all application directory partitions are fully instantiated on all replica servers.

**CrossRefValidation**
Verifies the validity of cross-references.

**VerifyReferences**
Verifies that certain system references are intact for the FRS and Replication infrastructure.

**VerifyEnterpriseReferences**
Verifies that certain system references are intact for the FRS and Replication infrastructure across all objects in the enterprise on each directory server.

**{/h | /?}**
Displays a syntax screen at the command prompt.

#### Examples

##### Example 1: A normal ADAM instance

To examine an ADAM instance to verify that it is healthy and functioning properly, type the following at the command prompt:

**dsdiag /s:adam1 /u:cohovineyard\administrator /p:password**

Output similar to the following appears:

```
Directory Service Diagnosis

Performing initial setup:

Done gathering initial info.

Doing initial required tests
```

```
Testing server: Default-First-Site-Name\adam1

Starting test: Connectivity

......................... adam1 passed test Connectivity

Doing primary tests

Testing server: Default-First-Site-Name\adam1

Starting test: Replications

......................... adam1 passed test Replications

Starting test: NCSecDesc

......................... adam1 passed test NCSecDesc

Starting test: NetLogons

......................... adam1 passed test NetLogons

Starting test: Advertising

......................... adam1 passed test Advertising

Starting test: KnowsOfRoleHolders

......................... adam1 passed test KnowsOfRoleHolders

Starting test: RidManager

......................... adam1 passed test RidManager

Starting test: MachineAccount

......................... adam1 passed test MachineAccount

Starting test: Services

......................... adam1 passed test Services

Starting test: ObjectsReplicated

......................... adam1 passed test ObjectsReplicated

Starting test: frssysvol

......................... adam1 passed test frssysvol

Starting test: kccevent

......................... adam1 passed test kccevent

Starting test: systemlog

An Error Event occured. EventID: 0xC25A001D

Time Generated: 12/21/2001 01:28:25

Event String: The time provider NtpClient is configured to

An Error Event occured. EventID: 0xC25A001D

Time Generated: 12/21/2001 01:40:30

Event String: The time provider NtpClient is configured to

An Error Event occured. EventID: 0xC25A001D

Time Generated: 12/21/2001 01:43:30

Event String: The time provider NtpClient is configured to

An Error Event occured. EventID: 0xC25A001D

Time Generated: 12/21/2001 01:58:46

Event String: The time provider NtpClient is configured to

An Error Event occured. EventID: 0xC25A001D

Time Generated: 12/21/2001 02:02:11
```

```
Event String: The time provider NtpClient is configured to

An Error Event occured. EventID: 0xC25A001D

Time Generated: 12/21/2001 02:05:11

Event String: The time provider NtpClient is configured to

An Error Event occured. EventID: 0xC25A001D

Time Generated: 12/21/2001 02:10:51

Event String: The time provider NtpClient is configured to

......................... adam1 failed test systemlog

Running partition tests on : Schema

Starting test: DeadCRTest

......................... Schema passed test DeadCRTest

Starting test: CheckSDRefDom

......................... Schema passed test CheckSDRefDom

Running partition tests on : Configuration

Starting test: DeadCRTest

......................... Configuration passed test DeadCRTest

Starting test: CheckSDRefDom

......................... Configuration passed test CheckSDRefDom

Running partition tests on : cohovineyard

Starting test: DeadCRTest

......................... cohovineyard passed test DeadCRTest

Starting test: CheckSDRefDom

......................... cohovineyard passed test CheckSDRefDom

Running enterprise tests on : cohovineyard.reskit.com

Starting test: Intersite

......................... cohovineyard.reskit.com passed test Intersite

Starting test: FsmoCheck

......................... cohovineyard.reskit.com passed test FsmoCheck
```

**Example 2: Unresponsive or inaccessible server**

To resolve replication problems, type the following at the command line:

**dsdiag /s:adam1 /u:cohovineyard\administrator /p:password /e**

Output similar to the following appears:

```
Directory Service Diagnosis

Performing initial setup:

Done gathering initial info.

Doing initial required tests

Testing server: Default-First-Site-Name\adam1

Starting test: Connectivity

......................... adam1 passed test Connectivity

Testing server: Default-First-Site-Name\RESKIT-DC2

Starting test: Connectivity

Server RESKIT-DC2 resolved to this IP address 172.26.220.34,
```

```
but the address couldn't be reached(pinged), so check the network.
The error returned was: Error due to lack of resources.
This error more often means that the targeted server is
shutdown or disconnected from the network
......................... RESKIT-DC2 failed test Connectivity
Doing primary tests
Testing server: Default-First-Site-Name\adam1
Starting test: Replications
[Replications Check,adam1] A recent replication attempt failed:
From RESKIT-DC2 to adam1
Naming Context: CN=Configuration,DC=cohovineyard,DC=reskit,DC=com
The replication generated an error (1722):
The RPC server is unavailable.
The failure occurred at 2001-12-21 02:19:04.
The last success occurred at 2001-12-21 01:57:43.
1 failures have occurred since the last success.
The source remains down. Please check the machine.
......................... adam1 passed test Replications
Starting test: NCSecDesc
......................... adam1 passed test NCSecDesc
Starting test: NetLogons
......................... adam1 passed test NetLogons
Starting test: Advertising
......................... adam1 passed test Advertising
Starting test: KnowsOfRoleHolders
......................... adam1 passed test KnowsOfRoleHolders
Starting test: RidManager
......................... adam1 passed test RidManager
Starting test: MachineAccount
......................... adam1 passed test MachineAccount
Starting test: Services
......................... adam1 passed test Services
Starting test: ObjectsReplicated
......................... adam1 passed test ObjectsReplicated
Starting test: frssysvol
......................... adam1 passed test frssysvol
Starting test: kccevent
......................... adam1 passed test kccevent
Starting test: systemlog
An Error Event occured. EventID: 0xC25A001D
Time Generated: 12/21/2001 01:28:25
Event String: The time provider NtpClient is configured to
```

```
An Error Event occured. EventID: 0xC25A001D
Time Generated: 12/21/2001 01:40:30
Event String: The time provider NtpClient is configured to
An Error Event occured. EventID: 0xC25A001D
Time Generated: 12/21/2001 01:43:30
Event String: The time provider NtpClient is configured to
An Error Event occured. EventID: 0xC25A001D
Time Generated: 12/21/2001 01:58:46
Event String: The time provider NtpClient is configured to
An Error Event occured. EventID: 0xC25A001D
Time Generated: 12/21/2001 02:02:11
Event String: The time provider NtpClient is configured to
An Error Event occured. EventID: 0xC25A001D
Time Generated: 12/21/2001 02:05:11
Event String: The time provider NtpClient is configured to
An Error Event occured. EventID: 0xC25A001D
Time Generated: 12/21/2001 02:10:51
Event String: The time provider NtpClient is configured to
An Error Event occured. EventID: 0xC25A001D
Time Generated: 12/21/2001 02:13:51
Event String: The time provider NtpClient is configured to
An Error Event occured. EventID: 0xC25A001D
Time Generated: 12/21/2001 02:18:58
Event String: The time provider NtpClient is configured to
......................... adam1 failed test systemlog
Testing server: Default-First-Site-Name\RESKIT-DC2
Skipping all tests, because server RESKIT-DC2 is
not responding to directory service requests
Running partition tests on : Schema
Starting test: DeadCRTest
......................... Schema passed test DeadCRTest
Starting test: CheckSDRefDom
......................... Schema passed test CheckSDRefDom
Running partition tests on : Configuration
Starting test: DeadCRTest
......................... Configuration passed test DeadCRTest
Starting test: CheckSDRefDom
......................... Configuration passed test CheckSDRefDom
Running partition tests on : cohovineyard
Starting test: DeadCRTest
......................... cohovineyard passed test DeadCRTest
Starting test: CheckSDRefDom
```

```
..... cohovineyard passed test CheckSDRefDom

Running partition tests on : reskit-sib

Starting test: DeadCRTest

..... reskit-sib passed test DeadCRTest

Starting test: CheckSDRefDom

..... reskit-sib passed test CheckSDRefDom

Running enterprise tests on : cohovineyard.reskit.com

Starting test: Intersite

..... cohovineyard.reskit.com passed test Intersite

Starting test: FsmoCheck

..... cohovineyard.reskit.com passed test FsmoCheck
```

**Formatting legend**

| **Format** | **Meaning** |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

# Dsmgmt

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

#### Dsmgmt

Provides management facilities for Active Directory Application Mode (ADAM).

To run **dsmgmt**, at an ADAM tools command prompt, type **dsmgmt**, and then press ENTER. At the **dsmgmt:** prompt, type one of the following commands:

- configurable settings
- ds behavior
- partition management
- LDAP policies
- metadata cleanup
- popups
- quit
- roles
- security account management

#### Configurable settings

Helps to modify the Time to Live (TTL) of dynamic data that is stored in ADAM. To reach the **configurable setting:** prompt, at the **dsmgmt:** prompt, type **configurable settings**, and then press ENTER. Then, at the **configurable setting:** prompt, use the following syntax.

##### Syntax

{**cancel changes**|**connections**|**list**|**set** *%s* **to** *%s*|**show values**}

##### Parameters

**cancel changes**
Cancels the changes that are made but that are not yet committed.

**Connections**
Invokes the **server connections** submenu.

**List**
Lists the names of the supported configurable settings.

**set *%s* to *%s***
Sets the configurable settings *%s1* to the value *%s2*.

**show values**
Displays values of configurable settings.

***%s***
An alphanumeric variable, such as an ADAM instance name.

**Quit**
Takes you back to the previous menu or exits the utility.

**? or help**
Displays help at the command prompt.

#### ds behavior

Enables administrators to specify whether password modifications can be made over unsecured (nonencrypted) connections. By default, password modifications over unsecured connections are not allowed. To view the **ds behavior** command-line options, type **?** at the **ds behavior:** command prompt.

#### partition management

Enables administrators to prepare cross-reference and server objects in the directory. To reach the **partition management:** prompt, at the **dsmgmt:** prompt, type **partition management**, and then press ENTER. Then, at the **partition management:** prompt, use the following syntax.

##### Syntax

{**add nc replica** *%s %s*|**connections**|**create nc** *%s %s*|**create nc** *%s %s %s*|**remove nc replica** *%s %s*|**list**|**list nc information** *%s*|**list nc replicas** *%s*|**precreate** *%s %s*|**delete NC** *%s*|**select operation targe**

##### Parameters

**add nc replica *%s %s***
Adds the ADAM instance *%s2* to the replica set for the application directory partition *%s1*. If *%s2* is not specified, the ADAM instance that you are connected to is used as the default.

**Connections**
Invokes the **Connections** submenu.

**create nc *%s %s***
This command only applies to Active Directory (and not to ADAM).

**create nc *%s %s %s***
This command applies to ADAM (and not to Active Directory). Creates the application directory partition with the distinguished name specified by *%s1*, of an object class specified by *%s2*, on the ADAM instance specified by the *computername:portnumber* specified by *%s3*. If *%s3* is not specified, the currently connected ADAM instance is used.

**remove nc replica *%s %s***
Removes the ADAM instance *%s2* from the replica set for the application directory partition *%s1*. If *%s2* is not specified, the ADAM instance that you are currently connected to is used.

**List**
Lists all the naming contexts (or, directory partitions) that exist in the configuration set.

**list nc information *%s***
Prints out the reference domain and replication delays for the application directory partition.

**list nc replicas *%s***
Prints the list of ADAM instances in the replica set for application directory partition *%s*. Remember that this is the list of ADAM instances that will eventually hold replicas of the application directory partition and that these replicas may not necessarily be fully replicated yet.

**precreate *%s %s***
Creates a cross-reference object for the application directory partition *%s1*, allowing an ADAM instance named *%s2* to hold a replica of that directory partition. The directory partition must be specified by a fully distinguished name, and the ADAM instance must be specified by *computername:portnumber*.

**delete nc *%s***
Removes the application directory partition *%s*. Before you remove a application directory partition, all the replicas must be removed and their removal must replicate back to the naming operations master.

**select operation target**
Invokes the **Select operation target** submenu.

**set nc reference domain *%s %s***
This command does not apply to ADAM.

**set nc replicate notification delay *%s %d %d***
Sets the notification delays for the application directory partition specified by *%s* to %d1 and %d2 for the delay in notifying the first ADAM instance of changes and the delay in notifying subsequent ADAM instances of changes, respectively.

***%s***
An alphanumeric variable, such as an ADAM instance name.

***%d***
A numeric variable, such as a replication delay time period.

**Quit**
Takes you back to the previous menu, or exits the utility.

**? or help**
Displays help at the command prompt.

**LDAP policies**
Sets the Lightweight Directory Access Protocol (LDAP) administration limits for the Default-Query Policy object. To reach the **LDAP policies:** prompt, at the **dsmgmt:** prompt, type **ldap policies**, and then press ENTER. At the **LDAP policies:** prompt, use the following syntax.

**Syntax**
{**cancel changes**|**commit changes**|**connections**|**list**|**set** *%s* **to** *%s*|**show values**}

**Parameters**
**cancel changes**
Cancels any uncommitted modifications of the LDAP administration limits to the default query policy.

**commit changes**
Commits all modifications of the LDAP administration limits to the default query policy.

**Connections**
Invokes the **server connections** submenu.

**List**
Lists all supported LDAP administration limits for the ADAM instance.

**set *%s1* to *%s2***
Sets the value of the LDAP administration limit *%s1* to the value *%s2*.

**show values**
Shows the current and proposed values for the LDAP administration limits.

***%s***
An alphanumeric variable, such as an ADAM instance name.

**Quit**
Takes you back to the previous menu or exits the utility.

**? or help**
Displays help at the command prompt.

**Remarks**

- The following table describes LDAP administration limits, with the default values noted in parentheses.

| Value | Description |
|---|---|
| **InitRecvTimeout** | Initial receive time-out (120 seconds) |
| **MaxConnections** | Maximum number of open connections (5000) |
| **MaxConnIdleTime** | Maximum amount of time a connection can be idle (900 seconds) |
| **MaxActiveQueries** | Maximum number of queries that can be active at one time (20) |
| **MaxNotificationPerConnection** | Maximum number of notifications that a client can request for a given connection (5) |
| **MaxPageSize** | Maximum page size supported for LDAP responses (1000 records) |
| **MaxQueryDuration** | Maximum length of time the ADAM instance can execute a query (120 seconds) |
| **MaxTempTableSize** | Maximum size of temporary storage allocated to execute queries (10,000 records) |
| **MaxResultSetSize** | Maximum size of the LDAP Result Set (262,144 bytes) |
| **MaxPoolThreads** | Maximum number of threads created by the ADAM instance for query execution (4 per processor) |
| **MaxDatagramRecv** | Maximum number of datagrams that can be processed by the ADAM instance simultaneously (1024) |

- To ensure that ADAM instances can support service level guarantees, specify operational limits for a number of LDAP operations. These limits prevent specific operations from adversely affecting the performance of a server, and they also make the server resilient to denial-of-service attacks.

  LDAP policies are implemented by using objects of the class **queryPolicy**. Query Policy objects can be created in the container Query Policies, which is a child of the Directory Service container in the configuration naming context. For example: CN=Query-Policies, CN=Directory Service, CN=Windows NT, CN=Services (configuration directory partition).

  An ADAM instance uses the following three mechanisms to apply LDAP policies:

  - An ADAM instance might refer to a specific LDAP policy. The **nTDSASettings** object includes an optional attribute **queryPolicyObject**, which contains the distinguished name of a Query Policy.

  - In the absence of a specific query policy being applied to an ADAM instance, the ADAM instance applies the Query Policy that has been assigned to the ADAM instance's site. The **ntDSSiteSettings** object includes an optional attribute **queryPolicyObject**, which contains the distinguished name of a Query Policy.

  - In the absence of a specific ADAM instance or a site Query Policy, an ADAM instance uses the default Query Policy that is named Default-Query Policy.

A Query Policy object includes the multivalued attributes **LDAPIPDenyList** and **LDAPAdminLimits**. An administrator can use **dsmgmt** to set the LDAP administration limits (but not for the IP Deny list) for the Default-Query Policy object.

#### metadata cleanup

Cleans up metadata for retired ADAM instances. To reach the **metadata cleanup:** prompt, at the **dsmgmt:** prompt, type **metadata cleanup**, and then press ENTER. At the **metadata cleanup:** prompt, use the following syntax.

##### Syntax

{**connections**|**remove selected domain**|**remove selected naming context**|**remove selected server**|**select operation target**}

##### Parameters

**Connections**
Removes metadata associated with the domain that is selected in the **Select operation target** submenu.

**remove selected domain**
This command does not apply to ADAM.

**remove selected naming context**
Removes directory service objects for the selected naming context.

**remove selected server**
Removes metadata associated with the ADAM instance that is selected in the **Select operation target** submenu.

**select operation target**
Invokes the **Select operation target** submenu.

**Quit**
Takes you back to the previous menu, or exits the utility.

**? or help**
Displays help at the command prompt.

##### Remarks

- Metadata is maintained for each ADAM instance in a configuration set. When an ADAM instance is uninstalled, any metadata about the uninstalled ADAM instance that has been replicated to other ADAM instances is also removed. However, in some cases, metadata for the uninstalled ADAM instance can remain on the other ADAM instances in the configuration set. To remove this stale metadata, connect to an ADAM instance that holds the stale metadata, and select the uninstalled ADAM instance as the operation target.

| **Caution** |
| --- |
| Do not delete the metadata of existing ADAM instances. |

#### popups {on|off}

Enables or disables interactive pop-up messages. To enable interactive pop-up messages, at the **dsmgmt:** prompt, type **popups on**. To disable pop-up messages, at the **dsmgmt:** prompt, type **popups off**.

#### quit

Type **quit** at the **dsmgmt:** prompt to close the **dsmgmt** command-line tool.

#### roles

Transfers and seizes operations master roles. To reach the **roles:** prompt, at the **dsmgmt:** prompt, type **roles**, and then press ENTER. At the **roles:** prompt, use the following syntax.

| **Note** |
| --- |
| Only the naming master and schema master roles apply to ADAM. Commands relating to the PDC, RID, and infrastructure masters do not apply to ADAM. |

##### Syntax

{**connections**|**seize naming master**|**seize infrastructure master**|**seize PDC**|**seize RID master**|**seize schema master**|**select operation target**|**transfer naming master**|**transfer infrastructure master**|

##### Parameters

**connections**
Invokes the **server connections** submenu.

**seize naming master**
Forces the ADAM instance to which you are connected to claim ownership of the naming operations master role without regard to the data that is associated with the role. Use only for recovery purposes.

**seize infrastructure master**
Does not apply to ADAM.

**seize PDC**
Does not apply to ADAM.

**seize RID master**
Does not apply to ADAM.

**seize schema master**
Forces the ADAM instance to which you are connected to claim ownership of the schema operations master role without regard to the data that is associated with the role. Use only for recovery purposes.

**select operation target**
Invokes the **Select operation target** submenu. From within the submenu, you can list the operations master (also known as flexible single master operations, or FSMO) roles that are held by the currently connected server. For a complete list of **select operation target** subcommands, type **?** at the **select operation target:** prompt.

**transfer naming master**
Instructs the ADAM instance to which you are connected to obtain the naming role by means of controlled transfer.

**transfer infrastructure master**
Does not apply to ADAM.

**transfer PDC**
Does not apply to ADAM.

**transfer RID master**
Does not apply to ADAM.

**transfer schema master**
Instructs the ADAM instance to which you are connected to obtain the schema operations master role by means of controlled transfer.

**quit**
Takes you back to the previous menu, or exits the utility.

**? or help**
Displays help at the command prompt.

##### Remarks

- Although ADAM is based on a multimaster administration model, some operations support only a single master. For multimaster operations, conflict resolution ensures that after the system finishes replicating, all replicas agree on the value for a given property on a given object. However, some data, for which adequate conflict resolution is not possible, is key to the operation of the system as a

whole. This data is controlled by individual ADAM instances, called operations masters.

The following two operations master roles apply to ADAM configuration sets.

- **Schema operations master.** There is a single schema operations master role for each ADAM configuration set. This role allows the operations master server to accept schema updates. There are other restrictions on schema updates.

- **Naming master.** There is a single naming master role for an ADAM configuration set. The naming master role allows the owner to define new cross-reference objects that represent directory partitions in the partitions container.

- An operations master role can only be moved by administrative involvement; it is not moved automatically. Additionally, moving a role is controlled by standard access controls. Therefore, an organization should tightly control the location and movement of operations master roles. For example, an organization with a strong information technology (IT) presence might place the schema role on a server in the IT group and configure its access control list (ACL) so that it cannot be moved at all.

  Operations master roles require two forms of management: controlled transfer and seizure.

  Use controlled transfer when you want to move a role from one server to another, perhaps to track a policy change with respect to role location or in anticipation of a server being shut down, moved, or decommissioned.

  Seizure is required when a server that is holding a role fails and you do not intend to restore it. Even in the case of a server that is recovered from a backup, the server does not assume that it owns a role (even if the backup tape says so), because the server cannot determine if the role was legitimately transferred to another server in the time period between when the backup was made and when the server failed and was recovered. The restored server assumes role ownership only if a quorum of existing servers is available during recovery and they all agree that the restored server is still the owner.

  The **roles** submenu in **dsmgmt** is used to perform controlled transfer and recovery of operations master roles. Controlled transfer is simple and safe. Because the source and destination servers are running, the system software guarantees that the operations master role token and its associated data is transferred safely. Operations master role seizure is equally simple but not as safe; you simply tell a particular ADAM instance that it is now the owner of a particular role.

| **Caution** |
| --- |
| Do not make a server a role owner by means of seizure commands if the real role holder exists on the network. Doing this can create irreconcilable conflicts for key system data. If an operations master role owner is temporarily unavailable, do not make another ADAM instance the role owner. This can result in a situation in which two computers function as the role owner, which might cause irreconcilable conflicts for key system data. |

**security account management**

This command does not apply to ADAM.

**Formatting legend**

| **Format** | **Meaning** |
| --- | --- |
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

# Ldp

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Ldp**

The **Ldp** dialog box consists of two panes: the console tree and the details pane. The console tree lists the base object and any child objects. The details pane lists the results of the Lightweight Directory Access Protocol (LDAP) operations. To start Ldp, click **Start**, point to **All Programs**, point to **ADAM**, click **ADAM Tools Command Prompt**, and then type **ldp** at the command prompt.

The following sections describe commands on the Ldp menus:

- Connection
- Browse
- View
- Options
- Utilities

**Connection**

The following table describes commands on the **Connection** menu.

| Command | Description |
|---|---|
| **Connect** | Opens a dialog box that you can use to open a session with a specified LDAP server. A connection must be established with an LDAP server before any other LDAP commands can be run. Type the appropriate port number for the service that you are connecting to. By default, LDAP uses TCP for a connection-oriented session. To use User Datagram Protocol (UDP) for a connectionless session, select the **Connectionless** check box. By default, a successful connection causes the appearance of **RootDSE** information in the details pane. |

| | |
|---|---|
| **Bind** | Opens a dialog box that you can use to authenticate a specified LDAP server. Type a user name and password of an account that has permissions to the LDAP server. If you do a simple bind with an empty password, you will be connected with anonymous credentials. As a shortcut, use the **Bind** command without using the **Connect** command to connect, and then authenticate with the server that you last connected to. Click **Advanced** on the **Bind** dialog box to open the **Bind Options** dialog box and configure authentication method options.<br><br>The following are the options in the **Bind Options** dialog box.<br><br>• **Function Type** - Specifies a category of authentication for Ldp to use when choosing authentication methods.<br><br>• **Generic** - Specifies use of a standard authentication protocol.<br><br>• **Simple** - Specifies use of no authentication protocol and sends the password in plaintext.<br><br>• **Extended** - Not available.<br><br>• **Method** - Specifies the type of authentication that Ldp uses when passing credentials.<br><br>• **Synchronous** - Specifies that the authenticating server must respond immediately to requests. This option works only with simple authentication.<br><br>• **Use auth. Identity** - Allows the use of alternate authentication credentials. All authentication methods except simple require synchronous calls. |
| **Disconnect** | Terminates an open session with a specified LDAP server. Closing Ldp automatically disconnects any open sessions. |
| **New** | Keeps the currently connected session, but clears the details pane. The keyboard shortcut for this action is CTRL+N. |
| **Save** | Saves changes to a previously saved file. |
| **Save as** | Saves the contents of the details pane to a text file. Use the **Open** command to view the contents of this file in the details pane later. |
| **Print** | Prints the contents of the details pane. |

**Browse**

The following sections describe commands on the **Browse** menu:

**Add Child**

Opens a dialog box that you can use to add objects to Active Directory. You must enter the full distinguished name of the object, as well as all the mandatory attributes for the class of object being added.

The following table describes options in the **Add** dialog box.

| Option | Description |
|---|---|
| **Dn** | Type the full distinguished name of the new object. |
| **Attribute** | Type the required or optional attribute. |
| **Values** | Type the values that are associated with the attribute. Separate multiple values for a single attribute with a semicolon. No spaces are required. |
| **Enter** | Adds the entered attribute and values to the **Entry List** box, and clears the **Attribute** and **Values** boxes. Continue entering attributes and values until all required and desired optional attributes are in the **Entry List** box. |
| **Insert file** | Opens a dialog box that you can use to open a text file with the appropriate attributes and values. |

| | |
|---|---|
| **Entry List** | Displays the attributes and values that you enter. |
| **Edit** | Opens a dialog box that you can use to enter changes to the selected entry in the **Entry List** box. |
| **Remove** | Deletes the selected entry from the **Entry List** box. |
| **Extended** | Select this check box if the object being added is part of an extended control. |
| **Synchronous** | If this check box is selected (which is the default), Ldp is required to wait for a response from the destination server before continuing. If you clear this check box, Ldp continues before a response is received. Clear this check box when slow wide area network (WAN) connections are causing Ldp commands to time out. |
| **Run** | Adds the current attributes and values in the Entry List to Active Directory. If Ldp encounters any errors, the object is not added and an error message appears in the details pane. |

**Delete**

This command opens a dialog box that you can use to delete an object from Active Directory. Attributes can be deleted only if they are defined as optional and if they contain no values. To delete an attribute's values, on the **Browse** menu in the **Add** dialog box, click **Edit**.

| Option | Description |
|---|---|
| **DN** | Type the full distinguished name of the new object. |
| **Extended** | Select this check box if the object being modified is part of an extended control. |
| **Synchronous** | If this check box is selected (which is the default), Ldp is required to wait for a response from the destination server before continuing. If you clear this check box, Ldp continues before a response is received. Clear this check box when slow WAN connections are causing Ldp commands to time out. |
| **Recursive (client)** | Deletes all objects in a container, but does not delete the container. |

**Modify**

This command opens a dialog box that you can use to change the attributes of an object that is stored in Active Directory.

The following table describes options in the **Modify** dialog box.

| Option | Description |
|---|---|
| **Dn** | Type the full distinguished name of the new object. |
| **Attribute** | Type the required or optional attribute. |
| **Values** | Type the values that are associated with the attribute. Separate multiple values for a single attribute with a semicolon. No spaces are required. |
| **Insert file** | Opens a dialog box that you can use to open a text file with the appropriate attributes and values. |
| **Enter** | Adds the entered attribute and values to the **Entry List** section of the dialog box, and clears the **Attribute** and **Values** fields. Continue entering attributes and values until all required and desired optional attributes are in the **Entry List** box. |
| **Operation** | **Add**, **Delete**, or **Replace**. To add a new value to an existing attribute, click **Add**. To permanently remove an attribute from the listed object, click **Delete**. Attributes that contain data cannot be deleted. Also, attempting to delete required attributes results in an error. To replace an existing value with another or to change listed values for an existing attribute, click **Replace**. |
| **Entry List** | Displays the existing attributes and values for an object. |
| **Edit** | Opens a dialog box that you can use to make changes to the selected entry in the **Entry List** box. |

| | |
|---|---|
| **Synchronous** | If this check box is selected (which is the default), Ldp is required to wait for a response from the destination server before continuing. If you clear this check box, Ldp continues before a response is received. Clear this check box when slow WAN connections are causing Ldp commands to time out. |
| **Extended** | Select this check box if the object being modified is part of an extended control. |
| **Run** | Sends the edited values in the **Entry List** box to Active Directory. |

**Modify DN**

Opens a dialog box that you can use to change the relative distinguished name of an object. This option is designed to modify leaf objects only. If you rename the container portion of the distinguished name, the object moves to the container that is named.

The following table describes options in the **Modify RDN** dialog box.

| Option | Description |
|---|---|
| **Old Dn** | Type the current distinguished name of the object. |
| **New Dn** | Type the new distinguished name for the object. |
| **Delete Old** | If this check box is selected (which is the default), the old distinguished name is removed from the LDAP directory. |
| **Synchronous** | If this check box is selected (which is the default), Ldp is required to wait for a response from the destination server before continuing. If you clear this check box, Ldp continues before a response is received. Clear this check box when slow WAN connections are causing Ldp commands to time out. |
| **Extended rename** | Select this check box if the object being renamed is part of an extended control. |
| **Run** | Sends the change to Active Directory. |

**Search**

Opens a dialog box that you can use to create a customized search filter and to perform the search on the directory information tree. The search base must be specified as a distinguished name, and the filter must be a valid LDAP filter. Items returned from a search are separated by **>>** characters.

The following table describes options in the **Search** dialog box.

| Option | Description |
|---|---|
| **Base Dn** | Type a distinguished name to specify where the search starts. |
| **Filter** | Type the search criteria, separated by LDAP search filters. Type attributes and values to find an object or set of objects. Note that LDAP search filters are defined in RFC 2254 and in the article 255602 at the Microsoft Knowledge Base. |
| **Scope** | Specifies how many levels the search encompasses. |
| **Base** | Searches the base object only. |
| **One Level** | Searches objects immediately subordinate to the base object, but does not search the base object. |
| **Subtree** | Searches the entire subtree, from the base object down to all child objects. |
| **Options** | Opens the **Search Options** dialog box. You can use these options to apply filters that allow some entries and exclude others from the search and that allow you to control the way the search is processed. |
| **Run** | Sends the search request to Active Directory. |

**Compare**

Opens a dialog box that you can use to compare the value of an attribute of an object with a specified value. The result

returned is either true or false.

The following table describes options in the **Compare** dialog box.

| Option | Description |
|---|---|
| **Dn** | Type the full distinguished name of the object whose values will be compared. |
| **Attribute** | Type the attribute to be compared. |
| **Value** | Type the value that will be compared with the existing value in Active Directory. Separate multiple values for a single attribute with a semicolon. No spaces are required. |
| **Synchronous** | If this check box is selected (which is the default), Ldp is required to wait for a response from the destination server before continuing. If you clear this check box, Ldp continues before a response is received. Clear this check box when slow WAN connections are causing Ldp commands to time out. |
| **Run** | Starts the comparison. |

**Extended Op**

Opens a dialog box that you can use to submit an extended operation to an LDAP directory by specifying an LDAP object identifier (also known as OID) and an applicable value.

The following table describes options in the **Extended Operations** dialog box.

| Option | Description |
|---|---|
| **Oid** | Type the object identifier number. |
| **Data** | Type the value of the object identifier attribute. |
| **Controls** | See **Controls Option** in the **Options** section. |
| **Send** | Submits the extended operation to Active Directory. |

**GetLastError**

Calls the LDAP **GetLastError** function.

**Security**

Opens a dialog box that you can use to submit an extended operation to an LDAP directory by specifying an LDAP object identifier and an applicable value.

The following table describes the **Security Descriptor** dialog box.

| Option | Description |
|---|---|
| **Security Descriptor** | Opens a dialog box that you can use to view access permissions on an object. |

**Replication**

Opens a dialog box that you can use to submit an extended operation to an LDAP directory by specifying an LDAP object identifier and an applicable value.

The following table describes the **Replication Metadata** dialog box.

| Option | Description |
|---|---|
| **Replication/View Metadata** | In **Object DN**, type the distinguished name of the object whose replication metadata you want to view. |

**Process Pending**

Opens a dialog box that shows the list of requests that are not finished processing.

**View**

The following table describes commands on the **View** menu.

| Command | Description |
|---|---|

| | |
|---|---|
| **Tree** | In **BaseDN**, type the distinguished name of the object to use as the base object in the navigation pane. |
| **Enterprise Configuration** | This menu option does not apply to ADAM. |
| **Status Bar** | Shows or hides the status bar, which is located along the bottom of the LDAP window. |

**Options**

The following sections describe the commands on the **Options** menu:

**Search**

The following table describes options in the **Search Options** dialog box.

| Option | Description |
|---|---|
| **Time limit** | Type the number of milliseconds that the search can take on the server. By default, the maximum is 120 seconds. |
| **Size limit** | Type the maximum number of bytes that the search can return. Typing a null value does not place a maximum size on the data that is returned. |
| **Timeout (s)** | Type the number of seconds that Ldp waits for the LDAP server to respond to a search request. |
| **Timeout (ms)** | Type the number of milliseconds that Ldp waits for the LDAP server to respond to a search request. |
| **Page size** | Type the maximum size, in bytes, of each page of returned data. |
| **Attributes** | Specifies which attributes to return in the search. Separate multiple attributes with a semicolon. Use the wildcard character (*****) to indicate all attributes. |
| **Search Call Type** | Specifies a call type to use in the search. If the search will take some time, you can click **Async.** so that you can perform other tasks while waiting for the search to complete. |
| **Attributes Only** | Select this check box to return only attributes of objects. The distinguished name is not returned. |
| **Chase referrals** | Performs a search for objects that are found in external LDAP directories. By default, the objects' trusts of external LDAP directories return only a referral instead of the actual object. |
| **Display Results** | Displays a detailed list of objects that are returned by the search. By default, only a success or failure and the number of objects found appear. |
| **Sort Keys** | Opens the **Sort Keys** dialog box. See the **Sort Keys** section below. |
| **Controls** | Opens the **Controls** dialog box. |

**Pending**

Opens a dialog box that you can use to place filters on the list of processes that have not yet completed.

The following table describes options in the **Pending Options** dialog box.

| Option | Description |
|---|---|
| **All search results** | Specifies that all search results display. |
| **Blocking** | Clear this check box to set a time limit. |
| **Time Limit (sec):** | Type a time limit in seconds. |
| **Time Limit (millisec):** | Type a time limit in milliseconds. |

**General**

The following table describes options in the **General Options** dialog box.

| Option | Description |
| --- | --- |
| **Value Parsing** | Specifies the display format of LDAP data. **Binary** displays the LDAP information in its native numerical format. **String** Converts the LDAP information from its native format to ASCII characters so that it is more readable when displayed. This is the default setting. Values that are too long to be converted are still displayed in binary form. |
| **LDAP Version** | Specifies which version of LDAP the server is using. The default is **version 3**. |
| **DN Processing** | Converts the distinguished names, displayed in component parts, by extending the data types that Ldp returns when performing a command. |
| **Buffer Size** | In **Number of lines**, type the number of returned lines to display per command. In **Chars per line**, type the number of returned characters to display per command. |
| **Auto default NC query** | Specifies that Ldp queries the default naming context when a connection to the LDAP server is made. The default naming context is the **RootDSE**. This setting is used when the distinguished name value in the **View\|Tree** dialog box is left blank. |
| **Virtual List View (VLV)** | Select **Auto VLV browse when** to display a virtual list view whenever the object count is greater than the value that is displayed in **container size is greater than**. The default value is 100. |

**Connection Options**

Opens a dialog box that you can use to change the value of any option.

The following table describes options in the **Connection Options** dialog box.

| Option | Description |
| --- | --- |
| **Option Name** | Type the name of the option whose value will be reset. |
| **Value** | Type the new value for the specified option. |
| **Set** | Sends the information to the LDAP directory. |

**Controls**

In the **Controls** dialog box, enter information to extend the functionality of LDAP.

The **Object Identifier** option must be specified when you implement a control. To obtain a list of object identifiers, view the **supportedControls** property in the **RootDSE** of a domain controller. Individual controls are described in the "Understanding LDAP" white paper at the Microsoft Web site.

### Notes

- Only server controls can be sent to a server. Client controls work only with LDAP application programming interfaces (APIs).
- To view a list of extended LDAP controls, see article 222560 at the Microsoft Knowledge Base.

**Sort Keys**

**Sort Keys** is a type of control that formats the display of search results. To format the display of your search results, in the **Sort Keys** dialog box, type an attribute type.

For more information, search for "sortKeyRequestControl" in the "Understanding LDAP" white paper at the Microsoft Web site.

**TLS**

The following table describes the subcommands for the **TLS** command.

| Option | Description |
| --- | --- |
| **StartTLS** or **StopTLS** | Starts or stops a secure session with the LDAP server that uses Transport Level Security (TLS). |

**Utilities**

The following table describes the commands on the **Utilities** menu.

| Option | Description |
| --- | --- |
| **Large Integer Converter** | To convert long integers into high and low parts, type a value in the **String** box. |
| **SID Lookup** | To determine the *domain\user* that is associated with a given security ID (SID), type a SID into **List of sids**, and then click **OK**. The associated *domain\user* appears in the details pane. |

# Repadmin

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Repadmin**
Performs tasks related to replication, including managing and modifing replication topology, forcing replication events, and displaying replication metadata and up-to-date vectors. **Repadmin** contains several subcommands, which are described below in detail.

### Notes

- Most **repadmin** subcommands apply to both Active Directory Application Mode (ADAM) and Active Directory, except as noted in the description for each subcommand.
- Not all general parameters apply to all subcommands.

**General syntax**
**repadmin** *subcommand* [*dsa*] [**/u:** *UserName*] [**/pw:** {*Password*|*****}] [**/rpc**] [**/ldap**] [**/homeserver:***dsaname*]

**General parameters**
***subcommand***
One of the **repadmin** subcommands described in the "Subcommands" section.

| Value | Description |
|---|---|
| *Dsaname* | One of the following:<br>• *computername*:*portnumber*<br>Specifies an ADAM instance or Active Directory domain controller, using the name of the computer on which the directory service is running and (for ADAM only) the port number used for communication, for example, **adam01:389** or **adam01:636**.<br>• *DNSname*:*portnumber*<br>Specifies an ADAM instance or Active Directory domain controller by Domain Name System (DNS) name of the computer on which the directory service is running and (for ADAM only) the port number used for communication, for example, **adam01.microsoft.com:389** or **adam01.microsoft.com:636**. |
| *dsaguid* | Specifies an ADAM instance or Active Directory domain controller using the globally unique identifier (GUID) of the directory service. When you use "*dsaguid*" against an ADAM configuration set, you must also use the **/homeserver:** parameter. |
| . | **Repadmin** attempts to select a directory server automatically. When you use "." against an ADAM configuration set, you must also use the **/homeserver:** parameter. |
| ***** | **Repadmin** selects all ADAM instances or Active Directory domain controllers in the ADAM configuration set or Active Directory forest. When you use "*" against an ADAM configuration set, you must also use the **/homeserver:** parameter. |
| **site**:*sitename* | Selects all ADAM instances or Active Directory domain controllers in the site called *sitename*, for example, **site:microsoftsite1**. When you use "site:" against an ADAM configuration set, you must also use the **/homeserver:** parameter. |
| **fsmo_dnm:** | Selects the directory server holding the naming operation master role (for ADAM), or the domain naming operation master role (for Active Directory). When you use **fsmo_dnm:** against an ADAM configuration set, you must also use the **/homeserver:** parameter. |

| | |
|---|---|
| **fsmo_schema:** | Selects the directory server holding the schema operation master role. When you use **fsmo_schema:** against an ADAM configuration set, you must also use the **/homeserver:** parameter. |

***dsa* {*dsaname*|*dsaguid*|.|*|site:*sitename*|fsmo_dnm:|fsmo_schema:}**

The *dsa* parameter represents the directory server (ADAM instance or Active Directory domain controller) to be targeted by the **repadmin** subcommand. The *dsa* parameter can be provided in a number of ways, as described in the following table. (Type **repadmin /listhelp** at the command line for additional information about the *dsa* parameter.)

**/u: *UserName***

Specifies the account name to use for binding to the directory. By default, **/u** uses the account name with which the user is currently logged on. You can use any of the following formats to specify an account name: *account name* (for example, **Linda**), *domain\account name* (for example, **widgets\Linda**), or *user principal name (UPN)* (for example, **Linda@cohovineyard.com**)

**/pw: {*Password*|*}**

Specifies the password to use for authentication. If you type *****, you are prompted for a password.

**/rpc**

Forces **repadmin** to communicate using a remote procedure call (RPC) session.

**/ldap**

Forces **repadmin** to communicate using a Lightweight Directory Access Protocol (LDAP) session. If LDAP communication fails, **repadmin** attempts to communicate using RPC. LDAP is the default communication method for **repadmin**.

**/homeserver: *dsaname***

Forces **repadmin** to run against a specific ADAM configuration set or Active Directory forest, determined by the configuration set or forest membership of the directory server that is represented by *dsaname*. *Dsaname* can be specified in the following formats for an ADAM instance: *computername:portnumber* or *DNSname:portnumber Dsaname* can be specified in the following formats for an Active directory domain controller: *computername* or *DNSname*. You must use this parameter when you run **repadmin** against an ADAM configuration set, and *dsaguid*, "*****", "**.**", "**site:***site*", "**fsmo_dnm:**", or "**fsmo_schema:**" are used as the value for the *dsa* parameter.

**Note**

- Not all **repadmin** subcommands require the *dsa* parameter.

**Subcommands**

To view the **repadmin** subcommand syntax, click a **repadmin** subcommand:

- repadmin bind
- repadmin bridgeheads
- repadmin checkprop
- repadmin dsaguid
- repadmin failcache
- repadmin istg
- repadmin kcc
- repadmin latency
- repadmin notifyopt
- repadmin queue
- repadmin replauthmode

- repadmin replicate
- repadmin replsingleobj
- repadmin showattr
- repadmin showcert
- repadmin showchanges
- repadmin showconn
- repadmin showctx
- repadmin showism
- repadmin showmsg
- repadmin showncsig
- repadmin showobjmeta
- repadmin showoutcalls
- repadmin showproxy
- repadmin showrepl
- repadmin showscp
- repadmin showtime
- repadmin showtrust
- repadmin showutdvec
- repadmin showvalue
- repadmin syncall
- repadmin viewlist
- repadmin writespn

**repadmin bind**

Connects to and displays the replication features for a directory server.

**Syntax**

**repadmin /bind** [*dsa*] [**/u:** *UserName*] [**/pw:** {*Password*|*****}] [**/rpc**] [**/ldap**]

**Parameters**

**/bind [*dsa*]**

Specifies the directory server for which information will be displayed. For details about the *dsa* parameter, see General parameters.

**Other parameters**

For details about other parameters, see General parameters.

#### repadmin bridgeheads

Lists the directory servers acting as bridgehead servers for a specified site.

##### Syntax

**repadmin /bridgeheads** [*dsa*] [**/verbose**] [**/u:** *UserName*] [**/pw:** {*Password*|*****}] [**/rpc**] [**/ldap**]

##### Parameters

**/bridgeheads *dsa***

Specifies a directory server that is located in the site for which bridgehead servers will be listed. For details about the *dsa* parameter, see General parameters.

**Other parameters**

For details about other parameters, see General parameters.

#### repadmin checkprop

Compares the properties of specified directory servers to determine if they are up to date with each other. The source directory server contains the original information that needs to be checked. The data on the destination directory server is compared to the data on the source directory server.

##### Syntax

**repadmin /checkprop** [*dsa*] *NamingContext OriginatingDCInvocationID OriginatingUSN*

##### Parameters

***Dsa***

Specifies a directory server. For details about the *dsa* parameter, see General parameters.

***NamingContext***

Specifies the distinguished name of the directory partition on the source directory server.

***OriginatingDCInvocationID***

Specifies the unique hexadecimal number that identifies an object on a source directory server. The *InvocationID* can be retrieved by using **showrepl**.

***OriginatingUSN***

Specifies the Update Sequence Number (USN) for the object on the source directory server. The USN is for the object whose *InvocationID* is already listed.

#### repadmin dsaguid

Returns a server name when given a GUID.

##### Syntax

**repadmin /dsaguid** [*dsa*] [*GUID*]

##### Parameters

***Dsa***

Specifies a directory server. For details about the *dsa* parameter, see General parameters.

***GUID***

Specifies the unique hexadecimal number that identifies the directory server. The GUID can be retrieved by using **showrepl**.

#### repadmin failcache

Displays a list of failed replication events that are detected by Knowledge Consistency Checker (KCC).

##### Syntax

**repadmin /failcache** [*dsa*]

##### Parameters

***Dsa***

Specifies a directory server. For details about the *dsa* parameter, see General parameters.

#### repadmin istg

Returns the computer name of the Intersite Topology Generator (ISTG) server for a specified site.

##### Syntax

**repadmin /istg** [*dsa*] [**/verbose**]

##### Parameters

***Dsa***
Specifies a directory server. For details about the *dsa* parameter, see General parameters.

**/verbose**
Lists detailed information.

**repadmin kcc**
Forces KCC to calculate replication topology for a specified directory server. By default, this calculation occurs every 15 minutes.

**Syntax**
**repadmin /kcc** [*dsa*] [**/async**]

**Parameters**
***Dsa***
Specifies a directory server. For details about the *dsa* parameter, see General parameters.

**/async**
Specifies that the operation will be asynchronous. This means that **repadmin** starts the replication event, but it does not expect an immediate response from the destination directory server. Use this parameter to start **kcc** without waiting for it to complete.

**repadmin latency**
Displays the amount of time between replications, using the ISTG Keep Alive time stamp.

### Notes

- This command does not apply to ADAM.
- The ISTG Keep Alive time stamp is not used in forests that are set to Windows Server 2003 forest functional level. Instead, in those environments, use **repadmin /showutdvec /latency**.

**Syntax**
**repadmin /latency** [*dsa*] [**/verbose**]

**Parameters**
***dsa***
Specifies a directory server. For details about the *dsa* parameter, see General parameters.

**/verbose**
Lists detailed information.

**repadmin notifyopt**
Displays or sets the notification timing settings for replication of a specified directory partition.

**Syntax**
**repadmin /notifyopt** [*dsa*] *NamingContext* [**/first:***Value*] [**/subs:***Value*]

**Parameters**
***dsa***
Specifies a directory server. For details about the *dsa* parameter, see General parameters.

***NamingContext***
Specifies the distinguished name of the directory partition on the source directory server.

**/first: *Value***
The number of seconds after a change is made before the directory server notifies its first replication partner that there is a change.

**/subs: *Value***
After the first replication partner is notified of a change, specifies the number of seconds to wait before notifying the next replication partner.

**repadmin queue**
Displays tasks that are waiting in the replication queue.

**Syntax**

**repadmin /queue** [*dsa*]

#### Parameters

***dsa***

Specifies a directory server. For details about the *dsa* parameter, see General parameters.

#### repadmin replauthmode

Modifies or displays the replication authentication mode in use by an ADAM configuration set.

#### Syntax

**repadmin /replauthmode** *dsa* [*newmode*]

#### Parameters

***Dsa***

Specifies a directory server. For details about the *dsa* parameter, see General parameters.

***newmode***

Specifies the replication authentication mode to be used by the configuration set. If this parameter is not specified, the command will display the current replication authentication mode. For information about replication authentication modes in ADAM, see Understanding ADAM replication and configuration sets.

### Other parameters

For details about other parameters, see General parameters.

#### repadmin replicate

Starts synchronization for the specified directory partition between the source and destination directory servers. The source GUID can be determined when viewing the replication partners by using **showrepl**.

### Syntax 1

**repadmin /replicate** *destination_dsa source_dsa* [**/force**] [**/async**] [**/full**] [**/addref**]

### Syntax 2

**repadmin /replicate** *destination_dsa* [**/force**] [**/async**] [**/full**] [**/addref**] **/allsources**

#### Parameters

***destination_dsa***

Specifies the host_name[:port] of the destination directory server with which you want to replicate. This parameter cannot be used with **/allsource**. For details about valid values for this parameter, see *dsa* in General parameters.

***source_dsa***

Specifies the host_name[:port] of the source directory server with which you want to replicate. This parameter cannot be used with **/allsource**. For details about valid values for this parameter, see *dsa* in General parameters.

**/force**

This parameter is used to override the **Disable Replication** option on a directory server. Do not use this parameter unless you are certain that replication has been disabled, and that you want to override this setting.

**/async**

Specifies that the operation will be asynchronous. This means that **repadmin** starts the replication event, but it does not expect an immediate response from the destination directory server. Use this parameter when there are slow links between directory servers.

**/full**

Forces a full replication of all objects from the destination directory server.

**/addref**

Directs the source to check for a notification entry on the source. If the source does not have a notification entry for this destination, one is added.

**/allsources**

A given destination can have multiple sources for the same naming context. Directs the destination to sync with all sources instead of just one. This parameter cannot be used with *source_dsa*.

**/readonly**

This parameter is ignored by the **/replicate** operation.

#### repadmin replsingleobj

Replicates a single object between any two directory servers that have partitions in common. The two directory servers do not have a replication agreement. Replication agreements can be shown by using the **repadmin /showrepl** command.

##### Syntax

**repadmin /replsingleobject** *dsa DsaSourceGUID ObjectDN*

##### Parameters

***dsa***

Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

***DsaSourceGUID***

Specifies the unique hexadecimal number that identifies the object that will be replicated. The **objectGUID** can be retrieved by using **showrepl**.

***ObjectDN***

Specifies the distinguished name of the object.

#### repadmin showattr

The **/showattr** operation displays objects and their attributes.

##### Syntax

**repadmin /showattr** *dsa OBJ_LIST* [*OBJ_LIST_OPTIONS*] [**/attr**|**/attrs:** *attribute attribute* ...] [**/allvalues**] [**/long**] [**/nolongblob**] [**/nolongblob**] [**/nolongfriendly**] [**/dumpallblob**]

##### Parameters

***Dsa***

Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters. This is a required parameter.

***OBJ_LIST***

This parameter takes a distinguished name or a special keyword that expands into a distinguished name. The keywords are as follows: **<***DistinguishedName***>** s**pecify the distinguished name of an object**; **Ncobj:config:**; **Ncobj:schema:**; **Ncobj:domain:** This keyword represents the distinguished name of the domain directory partition of the directory server; **Dsaobj:** This keyword represents the NTDS settings object of the directory server; and, **""** Use an empty string to specify the RootDSE.

***OBJ_LIST_OPTIONS***

The OBJ_LIST_OPTION parameter is required to perform a generic LDAP search from the command line. The parameter requires a **BaseDN**, with the ability to use a search modifier option. The valid search modifier options are as follows: **/filter:**<ldap_filter>; **/base**; **/subtree**; and, **/onelevel**. See the LDAP documentation for the LDAP filter syntax.

**/atts:<*attribute, attribute...*>**

Returns only the attributes that are specified. Separate each listed attribute with a comma. By default, returns all attributes.

**/allvalues**

For an attribute, the tool only displays 20 values unless this flag is specified, in which case it shows all values.

**/long**

Displays one value per line.

**/nolongblob**

This parameter is ignored by **/showattr**.

**/nolongfriendly**

This parameter is ignored by **/showattr**.

**/friendlyblob[:*attribute attribute* ...]**

If any attributes are specified, only those attributes that are named are displayed in friendly format. Otherwise, all known BLOBs are formatted.

**/nolongfriendly**

Does not format BLOBs.

**/dumpallblob**

Dumps the BLOB in a default byte-by-byte format if there is not a friendly formatted interpretation available for it.

**/nodumpallblob**
Does not show the BLOB if there is not a formatted interpretation available.

**/attr|/attrs:** ***attribute attribute ...***
The parameters **/attr** or **/attrs:**, if present, specify that only those named attributes are displayed.

| **Note** |
| --- |
| A BLOB in this context means an attribute that is not a simple type, like a string or an integer. A BLOB is a complex structured type that is stored as binary bytes. To make sense of the BLOB, a program must interpret it and format it. A friendly BLOB is a BLOB that the program knows about and can format in an understandable way. The program has a list of BLOBs that it understands. |

**repadmin showcert**
Displays the certificates (used with Simple Mail Transfer Protocol (SMTP)-based replication) that are loaded on a specified directory server.

| **Note** |
| --- |
| This command does not apply to ADAM. |

**Syntax**
**repadmin /showcert** *dsa*

**Parameters**
***dsa***
Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

**repadmin showchanges**
Displays changes from a specified directory partition or changes to a specified object. "Syntax 1" saves changes to a directory partition. If this information is saved to a file, the **getchanges** operation can be run again for comparison. "Syntax 2" lists changes to a specified object.

| **Note** |
| --- |
| For this command to run properly, the account under which the command is run must possess the **replication get changes** right on the specified directory partition. |

**Syntax 1**
**repadmin /showchanges** *source_dsa NamingContext* [**/cookie:** *File*] [**/atts:** *attribute,attribute,...*]

**Syntax 2**
**repadmin /showchanges** *dest_dsa SourcedsaObjectGUID NamingContext* [**/verbose**] [**/statistics**] [**/noincremental**] [**/objectsecurity**] [**/ancestors**] [**/atts:** *attribute1,attribute2,...*] [**/filter:** *ldap filter*]

**Parameters**
***dest_dsa***
Specifies the host name of the destination directory server from which to enumerate the host directory servers.

***source_dsa***
Specifies the host name of the directory server that hosts the directory partition whose changes you want to view.

***NamingContext***
Specifies the distinguished name of the directory partition.

**/cookie:** ***File***
Specifies a name for the file to which list changes are saved.

**/atts:** ***attribute1attributes...***
Returns only the attributes specified. Separate each listed attribute with a comma.

***SourcedsaObjectGUID***
Specifies the unique hexadecimal number that identifies the object whose changes will be listed. The **objectGUID** can be retrieved by using the **showrepl** operation.

**/verbose**
Lists detailed information.

**/statistics**
Displays a summary of information about changes, instead of a list of individual changes.

**/noincremental**
Returns changes in value change format, which lists current values for attributes as well what attributes have been added or deleted. If not specified, changes are returned in attribute change format, which shows only the current value of the attribute.

**/objectsecurity**
Overrides the need for the Get Changes right to the directory partition. By default, this right is needed to run the **GetChanges** parameter. However, only changes that the currently logged on user has the rights to view are displayed.

**/ancestors**
Returns changes in USN order.

**/filter: *ldap filter***
Returns only those changes that meet the filter requirements.

| **Note** |
| --- |
| The information from "Syntax 1" can be saved to a file for later comparison. |

#### repadmin showconn
Displays the connection objects for a specified directory server. The default is **local site**.

##### Syntax
**repadmin /showconn** [*dsa*] [*ServerRDN*|*ContainerDN*|*dsa_GUID*] [**/From:***ServerRDN*] [**/intersite**]

##### Parameters
***dsa***
Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

***ServerRDN***
Specifies the relative distinguished name of a server.

***ContainerDN***
Specifies the distinguished name of a container.

***dsa_GUID***
Specifies the unique hexadecimal number that identifies the directory server. The GUID can be retrieved by using the **showrepl** operation.

**/From: *ServerRDN***
Displays only the connection objects that exist to the specified server. By default, all connection objects are displayed.

**/intersite**
Displays only those connection objects that are between sites.

#### repadmin showctx
Displays a list of computers that have opened sessions with a specified directory server.

##### Syntax
**repadmin /showctx** [*dsa*] [**/nocache**]

##### Parameters
***dsa***
Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

**/nocache**
Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings.

#### repadmin showism
Queries the Intersite Messaging Service (ISM) for site routes. This operation cannot be executed remotely.

| Note |
| --- |
| This command does not apply to ADAM. |

**Syntax**

**repadmin /showism** [*TransportDN*] [**/verbose**]

**Parameters**

***TransportDN***

Specifies whether the mail server is using SMTP or RPC to send messages.

**/verbose**

Lists detailed information.

| Note |
| --- |
| This command cannot be run against a remote directory server. |

**repadmin showmsg**

Displays the error message for a given error number.

**Syntax**

**repadmin /showmsg** {*Win32Error*|*DSEventID* **/NTDSMSG**}

**Parameters**

***Win32Error***

Returns a short description of the given Win32 error code.

***DSEventID*** **/NTDSMSG**

Returns the actual event log text for the specified event ID.

**repadmin showncsig**

Each directory server maintains a directory partition signature list. This command displays a list of the removed application partition GUIDs. An application directory partition can be configured to be held or not held on a particular directory server using **dsdbutil** (for ADAM), or **ntdsutil** (for Active Directory).

**Syntax**

**repadmin /showncsig** [*dsa*]

**Parameters**

***dsa***

Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

**repadmin showobjmeta**

Displays the replication metadata for a specified object that is stored in the directory, including attribute ID, version number, originating and local USN, and originating server's GUID and Date and Time stamp. By comparing the replication metadata for the same object on different directory servers, an administrator can determine whether replication has taken place.

**Syntax**

**repadmin /showobjmeta** [*dsa*] *ObjectDN* [**/nocache**] [**/linked**]

**Parameters**

***dsa***

Specifies the host name of the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

***ObjectDN***

Specifies the distinguished name of the object.

**/nocache**

Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings.

**/linked**

Displays metadata that is associated with, but not stored with, the specified object.

**repadmin showoutcalls**

Displays calls that have been made by the specified directory server to other directory servers but not yet answered.

**Syntax**

**repadmin/showoutcalls** [*dsa*]

**Parameters**

***dsa***

Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

**repadmin showproxy**

Lists cross-domain move proxy objects. When an object is moved from one domain to another, a marker remains in the original domain. This marker is called a proxy.

| **Note** |
| --- |
| This command does not apply to ADAM. |

**Syntax 1**

**repadmin /showproxy** [*dsa*] [*NamingContext*] [**matchstring**]

**Syntax 2**

**repadmin /showproxy** [*dsa*] [*ObjectDN*] [**matchstring**] **/movedobject**

**Parameters**

***dsa***

Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

***NamingContext***

Specifies the distinguished name of the directory partition on the source directory server.

***ObjectDN***

Specifies the distinguished name of the object.

**matchstring**

Specifies a filter for the output. Type a string of characters that must be present in the distinguished name to display the object.

**movedobject**

Displays a history of information from the original domain on a moved object after it has reached the new domain.

**repadmin showrepl**

Displays replication information. Inbound replica links are displayed by default. Outbound links can also be shown, as well as connections corresponding to those links. The command also displays errors corresponding to replica links that cannot be created by KCC. This helps the administrator build a visual representation of the replication topology and see the role of each directory server in the replication process.

**Syntax**

**repadmin /showrepl** [*dsa*] [*SourceDCObjectGUID*] [*NamingContext*] [**/verbose**] [**/nocache**] [**/repsto**] [**/conn**] [**/csv**] [**/all**] [**/errorsonly**] [**/intersite**]

**Parameters**

***dsa***

Specifies the host name of the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

***SourceDCObjectGUID***

Specifies the unique hexadecimal number that identifies the object whose replication events will be listed.

***NamingContext***

Specifies the distinguished name of the directory partition.

**/verbose**

Lists detailed information.

**/nocache**

Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings.

**/repsto**
Lists the directory servers that pull replication information from the specified directory partition. To see the outbound neighbors, specify **/repsto** or **/all**.

**/conn**
Displays the connection objects that are associated with each link.

**/csv**
Displays the output of the **repadmin showrepl** operation in a Comma Separated Variable (CSV) format for viewing and analysis in Microsoft Excel. **Repadmin** supports redirection of screen output to a file.

**/all**
Displays all replication partners.

**/errorsonly**
Only shows the partnership if it has an error associated with it.

**/intersite**
Only shows this partnership if the source server belongs to a different site than the site of the server on which the command is being run.

#### repadmin showscp

Displays service connection point (SCP) objects in Active Directory that have been created by ADAM instances. This command must be run against an Active Directory domain controller.

**Syntax**

**repadmin /showscp** [*dsa*]

**Parameters**

***dsa***
Specifies the host name of a directory server that the object will be replicated to. For details about the *dsa* parameter, see General parameters.

### repadmin showsig

Displays the retired invocation IDs on a directory server. A directory server changes its invocation ID when it is restored or when it rehosts an application partition.

**Syntax**

**repadmin /showsig** [*dsa*]

**Parameters**

***dsa***
Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

#### repadmin showtime

Converts a directory service time value to string format for both the local and the Universal Time, Coordinated (UTC) time zones.

**Syntax**

**repadmin /showtime** [*DSTimeValue*]

**Parameters**

***DSTimeValue***
Specifies the time value to be converted.

| **Note** |
| --- |
| With parameters omitted, **repadmin /showtime** displays the current system time in both the directory service format and the string format. |

#### repadmin showtrust

Lists all Active Directory domains that are trusted by a specified Active Directory domain.

| Note |
| --- |
| This command does not apply to ADAM. |

**Syntax**

**repadmin /showtrust** [*dsa*]

**Parameters**

***dsa***

Specifies a directory server. For details about the *dsa* parameter, see General parameters.

**repadmin showutdvec**

Displays the highest USN for the specified directory server. This information shows how up to date a replica is with its replication partners.

**Syntax**

**repadmin /showutdvec** *dsa NamingContext* [**/nocache**] [**/latency**]

**Parameters**

***dsa***

Specifies the directory server to which the object will be replicated. For details about the *dsa* parameter, see General parameters.

***NamingContext***

Specifies the distinguished name of the directory partition.

**/nocache**

Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings.

**/latency**

Sorts the information by the time required to complete the replication. By default, the information is sorted by USN.

**repadmin showvalue**

Displays the values of the type, last modified time, originating directory server, and distinguished name of a specified object.

**Syntax**

**repadmin /showvalue** [*dsa*] *ObjectDN* [*AttributeName*] [*ValueDN*] [**/nocache**]

**Parameters**

***dsa***

Specifies a directory server. For details about the *dsa* parameter, see General parameters.

***ObjectDN***

Specifies the distinguished name of the object.

***AttributeName***

Specifies a single attribute whose value you want to display.

***ValueDN***

Specifies the distinguished name of the attribute that is displayed.

**/nocache**

Specifies that GUIDs are left in hexadecimal form. By default, GUIDs are translated into strings.

**repadmin syncall**

Synchronizes a specified directory server with all replication partners. This command contains several subcommands, which are described below.

| Note |
| --- |
| By default, if no directory partition is provided in the *namingcontext* parameter, the command performs its operations on the configuration directory partition. |

**Syntax**

**repadmin /syncall** *dsa* [*NamingContext*] [*Flags*]

**Parameters**

***dsa***
Specifies a directory server. For details about the *dsa* parameter, see General parameters.

***NamingContext***
Specifies the distinguished name of the directory partition.

***Flags***
Performs specific actions during the replication, as indicated in the following table.

| Flag | Description |
|---|---|
| **/a** | Abort if any server is unavailable. |
| **/A** | Synchronize all directory partitions that are held on the directory server represented by *dsa*. |
| **/d** | Identify servers by distinguished name (rather than GUID DNS) in messages. |
| **/e** | Synchronize the directory partition across all sites. (The default behavior is to synchronize the directory partition only in the site in which the directory server represented by *dsa* resides.) |
| **/h** | Display help for **repadmin /syncall**. |
| **/i** | Iterate indefinitely. |
| **/l** | Perform **repadmin /showrepl** on each directory server pair in the path, rather than performing synchronization. |
| **/j** | Synchronize adjacent directory servers only. |
| **/p** | Pause for possible user abort after every message. |
| **/P** | Push changes outward from the directory server that is represented by *dsa*. (The default behavior of the command is to pull, rather than to push, changes.) |
| **/q** | Run in quiet mode; suppress callback messages. |
| **/Q** | Run in very quiet mode; report fatal errors only. |
| **/s** | Perform a topology analysis and generate messages, but do not synchronize the directory partition. |
| **/S** | Skip initial directory server response checking, and assume that all directory servers are available. |

**repadmin viewlist**
Displays a list of directory servers.

**Syntax**
**repadmin /viewlist** [*dsa*] [*OBJ_LIST*]

**Parameters**
***dsa***
Specifies a directory server. For details about the *dsa* parameter, see General parameters.

***OBJ_LIST***
Accepts a distinguished name or a special keyword that expands into a distinguished name. The keywords include the following: **Ncobj:config:**; **Ncobj:schema:**; **Ncobj:domain:** This keyword is the domain partition distinguished name of the home server; and, **Dsaobj:** This keyword is the NTDS settings object of the home server.

**repadmin writespn**
Updates the service principal names (SPNs) in Active Directory.

| **Note** |
|---|
| This command does not create SPNs in ADAM. However, ADAM setup may generate scripts that utilize this command for writing SPNs for ADAM instances into Active Directory. For more information, see Administering ADAM service principal names. |

**Syntax**

**repadmin /writespn** [*dsa*] {**add|replace|delete**} *accountDN SPN*

**Parameters**

***dsa***

Specifies a directory server. For details about the *dsa* parameter, see General parameters.

**Formatting legend**

| **Format** | **Meaning** |
|---|---|
| *Italic* | Information that the user must supply |
| **Bold** | Elements that the user must type exactly as shown |
| Ellipsis (...) | Parameter that can be repeated several times in a command line |
| Between brackets ([]) | Optional items |
| Between braces ({}); choices separated by pipe (\|). Example: {even\|odd} | Set of choices from which the user must choose only one |
| Courier font | Code or program output |

# Glossary

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Symbols | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

**Symbols**

**1.2.840.113556.1.4.1851**

The object identifier that is assigned to Active Directory Application Mode (ADAM). You can use this object identifier programmatically to differentiate ADAM from other directory services.

See also: directory service; object identifier

**A**

**Active Directory**

The Windows-based directory service. Active Directory stores information about objects on a network and makes this information available to users and network administrators. Active Directory gives network users access to permitted resources anywhere on the network using a single logon process. It provides network administrators with an intuitive, hierarchical view of the network and a single point of administration for all network objects.

See also: directory service

**Active Directory Application Mode (ADAM)**

A stand-alone directory service that is designed specifically for use with directory-enabled applications. Active Directory Application Mode (ADAM) does not require or depend on Active Directory forests or domains. ADAM stores and replicates only application-related information. ADAM does not store or replicate network operating system (NOS)-related information.

See also: directory service; directory-enabled application

**Active Directory Application Mode (ADAM) instance**

For Active Directory Application Mode (ADAM), a single copy of the ADAM directory service, along with its associated directory store, assigned Lightweight Directory Access Protocol (LDAP) and Secure Sockets Layer (SSL) ports, and application event log. You can run multiple ADAM instances simultaneously on a single computer.

See also: Active Directory Application Mode (ADAM); directory service; Lightweight Directory Access Protocol (LDAP); Secure Sockets Layer (SSL)

**ADAM**

See definition for: Active Directory Application Mode (ADAM)

**ADAM instance**

See definition for: Active Directory Application Mode (ADAM) instance

**authentication**

The process for verifying that an entity or object is who or what it claims to be. Examples include confirming the source and integrity of information, such as verifying a digital signature or verifying the identity of a user or computer.

**authentication protocol**

The protocol by which an entity on a network proves its identity to a remote entity. Typically, identity is proved with the use of a secret key, such as a password, or with a stronger key, such as the key on a smart card. Some authentication protocols also implement mechanisms to share keys between client and server to provide message integrity or privacy.

See also: authentication

**B**

**binding**

A process by which software components and layers are linked together. When a network component is installed, the binding

relationships and dependencies for the components are established. Binding allows components to communicate with each other.

**C**

**configuration set**

For Active Directory Application Mode (ADAM), a set of ADAM instances that share and replicate a common schema partition and a common configuration partition.

See also: Active Directory Application Mode (ADAM); Active Directory Application Mode (ADAM) instance

**D**

**Digest Access protocol**

A lightweight authentication protocol for parties involved in communications that are based on Hypertext Transfer Protocol (HTTP) or Simple Authentication and Security Layer (SASL).

See also: Microsoft Digest; Simple Authentication and Security Layer (SASL)

**directory service**

Both the directory information source and the service that makes the information available and usable. A directory service enables the user to find an object when given any one of its attributes.

See also: Active Directory

**Directory Services Markup Language (DSML)**

An open, extensible, standards-based format for publishing directory service schemas and exchanging directory contents.

See also: directory service

**directory-enabled application**

An application that reads, writes, or modifies data that is stored in a directory service.

See also: directory service

**DLL**

See definition for: dynamic-link library (DLL)

**DSML**

See definition for: Directory Services Markup Language (DSML)

**dynamic-link library (DLL)**

An operating system feature that allows executable routines (generally serving a specific function or set of functions) to be stored separately as files with .dll extensions. These routines are loaded only when needed by the program that calls them.

**E**

There are no glossary terms that begin with this letter.

**F**

There are no glossary terms that begin with this letter.

**G**

**group**

A collection of users, computers, contacts, and other groups. Groups can be used as security or as e-mail distribution collections. Distribution groups are used only for e-mail. Security groups are used both to grant access to resources and as e-mail distribution lists.

**group account**

A collection of user accounts. By making a user account a member of a group, you give the related user all the rights and permissions granted to the group.

See also: group

**group memberships**

The groups to which a user account belongs. Permissions and rights granted to a group are also provided to its members. In most cases, the actions a user can perform in Windows are determined by the group memberships of the user account to which the user is logged on.

See also: group

**H**

**Hypertext Transfer Protocol (HTTP)**

The protocol used to transfer information on the World Wide Web. An HTTP address (one kind of Uniform Resource Locator (URL)) takes the following form: http://www.microsoft.com.

**I**

There are no glossary terms that begin with this letter.

**J**

There are no glossary terms that begin with this letter.

**K**

There are no glossary terms that begin with this letter.

**L**

**LDAP**

See definition for: Lightweight Directory Access Protocol (LDAP)

**Lightweight Directory Access Protocol (LDAP)**

The primary access protocol for Active Directory. LDAP is an industry-standard protocol, established by the Internet Engineering Task Force (IETF), that allows users to query and update information in a directory service. Active Directory supports both LDAP version 2 and LDAP version 3.

See also: Active Directory Application Mode (ADAM); directory service

**M**

**Microsoft Digest**

A security support provider (SSP) that implements the Digest Access protocol. Microsoft Digest provides a simple challenge response mechanism for authenticating clients, and it is intended for use by client/server applications that use communications that are based on Hypertext Transfer Protocol (HTTP) or Simple Authentication and Security Layer (SASL).

See also: authentication; Digest Access protocol; Hypertext Transfer Protocol (HTTP); security support provider (SSP); Simple Authentication and Security Layer (SASL)

**N**

**Network Service account**

A predefined local account that is used to start a service and provide the security context for that service. The name of the account is NT AUTHORITY\NetworkService. The Network Service account has limited access to the local computer and authenticated access (as the computer account) to network resources.

**O**

**object identifier**

A number that identifies an object class or attribute. Object identifiers (OIDs) are organized into an industry-wide global hierarchy. An object identifier is represented as a dotted decimal string, such as 1.2.3.4, with each dot representing a new branch in the hierarchy. National registration authorities issue root object identifiers to individuals or organizations, who manage the hierarchy below their root object identifier.

**P**

There are no glossary terms that begin with this letter.

**Q**

There are no glossary terms that begin with this letter.

**R**

**replica**

In Active Directory replication, one instance of a logical Active Directory partition that is synchronized by means of replication between domain controllers that hold copies of the same directory partition. *Replica* can also refer to an instance of an object or attribute in a distributed directory. In the File Replication service (FRS), a computer that has been included in the configuration of a specific replica set.

See also: Active Directory

**S**
**SASL**

See definition for: Simple Authentication and Security Layer (SASL)

**Secure Sockets Layer (SSL)**

A proposed open standard for establishing a secure communications channel to prevent the interception of critical information, such as credit card numbers. Primarily, it enables secure electronic financial transactions on the World Wide Web, although it is designed to work on other Internet services as well.

**security context**

The security attributes or rules that are currently in effect. For example, the rules that govern what a user can do to a protected object are determined by security information in the user's access token and in the object's security descriptor. Together, the access token and the security descriptor form a security context for the user's actions on the object.

**security ID (SID)**

A data structure of variable length that identifies user, group, and computer accounts. Every account on a network is issued a unique SID when the account is first created. Internal processes in Windows refer to an account's SID rather than the account's user or group name.

See also: group account

**security package**

The software implementation of a security protocol. Security packages are contained in security support provider dynamic-link libraries (DLLs) or in security support provider/authentication package DLLs.

See also: security support provider (SSP)

**security principal**

An account holder that is automatically assigned a security identifier (SID) to control access to resources. A security principal can be a user, group, service, or computer.

See also: group; security ID (SID)

**security protocol**

A specification that defines security-related data objects and rules about how the objects are used to maintain security on a computer system.

**security support provider (SSP)**

A dynamic-link library (DLL) that implements the Security Support Provider Interface (SSPI) by making one or more security packages available to applications. Each security package provides mappings between an application's SSPI function calls and an actual security model's functions. Security packages support security protocols such as Kerberos and NTLM.

See also: dynamic-link library; security package; security protocol; Security Support Provider Interface (SSPI)

**Security Support Provider Interface (SSPI)**

A common interface between transport-level applications, such as Microsoft Remote Procedure Call (RPC), and security support providers (SSPs), such as Windows Distributed Security. SSPI allows a transport application to call one of the SSPs to obtain an authenticated connection. These calls do not require extensive knowledge of the security protocol's details.

See also: security protocol; security support provider

**Service Principal Name (SPN)**

A way of referring to a service principal. SPN structures generally follow Internet Engineering Task Force (IETF) naming conventions, and they often include the name of the computer on which the service is running. SPNs may be used to request Kerberos tickets, and they are required for mutual authentication.

See also: authentication

**SID**

See definition for: security ID (SID)

**Simple Authentication and Security Layer (SASL)**

An open framework, described in Request for Comments (RFC) 2222, for adding authentication support to connection-based protocols.

See also: authentication

**SSL**

See definition for: Secure Sockets Layer (SSL)

**SSP**

See definition for: security support provider (SSP)

**SSPI**

See definition for: Security Support Provider Interface (SSPI)

**T**
There are no glossary terms that begin with this letter.

**U**
There are no glossary terms that begin with this letter.

**V**
There are no glossary terms that begin with this letter.

**W**
There are no glossary terms that begin with this letter.

**X**
**X.500**

A set of standards defining a distributed directory service, developed by the International Standards Organization (ISO).

See also: directory service

**Y**
There are no glossary terms that begin with this letter.

**Z**
There are no glossary terms that begin with this letter.

# Storage Manager for SANs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Storage Manager for SANs helps you create and manage logical unit numbers (LUNs) on Fibre Channel and iSCSI disk drive subsystems that support Virtual Disk Service (VDS) in your storage area network (SAN).

A LUN is a logical reference to a portion of a storage subsystem. A LUN can comprise a disk, a section of a disk, a whole disk array, or a section of a disk array in the subsystem. Using LUNs simplifies the management of storage resources in your SAN because they serve as logical identifiers through which you can assign access and control privileges.

- To install and start working with Storage Manager for SANs, see Deploying Storage Manager for SANs.
- For an overview of managing LUNs with Storage Manager for SANs, see Overview of Storage Manager for SANs.
- For a list of steps about how to create a LUN using Storage Manager for SANs, see Checklist: Prepare and Create a LUN.
- For information about the connections between the servers in your SAN and the LUNs that you create using Storage Manager for SANs, see Manage Server Connections.
- To create and manage LUNs, see Manage LUNs.
- To create and manage targets on iSCSI subsystems, see Manage iSCSI Targets.
- For information about Fibre Channel and iSCSI subsystems, see Manage Subsystems.
- For information about disk drives in the Fibre Channel and iSCSI subsystems, see Manage Drives.
- For information about DiskRAID, see DiskRAID.

# Deploying Storage Manager for SANs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

After you install a Windows Server 2003 R2 operating system, use **Add or Remove Programs** in Control Panel to install the Storage Manager for SANs component.

To install Storage Manager for SANs, you must be logged on under an account that has administrative rights on the local computer and have the installation CD for Windows Server 2003 R2 available.

| **Important** |
| --- |
| You may need to restart your server after the installation of Storage Manager for SANs has completed. You should plan accordingly. |

**To install Storage Manager for SANs**

1. In Control Panel, click **Add or Remove Programs**. Then click **Add/Remove Windows Components**.
2. From the list of components, in the **Windows Components Wizard** dialog box, select **Management and Monitoring Tools**, and click **Details**.
3. From the list of subcomponents, select **Storage Manager for SANs**, and then click **OK**.
4. Click **Next**, and after the configuration of the new component has completed, click **Finish**.

**Opening the Storage Manager for SANs snap-in**

You administer the Storage Manager for SANs component through a Microsoft Management Console (MMC) snap-in. To use Storage Manager for SANs, you must be logged on under an account that has administrative rights.

To open Storage Manager for SANs, in Control Panel, double-click **Administrative Tools**, and then double-click **Storage Manager for SANs**.

**See Also**

**Concepts**

Checklist: Prepare and Create a LUN

# Overview of Storage Manager for SANs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can use Storage Manager for SANs to create and manage logical unit numbers (LUNs) on both Fibre Channel and iSCSI disk storage subsystems that support Virtual Disk Service (VDS).

Because of hardware, protocol, and security differences, LUN configuration and management on Fibre Channel and iSCSI environments is different. This section explains those differences, lists the types of LUNs that can be created, and defines LUNs in the context of partitions and volumes.

**Managing LUNs in a Fibre Channel environment**
In a Fibre Channel environment, LUNs created on a disk storage subsystem are assigned directly to a server or cluster, which accesses the LUN through one or more Fibre Channel host bus adapter (HBA) ports. You only need to identify the server or cluster that will access the LUN, and then select which HBA ports on that server or cluster will be used for LUN traffic.

When a server or cluster is identified, Storage Manager for SANs automatically discovers the available Fibre Channel HBA ports on that server or cluster. You can also add ports manually by typing their World Wide Name (WWN).

**Managing LUNs in an iSCSI environment**
Unlike in a Fibre Channel environment, LUNs created on an iSCSI disk storage subsystem are not only assigned to a server or cluster. For iSCSI, LUNs are first assigned to logical entities called targets.

Targets are created in order to manage the connections between an iSCSI device and the servers that need to access it. A target defines the portals (IP addresses) that can be used to connect to the iSCSI device, as well as the security settings (if any) that the iSCSI device requires in order to authenticate the servers that are requesting access to its resources.

| **Note** |
| --- |
| In most cases, you can create and manage targets yourself. However, some iSCSI storage subsystems only support simple target configurations, where targets are automatically created when you create a LUN. With simple target configurations, you also cannot delete a target or manually assign LUNs to it. LUNs are automatically assigned when they are created. For this type of subsystem, you just need to identify the server or cluster that will access the LUN, and the iSCSI subsystem will enable access from that server or cluster to the LUN. |

To connect to a target, a server in the storage area network (SAN) uses an iSCSI initiator. An iSCSI initiator is a logical entity that enables the server to communicate with the target. The iSCSI initiator first logs on to the target, and only after access is granted by the target, the server can start reading and writing to LUNs assigned to that target. Each iSCSI initiator can have one or more network adapters through which communication is established.

As with Fibre Channel environments, you only need to identify the server or cluster that will access the LUN, and Storage Manager for SANs automatically discovers the iSCSI initiators on that server or cluster, and lists all the available adapters for those initiators. After the iSCSI initiator adapters have been discovered, you can select which adapters will be used for LUN traffic.

For more information, see Microsoft Storage Technologies - iSCSI at the Microsoft Web site.

**Types of LUNs**
Storage Manager for SANs supports the following types of LUNs.

| LUN type | Description |
| --- | --- |
| Simple | Simple LUNs use only one physical drive or one portion of a physical drive. This is the most basic type of LUN. |
| Spanned | Spanned LUNS are simple LUNs that span multiple physical drives. |

| | |
|---|---|
| Striped | Striped LUNs write data across multiple physical drives. Data is divided into blocks and spread among all the drives. Since striping writes data to all disks at the same rate, it provides increased I/O performance by distributing I/O requests across disks. Striped LUNs cannot be extended or mirrored, and do not offer fault tolerance. If one of the disks containing a striped LUN fails, the entire LUN fails.<br><br>Select this type of LUN when improved I/O performance is required. |
| Mirrored | Mirrored LUNs are fault-tolerant LUNs that provide data redundancy by creating identical copies of the LUN on two physical drives. All read and write operations happen concurrently on both drives. If one of the physical disks fails, the data on the failed disk becomes unavailable, but the LUN continues to be available using the unaffected disk.<br><br>Select this type of LUN when fault tolerance is required. |
| Striped with parity | Striped LUNs with parity are fault-tolerant LUNs with data and parity spread intermittently across three or more physical disks. If a portion of a physical disk fails, the data that was on the failed portion is recreated from the remaining data and parity information. This type of LUN provides better read performance than a mirrored LUN, but write performance is reduced by the parity calculation.<br><br>Select this type of LUN when fault tolerance is required and improved read performance is desired. |

**LUNs, partitions and volumes**

A LUN is a logical reference to a portion of a storage subsystem. A LUN can comprise a disk, a section of a disk, a whole disk array, or a section of a disk array in the subsystem. This logical reference, when it is assigned to a server in your SAN, acts as a physical disk drive that the server can read and write to. Using LUNs simplifies the management of storage resources in your SAN, because they serve as logical identifiers through which you can assign access and control privileges.

After a LUN has been assigned to a server, you can create one or more partitions on that LUN. Partitions define how much physical space is allocated for storage. For the operating system to start writing and reading data on partitions, you need to create volumes by formatting the partitions using a file system. Volumes define how much logical space is allocated for storage. They can expand over more than one partition.

**See Also**

**Concepts**

Manage Server Connections

# Checklist: Prepare and Create a LUN

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

| Step | Reference |
| --- | --- |
| Review the basic logical unit number (LUN) management concepts. | Overview of Storage Manager for SANs |
| Verify that all the storage subsystems in your storage area network (SAN) have been discovered. | Manage Subsystems |
| Configure the Fibre Channel and iSCSI connections of the servers to which your LUNs will be assigned. | Manage Server Connections |
| If you are going to create a LUN for an iSCSI subsystem, first create, configure and log on to the iSCSI targets in the subsystem. | Manage iSCSI Targets |
| If you are going to enable a server or cluster to access a LUN through multiple Fibre Channel ports or iSCSI initiator adapters, ensure that the server or cluster supports Multipath I/O (MPIO). | Support for Multipath I/O (MPIO) |
| To create a LUN, use the Create LUN Wizard. | Create a LUN |

# Manage Server Connections

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

With Storage Manager for SANs, you can configure the connections that servers in your SAN will use to access the LUNs that you create on your storage subsystems. Before you assign a LUN to a server or to a cluster, you must first configure the connections on each server that will access the LUN.

The configuration of Fibre Channel connections is different from that of iSCSI connections. This section explains those differences, lists the steps to create each type of connection, offers information about enabling more than one connection from a server or cluster to a LUN (Multipath I/O), and explains how to provide storage for server clusters.

**In this section**


- Add and Configure a Server with Fibre Channel Connections
- Add and Configure a Server with iSCSI Connections
- Define a Server Cluster in Storage Manager for SANs
- Support for Multipath I/O (MPIO)

# Add and Configure a Server with Fibre Channel Connections

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Logical unit numbers (LUNs) created on Fibre Channel storage subsystems are assigned directly to servers with Fibre Channel connections on the SAN. To configure the Fibre Channel connections, you only need to identify the server that will access the LUN, and then select which ports on that server will be enabled for LUN access.

Storage Manager for SANs automatically discovers Fibre Channel ports on a server. You can also add ports manually by entering their World Wide Name (WWN).

**To add and configure a server with Fibre Channel connections**

1. In **LUN Management**, click **Manage Server Connections**.
2. In the **Manage Server Connections** dialog box, click **Add**.
3. In the **Add Server** dialog box, type the name or browse to the server that you want to add and optionally, type a description for it.
4. Click **OK**. The server will now be listed in the **Manage Server Connections** dialog box, with all the ports that were automatically discovered on it listed on the **Fibre Channel Ports** tab.
5. To add a Fibre Channel port manually:
   a. On the **Fibre Channel Ports** tab, click **Add** and enter the WWN of the new port.
   b. Optionally, type a description for the new port.
   c. Click **OK**.
6. To enable Fibre Channel ports for LUN access:
   a. Select a server from the server list.
   b. On the **Fibre Channel Ports** tab, select all the ports that you want to enable on the selected server.

| ⚠**Caution** |
| --- |
| If you enable more than one Fibre Channel port and the server is not configured to use Multipath I/O, data corruption may occur. For more information about Multipath I/O, see Support for Multipath I/O (MPIO). |

7. Click **OK** when you are done configuring the Fibre Channel connections.

**See Also**
**Concepts**
Create a LUN
Assign a LUN

# Add and Configure a Server with iSCSI Connections

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Unlike in a Fibre Channel environment, logical unit numbers (LUNs) created on an iSCSI disk storage subsystem are not directly assigned to a server. For iSCSI, LUNs are assigned to logical entities called targets. Servers that require access to a LUN have to connect to the target to which the LUN is assigned. To connect to a target, a server in the storage area network (SAN) uses an iSCSI initiator.

| **Note** |
| --- |
| In most cases, you can create and manage targets yourself. However, some iSCSI storage subsystems only support simple target configurations, where targets are automatically created when you create a LUN. With simple target configurations, you also cannot delete a target or manually assign LUNs to it. LUNs are automatically assigned when they are created. For this type of subsystem, you just need to identify the server or cluster that will access the LUN, and the iSCSI subsystem will enable access from that server or cluster to the LUN. |

To configure the iSCSI connections, you first need to identify the server that will access the LUN and then select which iSCSI initiator adapters will be used for LUN access. iSCSI initiator adapters on a server are automatically discovered by Storage Manager for SANs.

**To add and configure a server with iSCSI connections**

1. In **LUN Management**, click **Manage Server Connections**.
2. In the **Manage Server Connections** dialog box, click **Add**.
3. In the **Add Server** dialog box, type the name or browse to the server that you want to add and optionally, type a description for it.
4. Click **OK**. The server will now be listed in the **Manage Server Connections** dialog box, with all the initiator adapters that were automatically discovered on it listed on the **iSCSI Initiator Adapters** tab.
5. To enable iSCSI initiator adapters for LUN access:
   a. Select a server from the server list.
   b. On the **iSCSI Initiator Adapters** tab, select all the initiator adapters that you want to enable on the selected server.

   | **Caution** |
   | --- |
   | If you enable more than one iSCSI initiator adapter and the server is not configured to use Multipath I/O, data corruption can occur. For more information about Multipath I/O, see Support for Multipath I/O (MPIO). |

6. Click **OK** when you are done configuring the iSCSI connections.

   | **Note** |
   | --- |
   | Optionally, you can change the iSCSI initiator name, by typing a new name for it in the text box provided on the **iSCSI Initiator Adapters** tab. |

**See Also**
**Concepts**
Create a LUN
Assign a LUN
Manage iSCSI Targets

# Define a Server Cluster in Storage Manager for SANs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

With Storage Manager for SANs, you can provide logical unit number (LUN) access to a group of servers (a server cluster) in your storage area network (SAN). Server clusters are created by defining them in Storage Manager for SANs and by enabling clustering services on each individual server in the cluster.

| **Caution** |
| --- |
| To prevent data loss, make sure that each server is a member of only one cluster and that clustering services are running on all cluster nodes. For more information about clustering services, see Clustering Services at the Microsoft Web site. |

After server clusters are created, you can assign LUNs to them in the same way you assign LUNs to individual servers.

**To define a server cluster in Storage Manager for SANs**

1. In **LUN Management**, click **Manage Server Connections**.
2. In the **Manage Server Connections** dialog box, click **Manage Clusters**.
3. In the **Manage Clusters** dialog box, click **Add**.
4. In the **Add New Cluster** dialog box, type a name for the cluster and optionally, type a general description that will help you identify the cluster.
5. From the server list, select all the servers that you want to include in the new cluster.

   | **Note** |
   | --- |
   | To add servers to a cluster you must first add them to the server list in **Manage Server Connections**. For more information see Manage Server Connections. |

6. Click **OK** to create the new cluster.
7. Click **OK** again, to go back to the **Manage Server Connections** dialog box.

The new cluster will now be listed in the **Manage Server Connections** dialog box. To configure the connections of the cluster, configure the Fibre Channel ports or iSCSI initiator adapters of each individual server included in the cluster by selecting the server from the list. All configuration changes you do on individual servers will take effect on the cluster.

| **Caution** |
| --- |
| If you are going to enable more than one connection (Fibre Channel ports and/or iSCSI initiator adapters) on a server in the cluster, and that server is not configured to use Multipath I/O, data corruption can occur. For more information about Multipath I/O, see Support for Multipath I/O (MPIO). |

**See Also**
**Concepts**
Assign a LUN

# Support for Multipath I/O (MPIO)

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Multipathing increases availability by providing multiple paths (path failover) from a server or cluster to a storage subsystem.

If a server in your storage area network (SAN) supports Microsoft Multipath I/O (MPIO), you can enable more than one read/write path to a logical unit number (LUN) by assigning multiple Fibre Channel ports or multiple iSCSI initiator adapters on that server to the same LUN. The same is true when accessing the LUN from a cluster.

| **Caution** |
| --- |
| To prevent data loss, ensure that the server or cluster supports MPIO before enabling access through multiple paths to the same LUN. If the server or cluster does not support MPIO, or if you are not sure, assign only a single Fibre Channel port or iSCSI initiator adapter to each LUN. |

For more information, see Microsoft Storage Technologies - Multipath I/O at the Microsoft Web site.

# Manage LUNs

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Storage Manager for SANs helps you create and manage logical unit numbers (LUNs) on Fibre Channel and iSCSI disk drive subsystems in your storage area network (SAN).

The **LUN Management** node lists all the LUNs created with Storage Manager for SANs and enables you to create new LUNs, extend the size of a LUN, assign and unassign LUNs, and delete LUNs.

**In this section**


Create a LUN

Extend a LUN

Delete a LUN

Assign a LUN

# Create a LUN

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Storage Manager for SANs includes the Create LUN Wizard, which can help you create logical unit numbers (LUNs).

| **Note** |
| --- |
| If you want to use the Create LUN Wizard to assign the LUN that you create to a server or cluster, you must first configure the storage area network (SAN) server connections by using the instructions in Manage Server Connections. |

| **Note** |
| --- |
| If the LUN that you are creating is on an iSCSI subsystem and you want to use the Create LUN Wizard to assign that LUN to an iSCSI target, you must first configure the iSCSI targets in your SAN by using the instructions in Manage iSCSI Targets. |

**To create a LUN**

1. In **LUN Management**, click **Create LUN**.
2. Follow the steps in the **Create LUN Wizard** pages:

| Wizard Page | Details |
| --- | --- |
| **Before You Begin** | Before you create a LUN, verify that the following steps have been completed:<br>• At least one subsystem is listed in the **Subsystems** node.<br>• Storage space is available on at least one of the available subsystems.<br>• If you want to use the wizard to assign the LUN to a server, the SAN server connections have been configured. |
| **Subsystem and Type** | Select the subsystem on which to create the LUN and specify the type of LUN to create:<br>• Fault tolerant LUN types require more disk space than simple LUNs. Therefore, the maximum size of the LUN is smaller than for a simple LUN.<br>• For more information about the different LUN types, see Overview of Storage Manager for SANs. |
| **Size and Name** | Type a name for the new LUN and specify the size of the LUN:<br>• The LUN name can have as many as 63 characters.<br>• Depending on the version of the Virtual Disk Service (VDS) provider that the subsystem is running, you may not be able to type a name for the LUN.<br>• Click the drop-down list to select between megabytes (MB), gigabytes (GB) and terabytes (TB). |
| **Target Access (iSCSI only)** | If iSCSI targets have already been created on the subsystem, you can assign the LUN to one of those targets now. You can also leave it unassigned for later use. |

| | |
|---|---|
| **Server Access** | If you have already configured your server connections, you can assign the LUN to a server or cluster now. You can also leave it unassigned for later use.<br><br>**Note**<br>If you are creating a LUN on an iSCSI subsystem, then this LUN must also be assigned to an iSCSI target. The server or cluster you select on the **Server Access** page will only have access to the LUN after the iSCSI initiator on that server or cluster logs on to the target the LUN has been assigned to. |
| **Create Volume (Fibre Channel only)** | If on the **Server Access** page you chose to assign the LUN to the server you are currently managing with the Storage Manager for SANs snap-in, you can create an NTFS volume for the LUN. If you select to format the volume:<br><br>• As a general rule, leave the **Allocation unit size** option as **Default**. The *allocation unit size* is the smallest amount of disk space that can be allocated to hold a file. Windows picks default sizes based on the size of the volume. These default sizes are selected to reduce the amount of space that is lost and the amount of fragmentation on the volume.<br><br>• The **Quick Format** option speeds the volume creation process on large LUNs. If you want to scan for bad sectors during the formatting process, unselect this option and a full file format will be done.<br><br>**Note**<br>If you prefer, you can select not to format the volume and can use Disk Management to format it later. For more information, see Disk Management at the Microsoft Web site. |
| **Create New LUN** | Before creating the LUN, the wizard lists all the properties you have selected for the new LUN. Click **Create LUN** to start the LUN creation process.<br><br>**Note**<br>Depending on the properties you selected, especially those related to the LUN size and the volume creation options, the LUN creation process may take some time to complete. |

# Extend a LUN

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

If there is storage space available in the subsystems where a logical unit number (LUN) was created, you can extend the size of that LUN.

| **Important** |
| --- |
| Extending a LUN does not automatically extend its file system partition. You must use Disk Manager or a similar tool to extend the file system partition from the server that has access to the LUN. |

**To extend a LUN**

1. In the console tree, click **LUN Management**.
2. In the results pane, select the LUN you want to extend.
3. In the action pane, click **Extend LUN**.
4. In the **Extend LUN** dialog box, type the new size for the LUN and select the units in megabytes (MB), gigabytes (GB) or terabytes (TB).
5. To extend the LUN, click **OK**.

**See Also**
**Concepts**
Manage Subsystems

# Delete a LUN

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

| ⚠Caution |
| --- |
| Deleting a logical unit number (LUN) permanently deletes all data on all volumes on the LUN. This operation cannot be undone. Make sure that you have shut down all applications that access the LUN. |

**To delete a LUN**

1. In the console tree, click **LUN Management**.
2. In the results pane, select the LUN you want to delete.
3. In the action pane, click **Delete LUN**.
4. Select **Check here if you want to delete this LUN**, and then click **OK**.

**See Also**
**Concepts**
Manage LUNs

# Assign a LUN

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

| **Caution** |
| --- |
| If two or more Fibre Channel ports or iSCSI initiator adapters have been enabled on the servers or clusters to which the logical unit number (LUN) will be assigned, and the servers or clusters are not configured to use Multipath I/O, data corruption can occur. For more information about Multipath I/O, see Support for Multipath I/O (MPIO). |

**To assign a LUN**

1. In the console tree, click **LUN Management**.
2. In the results pane, select the LUN you want to assign.
3. In the action pane, click **Assign LUN**.
4. Follow the steps in the **Assign LUN Wizard** pages:

| Wizard Page | Details |
| --- | --- |
| **Target Access (iSCSI only)** | If iSCSI targets have already been created on the subsystem, you can assign the LUN to one of those targets now. You can also leave it unassigned for later use. |
| **Server Access** | If you have already configured your server connections, you can assign the LUN to a server or cluster now. You can also leave it unassigned for later use.<br>**Note**<br>If you are creating a LUN on an iSCSI subsystem, then this LUN must also be assigned to an iSCSI target. The server or cluster you select on the **Server Access** page will only have access to the LUN after the iSCSI initiator on that server or cluster logs on to the target the LUN has been assigned to. |
| **Assign LUN** | Before assigning the LUN, the wizard lists any current assignments for that LUN and confirms the new assignment. To accept the new assignment and close the wizard, click **Assign LUN**.<br>**Note**<br>If a current LUN assignment exists, it will be lost when the new assignment takes place. |

**See Also**
**Concepts**
Manage LUNs

# Manage iSCSI Targets

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Logical unit numbers (LUNs) created on an iSCSI disk storage subsystem are not directly assigned to a server. For iSCSI, LUNs are assigned to logical entities called targets.

Targets are created in order to manage the connections between an iSCSI device and the servers that need to access it. A target defines the portals (IP addresses) that can be used to connect to the iSCSI device, as well as the security settings (if any) that the iSCSI device requires in order to authenticate the servers that are requesting access to its resources.

Servers that require access to a LUN have to connect to the target to which the LUN is assigned. To connect to a target, a server in the storage area network (SAN) uses an iSCSI initiator. An iSCSI initiator is a logical entity that enables the server to communicate with the target. The iSCSI initiator first logs on to the target. The target must grant access before the server can start reading and writing to all LUNs that are assigned to that target.

With Storage Manager for SANs, you can create targets for the iSCSI subsystems in your SAN, configure the security settings for these targets, and start logon sessions with them.

**Note**

In most cases, you can create and manage targets yourself. However, some iSCSI storage subsystems only support simple target configurations, where targets are automatically created when you create a LUN. With simple target configurations, you also cannot delete a target or manually assign LUNs to it. LUNs are automatically assigned when they are created. For this type of subsystem, you just need to identify the server or cluster that will access the LUN, and the iSCSI subsystem will enable access from that server or cluster to the LUN.

**Note**

For more information, see Microsoft Storage Technologies - iSCSI at the Microsoft Web site.

**In this section**


Create an iSCSI Target

Remove an iSCSI Target

Manage iSCSI Security

Log On to iSCSI Targets

# Create an iSCSI Target

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Targets are created in order to manage the connections between the iSCSI storage subsystems in your storage area network (SAN) and the servers that need to access them. You assign logical unit numbers (LUNs) to a target, and all servers that log on to that target will have access to the LUNs assigned to it.

| **Note** |
| --- |
| In most cases, you can create and manage targets yourself. However, some iSCSI storage subsystems only support simple target configurations, where targets are automatically created when you create a LUN. With simple target configurations, you also cannot delete a target or manually assign LUNs to it. LUNs are automatically assigned when they are created. For this type of subsystem, you just need to identify the server or cluster that will access the LUN, and the iSCSI subsystem will enable access from that server or cluster to the LUN. |

After iSCSI targets have been created, you must configure the iSCSI security settings for each of them. For more information about iSCSI security, see Manage iSCSI Security.

**To create an iSCSI target**

1. In **LUN Management**, click **Manage iSCSI targets**.

2. In the **Manage iSCSI Targets** dialog box, in **Select a subsystem to manage**, select the subsystem you want to create the target for, and then click **Add**.

3. In the **Add Target** dialog box, in **Target Friendly Name**, type a descriptive name for the target that will help you to distinguish it from other targets.

4. In **Select portals to enable for this target**, select the check box next to each IP address that the target should use, and then click **OK**.

**See Also**
**Concepts**
Manage iSCSI Security

# Remove an iSCSI Target

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

| ⚠Caution |
| --- |
| When you remove a target, all initiators currently logged on to that target will no longer be able to access the LUNs that are assigned to that target. This action cannot be undone. |

**To remove an iSCSI target**

1. In **LUN Management**, click **Manage iSCSI Targets**.
2. In the **Manage iSCSI Targets** dialog box, in **Select a subsystem to manage**, select the subsystem for which the target you want to remove was created.
3. In **Targets available on this subsystem**, select the target you want to remove, and then click **Remove**.
4. In the **Remove Target** dialog box, select **Check here if you want to remove this target**, and then click **OK** to remove the target.

**See Also**
**Concepts**
Manage iSCSI Targets

# Manage iSCSI Security

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

There are several levels of iSCSI security available with Storage Manager for SANs. The basic level is based on the Challenge Handshake Authentication Protocol (CHAP). CHAP is a protocol that is used to authenticate the peer of a connection and is based upon the peers sharing a *secret* (a security key that is similar to a password). IP security (IPsec) is a protocol that enforces authentication and data encryption at the IP packet layer, which provides an added level of security.

You must choose the security level that best fits the security policies of your organization:

- **One-way CHAP authentication**. With this level of security, only the target authenticates the initiator. The secret is set just for the target and all initiators that want to access that target need to use the same secret to start a logon session with the target.

- **Mutual CHAP authentication**. With this level of security, the target and the initiator authenticate each other. A separate secret is set for each target and for each initiator in the storage area network (SAN).

- **IPsec**. With this level of security, all IP packets sent during data transfers are encrypted and authenticated. A common key is set on all IP portals, allowing all peers to authenticate each other and negotiate packet encryption. For more information, see IPsec at the Microsoft Web site.

| **Caution** |
| --- |
| At a minimum, use one-way CHAP authentication between iSCSI initiators and targets. |

| **Note** |
| --- |
| The level of security that you can set for a storage subsystem depends on the hardware manufacturer. Not all subsystems support all levels of iSCSI security. You should contact your hardware manufacturer to verify what level of security is supported. |

| **Note** |
| --- |
| For more information about iSCSI security and to learn more about iSCSI in general, see Microsoft Storage Technologies - iSCSI at the Microsoft Web site. |

**To manage iSCSI security**

1. In **LUN Management**, click **Manage iSCSI Security**.

2. To configure one-way CHAP authentication, in the **Manage iSCSI Security** dialog box, configure the following settings on the **Targets** tab:

    a. If you want to configure different CHAP secrets for different targets, in the list of targets, select a target that you want to set the CHAP secret for, and click **Set Secret**.

    -Or-

    To use the same CHAP secret for a group of targets, select the targets from the list and click **Set Secret**.

    b. In the **Set Secret** dialog box, type and confirm the target CHAP secret.

    c. Optionally, select **Remember secret on local initiator** if you want to automatically pass the new secret to the local initiator.

d. To set the new secret, click **OK**.

3. To configure mutual CHAP authentication, you must first configure one-way CHAP authentication by following step 2. Then, enter the following configuration on the **Local Initiator** tab:

   a. Type and confirm the CHAP secret for the local initiator.

   b. Under mutual CHAP authentication, the initiator will only be able to log on to targets that know the initiator secret. To share the initiator secret with the targets that the server needs to access, in the list of targets, select each target that you want to authenticate on the initiator.

   c. To set the new secret for the local initiator and to share it with the selected targets, click **Apply Secret**.

4. To configure IPsec, in the **Manage iSCSI Security** dialog box, configure the following settings on the **Portals** tab:

   a. If you want to use different IPsec keys for different portals, in the list of portals, select a portal and click **Set IPsec Key**.

      -Or-

      To use the same IPsec key for a group of portals, select the portals from the list, and click **Set IPsec Key**.

   b. In the **Set IPsec Key** dialog box, type and confirm a new IPsec key.

   c. Optionally, select **Remember the IPsec key on local initiator** if you want to automatically pass the new key to the local initiator.

   d. To set the new IPsec key, click **OK**.

5. When you are done configuring iSCSI security, click **OK**.

**See Also**
**Concepts**
Log On to iSCSI Targets

# Log On to iSCSI Targets

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can log on to iSCSI targets using Storage Manager for SANs. This starts a logon session only between the server that is running Storage Manager for SANs and the selected target. Each server that requires access to the logical unit numbers (LUNs) that are assigned to a target must start a separate logon session with that target.

| **Note** |
| --- |
| If you have not yet configured iSCSI security for initiators and targets, you will not be able to log on to a target. For more information about iSCSI security, see Manage iSCSI Security. |

| **Note** |
| --- |
| You can also start a logon session between a server in your storage area network (SAN) and a target, using the local iSCSI initiator interface, if available (this depends on the iSCSI initiator software that is installed on that server). |

**To log on to iSCSI targets**

1. In **LUN Management**, click **Log On to iSCSI Targets**.
2. In the **Log On to iSCSI Targets** dialog box, select the target that you want to establish a logon session with, and click **Log On**.
3. In the **Log On to Target** dialog box, select the type of Challenge Handshake Authentication Protocol (CHAP) authentication to use for the logon session:
   - **Log on using one-way CHAP authentication**
   - **Log on using mutual CHAP authentication**
4. To establish the logon session, click **OK**.

**See Also**
**Concepts**
Manage iSCSI Security

# Manage Subsystems

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Storage subsystems are automatically discovered using the Virtual Disk Service (VDS). Only those subsystems in your storage area network (SAN) that support VDS can be administered using Storage Manager for SANs.

| **Note** |
| --- |
| For more information about VDS, see Virtual Disk Service Technical Reference. |

The **Subsystems** node lists all the storage subsystems currently discovered in your SAN and allows you to change the subsystem name that was returned by VDS.

**In this section**

Rename a Subsystem

# Rename a Subsystem

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

| Note |
| --- |
| Depending on the version of the Virtual Disk Service (VDS) provider that the storage subsystem is running, you may not be able to change the subsystem name. |

**To rename a subsystem**

1. In the console tree, click **Subsystems**.
2. In the results pane, select the subsystem you want to rename.
3. In the action pane, click **Rename Subsystem**.
4. In the **Rename Subsystem** dialog box, type a new name for the subsystem, and then click **OK**.

**See Also**
**Concepts**
Manage Subsystems

# Manage Drives

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The **Drives** node lists all the drives in the subsystems that are discovered in your storage area network (SAN). To help identify the drives you are working with, you can make the light on a drive blink.

| **Note** |
| --- |
| The ability to make the drive light blink depends on your hardware. |

**In this section**

Make a Drive Light Blink

# Make a Drive Light Blink

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

| **Note** |
| --- |
| The ability to make a drive light blink depends on your hardware. If your storage subsystem does not support this feature, the **Blink Drive Light** action will not be available. |

**To make a drive light blink**

1. In the console tree, click **Drives**.
2. In the results pane, click the drive whose light you want to blink.
3. In the action pane, click **Blink Drive Light**.
4. In the **Blink Drive LED** dialog box, click **Start**.

**See Also**
**Concepts**
Manage Drives

# DiskRAID

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

DiskRAID is a command-line tool that enables you to configure and manage redundant array of independent (or inexpensive) disks (RAID) storage subsystems.

RAID is a method used to standardize and categorize fault-tolerant disk systems. RAID levels provide various mixes of performance, reliability, and cost. RAID is usually used on servers. Some servers provide three of the RAID levels: Level 0 (striping), Level 1 (mirroring), and Level 5 (striping with parity).

A hardware RAID subsystem distinguishes physically addressable storage units from one another by using a Logical Unit Number (LUN). A LUN object must have at least one plex, and can have any number of additional plexes. Each plex contains a copy of the data on the LUN object. Plexes can be added to and removed from a LUN object.

Most DiskRAID commands operate on a specific host bus adapter (HBA) port, initiator adapter, initiator portal, provider, subsystem, controller, port, drive, LUN, target portal, target, or target portal group. You use the SELECT command to select an object. The selected object is said to have focus. Focus simplifies common configuration tasks, such as creating multiple LUNs within the same subsystem.

Windows Server 2003 operating systems include the Virtual Disk Service (VDS), which is a set of application programming interfaces (APIs) that provides a single interface for managing storage. Each hardware vendor writes a VDS hardware provider that translates the general-purpose VDS APIs into specific instructions for their hardware. With this abstraction layer provided by VDS, Windows Server 2003 simplifies storage management, because storage applications no longer need to take into account the specific hardware being targeted. Instead, you can use compatible storage applications, such as DiskRAID, to manage any hardware that has a VDS provider.

**Requirements**

Before you can use DiskRAID to manage your storage subsystems, you must:

- Have a computer running one of the Windows Server 2003 operating systems with Service Pack 1 (SP1) and Windows Server 2003 R2 installed.
- Be a member of the Backup Operators group or Administrators group on the local computer.
- Have at least one registered VDS hardware provider on your computer.

**DiskRAID commands**

To view the command syntax, click a command:

- add
- associate
- automagic
- break
- chap
- create
- delete
- detail
- dissociate

- exit
- extend
- flushcache
- help
- importtarget
- initiator
- invalidatecache
- ipsec
- lbpolicy
- list
- login
- logout
- maintenance
- name
- offline
- online
- recover
- reenumerate
- refresh
- rem
- remove
- replace
- reset
- select
- setflag
- standby
- unmask

**add**

Adds an existing LUN to the currently selected LUN, or adds an iSCSI target portal to the currently selected iSCSI target portal group.

**Syntax**

**add plex lun**=*n* [**noerr**]

**add tpgroup tportal=** *n* [**noerr**]

**Parameters**

**plex lun**=*n*

Specifies the LUN number to add as a plex to the currently selected LUN.

| ⚠**Caution** |
| --- |
| All data on the LUN being added as a plex will be deleted. |

**tpgroup tportal=** *n*

Specifies the iSCSI target portal number to add to the currently selected iSCSI target portal group.

**noerr**

Specifies that any failures that occur while performing this operation will be ignored. This is useful in script mode.

**associate**

Sets the specified list of controller ports as active for the currently selected LUN (other controller ports are made inactive), or adds the specified controller ports to the list of existing active controller ports for the currently selected LUN, or associates the specified iSCSI target for the currently selected LUN.

**Syntax**

**associate controllers** [**add**] *n*[**,***n* [**,**...]]

**associate ports** [**add**] *n*-*m*[**,***n*-*m*[**,**...]]

**associate targets** [**add**] *n*[**,***n* [**,**...]]

**Parameters**

**controllers**

For use with VDS 1.0 providers only. Adds to or replaces the list of controllers that are associated with the currently selected LUN.

**ports**

For use with VDS 1.1 providers only. Adds to or replaces the list of controller ports that are associated with the currently selected LUN.

**targets**

For use with VDS 1.1 providers only. Adds to or replaces the list of iSCSI targets that are associated with the currently selected LUN.

**add**

For VDS 1.0 providers, adds the specified controllers to the existing list of controllers associated with the LUN. If this parameter is not specified, the list of controllers replaces the existing list of controllers associated with this LUN.

For VDS 1.1 providers, adds the specified controller ports to the existing list of controller ports associated with the LUN. If this parameter is not specified, the list of controller ports replaces the existing list of controller ports associated with this LUN.

*n* **[,** *n* **[, ...]]**

For use with the **controllers** or **targets** parameter. Specifies the numbers of the controllers or iSCSI targets to set to active or associate.

*n* **-** *m* **[,** *n* **-** *m* **[,...]]**

For use with the **ports** parameter. Specifies the controller ports to set active using a controller number (*n*) and port number (*m*) pair.

**Example**
The following example shows how to associate and add ports to a LUN that uses a VDS 1.1 provider.

**automagic**
Sets or clears flags that give hints to providers on how to configure a LUN. Used with no parameters, the **automagic** operation displays a list of flags.

**Syntax**
**automagic** {**set** | **clear** | **apply**} **all** *flag*=*value* [*flag*=*value* [...]]

**Parameters**
**set**

Sets the specified flags to the specified values.

**clear**

Clears the specified flags. The **all** keyword clears all the automagic flags.

**apply**

Applies the current flags to the selected LUN.

*flag*

Flags are identified by three-letter acronyms.

| Flag | Description |
|---|---|
| FCR | Fast Crash Recovery Required |
| FTL | Fault Tolerant |
| MSR | Mostly Reads |
| MXD | Maximum Drives 0 |
| MXS | Maximum Size Expected 0B |
| ORA | Optimal Read Alignment0B |
| ORS | Optimal Read Size 0B |
| OSR | Optimize For Sequential Reads- |
| OSW | Optimize For Sequential Writes |
| OWA | Optimal Write Alignment 0B |
| OWS | Optimal Write Size0B |
| RBP | Rebuild Priority 0 |
| RBV | Read Back Verify Enabled |
| RMP | Remap Enabled |
| STS | Stripe Size0 |
| WTC | Write-Through Caching Enabled |
| YNK | Removable |

**break**
Removes the plex from the currently selected LUN. The plex and the data it contained are not retained, and the drive extents

may be reclaimed.

**Syntax**

**break plex**=*plex_number* [**noerr**]

**Parameters**

**plex**

Specifies the number of the plex to remove. The plex and the data it contained will not be retained, and the resources used by this plex will be reclaimed. The data contained on the LUN is not guaranteed to be consistent. If you want to retain this plex, use the Volume Shadow Copy Service (VSS).

**noerr**

Specifies that any failures that occur while performing this operation will be ignored. This is useful in script mode.

**Remarks**

| **Note** |
| --- |
| You must first select a mirrored LUN before using the **break** command. |

| **Caution** |
| --- |
| All data on the plex will be deleted. |

| **Caution** |
| --- |
| All data contained on the original LUN is not guaranteed to be consistent. |

**chap**

Sets the Challenge Handshake Authentication Protocol (CHAP) shared secret so that iSCSI initiators and iSCSI targets can communicate with one another.

**Syntax**

**chap initiator set secret**=[*secret*] [**target**=*target*]

**chap initiator remember secret**=[*secret*] **target**=*target*

**chap target set secret**=[*secret*] [**initiator**=*initiatorname*]

**chap target remember secret**=[*secret*] **initiator**=*initiatorname*

**Parameters**

**initiator set**

Sets the shared secret in the local iSCSI initiator service used for mutual CHAP authentication when the initiator authenticates the target.

**initiator remember**

Communicates the CHAP secret of an iSCSI target to the local iSCSI initiator service so that the initiator service can use the secret in order to authenticate itself to the target during CHAP authentication.

**target set**

Sets the shared secret in the currently selected iSCSI target used for CHAP authentication when the target authenticates the initiator.

**target remember**

Communicates the CHAP secret of an iSCSI initiator to the current in-focus iSCSI target so that the target can use the secret in order to authenticate itself to the initiator during mutual CHAP authentication.

**secret**

Specifies the secret to use. If empty the secret will be cleared.

**target**

Specifies a target in the currently selected subsystem to associate with the secret. This is optional when setting a secret on the initiator and leaving it out indicates that the secret will be used for all targets that do not already have an associated secret.

**initiatorname**

Specifies an initiator iSCSI name to associate with the secret. This is optional when setting a secret on a target and leaving it out indicates that the secret will be used for all initiators that do not already have an associated secret.

#### create

Creates a new LUN or iSCSI target on the currently selected subsystem, or creates a target portal group on the currently selected target. You can view the actual binding using the **DiskRAID list** command.

#### Syntax

**create lun simple** [**size=***n*] [**drives=***n*] [**noerr**]

**create lun stripe** [**size=***n*] [**drives=***n***,** *n* [**,**...]] [**stripesize=***n*] [**noerr**]

**create lun raid** [**size=***n*] [**drives=***n***,** *n* [**,**...]] [**stripesize=***n*] [**noerr**]

**create lun mirror** [**size=***n*] [**drives=***n***,** *n* [**,**...]] [**stripesize=***n*] [**noerr**]

**create lun automagic size=** *n* [**noerr**]

**create target name=** *name* [**iscsiname=***iscsiname*] [**noerr**]

**create tpgroup** [**noerr**]

#### Parameter

**simple**

Creates a simple LUN.

**stripe**

Creates a striped LUN.

**RAID**

Creates a striped LUN with parity.

**mirror**

Creates a mirrored LUN.

**automagic**

Creates a LUN using the *automagic* hints currently in effect. See the **automagic** sub-command for more information.

**size**=

Specifies the total LUN size in megabytes. If the **size=** parameter is not specified, the LUN created will be the largest possible size allowed by all the specified drives.

A provider typically creates a LUN at least as big as the requested size, but the provider may have to round up to the next largest size in some cases. For example, if size is specified as .99 GB and the provider can only allocate GB disk extents, the resulting LUN would be 1 GB.

To specify the size using other units, use one of the following recognized suffixes immediately after the size:

- **B** for byte.
- **KB** for kilobyte.
- **MB** for megabyte.
- **GB** for gigabyte.

- **TB** for terabyte.
- **PB** for petabyte.

**drives**=

Specifies the *drive_number* for the drives to use to create a LUN. If the **size=** parameter is not specified, the LUN created is the largest possible size allowed by all the specified drives. If the **size=** parameter is specified, providers will select drives from the specified drive list to create the LUN. Providers will attempt to use the drives in the order specified when possible.

**stripesize**=

Specifies the size in megabytes for a *stripe* or *RAID* LUN. The stripesize cannot be changed after the LUN is created.

To specify the size using other units, use one of the following recognized suffixes immediately after the size:

- **B** for byte.
- **KB** for kilobyte.
- **MB** for megabyte.
- **GB** for gigabyte.
- **TB** for terabyte.
- **PB** for petabyte.

**target**

Creates a new iSCSI target on the currently selected subsystem.

**name**

Supplies the friendly name for the target.

**iscsiname**

Supplies the iSCSI name for the target and can be omitted to have the provider generate a name.

**tpgroup**

Creates a new iSCSI target portal group on the currently selected target.

**noerr**

Specifies that any failures that occur while performing this operation will be ignored. This is useful in script mode.

**Remarks**

- Either the **size**= or the **drives**= parameter must be specified. They can also be used together.
- The stripe size for a LUN cannot be changed after creation.

**delete**

Deletes the currently selected LUN, iSCSI target (as long as there are not any LUNs associated with the iSCSI target) or iSCSI target portal group.

**Syntax**

**delete lun** [**noerr**]

**delete target** [**noerr**]

**delete tpgroup** [**noerr**]

#### Parameters

**lun**

Deletes the currently selected LUN and all data on it.

**target**

Deletes the currently selected iSCSI target if no LUNs are associated with the target.

**tpgroup**

Deletes the currently selected iSCSI target portal group.

**noerr**

Specifies that any failures that occur while performing this operation will be ignored. This is useful in script mode.

### detail

Displays detailed information about the currently selected object of the specified type.

#### Syntax

**Detail** {**hbaport** | **iadapter** | **iportal** | **provider** | **subsystem** | **controller** | **port** | **drive** | **lun** | **tportal** | **target** | **tpgroup**} [**verbose**]

#### Parameters

**hbaport**

Lists detailed information about the currently selected host bus adapter (HBA) port.

**iadapter**

Lists detailed information about the currently selected iSCSI initiator adapter.

**iportal**

Lists detailed information about the currently selected iSCSI initiator portal.

**provider**

Lists detailed information about the currently selected provider.

**subsystem**

Lists detailed information about the currently selected subsystem.

**controller**

Lists detailed information about the currently selected controller.

**port**

Lists detailed information about the currently selected controller port.

**drive**

Lists detailed information about the currently selected drive, including the occupying LUNs.

**lun**

Lists detailed information about the currently selected LUN, including the contributing drives. The output differs slightly depending on whether the LUN is part of a Fibre Channel or iSCSI subsystem. "If the Unmasked Hosts list contains only an asterisk, this means that the LUN is unmasked to all hosts.

**tportal**

Lists detailed information about the currently selected iSCSI target portal.

**target**

Lists detailed information about the currently selected iSCSI target.

**tpgroup**

Lists detailed information about the currently selected iSCSI target portal group.

**verbose**

For use only with the LUN parameter. Lists additional information, including its plexes.

#### dissociate

Sets specified list of controller ports as inactive for the currently selected LUN (other controller ports are not affected), or dissociates the specified list of iSCSI targets for the currently selected LUN.

#### Syntax

**dissociate controllers** *n*[**,***n*[**,**...]]

**dissociate ports** *n* **-** *m*[**,***n***-***m*[**,**...]]

**dissociate targets** *n*[**,***n*[**,**...]]

#### Parameter

**controllers**

For use with VDS 1.0 providers only. Removes controllers from the list of controllers that are associated with the currently selected LUN.

**ports**

For use with VDS 1.1 providers only. Removes controller ports from the list of controller ports that are associated with the currently selected LUN.

**targets**

For use with VDS 1.1 providers only. Removes targets from the list of iSCSI targets that are associated with the currently selected LUN.

*n*[**,***n*[**,**...]]

For use with the **controllers** or **targets** parameter. Specifies the numbers of the controllers or iSCSI targets to set as inactive or dissociate.

*n* **-** *m*[**,***n***-***m*[**,**...]]

For use with the **ports** parameter. Specifies the controller ports to set as inactive by using a controller number (*n*) and port number (*m*) pair.

#### Example

#### exit

Exits DiskRAID.

#### Syntax

**exit**

#### extend

Extends the currently selected LUN by adding sectors to the end of the LUN. Not all providers support extending LUNs. Does not extend any volumes or file systems contained on the LUN. After you extend the LUN, you should extend the associated on-disk structures using the **DiskPart extend** command.

#### Syntax

**extend lun** [**size**=*LUN_size*] [**drives**=*drive_number*, [*drive_number*, ...]] [**noerr**]

#### Parameters

**size=**

Specifies the size in megabytes to extend the LUN. If the **size=** parameter is not specified, the LUN is extended by the largest possible size allowed by all the specified drives. If the **size=** parameter is specified, providers select drives from the list specified by the **drives=** parameter to create the LUN.

To specify the size using other units, use one of the following recognized suffixes immediately after the size:

- **B** for byte.
- **KB** for kilobyte.
- **MB** for megabyte.
- **GB** for gigabyte.
- **TB** for terabyte
- **PB** for petabyte

**drives=**

Specifies the *drive_number* for the drives to use when creating a LUN. If the **size=** parameter is not specified, the LUN created is the largest possible size allowed by all the specified drives. Providers use the drives in the order specified when possible.

**noerr**

Specifies that any failures that occur while performing this operation should be ignored. This is useful in script mode.

**Remarks**

- Either the *size* or the *drive* parameter must be specified. They can also be used together.

**flushcache**
Clears the cache on the currently selected controller.

**Syntax**
**flushcache controller**

**help**
Displays a list of all DiskRAID commands.

**Syntax**
**help**

**importtarget**
Retrieves or sets the current Volume Shadow Copy Service (VSS) import target that is set for the currently selected subsystem.

**Syntax**
**importtarget subsystem [set target]**

**Parameter**
**set target**

If specified, sets the currently selected target to the VSS import target for the currently selected subsystem. If not specified, the command retrieves the current VSS import target that is set for the currently selected subsystem.

**initiator**
Retrieves information about the local iSCSI initiator.

**Syntax**
**initiator**

**invalidatecache**
Invalidates the cache on the currently selected controller.

**Syntax**
**invalidatecache controller**

**ipsec**

Lists and configures IPsec information about the currently selected iSCSI initiator portal or iSCSI target portal. Also sets the IPsec configuration of all the initiator portals in the local iSCSI initiator service and sets the group pre-shared key used for IPsec in the iSCSI target portals for the currently selected subsystem.

**Syntax**

**ipsec iportal tportal=** *tportal*

**ipsec iportal set key=**[*key*] **tportal=***tportal* [*flag* [*flag* [...]]]

**ipsec iportal set tunnel=** *tunneladdress* **destination=**destinationaddress

**ipsec tportal iportal=** *iportal*

**ipsec tportal set key=**[*key*] {**initiator** | **iportal=***iportal*} [*flag* [*flag* [...]]]

**ipsec tportal set tunnel=** *tunneladdress* **destination=** *destinationaddress*

**ipsec initiator set groupkey=**[*groupkey*]

**ipsec initiator set key=**[*key*] **tportal=***tportal* [*flag* [*flag* [...]]]

**ipsec initiator set tunnel=** *tunneladdress* **destination=** *destinationaddress*

**ipsec subsystem set groupkey=**[*key*]

**Parameters**

**iportal tportal=** *tportal*

Displays the IPsec configuration of the currently selected initiator portal that is associated with the specified target portal on the currently selected subsystem.

**tportal iportal=** *iportal*

Displays the IPsec configuration of the currently selected target portal that is associated with the specified initiator portal on the currently selected initiator adapter.

**iportal set**

Sets the IPsec configuration of the currently selected initiator portal with either the IPsec key and flags, or the IPsec tunnel address.

**tportal set**

Sets the IPsec configuration of the currently selected target portal with either the IPsec key and flags, or the IPsec tunnel address.

**initiator set**

Sets the IPsec configuration of all the initiator portals in the local iSCSI initiator service with the IPsec key and flags, the IPsec tunnel address, or group IPsec key.

**subsystem set**

Sets the group IPsec key of the target portals in the currently selected subsystem.

**key**

Specifies the pre-shared key to be used for IPsec. If not specified, the existing key is cleared.

**tportal**

Specifies the target portal in currently selected subsystem to associate with the key.

**iportal**

Specifies the initiator portal in the currently selected initiator adapter to associate the key with.

**initiator**

Associates the key with all initiator portals in the service.

*flag*

Flags are identified by three letter acronyms:

**IKE**: Enable IPsec/Internet Key Exchange (IKE)

**NMM**: Enable IPsec negotiating through main mode

**NAM**: Enable IPsec negotiating through aggressive mode

**PFS**: Enable perfect forward secrecy

**TRA**: Prefer transport mode

**TUN**: Prefer tunnel mode

**tunneladdress**

Sets the tunnel mode address for the currently selected initiator portal, currently selected target portal, or all initiator portals in the service. Specifies the outer IP address used to create a connection if tunnel mode is used.

**destinationaddress**

Specifies the destination address to associate with the tunnel mode address.

**groupkey**

Specifies the group pre-shared key to be used for IPsec in the local iSCSI initiator service (if **initiator set** is specified) or the currently subsystem (if **subsystem set** is specified). If empty the key will be cleared.

#### lbpolicy

Sets the load balance policy on the currently selected LUN.

#### Syntax

**lbpolicy set lun type=** *type* [**paths=***path***-**{**primary** | *weight*}[**,***path***-**{**primary** | *weight*}[**,**...]]]

**lbpolicy set lun paths=** *path* **-**{**primary** | *weight*}[**,***path***-**{**primary** | *weight*}[**,**...]]

#### Parameters

**type**

Specifies the load balance policy. If the type is not specified, then the **path** parameter must be specified. Type can be one of the following:

**FAILOVER**: Uses one primary path with other paths being backup paths.

**ROUNDROBIN**: Uses all paths in round-robin fashion, which tries each path sequentially.

**SUBSETROUNDROBIN**: Uses all primary paths in round-robin fashion; backup paths are used only if all primary paths fail.

**DYNLQD**: Uses the path with the least number of active requests.

**WEIGHTED**: Uses the path with the least weight (each path must be assigned a weight).

**LEASTBLOCKS**: Uses the path with the least blocks.

**VENDORSPECIFIC**: Uses a vendor-specific policy.

**paths**

Specifies whether a path is **primary** or has a particular *weight*. Any paths not specified are implicitly set as backup. Any paths listed must be one of the currently selected LUN's paths.

#### list

Displays a list of objects of the specified type.

#### Syntax

**List** {**hbaports** | **iadapters** | **iportals** | **providers** | **subsystems** | **controllers** | **ports** | **drives** | **LUNs** | **tportals** | **targets** | **tpgroups**}

#### Parameters

**hbaports**

Lists summary information about all HBA ports known to VDS. The currently selected HBA port is marked by an asterisk (*).

**iadapters**

Lists summary information about all iSCSI initiator adapters known to VDS. The currently selected initiator adapter is marked by an asterisk (*).

**iportals**

Lists summary information about all iSCSI initiator portals in the currently selected initiator adapter. The currently selected initiator portal is marked by an asterisk (*).

**providers**

Lists summary information about each provider known to VDS. The currently selected provider is marked by an asterisk (*).

**subsystems**

Lists summary information about each subsystem in the system. The currently selected subsystem is marked by an asterisk (*).

**controllers**

Lists summary information about each controller in the currently selected subsystem. The currently selected controller is marked by an asterisk (*).

**ports**

Lists summary information about each controller port in the currently selected controller. The currently selected port is marked by an asterisk (*).

**drives**

Lists summary information about each drive in the currently selected subsystem. The currently selected drive is marked by an asterisk (*).

**luns**

Lists summary information about each LUN in the currently selected subsystem. The currently selected LUN is marked by an asterisk (*).

**tportals**

Lists summary information about all iSCSI target portals in the currently selected subsystem. The currently selected target portal is marked by an asterisk (*).

**targets**

Lists summary information about all iSCSI targets in the currently selected subsystem. The currently selected target is marked by an asterisk (*).

**tpgroups**

Lists outputs summary information about all iSCSI target portal groups in the currently selected target. The currently selected portal group is marked by an asterisk (*).

### login

Logs the specified iSCSI initiator adapter into the currently selected iSCSI target.

#### Syntax

**login target iadapter=** *iadapter* [**type=**{**manual** | **persistent** | **boot**}] [**chap=**{**none** | **oneway** | **mutual**}] [**iportal=***iportal*] [**tportal=***tportal*] [*flag* [*flag* [...]]]

#### Parameters

**type**

Specifies the type of login to perform: **manual**, **persistent**, or **boot**. If unspecified, a manual login will be performed.

**manual** - Login manually.

**persistent** - Automatically use the same login when the computer is restarted.

**boot** - *This option is for future development and is not currently used.*

**chap**

Specifies the type of CHAP authentication to use: **none**, **oneway** CHAP, or **mutual** CHAP; if unspecified, no authentication will be used.

**tportal**

Specifies an optional target portal in the currently selected subsystem to use for the log in.

**iportal**

Specifies an optional initiator portal in the specified initiator adapter to use for the log in.

*flag*

Identified by three letter acronyms:

**IPS**: Require IPsec

**EMP**: Enable multipath

**EHD**: Enable header digest

**EDD**: Enable data digest

**logout**
Logs the specified iSCSI initiator adapter out of the currently selected iSCSI target.

**Syntax**
**logout target iadapter=** *iadapter*

**Parameters**
**iadapter**

**Specifies the initiator adapter with a login session to logout from.**

**maintenance**
Performs maintenance operations on the currently selected object of the specified type.

**Syntax**
**maintenance** *object operation* [**count**=*iteration*]

**Parameters**
*object*

Specifies the type of object on which to perform the operation. The *object* type can be a **subsystem**, **controller**, **port, drive** or **LUN**.

*operation*

Specifies the maintenance operation to perform. The *operation* type can be **spinup**, **spindown**, **blink**, **beep** or **ping**. An *operation* must be specified.

**count=**

Specifies the number of times to repeat the *operation*. This is typically used with **blink**, **beep**,or **ping**.

**name**
Sets the friendly name of the currently selected subsystem, LUN, or iSCSI target to the specified name.

**Syntax**
**name** {**subsystem** | **lun** | **target**} [*name*]

**Parameter**

*name*

Specifies a name for the subsystem, LUN, or target. The name must be less than 64 characters in length. If no name is supplied, the existing name, if any, is deleted.

**offline**
Sets the state of the currently selected object of the specified type to **offline**.

**Syntax**
**offline** *object*

**Parameter**
*object*

Specifies the type of object on which to perform this operation. The *object* type can be **subsystem**, **controller**, **drive**, **LUN**, or **tportal**.

**online**
Sets the state of the selected object of the specified type to **online**. If object is **hbaport**, changes the status of the paths to the currently selected HBA port to **online**.

**Syntax**
**online** *object*

**Parameter**
*object*

Specifies the type of object on which to perform this operation. The *object* type can be **hbaport**, **subsystem**, **controller**, **drive**, **LUN**, or **tportal**.

**recover**
Performs operations necessary, such as resynchronization or hot sparing, to repair the currently selected fault-tolerant LUN. For example, RECOVER might cause a hot spare to be bound to a RAID set that has a failed disk or other disk extent reallocation.

**Syntax**
**recover** *lun*

**reenumerate**
Reenumerates objects of the specified type.

**Syntax**
**reenumerate** {**subsystems** | **drives**}

**Parameters**
**subsystems**

Queries the provider to discover any new subsystems that were added in the currently selected provider.

**drives**

Queries the internal I/O buses to discover any new drives that were added in the currently selected subsystem.

**refresh**
Refreshes internal data for the currently selected provider.

**Syntax**
**refresh provider**

**rem**
Used to comment scripts.

**Syntax**
**Rem** *comment*

**remove**
Removes the specified iSCSI target portal from the currently selected target portal group.

**Syntax**

**remove tpgroup tportal=** *tportal* [**noerr**]

**Parameter**
**tpgroup tportal=** *tportal*

Specifies the iSCSI target portal to remove.

**noerr**

Specifies that any failures that occur while performing this operation should be ignored. This is useful in script mode.

**replace**
Replaces the specified drive with the currently selected drive.

**Syntax**
**replace drive**=*drive_number*

**Parameter**
**drive=**

Specifies the *drive_number* for the drive to be replaced.

**Remarks**

- The specified drive may not be the currently selected drive.

**reset**
Resets the currently selected controller or port.

**Syntax**
**Reset** {**controller** | **port**}

**Parameters**
**controller**

Resets the controller.

**port**

Resets the port.

**select**
Displays or changes the currently selected object.

**Syntax**
**Select** {**hbaport** | **iadapter** | **iportal** | **provider** | **subsystem** | **controller** | **port** | **drive** | **lun** | **tportal** | **target** | **tpgroup** } [*n*]

**Parameters**
**object**

Specifies the type of object to select. The *object* type can be **provider**, **subsystem**, **controller**, **drive**, or **LUN**.

**hbaport** [*n*]

Sets the focus to the specified local HBA port. If no HBA port is specified, the command displays the currently selected HBA port (if any). Specifying an invalid HBA port index results in no in-focus HBA port. Selecting an HBA port deselects any selected initiator adapters and initiator portals.

**iadapter** [*n*]

Sets the focus to the specified local iSCSI initiator adapter. If no initiator adapter is specified, the command displays the currently selected initiator adapter (if any). Specifying an invalid initiator adapter index results in no in-focus initiator adapter. Selecting an initiator adapter deselects any selected HBA ports and initiator portals.

**iportal** [*n*]

Sets the focus to the specified local iSCSI initiator portal within the selected iSCSI initiator adapter. If no initiator portal is specified, the command displays the currently selected initiator portal (if any). Specifying an invalid initiator portal index results

in no selected initiator portal.

**provider** [*n*]

Sets the focus to the specified provider. If no provider is specified, the command displays the currently selected provider (if any). Specifying an invalid provider index results in no in-focus provider.

**subsystem** [*n*]

Sets the focus to the specified subsystem. If no subsystem is specified, the command displays the subsystem with focus (if any). Specifying an invalid subsystem index results in no in-focus subsystem. Selecting a subsystem implicitly selects its associated provider.

**controller** [*n*]

Sets the focus to the specified controller within the currently selected subsystem. If no controller is specified, the command displays the currently selected controller (if any). Specifying an invalid controller index results in no in-focus controller. Selecting a controller deselects any selected controller ports, drives, LUNs, target portals, targets, and target portal groups.

**port** [*n*]

Sets the focus to the specified controller port within the currently selected controller. If no port is specified, the command displays the currently selected port (if any). Specifying an invalid port index results in no selected port.

**drive** [*n*]

Sets the focus to the specified drive, or physical spindle, within the currently selected subsystem. If no drive is specified, the command displays the currently selected drive (if any). Specifying an invalid drive index results in no in-focus drive. Selecting a drive deselects any selected controllers, controller ports, LUNs, target portals, targets, and target portal groups.

**lun** [*n*]

Sets the focus to the specified LUN within the currently selected subsystem. If no LUN is specified, the command displays the currently selected LUN (if any). Specifying an invalid LUN index results in no selected LUN. Selecting a LUN deselects any selected controllers, controller ports, drives, target portals, targets, and target portal groups.

**tportal** [*n*]

Sets the focus to the specified iSCSI target portal within the currently selected subsystem. If no target portal is specified, the command displays the currently selected target portal (if any). Specifying an invalid target portal index results in no selected target portal. Selecting a target portal deselects any controllers, controller ports, drives, LUNs, targets, and target portal groups.

**target** [*n*]

Sets the focus to the specified iSCSI target within the currently selected subsystem. If no target is specified, the command displays the currently selected target (if any). Specifying an invalid target index results in no selected target. Selecting a target deselects any controllers, controller ports, drives, LUNs, target portals, and target portal groups.

**tpgroup** [*n*]

Sets the focus to the specified iSCSI target portal group within the currently selected iSCSI target. If no target portal group is specified, the command displays the currently selected target portal group (if any). Specifying an invalid target portal group index results in no in-focus target portal group.

*n*

Specifies the *object number* to select. If the *object number* specified is not valid, any existing selections for objects of the specified type are cleared. If no *object number* is specified, the current object is displayed.

#### setflag

Sets the currently selected drive as a hot spare.

#### Syntax

**setflag drive hotspare**={**true** | **false**}

#### Parameters

**true**

Selects the currently selected drive as a hot spare.

### false

Unselects the currently selected drive as a hot spare.

#### Remarks

- Hot spares cannot be used for ordinary LUN binding operations. They are reserved for fault handling only. The drive must not be currently bound to any existing LUN.

#### standby

Changes the status of the paths to the currently selected host bus adapter (HBA) port to STANDBY.

#### Syntax

**standby hbaport**

#### Parameters

### hbaport

Changes the status of the paths to the currently selected host bus adapter (HBA) port to STANDBY.

#### unmask

Makes the currently selected LUNs accessible from the specified hosts.

#### Syntax

**unmask LUN** {**all** | **none** | [**add**] **wwn**=*hexadecimal_number* [;*hexadecimal_number* [;...]] | [**add**] **initiator=***initiator*[;*initiator*[;...]]} [**uninstall**]

#### Parameters

### all

Specifies that the LUN should be made accessible from all hosts. However, you cannot unmask the LUN to all targets in an iSCSI subsystem.

### none

Specifies that the LUN should not be accessible to any host.

### add

Specifies that the hosts specified must be added to the existing list of hosts that this LUN is accessible from. If this parameter is not specified, the list of hosts supplied replaces the existing list of hosts that this LUN is accessible from.

### WWN=

Specifies a list of hexadecimal numbers representing world-wide names from which the LUN or hosts should be made accessible. To mask/unmask to a specific set of hosts in a Fibre Channel subsystem, you can type a semicolon-separated list of WWN's for the ports on the host machines of interest.

### initiator=

Specifies a list of iSCSI initiators to which the currently selected LUN should be made accessible. To mask/unmask to a specific set of hosts in an iSCSI subsystem, you can type a semicolon-separated list of iSCSI initiator names for the initiators on the host computers of interest.

### uninstall

If specified, uninstalls the disk associated with the LUN on the local system before the LUN is masked.

#### Scripting DiskRAID

DiskRAID may be scripted on any Windows Server 2003 system with an associated VDS hardware provider. To invoke a DiskRAID script, at the command prompt type:

**diskraid /s** *script.txt*

By default, DiskRAID stops processing commands and returns an error code if there is a problem in the script. To continue script execution, include the NOERR parameter on the command. This permits such useful practices as using a single script to delete all the LUNs in a subsystem regardless of the total number of LUNs. Not all commands support the NOERR parameter.

Errors are always returned on command-syntax errors, regardless of whether you included the NOERR parameter,

**DiskRAID error codes:**

| Error Code | Error Description |
|---|---|
| 0 | No error occurred. The entire script ran without failure. |
| 1 | A fatal exception occurred. |
| 2 | The arguments specified on a DiskRAID command line were incorrect. |
| 3 | DiskRAID was unable to open the specified script or output file. |
| 4 | One of the services DiskRAID uses returned a failure. |
| 5 | A command syntax error occurred. The script failed because an object was improperly selected or was invalid for use with that command. |

**Example: Interactively View Status of Subsystem**
If you want to view the status of subsystem 0 on your computer, type the following at the command line:

**diskraid**

Press ENTER. The following is displayed:

```
Microsoft Diskraid version 5.2.xxxx
Copyright (c) 2003 Microsoft Corporation
On computer: COMPUTER_NAME
```

To select subsystem 0, type the following at the DiskRAID prompt:

**select subsystem 0**

Press ENTER. Output similar to the following is displayed:

```
Subsystem 0 is now the selected subsystem.

DISKRAID> list drives

  Drive ###  Status      Health          Size      Free    Bus  Slot  Flags
  ---------  ----------  ------------  --------  --------  ---  ----  -----
  Drive 0    Online      Healthy         107 GB    107 GB    0     1
  Drive 1    Offline     Healthy          29 GB     29 GB    1     0
  Drive 2    Online      Healthy         107 GB    107 GB    0     2
  Drive 3    Not Ready   Healthy          19 GB     19 GB    1     1
```

To exit DiskRAID, type the following at the DiskRAID prompt:

**exit**

Press ENTER. The following output is displayed:

```
Leaving Diskraid
```

# File Server Resource Manager

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

With the increasing demand on storage resources, as organizations rely more heavily on data than ever before, IT administrators face the challenge of overseeing a larger and more complex storage infrastructure, while at the same time, tracking the kind of information available in it. Managing storage resources not only includes data size and availability any more but also the enforcement of company policies and a very good understanding of how existing storage is utilized, allowing for sound strategic planning and proper response to organizational changes.

File Server Resource Manager is a suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers. By using File Server Resource Manager, administrators can place quotas on folders and volumes, actively screen files, and generate comprehensive storage reports. This set of advanced instruments not only helps the administrator to efficiently monitor existing storage resources but it also aids in the planning and implementation of future policy changes.

- Before working with File Server Resource Manager, to configure general notification and reporting options, see Setting File Server Resource Manager Options.
- To create quotas that set a soft or hard space limit on a volume or folder tree, see Quota Management.
- To create file screening rules that block files from a volume or a folder tree, see File Screening Management.
- To generate storage reports that can be used to monitor disk usage patterns, identify duplicate files and dormant files, track quota usage, and audit file screening, see Storage Reports.
- For problem-solving instructions, see Troubleshooting File Server Resource Manager.

# Deploying File Server Resource Manager

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

After you install Windows Server 2003 R2, use **Add or Remove Programs** in Control Panel to install the File Server Resource Manager component.

To install File Server Resource Manager, you must be logged on under an account that has administrative rights on the local computer and have the installation CD for Windows Server 2003 R2 at hand.

| **Important** |
| --- |
| You may need to reboot your server after the installation of File Server Resource Manager has completed. Please plan accordingly. |

**To install File Server Resource Manager**

1. In Control Panel, click **Add or Remove Programs**. Then click**Add/Remove Windows Components**.
2. From the list of components, in the **Windows Components Wizard** dialog box, select **Management and Monitoring Tools**, and click **Details**.
3. From the list of subcomponents, select **File Server Resource Manager**, and then click **OK**.
4. Click **Next**, and after the configuration of the new component has completed, click **Finish**.

**To open the File Server Resource Manager console**

The File Server Resource Manager component is administered through a Microsoft Management Console (MMC) snap-in console.

- On the **Start** menu, click **All Programs**, click **Administrative Tools**, and then click **File Server Resource Manager**.

# Setting File Server Resource Manager Options

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The general File Server Resource Manager options can be set in the **File Server Resource Manager Options** dialog box. These settings will be used throughout the different nodes and some of them can be modified when configuring specific actions.

**This section covers:**

- Configure E-mail Notifications
- Configure Storage Reports
- Configure Report Repositories
- Configure File Screen Audit

# Configure E-mail Notifications

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When you create quotas and file screens, you have the option of sending e-mail notifications to users when their quota limit is approaching or after they have attempted to save files that have been blocked. If you want to routinely notify certain administrators of quota and file screening events, you can configure one or more default recipients.

To send these notifications, you must specify the SMTP server to be used for forwarding the e-mail messages.

**To configure e-mail options**

1. In the console tree, right-click **File Server Resource Manager**, and then click **Configure options**. The **File Server Resource Manager Options** dialog box opens.

2. On the **E-mail Notifications** tab, under **SMTP server name or IP address**, type the host name or the IP address of the SMTP server that will forward e-mail notifications.

3. If you want to routinely notify certain administrators of quota or file screening events, under **Default administrator recipients**, type each e-mail address.

   Use the format *account@domain*. Use semicolons to separate multiple accounts.

4. To test your settings, click **Send Test E-mail**.

# Configure Storage Reports

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can configure the default parameters for storage reports. These default parameters are used for the incident reports generated when a quota or file screening event occurs. They are also used for scheduled and on-demand reports, but the default parameters can be changed when defining the specific properties of these.

| **Important** |
| --- |
| When you change the default parameters for a type of report, changes affect all incident reports. |

**To configure the default parameters for storage reports**

1. In the console tree, right-click **File Server Resource Manager**, and then click **Configure options**. The **File Server Resource Manager Options** dialog box opens.

2. On the **Storage Reports** tab, under **Configure default parameters**, select the type of report that you want to modify.

3. Click **Edit Parameters**.

4. Depending on the type of report that you select, different report parameters will be available for editing. Perform all necessary modifications, and then click **OK** to save them as default parameters for that type of report.

5. Repeat steps 2 through 4 for all types of reports that you want to edit.

**See Also**
**Concepts**
Storage Reports

# Configure Report Repositories

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

You can specify the location where storage reports will be saved on disk. A default path has been defined, but it can be changed if desired. There are three types of reports:

**Incident reports**
This type of report is automatically generated when a quota or file screening event occurs.

**Scheduled reports**
This type of report is generated when there is a schedule defined in the **Storage Reports** node.

**On-demand reports**
This type of report can be generated manually, in the **Storage Reports** node.

**To change the report repositories locations**

1. In the console tree, right-click **File Server Resource Manager**, and then click **Configure options**. The **File Server Resource Manager Options** dialog box opens.

2. On the **Report Locations** tab, type the path or browse to the location where you want the incident, scheduled, and on-demand reports to be saved.

3. Click **OK**.

**See Also**
**Concepts**
Storage Reports

# Configure File Screen Audit

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

By using File Server Resource Manager, you can record file screening activity in an auditing database. The information saved in this database is used to generate the File Screen Audit Report.

| **Important** |
| --- |
| If the **Record file screening activity in the auditing database** check box is cleared, the File Screen Audit Reports will not contain any information. |

**To configure file screen audit**

1. In the console tree, right-click **File Server Resource Manager**, and then click **Configure options**. The **File Server Resource Manager Options** dialog box opens.

2. On the **File Screen Audit** tab, select the **Record file screening activity in the auditing database** check box.

3. Click **OK**. All file screening activity will now be stored in the auditing database and can be viewed by running a File Screen Audit Report.

**See Also**
**Concepts**
Storage Reports

# Quota Management

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the Quota Management node of the File Server Resource Manager Microsoft Management Console (MMC) snap-in, you can perform the following tasks:

- Create quotas to limit the space allowed for a volume or folder and generate notifications when the quota limits are approached or exceeded.
- Generate auto quotas that apply to all existing folders in a volume or folder, as well as to any new subfolders created in the future.
- Define quota templates that can be easily applied to new volumes or folders and that can be used across an organization.

**For Example:**

- You can place a 200 megabytes (MB) limit on the personal folder of each user on a server, with a notification to you and the user when 180 MB of storage has been exceeded.
- A flexible 500 MB quota on a group's shared folder can be set. When this storage limit is reached, all users in the group are notified by e-mail that the storage quota has been temporarily extended to 520 MB, so that they can delete unnecessary files and comply with the preset 500 MB quota policy.
- You can receive a notification when a temporary folder reaches 2 gigabytes (GB) of usage yet not limit that folder's quota because it is necessary for a service running on your server.

**This section covers:**

- Create a quota
- Create a quota template
- Create an auto quota
- Edit quota template properties
- Edit auto quota properties

| **Note** |
| --- |
| To set e-mail notifications and certain reporting capabilities, you must first configure the general File Server Resource Manager options. |

**See Also**
**Concepts**
Setting File Server Resource Manager Options

# Create a quota

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Quotas can be created from a template or individually. The following procedure describes how to create a quota that is based on a template, which we recommend.

| **Note** |
| --- |
| By creating quotas exclusively from templates, you can centrally manage your quotas by updating the templates instead of the individual quotas. Then, you can apply changes to all quotas based on the modified template. This feature simplifies the implementation of storage policy changes by providing one central point where all updates can be made. |

**To create a quota that is based on a template**

1. In **Quota Management**, under **Quota Templates**, select the template on which you will base your new quota. Then, right-click the template, and click **Create quota from template**. This opens the **Create Quota** dialog box.

2. Under **Quota path**, select or type the path to the folder to which the quota will apply. Note that the quota properties will also apply to all of its subfolders.

3. To base the quota on a quota template, select the **Derive properties from this quota template** option, and select the template from the drop-down combo box. Note that each template properties will be displayed under **Summary of quota properties**.

4. Click **Create**.

**See Also**
**Concepts**
Create a quota template

# Create a quota template

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A *quota template* defines a space limit, the type of quota (hard or soft), and a set of notifications to be generated when the quota limit approaches.

| **Note** |
| --- |
| By creating quotas exclusively from templates, you can centrally manage your quotas by updating the templates instead of the individual quotas. This is a feature that simplifies the implementation of storage policy changes by providing one central point where all updates can be made. |

**To create a quota template**

1. In **Quota Management**, right-click **Quota Templates**, and click **Create quota template**. This opens the **Create Quota Template** dialog box.

2. If you want, you can copy the properties of an existing template by selecting it from the drop-down combo box and clicking **Copy**. These properties can be used as a base for the new template, or you can create an entirely new one.

3. To identify the new template, enter a template name and an optional quota label.

4. To configure the space limit for the quota:

   a. Under **Limit**, specify the amount of space to allow.

   b. Specify the type of quota to impose:

      A *hard quota* prevents users from saving files after the space limit is reached, and generates notifications when the volume of data reaches each configured threshold.

      A *soft quota* does not enforce the quota limit, but generates all configured notifications.

   *Notification thresholds* determine what happens when the quota limit approaches. You can send e-mail notifications, log an event, execute a command or a script, or generate reports. By default, no notifications are generated.

5. To configure notifications that will be generated when the quota limit is reached, in the **Notification thresholds** list, highlight the **Limit (100%)** threshold, and click **Edit**. This opens the **Limit Properties** dialog box.

6. To configure e-mail notifications, set the following options on the **E-mail Message** tab:

   - Select the **Send e-mail to the following administrators** check box, and enter the names of the administrative accounts that will receive notifications. Use the format *account@domain*; use semicolons to separate multiple accounts.

   - To contact the person who saved the file that exceeded the quota threshold, select the **Send e-mail to the user who exceeded the threshold** check box.

   - Optionally, edit the default subject line and message body that are provided.

     The text that is in brackets inserts variable information about the quota event that caused the notification. For example, the [Source Io Owner] variable inserts the user name of the user or application that wrote the file to disk. Use the **Insert Variable** button to insert additional variables in the text.

(You can use the **Additional E-mail Headers** button to configure additional headers, including "From," "Cc," "Blind Cc," and "Reply-To.")

7. To log an error to the event log, on the **Event Log** tab, select the **Send warning to event log** check box, and optionally edit the default log entry.

8. To run a command or script, on the **Command** tab, select the **Run this command or script** check box, and type the command or browse to the location where the script is stored. Optionally, enter command arguments, select a working directory for the command or script, or modify the command security setting.

9. To specify that one or more storage reports be generated automatically, on the **Report** tab, select the **Generate reports** check box, and select which reports to generate. You must choose at least one e-mail recipient for the report, either an administrator or the user.

10. After configuring each type of notification that you want to generate, click **OK** to save the threshold.

11. To configure additional notification thresholds, perform these steps:

    a. In the **Notification Thresholds** area, click **Add**.

    b. At the top of the **Add Threshold** dialog box, enter the percentage of the quota limit at which to generate the notifications. The default threshold for the first notification is 85 percent.

    c. Use the **E-Mail Message**, **Event Log**, **Command**, and **Reports** tabs to configure each type of notification that you want to generate, as you did in steps 6 through 10. Then click **OK**.

12. When you complete the quota properties, click **OK** to save the template.

| **Important** |
| --- |
| To send e-mail notifications and to configure the storage reports with parameters appropriate for your server environment, you must first set the general File Server Resource Manager options. |

**See Also**
**Concepts**
Setting File Server Resource Manager Options
Edit quota template properties

# Create an auto quota

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

By using auto quotas, you can assign a quota template to a folder and specify that quotas based on that template be automatically generated and applied to the existing subfolders, as well as to any new subfolders created in the future.

For example, you may define an auto quota for folders that are created on demand, for roaming profile users or for new users. Every time a new folder is created, a new quota entry is automatically generated for that folder, using the template defined in the auto quota profile of its parent folder. These automatically generated quota entries can then be viewed in the results pane, under the **Quotas** node, as individually created quotas, which can be edited and modified separately.

**To create an auto quota**

1. In **Quota Management** select the **Quotas** node.
2. Right-click **Quotas** and click **Create quota** or select **Create quota** from the actions pane. This opens the **Create Quota** dialog box.
3. Under **quota path** select or type the path to the folder to which the auto quota will apply. The auto quota will be applied to all subfolders (current and future) in this folder. Once the path is entered, select the **Auto apply template and create quotas on existing and new subfolders option.**
4. Under **Derive properties from this quota template** select the quota template that you wish to apply. You can review the properties of each quota template in the summary list box.
5. Click **Create**.

| **Note** |
| --- |
| You can verify all automatically generated quotas by opening the **Quotas** node. An individual quota for each subfolder will be listed. |

**See Also**
**Concepts**
Edit auto quota properties

# Edit quota template properties

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When making changes to a quota template, you have the option of extending those changes to quotas that were created using the original quota template. You can choose to modify only those quotas that still match the original template or all quotas that were derived from the original template, regardless of any modifications made to the quotas since they were created. This feature simplifies the process of updating the properties of quotas by providing one central point where all changes can be made.

| **Note** |
| --- |
| If you choose to apply the changes to all quotas that were derived from the original template, you will overwrite any custom quota properties that you may have created. |

**To edit quota template properties**

1. In **Quota Templates**, select the template that you want to modify.
2. Right-click the quota template and click **Edit template properties**, or select **Edit template properties** from the actions pane, under **Selected Quota Templates**. This opens the **Quota Template Properties** dialog box.
3. Perform all necessary changes. Optionally, you can copy the properties of a different template and modify them for this one.
4. When you are finished editing the template properties click **OK**. This will open the **Update Quotas Derived from Template** dialog box.
5. Select the type of update you want to apply:
   - If you have quotas that have been modified since they were created using the original template, and you do not want to change them, select **Apply template only to derived quotas that match the original template**. This option will update only those quotas that have not been edited and maintain the original template properties.
   - If you want to modify all existing quotas that were created using the original template, select **Apply template to all derived quotas**.
   - If you want to keep the existing quotas unchanged, select **Do not apply template to derived quotas**.
6. Click **OK**.

**See Also**
**Concepts**
Create a quota template

# Edit auto quota properties

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When making changes to an auto quota, you have the option of extending those changes to existing quotas in the auto quota path. You can choose to modify only those quotas that still match the original auto quota, or all quotas in the auto quota path, regardless of any modifications made to the quotas since they were created. This feature simplifies the process of updating the properties of quotas that were derived from an auto quota by providing one central point where all changes can be made.

| **Note** |
| --- |
| If you choose to apply the changes to all quotas in the auto quota path, you will overwrite any custom quota properties that you may have created. |

**To edit an auto quota**

1. In **Quotas**, select the quota that you want to modify.
2. Right-click the Quota entry and click **Edit quota properties** or select **Edit quota properties** from the actions pane, under **Selected Quotas**. This opens the **Edit Auto Apply Quota** dialog box.
3. Under **Derive properties from this quota template** select the new quota template that you want to apply. You can review the properties of each quota template in the summary list box.
4. Click **OK**. This will open the **Update Quotas Derived from Auto Apply Quota** dialog box.
5. Select the type of update you want to apply:
   - If you have quotas that have been modified since they were automatically generated, and you do not want to change them select **Apply auto apply quota only to derived quotas that match the original auto apply quota**. This option will update only those quotas in the Auto Apply Quota Path that have not been edited and maintain the original auto apply quota properties.
   - If you want to modify all existing quotas in the Auto Apply Quota Path, select **Apply auto apply quota to all derived quotas**.
   - If you want to keep the existing quotas unchanged, but make the modified auto quota effective for new folders in the Auto Apply Quota Path, select **Do not apply auto apply quota to derived quotas**.
6. Click **OK**.

**See Also**
**Concepts**
Create an auto quota

# File Screening Management

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the File Screening Management node of the File Server Resource Manager MMC snap-in, you can perform the following tasks:

- Create file screens to control the types of files that users can save and to send notifications when users attempt to save blocked files.
- Define file screening templates that can be easily applied to new volumes or folders and that can be used across an organization.
- Create file screening exceptions that extend the flexibility of the file screening rules.

**For example:**

- You could ensure that no music files are stored on personal folders on a server, yet you could allow storage of specific types of media files that support legal rights management or comply with company policies. In the same scenario, you may want to give a vice president in the company special privileges to store any type of files in his personal folder.
- A screening process could be implemented to notify you by e-mail when an executable file is stored on a shared folder, with information including the user that stored the file and the exact location of the file, so that the appropriate precautionary steps can be taken.

**This section covers:**

- Define file groups for screening
- Create a file screen
- Create a file screen exception
- Create a file screen template
- Edit file screen template properties

| **Note** |
| --- |
| To set e-mail notifications and certain reporting capabilities, you must first configure the general File Server Resource Manager options. |

**See Also**
**Concepts**
Setting File Server Resource Manager Options

# Define file groups for screening

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A *file group* is used to define a namespace for a file screen, file screen exception, or storage report. It consists of a set of file name patterns, which are grouped into files to include and files to exclude:

- **Files to include:** files that belong in the group.
- **Files to exclude:** files that do not belong in the group.

| **Note** |
| --- |
| For convenience, you can maintain file groups while you edit the properties of file screens, file screen exceptions, and file screen templates. Any file group changes that you make from these property sheets are not limited to the current item that you are working on. |

**To create a file group**

1. In **File Screen Management**, right-click **File Groups**, and then click **Create file group**.

   -Or-

   While you edit the properties of a file screen, file screen exception, or file screen template, under **Manage file groups**, click **Create**. This opens the **Create File Group Properties** dialog box.

2. Type a name for the file group.

3. Add files to include and files to exclude:

   - For each set of files that you want to include in the file group, in the **Files to include** box, enter a file name pattern, and click **Add**.

     Standard wildcard rules apply. For example, ***.exe** selects all executable files.

   - For each set of files that you want to exclude from the file group, in **Files to exclude**, add exclusions in the same way.

4. Click **OK**.

# Create a file screen

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When creating a new file screen, you can choose to save a file screen template based on the custom file screen properties that you define. The advantage of this is that a link is maintained between file screens and the template used to create them, so that in the future, changes to the template can be applied to all file screens that derive from it. This is a feature that simplifies the implementation of storage policy changes by providing one central point where all updates can be made.

**To create a file screen**

1. In **File Screening Management**, right-click **File Screens**, and then click **Create file screen**. This opens the **Create File Screen** dialog box.

2. Under **File screen path**, type or browse to the folder to which the file screen will apply. The file screen will apply to all subfolders of the selected folder.

3. Under **How do you want to configure file screen properties**, select the **Define custom file screen properties** option, and then click **Custom Properties**. This opens the **File Screen Properties** dialog box.

4. If you want to copy properties from an existing template, select the template that you want to use, and click **Copy**. You can edit the properties for the new file screen.

5. Under **Screening type**, select the type of screening to apply:

   - **Active screening** prevents users from saving files that are members of blocked file groups, and generates notifications when users try to save blocked files. (The file screen does not prevent users and applications from accessing files that were saved to the path before the file screen was created, regardless of whether the files are members of blocked file groups.)

   - **Passive screening** sends notifications, but does not prevent users from saving blocked files.

   To specify which files to screen, under **File Groups**, select each file group that you want to include. (To select the check box for the file group, click the file group label twice.)

   If you want to view the members and nonmembers of a file group, highlight the file group, and click **Edit**.

6. To configure e-mail notifications for the file screen, set the following options on the **E-mail Message** tab:

   - To notify administrators when a user or application attempts to save an unauthorized file, select **Send e-mail to the following administrators**, and enter the administrative accounts that will receive notifications. Use the format *account@domain*; use semicolons to separate multiple accounts.

   - To send a notice to the user who attempted to save an unauthorized file, select **Send e-mail to the user who attempted to save an unauthorized file**.

   - Optionally, edit the default subject line and message body.
     The text that is in brackets inserts variable information about the file screen event that caused the notification. For example, the [Source I/O Owner] variable inserts the user name of the user or application that attempted to write the file to disk. To insert additional variables in the text, click **Insert Variable**.

7. If you want to log an error to the event log when a user tries to save a blocked file, on the **Event Log** tab, select the **Send warning to event log** check box, and, optionally, edit the default log entry.

In addition, you can set options on the **Command** tab, to run a command or script when a user attempts a file screen violation, and set options on the **Report** tab, to specify one or more storage reports to be generated automatically.

8. When all file screen properties are complete, click **Create** to save the file screen.

   You will be asked if you want to save a file screen template based on the custom file screen properties you just defined. If you will use the current settings in other file screens, it is recommended that you save a template. The template will be applied to the new file screen. This will enable you to later update the file screen automatically when you update the template.

9. To save a template when you create the file screen, select **Save the custom properties as a template**, enter a name for the template, and click **OK**.

   -Or-

   If you do not want to save a template, select **Save the custom file screen without creating a template**, and click **OK**.

**See Also**
**Concepts**
Define file groups for screening
Edit file screen template properties

# Create a file screen exception

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Occasionally, you will need to be able to allow exceptions to file screening. For example, you might want to block video files from a file server, but you need to allow your training group to save the video files for their computer-based training. To allow files that other file screens are blocking, create a *file screen exception*.

A file screen exception is a special type of file screen that overrides any file screening that would otherwise apply to a folder, and all its subfolders, in a designated exception path. That is, it creates an exception to any rules derived from a parent folder.

| **Note** |
| --- |
| You cannot create a file screen exception on a folder for which you already have a file screen. You must assign the exception to a subfolder or make changes to the existing file screen. |

To determine which file types will be allowed with the file screen exception, file groups are assigned.

**To create a file screen exception**

1. In**File Screening Management**, right-click **File Screens**, and then click **Create file screen exception**. This opens the **Create File Screen Exception** dialog box.

2. Under **Exception path**, type or select the path to which the exception will apply. The exception will apply to the folder and all of its subfolders.

3. To determine which files to exclude from file screening, in **File groups**, select each file group that you want to include in the file screen exception.

   If you want to view the members and non-members of a file group, highlight the file group, and click **Edit**.

4. Click **OK**.

**See Also**
**Concepts**
Define file groups for screening

# Create a file screen template

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

A *file screen template* defines a group of file groups to block, the type of screening to perform (active or passive) and a set of notifications to be generated.

| **Note** |
| --- |
| By creating file screens exclusively from templates, you can centrally manage your file screens by updating the templates instead of the individual file screens. When you make changes to a template, you can choose to apply those changes to all file screens that are based on that template or only to those file screens whose properties match those in the template. This feature simplifies the implementation of storage policy changes, by providing one central point where all updates can be made. |

**To create a file screen template**

1. In **File Screening Management**, right-click **File Screen Templates**, and then click **Create file screen template**. This opens the **Create File Screen Template** dialog box.

2. To identify the new template, enter a template name.

3. Specify the screening type:

   An *Active Screening* will completely block users from saving unauthorized files.

   A *Passive Screening* will allow users to save unauthorized files, but will generate the notifications defined for this template.

4. Select the file groups to block. You can also choose to create new file groups or edit existing ones.

5. To configure e-mail notifications, set the following options on the **E-mail Message** tab:

   - Select the **Send e-mail to the following administrators** check box, and type the names of the administrative accounts that will receive notifications. Use the format *account@domain;* use semicolons to separate multiple accounts.

   - To contact the person who saved the file, select the **Send e-mail to the user who attempted to save an unauthorized file** check box.

   - Optionally, edit the default subject line and message body that are provided.

     The text that is in brackets inserts variable information about the quota event that caused the notification. For example, the [Source Io Owner] variable inserts the user name of the user or application that wrote the file to disk. To insert additional variables in the text, click **Insert Variable**.

     (To configure additional headers, click **Additional E-mail Headers**)

6. To log an error to the event log, on the **Event Log** tab, select the **Send warning to event log** check box, and optionally edit the default log entry.

7. To run a command or script when a screening event occurs, on the **Command** tab, select the **Run this command or script** check box, and type the command or browse to the location where the script is stored. Optionally, enter command arguments, select a working directory for the command or script, or modify the command security setting.

8. To specify that one or more storage reports be generated automatically, on the **Report** tab, select the **Generate reports** check box, and select which reports to generate. You must choose at least one e-mail recipient for the report, either an administrator or the user.

9. Click **OK**.

| **Important** |
| --- |
| To send e-mail notifications and to configure the storage reports with parameters appropriate for your server environment, you must first set the general File Server Resource Manager options. |

**See Also**
**Concepts**
Setting File Server Resource Manager Options
Edit file screen template properties

# Edit file screen template properties

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

When making changes to a file screen template, you have the option of extending those changes to file screens that were created using the original file screen template. You can choose to modify only those file screens that still match the original template, or all file screens that derive from the original template, regardless of any modifications done to them since they were created. This feature simplifies the process of updating the properties of file screens, by providing one central point where all changes can be made.

| **Note** |
| --- |
| If you choose to apply the changes to all file screens that derive from the original template, you will overwrite any custom file screen properties that you may have created. |

**To edit file screen template properties**

1. In **File Screen Templates**, select the template that you want to modify.
2. Right-click the file screen template and click **Edit template properties**, or select **Edit template properties** from the actions pane, under **Selected File Screen Templates**. This opens the **File Screen Template Properties** dialog box.
3. Perform all necessary changes. Optionally, you can copy the properties of a different template and modify them for this one.
4. Once you are done editing the template properties, click **OK**. This will open the **Update File Screens Derived from Template** dialog box.
5. Select the type of update you want to apply:
   - If you have file screens that have been modified since they were created using the original template, and you do not want to change them, select **Apply template only to derived file screens that match the original template**. This option will update only those file screens that have not been edited and maintain the original template properties.
   - If you want to modify all existing file screens that were created using the original template, select **Apply template to all derived file screens**.
   - If you want to keep the existing file screens unchanged, select **Do not apply template to derived file screens**.
6. Click **OK**.

**See Also**
**Concepts**
Create a file screen template

# Storage Reports

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On the Storage Reports node of the File Server Resource Manager MMC snap-in, you can perform the following tasks:

- Schedule periodic storage reports that allow you to identify trends in disk usage.
- Monitor attempts to save unauthorized files for all users or a selected group of users.
- Generate storage reports instantly.

**For example:**

- You can schedule a report that will run every Sunday at midnight, generating a list that includes the most recently used files of the previous two days. With this information, you can monitor weekend storage activity and plan server downtime that will have less impact on those users that connect from home over the weekend.
- A report can be run, at any given time, to identify all duplicate files in a volume on a server, so disk space can be quickly reclaimed without losing any data.
- You can run a *File by File Group* report, to identify how storage resources are segmented across different file groups; or a *File by Owner* report, to analyze how individual users are using shared storage resources.

**This section covers:**

- Schedule a set of reports
- Generate reports on demand

| **Note** |
| --- |
| To set e-mail notifications and certain reporting capabilities, you must first configure the general File Server Resource Manager options. |

**See Also**
**Concepts**
Setting File Server Resource Manager Options

# Schedule a set of reports

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

To generate a set of reports on a regular schedule, you schedule a *report task*. The report task specifies which reports to generate and what parameters to use; which volumes and folders to report on; how often to generate the reports; and which file formats to save them in.

| **Note** |
| --- |
| To minimize the impact of report processing on performance, generate multiple reports on the same schedule so that the data is only gathered once. To quickly add reports to existing report tasks, you can use the **Add/remove reports from report task** action. That action allows you to add or remove reports from multiple report tasks, and to edit the report parameters. To change schedules or delivery addresses, you must edit individual report tasks. |

**To schedule a report task**

1. Right-click the **Scheduled Report Tasks** node, and click **Schedule a new report task**. This opens the **Storage Reports Task Properties** dialog box.

2. To select a volume or folder on which to report:

   a. In **Scope**, click **Add**.

   b. Browse to the volume or folder that you want to report on, and click **OK**.

   This adds the selected volume or folder to the list. Add as many paths as you want to include.

3. In **Report data**, select each report that you want to generate. (By default, all reports are generated for a scheduled report task.)

4. To edit the parameters of a report:

   a. Highlight the report label, and click **Edit Parameters**.

   b. In the **Report Parameters** dialog box, edit the parameters as needed, and click **OK**.

5. If you want to deliver copies of the reports to administrators by e-mail, on the **Delivery** tab, click **E-mail report to administrators**, and enter each e-mail account. Use the format *account@domain*; use semicolons to separate multiple accounts.

6. To schedule the reports, on the **Schedule** tab, click **Create Schedule**. Then, in the **Schedule** dialog box, click **New**.

   This opens the **Schedule** dialog box, with a default schedule set for 9:00 A.M. daily, starting on the next day.

7. To change the default schedule:

   a. Under **Schedule Task**, select the interval at which to generate the report.

      You can schedule daily, weekly, or monthly reports, or generate the reports only once.

   b. In the options below this field, enter additional information, as needed, to determine when the reports will be generated.

      These options change based on the interval that you choose. For example, for a weekly report, you can specify how

many weeks between reports and which days of the week to generate reports on.

c. Under **Start time**, enter the time of day when you want to generate the report.

d. To save the schedule, click **OK**.

8. To save the report task, click **OK**.

The report task is added to the results pane for **Scheduled Report Tasks**. Tasks are identified by the reports to be generated, the namespace to be reported on, and report schedule.

In addition, you can view the current status of the report (whether or not the report is running), the last run time and the result of that run, and the next scheduled run time.

# Generate reports on demand

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

During daily operations, you may want to generate reports on demand to analyze different aspects of current disk usage on the server. Use the **Generate reports now** action to generate one or more reports. Before the reports are generated, current data is gathered.

When you generate reports on demand, the reports are saved in the reports repository, but no report task is created for later use. You can optionally view the reports immediately when they are complete.

| **Note** |
| --- |
| If you choose to open the reports immediately, you must wait while the reports are generated. Processing time varies, depending on the types of reports and the scope of the data. |

**To generate reports immediately**

1. Right-click the **Scheduled Report Tasks** node, and click **Generate reports now**. This opens the **Storage Reports Task Properties** dialog box.

2. In **Scope**, add each volume or folder that you want to report on.

3. In **Report data**, select and format each report that you want to include:

   a. To add a report to the task, select the check box for the report.

   b. To edit the parameters of a report, highlight the report label and click **Edit Parameters**. When you finish editing the parameters, click **OK**.

4. In **Report format**, select each file format that you want to save the report in.

5. To deliver the reports by e-mail, on the **Delivery** tab, select the **E-mail report to the following administrators** check box. Then, enter each e-mail address that you want the report delivered to. Use the format *account@domain*; use semicolons to separate multiple accounts.

6. To gather the data and generate the reports, click **OK**.

7. In the dialog box that opens, choose whether to open the reports when they are complete or generate the reports in the background, in which case the reports are not displayed.

   Both options save the reports in the formats that you selected. If you choose to open the reports after they are generated, each report opens in its own Internet Explorer window.

# Troubleshooting File Server Resource Manager

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section lists a few common issues you may encounter when using File Server Resource Manager.

**What problem are you having?**

- I am not receiving any e-mail notifications
- I am only receiving one e-mail notification, even though the event that caused it happened several times in a row
- My storage reports keep failing and no or little information is available in the Event Log regarding the source of the failure
- My File Audit reports do not contain any information
- The "Used" and "Available" values for some of the quotas I have created do not correspond to the actual "Limit" setting

| **Note** |
| --- |
| A good troubleshooting option is to look for event logs generated by File Server Resource Manager. All event log entries for File Server Resource Manager can be found in the Application event log under the source: **SRMSVC**. |

**I am not receiving any e-mail notifications.**
**Cause:**

The e-mail options have not been configured or have been configured incorrectly.

**Solution:**

Verify that the SMTP server and e-mail settings in the File Server Resource Manager options have been specified and are valid. Send a test e-mail to confirm that the information is correct and that the SMTP server used to send the notifications is working properly. For more information see Configure E-mail Notifications.

**I am only receiving one e-mail notification, even though the event that triggered that notification happened several times in a row.**
**Cause:**

When a user attempts several times to save a blocked file, or a file that exceeds a quota limit, and there is an e-mail notification configured for that file screening or quota event, then only one e-mail will be sent to the administrator for a period of 60 minutes. This prevents over-flooding of messages in the administrator's e-mail account.

**Solution:**

This is a feature of File Server Resource Manager.

**My storage reports keep failing and no or little information is available in the Event Log regarding the source of the failure.**
**Cause:**

The volume where the reports are being saved may be corrupted.

**Solution:**

Run **chkdsk** on the volume and try generating reports again.

**My File Audit reports do not contain any information.**
**Cause:**

The auditing database is empty (that is, no file screening activity has been reported) or the File Screen Audit properties have not been configured.

**Solution:**

In the **File Server Resource Manager Options** dialog box, on the **File Screen Audit** tab, verify that the **Record file screening activity in the auditing database**check box is selected. For more information, see Configure File Screen Audit.

**The "Used" and "Available" values for some of the quotas I have created do not correspond to the actual "Limit" setting.**
**Cause:**

You may have a nested quota, where the quota that applies to a folder will be a more restrictive quota that derives from one of its parent folders.

**Solution:**

Under **Quota Management**, go to **Quota**, and then select the quota entry that you are troubleshooting. Click **View quotas affecting folder** on the actions pane, and look for quotas applied to parent folders. This will allow you to identify which quotas have restrictive settings on the quota you have selected (that is, which parent folder quotas have a lower storage limit setting).

# Identity Management for UNIX: Welcome

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Identity Management for UNIX**
Identity Management for UNIX makes it easy to integrate users of Windows operating systems into existing UNIX environments. It provides manageability components that simplify network administration and account management across both platforms.

With Identity Management for UNIX, you can:

- Manage user accounts and passwords on Windows and UNIX systems using Network Information Service (NIS).
- Automatically synchronize passwords between Windows and UNIX operating systems.

For details about Identity Management for UNIX, see Identity Management for UNIX overview. For information about installable components, see Identity Management for UNIX components.

**In this section**


- Identity Management for UNIX overview
- Identity Management for UNIX components
- Connect to a computer you want to manage
- Identity Management for UNIX: Server for NIS
- Identity Management for UNIX: Password Synchronization
- Accessibility Features of Identity Management for UNIX

# Identity Management for UNIX overview

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Identity Management for UNIX overview**

Identity Management for UNIX makes it easy to integrate computers running Windows into your existing UNIX enterprise. Active Directory network administrators can use Server for NIS to manage Network Information Service (NIS) domains, and Password Synchronization automatically synchronizes passwords between Windows and UNIX operating systems.

With minor differences, Identity Management for UNIX is compliant with Internet Engineering Task Force (IETF) standard Request for Comments (RFC) 2307, meaning that your network's password and NIS attributes can be resolved by the Lightweight Directory Access Protocol (LDAP). For more information about RFC 2307, including a complete list of attributes specified in the RFC 2307 schema, see the RFC database Web site.

Identity Management for UNIX: Server for NIS completely integrates UNIX NIS networks with Windows Active Directory. Identity Management for UNIX includes an easy-to-use wizard that a Windows domain administrator can use to export NIS domain maps to Active Directory entries. Once this is done, an Active Directory domain controller running Server for NIS becomes the master server for the NIS domain.

Identity Management for UNIX: Password Synchronization automatically changes a user password on the UNIX network when the user changes his or her Windows password, and vice versa. This allows users to maintain just one password for both networks.

Password Synchronization supports Sun Solaris version 8 running on x86-based computers and Scalable Processor Architecture (SPARC)–based computers; Solaris version 9 running on SPARC–based computers; Hewlett Packard HP-UX version 11i; IBM AIX version 5L 5.2; and Red Hat Linux versions 8 and 9 running on x86-based computers and 64-bit AMD-based computers. Server for NIS should work with any operating system or product that uses LDAP.

# Identity Management for UNIX components

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Identity Management for UNIX components**

You can use the Identity Management for UNIX installation program to install the following:

- **Administration components**. Provides utilities for administering Identity Management for UNIX.
- **Password Synchronization**. Password Synchronization makes it easy for users to maintain one password for Windows domains and UNIX systems, synchronizing them when one of them changes.
- **Server for NIS**. Enables a Windows-based server to function as a Network Information Service (NIS) server, integrating NIS domains with Active Directory and giving Windows Active Directory administrators the ability to manage Windows and NIS domains together. Server for NIS can be installed only on an Active Directory domain controller.

# Connect to a computer you want to manage

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To connect to the computer you want to manage**

1. Open Identity Management for UNIX.
2. Do one of the following:
   - Right-click the **Microsoft Identity Management for UNIX** node at the top of the hierarchy pane, and then click **Connect to another computer**. On the **Select Computer** dialog box, select **Another Computer**.
   - In the **Actions** pane, click **Connect to another computer**, and then select **Another Computer**.
   - With the **Microsoft Identity Management for UNIX** node at the top of the hierarchy pane highlighted, click **Connect to another computer** on the **Action** menu. On the **Select Computer** dialog box, select **Another Computer**.
3. In the box, type the name of the computer to use, or click **Browse**.
4. When you have identified or found the computer you want, click **OK**.

| **Note** |
| --- |
| You must be logged on with an account that has administrator privileges on the remote computer.<br><br>There is no command-line method for this procedure.<br><br>To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**. |

# Identity Management for UNIX: Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Server for NIS**
Server for NIS enables a Microsoft Windows–based Active Directory domain controller to administer UNIX Network Information Service (NIS) networks.

- Before you start using Server for NIS, see Server for NIS Checklists.
- For tips about using Server for NIS, see Best practices for Server for NIS.
- For help with specific tasks, see How To.....
- For general background information, see Concepts of Server for NIS.
- For detailed reference information about command utilities for managing Server for NIS, see Server for NIS command-line utilities.
- For problem-solving instructions, see Server for NIS Troubleshooting.

# Server for NIS Checklists

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Checklists**

- Checklist: Setting up Server for NIS
- Checklist: NIS migration to Active Directory using the NIS Data Migration wizard

# Checklist: NIS migration to Active Directory using the NIS Data Migration wizard

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Checklist: NIS migration to Active Directory using the NIS Data Migration wizard**

| Step | Reference |
|---|---|
| Read more about Network Information Service (NIS) migration. | Understanding NIS Data Migration |
| Back up your Windows domain controllers, or use a test domain controller. | Windows documentation |
| Log on as a member of the Domain Admins group. | Windows documentation |
| Identify the NIS domain to migrate and how it will be migrated, keeping in mind the following:<br>• You can migrate only one NIS domain at a time, but you can merge the NIS domain with a previously migrated NIS domain.<br>• If you do not merge the domain, you can preserve the existing domain name or create a new name for the migrated domain.<br>• You should migrate an NIS domain to only one Active Directory domain. | NIS administrator |
| Make sure that the source files for the NIS maps that you want to migrate can be accessed by the computer running the NIS Data Migration wizard. (NIS map source files are the plain text files from which the NIS map databases are compiled.) You can make the files available to the Windows-based computer by using network file system (NFS), or by copying them to the Windows-based computer through File Transfer Protocol (FTP). All standard and nonstandard map source files should be in the same directory. | |
| Determine which standard NIS maps to migrate, and the directory paths to the map source files, keeping in mind the following:<br>• You can migrate maps in steps, that is, migrate one map first, then another, and so on.<br>• You should migrate the passwd map before migrating other maps (such as group or shadow maps) that rely on it.<br>• You must migrate the networks map before migrating netmasks. | NIS administrator<br>Managing standard and nonstandard NIS maps |
| Determine any nonstandard NIS maps to migrate, their formats, and the directory paths to their map source files. All standard and nonstandard map source files should be in the same directory. | NIS administrator |
| Examine NIS map source files for multiple entries having the same primary key (for example, more than one user with the same user identifier (UID) in the passwd file) and eliminate the duplicate primary keys. | NIS administrator |
| Determine Active Directory containers to which you will migrate the NIS domain. | |

| | |
|---|---|
| Decide on the strategy to migrate using one of the following two alternatives. If there are conflicts between the existing map entry and the entry being migrated during the migration process, you can choose one of the following methods:<br>• Overwrite the existing entry with data from the entry being migrated.<br>• Preserve data from the existing entry and ignore the entry being migrated.<br>• Resolve conflicts by changing the Windows account name in Active Directory. If objects of different types have the same name, the names of both objects are changed before the data is migrated. | Best practices for Server for NIS |
| Determine whether you want to run the wizard first without migrating to examine all events that will take place during migration (recommended), or migrate all data the first time the wizard is run. The migration wizard can be configured to log all the migration steps without actually migrating the data to Active Directory. | |
| Migrate NIS maps to Active Directory on a computer running Windows Server and Active Directory. | Migrate NIS domains |
| Review the log files after migration for errors or conflicts. Resolve log conflicts and migrate the resolved entries again. | Resolving migration conflicts |
| Configure any UNIX NIS subordinate (also known as slave) servers to use Server for NIS as the master NIS server. | Configuring a UNIX NIS server to use Server for NIS as the master server |
| Configure any UNIX NIS client computers to use Server for NIS as the master NIS server. | Configuring UNIX clients to use Server for NIS as the master server |
| Set the passwords of new Windows user accounts (created as a result of the NIS migration) to a known, temporary value, inform the users of the temporary password and instruct them to change their Windows password as soon as possible. | |

# Checklist: Setting up Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Checklist: Setting up Server for NIS**

| Step | Reference |
|---|---|
| Read about Server for NIS. | Concepts of Server for NIS |
| Migrate Network Information Service (NIS) maps to Active Directory. | Checklist: NIS migration to Active Directory using the NIS Data Migration wizard |
| Set how often maps are propagated to subordinate (also known as slave) NIS servers. | Change the frequency of map updates to UNIX subordinate (slave) NIS servers |
| Review and, if needed, add or remove UNIX subordinate servers in the domain. | Add or remove UNIX subordinate (slave) servers |
| As needed, install Server for NIS on other domain controllers, which will automatically be configured as subordinate servers. Install Password Synchronization on all other domain controllers. This will ensure that when a user's password is changed, the password will be updated across the domain. Server for NIS will receive the updated passwords when the data is replicated to the domain controller. | Install Server for NIS<br><br>Synchronizing passwords with an NIS domain |

# Best practices for Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The following information provides best practices for using Server for NIS.

**Best practices when migrating a Network Information Service (NIS) domain**

- Before performing the actual migration, carefully review Checklist: NIS migration to Active Directory using the NIS Data Migration wizard.
- Before starting the wizard, decide whether the domain you are migrating will remain a separate domain or will be merged with another domain on Server for NIS.
- If you migrate maps separately, you should migrate **passwd** maps before you migrate **group** or **shadow** maps. This ensures that the UNIX passwords will be stored properly.
- Be sure you know the structure of your nonstandard maps before starting the wizard. In particular, you should know the field separators, key field, file name, and map name. Nonstandard maps are accessed using only one key.
- Use the **Do not migrate (log only)** option first, so that the wizard performs all the steps but does not actually migrate the data. After analyzing the log files, run the wizard again, choosing the **Migrate and log** option.
- After examining the log, correct any problems before you migrate the data. If the same user name exists in both the Windows and NIS domains, determine whether the duplicate user names represent the person. If not, change the user name of one of the users. If the user names refer to the same person, determine whether the UNIX properties are the same. If not, determine which are correct. You can then preserve the existing entry in Active Directory, or overwrite it with the information in the UNIX map.
- If a conflict is reported during the actual migration, even though none was reported in the test migration, there might be conflicts within the NIS maps themselves. When the wizard is run with the **Log only** option, it reports only conflicts between NIS maps and Active Directory, not conflicts within the NIS maps themselves. If a conflict happens during the actual migration, resolve the conflict in the NIS maps and then migrate the NIS data that was not migrated using the command line **nis2ad –r yes** option.

**Make sure that subordinate (slave) servers are kept up-to-date**
If your NIS domain is active (that is, changes to the domain occur frequently), you should increase the frequency at which Server for NIS checks for changes. This will ensure that UNIX-based subordinate servers are updated soon after a change is registered on the master server. You can also click **Check for updates now** to immediately update subordinate servers.

**Do not migrate an NIS domain to more than one Windows Active Directory domain**
Although you can migrate an NIS domain to computers running Server for NIS in more than one Windows domain, this is strongly discouraged, because changes made in one Windows domain will not be replicated to the other domains.

**Discourage users from using yppasswd to change their NIS passwords**
Instead, users should change their NIS password by changing their Windows password. Server for NIS will then change the NIS password to match.

Server for NIS does not fully support the **yppasswd** utility available on UNIX systems. When a user runs **yppasswd**, Server for NIS changes the user's password in the NIS **passwd** map. Because **yppasswd** encrypts the new password before sending it to the NIS master server, however, Server for NIS cannot obtain a plain-text password to set the user's Windows password. Consequently, the Windows and UNIX passwords will no longer be the same. In addition, **yppasswd** can create security issues because it also sends the old password in plain text. Because the old password could also be the user's Windows password, this might expose the user's Windows password on the network.

Using Password Synchronization, you can provide users with a method for changing their NIS password using the **yppasswd** command. For more information, see Synchronizing passwords with an NIS domain.

# How To....

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

- Install Server for NIS
- Start Server for NIS
- Stop Server for NIS
- Add or remove UNIX subordinate (slave) servers
- Change a subordinate (slave) server to a master server
- Manage NIS Maps

# Install Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To install Server for NIS**

1. Click **Start**, click **Control Panel**, and then click **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. When the Windows Components Wizard opens, fill the **Active Directory Services** check box.
4. With **Active Directory Services** highlighted, click **Details**.
5. When the **Active Directory Services** dialog box opens, fill the **Identity Management for UNIX** check box.
6. With **Identity Management for UNIX** highlighted, click **Details**.
7. When the **Identity Management for UNIX** dialog box opens, fill the **Server for NIS** check box.

   Because Administration Components are required for Server for NIS operation, **Administration Components** are automatically selected for installation when you select Server for NIS.
8. Click **OK**.

   The Windows Components Wizard begins installing the Identity Management for UNIX components you selected.
9. If you are prompted to locate files required for Server for NIS installation, insert the Windows Server 2003 R2 product CD, or browse to the network location of your Windows Server 2003 R2 installation files.
10. If NIS map data compatible with Windows Services for UNIX 3.5 is stored in Active Directory, the Server for NIS installation process automatically migrates the data, and uses it to populate Server for NIS maps.
11. When installation completes, restart your computer to begin working with Server for NIS.

| **Note** |
| --- |
| To open **Add or Remove Programs**, click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.<br><br>Server for NIS can be installed only on an Active Directory domain controller. |

**See Also**
**Concepts**
Installing and uninstalling Server for NIS

# Start Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To start Server for NIS**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NIS**, and then click **Start**.

| **Note** |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**. |

**Using a command line**

- At a command prompt, type:

  **nisadmin** [*server*] **start** [**–u** *user* [**–p** *password*]]

| Argument | Description |
| --- | --- |
| *server* | The name of the server you want to start. |
| *user* | The name of the user who has administrative permissions on the server to be started, if different from the current user. |
| *password* | The password of the user who has administrator permissions on the server to be started, if different from the current user. If you type a user name but omit the password, you will be prompted for the password. |

| **Note** |
| --- |
| To view the complete syntax for this command, at a command prompt, type the following:<br><br>**nisadmin /?** |

**See Also**
**Concepts**
nisadmin
Starting and stopping Server for NIS
**Other Resources**
Connect to the computer you want to manage

# Stop Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To stop Server for NIS**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Right-click **Server for NIS**, and then click **Stop**.

| **Note** |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**. |

**Using a command line**

- At a command prompt, type:

  **nisadmin** [*server*] **stop** [**–u** *user* [**–p** *password*]]

| Argument | Description |
| --- | --- |
| *server* | The name of the server you want to stop. |
| *user* | The name of the user who has administrator permissions on the server to be stopped, if different from the current user. |
| *password* | The password of the user who has administrator permissions on the server to be stopped, if different from the current user. If you type a user name but omit the password, you will be prompted for the password. |

| **Note** |
| --- |
| To view the complete syntax for this command, at a command prompt, type:<br><br>**nisadmin /?** |

**See Also**
**Concepts**
nisadmin
Starting and stopping Server for NIS
**Other Resources**
Connect to the computer you want to manage

# Add or remove UNIX subordinate (slave) servers

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add or remove UNIX subordinate (slave) servers**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. In the console tree, expand **Server for NIS** and view the list of Network Information Service (NIS) domains.
4. Click the domain you want to modify.
5. In the details pane, click the **NIS Servers** tab, and then do one of the following:
   - To add a subordinate (also known as slave) server, in **Server Name**, type the name of the server you want to add, and then click **Add**.
   - To remove a subordinate server, in **Servers and modes**, click the server you want to remove, and then click **Remove**.
6. To save the settings, click **Apply**.

| **Note** |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**. |

**Using a command line**

- To add a UNIX subordinate (slave) server, at a command prompt, type:

  **nismap add -e "***NewSubordinate***" -a** *NISDomain* **ypservers**

  To remove a server, at a command prompt, type:

  **nismap del -k "***NewSubordinate***" -a** *NISDomain* **ypservers**

| Argument | Description |
| --- | --- |
| *NewSubordinate* | The new UNIX Network Information Service (NIS) server to be added to the NIS domain. |
| *NISDomain* | The name of the NIS domain. |

| **Note** |
| --- |

To view the complete syntax for this command, at a command prompt, type:

**nismap /?**

**See Also**

**Concepts**

nisadmin

Master and subordinate (slave) modes

**Other Resources**

Connect to a computer you want to manage

# Change a subordinate (slave) server to a master server

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To change a subordinate (slave) server to a master server**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. In the console tree, click **Server for NIS** and view the list of Network Information Service (NIS) domains.
4. Select the domain you want to modify.
5. In the details pane, click **NIS Servers**, select the Windows server you want to make the new master server, and then click **Promote**.

| **Note** |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.<br><br>Only an Active Directory domain controller can be promoted to be the master server.<br><br>When you change a subordinate server to a master server, the previous master server automatically changes to a subordinate server. |

**Using a command line**

Do one of the following:

- To change a master server to a subordinate server, at a command prompt, type:

  **nisadmin** [*server*] **mkslave –m** *newMaster***–d** *domainName* **–u** *username***–p** *password*

- Or, to change a subordinate server to a master server, at a command prompt, type:

  **nisadmin** [*server*] **mkmaster –d** *domainName***–u** *username* **–p** *password*

| Argument | Description |
| --- | --- |
| *server* | The domain controller server that is hosting Active Directory. If not specified, the current server is used. |
| *newMaster* | The name of the new master server, required for the **mkslave** command. |
| *domainName* | The name of the domain for which the change is being made. |
| *username* | The name of the user who has administrator permissions on the server to be changed, if different from the current user. |

| | |
|---|---|
| *password* | The password for the user who has administrator permissions on the server to be changed. If the password is not supplied, you will be prompted for it. |

**Note**

To view the complete syntax for this command, at a command prompt, type **nisadmin /?**.

Only an Active Directory domain controller can be promoted to be the master server.

**See Also**

**Concepts**

Master and subordinate (slave) modes

nisadmin

**Other Resources**

Connect to a computer you want to manage

# Manage NIS Maps

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage maps**

- Migrate NIS domains
- Add a group to an NIS domain
- Set the encryption method for a domain
- Change the frequency of map updates to UNIX subordinate (slave) NIS servers
- Propagate selected maps now
- Propagate changed maps now
- Use ypmatch to view an entry in an NIS map
- Use ypcat to print values in an NIS map
- Edit Maps
- Manage NIS Users
- Manage NIS Groups
- Manage NIS Computers

# Migrate NIS domains

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To migrate NIS domains**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open the NIS Data Migration wizard.
2. Follow the step-by-step directions in the wizard.

| **Note** |
|---|
| To start the NIS migration wizard, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**. Click **Server for NIS** in the navigation pane. With **Server for NIS** highlighted in the navigation pane, click the **NIS Data Migration Wizard** command in the **Actions** pane. You can also click **NIS Data Migration Wizard** on the **Action** menu, or right-click **Server for NIS**, then click **NIS Data Migration Wizard** on the shortcut menu.<br><br>Password file entries with names longer than eight characters will not be migrated.<br><br>Windows user accounts created as a result of the migration are disabled. After performing the migration, you must enable the accounts. For security reasons, it is recommended that you assign a temporary password to these accounts and instruct the affected users to change their Windows password as soon as possible. |

**Using a command line**

The following arguments are required.

| Argument | Description |
|---|---|
| **-y** *UNIXNISDomain* | Specifies the name of the Network Information Service (NIS) domain that contains the map to migrate. |
| **-a** *ActiveDirectoryNISDomain* | Specifies the NIS domain name in Active Directory. |
| *MapfileToMigrate* | Specifies the name of the NIS map source file to migrate. NIS map source files are the plain text files from which the NIS map databases are compiled. |

The **nis2ad** utility accepts the following options.

| Option | Description |
|---|---|
| **-m** | Perform the migration. If this option is omitted, the program finds and reports conflicts but does not actually perform the migration. |
| **-c** *FileName* | Specifies the file where conflict details are written. Uses a default file (%sfudir%\nis\conflicts.log) if not specified. |
| **-t** *TargetContainer* | Specifies the target container name. Applicable only when creating a new NIS domain. If not specified, uses the default or uses the container of the target domain. |
| **-f** *FileName* | Specifies the name of the log file. If not specified, **nis2ad** uses a default file (%sfudir%\nis\nis2ad.log). |

| | |
|---|---|
| **-r yes\|no** | Replace object in Active Directory with object being migrated. Default is no. |
| **-n** | Resolves conflicts by changing the Windows account name in Active Directory. If objects of different types have the same name, the name of both objects is changed before the data is migrated. |
| **-p** *Password* | If needed and if not specified, the user will be prompted. |
| **-d** *Directory* | Specifies the path of the directory that contains NIS map source files. |
| **-s** *Server* | Specifies the domain controller server hosting Active Directory. Otherwise use the current server. |
| **-u** *User* | Specifies the name of the user having administrator privileges on this computer. If not specified, **nis2ad** uses the current user. |

**Note**

To view the complete syntax for this command, at a command prompt type:

**nis2ad /?**

You can migrate only one map at a time using **nis2ad**. To migrate more than one map at a time, use the NIS Data Migration wizard.

Password file entries with names longer than eight characters will not be migrated.

Windows user accounts created as a result of the migration will be disabled. After performing the migration, you must enable the accounts. For security reasons, it is recommended that you assign a temporary password to these accounts and instruct the affected users to change their Windows password as soon as possible.

Even if you specify another user by using the **-u** option, the currently logged on user must have write permissions for the folder that will contain the log and conflict files. If necessary, modify the permissions on the folder to grant write access to the user who will be running the **nis2ad** utility before running the utility.

**See Also**

**Concepts**

Understanding NIS Data Migration

nis2ad

# Add a group to an NIS domain

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add a group to an NIS domain**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Users** folder.
3. In the details pane, right-click the group you want to add to the NIS domain, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **NIS Domain** list, click the domain to which you want to add the group.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.<br><br>The name of the group must comply with UNIX naming conventions. It must contain only lowercase alphabetic and numeric characters, no more than eight characters, and at least one alphabetic character. It must not include either a colon (:) or a newline character.<br><br>A group identifier (GID) is automatically assigned to the group. If your NIS domain has special requirements for GID values, you can change the assigned value. |

**See Also**
**Concepts**

Managing NIS groups
Change a group's ID (GID)
Add or remove members of a group
Edit passwd, group, and host map entries using Active Directory Users and Computers
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Set the encryption method for a domain

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To set the encryption method for a domain**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. In the console tree, expand **Server for NIS** and view the list of NIS domains.
4. Select the domain for which you want to set an encryption method.
5. In the details pane, click the **Maps** tab.
6. In the **UNIX passwords are encrypted using** list, click the encryption method used by all UNIX computers in the domain.

| **Note** |
|---|
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.<br><br>You can select the MD5 encryption method for a UNIX domain that consists exclusively of computers running Linux and using MD5 encryption. Domains that contain one or more computers using the crypt algorithm or that run any other operating system must use crypt. Although Linux versions 6.2 and later support MD5 encryption, Identity Management for UNIX is not supported for versions of Linux prior to version 8. |

**Using a command line**

- At a command prompt, type:

  **nisadmin** [*computer*] **encryptiontype -d** *domain* {**crypt** | **md5**} [**-u** *usr* [**-p** *pword*]]

| Argument | Description |
|---|---|
| *computer* | Specifies the remote computer you want to administer. You can specify the computer using a WINS or DNS name, or by Internet Protocol (IP) address. |
| *domain* | Specifies the name of the domain for which the change is being made. |
| *usr* | Specifies the user name of the user whose credentials are to be used. It might be necessary to add the domain name to the user name in the form *domain******username*. |
| *pword* | Specifies the password of the user specified using the **-u** option. If you specify the **-u** option but omit the **-p** option, you are prompted for the user's password. |

| **Note** |
|---|

To view the complete syntax for this command, at a command prompt, type: **nisadmin /?**

You can select the MD5 encryption method for a UNIX domain that consists exclusively of computers running Linux and using MD5 encryption. Domains that contain one or more computers using **crypt** or that run any other operating system must use **crypt**. Although Linux versions 6.2 and later support MD5 encryption, Identity Management for UNIX is not supported for versions of Linux prior to version 8.

Password encryption can only be set to one encryption method at a time. Changing the encryption method invalidates user passwords. All passwords must be reset if the encryption method is changed.

**See Also**
**Concepts**
Synchronizing user passwords
**Other Resources**
Connect to a computer you want to manage

# Change the frequency of map updates to UNIX subordinate (slave) NIS servers

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To change frequency of map updates to UNIX subordinate (slave) NIS servers**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Server for NIS**.
4. In the details pane, type the update interval you want, and then click **Apply**.

| **Note** |
|---|
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**. |

**Using a command line**

- At a command prompt, type:

  **nisadmin config pushint=**[[*days***:**]*hh***:**]*mm* [**-s** *Server*] [**-u** *User*] [**-p** *Password*]

| Argument | Description |
|---|---|
| [[*days***:**]*hh***:**]*mm* | Specifies the interval at which the service checks changes to Network Information Service (NIS) maps in Active Directory and propagates them to secondary NIS servers for all domains, in days, hours, and minutes. If *hh* is specified, *hh* must be in the range 0–23 and *mm* must be in the range 0–59. |
| *Server* | The master server for the domain. |
| *User* | The name of the user who has administrative privileges on the server to be started, if different from the current user. |
| *Password* | The password of the user who has administrative privileges on the server to be started, if different from the current user. If you type a user name but omit the password, you will be prompted for the password. |

| **Note** |
|---|
| •To view the complete syntax for this command, at a command prompt, type:<br><br>**nisadmin /?** |

**See Also**
**Concepts**
Sending periodic map updates to subordinate (slave) NIS servers

nisadmin

**Other Resources**

Connect to a computer you want to manage

# Propagate selected maps now

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To propagate selected maps now**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. In the console tree, expand **Server for NIS** and view the list of NIS domains.
4. Select the domain of interest.
5. In the details pane, click the **Maps** tab.
6. In the **NIS maps in this domain** list, click the maps to propagate, and then click **Propagate**. To select more than one map, press and hold the CTRL key while clicking the maps.

| **Note** |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**. |

**Using a command line**

- At a command prompt, type:

  **yppush** [**-d** *ActiveDirectoryNISDomain*] [**-q**] [**-t** *Timeout*] [**-h** *Hosts*] *MapName*

| Argument | Description |
| --- | --- |
| **-d** *ActiveDirectoryNISDomain* | Network Information Service (NIS) domain name in Active Directory. |
| **-q** | Quiet mode. Do not wait for response from subordinate (slave) servers and do not report errors. |
| **-t** *Timeout* | The number of seconds to wait for a response from the subordinate server before sending the next request. Must be greater than zero. The default value is **30**. |
| **-h** *Hosts* | The hosts to notify of changes. Default is all subordinate servers in the domain. Can be used multiple times for more than one computer. |
| *MapName* | The name of the NIS map to be transferred. |

**Note**

This command propagates a single map. Run the command again to propagate additional maps.

The **yppush** command pushes map data only from the Windows-based NIS master server to UNIX-based NIS subordinate servers. It does not push map data to NIS subordinate servers running Windows operating systems. Active Directory updates Windows-based NIS subordinate servers.

To view the complete syntax for this command, at a command prompt type:

**yppush /?**

**See Also**

**Concepts**

Pushing map updates immediately

yppush

**Other Resources**

Connect to a computer you want to manage

# Propagate changed maps now

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To propagate changed maps now**

- Using the Windows interface
- Using a command line

**Using the Windows interface**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Server for NIS**.
4. Click **Check Now**.

| **Note** |
| --- |
| ReplaceThisText<br><br>To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**. |

**Using a command line**

- At a command prompt, type:

  **nisadmin** [*server*] **syncall** [**–u** *user* [**–p** *password*]]

| Argument | Description |
| --- | --- |
| **syncall** | Propagate all maps. |
| *server* | The name of the server where the maps are stored. |
| *user* | The name of the user who has administrator permissions on the server, if different from the current user. |
| *password* | The password of the user who has administrator permissions on the server, if different from the current user. If you type a user name but omit the password, you will be prompted for the password. |

| **Note** |
| --- |
| The **nisadmin syncall** command propagates maps only on UNIX-based NIS subordinate servers. It does not propagate maps on NIS subordinate servers running Windows operating systems. Active Directory updates Windows-based NIS subordinate servers.<br><br>To view the complete syntax for this command, at a command prompt, type:<br><br>**nisadmin /?** |

**See Also**
**Concepts**

nisadmin
Pushing map updates immediately

# Use ypmatch to view an entry in an NIS map

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To use ypmatch to view an entry in an NIS map**

- At a command prompt, type:

  **ypmatch** [**-d** *Domain*] [**-h** *HostName*] [**-t**] [**-k**] *Key* [*Key* ...] *MapName*

  or

  **ypmatch -x**

| Argument | Description |
|---|---|
| *Domain* | A domain other than the default domain. |
| *HostName* | The name of the Network Information Service (NIS) server to which you want to send the request. |
| **-t** | Inhibits translation of nickname to map name. |
| **-k** | Prints a key followed by a colon before printing the value of the key. This is useful only if your keys do not have duplicate values, and you do not have an unreasonable number of keys to clutter the output. |
| *Key* | The key for which you want to print the value. |
| *MapName* | The name or the nickname of a map. |
| **-x** | Lists the nicknames as specified by the *MapName* parameter. |

| **Note** |
|---|
| There is no Windows interface method for this procedure.<br><br>To view the complete syntax for this command, at a command prompt, type:<br><br>**ypmatch /?** |

**See Also**
**Concepts**
ypmatch

# Use ypcat to print values in an NIS map

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To use ypcat to print values in an NIS map**

- At a command prompt, type:

  **ypcat** [**-k**] [**-d** *DomainName*] [**-t**] [**-h** *NISServer*] *MapName*

  or

  **ypcat -x**

| Argument | Description |
|---|---|
| **-k** | Displays keys and their values. |
| *DomainName* | Searches the specified domain for the specified map. |
| **-t** | Forces the **ypcat** command to bypass the nickname translation table and search only for the map specified by the *MapName* parameter. |
| *NISServer* | Name of the Network Information Service (NIS) server where the map is stored. |
| *MapName* | The name of the map to be listed. |
| **-x** | Lists the nicknames as specified by the *MapName* parameter. |

| **Note** |
|---|
| There is no Windows interface method for this procedure.<br><br>To view the complete syntax for this command, at a command prompt, type:<br><br>**ypcat /?** |

**See Also**
**Concepts**
ypcat

# Edit Maps

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Edit maps**

- Edit passwd, group, and host map entries using Active Directory Users and Computers
- Edit standard and nonstandard map entries using the command line interface
- Create the structure of a nonstandard map

# Edit passwd, group, and host map entries using Active Directory Users and Computers

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To edit passwd, group, and host map entries using Active Directory Users and Computers**

1. Open Active Directory Users and Computers.
2. In the console tree, click the Active Directory domain that contains the map you want to edit, and then do one of the following:
   - To edit passwd or group maps, click **Users**.
   - To edit the host map, click **Computers**.
3. Double-click the user or computer whose properties you want to edit.
4. Click the **UNIX Attributes** tab.
5. In the boxes provided, edit the properties you want to change.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Important** |
| --- |
| Do not use Active Directory Users and Computers to modify attributes of special users and groups, that is, users and groups with user identifiers (UIDs) or group identifiers (GIDs) of less than 100. If you do, the data returned by Server for NIS will be incorrect for those users and groups. Use the procedure described in "Using a command line" instead. |

| **Note** |
| --- |
| To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**. |

**See Also**
**Concepts**
Edit standard and nonstandard map entries using the command line interface
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Edit standard and nonstandard map entries using the command line interface

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To edit standard and nonstandard map entries using the command line**
At a command prompt, do one of the following:

- To add a user to a passwd map, type:

  **nismap add –e "***passwdEntry***" –a** *nisDomain***passwd**

- To modify a user in a passwd map, type:

  **nismap mod –e "***passwdEntry***" –k** *username***–a** *nisDomain* **passwd**

- To delete a user from a passwd map, type:

  **nismap del –k** *username* **–a** *nisDomain***passwd**

- To add a group to a group map, type:

  **nismap add –e "***groupEntry***" –a** *nisDomain***group**

- To modify a user from group map, type:

  **nismap mod –e "***groupEntry***" –k** *groupName***–a** *nisDomain* **group**

- To delete a user from group map, type:

  **nismap del –k** *groupName* **–a** *nisDomain***group**

| Argument | Description |
|---|---|
| *passwdEntry* | An entry from the passwd map. |
| *nisDomain* | The Network Information Service (NIS) domain. |
| *username* | The name of the user. |
| *groupEntry* | An entry from the group map. |
| *groupName* | The name of the group to which the user belongs. |

**Note**

To view the complete syntax for this command, at a command prompt, type:

**nismap /?**

When you use **nismap** to delete a user, group, or host object, only the UNIX attributes of the object are deleted. To delete the object from Active Directory, use the Active Directory administration tools.

**See Also**
**Concepts**
Edit passwd, group, and host map entries using Active Directory Users and Computers

Propagate changed maps now
nismap
Managing standard and nonstandard NIS maps

# Create the structure of a nonstandard map

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To create the structure of a nonstandard map**

- At a command prompt, type:

  **nismap create –i** *fieldNumber* **–g** "*separator*" *mapName*

| Argument | Description |
|---|---|
| *fieldNumber* | The number of the field that contains the key to the map. |
| " *separator* " | The character used to separate fields, in quotation marks. To specify a space as a separator, enclose the space in double quotation marks (" "). For example:<br><br>**nismap create –i 1 –g " " Phones**<br><br>creates a map called Phones in which the key field is the first field and the separator character is a space. Other white-space characters, such as tab, are also accepted. |
| *mapName* | The name of the map. |

**Notes**

| Note |
|---|
| There is no Windows interface method for this procedure.<br><br>After this structure is created, there is no way to remove it.<br><br>This structure applies to all NIS domains, so it is important to ensure that the format is consistent across all NIS domains.<br><br>Do not use the hash character (#) as a field separator because this character is used in standard maps to mark the beginning of a comment.<br><br>To view the complete syntax for this command, at a command prompt, type:<br><br>**nismap /?** |

**See Also**
**Concepts**
nismap
Managing standard and nonstandard NIS maps

# Manage NIS Users

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage NIS users**

- Add a user to an NIS domain
- Change a user's ID (UID)
- Change a user's default shell
- Change a user's home directory

# Add a user to an NIS domain

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add a user to an NIS domain**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Users** folder.
3. In the details pane, right-click the user you want to add to the NIS domain, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **NIS Domain** list, click the domain to which you want to add the user.
6. In the **Primary group name/GID** list, click the name of the user's primary group, and then click **OK**
7. In the details pane, right-click the group you selected as the user's primary group, and then click **Properties**.
8. Click the **UNIX Attributes** tab, and then click **Add**.
9. In the **Available NIS Users** list, click the name of the user you just added to the domain, click **Add**, and then click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.<br><br>The logon name of the user must comply with UNIX user-naming conventions. It must contain only lowercase alphabetic and numeric characters, no more than eight characters, and at least one alphabetic character. It must not include either a colon (**:**) or a newline character.<br><br>A user identifier (UID), default shell, and home directory are automatically assigned to the user. If your NIS domain has special requirements for these values, you can change them when you add the user to the domain or at any later time. |

**See Also**

**Concepts**

Managing NIS users
Change a user's ID (UID)
Change a user's default shell
Change a user's home directory
Edit passwd, group, and host map entries using Active Directory Users and Computers
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Change a user's ID (UID)

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To change a user's ID (UID)**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Users** folder.
3. In the details pane, right-click the user you want to administer, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **UID** box, type the new user identifier (UID) for the user.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**. |

**See Also**
**Concepts**

Managing NIS users
Add a user to an NIS domain
Change a user's default shell
Change a user's home directory
Edit passwd, group, and host map entries using Active Directory Users and Computers
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Change a user's default shell

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To change a user's default shell**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Users** folder.
3. In the details pane, right-click the user you want to administer, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **Login Shell** box, type the full path name of the new shell for the user.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**. |

**See Also**
**Concepts**

Managing NIS users
Add a user to an NIS domain
Change a user's ID (UID)
Change a user's home directory
Edit passwd, group, and host map entries using Active Directory Users and Computers
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Change a user's home directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To change a user's home directory**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Users** folder.
3. In the details pane, right-click the user you want to administer, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **Home Directory** box, type the full path name of the new home directory for the user.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**. |

**See Also**
**Concepts**

Managing NIS users
Add a user to an NIS domain
Change a user's ID (UID)
Change a user's default shell
Edit passwd, group, and host map entries using Active Directory Users and Computers
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Manage NIS Groups

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage NIS groups**

- Add a group to an NIS domain
- Change a group's ID (GID)
- Add or remove members of a group

# Add one or more groups to an NIS domain

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add a group to an NIS domain**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Users** folder.
3. In the details pane, right-click the group you want to add to the NIS domain, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **NIS Domain** list, click the domain to which you want to add the group.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.<br><br>The name of the group must comply with UNIX naming conventions. It must contain only lowercase alphabetic and numeric characters, no more than eight characters, and at least one alphabetic character. It must not include either a colon (:) or a newline character.<br><br>A group identifier (GID) is automatically assigned to the group. If your NIS domain has special requirements for GID values, you can change the assigned value. |

**See Also**
**Concepts**

Managing NIS groups
Change a group's ID (GID)
Add or remove members of a group
Edit passwd, group, and host map entries using Active Directory Users and Computers
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Change a group's ID (GID)

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To change a group's ID (GID)**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Users** folder.
3. In the details pane, right-click the group you want to administer, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **GID (Group ID)** box, type the new group identifier (GID) for the group.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**. |

**See Also**
**Concepts**
Managing NIS groups
Add a group to an NIS domain
Add or remove members of a group
Edit passwd, group, and host map entries using Active Directory Users and Computers
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Add or remove members of a group

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add or remove members of a group**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Users** folder.
3. In the details pane, right-click the group you want to administer, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. Do one of the following:
   - To add a user, click **Add**. In the **Available NIS users** list, click **Add**, and then click **OK**.
   - To remove a user, in the **Members** list, click the user you want to remove, and then click **Remove**.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**. |

**See Also**
**Concepts**
Managing NIS groups
Add a group to an NIS domain
Change a group's ID (GID)
Edit passwd, group, and host map entries using Active Directory Users and Computers
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Manage NIS Computers

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Manage NIS computers**

- Add a computer to a domain
- Specify the IP address of a computer in a domain
- Specify the alias of a computer in a domain

# Add a computer to a domain

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add a computer to a domain**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Computers** folder.
3. In the details pane, right-click the computer you want to add to the domain, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **NIS Domain** list, click the domain to which you want to add the computer.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**. |

**See Also**
**Concepts**

Managing NIS computers
Specify the IP address of a computer in a domain
Specify the alias of a computer in a domain
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Specify the IP address of a computer in a domain

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To specify the IP address of a computer in a domain**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Computers** folder.
3. In the details pane, right-click the computer you want to administer, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **IP Address** box, type the Internet Protocol (IP) address of the computer.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**. |

**See Also**
**Concepts**

Managing NIS computers
Add a computer to a domain
Specify the alias of a computer in a domain
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Specify the alias of a computer in a domain

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To specify the alias of a computer in a domain**

1. Open Active Directory Users and Computers.
2. In the console tree, click the **Computers** folder.
3. In the details pane, right-click the computer you want to administer, and then click **Properties**.
4. Click the **UNIX Attributes** tab.
5. In the **Alias** box, type the alias of the computer, and then click **Add**.
6. To save the changes and keep the dialog box open, click **Apply**. To save the changes and close the dialog box, click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure.<br><br>To open Active Directory Users and Computers, click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **Active Directory Users and Computers**.<br><br>To remove an alias of the computer, select the alias in the **List of Aliases** list, and then click **Remove**. |

**See Also**
**Concepts**
Managing NIS computers
Add a computer to a domain
Specify the IP address of a computer in a domain
Propagate changed maps now
Managing standard and nonstandard NIS maps

# Concepts of Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Concepts**

This section contains information about:

- Understanding Server for NIS
- Administering Server for NIS

# Understanding Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Understanding Server for NIS**
Server for NIS integrates Windows and Network Information Service (NIS) networks by giving a Windows–based Active Directory domain controller the ability to act as a master NIS server for one or more NIS domains. Server for NIS stores both standard and nonstandard NIS map data in Active Directory. (Standard maps consist of aliases, bootparams, ethers, hosts, group, netgroup, netid, netmasks, networks, passwd, protocols, rpc, services, pservers, and shadow; all other maps are non-standard.) Server for NIS thus creates a single name space for the Windows and NIS domains that can be managed by the Windows administrator using a single set of tools. The administrator can easily create, modify, and delete user accounts for Windows and UNIX domains at the same time. A user who has accounts on both Windows and UNIX can be managed from Active Directory with all attributes necessary for the respective domain and name space.

Server for NIS can also be installed on other domain controllers in the same domain as the master NIS server, giving them the ability to function as NIS subordinate (also known as slave) servers. In this case, the NIS data in Active Directory is automatically replicated among the Windows domain controllers running Server for NIS. In addition, UNIX-based computers can continue to function as subordinate servers in the NIS domain; Server for NIS uses **yppush** to propagate changes to the NIS data to UNIX-based NIS servers just as a UNIX-based master NIS server would.

After you install Server for NIS, your first major task is to migrate NIS maps from UNIX-based NIS servers to the computer running Server for NIS. You can use either a Windows-based migration wizard, or you can use a command-line utility to carry out this process. After the maps for an NIS domain have been migrated, Server for NIS is ready to act as a master server for the NIS domain. You can migrate multiple NIS domains to the same Windows–based Active Directory domain controller. These domains can be merged with each other or kept separate. You can migrate NIS map data in a staged manner by migrating NIS maps one at a time.

Server for NIS can support multiple NIS domains concurrently. Maps that belong to multiple domains can also reside concurrently in Active Directory. Requests from a client or another server in a domain receive data from a map in the same domain.

You can manage Server for NIS using the Windows-based Identity Management for UNIX. You can manage the passwd, group and hosts maps that you have migrated to Active Directory using Windows–based tools such as Active Directory Users and Computers. You can also use the command-line tools nismap and nisadmin to manage all standard and non-standard maps and to manage Server for NIS. You can monitor Server for NIS performance using Windows administrative tools, including using Event Viewer to audit significant events and errors logged by Server for NIS.

In addition to using administrative tools provided by Windows, you can also create additional administrative tools using Active Directory Service Interfaces (ADSI) or Lightweight Directory Access Protocol (LDAP).

Server for NIS implements the NIS remote procedure calls necessary to serve requests from UNIX-based and Windows-based NIS clients, as well as other subordinate NIS servers. Server for NIS supports both broadcast mode and the normal binding mode.

This section contains information about:

- Understanding NIS
- Master and subordinate (slave) modes
- Understanding NIS Data Migration

# Understanding NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Understanding NIS**

Formerly known as Yellow Pages (yp), Network Information Service (NIS) provides a simple network look-up service that consists of databases and processes. An NIS domain consists of a client and one or more servers. Clients use the NIS protocol to look up information stored in NIS databases, which are replicated among servers. A single master server is used to update databases; subordinate (also known as slave) servers provide read-only services. Databases are synchronized by copying them from master servers to subordinate servers periodically or upon change.

A database served by NIS is called an NIS map. The NIS lookup calls require a map (database) name and an NIS domain name. An NIS domain consists of a collection of such maps. In addition to map databases, the master NIS server maintains the time stamp that indicates when each database was last modified.

Each database is created from a source file. For example, passwd.byname and passwd.byuid databases are created from the source file named passwd. When source files change, administrators either execute a shell script called **ypmake** or a makefile to update the databases. In addition, the makefiles used to create the databases push the databases to subordinate NIS servers using **yppush**. The transfer uses ypxfr, a program that runs on the subordinate servers.

Clients can use many functions and remote procedure calls to connect to the network lookup service, including **yp_match**, **yp_first**, **yp_next**, **yp_all**, **yp_order**, **yp_master**, **yperr_string**, and **ypprot_err**. In addition there are usually many tools for administering NIS, such as **ypwhich**, **yppoll**, **ypset**, **ypcat**, and **domainname**.

| **Note** |
| --- |
| Server for NIS does not support **ypxfrd**, which might cause some UNIX subordinate (also known as slave) Network Information Service (NIS) servers to display warning messages. The warning messages originate in **ypinit** and do not indicate actual errors. Some UNIX subordinate NIS servers run **ypxfrd** on the master NIS server to replicate NIS maps. If they fail to connect to **ypxfrd**, they obtain copies of NIS maps by using **yp_all**. |

# Master and subordinate (slave) modes

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Master and subordinate (slave) modes**

Server for NIS can operate in a master mode or subordinate mode. (If Server for NIS is used in a domain, the master server in the domain must be Server for NIS.) In UNIX, the primary difference between the two is that both subordinate and master servers can read map data, while only the master server can update maps. Additionally, the master Network Information Service (NIS) server provides periodic updates of the maps to subordinate servers.

In Windows with Server for NIS, because the data is stored in Active Directory, both master and subordinate servers share the same replicated Active Directory database. This makes them peers as far as map updates or write access are concerned. Server for NIS in either master mode or subordinate mode will allow updates to Network Information Service (NIS) map data stored in Active Directory. Changes to Active Directory on any of these servers are propagated to other Servers for NIS in the domain by the Active Directory synchronization mechanism.

Although these servers run as peers for data updates, there must still be only one master server as far as the UNIX subordinate servers and UNIX clients are concerned. UNIX subordinate servers will continue to receive map updates from the Server for NIS that is designated as the master server. Similarly, they will continue to send passwd changes only to the master server.

For information about changing a subordinate server to a master server, see Change a subordinate (slave) server to a master server.

# Understanding NIS Data Migration

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Understanding NIS data migration**
Server for NIS allows Network Information Service (NIS) map data to be migrated to Active Directory.

After the migration, and upon receiving NIS requests from clients or other servers, Server for NIS searches Active Directory to reply to NIS queries.

Migration consists of the following three steps:

1. Migrate the NIS maps to a Windows-based server.

   - Using the command line

     Identity Management for UNIX includes a command-line tool called **nis2ad** to migrate maps from UNIX-based NIS servers to Active Directory–based Server for NIS.

   - Using the migration wizard

     The NIS Data Migration wizard in Identity Management for UNIX extracts the information necessary to perform the migration. Even when using the migration wizard, however, you must complete steps 2 and 3, which follow.

   The migration wizard and the **nis2ad** command read map data from NIS map source files, which are the plain text files from which the NIS map databases are compiled. These source files must be located in a location that can be accessed by the domain controller during migration, such as on a disk on the domain controller or in a shared directory accessible by the domain controller.

2. Inform the existing subordinate (also known as slave) NIS servers of the new master NIS server.

   After the migration, the original UNIX-based NIS server must send an update of maps to all subordinate NIS servers with the name of the new master server in the maps.

3. Disable the original NIS server.

UNIX-based subordinate NIS servers can continue to work as before; however, they will receive map updates from the Windows-based computer instead of the UNIX-based computer. Client computers will continue to get NIS maps or data from the new master.

This section provides information about the following:

- Migrating NIS to Active Directory
- Migrating standard and nonstandard maps
- Resolving migration conflicts
- Handling special users during migration to Active Directory

# Migrating NIS to Active Directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Migrating NIS to Active Directory**
During migration, Server for NIS can migrate each Network Information Service (NIS) domain and preserve it as a separate domain in Active Directory. It can also merge the data in an existing domain, however. This domain can be one of the previously migrated domains or a special NIS domain that is created when Server for NIS is installed.

One of the advantages of migrating NIS maps to Windows is to merge UNIX and Windows users and objects to create a common name space. Another advantage is that objects common in two networks can have a unique identity.

During migration, you must specify the Active Directory container in which migrated objects will be created. It is strongly recommended that you create these objects in the same container as the other Windows objects. If the container is not specified, new objects will be created in the default container that is the same as Windows.

If you migrate passwd, groups, or host maps to a container other than the default container, migrated entities will not appear by default in Windows Active Directory Users and Computers. To view all objects, select **Advanced features** in the **View** menu of the **Active Directory Users and Computers** console.

Windows user accounts created as a result of the migration will be disabled. After performing the migration, you must enable the accounts when you are ready for them to be used. For security reasons, it is recommended that you also assign a temporary password to these accounts and instruct the users to change their Windows password as soon as possible.

See Server for NIS Troubleshooting for information about dealing with migration problems.

**Staged migration**
After the maps are migrated to the Windows Active Directory domain controller, other UNIX-based subordinate (also known as slave) servers can also be migrated to other domain controllers in the same Windows domain. Migration of subordinate servers does not require migration of the data. Because of the architecture of Active Directory, other domain controllers in the Windows domain share the same Active Directory schema and database. Consequently, these domain controllers already have the NIS data. After the Server for NIS is installed on a domain controller in the same Windows domain, it is ready to act as a subordinate NIS server for the NIS domain.

The steps in the migration of a subordinate NIS server are as follows:

1. Make sure the master server of the NIS domain is migrated to Server for NIS.

2. Install Server for NIS on another Active Directory domain controller. The subordinate NIS server that you are migrating to Windows can continue to work as before. This will mean that each subnet can have more than one subordinate NIS server.

3. Disable the subordinate NIS server on the UNIX-based computer.

The NIS map source files of an NIS domain need not be migrated all at once. (NIS map source files are the plaintext files from which the NIS map databases are compiled.) In fact, each NIS map that belongs to a domain is migrated separately. The administrator can specify the domain name to which this map will be migrated. If a domain with that name does not exist, Server for NIS creates a new NIS domain.

After you migrate NIS maps to Server for NIS, there might be some delay before this data is available to Server for NIS.

**Migrating multiple NIS domains**
The migration utility or migration wizard allows migration of multiple NIS domains to Active Directory. During migration of a map, the administrator must specify the domain from which the data is being migrated. Server for NIS keeps the name of the domain with the map entries created in Active Directory. Subsequently, these map entries are returned to clients or subordinate servers from that domain only.

In addition, however, the migration wizard can also merge a map from one domain into another NIS domain. This can be one of the previously migrated NIS domains or a special domain that is created when Server for NIS is installed. The administrator must specify the name of the NIS domain to which the data will be migrated.

If you migrate multiple domains, you can manage individual domains through Identity Management for UNIX Administration.

Server for NIS supports a maximum of 960 unique combinations of NIS domains and maps. For example, Server for NIS cannot support more than 64 domains with 15 maps each (64×15=960).

# Migrating standard and nonstandard maps

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Migrating standard and nonstandard maps**

The command-line migration program and the NIS Data Migration wizard are both designed to migrate standard Network Information Service (NIS) maps, such as passwd, group, and services. The NIS Data Migration wizard provides a list of NIS map source files (that is, the plaintext files from which the NIS map databases are compiled) from which the administrator chooses the maps to migrate. The migration wizard then uses the standard format of these NIS map source files to migrate the data to Active Directory.

The following are considered standard maps. All other maps are nonstandard maps.

- aliases
- bootparams
- ethers
- hosts
- group
- netgroup
- netid
- netmasks
- networks
- passwd
- protocols
- rpc
- services
- ypservers
- shadow

Note that domain names are not case sensitive. Domains cannot be migrated if they have the same name, even if the case of the names is different.

Migration is handled in one of two ways, depending on whether the map being migrated has an equivalent in Active Directory. For standard maps such as passwd, group, and hosts, that have an equivalent Active Directory class, auxiliary classes are added. For example, an auxiliary class called posixAccount is added to the class user, which, in turn, adds attributes such as uidNumber and unixHomeDirectory. For standard NIS maps that have no equivalent in Active Directory, such as services and rpc, additional classes are created.

Because Identity Management for UNIX is compliant with Internet Engineering Task Force (IETF) standard Request for Comments (RFC) 2307, map attributes match those identified in the RFC 2307 schema, with the following exceptions:

- The Identity Management for UNIX equivalent of the RFC 2307 attribute **userPassword** is **unixUserPassword**.

- The Identity Management for UNIX equivalent of the RFC 2307 attribute **homeDirectory** is **unixHomeDirectory**.
- In Identity Management for UNIX, **posixgroup** is not a unique class as defined by RFC 2307. It is an auxiliary class.

For each map entry being migrated during the migration of standard maps, the migration wizard either creates a new object of the corresponding class or adds appropriate attributes to the auxiliary class.

Because the classes and attributes found in non-standard maps may not have matching correspondents in Active Directory, the migration wizard relies upon a key field to make a one-to-one match between entries in the existing map and entries created in Active Directory. Using the key field as a reference, the migration wizard copies to Active Directory other classes in the nonstandard map as they appear in the map source file.

| **Important** |
| --- |
| When you migrate passwd and group maps, users will not appear as members of their primary group. If you want users to be members of their primary group in Active Directory, you must use Active Directory Users and Computers to add users to their primary group. |

Before you migrate a nonstandard map, you must provide the syntax of the map. You must provide the name of the nonstandard map, the delimiter between the fields, and the field number of the key field.

For information about managing maps after migration, see Managing standard and nonstandard NIS maps.

**Migrating passwd maps**

If a new Windows user is created as a result of the migration, the user's Windows account is disabled. Before the user can log on to the Windows domain or a UNIX computer in the NIS domain, you must enable the Windows account. For security reasons, it is recommended that you change the user's Windows password to a known, temporary value, inform the user what the temporary password is, and then instruct the user to change the Windows password as soon as possible. When the user changes the Windows password, the new password will be propagated to UNIX-based NIS clients the next time the NIS maps are refreshed. After this refresh occurs, the user can log on to UNIX computers in the NIS domain.

If migration results in adding UNIX attributes to an existing Windows user, the UNIX password in Active Directory will be set to the existing UNIX password, and the Windows password will remain unchanged.

Server for NIS is designed so that whenever a user's Windows password is changed, the password in the NIS passwd map stored in Active Directory is changed. The reverse is not true, however. That is, when the user runs **yppasswd** on a UNIX client computer to change the user's NIS password, only the password in the NIS passwd map will be changed; the Windows password will not change. As a result, the UNIX and Windows passwords will differ. Users should change passwords only in Windows to ensure that Windows and UNIX passwords in the NIS map are identical. Note, however, that you can install and configure Password Synchronization to provide two-way password synchronization between Server for NIS and UNIX hosts in an NIS domain. For more information, see Synchronizing passwords with an NIS domain.

**Migrating ypservers maps**

Keys within the ypservers map are the ASCII names of the computers on which the NIS servers run. That is, the ypservers map can contain only a list of NIS servers, each one listed on a separate line. If your ypservers map contains aliases, before migrating the map, edit the ypservers map to contain only the host names of the servers, one per line. The following list is an example of a valid ypservers map:

```
ypserver_1
ypserver_2
ypserver_3
```

**Migrating Automount File System (AutoFS) maps**

You can migrate Automount File System (AutoFS) maps by importing the **auto_master** and **auto_home** files from a UNIX computer. You can use either command-line tools or the NIS Data Migration wizard.

Before performing the migration, you must copy the auto_master and auto_home files from the UNIX computer to any directory (such as %TEMP%) on the Windows-based computer.

When you use the wizard to migrate AutoFS maps, specify **auto_master** or **auto_home**, depending on which map you are migrating. Specify a blank space as the field separator and **1** (one) as the key field. Verify that the **Key is part of value** check box is not selected.

You can use the wizard to import other AutoFS maps as well. For example, you can migrate the auto.redirect map by copying and then importing the auto_redirect map file, making sure that the **Key is part of value** check box is not selected.

Using command-line tools to migrate AutoFS maps involves using the nismap utility to create the corresponding structure in the Active Directory schema and then running the nis2ad utility to import the map data.

For example, the following commands would migrate AutoFS maps from the NIS domain named NIS01 to the Windows domain named WIN01 by using files located in C:\Temp:

```
nismap create -i 1 -g " " -y auto.master
nismap create -i 1 -g " " -y auto.home
nis2ad -y NIS01 -a WIN01 -d C:\Temp -s localhost -r yes -m auto_master
nis2ad -y NIS01 -a WIN01 -d C:\Temp -s localhost -r yes -m auto_home
```

You can use the same method to migrate other AutoFS maps. For example, to migrate the auto.redirect map, use the **nismap** utility to create the structure for the auto.redirect map in Active Directory, copy the auto_redirect file to the Windows-based computer, and then run the **nis2ad** utility to import the data in the auto_redirect file.

**Migrating nonstandard maps**

Active Directory includes a class that applies to any NIS map. This class requires the name of the map, the key, and the rest of the entry.

To migrate nonstandard maps, the administrator must provide the name of the map source file and the format of the map entries to the migration wizard or command-line utility. The migration wizard uses the key field and the separator to migrate the data and create appropriate objects in Active Directory.

| **Important** |
| --- |
| If you migrate passwd, group, or hosts maps to a container other than the default container (**CN=users,**...), some migrated objects do not typically display in Windows Active Directory Users and Computers. To view all objects in Active Directory Users and Computers, select the domain and then click **Advanced Features** on the **View** menu. |

# Resolving migration conflicts

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Resolving migration conflicts**

During migration, conflicts may occur when Server for NIS attempts to merge incoming user accounts, groups, and computer names into classes that already exist in Active Directory. Conflicts can also occur because Server for NIS allows one Network Information Service (NIS) domain to merge with another NIS domain that has already been migrated to Server for NIS.

For example, consider a passwd map source file being migrated to Server for NIS. Suppose there is a UNIX user with a user name of johnwood. If there is no existing user with a user name of johnwood, migration of this user to Server for NIS is simple. If there is already a user with that user name previously migrated from another NIS domain, however, a conflict would occur.

When the NIS Data Migration wizard encounters a potential name conflict during migration, it resolves the conflict by prefixing the name being migrated. If the name being migrated is a user name, the prefix consists of the NIS domain name plus the characters _**u**_. If the name being migrated is a group name, the prefix consists of the NIS domain plus the characters _**g**_. This allows the migration to complete even if there is a conflict.

For example, if you are migrating a user named johnwood in an NIS domain named mktg, and a user with the same name was previously migrated from another NIS domain, then the user name of johnwood in the mktg domain will be changed to mktg_u_johnwood during the migration.

When such conflicts occur, you need to examine the migration logs and then determine how to handle each case. You need to consider whether the conflicting name represents the same user or group in both domains, or different users or groups. If not, you will probably want to rename one or both users or groups. If they do represent the same user or group, then you will need to decide whether the user or group really needs to be in both domains (in which case, at least one of them needs to be renamed) or whether the user or group in one of the domains can be deleted.

Rather than deal with the conflict after the domain migration has occurred, you can determine in advance whether a conflict will occur and then take steps to prevent the conflict (such as by changing the name of a user or group in the NIS domain).

By default, the NIS Data Migration wizard performs what is known as a dummy migration. During this process, the wizard goes through all the steps required to migrate an NIS domain to Active Directory, but does not actually change Active Directory. This lets you record the anticipated migration results in a log file and then, if potential conflicts are found, deal with those conflicts appropriately before actually performing the migration.

When identical names are detected during a dummy or actual migration, the conflict can be logged to a conflicts file. This file lists conflicts that have arisen during the migration of a particular map. If there was a conflict, the conflicts file lists the NIS entry being migrated and the existing entry in Active Directory. Consider the following two examples. In the first one, the conflicts file states that there were no conflicts in the migration of the passwd file.

```
--------------
## Tue Jun 1 16:22:47 1999 : Conflicts between entries from map file 'passwd' and
existing entries in Active Directory. ##
-------------
```

In the next example, the conflicts file states that there was a conflict in the migration of the map. It lists the existing entry in Active Directory and the new entry being migrated.

```
-------------
## Tue Jun 1 16:22:52 1999 : Conflicts between entries from map file 'aliases' and
existing entries in Active Directory. ##
EXISTS : having DN = 'CN=al1,CN=nisadmin,CN=DefaultMigrationContainer,DC=nis,
DC=sfu,DC=nttest,DC=microsoft,DC=com'
OLD : staff:wnj,mosher,sam
NEW : staff:pradeep,peter,wjs
-------------
```

According to the file, the "staff" map already exists in Active Directory. The entry in Active Directory is different from the entry being added. In such a case, you can change the name of the alias "staff" either in Active Directory or in the map source file, that is, the plain text file from which the NIS map database is compiled. You can also preserve or replace the existing entries by force.

In addition to a conflicts file, you can specify a log file where the entire migration steps will be logged. The following is a

sample output of such a log file:

```
## Start of NIS to Active Directory migration of 'passwd' @ Tue Jun 1 16:26:21 1999 ##
MESSAGE : Migrating 'passwd' entries from UNIX NIS domain 'nis01' to Active Directory
domain 'CorpDomain.'
SUCCESS : Migration of object 'nis0101' of class 'User' into
'LDAP://localhost/CN=Users,DC=nis,DC=sfu,DC=nttest,DC=microsoft,DC=com'.
SUCCESS : Migration of object 'nis0102' of class 'User' into
'LDAP://localhost/CN=Users,DC=nis,DC=sfu,DC=nttest,DC=microsoft,DC=com'.
## Start of NIS to Active Directory migration of 'passwd' @ Tue Jun 1 16:41:46 1999 ##
MESSAGE : Migrating 'passwd' entries from UNIX NIS domain 'conflicts' to Active
Directory domain 'conflicts'.
CONFLICT : Can't migrate 'nis0101' to
'LDAP://localhost/CN=Users,DC=nis,DC=sfu,DC=nttest,DC=microsoft,DC=com'. An object
having same attributes(name/uidNumber/gidNumber) exists at
'CN=nis0101,CN=Users,DC=nis,DC=sfu,DC=nttest,DC=microsoft,DC=com'.
```

For more information about dealing with migration problems, see Server for NIS Troubleshooting.

# Handling special users during migration to Active Directory

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Handling special users during migration**

During migration, for special users with user identifiers (UIDs) that are less than 20, Windows accounts are created such that the first nine characters of their Network Information Service (NIS) domain are appended to their UNIX user name to create a Windows logon name. Their UNIX user name is unaffected. For example, if you migrate root from physlab to Server for NIS, a new Windows user will be created called *physlab_u_root*. The UNIX user name is still root, and NIS queries from physlab domain will be unaffected.

A group called root from such a domain, if migrated, will create a group called *physlab_g_root*.

# Administering Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering Server for NIS**
This section provides information about:

- Installing and uninstalling Server for NIS
- Starting and stopping Server for NIS
- Configuring a UNIX NIS server to use Server for NIS as the master server
- Configuring UNIX clients to use Server for NIS as the master server
- Using the migration wizard
- Using the command-line utility for migration
- Managing standard and nonstandard NIS maps
- Managing NIS users
- Managing NIS groups
- Managing NIS computers
- Synchronizing user passwords
- Setting UNIX passwords to special values
- Synchronizing and propagating NIS changes
- Sending periodic map updates to subordinate (slave) NIS servers
- Pushing map updates immediately

# Installing and uninstalling Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Installing and uninstalling Server for NIS**

**The NIS domain**

Upon installation, Server for NIS creates a Network Information Service (NIS) domain with the same name as the Active Directory domain. Administrators can add users, groups, or computers to this domain. Moreover, administrators can migrate UNIX NIS domain data to Server for NIS and designate the Server for NIS computer as the master NIS server for the migrated domain.

**Uninstalling in master mode**

If you uninstall Server for NIS while it is running on a master server, you must verify that another server takes over the tasks of the master server. If other Windows-based subordinate (also known as slave) NIS servers are in the domains supported by the master server you uninstall, you must assign one of these servers the role of master server. Server for NIS cannot be a subordinate server to a master NIS server running on a UNIX operating system.

**Uninstalling in subordinate mode**

If you uninstall Server for NIS running as a subordinate server, the domain controller on which it was running will continue to receive NIS map updates through Active Directory Synchronization.

# Starting and stopping Server for NIS

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Starting and stopping Server for NIS**

You can use Identity Management for UNIX to start or stop Server for NIS. When Server for NIS is stopped, it will not respond to Network Information Service (NIS) requests from clients, including requests for password changes.

Even when Server for NIS is stopped, you can continue to make changes to NIS map data through Active Directory. UNIX-based computers respond differently in this situation, however. If the NIS server is down on UNIX, changes to maps undertaken through **make** will not be updated in the NIS database.

In addition, when Server for NIS is stopped, NIS map changes will not be propagated to subordinate (also known as slave) NIS servers running on UNIX. However, subordinate NIS servers running on Windows domain controllers will continue to receive updates.

# Configuring a UNIX NIS server to use Server for NIS as the master server

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Configuring a UNIX NIS server to use Server for NIS as the master server**
To change a UNIX-based Network Information Service (NIS) server from a master server to a subordinate (also known as slave) server, follow these steps:

1. Migrate maps from the UNIX master NIS server.

2. Transfer the maps from the old master server to other subordinate NIS servers by providing the name of the new Server for NIS for each map. At a command prompt, type:

   **ypxfr –h** *newserver mapname*

   where *newserver* is the name of the new NIS master server, and *mapname* is the name of the map to be transferred.

3. Run this command for each map on each of the subordinate servers. After this step, the UNIX subordinate servers will recognize the new Server for NIS master server.

# Configuring UNIX clients to use Server for NIS as the master server

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Configuring UNIX clients to use Server for NIS as the master server**

For each UNIX client computer on which you want to use Server for NIS as the master server, do the following:

1. If the domain name has changed, run the following command:

   **domainName** *newdomain*

2. Stop **ypbind** by running the following command:

   **kill –9 ypbindpid**

3. Restart **ypbind** by running the following command:

   **ypbind**

- Network Information Service (NIS) subordinate (also known as slave) servers require that the NIS master server name in the local /etc/hosts file match the master server name in the NIS maps exactly, including the case of the letters of these names. Make sure the case of the host name in /etc/hosts matches exactly the master server name as returned by the Server for NIS. Server for NIS typically returns the host computer name in all uppercase.

# Using the migration wizard

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Using the migration wizard**

Identity Management for UNIX provides a wizard that you can use to migrate your Network Information Service (NIS) system easily to Windows. The wizard takes you through a step-by-step process to complete the migration. Before using the wizard, review the following:

- Checklist: NIS migration to Active Directory using the NIS Data Migration wizard.
- Best practices for Server for NIS
- Migrating NIS to Active Directory

| **Caution** |
| --- |
| The migration of NIS to Active Directory is not reversible. Always back up your UNIX NIS systems and Windows-based server before performing an NIS migration. |

# Using the command-line utility for migration

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Using the command-line utility for migration**

Instead of using the migration wizard, you can use a command-line utility to migrate maps. For information about the command-line tool, see nis2ad.

| **Caution** |
| --- |
| The migration of Network Information Service (NIS) to Active Directory is not reversible. Make sure you back up your UNIX NIS systems and Windows server before doing the migration. |

# Managing standard and nonstandard NIS maps

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Managing standard and nonstandard maps**
You can use a Windows-based Network Information Service (NIS) server to manage both standard and nonstandard maps. You can manage maps that are part of a domain served by a Windows-based NIS server through two different mechanisms:

- The Identity Management for UNIX graphical user interface (GUI).
- Command-line utilities.

In many cases, you can use either method to perform a task. In some cases, however, only one method or the other can be used.

For information about migrating maps, see Migrating standard and nonstandard maps

**Managing passwd maps**
You can use Active Directory Users and Computers to manage passwd maps for Server for NIS. You can also use Active Directory Users and Computers to manage password file–related attributes for a user in the Windows domain, or for the users migrated from a given NIS domain.

**Managing group maps**
Data for group maps for Server for NIS are managed through Active Directory Users and Computers. Each group in the Windows domain or from the groups migrated from a given NIS domain can be managed by using Active Directory Users and Computers.

Note that, for a given group, the UNIX members are different from the Windows users for the same group.

**Managing hosts maps**
You can manage data for hosts maps for Server for NIS by using Active Directory Users and Computers. Each computer in the Windows domain, as well as the hosts migrated from a given NIS domain, can be managed through this snap-in.

**Managing other maps and nonstandard maps**
Server for NIS does not include GUI tools for managing standard NIS maps other than passwd, group, and hosts. Maps other than these can be managed using tools provided as part of Active Directory. In particular, administrators can use the adsiedit snap-in to view and edit NIS maps.

Standard maps include:

- passwd
- group
- hosts
- aliases
- bootparams
- ethers
- netgroup
- netid
- netmasks
- networks

- protocols
- rpc
- services
- ypservers
- shadow password maps

# Managing NIS users

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Managing NIS users**
UNIX systems store basic information about authorized users in the passwd file. This file consists of a series of lines, one for each authorized user. Each line can contain the following information about a user:

- User name
- Password (in encrypted form)
- User identifier (UID)
- Group identifier (GID) of the user's primary group
- User's full name and other information (known as the gecos field)
- Default login shell
- Home directory

Conventional Network Information Service (NIS) servers generate two maps from the passwd file, passwd.byname and passwd.byuid, which contain the same information as the passwd file, but are indexed to provide quick access to the records.

Instead of storing this information in files, Server for NIS maintains user information in Active Directory entries for each user. This allows you to manage the Windows and UNIX (NIS) information together by using the user properties dialog box in Active Directory Users and Computers. When adding a new user, for example, you have to enter the information shared by Windows and NIS (the user name, password, and full name and other gecos-type information) only once. In the case of UNIX-specific information (UID, primary GID, login shell, and UNIX home directory), you can enter this information using the **UNIX Attributes** tab of the same dialog box.

For more information about how to manage UNIX user accounts using Active Directory Users and Computers, see Manage NIS Users.

# Managing NIS groups

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Managing NIS groups**
UNIX systems store basic information about groups of users in the group file. This file consists of a series of lines, one for each group. Each line can contain the following information about a group:

- Group name
- Group identifier (GID)
- List of users belonging to the group

Conventional Network Information Service (NIS) servers generate two maps from the group file, group.byname and group.bygid, which contain the same information as the group file, but are indexed to provide quick access to the records.

Instead of storing this information in files, Server for NIS maintains group information in Active Directory entries for each group. This allows you to manage the Windows and UNIX (NIS) information together by using the group properties dialog box in Active Directory Users and Computers. When creating a new group, for example, you must enter the group name, which is shared by Windows and NIS, only once. In the case of UNIX-specific information (GID and members), you can enter this information using the **UNIX Attributes** tab of the same dialog box.

For more information about how to manage UNIX group accounts using Active Directory Users and Computers, see Manage NIS Groups.

# Managing NIS computers

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Managing NIS computers**
UNIX systems store basic information about host computers on the network in the hosts file. This file consists of a series of lines, one for each computer. Each line can contain the following information about a computer:

- Internet Protocol (IP) address
- Computer name
- List of aliases

A UNIX computer consults this list to determine the IP address of another computer when that computer is referred to by name or alias, rather than by address.

Instead of storing this information in a file, Server for NIS maintains information about network computers in Active Directory entries for each computer. This allows you to manage the Windows and UNIX (NIS) information together by using the computer properties dialog box in Active Directory Users and Computers. When adding a computer to the Windows and Network Information Service (NIS) domains, for example, you must enter the computer name, which is shared by Windows and NIS, only once. In the case of UNIX-specific information (IP address and aliases), you can enter this information using the **UNIX Attributes** tab of the same dialog box.

Note that you have to enter the IP address only if it is statically assigned. If the address is assigned by a DHCP server, the hosts map will be automatically updated whenever the computer's IP address changes.

For more information about how to manage UNIX computer name and address information using Active Directory Users and Computers, see Manage NIS Computers.

# Synchronizing user passwords

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Synchronizing user passwords**

Server for NIS provides limited support for keeping passwords synchronized between a user's Windows and UNIX accounts. Whenever a user's Windows password is changed, Password Synchronization (which you can install with Server for NIS for this purpose) captures the new password, encrypts it, and then stores the password in the passwd map in Active Directory. The new password is then propagated to UNIX-based Network Information Service (NIS) secondary servers and clients when Server for NIS performs its periodic map update.

If Server for NIS is installed in a Windows domain with multiple domain controllers, it is recommended that you install Password Synchronization on all domain controllers in the domain. This is because any domain controller can potentially respond to a request to change a user's password, and so it is recommended that Password Synchronization be running on that domain controller to ensure that the NIS passwd map is updated. If you install Password Synchronization on a domain controller solely to support Server for NIS in this fashion, it is not necessary to configure UNIX hosts to work with Password Synchronization running on the domain controller. On the other hand, you can configure Password Synchronization on a domain controller and selected UNIX hosts to provide two-way password synchronization. This will allow users of the UNIX hosts to use the **passwd** command (instead of **yppasswd**) on those hosts to change their NIS domain password as well.

When synchronizing passwords, Server for NIS can use either **crypt** or Message Digest 5 (MD5). Server for NIS can support different encryption methods for multiple domains, but all UNIX computers in a particular domain must use the same encryption method. See Set the encryption method for a domain for how to specify the encryption method.

# Setting UNIX passwords to special values

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Setting UNIX passwords to special values**

You can set the password for a UNIX user to a special value, such as NP (no password) or * (no login) without affecting the user's Windows password. To do this, use **nismap** to change the Network Information Service (NIS) password entry for specific user, as follows:

- To set password to no password (usually a blank):

  **nismap mod -e "username::***uid***:***gid***:***gecos***:/***home***/***username***:/***login_shell***" -k** *username* **passwd**

- To set password to * (no logon):

  **nismap mod -e "***username***:*:***uid***:***gid***:***gecos***:/***home***/***username***:/***login_shell***" -k** *username* **passwd**

- To set password to *LK* (locked):

  **nismap mod -e "username:*LK*:***uid***:***gid***:***gecos***:/***home***/***username***:/***login_shell***" -k** *username* **passwd**

# Synchronizing and propagating NIS changes

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Synchronizing and propagating NIS changes**
A Windows-based master Network Information Service (NIS) server can propagate changes in maps to subordinate (also known as slave) NIS servers in the same domain. These changes are propagated to subordinate NIS servers using the features of the NIS protocol.

The master NIS server supports the following transfer modes:

- Periodic transfer of NIS maps.
- On-demand transfer of NIS maps to subordinate servers. Typically, NIS maps get transferred to subordinate servers upon change through **make**. The Windows-based NIS server does not use **make**; instead, the server pushes immediately.
- Subordinate NIS servers can request transfer of maps at any time.

# Sending periodic map updates to subordinate (slave) NIS servers

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Sending periodic map updates to subordinate (slave) NIS servers**
In Network Information Service (NIS), a master NIS server sends periodic map updates to subordinate NIS servers to ensure that both the master and subordinate servers have the same view of the maps. This process uses **yppush**. Server for NIS also sends periodic map updates to subordinate NIS servers. You can set the frequency of map updates through Server for NFS Administration.

Server for NIS uses **yppush** to send map updates only to subordinate NIS servers running on UNIX.

Because Server for NIS uses Active Directory to store data, both master and subordinate servers share the same replicated Active Directory database. Changes to maps for a given domain are propagated to other Servers for NIS running in subordinate mode by Active Directory's synchronization mechanism. Thus, Windows-based subordinate NIS servers receive updates from the master NIS server with every change. The update frequency is the property of the Server for NIS and is fixed for all domains that server supports.

For information about how to change the frequency of map updates, see Change the frequency of map updates to UNIX subordinate (slave) NIS servers.

For information about using **yppush**, see yppush.

# Pushing map updates immediately

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Pushing map updates immediately**

In Network Information Service (NIS), whenever a change is made to map data, administrators run **make** to make corresponding changes to the NIS database. Typically, as a part of this process, **make** also propagates the changes to subordinate (also known as slave) servers through **yppush**. This provides a way for administrators to push changes immediately to subordinate servers.

Because Server for NIS does not use **make** to make changes to the database, it provides an explicit **yppush** command to immediately propagate changes to UNIX-based subordinate servers. Server for NIS also provides a control in Identity Management for UNIX to immediately push maps to subordinate NIS servers.

Note that you use **yppush** to push changes only to subordinate NIS servers running on UNIX. Because Server for NIS uses Active Directory to store data, both master and subordinate servers running Windows operating systems share the same replicated Active Directory database. Changes to Active Directory are immediately propagated to other Servers for NIS. Thus, it is not necessary to use **yppush** to push changes to subordinate NIS servers running Windows operating systems.

For information about propagating maps immediately, see Propagate selected maps now and Propagate changed maps now.

For information about using yppush, see yppush.

# Server for NIS command-line utilities

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Command-line Utilities**

The following command-line utilities are available to administer Server for NIS:

nis2ad

nisadmin

nismap

ypcat

ypclear

ypmatch

yppush

# yppush

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**yppush**
**NAME**
yppush - Force propagation of a changed Network Information Service (NIS) map

**SYNOPSIS**

```
yppush [-d adNisDomain] [-q] [-t timeout] [-h host[...]]mapname
```

**DESCRIPTION**
The **yppush**(1w) Windows command-line utility sends a request to UNIX-based subordinate (also known as slave) NIS servers to transfer a new version of the NIS map specified by *mapname* from the master NIS server. It is run only on the master NIS server after the NIS map has been changed. If the **-h** option is specified, the transfer request is sent only to the specified server. Otherwise, **yppush** first constructs a list of NIS server hosts by reading the ypservers NIS map within the domain. For more information about the correct format of the ypservers map, see "Note" later in this topic.

The **yppush** command sends a transfer map request from the NIS server to the transfer agent on each subordinate host computer. A host's transfer agent is the utility that copies map data from the master server. When the transfer has completed, whether the transfer operation succeeded or failed, the transfer agent returns a status message, the contents of which can be printed to standard output. Transfer agents also return messages when a transfer is not possible; for example, when the **yppush** request is undeliverable to one or more host computers, or when the time-out period on responses from host computers has expired.

The **yppush** command accepts the following options and arguments:

| Term | Definition |
|---|---|
| **-d** *adNisDomain* | Specifies the NIS domain. If this option is not used, the NIS domain is assumed to be the domain named in the **DEFAULT_NIS_DOMAIN** environment variable. |
| **-q** | Quiet mode. Does not wait for a response from subordinate servers, and does not report errors. |
| **-t** *timeout* | Specifies a time-out value, in seconds, that **yppush** will wait for a response from the subordinate server. This value must be greater than zero. The default value is **30 seconds**. |
| **-h** *host*[...] | The name of the host to notify of the map change. You can specify multiple host names, separated by spaces. If you do not specify this option, the notification is sent to all subordinate servers in the domain. |

**Note**

Keys within the ypservers map are the ASCII names of the computers on which the NIS servers run. That is, the ypservers map should contain a list of NIS servers, each one listed on a separate line. If your ypservers map contains aliases, edit the ypservers map to contain only the host names of the servers, one per line. The following list is an example of a valid ypservers map:

```
ypserver_1
ypserver_2
ypserver_3
```

**See Also**
**Concepts**
Identity Management for UNIX: Server for NIS

# ypclear

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ypclear**
**NAME**
ypclear - Windows command-line utility to clear the Network Information Service (NIS) server cache for an NIS map and its dependent maps

**SYNOPSIS**

```
ypclear [-d domainName] [-t] [-h nisServer] mapname
ypclear -x
```

**DESCRIPTION**
The **ypclear**(1w) Windows command-line utility tells the specified Network Information Service (NIS) server in a specified domain that a map has changed and that the server should clear its cache. Without this notification, the NIS server can sometimes serve out-of-date information.

The **ypclear** command can be used to rebuild the map cache of an NIS server that still displays obsolete map data, even when Active Directory shows a current and accurate map of the network. Typically, using **ypclear** corrects the problem.

The **ypclear** command accepts the following options and arguments:

| Term | Definition |
|---|---|
| **-d** *domain* | Specifies the NIS domain. If this option is not used, the NIS domain is assumed to be the domain named in the **DEFAULT_NIS_DOMAIN** environment variable. |
| **-t** | Inhibits map nickname translation. |
| **-h** *nisServer* | Specifies the name of the NIS server. If this option is not used, the NIS server is assumed to be the server named in the **NIS_SERVER** environment variable. |
| **-x** | Displays the map nickname translation table. |

**See Also**
**Concepts**
Identity Management for UNIX: Server for NIS

# ypmatch

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ypmatch**
**NAME**
ypmatch - Windows command-line utility to display the values of one or more keys from a Network Information Service (NIS) map

**SYNOPSIS**

```
ypmatch [-k] [-t] [-d domain] [-h nisServer]key[...]mapname
ypmatch -x
```

**DESCRIPTION**
The **ypmatch**(1w) Windows command-line utility displays the values one or more keys from the Network Information Service (NIS) map specified by *mapname*, which may be a map name or a map nickname.

The **ypmatch** command accepts the following options and arguments:

**-k**
Displays map keys. This option is useful with maps in which the values are null or the key is not part of the value.

**-t**
Inhibits map nickname translation.

**-d *domain***
Specifies the NIS domain. If this option is not used, the NIS domain is assumed to be the domain named in the *DEFAULT_NIS_DOMAIN* environment variable.

**-h *nisServer***
Specifies the name of the NIS server. If this option is not used, the NIS server is assumed to be the server named in the *NIS_SERVER* environment variable. If *NIS_SERVER* is not set, **ypmatch** broadcasts the request.

***key***
The key for which you want to display the value. To display the value of more than one key, specify the keys separated by spaces. If any key begins with a slash mark (/), backslash (\) or hyphen (-), these special characters must be preceded by a backslash (\). For example, if the key name is **\mykey**, then use **\\mykey**.

**-x**
Displays the map nickname translation table.

**See Also**
**Concepts**
ypcat
Identity Management for UNIX: Server for NIS

# ypcat

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**ypcat**
**NAME**
ypcat - Windows command-line utility to print values of all keys in an NIS map

**SYNOPSIS**

```
ypcat [-k] [-t] [-h nisServer] [-d domain]mapname
ypcat -x
```

**DESCRIPTION**
The **ypcat**(1w) Windows command-line utility prints the values of all keys from the Network Information Service (NIS) database specified by *mapname*, which can be a map name or a map nickname.

The **ypcat** command accepts the following options and arguments:

| Term | Definition |
|---|---|
| **-k** | Displays map keys. This option is useful with maps in which the values are null or the key is not part of the value. |
| **-t** | Inhibits map nickname translation. |
| **-d** *domain* | Specifies the NIS domain. If this option is not used, the NIS domain is assumed to be the domain named in the *DEFAULT_NIS_DOMAIN* environment variable. |
| **-h** *nisServer* | Sends a request to the server name represented by the *nisServer* variable. The **-h** option broadcasts the request to all other NIS servers in the domain specified with the **-d** option, if no server name is provided. |
| **-x** | Displays the map nickname translation table. |

**See Also**
**Concepts**
ypmatch
Identity Management for UNIX: Server for NIS

# nisadmin

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**nisadmin**
**NAME**
nisadmin - Windows command-line utility to administer Server for NIS

**SYNOPSIS**

```
nisadmin [computer] [-u usr[-p pword]]
nisadmin [computer] [-u usr[-p pword]] mkmaster -d domain
nisadmin [computer] [-u usr[-p pword]] mkslave -d domain -m server
nisadmin [computer] [-u usr[-p pword]] config option[...]
nisadmin [computer] [-u usr[-p pword]] syncall
nisadmin [computer] [-u usr[-p pword]] start
nisadmin [computer] [-u usr[-p pword]] stop
nisadmin [computer] [-u usr[-p pword]] pause
nisadmin [computer] [-u usr[-p pword]] continue
nisadmin [computer] [-u usr[-p pword]] encryptiontype
         -d domain{crypt | md5}
```

**DESCRIPTION**
The **nisadmin** Windows command-line utility administers Server for NIS on the local computer or on a remote computer. If you are logged on with an account that does not have the required privileges, you can specify a user name and password of an account that does. The action performed by **nisadmin** depends on the command argument you supply.

In addition to specific command arguments, **nisadmin** accepts the following options and arguments:

| Term | Definition |
|---|---|
| *computer* | Specifies the remote computer you want to administer. You can specify the computer using a Windows Internet Name Service (WINS) or Domain Name System (DNS) name, or by Internet Protocol (IP) address. |
| **-u** *usr* | Specifies the user name of the user whose credentials are to be used. It might be necessary to add the domain name to the user name in the form *domain\username*. |
| **-p** *pword* | Specifies the password of the user specified using the **-u** option. If you specify the **-u** option but omit the **-p** option, you are prompted for the user's password. |

The specific action that **nisadmin** performs depends on the command argument you specify:

| Term | Definition |
|---|---|
| **mkmaster** | Changes a subordinate (also known as slave) server to a master server. The following option and argument are required when you use **mkmaster**:<br>• **-d** *domain* Specifies the name of the domain for which the change is being made. |
| **mkslave** | Changes a master server to a subordinate (slave) server. The following options and arguments are required when you use **mkslave**:<br>• **-d** *domain* Specifies the name of the domain for which the change is being made.<br>• **-m** *server* Specifies the name of the domain controller to be promoted to master server. |

| | |
|---|---|
| **config** | Configures service parameters. You must supply one or more of the following options when you use **config**:<br>• **pushint =**[[*days***:**]*hh***:**]*mm* Specifies the interval at which the service checks changes to Network Information Service (NIS) maps in Active Directory and propagates them to secondary NIS servers for all domains, in days, hours, and minutes. If *hh* is specified, *hh* must be in the range 0-23 and *mm* must be in the range 0-59.<br>• **logging =**{**n**[**ormal**] \| **v**[**erbose**]} Specifies the amount of information to be available in event logs. Specify **normal** (or **n**) to log only failure events; specify **verbose** (or **v**) to log failure, warning, and success events. |
| **syncall** | Synchronizes changed maps with those on UNIX-based subordinate servers immediately. |
| **start** | Starts the Server for NIS service. |
| **stop** | Stops the Server for NIS service. |
| **pause** | Pauses the Server for NIS service. |
| **continue** | Continues running the Server for NIS service. |
| **encryptiontype** | Specifies the method used to encrypt UNIX passwords for the specified *domain*. For domains whose UNIX-based computers are all running Linux with MD5 encryption, this value can be **md5**. For domains containing one or more UNIX-based computers running any other operating system or encryption method, this value must be **crypt** (the default). All computers in the domain must use the same method.<br>Password encryption can only be set to one encryption method at a time. Changing the encryption method invalidates user passwords. All passwords must be reset if the encryption method is changed. |

If you do not supply a command argument, **nisadmin** displays the current Server for NIS settings.

**See Also**
**Concepts**
Identity Management for UNIX: Server for NIS

# nis2ad

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**nis2ad**
**NAME**
nis2ad - Windows command-line utility to migrate maps from Network Information Service (NIS) to Active Directory directory service

**SYNOPSIS**

```
nis2ad.exe -y NIS_Domain -a AD_domain [options]mapfile
```

**DESCRIPTION**
The **nis2ad** Windows command-line utility transfers the entries in the Network Information Service (NIS) map source file specified by *mapfile* located in the UNIX NIS domain specified by *NIS_domain* to the Active Directory domain specified by *AD_domain*. NIS map source files are the plain text files from which the NIS map databases are compiled.

The **-m** option must be specified to perform the actual migration; otherwise, **nis2ad** performs the actions it would perform in an actual migration without actually changing Active Directory, including logging any conflicting entries.

The **nis2ad** utility accepts the following options and arguments:

| Term | Definition |
|---|---|
| **-d** *directoryname* | Specifies the name of the directory that contains NIS map source files. |
| **-s** *server* | Specifies the domain controller server hosting Active Directory; otherwise, uses the current server. |
| **-u** *user* | Specifies the name of a user having administrator permissions on this computer; otherwise, uses the current user. |
| **-p** *password* | If a password is needed but is not specified, **nis2ad** prompts the user for the password. |
| **-f** *logfile* | Specifies the name of the log file. The log file must be on your local computer. The default is %sfudir%\nis\nis2ad.log. |
| **-c** *conflictfile* | Specifies the file where conflict details are written. The default is %sfudir%\nis\conflicts.log. |
| **-r** {**yes**\|**no**} | Replaces the object in Active Directory with the object being migrated. The default is **no**. |
| **-n** | Resolves conflicts by changing the Windows account name in Active Directory. If objects of different types have the same name, the name of both objects is changed before the data is migrated. |
| **-t** *targetContainer* | Specifies the target container name. Applicable only when creating a new NIS domain. If not specified, uses default or uses the container of the target domain. |
| **-m** | Performs the migration. If this option is omitted, the program finds and reports conflicts, but does not actually perform the migration. |
| **-h/-?** | Displays usage information for the command. |

**Note**

The currently logged on user must have write permissions for the folder that will contain the log and conflict files, even if another user is specified with the **-u** option. If necessary, modify the permissions on the folder to grant write access to the user who will be running the **nis2ad** utility before running the utility.

**See Also**
**Concepts**

# nismap

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**nismap**
**NAME**
nismap - Windows command-line utility to manage Network Information Service (NIS) maps

**SYNOPSIS**

```
nismap add [common_option[...]] -e "mapentry"
           [-r {yes | no}] [-c conflictfile]mapname
nismap mod [common_option[...]] -e "mapentry" -k key mapname
nismap del [common_option[...]] -k key mapname
nismap create [common_option[...]] -i fieldnum -g separator [-y]mapname
```

**DESCRIPTION**
The **nismap** Windows command-line utility manages the Network Information Service (NIS) map specified by *mapname* maintained by Server for NIS on a local computer or a remote computer. If you are logged on with an account that does not have the required privileges, you can specify the user name and password of an account that does. The specific action that **nismap** performs depends on the command argument you specify.

In addition to specific command arguments, **nismap** accepts the following common options and arguments, represented by *common_option* in the command synopsis:

| Term | Definition |
|---|---|
| **-a** *AD_domain* | The name of the NIS domain in Active Directory. This option is not available for the **create** command argument. |
| **-f** *logfile* | Specifies the fully qualified path and name of the log file. Logged events are appended to this file. The default path and file name is %temp%\conflicts.log. |
| **-s** *server* | Specifies the name of the domain controller. The default is the domain controller of the domain of the computer on which the command is executed. |
| **-u** *usr* | Specifies the user name of the user whose credentials are to be used. It might be necessary to add the domain name to the user name in the form *domain******username*. |
| **-p** *pword* | Specifies the password of the user specified using the **-u** option. If you specify the **-u** option but omit the **-p** option, you are prompted for the user's password. |
| **-h/-?** | Displays usage information for the command. |

The following command arguments are accepted by **nismap**:

| Term | Definition |
|---|---|
| **add** | Adds a map entry. The following options and arguments are used with **add**:<br>• **-e "** *mapentry* **"** Specifies the new map entry as a string in the appropriate NIS map format. The string must be enclosed in quotation marks (**" "**).<br>• **-r {yes \| no}** Specifies whether the existing object in Active Directory is to be replaced with this entry. The default is **no**.<br>• **-c** *conflictfile* Specifies the name of the file in which conflict information is to be recorded. If this option is not provided, a default file name is used. |

| | |
|---|---|
| **mod** | Modifies an existing map entry. The following options and arguments are required with **mod**<br>• **-e "** *mapentry* **"** Specifies the replacement map entry as a string in the appropriate NIS map format. The string must be enclosed in quotation marks (**" "**).<br>• **-k** *key* Specifies the search key that identifies the map entry to be modified. |
| **del** | Deletes an existing map entry. The following option and argument are required with **del**:<br>• **-k** *key* Specifies the search key that identifies the map entry to be deleted. |
| **create** | Creates a structure for a nonstandard map. The following options and arguments are required with **create**:<br>• **-i** *fieldnum* Specifies the field number of the key field.<br>• **-g** *separator* Specifies the character that separates fields in the map.<br>• **-y** Specifies that the key is not part of the value for this map. |

**Note**

When you use **nismap** to delete a user, group, or host object, only the UNIX attributes of the object are deleted from Active Directory. To delete the object from Active Directory, use the Active Directory administration tools.

Do not use **nismap** to change a user's password. Use net.exe or the Windows interface (CTRL+ALT+DELETE) instead. The **nismap** utility will fail if you change the passwd entry using the **nismap mod** command. This value is the encrypted string of the user's password. If you need to change the passwd entry, use ypcat(1w) or ypmatch(1w) to retrieve the user's current passwd entry and to ensure that the password field is identical.

If you try to modify a mail alias using the **nismap mod** command, Server for NIS adds the new alias to the original one instead of replacing it. To change the alias, first use **nismap del** to delete the original alias, and then use **nismap add** to add the modified alias.

**See Also**
**Concepts**
Identity Management for UNIX: Server for NIS

# Server for NIS Troubleshooting

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Troubleshooting**
What trouble are you having?


- NIS Data Migration wizard or command-line utility completes successfully, but no maps are added to Server for NIS
- Migration conflicts occur even though test migration with "Do Not Migrate" option succeeded
- A conflict occurs while trying to migrate more than one user with the same user identifier (UID)
- NIS Data Migration wizard or command-line utility reports Network Information Service (NIS) name conflicts with system account
- The ypmatch utility on an HP-UX computer fails
- The first request to Server for NIS fails
- Objects migrated to nonstandard containers do not appear in Windows Active Directory Users and Computers
- The ypcat utility sometimes displays incorrect passwd information, but ypmatch provides accurate information
- Users with new Windows accounts created by NIS migration cannot log on to Windows and cannot log on to NIS client computers
- A user account disabled in Active Directory is not also disabled on UNIX hosts
- A newly migrated user is not authenticated from a UNIX client


**NIS Data Migration wizard or command-line utility completes successfully, but no maps are added to Server for NIS.**
**Cause:** The "Do Not Migrate (log only)" option was in effect.

**Solution:** Because migration is irreversible, the "Do Not Migrate (log only)" option is the default. You can override the default when running the wizard by selecting either **Overwrite** or **Preserve**. To override the default when you are running the command-line utility (Nis2ad.exe), specify the **-m** option on the command line.

**Migration conflicts occur even though test migration with "Do Not Migrate" option succeeded.**
**Cause:** If data in the Network Information Service (NIS) maps being migrated conflicts, the conflict might not be reported until the maps are actually migrated because only conflicts between NIS maps and Active Directory are reported.

**Solution:** Resolve the conflicts in the original NIS maps and then migrate the NIS data that was not migrated using the **nis2ad –r yes** option.

**A conflict occurs while trying to migrate more than one user with the same user identifier (UID).**
**Cause:** The NIS Data Migration wizard cannot migrate users that have the same UID. Only the first user will be successfully migrated.

**Solution:** Make sure all users have unique UIDs before migrating NIS maps.

**NIS Data Migration wizard or command-line utility reports Network Information Service (NIS) name conflicts with system account.**
**Cause:** A user name in NIS is identical to a reserved Windows account name. Reserved account names include such names as Network, System, Users, and Computers.

**Solution:** Resolve the conflict by renaming the UNIX account, and then remigrate the entry.

**The ypmatch utility on an HP-UX computer fails.**
**Cause:** A map contains keys in mixed case. HP-UX converts all keys to lowercase before sending the NIS request. This is an

issue with HP-UX that is beyond the control of Server for NIS.

**Solution:** Convert keys to lowercase before migrating them.

**The first request to Server for NIS fails.**
**Cause:** If the number of objects migrated to Server for NIS was very large, Server for NIS can take a long time to build the map cache.

**Solution:** Make the request again after waiting 30 minutes or more. Subsequent requests might succeed.

**Objects migrated to nonstandard containers do not appear in Windows Active Directory Users and Computers.**
**Cause:** Active Directory Users and Computers does not display nonstandard containers by default.

**Solution:** On the **View** menu, click **Advanced Features**.

**The ypcat utility sometimes displays incorrect passwd information, but ypmatch provides accurate information.**
**Cause:** Server for NIS data has not yet updated the map cache after data in Active Directory has changed. The **ypcat** utility takes its information from this cache; the **ypmatch** utility takes its information directly from Active Directory.

**Solution:** Decrease the refresh interval between updates to the map cache, or check that the updates have been cached immediately upon making changes in Active Directory. For more information, see Change the frequency of map updates to UNIX subordinate (slave) NIS servers and Propagate changed maps now.

**Users with new Windows accounts created by NIS migration cannot log on to Windows and cannot log on to NIS client computers.**
**Cause:** To avoid exposing new accounts to misuse, the migration wizard disables the new accounts.

**Solution:** After completing migration, enable new accounts only when users are ready to use them. For security reasons, it is recommended that you change the password on all newly created Windows user accounts to temporary values. Notify users of the temporary passwords and instruct them to change their Windows passwords as soon as possible. Inform them of how often map changes are propagated (see Sending periodic map updates to subordinate (slave) NIS servers) so they will know when they can expect their UNIX passwords to be updated.

**A user account disabled in Active Directory is not also disabled on UNIX hosts.**
**Cause:** Windows and UNIX accounts must be disabled separately because the mechanism for doing so is fundamentally different, and the method for disabling UNIX accounts differs based on system version and administrator preference.

**Solution:** After disabling an account using Active Directory Users and Groups, you must also disable or lock the corresponding UNIX account by the usual means (such as by changing the password to a value unknown to the user, by adding a special character to the user's password in the **passwd** file, by changing the account's login shell, or all three).

**A newly migrated user is not authenticated from a UNIX client.**
**Cause:** Possible configuration errors on the client computer.

**Solution:** Perform the following troubleshooting steps for the indicated client operating system.

**On a Solaris client:**

- Ensure that the Windows server that is running Server for NIS is added to the **/etc/hosts** file on the Solaris client. Then run the **ypwhich** command on the client and check that it shows the correct Windows server.

- Ensure that the **/etc/nsswitch.conf** file contains an entry similar to the following:

  **passwd: files [NOTFOUND=continue] nis**

- Ensure that the user is not listed in the local **passwd** file.

- Add the following to the end of the **/etc/passwd** file on the client computer:

  **+::::::**

- Ensure that shadow files, if any, are also migrated to Server for NIS.

- Ensure that the **ypcat password** command does not show **X** in the user's password field.

- Ensure that the file **/var/yp/securenets** is configured to allow authentication from the client computer.

**On a Linux client**

- Ensure that the Windows server running Server for NIS has been added to the **/etc/hosts** file on the client computer and that the **ypwhich** command on the client computer correctly shows the name of the Windows-based computer.

- Ensure that the **USENIS** entry in the **/etc/sysconfig/authconfig** is set to **USENIS=yes**.

- Ensure that the **nsswitch.conf** file is properly configured.

# Identity Management for UNIX: Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Password Synchronization makes it easy for users to maintain one user name and password for Windows domains and UNIX systems, synchronizing the passwords when one of them changes.

Depending on how Password Synchronization and the UNIX servers are configured, synchronization can be one way or two way.

- Before you begin using Password Synchronization, see Checklists for Password Synchronization.
- For information about how to implement Password Synchronization most effectively, see Best practices for Password Synchronization.
- For help with specific tasks, see How To .....
- For general background information, see Password Synchronization Concepts.
- For problem-solving instructions, see Troubleshooting Password Synchronization.
- For information about the Password Synchronization command-line utility, see Command-line utility.

# Checklists for Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Checklists**

- Checklist: Setting up Password Synchronization for use with standalone UNIX hosts
- Checklist: Setting up Password Synchronization for use with an NIS domain (UNIX master server)
- Checklist: Setting up Password Synchronization for use with an NIS domain (Server for NIS master server)

# Checklist: Setting up Password Synchronization for use with standalone UNIX hosts

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Checklist: Setting up Password Synchronization for use with standalone UNIX hosts**

**Configuring Windows systems**

| Step | Reference |
|---|---|
| Read about Password Synchronization. | Password Synchronization Concepts |
| Log on as a member of the Administrators group. | |
| Install Password Synchronization on the appropriate Windows-based computers. If the passwords of local accounts on a server are to be synchronized, install Password Synchronization on the server. If Windows domain passwords are to be synchronized, install Password Synchronization on all domain controllers. | Install Password Synchronization |
| Change the encryption key. | Set the default encryption key |
| Change other settings, as needed. | Configure Password Synchronization |
| Add UNIX computers with which passwords will be synchronized. If you want to use nondefault values, specify values for the port number, encryption key, or both. | Add and remove computers for synchronization |
| Ensure that the Password Synchronization configurations on all domain controllers in the domain are identical. | |

**Configuring UNIX systems**

| Step | Reference |
|---|---|
| Install and configure the Password Synchronization single sign-on daemon (SSOD) on all UNIX computers with which passwords will be synchronized. Be sure to change the default encryption key in the sso.conf file to match the Password Synchronization encryption key set in previous steps before copying it to the UNIX computers. | Install the Password Synchronization daemon |
| Specify which users will and will not be allowed to synchronize passwords. | Controlling password synchronization for user accounts |
| Start the Password Synchronization daemon. | Start the Password Synchronization daemon |
| Install and configure the Password Synchronization pluggable authentication module (PAM) on all UNIX computers from which password changes are to be synchronized with Windows passwords. | Configure UNIX Computers for UNIX-to-Windows Synchronization |

# Checklist: Setting up Password Synchronization for use with an NIS domain (UNIX master server)

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Checklist: Setting up Password Synchronization for use with an NIS domain (UNIX master server)**

**Configuring Windows systems**

| Step | Reference |
|---|---|
| Read about Password Synchronization. | Password Synchronization Concepts |
| Log on as a member of both the Schema Administrators and Enterprise Administrators groups. | |
| Install Password Synchronization on the appropriate Windows-based computers. If the passwords of local accounts on a server are to be synchronized, install Password Synchronization on the server. If Windows domain passwords are to be synchronized, install Password Synchronization on all domain controllers. | Install Password Synchronization |
| Change the encryption key. | Set the default encryption key |
| Change other settings, as needed. Be sure to select the **Synchronize password changes from computers that run UNIX to computers that run Windows** check box. | Configure Password Synchronization |
| Add the Network Information Service (NIS) master server to the list of computers with which the Windows-based computer will synchronize passwords. Select the NIS master server in the list, click **Configure**, select both the **Synchronize password changes to** and **Synchronize password changes from** check boxes, and then click **OK**. | Add and remove computers for synchronization |
| Add UNIX computers with which passwords will be synchronized. For each computer, select the computer in the list, click **Configure**, clear the **Synchronize password changes to** check box, select the **Synchronize password changes from** check box, and then click **OK**. If you want to use nondefault values, specify values for the port number, encryption key, or both. | Add and remove computers for synchronization |
| Specify which users will and will not be allowed to synchronize passwords. | Controlling password synchronization for user accounts |
| Ensure that the Password Synchronization configurations on all domain controllers in the domain are identical. | |

**Configuring UNIX systems**

| Step | Reference |
|---|---|
| Install and configure the Password Synchronization single sign-on daemon (SSOD) on the NIS master server. Be sure to change the default encryption key in the sso.conf file to match the Password Synchronization encryption key set in previous steps before copying it to the server, and edit it to specify the following:<br>• **USE_NIS=1**<br>• **NIS_UPDATE_PATH=** *Makefile_path*, where *Makefile_path* is the path and name of the NIS makefile, such as /var/yp/Makefile | |

| | |
|---|---|
| Install and configure the Password Synchronization pluggable authentication module (PAM) on all UNIX computers from which password changes are to be synchronized with Windows passwords. Typically, this would be any computer on which users would run **yppasswd** and any standalone computers (computers that do not belong to the domain). | Configure UNIX Computers for UNIX-to-Windows Synchronization |
| Copy the sso.conf file from the NIS master server to the /etc directory of each computer on which the Password Synchronization PAM module is installed. | |
| On each NIS client on which you installed the Password Synchronization PAM module, replace the yppasswd binary file with a link to the passwd binary file, and then edit the /etc/nsswitch.conf file to change the passwd and shadow lines, as shown:<br>`passwd:  files [NOTFOUND=continue] nis`<br>`shadow:  files [NOTFOUND=continue] nis` | |
| Start the Password Synchronization daemon on the NIS master server. | Start the Password Synchronization daemon |

# Checklist: Setting up Password Synchronization for use with an NIS domain (Server for NIS master server)

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Checklist: Setting up Password Synchronization for use with an NIS domain (Server for NIS master server)**
**Configuring Windows systems**

| Step | Reference |
|---|---|
| Read about Password Synchronization. | Password Synchronization Concepts |
| Log on as a member of both the Schema Administrators and Enterprise Administrators groups. | |
| Install Password Synchronization on all domain controllers. | Install Password Synchronization |
| Change the encryption key. | Set the default encryption key |
| Change other settings, as needed. Be sure to select the **Synchronize password changes from computers that run UNIX to computers that run Windows** check box. | Configure Password Synchronization |
| Add UNIX computers with which passwords will be synchronized if they are not members of the Network Information Service (NIS) domain. For each computer, select the computer in the list, click **Configure**, clear the **Synchronize password changes to** check box, select the **Synchronize password changes from** check box, and then click **OK**. If you want to use nondefault values, you can also specify values for the port number, encryption key, or both. | Add and remove computers for synchronization |
| Ensure that the Password Synchronization configurations on all domain controllers in the domain are identical. | |
| Specify which users will and will not be allowed to synchronize passwords. | Controlling password synchronization for user accounts |

**Configuring UNIX systems**

| Step | Reference |
|---|---|
| Install and configure the Password Synchronization pluggable authentication module (PAM) on all UNIX computers from which password changes are to be synchronized with Windows passwords. Typically, this would be any computer on which users would run **yppasswd** and any computers that do not belong to the NIS domain. | Configure UNIX Computers for UNIX-to-Windows Synchronization |

| | |
|---|---|
| Using a binary file-copy method such as File Transfer Protocol (FTP) to avoid corrupting CR/LF (carriage-return/line-feed) pairs, copy Sso.cfg from IDMU\Unix\Bins on the Windows Server 2003 R2 CD to the /etc directory of each computer on which the Password Synchronization PAM module is installed, and change its name to sso.conf. Open the file with a text editor, and then perform the following steps:<br><br>1. Edit the following line to specify the encryption key. This value must match the default key specified on all domain controllers with which this computer will synchronize passwords:<br><br>**ENCRYPT_KEY=***encryptionKey*<br><br>2. If you have changed the default port, edit the following line to specify the new port. This value must match the port number specified on all domain controllers with which this computer will synchronize passwords.<br><br>**PORT_NUMBER=***portNumber*<br><br>3. Edit the following line to specify the computer running Server for NIS. If you have specified a nondefault port number or encryption key for the UNIX computer when configuring Password Synchronization on the Windows domain controller, specify that value where indicated; otherwise, leave the value blank:<br><br>**SYNC_HOSTS=(***domainController*[**,** *portNumber* [**,** *encryptionKey*]]**)** | |
| On each NIS client on which you installed the Password Synchronization PAM module, replace the yppasswd binary file with a link to the passwd binary file, and then edit the /etc/nsswitch.conf file to change the passwd and shadow lines, as shown:<br><br>`passwd:  files [NOTFOUND=continue] nis`<br>`shadow:  files [NOTFOUND=continue] nis` | |

# Best practices for Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Best practices**

- **Install Password Synchronization on appropriate domain controllers** To ensure consistent synchronization of domain passwords with UNIX passwords, Password Synchronization must be installed on the primary domain controller and, in the case of a Windows 2000 domain, all domain controllers in a domain.

    - If you add a domain controller to a domain, you should install Password Synchronization on the new domain controller as soon as possible and configure it to match the other domain controllers.

    - If you need to remove Password Synchronization from any domain controller, you should demote the server to a member server before uninstalling Password Synchronization.

- **Ensure consistent password policies** If you are providing only for one-way password synchronization, make sure that the password policy on the computer from which passwords will be synchronized is at least as restrictive in all areas as the policy on the computer to which passwords will by synchronized. For example, if you configure Windows-to-UNIX synchronization, the Windows password policy must be at least as restrictive as the policy of the UNIX computers with which it will synchronize passwords. If you are supporting two-way synchronization, the password policies must be equally restrictive on both systems. Failure to ensure that password policies are consistent can result in synchronization failure when a user changes a password on the less restrictive system, or the password might be changed on the more restrictive system even though it does not conform to the system's policies.

  Also make sure that Windows users are aware of any special password restrictions on the UNIX systems with which their passwords will be synchronized. For example, some versions of UNIX support a maximum password length of eight characters. For maximum compatibility with the default Windows password policy and these UNIX limitations, passwords should be seven or eight characters long unless you are sure that all UNIX systems can support longer passwords.

- **Configure Password Synchronization to provide the maximum protection for your users' passwords**

  Follow these recommendations to maintain optimal security:

    - **Explicitly list the users whose passwords are to be synchronized** To provide maximum control over which users can synchronize passwords, do not use the ALL keyword with the SYNC_USERS list in sso.conf on the UNIX host. Instead, you should explicitly list each user for whom password synchronization is allowed or blocked. On the Windows-based computer running Password Synchronization, create the PasswordPropAllow group and add the accounts of users whose passwords you want to synchronize. For more information, see Controlling password synchronization for user accounts.

    - **Do not synchronize passwords for disabled UNIX accounts** On some versions of UNIX, changing the password of a disabled user account activates that account. Consequently, if a user has a disabled account on a UNIX computer that is configured to synchronize passwords with a Windows-based computer, the user or an administrator can activate the UNIX account by changing the user's Windows password. To prevent this, use the PasswordPropDeny group to block synchronization for disabled UNIX accounts.

      Also, when an administrator disables a UNIX account, the administrator should use the SYNC_USERS entry in sso.conf to block password synchronization for the account.

    - **Avoid synchronizing administrator passwords** Do not synchronize passwords for members of the Windows Administrators groups or the passwords of UNIX superuser or root accounts.

    - **Perform the Windows Server 2003 Service Pack 1 (SP1) compatibility check** when selecting **Enable Windows to NIS (AD) Password Sync** in the **Password Synchronization Properties** dialog box, **Configuration** tab. To protect the security of user account passwords in your enterprise, it is strongly

recommended that you allow Password Synchronization to identify all domain controllers in the forest that are not running Windows Server 2003 SP1 or later releases.

Password Synchronization prompts you to allow the compatibility check when you select the **Enable Windows to NIS (AD) Password Sync** option. With Windows Server 2003 SP1, or a later release installed on all the domain controllers in a forest, the risk of exposing user password hashes to unauthorized viewers is greatly reduced. When Windows Server 2003 SP1 is not the minimum functional level of all the domain controllers in a forest, it is possible for any authenticated user on the domain to view the password hash for any UNIX user whose account has been migrated to Active Directory.

In the event an unauthorized user breaks the password hash for a UNIX-based user account in Active Directory, the Windows-based password for the account is no longer secure.

When Password Synchronization is installed, members of the local Administrators group and the Domain Admins group are added to the PasswordPropDeny group, which prevents their passwords from being synchronized. If you add a user to either the Administrators or Domain Admins group, be sure to add the user to the PasswordPropDeny group as well.

Use the SYNC_USERS statement in the sso.conf file on all UNIX systems to prevent the passwords of superusers from being synchronized.

- **Do not use the default port number and encryption key** If you use the default port number and encryption key, you make it possible for an attacker to set up an impostor UNIX host to capture passwords. You should protect the port number and encryption keys used to synchronize passwords as carefully as the passwords themselves.

- **Secure the sso.conf file** The sso.conf file on each UNIX host contains important configuration information that could be used to compromise security. It is recommended that you set the mode bit mask of the file to 600.

- **Ensure that the directory identified by TEMP_FILE_PATH is properly protected** The temporary files created on UNIX hosts by Password Synchronization contain information that could be used by an attacker to compromise system security. For this reason, you should ensure that any directory referenced by TEMP_FILE_PATH in sso.conf has read access only for the root account and cannot be accessed by any other users.

- **Ensure log files are appropriately protected** On the UNIX host, Password Synchronization uses the syslogd daemon to log messages that result from synchronization operations. The resulting logs contain such information as the names of users whose passwords are being synchronized and with which computers, propagation errors, and so on. These log files should be protected to ensure that only administrators can read them.

- **Restart the Password Synchronization daemon after changing its configuration** When you make changes to the Password Synchronization daemon configuration file (sso.conf), you must stop and restart the daemon to make sure that the configuration changes take effect.

- **Configure systems to handle user name case sensitivity correctly** Unless you rigorously enforce a policy to ensure that Windows and UNIX user names match in both spelling and case, make sure that the CASE_IGNORE_NAME option in the sso.conf file is set to 1 (the default). UNIX user names are case sensitive; therefore, passwords might not synchronize properly if the user names do not match exactly because the Password Synchronization daemon is unable to associate the Windows user name with the corresponding UNIX user name.

- **Make sure that password file type and name are consistent** When you configure the Password Synchronization daemon, make sure that the password file type (specified by USE_SHADOW) and path name (set by FILE_PATH) are appropriate for each other. For example, on most systems, if USE_SHADOW is set to 0 (to indicate that the passwd file is used for synchronization), then the FILE_PATH option should be set to /etc/passwd. However, if USE_SHADOW is set to 1 (to indicate that the shadow file is used instead), then the FILE_PATH option should be set to /etc/shadow. (On AIX systems, the path and name of the shadow file is /etc/security/passwd.)

# How To ....

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

- Use UNIX Components
- Use Windows Components

# Use UNIX Components

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Use UNIX Components**

- Install the Password Synchronization daemon
- Start the Password Synchronization daemon
- Stop the Password Synchronization daemon
- Configure UNIX Computers for UNIX-to-Windows Synchronization

# Install the Password Synchronization daemon

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To install the Password Synchronization daemon**

1. Copy the appropriate source binary file from IDMU\Unix\Bins on the Windows Server 2003 R2 CD to /usr/bin or /usr/local/bin on the UNIX computer, and change its name to ssod. The name of the source binary file depends on the version of UNIX you are using.

   - If the computer is running Hewlett-Packard HP-UX, the source binary file name is ssod.hpx.
   - If the computer is running Red Hat Linux, the source binary file name is ssod.rhl.
   - If the computer is running Sun Microsystems Solaris, the source binary file name is ssod.sol.
   - If the computer is running IBM AIX, the source binary file name is ssod.aix.

2. Using a binary file-copy method such as File Transfer Protocol (FTP) to avoid corrupting CR/LF (carriage-return/line-feed) pairs, copy Sso.cfg from \Unix\Bins on the Windows Server 2003 R2 CD to /etc on the UNIX computer, and change its name to sso.conf.

3. Open sso.conf with a text editor.

4. If you have changed the default encryption key, edit the following line to specify the new default key. This value must match the default key specified on all domain controllers with which this computer will synchronize passwords:

   **ENCRYPT_KEY=***encryptionKey*

5. If you have changed the default port, edit the following line to specify the new port. This value must match the port number specified on all domain controllers with which this computer will synchronize passwords.

   **PORT_NUMBER=***portNumber*

6. Edit the following line to specify one domain controller in each Windows domain with which the computer is to synchronize passwords. If you have specified a nondefault port number or encryption key for the UNIX computer when configuring Password Synchronization on the Windows domain controllers, specify that value where indicated; otherwise, leave the value blank:

   **SYNC_HOSTS=(***domainController*[**,** *portNumber* [**,** *encryptionKey*]]**)** ...

   Each entry in the list must be enclosed by parentheses (the "(" and ")" characters) and separated from the next entry by a blank space.

7. If the computer is a Network Information Service (NIS) master server, and if you want passwords to be synchronized throughout the NIS domain, edit the following line as shown to enable NIS synchronization:

   **USE_NIS=1**

   Also, if required, edit the following line to specify the location of the NIS makefile:

   **NIS_UPDATE_PATH=***makefilePath*

8. Set the file permissions of sso.conf to read/write for the root user only, and deny access to all other users.

9. If the computer is running Linux, copy /etc/pam.d/system-auth to /etc/pam.d/ssod.

**Important**

The sso.conf file contains encryption keys and other sensitive information. For this reason, it must be accessible only by system administrators.

**Note**

This daemon program must be installed on the computer running UNIX to enable Password Synchronization to change users' passwords on that computer.

Password Synchronization supports synchronization with UNIX computers running any of the following operating systems:

- Hewlett-Packard HP-UX version 11i
- IBM AIX version 5L 5.2
- Red Hat Linux versions 8 and 9
- Sun Solaris version 8 running on x86-based computers and Scalable Processor Architecture (SPARC)–based computers, and Solaris version 9 running on SPARC–based computers

**See Also**

**Concepts**

Understanding Password Synchronization
Using sso.conf to configure Password Synchronization on the UNIX computer
Start the Password Synchronization daemon
Install the pluggable authentication module (PAM) on AIX
Install the pluggable authentication module (PAM) on HP-UX
Install the pluggable authentication module (PAM) on Linux
Install the pluggable authentication module (PAM) on Solaris

# Start the Password Synchronization daemon

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

- Run the single sign-on daemon (SSOD).

| Note |
| --- |
| There is no other method for this procedure.<br><br>The sso.conf file must be in the /etc/ directory.<br><br>If you see no errors, the daemon is running. If anything goes wrong during the initialization of the daemon, an error message is displayed on the console. Some errors might also be logged to the Syslog file.<br><br>If you want your service to start automatically, add the single sign-on daemon (SSOD) to your system startup files (typically Rc.local). This will start SSOD as root. |

**See Also**
**Concepts**
Install the Password Synchronization daemon
Understanding Password Synchronization

# Stop the Password Synchronization daemon

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

- Send the process a TERM signal as follows:

  **kill –term** *pid*

  where *pid* is the process identifier of the Password Synchronization single sign-on daemon (SSOD).

**See Also**
**Concepts**
Understanding Password Synchronization

# Configure UNIX Computers for UNIX-to-Windows Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Configure UNIX computers for UNIX-to-Windows synchronization**

- Install the pluggable authentication module (PAM) on AIX
- Install the pluggable authentication module (PAM) on HP-UX
- Install the pluggable authentication module (PAM) on Linux
- Install the pluggable authentication module (PAM) on Solaris

# Install the pluggable authentication module (PAM) on AIX

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To install the pluggable authentication module (PAM) on AIX**

1. Copy pam_sso.aix from IDMU\Unix\Bins on the Windows Server 2003 R2 CD to /usr/lib/ on the UNIX computer, and change its name to pam_sso.aix.1.

2. On the UNIX computer, log on as root, and then enter the following commands:

   **chown root /usr/lib/pam_sso.aix.1 chmod 555 /usr/lib/pam_sso.aix.1**

3. If necessary, create the /etc/pam.conf file according to your network requirements, setting the owner to root and the base permissions to 644. For more information about creating the pam.conf file, see "Pluggable Authentication Modules" in *System Management Guides: Security Guide* in your AIX documentation.

   Sample pam.conf file

```
#
# Authentication management
#
OTHER   auth     required       /usr/lib/security/pam_aix

#
# Account management
#
OTHER   account  required       /usr/lib/security/pam_aix

#
# Session management
#
OTHER   session  required       /usr/lib/security/pam_aix
```

4. Open /etc/pam.conf with a text editor.

5. In the **Password management** section, add the following line:

   **passwd password required /usr/lib/security/pam_sso.aix.1**

   Sample pam.conf file with this line added

```
#
# Authentication management
#
OTHER   auth     required       /usr/lib/security/pam_aix

#
# Account management
#
```

```
OTHER   account  required       /usr/lib/security/pam_aix

#
# Session management
#
OTHER   session  required       /usr/lib/security/pam_aix

#
# Password management
#
passwd   password required       /usr/lib/security/pam_sso.aix.1
```

6. Open /usr/lib/security/methods.cfg with a text editor and add the following lines at the end of the file:

   **PAM:  program = /usr/lib/security/PAM**

   **PAMfiles:  options = auth=PAM,db=BUILTIN**

7. Open /etc/security/user with a text editor and add authentication information for the specific users whose passwords you want to synchronize. For example:

   user1:  admin = false  SYSTEM = PAMfiles[*] AND "compat"  registry = PAMfiles

| **Note** |
| --- |
| You can choose to change the default section of /etc/security/user to allow all users to synchronize their passwords. In this case, to restrict access to Password Synchronization, you can use the SYNC_USERS attribute in the /etc/sso.conf file to restrict access. For more information, see Using sso.conf to configure Password Synchronization on the UNIX computer.<br><br>To disable UNIX-to-Windows password synchronization, remove the entry in /etc/pam.conf that you added in step 5. |

**See Also**
**Concepts**
Understanding Password Synchronization
Implementing Password Synchronization

# Install the pluggable authentication module (PAM) on HP-UX

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To install the pluggable authentication module (PAM) on HP-UX**

1. Copy pam_sso.hpx from IDMU\Unix\Bins on the Windows Server 2003 R2 CD to /usr/lib/security on the UNIX computer, change its name to pam_sso.hp.1, and then set its file-mode bits to 544.

2. On the UNIX computer, open /etc/pam.conf with a text editor.

3. In the **Password management** section, locate the following line:

   **other password required /usr/lib/security/libpam_unix.1**

4. Immediately following the line located in the previous step, add the following line:

   **other password required /usr/lib/security/pam_sso.hp.1**

| **Note** |
| --- |
| To disable UNIX-to-Windows password synchronization, remove the entry in /etc/pam.conf that you added in step 4.<br><br>Before installing the pam_sso module, make sure that PAM support is properly installed and configured on the UNIX computer. The following file samples show a typical configuration. Actual contents of these files may differ, depending on your system configuration.<br><br>Sample HP-UX PAM configuration file |

```
#
# PAM configuration
#
# Authentication management
#
login    auth required  /usr/lib/security/libpam_unix.1
su       auth required  /usr/lib/security/libpam_unix.1
dtlogin  auth required  /usr/lib/security/libpam_unix.1
dtaction auth required  /usr/lib/security/libpam_unix.1
ftp      auth required  /usr/lib/security/libpam_unix.1
OTHER    auth required  /usr/lib/security/libpam_unix.1
#
# Account management
#
login    account required       /usr/lib/security/libpam_unix.1
su       account required       /usr/lib/security/libpam_unix.1
dtlogin  account required       /usr/lib/security/libpam_unix.1
dtaction account required       /usr/lib/security/libpam_unix.1
ftp      account required       /usr/lib/security/libpam_unix.1
#
OTHER    account required       /usr/lib/security/libpam_unix.1
#
# Session management
#
login    session required       /usr/lib/security/libpam_unix.1
dtlogin  session required       /usr/lib/security/libpam_unix.1
dtaction session required       /usr/lib/security/libpam_unix.1
OTHER    session required       /usr/lib/security/libpam_unix.1
#
# Password management
#
login    password required      /usr/lib/security/libpam_unix.1
dtlogin  password required      /usr/lib/security/libpam_unix.1
```

```
dtaction password required      /usr/lib/security/libpam_unix.1
other    password required      /usr/lib/security/libpam_unix.1
other    password required      /usr/lib/security/pam_sso.hp.1
```

The file-mode bits for pam_sso.hp.1 must be set to 544 (o:r-x,g:r--,w:r--) or it will not function properly.

**See Also**
**Concepts**
Understanding Password Synchronization
Implementing Password Synchronization

# Install the pluggable authentication module (PAM) on Linux

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To install the pluggable authentication module (PAM) on Linux**

1. Copy pam_sso.rhl from IDMU\Unix\Bins on the Windows Server 2003 R2 CD to /lib/security on the UNIX computer, and change its name to pam_sso.so.1.

2. On the UNIX computer, copy /etc/pam.d/system-auth to /etc/pam.d/ssod.

3. Open /etc/pam.d/system-auth with a text editor, and locate the following line:

   **passwordrequired/lib/security/pam_cracklib.soretry=3**

4. After the line in the previous step, add the following line:

   **password required /lib/security/pam_sso.so.1**

5. Locate and delete the following line:

   **passwordrequired/lib/security/pam_deny.so**

6. Save the modified file.

| **Note** |
| --- |
| These instructions apply to the typical Linux configuration. If you have configured PAM support differently, you might have to adjust these instructions to your specific configuration.<br><br>To disable UNIX-to-Windows password synchronization, remove the entry in /etc/pam.d/system-auth that you added in step 4.<br><br>Before installing the pam_sso module, make sure that PAM support is properly installed and configured on the UNIX computer. The following file samples show a typical configuration. Actual contents of these files may differ, depending on your system configuration.<br><br>Sample Linux PAM configuration file: |

```
/etc/pam.d/passwd

#%PAM-1.0
auth       required     /lib/security/pam_stack.so service=system-auth
account    required     /lib/security/pam_stack.so service=system-auth
password   required     /lib/security/pam_stack.so service=system-auth


/etc/pam.d/ssod

#%PAM-1.0
# This file is auto-generated.
# User changes will be destroyed the next time authconfig is run.
auth        required      /lib/security/pam_env.so
auth        sufficient    /lib/security/pam_unix.so likeauth nullok
auth        required      /lib/security/pam_deny.so

account     required      /lib/security/pam_unix.so

password    required      /lib/security/pam_cracklib.so retry=3 type=
password    sufficient    /lib/security/pam_unix.so nullok use_authtok shadow
```

```
password    required      /lib/security/pam_deny.so

session     required      /lib/security/pam_limits.so
session     required      /lib/security/pam_unix.so
```

```
/etc/pam.d/system-auth

#%PAM-1.0
# This file is auto-generated.
# User changes will be destroyed the next time authconfig is run.
auth        required      /lib/security/pam_env.so
auth        sufficient    /lib/security/pam_unix.so likeauth nullok
auth        required      /lib/security/pam_deny.so

account     required      /lib/security/pam_unix.so

password    required      /lib/security/pam_cracklib.so retry=3 type=
password    required      /lib/security/pam_sso.so.1
password    sufficient    /lib/security/pam_unix.so nullok use_authtok shadow
password    required      /lib/security/pam_deny.so

session     required      /lib/security/pam_limits.so
session     required      /lib/security/pam_unix.so
```

**See Also**

**Concepts**

Understanding Password Synchronization

Implementing Password Synchronization

# Install the pluggable authentication module (PAM) on Solaris

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To install the pluggable authentication module (PAM) on Solaris**

1. Copy pam_sso.sol from the IDMU\Unix\Bins folder on the Windows Server 2003 R2 CD to the /usr/lib/security directory on the UNIX computer, and change its name to pam_sso.so.1.

2. On the UNIX computer, open /etc/pam.conf with a text editor.

3. In the **Password management** section, locate the following line:

   **other password required /usr/lib/security/$ISA/pam_unix.so.1**

4. Immediately following the line located in the previous step, add the following line:

   **other password required /usr/lib/security/$ISA/pam_sso.so.1**

| **Note** |
| --- |
| To disable UNIX-to-Windows password synchronization, remove the entry in /etc/pam.conf that you added in step 4.<br><br>Before installing the pam_sso module, make sure that PAM support is properly installed and configured on the UNIX computer. The following file samples show a typical configuration. Actual contents of these files may differ, depending on your system configuration.<br><br>Sample Solaris PAM configuration file |

```
#
#ident  "@(#)pam.conf   1.14    99/09/16 SMI"
#
# Copyright (c) 1996-1999, Sun Microsystems, Inc.
# All Rights Reserved.
#
# PAM configuration
#
# Authentication management
#
login   auth required   /usr/lib/security/$ISA/pam_unix.so.1
login   auth required   /usr/lib/security/$ISA/pam_dial_auth.so.1
#
rlogin  auth sufficient /usr/lib/security/$ISA/pam_rhosts_auth.so.1
rlogin  auth required   /usr/lib/security/$ISA/pam_unix.so.1
#
dtlogin auth required   /usr/lib/security/$ISA/pam_unix.so.1
#
rsh     auth required   /usr/lib/security/$ISA/pam_rhosts_auth.so.1
other   auth required   /usr/lib/security/$ISA/pam_unix.so.1
#
# Account management
#
login   account requisite       /usr/lib/security/$ISA/pam_roles.so.1
login   account required        /usr/lib/security/$ISA/pam_unix.so.1
#
dtlogin account requisite       /usr/lib/security/$ISA/pam_roles.so.1
dtlogin account required        /usr/lib/security/$ISA/pam_unix.so.1
#
other   account requisite       /usr/lib/security/$ISA/pam_roles.so.1
other   account required        /usr/lib/security/$ISA/pam_unix.so.1
#
# Session management
```

```
#
other   session required        /usr/lib/security/$ISA/pam_unix.so.1
#
# Password management
#

other   password required       /usr/lib/security/$ISA/pam_unix.so.1
other  password required        /usr/lib/security/$ISA/pam_sso.so.1
dtsession auth required /usr/lib/security/$ISA/pam_unix.so.1

#
# Support for Kerberos V5 authentication (uncomment to use Kerberos)
#
#rlogin auth optional   /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
#login  auth optional   /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
#dtlogin        auth optional   /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
#other  auth optional   /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
#dtlogin        account optional /usr/lib/security/$ISA/pam_krb5.so.1
#other  account optional /usr/lib/security/$ISA/pam_krb5.so.1
#other  session optional /usr/lib/security/$ISA/pam_krb5.so.1
#other  password optional /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
```

**See Also**

**Concepts**

Understanding Password Synchronization

Implementing Password Synchronization

# Use Windows Components

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Use Windows components**

- Install Password Synchronization
- Configure Password Synchronization

# Install Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To install Password Synchronization**

1. Click **Start**, click **Control Panel**, and then click **Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. When the Windows Components Wizard opens, fill the **Active Directory Services** check box.
4. With **Active Directory Services** highlighted, click **Details**.
5. When the **Active Directory Services** dialog box opens, fill the **Identity Management for UNIX** check box.
6. With **Identity Management for UNIX** highlighted, click **Details**.
7. When the **Identity Management for UNIX** dialog box opens, fill the **Password Synchronization** check box.
8. Click **OK**.

   The Windows Components Wizard begins installing the Identity Management for UNIX components you selected.
9. If you are prompted to locate files required for Password Synchronization installation, insert the Windows Server 2003 R2 product CD, or browse to the network location of your Windows Server 2003 R2 installation files.

| **Note** |
| --- |
| To open **Add or Remove Programs**, click **Start**, point to **Control Panel**, and then click **Add or Remove Programs**.<br><br>Password Synchronization cannot be run from a network server. All files must be installed on the local computer.<br><br>You must be logged on to the computer as a member of both the Schema Administrators and Enterprise Administrators groups to install Password Synchronization. The Password Synchronization installation process verifies that users have both Schema Administrator and Enterprise Administrator access permissions before installation can progress. |

**See Also**
**Concepts**
Install the Password Synchronization daemon

# Configure Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Configure Password Synchronization**


- Set default synchronization
- Set the default encryption key
- Set the default port
- Add and remove computers for synchronization
- Set computer-specific synchronization properties
- Set auditing options

# Set default synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To set default synchronization**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**, and then click the **Default** tab.
4. To allow password synchronization from UNIX-based computers to Windows-based computers, click **Synchronize password changes from computers that run UNIX to computers that run Windows.**
5. To allow password synchronization from Windows-based computers to UNIX-based computers, click **Synchronize password changes from computers that run Windows to computers that run UNIX.**
6. To save the new settings, click **Apply**.

**Note**

To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.

To perform this task in the command line environment, see psadmin.

These settings affect the default synchronization for UNIX hosts when they are added for synchronization. It does not affect computers that have already been added for synchronization.

**See Also**

**Concepts**

Add and remove computers for synchronization

Password Synchronization administration

**Other Resources**

Connect to a computer you want to manage

# Set the default encryption key

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To set the default encryption key**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**, and then click the **Default** tab.
4. In the **Encryption key** text box, type the key you want to use, or, to have the program produce a key for you, click **New Key**.
5. To save the new settings, click **Apply**.

**Note**

To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.

To perform this task in the command line environment, see psadmin.

This setting affects the default encryption key for UNIX hosts when they are added for synchronization as well as the port used for UNIX-to-Windows synchronization. If you change this setting, you must edit the /etc/sso.conf file to specify the same encryption key on UNIX hosts that are configured for UNIX-to-Windows password synchronization with this computer.

For maximum security, you should use a key that is the maximum 21 characters in length.

**See Also**

**Concepts**

Encryption key requirements
Add and remove computers for synchronization
Password encryption
Password Synchronization administration
Using sso.conf to configure Password Synchronization on the UNIX computer

**Other Resources**

Connect to the computer you want to manage

# Set the default port

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To set the default port**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**, and then click the **Default** tab.
4. To use a port other than 6677, in the **Port number** box, type the port number you want.
5. To save the new settings, click **Apply**.

**Note**

To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.

This setting affects the default port number for UNIX hosts when they are added for synchronization as well as the port used for UNIX-to-Windows synchronization. If you change this setting, you must edit the /etc/sso.conf file to specify the same port on UNIX hosts that are configured for UNIX-to-Windows password synchronization with this computer.

For maximum security, use a port number other than the default (6677).

To perform this task in the command line environment, see psadmin.

**See Also**

**Concepts**

Add and remove computers for synchronization
Password Synchronization administration
Using sso.conf to configure Password Synchronization on the UNIX computer

**Other Resources**

Connect to a computer you want to manage

# Add and remove computers for synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To add and remove computers for synchronization**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**.
4. Click the **Advanced** tab, and then do one of the following:
   - To add a computer to the list of participating computers, in **Computer name**, type the name of the UNIX-based computer you want to add, and then click **Add**.
   - To remove a computer, in the **Current computers** list, click the UNIX-based computer you want to remove, and then click **Remove**.
5. To save the new settings, click **Apply**.

**Note**

To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.

To perform this task in the command line environment, see psadmin.

In addition to adding a UNIX-based computer to the list, if you want to change the user's password on the UNIX computer when the corresponding Windows user's password is changed, you must install the Password Synchronization single sign-on daemon (SSOD) on the UNIX-based computer . If you want to change the Windows user's password when the corresponding UNIX-based computer user's password is changed, you must install the pluggable authentication module (PAM) on the UNIX-based computer.

**See Also**

**Concepts**

Set computer-specific synchronization properties
Install the Password Synchronization daemon
Configure UNIX Computers for UNIX-to-Windows Synchronization
Password Synchronization administration

**Other Resources**

Connect to a computer you want to manage

# Set computer-specific synchronization properties

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To set computer-specific synchronization properties**

1. Open **Identity Management for UNIX**.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**, and then click the **Advanced** tab.
4. In the **Current computers** list, click the one for which you want to set properties, and then click **Configure**.
5. Set the properties you want to apply to the selected computer.
6. To save the new settings, click **Apply**.

| **Note** |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.<br><br>To perform this task in the command line environment, see psadmin.<br><br>Before setting computer-specific synchronization properties in Password Synchronization, edit the sso.conf file on the UNIX-based computer to specify the same settings. For information about assigning computer-specific settings in the sso.conf file, see Install the Password Synchronization daemon. |

**See Also**

**Concepts**

Set default synchronization
Set the default encryption key
Set the default port
Password Synchronization administration

**Other Resources**

Connect to a computer you want to manage

# Set auditing options

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**To set auditing options**

1. Open Identity Management for UNIX.
2. If necessary, connect to the computer you want to manage.
3. Click **Password Synchronization**, and then click the **Default** tab.
4. To have information events as well as warnings and errors logged to Event Viewer, click **Enable extensive logging**.
5. To save the new settings, click **Apply**.

| **Note** |
| --- |
| To open Identity Management for UNIX, click **Start**, point to **All Programs**, point to **Identity Management for UNIX**, and then click **Identity Management for UNIX**.<br><br>To perform this task in the command line environment, see psadmin.<br><br>When the log file reaches the limit you set, logging stops and a note is entered in Event Viewer. |

**See Also**
**Concepts**
Password Synchronization administration
**Other Resources**
Connect to a computer you want to manage

# Password Synchronization Concepts

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

This section provides background information about:

- Understanding Password Synchronization
- Administering Password Synchronization

# Understanding Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Understanding Password Synchronization**
Password Synchronization helps integrate Windows and UNIX networks by simplifying the process of maintaining secure passwords in both environments. Users are freed of the difficulty of maintaining separate passwords for their Windows and UNIX accounts or having to remember to change the password wherever it is used. With Password Synchronization, whenever a user's password is changed on a Windows-based computer or domain, the password can also be automatically changed on every UNIX host for which the user has an account. Password Synchronization can also be configured to change the user's Windows password when the user's UNIX password is changed.

This makes it possible to administer passwords from a single computer, which simplifies things for administrators as well as individual users. Password Synchronization is also flexible: administrators can exclude specific users and systems from being synchronized. Password Synchronization can synchronize passwords on stand-alone Windows-based computers (such as computers running Windows 2000 Server that do not belong to a domain) or for an entire Windows domain. Similarly, Password Synchronization can be used to manage passwords on individual UNIX hosts or on all computers in a Network Information Service (NIS) domain.

Password Synchronization propagates passwords securely by transmitting only encrypted passwords over TCP/IP sockets. This eliminates the need to use nonsecure methods (such as scripts) to administer passwords remotely. Passwords are also synchronized immediately. This means that, unlike methods such as **rdist**, which batches the password propagation, there is no appreciable delay between the time that a password is changed on one system and when it is changed on all other affected systems. This eliminates confusion and frustration for active users. Even more important, perhaps, it eliminates a potential security risk if a password needs to be changed to block a user's access to the network. To enhance network security even further, different encryption keys can be used for each Windows-based computer/UNIX host pair.

Password Synchronization is a combination of three software components:

- The Password Synchronization service running on one or more Windows-based computers
- The Password Synchronization daemon running on one or more UNIX computers
- The Password Synchronization pluggable authentication module (PAM) installed on one or more UNIX computers.

When Password Synchronization is configured for Windows-to-UNIX synchronization and a password is changed on a Windows-based computer running Password Synchronization, the Password Synchronization service determines whether the user's password is to be synchronized on UNIX computers. If it is, the service encrypts the password and sends it to the Password Synchronization daemon on each computer with which the Windows-based computer is configured to be synchronized. The daemon then decrypts the password and changes the password on the UNIX host. If the UNIX host is an NIS master server and it is configured to do so, the daemon also runs **make** to propagate the password change throughout the NIS domain.

When Password Synchronization is configured for UNIX-to-WINDOWS synchronization, passwords that are changed on UNIX hosts are synchronized on Windows-based computers and domains. The Password Synchronization PAM module makes this possible by intercepting the password change request on the UNIX host, encrypting the password, and then sending the password change request to the Password Synchronization service running on the Windows-based computers with which it is configured to be synchronized.

For information about implementing Password Synchronization, see Implementing Password Synchronization.

# Administering Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Administering Password Synchronization**
This section contains information about:

- Implementing Password Synchronization
- Synchronizing passwords with an NIS domain
- Password Synchronization administration
- Password encryption
- Encryption key requirements
- Controlling password synchronization for user accounts
- Using sso.conf to configure Password Synchronization on the UNIX computer

# Implementing Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Implementing Password Synchronization**

Password Synchronization changes a user's UNIX password whenever the user's password is changed on a Windows-based computer or domain. In addition, Password Synchronization and the affected UNIX hosts can be configured to change the user's Windows password whenever the UNIX password is changed.

Two-way password synchronization is supported on UNIX computers running any of the following operating systems:

- Hewlett-Packard HP-UX version 11i, 32-bit environment
- IBM AIX version 5L 5.2, 32-bit environment
- Red Hat Linux versions 8 and 9, 32-bit environment
- Sun Solaris version 8 running on x86-based computers and Scalable Processor Architecture (SPARC)–based computers, and Solaris version 9 running on SPARC–based computers. The 32-bit environment has been tested on SPARC-based computers.

For more information about how Password Synchronization works on Windows and UNIX computers, see Understanding Password Synchronization.

The remainder of this topic provides information about how to implement password synchronization on Windows and UNIX computers.

**Installing Password Synchronization in Windows**

You can use Password Synchronization to synchronize passwords between a Windows domain and one or more UNIX hosts, or you can use it to synchronize passwords between a stand-alone computer running Windows and one or more UNIX hosts.

To synchronize local account passwords on a Windows-based computer, install Password Synchronization on that computer only. To synchronize Windows domain passwords, you must install Password Synchronization on all domain controllers for that domain. This will ensure that when a domain controller processes a password change request, the Password Synchronization service on that domain controller will be able to synchronize the new password with the appropriate UNIX hosts. For this reason, before you remove Password Synchronization from a domain controller, you should demote the domain controller to a member server to prevent password discrepancies between the Windows domain and the UNIX hosts.

For information about installing Password Synchronization on Windows-based computers, see Install Password Synchronization.

**Installing the Password Synchronization daemon on UNIX hosts**

To allow synchronization of Windows passwords with UNIX hosts, you must install the Password Synchronization daemon on each UNIX host on which passwords are to be synchronized. For information about installing the daemon, see Install the Password Synchronization daemon.

When Password Synchronization receives a request for a password change, it encrypts the password and sends it to all UNIX hosts that are to be synchronized with the Windows-based computer or domain. To process the password change request, the UNIX host must be running the Password Synchronization daemon. This daemon receives the request and changes the password on the UNIX host. In addition, if the UNIX host is a master Network Information Service (NIS) server, the Password Synchronization daemon runs **make** to rebuild the NIS passwd map so it can be replicated to subordinate (slave) servers in the NIS domain.

The Password Synchronization daemon performs event logging through the syslogd daemon running on the UNIX host.

**Installing the pluggable authentication module on UNIX hosts**

Pluggable authentication modules (PAMs) allow a UNIX computer to support multiple authentication technologies. Password Synchronization uses this facility to provide UNIX-to-Windows password synchronization.

To allow passwords on Windows-based computers or domains to be changed when users change their UNIX password, the

Password Synchronization PAM module (pam_sso) must be installed on each UNIX host where users can change their passwords. Much like Password Synchronization running on a Windows-based computer, the Password Synchronization PAM module on a UNIX computer intercepts the password change request, encrypts the password, and then transmits the request to the appropriate Windows-based computers running Password Synchronization.

Like the Password Synchronization daemon, the Password Synchronization PAM module performs event logging through the syslogd daemon running on the UNIX host.

For information about configuring UNIX computers for UNIX-to-Windows synchronization, see Configure UNIX Computers for UNIX-to-Windows Synchronization.

**Synchronizing passwords between Windows and NIS domains**

In addition to synchronizing passwords between Windows-based computers and standalone UNIX hosts, you can also use Password Synchronization to provide one-way (Windows-to-UNIX) synchronization as well as two-way synchronization. For more information, see Synchronizing passwords with an NIS domain.

For security reasons, it is recommended that you perform the Windows Server 2003 Service Pack 1 (SP1) compatibility check when selecting **Enable Windows to NIS (AD) Password Sync** in the Password Synchronization Properties dialog box, Configuration tab. For more information about the compatibility check, see Best practices for Password Synchronization.

**Coordinating account names and password policies**

The password policies on both systems must be similar. If the policy on one system is stronger (more restrictive) than the policy on the other system, Password Synchronization might fail to synchronize passwords, and the failure might not be reported.

In addition, Password Synchronization can only synchronize the passwords of accounts with identical user names. Windows and UNIX administrators must ensure that the user names for the Windows and UNIX accounts of a given user match exactly (including case).

If you are configuring Password Synchronization for one-way (Windows-to-UNIX) synchronization, you should consider disabling the ability of users to change passwords on the UNIX hosts that are to be synchronized with the Windows-based computers. Otherwise, if users change their UNIX passwords, their passwords will no longer be synchronized.

# Synchronizing passwords with an NIS domain

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Synchronizing passwords with an NIS domain**

Using Password Synchronization, you can provide one-way (Windows-to-UNIX) and two-way password synchronization between Windows domains and Network Information Service (NIS) domains. You can do this regardless of whether the master server of the NIS domain is running on UNIX or on Windows (Server for NIS).

If the NIS master server is running UNIX, all that is required to provide one-way synchronization is to install Password Synchronization on all Windows-based computers (for example, on the domain controllers) from which you want to synchronize passwords, and then install the single sign-on daemon (SSOD) on the NIS master server. You then edit the sso.conf file on the NIS master server to do the following:

- Set USE_NIS to 1
- Set NIS_UPDATE_PATH to specify the location of the NIS makefile

This instructs the SSOD to run the makefile and push the changed maps whenever a password change request is received from the Windows domain. For more information and additional instructions, see Install the Password Synchronization daemon.

If Server for NIS is acting as the master server for the NIS domain, no action is required to provide one-way password synchronization because whenever a Windows user's password is changed, Server for NIS automatically updates the UNIX password for NIS clients. If you also need to synchronize passwords with UNIX computers that are not part of the NIS domain, you can install Password Synchronization on the Windows domain controllers and configure the UNIX computers as described earlier in this topic.

Providing UNIX-to-Windows synchronization is similar for both types of NIS domains, and is accomplished by performing the following operations:

- If the NIS master server is running UNIX, set it up for one-way synchronization as described earlier in this topic.
- Install Password Synchronization on all domain controllers. If the NIS master server is a UNIX computer, configure Password Synchronization on Windows for two-way synchronization with the master server. Finally, add each NIS client to the list of computers that Password Synchronization will synchronize with, being sure to enable UNIX-to-Windows synchronization and disabling Windows-to-UNIX synchronization. Windows-to-UNIX synchronization should be enabled *only* for the NIS master server. For more information about adding and configuring computers, see Add and remove computers for synchronization and Set computer-specific synchronization properties.
- Install the Password Synchronization pluggable authentication module (PAM) on each NIS client, and then copy the sso.conf file from the master server to the /etc directory of those clients.
- If the NIS master server is a Windows-based computer running Server for NIS, copy Sso.cfg to one of the NIS clients, set SYNC_HOSTS to specify the computer running Server for NIS as the Windows-based computer to synchronize passwords with, and then copy that file to the other UNIX clients. See Using sso.conf to configure Password Synchronization on the UNIX computer to learn more about settings in this file.
- Configure each UNIX computer to allow users to use the **yppasswd** command to change their passwords. To do this, replace the yppasswd binary file on the UNIX computer with a link to the passwd binary file, and then edit the /etc/nsswitch.conf file to replace the passwd and shadow lines with the following:

  ```
  passwd:  files [NOTFOUND=continue] nis
  shadow:  files [NOTFOUND=continue] nis
  ```

  After you do this, when a user runs the yppasswd command to change the user's password, it is actually the passwd binary file that is run to change the password. If the user's passwd entry is not found in the local passwd and shadow files, the NIS password is changed instead.

# Password Synchronization administration

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Password Synchronization administration**
When you have installed Password Synchronization on the appropriate Windows–based computer, you can administer it with Password Synchronization administration (part of Identity Management for UNIX). You must have administrator privileges on the computer you want to administer.

You can use Password Synchronization administration to:

- Select the server to be administered.
- Set the default settings that apply to the entire configuration of UNIX-based computers defined for the Windows–based computer or domain. The settings determine what the Event Log displays, the maximum number of times to resend the failed password update, and the time the service waits before resending a password update that has failed.
- Add or remove a UNIX-based computer from the list of computers designated to receive password updates.
- Create or modify the configuration for the UNIX-based computer, including the custom settings (the default settings applied to that single computer) and the encryption settings for secure communication. The order in which the names of the UNIX hosts appear in the list determines their order in the registry and the order in which they are processed for password synchronization.

# Password encryption

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Password encryption**

A Windows-based computer can send and receive updated passwords to and from a UNIX-based computer as encrypted text only. The Password Synchronization single sign-on daemon (SSOD) receives the encrypted password and decrypts it before requesting the password change on the UNIX host. Similarly, if Password Synchronization is configured to support UNIX-to-Windows synchronization, the pluggable authentication module (PAM) encrypts the password before sending it to Password Synchronization on the Windows-based computer, which then decrypts the password before requesting the password change on the Windows-based computer.

The password can be successfully decrypted only if Password Synchronization and the SSOD or PAM module use the same encryption key to encrypt and decrypt the password. Before installing the SSOD on any UNIX computer, you must first set the default encryption key. You must then specify the same key in the sso.conf file when you install the SSOD on each UNIX host. This will ensure that Password Synchronization and the SSOD on the UNIX hosts will use the same encryption key. For more information about setting the default encryption key, see Set the default encryption key. For information about installing and configuring the SSOD, see Install the Password Synchronization daemon.

For added security, you can specify an encryption key that is used only between a particular Windows-based computer and a UNIX host. For information about configuring Password Synchronization to use a computer-specific encryption key, see Set computer-specific synchronization properties. For information about setting the computer-specific encryption key on the UNIX computer, see Install the Password Synchronization daemon.

# Encryption key requirements

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Encryption key requirements**

The encryption key must meet the following requirements:

- It must be 16 to 21 characters long (21 is recommended).
- It must contain characters from at least three of the following four groups:
  - Uppercase English letters (A–Z)
  - Lowercase English letters (a–z)
  - Westernized Arabic numerals (0–9)
  - Punctuation symbols ` ~ ! @ # $ % ^ & * _ – + = | \ { } [ ] : ; \ " ' < > . ?
- It must not contain a left or right parentheses (that is a "(" or ")" character), a comma (,), or a blank space ( ).

# Controlling password synchronization for user accounts

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Controlling password synchronization for user accounts**

You can control which users' passwords are synchronized by creating two local user groups: PasswordPropAllow and PasswordPropDeny. (Use the Windows User Manager to create these groups.)

In the PasswordPropAllow group, add the user names for which passwords should be synchronized. In the PasswordPropDeny group, add the user names for which passwords should not be synchronized.

Passwords are synchronized for users who are in PasswordPropAllow and are not in PasswordPropDeny.

If PasswordPropAllow does not exist, the effect is the same as if it did exist with all user names in it. If PasswordPropDeny does not exist, the effect is the same as if it did exist with no user names in it.

These rules apply to synchronization from Windows to UNIX and from UNIX to Windows. If a user's password cannot be synchronized from Windows to UNIX, it cannot be synchronized from UNIX to Windows.

The UNIX administrator can ensure that the passwords for certain users are never synchronized, even if synchronization is allowed by the Password Synchronization server. To ensure that a UNIX user account will never have its password synchronized with the Windows password, edit the sso.conf file to place the user name of the account, preceded by a minus sign (–), after **SYNC_USERS=**. For example, to ensure that the password of the root account is never synchronized with a Windows account by that name, make sure that the following line appears in sso.conf:

**SYNC_USERS=–root**

# Using sso.conf to configure Password Synchronization on the UNIX computer

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Using sso.conf to configure Password Synchronization on the UNIX computer**

To change settings and customize how Password Synchronization works on a UNIX-based computer, you change settings in the sso.conf file. For information about installing the sso.conf file, see Install the Password Synchronization daemon.

The following table describes the values you can set in the sso.conf file.

| Value | Description |
|---|---|
| CASE_IGNORE_NAME | Specifies whether Password Synchronization will ignore differences in uppercase and lowercase letters in user names when it compares Windows and UNIX user names. To allow case-insensitive comparisons, set this entry to 1 (the default). To force Password Synchronization to use case-sensitive comparisons, set this entry to 0. |
| ENCRYPT_KEY | Specifies the default key used to encrypt passwords exchanged with Windows servers. You can use settings in the SYNC_HOSTS value to specify a different encryption key for a specific Windows server. For information about valid encryption keys, see Encryption key requirements. |
| FILE_PATH | Specifies the full path and name of the passwd or shadow file (such as /etc/passwd). This file must contain the encrypted passwords for users, and the type of file (passwd or shadow) must be the same as specified by USE_SHADOW. On AIX systems, the path and name of the shadow file is /etc/security/passwd. |
| IGNORE_PROPAGATION_ERRORS | When set to **1**, specifies that the Password Synchronization PAM module is to ignore any error that occurs when a Windows password is being changed and to continue synchronization with other hosts specified in SYNC_HOSTS. |
| NIS_UPDATE_PATH | Specifies the full path to the NIS makefile. This value is ignored unless USE_NIS is set to **1**. |
| PORT_NUMBER | Specifies the default number of the port on which the Password Synchronization daemon will listen for password changes from Windows servers. You can use settings in the SYNC_HOSTS value to specify a different port number for a specific Windows server. |
| SYNC_DELAY | Specifies the number of seconds the Password Synchronization PAM module will wait between synchronization attempts. |

| | |
|---|---|
| SYNC_HOSTS | Specifies the Windows servers or domain controllers with which passwords are to be synchronized. In addition, you can specify a port number or encryption key (or both) for a specific server. Enclose each entry in parentheses and separate items with a blank space. You can have multiple entries on separate lines, each of which cannot exceed 269 characters. The total list of servers or domain controllers is generated by concatenating all entries. For example:<br><br>`SYNC_HOSTS=(Marketing)`<br>`SYNC_HOSTS=(Sales,ASDFhjkl4321ZyXw) (Accounting,6678)`<br>`SYNC_HOSTS=(Shipping,6678,ASDFhjkl4321ZyXw)`<br><br>synchronizes passwords with the Marketing server using the default port and encryption key, with the Sales server using ASDFhjkl4321ZyXw as the encryption key, with the Accounting server using 6678 as the port number, and with the Shipping server using 6678 as the port number and ASDFhjkl4321ZyXw as the encryption key.<br><br>If you use server-specific port number or encryption key settings, you must use the same values to configure Password Synchronization on the Windows server, or the passwords will not be synchronized. For information about configuring Password Synchronization on the Windows server, see Configure Password Synchronization. |
| SYNC_RETRIES | Specifies the number of times that the Password Synchronization PAM module will attempt to synchronize a password change with a Windows server or domain controller. |
| SYNC_USERS | Specifies UNIX users whose passwords are to be synchronized. You can specify ALL to synchronize passwords for all users or NONE to disable password synchronization for users. You can also specify particular users. If you specify one or more users preceded by plus signs (**+**), only those users will have their passwords synchronized. If you specify one or more users preceded by minus signs (**–**), all users except the specified users will have their passwords synchronized. For example, to allow only users bobg and kimr to synchronize their passwords, specify:<br><br>`SYNC_USERS=+bobg +kimr`<br><br>To prevent only root and bobg from having their passwords synchronized, specify:<br><br>`SYNC_USERS=-root -bobg`<br><br>The minus sign always takes precedence, regardless of the order in which entries appear. For example, the following specifies that the password for user chrisa will not be synchronized:<br><br>`SYNC_USERS=+chrisa -chrisa +chrisa` |
| TEMP_FILE_PATH | Specifies the full path of the directory to be used to hold a temporary file while the passwd or shadow file is updated. This should be the same directory in which the passwd or shadow file is located. For security reasons, only the administrator should have access to this directory. |
| USE_NIS | Set to **0** if Password Synchronization is not synchronizing with an Network Information Service (NIS) domain; set to **1** if Password Synchronization is synchronizing with an NIS domain. If USE_NIS is set to **1**, specify a valid path for NIS_UPDATE_PATH. |
| USE_SHADOW | Set to **0** if the passwd file is to be used for synchronization; set to **1** if the shadow file is to be used. |

# Troubleshooting Password Synchronization

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Troubleshooting**
What trouble are you having?

- User cannot log on to UNIX system after changing Windows password
- Passwords fail to synchronize in a Windows domain, seemingly at random
- User's password is changed on some, but not all, computers
- Error ID 4104 is recorded in the event log for a system account, which usually has a name ending with a dollar sign ($)
- An error message about the encryption key is recorded in Event Viewer after Password Synchronization installation completes
- I cannot stop the single sign-on daemon (SSOD) using kill -TERM on Linux

**User cannot log on to UNIX system after changing Windows password**
**Cause:** Windows does not report a problem if an attempt to synchronize a UNIX password with a new Windows password fails.

**Solution:** Check the Windows event log to determine why the attempt to change the password on the UNIX system failed.

**Passwords fail to synchronize in a Windows domain, seemingly at random**
**Cause:** Password Synchronization is not configured identically on all domain controllers in the domain. As a result, if a nonconforming domain controller accepts a user's password change, it might not be able to change the password on UNIX computers.

**Solution:** Ensure that Password Synchronization is configured identically on all domain controllers, particularly host settings and default settings for encryption keys and ports.

**User's password is changed on some, but not all, computers**
**Cause:** Password policies are more restrictive on some computers, user names do not match between Windows and UNIX computers, or the user changed the password on a UNIX computer when two-way synchronization is not set up.

**Solution:** Ensure that password policies on Windows and UNIX computers that synchronize passwords are similar. Otherwise, if the user changes the password on the less restrictive computer, the more restrictive system might not accept the new password. Password policies that govern minimum and maximum length, character case and alphanumeric mix, expiration, and reuse must be as close as possible between Windows and UNIX computers that synchronize passwords. Also, Windows and UNIX system administrators must ensure that that user names, including case, are identical on the Windows and UNIX computers.

**Error ID 4104 is recorded in the event log for a system account, which usually has a name ending with a dollar sign ($)**
**Cause:** This error does not indicate a problem. It is logged when a backup domain controller or domain member server resets its secure channel with the domain. When this happens, the server password is also reset. Password Synchronization intercepts these password change requests; because they are for computer accounts rather than for user or group accounts, Password Synchronization logs error number 4104.

**Solution:** No corrective measures are necessary.

**An error message about the encryption key is recorded in Event Viewer after Password Synchronization installation completes**
**Cause:** This error does not indicate a problem; it is a reminder for the Password Synchronization administrator to change the default encryption key. Changing the default encryption key is a security best practice for Password Synchronization, and helps prevent unauthorized users from obtaining passwords. Typically, the description text of this error message is "Default encryption key is insecure. Please generate new encryption key."

**Solution:** Set a new encryption key by using the procedure in the topic Set the default encryption key.

**I cannot stop the single sign-on daemon (SSOD) using kill -TERM on Linux**

**Cause:** This is a known limitation.

**Solution:** Use **kill 9** *SSOD_PID* instead.

# Command-line utility

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**Command-line Utility**
The following command-line utility is available to administer Password Synchronization:

psadmin

# psadmin

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

**psadmin**
**NAME**
psadmin - Windows command-line utility to manage Password Synchronization

**SYNOPSIS**

```
psadmin [computername] [common_option] [add | delete | list]
psadmin [computername] [common_option] config [config_option]
```

**DESCRIPTION**
The **psadmin** Windows command-line utility manages the Password Synchronization component of Identity Management for UNIX, either on a specified computer, or globally. The specific action that **psadmin** performs depends on the command argument you specify.

In addition to specific command arguments, **psadmin** accepts the following common options and arguments, represented by *common_option* in the command synopsis:

| Term | Definition |
|---|---|
| **-u** *username* | The account name of the user whose password synchronization behavior you want to modify or view. |
| **-p** *password* | The password for the user account whose password synchronization behavior you want to modify or view. |
| **-?** | Displays usage information for the command. |

The following configuration options are accepted by **psadmin**:

| Term | Definition |
|---|---|
| **-comp** *name* | Computer to which configuration options are applied. If **-comp** is unspecified, Password Synchronization modifies the default configuration settings. If **-comp** is the only option specified, then Password Synchronization configuration of the specified computer is displayed. |
| **-enable** *direction* | Specifies the direction of password synchronization. The variable *direction* can contain one of the following values:<br><br>**WintoUnix**: Synchronize password changes from computers that run Windows operating systems to computers that run UNIX operating systems.<br><br>**UnixToWin**: Synchronize password changes from computers that run UNIX operating systems to computers that run Windows operating systems.<br><br>**BothDir**: Enable two-way password synchronization. |
| **-key** *keyvalue* | Sets the encryption and decryption key for the computer specified by **-comp**. If *keyvalue* is random, Password Synchronization uses a random encryption key. |
| **-port** *number* | Sets the port number for the specified computer. |
| **-retry** *number* | Specifies the number of retries allowed. Because this option is a global setting, it can be used only when **-comp** is not used. |
| **-interval** *secs* | Specifies the elapsed time period, in seconds, between retries. Because this option is a global setting, it can be used only when **-comp** is not used. |
| **-log** [**yes** \| **no**] | Enables or disables logging. Because this option is a global setting, it can be used only when **-comp** is not used. |

| | |
|---|---|
| **-?** | Displays **psadmin** usage and arguments. |

The following command arguments are accepted by **psadmin**:

| Term | Definition |
|---|---|
| **add** *computername* | Adds the specified computer to the list of computers participating in password synchronization. |
| **delete** *computername* | Deletes the specified computer from the list of computers participating in password synchronization. |
| **list** | Displays the list of computers participating in Password Synchronization. |

**See Also**
**Concepts**
Configure Password Synchronization

# Accessibility Features of Identity Management for UNIX

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Products in the Windows Server 2003 family include a number of accessibility features that can improve display, sound, mouse, and keyboard settings for users who are blind, have motion disabilities, or are deaf or hard of hearing. Many accessibility features are also useful to people without disabilities.

**Accessibility Features and Options**
Because Identity Management for UNIX tasks can be completed using either a Windows-based interface or shell commands, accessibility options are available for both computing environments.

**Accessibility Features and Options for Windows-based Tools in Identity Management for UNIX**
For more information about accessibility features and options, see Accessibility for People with Disabilities on the Microsoft Windows Server 2003 TechCenter.

Because the Windows-based management tool for Identity Management for UNIX is a Microsoft Management Console (MMC) 3.0 snap-in, Identity Management for UNIX shares applicable MMC 3.0 accessibility features. For more information about accessibility in MMC, see Accessibility for MMC.

**Accessibility Features and Options for the Command-line Environment**
For information about keyboard shortcuts that can provide greater accessibility to users of command-line and UNIX-based shell computing environments, refer to the Help and release notes provided with the Utilities and Software Development Kit (SDK) for UNIX-based Applications installation package, available for download from the Microsoft Web site.

# Welcome to Subsystem for UNIX-based Applications

Updated: March 13, 2006

Applies To: Windows Server 2003 R2

Subsystem for UNIX-based Applications (SUA) is a source-compatibility subsystem for compiling and running custom UNIX-based applications on a computer running a Windows server-class operating system. You can make your UNIX applications fully interoperable with Windows in SUA with little or no change to your original source code.

Subsystem for UNIX-based Applications provides an operating system for Portable Operating System Interface (POSIX) processes. SUA, along with its package of support utilities (such as shells and a Telnet client) available for download on the Microsoft Web site, provides a complete UNIX environment. The download package includes a comprehensive set of scripting utilities, and a software development kit (SDK) designed to fully support the development capabilities of SUA and to provide a complete UNIX-based application development experience.

SUA also supports case-sensitive file names, job control, compilation tools, and the use of over 300 UNIX commands, utilities, and shell scripts. Because the subsystem installs separately from the Windows kernel, it offers true UNIX functionality without any emulation.

**New features for Windows Server 2003 Release 2**
New features for Windows Server 2003 Release 2 (R2) include the following:

- **Database (OCI/ODBC) library connectivity.** SUA supports connectivity to Oracle and SQL Server from database applications by using the Oracle Call Interface (OCI) and the Open Database Connectivity (ODBC) standard.
- **Microsoft Visual Studio Debugger Extension for debugging POSIX applications.** SUA includes support for debugging your POSIX processes using Visual Studio IDE.
- **Utilities based on SVR-5 and BSD UNIX environments.** The SUA download package supports two different UNIX environments: SVR-5 and BSD.
- **Support for 64-bit applications.** Using a process called thunking, SUA provides support not only for 64-bit applications running on a 64-bit operating system, but also default support for 32-bit binaries running on a 64-bit operating system.

For more information, and to download the scripting utilities and SDK, see the Microsoft Web site.

**In this topic**


- Subsystem for UNIX-based Applications Overview
- Installing and Using Utilities and SDK for Subsystem for UNIX-based Applications
- Case Sensitivity in Subsystem for UNIX-based Applications
- Setting setuid in Subsystem for UNIX-based Applications
- POSIX Compliance in Subsystem for UNIX-based Applications
- System Calls Supported by Subsystem for UNIX-based Applications

# Subsystem for UNIX-based Applications Overview

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Subsystem for UNIX-based Applications (SUA) is the next-generation of Microsoft's POSIX subsystem, similar to the Interix subsystem that formerly shipped with Windows Services for UNIX 3.5 or previous POSIX subsystems that shipped with Windows 2000 and NT 4.0. Subsystem for UNIX-based Applications and its accompanying utilities provide you with an environment that resembles any other UNIX system. It also includes case-sensitive file names, job control, compilation tools, and the use of over 300 UNIX commands and utilities and shell scripts. Because the Subsystem for UNIX-based Applications is layered on top of the Windows kernel, it offers true UNIX functionality without any emulation.

**Installing Subsystem for UNIX-based Applications**
Subsystem for UNIX-based Applications is available with Windows Server 2003 Release 2 (R2). You can install SUA using the procedures provided in the topic Install an optional component by using Add/Remove Windows Components.

**Computing Environments with Subsystem for UNIX-based Applications**
A computer running SUA provides two different command-line environments: the UNIX environment and the Windows environment. Applications run on specific subsystems and in specific environments. When you load Subsystem for UNIX-based Applications, you are using a UNIX environment; when you run applications on the Windows subsystem, you are using a Windows environment.

For example, when you use the Korn shell or C shell available with the download package Utilities and Software Development Kit (SDK) for UNIX-based Applications, or any other application that was compiled to run on SUA, you are using a UNIX environment. Among other things, this environment features case-sensitive file names and path names that use the format /usr/examples. For more information about the commands and utilities that are available in the Subsystem for UNIX-based Applications, see the Help files that ship with the download package Utilities and Software Development Kit (SDK) for UNIX-based Applications. This download package is available from the Microsoft Web site.

When you use the command processor **cmd.exe**, you are using the Windows environment. This environment includes case-insensitive file names, and path names that include drive letters, as shown in the following example: C:\SUA\usr\examples. For more information about path names in the Subsystem for UNIX-based Applications and Windows, see the Help files that ship with the download package Utilities and Software Development Kit (SDK) for UNIX-based Applications.

**Porting Applications by Using Subsystem for UNIX-based Applications**
Subsystem for UNIX-based Applications provides both the tools and the application programming interface (API) libraries for porting applications to run on Windows-based computers. With the Software Development Kit (SDK), which provides a front end for Microsoft Visual C++, you can have a UNIX environment for development, but still have the benefits of the native compiler for Windows. For more information about path names in Subsystem for UNIX-based Applications and Windows, see the Help files that ship with the download package Utilities and Software Development Kit (SDK) for UNIX-based Applications.

**See Also**
**Other Resources**
Installing and Using Utilities and SDK for UNIX-based Applications

# Installing and Using Utilities and SDK for Subsystem for UNIX-based Applications

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The download package available from the Microsoft Web site, **Utilities and Software Development Kit (SDK) for UNIX-based Applications**, is designed to work with the Subsystem for UNIX-Based Applications (SUA) component that ships with Microsoft Windows Server 2003 Release 2 (R2).

**Components Included with Download Package**
The download package contains the following components:

- Base Utilities
- SVR-5 Utilities
- Base SDK
- GNU SDK
- GNU Utilities
- UNIX Perl
- Visual Studio Debugger Extension

In addition to the features provided in Windows Services for UNIX 3.5, the download package supports 64-bit application development with SUA. The download also enables development and porting of custom UNIX applications using the Windows Oracle Call Interface (OCI) and Windows ODBC libraries (collectively referred to as "mixed mode" throughout this document).

**Base Utilities, GNU Utilities, SDK and UNIX Perl**
These components include all utilities that formerly installed with Windows Services for UNIX 3.5. This consists of the nearly 300 BSD-based and optional SUA GNU utilities. Many of these utilities (approximately 80) are also available in 64-bit versions.

GNU components include utilities and SDK tools distributed under the terms of the GNU General Public License (GPL).

See the Help and release notes provided with the installation package for a complete listing of base utilities.

**SVR-5 Utilities**
This set consists of a Korn shell and approximately 150 utilities, most of which support SVR-5. These utilities currently are available in 32-bit versions, and may fail in specific 64-bit scenarios; however, the utilities run on WOW 64-bit platforms.

See the Help and release notes provided with the installation package for a complete listing of SVR-5 utilities.

**Base SDK**
The current release supports most of the libraries and headers installed with Windows Services for UNIX 3.5. These support both 32-bit and 64-bit SUA-related software migration and development.

This SDK consists of the following libraries and corresponding header files: Standard C Library, Standard C++ Library and Math Library. The C++ library and the math library are functional equivalents of those supplied by Microsoft as part of their Visual Studio environment. Standard C Library and Math Library are available for both 32-bit and 64-bit platforms. Standard C++ Library is available for both 32-bit and 64-bit binaries, and supports multithreaded applications.

**Compiler Compatibility**
The Base SDK package can be used with either the gcc (GNU C compiler) and g++ (GNU C++ compiler) compilers, or the Visual Studio 2005 compiler. Linking to shared objects is not supported with the Visual Studio 2005 compiler. The C++ library supplied with this package is compatible with the Visual Studio 2005 compiler, and can be used with c89.

The g++ compiler uses its own library for C++ compilation.

**System Requirements**
The Utilities and SDK for Subsystem for UNIX-based Applications package installs only on Windows Server 2003 R2.

The minimum system requirements for installing Utilities and SDK for Subsystem for UNIX-based Applications depends on which components you are installing and the file system of the disk where the components will be installed.

The maximum disk space required to install all Utilities and SDK for Subsystem for UNIX-based Applications components is approximately 350 MB.

**Additional Requirements**

- To receive product support, you must install the latest Windows service pack before installing Utilities and SDK for Subsystem for UNIX-based Applications, and then continue to install additional Windows service packs as they become available.
- You must install Utilities and SDK for Subsystem for UNIX-based Applications on a partition that is formatted with the NTFS file system. File system operations on FAT partitions are not supported.
- It is recommended that the computer have at least 16 MB of RAM in addition to the recommended minimum configuration for the operating system.

**Installation Options**
You can choose one of the following installation options when installing Utilities and SDK for Subsystem for UNIX-based Applications:

- Standard Installation
- Custom Installation

The following table shows which components are installed on a computer as part of the standard installation (Standard) and which additional components are available as part of a customized installation (Custom). If you select the customized installation, you can choose not to install components that are automatically installed as part of the standard installation. These components can only be installed on Windows Server 2003 R2.

| Component | Type of Installation |
|---|---|
| Base utilities | Standard |
| Base SDK | Standard |
| SVR-5 utilities | Custom |
| GNU utilities | Custom |
| GNU SDK | Custom |
| UNIX Perl | Custom |
| Visual Studio Debugger Extension | Custom |

**Note**
If you are upgrading from Windows Services for UNIX 3.5, existing settings are saved during R2 installation, and are restored during the setup process for Utilities and SDK for Subsystem for UNIX-based Applications.

**Installing Utilities and SDK for Subsystem for UNIX-based Applications by Using the Windows Interface**
The following installation options are available:

- Standard Installation
- Custom Installation

#### Standard Installation

### To perform a standard installation

1. Download the Utilities and SDK for Subsystem for UNIX-based Applications from the Microsoft Web site
2. In the Utilities and SDK for Subsystem for UNIX-based Applications Wizard dialog box, click **Next**.
3. In the **User name** text box, type your name. If the name of your organization does not show in the **Organization** field, enter the name of your organization.
4. Read the End User License Agreement (EULA) carefully. If you accept the terms of the agreement, click **I accept the terms in the License Agreement**, and then click **Next** to continue installation. If you click **I do not accept the License Agreement (Exit Setup)**, the installation procedure terminates.
5. To install default Utilities and SDK for Subsystem for UNIX-based Applications components in the default directory, click **Standard Installation**, and then click **Next**. To customize your installation, see Custom Installation.

#### Custom Installation

### To perform a custom installation

1. Download the Utilities and SDK for Subsystem for UNIX-based Applications from the Microsoft Web site .
2. In the Utilities and SDK for Subsystem for UNIX-based Applications Wizard dialog box, click **Next**.
3. In the **User name** text box, type your name. If the name of your organization does not show in the **Organization** field, enter the name of your organization.
4. Read the End User License Agreement (EULA) carefully. If you accept the terms of the agreement, click **I accept the terms in the License Agreement**, and then click **Next** to continue installation. If you click **I do not accept the License Agreement (Exit Setup)**, the installation procedure terminates.
5. To install default Utilities and SDK for Subsystem for UNIX-based Applications components in the default directory, click **Standard Installation**, and then click **Next** (see Standard Installation in the preceding section). If you want to specify a different set of components or a different installation location, click **Custom Installation**, and then click **Next**.
6. Click to fill the check boxes for each component that you want to install, and then click the appropriate option. If you do not want to install a component, click the icon next to the component, and then click **Entire feature will not be available**. When you are finished specifying the components to install, click **Next**.
7. If you are installing the GNU Software Development Kit (SDK), read the information about the GNU Library General Public License carefully. If you want to install the GNU SDK, click **Next** to continue installation. To make changes to the components you have selected click **Back**.

**Note**

After you have installed Utilities and SDK for Subsystem for UNIX-based Applications, you can install additional components or remove installed components.

Utilities and SDK for Subsystem for UNIX-based Applications components cannot be run from a network server. All files must be installed on the local computer.

Standard installation does not provide a complete UNIX development environment. To obtain all of the GNU compilers, build tools, and other SDK components, you must select **Custom Installation** and then select **GNU SDK**.

Once Utilities and SDK for Subsystem for UNIX-based Application is installed, **setuid** behavior is disabled by default. This affects the behavior of certain SUA programs.

In this installation, case sensitivity will not be enabled for object names (such as file names) in non-Win32 subsystems.

#### Installing Utilities and SDK for Subsystem for UNIX-based Applications by Using the Command Line

Perform the following step to start a command-line installation of the Utilities and SDK for Subsystem for UNIX-based Applications.

**To perform a command-line installation**

- At a command prompt, enter the following:

  **msiexec /i** [path\]**setup.msi** [**ADDLOCAL="**component[,component...]**"**] {/**qb**|/**q**}

The following table describes the arguments and options for this command.

| Argument/Option | Description |
|---|---|
| path | The fully qualified path of the Utilities and SDK for Subsystem for UNIX-based Applications setup.msi file. If you omit the path, the setup.msi file must be in the current directory. |
| /qb | Installs Utilities and SDK for Subsystem for UNIX-based Applications with a basic user interface. You must specify either the /q or /qb option; otherwise, the results will be unpredictable. |
| /q | Installs Utilities and SDK for Subsystem for UNIX-based Applications with no user interface or display. This option allows you to perform an unattended installation. You must specify either the /q or /qb option; otherwise, the results will be unpredictable. |

**Note**

This command completes only a standard installation.

For security purposes, all SUA services (also called daemons) are disabled by default. To manage SUA services, edit the **/etc/inetd.conf** file.

**See Also**
**Other Resources**
Microsoft Windows Server 2003 TechCenter
MSDN Servers and Enterprise Development

# Case Sensitivity in Subsystem for UNIX-based Applications

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

In Windows-based operating systems, the names of most objects (such as files and directories) are case preserving, but case insensitive. That means you can use uppercase and lowercase characters when naming such objects, but Windows does not distinguish between names based on case alone. For example, you cannot have two files in the same directory named sample.txt and Sample.txt because Windows regards the names to be identical for the purposes of identifying files.

**Case Sensitivity and Subsystem for UNIX-based Applications**

The UNIX operating system, by contrast, is fully case sensitive, and so UNIX computers distinguish between object names when the only difference between those names is the case of characters used in the object names. On most UNIX computers, therefore, sample.txt and Sample.txt could be stored in the same directory, and the UNIX computer would distinguish between them when performing operations on the files. For example, the command **rm S*.txt** would delete Sample.txt, but not sample.txt.

In order to implement typical UNIX behavior, the Subsystem for UNIX-based Applications is normally case sensitive when working with file names. This can present security issues, particularly for Windows users who are accustomed to the case-insensitive conventions of Windows. For example, a Trojan horse version of edit.exe named EDIT.EXE could be stored in the same directory as edit.exe. If a user were to type **edit** at a Windows command prompt, the Trojan horse version (EDIT.EXE) could be executed instead of the standard version. If case sensitivity is enabled, Windows users should be made aware of this possibility.

**Restoring Case Sensitivity Defaults**

In Windows XP and in Windows Server 2003 operating systems, the default behavior of subsystems other than the Win32 subsystem is to be case preserving but case insensitive; in previous versions of Windows, such subsystems were fully case sensitive by default. To support standard UNIX behavior, you can change the default Windows XP and Windows Server 2003 behavior for non-Win32 subsystems when installing the base utilities that are part of the download package Utilities and SDK for Subsystem for UNIX-based Applications. If you enable case sensitivity and then subsequently uninstall the base utilities, setup restores the default, case-insensitive behavior of non-Win32 subsystems.

For more information about changing case sensitivity settings, see Configuring case sensitivity for file and folder names.

# Setting setuid in Subsystem for UNIX-based Applications

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

On many systems, a root user or superuser has complete control of the system. Subsystem for UNIX-based Applications (SUA) does not recognize a root user. When the Portable Operating System Interface (POSIX) standard was developed, the concept of a root user was considered an administrative issue. Instead of a root user, the POSIX standard defines appropriate privileges for some operations.

**In This Topic**

setuid and setgid

Subsystem for UNIX-based Applications and Identities

Setup Considerations for setuid in Subsystem for UNIX-based Applications

**setuid and setgid**

The **setuid** and **setgid** mechanisms allow a program, when run, to adopt certain aspects of a security principal other than that of the user running the program. They allow the program to alternate these aspects between those of the invoking user and those of another security principal.

Appropriate privileges are distinct from this. The setuid and setgid mechanisms allow an application to control the answer to the question: "Who are you?" Appropriate privileges answer the question: "What can you do?"

**Subsystem for UNIX-based Applications and Identities**

In typical UNIX environments, there is exactly one user to whom all privileges are granted. That user, typically called **root**, has **uid == 0**. In SUA, all supported privileges on a given system are granted to users who are members of either the Administrators group or the Domain Administrators group on that system. You do not need to log on as Administrator. Any user with membership in either the Administrators group or the Domain Administrators group possesses **root** privileges. Not all privileges defined in POSIX or UNIX are available in SUA; that is, there are certain privileges that are not granted to any user.

Appropriate privileges are typically required for a variety of actions. These include accessing the file system, sending signals to other processes (process control), or changing the effective user identifier (UID) or group identifier (GID) for a process in order to change the ability of that process to perform certain actions.

According to the POSIX standard, a file has permissions that include bits to set a UID (**setuid**) and set a GID (**setgid**). If either or both bits are set on a file, and a process executes that file, the process gains the UID or GID of the file. When used carefully, this mechanism allows a nonprivileged user to execute programs that run with the higher privileges of the file's owner or group. When used incorrectly, however, this can present security risks by allowing nonprivileged users to perform actions that should only be performed by an administrator. For this reason, SUA does not support this mechanism by default. Instead, if you attempt to execute a file with either the **setuid** or **setgid** bit set, SUA does not execute the file and returns the error code **ENOSETUID**.

**Setup Considerations for setuid in Subsystem for UNIX-based Applications**

If you rely on an application that requires standard POSIX behavior, you can configure SUA to execute files with the **setuid** or **setgid** bits set. If SUA is configured in this way, when a process executes a file that has the **setuid** or **setgid** bit set, SUA constructs local security tokens for the process with the privileges assigned to the owner or group of the file. Because the tokens are local, they are not recognized by other computers on the network. This means that, even if the file is owned by a member of the Domain Administrators group, for example, the process does not have trusted access through Windows networking to other computers in the domain. Instead, the privileges will be effective only on the computer on which the process is running.

For example, a process executes a program file with its **setuid** bit set and owned by a member of the Domain Administrators group. If that program attempts to change a domain user's password, that attempt will fail because the security tokens of the process are local and so are not recognized by other systems in the domain. If, on the other hand, the program attempts to change a local user's password, the attempt will succeed because the file's owner is a member of the Domain Administrators group, which typically belongs to the local computer's Administrators group. For information about how to configure SUA to execute files with the **setuid** or **setgid** bits set, see the Help files that ship with the download package Utilities and Software Development Kit (SDK) for UNIX-based Applications.

For the **setuid(2)**, **setgid(2)**, and **chroot(2)** functions, a process must have the effective UID that maps back to the system account, the Administrator account for the local domain, or the Administrator account for the principal domain. The local

Administrator account can only change the UID or GID to another ID in the same domain.

For **chown(1)**, **chmod(1)**, and **chgrp(1)**, an account must have the Windows privileges SE_BACKUP and SE_RESTORE to perform **make** operations on files owned by another user or group, or to modify permissions on a file owned by another user. These permissions typically belong to the Administrator and Backup Operator accounts.

# POSIX Compliance in Subsystem for UNIX-based Applications

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

Subsystem for UNIX-based Applications (SUA) provides a Portable Operating System Interface-compliant (POSIX-compliant) subsystem that runs UNIX shells, utilities, and applications natively alongside Windows applications. This topic describes some of the working considerations for users requiring a POSIX-compliant operating system.

**POSIX Shell Compatibility**

The SUA Korn shell follows traditional Korn shell behavior, which is almost identical to POSIX behavior. If you need strict POSIX conformance, which some work environments require, see the differences between the Korn shell and POSIX in the **ksh(1)** documentation in the Help files found in the download package Utilities and Software Development Kit (SDK) for UNIX-based Applications.

To get full POSIX.2 conformance from the Korn shell, you must run it in POSIX mode. To run the Korn shell in POSIX mode, either open the Korn shell with the **-o posix** option, run the command set **-o posix** from within the shell, or set the **POSIXLY_CORRECT** shell parameter.

POSIX conformance is not supported in the C shell. Unlike the Korn shell, the C shell has no POSIX-mode command options.

**Traverse Checking of Directory Permissions**

By default, Windows users have permissions to move through a directory tree, though they might not have permissions to access each directory in the path. On POSIX systems, you cannot access a directory unless you have permission to access every directory in the path to the directory. This POSIX behavior is known as traverse checking.

For example, if you do not have permission to access a directory, you cannot change the current directory to any of its subdirectories, even if you have permission to access the subdirectories. By default in Windows, you can change the current directory to any subdirectories for which you have permission.

In Windows, there is a special permission called the Traverse Folder/Execute File permission, which you can set for a user or group. You can configure this permission to allow or deny movement through a directory tree. The Traverse Folder permission takes effect only when the group or user is not granted the Bypass Traverse Checking user right in the Group Policy snap-in. In Windows 2000, by default, the **Everyone** group is given the Bypass Traverse Checking user right. In Windows XP and the Windows Server 2003 operating systems, the **Administrators**, **Backup Operators**, **Everyone**, **Power Users**, and **Users** groups are all given this right. For strict conformance with POSIX standards, remove the Bypass Traverse Checking right from all users and groups except **Backup Operators**. To conform to privileges granted to the root user on many other UNIX implementations, remove the Bypass Traverse Checking right from all users except **Administrators** and **Backup Operators**.

# System Calls Supported by Subsystem for UNIX-based Applications

Updated: August 22, 2005

Applies To: Windows Server 2003 R2

The **syscall**(2) function performs the system call whose assembly language interface has the specified number with the specified arguments. Symbolic constants for system calls can be found in the header file **<sys/syscall.h>**.

For more information about the **syscall** function, see the Help and release notes provided with the Utilities and Software Development Kit (SDK) for UNIX-based Applications installation package, available for download from the Microsoft Web site.

**Supported System Calls**

The following system calls are supported by Subsystem for UNIX-based Applications:

| | | | | | | |
|---|---|---|---|---|---|---|
| _control87 | _controlfp | _exit | _getInstallPath | _getInstallPath_Win | _prefixInstallPath | _prefixInstallPath_Win |
| _tolower | _toupper | Aa64l | abort | abs | accept | access |
| acos | acosh | alarm | alloca | alphasort | arc4random_addrandom | arc4random |
| arc4random_stir | asctime | asctime_r | asin | asinh | assert | atan |
| atan2 | atan2f | atanf | atanh | atexit | atof | atoi |
| atol | auth_destroy | authenticateuser | authnone_create | authunix_create | authunix_create_default | Bbasename |
| bcmp | bcopy | bind | bindresvport | blf_cbc_decrypt | blf_cbc_encrypt | blf_dec |
| blf_ecb_decrypt | blf_ecb_encrypt | blf_enc | blf_key | brk | bsearch | bzero |
| Ccabs | calloc | callrpc | catclose | catgets | catopen | cbrt |
| ceil | cfgetispeed | cfgetospeed | cfsetispeed | cfsetospeed | chdir | chmod |
| chown | chpass | chroot | clearerr | clnt_control | clnt_create | clnt_destroy |
| clnt_freeres | clnt_geterr | clnt_pcreateerror | clnt_perrno | clnt_perror | clnt_spcreateerror | clnt_sperrno |
| clnt_sperror | clnt_statclnt_broadcast | clnt_statclnt_call | clnt_statpmap_rmtcall | clntraw_create | clnttcp_create | clntudp_bufcreate |
| clntudp_create | clock | close | closedir | closelog | confstr | connect |
| copysign | copysignf | cos | cosh | creat | crypt | ctermid |
| ctime | ctime_r | cuserid | Ddaemon | dbm_clearerr | dbm_close | dbm_delete |
| dbm_error | dbm_fetch | dbm_firstkey | dbm_nextkey | dbm_open | dbm_store | des_cipher |
| des_setkey | difftime | dirfd | dirname | div | dlclose | dlerror |
| dlopen | dlsym | dn_comp | dn_expand | drand48 | drem | dup |
| dup2 | Eecvt | el_deletestr | el_end | el_getc | el_gets | el_init |
| el_insertstr | el_line | el_parse | el_push | el_reset | el_resize | el_set |
| el_source | encrypt | endgrent | endhostent | endnetent | endprotoent | endpwent |
| endservent | endutxent | env | env_alloc | env_array | env_cron | env_expand_win |
| env_free | env_get | env_login | env_put | env_putarray | env_set | env_strfree |
| env_t | env_unset | env_winlogin | erand48 | erf | erfc | err |
| errno | errx | exec | exec_asuser | execl | execl_asuser | execle |
| execle_asuser | execlp | execlp_asuser | execv | execv_asuser | execve | execve_asuser |
| execvp | execvp_asuser | exit | exp | expm1 | Ffabs | fabsf |
| fchdir | fchmod | fchown | fclose | fcntl | fcvt | fdopen |
| feof | ferror | fflush | ffs | fgetc | fgetln | fgetpos |
| fgets | fgetwc | fgetws | fileno | finite | finitef | flock |
| flockfile | floor | floorf | fmod | fnmatch | fopen | fork |
| forkpty | fp_nquery | fp_resstat | fpathconf | fprintf | fpurge | fputc |
| fputs | fputwc | fputws | fread | free | freeaddrinfo | freopen |
| frexp | fscanf | fseek | fsetpos | fstat | fstatvfs | fsync |
| ftell | ftime | ftok | ftruncate | ftrylockfile | fts_children | fts_close |
| fts_open | fts_read | fts_set | ftw | funlockfile | fwrite | Ggai_strerror |
| gamma | gamma_r | gcvt | get_myaddress | getaddrinfo | getbsize | getc |
| getcwd | getdate | getdtablesize | getegid | getenv | geteuid | getgid |
| getgrent | getgrent_nomembers | getgrgid | getgrgid_r | getgrgid_ex | getgrgid_nomembers | getgrnam |
| getgrnam_r | getgrnam_nomembers | getgroups | gethostbyaddr | gethostbyname | gethostent | gethostname |
| getitimer | getlogin | getlogin_r | getmode | getnameinfo | getnetbyaddr | getnetbyname |
| getnetent | getopt | getpass | getpdomain | getpeername | getpgrp | getpid |
| getppid | getpriority | getprotobyname | getprotobynumber | getprotoent | getpwent | getpwnam |

| | | | | | | |
|---|---|---|---|---|---|---|
| getpwnam_r | getpwuid | getpwuid_ex | getpwuid_r | getrlimit | getrpcbyname | getrpcbynumber |
| getrpcent | getrpcport | getrusage | gets | getsecret | getservbyname | getservbyport |
| getservent | getsockname | getsockopt | getsubopt | gettimeofday | gettzenv | getuid |
| getutxent | getutxid | getutxline | getwc | getw | getwchar | glob |
| globfree | gmtime | gmtime_r | grantpt | group_from_gid | Hhcreate | hcreate_r |
| hdestroy | hdestroy_r | herror | hesiod | hesiod_end | hesiod_free_list | hesiod_init |
| hesiod_resolve | hesiod_to_bind | history | history_end | history_init | hostalias | hsearch |
| hsearch_r | hstrerror | hstrerror | htonl | htons | hypot | liconv |
| iconv_close | iconv_open | ieee | ieee_test | ilogb | index | inet_addr |
| inet_aton | inet_cidr_ntop | inet_cidr_pton | inet_lnaof | inet_makeaddr | inet_netof | inet_network |
| inet_ntoa | infnan | initstate | ioctl | iruserok | isalnum | isalpha |
| isascii | isatty | isblank | iscntrl | isdigit | isgraph | isinf |
| isinff | islower | isnan | isnanf | isprint | ispunct | isspace |
| isupper | isxdigit | iswalnum | iswalpha | iswblank | iswcntrl | iswctype |
| iswdigit | iswgraph | iswlower | iswprint | iswpunct | iswspace | iswupper |
| iswxdigit | Jj0 | j1 | jn | jrand48 | Kkill | killpg |
| LI64a | labs | lchown | lcong48 | ldexp | ldiv | lfind |
| lgamma | lgamma_r | libm | link | listen | localeconv | localtime |
| localtime_r | lockf | log | log10 | log1p | logb | longjmp |
| lrand48 | lsearch | lseek | lstat | Mmadvise | major | malloc |
| math | mblen | mbsntowcs | mbstowcs | mbtowc | mctl | MD4Data |
| MD4End | MD4File | MD4Final | MD4Init | MD4Transform | MD4Update | MD5Data |
| MD5End | MD5File | MD5Final | MD5Init | MD5Transform | MD5Update | memccpy |
| memchr | memcmp | memcntl | memcpy | memmove | memset | minor |
| mkdev | mkdir | mkfifo | mknod | mktemp | mktime | mmap |
| modf | mprotect | mrand48 | msgctl | msgget | msgrcv | msgsnd |
| msync | munmap | Nnextafter | nftw | nice | nl_langinfo | nrand48 |
| ns_find_tsig | ns_sign | ns_sign_tcp | ns_sign_tcp_init | ns_verify | ns_verify_tcp | ns_verify_tcp_init |
| ntohl | ntohs | Oopen | opendir | openlog | openpty | Pp_query |
| path_casesensitive | pathconf | pause | pclose | perror | pipe | pmap_getmaps |
| pmap_getport | pmap_set | pmap_unset | poll | popen | pow | pread |
| printf | pthread_atfork | pthread_attr_destroy | pthread_attr_getdetachstate | pthread_attr_getguardsize | pthread_attr_getinheritsched | pthread_attr_getsched |
| pthread_attr_getschedpolicy | pthread_attr_getscope | pthread_attr_getstack | pthread_attr_getstackaddr | pthread_attr_getstacksize | pthread_attr_init | pthread_attr_setdetach |
| pthread_attr_setguardsize | pthread_attr_setinheritsched | pthread_attr_setschedparam | pthread_attr_setschedpolicy | pthread_attr_setscope | pthread_attr_setstack | pthread_attr_setstacka |
| pthread_attr_setstacksize | pthread_cancel | pthread_cleanup_pop | pthread_cleanup_push | pthread_cond_broadcast | pthread_cond_destroy | pthread_cond_init |
| pthread_cond_signal | pthread_cond_timedwait | pthread_cond_wait | pthread_condattr_destroy | pthread_condattr_getpshared | pthread_condattr_init | pthread_condattr_setp |
| pthread_cond_init | pthread_create | pthread_detach | pthread_equal | pthread_exit | pthread_getconcurrency | pthread_getschedpara |
| pthread_getspecific | pthread_join | pthread_key_create | pthread_key_delete | pthread_kill | pthread_mutex_destroy | pthread_mutex_init |
| pthread_mutex_lock | pthread_mutex_timedlock | pthread_mutex_trylock | pthread_mutex_unlock | pthread_mutexattr_destroy | pthread_mutexattr_getpshared | pthread_mutexattr_ge |
| pthread_mutexattr_init | pthread_mutexattr_setpshared | pthread_mutexattr_settype | pthread_once | pthread_rwlock_destroy | pthread_rwlock_init | pthread_rwlock_rdlock |
| pthread_rwlock_timedrdlock | pthread_rwlock_timedwrlock | pthread_rwlock_tryrdlock | pthread_rwlock_trywrlock | pthread_rwlock_unlock | pthread_rwlock_wrlock | pthread_rwlockattr_de |
| pthread_rwlockattr_getpshared | pthread_rwlockattr_init | pthread_rwlockattr_setpshared | pthread_self | pthread_setcancelstate | pthread_setcanceltype | pthread_setconcurren |
| pthread_setschedparam | pthread_setschedprio | pthread_setspecific | pthread_sigmask | pthread_spin_destroy | pthread_spin_init | pthread_spin_lock |
| pthread_spin_trylock | pthread_spin_unlock | pthread_testcancel | ptsname | putc | putchar | putenv |
| puts | pututxline | putw | putwc | putwchar | putws | pwrite |
| Qqsort | Rradixsort | raise | rand | rand_r | rand48 | random |
| rcmd | read | readdir | readdir_r | readlink | readv | realloc |
| realpath | recv | recvfrom | regcomp | regerror | regexec | regfree |
| registerrpc | remainder | remove | rename | renamewtmpx | res_close | res_findzonecut |
| res_getservers | res_hostalias | res_init | res_isourserver | res_mkquery | res_nclose | res_ninit |
| res_nmkquery | res_nmkupdate | res_nquery | res_nquerydomain | res_nsearch | res_nsend | res_nsendsigned |
| res_nupdate | res_ourserver_p | res_pquery | res_query | res_querydomain | res_search | res_send |
| res_setservers | res_update | rewind | rewinddir | rexec | rindex | rint |

| | | | | | | |
|---|---|---|---|---|---|---|
| RMD160Data | RMD160End | RMD160File | RMD160Final | RMD160Init | RMD160Transform | RMD160Update |
| rmdir | rpc | rresvport | ruserok | Ssbrk | scalb | scalbn |
| scalbnf | scandir | scanf | sched_get_priority_max | sched_get_priority_min | sched_yield | seed48 |
| seekdir | select | sem_destroy | sem_getvalue | sem_init | sem_post | sem_timedwait |
| sem_trywait | sem_wait | semctl | semget | semop | send | sendto |
| setbuf | setbuffer | setegid | setenv | seteuid | setgid | setgrent |
| sethostent | setitimer | setjmp | setkey | setlinebuf | setlocale | setlogmask |
| setmode | setnetent | setpgid | setpriority | setprotoent | setpwent | setregid |
| setreuid | setrlimit | setsecret | setservent | setsid | setsockopt | setstate |
| settimeofday | setuid | setuser | setutxent | setvbuf | SHA1Data | SHA1End |
| SHA1File | SHA1Final | SHA1Init | SHA1Transform | SHA1Update | shmat | shmctl |
| shmdt | shmget | shutdown | sigaction | sigaddset | sigblock | sigdelset |
| sigemptyset | sigfillset | sighold | sigignore | sigismember | siglongjmp | sigmask |
| signal | significand | sigpause | sigpending | sigprocmask | sigrelse | sigset |
| sigsetjmp | sigsetmask | sigsetops | sigsuspend | sigvec | sigwait | sin |
| sinf | sinh | sleep | snprintf | sockaddr_in | socket | socketpair |
| sprintf | sqrt | sqrtf | srand | srand48 | srandom | sscanf |
| stat | statvfs | stdarg | strcasecmp | strcasestr | strcat | strchr |
| strcmp | strcoll | strcpy | strcspn | strdup | strerror | strerror_r |
| strfmon | strftime | stripdomainprefix | strlcat | strlcpy | strlen | strmode |
| strncasecmp | strncasestr | strncat | strncmp | strncpy | strpbrk | strptime |
| strrchr | strsep | strsignal | strsigname | strspn | strstr | strtod |
| strtok | strtok_r | strtol | strtoul | strunvis | strvis | strvisx |
| strxfrm | svc_destroy | svc_freeargs | svc_getargs | svc_getreq | svc_getreqset | svc_register |
| svc_run | svc_sendreply | svc_unregister | svcerr_auth | svcerr_decode | svcerr_noproc | svcerr_noprog |
| svcerr_progvers | svcerr_systemerr | svcerr_weakauth | svcfd_create | svcraw_create | svctcp_create | svcudp_bufcreate |
| swab | symlink | syscall | sysconf | syslog | system | Tt_accept |
| t_alloc | t_bind | t_close | t_connect | t_error | t_free | t_getinfo |
| t_getprotaddr | t_getstate | t_listen | t_look | t_open | t_optmgmt | t_rcv |
| t_rcvconnect | t_rcvdis | t_rcvrel | t_rcvudata | t_rcvuderr | t_snd | t_snddis |
| t_sndrel | t_sndudata | t_strerror | t_sync | t_unbind | tan | tanf |
| tanh | tcdrain | tcflow | tcflush | tcgetattr | tcgetpgrp | tcsendbreak |
| tcsetattr | tcsetpgrp | tdelete | telldir | tempnam | tfind | tgetent |
| tgetflag | tgetnum | tgetstr | tgoto | time | times | tmpfile |
| tmpnam | toascii | tolower | toupper | towlower | towupper | tputs |
| truncate | tsearch | tsig | ttyname | ttyname_r | ttyslot | twalk |
| tzset | Uualarm | uidtontsid | ulimit | umask | uname | ungetc |
| ungetwc | unixpath2win | unlink | unlockpt | unsetenv | unvis | user_from_uid |
| usleep | utime | Vvarargs | verr | verrx | vfork | vfprintf |
| vfscanf | vis | vprintf | vscanf | vsnprintf | vsprintf | vsscanf |
| vsyslog | vwarn | vwarnx | Wwait | waitpid | warn | warnx |
| wcs_access | wcs_chdir | wcs_chmod | wcs_chown | wcs_chroot | wcs_fnmatch | wcs_lchown |
| wcs_link | wcs_lstat | wcs_mkdir | wcs_mkfifo | wcs_mknod | wcs_open | wcs_opendir |
| wcs_pathconf | wcs_readlink | wcs_regcomp | wcs_regexec | wcs_remove | wcs_rename | wcs_renamewtmpx |
| wcs_rmdir | wcs_stat | wcs_statvfs | wcs_symlink | wcs_unlink | wcs_utime | wcscasecmp |
| wcscasewcs | wcscat | wcschr | wcscmp | wcscoll | wcscpy | wcscspn |
| wcsdup | wcsftime | wcslcat | wcslcpy | wcslen | wcsncasecmp | wcsncat |
| wcsncmp | wcsncpy | wcspbrk | wcsrchr | wcssep | wcsspn | wcsstr |
| wcstod | wcstok | wcstol | wcsntombs | wcstombs | wcstoul | wcswcs |
| wcswidth | wcsxfrm | wctomb | wctype | wcwidth | winpath2unix | write |
| writev | Xxdr | xdr_accepted_reply | xdr_array | xdr_authunix_parms | xdr_bool | xdr_bytes |
| xdr_callhdr | xdr_callmsg | xdr_char | xdr_destroy | xdr_double | xdr_enum | xdr_float |
| xdr_free | xdr_getpos | xdr_inline | xdr_int | xdr_long | xdr_opaque | xdr_opaque_auth |

| | | | | | | |
|---|---|---|---|---|---|---|
| xdr_pmap | xdr_pmaplist | xdr_pointer | xdr_reference | xdr_rejected_reply | xdr_replymsg | xdr_setpos |
| xdr_short | xdr_string | xdr_u_char | xdr_u_long | xdr_u_short | xdr_union | xdr_vector |
| xdr_void | xdr_wrapstring | xdrmem_create | xdrrec_create | xdrrec_endofrecord | xdrrec_eof | xdrrec_skiprecord |
| xdrstdio_create | xprt_register | xprt_unregister | Yy0 | y1 | yn | Zz_abs |

**See Also**
**Other Resources**
Microsoft Web site

# Anonymous Users and Traverse Checking Settings

Applies To: Windows Server 2003 R2

On computers running Windows XP and Windows Server 2003 operating systems, the Everyone group no longer includes anonymous users by default. This change reduces the number of network resources available by default to anonymous users and simplifies how network administrators can control access by anonymous users.

**Implications of Limiting Anonymous Access**
With the default denial of anonymous user access, it is easier for administrators to configure a secure system.

The default access control lists (ACLs) on earlier versions of Windows that granted the Everyone group access to resources, and potentially exposed them to attack, no longer grant this access to anonymous users after the computer has been upgraded to Windows XP or Windows Server 2003 operating systems.

Anonymous users cannot accidentally be granted access to resources as in the past, when administrators may not have been aware that anonymous users were included in the Everyone group.

This change affects anonymous users attempting to access resources hosted on computers running Windows XP and Windows Server 2003 operating systems. When a Windows 2000-based system is upgraded to Windows XP or Windows Server 2003, resources with ACLs that grant access to everyone (and not explicitly to Anonymous Logon), are no longer available to anonymous users after the upgrade. In most cases, this is an appropriate restriction of anonymous access.

You can still allow anonymous access to selected shared directories and files by adding the Anonymous Logon group to the discretionary access control lists (DACLs) protecting those resources. In addition, you should grant the **Bypass Traverse Checking** user right to the Anonymous Logon group. For more information, see. To bypass traverse checking for anonymous users in Windows XP and Windows Server 2003 operating systems.

In some situations, it may be difficult to determine which resources must grant anonymous access, or to modify the permissions on all the necessary resources. If so, you can configure Windows XP and Windows Server 2003 to permit anonymous access by the Everyone group. For more information, see To allow anonymous access by the Everyone group in Windows XP and Windows Server 2003 operating systems.

**To bypass traverse checking for anonymous users in Windows XP and Windows Server 2003 operating systems**

1. Click **Start**, point to **Control Panel**, point to **Administrative Tools**, and then click **Local Security Policy** (or **Domain Security Policy** on a domain controller).

2. In the console tree, open **Security Settings**, open either **Local Policies** or **Domain Policies**, and then click **User Rights Assignment**.

3. In the details pane, right-click **Bypass traverse checking**, and then click **Properties**.

4. Click **Add User or Group**.

5. In the **Select Users, Computers, or Groups** dialog box, type **Anonymous Logon** in the **Enter the object names to select** list box.

6. Click **Check names** to verify that your entry is valid, and then click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure. |

**To allow anonymous access by the Everyone group in Windows XP and Windows Server 2003 operating systems**

1. Click **Start**, point to **Control Panel**, point to **Administrative Tools**, and then click **Local Security Policy** (or **Domain Security Policy** on a domain controller).

2. In the console tree, open **Security Settings**, open either **Local Policies** or **Domain Policies**, and then click **Security**

**Options**.

3. In the details pane, right-click **Network access: Let Everyone permissions apply to anonymous users**, and then click **Properties**.

4. To allow permissions applied to the **Everyone** group to apply to anonymous users, click **Enabled**.

5. Or, to prevent permissions applied to the **Everyone** group from applying to anonymous users, click **Disabled**.

6. Click **OK**.

| **Note** |
| --- |
| There is no command-line method for this procedure. |

# Print Management feature overview

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2

Print Management in the Microsoft® Windows Server™ 2003 R2 operating system provides a single interface that administrators can use to perform printer and print server management tasks efficiently with detailed control. You can use Print Management from any computer running Windows Server 2003 R2, and you can manage all network printers on print servers running Windows 2000 Server, Windows Server 2003, or Windows Server 2003 R2.

Use Print Management to perform the following tasks:

- Add printers to the local print server automatically
- Add and remove print servers
- Create a new printer filter
- Perform printer management tasks
- View extended features for your printer
- Deploy printers by using Group Policy
- Resources for Print Management

© 2005 Microsoft Corporation. All rights reserved.

# Add printers to the local print server automatically

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Vista

Print Management (Printmanagement.msc) can automatically detect all printers located on the same subnet as the print server, install the appropriate printer drivers, set up the queues, and share the printers. Unless a printer driver cannot be found, no intervention is needed.

If you want to automatically detect printers for a branch office or another location beyond the subnet of your server, use Remote Desktop to log on to the remote print server and start the process from that computer. You need to have administrator permissions on the computer where the printers will be added. Once you connect to the remote server by using Remote Desktop, you can search for network printers by starting Print Management or by running `\windows\pmcsnap\fnprinters.exe` from the command line.

| **Note** |
| --- |
| The print server role must be installed and you must be a member of the **Administrators** group to perform these procedures. |

**To add network printers to the local printer server automatically**

1. Open the Administrative Tools folder, and then double-click **Print Management**.
2. In the Printer Management tree, right-click *ServerName* **(local)**, click **Automatically Add Network Printers**, and then click **Start** to automatically locate and install printers found on this server's local subnet.

| **Note** |
| --- |
| You may be prompted to specify which driver you want to install for a printer, but otherwise, the process is automatic. |

# Add and remove print servers

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Vista

Print Management (Printmanagement.msc) allows you to manage printers that are running on print servers running Windows 2000 or later.

| **Note** |
| --- |
| The print server role must be installed and you must be a member of the Administrators group to perform these procedures. |

**To add print servers to Print Management**

1. Open the Administrative Tools folder, and then double-click **Print Management**.
2. In the Print Management tree, right-click **Print Management**, and then click **Add/Remove Servers**.
3. In the **Add/Remove Servers** dialog box, under **Specify print server**, in **Add server**, do one of the following:
   - Type the name.
   - Click **Browse** to locate and select the print server.
4. Click **Add to List**.
5. Add as many print servers as you want, and then click **OK**.

| **Note** |
| --- |
| You can add the local server on which you are working by clicking **Add the Local Server**. |

**To remove print servers from Print Management**

1. Open the Administrative Tools folder, and then double-click **Print Management**.
2. In the Print Management tree, right-click **Print Management**, and then click **Add/Remove Servers**.
3. In the **Add/Remove Servers** dialog box, under **Print servers**, select one or more servers, and click **Remove**.

# Create a new printer filter

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, Windows Vista

Filters display only those printers that meet a certain set of criteria. For example, it might be helpful to filter for printers with certain error conditions or those printers in a group of buildings regardless of the print server they use. Filters are stored in the Custom Printer Filters folder in the Print Management tree and are dynamic, so the data is always current.

Two default filters are provided with Print Management (Printmanagement.msc). For each filter that you create, you have the option to set up an e-mail notification or to run a script when the conditions of the filter are met. This is useful when you want to be alerted about printer problems, particularly in an organization with multiple buildings and administrators.

For example, you can set up a filter of all printers managed by a particular print server where the status does not equal Ready. Then, if a printer changes from the Ready status to any other status, the administrator could receive a notification e-mail from Print Management.

| **Note** |
| --- |
| The print server role must be installed and you must be a member of the **Administrators** group to perform these procedures. |

**To set up and save a filtered view**

1. Open the Administrative Tools folder, and then double-click **Print Management**.
2. In the Print Management tree, right-click the **Custom Printer Filters** folder, and then click **Add New Printer Filter**. This will launch the **New Printer Filter Wizard**.
3. On the **Printer Filter Name and Description** wizard page, type a name for the printer filter. The name will appear in the **Custom Printer Filters** folder in the Print Management tree.
4. In **Description**, type an optional description.
5. To display the number of printers that satisfy the conditions of a filter, select the **Display the total number of printers next to the name of the printer filter** check box
6. Click **Next**.
7. On the **Define a printer filter** wizard page, do the following:
   a. In the **Field** list, click the print queue or printer status characteristic.
   b. In the **Condition** list, click the condition.
   c. In the **Value** box, type a value.
   d. Continue adding criteria until your filter is complete, and then click **Next**.
8. On the **Set Notifications (Optional)** wizard page, do one or both of the following:
   - To set an e-mail notification, select the **Send e-mail notification** check box, and type one or more recipient and sender e-mail addresses. An SMTP server must be specified to route the message. Use the format *account@domain* and semicolons to separate multiple accounts.
   - To set a script to run, select the **Run script** check box, and then type the path where the script file is located. To add more arguments, type them in **Additional arguments**.
9. Click **Finish**.

| **Note** |
| --- |

To set notifications on existing printer filters, right-click a filtered view, and then click **Set Notifications**.

**Setting Optional Notifications**

When you create or modify a filter, you have the option of sending an automatic e-mail notification to someone or running a script when the conditions of the filter are met. This is useful for resolving printer problems, particularly in an organization with multiple buildings and administrators.

For example, you can set up a view of all printers managed by a particular print server where the status does not equal Ready. Then, if a printer changes from the Ready status to another status, the administrator could receive a notification e-mail from Print Management.

In addition to setting notifications on a custom set of printers, you can set notifications on printer server objects. For example, if the server is offline, or the spooler goes down, an e-mail notification can be sent.

Setting script notifications are useful for resolving printer problems and troubleshooting. For example, you could automatically run a script to restart a spooler when printers go offline. You could also automatically run a script that prints a test page.

Scripts can be written in Visual Basic Script (.vbs) or any scripting language available on the computer. The script has to be on the computer with Print Management. The script should be running with your credentials and you need to have the permissions to do whatever it is that you want the script to do.

# Perform printer management tasks

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Vista

You can use Print Management (Printmanagement.msc) to manage all of the printers in your enterprise, including printers in branch offices. You can use the same interface to perform bulk operations, update drivers, and list or remove printers from Active Directory.

| **Note** |
| --- |
| The print server role must be installed and you must be a member of the **Administrators** group to perform these procedures. |

**Perform Bulk Operations**

Bulk operations can be performed on all the printers on a particular server or all the printers under a particular filter.

The following is a list of the tasks that are possible to perform on multiple printers simultaneously:

- Pause Printing
- Resume Printing
- Cancel All Jobs
- List in Directory
- Remove from Directory
- Delete

**Update Device Drivers**

Printer driver details are shown separately for each server that you have added to Print Management. By default, the driver name, driver version, environment, and provider columns are visible. You can add columns to show more details such as the config file path, data file path, default datatype, driver file path, manufacturer, and so on.

You can export the list of details to a text file for use outside of Print Management.

### To manage device drivers

1. Open the Administrative Tools folder, and then double-click **Print Management**.
2. In the Print Management tree, under the appropriate print server, right-click **Drivers**, and then click **Manage Drivers**.
3. Do one or more of the following:
    - To add a printer driver, click **Add**, and then complete the **Add Printer Driver Wizard**.
    - To remove a printer driver, select a driver from the list, and then click **Remove**.
    - To reinstall a printer driver, select a driver, and then click **Reinstall**.
    - To view the details for a printer driver, select the driver, and then click **Properties**.

**List and Remove Printers in Active Directory**

Listing printers in Active Directory makes it easier for users to locate and install printers. After you install printers on a printer server, you can use Print Management to list them in Active Directory.

In fact, you can list more than one printer simultaneously. You may want to set up a filter to show all of the printers to list or remove, so that you can easily select all of the printers at the same time.

**To list or remove printers in Active Directory**

1. Open the Administrative Tools folder, and then double-click **Print Management**.
2. In the Print Management tree, under the appropriate print server, click **Printers**.
3. In the results pane, right-click the printer you want to list or remove, and then click **List in Directory** or **Remove from Directory**.

# View extended features for your printer

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Vista

In Print Management (Printmanagement.msc), extended view displays beneath the columns in the right pane, separated by a splitter bar. Extended view is useful when you want more information about the status of a print job, its owner, the number of pages, the size of the job, when it was submitted, its port, its priority, and so on.

In addition, when the printer has a Web page, extended view displays a **Printer Web Page** tab. The Web page provides details about the physical properties of the printer and specifications, and sometimes allows remote administration.

| **Note** |
| --- |
| The print server role must be installed and you must be a member of the **Administrators** group to perform these procedures.. |

**To show extended view**

1. Open the Administrative Tools folder, and then double-click **Print Management**.
2. In the Print Management tree, under any print server, right-click **Printers**, and then click **Show Extended View**.

| **Note** |
| --- |
| To display or hide columns in extended view, select a printer. On the **Jobs** tab, right-click the column heading row, and then click the name of the column that you want to display or hide. |

# Deploy printers by using Group Policy

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Vista

Print Management (Printmanagement.msc) can be used with Group Policy to automatically add printer connections to a computer's Printers and Faxes folder.

To do this, you use the **Deploy with Group Policy** dialog box to automatically add a printer connection setting to an existing Group Policy object (GPO) in Active Directory. When Group Policy processing runs on client computers, the printer connection settings are applied to the users or computers associated with the GPO. This is called deploying printer connections. Printers you deploy by using this method appear in the **Deployed Printers** object of Print Management tree when the print server they are connected to is being monitored.

This method of installing a printer is useful in a laboratory, classroom, or branch office setting where every computer in the room or office needs access to the same printer. It is also useful in large organizations, where computers and printers are often separated by function, workgroup, or department, such as marketing or human resources.

A printer connection that has been installed by using a per-user connection is available to the user no matter what computer the user logs on to in the network. A printer connection that has been installed by using a per-machine connection appears in the Printers and Faxes folder, ready for use by any user of that computer.

| **Note** |
| --- |
| Only per-user printer connections are supported on computers running Windows 2000. Per-machine printer connections are supported on computers running Windows XP or later. |

To deploy printer connections using Group Policy, add the PushPrinterConnections.exe utility to a machine startup script (for per-machine connections) or to a user logon script (for per-user connections). The utility reads the printer connection settings made in the GPO and adds the printer connection.

| **Important** |
| --- |
| In order for the PushPrinterConnections.exe utility to work, you must update your Active Directory schema with the Windows Server 2003 R2 changes. For more information about these schema updates, see the Microsoft Web site (http://go.microsoft.com/fwlink?linkid=51166). |

Once you deploy printer connection settings to a GPO, the PushPrinterConnections.exe utility will run on the client computers upon logon or restart. The utility reads the settings from the GPO and adds the deployed printer connections to the users or computers to which the GPO applies. You can reuse the GPO used to deploy per-user or per-machine printer connections. For more information about Group Policy and Printers, see Set Group Policy for Printers on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50998).

It is a good idea to use the same GPO for deploying the printer connection settings and the PushPrinterConnections.exe computer startup or user logon script. This ensures that only users (or computers) that receive the printer connection settings will run the PushPrinterConnections.exe utility.

| **Important** |
| --- |
| The print server role must be installed and you must be a member of the **Administrators** group to perform these procedures. You must also have write access to the Group Policy object to use it to manage printers. Before you install printers by using Group Policy, you must have a GPO for your printer connections settings that is assigned to the appropriate users and computers. You can use the Group Policy Object Editor or Active Directory Users and Computers to create a GPO. |

For more information about using Group Policy, see Enterprise Management with the Group Policy Management Console on the Microsoft Web site (http://go.microsoft.com/fwlink?linkid=22814).

**To install printers to groups of users or computers by using Group Policy**

1. Open the Administrative Tools folder, and then double-click **Print Management**.

2. In the Print Management tree, under the appropriate print server, click **Printers**.
3. In the results pane, right-click the printer you want to deploy, and then click **Deploy with Group Policy**.
4. In the **Deploy with Group Policy** dialog box, click **Browse**, and then choose a Group Policy object.
5. Click **OK**.
6. To assign the printer connection setting to the GPO, do one or both of the following:
   - As a per-user setting, select the **The users that this GPO applies to (per user)** check box.
   - As a per-machine setting, select the **The computers that this GPO applies to (per machine)** check box.
7. Click **Add**.
8. Repeat steps 3 to 6 to add the printer connection setting to another GPO.
9. Click **OK**.

**To use the PushPrinterConnections.exe file**

1. Using Group Policy Management console (gpmc.msc), right-click the GPO with your printer connections settings, and then click **Edit**.
2. In the Group Policy Object Editor tree, navigate to one following locations:
   - If the printer connections are deployed per-machine, go to **Computer Configuration**, **Windows Settings**, **Scripts (Startup/Shutdown)**.
   - If the printer connections are deployed per-user, go to **User Configuration**, **Windows Settings**, **Scripts (Logon/Logoff)**.
3. Right-click **Startup** or **Logon**, and then click **Properties**.
4. In the **Logon Properties** or **Startup Properties** dialog box, click **Show Files**.
5. Copy the PushPrinterConnections.exe file to this location and then close the window.
6. In the **Logon Properties** or **Startup Properties** dialog box, click **Add**.
7. Type **PushPrinterConnections.exe** in the **Script Name** box.
8. If you want to enable logging, type **–log** in the **Script Parameters** box. Log files are written to %windir%\temp\ppcMachine.log (for per-computer connections) and %temp%\ppcUser.log (for per-user connections) on the computer on which the policy is applied.
9. Click **OK**.

| **Note** |
| --- |
| For per-computer connections, the printer connections will be added when the client computer restarts. For per-user connections, the printer connections will be added when the user logs on. If you remove the printer connection settings from the GPO, PushPrinterConnections.exe will remove the corresponding printers from the client computer on the next restart or user logon. |

# Resources for Print Management

Updated: August 22, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Vista

The following resources provide additional information about Print Management.

- For information about using scripts, see Script Repository: Printing on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50001).

- For more information about Print Services, see Print Services on the Microsoft Web site (http://go.microsoft.com/fwlink/?LinkId=50121).

# Windows Server 2003: Security and Protection

This collection contains information about security technologies in Windows Server 2003.

- Auditing

  Security auditing is one of the most powerful tools to help maintain the security of your system. Auditing should identify attacks, either successful or not, that pose a threat to your network, or attacks against resources that you have determined to be valuable in your risk assessment.

- Authorization and Access Control

  Access control is the process of authorizing users, groups, and computers to access objects on the network by using permissions, user rights, and object auditing.

- Authorization Manager

  Authorization Manager is a Microsoft Management Console (MMC) snap-in that can help provide effective control of access to resources.

- Kerberos Authentication

  Kerberos is an authentication mechanism used to verify the identity of a user or host.

- Logon and Authentication

  Logon and authentication includes passwords, Kerberos, NTLM, Transport Layer Security/Secure Sockets Layer (TLS/SSL), and Digest. In addition, some protocols are combined into authentication packages, such as Negotiate and Schannel, as part of an extensible authentication architecture.

- Public Key Infrastructure

  A public key infrastructure (PKI) is a system of digital certificates, certification authorities (CAs), and other registration authorities that verify and authenticate the validity of each party involved in an electronic transaction.

- Security Configuration Wizard

  Security Configuration Wizard (SCW) is an attack-surface reduction tool that guides administrators in creating security policies based on the minimum functionality required for a server's role or roles.

- Security Policy

  Security policy is the configurable set of rules that the operating system follows when determining the permissions to grant in response to a request for access to resources.

- Smart Cards

  Smart cards are used to securely store certificates, public and private keys, passwords, and other types of personal information.

**Additional Resources**

- Secure Windows Server 2008

- Security Tools [LHS Nav page]

- Threats and Countermeasures Guide: Security Settings in Windows Server 2003 and Windows Vista
- Using Windows Server 2003 with Service Pack 1 in a Managed Environment: Controlling Communication with the Internet
- Windows Server 2003 Security Compliance Management Toolkit
- Windows Vista: Security and Protection [RTM]
- Internet Explorer 8 Enhanced Security Configuration

# Auditing

Updated: March 16, 2008

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This page lists resources for security auditing in Windows Server 2003. Security auditing is one of the most powerful tools to help maintain the security of your system. Auditing should identify attacks, either successful or not, that pose a threat to your network, or attacks against resources that you have determined to be valuable in your risk assessment.

**Getting Started**

- Checklist: Setting up object access auditing

  This checklist provides an overview of the steps needed to set up object access auditing.

**Planning and Architecture**

- Auditing Security Events Best practices

  This topic lists best practices for auditing security events.

**Operations**

- Auditing Security Events How To...

  Procedures are available for the following auditing tasks:

  - Define or modify auditing policy settings for an event category
  - Apply or modify auditing policy settings for a local file or folder
  - Apply or modify auditing policy settings for an object using Group Policy
  - View the security log

**Technical Reference**

- Auditing Policy

  This reference describes the event categories that you can choose to audit.

**Additional Resources**

- For information about auditing in Windows Server 2008, see Security Auditing.

# Authorization and Access Control

Updated: March 16, 2008

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This page lists resources for authorization and access control in Windows Server 2003. Access control is the process of authorizing users, groups, and computers to access objects on the network by using permissions, user rights, and object auditing.

**Getting Started**

- Security identifiers

  This topic provides information about well-known security identifiers.

**Planning and Architecture**

- Best practices for assigning permissions on Active Directory objects

  This topic lists best practices for assigning permissions on Active Directory objects.

- Best practices for permissions and user rights

  This topic lists best practices for permissions and user rights.

**Deployment**

- Designing a Resource Authorization Strategy

  This deployment guide provides information about establishing a resource authorization method, defining policies for security group management, and delegating security group maintenance.

**Operations**

- Access Control How To...

  Procedures are available for the following access control tasks:

  - Set, View, Change, or Remove Permissions on an Object
  - Secure a file or folder
  - Apply or modify permission entries for objects using Group Policy
  - Take ownership of a file or folder
  - View effective permissions on files and folders

**Technical Reference**

- Authorization and Access Control Technologies

  The following technical references are available for authorization and access control technologies:

  - Security Descriptors and Access Control Lists Technical Reference

    - Access Tokens Technical Reference
    - Permissions Technical Reference
    - Security Principals Technical Reference
    - Security Identifiers Technical Reference
- Data Security Technologies

  The following technical reference is available for data security technologies:

    - Encrypting File System Technical Reference

**Troubleshooting**

- Troubleshooting Access Control

  This guide provides troubleshooting information for access control in Windows XP and Windows Server 2003. It is designed to help you identify and resolve problems that are related to access control.

**Additional Resources**

- For changes in authorization and access control from Windows Server 2003 to Windows Server 2008, see What's New for Access Control in Windows Server 2008 in Changes in Functionality in Windows Server 2008.
- For more information about authorization and access control in Windows Server 2008, see Access Control.

# Authorization Manager

Updated: March 16, 2008

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This page lists resources for Authorization Manager in Windows Server 2003. Authorization Manager provides a flexible framework for integrating role-based access control into applications. It enables administrators who use those applications to provide access through assigned user roles that relate to job functions. Authorization Manager applications store authorization policy in the form of authorization stores that are stored in Active Directory or XML files and apply authorization policy at run time.

**Getting Started**

- Checklist: Before you start using Authorization Manager

  This checklist provides an overview of the steps you should take before you start using Authorization Manager.

**Planning and Architecture**

- Role-Based Access Control for Multi-tier Applications Using Authorization Manager

  This white paper provides information about how Authorization Manager can be used to provide role-based access control for multi-tier applications.

**Operations**

- Authorization Manager How To...

  Procedures are available for the following Authorization Manager tasks:

  - Start Authorization Manager
  - Set Authorization Manager options
  - Assign an application group to a role
  - Assign a Windows user or group to a role
  - Add an authorization rule
  - Refresh Authorization Manager
  - Reload an authorization store
  - Work with Authorization Stores
  - Work with Applications
  - Work with Role, Task, and Operation Definitions
  - Work with Scopes

**Additional Resources**

- For changes in Authorization Manager from Windows Server 2003 to Windows Server 2008, see Authorization Manager in Changes in Functionality in Windows Server 2008.

- For more information about Authorization Manager in Windows Server 2008, see Authorization Manager.

# Kerberos Authentication

Updated: March 24, 2009

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Kerberos is an authentication mechanism used to verify user or host identity. This page contains the information you need to evaluate, plan, and deploy Kerberos, the preferred authentication method for services in Windows Server 2003.

**Product Evaluation**

- What's New in Windows Server 2003 Kerberos Authentication

  This page outlines the new features in Kerberos authentication and provides the basic information administrators need to begin using these features.

**Planning and Architecture**

- Kerberos Protocol Transition and Constrained Delegation

  This document examines ways to authenticate the users of Web applications and discusses how the new extensions to the Kerberos authentication protocol can satisfy these requirements.

- Planning and Implementing Federated Forests in Windows Server 2003

  This paper examines common scenarios where you can deploy Active Directory forest trusts and describes the related technologies in Windows Server 2003.

**Deployment**

- Server Clusters: Security Best Practices, Windows 2000 and Windows Server 2003, General Assumptions

  To ensure security in server cluster environments, organizations can implement these general assumptions and operational best practices for the infrastructure.

- Kerberos authentication for load balanced web sites

  This document explains how to configure Kerberos in a way that enables applications to use its authentication capabilities in a cluster that also uses network load balancing.

- Planning and Implementing Federated Forests in Windows Server 2003

  This paper examines common scenarios where you can deploy Active Directory forest trusts and describes the related technologies in Windows Server 2003.

- Planning and Implementing Multitier Applications Using Windows Server 2003 Security Services

  This white paper introduces Windows Server 2003 security technologies and describes how they operate and interact with each other to allow you to build multitier applications that use these technologies to provide enhanced security.

**Interoperability**

- Solution Guide: Building Security and Directory Solutions for UNIX Using the Windows Server 2003 Active Directory Kerberos and LDAP Services

  This guide covers evaluating, planning, building, and deploying a security and directory infrastructure based on Windows Server 2003 by using Active Directory, Kerberos, and LDAP services.

**Operations**

- Troubleshooting Kerberos Errors

  This white paper helps you troubleshoot Kerberos authentication problems by outlining simple troubleshooting basics, explaining the causes of common Kerberos errors, and summarizing common troubleshooting tools.

- Troubleshooting Kerberos Delegation

  This white paper explains how to troubleshoot delegation issues that can arise in Kerberos authentication scenarios, summarizes required infrastructure, and describes Windows authentication scenarios.

- SPN Query

  This script queries the Active Directory Global Catalog for a security principal with a specified service principal name (SPN).

**Technical Reference**

- Kerberos Authentication Technical Reference

  This reference explains Kerberos authentication and how the Kerberos version 5 protocol and extensions supported by Windows Server 2003 work.

- Utilizing the Windows Authorization Data in Kerberos Tickets for Access Control to Resources

  This document describes the structure of the Windows authorization data that is used by servers in performing access control. The information applies to Microsoft Windows 2000 Server and Windows Server 2003.

**IT Developer**

- HTTP-Based Cross-Platform Authentication via the Negotiate Protocol
  - Part I—Network Infrastructure
  - Part II—SPNEGO Tokens and the Negotiate Protocol
  - Part III—SPNEGO Token Handler API
- Kerberos Protocol Transition and Constrained Delegation
- Sample Scenario Source Files
- .NET Remoting Security Solution, Part 1: Microsoft.Samples.Security.SSPI Assembly
- Planning and Implementing Multitier Applications Using Windows Server 2003 Security Services
- Contoso Sample ASP.NET Extranet Application
- MSDN Library: Win32 and COM Development: Security: Authentication: SSPI
- Server Clusters: Security Best Practices, Windows 2000 and Windows Server 2003, General Assumptions
- SSPI/Kerberos Interoperability with GSSAPI

**Additional Resources**
**Columns**

- Professor Windows: Windows Server 2003 Trust Enhancements

**Tools**

- Windows Server 2003 Resource Kit Tools

**Product Support**

- Help and Support

**Webcasts**

- Introduction to Kerberos

**Kerberos Resources for Windows 2000 Server**
The information in the following articles also applies to Windows Server 2003–based Kerberos:

- Windows 2000 Kerberos Authentication
- Step-by-Step Guide to Kerberos 5 Interoperability
- More Windows 2000 Kerberos Resources

# Logon and Authentication

Updated: March 16, 2008

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This page lists resources for logon and authentication in Windows Server 2003, which includes passwords, Kerberos, NTLM, Transport Layer Security/Secure Sockets Layer (TLS/SSL), and Digest. In addition, some protocols are combined into authentication packages, such as Negotiate and Schannel, as part of an extensible authentication architecture.

**Planning and Architecture**

- Password Best practices

  This topic lists best practices for passwords.

- Stored User Name and Password Best practices

  This topic lists best practices for stored user names and passwords.

**Deployment**

- Designing an Authentication Strategy

  This deployment guide provides information about creating a foundation for authentication, securing the authentication process, extending your authentication framework, enabling supplemental authentication strategies, and educating users about the authentication process.

- Account Passwords and Policies in Windows Server 2003

  This white paper describes the reasons why you should take a structured approach to setting the account and password policy features and provides information about the tools and log files that you can use to troubleshoot account lockouts. In addition, it shows you how to determine from which computer the account lockouts are being generated, as well as the program or service that is generating the lockout.

**Operations**

- Authentication Protocol How To...

  Procedures are available for the following authentication tasks:

  - Configure Kerberos policy
  - Allow a computer to be trusted for delegation for specific services
  - Allow a user to be trusted for delegation for specific services
  - Allow a computer to be trusted for delegation
  - Allow a user to be trusted for delegation
  - Map an account from a trusted non-Windows kerberos realm to a user account

- Password How To...

  Procedures are available for the following password tasks:

- Change an existing password
- Reset a password
- Reset a local user account password with a password reset disk
- Create a password reset disk
- Apply or modify password policy
- Apply or modify account lockout policy
- Create or update a system key

**Technical Reference**

- Logon and Authentication Technologies

  The following technical references are available for logon and authentication technologies:

  - Digest Authentication Technical Reference
  - Interactive Logon Technical Reference
  - Kerberos Authentication Technical Reference
  - TLS/SSL Technical Reference

**Troubleshooting**

- Password Troubleshooting
- Stored User Name and Password Troubleshooting

**Additional Resources**

- For changes in logon and authentication from Windows Server 2003 to Windows Server 2008, see What's New for Identity Management in Windows Server 2008.
- For more information about logon and authentication in Windows Server 2008, see Windows Authentication.
- For information about Kerberos authentication, see Kerberos Authentication.

# Public Key Infrastructure

Updated: March 16, 2008

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This page lists resources for public key infrastructure (PKI) in Windows Server 2003. A PKI is a system of digital certificates, certification authorities (CAs), and other registration authorities that verify and authenticate the validity of each party involved in an electronic transaction.

**Getting Started**

- Checklist: Configuring certificate autoenrollment

  This checklist provides an overview of the steps needed to configure certificate autoenrollment.

- Checklist: Creating a certification hierarchy with an offline root certification authority

  This checklist provides an overview of the steps needed to create a certification hierarchy with an offline root CA.

- Checklist: Credential roaming

  This checklist provides an overview of the steps needed to use credential roaming.

- Checklist: Decommissioning a certification authority

  This checklist provides an overview of the steps needed to decommission a CA.

**Planning and Architecture**

- Certificate Services Best practices

  This topic provides a list of best practices for Certificate Services.

- Certificate Templates Best practices

  This topic provides a list of best practices for certificate templates.

- Implementing and Administering Certificate Templates in Windows Server 2003

  This white paper discusses the best practices in designing, administering, and implementing version 2 certificate templates by using Windows Server 2003 Enterprise Edition and enterprise CAs.

- Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003

  This white paper provides a technical reference and planning guide for PKI administrators who want to perform PKI cross-certification, deploy bridge CAs, and understand how to implement qualified subordination in Windows Server 2003.

**Deployment**

- Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure

  This document is a quick start guide that you can use to set up a Windows Server 2003 PKI.

- Certificate Autoenrollment in Windows Server 2003

  This white paper includes information about autoenrollment, including configuration, certificate renewal, autoenrollment

functions, advanced features, supported hardware, and troubleshooting.

- Designing a Public Key Infrastructure

  This deployment guide provides information about defining certificate requirements, designing and extending your PKI, defining certificate configuration options, creating a certificate management plan, and deploying the PKI.

**Operations**

- Advanced Certificate Enrollment and Management

  This white paper explains several remote deployment scenarios along with the step-by-step procedures to perform X.509 certificate enrollment to implement a secure infrastructure.

- Key Archival and Management in Windows Server 2003

  This white paper covers best practices and procedures in a key recovery strategy as well as migration procedures for moving from a Microsoft Exchange Key Management Server (KMS) environment to a Windows Server 2003–based CA.

- Windows Server 2003 PKI Operations Guide

  This guide provides information about configuring and operating a Windows-based CA, including various operational scenarios, custom configuration information, sample commands, and best practices.

**Technical Reference**

- PKI Technologies

  The following technical references are available for PKI technologies:

  - CA Certificates Technical Reference
  - Certificates Technical Reference
  - Certificate Services Technical Reference

**Troubleshooting**

- Certificate Services Troubleshooting
- Certificate Templates Troubleshooting
- Certificates Troubleshooting

**Additional Resources**

- For changes in PKI from Windows Server 2003 to Windows Server 2008, see Active Directory Certificate Services Role in Changes in Functionality in Windows Server 2008.
- For more information about Active Directory Certificate Services in Windows Server 2008, see Active Directory Certificate Services.

# Security Configuration Wizard

Updated: March 16, 2008

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This page lists resources for the Security Configuration Wizard (SCW) in Windows Server 2003. SCW is an attack-surface reduction tool that guides administrators in creating security policies based on the minimum functionality required for a server's role or roles.

**Getting Started**

- Security Configuration Wizard Quick Start Guide

  This guide provides system requirements, installation instructions, steps for getting started with SCW, and instructions for troubleshooting simple problems.

**Deployment**

- Deployment Guide for the Security Configuration Wizard

  This deployment guide shows you how to deploy SCW and SCW policies to your computers running Windows Server 2003 with Service Pack 1 (SP1).

**Additional Resources**

- For changes in SCW from Windows Server 2003 to Windows Server 2008, see Security Configuration Wizard in Changes in Functionality in Windows Server 2008.

- For more information about SCW in Windows Server 2008, see Security Configuration Wizard.

# Security Policy

Updated: March 16, 2008

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This page lists resources for security policy in Windows Server 2003. Security policy is the configurable set of rules that the operating system follows when determining the permissions to grant in response to a request for access to resources.

**Getting Started**

- Checklist: Using software restriction policies to help protect your computer against an e-mail virus

  This checklist provides an overview of the steps needed to use software restriction policies to help protect your computer against an e-mail virus.

**Planning and Architecture**

- Software Restriction Policies Best Practices

  This topic lists best practices for software restriction policies.

**Operations**

- Software Restriction Policies How To...

  Procedures are available for the following software restriction policies tasks:

  - Create new software restriction policies
  - Open Software Restriction Policies
  - Apply software restriction policies to DLLs
  - Prevent software restriction policies from applying to local administrators
  - Enable certificate rules
  - Create a certificate rule
  - Create a hash rule
  - Create an Internet zone rule
  - Create a path rule
  - Create a registry path rule
  - Add or delete a designated file type
  - Change the default security level of software restriction policies
  - Set trusted publisher options

**Technical Reference**

- Security Policy Settings

  This provides a reference of security policy settings in Windows Server 2003.

**Troubleshooting**

- Software Restriction Policies Troubleshooting

**Additional Resources**

- For information about security policy management in Windows Server 2008, see Server Security Policy Management.

# Smart Cards

Updated: March 16, 2008

Applies To: Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

This page lists resources for smart cards in Windows Server 2003. Smart cards are used to securely store certificates, public and private keys, passwords, and other types of personal information.

**Getting Started**

- Checklist: Deploying smart cards for logging on to Windows

  This checklist provides an overview of the steps needed to deploy smart cards for use in Windows logon.

**Deployment**

- Planning a Smart Card Deployment

  This deployment guide provides information about deploying smart cards, including creating a plan for smart card use, selecting smart card hardware, creating a smart card deployment plan, and planning for ongoing smart card support.

**Additional Resources**

- For changes in smart cards from Windows Server 2003 to Windows Server 2008, see What's New for Identity Management in Windows Server 2008.

# Windows Server 2003 Glossary

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - #**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**$$Rename.txt**
A file that lists all files contained in a particular folder that need to be renamed from an 8.3 format name to a long file name during Setup. All file names must be in 8.3 format if you run Setup using Winnt.exe, or if you copy a set of files from one computer to another using MS-DOS.

**.cab file**
Cabinet file. A single cabinet file that stores multiple compressed files. These files are commonly used in software installation and to reduce the file size and the associated download time for Web content.

**.cer file format**
A file format that is used for the storage of public key certificates.

See also: public key

**.inf**
The file name extension for files that contain device information or scripts to control hardware operations.

See also: device

**.NET Passport**
A service that is part of Microsoft .NET Framework. .NET Passport includes a passport authentication protocol, and it provides the ability to map passport IDs to Active Directory accounts.

See also: Active Directory  authentication  passport authentication  service

**.pfx file format**
A file format that is used for the storage of public key certificates and their associated private keys.

See also: private key  public key

**1.2.840.113556.1.4.1851**
The object identifier that is assigned to Active Directory Application Mode (ADAM). You can use this object identifier programmatically to differentiate ADAM from other directory services.

**100baseT**
The Ethernet standard for local area networks using twisted-pair cable carrying data at 100 megabits per second (Mbps).

**3DES**
See other term: Triple DES (3DES)

**802.1X authentication**
An Institute of Electrical and Electronic Engineers (IEEE) standard for port-based network access control that provides authenticated network access to Ethernet networks and wireless 802.11 local area networks (LANs).

See also: access control  authentication  local area network (LAN)  port  port based network access control

# Windows Server 2003 Glossary - A

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - A**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**AAAA (quad-A) resource record**
A resource record used to map a DNS domain name to a host Internet Protocol version 6 (IPv6) address on the network.

See also: domain name  Domain Name System (DNS)  IP address  resource record (RR)

**AAL**
See other term: ATM adaptation layer (AAL)

**ABR**
See other term: area border router (ABR)

**access control**
A security mechanism that determines which operations a user, group, service, or computer is authorized to perform on a computer or on a particular object, such as a file, printer, registry subkey, or directory service object.

See also: group  object  permission  registry  service

**access control entry (ACE)**
An entry in either a securable object's discretionary access control list (DACL) or an object's system access control list (SACL). In a DACL, the entry grants or denies permissions to a user or group. In a SACL, the entry specifies which security events to audit for a particular user or group or controls the Windows Integrity Level for the object.

See also: access control  access control list (ACL)  discretionary access control list (DACL)  object  permission  security descriptor  system access control list (SACL)

**access control list (ACL)**
A list of security protections that apply to an entire object, a set of the object's properties, or an individual property of an object. There are two types of access control lists: discretionary and system.

See also: access control entry (ACE)  discretionary access control list (DACL)  object  security descriptor  system access control list (SACL)

**access number**
A phone number that subscribers can dial to reach online services.

**access token**
A data structure that contains authentication and authorization information for a user. Windows creates the access token when the user logs on and the user's identity is confirmed. The access token contains the user's security ID (SID), the list of groups that the user is a member of, and the list of privileges held by that user. Each process or thread started for the user inherits a copy of the access token. In some cases a user may have more than one access token, with different levels of authority.

See also: privilege  security ID (SID)  security principal

**account lockout**
A Windows security feature that locks a user account if a number of failed logon attempts occur within a specified amount of time, based on security policy lockout settings. Locked accounts cannot log on.

See also: policy  user account

**account logon event category**
In auditing, a group of events that are logged when a user logs on to a computer that uses another computer to validate the user account. Events in this event category are logged on the computer that validates the user account.

See also: auditing  event  user account

**account management event category**
In auditing, a group of events that are logged when user or group account information is modified.

See also: auditing  event  group account  user account

**account partner**
A federation partner that is trusted by the Federation Service to provide security tokens. The account partner issues these tokens to its users - that is, users in the account partner realm - so that they can access Web-based applications in the resource partner.

**ACE**
See other term: access control entry (ACE)

**ACID**
See definition for: atomicity, consistency, isolation, durability (ACID)

**ACK (acknowledgment)**
A message transmitted to indicate that data has been received correctly. The Transmission Control Protocol (TCP) requires that the recipient acknowledge successful receipt of data. Such acknowledgments (ACKs) generate additional network traffic, decreasing the rate at which data passes but increasing reliability. To reduce the impact on performance, most hosts send an acknowledgment for every other segment or when a specified time interval has passed.

See also: host

**acknowledgment message**
For Message Queuing, a message class that indicates that a message arrived or was retrieved by the target application (a positive acknowledgment), or that an error occurred before the original message could be retrieved (a negative acknowledgment). These messages are sent to administration queues on the source computer.

See also: administration queue  Message Queuing

**ACL**
See other term: access control list (ACL)

**ACPI**
See other term: Advanced Configuration and Power Interface (ACPI)

**active**
Describes the window or icon that you are currently using or that is currently selected. The operating system always applies the next keystroke or command you choose to the active window. Windows or icons on the desktop that are not selected are inactive.

**Active Accessibility**
An application programming interface (API) built into a variety of Windows operating systems, including Windows 2000, Windows XP, and products in the Windows Server 2003 family. Active Accessibility allows an accessibility aid (or other Active Accessibility client) to collect meaningful information from an application`s user interface and convey that information to the user. For example, Microsoft Active Accessibility allows a screen reader to convey to the user that an application is displaying a pop-up message asking whether to save or discard file changes. Because Active Accessibility reduces the need for application-specific code, it reduces overall maintenance costs and allows software developers to innovate in their user interface without sacrificing compatibility with accessibility aids.

See also: application programming interface (API)

**active caching**
A feature in Microsoft Internet Security and Acceleration (ISA) Server that retrieves files for Web pages that are accessed frequently.

See also: caching

**active content**
Interactive or animated content used on the Internet. Active content includes ActiveX controls and web browser add-ons.

See also: channel

**Active Directory**
The Windows-based directory service. Active Directory stores information about objects on a network and makes this

information available to users and network administrators. Active Directory gives network users access to permitted resources anywhere on the network using a single logon process. It provides network administrators with an intuitive, hierarchical view of the network and a single point of administration for all network objects.

See also: directory partition  directory service  domain  forest  object  replication

**Active Directory Application Mode (ADAM)**
A stand-alone directory service that is designed specifically for use with directory-enabled applications. Active Directory Application Mode (ADAM) does not require or depend on Active Directory forests or domains. ADAM stores and replicates only application-related information. ADAM does not store or replicate network operating system (NOS)-related information.

**Active Directory Application Mode (ADAM) instance**
For Active Directory Application Mode (ADAM), a single copy of the ADAM directory service, along with its associated directory store, assigned Lightweight Directory Access Protocol (LDAP) and Secure Sockets Layer (SSL) ports, and application event log. You can run multiple ADAM instances simultaneously on a single ADAM server.

See also: Active Directory Application Mode (ADAM)  directory service  Lightweight Directory Access Protocol (LDAP)  Secure Sockets Layer (SSL)

**Active Directory Federation Services (ADFS)**
A Windows Server 2003 R2 component that provides Web single-sign-on (SSO) technologies to authenticate a user to multiple Web applications over the life of a single online session. ADFS accomplishes this by securely sharing digital identity and entitlement rights across security and enterprise boundaries. ADFS in Windows Server 2003 R2 supports the WS-Federation Passive Requestor Profile (WS-F PRP).

**Active Directory Installation Wizard**
The tool that is used to install and remove Active Directory.

See also: Active Directory

**Active Directory replication**
The synchronization of Active Directory partition replicas between domain controllers. Replication automatically copies the changes that originate on a writable directory partition replica to all other domain controllers that hold the same directory partition replica. More specifically, a destination domain controller pulls these changes from the source domain controller.

See also: Active Directory  domain controller  global catalog  replica  replication

**Active Directory Service Interfaces (ADSI)**
A directory service model and a set of Component Object Model (COM) interfaces. ADSI enables Windows applications and Active Directory clients to access several network directory services, including Active Directory. ADSI is supplied as a software development kit (SDK).

See also: Active Directory  Component Object Model (COM)  directory service

**Active Directory Users and Computers**
An administrative tool used by an administrator to perform day-to-day Active Directory administration tasks. The tasks that can be performed with this tool include creating, deleting, modifying, moving, and setting permissions on objects stored in the directory. Examples of objects in Active Directory are organizational units, users, contacts, groups, computers, printers, and shared file objects.

See also: Active Directory  object  permission

**Active Directory-integrated zone**
A primary Domain Name System (DNS) zone that is stored in Active Directory so that it can use multimaster replication and Active Directory security features.

See also: Active Directory  Domain Name System (DNS)  multimaster replication  primary zone

**active partition**
A partition from which a computer starts up. The active partition must be a primary partition on a basic disk. If you use Windows exclusively, the active partition can be the same as the system volume.

See also: basic disk  primary partition  system partition  system volume  x86

**active volume**
The volume from which the computer starts up. The active volume must be a simple volume on a dynamic disk. You cannot

mark an existing dynamic volume as the active volume, but you can upgrade a basic disk containing the active partition to a dynamic disk. After the disk is upgraded to dynamic, the partition becomes a simple volume that is active.

See also: active partition  basic disk  dynamic disk  dynamic volume  simple volume

**ActiveX**
A set of technologies that allows software components to interact with one another in a networked environment, regardless of the language in which the components were created.

See also: ActiveX component

**ActiveX component**
A reusable software component that can be used to incorporate ActiveX technology.

See also: ActiveX

**ADAM**
See other term: Active Directory Application Mode (ADAM)

**ADAM instance**
See other term: Active Directory Application Mode (ADAM) instance

**address**
A unique identifier that identifies a network node to other nodes on the network. Also known as the net address or MAC address.

See also: media access control (MAC) address  node

**address (A) resource record**
A resource record (RR) used to map a DNS domain name to a host Internet Protocol version 4 (IPv4) address on the network.

See also: domain name  Domain Name System (DNS)  IP address  resource record (RR)

**address class**
A predefined grouping of Internet addresses that defines a network of a certain size. The range of numbers that can be assigned for the first octet in the IP address is based on the address class. Class-based IP addressing has been superceded by Classless Interdomain Routing (CIDR).

See also: Class A IP address  Class B IP address  Class C IP address  Classless Interdomain Routing (CIDR)

**address pool**
The addresses within a DHCP scope range of addresses that are available for leased distribution to clients.

See also: address  Dynamic Host Configuration Protocol (DHCP)  lease  scope

**Address Resolution Protocol (ARP)**
In TCP/IP, a protocol that uses broadcast traffic on the local network to resolve a logically assigned Internet Protocol version 4 (IPv4) address to its physical hardware or media access control (MAC) layer address. In asynchronous transfer mode (ATM), ARP is used two different ways. For classical IPv4 over ATM (CLIP), ARP is used to resolve addresses to ATM hardware addresses. For ATM LAN emulation (LANE), ARP is used to resolve Ethernet/802.3 or Token Ring addresses to ATM hardware addresses.

See also: asynchronous transfer mode (ATM)  classical IP over ATM (CLIP)  Internet Protocol (IP)  IP address  LAN emulation (LANE)  media access control (MAC) address  Message Authentication Code (MAC)  packet  Transmission Control Protocol/Internet Protocol (TCP/IP)

**Address Windowing Extensions (AWE) API**
A set of application programming interfaces (APIs) developed by Microsoft that allows software developers to create applications that use up to 64 GB of physical nonpaged memory in a 32-bit virtual address space on 32-bit platforms. This technology allows for windowed views to this physical memory from within the application`s virtual address space.

See also: application programming interface (API)  virtual address

**adjacency**
A relationship formed between selected neighboring Open Shortest Path First (OSPF) routers for the purpose of exchanging routing information. When the link state databases of two neighboring routers are synchronized, the routers are said to be

adjacent. Not every pair of neighboring routers becomes adjacent.

See also: link state database  Open Shortest Path First (OSPF)  router

**administration queue**
For Message Queuing, a queue that stores acknowledgment messages.

See also: acknowledgment message  Message Queuing  queue

**administrative alerts**
Alerts that relate to server and resource use. They notify users about problems in areas such as security and access, user sessions, server shutdown due to power loss (when an uninterruptible power supply (UPS) is available), directory replication, and printing. When a computer generates an administrative alert, a message is sent to a predefined list of users and computers.

See also: Alerter service  uninterruptible power supply (UPS)

**administrative credentials**
Logon information that is used to identify a member of an administrative group. Groups that use administrative credentials include Administrators, Domain Admins, and DNS Admins. Most system-wide or domain-wide tasks require administrative credentials.

See also: Administrators group  group

**administrator**
In the Windows Server 2003 family, a person who is responsible for setting up and managing local computers, stand-alone servers, member servers, or domain controllers. An administrator sets up user and group accounts, assigns passwords and permissions, and helps users with networking problems. Administrators can be members of the Administrators group on local computers or servers. A person who is a member of the Administrators group on a local computer or server has full access to that computer or server and can assign access control rights to users as necessary. Administrators can also be members of the Domain Admins group on domain controllers and have full control over user and computer accounts residing in that domain.

See also: access control  computer account  domain  domain controller  group account  member server  network administrator  stand-alone server  user account

**administrator account**
On a local computer, the first account that is created when you install an operating system on a new workstation, stand-alone server, or member server. By default, this account has the highest level of administrative access to the local computer, and it is a member of the Administrators group. In an Active Directory domain, the first account that is created when you set up a new domain by using the Active Directory Installation Wizard. By default, this account has the highest level of administrative access in a domain, and it is a member of the Administrators, Domain Admins, Domain Users, Enterprise Admins, Group Policy Creator Owners, and Schema Admins groups.

See also: Active Directory  Administrators group  local computer  member server  stand-alone server  user account

**Administrators group**
On a local computer, a group whose members have the highest level of administrative access to the local computer. Examples of administrative tasks that can be performed by members of this group include installing programs; accessing all files on the computer; auditing access control; and creating, modifying, and deleting local user accounts. In an Active Directory domain, a group whose members have the highest level of administrative access in the domain. Examples of administrative tasks that can be performed by members of this group include setting domain policy; assigning and resetting domain user account passwords; setting up and managing domain controllers; and creating, modifying, and deleting domain user accounts.

See also: access control  Active Directory  auditing  domain  domain controller  group  local computer  object  user account

**ADSI**
See other term: Active Directory Service Interfaces (ADSI)

**Advanced Configuration and Power Interface (ACPI)**
An open industry specification that defines power management on a wide range of mobile, desktop, and server computers and peripherals. ACPI is the foundation for the OnNow industry initiative that allows system manufacturers to deliver computers that start at the touch of a keyboard. ACPI design is essential to take full advantage of power management and Plug and Play.

**affinity**

For Network Load Balancing, the method used to associate client requests to cluster hosts. When no affinity is specified, all network requests are load balanced across the cluster without respect to their source. Affinity is implemented by directing all client requests from the same IP address to the same cluster host.

**agent**
An application that runs on a Simple Network Management Protocol (SNMP) managed device. The agent application is the object of management activities. A computer running SNMP agent software is also sometimes referred to as an agent.

See also: Simple Network Management Protocol (SNMP)

**Alerter service**
A service used by the server and other services to notify selected users and computers of administrative alerts that occur on a computer. The Alerter service requires the Messenger service.

See also: administrative alerts  Messenger service  service

**algorithm**
In cryptography, a mathematical process that is used in cryptographic operations such as the encryption and digital signing of data. An algorithm is commonly used with a cryptographic key to enhance security.

See also: encryption  public key encryption  symmetric encryption

**allocation unit**
The smallest amount of disk space that can be allocated to hold a file. All file systems used by Windows organize hard disks based on allocation units. The smaller the allocation unit size, the more efficiently a disk stores information. If you do not specify an allocation unit size when formatting the disk, Windows picks default sizes based on the size of the volume. These default sizes are selected to reduce the amount of space that is lost and the amount of fragmentation on the volume. Also called a cluster.

See also: file system  volume

**American Standard Code for Information Interchange (ASCII)**
A standard single-byte character encoding scheme used for text-based data. ASCII uses designated 7-bit or 8-bit number combinations to represent either 128 or 256 possible characters. Standard ASCII uses 7 bits to represent all uppercase and lowercase letters, the numbers 0 through 9, punctuation marks, and special control characters used in U.S. English. Most current x86-based systems support the use of extended (or "high") ASCII. Extended ASCII allows the eighth bit of each character to identify an additional 128 special symbol characters, foreign-language letters, and graphic symbols.

See also: Unicode

**Anonymous access**
An authentication mechanism by which users who are able to connect to an Internet site without credentials are assigned to the IUSR_ComputerName account and granted the access rights that are assigned to that account.

See also: access control  Anonymous authentication  authentication

**Anonymous authentication**
An authentication mechanism that does not require user accounts and passwords. Anonymous authentication grants remote users the identity IUSR_ComputerName. Anonymous authentication is used on the Internet to grant visitors restricted access to predefined public resources.

See also: Anonymous access  authentication  password  user account

**answer file**
A text file used to automate Setup or other installation processes. Using this text file, you can provide custom answers to Setup-related questions. Typically, you must point the Setup program to use the answer file at the same time Setup is started. Answer files can only be used on applications and operating systems that support them.

See also: Setup

**API**
See other term: application programming interface (API)

**APIPA**
See other term: Automatic Private IP Addressing (APIPA)

**AppleShare**

Client software that is shipped with all Macintosh computers and with Apple Computer server software. Macintosh computers can use their AppleShare client software to connect to computers running Services for Macintosh.

See also: client  server

**AppleTalk Control Protocol (ATCP)**
A network control protocol defined within Point-to-Point Protocol (PPP). ATCP allows clients to connect to a Windows server by running PPP over AppleTalk. ATCP negotiates the AppleTalk parameters to use during the PPP session.

See also: Network Control Protocol  Point-to-Point Protocol (PPP)

**AppleTalk Filing Protocol**
The presentation-layer protocol that manages access to remote files in an AppleTalk network.

See also: protocol

**AppleTalk Phase 2**
The extended AppleTalk Internet model designed by Apple Computer that supports multiple zones within a network and extended addressing capacity.

See also: zone

**AppleTalk Protocol**
The set of network protocols on which AppleTalk network architecture is based. The AppleTalk Protocol is installed with Services for Macintosh to help users access resources on a network.

See also: protocol  resource

**AppleTalk Transport**
The layer of AppleTalk Phase 2 protocols that delivers data to its destination on the network.

See also: AppleTalk Phase 2

**application assignment**
A process that uses Software Installation (an extension of Group Policy) to assign programs to groups of users. The programs appear on the users` desktop when they log on.

See also: Group Policy

**application directory partition**
An Active Directory directory partition that stores application-specific data that can be dynamic (subject to Time to Live restrictions). Application directory partitions can store any type of object except security principals and are not replicated to the global catalog. The replication scope of an application directory partition can be configured to include any set of domain controllers in the forest.

See also: Active Directory  directory partition  Time to Live (TTL)

**application memory tuning**
A feature of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. Application memory tuning can be enabled on computers with between 2 gigabytes (GB) and 4 GB of physical RAM, and it allows applications running on these computers to address 3 GB of virtual memory instead of the 2 GB normally allocated to each application process. Also known as 4-gigabyte tuning (4GT).

See also: virtual memory

**application programming interface (API)**
A set of routines that an application uses to request and carry out lower-level services performed by a computer`s operating system. These routines usually carry out maintenance tasks such as managing files and displaying information.

**area border router (ABR)**
An Open Shortest Path First (OSPF) router that is attached to multiple areas. ABRs maintain separate topological databases for each area.

See also: Open Shortest Path First (OSPF)  topological database

**ARP**
See other term: Address Resolution Protocol (ARP)

**AS**
See other term: autonomous system (AS)

**ASBR**
See other term: autonomous system boundary router (ASBR)

**ASCII**
See other term: American Standard Code for Information Interchange (ASCII)

**ASR**
See other term: Automated System Recovery (ASR)

**asynchronous communication**
A form of data transmission in which information is sent and received at irregular intervals, one character at a time. Because data is received at irregular intervals, the receiving modem must be signaled to let it know when the data bits of a character begin and end. This is done by means of start and stop bits.

See also: modem (modulator/demodulator)

**asynchronous transfer mode (ATM)**
A high-speed, connection-oriented, virtual circuit-based packet switching protocol used to transport many different types of network traffic. ATM packages data in 53-byte, fixed-length cells that can be switched quickly between logical connections on a network.

**ATCP**
See other term: AppleTalk Control Protocol (ATCP)

**ATM**
See other term: asynchronous transfer mode (ATM)

**ATM adaptation layer (AAL)**
The layer of the asynchronous transfer mode (ATM) protocol stack that parses data into the payload portion of the ATM cell for transport across an ATM network.

See also: asynchronous transfer mode (ATM)

**atomicity, consistency, isolation, durability (ACID)**
The four essential properties of an electronic transaction. Atomicity is the ability to do or undo a transaction completely; consistency is the ability of a transaction to change a system from one consistent state to another consistent state; isolation is the ability for a transaction to occur independently of other transactions made at the same time; and durability is the ability of a transaction to persist, even during system failure.

**attribute**
For files, information that indicates whether a file is read-only, hidden, ready for archiving (backing up), compressed, or encrypted, and whether the file contents should be indexed for fast file searching.

See also: Active Directory  class  object  object class  schema

**audit entry**
Data that is recorded in the event log when specified system, application, and security-related events take place. Audit entries provide valuable data about system operations that can be used to identify system use and to diagnose system behavior.

See also: auditing  event

**audit logging**
Enhanced logging capabilities that monitor and manage the growth and size of log files used by DHCP on an ongoing daily basis.

See also: Dynamic Host Configuration Protocol (DHCP)

**audit policy**
Policy that determines the security events to be reported to the network administrator.

See also: policy  security

**auditing**
The process that tracks the activities of users by recording selected types of events in the security log of a server or a

workstation.

See also: security log

**authentication**
The process of verifying the identity of a user, computer, process, or other entity by validating the credentials provided by the entity. Common forms of credentials are digital signatures, smart cards, biometric data, and a combination of user names and passwords.

See also: cryptography  Kerberos V5 authentication protocol  nonrepudiation  NTLM authentication protocol  smart card  trust relationship

**Authentication Header (AH)**
A header that provides authentication, integrity, and anti-replay for the entire packet (the Internet Protocol (IP) header and the data payload carried in the packet).

See also: authentication  Internet Protocol (IP)  packet

**authentication protocol**
The protocol by which an entity on a network proves its identity to a remote entity. Typically, identity is proved with the use of a secret key, such as a password, or with a stronger key, such as the key on a smart card. Some authentication protocols also implement mechanisms to share keys between client and server to provide message integrity or privacy.

See also: authentication  identity  Kerberos V5 authentication protocol  key  NTLM authentication protocol  protocol  smart card

**authoritative**
Describes a DNS server that hosts a primary or secondary copy of a DNS zone.

See also: authoritative restore  DNS server  Domain Name System (DNS)  resource record (RR)  zone

**authoritative restore**
In Backup, a type of restore operation performed on an Active Directory domain controller in which the objects in the restored directory are treated as authoritative, replacing (through replication) all existing copies of those objects.

See also: Active Directory  domain controller  nonauthoritative restore  object  replication

**authorization**
The process that determines what a user is permitted to do on a computer system or network.

See also: authentication

**authorization store**
A database that stores Authorization Manager policy.

**automated library**
A tape or disk library with an automatic mechanism, such as a robotic arm, that loads and unloads media into a drive or drives without manual intervention from the user. Automated libraries are sometimes known as autoloaders, changers, jukeboxes, media changers, and robotic libraries.

**Automated System Recovery (ASR)**
A feature that helps you recover a system that will not start. To use Automated System Recovery, you must first use the Automated System Recovery Preparation Wizard (part of Backup). This wizard backs up the partition used by the operating system, but it does not back up other partitions, such as program or data partitions. Those partitions must be backed up using Backup or other standard routines.

See also: partition

**automatic file replication**
The replication of files and directories between computers as performed automatically by the File Replication service (FRS). In a domain distributed file system, FRS provides automatic file and directory replication between targets. The user configures automatic replication through the DFS administration tool. Files in a stand-alone distributed file system must be replicated between targets manually or by using other utilities.

See also: Distributed File System (DFS)  domain DFS  File Replication service (FRS)  replication

**Automatic Private IP Addressing (APIPA)**

A TCP/IP feature in Windows XP and Windows Server 2003 that automatically configures a unique IP address from the range 169.254.0.1 through 169.254.255.254 with a subnet mask of 255.255.0.0 when the TCP/IP protocol is configured for dynamic addressing and a DHCP server is not available. The APIPA range of IP addresses is reserved by the Internet Assigned Numbers Authority (IANA) for use on a single subnet, and IP addresses within this range are not used on the Internet.

See also: DHCP server  Dynamic Host Configuration Protocol (DHCP)  IP address  Transmission Control Protocol/Internet Protocol (TCP/IP)

**autonomous system (AS)**
A group of routers exchanging routing information by using a common routing protocol.

See also: autonomous system boundary router (ASBR)  router

**autonomous system boundary router (ASBR)**
A router that exchanges routing information with routers that belong to other autonomous systems. The ASBR then advertises external routes throughout the autonomous system. ASBRs can be internal or area border routers, and they might or might not be connected to the backbone.

See also: area border router (ABR)  autonomous system (AS)  backbone  router

**auto-static updates**
The process of adding static routes to the routing table automatically. When you configure an interface to use auto-static update mode, the router sends a request to other routers and inherits routes. The routes are saved in the routing table as auto-static routes and are kept even if the router is restarted or the interface goes down. Auto-static updates are supported in Routing Information Protocol (RIP) for Internet Protocol (IP) and in RIP for Internetwork Packet Exchange (IPX), but they are not available for use with Open Shortest Path First (OSPF).

See also: Open Shortest Path First (OSPF)  routing  Routing Information Protocol over IPX (RIPX)  static routes

**availability**
A level of service provided by applications, services, or systems. Highly available systems have minimal downtime, whether planned or unplanned. Availability is often expressed as the percentage of time that a service or system is available, for example, 99.9 percent for a service that is down for 8.75 hours a year.

**available bit rate (ABR)**
An asynchronous transfer mode (ATM) service type that supports available-bit-rate traffic, minimum guaranteed transmission rate, and peak data rates. ABR also allows bandwidth allocation depending on availability, and it uses flow control to communicate bandwidth availability to the end node.

See also: asynchronous transfer mode (ATM)  bandwidth  node

**AXFR**
See other term: full zone transfer (AXFR)

# Windows Server 2003 Glossary - B

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - B**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**backbone**
In Open Shortest Path First (OSPF), an area common to all other OSPF areas that is used as the transit area for inter-area traffic and for distributing routing information between areas. The backbone must be contiguous.

See also: Open Shortest Path First (OSPF)

**backbone router**
In Open Shortest Path First (OSPF), a router that is connected to the backbone area. This includes routers that are connected to more than one area (area border routers (ABRs)). However, backbone routers do not have to be ABRs. Routers that have all networks connected to the backbone are internal routers.

See also: area border router (ABR) backbone Open Shortest Path First (OSPF) router

**background program**
A program that runs while the user is working on another task. The computer`s microprocessor assigns fewer resources to background programs than foreground programs.

See also: foreground program

**backup**
A duplicate copy of a program, a disk, or data, made either for archiving purposes or for safeguarding valuable files from loss in case the active copy is damaged or destroyed. Some application programs automatically make backup copies of data files, maintaining both the current version and the preceding version.

**backup domain controller (BDC)**
A domain controller running Windows NT Server 4.0 or earlier that receives a read-only copy of the directory database for the domain. The directory database contains all account and security policy information for the domain.

See also: Active Directory mixed mode primary domain controller (PDC)

**backup media pool**
A logical collection of data-storage media that has been reserved for use by Microsoft Windows Backup. Backup uses Removable Storage to control access to specific media within a library.

See also: library media pool Removable Storage

**Backup Operators group**
A type of local or global group that contains the user rights you need to back up and restore files and folders. Members of the Backup Operators group can back up and restore files and folders regardless of ownership, permissions, encryption, or auditing settings.

See also: auditing global group local group user rights

**backup set**
A collection of files, folders, and other data that has been backed up and stored in a file or on one or more tapes.

See also: backup set catalog on-disk catalog on-media catalog

**backup set catalog**
A summary of the files and folders that have been saved in a backup set.

See also: backup set

**backup types**

A type that determines which data is backed up and how it is backed up. There are five backup types: copy, daily, differential, incremental, and normal.

See also: copy backup  daily backup  differential backup  incremental backup  normal backup

**bandwidth**
The data transfer capacity of a transmission medium. In digital communications, the transfer capacity expressed in bits per second (bps) or megabits per second (Mbps). For example, Ethernet accommodates a bandwidth of 10,000,000 bps or 10 Mbps. In analog communications, the difference between the highest and lowest frequencies in a specific range. For example, an analog telephone line accommodates a bandwidth of 3,000 hertz (Hz), the difference between the lowest (300 Hz) and highest (3,300 Hz) frequencies that it can carry.

See also: bits per second (bps)

**Bandwidth Allocation Protocol (BAP)**
A Point-to-Point Protocol (PPP) control protocol that is used on a multiprocessing connection to dynamically add and remove links.

See also: Point-to-Point Protocol (PPP)

**BAP**
See other term: Bandwidth Allocation Protocol (BAP)

**base priority**
A precedence ranking that determines the order in which the threads of a process are scheduled for the processor. Use Task Manager to view and change base priorities. For Message Queuing, a property that specifies the queue`s priority in a public queue. You can set the base priority from -32,768 to 32,767; the default priority is 0. Private queues do not support base priority. Message Queuing routes and delivers messages first by base priority, then by message priority.

See also: Message Queuing  private queue  public queue

**Basic authentication**
An authentication mechanism that is supported by most browsers, including Internet Explorer. Basic authentication encodes user name and password data before transmitting it over the network. Note that encoding is not the same as encryption. Also known as plaintext authentication.

See also: Anonymous authentication  authentication  Digest authentication  encryption  password  user name

**basic disk**
A physical disk that can be accessed by MS-DOS and all Windows-based operating systems. Basic disks can contain up to four primary partitions, or three primary partitions and an extended partition with multiple logical drives.

See also: dynamic disk  extended partition  logical drive  MS-DOS (Microsoft Disk Operating System)  primary partition

**basic input/output system (BIOS)**
The set of essential software routines that test computer hardware at startup, start the operating system, and support the transfer of information between hardware devices. The BIOS is stored in read-only memory (ROM) so that it can be run when the computer is turned on. Although critical to performance, the BIOS is usually invisible to computer users.

See also: Extensible Firmware Interface (EFI)  read-only memory (ROM)  x86

**basic storage**
A storage method in MS-DOS and Windows for primary partitions, extended partitions, and logical drives.

See also: dynamic storage  extended partition  logical drive

**basic volume**
A primary partition or logical drive that resides on a basic disk.

See also: basic disk  logical drive  primary partition

**batch program**
An ASCII (unformatted text) file that contains one or more operating system commands. A batch program`s file name has a .cmd or .bat extension. When you type the file name at the command prompt, or when the batch program is run from another program, its commands are processed sequentially. Also called batch files.

See also: American Standard Code for Information Interchange (ASCII)  logon script

**baud rate**
The speed at which a modem communicates. Baud rate refers to the number of times the condition of the line changes. This is equal to bits per second only if each signal corresponds to one bit of transmitted data. Modems must operate at the same baud rate in order to communicate with each other. If the baud rate of one modem is set higher than that of the other, the faster modem usually alters its baud rate to match that of the slower modem.

See also: bits per second (bps)  modem (modulator/demodulator)

**B-channel**
A single channel of an Integrated Services Digital Network (ISDN) line that is used to carry either voice or data information. ISDN Basic Rate Interface (BRI) has 2 B-channels. ISDN Primary Rate Interface (PRI) in North America has 23 B-channels. ISDN Primary Rate Interface (PRI) in Europe has 23 B-channels. Also called bearer channel.

See also: Integrated Services Digital Network (ISDN)

**BDC**
See other term: backup domain controller (BDC)

**Berkeley Internet Name Domain (BIND)**
An implementation of Domain Name System (DNS) written and ported to most available versions of the UNIX operating system. The Internet Software Consortium maintains the BIND software.

See also: BIND boot file  Domain Name System (DNS)

**BGP**
See other term: Border Gateway Protocol (BGP)

**binary**
A base-2 number system in which values are expressed as combinations of two digits, 0 and 1.

**BIND**
See other term: Berkeley Internet Name Domain (BIND)

**BIND boot file**
A configuration file used by DNS servers running under versions of the Berkeley Internet Name Domain (BIND) software implementation. The BIND boot file is a text file, Named.boot, where individual lines in the file list boot directives used to start a service when the DNS server is started. By default, Microsoft DNS servers use DNS service parameters stored in the registry, but they allow the use of a BIND boot file as an alternative for reading boot configuration settings.

See also: Berkeley Internet Name Domain (BIND)  DNS server  registry boot

**bindery**
A database in Novell NetWare 3.x that contains organizational and security information about users and groups.

**BIOS**
See other term: basic input/output system (BIOS)

**B-ISDN**
See other term: broadband integrated services digital network (B-ISDN)

**bit (binary digit)**
The smallest unit of information handled by a computer. One bit expresses a 1 or a 0 in a binary numeral, or a true or false logical condition. A group of 8 bits makes up a byte, which can represent many types of information, such as a letter of the alphabet, a decimal digit, or other character. Also known as a binary digit.

See also: binary

**bits per second (bps)**
The number of bits transmitted every second, used as a measure of the speed at which a device, such as a modem, can transfer data.

See also: modem (modulator/demodulator)

**blank password**
A password that does not contain any characters. Also called a null password.

See also: password policy  strong password  weak password

**block inheritance**
In Group Policy, the method of preventing a child container from inheriting the policy settings from all parent containers. Block inheritance does not apply to Group Policy object (GPO) links that are enforced.

See also: enforce  Group Policy Management console (GPMC)  Group Policy object link  inheritance  precedence

**boot**
The process of starting or resetting a computer. When first turned on (cold boot) or reset (warm boot), the computer runs the software that loads and starts the computer's operating system, which prepares it for use.

**boot files**
The system files needed to start Windows. The boot files include Ntldr and Ntdetect.com.

See also: partition boot sector

**Boot Information Negotiation Layer service (BINLSVC)**
A service that runs on a Remote Installation Services (RIS) server that acts on client boot requests. The display name of BINLSVC is Remote Installation.

See also: Remote Installation Services (RIS)

**Boot Logging**
A process in which a computer that is starting (booting) creates a log file that records the loading of each device and service. The log file is called Ntbtlog.txt, and it is saved in the systemroot directory.

See also: systemroot

**boot partition**
The partition that contains the Windows operating system and its support files. The boot partition can be, but does not have to be, the same as the system partition.

See also: partition  Primary disk  system partition

**boot volume**
The volume that contains the Windows operating system and its support files. The boot volume can be, but does not have to be, the same as the system volume.

See also: system volume  volume

**bootable CD**
A CD-ROM that can be used to start a computer. An automated installation uses a bootable CD to start a computer.

**BOOTP**
See other term: bootstrap protocol (BOOTP)

**BOOTP extensions**
A set of optional information types defined originally in RFC 1497 for use with BOOTP service and later supported by DHCP. In DHCP, these extensions form the legacy core set of client parameters available and supported by most standard DHCP and BOOTP servers.

See also: bootstrap protocol (BOOTP)  Dynamic Host Configuration Protocol (DHCP)  Request for Comments (RFC)

**bootstrap protocol (BOOTP)**
A protocol used primarily on TCP/IP networks to configure diskless workstations. RFCs 951 and 1542 define this protocol. DHCP is a later boot configuration protocol that uses this protocol. The Microsoft DHCP service provides limited support for BOOTP service.

See also: Dynamic Host Configuration Protocol (DHCP)  Request for Comments (RFC)  Transmission Control Protocol/Internet Protocol (TCP/IP)

**Border Gateway Protocol (BGP)**
A routing protocol designed for use between autonomous systems. BGP is especially useful for detecting routing loops.

See also: autonomous system (AS)  routing  routing protocol

**bps**
See other term: bits per second (bps)

**bridgehead server**
In Active Directory replication, a single server in a site that is designated to perform site-to-site replication for a specific domain and transport. In a messaging system, a server that receives and forwards e-mail traffic at each end of a connection agreement, similar to the task a gateway performs.

See also: Active Directory  Active Directory replication  gateway  replication  site

**broadband integrated services digital network (B-ISDN)**
An International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T) communication standard for high-speed networking that provides new services, including voice, video, and data on the same network.

See also: International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T)

**broadcast**
The transmission of packets by an Internet Protocol version 4 (IPv4) host to all computers on the subnet.

See also: broadcast address  Internet Protocol (IP)  packet  subnet

**broadcast address**
An Internet Protocol version 4 (IPv4) address that is destined for all hosts on a particular network segment.

See also: broadcast  host  Internet Protocol (IP)  IP address

**broadcast and unknown server (BUS)**
A multicast service on an emulated local area network (ELAN) that forwards broadcast, multicast, and initial unicast data traffic sent by a LAN emulation client (LEC).

See also: emulated local area network (ELAN)  LAN emulation client (LEC)  service

**broadcast message**
A network message sent from a single computer that is distributed to all other devices on the same segment of the network as the sending computer.

**broadcast network**
A network that supports more than two attached routers and has the ability to address a single physical message to all of the attached routers (broadcast). Ethernet is an example of a broadcast network.

See also: router

**browsing**
The process of creating and maintaining an up-to-date list of computers and resources on a network or part of a network by one or more designated computers running the Computer Browser service.

See also: Computer Browser service

**bucket brigade attack**
See other term: man-in-the-middle attack

**buffer**
An area of memory reserved for temporarily holding data before that data is used by a receiving device or application. Buffering protects against the interruption of data flow.

**built-in groups**
The default security groups installed with the operating system. Built-in groups have been granted useful collections of rights and built-in abilities. In most cases, built-in groups provide all the capabilities needed by a particular user. For example, members of the built-in Backup Operators group can back up and restore files and folders. To provide a needed set of capabilities to a user account, assign it to the appropriate built-in group.

See also: group

**bus**
A communication line used for data transfer among the components of a computer system. A bus essentially allows different parts of the system to share data. For example, a bus connects the disk-drive controller, memory, and input/output ports to the microprocessor.

See also: expansion slot  universal serial bus (USB)

# Windows Server 2003 Glossary - C

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - C**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**CA**
See other term: certification authority (CA)

**cache**
A special memory subsystem in which frequently used data values are duplicated for quick access.

See also: cache file  file system cache

**cache file**
A file used by DNS servers and clients to store responses to DNS queries. For Windows DNS servers, the cache file is named Cache.dns by default.

See also: authoritative  cache  DNS client  DNS server  Domain Name System (DNS)

**caching**
The process of temporarily storing recently accessed information in a special memory subsystem for quicker access.

See also: cache  caching resolver

**caching resolver**
A program that extracts information from DNS servers in response to client requests.

See also: cache  cache file  caching  DNS server

**caching-only server**
A DNS sever that does not host any DNS zones but that performs name resolution and stores the results in its cache.

See also: caching  DNS server  DNS zone  name resolution

**Callback Control Protocol (CBCP)**
The Network Control Protocol that negotiates the use of callback over PPP links.

See also: Point-to-Point Protocol (PPP)

**callback number**
The number that a remote access server uses to call back a user. This number can be preset by the administrator or specified by the user at the time of each call, depending on how the administrator configures the user`s callback options. The callback number should be the number of the phone line to which the user`s modem is connected.

See also: preset-to callback  remote access server  set-by-caller callback

**callback security**
A form of network security in which a remote access server calls a user back at a preset number after the user has made an initial connection and has been authenticated.

See also: preset-to callback  remote access server

**called subscriber ID (CSID) string**
A string that specifies the called subscriber ID transmitted by the receiving fax machine when receiving an inbound fax. This string is usually a combination of the fax or telephone number and the name of the business. It is often the same as the transmitter subscriber ID.

See also: string  transmitting station ID (TSID) string

**canonical (CNAME) resource record**

A resource record used to map an alternate alias name to a primary canonical DNS domain name used in the zone.

See also: resource record (RR)

**canonical name**
An object`s distinguished name presented with the root first and without the Lightweight Directory Access Protocol (LDAP) attribute tags (such as: CN=, DC=). The segments of the name are delimited with forward slashes (/). For example, CN=MyDocuments,OU=MyOU,DC=Microsoft,DC=Com is presented as microsoft.com/MyOU/MyDocuments in canonical form.

See also: distinguished name  Lightweight Directory Access Protocol (LDAP)

**capture**
To assign a port to a printer. Documents that you print are sent to the printer through the captured port. For Network Monitor, the process by which frames are copied.

See also: frame

**capture password**
A password that allows the user to capture statistics from the network and to display captured data.

See also: display password  password

**capture trigger**
A set of conditions that, when met, stop the capture or execute a program or command file to perform another type of action related to the capture.

**CardBus**
A 32-bit PC Card.

**cartridge font**
A font contained in a plug-in cartridge and used to add fonts to laser, ink-jet, or high-end dot-matrix printers. Cartridge fonts are distinguished both from internal fonts, which are contained in read-only memory (ROM) in the printer and are always available, and from downloadable (soft) fonts, which reside on disk and which can be sent to the printer as needed.

See also: downloadable fonts  font  font cartridge  read-only memory (ROM)

**cascading hubs**
A network configuration in which hubs are connected to other hubs.

See also: hub

**catalog**
For Indexing Service, a collection of all index information and stored properties for a particular group of file system directories. By default, Indexing Service indexes the System and Web catalogs on your hard drive.

See also: Indexing Service  property

**CBCP**
See other term: Callback Control Protocol (CBCP)

**CBR**
See other term: constant bit rate (CBR)

**CCP**
See other term: Compression Control Protocol (CCP)

**CDFS**
See other term: Compact Disc File System (CDFS)

**certificate**
A digital document that is commonly used for authentication and to help secure information on a network. A certificate binds a public key to an entity that holds the corresponding private key. Certificates are digitally signed by the certification authority that issues them, and they can be issued for a user, a computer, or a service.

**certificate revocation list (CRL)**
A document maintained and published by a certification authority that lists certificates that have been revoked.

See also: certification authority (CA)

**Certificate Services**
A software service that issues certificates for a particular certification authority (CA). It provides customizable services for issuing and managing certificates for the enterprise. Certificates can be used to provide authentication support, including secure e-mail, Web-based authentication, and smart-card authentication.

See also: authentication certification authority (CA) Internet Authentication Service (IAS) service

**certificate store**
A permanent storage where certificates, certificate revocation lists, and certificate trust lists are stored.

See also: certificate revocation list (CRL) certificate trust list (CTL)

**certificate template**
A Windows construct that specifies the format and content of certificates based on their intended usage. When requesting a certificate from a Windows enterprise certification authority (CA), certificate requestors can select from a variety of certificate types that are based on certificate templates.

See also: certification authority (CA)

**certificate trust list (CTL)**
A signed list of root certification authority certificates that an administrator considers reputable for designated purposes, such as client authentication or secure e-mail.

See also: certification authority (CA) root certificate

**certification authority (CA)**
An entity responsible for establishing and vouching for the authenticity of public keys belonging to subjects (usually users or computers) or other certification authorities. Activities of a certification authority can include binding public keys to distinguished names through signed certificates, managing certificate serial numbers, and certificate revocation.

See also: public key root certification authority

**certification hierarchy**
A model of trust for certificates in which certification paths are created by means of the establishment of parent-child relationships between certification authorities.

See also: certification authority (CA) certification path

**certification path**
An unbroken chain of trust, consisting of certificates from trusted certification authorities, from a specific certificate to the root certification authority in a certification hierarchy.

See also: certification authority (CA) certification hierarchy public key

**Challenge Handshake Authentication Protocol (CHAP)**
A challenge-response authentication protocol for Point-to-Point Protocol (PPP) connections described in RFC 1994. It uses the industry-standard MD5 hashing algorithm to hash the combination of a challenge string issued by the authenticating server and the user`s password in the response.

See also: MD5 message digest Point-to-Point Protocol (PPP)

**change journal**
A feature that tracks changes to NTFS volumes, including additions, deletions, and modifications. The change journal exists on the volume as a sparse file.

See also: NTFS file system sparse file volume

**channel**
A path or link through which noncontrol information passes between two devices. A single Basic Rate Interface (BRI) connection, for example, has one physical connection but two channels for exchanging information between devices. This is often called a bearer channel, implying a channel that carries information. On the Internet, a Web site designed to deliver content from the Internet to your computer, similar to subscribing to a favorite Web site.

See also: active content B-channel D-channel

**character mode**
A display mode in which the monitor can display letters, numbers, and other text characters, but no graphical images or character formatting (italics, superscript, and so on).

See also: graphics mode

**child domain**
For DNS and Active Directory, a domain located in the namespace tree directly beneath another domain (the parent domain). For example, example.microsoft.com would be a child domain of the parent domain microsoft.com. Also known as a subdomain.

See also: Active Directory domain Domain Name System (DNS) parent domain

**child object**
An object that resides in another object. A child object implies relation. For example, a file is a child object that resides in a folder, which is the parent object.

See also: object parent object

**Chooser**
The Macintosh desk accessory with which users select the network server and printers they want to use.

**Chooser Pack**
A collection of files, some of which contain PostScript information. When a Macintosh computer sends a print job to a PostScript printer, the printer uses a Chooser Pack to interpret PostScript commands in the print job.

See also: PostScript PostScript printer

**cipher**
The particular algorithm used to transform a readable message (called plaintext or cleartext) into an unreadable, scrambled, or hidden message (called ciphertext).

**circuit switching**
A technology in which a connection is established by closing a switch, for example, by dialing a telephone. This creates a circuit in which all the information transmitted over the connection follows the same route to the destination. The circuit is dedicated to the connection until it is closed.

**claim**
A statement that a server makes (for example, name, identity, key, group, privilege, or capability) about a client.

**claim mapping**
The act of mapping, removing or filtering, or passing claims between various claim sets.

**claims-aware application**
A Microsoft ASP .NET application that performs authorization based on the claims that are present in an Active Directory Federation Services (ADFS) security token.

**class**
A category of objects that share a common set of characteristics. Each object in the directory is an instance of one or more classes in the schema.

See also: object schema

**Class A IP address**
A unicast IP address that ranges from 1.0.0.1 through 127.255.255.254. The first octet indicates the network, and the last three octets indicate the host on the network. Class-based IP addressing has been superceded by Classless Interdomain Routing (CIDR).

See also: Class B IP address Class C IP address Classless Interdomain Routing (CIDR) IP address

**Class B IP address**
A unicast IP address that ranges from 128.0.0.1 through 191.255.255.254. The first two octets indicate the network, and the last two octets indicate the host on the network. Class-based IP addressing has been superceded by Classless Interdomain Routing (CIDR).

See also: Class A IP address Class C IP address Classless Interdomain Routing (CIDR) IP address

**Class C affinity**
Specifies that Network Load Balancing direct multiple requests from the same TCP/IP Class C address range to the same cluster host.

See also: affinity  Network Load Balancing  single affinity

**Class C IP address**
A unicast IP address that ranges from 192.0.0.1 to 223.255.255.254. The first three octets indicate the network, and the last octet indicates the host on the network. Network Load Balancing provides optional session support for Class C IP addresses (in addition to support for single IP addresses) to accommodate clients that make use of multiple proxy servers at the client site. Class-based IP addressing has been superceded by Classless Interdomain Routing (CIDR).

See also: Class A IP address  Class B IP address  Classless Interdomain Routing (CIDR)  IP address  Network Load Balancing  proxy server

**classical IP over ATM (CLIP)**
A proposed Internet standard, described in RFC 2225, that allows Internet Protocol (IP) communication directly on the asynchronous transfer mode (ATM) layer, bypassing an additional protocol (such as Ethernet or Token Ring) in the protocol stack.

See also: asynchronous transfer mode (ATM)  Internet Protocol (IP)

**Classless Interdomain Routing (CIDR)**
An IP address and routing management method that allocates IP addresses in a way that reduces the number of routes stored on any individual router, while also increasing the number of available IP addresses. CIDR replaces class-based IP address allocation.

See also: IP address  router

**clear**
To turn off an option by removing the X or check mark from a check box. You clear a check box by clicking it, or by selecting it and then pressing the SPACEBAR.

**cleartext**
See other term: plaintext

**cleartext passwords**
Passwords that are not scrambled, thus making them more susceptible to network sniffers.

See also: password  plaintext

**client**
Any computer or program connecting to, or requesting the services of, another computer or program.

See also: server

**client access**
The ability of client computers to connect to a server or resource, for example, a file on a server or a cluster resource.

See also: client  cluster  resource  server

**client account partner discovery Web page**
The Web page that is used to interact with a user to determine which account partner the user belongs to when Active Directory Federation Services (ADFS) cannot automatically determine which of the account partners should authenticate the user.

**client application**
A Windows-based application that can display and store linked or embedded objects. For distributed applications, the application that imitates a request to a server application.

See also: server application

**client authentication**
A method of authentication by which the client in a client-server communication proves its identity to the server.

See also: authentication  client  identity  server

**client failback**
The process in which a namespace client, after failing over to another target, fails back to the previous target.

**client failover**
The process in which a namespace client attempts to access subsequent targets in a referral if the first target is (or becomes) unavailable.

**client logoff Web page**
When Active Directory Federation Services (ADFS) performs a logoff operation, a Web page that is executed to provide visual feedback to the user that the logoff has occurred.

**client logon Web page**
When Active Directory Federation Services (ADFS) collects client credentials, a Web page that is executed to perform the user interaction. The client logon Web page may use any necessary business logic to determine the type of credentials to collect.

**client request**
A service request from a client computer to a server computer or a cluster of server computers.

See also: client  server

**Client Service for NetWare**
A service included that allows clients to make direct connections to resources on computers running NetWare 2.x, 3.x, 4.x, or 5.x server software by using the Internetwork Packet Exchange (IPX) protocol only. This service is included with Windows XP Professional and the Windows Server 2003 family.

See also: Internetwork Packet Exchange (IPX)  service

**CLIP**
See other term: classical IP over ATM (CLIP)

**ClipBook Server**
A system service that supports ClipBook Viewer, which allows pages to be seen by remote ClipBooks.

**cluster**
In data storage, the smallest amount of disk space that can be allocated to hold a file. All file systems used by Windows organize hard disks based on clusters, which consist of one or more contiguous sectors. The smaller the cluster size, the more efficiently a disk stores information. If no cluster size is specified during formatting, Windows picks defaults based on the size of the volume. These defaults are selected to reduce the amount of space that is lost and the amount of fragmentation on the volume. Also called an allocation unit. In computer networking, a group of independent computers that work together to provide a common set of services and present a single-system image to clients. The use of a cluster enhances the availability of the services and the scalability and manageability of the operating system that provides the services.

See also: availability  client  file system  scalability  volume

**cluster adapter**
The adapter that, when using multiple network adapters in each host of a Network Load Balancing cluster, handles the network traffic for cluster operations (the traffic for all hosts in the cluster). This adapter is programmed with the host`s cluster IP address.

See also: cluster  dedicated adapter  host  IP address  Network Load Balancing

**Cluster Administrator**
An application that configures a cluster and its nodes, groups, and resources. Cluster Administrator can run on any member of the trusted domain regardless of whether the computer is a cluster node.

See also: cluster  Cluster.exe  node  server cluster

**Cluster Administrator extension**
A software component that implements the Cluster Administrator extension application programming interface (API) for allowing Cluster Administrator to configure a new resource type.

See also: application programming interface (API)  cluster  Cluster Administrator

**Cluster API**
A collection of functions that are implemented by the cluster software and used by a cluster-aware client or server application, a cluster management application, or a Resource DLL. The Cluster API is used to manage the cluster, cluster objects, and the

cluster database.

See also: cluster  Resource DLL  server cluster

**cluster full Internet name**
A full Internet name for the Network Load Balancing cluster (for example, cluster.microsoft.com). This name is used for the cluster as a whole and should be the same for all hosts in the cluster.

See also: affinity  Network Load Balancing  single affinity

**cluster IP address**
The Network Load Balancing cluster`s IP address in standard Internet dotted notation (for example, w.x.y.z). The address is a virtual IP address used to address the cluster as a whole, and it should be the IP address that maps to the full Internet name that you specify for the cluster. In a Network Load Balancing cluster, this address must be set identically for all hosts in the cluster.

See also: Network Load Balancing  Network Load Balancing cluster

**cluster log**
An optionally enabled trace record of Cluster service events on a node. Not synonymous with quorum log.

See also: cluster  node  quorum log

**cluster network address**
The network (media access control) address for the network adapter that is to be used for handling client-to-cluster traffic in a Network Load Balancing cluster.

See also: Network Load Balancing  Network Load Balancing cluster

**Cluster service**
The essential software component that controls all aspects of server cluster operation and manages the cluster database. Each node in a server cluster runs one instance of the Cluster service.

See also: cluster  node  server cluster

**cluster storage**
Storage where one or more attached disks hold data used either by server applications running on the cluster or by applications for managing the cluster. Each disk on the cluster storage is owned by only one node of the cluster. The ownership of disks moves from one node to another when the disk group fails over or moves to the other node.

See also: cluster  failover  node

**Cluster.exe**
An alternative to using Cluster Administrator to administer clusters from the command prompt. You can also call Cluster.exe from command scripts to automate many cluster administration tasks.

See also: cluster  Cluster Administrator

**cluster-aware application**
An application that can run on a cluster node and that can be managed as a cluster resource. Cluster-aware applications use the Cluster API to receive status and notification information from the server cluster.

See also: cluster  Cluster API  cluster-unaware application  server cluster

**cluster-unaware application**
An application that can run on a cluster node and be managed as a cluster resource but that does not support the Cluster API.

See also: cluster  Cluster API  cluster-aware application  node

**CMYK color space**
Multidimensional color space consisting of the cyan, magenta, yellow, and black intensities that make up a given color. Commercial color printing devices generally use this system of four-color process inks.

See also: color management  color space

**cn (Common-Name)**
The descriptive relative distinguished name for the schema object.

See also: distinguished name  schema

**code page**
A means of providing support for character sets and keyboard layouts for different countries or regions. A code page is a table that relates the binary character codes used by a program to keys on the keyboard or to characters on the display.

**codec**
Hardware that can convert audio or video signals between analog and digital forms (coder/decoder), hardware or software that can compress and uncompress audio or video data (compression/decompression), or the combination of coder/decoder and compression/decompression. Generally, a codec compresses uncompressed digital data so that the data uses less memory.

**color gamut**
The particular range of colors that a device is able to produce. A device such as a scanner, monitor, or printer can produce a unique range of colors, which is determined by the characteristics of the device itself.

See also: color profile  rendering intent

**color management**
Process of producing accurate, consistent color among a variety of input and output devices. A color management system (CMS) maps colors between devices such as scanners, monitors, and printers; transforms colors from one color space to another (for example, RGB to CMYK); and provides accurate on-screen or print previews.

See also: CMYK color space  RGB color space

**color profile**
A profile that contains the data needed for translating the values of a color gamut. This data includes information about color, hue, saturation, and brightness.

See also: color gamut  hue  saturation

**color space**
A set of three values that defines how a color can be represented on computer devices such as monitors, scanners, and printers. For example, in the LAB color space, the terms luminance or whiteness (L), redness-greenness (A), and yellowness-blueness (B) are used; in the HVC system, the terms are hue (H), value (V), and chroma (C). Color space refers to the three-dimensional space that is defined by the respective values, such as L, A, and B.

See also: CMYK color space  color space  RGB color space

**COM**
See other term: Component Object Model (COM)

**COM Internet Services Proxy**
Software that supports distributed applications that use Hypertext Transfer Protocol (HTTP) to communicate through Internet Information Services (IIS).

See also: Hypertext Transfer Protocol (HTTP)  Internet Information Services (IIS)

**Command Prompt window**
A window displayed on the desktop used to interface with the MS-DOS operating system. MS-DOS commands are typed at an entry point identified by a blinking cursor.

See also: MS-DOS (Microsoft Disk Operating System)

**common groups**
Groups that appear in the program list on the Start menu for all users who log on to the computer. Only administrators can create or change common groups.

See also: group

**communication port**
A port on a computer that allows asynchronous communication of one byte at a time. Also called a serial port.

See also: asynchronous communication  serial port

**communication settings**
Operating parameters, such as bits per second (bps) and modem type, that apply to serial ports on a computer.

See also: bits per second (bps)  modem (modulator/demodulator)  serial port

**community name**
A name that identifies groups of Simple Network Management Protocol (SNMP) hosts. This name is placed in SNMP messages that are sent between SNMP-managed devices such as servers running Windows Server 2003 and SNMP management stations. Typically, all hosts belong to Public, which is the standard name for a common community of all SNMP hosts.

See also: host  Simple Network Management Protocol (SNMP)  trap

**Compact Disc File System (CDFS)**
A 32-bit protected-mode file system that controls access to the contents of CD-ROM drives.

**Component Object Model (COM)**
An object-based programming model designed to promote software interoperability; it allows two or more applications or components to easily cooperate with one another, even if they were written by different vendors, at different times, in different programming languages, or if they are running on different computers running different operating systems. OLE technology and ActiveX are both built on top of COM.

See also: ActiveX  OLE

**Compression Control Protocol (CCP)**
A protocol used in the negotiation process in a Point-to-Point Protocol (PPP) connection. CCP is one type of Network Control Protocol. Network Control Protocols are used to establish and configure different network protocol parameters for Internet Protocol (IP), Internetwork Packet Exchange (IPX), and NetBIOS Extended User Interface (NetBEUI).

See also: Internet Protocol (IP)  Internetwork Packet Exchange (IPX)  NetBIOS Extended User Interface (NetBEUI)  Network Control Protocol  Point-to-Point Protocol (PPP)

**computer account**
An account that is stored in Active Directory and that uniquely identifies a computer in a domain. A computer account uses the same name as the computer joining the domain.

See also: Active Directory  domain

**Computer Browser service**
A service that maintains an up-to-date list of computers that share resources on your network and that supplies the list to programs that request it. The Computer Browser service is used to view a list of available network resources.

See also: resource  service

**Computer Management**
A component you can use to view and control many aspects of the computer configuration. Computer Management combines several administration utilities into a single console tree, providing easy access to a local or remote computer`s administrative properties and tools.

See also: console tree

**computer name**
A unique name of up to 15 uppercase characters that identifies a computer to the network. The name cannot be the same as any other computer or domain name in the network.

See also: domain name

**computer quota**
For Message Queuing, the storage size limit for messages on a computer, based on the total size of the messages. When a computer quota is reached, Message Queuing can no longer send messages to that computer until one or more messages are removed from queues. Message Queuing enforces the computer quota before it enforces the queue quota on a computer.

See also: Message Queuing  queue quota

**computer-telephony integration (CTI)**
The practice of using a computer to control one or more telephone and communications functions.

**configuration set**
For Active Directory Application Mode (ADAM), a set of ADAM instances that share and replicate to a common schema partition and a common configuration partition.

**conflict detection**
For DHCP, an optional server-side mechanism for detecting if a scope IP address is in use on the network. When enabled, the DHCP server will ping an address first before offering that address to clients, and then it will briefly await a response. If the pinged address responds, a conflict is registered and that address is not offered to clients obtaining a lease from the server.

See also: DHCP server  Dynamic Host Configuration Protocol (DHCP)  lease  ping  scope

**connect**
To assign a drive letter, port, or computer name to a shared resource so that you can use it.

See also: shared resource

**connected user**
A user who has access to a computer or a resource across the network.

See also: resource

**connected, authenticating user**
A user's status when a telephone connection has been established but authentication has not yet taken place. The user may be trying to prove security clearance, or the system may be idle. If this condition occurs, followed by the Waiting for Call phase, then the user was unable to provide a correct user name or password. If this phase is repeated, followed by the Waiting for Call phase, an unauthorized attempt to access the network may be under way.

See also: authentication  Waiting for Call

**connection**
For DFS Replication, a one-way replication path between two members of a replication group.

**connection object**
An Active Directory object that represents a replication connection from one domain controller to another. The connection object is a child of the replication destinations NTDS Settings object and identifies the replication source server, contains a replication schedule, and specifies a replication transport. Connection objects are created automatically by the Knowledge Consistency Checker (KCC), but they can also be created manually. Automatically generated connections must not be modified by the user unless they are first converted into manual connections.

See also: Active Directory  Knowledge Consistency Checker (KCC)  replication

**connection request policy**
A set of conditions and profile settings that network administrators use to specify how Internet Authentication Service (IAS) servers handle incoming authentication and accounting request messages.

See also: authentication  Internet Authentication Service (IAS)  remote access policy

**connectionless**
A network protocol in which a sender broadcasts traffic on the network to an intended receiver without first establishing a connection to the receiver.

See also: protocol

**connection-oriented**
A type of network protocol that requires an end-to-end virtual connection between the sender and receiver before communicating across the network.

See also: protocol

**connector application**
For Message Queuing, an application that enables Message Queuing computers to communicate with computers that use other messaging systems.

See also: connector queue  foreign computer  Message Queuing

**connector queue**
For Message Queuing, a queue created on servers running a connector application. You can use the connector application to exchange messages with computers that are running other message-queuing products.

See also: connector application  foreign computer  Message Queuing

**console**
A framework for hosting administrative tools, such as Microsoft Management Console (MMC). A console is defined by the items in its console tree, which might include folders or other containers, World Wide Web pages, and other administrative items. A console has windows that can provide views of the console tree and the administrative properties, services, and events that are acted on by the items in the console tree.

See also: console tree  details pane  Microsoft Management Console (MMC)

**console tree**
The left pane in Microsoft Management Console (MMC) that displays the items contained in the console. The items in the console tree and their hierarchical organization determine the capabilities of a console.

See also: details pane  Microsoft Management Console (MMC)

**constant bit rate (CBR)**
An asynchronous transfer mode (ATM) service type that supports constant bandwidth allocation. This service type is used for voice and video transmissions that require little or no cell loss and rigorous timing controls during transmission.

See also: asynchronous transfer mode (ATM)  bandwidth

**constrained delegation**
The ability to specify that a service or computer account can perform Kerberos delegation to a limited set of services.

See also: delegation  Kerberos V5 authentication protocol  service

**container object**
An object that can logically contain other objects. For example, a folder is a container object.

See also: Active Directory  noncontainer object  object

**convergence**
The process of stabilizing a system after changes occur in the network. For dynamic routing, if a route becomes unavailable, routers send update messages throughout the network, reestablishing information about preferred routes. For Network Load Balancing, a process by which hosts exchange messages to determine a new, consistent state of the cluster and to elect the default host. During convergence, a new load distribution is determined for hosts that share the handling of network traffic for specific Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports.

See also: cluster  default host  host  Network Load Balancing  routing  User Datagram Protocol (UDP)

**copy backup**
A backup that copies all selected files but does not mark each file as having been backed up (in other words, the archive attribute is not cleared). Copying is useful if you want to back up files between normal and incremental backups because copying does not affect these other backup operations.

See also: daily backup  differential backup  incremental backup  normal backup

**copy set**
A duplicate of the media master that Remote Storage copies to a different tape or disk. Copy sets are typically used for backup purposes.

See also: media master  Remote Storage

**count to infinity**
A loop that happens when a link in a network goes down and routers on the network update their routing tables with incorrect hop counts. For example, a loop can develop if the link to Router C goes down. Router B then advertises that the link is down and that it has no route to C. Because Router A has a route to C with a metric of 2, it responds to Router B and sends its link to C. Router B then updates its table to include a link with metric 3, and the routers continue to announce and update their links to C until they reach the number 16. This is a count to infinity.

See also: router

**CPU Time**
In Task Manager, the total processor time, in seconds, used by a process since it started.

See also: Task Manager

**CPU Usage**

In Task Manager, the percentage of time that a process used the CPU since the last update. This percentage is displayed in Task Manager on the Processes tab, under the CPU column heading.

See also: Task Manager

**crash consistency**
A feature of shadow copy backups that ensures all files are backed up, regardless of their state.

See also: volume shadow copy

**Crash Dump Type**
Specifies the file format Dr. Watson will use to store the information. The Full format contains the entire memory space of the program, as well as the program image itself, the handle table, and other information that will be useful to the debugger. The Mini format may include the full memory and handle table, or it may simply contain information about a single thread. The Windows NT 4.0-compatible Full format provides you with the opportunity to use older tools to analyze the dump file. Crash Dump Type is only available when you have selected the Create Crash Dump File check box.

**CRC errors**
Errors caused by the failure of a cyclic redundancy check. A CRC error indicates that one or more characters in the data packet received were found garbled on arrival.

**credentials**
Information that includes identification and proof of identification that is used to gain access to local and network resources. Examples of credentials are user names and passwords, smart cards, and certificates.

See also: password  smart card

**CRL**
See other term: certificate revocation list (CRL)

**CRL distribution point**
An optional extension in an X.509 v3 certificate that identifies how information is obtained. Also, a directory entry or other distribution source for certificate revocation lists.

See also: certificate revocation list (CRL)  X.509 v3 certificate

**cross-reference object**
An object in which Active Directory stores information about directory partitions and external directory services. An example of an external directory service is another LDAP-compliant directory.

See also: Active Directory  directory partition  Lightweight Directory Access Protocol (LDAP)

**CryptoAPI**
An application programming interface (API) that is provided as part of Microsoft Windows. CryptoAPI provides a set of functions that allows applications to encrypt or digitally sign data in a flexible manner while providing protection for the user`s sensitive private key data. Actual cryptographic operations are performed by independent modules known as cryptographic service providers (CSPs).

See also: application programming interface (API)  cryptographic service provider (CSP)  private key

**cryptographic checksum**
A mathematical value that is a result of a calculation involving a cryptographic algorithm and a cryptographic key. Also known as an integrity check value (ICV) or a message integrity code (MIC).

See also: cryptography

**cryptographic service provider (CSP)**
The code that performs authentication, encoding, and encryption services that Windows-based applications access through CryptoAPI. A CSP is responsible for creating keys, destroying them, and using them to perform a variety of cryptographic operations. Each CSP provides a different implementation of the CryptoAPI. Some provide stronger cryptographic algorithms, while others use hardware components, such as smart cards.

See also: authentication  CryptoAPI  encryption  service  smart card

**cryptography**
The study or analysis of codes and encoding methods used to secure information. Cryptography is used to provide confidentiality, data integrity, authentication (entity and data origin), and nonrepudiation.

See also: authentication  nonrepudiation

**CSP**
See other term: cryptographic service provider (CSP)

**CTI**
See other term: computer-telephony integration (CTI)

**CTL**
See other term: certificate trust list (CTL)

**custom file type**
Typically, files with extensions that have been created for special kinds of files. Custom file types are not tracked by the system registry.

See also: registry

# Windows Server 2003 Glossary - D

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - D**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**DACL**

See other term: discretionary access control list (DACL)

**daily backup**

A backup that copies all selected files that have been modified the day the daily backup is performed. The backed-up files are not marked as having been backed up (in other words, the archive attribute is not cleared).

See also: copy backup  differential backup  incremental backup  normal backup

**Data Communications Equipment (DCE)**

One of two types of hardware that are connected by an RS-232-C serial connection, the other being a Data Terminal Equipment (DTE) device. A DCE is an intermediary device that often transforms input from a DTE before sending it to a recipient. A modem, for example, is a DCE that modulates data from a microcomputer (DTE) and sends it along a telephone connection.

See also: Data Terminal Equipment (DTE)  RS-232-C standard

**Data Encryption Standard (DES)**

An encryption algorithm that uses a 56-bit key and maps a 64-bit input block to a 64-bit output block. The key appears to be a 64-bit key, but one bit in each of the eight bytes is used for odd parity, resulting in 56 bits of usable key.

See also: key

**data fork**

One of two forks (resource fork and data fork) that make up each Macintosh file. The data fork holds most of the file's information and is shared between Macintosh and personal computer clients.

See also: resource fork

**data integrity**

A property of secure communications that allows a computer to verify that data has not been modified in transit from its source. For example, Internet Protocol security (IPSec)-protected data contains a cryptographic checksum that incorporates a secret key that is known only to the IPSec peers that are communicating. An intermediate node can modify the data, but without knowledge of the secret key, the node cannot recompute a correct cryptographic checksum.

See also: cryptographic checksum  Internet Protocol security (IPsec)

**Data Link Control (DLC)**

An address that uniquely identifies a node on a network. Every network adapter has a DLC address or DLC identifier (DLCI). Some network protocols, such as Ethernet and Token Ring, use DLC addresses exclusively. Other protocols, such as TCP/IP, use a logical address at the OSI Network layer to identify nodes. However, all network addresses must eventually be translated to DLC addresses. In TCP/IP networks, this translation is performed by the Address Resolution Protocol (ARP).

See also: Address Resolution Protocol (ARP)  node  Open Systems Interconnection (OSI) reference model

**data origin authentication**

A property of secure communications that allows a computer to verify that data originated from a valid source. For example, IPSec-protected data contains a cryptographic checksum that incorporates a secret key that is known only to the IPSec peers that are communicating. An invalid source can send data, but without knowledge of the secret key, the source cannot compute a correct cryptographic checksum.

See also: cryptographic checksum  cryptography  encryption  Internet Protocol security (IPsec)  key

**Data Terminal Equipment (DTE)**
In the RS-232-C hardware standard, any device, such as a remote access server or client, that has the ability to transmit information in digital form over a cable or a communications line.

See also: Data Communications Equipment (DCE)  remote access server  RS-232-C standard

**datagram**
One packet, or unit, of information that includes relevant delivery information, such as the destination address, that is sent through a packet-switching network.

See also: packet

**data-overrun error**
A state in which the sending computer is transmitting characters faster than the receiving computer can accommodate them. If this problem persists, reduce the bits-per-second (bps) rate.

See also: bits per second (bps)

**Dcdiag**
A Windows Support Tool that analyzes the state of domain controllers in a forest or enterprise and provides detailed information about how to identify abnormal behavior in a system. Domain controllers are identified and tested according to directives entered by the user at the command line. Also known as Domain Controller Diagnostic Tool.

See also: Command Prompt window  domain controller

**DCE**
See other term: Data Communications Equipment (DCE)

**D-channel**
A separate channel of an Integrated Services Digital Network (ISDN) line that is used for ISDN signaling. For ISDN Basic Rate Interface (BRI), the D-channel is 16 kilobits per second (Kbps). For ISDN Primary Rate Interface (PRI), the D-channel is 64 Kbps. Also called data channel.

See also: Integrated Services Digital Network (ISDN)  switch type

**DCOM**
See other term: Distributed Component Object Model (DCOM)

**DDE**
See other term: dynamic data exchange (DDE)

**dead-letter queue**
For Message Queuing, a queue that stores nontransactional messages that are undeliverable or expired. These queues store failed messages on the computer on which the message expired. Messages in these queues are written to disk and are therefore recoverable.

See also: Message Queuing  queue  transactional dead-letter queue  transactional message

**debugger**
A program designed to aid in detecting, locating, and correcting errors in another program by allowing the programmer to step through the program, examine the data, and monitor conditions such as the values of variables.

See also: variable

**dedicated adapter**
The network adapter that, when using multiple network adapters in each host of a Network Load Balancing cluster, handles network traffic not related to cluster operations (the traffic for individual hosts on the network). This adapter is programmed with the host`s dedicated IP address.

See also: cluster adapter  IP address  Network Load Balancing

**dedicated connection**
A communications channel that connects two or more geographic locations. Dedicated connections are private or leased lines, rather than public lines.

See also: channel

**dedicated IP address**
The IP address of a Network Load Balancing host used for network traffic that is not associated with the Network Load Balancing cluster (for example, Telnet access to a specific host within the cluster). This IP address is used to individually address each host in the cluster and therefore should be unique for each host.

See also: host  IP address  Network Load Balancing  Network Load Balancing cluster

**default button**
In some dialog boxes, the command button that is selected or highlighted when the dialog box is initially displayed. The default button has a bold border, indicating that it will be chosen automatically if you press ENTER. You can override a default button by clicking Cancel or another command button.

**default host**
The host with the highest host priority for which a drainstop command is not in progress. After convergence, the default host handles all of the network traffic for Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports that are not otherwise covered by port rules.

See also: convergence  drainstop  host  host priority  port rule  User Datagram Protocol (UDP)

**default network**
In the Macintosh environment, the physical network on which the processes of a server reside as nodes and on which the server appears to users. The default network of the server must be one to which that server is attached. Only servers on AppleTalk Phase 2 internets have default networks.

See also: AppleTalk Phase 2  Internet  node

**default printer**
The printer to which a computer sends documents if you select the Print command without first specifying which printer you want to use with a program. You can have only one default printer; it should be the printer you use most often.

See also: printer

**default security level**
For software restriction policies, the default setting that defines whether software in a Group Policy object (GPO) is allowed to run. Exceptions can be made to the default security level by means of software restriction policies rules.

See also: Group Policy object (GPO)  software restriction policies  software restriction policies rule

**default tree and context**
The position of a User object in the file system of a Novell Directory Services (NDS) tree on a Novell NetWare network. The default tree and context is the default logon location in an NDS tree for a Microsoft Windows computer that is running Client Service for NetWare.

See also: Client Service for NetWare  Novell Directory Services (NDS)  object

**default user**
The profile that serves as a basis for all user profiles. Every user profile begins as a copy of the default user profile.

See also: user profile

**default zone**
The zone to which all Macintosh clients on the network are assigned by default.

See also: zone

**defragmentation**
The process of rewriting parts of a file to contiguous sectors on a hard disk to increase the speed of access and retrieval. In Active Directory, defragmentation rearranges how the data is written in the directory database file to compact it.

See also: Active Directory  fragmentation

**delegated authentication**
A method of authentication by which a trusted server (or service) is enabled to act as the client when it connects to third-tier application resources.

See also: authentication  client  server  service

**delegation**
An assignment of administrative responsibility to a user, computer, group, or organization. For Active Directory, an assignment of responsibility that allows users without administrative credentials to complete specific administrative tasks or to manage specific directory objects. Responsibility is assigned through membership in a security group, the Delegation of Control Wizard, or Group Policy settings. For DNS, an assignment of responsibility for a DNS zone. Delegation occurs when a name server (NS) resource record in a parent zone lists the DNS server that is authoritative for a child zone.

See also: Active Directory  administrative credentials  DNS server  Domain Name System (DNS)  Group Policy  security group  zone

**delta CRL**
A document that lists certificates that have been revoked after the last full certificate revocation list (CRL) has been published. Delta CRLs are maintained and published by a certification authority (CA), and they are normally much smaller than a full CRL.

See also: certificate revocation list (CRL)  certification authority (CA)

**demand-dial connection**
A connection, typically using a circuit-switched wide area network link, that is initiated when data needs to be forwarded. The demand-dial connection is typically terminated when there is no traffic.

See also: wide area network (WAN)

**demand-dial routing**
Routing that makes dial-up connections to connect networks based on need. For example, a branch office with a modem that dials and establishes a connection only when there is network traffic from one office to another.

See also: dial-up connection  modem (modulator/demodulator)

**demilitarized zone (DMZ)**
See other term: perimeter network

**denial-of-service attack**
A situation in which a person exploits a weakness or a design limitation of a network service to overload or halt the service, so that the service is not available for use. This type of attack is typically started to prevent other users from using a network service such as a web server or a file server.

See also: service  Web server

**dependency**
A relationship of reliance between two resources that makes it necessary for them to run in the same group on the same node. For example, an application is dependent on the disks that contain its data resources.

See also: node  resource

**dependency tree**
A diagram for visualizing the dependency relationships between resources.

See also: dependency  resource

**dependent client**
For Message Queuing, a computer that requires synchronous access to a Message Queuing server to perform all standard message queuing operations, such as sending and receiving messages and creating queues.

See also: independent client  Message Queuing server

**DES**
See other term: Data Encryption Standard (DES)

**descendent key**
All the subkeys that appear when a key in the registry is expanded. A descendent key is the same as a subkey.

See also: key  subkey

**desktop**
The on-screen work area on which windows, icons, menus, and dialog boxes appear.

**desktop pattern**

A design that appears across your desktop. You can create your own pattern or select a pattern provided by Windows.

See also: desktop

**destination document**
The document into which a package or a linked or embedded object is being inserted. For an embedded object, this is sometimes also called the container document.

See also: embedded object

**details pane**
The right pane in Microsoft Management Console (MMC) that displays details for the selected item in the console tree. The details can be a list of items or they can be administrative properties, services, and events that are acted on by a snap-in.

See also: console tree  Microsoft Management Console (MMC)  service  snap-in

**device**
Any piece of equipment that can be attached to a network or computer, for example, a computer, printer, joystick, adapter, or modem card, or any other peripheral equipment. Devices normally require a device driver to function with Windows. For Windows licensing, devices can be electronic equipment such as computers, workstations, terminals, and handheld computers that can access or use the services of Windows operating systems, including file sharing, print sharing, remote access, and authentication.

See also: device driver

**device conflict**
A conflict that occurs when the same system resources have been allocated to two or more devices. System resources include interrupt request (IRQ) lines, direct memory access (DMA) channels, input/output (I/O) ports, and memory addresses.

See also: direct memory access (DMA)  input/output (I/O) port  interrupt request (IRQ) lines  memory address  resource

**device driver**
Software that allows Windows to communicate with a specific hardware device. Before Windows can use any new hardware, a device driver must be installed.

See also: device  signed driver  unsigned driver  Windows Catalog

**device fonts**
See other term: printer fonts

**Device Manager**
An administrative tool that you can use to manage the devices on your computer. Using Device Manager, you can view and change device properties, update device drivers, configure device settings, and uninstall devices.

See also: device  uninstall

**DFS**
See other term: Distributed File System (DFS)

**DFS link**
An element in the Distributed File System (DFS) namespace that lies below the root and maps to one or more targets, each of which corresponds to a shared folder or another DFS root.

See also: DFS root  Distributed File System (DFS)  domain DFS

**DFS namespace**
A namespace consisting of a root and many links and targets. The namespace starts with a root that maps to one or more root targets. Below the root are links that map to their own targets. The Distributed File System (DFS) namespace provides the user with a logical view of distributed network shares.

See also: Distributed File System (DFS)  root

**DFS path**
The combination of a Distributed File System (DFS) root and a DFS link. An example of a DFS path is \\server\dfs\a\b\c\link, where \\server\dfs is the DFS root, and \a\b\c\ is the DFS link.

See also: DFS link  DFS root  Distributed File System (DFS)

**DFS root**
The starting point of the Distributed File System (DFS) namespace. The root is often used to refer to the namespace as a whole. A root maps to one or more root targets, each of which corresponds to a shared folder on a server.

See also: DFS link  DFS namespace  Distributed File System (DFS)

**DFS topology**
The overall logical hierarchy of the Distributed File System (DFS), including elements such as roots, links, shared folders, and replica sets, as depicted in the DFS administrative console. This is not to be confused with the DFS namespace, which is the logical view of shared resources seen by users.

See also: DFS link  DFS namespace  DFS root  Distributed File System (DFS)  domain DFS

**DHCP**
See other term: Dynamic Host Configuration Protocol (DHCP)

**DHCP acknowledgment message (DHCPACK)**
A message sent by the DHCP server to a client to acknowledge and complete a client`s request for leased configuration. This message will contain a committed IP address for the client to use for a stated period of time along with other optional client parameters. The DHCP acknowledgment message name is DHCPACK.

See also: DHCP server  Dynamic Host Configuration Protocol (DHCP)

**DHCP class identifier**
A special reserved option type used by DHCP clients to optionally identify membership in a specific DHCP option class, either a vendor or user class. For vendor class identification, hardware vendors can choose to predefine specific identifier values; for example, to identify a client`s hardware configuration. For user class identification, values can be defined administratively to identify a logical group of DHCP clients, such as all clients in a particular building and floor location.

See also: DHCP client  DHCP option  Dynamic Host Configuration Protocol (DHCP)

**DHCP client**
Any network-enabled device that supports the ability to communicate with a DHCP server for the purpose of obtaining dynamic leased IP configuration and related optional parameters information.

See also: DHCP server  Dynamic Host Configuration Protocol (DHCP)  lease

**DHCP client alternate configuration**
An alternate static configuration option for TCP/IP network connections that provides simplified computer migration between networks.

See also: Transmission Control Protocol/Internet Protocol (TCP/IP)

**DHCP decline message (DHCPDECLINE)**
A message sent by a DHCP client to the DHCP server to decline the offer of an IP address on the network. This message is used when the client detects a potential conflict because the IP address is found to be already in use on the network. The DHCP decline message name is DHCPDECLINE.

See also: DHCP client  DHCP server  Dynamic Host Configuration Protocol (DHCP)

**DHCP information message (DHCPINFORM)**
A reserved DHCP message type used by computers on the network to request and obtain information from a DHCP server for use in their local configuration. When this message type is used, the sender is already externally configured for its IP address on the network, which may or may not have been obtained using DHCP. The DHCP information message name is DHCPINFORM.

See also: DHCP server  Dynamic Host Configuration Protocol (DHCP)

**DHCP negative acknowledgment message (DHCPNAK)**
A message sent by a DHCP server to a client to indicate that the IP address that the client requested is not correct for the local IP network served by the DHCP server. This message is most often used when the client computer was moved to a new location, but it could also indicate that the client`s lease with the server has expired. The DHCP negative acknowledgment message name is DHCPNAK.

See also: DHCP server  Dynamic Host Configuration Protocol (DHCP)  lease

**DHCP offer message (DHCPOFFER)**

A message used by DHCP servers to offer the lease of an IP address to a DHCP client when it starts on the network. When this message is used, a client can receive more than one offer if multiple DHCP servers are contacted during the DHCP discovery phase, but the client will typically select the first address it is offered. The DHCP offer message name is DHCPOFFER.

See also: DHCP client  DHCP server  Dynamic Host Configuration Protocol (DHCP)

**DHCP option**
Address configuration parameters that a DHCP service assigns to clients. Most DHCP options are predefined, based on optional parameters defined in Request for Comments (RFC) 1542, although extended options can be added by vendors or users.

See also: DHCP service  Dynamic Host Configuration Protocol (DHCP)

**DHCP release message (DHCPRELEASE)**
A message sent by clients to the DHCP server to indicate release of its leased IP address. The client uses this message to cancel its currently active lease. You can perform address release manually using the ipconfig /release command at a command prompt. The DHCP release message name is DHCPRELEASE.

See also: DHCP server  Dynamic Host Configuration Protocol (DHCP)  lease

**DHCP request message (DHCPREQUEST)**
A message sent by clients to the DHCP server to request or renew lease of its IP address. The client uses this message to select and request a lease from a specific DHCP server, to confirm a previously leased IP address after the client system is restarted, or to extend the current IP address lease for the client. The DHCP request message name is DHCPREQUEST.

See also: DHCP server  Dynamic Host Configuration Protocol (DHCP)  lease

**DHCP server**
A computer running the Microsoft DHCP service that offers dynamic configuration of IP addresses and related information to DHCP-enabled clients.

See also: DHCP service  Dynamic Host Configuration Protocol (DHCP)  IP address

**DHCP service**
A service that enables a computer to function as a DHCP server and configure DHCP-enabled clients on a network. DHCP runs on a server, enabling the automatic, centralized management of IP addresses and other TCP/IP configuration settings for network clients.

See also: Dynamic Host Configuration Protocol (DHCP)  IP address  service

**DHCP service resource**
A resource type that provides DHCP services from a cluster.

See also: cluster  Dynamic Host Configuration Protocol (DHCP)

**DHCP/BOOTP Relay Agent**
The agent program or component responsible for relaying Dynamic Host Configuration Protocol (DHCP) and bootstrap protocol (BOOTP) broadcast messages between a DHCP server and a client across an Internet Protocol (IP) router. A DHCP relay agent supports DHCP/BOOTP message relay as defined in RFCs 1541 and 2131. The DHCP Relay Agent routing protocol component is managed using the Routing and Remote Access snap-in.

See also: bootstrap protocol (BOOTP)  DHCP server  Dynamic Host Configuration Protocol (DHCP)  Request for Comments (RFC)

**DHCPACK**
See other term: DHCP acknowledgment message (DHCPACK)

**DHCPDECLINE**
See other term: DHCP decline message (DHCPDECLINE)

**DHCPINFORM**
See other term: DHCP information message (DHCPINFORM)

**DHCPNAK**
See other term: DHCP negative acknowledgment message (DHCPNAK)

**DHCPOFFER**
See other term: DHCP offer message (DHCPOFFER)

**DHCPRELEASE**
See other term: DHCP release message (DHCPRELEASE)

**DHCPREQUEST**
See other term: DHCP request message (DHCPREQUEST)

**dial location**
The country code, area code, and specific dialing requirements for the place you are dialing from. Once you have created a dial location, you can select it to apply the dialing requirements to all your calls. To change dialing locations, select or create a different one.

**dial-in constraints**
Settings in a remote access policy that permit or deny access to remote access clients.

See also: remote access policy

**dial-up connection**
The connection to your network if you use a device that uses the telephone network. This includes modems with a standard telephone line, ISDN cards with high-speed ISDN lines, or X.25 networks. If you are a typical user, you might have one or two dial-up connections, for example, to the Internet and to your corporate network. In a more complex server situation, multiple network modem connections might be used to implement advanced routing.

See also: Integrated Services Digital Network (ISDN)  modem (modulator/demodulator)

**dial-up line**
A standard dial-up connection, such as telephone and Integrated Services Digital Network (ISDN) lines. Also called switched circuit.

See also: Integrated Services Digital Network (ISDN)

**differential backup**
A backup that copies files created or changed since the last normal or incremental backup. It does not mark files as having been backed up (in other words, the archive attribute is not cleared). If you are performing a combination of normal and differential backups, restoring files and folders requires that you have the last normal as well as the last differential backup.

See also: copy backup  daily backup  incremental backup  normal backup

**differential data**
Saved copies of changed data that can be applied to an original volume to generate a volume shadow copy.

See also: volume  volume shadow copy

**Diffie-Hellman key agreement protocol**
A cryptographic mechanism that allows two parties to establish a shared secret key without having any preestablished secrets between them. Diffie-Hellman is frequently used to establish the shared secret keys that are used by common applications of cryptography, such as Internet Protocol security (IPSec). It is not normally used for data protection.

See also: cryptography  Internet Protocol security (IPsec)  key

**digest**
See other term: Digest authentication

**Digest Access protocol**
A lightweight authentication protocol for parties involved in communications that are based on Hypertext Transfer Protocol (HTTP) or Simple Authentication and Security Layer (SASL).

**Digest authentication**
An authentication mechanism that hashes user name, password, and other data before transmitting it over the network.

See also: authentication  Basic authentication  encryption  hash  password  user name

**digital signature**
Data that binds a sender's identity to the information being sent. A digital signature may be bundled with any message, file, or other digitally encoded information, or transmitted separately. Digital signatures are used in public key environments and provide authentication and integrity services.

**Digital Signature Standard (DSS)**

A standard that uses the Digital Signature Algorithm (DSA) for its signature algorithm and Secure Hash Algorithm (SHA-1) as its message hash algorithm. DSA is a public-key cipher that is used only to generate digital signatures and cannot be used for data encryption.

See also: Secure Hash Algorithm (SHA-1)

**digital subscriber line (DSL)**
A special communication line that uses modulation technology to maximize the amount of data that can be sent over a local-loop, copper-wire connection to a central phone office. DSL is used for connections from telephone switching stations to a subscriber rather than between switching stations.

**digital video disc (DVD)**
A type of optical disc storage technology. A digital video disc (DVD) looks like a CD-ROM disc, but it can store greater amounts of data. DVDs are often used to store full-length movies and other multimedia content that requires large amounts of storage space.

See also: DVD decoder  DVD drive

**direct cable connection**
A link between the input/output (I/O) ports of two computers created with a single cable rather than a modem or other interfacing devices. In most cases, a direct cable connection is made with a null modem cable.

See also: input/output (I/O) port  null modem cable

**direct hosting**
For Microsoft networking, the sending of file and print sharing traffic using the Server Message Block (SMB) protocol (also known as the Common Internet File System (CIFS) protocol) without the use of network basic input/output system (NetBIOS). Direct hosting for the Microsoft redirector (the Workstation service) and file server (the Server service) is supported over both TCP/IP and Internetwork Packet Exchange (IPX). Although direct hosting may be more efficient, a direct hosting client can connect only to a direct hosting server.

See also: Internetwork Packet Exchange (IPX)  network basic input/output system (NetBIOS)  NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)  Server Message Block (SMB)  Transmission Control Protocol/Internet Protocol (TCP/IP)

**direct memory access (DMA)**
Memory access that does not involve the microprocessor. DMA is frequently used for data transfer directly between memory and a peripheral device such as a disk drive.

See also: hardware configuration

**directory**
An information source that contains information about users, computer files, or other objects. In a file system, a directory stores information about files. In a distributed computing environment (such as a Windows domain), the directory stores information about objects such as printers, fax servers, applications, databases, and other users.

See also: domain  object

**directory database**
The physical storage for each replica of Active Directory. Also called the store.

See also: Active Directory  replica

**directory partition**
A contiguous subtree of Active Directory that is replicated as a unit to other domain controllers in the forest that contain a replica of the same subtree. In Active Directory, a single domain controller always holds at least three directory partitions: schema (class and attribute definitions for the directory), configuration (replication topology and related metadata), and domain (subtree that contains the per-domain objects for one domain). Domain controllers running Windows Server 2003 can also store one or more application directory partitions.

See also: Active Directory  application directory partition  attribute  domain  global catalog  replica  replication  schema

**directory service**
Both the directory information source and the service that makes the information available and usable. A directory service enables the user to find an object when given any one of its attributes.

See also: Active Directory  attribute  directory  object

**Directory Services Markup Language (DSML)**
An open, extensible, standards-based format for publishing directory service schemas and exchanging directory contents.

**directory store**
The physical storage for Active Directory directory partition replicas on a specific domain controller. The store is implemented using the Extensible Storage Engine.

See also: Active Directory  directory partition  domain controller  replica

**DirectX**
An extension of the Microsoft Windows operating system. DirectX technology helps games and other programs use the advanced multimedia capabilites of your hardware.

**dirty bit**
A bit that is used to mark modified data in a cache so that the modifications may be carried over to primary memory. For Windows file systems, a bit that is used to indicate that the file system is in an inconsistent and possibly corrupted state. When it is set, the file system`s dirty bit triggers the chkdsk command to run when the computer is restarted.

See also: bit (binary digit)  file system

**disable**
To make a device nonfunctional. For example, if you disable a device in a hardware configuration, you cannot use the device when your computer uses that hardware configuration. Disabling a device frees the resources that were allocated to the device.

See also: device  enable  hardware configuration

**discretionary access control list (DACL)**
The part of an object`s security descriptor that grants or denies specific users and groups permission to access the object. Only the owner of an object can change permissions granted or denied in a DACL; thus, access to the object is at the owner`s discretion.

See also: access control entry (ACE)  distribution group  object  security descriptor  security group  system access control list (SACL)

**disk**
A storage device that is attached to a computer.

See also: basic disk  dynamic disk

**disk configuration information**
Information in the Windows registry on assigned drive letters, simple volumes, striped volumes, mirrored volumes, spanned volumes, and RAID-5 volumes. You can change the disk configuration by using Disk Management.

See also: mirrored volume  RAID-5 volume  registry  simple volume  spanned volume  striped volume  volume

**disk mirroring**
A set of software processes that maintains a backup copy of a volume at all times. Each mirror of a volume resides on a different disk; ideally, each disk has its own controller. If one mirror becomes unavailable (due to a disk failure, for example), you can use the other mirror to gain access to the volume`s data.

See also: mirror  volume

**dismount**
To remove a removable tape or disc from a drive.

See also: library  mount

**display adapter**
See other term: video adapter

**display password**
A password that allows the user to open only previously saved capture (.cap) files.

See also: capture password  password

**distinguished name**
A name that uniquely identifies an object by using the relative distinguished name for the object, plus the names of container

objects and domains that contain the object. The distinguished name identifies the object as well as its location in a tree. Every object in Active Directory has a distinguished name. A typical distinguished name might be CN=MyName,CN=Users,DC=Microsoft,DC=Com. This identifies the MyName user object in the microsoft.com domain.

See also: Active Directory container object domain object relative distinguished name

**Distinguished Name Tag (DNT)**
An integer column that is maintained by the Extensible Storage Engine (ESE). Domain controllers use DNTs when they create objects, either locally or through replication. Each domain controller creates and maintains its own unique DNTs when it creates objects. DNTs are not shared or replicated between domain controllers. A domain controller can create a maximum of approximately 2 billion DNTs over its lifespan. Domain controllers that are installed with the Install from Media (IFM) option inherit the DNT values from the domain controller that was used to create the IFM backup. Therefore, using IFM installations can reduce the total number of DNTs that the newly installed domain controller can create over its lifespan, depending on the number of objects that are created by the domain controller from which the IFM backup was created.

**Distributed Component Object Model (DCOM)**
The Microsoft Component Object Model (COM) specification that defines how components communicate over Windows-based networks. Use the DCOM Configuration tool to integrate client/server applications across multiple computers. DCOM can also be used to integrate robust Web browser applications.

See also: Component Object Model (COM)

**Distributed File System (DFS)**
A service that allows system administrators to organize distributed network shares into a logical namespace, enabling users to access files without specifying their physical location and providing load sharing across network shares.

See also: service

**distribution group**
A group that is used solely for e-mail distribution and that is not security-enabled. Distribution groups cannot be listed in discretionary access control lists (DACLs) used to define permissions on resources and objects. Distribution groups can be used only with e-mail applications (such as Microsoft Exchange) to send e-mail to collections of users. If you do not need a group for security purposes, create a distribution group instead of a security group.

See also: discretionary access control list (DACL) security group

**DLC**
See other term: Data Link Control (DLC)

**DLL**
See other term: dynamic-link library (DLL)

**DMA**
See other term: direct memory access (DMA)

**DMZ**
See other term: perimeter network

**DNS**
See other term: Domain Name System (DNS)

**DNS client**
A client computer that queries DNS servers in an attempt to resolve DNS domain names. DNS clients maintain a temporary cache of resolved DNS domain names.

See also: client DNS server Domain Name System (DNS)

**DNS dynamic update protocol**
See other term: dynamic update

**DNS Security Extensions (DNSSEC)**
DNS extensions that use digital signatures to provide data authentication and integrity to compliant DNS resolvers and applications.

See also: authentication Domain Name System (DNS) resolver resource record (RR) secure zone

**DNS server**

A server that maintains information about a portion of the DNS database and that responds to and resolves DNS queries.

See also: DNS client  Domain Name System (DNS)  server

**DNS suffix**
For DNS, a character string that represents a domain name. The DNS suffix shows where a host is located relative to the DNS root, specifying a host's location in the DNS hierarchy. Usually, the DNS suffix describes the latter portion of a DNS name, following one or more of the first labels of a DNS name.

See also: domain name  Domain Name System (DNS)

**DNS zone**
In a DNS database, a contiguous portion of the DNS tree that is administered as a single, separate entity by a DNS server. The zone contains resource records for all the names within the zone.

See also: DNS server  Domain Name System (DNS)  resource record (RR)

**DNSSEC**
See other term: DNS Security Extensions (DNSSEC)

**docking station**
A unit for housing a portable computer that contains a power connection, expansion slots, and connections to peripherals, such as a monitor, printer, full-sized keyboard, and mouse. The docking station turns the portable computer into a desktop computer.

**document**
Any self-contained piece of work created with an application program and, if saved on disk, given a unique file name by which it can be retrieved.

See also: filter  property cache

**domain**
In Active Directory, a collection of computer, user, and group objects defined by the administrator. These objects share a common directory database, security policies, and security relationships with other domains. In DNS, any tree or subtree within the DNS namespace. Although the names for DNS domains often correspond to Active Directory domains, DNS domains should not be confused with Active Directory domains.

See also: Active Directory  directory database  Domain Name System (DNS)  object

**domain administrator**
A person who is a member of the Domain Admins group. Domain administrators can create, delete, and manage all objects that reside within the domain in which they are administrators. They can also assign and reset passwords and delegate administrative authority for network resources to other trusted users.

See also: domain  object  resource

**domain controller**
In an Active Directory forest, a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. Administrators can manage user accounts, network access, shared resources, site topology, and other directory objects from any domain controller in the forest.

See also: Active Directory  authentication  directory  forest  shared resource

**domain controller locator (Locator)**
An algorithm running in the context of the Net Logon service that enables a client to locate a domain controller. Locator can find domain controllers by using DNS or network basic input/output system (NetBIOS) names. The DNS service (SRV) resource records registered by Locator on behalf of domain controllers are also known as domain controller locator (Locator) resource records.

See also: domain controller  Domain Name System (DNS)  network basic input/output system (NetBIOS)  service (SRV) resource record

**domain DFS**
An implementation of Distributed File System (DFS) in which DFS topological information is stored in Active Directory. Because this information is made available on multiple domain controllers in the domain, domain DFS provides fault tolerance for any distributed file system in the domain.

See also: Active Directory DFS topology Distributed File System (DFS) fault tolerance

**domain functionality**
The functional level of an Active Directory domain that has one or more domain controllers running Windows Server 2003. The functional level of a domain can be raised to enable new Active Directory features that will apply to that domain only. There are four domain functional levels: Windows 2000 mixed, Windows 2000 native, Windows Server 2003 interim, and Windows Server 2003. The default domain functional level is Windows 2000 mixed. When the domain functional level is raised to Windows 2000 native, Windows Server 2003 interim, or Windows Server 2003, advanced domain-wide Active Directory features are available.

See also: Active Directory domain domain controller

**domain hierarchy**
The parent/child tree structure of domains.

See also: domain

**domain local group**
A security or distribution group that can contain universal groups, global groups, other domain local groups from its own domain, and accounts from any domain in the forest. Domain local security groups can be granted rights and permissions on resources that reside only in the same domain where the domain local group is located.

See also: distribution group domain tree forest global group security group universal group

**domain name**
The name given by an administrator to a collection of networked computers that share a common directory. Part of the DNS naming structure, domain names consist of a sequence of name labels separated by periods.

See also: domain Domain Name System (DNS) label

**Domain Name System (DNS)**
A hierarchical, distributed database that contains mappings of DNS domain names to various types of data, such as IP addresses. DNS enables the location of computers and services by user-friendly names, and it also enables the discovery of other information stored in the database.

See also: domain name IP address ping service Transmission Control Protocol/Internet Protocol (TCP/IP)

**domain namespace**
The database structure used by DNS.

See also: Domain Name System (DNS)

**domain naming master**
A domain controller that holds the domain naming operations master role in Active Directory. The domain naming master controls the addition or removal of domains in the forest. At any time, the domain naming master role can be assigned to only one domain controller in the forest.

See also: Active Directory domain controller operations master

**domain of origin**
The parent DNS domain name that is used to root either a zone or a resource record within a zone. This name is joined to the end of unqualified or relative domain names to form a fully qualified domain name (FQDN) within the zone. In DNS Manager, the domain of origin will correspond to the zone name as it appears in the Add Zone Wizard or the name that appears in the Parent domain name field for any resource records created within the zone.

See also: domain domain name Domain Name System (DNS) fully qualified domain name (FQDN) relative name resource record (RR) zone

**domain root**
A DFS namespace, for which the configuration information is stored in Active Directory. The path to access the root or a link starts with the host domain name. A domain root can have multiple root targets, which offer fault tolerance and load sharing at the root level.

See also: Active Directory DFS namespace root target

**domain suffix**
For DNS, an optional parent domain name that can be appended to the end of a relative domain name used in a name query

or host lookup. The domain suffix can be used to complete an alternate fully qualified DNS domain name to be searched when the first attempt to query a name fails.

See also: Domain Name System (DNS) fully qualified domain name (FQDN) parent domain

**domain tree**
In DNS, the inverted hierarchical tree structure that is used to index domain names. Domain trees are similar in purpose and concept to the directory trees used by computer filing systems for disk storage. For example, when numerous files are stored on disk, directories can be used to organize the files into logical collections. When a domain tree has one or more branches, each branch can organize domain names used in the namespace into logical collections. In Active Directory, a hierarchical structure of one or more domains, connected by transitive, bidirectional trusts, that forms a contiguous namespace. Multiple domain trees can belong to the same forest.

See also: Active Directory domain domain name Domain Name System (DNS) forest transitive trust two-way trust

**dots per inch (DPI)**
The standard used to measure screen and printer resolution, expressed as the number of dots that a device can display or print per linear inch. The greater the number of dots per inch, the better the resolution.

**downcasing**
The case-insensitive handling of DNS domain names. Windows DNS servers use downcasing to convert any uppercase letters used in domain names to equivalent lowercase letters. For example, the name HOST.example.Microsoft.com would be downcased to host.example.microsoft.com.

See also: DNS server domain name Domain Name System (DNS) upcasing

**downloadable fonts**
A set of characters stored on disk and sent (downloaded) to a printer`s memory when needed for printing a document. Downloadable fonts are most commonly used with laser printers and other page printers, although many dot-matrix printers can accept some of them. Also called soft fonts.

See also: font font cartridge PostScript fonts

**DPI**
See other term: dots per inch (DPI)

**drag**
To move an item on the screen by selecting the item and then pressing and holding down the mouse button while moving the mouse. For example, you can move a window to another location on the screen by dragging its title bar.

**drain**
For Network Load Balancing, a command that disables new traffic handling for the rule whose port range contains the specified port. All ports specified by the port rule are affected.

See also: cluster drainstop Network Load Balancing port port rule

**drainstop**
For Network Load Balancing, a command that disables all new traffic handling on the specified hosts. The hosts then enter draining mode to complete existing connections. While draining, hosts remain in the cluster and stop their cluster operations when there are no more active connections. To terminate draining mode, explicitly stop cluster mode with the stop command, or restart new traffic handling with the start command. To drain connections from a specific port, use the drain command.

See also: drain host Network Load Balancing

**drive**
An area of storage that is formatted with a file system and has a drive letter. The storage can be a floppy disk, a CD or DVD, a hard disk, or another type of disk.

See also: drive letter file system volume

**drive letter**
The naming convention Windows uses to identify disk drives or partitions on a computer. Drives are named by letter, beginning with A, followed by a colon.

See also: drive

**drop folder**

In the Macintosh environment, a folder for which you have the Make Changes permission but not the See Files or See Folders permission. You can copy files into a drop folder, but you cannot see what files and subfolders the drop folder contains.

See also: Make Changes

**DSL**
See other term: digital subscriber line (DSL)

**DSML**
See other term: Directory Services Markup Language (DSML)

**DSS**
See other term: Digital Signature Standard (DSS)

**DTE**
See other term: Data Terminal Equipment (DTE)

**dual boot**
A computer configuration that can start two different operating systems.

See also: boot  multiple boot  startup environment

**duplex**
A system capable of transmitting information in both directions over a communications channel.

See also: full-duplex  half-duplex

**DVD**
See other term: digital video disc (DVD)

**DVD decoder**
A hardware or software component that allows a digital video disc (DVD) drive to display movies on your computer screen.

See also: digital video disc (DVD)  DVD drive  hardware decoder  software decoder

**DVD drive**
A disk storage device that uses digital video disc (DVD) technology. A DVD drive reads both CD-ROM and DVDs; however, you must have a DVD decoder to display DVD movies on your computer screen.

See also: digital video disc (DVD)  DVD decoder

**DWORD**
A data type that is composed of hexadecimal data with a maximum allotted space of 4 bytes.

**dynamic data exchange (DDE)**
A form of interprocess communication (IPC) implemented in the Microsoft Windows family of operating systems. Two or more programs that support dynamic data exchange (DDE) can exchange information and commands.

See also: Network DDE service

**dynamic disk**
A physical disk that provides features that basic disks do not, such as support for volumes that span multiple disks. Dynamic disks use a hidden database to track information about dynamic volumes on the disk and other dynamic disks in the computer. You convert basic disks to dynamic by using the Disk Management snap-in or the DiskPart command-line tool. When you convert a basic disk to dynamic, all existing basic volumes become dynamic volumes.

See also: active volume  basic disk  basic volume  dynamic volume  partition  volume

**Dynamic Host Configuration Protocol (DHCP)**
A TCP/IP service protocol that offers dynamic leased configuration of host IP addresses and distributes other configuration parameters to eligible network clients. DHCP provides safe, reliable, and simple TCP/IP network configuration, prevents address conflicts, and helps conserve the use of client IP addresses on the network. DHCP uses a client/server model where the DHCP server maintains centralized management of IP addresses that are used on the network. DHCP-supporting clients can then request and obtain lease of an IP address from a DHCP server as part of their network boot process.

See also: DHCP client  DHCP server  IP address  lease  service  Transmission Control Protocol/Internet Protocol (TCP/IP)

**dynamic routing**
The use of routing protocols to update routing tables. Dynamic routing responds to changes in the internetwork topology.

See also: protocol  routing

**dynamic storage**
A storage method in Windows that allows disk and volume management without requiring operating system restart.

See also: basic storage

**dynamic update**
An update to the Domain Name System (DNS) standard that permits DNS clients to dynamically register and update their resource records in zones.

See also: DNS client  DNS server  Domain Name System (DNS)  resource record (RR)  zone

**dynamic volume**
A volume that resides on a dynamic disk. Windows supports five types of dynamic volumes: simple, spanned, striped, mirrored, and RAID-5. A dynamic volume is formatted by using a file system, such as file allocation table (FAT) or NTFS, and has a drive letter assigned to it.

See also: basic disk  basic volume  dynamic disk  mirrored volume  RAID-5 volume  simple volume  spanned volume  striped volume  volume

**dynamic-link library (DLL)**
An operating system feature that allows executable routines (generally serving a specific function or set of functions) to be stored separately as files with .dll extensions. These routines are loaded only when needed by the program that calls them.

See also: Resource DLL

# Windows Server 2003 Glossary - E

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - E**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**EAP**
See other term: Extensible Authentication Protocol (EAP)

**effective permissions**
Overall permissions that a user or group has on an object, taking into account group membership as well as inheritance from the parent object.

See also: group  inheritance  object  parent object  permission

**EFI**
See other term: Extensible Firmware Interface (EFI)

**EFI system partition**
A small portion on a GUID partition table (GPT) disk that is formatted with the file allocation table (FAT) file system and contains the files necessary to start the computer and pass execution to an operating system, which resides on another partition.

See also: Extensible Firmware Interface (EFI)  file allocation table (FAT)  GUID partition table (GPT)  Itanium  Microsoft Reserved (MSR) partition  system volume  x86

**EFS**
See other term: Encrypting File System (EFS)

**ELAN**
See other term: emulated local area network (ELAN)

**embedded object**
Information created in another program that has been pasted inside your document. When information is embedded, you can edit the information in the new document using toolbars and menus from the original program. To edit the embedded information, double-click it and the toolbars and menus from the program used to create the information appear. Embedded information is not linked to the original source. If you change information in one place, it is not updated in the other.

See also: linked object  OLE  package  source document

**emulated local area network (ELAN)**
A logical ATM network that emulates the services of an Ethernet or Token Ring LAN.

See also: asynchronous transfer mode (ATM)  local area network (LAN)  Token Ring

**enable**
To make a device functional. For example, if a device in your hardware configuration settings is enabled, the device is available for use when your computer uses that hardware configuration.

See also: disable  hardware configuration

**encapsulated PostScript (EPS) file**
A file that prints at the highest possible resolution for your printer. An EPS file may print faster than other graphical representations. Some Windows-based and non-Windows-based graphical programs can import EPS files.

See also: PostScript

**Encapsulating Security Payload (ESP)**
An Internet Protocol security (IPSec) protocol that provides confidentiality, in addition to authentication, integrity, and anti-replay. ESP can be used alone, in combination with Authentication Header (AH), or nested with the Layer Two Tunneling

Protocol (L2TP). ESP does not normally sign the entire packet unless it is being tunneled. Ordinarily, just the data payload is protected, not the IP header.

See also: authentication  Authentication Header (AH)  Internet Protocol security (IPsec)  Layer Two Tunneling Protocol (L2TP)

**encapsulation**
The method used to pass data from one protocol over a network within a different protocol. Data from one protocol is wrapped with the header of a different protocol. Encapsulation is described in RFC 1483.

See also: protocol  Request for Comments (RFC)

**encrypted channel**
A path or link through which encrypted information passes between two devices. This information typically includes the primary message content, not session configuration information.

See also: device  encryption

**encrypted password**
A password that is scrambled. Encrypted passwords are more secure than plaintext passwords, which are susceptible to network sniffers.

See also: encryption  password

**Encrypting File System (EFS)**
A Microsoft file-based encryption technology that enables users to encrypt files and folders on NTFS volumes. EFS helps protect the confidentiality of data by ensuring that only authorized users can decrypt the encrypted files or folders.

See also: NTFS file system  recovery agent

**encryption**
The process of disguising a message or data in such a way as to hide its substance.

See also: public key encryption  symmetric encryption

**enforce**
In Group Policy Management console (GPMC), to force the Group Policy object (GPO) links of the parent container to take precedence over the GPO links of the child containers. By default, the GPO links that are closest to the user or computer (links to the child container) have higher precedence than GPO links higher up (links to the parent container). Enforcement causes GPOs that are linked in parent containers to take precedence instead. In operating systems in the Windows Server 2003 family without GPMC, No Override is used instead of Enforce.

See also: block inheritance  Group Policy Management console (GPMC)  Group Policy object link  inheritance  No Override  precedence

**enhanced small device interface (ESDI)**
A standard that can be used with high-capacity hard disks, floppy disk drives, and tape drives to allow these devices to communicate with a computer at high speeds.

**enterprise certification authority**
A certification authority (CA) that is fully integrated with Active Directory.

See also: Active Directory  certification authority (CA)

**Enterprise Memory Architecture**
A feature of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, that allows you to run applications that take advantage of large amounts of physical memory. Enterprise Memory Architecture supports two types of memory enhancement: application memory tuning, also known as 4-gigabyte tuning (4GT), and Physical Address Extension (PAE) X86.

**entry**
The lowest level element in the registry. Entries appear in the details pane of a Registry Editor window. Each entry consists of an entry name, its data type, and its value. Entries store the configuration data that affects the operating system and the programs that run on the system. As such, they are different from registry subtrees, keys, and subkeys, which are containers.

See also: key  registry  subkey  subtree

**environment variable**
A string consisting of environment information, such as a drive, path, or file name, associated with a symbolic name that can be used by Windows. You use System in Control Panel or the set command from the command prompt to define environment variables.

See also: string  variable

**error detection**
A technique for detecting when data is lost during transmission. This allows the software to recover lost data by notifying the transmitting computer that it needs to retransmit the data.

**ESDI**
See other term: enhanced small device interface (ESDI)

**ESP**
See other term: Encapsulating Security Payload (ESP)

**event**
Any significant occurrence in the system or an application that requires users to be notified or an entry to be added to a log.

**Event Log service**
A service that records events in the system, security, and application logs. The Event Log service is located in Event Viewer.

See also: event  event logging  Event Viewer  service

**event logging**
The process of recording an audit entry in the audit trail whenever certain events occur, such as services starting and stopping or users logging on and off and accessing resources.

See also: auditing  event  Event Viewer  service

**Event Viewer**
A component you can use to view and manage event logs, gather information about hardware and software problems, and monitor security events. Event Viewer maintains logs about program, security, and system events.

See also: event  event logging

**everyone category**
In the Macintosh environment, one of the user categories to which you assign permissions for a folder. Permissions granted to everyone apply to all users who use the server, including guests.

See also: permission

**exclusion range**
A small range of one or more IP addresses within a DHCP scope excluded from the DHCP service. Exclusion ranges ensure that these scope addresses will never be offered to clients by the DHCP server.

See also: DHCP server  DHCP service  scope

**expanded memory**
Type of memory that can be added to IBM personal computers. The use of expanded memory is defined by the Expanded Memory Specification (EMS), which supports memory boards containing RAM that can be enabled or disabled by software.

See also: extended memory

**expansion slot**
A socket in a computer, designed to hold expansion boards and connect them to the system bus.

See also: bus

**expire interval**
For DNS, the number of seconds that DNS servers operating as secondary masters for a zone will use to determine if zone data should be expired when the zone is not refreshed and renewed.

See also: DNS server  Domain Name System (DNS)  zone

**explicit permissions**
Object permissions that are defined when the object is created, specifically assigned, or changed by the owner of the object.

See also: object  permission

**express message**
For Message Queuing, a message that uses fewer resources and is faster than a recoverable message. However, because express messages are mapped to memory, they are lost if the computer storing them fails.

See also: Message Queuing  recoverable message

**extended memory**
Memory beyond one megabyte in 80286, 80386, 80486, and Pentium computers.

See also: expanded memory

**extended partition**
A type of partition that you can create only on basic master boot record (MBR) disks. Extended partitions are useful if you want to create more than four volumes on a basic MBR disk. Unlike primary partitions, you do not format an extended partition with a file system and then assign a drive letter to it. Instead, you create one or more logical drives within the extended partition. After you create a logical drive, you format it and assign it a drive letter. An MBR disk can have up to four primary partitions or three primary partitions, one extended partition, and multiple logical drives.

See also: basic disk  drive letter  logical drive  master boot record (MBR)  partition  primary partition  unallocated space  volume

**Extensible Authentication Protocol (EAP)**
An extension to the Point-to-Point Protocol (PPP) that allows for arbitrary authentication mechanisms to be employed for the validation of a PPP connection.

See also: Point-to-Point Protocol (PPP)

**Extensible Firmware Interface (EFI)**
A type of interface between a computer, firmware, hardware, and the operating system. EFI defines a new partition style called GUID partition table (GPT). EFI serves the same purpose as the basic input/output system (BIOS) found in most x86-based computers.

See also: basic input/output system (BIOS)  GUID partition table (GPT)  Itanium  x86

**Extensible Markup Language (XML)**
A meta-markup language that provides a format for describing structured data. This facilitates more precise declarations of content and more meaningful search results across multiple platforms. In addition, XML enables a new generation of Web-based data viewing and manipulation applications.

See also: Hypertext Markup Language (HTML)

**extension-type association**
The association of a file name extension with a Macintosh file type and file creator. By creating extension-type associations, you can choose which program starts automatically when you open a file with a particular extension. Services for Macintosh has many predefined extension-type associations.

**external network number**
A 4-byte hexadecimal number used for addressing and routing purposes on NetWare networks. The external network number is associated with physical network adapters and networks.

See also: internal network number  Internetwork Packet Exchange (IPX)

**external trust**
A trust that is manually created between two Active Directory domains that are located in different forests or between an Active Directory domain and a Windows NT 4.0 or earlier domain. External trusts are nontransitive and one-way.

See also: Active Directory  domain  forest  nontransitive trust  one-way trust  trust relationship

**extranet**
A limited subset of computers or users on a public network, typically the Internet, that can access an organization`s internal network. For example, the computers or users might belong to a partner organization.

# Windows Server 2003 Glossary - F

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - F**


# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z


**failback**
The process of moving resources, either individually or in a group, back to their preferred node after the node has failed and come back online.

See also: failback policy  node  resource

**failback policy**
Parameters that an administrator can set using Cluster Administrator that affect failback operations.

See also: Cluster Administrator  failback

**failed**
A state that applies to a resource or a node in a cluster. A resource or a node is placed in the failed state after an unsuccessful attempt has been made to bring it online.

See also: cluster  node  resource

**failover**
In server clusters, the process of taking resource groups offline on one node and bringing them online on another node. When failover occurs, all resources within a resource group fail over in a predefined order; resources that depend on other resources are taken offline before, and are brought back online after, the resources on which they depend.

See also: failover policy  node  offline  possible owner  server cluster

**failover policy**
Parameters that an administrator can set using Cluster Administrator that affect failover operations.

See also: Cluster Administrator  failover

**FAT**
See other term: file allocation table (FAT)

**FAT32**
A system used to store files on a computer drive. FAT32 is based on the file allocation table (FAT) file system, but it uses 32-bit values for storing files instead of the 16-bit values used by the original FAT file system. FAT32 offers more efficient drive space allocation by creating smaller clusters than FAT and supports volumes of up to 2 terabytes (TB) of size.

**fault tolerance**
The ability of computer hardware or software to ensure data integrity when hardware failures occur. Fault-tolerant features appear in many server operating systems and include mirrored volumes, RAID-5 volumes, and server clusters.

See also: cluster  mirrored volume  RAID-5 volume

**Fax Service**
A system service that provides fax services to local and remote network clients. Fax services include receiving faxes and faxing documents, fax wizard messages, and e-mail messages.

See also: service

**FCB**
See other term: file control block (FCB)

**Federal Information Processing Standard (FIPS)**
A standard entitled Security Requirements for Cryptographic Modules. FIPS 140-1 (1994) and FIPS 140-2 (2001) describe

government requirements for hardware and software cryptomodules used in the U.S. government.

See also: cryptography

**federation**
A pair of realms or domains that have established a federation trust.

**Federation Service**
A security token service that is built into Windows Server 2003 R2. The Federation Service provides tokens in response to requests for security tokens.

**Federation Service Proxy**
A proxy to the Federation Service in the perimeter network (also known as screened subnet). The Federation Service Proxy uses WS-Federation Passive Requestor Profile (WS-F PRP) protocols to collect user credential information from browser clients and Web applications and send the information to the Federation Service on their behalf.

**FEK**
See other term: file encryption key (FEK)

**FEP**
See other term: front-end processor (FEP)

**Fibre Channel**
A networking standard developed to connect devices that require the transmission of large volumes of data at a very high speed. A leading implementation of Fibre Channel technology has been in storage area networks (SANs. Although the term Fibre Channel implies the use of fiber-optic technology, copper coaxial cable is also supported.

**file allocation table (FAT)**
A file system used by MS-DOS and other Windows operating systems to organize and manage files. The file allocation table is a data structure that Windows creates when you format a volume by using FAT or FAT32 file systems. Windows stores information about each file in the file allocation table so that it can retrieve the file later.

See also: file system  NTFS file system

**file control block (FCB)**
A small block of memory temporarily assigned by a computer's operating system to hold information about a file that has been opened for use. An FCB typically contains such information as the file's identification, its location on disk, and a pointer that marks the user's current (or last) position in the file.

**file creator**
A four-character sequence that identifies which program was used to create a file. With Services for Macintosh, you can associate file name extensions with file creators and file types to specify which program starts automatically when you open a file with a particular extension.

See also: extension-type association

**file encryption key (FEK)**
A pseudo-random cryptographic key that Encrypting File System (EFS) uses to encrypt a file. The FEK is encrypted by the public key of the user performing the encryption, and it is typically different for each encrypted file.

See also: Encrypting File System (EFS)  encryption  key  public key

**file fork**
One of two subfiles of a Macintosh file. When Macintosh files are stored on a computer running Services for Macintosh, each fork is stored as a separate file. Each fork can be independently opened by Macintosh users.

**file group**
A File Server Resource Manager option that is used to define a namespace for a file screen, file screen exception, or storage report. It consists of a set of file name patterns, which in turn determine whether files are included or excluded from a group.

**File Replication service (FRS)**
A service that provides multimaster file replication for designated directory trees between designated servers running Windows Server 2003. The designated directory trees must be on disk partitions formatted with the version of NTFS used with the Windows Server 2003 family. FRS is used by Distributed File System (DFS) to automatically synchronize content between assigned replicas and by Active Directory to automatically synchronize content of the system volume information across domain controllers.

See also: Active Directory  NTFS file system  replica  replication  service

**file screen**
A File Server Resource Manager option that is used to block certain files from being saved on a volume or in a folder tree. A file screen is applied at the folder level and affects all folders and subfolders in the designated path.

**File Server for Macintosh**
A service that allows users of Macintosh computers to store, access, and share files on servers running Services for Macintosh. Also called MacFile.

See also: service

**File Server Resource Manager**
A suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers.

**file share**
In a server cluster, any folder that has an associated File Share resource and is managed by the Cluster service. The file share can fail over from one node to another, but to the end user, the folder looks like a regular folder that remains in one location. Multiple users can access a file share.

See also: cluster  resource  service

**File Share resource**
A file share accessible by a network path that is supported as a cluster resource by a Resource DLL.

See also: Resource DLL

**file system**
In an operating system, the overall structure in which files are named, stored, and organized. NTFS, FAT, and FAT32 are types of file systems.

See also: FAT  FAT32  NTFS file system

**file system cache**
An area of physical memory that holds frequently used pages. It allows applications and services to locate pages rapidly and reduces disk activity.

See also: cache

**File Transfer Protocol (FTP)**
A member of the TCP/IP suite of protocols, used to copy files between two computers on the Internet. Both computers must support their respective FTP roles: one must be an FTP client and the other an FTP server.

See also: Transmission Control Protocol/Internet Protocol (TCP/IP)

**file type**
In the Windows environment, a designation of the operational or structural characteristics of a file. The file type identifies the program, such as Microsoft Word, that is used to open the file. File types are associated with a file name extension. For example, files that have the .txt or .log extension are of the Text Document type and can be opened using any text editor. In the Macintosh environment, a four-character sequence that identifies the type of a Macintosh file. The Macintosh Finder uses the file type and file creator to determine the appropriate desktop icon for that file.

**filter**
For Indexing Service, software that extracts content and property values from a document to index them. For Internet Protocol security (IPSec), a specification of Internet Protocol (IP) traffic that provides the ability to trigger security negotiations for a communication based on the source, destination, and type of IP traffic. For Internet Information Services (IIS), a feature of Internet Server Application Programming Interface (ISAPI) that allows preprocessing of requests and postprocessing of responses, permitting site-specific handling of Hypertext Transfer Protocol (HTTP) requests and responses. In IP and Internetwork Packet Exchange (IPX) packet filtering, a definition in a series of definitions that indicates to the router the type of traffic allowed or disallowed on each interface.

See also: Indexing Service  Internet Information Services (IIS)  Internet Protocol (IP)  Internet Protocol security (IPsec)  Internet Server Application Programming Interface (ISAPI)  Internetwork Packet Exchange (IPX)

**filtering mode**

For Network Load Balancing, the method by which network traffic inbound to a cluster is handled by the hosts within the cluster. Traffic can either be handled by a single server, load balanced among the hosts within the cluster, or disabled completely.

See also: cluster  host  load balancing  Network Load Balancing

**FIPS**
See other term: Federal Information Processing Standard (FIPS)

**firewall**
A security solution that segregates one portion of a network from another portion, allowing only authorized network traffic to pass through according to traffic filtering rules.

**Firewire**
See other term: IEEE 1394

**firmware**
Software routines and low-level input/output instructions stored in read-only memory (ROM). Unlike random-access memory (RAM), read-only memory stays intact even in the absence of electrical power.

See also: random access memory (RAM)  read-only memory (ROM)

**flexible single-master operations (FSMO)**
See other term: operations master

**folder**
A container for programs and files in graphical user interfaces, symbolized on the screen by a graphical image (icon) of a file folder. A folder is a means of organizing programs and documents on a disk and can hold both files and additional folders. For DFS Namespaces, any folder that appears after \\ServerOrDomainName\RootName. A folder can have optional folder targets.

**folder target**
A Universal Naming Convention (UNC) path of a shared folder or another namespace that is associated with a folder in a namespace.

**font**
A graphic design applied to a collection of numbers, symbols, and characters. A font describes a certain typeface, along with other qualities such as size, spacing, and pitch.

See also: OpenType font  PostScript fonts  screen font  Type 1 fonts

**font cartridge**
A plug-in unit available for some printers that contains fonts in several styles and sizes. As with downloadable fonts, printers using font cartridges can produce characters in sizes and styles other than those created by the fonts built into it.

See also: downloadable fonts  font

**foreground program**
The program that runs in the active window (the uppermost window with the highlighted title bar). The foreground program responds to commands issued by the user.

See also: background program  title bar

**foreign computer**
A computer that uses another message queuing system but, through a connector application, can exchange messages with computers that run Message Queuing.

See also: connector application  Message Queuing

**foreign security principal**
An object in a domain that represents a security principal that exists in a trusted domain located in a different forest. Foreign security principals are necessary for users in a domain to access resources that exist in a different forest.

See also: domain  forest  object  resource  security principal

**forest**
One or more Active Directory domains that share the same class and attribute definitions (schema), site and replication information (configuration), and forest-wide search capabilities (global catalog). Domains in the same forest are linked with

two-way, transitive trust relationships.

See also: Active Directory domain global catalog schema transitive trust two-way trust

**forest functionality**
The functional level of an Active Directory forest that has one or more domain controllers running Windows Server 2003. The functional level of a forest can be raised to enable new Active Directory features that will apply to every domain in the forest. There are three forest functional levels: Windows 2000, Windows Server 2003 interim, and Windows Server 2003. The default forest functional level is Windows 2000. When the forest functional level is raised to Windows Server 2003 interim or Windows Server 2003, advanced forest-wide Active Directory features are available.

See also: Active Directory domain domain controller forest

**forest root domain**
The first domain created in a new forest. The forest-wide administrative groups, Enterprise Admins and Schema Admins, are located in this domain. As a best practice, new domains are created as children of the forest root domain.

See also: child domain domain domain hierarchy forest

**forest trust**
A trust between two Windows Server 2003 forests that forms trust relationships between every domain in both forests. A forest trust can be created only between the forest root domains in each forest. Forest trusts are transitive, and they can be one-way or two-way. An administrator must manually establish a forest trust, unlike an automatically established trust, such as a parent-child trust.

See also: domain forest one-way trust parent-child trust root domain transitive trust trust relationship two-way trust

**form**
The specification of physical characteristics such as paper size (that is, letter or legal) and printer area margins of paper or other print media. For example, by default, the Letter form has a paper size of 8.5 inches by 11 inches and does not reserve space for margins.

**FORTEZZA**
A family of security products including PCMCIA-based cards, compatible serial port devices, combination cards (such as FORTEZZA/Modem and FORTEZZA/Ethernet), server boards, and others. FORTEZZA is a registered trademark held by the U.S. National Security Agency.

See also: serial port

**forward lookup**
A DNS query for a DNS name.

See also: Domain Name System (DNS)

**forwarder**
A DNS server designated by other internal DNS servers to be used to forward queries for resolving external or offsite DNS domain names.

See also: DNS server domain name Domain Name System (DNS)

**FQDN**
See other term: fully qualified domain name (FQDN)

**fragmentation**
The scattering of parts of the same disk file over different areas of the disk. Fragmentation occurs as files on a disk are deleted and new files are added. It slows disk access and degrades the overall performance of disk operations, although usually not severely.

See also: defragmentation

**frame**
In synchronous communication, a package of information transmitted as a single unit from one device to another.

See also: capture

**frame type**

The way in which a network type, such as Ethernet, formats data to be sent over a network. When multiple frame types are allowed for a particular network type, the packets are structured differently and are, therefore, incompatible. All computers on a network must use the same frame type to communicate. Also called frame format.

See also: frame  packet

**free media pool**
A logical collection of unused data-storage media that can be used by applications or other media pools. When media are no longer needed by an application, they are returned to a free media pool so that they can be used again.

See also: media pool  Removable Storage

**free space**
Available space that you use to create logical drives within an extended partition.

See also: extended partition  logical drive  unallocated space

**front-end processor (FEP)**
In communications, a computer that is located between communications lines and a main (host) computer and used to relieve the host of tasks related to communications; sometimes considered synonymous with communications controller. A front-end processor is dedicated entirely to handling transmitted information, including error detection and control; receipt, transmission, and possibly encoding of messages; and management of the lines running to and from other devices.

See also: host

**FRS**
See other term: File Replication service (FRS)

**FSMO**
See other term: operations master

**FTP**
See other term: File Transfer Protocol (FTP)

**full computer name**
A fully qualified domain name (FQDN). The full computer name is a concatenation of the computer name (for example, client1) and the primary DNS suffix of the computer (for example, reskit.com.). The same computer could be identified by more than one FQDN. However, it has only one full computer name.

See also: DNS suffix  fully qualified domain name (FQDN)

**Full Control**
An access control entry (ACE) that assigns all applicable rights to a file system or directory service object.

See also: access control entry (ACE)  object  permission

**full name**
A user`s complete name, usually consisting of the last name, first name, and middle initial. The full name is information that Local Users and Groups or Active Directory Users and Computers can maintain as part of the information identifying and defining a user account.

See also: Active Directory Users and Computers  user account

**full zone transfer (AXFR)**
The standard query type supported by all DNS servers to update and synchronize zone data when the zone has been changed. When a DNS query is made using AXFR as the specified query type, the entire zone is transferred as the response.

See also: DNS server  zone

**full-duplex**
A system capable of simultaneously transmitting information in both directions over a communications channel.

See also: duplex  half-duplex

**fully qualified domain name (FQDN)**
A DNS name that has been stated to indicate its absolute location in the domain namespace tree. In contrast to relative names, an FQDN has a trailing period (.) to qualify its position to the root of the namespace (host.example.microsoft.com.).

See also: domain name  Domain Name System (DNS)  domain namespace  relative name

# Windows Server 2003 Glossary - G

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - G**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**gateway**
A dedicated device (or a set of services running on a dedicated computer) that routes network traffic and enables communication between different networking protocols. A gateway is a multiprotocol Internet Protocol (IP) router that translates between different transport protocols or data formats.

See also: device  Internet Protocol (IP)

**GDI objects**
Objects from the Graphics Device Interface (GDI) library of application programming interfaces (APIs) for graphics output devices. In Task Manager, the number of GDI objects currently used by a process.

See also: application programming interface (API)  Task Manager

**Generic Application resource**
An application that is supported as a cluster resource by a Resource DLL.

See also: Resource DLL

**Generic Security Services Application Programming Interface (GSS-API)**
A generic API for performing client-server authentication that can be implemented over any authentication system. GSS-API makes it possible for applications to use the same API with different authentication mechanisms.

**Generic Service resource**
A Windows service that is supported as a cluster resource by a Resource DLL.

See also: Resource DLL  service

**GID**
See other term: group identifier (GID)

**global account**
In an Active Directory network, a normal user account in a user`s domain. Most user accounts are global accounts. If there are multiple domains in the network, it is best if each user in the network has only one user account in only one domain, and each user`s access to other domains is accomplished through the establishment of domain trust relationships. In Microsoft Provisioning System, the Exchange server maintains a list of global catalogs, and it maintains a load balance across global catalogs.

See also: Active Directory  domain  global catalog

**global catalog**
A directory database that applications and clients can query to locate any object in a forest. The global catalog is hosted on one or more domain controllers in the forest. It contains a partial replica of every domain directory partition in the forest. These partial replicas include replicas of every object in the forest, as follows: the attributes most frequently used in search operations and the attributes required to locate a full replica of the object. In Microsoft Provisioning System, the Exchange server maintains a list of global catalogs, and it maintains a load balance across global catalogs.

See also: Active Directory  attribute  domain controller  forest  replication

**global group**
A security or distribution group that can contain users, groups, and computers from its own domain as members. Global security groups can be granted rights and permissions for resources in any domain in the forest.

See also: group  local group  member server  permission  user account

**globally unique identifier (GUID)**
A 16-byte value generated from the unique identifier on a device, the current date and time, and a sequence number. A GUID is used to identify a particular device or component.

See also: device

**glue chasing**
In DNS, queries to resolve delegation name server (NS) resource records that do not have corresponding glue address (A) resource records in the same zone.

See also: address (A) resource record  delegation  Domain Name System (DNS)  glue record  name server (NS) resource record  zone

**glue record**
In DNS, a delegation resource record used for locating the authoritative DNS servers for a delegated zone. These records are used to glue zones together and provide an effective delegation and referral path for other DNS servers to follow when resolving a name.

See also: authoritative  delegation  DNS server  Domain Name System (DNS)  glue chasing  resource record (RR)  zone

**GPMC**
See other term: Group Policy Management console (GPMC)

**GPO**
See other term: Group Policy object (GPO)

**GPO link**
See other term: Group Policy object link

**GPT**
See other term: GUID partition table (GPT)

**graphics mode**
A display mode in which lines and characters on the screen are drawn pixel by pixel. Graphics mode displays images by grouping individual dots into shapes, such as the arrowhead of a mouse pointer. It can also preview character formatting, such as boldface and italics, as it will appear in print.

See also: character mode

**gratuitous ARP**
An Address Resolution Protocol (ARP) Request frame sent by a host for the host`s own Internet Protocol version 4 (IPv4) address when the TCP/IP protocol obtains addressing information. Gratuitous ARPs are used to check for duplicate IPv4 addresses on the subnet. If the host receives a reply to the gratuitous ARP request from another host or computer, it detects a conflict for its configured IPv4 address and will not use it.

See also: Address Resolution Protocol (ARP)  frame  host  IP address  subnet

**group**
A collection of users, computers, contacts, and other groups. Groups can be used as security or as e-mail distribution collections. Distribution groups are used only for e-mail. Security groups are used both to grant access to resources and as e-mail distribution lists.

See also: domain  global group  local group

**group account**
A collection of user accounts. By making a user account a member of a group, you give the related user all the rights and permissions granted to the group.

See also: group  user account

**group identifier (GID)**
An identifier in UNIX that associates a user with a group of other users that have something in common. A user can be a member of one or more groups.

**group memberships**
The groups to which a user account belongs. Permissions and rights granted to a group are also provided to its members. In most cases, the actions a user can perform in Windows are determined by the group memberships of the user account to

which the user is logged on.

See also: group  user account

**Group Policy**
The infrastructure within Active Directory directory service that enables directory-based change and configuration management of user and computer settings, including security and user data. You use Group Policy to define configurations for groups of users and computers. With Group Policy, you can specify policy settings for registry-based policies, security, software installation, scripts, folder redirection, remote installation services, and Internet Explorer maintenance. The Group Policy settings that you create are contained in a Group Policy object (GPO). By associating a GPO with selected Active Directory system containers—sites, domains, and organizational units—you can apply the GPO`s policy settings to the users and computers in those Active Directory containers. To create an individual GPO, use the Group Policy Object Editor. To manage Group Policy objects across an enterprise, you can use the Group Policy Management console.

See also: Active Directory  Group Policy Management console (GPMC)  Group Policy object (GPO)  Group Policy Object Editor  Group Policy object link

**Group Policy Management console (GPMC)**
An optional tool that unifies and centralizes administration of Group Policy.

See also: Group Policy

**Group Policy object (GPO)**
A collection of Group Policy settings. GPOs are essentially the documents created by the Group Policy Object Editor. GPOs are stored at the domain level, and they affect users and computers that are contained in sites, domains, and organizational units. In addition, each computer has exactly one group of policy settings stored locally, called the local Group Policy object.

See also: Group Policy  Group Policy Management console (GPMC)  Group Policy Object Editor  Group Policy object link

**Group Policy Object Editor**
The Microsoft Management Console (MMC) snap-in that is used to edit Group Policy objects (GPOs).

See also: Group Policy Management console (GPMC)  Group Policy object (GPO)  Group Policy object link  Microsoft Management Console (MMC)

**Group Policy object link**
A method of applying settings in a Group Policy object (GPO) to an Active Directory container (site, domain, or organizational unit). Linking a GPO applies the settings of that GPO to the users and computers in a site, domain, or organizational unit and, by default, to the users and computers in all child containers.

See also: Active Directory  Group Policy Management console (GPMC)  Group Policy object (GPO)  inheritance  precedence

**Group Policy Security Settings**
The subtrees of the Group Policy Object Editor that allow a security administrator to manually configure security levels assigned to a Group Policy object (GPO) or local computer policy.

See also: Group Policy object (GPO)  Group Policy Object Editor

**GSS-API**
See other term: Generic Security Services Application Programming Interface (GSS-API)

**guest**
A user who does not have a user account or who does not provide a password.

See also: password  user account

**Guest account**
A built-in account used to log on to a computer running Windows when a user does not have an account on the computer or domain or in any of the domains trusted by the computer`s domain.

See also: domain  user account

**GUI mode**
The portion of Setup that uses a graphical user interface (GUI).

See also: Setup  text mode

**GUID**
See other term: globally unique identifier (GUID)

**GUID partition table (GPT)**
A disk-partitioning scheme that is used by the Extensible Firmware Interface (EFI). GPT offers more advantages than master boot record (MBR) partitioning because it allows up to 128 partitions per disk, provides support for volumes up to 18 exabytes in size, allows primary and backup partition tables for redundancy, and supports unique disk and partition IDs (GUIDs).

See also: Extensible Firmware Interface (EFI)  Itanium  master boot record (MBR)

# Windows Server 2003 Glossary - H

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - H**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**H.323**
The International Telecommunication Union - Telecommunication (ITU-T) standard for packet-based multimedia communications.

See also: International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T)  packet

**half-duplex**
A system capable of transmitting information in only one direction at a time over a communications channel.

See also: duplex  full-duplex

**handle**
In the user interface, an interface added to an object that facilitates moving, sizing, reshaping, or other functions pertaining to an object. In programming, a pointer to a pointer, that is, a token that lets a program access an identified resource.

**handle count**
In Task Manager, the number of object handles in a process`s object table.

See also: Task Manager

**handshaking**
A series of signals acknowledging that communication can take place between computers or other devices. A hardware handshake is an exchange of signals over specific wires (other than the data wires), in which each device indicates its readiness to send or receive data. A software handshake consists of signals transmitted over the same wires used to transfer data, as in modem-to-modem communications over telephone lines.

See also: device  modem (modulator/demodulator)

**Hardware Compatibility List (HCL)**
A hardware list that Microsoft compiled for specific products, including Windows 2000 and earlier versions of Windows. The list for a specific product, such as Windows 2000, includes the hardware devices and computer systems that are compatible with that version of the product. For products in the Windows Server 2003 family, you can find the equivalent information on the Windows Catalog Web site.

See also: device  Windows Catalog

**hardware compression**
A feature available on some tape devices that automatically compresses the data that is being stored on the device. This is usually an option that is turned on or off in a backup program.

See also: backup  device

**hardware configuration**
Resource settings that have been allocated for a specific device. Each device on your computer has a hardware configuration, which can consist of interrupt request (IRQ) lines, direct memory access (DMA), an input/output (I/O) port, or memory address settings.

See also: device  direct memory access (DMA)  input/output (I/O) port  interrupt request (IRQ) lines  memory address

**hardware decoder**
A type of digital video disc (DVD) decoder that allows a DVD drive to display movies on your computer screen. A hardware decoder uses both software and hardware to display movies.

See also: DVD decoder  DVD drive  software decoder

**hardware profile**
Data that describes the configuration and characteristics of specific computer equipment. This information can be used to configure computers for using peripheral devices.

See also: device

**hardware type**
A classification for similar devices. For example, Imaging Device is a hardware type for digital cameras and scanners.

See also: device

**hash**
A fixed-size result that is obtained by applying a one-way mathematical function (sometimes called a hash algorithm) to an arbitrary amount of data. If there is a change in the input data, the hash changes. The hash can be used in many operations, including authentication and digital signing. Also called a message digest.

See also: authentication  hash algorithm

**hash algorithm**
An algorithm that produces a hash value of some piece of data, such as a message or session key. With a good hash algorithm, changes in the input data can change every bit in the resulting hash value; for this reason, hashes are useful in detecting any modification in a data object, such as a message. Furthermore, a good hash algorithm makes it computationally infeasible to construct two independent inputs that have the same hash. Typical hash algorithms include MD2, MD4, MD5, and SHA-1. Also called a hash function.

See also: Hash-based Message Authentication Mode (HMAC)  MD2  MD4  MD5  message digest  Secure Hash Algorithm (SHA-1)

**Hash-based Message Authentication Mode (HMAC)**
A mechanism for message authentication using cryptographic hash functions. HMAC can be used with any iterative cryptographic hash function (for example, MD5 and SHA-1) in combination with a secret shared key. The cryptographic strength of HMAC depends on the properties of the underlying hash function.

See also: hash algorithm  MD5  Secure Hash Algorithm (SHA-1)

**HCL**
See other term: Hardware Compatibility List (HCL)

**header error control (HEC)**
The fifth byte in the asynchronous transfer mode (ATM) cell header used to detect and correct errors in the ATM header.

See also: asynchronous transfer mode (ATM)

**headless server**
See other term: remotely administered server

**heartbeat**
A signal emitted at regular intervals by software to indicate that it is still running.

See also: Network Load Balancing  Network Load Balancing cluster  server cluster

**HEC**
See other term: header error control (HEC)

**heterogeneous environment**
An internetwork with servers and workstations running different operating systems, such as Microsoft Windows, Apple Macintosh, or Novell NetWare, using a mix of different transport protocols.

See also: protocol  server

**hexadecimal**
A base-16 number system represented by the digits 0 through 9 and the uppercase or lowercase letters A (equivalent to decimal 10) through F (equivalent to decimal 15).

**high availability**
For server clusters, the restarting of a failed application or the dispersion of the work to remaining computers when a computer or application in the server cluster fails.

See also: server cluster

**histogram**
A chart consisting of horizontal or vertical bars, the widths or heights of which represent the values of certain data.

**hive**
A file in which the system stores a portion of the registry (named for their resemblance to the cellular structure of a beehive). A hive is backed by a single file and a .log file, which are in either the %SYSTEMROOT%\System32\Config folder or the %SYSTEMROOT%\Profiles\%USERNAME% folder. By default, most hive files (Default, SAM, Security, and System) are stored in the %SYSTEMROOT%\System32\Config folder. The %SYSTEMROOT%\Profiles folder contains the user profile for each user of the computer. Because a hive is a file, it can be moved from one system to another. However, you must use the Registry Editor, Regedit.exe, to edit the file.

See also: key  registry  systemroot

**HMAC**
See other term: Hash-based Message Authentication Mode (HMAC)

**home folder**
A folder (usually on a file server) that administrators can assign to individual users or groups. Administrators use home folders to consolidate user files onto specific file servers for easy backup. Home folders are used by some programs as the default folder for the Open and Save As dialog boxes. Sometimes referred to as home directories.

See also: group

**hop count**
The value in the Transport Control field that indicates the number of Internetwork Packet Exchange (IPX) routers that have processed the IPX packet.

See also: Internetwork Packet Exchange (IPX)  packet  router

**host**
Any device on a TCP/IP network that has an Internet Protocol (IP) address. Examples of hosts include servers, workstations, network-interface print devices, and routers. Sometimes used to refer to a specific network computer that is running a service used by network or remote clients. For Network Load Balancing, a cluster consists of multiple hosts connected over a local area network (LAN).

See also: client  cluster  local area network (LAN)  Network Load Balancing  server  service  Transmission Control Protocol/Internet Protocol (TCP/IP)

**host domain**
The domain in which a DFS namespace has been configured.

See also: DFS namespace  domain

**host ID**
The portion of the IP address that identifies a computer within a particular network ID.

See also: IP address  network ID

**host name**
The DNS name of a device on a network. These names are used to locate computers on the network. To find another computer, its host name must either appear in the Hosts file or be known by a DNS server. For most Windows-based computers, the host name and the computer name are the same.

See also: DNS server  Domain Name System (DNS)

**host priority**
For Network Load Balancing, a host`s precedence for handling default network traffic for TCP and UDP ports. It is used if a host within the cluster goes offline, and it determines which host within the cluster will assume responsibility for the traffic previously handled by the offline host.

See also: cluster  host  Network Load Balancing  User Datagram Protocol (UDP)

**host server**
The server on which a root target is located.

See also: root  root target  server

**Hosts file**
A local text file in the same format as the 4.3 Berkeley Software Distribution (BSD) UNIX /etc/hosts file. This file maps host names to IP addresses, and it is stored in the \%Systemroot%\System32\Drivers\Etc folder.

See also: systemroot

**HTML**
See other term: Hypertext Markup Language (HTML)

**HTTP**
See other term: Hypertext Transfer Protocol (HTTP)

**HTTPS**
See other term: Secure Hypertext Transfer Protocol

**hub**
A common connection point for devices in a network. Typically used to connect segments of a local area network (LAN), a hub contains multiple ports. When data arrives at one port, it is copied to the other ports so that all segments of the LAN can see the data.

See also: local area network (LAN)  port  switching hub

**hubbed mode**
A mode in which the Address Resolution Protocol/multicast address resolution service (ARP/MARS) provides asynchronous transfer mode (ATM) addresses to requesting clients in the form of a multicast server (MCS) list value. In this mode, ARP/MARS acts as a multicast server, providing active forwarding of all multicast and broadcast traffic destined for IP addresses contained within the ranges specified in the list.

See also: Address Resolution Protocol (ARP)  asynchronous transfer mode (ATM)  IP address  multicast address resolution service (MARS)  multicast server (MCS)  nonhubbed mode

**hue**
The position of a color along the color spectrum. For example, green is between yellow and blue. This attribute can be set using Display in Control Panel.

See also: saturation

**Hypertext Markup Language (HTML)**
A simple markup language used to create hypertext documents that are portable from one platform to another. HTML files are simple ASCII text files with codes embedded (indicated by markup tags) to denote formatting and hypertext links.

See also: American Standard Code for Information Interchange (ASCII)

**Hypertext Transfer Protocol (HTTP)**
The protocol used to transfer information on the World Wide Web. An HTTP address (one kind of Uniform Resource Locator (URL)) takes the following form: http://www.microsoft.com.

See also: protocol

# Windows Server 2003 Glossary - I

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - I**


# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z


**IAS**
See other term: Internet Authentication Service (IAS)

**ICMP**
See other term: Internet Control Message Protocol (ICMP)

**IDE**
See other term: integrated device electronics (IDE)

**identity**
A person or entity that must be verified by means of authentication, based on criteria such as a password or a certificate.

See also: authentication  password

**Identity Management for UNIX**
A feature of Windows that integrates computers running Windows into an existing UNIX enterprise.

**IEEE**
Institute of Electrical and Electronics Engineers, founded in 1963. IEEE is an organization composed of engineers, scientists, and students, best known for developing standards for the computer and electronics industry.

**IEEE 1394**
A standard for high-speed serial devices such as digital video and digital audio editing equipment.

See also: device

**IETF**
See other term: Internet Engineering Task Force (IETF)

**IGMP**
See other term: Internet Group Management Protocol (IGMP)

**IIS**
See other term: Internet Information Services (IIS)

**IIS Server Instance resource**
A server-instance designation used with Internet Information Services (IIS) that supports the WWW and FTP services. IIS server instances are supported as cluster resources by a Resource DLL. IIS Server Instance resources can have dependencies on IP Address resources, Network Name resources, and Physical Disk resources. Access information for server instances does not fail over.

See also: dependency  failover  Internet Information Services (IIS)  Resource DLL

**IKE**
See other term: Internet Key Exchange (IKE)

**ILMI**
See other term: integrated local management interface (ILMI)

**Image Name**
The name of a process as displayed in Task Manager.

See also: Task Manager

**impersonation**
The ability of a thread to run in the security context of a security principal different from the security principal that started the process. This is usually so that a process can gain access to resources on behalf of a user.

**impersonation token**
An access token that captures the security information of a client process, allowing a service to "impersonate" the client process in security operations.

See also: access token

**import media pool**
A logical collection of data-storage media that has not been cataloged by Removable Storage. Media in an import media pool is cataloged as soon as possible so that they can be used by an application.

See also: media pool  Removable Storage

**in-addr.arpa domain**
A special top-level DNS domain reserved for reverse mapping of IP addresses to DNS host names.

See also: Domain Name System (DNS)  reverse lookup  top-level domains

**in-band connection**
A connection between two computers that relies on a standard network, such as a local area network (LAN) or the Internet, and standard remote administration tools, such as Remote Desktop or Telnet. An in-band connection can only be used to manage computers remotely if both the local and remote computers are in a functional state and accessible on the network.

See also: local area network (LAN)  out-of-band connection  Telnet

**Incoming Forest Trust Builders group**
A group whose members can create incoming, one-way forest trusts to the forest-root domain. For example, members of this group residing in forest A can create a one-way incoming forest trust from forest B. This one-way incoming forest trust allows users in forest A to access resources that are located in forest B. Members of this group are assigned the permission Create Inbound Forest Trust on the forest-root domain. This group has no default members.

See also: forest  forest root domain  forest trust  group  one-way trust  permission  resource

**incremental backup**
A backup that copies only those files created or changed since the last normal or incremental backup. It marks files as having been backed up (in other words, the archive attribute is cleared). If you use a combination of normal and incremental backups to restore your data, you will need to have the last normal backup and all incremental backup sets.

See also: copy backup  daily backup  differential backup  normal backup

**incremental zone transfer (IXFR)**
In DNS, a zone transfer request involving only incremental resource record changes between each version of the zone. An IXFR contrasts with a full zone transfer (AXFR) request for all resource records.

See also: DNS server  Domain Name System (DNS)  full zone transfer (AXFR)  zone  zone transfer

**independent client**
A computer with Message Queuing installed that can host queues and store messages locally. Independent clients do not require synchronous access to a Message Queuing server to send and receive messages, but they can use Message Queuing servers with routing enabled for efficient message routing.

See also: dependent client  Message Queuing server  routing services

**indexed sequential access method (ISAM)**
A technique for indexing database records that allows sequential or random access of records. When records are accessed sequentially they are accessed in the order in which they were entered in the database, and when records are accessed randomly, the records are accessed through an index.)

**Indexing Service**
Software that provides search functions for documents stored on disk, allowing users to search for specific document text or properties.

**infrared (IR)**
Light that is beyond red in the color spectrum. While the light is not visible to the human eye, infrared transmitters and

receivers can send and receive infrared signals.

See also: Infrared Data Association (IrDA)  infrared device  infrared port

**Infrared Data Association (IrDA)**
The industry organization of computer, component, and telecommunications vendors who establish the standards for infrared communication between computers and peripheral devices, such as printers.

See also: infrared (IR)

**infrared device**
A computer, or a computer peripheral such as a printer, that can communicate by using infrared light.

See also: infrared (IR)

**infrared file transfer**
Wireless file transfer between a computer and another computer or device using infrared light.

See also: infrared (IR)

**infrared network connection**
A direct or incoming network connection to a remote access server using an infrared port.

See also: infrared port  remote access server

**infrared port**
An optical port on a computer that enables communication with other computers or devices by using infrared light, without cables. Infrared ports can be found on some portable computers, printers, and cameras.

See also: infrared (IR)  infrared device  port

**infrastructure master**
A domain controller that holds the infrastructure operations master role in Active Directory. The infrastructure master updates the group-to-user reference whenever group memberships change and replicates these changes across the domain. At any time, the infrastructure master role can be assigned to only one domain controller in each domain.

See also: Active Directory  domain controller  operations master

**inheritance**
In security, a mechanism that allows a specific access control entry (ACE) to be copied from the container where it was applied to all children of the container. Inheritance can be used to manage access to a whole subtree of objects in a single update operation. In Active Directory, the ability to build new object classes from existing object classes. The new object is defined as a subclass of the original object class. The original object class becomes a superclass of the new object. A subclass inherits the attributes of the superclass, including structure rules and content rules. In Group Policy, a mechanism that allows policy settings in Group Policy objects (GPOs) that are linked to parent containers to be applied to objects in child containers.

See also: access control entry (ACE)  Active Directory  block inheritance  delegation  enforce  Group Policy Management console (GPMC)  Group Policy object (GPO)  object class  precedence

**inherited permissions**
Permissions on an object that are automatically inherited from its parent object. Inherited permissions cannot be modified.

See also: object  parent object  permission

**initialize**
In Disk Management, the process of detecting a disk or volume and assigning it a status (for example, healthy) and a type (for example, dynamic).

See also: basic disk  basic volume  dynamic disk  dynamic volume

**inode**
A file identifier in UNIX. Each inode is identified by a unique number and contains user and group ownership information, access settings, file type, size, links, and modification information.

**input language**
A setting that specifies the language in which you enter text. When you add a new input language, you also select a keyboard layout and other input methods for that language.

See also: Input Method Editor (IME)

**Input Method Editor (IME)**
A program used to enter the thousands of different characters in written Asian languages with a standard 101-key keyboard. An IME consists of both an engine that converts keystrokes into phonetic and ideographic characters and a dictionary of commonly used ideographic words. As the user enters keystrokes, the IME engine attempts to identify which character or characters that the keystrokes should be converted into.

**input/output (I/O) port**
A channel through which data is transferred between a device and the microprocessor. The port appears to the microprocessor as one or more memory addresses that it can use to send or receive data.

See also: device  memory address  port

**install**
When referring to software, to add program files and folders to your hard disk and related data to your registry so that the software runs properly. Installing contrasts with upgrading, where existing program files, folders, and registry entries are updated to a more recent version. When referring to hardware, to physically connect the device to your computer, to load device drivers onto your computer, and to configure device properties and settings.

See also: device driver  registry  uninstall  upgrade

**integrated device electronics (IDE)**
A type of disk-drive interface in which the controller electronics reside on the drive itself, eliminating the need for a separate adapter card. IDE offers advantages such as look-ahead caching to increase overall performance.

**integrated local management interface (ILMI)**
A limited set of Simple Network Management Protocol (SNMP) functions included in the asynchronous transfer mode (ATM) specification for the ATM user network interface (UNI).

See also: asynchronous transfer mode (ATM)  Simple Network Management Protocol (SNMP)  user-to-network interface (UNI)

**Integrated Services Digital Network (ISDN)**
A digital phone line used to provide higher bandwidth. ISDN in North America is typically available in two forms: Basic Rate Interface (BRI) consists of 2 B-channels at 64 kilobits per second (Kbps) and a D-channel at 16 Kbps; Primary Rate Interface (PRI) consists of 23 B-channels at 64 Kbps and a D-channel at 64 Kbps. An ISDN line must be installed by the phone company at both the calling site and the called site.

See also: B-channel  D-channel  multilink dialing  Service Profile Identifier (SPID)  switch type

**Integrated Services over Slow Links (ISSLOW)**
A queuing mechanism used to optimize slow (low-capacity) network interfaces by reducing latency. In particular, it is designed for interfaces that forward traffic to modem links, Integrated Services Digital Network (ISDN) B-channels, and sub-T1 links.

See also: B-channel  Integrated Services Digital Network (ISDN)

**Integrated Windows authentication**
A configuration setting that enables negotiation of authentication protocols in Internet Information Services (IIS).

See also: authentication protocol  Internet Information Services (IIS)

**IntelliMirror**
A set of change and configuration management features based on Active Directory that enables management of user and computer data and settings, including security data. IntelliMirror also provides limited ability to deploy software to Windows 2000 and later workstations or servers.

See also: Active Directory

**interactive dialog box**
A dialog box that requires a response from the user. Intermediary devices such as a security host require such a dialog box as an added layer of security between the client and the remote access server. In such dialog boxes, the user types an access code or a user name and password on the remote access terminal screen.

See also: intermediary device  remote access server  static dialog box

**interactive logon**

The process of logging on to a local computer by using a keyboard.

**interconnect**
A private network that connects nodes in a cluster.

See also: cluster  node

**intermediary device**
A device other than a modem or X.25 PAD, located between a network connection and the remote access server. This device is typically a modem-pool switch or security host and requires either a static or interactive dialog box between the client and itself.

See also: interactive dialog box  PAD (packet assembler/disassembler)  remote access server  static dialog box

**internal network number**
A 4-byte hexadecimal number used for addressing and routing purposes. The internal network number identifies a virtual network inside a computer. The internal network number must be unique to the IPX internetwork. Also called virtual network number.

See also: external network number  Internetwork Packet Exchange (IPX)

**internal private queue**
For Message Queuing, a queue that stores various types of administrative messages, or an interim queue for storing and forwarding messages in transit to a destination queue. Internal private queues are not displayed in Microsoft Management Console (MMC) snap-ins, and they cannot be deleted.

See also: Message Queuing  private queue  queue

**internal router**
A router for which all the networks that it is connected to belong in the same area.

See also: router

**international prefix**
Digits dialed before the country code to access the international phone service. The actual digits depend on the country or region in which you are dialing an international number. For example, in the United States of America, the prefix for international dialing is 011. To dial from the United States of America to Honduras, which has the country code 504, you would dial: (011) (504) (000) 000-0000

**International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T)**
The sector of the International Telecommunication Union (ITU) responsible for telecommunication standards. ITU-T replaces the Comite Consultatif International Telegraphique et Telephonique (CCITT). Its responsibilities include standardizing modem design and operations, and standardizing protocols for networks and facsimile transmission. ITU is an international organization within which governments and the private sector coordinate global telecom networks and services.

**Internet**
A worldwide network of computers. If you have access to the Internet, you can retrieve information from millions of sources, including schools, governments, businesses, and individuals.

See also: router  World Wide Web (WWW)

**Internet address**
An address for a resource on the Internet that is used by Web browsers to locate Internet resources. An Internet address typically starts with a protocol name, followed by the name of the organization that maintains the site; the suffix identifies the kind of organization it is. For example, the address http://www.yale.edu/ provides the following information: http: This Web server uses the Hypertext Transfer Protocol. www: This site is on the World Wide Web. edu: This is an educational institution. Internet address is also called Uniform Resource Locator (URL).

See also: protocol  Web server

**Internet Authentication Service (IAS)**
The Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server, which provides authentication and accounting for network access, and proxy, which provides forwarding of RADIUS messages.

See also: authentication  Certificate Services  Remote Access Service (RAS)  Remote Authentication Dial-In User Service (RADIUS)  service  virtual private network (VPN)

**Internet Control Message Protocol (ICMP)**
A required maintenance protocol in the TCP/IP suite that reports errors and provides simple diagnostic capabilities. ICMP is used by the Ping tool to perform TCP/IP troubleshooting.

See also: Internet Protocol (IP) protocol Transmission Control Protocol/Internet Protocol (TCP/IP)

**Internet Engineering Task Force (IETF)**
An open community of network designers, operators, vendors, and researchers concerned with the evolution of Internet architecture and the smooth operation of the Internet. Technical work is performed by working groups organized by topic areas (such as routing, transport, and security) and through mailing lists. Internet standards are developed in IETF Requests for Comments (RFCs), which are a series of notes that discuss many aspects of computing and computer communication, focusing on networking protocols, programs, and concepts.

See also: Request for Comments (RFC)

**Internet Group Management Protocol (IGMP)**
A protocol used by Internet Protocol version 4 (IPv4) hosts to report their multicast group memberships to any immediately neighboring multicast routers.

See also: Internet Protocol multicasting Multicast Listener Discovery (MLD) protocol

**Internet group name**
A name known by a DNS service that includes a list of the specific addresses of systems that have registered the name.

See also: Domain Name System (DNS)

**Internet Information Services (IIS)**
Software services that support Web site creation, configuration, and management, along with other Internet functions. Internet Information Services include Network News Transfer Protocol (NNTP), File Transfer Protocol (FTP), and Simple Mail Transfer Protocol (SMTP).

See also: File Transfer Protocol (FTP) Network News Transfer Protocol (NNTP) Simple Mail Transfer Protocol (SMTP)

**Internet Key Exchange (IKE)**
A protocol that establishes the security association and shared keys necessary for two parties to communicate by using Internet Protocol security (IPSec).

See also: Internet Protocol security (IPsec) protocol

**Internet Printing Protocol (IPP)**
The protocol that uses the Hypertext Transfer Protocol (HTTP) to send print jobs to printers throughout the world. Windows 2000, Windows XP, and the Windows Server 2003 family support Internet Printing Protocol (IPP) version 1.0.

See also: protocol

**Internet Protocol (IP)**
A routable protocol in the TCP/IP protocol suite that is responsible for IP addressing, routing, and the fragmentation and reassembly of IP packets.

See also: Compression Control Protocol (CCP) packet Transmission Control Protocol/Internet Protocol (TCP/IP) voluntary tunnel

**Internet Protocol multicasting**
The extension of local area network multicasting technology to a TCP/IP network. Hosts send and receive multicast datagrams, the destination fields of which specify IP host group addresses rather than individual IP addresses. A host indicates that it is a member of a group by means of the Internet Group Management Protocol (IGMP).

See also: Internet Group Management Protocol (IGMP) Transmission Control Protocol/Internet Protocol (TCP/IP)

**Internet Protocol security (IPsec)**
A set of industry-standard, cryptography-based protection services and protocols. IPsec protects all protocols in the TCP/IP protocol suite except Address Resolution Protocol (ARP). For virtual private network (VPN) connections, IPsec is used in conjunction with Layer Two Tunneling Protocol (L2TP).

See also: Address Resolution Protocol (ARP) Layer Two Tunneling Protocol (L2TP) protocol Transmission Control Protocol/Internet Protocol (TCP/IP) virtual private network (VPN)

**Internet SCSI (iSCSI)**
An industry standard developed to enable transmission of SCSI commands over the existing Internet Protocol (IP) network by using the TCP/IP protocol. iSCSI offers the possibility of delivering both messaging traffic and block-based storage over IP networks without installing a separate Fibre Channel network.

**Internet Server Application Programming Interface (ISAPI)**
An application programming interface (API) that resides on a server computer for initiating software services tuned for Windows operating systems. In Microsoft Provisioning System, ISAPI resides on the Web server.

See also: application programming interface (API)

**Internet service provider (ISP)**
A company that provides individuals or companies access to the Internet and the World Wide Web. An ISP provides a telephone number, a user name, a password, and other connection information so users can connect their computers to the ISP's computers. An ISP typically charges a monthly or hourly connection fee.

See also: Web server

**Internetwork Packet Exchange (IPX)**
A network protocol native to NetWare that controls addressing and routing of packets within and between local area networks (LANs). IPX does not guarantee that a message will be complete (no lost packets).

See also: Compression Control Protocol (CCP) Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) local area network (LAN) Routing Information Protocol over IPX (RIPX)

**Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX)**
Transport protocols used in Novell NetWare networks, which together correspond to the combination of TCP and IP in the TCP/IP protocol suite. Windows implements IPX through NWLink.

See also: Internetwork Packet Exchange (IPX) NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) Transmission Control Protocol/Internet Protocol (TCP/IP)

**interrupt**
A request for attention from the processor. When the processor receives an interrupt, it suspends its current operations, saves the status of its work, and transfers control to a special routine known as an interrupt handler, which contains the instructions for dealing with the particular situation that caused the interrupt.

See also: interrupt request (IRQ)

**interrupt request (IRQ)**
A signal sent by a device to get the attention of the processor when the device is ready to accept or send information. Each device sends its interrupt requests over a specific hardware line. Each device must be assigned a unique IRQ number.

See also: interrupt

**interrupt request (IRQ) lines**
Hardware lines over which devices can send signals to get the attention of the processor when the device is ready to accept or send information. Each device must have a unique IRQ line.

See also: device interrupt request (IRQ)

**intersite messaging service (ISM)**
A service that supports transports for asynchronous, site-to-site messaging. Each transport serves two major roles: send/receive and topology queries (such as, what are the various sites connected by this transport, and at what cost?). The intersite messaging services shipped in Windows are remote procedure call (RPC) and Simple Mail Transfer Protocol (SMTP) (mail).

See also: remote procedure call (RPC) service Simple Mail Transfer Protocol (SMTP)

**intersite replication**
In Active Directory, the replication of directory partition updates between sites. Intersite replication occurs between bridgehead servers that store the same domain or application directory partition. One bridgehead server per domain or application directory partition and per replication transport is designated automatically in each site. If two sites have no domains or application directory partitions in common, a bridgehead server in each site replicates both the configuration and schema directory partitions between the two sites.

See also: Active Directory bridgehead server directory partition intrasite replication replication site

**Intersite Topology Generator**
An Active Directory process that runs on one domain controller in a site that considers the cost of intersite connections, checks if previously available domain controllers are no longer available, and checks if new domain controllers have been added. The Knowledge Consistency Checker (KCC) process then updates the intersite replication topology accordingly.

See also: Active Directory domain controller Knowledge Consistency Checker (KCC) replication topology site

**intersite topology generator (ISTG)**
A server in an Active Directory architecture that is responsible for managing the inbound replication connection objects for all the server contact points in the server's site.

**intrasite replication**
In Active Directory, the replication of directory partition updates that occurs between two or more domain controllers that store the same domain or application directory partition and that reside within the same site.

See also: Active Directory directory partition domain controller intersite replication site

**IP**
See other term: Internet Protocol (IP)

**IP address**
For Internet Protocol version 4 (IPv4), a 32-bit address used to identify an interface on a node on an IPv4 internetwork. Each interface on the IP internetwork must be assigned a unique IPv4 address, which is made up of the network ID, plus a unique host ID. This address is typically represented with the decimal value of each octet separated by a period (for example, 192.168.7.27). You can configure the IP address statically or dynamically by using Dynamic Host Configuration Protocol (DHCP). For Internet Protocol version 6 (IPv6), an identifier that is assigned at the IPv6 layer to an interface or set of interfaces and that can be used as the source or destination of IPv6 packets.

See also: Internet Protocol (IP) node scope

**IP Address resource**
A 32-bit number in dotted decimal format that represents an Internet Protocol (IP) address and is supported as a cluster resource by a Resource DLL provided with Windows.

See also: IP address Resource DLL

**IPP**
See other term: Internet Printing Protocol (IPP)

**IPsec**
See other term: Internet Protocol security (IPsec)

**IPsec policy**
Configuration policy that defines which traffic Internet Protocol security (IPsec) examines, how that traffic is secured and encrypted, and how IPsec peers are authenticated.

See also: authentication encryption Internet Protocol security (IPsec) policy

**IPX/SPX**
See other term: Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX)

**IR**
See other term: infrared (IR)

**IrDA**
See other term: Infrared Data Association (IrDA)

**IRQ**
See other term: interrupt request (IRQ)

**ISAM**
See other term: indexed sequential access method (ISAM)

**ISAPI**
See other term: Internet Server Application Programming Interface (ISAPI)

**iSCSI**

See other term: Internet SCSI (iSCSI)

**iSCSI initiator**
A logical entity that enables a server to communicate with an iSCSI device over an Internet Protocol (IP) network. Each iSCSI initiator can have one or more network adapters through which communication is established.

**iSCSI target**
A logical entity created in order to manage the connections between an iSCSI device and the servers that need to access it over an Internet Protocol (IP) network. A target defines the portals (IP addresses) that can be used to connect to the iSCSI device, as well as the security settings (if any) that the iSCSI device requires in order to authenticate the servers that are requesting access to its resources.

**ISDN**
See other term: Integrated Services Digital Network (ISDN)

**ISM**
See other term: intersite messaging service (ISM)

**ISP**
See other term: Internet service provider (ISP)

**ISTG**
See other term: intersite topology generator (ISTG)

**Itanium**
An Intel microprocessor that uses explicitly parallel instruction set computing and 64-bit memory addressing. "Itanium-based" refers to systems or platforms that are based on the Itanium processor. "Itanium 2-based" refers to systems or platforms that are based on the Itanium 2 processor. "Itanium architecture-based" refers to systems or platforms that are based on the Itanium and Itanium 2 processors.

**iteration**
See other term: iterative query

See also: DNS server  Domain Name System (DNS)  iterative query  recursive resolution  referral

**iterative query**
A query made to a DNS server for the best answer the server can provide without seeking further help from other DNS servers. Also called a nonrecursive query.

See also: DNS server  Domain Name System (DNS)  recursive query  referral

**ITU-T**
See other term: International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T)

**IXFR**
See other term: incremental zone transfer (IXFR)

# Windows Server 2003 Glossary - J

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - J**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**job object**
A system-level structure that allows processes to be grouped together and managed as a single unit.

See also: process

**join**
In Active Directory Application Mode (ADAM) or Active Directory Lightweight Directory Services (AD LDS), to create a replica of an existing ADAM or AD LDS instance in the same configuration set. When you create a replica of an existing ADAM or AD LDS instance, you join it to the same configuration set that hosts the original (existing) instance.

**journal**
For Message Queuing, a queue where journal messages are stored.

See also: journal quota  Message Queuing  queue  source journaling  target journaling

**journal quota**
For Message Queuing, a property that specifies the cumulative limit for journal messages in a journal queue. The limit is based on cumulative journal message size. When a journal queue`s quota is reached, messages are no longer stored.

See also: journal  Message Queuing  property

**junction point**
A physical location on a hard disk that points to data located at another location on your hard disk or another storage device. Junction points are created when you create a mounted drive. You can also create a junction point using the linkd command.

See also: mounted drive

# Windows Server 2003 Glossary - K

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - K**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**KDC**

See other term: Key Distribution Center (KDC)

**Kerberos V5 authentication protocol**

An authentication mechanism used to verify user or host identity. The Kerberos V5 authentication protocol is the default authentication service. Internet Protocol security (IPSec) can use the Kerberos protocol for authentication.

See also: Internet Protocol security (IPsec) Key Distribution Center (KDC) realm service ticket ticket-granting service (TGS) ticket-granting ticket (TGT)

**kernel mode**

A highly privileged mode of operation where program code has direct access to all memory, including the address spaces of all user-mode processes and applications, and to hardware. Also known as supervisor mode, protected mode, or Ring 0.

**key**

In Registry Editor, a folder that appears in the left pane of the Registry Editor window. A key can contain subkeys and entries. For example, Environment is a key of HKEY_CURRENT_USER. In IP security (IPSec), a value used in combination with an algorithm to encrypt or decrypt data. Key settings for IPSec are configurable to provide greater security.

See also: entry Internet Protocol security (IPsec) registry subkey

**Key Distribution Center (KDC)**

A network service that supplies session tickets and temporary session keys used in the Kerberos V5 authentication protocol.

See also: Kerberos V5 authentication protocol ticket-granting service (TGS) ticket-granting ticket (TGT)

**keyboard layout**

The arrangement that accommodates the special characters and symbols used in different languages. Keyboard layouts affect which characters appear when users press the keys on their keyboards. After the keyboard layout is changed, the characters that appear on the screen may no longer correspond to the characters that are printed on the keyboard keys.

**Knowledge Consistency Checker (KCC)**

A built-in process that runs on all domain controllers and generates the replication topology for the Active Directory forest. At specified intervals, the KCC reviews and makes modifications to the replication topology to ensure propagation of data either directly or transitively.

See also: Active Directory domain controller forest replication topology

# Windows Server 2003 Glossary - L

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - L**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**L2TP**

See other term: Layer Two Tunneling Protocol (L2TP)

**label**

In DNS, each part of a DNS domain name that represents a node in the domain namespace tree. For example, the three labels "example," "microsoft," and "com" make up the DNS domain name example.microsoft.com. Each label used in a DNS name cannot exceed 63 octets, 255 bytes including the terminating dot, for the fully qualified domain name (FQDN).

See also: domain name  Domain Name System (DNS)  fully qualified domain name (FQDN)  node

**LAN**

See other term: local area network (LAN)

**LAN emulation (LANE)**

A group of software components that allows asynchronous transfer mode (ATM) to work with legacy networks and applications. With LANE, you can run your traditional local area network (LAN)-aware applications and protocols on an ATM network without modification. LANE makes the ATM protocol layers appear to be an Ethernet or Token Ring LAN to overlying protocols and applications. LANE provides an intermediate step between fully exploiting ATM and not using ATM at all.

See also: asynchronous transfer mode (ATM)  local area network (LAN)  protocol  Token Ring

**LAN emulation client (LEC)**

The client on an emulated local area network (ELAN) that performs data forwarding, address resolution, and other control functions. The LEC resides on end stations in an ELAN.

See also: client  emulated local area network (ELAN)  LAN emulation (LANE)

**LAN emulation configuration server**

The service that assigns individual local area network emulation (LANE) clients to particular emulated local area networks (ELANs) by directing them to the LAN emulation server (LES).

See also: emulated local area network (ELAN)  LAN emulation (LANE)  LAN emulation server (LES)  service

**LAN emulation server (LES)**

The central control point for an emulated local area network (ELAN). LES enables LAN emulation (LANE) clients to join the ELAN, and it resolves local area network (LAN) addresses to asynchronous transfer mode (ATM) addresses.

See also: asynchronous transfer mode (ATM)  emulated local area network (ELAN)  LAN emulation (LANE)  local area network (LAN)

**LANE**

See other term: LAN emulation (LANE)

**Last Known Good Configuration**

A hardware configuration that is available by pressing F8 during startup. If the current hardware settings prevent the computer from starting, Last Known Good Configuration can allow you to start the computer and examine the configuration. When Last Known Good Configuration is used, later configuration changes are lost.

See also: Safe mode

**latency**

In Active Directory replication, the delay between the time an update is applied to a specific replica and the time it is applied to

some other replica. Also known as propagation delay.

See also: Active Directory replica replication

**Layer Two Tunneling Protocol (L2TP)**
An industry-standard Internet tunneling protocol that provides encapsulation for sending Point-to-Point Protocol (PPP) frames across packet-oriented media. For IP networks, L2TP traffic is sent as User Datagram Protocol (UDP) messages. In Microsoft operating systems, L2TP is used in conjunction with Internet Protocol security (IPSec) as a virtual private network (VPN) technology to provide remote access or router-to-router VPN connections. L2TP is described in RFC 2661.

See also: frame Internet Protocol security (IPsec) Point-to-Point Protocol (PPP) tunnel User Datagram Protocol (UDP)

**Layer Two Tunneling Protocol/Internet Protocol security (L2TP/IPsec)**
A virtual private network (VPN) connection method that provides session authentication, address encapsulation, and strong encryption of private data between remote access servers and clients. L2TP provides address encapsulation and user authentication, and Internet Protocol security (IPsec) provides computer authentication and encryption of the L2TP session.

See also: authentication encapsulation encryption Internet Protocol security (IPsec) Layer Two Tunneling Protocol (L2TP) virtual private network (VPN)

**LDAP**
See other term: Lightweight Directory Access Protocol (LDAP)

**LDAP Data Interchange Format (LDIF)**
A draft Internet standard for a file format that can be used to perform batch operations on directories that conform to Lightweight Directory Access Protocol (LDAP) standards.

See also: batch program Lightweight Directory Access Protocol (LDAP)

**LDIF**
See other term: LDAP Data Interchange Format (LDIF)

**lease**
The length of time for which a DHCP client can use a dynamically assigned IP address configuration. Before the lease time expires, the client must either renew or obtain a new lease with DHCP.

See also: DHCP client Dynamic Host Configuration Protocol (DHCP) IP address

**LEC**
See other term: LAN emulation client (LEC)

**LES**
See other term: LAN emulation server (LES)

**library**
A data-storage system, usually managed by Removable Storage. A library consists of removable media (such as tapes or discs) and a hardware device that can read from or write to the media. There are two major types of libraries: robotic libraries (automated multiple-media, multidrive devices) and stand-alone drive libraries (manually operated, single-drive devices). A robotic library is also called a jukebox or changer.

See also: Removable Storage

**Lightweight Directory Access Protocol (LDAP)**
The primary access protocol for Active Directory. LDAP is an industry-standard protocol, established by the Internet Engineering Task Force (IETF), that allows users to query and update information in a directory service. Active Directory supports both LDAP version 2 and LDAP version 3.

See also: Active Directory directory service Internet Engineering Task Force (IETF) protocol

**Line Printer Daemon (LPD)**
A service on a print server that receives print jobs from Line Printer Remote (LPR) tools that are running on client computers.

See also: Line Printer Remote (LPR) print job print server service

**Line Printer Remote (LPR)**
A connectivity tool that runs on client computers and that is used to print files to a computer running a Line Printer Daemon (LPD) server.

See also: Line Printer Daemon (LPD)

**link state database**
A map of an area maintained by Open Shortest Path First (OSPF) routers. It is updated after any change in the network topology. The link state database is used to compute Internet Protocol (IP) routes, which must be computed again after any change in the topology.

See also: adjacency  Internet Protocol (IP)  Open Shortest Path First (OSPF)

**linked object**
An object that is inserted into a document but still exists in the source file. When information is linked, the new document is updated automatically if the information in the original document changes. If you want to edit the linked information, double-click it. The toolbars and menus from the original program will appear. If the original document is on your computer, changes that you make to the linked information will also appear in the original document.

See also: embedded object  OLE  package  source document

**LIS**
See other term: logical IP subnet (LIS)

**Lmhosts file**
A local text file that maps network basic input/output (NetBIOS) names (commonly used for computer names) to IP addresses for hosts that are not located on the local subnet. In this version of Windows, this file is stored in the systemroot\System32\Drivers\Etc folder.

See also: host  IP address  network basic input/output system (NetBIOS)  systemroot

**load balancing**
A technique used by Windows Clustering to scale the performance of a server-based program (such as a Web server) by distributing its client requests across multiple servers within the cluster. Each host can specify the load percentage that it will handle, or the load can be equally distributed across all the hosts. If a host fails, Windows Clustering dynamically redistributes the load among the remaining hosts.

See also: cluster  host

**load weight**
A Network Load Balancing parameter that specifies the relative amount of load-balanced network traffic that the host handles for the associated port rule. Allowed values range from 0 (zero) to 100. To prevent a host from handling any network traffic, set the load weight to 0 (zero).

See also: Network Load Balancing  Network Load Balancing cluster

**local area network (LAN)**
A network of computers, printers, and other devices located within a relatively limited area (for example, a building). A LAN enables any connected device to interact with any other on the network.

See also: NetBIOS Extended User Interface (NetBEUI)  network basic input/output system (NetBIOS)  virtual local area network (VLAN)  workgroup

**local computer**
The computer that you are currently logged on to as a user. More generally, a local computer is a computer that you can access directly without using a communications line or a communications device, such as a network adapter or a modem.

See also: remote computer

**local group**
A security group that can be granted rights and permissions only to resources on the computer on which the group is created. Local groups can have any user accounts that are local to the computer as members, as well as users, groups, and computers from a domain to which the computer belongs.

See also: global group  member server  user account

**local printer**
A printer that is directly connected to one of the ports on a computer.

See also: port

**Local Security Authority (LSA)**
A protected subsystem that authenticates and logs users on to the local computer. In addition, the LSA maintains information about all aspects of local security on a computer (collectively known as the local security policy), and it provides various services for translation between names and identifiers.

See also: local security policy  security

**Local Security Authority Subsystem Service (LSASS)**
An interface for managing local security, domain authentication, and Active Directory processes. It handles authentication for the client and for the server. It also contains features that are used to support Active Directory utilities.

**local security policy**
Security information about all aspects of local security on a computer. The local security policy identifies who is assigned privileges and what security auditing is to be performed.

See also: auditing  security

**local service**
A service that performs a specific operating system function for other local processes on the computer.

See also: service

**Local Service account**
A predefined local account that is used to start a service and provide the security context for that service. The name of the account is NT AUTHORITY\LocalService. The Local Service account has limited access to the local computer and Anonymous access to network resources.

See also: access control list (ACL)  Local System account  Network Service account  service

**local storage**
For the Windows Server 2003 family, NTFS file system disk volumes used as primary data storage. Such disk volumes can be managed by Remote Storage by copying infrequently accessed files to remote (secondary) storage.

See also: NTFS file system  Remote Storage  volume

**Local System account**
A predefined local account that is used to start a service and provide the security context for that service. The name of the account is NT AUTHORITY\System. This account does not have a password, and any password information that you supply is ignored. The Local System account has full access to the system, including the directory service on domain controllers. Because the Local System account acts as a computer on the network, it has access to network resources.

See also: directory service  domain controller  Local Service account  Network Service account  service

**local user profile**
A computer-based record about an authorized user that is created automatically on the computer the first time a user logs on to a workstation or server computer.

**LocalTalk**
The Apple networking hardware built into every Macintosh computer. LocalTalk includes the cables and connector boxes that connect components and network devices that are part of the AppleTalk network system. Formerly known as AppleTalk Personal Network.

**log file**
A file that stores messages generated by an application, service, or operating system. These messages are used to track the operations performed. For example, Web servers maintain log files listing every request made to the server. Log files are usually plain text (ASCII) files and often have a .log extension. In Backup, a file that contains a record of the date the tapes were created and the names of files and directories successfully backed up and restored. The Performance Logs and Alerts service also creates log files.

See also: American Standard Code for Information Interchange (ASCII)  service

**logical drive**
A volume that you create within an extended partition on a basic master boot record (MBR) disk. Logical drives are similar to primary partitions, except that you are limited to four primary partitions per disk, whereas you can create an unlimited number of logical drives per disk. A logical drive can be formatted and assigned a drive letter.

See also: basic disk  basic volume  drive letter  extended partition  master boot record (MBR)  primary partition  volume

**logical IP subnet (LIS)**
A group of Internet Protocol (IP) hosts that belong to the same IP subnet.

See also: host  Internet Protocol (IP)  subnet

**logical printer**
The software interface between the operating system and the printer in Windows. While a printer is the device that does the actual printing, a logical printer is its software interface on the print server. This software interface determines how a print job is processed and how it is routed to its destination (to a local or network port, to a file, or to a remote print share). When you print a document, it is spooled (or stored) on the logical printer before it is sent to the printer itself.

See also: printer  spooling

**logical unit number (LUN)**
A logical reference to a portion of a storage subsystem. For example, in a disk storage subsystem, a LUN can comprise a disk, a section of a disk, a whole disk array, or a section of a disk array in the subsystem.

**logon event category**
In auditing, a group of events that are logged when a user attempts to log on to or log off from a computer.

See also: auditing  event

**logon rights**
User rights that are assigned to a user and that specify the ways in which a user can log on to a system. An example of a logon right is the right to log on to a system remotely.

See also: user rights

**logon script**
A file, typically a batch file, that runs automatically every time a user logs on to a computer or network. It can be used to configure a user`s working environment whenever a user logs on, and it allows an administrator to influence a user`s environment without managing all aspects of it. A logon script can be assigned to one or more user accounts.

See also: batch program  logon script path  user account

**logon script path**
A sequence of directory names that specifies the location of the logon script. When a user logs on, the authenticating computer locates the specified logon script (if one has been assigned to that user account) by following that computer`s local logon script path (usually systemroot\System32\Repl\Import\Scripts).

See also: logon script  systemroot  user account

**long distance operator**
Digits dialed before the area or city code to access a long distance phone service. For example, in the United States of America, you dial a 1 before the area or city code: (1) (206) 000-0000. The actual digits depend on the country/region or phone service. Area codes and long distance operators are not used in all countries/regions.

**long name**
A folder name or file name longer than the 8.3 file name standard (up to eight characters followed by a period and an extension of up to three characters) of the file allocation table (FAT) file system. This version of Windows supports file names up to 255 characters and automatically translates long names of files and folders to 8.3 names for MS-DOS and Windows 3.x users. In a Macintosh environment, users can assign names up to 31 characters, excluding colons, to files and folders.

See also: file allocation table (FAT)  MS-DOS (Microsoft Disk Operating System)

**loopback**
An advanced Group Policy setting that is useful in certain closely managed environments, such as laboratories, classrooms, public kiosks, and reception areas. You can find the setting and its description in the Group Policy console tree under Computer Configuration\Administrative Templates\System\Group Policy. The full name of this setting is User Group Policy loopback processing mode.

See also: console tree  Group Policy

**loose name checking**
A form of domain name checking DNS uses that examines characters in DNS names for valid American National Standards Institute (ANSI) character compliance. However, loose name checking does not check for compliance with DNS naming requirements and valid character usage for Internet host names, as specified in RFC 1123, Requirements for Internet Hosts -

Applications and Support. For RFC compliance, DNS domain names will use name labels made up only of valid uppercase and lowercase letters, number characters, and hyphens (A through Z, a through z, 0 through 9, and -) separated by periods.

See also: domain name  Domain Name System (DNS)  label  Request for Comments (RFC)

**LPD**
See other term: Line Printer Daemon (LPD)

**LPR**
See other term: Line Printer Remote (LPR)

**LSASS**
See other term: Local Security Authority Subsystem Service (LSASS)

**luminosity**
The brightness of a color based on a scale from black to white on your monitor.

**LUN**
See other term: logical unit number (LUN)

# Windows Server 2003 Glossary - M

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - M**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**machine-centric**
A service model that requires specific computers and services be used to complete a task. If a user must connect to a specific computer to access a service, then that service is considered machine-centric.

See also: service  service-centric

**Macintosh-accessible volume**
Storage space on the server used for folders and files of Macintosh users. A Macintosh-accessible volume is equivalent to a shared folder for Windows users. Each Macintosh-accessible volume on a computer running Services for Macintosh will correspond to a folder. Both Windows users and Macintosh users can be given access to files located in a folder that is designated as both a shared folder and a Macintosh-accessible volume.

See also: shared folder  volume

**Macintosh-style permissions**
Folder and volume permissions that are similar to the access privileges used on a Macintosh.

See also: permission  privilege  volume

**majority node set server cluster**
A cluster configuration that has two or more nodes and that is configured so that the nodes may or may not be attached to one or more cluster storage devices. The cluster configuration data is stored on multiple disks across the cluster, and the Cluster service makes sure that this data is kept consistent across the different disks. There are advantages and limitations for each cluster configuration (single node server cluster, single quorum device server cluster, and majority node set server cluster).

See also: cluster  cluster storage  node  single node server cluster  single quorum device server cluster

**Make Changes**
The Macintosh-style privilege that allows you to change the contents of folders for which you have this privilege. If you have the Make Changes privilege, you can modify, rename, move, create, and delete files in folders for which you have that privilege. When Services for Macintosh translates Macintosh-style privileges into Windows permissions, you are granted Write and Delete permissions for any folders for which you have the Make Changes privilege.

See also: privilege

**malicious program**
Any program that is created to do intentional harm to or compromise the security of a computer. Examples of malicious programs include trojan horses and computer viruses.

See also: Trojan horse  virus

**malicious user**
(1) A user who intentionally accesses a system with the intent to cause harm to the system or to use it in an unauthorized manner. (2) A person who has legitimate access to a system and poses a security threat to it, such as someone who tries to elevate their user rights to gain access to unauthorized data.

**managed volume**
A local NTFS file system 5.0 volume whose disk space is managed by Remote Storage. Remote Storage frees up disk space by automatically moving infrequently accessed files to a remote storage device.

See also: local storage  NTFS file system  Remote Storage  volume

**Management and Monitoring Tools**

Software components that include utilities for network management and monitoring, along with services that support client dialing and the updating of client phone books. Also included is the Simple Network Management Protocol (SNMP).

See also: Simple Network Management Protocol (SNMP)

**management information base (MIB)**
A set of objects that represent various types of information about a device, used by Simple Network Management Protocol (SNMP) to manage the device. Because different network management services are used for different types of devices and protocols, each service has its own set of objects.

See also: service  Simple Network Management Protocol (SNMP)

**management system**
A network-enabled host running Simple Network Management Protocol (SNMP) management software. This software requests information from SNMP agents. Also called a management console.

See also: host  Simple Network Management Protocol (SNMP)

**mandatory user profile**
A user profile that is not updated when the user logs off. It is downloaded to the user's desktop each time the user logs on, and it is created by an administrator and assigned to one or more users to create consistent or job-specific user profiles. Only members of the Administrators group can change profiles.

See also: roaming user profile  user profile

**man-in-the-middle attack**
A security attack in which an attacker intercepts and possibly modifies data that is transmitted between two users. The attacker pretends to be the other person to each user. In a successful man-in-the-middle attack, the users are unaware that there is an attacker between them, intercepting and modifying their data. Also referred to as a bucket brigade attack.

See also: security

**map**
To translate one value into another. In virtual memory systems, a computer might map a virtual address into a physical address.

See also: virtual memory

**MAPI**
See other term: Messaging API (MAPI)

**MARS**
See other term: multicast address resolution service (MARS)

**master boot record (MBR)**
The first sector on a hard disk, which begins the process of starting the computer. The MBR contains the partition table for the disk and a small amount of executable code called the master boot code.

See also: partition boot sector  Recovery Console

**master file table (MFT)**
An NTFS system file on NTFS-formatted volumes that contains information about each file and folder on the volume. The MFT is the first file on an NTFS volume.

See also: file allocation table (FAT)  NTFS file system  volume

**master server**
In Network Information Service (NIS), a server that provides periodic updates of NIS maps to subordinate servers. A Windows-based master server running Server for NIS can push periodic updates to both UNIX-based and Windows-based subordinate (also known as slave) servers. Master servers running Server for NIS update Server for NIS subordinate servers through an Active Directory database shared with the subordinate server.

**master server**
An authoritative DNS server for a zone. Master servers can vary and are one of two types (either primary or secondary masters), depending on how the server obtains its zone data.

See also: authoritative  DNS server  primary master  zone  zone transfer

**maximize**
To increase the size of a window as much as possible. To maximize a window, double-click the title bar, click the Maximize button (the middle button on the right side of the title bar), or press ALT+SPACEBAR and then press X.

See also: minimize  title bar

**maximum password age**
The period of time a password can be used before the system requires the user to change it.

See also: password

**MBR**
See other term: master boot record (MBR)

**MCS**
See other term: multicast server (MCS)

**MD2**
A hash algorithm that creates a 128-bit hash value and was developed by RSA Data Security, Inc.

See also: hash algorithm

**MD4**
A hash algorithm that creates a 128-bit hash value and was developed by RSA Data Security, Inc.

See also: hash algorithm

**MD5**
An industry-standard one-way, 128-bit hashing scheme, developed by RSA Data Security, Inc., and used by various Point-to-Point Protocol (PPP) vendors for encrypted authentication. A hashing scheme is a method for transforming data (for example, a password) in such a way that the result is unique and cannot be changed back to its original form. The Challenge Handshake Authentication Protocol (CHAP) uses challenge-response with one-way MD5 hashing on the response. In this way, you can prove to the server that you know your password without actually sending the password over the network.

See also: Challenge Handshake Authentication Protocol (CHAP)  hash algorithm  MD5  message digest

**media**
Any fixed or removable objects that store computer data. Examples include hard disks, floppy disks, tapes, and compact discs.

**media access control (MAC) address**
The address that is used for communication between network adapters on the same subnet. Each network adapter has an associated MAC address.

See also: multicast media access control (MAC) address  network adapter  subnet

**media master**
The original set of files that are copied to a tape or disk from local storage. Remote Storage can make up to three copies of the media master on different storage media. Each copy of the media master is called a copy set.

See also: copy set  media  Remote Storage

**media pool**
A logical collection of removable media that have the same management policies. Media pools are used by applications to control access to specific tapes or discs within libraries managed by Removable Storage.

See also: backup media pool  media  Removable Storage

**media sensing**
A feature that, when it detects a network cable connection failure, removes the bound protocols from the failed network adapter until the cable connection is reestablished. Without these bound protocols, the network interface connected through the failed network adapter is not available for cluster communication.

See also: network adapter  protocol

**member**
For DFS Replication, a server that is part of a replication group.

**member server**

A server that is joined to a domain but is not a domain controller. Member servers typically function as file servers, application servers, database servers, Web servers, certificate servers, firewalls, or remote access servers.

See also: domain  domain controller  stand-alone server  Web server

**membership**
Defines the relationship between each replicated folder/member pair. Each membership has a status, either enabled or disabled.

**memory address**
A portion of computer memory that can be allocated to a device or used by a program or the operating system. Devices are usually allocated a range of memory addresses.

See also: device

**Memory Usage**
In Task Manager, the current working set of a process, in kilobytes. The current working set is the number of pages currently resident in memory. On the Task Manager Processes tab, the column heading is Mem Usage.

See also: Task Manager

**Memory Usage Delta**
In Task Manager, the change in memory, in kilobytes, used since the last update.

See also: Task Manager

**message**
For Message Queuing, a unit of information sent between computers running Message Queuing. The message can contain text or binary data as defined by the sending application. All messages, including status messages, are stored in queues on Message Queuing computers.

See also: Message Queuing

**Message Authentication Code (MAC)**
An algorithm that ensures the quality of a block of data.

**message digest**
See other term: hash

**Message Queuing**
A message queuing and routing system for Windows that enables distributed applications running at different times to communicate across heterogeneous networks and with computers that may be offline. Message Queuing provides guaranteed message delivery, efficient routing, security, and priority-based messaging. Formerly known as MSMQ.

See also: Microsoft Distributed Transaction Coordinator (MS DTC)  MSMQ

**Message Queuing resource**
A resource type that can use Message Queuing in a cluster.

See also: cluster  Message Queuing  resource types

**Message Queuing server**
For Message Queuing (formerly known as MSMQ), a computer that can provide message queuing, routing, and directory services to client computers. Message Queuing servers can be used to provide message routing and session concentration for independent clients, provide message routing between sites over routing links, create queues and store messages for dependent clients, and access information in Active Directory (if installed on a Windows domain controller).

See also: Active Directory  dependent client  independent client  MSMQ  routing link  routing services  session concentration

**Messaging API (MAPI)**
The application programming interface (API) for which Active Directory provides support for backward compatibility with Microsoft Exchange applications. New applications should use Active Directory Service Interfaces (ADSI) for accessing Active Directory.

See also: Active Directory  Active Directory Service Interfaces (ADSI)  application programming interface (API)

**Messenger service**
A service that sends and receives messages sent by administrators or by the Alerter service.

See also: Alerter service  performance alert  service

**metadata**
Data about data. For example, the title, subject, author, and size of a file constitute the file`s metadata.

**metric**
A number used to indicate the cost of a route in the Internet Protocol (IP) routing table that enables the selection of the best route among possible multiple routes to the same destination.

See also: Internet Protocol (IP)  routing

**MFT**
See other term: master file table (MFT)

**MIB**
See other term: management information base (MIB)

**Microcom Networking Protocol Five (MNP5)**
A data-compression standard that allows modems to increase throughput by compressing data before transmission. Data can be compressed with a ratio of up to 2:1. MNP5 sometimes expands data that has already been compressed, resulting in poorer performance in those cases. If you have an MNP5 modem, do not turn on modem compression and software compression at the same time. To turn on MNP5 compression, you must also turn on MNP4 error control.

See also: Microcom Networking Protocol Four (MNP4)

**Microcom Networking Protocol Four (MNP4)**
An industry-standard communication protocol that allows modems to automatically retransmit corrupted data, assuring that only error-free data passes through the modem. MNP2 and MNP3 standards are included in MNP4.

See also: Microcom Networking Protocol Five (MNP5)

**Microsoft Digest**
A security support provider (SSP) that implements the Digest Access protocol. Microsoft Digest provides a simple challenge response mechanism for authenticating clients, and it is intended for use by client/server applications that use communications that are based on Hypertext Transfer Protocol (HTTP) or Simple Authentication and Security Layer (SASL).

**Microsoft Distributed Transaction Coordinator (MS DTC)**
A transaction manager that coordinates transactions that span multiple resource managers, such as Message Queuing and Microsoft SQL Server. MS DTC is automatically installed when Message Queuing is installed.

See also: Message Queuing  transaction

**Microsoft Management Console (MMC)**
A framework for hosting administrative tools called snap-ins. A console might contain tools, folders or other containers, World Wide Web pages, and other administrative items. These items are displayed in the left pane of the console, called a console tree. A console has one or more windows that can provide views of the console tree. The main MMC window provides commands and tools for authoring consoles. The authoring features of MMC and the console tree itself might be hidden when a console is in User Mode.

See also: console tree  details pane  snap-in

**Microsoft Point-to-Point Encryption (MPPE)**
A 128-bit key or 40-bit key encryption algorithm using RSA RC4. MPPE provides for packet confidentiality between the remote access client and the remote access or tunnel server, and it is useful where Internet Protocol security (IPSec) is not available. MPPE 40-bit keys are used to satisfy current North American export restrictions. MPPE is compatible with Network Address Translation.

See also: remote access server  tunnel server

**Microsoft Reserved (MSR) partition**
A required partition on every GUID partition table (GPT) disk. System components can allocate portions of the MSR partition into new partitions for their own use. For example, when you convert a basic GPT disk to dynamic, the system allocates a portion of the MSR partition to be used as the Logical Disk Manager (LDM) metadata partition. The MSR partition varies in size based on the size of the GPT disk. For disks smaller than 16 GB, the MSR partition is 32 MB. For disks larger than 16 GB, the

MSR partition is 128 MB. The MSR partition is not visible in Disk Management, and you cannot store data on the MSR partition or delete it.

See also: GUID partition table (GPT)  partition

**MIDI**
See other term: Musical Instrument Digital Interface (MIDI)

**MIDI setup**
The setup that specifies the type of Musical Instrument Digital Interface (MIDI) device you are using, the channel settings needed to play MIDI files, and the port your device is using.

See also: Musical Instrument Digital Interface (MIDI)

**minimize**
To reduce a window to a button on the taskbar. To minimize a window, click the Minimize button (leftmost button on the right side of the title bar) or press ALT+SPACEBAR and then press N.

See also: maximize  title bar

**minimum password length**
The fewest characters a password can contain.

See also: password

**minimum TTL**
In DNS, a default Time to Live (TTL) value that is set in seconds and used with all resource records in a zone. This value is set in the start of authority (SOA) resource record for each zone. By default, the DNS server includes this value in query responses. It is used to inform recipients how long they can store and use resource records, which are provided in the query answer, before they must expire the stored records data. When TTL values are set for individual resource records, those values override the minimum TTL.

See also: DNS server  Domain Name System (DNS)  resource record (RR)  start-of-authority (SOA) resource record  Time to Live (TTL)  zone

**mirror**
One of the two volumes that make up a mirrored volume. Each mirror of a mirrored volume resides on a different disk. If one mirror becomes unavailable (due to a disk failure, for example), Windows can use the remaining mirror to gain access to the volume`s data.

See also: fault tolerance  mirrored volume  volume

**mirrored volume**
A fault-tolerant volume that duplicates data on two physical disks. A mirrored volume provides data redundancy by using two identical volumes, which are called mirrors, to duplicate the information contained on the volume. A mirror is always located on a different disk. If one of the physical disks fails, the data on the failed disk becomes unavailable, but the system continues to operate in the mirror on the remaining disk. You can create mirrored volumes only on dynamic disks on computers running the Windows 2000 Server or Windows Server 2003 families of operating systems. You cannot extend mirrored volumes.

See also: dynamic disk  dynamic volume  fault tolerance  RAID-5 volume  volume

**mixed mode**
In a Windows 2000 domain, the default domain mode setting. Mixed mode enables Windows NT-based backup domain controllers to coexist with Windows 2000-based domain controllers. Mixed mode does not support universal groups or the nesting of groups. You can change the domain mode setting to native mode when all Windows NT-based domain controllers are removed from a domain. In Windows Server 2003 domains, mixed mode is referred to as "Windows 2000 mixed," and it is one of three domain functional levels available.

See also: Active Directory  backup domain controller (BDC)  domain functionality  native mode  universal group

**MMC**
See other term: Microsoft Management Console (MMC)

**MNP4**
See other term: Microcom Networking Protocol Four (MNP4)

**MNP5**
See other term: Microcom Networking Protocol Five (MNP5)

**modem (modulator/demodulator)**
A device that enables computer information to be transmitted and received over a telephone line. The transmitting modem translates digital computer data into analog signals that can be carried over a telephone line. The receiving modem translates the analog signals back to digital form.

See also: modem compression  null modem cable  port  Telephony API (TAPI)  Waiting for Call

**modem compression**
A technique used to reduce the number of characters transmitted without losing data content. The transmitting modem compresses the data and the receiving computer or modem decompresses the data back to its original state.

See also: modem (modulator/demodulator)

**modulation standards**
Protocols that determine how modems convert digital data into analog signals that can be transmitted over telephone lines. Initially, Bell created modulation standards used in the United States, and the CCITT created international recommendations. The ITU-T (formerly called the CCITT) now makes recommendations generally adopted by modem manufacturers both internationally and in the United States. The ITU-T V series recommendations (such as V.34 and V.90) define data communication over the telephone network. The suffixes -bis and -ter (for example, V.32bis) indicate later versions.

See also: International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T)  modem (modulator/demodulator)  protocol  V.34  V.90

**mount**
To place a removable tape or disc into a drive.

See also: dismount  library  mounted drive

**mounted drive**
A drive attached to an empty folder on an NTFS volume. Mounted drives function the same as any other drive, but are assigned a label or name instead of a drive letter. The mounted drive`s name is resolved to a full file system path instead of just a drive letter. Members of the Administrators group can use Disk Management to create mounted drives or reassign drive letters.

See also: drive  mount  NTFS file system  volume

**MPIO**
See other term: Multipath I/O (MPIO)

**MPPE**
See other term: Microsoft Point-to-Point Encryption (MPPE)

**MPR**
See other term: MultiProtocol Routing (MPR)

**MS DTC**
See other term: Microsoft Distributed Transaction Coordinator (MS DTC)

**MS-DOS (Microsoft Disk Operating System)**
A Microsoft computer operating system with a command-line interface that is used to control many internal computer functions, such as running programs and organizing and maintaining files.

See also: Command Prompt window  MS-DOS-based program

**MS-DOS-based program**
A program that is designed to run with MS-DOS and therefore may not be able to take full advantage of all Windows features.

See also: MS-DOS (Microsoft Disk Operating System)  program information file (PIF)

**MSMQ**
The message queuing and routing system for Windows NT 4.0, Windows 95, and Windows 98. For Windows 2000, Windows XP, and the Windows Server 2003 family, this feature is called Message Queuing.

See also: Message Queuing

**multibyte**
A character set that may consist of both one-byte and two-byte characters. A multibyte-character string may contain a mixture of single-byte and double-byte characters. Windows Server 2003 DNS uses the Unicode Transformation Format 8 (UTF-8) encoding scheme described in RFC 2044 to interpret and transform multibyte characters into single-byte characters of 8-bit length.

See also: Domain Name System (DNS)  UCS Transformation Format 8 (UTF-8)

**multicast address resolution service (MARS)**
A service for resolving multicast IP addresses to the asynchronous transfer mode (ATM) addresses of the clients that have joined that multicast group. MARS can work in conjunction with the multicast server (MCS) and clients to distribute multicast data through point-to-multipoint connections.

See also: asynchronous transfer mode (ATM)  IP address  multicast server (MCS)  multicasting

**multicast datagram**
An Internet Protocol (IP) datagram sent to a select multicast group of hosts. Each datagram is sent from a single host source for transmission and delivery to multiple host destinations using the Class D IP address reserved for a specified multicast group.

See also: datagram  host  Internet Protocol (IP)  multicasting

**Multicast Listener Discovery (MLD)**
A protocol used by Internet Protocol version 6 (IPv6) hosts to report their multicast group memberships to any immediately neighboring multicast routers.

See also: multicasting  protocol

**multicast media access control (MAC) address**
A type of media access control address used by multiple, networked computers to receive the same incoming network frames concurrently. Network Load Balancing uses multicast MAC addresses to efficiently distribute incoming network traffic to cluster hosts.

See also: media access control (MAC) address  multicasting  Network Load Balancing

**multicast scope**
A range of multicast group IP addresses in the Class D address range that are available to be leased or assigned to multicast DHCP clients by DHCP.

See also: address class  Dynamic Host Configuration Protocol (DHCP)  IP address  lease  multicasting

**multicast server (MCS)**
A service that manages zero or more multicast groups and distributes multicast data sent to it by clients of those multicast groups through point-to-multipoint connections.

See also: multicasting  nonhubbed mode  service

**multicasting**
The process of sending a message simultaneously to more than one destination on a network.

See also: multicast server (MCS)

**multihomed computer**
A computer that has multiple network adapters or that has been configured with multiple IP addresses for a single network adapter.

See also: IP address  network adapter  virtual IP address

**multilink dialing**
The combination of two or more physical communications links` bandwidth into a single logical link to increase your remote access bandwidth and throughput by using remote access Multilink. Based on the Internet Engineering Task Force (IETF) standard RFC 1990, Multilink combines analog modem paths, Integrated Services Digital Network (ISDN) B-channels, and mixed analog and digital communications links on both your client and server computers. This increases your Internet and intranet access speed and decreases the amount of time you are connected to a remote computer.

See also: bandwidth  B-channel  Integrated Services Digital Network (ISDN)

**multimaster replication**
A replication model in which any domain controller accepts and replicates directory changes. This model differs from single-master replication models, in which one domain controller stores the single modifiable copy of the directory and other domain controllers store backup copies.

See also: domain controller  replication

**multimedia streaming**
Playing audio and video immediately as it is downloaded from a network, rather than storing it in a file first.

**Multipath I/O (MPIO)**
A Microsoft multipathing solution that enables more than one read/write path to a logical unit number (LUN). With MPIO, multipathing is achieved by assigning multiple Fibre Channel host bus adapter ports or multiple iSCSI initiator adapters on a server to the same LUN, and then configuring that server to use MPIO.

**multipathing**
In storage technologies, a solution that increases availability and path failover by providing multiple paths from a server or cluster to a storage subsystem. In the event that the path being used fails, multipathing logic uses an alternate path for I/O so that applications can still access their data.

**multiple boot**
A computer configuration that runs two or more operating systems.

See also: dual boot  startup environment

**MultiProtocol Routing (MPR)**
A protocol that enables routing over Internet Protocol (IP) and Internetwork Packet Exchange (IPX) networks by connecting local area networks (LANs) or by connecting LANs to wide area networks (WANs).

See also: Internet Protocol (IP)  Internetwork Packet Exchange (IPX)  local area network (LAN)  wide area network (WAN)

**Musical Instrument Digital Interface (MIDI)**
A serial interface standard that allows for the connection of music synthesizers, musical instruments, and computers. The MIDI standard is based partly on hardware and partly on a description of the way in which music and sound are encoded and communicated between MIDI devices. The information transmitted between MIDI devices is in a form called a MIDI message, which encodes aspects of sound, such as pitch and volume, as 8-bit bytes of digital information. MIDI devices can be used for creating, recording, and playing back music. Using MIDI, computers, synthesizers, and sequencers can communicate with each other, either keeping time or actually controlling the music created by other connected equipment.

See also: MIDI setup

**My Documents**
A folder that provides you with a convenient place to store documents, graphics, or other files you want to access quickly. When you save a file in a program such as WordPad or Paint, the file is automatically saved in My Documents, unless you choose a different folder.

See also: home folder

# Windows Server 2003 Glossary - N

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - N**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**name**
The identifier of your computer on the network.

**name resolution**
The process of having software translate between names that are easy for users to work with and numerical IP addresses, which are difficult for users but necessary for TCP/IP communications. Name resolution can be provided by software components such as DNS or WINS.

See also: Domain Name System (DNS)  Transmission Control Protocol/Internet Protocol (TCP/IP)  Windows Internet Name Service (WINS)

**name resolution service**
A service, such as that provided by WINS or DNS, that allows friendly names to be resolved to an address, or other specially defined resource data used to locate network resources of various types and purposes.

See also: Domain Name System (DNS)  service  Windows Internet Name Service (WINS)

**name server (NS) resource record**
A resource record used in a zone to designate the DNS domain names for authoritative DNS servers for the zone.

See also: DNS server  Domain Name System (DNS)  resource record (RR)  zone

**named pipe**
A portion of memory that can be used by one process to pass information to another process, so that the output of one is the input of the other. The second process can be local (on the same computer as the first) or remote (on a networked computer).

See also: shared resource

**namespace**
A naming convention that defines a set of unique names for resources in a network. For DNS, a hierarchical naming structure that identifies each network resource and its place in the hierarchy of the namespace. For WINS, a flat naming structure that identifies each network resource using a single, unique name. For DFS Namespaces, a virtual tree of folders that begins with \\ServerOrDomainName\RootName.

See also: Domain Name System (DNS)  Windows Internet Name Service (WINS)

**namespace path**
A Universal Naming Convention (UNC) path to the namespace root (\\ServerOrDomainName\RootName) or to a folder under the namespace root (\\ServerOrDomainName\RootName\Subfolder\Folder). The namespace path of a domain-based namespace begins with a domain name, whereas the namespace path of a stand-alone namespace begins with the name of the server that hosts the root target.

**namespace root**
The top-level folder in the namespace. For example, \\ServerOrDomainName\RootName is the namespace root.

**namespace server**
A server that hosts a namespace. A namespace server can be a domain controller or member server.

**naming context**
See other term: directory partition

**NAS**
See other term: network access server (NAS)

**native mode**
In Windows 2000 domains, the domain mode in which all domain controllers in a domain are running Windows 2000 and a domain administrator has switched the domain operation mode from mixed mode to native mode. Native mode supports universal groups and nesting of groups. In native mode, domain controllers running Windows NT 4.0 or earlier are not supported. In Windows Server 2003 domains, native mode is referred to as "Windows 2000 native," and it is one of three domain functional levels available.

See also: Active Directory  domain functionality  mixed mode  universal group

**NCP**
See other term: NetWare Core Protocol

**NDIS**
See other term: Network Driver Interface Specification (NDIS)

**NDS**
See other term: Novell Directory Services (NDS)

**negative caching**
In DNS, client caching of failed responses to a query. Negative caching improves the response time for successive queries for the same name.

See also: cache  Domain Name System (DNS)

**Net Logon service**
A user-mode service that runs in the Windows security subsystem. The Net Logon service passes the user`s credentials through a secure channel to the domain database and returns the domain security identifiers and user rights for the user. In addition, the Net Logon service performs a variety of other functions related to the user logon process, such as periodic password updates for computer accounts and domain controller discovery.

See also: user mode  user rights

**NetBEUI**
See other term: NetBIOS Extended User Interface (NetBEUI)

**NetBIOS**
See other term: network basic input/output system (NetBIOS)

**NetBIOS Extended User Interface (NetBEUI)**
A network protocol native to Microsoft Networking. It is usually used in small, department-size local area networks (LANs) of 1 to 200 clients. NetBEUI can use Token Ring source routing as its only method of routing. NetBEUI is the Microsoft implementation of the NetBIOS standard.

See also: Compression Control Protocol (CCP)  local area network (LAN)  network basic input/output system (NetBIOS)  protocol  Token Ring

**NetBIOS name**
A 16-byte name of a process using network basic input/output system (NetBIOS). The NetBIOS name is a name that is recognized by WINS, which maps the name to an IP address.

See also: IP address  network basic input/output system (NetBIOS)  Windows Internet Name Service (WINS)

**Netdom**
A tool that allows management of Windows domains and trust relationships from the command line.

See also: domain  trust relationship

**netmask ordering**
A method DNS uses to give ordering and preference to IP addresses on the same network when a requesting client queries for a host name that has multiple host address (A) type resource records. This is designed so that the client program will attempt to connect to a host using the closest (and fastest) IP address available.

See also: address (A) resource record  Domain Name System (DNS)  host  IP address

**NetWare Core Protocol**
The file-sharing protocol that governs communications about resources (such as disk and printer), bindery, and Novell Directory Services (NDS) operations between server and client computers on a Novell NetWare network.

See also: bindery  Internetwork Packet Exchange (IPX)  Novell Directory Services (NDS)  protocol

**network access server (NAS)**
The device that accepts Point-to-Point Protocol (PPP) connections and places clients on the network that the NAS serves.

See also: Point-to-Point Protocol (PPP)

**network adapter**
A device that connects your computer to a network. Sometimes called a network interface card (NIC).

See also: device

**network adapter driver**
A device driver that works directly with the network adapter, acting as an intermediary between the adapter and the protocol driver.

See also: device driver  network adapter

**network address translation (NAT)**
A process that translates the internal Internet Protocol (IP) address of each computer on a local area network (LAN) into an external IP address when that computer connects to the Internet. This process improves the security of the network because internal addresses are not made public.

See also: Internet Protocol (IP)

**network administrator**
A person responsible for planning, configuring, and managing the day-to-day operation of the network. Also called a system administrator.

See also: administrator

**network basic input/output system (NetBIOS)**
An application programming interface (API) that can be used by programs on a local area network (LAN). NetBIOS provides programs with a uniform set of commands for requesting the lower-level services required to manage names, conduct sessions, and send datagrams between nodes on a network.

See also: application programming interface (API)  basic input/output system (BIOS)  datagram  local area network (LAN)  node  service

**network bridge**
A device that connects networks using the same communications protocols so that information can be passed from one network to the other. Also, a device that connects two local area networks (LANs), whether or not they use the same protocols. A network bridge operates at the ISO/OSI data-link layer.

See also: local area network (LAN)  protocol

**Network Connections**
A component you can use to gain access to network resources and functionality, whether you are physically at the network location or in a remote location. By using the Network Connections folder you can create, configure, store, and monitor connections. Formerly called Network and Dial-up Connections or Dial-Up Networking.

**Network Control Protocol**
A protocol within the Point-to-Point Protocol (PPP) suite that negotiates the parameters of an individual local area network (LAN) protocol such as TCP/IP or Internetwork Packet Exchange (IPX).

See also: Internetwork Packet Exchange (IPX)  local area network (LAN)  Point-to-Point Protocol (PPP)  Transmission Control Protocol/Internet Protocol (TCP/IP)

**Network DDE service**
A service that provides network transport and security for dynamic data exchange (DDE) conversations.

See also: dynamic data exchange (DDE)  service

**Network Driver Interface Specification (NDIS)**
A Microsoft/3Com specification establishing a common shared interface for Microsoft operating systems to support protocol-independent transport of multiple network transport protocols (such as TCP/IP, NetBEUI, IPX/SPX, and AppleTalk). NDIS allows more than one transport protocol to be bound and to operate simultaneously over a single network adapter.

See also: Internetwork Packet Exchange (IPX) Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) NetBIOS Extended User Interface (NetBEUI) Transmission Control Protocol/Internet Protocol (TCP/IP)

**Network File System (NFS)**
A service for distributed computing systems that provides a distributed file system, eliminating the need for keeping multiple copies of files on separate computers.

See also: service

**network gateway**
A device that connects networks using different communications protocols so that information can be passed from one network to the other. A network gateway both transfers information and converts it to a form that is compatible with the protocols being used by the receiving network.

See also: device protocol

**network ID**
A number used to identify the systems that are located on the same physical network bounded by routers. The network ID should be unique to the internetwork.

See also: router

**Network Information Service (NIS)**
A service for distributed computing systems that provides a distributed database system for common configuration files.

See also: service

**Network Information Service (NIS)**
A network naming and administration system that was developed by Sun Microsystems and is constructed on the client-server model of computing. Formerly known as Yellow Pages (yp), NIS provides a simple network look-up service that consists of databases and processes.

**Network Load Balancing**
A Windows network component that uses a distributed algorithm to load-balance Internet Protocol (IP) traffic across a number of hosts, helping to enhance the scalability and availability of mission-critical, IP-based services, such as Terminal Services, Web services, virtual private networking, and streaming media. It also provides high availability by detecting host failures and automatically redistributing traffic to the surviving hosts.

See also: availability cluster host scalability

**Network Load Balancing cluster**
Up to 32 Web servers from which Network Load Balancing presents a single IP address to Web clients and among which Network Load Balancing distributes incoming Web requests.

See also: IP address Network Load Balancing Web server

**network media**
The type of physical wiring and lower-layer protocols used for transmitting and receiving packets, for example, Ethernet, Fiber Distributed Data Interface (FDDI), and Token Ring.

See also: packet protocol Token Ring

**Network Name resource**
The name of a device that exists on a network and is supported as a cluster resource by a Resource DLL provided with Windows.

See also: Resource DLL

**Network News Transfer Protocol (NNTP)**
A member of the TCP/IP suite of protocols used to distribute network news messages to NNTP servers and clients (newsreaders) on the Internet. NNTP is designed so that news articles are stored on a server in a central database, thus enabling a user to select specific items to read.

See also: Transmission Control Protocol/Internet Protocol (TCP/IP)

**network number**
In the Macintosh environment, the routing address or range of addresses assigned to the physical network that AppleTalk

Phase 2 routers use to direct information to the appropriate network. Also called network range and cable range.

See also: router

**network partition**
A state in which one or more of the nodes in a cluster cannot communicate with the other cluster nodes.

See also: cluster  node

**network place**
A folder on a Web server. You can view files and folders on Web servers just as you would view files and folders on network servers. However, when you save a file to a network place, the file is saved on a Web server, not on your computer`s hard disk. You can create network places by using the Add Network Place Wizard, which is located in My Network Places. Network places are available only on Web servers that support Web Extender Client (WEC), FrontPage extensions, and Distributed Authoring and Versioning (DAV) protocols.

See also: Web server

**Network Service account**
A predefined local account that is used to start a service and provide the security context for that service. The name of the account is NT AUTHORITY\NetworkService. The Network Service account has limited access to the local computer and authenticated access (as the computer account) to network resources.

See also: access control list (ACL)  Local Service account  Local System account  service

**network/bit-masked identifier**
An IP address that includes a subnet bit-mask. Active Directory uses the subnet bit-mask to recognize sites and establish how replication of directory information will occur.

See also: Active Directory  IP address  replication  site  subnet

**network-to-network interface (NNI)**
An interface that defines how switches in a network communicate with each other.

**NFS**
See other term: Network File System (NFS)

**NFS client**
A computer running client software to access shared Network File System (NFS) resources.

**NFS server**
A computer providing Network File System (NFS) disk resources to NFS clients.

**NIS**
See other term: Network Information Service (NIS)

**NIS**
See other term: Network Information Service (NIS)

**NIS map**
A database served by Network Information Service (NIS). The NIS lookup calls require a map (database) name and an NIS domain name. An NIS domain consists of a collection of such maps.

**nis2ad**
A Windows command-line utility that transfers the entries in the Network Information Service (NIS) map source file specified by mapfile located in the UNIX NIS domain specified by NIS_domain to the Active Directory domain specified by AD_domain.

**nisadmin**
A Windows command-line utility that administers Server for NIS on the local computer or on a remote computer.

**nismap**
A Windows command-line utility that manages the map specified by mapname maintained by Server for NIS on a local computer or a remote computer.

**NNI**
See other term: network-to-network interface (NNI)

**NNTP**
See other term: Network News Transfer Protocol (NNTP)

**No Override**
In Group Policy in the Windows Server 2003 family, to force the Group Policy object (GPO) links of the parent container to take precedence over the GPO links of the child containers. By default, the GPO links that are closest to the user or computer (links to the child container) have higher precedence than GPO links higher up (links to the parent container). Enforcement causes GPOs that are linked in parent containers to take precedence instead. In Group Policy Management console (GPMC), Enforce is used instead of No Override.

See also: block inheritance  enforce  Group Policy  Group Policy Management console (GPMC)  Group Policy object link  inheritance  precedence

**node**
For tree structures, a location on the tree that can have links to one or more items below it. For local area networks (LANs), a device that is connected to the network and is capable of communicating with other network devices. For server clusters, a computer system that is an active or inactive member of a cluster.

See also: local area network (LAN)  server cluster

**nonauthoritative restore**
A restore operation performed on an Active Directory domain controller in which the objects in the restored directory are not treated as authoritative. The restored objects are updated with changes held on other domain controllers in the domain.

See also: Active Directory  authoritative restore

**noncontainer object**
An object that cannot logically contain other objects. For example, a file is a noncontainer object.

See also: container object  object

**noncontiguous namespace**
A namespace based on different DNS root domain names, such as that of multiple trees in the same forest.

See also: domain name  Domain Name System (DNS)  forest  namespace

**nonhubbed mode**
A mode in which asynchronous transfer mode (ATM) Address Resolution Protocol/multicast address resolution service (ARP/MARS) does not forward multicast and broadcast traffic for multicast group clients. In this mode, the service returns a dynamic listing of ATM hosts currently registered for the multicast group address to requesting clients. Clients then use this list to initiate and establish their own point-to-multipoint virtual connections with each of the members in the multicast list.

See also: Address Resolution Protocol (ARP)  asynchronous transfer mode (ATM)  hubbed mode  multicast address resolution service (MARS)  multicast server (MCS)

**nonpaged memory**
Memory that cannot be paged to disk. Paging is the moving of infrequently used parts of a program`s working memory from random access memory (RAM) to another storage medium, usually the hard disk.

See also: paging file  random access memory (RAM)

**nonpaged pool**
Operating system memory that is never paged to disk. Paging is the moving of infrequently used parts of a program`s working memory from RAM to another storage medium, usually the hard disk. In Task Manager, the amount of memory used by a process, in kilobytes.

See also: paging file  Task Manager

**nonrepudiation**
The use of cryptography to ensure that someone performing an action on a computer cannot falsely deny that they performed that action. A system with nonrepudiation provides undeniable proof that a user took a specific action such as transferring money, authorizing a purchase, or sending a message.

See also: cryptography

**nontransitive trust**
A trust relationship in a multiple-domain environment that is restricted to just two domains. For example, if domain A has a

nontransitive trust with domain B, and domain B trusts domain C, then there is no trust relationship between domain A and domain C. Nontransitive trusts can be one-way or two-way.

See also: forest  nontransitive trust  one-way trust  trust relationship  two-way trust

**non-Windows**
Operating systems that are not Windows-brand operating systems.

**normal backup**
A backup that copies all selected files and marks each file as having been backed up (in other words, the archive attribute is cleared). With normal backups, you need only the most recent copy of the backup file or tape to restore all of the files. You usually perform a normal backup the first time you create a backup set.

See also: copy backup  daily backup  differential backup  incremental backup

**notification area**
The area on the taskbar to the right of the taskbar buttons. The notification area displays the time and can also contain shortcuts that provide quick access to programs, such as Volume Control and Power Options. Other shortcuts can appear temporarily, providing information about the status of activities. For example, the printer shortcut icon appears after a document has been sent to the printer and disappears when printing is complete.

See also: taskbar

**notify list**
A list maintained by the primary master for a zone of other DNS servers that should be notified when zone changes occur. The notify list is made up of IP addresses for DNS servers configured as secondary masters for the zone. When the listed servers are notified of a change to the zone, they will initiate a zone transfer with another DNS server and update the zone.

See also: DNS server  IP address  primary master  zone  zone transfer

**Novell Directory Services (NDS)**
On networks running Novell NetWare 4.x and higher, a distributed database that maintains information about every resource on the network and that provides access to these resources.

See also: resource

**Nslookup**
A command-line tool used to diagnose DNS infrastructure.

See also: Domain Name System (DNS)

**NTFS file system**
An advanced file system that provides performance, security, reliability, and advanced features that are not found in any version of file allocation table (FAT). For example, NTFS guarantees volume consistency by using standard transaction logging and recovery techniques. If a system fails, NTFS uses its log file and checkpoint information to restore the consistency of the file system. NTFS also provides advanced features, such as file and folder permissions, encryption, disk quotas, and compression.

See also: file allocation table (FAT)  file system

**NTFS permissions**
The settings that administrators apply to access control entries (ACEs) for managing access to files and folders under the NTFS file system. Take Ownership is an example of an NTFS permission.

See also: access control entry (ACE)  access control list (ACL)  NTFS file system  permission

**NTGATEWAY**
A group of NetWare servers. Accounts used to create a gateway must belong to this group.

See also: gateway

**NTLM authentication protocol**
A challenge/response authentication protocol. The NTLM authentication protocol is supported in Windows 2000, Windows XP, and the Windows Server 2003 family, but it is not the default. It was the default authentication protocol for earlier versions of Windows.

See also: authentication  protocol

**null modem cable**
Special cabling that eliminates the modem`s need for asynchronous communications between two computers over short distances. A null modem cable emulates modem communication.

See also: modem (modulator/demodulator)

**null password**
See other term: blank password

**NWLink**
See other term: NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)

**NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)**
The Microsoft implementation of the Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) protocol used on NetWare networks. NWLink allows connectivity between Windows-based computers and NetWare networks running IPX/SPX. NWLink also provides network basic input/output system (NetBIOS) functionality and the Routing Information Protocol (RIP).

See also: Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) NetWare network basic input/output system (NetBIOS) Routing Information Protocol over IPX (RIPX)

# Windows Server 2003 Glossary - O

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - O**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**object**
An entity, such as a file, folder, shared folder, printer, or Active Directory object, described by a distinct, named set of attributes. For example, the attributes of a File object include its name, location, and size; the attributes of an Active Directory User object might include the user's first name, last name, and e-mail address. For OLE and ActiveX, an object can also be any piece of information that can be linked to, or embedded into, another object.

See also: Active Directory attribute attribute child object OLE parent object

**object access event category**
In auditing, a group of events that are logged when an object, such as a file or folder, is accessed by means of an action for which auditing has been enabled. For example, object access events might be logged when a file is opened, modified, or deleted.

See also: auditing event object

**object class**
A distinct, named set of attributes that represents a specific type of entity stored in the directory, such as users, printers, or applications. The attributes include data describing the thing that is identified by the directory object. Attributes of a user might include the user`s first name, last name, and e-mail address.

See also: directory

**object identifier**
A number that identifies an object class or attribute. Object identifiers (OIDs) are organized into an industry-wide global hierarchy. An object identifier is represented as a dotted decimal string, such as 1.2.3.4, with each dot representing a new branch in the hierarchy. National registration authorities issue root object identifiers to individuals or organizations, who manage the hierarchy below their root object identifier.

See also: attribute object class

**ODBC**
See other term: Open Database Connectivity (ODBC)

**offline**
A state that marks a component in a cluster as unavailable. A node in an offline state is either inactive or not running. Resources and groups also have an offline state.

See also: group node online pending resource

**offline mode**
For Message Queuing, a condition in which a computer that belongs to a domain is temporarily unable to communicate with a domain controller. This occurs when the computer itself is offline, all domain controllers in the site are offline, or when an attempt is made to access a remote computer and the remote computer is temporarily unable to query a domain controller for authentication.

See also: Message Queuing

**OID**
See other term: object identifier

**OLE**
A way to transfer and share information between applications by pasting information created in one application into a document created in another application, such as a spreadsheet or word processing file.

See also: embedded object  linked object  package

**on-disk catalog**
Information stored on a local disk drive. The on-disk catalog contains a list of files and folders that have been backed up in a backup set.

See also: backup set  on-media catalog

**one-way trust**
A trust relationship between two domains in which only one of the two domains trusts the other domain. For example, domain A trusts domain B, and domain B does not trust domain A. One-way trusts are often used to enable authenticated access to resource domains.

See also: authentication  domain  trust relationship  two-way trust

**online**
A state that marks a component in a cluster as available. When a node is online, it is an active member of the cluster and can own and run groups as well as honor cluster database updates, contribute votes to the quorum algorithm, and maintain heartbeats. Resources and groups also have an online state.

See also: cluster  group  heartbeat  node  offline  pending  resource

**on-media catalog**
Information stored on backup storage media. The on-media catalog contains a list of files and folders that have been backed up in a backup set.

See also: backup set  on-disk catalog

**Open Database Connectivity (ODBC)**
An application programming interface (API) that enables database applications to access data from a variety of existing data sources.

See also: application programming interface (API)

**Open Shortest Path First (OSPF)**
A routing protocol used in medium-sized and large networks. This protocol is more complex than Routing Information Protocol (RIP), but it allows better control and is more efficient in propagation of routing information.

See also: protocol  routing  Routing Information Protocol (RIP)

**Open Systems Interconnection (OSI) reference model**
A networking model introduced by the International Organization for Standardization (ISO) to promote multi-vendor interoperability. OSI is a seven-layered conceptual model consisting of the application, presentation, session, transport, network, data-link, and physical layers.

See also: Transmission Control Protocol/Internet Protocol (TCP/IP)

**OpenType font**
A type of computer font that can be rotated or scaled to any size. OpenType fonts are clear and readable in all sizes and can be sent to any printer or other output device that is supported by Windows.

See also: font  TrueType font

**operations master**
A domain controller that has been assigned one or more special roles in an Active Directory domain. The domain controllers assigned these roles perform operations that are single-master (not permitted to occur at different places on the network at the same time). Examples of these operations include resource identifier allocation, schema modification, PDC emulation, adding and removing domains to and from the forest, and tracking changes to security principals across all domains in a forest.

See also: Active Directory  domain controller  domain naming master  infrastructure master  PDC emulator master  schema master

**operator**
In mathematics and in programming and computer applications, a symbol or other character indicating an operation that acts on one or more elements. You can use the following four operators in standard calculations: / (divide); * (multiply); - (subtract); and + (add). For Indexing Service, a word or character that specifies a relationship in a query.

See also: Indexing Service

**ordering method**
The order in which targets are listed in a referral. Targets in a client's site are always listed first in a referral (unless a target priority setting overrides this). Targets outside of a client's site can be listed in order of lowest cost or in random order, or they can be excluded altogether from the referral.

**organization claims**
Claims in intermediate or normalized form within an organization's namespace.

**organizational unit**
An Active Directory container object used within domains. An organizational unit is a logical container into which users, groups, computers, and other organizational units are placed. It can contain objects only from its parent domain. An organizational unit is the smallest scope to which a Group Policy object (GPO) can be linked, or over which administrative authority can be delegated.

See also: Active Directory  container object  Group Policy object (GPO)  parent domain

**original equipment manufacturer (OEM)**
A company that typically purchases computer components from other manufacturers, uses the components to build a personal computer, preinstalls Windows onto that computer, and then sells the computer to the public.

**orphan**
A member of a mirrored volume or a RAID-5 volume that has failed due to a severe cause, such as a loss of power or a complete hard-disk head failure. When this happens, the fault-tolerant driver determines that it can no longer use the orphaned member and directs all new reads and writes to the remaining members of the fault-tolerant volume.

See also: fault tolerance  mirrored volume  RAID-5 volume

**orphan file**
A file that is stored inside My Briefcase and not linked to any file outside My Briefcase. When you update files, the orphan file is not synchronized with any other file.

**OS/2**
A protected-mode, virtual memory, multitasking operating system for personal computers based on the Intel 80286, 80386, i486, and Pentium processors. OS/2 can run most MS-DOS-based programs and can read all MS-DOS disks.

See also: MS-DOS-based program

**OSChooser**
The client-side wizard that walks a user through the installation of an operating system or provides access to maintenance and troubleshooting utilities.

See also: client

**OSI**
See other term: Open Systems Interconnection (OSI) reference model

**OSPF**
See other term: Open Shortest Path First (OSPF)

**out-of-band connection**
A connection between two computers that relies on a nonstandard network connection, such as a serial port connection, and nonstandard remote administration tools, such as Special Administration Console (SAC). An out-of-band connection is usually used only when a remote computer cannot access a network or is not in a functional state because of hardware or software failure.

See also: in-band connection  remote computer  serial port

**owner**
In a Windows environment, the person who controls how permissions are set on objects and can grant permissions to others. In the Macintosh environment, an owner is the user responsible for setting permissions for a folder on a server. A Macintosh user who creates a folder on the server automatically becomes the owner of the folder, and can then transfer ownership to someone else. Each Macintosh-accessible volume on the server also has an owner.

See also: permission

**owner category**
In the Macintosh environment, the user category to which you assign permissions for the owner of a folder or a Macintosh volume.

See also: permission  volume

# Windows Server 2003 Glossary - P

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - P**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**package**
An icon that represents embedded or linked information. That information may consist of a complete file, such as a Paint bitmap, or part of a file, such as a spreadsheet cell. When you choose the package, the application used to create the object either plays the object (for example, a sound file) or opens and displays the object. If you change the original information, linked information is automatically updated. However, you must manually update embedded information.

See also: embedded object  linked object  OLE

**packet**
An Open Systems Interconnection (OSI) network layer transmission unit that consists of binary information representing both data and a header containing an identification number, source and destination addresses, and error-control data.

See also: Open Systems Interconnection (OSI) reference model

**packet filtering**
Prevents certain types of network packets from either being sent or received. This can be employed for security reasons (to prevent access from unauthorized users) or to improve performance by disallowing unnecessary packets from going over a slow connection.

See also: packet

**packet header**
In network protocol communications, a specially reserved field of a defined bit length that is attached to the front of a packet for carry and transfer of control information. When the packet arrives at its destination, the field is then detached and discarded as the packet is processed and disassembled in a corresponding reverse order for each protocol layer.

See also: packet

**packet switching**
A technology for breaking data into packets and then sending the packets over a network. Each packet has a header containing its source and destination, a sequence number to reassemble the information, a block of data content, and an error-checking code. The data packets may take different routes to their destination, where the original information is reassembled after the packets arrive. The international standard for packet switching networks is X.25.

See also: packet

**PAD (packet assembler/disassembler)**
A device that connects a non-X.25 device such as a modem to an X.25 packet switching network.

See also: packet

**PAE**
See other term: Physical Address Extension (PAE)

**page**
In virtual memory systems, a unit of data storage that is brought into random access memory (RAM), typically from a hard drive, when a requested item of data is not already in RAM.

See also: virtual memory

**page fault**
The interrupt that occurs when software attempts to read from or write to a virtual memory location that is marked not present. In Task Manager, page fault is the number of times data has to be retrieved from disk for a process because it was not

found in memory. The page fault value accumulates from the time the process started.

See also: Page Faults Delta  Task Manager  virtual memory

**Page Faults Delta**
In Task Manager, the change in the number of page faults since the last update.

See also: page fault  Task Manager

**paged pool**
The system-allocated virtual memory that has been charged to a process and that can be paged. Paging is the moving of infrequently used parts of a program`s working memory from random access memory (RAM) to another storage medium, usually the hard disk. In Task Manager, the amount of system-allocated virtual memory, in kilobytes, used by a process.

See also: registry size limit (RSL)  Task Manager  virtual memory

**page-description language (PDL)**
A computer language that describes the arrangement of text and graphics on a printed page.

See also: PostScript  PostScript fonts  Printer Control Language (PCL)  Printer Job Language (PJL)

**paging file**
A hidden file or files on the hard disk that Windows uses to hold parts of programs and data files that do not fit in memory. The paging file and physical memory, or random access memory (RAM), comprise virtual memory. Windows moves data from the paging file to memory as needed and moves data from memory to the paging file to make room for new data. Also known as a swap file.

See also: Peak Memory Usage  registry size limit (RSL)  virtual memory

**PAP**
See other term: Password Authentication Protocol (PAP)

**paper source**
The location (such as Upper Paper Tray or Envelope Feeder) of the paper at the printer.

See also: printer

**parallel port**
The input/output connector for a parallel interface device. Printers are generally plugged into a parallel port.

See also: port  serial port

**parent domain**
For DNS and Active Directory, domains that are located in the namespace tree directly above other derivative domain names (child domains). For example, microsoft.com would be the parent domain for example.microsoft.com, a child domain.

See also: Active Directory  child domain  domain  Domain Name System (DNS)

**parent object**
An object in which another object resides. For example, a folder is a parent object in which a file, or child object, resides. An object can be both a parent and a child object. For example, a subfolder that contains files is both the child of the parent folder and the parent folder of the files.

See also: child object  object

**parent-child trust**
A trust that is automatically established when a new domain (the child domain) is added, or becomes subordinate, to an existing domain (the parent domain). Parent-child trusts are transitive and two-way.

See also: child domain  parent domain  transitive trust  two-way trust

**parity**
A calculated value that is used to reconstruct data after a failure. RAID-5 volumes stripe data and parity intermittently across a set of disks. When a disk fails, some server operating systems use the parity information together with the data on good disks to recreate the data on the failed disk.

See also: fault tolerance  RAID-5 volume  striped volume

**parity bit**
In asynchronous communications, an extra bit used in checking for errors in groups of data bits transferred within or between computer systems. In modem-to-modem communications, a parity bit is often used to check the accuracy with which each character is transmitted.

See also: parity

**partition**
A portion of a physical disk that functions as though it were a physically separate disk. After you create a partition, you must format it and assign it a drive letter before you can store data on it. On basic disks, partitions are known as basic volumes, which include primary partitions and logical drives. On dynamic disks, partitions are known as dynamic volumes, which include simple, striped, spanned, mirrored, and RAID-5 volumes.

See also: basic disk  basic volume  drive letter  dynamic volume  extended partition  primary partition  system partition

**partition boot sector**
A portion of a hard disk partition that contains information about the disk`s file system and a short machine language program that loads the Windows operating system.

See also: partition

**passive client**
A Hypertext Transfer Protocol (HTTP) browser, capable of broadly supported HTTP, that can make use of cookies. Active Directory Federation Services (ADFS) in Windows Server 2003 R2 supports only passive clients, and it adheres to the WS-Federation Passive Requestor Profile (WS-F PRP) specification.

**passport authentication**
A user authentication mechanism that uses standard Web technologies to enable single sign-on access to multiple Web resources.

See also: .NET Passport  authentication

**password**
A security measure used to restrict logon names to user accounts and access to computer systems and resources. A password is a string of characters that must be provided before a logon name or an access is authorized. A password can be made up of letters, numbers, and symbols, and it is case-sensitive.

See also: maximum password age  minimum password length  user account

**Password Authentication Protocol (PAP)**
A simple, plaintext authentication scheme for authenticating Point-to-Point Protocol (PPP) connections. The user name and password are requested by the remote access server and returned by the remote access client in plaintext.

See also: authentication  password  plaintext  Point-to-Point Protocol (PPP)  remote access server

**password policy**
A collection of policy settings that define the password requirements for a Group Policy object (GPO).

See also: Group Policy object (GPO)  password

**Password Synchronization**
A feature of Windows that helps integrate Windows and UNIX networks by simplifying the process of maintaining secure passwords in both environments.

**path**
A sequence of directory (or folder) names that specifies the location of a directory, file, or folder within the Windows directory tree. Each directory name and file name within the path must be preceded by a backslash (\). For example, to specify the path of a file named Readme.doc located in the Windows directory on drive C, type C:\Windows\Readme.doc.

See also: directory

**paused**
A state that applies to a node in a cluster. The node is a fully active member in the cluster but cannot accept new resource groups (for example, a resource group cannot fail over or fail back to a paused node). You can administer and maintain a paused node.

See also: cluster  failback  failover  node  offline

**PC Card**
A removable device, approximately the size of a credit card, that can be plugged into a Personal Computer Memory Card International Association (PCMCIA) slot in a portable computer. PCMCIA devices can include modems, network cards, and hard disk drives.

See also: modem (modulator/demodulator)

**PCL**
See other term: Printer Control Language (PCL)

**PDC**
See other term: primary domain controller (PDC)

**PDC emulator master**
A domain controller that holds the PDC emulator operations master role in Active Directory. The PDC emulator services network clients that do not have Active Directory client software installed, and it replicates directory changes to any Windows NT backup domain controllers (BDCs) in the domain. The PDC emulator handles password authentication requests involving passwords that have recently changed and not yet been replicated throughout the domain. At any time, the PDC emulator master role can be assigned to only one domain controller in each domain.

See also: Active Directory  backup domain controller (BDC)  domain controller  operations master  primary domain controller (PDC)

**PDL**
See other term: page-description language (PDL)

**Peak Memory Usage**
In Task Manager, the peak amount of physical memory resident in a process since it started.

See also: Task Manager

**pending**
A state that refers to a resource in a cluster when the resource is in the process of being brought online or taken offline.

See also: offline  online  resource

**Per Device or Per User Licensing**
A licensing mode that requires a separate device-based or user-based Client Access License (CAL) for each device or user that accesses a server running Windows Server 2003, regardless of the number of concurrent connections.

See also: client  device  Per Server Licensing

**Per Server Licensing**
A licensing mode that requires a separate device-based or user-based Client Access License (CAL) for each concurrent connection to a server running Windows Server 2003.

See also: client  device  Per Device or Per User Licensing

**performance alert**
A feature that detects when a predefined counter value rises above or falls below the configured threshold and notifies a user by means of the Messenger service.

See also: Messenger service

**performance counter**
In System Monitor, a data item that is associated with a performance object. For each counter selected, System Monitor presents a value corresponding to a particular aspect of the performance that is defined for the performance object.

See also: performance object

**performance object**
In System Monitor, a logical collection of counters that is associated with a resource or service that can be monitored.

See also: performance counter

**performance object instance**
In System Monitor, a term used to distinguish between multiple performance objects of the same type on a computer.

See also: performance object

**perimeter network**
An Internet Protocol (IP) network segment that contains resources, such as Web servers and virtual private network (VPN) servers, that are available to Internet users. Also known as screened subnet.

See also: Internet Authentication Service (IAS)  Internet Protocol (IP)  virtual private network (VPN)

**permanent virtual circuit (PVC)**
A virtual circuit assigned to a preconfigured static route.

See also: static routes

**permission**
A rule associated with an object to regulate which users can gain access to the object and in what manner. Permissions are assigned or denied by the object`s owner.

See also: object  privilege

**personal identification number (PIN)**
A secret identification code that is used to protect smart cards from misuse. The PIN is similar to a password and is known only to the owner of the card. The smart card can be used only by someone who possesses the smart card and knows the PIN.

See also: smart card

**Physical Address Extension (PAE)**
A feature of x86-based servers that allows computers running Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, to support more than 4 gigabytes (GB) of physical memory. Physical Address Extension (PAE) allows up to 64 GB of physical memory to be used as regular 4-kilobyte (KB) pages and expands the number of bits that can be used by the kernel to address physical memory from 32 to 36.

See also: x86

**Physical Disk resource**
A disk on a cluster storage device. Physical disks are supported as cluster resources by a Resource DLL.

See also: Resource DLL

**physical security**
The protection of critical data repositories and paths through the use of physical barriers. Physical barriers range in effectiveness and complexity from locked doors to multilayer access control systems involving proof of identity and explicit access permissions.

See also: security

**PID**
See other term: process identifier (PID)

**PIF**
See other term: program information file (PIF)

**ping**
A utility that verifies connections to one or more remote hosts. The ping command uses Internet Control Message Protocol (ICMP) echo request and echo reply packets to determine whether a particular Internet Protocol (IP) system on a network is functional. Ping is useful for diagnosing IP network or router failures.

See also: host  Internet Control Message Protocol (ICMP)  Internet Protocol (IP)  packet

**pixel**
Short for picture element, one spot in a rectilinear grid of thousands of such spots that form an image produced on the screen by a computer or on paper by a printer. A pixel is the smallest element that display or print hardware and software can manipulate to create letters, numbers, or graphics. Also called a pel.

See also: screen resolution

**PJL**
See other term: Printer Job Language (PJL)

**PKCS #10**
The Certification Request Syntax Standard, developed and maintained by RSA Data Security, Inc.

See also: Public Key Cryptography Standards (PKCS)  Rivest-Shamir-Adleman (RSA) cryptographic algorithms

**PKCS #12**
Personal Information Exchange Syntax Standard, developed and maintained by RSA Data Security, Inc. This standard primarily provides a format for storing keys or a certificate in a file.

See also: Public Key Cryptography Standards (PKCS)  Rivest-Shamir-Adleman (RSA) cryptographic algorithms

**PKCS #7**
The Cryptographic Message Syntax Standard. It is a general syntax, developed and maintained by RSA Data Security, Inc., for data to which cryptography may be applied, such as digital signatures and encryption. It also provides a syntax for disseminating certificates or certificate revocation lists.

See also: certificate revocation list (CRL)  encryption  Public Key Cryptography Standards (PKCS)

**PKI**
See other term: public key infrastructure (PKI)

**plaintext**
Data in its unencrypted or decrypted form. Sometimes also called cleartext.

See also: encryption

**plotter font**
A font created by a series of dots connected by lines. Plotter fonts can be scaled to any size and are most often printed on plotters. Some dot-matrix printers also support plotter fonts.

See also: font

**Plug and Play**
A set of specifications developed by Intel Corporation that enables a computer to detect and configure a device automatically and install the appropriate device drivers.

See also: universal serial bus (USB)

**PNNI**
See other term: private or public network-to-network interface (PNNI)

**p-node**
A network basic input/output system (NetBIOS) implementation that uses point-to-point communication with a WINS server to resolve names as IP addresses.

See also: IP address  network basic input/output system (NetBIOS)  Windows Internet Name Service (WINS)

**point of presence (POP)**
The local access point for a network provider. Each POP provides a telephone number that allows users to make a local call for access to online services.

**point of termination station (POTS)**
Basic dial telephone connections to the public switched network, without any added features or functions. Also called plain old telephone service.

**pointer (PTR) resource record**
A DNS resource record used in a reverse lookup zone to map an IP address to a DNS name.

See also: Domain Name System (DNS)  in-addr.arpa domain  IP address  resource record (RR)  reverse lookup  zone

**Point-to-Point Protocol (PPP)**
An industry standard suite of protocols for the use of point-to-point links to transport multiprotocol datagrams. PPP is documented in RFC 1661.

See also: Compression Control Protocol (CCP)  remote access  Request for Comments (RFC)  Transmission Control Protocol/Internet Protocol (TCP/IP)  voluntary tunnel

**Point-to-Point Tunneling Protocol (PPTP)**

Networking technology that supports multiprotocol virtual private networks (VPNs), enabling remote users to access corporate networks securely across the Internet or other networks by dialing into an Internet service provider (ISP) or by connecting directly to the Internet. PPTP tunnels, or encapsulates, Internet Protocol (IP) or Internetwork Packet Exchange (IPX) traffic inside IP packets. This means that users can remotely run applications that depend on particular network protocols. PPTP is described in RFC 2637.

See also: Internet Protocol (IP) Internetwork Packet Exchange (IPX) NetBIOS Extended User Interface (NetBEUI) packet Request for Comments (RFC) tunnel virtual private network (VPN)

**poison reverse**
A process that, used with split horizon, improves Routing Information Protocol (RIP) convergence over simple split horizon by advertising all network IDs. However, the network IDs learned in a given direction are advertised with a hop count of 16, indicating that the network is unavailable.

See also: hop count Routing Information Protocol (RIP) split horizon

**policy**
The mechanism by which computer settings are configured automatically, as defined by the administrator. Depending on context, this can refer to Group Policy or Windows NT 4.0 System Policy.

See also: Group Policy Group Policy object (GPO) System Policy

**polling intervals**
The frequency the Resource Monitor checks that the resource is available and operating. There are two levels of polling: Looks Alive and Is Alive. The server cluster requests a more thorough check of the resource`s state at each Is Alive interval than it does at each Looks Alive interval; therefore, the Is Alive polling interval is typically longer than the Looks Alive polling interval. You can specify the two polling intervals and a time-out value for resources.

See also: resource Resource Monitor

**POP**
See other term: point of presence (POP)

**POP3 (Post Office Protocol 3)**
A popular protocol used for receiving e-mail messages. This protocol is often used by Internet service providers (ISPs). POP3 servers allow access to a single Inbox in contrast to Internet Message Access Protocol (IMAP) servers, which provide access to multiple server-side folders.

See also: Internet service provider (ISP) Simple Mail Transfer Protocol (SMTP)

**port**
A connection point on your computer where you can connect devices that pass data into and out of a computer. For example, a printer is typically connected to a parallel port (also called an LPT port), and a modem is typically connected to a serial port (also called a COM port).

See also: parallel port serial port universal serial bus (USB)

**port based network access control**
A form of network access control that uses the physical characteristics of a switched local area network (LAN) infrastructure to authenticate devices that are attached to a LAN port and to prevent access to that port if the authentication process fails. This form of network access control is used in 802.1X authentication.

See also: 802.1X authentication access control authentication local area network (LAN) port

**port ID**
The method that Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) use to specify which program running on the system is sending or receiving the data.

See also: User Datagram Protocol (UDP)

**port rule**
For Network Load Balancing, a set of configuration parameters that determine the filtering mode to be applied to a range of ports.

See also: filtering mode Network Load Balancing port

**port scanning**

A method that is used to identify services or programs that respond to service requests that are made over a network port.

See also: port  service

**Portable Operating System Interface for UNIX (POSIX)**
An Institute of Electrical and Electronics Engineers (IEEE) standard that defines a set of operating-system services. Programs that adhere to the POSIX standard can be easily ported from one system to another. POSIX was based on UNIX system services, but it was created in a way that allows it to be implemented by other operating systems.

See also: service

**POSIX**
See other term: Portable Operating System Interface for UNIX (POSIX)

**possible owner**
A node in a cluster that can run a specific resource. By default, all nodes appear as possible owners, so the resource can run on any node. In most cases, it is appropriate to use this default setting. If you want the resource to be able to fail over, at least two nodes must be designated as possible owners.

See also: cluster  failover  node  resource

**PostScript**
A page-description language (PDL), developed by Adobe Systems for printing on laser printers. PostScript offers flexible font capability and high-quality graphics. It is the standard for desktop publishing because it is supported by imagesetters, the high-resolution printers used by printing services for commercial typesetting.

See also: page-description language (PDL)  PostScript fonts  PostScript printer  Printer Control Language (PCL)  Printer Job Language (PJL)  service  Type 1 fonts

**PostScript fonts**
Fonts that are defined in terms of the PostScript page-description language (PDL) rules and are intended to be printed on a PostScript-compatible printer. When a document displayed in a screen font is sent to a PostScript printer, the printer uses the PostScript version if the font exists. If the font doesn`t exist but a version is installed on the computer, that font is downloaded to the printer. If there is no PostScript font installed in either the printer or the computer, the bit-mapped (raster) font is translated into PostScript and the printer produces text using the bit-mapped font. PostScript fonts are distinguished from bit-mapped fonts by their smoothness, detail, and faithfulness to standards of quality established in the typographic industry.

See also: downloadable fonts  font  page-description language (PDL)  PostScript  raster fonts

**PostScript printer**
A printer that uses the PostScript page-description language (PDL) to create text and graphics on the output medium, such as paper or overhead transparency. Examples of PostScript printers include the Apple LaserWriter, the NEC LC-890, and the QMS PS-810.

See also: page-description language (PDL)  PostScript  virtual printer memory

**POTS**
See other term: point of termination station (POTS)

**power conditioning**
A feature of an uninterruptible power supply (UPS) that removes spikes, surges, sags, and noise from the power supply.

See also: uninterruptible power supply (UPS)

**Power Users group**
A group whose members can manage accounts, resources, and applications that are installed on a workstation, stand-alone server, or member server. This group does not exist on domain controllers. Administrative tasks that can be performed by members of this group include creating local users and groups; modifying and deleting accounts that they have created; removing users from the Power Users, Users, and Guests groups; installing most applications; and creating and deleting file shares.

See also: domain controller  group  member server  stand-alone server

**PPP**
See other term: Point-to-Point Protocol (PPP)

**PPTP**

See other term: Point-to-Point Tunneling Protocol (PPTP)

**precedence**
In Group Policy, the application hierarchy, based on the order in which Group Policy objects (GPOs) are applied. The GPO that is applied last takes precedence over all other objects.

See also: Group Policy Management console (GPMC)  Group Policy object (GPO)  Group Policy object link  inheritance

**predefined key**
A key that represents one of the main divisions of the registry. Each predefined key is displayed in a separate Registry Editor window, with the key's name appearing in the window`s title bar. For example, HKEY_CLASSES_ROOT is a predefined key.

See also: hive  key  registry

**preferred owner**
The node on which you prefer each group to run. For example, the static load balancing model performs best when groups are appropriately balanced between two nodes. When a node fails, the remaining node takes over the groups from the failed node, but performance is diminished. By setting those groups to fail back to their preferred server (the failed node), you automatically restore maximum performance when failback occurs. A group does not fail back if a preferred owner is not selected. You will not always choose a preferred owner because it may not matter where the group resides; all that matters is that the group is still running on one of the two nodes. Or, the nodes may be equally capable of handling the load required to use some or all of the resources.

See also: failback  group  node  resource

**preferred server**
The NetWare server that you connect to by default when you log on to your computer. The preferred server validates your user credentials and is queried when you request information about resources available on the NetWare network.

See also: credentials  resource

**preset-to callback**
A form of security in which a remote access server verifies users by calling them back at numbers supplied by the network administrator at the time user privileges are granted. Only a network administrator can change a preset callback number. This ensures that no one can borrow a user`s password and connect to the server from a location other than the user`s normal one.

See also: callback number  connected, user authenticated  remote access server

**preshared key**
An Internet Protocol security (IPSec) technology in which a shared, secret key is used for authentication in IPSec policy.

See also: authentication  Internet Protocol security (IPsec)  IPsec policy  key

**prestage**
To create a computer account in Active Directory before an operating system has been installed on the computer that will use the account. Prestaging is a part of supporting Remote Installation Services (RIS) for the client computer, and it ensures smooth client installations because the newly installed client computer is immediately recognized as having an account in Active Directory. The computer account information must include the globally unique identifier (GUID), which RIS uses to identify the client computer. An administrator can also use prestaging to control which computers can be installed through RIS.

See also: Active Directory  computer account  globally unique identifier (GUID)  Remote Installation Services (RIS)

**Primary disk**
The hard disk drive that contains the system and boot partitions used to start Windows.

See also: boot partition  system partition

**primary domain controller (PDC)**
In a Windows NT domain, a domain controller running Windows NT Server 4.0 or earlier that authenticates domain logon attempts and updates user, computer, and group accounts in a domain. The PDC contains the master read-write copy of the directory database for the domain. A domain has only one PDC. In a Windows 2000 or Windows Server 2003 domain, the PDC emulator master supports compatibility with client computers that are not running Windows 2000 or Windows XP Professional.

See also: Active Directory  backup domain controller (BDC)  mixed mode  PDC emulator master

**primary group**

The group with which a Macintosh user usually shares documents stored on a server. You specify a user`s primary group in the user`s account. When a user creates a folder on the server, the user`s primary group is set by default as the folder`s associated group.

See also: group

**primary master**
An authoritative DNS server for a zone that can be used as a point of update for the zone. Only primary masters have the ability to be updated directly to process zone updates, which include adding, removing, or modifying resource records that are stored as zone data. Primary masters are also used as the first sources for replicating the zone to other DNS servers.

See also: DNS server  replication  resource record (RR)  zone

**primary partition**
A type of partition that you can create on basic disks. A primary partition is a portion of a physical disk that functions as though it were a physically separate disk. On basic master boot record (MBR) disks, you can create up to four primary partitions on a basic disk, or three primary partitions and an extended partition with multiple logical drives. On basic GUID partition table (GPT) disks, you can create up to 128 primary partitions. Also known as a volume.

See also: basic disk  extended partition  GUID partition table (GPT)  logical drive  master boot record (MBR)  partition  volume

**primary zone**
A copy of the zone that is administered locally.

See also: secondary zone  zone

**principle of least privilege**
The security guideline that a user should have the minimum privileges necessary to perform a specific task. This helps to ensure that, if a user is compromised, the impact is minimized by the limited privileges held by that user. In practice, a user runs within the security context of a normal user. When a task requires additional privileges, the user can use a tool such as Run as to start a specific process with those additional privileges or to log on as a user with the necessary privileges.

See also: privilege  security

**print job**
The source code that contains both the data to be printed and the commands for print. Print jobs are classified into data types based on what modifications, if any, the spooler must make to the job for it to print correctly.

See also: print spooler  Printer window  printing pool

**print processor**
The component that, working in conjunction with the printer driver, receives and alters print jobs, as necessary, according to their data type to ensure that the jobs print correctly.

See also: print job  printer driver

**print server**
A computer that is dedicated to managing the printers on a network. The print server can be any computer on the network.

See also: printer

**Print Server for Macintosh**
A service that enables Macintosh clients to send and spool documents to printers attached to a computer running Windows NT Server; Windows 2000 Server; or an operating system in the Windows Server 2003 family, excluding 64-bit editions, and that enables clients to send documents to printers anywhere on an AppleTalk network. Also known as MacPrint.

See also: service

**print spooler**
Software that accepts a document sent to a printer and then stores it on disk or in memory until the printer is ready for it.

See also: printer  spooling

**Print Spooler resource**
Printer queues providing access to a network printer connected to the network by an IP address rather than by an individual

name. Print spoolers are supported as cluster resources by a Resource DLL.

See also: IP address  print spooler  Resource DLL

**printer**
A device that puts text or images on paper or other print media. Examples are laser printers or dot-matrix printers.

See also: logical printer  print spooler  printing pool

**Printer Control Language (PCL)**
The page-description language (PDL) developed by Hewlett-Packard for their laser and inkjet printers. Because of the widespread use of laser printers, this command language has become a standard in many printers.

See also: page-description language (PDL)  PostScript  Printer Job Language (PJL)

**printer driver**
A program designed to allow other programs to work with a particular printer without concerning themselves with the specifics of the printer`s hardware and internal language. By using printer drivers that handle the subtleties of each printer, programs can communicate properly with a variety of printers.

See also: device driver

**printer fonts**
Fonts residing in or intended for a printer. A printer font, usually located in the printer`s read-only memory (ROM), can be internal, downloaded, or on a font cartridge.

See also: device fonts  downloadable fonts  font  font cartridge

**Printer Job Language (PJL)**
The printer command language developed by Hewlett Packard that provides printer control at the print-job level. Using PJL commands, you can change default printer settings such as number of copies to print. PJL commands also permit switching printer languages between print jobs without action by the user. If bi-directional communication is supported, a PJL-compatible printer can send information such as printer model and job status to the print server.

See also: page-description language (PDL)  PostScript  Printer Control Language (PCL)

**printer permissions**
Permissions that specify the type of access that a user or group has to a printer. The printer permissions are Print, Manage Printers, and Manage Documents.

See also: permission

**Printer window**
A window that shows information about any pending print jobs for the printer. For each printer you have installed or to which you are connected, you can view information such as how many documents are waiting to be printed, who owns them, and how large they are. Also called the queue view.

See also: print job

**Printers and Faxes**
The folder in Control Panel that contains the Add Printer Wizard and icons for all the printers installed on your computer.

See also: printer

**printing pool**
Two or more identical printers that are connected to one print server and act as a single printer. In this case, when you print a document, the print job will be sent to the first available printer in the pool.

See also: print job  printer

**private branch exchange (PBX)**
An automatic telephone switching system that enables users within an organization to place calls to each other without going through the public telephone network. Users can also place calls to outside numbers.

**private key**
The secret half of a cryptographic key pair that is used with a public key algorithm. Private keys are typically used to decrypt a symmetric session key, digitally sign data, or decrypt data that has been encrypted with the corresponding public key.

See also: public key  public key encryption

**private network**
A cluster network that supports only node-to-node communication.

See also: cluster  node

**private or public network-to-network interface (PNNI)**
The ATM Forum standard that defines the interface between asynchronous transfer mode (ATM) switches in a private network or a public network.

See also: asynchronous transfer mode (ATM)

**private queue**
For Message Queuing, a queue that is not published in Active Directory and that can be accessed only by applications that have access to the full format name of the queue.

See also: Active Directory  internal private queue  Message Queuing  public queue

**private volume**
A Macintosh-accessible volume that is accessible by only one Macintosh user. For a volume to be a private volume, the permissions on its root directory must give the volume`s owner all three permissions (Make Changes, See Files, and See Folders), while giving the primary group and everyone categories no permissions at all. When a private volume`s owner uses the Chooser to view the volumes available on the server, the private volume is listed; however, no other users can see the private volume when viewing the volumes available on the server.

See also: Chooser  permission  volume

**privilege**
A user's right to perform a specific task, usually one that affects an entire computer system rather than a particular object. Privileges are assigned by administrators to individual users or groups of users as part of the security settings for the computer.

See also: object  user rights

**process**
The virtual address space and the control information necessary for the execution of a program.

See also: socket

**process identifier (PID)**
A numerical identifier that uniquely distinguishes a process while it runs. Use Task Manager to view PIDs.

See also: process  Task Manager

**program**
A set of instructions that a computer uses to perform a specific task, such as word processing, accounting, or data management. Also called an application.

**program information file (PIF)**
A file that provides information to Windows about how best to run MS-DOS-based programs. When you start an MS-DOS-based program, Windows looks for a PIF to use with it. PIFs contain such items as the name of the file, a start-up directory, and multitasking options.

See also: MS-DOS-based program

**property**
A characteristic or parameter of a class of objects or devices. For example, properties of Microsoft Word files include Size, Created, and Characters.

See also: device  object  property cache  property value  tag

**property cache**
For Indexing Service, a file that stores values for document properties.

See also: document  Indexing Service  property

**property value**

A specific characteristic or parameter that defines a property. For example, property values of a specific Microsoft Word document could include Size = 10,000 bytes, Created = Jan 2, 1999, and Characters = 5,250.

See also: property

**Protective MBR**
The first sector of a GUID partition table (GPT) disk that is structured like the first sector of a master boot record (MBR) disk to prevent x86-based disk utilities from destroying GPT partitions. The Protective MBR contains one partition that reserves the entire space used on the disk by GPT partitions.

See also: Extensible Firmware Interface (EFI)  GUID partition table (GPT)  master boot record (MBR)  x86

**protocol**
A set of rules and conventions for sending information over a network. These rules govern the content, format, timing, sequencing, and error control of messages exchanged among network devices.

See also: Internet Protocol (IP)  Transmission Control Protocol/Internet Protocol (TCP/IP)

**protocol parser**
A dynamic-link library (DLL) that identifies the protocols used to send a frame onto the network.

See also: dynamic-link library (DLL)  frame  protocol

**proxy server**
A firewall component that manages Internet traffic to and from a local area network (LAN) and that can provide other features, such as document caching and access control. A proxy server can improve performance by supplying frequently requested data, such as a popular Web page, and it can filter and discard requests that the owner does not consider appropriate, such as requests for unauthorized access to proprietary files.

See also: local area network (LAN)

**psadmin**
A Windows command-line utility that manages the Password Synchronization component of Identity Management for UNIX, either on a specified computer, or globally.

**PSTN**
See other term: public switched telephone network (PSTN)

**public key**
The nonsecret half of a cryptographic key pair that is used with a public key algorithm. Public keys are typically used when encrypting a session key, verifying a digital signature, or encrypting data that can be decrypted with the corresponding private key.

See also: key  private key  Public Key Cryptography Standards (PKCS)  public key encryption  recovery agent

**public key cryptography**
See other term: public key encryption

**Public Key Cryptography Standards (PKCS)**
A family of standards for public key cryptography that includes RSA encryption, Diffie-Hellman key agreement, password-based encryption, extended-syntax, cryptographic message syntax, private key information syntax, and certificate request syntax, as well as selected attributes. Developed, owned, and maintained by RSA Data Security, Inc.

See also: public key cryptography

**public key encryption**
A method of encryption that uses two encryption keys that are mathematically related. One key is called the private key and is kept confidential. The other is called the public key and is freely given out to all potential correspondents. In a typical scenario, a sender uses the receiver's public key to encrypt a message. Only the receiver has the related private key to decrypt the message. The complexity of the relationship between the public key and the private key means that, provided the keys are long enough, it is computationally infeasible to determine one from the other. Also called asymmetric encryption.

See also: encryption  private key  public key  symmetric encryption

**public key infrastructure (PKI)**
The laws, policies, standards, and software that regulate or manipulate certificates and public and private keys. In practice, it is a system of digital certificates, certification authorities, and other registration authorities that verify and authenticate the validity

of each party involved in an electronic transaction. Standards for PKI are still evolving, even though they are being widely implemented as a necessary element of electronic commerce.

See also: certification authority (CA)  public key

**public network**
A cluster network that supports client-to-cluster communication (either with or without supporting node-to-node communication).

See also: cluster  node

**public queue**
For Message Queuing, a queue that is published in Active Directory and replicated throughout a Windows enterprise. Public queues can, therefore, be located by any computer running Message Queuing within the enterprise.

See also: Active Directory  Message Queuing  private queue  queue

**public switched telephone network (PSTN)**
Standard analog telephone lines, available worldwide.

**pull partner**
A WINS component that requests replication of updated WINS database entries from its push partner.

See also: push partner  replica  Windows Internet Name Service (WINS)

**pulse dialing**
A form of dialing that enters a phone number by means of pulse frequencies. The user typically hears a series of clicking sounds when dialing. Old-fashioned rotary dial phones use pulse dialing.

See also: touch-tone dialing

**push partner**
A WINS component that notifies its pull partner when updated WINS database entries are available for replication.

See also: pull partner  replica  Windows Internet Name Service (WINS)

**PVC**
See other term: permanent virtual circuit (PVC)

# Windows Server 2003 Glossary - Q

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - Q**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**QoS**
See other term: Quality of Service (QoS)

**qualified subordination**
An extension of a standard certification authority (CA) that allows you to place certificate issuance constraints on subordinate CAs and to place usage constraints on the certificates that are issued by subordinate CAs.

See also: certification authority (CA)  public key infrastructure (PKI)

**Quality of Service (QoS)**
A set of quality assurance standards and mechanisms for data transmission, implemented in this version of Windows.

**query**
For Indexing Service, a structured statement that specifies the documents you want to find. The simplest query is a single word.

See also: Indexing Service  tag  vector  wildcard character

**queue**
A list of programs or tasks waiting for execution. In Windows printing terminology, a queue refers to a group of documents waiting to be printed. In NetWare and OS/2 environments, queues are the primary software interface between the application and print device; users submit documents to a queue. With Windows, however, the printer is that interface; the document is sent to a printer, not a queue.

See also: printer  transactional message

**queue quota**
For Message Queuing, the storage size limit for messages in public queues. When a queue quota is reached, Message Queuing can no longer send messages to that queue until one or more messages are removed from the queue. Message Queuing enforces the computer quota before it enforces the queue quota on a computer.

See also: computer quota  Message Queuing  queue

**queue type**
For Message Queuing, a globally unique identifier (GUID) specified by the application that created the queue.

See also: globally unique identifier (GUID)  Message Queuing  queue

**quiet answer**
A telephone-answering protocol in which incoming calls are answered with silence instead of a tone signal. Some telephone-switching systems use quiet answering. These switching systems expect the caller to provide another phone number, code, or extension after the quiet answer.

See also: protocol

**quorum log**
The log where the quorum resource stores data. This data is maintained by the clustering software. Also known as the recovery log or change log.

See also: quorum resource  server cluster

**quorum resource**
The quorum-capable resource selected to maintain the configuration data necessary for recovery of the cluster. This data contains details of all of the changes that have been applied to the cluster database. The quorum resource is generally accessible to other cluster resources so that any cluster node has access to the most recent database changes. By default there

is only one quorum resource per cluster.

See also: node  server cluster

**quorum traffic**
In a majority node set server cluster, the replication traffic between nodes in a cluster. This replication traffic contains the cluster configuration data and ensures that the cluster quorum information is kept in sync across all nodes.

See also: cluster  majority node set server cluster  node  quorum resource  replication

**quota limit**
The amount of disk space available to a user.

See also: warning level

**quota template**
A template that defines a space limit on the physical drive, the type of quota (hard or soft), and a set of notifications to be generated when the quota limit is being approached.

# Windows Server 2003 Glossary - R

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - R**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**RA**
See other term: registration authority (RA)

**RADIUS**
See other term: Remote Authentication Dial-In User Service (RADIUS)

**RAID-5 volume**
A fault-tolerant volume with data and parity striped intermittently across three or more physical disks. Parity is a calculated value that is used to reconstruct data after a failure. If a portion of a physical disk fails, Windows recreates the data that was on the failed portion from the remaining data and parity. You can create RAID-5 volumes only on dynamic disks on computers running the Windows 2000 Server or Windows Server 2003 families of operating systems. You cannot mirror or extend RAID-5 volumes. In Windows NT 4.0, a RAID-5 volume was known as a striped set with parity.

See also: dynamic disk  dynamic volume  fault tolerance  parity  volume

**RAM**
See other term: random access memory (RAM)

**random access memory (RAM)**
Memory that can be read from or written to by a computer or other devices. Information stored in RAM is lost when the computer is turned off.

See also: virtual memory

**raster fonts**
Fonts that are stored as bitmaps. Raster fonts are designed with a specific size and resolution for a specific printer and cannot be scaled or rotated. If a printer does not support raster fonts, it will not print them. The five raster fonts are Courier, MS Sans Serif, MS Serif, Small, and Symbol. Also called bit-mapped fonts.

See also: font  printer

**raw socket**
A socket that provides direct access to lower-level network protocols.

See also: socket

**read-only memory (ROM)**
A semiconductor circuit that contains information that cannot be modified.

**realm**
A set of security principles, in a non-Windows networked environment, that are subject to Kerberos authentication.

See also: Kerberos V5 authentication protocol

**realm name**
An identifying prefix or suffix appended to a user name to enable appropriate routing and authentication during a remote logon process.

See also: authentication  routing  user name

**realm trust**
A trust between non-Windows Kerberos V5 realms, such as a UNIX realm, and Active Directory domains. Realm trusts can be transitive, nontransitive, one-way, or two-way.

See also: Active Directory  domain  Kerberos V5 authentication protocol  nontransitive trust  one-way trust  realm  transitive trust  two-way trust

**Real-Time Transport Protocol (RTP)**
An Internet protocol for transporting real-time traffic over multicast and unicast network services.

See also: protocol  unicast

**rebinding state**
A state used by DHCP clients to extend and renew their address lease when the current lease is close to expiring. In this state, the client broadcasts to the network to locate any DHCP server that can either renew or replace its currently leased configuration. The rebinding state begins when 87.5 percent of the client's lease time has elapsed.

See also: Dynamic Host Configuration Protocol (DHCP)  lease

**receiving member**
Between a pair of replicating servers, the member that is receiving replicated content over a connection.

**reconciliation**
The process of making the client lease information stored in the DHCP server database consistent with a duplicate copy of this same information stored in the Windows registry. Using the DHCP Microsoft Management Console (MMC), reconciliation helps restore or recover the contents of the DHCP server database file, Dhcp.mdb, located by default in the %Systemroot%\System32\Dhcp folder.

See also: Dynamic Host Configuration Protocol (DHCP)  lease  Microsoft Management Console (MMC)  systemroot

**recoverable message**
For Message Queuing, a message that can be recovered no matter which computer fails, but that uses more resources and is slower than an express message.

See also: express message  Message Queuing

**recovery agent**
A person who is issued a public key certificate for the purpose of recovering user data that is encrypted with Encrypting File System (EFS).

See also: Encrypting File System (EFS)  public key

**Recovery Console**
A command-line interface that provides a limited set of administrative commands that are useful for repairing a computer.

See also: NTFS file system

**recovery policy**
A type of public key Group Policy object (GPO) used by Encrypting File System (EFS) that provides for one or more user accounts to be designated as recovery agents.

See also: Encrypting File System (EFS)  Group Policy object (GPO)  public key  recovery agent  user account

**recursive query**
A query made to a DNS server in which the requester asks the server to assume the full workload and responsibility for providing a complete answer to the query. The DNS server will then use separate iterative queries to other DNS servers on behalf of the requester to assist in completing an answer for the recursive query.

See also: DNS server  iterative query  recursive resolution

**recursive resolution**
One of the two process types (iterative and recursive) for DNS name resolution. In this process, a resolver (a DNS client) will request that a DNS server provide a complete answer to a query that does not include pointers to other DNS servers. When a client makes a query and requests that the server use recursive resolution to answer, it effectively shifts the workload of resolving the query from the client to the DNS server. If the DNS server supports and uses recursive resolution, it contacts other DNS servers as necessary (using iterative queries on behalf of the client) until it obtains a definitive answer to the query. This type of resolution allows the client resolver to be small and simple.

See also: DNS server  Domain Name System (DNS)  iteration  iterative query  recursive query

**Recycle Bin**

The place in which Windows stores deleted files. You can retrieve files you deleted in error, or you can empty the Recycle Bin to create more disk space.

**Redundant Array of Independent Disks (RAID)**
A method used to standardize and categorize fault-tolerant disk systems. RAID levels provide various mixes of performance, reliability, and cost. Some servers provide three of the RAID levels: Level 0 (striping), Level 1 (mirroring), and Level 5 (RAID-5).

See also: fault tolerance  RAID-5 volume

**referral**
An ordered list of targets that a client computer receives from a domain controller or namespace server when the user accesses a namespace root or folder with targets in the namespace.

See also: DFS namespace  Distributed File System (DFS)  iteration

**referral server**
A server that responds to requests for service from remote installation clients. It refers clients to other remote installation servers where they can install Windows desktop operating systems. The referral server does not directly host client computers for installing the operating system.

**refresh**
In PerformancePoint Planning, to update displayed information with current data. In PerformancePoint Monitoring, the activity of synchronizing dashboards and dashboard elements between the local workspace and Monitoring Server.

See also: refresh rate

**refresh interval**
An interval of time used by secondary masters of a zone to determine how often to check if their zone data needs to be refreshed. When the refresh interval expires, the secondary master checks with its source for the zone to see if its zone data is still current or if it needs to be updated using a zone transfer. This interval is set in the start-of-authority (SOA) resource record for each zone.

See also: start-of-authority (SOA) resource record  zone  zone transfer

**refresh rate**
The frequency with which the video screen is retraced to prevent the image from flickering. The entire image area of most monitors is refreshed approximately 60 times per second.

See also: refresh

**region**
Contiguous chunks of storage on a disk.

See also: disk

**registered file type**
File types that are tracked by the system registry and are recognized by the programs you have installed on your computer.

See also: file type  registry

**registration authority (RA)**
A computer that is configured for an administrator to request and retrieve issued certificates on behalf of other users. An RA does not require that a certification authority be installed on the same computer.

See also: certification authority (CA)

**registry**
A database repository for information about a computer's configuration. The registry contains information that Windows continually references during operation, such as: profiles for each user; the programs installed on the computer and the types of documents that each can create; property settings for folders and program icons; what hardware exists on the system; which ports are being used. The registry is organized hierarchically as a tree, and it is made up of keys and their subkeys, hives, and entries.

See also: entry  hive  key  registry size limit (RSL)  subkey

**registry boot**
The default boot option used by most Windows DNS servers. When registry boot is used, DNS is started and initialized using

DNS parameters and their values as they are stored in the Windows registry. You can use a Berkeley Internet Name Domain (BIND) boot file as an alternative to this method of boot configuration for DNS.

See also: BIND boot file  DNS server  Domain Name System (DNS)

**registry size limit (RSL)**
A universal maximum for the space that registry files (hives) can consume in the paged pool. This maximum prevents an application from filling the paged pool with registry data.

See also: hive  paged pool  registry

**relative distinguished name**
The part of the object name that identifies the object as unique from its siblings at its level in the naming hierarchy. For example, in the distinguished name CN=My Name,CN=Users,DC=Microsoft,DC=Com the relative distinguished name of the user object is My Name. The relative distinguished name of the user object`s parent object is Users.

See also: object

**relative ID (RID)**
The part of a security ID (SID) that uniquely identifies an account or group within a domain.

See also: domain  forest  group  RID master  security ID (SID)

**relative name**
The partial DNS domain name configured in individual resource records to locate and qualify the record within a zone. The relative name is joined to the front of the parent domain (domain of origin) for each resource record to form a fully qualified domain name (FQDN) within the zone. In DNS Manager, the relative name will correspond to fields that use record-specific name properties, such as the Host computer name field used in an address (A) resource record.

See also: address (A) resource record  Domain Name System (DNS)  fully qualified domain name (FQDN)  resource record (RR)  zone

**remote access**
Part of the integrated Routing and Remote Access service that provides remote networking for telecommuters, mobile workers, and system administrators who monitor and manage servers at multiple branch offices. Users can use Network Connections to dial in to remotely access their networks for services such as file and printer sharing, electronic mail, scheduling, and SQL database access.

See also: remote access server  response  Serial Line Internet Protocol (SLIP)  service

**remote access policy**
A set of conditions and connection parameters that define the characteristics of the incoming connection and the set of constraints imposed on it. Remote access policy determines whether a specific connection attempt is authorized to be accepted.

**remote access server**
A Windows-based computer running the Routing and Remote Access service and configured to provide remote access.

See also: remote access  set-by-caller callback

**Remote Access Service (RAS)**
A Windows NT 4.0 service that provides remote networking for telecommuters, mobile workers, and system administrators who monitor and manage servers at multiple offices.

**remote administration**
The management of one computer by an administrator working at another computer that is connected to the first computer across a network.

See also: remotely administered server

**Remote Authentication Dial-In User Service (RADIUS)**
A security authentication protocol based on a client/server model and widely used by Internet service providers (ISPs). RADIUS is the most popular means of authenticating and authorizing dial-up, virtual private network (VPN), wireless, and authenticating switch clients today. A RADIUS client is included in the Routing and Remote Access service that ships with the Windows Server 2003 family. A RADIUS server and proxy, named Internet Authentication Service (IAS), is included in Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows

Server 2003, Datacenter Edition.

See also: authentication  authorization  Internet Authentication Service (IAS)  tunnel  virtual private network (VPN)

**remote computer**
A computer that you can access only by using a communications line or a communications device, such as a network card or a modem.

See also: local computer

**remote differential compression (RDC)**
A protocol that can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, and re-arrangements of data in files.

**Remote Installation Services (RIS)**
Software services that allow an administrator to set up new client computers remotely, without having to visit each client. The target clients must support remote booting.

See also: Single Instance Store (SIS)

**remote procedure call (RPC)**
A message-passing facility that allows a distributed application to call services that are available on various computers on a network. Used during remote administration of computers.

See also: service

**Remote Storage**
A data management service used to migrate infrequently accessed files from local storage to remote storage. Migrated files are recalled transparently when the user opens the file.

See also: local storage  service  validation

**remotely administered server**
A server that you can administer by using a different computer. You typically access this type of server by using a network connection. A remotely administered server can have a local keyboard, mouse, or video card and monitor. If it does not have these peripherals attached, it is also known as a headless server. Such servers are often housed in a physically secure location.

See also: physical security  remote administration  server

**Removable Storage**
A service used for managing removable media (such as tapes and discs) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

See also: library  service

**rendering intent**
In color management, the approach used to map the colors specified in an image file to the color gamut of your monitor or printer. The color gamut is the range of color that a device can produce.

See also: color gamut

**rendezvous**
Client service referral behavior performed using a DNS query to resolve domain names to locate host computers that are service providers for well-known TCP/IP-based network services. Rendezvous for clients uses queries for resource record types such as mail exchanger (MX) or service (SRV) records.

See also: Domain Name System (DNS)  resource record (RR)  service

**reparse points**
NTFS file system objects that have a definable attribute containing user-controlled data and that are used to extend functionality in the input/output (I/O) subsystem.

See also: attribute  NTFS file system  object

**replica**
In Active Directory replication, one instance of a logical Active Directory partition that is synchronized by means of replication between domain controllers that hold copies of the same directory partition. <i>Replica</i> can

also refer to an instance of an object or attribute in a distributed directory. In the File Replication service (FRS), a computer that has been included in the configuration of a specific replica set.

See also: Active Directory File Replication service (FRS) replica set replication topology

**replica set**
One or more shared folders that participates in replication.

See also: replication shared folder

**replicated folder**
A folder that is kept synchronized on members of a replication group. A replicated folder has an associated local path on each member.

**replication**
The process of copying updated data from a data store or file system on a source computer to a matching data store or file system on one or more destination computers to synchronize the data.

See also: Active Directory DFS root directory partition Distributed File System (DFS) File Replication service (FRS) multimaster replication replica replication policy topology

**replication group**
A set of servers that participates in the replication of one or more replicated folders.

**replication latency**
In Active Directory replication, the delay between the time an update is applied to a given replica of a directory partition and the time it is applied to some other replica of the same directory partition. A server receives changes no sooner than either it is notified of a change from its neighbor in the same site or its periodic replication timer expires. Sometimes referred to as propagation delay.

See also: Active Directory replication directory partition replica

**replication partner**
A domain controller that acts as a replication source for a given domain controller. The Knowledge Consistency Checker (KCC) determines which servers are best suited to replicate with each other, and it generates the list of domain controllers that are candidates for replication partners from the list of domain controllers in the site on the basis of connectivity, history of successful replication, and matching of full and partial replicas. A domain controller has some number of direct replication partners with which it replicates for a given directory partition. The other domain controllers in the site replicate transitively with this domain controller.

See also: directory partition domain controller Knowledge Consistency Checker (KCC) site

**replication policy**
Rules that define how and when replication is performed.

See also: replication

**replication topology**
In Active Directory replication, the set of physical connections that domain controllers use to replicate directory updates among domain controllers within sites and between sites. In the File Replication service (FRS), the interconnections between replica set members. These interconnections determine the path that data takes as it replicates to all replica set members.

See also: Active Directory Distributed File System (DFS) domain controller File Replication service (FRS) replica replica set replication

**report message**
For Message Queuing, a message that contains status information and is sent to report queues on the source computer. Examples include test messages and route tracking messages.

See also: Message Queuing report queue

**report queue**
For Message Queuing, a queue on the source computer that contains report messages. Report queues can be used when sending test messages and tracking message routes.

See also: Message Queuing report message

**Request for Comments (RFC)**
An official document of the Internet Engineering Task Force (IETF) that specifies the details for protocols included in the TCP/IP family.

See also: Internet Engineering Task Force (IETF) protocol Transmission Control Protocol/Internet Protocol (TCP/IP)

**reservation**
A specific IP address within a scope permanently reserved for leased use to a specific DHCP client. Client reservations are made in the DHCP database using DHCP Manager and based on a unique client device identifier for each reserved entry.

See also: Dynamic Host Configuration Protocol (DHCP) IP address lease scope

**resolver**
DNS client programs used to look up DNS name information. Resolvers can be either a small stub (a limited set of programming routines that provide basic query functionality) or larger programs that provide additional lookup DNS client functions, such as caching.

See also: caching caching resolver Domain Name System (DNS)

**resource**
Generally, any part of a computer system or network, such as a disk drive, printer, or memory, that can be allotted to a running program or a process. For Device Manager, any of four system components that control how the devices on a computer work. These four system resources are interrupt request (IRQ) lines, direct memory access (DMA) channels, input/output (I/O) ports, and memory addresses. For server clusters, a physical or logical entity that is capable of being managed by a cluster, brought online and taken offline, and moved between nodes. A resource can be owned only by a single node at any point in time.

See also: direct memory access (DMA) input/output (I/O) port interrupt request (IRQ) lines memory address node offline online server cluster

**resource allocation**
The process of distributing a computer system's facilities to different components of a job in order to perform the job.

**Resource DLL**
A dynamic-link library (DLL) containing an implementation of the Resource application programming interface (API) for a specific type of resource. The Resource DLL is loaded into the address space of its Resource Monitor.

See also: application programming interface (API) dynamic-link library (DLL)

**resource fork**
One of two forks (resource fork and data fork) that make up each Macintosh file. The resource fork holds Macintosh operating system resources, such as code, menu, font, and icon definitions. Resource forks have no relevance to personal computers, so the resource forks of files on the server are never accessed by personal computer clients.

See also: data fork

**resource group**
In a server cluster, a defined collection of resources. Resources that are dependent on each other are typically placed within the same resource group.

See also: node resource server cluster

**Resource Monitor**
A cluster software component that facilitates communication between a node`s server cluster and one or more of its resources.

See also: node resource server cluster

**resource partner**
A federation partner that trusts the Federation Service to issue claims-based security tokens. The resource partner contains published Web-based applications that users in the account partner can access.

**resource record (RR)**
A standard DNS database structure containing information used to process DNS queries. For example, an address (A) resource record contains an IP address corresponding to a host name. Most of the basic resource record types are defined in RFC 1035, but additional RR types have been defined in other RFCs and approved for use with DNS.

See also: Domain Name System (DNS) Request for Comments (RFC) retry interval zone

**resource record set (RRset)**
A collection of more than one resource record returned in a query response by a DNS server. RRsets are used in responses where more than one record is part of the answer.

See also: DNS server  resource record (RR)

**resource types**
Classes of resources, such as print spooler, physical disk, and file share, organized by the server cluster. Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, provide Resource DLL files for the most common types of resources. Using the application programming interface (API) provided in the Microsoft Platform Software Development Kit (SDK), other vendors can add support for other resource types.

See also: print spooler  Resource DLL  server cluster  storage-class resource

**response**
In Windows remote access, strings expected from the device, which can contain macros.

See also: remote access  string

**response message**
For Message Queuing, a message sent by a receiving application to the response queue specified by a sending application. Any available queue can be specified as a response queue.

See also: Message Queuing  queue  response queue

**response queue**
For Message Queuing, a queue that is created by the sending application and used by the receiving application to reply to messages. For example, an application might send a response message to a response queue every time the application receives a message.

See also: Message Queuing  queue  response message

**restricted token**
An access token with restricted access to securable objects or performing privileged tasks.

See also: access token  object  privilege  security ID (SID)

**Resultant Set of Policy (RSoP)**
A feature that simplifies Group Policy implementation and troubleshooting. RSoP uses Windows Management Instrumentation (WMI) to determine how policy settings are applied to users and computers. RSoP has two modes: logging mode and planning mode. Logging mode determines the resultant effect of policy settings that have been applied to an existing user and computer based on a site, domain, and organizational unit. Planning mode simulates the resultant effect of policy settings that are applied to a user and computer.

See also: Group Policy  Group Policy Management console (GPMC)  Microsoft Management Console (MMC)  Windows Management Instrumentation (WMI)

**retry interval**
The time, in seconds after the refresh interval expires, used by secondary masters of a zone to determine how often to try and retry contacting its source for zone data to see if its replicated zone data needs to be refreshed. This interval is set in the start-of-authority (SOA) resource record for each zone.

See also: refresh interval  replication  start-of-authority (SOA) resource record  zone

**reverse lookup**
A DNS query for a pointer (PTR) resource record that resolves an IP address to its corresponding fully qualified domain name (FQDN).

See also: Domain Name System (DNS)  fully qualified domain name (FQDN)  IP address  pointer (PTR) resource record

**reversible encryption**
A mechanism that stores an encrypted password in such a way that the original password can be unencrypted and retrieved. Some applications require the unencrypted password so that they can perform certain tasks.

See also: encrypted password

**RFC**

See other term: Request for Comments (RFC)

**RGB color space**
Multidimensional color space consisting of the red, green, and blue intensities that make up a given color. This system is typically used in scanners, digital cameras, computer monitors, and computer printers.

See also: color management  color space

**RID**
See other term: relative ID (RID)

**RID master**
A domain controller that holds the RID operations master role in Active Directory. The RID master is assigned to allocate unique sequences of relative IDs to each domain controller in its domain. As the domain controllers use the IDs allocated, they contact the RID master and are allocated additional sequences as needed. At any time, the RID master role can be assigned to only one domain controller in each domain.

See also: Active Directory  domain controller  operations master  relative ID (RID)  security ID (SID)  security principal

**RIP**
See other term: Routing Information Protocol (RIP)

**RIPX**
See other term: Routing Information Protocol over IPX (RIPX)

**RIS**
See other term: Remote Installation Services (RIS)

**Rivest-Shamir-Adleman (RSA) cryptographic algorithms**
A widely used set of public key algorithms that were published by RSA Data Security, Inc. The RSA cryptographic algorithms are supported by the Microsoft Base Cryptographic Service Provider and the Microsoft Enhanced Cryptographic Service Provider.

See also: cryptographic service provider (CSP)  public key cryptography

**roaming user profile**
A server-based user profile that is downloaded to the local computer when a user logs on and that is updated both locally and on the server when the user logs off. A roaming user profile is available from the server when logging on to a workstation or server computer. When logging on, the user can use the local user profile if it is more current than the copy on the server.

See also: local user profile  mandatory user profile  user profile

**role-based administration**
A management technique that is characterized by the use of collections of settings that are based on an object's role within an organization. Role-based administration can be used to manage users, computers, and other file system and directory service objects.

See also: access control entry (ACE)  authorization  object

**rolling upgrade**
In a cluster, the process of systematically upgrading cluster nodes one node at a time. During the upgrade, the remaining cluster nodes continue to provide service.

See also: cluster  node

**ROM**
See other term: read-only memory (ROM)

**root**
The highest or uppermost level in a hierarchically organized set of information. The root is the point from which further subsets are branched in a logical sequence that moves from a broad or general focus to narrower perspectives.

See also: DFS root  root certificate  root certification authority  root hints  root servers  root target

**root certificate**
A self-signed certification authority certificate. It is called a root certificate because it is the certificate for the root authority. The root authority must sign its own certificate because by definition there is no higher certifying authority in the certification

hierarchy.

See also: certification authority (CA)  certification hierarchy  root certification authority

**root certification authority**
The most trusted certification authority (CA), which is at the top of a certification hierarchy. The root CA has a self-signed certificate. Also called the root authority.

See also: certification authority (CA)  certification hierarchy  root

**root directory**
The top-level directory (or folder) on a computer, partition or volume, or Macintosh-accessible volume.

See also: Macintosh-accessible volume  partition  volume

**root domain**
The beginning of the DNS namespace. In Active Directory, the initial domain in an Active Directory tree. Also, the initial domain of a forest.

See also: Active Directory  domain  Domain Name System (DNS)  forest  namespace

**root hints**
DNS data stored on a DNS server that identifies the authoritative DNS servers for the root zone of the DNS namespace. The root hints are stored in the file Cache.dns, located in the systemroot\System32\Dns folder.

See also: DNS server  Domain Name System (DNS)  root servers  systemroot

**root servers**
DNS servers that are authoritative for the root of the namespace.

See also: authoritative  DNS server  root

**root target**
A Universal Naming Convention (UNC) path of a shared folder that is associated with the namespace root. A domain-based namespace can have multiple root targets, whereas a stand-alone namespace can have a single root target only.

See also: DFS root  target

**round robin**
A simple mechanism used by DNS servers to share and distribute loads for network resources. Round robin is used to rotate the order of resource records (RRs) returned in a response to a query when multiple RRs of the same type exist for a queried DNS domain name.

See also: DNS server  resource record (RR)

**Routemon utility**
A scripting utility for the Routing and Remote Access service that is a command-line alternative to the router administration user interface available through the Routing and Remote Access Manager.

**router**
A device or computer that forwards packets between interfaces based on a network layer destination address. For example, an Internet Protocol (IP) router forwards IP packets based on the destination IP address in the IP header. Routers typically use a routing table, which contains a series of entries for destinations and the corresponding next-hop address and interface to use to forward a packet to its eventual destination.

See also: local area network (LAN)  packet header  routing  Routing Information Protocol over IPX (RIPX)  static routes  Token Ring  wide area network (WAN)

**routing**
The process of forwarding a packet through an internetwork from a source host to a destination host.

See also: host  packet

**Routing Information Protocol (RIP)**
An industry standard, distance vector routing protocol used in small- to medium-sized Internet Protocol (IP) and Internetwork Packet Exchange (IPX) internetworks.

See also: Internet Protocol (IP)  Internetwork Packet Exchange (IPX)  protocol

**Routing Information Protocol over IPX (RIPX)**
A protocol used by routers to exchange information between routers on an Internetwork Packet Exchange (IPX) network and by hosts to determine the best router to use when forwarding IPX traffic to a remote IPX network.

See also: host Internetwork Packet Exchange (IPX) NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) protocol router

**routing link**
For Message Queuing, a communications link established between Windows sites for routing messages. Specially configured Message Queuing servers with routing services enabled are used to create a routing link between sites.

See also: Message Queuing routing services routing-link cost site

**routing protocol**
Any of several protocols that enable the exchange of routing table information between routers. Typically, medium- to large-sized TCP/IP internetworks implement routing protocols to simplify the administration of routing tables.

See also: router routing

**routing services**
For Message Queuing, a service on a Message Queuing server that provides message routing services. If so configured, this feature can be used on a Message Queuing server to enable computers that use different network protocols to communicate, reduce the number of sessions by acting as a gateway for all incoming or outgoing messages for independent clients, and route messages between sites over a routing link.

See also: gateway independent client Message Queuing Message Queuing server routing link service session concentration site

**routing-link cost**
For Message Queuing, a number used to determine the route messages can take between two sites. This number represents the relative monetary cost of communication over a link. A routing link has a default routing-link cost of 1 and should not be changed unless you have multiple routing links between two sites and you want to enforce message routing over a specific routing link.

See also: Message Queuing routing link site

**RPC**
See other term: remote procedure call (RPC)

**RR**
See other term: resource record (RR)

**RRset**
See other term: resource record set (RRset)

**RS-232-C standard**
An accepted industry standard for serial communication connections. Adopted by the Electric Industries Association, this Recommended Standard (RS) defines the specific lines and signal characteristics used by erial communications controllers to standardize the transmission of serial data between devices. The letter C signifies the third in a series.

See also: device

**RSA**
See other term: Rivest-Shamir-Adleman (RSA) cryptographic algorithms

**RSL**
See other term: registry size limit (RSL)

**RSoP**
See other term: Resultant Set of Policy (RSoP)

**RTP**
See other term: Real-Time Transport Protocol (RTP)

**Run as**
A feature that provides users with a secondary logon capability. By using Run as, users can run applications or commands in a different security context without having to log off. Run as prompts the user for different credentials before running the

application or command.

See also: secondary logon  security context

**runaway recall limit**
In Remote Storage, the setting that prohibits a user or an application from recalling files from Remote Storage more than the specified number of times in succession with less than 10 seconds between each recall. Otherwise, applications such as those for virus checking, data backup, indexing, and searching typically read all files on your computer and can cause excessive file recalls and slow performance.

See also: Remote Storage

# Windows Server 2003 Glossary - S

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - S**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**S/MIME**
See other term: Secure/Multipurpose Internet Mail Extensions (S/MIME)

**SA**
See other term: security association (SA)

**SACL**
See other term: system access control list (SACL)

**Safe mode**
A method of starting Windows using only basic files and drivers. Safe mode is available by pressing the F8 key during startup. This allows you to start the computer when a problem prevents it from starting correctly.

See also: Safe Mode with Command Prompt  Safe Mode with Networking

**Safe Mode with Command Prompt**
A method of starting Windows using basic files and drivers only, without networking, and with only a command prompt displayed. This mode is available by pressing the F8 key when prompted during startup. This allows you to start your computer when a problem prevents it from starting normally.

See also: Safe mode  Safe Mode with Networking

**Safe Mode with Networking**
A method of starting Windows using only basic files, drivers, and networking. This mode is available by pressing the F8 key when prompted during startup. This allows you to start your computer when a problem prevents it from starting normally.

See also: Safe mode  Safe Mode with Command Prompt

**SAM**
See other term: Security Accounts Manager (SAM)

**SAM account name**
A domain-unique security principal name in Windows NT 4.0 and earlier.

See also: Security Accounts Manager (SAM)  security principal name

**SAN**
See other term: system area network (SAN)

**SAP**
See other term: Service Advertising Protocol (SAP)

**SASL**
See other term: Simple Authentication and Security Layer (SASL)

**saturation**
In color management, the purity of a color's hue, moving from gray to the pure color.

See also: hue

**SBM**
See other term: subnet bandwidth management (SBM)

**scalability**

A measure of how well a computer, service, or application can grow to meet increasing performance demands. For server clusters, the ability to incrementally add one or more systems to an existing cluster when the overall load of the cluster exceeds its capabilities.

See also: server cluster

**scavenging**
The process of cleaning and removing extinct or outdated names data from the WINS database.

See also: Windows Internet Name Service (WINS)

**Schannel**
See other term: secure channel (Schannel)

**schema**
The set of definitions for the universe of objects that can be stored in a directory. For each object class, the schema defines which attributes an instance of the class must have, which additional attributes it can have, and which other object classes can be its parent object class.

See also: attribute  directory partition  object  object class  parent object

**schema master**
A domain controller that holds the schema operations master role in Active Directory. The schema master performs write operations to the directory schema and replicates updates to all other domain controllers in the forest. At any time, the schema master role can be assigned to only one domain controller in the forest.

See also: Active Directory  domain controller  operations master  schema

**scope**
A range of IP addresses that are available to be leased or assigned to DHCP clients by the DHCP service.

See also: Dynamic Host Configuration Protocol (DHCP)  IP address  lease

**scope of influence**
In a domain environment, a site, domain, or organizational unit; in a workgroup environment, the local disk.

See also: domain  organizational unit  site

**scope of management (SOM)**
In Group Policy, any Active Directory container to which you can link a Group Policy object (GPO). These containers can be sites, domains, or organizational units.

See also: Active Directory  Group Policy  Group Policy object (GPO)

**screen font**
A typeface designed for display on a computer monitor screen. A screen font often has an accompanying PostScript font for printing to PostScript-compatible printers.

See also: font  PostScript

**screen resolution**
The setting that determines the amount of information that appears on your screen, measured in pixels. Low resolution, such as 640 x 480, makes items on the screen appear large, although the screen area is small. High resolution, such as 1024 x 768, makes the overall screen area large, although individual items appear small.

See also: pixel

**screen saver**
A moving picture or pattern that appears on a computer screen when the mouse or keyboard has not been used for a specified period of time.

**SCSI**
See other term: small computer system interface (SCSI)

**SDP**
See other term: Session Description Protocol (SDP)

**secondary cluster**
A specific virtual IP address assigned to a Network Load Balancing cluster (the "primary cluster"). The secondary cluster`s virtual IP address is different than the primary cluster`s virtual IP address. Secondary clusters allow you to configure an independent set of port rules for each virtual IP address in your Network Load Balancing (primary) cluster. Also known as a virtual cluster.

See also: cluster  Network Load Balancing  virtual IP address

**secondary logon**
The practice of logging on by using one security context and then, within the initial logon session, authenticating and using a second account. In Windows 2000, Windows XP Professional, and the Windows Server 2003 family, secondary logon is enabled by the RunAs.exe program and service.

See also: authentication

**secondary zone**
A read-only copy of a DNS zone that is transferred from an authoritative DNS server to another DNS server to provide redundancy.

See also: authoritative  DNS server  DNS zone  Domain Name System (DNS)  primary zone  zone

**second-level domain**
A DNS domain name that is rooted hierarchically at the second tier of the domain namespace, directly beneath the top-level domain names. Top-level domain names include .com and .org. When DNS is used on the Internet, second-level domains are names that are registered and delegated to individual organizations and businesses.

See also: domain name  Domain Name System (DNS)  parent domain  top-level domains

**sector**
A 512-byte unit of physical storage on a hard disk. Windows file systems allocate storage in clusters, where a cluster is one or more contiguous sectors.

See also: cluster  file system

**secure channel (Schannel)**
A Security Support Provider (SSP) that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) Internet standard authentication protocols.

See also: authentication protocol  Secure Sockets Layer (SSL)  Transport Layer Security (TLS)

**secure dynamic update**
The process in which a DNS client submits a dynamic update request to a DNS server and the DNS server performs the update only if the client is authenticated.

See also: authentication  DNS client  DNS server  dynamic update

**Secure Hash Algorithm (SHA-1)**
An algorithm that generates a 160-bit hash value from an arbitrary amount of input data. SHA-1 is used with the Digital Signature Algorithm (DSA) in the Digital Signature Standard (DSS), among other places.

See also: Digital Signature Standard (DSS)  hash algorithm

**Secure Hypertext Transfer Protocol**
A protocol that provides a secure Hypertext Transfer Protocol (HTTP) connection.

See also: Hypertext Transfer Protocol (HTTP)  protocol

**Secure Sockets Layer (SSL)**
A proposed open standard for establishing a secure communications channel to prevent the interception of critical information, such as credit card numbers. Primarily, it enables secure electronic financial transactions on the World Wide Web, although it is designed to work on other Internet services as well.

See also: Internet

**secure zone**
A DNS zone that is stored in Active Directory and to which access control list (ACL) security features are applied.

See also: access control list (ACL)  Active Directory  DNS zone  Domain Name System (DNS)

**Secure/Multipurpose Internet Mail Extensions (S/MIME)**
An extension of MIME that supports secure mail. It enables message originators to digitally sign e-mail messages to provide proof of message origin and data integrity. It also enables messages to be transmitted in encrypted format to provide confidential communications.

See also: public key encryption

**security**
On a network, protection of a computer system and its data from harm or loss, implemented especially so that only authorized users can gain access to shared files.

See also: authorization

**Security Accounts Manager (SAM)**
A Windows service used during the logon process. SAM maintains user account information, including groups to which a user belongs.

See also: group  service  user account

**security association (SA)**
A combination of identifiers, which together define Internet Protocol security (IPSec) that protects communication between sender and receiver. An SA is identified by the combination of a Security Parameters Index (SPI), destination IP address, and security protocol (Authentication Header (AH) or Encapsulating Security Payload (ESP)). An SA must be negotiated before secured data can be sent.

See also: Authentication Header (AH)  Encapsulating Security Payload (ESP)  Internet Protocol security (IPsec)  IP address  Security Parameters Index (SPI)

**security context**
The security attributes or rules that are currently in effect. For example, the rules that govern what a user can do to a protected object are determined by security information in the user`s access token and in the object`s security descriptor. Together, the access token and the security descriptor form a security context for the user`s actions on the object.

See also: access token  object  security descriptor

**security descriptor**
A data structure that contains security information associated with a protected object. Security descriptors include information about who owns the object, who can access it and in what way, and what types of access are audited.

See also: discretionary access control list (DACL)  group  object  permission  security context  system access control list (SACL)

**security event**
An event that is logged in the security log in Event Viewer. All events that are logged in the security log are auditing events, and they can be divided into the following event categories: account logon, account management, directory service access, logon, object access, policy change, privilege use, process tracking, or system.

See also: auditing  event  Event Viewer  security log

**security filtering**
A method of applying the settings of a Group Policy object (GPO) based on security principals.

See also: Group Policy object (GPO)  security principal

**security group**
A group that can be listed in discretionary access control lists (DACLs) used to define permissions on resources and objects. A security group can also be used as an e-mail entity. Sending an e-mail message to the group sends the message to all the members of the group.

See also: discretionary access control list (DACL)  group

**security host**
An authentication device, supplemental to standard Windows and remote access server security, that verifies whether a caller from a remote client is authorized to connect to the remote access server.

See also: authentication  remote access  remote access server

**security ID (SID)**
A data structure of variable length that identifies user, group, and computer accounts. Every account on a network is issued a unique SID when the account is first created. Internal processes in Windows refer to an account`s SID rather than the account's user or group name.

See also: group account  user account  user name

**security log**
An event log containing information about security events that are specified in the audit policy.

See also: audit policy  event

**security package**
The software implementation of a security protocol. Security packages are contained in security support provider dynamic-link libraries (DLLs) or in security support provider/authentication package DLLs.

**Security Parameters Index (SPI)**
A unique, identifying value in the security association (SA) used to distinguish among multiple SAs existing at the receiving computer.

See also: security association (SA)

**security principal**
In Windows-based computers, an account (such as a user, security group, device, or computer) that can be granted or denied access to resources.

See also: group  resource  security ID (SID)  service

**security principal name**
A name that uniquely identifies a user, group, or computer within a single domain. This name is not guaranteed to be unique across domains.

See also: domain  group  security principal

**security protocol**
A specification that defines security-related data objects and rules about how the objects are used to maintain security on a computer system.

**security support provider (SSP)**
A dynamic-link library (DLL) that implements the Security Support Provider Interface (SSPI) by making one or more security packages available to applications. Each security package provides mappings between an application`s SSPI function calls and an actual security model`s functions. Security packages support security protocols such as Kerberos and NTLM.

**Security Support Provider Interface (SSPI)**
A common interface between transport-level applications, such as Microsoft Remote Procedure Call (RPC), and security support providers (SSPs), such as Windows Distributed Security. SSPI allows a transport application to call one of the SSPs to obtain an authenticated connection. These calls do not require extensive knowledge of the security protocol`s details.

**security template**
A physical file representation of a security configuration that can be applied to a local computer or imported to a Group Policy object (GPO) in Active Directory. When you import a security template to a GPO, Group Policy processes the template and makes the corresponding changes to the members of that GPO, which can be users or computers.

See also: Active Directory  Group Policy  Group Policy object (GPO)

**security token**
A cryptographically signed data unit that expresses one or more claims.

**security token service**
A Web service that issues security tokens. A security token service makes assertions based on evidence that it trusts to whoever trusts it. To communicate trust, this service requires proof, such as a security token or set of security tokens, and it issues a security token with its own trust statement. (Note that for some security token formats, this can simply be a reissuance or cosignature.) In Active Directory Federation Services (ADFS), the Federation Service is a security token service.

**See Files**

The Macintosh-style privilege that allows you to see any files in the folders for which you have this privilege. When Services for Macintosh translates Macintosh-style privileges into Windows permissions, you are granted Read permission for any folders for which you have the See Files and See Folders privileges.

See also: See Folders

**See Folders**
The Macintosh-style privilege that allows you to see folders for which you have this privilege but not any files those folders might contain. When Services for Macintosh translates Macintosh-style privileges into Windows permissions, you are granted Read permission for any folders for which you have the See Files and See Folders privileges.

See also: See Files

**seed router**
A router that initializes and broadcasts network numbers and zones about one or more physical AppleTalk networks. Servers that are running the AppleTalk protocol can function as seed routers. You can also use hardware routers from other companies as seed routers.

See also: router  zone

**sending member**
Between a pair of replicating servers, the member that is sending replicated content over a connection.

**Serial Line Internet Protocol (SLIP)**
An older industry standard that is part of Windows remote access client to ensure interoperability with other remote access software.

See also: remote access

**serial port**
An interface on the computer that allows asynchronous transmission of data characters one bit at a time. Also called a communication port or COM port.

See also: communication port  port

**server**
A computer that provides shared resources, such as files or printers, to network users.

See also: client  shared resource

**server application**
Application software running on a cluster node, regardless of whether it does service registration.

See also: node

**server cluster**
A group of computers, known as nodes, working together as a single system to ensure that mission-critical applications and resources remain available to clients. A server cluster presents the appearance of a single server to a client.

See also: cluster  node

**server farm**
In Active Directory Federation Services (ADFS), a collection of load-balanced federation servers, federation server proxies, or Web servers hosting the ADFS Web Agent.

**Server for NIS**
A feature of Windows that enables a Windows–based Active Directory domain controller to administer Network Information Service (NIS) networks.

**Server Message Block (SMB)**
A file-sharing protocol designed to allow networked computers to transparently access files that reside on remote systems over a variety of networks. The SMB protocol defines a series of commands that pass information between computers. SMB uses four message types: session control, file, printer, and message.

See also: protocol

**Server Operators group**

A group whose members can manage all domain controllers in a single domain. This group does not exist on workstations, stand-alone servers, or member servers. Administrative tasks that can be performed by members of this group include logging on locally, creating and deleting network shared resources, starting and stopping services, backing up and restoring files, formatting the hard disk of the computer, and shutting down the computer.

See also: domain  domain controller  group  member server  service  stand-alone server

**server zone**
The AppleTalk zone on which a server appears. On a Phase 2 network, a server appears in the default zone of the server`s default network.

See also: server  zone

**Server-Gated Cryptography (SGC)**
An extension of Secure Sockets Layer (SSL) that enables organizations, such as financial institutions, that have export versions of Internet Information Services (IIS) to use strong encryption (for example, 128-bit encryption).

See also: cryptography  Internet Information Services (IIS)  Secure Sockets Layer (SSL)

**service**
A program, routine, or process that performs a specific system function to support other programs, particularly at a low (close to the hardware) level. When services are provided over a network, they can be published in Active Directory, facilitating service-centric administration and usage. Some examples of services are the Security Accounts Manager service, File Replication Service (FRS), and Routing and Remote Access service.

See also: File Replication service (FRS)  Security Accounts Manager (SAM)

**service (SRV) resource record**
A DNS resource record used to identify computers that host specific services, specified in RFC 2782. SRV resource records are used to locate domain controllers for Active Directory.

See also: Active Directory  domain controller  Domain Name System (DNS)  Request for Comments (RFC)  resource record (RR)

**Service Advertising Protocol (SAP)**
A NetWare protocol used to identify the services and addresses of servers attached to the network. When a server starts, it uses the protocol to advertise its service. When the same server goes offline, it uses the protocol to announce that it is no longer available. NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink) uses SAP to locate NetWare servers and services.

See also: NWLink IPX/SPX/NetBIOS Compatible Transport Protocol (NWLink)  protocol

**Service Pack**
A software upgrade to an existing software distribution that contains updated files consisting of patches and hot fixes.

**Service Principal Name (SPN)**
A way of referring to a service principal. SPN structures generally follow Internet Engineering Task Force (IETF) naming conventions, and they often include the name of the computer on which the service is running. SPNs may be used to request Kerberos tickets, and they are required for mutual authentication.

See also: authentication  Internet Engineering Task Force (IETF)  service  ticket

**Service Profile Identifier (SPID)**
An 8-digit to 14-digit number that identifies the services that you order for each B-channel. For example, when you order Primary Rate ISDN, you obtain two phone numbers and two SPIDs from your Integrated Services Digital Network (ISDN) provider. Typical ISDN adapters cannot operate without configuring SPIDs.

See also: B-channel  Integrated Services Digital Network (ISDN)  service

**Service Provider Interface**
Calling conventions that back-end services use to make themselves accessible to front-end applications.

See also: service

**service ticket**
A ticket issued by the Kerberos V5 ticket-granting service (TGS) that allows a user to authenticate to a specific service in the domain.

See also: Kerberos V5 authentication protocol ticket ticket-granting service (TGS)

**service-centric**
A service model that does not require that a specific computer be used to complete a task. In a service-centric environment, users need to know only what service they want to use. They do not need to know what computer that service resides on to use it successfully.

See also: machine-centric service

**session**
A logical connection created between two hosts to exchange data. Typically, sessions use sequencing and acknowledgments to send data reliably. In the context of load balancing TCP/IP traffic, a set of client requests directed to a server. These requests can be invoked with multiple, possibly concurrent, TCP connections. The server program sometimes maintains state information between requests. To preserve access to the server state, Network Load Balancing needs to direct all requests within a session to the same cluster host when load balancing.

See also: client request host load balancing Network Load Balancing server Transmission Control Protocol/Internet Protocol (TCP/IP)

**session concentration**
For Message Queuing, a feature that typically reduces network bandwidth within a site and the number of sessions between sites. Specially configured Message Queuing servers with routing services provide session concentration.

See also: bandwidth Message Queuing Message Queuing server routing services session site

**Session Description Protocol (SDP)**
A protocol that Telephony API (TAPI) uses to advertise Internet Protocol (IP) multicast conferences. This protocol describes multimedia sessions for the purposes of session announcement, session invitation, and other forms of session initiation. SDP descriptors are stored in Active Directory. SDP is described in RFC 2327 of the Internet Engineering Task Force (IETF).

See also: Active Directory Internet Engineering Task Force (IETF) protocol Request for Comments (RFC) Telephony API (TAPI)

**session key**
In Internet Protocol security (IPSec), a value that is used in combination with an algorithm to encrypt or decrypt data that is transferred between computers. A session key is created for every pair of computers to provide enhanced security on computers that have multiple simultaneous active sessions.

See also: algorithm encryption Internet Protocol security (IPsec) key session

**set-by-caller callback**
In Network Connections, a form of callback in which the user supplies the telephone number that the remote access server uses for callback. This setting spares the user any long-distance telephone charges.

See also: Network Connections remote access server

**Setup**
The program that installs Windows. Also known as unattended installation, Winnt32.exe, and Winnt.exe.

**SGC**
See other term: Server-Gated Cryptography (SGC)

**SHA-1**
See other term: Secure Hash Algorithm (SHA-1)

**share**
To make resources, such as folders and printers, available to others.

See also: resource shared folder

**shared folder**
A folder on another computer that has been made available for other people to use on the network.

See also: share

**shared folder permissions**
Permissions that restrict a shared resource`s availability over the network to only certain users.

See also: permission  shared resource

**shared printer**
A printer that receives input from more than one computer. For example, a printer attached to another computer on the network can be shared so that it is available for you to use. Also called a network printer.

See also: printer  share

**shared resource**
Any device, data, or program that is used by more than one program or one other device. For Windows, shared resource refers to any resource that is made available to network users, such as folders, files, printers, and named pipes. Shared resource can also refer to a resource on a server that is available to network users.

See also: device  named pipe  resource  server  share

**shortcut**
A link to any item accessible on your computer or on a network, such as a program, file, folder, disk drive, printer, or another computer. You can put shortcuts in various areas, such as on the desktop, on the Start menu, or in specific folders.

See also: desktop

**shortcut trust**
A trust that is manually created between two domains in the same forest. The purpose of a shortcut trust is to optimize the interdomain authentication process by shortening the trust path. Shortcut trusts are transitive and can be one-way or two-way.

See also: Active Directory  domain  forest  one-way trust  trust path  trust relationship  two-way trust

**SID**
See other term: security ID (SID)

**signaling protocol**
A protocol that uses signaling to request, route, accept, and tear down virtual circuits.

See also: protocol

**signed driver**
A driver that meets the criteria of the Windows Logo Program. For performance and stability, Microsoft strongly recommends that you use only signed drivers for new or updated drivers.

See also: unsigned driver  Windows Logo Program

**Simple Authentication and Security Layer (SASL)**
An open framework, described in Request for Comments (RFC) 2222, for adding authentication support to connection-based protocols.

**Simple Mail Transfer Protocol (SMTP)**
A member of the TCP/IP suite of protocols that governs the exchange of electronic mail between message transfer agents.

See also: protocol  Transmission Control Protocol/Internet Protocol (TCP/IP)

**Simple Network Management Protocol (SNMP)**
A network protocol used to manage TCP/IP networks. In Windows, the SNMP service is used to provide status information about a host on a TCP/IP network.

See also: agent  protocol  service  Transmission Control Protocol/Internet Protocol (TCP/IP)

**Simple Object Access Protocol (SOAP)**
An XML/HTTP-based protocol for platform-independent access to objects and services on the Web. SOAP defines a message format in XML that travels over the Internet using Hypertext Transfer Protocol (HTTP). By using existing Web protocols (HTTP) and languages (XML), SOAP runs over the existing Internet infrastructure without being tied to any operating system, language, or object model.

See also: Extensible Markup Language (XML)  Hypertext Transfer Protocol (HTTP)  object  service

**Simple TCP/IP Services**
Four TCP/IP services: Character Generator, Daytime Discard, Echo, and Quote of the Day.

See also: service  Transmission Control Protocol/Internet Protocol (TCP/IP)

**simple volume**
A dynamic volume made up of disk space from a single dynamic disk. A simple volume can consist of a single region on a disk or multiple regions of the same disk that are linked together. If the simple volume is not a system volume or boot volume, you can extend it within the same disk or onto additional disks. If you extend a simple volume across multiple disks, it becomes a spanned volume. You can create simple volumes only on dynamic disks. Simple volumes are not fault tolerant, but you can mirror them to create mirrored volumes on computers running the Windows 2000 Server or Windows Server 2003 families of operating systems.

See also: dynamic disk  dynamic volume  fault tolerance  mirrored volume  spanned volume  volume

**single affinity**
Specifies that Network Load Balancing should direct multiple requests from the same client IP address to the same cluster host. This is the default setting for affinity.

See also: affinity  Class C affinity  Network Load Balancing

**Single Instance Store (SIS)**
A component that saves disk space on the server by maintaining a single physical copy of all identical files found. If SIS finds a duplicate file on the server, it copies the original file into the SIS store and leaves a link where the original resided. This technology is used only with Remote Installation Services.

See also: Remote Installation Services (RIS)

**single node server cluster**
A cluster configuration that has one node and that can be configured with or without external cluster storage devices. For a single node cluster without an external cluster storage device, the local disk is configured as the cluster storage device. There are advantages and limitations for each cluster configuration (single node server cluster, single quorum device server cluster, and majority node set server cluster).

See also: cluster  cluster storage  majority node set server cluster  node  single quorum device server cluster

**single quorum device server cluster**
A cluster configuration that has two or more nodes and that is configured so that every node is attached to one or more cluster storage device. The cluster configuration data is stored on a single cluster storage device. There are advantages and limitations for each cluster configuration (single node server cluster, single quorum device server cluster, and majority node set server cluster).

See also: cluster  cluster storage  majority node set server cluster  node  single node server cluster

**single sign-on**
A process that enables a user with a domain account to log on to a network once, using a password or smart card, and to gain access to any computer in the domain.

See also: domain  smart card

**single sign-on (SSO)**
An optimization of the authentication sequence to remove the burden of repeated logon actions by an end user

**single sign-on daemon (SSOD)**
A daemon on a Windows-based computer that receives an encrypted password and decrypts it before requesting a password change on the UNIX host.

**SIS**
See other term: Single Instance Store (SIS)

**site**
One or more well-connected (highly reliable and fast) TCP/IP subnets. A site allows administrators to configure Active Directory access and replication topology to take advantage of the physical network.

See also: Active Directory  replication topology  subnet  Transmission Control Protocol/Internet Protocol (TCP/IP)

**site link**
An Active Directory object that represents a set of sites that can communicate at uniform cost through some intersite transport. For Internet Protocol (IP) transport, a typical site link connects just two sites and corresponds to an actual wide area network (WAN) link. An IP site link connecting more than two sites might correspond to an asynchronous transfer mode (ATM) backbone connecting more than two clusters of buildings on a large campus or several offices in a large metropolitan area

connected via leased lines and IP routers.

See also: Active Directory  asynchronous transfer mode (ATM)  Internet Protocol (IP)  site  site link bridge  wide area network (WAN)

**site link bridge**
An Active Directory object that represents a set of site links, all of whose sites can communicate via some transport. Typically, a site link bridge corresponds to a router (or a set of routers) in an Internet Protocol (IP) network. By default, the Knowledge Consistency Checker (KCC) may form a route through any and all site links in a transitive manner. If this behavior is turned off, each site link represents its own distinct and isolated network. Sets of site links that can be treated as a single route are expressed through a site link bridge. Each bridge represents an isolated communication environment for network traffic.

See also: Active Directory  Internet Protocol (IP)  Knowledge Consistency Checker (KCC)  site  site link

**sleep mode**
An energy-saving mode in which the Windows operating system shuts down all unnecessary components, such as the display screen and the disk drive. The computer returns to its former operating status when it is awakened.

See also: wake-on-LAN

**SLIP**
See other term: Serial Line Internet Protocol (SLIP)

**small computer system interface (SCSI)**
A standard high-speed parallel interface defined by the American National Standards Institute (ANSI). A SCSI interface is used for connecting microcomputers to peripheral devices, such as hard disks and printers, and to other computers and local area networks (LANs).

See also: device  local area network (LAN)

**smart card**
A credit card–sized device that is used with an access code to enable certificate-based authentication and single sign-on to the enterprise. Smart cards securely store certificates, public and private keys, passwords, and other types of personal information. A smart card reader attached to the computer reads the smart card.

See also: authentication  single sign-on  smart card reader

**smart card reader**
A device that is installed in computers to enable the use of smart cards for enhanced security features.

See also: smart card

**SMS**
See other term: Systems Management Server (SMS)

**SMTP**
See other term: Simple Mail Transfer Protocol (SMTP)

**snap-in**
A type of tool that you can add to a console supported by Microsoft Management Console (MMC). A stand-alone snap-in can be added by itself; an extension snap-in can be added only to extend the function of another snap-in.

See also: Microsoft Management Console (MMC)

**sniffer**
An application or device that can read, monitor, and capture network data exchanges and read network packets. If the packets are not encrypted, a sniffer provides a full view of the data inside the packet.

See also: device  packet

**SNMP**
See other term: Simple Network Management Protocol (SNMP)

**socket**
An identifier for a particular service on a particular node on a network. The socket consists of a node address and a port number, which identifies the service. For example, port 80 on an Internet node indicates a Web server. There are two kinds of sockets: streams (bidirectional) and datagrams.

See also: datagram  node  port  process  raw socket

**software decoder**
A type of digital video disc (DVD) decoder that allows a DVD drive to display movies on your computer screen. A software decoder uses only software to display movies.

See also: DVD decoder  DVD drive  hardware decoder

**software restriction policies**
A collection of policy settings that define what software can run on a computer, based on the default security level for a Group Policy object (GPO). Exceptions to that default security level can then be defined by certificate rules, hash rules, path rules, registry path rules, and Internet zone rules.

See also: default security level  Group Policy  Group Policy object (GPO)  software restriction policies rule

**software restriction policies rule**
A rule that creates an exception to the default security level that is defined by software restriction policies. The following types of rules can be created: certificate rules, which recognize software that is digitally signed by an Authenticode software publisher certificate; hash rules, which recognize specific software based on a hash of the software; path rules, which recognize software based on the location in which the software is stored; registry path rules, which recognize software based on the location of the software as it is stored in the registry; and Internet zone rules, which recognize software based on the zone of the Internet from which the software is downloaded.

See also: default security level  hash  path  registry  software restriction policies

**SOM**
See other term: scope of management (SOM)

**source document**
The document where a linked or embedded object was originally created.

See also: embedded object  linked object

**source journaling**
For Message Queuing, the process of storing a copy of an outgoing message. Source journaling is configured on a message basis, and it is set by the sending application. When source journaling is enabled, a copy of the message is put in the source journal of the source computer when the message arrives at the destination (target) queue.

See also: journal  Message Queuing  target journaling

**spanned volume**
A dynamic volume consisting of disk space on more than one physical disk. You can increase the size of a spanned volume by extending it onto additional dynamic disks. You can create spanned volumes only on dynamic disks. Spanned volumes are not fault tolerant and cannot be mirrored.

See also: dynamic disk  dynamic volume  fault tolerance  mirrored volume  simple volume  volume

**sparse file**
A file that is handled in a way that requires much less disk space than would otherwise be needed. Sparse support allows an application to create very large files without committing disk space for those regions of the file that contain only zeros. For example, you can use sparse support to work with a 42-GB file in which you need to write data only to the first 64 KB (the rest of the file is zeroed).

**special permissions**
On NTFS volumes, a custom set of permissions. You can customize permissions on files and directories by selecting the individual components of the standard sets of permissions.

See also: NTFS file system  permission  volume

**SPID**
See other term: Service Profile Identifier (SPID)

**split horizon**
A route-advertising algorithm that prevents the advertising of routes in the same direction in which they were learned. Split horizon helps prevent routing loops.

See also: poison reverse  routing

**splitting**
A Briefcase command that separates the copy of the file inside Briefcase from the copy outside Briefcase.

**spooling**
A process on a server in which print documents are stored on a disk until a printer is ready to process them. A spooler accepts each document from each client, stores it, then sends it to a printer when the printer is ready.

See also: print spooler

**SQL**
See other term: structured query language (SQL)

**SSL**
See other term: Secure Sockets Layer (SSL)

**SSP**
See other term: security support provider (SSP)

**SSPI**
See other term: Security Support Provider Interface (SSPI)

**stand-alone certification authority**
A certification authority (CA) that is not integrated with Active Directory.

See also: Active Directory  certification authority (CA)

**stand-alone root**
A DFS namespace, the configuration information for which is stored locally on the host server. The path to access the root or a link starts with the host server name. A stand-alone root has only one root target. There is no root-level fault tolerance. Therefore, when the root target is unavailable, the entire DFS namespace is inaccessible.

See also: DFS namespace  root target

**stand-alone server**
A server that runs Windows 2000 or Windows Server 2003, but does not participate in a domain. A stand-alone server has only its own database of users, and it processes logon requests by itself. A stand-alone server does not share account information with other computers and cannot provide access to domain accounts, but it can participate in a workgroup.

See also: domain  member server  workgroup

**standard NIS map**
In Network Information Service (NIS), a map that consists of aliases, bootparams, ethers, hosts, group, netgroup, netid, netmasks, networks, passwd, protocols, rpc, services, pservers, and shadow; all other maps are nonstandard.

**start-of-authority (SOA) resource record**
A record that indicates the starting point or original point of authority for information stored in a zone. The SOA resource record (RR) is the first RR created when adding a new zone. It also contains several parameters used by other computers that use DNS to determine how long they will use information for the zone and how often updates are required.

See also: authoritative  Domain Name System (DNS)  resource record (RR)  zone

**Startrom.com**
A program used by Remote Installation Services (RIS) that is the first file downloaded to the client using Trivial File Transfer Protocol (TFTP). Startrom.com is a small program that displays the Press F12 for Network Service Boot prompt. If F12 is pressed within three seconds, the Client Installation Wizard (OSChooser) is downloaded to begin the remote installation process.

See also: Remote Installation Services (RIS)  Trivial File Transfer Protocol (TFTP)

**startup environment**
In dual-boot or multiple-boot systems, the configuration settings that specify which system to start and how each system should be started.

See also: dual boot  multiple boot

**stateless**

As related to servers, not involving the update of a server-side database based on a client request. As related to the handling of files, the content of the file is not modified or noticed. For Web servers, a stateless client request, which members of a Network Load Balancing cluster can process, is one that returns a static Web page to the client.

See also: Network Load Balancing cluster

**static dialog box**
A scripted dialog box between the client computer and an intermediary device. This kind of dialog box requires no response from the user.

See also: client

**static load balancing**
The process of manually moving a group between nodes to balance the load across the nodes.

See also: group  node

**static routes**
Routes in a routing table that are permanent until changed by a network administrator or by an automatically scheduled auto-static update.

See also: router  routing

**status area**
See other term: notification area

**status bar**
A line of information related to the current program. The status bar is usually located at the bottom of a window. Not all windows have a status bar.

**Stop error**
A serious error that affects the operating system and that could place data at risk. The operating system generates an obvious message, a screen with the Stop error, rather than continuing on and possibly corrupting data. Also called a fatal system error.

See also: Stop screen

**Stop screen**
A blue character-mode screen that appears when Windows encounters a condition that compromises safe system operation and the system stops. Commonly referred to as a blue screen or bug check.

See also: Stop error

**storage report**
A reporting tool that alerts administrators to current disk use trends, as well as to attempts by certain users or groups to save unauthorized files.

**storage subsystem**
A stand-alone hardware appliance that hosts one or more storage devices (such as disk drives, tape drives, optical drives), and is peripheral to the server or servers that control access to it. Storage subsystems are used to create centralized data repositories, while freeing computing and storage resources on servers.

**storage-class resource**
A required dependency for many resource types. Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, provide one storage-class resource: Physical Disk. However, your vendor or reseller might supply other storage-class resource types.

See also: dependency  resource types

**strict RFC checking**
For DNS, a form of domain name checking that examines characters used in DNS names for compliance with DNS naming requirements and valid character usage as specified in RFC 1123, Requirements for Internet Hosts - Applications and Support. For strict RFC compliance, DNS domain names will use name labels made up only of valid uppercase and lowercase letters, number characters, and hyphens (A through Z, a through z, 0 through 9, -), separated by periods.

See also: domain name  Domain Name System (DNS)  label  loose name checking  Request for Comments (RFC)

**string**

A group of characters or character bytes handled as a single entity. Computer programs use strings to store and transmit data and commands. Most programming languages consider strings (such as 2674:gstmn) as distinct from numeric values (such as 470924).

See also: transmitting station ID (TSID) string

**striped volume**
A dynamic volume that stores data in stripes on two or more physical disks. Data in a striped volume is allocated alternately and evenly (in stripes) across the disks. Striped volumes offer the best performance of all the volumes that are available in Windows, but they do not provide fault tolerance. If a disk in a striped volume fails, the data in the entire volume is lost. You can create striped volumes only on dynamic disks. Striped volumes cannot be mirrored or extended.

See also: dynamic disk  dynamic volume  fault tolerance  mirrored volume  volume

**strong password**
A password that cannot be easily guessed or cracked. A strong password is at least six characters long, does not contain all or part of the user's account name, and contains at least three of the four following categories of characters: uppercase characters, lowercase characters, numbers, and symbols found on the keyboard (such as !, @, #).

See also: password  password policy  weak password

**structured query language (SQL)**
A widely accepted standard database sublanguage used in querying, updating, and managing relational databases.

**stub area**
An Open Shortest Path First (OSPF) area that does not advertise individual external networks. A stub area blocks external routes and therefore reduces the amount of memory required on the internal routers located in the stub area. To keep the topology database size small, routing to all external networks in a stub area is done through a default route (destination 0.0.0.0 with the subnet mask of 0.0.0.0). In OSPF, any destination that you cannot reach through an intra-area or inter-area router is reachable through the default route.

See also: Open Shortest Path First (OSPF)  routing

**stub zone**
A copy of a zone that contains only the resource records required to identify the authoritative DNS servers for that zone. A DNS server that hosts a parent zone and a stub zone for one of the parent zone`s delegated child zones can receive updates from the authoritative DNS servers for the child zone.

See also: authoritative  DNS server  resource record (RR)  zone

**subdomain**
A DNS domain located directly beneath another domain name (the parent domain) in the namespace tree. For example, example.microsoft.com would be a subdomain of the domain microsoft.com. Also called child domain.

See also: child domain  domain  Domain Name System (DNS)  parent domain

**subject**
In public key cryptography, an entity that requests or holds a certificate. A subject can be a user, a computer, or any other device capable of requesting or using a certificate.

See also: public key cryptography

**subkey**
An element of the registry that contains entries or other subkeys. A tier of the registry that is immediately below a key or a subtree (if the subtree has no keys).

See also: descendent key  entry  key  registry  subtree

**subnet**
A subdivision of an Internet Protocol (IP) network. Each subnet has its own unique subnetted network ID.

See also: Internet Protocol (IP)  network ID

**subnet bandwidth management (SBM)**
An Internet Engineering Task Force (IETF) standard that enables administrative control at the subnet level.

See also: Internet Engineering Task Force (IETF)

**subordinate server**
A server that does not attempt to resolve queries on its own. Instead, it sends all queries to forwarders. Formerly known as slave server.

**subtree**
Any node within a tree, along with any selection of connected descendant nodes. The highest level of the registry (for example, HKEY_LOCAL_MACHINE).

See also: key  node  registry  subkey

**superscope**
An administrative grouping feature that supports a DHCP server`s ability to use more than scope for each physical interface and subnet. Superscopes are useful under the following conditions: If more DHCP clients must be added to a network than were originally planned, if an Internet Protocol (IP) network is renumbered, or if two or more DHCP servers are configured to provide scope redundancy and fault-tolerant design DHCP service for a single subnet. Each superscope can contain one or more member scopes (also known as child scopes).

See also: DHCP client  DHCP server  fault tolerance  scope  subnet

**SVC**
See other term: switched virtual circuit (SVC)

**switch type**
The type of interface to which your Integrated Services Digital Network (ISDN) device is being attached. Also called switch.

See also: B-channel  D-channel  Integrated Services Digital Network (ISDN)

**switched virtual circuit (SVC)**
A connection established dynamically between devices on an asynchronous transfer mode (ATM) network through the use of signaling.

See also: asynchronous transfer mode (ATM)

**switching hub**
A central network device (multiport hub) that forwards packets to specific ports rather than, as in conventional hubs, broadcasting every packet to every port. In this way, the connections between ports deliver the full bandwidth available.

See also: hub  packet  port

**symmetric encryption**
An encryption algorithm that requires the same secret key to be used for both encryption and decryption. Because of its speed, symmetric encryption is typically used when a message sender needs to encrypt large amounts of data. Also called secret key encryption.

See also: public key encryption

**symmetric key**
A single key that is used with symmetric encryption algorithms for both encryption and decryption.

See also: encryption  key  symmetric encryption

**synchronous**
Each application or command runs in the order listed, and each item must finish before the next command is run.

**syntax**
The order in which a command must be typed, and the elements that follow the command.

See also: Command Prompt window

**Syskey**
A tool used to configure the startup key, a random, 128-bit, symmetric cryptographic key created at system startup and used to encrypt all of the user`s symmetric cryptographic keys.

See also: encryption  symmetric key

**system access control list (SACL)**
The part of an object`s security descriptor that specifies which events are to be audited per user or group. Examples of auditing

events are file access, logon attempts, and system shutdowns.

See also: auditing  discretionary access control list (DACL)  event  object  security descriptor

**system area network (SAN)**
A network configuration, usually on a separate Internet Protocol (IP) subnet, that gives data a direct path to system hardware.

See also: Internet Protocol (IP)  subnet

**system default profile**
The user profile that is loaded when Windows is running and no user is logged on. When the Begin Logon dialog box is visible, the system default profile is loaded.

See also: user profile

**system disk**
A disk that contains the MS-DOS system files necessary to start MS-DOS.

See also: MS-DOS (Microsoft Disk Operating System)

**system files**
Files used by Windows to load, configure, and run the operating system. Generally, system files must never be deleted or moved.

**System menu**
A menu that contains commands you can use to manipulate a window or close a program. You click the program icon at the left of the title bar to open the System menu.

**system partition**
The partition that contains the hardware-specific files needed to load Windows (for example, Ntldr, Osloader, Boot.ini, Ntdetect.com). The system partition can be, but does not have to be, the same as the boot partition.

See also: boot partition  partition

**System Policy**
A Windows NT 4.0-style policy based on registry settings made using Poledit.exe, the System Policy Editor.

See also: policy  registry

**System State**
In Backup, a collection of system-specific data maintained by the operating system that must be backed up as a unit. It is not a backup of the entire system. The System State data includes the registry, COM+ Class Registration database, system files, boot files, and files under Windows File Protection. For servers, the System State data also includes the Certificate Services database (if the server is a certificate server). If the server is a domain controller, the System State data also includes the Active Directory database and the SYSVOL directory. If the server is a node in a cluster, it includes the Cluster database information. The IIS Metabase is included if Internet Information Services (IIS) is installed.

See also: Active Directory  cluster  domain controller  Internet Information Services (IIS)  node  registry  SYSVOL

**system variables**
Storage locations for data that are defined by the operating system and that are the same regardless of who is logged on at the computer. (Users who are also members of the Administrators group can add new variables or change the values.)

See also: Administrators group

**system volume**
The volume that contains the hardware-specific files that are needed to load Windows on x86-based computers with a basic input/output system (BIOS). The system volume can be, but does not have to be, the same volume as the boot volume.

See also: basic input/output system (BIOS)  boot volume  volume  x86

**systemroot**
The path and folder name where the Windows system files are located. Typically, this is C:\Windows, although you can designate a different drive or folder when you install Windows. You can use the value %systemroot% to replace the actual location of the folder that contains the Windows system files. To identify your systemroot folder, click Start, click Run, type %systemroot%, and then click OK.

**Systems Management Server (SMS)**
A Microsoft product that includes inventory collection, software deployment, and diagnostic tools. SMS automates the task of upgrading software, allows remote problem solving, provides asset management information, and monitors software usage, computers, and networks.

**SYSVOL**
A shared directory that stores the server copy of the domain`s public files, which are replicated among all domain controllers in the domain.

See also: directory  domain  domain controller

# Windows Server 2003 Glossary - T

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - T**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**T.120**
The International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T) standard for multipoint data conferencing. T.120 provides the protocols for establishing and managing data flow, connections, and conferences. Support for T.120 enables data transfer from conferencing applications, such as file transfer and application sharing, to operate in conjunction with H.323 connections.

See also: H.323  International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T)

**tag**
For Indexing Service, one or more terms that identify an element in a query, such as weight, phrase, property, or regular expression. For example, the tag {prop name=created} specifies the Created property in a query.

See also: property  query

**TAPI**
See other term: Telephony API (TAPI)

**TAPI server security file**
A file that stores information about the telephony devices for each installed telephony service provider, including the addresses that are assigned to each device and the users who are authorized to use each device.

See also: Telephony API (TAPI)

**target**
A Universal Naming Convention (UNC) path that corresponds to a namespace root or folder.

See also: DFS link  DFS root  Distributed File System (DFS)

**target journaling**
For Message Queuing, the process of storing copies of incoming messages. Target journaling is configured on a queue basis. When target journaling is enabled, a copy of each incoming message is placed in the target journal when the message is retrieved (read) from the destination queue.

See also: journal  Message Queuing  source journaling

**target priority**
A setting that specifies whether a specific target in a referral should be listed first or last, or first or last among targets of equal cost.

**Task Manager**
A tool that provides information about programs and processes running on the computer. Using Task Manager, you can end or run programs, end processes, and display a dynamic overview of your computer`s performance.

See also: process  program

**taskbar**
The bar that contains the Start button and appears by default at the bottom of the desktop. You can click the taskbar buttons to switch between programs. You can also hide the taskbar, move it to the sides or top of the desktop, and customize it in other ways.

See also: desktop  notification area

**taskpad**
A details pane view that displays shortcuts to commands that are available for the selected snap-in in the console tree. Users

can run a command by clicking a task.

See also: console tree  details pane  snap-in

**TCP/IP**
See other term: Transmission Control Protocol/Internet Protocol (TCP/IP)

**TDI**
See other term: Transport Driver Interface (TDI)

**teaming network adapters**
A configuration method that provides fault tolerance by grouping network adapters on multiple ports to a single physical network segment. For example, if connections through one port fail (due to failure of the adapter, cable, switch port, or switch), another port is activated automatically. Teaming network adapters work transparently to the operating system and other devices on the network.

See also: device  fault tolerance  network adapter  port

**Telephony API (TAPI)**
An application programming interface (API) used by communications programs to work with telephony and network services.

See also: application programming interface (API)  Internet Protocol (IP)  modem (modulator/demodulator)  service

**telephony switch**
A computer or electromechanical device that controls the routing and operation of a signal path.

**Telnet**
A protocol that enables an Internet user to log on to and enter commands on a remote computer linked to the Internet, as if the user were using a text-based terminal directly attached to that computer. Telnet is part of the TCP/IP suite of protocols. The term telnet also refers to the software (client or server component) that implements this protocol.

See also: protocol  Transmission Control Protocol/Internet Protocol (TCP/IP)

**temporary address**
An Internet Protocol version 6 (IPv6) address that uses an interface identifier that changes over time and that is not derived from a hardware property, such as a media access control (MAC) address. A temporary address makes it difficult to track the identity of a computer on the Internet across sessions.

See also: Internet Protocol (IP)  media access control (MAC) address

**Terminal Services**
The underlying technology that enables Remote Desktop, Remote Assistance, and Terminal Server.

See also: Terminal Services Licensing

**Terminal Services Licensing**
Software that provides the ability to register and track licenses for Terminal Services clients.

See also: Terminal Services

**terminate-and-stay-resident (TSR) program**
A program running under MS-DOS that remains loaded in memory even when it is not running, so that it can be quickly invoked for a specific task performed while any other application is operating.

See also: MS-DOS (Microsoft Disk Operating System)

**test queue**
For Message Queuing, a queue that stores sent test messages.

See also: Message Queuing  queue

**text box**
In a dialog box, a box in which you type information needed to carry out a command. The text box may be blank or may contain text when the dialog box opens.

**text mode**
The portion of Setup that uses a text-based interface.

See also: GUI mode  Setup

**TFTP**
See other term: Trivial File Transfer Protocol (TFTP)

**TGS**
See other term: ticket-granting service (TGS)

**TGT**
See other term: ticket-granting ticket (TGT)

**Thread Count**
In Task Manager, the number of threads running in a process.

See also: process  Task Manager

**thumbnail**
A miniature version of an image that is often used for quick browsing through multiple images.

**ticket**
A set of identification data for a security principal, issued by a domain controller for purposes of user authentication. Two forms of tickets in Windows are ticket-granting tickets (TGTs) and service tickets.

See also: authentication  domain controller  service ticket  ticket-granting ticket (TGT)

**ticket-granting service (TGS)**
A Kerberos V5 service provided by the Kerberos V5 Key Distribution Center (KDC) service that issues service tickets that allow users to authenticate to services in a domain.

See also: Kerberos V5 authentication protocol  Key Distribution Center (KDC)  service ticket  ticket-granting ticket (TGT)

**ticket-granting ticket (TGT)**
A credential issued to a user by the Kerberos Key Distribution Center (KDC) when the user logs on. The user must present the TGT to the KDC when requesting session tickets for services. Because a TGT is normally valid for the life of the user`s logon session, it is sometimes called a user ticket.

See also: Kerberos V5 authentication protocol  Key Distribution Center (KDC)  ticket-granting service (TGS)

**time server**
A computer that periodically synchronizes the time on all computers within a network. This ensures that the time used by network services and local functions remains accurate.

See also: server

**time slice**
A brief period of time during which a particular task is given control of the microprocessor in a time-sharing multitasking environment. A computer`s processor is allocated to an application, usually measured in milliseconds. Also called quantum.

**time stamp**
A certification specifying that a particular message existed at a specific time and date. In a digital context, trusted third parties generate a trusted time stamp for a particular message by having a time stamping service append a time value to a message and then digitally signing the result.

See also: service

**Time to Live (TTL)**
For Internet Protocol (IP), a field in the IP header of an IP packet that indicates the maximum number of links over which the packet can travel before being discarded by a router. For DNS, TTL values are used in resource records within a zone to determine how long requesting clients should cache and use this information when it appears in a query response answered by a DNS server for the zone.

See also: DNS server  Domain Name System (DNS)  Internet Protocol (IP)  packet  resource record (RR)  Transmission Control Protocol/Internet Protocol (TCP/IP)  zone

**time-out error**
A condition where an expected character is not received in time. When this condition occurs, the software assumes that the data has been lost and requests that it be resent.

**title bar**
The horizontal bar at the top of a window that contains the name of the window. On many windows, the title bar also contains the program icon, the Maximize, Minimize, and Close buttons, and the optional ? button for context-sensitive Help. To display a menu with commands such as Restore and Move, right-click the title bar.

See also: maximize  minimize

**TLS**
See other term: Transport Layer Security (TLS)

**token**
Any nonreducible textual element in data that is being parsed. For example, the use in a program of a variable name, a reserved word, or an operator. Storing tokens as short codes shortens program files and speeds execution. For networking, a unique structured data object or message that circulates continuously among the nodes of a token ring and describes the current state of the network. Before any node can send a message on the network, it must first wait to control the token.

See also: Token Ring

**Token Ring**
The Institute of Electrical and Electronics Engineers (IEEE) 802.5 standard that uses a token-passing technique for media access control (MAC). Token Ring supports media of both shielded and unshielded twisted pair wiring for data rates of 4 megabits per second (Mbps) and 16 megabits per second.

See also: client  token

**token-signing certificate**
An X.509 certificate whose associated public/private key pair is used to provide integrity for security tokens.

**tombstone**
In Active Directory, an object that is removed from the directory but not yet deleted.

See also: Active Directory  directory  object

**tombstone lifetime**
The length of time that an object lives as a tombstone in the directory before being collected as garbage.

See also: directory  object  tombstone

**top-level domains**
Domain names that are rooted hierarchically at the first tier of the domain namespace directly beneath the root (.) of the DNS namespace. On the Internet, top-level domain names such as .com and .org are used to classify and assign second-level domain names (such as microsoft.com) to individual organizations and businesses according to their organizational purpose.

See also: domain  domain name  Domain Name System (DNS)  domain namespace  root  second-level domain

**topological database**
An overview of networks and how they relate to routers. Routers in the same area have the same topological database.

See also: router

**topology**
The physical layout of computers, cables, switches, routers, and other components of a network. Topology also refers to the underlying network architecture, such as Ethernet or Token Ring. In Active Directory replication, the set of connections that domain controllers use to replicate information among themselves.

See also: Active Directory replication  domain controller

**touch-tone dialing**
A form of dialing that uses multiple-tone signaling. The user hears a series of tones (beeps) when dialing. Push-button telephones usually use touch-tone dialing.

See also: pulse dialing

**trace log**
A type of log generated when the user selects a trace data provider using the Performance tool. Trace logs differ from counter-data logs in that they measure data continuously rather than taking periodic samples.

**transaction**
For Message Queuing, the pairing of two or more actions that are performed together as a single action; the action succeeds or fails as a whole. Using Microsoft Distributed Transaction Coordinator (MS DTC) ensures that either both actions succeed or neither is executed.

See also: Message Queuing  Microsoft Distributed Transaction Coordinator (MS DTC)  transactional dead-letter queue  transactional message

**transactional dead-letter queue**
For Message Queuing, a queue that stores transactional messages that cannot reach their destination queue. Transactional dead-letter queues store failed messages on the computer on which the message expired. Messages in these queues are written to disk and are therefore recoverable.

See also: dead-letter queue  Message Queuing  queue  transaction

**transactional message**
For Message Queuing, a message that can be sent and received only from within a transaction. This type of message returns to its prior state when a transaction is terminated abruptly. A transactional message is removed from a queue only when the transaction is committed; otherwise, it remains in the queue and can be subsequently read during another transaction.

See also: Message Queuing  queue  transaction

**transitive trust**
A trust relationship that flows throughout a set of domains, such as a domain tree, and forms a relationship between a domain and all domains that trust that domain. For example, if domain A has a transitive trust with domain B, and domain B trusts domain C, then domain A trusts domain C. Transitive trusts can be one-way or two-way, and they are required for Kerberos-based authentication and Active Directory replication.

See also: Active Directory  Active Directory replication  domain tree  Kerberos V5 authentication protocol  nontransitive trust  one-way trust  trust relationship  two-way trust

**Transmission Control Protocol/Internet Protocol (TCP/IP)**
A set of networking protocols widely used on the Internet that provides communications across interconnected networks of computers with diverse hardware architectures and various operating systems. TCP/IP includes standards for how computers communicate and conventions for connecting networks and routing traffic.

See also: Internet Protocol (IP)  protocol

**transmitting station ID (TSID) string**
A string that specifies the transmitter subscriber ID sent by the fax machine when sending a fax to a receiving machine. This string is usually a combination of the fax or telephone number and the name of the business. It is often the same as the called subscriber ID.

See also: called subscriber ID (CSID) string  string

**Transport Driver Interface (TDI)**
A common set of routines for network layer components that communicate with the session layer of the Open Systems Interconnection (OSI) model. These routines allow software components above and below the transport layer to be mixed and matched without reprogramming.

See also: Open Systems Interconnection (OSI) reference model

**Transport Layer Security (TLS)**
A protocol that provides communications privacy and security between two applications communicating over a network. TLS encrypts communications and enables clients to authenticate servers and, optionally, servers to authenticate clients. TLS is a more secure version of the Secure Sockets Layer (SSL) protocol.

See also: authentication  protocol  Secure Sockets Layer (SSL)

**transport provider**
The driver and support files that provide transport services in a networking environment.

**trap**
In Simple Network Management Protocol (SNMP), a message sent by an agent to a management system indicating that an event has occurred on the host running the agent.

See also: agent  event  host  Simple Network Management Protocol (SNMP)

**tree root trust**
A trust that is automatically established when you add a new domain tree (the tree root domain) to an Active Directory forest (the forest root domain). Tree root trusts are transitive and two-way.

See also: Active Directory domain tree forest root domain transitive trust trust relationship two-way trust

**tree view**
A hierarchical representation of the folders, files, disk drives, and other resources connected to a computer or network. For example, Windows Explorer uses a tree view to display the resources that are attached to a computer or a network.

See also: resource

**triggered update**
A type of Routing Information Protocol (RIP) announcement that occurs when network topology changes. With triggered updates, the update announcing network topology changes is sent almost immediately rather than waiting for the next periodic announcement. Triggered updates improve the convergence time (the time it takes for a router to update its routing tables) of RIP internetworks, but at the cost of additional broadcast traffic while the triggered updates are propagated.

See also: routing Routing Information Protocol (RIP)

**Triple DES (3DES)**
An implementation of Data Encryption Standard (DES) encryption that employs three iterations of cryptographic operations on each segment of data. Each iteration uses a 56-bit key for encryption, which yields 168-bit encryption for the data. Although 3DES is slower than DES because of the additional cryptographic calculations, its protection is far stronger than DES.

See also: cryptography Data Encryption Standard (DES) encryption

**Trivial File Transfer Protocol (TFTP)**
A protocol used to download the initial files needed to begin the installation process.

See also: protocol

**Trojan horse**
A program that appears to be useful or harmless but that contains hidden code designed to exploit or damage the system on which it is run. Trojan horse programs are most commonly delivered to users through e-mail messages that misrepresent the program's purpose and function. Also called Trojan code.

See also: virus

**TrueType font**
A type of computer font that can be scaled to any size. TrueType fonts are clear and readable in all sizes and can be sent to any printer or other output device that is supported by Windows.

See also: font

**trust path**
A series of trust relationships that authentication requests must follow between domains. Domain controllers determine the trust path for all authentication requests between a domain controller in the trusting domain and a domain controller in the trusted domain.

See also: authentication domain domain controller trust relationship

**trust relationship**
A logical relationship established between domains to allow pass-through authentication, in which a trusting domain honors the logon authentications of a trusted domain. User accounts and global groups defined in a trusted domain can be given rights and permissions in a trusting domain, even though the user accounts or groups don`t exist in the trusting domain`s directory.

See also: authentication domain global group group permission user account

**trusted domain object (TDO)**
An object that represents one direction of a trust relationship. When trust relationships are established, unique TDOs are created and stored in the domain. For example, when a two-way trust is established, two TDOs are created.

See also: domain object trust relationship two-way trust

**TTL**

See other term: Time to Live (TTL)

**TTY**
A communications device that consists of a keyboard and a printer. Each keystroke on the sending machine generates a character code that is sent to the receiving machine, which prints the character. TTY is usually associated with a video display that is treated like a teletypewriter or that emulates one.

**tunnel**
A logical connection over which data is encapsulated. Typically, both encapsulation and encryption are performed, and the tunnel is a private, secure link between a remote user or host and a private network.

See also: encryption  host  tunnel server  voluntary tunnel

**tunnel server**
A server or router that terminates tunnels and forwards traffic to the hosts on the target network.

See also: host  router  server  tunnel

**tunneling protocol**
A communication standard used to manage tunnels and encapsulate private data. Data that is tunneled must also be encrypted to be a virtual private network (VPN) connection. Two commonly used tunneling protocols are the Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP).

See also: Layer Two Tunneling Protocol (L2TP)  Point-to-Point Tunneling Protocol (PPTP)  virtual private network (VPN)

**two-way trust**
A trust relationship between two domains in which both domains trust each other. For example, domain A trusts domain B, and domain B trusts domain A. All parent-child trusts are two-way.

See also: domain  one-way trust  parent-child trust  trust relationship

**Type 1 fonts**
Scalable fonts designed to work with PostScript devices.

See also: font  PostScript

# Windows Server 2003 Glossary - U

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - U**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**UAM**
See other term: user authentication module (UAM)

**UBR**
See other term: unspecified bit rate (UBR)

**UCS**
See other term: Universal Character Set (UCS)

**UCS Transformation Format 8 (UTF-8)**
A character set for protocols evolving beyond the use of ASCII. The UTF-8 protocol provides for support of extended ASCII characters and translation of UCS-2, an international 16-bit Unicode character set. UTF-8 enables a far greater range of names than can be achieved using ASCII or extended ASCII encoding for character data.

See also: American Standard Code for Information Interchange (ASCII) Unicode Universal Character Set (UCS)

**UDDI**
See other term: Universal Description, Discovery, and Integration (UDDI)

**UDP**
See other term: User Datagram Protocol (UDP)

**UDP socket**
A socket that transmits datagrams over the User Datagram Protocol (UDP).

See also: datagram socket User Datagram Protocol (UDP)

**UID**
See other term: user identifier (UID)

**unallocated space**
Available disk space that is not allocated to any volume. The type of volume that you can create on unallocated space depends on the disk type. On basic disks, you can use unallocated space to create primary or extended partitions. On dynamic disks, you can use unallocated space to create dynamic volumes.

See also: basic disk dynamic disk extended partition logical drive object partition primary partition volume

**unattended Setup**
An automated, hands-free method of installing Windows. During installation, unattended Setup uses an answer file to supply data to Setup instead of requiring that an administrator or end user interactively provide the answers.

See also: Setup

**UNC**
See other term: Universal Naming Convention (UNC)

**UNC (Universal Naming Convention) name**
The full name of a resource on a network. It conforms to the \\servername\sharename syntax, where servername is the name of the server and sharename is the name of the shared resource. UNC names of directories or files can also include the directory path under the share name, with the following syntax: \\servername\sharename\directory\filename

See also: resource

**UNI**

See other term: user-to-network interface (UNI)

**unicast**
An address that identifies a specific, globally unique host.

See also: host

**Unicode**
A character encoding standard developed by the Unicode Consortium that represents almost all of the written languages of the world. The Unicode character repertoire has multiple representation forms, including UTF-8, UTF-16, and UTF-32. Most Windows interfaces use the UTF-16 form.

See also: American Standard Code for Information Interchange (ASCII) UCS Transformation Format 8 (UTF-8) Universal Character Set (UCS)

**Uniform Resource Identifier (URI)**
A compact string of characters that identifies an abstract resource or physical resource. URIs are explained in RFC 2396. In Active Directory Federation Services (ADFS), URIs are used to uniquely identify partners and account stores.

**uninstall**
When referring to software, the act of removing program files and folders from your hard disk and removing related data from your registry so the software is no longer available. When referring to a device, the act of removing the corresponding device drivers from your hard disk and physically removing the device from your computer.

See also: device device driver install

**uninterruptible power supply (UPS)**
A device that connects a computer and a power source to ensure that electrical flow is not interrupted. UPS devices use batteries to keep the computer running for a period of time after a power failure. UPS devices usually provide protection against power surges and brownouts as well.

**Universal Asynchronous Receiver/Transmitter (UART)**
An integrated circuit (silicon chip) that is commonly used in microcomputers to provide asynchronous communication. The UART provides parallel-to-serial conversion of data to be transmitted and serial-to-parallel conversion of data received.

See also: asynchronous communication

**Universal Character Set (UCS)**
An international standard character set reference that is part of the Unicode standard. The most widely held existing version of the UCS standard is UCS-2, which specifies 16-bit character values currently accepted and recognized for use to encode most of the world`s languages.

See also: American Standard Code for Information Interchange (ASCII) UCS Transformation Format 8 (UTF-8) Unicode

**Universal Description, Discovery, and Integration (UDDI)**
An industry specification for publishing and locating information about Web services. UDDI defines a standards-based way to store and retrieve information about Web services, Web service providers, binding information, and technical interface definitions—all classified using a set of standard or custom classification schemes.

See also: service

**universal group**
A security or distribution group that can contain users, groups, and computers from any domain in its forest as members. Universal security groups can be granted rights and permissions on resources in any domain in the forest.

See also: distribution group domain forest security group

**Universal Naming Convention (UNC)**
A convention for naming files and other resources beginning with two backslashes (\), indicating that the resource exists on a network computer. UNC names conform to the \\%SERVERNAME%\%SHARENAME% syntax, where %SERVERNAME% is the server's name and %SHARENAME% is the name of the shared resource. The UNC name of a directory or file can also include the directory path after the share name, by using the following syntax: \\%SERVERNAME%\%SHARENAME\%DIRECTORY%\%FILENAME%.

**universal serial bus (USB)**
An external bus that supports Plug and Play installation. Using USB, you can connect and disconnect devices without shutting

down or restarting your computer. You can use a single USB port to connect up to 127 peripheral devices, including speakers, telephones, CD-ROM drives, joysticks, tape drives, keyboards, scanners, and cameras. A USB port is usually located on the back of your computer near the serial port or parallel port.

See also: bus  Plug and Play  port

**UNIX file permissions**
A set of permissions that describes the access rights to a UNIX file or directory for the file owner, the group owner, and other users. Every UNIX file and directory possesses a set of permissions, and each user can have any combination of read, write, or execute access.

**unsigned driver**
A driver that is not supported and that is not certified to work with a particular version of Windows.

See also: signed driver  Windows Logo Program

**unspecified bit rate (UBR)**
An asynchronous transfer mode (ATM) service type that does not include any service guarantees.

See also: asynchronous transfer mode (ATM)

**upcasing**
The method of case-insensitive handling for network basic input/output system (NetBIOS) names. Upcasing is used by Windows to convert any lowercase letters entered in NetBIOS names to equivalent uppercase letters. For example, the name host1 is case-converted to HOST1.

See also: downcasing  network basic input/output system (NetBIOS)

**upgrade**
When referring to software, to update existing program files, folders, and registry entries to a more recent version. Upgrading, unlike performing a new installation, leaves existing settings and files in place.

See also: install  registry

**UPN**
See other term: user principal name (UPN)

**UPS**
See other term: uninterruptible power supply (UPS)

**UPS service**
A service that manages an uninterruptible power supply (UPS) connected to a computer.

See also: service  uninterruptible power supply (UPS)

**USB**
See other term: universal serial bus (USB)

**user account**
In Active Directory, an object that consists of all the information that defines a domain user, which includes user name, password, and groups in which the user account has membership. User accounts can be stored in either Active Directory or on your local computer. For computers running Windows XP Professional and member servers running Windows Server 2003, use Local Users and Groups to manage local user accounts. For domain controllers running Windows Server 2003, use Active Directory Users and Computers to manage domain user accounts.

See also: Active Directory  group  member server  object  password  user name

**user authentication module (UAM)**
A software component that prompts clients for their user names and passwords.

See also: password  user name

**user class**
An administrative feature that allows DHCP clients to be grouped logically according to a shared or common need. For example, a user class can be defined and used to allow similar DHCP leased configuration for all client computers in a specific building or site location.

See also: DHCP client  Dynamic Host Configuration Protocol (DHCP)  lease

**User Datagram Protocol (UDP)**
A transport layer protocol that offers a connectionless datagram service that guarantees neither delivery nor correct sequencing of delivered packets (much like Internet Protocol (IP)), but provides a payload checksum and upper layer protocol identification that uses source and destination ports.

See also: datagram  Internet Protocol (IP)  packet  service  Transmission Control Protocol/Internet Protocol (TCP/IP)

**user identifier (UID)**
An identifier assigned to a user within UNIX. A user can only have one UID.

**user mode**
The processing mode in which applications run.

**user name**
A unique name identifying a user account to Windows. An account's user name must be unique among the other group names and user names within its own domain or workgroup.

See also: domain  user account  workgroup

**USER object**
An object from Window Manager, which includes windows, menus, cursors, icons, hooks, accelerators, monitors, keyboard layouts, and other internal objects. In Task Manager, the number of USER objects currently being used by a process.

See also: Task Manager

**user password**
The password stored in each user's account. Each user generally has a unique user password and must type that password when logging on or accessing a server.

See also: password  user account

**user principal name (UPN)**
A user account name (sometimes referred to as the user logon name) and a domain name identifying the domain in which the user account is located. This is the standard usage for logging on to a Windows domain. The format is as follows: someone@example.com (as for an e-mail address).

See also: domain  domain name  user principal name (UPN) suffix

**user principal name (UPN) suffix**
The part of the user principal name (UPN) to the right of the @ character. The default UPN suffix for a user account is the DNS domain name of the domain that contains the user account. Alternative UPN suffixes may be added to simplify administration and user logon processes by providing a single UPN suffix for all users. The UPN suffix is only used within the Active Directory forest, and it is not required to be a valid DNS domain name.

See also: Active Directory  domain  domain name  Domain Name System (DNS)  forest  user account  user principal name (UPN)

**user profile**
A file that contains configuration information for a specific user, such as desktop settings, persistent network connections, and application settings. Each user's preferences are saved to a user profile that Windows uses to configure the desktop each time that a user logs on.

See also: desktop

**user rights**
Tasks that a user is permitted to perform on a computer system or domain. There are two types of user rights: privileges and logon rights. An example of a privilege is the right to shut down the system. An example of a logon right is the right to log on to a computer locally. Both types are assigned by administrators to individual users or groups as part of the security settings for the computer.

See also: domain  group  logon rights  privilege

**user rights policy**
Security settings that manage the assignment of rights to groups and user accounts.

See also: group  user account  user rights

**user-to-network interface (UNI)**
The interface between asynchronous transfer mode (ATM) users or end stations and an ATM switch or network. The UNI is defined in the ATM Forum UNI documents.

See also: asynchronous transfer mode (ATM)

**UTF-8**
See other term: UCS Transformation Format 8 (UTF-8)

# Windows Server 2003 Glossary - V

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - V**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**V.34**
A data transmission standard that provides for up to 33,600 bits per second (bps) communications over telephone lines. It defines a full-duplex (two-way) modulation technique and includes error-correcting and negotiation.

See also: bits per second (bps) full-duplex modulation standards V.90

**V.90**
A data transmission standard that provides for up to 56,000 bits per second (bps) communications over telephone lines. The transmission speed from the client-side modem is 33,600 bps, the same as for V.34. The transmission speed from the host-side modem, such as an Internet service provider (ISP) or corporate network, is up to 56,000 bps, with an average speed of 40,000 to 50,000 bps. When the host-side modem does not support this standard, the alternative is V.34.

See also: bits per second (bps) client host Internet service provider (ISP) modem (modulator/demodulator) modulation standards V.34

**validation**
The process of comparing files on local volumes with their associated data in secondary storage by Remote Storage. Volumes that are validated ensure that the correct data is recalled from remote storage when a user attempts to open the file from a local volume.

See also: Remote Storage volume

**variable**
In programming, a named storage location capable of containing a certain type of data that can be modified during program execution. System environment variables are defined by the operating systems in the Windows Server 2003 family, and they are the same no matter who is logged on to the computer. Administrator group members can add new variables or change the values, however. User environment variables can be different for each user of a particular computer. They include any environment variables you want to define or variables defined by your applications, such as the path where application files are located.

See also: environment variable

**variable bit rate (VBR)**
An asynchronous transfer mode (ATM) service type that guarantees service based on average and peak traffic rates. VBR is used for traffic that requires little or no cell loss. It transmits data in spurts, or bursts, rather than in a continuous stream.

See also: asynchronous transfer mode (ATM) service

**VBR**
See other term: variable bit rate (VBR)

**VCI**
See other term: virtual channel identifier (VCI)

**VDM**
See other term: virtual DOS machine (VDM)

**VDS**
See other term: Virtual Disk Service (VDS)

**vector**
For Indexing Service, an ordered series of words or numbers used in a query. For example, a single document can have (Scott Cooper; Don Hall; Amy Egert) as a vector of coauthors.

See also: Indexing Service  query

**vector font**
A font rendered from a mathematical model, in which each character is defined as a set of lines drawn between points. Vector fonts can be cleanly scaled to any size or aspect ratio.

See also: font  plotter font

**vendor class**
An administrative feature that allows DHCP clients to be identified and leased according to their vendor and hardware configuration type. For example, assigning a vendor class of HP to a printer vendor such as Hewlett-Packard would allow all Hewlett-Packard printers to be managed as a single unit so they could all obtain a similar DHCP leased configuration.

See also: DHCP client  Dynamic Host Configuration Protocol (DHCP)  lease

**video adapter**
An expansion board that plugs into a personal computer to give it display capabilities. A computer`s display capabilities depend on both the logical circuitry (provided in the video adapter) and the monitor. Each adapter offers several different video modes. The two basic categories of video modes are text and graphics. Within the text and graphics modes, some monitors also offer a choice of resolutions. At lower resolutions a monitor can display more colors. Modern adapters contain memory, so that the computer`s random access memory (RAM) is not used for storing displays. In addition, most adapters have their own graphics coprocessor for performing graphics calculations. These adapters are often called graphics accelerators.

See also: graphics mode

**virtual address**
In a virtual memory system, the address the application uses to reference memory. The kernel and the memory management unit (MMU) translate this address into a physical address before the memory is actually read or written.

See also: virtual memory

**virtual channel identifier (VCI)**
A section of the asynchronous transfer mode (ATM) cell header that contains the virtual channel address over which the cell is to be routed.

See also: asynchronous transfer mode (ATM)

**virtual cluster**
A Network Load Balancing cluster that you create by assigning specific port rules to specific virtual IP addresses. With virtual clusters, you can use different port rules for different Web sites or applications hosted on the cluster, provided each Web site or application has a different virtual IP address.

See also: Network Load Balancing  Network Load Balancing cluster  port rule  virtual IP address

**virtual container**
A container that allows any Lightweight Directory Access Protocol (LDAP)-compliant directory to be accessed through Active Directory.

See also: Active Directory  directory  Lightweight Directory Access Protocol (LDAP)

**Virtual Disk Service (VDS)**
A set of application programming interfaces (APIs) that provides a single interface for managing disks. VDS uses two sets of providers to manage storage devices. The built-in VDS software providers enable the management of disks and volumes at the operating system level. The VDS hardware providers enable the management of disk storage subsystems and are supplied by the hardware vendor.

**virtual DOS machine (VDM)**
A protected environment subsystem that emulates MS-DOS and 16-bit Windows within 32-bit versions of the Windows Server 2003 family. A VDM is created whenever a user starts an MS-DOS program in 32-bit versions of the Windows Server 2003 family. Note that 64-bit versions of the Windows Server 2003 family cannot run MS-DOS and 16-bit Windows programs.

See also: MS-DOS (Microsoft Disk Operating System)

**virtual IP address**
An IP address that is shared among the hosts of a Network Load Balancing cluster. A Network Load Balancing cluster might

also use multiple virtual IP addresses, for example, in a cluster of multihomed Web servers.

See also: host  IP address  multihomed computer  Network Load Balancing cluster

**virtual link**
A logical link between a backbone area border router and an area border router that is not connected to the backbone.

See also: area border router (ABR)  backbone router

**virtual local area network (VLAN)**
A logical grouping of hosts on one or more local area networks (LANs) that allows communication to occur between hosts as if they were on the same physical LAN.

See also: host  local area network (LAN)

**virtual memory**
Temporary storage used by a computer to run programs that need more memory than it has. For example, programs could have access to 4 gigabytes (GB) of virtual memory on a computer`s hard drive, even if the computer has only 32 megabytes (MB) of random access memory (RAM). The program data that does not currently fit in the computer`s memory is saved into paging files.

See also: paging file  random access memory (RAM)  Virtual Memory Size  virtual printer memory

**Virtual Memory Size**
In Task Manager, the amount of virtual memory, or address space, committed to a process.

See also: Task Manager  virtual memory

**virtual network**
A logical network that exists inside Novell NetWare and NetWare-compatible servers and routers but that is not associated with a physical adapter. The virtual network appears to a user as a separate network. On a computer running Windows 2000 Server or on an operating system in the Windows Server 2003 family, programs advertise their location on a virtual network, not a physical network. The internal network number identifies a virtual network inside a computer.

See also: external network number  internal network number

**virtual path identifier (VPI)**
A section of the asynchronous transfer mode (ATM) cell header that contains the virtual path address over which the cell is to be routed.

See also: asynchronous transfer mode (ATM)

**virtual printer memory**
In a PostScript printer, a part of memory that stores font information. The memory in PostScript printers is divided into two areas: banded memory and virtual memory. The banded memory contains graphics and page-layout information needed to print your documents. The virtual memory contains any font information that is sent to your printer either when you print a document or when you download fonts.

See also: PostScript printer  virtual memory

**virtual private network (VPN)**
The extension of a private network that encompasses encapsulated, encrypted, and authenticated links across shared or public networks. VPN connections typically provide remote access and router-to-router connections to private networks over the Internet.

See also: authentication  encryption  remote access  routing  tunnel

**virtual server**
In a server cluster, a collection of services that appear to clients as a physical Windows-based server but are not associated with a specific server. A virtual server is typically a resource group that contains all of the resources needed to run a particular application and that can be failed over like any other resource group. All virtual servers must include a Network Name resource and an IP Address resource.

See also: failover  IP Address resource  Network Name resource  resource  server cluster

**virus**

A malicious computer program designed specifically to replicate itself and spread from computer to computer. A virus may damage hardware, software, or data.

See also: Trojan horse

**VLAN**
See other term: virtual local area network (VLAN)

**VoIP (Voice over Internet Protocol)**
A method for sending voice over a local area network (LAN), a wide area network (WAN), or the Internet using TCP/IP packets.

See also: local area network (LAN)  Transmission Control Protocol/Internet Protocol (TCP/IP)  wide area network (WAN)

**volume**
An area of storage on a hard disk. A volume is formatted by using a file system, such as file allocation table (FAT) or NTFS, and has a drive letter assigned to it. You can view the contents of a volume by clicking its icon in Windows Explorer or in My Computer. A single hard disk can have multiple volumes, and volumes can also span multiple disks.

See also: disk  drive letter  file allocation table (FAT)  NTFS file system  simple volume  spanned volume

**volume password**
An optional, case-sensitive password you can assign to a Macintosh-accessible volume when you configure the volume. To access the volume, a user must type the volume password.

See also: password  volume

**volume set**
A volume that consists of disk space on one or more physical disks. A volume set is created by using basic disks and is supported only in Windows NT 4.0 or earlier. Volume sets were replaced by spanned volumes, which use dynamic disks.

See also: basic disk  dynamic disk  partition  spanned volume  volume

**volume shadow copy**
A volume that represents a duplicate of the original volume taken at the time the copy began.

See also: crash consistency  differential data  volume

**voluntary tunnel**
A tunnel that is initiated by the client. A voluntary tunnel tunnels Point-to-Point Protocol (PPP) over Internet Protocol (IP) from the client to the tunnel server, and then the data is forwarded to the target host by the tunnel server.

See also: client  host  Internet Protocol (IP)  Point-to-Point Protocol (PPP)  tunnel  tunnel server

**VPI**
See other term: virtual path identifier (VPI)

**VPN**
See other term: virtual private network (VPN)

**VPN server**
A computer that accepts virtual private network (VPN) connections from VPN clients. A VPN server can provide a remote access VPN connection or a router-to-router VPN connection.

See also: client  router  server  virtual private network (VPN)

# Windows Server 2003 Glossary - W

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - W**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**Waiting for Call**
A telephony signal that Network Connections has put the modem in Listen mode and is waiting for incoming calls.

See also: modem (modulator/demodulator)  Network Connections

**wake-on-LAN**
A power management feature that allows a device to bring the operating system out of standby or sleep mode if there is network activity.

See also: sleep mode

**WAN**
See other term: wide area network (WAN)

**warning level**
The point at which users are nearing their quota limit.

See also: quota limit

**weak password**
A password that does not provide an effective defense against unauthorized access to a resource. A weak password can be less than six characters long, contain all or part of a users account name, or contain less than three of the four following categories of characters: uppercase characters, lowercase characters, base 10 digits, and symbols found on the keyboard (such as !, @, #).

See also: password  password policy  strong password

**Web Distributed Authoring and Versioning (WebDAV)**
An application protocol related to Hypertext Transfer Protocol (HTTP) 1.1 that allows clients to transparently publish and manage resources on the World Wide Web.

See also: Hypertext Transfer Protocol (HTTP)  protocol

**Web server**
A computer that is maintained by a system administrator or Internet service provider (ISP) and that responds to requests from a user's browser.

See also: Internet service provider (ISP)

**Web Services (WS-*)**
The specifications for a Web Services Architecture that is based on industry standards such as Simple Object Access Protocol (SOAP); XML; Web Service Description Language (WSDL); and Universal Description, Discovery, and Integration (UDDI). WS-* provides a foundation for delivering complete, interoperable business solutions for the extended enterprise, including the ability to manage federated identity and security.

**Web Services Security (WS-Security)**
A series of specifications that describe how to attach signature and encryption headers to Simple Object Access Protocol (SOAP) messages. In addition, WS-Security describes how to attach security tokens, including binary security tokens such as X.509 certificates and Kerberos tickets, to messages. In Active Directory Federation Services (ADFS), WS-Security is used when Kerberos signs security tokens.

**WebDAV**
See other term: Web Distributed Authoring and Versioning (WebDAV)

**well-connected**

Sufficient connectivity to make your network and Active Directory useful to clients on your network. The precise meaning of well-connected is determined by your particular needs.

See also: Active Directory  client

**well-known services**
Services that are so pervasive in the computer industry that standard bindings are universally recognized. These services often maintain a machine-centric status and are frequently legacy services such as those developed for use in TCP/IP networks, for example, Telnet and File Transfer Protocol (FTP). Well-known services can published in Active Directory, but because their bindings are widely known, the information is superfluous.

See also: Active Directory  File Transfer Protocol (FTP)  machine-centric  service  Transmission Control Protocol/Internet Protocol (TCP/IP)

**wide area network (WAN)**
A network that connects geographically separated locations by using telecommunications services.

See also: device  local area network (LAN)

**wildcard character**
A keyboard character that can be used to represent one or many characters when conducting a query. The question mark (?) represents a single character, and the asterisk (*) represents one or more characters.

**wildcarding**
In DNS, the supported use of wildcard characters such as the asterisk (*) in domain names for DNS queries that resolve to multiple names. When wildcarding is used, DNS servers must support resolving name queries that use wildcard characters, and resolvers must be able to parse multiple records from any resource record sets (RRsets) issued in wildcard query responses.

See also: DNS server  domain name  Domain Name System (DNS)  resource record set (RRset)  wildcard character

**Windows Catalog**
A Web site that lists hardware and software that is designed for use with Windows XP, Windows 2000 Server products, and products in the Windows Server 2003 family. You can use this site to help you decide whether to purchase a particular device or program, to help you evaluate whether a particular computer would support an upgraded operating system, or for similar decisions about hardware and software.

See also: device  Hardware Compatibility List (HCL)

**Windows Clustering**
A feature of Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, that provides the complementary clustering technologies, Network Load Balancing clusters, and server clusters.

See also: cluster  Network Load Balancing cluster  server cluster

**Windows Hardware Quality Labs (WHQL)**
Provides compatibility testing services to test hardware and drivers for Windows operating systems.

**Windows Internet Name Service (WINS)**
A Windows name resolution service for network basic input/output system (NetBIOS) names. WINS is used by hosts running NetBIOS over TCP/IP (NetBT) to register NetBIOS names and to resolve NetBIOS names to Internet Protocol (IP) addresses.

See also: IP address  network basic input/output system (NetBIOS)  resource  service  WINS proxy  WINS resource

**Windows Logo Program**
A certification program to help customers identify systems, hardware, and software that meet a baseline definition of platform features and quality goals and ensure a good user experience of Windows.

See also: signed driver  unsigned driver

**Windows Management Instrumentation (WMI)**
A management infrastructure in Windows that supports monitoring and controlling system resources through a common set of interfaces and provides a logically organized, consistent model of Windows operation, configuration, and status.

See also: resource

**Windows Management Instrumentation filter**

A query that is based on Windows Management Instrumentation (WMI) and used to filter the effect of a Group Policy object (GPO). WMI filters are written in WMI Query Language (WQL) and are evaluated on the target computer or user. If the filter evaluates to true, then the GPO is applied. Otherwise, the GPO is not applied. WMI filters are useful for exception management.

See also: Group Policy Management console (GPMC)  Group Policy object (GPO)  Group Policy object link

**Windows Media Services**
Software that provides multimedia support, allowing you to deliver content using Advanced Streaming Format over an internet or the Internet.

See also: multimedia streaming

**Windows Open Services Architecture (WOSA)**
A common set of interfaces that connect front-end applications with back-end services.

See also: service

**Windows Support Tools**
Tools that administrators, developers, or support personnel can use to diagnose and troubleshoot operating system configuration problems. Although Windows Support Tools are included on the operating system CD, they are not guaranteed or supported by Microsoft, and they must be installed separately from the operating system.

**Windows NT token–based application**
A Windows application that relies on a Windows NT token to perform authorization of users.

**WINS**
See other term: Windows Internet Name Service (WINS)

**WINS proxy**
A computer that listens to name query broadcasts and responds for those names not on the local subnet. The proxy communicates with a WINS server to resolve names and then caches them for a specific time period.

See also: subnet  Windows Internet Name Service (WINS)

**WINS resource**
A resource type that provides Windows Internet Name Service (WINS) from a cluster.

See also: cluster  resource types  Windows Internet Name Service (WINS)

**Winsock**
Windows Sockets. An application programming interface (API) standard for software that provides a TCP/IP interface under Windows.

See also: application programming interface (API)  Transmission Control Protocol/Internet Protocol (TCP/IP)

**wireless communication**
Communication between a computer and another computer or device without wires. The form of wireless communication provided as part of the Windows operating system uses infrared (IR) light to transmit files. Radio frequencies, as used by cellular and cordless telephones, are another form of wireless communication.

See also: infrared (IR)  infrared device  infrared port

**WMI**
See other term: Windows Management Instrumentation (WMI)

**WMI filter**
See other term: Windows Management Instrumentation filter

**workgroup**
A simple grouping of computers, intended only to help users find such things as printers and shared folders within that group. Workgroups in Windows do not offer the centralized user accounts and authentication offered by domains.

See also: authentication  domain  user account

**working set**
For a process, the amount of physical memory assigned to the process by the operating system.

See also: process

**World Wide Web (WWW)**
A system for exploring the Internet by using hyperlinks. When you use a Web browser, the Web appears as a collection of text, pictures, sounds, and digital movies.

See also: Internet

**WOSA**
See other term: Windows Open Services Architecture (WOSA)

**WS-Federation**
A specification that defines a model and set of messages for brokering trust and the federation of identity and authentication information across different trust realms. The WS-Federation specification identifies two sources of identity and authentication requests across trust realms: active requestors, such as Simple Object Access Protocol (SOAP)-enabled applications, and passive requestors, which are defined as HTTP browsers capable of supporting broadly supported HTTP, for example, HTTP 1.1.

**WS-Federation Passive Requestor Profile (WS-F PRP)**
An implementation of the WS-Federation specification that proposes a standard protocol for how passive clients (such as Web browsers) apply the federation framework. Within this protocol, Web service requestors are expected to understand the new security mechanisms and be capable of interacting with Web service providers.

# Windows Server 2003 Glossary - X

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - X**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**X.500**
A set of standards defining a distributed directory service, developed by the International Standards Organization (ISO).

See also: directory service

**X.509 v3 certificate**
Version 3 of the International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T) recommendation X.509 for certificate syntax and format. This is the standard certificate format used by Windows certificate-based processes. An X.509 certificate includes the public key and information about the person or entity to whom the certificate is issued, information about the certificate, plus optional information about the certification authority (CA) issuing the certificate.

See also: certification authority (CA)  International Telecommunication Union - Telecommunication [Standardization Sector] (ITU-T)  public key

**x86**
Refers to microprocessors that have or emulate the 32-bit Intel processor architecture.

**XML**
See other term: Extensible Markup Language (XML)

**XOR (exclusive OR)**
A Boolean operation that yields true if and only if one of its operands is true and the other is false. For some server operating systems, the exclusive OR is a Boolean operation in which the RAID-5 volume form of fault tolerance maintains an XOR of the total data to provide data redundancy. This is used to reconstruct missing data on a failed disk or sector from the remaining disks in the RAID-5 volume.

See also: fault tolerance  RAID-5 volume

# Windows Server 2003 Glossary - Y

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - Y**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**ypcat**
A Windows command-line utility that prints the values of all keys from the Network Information Service (NIS) database specified by mapname, which can be a map name or a map nickname.

**ypclear**
A Windows command-line utility that tells the specified Network Information Service (NIS) server in a specified domain that a map has changed and that the server should clear its cache.

**ypmatch**
A Windows command-line utility that displays the values of one or more keys from the Network Information Service (NIS) map specified by mapname, which may be a map name or a map nickname.

**yppush**
A Windows command-line utility that sends a request to a UNIX-based subordinate (also known as slave) Network Information Service (NIS) server to transfer a new version of the NIS map specified by mapname from the master NIS server.

# Windows Server 2003 Glossary - Z

Updated: March 7, 2008

Applies To: Windows Server 2000, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

For more Windows Server terms, see the Windows Server 2008 Glossary.

**Glossary - Z**

# A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**zone**

In the Macintosh environment, a logical grouping that simplifies browsing the network for resources, such as servers and printers. In a DNS database, a manageable unit of the DNS database that is administered by a DNS server. A zone stores the domain names and data of the domain with a corresponding name, except for domain names stored in delegated subdomains.

See also: DNS server  domain  Domain Name System (DNS)  partition  resource

**zone list**

In the Macintosh environment, a list that includes all of the zones associated with a particular network. A zone list is not to be confused with a DNS zone.

See also: DNS zone  Domain Name System (DNS)  zone

**zone transfer**

The synchronization of authoritative DNS data between DNS servers. A DNS server configured with a secondary zone periodically queries the master DNS servers to synchronize its zone data.

See also: authoritative  DNS server  primary zone  secondary zone  zone

# IIS 6.0 Glossary

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

**Symbols**

**3DES** See definition for Triple DES (3DES).

**A**

**abstract classes** Templates used only to derive new Structural classes. Abstract classes cannot be instantiated in the directory.

**access control** A security mechanism that determines which operations a user, group, service, or computer is authorized to perform on a computer or on a particular object, such as a file, printer, registry subkey, or directory service object. See also group; object; permission; registry.

**access control list (ACL)** A list of security protections that apply to an entire object, a set of the objects properties, or an individual property of an object. There are two types of access control lists: discretionary and system. See also object.

**Active Directory** The Windows-based directory service. Active Directory stores information about objects on a network and makes this information available to users and network administrators. Active Directory gives network users access to permitted resources anywhere on the network using a single logon process. It provides network administrators with an intuitive, hierarchical view of the network and a single point of administration for all network objects. See also directory partition; directory service; domain; forest; object.

**Active Directory Service Interfaces (ADSI)** A directory service model and a set of Component Object Model (COM) interfaces. ADSI enables Windows applications and Active Directory clients to access several network directory services, including Active Directory. ADSI is supplied as a software development kit (SDK). See also Active Directory; Component Object Model (COM); directory service.

**ActiveX** A set of technologies that allows software components to interact with one another in a networked environment, regardless of the language in which the components were created.

**Address Resolution Protocol (ARP)** In TCP/IP, a protocol that uses broadcast traffic on the local network to resolve a logically assigned Internet Protocol version 4 (IPv4) address to its physical hardware or media access control (MAC) layer address. In asynchronous transfer mode (ATM), ARP is used two different ways. For classical IPv4 over ATM (CLIP), ARP is used to resolve addresses to ATM hardware addresses. For ATM LAN emulation (LANE), ARP is used to resolve Ethernet/802.3 or Token Ring addresses to ATM hardware addresses. See also asynchronous transfer mode (ATM); Internet Protocol (IP); IP address; packet; Transmission Control Protocol/Internet Protocol (TCP/IP).

**administrative credentials** Logon information that is used to identify a member of an administrative group. Groups that use administrative credentials include Administrators, Domain Admins, and DNS Admins. Most system-wide or domain-wide tasks require administrative credentials. See also Administrators group; group.

**Administrators group** On a local computer, a group whose members have the highest level of administrative access to the local computer. Examples of administrative tasks that can be performed by members of this group include installing programs; accessing all files on the computer; auditing access control; and creating, modifying, and deleting local user accounts. In an Active Directory domain, a group whose members have the highest level of administrative access in the domain. Examples of administrative tasks that can be performed by members of this group include setting domain policy; assigning and resetting domain user account passwords; setting up and managing domain controllers; and creating, modifying, and deleting domain user accounts. See also access control; Active Directory; auditing; domain; domain controller; group; object.

**ADSI** See definition for Active Directory Service Interfaces (ADSI).

**ADSI provider** COM objects that implement ADSI for a particular namespace (for example, an LDAP namespace such as Active Directory).

**agent** An application that runs on a Simple Network Management Protocol (SNMP) managed device. The agent application is the object of management activities. A computer running SNMP agent software is also sometimes referred to as an agent.

**aggregation** A composition technique for implementing component objects in which a new object can be built by using one or more existing objects that support some or all of the new objects required interfaces.

**American Standard Code for Information Interchange (ASCII)** A standard single-byte character encoding scheme used

for text-based data. ASCII uses designated 7-bit or 8-bit number combinations to represent either 128 or 256 possible characters. Standard ASCII uses 7 bits to represent all uppercase and lowercase letters, the numbers 0 through 9, punctuation marks, and special control characters used in U.S. English. Most current x86-based systems support the use of extended (or high) ASCII. Extended ASCII allows the eighth bit of each character to identify an additional 128 special symbol characters, foreign-language letters, and graphic symbols.

**Anonymous access** An authentication mechanism by which users who are able to connect to an Internet site without credentials are assigned to the IUSR_*ComputerName* account and granted the access rights that are assigned to that account. See also access control; Anonymous authentication; authentication.

**Anonymous authentication** An authentication mechanism that does not require user accounts and passwords. Anonymous authentication grants remote users the identity IUSR_*ComputerName*. Anonymous authentication is used on the Internet to grant visitors restricted access to predefined public resources. See also Anonymous access; authentication.

**Anonymous FTP authentication** A protocol that makes it possible for a user to retrieve documents, files, programs, and other archived data from anywhere on the Internet without having to establish a logon name and password.

**apartment-threaded** A threading model in which each method of a component will execute on a thread that is associated with that component. See also multithreaded apartment (MTA); single-threaded apartment (STA).

**API** See definition for application programming interface (API).

**application** A computer program, such as a word processor or electronic spreadsheet, or a group of Active Server Pages (ASP) scripts and components that perform such tasks.

**application isolation** The separation of applications by process boundaries that prevent the applications from affecting one another. Application isolation is configured differently for each of the two Internet Information Services (IIS) isolation modes. See also IIS 5.0 isolation mode; worker process isolation mode.

**application pool** A grouping of one or more URLs served by a worker process.

**application programming interface (API)** A set of routines that an application uses to request and carry out lower-level services performed by a computers operating system. These routines usually carry out maintenance tasks such as managing files and displaying information.

**application root** The root directory for an application. All directories and files contained within the application root are considered part of the application. Also called an application starting-point directory.

**application scope** A way of making data available to all users of an application from all pages of a Web application. A variable or an object instance is given application scope by being stored in the Active Server Pages (ASP) application object. Application scope is useful for global data, such as a global counter.

**argument** A constant, variable, or expression that is passed to a procedure.

**array** A list of data values—all of the same type—any element of which can be referenced by an expression that consists of the array name followed by an indexing expression. Arrays are part of the fundamentals of data structures, which, in turn, are a major fundamental of computer programming.

**ASCII (American Standard Code for Information Interchange)** See definition for American Standard Code for Information Interchange (ASCII).

**ASP buffering** Functionality of Active Server Pages (ASP) that temporarily stores all output that is generated by a script until script execution is complete and then sends the output to a client.

**association** In file name extension mapping, the linking of a file extension, such as .asp, to an application, such as asp.dll. In Windows Management Instrumentation (WMI), an association class represents a relationship between two specific WMI classes. The properties of an association class include pointers, or references, to the two classes or instances.

**asymmetric key algorithm** See definition for public-key algorithm.

**asynchronous transfer mode (ATM**) A high-speed, connection-oriented, virtual circuit-based packet switching protocol used to transport many different types of network traffic. ATM packages data in 53-byte, fixed-length cells that can be switched quickly between logical connections on a network. See also protocol.

**ATM** See definition for asynchronous transfer mode (ATM).

**attribute** For files, information that indicates whether a file is read-only, hidden, ready for archiving (backing up), compressed, or encrypted, and whether the file contents should be indexed for fast file searching. In Active Directory, a

property of an object. For each object class, the schema defines which attributes an instance of the class must have and which additional attributes it might have. See also Active Directory; class; object.

**auditing** The process that tracks the activities of users by recording selected types of events in the security log of a server or a workstation.

**authentication** The process for verifying that an entity or object is who or what it claims to be. Examples include confirming the source and integrity of information, such as verifying a digital signature or verifying the identity of a user or computer. See also cryptography; Kerberos V5 authentication protocol.

**authorization** The process that determines what a user is permitted to do on a computer system or network. See also authentication.

**Automation** A Component Object Model (COM) based technology that allows for interoperability among ActiveX components, including OLE components. Formerly referred to as *OLE Automation*.

**availability** A level of service provided by applications, services, or systems. Highly available systems have minimal downtime, whether planned or unplanned. Availability is often expressed as the percentage of time that a service or system is available, for example, 99.9 percent for a service that is down for 8.75 hours a year.

**B**

**bandwidth** The data transfer capacity of a transmission medium. In digital communications, the transfer capacity expressed in bits per second (bps) or megabits per second (Mbps). For example, Ethernet accommodates a bandwidth of 10,000,000 bps or 10 Mbps. In analog communications, the difference between the highest and lowest frequencies in a specific range. For example, an analog telephone line accommodates a bandwidth of 3,000 hertz (Hz), the difference between the lowest (300 Hz) and highest (3,300 Hz) frequencies that it can carry. See also bits per second (bps).

**bandwidth throttling** Setting the maximum portion of total network capacity that a service is allowed to use. An administrator can deliberately limit a servers Internet workload by not allowing it to receive requests at full capacity, thus saving resources for other programs, such as e-mail.

**baseline** A range of measurements derived from performance monitoring that represents acceptable performance under typical operating conditions.

**Basic authentication** An authentication mechanism that is supported by most browsers, including Internet Explorer. Basic authentication encodes user name and password data before transmitting it over the network. Note that *encoding* is not the same as *encryption*. Also known as *plaintext authentication*. See also Anonymous authentication; authentication; Digest authentication; encryption.

**baud rate** The speed at which a modem communicates. Baud rate refers to the number of times the condition of the line changes. This is equal to bits per second only if each signal corresponds to one bit of transmitted data. Modems must operate at the same baud rate in order to communicate with each other. If the baud rate of one modem is set higher than that of the other, the faster modem usually alters its baud rate to match that of the slower modem. See also bits per second (bps); modem (modulator/demodulator).

**Berkeley Internet Name Domain (BIND)** An implementation of Domain Name System (DNS) written and ported to most available versions of the UNIX operating system. The Internet Software Consortium maintains the BIND software. See also Domain Name System (DNS).

**binary** A base-2 number system in which values are expressed as combinations of two digits, 0 and 1.

**BIND** See definition for Berkeley Internet Name Domain (BIND).

**binding** A process by which software components and layers are linked together. When a network component is installed, the binding relationships and dependencies for the components are established. Binding allows components to communicate with each other.

**bitmask** A value that is used with bit-wise operators (And, Eqv, Imp, Not, Or, Xor) to test the state of individual bits in a particular bit-field value. See also bitmask identifier.

**bitmask identifier** For the metabase, a name assigned to a bitmask to help identify its purpose. For example, In IIS 6.0, bitmask 512 is assigned the identifier MD_ACCESS_SCRIPT. See also bitmask.

**bits per second (bps)** The number of bits transmitted every second, used as a measure of the speed at which a device, such as a modem, can transfer data. See also modem (modulator/demodulator).

**Boolean data type** A data type with only two passable values, True (-1) or False (0). Boolean variables are stored as 16-bit

(2-byte) numbers.

**both-threaded** A threading model in which the object has the characteristics of an apartment-threaded object as well as a free-threaded object. See also apartment-threaded.

**bps** See definition for bits per second (bps).

**browser** Software that interprets the markup of files in HTML, formats them into Web pages, and displays them to the end user. Some browsers also permit end users to send and receive e-mail, read newsgroups, and play sound or video files embedded in Web documents.

**built-in groups** The default security groups installed with the operating system. Built-in groups have been granted useful collections of rights and built-in abilities. In most cases, built-in groups provide all the capabilities needed by a particular user. For example, members of the built-in Backup Operators group can back up and restore files and folders. To provide a needed set of capabilities to a user account, assign it to the appropriate built-in group. See also group.

**bulk encryption** A process in which large amounts of data, such as files, e-mail messages, or online communications sessions, are encrypted for confidentiality. It is usually done with a symmetric key algorithm. See also encryption.

**C**

**CA** See definition for certification authority (CA).

**cache** A special memory subsystem in which frequently used data values are duplicated for quick access.

**call** To transfer program execution to some section of code (usually a subroutine) while saving the necessary information to allow execution to resume at the calling point when the called section has completed execution. When a subroutine call occurs, one or more values (known as arguments or parameters) are often passed to the subroutine, which can then use and sometimes modify these values.

**callback function** A function provided by Internet Information Services (IIS) that allows an Internet Server API (ISAPI) extension or filter to access IIS services.

**certificate** A digital document that is commonly used for authentication and to secure information on open networks. A certificate securely binds a public key to the entity that holds the corresponding private key. Certificates are digitally signed by the issuing certification authority (CA), and they can be issued for a user, a computer, or a service. See also certification authority (CA); private key; public key.

**certificate revocation list (CRL)** A document maintained and published by a certification authority that lists certificates that have been revoked. See also certificate; certification authority (CA).

**certificate trust list (CTL)** A signed list of root certification authority certificates that an administrator considers reputable for designated purposes, such as client authentication or secure e-mail. See also certificate; certification authority (CA).

**certificate, client** See definition for client certificate.

**certification authority (CA)** An entity responsible for establishing and vouching for the authenticity of public keys belonging to subjects (usually users or computers) or other certification authorities. Activities of a certification authority can include binding public keys to distinguished names through signed certificates, managing certificate serial numbers, and certificate revocation. See also certificate; public key.

**CGI** See definition for common gateway interface (CGI).

**class** A category of objects that share a common set of characteristics. Each object in the directory is an instance of one or more classes in the schema. See also object.

**client** Any computer or program connecting to, or requesting the services of, another computer or program. Client can also refer to the software that enables the computer or program to establish the connection. For a local area network (LAN) or the Internet, a computer that uses shared network resources provided by another computer (called a *server*).See also server.

**client certificate** A digital certificate that functions in a manner that is similar to a driver's license or passport. Client certificates can contain detailed identification information about the user and organization that issued the certificate.

**client tier** In the three-tier Web application model, the application or process that requests services from the middle tier, which typically includes a Web server and business processes. See also data source tier; middle tier.

**client/server architecture** A model of computing in which client applications running on a desktop or personal computer access information on remote servers or host computers. The client portion of the application is typically optimized for user

interaction, whereas the server portion provides centralized, multi-user functionality.

**cluster** In data storage, the smallest amount of disk space that can be allocated to hold a file. All file systems used by Windows organize hard disks based on clusters, which consist of one or more contiguous sectors. The smaller the cluster size, the more efficiently a disk stores information. If no cluster size is specified during formatting, Windows picks defaults based on the size of the volume. These defaults are selected to reduce the amount of space that is lost and the amount of fragmentation on the volume. Also called an *allocation unit*. In computer networking, a group of independent computers that work together to provide a common set of services and present a single-system image to clients. The use of a cluster enhances the availability of the services and the scalability and manageability of the operating system that provides the services. See also availability; client; scalability.

**code page** A means of providing support for character sets and keyboard layouts for different countries or regions. A code page is a table that relates the binary character codes used by a program to keys on the keyboard or to characters on the display.

**COM** See definition for Component Object Model (COM).

**COM+** An extension of the COM (Component Object Model) programming architecture that includes a runtime or execution environment and extensible services, including transaction services, security, load balancing, and automatic memory management. See also Component Object Model (COM).

**common gateway interface (CGI)** A server-side interface for initiating software services. For example a set of interfaces that describe how a Web server communicates with software on the same computer. Any software can be a CGI program if it handles input and output according to the CGI standard.

**Component Object Model (COM)** An object-based programming model designed to promote software interoperability; it allows two or more applications or components to easily cooperate with one another, even if they were written by different vendors, at different times, in different programming languages, or if they are running on different computers running different operating systems. OLE technology and ActiveX are both built on top of COM. See also ActiveX.

**concrete class** In Windows Management Instrumentation (WMI), a class from which you can create an instance because it has a full implementation.

**concurrency** The appearance of simultaneous execution of processes or transactions by interleaving the execution of multiple pieces of work.

**connected user** A user who has access to a computer or a resource across the network.

**console tree** The left pane in Microsoft Management Console (MMC) that displays the items contained in the console. The items in the console tree and their hierarchical organization determine the capabilities of a console. See also details pane.

**cookie** A block of data that a Web server stores on a client system. When a user returns to the same Web site, the browser sends a copy of the cookie back to the server. Cookies identify users, instruct the server to send a customized version of the requested Web page, and submit account information for the user.

**credentials** A set of information that includes identification and proof of identification that is used to gain access to local and network resources. Examples of credentials are user names and passwords, smart cards, and certificates. See also certificate.

**CRL** See definition for certificate revocation list (CRL).

**CryptoAPI** An application programming interface (API) that is provided as part of Microsoft Windows. CryptoAPI provides a set of functions that allows applications to encrypt or digitally sign data in a flexible manner while providing protection for the user's sensitive private key data. Actual cryptographic operations are performed by independent modules known as *cryptographic service providers (CSPs)*. See also application programming interface (API); cryptographic service provider (CSP); private key.

**cryptographic service provider (CSP)** The code that performs authentication, encoding, and encryption services that Windows-based applications access through CryptoAPI. A CSP is responsible for creating keys, destroying them, and using them to perform a variety of cryptographic operations. Each CSP provides a different implementation of the CryptoAPI. Some provide stronger cryptographic algorithms, while others use hardware components, such as smart cards. See also authentication; CryptoAPI; encryption.

**cryptography** The processes, art, and science of keeping messages and data secure. Cryptography is used to enable and ensure confidentiality, data integrity, authentication (entity and data origin), and nonrepudiation. See also authentication.

**CSP** See definition for cryptographic service provider (CSP).

**CTL** See definition for certificate trust list (CTL).

**custom property** A metabase property that is not predefined in the metabase schema (MBSchema.xml) and is created programmatically to meet a specific need. See also schema.

**cycle** In logging, to close an existing log file and start a new one.

**D**

**daemon** A networking program, usually associated with UNIX systems, that runs in the background performing tool functions such as housekeeping or maintenance without user intervention or awareness. Pronounced demon.

**Data Encryption Standard (DES)** An encryption algorithm that uses a 56-bit key and maps a 64-bit input block to a 64-bit output block. The key appears to be a 64-bit key, but one bit in each of the eight bytes is used for odd parity, resulting in 56 bits of usable key. See also key.

**data source tier** A logical layer that represents a computer running a Database Management System (DBMS), such as a SQL Server database. See also client tier; middle tier.

**datagram** One packet, or unit, of information that includes relevant delivery information, such as the destination address, that is sent through a packet-switching network. See also packet.

**DCOM** See definition for Distributed Component Object Model (DCOM).

**deadlock** A situation in which a thread will not relinquish its exclusive access to a critical section.

**debugger** A program designed to aid in detecting, locating, and correcting errors in another program by allowing the programmer to step through the program, examine the data, and monitor conditions such as the values of variables.

**decryption** The process of making encrypted data readable again by converting ciphertext to plaintext. See also encryption; plaintext.

**default document** The file that is sent by a Web server when it receives a request for a Uniform Resource Locator (URL) that does not specify a file name. This document can be generated automatically by the server, or it can be a custom file that is placed in that directory by the administrator. Sometimes called a default home page.

**default gateway** A configuration item for the TCP/IP protocol that is the IP address of a directly reachable IP router. Configuring a default gateway creates a default route in the IP routing table. See also Internet Protocol (IP); IP address; Transmission Control Protocol/Internet Protocol (TCP/IP).

**DES** See definition for Data Encryption Standard (DES).

**details pane** The right pane in Microsoft Management Console (MMC) that displays details for the selected item in the console tree. The details can be a list of items or they can be administrative properties, services, and events that are acted on by a snap-in. See also console tree; snap-in.

**DHCP** See definition for Dynamic Host Configuration Protocol (DHCP).

**dial-up connection** The connection to your network if you use a device that uses the telephone network. This includes modems with a standard telephone line, ISDN cards with high-speed ISDN lines, or X.25 networks. If you are a typical user, you might have one or two dial-up connections, for example, to the Internet and to your corporate network. In a more complex server situation, multiple network modem connections might be used to implement advanced routing. See also Integrated Services Digital Network (ISDN); modem (modulator/demodulator).

**dial-up networking (DUN)** Connecting to a remote network or the Internet through a dial-up connection, such as a modem.

**Digest authentication** An authentication mechanism that hashes user name, password, and other data before transmitting it over the network. See also authentication; Basic authentication; encryption; hash.

**digital certificate** An electronic certification issued by certification authorities that shows where a program comes from and proves that the installation package has not been altered. Administrators should sign their code with a digital certificate if planning to distribute an Internet Explorer package over the Internet. See also certification authority (CA).

**digital signature** The part of a digital certificate that contains an encryption key that uniquely identifies the holder of the certificate. See also certificate; client; key pair.

**directory** An information source that contains information about users, computer files, or other objects. In a file system, a directory stores information about files. In a distributed computing environment (such as a Windows domain), the directory stores information about objects such as printers, fax servers, applications, databases, and other users. See also domain; object.

**directory browsing** A feature that automatically provides a default Web page of available directories and files to browsers that submit a Uniform Resource Locator (URL) that does not specify a particular file.

**directory partition** A contiguous subtree of Active Directory that is replicated as a unit to other domain controllers in the forest that contain a replica of the same subtree. In Active Directory, a single domain controller always holds at least three directory partitions: schema (class and attribute definitions for the directory), configuration (replication topology and related metadata), and domain (subtree that contains the per-domain objects for one domain). Domain controllers running Windows Server 2003 can also store one or more application directory partitions. See also Active Directory; attribute; domain.

**directory replication** The copying of a master set of directories from a server (called an export server) to specified servers or workstations (called import computers) in the same or other domains. Replication simplifies the task of maintaining identical sets of directories and files on multiple computers because only a single master copy of the data must be maintained. Files are replicated when they are added to an exported directory and every time a change is saved to the file.

**directory service** Both the directory information source and the service that makes the information available and usable. A directory service enables the user to find an object when given any one of its attributes. See also Active Directory; attribute; directory; object.

**Distributed Component Object Model (DCOM)** The Microsoft Component Object Model (COM) specification that defines how components communicate over Windows-based networks. Use the DCOM Configuration tool to integrate client/server applications across multiple computers. DCOM can also be used to integrate robust Web browser applications. See also Component Object Model (COM).

**distributed processing** A computing environment that contains a client and a server. This structure allows the workload to be divided into parts yet appear as a single process.

**DLL** See definition for dynamic-link library (DLL).

**DNS** See definition for Domain Name System (DNS).

**domain** In Active Directory, a collection of computer, user, and group objects defined by the administrator. These objects share a common directory database, security policies, and security relationships with other domains. In DNS, any tree or subtree within the DNS namespace. Although the names for DNS domains often correspond to Active Directory domains, DNS domains should not be confused with Active Directory domains. See also Active Directory; Domain Name System (DNS); object.

**domain controller** In an Active Directory forest, a server that contains a writable copy of the Active Directory database, participates in Active Directory replication, and controls access to network resources. Administrators can manage user accounts, network access, shared resources, site topology, and other directory objects from any domain controller in the forest. See also Active Directory; authentication; directory; forest; shared resource.

**domain name** The name given by an administrator to a collection of networked computers that share a common directory. Part of the DNS naming structure, domain names consist of a sequence of name labels separated by periods. See also domain; Domain Name System (DNS).

**Domain Name System (DNS)** A hierarchical, distributed database that contains mappings of DNS domain names to various types of data, such as IP addresses. DNS enables the location of computers and services by user-friendly names, and it also enables the discovery of other information stored in the database. See also domain name; IP address; ping; Transmission Control Protocol/Internet Protocol (TCP/IP).

**DWORD** A data type that is composed of hexadecimal data with a maximum allotted space of 4 bytes.

**dynamic binding** Binding (converting symbolic addresses in the program to storage-related addresses) that occurs during program execution. The term often refers to object-oriented applications that determine, during run time, which software routines to call for particular data objects. Also called late binding.

**Dynamic Host Configuration Protocol (DHCP)** A TCP/IP service protocol that offers dynamic leased configuration of host IP addresses and distributes other configuration parameters to eligible network clients. DHCP provides safe, reliable, and simple TCP/IP network configuration, prevents address conflicts, and helps conserve the use of client IP addresses on the network. DHCP uses a client/server model where the DHCP server maintains centralized management of IP addresses that are used on the network. DHCP-supporting clients can then request and obtain lease of an IP address from a DHCP server as part of their network boot process. See also IP address; Transmission Control Protocol/Internet Protocol (TCP/IP).

**Dynamic HTML** A collection of features that extends the capabilities of traditional HTML, giving Web authors more flexibility, design options, and creative control over the appearance and behavior of Web pages.

**dynamic page** A Hypertext Markup Language (HTML) document that contains animated GIFs, Java applets, ActiveX Controls, or dynamic HTML (DHTML). Also, a Web page that is created automatically, based on information that is provided by the user, or that is generated on the fly with Active Server Pages (ASP).

**dynamic-link library (DLL)** An operating system feature that allows executable routines (generally serving a specific function or set of functions) to be stored separately as files with .dll extensions. These routines are loaded only when needed by the program that calls them.

**E**

**early binding** Occurs when an object is assigned to a variable that is declared to be of a specific object type. Early bound objects allow the compiler to allocate memory and perform other optimizations before an application executes.

**encapsulation** The method used to pass data from one protocol over a network within a different protocol. Data from one protocol is wrapped with the header of a different protocol. Encapsulation is described in RFC 1483. See also protocol.

**encryption** The process of disguising a message or data in such a way as to hide its substance. See also public key encryption; symmetric encryption.

**encryption key** A bit string that is used in conjunction with an encryption algorithm to encrypt and decrypt data. See also private key; public key.

**Ethernet** The IEEE 802.3 standard that uses Carrier Sense Multiple Access with Collision Detection (CSMA/CD) as the medium access control. Ethernet supports different mediums, such as coaxial cable, fiber-optic cable, and twisted-pair wiring, and different data rates, such as 10 megabits per second (Mbps).

**event** Any significant occurrence in the system or an application that requires users to be notified or an entry to be added to a log.

**Event Log service** A service that records events in the system, security, and application logs. The Event Log service is located in Event Viewer. See also event; event logging; Event Viewer.

**event logging** The process of recording an audit entry in the audit trail whenever certain events occur, such as services starting and stopping or users logging on and off and accessing resources. See also auditing; event; Event Viewer.

**Event Viewer** A component you can use to view and manage event logs, gather information about hardware and software problems, and monitor security events. Event Viewer maintains logs about program, security, and system events. See also event; event logging.

**extended partition** A type of partition that you can create only on basic master boot record (MBR) disks. Extended partitions are useful if you want to create more than four volumes on a basic MBR disk. Unlike primary partitions, you do not format an extended partition with a file system and then assign a drive letter to it. Instead, you create one or more logical drives within the extended partition. After you create a logical drive, you format it and assign it a drive letter. An MBR disk can have up to four primary partitions or three primary partitions, one extended partition, and multiple logical drives. See also logical drive; partition.

**Extensible Markup Language (XML)** A meta-markup language that provides a format for describing structured data. This facilitates more precise declarations of content and more meaningful search results across multiple platforms. In addition, XML enables a new generation of Web-based data viewing and manipulation applications. See also Hypertext Markup Language (HTML).

**extranet** A limited subset of computers or users on a public network, typically the Internet, that can access an organizations internal network. For example, the computers or users might belong to a partner organization.

**F**

**failback** The process of moving resources, either individually or in a group, back to their preferred node after the node has failed and come back online. See also node.

**failover** In server clusters, the process of taking resource groups offline on one node and bringing them online on another node. When failover occurs, all resources within a resource group fail over in a predefined order; resources that depend on other resources are taken offline before, and are brought back online after, the resources on which they depend. See also node; server cluster.

**FAT** See definition for file allocation table (FAT).

**fat server** In a client/server architecture, a server that performs most of the processing, with the client performing little or no processing.

**FAT32** A derivative of the file allocation table (FAT) file system. FAT32 supports smaller cluster sizes and larger volumes than FAT, which results in more efficient space allocation on FAT32 volumes. See also file allocation table (FAT).

**fault tolerance** The ability of computer hardware or software to ensure data integrity when hardware failures occur. Fault-tolerant features appear in many server operating systems and include mirrored volumes, RAID-5 volumes, and server clusters. See also cluster.

**file allocation table (FAT)** A file system used by MS-DOS and other Windows operating systems to organize and manage files. The file allocation table is a data structure that Windows creates when you format a volume by using FAT or FAT32 file systems. Windows stores information about each file in the file allocation table so that it can retrieve the file later. See also FAT32; NTFS file system.

**File Transfer Protocol (FTP)** A member of the TCP/IP suite of protocols, used to copy files between two computers on the Internet. Both computers must support their respective FTP roles: one must be an FTP client and the other an FTP server. See also Transmission Control Protocol/Internet Protocol (TCP/IP).

**filter** For Indexing Service, software that extracts content and property values from a document to index them. For Internet Protocol security (IPsec), a specification of Internet Protocol (IP) traffic that provides the ability to trigger security negotiations for a communication based on the source, destination, and type of IP traffic. For Internet Information Services (IIS), a feature of Internet Server Application Programming Interface (ISAPI) that allows preprocessing of requests and postprocessing of responses, permitting site-specific handling of Hypertext Transfer Protocol (HTTP) requests and responses. In IP and Internetwork Packet Exchange (IPX) packet filtering, a definition in a series of definitions that indicates to the router the type of traffic allowed or disallowed on each interface. See also Internet Information Services (IIS); Internet Protocol (IP); Internet Protocol security (IPsec); Internet Server Application Programming Interface (ISAPI).

**firewall** A combination of hardware and software that provides a security system for the flow of network traffic, usually to prevent unauthorized access from outside to an internal network or intranet. Also called a *security-edge gateway*. See also proxy server.

**forest** One or more Active Directory domains that share the same class and attribute definitions (schema), site and replication information (configuration), and forest-wide search capabilities (global catalog). Domains in the same forest are linked with two-way, transitive trust relationships. See also Active Directory; domain.

**FORTEZZA** A family of security products including PCMCIA-based cards, compatible serial port devices, combination cards (such as FORTEZZA/Modem and FORTEZZA/Ethernet), server boards, and others. FORTEZZA is a registered trademark held by the U.S. National Security Agency.

**frame** In synchronous communication, a package of information transmitted as a single unit from one device to another.

**FTP** See definition for File Transfer Protocol (FTP).

**G**

**gateway** A dedicated device (or a set of services running on a dedicated computer) that routes network traffic and enables communication between different networking protocols. A gateway is a multiprotocol Internet Protocol (IP) router that translates between different transport protocols or data formats. See also Internet Protocol (IP).

**globally unique identifier (GUID)** A 16-byte value generated from the unique identifier on a device, the current date and time, and a sequence number. A GUID is used to identify a particular device or component.

**graphical user interface (GUI)** A display format, like that of Windows, that represents a program's functions with graphic images such as buttons and icons. GUIs enable a user to perform operations and make choices by pointing and clicking with a mouse.

**group** A collection of users, computers, contacts, and other groups. Groups can be used as security or as e-mail distribution collections. Distribution groups are used only for e-mail. Security groups are used both to grant access to resources and as e-mail distribution lists. See also domain; local group.

**group account** A collection of user accounts. By making a user account a member of a group, you give the related user all the rights and permissions granted to the group. See also group.

**Group Policy** The infrastructure within Active Directory directory service that enables directory-based change and configuration management of user and computer settings, including security and user data. You use Group Policy to define configurations for groups of users and computers. With Group Policy, you can specify policy settings for registry-based policies, security, software installation, scripts, folder redirection, remote installation services, and Internet Explorer maintenance. The Group Policy settings that you create are contained in a Group Policy object (GPO). By associating a GPO with selected Active Directory system containers—sites, domains, and organizational units—you can apply the GPOs policy

settings to the users and computers in those Active Directory containers. To create an individual GPO, use the Group Policy Object Editor. To manage Group Policy objects across an enterprise, you can use the Group Policy Management console. See also Active Directory.

**Guest account** A built-in account used to log on to a computer running Windows when a user does not have an account on the computer or domain or in any of the domains trusted by the computers domain. See also domain.

**GUID** See definition for globally unique identifier (GUID).

**H**

**hash** A fixed-size result that is obtained by applying a one-way mathematical function (sometimes called a *hash algorithm*) to an arbitrary amount of data. If there is a change in the input data, the hash changes. The hash can be used in many operations, including authentication and digital signing. Also called a *message digest*. See also authentication; hash algorithm.

**hash algorithm** An algorithm that produces a hash value of some piece of data, such as a message or session key. With a good hash algorithm, changes in the input data can change every bit in the resulting hash value; for this reason, hashes are useful in detecting any modification in a data object, such as a message. Furthermore, a good hash algorithm makes it computationally infeasible to construct two independent inputs that have the same hash. Typical hash algorithms include MD2, MD4, MD5, and SHA-1. Also called a *hash function*. See also MD5; Secure Hash Algorithm (SHA-1).

**headless server** See definition for remotely administered server.

**heaps** A portion of memory reserved for a program to use for the temporary storage of data structures whose existence or size cannot be determined until the program is running.

**Help and Support Center** A unified place where a user can access all Help and Support content and services from both Microsoft and the OEM.

**hexadecimal** A base-16 number system represented by the digits 0 through 9 and the uppercase or lowercase letters A (equivalent to decimal 10) through F (equivalent to decimal 15).

**home directory** The root directory for a Web site, where the content files are stored. Also called a document root or Web root. In Internet Information Services (IIS), the home directory and all its subdirectories are available to users by default. Also, the root directory for an IIS service. Typically, the home directory for a site contains the home page. See also home page.

**home page** In the context of Internet Explorer, the home page is the first page users see when they start the browser. Home page is also a more general term for the main page of a Web site, which usually contains a main menu or table of contents with links to other pages within the site.

**host** Any device on a TCP/IP network that has an Internet Protocol (IP) address. Examples of hosts include servers, workstations, network-interface print devices, and routers. Sometimes used to refer to a specific network computer that is running a service used by network or remote clients. For Network Load Balancing, a cluster consists of multiple hosts connected over a local area network (LAN).See also client; cluster; local area network (LAN); Network Load Balancing; server; Transmission Control Protocol/Internet Protocol (TCP/IP).

**host name** The DNS name of a device on a network. These names are used to locate computers on the network. To find another computer, its host name must either appear in the Hosts file or be known by a DNS server. For most Windows-based computers, the host name and the computer name are the same. See also Domain Name System (DNS).

**hotfix** An update to address an issue identified after a software product has been distributed. Hotfix distribution is limited by its licensing terms.

**HTML** See definition for Hypertext Markup Language (HTML).

**HTTP** See definition for Hypertext Transfer Protocol (HTTP).

**HTTP header** An informational listing at the top of a Hypertext Transfer Protocol (HTTP) request or response.

**HTTPS** See definition for Secure Hypertext Transfer Protocol.

**Hypertext Markup Language (HTML)** A simple markup language used to create hypertext documents that are portable from one platform to another. HTML files are simple ASCII text files with codes embedded (indicated by markup tags) to denote formatting and hypertext links. See also American Standard Code for Information Interchange (ASCII).

**Hypertext Transfer Protocol (HTTP)** The protocol used to transfer information on the World Wide Web. An HTTP address (one kind of Uniform Resource Locator (URL)) takes the following form: http://www.microsoft.com. See also protocol.

**I**

**ICMP** See definition for Internet Control Message Protocol (ICMP).

**identities, multiple** See definition for multiple identities.

**identity** A person or entity that must be verified by means of authentication, based on criteria such as a password or a certificate. See also authentication; certificate.

**IEEE** Institute of Electrical and Electronics Engineers, founded in 1963. IEEE is an organization composed of engineers, scientists, and students, best known for developing standards for the computer and electronics industry.

**IETF** See definition for Internet Engineering Task Force (IETF).

**IIS** See definition for Internet Information Services (IIS).

**IIS 5.0 isolation mode** Internet Information Services (IIS) 6.0 isolation mode that simulates the IIS 5.0 Web process model.

**IIS Admin Objects** A set of methods, provided by Internet Information Services (IIS), that allow applications to access and modify configuration settings in the metabase.

**IIS Server Instance resource** A server-instance designation used with Internet Information Services (IIS) that supports the WWW and FTP services. IIS server instances are supported as cluster resources by a Resource DLL. IIS Server Instance resources can have dependencies on IP Address resources, Network Name resources, and Physical Disk resources. Access information for server instances does not fail over. See also failover; Internet Information Services (IIS).

**impersonation** A circumstance that occurs when Windows allows one process to take on the security attributes of another. See also attribute; security.

**impersonation token** An access token that captures the security information of a client process, allowing a service to impersonate the client process in security operations.

**in-memory metabase** An image of the Internet Information Services (IIS) metabase that has been loaded from disk into the computers RAM memory and is used while IIS is running. See also metabase.

**in-process** Internet Server API (ISAPI) extensions that are hosted in the worker process address space. See also Internet Server Application Programming Interface (ISAPI).

**in-schema property** A metabase property predefined in the metabase schema (MBSchema.xml) file. See also metabase schema.

**inheritance** In security, a mechanism that allows a specific access control entry (ACE) to be copied from the container where it was applied to all children of the container. Inheritance can be used to manage access to a whole subtree of objects in a single update operation. In Active Directory, the ability to build new object classes from existing object classes. The new object is defined as a subclass of the original object class. The original object class becomes a superclass of the new object. A subclass inherits the attributes of the superclass, including structure rules and content rules. In Group Policy, a mechanism that allows policy settings in Group Policy objects (GPOs) that are linked to parent containers to be applied to objects in child containers. See also Active Directory.

**input/output (I/O) port** A channel through which data is transferred between a device and the microprocessor. The port appears to the microprocessor as one or more memory addresses that it can use to send or receive data.

**Integrated Services Digital Network (ISDN)** A digital phone line used to provide higher bandwidth. ISDN in North America is typically available in two forms: Basic Rate Interface (BRI) consists of 2 B-channels at 64 kilobits per second (Kbps) and a D-channel at 16 Kbps; Primary Rate Interface (PRI) consists of 23 B-channels at 64 Kbps and a D-channel at 64 Kbps. An ISDN line must be installed by the phone company at both the calling site and the called site.

**Integrated Windows authentication** A configuration setting that enables negotiation of authentication protocols in Internet Information Services (IIS). See also Internet Information Services (IIS).

**internet** Two or more network segments connected by routers. Another term for *internetwork*.

**Internet** A worldwide network of computers. If you have access to the Internet, you can retrieve information from millions of sources, including schools, governments, businesses, and individuals.

**Internet Control Message Protocol (ICMP)** A required maintenance protocol in the TCP/IP suite that reports errors and provides simple diagnostic capabilities. ICMP is used by the Ping tool to perform TCP/IP troubleshooting. See also Internet Protocol (IP); protocol; Transmission Control Protocol/Internet Protocol (TCP/IP).

**Internet Engineering Task Force (IETF)** An open community of network designers, operators, vendors, and researchers concerned with the evolution of Internet architecture and the smooth operation of the Internet. Technical work is performed by working groups organized by topic areas (such as routing, transport, and security) and through mailing lists. Internet standards are developed in IETF Requests for Comments (RFCs), which are a series of notes that discuss many aspects of computing and computer communication, focusing on networking protocols, programs, and concepts.

**Internet Information Services (IIS)** Software services that support Web site creation, configuration, and management, along with other Internet functions. Internet Information Services include Network News Transfer Protocol (NNTP), File Transfer Protocol (FTP), and Simple Mail Transfer Protocol (SMTP). See also File Transfer Protocol (FTP).

**Internet Protocol (IP)** A routable protocol in the TCP/IP protocol suite that is responsible for IP addressing, routing, and the fragmentation and reassembly of IP packets. See also packet; Transmission Control Protocol/Internet Protocol (TCP/IP).

**Internet Protocol security (IPsec)** A set of industry-standard, cryptography-based protection services and protocols. IPsec protects all protocols in the TCP/IP protocol suite except Address Resolution Protocol (ARP). For virtual private network (VPN) connections, IPsec is used in conjunction with Layer Two Tunneling Protocol (L2TP). See also Address Resolution Protocol (ARP); protocol; Transmission Control Protocol/Internet Protocol (TCP/IP).

**Internet Server Application Programming Interface (ISAPI)** An application programming interface (API) that resides on a server computer, that initiates software services tuned for Windows operating systems. In Microsoft Provisioning System, ISAPI resides on the Web server. See also application programming interface (API).

**Internet service** One of any of a number of technologies for making information accessible to users over the Internet. Each Internet service is defined by a protocol, such as Hypertext Transfer Protocol (HTTP), and each is enabled using client/server applications, such as Web browsers and Web servers. Internet protocols are defined in the Request for Comments (RFC) documents that are published by the Internet Engineering Task Force (IETF).

**Internet service provider (ISP)** A company that provides individuals or companies access to the Internet and the World Wide Web. An ISP provides a telephone number, a user name, a password, and other connection information so users can connect their computers to the ISP's computers. An ISP typically charges a monthly or hourly connection fee.

**intranet** A network within an organization that uses Internet technologies and protocols, but is available only to certain people, such as employees of a company. Also called a *private network*. See also internet.

**IP** See definition for Internet Protocol (IP).

**IP address** For Internet Protocol version 4 (IPv4), a 32-bit address used to identify an interface on a node on an IPv4 internetwork. Each interface on the IP internetwork must be assigned a unique IPv4 address, which is made up of the network ID, plus a unique host ID. This address is typically represented with the decimal value of each octet separated by a period (for example, 192.168.7.27). You can configure the IP address statically or dynamically by using Dynamic Host Configuration Protocol (DHCP). For Internet Protocol version 6 (IPv6), an identifier that is assigned at the IPv6 layer to an interface or set of interfaces and that can be used as the source or destination of IPv6 packets. See also Dynamic Host Configuration Protocol (DHCP); Internet Protocol (IP); node; octet.

**IPsec** See definition for Internet Protocol security (IPsec).

**ISAPI** See definition for Internet Server Application Programming Interface (ISAPI).

**ISDN** See definition for Integrated Services Digital Network (ISDN).

**ISP** See definition for Internet service provider (ISP).

**J**

**JIT** See definition for just-in-time (JIT) activation.

**just-in-time (JIT) activation** The ability of a Component Object Model (COM) object to be activated only as needed for executing requests from its client. Objects can be deactivated even while clients hold references to them, allowing otherwise idle server resources to be used more productively.

**K**

**Keep-Alive connection** A Hypertext Transfer Protocol (HTTP) connection that is not closed after an exchange is completed.

**Kerberos V5 authentication protocol** An authentication mechanism used to verify user or host identity. The Kerberos V5 authentication protocol is the default authentication service. Internet Protocol security (IPsec) can use the Kerberos protocol for authentication. See also Internet Protocol security (IPsec).

**kernel** The core of layered architecture that manages the most basic operations of the operating system and the computer's

processor. The kernel schedules different blocks of executing code, called threads, for the processor; this efficiently keeps the processor as busy as possible and coordinates multiple processors to optimize performance. The kernel also synchronizes activities among executive-level subcomponents, such as I/O Manager and Process Manager, and handles hardware exceptions and other hardware-dependent functions. The kernel works closely with the hardware abstraction layer.

**kernel mode** A highly privileged mode of operation where program code has direct access to all memory, including the address spaces of all user-mode processes and applications, and to hardware. Also known as *supervisor mode*, *protected mode*, or *Ring 0*.

**key** In Registry Editor, a folder that appears in the left pane of the Registry Editor window. A key can contain subkeys and entries. For example, Environment is a key of **HKEY_CURRENT_USER**. In IP security (IPsec), a value used in combination with an algorithm to encrypt or decrypt data. Key settings for IPsec are configurable to provide greater security. See also Internet Protocol security (IPsec); registry; subkey.

**key pair** A private key and its related public key. See also private key; public key.

**L**

**late binding** See definition for dynamic binding.

**LCID** See definition for Locale Identifier (LCID).

**LDAP** See definition for Lightweight Directory Access Protocol (LDAP).

**Lightweight Directory Access Protocol (LDAP)** The primary access protocol for Active Directory. LDAP is an industry-standard protocol, established by the Internet Engineering Task Force (IETF), that allows users to query and update information in a directory service. Active Directory supports both LDAP version 2 and LDAP version 3. See also Active Directory; directory service; Internet Engineering Task Force (IETF); protocol.

**load balancing** A technique used by Windows Clustering to scale the performance of a server-based program (such as a Web server) by distributing its client requests across multiple servers within the cluster. Each host can specify the load percentage that it will handle, or the load can be equally distributed across all the hosts. If a host fails, Windows Clustering dynamically redistributes the load among the remaining hosts. See also cluster; host.

**local area network (LAN)** A communications network connecting a group of computers, printers, and other devices located within a relatively limited area (for example, a building). A LAN enables any connected device to interact with any other on the network.

**local group** A security group that can be granted rights and permissions on only resources on the computer on which the group is created. Local groups can have any user accounts that are local to the computer as members, as well as users, groups, and computers from a domain to which the computer belongs.

**Locale Identifier (LCID)** A unique integer that represents a locale for the formatting style of dates, times, currencies and other values, which is different for each geographical location. In Internet Information Services (IIS), the default LCID is the same as the system locale. See also LOCALE_SYSTEM_DEFAULT.

**LOCALE_SYSTEM_DEFAULT** The default system locale. There is also a default user locale. See also Locale Identifier (LCID).

**LocalHost** A placeholder for the name of the computer on which a program is running. LocalHost uses the reserved loopback IP address (127.0.0.1 in IPv4 and ::1 in IPv6).

**log file** A file that stores messages generated by an application, service, or operating system. These messages are used to track the operations performed. For example, Web servers maintain log files listing every request made to the server. Log files are usually plain text (ASCII) files and often have a .log extension. In Backup, a file that contains a record of the date the tapes were created and the names of files and directories successfully backed up and restored. The Performance Logs and Alerts service also creates log files. See also American Standard Code for Information Interchange (ASCII).

**logical drive** A volume that you create within an extended partition on a basic master boot record (MBR) disk. Logical drives are similar to primary partitions, except that you are limited to four primary partitions per disk, whereas you can create an unlimited number of logical drives per disk. A logical drive can be formatted and assigned a drive letter. See also extended partition.

**M**

**Mail or Messaging Application Programming Interface (MAPI)** An open and comprehensive messaging interface that is used by developers to create messaging and workgroup applications, such as e-mail, scheduling, calendars, and document management. In a distributed client/server environment, MAPI provides enterprise messaging services within Windows Open Services Architecture (WOSA).

**malicious user** A person who has legitimate access to a system and poses a security threat to it, such as someone who tries to elevate their user rights to gain access to unauthorized data. See also security; user rights.

**Management Information Base (MIB)** Information about aspects of a network that can be managed by using the Simple Network Management Protocol (SNMP). This information is formatted in MIB files that are provided for each service that can be monitored. Most third-party monitors (clients) use SNMP and MIB files to monitor Web, File Transfer Protocol (FTP), and other Windows services. Using SNMP, developers or system administrators can write their own custom monitoring applications.

**MAPI** See definition for Mail or Messaging Application Programming Interface (MAPI).

**master properties** In Internet Information Services (IIS), properties that are set at the computer level that become default settings for all Web or File Transfer Protocol (FTP) sites on that computer. See also inheritance.

**MD5** An industry-standard one-way, 128-bit hashing scheme, developed by RSA Data Security, Inc., and used by various Point-to-Point Protocol (PPP) vendors for encrypted authentication. A hashing scheme is a method for transforming data (for example, a password) in such a way that the result is unique and cannot be changed back to its original form. The Challenge Handshake Authentication Protocol (CHAP) uses challenge-response with one-way MD5 hashing on the response. In this way, you can prove to the server that you know your password without actually sending the password over the network. See also hash algorithm; MD5.

**MDAC** See definition for Microsoft Data Access Components (MDAC).

**Message Queuing** A message queuing and routing system for Windows that enables distributed applications running at different times to communicate across heterogeneous networks and with computers that may be offline. Message Queuing provides guaranteed message delivery, efficient routing, security, and priority-based messaging. Formerly known as *MSMQ*.

**metabase** A hierarchical store of configuration information and schema that is used to configure Internet Information Services (IIS). The metabase performs some of the same functions as the system registry, but it uses less disk space. In physical terms, the metabase is a combination of the MetaBase.xml and MBSchema.xml files and the in-memory metabase.

**metabase configuration file** A file that stores Internet Information Services (IIS) configuration settings to disk. This file is named MetaBase.xml by default. When IIS is started or restarted, the configuration settings are read from MetaBase.xml into the IIS cache in memory, which is called the in-memory metabase.

**metabase schema** The master configuration file (MBSchema.xml) supplied with Internet Information Services (IIS) that contains all of the predefined properties from which metabase entries are derived.

**metadata** Data that is used to describe other data. For example, Indexing Service must maintain data that describes the data in the content index.

**method** A procedure (function) that acts on an object.

**MIB** See definition for Management Information Base (MIB).

**Microsoft Data Access Components (MDAC)** Consists of ActiveX Data Objects (ADO), the Remote Data Service (RDS), Microsoft OLE DB Provider for ODBC, Open Database Connectivity (ODBC), ODBC drivers for Microsoft SQL Server, Microsoft Access and other desktop databases, as well as Oracle databases.

**middle tier** The logical layer between a user interface or Web client and the database. This is typically where the Web server resides and where business objects are instantiated. Also known as *application server tier*. See also client tier; data source tier.

**modem (modulator/demodulator)** A device that enables computer information to be transmitted and received over a telephone line. The transmitting modem translates digital computer data into analog signals that can be carried over a telephone line. The receiving modem translates the analog signals back to digital form.

**MSMQ** See definition for Message Queuing.

**MTA** See definition for multithreaded apartment (MTA).

**multiple identities** Multiple Web sites that are hosted on one computer. Also called virtual servers.

**Multipurpose Internet Mail Extensions mapping (MIME mapping)** A method of configuring browsers to view files that are in multiple formats. An extension of the Internet mail protocol that enables the sending of 8-bit-based e-mail messages, which are used to support extended character sets, voice mail, facsimile images, and so on.

**multithreaded apartment (MTA)** A form of multithreading that is supported by Component Object Model (COM). In a multithreaded apartment model, all of the threads in the process that have been initialized as free-threaded reside in a single

apartment.

**multithreading** Running several processes in rapid sequence within a single program, regardless of which logical method of multitasking is being used by the operating system. Because the user's sense of time is much slower than the processing speed of a computer, multitasking appears to be simultaneous, even though only one task at a time can use a computer processing cycle.

**N**

**name resolution** The process of having software translate between names that are easy for users to work with and numerical IP addresses, which are difficult for users but necessary for TCP/IP communications. Name resolution can be provided by software components such as DNS or WINS. See also Domain Name System (DNS); Transmission Control Protocol/Internet Protocol (TCP/IP); Windows Internet Name Service (WINS).

**namespace** A naming convention that defines a set of unique names for resources in a network. For DNS, a hierarchical naming structure that identifies each network resource and its place in the hierarchy of the namespace. For WINS, a flat naming structure that identifies each network resource using a single, unique name. See also Domain Name System (DNS); Windows Internet Name Service (WINS).

**Network File System (NFS)** A service for distributed computing systems that provides a distributed file system, eliminating the need for keeping multiple copies of files on separate computers.

**network latency** The time it takes for information to be transferred between computers in a network.

**Network Load Balancing** A Windows network component that uses a distributed algorithm to load-balance Internet Protocol (IP) traffic across a number of hosts, helping to enhance the scalability and availability of mission-critical, IP-based services, such as Terminal Services, Web services, virtual private networking, and streaming media. It also provides high availability by detecting host failures and automatically redistributing traffic to the surviving hosts. See also availability; cluster; host; scalability.

**Network News Transfer Protocol (NNTP)** A protocol that is used to distribute network news messages to NNTP servers and to NNTP clients (news readers) on the Internet. NNTP provides for the distribution, inquiry, retrieval, and posting of news articles by using a reliable, stream-based transmission of news on the Internet. NNTP is designed in such a way that news articles are stored on a server in a central database, so that users can select specific items to read. Indexing, cross-referencing, and expiration of old messages are also provided. NNTP is defined in RFC 977.

**NNTP** See definition for Network News Transfer Protocol (NNTP).

**node** For tree structures, a location on the tree that can have links to one or more items below it. For local area networks (LANs), a device that is connected to the network and is capable of communicating with other network devices. For server clusters, a computer system that is an active or inactive member of a cluster. See also local area network (LAN); server cluster.

**NTFS file system** An advanced file system that provides performance, security, reliability, and advanced features that are not found in any version of file allocation table (FAT). For example, NTFS guarantees volume consistency by using standard transaction logging and recovery techniques. If a system fails, NTFS uses its log file and checkpoint information to restore the consistency of the file system. NTFS also provides advanced features, such as file and folder permissions, encryption, disk quotas, and compression. See also FAT32; file allocation table (FAT).

**O**

**object** An entity, such as a file, folder, shared folder, printer, or Active Directory object, described by a distinct, named set of attributes. For example, the attributes of a File object include its name, location, and size; the attributes of an Active Directory User object might include the user's first name, last name, and e-mail address. For OLE and ActiveX, an object can also be any piece of information that can be linked to, or embedded into, another object. See also Active Directory; attribute.

**object identifier (OID)** An extensible, unique identification number for attributes and classes. Performance counter names have their own object identifiers, which are listed in Management Information Base (MIB) files, to provide performance monitoring applications with access to the counters. See also Management Information Base (MIB); Simple Network Management Protocol (SNMP).

**Object Linking and Embedding (OLE)** A set of integration standards for transferring and sharing information among client applications. Also, a protocol that enables the creation of compound documents with embedded links to applications, so that a user does not have to switch among applications to make revisions. OLE is based on the Component Object Model (COM), and it enables the development of reusable objects that operate across multiple applications.

**Object Linking and Embedding Database (OLE DB)** Data-access interfaces that provide consistent access to Structured Query Language (SQL) data sources and non-SQL data sources across an organization and the Internet. See also structured query language (SQL).

**object-cache scavenger** The code that periodically scans the cache for objects to be discarded. It deletes files that have not been used recently and therefore are unlikely to be used again in the near future.

**octet** In programming, an octet refers to eight bits or one byte. IP addresses, for example, are typically represented in dotted-decimal notation; that is, with the decimal value of each octet of the address separated by a period.

**ODBC** See definition for Open Database Connectivity (ODBC).

**OID** See definition for object identifier (OID).

**OLE** See definition for Object Linking and Embedding (OLE).

**OLE DB** See definition for Object Linking and Embedding Database (OLE DB).

**Open Database Connectivity (ODBC)** An application programming interface (API) that enables applications to access data from a variety of existing data sources. A standard specification for cross-platform database access.

**out-of-process** For IIS 5.0 isolation mode, ISAPI extensions that are hosted in a surrogate process called DLLHOST.exe, which is managed by COM+. See also IIS 5.0 isolation mode.

**out-of-process component** A Component Object Model (COM) component that runs in a separate process space from its client. See also process isolation.

**out-of-process, pooled** See definition for pooled out-of-process.

**P-Q**

**packet** An Open Systems Interconnection (OSI) network layer transmission unit that consists of binary information representing both data and a header containing an identification number, source and destination addresses, and error-control data.

**page** See definition for Web page.

**parameter** A value that is passed in a function call.

**partition** A portion of a physical disk that functions as though it were a physically separate disk. After you create a partition, you must format it and assign it a drive letter before you can store data on it. On basic disks, partitions are known as basic volumes, which include primary partitions and logical drives. On dynamic disks, partitions are known as dynamic volumes, which include simple, striped, spanned, mirrored, and RAID-5 volumes. See also extended partition.

**password authentication** See definition for authentication.

**path, physical** See definition for physical path.

**path, relative** See definition for relative path.

**performance counter** In System Monitor, a data item that is associated with a performance object. For each counter selected, System Monitor presents a value corresponding to a particular aspect of the performance that is defined for the performance object.

**Perl (Practical Extraction and Report Language)** An interpreted language that is based on C and several UNIX utilities. Perl has powerful string-handling features for extracting information from text files. Perl can assemble a string and send it to the shell as a command; therefore, it is often used for system administration tasks. A program in Perl is known as a script. Perl was devised by Larry Wall at the NASA Jet Propulsion Laboratory. See also script.

**permission** A rule associated with an object to regulate which users can gain access to the object and in what manner. Permissions are assigned or denied by the objects owner. See also object.

**physical path** A universal naming convention (UNC) directory path. See also relative path.

**ping** A utility that verifies connections to one or more remote hosts. The ping command uses Internet Control Message Protocol (ICMP) echo request and echo reply packets to determine whether a particular Internet Protocol (IP) system on a network is functional. **Ping** is useful for diagnosing IP network or router failures. See also host; Internet Control Message Protocol (ICMP); Internet Protocol (IP); packet.

**plaintext** Data that is not encrypted. Sometimes also called *cleartext*. See also encryption.

**Point-to-Point Protocol (PPP)** A set of industry-standard framing and authentication protocols that are included with Windows to ensure interoperability with other remote access software. PPP negotiates configuration parameters for multiple

layers of the Open Systems Interconnection (OSI) model. The Internet standard for serial communications, PPP defines how data packets are exchanged with other Internet-based systems using a modem connection.

**Point-to-Point Tunneling Protocol (PPTP)** A specification for virtual private networks (VPNs) in which some nodes of a local area network (LAN) are connected through the Internet. PPTP is an open industry standard that supports widely used networking protocols: Internet Protocol (IP), Internetwork Packet Exchange (IPX), and Microsoft NetBIOS Extended User Interface (NetBEUI). Organizations can use PPTP to outsource their remote dial-up needs to an Internet service provider or to other carriers to reduce cost and complexity.

**policies** Conditions that are set by the system administrator, for example, how quickly account passwords expire and how many unsuccessful logon attempts are allowed before a user is locked out. Policies manage accounts to help prevent exhaustive or random password attacks.

**pooled out-of-process** For IIS 5.0 isolation mode, a special Web Application Manager (WAM) package that hosts all out-of-process ISAPI extensions that are set to medium isolation within the same DLLHOST.exe process. See also IIS 5.0 isolation mode; out-of-process; Web Application Manager (WAM).

**port number** A number that identifies a certain Internet application. For example, the default port number for the WWW service is 80.

**PPP** See definition for Point-to-Point Protocol (PPP).

**PPTP** See definition for Point-to-Point Tunneling Protocol (PPTP).

**private key** The secret half of a cryptographic key pair that is used with a public key algorithm. Private keys are typically used to decrypt a symmetric session key, digitally sign data, or decrypt data that has been encrypted with the corresponding public key. See also public key; public key encryption.

**process** An operating system object that consists of an executable program, a set of virtual memory addresses, and one or more threads. When a program runs, a process is created.

**process accounting** A feature of Internet Information Services (IIS) that administrators can use to monitor and log resource consumption of Common Gateway Interface (CGI) scripts and out-of-process applications.

**process isolation** Running an application or component out of process. See also out-of-process component.

**property inheritance** See definition for inheritance.

**protocol** A set of rules and conventions for sending information over a network. These rules govern the content, format, timing, sequencing, and error control of messages exchanged among network devices. See also Internet Protocol (IP); Transmission Control Protocol/Internet Protocol (TCP/IP).

**provider** See definition for WMI provider.

**proxy** A software program that connects a user to a remote destination through an intermediary gateway.

**proxy server** A firewall component that manages Internet traffic to and from a local area network (LAN) and that can provide other features, such as document caching and access control. A proxy server can improve performance by supplying frequently requested data, such as a popular Web page, and it can filter and discard requests that the owner does not consider appropriate, such as requests for unauthorized access to proprietary files. See also firewall; local area network (LAN).

**public key** The non-secret half of a cryptographic key pair that is used with a public key algorithm. Public keys are typically used when encrypting a session key, verifying a digital signature, or encrypting data that can be decrypted with the corresponding private key. See also key; private key; public key encryption.

**public key encryption** A method of encryption that uses two encryption keys that are mathematically related. One key is called the *private key* and is kept confidential. The other is called the *public key* and is freely given out to all potential correspondents. In a typical scenario, a sender uses the receiver's public key to encrypt a message. Only the receiver has the related private key to decrypt the message. The complexity of the relationship between the public key and the private key means that, provided the keys are long enough, it is computationally infeasible to determine one from the other. Also called *asymmetric encryption*. See also encryption; private key; public key; symmetric encryption.

**public-key algorithm** An asymmetric cipher that uses two keys, one for encryption, the public key, and the other for decryption, the private key. See also asymmetric key algorithm; decryption; encryption; private key; public key.

**R**

**RAID** See definition for Redundant Array of Independent Disks (RAID).

**random access memory (RAM)** Memory that can be read from or written to by a computer or other devices. Information stored in RAM is lost when the computer is turned off.

**realm** A term that is sometimes used for domain, in this case to refer to user domains that are established for security reasons, not Internet domains. For password-protected files, the name of the protected resource or area on the server. If the user tries to access the protected resource while browsing, the name of the realm usually appears in the dialog box that asks for a user name and password.

**redirection** Redirection can be used to automatically send a user from an outdated Uniform Resource Locator (URL) to a new URL.

**Redundant Array of Independent Disks (RAID)** A data storage method in which data, along with information used for error correction, such as parity bits, is distributed among two or more hard disk drives to improve performance and reliability. The hard disk array is governed by array management software and a disk controller, which handles the error correction. RAID is generally used on network servers. Several defined levels of RAID offer differing trade-offs among access speed, reliability, and cost. Windows includes three of the RAID levels: Level 0, Level 1, and Level 5.

**registry** A database repository for information about a computer's configuration. The registry contains information that Windows continually references during operation, such as:

- Profiles for each user.
- The programs installed on the computer and the types of documents that each can create.
- Property settings for folders and program icons.
- What hardware exists on the system.
- Which ports are being used.

The registry is organized hierarchically as a tree, and it is made up of keys and their subkeys, hives, and entries. See also key; subkey.

**relative path** A universal naming convention (UNC) directory path with placeholders, or wildcards, at some levels. Also, the physical path that corresponds to a Uniform Resource Locator (URL). See also physical path.

**remote procedure call (RPC)** In programming, a call by one program to a second program on a remote system. The second program usually performs a task and returns the results of that task to the first program.

**remotely administered server** A server that you can administer by using a different computer. You typically access this type of server by using a network connection. A remotely administered server can have a local keyboard, mouse, or video card and monitor. If it does not have these peripherals attached, it is also known as a *headless server*. Such servers are often housed in a physically secure location. See also server.

**replication** The copying from one server node to another of either content or the configuration metabase, or both. This copying can be done either manually or automatically by using replication software. Replication is a necessary function of clustering that ensures fault tolerance. See also fault tolerance.

**Request for Comments (RFC)** The document series, begun in 1969, that describes the Internet suite of protocols and related experiments. Not all (in fact, very few) RFCs describe Internet standards, but all Internet standards are written up as RFCs. The RFC series of documents is unusual in that the proposed protocols are forwarded by the Internet research and development community, acting on its own behalf, as opposed to the formally reviewed and standardized protocols that are promoted by organizations such as the American National Standards Institute (ANSI).

**router** An intermediary device on a communications network that expedites message delivery. On a single network linking many computers through a mesh of possible connections, a router receives transmitted messages and forwards them to their correct destinations over the most efficient available route. On an interconnected set of local area networks (LANs) using the same communications protocols, a router serves the somewhat different function of acting as a link between LANs, enabling messages to be sent from one LAN to another.

**RPC** See definition for remote procedure call (RPC).

**RSA** A public-key encryption standard for Internet security. This acronym derives from the last names of the inventors of the technology: Rivest, Shamir, and Adleman.

**S**

**scalability** A measure of how well a computer, service, or application can grow to meet increasing performance demands. For server clusters, the ability to incrementally add one or more systems to an existing cluster when the overall load of the cluster exceeds its capabilities. See also server cluster.

**schema** A representation of the structure of something. Classes in Visual Basic and C++ can be said to be schemas of objects, and objects are instances of classes. In Internet Information Services (IIS), the metabase schema represents the structure of the metabase configuration file.

**scope** In programming, the extent to which an identifier, such as a constant, data type, variable, or routine, can be referenced within a program. Scope can be global or local. Scope can also be affected by the redefinition of identifiers, for example, by giving the same name to both a global variable and a local variable.

**script** A kind of program that consists of a set of instructions for an application or utility program. A script can be embedded in a Web page. See also ActiveX; common gateway interface (CGI).

**scripting engine** A program that interprets and executes a script. See also script.

**search expression** See Other Definition

**search interface** See Other Definition

**search string** See Other Definition

**Secure Hash Algorithm (SHA-1)** An algorithm that generates a 160-bit hash value from an arbitrary amount of input data. SHA-1 is used with the Digital Signature Algorithm (DSA) in the Digital Signature Standard (DSS), among other places. See also hash algorithm.

**Secure Hypertext Transfer Protocol** A protocol that provides a secure Hypertext Transfer Protocol (HTTP) connection. See also Hypertext Transfer Protocol (HTTP); protocol.

**Secure Sockets Layer (SSL)** A protocol that supplies secure data communication through data encryption and decryption. SSL uses RSA public-key encryption for specific TCP/IP ports. It is intended for handling commerce payments. An alternative method is Secure-HTTP (S-HTTP), which is used to encrypt specific Web documents, rather than the entire session. SSL is a general-purpose encryption standard. SSL can also be used for Web applications that require a secure link, such as e-commerce applications, or for controlling access to Web-based subscription services.

**security** On a network, protection of a computer system and its data from harm or loss, implemented especially so that only authorized users can gain access to shared files. See also authorization.

**security context** The security attributes or rules that are currently in effect. For example, the rules that govern what a user can do to a protected object are determined by security information in the user's access token and in the object's security descriptor. Together, the access token and the security descriptor form a security context for the user's actions on the object. See also object.

**Selectable Cryptographic Service Provider** See definition for cryptographic service provider (CSP).

**server** In general, a computer that provides shared resources to network users. See also client; shared resource.

**server certificate** A unique digital identification that forms the basis of a Web server's Secure Sockets Layer (SSL) security features. Server certificates are obtained from a mutually trusted, third-party organization, and they provide a way for users to authenticate the identity of a Web site.

**server cluster** A group of computers, known as *nodes*, working together as a single system to ensure that mission-critical applications and resources remain available to clients. A server cluster presents the appearance of a single server to a client. See also cluster; node.

**server node** An individual computer in a server cluster.

**server process** A process that hosts Component Object Model (COM) components. A COM component can be loaded into a surrogate server process, either on the client computer (local) or on another computer (remote). It can also be loaded into a client application process (in-process).

**server-side include** (SSI) A mechanism for including dynamic text in World Wide Web documents. Server-side includes are special command codes that are recognized and interpreted by the server; their output is placed in the document body before the document is sent to the browser. Server-side includes can be used, for example, to include the Date and Time stamp in the text of the file.

**session key** A digital key that is created by a client, encrypted, and then sent to a server. This key is used to encrypt data that is sent by the client.

**shared resource** Any device, data, or program that is used by more than one program or one other device. For Windows, *shared resource* refers to any resource that is made available to network users, such as folders, files, printers, and named pipes. *Shared resource* can also refer to a resource on a server that is available to network users. See also server.

**Simple Mail Transfer Protocol (SMTP)** A TCP/IP protocol for sending messages from one computer to another on a network. This protocol is used on the Internet to route e-mail.

**Simple Network Management Protocol (SNMP)** The network management protocol of TCP/IP. In SNMP, agents or clients monitor the activity of various devices on the network and report to the network console workstation. The agents or clients can be hardware as well as software. Control information about each device or service is maintained in a structure known as a management information block. One way to access this information is with Performance Counters. See also Management Information Base (MIB).

**single-threaded apartment (STA)** A form of threading that is supported by Component Object Model (COM). In a single-threaded apartment model, all objects are executed on a single thread and each thread resides within its own apartment.

**SMTP** See definition for Simple Mail Transfer Protocol (SMTP).

**snap-in** A type of tool that you can add to a console supported by Microsoft Management Console (MMC). A stand-alone snap-in can be added by itself; an extension snap-in can be added only to extend the function of another snap-in.

**sniffer** An application or device that can read, monitor, and capture network data exchanges and read network packets. If the packets are not encrypted, a sniffer provides a full view of the data inside the packet. See also packet.

**SNMP** See definition for Simple Network Management Protocol (SNMP).

**socket** An identifier for a particular service on a particular node on a network. The socket consists of a node address and a port number, which identifies the service. For example, port 80 on an Internet node indicates a Web server. There are two kinds of sockets: streams (bidirectional) and datagrams. See also datagram; node.

**SSI** See definition for server-side include (SSI).

**SSL** See definition for Secure Sockets Layer (SSL).

**STA** See definition for single-threaded apartment (STA).

**stateful object** An object that holds private state accumulated from the execution of one or more client calls.

**stateless object** An object that does not hold private state accumulated from the execution of one or more client calls.

**static page** A Hypertext Markup Language (HTML) page that is prepared in advance of a request for it and that is sent to the client upon request. This page takes no special action when it is requested. See also dynamic page.

**structured query language (SQL)** A widely accepted standard database sublanguage used in querying, updating, and managing relational databases.

**subkey** An element of the registry that contains entries or other subkeys. A tier of the registry that is immediately below a key or a subtree (if the subtree has no keys). See also key; registry.

**subnet mask** A 32-bit value that enables the recipient of Internet Protocol version 4 (IPv4) packets to distinguish the network ID and host ID portions of the IPv4 address. Typically, subnet masks use the format 255.x.x.x. IPv6 uses network prefix notations rather than subnet masks. See also IP address.

**symmetric encryption** An encryption algorithm that requires the same secret key to be used for both encryption and decryption. Because of its speed, symmetric encryption is typically used when a message sender needs to encrypt large amounts of data. Also called *secret key encryption*. See also public key encryption.

**System Data Source Name (DSN)** A name that can be used by any process on the computer. Internet Information Services (IIS) uses system DSNs to access Open Database Connectivity (ODBC) data sources.

**systemroot** The path and folder name where the Windows system files are located. Typically, this is C:\Windows, although you can designate a different drive or folder when you install Windows. You can use the value *%systemroot%* to replace the actual location of the folder that contains the Windows system files. To identify your systemroot folder, click **Start**, click **Run**, type **%systemroot%**, and then click **OK**.

**Systems Network Architecture (SNA)** A communications framework developed by IBM to define network functions and establish standards for enabling computers to share and process data.

**T**

**T1** A U.S. telephone standard for a transmission facility at digital signal level 1 (DS1) with 1.544 megabits per second in North America and 2.048 megabits per second in Europe. The bit rate is the equivalent bandwidth of approximately twenty-four 56-kilobits-per-second lines. A T1 circuit is capable of serving a minimum of 48 modems at 28.8 kilobits per second or 96 modems at 14.4 kilobits per second. T1 circuits are also used for voice telephone connections. A single T1 line carries 24 telephone connections with 24 telephone numbers. When it is used for voice transmission, a T1 connection must be split into 24 separate circuits.

**T3** A U.S. telephone standard for a transmission facility at digital signal level 3 (DS3). T3 is equivalent in bandwidth to 28 T1s, and the bit rate is 44.736 megabits per second. T3 is sometimes called a 45-meg circuit.

**TCP/IP** See definition for Transmission Control Protocol/Internet Protocol (TCP/IP).

**Telnet** A protocol that enables an Internet user to log on to and enter commands on a remote computer linked to the Internet, as if the user were using a text-based terminal directly attached to that computer. Telnet is part of the TCP/IP suite of protocols. The term *telnet* also refers to the software (client or server component) that implements this protocol. See also protocol; Transmission Control Protocol/Internet Protocol (TCP/IP).

**thread** A type of object within a process that runs program instructions. Using multiple threads allows concurrent operations within a process and enables one process to run different parts of its program on different processors simultaneously. A thread has its own set of registers, its own kernel stack, a thread environment block, and a user stack in the address space of its process. See also kernel.

**throttling** Controlling the maximum amount of bandwidth that is dedicated to Internet traffic on a server. This feature is useful if there are other services (such as e-mail) sharing the server over a busy link.

**time-out** A setting that automatically cancels an unanswered client request after a certain period of time.

**TP** See definition for transaction processing (TP).

**transaction processing (TP)** The real-time handling of computerized business transactions as they are received by the system. Also called online transaction processing (OLTP).

**Transmission Control Protocol/Internet Protocol (TCP/IP)** A set of networking protocols widely used on the Internet that provides communications across interconnected networks of computers with diverse hardware architectures and various operating systems. TCP/IP includes standards for how computers communicate and conventions for connecting networks and routing traffic. See also Internet Protocol (IP); protocol.

**Transport Layer Security (TLS) encryption** A generic security protocol similar to Secure Sockets Layer (SSL), used with Simple Mail Transfer Protocol (SMTP). See also Secure Sockets Layer (SSL).

**Triple DES (3DES)** An implementation of Data Encryption Standard (DES) encryption that employs three iterations of cryptographic operations on each segment of data. Each iteration uses a 56-bit key for encryption, which yields 168-bit encryption for the data. Although 3DES is slower than DES because of the additional cryptographic calculations, its protection is far stronger than DES. See also cryptography; Data Encryption Standard (DES); encryption.

**two-tier architecture** See definition for client/server architecture.

**type library** A binary file that describes a components methods, properties, and data structure.

**U**

**UNC (Universal Naming Convention) name** The full name of a resource on a network. It conforms to the *servername**sharename* syntax, where *servername* is the name of the server and *sharename* is the name of the shared resource. UNC names of directories or files can also include the directory path under the share name with the following syntax: *servername**sharename**directory**filename*

**Uniform Resource Locator (URL)** A naming convention that uniquely identifies the location of a computer, directory, or file on the Internet. A URL also specifies the appropriate Internet protocol, such as Hypertext Transfer Protocol (HTTP) or File Transfer Protocol (FTP). An example of a URL is http://www.microsoft.com.

**Universal Naming Convention (UNC)** A convention for naming files and other resources beginning with two backslashes (\), indicating that the resource exists on a network computer. UNC names conform to the *servername**sharename* syntax, where *servername* is the server's name and *sharename* is the name of the shared resource. The UNC name of a directory or file can also include the directory path after the share name by using the following syntax:

*\\servername\sharename\directory\filename*.

**upload** In communications, the process of transferring a copy of a file from a local computer to a remote computer by means of a modem or network. With a modem-based communications link, the process generally involves instructing the remote computer to prepare to receive the file on its disk and then wait for the transmission to begin.

**URL** See definition for Uniform Resource Locator (URL).

**URL directory** See definition for virtual directory.

**URL mapping** The process of associating a Uniform Resource Locator (URL) with a physical directory. See also virtual directory.

**usage data** Information that an administrator can use to learn how other people access and use a site. By analyzing this data, an administrator can identify a site's most popular (or unpopular) areas and clarify the most common navigational paths through the site.

**user rights** Tasks that a user is permitted to perform on a computer system or domain. There are two types of user rights: privileges and logon rights. An example of a privilege is the right to shut down the system. An example of a logon right is the right to log on to a computer locally. Both types are assigned by administrators to individual users or groups as part of the security settings for the computer. See also domain; group.

**user type** A DWORD that specifies how data is used. A user type is assigned to an identifier in the metabase.

**UTF-8** A method of character encoding that allows for both single and multibyte characters in one string. UTF-8 files take up more space than files that are stored in an American National Standards Institute (ANSI) format. Internet Information Services (IIS) supports Web files that are saved in UTF-8 format or in ANSI format. See also code page.

**V**

**virtual directory** A directory name, used in an address, that corresponds to a physical directory on the server. Sometimes called URL mapping.

**virtual server** A virtual computer that resides on a Hypertext Transfer Protocol (HTTP) server but appears to the user as a separate HTTP server. Several virtual servers can reside on one computer, each capable of running its own programs and each with individualized access to input and peripheral devices. Each virtual server has its own domain name and IP address, and each appears to the user as an individual Web site or File Transfer Protocol (FTP) site. Some Internet service providers (ISPs) use virtual servers for those clients who want to use their own domain names. Also called a Web site.

**volatile objects** Typically, files that a Web site administrator updates frequently.

**W-Z**

**W3C** See definition for World Wide Web Consortium (W3C).

**W3SVC** See definition for World Wide Web Publishing Service (WWW service).

**WAM** See definition for Web Application Manager (WAM).

**WBEM** See definition for Web-Based Enterprise Management (WBEM).

**Web application** A software program that uses Hypertext Transfer Protocol (HTTP) for its core communication protocol and that delivers Web-based information to the user in the Hypertext Markup Language (HTML) language. Also called a Web-based application.

**Web Application Manager (WAM)** For IIS 5.0 isolation mode, a COM+ application package that works with DLLHOST.exe to host out-of-process ISAPI extensions. Provides communication between DLLHOST.exe and INETINFO.exe. See also IIS 5.0 isolation mode.

**Web Distributed Authoring and Versioning (WebDAV)** An extension to the Hypertext Transfer Protocol (HTTP) 1.1 standard that facilitates access to files and directories through an HTTP connection. Remote authors can add, search, delete, or change directories and documents and their properties.

**Web farm** A Network Load Balancing cluster of IIS servers that support client Web site requests.

**Web garden** An application pool served by more than one worker process.

**Web page** A World Wide Web document. A Web page typically consists of a Hypertext Markup Language (HTML) file, with associated files for graphics and scripts, in a particular directory on a particular computer. It is identified by a Uniform Resource Locator (URL).

**Web server** In general, a computer that is equipped with server software that uses Internet protocols such as Hypertext Transfer Protocol (HTTP) and File Transfer Protocol (FTP) to respond to Web client requests on a TCP/IP network.

**Web service extensions** ISAPIs and CGIs that extend Internet Information Services (IIS) functionality beyond serving static pages.

**Web-Based Enterprise Management (WBEM)** An industry initiative to develop a standard technology for accessing management information about systems in an enterprise environment. The Microsoft implementation of Web-based Enterprise Management is Windows Management Instrumentation (WMI).

**WebDAV** See definition for Web Distributed Authoring and Versioning (WebDAV).

**wide area network (WAN)** A communications network connecting geographically separated locations that uses long-distance links of third-party telecommunications vendors. See also local area network (LAN).

**wildcard character** A keyboard character that can be used to represent one or many characters when conducting a query. The question mark (?) represents a single character, and the asterisk (*) represents one or more characters.

**Windows Internet Name Service (WINS)** A Windows name resolution service for network basic input/output system (NetBIOS) names. WINS is used by hosts running NetBIOS over TCP/IP (NetBT) to register NetBIOS names and to resolve NetBIOS names to Internet Protocol (IP) addresses. See also IP address.

**Windows Management Instrumentation (WMI)** The Microsoft implementation of Web-Based Enterprise Management (WBEM), which is an industry-wide standard technology for accessing management information about systems in an enterprise environment. WMI uses the Common Information Model (CIM) industry standard to represent managed components in a system. A system developer can develop a WMI interface that allows programmatic access to a system, so that users can write command-line administration scripts and tools. See also Web-Based Enterprise Management (WBEM).

**Windows Script Host (WSH)** A language-independent scripting host for ActiveX scripting engines on 32-bit Windows platforms.

**WMI** See definition for Windows Management Instrumentation (WMI).

**WMI provider** In Windows Management Instrumentation (WMI), a set of interfaces that provide programmatic access to management information in a system. Internet Information Services (IIS) implements a WMI provider in the namespace called MicrosoftIISv2to provide programmatic access to metabase properties and system settings.

**worker process** The implementation of the core Web server in Internet Information Services (IIS). Worker processes run in W3wp.exe.

**worker process isolation mode** The new Web process model for Internet Information Services (IIS) 6.0.

**worker thread** A thread that is created by a component, Internet Server API (ISAPI) extension, or filter to perform asynchronous processing. Using worker threads frees up Internet Information Services (IIS) I/O threads to process additional requests.

**working directory** The directory in which Web server software is installed. The RAM that is allocated to a process in the Windows operating system.

**World Wide Web (WWW)** A set of services that run on top of the Internet and provide a cost-effective way of publishing information, supporting collaboration and workflow, and delivering business applications to connected users all over the world. The Web is a collection of Internet host systems that make these services available on the Internet, using the Hypertext Transfer Protocol (HTTP). Web-based information is usually delivered in the form of hypertext and hypermedia, using Hypertext Markup Language (HTML). The most graphical service on the Internet, the Web also has the most sophisticated linking abilities.

**World Wide Web Consortium (W3C)** An international industry consortium that is jointly hosted by the Massachusetts Institute of Technology Laboratory for Computer Science (MIT/CS) in North America, by the Institut National de Recherche en Informatique et en Automatique (INRIA) in Europe, and by the Keio University Shonan Fujisawa Campus in Asia. W3C was founded in 1994 to develop common standards for the World Wide Web. Initially, the W3C was established in collaboration with CERN, where the Web originated, with support from the Defense Advanced Research Projects Agency (DARPA) and the European Commission.

**World Wide Web Publishing Service (WWW service)** The service that manages the Internet Information Services (IIS) core components that process HTTP requests and configure and manage Web applications. Formerly known as W3SVC.

**WSH** See definition for Windows Script Host (WSH).

**WWW** See definition for World Wide Web (WWW).

**WWW Service Administration and Monitoring component** A component of the World Wide Web Publishing Service (WWW service) in Internet Information Services (IIS) that is responsible for configuration, by means of the metabase, and for worker process management.

**XML** See definition for Extensible Markup Language (XML).